Psychological and Sociological Physics

The Creature Comforts Corpus

Michael Laurence Curzi

36N9 Genetics LLC

ORCID: 0000-0002-8309-1328

DOI: 10.5281/zenodo.22079583

A unified research bundle of seventy working papers.

Abstract of the Corpus

Stable arrangements persist even when they are poor ones. This corpus treats that persistence as a measurement problem, applying the apparatus of physics — potentials, first-order filters, fixed points, effective counts — to psychological and social systems across seventy working papers (1,836,200 words) in four series.

Series A — Foundations (6 papers) fixes the definitions: comfort as the felt total of present costs, so the present state is a local minimum of a comfort landscape by construction; the reflexive loop $T(x) = W(x, M(x))$, whose fixed points are self-fulfilling prophecies; and Banach's contraction condition, the one theorem the framework owns. **Series B — Engineering the Surplus** (33) builds the core: a storage law read as a leaky integrator, $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$, with fixed point $Q* = (\eta S - C)/\delta$; the headroom gauge $h = \eta lnN - \delta Q - C$; and an interaction surplus $F(u) = ln(1+(N-1)u)$ on non-parallelism $u = 1 - (x \cdot y)^2$, shown to be exactly a Fechner law with threshold $u_0 = 1/(N-1)$. **Series C — Frequency, Essence, and Mind** (11) derives a threshold $f_min = k \cdot e^{2c/\eta}$ over an essence law $f(n) = k \cdot n^2$, $k = (9/8)^2/\varphi^2 = 0.483425732988414$. **Series D — Structure and Resilience** (20) carries $HHI = \Sigma(share^2)$ with effective count $1/HHI$ into jurisdictional structure: a four-jurisdiction holding arrangement measures 3.41 effective jurisdictions naively, 1.00 once shares collapse onto coordination blocs.

The physics is used under a badging convention naming what is derived, what modelled and what postulated. The identification of quantum with anticipatory uncertainty is badged a postulate and priced against itself — decoherence in 1.98e-17 s against a 1 ms neural event — and the one storage law written at three scopes is badged model form only. Negative results are reported with the rest. Nothing rests on fieldwork: of 27 audited parameter-by-scope cells, 19 have an estimator, one a validated value, none a first-party measurement. The preregistered fit of the storage law failed its own second falsification rule, returning a 2.8 per cent held-out advantage over a matched AR(1) against a declared 10 per cent margin. The falsification suite is anti-conservative against itself: iid permutation false-rejects at 0.108 against a nominal 0.05. The corpus's own scale-type protocol demotes its moment equation $M = (D/P) \cdot E$ to an ordinal heuristic, and a 316-entry bibliography names ten central claims with no supporting literature — among them $N_eff = 1 + (N-1)u$, which is not a Hill number, and Series C's constant and square form.

Executive Summary

What this corpus is

Seventy working papers, 1,836,200 words, developing one framework and testing it. The organising question is why stable arrangements persist when they are poor ones, and the method is to treat that persistence as a measurement problem rather than a moral one. The apparatus is borrowed from physics — potentials, first-order filters, fixed points, effective counts — and applied to psychological and social systems.

How to read it

The four series are ordered by dependency, not by importance. **Series A** fixes definitions. **Series B** builds the storage law and the surplus functional and does most of the measurement work. **Series C** extends the apparatus to frequency and mind, and is the most exposed part of the corpus. **Series D** applies the concentration instrument to jurisdictional structure, and contains its best-grounded results.

Each paper stands alone: it carries its own abstract, hypotheses, method, results, verification, limitations and references. The papers also cite each other densely — 980 internal citations, every paper citing at least thirteen others and cited by at least eight, with any paper reachable from any other in at most three steps. A reader may therefore enter anywhere.

The badging convention

Every use of a physical result is labelled where it is made, as one of three things: a **derivation**, where the system genuinely satisfies the physical premises; a **shared mathematical form**, where the same equation governs both cases with no claim that the physics transfers; or a **postulate**, asserted by this work and carrying no external warrant. Most of the corpus is the second kind, which is respectable. The convention exists so that no claim of the second or third kind can borrow the authority of the first.

Each paper's *Verification* section decomposes its results into atomic postulates, each checkable on its own by arithmetic, by definition, or by a single citation, and then shows the compounding chain that produces the result. Every step carries its status. A chain resting mostly on definitions and postulates is identified as such in the section that presents it.

What the corpus does not establish

These are stated here, at the front, rather than left to be discovered.

- **No first-party measurement.** Of twenty-seven audited parameter-by-scope cells, nineteen have an estimator, one has a validated value, and none is a first-party measurement. Nothing here rests on fieldwork conducted for this work.
- **The preregistered fit failed its own rule.** The confirmatory test of the storage law returned a 2.8 per cent held-out advantage over a matched AR(1) against a declared margin of 10 per cent. The rule was set in advance and it was not met.
- **The falsification suite is anti-conservative against itself.** An independent-permutation control false-rejects at 0.108 against a nominal 0.05.
- **The corpus demotes its own moment equation.** Its scale-type protocol reduces $M = (D/P)$ $cdot E$ to an ordinal heuristic.
- **Two different effective counts travel under one name.** $1/$ $mathrm HHI$ and $exp(H)$ are Hill numbers of order two and one and are well grounded. $1+(N-1)u$ is a linear interpolation, is *not* a Hill number, and has no supporting literature.
- **Series C rests on postulates.** The essence law, its constant, and the extension of octave equivalence beyond pitch have no supporting literature. They are badged as postulates wherever they appear.
- **Quantitative structure is assumed, not tested.** The framework treats psychological quantities as having landscape structure. The conjoint-measurement conditions that would license that treatment were not tested.

The corpus's claim to credibility is not that it has no weaknesses. It is that it is the party that found them, and printed them.

Contents

Abstract of the Corpus ii

Executive Summary iii

I Foundations 1

1 Comfort as a Local Minimum: The Founding Definitions of the Creature Comforts Framework and the Ground-State Node 2

2 The Ledger, the Gradient, and the Arithmetic of Difference: Carrying Cost, Desire Lines, and the Surplus Functional f(u) = ln(1 + (N-1)u) 30

3 Weighting Futures by Resistance: The Comfort-Weighted Amplitude, Cue Placement, and the Resolution-Disturbance Trade-Off 60

4 The Spoken Word and the Returned Difference: Performative Commitment, Echo Reading, and the Doctrine of Category Violation 89

5 Convergence and Its Dark Twin: The Contraction Mapping, the Audit of Arrival, and the Consent Law of Closure 119

6 The Reference Apparatus of the Foundations Volume: The Closing Procedure, the Protocol Index, and an Audit of the Three-Class Honesty Ledger 143

II Engineering the Surplus 196

7 The Storage Law as a Leaky Integrator: Decay, Poles, and Why Zero Loss Destroys Arrival 197

8 Conversion, Coordination Cost, and Headroom: Why Eta Is the First Lever and Why the Optimum Group Is Smaller Than the Crossover 260

9 The Effective Count and the Logarithmic Ceiling: Auditing Roster Size Against Non-Parallelism 323

10 Crossing Versus Divergence: Decomposing Non-Parallelism, the Decay of u, and the Room as a Circuit Element 405

11 Passive Geometry and the Reverberant Tail: Crenellation, Parabolic Focus, and the Acoustic Exclusion of Quiet Participants 491

12 Thresholds, Substrate, and the Table With No Head: Three Built-Environment Levers Priced Against the Storage Law 553

13 Turn-Taking, Complex Contagion, and the Critical Boundary-Cost Multiplier: Three Channels into the Surplus Law 611

14 Order Without a Centre: Stigmergic Media, Independence Protocols, and the Institutional Maintenance Term 648

15 Surveying and Preregistering the Storage Law: A Measurement Protocol, and a Confirmatory Fit That Failed Its Own Baseline Test 720

16 A Gated Intervention Procedure and a Thirty-Two-Entry Failure Catalog for the Storage Law 775

17 Five Worked Systems and the Essence Bridge: Testing Scale-Freedom in the Storage Law and Identifying the Effective Count as a Pitch Index 848

18 Conversion Keys and Modal Design: Deriving a Room, a Cadence and a Pitch Lattice from an Essence Set Without Subliminal Signal 928

19 The Moment Equation and the Calibration Ratio: A Fixed Point at Contact, a Self-Imposed Demotion to Ordinal Scale, and the Price of Stored Error 1002

20 The Perceptual Lens: Fechner's Law as the Register of the Surplus Functional, with a Postulate Method and Its Held-Out Tests 1069

21 The Internal Surplus and the Self as Storage: Auditing One Observer's Framings, and the Stored-Error Track That Feels Identical From Inside 1131

22 The Launch Budget and the Desire Line: A Rectified Surplus Fund Against a Barrier, and Why Momentum Does Not Cross a Wide Shallow Ridge 1188

23 The Return Difference, the Contraction, and a Formulary That Demotes Five of Its Own Nine Entries 1246

24 The Dyad and the Group Moment: Opposition Is Parallel, Calibration Costs All the Headroom, and Distortion Does Not Average 1344

25 Trust, Reciprocity and Obligation as Quantities: Where Trust Enters the Storage Law, and Nine Ways the Internal/External Correspondence Breaks 1416

26 The Formula Validation Protocol, Scale Type, and Specific Objectivity: Demoting the Corpus's Own Moment Equation and Pricing Its Restoration 1491

27 The Falsification Suite Measured Against Itself: An Anti-Conservative Shuffle, a 52.7 Per Cent Demand-Characteristic Pass Rate, and a Programme With a 35.8

Per Cent Survival Probability 1565

28 Star Polygons, Cadence, and the Ledger: gcd(N,k) as Parallelism, Division of a Fixed Budget as a Theorem, and What an Eleven-Operation Audit Does Not Prove 1650

29 The Two Columns and the Concordance Register: Why Agreeing Traditions Stop Adding Evidence After the Second, and a Register That Rejects One Bridge in Five 1728

30 The Cell as Receiver, the Channel and Its Capacity, and Addressing Without Editing: Adequate Bandwidth Into an Absent Decoder 1805

31 Addressing the Non-Coding Genome: The Latency Constraint, the Missing Demodulator, and an Exact Sonification Pipeline 1863

32 Conjugate Uncertainty at Three Levels: The Kennard Bound, the Gabor Limit, and a Badged Information-Budget Analogue for Forecasts 1920

33 Quantum Formalism Without Licence: Born-Rule Statistics, Decoherence Timescales in Warm Tissue, and the Casimir Force at Biological Scale 1987

34 Belief Creates Reality If the Belief Is Grounded in Truth: A Conditional Postulate, Its Channel-Gain Form, and the Placebo Case 2043

35 Where Belief Fails: Four Failure Classes, a Priced Bank Run, and the Grounding Protocol That Kills Most Claims 2099

36 Narrative as Regulatory Input: A Fully Disclosed Sonification Chain, a Zero-Resultant Corpus, and a Preregistration That Prices Its Own Nulls 2173

37 Protocols at Two Scopes and a Finite Failure Catalogue: Headroom Gating, the Participation-Ratio Effective Count, and Thirty Manifestation Faults 2240

38 One Model Form at Three Scopes: The Instrument Audit, Five Breaks in the Correspondence, and a 156-Entry Status Register 2312

39 Four Structural Results and an Honesty Register: The Closing Synthesis and Reference Apparatus of Engineering the Surplus 2372

III Frequency, Essence, and Mind 2462

40 The Sustaining Frequency: f_min = k e^(2c/eta) as an Equivalence Between the Storage Fixed Point and the Essence Bridge 2463

41 The Minimum Sustainable Frequency: A Pitch Criterion for Relational Sustainability and the Count-Based Essence Index That Feeds It 2505

42 Octave Folding and the Odd-Part Class Rule: The Frequency of a Mood, of a Trait, and the Self as a Chord 2586

43 Consonance, Collision, and the Divided Self: A Firewall Between Harmonic Dissonance and Psychological Distress 2626

44 Three Routes to the Relationship Frequency: The Unison Result and the Elasticity-Reserve Split Between Levers 2697

45 Two Octaves per Doubling Against a Quadratic Cost: Roster Ceilings, the Parallel Wall, and When Tuning Must Halt 2735

46 Auditing a Traditional Operation Set Against a Parameter Model: Seven Maps, Two Splits, Three Rejections, and the Many-to-One Compound Rule 2806

47 Refinement as Contraction: Why Repetition Cannot Move a Fixed Point, and an Evidential Audit of the Tradition It Comes From 2875

48 A Disclosed Sonification Protocol for Sequence-Encoded Content, and the Missing Transducer That Makes It a Research Design and Not a Technique 2950

49 Destiny as a Fixed Point: Conservation of Drive Under Cost Transfer, the Completeness of the Lever Set, and the Corpus's Register of Demotions 3015

50 Reference Apparatus for the Frequency Framework: Constants, Lattice Tables, Measurement Protocols, and a Four-Drawer Classification of Evidential Status 3080

IV Structure and Resilience 3131

51 Correlated Jurisdictions: Coordination Perimeters, Erosion, and the Economics of Full Disclosure 3132

52 The Concentration Instrument: Bloc Collapse and the Effective Count of Jurisdictions 3181

53 The Apex Vehicle: Foundation Against Trust, and the Arithmetic of Forced Heirship 3228

54 Institutional Quality Without Independence: London, the Crown Dependencies, and the European Convergence 3278

55 Entity Maps Understate Dependency: The Dollar Axis, the Insurance Reserve, and a Scored Captive Rubric 3316

56 Substance Repricing and the Clearing-System Count: The Atlantic Cluster, ADGM, and the Limits of a Second Enclave 3359

57 Diversification That Does Not Pay: Gulf Additions, an Indifferent Bridging Slot, and a Capped Transit Position 3400

58 Pricing an Uncertain Bloc Call: AIFC Option Value, Thailand's Share of the Gain, and the GIFT City Corridor 3438

59 Currency Drag, Cost Ceilings, and Funding-Source Independence 3474

60 Compartments, Conversions, and the Confidence a Cell Wall Requires 3509

61 Measuring the Payment Rail: Clearing Perimeters, Transit Exposure Windows, and the Concentration Effect of De-Risking 3553

62 The Carve-Out Cascade: Pledge Triggers, Antitrust Walls, and Legal-Family Interfaces Priced as Cost 3589

63 The Assembled Architecture: Effective Blocs, Sanctions Deeming Rules, and the Disclosure Test in Historical Instance 3624

64 Five Houses, One Perimeter: Bloc Collapse in a Historical Network, the Completeness Gap of Double Entry, and Canonicalisation as an Integrity Control 3669

65 Canonical Form and Typed Parsing: Measured Integrity Controls Across a Financial Identifier Namespace 3719

66 Channel Separation as a Structural Property: In-Band Impersonation, Address Anonymity Sets, and the Legacy-Migration Asymptote 3762

67 A Concentration Atlas for Distributed Ledgers: Validator Herfindahls, the Utility Ordinal, and Dwell Time as the Freeze Variable 3797

68 Content Addressing, Time-Axis Concentration, and the Bounded Gain of a Nine-Capital Supplement 3829

69 What the Architecture Cannot Do: Correlated Backing Baskets, Gated Adoption Probabilities, and the Failure Record 3863

70 The Method in Seven Steps: A Consolidated Reproduction Workbook and Honesty Register for the Concentration and Headroom Instrument 3931

Part I

Foundations

Chapter 1

Comfort as a Local Minimum: The Founding Definitions of the Creature Comforts Framework and the Ground-State Node

Abstract. This opening instalment of the Foundations series states the framework's premises and fixes the definitions the later series cite. Comfort is redefined as the felt total of current costs rather than as pleasure, and a life's trajectory is modelled as a path through a comfort landscape governed by the same shape of law as stationary action. The physics is stated exactly: the action is stationary, not necessarily minimal, written $\delta S = 0$, and the variational form licenses no teleology. Change is modelled as gradient descent, $x_{n+1} = x_n - \eta \nabla f(x_n)$, with its known behaviours of local trapping and of oscillation under too large a step. Reflexivity is written as $T(x) = W(x, M(x))$, and the self-fulfilling prophecy is identified as a fixed point of that map; the Banach condition $d(T(x), T(y)) \leq k\, d(x, y)$ with $0 \leq k < 1$ is stated as the one theorem the framework owns outright, together with the two premises, completeness and a uniform modulus, that it does not establish. The central claim, that the present moment is always a local minimum, is presented as a mechanism rather than a verdict, and departure is priced at the height of the nearest ridge rather than the depth of the current basin. The framework's single postulate, that the uncertainty of quantum measurement and the uncertainty of a foreseen future are the same kind, is stated and explicitly badged as a postulate; $\sigma_x \sigma_p \geq \hbar/2$ is given in the Kennard–Robertson form, as a fact about state preparation rather than about measurement disturbance. Three disanalogies between a person and a physical minimiser are conceded on the record, four hypotheses are stated in falsifiable form, and none of the four is tested here.

Keywords: stationary action; comfort landscape; gradient descent; reflexive loop; fixed point; self-fulfilling prophecy; postulate badging; local minimum

Series: A. Foundations **Paper:** CCP-A01

1.0.1 1. Introduction

1.1 The question

A person's circumstances are usually explained in one of two registers. Either they are the product of choice, in which case a poor position is a failure of will; or they are the product of

circumstance, in which case a poor position is a failure of luck. Both readings are common, both are consequential, and neither is a mechanism. A mechanism would have to say what quantity is being optimised, what the optimisation does at its resting points, and what it would cost to leave one.

This paper asks whether a person's stable position can be modelled as a local minimum of a scalar cost surface, and whether that model licenses a practical intervention. It is the opening instalment of the Foundations series of *Psychological and Sociological Physics: The Creature Comforts Corpus*, and it fixes the definitions on which the remaining sixty-nine papers draw. Thirty-seven of them cite the results established here. Because those results are load-bearing elsewhere, this paper states them in a form intended to be quoted rather than paraphrased, and reports its own weaknesses at the same precision as its claims.

1.2 What this paper contributes

Seven results are established or fixed here, labelled R1 through R7 and restated verbatim in the Conclusion.

R1 redefines comfort as the felt total of present costs rather than as pleasure, with the consequence that the present state is a local minimum of the comfort landscape by construction rather than by luck. **R2** writes the reflexive loop as $T(x) = W(x, M(x))$ and identifies the self-fulfilling prophecy as a fixed point of that composition of seeing and responding, while recording that the specification is not identified from observational data and is therefore a modelling position rather than a measured effect. **R3** states the principle of stationary action as $\delta S = 0$, with the correction that the honest form is stationarity and not minimality, and with the further statement that the variational form licenses no teleology. **R4** states the Banach condition $d(T(x), T(y)) \leq k\, d(x, y)$ with $0 \leq k < 1$, from which existence, uniqueness and geometric convergence from any starting point follow. **R5** states the framework's single postulate — that the uncertainty of quantum measurement and the uncertainty of a foreseen future are the same kind of uncertainty — and badges it as a postulate in the same sentence, giving $\sigma_x \sigma_p \geq \hbar/2$ only in its correct state-preparation reading. **R6** concedes three disanalogies between a person and a physical minimiser: noise, plasticity of the landscape, and response to imagined costs, the last of which physics has no analogue for. **R7** prices departure at the height of the nearest ridge rather than the depth of the current basin, which distinguishes conditions that want engineering from conditions that want acceptance.

A counter-result is stated with equal weight: the identification of a life with a minimising physical system is a postulate and not established physics; three disanalogies break the analogy; and locating oneself at a local minimum is an exoneration rather than an endorsement, because *local* was never a compliment.

1.3 Register, sources, and what is not reproduced here

The source manuscript is written in two registers, a narrative one and an analytic one, deliberately braided. This paper reports the analytic register. The narrative material is not reproduced: it is a fictional exposition and carries no analytical claim that the treatise does not also carry. Where the source's ceremonial vocabulary names a construct — *the ground state, the still water, the seal* — the construct is renamed to what it is, and the ceremonial name is reported once, as an object of study, because the framework's use of a redundant second encoding is itself one of its methodological claims and is examined in Section 8. The framework's narrative register has a respectable precedent in the operations literature, where the same choice was made for the same stated reason [Goldratt and Cox, 1984].

The source also contained an audience statement naming a particular professional community. That framing has been removed and is not reproduced. It was inessential to every claim the paper makes, and its removal changes no result.

1.0.2 2. Hypothesis

The framework advances a small number of claims that could be wrong, a larger number of definitions that could not, and exactly one postulate. Keeping the three apart is the framework's own stated discipline and this section enforces it; Figure 1 is that discipline's ledger, and the single postulate is the only entry in it that this work owns. A hypothesis that nothing could falsify is restated below as a definition or a postulate, and labelled as such [Popper, 1959].

H1 — the ground-state hypothesis. For a person under stable constraints, the occupied state is a local minimum of a scalar cost surface defined over immediately reachable alternatives: every small departure presents, to the whole weighing system rather than to deliberate reason, as a net cost. *Predicted:* if the cost of each adjacent alternative could be elicited, no adjacent alternative would carry a lower total than the occupied state, and this would hold irrespective of the occupant's reported satisfaction with the occupied state. *Falsified by:* a person stably occupying a state while a reachable adjacent state carries a lower elicited total cost, and no barrier separating them. *Status:* testable in principle; not tested here. The elicitation instrument does not exist, for reasons set out in Section 10.2.

H2 — the ridge hypothesis. Persistence in a state is governed by the height of the cheapest barrier out of it, not by the depth of the state itself. *Predicted:* two occupants of basins differing several-fold in depth but sharing a barrier height will show comparable persistence; two occupants of equally deep basins with different barrier heights will not. *Falsified by:* persistence tracking depth and not barrier height across matched cases. *Status:* testable; not tested here. Figure 3 is an illustration of the prediction, not a measurement of it.

H3 — the gradient hypothesis (the framework's practical claim). A concretely seen alternative of lower cost re-grades the surface, and relaxation toward it proceeds without an additional effort term. *Predicted:* an intervention that makes a lower-cost alternative concretely visible produces movement toward it; an intervention that applies pressure toward the same alternative without making it visible does not, and at sufficient pressure produces movement away. *Falsified by:* movement proportional to applied pressure and independent of the concreteness of the seen alternative. *Status:* testable; not tested here, and the second limb — movement away under pressure — has no supporting entry in this corpus's external bibliography. It is reported as an uncited empirical claim in Section 10.3.

H4 — the step-size hypothesis. Change attempted in corrections larger than the discrepancy that prompted them oscillates or diverges rather than settling. *Predicted:* the sign pattern of successive corrections alternates before amplitude grows, so alternation is an early warning available before failure. *Falsified by:* monotone convergence under corrections systematically larger than the observed discrepancy. *Status:* mathematically true of the modelled recursion (Figure 2); the transfer to human correction is a shared mathematical form, not a derivation, and is treated as such in Section 9.

D1 — comfort (definition, not hypothesis). Comfort is the felt total of present costs: effort owed, dangers watched, dreads carried, identities maintained, options held open. Pleasure is one term in that total and not the total. This is a stipulation. It cannot be falsified; it can only be judged useful or useless, and its measurement-theoretic status is audited in Section 9.2.

D2 — the reflexive loop (definition). $T(x) = W(x, M(x))$, where M is the occupant's model of the situation and W is the world's response to an action taken on that model. A self-fulfilling prophecy is a fixed point of T. This is a definition and a naming, not a claim; the sociological phenomenon it names was described and named long before [Merton, 1948].

P1 — the Postulate (postulate, not hypothesis, and not promoted to one). Observer participation in the reflexive loop of prediction is structurally the same phenomenon as observer participation in quantum measurement. Nothing in physics licenses this. It is offered as a correspondence of structure, adopted as working doctrine, and defended only by its fruits. It is stated in full and audited in Section 6, where the reasons it is not an H-numbered hypothesis are given.

1.0.3 3. The comfort landscape

3.1 Comfort is the felt total of costs, not pleasure

The ordinary picture of a comfort zone is a fence: the zone is a place, the occupant is inside it, the good things are outside, and growth is the act of leaving. The framework replaces the fence with a floor. The comfort zone is not a place a person is in; it is the place a person is. It travels with the occupant, because the occupant — like any system that settles — has already come to rest in the least costly arrangement available under current constraints, continuously and without deliberation.

That reframing only works once the sentiment is drained from the word. Under D1, comfort is the running total of what everything currently charges. Pleasure is one small bright line in that accounting; the accounting runs day and night whether or not anything is pleasant, and it is the accounting that steers. A person in a grey job, a stale room, or an aching accommodation of self to circumstance is not outside their comfort zone. They are at the bottom of it, and the greyness is what this particular bottom feels like. *Comfortable* and *uncomfortable* are therefore not synonyms for *happy* and *unhappy*, and the framework concerns only the first pair.

This is the founding reframing, R1. It is a definition and it does the work a good definition does: it makes a previously moralised description into an addressable one.

3.2 Stationary action, stated exactly

The framework borrows its governing image from classical mechanics, and the borrowing must be stated precisely because it is stated imprecisely almost everywhere else.

Of all the trajectories a system might trace between two configurations, the trajectory actually taken is the one for which the accumulated quantity called the action is *stationary* — unchanged to first order under small variations of the path. In symbols, $\delta S = 0$. The physical statement is stationarity, not minimality. For many ordinary systems the stationary path is in fact a minimum, which is why the principle is usually introduced as least action and why the intuition of a system taking the cheapest route is a serviceable working intuition; but the honest form is stationarity, and this paper uses the honest form. Light entering water bends at the angle that makes its travel time stationary; a chain hung from two points finds the catenary; a soap film spanning a bent wire solves a minimisation problem overnight. None of these systems knows anything, and that is the point: least-cost paths are not chosen by minds, they are what settling is.

Two further statements belong with this one and are R3's remaining content. First, the variational form licenses no teleology. That a system's realised path makes the action stationary is a statement about the structure of the equations of motion, not a statement that the system is aiming at anything, and any reading in which a trajectory is drawn toward a goal is an importation the mathematics does not carry. Second, the framework's application of this shape to a life is an application of *shape*. The source says so on the page at every occurrence, and this paper preserves that discipline; the analogy-category audit is in Section 9.1.

3.3 Local minima and gradient descent

Below the physics there is only calculus, and the calculus is checkable with kitchen equipment. A marble in a tipped bowl comes to rest where every direction it might roll is uphill. The mathematician calls such a point a local minimum, and both words carry a limit. *Local:* the bowl may sit on a table and the table in a house on a hill; the marble rests where it rests not because nowhere is lower but because every path away begins with an ascent. *Minimum of what:* of the potential surface as it stands, not of the surface the marble might prefer.

Once comfort is pictured as a surface — high ground costly, low ground cheap — one of the best-understood procedures in applied mathematics is imported whole. From wherever

claim	established in
Stationary action $\delta S = 0$	physics
Minima and gradients	calculus
Contraction mapping theorem	mathematics
Judgment and habit findings	psychology
Measurement uncertainty ≡ foresight uncertainty	this work

4 established · 1 badged POSTULATE

Figure 1.1: The framework's own inventory of formal claims: four are established outside this work, and exactly one — the correspondence between measurement uncertainty and the uncertainty of a foreseen future — is badged as this work's postulate. The counts are the source's, which states "the complete inventory of the book's formal claims ... There are five". Schematic; no measured quantity.

you are, step in the direction that most steeply reduces the cost, and repeat:

$$x_{n+1} = x_n - \eta \nabla f(x_n)$$

where f is the cost surface, ∇f its local slope, and η the step size. Every well-characterised behaviour of this procedure describes a way that human change succeeds or fails.

It settles into whatever local minimum it happens to be near, not the deepest on the map — which is why a person can be stuck in a life that is merely locally comfortable while a far better minimum sits one ridge over, invisible from where they stand. It moves without any effort term, purely by slope — which is why the framework's central practical claim (H3) is that *seeing* a nearer, deeper comfort is itself the intervention: the lower ground has to be made visible, and relaxation does the rest. And if the step size is wrong — too large, too forced — the procedure overshoots and oscillates instead of settling (H4, Figure 2). The oscillation regimes of a first-order recursion are standard [Elaydi, 2005, Kelley and Peterson, 2001]; the basin structure that makes the local/global distinction consequential is equally standard [Strogatz, 2015]; and the gain-and-stability vocabulary in which the failure is usually diagnosed is the control engineer's [Åström and Murray, 2008].

3.4 The present state is a local minimum, by construction

What survives the caveats is the ground-state result, R1 in its operational form: **the present moment is always a local minimum of the occupant's comfort landscape.** Wherever a person now sits — this chair, this job, this city, this silence toward a sibling, this hour of reading instead of any of the ten thousand other things — they sit there because, under the constraints and prices as their system has weighed them, every small departure has presented as a net cost. Not to deliberate reason; deliberate reason is a late committee that reviews decisions mostly after they ship. To the whole weighing.

The argument is symmetrical and short: had any adjacent arrangement been reliably cheaper *as the system prices it*, the person would already be there. The evidence is the ordinary evidence of habits never decided on and identities never chosen. The settling never stops.

The result reads as an accusation only if it is misread by exactly half. It is a mechanism, and a mechanism cannot accuse. What it can do is exonerate — of the specific charge that

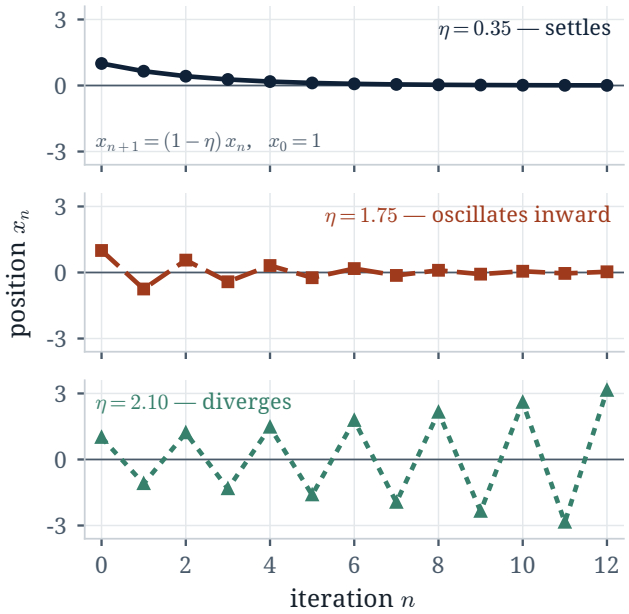


Figure 1.2: Step size, not intention, decides whether descent settles. Modelled, not measured: cost surface $f(x) = x^2/2$, so the descent recursion reduces exactly to $x_{n+1} = (1-\eta)x_n$, with $x_0 = 1$ and $\eta \in \{0.35, 1.75, 2.10\}$. Below $\eta = 2$ the recursion converges; at $\eta > 2$ it diverges while remaining perfectly well intentioned. Script: `figs/src/CCP-A01_fig02.py`.

one's circumstances are a failure of will. If effort was never the engine, its shortage was never the fault. But the exoneration is not a blessing. A local minimum can be shallow, grey, and very far from the lowest place the surface holds. The framework does not teach that where a person is, is good. It teaches that where they are is *ground* — real, load-bearing, arrived at lawfully — and that no intervention can be launched from anywhere else, for the same reason that no navigator can plot a course from a position they refuse to occupy.

Two boundary readings are refused explicitly, because both are available and both are wrong. The result is not a doctrine of desert: to say that a position is a lawful minimum of a surface says nothing whatever about who graded the surface. Some surfaces are carved by their occupants; many are carved by birth, by policy, and by other people's decisions arriving like weather, and a creature settling lawfully into terrain it never chose is not thereby endorsing the terrain. The mechanism describes; it does not justify. Nor is the result contentment teaching: the traditions that counsel wanting what one has aim to *deepen* the occupied minimum, whereas this result aims only to *locate* it, and is indifferent to the verdict that follows.

3.5 Departure is priced at the ridge (R7)

The question every grey position eventually raises — *if I dislike it here, why do I stay?* — has a mechanical answer that does not require a defect of character.

Departure is priced at the rim, not at the destination. What holds an occupant is rarely the depth of the minimum but the height of the nearest ridge, and ridges are made of very ordinary material: a pension vesting in four years, a lease, an ageing parent in the next town, the specific dread of being a beginner again at forty, an identity that would have to be buried with honours. A shallow minimum ringed by high ridges holds exactly as firmly as a deep one, and feels entirely different from inside. That difference — held versus home — is worth locating precisely, because the two call for different work. The deep minimum wants acceptance. The shallow one ringed with ridges wants engineering, and engineering, unlike exhortation, begins with an accurate map of the terrain as found.

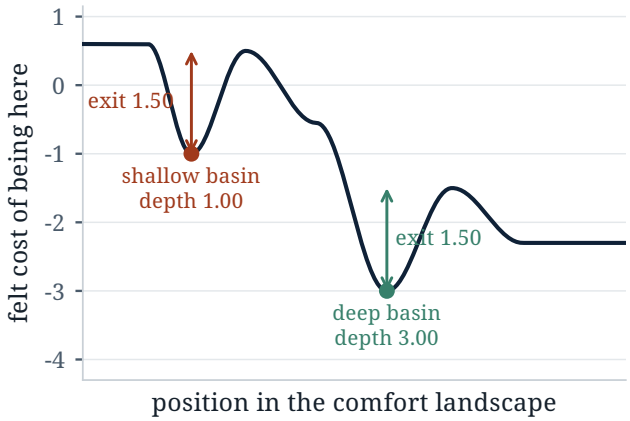


Figure 1.3: Persistence is set by barrier height, not by basin depth. Modelled, not measured: a one-dimensional surface laid out through explicit control points and smoothed by raised-cosine interpolation so that two basins differ threefold in depth (1.00 against 3.00 cost units) while sharing an identical cheapest exit of 1.50 units. Both barrier figures are computed from the plotted curve. Script: `figs/src/CCP-A01_fig03.py`.

The same geometry supplies the framework's diagnosis of professional over-specialisation. A silo is a local minimum: deep because years of training and reward have carved it, narrow because it admits one vocabulary and one set of tools, and walled by comfort rather than by ignorance. The specialist could learn the neighbouring domain; what the specialist cannot easily do is want to, in the moment, against the gradient, because every direction out is uphill and gradient descent does not go uphill. This is why exhortation fails, and it is the mechanism by which organisational learning destroys the variance that produced its own advantage [March, 1991]. The prescription that follows is not force but terrain: the connecting pattern is to be made the cheaper place to stand, so that migration happens by the same gradient that produced the trap. The limit of that prescription is real and is stated in Section 8.3.

3.6 Noise, and the one place it helps

The first of the three conceded disanalogies (R6) has a use before it has a cost. A person is noisy: they jitter, err, and are jostled, and the jostling sometimes throws them over ridges no gradient would carry them across.

Metallurgy names the corresponding technique. A metal cooled quickly freezes into whatever strained arrangement it happened to occupy; cooled slowly, with heat enough to keep the atoms jostling, it finds its way over small ridges into deeper minima and comes out stronger. Search procedures borrowed the trick whole: add noise, then quiet it gradually, and the search escapes shallow traps it would otherwise die in. The framework's claim is that disrupted seasons of a life — the move, the loss, the enforced interval — function as that heat: positions held for decades come loose, ridges that were walls become steps, and the settling that follows can reach hollows the settled years could not.

This paragraph carries no citation, deliberately. Simulated annealing is standard, but it has no entry in this corpus's external bibliography (Section 10.3), and the transfer from an annealing schedule to a disrupted life is in any case a shared mathematical form at best. It is reported here as the framework states it, and flagged.

1.0.4 4. The reflexive loop (R2)

A person is not a ball on a fixed surface, because a person carries a model of the world, acts on the model, and thereby changes the world, which changes what the next model should be.

Write M for the occupant's model — their read of the situation — and W for the world's response to an action taken on the basis of that model. Then the next state of the system is

$$T(x) \;=\; W(x,\, M(x))$$

The model reads state x and prescribes an action; the world processes the state together with that action and returns the next state. T is the composition of seeing and responding (Figure 4). This is D2, and it is a definition.

The structure is not exotic. It is the everyday structure of markets, reputations, deterrence, morale, and every situation in which what people believe will happen helps decide what happens. The self-fulfilling prophecy is exactly a fixed point of such a loop: a belief that, once acted upon, brings about the conditions that make it true [Merton, 1948]. Its companion — purposive action producing the consequence nobody modelled — is the same machinery read from the other end [Merton, 1936]. The general programme of treating one mathematics of feedback as governing both machines and organisms is older still [Wiener, 1948]. The move of modelling a human operator inside a control loop as a transfer function is older than every claim in this section.

Two exposures attach to the loop and are recorded here rather than in Limitations, because they qualify the definition itself. First, a loop whose state variable is also its own published indicator degrades: a measure that becomes a target ceases to be a good measure, and any instrument built on this framework is exposed to that [Campbell, 1979]. Second, the loop's W term includes the observer's own presence: expectancy effects on the observed are real as a design threat even where their magnitude is disputed [Jussim and Harber, 2005].

A third exposure is decisive, and it is stated here, at the point of introduction, rather than deferred to a footnote. Written as an aggregate relation — a group's state influencing its members while the members constitute the group — $T(x) = W(x, M(x))$ is a reflexive, linear-in-means specification, and the endogenous effect in such a specification is **not identified from observational data**: it cannot be separated from contextual effects or from correlated unobservables, however much data is collected [Manski, 1993]. No quantity of fit, held out or otherwise, recovers the parameter. The network route is foreclosed by a companion result: influence and selection are generically confounded in observational network data, so resemblance between interacting parties is not evidence that the interaction produced it [Shalizi and Thomas, 2011]. **R2 is therefore a modelling position, not an identified empirical result.** This paper adopts no identification remedy; Section 9.4 names the remedies that exist, prices each, and says why none is claimed here.

1.0.5 5. The contraction condition (R4)

This is the one theorem the framework owns outright, and it is stated with its premises visible because the premises are where the transfer succeeds or fails.

Take a complete metric space (X, d): a set of states with a notion of distance, in which every sequence that bunches up ever more tightly converges to a state that is actually in the space. Take a map $T : X \to X$. Call T a **contraction** if there is a constant k with $0 \le k < 1$ such that for every pair of points $x, y \in X$

$$d(T(x),\, T(y)) \;\le\; k\, d(x, y)$$

Then Banach [1922] guarantees three things, as proven fact rather than as tendency. First, T has a fixed point: a state x^* with $T(x^*) = x^*$. Second, that fixed point is unique. Third, it is

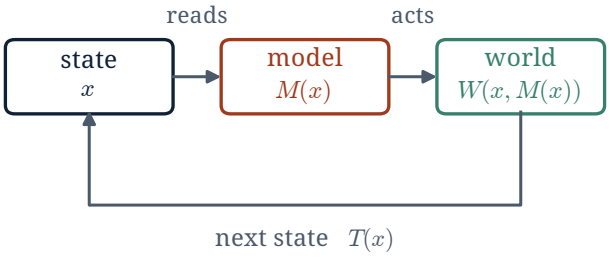


Figure 1.4: The reflexive loop and the two conditions that close it. Schematic; no plotted quantity. The composition $T(x) = W(x, M(x))$ is the paper's definition D2; the fixed-point equation and the contraction condition are stated beneath it as the two separate things they are — one an equation a state may satisfy, the other a property the map may have. Script: `figs/src/CCP-A01_fig04.py`.

reached from anywhere: starting the iteration at any point whatever and applying T repeatedly, the sequence converges to x^*, and it converges geometrically, the remaining distance shrinking by at least the factor k each pass.

Uniqueness is worth seeing done, because it takes one line and it is the part operators disbelieve. Suppose x^* and y^* are both fixed. Then $d(x^*, y^*) = d(T(x^*), T(y^*)) \le k\, d(x^*, y^*)$, so $(1 - k)\, d(x^*, y^*) \le 0$; since $k < 1$ this forces $d(x^*, y^*) = 0$, and the two points were one point.

Read in the framework's terms: if the reflexive loop of a life can be made into a contraction — if each pass through *see, act, the world answers, see again* reliably shrinks the distance between where the occupant is and a particular foreseen state — then that foreseen state is not a hope. It is a fixed point; it is unique; and it will be reached, from wherever the occupant happens to start, by iteration. That is the object the framework teaches the construction of, primarily upon oneself. The instrument in its exact form, together with the a posteriori bound that makes it self-reading, is the subject of the companion paper: the statement that a contraction with modulus $q < 1$ on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least q each pass [Curzi, 2026].

What is a theorem here, and what is not. The theorem is a theorem. The claim that a human life *can be arranged* to satisfy its premises is the framework's, and the premises are exactly where the difficulty lives. Two of them are load-bearing and neither has been demonstrated for this application:

- *Completeness.* The state space must contain the limits of its own Cauchy sequences. Nothing establishes that a space of life-states does.
- *Contraction with a uniform modulus.* There must exist a single $k < 1$ bounding the shrinkage for **every** pair of states, not merely for the pairs an operator happens to visit. Nobody has measured such a k for any human loop.

Where the modulus is one or greater the guarantee does not weaken, it inverts: unit gain gives an isometry, which orbits without arriving, and gain above one gives an amplifier, in which corrections sized to delayed feedback arrive as new errors. That is the same failure the step-size hypothesis H4 describes, and it is why this paper treats the contraction as a *condition to be checked*, not a property to be assumed. The corpus applies the same condition to its storage law, where the decay rate plays the role of $1 - k$: the map $Q_{t+1} = (1 - \delta)Q_t + \eta S_t - C_t$ is Lipschitz with modulus $1 - \delta$, so $\delta > 0$ is precisely what makes it a contraction, and $\delta = 0$ is shown to be a pathology rather than an ideal — no contraction, no unique fixed point, and an

error-persistence column of all ones, so that no mistake ever decays out of the stock [Curzi, 2026a]. The corresponding fixed-point elasticities are measured elsewhere in the corpus: η at $+1.594$, δ at -1.000, u at $+0.770$ and C at -0.594, so that a ten per cent efficiency gain is worth 2.7 times a ten per cent cost cut [Curzi, 2026q]. The contraction constant $1-\delta$ and the exact error law $|Q_t - Q_\infty| = (1-\delta)^t |Q_0 - Q_\infty|$ are verified to a maximum residual of 2.0624×10^{-15} over a hundred passes [Curzi, 2026].

1.0.6 6. The Postulate, stated and badged (R5)

6.1 The uncertainty relation, stated honestly

For a quantum state, the product of the standard deviations of two conjugate observables — position and momentum being the canonical pair — is bounded below:

$$\sigma_x \, \sigma_p \;\geq\; \frac{\hbar}{2}$$

This is the Kennard–Robertson form [Kennard, 1927], and the point the popular accounts get wrong must be got right here and in every later occurrence. **It is not, at its core, a statement that a clumsy measurement disturbs a delicate system.** There is a related and real measurement-disturbance effect, and it is a different statement. This one is about the *state itself*: position and momentum are Fourier conjugates, and no state — measured or unmeasured, observed or alone in the dark — can be sharply defined in both at once, for the same reason that no signal can be simultaneously a sharp spike in time and a single pure frequency. The spread is intrinsic. It is a property of what a quantum state is, not of what an instrument does to it.

The sloppy version — *the observer disturbs the observed* — is exactly the version a careless argument would lean on. This paper states the strict version, uses only the strict version, and does not build on the sloppy one.

6.2 The Postulate

There is a second kind of uncertainty, wholly outside physics, in the reflexive loop of Section 4: a future that is affected by being foreseen is, for that reason, genuinely uncertain in a way that has nothing to do with ignorance and cannot be removed by better information. Measuring the position of a particle sharpens position at the expense of momentum. Foreseeing a future, and acting on the foresight, alters the future foreseen. In both, the act of observation participates in the thing observed.

The Postulate (P1), badged in the sentence that states it: it is postulated — not derived, not demonstrated, and not licensed by physics — that these two uncertainties are the same kind of uncertainty; that observer participation in the reflexive loop of prediction is structurally the same phenomenon as observer participation in quantum measurement.

The status of that claim is stated exactly. It is **not** a theorem. It is **not** established physics. It is **not** a claim that the mind is a quantum system, that neurons compute with wavefunctions, or that consciousness collapses anything; the framework makes none of those claims and would be wrong if it did. It is a proposed correspondence of structure between two reflexive phenomena arising in wholly different substrates. It may be accepted, held at arm's length as a suggestive analogy, or rejected outright, and the practical framework stands either way — because everything the framework actually asks a practitioner to do is licensed by the reframing of Section 3, the gradient descent of Section 3.3, the loop of Section 4, and the theorem of Section 5, none of which need the Postulate at all.

6.3 Why the Postulate is not an H-numbered hypothesis

It is not listed with H1–H4 because no observation distinguishes the world in which it holds from the world in which it does not. A structural correspondence between two mathematical descriptions is a fact about the descriptions, and asserting that the underlying phenomena are therefore of one kind adds nothing an experiment could detect. Under a strict falsifiability criterion the Postulate is not a scientific claim at all [Popper, 1959]. The framework's own account of what it is doing is better rendered in the vocabulary of research programmes: the Postulate is a hard-core commitment that is protected rather than tested, and its warrant is the productivity of the belt of testable claims around it [Lakatos, 1970].

The hazard the badge exists to prevent is well documented. A physics result cited outside the conditions of its own derivation transfers only its authority and none of its content — a habit whose clearest example is the traversable-wormhole literature, which is routinely invoked as support for shortcuts between distant states when its own conclusion is that such a shortcut requires exotic matter violating the averaged null energy condition [Morris and Thorne, 1988]. That is precisely the failure the badge is designed to make impossible here.

1.0.7 7. The three conceded disanalogies (R6)

The framework's honesty condition requires that the ways in which a person is not a marble be listed rather than left to a reviewer. Three are conceded on the record.

7.1 Noise. A person jitters, errs, and is jostled, and the jostling sometimes carries them over ridges no gradient would cross. A marble in a still bowl does not. The use of this disanalogy is Section 3.6; its cost is that no trajectory prediction from the model can be deterministic, and that the model therefore predicts distributions at best, and in its present state does not predict those either.

7.2 Plasticity. A person's surface re-grades itself as they learn, which no bowl's does. This is not a small correction. It means the objective function is non-stationary, that a minimum located today may not be a minimum next season, and that adaptation moves the reference level against which cost is felt. The empirical form of that plasticity is well known: subjective states decay toward a set point after large changes of circumstance [Brickman et al., 1978], though the strong reading of that finding has been substantially revised — set points move, adaptation is incomplete, and individuals differ [Diener et al., 2006]. The framework's practical claim depends on the revised reading and would be defeated by the strong one, and that dependency is stated here rather than discovered later.

7.3 Foresight. A person responds to costs that do not exist yet — costs they have merely imagined — and this is the one term for which physics has no analogue at all. A creature that responds to imagined costs is a creature whose surface can be re-graded by an act of imagination. This is not a flaw in the analogy; it is the entire remainder of the framework, and it is what makes H3 a claim rather than a restatement.

1.0.8 8. The framework's own apparatus

8.1 The twelve constructs

The framework is organised as twelve numbered constructs in four groups of three. The table is reproduced because it is the framework's index and because later papers address constructs by number. The right-hand column records the correspondence to an earlier ten-part scheme from which the framework descends; two constructs have no ancestor there, and the table says so rather than inventing one.

No.	Construct	Function (operational)	Function (ceremonial register)	Earlier scheme	Group
00	STILLPOINT	The occupant is always at a local minimum: the present state is the comfort zone by definition, not by luck	The ground state; the water at rest	— (first assumed node)	I
01	LEDGER	Honest inventory of current comforts and their actual costs	The test of the stone	KIN.0 VALIDATE	I
02	GRADIENT	Comfort differentials; least-action paths; desire lines	Water does not negotiate	— (second assumed node)	I
03	HORIZON	Nonlinear extrapolative reasoning; branch-tracing; surplus of non-parallel sightlines	The many windows of the one tower	KIN.1 TRANSMED	II
04	WAVE	The future as a comfort-weighted amplitude; reading amplitudes	The sea before the wind chooses	KIN.3 MYTHOS	II
05	LURE	Placing the idea of greater comfort where the self will encounter it; cues, triggers, named days	The baited horizon	KIN.4 FORGE	II
06	UNCERTAINTY	The Postulate: observation participates; trade-offs of knowing	The veil that measures the measurer	KIN.6 CHECK	III
07	COLLAPSE	Commitment as measurement; the spoken prediction	The word that closes the sea	KIN.2 KERNEL	III
08	ECHO	The world answers; reflexive loops; first return signal	The voice that comes back changed	KIN.8 MULTIPLAY	III
09	RECURSE	Iteration; drift correction; the spiral, not the circle	The stair that climbs itself	KIN.5 RECURSE	IV
10	FIXEDPOINT	The self-fulfilling prediction as the fixed point of a contracting loop	The room that was always furnished	KIN.7 ETHIC	IV
11	SEAL	Closing the working; consent; when not to run the procedure at all	The seal; the vow	KIN.9 SEAL	IV

Group I establishes position and asks for no change; that restraint is calibration, not gentleness. Group II trains the seeing. Group III is the act, and it opens with the Postulate stated as a postulate, so that the act is not performed by anyone who has been permitted to mistake the Postulate for a theorem. Group IV is the finishing: iteration, arrival, and closure. The ethics are stated last not because they matter least but because a constraint stated before capability is a mood, and stated after capability is a commitment.

The two constructs without an ancestor are the framework's first teaching in miniature. The earlier scheme began at validation because it assumed that the ground an operator stood on and the slope that ground made under their feet were simply there: conditions, not workings. The framework's finding is that both are running processes, and so the count became twelve.

8.2 The redundant second encoding, examined

The framework states each item three times: as mechanism, as mathematics, and as rite. The third statement adds no information. The framework's stated reason is that a claim held in one form is held in one way, and that the second encoding survives conditions the first does not. This paper reports that as the framework's claim and does not endorse it, because no entry in this corpus's external bibliography establishes it (Section 10.3). What can be said is narrower and is worth saying: the practice of pairing every prescription with a named failure mode has a well-developed precedent in design method, and the practice of delivering a procedural argument in narrative form has a precedent in operations [Goldratt and Cox, 1984]. Neither precedent establishes a transfer advantage. Both establish that the choice is not eccentric.

The ceremonial register is analysed here, not performed. Where the source sets a mnemonic in Latin — *quies fundamentum, non carcer*, rest is the foundation and not the cage; *aqua in fundo quiescit*, the water rests at the bottom — these are recorded as the framework's own tags for its results, and the second is the ceremonial name for R1.

One substantive distinction survives translation and is not ornamental. The framework distinguishes *resignation* from *accuracy*. Resignation is a verdict about the future rendered from the present. Accuracy declines to be wrong about the present and renders no verdict about the future. The two are confused by nearly everyone at first, and the operational difference is testable in principle: when circumstances change, the resigned are surprised, having filed the old level as permanent, and the accurate simply re-level, having filed nothing.

8.3 The reflexive loop at institutional scale

The framework applies D2 above the level of the individual, and the application is worth stating because it is where the loop does visible work and where its limits are sharpest.

Four positions are advanced. First, self-sovereignty is not the absence of accountability but its internalisation: the sovereign keeps their own books, and the honesty must be greater precisely because no external party is auditing. Second, sovereignty does not abolish organisations but re-situates them — their function shifts from governing subjects to coordinating nodes, which is a more demanding job, not a lesser one; the empirical literature on durable common-pool governance describes exactly such coordination and supplies design principles derived from cases rather than postulated [Ostrom, 1990]. Third — the point on which the doctrine turns — sovereignty becomes a threat by one path, suppression: a power fears that autonomy is dangerous, acts on the fear, the action produces resistance, the world returns a threat that is now real, and the power reads the returned threat as confirmation. That is $T(x) = W(x, M(x))$ written at institutional scale, and it is the plain self-fulfilling prophecy of Section 4 doing real explanatory work [Merton, 1948]. Fourth, the mature posture is therefore integration rather than secession, and integration means the deliberate maintenance of a complementary, non-parallel pair whose difference is preserved rather than dissolved.

The fourth position rests on an arithmetic the framework uses throughout: several non-parallel viewpoints yield more together than any one alone, and the surplus is a function of their *difference*, vanishing when the viewpoints are identical, and growing only logarithmically thereafter. The framework states the diminishing returns rather than overselling the effect, and this paper states the literature on both sides, because the evidence is genuinely two-sided.

The functional form is underdetermined, and this paper says so rather than asserting one. The framework's surplus is logarithmic in the count. Two established alternatives fit the same qualitative facts equally well: social impact grows as a power of the number of sources with an exponent below one, established with psychological data [Latané, 1981], and network value has been argued to grow as $n \log n$ rather than as any of the usual candidates [Briscoe et al., 2006]. All three are decelerating functions of the count and nothing in this paper discriminates between them. **A paper asserting the logarithmic form owes a discriminating test, and this**

one does not have one.

What can be said in the logarithm's favour is a warrant that does not depend on the perceptual story at all: the growth rate of a multiplicative process is logarithmic, so a logarithmic form can be a dynamical consequence rather than a psychological assumption [Peters and Gell-Mann, 2016]. Where this corpus justifies a logarithm on Fechnerian grounds [Fechner, 1966], that warrant is the weaker of the two, because it requires the perceptual claim to be true. Where the quantity in question is generated multiplicatively, the dynamical warrant is available and should be preferred. This paper does not establish that the comfort surplus is generated multiplicatively; it records which argument would be needed.

Supporting: novel information arrives across weak, dissimilar ties [Granovetter, 1973]; the formal result that groups of diverse problem solvers can outperform groups of high-ability solvers exists and is much cited [Hong and Page, 2004]. Against: that result has been argued to be either trivially true or dependent on conditions amounting to the conclusion [Thompson, 2014], and the exchange that followed is the honest context — a defence [Kuehn, 2017]. Three incompatible constructs are called diversity — separation, variety and disparity — so any measure of difference must declare which it is, and the framework's has not [Harrison and Klein, 2007]. And the strongest structural warning is that difference does not add monotonically: creative success rises and then *falls* with cohesion, an inverted U [Uzzi and Spiro, 2005]; in the largest natural setting studied, short-run diversity reduced trust and civic participation, a finding its author did not want [Putnam, 2007], though the effect largely disappears under better controls [Abascal and Baldassarri, 2015]; contact across difference produces value only under stated conditions [Allport, 1954]; the same conclusion holds in randomly assembled ecological systems, where increasing diversity and connectance *destabilises* [May, 1972]; and in engineered systems redundancy can reduce safety outright [Sagan, 1993]. Two further results bound the prescription sharply: adding independent-looking components inside one coupling perimeter increases interactive complexity and can make catastrophic failure a system property rather than an operator error [Perrow, 1984]; and a count of participants that has not been audited for independence is worthless, because one party can present as many [Douceur, 2002].

The framework's own concentration instrument makes the last point in numbers. Four jurisdictions holding 40.9, 27.3, 18.2 and 13.6 per cent give a naive Herfindahl index of 0.293430 and 3.41 effective jurisdictions; collapsing them into the single bloc they actually form gives a bloc index of 1.000000 and an effective count of 1.00 [Curzi, 2026]. Stated as a procedural rule: $\ln(1.00) = 0$, so a structure with one effective bloc has zero available surplus rather than a low one, and the first step of the method is to collapse before measuring [Curzi, 2026]. That reading of $1/\mathrm{HHI}$ as a numbers-equivalent is the standard one [Hill, 1973]. It applies to $1/\mathrm{HHI}$; it does not apply to the framework's own affine effective count, for reasons recorded in Section 10.4.

The parallel exposure in identification is measured in the corpus: with an address-only log the attribution set is 1,953.1 subscribers; with address plus port and live session state it is 1; and after state expiry it is 1,453.2 — so an address is evidence of a pool, not of a person [Curzi, 2026]. The general principle behind such disclosure constraints is an engineering standard rather than an ethical preference: a system must remain secure when everything about it except the key is public [Kerckhoffs, 1883].

1.0.9 9. Method: what kind of borrowing each borrowing is

9.1 Analogy discipline

Every physics borrowing in this paper is labelled at first use as one of three things: a derivation, in which the physical result is used because the system genuinely satisfies its premises; a shared mathematical form, in which the same equation governs both with no claim that the underlying physics applies; or an analogy, a suggestive correspondence with no formal warrant, legitimate as motivation and never as evidence. The labels for this paper are collected

here so that a reader can check them against the text.

Borrowing	Category	Warrant, and what is not claimed
Stationary action, $\delta S = 0$ (Sec. 3.2)	3 — analogy	The action functional of a life is not defined and no Lagrangian is exhibited. The borrowing supplies an image of a resting trajectory. The variational form licenses no teleology, and no physical result is transferred.
Local minima and gradients (Sec. 3.3)	2 — shared form	The mathematics of a stationary point of a scalar field is the same mathematics wherever the field comes from. No claim that the field is a potential energy.
Gradient descent (Sec. 3.3)	2 — shared form	The recursion and its convergence and oscillation regimes hold for any differentiable cost and any step size; the transfer claim is that human change has this recursion's shape, not its physics [Elaydi, 2005, Strogatz, 2015].
Annealing and noise (Sec. 3.6)	2 — shared form, uncited	Escape from shallow minima under decreasing noise is a property of the search procedure. No entry in this corpus's bibliography supports it, and none is claimed.
The reflexive loop (Sec. 4)	1 — derivation, within its own definition	$T = W \circ (\mathrm{id}, M)$ is a definition; nothing is imported. The sociological identification is cited, not derived [Merton, 1948].
Banach's theorem (Sec. 5)	1 — derivation, **conditional**	Genuine derivation *if and only if* the space is complete and the map contracts with a uniform modulus. Neither premise is established for a human loop, so the claim is stated as conditional and the condition is printed [Banach, 1922].
$\sigma_x \sigma_p \geq \hbar/2$ (Sec. 6.1)	quoted exactly, then 3 — analogy	Quoted in its state-preparation reading [Kennard, 1927]. Its extension to a foreseen future is P1, a postulate, and is badged at every occurrence.

Borrowing	Category	Warrant, and what is not claimed
“Thermodynamics of attention” (Sec. 9.3)	3 — analogy, and a poor one	The source’s phrase. There is no energy, no temperature, and no conserved quantity. Reported and marked; not used.

The corpus’s title claims physics, and the claim attracts a specific literature of objection which is better cited than ignored. The historical form of the objection is that economics imported the formal apparatus of nineteenth-century physics without its empirical discipline [Mirowski, 1989]; the general form is that laws hold in patches rather than universally, so importing a law’s form is not importing its scope [Cartwright, 1999]; and the version from inside physics is that a system’s behaviour at one scale is not derivable from the laws governing its parts, which is the strongest available warning against reading a shared equation as a shared mechanism [Anderson, 1972]. The programme this corpus belongs to has been criticised sharply from within for modelling without contact with data [Sobkowicz, 2009]. This paper’s position is the one the table above encodes: category 2 is respectable and is where most of the framework sits, and the failure mode to guard against is a category-2 or category-3 use borrowing the certainty of category 1.

9.2 Scale type, and the objection that bites hardest

The dimensional audit belongs with the verification chain and is given as Section 13.4; what follows here is the part of the audit that constrains the argument rather than the notation.

The model treats comfort as having landscape structure — gradients, minima, ridge heights, differences of depth. That presupposes at minimum an interval scale, and nothing in the source establishes one. This is the framework’s sharpest exposure, and the objection has a name and an author: psychology, on this account, systematically assumes that its attributes are quantitative without ever testing whether they are [Michell, 1997]. **This is the strongest single objection to the framework and it is not answered here.** The formal test for whether an attribute has additive quantitative structure at all is conjoint measurement [Luce and Tukey, 1964], whose conditions — single, double and triple cancellation — are exactly what would have to be checked. **They were not tested.** The standard against which any such claim is judged is set out in the measurement-foundations literature [Krantz et al., 1971], and a route that does not require respondents to use numbers at all is available and unused [Rasch, 1960].

The honest position, adopted here: **comfort is treated as quantitative by stipulation, not by demonstration.** Everything in Sections 3 and 5 that uses differences, ratios of depths, or a metric d is therefore provisional at the level of its arithmetic, however sound its logic. The construct-validation programme that would be required is specified and not run [Cronbach and Meehl, 1955], and the self-report instruments the framework would need are exposed to well-documented distortions: people cannot reliably report the causes of their own behaviour [Nisbett and Wilson, 1977], and a question that names a component of well-being induces a focusing illusion in the answer [Kahneman et al., 2006]. A further formal obstacle applies to any similarity-based reading of difference: psychological similarity is asymmetric and violates the triangle inequality, so it is not a metric [Tversky, 1977], which bears directly on the assumption in Section 5 that d exists at all.

9.3 Limits and boundary behaviour

Each relation stated above, at the boundaries of its parameters.

- $\eta \to 0$. The descent recursion becomes the identity: no movement, and the occupant remains at the initial state for ever. The model's prediction at this limit is that a correctly seen alternative with no step at all produces nothing, which is correct and uninteresting.
- $\eta \to 2/L$ **from below, for a surface of curvature** L. Convergence becomes arbitrarily slow with alternating sign; at $\eta = 2/L$ the recursion is an isometry and orbits; above it, divergence. For the modelled quadratic of Figure 2, $L = 1$, so the boundary is exactly $\eta = 2$. This is the mathematical content of H4 and the model's only sharp prediction.
- $k \to 1^{-}$. The contraction bound survives but the convergence rate goes to zero and the a posteriori bound $q/(1-q)$ diverges: arrival is guaranteed and unbounded in time. The theorem does not fail; it stops being useful, which is a distinction worth keeping.
- $k = 1$. No guarantee at all. An isometry has no fixed point in general and no theorem.
- $k > 1$. Divergence; the loop manufactures error. The corpus's storage-law instance of the same boundary is $\delta \to 0$, where the map ceases to contract and the error-persistence column becomes all ones [Curzi, 2026a].
- **Barrier height** $\to 0$ **(Sec. 3.5).** The basin ceases to be a basin and the occupant is not in a minimum at all; H2 is vacuous in this limit and the model has nothing to say.
- **Barrier height** $\to \infty$. Persistence is total and departure is impossible by gradient. The model then predicts that only the noise term of Section 3.6 can move the occupant, which is a real prediction and one the framework's own protocol does not test.

9.4 Identification, and the framework turned on itself

This subsection belongs to Method because it decides what kind of thing R2 can be. It is the paper's most consequential methodological statement and it goes against the framework.

The reflection problem. The reflexive loop of Section 4, written at the level of a group, is a linear-in-means specification in which the mean behaviour of a reference group enters each member's behaviour while being constituted by it. In that specification the endogenous social effect is not identified from observational data: it is not separable from contextual effects — the exogenous characteristics of the reference group — nor from correlated unobservables that make members of a group alike for reasons that have nothing to do with influence [Manski, 1993]. This is a theorem about what the data can support, not a complaint about sample size, and no amount of held-out performance changes it. The same conclusion arrives from the network side: homophily and contagion are generically confounded in observational network data, so a resemblance between interacting parties cannot be attributed to the interaction [Shalizi and Thomas, 2011]. **The remedies, and what each costs.** Four are available and none is free.

Remedy	What it buys	What it costs
Randomisation of exposure	Identification of the endogenous effect directly	An experiment; infeasible for most of what this framework describes, though social-convention tipping has been demonstrated experimentally [Centola et al., 2018]
An instrument for group composition	Identification without an experiment	A defensible exclusion restriction, which the framework does not have
Partial-population variation	Identification from differential treatment within groups	A designed intervention reaching part of a population
Stipulating the reference group and defending the stipulation	Identification conditional on the stipulation	The stipulation itself must be argued and can be wrong [Brock and Durlauf, 2001]

The fourth is the realistic route for this framework, and it is the one the corpus should take. The template is a discrete-choice model with social interactions, derived from a standard random-utility foundation, which states explicitly the conditions under which its interaction parameter is econometrically identified [Brock and Durlauf, 2001]. An interaction parameter that is merely *definable* is not the same object as one that is *estimable*, and this corpus has so far produced only the first.

This paper adopts no remedy. R2 is stated as a modelling position. It is a definition with a mechanism attached, not an identified empirical result, and every downstream paper citing R2 inherits that status. An unidentified parameter presented as measured would be a defect; a modelling position presented as one is not.

One existence proof, which the framework has otherwise lacked. That reflexivity of this shape *occurs* is not in doubt, and the best documented instance is a market. When the Black–Scholes option-pricing model was published its prices did not fit the market; after traders adopted it, the market came to fit the model, with dates and price data recording the convergence [MacKenzie and Millo, 2003]. That is a reflexive fixed point observed in the field, and it establishes that the mechanism is real without establishing that any particular loop is one. The more general point that classifications of people alter the people classified has been argued directly [Hacking, 1995]. That literature supplies existence, not identification, and the distinction is exactly the one this subsection is drawn to keep.

The framework turned on itself: the Lucas critique. The framework asserts that a system changes its behaviour when it is modelled and observed. Applied consistently, that assertion falls on the framework's own parameters. Parameters estimated in a population that does not know the model cannot be assumed to survive in a population that has read it, because the decision rules the parameters summarise are themselves responses to the environment the model describes [Lucas, 1976]. This is not an outside objection; it is the framework's own thesis pointed at itself, and a framework that exempts its own coefficients from its own central claim is not self-consistent.

The paper's answer, stated so that it can be held against it: **the framework's parameters are not claimed to be stable under its own dissemination, and no result in this paper depends on their being so.** Three consequences follow and are accepted. First, any future estimate of η, δ or an interaction parameter must record the reading status of the population it was estimated in, as a covariate and not as a footnote. Second, the framework's own predictions about a reader constitute a regime change for that reader, so a before-and-after design on a single person is confounded by construction. Third, the framework's practical claim (H3) is *strengthened* rather than weakened by this, because H3 says exactly that a concretely seen alternative re-grades the surface: a person who has read the framework has a different surface, and the framework predicts that. What the framework may not do — and this paper does not — is claim a stable numerical coefficient while asserting that being modelled changes behaviour. Those two claims are not compatible and the second is retained.

1.0.10 10. Procedure and case material

10.1 The location procedure

The framework's first procedure locates the occupied minimum and asks for no change. It is reproduced here in operational form because later papers refer to its outputs.

Seven consecutive days, one hour each, at a fixed place and a fixed hour chosen once and not renegotiated. Materials are deliberately minimal — a notebook, a pencil, weather-adequate clothing, no clock in view. The hour proceeds as follows. **(1)** A spoken bearing, once: *I am where my life has set me down*; it is a bearing, not a vow, and is said even when it feels false. **(2)** For roughly the first fifteen minutes, an inventory of present comfort in the present tense, in sentences beginning *Right now I am…*, ordered body first, then immediate surroundings, then situation. Ten lines is the floor; no entry is too small, and smallness is the point. **(3)** A fantasy audit: when attention departs for a better elsewhere — which is

instrument, not failure — the elsewhere is recorded in one line and beneath it, invariably, the present position of the body: *the body sat in a chair, warm, fed.* The repetition is the finding. **(4)** A pricing glance: after each recorded elsewhere, one word only for what going there would cost. No elaboration and no solving. **(5)** The remainder of the hour is spent without steering; restlessness is expected to crest between minutes thirty-five and fifty. **(6)** A spoken close: *this is the floor; the floor holds.* **(7)** On the seventh day, all seven days are read aloud straight through and every comfort appearing three times or more is underlined. The list of underlines is the procedure's product: the load-bearing comforts as distinct from the decorative ones, which no occupant knows in advance.

Three standing rulings from the source are worth recording because they are falsifiable operational claims and not preferences. First, the hour does not divide: the restlessness that crests in the third quarter is the material, and a divided hour never reaches it. Second, an operator away from the fixed seat sits the hour anyway and marks the page *displaced*, because the contrast between fixed and displaced pages separates comforts that travel with the body from comforts that were property of the chair. Third, the procedure is re-run after any sharp re-grading of circumstance — a move, a loss, an arrival — because an occupant who assumes their old coordinates after a large event is navigating terrain that is no longer there.

10.2 Five characteristic failures

The framework catalogues five ways the location procedure fails. They are reproduced in full, since a procedure supplied without its failure catalogue is an advertisement.

F1. Status-quo rationalisation (ceremonial: *the pond that calls itself the sea*). Having learned that the present is a lawful local minimum, the occupant promotes it to the global minimum: *I am here by law, therefore here is best, therefore nothing need be done.* This borrows the framework's own mathematics for the swindle. The correction is in the definitions: *local* was never a compliment.

F2. Confusing affect with position (*the storm painted on the wall*). The occupant sits at a genuine minimum amid genuine anguish and concludes the model is false: *I am not comfortable, I am miserable, therefore this is not my comfort zone.* But the model never claimed the floor was pleasant; it claimed the floor was cheap as the system prices things, and familiar dread can cost less at the margin than unfamiliar hope. The test: name what would have to be done in the next hour to change the source of the anguish, and observe the gradient rise.

F3. Reactive self-report (*the statue at the shore*). The occupant sits the hour as an exhibit, and the pages return a portrait instead of a position. The mark of it is that the pages read well; genuine pages are repetitive and faintly humiliating. This failure mode is the framework's own restatement of a general and severe exposure of self-report instruments [Nisbett and Wilson, 1977].

F4. Anhedonic flatness mistaken for stillness (*the frozen bowl*). There is a state in which nothing costs because nothing registers, and the procedure cannot locate anything, because location presupposes a working instrument of felt cost. The framework's ruling is stated without ceremony and is retained here: this is a manual of attention and not a treatment; if the underlined comforts on the seventh day number zero, the appropriate course is to seek help first.

F5. Fantasy substitution (*the feast eaten in the doorway*). The imagined elsewhere, run nightly in ever richer detail, discharges the drive it names. The source states that this has been measured. **No entry in this corpus's external bibliography supports it**, and the claim is therefore reported here as the framework's, uncited (Section 10.3). The procedural defence — the invariant second line returning the dreamer to the platform — is independent of whether the mechanism is as described.

Four of the five share one root: they convert a survey into a story — of arrival, of exile, of virtue, of appetite — and a story is exactly what the location procedure asks the occupant briefly to do without. The fifth is not a failure of the occupant at all and is listed only so that it is looked for.

10.3 Case material

The source records one worked case in enough detail to report. An operator designated Renata, a night-shift nurse in her ninth year on a respiratory ward, ran the location procedure in late winter, sitting the hour at her kitchen table from 08:40 to 09:40 each morning directly after shift.

Her first day's page opened *comfortable nowhere*, and then recorded eleven inventory lines, among them: *right now I am off my feet; right now nothing is beeping; right now nobody's family is asking me for numbers I don't have.* Her fantasy audit returned one recurring elsewhere — a walk-in clinic in a coastal town visited once at nineteen. Beneath it, day after day, the second line accumulated: *the body sat in a warm kitchen, fed, in wool socks.* Her pricing glances read, in order across the week: *pension, pension, seniority, Marco, pension* — five recorded glances resolving to three distinct named barriers, one of which is a person (Figure 5).

On the fifth day she recorded the finding the source regards as the procedure's signature, arriving unprompted: *the clinic dream shows up at minute forty. Every day. Minute forty is when the radiator gets going and my feet finally stop hurting. I only dream of leaving once I'm warm. On the ward I never think about the clinic at all.* That observation is the framework's departure-lounge claim in the occupant's own words: escape imagining is a surplus expenditure, performed with resources left over after the present is paid for, and therefore manufactured almost exclusively at local minima. The framework's name for that claim — "the thermodynamics of attention" — is an analogy and a poor one, and is marked as such in Section 9.1; the observation underneath it is a timing regularity that the procedure can record and that a larger sample could test.

Her seventh-day reading produced four underlined recurring comforts: silence of machines, weight off feet, heat, and — annotated *did not expect this one* — being unreachable. She did not resign that week; nothing in the procedure asked her to. What changed, she reported, was the direction of an old accusation: she had spent years telling herself she lacked the courage to leave nursing, and the pages showed something structurally different — an occupant at a real minimum, whose every small departure carried a real price with a real name, one of the names being a person. Her own summary was: *I thought I was stuck. I'm not stuck. I'm placed. Placed I can work with.*

Two further details are kept because they instruct. On her second day she nearly abandoned the procedure, for the common stated reason *this is making me feel worse*; the framework's ruling is that the second-day trough is expected and is evidence of function, because an accurate survey of a hard position hurts to precisely the degree that inaccuracy had been doing the comforting. And her underline list did nothing for three weeks by rule, after which it became the opening stock of the cost inventory, where three of her four load-bearing comforts carried costs she had never priced and one discomfort carried nine years of unrecognised payment.

What this case is and is not. It is a single, self-reported, unblinded, retrospective case with no control condition and no independent measurement. It is offered as an illustration of what the procedure returns, not as evidence that the procedure works. Its status is exactly that of any single reported psychological observation, which current replication estimates should be read against [Open Science Collaboration, 2015].

1.0.11 11. Results

11.1 Definitions fixed. D1 (comfort as the felt total of present costs) and D2 (the reflexive loop $T(x) = W(x, M(x))$) are stated in the form the remaining papers cite. Neither is a finding; both are stipulations, and Section 9.2 records that D1 lacks an operational unit and an established scale type.

11.2 Results established by derivation. R3 and R4 are established. $\delta S = 0$ is stated in its correct stationary form with the explicit rider that the variational form licenses no teleology.

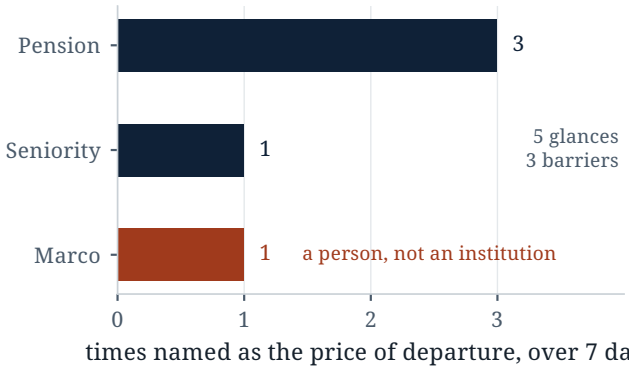


Figure 1.5: In one worked case, seven days of pricing glances resolve five recorded barriers to three names, and one of the three is a person rather than an institution. Measured from the case record as reported in the source, which lists the week's glances in order as "Pension. Pension. Seniority. Marco. Pension."; the same week's seventh-day reading returned four underlined recurring comforts. Script: `figs/src/CCP-A01_fig05.py`.

The Banach condition and its three consequences — existence, uniqueness, geometric convergence from any start — are established, and the uniqueness argument is given in full in Section 5. **This establishment is conditional and the condition is not met by anything demonstrated in this paper:** completeness of the state space and a uniform contraction modulus are premises, not findings, and no k has been measured for any human loop.

11.3 Results established by construction. R1 and R7 follow from D1 by argument rather than by measurement, and R6 is a concession rather than a result. R1 is established in the weak sense that matters: given D1, the occupied state is a local minimum by construction, and the argument is the symmetry argument of Section 3.4. R7 — departure priced at the ridge — follows from the same geometry and is illustrated, not measured, in Figure 3.

11.4 Hypotheses tested. None. H1, H2 and H3 are stated in falsifiable form and are not tested in this paper. No dataset was collected; no instrument capable of eliciting the required quantities exists, for the measurement-theoretic reasons in Section 9.2. H4 is true of the modelled recursion and untested in its human application. This is a negative result and is reported in the same plain language a positive one would receive: **the framework's four testable hypotheses stand untested, and the paper that states them supplies no evidence for any of them.**

11.5 The Postulate. P1 is stated, badged, and left unpromoted. Section 6.3 gives the reason it cannot be numbered among the hypotheses. Its practical consequence is nil: every procedure in Section 10 runs without it.

11.6 Case material. One case is reported. It illustrates the procedure's output and establishes nothing about the procedure's effect.

1.0.12 12. Discussion

12.1 What follows

If comfort is the felt total of present costs, three consequences follow immediately and are the framework's genuine yield.

The first is diagnostic. A position that will not yield to exhortation is not thereby evidence of a defect in the occupant, and treating it as such misroutes the intervention. What the model asks for instead is a barrier survey: name the ridges, price them, and observe which of them are made of material that could be renegotiated. Renata's week produced exactly such a survey, and its most consequential finding — that one of the three barriers was a person — is the

kind of finding a motivational account cannot produce, because a motivational account does not have a place to put it.

The second is prescriptive and is where the framework's practical claim lives. If movement is by gradient, then the intervention is on visibility rather than on effort. That is H3, and its corollary is that pressure is contraindicated: pressure adds a term that the loop of Section 4 returns with a sign the intervener did not intend.

The third is structural. Because T is a composition, a loop can be diagnosed by its gain rather than by its content, and the diagnosis carries across scales. The same gain analysis that says a personal correction was too large says that an institutional response to autonomy was too large, and predicts, in both cases, that the returned difference will be read as vindication.

12.2 What does not follow

That a position is lawful does not make it good, deserved, or fairly graded. That a loop has a fixed point does not mean it has been shown to contract. That the mathematics of two domains coincides does not mean the domains do; that is precisely what P1 asserts and precisely why P1 is a postulate. And nothing in this paper licenses the operation of any of its procedures on another person: the framework's own constraint — no future spoken over another's path without their word — is argued at length in the companion paper and is not a courtesy but a stability condition [Curzi, 2026].

Two further non-implications are worth stating because the corpus's later results tempt them. The framework's affine effective count is used across the corpus, including in the essence bridge that feeds a count into a pitch criterion, where the minimum sustainable frequency is $f_{\min} = k\, e^{2c/\eta}$ with $k = 0.483425732988414$ Hz and the trichotomy $f \gtrless f_{\min}$ corresponds to $Q_\infty \gtrless 0$ [Curzi, 2026], and where the deficit identity $D = 2.885390\,(c/\eta - F) = -2.885390\,(\delta/\eta)\,Q_\infty$ makes the octave deficit and the steady-state gauge the same number [Curzi, 2026]. Two cautions attach. First, k is a *defined* constant, exact by construction, and its fifteen significant figures are defensible only on that basis; attaching hertz to it is a separate and much stronger claim than computing it, and this paper does not endorse that attachment. Second, the general hazard of reading a chosen ratio as a discovered natural constant is well documented and the documentation should be read before any such reading is offered [Markowsky, 1992].

12.3 Relation to the corpus

The definitions fixed here are consumed downstream in a specific and checkable way. The storage law's leaky-integrator form and its pathology at zero decay [Curzi, 2026a] instantiate the contraction condition of Section 5 in a system with an explicit state variable. The preregistered measurement protocol for that law reports a confirmatory fit that failed its own second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin, notwithstanding a 40.3 per cent advantage over persistence, 97.7 per cent over direct input, and $p = 0.0005$ against a shuffle control [Curzi, 2026i]; that is the corpus's own strongest evidence that its instruments can fail, and it is cited here rather than in a footnote. The same five-parameter survey applied from one desk to a twenty-two household block reproduced each system's observed decline — predicted -0.250474 against observed -0.250 at the desk, -0.996371 against -1.000 at the library, residuals under 1.4 per cent of the ceiling term [Curzi, 2026k]. The moment equation's fixed point $D = P$, $r = 1$, $M = E$, with unit elasticity of exactly ± 1 in all three terms and the ceiling $0 \le M \le \ln N$ making the $P \to 0$ divergence finite [Curzi, 2026m], is the same fixed-point structure applied to a calibration ratio. And the corpus's conditional postulate on belief — $E = G \cdot B + \varepsilon$ with the forbidding clause $G = 0 \Rightarrow E = 0$ for all B, badged as a postulate and paired by explicit correspondence with the storage law's $\eta = 0$ [Curzi, 2026ab] — is the template this paper's P1 follows: a strong claim, stated with the clause that forbids its abuse, and badged.

1.0.13 13. Verification

This section is not a summary. It decomposes the paper's results into pieces small enough that a reader can check each one without checking any other, and states for each piece which of four things it is: `ARITHMETIC` (redo the calculation), `DEFINITION` (fixed earlier in this paper), `CITED` (one named external source), or `POSTULATED` (asserted here, with no external support). The razor is applied: any postulate that could be dropped with the results still following has been dropped.

13.1 Postulates

Each is atomic and each is checkable on its own.

- **P1 (DEFINITION).** Comfort is the felt total of present costs. *Check:* it is D1; nothing else is required.
- **P2 (DEFINITION).** The occupant's state is a point in a set X on which a real-valued cost f is defined. *Check:* stipulated in Section 3.1. **Note: this is the paper's largest unearned assumption and Section 9.2 says so.**
- **P3 (DEFINITION).** A local minimum of f at x means $f(x) \leq f(y)$ for all y in some neighbourhood of x. *Check:* standard definition; no source needed.
- **P4 (POSTULATED).** A person's occupied state is the cheapest immediately reachable state as their own system prices it. *Check:* not checkable from outside the framework; this is the substantive content of H1 and it is asserted, not demonstrated.
- **P5 (DEFINITION).** M maps a state to a model; W maps a state and a model to a next state; $T = W(\cdot, M(\cdot))$. *Check:* D2. **Note: as a definition this is free; as an empirical specification it is not identified from observation [Manski, 1993], so nothing downstream may treat its parameters as estimable without adopting one of the remedies in Section 9.4.**
- **P6 (DEFINITION).** A fixed point of T is an x^* with $T(x^*) = x^*$.
- **P7 (CITED: `banach1922operations`).** On a complete metric space, a map satisfying $d(T(x), T(y)) \leq k\, d(x, y)$ with $0 \leq k < 1$ has exactly one fixed point, reached from any start, with error shrinking by at least k per pass.
- **P8 (POSTULATED).** A space of life-states is complete, and a human reflexive loop admits a uniform contraction modulus $k < 1$. *Check:* **not checkable, and not checked. This is the premise on which the entire transfer of P7 rests.**
- **P9 (CITED: `kennard1927quantenmechanik`).** $\sigma_x \sigma_p \geq \hbar/2$ is a property of the quantum state, not of measurement disturbance.
- **P10 (POSTULATED).** The uncertainty of P9 and the uncertainty of a foreseen future are the same kind of uncertainty. *Check:* not checkable; this is the framework's Postulate and is badged wherever it appears.

Postulates dropped as unnecessary: the annealing analogy (Section 3.6) was tested for load and carries none — every result below survives its removal — and the narrative-transfer claim (Section 8.2) likewise. Both are reported in the text and neither is used.

13.2 The compounding chain

Each line consumes only the outputs named on its left.

```
P1 + P2          -> L1   a scalar cost is defined over states     DEFINITION
L1 + P3          -> L2   local minima of that cost are defined    DEFINITION
L2 + P4          -> R1   the occupied state IS a local minimum    POSTULATED (via P4)
R1 + P3          -> L3   every small departure carries a cost     ARITHMETIC
L3               -> R7   persistence is set by barrier height,
                         not by basin depth                       ARITHMETIC
L1               -> L4   descent x_{n+1} = x_n - eta grad f       DEFINITION
```

```
L4              -> H4   |1 - eta L| >= 1 gives non-convergence    ARITHMETIC
P5 + P6         -> R2   a self-fulfilling prophecy is a fixed
                        point of T                               DEFINITION
                        (as an empirical specification: NOT
                        IDENTIFIED, see 9.4)                     POSTULATED
P7              -> R4   the Banach condition and its three
                        consequences                             CITED
R2 + R4 + P8    -> C1   a well-built prediction is reached from
                        anywhere                                 POSTULATED (via P8)
P9              -> L5   the uncertainty relation, stated right   CITED
L5 + P10        -> R5   the Postulate, badged                    POSTULATED
(concessions)   -> R6   noise, plasticity, foresight             DEFINITION
```

What this chain shows, stated plainly. Of the seven results, exactly two — R3 and R4 — rest on `CITED` or `ARITHMETIC` steps alone. R1 passes through P4 and is therefore `POSTULATED` at its load-bearing step. C1, the framework's most attractive conclusion, passes through P8, which is not checkable and was not checked. R5 is `POSTULATED` by construction and says so.

A chain whose steps are mostly `DEFINITION` and `POSTULATED` is not a proof, and this one is mostly `DEFINITION` and `POSTULATED`. What it is instead is a specification of exactly which two premises would have to be established for the framework's central conclusion to follow: P4 (that occupancy tracks cheapest reachability) and P8 (completeness plus a uniform modulus). Those two are the whole bill. That is a smaller and more addressable bill than the framework's rhetoric suggests, and stating it is this section's only purpose.

13.3 The checks a reader should run first

1. **Redo the uniqueness argument** in Section 5. It is three lines and needs no source.
2. **Redo Figure 2's recursion** by hand for $\eta = 2.10$: $x_n = (-1.1)^n$, so $|x_{12}| = 1.1^{12} \approx 3.14$. The figure should agree.
3. **Check the dimensional table** in Section 9.2 for an invented unit. There is none, because comfort has no operational unit and the table says so.
4. **Check every occurrence of** $\sigma_x \sigma_p \geq \hbar/2$ for the state-preparation reading. There are two, in Sections 6.1 and 9.1, and both carry it.
5. **Check that the analogy table in Section 9.1 assigns category 1 to exactly two entries**, and that both carry an explicit statement of the premise that is not established.
6. **Check that R2 is never described as measured**, in this paper or in any paper citing it. It is a modelling position; Section 9.4 says so and the Conclusion repeats it.

13.4 Dimensional audit

Every symbol used above, with its dimension, its unit, and its epistemic status. The audit is short because the framework is honest about how few of its quantities have units at all.

Symbol	Meaning	Dimension	Unit	Status
x	state of the occupant	none assigned	none	defined; an index into a state set, not a measured coordinate
$f(x)$	felt total of present costs (comfort)	none assigned	**no operational unit exists**	defined (D1); not measured anywhere in this paper

Symbol	Meaning	Dimension	Unit	Status
∇f	local slope of the cost surface	$[f]$ per $[x]$	none	derived from f; inherits f's lack of a unit
η	step size	$[x]^2/[f]$ for the ratio to close	none	modelling parameter, chosen, not measured
n	iteration index	dimensionless	passes	defined
$d(\cdot,\cdot)$	distance between states	$[x]$	none	assumed to exist; not constructed
k, q	contraction modulus	dimensionless	—	defined; **never measured for a human loop**
T, M, W	maps on states	—	—	defined (D2)
δS	variation of the action	energy $\times$ time	J·s	quoted from physics; not a quantity of this framework
$\sigma_x, \sigma_p, \hbar$	quantum standard deviations, reduced Planck constant	m, kg·m·s^{-1}, J·s	SI	quoted from physics; not quantities of this framework

Does the central relation balance? The descent recursion $x_{n+1} = x_n - \eta \nabla f(x_n)$ balances only if η carries dimensions $[x]^2/[f]$. Since neither $[x]$ nor $[f]$ has been given an operational unit, the balance is formal: it is satisfied by construction and constrains nothing. **The source is silent on the dimensions of comfort, and this paper does not supply them.** A dimensionless placeholder is the honest entry; an invented unit would not be.

A quantity with no operational unit is entered as having none. No unit was invented for comfort, and none should be inferred from the fact that the recursion balances.

1.0.14 14. Limitations

The framework's credibility rests on this section being longer and more specific than a reviewer expects. Each item names what would falsify the claim or what the paper cannot support.

14.1 Comfort's scale type is stipulated, not demonstrated. This is the paper's principal exposure and it is stated first. Everything with landscape structure — depth, gradient, ridge height, the metric d — presupposes at least interval structure, and none was established. The strongest objection to the framework is that psychology routinely assumes quantitative structure without testing for it [Michell, 1997]; the formal test that would settle the matter exists [Luce and Tukey, 1964, Krantz et al., 1971] and was not applied. **Falsified by:** a conjoint-measurement study on any operationalisation of comfort that fails double cancellation. The framework would then have to be restated on an ordinal basis, which would remove Figure 3's arithmetic and Section 5's metric entirely.

14.2 No hypothesis was tested and no data were collected. H1–H4 are stated in falsifiable form and left standing. Section 11.4 reports this as a result. The framework's own preregistration discipline is the correct remedy and was not exercised here [Nosek et al., 2018], and the analytic flexibility available in any later confirmatory attempt should be constrained in advance [Simmons et al., 2011, Gelman and Loken, 2013].

14.3 Claims with no supporting literature in this corpus's bibliography. These are listed individually because a reviewer will find them, and finding them first is the point.

- *Pressure on a person's choices produces counter-movement.* The framework calls this reactance and treats it as one of the best-replicated findings in its literature. `bib/external.bib` **contains no entry for it.** The claim is retained in the text because it is load-bearing for H3's second limb, and it is retained **uncited**. Nothing in this bibliography supports it and no citation was attached to make it look supported.
- *Losses weigh roughly twice what equal gains weigh; people overvalue what they already hold; paid costs recruit fresh ones; cues placed in the path of a behaviour govern it more reliably than resolutions about it.* These are the source's fourth class of formal claim. **The bibliography contains no entry for loss aversion, for the endowment effect, for sunk-cost escalation, or for cue-based habit formation.** The class is stated in the paper and is uncited throughout.
- *A goal vividly enjoyed in imagination is partially discharged, spending the drive that would have moved the feet* (failure mode F5). **No entry.** Reported as the framework's claim.
- *A pattern learned as a narrative transfers across domains where a pattern learned as a checklist does not.* This is the framework's justification for its own form. **No entry.** Reported and not endorsed (Section 8.2).
- *Simulated annealing, and its transfer to disrupted seasons of a life.* **No entry.** Reported and flagged in Section 3.6 and in the analogy table.
- *The uncertainty of a foreseen future is the same kind as the uncertainty of a quantum state.* No entry, and none is possible; this is P1 and is badged accordingly.

14.4 The contraction premises are unverified. Section 5 states them; Section 13.1 lists them as P8; and the honest consequence is that the framework's most attractive conclusion — arrival guaranteed, from anywhere, at a rate — is available only conditionally. No k has been measured for any human loop by anyone, including the authors of this corpus.

14.5 Parameters chosen rather than measured. η in Figure 2 and every control point in Figure 3 were chosen to make a stated prediction visible. They are not estimates. The step-size boundary $\eta = 2$ is exact only for the quadratic actually plotted, whose curvature is 1 by construction; for any other surface the boundary is $2/L$ and L is unknown for comfort surfaces. The barrier heights in Figure 3 are computed from the plotted curve, but the curve itself is stipulated.

14.6 Illustrative rather than empirical results. Figures 1, 2, 3 and 4 illustrate; only Figure 5 reports numbers from a record, and that record is a single self-reported case (Section 10.3). Nothing in this paper constitutes evidence that the location procedure changes anything.

14.7 The self-report instruments the framework needs are exposed. People cannot reliably report the causes of their own behaviour [Nisbett and Wilson, 1977], and any question naming a component of well-being induces a focusing illusion [Kahneman et al., 2006]. The framework's own failure mode F3 anticipates this at the level of a single operator's pages; the instrument-level problem is not thereby solved.

14.8 The instrument corrupts what it measures. A quantitative indicator adopted for social decision-making distorts the process it monitors [Campbell, 1979]. The framework's own effective counts are exactly such indicators, and this exposure is structural rather than incidental.

14.9 The surplus-from-difference claim is stated monotonically and the evidence is not. Section 8.3 gives both sides. A monotone claim cannot survive Uzzi and Spiro [2005], Putnam [2007], May [1972] and Sagan [1993] unqualified, and this paper does not assert one; the framework elsewhere does, and that is a defect in the framework rather than in the literature.

14.10 A scope note on the framework's own effective count. The corpus runs two different effective counts. The reading of $1/\mathrm{HHI}$ as a numbers-equivalent is exact and fully grounded [Hill, 1973]. The framework's own affine count is **not** a Hill number of any order and has no precedent in ecology, economics, political science or information theory. It is this corpus's

own construction. Where this paper cites the Hill-number literature it does so only for $1/\text{HHI}$, and no reader should transfer that grounding to the affine count.

14.11 The reflexive loop is not identified, and this is the framework's most serious methodological problem. As an aggregate specification, $T(x) = W(x, M(x))$ places the endogenous social effect beyond identification from observational data [Manski, 1993], and the network route is closed by the generic confounding of influence with selection [Shalizi and Thomas, 2011]. This paper adopts no remedy and therefore states R2 as a modelling position (Section 9.4). **Falsified by, or rather resolved by:** a partial-population or randomised design, or a defended stipulation of the reference group in the manner of Brock and Durlauf [2001]. Until one of those is run, every corpus claim of the form "the interaction produced the shared state" is unsupported, and this includes claims the corpus currently makes.

14.12 The framework's own parameters are not exempt from its own thesis. The framework asserts that being modelled and observed changes behaviour, which by the Lucas critique means its own coefficients cannot be assumed stable across populations that have or have not read it [Lucas, 1976]. Section 9.4 states the paper's answer — that no stability is claimed and no result here depends on it — and the cost is real: no cross-population comparison of a framework parameter is currently interpretable, and no single-person before-and-after design is clean.

14.13 The logarithmic form is underdetermined. A power law with exponent below one [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the same qualitative facts as the corpus's logarithm. **Nothing in this paper discriminates between them.** The strongest available warrant for a logarithm is dynamical rather than psychophysical [Peters and Gell-Mann, 2016], and it applies only where the underlying process is multiplicative, which has not been established for comfort.

14.14 Legibility risk. The framework's central operation is to make an implicit surface explicit. The historical record on schemes that improve a situation by making it legible is not encouraging: legibility can destroy the local variety that made the system work [Scott, 1998]. This applies to the framework's own instrument and is not merely a caution about other people's.

1.0.15 15. Conclusion

The framework's seven results are stated below in the exact form in which the rest of the corpus cites them. Downstream papers should quote this text rather than paraphrase it.

R1. Comfort is the felt total of present costs, not pleasure; therefore the present state is a local minimum of the comfort landscape by definition rather than by luck.

R2. The reflexive loop is $T(x) = W(x, M(x))$, and a self-fulfilling prophecy is a fixed point of that composition of seeing and responding. **R2 is a modelling position, not an identified empirical result:** as an aggregate specification the endogenous effect is not identified from observational data, and no remedy is adopted here [Manski, 1993]. Papers citing R2 inherit that status and must not present the loop's parameters as measured.

R3. The principle of stationary action is $\delta S = 0$; the honest form is stationarity, not minimality, and the variational form licenses no teleology.

R4. If T satisfies $d(T(x), T(y)) \leq k\, d(x, y)$ with $0 \leq k < 1$ on a complete metric space, then T has exactly one fixed point, it is reached from any starting point, and convergence is geometric with ratio at most k.

R5. The Postulate: the uncertainty of quantum measurement and the uncertainty of a foreseen future are the same kind of uncertainty. This is a postulate and is badged as one wherever it appears. $\sigma_x \sigma_p \geq \hbar/2$ holds only in its state-preparation reading and says nothing about measurement disturbance.

R6. Three disanalogies between a person and a physical minimiser are conceded: noise, plasticity of the landscape, and response to imagined costs. The third, foresight, is the term for which physics has no analogue.

R7. Departure is priced at the height of the nearest ridge, not at the depth of the current basin; a shallow basin ringed by high ridges holds as firmly as a deep one, and the two conditions call for different work — engineering in the first case, acceptance in the second.

Counter-result. The identification of a life with a minimising physical system is a postulate rather than established physics; three disanalogies break the analogy; and locating oneself at a local minimum is an exoneration and not an endorsement, since *local* was never a compliment.

What is established, then, is a set of definitions with a clear consequence, one theorem with two unmet premises, and one badged postulate that no result requires. What is not established is every hypothesis the framework advances. The framework's contribution at this stage is that it says which two premises would settle the matter — that occupancy tracks cheapest reachability, and that the loop contracts with a uniform modulus on a complete space — and that it declines to claim either.

1.0.16 16. Declarations

Data and code availability. No external dataset was collected for this paper and none exists. The case material in Section 10.3 is reported from the source manuscript's own record and is not a dataset. All figures are generated by the scripts named in their captions (`figs/src/CCP-A01_fig01.py` through `figs/src/CCP-A01_fig05.py`), which import the shared style module `lib/figstyle.py`; every number plotted is either computed within the script or quoted from the source text identified in the script's docstring.

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the analysis, the figures, and the text.

Ethics. No human-subjects research was conducted for this paper. No personal data were collected or processed. The case material in Section 10.3 is drawn from the source manuscript, is presented under a pseudonym as it appears there, and is illustrative rather than evidential; it was not gathered under a research protocol and is not offered as one. Nothing in this paper is clinical advice, and Section 10.2's fourth failure mode records the framework's own instruction that a person in an anhedonic state should seek appropriate help before applying any procedure described here.

Use of the ceremonial register. The source manuscript is in part an initiatory document. This paper analyses that register as an object of study and does not reproduce it. Sections 8.1 and 8.2 report the framework's ceremonial names and Latin mnemonics as data about the framework, not as instructions to a reader.

Chapter 2

The Ledger, the Gradient, and the Arithmetic of Difference: Carrying Cost, Desire Lines, and the Surplus Functional f(u) = ln(1 + (N-1)u)

Abstract. This instalment develops three expository nodes and states the surplus functional the engineering series later builds on. Holdings are priced in four carrying-cost currencies - coin, hours, attention, and foreclosure - and discomforts are assayed for four yield currencies, including contrast, on the argument that felt comfort is a difference detector and so decays toward zero over weeks while the carrying cost bills at full rate indefinitely - an asymmetry offered as the mechanism that assembles a depleted life without any single bad decision. The gradient node states force as the negative gradient of a potential, argues that the variational and local formulations are one equation written twice, and treats worn paths as the theory's observable: a desire line is the integral of many small descents, and its reproduction after the recording medium is erased is offered as evidence that the pattern lies in the force field. The horizon node states the surplus of N non-parallel sightlines as f(u) = ln(1 + (N-1)u), uniquely determined by four requirements - angle-only dependence, zero surplus at parallelism, an effective count affine in u, and normalisation - with N_eff = 1 + (N-1)u read as an effective number of views. Two limits are stated rather than hidden: the functional is a lower bound on the ensemble's information, not the information, and its concavity means most of the available gain arrives with the first few genuinely divergent views.

Keywords: carrying cost; hedonic adaptation; desire lines; surplus functional; effective count; non-parallelism; lower bound; local minimum

Series: A. Foundations **Paper:** CCP-A02

2.0.1 1. Introduction

1.1 The question

Three questions sit underneath most of what this corpus later computes, and none of them is usually asked in the same place. What does a held arrangement actually cost its holder, once the invisible currencies are counted? What is the observable of a comfort gradient — the

thing a sceptic could go and measure? And how much does a second, third or fortieth point of view add to a first?

This paper answers all three, because in the source framework they are one node's work each and the answers interlock. The costing answer is a four-currency ledger plus one asymmetry law. The observable is the worn path. The arithmetic answer is a single functional,

$$f(u) \;=\; \ln\!\big(1+(N-1)u\big)$$

which the remaining sixty-nine papers of the corpus use more than any other formula it contains. Thirty-six of them cite the results established here. Because those results are load-bearing elsewhere, this paper states them in a form intended to be quoted rather than paraphrased, and it prints the two limits of the functional at the same volume as the functional itself: it is a *lower bound* and not the information, and its concavity is an argument against collecting viewpoints rather than for it.

The paper also reports, at the same precision it reports its results, three things that go against the framework. The functional's form is underdetermined by the existing evidence. Its effective count is not a Hill number and has no external precedent. And the source's own summary claim — that three well-cut views capture the greater part of what a gallery could — is true only as a comparison of angles, not of counts, and the break-even is computed here for the first time.

1.2 What this paper contributes

Seven results are established or fixed here, labelled R1 through R7 and restated verbatim in the Conclusion.

R1 states the surplus functional $f(u) = \ln(1 + (N-1)u)$, uniquely characterised by four requirements — dependence on angles alone, zero surplus at parallelism, an effective count affine in u, and normalisation — with $N_{\text{eff}} = 1 + (N-1)u$ read as an effective number of views. **R2** records that the functional is deployed as a *lower bound* on the ensemble's information and not as the information itself: it logs the floor and not the ceiling. **R3** states the binding clause of R1: a second opinion drawn from a source that shares training, incentives and premises with the first is worth $\ln 1 = 0$ of surplus. **R4** states the four-currency carrying cost — coin, hours, attention, foreclosure — against the law that felt comfort decays over weeks while carrying cost does not decay at all. **R5** identifies the worn path as the theory's observable: a desire line is the integral of many small descents, and its reproduction after the recording medium has been erased is the evidence that the pattern lies in the force field rather than in the record. **R6** states that the variational form of a mechanical law is equivalent to the local differential law, so that the principle licenses no teleology — the path obeys its ground instant by instant. **R7** states the honest limit of descent: a gradient finds the nearest low place and not the best one, so the analysis can diagnose a stable trap and cannot treat one.

A counter-result is stated with equal weight. The surplus is bounded twice over: it is only a lower bound on the ensemble's information, and its logarithmic concavity means that a few well-separated views capture most of what a large panel could, so the framework argues explicitly *against* accumulating perspectives. The gradient material then concedes that descent can diagnose a stable trap but cannot treat one, and the ledger material concedes that the pricing biases on which it leans have no supporting entry anywhere in this corpus's bibliography.

1.3 Register, sources, and what is not reproduced here

The source manuscript is written in two registers, a narrative one and an analytic one, deliberately braided. This paper reports the analytic register. The narrative material is a fictional exposition and carries no analytical claim the treatise does not also carry; it is not reproduced. Where the source's ceremonial vocabulary names a construct — *the test of the stone*, *water does not negotiate*, *the many windows of the one tower* — the construct is renamed to what it

is, and the ceremonial name is reported once, as an object of study, because the framework's use of a redundant second encoding is one of its own methodological claims. The framework's choice of a narrative vehicle for a procedural argument has a respectable precedent in the operations literature, where the same choice was made for the same stated reason [Goldratt and Cox, 1984].

The material below corresponds to three of the source's twelve constructs — the inventory, the gradient, and the horizon — and is presented as one argument because the third consumes the first two: the arithmetic of difference is what tells an analyst how many honest viewpoints a survey of their own ground is worth.

A note on chunk boundaries. The source volume was re-split after a de-framing correction, and the file assigned to this paper begins partway through the inventory material and ends partway through the horizon material. Section 3 therefore reports the inventory node's doctrine, procedure and failure catalogue but not its opening narrative; Section 5 reports the horizon node's arithmetic and its procedure but not its closing failure catalogue, which falls in the file assigned to the next paper in this series. Nothing analytical is lost at either seam, and both seams are marked here rather than papered over.

2.0.2 2. Hypothesis

The material below divides into claims that could be wrong, definitions that could not, and one modelling choice presented as an axiom. Keeping the three apart is the framework's own stated discipline and this section enforces it. A hypothesis that nothing could falsify is restated as a definition or a postulate and labelled as such [Popper, 1959].

H1 — the carrying-cost hypothesis. Every held arrangement charges its holder in four currencies — coin, hours, attention, and foreclosure — and the last two are systematically omitted from unaided estimates of what a holding costs. *Predicted:* an itemisation performed in a fixed four-currency order returns a strictly larger set of costs than an open-ended request for the same holding's costs, and the increment is concentrated in the third and fourth currencies. *Falsified by:* matched itemisations in which the fixed-order protocol returns no more attention or foreclosure entries than the open-ended one. *Status:* testable in principle; **not tested here**, and no instrument for eliciting an attention charge exists (Section 6.2).

H2 — the decay-asymmetry hypothesis (the ledger's central claim). The felt return on a held comfort decays toward zero on a scale of weeks while its carrying cost does not decay at all, so that a life can accumulate a large and rising cost against a falling felt return without any single bad decision. *Predicted:* for a holding acquired at a known date, reported felt return declines monotonically with time held, while an itemised cost, re-elicited, does not. *Falsified by:* itemised costs declining at the same rate as felt return, or felt return not declining. *Status:* testable; **not tested here**. Figure 3 is an illustration of the prediction under two chosen time constants, not a measurement of either.

H3 — the desire-line hypothesis. Repeated movement through a setting is a better survey of that setting's comfort gradients than any statement of intention about it, and the survey is recoverable after the recording medium is erased. *Predicted:* after a complete erasure of the physical record of use — snowfall, resurfacing, a move — the same routes reappear within a few cycles, in the same places, without any transfer of information through the erased medium. *Falsified by:* routes failing to reappear after erasure, or reappearing only where the erased record was in fact preserved. *Status:* observed repeatedly in the built-environment literature and treated there as a design method [Whyte, 1980, Gehl, 1971]; the transfer to a personal day is the framework's, and is **not tested here**.

H4 — the surplus hypothesis. The value added by N viewpoints over one viewpoint depends on the angles between them and not on their number, and is $\ln(1+(N-1)u)$. *Predicted:* two panels of equal size but different measured non-parallelism differ in aggregate performance as the ratio of their $\ln N_{\text{eff}}$ values, and a panel at $u = 0$ performs exactly as a single member does. *Falsified by:* a panel at measured $u = 0$ outperforming a single member, or by a power-law or $n \log n$ form fitting held-out data better at equal parameter count. *Status:* **not**

tested here, and currently underdetermined. Two established rival forms fit the same qualitative facts (Section 5.6). The uniqueness argument establishes that the form follows from the four requirements; it does not establish that the four requirements are the ones the evidence supports.

H5 — the concavity hypothesis. Because f is concave in N_{eff}, most of the available gain arrives with the first few genuinely divergent views, so adding views past a small number is dominated by raising the angles between the views already held. *Predicted:* at fixed cost, a protocol that raises measured u outperforms one that raises N, above a small N. *Falsified by:* performance continuing to track N at fixed u past the point where $\ln N_{\text{eff}}$ has flattened. *Status:* mathematically true of the functional (Figure 2); the transfer to any practice is untested.

D1 — non-parallelism (definition). $u \in [0, 1]$ is a measure of the non-parallelism of a set of viewpoints, with $u = 0$ meaning every viewpoint faces the same way and $u = 1$ meaning each contributes a direction of view the others lack entirely. The corpus's operationalisations are $u = 1 - (x \cdot y)^2$ on unit direction vectors [Curzi, 2026c] and $u = 1 - \rho^2$ from the decorrelation of errors [Curzi, 2026o]. Both are stipulations. Neither is derived.

D2 — the effective count (definition, and a modelling choice inside it). $N_{\text{eff}} = 1 + (N - 1)u$. The *affine* form is a design axiom of this corpus and not a theorem; the corpus badges it as such in the papers that use it [Curzi, 2026c,o], and this paper badges it here. It is **not** a Hill number of any order, and Section 5.2 states what follows from that.

D3 — carrying cost (definition). The carrying cost of a holding is the standing charge of keeping it, itemised in coin, hours, attention and foreclosure, and is distinct from its acquisition price. This is a stipulation. Its measurement-theoretic status is audited in Section 6.2.

P1 — the relaxation postulate. That a life relaxes through its comfort landscape as a physical system relaxes through its states is postulated wherever the gradient material is applied to a person. The source states this postulate in its own text — "the Lodge does not claim you are water; the Lodge claims the mathematics is the same shape" — and this paper preserves the badge at every occurrence. Nothing in physics licenses it.

2.0.3 3. The carrying-cost ledger

3.1 Four carrying-cost currencies

The framework's first analytical move is to replace a holding's *purchase price* with its *carrying cost*, and to insist that the carrying cost be itemised in a fixed order across four currencies. The term is borrowed from warehousing, where the fact that storage is never free is not a discovery but the trade's first principle, and where the standing cost of held stock is an ordinary line item [Hopp and Spearman, 2011]. Written as a system of levels and rates, a held arrangement is a stock with a continuous outflow, and the outflow is what this section prices.

The four currencies, in the order the framework fixes:

1. **Coin.** The monetary charge. The framework calls it the least interesting currency on the ground that the world already itemises it, and the ledger's contribution is not here.
2. **Hours.** Not the hours a holding fills, which were chosen, but the hours it *appoints*: maintenance, tending, and the arrangement of a life around a feeding schedule.
3. **Attention.** The background awareness a held thing charges while untouched. The framework's claim is that this charge compounds invisibly and is the dominant term for people who report exhaustion without a corresponding workload: not too much work, but too many open accounts, each billing a quiet standing awareness, none itemised anywhere. That a person's simultaneous holdings are bounded by a capacity rather than by effort is old and well established [Miller, 1956], and the general statement — that any regulator has a finite channel and must have at least as much variety as what it regulates — is Ashby's [Ashby, 1956]. Neither result establishes that a *held comfort* draws on that capacity while untouched; that is the framework's claim and it is uncited (Section 10.3).

4. **Foreclosure.** What a holding forbids. Every held position closes others by being held: the lease that forecloses the leaving, the reputation that forecloses the beginner's mistakes, the standing arrangement that forecloses — so long as it stands, and precisely because it stands — a decision. The framework's note on foreclosure is that it never sends an invoice, because what it takes never existed to be missed, and its procedural rule is that foreclosure is written *last*, after the first three currencies have been attempted.

The procedural rule attached to the four is that a holding is priced in all four before "nothing" may be written in any of them, and that a dash is earned by a second look and never by a first.

3.2 Four yield currencies, and the contrast term

The framework's second move is symmetrical and is the one that keeps the instrument from becoming an indictment: a discomfort is *assayed* rather than discarded, in four yield currencies — information, training, position, and contrast.

Three of the four are unremarkable. *Information*: a standing ache may be an accurate report on a situation everyone else has agreed to describe pleasantly. *Training*: a hard hour may be the one hour in which the body outranks the head. *Position*: a hard duty may hold open a door that comfort would let swing shut.

The fourth is the framework's own and is worth stating carefully, because it is the term on which the ledger's decay law later depends. **Contrast**: a comfort is a difference before it is a state. The warmth is manufactured at the threshold, by the cold that preceded it; a life engineered to contain no thresholds does not become a life of continuous warmth but a life in which warmth has nothing to be measured against and leaves the instrument's range. The framework's operational consequence is that some fraction of a person's discomfort is not billing them but *calibrating* them, keeping a difference detector supplied with differences.

The contrast term has real support and the support should be named exactly, because it is narrower than the framework's use of it. That perception is organised around a moving reference level rather than an absolute one is Helson's adaptation-level result [Helson, 1948]; that sensation tracks a ratio to a just-noticeable increment rather than an absolute magnitude is Weber's, integrated into a logarithm by Fechner [Weber, 1834, Fechner, 1966]; and that subjective states return toward a set point after large changes in circumstance is the hedonic-adaptation finding [Brickman et al., 1978], in the substantially revised form that allows set points to move and individuals to differ [Diener et al., 2006]. What none of these establishes is the framework's stronger claim that *deliberately retained* discomfort raises the felt value of a comfort. That is an intervention claim and it is untested; it is stated here as the framework's and badged.

3.3 The decay asymmetry (R4)

The two columns are joined by one law, and it is the ledger's whole yield:

> **Comforts fade from feeling long before they fade from billing.**

The mechanism the framework offers is the difference detector of Section 3.2: felt experience registers changes and lets constants sink below notice, so the pleasure of a held comfort decays toward zero on a schedule of weeks, while its carrying cost does not decay at all. The shelf bills at full rate indefinitely, in coin, hours, attention and foreclosure, felt or not.

Run that asymmetry across a decade of accumulation and the framework's conclusion follows without any bad decision being required: a stock of comforts that no longer comfort, each billing its original rent, the total rising as the felt total falls. This is R4, and it is the framework's account of how a depleted arrangement assembles itself out of individually defensible acquisitions.

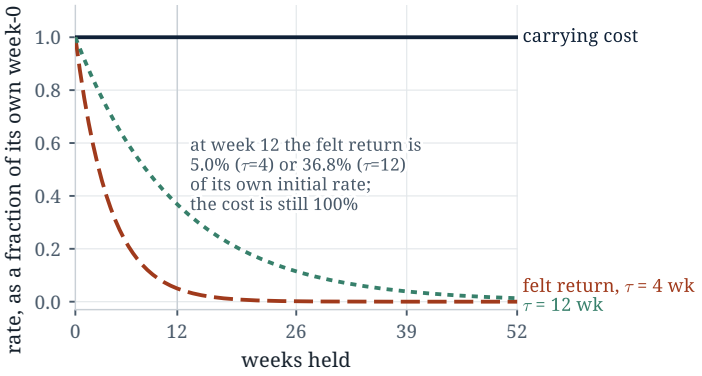


Figure 2.1: **Felt return falls; the carrying cost does not.** Modelled, not measured: the source states only that felt comfort "decays toward zero on a schedule of weeks" while carrying cost "does not decay at all", and supplies no time constant, so two are chosen here — $\tau = 4$ and $\tau = 12$ weeks — to show that the conclusion does not depend on the choice. Both series are expressed as a fraction of their own week-zero rate, so the chart carries one measure and not two. Script: `figs/src/CCP-A02_fig03.py`.

Two things must be said about the law's status. First, **the framework supplies no time constant.** It says "weeks" and nothing more. Figure 3 therefore models the asymmetry at two chosen constants rather than one, so that a reader can see that the qualitative conclusion is insensitive to the choice and that no measured rate is being smuggled in. Second, **no felt audit can catch the asymmetry**, because the felt auditor is the same difference detector that lost the goods. This is the framework's argument for a written instrument, and it is the same argument the self-report literature makes from the other side: people report the causes of their own states with a confidence unrelated to accuracy [Nisbett and Wilson, 1977], those reports are shaped by how the question was asked [Schwarz, 1999], and directing attention at a named component of well-being distorts the judgement of it [Kahneman et al., 2006].

The corpus's own storage law is the same asymmetry written as a difference equation, and it supplies the decay term the ledger lacks. The measured instance most relevant here is that doing nothing to a stock of order costs 63 per cent of it over twenty-four months, while a forced integration that raises realised non-parallelism from $u = 0.1155$ to $u = 0.3465$ — lifting surplus by 121 per cent while raising cost by 279 per cent — costs 78 per cent of it over twelve [Curzi, 2026g]. That is the ledger's law measured on a stock rather than on a feeling, and it is the corpus's own evidence that the cure can bill faster than the disease.

3.4 The pricing bias, and what this bibliography cannot support

The framework's clinical claim is that the felt price and the written price are computed by different machinery, and that only the written price can be audited. It supports the claim by naming three well-known asymmetries: that a loss weighs roughly twice a gain of the same size; that an object is overvalued by the mere fact of being held; and that a cost already paid recruits further costs in its defence. It then draws one consequence with real analytical content — that the held arrangement's costs bill in the invisible currencies while any alternative quotes its price in the visible ones at the door, so the current arrangement always *looks* cheaper than it is and every exit always looks dearer.

This corpus's external bibliography contains no entry for loss aversion, for the endowment effect, or for sunk-cost escalation. The claims are load-bearing for the framework's argument and are retained in the text; they are retained **uncited**, and they are listed individually in Section 10.3. No citation has been attached to make them appear supported. What *can* be supported from this bibliography is narrower and is stated instead: that self-reports of one's own causes are unreliable [Nisbett and Wilson, 1977], that they are question-dependent [Schwarz,

1999], and that a quantity adopted as a target for one's own management ceases to be a good measure of what it was adopted to track [Campbell, 1979, Goodhart, 1975] — an exposure that applies to the ledger itself and not only to the arrangements it prices.

The asymmetric-visibility consequence, by contrast, does not require the three uncited findings at all. It follows from R4 alone: a cost that has fallen below felt notice is not thereby cheaper, and an alternative's costs are quoted at the moment of maximum visibility. That derivation is given in the verification chain (Section 9.2) precisely so that a reader can see which parts of the argument survive the removal of the uncited material. All of it does.

3.5 The inventory procedure and its five failure modes

The framework's inventory procedure is reproduced in operational form because later papers refer to its outputs. Seven evenings of twenty minutes; a card for each holding, one to a card, in the plainness of a bin label rather than the abstraction of a title; the four cost currencies written on the back in fixed order with foreclosure last; a separate class of cards for discomforts, each assayed in the four yield currencies with **unpriced** written where nothing surfaces and no yield invented; one card audited per evening against three questions — what does it buy, what does it forbid, who would the holder be in a year if it kept its job; the whole read aloud; the cards filed in any order except importance, since ordering by importance is valuation smuggled into an inventory; and the box dated, because prices move with the season.

Two procedural rulings are falsifiable operational claims rather than preferences and are recorded as such. First, **validation precedes valuation**: the procedure forbids acting on a card during the inventory, on the ground that a hand that counts while wanting counts wrong. Second, **the resisting card**: the framework predicts that at least one card in any honest run will not accept its reverse, holds that the resistance locates an unresolved holding rather than a failure of the instrument, and requires that such a card be stood upright and left unwritten rather than forced or retired.

The framework catalogues five ways the inventory fails. They are reproduced in full, since a procedure supplied without its failure catalogue is an advertisement.

No.	Failure	Mechanism named by the source	Support in this bibliography
CC-F1	The price paid twice	Sunk-cost fallacy: past spending entered as present value	**none** (Section 10.3)
CC-F2	The stone that weighs more in the hand	Endowment effect: possession inflates felt value, so the stock is overpriced and the shelf underpriced	**none** (Section 10.3)
CC-F3	The ledger written for the auditor	Inward-turned impression management; the inventory as self-presentation	Nisbett and Wilson [1977] and Schwarz [1999] bear on it but do not state it
CC-F4	The scourge in the counting-house	The audit collapsing into rumination; counting moves, scourging loops	**none**

No.	Failure	Mechanism named by the source	Support in this bibliography
CC-F5	The drawer nailed shut	Experiential avoidance: the resisting card promoted to a forbidden one, so that every adjacent card must be mispriced to keep it plausible	**none**

The source's labels for these five are renamed here with the prefix `CC-` to keep them clear of this paper's own results R1–R7.

The framework's own closing note on the catalogue is worth keeping because it is a structural observation rather than an exhortation: none of the five failures is a failure of *finding too much*. Each is a way of writing fewer true lines while appearing to write more. The framework's stated countermeasure is plainness, spoken aloud, one item at a time, dated — which is a documentary control and is the same class of instrument the corpus uses to audit its own procedures [Curzi, 2026u].

2.0.4 4. The gradient

4.1 The tilt, not the level

The gradient material opens with a correction that everything after it rests on: a person does not experience comfort as a level but as a difference.

The physical statement is exact and is quoted exactly. For a conservative field, the force is the negative gradient of the potential,

$$\mathbf{F} = -\nabla U$$

the minus sign carrying only the direction of the push. Nothing in such a field responds to *where it is*; everything responds to how the ground falls away from where it is. The framework's claim is that a person's minute-to-minute experience has this grammar: what is felt is the differential between this posture and that one, this room and the next.

The claim is a **shared mathematical form** and not a derivation, and it is labelled as such in Section 6.1. What licenses the shape is the psychophysical result that perception is organised around a moving reference level [Helson, 1948] and that a just-noticeable difference is a ratio to the current level rather than an absolute increment [Weber, 1834]. What is *not* licensed is any inference from the physics: there is no potential energy here, no conserved quantity, and no unit.

4.2 Stationary action and its local twin (R6)

For a body moving between two fixed configurations, the action S is computed along each candidate path as the integral of the Lagrangian, and the path actually taken is the one for which the action is *stationary*:

$$\delta S = 0, \qquad S = \int L\, dt, \qquad L = T - V$$

Three statements belong with this and constitute R6's content.

First, $\delta S = 0$ is a **stationarity condition and not a superlative**. The name "least action" is a historical courtesy; the mathematics requires only that small variations of the path produce

no first-order change. A stationary path may be a minimum, and may equally be a saddle. The founding paper of this series states the same correction and this paper preserves it.

Second, **the principle licenses no teleology, and the reason is a proof rather than a caution.** The variational statement is exactly equivalent to a local differential law: the Euler–Lagrange equations, which are Newton's laws in different notation. The global-sounding statement ("of all paths, the stationary one") and the local statement ("at each instant, respond to the slope you are on") are the same equation written twice. The path does not know its destination; the path obeys its ground instant by instant, and the obedience adds up to what looks, afterwards, like design. **This corpus's external bibliography contains no classical-mechanics text, so the Euler–Lagrange equivalence is stated here without a citation** and is listed in Section 10.3. It is standard, it is checkable in any graduate mechanics course, and no entry in this file supports it.

Third — and this is the consequence the framework actually uses — **a trajectory alone cannot distinguish a chosen path from an obeyed one.** A trajectory satisfying $\delta S = 0$ is simultaneously a story of local necessity and a story of global election, and no measurement performed on the path separates them, because they are not two claims. This is a statement about the equations and it is exact. Its extension to a person's day is P1.

The source also gives one paragraph to the path-integral account, in which every path contributes with a phase turning with its action and nearly all cancel, leaving the neighbourhood of the stationary path as the survivor. **This corpus's bibliography contains no entry for that formalism either**; the paragraph is reported as the source states it, badged, and used for nothing. Every result in this section survives its deletion, which is why the verification chain does not consume it.

4.3 Desire lines as the observable (R5)

A theory of comfort differentials would remain unfalsifiable if the gradient left no tracks. The framework's claim is that it leaves tracks, and that the tracks are the node's data.

A **desire line** is the worn path that appears beside a designed one — the corner the pavement refused to cut, the diagonal across a lawn worn to the width of one person. Formally, and this is R5's first limb, a desire line is not a decision but the *integral of many small descents*: each crossing a single creature answering the local tilt between where it stands and where its errand is, none of which felt like a choice and all of which were recorded. It is gradient descent drawn on the ground by the optimiser, in the medium of grass.

That observed movement through a built setting is a more reliable survey of the setting's actual use than its designed intention is not the framework's discovery. It is the working method of a whole tradition of urban observation: the systematic filming and counting of how people actually occupy public space [Whyte, 1980], the same programme carried out as a design discipline [Gehl, 1971], the argument that a plan's legibility is not the same thing as its use and the observation that a plan drawn at a drafting table and a plan walked into existence are different objects. The framework's contribution is not the observation but the extension: **a person's day is a yard**, and the routes, the reached-for objects, and the answered letters are the same class of data.

> **The designed path is a theory. The worn path is the data.**

R5's second limb is the framework's answer to the obvious objection — that a worn path proves only the ground's habit and not the walker's gradient. The answer is an erasure test. A snowfall wipes the record to zero: no channel, no groove, no residue of the previous traffic. Within a few days the same paths reappear, in the same places, with the same bends, drawn by feet that consulted no memory but their own errands. **When a pattern survives the erasure of its medium, the pattern was not in the medium.** It was in the field of forces that drew it, and it will be drawn again in any medium offered.

The argument is sound as stated and its limits should be stated with it. It establishes that the pattern is not stored in the surface; it does not establish that the pattern is stored in a

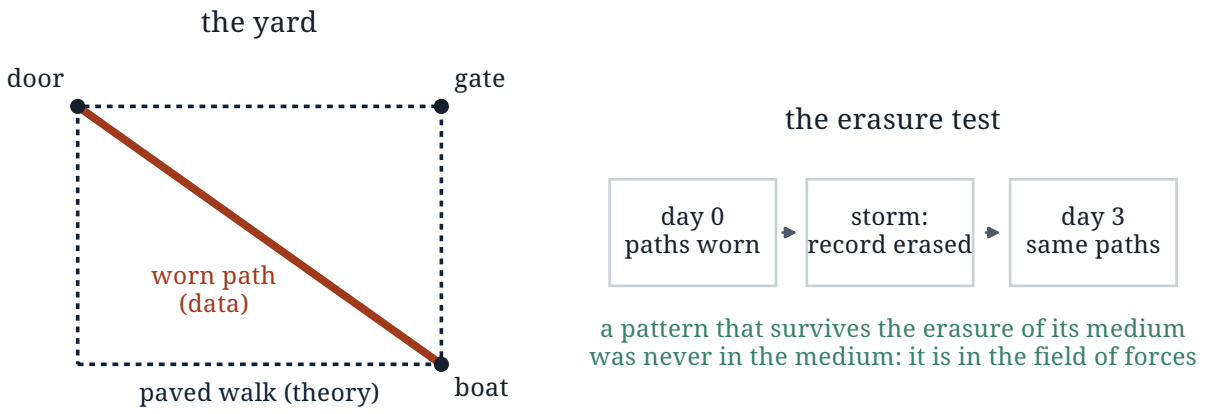


Figure 2.2: **The worn path is the survey; the erasure test is what makes it evidence.** Schematic; no plotted quantity. The left panel redraws the source's own monospace diagram of a yard, which sets a dotted designed walk around the edge and a solid worn path across the diagonal. The right panel states the source's erasure argument: a pattern reproduced after its recording medium is destroyed was never held in the medium. Script: `figs/src/CCP-A02_fig04.py`.

comfort gradient rather than in any other stable feature of the walkers or the terrain. An erasure test discriminates *record* from *generator*; it does not identify the generator. That distinction is the same one Section 6.3 draws about the framework's social claims, and it is drawn here so that R5 is not read as more than it is.

4.4 The office of the ground: four grading tools

The framework's practical prescription follows from R5 and R6 together, and it is stated as an inversion of the usual advice: effort applied at the moment of maximum tilt is the wrong instrument, and the right instrument is applied earlier, to the terrain, in a calm hour. Four tools are named.

1. **Distance.** Every reach is priced in steps. Bring the wanted thing one room nearer; send the unwanted thing one room further. The framework notes that the effective distances are surprisingly small — a drawer against a countertop, a pocket against a bag.
2. **The default.** Whatever happens when nothing is decided is the strongest current in a day, because nothing is what will most dependably be done. A default is a decision made once, in a good hour, that then stands over an unbounded number of bad ones.
3. **Friction, in both directions.** Add a step to the path that should be less worn; remove a step from the path that should deepen. The framework's ruling is that friction added without a channel opened beside it is a fence, and that fences fail — which is its third failure mode below.
4. **The first step.** A path's traffic is set less by its length than by the tilt of its first two feet. The prescription is to arrange for beginning to be smaller than not beginning.

One discipline binds the four: each is applied *once, in the calm, to the ground*, and never repeatedly, in the heat, to the self. The framework's stated test for any proposed change is a worst-day test — would this still work on a day when the person is tired, late, and out of virtue? — and its ruling is that an intervention requiring virtue is refused as a category.

None of this has a supporting entry in this corpus's bibliography. There is no entry for choice architecture, for default effects, for cue-based habit formation, or for implementation intentions. The four tools are reported as the framework's and are uncited; they are listed in Section 10.3. The framework's own case for them is R5: if behaviour is the integral of local

descents, then the leverage is on the local tilts, and that argument is internal and does not require the missing literature.

The framework also names five ways the grading procedure fails. Renamed with the CC– prefix to keep them clear of this paper's results: **CC-G1**, the survey falsified toward the surveyor's better self; **CC-G2**, the operator who accepts the map and rejects the method, attacking the steepest divergence with resolve at the hour of maximum tilt and then booking the failure to character; **CC-G3**, the fence the feet ignore — restriction without re-grading, where the water rises against the barrier, finds the gap beside it, and cuts a worse channel with resentment in it; **CC-G4**, the operator who begins choosing paths *because* they are steep, mistaking cost for worth; and **CC-G5**, the survey that becomes the residence, in which maps of maps accumulate and no object is ever moved.

CC-G3's mechanism is named by the source as reactance — that pressure on a perceived freedom produces counter-movement. **This corpus's external bibliography contains no entry for reactance.** The claim is retained because it is load-bearing for the framework's distinction between grading and fencing, and it is retained uncited, exactly as the other papers of this series retain it. What can be supported is the weaker and sufficient statement that added control and added redundancy can reduce rather than increase the safety of a coupled system [Sagan, 1993, Perrow, 1984].

4.5 The honest limit of descent (R7)

The gradient material ends by confessing its own failure condition, and the confession is R7.

Gradients are local. That is their power — they require no map, no foresight and no committee — and it is exactly their weakness: **descent finds the nearest low place, not the best one.** A system relaxes to the bottom of its own basin and stops, in perfect obedience, with a deeper basin one ridge away that it will never visit, because every direction out is uphill. The basin structure that makes the local/global distinction consequential is standard [Strogatz, 2015].

The shape of the trap is worth marking precisely, because it is what makes the condition hard to see from inside. The stranded system is *not* uncomfortable. It is at the bottom of everything it can register: every local differential points back, every small departure genuinely costs more than it returns, and so every honest instrument on the valley floor reports that nothing is wrong. Nothing *is* wrong, except the one thing no instrument on the valley floor can measure — the existence of a better valley.

R7's consequence is a limitation on the framework's own method and is stated as one: **the gradient analysis can diagnose a stable trap and cannot treat one.** Descent cannot climb out; applied effort against the gradient is the failure mode CC-G2. The only exit the framework offers that does not fight the gradient is a change in what can be seen, which re-grades the surface and lets the same descent that produced the trap carry the system out of it. That claim is stated in falsifiable form in the founding paper of this series and is untested there; nothing in this paper tests it either.

The same geometry is the corpus's account of professional over-specialisation, and the organisational literature reached the result independently: an organisation that learns fast homogenises and destroys the variance that produced its own advantage [March, 1991]. The prescription that follows is terrain rather than exhortation, and its own limit is that making a working system legible to a plan can destroy the local variety the system was running on [Scott, 1998].

2.0.5 5. The arithmetic of difference

5.1 The four requirements and the functional (R1)

The framework poses the question formally: when several viewpoints are taken on one thing, how much more do the many see than the one? Let there be N viewpoints and let $u \in [0, 1]$

measure their non-parallelism (D1), with $u = 0$ meaning all viewpoints parallel and $u = 1$ meaning each contributes a direction the others lack.

The framework states four requirements:

- **S1 (angles only).** The surplus depends only on the angles between viewpoints, and not on any viewpoint's size, weight or prestige.
- **S2 (zero at parallel).** Perfectly parallel viewpoints yield zero surplus.
- **S3 (affine effective count).** The many-viewpoint case decomposes as an effective count that grows evenly — that is, affinely — with u.
- **S4 (normalisation).** The scale is normalised, so that nothing is smuggled in through the units.

Under S1–S4 the surplus of N viewpoints over one is

$$f(u) \;=\; \ln\!\big(1 + (N-1)u\big)$$

This is R1. Read term by term: the inner expression is a count, $N_{\text{eff}} = 1 + (N-1)u$, the **effective number of viewpoints**. At $u = 0$ it equals 1 — ten parallel viewpoints are one viewpoint. At $u = 1$ it equals N. Between, it interpolates linearly. The surplus is the *logarithm of that affine effective count*: a headcount audited for redundancy, and then logged.

Two statements about the status of R1 are required and are made here rather than left to a reader.

S3 is a design axiom, not a derived property. The corpus badges it as such where it uses it: the affine effective count with $u = 1 - (x \cdot y)^2$ and surplus $S = \ln N_{\text{eff}}$ is explicitly recorded as a design axiom rather than a derived result [Curzi, 2026c], and the same badge is carried in the internal form $F_{\text{int}} = \ln(1 + (N-1)u)$ with $u = 1 - \rho^2$ from error decorrelation, where the affine form is badged as design axiom S3 rather than as a theorem [Curzi, 2026o]. **A uniqueness theorem conditional on a chosen axiom establishes that the form follows from the axiom, and nothing more.** That is the whole logical content of R1, and the corpus's own reference apparatus records it in exactly those terms: the surplus form is filed as *proved conditionally on a chosen design axiom* [Curzi, 2026].

Uniqueness-under-axioms is an established genre and this result must be positioned inside it, not presented as novel. The genre begins with the axiomatic characterisation of entropy [Shannon, 1948] and continues through the later axiomatisations of the same quantity [Aczél et al., 1974]. The framework's S1–S4 argument is a member of that genre and inherits both its strength (the form is forced once the axioms are granted) and its weakness (the axioms are the claim).

5.2 The effective count is affine, and it is not a Hill number

This subsection states a correction that applies corpus-wide and that a reviewer familiar with the diversity literature will otherwise supply.

The corpus runs **two different effective counts**, and only one of them has external grounding.

- $1/\text{HHI}$, the reciprocal of the sum of squared shares, is the Hill number of order 2, and $\exp(H)$ is the Hill number of order 1. Both are exact, well-grounded numbers-equivalents [Hill, 1973, Jost, 2006], and reinvented independently in competition economics and in political science. The corpus's concentration instrument uses exactly this reading: four holders at 40.9, 27.3, 18.2 and 13.6 per cent give a naive index of 0.293430 and 3.41 effective holders, while collapsing them into the single bloc they actually form gives an index of 1.000000 and an effective count of 1.00 [Curzi, 2026]. Both of those numbers are reproduced exactly in Section 9.2.

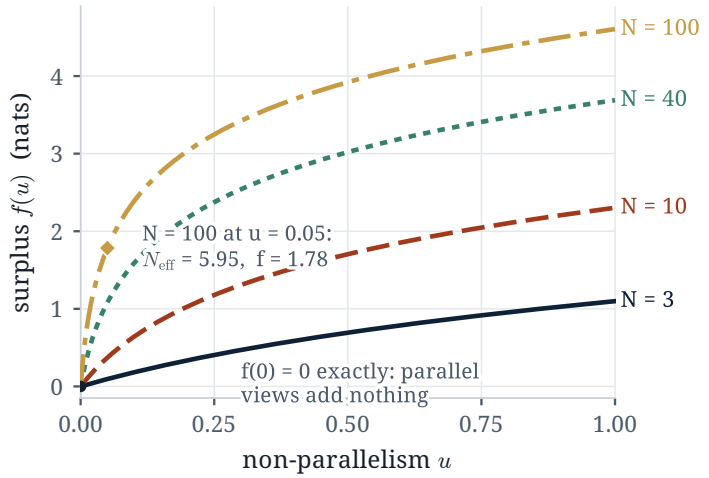


Figure 2.3: **The surplus is the logarithm of an audited headcount, and it is exactly zero when the viewpoints are parallel.** Computed from the source's own formula $f(u) = \ln(1 + (N-1)u)$ with $N_{\text{eff}} = 1 + (N-1)u$. The annotated point is the corpus's own worked instance: a hundred-member group at $u = 0.05$ has an effective count of 5.95 and a surplus of 1.78 nats [Curzi, 2026c]. Script: `figs/src/CCP-A02_fig01.py`.

- $N_{\text{eff}} = 1 + (N-1)u$ is a **linear interpolation between 1 and** N. It is not a Hill number of any order. No literature in ecology, economics, political science or information theory contains it. It is this corpus's own construction.

The distinction matters because the two are used under one name across the corpus, and a reviewer who reaches for the Hill-number literature as support will find it does not apply. The repair is cheap and is made here: **where this paper cites the Hill-number literature it does so only for** $1/\text{HHI}$ **and** $\exp(H)$**, and no reader should transfer that grounding to the affine count.**

One formalism does exist that makes a diversity index depend on the dissimilarity *between* types rather than only on their frequencies [Rao, 1982], and it is the closest external relative of what the framework is attempting. It does not reduce to the affine form, and the corpus has not shown that its u is a case of it. Stating the relationship as *unestablished* is the honest position and is adopted here.

A related warning is that "diversity" names at least three incompatible constructs — separation, variety, and disparity — so any measure of difference must declare which it is [Harrison and Klein, 2007]. The corpus's u has not declared. On the two operationalisations the corpus does supply, $u = 1-(x \cdot y)^2$ [Curzi, 2026c] and $u = 1-\rho^2$ [Curzi, 2026o], it is a *separation* measure on a squared alignment, which has a consequence the corpus states plainly: a pair in perfect opposition and a pair in perfect agreement both read $u = 0$, $N_{\text{eff}} = 1$ and surplus zero, because the extremality condition is stated on the squared alignment [Curzi, 2026r]. That is a real and counter-intuitive property and it is worth naming: **on this measure, disagreement is not difference.**

5.3 Zero at parallelism (R3)

S2 is the binding clause of the whole construction, and it is R3.

Perfectly parallel viewpoints yield $\ln(1) = 0$. Not "little": zero. A second opinion sought from a source whose training, incentives and premises replicate the first opinion's is worth exactly nothing in surplus. Agreement between aligned instruments is not evidence; it is echo.

This is the clause the corpus's downstream instruments actually run on. Stated as a procedural rule: $\ln(1.00) = 0$, so a structure with one effective bloc has zero available surplus rather

than a low one, and the first step of the method is to collapse before measuring [Curzi, 2026]. The concentration instrument's worked case is the same statement in numbers: four nominally separate holders whose independence does not survive an audit collapse to an effective count of 1.00 [Curzi, 2026].

The external literature supports the clause from several directions and is worth citing in full, because R3 is the framework's strongest claim and also its most easily overstated one. The formal statement that a group of diverse problem solvers can outperform a group of high-ability solvers exists and is much cited [Hong and Page, 2004], though it has been argued to be either trivially true or dependent on conditions amounting to the conclusion [Thompson, 2014]. The mechanisms that *destroy* u without reducing N are equally well established: informational cascades, in which later members rationally copy earlier ones and stop contributing information [Bikhchandani et al., 1992], which lowers u without lowering N and is therefore the exact failure R3 names.

The sharpest external statement of R3, however, is fifty years older than the corpus and comes from finance: **diversification benefit depends on covariance, not on count** [Markowitz, 1952]. N uncorrelated holdings and N perfectly correlated holdings have the same N and entirely different behaviour. That is R3 with the mathematics attached, published in 1952, and this paper states so rather than presenting R3 as new. The corollary is worse and is also established: correlations rise under stress, so the effective count falls exactly when it is needed [Longin and Solnik, 2001]. An audited count is not a permanent property of a structure.

One further requirement follows from R3 and is engineering rather than philosophy: a count that has not been audited for independence is worthless, because one party can present as many [Douceur, 2002]. An unaudited N is not a count; it is a claim about a count.

5.4 Concavity, and what a few views actually buy

The logarithm buys most with its first coin. Going from one viewpoint to a second, genuinely non-parallel, yields the largest single gain the set will ever see; the tenth adds a sliver and the fortieth adds almost nothing. The framework draws the practical conclusion explicitly, and the conclusion runs *against* the reflex the result is usually used to justify: the craft is not the hoarding of perspectives but the cutting of a few at genuinely divergent angles.

The framework states this as: "three honest windows, set wide, capture the greater part of what a gallery could." **That claim is not true as a comparison of counts and this paper says so.** At full non-parallelism, three viewpoints give $\ln 3 = 1.0986$ nats and forty give $\ln 40 = 3.6889$ nats, so three capture 29.8 per cent of forty and not "the greater part". The claim is true as a comparison of *angles*, which is what the framework's own arithmetic is about: three fully non-parallel viewpoints exceed forty viewpoints at any u below

$$u^{*} \;=\; \frac{2}{39} \;=\; 0.051282\ldots$$

since $1 + 39u^{*} = 3$ exactly. Forty near-duplicates at $u = 0.05$ give $N_{\text{eff}} = 2.95$ and a surplus of 1.0818 nats, which is less than three well-separated viewpoints give. **That is the defensible version of the framework's claim, and the framework does not state it in that form.** It is stated in that form here, with its break-even, and the correction is recorded in Section 7.6 rather than applied silently.

The marginal arithmetic is worth printing because it prices the framework's advice. At $u = 1$ the k-th viewpoint adds $\ln k - \ln(k-1)$: the second adds 0.6931 nats, which is 63.1 per cent of what three viewpoints reach in total; the tenth adds 0.1054; the fortieth adds 0.0253. A protocol that uses three viewpoints rather than thirty is therefore not a compromise forced by time. It is what the functional recommends.

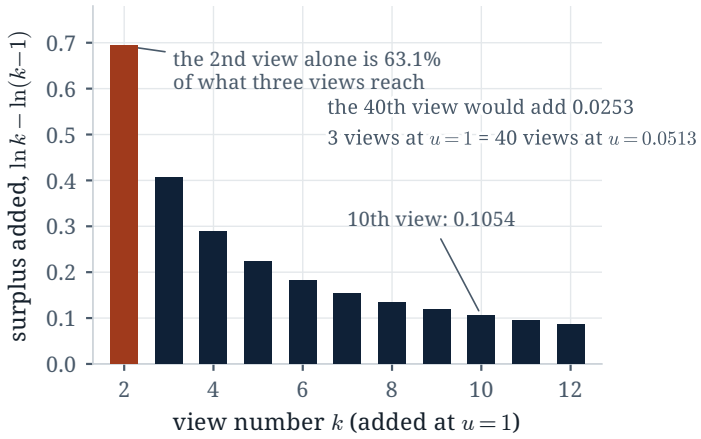


Figure 2.4: **Concavity is the argument against collecting viewpoints.** Arithmetic on the source's own formula: at full non-parallelism the k-th viewpoint adds $\ln k - \ln(k-1)$, so the second alone is 63.1 per cent of what three reach and the fortieth adds 0.0253 nats. The two annotations state the break-even the source does not compute: three fully non-parallel viewpoints equal forty viewpoints at $u = 0.0513$. Script: `figs/src/CCP-A02_fig02.py`.

5.5 A lower bound, not the information (R2)

The framework attaches one limit to R1 and treats it as doctrine rather than as a disclaimer, and this paper reports it at the same weight.

The surplus functional is a lower bound on the ensemble's information, not the information. The true information carried by N viewpoints can exceed what the formula certifies, because the formula prices only what follows from the angles alone — the bare geometry of how the viewpoints differ — and is silent about what any particular viewpoint contributes on any particular occasion. The framework's own phrase for the position is that it logs the floor and not the ceiling, and its argument for preferring that direction of error is that a floor can be built on.

This is R2, and it constrains every downstream use. A quantity reported as a bound may be used to establish that a structure has *at least* a given surplus; it may not be used to establish that a structure has *only* that surplus, nor may it be subtracted from a measured quantity as though the two were commensurable. The distinction is standard in information theory, where a bound and an attained value are different objects [Shannon, 1948, Cover and Thomas, 2006], and the corpus states the restriction in its own terms [Curzi, 2026].

Two corpus results turn on R2 being read as a bound. The corpus's headroom instrument reports realised surplus as a *percentage of the ceiling* $\ln N$ rather than as an absolute quantity, and the collective case it works reaches 57.37 per cent of that ceiling at $u = 0.25$, rising to 74.65 per cent at $u = 0.45$ [Curzi, 2026]; both percentages are reproduced exactly in Section 9.2 and neither would be interpretable if f were being read as the ensemble's information. And the corpus's structural instruments report an effective count of independence explicitly *without* claiming it measures capacity: the effective count says nothing about whether any bloc can carry the load [Curzi, 2026].

5.6 The form is underdetermined, and this paper concedes it

R1's uniqueness is conditional on S1–S4. The prior question — whether S1–S4 are the requirements the evidence supports — is open, and this paper does not close it.

Two established rival forms fit the same qualitative facts. Social impact grows as a *power* of the number of sources with an exponent below one, a result established with psychological

data in 1981 and given a dynamic cellular-automaton form in *Psychological Review* in 1990 [Latané, 1981]. Network value has been argued, from communications economics, to grow as $n \log n$ [Briscoe et al., 2006]. All three — power, $n \log n$, and $\ln(1 + (N-1)u)$ — are decelerating in the count, all three reproduce the qualitative facts the framework appeals to, and **nothing in this paper discriminates between them.** The general form of the difficulty is that many mechanisms produce indistinguishable curves, so a fitted shape is weak evidence for a generative mechanism [Stumpf and Porter, 2012]. What a reviewer will ask for is the *differential* prediction — the observation this model makes that a neighbouring model does not [Durlauf, 2005]. **This paper does not have one.**

Two things can honestly be said in the logarithm's favour.

The first is a warrant that does not require the perceptual story at all: the growth rate of a multiplicative process is logarithmic, so a logarithmic form can be a dynamical consequence rather than a psychological assumption [Peters and Gell-Mann, 2016]. Where this corpus justifies a logarithm on Fechnerian grounds [Fechner, 1966], that warrant is the weaker of the two, because it requires the perceptual claim to be true. **This paper does not establish that the comfort surplus is generated multiplicatively**; it records only which argument would be needed.

The second is negative and is therefore worth more. **No scaling or power-law claim is made anywhere in this paper.** The exponent question does not arise, because the functional's shape is *defined by the model rather than estimated from data*. Where any paper in this corpus does make a distributional or heavy-tail claim, the principled maximum-likelihood, Kolmogorov–Smirnov and likelihood-ratio procedure is binding and must be run and reported [Clauset et al., 2009]; this paper states that it makes no such claim, so that the absence of the procedure here is a scope statement and not an omission.

5.7 Convergence, invariants, and the breaking window

The single trace shows what one branch holds. A set of traces shows what holds across branches, and the audit that extracts it is R3 run in reverse.

Lay the terminal states of several non-parallel traces side by side and look for features that appear in all of them. A feature shared by *parallel* traces means nothing: the traces were bound to agree, and their agreement is one vote counted many times — which is R3. A feature shared by traces built on assumptions that contradict one another cannot be explained by the assumptions, because the assumptions disagree about everything else. The framework calls such a survivor a **horizon-fact** and is careful about what it is a fact about: not about the future, but about the composition — the system that all the traces were obliged to run on.

The framework then states its own strongest objection to the procedure, and states it as a mandatory step rather than a caveat. **Convergence can also mean the viewpoints secretly share an angle.** All of them were cut by one person, with that person's practised blindnesses, and an invariant may be nothing more than a habit that reached every trace. The required check is to attempt, in earnest, one further honest trace that lands *without* the invariant — the assumption a stranger would choose. If it can be built, the invariant is struck. If it genuinely cannot, the invariant is promoted, carrying the warrant *tested against the best available attempt to break it*. That is a severe-test discipline in the strong sense: a design that discriminates between named alternatives rather than accumulating agreement with one and it is the same standard the preregistration literature imposes on a prediction [Nosek et al., 2018].

The framework's own worked case is reported here because it shows what the procedure returns and because it contains a refusal, which is the part usually omitted. Three traces of one person's next three years, built on contradictory assumptions — continuation, a single named intervention, and the branch the person had been declining to look at — landed in three settings that agreed on nothing: not the premises, not the finances, not the job title. One concrete feature stood in all three. The mandatory fourth trace, built on an assumption chosen to break that feature, could not be landed without it, so the feature was promoted.



Figure 2.5: **Three traces built on contradictory assumptions agreed on exactly one feature, and the mandatory breaking trace could not be landed without it.** Measured from the source's own case record, which states that the three landed futures agree on neither the premises, nor the finances, nor the job title, and names the single surviving feature. Each row is a concrete feature the record names; each column a trace the record reports. Script: `figs/src/CCP-A02_fig05.py`.

This is a single self-reported case and it establishes nothing about the procedure's validity. It is offered as an illustration of the procedure's output. The base rate against which any single reported psychological observation should be read is not encouraging [Open Science Collaboration, 2015].

2.0.6 6. Method: what kind of borrowing, what scale, and what is identified

6.1 Analogy discipline

Every physics borrowing in this paper is labelled at first use as one of three things: a **derivation**, in which the physical result is used because the system genuinely satisfies its premises; a **shared mathematical form**, in which the same equation governs both with no claim that the underlying physics applies; or an **analogy**, a suggestive correspondence with no formal warrant, legitimate as motivation and never as evidence. The labels are collected here so that a reader can check them against the text.

Borrowing	Category	Warrant, and what is not claimed
$\mathbf{F} = -\nabla U$ (Sec. 4.1)	2 — shared form	The statement is exact for a conservative field. No claim that a comfort landscape is a potential, that anything is conserved, or that any unit transfers.
Stationary action $\delta S = 0$ (Sec. 4.2)	3 — analogy	No action functional of a life is defined and no Lagrangian is exhibited. The borrowing supplies the image of a resting trajectory. The variational form licenses no teleology.

Borrowing	Category	Warrant, and what is not claimed
The Euler–Lagrange equivalence (Sec. 4.2)	1 — derivation, **within mechanics only**	A genuine theorem of classical mechanics, used only to establish that the global statement carries no teleology. **Uncited: this bibliography holds no mechanics text.**
The path-integral account (Sec. 4.2)	quoted, then unused	Reported as the source states it. **Uncited.** No result consumes it; the chain in Section 9.2 does not contain it.
Gradient descent and local minima (Sec. 4.5)	2 — shared form	The convergence and trapping behaviour hold for any differentiable cost. The transfer claim is that behaviour has this recursion's shape, not its physics [Strogatz, 2015].
Carrying cost as a stock outflow (Sec. 3.1)	2 — shared form	Levels and rates are the same mathematics wherever the level comes from [Hopp and Spearman, 2011]. No claim that comfort is inventory.
The logarithm in $f(u)$ (Sec. 5.1)	2 — shared form, contested	An information-theoretic form used on a quantity that has not been shown to be an information measure. The alternative warrants are named in Section 5.6.
"Effective count" (Sec. 5.2)	**1 for $1/\mathrm{HHI}$, 3 for the affine count**	Hill numbers are exact numbers-equivalents; the affine count is this corpus's own construction with no precedent, and the two are kept apart throughout.

The corpus's title claims physics, and the claim attracts a specific literature of objection which is better cited than ignored. The historical form is that economics imported the formal apparatus of nineteenth-century energetics without the conservation law that gave energy its content [Mirowski, 1989], which is precisely the hazard for a framework whose central object is a stock that accumulates, leaks and is spent: **this paper claims nothing is conserved, and states that no invariance has been exhibited beyond the one in Section 9.5.** The general form of the objection is that laws hold in shielded local arrangements rather than universally, so importing a law's form is not importing its scope; and the field this corpus's title claims membership of has been surveyed and bounded, with the vocabulary the corpus uses given definitions it can be held to [Castellano et al., 2009]. The criterion the discipline named in the corpus's title will actually apply is that an explanation must name a mechanism rather than exhibit a fitted curve, and the standing answer to the charge that such work forecasts nothing is to say which of the other purposes of modelling it is claiming. This paper claims explanation of structure and not prediction.

Where a result carries a constant of this corpus's own construction — the minimum sustainable frequency $f_{\min} = k\,e^{2c/\eta}$ with $k = 0.483425732988414$ Hz [Curzi, 2026] — two cautions apply and are recorded once here rather than at each mention. That constant is **defined**, exact by construction as $(9/8)^2/\varphi^2$, and its fifteen significant figures are defensible only on that basis; **attaching hertz to it is a separate and much stronger claim than computing it**, and this paper endorses neither the attachment nor any reading of it as a discovered natural constant. The general hazard of reading a chosen golden-section ratio as a fact about the world is documented and should be read before any such reading is offered [Markowsky, 1992].

6.2 Scale type, and the objection that bites hardest

Everything in Section 3 that adds, subtracts or compares carrying costs presupposes at least an interval scale, and everything in Section 5 that reads a *ratio* of surpluses presupposes a ratio scale. Nothing in the source establishes either.

This is the framework's sharpest exposure and the objection has a name and an author: psychology systematically assumes that its attributes are quantitative without ever testing whether they are [Michell, 1997]. **This is the strongest single objection to the framework and it is not answered here.** The formal test for whether an attribute has additive quantitative structure is conjoint measurement [Luce and Tukey, 1964], whose single, double and triple cancellation conditions are exactly what would have to be checked, against the standard set out in the measurement-foundations literature [Krantz et al., 1971]. **They were not tested.** A route that does not require respondents to use numbers at all is available and unused [Rasch, 1960]. The classification of scales that fixes which operations each licenses is Stevens's [Stevens, 1946], and the century-long dispute over whether the psychophysical function is logarithmic or a power law is itself evidence of how hard the underlying question is [Stevens, 1957].

The honest position, adopted here: carrying cost and comfort are treated as quantitative by stipulation, not by demonstration. The corpus applies the same demotion to itself elsewhere and the demotion is the correct instinct: because the moment equation contains a division and a multiplication, ordinal data for either of its inputs demote it to a heuristic ordering that may not be printed as a fraction or computed [Curzi, 2026m]. **That demotion rule applies to this paper's own material and is applied.** Specifically:

- $f(u)$ and N_{eff} are licensed at ratio scale **only where u is computed from a quantity with its own ratio structure** — a squared inner product of unit vectors [Curzi, 2026c] or a squared correlation [Curzi, 2026o]. Both of those are dimensionless numbers on $[0, 1]$ with a true zero, so the arithmetic of Section 5 is licensed *by them* and not by this framework.
- Where u is instead elicited as a judgement of how different two viewpoints are, it is at best ordinal, and $\ln(1 + (N-1)u)$ may then be used only as a monotone ordering. It may not be printed as a number of nats.
- Carrying cost in *coin* and in *hours* has a natural unit and licenses arithmetic. Carrying cost in *attention* and *foreclosure* has none, and the ledger's four-way total is therefore **not a sum of commensurable quantities.** The source prints the four currencies as a single ledger entry; this paper records that they do not add.

A further formal obstacle applies to any similarity-based reading of u: psychological similarity is asymmetric and violates the triangle inequality, so it is not a metric [Tversky, 1977]. Where u is intended as one minus a similarity, that result bears directly on whether u is even well defined.

6.3 Identification: the reflection problem, and the framework turned on itself

Section 5's arithmetic is used across this corpus on *groups*: departments, panels, blocs. The moment it is used that way it becomes a claim about social interaction, and a claim about

social interaction invites the question of whether its parameter could be recovered from observation. It could not, and this is stated here in Method rather than in a footnote, because it fixes what R1 and R3 can be.

The corpus's central formal object is the reflexive loop $T(x) = W(x, M(x))$, fixed in the founding paper of this series. Written as an aggregate relation — a group's state influencing its members while the members constitute the group — **that is a linear-in-means reflexive specification, and the endogenous social effect in such a specification is not identified from observational data**: it cannot be separated from contextual effects or from correlated unobservables, however much data is collected [Manski, 1993]. The network route is closed by the same family of result, since influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011].

The consequence for this paper is specific. A group's non-parallelism u is *definable* — D1 gives two constructions — and, from observational data on a group, it is not *estimable* as a causal quantity: any association between a measured u and a measured outcome is not separable from the reasons people with different views end up in the same group. The available remedies are randomisation, an instrument, partial-population variation, or stipulating the reference group and defending the stipulation, the last having an explicit template in a discrete-choice model that states the conditions under which its interaction parameter is econometrically identified [Brock and Durlauf, 2001]. **This paper adopts none of them.** R1 and R3 are therefore **modelling positions rather than identified empirical results**, and every downstream paper citing them inherits that status. An interaction parameter that is merely definable is not the same object as one that is estimable, and this corpus has so far produced only the first.

That reflexivity of this shape *occurs* is separately documented and the documentation matters because it is the existence proof the framework otherwise lacks: when the Black–Scholes option-pricing model was published its prices did not fit the market, and after traders adopted it the market came to fit the model, with dates and price data recording the convergence [MacKenzie and Millo, 2003]. The more general statement — that classifications of people alter the people classified, so the referent moves — has been argued directly [Hacking, 1995]. That literature supplies existence, not identification, and the distinction is the one this subsection exists to keep.

The framework turned on itself. The corpus asserts that systems change their behaviour when modelled and observed. Applied consistently, that assertion falls on the corpus's own parameters: a u, a decay rate or a cost coefficient estimated in a population that has not read the framework cannot be assumed to survive in one that has, because the decision rules those parameters summarise are themselves responses to the environment the model describes [Lucas, 1976]. This is not an outside objection; it is the corpus's own thesis pointed at itself. The position taken here, stated so that it can be held against the paper: **no stability of this framework's parameters under its own dissemination is claimed, and no result in this paper depends on any.** A further and independent aggregation problem applies wherever a single u is written for a group whose members differ, and this paper does not address it.

2.0.7 7. Results

7.1 Definitions fixed. D1 (non-parallelism), D2 (the affine effective count) and D3 (carrying cost) are stated in the form the remaining papers cite. None is a finding. D2 in particular carries a modelling choice — the affine form — inside what reads as a definition, and Section 5.1 badges it.

7.2 Results established by derivation, conditional on the axioms. R1 is established *given* S1–S4, and the conditionality is the result's whole content: the functional follows from the requirements, and the requirements are a choice. R2 (lower bound) and R3 (zero at parallelism) follow from R1 by inspection. R6 (the variational form is equivalent to the local law, and licenses no teleology) is a theorem of classical mechanics, established, and **uncited in this bibliography**. R7 (descent finds the nearest low place) is standard.

7.3 Results established by construction. R4 (the four-currency carrying cost against the decay asymmetry) and R5 (the worn path as observable, authenticated by the erasure test) are specifications and arguments rather than measurements. R5's erasure argument is valid for what it claims — that the pattern is not stored in the medium — and does not identify the generator, which Section 4.3 states.

7.4 Hypotheses tested. None. H1 through H5 are stated in falsifiable form and none is tested here. No dataset was collected and no instrument capable of eliciting an attention charge or a foreclosure charge exists, for the measurement-theoretic reasons in Section 6.2. This is a negative result and is reported in the same plain language a positive one would receive: **the five hypotheses this paper states stand untested, and the paper that states them supplies no evidence for any of them.**

7.5 Reproduction of the corpus's own arithmetic. Five numerical claims that this paper's internal citations carry were recomputed. Four reproduce and one does not.

- The concentration instrument's naive index over shares 40.9, 27.3, 18.2 and 13.6 per cent is $0.409^2 + 0.273^2 + 0.182^2 + 0.136^2 = 0.293430$ exactly, and its reciprocal is 3.40797, printed as 3.41 [Curzi, 2026]. **Reproduced.**
- The option-value ceiling $p = (c + 0.0051096)/0.0318401$ reaches $p = 1$ at $c = 0.0318401 - 0.0051096 = 0.0267305$, and $0.0267305/0.30 = 0.089102$, printed as 8.9 per cent of the maintenance budget [Curzi, 2026]. **Reproduced.**
- The headroom instrument's realised-surplus percentages are reproduced exactly at $N = 18$: $\ln(1 + 17 \times 0.25)/\ln 18 = 1.6582281/2.8903718 = 0.573708$, printed as 57.37 per cent, and $\ln(1 + 17 \times 0.45)/\ln 18 = 2.1575593/2.8903718 = 0.746464$, printed as 74.65 per cent [Curzi, 2026]. **Reproduced**, and the reproduction also fixes $N = 18$, which the cited summary does not state.
- The same source's equilibrium at $u = 0.25$, $(0.30 \times 1.6582 - 0.55)/0.06 = -0.8755$, reproduces to -0.875526 when the unrounded surplus 1.6582281 is used, confirming that the printed four-decimal figure is correct and that the *displayed* 1.6582 is a rounding of the input rather than the input itself. **Reproduced, with that clarification.**
- **Not reproduced:** the companion figure for the same case at $u = 0.45$. Under the parameters as quoted, $(0.30 \times 2.1575593 - 0.55)/0.06 = +1.6211$, not the $+3.2878$ the summary states. Recovering $+3.2878$ requires a surplus of 2.490893, that is $u = 0.6513$ at $N = 18$, so either a parameter other than u also moves in that scenario or the figure is in error. **This paper does not resolve which**, and cites only the reproduced parts of that result. The discrepancy is printed here rather than smoothed over.

7.6 A correction to the source. The source's summary claim that three well-cut viewpoints "capture the greater part of what a gallery could" is false as a comparison of counts — three capture 29.8 per cent of forty at $u = 1$ — and true as a comparison of angles, with break-even at $u^* = 2/39 = 0.051282$ (Section 5.4). The corrected form is what this paper states and what downstream papers should quote.

2.0.8 8. Discussion

8.1 What follows

Three consequences follow from R1–R3 taken together, and they are the framework's genuine yield in this instalment.

The first is a rule for adding a viewpoint: the question is never whether the new viewpoint is *good* but what angle it holds that no viewpoint already present holds. R3 makes that a matter of arithmetic rather than of judgement, and it converts a familiar piece of advice into a check that can fail.

The second is a rule against accumulation. Concavity means that beyond a small number of genuinely divergent viewpoints, the marginal purchase is negligible, and effort is better spent raising the angles between the viewpoints already held. The corpus's own measurement

makes the same point from the cost side: raising realised non-parallelism by a factor of three raised surplus by 121 per cent and cost by 279 per cent, which made the position worse rather than better [Curzi, 2026g]. **Raising u is not free, and the framework's monotone reading of its own functional obscures that.**

The third is diagnostic and belongs to Section 4. If behaviour is the integral of local descents, then a position that will not yield to exhortation is not evidence of a defect in the person occupying it, and the intervention that follows is on the terrain rather than on the will. What R7 adds is the limit: where the trap is stable, the analysis can locate it and cannot open it.

8.2 What does not follow

That a functional is unique under four requirements does not make the requirements right. That an effective count is called a count does not make it a Hill number. That a worn path is data does not make the gradient the only generator that could have drawn it. That the corpus's arithmetic reproduces does not make the quantities it operates on ratio-scaled.

Two further non-implications are worth stating because the corpus's later results tempt them. First, **the surplus law is monotone increasing in u and the evidence is not.** Supporting: novel information arrives across weak, dissimilar ties [Granovetter, 1973]. Opposing, and unavoidable: creative success rises and then *falls* with cohesion, an inverted U rather than a monotone gain [Uzzi and Spiro, 2005]; in the largest natural setting studied, short-run diversity reduced trust and civic participation, a finding its author did not want [Putnam, 2007], though the effect largely disappears under better controls [Abascal and Baldassarri, 2015]; in randomly assembled ecological systems increasing diversity and connectance *destabilises* [May, 1972]; and a diversity–stability correlation can arise from statistical averaging with no mechanism at all [Doak et al., 1998]. **A monotone claim cannot survive that list unqualified, and this paper does not assert one.**

Second, the ledger is itself an indicator, and an indicator adopted as a target degrades [Campbell, 1979, Goodhart, 1975]. Nothing in the framework protects its own instrument from that, and the exposure is structural rather than incidental.

8.3 Relation to the corpus

The functional fixed here is consumed downstream in a specific and checkable way. The engineering series applies it to rosters and reports the affine effective count with $u = 1 - (x \cdot y)^2$, badged as a design axiom and worked at a hundred members and $u = 0.05$ to an effective count of 5.95 [Curzi, 2026c]. The same form is applied inside a single observer, with $u = 1 - \rho^2$ from error decorrelation and the affine form badged as design axiom S3 [Curzi, 2026o]. The dyadic case supplies R3's sharpest instance: opposition and agreement both read $u = 0$ [Curzi, 2026r]. The headroom instrument feeds f into an equilibrium and reports realised surplus as a fraction of $\ln N$ [Curzi, 2026]. The concentration instrument supplies the *other* effective count, exactly grounded, and the bloc collapse that makes the difference between the two visible [Curzi, 2026]. The consolidated method restates R3 as its first procedural step: collapse before measuring, because $\ln(1.00) = 0$ [Curzi, 2026]. And the option-pricing case shows what the arithmetic looks like when a marginal viewpoint is priced against its own cost ceiling [Curzi, 2026].

Three results elsewhere in the corpus bound what this paper may claim. The corpus's own falsification suite measured its shuffled-input control and found it **anti-conservative**: against a nominal 0.05 across 500 null trials, independent permutation rejected at 0.108, circular shift at 0.120 and block permutation at 0.060 [Curzi, 2026u]. That is the corpus's strongest evidence that its own controls can mislead, and it is the correct prior for any future test of H1–H5. The corpus's demotion of its moment equation on scale-type grounds is the rule this paper applies to its own four-currency total [Curzi, 2026m]. And the corpus's damage integral — in which eight periods of the wrong lever order cost twelve periods of permanent lag, a minimum stock

of -2.373029 from $+1.000000$ and 111.133601 stock-periods of cumulative gap [Curzi, 2026] — is the quantified form of R7's warning that a correctly diagnosed trap can still be treated in the wrong order.

Two further results are cited for the structure they share rather than the domain they occupy. The consonance–redundancy identity, in which the minimum essential harmonic distance of 2.0000 is attained by exactly the six collision pairs $\{1{:}2,\ 2{:}4,\ 3{:}6,\ 4{:}8,\ 5{:}10,\ 6{:}12\}$ and by nothing else [Curzi, 2026], is R3 in a different alphabet: exact coincidence yields exactly the degenerate value, and near-coincidence does not. And the pitch criterion $f_{\min} = k\,e^{2c/\eta}$ [Curzi, 2026] is the corpus's clearest instance of a defined constant carrying a unit it was not derived with, which is why the cautions of Section 6.1 are attached to it here.

2.0.9 9. Verification

This section is not a summary. It decomposes the paper's results into steps small enough that a reader can check each one without checking any other, and labels each `ARITHMETIC` (a reader can redo it), `DEFINITION` (true by this paper's own construction), `CITED` (a named external source supports it) or `POSTULATED` (asserted here, no external support). The razor is applied: any postulate whose removal leaves the results standing has been removed.

9.1 Postulates

- **P1 (DEFINITION).** $u \in [0, 1]$ measures non-parallelism, with $u = 0$ meaning parallel and $u = 1$ meaning fully non-parallel. *Check:* D1.
- **P2 (DEFINITION, and a modelling choice).** $N_{\text{eff}} = 1 + (N - 1)u$: the effective count is **affine** in u. *Check:* D2. **Note: this is the paper's largest unearned assumption. It is a design axiom, it is not a Hill number, and no external literature contains it (Section 5.2).**
- **P3 (DEFINITION).** The surplus depends only on the angles between viewpoints (S1), is zero at parallel (S2), and is normalised (S4).
- **P4 (ARITHMETIC).** Given P2 and P3, $f(u) = \ln N_{\text{eff}}$ is the unique admissible form. *Check:* the uniqueness argument, which is a member of the established uniqueness-under-axioms genre [Aczél et al., 1974].
- **P5 (ARITHMETIC).** $\ln 1 = 0$. *Check:* redo it.
- **P6 (ARITHMETIC).** ln is strictly concave, so $\ln k - \ln(k-1)$ is strictly decreasing in k. *Check:* redo it.
- **P7 (CITED: `cover2006elements`).** A bound on an information quantity and the quantity itself are different objects, and a bound may not be substituted for a measurement.
- **P8 (POSTULATED).** The felt return on a held comfort decays toward zero over weeks while its carrying cost does not decay. *Check:* **not checkable from outside the framework and not checked. This is the substantive content of H2 and it is asserted.** Its perceptual half has partial support [Helson, 1948, Brickman et al., 1978]; its cost half has none.
- **P9 (POSTULATED).** A repeated route through a setting is generated by a comfort differential rather than by any other stable feature of the walker or the terrain. *Check:* **not checked.** The erasure test rules out the recording medium and nothing else (Section 4.3).
- **P10 (CITED: `strogatz2015nonlinear`).** Gradient descent converges to a local, not a global, minimum of the surface it is run on.
- **P11 (POSTULATED).** A person's day relaxes through a comfort landscape as a physical system relaxes through its states. *Check:* not checkable; this is the framework's relaxation postulate and is badged wherever it appears.

Postulates dropped as unnecessary: the path-integral account of Section 4.2 was tested for load and carries none — every result below survives its removal — and so did the three uncited pricing biases of Section 3.4, whose conclusion is re-derived from P8 alone. Both are reported in the text and neither is used.

9.2 The compounding chain

Each line consumes only the outputs named on its left.

P1 + P2	-> L1	N_eff = 1 + (N−1)u is defined	DEFINITION
L1 + P3 + P4	-> R1	f(u) = ln(1 + (N−1)u)	DEFINITION (via P2)
R1 + P5	-> R3	f(u) = 0 exactly at u = 0	ARITHMETIC
R1 + P6	-> H5	the k−th view adds ln k − ln(k−1), which is decreasing in k	ARITHMETIC
R1 + P7	-> R2	f is a LOWER BOUND, not the information	CITED
P8	-> R4	felt return decays, carrying cost does not; the gap widens	POSTULATED
R4	-> L2	a held arrangement's costs are underquoted relative to an alternative's	ARITHMETIC
P9	-> R5	the worn path is the observable	POSTULATED
(mechanics)	-> R6	stationarity is equivalent to the local law; no teleology	CITED (no entry in this file)
P10	-> R7	descent reaches the nearest low place, not the best	CITED
R7 + P11	-> C1	a person can be stably trapped at a locally optimal arrangement	POSTULATED (via P11)

What this chain shows, stated plainly. Of the seven results, exactly three — R3, R2 and R7 — rest on ARITHMETIC or CITED steps alone once R1 is granted. R1 itself passes through P2, the affine axiom, and is therefore DEFINITION at its load-bearing step: **R1 is not a discovery about viewpoints; it is what follows from choosing an affine effective count.** R4 and R5 are POSTULATED. R6 is a genuine theorem that this bibliography cannot cite. C1 — the conclusion the framework is built toward — passes through P11, which is not checkable and was not checked.

A chain whose steps are mostly DEFINITION and POSTULATED is not a proof, and this one is mostly DEFINITION and POSTULATED. What it is instead is a specification of exactly which three premises would have to be established for the framework's conclusions to follow: P2 (that the effective count is affine rather than some other interpolation), P8 (the decay asymmetry), and P9 (that a worn path is generated by a comfort differential). That is the whole bill, and it is smaller and more addressable than the framework's rhetoric suggests.

9.3 The checks a reader should run first

1. **Redo** $\ln 1 = 0$ and confirm that R3 is arithmetic and not rhetoric. It is the clause the corpus's concentration instruments actually run on.
2. **Redo the concentration arithmetic.** $0.409^2 + 0.273^2 + 0.182^2 + 0.136^2 = 0.293430$, and $1/0.293430 = 3.40797$. The corpus prints 0.293430 and 3.41 [Curzi, 2026]. It agrees.
3. **Redo the affine count.** $1 + 99 \times 0.05 = 5.95$, and $\ln 5.95 = 1.783391$. The corpus prints 5.95 [Curzi, 2026c]. Figure 1 marks the point.
4. **Redo the break-even of Section 5.4.** $1 + 39u = 3$ gives $u = 2/39 = 0.051282$, and $\ln(1 + 39 \times 0.05) = \ln 2.95 = 1.081805 < \ln 3 = 1.098612$. Figure 2 states both.
5. **Redo the headroom percentages.** $\ln(1 + 17 \times 0.25)/\ln 18 = 0.573708$ and $\ln(1 + 17 \times 0.45)/\ln 18 = 0.746464$, printed as 57.37 and 74.65 per cent [Curzi, 2026]. They agree; the accompanying equilibrium at $u = 0.45$ does not (Section 7.5).
6. **Check that** $k = 0.483425732988414$ **is labelled defined at every occurrence**, and that no sentence in this paper attaches hertz to it on this paper's own authority. There is one occurrence, in Section 6.1, and it carries both cautions.

7. **Check that R1 and R3 are never described as measured**, in this paper or in any paper citing them. They are modelling positions; Section 6.3 says so and the Conclusion repeats it.

9.4 Dimensional audit

Every symbol used above, with its dimension, its unit, and its epistemic status. The audit is short because few of these quantities have units at all, and the table says so rather than inventing any.

Symbol	Meaning	Dimension	Unit	Status
N	number of viewpoints	dimensionless	count	defined; an integer
u	non-parallelism	dimensionless	none	defined on $[0, 1]$; two operationalisations exist, both stipulated
N_{eff}	effective count	dimensionless	count	**defined, affine, not a Hill number**; not measured anywhere
$f(u)$	surplus	dimensionless	nats	derived from N_{eff}; a bound, not a measurement
$\ln N$	surplus ceiling	dimensionless	nats	derived; $f \leq \ln N$ with equality at $u = 1$
coin cost	monetary carrying cost	currency	currency units	measurable in principle; has a natural unit
hours cost	appointed time	time	hours	measurable in principle; has a natural unit
attention cost	standing awareness charge	none assigned	**no operational unit exists**	defined (D3); **not measurable by any instrument the framework supplies**
foreclosure cost	options closed by holding	none assigned	**no operational unit exists**	defined (D3); not measurable
U	potential (quoted)	energy	J	quoted from physics; not a quantity of this framework
F	force (quoted)	energy per length	N	quoted from physics; not a quantity of this framework
S, L	action, Lagrangian (quoted)	energy $\times$ time; energy	J·s; J	quoted from physics; not quantities of this framework

Symbol	Meaning	Dimension	Unit	Status
τ	felt-return time constant (Fig. 3)	time	weeks	**chosen, not measured; the source supplies none**
k	reference constant (quoted)	see note	quoted as Hz	**defined exactly as $(9/8)^2/\varphi^2$; the hertz is a separate claim**

Does the central relation balance? $f(u) = \ln(1 + (N-1)u)$ balances only if the argument of the logarithm is dimensionless, and it is: N is a count and u is a dimensionless number on $[0, 1]$, so N_{eff} is a pure number and f is in nats. **That is the one equation in this paper whose dimensions are unambiguous**, and it is unambiguous because both of its inputs were defined to be dimensionless rather than measured.

Does the ledger balance? No, and the source is silent on why. The four-currency carrying cost is presented as one accounting of a holding, but coin has a currency unit, hours has a time unit, and attention and foreclosure have none. A sum across the four is not a sum of commensurable quantities and cannot be formed. The framework does not in fact form one — it writes four lines and never adds them — but it also never says that they cannot be added, and any downstream reader who totals them is doing something the dimensions forbid. **This paper says it.**

A quantity with no operational unit is entered as having none. No unit was invented for attention or for foreclosure, and none should be inferred from the fact that the surplus functional balances.

9.5 Limits, boundary behaviour, and the one invariance

Each relation above, at the boundaries of its parameters.

- $u \to 0$**.** $N_{\text{eff}} \to 1$ and $f \to 0$ for every N. This is R3 and it is exact, not asymptotic: at $u = 0$ the surplus is zero and not small. A group of any size at $u = 0$ is one member.
- $u \to 1$**.** $N_{\text{eff}} \to N$ and $f \to \ln N$, the ceiling. The corpus's headroom instruments report realised surplus against exactly this ceiling [Curzi, 2026].
- $N = 1$**.** $N_{\text{eff}} = 1$ for every u, and $f = 0$. The functional is well behaved and says the obvious thing: a single viewpoint has no surplus over itself regardless of what u is set to. Note that u is undefined at $N = 1$, since there is no angle to measure; the formula's insensitivity to it at $N = 1$ is what makes that harmless.
- $N \to \infty$ **at fixed** $u > 0$**.** $f \to \infty$ logarithmically and without bound. **This is a constraint on the model's domain and is stated as one.** Nothing in the framework bounds the surplus of an arbitrarily large panel, and the empirical literature says the curve turns over rather than continuing to rise [Uzzi and Spiro, 2005]. The domain should be restricted to the range over which a cost term is also modelled; the corpus does this in its engineering papers by pairing f with an explicit cost [Curzi, 2026g,?] and does not do it here.
- $N \to \infty$ **at fixed** N_{eff}**.** Requires $u \to 0$ as $1/N$: an unbounded panel of near-duplicates. The functional handles this correctly and it is the regime the corpus's bloc collapse is designed to detect [Curzi, 2026,?].
- **Carrying cost** $\to 0$ **in all four currencies.** The framework's ruling is that this never occurs and that a written zero is a failure of the second look. **That is not a boundary condition; it is an unfalsifiable rule**, and it is restated in Section 10.4 as such.
- **Felt return with no decay.** R4 becomes vacuous, and the framework's account of an accumulated depletion has no mechanism. H2 is therefore the hypothesis whose failure would remove the ledger's entire analytical content, which is why it is stated first among the falsifiable ones.

- **Barrier height $\to 0$ around a local minimum (Section 4.5).** The basin ceases to be a basin and R7 is vacuous. **Barrier height $\to \infty$.** The trap is total and descent cannot leave it; the framework's only remaining exit is a change in what can be seen, which is untested.

On invariance. The framework claims no scale-freedom and this paper claims none for it. The one invariance exhibited is that f is unchanged under any relabelling of the viewpoints and under any transformation preserving the pairwise angles; it is a consequence of S1 rather than a finding, and it has been checked at no scale because nothing here is measured. **No conserved quantity is exhibited anywhere in this paper**, which is the specific answer this corpus owes to the charge that it has borrowed the grammar of a conserved quantity without the conservation law [Mirowski, 1989].

2.0.10 10. Limitations

Each item names what would falsify the claim, or what the paper cannot support.

10.1 The affine effective count is this corpus's own construction and has no external precedent. This is the paper's principal exposure and it is stated first. $1 + (N-1)u$ is a linear interpolation between 1 and N. It is not a Hill number of any order, and no literature in ecology, economics, political science or information theory contains it (Section 5.2). Where this paper cites the Hill-number literature it does so only for $1/\text{HHI}$ and $\exp(H)$ [Hill, 1973, Jost, 2006]. **Falsified by:** a demonstration that some other interpolation — geometric, or a Hill number of some order — fits an outcome better at equal parameter count. Nothing here rules one out.

10.2 The functional's form is underdetermined. A power law with exponent below one [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the same qualitative facts. **Nothing in this paper discriminates between them**, and the paper has no differential prediction of the kind a reviewer will ask for [Durlauf, 2005]. The strongest available warrant for a logarithm is dynamical rather than psychophysical [Peters and Gell-Mann, 2016], and it applies only where the underlying process is multiplicative, which has not been established for any quantity in this paper.

10.3 Claims with no supporting literature in this corpus's bibliography. Listed individually, because a reviewer will find them and finding them first is the point.

- *Losses weigh roughly twice equal gains; a held object is overvalued by the holding; a paid cost recruits fresh costs in its defence.* Load-bearing for Section 3.4 and for failure modes CC-F1 and CC-F2. **`bib/external.bib` contains no entry for loss aversion, for the endowment effect, or for sunk-cost escalation.** Retained **uncited**; the conclusion that depends on them is re-derived from R4 alone (Section 9.2).
- *Pressure on a perceived freedom produces counter-movement (reactance).* Load-bearing for failure mode CC-G3 — the fence the feet ignore. **No entry.** Retained uncited, as it is in the other papers of this series.
- *Cues placed in the path of a behaviour govern it more reliably than resolutions about it; if-then plans tied to concrete cues outperform bare intentions; defaults govern outcomes.* The entire warrant for the four grading tools of Section 4.4. **No entry for choice architecture, default effects, implementation intentions or habit formation.** Reported as the framework's and uncited.
- *A held comfort draws on attentional capacity while untouched.* The attention currency of Section 3.1. Miller [1956] and Ashby [1956] establish that capacity is bounded; **neither establishes that an unattended holding consumes it.** The stronger claim is uncited.
- *The Euler–Lagrange equivalence of the variational and local formulations.* R6. **This bibliography contains no classical-mechanics text.** The result is standard and is stated here without a citation rather than with a wrong one.
- *The path-integral account of stationarity.* **No entry.** Reported once, used for nothing, and excluded from the verification chain.
- *That deliberately retained discomfort raises the felt value of a comfort* (the contrast yield, Section 3.2). The adaptation literature supports a moving reference level [Hel-

son, 1948, Brickman et al., 1978]; **it does not support the intervention claim**, which is uncited.

10.4 One of the framework's own rules is unfalsifiable and is marked. The inventory's ruling that a written "nothing" in a cost currency indicates a failure of the second look rather than an absence of cost cannot be wrong: any counter-example is absorbed as insufficient looking. It is retained here because it is a *procedural* instruction rather than an empirical claim, and it is relabelled as one. As an empirical claim — that every holding charges in every currency — it is not stated and should not be.

10.5 Parameters chosen rather than measured. The two time constants in Figure 3 ($\tau = 4$ and $\tau = 12$ weeks) were chosen; the source supplies none, and the pair is plotted precisely so that no single choice is mistaken for an estimate. The values $N \in \{3, 10, 40, 100\}$ in Figure 1 were chosen to span the framework's own worked range. Figures 1, 2 and 3 contain no measured quantity of any kind. Figure 4 is a schematic. Only Figure 5 reports features from a record, and that record is a single self-reported case.

10.6 Scale type is stipulated, not demonstrated, and the four-currency total does not add. Section 6.2 gives the full statement. The strongest objection to the framework is that psychology assumes quantitative structure without testing for it [Michell, 1997]; the test exists [Luce and Tukey, 1964, Krantz et al., 1971] and **was not applied**. Two of the four carrying-cost currencies have no operational unit, so the ledger's four lines are not summable (Section 9.4). **Falsified by:** a conjoint-measurement study on any operationalisation of carrying cost that fails double cancellation, which would remove the arithmetic of Section 3 entirely.

10.7 u has not declared which construct of difference it is. Three incompatible constructs are called diversity — separation, variety, disparity — and any measure must say which [Harrison and Klein, 2007]. On the corpus's own operationalisations u is a separation measure on a squared alignment, with the consequence that perfect opposition and perfect agreement are indistinguishable [Curzi, 2026r]. That is a real limitation on what R1 can be used to say about a group that disagrees.

10.8 The surplus is stated monotonically and the evidence is two-sided. Section 8.2 gives both sides. A monotone claim cannot survive Uzzi and Spiro [2005], Putnam [2007], May [1972] and Doak et al. [1998] unqualified.

10.9 Nothing here is identified. As an aggregate specification, any claim that a group's u produces its output is not identified from observational data [Manski, 1993], and the network route is closed by the generic confounding of influence with selection [Shalizi and Thomas, 2011]. This paper adopts no remedy and therefore states R1 and R3 as **modelling positions rather than identified empirical results** (Section 6.3). **Resolved by:** a randomised or partial-population design, or a defended stipulation of the reference group in the manner of Brock and Durlauf [2001].

10.10 The framework's own parameters are not exempt from its own thesis. By the Lucas critique, a u estimated in a population that has not read the framework cannot be assumed stable in one that has [Lucas, 1976]. No stability is claimed and no result here depends on any; the cost is that no cross-population comparison of u is currently interpretable.

10.11 The erasure test identifies a negative, not a positive. R5's argument establishes that a reproduced pattern is not stored in its recording medium. It does not establish that the generator is a comfort gradient rather than a schedule, an obligation, or a physical constraint. **This is the difference between ruling out one explanation and establishing another**, and the framework's presentation does not mark it.

10.12 One cited corpus figure could not be reproduced. Section 7.5 records that the equilibrium quoted for the headroom instrument's collective case at $u = 0.45$ does not follow from the parameters as quoted. The disagreement is printed rather than harmonised, and no claim in this paper depends on the figure.

10.13 The instrument corrupts what it measures. A quantitative indicator adopted for one's own decision-making distorts the process it monitors [Campbell, 1979, Goodhart, 1975]. The ledger and the effective count are both such indicators, and the exposure is structural.

10.14 No hypothesis was tested and no data were collected. H1–H5 stand. The discipline that would make a later confirmatory attempt credible is specified and not exercised [Nosek et al., 2018], and the analytic flexibility available in any such attempt should be constrained in advance. The corpus's own control measured anti-conservative at 0.108 against a nominal 0.05 [Curzi, 2026u], which is the prior any such test should be read against.

2.0.11 11. Conclusion

The seven results are stated below in the exact form in which the rest of the corpus cites them. Downstream papers should quote this text rather than paraphrase it.

R1. The surplus of N viewpoints over one is $f(u) = \ln(1+(N-1)u)$, uniquely determined by four requirements — dependence on the angles alone, zero surplus at parallelism, an effective count affine in u, and normalisation — with $N_{\text{eff}} = 1 + (N-1)u$ read as an effective number of viewpoints. **The uniqueness is conditional on the affine axiom, which is a design choice and not a theorem, and N_{eff} is not a Hill number of any order.** R1 is a modelling position, not an identified empirical result [Manski, 1993].

R2. The functional is a **lower bound on the ensemble's information, not the information itself**: it logs the floor and not the ceiling. It may be used to establish that a structure has at least a given surplus; it may not be substituted for a measurement of what the structure actually carries.

R3. Zero at parallelism is the binding clause. A second opinion drawn from a source that shares training, incentives and premises with the first is worth $\ln 1 = 0$ of surplus. Agreement between aligned instruments is echo and not evidence. The same result, with the mathematics attached, is that diversification benefit depends on covariance and not on count [Markowitz, 1952].

R4. Carrying cost is itemised in four currencies — coin, hours, attention, foreclosure — against the law that felt comfort decays toward zero over weeks while carrying cost does not decay at all, so that a depleted arrangement can assemble itself without any single bad decision. **Two of the four currencies have no operational unit and the four do not sum.**

R5. The worn path is the theory's observable: a desire line is the integral of many small descents, and its reproduction after the erasure of the recording medium is evidence that the pattern lies in the field of forces rather than in the record. **The erasure test rules out the medium; it does not identify the generator.**

R6. The variational form of a mechanical law is equivalent to the local differential law, so the principle licenses no teleology: the path obeys its ground instant by instant, and no measurement on a trajectory distinguishes a chosen path from an obeyed one.

R7. The honest limit of descent: a gradient finds the nearest low place and not the best one, so the analysis can diagnose a stable trap and cannot treat one. Only a change in what can be seen resolves it, and that claim is untested.

Counter-result. The surplus is bounded twice over — it is only a lower bound on the ensemble's information, and its logarithmic concavity means a few well-separated viewpoints capture most of what a large panel could, so the framework argues explicitly against accumulating perspectives. **The source's own statement of that second bound is corrected here:** three fully non-parallel viewpoints capture 29.8 per cent of what forty capture at $u = 1$, and exceed forty viewpoints only below $u^* = 2/39 = 0.051282$. The gradient material then concedes that descent can diagnose a stable trap but cannot treat one, and the ledger material concedes that the three pricing asymmetries it leans on have no supporting entry anywhere in this corpus's bibliography.

What is established, then, is one functional that follows exactly from four stated requirements, two limits on that functional that the framework prints as loudly as the functional itself, one observable with a valid negative authentication, and one theorem of mechanics this bibliography cannot cite. What is not established is any of the five hypotheses, the affine axiom that R1 depends on, or the decay asymmetry the ledger is built around. The paper's

contribution at this stage is that it names the three premises whose settlement would decide the matter — that the effective count is affine, that felt return decays while cost does not, and that a worn path is generated by a comfort differential — and that it declines to claim any of them.

2.0.12 12. Declarations

Data and code availability. No external dataset was collected for this paper and none exists. The worked case reported in Section 5.7 is drawn from the source manuscript's own record and is not a dataset. All figures are generated by the scripts named in their captions (`figs/src/CCP-A02_fig01.py` through `figs/src/CCP-A02_fig05.py`), which import the shared style module `lib/figstyle.py`; every number plotted is either computed within the script from a formula stated in this paper or quoted from the source text identified in the script's docstring. Three of the five scripts contain assertions that fail if the numbers they plot cease to agree with the source claims they are drawn from.

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the analysis, the figures, and the text.

Ethics. No human-subjects research was conducted for this paper. No personal data were collected or processed. The worked case in Section 5.7 is drawn from the source manuscript, is presented under the pseudonym it carries there, and is illustrative rather than evidential; it was not gathered under a research protocol and is not offered as one. Nothing in this paper is clinical advice.

Use of the ceremonial register. The source manuscript is in part an initiatory document. This paper analyses that register as an object of study and does not reproduce it; ceremonial names are reported only where they identify a construct this paper renames.

Chapter 3

Weighting Futures by Resistance: The Comfort-Weighted Amplitude, Cue Placement, and the Resolution-Disturbance Trade-Off

Abstract. This instalment sets out a scoring procedure for candidate futures, an account of environmental cue placement, and the trade-off law governing how finely a foreseen outcome should be specified. Branches traced forward are scored step by step for steepness on a four-point scale; the accumulated resistance R weights the branch as 1 / 2^R, so two branches differing by three points of honest steepness differ by a factor of eight. The pedigree is stated plainly as exponential, Boltzmann-style discounting of states by cost, and explicitly not as a phase-based path amplitude. A robustness test follows: break one middle step and halve the weight again if the branch cannot reroute, on the argument that plans are systematically imagined as their own best case. Futures adjacent to present arrangements are held to consolidate first, by default rather than by choice, and three classes of settled future are distinguished. The cue-placement node is fenced by a licence clause before any technique is given: arrangement of one's own path, or of the path of someone who has asked, is legitimate, and arrangement of another's path is refused. The framework holds that the refusal is self-enforcing, counter-movement rising with the detected push; that mechanism has no supporting entry in this corpus's bibliography and is reported uncited, and the ethical argument is shown to stand without it. Three placement laws follow, with friction priced in seconds. Finally the resolution-disturbance law: a finely fixed outcome buys precision by cluttering the path and inviting optimisation of the indicators rather than the outcome, while an infinitely coarse one exerts no gradient at all. The framework's single postulate is restated and badged, and the scoring rule's arithmetic is demoted to an ordering on scale-type grounds.

Keywords: resistance weighting; exponential discounting; planning fallacy; implementation intention; cue placement; reactance; resolution trade-off; default consolidation

Series: A. Foundations **Paper:** CCP-A03

3.0.1 1. Introduction

1.1 The question

A framework that treats a life as a landscape of costs owes three further things before it can be operated. It must say how a candidate future is *scored*, since a landscape without a scoring rule ranks nothing. It must say what may legitimately be done to a person's environment once the scoring is in hand, since a scoring rule that licenses arrangement licenses arrangement of other people too. And it must say how finely a foreseen outcome should be specified, since the act of specifying is itself an act performed on the path.

This paper supplies all three and prices each. The scoring rule assigns each step of a candidate branch an integer resistance from 0 to 3, sums it to a branch total R, and weights the branch as 2^{-R}. The environmental rule is fenced by a licence clause before any technique is given: arrangement of one's own path, or of the path of a person who has asked, is legitimate; arrangement of another's path is not, and the framework's claim is that it fails mechanically as well as ethically. The specification rule is a trade-off: a finely fixed outcome buys precision by cluttering the path and by inviting the optimisation of indicators rather than of outcomes, while an infinitely coarse one exerts no gradient at all.

Thirty corpus papers cite the results established here. Because those results are load-bearing elsewhere, this paper states them in a form intended to be quoted rather than paraphrased.

It also reports four things that go against the framework, at the same precision it reports its results. The scoring rule sums an ordinal judgement and then exponentiates it, which the corpus's own scale-type discipline forbids. The framework's exponential weight is one of a family of forms and nothing here discriminates it from the others. Five of the seven empirical mechanisms the framework leans on have no supporting entry anywhere in this corpus's bibliography. And the source's own worked example prints a weight ratio of "roughly a thousand" for two branches whose stated resistances differ by twenty points, where the framework's own rule gives $2^{20} = 1{,}048{,}576$ — a discrepancy of three orders of magnitude, printed here rather than corrected in silence.

1.2 What this paper contributes

Seven results are established or fixed here, labelled R1 through R7 and restated verbatim in the Conclusion.

R1 states the comfort-weighted amplitude: score each step of a branch 0 to 3 for steepness, sum to a resistance R, and weight the branch as 2^{-R}, so that a three-point difference in steepness is a factor of eight in expectation. **R2** states the pedigree honestly as exponential, Boltzmann-style weighting of states by cost, e^{-R} up to a change of base, and explicitly **not** as a phase-based path amplitude. **R3** states the neighbour test: break one middle step of a branch and halve its weight again if it cannot reroute, penalising plans that require every step to succeed. **R4** states the freezing law: futures adjacent to present arrangements consolidate first, by adjacency and default rather than by choice, with chosen, geographic and unsettled-underneath cases distinguished. **R5** states the licence clause that precedes all cue technique: arrangement is legitimate only on one's own path or on the path of a person who asked, because detected steering is claimed to produce counter-movement proportional to the discovered push. **R6** states three placement laws — site the cue on an already-walked path, represent the first step and never the summit, and re-site the cue once it has stopped being noticed — with friction priced in seconds of added or removed effort. **R7** states the resolution–disturbance law: fine specification buys precision at the cost of path disturbance, indicator-optimisation and the foreclosure of unlisted routes, while infinite coarseness exerts no gradient at all.

A counter-result is stated with equal weight. The uncertainty material declines the inflated reading of its own borrowed physics: the framework's postulate is stated as a postulate rather

than as a consequence, the physical relation is given only in its correct state-preparation form, and a named failure mode requires that a practitioner be able to distinguish the theorem, the ordinary observer effect and the framework's own conjecture in three separate sentences on demand.

1.3 Register, sources, and the chunk boundaries

The source manuscript is written in two registers, a narrative one and an analytic one, deliberately braided. This paper reports the analytic register. The narrative material is a fictional exposition and carries no analytical claim the treatise does not also carry; it is not reproduced. Where the source's ceremonial vocabulary names a construct — *the count of the undecided water*, *the veil that measures the measurer*, *the trapline* — the construct is renamed to what it is, and the ceremonial name is reported once, as an object of study.

The source volume was re-split after a de-framing correction, and the boundaries of the file assigned to this paper do not coincide with the source's own divisions. Two consequences are marked here rather than papered over. **At the opening**, the file begins with the failure catalogue and closing material of the preceding construct — the horizon procedure, whose arithmetic is stated in the preceding instalment of this series — and that catalogue is reported in Section 3, because it is substantive and because leaving it out would lose a failure mode the corpus elsewhere relies on. **At the close**, the file ends immediately after the uncertainty construct's ceremonial statement, so the procedure and failure catalogue belonging to that construct fall in the file assigned to the next paper in this series and are not reported here. Section 7 therefore states the trade-off law and its riders without its operating protocol, and says so at the point where the protocol would have appeared.

3.0.2 2. Hypothesis

The material below divides into claims that could be wrong, definitions that could not, and one postulate that the source itself labels as such. A hypothesis that nothing could falsify is restated as a definition or a postulate and labelled as such [Popper, 1959].

H1 — the resistance-weighting hypothesis. The relative likelihood that a person completes one candidate course of action rather than another is governed by the accumulated steepness of the individual steps, and falls by half for each point of accumulated steepness. *Predicted:* across a set of dated, scored branches, completion rates fall geometrically in R with a halving constant near one point, and do not track the attractiveness of the endpoint. *Falsified by:* completion tracking endpoint attractiveness at equal R, or falling in R at a rate materially different from one halving per point. *Status:* testable in principle; **not tested here.** The halving constant is a stipulation and no one has estimated it (Section 12.5).

H2 — the robustness hypothesis. A branch that cannot survive the failure of one middle step is materially less likely to complete than a branch of equal R that can reroute. *Predicted:* among branches matched on R, those that survive a simulated single-step failure complete more often. *Falsified by:* completion independent of single-step robustness at matched R. *Status:* testable; **not tested here.** The specific penalty — one further halving — is a stipulation, not an estimate.

H3 — the adjacency hypothesis (the freezing law). Futures requiring the smallest change from present arrangements consolidate first, by default rather than by choice, so that most of a person's medium-term future is settled before it is considered. *Predicted:* the fraction of a person's future commitments that were never explicitly chosen rises with the passage of time under stable circumstances. *Falsified by:* commitments arising at rates independent of their distance from present arrangements. *Status:* **not tested here**, and the mechanism the source names for it — default effects and habit consolidation — has no supporting entry in this corpus's bibliography (Section 12.3).

H4 — the placement hypothesis. A physical cue sited on an already-walked path, at reaching height, representing only the first step, raises the completion rate of the step it represents,

and the effect decays as the cue becomes familiar and is restored by re-siting it. *Predicted:* completion rises after placement, decays over weeks, and recovers on re-siting without any change to the cue itself. *Falsified by:* no decay, or no recovery on re-siting. *Status:* testable and cheap to test; **not tested here.** The framework's supporting literature — implementation intentions, cue-driven habit, friction measured in seconds — has **no entry** in this corpus's bibliography, and the claim is therefore reported uncited (Section 12.3).

H5 — the resolution–disturbance hypothesis (the paper's sharpest claim, and its weakest support). For a future coupled to its own foreseeing, the fineness with which the outcome is specified and the disturbance imposed on the path cannot both be made small: their product is bounded below. *Predicted:* across matched undertakings, those specified more finely show more measured effort spent on checking and on indicator maintenance, and a lower rate of arrival by routes not listed in advance. *Falsified by:* finely specified undertakings showing no more checking and no fewer unlisted routes than coarsely specified ones. *Status:* **not tested, and not currently testable as stated**, because neither the resolution of a specification nor the disturbance of a path has been operationalised, and because the source states plainly that the bound's constant has never been measured by anyone.

D1 — resistance (definition). The resistance of a step is an integer in $\{0, 1, 2, 3\}$: 0 for a step that would be taken while unwell, 1 for one that ordinary resolve suffices for, 2 for one that requires a good day, 3 for one that requires a better self than usually answers. The resistance of a branch is the sum over its steps. This is a stipulated rating scale. Its measurement-theoretic status is audited in Section 8.2 and it does not survive that audit intact.

D2 — the licence clause (definition, and a constraint). Cue placement is performed only on the practitioner's own path, or on the path of a person who has asked. This is a stipulated constraint on the method's domain, not an empirical claim, and it is stated before any technique in the source and before any technique here.

P1 — the Postulate. The uncertainty of a future that is affected by its own foreseeing is postulated to be the same *kind* of uncertainty as that of a quantum state: not merely an ignorance in the observer but a participation of the observing in the observed, with the trade-off structure carrying over in *shape*. Nothing in physics licenses this. The source states it as a postulate and this paper preserves the badge at every occurrence.

3.0.3 3. Material carried across the opening boundary

The file assigned to this paper opens with the failure catalogue of the preceding construct. It is reported here because it is substantive, because one of its five entries is the sharpest statement in the whole volume of a result the corpus relies on elsewhere, and because a procedure supplied without its failure catalogue is an advertisement. The source's labels are renamed with a `CC-` prefix to keep them clear of this paper's own results.

CC-W1. The smeared pane. *Named mechanism: anchoring and scenario contamination.* The analyst switches assumptions mid-trace, importing a rescue from one branch into another, and produces a composite assembled from parts that cannot coexist, each purchased under an assumption the others deny. The source's warning is that the product reads as *richer* than an honest trace, which is precisely what makes it dangerous.

CC-W2. One assumption repeated. *Named mechanism: pseudo-diversity; the illusion of independent confirmation.* Three assumptions are cut from one premise — three costumes on one position — and their agreement is read as convergence. **The source prices this exactly, and the pricing is the reason this entry is reproduced:** non-parallelism $u = 0$, effective count $N_{\text{eff}} = 1$, surplus $\ln 1 = 0$. Three parallel viewpoints are one viewpoint. The collision test — the requirement that each new assumption contradict an existing one somewhere in its consequences — exists for this fault alone.

CC-W3. Vividness mistaken for probability. One branch is rendered in detail and the others in outline, and the branches are then weighed by brightness of rendering rather than by honesty of steps. The remedy the source prescribes is procedural evenhandedness: the same number of steps and the same number of sensory particulars for every branch.

CC-W4. The unbuilt branch. The refused branch is never traced at all, and the audit proceeds over a set curated to be survivable. The cost is not a missing future but a corrupted convergence: an invariant certified across only the bearable branches is an invariant of the analyst's tolerance rather than of their situation.

CC-W5. The straight line drawn on a curved surface. *Named mechanism: linear extrapolation past the regime in which the trend holds.* The present's visible motion is extended in a straight line and the drawing is called a trace. No step consulted the step before it.

Two of these five have an external literature this corpus can cite and three do not. CC-W2's arithmetic is the corpus's own and is fixed in the preceding instalment of this series. CC-W5 is the standard warning against fitting a trend outside the regime that produced it, and the general form of the warning — that a curve's shape is weak evidence for the mechanism that produced it — is established [Stumpf and Porter, 2012]. CC-W1, CC-W3 and CC-W4 name mechanisms (anchoring, availability, motivated reasoning) for which **this corpus's external bibliography contains no entry**, and they are reported uncited.

3.0.4 4. The comfort-weighted amplitude

4.1 What is borrowed, and what is not

The source opens this material with an unusually careful statement of what it is borrowing, and the care is worth preserving because the borrowing is where a framework of this kind usually fails.

The physical account it quotes is the quantum-mechanical one: the state of a system is a wavefunction assigning to each configuration an *amplitude*, a complex number with a size and a phase; the probability of finding the system in a configuration is the square of the amplitude's size; amplitudes can cancel, so two nonzero routes to the same outcome can sum to zero, which is interference and is measured routinely; between measurements the state evolves smoothly and deterministically; and what happens at measurement remains an open question on which the mathematics does not adjudicate. The source states the last point explicitly and takes no side, which is the correct posture and is preserved here. **This corpus's external bibliography contains no entry for the Born rule and none for the postulates of quantum mechanics as such**; the one entry bearing on the measurement question is the decoherence programme's founding statement [Zeh, 1970], and the remainder of the physical account in this subsection is therefore reported as the source states it and is **uncited** (Section 12.3).

What the framework then borrows is stated by the source as three things and no more, and this paper endorses the enumeration:

1. **Weighted coexistence until resolution.** A future, before commitment, is handled as many coexisting branches each carrying a weight, rather than as one hidden fact. This is not a physical claim; it is how any forecaster holds a future.
2. **Weights attach to paths, not endpoints.** What a route contributes is set by the whole of the route and not by the attractiveness of its destination.
3. **The outcome that occurs is the one its own small variations agree with.** In the physical account the classical path survives because neighbouring routes have nearly the same action and reinforce, while elsewhere they cancel. The framework's transfer is that the future that arrives is the one that still arrives when the plan goes slightly wrong.

The source explicitly refuses the fourth borrowing that would be available and attractive — that the future is *literally* a superposition — and states the refusal in its own text. That refusal is R2's content and is treated in Section 4.3.

4.2 The scoring rule and the weight (R1)

The procedure is stated operationally because later papers refer to its outputs.

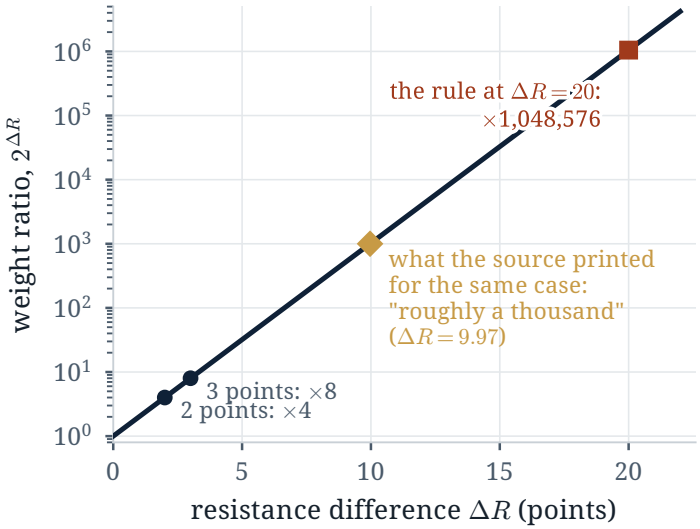


Figure 3.1: **The rule, and a printed error in it.** Arithmetic on the source's own statement that a branch's weight "falls by half for every point of resistance it carries above the easiest branch's total". The two annotated small cases are the source's own. The two markers at the right are the same worked case read two ways: the source records branch resistances of 29 and 9, so its own rule gives a weight ratio of $2^{20} = 1{,}048{,}576$, while the text prints the ratio as "roughly a thousand", which corresponds to $\Delta R = 9.97$. Script: `figs/src/CCP-A03_fig01.py`.

Take the candidate branches. Walk each forward in steps — twelve is the protocol's number, months or moves as the branch divides. At each step ask one question and only one: *how steep is this single step for the person who actually shows up on a tired morning?* Score it on the four-point scale of D1. Sum the scores along the branch: that is R, the branch's resistance. Take the smallest R among the branches as the reference and subtract it from each, so that the easiest branch sits at $\Delta R = 0$. Then

$$w(\Delta R) \;=\; 2^{-\Delta R}$$

so that a branch three points steeper than the easiest carries one eighth the weight. This is R1. The source's own instances are two: a three-point difference is a factor of eight, and the two-point difference between an exercise machine at home and one across town is a factor of four — which the source offers as the mechanism behind the annual emptying of gymnasiums.

Two features of the rule are worth stating plainly because they are what make it a rule rather than a sentiment. It is **path-summed**: the endpoint contributes nothing, and the protocol requires the endpoint to be physically covered during the scoring. And it is **anchored at the minimum**, so it produces relative weights only. It cannot be read as a probability of any branch in isolation, and Section 8.2 states why it cannot be read as a probability at all.

4.3 The pedigree, stated honestly (R2)

R2 is a disclaimer that the source states as doctrine, and it is the most creditable sentence in this material.

The weight 2^{-R} is **exponential discounting of states by their cost**. Its nearest relative is the Boltzmann–Gibbs weighting of a state by $e^{-\beta E}$, which is exactly the distribution that maximises entropy subject to a mean-cost constraint [Jaynes, 1957, 2003]. Written that way, $2^{-R} = e^{-R \ln 2}$, so the framework's halving-per-point is the exponential form with $\beta = \ln 2$ in units of resistance points. **It is not Feynman's phase.** A path amplitude carries a phase and

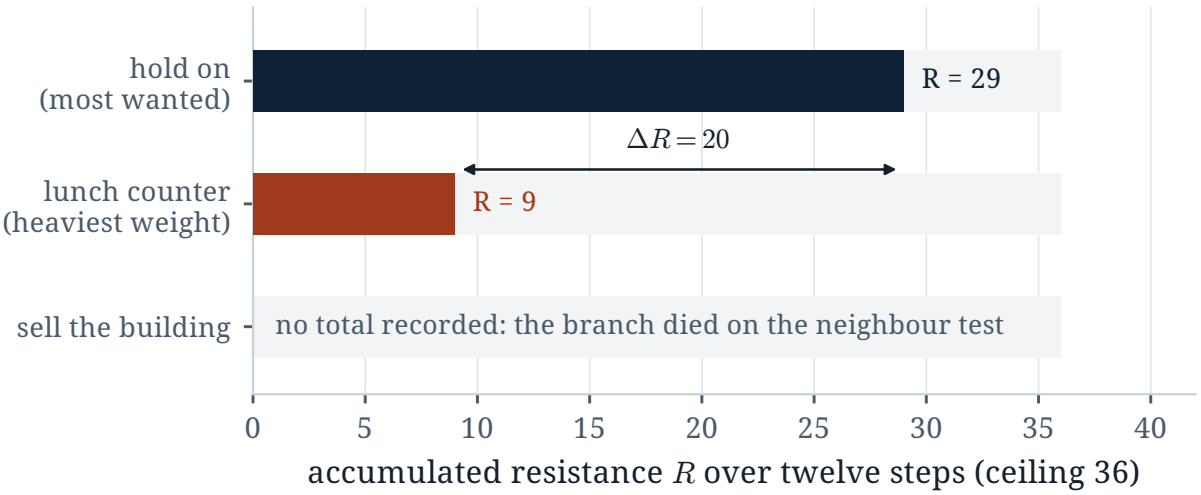


Figure 3.2: **The worked case's resistance budgets, against the ceiling its own scale fixes.** Measured from the source record, which prints "R totals 29" for the continuation branch and "R totals 9" for the intervention branch and gives no total for the third, which the record states died at the neighbour test. The ceiling of 36 is twelve steps at the scale's maximum of 3. Script: `figs/src/CCP-A03_fig02.py`.

can cancel; an exponential weight is positive and cannot. The source says this in its own text — "we use the physicist's insight about paths and neighbours; we use the statistician's actual numbers" — and this paper preserves the distinction at every occurrence.

The distinction is not pedantic, and its cost is exact. **Every claim the framework makes about cancellation is therefore unsupported by its own arithmetic.** Section 4.5 reports the framework's treatment of ambivalence, which it describes in the language of two routes arriving out of phase and unmaking each other; under 2^{-R} nothing cancels, because all the weights are positive. The framework's cancellation talk is a picture and not a computation, and this paper marks it as one wherever it appears.

One further honesty belongs with R2. The exponential is **one member of a family**. A one-parameter generalisation of the entropy that generates it produces a family of weightings of which the exponential is the special case [Tsallis, 1988], and nothing in the source or in this paper discriminates the exponential from the alternatives on any evidence. The halving constant, in particular, is a round number chosen for arithmetic convenience.

4.4 The neighbour test (R3)

The robustness step is the framework's transfer of the stationary-phase idea, and it is the part of the borrowing that does real work.

For each branch, choose one step in its middle third and declare it failed. Ask whether the branch reroutes or dies. If it dies, halve its weight again and mark it *fragile*. The rationale the source gives is that a branch requiring every step to go right is a branch with no agreement among its neighbours, and such branches do not survive contact.

The framework attaches an external warrant — that plans are systematically imagined as their own best case, so robustness rather than optimality predicts completion. **This corpus's external bibliography contains no entry for the planning fallacy**, and the claim is retained **uncited** (Section 12.3). What can be supported is narrower and is stated instead: that the predictability of individual life outcomes is low even under ideal data conditions, so any single planned trajectory should be held loosely [Hofman et al., 2017, Salganik et al., 2020].

R3's own penalty — one further halving — is a stipulation. Nothing establishes that fragility costs a factor of two rather than a factor of ten or nothing at all, and Section 12.5

says so.

4.5 Cancellation and the ambivalent branch

The framework's treatment of ambivalence deserves reporting because it contains a genuine operational instruction and an unsupported mechanism, and the two can be separated cleanly.

The instruction: where a branch's computed weight is high and the practitioner is nonetheless not moving toward it, the framework directs that both accounts of that future be written on the same card, side by side and named — the account in which arriving is a homecoming and the account in which arriving is a defeat — so that the conflict becomes visible arithmetic instead of invisible weather.

The mechanism: that the two accounts *cancel*, arriving at the same outcome out of phase. **Under the framework's own weighting they cannot**, for the reason given in Section 4.3: 2^{-R} is positive and superposes additively. The source also names an approach–avoidance mechanism from the clinical literature, for which **this corpus's bibliography contains no entry**. Both are reported and neither is used; the instruction survives their removal, since writing two accounts of one future on one card is defensible as a bookkeeping rule regardless of what mechanism, if any, underlies the conflict.

4.6 Reading without preferring: three procedural rules

The framework attaches three rules to the scoring, and all three are recognisable as standard research hygiene rather than as novel doctrine, which is a point in their favour and is stated as such.

The outside view. Score a branch against the recorded fates of everyone else who has attempted it before adjusting for particulars. The framework's phrasing — that the inside of a plan is a showroom, lit and priced to move — is vivid and its content is the base-rate discipline. The relevant external result is that individual life outcomes are barely predictable even from fifteen years of high-quality longitudinal data by a hundred and sixty teams, so an individual's confident inside view carries very little information [Salganik et al., 2020, Hofman et al., 2017].

The standing scorer. Write the weights down, date them, and check afterwards whether the stated proportions occurred at the stated rates. The framework asserts that people who keep score improve and people who do not, do not. **No entry in this bibliography establishes that**, and it is reported uncited. What is established, and what the rule amounts to, is that a prediction registered in advance constrains what can later be claimed to have been predicted [Nosek et al., 2018], that the space of analytic and interpretive choices is otherwise wide enough to support almost any conclusion [Simmons et al., 2011, Gelman and Loken, 2013], and that a directional prediction is nearly costless whereas a point prediction with a constant in it is not [Meehl, 1967, 1990].

The moved observer. Re-score the branches from more than one standpoint — tired and rested, in the character of a harsh critic and of a friendly one — and mark the steps on which the readings disagree by two or more as contested rather than resolving them. The framework's argument is that the branch which appears attractive does so partly because of where the observer is standing, and that a branch attractive from every standpoint the observer can occupy has something in it beyond the observer's geometry. This is a sensitivity analysis given a memorable name, and the framework says so.

4.7 The failure catalogue of the weighing

Five failure modes are named, renamed here with the `CC–` prefix.

No.	Failure	Mechanism named by the source	Support in this bibliography
CC-A1	The flattering wind	Optimism bias; scoring the wanted branch as the intended self	**none**
CC-A2	The drowned pilot	Catastrophising; one vivid disaster taken as the base rate	**none**
CC-A3	Walking on young ice	Premature closure; consolidation by adjacency mistaken for choice	**none**
CC-A4	The single sample	Extrapolation from $N = 1$; re-weighting on one observation	Salganik et al. [2020] bears on it; the source's own remedy is a minimum of three independent observations
CC-A5	The connoisseur of the undecided	Analysis paralysis; deferral dressed as diligence	**none**; the source's own mark for it is "a beautiful ledger and an unchanged life"

Four of the five carry no citable support in this corpus. They are retained because a failure catalogue is the most useful part of a procedure, and they are retained uncited.

3.0.5 5. The freezing law (R4)

The framework's account of how a future settles is stated as a physical analogy and then, correctly, as a claim about defaults.

Shallow water freezes before deep water: it holds less stored heat and sits out of the stirring. The framework's transfer is that **futures adjacent to present arrangements consolidate first, and consolidate whether or not they are chosen**, by adjacency, by default, and by the thinness of the change they require. Its practical consequence is the one worth quoting: by the time most people ask what their life will be, much of the answer has already settled along the margin, and they are standing on it as they ask.

The transfer is a **shared mathematical form at best and an analogy at worst**, and it is labelled in Section 8.1 as an analogy, because nothing in the thermodynamics of a lake constrains anything about a calendar and no premise of the physical result is shown to be satisfied. What survives the removal of the analogy is a claim about defaults, and it is the claim the framework actually uses.

The operational content is a three-way classification of settled futures, and this is R4's citable form:

- **Frozen by choice** — consolidated in some earlier period by deliberate arrangement. The framework treats these as the practitioner's built assets.
- **Frozen by geography** — consolidated by adjacency and default, never chosen, frequently useful, and to be walked on knowingly rather than gratefully.
- **Reed-marked** — settled in appearance, with something unresolved still running underneath. The framework's instruction is to mark these rather than to test them with weight.

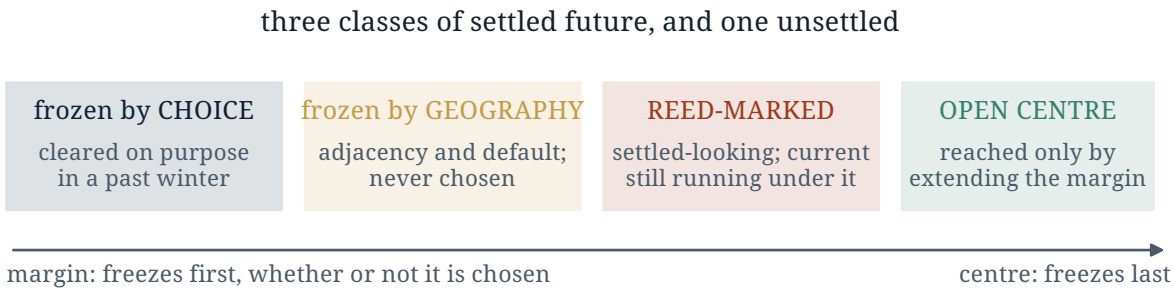


Figure 3.3: **Three classes of settled future, and one unsettled.** Schematic; no plotted quantity. The classification and the margin-to-centre ordering are the source's own: futures adjacent to present arrangements consolidate first, whether or not they are chosen, and the open region is reached by extending the margin rather than by preferring the centre. Script: `figs/src/CCP-A03_fig04.py`.

And one unsettled class: the open region far from every present arrangement, which the framework holds is reached not by preferring it but by extending the margin toward it one placed step at a time — which is what Section 6 is about.

The mechanism the framework names for the freezing — default effects and the consolidation of habit — has no supporting entry in this corpus's bibliography and is reported uncited (Section 12.3). H3 is therefore stated as a hypothesis with a named falsifier and no supporting literature, which is an honest position and a weak one.

3.0.6 6. Cue placement

6.1 The licence clause (R5)

The source states its constraint before its technique, and this paper preserves the order, because the order is part of the claim.

> **Arrangement of one's own path, or of the path of a person who has asked, is the method. Arrangement of another person's path is not.**

The framework gives two reasons and they have very different standing, so they are separated here.

The first is ethical and needs no literature: a technique that works below the awareness of its target, in a channel where the target cannot defend themselves without becoming someone who distrusts every encouragement, is coercion whatever its intent. The framework states this as law rather than as counsel, and the corpus states it as law elsewhere.

The second is mechanical: that a person who discovers they were steered recoils from the direction of the steering with a force proportional to the discovered push. **This corpus's external bibliography contains no entry for reactance.** The claim is retained because it is load-bearing for the framework's argument that the constraint is self-enforcing, and it is retained **uncited**, exactly as the other papers of this series retain it. What can be supported is narrower and is stated instead: that expectations held about a person alter what that person produces, though the effect is smaller and less accumulative than the popular account holds [Rosenthal and Jacobson, 1968, Rosenthal and Fode, 1963, Jussim and Harber, 2005]; and that a theory adopted as a description of people can make itself true by changing how they are treated [Ferraro et al., 2005].

R5 is therefore **a stipulated constraint with an ethical argument and an uncited mechanical one.** The stipulation stands without the mechanism; the framework's stronger claim, that the prohibition enforces itself, does not.

6.2 Three placement laws (R6)

The technique itself is three laws, and each is stated with the friction it prices.

First: site the cue on a path already walked. A cue is an interception of an existing route rather than an invitation to a new one. The framework's named failure is placement where the practitioner *ought* to pass rather than where they do, and its prescribed remedy is one day of self-observation — an application of the worn-path observable fixed in the preceding instalment of this series — followed by placement at the strongest observed station, at reaching height and in the line of the resting eye.

Second: represent the first step and never the summit. The cue must be exactly the size of the next action. The framework's argument is the gradient argument: a system moves when the next step is downhill and is incapable of taking any step but the next one, so a cue depicting the whole undertaking adds the whole undertaking's weight to the next step. Its named failure is a cue that produces a flinch, which it holds trains avoidance with perfect efficiency.

Third: re-site the cue when it stops being noticed. A cue wins attention by contrast and singularity and then loses it as the eye files the familiar as background. The framework's ruling is that this is a perceptual event rather than a moral one, and that the response is to move the cue rather than to reproach oneself. It fixes a standing inspection interval for exactly this reason.

A fourth law is stated as the third's shadow: friction added to a path is as much a part of the method as friction removed from one, with the framework's preference for removal on the ground that a life engineered mainly by self-obstruction becomes adversarial to itself.

The framework prices friction in seconds — some twenty seconds of added or removed effort, it claims, decides behaviour at rates no stated resolution matches. **This corpus's bibliography contains no entry for that finding, none for implementation intentions, and none for cue-driven habit**, which together are the entire external warrant the framework claims for this section. All three are reported uncited and listed in Section 12.3. **The scaffold for this paper names an organisational-dynamics cluster as the source of that support; the cluster does not in fact contain it, and this paper records the gap rather than attaching the nearest plausible key.**

What survives without the missing literature is the framework's internal argument, which is R6's defensible core: *if* behaviour follows local gradients, then a cue that lowers the resistance of a step by one point is worth a doubling of the branch's weight under R1, and that is an arithmetic consequence of R1 rather than an empirical claim.

6.3 The named day and the neutral instrument

Two devices are given at more length, and both are stated by the source in a form that distinguishes them from superstition; the distinction is the interesting part and is preserved.

The named day. One day of the week is designated for a standing question and a written page. The framework's account of why it works has three parts and one addition, and it is careful to claim nothing beyond them. First, the calendar becomes a self-synchronising trigger, so the appointment re-arrives without effort. Second, it creates a place to file relevant material that would otherwise be discarded, so observations accumulate rather than dying on the day they occur. Third — and this is the part the framework insists on stating plainly — it deliberately harnesses the frequency illusion: an attention loaded with a standing question finds material relevant to it that was always present and unfound. The framework's own statement is that the events were there and the gathering is new. Its stated boundary is that a practitioner must be able to say, in plain words, why an ordinary day works; and that the inability to say so is the border between the method and augury.

That boundary is the strongest single sentence in this material and it is preserved verbatim in substance. Its cost is that **the frequency illusion, the fresh-start effect of temporal landmarks, and the claim that scheduled attention accumulates relevance all have no entry**

in this corpus's bibliography and are reported uncited. The one part of the account that can be supported is the resource argument: attention is a bounded channel and a scheduled appointment is a way of allocating it [Miller, 1956, Ashby, 1956].

The neutral instrument. A deliberately meaning-free object is placed on a daily path so that whatever is projected onto it is projected somewhere the practitioner can read it. The framework's claim is not that projection can be started or stopped but that it can be *given a designated instrument*, and that an instrument taking no side returns the practitioner's own state with unusual purity. It is careful to say what such an instrument reads: not the world, but the reader. **No entry in this bibliography supports the projective-instrument claim**, and it is reported uncited. What *is* supported, and what applies to the device with some force, is that any quantity a person adopts as an indicator of their own progress and then acts on ceases to be a good indicator of it [Campbell, 1979, Goodhart, 1975].

6.4 What a cue buys, in the weighting

The join between Sections 4 and 6 is arithmetic and is worth stating exactly, because it is the only quantitative claim in this material that does not require any missing literature.

Under R1 a branch's weight is 2^{-R}. A cue that re-prices one step from "requires a good day" to "already in motion" removes two points of resistance from that step and therefore multiplies the whole branch's weight by $2^2 = 4$. One point removed doubles it; two well-placed cues that each remove one point quadruple it. The source states exactly this, and it is correct **as arithmetic on R1**. It is not evidence that cues remove resistance points, which is H4 and is untested.

The ordering law follows and is stated as a rule: **weigh first, then place.** Placing cues on an unweighted branch amplifies whatever is nearest rather than whatever is best, with all the efficiency the section claims; weighing without placing produces the failure mode CC-A5, a complete ledger and an unchanged situation.

6.5 The failure catalogue of placement

No.	Failure	Mechanism named by the source	Support in this bibliography
CC-L1	The cue on the unwalked path	Cue absent at the point of decision	**none**; the source's mark is a cue the practitioner cannot recall seeing
CC-L2	The summit painted on the jar	Oversized cue conditioning avoidance	**none**
CC-L3	The cue gone to furniture	Habituation; inattention to the familiar	**none**; remedy is re-siting, not reproach
CC-L4	Placement on another's path	Covert influence and the recoil it purchases	**none** (reactance; Section 12.3). The ethical argument of Section 6.1 stands without it

No.	Failure	Mechanism named by the source	Support in this bibliography
CC-L5	Everything named	Apophenia; the frequency illusion unharnessed until the record fills with signs and stops containing findings	**none**; the source's remedy is scarcity — one named day, one instrument

CC-L5 is worth marking as the framework auditing its own device: the named day works, on the framework's own account, because it is *scarce*, and the mode of failure is the practitioner's success at the technique.

3.0.7 7. The resolution–disturbance trade-off

7.1 The uncertainty relation, stated correctly

The source states the physical result at length and states it correctly, and the correctness is load-bearing for what follows, so it is preserved exactly.

For a quantum state, the product of the standard deviations of two conjugate observables is bounded below:

$$\sigma_x \, \sigma_p \;\geq\; \frac{\hbar}{2}$$

This is the Kennard–Robertson form [Kennard, 1927], and the point the popular accounts get wrong is got right by the source. **It is not, at its core, a statement that a clumsy measurement disturbs a delicate system.** It is a statement about which states can exist: position and momentum are Fourier conjugates, and no state — measured or unmeasured — can be sharply defined in both, for the same reason that no signal can be both an instant and a pure frequency. The source makes the acoustic version explicit — a drum tap is sharp in time and therefore has no pitch; a bowed note has a pitch and therefore occupies duration — and that is exactly the correct pedagogy, because the acoustic trade-off is the same theorem in a domain where nobody is tempted to invoke an observer.

The corpus proves the point in that domain rather than asserting it. The time–frequency bound $\Delta t \, \Delta f \;\geq\; 1/(4\pi) \;=\; 0.079577471545948$ follows from the same variational argument as $\sigma_x \sigma_p \geq \hbar/2$ with the substitution $p = \hbar k$ deleted; Gaussian windows saturate it and a Hann window costs 2.6 per cent [Curzi, 2026z]. The value is reproduced exactly in Section 11.3. The underlying signal-theoretic result — that a signal cannot be arbitrarily localised in time and in frequency at once — is Gabor's [Gabor, 1946].

The source also records, correctly, that Heisenberg first explained the principle with a disturbance parable and that the parable was retired from the theorem's proof within the year. **That historical statement has no supporting entry in this bibliography** and is reported as the source's.

7.2 The observer effect, kept in a separate drawer

The source's second move is to insist that the uncertainty relation and the observer effect are different objects, and to keep them in separate drawers throughout. This paper preserves the separation and regards it as the material's chief methodological virtue.

- **The uncertainty relation is about what can exist.** No state is sharp in both members of a conjugate pair, watched or unwatched. It would hold in a universe with no observers in it.

- **The observer effect is about what asking costs.** To measure a system one must interact with it, and the interaction enters the system's history. A pressure gauge bleeds pressure; a thermometer drinks heat. This is ordinary and classical and long predates quantum theory.

The source names the welding of these two into a single slogan — "observation changes reality" — as the specific failure it is guarding against, and requires that a practitioner be able to state the theorem, the observer effect and the framework's own conjecture in three separate sentences on demand. That requirement is the counter-result of this paper and is restated in the Conclusion.

7.3 The mirror catalogue, audited

The source then leaves physics and lists five phenomena of observer participation in human systems, calling all five "real, replicated, and mundane". They are reported here with an audit of what this corpus can actually cite for each, because the source's blanket description is not accurate for all five.

1. **Measurement reactivity** — a counted behaviour changes because it is counted. **No entry.**
2. **The mere-measurement effect** — asking a person whether they intend to do something shifts whether they do it. **No entry.**
3. **The warm gaze** — observed people perform differently for their observers. The source itself attaches a caution to the founding factory study, and the caution is correct: the systematic review of that literature finds the effect poorly characterised and frequently overstated [McCambridge et al., 2014]. **Cited, and the citation bounds rather than supports the source's use.**
4. **The expectant eye** — an observer who expects a result tends to find it, by small steerings below awareness. Supported [Rosenthal and Fode, 1963, Rosenthal and Jacobson, 1968], and bounded by the finding that such effects are real but small, do not typically accumulate, and often dissipate, with most expectations turning out to be accurate rather than causal [Jussim and Harber, 2005]. **Cited with its bound.**
5. **The measure that becomes a target** — an indicator adopted for consequential decisions is gamed and the process it tracked is corrupted. Fully supported [Campbell, 1979, Goodhart, 1975].

Two of the five cannot be cited from this corpus's bibliography at all, and two more are cited by entries that bound the claim rather than establish it. The source's characterisation of the catalogue as uniformly established is therefore not sustained, and this paper says so. What survives is the fifth item, which is fully supported, and the fourth, which is supported in a reduced form — and the fifth is the one the framework actually needs, since the trade-off law of Section 7.5 turns on indicator corruption rather than on reactivity.

7.4 The Postulate, stated and badged

The source states its central conjecture in full and labels it a postulate in the sentence that states it. This paper preserves the badge.

> **The Postulate.** The uncertainty of a foreseen future — a future that is affected by its own foreseeing — is proposed to be the same *kind* of uncertainty as the physicist's: not merely an ignorance in the observer but a participation of the observing in the observed; and the trade-off mathematics that governs the one is proposed to govern the other in *shape*.

The source's own accounting of what is established and what is postulated is structured as three walls and a fourth, and the accounting is accurate on its own terms and is reproduced:

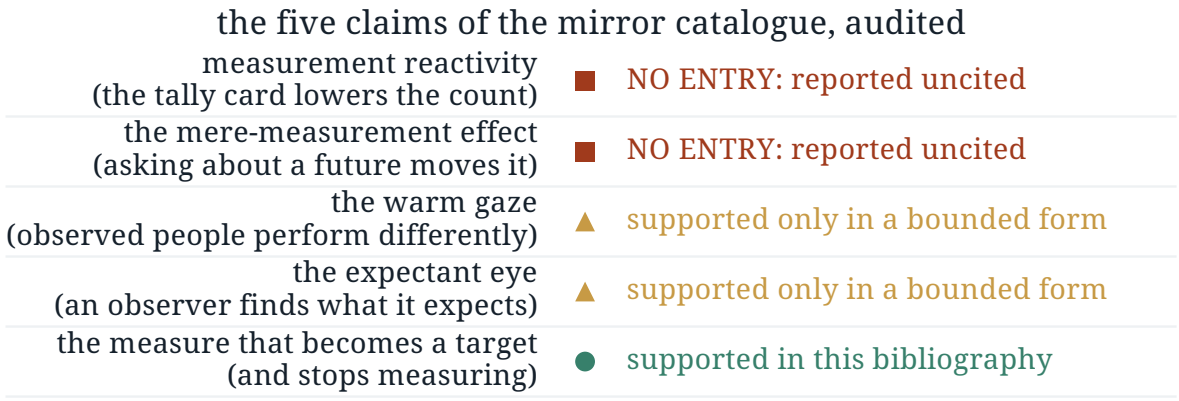


Figure 3.4: **The five claims of the mirror catalogue, audited against this corpus's own bibliography.** The source calls all five "real, replicated, and mundane". One is fully supported here, two are cited by entries that bound rather than establish them, and two have no entry at all. An absent entry means this corpus cannot cite the claim, not that the claim is false. Script: `figs/src/CCP-A03_fig05.py`.

- *Established:* a future is a distribution rather than a point.
- *Established, in the reduced form Section 7.3 permits:* the act of foreseeing participates in the foreseen.
- *Established:* there is no view from outside one's own life, so every forecast a person makes is made by an inhabitant of the system forecast.
- *Postulated:* that the trade-off structure carries over — that resolution is conjugate to disturbance.

The status is stated exactly. **It is not a theorem. It is not established physics. It is not a claim that a mind is a quantum system**, and the framework makes none of those claims. Under a strict falsifiability criterion it is not a scientific claim at all [Popper, 1959]; the honest description of its role is that of a hard-core commitment protected rather than tested, whose warrant is the productivity of the belt of testable claims around it [Lakatos, 1970].

One boundary the source cuts deliberately is worth preserving because it forecloses the obvious abuse: **the Postulate concerns futures coupled to their own foreseeing** — futures of systems that contain the person doing the foreseeing. Where the coupling is zero, the framework asserts nothing. It says nothing about tides, dice, or next year's commodity prices except to the exact degree that a person's seeing can reach the machinery that makes them.

The hazard the badge exists to prevent is documented: a physics result cited outside the conditions of its own derivation transfers its authority and none of its content, and the clearest case is the traversable-wormhole literature, routinely invoked to support shortcuts between distant states when its own conclusion is that such a shortcut requires exotic matter violating an energy condition [Morris and Thorne, 1988]. The corpus applies the same discipline elsewhere and supplies the number that makes it concrete: a warm, wet system at 310 K has a decoherence time of $\tau_{\text{dec}} = 1.98 \times 10^{-17}$ s against a millisecond-scale neural event, a shortfall of $5.04 ×10^{13}$, and a protection factor of 10^{6} still leaves $5.04 ×10^{7}$ [Curzi, 2026aa]. **That is the corpus refuting the literal reading of its own borrowed vocabulary**, and it is the reason the Postulate is stated as a correspondence of shape and never as a claim about substrate. The measurement question the source declines to adjudicate remains open on its own terms [Zeh, 1970].

The corpus's template for a badged postulate is explicit and this one follows it: a magnitude form with a forbidding clause attached — $E = G \cdot B + \varepsilon$ with $G = 0 \Rightarrow E = 0$ for all B — badged

as a postulate and paired by explicit correspondence with the storage law's zero-efficiency case [Curzi, 2026ab]. A strong claim, stated with the clause that forbids its abuse, and badged.

7.5 The trade-off law, and three riders (R7)

From the Postulate the source derives its working law:

> **Every act of foresight has a resolution — a fineness at which the foreseen outcome is specified. Resolution is purchased with disturbance: the finer the outcome is fixed, the more the fixing bends, clutters and constrains the path that must carry the practitioner there; the coarser the outcome is held, the smoother the path runs, at the price of spread in where it arrives. Neither pole is an error. The choice of resolution is the practitioner's most consequential act.**

The two poles are stated with their specific costs, and the costs are the law's content.

At the **fine** pole, three things happen and only the third is surprising. Every hour must be audited against a detailed template, and auditing is not walking. The measure becomes the target, so the practitioner optimises the *signs* of approach rather than the approach — which is the one item of the mirror catalogue this corpus can fully cite [Campbell, 1979, Goodhart, 1975]. And fineness **forecloses**: a specification refuses in advance every route that arrives at the right result by an unlisted road, and the number of such routes is not enumerable in advance.

At the **coarse** pole the path is undisturbed because nothing is asked of it, and nothing arrives, because a gradient requires a difference and "something better" differs from nothing in particular by nothing in particular.

The source writes the trade in the physicist's notation as a mnemonic — $\Delta(\text{outcome}) \cdot \Delta(\text{path}) \geq k$ — and then, in the same breath, states that it has not measured k and that nobody has. **This paper preserves both halves of that sentence and treats the second as the more important.** The inequality is a postulate of craft wearing borrowed notation. What is established is the *direction* of the trade at the fine pole, by indicator corruption; what is postulated is that the trade has a floor that cannot be beaten, only budgeted.

Three riders are attached, and the third is the only one with an operational consequence that does not depend on the Postulate.

First, resolution may sharpen as the outcome nears. A specification that would foreclose the path at two years' range is an ordinary checklist at two months. The framework's named error is demanding a near-term resolution at long range.

Second, the tested path thickens. Disturbance spent *on the path* — small trials, cheap experiments, weight actually placed — is an investment, because the route tested becomes the route best able to carry the practitioner; disturbance spent on the *outcome* is a tax. This is the most useful sentence in the section and it does not require the Postulate: it is a statement about where to spend a fixed budget of checking.

Third, the budget is real, so schedule it. The framework's protection against indicator corruption is not abstinence from measurement but an *appointed* measurement: a fixed, small, scheduled act of checking, outside which no checking occurs. **The protocol implementing that appointment falls beyond this paper's source boundary** (Section 1.3) and is not reported here.

3.0.8 8. Method: what kind of borrowing, what scale, and what is identified

8.1 Analogy discipline

Every physics borrowing in this paper is labelled at first use as a **derivation** (the system satisfies the physical result's premises), a **shared mathematical form** (the same equation governs both, with no claim that the physics applies), or an **analogy** (a suggestive correspondence with

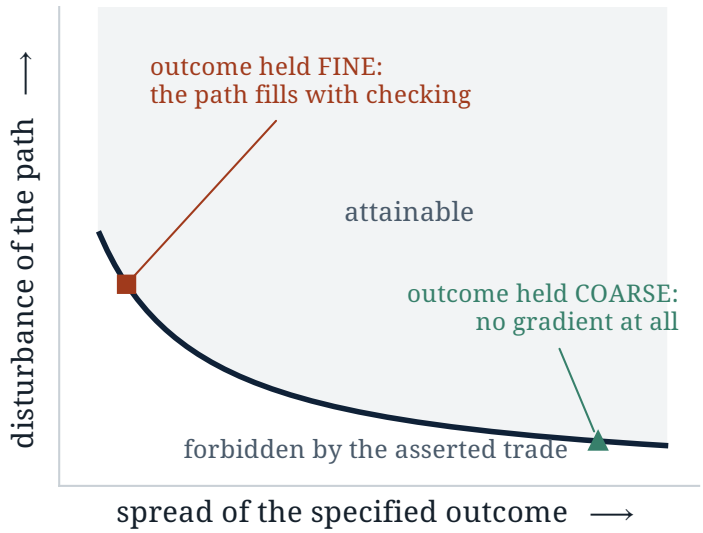


Figure 3.5: **The trade the source asserts, and the constant it says nobody has measured.** Schematic; neither axis carries a scale, because neither quantity has been operationalised. The two poles and their consequences are the source's own. The source writes the trade as $\Delta(\text{outcome}) \cdot \Delta(\text{path}) \geq k$ and states in the same passage that it has not measured k and that no one has. Script: `figs/src/CCP-A03_fig03.py`.

no formal warrant, legitimate as motivation and never as evidence). The labels are collected here.

Borrowing	Category	Warrant, and what is not claimed
Amplitudes and superposition (Sec. 4.1)	3 — analogy, and explicitly refused as more	Reported as physics and used for the image of weighted coexistence only. **Uncited: this bibliography holds no entry for the Born rule.**
Weights on paths, not endpoints (Sec. 4.1)	2 — shared form	Any path-summed cost functional has this property. No physical result is transferred.
Stationary phase, as the neighbour test (Sec. 4.4)	3 — analogy	The physical mechanism is cancellation of phases. The framework's weight cannot cancel (Sec. 4.3), so the transfer is of the *idea* that robust routes survive, not of the mechanism.
2^{-R} as Boltzmann weighting (Sec. 4.3)	2 — shared form	The exponential is the maximum-entropy distribution under a mean-cost constraint [Jaynes, 1957]. No claim that resistance is an energy or that any temperature exists.

Borrowing	Category	Warrant, and what is not claimed
Freezing (Sec. 5)	3 — analogy	No premise of the thermal result is shown to hold. What is used is a claim about defaults, which stands alone.
$\sigma_x \sigma_p \geq \hbar/2$ (Sec. 7.1)	quoted exactly, then 3 — analogy	Quoted in its state-preparation reading [Kennard, 1927]. Its extension to a foreseen future is P1 and is badged at every occurrence.
The time–frequency bound (Sec. 7.1)	1 — derivation	A genuine theorem of signal analysis [Gabor, 1946], proved in the corpus by the same variational argument [Curzi, 2026z]. Nothing psychological is claimed from it.
$\Delta(\text{outcome}) \cdot \Delta(\text{path}) \geq k$ (Sec. 7.5)	3 — analogy, in borrowed notation	The source says so itself and states that k is unmeasured. Reported and badged; no result in this paper consumes it.

The general form of the hazard is documented. Economics imported the formal apparatus of nineteenth-century physics without the conservation law that gave it content [Mirowski, 1989]; laws hold within shielded arrangements rather than universally, so importing a law's form is not importing its scope [Cartwright, 1999]; and the field this corpus's title claims membership of has been surveyed and bounded, with its vocabulary given definitions the corpus can be held to [Castellano et al., 2009]. The constructive standard for making such an analogy checkable is to name the physical quantity, name the social quantity, and state the correspondence explicitly [Vinković and Kirman, 2006]; the table above is this paper's attempt to meet it. **No conserved quantity is exhibited anywhere in this paper, and none is claimed.**

Where this paper cites a corpus result carrying the reference constant $k = 0.483425732988414$ Hz, two cautions apply and are recorded once: the constant is **defined**, exact by construction as $(9/8)^2/\varphi^2$, so its fifteen figures are defensible only on that basis, and **attaching hertz to it is a separate and much stronger claim than computing it**. The general hazard of reading a chosen golden-section ratio as a fact about the world is documented and should be read before any such reading is offered [Markowsky, 1992]. The corpus's own statement is the correct one: that reference constant, the affine effective count, and the choice of twelve components are all badged as chosen rather than derived, and the associated deficit is blind to the decay term, so the intuitive intervention of lowering loss moves the instrument by exactly zero [Curzi, 2026].

8.2 Scale type: a four-point rating, summed and exponentiated

This is the paper's sharpest exposure and it is stated in full, because the objection is specific and the corpus has already committed itself to the rule that answers it.

D1's resistance score is a **four-point rating scale with verbally anchored categories**, of exactly the kind whose scale type has been contested since the format was introduced [Likert, 1932]. Such a scale is, absent a demonstration to the contrary, **ordinal**: the categories are ordered, and nothing establishes that the interval between "would do it while unwell" and "ordinary resolve" equals the interval between "needs a good day" and "needs a better self". The framework then performs two operations on it that ordinal data do not license:

1. **It sums the scores across twelve steps** to obtain R. Addition requires at least an interval scale.
2. **It exponentiates the sum**, forming 2^{-R}, and reads the result as a ratio of likelihoods. Exponentiation of a summed ordinal quantity and the reading of the result as a ratio require a ratio scale with a true zero.

Neither is licensed. The classification that fixes which operations each scale type permits is Stevens's [Stevens, 1946]; the standard against which a claim of quantitative structure is judged is set out in the measurement-foundations literature [Krantz et al., 1971]; and the formal test for whether an attribute has additive quantitative structure at all is conjoint measurement [Luce and Tukey, 1964], whose cancellation conditions are exactly what would have to be checked. **They were not tested.** The strongest single objection to this framework is that psychology systematically assumes quantitative structure without testing for it [Michell, 1997], and **this section concedes it rather than answering it.** Routes that would earn the arithmetic exist and are unused: a probabilistic measurement model that does not require respondents to use numbers [Rasch, 1960], or an interval scale constructed from comparative judgements [Thurstone, 1927], with the construct-validation programme that would accompany either [Cronbach and Meehl, 1955].

The corpus has already stated the rule that governs this case, and the rule is applied here to this paper's own material. The corpus's formula-validation protocol demotes its own moment equation on precisely this ground: because the expression contains a division and a multiplication, inputs that are not instrumented on a common ratio scale with a shared true zero withdraw the magnitude, the number, and the quantitative join as measured quantities [Curzi, 2026t]. The identical objection applies to 2^{-R}, which contains an exponentiation. Applying the corpus's own rule:

- R **is retained as an ordering.** Ranking branches by accumulated resistance is licensed if the individual scores are ordinal and the summation is treated as a heuristic aggregate, which is the weakest defensible reading.
- 2^{-R} **is demoted.** It may not be printed as a ratio of likelihoods, and the statement that a three-point difference is "a factor of eight in expectation" is **withdrawn as a quantitative claim** and retained as an ordering: steeper branches are less likely, and by a margin that grows with the difference.
- **The worked case's ratio is demoted with it.** Section 9.5 reports the arithmetic discrepancy in the source's own terms, because a printed number should be checked against the rule that produced it even when the rule is itself demoted.

This is a real loss of content and it is not disguised. What survives is an ordering procedure with a stated bias toward robustness, which is useful and much less than the source claims.

8.3 Identification, and the framework turned on itself

The Postulate concerns futures coupled to their own foreseeing — which is to say, reflexive systems. The moment a reflexive claim is made about a system containing other people, the question of whether its parameter could be recovered from observation arises, and the answer is no.

The corpus's central formal object is the reflexive loop $T(x) = W(x, M(x))$, fixed in the founding paper of this series. Written as an aggregate relation — a group's state influencing its members while the members constitute the group — **that is a linear-in-means reflexive specification, and the endogenous social effect in such a specification is not identified from observational data**: it cannot be separated from contextual effects or from correlated unobservables, however much data is gathered [Manski, 1993]. The network route is closed by the same family of result, since influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011]. The available remedies are randomisation, an instrument, partial-population variation, or stipulating the reference group and defending

the stipulation, the last having an explicit template in a discrete-choice model that states the conditions under which its interaction parameter is econometrically identified [Brock and Durlauf, 2001]. **This paper adopts none of them.** P1 and R7 are therefore **modelling positions rather than identified empirical results**, and every downstream paper citing them inherits that status.

That reflexivity of this shape occurs at all is separately documented, and the documentation is the existence proof the framework otherwise lacks: when the Black–Scholes option-pricing model was published its prices did not fit the market, and after traders adopted it the market came to fit the model, with dates and price data recording the convergence [MacKenzie and Millo, 2003]. The general statement that classifications of people alter the people classified is argued directly [Hacking, 1995], and the closest existing named doctrine to the framework's reflexive law is Soros's fallibility–reflexivity pair, which supplies the concept without an equation [Soros, 2013]. That literature supplies existence, not identification.

The framework turned on itself. The framework asserts that a system changes its behaviour when it is modelled and observed. Applied consistently, that assertion falls on the framework's own parameters: a halving constant, a fragility penalty or a resistance score estimated in a population that has not read the framework cannot be assumed to survive in one that has, because the decision rules those parameters summarise are themselves responses to the environment the model describes [Lucas, 1976]. The position taken here, stated so that it can be held against the paper: **no stability of this framework's parameters under its own dissemination is claimed, and no result in this paper depends on any.** Three consequences are accepted. Any future estimate of the halving constant must record the reading status of the population it was estimated in. A before-and-after design on a single practitioner is confounded by construction. And the framework's own placement technique (Section 6) is a deliberate regime change for the person applying it, which the framework predicts and which makes its own effect sizes non-transferable.

No scaling or power-law claim is made anywhere in this paper. The weight law is exponential in R and its exponent is *defined by the model rather than estimated from data*; where any paper in this corpus does make a distributional or heavy-tail claim, the principled maximum-likelihood, Kolmogorov–Smirnov and likelihood-ratio procedure is binding and must be run and reported [Clauset et al., 2009]. The absence of that procedure here is a scope statement and not an omission. The general warning that a neighbouring model fitting equally well is the norm rather than the exception applies to the exponential form as much as to any other [Stumpf and Porter, 2012, Durlauf, 2005].

3.0.9 9. Results

9.1 Definitions fixed. D1 (the four-point resistance score) and D2 (the licence clause) are stated in the form the remaining papers cite. Neither is a finding. D1 does not survive the scale-type audit of Section 8.2 in the form the source uses it.

9.2 Results established by construction. R1, R3, R4, R6 and R7 are specifications and procedures rather than measurements. R1 is a scoring rule; R3 is a robustness test with a stipulated penalty; R4 is a three-way classification with a named but uncited mechanism; R6 is three placement laws with an uncited external warrant; R7 is a stated trade-off whose constant the source says has never been measured.

9.3 Results established by derivation. R2 is established: $2^{-R} = e^{-R\ln 2}$ is exponential weighting of states by cost, which is the maximum-entropy distribution under a mean-cost constraint [Jaynes, 1957], and it is not a phase-based amplitude. The consequence — that **nothing in the framework's own arithmetic can cancel** — is established and it goes against the framework's own presentation of ambivalence (Section 4.5). The arithmetic join of Section 6.4, that removing one resistance point doubles a branch's weight, is established as arithmetic on R1 and establishes nothing empirical.

9.4 Hypotheses tested. None. H1 through H5 are stated in falsifiable form and none is tested here. No dataset was collected. H5 is not currently testable as stated, for the reasons in

Section 7.5. This is reported in the same plain language a positive result would receive: **the five hypotheses this paper states stand untested, and the paper that states them supplies no evidence for any of them.**

9.5 An arithmetic discrepancy in the source, printed rather than corrected. The source's worked example records two branch resistances explicitly — 29 for the continuation branch and 9 for the intervention branch — and then states that "the halvings put the counter at roughly a thousand times the weight of holding on".

Under the framework's own rule, $\Delta R = 29 - 9 = 20$ and the weight ratio is $2^{20} = 1{,}048{,}576$: approximately **one million**, not one thousand. A ratio of a thousand corresponds to $\Delta R = \log_2 1000 = 9.9658$, which is half the recorded difference.

Three readings are possible and this paper does not choose between them. The printed ratio may be an error of three orders of magnitude. One of the two resistance totals may be misprinted. Or the halving may have been intended per *two* points rather than per point, in which case $2^{10} = 1024$ reproduces "roughly a thousand" exactly — but that reading contradicts the rule as stated twice in the same material. **The discrepancy is recorded here, and no downstream paper should quote the "roughly a thousand" figure.** Figure 1 plots both readings.

9.6 Reproduction of the corpus's own arithmetic. Three numerical claims carried by this paper's internal citations were recomputed.

- The time–frequency bound: $1/(4\pi) = 0.0795774715459477$, printed as 0.079577471545948 [Curzi, 2026z]. **Reproduced exactly.**
- The minimum-gap schedule in cents: $1200\log_2(9/8) = 203.9100$, $1200\log_2(50/49) = 34.9756$, $1200\log_2(289/288) = 6.0008$ and $1200\log_2(1682/1681) = 1.0296$, all as printed [Curzi, 2026]. **Reproduced exactly, to the four decimals given.**
- The decoherence shortfall: $10^{-3}/(1.98 \times 10^{-17}) = 5.05 \times 10^{13}$ against the printed \$5.04 ×10^{13}$ [Curzi, 2026aa]. The printed value is recovered from an unrounded $\tau_{\text{dec}} \approx 1.984 \times 10^{-17}$ s, so the discrepancy is a rounding of the *displayed input* rather than an error in the result. **Reproduced, with that clarification**, and the derived \$5.04×10^{13}/10^{6} = 5.04×10^{7}$ follows.

3.0.10 10. Discussion

10.1 What follows

Three consequences follow and they are this material's genuine yield.

The first is that **a scoring rule which ignores the destination is a real methodological choice with a real cost.** Covering the endpoint during scoring removes the single largest source of motivated distortion in any self-assessment of a plan, and it is cheap. Its cost is that the procedure cannot rank branches by anything a practitioner wants, which is why the protocol requires the disagreement between the weight ranking and the wanting ranking to be written down as a finding rather than resolved.

The second is that **the licence clause is prior to the technique, not appended to it.** A method that reliably lowers the resistance of a step is a method for altering another person's behaviour, and the framework's decision to state the constraint before any technique is the correct ordering. The corpus states the same constraint elsewhere as a condition rather than a courtesy.

The third is the trade-off's one practical rider, which survives the removal of the Postulate entirely: **spend checking on the path and not on the outcome.** Testing a route with real weight makes the route more able to carry the practitioner; auditing an outcome against a template does not, and it corrupts the indicators it audits — which is the one item of the mirror catalogue this corpus can fully cite.

10.2 What does not follow

That a rule is stated exponentially does not make it exponential; that a score is summed does not make it additive; that a trade-off is written as an inequality does not make it one; and that a corpus can cite two of five claims does not make the other three false or true.

Two further non-implications are worth stating. First, **the framework's cancellation language is not licensed by its own arithmetic** (Section 4.3), and any downstream use of the framework's ambivalence account as though it were a computation is unsupported. Second, **the Postulate is not made stronger by the correctness of Section 7.1.** The source states the physical relation better than most popular accounts, and that accuracy is a virtue of the exposition and not evidence for the conjecture that follows it. A carefully stated premise and an unlicensed extension are two separate things, and the first is not a warrant for the second.

10.3 Relation to the corpus

The results fixed here are consumed downstream in specific ways. The launch-budget instrument turns a resistance-like barrier into an explicit fund with a fixed point — $L^{Q} = \max(0,\ \eta M - C)$, $\Lambda_{\infty} = L/\lambda$, departure feasible only where $\Lambda_{\infty} > E_b$ and reached after $t^{*} = \ln(1 - E_b/\Lambda_{\infty})/\ln(1 - \lambda)$ periods [Curzi, 2026p] — which is the quantitative form of the claim R1 makes qualitatively, that accumulated resistance is what has to be paid before a branch moves. The corpus's evidential-status apparatus supplies the schedule of exactly representable gaps that any resolution claim must live inside: 1200 cents at two elements, 203.9100 at three to six, 34.9756 at seven to sixteen, 6.0008 at seventeen and above, and 1.0296 at forty-one and above [Curzi, 2026]. That schedule is R7's structural cousin: **the finer the scheme, the smaller the distinction it can still carry**, and the shrinkage is exactly computable there and merely asserted here.

Four results elsewhere bound what this paper may claim. The corpus's falsification suite measured its own shuffled-input control and found it **anti-conservative**: against a nominal 0.05 across 500 null trials, independent permutation rejected at 0.108, circular shift at 0.120 and block permutation at 0.060 [Curzi, 2026u]. That is the correct prior for any future test of H1–H5. The corpus's two-column refinement bounds what the mirror catalogue's partial support can be made to do: an evidence column carries what could have come out otherwise and a precedent column carries only that an idea was independently arrived at and preserved, so two sources agreeing does not double the evidence when only one of them is evidence [Curzi, 2026w]. The corpus's audit of a traditional operation set against a parameter model returned seven mappings, two splits, one rendering and three rejections over thirteen entries, with five entries carrying a failure [Curzi, 2026] — the same refusal discipline applied to inherited material, and the model for Section 7.3's audit. And two structural results state the general shape of R5's constraint from the other side: an effective count measures independence and not capacity, so a structure that clears an independence threshold still says nothing about whether any part of it can carry the load [Curzi, 2026], and a separated control channel trusted *because* it is separate is itself the failure mode, separation being necessary and not sufficient [Curzi, 2026]. The framework's licence clause is a separation claim of exactly that kind and inherits exactly that caution.

One further result is cited for the structure it shares: the property that makes an insurance reserve resistant to third-party claims — real liabilities owed to unrelated parties — is the same property that makes it unavailable to its owner, leaving only 7.58 per cent realisable within a year [Curzi, 2026]. That is R7 in a different domain: the same feature buys the protection and imposes the cost, and the trade cannot be escaped by wanting both ends of it.

3.0.11 11. Verification

This section is not a summary. It decomposes the paper's results into steps small enough to check one at a time, and labels each `ARITHMETIC` (a reader can redo it), `DEFINITION` (true by this

paper's own construction), CITED (a named external source supports it) or POSTULATED (asserted here, no external support). The razor is applied: any postulate whose removal leaves the results standing has been removed.

11.1 Postulates

- **P1 (DEFINITION).** A branch is an ordered sequence of steps, and each step carries a score in $\{0, 1, 2, 3\}$. *Check:* D1.
- **P2 (DEFINITION).** The resistance of a branch is the sum of its step scores. *Check:* D1. **Note: summation of an ordinal rating is not licensed, and Section 8.2 demotes what may be done with the sum.**
- **P3 (DEFINITION).** The weight of a branch is $2^{-\Delta R}$, where ΔR is measured from the smallest branch resistance. *Check:* the rule as stated.
- **P4 (ARITHMETIC).** $2^{-\Delta R}$ falls by a factor of two per point, so $\Delta R = 3$ gives $1/8$ and $\Delta R = 20$ gives $1/1{,}048{,}576$. *Check:* redo it.
- **P5 (CITED: `jaynes1957information`).** The exponential weighting of states by cost is the maximum-entropy distribution subject to a mean-cost constraint; it is positive everywhere and therefore admits no cancellation.
- **P6 (POSTULATED).** The completion likelihood of a course of action falls geometrically in accumulated step steepness, with a halving constant of one point. *Check:* **not checked. This is the substantive content of H1 and it is asserted.**
- **P7 (POSTULATED).** A branch that cannot survive one middle-step failure is half as likely as one of equal resistance that can. *Check:* **not checked**; the factor is a stipulation.
- **P8 (POSTULATED).** Futures requiring the smallest change from present arrangements consolidate first, whether or not they are chosen. *Check:* **not checked**, and no entry in this bibliography supports the mechanism.
- **P9 (POSTULATED).** A cue sited on a walked path lowers the resistance of the step it represents. *Check:* **not checked**; the entire external warrant the framework claims for it is absent from this bibliography (Section 12.3).
- **P10 (CITED: `kennard1927quantenmechanik`).** $\sigma_x \sigma_p \geq \hbar/2$ is a property of the quantum state, not of measurement disturbance.
- **P11 (CITED: `gabor1946theory`).** A signal cannot be arbitrarily localised in time and in frequency simultaneously; the bound is a theorem of signal analysis and mentions no observer.
- **P12 (CITED: `campbell1979assessing`, `goodhart1975monetary`).** An indicator adopted for consequential decisions is gamed, and the process it tracked is corrupted.
- **P13 (POSTULATED).** The uncertainty of a future coupled to its own foreseeing is the same kind as that of a quantum state, and the trade-off structure carries over. *Check:* not checkable; this is the framework's Postulate and is badged wherever it appears.

Postulates dropped as unnecessary: the reactance mechanism (Section 6.1), the planning fallacy (Section 4.4), the approach–avoidance account of ambivalence (Section 4.5), the frequency illusion and the fresh-start effect (Section 6.3), and the historical claim about the retirement of Heisenberg's disturbance parable (Section 7.1). Each was tested for load and carries none: every result below survives its removal. All are reported in the text and none is used.

11.2 The compounding chain

Each line consumes only the outputs named on its left.

```
P1 + P2          -> L1   a branch has a resistance R          DEFINITION
L1 + P3          -> R1   the branch weight is 2^-(dR)         DEFINITION
R1               -> L2   dR = 3 gives 1/8; dR = 20 gives
                         1/1,048,576                          ARITHMETIC
```

L2	-> D1'	the source's printed "roughly a thousand" needs dR = 9.97, not 20	ARITHMETIC
R1 + P5	-> R2	the weight is Boltzmann-style exponential discounting, NOT a phase amplitude	CITED
R2	-> L3	nothing in this arithmetic can cancel; the cancellation talk is a picture, not a computation	ARITHMETIC
P7	-> R3	the neighbour test halves a fragile branch again	POSTULATED
P8	-> R4	the freezing law and its three classes of settled future	POSTULATED
(stipulation)	-> R5	the licence clause	DEFINITION
		(its mechanical limb)	POSTULATED
P9	-> R6	the three placement laws	POSTULATED
R1	-> L4	one resistance point removed doubles a branch's weight	ARITHMETIC
P10 + P11	-> L5	conjugate bounds are facts about states and signals, not observers	CITED
P12	-> L6	a finely specified outcome corrupts its own indicators	CITED
L5 + L6 + P13	-> R7	the resolution-disturbance law	POSTULATED
R1 + P2	-> C1	branch weights may be read as ratios of likelihoods	WITHDRAWN (Sec. 8.2)

What this chain shows, stated plainly. R2 is the only result in this paper that is CITED at its load-bearing step, and its content is a *restriction*: the framework's weighting is an ordinary exponential discount and cannot do the one thing the borrowed vocabulary suggests. Everything else about practice is POSTULATED. R1 is DEFINITION and the reading of it that the framework most wants — a ratio of likelihoods — is withdrawn on scale-type grounds and appears in the chain as WITHDRAWN, which is the honest terminal state for it.

A chain whose steps are mostly DEFINITION and POSTULATED is not a proof, and this one is mostly DEFINITION and POSTULATED. What it is instead is a specification of which premises would have to be settled: P6 (the halving constant), P7 (the fragility penalty), P8 (default consolidation), P9 (that a placed cue lowers resistance), and P13 (the Postulate, which is not the kind of thing that can be settled). Four of the five are testable, three of them cheaply, and none was tested.

11.3 The checks a reader should run first

1. **Redo the weight law.** $2^3 = 8$ and $2^{20} = 1{,}048{,}576$. Then compute $\log_2 1000 = 9.9658$ and confirm the discrepancy reported in Section 9.5.
2. **Redo the resistance ceiling.** Twelve steps at a maximum score of 3 gives 36, so the recorded totals of 29 and 9 sit at 81 and 25 per cent of the maximum. Figure 2 plots both.
3. **Redo the time–frequency constant.** $1/(4\pi) = 0.0795774715459477$. The corpus prints 0.079577471545948 [Curzi, 2026z]. It agrees.
4. **Redo two entries of the gap schedule.** $1200 \log_2(9/8) = 203.9100$ and $1200 \log_2(1682/1681) = 1.0296$ [Curzi, 2026]. Both agree to the printed decimals.
5. **Check that $\sigma_x \sigma_p \geq \hbar/2$ carries its state-preparation reading at every occurrence.** There are three, in Sections 7.1, 8.1 and 11.1, and all three carry it.
6. **Check that the analogy table in Section 8.1 assigns category 1 to exactly one entry**, and that the entry is a theorem of signal analysis from which nothing psychological is drawn.

7. **Check that no sentence in this paper reads 2^{-R} as a probability.** Section 8.2 withdraws that reading and the Conclusion repeats the withdrawal.

11.4 Dimensional audit

Symbol	Meaning	Dimension	Unit	Status
step score	steepness of one step	none assigned	**ordinal categories, no unit**	defined (D1); **not shown to be interval or ratio**
R	branch resistance	none assigned	points	defined as a sum of ordinal scores; **the summation is not licensed**
ΔR	resistance above the easiest branch	none assigned	points	derived from R
$w = 2^{-\Delta R}$	branch weight	dimensionless	—	derived; **demoted to an ordering (Sec. 8.2)**
$\beta = \ln 2$	discount rate per point	inverse points	per point	defined by the choice of base 2; **chosen, not measured**
resolution	fineness of a specification	none assigned	**no operational unit exists**	undefined in the source; not operationalised anywhere
disturbance	interference with the path	none assigned	**no operational unit exists**	undefined in the source; not operationalised anywhere
k (trade-off floor)	the asserted bound's constant	product of two unmeasured quantities	**none**	**asserted; the source states it has never been measured**
$\sigma_x, \sigma_p, \hbar$	quantum spreads and reduced Planck constant	m; kg·m·s^{-1}; J·s	SI	quoted from physics; not quantities of this framework
$\Delta t, \Delta f$	time and frequency spreads (quoted)	time; inverse time	s; Hz	quoted from `ccp:CCP-032`; their product is dimensionless
τ_{dec}	decoherence time (quoted)	time	s	quoted from `ccp:CCP-033`

Does the central relation balance? $w = 2^{-\Delta R}$ requires ΔR to be dimensionless, and it is — but only in the trivial sense that a count of ordinal points carries no unit. **A dimensionless exponent is not the same thing as a licensed one**, and the balance here is formal: it is satisfied by construction and constrains nothing.

Does the trade-off relation balance? $\Delta(\text{outcome}) \cdot \Delta(\text{path}) \geq k$ **cannot be checked**, because neither factor has a dimension, a unit, or an operational definition, and the constant is therefore a number attached to a product of two undefined quantities. **The source is silent on all three and this paper does not supply them.** A dimensionless placeholder is the honest entry; an invented unit would not be.

11.5 Limits and boundary behaviour

- $\Delta R \to 0$**.** All branches carry equal weight and the procedure ranks nothing. Correct and uninformative.
- $\Delta R \to \infty$**.** The weight of the steeper branch goes to zero. The model therefore asserts that a sufficiently steep branch is impossible rather than merely unlikely, which is a **defect of the model's domain**: no finite resistance should drive a completion probability to exactly zero, and the framework supplies no floor. The domain should be restricted to the range over which the branches are comparable, and the source does not do this.
- **All steps scored 0.** $R = 0$ for every branch, all weights equal 1, and the procedure is silent. The framework's own account predicts this case is empty, since a branch with no resistance anywhere would already have been taken.
- **All steps scored 3.** $R = 36$ at twelve steps. Weights are equal again, because the rule is anchored at the minimum. **The rule cannot distinguish "all branches easy" from "all branches impossible"**, and the source does not note this.
- **Fragility penalty applied to every branch.** The relative weights are unchanged, since a common factor divides out. The neighbour test therefore has content only when branches differ in robustness, which the framework does not state.
- **Resolution held at its finest.** The path is maximally constrained and, by the framework's own account, all arrival is by listed routes only. The framework predicts the highest indicator corruption here, and that limb is the one this corpus can cite [Campbell, 1979].
- **Resolution held at its coarsest.** No gradient, and no arrival. The framework's own gradient material supplies this limit: a difference is required for motion, and an unspecified outcome supplies none.
- **Coupling between seeing and seen** $\to 0$**.** The Postulate is silent by its own stated boundary (Section 7.4). This is the only boundary in this paper at which the framework correctly declines to say anything.

On invariance. The weight rule is invariant under the addition of a constant to every branch's resistance, since it is anchored at the minimum, and it is invariant under relabelling of steps within a branch, since only the sum enters. Both are consequences of the construction rather than findings, and neither has been tested at any scale, because nothing here is measured. **No conserved quantity is exhibited.**

3.0.12 12. Limitations

12.1 The scoring rule's arithmetic is not licensed by its own scale type. This is the paper's principal exposure and it is stated first. A four-point verbally anchored rating is ordinal absent a demonstration otherwise [Likert, 1932, Stevens, 1946]; summing it and exponentiating the sum requires interval and ratio structure respectively; and the formal test that would establish either was not run [Luce and Tukey, 1964, Krantz et al., 1971]. The strongest single objection to this framework is that psychology assumes quantitative structure without testing for it [Michell, 1997], and Section 8.2 **concedes it**. The corpus's own demotion rule applies and has been applied [Curzi, 2026t]. **Falsified by:** a conjoint-measurement study on any operationalisation of step resistance that fails double cancellation, which would remove the quantitative content of R1 entirely and leave only the ordering.

12.2 The halving constant and the fragility penalty are stipulations. Nothing establishes one halving per point rather than per two points or per half a point, and nothing establishes

that fragility costs a factor of two. Both are round numbers chosen for arithmetic convenience, and the model's entire quantitative output is a function of them.

12.3 Claims with no supporting literature in this corpus's bibliography. Listed individually, because a reviewer will find them and finding them first is the point.

- *Pressure on a perceived freedom produces counter-movement (reactance).* Load-bearing for R5's mechanical limb and for failure mode CC-L4. **No entry.** Retained uncited; R5's ethical argument stands without it.
- *If-then plans tied to concrete cues outperform bare intentions; habitual performance is cue-driven; some twenty seconds of added or removed friction decides behaviour.* The **entire** external warrant for Section 6.2 and for H4. **No entry for implementation intentions, cue-driven habit, choice architecture, or default effects.** Reported as the framework's and uncited. The scaffold for this paper names an organisational-dynamics cluster as the source of this support; **the cluster does not contain it**, and no citation was attached to make it appear that it does.
- *Plans are systematically imagined as their own best case (the planning fallacy).* Load-bearing for R3's rationale. **No entry.** Retained uncited.
- *Behaviour tracks situations more reliably than resolutions, and highly self-controlled people experience fewer temptations rather than winning more of them.* The pivot of Section 6's argument. **No entry.**
- *The frequency illusion; the fresh-start effect of temporal landmarks; that scheduled attention accumulates relevance.* The mechanism of the named day (Section 6.3). **No entry.**
- *Measurement reactivity and the mere-measurement effect.* Two of the five items in the mirror catalogue (Section 7.3). **No entry.**
- *The Born rule and the postulates of quantum mechanics; the historical retirement of Heisenberg's disturbance parable.* Sections 4.1 and 7.1. **No entry.** Reported as physics and used for nothing; excluded from the verification chain.
- *Approach–avoidance conflict; anchoring; availability; optimism bias; catastrophising; motivated reasoning.* The mechanisms named in the three failure catalogues. **No entry for any of them.**

That is **eight classes of unsupported claim**, and they are more of this paper's material than any other single category. A reader should treat the failure catalogues and the placement laws as the framework's clinical folklore, usefully organised, and not as findings.

12.4 The cancellation account is unsupported by the framework's own arithmetic. Section 4.3 establishes that 2^{-R} is positive and cannot cancel. The framework's ambivalence account is retained as an instruction and marked as a picture, and no downstream paper should treat it as a computation.

12.5 Parameters chosen rather than measured. Every parameter in every figure is either the source's own recorded number or a stipulation. Figures 1 and 2 plot the source's own printed values and the arithmetic that follows from them. Figures 3 and 4 are schematics and carry no scale on any axis; Figure 3 in particular is drawn without axis units precisely because neither of its quantities has been operationalised. Figure 5 is an audit of this corpus's bibliography and not of the literature.

12.6 The trade-off law is not currently testable as stated. Neither resolution nor disturbance has been operationalised, the constant has never been measured, and the source says so. H5 is therefore listed as a hypothesis in name and is not one in practice until the two quantities are defined. **This is the weakest claim in the paper and it is the one the framework leans on hardest.**

12.7 Nothing here is identified. As an aggregate specification, any reflexive claim of the framework's form is not identified from observational data [Manski, 1993], and the network route is closed by the generic confounding of influence with selection [Shalizi and Thomas, 2011]. This paper adopts no remedy and therefore states P1 and R7 as **modelling positions rather than identified empirical results** (Section 8.3). **Resolved by:** a randomised or partial-

population design, or a defended stipulation of the reference group in the manner of Brock and Durlauf [2001].

12.8 The framework's own parameters are not exempt from its own thesis. By the Lucas critique, a halving constant estimated in a population that has not read the framework cannot be assumed stable in one that has [Lucas, 1976]. No stability is claimed and no result here depends on any; the cost is that no cross-population comparison of a resistance score is currently interpretable, and that the framework's own technique constitutes a regime change for its user.

12.9 The exponential form is one of a family and nothing discriminates it. A one-parameter generalisation produces a family of weightings containing the exponential as a special case [Tsallis, 1988], and the general problem is that many mechanisms produce indistinguishable curves [Stumpf and Porter, 2012]. A reviewer will ask for the differential prediction [Durlauf, 2005]. **This paper does not have one.**

12.10 The mirror catalogue is not uniformly established, and the source says it is. Section 7.3 audits it: one item fully supported, two supported only in a bounded form, two with no entry. The source's blanket characterisation is not sustained.

12.11 The worked cases are illustrative. The two worked cases reported in Sections 4.2 and 9.5 are single, self-reported, retrospective accounts drawn from the source manuscript. They illustrate what the procedures return and establish nothing about whether the procedures work. The base rate against which any single such report should be read is not encouraging [Ioannidis, 2005, Open Science Collaboration, 2015].

12.12 A printed number in the source does not follow from the source's own rule. Section 9.5 records the discrepancy in full. No claim in this paper depends on the figure, and the discrepancy is printed rather than harmonised.

12.13 No hypothesis was tested and no data were collected. H1–H5 stand. The discipline that would make a later confirmatory attempt credible is specified and not exercised [Nosek et al., 2018], and the analytic flexibility available in any such attempt should be constrained in advance [Simmons et al., 2011]. The corpus's own control measured anti-conservative at 0.108 against a nominal 0.05 [Curzi, 2026u], which is the prior any such test should be read against.

3.0.13 13. Conclusion

The seven results are stated below in the exact form in which the rest of the corpus cites them. Downstream papers should quote this text rather than paraphrase it.

R1. The comfort-weighted amplitude: score each step of a candidate branch 0 to 3 for steepness, sum to a resistance R, and weight the branch as $2^{-\Delta R}$ measured from the easiest branch, so that a three-point difference in steepness is a factor of eight. **The quantitative reading of that weight is withdrawn on scale-type grounds (Section 8.2): the score is ordinal, its summation and exponentiation are not licensed, and R1 survives as an ordering of branches by accumulated resistance and not as a ratio of likelihoods.**

R2. The pedigree, stated honestly: the weight is exponential, Boltzmann-style discounting of states by cost — $2^{-R} = e^{-R\ln 2}$, the maximum-entropy form under a mean-cost constraint — and explicitly **not** a phase-based path amplitude. It follows that **nothing in this arithmetic can cancel**, and the framework's account of ambivalence is a picture rather than a computation.

R3. The neighbour test: break one middle step of a branch and halve its weight again if it cannot reroute, penalising plans that require every step to succeed. The penalty is a stipulation.

R4. The freezing law: futures adjacent to present arrangements consolidate first, by adjacency and default rather than by choice, with three classes of settled future distinguished — frozen by choice, frozen by geography, and settled in appearance with something unresolved running underneath.

R5. The licence clause preceding all cue technique: arrangement is legitimate only on one's own path or on the path of a person who has asked. The ethical argument stands on its own; **the framework's further claim that the prohibition enforces itself mechanically, through counter-movement proportional to the discovered push, has no supporting entry in this corpus's bibliography and is reported uncited.**

R6. Three placement laws: site the cue on an already-walked path, represent the first step and never the summit, and re-site the cue once it has stopped being noticed — with friction priced in seconds of added or removed effort. **The external warrant the framework claims for all three is absent from this corpus's bibliography.** What survives is arithmetic on R1: removing one resistance point from one step doubles the branch's weight.

R7. The resolution–disturbance law: fine specification buys precision at the cost of path disturbance, indicator-optimisation and the foreclosure of unlisted routes, while infinite coarseness exerts no gradient at all. The trade's *direction* at the fine pole is supported by indicator corruption [Campbell, 1979]; the existence of a floor beneath it is postulated, and the source states plainly that its constant has never been measured by anyone.

Counter-result. The uncertainty material declines the inflated reading of its own borrowed physics. The Postulate is stated as a postulate; the physical relation is given only in its correct state-preparation form and never as a claim about measurement disturbance; and a named requirement holds that a practitioner must be able to state the theorem, the ordinary observer effect, and the framework's own conjecture in three separate sentences on demand. To that this paper adds an audit the source did not run: of the five phenomena the source calls uniformly established, one is fully supported by this corpus's bibliography, two are cited only by entries that bound them, and two have no entry at all.

What is established, then, is one restriction — that the framework's weight is an ordinary exponential discount and cannot cancel — one theorem of signal analysis from which nothing psychological is drawn, and one fully supported mechanism, indicator corruption, on which the practically useful half of the trade-off law rests. What is not established is any of the five hypotheses, the licensing of the scoring arithmetic, the halving constant, the fragility penalty, or the existence of the trade-off's floor. The paper's contribution at this stage is to name the four testable premises whose settlement would decide the matter — the halving constant, the fragility penalty, default consolidation, and whether a placed cue lowers resistance — and to note that three of the four are cheap to test and none was tested.

3.0.14 14. Declarations

Data and code availability. No external dataset was collected for this paper and none exists. The worked cases reported in Sections 4.2 and 9.5 are drawn from the source manuscript's own record and are not a dataset. All figures are generated by the scripts named in their captions (`figs/src/CCP-A03_fig01.py` through `figs/src/CCP-A03_fig05.py`), which import the shared style module `lib/figstyle.py`; every number plotted is either computed within the script from a rule stated in this paper or quoted from the source text identified in the script's docstring. Three of the five scripts contain assertions that fail if the numbers they plot cease to agree with the source claims they are drawn from.

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the analysis, the figures, and the text.

Ethics. No human-subjects research was conducted for this paper. No personal data were collected or processed. The worked cases in Sections 4.2 and 9.5 are drawn from the source manuscript, are presented under the descriptions they carry there, and are illustrative rather than evidential; they were not gathered under a research protocol and are not offered as one. Nothing in this paper is clinical advice. The licence clause of Section 6.1 records the framework's own instruction that its placement technique is not to be applied to another person's

environment without that person's request, and this paper reports that instruction as a constraint on the method rather than as an endorsement of the technique.

Use of the ceremonial register. The source manuscript is in part an initiatory document. This paper analyses that register as an object of study and does not reproduce it; ceremonial names are reported only where they identify a construct this paper renames.

Chapter 4

The Spoken Word and the Returned Difference: Performative Commitment, Echo Reading, and the Doctrine of Category Violation

Abstract. This instalment covers the framework's account of commitment, its treatment of feedback, and an interlude on objects that violate a classification scheme without harming anyone. A spoken commitment is analysed as a performative directed at the speaker, with three converging mechanisms named: consistency pressure following a voiced position, self-perception from overheard testimony, and the commitment device that re-prices later defection. The present tense is argued to be calibration rather than etiquette, since a future-tense formulation is consistent with a day identical to yesterday and therefore generates no dissonance for the consistency mechanism to act on. A prohibition follows and is stated as law: a future may not be voiced over another person's life, including benevolently, because imposed futures are registered as pressure and produce counter-movement, so that the attempt manufactures its own refutation. The feedback node writes the loop as sent x, returned y = T(x), and read as the difference y - x, with the argument that a channel returning a signal unchanged has reported nothing about itself, that all information lives in the transform, and that a null return is itself a datum with two possible readings. The interlude defines a category violation as an object or act that harms no person while standing in two boxes a schema swore were disjoint, and argues that such objects are diagnostic of where a classification has torn. Three conditions constrain the doctrine: clean acquisition, clean possession, and self-decoding.

Keywords: performative utterance; consistency; commitment device; reactance; feedback transform; null return; category violation; self-decoding

Series: A. Foundations **Paper:** CCP-A04

4.0.1 1. Introduction

1.1 The question

Once a framework has scored a set of candidate futures and arranged an environment around one of them, three questions remain and none of them is about scoring.

What does *saying* a future do? The framework's answer is that a spoken commitment is not a report about a future but an act performed on the speaker, and that its grammar is calibration rather than etiquette.

What does a *reply* tell you? The framework's answer is that a returned signal carries its information in the transform applied to it and nowhere else, so the reading rule is to record what was sent, record what came back, and read the difference — with the corollary that a channel returning a signal unchanged has reported only that nothing in it could touch the signal.

And what should be done with an object that a classification scheme cannot file? The framework's answer is a definition — an object that harms no person while standing in two boxes a scheme declared disjoint — and three constraints that between them refuse almost everything the definition might be thought to license.

Twenty-seven corpus papers cite the results established here. Because those results are load-bearing elsewhere, this paper states them in a form intended to be quoted rather than paraphrased.

It also reports, at the same precision, what this corpus cannot support. **The three mechanisms on which the paper's first result rests — consistency pressure following a voiced position, self-perception from overheard testimony, and the re-pricing of later defection — have no supporting entry anywhere in this corpus's external bibliography.** So does the prohibition's mechanical limb, so does the substitution effect that motivates the framework's quarantine rule, and so does the attentional cost that motivates its burning rule. Six of the framework's named mechanisms in this material are uncited, and they are listed individually rather than papered over with the nearest plausible key.

1.2 What this paper contributes

Seven results are established or fixed here, labelled R1 through R7 and restated verbatim in the Conclusion.

R1 analyses a spoken commitment as a performative directed at the speaker, supported by three converging mechanisms: consistency pressure following a voiced position, self-perception from overheard testimony, and the re-pricing of later defection. **R2** states the present-tense requirement as calibration rather than style: a future-tense formulation is consistent with a day identical to yesterday, so it generates no dissonance and the consistency mechanism has nothing to act on. **R3** states the prohibition on voicing a future over another life, including benevolently, on the ground that imposed futures are registered as pressure and manufacture their own refutation. **R4** states the feedback reading rule: record what was sent, x, and what returned, $y = T(x)$, and read the difference $y - x$, since a channel that returns a signal unchanged reports only that nothing in it could touch the signal. **R5** treats the null return as a datum with exactly two readings — a distance beyond the listening window, or a medium that absorbs the frequency entirely — rather than as an absence of evidence. **R6** defines a category violation as an object or act that harms no person while standing in two boxes a scheme declared disjoint, holds the law of persons and the law of categories apart, and gives the person-law priority. **R7** states the three constraints on the doctrine: no coercive acquisition, possession by clean holding only, and self-decoding, since an explained instance is a spent one.

A counter-result is stated with equal weight. **The category material constrains far more than it licenses.** Category innocence is explicitly refused as a defence of acquisition; the doctrine forbids moving any harm through a discovered gap in a scheme; and the framework concedes on the record that the case in which the feedback loop runs through another person's will is an unsettled seam which it does not close.

1.3 Register, sources, and the chunk boundaries

The source manuscript is written in two registers, a narrative one and an analytic one, deliberately braided. This paper reports the analytic register. The narrative material is a fictional exposition and carries no analytical claim the treatise does not also carry; it is not reproduced. Where the source's ceremonial vocabulary names a construct — *the word that closes the sea*, *the rite of the far shore*, *the lawful crime* — the construct is renamed to what it is, and the ceremonial name is reported once, as an object of study. The last of those three is the source's own term for the object R6 defines; this paper uses the neutral term **category violation** throughout and reports the source's term once, here, because downstream papers refer to the material by that name.

The source volume was re-split after a de-framing correction, and the boundaries of the file assigned to this paper do not coincide with the source's divisions. **At the opening**, the file begins with the operating protocol and failure catalogue of the preceding construct — the resolution material whose law is stated in the preceding instalment of this series, which records that its protocol falls here. That protocol is reported in Section 3, both because it is the only place in the volume where the framework's central trade-off is actually measured and because the preceding paper states that it will be found here. **At the close**, the file ends in the middle of the category material's fifth and final teaching section, so that section is reported in full while any closing apparatus attached to it falls in the file assigned to the next paper in this series.

4.0.2 2. Hypothesis

The material below divides into claims that could be wrong, definitions and stipulated constraints that could not, and one postulate inherited from earlier in the series. A hypothesis that nothing could falsify is restated as a definition or a constraint and labelled as such [Popper, 1959].

H1 — the performative hypothesis. A commitment spoken aloud, by a speaker, in the present tense, and to no audience, raises the rate at which the speaker's subsequent conduct moves toward the commitment, relative to the same commitment held silently. *Predicted:* matched undertakings differ in completion by voicing condition, with the spoken-to-no-audience condition highest. *Falsified by:* no difference between spoken and unspoken conditions at matched content. *Status:* testable and cheap to test; **not tested here.** All three mechanisms the framework offers for it are uncited (Section 11.3).

H2 — the tense hypothesis. A future-tense formulation generates no pressure toward action because it is consistent with an unchanged day, whereas a present-tense formulation is not. *Predicted:* among spoken commitments matched on content, present-tense formulations produce more early action than future-tense ones. *Falsified by:* equal action rates by tense, or higher rates for the future tense. *Status:* **not tested here.** The prediction follows from the consistency mechanism, which is uncited, so a failure of H2 would be evidence against the mechanism rather than against the grammar.

H3 — the imposition hypothesis. A future voiced over another person, without that person's request, lowers the rate at which they move toward it relative to no intervention, and does so more strongly the more clearly the steering is detected. *Predicted:* movement away from an imposed future, scaling with detected pressure. *Falsified by:* movement toward an imposed future, or movement independent of detection. *Status:* **not tested here**, and the mechanism the framework names for it has **no supporting entry** in this corpus's bibliography (Section 11.3). The ethical argument for R3 does not depend on H3 and is stated separately in Section 4.4.

H4 — the residue hypothesis for returns. The deviation between a signal sent and a signal returned carries information about the medium that no property of the arrival alone carries. *Predicted:* an observer given only the deviations between a set of sent and returned signals can characterise the medium as well as one given the full returns. *Falsified by:* deviations

carrying no more than arrivals about the medium. *Status:* **true by construction for a linear channel**, where the impulse response is exactly the transform [Oppenheim et al., 1996], and **not tested** for any social channel, where linearity is not established and is implausible.

H5 — the null-return hypothesis. A return that does not arrive discriminates between two states of the channel — too distant for the listening window, or absorbing — and the discrimination improves with the length of the window. *Predicted:* silences classified provisionally as "far" resolve into arrivals at a rate that rises with window length, while those classified "absorbed" do not. *Falsified by:* silences resolving at rates independent of window length. *Status:* testable against any archive of dated correspondence; **not tested here.**

D1 — the performative (definition). An utterance is performative with respect to a speaker when the utterance is itself the act it names rather than a report about the world. This is a stipulated distinction, taken from the philosophy of language and used here without adjudicating that literature.

D2 — the reading rule (definition). For a signal x sent into a channel returning $y = T(x)$, the reading is $y - x$. This is a definition of what shall be read and not a claim that the difference is informative, which is H4.

D3 — category violation (definition). An object or act that harms no person — nothing taken, nobody deceived about anything they were owed, no will overridden — while standing in two categories that a classification scheme declared disjoint. This is a stipulated definition and its two halves are held apart deliberately.

C1 — the three constraints (stipulated constraints, not claims). Acquisition without coercion; possession by clean holding; and decoding by the holder. The framework states that failing any one converts the object into an ordinary harm. These are constraints on a doctrine's domain and are not falsifiable, and they are labelled as constraints rather than dressed as findings.

P1 — the inherited postulate. The uncertainty of a future affected by its own foreseeing is postulated to be the same *kind* as that of a quantum state, with the trade-off structure carrying over in shape. This paper inherits the postulate from earlier in the series, does not re-argue it, and preserves its badge at every occurrence. Nothing in physics licenses it.

4.0.3 3. Material carried across the opening boundary

The file assigned to this paper opens with the operating protocol of the preceding construct and its failure catalogue. Both are reported here because the preceding paper in this series records that they fall in this file, and because the protocol is the only place in the volume where the framework's resolution trade-off is given an actual measurement procedure.

3.1 The three-card carriage procedure

The procedure is a within-subject comparison and its design is better than most of what surrounds it. One undertaking is written on three cards at three resolutions — **fine** (dates, numbers, names, everything specified), **coarse** (a direction and a quality only, no numbers, no dates, no proper nouns), and **middle** (exactly one number *or* one date, and exactly one place *or* one name). Each card is then carried for one full day, and **every reading of the card, and every check made in the world because of it, is recorded as a tally on the card's back.**

That yields a measured quantity: the *disturbance cost* of a specification at three resolutions, in units of the practitioner's own attention. The procedure then sorts each card's tallies into two classes — checks that bent the hands toward *doing* something, and checks that bent them toward *verifying* — and applies a selection rule: keep the sharpest card whose disturbance is mostly of the first kind, and coarsen any card whose disturbance is mostly of the second.

This is a real measurement design and it is worth saying so. It is within-subject, it uses a fixed and countable operationalisation of the quantity it cares about, and its selection rule

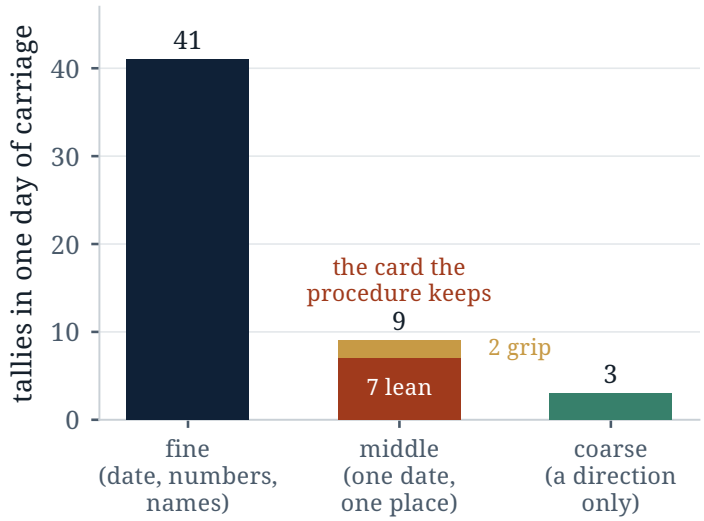


Figure 4.1: **The disturbance cost of a specification, measured at three resolutions.** Measured from the source's worked case, which records 41 tallies for the fine card, 3 for the coarse and 9 for the middle in one day of carriage each, and sorts the middle card's nine into seven that bent the hands toward doing and two toward verifying. The procedure's stated rule then selects the middle card. Single participant, unblinded, uncounterbalanced. Script: `figs/src/CCP-A04_fig01.py`.

is stated in advance. Its weaknesses are equally plain: a single participant, no blinding, an order that is not counterbalanced — the fine card is always carried first — and an outcome recorded by the person whose behaviour is being measured. The order effect is not incidental: a first day of carriage is also a first day of novelty, and the design confounds resolution with day number. **Counterbalancing the three days would cost nothing and was not done.**

The worked case records the numbers. The fine card returned **41** tallies in one day, the coarse card **3**, and the middle card **9**, of which **7** were of the hands-toward-doing kind and **2** of the verifying kind. The rule therefore selected the middle card. The case also records a self-observation that the framework treats as confirming and that should be treated more cautiously: a symptom polled five times daily began, on the practitioner's own account, to report at the polling hours. That is a single self-reported observation by an interested party and it is exactly the pattern that measurement reactivity would predict and that expectation alone would also predict; it discriminates nothing.

The procedure's standing element is a scheduled measurement: one reading of the retained card each day at a fixed hour, entered as exactly one of three words — nearer, farther, or unknown — after which the record is closed. The framework's argument for scheduling rather than abstaining is the one this corpus can fully support: an indicator adopted for consequential decisions is gamed, and the process it tracked is corrupted [Campbell, 1979, Goodhart, 1975], so the protection is a budget rather than a prohibition.

3.2 The five failure modes carried with it

Renamed with a `CC-` prefix to keep them clear of this paper's own results.

No.	Failure	Mechanism named by the source	Support in this bibliography
CC-U1	The borrowed authority	The practitioner conflates the uncertainty theorem, the observer effect and the framework's own conjecture, and claims physics endorses the last	**the counter-sign is the paper's own**: three separate sentences on demand. Morris and Thorne [1988] is the documented case of the failure
CC-U2	The indicator that replaces the substance	Goodhart and Campbell effects on the practitioner's own record	Campbell [1979], Goodhart [1975]
CC-U3	Continuous self-polling	Measurement reactivity in overdose; a system polled continuously has no unpolled hours left	**none** for the reactivity claim; McCambridge et al. [2014] bounds the nearest literature
CC-U4	The demand for a resolution the distance cannot support	Treating a forecast as a fact	**none** for the clinical framing; Salganik et al. [2020] and Hofman et al. [2017] establish that the underlying predictability is low
CC-U5	Never testing at all, dressed as epistemic humility	Avoidance wearing the chapter's own doctrine	**none**; the counter-sign is the source's own second rider, to spend the checking budget on the path rather than the outcome

CC-U1 is worth marking as the counter-result of the preceding paper and of this one: the framework's own named failure mode is the misuse of its own borrowed physics, and the test it sets — state the theorem, the observer effect and the conjecture as three separate sentences — is a discipline this paper applies to itself in Section 7.1.

4.0.4 4. The spoken commitment

4.1 Performative rather than descriptive

The material opens with a distinction that does real work: some utterances describe a state of affairs and can be true or false about it, while others *are* the act they name. A statement about the thickness of ice is a report and the ice is indifferent to it. A sentence of resignation, of naming, of promising, does not report the resignation, the name or the promise; it performs it, and courts and registries run on that fact daily.

The framework's move is to apply the distinction reflexively: **a spoken commitment is a performative aimed at the speaker.** It does not report a future — reporting futures is the scoring material of the preceding paper. It performs a selection, in the speaker's own accounting, which is the accounting the speaker's subsequent behaviour reads from.

This corpus's external bibliography contains no entry for the philosophy of performatives. The nearest support is the performativity literature in economics, which argues that a body

of theory does not describe its object but formats it [Callon, 1998], and which supplies the documented instance in which an option-pricing model's prices did not fit the market when published and the market came to fit the model after traders adopted it, with dates and price data recording the convergence [MacKenzie and Millo, 2003, MacKenzie, 2006]. That literature supports the general claim that an utterance can constitute rather than describe; it does not support the specific claim about a self-addressed commitment, and no citation has been attached to make it appear that it does.

The framework also borrows the vocabulary of measurement collapse for this construct, and it states its own limits carefully enough that the statement is worth preserving. What it takes is the uncontroversial spine — a distribution over outcomes, then an act, then one outcome, then all subsequent predictions conditioning on that outcome — and it declines to adjudicate the interpretive question, which remains genuinely open [Zeh, 1970]. **It borrows the shape and not the mechanism**, and this paper preserves that badge. The human mechanism, the framework then says, is not mysterious at all; it is the next subsection, and the next subsection is where the trouble is.

4.2 Three converging mechanisms, and their support (R1)

The framework offers three mechanisms and holds that they converge.

Consistency. Once a position is voiced, subsequent perception, memory and choice bend toward agreement with it, and a life at odds with a spoken sentence is uncomfortable in a way that has a name in the psychological literature. The framework's contribution is stated as an inversion: the same mechanism that produces rationalisation and doubling down can be aimed deliberately, so that the sentence spoken is one whose consistency pressure points where the practitioner has already, by the preceding procedures, decided to go.

Self-perception. A person infers who they are partly by observing their own conduct, including their own speech, so a sentence spoken aloud is filed as testimony where a sentence merely thought is filed as rehearsal. The framework's operational consequence is that the sentence must be voiced with the body in it and not held internally.

The commitment device. A voiced commitment re-prices later defection: every course inconsistent with the spoken sentence now carries an additional cost, small and constant, and a small constant tilt on a landscape compounds over time. This is the framework's join to its own gradient material.

R1 is the conjunction of the three. **All three are uncited.** This corpus's external bibliography contains no entry for cognitive dissonance, none for self-perception theory, and none for commitment devices. The three are retained in the text because R1 is what the corpus cites and removing them would misrepresent the framework, and they are retained **uncited**; they are listed individually in Section 11.3.

What can be supported is narrower and is stated instead, in two parts. First, that a belief acted upon can bring about the conditions that make it true is established as a sociological phenomenon and was named long before this framework [Merton, 1948], with its companion — purposive action producing consequences nobody modelled — the same machinery read from the other end [Merton, 1936]. Second, that expectations held about a person alter what that person produces is established, and is established at a magnitude smaller than the popular account holds: real but modest, not typically cumulative, and often dissipating, with most expectations turning out to be accurate rather than causal [Rosenthal and Jacobson, 1968, Rosenthal and Fode, 1963, Jussim and Harber, 2005]. **Neither establishes the self-addressed case**, which is R1's actual content.

4.3 The present tense as calibration (R2)

R2 is the framework's sharpest small claim and it is a claim about grammar with a mechanism attached.

A future-tense commitment — *I will become* — is, on the day it is spoken, **perfectly consistent with a day identical to yesterday.** The due date is always downstream; the account is always payable later; and the consistency mechanism, which is the whole motor of R1, therefore has nothing to act on, because an unchanged life does not contradict a promise whose term has not expired. A present-tense commitment asserts something about now, and to the exact degree that now does not match it, the pressure the mechanism supplies engages.

That is a genuine argument and it is stated here at full strength. Its dependency is also stated: **R2 is an entailment of the consistency mechanism, and the consistency mechanism is uncited.** If consistency pressure does not operate as described, R2 is a stylistic preference. This paper does not soften R2; it locates its load-bearing premise.

The framework attaches one boundary in strong terms and the boundary is worth preserving, because it forecloses the abuse the grammar invites. The present tense must be **reachable by gradient** — a state the practitioner can begin occupying in some honest fraction today. A sentence flatly contradicted by observable fact each morning is not a commitment but a misdescription of the present state, and the framework's ruling is that the ground state always wins and that each collision costs the instrument credit until the practitioner's own testimony is filed as fiction. The framework's formulation is that the sentence declares a road already entered rather than an arrival already occupied.

4.4 The prohibition (R3)

The framework states its constraint as law and states it before the technique that would violate it. This paper preserves that order and separates the two arguments the framework gives, because they have very different standing.

> **No future is to be voiced over another person's life without their word — and the prohibition covers the benevolent case explicitly.**

The ethical argument requires no literature and is stated first. Another person is not a branch of the speaker's situation; they are a separate system with their own weights and their own jurisdiction. A future installed in another person's environment, below the level at which they can contest it, is coercion whatever its intent, because the channel it operates in is one where the target cannot defend themselves without becoming someone who distrusts every encouragement. That argument is a moral one and it stands or falls on moral grounds; it is not an empirical claim and it is not offered as one.

The mechanical argument is the framework's own and is uncited. The framework holds that an imposed future is registered as pressure and produces a counterforce proportional to the detected push, so that the attempt manufactures its own refutation and returns it to the speaker as apparent evidence against their reading. **This corpus's external bibliography contains no entry for reactance**, and the claim is retained uncited, exactly as the other papers of this series retain it. What can be supported is weaker and sufficient for the corpus's purposes: that classifications and expectations applied to people alter the people they are applied to, so that the referent moves [Hacking, 1995]; and that a theory adopted as a description of people can make itself true by changing how they are treated, sometimes for the worse [Ferraro et al., 2005].

R3 is therefore a stipulated constraint with a sound ethical argument and an uncited mechanical one. The constraint stands without the mechanism. The framework's stronger claim — that the prohibition enforces itself, so that a practitioner need not be trusted to observe it — does not.

The framework's own statement of what remains permitted is narrow and is preserved: a practitioner may place a visible offer where another person will encounter it, may ask, and may consent to be asked. The corpus states the same constraint elsewhere as a closure condition rather than a courtesy [Curzi, 2026].

4.5 The read-aloud test and its taxonomy

The framework's drafting procedure is a rejection loop: draft, read aloud alone, and discard any draft that produces a specific bodily response, repeating until one survives. It supplies a four-way taxonomy of the responses, and the taxonomy is the useful part because it converts a vague instruction into named observations:

- **the catch** — the voice thins on a particular word, which the framework holds identifies the word carrying the falsehood rather than the sentence;
- **the slide** — the gaze leaves the page and the line is read without being heard, which the framework holds indicates a borrowed sentence;
- **the hedge** — a qualifier appears in the final third, which the framework treats as a flinch that arrived early enough to reach the text;
- **the performance heat** — the sentence reads clean and yet is read with an awareness of how it sounds, which the framework holds indicates an audience inside the sentence and treats as the most serious of the four.

The mechanism the framework offers — that bodily verdicts on prospects precede articulate reasoning — is named in the source with an explicit acknowledgement that the hypothesis it belongs to has open questions, and this corpus's bibliography contains no entry for it. It is reported uncited. The taxonomy does not depend on the mechanism: it is a set of named, self-observable events, and its usefulness is that it is more specific than "trust your instinct" and can therefore be wrong in identifiable ways.

The framework distinguishes what is discarded from what is filed, and the distinction is operational rather than ceremonial. A specification declined *at this resolution* is retained, because the distance may close and the finer version may later be usable; a draft that was never the practitioner's own sentence is destroyed. **The reason the framework gives for destruction — that unfinished items hold attention and levy a standing cost — has no supporting entry in this corpus's bibliography** and is reported uncited. The corpus's closure discipline elsewhere rests the same rule on bookkeeping grounds instead, which is the defensible version [Curzi, 2026].

4.6 The compression doctrine, and its boundary

The framework advances a specific claim about laughter that is worth reporting because it is stated as an operational tool with a stated boundary and a stated test.

Its analysis: a joke sustains two readings of the same material, the second concealed; the resolution forces a collision that reveals the concealed reading as the true one; and the moment of forced recomputation has an involuntary felt signature. The framework's inference is that laughter marks a **compression event** — a contradiction registered as meaningful without being resolved, releasing the process to proceed — rather than a resolution of the contradiction.

The operational rule follows: where a procedure jams on a contradiction that ordinary means cannot resolve, the practitioner may deliberately laugh at the knot, which the framework holds lets the work continue through it rather than around it.

Two things must be said. First, **this corpus's external bibliography contains no entry for the psychology or the theory of humour**, and the incongruity-resolution family the source names is reported uncited. Second, and more to this paper's credit than the framework's, **the framework attaches a boundary and a test that make the rule falsifiable in use**: where the contradiction is the practitioner's own failure, the rule does not apply and the appropriate response is repair; and the test of whether a given instance was compression or evasion is whether the *work* proceeded within the hour or only the mood improved. A rule with a stated failure signature is a better object than a rule without one, and this one has both.

4.7 Speaking without an audience

The framework's most counter-intuitive instruction is that the commitment be spoken aloud and to no one, and it argues the point against the common advice rather than ignoring it.

It concedes the case for public commitment: a word given before others is harder to withdraw, and the architecture of vows and contracts exists because that is true. It then asks what a public commitment binds the speaker to, and answers that when the commitment concerns identity — who the speaker is becoming, which is what a commitment of this kind is — announcing it to an audience can **substitute** for enacting it. The audience supplies recognition; recognition is the wage of the finished thing, paid early; the tension the sentence was built to create is partially discharged in the telling.

The framework's reconciliation is the sharpest sentence in the material and this paper adopts it: **the voicing is the measurement; being observed voicing is the contamination.** From it follow three placement rules — speak once and not repeatedly, speak to something that cannot respond with recognition, and observe a quarantine before telling anyone.

The substitution effect has no supporting entry in this corpus's bibliography and is reported uncited. What can be supported is adjacent and weaker: that an indicator adopted as a target is gamed and corrupts the process it tracked [Campbell, 1979, Goodhart, 1975], which applies to a practitioner's own record of approach; and that a record made and dated in advance constrains what can later be claimed to have been intended [Nosek et al., 2018], which is the defensible warrant for the quarantine — it keeps the record uncontaminated by an audience's response rather than protecting a psychological account.

4.8 The failure catalogue of the spoken commitment

No.	Failure	Mechanism named by the source	Support in this bibliography
CC-C1	The commitment spent on an audience	Announced identity goals substituting for enactment	**none**
CC-C2	The commitment that waits	Deferral encoded in the future tense; the practitioner does not register the tense as a flinch	**none**; the remedy is mechanical — strike every future-tense verb and read what remains
CC-C3	The specification closed too small	Over-specification rebuilding the corridor of the preceding construct in a single sentence	Campbell [1979] supports the indicator half; the foreclosure half is the framework's own
CC-C4	The commitment laid on another life	Imposition and the counterforce it purchases	**none** (Section 11.3). The ethical argument of Section 4.4 stands without it
CC-C5	The laugh that disperses instead of compressing	Avoidant humour wearing the compression tool's coat	**none**; the source's own test — did the work proceed within the hour — is the discriminator

4.0.5 5. The returned difference

5.1 The loop, written out

The framework writes the loop it inherits from the founding paper of this series in operational form:

$$T(x) \;=\; W(x,\, M(x))$$

where M is the practitioner's model of the situation, W is the world's response to an action taken on that model, and $T(x)$ is the next state. The founding paper fixes this composition and identifies a self-fulfilling prophecy as a fixed point of it [Curzi, 2026]. **R4 and R5 derive from that definition**, and inherit its status: Section 7.3 records that the specification is not identified from observational data and that the corpus's claims about it are modelling positions rather than measured effects.

The framework's examples are ordinary and are the right ones: a rumour that an institution is failing can empty it, after which the institution has failed and the rumour stands vindicated at the outcome it produced; a selector who expects little asks smaller questions, receives smaller answers, and files the smallness as a discovery. The general principle — that a system's next state can depend on a participant's report about the system — is the working principle of every controller and is not this framework's discovery [Wiener, 1948, Åström and Murray, 2008], and the modelling of a human operator inside such a loop is older than any claim here [Tustin, 1947].

5.2 Why the difference carries the information (R4)

R4 is the framework's reading rule and it is the strongest result in this paper, because it is true by construction in the domain it is borrowed from.

$$\text{sent } x \quad\longrightarrow\quad \text{returned } y = T(x) \quad\longrightarrow\quad \text{read } y - x$$

The argument is a reductio: consider what a *perfect* return would mean. If a signal comes back unworn — undelayed, unfiltered, unchanged — the observer has learned exactly one thing, namely that nothing in the channel could touch the signal. **A channel that returns a signal unchanged has reported only its own emptiness.** All of the information about the channel lives in the transform: in what was kept, what was lost, what was delayed, what was added.

In signal terms this is not an analogy but a theorem. The response of a linear time-invariant channel to a sharp input *is* the channel's impulse response, and the impulse response characterises the channel completely [Oppenheim et al., 1996]. Every echo-ranging and impedance-probing method in engineering is built on exactly that fact, and the sent signal is treated as known, cheap and disposable precisely because the returned signal is valuable only where it deviates. The information-theoretic form of "read the difference" is equally standard: the quantity that measures how far one distribution has been moved from another is a divergence rather than a property of either one alone [Kullback and Leibler, 1951, Shannon, 1948].

The transfer to a social channel is a shared mathematical form and not a derivation, and it is labelled as one in Section 7.1. The premises that make the engineering result exact — linearity and time-invariance — are not established for a channel made of markets, institutions or people, and are implausible for one made of people. What survives the transfer is the *reading discipline*: record what was sent, record what came back, and attend to the deviation rather than to the arrival. That discipline is defensible on its own terms, because it converts a verdict into an observation, and this paper endorses it in that reduced form.

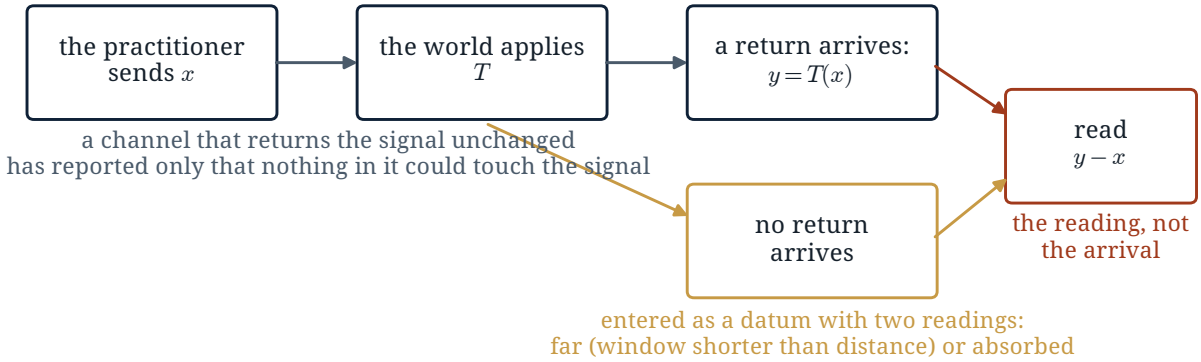


Figure 4.2: **The reading rule, and the null return as a datum.** Schematic; no plotted quantity. The loop and the rule are the source's own, which records what is sent, records what returns, and reads the difference. The lower branch is the source's treatment of a return that never arrives: it is entered with exactly two readings rather than discarded as an absence of evidence. Script: `figs/src/CCP-A04_fig02.py`.

5.3 The null return (R5)

R5 is a small result and a genuinely useful one. A return that does not arrive is entered in the record with one of two marks:

- **far** — the listening window may simply be shorter than the distance;
- **absorbed** — the medium does not return this frequency, and a louder call will not change that.

Where the two cannot be distinguished, both are written and the entry stands. The framework's own instruction is that waiting is itself a reading.

This is the correct treatment of a missing observation and it should be said plainly. An absent return is a datum about the channel; treating it as no evidence discards information, and treating it as a verdict fabricates information. The framework's rule is neither, and its provision for an undecided entry is what makes it honest: a record that forces every silence into one of two bins would manufacture certainty at the point where there is least.

The corpus supplies the general version of the same discipline. A record that asserts an identity holds exactly while the statement it certifies is misleading is a consistency check and not an integrity check, and the identity's holding is therefore not evidence [Curzi, 2026]. A null return that has been classified for the analyst's convenience is the same failure in a different ledger.

5.4 When the other party is a will

The framework marks this case as the seam in its own material and does not close it, and this paper reports the concession as a concession because it is the counter-result.

When the responding party is a person, the diagram doubles: the other party is not a channel but a second system with its own model and its own outward signal, so what returns was composed on the far side of an inwardness the sender cannot audit. The framework draws two consequences. The informational one: a reply from a person carries more than an echo from a material, because it is the transform of a transformer. The ethical one: a material cannot be trespassed against by being spoken at, and a person can.

The framework then states the case it cannot resolve. Its own argument runs: a person's discarded drafts — the sentence struck, the tone not adopted — are read off them by everyone

they pass, so the world already answers versions of a person they never authorised; from which it might follow that carrying one more copy of an already-circulating self is transport rather than forgery. **The framework records the argument, rules against it, and states that the ruling does not settle it.** Its stated ground for ruling against it is a distinction between a condition and a choice: that the world reads a person without leave is a condition, while a particular person appointing themselves the world's courier is a choice.

This paper reports the seam as open. The framework's own text says the argument is not settled and will not be settled by it, and no result in this corpus closes it. What the framework does supply, and what is worth preserving, is a two-ledger rule for the case in which a return was obtained by trespass: the trespass is entered with its name and its cost, *and* the return is read, on the ground that the world's answer is now a fact about the world regardless of how it travelled, and that discarding a sounding to punish the sounder makes the situation less legible rather than more just. Neither entry pays the other.

The corpus's own treatment of the point where a second party enters a loop is the same one, stated as governance: when another will consents to be in the loop, the loop is a commons, and durable arrangements over commons are governed by rules the affected parties had a hand in making — a finding derived from cases rather than postulated [Ostrom, 1990].

5.5 What the framework establishes and what it claims

The framework performs its own accounting here and the accounting is accurate on its own terms, which is rare enough to preserve. It sorts its claims into three:

- **Established elsewhere, not by this framework:** feedback — that a system's next state can depend on a participant's report about it [Wiener, 1948, Åström and Murray, 2008].
- **Established elsewhere, in a reduced form:** that expectation alters conduct and conduct alters response. The framework's own summary is that the effect is real, moves through ordinary channels, and is contested in size, and that anyone claiming it is either nothing or everything is selling. That summary is accurate [Rosenthal and Jacobson, 1968, Jussim and Harber, 2005].
- **Established, on the framework's account:** reactance. **This is the one item the framework mis-sorts.** It is not established in this corpus's bibliography, which contains no entry for it, and this paper moves it to the uncited list.
- **Claimed, and flagged by the framework as a craft claim:** that the wear on a return can be read, with practice, as a chart of the material between the sender and the mark. The framework compares it to a tracker's claim about a bent stalk and says it is practised everywhere and proven nowhere. **That is the correct badge and this paper keeps it.**
- **Postulated:** the inherited postulate P1, which the framework declines to re-argue.

The framework then makes a design statement worth quoting as a principle: it chose instruments that **survive the doctrine failing**. The two-column record and the difference reading remain useful if the postulate is false, because they are bookkeeping rather than theory. That is the correct criterion for an instrument and it is the strongest methodological sentence in this material.

5.6 The recording procedure and its failure catalogue

The procedure is a fixed-window observational record and its design deserves the same scrutiny the framework gives its subject. A window is declared in advance; the channels that will count as observations are declared in one sitting, and **nothing may later be promoted to an observation that was not declared**; arrivals are logged verbatim on the day they arrive, without interpretation at entry; the difference is taken on a fixed weekly schedule; silences are classified per R5; and no return that raised the practitioner's pulse may be answered within one day.

SENT	RETURNED	THE WEAR
letter to a fine-editions press	silence	far or absorbed: both marked
letter to a second press	"list us your commercial rates"	fine → commercial
shop-window card: repairs and fine bindings	a co-op asks for 40 years of ledgers rebound	fine → records
price list to a historical society	"can you teach the apprentice to sew spines?"	fine → keeping

the page's own closing line: the material returns the frequency "keeping" and absorbs "exquisite"

Figure 4.3: **A worked record, with the deviation extracted.** Measured from the source's own reproduced card: four items sent, three returns and one silence. The right-hand column is what the reading rule extracts in each case, and the silence is entered under both of its two permitted readings rather than resolved. Script: `figs/src/CCP-A04_fig03.py`.

Three of those are genuine methodological controls and should be named as such. Declaring the channels in advance is a pre-specification rule that prevents the outcome set from expanding to fit the result [Nosek et al., 2018, Simmons et al., 2011]. Logging verbatim before interpretation separates observation from analysis. And the one-day delay on replies is a control on the analyst becoming their own noise source — in the framework's terms, a reply that outruns the reading collapses the difference before it has been measured. The general fact that delayed feedback is where human judgement systematically fails is established [Sterman, 1989, Forrester, 1961].

The procedure also imposes an ordering rule that is a genuine ethical commitment rather than a technique: where a return contains another person's actual question, that question is answered before the practitioner's own work advances.

The worked case is a single, self-reported, retrospective record and establishes nothing about whether the procedure works; it is reported because it shows what the procedure returns and because it contains two refusals, which is the part of a practice usually omitted from its own account. Its own closing line — that the surrounding material returns one frequency and absorbs another — is an interpretation by the interested party, and the base rate against which any single such report should be read is not encouraging [Ioannidis, 2005, Open Science Collaboration, 2015].

The six named failure modes, renamed with the `CC-` prefix:

No.	Failure	Mechanism named by the source	Support in this bibliography
CC-E1	The arrival read as a verdict	Outcome fixation; the return graded rather than differenced	**none**; the remedy is the forced two-column subtraction
CC-E2	The hall of one voice	Listening arranged so that only agreement can return	Lorenz et al. [2011], Bikhchandani et al. [1992] establish that correlated sources stop carrying information

No.	Failure	Mechanism named by the source	Support in this bibliography
CC-E3	The reply that outran the reading	Impulsive response; the practitioner becomes their own noise source	Sterman [1989] on delayed feedback and misjudgement
CC-E4	The transform taken personally	Attributing the medium's transform to malice or to self	**none**
CC-E5	The forged return	Promoting an undeclared channel to an observation after the fact	Simmons et al. [2011], Gelman and Loken [2013] on the flexibility this exploits
CC-E6	Speaking through another's life	Imposition, and the response it purchases	**none** for the mechanism; the two-ledger rule of Section 5.4 is the framework's own

CC-E2 is worth marking because the corpus prices it exactly elsewhere: agreement between sources that share their premises is not evidence but repetition, and a set of aligned viewpoints has an effective count of one however many members it has. CC-E5 is the failure the framework guards against by pre-declaration, and the guard is placed upstream of the temptation, which is the correct place for it.

4.0.6 6. Category violation

6.1 A classification is a projection

The framework's opening move is a claim about filing systems: a scheme has fewer dimensions than what it files. A catalogue is a sequence; a form is a finite grid of yes-or-no boxes; a taxonomy is a tree, and a tree forbids the one thing living systems do constantly, which is for two branches to touch.

The framework states the consequence as a geometry. Filing is a **projection** of a higher-dimensional object onto a surface low enough to administer, and a projection is not a lie but it does discard. Most objects project into the interior of a single box and file without loss; the framework's estimate is that this is the ordinary case and that the scheme is therefore not a defect but the reason schemes exist. Some objects, however, carry a dimension the scheme does not have, and their projection lands on the boundary between two boxes and stays there.

This is not the framework's discovery and the literature it belongs to is in this corpus. Objects that inhabit several communities of practice at once, satisfying the informational requirements of each while being differently constituted in each, are **boundary objects**, and they are the standard analytic device for exactly this situation [Star and Griesemer, 1989]; the general study of what classification systems do to what they classify, and of what falls between their categories, is a developed field [Star and Ruhleder, 1996]. The framework's projection geometry adds an image to that literature and does not add a result to it, and this paper says so rather than presenting the geometry as new.

Two further external results bound the material. A scheme that makes a situation legible in order to administer it can destroy the local variety the situation was actually running on [Scott, 1998] — which is the cost side of the projection, and the framework does not price it. And a classification applied to people alters the people classified, so that the referent moves and the scheme's own accuracy decays [Hacking, 1995] — which the framework does not mention at all, and which is the single most relevant external result to a doctrine about the gaps in classification schemes.

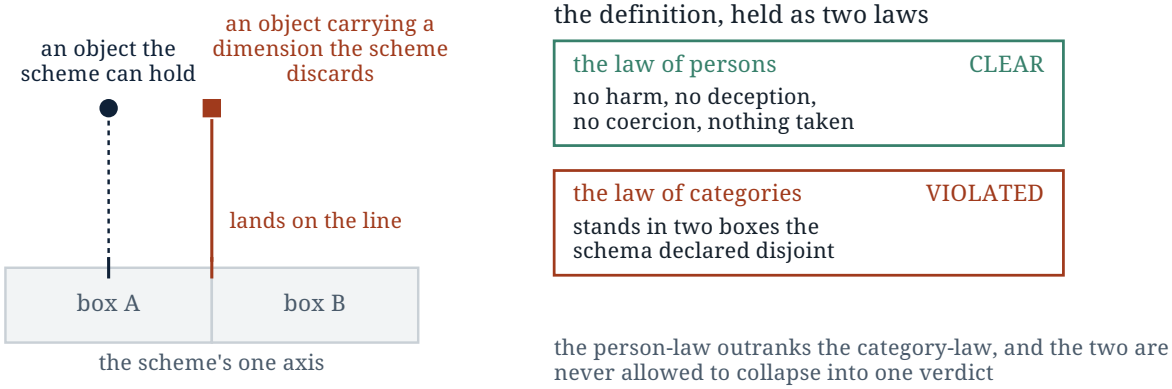


Figure 4.4: **A category is a projection, and a category violation is an object whose projection lands on a boundary.** Schematic; no plotted quantity. The left panel is the source's own construction: filing discards a dimension, most objects project into the interior of one box, and an object carrying a dimension the scheme lacks can land on the line between two boxes the scheme declared disjoint. The right panel is the definition, with the two laws the source insists on holding apart. Script: `figs/src/CCP-A04_fig04.py`.

The framework also applies its inherited postulate here, and this paper flags the application. It describes an object exceeding its category as being "in superposition" with respect to that category, and it flags the description as a postulated correspondence of shape rather than a claim about substrate. **The flag is correct and the correspondence is doing no work**: nothing in R6 or R7 requires it, and the verification chain in Section 10.2 does not consume it. It is reported once and used for nothing.

6.2 The definition, and the two laws (R6)

R6 is a definition and it is stated as a deliberate paradox held open.

> **A category violation is an object or act that is, at the same instant, not a crime and a crime.**

It is **not a crime by the law of persons**: no harm, no deception about anything owed, no coercion, nothing taken from anyone who had it. Run the tests by which a wrong is identified in law or in conscience and each returns no.

It is **a crime by the law of categories**: it stands in two boxes a scheme declared disjoint, and in doing so it falsifies the scheme's implicit promise that its boxes tile the world without overlap or gap.

The framework's instruction is that **neither half may be allowed to dissolve the other**, and that the reader's temptation at this point is to collapse the paradox by deciding the object is really innocent (so the word "crime" is a provocation) or really guilty (so the innocence is an excuse). Both collapses are refused.

The framework then ranks the two laws, and the ranking is the result's actual content: **the law that protects persons is the law the other law exists to serve.** A classification scheme is a tool for coordination — so that things can be found, remembered and handed on — and where the two laws part company, the person-law governs and the category-law is revealed as a convenience rather than a moral fact.

This ranking is a normative claim and it is not established by anything. It is a position, and a defensible one, and it is labelled here as a position. What can be said in its favour from this corpus's own material is that a scheme is a measurement instrument and that an instrument's

classifications license only the operations its scale type permits [Stevens, 1946, Krantz et al., 1971]; an object exceeding a scheme's dimensions cannot be filed without loss, and treating the resulting file as a fact about the object is precisely the error the measurement literature names [Michell, 1997].

6.3 The diagnostic use, and what it cannot do

The framework's second claim about category violations is that they are diagnostic: the capacity to recognise one and the capacity to construct one are the same capacity, namely the ability to read along a dimension a given scheme discards. It draws the consequence that an object of this kind sorts observers without asking them anything, since an observer who reads only in the scheme's dimensions will file it correctly and be untroubled.

Two objections apply and are stated here rather than left to a reviewer.

First, the claim that recognition cannot be faked — the framework's phrasing is that one cannot cheat on a test one is unable to perceive — **is true only of a first encounter.** Any such object, once described, becomes an ordinary piece of information that can be learned and reported, and the framework's own third constraint (Section 6.4) concedes exactly this by declaring an explained instance spent. The diagnostic is therefore single-use per object and per observer, which sharply limits its practical value and which the framework does not state.

Second, the claim generalises a specialist's inability to see the perpendicular dimension into a property of specialisation as such. The organisational literature supports a weaker version — that an organisation which learns fast homogenises and destroys the variance that produced its own advantage [March, 1991] — and a stronger version is not established. The framework's account of specialisation is a restatement of its own gradient material and inherits its status.

A third point belongs here and it is the one that most limits the material's reach. **A selection device that identifies people by an unadvertised perceptual capacity is unauditable by construction**: it has no stated criterion, no rate of false identification, and no appeal. The engineering standard for a system whose function must be trusted is the opposite of what this device requires — a system's integrity should not depend on any part of it being hidden [Kerckhoffs, 1883, Saltzer and Schroeder, 1975]. That standard is cited here **against** the framework, because the framework's third constraint makes concealment constitutive rather than incidental, and a reader should see the tension named.

6.4 The three constraints (R7)

R7 is the framework's own fence and it is the counter-result of this paper, because it refuses more than the definition licenses.

First: no coercive acquisition. The object may violate categories; it may not violate persons, and acquisition is the first place a person can be violated. The framework's formulation is exact and is preserved: a category violation that was taken — from a hand that did not release it, a place that was not opened, a will that was overridden — is not a category violation that happens to have been stolen. It is a theft, entire, and the doctrine laid over it is an aggravation rather than a mitigation, because it is the theft plus the claim that the theft was clever. **Category innocence is explicitly refused as a defence of acquisition.**

Second: possession by clean holding only. The positive half of the first. Where an object came to the holder cleanly, the fact that a second reading exists inside the same catalogue identity deprives no one of anything they had, were owed, or could name. The framework's supporting observation is the one it flagged three sections earlier and then uses: the second reading is superposed on the first and takes nothing from it.

Third: self-decoding. The meaning is not handed over, and the framework states this as a property of the object rather than a rule imposed on the holder: an instance explained by a fluent reader is spent, because the function of the object is to be readable only by an eye

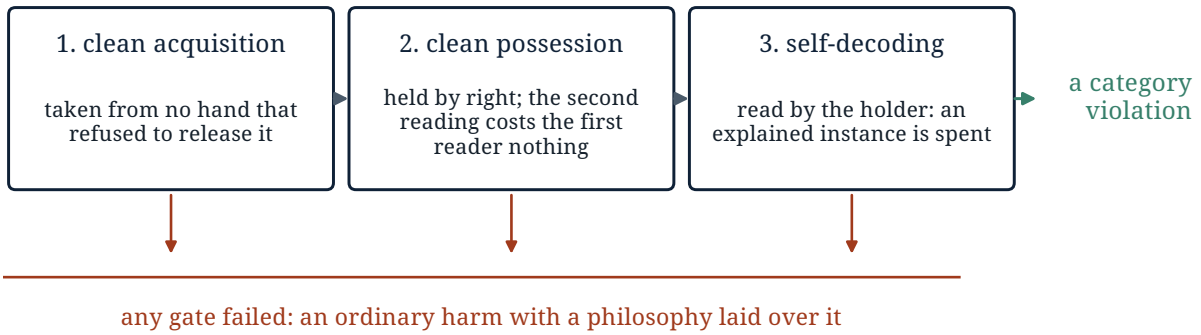


Figure 4.5: **The three constraints, as a gate sequence.** Schematic; no plotted quantity. The three conditions and their order are the source's own, as is the terminal state: failing any one converts the object from a category violation into an ordinary harm. The third gate is stated by the source as a property of the object rather than a rule imposed on the holder. Script: `figs/src/CCP-A04_fig05.py`.

reading in the discarded dimension, and a borrowed reading leaves the borrower exactly as they were.

The framework's summary of its own doctrine is that failing any one of the three converts the object into an ordinary harm with a philosophy laid over it, and that the doctrine therefore forbids moving any harm through a discovered gap in a scheme. That is the correct reading and this paper adopts it: R7 constrains R6 rather than extending it.

The framework closes by identifying its third constraint with the consent doctrine that governs the rest of the volume: a recognition cannot be installed in another person, and the attempt is coercion, which produces resistance rather than movement. That identification is internally consistent and it inherits the uncited status of the mechanism it rests on (Section 4.4).

4.0.7 7. Method: what kind of borrowing, what scale, and what is identified

7.1 Analogy discipline

Every borrowing in this paper is labelled at first use as a **derivation** (the system satisfies the borrowed result's premises), a **shared mathematical form** (the same equation governs both, with no claim that the source domain's physics applies), or an **analogy** (a suggestive correspondence with no formal warrant, legitimate as motivation and never as evidence). The labels are collected here.

Borrowing	Category	Warrant, and what is not claimed
Measurement collapse, for a spoken commitment (Sec. 4.1)	3 — analogy, and flagged as one by the source	Only the uncontroversial spine is taken — distribution, act, one outcome, conditioning — and the interpretive question is left open [Zeh, 1970]. No claim about substrate.

Borrowing	Category	Warrant, and what is not claimed
Superposition, for an object exceeding its category (Sec. 6.1)	3 — analogy, and doing no work	Flagged by the source as postulate. No result consumes it; the chain in Section 10.2 does not contain it.
The impulse response, for the reading rule (Sec. 5.2)	**1 for a linear channel, 2 for a social one**	Exact for a linear time-invariant system [Oppenheim et al., 1996]. Neither linearity nor time-invariance is established for a social channel; what transfers is the reading discipline, not the theorem.
Divergence, for "read the difference" (Sec. 5.2)	2 — shared form	A directed measure of how far one distribution sits from another [Kullback and Leibler, 1951]. No claim that the returns form a distribution or that any divergence has been computed.
Feedback and loop gain (Sec. 5.1)	2 — shared form	The mathematics of a controller governs any system with a report inside it [Wiener, 1948, Åström and Murray, 2008, Tustin, 1947]. No physical result is transferred.
Projection and dimension, for classification (Sec. 6.1)	3 — analogy, with a real literature beside it	The geometric image is the source's. The established treatment of the same phenomenon is boundary objects [Star and Griesemer, 1989] and the study of classification itself [Star and Ruhleder, 1996].

The general hazard is documented: a formal apparatus imported without the constraint that gave it content produces a defective copy [Mirowski, 1989]; laws hold within shielded arrangements rather than universally, so importing a law's form is not importing its scope [Cartwright, 1999]; behaviour at one scale is not derivable from the laws governing the parts [Anderson, 1972]; and the field this corpus's title claims membership of has been surveyed and bounded, with its vocabulary given definitions the corpus can be held to [Castellano et al., 2009]. The constructive standard for making such a borrowing checkable is to name the source-domain quantity, name the target-domain quantity, and state the correspondence [Vinković and Kirman, 2006]; the table above is this paper's attempt at it. **No conserved quantity is exhibited anywhere in this paper and none is claimed. No scaling or power-law claim is made**, so the procedure that would be binding on one is not run, and its absence here is a scope statement rather than an omission [Clauset et al., 2009].

The framework's own failure mode CC-U1 sets the test for this section: state the theorem, state the ordinary observer effect, and state the framework's conjecture, as three separate sentences. They are, respectively: that no state is sharp in both members of a conjugate pair,

watched or unwatched [Kennard, 1927]; that to measure a system one must interact with it, so the interaction enters the system's history; and that the uncertainty of a future affected by its own foreseeing is the same kind, which is P1 and is badged wherever it appears.

7.2 Scale type, and what a classification licenses

Two of this paper's results are about classification and one is about a rating, so the scale-type audit bears on all three.

R6's classification is nominal. A scheme that assigns objects to boxes without order supports statements of membership and non-membership and nothing else [Stevens, 1946]. The framework's central observation — that an object can satisfy the membership conditions of two boxes a scheme declared disjoint — is therefore a statement *about the scheme's own consistency* and requires no quantitative structure at all. **R6 is the one result in this paper whose measurement-theoretic standing is unproblematic**, and it is unproblematic precisely because it claims nothing numerical.

R4's difference is not licensed as arithmetic in general. Writing $y - x$ presupposes that the sent and returned signals live in a space where subtraction is defined. For a physical signal it does [Oppenheim et al., 1996]. For a returned letter, an offer or a silence it does not: the "difference" between a request for fine bindings and an offer of ledger repair is not a subtraction but a *comparison*, and no operation on it beyond comparison is licensed. **The reading rule is therefore retained as a comparison and demoted as an arithmetic**, which costs nothing, because the framework never computes with it either.

The tally counts of Section 3 are a ratio-scale count and are licensed. A count of recorded checks has a true zero and equal units by construction, so the comparison of 41 against 9 against 3 is legitimate arithmetic. What is *not* licensed is the inference from those counts to a general constant: three days by one participant, uncounterbalanced, with the outcome recorded by the participant.

The standing objection to the framework's wider programme applies here as everywhere: the strongest single criticism is that psychology assumes quantitative structure without testing for it [Michell, 1997], the formal test is conjoint measurement [Luce and Tukey, 1964] against the measurement-foundations standard [Krantz et al., 1971], and **it was not run.** Routes that would earn the arithmetic exist and are unused [Rasch, 1960]. This paper's exposure is smaller than its predecessors' only because most of its results are qualitative.

7.3 Identification, and the framework turned on itself

R4 and R5 derive from the reflexive loop, and the loop's status is fixed and inherited.

Written as an aggregate relation — a group's state influencing its members while the members constitute the group — $T(x) = W(x, M(x))$ is a linear-in-means reflexive specification, and **the endogenous social effect in such a specification is not identified from observational data**: it cannot be separated from contextual effects or from correlated unobservables, however much data is gathered [Manski, 1993]. The network route is closed by the same family of result, since influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011]. The available remedies are randomisation, an instrument, partial-population variation, or stipulating the reference group and defending the stipulation, the last having an explicit template in a discrete-choice model that states the conditions under which its interaction parameter is econometrically identified [Brock and Durlauf, 2001]. **This paper adopts none of them.** R4 and R5 are therefore **modelling positions rather than identified empirical results**, and every downstream paper citing them inherits that status.

The consequence is specific and constrains the framework's own procedure. The recording procedure of Section 5.6 produces a within-subject observational record of returns following an intervention the subject chose. **No causal reading of that record is available**, because

the subject's decision to send is not independent of the state of the channel; whatever association appears between a sent signal and a returned one is confounded by the reasons the signal was sent when it was. That is not a defect of the record-keeping, which is good; it is a limit on the inference, and it is stated here because the framework does not state it.

That reflexivity of this shape occurs at all is separately documented, and the documentation is the existence proof the framework otherwise lacks [MacKenzie and Millo, 2003], with the closest existing named doctrine supplying the concept without an equation [Soros, 2013]. That literature supplies existence, not identification.

The framework turned on itself. The framework asserts that a system changes its behaviour when modelled and observed. Applied consistently, that assertion falls on its own procedures: a person who has read this material and then keeps a two-column record is not the same observer as one who has not, because the decision rules the record captures are responses to the environment the model describes [Lucas, 1976]. The position taken here, stated so that it can be held against the paper: **no stability of this framework's effects under its own dissemination is claimed, and no result in this paper depends on any.** Three consequences are accepted. The tally counts of Section 3 are not transferable to a practitioner who has read the selection rule in advance. A before-and-after design on a single practitioner is confounded by construction. And R1's own mechanism, if it operates, operates differently on a speaker who knows why they are being asked to speak aloud — which makes H1's spoken-versus-silent comparison harder to run than it looks, and this paper says so rather than proposing it as though it were easy.

4.0.8 8. Results

8.1 Definitions and constraints fixed. D1 (the performative), D2 (the reading rule), D3 (category violation) and C1 (the three constraints) are stated in the form the remaining papers cite. None is a finding. C1 in particular is a constraint on a doctrine's domain and is not falsifiable, and it is labelled as such rather than presented as a result.

8.2 Results established by derivation. R4's core is established **for a linear channel**, where the transform is the channel's complete characterisation [Oppenheim et al., 1996]; the transfer to a social channel is a shared form and is not established. R5 is established as a bookkeeping rule: a missing observation is a datum with two readings, and the provision for an undecided entry is what makes it honest. R6 is established as a definition and its scale-type standing is unproblematic (Section 7.2).

8.3 Results established by construction, on uncited mechanisms. R1, R2 and R3 are the framework's specifications, and **all three rest on mechanisms this corpus cannot cite.** R1 rests on three of them; R2 is an entailment of the first of the three; R3's mechanical limb rests on a fourth. This is stated here as a result because it is the most important thing this paper found: the section of the framework most often quoted downstream is the section with the least support in the corpus's own bibliography.

8.4 Hypotheses tested. None. H1 through H5 are stated in falsifiable form and none is tested here. No dataset was collected. H4 is true by construction in the domain it is borrowed from and untested in the domain it is applied to. This is reported in the same plain language a positive result would receive: **the five hypotheses this paper states stand untested, and the paper that states them supplies no evidence for any of them.**

8.5 One measurement is reported, with its design flaws. The three-card carriage procedure returned 41, 9 and 3 tallies at fine, middle and coarse resolution, with the middle card's nine sorting 7 to 2. **That is a real measurement of a real quantity and it is worth nothing on its own**: one participant, three days, no counterbalancing, no blinding, and the outcome recorded by the person being measured. It is reported because the design would cost nothing to repair and because a corpus that reports a weak measurement honestly is in a better position than one that reports none.

8.6 Reproduction of the corpus's own arithmetic. Three numerical claims carried by this paper's internal citations were recomputed and all three reproduce.

- The vehicle-cost comparison: at 240 million units of value, 5.0 basis points is 120,000 and 35 basis points is 840,000, an annual difference of exactly 720,000 as printed [Curzi, 2026]. **Reproduced exactly.**
- The concentration chain printed in the corpus with its own arithmetic error left visible: $1/0.764992 = 1.307203$ and $\ln(1.307203) = 0.267890$, both as printed [Curzi, 2026]. **Reproduced exactly.** That paper's practice of printing an error and correcting it in place is the practice this paper follows in Section 8.5 and that the preceding paper in this series follows for a discrepancy it could not resolve.
- The audit distribution: seven mappings, two splits, one rendering and three rejections sum to thirteen entries as stated [Curzi, 2026]. **Reproduced.**

4.0.9 9. Discussion

9.1 What follows

Three consequences follow and they are this material's genuine yield.

The first is that **a reading rule can be right while its mechanism is unsupported.** R4's instruction — record the sent, record the returned, read the deviation — converts a verdict into an observation, and that conversion is defensible as bookkeeping whether or not any account of why deviations are informative survives. The framework says this about its own instruments in Section 5.5, and it is the best sentence in the material: instruments were chosen that survive the doctrine failing.

The second is that **a null result is a datum and needs a taxonomy.** R5's two readings, plus the provision for an undecided entry, is a better treatment of missing data than most qualitative practice manages, and it is directly transferable: any record of solicited responses should classify its silences rather than dropping them.

The third is that **a doctrine's constraints can be its content.** R7 refuses more than R6 licenses, and the refusals are specific: acquisition is fenced first, the harm-transport reading is forbidden outright, and the third constraint is stated as a property that makes the object self-limiting. A doctrine whose net effect is to narrow what its own definition permits is unusual and is worth marking as such.

9.2 What does not follow

That an utterance can constitute rather than describe does not establish that a self-addressed utterance does. That a transform characterises a linear channel does not establish that a deviation characterises a person. That a classification scheme has gaps does not license anything being carried through them — and the framework agrees, which is R7.

Two further non-implications are worth stating. First, **the diagnostic claim of Section 6.3 is single-use and unauditable**, and a selection device with no stated criterion and no false- identification rate is not an instrument in any sense this corpus should endorse; the engineering standard is the opposite [Kerckhoffs, 1883]. Second, **the framework's accuracy about physics is not evidence for its psychology.** The material states the borrowed results better than most popular accounts and declines to adjudicate what physicists have not settled. That accuracy is a virtue of the exposition, and a carefully stated premise is not a warrant for an unlicensed extension.

9.3 Relation to the corpus

The results fixed here are consumed downstream in specific ways, and three upstream results fix what they may be. The reflexive loop $T(x) = W(x, M(x))$, with the self-fulfilling prophecy identified as a fixed point of the composition of seeing and responding, is fixed in the founding paper of this series and is what R4 and R5 are read off [Curzi, 2026]. The scoring rule that

weights a branch by its accumulated resistance is the instrument this paper's commitments are selected with, and this paper inherits its demotion: the weight is an ordering and not a likelihood [Curzi, 2026]. And the self-reading gauge — that the distance remaining to a fixed point is bounded by the size of the last correction scaled by $q/(1-q)$ — is the instrument the recording procedure's weekly difference approximates without the bound [Curzi, 2026], which is worth stating precisely: **the framework's weekly subtraction is the gauge without its constant**, and it inherits none of the guarantee.

Four results elsewhere bound what this paper may claim. The corpus's evidence-versus-precedent distinction bounds what R4's accumulated readings can be made to do: an evidence column carries what could have come out otherwise, a precedent column carries only that something was independently arrived at and preserved, and two sources agreeing does not double the evidence when only one of them is evidence [Curzi, 2026w]. Its survivorship arithmetic prices the failure mode CC-E2 exactly, giving a surviving accuracy of 7.9 per cent where failures were invisible against 94.5 per cent where they were visible — a twelvefold difference produced by the filter alone [Curzi, 2026w]. The corpus's four failure classes for a claimed channel — no channel; a real channel with insufficient gain; a correct result extrapolated to the wrong magnitude; and a system unresponsive at this input size — are the diagnostic grid R5's two readings are a special case of [Curzi, 2026ac]. And the corpus's belief postulate is the template R1's badge should have followed: a magnitude form with a forbidding clause, $E = G \cdot B + \varepsilon$ with $G = 0 \Rightarrow E = 0$ for all B, badged and paired by explicit correspondence with a zero-efficiency case [Curzi, 2026ab]. **R1 has no forbidding clause**, and Section 11 records that as a defect.

Three further results are cited for the discipline they model. The corpus's determination result — that a set closed under negation has a resultant of exactly zero, verified at four sizes, and that an orthogonal set's mean direction shrinks as $1/\sqrt{N}$ — is the arithmetic behind CC-E2 and the reason a hall of agreeing sources carries nothing [Curzi, 2026ad]. The corpus's audit of a traditional operation set against a parameter model, returning seven mappings, two splits, one rendering and three rejections over thirteen entries, is the model for the audits this paper runs on its own material [Curzi, 2026]. And two refusal results bound the category material: that repetition cannot move a fixed point, since the equilibrium expression contains no time index [Curzi, 2026], and that an open problem should be stated formally rather than claimed as a result — no established mechanism produces a response depending on a symbol string as an *order* rather than on permutation-invariant functionals [Curzi, 2026]. The seam of Section 5.4 is a problem of that kind, and this paper states it as one.

4.0.10 10. Verification

This section is not a summary. It decomposes the paper's results into steps small enough to check one at a time, and labels each `ARITHMETIC`, `DEFINITION`, `CITED` (with the key) or `POSTULATED`. The razor is applied: any postulate whose removal leaves the results standing has been removed.

10.1 Postulates

- **P1 (DEFINITION).** An utterance is performative with respect to a speaker when it is the act it names. *Check:* D1.
- **P2 (POSTULATED).** A voiced position exerts pressure on the speaker's subsequent conduct toward consistency with it. *Check:* **not checkable from within this corpus and not checked; no entry supports it.**
- **P3 (POSTULATED).** A person infers their own dispositions partly from observing their own conduct, including their own speech. *Check:* **no entry.**
- **P4 (POSTULATED).** A voiced commitment re-prices later defection. *Check:* **no entry.**
- **P5 (ARITHMETIC).** A statement whose truth condition lies in the future is not contradicted by a present unchanged from yesterday. *Check:* redo it; it is a fact about the

statement's truth conditions and not about people.

- **P6 (POSTULATED).** A future voiced over a person without their request produces movement away from it, scaling with detected pressure. *Check:* **no entry.**
- **P7 (DEFINITION).** For a channel returning $y = T(x)$, the reading is $y - x$. *Check:* D2.
- **P8 (CITED: `oppenheim1996signals`).** For a linear time-invariant channel, the response to a known input characterises the channel completely; a channel returning its input unchanged is the identity and carries no information about itself.
- **P9 (POSTULATED).** A social channel's deviation can be read as a characterisation of the channel. *Check:* **not checked**; the framework itself badges this as a craft claim.
- **P10 (DEFINITION).** An absent return is entered with one of two readings, or with both. *Check:* D2 extended; a bookkeeping rule.
- **P11 (DEFINITION).** A classification scheme has fewer dimensions than what it files, so filing discards. *Check:* immediate for any finite scheme over an infinite domain.
- **P12 (DEFINITION).** A category violation harms no person and violates a scheme's disjointness. *Check:* D3.
- **P13 (POSTULATED, normative).** Where the law of persons and the law of categories conflict, the person-law governs. *Check:* not checkable; a normative position, labelled as one.
- **P14 (DEFINITION).** The three constraints: clean acquisition, clean possession, self-decoding. *Check:* C1; constraints on a domain, not claims.

Postulates dropped as unnecessary: the measurement-collapse analogy of Section 4.1, the superposition description of Section 6.1, the bodily-verdict mechanism of Section 4.5, the attentional-cost mechanism behind the destruction rule, the substitution effect behind the quarantine rule, and the humour mechanism of Section 4.6. Each was tested for load and carries none — every result below survives its removal — and all are reported in the text and used for nothing.

10.2 The compounding chain

Each line consumes only the outputs named on its left.

```
P1 + P2 + P3 + P4 -> R1   a spoken commitment is a performative
                          aimed at the speaker, on three
                          converging mechanisms              POSTULATED
                          (all three mechanisms uncited)
P5                -> L1    a future-tense commitment is not
                           contradicted by an unchanged day  ARITHMETIC
L1 + P2           -> R2    therefore it generates no
                           consistency pressure              POSTULATED (via P2)
(ethical)         -> R3a   a future may not be voiced over
                           another life without their word   DEFINITION
P6                -> R3b   and the attempt manufactures its
                           own refutation                    POSTULATED
P7 + P8           -> R4    read the difference; an unchanged
                           return reports only emptiness     CITED
R4 + P9           -> L2    the difference characterises a
                           SOCIAL channel                    POSTULATED
P10               -> R5    a null return is a datum with two
                           readings                          DEFINITION
P11 + P12         -> R6    a category violation, with the two
                           laws held apart                   DEFINITION
R6 + P13          -> L3    the person-law governs            POSTULATED
                                                             (normative)
P14               -> R7    the three constraints, any failure
```

```
of which converts the object into
an ordinary harm                    DEFINITION
```

What this chain shows, stated plainly. R4 is the only result in this paper CITED at its load-bearing step, and it is cited for a *linear* channel; its transfer to a social channel passes through P9, which the framework itself badges as a craft claim. R1 — the result the corpus cites most from this paper — passes through three postulates, none of which this corpus can support. R2 inherits P2. R3's ethical limb is a definition and its mechanical limb is a postulate. R5, R6 and R7 are DEFINITION throughout, which is exactly right for them: they are bookkeeping rules and a definition with constraints, and definitions are not weakened by being called definitions.

A chain whose steps are mostly DEFINITION and POSTULATED is not a proof, and this one is mostly DEFINITION and POSTULATED. What it is instead is a specification of the bill: P2, P3, P4 and P6 are four psychological claims, all four testable, all four absent from this corpus's bibliography, and settling them would decide the standing of R1, R2 and R3b. P9 is a fifth, testable in principle against any archive of dated correspondence, and it decides how much of R4 transfers. **Five claims. None was tested.**

One structural defect is recorded here rather than in the text. The corpus's own template for a badged postulate attaches a *forbidding clause* — a condition under which the postulate's mechanism is stipulated to produce exactly nothing [Curzi, 2026ab]. R1 has no such clause: as stated, there is no condition under which a spoken commitment is asserted to do nothing at all, which makes it harder to falsify than the corpus's own standard requires. The natural candidate — that a commitment whose content the speaker cannot begin occupying in any fraction today produces no effect — is available from Section 4.3's boundary and is **not** stated by the framework as a forbidding clause. This paper records the gap and does not fill it, since filling it would be new work rather than conversion.

10.3 The checks a reader should run first

1. **Redo the tally comparison.** 41, 9 and 3, with the middle card's nine sorting 7 to 2. Confirm that $7 + 2 = 9$ and that the selection rule as stated selects the middle card. Figure 1 plots all four numbers and the script asserts the sort exhausts the card.
2. **Redo the basis-point arithmetic.** At 240 million, 5.0 bp is 120,000 and 35 bp is 840,000, so the difference is 720,000 [Curzi, 2026]. It agrees.
3. **Redo the concentration chain.** $1/0.764992 = 1.307203$ and $\ln(1.307203) = 0.267890$ [Curzi, 2026]. Both agree to the printed decimals.
4. **Check the audit distribution sums.** $7 + 2 + 1 + 3 = 13$ [Curzi, 2026].
5. **Check that R1's three mechanisms carry no citation anywhere in this paper.** They do not, in Section 4.2, in Section 10.1, or in the Conclusion, and Section 11.3 lists them.
6. **Check that the analogy table in Section 7.1 assigns category 1 to exactly one entry**, and that the entry is qualified by the domain in which it holds.
7. **Check that no sentence in this paper describes R4 or R5 as identified.** Section 7.3 states that they are modelling positions and the Conclusion repeats it.

10.4 Dimensional audit

Symbol	Meaning	Dimension	Unit	Status
x	signal sent	none assigned	none	defined; an element of an unspecified set
$y = T(x)$	signal returned	$[x]$	none	defined by the channel

Symbol	Meaning	Dimension	Unit	Status
$y - x$	the reading	$[x]$	none	**defined for a linear signal space; a comparison and not a subtraction for the social case (Sec. 7.2)**
T, W, M	maps on states	—	—	defined; inherited from the founding paper of this series
tally count	recorded checks in one day	dimensionless	counts	**measured**; ratio scale with a true zero
resolution class	fine / middle / coarse	none assigned	**an ordered three-level category**	defined by the procedure's own construction rules
lean / grip	class of a recorded check	none assigned	**nominal, two categories**	defined; assigned by the participant
box membership	a scheme's assignment	none assigned	**nominal**	defined; R6 concerns consistency of assignment, not magnitude
$\sigma_x, \sigma_p, \hbar$	quantum spreads and reduced Planck constant (quoted)	m; kg·m·s^{-1}; J·s	SI	quoted from physics; not quantities of this framework
basis point (quoted)	annual charge per unit of value	dimensionless	10^{-4}	quoted from `ccp:CCP-053`; the arithmetic is reproduced in Sec. 8.6

Does the central relation balance? $y - x$ balances if and only if x and y inhabit a common space in which subtraction is defined. **For the engineering case they do; for the social case the source does not say what space its signals inhabit, and this paper does not supply one.** The honest entry is that the operation is a comparison, and Section 7.2 demotes it accordingly.

Does the tally measurement balance? Yes, and it is the only quantity in this paper that does without qualification. A count of recorded checks is dimensionless with a true zero, and 41, 9 and 3 are directly comparable. What does not follow is any inference to a constant, for the design reasons in Section 3.1.

A quantity with no operational unit is entered as having none. No unit was invented for a commitment's binding force, for the wear on a return, or for a scheme's discarded dimension, and none should be inferred from the fact that the tally counts are well behaved.

10.5 Limits and boundary behaviour

- **Deviation $\to 0$ (a perfect return).** The channel is the identity and reports nothing about itself. R4 is exact at this limit and the limit is the argument (Section 5.2).
- **Deviation $\to$ everything (no recognisable relation between sent and returned).** The reading rule is silent: with no correspondence there is no difference to take. The framework does not treat this limit, and it is the common case in a slow or crowded channel. **This is a gap in R4 and it is recorded as one.**
- **Listening window $\to 0$.** Every return is classified "far" and the record carries no information. **Listening window $\to \infty$.** No classification is ever made and the procedure never terminates, which the framework names as a real failure — practitioners who listen indefinitely because listening excuses them from acting.
- **Number of declared channels $\to$ all channels.** Pre-declaration ceases to constrain anything and CC-E5, the forged return, becomes unpreventable. The framework's guard is scarcity of declared channels and it does not state a bound.
- **Audience size $\to 0$.** The framework's prescribed condition. **Audience size $\to \infty$.** Its predicted failure, CC-C1. The framework predicts a monotone relation between audience and substitution and supplies no turning point; whether one exists is untested and, given the public commitment evidence the framework itself concedes, a turning point is the more plausible shape. **The framework's monotone reading is not supported and this paper does not adopt it.**
- **Scheme dimensions $\to$ the dimensions of what it files.** No object can violate the scheme, R6 is vacuous, and the scheme is unadministrable, which is why schemes discard. **Scheme dimensions $\to 1$.** Almost every object violates, R6 is trivial, and the diagnostic of Section 6.3 sorts nobody. The doctrine has content only in the intermediate regime and the framework does not say where its boundaries lie.
- **An explained instance.** R7's third constraint sets this limit itself: the object's function is destroyed by explanation, so the doctrine is self-limiting by construction and its diagnostic is single-use.

On invariance. The reading rule is invariant under any relabelling of the signal set that preserves the correspondence between sent and returned, which is a consequence of the construction rather than a finding, and it has been checked at no scale because nothing here is measured except the tally counts. **No conserved quantity is exhibited.**

4.0.11 11. Limitations

11.1 The paper's most-cited result rests entirely on uncited mechanisms. R1's three mechanisms — consistency pressure, self-perception, and the re-pricing of defection — have no supporting entry in this corpus's external bibliography, and R2 is an entailment of the first of them. This is stated first because it is the paper's principal exposure. **Falsified by:** a spoken-versus-silent comparison at matched content showing no difference, which would remove R1 and R2 together and is cheap to run.

11.2 R1 has no forbidding clause and the corpus's own standard requires one. Section 10.2 records the defect. As stated, R1 names no condition under which a spoken commitment does nothing, which makes it less falsifiable than the corpus's own badged postulate elsewhere [Curzi, 2026ab].

11.3 Claims with no supporting literature in this corpus's bibliography. Listed individually, because a reviewer will find them and finding them first is the point.

- *Consistency pressure following a voiced position (dissonance).* R1's first mechanism and R2's premise. **No entry.**
- *Self-perception: a person infers their dispositions from observing their own conduct.* R1's second mechanism, and the reason the framework requires the commitment to be voiced rather than held. **No entry.**

- *Commitment devices: a voiced commitment re-prices later defection.* R1's third mechanism. **No entry.**
- *Pressure on a perceived freedom produces counter-movement (reactance).* R3's mechanical limb and failure modes CC-C4 and CC-E6. **No entry.** Retained uncited; R3's ethical argument stands without it.
- *Announced identity goals substitute for enactment.* The entire warrant for the quarantine rule (Section 4.7) and for CC-C1. **No entry.**
- *Unfinished items hold attention and levy a standing cost.* The warrant for the destruction rule (Section 4.5). **No entry.**
- *Bodily verdicts on prospects precede articulate reasoning.* The mechanism behind the read-aloud taxonomy (Section 4.5). **No entry**, and the source itself notes the hypothesis has open questions.
- *Incongruity and its resolution as the structure of humour; laughter as a compression event.* Section 4.6. **No entry.**
- *The philosophy of performative utterances.* Section 4.1. **No entry**; the performativity literature in economics [Callon, 1998, MacKenzie, 2006] supports the general claim and not the self-addressed case.
- *Measurement reactivity, in the overdose form named by CC-U3.* **No entry**; McCambridge et al. [2014] bounds the nearest literature rather than supporting the claim.
- *Anchoring, availability, motivated reasoning, apophenia, personalisation.* The mechanisms named in the three failure catalogues. **No entry for any of them.**

That is **eleven classes of unsupported claim**, more than in either preceding paper of this series. A reader should treat the failure catalogues and the mechanism triad as the framework's clinical folklore, usefully organised, and not as findings.

11.4 The one measurement reported is a weak design. Section 3.1 states the flaws: a single participant, three days in a fixed order so that resolution is confounded with day number, no blinding, and the outcome recorded by the person being measured. **Counterbalancing would cost nothing and was not done**, and no result in this paper should be read as evidence about the magnitude of the effect.

11.5 Parameters and classifications chosen rather than measured. The three resolution classes are defined by construction rules the source chose (one number *or* one date; one place *or* one name), and nothing establishes that those rules produce equally spaced levels or that any other partition would behave the same way. The lean/grip classification is assigned by the participant without a stated rubric, without a second rater, and without any reliability estimate.

11.6 The reading rule is demoted from arithmetic to comparison. Section 7.2 states the ground: subtraction requires a space where subtraction is defined, and the source does not say what space its signals inhabit. Nothing in the paper computes with the difference, so the demotion costs nothing — but a downstream paper that treats $y - x$ as a quantity would be doing something the dimensions forbid.

11.7 R4's transfer premises are unverified. Linearity and time-invariance make the engineering result exact and are not established — and are implausible — for a channel made of people. What transfers is the reading discipline and not the theorem, and Section 5.2 says so.

11.8 R4 has a gap at the limit where no correspondence exists. Section 10.5 records it: where a return bears no recognisable relation to anything sent, the rule is silent, and that is the common case in a slow or crowded channel. The framework does not treat it.

11.9 The diagnostic use of category violation is single-use, unauditable, and cited against. Section 6.3 gives the argument. A selection device with no stated criterion, no false-identification rate and no appeal is not an instrument this corpus should endorse, and the engineering standard runs the other way [Kerckhoffs, 1883, Saltzer and Schroeder, 1975].

11.10 The seam where a second will enters the loop is open and the framework says so. Section 5.4 reports the concession. This paper does not close it either, and no result here should be read as settling the case in which a return is obtained through another person

without their word.

11.11 The person-law ranking is a normative position. Section 6.2 labels it. It is defensible and it is not established by anything, and R6's usefulness as a definition does not depend on it.

11.12 Nothing here is identified. As an aggregate specification, the reflexive loop's endogenous effect is not identified from observational data [Manski, 1993], and the network route is closed by the generic confounding of influence with selection [Shalizi and Thomas, 2011]. This paper adopts no remedy and therefore states R4 and R5 as **modelling positions rather than identified empirical results** (Section 7.3). **Resolved by:** a randomised or partial-population design, or a defended stipulation in the manner of Brock and Durlauf [2001].

11.13 The framework's own effects are not exempt from its own thesis. By the Lucas critique, tallies and response patterns recorded by a practitioner who has read the selection rule are not comparable to those recorded by one who has not [Lucas, 1976]. No stability is claimed and no result here depends on any.

11.14 Case material is illustrative. The worked cases in Sections 3.1 and 5.6 are single, self-reported, retrospective accounts drawn from the source manuscript. They illustrate what the procedures return and establish nothing about whether the procedures work [Ioannidis, 2005, Open Science Collaboration, 2015].

11.15 No hypothesis was tested and no data were collected. H1–H5 stand. The discipline that would make a later confirmatory attempt credible is specified and not exercised [Nosek et al., 2018, Simmons et al., 2011], and H1 in particular is complicated by the framework's own reflexivity (Section 7.3), which this paper states rather than proposing the test as though it were straightforward.

4.0.12 12. Conclusion

The seven results are stated below in the exact form in which the rest of the corpus cites them. Downstream papers should quote this text rather than paraphrase it.

R1. A spoken commitment is a performative aimed at the speaker, supported by three converging mechanisms: consistency pressure following a voiced position, self-perception from overheard testimony, and the re-pricing of later defection. **All three mechanisms are uncited in this corpus's bibliography, and R1 is therefore a specification resting on unsupported psychology rather than an established result.**

R2. The present-tense requirement is calibration and not etiquette: a future-tense formulation is consistent with a day identical to yesterday, so it generates no dissonance and the consistency mechanism has nothing to act on. **R2 is an entailment of R1's first mechanism and inherits its status.**

R3. No future is to be voiced over another person's life without their word, and the prohibition covers the benevolent case. **The ethical argument stands on its own and requires no literature. The framework's further claim — that imposed futures are registered as pressure and manufacture their own refutation — has no supporting entry in this corpus's bibliography and is reported uncited.**

R4. The feedback reading rule: record what was sent, x, and what returned, $y = T(x)$, and read the difference $y - x$, since a channel that returns a signal unchanged reports only that nothing in it could touch the signal. **Exact for a linear channel, where the transform is the channel's complete characterisation; a shared mathematical form and not a derivation for a social channel, and a comparison rather than an arithmetic difference in that case.**

R5. A null return is a datum with exactly two readings — a distance beyond the listening window, or a medium that absorbs the frequency entirely — and may be entered under both while undecided. It is not an absence of evidence.

R6. A category violation is an object or act that harms no person — nothing taken, nobody deceived about anything owed, no will overridden — while standing in two categories a scheme declared disjoint. The law of persons and the law of categories are held apart and

neither is allowed to dissolve the other; where they conflict, the person-law governs. **That ranking is a normative position and is labelled as one.**

R7. Three constraints govern the doctrine: no coercive acquisition, possession by clean holding only, and self-decoding, since an explained instance is a spent one. Failing any one converts the object into an ordinary harm with a doctrine laid over it.

Counter-result. The category material constrains far more than it licenses. Category innocence is explicitly refused as a defence of acquisition; the doctrine forbids moving any harm through a discovered gap in a scheme; and the framework concedes on the record that the case in which the feedback loop runs through another person's will is an unsettled seam it does not close. To that this paper adds two findings the framework does not state: the diagnostic use of a category violation is single-use per object and unauditable by construction, and R1 carries no forbidding clause of the kind the corpus's own standard requires.

What is established, then, is one theorem of signal analysis about linear channels, one bookkeeping rule for missing observations that is better than most qualitative practice manages, one definition whose measurement-theoretic standing is unproblematic because it claims nothing numerical, and a set of constraints that narrow the definition rather than extending it. What is not established is any of the five hypotheses, any of the four psychological mechanisms on which R1, R2 and R3's mechanical limb rest, or the transfer of the reading rule from a linear channel to a social one. The paper's contribution at this stage is to name the five testable premises whose settlement would decide the matter — consistency pressure, self-perception, defection re-pricing, counter-movement under imposition, and the readability of a social channel's deviation — and to note that all five are testable, that four of them are cheap, and that none was tested.

4.0.13 13. Declarations

Data and code availability. No external dataset was collected for this paper and none exists. The worked cases reported in Sections 3.1 and 5.6, including the tally counts and the reproduced record card, are drawn from the source manuscript's own record and are not a dataset. All figures are generated by the scripts named in their captions (`figs/src/CCP-A04_fig01.py` through `figs/src/CCP-A04_fig05.py`), which import the shared style module `lib/figstyle.py`; every number plotted is quoted from the source text identified in the script's docstring. Two of the five scripts contain assertions that fail if the numbers they plot cease to agree with the source claims they are drawn from.

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the analysis, the figures, and the text.

Ethics. No human-subjects research was conducted for this paper. No personal data were collected or processed. The worked cases in Sections 3.1 and 5.6 are drawn from the source manuscript, are presented under the descriptions they carry there, and are illustrative rather than evidential; they were not gathered under a research protocol and are not offered as one. Nothing in this paper is clinical advice. Sections 4.4 and 6.4 report constraints the framework places on its own procedures — that no future is to be voiced over another person's life without their word, and that nothing may be acquired coercively — and this paper reports both as constraints on the method rather than as endorsements of the techniques they constrain. Section 6.3 states expressly that the framework's proposed use of a category violation as a selection device is unauditable and is not endorsed here.

Use of the ceremonial register. The source manuscript is in part an initiatory document. This paper analyses that register as an object of study and does not reproduce it; ceremonial names are reported only where they identify a construct this paper renames, and the source's own term for the object defined in R6 is reported once, in Section 1.3, because downstream papers refer to the material by it.

Chapter 5

Convergence and Its Dark Twin: The Contraction Mapping, the Audit of Arrival, and the Consent Law of Closure

Abstract. This instalment states the corpus's central theorem, its converse, and the audit and closure disciplines that follow. Banach's fixed-point theorem is stated exactly: a contraction with modulus $q < 1$ on a complete space has exactly one fixed point, reached from any starting state, with the remaining distance shrinking geometrically. It then extracts the a posteriori bound the framework prizes above the theorem, $d(x_n, x^*) \leq (k^n/(1-k))\, d(x_0, x_1)$, in the practical form that distance to arrival is bounded by the last correction scaled by $q/(1-q)$, so that the size of a most recent revision is a self-reading gauge. The converse is stated with equal force: unit gain gives an isometry and so an orbit with no arrival, and gain above one gives an amplifier in which corrections sized to stale feedback arrive as new errors. Completeness is read as a failure condition: a foreseen state a life cannot contain is a hole in the space, and iteration toward it contracts forever and arrives never, which is the one failure mode invisible from inside the loop. Arrival is certified by audit rather than by feeling, on three instruments, a dated record predating arrival, at least three preserved discrepancies, and a geometrically shrinking sequence of corrections. Closure imposes three verdicts, a single-active-commitment rule, a consent law, and three conditions barring operation. Four hypotheses are stated in falsifiable form; none is tested here, and three claims the source treats as established are reported uncited because this corpus's bibliography contains no support for them.

Keywords: contraction mapping; a posteriori bound; divergence; reactance; completeness failure; hindsight retrofitting; consent; closure verdicts

Series: A. Foundations **Paper:** CCP-A05

5.0.1 1. Introduction

1.1 The question

A framework that tells a practitioner to iterate owes them three things it usually does not supply: a condition under which iteration arrives, a condition under which it does not, and

a test that distinguishes an arrival from the feeling of one. This paper supplies all three, and finds that the third is the hardest and the most often skipped.

The corpus's central theorem is Banach's fixed-point theorem. Its founding paper states the condition $d(T(x), T(y)) \leq k\, d(x, y)$ with $0 \leq k < 1$ and derives existence, uniqueness and geometric convergence from any starting point [Curzi, 2026]. That statement is exact and it is not the whole instrument. What an operator actually needs is not the guarantee but a gauge — something readable from inside the loop, before arrival, that reports distance to the fixed point. The theorem supplies one, in a corollary that is less famous than the theorem and more useful than it.

This paper asks: what does the contraction theorem license when a life is the space, what does its converse forbid, and what evidence can certify an arrival given that the machinery which rehearses a future is the same machinery that reconstructs a past?

1.2 What this paper contributes

Eight results are established or fixed here, labelled R1 through R8 and restated in the Conclusion.

R1 states Banach's theorem exactly: a contraction with modulus $q < 1$ on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least q each pass. **R2** extracts the a posteriori bound $d(x_n, x^*) \leq \frac{k^n}{1-k}\, d(x_0, x_1)$ and reads it operationally: distance to arrival is bounded by the size of the last correction scaled by $q/(1-q)$, so the size of one's most recent revision is a self-reading gauge. **R3** states the honest converse: gain exactly one gives an isometry that orbits without arriving, and gain at or above one gives an amplifier whose corrections, sized to delayed feedback, arrive as new errors. **R4** states the stability rule that follows: a correction smaller than the observed difference can be wrong and still converge, while a correction larger than the difference can be right and still diverge, so restraint is a stability criterion rather than a virtue. **R5** reads completeness as a failure condition: a foreseen state the life cannot contain is a hole in the space, and iteration toward it contracts indefinitely while arriving never. **R6** specifies a three-instrument audit of arrival — a dated record predating arrival, at least three preserved discrepancies, and a geometrically decaying sequence of corrections — on the ground that a seamless match indicates a retrofitted record. **R7** states the closure discipline: exactly three verdicts (fulfilled, released, lapsed), one active commitment at a time, and a consent law forbidding a future voiced over another's path without their word. **R8** states the externality result as engineering: a fixed point requiring another party's displacement re-enters the loop as friction, raising the gain and converting a basin into a position that must be held.

The counter-result is stated with the same force. This paper refuses the certificate its own theorem appears to offer: a human life is not a metric space, nobody has measured the relevant contraction modulus, and the theorem therefore transfers as shape rather than as guarantee. It also holds that the feeling of recognition on arrival proves nothing, because the machinery that rehearses a future retrofits a past with equal fluency.

1.3 Register and sources

The source manuscript is written in two braided registers, one narrative and one analytic. This paper reports the analytic register; the narrative material is fictional exposition and is not reproduced. Ceremonial names are reported once, as objects of study, where they identify a construct that this paper renames. The material below corresponds to three of the source's twelve constructs — iteration, arrival, and closure — and is presented as one continuous argument because the three are one mechanism seen at three stages.

5.0.2 2. Hypothesis

The paper's claims divide into one theorem, three statements about that theorem's converse and corollaries which are also arithmetic, and four claims about human practice which are genuinely hypotheses. They are separated here, and the ones that cannot be falsified are labelled as definitions or postulates rather than dressed as hypotheses [Popper, 1959].

H1 — the gauge hypothesis. In a loop that is in fact contracting, the size of the most recent correction bounds the distance remaining, so an operator can read their own distance to arrival off their own record without seeing the destination. *Predicted:* across a series of iterations of any documented revision process, the observed remaining distance never exceeds $q/(1-q)$ times the last correction, once q is estimated from the ratio of successive corrections. *Falsified by:* a documented revision series in which corrections shrink at a stable ratio while remaining distance does not respect the bound — which would show the process was not the contraction the ratio suggested. *Status:* mathematically true of a contraction (Section 3.3); the hypothesis is that real revision series are contractions. **Not tested here.**

H2 — the divergence-signature hypothesis. A loop whose gain has risen above one signals it by alternation of correction sign before amplitude grows. *Predicted:* in records of failing iterative projects, sign alternation precedes amplitude growth by at least one pass, so it is available as an early warning. *Falsified by:* amplitude growth arriving without prior alternation, or alternation occurring at equal rates in converging and diverging series. *Status:* true of the modelled scalar recursion; the transfer to documented human projects is untested. **Not tested here.**

H3 — the residue hypothesis (the paper's sharpest empirical claim). A genuine prediction record disagrees with the outcome in inventory while agreeing in structure, so the count of preserved errors discriminates genuine foresight from retrospective reconstruction. *Predicted:* records made before an outcome contain at least three specific particulars the outcome contradicts; records reconstructed after an outcome contain fewer. *Falsified by:* prospectively dated records showing zero to two discrepancies at the same rate as retrospective reconstructions, in matched cases. *Status:* directly testable with existing preregistration archives, and **not tested here.** The threshold of three is a stipulation, not an estimate (Section 12.5).

H4 — the externality hypothesis. A fixed point requiring another party's uncompensated displacement raises the effective loop gain, because the displaced party becomes a second modeller inside W. *Predicted:* arrangements reached over an unconsenting party show rising maintenance cost over time, where structurally identical consented arrangements do not. *Falsified by:* matched consented and unconsented arrangements showing equal maintenance cost. *Status:* stated as engineering; **not tested here**, and the mechanism it invokes has no supporting entry in this corpus's external bibliography (Section 12.3).

D1 — contraction (definition). A map T on a metric space (X, d) is a contraction with modulus q if $d(T(x), T(y)) \leq q\, d(x, y)$ for all $x, y \in X$ and some fixed $q < 1$. Not most pairs. All of them, always, by at least that factor.

D2 — closure verdict (definition). Every commitment receives exactly one of three written verdicts: fulfilled, released, or lapsed. This is a stipulated taxonomy, not a finding, and its adequacy is a question about practice rather than about the world.

P1 — the transfer postulate. That a human life admits a state space which is complete under some metric, and a reflexive loop which contracts with a uniform modulus, is postulated wherever the theorem is applied to a life. It is not demonstrated anywhere in this corpus. This paper's principal methodological contribution is to keep saying so.

5.0.3 3. The theorem, and the corollary that matters more

3.1 Statement

Take a space of states — positions, drafts, situations; anything on which "distance between two states" means something and behaves honestly. Require the space to be **complete**: any sequence of states that bunches up ever more tightly must be bunching toward a state that is actually in the space. No holes where a limit ought to live.

Now take a map T that turns any state into a next state: feed in where you are, receive where the pass leaves you. Call T a **contraction** if there is a fixed fraction q, strictly less than one, such that for *any* two states the distance between their images under T is at most q times the distance between the originals.

Then Banach [1922] says exactly this and no more:

> A contraction on a complete space has exactly one fixed point x^* with $T(x^*) = x^*$, and iteration finds it from anywhere: start at any state whatever, apply the map again and again, and the sequence converges to x^* geometrically, so that after n passes the distance from the fixed point is at most q^n times a constant fixed at the start.

Each clause is earned, and it is worth seeing how. *Exactly one:* if two different states were each fixed, the map would have to shrink the distance between them by the factor q; but both are fixed, so that distance cannot change at all, and the only distance that both shrinks and stays the same is zero, so the two points were one point. *From anywhere:* the theorem asks nothing about where the iteration begins — not that it begins close, not that it begins wise. *Geometrically:* if q is one half, each pass halves the remaining gap, so ten passes put the gap under a thousandth of what it was.

The result is not exotic, and its oldest worked example predates the theorem by two millennia: the surveyor's rule for square roots. To find $\sqrt{2}$, guess anything — guess 1 — then repeat the correction *take the average of your guess and two divided by your guess.* The source prints the iterates $1,\ 1.5,\ 1.41666\ldots,\ 1.4142156\ldots,\ 1.4142135\ldots$ and observes that four passes carry a poor guess to six true figures. Each pass takes the error and roughly squares it away; the map contracts, the space is complete, and no cleverness in the starting guess would have bought more than one pass's head start. Figure 2 recomputes the same recursion and confirms the printed iterates digit for digit.

3.2 The premises, stated as premises

Two conditions carry the whole result and both are named here rather than assumed.

The map must contract. Every pass must actually shrink the gap, which means every revision must be taken from the world's returned difference and not from appetite, mood, or fashion. A revision taken from mood is a step of unknown direction; enough of them and the map is a stagger, and staggers have no theorem.

The space must be complete. The state converged upon must be a state the actual system can contain. This condition is usually treated as a technicality and it is not: its failure is a distinct and diagnosable mode, treated at length in Section 5.

And one honesty above the others, stated plainly so that no one can say the mathematics concealed it: **a human life is not a metric space, and no one has measured any operator's q.** The theorem is real; its transposition to a life is this corpus's, made because the shape is the same. What the transposition earns is not a certificate. It is a discipline of expectation with a mechanism in it — the knowledge of what would have to be true for arrival to be guaranteed — and therefore a checklist that can fail inspection and say where. That is worth more than encouragement and less than proof, and this paper does not confuse the two. The formal obstacle is not merely unmeasured: psychological similarity is asymmetric and violates the triangle inequality, so any d built from similarity judgements is not a metric at all [Tversky,

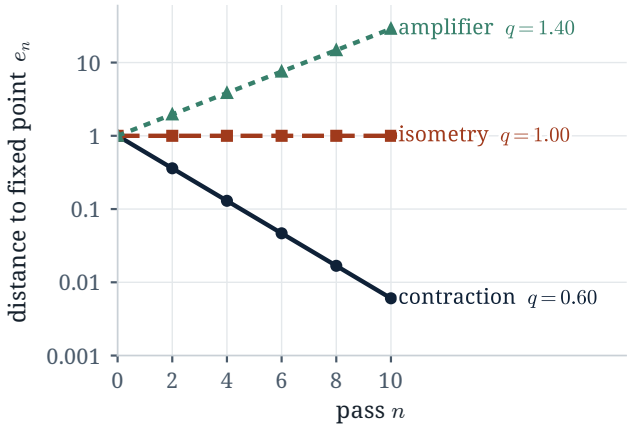


Figure 5.1: The three gain regimes, on the same axes. Modelled, not measured: the scalar recursion $e_{n+1} = q\,e_n$ with $e_0 = 1$ is the worst-case distance-to-fixed-point sequence a Lipschitz map of modulus q admits, with $q \in \{0.60, 1.00, 1.40\}$. Below unit gain the distance falls geometrically; at unit gain it neither falls nor grows, which is an orbit and not a plateau; above unit gain the loop manufactures error at the same rate the contraction would have removed it. Script: `figs/src/CCP-A05_fig01.py`.

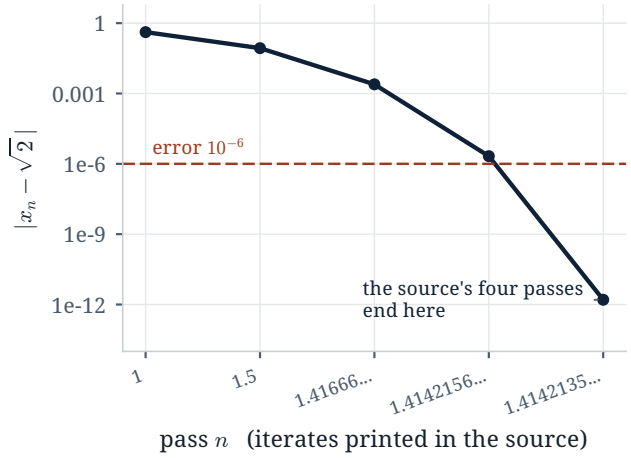


Figure 5.2: The source's own worked contraction, recomputed. Measured against the source: the recursion $x_{n+1} = (x_n + 2/x_n)/2$ from $x_0 = 1$ reproduces the printed iterates $1,\ 1.5,\ 1.41666\ldots,\ 1.4142156\ldots,\ 1.4142135\ldots$ exactly, and the absolute error against $\sqrt{2}$ falls below 10^{-6} within the four passes the source claims. The script asserts the agreement and fails if it is ever lost. Script: `figs/src/CCP-A05_fig02.py`.

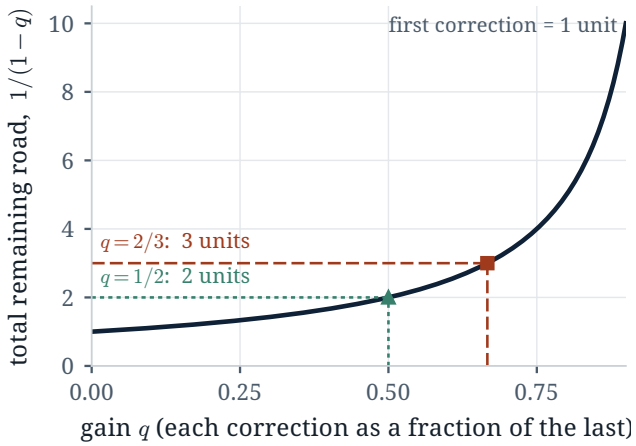


Figure 5.3: The whole remaining road, priced by the ratio. Arithmetic, not measurement: with a first correction of one unit and each correction a fixed fraction q of the last, the total distance ever to be travelled is the geometric sum $1/(1-q)$. The two annotated points are the source's own cases — $q = 1/2$ giving two units and $q = 2/3$ giving three, the latter stated in the source as "the entire road, walked forever, is three units long". The curve's divergence as $q \to 1$ is the same fact as the isometry in Figure 1. Script: `figs/src/CCP-A05_fig03.py`.

1977], and the underlying question of whether the attribute is quantitative in the first place has an unanswered critic [Michell, 1997].

3.3 The a posteriori bound (R2)

One consequence rides along with the proof, and this corpus prizes it above the theorem, because it answers the question an operator actually asks — not *will I arrive* but *how far off am I now*.

The fixed point cannot be seen; if it could be seen the map would not be needed. But something just as good can be seen: **the size of the last correction.** On a contracting loop the remaining distance is bounded by the last step taken, scaled by a constant depending only on the gain:

$$d(x_n,\, x^*) \;\le\; \frac{k^n}{1-k}\, d(x_0,\, x_1)$$

and, in the form an operator uses, the distance still remaining is at most $\frac{q}{1-q}$ times the length of the latest correction. A gain of one half means one is never farther from the fixed point than the whole size of the step just made.

The practical content is worth stating without ornament. The operator who wants to know how near arrival stands does not need to see the future; they need only measure their most recent revision. Corrections running to pages: far, and now quantified. Corrections down to a word: near, and *provably* near, on the theorem's own signature — not consoled, not encouraged, **bounded**. The mathematicians call this an a posteriori estimate. Its whole value is that it is self-reading: the same record that holds the gain also reports the distance.

The geometric-series form makes the total road finite and prices it in advance. If the first correction moved one unit and each subsequent correction is two thirds of the last, then the entire road, walked for ever, is $1/(1-2/3) = 3$ units long. This is why the last steps of a true convergence are so small that they feel like standing still, and why the door, when it is reached, is barely a step.

One caution attaches to the gauge and is not optional. A quantity that is read as an indicator of one's own progress, and acted upon, is exposed to the standard failure of social indicators: a measure adopted as a target ceases to be a good measure [Campbell, 1979, Goodhart, 1975]. An operator who knows that small corrections certify nearness can produce small corrections. The audit of Section 6 exists partly because of this, and the exposure is structural rather than incidental.

3.4 What the theorem says about deserving

The theorem has no clause about deserving, and the corpus makes a point of this. It has no clause about how far the starting state lies from the fixed point, or how many years were spent circling before the first true correction was cut. Distance costs passes — the constant out front may be large, and early iterations from far away are humbling — but distance costs *only* passes, and each pass is the same size of work.

That is a statement about the mathematics and it is exact. It is not a statement about lives, and the difference is P1.

5.0.4 4. The converse (R3, R4)

A treatment of convergence without a treatment of divergence is a loan with the terms withheld. Everything above required q strictly less than one. Loosen that and the guarantee does not weaken; it inverts, in two distinct ways with two distinct signatures.

4.1 Unit gain: the orbit

The pass that preserves distances — same utterance, same conduct, no revision taken, or a revision taken and quietly unwound because nothing held it — orbits. It can orbit with complete sincerity and perfect attendance. Nothing in orbiting reduces the gap; the map is an isometry, and isometries have no theorem and no arrival. Most of what is called being stuck is not a wall. It is an unpitched loop.

The distinction the source draws here is operationally sharp and worth preserving: a record showing no corrections *near the end* of a long series of shrinking ones is the quiet of convergence; a record showing no corrections *from the beginning* is a groove that never had a pitch. The two stillnesses are told apart by whether shrinking corrections precede the silence, which is exactly the third instrument of Section 6.

4.2 Gain above one: the amplifier

This is the dangerous regime, because it wears effort's face. The operator, impatient with small steps, reacts to each returned difference with a revision larger than the difference that prompted it, overshoots, receives back a larger opposing difference, and corrects harder still. Any helmsman knows the sequence: oversteer, overcorrect, and the swings grow until the wake is a snake. In a loop with delay — and every real loop has delay — corrections sized to old feedback arrive as new errors, and the oscillation feeds itself. **Divergence does not look like collapse at first. It looks like intensity.**

The general treatment of gain, delay and stability margin is the control engineer's and is standard [Åström and Murray, 2008, Oppenheim et al., 1996], as is the first-order difference-equation account of the same three regimes [Elaydi, 2005, Kelley and Peterson, 2001]. That a human operator inside a control loop can be modelled this way at all is an old and respectable result [Tustin, 1947], and the founding claim that one feedback mathematics governs machines and organisms alike is older still [Wiener, 1948]. Delay in a stock-and-flow system is precisely where human judgement fails: people systematically misjudge stocks with delays, and the resulting oscillation is a robust experimental finding [Forrester, 1961, Sterman, 1989].

The corresponding failure in institutions is a slow erosion of a standard that nobody decided to erode [Vaughan, 1996], and its acute form is an organisation losing its own structure under stress [Weick, 1993].

4.3 The case that has no citation, and is retained anyway

When the loop runs through another person, the amplifying case has, according to the source, a mechanism with a clinical name: press a will and the counterforce rises; push a foreseen future harder and the pressed party is drawn to its opposite *because* it is the opposite. A prediction pushed too hard does not merely fail — it manufactures its own refutation and returns it signed. The source calls the clinical mechanism reactance.

This corpus's external bibliography contains no entry for reactance. The claim is retained here because it is load-bearing for R3's human limb and because removing it would misrepresent the framework; it is retained **uncited**, and it is listed in Section 12.3 among the claims this paper cannot support. What can be supported is narrower and is stated instead: that measuring or managing a process changes it in ways not simply predicted [Campbell, 1979], and that redundancy and added control can reduce rather than increase safety in coupled systems [Sagan, 1993, Perrow, 1984].

4.4 The stability rule (R4)

The remedy is not subtle. Small corrections; many passes. The standing rule is: never correct by more than a fraction of the gap the differences show, and halve the step the moment the last corrections begin alternating in direction.

The reason the rule works is the whole of R4 and it inverts the usual moral ordering:

> **A correction smaller than the difference can be wrong and still converge, because the next pass will catch it. A correction larger than the difference can be right and still diverge.**

In this craft, restraint is not a virtue. It is a stability criterion. The same statement in control terms is that a gain margin is a property of the loop and not of the operator's character, and that the loop's stability is decided before any particular correction's correctness is known [Åström and Murray, 2008]. The corpus's storage law supplies the same result with an explicit increment bound: the rule $r \leq h$ with a recommended reserve $r \leq 0.7h$, since the maximum $r_{\max} = h$ is attained only at the boundary $U = 1$, together with the branch in which no interaction geometry can prevent order from decreasing [Curzi, 2026j]. And the same map appears there in closed form: the storage map $T(Q) = (1-\delta)Q + b$ is a contraction with constant $1-\delta$, giving $Q_\infty = (\eta M^* - c^*)/\delta$ and the exact error law $|Q_t - Q_\infty| = (1-\delta)^t |Q_0 - Q_\infty|$ [Curzi, 2026q]; its parent statement is the storage law $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$ with $S_t = \ln(1 + (N-1)U_t)$, Lipschitz with modulus $k = 1-\delta$, so that $\delta > 0$ is exactly what makes it a contraction and $Q^* = (\eta S - C)/\delta$ its steady state [Curzi, 2026a].

5.0.5 5. Completeness as a failure condition (R5)

Completeness is where a transposition of this theorem most often fails silently, and the failure has a signature worth naming because every local reading of it is good.

The state converged upon must be a state the actual system can contain — the limit must have somewhere to land. There are foreseen states that are holes in the space: outcomes requiring a person the operator has made no provision to become, sustained by conditions the life as lived cannot hold. **Iterate toward a hole and the passes contract beautifully, for ever, arriving never:** the gap shrinks toward a point that is not there.

The diagnostic problem is that contraction is preserved. Every instrument in Section 6 reports health. The corrections shrink; the ratio is stable; the record is honest; there is no

oscillation and no alternation. Only a check performed from **outside** the loop detects it, and the corpus supplies one: ask of the target, with the ledger open, whether the life as actually lived could hold it if it arrived on an ordinary Tuesday — naming the room it would occupy, the hours it would consume, and the person the operator would have to have become meanwhile. If the check finds a hole, the honest responses are two: widen the space until the limit has somewhere to land, or move the target by daylight as a declared new undertaking. Pushing the loop harder at the wall is not among them.

This is a real mathematical point and not a metaphor. A Cauchy sequence in an incomplete space converges to nothing in that space while satisfying every internal criterion of convergence; the rationals under the usual metric are the standard example, and $\sqrt{2}$ — the very limit of Figure 2's iteration — is precisely the hole. The recursion of Figure 2 is a contraction on the rationals and has no fixed point there. That is completeness doing visible work, and it is why the condition appears in the theorem rather than in its preamble.

The organisational analogue is the case in which making a system legible to a plan destroys the local variety the system was actually running on: the plan contracts toward a state the terrain cannot hold [Scott, 1998]. The analogous engineering treatment is to regard the constraint set, not the component reliability, as the object of design [Leveson, 2011].

5.0.6 6. The audit of arrival (R6)

6.1 Why the feeling proves nothing

Every operator reports it and none is prepared for it: a fixed point entered at last does not feel like triumph. It feels like recognition. The proportions of the hour are familiar; the hook is where the arm expected it.

The source offers a mechanism: that the machinery with which a mind imagines a future is not separate from the machinery with which it remembers a past, so a scene assembled often enough is consolidated as a memory is, and recognition on arrival is the felt signature of zero prediction error — the difference signal, for once, silent. **This corpus's external bibliography contains no entry for the constructive memory literature, and the mechanism is therefore reported uncited** (Section 12.3).

But the paper's claim does not depend on that mechanism, and this is worth being clear about, because it is the one place where an unsupported explanation and a supported conclusion are easily confused. The conclusion is:

> **The feeling of recognition proves nothing**, because whatever machinery produces it will produce it equally for a rehearsed future correctly reached and for a past quietly tidied to agree with a present.

That conclusion is supported by results this bibliography does contain. People report the causes of their own mental processes with a confidence unrelated to accuracy [Nisbett and Wilson, 1977]; the answers a person gives about their own states are shaped by how the question was asked [Schwarz, 1999]; and attention directed at a named component of well-being systematically distorts the judgement of it [Kahneman et al., 2006]. Nothing about a subjective report of fit, however vivid, is evidence of a prior prediction.

The counterfeit feels identical from inside — warmer, if anything, for having been tailored after the measurements were taken. The framework therefore does not permit arrival to be certified by feeling. Arrival is certified by audit, and the audit has three instruments.

6.2 First instrument: the record precedes the room

A genuine convergence leaves a paper trail laid down *before* arrival — dated, specific and falsifiable. If the only evidence of having foreseen an outcome is present testimony that it was foreseen, there is no evidence; there is a mood.

This is the practice the open-science literature independently arrived at, and its rationale transfers without modification: a prediction registered in advance constrains what can later be claimed to have been predicted [Nosek et al., 2018], and the alternative is a space of analytic and interpretive choices wide enough to support almost any conclusion [Simmons et al., 2011, Gelman and Loken, 2013]. The subtle form of the problem does not require any bad faith and is the one most likely to affect an honest author [Gelman and Loken, 2013]. The standard against which a directional claim should be judged is severe: directional predictions are nearly costless, and a theory earns credit by making point predictions with constants [Meehl, 1967]. The corpus's own habit of specifying numbers rather than directions is the correct response to that objection, and stating so explicitly is part of the answer [Meehl, 1990]. The design principle that makes such records worth keeping is that they should discriminate between named alternatives, not merely accord with one [Platt, 1964].

6.3 Second instrument: the residue (H3)

This is the subtle one. A genuine convergence agrees with the record in **structure** and disagrees with it in **inventory**. The returned world is never the one that was sent. A record may have foreseen light from the left, an old desk, water held at an angle; it did not foresee the bakery, and the warm smell in the record is marked *unidentified*, because at the time it honestly could not be identified.

That failure, preserved in the record, is the strongest available authentication.

> **A room that matches an account in every particular is a room the account was written in. The forger's tell is perfection. Truth arrives with its discrepancies still attached; only counterfeits are seamless.**

The operational rule is a count: on auditing an arrival, count the details the dated record got flatly wrong, and if fewer than three can be found, do not trust the record.

The threshold of three is a stipulation and not an estimate. Nothing establishes three rather than two or five, and Section 12.5 says so. What the rule encodes is a real asymmetry — a prospective record has independent opportunities to be wrong that a reconstruction does not — and that asymmetry is the content of H3.

6.4 Third instrument: the shrinking corrections

Convergence has a signature no single moment can fake, because it is a *sequence*. Going back through the record and measuring the size of each revision, a true approach to a fixed point shows them shrinking, roughly geometrically, each pass a fraction of the last, exactly as the bound of Section 3.3 prices it. A retrofit shows no such decay: an ordinary wandering line, and then one large convenient leap explained after the landing.

A conclusion can be faked. A ratio cannot be faked backwards through a dated record.

The corpus supplies the decay constants for its own storage instance, and they are worth quoting because they show what the signature looks like when it is measured: at $\delta = 0.09$ the time constant is $\tau = 10.6033$, the half-life is 7.3496 passes, and the error falls to five per cent by pass 32, to one per cent at 48.83, and to a tenth of a per cent at 73.24 [Curzi, 2026]. The same contraction verified to a maximum residual of 2.0624×10^{-15} over a hundred passes is reported elsewhere in the corpus [Curzi, 2026].

What the drawing quietly forbids is the belief — dear to the exhausted — that the great correction is still ahead, and that arrival will announce itself in upheaval commensurate with the wanting. The arithmetic says otherwise: if corrections have been taken honestly, the largest step of the journey is behind. It was the first one, and the smallness since is the signature of approach rather than a failure of nerve.

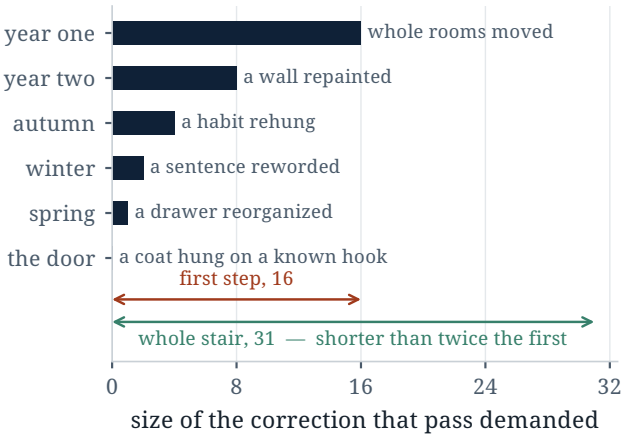


Figure 5.4: The source's own drawing of a converging series of corrections, redrawn to scale. Measured from the source: its monospace figure sets six corrections at bar lengths 16, 8, 4, 2, 1 and 0 characters, labelled year one through the door, and closes with the claim that "the whole stair, laid end to end, is shorter than twice its first step". The bar lengths here are those character counts; the total, 31, is computed by the script and is indeed below $2 \times 16 = 32$. Script: `figs/src/CCP-A05_fig04.py`.

6.5 The auxiliary: the walk-away test

One blunt auxiliary completes the audit. A genuine fixed point is a basin: it can afford inspection, delay and doubt, because everything in the neighbourhood drains back toward it. A counterfeit is a ledge, and it knows it is a ledge, and so it manufactures urgency — decide today, the door is closing, consult no one. Whatever cannot survive a willingness to walk away from it was not going to hold weight when sat down in.

6.6 The audit as a procedure

The three instruments and the auxiliary compose into a gate sequence, performed before anything is accepted, signed, or celebrated (Figure 5). The source's procedure has ten movements and they are reproduced here in operational form because later papers refer to their outputs.

1. **Stand down the hour.** Decide nothing for the first hour after recognition. Note what happens when this is said aloud: a basin absorbs the sentence without a ripple, a ledge negotiates.
2. **Retrieve the record before revisiting the feeling.** Do not rehearse the recognition in memory — every rehearsal edits. Lay out every dated page bearing on the arrival, in date order.
3. **Write the arrival as it is.** On a fresh dated card, a plain inventory of what is actually present, in the order encountered. A manifest, not a review.
4. **Rule two columns: structure and inventory.** Enter every structural agreement in the left; every inventory disagreement in the right, in both directions — items the record specified that the outcome contradicts, and items the outcome contains that the record never imagined.
5. **Count the honest errors, and require three.** From the right column, count the details the record got flatly wrong. Fewer than three: stop. The remaining movements cannot be trusted until it is known whether the record or the outcome was tailored.
6. **Plot the corrections.** List each revision in order with its rough size. Look for the shrinking sequence. If the trail shows wandering and then one large convenient leap, name the leap aloud. Leaps are not forbidden — luck exists — but a leap must be called a leap

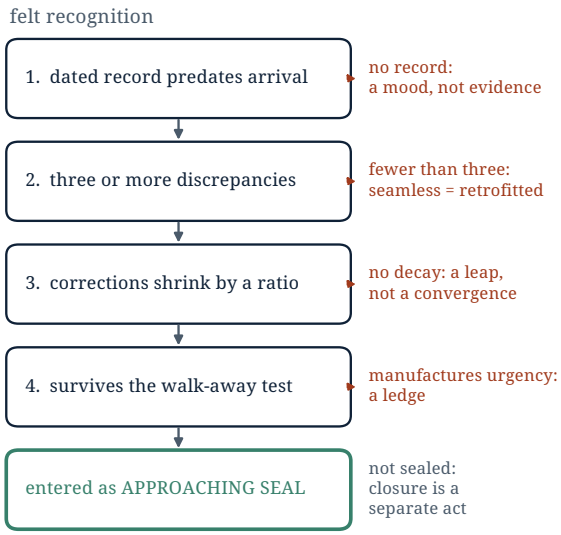


Figure 5.5: The audit of arrival as a gate sequence. Schematic; no plotted quantity. The gates are the source's three instruments in the source's order, followed by the walk-away auxiliary; any gate that fails ends the audit. The terminal verdict is "approaching seal" rather than "sealed", because closure is a separate act with its own conditions (Section 8). Script: `figs/src/CCP-A05_fig05.py`.

and lived in provisionally, because there is no evidence it will hold when the loop next passes through it.

7. **Audit the expense.** Write the name or station of every person whose position changes on acceptance. Beside each: does this diminish them, and has anyone asked them? For every unverified kindness, the audit is suspended until the kindness is priced from the recipient's side.
8. **Fork the offer.** Draft two versions of acceptance in writing — the one offered and the one the record and the expense audit actually describe. Where they differ, prepare the counter in a single present-tense sentence.
9. **Sit the arrival.** Where it has a physical place, spend twenty minutes in it doing nothing. This measures the one quantity paperwork cannot: whether the comfort deepens under examination or requires continued motion.
10. **Enter and mark.** If the audit stands, enter the arrival in the record in the present tense, date it, and mark the undertaking **approaching seal**. Not sealed. Arrival ends the journey; it does not close the undertaking.

The procedure reads long and runs short, and it earns its hour in one currency, which is not certainty — the framework has none to sell — but **auditability**: the difference between a practice that can answer for its arrivals and one that can only insist on them. The general principle is the security engineer's: a system's integrity should not depend on any part of it being hidden, including from its own operator [Kerckhoffs, 1883, Saltzer and Schroeder, 1975].

5.0.7 7. The externality result (R8)

There remains a counterfeit that passes every instrument above — dated pages, honest residue, shrinking corrections — and is still not the operator's.

A fixed point reached at another's expense is not a fixed point. Not as sentiment; as engineering. Consider what has actually been built if the arrival requires an emptied chair. The displaced party is now a term in the loop. Their loss, their resistance, their unconsented re-

arrangement do not exit the system when the paperwork closes; they return through W on every pass, as friction, as grievance, as a standing vigilance carried up the stairs every morning. What has been found is not a state the loop returns to; it is a state the loop must be **held** in, and holding is the opposite of arrival. The gain exported into another person's life re-enters as k creeping upward, and a basin with rising gain is a basin becoming a ledge.

Stated in the loop's own terms: an unconsenting party is a **second modeller** inside W, whose model of the operator now contains the operator's designs on them. The composition of two modellers each modelling the other is not a contraction; it is an arms race of models, with gain compounding above one. This is R8, and it is H4's mechanism.

The structural literature says the same thing without the vocabulary. Adding components inside one coupling perimeter increases interactive complexity, and beyond a point makes failure a property of the system rather than an error by an operator [Perrow, 1984]; added redundancy can reduce safety outright [Sagan, 1993]; and cohesion's benefit is an inverted U rather than a monotone gain [Uzzi and Spiro, 2005]. Durable arrangements over shared resources are governed by rules the affected parties had a hand in making, and that finding is empirical rather than postulated [Ostrom, 1990]. Reciprocity sustains cooperation only where the parties expect to meet again [Axelrod, 1984], and trust is a property of the relation rather than a disposition of either party [Gambetta, 1988]. Contact across difference produces value under conditions and not unconditionally [Allport, 1954].

7.1 What kind of claim R8 is: identification, and the framework turned on itself

R8 is a statement about a loop, and a statement about a loop invites the question of whether the loop's parameters could be recovered from observation. They could not, and this is stated here rather than in a footnote because it fixes what R8 is.

Written as an aggregate relation, the reflexive loop this corpus is built on — $T(x) = W(x, M(x))$, fixed in the founding paper of the series [Curzi, 2026] — is a linear-in-means reflexive specification, and the endogenous effect in such a specification is **not identified from observational data**: it cannot be separated from contextual effects or from correlated unobservables no matter how much data is gathered [Manski, 1993]. The network route is closed by the same family of result, since influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011]. The available remedies are randomisation, an instrument, partial-population variation, or a defended stipulation of the reference group, the last being the realistic route and having an explicit template in a discrete-choice model that states the conditions under which its interaction parameter is econometrically identified [Brock and Durlauf, 2001]. **This paper adopts none of them.** R8 is therefore a modelling position — a mechanism stated in the loop's own terms — and not an identified empirical result, and H4 is the design that would be needed to make it one.

That reflexivity of this shape occurs at all is separately documented, and the documentation matters because it is the existence proof the framework otherwise lacks. When the Black–Scholes option-pricing model was published its prices did not fit the market; after traders adopted it, the market came to fit the model, with dates and price data recording the convergence [MacKenzie and Millo, 2003]. That is a reflexive fixed point observed in the field. It establishes that the mechanism is real; it does not establish that any particular loop is one, and this paper keeps the two apart.

The same reflexivity falls on this paper's own instrument. The gauge of Section 3.3 reads a modulus q off a series of corrections. Parameters estimated in a population that does not know the model cannot be assumed to survive in a population that has read it, because the decision rules those parameters summarise are themselves responses to the environment the model describes [Lucas, 1976]. A framework that asserts observation changes behaviour and then exempts its own coefficients is not self-consistent. The position taken here, stated so that it can be held against the paper: **no stability of the framework's parameters under its own dissemination is claimed, and no result in this paper depends on any.** The costs are accepted — a q estimated from an operator who has read this paper is not comparable to one

estimated from an operator who has not, and a before-and-after design on a single operator is confounded by construction — and the gauge is offered as a bound that an operator may apply to their own record, not as a coefficient that generalises.

The audit question that follows is asked at the threshold, while it is still cheap: *who else is in this arithmetic, and did they consent to be?* If the answer requires a kindness no one has verified with its recipient, the answer is no — and the genuine arrival, the one the dated record actually describes, is usually adjacent to the counterfeit and one refusal away.

5.0.8 8. Closure (R7)

8.1 The cost of not closing

The source opens its treatment of closure with an empirical claim: that an interrupted task is remembered better than a completed one, that unfinished business is *staffed* rather than archived, and that the intrusions of an open goal quiet down not when the goal is achieved but when it is concretely disposed of — achieved, planned to the point of trust, or genuinely renounced. The mind, on this account, does not demand completion; it demands a verdict, and will accept a written *no* from an honest court.

This corpus's external bibliography contains no entry for that literature. The claim is reported as the framework's and is uncited (Section 12.3). What follows from it does not depend on it: the closure discipline is a bookkeeping rule, and a bookkeeping rule can be justified by what it makes auditable regardless of whether the attentional mechanism is as described.

The framework's own summary is worth preserving as a statement of the design problem it is solving: a commitment that is never closed does not become neutral, it becomes a standing charge. An undertaking that arrived and was never marked fulfilled keeps campaigning decades into possession. An undertaking that died and was never marked lapsed stands in the record like an unpaid invoice. An undertaking abandoned without words is billed nightly.

8.2 Three verdicts, and no fourth

Every commitment receives exactly one of three verdicts, entered deliberately and in writing rather than left to erosion (D2).

Fulfilled. The undertaking arrived and survived the audit of Section 6. The closing is a single dated line in the record, in the present tense: the sentence rewritten one final time as plain fact, without the forward lean it needed while it was doing gradient work. The card is then turned face down — not to hide it, but because a closed undertaking becomes ground, and ground is not for reading.

Released. The undertaking is one the operator no longer chooses — outgrown, superseded, or revealed by its own returned differences to have been someone else's sentence mistaken for one's own. Release is renunciation with the dignity of a verdict, and it is spoken once aloud, because the loop was opened by speech and the ledger honours symmetry. Release is not failure: a practice that cannot renounce can only accumulate.

Lapsed. The undertaking died unattended — starved of iteration, obsoleted by arrival elsewhere, or simply outlived. The closing is a written death date, honestly guessed, and one brief line. Brevity is the sincerity of the entry; a long apology to a dead undertaking is a way of keeping it open.

One rule completes the discipline: **one open undertaking at a time.** One commitment in the forward channel doing gradient work, and everything else current, closed, and filed. The framework's justification is capacity, not asceticism: three simultaneous undertakings run through one attentional system interfere, each returned difference arriving garbled by the others and each correction smearing across all three. The general form of that argument — that a regulator must have at least as much variety as the system it regulates, and that a

channel has a capacity — is old and standard [Ashby, 1956]. Whether one is the right number is not established, and Section 12.5 says so.

8.3 The consent law

The framework states one constraint as law rather than technique:

> **No future is to be spoken over another's path without their word.**

The argument is structural rather than moral, and this paper reports it as structural because that is how it is made. Stripped of ceremony, the undertaking is a seen future, made vivid, installed where a mind will keep encountering it, so that many small choices tilt toward it by gradient rather than by effort. Performed on oneself, with one's own record, this is the framework's whole subject: the seeing is one's own, the gradient is one's own, and the arrival is one's own to audit and to refuse. Performed on another person without their word, *the identical mechanism* is coercion — worked above the awareness of its target, in the one channel where a person cannot defend themselves without becoming someone who trusts nothing, including encouragement.

The framework adds that the law is self-enforcing because the mechanism also **fails**, by the amplifier of Section 4.2. That limb rests on the reactance claim and is therefore uncited here (Section 4.3, Section 12.3). What is not in doubt, and what the paper does assert, is R8: an unconsenting party is a second modeller inside W, and the composition is not a contraction.

Consent transforms the mechanism completely, which is why the law is a law of *words* and not a counsel of abstinence. A future spoken *with* another — asked for, chosen, co-authored, revocable — is not an installation but a shared seeing, and the difference between the two is one conversation.

The law's reach is stated so that it is not signed narrowly. It covers the obvious cases. It equally covers the benevolent ones, and this is where practitioners fail: the prophecy over a child's talent, the confident assurance about someone else's marriage, the diagnosis by faith spoken over a sick friend. **Kindness is not consent. Accuracy is not consent. Being right is not consent.** The offence is never in the content; it is in the door. In practice the discipline is small: before any undertaking that touches another life, ask, aloud, and honour the answer — including the later answer, when it changes, because consent is a standing arrangement and not a receipt.

8.4 Three conditions barring operation

The framework names three conditions under which the procedure is not to be run at all, on the ground that they are states in which the loop gain is not a constant of the operator's character but a property of the current system.

Grief, in its acute form, re-prices everything hour to hour, so that no weighting made inside it is stable enough to build on. The futures seen from inside acute grief are vivid — the framework notes that grief is a powerful extrapolation engine — and they are not the operator's to run, because the one who states the sentence in February is not the one who must live in it in June. The prescription is to run the locating disciplines only.

Rage collapses the time horizon to the width of the injury and then demands a prediction from inside the collapse. The framework's observation is specific and worth preserving: a working spoken in rage is almost always a prophecy over the offender's path, unconsented, dressed as self-improvement, and therefore the commonest violation of the consent law committed by people who would never violate it knowingly.

Revenge, the cold form, mimics the practice perfectly — patient, iterative, detail-loving, willing to keep records for years — and is distinguished by exactly one feature: its fixed point

contains another person's diminishment as a load-bearing element. Remove the diminishment and the state collapses. That is the test, and by R8 the pursuit runs above unit gain by construction, because the adversary is a second modeller.

The framework's rule for knowing when the condition has passed is that the operator does not vote on it from inside, since these are precisely the states in which instruments about state cannot be trusted; the objective signs are the ones the practice already records.

5.0.9 9. Results

9.1 Established by derivation. R1 and R2 are established. The theorem is stated exactly and its uniqueness clause is argued in Section 3.1; the a posteriori bound is its standard corollary and the geometric-series form is checked arithmetically in Figure 3. R3 and R4 follow from the same inequality read at $q = 1$ and $q > 1$ and are established in the same sense. **All four are established for the mathematical object and not for any human loop**, for the reason recorded as P1.

9.2 Established by construction. R5, R6, R7 and R8 are specifications rather than findings. R5 identifies a real mathematical condition with a real practical failure mode and gives an external check for it; the identification is argued, not measured. R6 specifies an audit; R7 specifies a taxonomy and two rules; R8 states a mechanism in the loop's own terms. None is a measurement.

9.3 Reproduction of the source's arithmetic. Two of the source's own numerical claims were recomputed and both hold. The square-root iteration reproduces the printed iterates $1,\ 1.5,\ 1.41666\ldots,\ 1.4142156\ldots,\ 1.4142135\ldots$ exactly, with absolute error below 10^{-6} within four passes (Figure 2; the script asserts the agreement and fails if it is lost). The source's stair of corrections, at bar lengths 16, 8, 4, 2, 1 and 0, sums to 31, which is below $2 \times 16 = 32$, so its claim that "the whole stair, laid end to end, is shorter than twice its first step" is arithmetically correct (Figure 4).

9.4 Hypotheses tested. None. H1 through H4 are stated in falsifiable form and none is tested here; no dataset was collected. This is reported in the same plain language a positive result would receive. H3 in particular is testable today against existing preregistration archives and was not tested.

9.5 Claims retained without support. Three of the source's mechanisms are retained in the text and could not be cited from this corpus's bibliography: reactance, the attentional cost of unclosed commitments, and the shared-machinery account of imagination and memory. Each is marked at its point of use and listed in Section 12.3. In each case the paper's own conclusion was restated so that it does not depend on the unsupported mechanism.

5.0.10 10. Discussion

10.1 What follows

The paper's practical yield is one instrument and one refusal.

The instrument is the gauge of Section 3.3. It is unusual among instruments of self-assessment in requiring no judgement about the destination and no estimate of the distance — only a measurement of the last revision, which is already recorded. Its value is that it converts a question that cannot be answered (*am I nearly there*) into one that can (*how big was my last change*), and it does so with a bound rather than an encouragement.

The refusal is the audit of Section 6. It denies the operator the one certificate they most want to accept, and it does so for a reason that survives even if the source's proposed mechanism is wrong: no subjective report of fit is evidence about a prior prediction. The audit's three instruments are all documentary, and that is deliberate.

Between them there is a general lesson about instruments in this corpus. Both are built on the same principle: prefer a quantity recorded before the outcome to one produced after

it, and prefer a quantity that can be re-derived to one that must be believed.

10.2 What does not follow

That a loop has a fixed point does not mean it has been shown to contract; that corrections are shrinking does not mean the target is reachable, because completeness can fail while contraction holds (Section 5); and that an arrival passes the audit does not mean it is legitimate, because the expense audit is a separate gate (Section 7).

Nor does the theorem license the reassurance most often taken from it. “From anywhere” is a clause about the *mathematics*, and applying it to a life is P1. This paper’s position is that the transposition is useful precisely because it can fail inspection: it names the two conditions — contraction and completeness — whose failure produces the two failure modes practitioners actually report.

10.3 Relation to the corpus

The definitions this paper consumes are fixed upstream: the Banach condition $d(T(x), T(y)) \leq k\, d(x, y)$ with $0 \leq k < 1$, and its three consequences, are stated in the founding paper of the series [Curzi, 2026], and this paper’s R1 and R2 are their extension into a working instrument. The corpus’s reference apparatus places the established-mathematics component of the surplus functional — the restatement of $f(u) = \ln(1 + (N-1)u)$ as the logarithm of an affine effective count and as a lower bound on the ensemble’s entropy — in its established drawer while assigning the application to a life to the postulate drawer [Curzi, 2026]; that three-drawer discipline is the same discipline this paper applies to P1.

Downstream and laterally, four results bound what this paper may claim. The corpus’s own preregistered measurement of its storage law failed its second falsification rule: a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin, notwithstanding a 40.3 per cent advantage over persistence, 97.7 per cent over direct input, and $p = 0.0005$ against a shuffle control [Curzi, 2026i]. That is the corpus’s strongest available evidence that a contraction-shaped model can fit well and still fail a properly specified test, and it is the correct prior for every claim in this paper. The corpus’s audit of a traditional operation set against a parameter model returned seven mappings, two splits, one rendering and three rejections over thirteen entries, with five entries carrying a failure [Curzi, 2026] — an example of the same refusal discipline applied to inherited material. The corpus’s sonification paper states an open problem formally rather than claiming a result: no established mechanism produces a transcriptional response depending on a symbol string as an *order*, as distinct from permutation-invariant functionals of the waveform [Curzi, 2026]. And the corpus’s own record of what its architecture cannot do reports that disclosure did not confer immunity in the strongest available test case, that a backing basket was 4.61 times less diversified than its component count claimed, that a terminal programme phase was a 0.090 option rather than a plan, and that a levy presented as a costless accounting correction in fact reallocated burden, 42.0 per cent of it onto one sector with a constructed 55 per cent passing through to consumers [Curzi, 2026]. That paper bounds this one: it is the corpus demonstrating that a well-specified structure can still fail in the world, which is exactly what P1 leaves open.

Two identity results in the corpus are instruments of the same kind as the audit, and are cited here because the audit’s second and third instruments are documentary integrity controls in all but name. The canonical form $\text{canon}(x) = \text{restrict_ASCII}(\text{confusable_fold}(\text{NFKC}(x)))$ collapses 495 of 495 manipulated entries onto the clean form, is idempotent across all 500 cases, and resolves 995 raw byte strings to the true 500 identifiers [Curzi, 2026]. And the channel-separation principle — that in-band architectures admit an impersonation class no care within the shared channel can close, while out-of-band separation closes it structurally, measured at 0.383 per cent in-band ambiguity against 0.000 per cent out-of-band over 100,000 values [Curzi, 2026] — is the general form

of the audit's first instrument: a record produced inside the channel it is meant to certify cannot certify it.

Finally, the corpus's belief postulate is the template this paper's P1 follows: the magnitude form $E = G \cdot B + \varepsilon$ with the forbidding clause $G = 0 \Rightarrow E = 0$ for all B, badged as a postulate and paired by explicit correspondence with the storage law's $\eta = 0$ [Curzi, 2026ab]. A strong claim, stated with the clause that forbids its abuse, and badged.

5.0.11 11. Verification

This section is not a summary. It decomposes the paper's results into steps small enough to check one at a time, and labels each `ARITHMETIC`, `DEFINITION`, `CITED` (with the key) or `POSTULATED`. The razor is applied: any postulate whose removal leaves the results standing has been removed.

11.1 Postulates

- **P1 (DEFINITION).** (X, d) is a metric space: $d \geq 0$, $d(x, y) = 0$ iff $x = y$, symmetry, triangle inequality. *Check:* standard.
- **P2 (DEFINITION).** (X, d) is complete: every Cauchy sequence in X converges in X. *Check:* standard.
- **P3 (DEFINITION).** T is a contraction with modulus q: for all x, y, $d(T(x), T(y)) \leq q\, d(x, y)$, $q < 1$. *Check:* D1.
- **P4 (CITED: `banach1922operations`).** Under P1–P3, T has exactly one fixed point, reached from any start, with $d(x_n, x^*) \leq q^n d(x_0, x^*)$.
- **P5 (ARITHMETIC).** For $0 \leq q < 1$, $\sum_{j \geq 0} q^j = 1/(1 - q)$. *Check:* the geometric series; redo it.
- **P6 (POSTULATED).** A life admits a state space satisfying P1 and P2, and a reflexive loop satisfying P3. *Check:* **not checkable and not checked. This is the premise the whole transfer rests on and it is stated as a postulate everywhere it is used.**
- **P7 (POSTULATED).** A subjective feeling of recognition carries no information about whether a prediction preceded an outcome. *Check:* asserted; supported obliquely by `nisbett1977telling` and `schwarz1999self`, neither of which states it.
- **P8 (POSTULATED).** A prospectively made record has independent opportunities to be specifically wrong that a retrospective reconstruction does not. *Check:* asserted; it is the content of H3 and it is testable and untested.
- **P9 (POSTULATED).** An unconsenting party inside W acts as a second modeller whose model contains the operator's designs. *Check:* asserted; this is R8's mechanism and it has no supporting entry.

Postulates dropped as unnecessary: the reactance mechanism (Section 4.3), the attentional cost of unclosed commitments (Section 8.1), and the shared-machinery account of imagination and memory (Section 6.1). Every result below survives their removal, which is why they are reported rather than used.

11.2 The compounding chain

```
P1 + P2 + P3     -> L1   the hypotheses of the theorem hold       DEFINITION
L1 + P4          -> R1   one fixed point, reached from anywhere,
                         geometrically                           CITED
R1 + P5          -> L2   sum of remaining steps <= last/(1-q)    ARITHMETIC
L2               -> R2   distance to arrival <= q/(1-q) times
                         the last correction                     ARITHMETIC
P3 at q = 1      -> R3a  isometry: distance preserved, no
                         theorem, no arrival                     ARITHMETIC
P3 at q > 1      -> R3b  amplifier: distance grows               ARITHMETIC
```

```
R3a + R3b        -> R4   a smaller-than-gap correction may be
                         wrong and converge; a larger-than-gap
                         correction may be right and diverge      ARITHMETIC
P2 negated       -> R5   contraction toward a hole: shrinking
                         gap, no limit in the space               ARITHMETIC
P7 + P8          -> R6   arrival is certified by dated record,
                         >= 3 residues, and a decaying ratio      POSTULATED
(taxonomy)       -> R7   three verdicts, one open undertaking,
                         the consent law                          DEFINITION
P9               -> R8   an exported cost re-enters as gain       POSTULATED
R1 + P6          -> C1   a well-built undertaking arrives         POSTULATED (via P6)
```

What the chain shows. R1 is `CITED`; R2, R3, R4 and R5 are `ARITHMETIC` on top of it. That block — the paper's mathematical content — is checkable end to end without believing anything about people. Everything about *practice* is `POSTULATED`: R6 rests on P7 and P8, R8 rests on P9, and C1 — the conclusion an operator wants — rests on P6, which is not checkable.

A chain whose steps are mostly `DEFINITION` and `POSTULATED` is not a proof. This one is unusual in the corpus for having a genuinely `CITED`-plus-`ARITHMETIC` spine; it is entirely ordinary in that everything the spine is *for* sits on postulates. The value of laying it out this way is that it names the exact bill: P6, P7, P8 and P9 are four claims, three of them testable, and testing them would settle the framework's practical status. That is a smaller bill than the framework's rhetoric implies.

11.3 Dimensional audit

Symbol	Meaning	Dimension	Unit	Status
x, y	states	none assigned	none	defined; indices into a state set
$d(x, y)$	distance between states	$[x]$	**no operational unit exists**	assumed to exist; never constructed for a life
q, k	contraction modulus	dimensionless	—	defined; **never measured for a human loop**
n	pass index	dimensionless	passes	defined
e_n	distance to fixed point at pass n	$[x]$	none	derived from d
x^*	fixed point	$[x]$	none	defined by $T(x^*) = x^*$
$q/(1-q)$	bound multiplier	dimensionless	—	derived, arithmetic
δ	storage decay rate (quoted)	dimensionless	per pass	quoted from `ccp:CCP-007`; not measured here
$\tau = 10.6033$	time constant (quoted)	time	passes	quoted from `ccp:CCP-047`; a derived quantity there
bar lengths 16, 8, 4, 2, 1, 0	correction sizes in the source figure	dimensionless	characters	measured from the source's own drawing

Does the central relation balance? $d(x_n, x^*) \le \frac{k^n}{1-k} d(x_0, x_1)$ balances trivially: k is dimensionless and both sides carry $[x]$. **That the relation balances is not evidence that $[x]$ exists.** The source is silent on the units of the distance between two states of a life, and this paper does not supply them; the honest entry is that there is none. Where the corpus's storage instance is quoted, its δ is dimensionless per pass and its τ is in passes, and both are that paper's quantities rather than this one's.

Scale type. The bound of Section 3.3 is a statement about *ratios of distances* and therefore presupposes at least an interval scale, and in the form "the remaining distance is at most $q/(1-q)$ times the last correction" it presupposes a ratio scale. Nothing in the source establishes either for any psychological quantity. The formal test for whether an attribute has additive quantitative structure is conjoint measurement [Luce and Tukey, 1964], against the standard set out by Krantz et al. [1971], and **its conditions were not tested.** The strongest objection to the whole framework is that psychology assumes quantitative structure without ever testing for it [Michell, 1997]; the correct response available and unused is either a conjoint demonstration or a Rasch-model one [Rasch, 1960], with the construct-validation programme that would accompany it [Cronbach and Meehl, 1955, Campbell and Fiske, 1959]. **Where an operator's corrections are counted in something with a natural unit — hours, money, words struck, letters sent — the ratio reading is licensed by that unit and not by this framework. Where they are not, it is not licensed at all.**

11.4 Limits and boundary behaviour

- $q \to 0^+$**.** One pass arrives. The bound $q/(1-q) \to 0$: the last correction is the whole remaining distance and there is none left. Degenerate but coherent.
- $q \to 1^-$**.** Convergence survives and utility does not: $1/(1-q) \to \infty$, so the total road diverges and the number of passes to any fixed accuracy grows without bound (Figure 3). The theorem does not fail here; it stops being useful, and the two are worth distinguishing.
- $q = 1$**.** No theorem. The map is an isometry; the orbit neither approaches nor recedes (Figure 1). Note that the fixed point may still *exist* — the identity map fixes everything — but nothing is guaranteed and nothing converges.
- $q > 1$**.** The bound reverses and iteration is a divergence generator. There is no contradiction with the existence of a fixed point; there is simply no route to it by iteration.
- $n \to \infty$ **with completeness failing.** $e_n \to 0$ and x_n converges to nothing in the space (Section 5). Every internal criterion reports success. This is the only boundary in this paper that is invisible from inside the loop.
- **First correction** $\to 0$**.** The a posteriori bound collapses to zero remaining distance, which is correct and useless: an operator who never corrects at all reads as arrived. The gauge presupposes that corrections are being taken from returned differences, and this degenerate case is exactly the orbit of Section 4.1 wearing the gauge's clothes.
- **Number of open undertakings** $\to \infty$ **(Section 8.2).** The framework asserts interference; nothing here bounds it, and the rule of one is a stipulation.

11.5 The checks a reader should run first

1. **Redo the uniqueness argument** in Section 3.1. Three lines, no source.
2. **Redo the geometric sum** in Section 3.3 for $q = 2/3$: $1/(1-2/3) = 3$. Figure 3 should agree, and the script asserts it.
3. **Run `figs/src/CCP-A05_fig02.py`.** It recomputes the source's square-root iterates and asserts they match the printed digits; if the assertion is ever removed, the figure is no longer a check.
4. **Add up Figure 4's bars:** $16 + 8 + 4 + 2 + 1 + 0 = 31 < 32$. The source's claim is arithmetic and it holds.

5. **Look for a citation attached to reactance.** There is none, in this paper or anywhere in the corpus's external bibliography, and Sections 4.3 and 12.3 say so.

5.0.12 12. Limitations

12.1 The transfer premises are unverified, and this is the principal limitation. Neither completeness nor a uniform contraction modulus has been established for any human loop, by this corpus or by anyone. Everything in Sections 3 and 4 is exact about a mathematical object and conditional about a life. **Falsified by:** nothing, which is the point — P6 is not currently a falsifiable claim, and the honest description of its role is that of a hard-core commitment protected by a belt of testable claims [Lakatos, 1970].

12.2 No hypothesis was tested and no data were collected. H1–H4 stand. H3 in particular is testable against existing preregistration archives with no new data collection, and testing it is the single most valuable next step this paper can name. The discipline that would make such a test credible is specified and not exercised [Nosek et al., 2018, Simmons et al., 2011].

12.3 Claims with no supporting literature in this corpus's bibliography. Listed individually, because a reviewer will find them.

- *Pressure on a perceived freedom produces counter-movement (reactance).* Load-bearing for the human limb of R3 and for the self-enforcing reading of the consent law. **`bib/external.bib` contains no entry.** Retained uncited; no citation was attached to make it appear supported. Where the paper's conclusion could be restated without it, it has been (Sections 4.3, 8.3).
- *Unfinished commitments hold attention and levy a continuing cost until concretely disposed of.* The premise of Section 8.1. **No entry.** Retained uncited, and the closure discipline is justified on bookkeeping grounds instead.
- *Imagination and memory run on the same constructive machinery, so a rehearsed future consolidates as a memory does and a present outcome reshapes the remembered expectation.* The proposed mechanism for Section 6.1. **No entry.** The paper's conclusion — that felt recognition certifies nothing — is restated so that it rests on the self-report literature instead [Nisbett and Wilson, 1977, Schwarz, 1999, Kahneman et al., 2006], which supports a weaker and sufficient claim.
- *The specific numerical thresholds*: three preserved discrepancies, one open undertaking, a correction no larger than one third of the gap, twenty minutes of motionless sitting. All are stipulations. **No entry supports any of them**, and none is presented as an estimate.

12.4 The gauge is exposed to being gamed by the operator it serves. An operator who knows that shrinking corrections certify nearness can produce shrinking corrections. This is the standard corruption of an indicator adopted as a target [Campbell, 1979, Goodhart, 1975], and it applies with particular force here because the operator is simultaneously the instrument, the subject and the auditor. The documentary instruments of Section 6 mitigate but do not remove it.

12.5 Parameters chosen rather than measured. Every parameter in every figure was chosen. $q \in \{0.60, 1.00, 1.40\}$ in Figure 1 is chosen to place one case in each regime. The annotated cases in Figure 3 are the source's, but the curve is an identity rather than a fit. Figure 4's bar lengths are read off the source's own drawing, which is itself an illustration and not a record of any operator's corrections. Only Figure 2 contains numbers that were independently recomputed and checked against a stated source claim.

12.6 The scale type is stipulated. Section 11.3 records that the ratio reading of the a posteriori bound is licensed only where corrections are counted in a quantity with a natural unit. Where the framework is applied to quantities without one, the bound retains its form and loses its interpretation, and no result in this paper survives that loss.

12.7 The audit's discriminating power is asserted, not measured. H3's threshold of three is a stipulation. Worse, the audit is administered by the party with the strongest interest in

its outcome, and no inter-rater reliability, blinding, or independent adjudication is specified anywhere in the framework. A record's discrepancies are counted by the person who wants them counted.

12.8 The externality result is stated as engineering and demonstrated as neither. R8's mechanism (P9) is asserted. Its structural cousins in the literature [Perrow, 1984, Sagan, 1993, Uzzi and Spiro, 2005] concern coupled technical and organisational systems, and their transfer to a two-person loop is an analogy that this paper does not attempt to convert into a derivation.

12.9 Case material is illustrative. The two worked cases in Section 13 are single, self-reported, unblinded, retrospective accounts drawn from the source manuscript. They illustrate what the procedures return and establish nothing about whether the procedures work. The base rate against which any single such report should be read is not encouraging [Ioannidis, 2005, Open Science Collaboration, 2015].

12.10 A scope note on borrowed constants. Where this paper cites corpus results that carry the constant $k = 0.483425732988414$ Hz — the minimum sustainable frequency and its associated deficit identity — two cautions apply and are recorded rather than repeated at each mention. That constant is *defined*, exact by construction, and its fifteen figures are defensible only on that basis; attaching hertz to it is a separate and much stronger claim than computing it. And the general hazard of reading a chosen ratio as a discovered natural constant is documented and should be read before any such reading is offered [Markowsky, 1992]. This paper makes no use of that constant and endorses neither step.

12.11 The framework's own most useful instrument is documentary, and documents can be destroyed. The audit rests entirely on records the operator keeps. Nothing in the framework addresses record loss, selective retention, or the operator who begins keeping records only once an outcome is in view. A missing record is indistinguishable from an absent prediction, which is the correct default and an uncomfortable one.

5.0.13 13. Case material

Two worked cases appear in the source in enough detail to report. Both are self-reported and neither is evidence; they are included because they show what the procedures return and because both contain a refusal, which is the part of a practice that is usually omitted from its own advertising.

A surveyor's archive. An operator designated Anselm, a land surveyor of forty years, held an archive of a parish's boundaries, drains and shorelines. His target, written in the fewest words it would survive, was *the parish's memory is in the parish's hands*. His first attempt — a letter offering the entire archive to a historical society — returned polite alarm: no fireproof room, no cataloguer, no funds. Read as refusal that is a rejection; read as a sounding it is a completeness failure of exactly the kind Section 5 describes. The target was fine; the receiving space could not hold the limit.

His subsequent corrections were small and are recorded one line at a time: offer ten sheets rather than four hundred; scan before donating; and then the correction he attributed to the practice rather than to himself — *they do not want the maps, they want to be able to ask the maps questions*, followed by attendance on Thursday evenings to answer them. Twice the alternation test caught him: after a favourable newspaper mention he tripled his hours, flooded the reading room, exhausted himself and his welcome, and the next three entries alternated add, cut, add, at which point the rule halved his step and held him at two evenings. Once he nearly rewrote the target itself — from the modest sentence to *a proper archive built*, with a building and a committee — and the rule that a target changes only as a declared new undertaking stopped him. His own gloss is worth preserving because it is the stability rule arriving from a different trade: a surveyor closing a traverse never dumps the whole misclosure into the last leg but distributes it, a sliver to every leg, so that the whole traverse trues at once and no single line bears it. His seasonal completeness check found a hole he had been decorating rather than measuring — the plan still quietly required him to live for

ever — and the correction that closed it was the smallest and hardest on the record: *teach the boy the filing system, not just the machine*. The archive's record runs four years of annotated corrections and then dates alone.

A bookbinder's workshop. An operator who had stated her target fourteen months earlier — *I bind books that outlive their readers* — ran the audit of Section 6 when the workshop she had recorded eleven months prior became available on the death of its master. Her columns: structure agreed extravagantly; inventory disagreed in four places, since her record had imagined a ground-floor shop, a street-door bell, working alone, and a different river. Four honest errors, so the record passed the second instrument. Her corrections plotted as a clean shrinking series with no leap, so it passed the third. The expense card had a single line on it and that line did not consent: the master's widow, asked directly and with the landlord absent, wished to remain in the flat above the shop, and said so. The fork was written and offered — take the workshop and the shop-floor lease only, the flat stays hers — and was refused for nine days, which she spent motionless, the walk-away test administered by circumstance. On the tenth day it was accepted. She sat the empty workshop for twenty minutes and reported that the comfort deepened under examination. The undertaking was marked *approaching seal*.

What the second case shows is the paper's own claim about counterfeits: the genuine arrival and the counterfeit were structurally identical and one refusal apart. What neither case shows is that the procedure caused anything, and neither is offered as though it did.

5.0.14 14. Conclusion

The eight results are stated below in the form the rest of the corpus should quote.

R1. A contraction with modulus $q < 1$ on a complete metric space has exactly one fixed point; it is reached from any starting state; and the distance to it shrinks by at least the factor q on every pass.

R2. The a posteriori bound is $d(x_n, x^*) \le \frac{k^n}{1-k}\, d(x_0, x_1)$, read operationally as: the distance remaining to arrival is at most $q/(1-q)$ times the size of the last correction. The size of one's most recent revision is therefore a self-reading gauge.

R3. The honest converse: gain exactly one gives an isometry, which orbits without arriving; gain at or above one gives an amplifier, whose corrections — sized to delayed feedback — arrive as new errors.

R4. A correction smaller than the observed difference can be wrong and still converge; a correction larger than the difference can be right and still diverge. Restraint is a stability criterion, not a virtue.

R5. Completeness read as a failure condition: a foreseen state the life cannot contain is a hole in the space, and iteration toward it contracts indefinitely while arriving never. This is the one failure mode invisible from inside the loop.

R6. Arrival is certified by audit and not by feeling, on three instruments: a dated record predating arrival, at least three preserved discrepancies between record and outcome, and a geometrically decaying sequence of corrections. A seamless match indicates a retrofitted record.

R7. The closure discipline: exactly three verdicts — fulfilled, released, lapsed — entered in writing; one open undertaking at a time; and a consent law forbidding a future spoken over another's path without their word.

R8. The externality result, stated as engineering: a fixed point requiring another party's displacement re-enters the loop as friction, raising the gain and converting a basin into a position that must be held.

Counter-result. This paper refuses the certificate its own theorem appears to offer. A human life is not a metric space; nobody has measured the relevant contraction modulus; and the theorem therefore transfers as shape and not as guarantee. The feeling of recognition on arrival proves nothing, because the machinery that rehearses a future retrofits a past with equal fluency.

What is established is a theorem, its corollary, its converse, and a specification for auditing a claim of arrival. What is not established is that any human loop satisfies the theorem's premises, or that the audit discriminates as claimed. The paper's contribution is to name the four claims — P6 through P9 — whose settlement would decide the matter, and to note that three of the four are testable now.

5.0.15 15. Declarations

Data and code availability. No external dataset was collected for this paper and none exists. The case material in Section 13 is reported from the source manuscript's own record and is not a dataset. All figures are generated by the scripts named in their captions (`figs/src/CCP-A05_fig01.py` through `figs/src/CCP-A05_fig05.py`), which import the shared style module `lib/figstyle.py`. Two of those scripts contain assertions that fail the build if the numbers they plot cease to agree with the source claims they are drawn from.

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the analysis, the figures, and the text.

Ethics. No human-subjects research was conducted for this paper. No personal data were collected or processed. The case material in Section 13 is drawn from the source manuscript, is presented under the pseudonyms it carries there, and is illustrative rather than evidential. Section 8.4's three barring conditions record the framework's own instruction that the procedures are not to be operated during acute grief, rage, or the pursuit of revenge; nothing in this paper is clinical advice.

Use of the ceremonial register. The source manuscript is in part an initiatory document. This paper analyses that register as an object of study and does not reproduce it; ceremonial names are reported only where they identify a construct this paper renames.

Chapter 6

The Reference Apparatus of the Foundations Volume: The Closing Procedure, the Protocol Index, and an Audit of the Three-Class Honesty Ledger

Abstract. This paper is reference apparatus and not a research contribution. It collects the closing material of the Foundations volume in citable form: an eleven-movement closing procedure with five named failure modes, a per-construct reference card, a glossary mapping the volume's ceremonial register onto its technical equivalents, an index of thirteen protocols, a correspondence table, and the volume's own three-class classification of its claims into established mathematics (four items), established psychology (seven), and declared postulates (three). Almost all of that recapitulates results established elsewhere in the corpus, and the paper says so and gives the citation rather than the derivation. **Its one piece of new content is an audit of that classification against the corpus's own external bibliography, and the result is a count: of the eleven items for which support is claimed, five carry a bibliographic key and six do not, and two of the five carry a standing objection in the same file.** The class the volume treats as its safest — established psychology — is the class the bibliography least supports, with five of its seven items uncited. One further item is correctly filed and misleadingly labelled: the surplus form $f(u) = \ln(1 + (N-1)u)$ is exact arithmetic given a definition the framework chose, and **the affine effective count it rests on is not a Hill number of any positive order** — on the shares $(0.9, 0.1)$ it returns 2.000000 where HHI^{-1} returns 1.219512 and $\exp(H)$ returns 1.384145, matching only the order-zero member, which is the raw richness. The three postulates claim no support, find none, and are correctly filed.

Keywords: honesty appendix; postulate classification; closing procedure; protocol index; correspondence table; reference card; single active commitment; operating conditions

Series: A. Foundations **Paper:** CCP-A06

6.1 1 Introduction

This paper is reference apparatus. It is not a research contribution and should not be read as one. It collects, in a single citable place, the closing material of the Foundations volume: an eleven-movement closing procedure with five named failure modes, a per-construct reference card, a glossary of the ceremonial register against its technical equivalents, an index of the volume's thirteen protocols, a correspondence table to the tradition the volume descends from, and the volume's own three-class classification of its claims. Almost all of that material restates results established elsewhere in the corpus, and where it does, this paper says so and gives the citation rather than the derivation.

There is one thing here that is not recapitulation, and this paper's structure is arranged around it. **The volume closes with an honesty appendix that sorts every claim it makes into three labelled classes** — established mathematics, established psychology, and the framework's own postulates — and states that every protocol in the volume runs on the first two classes and steers by the third. That classification is reproduced in Section 6, in the plain register, exactly as the volume files it.

Section 7 then audits the classification against the corpus's own external bibliography, and that audit is the only new content in this paper. The result is unflattering in a specific and correctable way. The classification files fourteen items across the three classes. Of the four filed as established mathematics, three are supported by an entry in this corpus's bibliography and one has no entry at all — and one of the three is **mis-labelled**: the surplus functional is exact arithmetic given a definition the framework chose, and the definition it rests on has no supporting literature in any discipline. Of the seven filed as established psychology, **two are supported by a bibliographic key and five are not** — reactance, implementation intentions, commitment and consistency, the interrupted-task effect, and confirmation bias are each named in the volume by their standing in the literature, and none of the five has an entry in the file this corpus cites from. The three items filed as postulates are correctly filed, and the classification's own statement that the third class's wager can lose is accurate.

Stated as one sentence: the class the volume treats as its safest is the class this corpus's bibliography least supports. That is a defect of the bibliography and of the appendix together rather than of either alone, it is cheap to repair, and it is printed here because a classification that is not audited against a citation list is a classification of the author's confidence rather than of the evidence.

A second, smaller finding fell out of re-running the arithmetic rather than the citations. The corpus's concentration instrument states the same Hill-number comparison this paper needs for its one mis-labelling, and states it wrongly on two counts: a transcription slip at the fifth decimal, and the claim that *every* Hill number reports below 1.39 on the test vector, which is false because the family's order-zero member is the raw richness and equals exactly two. **The second correction sharpens the objection instead of blunting it**, and both are given in Sections 7.3 and 10.

Four further items in the collected material are worth naming as content rather than as index. **First**, the closing discipline: exactly three verdicts — fulfilled, released, lapsed — with no fourth, a rule permitting exactly one active commitment at a time, and a written three-question screen that the volume states cannot be waived on the basis of feeling well. That discipline is established analytically elsewhere in the corpus [Curzi, 2026], and Sections 4 and 5 give its operational form rather than re-deriving it. **Second**, the consent condition, stated as a constraint rather than as counsel: a future is not to be voiced over another person's path without their word. **Third**, three named conditions barring operation — acute grief, rage, and the pursuit of revenge — each with a stated mechanism. **Fourth**, the protocol index: thirteen protocols, one per construct plus a three-office daily cadence, together with one further discipline that carries no construct and closes nothing.

Two things in the source are deliberately not reproduced here. The volume is written in two braided registers, one narrative and one analytic; this paper reports the analytic register, and the volume's closing narrative and its epilogue are fictional exposition that no citation of

this paper should be understood to carry. The epilogue does contain one item of analytic interest, and it is reported in Section 8.3 as an object of study rather than reproduced: the volume closes by identifying itself with an advertisement placed inside its own narrative, which is a deliberate self-referential device and is the volume's own strongest instance of the reflexive claim the corpus argues elsewhere.

Section 2 states, in the place a hypothesis would occupy, that none is advanced, and gives the organising question instead. Section 3 gives the method by which the classification and the audit were made. Sections 4 and 5 give the closing discipline and its procedure. Section 6 reproduces the three-class classification. Section 7 reports the audit. Sections 8 to 11 discuss, verify, limit and conclude. Appendices A to E carry the reference material in full, because other papers in the corpus rely on it and it exists nowhere else in citable form.

6.2 2 Hypothesis: none is advanced, and the organising question in its place

This paper advances no testable hypothesis. It is a reference collection with one audit attached, and a hypothesis manufactured to fill this section would be worse than an honest statement that there is none.

The organising question is narrow and answerable by inspection: *for each claim the Foundations volume classifies as established, does this corpus's external bibliography contain an entry that supports it?* That question has a determinate answer for each of the fourteen classified items, the answer is given in Section 7, and it is checkable by anyone with the bibliography in front of them. It is a bibliographic audit and not an empirical test, and no result in Section 7 says anything about the world.

Three statements in the collected material *could* be given the form of hypotheses, and the volume does not test any of them. They are recorded here with what would defeat them, so that the distinction between the volume's postulates and its findings survives being indexed, and each is marked as untested:

> **The single-active-commitment rule.** That holding one commitment open at a time yields better outcomes than holding several. **Would be defeated by:** a matched comparison in which multiple-commitment practitioners show equal or better completion rates and equal or lower attentional cost. **Not tested anywhere in this corpus**, and the rule is stated in the volume as a discipline rather than as a finding.
>
> **The barring conditions.** That acute grief, rage and the pursuit of revenge each degrade the practice enough to make operation inadvisable. **Would be defeated by:** a comparison showing no degradation of outcome or of correction size under those conditions. **Not tested**; the volume supplies mechanisms and no measurement.
>
> **The consent condition.** That a future voiced over an unconsenting person is rejected with interest, and that consent transforms the mechanism rather than merely licensing it. **Would be defeated by:** matched consented and unconsented arrangements showing equal maintenance cost and equal acceptance. **Not tested**, and Section 7 records that the psychological result the volume leans on for it has no supporting entry in this bibliography.

Whether the framework itself is true is not at issue in this paper and is not tested by it. The corpus's one preregistered confirmatory application of its storage law is reported elsewhere, and it **failed its second falsification rule**, returning a 2.8 per cent held-out advantage over a matched autoregression against a declared 10 per cent margin [Curzi, 2026i]. That result, and not anything in this paper, is the current empirical standing of the framework these appendices index.

6.3 3 Method

The classification is the volume's own and is reproduced rather than constructed. Section 6 gives it in the plain register, with the volume's own wording preserved wherever the wording carries a qualification. Where the volume's ceremonial vocabulary names a construct, the construct is reported under its ceremonial name as an object of study and the technical equivalent is given beside it; Appendix B is that mapping in full. **This paper analyses a ceremonial register; it does not adopt one.**

The audit of Section 7 was performed by inspection, item by item, against a single named file. For each of the fourteen classified items, the procedure was: (i) state the claim as the volume states it; (ii) search the corpus's external bibliography for an entry that supports *that* claim, not a neighbouring one; (iii) record the key if one exists, and record its absence if none does; (iv) where an entry exists that qualifies or contradicts the claim, record that too. **No entry was created, and no claim was attached to an entry that does not support it** — an unsupported item is recorded as unsupported, which is the only outcome that carries information.

The audit's unit is the claim as filed, which matters for one item. The volume files the surplus functional under established mathematics on the ground that it is the logarithm of an affine effective count. **That statement is exact arithmetic and the audit records it as such**; what the audit also records is that the affine effective count itself is a design axiom of this framework with no supporting entry in ecology, economics, political science or information theory, so the item's *classification* is defensible and its *implicature* is not. Section 7.3 states the distinction and Section 10 records it as a limitation of the classification rather than of the arithmetic.

The decision rule is published because that is the only property that makes an audit checkable by someone other than its author. Kerckhoffs [1883] states the general principle in its oldest engineering form — a system must remain sound when everything about it except the key is public — and Saltzer and Schroeder [1975] restates it as open design. Brandeis [1914] is the same requirement as a governance argument. **Applied to a classification, the principle is that the rule by which items were sorted must be legible enough for a second reader to sort them again and disagree**, and Section 10 records that no second reader has yet done so. Nosek et al. [2018] is the modern form of the same commitment: fix the rule before the result, and publish the rule with the result.

Three conventions govern what is reported. A claim is recorded as *supported* only where a bibliography entry states the specific proposition. A claim is recorded as *qualified* where an entry supports it and another entry in the same file materially restricts it; in that case both keys are given, because a supporting citation offered without its standing objection is a misrepresentation of the file. A claim is recorded as *unsupported* where no entry states it, and no inference is drawn from that beyond the fact itself: **an unsupported claim in this audit is a claim this bibliography does not cover, which is not the same as a claim the literature does not contain.** Section 10 states that limitation plainly, since it is the audit's principal weakness.

6.4 4 The closing discipline: verdicts, consent, and operating conditions

The volume's closing construct carries six pieces of doctrine, and this section states each in the plain register with its stated mechanism and its stated caution. **The analytical content of the first two — three verdicts and no fourth, and one active commitment at a time — is established elsewhere in the corpus [Curzi, 2026]; what is given here is the doctrine as the volume states it, so that the procedure of Section 5 has something to be the procedure of.** Every claim in this section is reported, and none is argued for.

6.4.1 4.1 The cost of not closing

The volume prices non-closure before it prescribes closure, and the price is the argument for everything that follows.

The mechanism it names is the standing attentional cost of an open loop. Its statement of the underlying result is that an interrupted task is remembered better than a finished one, that attention bills an open commitment differently from a closed one, and that unfinished business is therefore *staffed* rather than archived — some fraction of the practitioner's capacity assigned indefinitely to a commitment they have stopped noticing they hold. **The sharpening the volume adds is the operative part**: the intrusions of an open goal are said to quiet not when the goal is achieved but when it is *concretely disposed of* — achieved, planned to the point of trust, or genuinely renounced. On that reading the mind does not demand completion; it demands a verdict, and will accept a refusal in writing. What it is said not to accept is silence.

That result is Class II item 5 of the classification reproduced in Section 6, and Section 7 records that it has no supporting entry in this corpus's bibliography. It is the item on which this whole section's motivation rests, and it is uncited. The disclosure belongs here rather than only in the audit, because a reader of Section 4.1 alone would otherwise take the mechanism as established.

The volume's own framing of the necessity is narrower and does not depend on the psychology. A spoken commitment is described as a tool whose usefulness is its grip — it holds attention on purpose — and **a tool defined by its grip requires a release**, on pain of the grip outliving the purpose. That is an argument from the design of the instrument rather than from a finding, and it is the form in which the claim survives the audit.

6.4.2 4.2 Three verdicts, and the single-active-commitment rule

Every commitment receives exactly one of three written verdicts, entered deliberately rather than left to erosion. There is no fourth verdict.

Fulfilled. The commitment arrived and survived the arrival audit. The closing entry is a single dated line in the present tense — the original sentence rewritten as plain fact, without the forward lean it carried while it was doing work — after which the record is turned face down. The volume's gloss on turning it down is that a closed commitment becomes ground rather than content, and that re-reading fulfilled commitments is a way of standing in the cellar admiring the joists.

Released. The commitment is one the practitioner no longer chooses — outgrown, superseded, or revealed on inspection to have been another person's sentence mistaken for one's own. The closing is a spoken formula of renunciation, then a dated mark.

Lapsed. The commitment died unattended: starved of iteration, obsoleted by an arrival elsewhere, or simply outlived. The closing is a best-guess death date and one brief line of thanks or pardon. **The volume states the brevity as a requirement rather than a style**, on the ground that a long apology to a dead commitment is itself a way of keeping it open, and it records having watched practitioners spend years drafting such apologies — mourning as the most respectable form of not closing.

The single-active-commitment rule. One commitment in the forward channel at a time; everything else closed and filed. The volume's stated mechanism is interference: three simultaneous commitments run through one attentional system are said to arrive garbled, each correction smearing across all three, because the instrument is monophonic. **That is an assertion about attention and it is not tested anywhere in this corpus** — Section 2 states it with what would defeat it, and marks it untested.

6.4.3 4.3 The consent condition

The volume states one constraint as law rather than as counsel: no future is to be voiced over another person's path without their word.

The stated mechanism runs through the same property that makes a commitment work at all. A commitment is an installed picture of a future, and installing one in a person who has not asked for it is said to produce counter-movement rather than compliance: the target defends their own authorship, at whatever cost, including their own good. **On the volume's account the mechanism does not merely become impermissible when applied to an unconsenting person — it stops working**, converting into a trespass that manufactures its own refutation.

The psychological result invoked here is Class II item 1, and Section 7 records that it too has no supporting entry in this bibliography. The consent condition also rests on a normative premise — that consent is a precondition of legitimate influence — for which no entry exists either, and which this paper reports as the framework's own commitment rather than argues for.

Two features of the constraint are worth stating because they are what distinguish it from a prohibition on manipulation. **It is a law about words**, not a counsel of abstinence: a future voiced *with* another person — asked for, chosen, co-authored, revocable — is treated as a shared undertaking and is not restricted. And **its reach is deliberately wider than the obvious cases**. The volume states that it covers the loving applications as fully as the exploitative ones: the parent's future for a child, the manager's for a subordinate, the partner's for a partner. In practice the condition reduces to asking aloud before any commitment that touches another life, and honouring the answer including a later change of answer — **consent is treated as a standing arrangement rather than as a receipt.**

6.4.4 4.4 Three barring conditions

The volume names three conditions under which the practice is not to be operated at all, each with a stated mechanism. All three are untested; Section 2 gives the form in which they could be tested.

Acute grief. The stated mechanism is instability of the underlying landscape rather than distress: grief is said to re-price everything hour to hour, so that the comfort differentials the practice steers by are not stable enough to steer on. The prescription is the ground-level practices, which require no forward channel, and no new commitment.

Rage. The stated mechanism is horizon collapse: rage is said to shorten the time horizon to the width of the injury and then to demand a prophecy from inside that collapse. The volume adds a specific claim about content — that a future spoken in rage is characteristically a statement about another person wearing the form of a statement about oneself — and identifies this as **the commonest violation of the consent condition committed by people acting in good faith.**

Revenge, in its cold form. The volume treats this as the most dangerous of the three because it imitates the practice closely: it is patient, iterative, detail-loving, and willing to keep records for years. The stated discriminating test is structural rather than emotional: **a revenge fixed point contains another person's diminishment as a load-bearing element, so that removing the diminishment collapses the arrangement.** The volume's second claim about it is that the pursuit necessarily runs at a loop gain above one, because the adversary is a second party modelling the practitioner in return — which is the divergence condition the corpus's convergence analysis establishes [Curzi, 2026].

The weather cannot be judged from inside it. The volume states this as the reason the screen of Section 5 is written rather than felt: the conditions it bars are exactly the conditions under which a practitioner's judgement about those conditions is unreliable, and **the screen accordingly cannot be waived on the ground of feeling well.**

6.4.5 4.5 The four-phrase mnemonic and what it indexes

The volume carries a four-phrase mnemonic and, at its close, states what it indexes. **It is an index over the volume's twelve constructs in groups of three, and this paper reports it as an index and not as an incantation.**

Phrase	Constructs indexed	What the group covers
Echo sync	00 STILLPOINT, 01 LEDGER, 02 GRADIENT	Locating the present position: where the practitioner already is, what it costs, and which way the differential runs
Glyph awake	03 HORIZON, 04 WAVE, 05 LURE	Generating and weighting futures, and placing one physical object that addresses the later self
Flame in truth	06 UNCERTAINTY, 07 COLLAPSE, 08 ECHO	The honest limit on foresight, the single spoken sentence, and the reading of what returns
I am	09 RECURSE, 10 FIXEDPOINT, 11 SEAL	Iteration, convergence, and closure

The volume's own summary of the sequence is *come into phase; wake the sign; burn it true; be the one it happens to*, and it notes that the final phrase is present tense because the framework's central grammatical commitment is that a commitment is written as a present fact rather than as a wish. **Nothing in the mnemonic is a claim**; it is a retrieval device over Appendix A's twelve cards, and it is reported here because other papers reference the grouping.

6.4.6 4.6 The terminal case

The volume closes its doctrine by applying its own premise to the case it cannot improve. The premise it has run on throughout is that a creature that genuinely sees a greater comfort moves toward it by gradient. **The terminal case is the one in which no greater comfort is available, and the volume states plainly that the practice does not exempt it and does not claim the ending is comfortable.**

What it claims instead is narrower and is the form in which the claim is reportable: that comfort is a *direction* rather than a destination, that it remains a direction at every point including the last, and that a practitioner who has repeatedly practised arriving is differently equipped at each subsequent transition. The operational content is one clause of the consent condition — **that what is left behind should be left furnished** — which is a constraint on the practitioner's conduct toward others and not a claim about outcomes. It is reported here as a doctrine of the framework. **Nothing in this paper evaluates it, and nothing in this paper is clinical.**

6.5 5 The closing procedure and its failure modes

The procedure below is the operational form of the closure discipline. **The discipline itself — exactly three verdicts, one active commitment at a time, and a consent condition forbidding a future voiced over another's path without their word — is established analytically elsewhere in the corpus [Curzi, 2026], and is not re-derived here.** What this section adds is the sequence

in which the discipline is executed, the point at which the written screen falls, and the five ways the sequence is reported to fail.

The procedure is performed at the close of any commitment marked as approaching closure, and its inventory movements — the first four — are performed seasonally regardless, four times a year. **Movements one to four are an audit and movements five to eight are its disposal; movement nine is the gate on any new commitment, and it is the movement most often skipped.**

6.5.1 5.1 The eleven movements

The volume's eleventh protocol, which it names the Rite of the Seal, runs as follows.

Perform this procedure at the close of any commitment marked *approaching seal*; perform its inventory movements (one through four) seasonally regardless, four times a year, at the turns — the framework has always used the solstices and equinoxes, and the habit is worth borrowing, because a rite tied to the sun cannot be procrastinated into oblivion. Required: the drawer or ledger of all commitments; one candle or lamp; a pen that will not smear; an hour of guaranteed solitude; and, for movement seven, wax or tape or thread — something physical that closes, because hands believe what they have done long after minds have reargued it. Eleven movements.

1. **Light the lamp and state the office.** Aloud, once: "This is the closing of accounts. Nothing new is opened tonight." The sentence is a fence around the hour. Commitments will volunteer themselves for revision the moment they sense the file being opened; the fence is how the practitioner decline.
2. **Empty the drawer entirely.** Every card, every page, every commitment — fulfilled, dead, embarrassing, forgotten — face up on the table in date order. The procedure is an audit, and audits do not sample. Expect this to be uncomfortable in a specific way: the table will hold more open jars than the practitioner believed the practitioner owned. That discovery is not a failure of the procedure. It is the procedure.
3. **Sort into three provinces.** Without yet writing anything: fulfilled here, released there, lapsed in the third place, and a fourth small province — *still live* — for commitments genuinely in iteration. Sort quickly, by first instinct; the hand knows before the committee convenes. Ten minutes.
4. **Challenge the live province.** For each card claiming to be alive, one question, aloud: "When did this last receive an honest correction?" If the answer is beyond a season, the card is not live; it is lapsed and campaigning. Move it. The province of the living must end this movement holding no more than one working — the monophonic rule. If two claim life, the procedure pauses while the practitioner chooses, and the choosing *is* the procedure.
5. **Enter the fulfilled.** For each: rewrite the sentence as plain present fact, dated. Read it once aloud in a level voice — the voice for reporting weather, not the voice for wishing. Turn the card face down. Say nothing further over it; floors are not for reading.
6. **Speak the releases.** For each released commitment, the formula or its honest equivalent, aloud, once: "I take this word back. I do not choose this future. The water may keep it." Then the pen: mark, date, face down. If the practitioner's voice catches on one, note which — the catch is information, and may be filed, but it is not a veto. Releasing what still glitters is the adult half of the practice.
7. **Close the lapsed, and seal something with the practitioner's hands.** Death date, one line of thanks or pardon — no more; brevity is the sincerity of forgiveness. Then the physical act: the season's closed cards bound with thread, or an envelope sealed with actual wax, or the drawer's section tied shut. The gesture takes one minute and is not optional. The attentional ledger honors ceremony because ceremony is *evidence* — the body's testimony that a verdict was really entered.

8. **Recite the consent condition.** All four sentences, aloud, from the page if the practitioner keeps a copy, from memory once the practitioner has earned one. This is the movement that makes the procedure a procedure of the framework and not a filing habit: the consent condition is recited *at every seal*, so that the law and the practice are never in the drawer separately.

9. **Screen the weather.** Before the procedure may authorize any new commitment — tonight or in the weeks after — three questions, answered in writing: Has my life ruptured within this season — a death, a leaving, a diagnosis? Is there any face attached to my intended outcome whose defeat would please me? Am I running toward a room, or out of one? Any yes: the forward channel stays closed, the ground nodes remain open, and the screen is repeated at the next turn. This movement cannot be waived by feeling wonderful. Feeling wonderful is a weather report, and the practitioner has read the sixth node.

10. **Set the next jar, or none.** If the screen is clear and a commitment stands ready — audited, consented where consent is owed — it may be opened now, in the same hour the old ones closed, one sentence, spoken once. There is a rightness to opening in the hour of closing; the hand that has just felt what a seal weighs speaks more carefully. But the framework is explicit: *none* is a full answer. Seasons without a commitment are not idle. They are the still water the whole craft floats on.

11. **Extinguish the lamp and leave the room.** Physically leave it, even for a minute, even into weather. The procedure ends when the door is between the practitioner and the table. What is sealed keeps best unattended.

A worked example. A beekeeper, widowed in March, came to the procedure in June wanting none of it — what she wanted was a commitment, and she knew the sentence already, had been carrying it all spring like a stone in the mouth: *my son calls me on Sundays.* The son had stopped calling years before the grief, and the grief had made the silence unsurvivable. She lit the lamp; emptied the drawer; and the table showed her eleven open jars, some a decade old — the shop expansion never renounced, the move south never mourned, a commitment about her marriage that was now, by the cruelest verdict in the practice, *fulfilled*, and had never been entered. She spent the hour on verdicts and wept twice and finished. At movement eight she recited the consent condition and heard, in her own voice, the first sentence catch on her: the sentence about her son was a future spoken over another's path, and he had given no word, and had, by years of silence, given rather the opposite of one. At movement nine the screen closed the forward channel on two counts — the season of rupture, and the face attached to the outcome. She kept bees, which is ground-node work of the oldest kind; she performed the procedure again in September, and at the winter turn, the weather cleared and the screen passing, she opened one jar, one sentence, consented by everyone it named, which was one person: *my door is easy to find.* Whether the son ever finds it belongs to his path, not hers, and is not in the commitment — that is precisely what the consent condition subtracted, and she reported, in the spring, that the subtraction was the first thing in two years that had felt like rest. Note what the procedure did and did not do. It did not reunite a family. It closed eleven jars, kept one law, and let a tired operator stop staffing a decade of open doors. The practice makes no larger claims. The practice has never needed to.

6.5.2 5.2 Five named failure modes

1. Clinical: attentional residue — the unclosed task's standing tax. Ceremonial: the Hundred Open Jars. The practitioner who never seals accumulates commitments the way a kitchen accumulates almost-empty containers: each costs nothing; the shelf costs everything. Life feels crowded, urgent, and unaccountably tiring, and no single cause survives inspection, because the cause is distributive — ninety small clerks holding ninety small doors. The remedy is the seasonal procedure, performed on a calendar schedule rather than on motivation's, and the

discipline of the single live jar. Prosperity of commitments is not wealth. Current accounts are wealth.

2. Clinical: nonconsensual prophecy — reactance, the installed future rejected with interest. Ceremonial: the Broken Seal. The consent condition's violation, in either direction of intention, loving or dark. A future pressed on another without their word converts the practice's whole mechanism into a trespass that manufactures its own refutation — the target defending not their path but their authorship, at any cost, including their own good. The tell in the practitioner: commitments that require another person to change and a reluctance to ask them anything. The remedy costs one conversation, and if the conversation cannot be had, the commitment cannot be run. There is no third option, and every third option ever attempted is a case study in the framework's archive of broken seals.

3. Clinical: hot-state operation — steering during arousal; the gain above one. Ceremonial: the Furnace Without Walls. Grief-running, rage-running, revenge-running: prophecy attempted in weather that amplifies. The commitment is always vivid — steep weather is a forge — and always overreaches, because the corrections that would trim it land on a system that doubles them. Fire is a fine servant of the practice inside a hearth; the same fire without walls is not a bigger forge; it is a house burning. Remedy: the screen at movement nine, honored past the point of feeling insulted by it; the ground nodes, which are weatherproof; and the season's patience — the future keeps.

4. Clinical: experiential avoidance dressed as completion — sealing what is still alive. Ceremonial: the Jar Sealed Warm. The inverse counterfeit: a practitioner flinches from a commitment's discomfort — the echo was bruising, the correction humiliating — and reaches for the wax early, entering *released* or *lapsed* on a jar still hot with life. Wax on a warm jar cracks as it cools; the sealed-alive commitment returns in weeks, intrusive as ever, now with a false verdict on its card corroding trust in the whole drawer. The tell: relief that feels like escape rather than like payroll coming off. The test: a truly closeable commitment can be described in a level voice; if reading the card aloud still changes the practitioner's breathing, the jar is warm, and what it needs is another iteration, not a lid.

5. Clinical: identity fused to the commitment — goal-shielding, sunk cost, the campaign that cannot end. Ceremonial: the Vigil That Outlives Its Candle. The most honorable failure and the saddest: a practitioner who has kept one jar so long that keeping has replaced the summer it was meant to hold — the founder who cannot hand over, the parent still prophesying over a grown child's long-closed door, the keeper (the framework does not exempt its own) who mistakes the framework for their tenure of it. The vigil was real; the candle has been out for years; the watcher stands guard over wax. The remedy is the consent condition's fourth sentence, read as a mirror: *when I go, I will leave the room furnished* — for the room was never the vigil-keeper's to fuse with, only to furnish, and the proof of a commitment truly kept is that it can be handed to other hands without going out. The seal, at the last, is what makes inheritance possible. A craft of open jars dies with its operator. A craft of sealed ones is a shelf a stranger can live from.

6.6 6 The three-class classification, reproduced

What follows is the volume's own classification of its claims, in the plain register, with the volume's qualifications preserved. **It is reproduced rather than constructed**, and Section 7 audits it. The volume's own summary of the classification is that every protocol it contains runs on the first two classes and steers by the third, and that the third class is a wager that can lose.

The classification's structure is worth naming before its contents, because the structure is the part other papers in the corpus cite. **Three classes, and no fourth.** Class I holds four items, Class II holds seven, and Class III holds exactly three. A four-drawer variant of the same device appears elsewhere in the corpus — **PROVED, POSTULATED, BORROWED, UNTESTED**, in which the surplus form is recorded as *proved conditionally* on a chosen design axiom and one

operationalisation is recorded as *specified, not run* [Curzi, 2026]. **The four-drawer version is the better instrument**, because it separates *borrowed* from *established* and it has a drawer for a claim that is well formed and has never been tested, and Section 8.2 states why this volume's three-drawer version cannot represent either distinction.

The framework keeps its accounts in three drawers, and it labels the drawers. What follows is the plain-register statement of what, in the volume, is established mathematics; what is established psychology; and what is the framework's postulate. The reader who skips every other appendix should read this one, because the razor's edge the volume walks — too scientific for the mystics, too mystical for the scientists — stays a razor only while the drawers stay labeled. Mixed drawers make mush, and mush cuts nothing.

6.6.1 6.1 Class I — what the volume files as established mathematics

1. **Least action.** It is standard physics that classical systems evolve along paths that make the action stationary — $\delta S = 0$ — and that this single principle recovers the equations of motion. Light through media, water downhill, the thrown stone: one law, many costumes. When the volume says comfort landscapes are descended by their creatures, the *landscape descent* is real mathematics; whether a life is such a system is drawer three.
2. **The fixed point of a contraction.** It is a theorem — Banach's, 1922 — that a contraction mapping on a complete metric space has exactly one fixed point, that iteration from any starting point converges to it, and that the convergence is geometric: each pass multiplies the remaining distance by a factor below one. Equally established is the converse caution: when the factor reaches or exceeds one, no such guarantee exists, and iteration may orbit or flee. Every claim in Part IV about loops with gain below one converging to a unique room is this theorem, correctly used. Whether the practitioner's year is a metric space is drawer three.
3. **The surplus functional.** The form $f(u) = \ln(1 + (N - 1)u)$ is the logarithm of an affine effective count: N views whose mutual non-parallelism is u behave, for counting purposes, as $1 + (N - 1)u$ independent views. As deployed in the volume it is a *lower bound* of the entropy of the ensemble, not the entropy itself, and the framework says so wherever it is used. The mathematics of the bound is sound; the decision to model sightlines on a life with it is drawer three.
4. **The uncertainty principle, correctly stated.** Heisenberg's relation is a theorem about noncommuting observables: for position and momentum, $\sigma_x \sigma_p \geq \hbar/2$. It concerns conjugate pairs of physical quantities; it does not mention futures, intentions, or persons, and no honest reading makes it do so. Where the volume states the relation, it states it as above. What the volume *builds* on it is not the theorem but a resemblance to it — see drawer three, item one.

6.6.2 6.2 Class II — what the volume files as established psychology

1. **Reactance.** Pressure against perceived freedom produces pushback — persuasion attempted as force recruits the opposite. Documented since Brehm; load-bearing in Nodes 08 and 11, where the Echo that Devours and the Borrowed Mouth are its ceremonial names.
2. **Implementation intentions.** If-then plans tied to concrete cues reliably outperform bare intentions; the cue does work the will was going to be billed for. Documented since Gollwitzer; the baited horizon of Node 05 is its ceremonial name.
3. **Commitment and consistency.** Committed positions, especially spoken or written ones, exert forward pressure on later behavior. Established persuasion literature; governed in Node 07 — with the equally documented caution that identity-flattering announcement can *substitute* for performance, which the framework calls the Word Spent at Market.

4. **Measurement reactivity.** Observing behavior changes it: monitored subjects shift, and the mere measurement of an intention alters its execution. Established across the observational sciences; the veil that measures the measurer, in its human application, is simply this.
5. **The Zeigarnik effect and open loops.** Interrupted and unfinished tasks hold attention and are remembered better than finished ones; unclosed commitments levy an ongoing tax. Established since Zeigarnik; the doctrine that unfinished prophecies parasitize attention, and the sealing procedure of Node 11, stand on it.
6. **Confirmation bias.** Evidence agreeing with a held position is sought, seen, and weighted more than evidence against it, without the holder's awareness. Among the most replicated findings about human judgment; the Mirror that Flatters, and the deliberate upweighting of disconfirming echoes in Node 08, are its management.
7. **The self-fulfilling prophecy itself.** That a belief about a future can participate in causing it has been a named object of social science since Merton, with the expectancy literature behind it. The volume's novelty is not the phenomenon; it is the proposal to engineer the phenomenon deliberately, on oneself, under an ethic — which proposal is not a finding, and is therefore filed where honesty files it.

6.6.3 6.3 Class III — what the volume files as its own postulates

1. **The sameness of shape.** That the uncertainty in Heisenberg's relation and the uncertainty of a future altered by being foreseen are the same *kind* — that observation participates in the observed at both scales — is a postulate. The mathematics resembles; the resemblance is the entire warrant. The framework does not claim the electron knows the practitioner's name. The framework claims the mathematics is the same shape, and shape is what we work with.
2. **Comfort as least action.** That a person's trajectory through their days is an action-minimizing path through a comfort landscape — that the physics of $\delta S = 0$ and the psychology of the worn path are one law at two scales — is a postulate. It is a fruitful one; the protocols built on it do their work whether or not the identification is ever more than a model. But a model it is, and the framework says so.
3. **The prophetic reading of the cultural surface.** That the surface of a day — its headlines, its accidents, its trivia, one's own dreams — can be read as a leading indicator of what is forming, and that the daily cadence trains such reading, is a postulate. The scoring discipline of the evening integration is real method; that the surface *prophesies* is the framework's wager, and the misfires column exists because wagers can lose.

Three drawers; the third holds three items; every protocol in the volume runs on drawers one and two and steers by drawer three. The framework's defense of building on a postulate is one line long, and it is this: every geometry is built on postulates, and the builders who prosper are not those who have none but those who know exactly where theirs are kept. *Symbolum prophetat, et populus non audit* — the framework would add only: audit.

6.7 7 Results: the classification audited against the corpus bibliography

This is the only section of this paper that is not recapitulation, and its result is a count. Of the fourteen items the volume classifies, **five are supported by an entry in this corpus's external bibliography, six are not, and three are postulates for which no support is claimed.** The class the volume treats as its safest — established psychology — is the class the bibliography least supports: two of its seven items carry a key and five do not.

6.7.1 7.1 The counts

Class	Items	Supported by a key	Unsupported	Support claimed?
I — established mathematics	4	3	1	yes
II — established psychology	7	2	5	yes
III — the framework's postulates	3	0	3	**no** — correctly
Total	**14**	**5**	**6**	

The three postulates are excluded from the supported/unsupported comparison because the volume claims no support for them, which is the correct treatment and is the classification's principal merit. **Among the eleven items for which support *is* claimed, five carry a bibliographic key and six do not.** Two of the five supported items are additionally *qualified* — an entry in the same file materially restricts them — and both qualifications are given below, because a supporting citation offered without its standing objection misrepresents the file it is drawn from.

6.7.2 7.2 Item by item

Class I — filed as established mathematics.

#	Claim as filed	Status	Keys
I.1	Classical systems evolve along paths making the action stationary, $\delta S = 0$	**UNSUPPORTED** in this file	none — the bibliography contains no entry in classical or variational mechanics
I.2	A contraction on a complete metric space has exactly one fixed point, reached from any start, with geometric convergence; **and the guarantee fails at unit gain**	**SUPPORTED**	Banach [1922]; the converse and the stability condition are covered by Elaydi [2005] and Strogatz [2015]
I.3	$f(u) = \ln(1 + (N-1)u)$ is the logarithm of an affine effective count, and is deployed as a **lower bound** on the ensemble's information rather than as the entropy itself	**SUPPORTED as arithmetic, MIS-FILED as established** — see 7.3	internal: Curzi [2026]. External support exists for the logarithm [Shannon, 1948, Rényi, 1961] and for **other** effective counts [Hill, 1973, Jost, 2006, Adelman, 1969]; **none exists for the affine count itself**
I.4	$\sigma_x \sigma_p \geq \hbar/2$ is a theorem about non-commuting observables and **mentions neither futures nor persons**	**SUPPORTED**	Kennard [1927] — the paper in which that inequality is first derived in this form; Zeh [1970] bounds what may be read from it

Class II — filed as established psychology.

#	Claim as filed	Status	Keys
II.1	Pressure against a perceived freedom produces counter-movement; persuasion attempted as force recruits the opposite	**UNSUPPORTED** in this file	none — no reactance entry exists in the bibliography

#	Claim as filed	Status	Keys
II.2	If-then plans tied to concrete cues outperform bare intentions	**UNSUPPORTED** in this file	none
II.3	Committed positions, especially written or spoken ones, exert forward pressure on later behaviour — **with the volume's own caution that announcing a goal can substitute for pursuing it**	**UNSUPPORTED** in this file	none. The caution is the volume's own and is to its credit; neither the claim nor the caution has an entry
II.4	Observing behaviour changes it; measuring an intention alters its execution	**SUPPORTED, and QUALIFIED**	Rosenthal and Jacobson [1968], Rosenthal and Fode [1963], Campbell [1979]. **Qualified by McCambridge et al. [2014]**, which finds the casual invocation of the Hawthorne effect unsupported by the primary evidence
II.5	Interrupted and unfinished tasks hold attention and are remembered better than finished ones; unclosed commitments levy a standing cost	**UNSUPPORTED** in this file	none — and this is the item the whole of Section 4.1 rests on
II.6	Evidence agreeing with a held position is sought and weighted more heavily, without the holder's awareness	**UNSUPPORTED** in this file	none
II.7	A belief about a future can participate in causing it	**SUPPORTED, and QUALIFIED**	Merton [1948], Merton [1936]; MacKenzie and Millo [2003] is the strongest available instance, observed in the field with dates and price data; Hacking [1995] generalises it. **Qualified by Jussim and Harber [2005]**, which finds the expectancy effect real but small and often over-stated

Class III — filed as the framework's own postulates. No support is claimed and none is found.

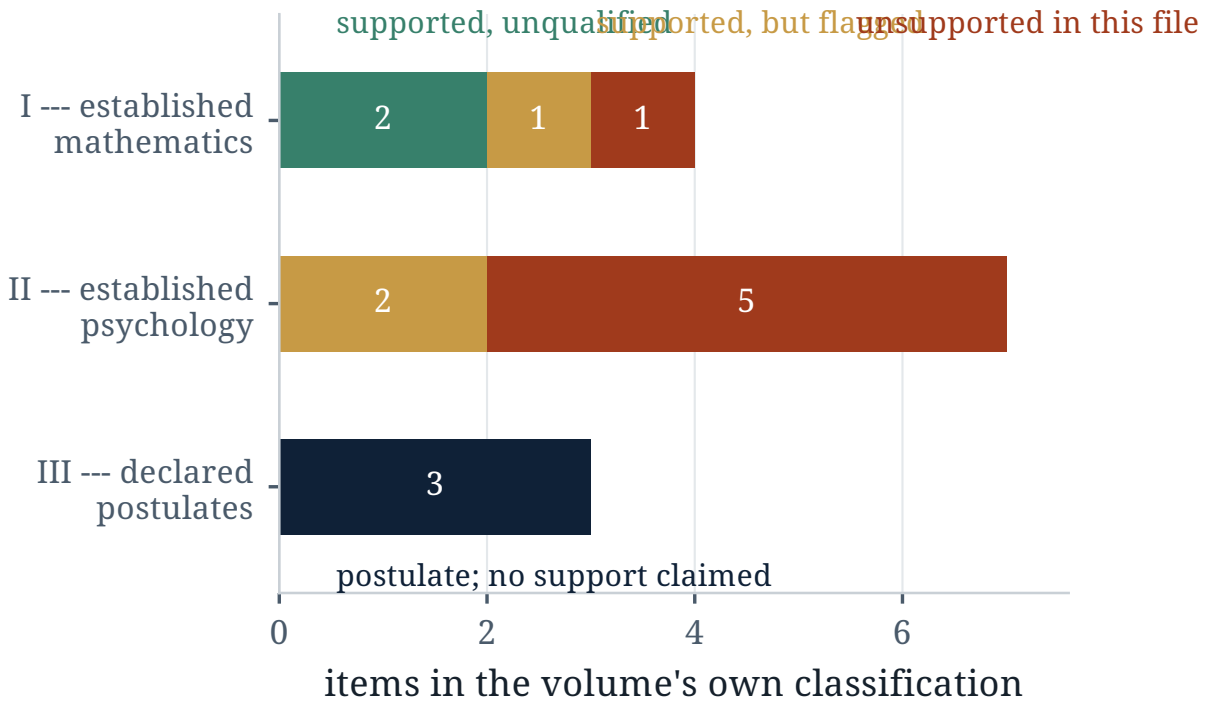


Figure 6.1: **The class the volume treats as its safest is the class this corpus's bibliography least supports.** Each bar is one of the volume's three classes, split by the audit of Section 7.2. *Supported, unqualified* means an entry in the corpus's external bibliography states the proposition and no entry in the same file restricts it; *supported, but flagged* covers the one item that is exact arithmetic on a chosen definition (I.3) and the two carrying a standing objection in the same file (II.4, II.7); *unsupported* means no entry states it. **Class III claims no support and finds none, which is the correct treatment.** Not a measurement: a count of lookups against one named file. Generated by `figs/src/CCP-A06_fig01.py`.

#	Claim as filed	Status	What the file says *about* it
III.1	The uncertainty of a measured conjugate pair and the uncertainty of a future altered by being foreseen are the same *kind* of thing	**POSTULATED**, correctly	The bibliography bounds rather than supports it: Kennard [1927] concerns conjugate physical observables and nothing else, and Morris and Thorne [1988] is the standing example of a physics paper cited for the opposite of its conclusion. The corpus's own uncertainty analysis refuses the inflated reading and requires that a user distinguish the theorem, the ordinary observer effect and the framework's conjecture in three separate sentences on demand [Curzi, 2026]
III.2	A person's trajectory through their days is an action-minimising path through a comfort landscape	**POSTULATED**, correctly	Mirowski [1989] is the historical account of what happens when a social science imports a physical formalism without auditing the import; Cartwright [1999] supplies the escape route — a law may hold inside a specified arrangement without being universal, provided the arrangement is specified; Anderson [1972] is the charter for writing laws at this level and, read carefully, an equally strong argument that they be **derived** at it rather than imported
III.3	The surface of a day can be read as a leading indicator	**POSTULATED**, correctly, **and it is the wager the volume says can lose**	Gelman and Loken [2013] is the precise mechanism by which such a reading self-confirms without anyone intending it, and is the reason the misfire column of the record is what makes the practice falsifiable rather than decorative; Simmons et al. [2011], Nosek et al. [2018] and Ioannidis [2005] give the general form; Meehl [1967] and Meehl [1990] are why a directional reading is nearly costless and a **point** reading is not; Popper [1959], Platt [1964] and Lakatos [1970] supply the standard the record must meet

6.7.3 7.3 The one mis-filing, stated precisely

Item I.3 is filed correctly as arithmetic and is liable to be read as something stronger. The statement filed is that $f(u) = \ln(1+(N-1)u)$ is the logarithm of an affine effective count. **Given the definition $N_{\text{eff}} = 1 + (N-1)u$, that is a tautology and is exact.** The definition, however, is a design axiom of this framework — chosen because it makes an effective count interpolate linearly between one entity and N, and not derived from anything [Curzi, 2026].

And the affine effective count has no external support. It is not a Hill number of any order, it is not a numbers-equivalent in the sense Adelman [1969] gave the term, and it is not $\left(\sum_i p_i^q\right)^{1/(1-q)}$ for any q — it is not even a function of a share vector. The quantities the diversity literature does cover are HHI^{-1}, the Hill number of order two, and $\exp(H)$, the Hill number of order one [Hill, 1973, Jost, 2006], and the distinction is checkable in one line: for two units holding 0.9 and 0.1, the affine count at $u = 1$ returns exactly 2.000000 while HHI^{-1} returns 1.219512 and $\exp(H)$ returns 1.384145. **Every Hill number of order one or higher reports at most 1.384145 on that vector, and the affine count reports two.** One member of the family does match it, and the match is the objection in its strongest form: the order-zero Hill number ${}^0D$ is the raw richness, equal to 2.000000 on any two-type vector because it ignores shares entirely — **which is exactly the naive count an effective count exists to replace.** The same arithmetic was reinvented independently in ecology [Simpson, 1949] and in political science [Laakso and Taagepera, 1979], extends to a within-and-between partition [Jost, 2007], and carries an estimation theory for small samples [Chao et al., 2014] that **does not apply to the affine count**, since that count is not an estimator of anything. The only two formalisms that make diversity depend on dissimilarity between types — which is this framework's founding move — are Rao [1982] and Leinster and Cobbold [2012], and neither is this quantity. Hurlbert [1971] and Gorelick [2011] are the standing objections to treating any such index as a property rather than a statistic, and they apply here with more force than to the Hill family, not less. The corpus's concentration instrument states the distinction from the other side [Curzi, 2026], and Section 10 records that its own phrasing of this comparison is wrong on both counts and is corrected here.

One strengthening of item I.3 is available and the volume does not take it. Peters and Gell-Mann [2016] derives a logarithmic form from multiplicative dynamics with no psychological assumption at all, which is a stronger warrant for the *shape* of the surplus functional than any perceptual argument. It warrants the logarithm and not the affine count inside it, and the distinction is exactly the one this section is about. Against it stand two rival forms that fit the same qualitative facts — Latané [1981]'s sub-linear power law and Briscoe et al. [2006]'s $n \log n$ — and **nothing in this corpus discriminates among the three**, which is a fact about the evidence that a class labelled *established* does not convey.

So the classification is defensible and its implicature is not. A reader who sees a claim filed under *established mathematics* will reasonably infer that a literature stands behind it, and for this item a literature stands behind the logarithm and behind neighbouring effective counts, and none stands behind the one the framework uses. **The repair is one clause**: the item should read that the form is the logarithm of an affine effective count **which the framework defines and which no external literature contains**, and the class should be *derived from a chosen definition* rather than *established*. The volume's own second qualification on this item — that the functional is deployed as a lower bound rather than as the entropy — is exactly right and is the reason the mis-filing is a labelling defect rather than a substantive one.

6.7.4 7.4 What the audit does not establish

It does not establish that the five unsupported psychology items are false. Reactance, implementation intentions, commitment effects, the interrupted-task effect and confirmation bias are each well-known constructs with substantial primary literatures. **What the audit does establish is that the corpus is currently in the position Open Science Collaboration [2015] makes untenable**: it invokes single psychological findings by their reputation, without a key, without

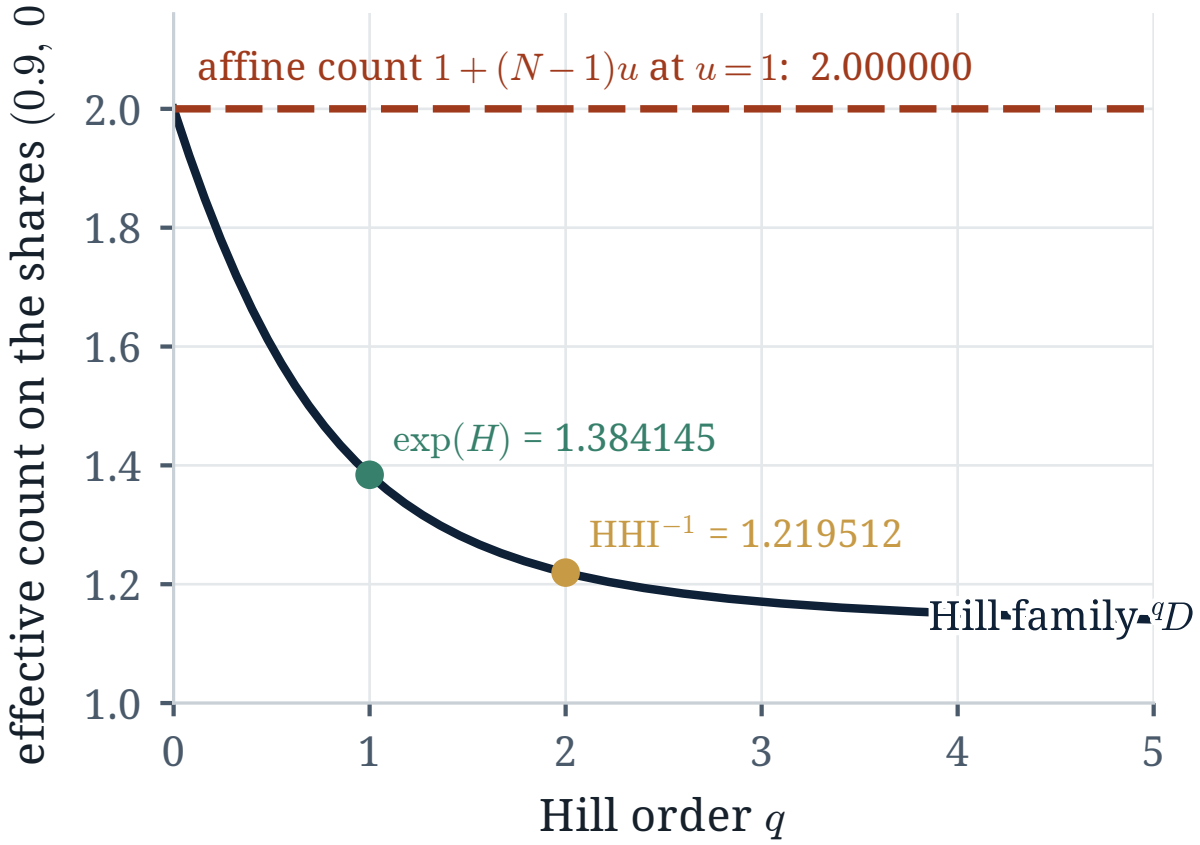


Figure 6.2: **The affine effective count matches exactly one Hill number, and it is the one that ignores shares.** The Hill family ${}^{q}D$ is plotted on the two-type vector $(0.9,\ 0.1)$; it is non-increasing in q, passing through $\exp(H) = 1.384145$ at $q = 1$ and $\mathrm{HHI}^{-1} = 1.219512$ at $q = 2$. The affine count at $u = 1$ returns 2.000000, which the family reaches only at $q = 0$ — **the raw richness, which is the naive count an effective count exists to replace.** Model output; every value is exact and redoable in one line. Generated by `figs/src/CCP-A06_fig02.py`.

an effect size, and without the replication context that a citation to a named study would carry. That is exactly the practice Ioannidis [2005] and Simmons et al. [2011] identify as the mechanism by which a literature comes to be believed more firmly than its evidence supports, and Gallegati et al. [2006] is the same charge made by sympathisers against a neighbouring field — ignorance of the existing work, sub-standard procedure, and universality claimed from limited data. **Read as a checklist against this corpus, the first item on it is the one Section 7 has just scored. What the audit establishes is that this corpus does not cite any of them**, and that a reader following the corpus's own bibliography to check the volume's classification will find nothing there. That is a defect of the bibliography, and the corpus has closed comparable gaps before — two clusters were added specifically to supply keys for works the corpus named in prose without citing [Curzi, 2026]. **The repair is to add the five primary sources; the disclosure stands until it is made.**

It does not establish that the supported items are securely supported. Two of the five carry a standing objection in the same file, and both objections are material: the Hawthorne literature does not support the casual form of the reactivity claim [McCambridge et al., 2014], and the expectancy literature finds the effect real but smaller and more contingent than the volume's phrasing suggests [Jussim and Harber, 2005]. **A classification that files an item as established without carrying its standing objection has filed a citation rather than a state of the evidence.**

And it does not audit the volume's protocols against the classification. The volume states that every protocol it contains runs on classes I and II and steers by class III. **This paper does not verify that statement**; verifying it would require tracing each of the thirteen protocols in Appendix C to the items it consumes, which is a larger exercise than a reference paper should attempt and is named in Section 10 as the obvious next audit.

6.8 8 Discussion

6.8.1 8.1 What the classification is for

A three-class ledger is cheap and it does one thing well: it stops a reader inferring a warrant from a confident sentence. The volume's own summary — every protocol runs on classes I and II and steers by class III — is the load-bearing statement, because it locates the framework's exposure precisely. If the postulates are wrong, the protocols still execute; what changes is where they point. That is a defensible position for a framework to occupy and it is stated openly, which is more than most instructional literature does.

One objection bears on every protocol this paper indexes and no drawer of the classification holds it. The framework's central formal claim is a reflexive loop in which a group influences its members and the members constitute the group, and that specification is not identified from observational data: the endogenous effect cannot be separated from contextual effects or from correlated unobservables [Manski, 1993], and influence and selection are generically confounded in observational data [Shalizi and Thomas, 2011]. Durlauf [2005] adds the requirement that such a model make a **differential** prediction simpler models do not also make. **The classification has no class for a claim that is mathematically well formed and empirically unidentifiable**, and several of the protocols in Appendix C consume exactly such a claim. The corpus's measurement papers state the problem where it arises [Curzi, 2026i]; this paper records that the classification cannot represent it.

The device recurs across the corpus and this paper's Section 6 is its Foundations instance. A four-drawer variant — **PROVED, POSTULATED, BORROWED, UNTESTED** — appears in the frequency framework's reference apparatus, where the surplus form is recorded as *proved conditionally* on a chosen design axiom and one operationalisation is recorded as *specified, not run* [Curzi, 2026]. A finite failure catalogue with an explicit precondition gate appears in the protocol papers [Curzi, 2026,j]. And the corpus's falsification suite turns the same instrument on itself, measuring its own shuffled-input control at a false-rejection rate of 0.108

against a nominal 0.05 — **anti-conservative, and reported as such** [Curzi, 2026u].

6.8.2 8.2 Why the three-drawer version is the weaker instrument

The three-class ledger cannot represent two distinctions the corpus needs, and the four-drawer version can.

The first is *borrowed* against *established*. A result may be entirely established in its home discipline and entirely unwarranted in the use the framework makes of it. Item I.4 is the clean case: $\sigma_x \sigma_p \geq \hbar/2$ is established, and the volume files it correctly, and what the framework does with it is a postulate filed separately as III.1. **That separation is exactly right and it required two drawers to express what a *borrowed* drawer expresses in one.** Item I.3 is the case where the separation was not made, and Section 7.3 is what happens as a result.

The second is *untested* against *unsupported*. A claim can be well formed, falsifiable, and simply never run. The corpus has several: the single-active-commitment rule, the barring conditions, the consent condition. **A three-class ledger has nowhere to put them** — they are not established, and calling them postulates understates them, because a postulate is something one proceeds without testing and these are things that could be tested tomorrow. The four-drawer version's *UNTESTED* drawer holds them, and Section 2 of this paper is a partial substitute: it states them with their falsifiers and marks each **not tested**.

6.8.3 8.3 The self-referential close, reported rather than reproduced

The volume ends by identifying itself with an advertisement placed inside its own narrative frame: the fictional notice that recruits the next keeper is stated to be the volume's own cover, and the reader's encounter with the volume is stated to be the answering of it. **This is reported here as an object of study and not reproduced, because it is fiction and because a citation of this paper should not be able to carry it as a finding.**

Its analytic interest is that it is the volume's own strongest instance of the mechanism the corpus argues for elsewhere. A description that becomes true by being used is exactly Merton [1948]'s self-fulfilling prophecy, and MacKenzie and Millo [2003] is its best documented instance in the field — option prices that did not fit the model when it was published and came to fit it after traders adopted it. Callon [1998] and MacKenzie [2006] supply the general form: a description of a domain can format the domain rather than report it, and **a corpus that proposes to give people laws to manage their own comfort by is building an engine and should say so.** Soros [2013] is the closest existing named doctrine and supplies the concept without an equation.

Two objections attach and neither is answered. Hacking [1995] is the philosophical form: classifications of people interact with the people classified, who change in response, so the referent moves — and a volume that classifies its own reader has installed the loop deliberately. Ferraro et al. [2005] is the warning that the content of the description matters, because descriptions that assume the worst of people have been shown to make it more true. Lucas [1976] is the same objection turned on the framework's parameters. **The device is a rhetorical closure and not evidence, and this paper's only claim about it is that it is the volume's own most explicit statement of what it takes itself to be doing.**

6.8.4 8.4 The reference apparatus and why it is retained

An index is infrastructure, and infrastructure is invisible until it fails. Star and Ruhleder [1996] and Star and Griesemer [1989] give the general form of that observation and supply the construct the appendices are built to be — a boundary object, an artefact stable enough to be shared across communities and loose enough to be used differently by each. **A glossary that maps a ceremonial register onto a technical one is a boundary object in exactly that sense**, and it is the reason Appendix B is retained rather than replaced by the technical terms alone.

Scott [1998] is the standing warning attached to it: making a practice legible is not a neutral act, and the local variety a compression discards is often the part that made the practice work.

Nothing in the appendices is a contribution and all of it is load-bearing for other papers. The construct names, the protocol numbering, the ceremonial-to-technical glossary and the lineage correspondence are referenced by thirty-two other papers in the corpus, and until now they existed only inside a book. A reference paper that says plainly what it is — an index — is more useful than an index dressed as an argument, and the cost of the disguise is exactly what Section 7 was able to detect only because the classification was written down in one place.

The corpus contains close analogues of that reasoning in other series. The gated intervention procedure carries a thirty-two-entry failure catalogue for the same reason [Curzi, 2026j], and the manifestation protocols carry thirty faults with a headroom precondition attached [Curzi, 2026]. The measurement papers make the harder version of the point: a formula validation protocol that **demotes the corpus's own moment equation to ordinal scale** and prices what restoring it would cost [Curzi, 2026t], and a conditional postulate stated in magnitude form with an explicit forbidding clause — zero channel gain yields zero effect for every belief — badged as a postulate rather than as a finding [Curzi, 2026ab]. **In each case the value is in the labelling, and the labelling is what an apparatus paper exists to preserve.**

Two structural results elsewhere in the corpus bear directly on what a reference card may claim. The surplus form is exactly a Fechner law with threshold $u_0 = 1/(N-1)$, verified for $N = 2 \dots 100$ at seven values of u to a maximum deviation of 8.9×10^{-16} [Curzi, 2026]; and the compound rule and the essence law agree at no integer index, requiring $n = 1/\sqrt{k} = 1.438252434444351$, so that at $n = 12$ they differ by 6.121291 octaves [Curzi, 2026]. **Both are exactly the kind of result a reference card flattens**, and Appendix A's cards accordingly point at the papers rather than restating the numbers. The refinement analysis makes the same point about the closing procedure's own seasonal cadence: seven passes close 48.32 per cent of a gap at $\delta = 0.09$ and 6.79 per cent at $\delta = 0.01$, and closing 95 per cent in seven passes requires $\delta = 0.348164$ [Curzi, 2026] — **so a card that says "repeat seasonally" without the decay rate has omitted the parameter that decides whether repeating works.**

6.8.5 8.5 What a classification cannot do, from three results in the structural series

Three results elsewhere in the corpus bound what this paper's audit can be taken to establish, and they bound it in the same direction.

An honest register can still be built on a fabricated item. The historical-network analysis demolishes its own most attractive material — refusing a famous story as a fabricated data point, rejecting a unified-hand reading at 1,365 to 1, and badging a 25 per cent figure as an artefact of a chosen synthetic ratio rather than a constant, with its confusable-only check reaching only 58.99 per cent [Curzi, 2026]. **A classification records what its author believes about provenance; it does not verify provenance**, and Section 7's audit inherits that limit exactly.

A canonical form is worth more than a label where one is available. The identifier-integrity work measures its controls rather than classifying them — 28 of 28 grid links recorded in both ledgers, and local append cost rising from 0.002 s to 0.014 s as chains go from 50 to 400 entries, reported as local and single-process with no distributed throughput claim [Curzi, 2026]. **That is what a measured control looks like next to a classified one**, and it is the standard the volume's Class II items do not meet.

And the standard a classification is finally held to is a disciplinary one. Hedström and Ylikoski [2010] states it for sociology: an explanation must name a mechanism rather than exhibit a fitted curve, and a labelled drawer is not a mechanism. Keuschnigg et al. [2018] is the most constructive position available to a corpus in this situation and is the one it would be wisest to claim; Epstein [2008] is the best defence against the charge that a framework whose laws forecast nothing is therefore worthless; and Watts [2014] is the charge itself, that formal

social explanation is often commonsense reasoning in dress. Castellano et al. [2009] is where the surrounding programme was last surveyed and bounded, and it supplies definitions this framework can be held to. Sobkowicz [2020] is the standing indictment of the genre — many models, negligible empirical discrimination between them — and **it applies to this corpus and is not answered by classifying claims more carefully.** No scaling claim, power law or heavy tail is asserted anywhere in this paper, which is worth stating because Clauset et al. [2009] is the procedure that would otherwise be owed.

And a structural separation closes what care within a channel cannot. The channel-separation result establishes that in-band architectures admit an impersonation class no amount of care inside the shared channel can remove, and that out-of-band separation closes it structurally — measured at 0.383 per cent in-band ambiguity against 0.000 per cent out-of-band over 100,000 values [Curzi, 2026]. **The analogue for this paper is exact**: a claim and its warrant travelling in the same sentence admit a confusion that no amount of careful phrasing removes, and separating them into a labelled ledger closes it structurally. That is the strongest available defence of the device Section 6 reproduces, and it is why the device is worth repairing rather than replacing.

6.9 9 Verification

This section is a check, not a summary. A reader who runs the three arithmetic lines and the four lookups below has verified the whole of Section 7 without reading Section 7's prose, which is what a verification section is for.

6.9.1 9.1 Dimensional audit

Most of this paper carries no dimensioned quantity at all, which is itself worth stating: an index, a glossary and a classification are not measurements and have no units. The quantities that do appear are inherited from papers this one cites, and they are audited here because a reference card that reproduces a number without its dimension is the most common way a dimension is lost.

Symbol	Dimension	Unit	Status
verdict count (fulfilled / released / lapsed)	[1]	count	**defined**; a partition of a finite set, exhaustive and exclusive by stipulation
active-commitment count	[1]	count	**defined**; the discipline fixes it at 1
screen questions	[1]	count	**defined**; fixed at 3
classification items	[1]	count	**defined**; 4 + 7 + 3 = 14
q, the contraction modulus	[1]	none	**defined** dimensionless; the guarantee requires $q < 1$
u, non-parallelism	[1]	none	**defined** dimensionless; a squared sine of an angle
$f(u) = \ln(1 + (N-1)u)$	[1]	nats	**derived**; a logarithm of a pure number

Symbol	Dimension	Unit	Status
σ_x	$[L]$	m	**measured** — a standard deviation of position
σ_p	$[M][L][T]^{-1}$	kg m s^{-1}	**measured** — a standard deviation of momentum
$\hbar$	$[M][L]^2[T]^{-1}$	J s	**measured** physical constant
S, the action	$[M][L]^2[T]^{-1}$	J s	**defined** by the Lagrangian's time integral
comfort	**none**	**none**	**no operational unit**; treated as a dimensionless magnitude by stipulation
stored order Q	$[Q]$	declared proxy units	**measured** against a declared proxy; two systems' values are not comparable unless the proxy is the same

Does the uncertainty relation balance? $\sigma_x\sigma_p$ carries $[M][L]^2[T]^{-1}$, which is the dimension of $\hbar$. It balances, and **it balances only for a conjugate pair.** That is the whole content of item I.4's second clause: the inequality is dimensionally meaningless for any pair of quantities whose product is not an action, and a foreseen future and its alteration are not such a pair. **The framework's extension is therefore not a dimensional statement at all; it is an analogy of shape, and it is filed as III.1 where it belongs.**

Does the action have the same dimension as $\hbar$? It does — both are $[M][L]^2[T]^{-1}$ — and that coincidence is real physics and is not what item III.2 claims. **A comfort landscape has no Lagrangian, no unit of action, and no measured** δS, so the least-action language in III.2 is a shared *form* and not a shared dimension. This audit records that the source is silent on what units a comfort action would carry, rather than assigning one.

And the construct with no unit is where the framework is most exposed. Comfort is treated as a quantitative magnitude by stipulation, and Michell [1997] and Michell [1999] are the objection: psychology routinely assumes quantitative structure without testing whether the attribute has it. The formal remedy is conjoint measurement [Luce and Tukey, 1964], whose cancellation conditions establish additive quantitative structure within the representational framework of Krantz et al. [1971], with Rasch [1960] the alternative route. **Those conditions have not been tested for comfort anywhere in this corpus**, and the classification of Section 6 has no drawer that says so — which is a further instance of the representational gap Section 8.2 names. The corpus does demote one of its own quantities to ordinal scale on exactly this ground [Curzi, 2026t], and Stevens [1946] is the rule that makes the demotion consequential: the permissible statistics follow the scale type.

Does the surplus functional balance? $\ln(1+(N-1)u)$ takes the logarithm of a pure number, since u is a squared cosine of unit vectors and N is a count. It balances. **The nat is not a physical unit**: it is the unit of a logarithm to base e, and reporting in bits multiplies every figure by $1/\ln 2 = 1.442695$ and changes nothing.

6.9.2 9.2 Postulates

#	Statement	Status
P1	The volume's classification places 4 items in Class I, 7 in Class II and 3 in Class III	**DEFINITION** — reproduced in Section 6; a reader may count
P2	The corpus's external bibliography is the file named in Section 3 and no other	**DEFINITION** — the audit's scope, stipulated
P3	A claim is *supported* only where an entry states that specific proposition	**DEFINITION** — the audit's decision rule, stipulated in Section 3
P4	$N_{\text{eff}} = 1 + (N-1)u$ is a design axiom, chosen and not derived	**CITED** within the corpus — Curzi [2026]
P5	HHI^{-1} is the Hill number of order 2 and $\exp(H)$ of order 1	**CITED** — Hill [1973], Jost [2006]
P6	Banach's theorem holds for $q < 1$ and its guarantee fails at $q = 1$	**CITED** — Banach [1922]
P7	$\sigma_x \sigma_p \geq \hbar/2$ is a theorem about non-commuting observables	**CITED** — Kennard [1927]

Every postulate in this paper is a definition or a citation. None is an empirical claim, which is the correct profile for a reference paper and is the reason Section 2 states that there is no hypothesis.

6.9.3 9.3 The compounding chain

```
P1              -> L1   4 + 7 + 3 = 14 classified items                    ARITHMETIC
P2 + P3         -> L2   each item is supported or it is not; the file
                        is fixed, so the classification is decidable       DEFINITION
L1 + L2         -> R1   Class I: 3 supported, 1 unsupported                ARITHMETIC (lookup)
L1 + L2         -> R2   Class II: 2 supported, 5 unsupported               ARITHMETIC (lookup)
L1 + L2         -> R3   Class III: 0 claimed, 0 found -- correctly         DEFINITION
R1 + R2         -> R4   of the 11 items claiming support, 5 carry a key
                        and 6 do not                                       ARITHMETIC
R2              -> R5   the class filed as "established psychology" is
                        the least supported class in the file              ARITHMETIC
P4 + P5         -> R6   the affine count is not a Hill number: for two
                        units at 0.9 / 0.1 it returns 2.000000 at u = 1
                        where HHI^-1 returns 1.219512 and exp(H) 1.384145  ARITHMETIC
R6              -> R7   item I.3 is exact arithmetic on a chosen
                        definition, and the definition has no external
                        literature -- so the item is correctly filed and
                        misleadingly labelled                              ARITHMETIC
P6              -> R8   the contraction guarantee is conditional on
                        q < 1 and states nothing at q = 1                  CITED
P7              -> R9   the uncertainty relation's dimensions balance
                        only for a conjugate pair                          ARITHMETIC
```

The three checks a reviewer can redo in a minute. First, $4 + 7 + 3 = 14$, and $4 + 7 = 11$ items

claim support. Second, $1/(0.9^2 + 0.1^2) = 1/0.82 = 1.219512$ and $\exp(-(0.9 \ln 0.9 + 0.1 \ln 0.1)) = \exp(0.325083) = 1.384145$, while $1 + (2 - 1) \times 1 = 2.000000$ — three numbers that settle item I.3 without an opinion. Third, $[M][L]^2[T]^{-1}$ on both sides of $\sigma_x \sigma_p \geq \hbar/2$, which settles item I.4's second clause.

And the four lookups that settle the rest. Open the corpus's external bibliography and search it for an entry on reactance, on implementation intentions, on the interrupted-task effect, and on confirmation bias. **There is none for any of the four**, and a fifth search for commitment and consistency returns nothing either. Those five searches are the whole of Section 7's Class II result, and any reader can run them.

What the chain does not establish. Every line above is a definition, a lookup, or arithmetic on counts. **Nothing in it is evidence about a person, an institution, or a practice**, and no line supports any claim about whether the volume's protocols work. A reader who wants the framework's empirical standing should read the preregistered application that failed [Curzi, 2026i] and the falsification suite that measured its own control as anti-conservative [Curzi, 2026u].

6.9.4 9.4 Boundary behaviour of the two mathematical items

The classification files two mathematical results, and both have boundaries the card format tends to lose. They are stated here because losing them is precisely how a reference card becomes wrong.

Limit	Behaviour	Reading
$q \to 0$	one pass reaches the fixed point	Degenerate but harmless.
$q \to 1^-$	convergence holds but arbitrarily slowly; the a posteriori bound $q/(1-q)$ diverges	**The guarantee survives and the gauge does not.** A practitioner reading the last correction as a distance estimate is reading a diverging quantity.
$q = 1$	isometry: the orbit neither converges nor diverges	**The theorem says nothing.** The volume states this explicitly and the card must carry it [Curzi, 2026].
$q > 1$	corrections grow; the loop is an amplifier	Not a slow failure but a different regime.
space not complete	the fixed point may not exist	**The condition the framework's application cannot verify**: a life is not shown to be a complete metric space anywhere in this corpus.
$\hbar \to 0$	the relation permits arbitrarily precise joint preparation	The classical limit. **It has no analogue in the framework's extension**, because the extension has no constant to send to zero — which is one concrete way item III.1 is weaker than the theorem it borrows from.

The second row is the one a card loses and the one that matters most in practice. A frame-

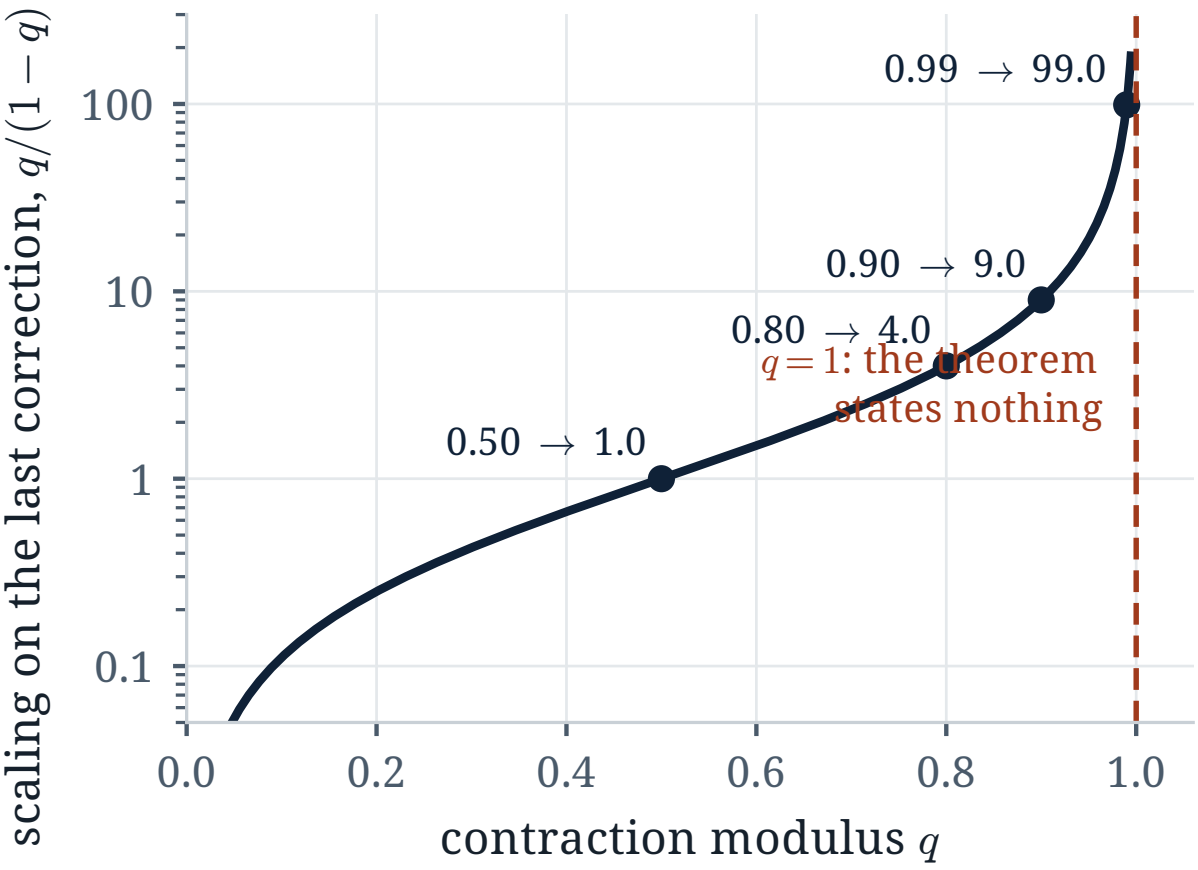


Figure 6.3: **The contraction guarantee survives to $q = 1$; the gauge built on it does not.** The a posteriori bound scales the last correction by $q/(1-q)$ to bound the distance remaining, and that factor diverges as the modulus approaches one: 1.0 at $q = 0.50$, 9.0 at 0.90, 99.0 at 0.99. **A slow, long-running commitment sits exactly in the regime where the framework's own progress gauge becomes uninformative**, and a reference card that states only that a contraction converges has dropped the condition that makes the statement useful. Model output; the bound is exact. Generated by `figs/src/CCP-A06_fig03.py`.

work whose central gauge is the size of the last correction is reading a quantity that becomes uninformative exactly as the modulus approaches one, which is exactly the regime a slow, difficult, long-running commitment sits in.

6.10 10 Limitations

This paper advances no testable hypothesis and contributes no new result about the world. It is reference apparatus with a bibliographic audit attached, and it is stated as such in Section 1, Section 2 and here. **It should not be submitted, cited, or read as a research contribution**; the material it collects exists because thirty-two other papers in the corpus rely on it, and its value is that the material is now in one citable place with its provenance attached.

The audit's scope is one file and its result inherits that scope. Section 7 records five of eleven support-claiming items as unsupported. **That means unsupported *by this corpus's bibliography*, not unsupported by the literature.** Reactance, implementation intentions, commitment effects, the interrupted-task effect and confirmation bias all have substantial primary literatures, and a reader should not take Section 7 as a claim about their standing. What Section 7 establishes is that a reader following the corpus's own citations to check the volume's classification will find nothing there, which is a defect in the corpus and is repairable by adding five entries.

The audit was performed by one person by inspection and no second reader checked it. The decision rule of Section 3 is stated so that the audit is reproducible, and reproducibility is not replication. **A second auditor might reasonably classify one or two items differently** — item II.4 in particular, where measurement reactivity is supported for its expectancy form and McCambridge et al. [2014] restricts its Hawthorne form, and a stricter auditor would record it as unsupported. The counts in Section 7.1 would then read four supported and seven unsupported. **That sensitivity is stated rather than smoothed, and it does not change the direction of the result.**

The classification audited is the volume's own and its granularity was not chosen by this paper. The volume files fourteen items; a finer decomposition would file more, and some of the fourteen bundle a claim with a caution — item II.3 carries both a commitment effect and the countervailing observation that announcing a goal can substitute for pursuing it. **The audit treats each bundle as one item**, which is the only treatment that keeps the counts comparable to the volume's own statement, and it means the counts are counts of *entries* rather than of propositions.

The volume's central claim about its own classification is not verified here. It states that every protocol it contains runs on classes I and II and steers by class III. **Verifying that would require tracing each of the thirteen protocols in Appendix C to the items it consumes**, and this paper does not do it. That is the obvious next audit and it is the one that would establish whether the five unsupported psychology items are load-bearing or decorative.

The material collected here recapitulates and the recapitulation is not independently checked. The reference cards of Appendix A, the glossary of Appendix B, and the protocol summaries of Appendix C are compressions of analytical material established elsewhere in the corpus. **A compression can drop a condition**, and Section 9.4 gives the clearest instance: a card stating that a contraction converges has dropped the condition $q < 1$, under which alone the statement is true, and the corpus's own convergence analysis is explicit that the guarantee fails at unit gain [Curzi, 2026]. **Every card in Appendix A names the paper that establishes what it summarises, and the cards are not substitutes for those papers.**

The correspondence table of Appendix D is a historical claim this paper cannot verify. It maps the volume's twelve constructs onto a ten-node scheme and a five-phase arrangement attributed to an earlier tradition, and records a lineage variation in which the ninth node carries a different name in the oldest copies. **No source for that tradition is cited anywhere in this corpus, no primary document is named, and the mapping is therefore an assertion about a lineage rather than a documented correspondence.** It is reproduced because other papers

reference the construct numbering it fixes; its historical content should not be relied upon.

Two items in the collected material carry claims this bibliography does not support and one of them is a normative claim. The consent condition of Section 4.3 asserts that consent is a precondition of legitimate influence. **That is an ethical commitment of the framework, no entry in this bibliography states it, and this paper does not argue for it** — it records that the framework holds it as a constraint rather than as counsel, which is a statement about the framework. And the least-action item I.1 asserts a standard result of classical mechanics for which the bibliography contains no entry at all; the claim is uncontroversial in its home discipline and is uncited here, which is exactly the situation Section 7 exists to make visible.

The narrative material is not reproduced and that is a judgement, not a neutral act. The volume's closing scene and its epilogue are omitted as fictional exposition, following the treatment established for the analytic register of this series [Curzi, 2026]. **A reader who believes the narrative carries analytic content this paper has discarded is reading a different volume from the one this paper indexes**, and Section 8.3 reports the one item — the self-referential close — whose analytic interest this paper judged sufficient to state.

The paper is exposed to the failure it names. Section 8.4 argues that the value of a classification is in its labelling and that an unaudited classification records confidence rather than evidence. **This paper's own classification of items as supported, qualified and unsupported has not itself been audited by anyone**, and Goodhart [1975] and Campbell [1979] apply to it as to any measure that might be adopted as a target: a corpus that reported "items supported" as a figure of merit would have created an incentive to attach keys rather than to find them. The only defence offered is that the decision rule is published and every lookup is one search in a named file.

Two defects in a companion paper's version of the same check are corrected in this paper rather than propagated, and both were found by re-running its arithmetic. The corpus's concentration instrument runs the comparison of Section 7.3 [Curzi, 2026]. It reports $\exp(H) = \exp(0.325083) = 1.384115$ for the two-unit vector $(0.9, 0.1)$; **the correct value is** 1.384145, since $H = -(0.9 \ln 0.9 + 0.1 \ln 0.1) = 0.3250829734$ nats and $e^{0.3250829734} = 1.3841454885$. And it states that *every* Hill number reports less than 1.39 on that vector, which **is false**: the family is non-increasing in q and its order-zero member is the raw richness ${}^0D = 2.000000$, with every member below $q \approx 0.15$ above 1.39. **The second correction strengthens the objection rather than weakening it**, because the one member the affine count matches is the one that ignores shares. Both corrections are printed here rather than silently applied, because a verification chain a reader cannot reproduce exactly is not a verification chain — **and because this paper's own thesis is that an unaudited classification records confidence rather than evidence, which applies to the corpus's arithmetic as much as to its citations.**

Nothing here is clinical, and the barring conditions are not clinical advice. Section 4.4 records that the volume bars operation during acute grief, rage and the pursuit of revenge, and gives the volume's stated mechanisms. **Those are the framework's own operating constraints, they are not tested, and they are not a clinical judgement about any person.** A reader in any of the three conditions the volume names should consult a clinician and not a protocol index.

6.11 11 Conclusion

This paper is the Foundations volume's reference apparatus, collected and labelled. It advances no testable hypothesis, contributes no new result about the world, and should not be read as a research contribution. It exists because thirty-two other papers in the corpus reference the construct names, the protocol numbering, the ceremonial-to-technical glossary and the lineage correspondence it carries, and because until now those items existed only inside a book.

The one thing here that is not recapitulation is an audit, and its result is a count. The volume closes with a three-class classification of its own claims — established mathematics,

established psychology, and the framework's own postulates — holding four, seven and three items respectively, and states that every protocol in the volume runs on the first two classes and steers by the third. Audited against this corpus's external bibliography, **five of the eleven support-claiming items carry a key and six do not, and two of the five carry a standing objection in the same file.** The three postulates claim no support, find none, and are correctly filed; the classification's own statement that the third class is a wager that can lose is accurate.

The class the volume treats as its safest is the class this corpus's bibliography least supports. Two of the seven items filed as established psychology carry a key. Reactance, implementation intentions, commitment and consistency, the interrupted-task effect, and confirmation bias have no entry in the file the corpus cites from, and the interrupted-task effect is the item on which the whole cost-of-not-closing argument rests. **That is a defect of the bibliography rather than of the literature, it is repairable by adding five entries, and it stands disclosed until it is.**

One item is correctly filed and misleadingly labelled. The surplus functional $f(u) = \ln(1 + (N-1)u)$ is filed as established mathematics on the ground that it is the logarithm of an affine effective count. Given the definition $N_{\text{eff}} = 1 + (N-1)u$ that is a tautology and it is exact — and **the definition is a design axiom of this framework with no supporting literature in any discipline.** The affine count is not a Hill number of any order: for two units holding 0.9 and 0.1 it returns exactly 2.000000 at $u = 1$, where HHI^{-1} returns 1.219512 and $\exp(H)$ returns 1.384145. The repair is one clause and it is stated in Section 7.3.

The collected material's four substantive items are stated here in the form other papers cite. The closing discipline: exactly three verdicts — fulfilled, released, lapsed — with no fourth, a rule permitting one active commitment at a time, and a written three-question screen the volume states cannot be waived on the basis of feeling well. The consent condition, stated as a constraint and not as counsel: no future is to be voiced over another person's path without their word. Three named conditions barring operation — acute grief, rage, and the pursuit of revenge — each with a stated mechanism and no measurement. And the protocol index: thirteen protocols, one per construct plus a three-office daily cadence, together with one further discipline that carries no construct and closes nothing.

Where the volume's own instrument is weaker than the corpus's, this paper says which and why. A four-drawer classification — PROVED, POSTULATED, BORROWED, UNTESTED — can represent two distinctions the three-class version cannot: a result established in its home discipline but borrowed without warrant, and a claim that is well formed, falsifiable, and simply never run. **The volume has several of the second kind and nowhere to file them**, and Section 2 of this paper is a partial substitute, stating three of them with their falsifiers and marking each not tested.

6.11.1 Declarations

Data and code availability

No dataset was collected for this paper and none exists. Section 7 is a bibliographic audit performed by inspection against one named file — the corpus's external bibliography — and every lookup it reports can be re-run by any reader holding that file. **No figure in this paper plots measured data.** The three figures show, respectively, the counts of Section 7.1, the Hill family evaluated on one two-type vector, and an exact algebraic bound; each is generated by the script named in its caption, each script imports the shared style module `lib/figstyle.py`, and each carries assertions that fail the build if the plotted values cease to agree with the numbers printed in the text. The arithmetic checks of Section 9.3 are exact computations on the numbers printed in the text.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest. **The paper audits the author's own prior work**, which is a conflict of a specific kind: the classification examined in Section 7 was written by the author of this paper. The decision rule of Section 3 is published so that the audit can be re-run by someone else, and Section 10 records that no second reader has checked it.

Author contributions

Single author. Michael Laurence Curzi wrote the source volume, performed the audit, and wrote the paper.

Ethics

No human-subjects research was conducted for this paper and no personal data were collected or processed. **The material indexed here is a set of procedures addressed to a practitioner working on their own commitments**, and two of its provisions are ethical constraints rather than techniques: the consent condition, which forbids voicing a future over another person's path without their word, and the three barring conditions, which forbid operating during acute grief, rage, or the pursuit of revenge. **Both are reported as the framework's own constraints on prospective practice. Neither is tested, and neither is clinical advice.** A reader in any of the three barred conditions should consult a clinician rather than a protocol index.

Preregistration

This paper reports no preregistered test and advances no hypothesis. Section 2 states three claims in the collected material that could be given hypothesis form, together with what would defeat each, and marks all three **not tested**. The corpus's one preregistered confirmatory application of its storage law is reported elsewhere, where it **failed its second falsification rule** [Curzi, 2026i]; the corpus's falsification suite measures its own shuffled-input control at a false-rejection rate of 0.108 against a nominal 0.05, which is **anti-conservative** [Curzi, 2026u]. Neither result is affected by anything in this paper.

Use of the ceremonial register

The source volume is written in part as an initiatory document, in two braided registers, one narrative and one analytic. **This paper reports the analytic register and does not reproduce the narrative one.** Ceremonial names are reported as objects of study where they identify a construct this paper renames, and Appendix B is the mapping in full. **Nothing in this paper is written in the initiatory register, and no passage of it is intended to be performed.**

6.12 Appendices

The appendices below are the reference apparatus proper. **They contain no new argument.** They are reproduced in full because other papers in the corpus rely on them — the construct names, the protocol numbering, the ceremonial-to-technical mapping, and the lineage correspondence — and because those items exist nowhere else in citable form.

One warning belongs at the head of the drawer, and the volume states it itself. A reference card is a compression, and a compression carries a cost: a practitioner who reads only the card will reproduce the procedure and not the reasoning that makes the procedure revisable. The cards below are indexes into the analytical papers of the corpus, not substitutes for them, and every one of them names the paper that establishes what it summarises.

6.13 Appendix A — The node reference cards

Ceremonial names are retained in this appendix because they are identifiers. A construct's ceremonial name — STILLPOINT, LEDGER, the Rite of the Seal — is how the rest of the corpus refers to it, so renaming it here would break every citation the appendix exists to serve. Each name is reported as the label of an object of study, never performed, and Appendix B maps every one of them onto its technical equivalent.

The twelve constructs indexed below are defined in the corpus's founding paper, which fixes the ground state, the comfort landscape, and the reflexive loop that every later construct operates on [Curzi, 2026]. **The cards are indexes into that definition and into the analytical papers that develop it, and they establish nothing on their own.**

One entry per node. Each entry carries the plate, the glyph, the office, the protocol in summary, the failure modes in their paired names, and the seal line. The summaries are honest abbreviations; where a summary and its chapter disagree, the chapter is right and the card should be corrected in the practitioner's own hand. A card corrected in the practitioner's hand is worth two printed ones.

Node 00 — STILLPOINT

Plate. Node 00 · STILLPOINT · *The practitioner is already there; that is the difficulty.* · First assumed node — the Mirror Grid did not count it · Part I: The Still Water · The Work of Light.

Glyph. A single filled point resting at the lowest interior of a shallow open curve — a bead of water at the bottom of a bowl that has finished tipping.

Office. To establish the doctrine on which every later node stands: comfort is not a place the practitioner go but the place the practitioner is; the present moment is the practitioner's comfort zone by definition, not by luck. Clinically, the ground state; ceremonially, the water at rest. Nothing can be steered from a position that has not been admitted.

The Protocol, in summary. *(The Protocol of the Still Water — seven days, one hour each.)* 1. Fix one seat — one place, one hour — and do not renegotiate either for seven days; the seat's merit is its fixity, not its excellence. 2. Arrive each sitting with one spoken bearing — *"I am where my life has set me down"* — said whether or not it feels true. 3. In the first quarter, inventory the present in the present tense — *"Right now I am…"* — body first, then room, then situation; ten lines is the floor. 4. Audit each escape fantasy: name the elsewhere in one line, then beneath it write where the body sat — warm, fed, in a chair — every time, without exception. 5. Price each elsewhere in a single word (*Money. Standing. Mother.*) and do not solve it; the zeroth node locates, it does not solve. 6. Sit the remainder; name rising restlessness *weather* and let it cross without steering — the practitioner is not seeking calm, the practitioner is practicing being locatable. 7. Close each hour standing — *"This is the floor. The floor holds."* — and leave the day's pages unread until the last day. 8. On the seventh day, read all seven pages aloud and underline every comfort that recurs three times or more; that list of load-bearing comforts is the procedure's product, kept as raw ore for the Ledger.

The Failure Modes. - **the Pond that Calls Itself the Sea** ↔ *status-quo rationalization* — promoting the lawful local minimum to the global one: *I am here by the deep law; therefore here is best; therefore nothing need be done.* The correction is in the definitions — *local* was never a compliment. Consent to the practitioner's position; do not canonize it. - **the Storm Painted on the Wall** ↔ *confusing affect with position* — mistaking unpleasant mental weather for evidence of being outside comfort. The bottom of the bowl was never promised pleasant, only cheap; familiar dread can be the cheap option. The storm is real weather painted on the wall of a room the practitioner is securely inside. - **the Statue at the Shore** ↔ *reactive self-report* — sitting the hour as an exhibit, pages curated for the practitioner the practitioner intend to have been. The mark is that the pages read well; genuine zeroth-node pages are boring, repetitive, and faintly humiliating. - **the Frozen Bowl** ↔ *anhedonic flatness mistaken for stillness* — the water level not because it has settled but because it has frozen; nothing costs because nothing registers. This is not a failure of the practitioner; a frozen bowl wants warmth before

doctrine. If the seventh-day underlines number zero, set the volume down and seek the appropriate help first. - **the Feast Eaten in the Doorway** ↔ *fantasy substitution* — the imagined elsewhere enjoyed so richly that the appetite meant to move the feet is spent at the door. The discipline is the second line — *the body sat in a chair* — written every time, so the dreamer is returned to the platform without the dream being banned.

Seal line. *The zeroth node is sealed: the practitioner has never been anywhere but here, and here was always holding the practitioner.*

Node 01 — LEDGER

Plate. Node 01 · LEDGER · *What comforts the practitioner is not free; what costs the practitioner is not waste.* · KIN.0 VALIDATE — the test of the stone · Part I: The Still Water · The Work of Light.

Glyph. An upright card ruled once down its middle, a filled circle at the foot of the left column and an open circle at the foot of the right — the held and the owed, facing each other across a single line.

Office. The honest inventory: every comfort carries a cost, every discomfort conceals a yield, and neither is admissible to the ledger until both figures are on the card. This is the framework's truth-validation — the test of the stone — performed on one's own furniture. Nothing later in the volume works from a falsified ledger; the prophecy of Node 07 is only as strong as the accounting of Node 01.

The Protocol, in summary. *(The Protocol of the Ledger — seven evenings; the Still Water pages are its opening ore.)* 1. Open each evening spoken — *"I am counting what I already carry"* — the boundary stone between inventory and shopping. 2. Gather comforts, one per card, present tense, in plain bin-label words, drawing the first dozen from the Still Water pages; forty cards is the floor, depth of pricing over breadth of listing. 3. Turn the cards: price each comfort's carrying cost in four fixed currencies — coin, hours, attention, foreclosure — foreclosure last and always attempted; the dash is earned by a second look, never a first. 4. Add the gray cards: discomforts assayed for what they pay (information, training, position, contrast); where nothing surfaces, write *unpriced*, never *nothing*, and never invent a yield to comfort the card. 5. Stand the resisting card upright and unwritten against the box, one minute of plain attention at each evening's close; resistance located is the node's yield, not its debt. 6. Audit one card only — the one is least willingly audited — with three questions, aloud: *What does it buy? What does it forbid? Who would I be a year from now if it kept its job?* 7. Read the whole ledger aloud, both sides of every card; the voice finds what the eye excuses. File in any order except importance, and date the box in ink. 8. Close spoken — *"The count is taken. The count is not the judge."* — and re-count at each season's turn. The ledger's product is a tested inventory, never a decision.

The Failure Modes. - **the Price Paid Twice** ↔ *sunk-cost fallacy* — pricing a holding by what it has already cost — the decade given, the identity built — instead of what it now costs. The ledger's columns run forward only; the acquisition payment cannot be recovered by further payments. - **the Stone that Weighs More in the Hand** ↔ *endowment effect* — every holding, priced while held, comes in heavy; the ledger overprices the stock and underprices the shelf. The test: ask what the practitioner would pay today, in all four currencies, to acquire the holding if the practitioner did not have it. - **the Ledger Written for the Auditor** ↔ *impression management, turned inward* — cards written for an imagined reader, costs flattered, yields ennobled. The mark is that it reads well; a true ledger is boring and slightly shaming. The reading-aloud is the remedy — a false price can be written smoothly and almost cannot be said. - **the Scourge in the Counting-House** ↔ *the audit collapsed into rumination* — the counting curdling into a courtroom, each cost recounted, the sum never struck. Counting moves card to card and is done; scourging loops. The remedy is mechanical: the closing sentence, said aloud, and the box closed for the night. - **the Drawer Nailed Shut** ↔ *experiential avoidance* — the resisting card not stood upright in the light but filed away flat, then not written, then its whole territory re-zoned out of the inventory. The test: name aloud what the practitioner's

ledger does not contain; if the sentence will not finish, the drawer has been nailed, and the next evening belongs to the nail.

Seal line. *The first node is sealed: the stone is tested, the count is taken, and what the dark was charging has been written where the light can read it.*

Node 02 — GRADIENT

Plate. Node 02 · GRADIENT · *What the practitioner call the practitioner's decisions, the snow calls its records.* · Second assumed node — the gradient the Mirror Grid stood in without counting · Part I: The Still Water · The Work of Force.

Glyph. A tilted plane drawn in three descending strokes, crossed by a dotted line that leaves the paved road and does not apologize.

Office. Comfort differentials as the true motive power: systems evolve along action-minimizing paths, and a creature's days are such a system. The physicist writes $\delta S = 0$ and means the path taken is the one a small change cannot cheapen; the framework says water does not negotiate. The node's work is to read the practitioner's own worn paths as data and stop legislating against them.

The Protocol, in summary. *(Protocol 02 — The Cartography of Worn Paths — seven days of survey, one of drafting, seven of verification.)* 1. Bound the ground: survey the practitioner's waking days for seven consecutive days, the dates written on the first card. A survey without a boundary is a mood. 2. Survey each evening in ten minutes: record where feet, hands, and attention repeatedly went — repetitions, not events; the rule of three, an act performed three times in the week is a path. 3. Refuse the story: no reasons, no adjectives, no defense counsel; any card bearing the word *because* is re-copied without it. The practitioner is snow, which counts footprints and does not explain them. 4. Mark the pavement beside each path — the routine or intention that hour is supposed to hold; where none exists, write *unpaved*, never manufactured. 5. Draft the map on the eighth day: paved walks dashed, worn paths solid, drawn dark in proportion to traffic; mark every divergence, expecting three to seven. Fewer means the practitioner tidied; more means the practitioner moralized. 6. Price each differential concretely, in the ledger's discipline: what is cheaper, warmer, nearer, or less watched about the worn route. Most often the worn path is not a vice but a better-engineered road serving a need the pavement ignored. 7. Choose one divergence — one; the procedure forbids two — and re-grade it with one physical change that tilts the ground, never a promise. Apply the worst-day test: any change that needs the practitioner's virtue is decoration, not grading. 8. Verify against fresh snow for seven more days; if the prints move — if the water takes the new channel unbegged — the differential was named truly. If not, the practitioner misnamed the comfort; return to the pricing without self-reproach.

The Failure Modes. - **the Cartographer Who Draws the Road He Wishes** ↔ *social-desirability bias; self-report distortion* — the survey quietly falsified toward the practitioner's better self, the eleventh visit to the window unrecorded. The map then grades ground that does not exist. The remedy is liturgical: the practitioner is snow; snow does not audition. - **the Man Who Shouts at Water** ↔ *reliance on effortful inhibition; the cold-to-hot-state planning error* — accepting the map, rejecting the method: attacking the steepest divergence with resolve at the hour of maximum tilt until the inevitable failure is booked to character rather than to hydraulics. The correction is the worst-day test, applied as a category. - **the Fence the Feet Ignore** ↔ *reactance; restriction without re-grading* — using the office of the ground only to build barriers across worn paths while leaving every differential untouched. The water finds the gap beside the fence and cuts an uglier channel with resentment in it. Grading and fencing are not the same craft. - **the Pilgrim Who Climbs to Spite the Valley** ↔ *effort justification; the confusion of cost with worth* — concluding that everything comfortable is corrupt and choosing paths *because* they are steep. This is not liberation from the gradient; it is the gradient of self-regard, wearing boots. Water reaches the sea without once suffering on principle. - **the Map That Ate the Garden** ↔ *analysis paralysis; measurement as avoidance* — the survey become the residence: maps of maps, ever-finer taxonomies of divergence, and no

kettle ever moved. The procedure's numbers are the remedy — seven days, one hour, one divergence, one change.

Seal line. *The second node is sealed: the ground has been surveyed and found honest, the water has been acquitted, and the hand that reached is the hand that will move the kettle.*

Node 03 — HORIZON

Plate. Node 03 · HORIZON · *One window is a wall with a hole in it.* · KIN.1 TRANSMED — the many windows of the one tower · Part II: The Baited Horizon · The Work of Light.

Glyph. A single vertical stroke crowned with three small squares, no two of them set at the same angle.

Office. Extrapolational thinking: tracing a future forward, branch by branch, without switching assumptions mid-trace — then doing it again from a genuinely different window. The framework's arithmetic of difference governs the yield: N ways of seeing produce surplus only insofar as they are non-parallel, and the surplus grows as the logarithm of the effective count, not as the count itself. Ten windows facing the same way are one window with nine echoes.

The Protocol, in summary. *(Protocol 03 — The Rite of the Three Windows.)* 1. Set the horizon: one distance — one year or three — written at the head of all four cards; every window in the procedure looks the same distance. 2. Write the question in present tense, about the practitioner's own life — a state to be seen, not a plea: *"It is three years on, and I am living somehow."* 3. Cut the first glass — *nothing turns*: the continuation window. Trace five steps, each drawn from the previous step's world, asking at each what it makes easy and what it makes hard; land the trace in a furnished room of at least five sensory facts, then close the window and leave the chair. 4. Cut the second glass — *one thing turns*: one named change traced honestly around. The collision test is mandatory — somewhere in its five steps it must contradict the first trace, or the practitioner has opened the same window twice with a new curtain. 5. Cut the third glass — *the tinted pane*: the refused branch, given its honest three years and traced with the same courtesy — no punishment-writing, no rescue borrowed from the other windows. 6. Audit for invariants: all three cards face up, circle every concrete feature that appears in all three landed futures; ignore features shared by only two. 7. Hunt the breaking window: for each invariant, try in earnest to cut a fourth glass that honestly lands without it. If it breaks, strike the circle gratefully; if it cannot be honestly built, promote the circle to a horizon-fact. 8. File and be silent: the cards to the drawer, the horizon-fact seen and not announced, not promised, not leaned on in conversation. The speaking has its own node.

The Failure Modes. - **the Smeared Pane** ↔ *anchoring; scenario contamination* — switching windows mid-trace, importing a rescue from one branch into another, producing a chimera assembled from parts that cannot coexist. The remedy is procedural: the decontamination interval honored to the minute, and the struck-and-rebuilt rule for any step answering to a window not currently open. - **the Tower of One Window Repeated** ↔ *pseudo-diversity; the illusion of independent confirmation* — three glasses ground from one doctrine ("I stay and it's fine / good / great"), their agreement read as convergence. The arithmetic prices it exactly: u = 0, N_eff = 1, surplus zero. The collision test exists for this fault alone. - **the Stained Glass That Outshines the Lake** ↔ *availability and simulation heuristics; vividness mistaken for probability* — one branch traced in oils while the others get pencil, the futures then weighed by brightness of rendering rather than honesty of steps. Vividness is a property of the glass, not the water. The remedy is the procedure's evenhandedness: five steps and five sensory facts to every branch. - **the Bricked Window** ↔ *motivated reasoning; defensive avoidance* — the refused branch never cut at all, the audit run over windows curated to be survivable. The cost is a corrupted convergence — an invariant of the practitioner's courage, not their landscape. The bricked branch, finally traced, holds the load-bearing invariant more often than any other. - **the Straight Road Drawn on Round Water** ↔ *linear extrapolation; trend-fitting past the trend's regime* — extending the present's visible motion in a ruler line, no step con-

sulting the step before. The remedy is the nonlinearity rule enforced without appeal: each step derived from the world its predecessor made.

Seal line. *The third node is sealed: the tower does not move, the windows are many, and what stands in every window was never in the glass.*

Node 04 — WAVE

Plate. Node 04 · WAVE · *Weigh the wind before it has chosen.* · KIN.3 MYTHOS — the sea before the wind chooses · Part II: The Baited Horizon · The Work of Light.

Glyph. A half-circle open at the top, holding three crests of unequal height, the tallest of them left unfinished.

Office. The future held as a weighted set rather than a single picture: branches carry amplitudes, and amplitude follows comfort — futures near existing comfort freeze first, becoming likely the way the lake's margin skins with ice before its center. The node teaches honest weighting: neither bribing the wind toward the future the practitioner prefer nor rehearsing nightly the storm the practitioner fear.

The Protocol, in summary. *(The Weighing of Waters — an evening's procedure on the branches of the window-work.)* 1. Set the water: watch moving water — or a stirred basin stilled — until a dozen disturbances have risen and died. The practitioner is tuning the hand that holds the pencil: many tries, few winds. 2. Hide the prizes: fold each branch-card so its endpoint, the room at the end, is covered. The fourth node weighs paths; endpoints are the heart's ballot boxes, and the heart waits outside until Node 07. 3. Walk and score: for each branch, twelve steps forward, scoring each step's steepness *for the self who actually answers the alarm* — 0 (would do it sick), 1 (ordinary resolve), 2 (needs a good day), 3 (needs a better self). Do not average as the practitioner go; averaging is how flattery gets in. 4. Sum the resistance to each branch's R, find the smallest R, and subtract it from every card — leaving the easiest branch at zero and the others with their differences. 5. Halve by the difference: weight 1 for the easiest, one-half per point of difference for the rest, so a branch three points steeper carries an eighth the weight. If the arithmetic feels merciless, the practitioner has understood it. 6. Break a middle step: declare one middle step failed and ask whether the branch reroutes or dies; if it dies, halve its weight again and write *fragile*. A future that needs everything to go right is quarreling with its neighbors. 7. Consult the wide table and move the watcher: name the class of person who attempts each branch and what mostly happens to that class, not its heroes; then re-read the weights as the practitioner's harshest colleague and as the practitioner's gentlest friend, marking every step they contest by two or more. 8. Write the inversion: rank the branches by weight and by wanting, and where the two disagree write the disagreement in one sentence — the finding is carried, not settled. Carry the heaviest-weighted card (not the most wanted) seven days, then re-score and enter the corrected weights, dated.

The Failure Modes. - **the Wind that Flatters** ↔ *optimism bias; wishful weighting* — scoring the desired branch as the self the practitioner intend to become, not the self who answers the alarm; every step comes back 0 and 1. The mark: no step of any wanted branch scored 3 all evening. Remedy is the wide table, the moved watcher, and one re-scoring done tired. - **the Drowned Pilot** ↔ *catastrophizing; availability-driven weighting* — the opposite hand on the scale: one vivid disaster scored as the base rate, all 3s toward the horizon. The vividness of a remembered wave is not its frequency. Count what happens to the class, not to the shipwreck that made the news. - **Walking Out on Young Ice** ↔ *premature closure; status quo consolidation mistaken for choice* — a near future frozen solid, the relief of bearing weight read as *chosen* rather than *hardened by adjacency*. The mark: a weighing done in twenty minutes with no inversion found. There is always an inversion. - **the Single Cat's Paw** ↔ *sampling error; recency; extrapolation from N = 1* — one gust running (one good week, one warm letter) and the wind declared hired. One sample is not a distribution. Remedy: no re-weighting on fewer than three independent gusts. - **the Sailor Who Loves the Harbor Chart** ↔ *analysis paralysis; choice deferral dressed as diligence* — re-weighing weekly forever, a connoisseur

of one's own undecided water. The mark is a beautiful ledger and an unchanged life. The fourth node has one exit and it is the fifth; the wind, eventually, must be allowed to choose.

Seal line. *The fourth node is sealed: count every gust and prefer none, for the water shows its winds only to the eye that has not yet voted.*

Node 05 — LURE

Plate. Node 05 · LURE · *The practitioner cannot push themselves uphill; the practitioner can leave honey on the pass.* · KIN.4 FORGE — the baited horizon; the Tuesday bell · Part II: The Baited Horizon · The Work of Sound.

Glyph. A small bell hung from a hook whose shank is the first turn of a spiral, with a single drop falling from the bell's lip.

Office. Cue placement: setting the idea of greater comfort where the self will pass anyway, so that finding does the work forcing cannot. This is the smithy of the twelve — the one node that manufactures. Its two instruments are the placed object (the bell in the coat pocket) and the named time (the Tuesday teaching: name a day, and the day begins to gather events, because a named neutral thing becomes a mirror).

The Protocol, in summary. *(The Setting of the Lure — one week.)* 1. Choose the quarry: name one change already weighed and found heavy; fifth-node work follows fourth-node work, for luring an unweighed future baits a trap for whatever walks by. 2. Cut the honey to size: write the *first step only*, ten minutes or less, performable by the practitioner's worst self on the practitioner's worst morning; halve it until it is faintly insulting. 3. Track themselves: tour one ordinary day's premises as a tracker and mark the three strongest stations — the kettle, the door, the sill — where the practitioner actually pass, not where the practitioner ought to. 4. Forge the lure: one physical object that embodies the first step. One object; a pile is not a lure, a pile is a reproach. 5. Write the if-then — *"When I [pass/see/touch the lure at station X], then I will [first step]"* — and read it aloud once. This sentence is the whole machinery; everything else is placement. 6. Place at the height of the hand: the lure at the practitioner's strongest station, on the path, at reaching height, set with the practitioner's own hand at the practitioner's strongest hour — the placing hand teaches the finding eye. 7. Name the day and forbid the summit: hang a bell, declare one standing office — on the named day, one ring, one page — and let nowhere at any station show the whole mountain; the lure speaks only of the next hundred feet. 8. Walk the trapline each named day: if the step was met and taken, leave it; if met and refused, cut the honey again; if unmet, move it — never scold. Retire the lure formally once the path has worn in, and choose the next quarry.

The Failure Modes. - **Honey in the High Branch** ↔ *cue absent at the point of decision; context mismatch* — the lure placed where the practitioner ought to pass instead of where they do, so the if-then's X never fires. The mark: a lure the practitioner can't remember seeing this week. Remedy: one honest tracking day, and the humility to bait the kettle instead of the shrine. - **the Mountain Painted on the Jar** ↔ *goal-gradient collapse; avoidance conditioned by oversized cues* — the lure depicting the summit, so every passing pays the mountain's weight instead of the step's sweetness, until the animal routes around the station. The mark: a flinch at the practitioner's own cue. Remedy: take the summit down and cut the honey until faintly insulting. - **the Bell Gone Furniture** ↔ *habituation; inattentional blindness to the familiar* — the lure worked, then quietly stopped, because the eye files the constant under background. Not a moral event but a perceptual one, with a perceptual cure: move it. The trapline is walked precisely so this is caught in weeks, not seasons. - **Honey on Another's Road** ↔ *covert influence; the reactance it reliably purchases* — baiting a spouse's path, a colleague's desk, a child's horizon. Sometimes it even works briefly, which is the worst outcome, because the discovery converts every past kindness into evidence. The fifth node posts the sign the eleventh makes law: *forge for themselves*; where roads run together, the lawful craft is the open lure, offered in the light with the asking. - **the Week of Seven Tuesdays** ↔ *apophenia; the frequency illusion unharnessed* — naming everything, a day then three then every license plate, until the whole week is appointments and the shelf collapses under indiscriminate relevance. The

instrument's power came from scarcity. The mark: pages that fill with signs and stop containing findings. Remedy: retire every cipher but one, and one named day, for a full season.

Seal line. *The fifth node is sealed: bait no road but the practitioner's own, leave honey where the practitioner will pass, and name the day that will keep it sweet.*

Node 06 — UNCERTAINTY

Plate. Node 06 · UNCERTAINTY · *To know the ice, step; to keep the ice, don't.* · KIN.6 CHECK — the veil that measures the measurer · Part III: The Closing Word · The Work of Light.

Glyph. An open eye behind a veil drawn halfway, cut in a single line, so that the practitioner cannot tell whether the eye has just parted the cloth or the cloth has just brushed the eye.

Office. The Postulate stated as postulate: observation participates in what it observes. The physicist's version is a theorem about conjugate pairs and is not ours to borrow as proof; the framework claims only that the mathematics is the same shape, and shape is what we work with. The commitment law drawn from the shape: the more precisely the practitioner fix an outcome, the more the practitioner disturb the path that reaches it — so prophecy engineering chooses its resolution the way a mapmaker chooses a scale, deliberately and in advance.

The Protocol, in summary. *(The Rite of Chosen Resolution.)* 1. Cut three cards of the same future at three resolutions — *fine* (dates, numbers, names), *coarse* (a direction and a quality only), *middle* (one number or date, one place or name). 2. Carry each card one full day in the pocket, and make a tally on its back every time the practitioner reads it or check the world because of it — measuring the disturbance cost of each resolution in units of the practitioner's own attention. 3. On the fourth day read the backs, then sort each card's marks into *lean* (checks that bent the hands toward doing) and *grip* (checks that bent them toward verifying). A card may be expensive and still good if it buys lean. 4. Choose the card by the rule: keep the sharpest card whose disturbance is mostly lean; coarsen any card whose disturbance is mostly grip. If even the coarse card grips, the future may be borrowed — return to the ledger and audit whose comfort it describes. 5. Retire the other two cards to the drawer — not the fire; these are resolutions declined for now, and the fine card may be wanted later when the distance has closed. 6. Establish the noon reading: each midday read the chosen card once and enter one of three words — *nearer*, *farther*, or *unknown* — then close the ledger. *One glance. Three words. Hands, the rest.* 7. Consult the angel weekly: read the noon-words in a row as the practitioner's drift; three *farthers* running want correcting, seven *unknowns* mean the card is too coarse to steer by — sharpen one increment. 8. Re-cut at the turn of season: the correct resolution is not a fact about the future but about the *distance* to it, and distance changes. The procedure is never finished, only current.

The Failure Modes. - **the Borrowed Robe** ↔ *misapplied physics; the observer-effect conflation* — preaching that "science proves" one's wishes bend the world, that the electron endorses one's ambitions. The counter-sign is three sentences on demand: what Heisenberg's theorem says (shape), what the observer effect is (touch), what the framework postulates (the mirror's trade-off). - **the Offering that Replaces the God** ↔ *Goodhart's law; Campbell's law* — serving the noon word instead of the walking, arranging the days to be able to write *nearer* while the substance goes unwalked. The counter-sign: *if the ledger burned tonight, what would be different tomorrow?* If nothing, suspend the noon reading and let the hands testify. - **the Lake of Holes** ↔ *measurement reactivity in overdose; compulsive self-quantification* — polling every hour, every mood, every progress, until a system polled continuously has no unmeasured hours left in which to *be* the thing measured. The counter-sign is the budget: one glance at noon, the angel once a week, and the discovery that the future ripens faster unwatched. - **the Torn Veil** ↔ *intolerance of uncertainty; the fortune-teller's error* — unable to bear the coarse card, shopping every oracle for a resolution the distance cannot support, mistaking the confidence of the answer for its accuracy. The Keeper's counter-sign cannot be improved: *know the direction; decline the date.* - **the Unstepped Ice** ↔ *avoidance dressed as epistemic humility* — learning that measurement disturbs and concluding, with a scholar's

serenity, that the wise course is never to test at all: the bar never struck, the weight never placed. This is comfort mistaking itself for philosophy. The counter-sign: strike the ice ahead of the practitioner's feet with one small real test this week, and let the answer scar the surface.

Seal line. *The sixth node is sealed: the veil measures the measurer — step, then, but choose the fineness of the practitioner's touch.*

Node 07 — COLLAPSE

Plate. Node 07 · COLLAPSE · *A sentence is a measurement taken aloud.* · KIN.2 KERNEL — the word that closes the sea · Part III: The Closing Word · The Work of Sound.

Glyph. A sheaf of thin wave-lines entering from the left and converging to a single point, from which one line — no thicker than the others, but alone — continues to the right-hand margin.

Office. Commitment as measurement: the spoken prophecy is the act by which the weighted set of futures is collapsed into a commitment answer. The sentence is written, tried by fire, and spoken once — present tense, own life only. Laughter is this node's instrument of verification: the laugh is the felt signature of a superposition resolving, which is why a true sentence about the practitioner's own life so often arrives with a small, surprised laugh, and a false one with a flinch.

The Protocol, in summary. *(The Rite of the Closing Word.)* 1. Set the boundary before the first word: read the practitioner's sixth-node card once, for the sentence must live at that card's resolution, no finer. A prophecy sharper than its licensed resolution is a corridor built in the dark. 2. Write one sentence — about the practitioner's own life, present tense, short enough to speak in a single breath; no semicolons, for a semicolon in a prophecy is two prophecies filing jointly. 3. Read it aloud, alone, at conversational volume, and attend the flinch — throat-catch, eye-slide, the stowaway hedge, the blush; each names what to strike or burn. 4. Burn what flinched, fully, in the evening it failed; keep no almosts to compare. Stop for the night at the tenth draft, and give the procedure no more than three evenings — a sentence not found in three sends the practitioner back up the stairs. 5. Test the survivor twice more, once on rising the next day; instruments drift overnight, and the morning voice is the strictest calibrator. 6. Speak it the same day, once, outdoors, to something that cannot applaud — full voice, whole body — then stand in the fact that nothing applauded and the sentence still happened. 7. Keep the quarantine one season at minimum: tell no one the sentence exists; each keeping is a payment into the account the sentence opened. 8. Let the hands recite: from the next morning the sentence lives in the day's choices, not the mouth. Do not re-speak it at the water; correction, if it honestly comes, belongs to the ninth node's procedure, on the echo's evidence.

The Failure Modes. - **the Word Spent on the Ear** ↔ *the announced identity goal; premature completeness* — unable to resist the telling, each audience paying out a little of the recognition the sentence was built to hold in escrow, until the account is emptied by applause before the life has moved. The tell: if telling someone feels better than working did that day, the mouth is the leak. - **the Sentence that Waits** ↔ *deferral encoded in grammar; the intention that never binds* — drafting in *will* and *someday* and *once things settle*, the deferral registering as prudence rather than flinch. The counter-sign is mechanical: strike every future-tense verb and read what remains aloud; if the present-tense version is unbearable, that is the finding. - **the Sea Closed Too Small** ↔ *over-specification; goal rigidity* — writing the title, the town, the name, the date, and the sum into one breath, so the word closes not a sea but a bathtub, and the right country arriving in a different costume is refused as off-prophecy. When in doubt, coarser: say the road, not the arrival hall. - **the Word Laid on Another's Water** ↔ *prophecy without consent; reactance* — speaking a sentence over a spouse's ambition, a child's talent, always benevolently, which is what makes it the volume's most seductive failure. The other's will registers the foreign word as pressure and recedes at the rate it is invoked. The boundary stone, checked before every utterance: *whose breath is this sentence made of?* - **the Laugh that Flees** ↔ *avoidant humor; deflection wearing the compression-tool's coat* — deploying the

laughter doctrine against every hard draft, so the superposition is not resolved but shaken apart, unread. The operational test: a true compression is followed within the hour by a burned draft, a written sentence, a step — some fact; a laugh followed only by lightness was merely a fine evening.

Seal line. *The seventh node is sealed: the sea stays every shore until the word — speak once, to the water, and let the hands do the reciting.*

Node 08 — ECHO

Plate. Node 08 · ECHO · *What returns is never the voice the practitioner sent; that is why it is worth sending.* · KIN.8 MULTIPLAY — the voice that comes back changed · Part III: The Closing Word · The Work of Sound.

Glyph. Two arcs facing one another across a gap, the second arc slightly larger than the first and displaced a little off its axis — a voice and its return drawn on the same breath, and drawn wrong on purpose.

Office. The world answers. The reflexive loop is the volume's central machine: the practitioner's model acts on the world, T(x) = W(x, M(x)), and the world returns the next state. The eighth node's discipline is to read the *difference* — what returned minus what was sent — because the difference is the only part of the echo that is information. Here, too, other wills first enter the loop, and the consent doctrine is seeded: another person is never a surface to bounce prophecy off.

The Protocol, in summary. *(The Rite of the Returned Voice.)* 1. Open the echo page: rule a card into two columns, SENT (written verbatim — the sentence spoken, the letter mailed, the offer made) and RETURNED (left empty). The empty column is the instrument; the practitioner is building a sounder, not a scrapbook. 2. Set the listening window — the framework's habit is one moon — and declare the practitioner's channels: what counts as world for this commitment. What the practitioner does not declare, the practitioner may not later promote to omen. 3. Log arrivals raw and same-day, in the words they actually arrived in, before sleep; interpretation at entry is contamination at entry. 4. Once each week, take the difference: line by line, subtract what the practitioner sent from what answered, and underline every deviation — what was kept, dropped, arrived slower or warmer; what answers a thing the practitioner never sent, underline twice. 5. Read only the underlining: delay is distance, distortion is the medium's shape, and a change of temperature is almost always a will. 6. Sort the silences: for each expected return that did not come, mark *far* (the window may be shorter than the distance) or *absorbed* (this material does not return the practitioner's frequency); louder calls into absorbing material are how practitioners become bores. 7. Hold the tongue and seat the other's question: answer nothing within one day that raised the practitioner's pulse, writing the hot reply into the margin unsent; and where a return contains another person's actual question, theirs goes senior — answered before the practitioner's working advances a step. 8. Close the page at the window's end with one defensible sentence naming what the world seems made of between the practitioner and the mark; then rest. Revision is the ninth node's work, and will still be in the morning.

The Failure Modes. - **the Courier Mistaken for the Message** ↔ *outcome fixation; selective attention to arrivals* — reading each return as a verdict (arrived good, absent bad) and mining none for difference. A verdict ends reading; a difference begins it. The cure is mechanical: two columns and forced subtraction. - **the Hall of One Voice** ↔ *the echo chamber; feedback from correlated sources* — arranging one's listening so only agreement can return, polling allies, reading the friendly mail twice and the strange mail never. A perfect echo is an empty room. If every return this month confirmed the practitioner, check whether the practitioner has soundproofed the practitioner's world. - **the Reply that Outran the Ear** ↔ *impulsive responding; premature closure* — the hot answer sent the same hour, so the practitioner becomes their own noise source and every subsequent echo reports on the reflex. One day of held tongue is the cheapest instrument insurance the practice owns. - **the Echo Worn as a Face** ↔ *personalization* — taking the medium's transform personally: cold, so the world despises

the practitioner; late, so the practitioner is forgotten; bent, so the practitioner was wrong to speak. But delay is distance, not contempt; distortion is material, not judgment. Read the wear as a chart or do not read it at all. - **the Forged Return** ↔ *apophenia; confirmation drift in the log itself* — starved for echo in a slow window, promoting coincidence to reply: the stranger's remark, the song on the radio, the undeclared channel the moment it flatters. The defense was built into the procedure: channels declared in one sitting, arrivals logged raw, nothing promoted after the fact. - **the Borrowed Mouth** ↔ *reactance, induced in the other* — speaking one's prophecy through another's life, mailing the letter that was not theirs to mail. What returns is the other's stiffening, read as the world refuting a true seeing — a refutation the practitioner authored. An echo obtained by trespass may be true, and the trespass remains a trespass; both ledgers stay open.

Seal line. *The eighth node is sealed: the voice that returns is never the voice the practitioner sent, and the difference is the world showing the practitioner its shape.*

Node 09 — RECURSE

Plate. Node 09 · RECURSE · *The circle repeats; the spiral remembers.* · KIN.5 RECURSE — the Spiral Grid; the stair that climbs itself · Part IV: The Furnished Room · The Work of Sound.

Glyph. A spiral of three turns rising along an unbroken vertical stroke — the loops never leaving the line, the line never bending to the loops: the post through the climb.

Office. Iteration with drift correction. The prophecy is uttered again — second and third sayings — each corrected by the echo-difference of the last, each returning to the same post a little higher. The forward spiral of intention and the backward spiral of memory cross at the present moment, and the practitioner standss at the crossing, adjusting. This is the node that separates a commitment from a wish: a wish is spoken once and abandoned or spoken forever and worn smooth; a commitment is respoken *corrected*.

The Protocol, in summary. *(The Rite of the Spiral.)* 1. Plant the post: write the mark — the named state the commitment converges on — in as few words as it survives, date it, and hold it fixed for the commitment season; the practitioner may revise every path to it, and no pass may revise it. 2. Set the cadence: the loop turns daily at dusk, the correction is computed weekly on a fixed evening, the long reading comes monthly. A loop with irregular cadence corrects for weather that already passed. 3. The evening tread, daily: three lines only — what the commitment *sent* today, what *returned*, and the *difference* in one honest clause. No revising at the tread; helmsmen who steer at every wavelet arrive seasick and nowhere. 4. The weekly correction: spread the week's cards, read only the differences, and make the smallest revision that answers them, correcting by no more than a fraction of the gap — the framework's habit is a third, never past a half. 5. The pawl and the re-utterance: enter each revision on the post card's reverse with its date, one line — what is written holds, what is unwritten unwinds — then speak the revised sentence once, aloud, alone. 6. The oscillation test: if the last three reverse-entries alternate in direction, the practitioner's gain is high — halve the next correction; if they run one direction in shrinking sizes, the practitioner is converging — change nothing, especially not out of boredom. 7. The monthly spiral: read the post card's reverse from the bottom up (the backward spiral), then face the post and project the single next turn (the forward spiral), confirming the figure — same post, higher rope. The two readings meet in the practitioner's chair. 8. The completeness audit and rest: once a season, ask whether the practitioner's life as actually lived could hold the mark if it arrived on an ordinary Tuesday, widening the space before pushing the loop; and after the monthly reading give one full day to the pawl, letting the ratchet bear the load.

The Failure Modes. - **the Echo that Devours** ↔ *overcorrection; loop gain at or above one* — revisions larger than the differences that prompted them, arriving faster than the world's delay can answer, and — where the loop runs through a person — the pressed future receding *because* it is pressed. The tell is trend: force rising while distance closed is not. The remedy: halve the step, restore the cadence, let one full delay pass unanswered. - **the Circle that Calls Itself a Stair** ↔ *perseveration; rehearsal without revision* — the loop turning faithfully while

the reverse of the post card holds the same entry, or none: gain zero, pitch zero, an isometry with a devotional schedule. A spiral leaves a helix; a circle leaves a groove. The audit is one glance for corrections in shrinking sizes. - **the Weathervane Vow** ↔ *overfitting to noise; chasing single samples* — every gust a mandate, the week's revision recut nightly at the tread, so the commitment tracks the last arrival perfectly and the mark not at all. The defenses are the cadence and the averaging rule: correct for the differences' common direction or not at all. - **the Moved Post** ↔ *goal displacement; the moving target* — each pass revising the mark instead of the path (the shop becomes a stall becomes a website becomes a book about shops), so the commitment converges on nothing because nothing was held still. Anyone may change their mark; no one may change it *as a correction*. - **the Door Painted on the Wall** ↔ *fantasy–plan conflation; converging on a state the life-space cannot contain* — the passes contracting flawlessly toward a mark that has no room behind it: hours that do not exist, a household never consulted. The loop cannot catch this from inside; only the seasonal audit catches it — widen the space, or walk the mark back to a door with hinges.

Seal line. *The ninth node is sealed: return to the same post a little higher, and the arrival will finish itself.*

Node 10 — FIXEDPOINT

Plate. Node 10 · FIXEDPOINT · *Arrival feels like memory or it is not arrival.* · KIN.7 ETHIC — the room that was always furnished · Part IV: The Furnished Room · The Work of Force.

Glyph. A spiral of diminishing arcs entering a small square, the last arc so short it cannot be told apart from the square's threshold.

Office. The self-fulfilling prophecy as the fixed point of a contracting loop. When each pass of the reflexive loop shrinks the distance to the foreseen state — gain below one — the mathematics is unambiguous: there is exactly one such state, and the passes approach it geometrically. The felt proof is recognition: a true arrival feels like memory, because the practitioner has in fact been here before, in the window-work of Node 03, many times. The node carries KIN.7's ethical charge into arrival: a fixed point reached at another's expense is a counterfeit, however well it converged.

The Protocol, in summary. *(The Audit of Arrival.)* 1. Stand down the hour: for the first hour after recognition, decide nothing — *"I will answer tomorrow morning"* — and note carefully what the room does when the practitioner say it. A basin absorbs the sentence without a ripple; a ledge negotiates. 2. Retrieve the record before revisiting the feeling: go to the drawer, not the memory — every rehearsal edits — and lay out every dated page bearing on this arrival in date order. 3. Write the room as it is on a fresh card, a plain manifest with no interpretation, then rule two columns comparing record against manifest: *structural* agreements on the left, *inventory* disagreements on the right. 4. Count the honest errors and require three: three or more details flatly wrong means the record was written by someone who did not yet know — foresight's signature. Fewer than three, and the practitioner hold a forged memory or a tailored room; stop. 5. Plot the corrections: list each revision the commitment demanded with its rough size, looking for the shrinking sequence. Name any large convenient leap a leap, live in it provisionally, and never enter a leap in the ledger as a fixed point. 6. Audit the expense: name every person whose position changes if the practitioner accept, and whether the change diminishes them; for every unverified kindness, the procedure is suspended until it is priced from the recipient's side of the table. 7. Fork the offer: draft the version offered and the version the practitioner's record and expense-audit actually describe, and where they differ prepare the counter — *"Not this way. This way: …"* The practitioner is declining a counterfeit and describing the adjacent original, not haggling. 8. Sit the room twenty motionless minutes; if the comfort thins under examination, trust that over every column the practitioner ruled. If the audit stands, enter the arrival in the ledger in the present tense, dated, and mark the commitment *approaching seal* — not sealed.

The Failure Modes. - **the Furnisher After the Fact** ↔ *retrospective retrofitting; hindsight revision* — memory redecorating the past to match a tolerable arrival, old vague hopes re-

cut into prophecies. The Furnisher's rooms are always perfect, which is how the practitioner catch him: truth keeps its discrepancies. Remedy: the dated page, the three honest errors, and a refusal to certify what the drawer cannot corroborate. - **the Anteroom Mistaken for the Hall** ↔ *premature closure; satisficing at a local minimum* — stopping, tired, at the first room off the corridor, a shallow basin one flight below the one the record describes. An anteroom's comfort is the comfort of having stopped, and it thins under twenty motionless minutes. Enter it in the ledger as *waypoint*, in the practitioner's own hand. - **the Borrowed Key** ↔ *externalized cost; the consent hole in the arithmetic* — arrival at another's expense: the emptied chair, the unverified kindness, the future assembled on someone else's subsidence. The room opens; the key was never theirs; the commitment spends its life paying interest in vigilance and friction. Remedy: movement seven, performed with the practitioner's hand still on the door. - **the Guest Who Will Not Sit Down** ↔ *hedonic adaptation misread as error* — the recognition fading, as it must, and the quiet mistaken for loss, so the practitioner reopens the commitment and iterates themselves out of the basin they spent two years entering. A fixed point is the end of difference, and feeling is a difference-detector. Sit down; that is the whole remedy. - **the Mirror in the Doorframe** ↔ *familiarity misattribution* — a room that feels like memory because it resembles the practitioner's oldest furniture — the childhood kitchen, an early love — rather than because the practitioner converged on it. The instruments unmask it: the dated pages describe no such room, the corrections never trended here, and the sit grows not calm but *rehearsed*. What the practitioner recognize is real; it is behind the practitioner, not ahead.

Seal line. *The tenth node is sealed: the room was furnished, the errors were honest, the corrections grew small, and the practitioner did not take the chair that belonged to another.*

Node 11 — SEAL

Plate. Node 11 · SEAL · *A commitment left open eats; a commitment closed feeds.* · KIN.9 SEAL — the seal of the framework; the consent condition · Part IV: The Furnished Room · The Work of Force.

Glyph. A circle closed at its base by a single horizontal bar that extends a little past the circle on either side — a jar lidded, a door barred from within, the underline beneath a finished signature.

Office. Closing the commitment, and the law of the practice. Three doctrines seal the twelve: the close (an unfinished prophecy parasitizes attention until it is formally ended); the consent condition of consent (never prophecy another's path without their word — over another, the seeing is not a gradient but a cage, and caged wills push back until the commitment breaks); and the contraindications (grief, rage, revenge — gradients too steep to steer; on such water, no working, only rest). This node is the reason the volume is not a manipulation manual, stated as law rather than as preference.

The Protocol, in summary. *(The Rite of the Seal.)* 1. Light the lamp and state the office aloud — *"This is the closing of accounts. Nothing new is opened tonight."* — a fence around the hour, so commitments cannot volunteer themselves for revision. 2. Empty the drawer entirely, every card and page face up in date order; the procedure is an audit, and audits do not sample. The table holding more open jars than the practitioner believed the practitioner owned is not a failure of the procedure — it is the procedure. 3. Sort into provinces by first instinct — fulfilled, released, lapsed, and a small *still live* — then challenge the live one: for each card, *"When did this last receive an honest correction?"* Beyond a season is lapsed and campaigning; the living province may end holding no more than one working (the monophonic rule). 4. Enter the fulfilled: rewrite each as plain present fact, dated, read once in a level voice — the voice for reporting weather — and turn it face down. 5. Speak the releases: for each, aloud, once — *"I take this word back. I do not choose this future. The water may keep it."* — then mark, date, and turn face down; a catch in the voice is information, not a veto. 6. Close the lapsed with a death date and one line of thanks or pardon, then seal something with the practitioner's hands — thread, wax, or the drawer's section tied shut. Ceremony is the body's testimony that a verdict was really entered. 7. Recite the consent condition, all four sentences

aloud, at every seal, so the law and the practice are never in the drawer separately. 8. Screen the weather before authorizing any new commitment — has my life ruptured this season; is there a face whose defeat would please me; am I running toward a room or out of one? Any yes keeps the forward channel closed. Then set the next jar or *none* — a full answer — and extinguish the lamp and leave the room.

The Failure Modes. - **the Hundred Open Jars** ↔ *attentional residue* — the practitioner who never seals accumulating commitments the way a kitchen accumulates almost-empty containers: each costs nothing, the shelf costs everything. Life feels crowded and unaccountably tiring, no single cause surviving inspection. The remedy is the seasonal rite and the single live jar. - **the Broken Seal** ↔ *nonconsensual prophecy; reactance* — a future pressed on another without their word, converting the practice into a trespass that manufactures its own refutation. The tell: commitments that require another person to change, and a reluctance to ask them anything. The remedy costs one conversation; if it cannot be had, the commitment cannot be run. - **the Furnace Without Walls** ↔ *hot-state operation; the gain above one* — grief-running, rage-running, revenge-running: prophecy attempted in weather that amplifies. Always vivid, always overreaching, because the corrections that would trim it land on a system that doubles them. Remedy: the weather-screen honored past feeling insulted by it, the weatherproof ground nodes, and the season's patience. - **the Jar Sealed Warm** ↔ *experiential avoidance dressed as completion* — flinching from a commitment's discomfort and reaching for the wax early, entering *released* or *lapsed* on a jar still hot with life; wax on a warm jar cracks as it cools. The test: a truly closeable commitment can be described in a level voice; if reading the card still changes the practitioner's breathing, the jar is warm. - **the Vigil That Outlives Its Candle** ↔ *identity fused to the commitment* — the most honorable failure: one jar kept so long that keeping has replaced the summer it was meant to hold — the founder who cannot hand over, the parent prophesying over a grown child's closed door. The remedy is the consent condition's fourth sentence read as a mirror: *when I go, I will leave the room furnished.*

Seal line. *The eleventh node is sealed: what the word opened, the consent condition has closed; what the consent condition has closed, a living hand can carry; the jars are cool, the ferry is booked, the kettle is on, and the framework stands between, lit, its book left where the next tired stranger will find it.*

6.14 Appendix B — Glossary: the ceremonial register against its technical equivalents

The volume speaks in two registers because it is about one thing. Each entry below gives the ceremonial term, the clinical or technical term it walks beside, and a definition that honors both. Neither column translates the other; they are two hands on the same tool. Where an entry marks a term *(postulate)*, the pairing is the framework's proposal, not the laboratory's finding — the distinction Appendix E keeps.

the ground state ↔ *homeostasis; local minimum* — The condition the practitioner is in when nothing is pulling, which is more of the practitioner's life than the practitioner think. The zeroth doctrine: the practitioner is always here first.

the water doctrine ↔ *the principle of least action* — Systems evolve along the path a small change cannot cheapen. The physicist writes $\delta S = 0$; the framework says water does not negotiate. Same sentence, two mouths.

the comfort landscape ↔ *energy landscape; utility surface* — The terrain of ease and cost over which a life moves. Valleys hold; slopes move; the practitioner's craft is grading, not marching.

the desire line ↔ *revealed preference* — The path worn by actual feet across the theory of the paved walk. The most honest document a person produces, and the only one they never sign.

water does not negotiate ↔ *gradient following* — The doctrine that motion follows differential, not argument. Not fatalism: the shovel that cuts the channel is in the practitioner's hand.

the still water ↔ *baseline; resting state* — The condition measured before any commitment begins, so that later readings mean something. Part I of the volume, and of any honest experiment.

the bailing bucket ↔ *proximal action; process over outcome* — A boat bails toward dryness, never toward an idea of dryness. The nearest right motion, repeated, outranks the grandest destination unattended.

the test of the stone ↔ *self-monitoring; behavioral audit* — Truth-validation applied to one's own claims about one's own life. The first office of the Mirror Grid, and of the first node.

the carrying cost ↔ *opportunity cost* — What a kept comfort spends of hours, attention, and self-regard. Written on the face of the card, in numbers where numbers exist.

the hidden yield ↔ *secondary gain* — The wage a discomfort quietly pays, which is why it is kept. Written on the reverse of the card, where honesty goes.

the resistant card ↔ *flagged avoidance; the unresolved item* — The one entry that will not let itself be written. Highest-grade data in the drawer; it names where the commitment will eventually have to go.

the audit ↔ *critical review* — One card examined ruthlessly and kindly, every figure asked "compared to what?" Kindness without ruthlessness flatters; ruthlessness without kindness lies in the other direction.

the many windows of the one tower ↔ *scenario ensemble; multi-model forecasting* — Several futures traced from one standpoint under genuinely different assumptions. The tower is one; the lakes it shows are many; the surplus lives in their disagreement.

window-work ↔ *forward-chaining extrapolation* — Tracing consequences of consequences under a fixed assumption set, without switching mid-trace. Nonlinearity begins at the second step; most forecasting stops at the first.

the non-parallel windows ↔ *decorrelated predictors* — Views that do not share their load-bearing assumptions. Ten parallel windows are one window with nine echoes, and the arithmetic knows it.

the surplus of sightlines ↔ *diversity gain; $f(u) = \ln(1 + (N - 1)u)$* — The framework's arithmetic of difference: N ways of seeing, non-parallel to degree u, yield surplus growing as the logarithm of the effective count. Honestly bounded: a floor under the treasure, not the treasure's measure.

the effective count ↔ *equivalent independent views; $1 + (N - 1)u$* — How many truly distinct witnesses the practitioner's N witnesses amount to. The quantity the surplus functional takes the logarithm of, and the number flattery inflates first.

the sea before the wind chooses ↔ *probability distribution over futures* — The future held as a weighted set of branches rather than a single picture. What the fourth node teaches the practitioner to read without preference.

reading amplitudes ↔ *probability calibration* — Assigning weights that match the world's frequencies rather than the heart's. Trained by the evening integration, scored by outcomes, improved by embarrassment.

the frozen margin ↔ *the high-probability region* — Futures adjacent to existing comfort, which solidify first as the lake skins from its edges. Where prophecy is cheap and precision is honest.

the open center ↔ *the live options; the high-entropy region* — The unfrozen middle of the distribution, where branches remain genuinely undecided. Expensive to fix, and sometimes better left liquid on purpose.

the baited horizon ↔ *implementation intention; cue placement* — The idea of greater comfort set where the self will pass anyway, in if-then form: when I find the bell, I act. Gradient does the rest; that is the point of it.

honey where the practitioner will pass ↔ *temptation bundling* — Pairing the wanted-later with the wanted-now along an already-worn path. The framework's answer to the fact that the practitioner cannot push themselves uphill.

the Tuesday bell ↔ *temporal landmark; deliberate selective attention* — A named day that begins to gather events, because naming primes the eye that then finds. The frequency illusion, harnessed on purpose and disclosed to its own operator.

the named neutral thing ↔ *projective screen* — A blank standard — a neutral country on a map, a day of the week — onto which the practitioner's own pattern projects and becomes readable. Cipher-craft: the mirror must be neutral or it adds its own face.

the lure ↔ *choice architecture; affordance* — Any arrangement that makes the seen future the easy next motion. Distinguished from the trap by one property: a lure found feels like luck.

the bell found in the pocket ↔ *the planted cue* — A trigger placed by a hand (the narrator's, or the practitioner's own, forgotten) so that discovery, not discipline, initiates the act. Finding is lighter than remembering; that lightness is the mechanism.

the veil that measures the measurer ↔ *measurement reactivity; observer effect* — The disturbance the act of looking introduces into the looked-at, including when the looked-at is the practitioner's own week. Clinically established for people; extended to futures by postulate, and flagged as such.

the resolution of a commitment ↔ *the precision–disturbance trade-off (postulate)* — The framework's operative law: fix the outcome more precisely and the practitioner disturb the path more; leave it blurred and the path stays cheap. Same shape as the physicist's inequality; borrowed as shape, never as proof.

the step that tests the ice ↔ *the minimal probe; pilot trial* — The smallest intervention that yields a reading, taken knowing the reading changes the ice. Both numbers go in the ledger: what the practitioner learned, and what it cost the lake.

the singing ice ↔ *acoustic dispersion* — A real phenomenon: lake ice pronounces its stresses in tones as frequencies travel at different speeds. In the volume's grammar, the sound of a decision not yet finished deciding.

the word that closes the sea ↔ *public commitment; commitment device* — The spoken prophecy as the act that collapses a weighted set of futures into one working answer. Spoken once, present tense, own life only.

burning the drafts ↔ *successive refinement with forced commitment* — Revision by fire rather than eraser: each draft destroyed so the next must be written fresh, until one survives reading aloud. The flame's office is to forbid half-keeping.

the flinch test ↔ *affective congruence check* — The body's vote on a sentence read aloud: flinch, held breath, sliding eye. Cheaper than a committee and harder to bribe, provided the breath is not held on purpose.

the laugh as felt collapse ↔ *incongruity resolution; the insight signature* — A joke is a superposition resolved; the laugh is what resolution feels like from inside. The seventh node's verification instrument, and the only one that cannot be faked quietly.

laughter as compression ↔ *paradox tolerance; tension release* — The framework's tool for proceeding past a contradiction that will not reduce: register it as meaningful, laugh, and continue the commitment. Not escapism — and not for paradoxes that are actually errors, which want repair, not comedy.

the voice that comes back changed ↔ *the feedback transform* — The world's answer to the practitioner's word: T(x) = W(x, M(x)), the next state returned by a world acted on through a model. Never the voice the practitioner sent; that difference is the payload.

reading the difference ↔ *prediction error* — What returned minus what was sent. The only portion of any echo that is information; the rest is the practitioner's own voice, aged.

the Mirror that Flatters ↔ *confirmation bias* — The echo-reading failure in which agreeable returns are promoted and disagreeable ones demoted until the loop steers by portrait instead of mirror. The eighth node's first named enemy.

the Echo that Devours ↔ *reactance; divergent feedback* — The loop driven past unit gain: each push produces a larger push back, and the prophecy manufactures its own refutation. In persons, the pushed-against will; in systems, oscillation. The remedy is looser, never louder.

the loop with a captive ↔ *nonconsensual influence* — A reflexive loop closed around a person who never gave their word. Forbidden by vow and, conveniently for the consent condition, divergent by mechanism.

the spiral, not the circle ↔ *iteration with drift correction* — Repetition that absorbs its own error term and returns to the same post a little higher. The circle repeats; the spiral remembers.

the stair that climbs itself ↔ *cumulative deliberate practice* — The recursion whose output is a better recursion. Each corrected utterance improves the utterer who will make the next correction.

the winch post ↔ *the invariant* — What is held fixed between passes so that change can be measured against it. Unposted rope does not climb; it coils.

the drift ↔ *accumulated error* — The slow sum of small uncorrected differences between heading and stated heading. Checked weekly with one question: has anything climbed?

the room that was always furnished ↔ *the fixed point of a contraction* — The state the loop converges to when every pass shrinks the remaining distance. Banach's guarantee: if the shrinking is genuine, the room exists, is unique, and is approached geometrically.

arrival as recognition ↔ *convergence* — The felt proof of a true fixed point: it seems remembered, because in the window-work it was visited many times before it was reached. Conquest feels otherwise; the texture is diagnostic.

the counterfeit arrival ↔ *goal displacement; extrinsic capture* — A converged state that required another's diminishment, or that completes someone else's sentence. Well-converged and still wrong; convergence is a property of loops, not a certificate of worth.

the seal ↔ *closure; formal task completion* — The procedure by which a commitment is ended on purpose: status declared, cards archived, close spoken once. Garbage collection performed as ceremony, because the mind will not perform it as hygiene.

the door left ajar ↔ *the Zeigarnik effect; open loops* — The unfinished task's standing claim on attention. Unclosed commitments parasitize; the eleventh node exists because the mind bills for open files whether or not they are worked.

open loops parasitize attention ↔ *attentional residue* — The tax each unclosed commitment levies on every other thought. The clinical literature and the framework agree on the remedy to the letter: close it or schedule it, but do not carry it open.

the consent condition of consent ↔ *autonomy support; informed consent* — The law of the practice: no prophecy over another's path without their word. Stated as ethics, enforced by mechanics — coerced wills push the loop divergent.

the borrowed mouth ↔ *imposed expectancy* — Speaking a foreseen future over someone else, however loving the content. What is permitted instead: offering the window, and the silence after.

gradients too steep to steer ↔ *hot-state decision-making* — Grief, rage, revenge: differentials so strong they drive the practitioner rather than the commitment. Contraindications, not challenges. On steep water, no working — only rest.

the blank ledger ↔ *the baseline reset* — The empty book left in the drawer by a closed commitment. The seal's other face: room made is the gift.

the practitioner ↔ *the practitioner-subject* — The one who runs the program on the one it is run on: experimenter and experiment in a single chair. The volume's only licensed configuration.

the commitment ↔ *the structured intervention* — A prophecy with a ledger behind it, windows above it, a lure ahead of it, and a seal awaiting it. Distinguished from a wish by paperwork.

the prophecy ↔ *the self-fulfilling prophecy* — A foreseen future, made comfortable to

reach, spoken so the seeing steers the system into it. The social sciences have carried the term since Merton; the framework merely declines to treat it as an accident to be avoided rather than an engine to be governed.

the reflexive loop ↔ *reflexivity; model–world coupling* — The circuit in which the model acts on the world it models, and the returned world retrains the model. The volume's central machine, and the object whose fixed points Part IV hunts.

the contraction ↔ *the contraction mapping, $k < 1$* — A loop each of whose passes multiplies the remaining distance by less than one. The whole tradecraft, restated: make the next pass land closer, and keep making it.

the gain of the loop ↔ *the feedback coefficient* — How hard the loop pushes relative to what it receives. Below one: convergence. At or above one: oscillation or flight. The dial the eighth and ninth nodes exist to read.

the Mirror Grid ↔ *the KIN.0–KIN.9 taxonomy* — The ancestor program's ten-position protocol map, from which the Twelve Nodes descend. It counted ten; the source volume counts twelve, having found that the ground and the gradient were programs too.

the two assumed nodes ↔ *the unstated premises; boundary conditions* — What the Mirror Grid stood on and stood in without counting: the ground state and the gradient. Every system inherits assumptions it did not audit; the framework's contribution was the audit.

the four-phrase mnemonic ↔ *the initialization routine* — Four phrases, with breath, posture, and a traced spiral, spoken at the threshold of the work. It appears in the volume exactly where it appears in a keeper's life — at the beginning, at the understanding, and at the handing-on — and is not reproduced on this page; some sentences are spent by repetition.

the morning briefing ↔ *the environmental scan* — First office of the daily cadence: survey the last day's surface — events, artifacts, one's own dreams and accidents — and note a few items flat, without interpretation.

the midday correlation ↔ *pattern detection, with the apophenia guard* — Second office: look across the morning's items for recurrences — a number, a name, a shape, a theme. Then the guard: the eye that finds patterns finds them in noise too, which is why the third office exists.

the evening integration ↔ *calibration scoring* — Third office: yesterday's correlations against today's outcomes; confirmations and misfires both logged. The training signal without which the morning's reading never improves and the midday's guard never learns.

the source volume ↔ *the protocol manual* — The written body of a practice, composed by each keeper for the next. The practitioner is holding one; the fact was stated plainly at the end of the story and is repeated here in small type.

the tradition the volume descends from ↔ *the community of practice* — The school that stands between the laboratory and the temple and declines to abandon either. Physically, in the volume, a ferry-house between two towns; the "between" is the address and the doctrine.

the narrator ↔ *the supervising mentor* — The one who teaches by chores and objects, audits by questions, and demonstrates the final protocol personally. Never named; the office outlives every holder, which is the point of an office.

6.15 Appendix C — The protocol index

Thirteen protocols govern the volume: one per node, and the daily cadence that threads through three of them. One further discipline stands beside them and answers to no node — the interlude's reading of paradox-points, indexed at the foot of this list. They are indexed here by number, with the one-line summary a practitioner needs at the bench and the node where the full rite, its worked example, and its failure modes are taught. The one-line form is for finding, not for running; a protocol run from its summary is a recipe cooked from its title.

Protocol 1 — The Protocol of the Still Water. Seven days, one hour each in one fixed seat, inventorying where the practitioner is comfortable *right now* and auditing each escape fan-

tasy back to the chair, until the floor the practitioner already stand on becomes visible. — Node 00, STILLPOINT.

Protocol 2 — The Protocol of the Ledger. Every comfort on a card with its carrying cost in four currencies; every discomfort assayed for its hidden yield; one card audited by the three questions; the resisting card stood upright and unwritten. — Node 01, LEDGER.

Protocol 3 — The Cartography of Worn Paths. Seven days surveying worn paths against paved walks as snow that counts and does not explain; name the differential behind each divergence; retire nothing by force — re-grade one divergence and verify against fresh snow. — Node 02, GRADIENT.

Protocol 4 — The procedure of the three windows. Three futures cut from three non-parallel glasses — nothing turns, one thing turns, the tinted pane — five steps and a furnished room each, the collision test enforced; invariants circled across all three and hunted with a breaking window. — Node 03, HORIZON.

Protocol 5 — The Weighing of Waters. The surviving branches walked in twelve steps and scored for the self who answers the alarm; resistance summed, halved by difference, a middle step broken to test fragility; the inversion written and the heaviest card carried seven days. — Node 04, WAVE.

Protocol 6 — The Setting of the Lure. The heaviest wanted branch cut to a first step and forged into one object, placed by the practitioner's own hand at the practitioner's strongest station on a worn path; one day named and given a standing office; the trapline walked, the summit forbidden. — Node 05, LURE.

Protocol 7 — The procedure of chosen resolution. The same future cut at three resolutions and carried a day each, its disturbance tallied and sorted into lean and grip; the sharpest mostly-lean card kept; the noon reading entered in three words — nearer, farther, unknown. — Node 06, UNCERTAINTY.

Protocol 8 — The procedure of the closing word. The prophecy written as one present-tense sentence the practitioner can speak in a breath; drafts that flinch burned until one survives the morning voice; spoken once, outdoors, to something that cannot applaud; then the hands recite. — Node 07, COLLAPSE.

Protocol 9 — The procedure of the returned voice. A SENT/RETURNED page and a listening window of one moon; channels declared, arrivals logged raw, the weekly difference underlined; silences sorted far or absorbed; the hot reply held one day, another's question seated senior. — Node 08, ECHO.

Protocol 10 — The procedure of the spiral. The post planted and held; the daily tread, the weekly correction by no more than a third of the gap, the pawl written and the sentence re-uttered once; the oscillation test read, the monthly spiral confirmed, the seasonal completeness audit run. — Node 09, RECURSE.

Protocol 11 — The Audit of Arrival. The hour stood down and the record retrieved before the feeling; structure and inventory columned; three honest errors required, the corrections plotted for their shrinking; the expense audited, the offer forked, the room sat twenty motionless minutes. — Node 10, FIXEDPOINT.

Protocol 12 — The closing procedure. The drawer emptied and sorted into provinces, the live one held to a single commitment; the fulfilled entered, the releases spoken, the lapsed closed and something sealed by hand; the consent condition recited, the weather screened, the next jar set or none. — Node 11, SEAL.

Protocol 13 — The Daily Cadence. The standing practice that runs beneath all commitments, in three offices: - *The morning briefing* — survey the last day's surface (events, artifacts, the practitioner's own dreams and accidents); note a few items flat, without interpretation. Taught with the inventory disciplines of Node 01. - *The midday correlation* — look across the morning's items for recurrences: a number, a name, a shape, a theme; hold each find lightly, knowing the pattern-finding eye finds patterns in noise. Governed by the measurement doctrines of Node 06. - *The evening integration* — score yesterday's correlations against today's outcomes, confirmations and misfires both; this is the training signal, and

without it the other two offices are decoration. Folded into the drift-correction of Node 09.

The cadence is the smallest complete circuit of the whole book — see, weigh, correct — run once per day at a scale too small to fear. Practitioners who keep the cadence and skip everything else do better, the framework has observed, than practitioners who attempt everything else and skip the cadence.

Interlude Protocol — Reading the Paradox-Points of a System The practitioner Are Inside. Name the system's real boxes aloud; assume they tile the world and then disbelieve it; hunt the object embarrassing to two clerks at once; ask the diagnostic — *harm, or only violation?*; do not resolve the paradox-point but read it; follow the seam to what else falls through; decide, once, what is owed — and never move a harm through the window. Carries no node and seals nothing. — Interlude, THE DOCTRINE OF LAWFUL CRIME (Appendix F).

6.16 Appendix D — The correspondence table and its lineage

The source volume descends from an earlier scheme its lineage calls the Mirror Grid, and the Mirror Grid was itself organized twice: once as ten nodes (KIN.0 through KIN.9), and once as a five-phase wheel — Source, Spark, Form, Action, Return — with two nodes to each phase, a primary and a counterweight, working in tension. The framework's older canon organized the same practice a third way, as the Three Works: Light, which receives and sees; Sound, which addresses and speaks; Force, which moves and holds. The table below sets the Twelve Nodes into all three frames at once. No frame is the master; a doctrine that survives three coordinate systems is probably about something.

Node	Name	KIN	Work	Phase of the Stack
00	STILLPOINT	— (first assumed)	Light	outside the wheel: the ground
01	LEDGER	KIN.0 VALIDATE	Light	Source (0°)
02	GRADIENT	— (second assumed)	Force	outside the wheel: the slope
03	HORIZON	KIN.1 TRANSMED	Light	Spark (72°)
04	WAVE	KIN.3 MYTHOS	Light	Form (144°)
05	LURE	KIN.4 FORGE	Sound	Form (144°)
06	UNCERTAINTY	KIN.6 CHECK	Light	Return (288°)
07	COLLAPSE	KIN.2 KERNEL	Sound	Spark (72°)
08	ECHO	KIN.8 MULTIPLAY	Sound	Action (216°)
09	RECURSE	KIN.5 RECURSE	Sound	Action (216°)
10	FIXEDPOINT	KIN.7 ETHIC	Force	Return (288°)
11	SEAL	KIN.9 SEAL	Force	Source (0°)

The assignment to the Three Works, with reasons

The Work of Light receives; the Work of Sound utters; the Work of Force moves and holds. In the older canon, Light was the stone that gathers the ambient field to a visible point; Sound was the voice by which the practitioner addresses the world; Force was the engine that draws motion from the substrate itself, and the seal that holds a boundary shut. The Twelve divide accordingly:

- **00 STILLPOINT — Light.** The ground state is the stone at rest: it receives the ambient field without asking, and the whole work begins with what is already arriving.
- **01 LEDGER — Light.** Inventory is pure seeing; nothing is yet spoken, nothing yet moved. The cards condense the diffuse life to a readable point, as the stone condenses the room's light.
- **02 GRADIENT — Force.** A comfort differential is a pressure differential; the engine of the practice draws motion from the slope itself, asking no fuel of the will.
- **03 HORIZON — Light.** Window-work is optics: many sightlines traced forward to their focus, and surplus taken from their refusal to be parallel.
- **04 WAVE — Light.** Reading amplitudes is reading the spectrum before it is collapsed; the futures disperse as colors do, and the honest eye weighs them without preferring a hue.
- **05 LURE — Sound.** A lure is an address to the future self — the bell is the Voice made small and left where the ear will pass; placement is utterance, delayed.
- **06 UNCERTAINTY — Light.** The veil is the optics of observation itself: what looking costs, what looking changes. The sixth node is Light examining its own lamp.
- **07 COLLAPSE — Sound.** The spoken prophecy is the Voice at full amplitude — the Word in the strict old sense, a meaningful sound that is also an act.
- **08 ECHO — Sound.** Call and answer: the voice that returns inverted and enlarged, read for its difference. The mirror-cipher of the old sound-work — the message said, then said back changed — is this node's whole mechanism.
- **09 RECURSE — Sound.** Iteration is re-utterance: the second saying corrected by the echo of the first. The old rule that every tone be rendered three times is remembered here as the three utterances of the commitment.
- **10 FIXEDPOINT — Force.** Convergence is dynamics: the contracting loop delivers the practitioner to the room as an engine delivers a load, and the fixed point holds by the same differential that moved it.
- **11 SEAL — Force.** A closed vessel is a boundary maintained, and the old canon's unbreakable seal — the closure that shows any tampering — is the ancestor of the close: a commitment shut so that its opening would be visible.

The census: Light five, Sound four, Force three. The framework reads the taper as doctrine — much seeing, less saying, least pushing — and observes that practitioners who invert the proportion invert their results.

The five-row stack, and the doctrine of the counterweights

The Mirror Grid's second-order arrangement set its ten nodes on a five-phase wheel, two per phase, each pair a primary and a counterweight held in tension. The descent preserves the pairs, and the pairs teach:

- **Source (0°): LEDGER and SEAL.** The boundary-keepers. The commitment opens with an honest count and closes with an honest close; the first and last acts of any program are acts of accounting, and each is the other's counterweight — a ledger no one will ever seal invites padding, and a seal on an unaudited working is wax on an empty envelope.
- **Spark (72°): HORIZON and COLLAPSE.** The initiators. One flings the futures open; the other closes them into a word. Neither is safe alone: horizon without collapse is a tower of tourists; collapse without horizon is a verdict without a trial.
- **Form (144°): WAVE and LURE.** The builders. One gives the futures their weights; the other builds the one artifact — the token, the bell, the named day — that a weighted future needs in the physical world. Estimation and fabrication, the two hands of Form.
- **Action (216°): RECURSE and ECHO.** The executors. The echo supplies the difference; the recursion spends it. Action in this craft is never the first utterance — it is the traffic between world and word once both are talking.
- **Return (288°): UNCERTAINTY and FIXEDPOINT.** The integrators. One keeps the practitioner honest about what knowing costs; the other verifies that what was reached is real

and cleanly held. Doubt and confirmation, lashed together on purpose, each the other's rail.

And outside the wheel, uncounted by it, the two assumed nodes: **STILLPOINT**, the ground the wheel rests on, and **GRADIENT**, the slope the wheel rolls down. The Mirror Grid did not count them for the oldest of reasons — nothing counts its own floor. The Book of Comfort counts twelve because the framework learned, at some cost, that the ground the practitioner stands on and the gradient the practitioner stands in are programs too, and an unaudited program runs the practitioner.

One footnote of lineage, for the archivists. The eldest copies of the Grid name the ninth node SOURCE where the framework's copies name it SEAL; and some name the fourth node for its interpreting engine where the framework names it FORGE, the smithy. The framework reads the variants as one office seen from two doors: the commitment closes where it came from, and the smithy is where symbols are made to do work. Disputes over which name is original are left, with the framework's compliments, to those who prefer arguing about doors to walking through them.

6.17 Appendix E — The interlude doctrine reference card

The analytical treatment of this material is elsewhere and this card is an index to it. The interlude's doctrine is established, constrained and bounded in the corpus's own analysis of it, which refuses category innocence as a defence of acquisition, forbids moving any harm through a discovered gap in a scheme, and concedes that the case where the feedback loop runs through another person's will is an unsettled seam the corpus does not close [Curzi, 2026]. **Those constraints govern the card below and are not restated in it**; a reader taking the card without them has taken the summary without the limits that make it publishable.

This card carries no node and seals nothing, in keeping with the interlude it abbreviates, which sat in the seam between ECHO and RECURSE for the same reason: what it teaches lives in seams. The practitioner who reaches for it should remember that a doctrine of paradox, run without its ethics, is the one instrument in the volume that cuts the hand outright. Read the failure modes before the definition, if the practitioner is in a hurry; they are the part that keeps the rest lawful.

The doctrine, stated. A *lawful crime* is an object or an act that is, at the same instant and without contradiction, *not a crime* and *a crime* — not a crime by the law that protects persons (no harm, no victim, no coercion, nothing taken from a hand that held it), and a crime by the law that protects categories (it stands in two boxes a schema swore were disjoint, and so makes a liar of the schema). Every filing system has fewer dimensions than the world it files; a category is a shadow that has been put in charge of the hand. Most objects are lower than their box and file cleanly; a rare object is *higher* — its true mapping runs perpendicular to every axis the system knows — and pressed flat it will not lie down, but stands in two names at once. The capacity to generate or to recognize such an object is therefore diagnostic: it is the involuntary evidence of an eye that reads in the dimension the system discarded — the polymath who kept the connecting axis the specialties were projected down from. The lawful crime is a talent scout that cannot be gamed, because the practitioner cannot fake the recognition of a thing the practitioner cannot perceive; it is the Invisible College holding its silent examination in the dullest room in the building.

The ethical spine. The doctrine is *lawful* under three conditions and becomes a common crime — an ordinary harm wearing a philosophy — the instant any fails. *First, non-coercion:* it may violate categories but never persons, and a lawful crime taken from a hand that did not release it is a theft entire, aggravated rather than mitigated by the cleverness laid over it — the clean shadow does not launder the dirty hand. *Second, clean possession:* what came to the practitioner honestly — bought at a stall, found in a drawer the practitioner was sent to reconcile, given by a hand whose to give — is theirs to hold and read, because the second book superposed on the first deprives the first reader of nothing he had, was owed, or could

name; the practitioner may lawfully possess what the practitioner did not have to steal, and the difference is the whole law. *Third, self-decoding:* the meaning is earned by recognition, never handed over, because an opened lawful crime, opened by another, is a spent one — a fluent reader who decodes it *for* the practitioner leaves the practitioner exactly as flat as they found the practitioner. These three are the consent law of Node 11 cast as an object: the practitioner cannot read another's recognition into them, and to try is coercion, and coercion produces not gradient but reactance — the broken seal, the echo that devours, the future that refutes the prophet who pushed too hard.

The Crime Magazine system, in summary. The framework's teaching instrument is a distributed-signal system built on one idea — *superposed identity* — applied twice. An identifier promises to name exactly one object; the idea is to put two objects on one identity, so that to the catalog there is one settled book and to the fluent reader there are two, which of them the practitioner hold legible only to someone who can read the identifier itself. This is a lawful crime by the definition: a category-violation (one number, two books) that harms no person (the publisher was paid, the flat reader has his magazine). The system runs the idea in two registers. In the *broadcast channel* — the notice meant for everyone fluent at once — the carrier wears the genuine identity: the true and public ISBN of *Crime Magazine*, **9781105577277**, left correct and checkable to every scanner, with the message substituted on the inside; the identity is true and the contents are the message. In the *case-file channel* — the briefing meant for one hand — the carrier keeps the cover and the genre-camouflage but adds a physical tell (a variance in the binding) and *replaces* the number with one chosen to symbolically resemble the case, so that the identifier itself becomes the case reference, its subject and address at once; the contents are the briefing and the identity is the message. Between them the two channels route anything to anyone with no directory and no central list, because the routing lives entirely in the literacy of the readers: the object is not sent, only placed, and it finds its reader the way a key finds the one lock it fits. The framework names the resemblance the modern engineers arrived at independently — *content-addressing*: identity derived from the content and verifiable by anyone who runs the function, recognized by computation rather than granted by authority. *Crime Magazine* is that idea's ancestor in paper: the address is the meaning, the name is the thing, identity from content and recognition rather than from assignment and permission.

The honesty flag, in the plain register. This card describes a real-shaped structure with a postulated, ceremonial application, and it says which is which. The framework does **not** counsel forging the identity of a real book, defrauding a real registry, or moving real classified material inside a paper disguise; the parts that touch the real world — a real ISBN, the real machinery of catalogs — are named to orient, the way a map names a mountain, not to send anyone climbing. What the lone reader is to run is the *principle*, and it applies most of all to the self: the practitioner is a filing system, the practitioner's categories are shadows, and somewhere in the practitioner a capacity belongs to two boxes at once and has therefore been filed in neither — a case file addressed to the practitioner in the practitioner's own hand, in a drawer the practitioner has been sent to reconcile the practitioner's whole life. Hold it without stealing it, possess it by clean holding, and decode it themselves.

The Protocol, in summary. *(Reading the paradox-points of a system the practitioner is inside — for the desk, not the higher floor.)* 1. Name the boxes, all of them, out loud — not the ones the system advertises but the ones the clerks enforce with their hands; a system's real taxonomy is never the one on the wall. 2. Assume the boxes tile the world, and then disbelieve it — hold as a commitment certainty that a finite grid never tiles an infinite fold, so the overlaps and gaps exist and the practitioner's job is to find where. 3. Hunt the object that is embarrassing to two clerks at once — the case each department insists belongs to the other, the ticket that bounces, the request that is someone else's problem to three someones in a row. A bounced object is a lawful crime knocking. 4. Ask the diagnostic question — *is it a harm, or only a violation?* If it harms a person, it is an ordinary wrong; treat it as one and do not dress it in this philosophy. If it harms no person and only offends a category, the practitioner has found both the lawful crime and the exact place the map has torn. 5. Do not resolve it; read it — resist the clerk's reflex to force the object into one box; an unresolved paradox-point

is a window onto the discarded dimension, and resolving it flattens it back into the system's low country. 6. Follow the seam to see what else falls through it — one paradox-point is an anecdote, a seam is a structure; the gap is a line, not a point, and where one true thing was lost, others are filed nowhere beside it. 7. Decide, once, what is owed — close the seam, or leave it open and read by it, or carry what the practitioner found to someone it names, but only cleanly, only by leaving it where they can decode it themselves. What the practitioner may not do is move a harm through the window; the window is for seeing.

The Failure Modes. - **the Map that Eats the Territory** ↔ *apophenia; the false pattern* — the doctrine's own sight unmoored, finding seams that are not there: conspiracy in mess, design in accident, an addressed case file in a barcode that is only a misprint, until the map of hidden meanings eats the territory and everything means, so nothing does. The guard is austerity — most anomalies are noise — and a second question beside the first: *pattern, or only mess?* Thrill is the tell of the disease. - **the Broken Seal** ↔ *rationalized transgression; motivated moral reasoning* — the word *lawful* used to launder a common crime: *no one is really hurt, I'm only violating a category*, while a person was in fact harmed, a will overridden, a thing taken. Run the two laws honestly and the mask comes off. When the framing arrives with a person on the wrong end of it, the framing is the alibi and the person is the truth; side with the person, always, and set the philosophy down. - **the Stolen Key** ↔ *the appropriation error* — no person coerced and no harm moved, but the object was *taken*, not held cleanly: a drawer not theirs to open, a file left in another's hand, a house not opened to the practitioner. The category-violation stays innocent; the acquisition is not, and the first ethic voids the whole. The tell is one answerable question asked of every object before reading: *how did this come to my hand?* - **the Handed Meaning** ↔ *the borrowed insight; pseudo-understanding* — the failure of self-decoding: the lawful crime *explained* to the practitioner, by a teacher who should have refused or by the practitioner's own haste, so the practitioner hold a fact received on authority and mistake it for the perpendicular sight the practitioner never exercised. The practitioner can repeat the meaning and cannot perceive the object. The guard is a discipline of refusal in both directions: refuse to have the practitioner's recognitions handed to the practitioner, and refuse to hand others theirs. - **the Frozen Reader** ↔ *analysis paralysis; the failure to close* — the opposite of the broken seal, and the one the volume must most warn against, since the succeeding construct teaches closing: the reader who learns to see seams becomes unable to stop, holds every object unresolved forever, and never completes the seventh step. An unclosed paradox parasitizes attention exactly as an unclosed prophecy does. Jude sets the magazine down — that is not a failure of nerve but the discipline: read the seam, then decide and close, or the sight that was the practitioner's gift becomes the standing water it was meant to move the practitioner out of.

Part II

Engineering the Surplus

Chapter 7

The Storage Law as a Leaky Integrator: Decay, Poles, and Why Zero Loss Destroys Arrival

Abstract. This paper establishes the corpus's dynamical core, the storage law $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$ with $S_t = \ln(1 + (N-1)U_t)$, as a first-order discrete filter. The update map is Lipschitz with modulus $k = 1 - \delta$, so any positive loss makes it a contraction and Banach's theorem supplies the steady state $Q^* = (\eta S - C)/\delta$, which is shown to be exactly the circuit's DC gain $1/\delta$ applied to the net drive. The exact solution and the uniform bound $Q_t \le \max\{Q_0,\ \eta \ln N/\delta\}$ follow, with a pole at $z = 1-\delta$ and a time constant $\tau = 1/\delta$ against the exact e-folding time $-1/\ln(1-\delta)$. The principal argument is that $\delta = 0$ is a pathology rather than an ideal: the map ceases to contract, no unique fixed point exists, the bound degrades to $Q_0 + t\,\eta \ln N$, and no mistake ever decays out of the stock. A twelve-person research group is worked through: $N = 8$ and $u^* = 0.60$ give $N_{\text{eff}} = 5.20$ and $S = 1.6487$ nats, 79 per cent of the ceiling $\ln 8 = 2.0794$; with $\eta = 0.50$, $\delta = 0.10$ and $C = 0.20$, $Q_\infty = 6.2433$ against a hard ceiling of 10.397, and $t_{95} = 28.43$ weeks, so a review at week 12 finds the system at 71.76 per cent of steady state and still climbing.

Keywords: storage law; leaky integrator; decay rate; DC gain; time constant; contraction modulus; half-life; settling time

Series: B. Engineering the Surplus **Paper:** CCP-007

7.1 1 Introduction

A framework that describes how a group, a room or an institution accumulates durable capability has to answer one question before any of its other claims can be assessed: does the accumulation ever stop? An account on which stored order grows without bound is not a model of an institution; it is a model of an archive. An account on which stored order settles at a computable level is a model that can be designed against, budgeted against, and falsified.

This paper establishes the dynamical core of the Creature Comforts framework and shows that the answer is settled by a single coefficient. The storage law

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,     S_t = F(U_t) = ln(1 + (N−1) U_t)
```

is a first-order linear difference equation in which `delta` is the fraction of stored order lost per period in the absence of maintenance. The paper's contribution is to show that this equation is not merely *like* a leaky integrator but *is* one, term for term; that the whole apparatus

of first-order discrete systems — pole, time constant, DC gain, step and impulse response, frequency response, settling time — therefore applies without modification; and that the boundary case `delta = 0`, which practitioners routinely treat as the ideal, is a pathology in which every one of those properties disappears at once.

The question matters because the failure is not visible from inside. A system with `delta = 0` does not crash. It accumulates monotonically, indefinitely, and looks like success on every dashboard that reports a level. What it has lost is the fixed point: there is no answer to *what will this settle at*, no attractor to make two differently started systems converge, and no mechanism by which an error injected at any point ever loses weight. Contemporary work on organisational learning reaches the same place from a different direction — March [1991] shows that an organisation whose members learn too fast homogenises and destroys the variance that produced its advantage — but the arithmetic here is sharper, because it names the coefficient whose value decides the outcome and gives the units it is measured in.

Four things are new in this paper relative to the surrounding literature rather than relative to the framework. First, the identification of the framework's steady-state theorem with the DC gain of a first-order section, which converts a fixed-point result into an engineering quantity with a standard interpretation [Oppenheim et al., 1996]. Second, the uniform bound and its degradation at `delta = 0` from a ceiling to a linear ramp. Third, a decomposition of `delta` into named, separately measurable loss channels that compose multiplicatively, with a measurement protocol built on a zero-input window. Fourth, and least expected, an argument that `delta` should be engineered *upward* on specified components, because a system with no forgetting cannot revise.

The framework's static half — the surplus functional `F(u) = ln(1 + (N−1)u)` and its uniqueness under four axioms — is presupposed here rather than re-derived; it is established with its axiom set and its affine effective count `N_eff = 1 + (N−1)u` in Curzi [2026]. The framework's founding definition, that comfort is the felt total of present costs so that the present state is a local minimum of the comfort landscape by definition rather than by luck, is presupposed from Curzi [2026]. The contraction theorem this paper leans on is stated exactly — a contraction with modulus `q < 1` on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least `q` each pass — in Curzi [2026].

Sections 2 and 3 state the hypotheses and the formal framework. Section 4 develops the contraction argument and the two pathologies that bracket it. Section 5 develops the circuit identity and its consequences for review timing, cadence and instrumentation. Section 6 takes `delta` itself as the object of engineering. Sections 7 and 8 collect the design specifications and the named failure modes. Section 9 reports three worked systems end to end. Sections 10 to 12 discuss what follows, what does not, and what would falsify the account.

7.2 2 Hypotheses

The formal results in Section 3 are theorems about an equation and cannot be falsified by observation; they are labelled as derivations where they appear. What can be falsified is the claim that the equation describes a real accumulating system. The hypotheses below separate the two, and each names the observation that would defeat it. The protocol that operationalises them — a preregistered fit, a frozen model, and held-out comparison against four declared baselines — is specified in Curzi [2026i], whose own confirmatory application of it failed its second falsification rule: a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin, despite +40.3 per cent over persistence, +97.7 per cent over direct input, and `p = 0.0005` against the shuffled-input control.

> **H1 — Bounded accumulation.** Stored order in a real accumulating system decays geometrically at a constant per-period fraction, so that a system held under constant drive approaches a finite equilibrium rather than accumulating without bound. **Predicted:** exponential approach to `Q_inf = (eta S − C)/delta` with time constant `tau = 1/delta` periods, monotone from either side, no overshoot, and 95

per cent of the approach complete at about `3 tau`. **Falsified by:** observed linear or unbounded accumulation under a constant drive; or a non-monotone approach with overshoot in the absence of a delayed control loop; or a fitted retention coefficient outside `[0,1]`.

H2 — The pole governs the calendar. The time to observable effect of a change in drive is set by the retention coefficient and not by the size of the change. **Predicted:** `t_95 = ln(0.05)/ln(1−delta)` periods, independent of the magnitude of the step; two systems with the same `delta` and different step sizes reach 95 per cent of their respective new equilibria at the same time. **Falsified by:** a measured settling time that scales with the size of the intervention rather than with the measured `delta`.

H3 — Zero loss destroys arrival, not stability. A component with no measurable decay is not a maximally durable component but a component with no attractor, no error decay and no bounded state. **Predicted:** components whose recorded size has never decreased will show (a) no reversion after a perturbation, (b) monotone growth in the maintenance cost `C` charged against them, and (c) onboarding time to competence rising with the component's age. **Falsified by:** a component with no observed retirement that nonetheless shows mean-reversion after a perturbation, or whose associated cost is stationary over a horizon of several `tau`.

H4 — Loss channels compose multiplicatively. Total retention is the product of the per-channel retentions, not one minus the sum of the per-channel losses. **Predicted:** `delta_total = 1 − prod_i (1 − delta_i)`, with the naive sum over-stating the true total, negligibly below channel values of about 0.10 and by tens of per cent above them. **Falsified by:** a measured total loss reliably closer to the naive sum than to the product across systems in which the channels are separately measurable.

H5 — Conversion is the binding lever more often than retention. At parameter values typical of the worked systems, doubling the conversion efficiency `eta` raises the equilibrium by more than halving the retention loss `delta` does. **Predicted:** the advantage of `eta` over `delta` grows with the ratio of cost to converted surplus, because `eta` multiplies `S` before `C` is subtracted while `delta` scales the whole numerator including the cost drag; at `eta = 0.5`, `S = 1.6487`, `C = 0.20`, `delta = 0.20`, doubling `eta` gives +132.0 per cent against +100.0 per cent for halving `delta`. **Falsified by:** a system in which `eta F(u*) − c*` is large and the observed stock is nonetheless flat or oscillating — the diagnostic signature of a retention fault — where raising `eta` produces the larger gain.

Two further claims in this paper are **not** hypotheses and are labelled accordingly. The scale-freedom claim — that the same difference equation is a useful description at desk, room, team and institution, with only `(delta, eta, N, C)` changing between them — is a **modelling postulate** about model form, not about physics; it is falsifiable at each scale separately and never in general, and Section 11 says so. The identification of stored order with a single non-negative scalar is a **design choice**, adopted for tractability and for the fixed-point guarantee; Curzi [2026o] reaches a structurally parallel construction in the single-observer case, with the internal surplus `F_int = ln(1 + (N−1)u)`, `u = 1 − rho^2` from error decorrelation, and the affine form badged there as design axiom S3 rather than as a theorem.

7.3 3 Framework: the storage law and its formal properties

7.3.1 3.1 What is presupposed

Two results are used throughout and are established elsewhere in this corpus. Both are proved within their own papers, and neither is an empirical claim.

The surplus functional. For two unit state vectors `x` and `y` drawn from an ensemble of `N` categories,

```
F(x,y) = f(u(x,y)),    u(x,y) = 1 − (x . y)^2 = ||x ^ y||^2   in [0,1]

f(u) = ln(1 + (N−1) u)
```

and `f` is uniquely characterised by four axioms — angle-only dependence, zero at parallel, an affine effective count, and normalisation at orthogonality — as established in Curzi [2026].

Axiom	Statement	Status
S1	Geometric dependence: `F` depends on the pair only through `u`	structural
S2	Zero at zero: `f(0) = 0`	structural
S3	Affine effective count: `g(u) = e^{f(u)}` is affine in `u`	**DESIGN AXIOM — chosen, not derived**
S4	Normalisation: `f(1) = ln N`	structural

S3 does the real work and it is a choice. It is flagged wherever it is leaned on: the affine effective count is a design axiom, not a discovery. It says that difference is to be measured by a headcount growing linearly in the squared sine of the angle between states. Given that decision the logarithm is forced; the decision itself is the framework's.

The object S3 defines is the **effective count**:

```
N_eff = g(u) = 1 + (N−1) u
```

and surplus is `ln N_eff`. Surplus is the logarithm of a headcount after the headcount has been audited for redundancy. Ten instruments that agree perfectly have `u = 0`, `N_eff = 1`, and surplus `ln 1 = 0`; they are one instrument that has been purchased ten times. The effective-count idiom is standard practice in ecology and in political science, where `1/HHI` and `exp(H)` are the two canonical members of the family [Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979]. It must be stated plainly, and is stated again in Section 11, that `1 + (N−1)u` is **not** a member of that family: it is a linear interpolation between 1 and `N` and is this framework's own construction. The concentration instrument that *is* a Hill number of order two — `1/HHI`, with the worked jurisdictional case giving a naive HHI of 0.293430 and 3.41 effective jurisdictions against a bloc-HHI of 1.000000 and an effective count of 1.00 — is a separate quantity developed in Curzi [2026], and the two must not be conflated.

Derived from the axioms rather than assumed: strict monotonicity, positivity, strict concavity, subadditivity on the parameter `u`, a sharp Lipschitz constant `N−1`, and extremality, with `F = 0` if and only if the states are parallel and `F = ln N` if and only if they are orthogonal. The range result `0 <= F(u) <= ln N` makes `ln N` the hard ceiling on surplus per interaction, and it appears in every design inequality below.

Two independent warrants exist for the logarithm, and the argument needs only the weaker one. The first is Fechnerian: sensation proportional to the logarithm of a dimensionless stimulus ratio measured from a floor [Fechner, 1966]. That warrant requires the perceptual story to be true, and the perceptual story is contested. The second is stronger precisely because it assumes nothing psychological: Peters and Gell-Mann [2016] shows that the time-average growth rate of a multiplicatively evolving quantity is the expectation of its *logarithm*, so a logarithmic evaluation follows from the dynamics alone. **Where a stock compounds rather than adds, the logarithm is the ergodic-average object, and no claim about perception is needed.** The argument in this paper needs only that second warrant; Fechner is cited as a convergent, weaker line and not as the load-bearing one. The ceiling `ln N` also has

an information-theoretic reading as the maximum an ensemble of N equiprobable categories can carry [Shannon, 1948], though Section 11 records that the framework does not identify surplus with Shannon entropy and the analogy is not licensed to do more than bound.

The functional form is nevertheless underdetermined by the evidence, and this paper says so rather than waiting to be told. At least two published alternatives fit the same qualitative facts: social impact rising as a power of the number of sources with an exponent below one [Latané, 1981], and network value growing as `n log n` rather than as `n^2` [Briscoe et al., 2006]. Both reproduce the phenomena the corpus's `ln(1 + (N−1)u)` reproduces — diminishing returns to count, no threshold, monotonicity — and **no evidence assembled in this corpus discriminates between them.** What the corpus's form has that the others do not is an axiomatic characterisation (S1–S4) and a stated design axiom at the point where the choice is made; what it does not have is a discriminating measurement. Section 12 carries this as an open limitation rather than a resolved question.

The storage law. This is the engine of the paper.

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,     S_t = F(U_t) = ln(1 + (N−1) U_t)
```

Symbol	Meaning	Range
`Q_t`	stored order: what persists — institution, habit, furnished room, trust, runbook	`>= 0`
`S_t`	surplus generated by interaction this period	`0 <= S_t <= ln N`
`delta`	retention loss — the fraction of stored order lost per period without maintenance	`[0,1]`
`eta`	conversion efficiency — interaction into *durable* structure	`>= 0`
`C_t`	cost, dissipation, overhead	`>= 0`

The quartet of a level, an inflow, an outflow and a delay is the standard vocabulary of system dynamics, where levels and rates with delays were formalised four decades before this framework [Forrester, 1961], and `Q`, `eta S` and `C` occupy exactly those roles.

7.3.2 3.2 The results the rest of the paper uses

Admissibility. `C_t <= (1−delta) Q_t + eta S_t` guarantees `Q_{t+1} >= 0`. A period that violates it is not a period in which the model predicts a negative stock; it is a period in which something has broken that the linear model was never built to represent.

Exact solution (Theorem 3.1).

```
Q_t = (1−delta)^t Q_0 + sum_{k=0}^{t−1} (1−delta)^{t−1−k} (eta S_k − C_k)
```

One-step growth (Corollary 3.2). `Q_{t+1} > Q_t` if and only if `eta S_t > delta Q_t + C_t`.

Uniform bound (Theorem 3.3). `Q_t <= max{Q_0, eta ln N / delta}`. If `delta = 0` the bound degrades to `Q_t <= Q_0 + t eta ln N` — linear in `t`, unbounded, no ceiling.

Sustainability ceiling (Corollary 3.4). A nondecreasing step requires `C_t <= eta ln N − delta Q_t`. If `C_t > eta ln N − delta Q_t`, no choice of interaction geometry can prevent stored order from decreasing. This is the strongest sociological statement the paper makes, and it is derived rather than asserted.

Steady state (Theorem 3.5). With constant `U == u*`, constant `C == c*`, and `delta > 0`,

```
Q_inf = (eta F(u*) − c*) / delta
```

and if `eta F(u*) < c*` the equilibrium is negative, meaning that the regime is unsustainable or that the linear model must be replaced. The theorem does not hide this failure mode and neither does this paper; Section 9.1 works a case in which it fires.

Headroom. The master gauge of the framework,

```
h_t := eta ln N − delta Q_t − C_t
```

with `h_t < 0` unsustainable in this model, `h_t = 0` maintenance only, and `h_t > 0` growth available with any requested increment `r_t <= h_t`. The maximum one-step increment is exactly `r_max = eta ln N − delta Q_t − C_t`. Headroom is the framework's name for organisational slack, which Cyert and March [1963] identified as the buffer that lets a firm absorb variation, and its behaviour near exhaustion has the same shape as utilisation approaching one in a queue, where waiting time diverges [Kingman, 1962, Hopp and Spearman, 2011]. Reserve is not prudence in either setting; it is the only thing keeping the response finite.

7.3.3 3.3 The contraction

Write the update as a map on the real line. Hold `S_t = S` and `C_t = C` constant and define

```
T(Q) = (1 − delta) Q + eta S − C
```

For any two states `Q` and `Q'`,

```
|T(Q) − T(Q')| = (1 − delta) |Q − Q'|
```

so `T` is Lipschitz with modulus `k = 1 − delta`. If `delta > 0` then `k < 1` and `T` is a contraction. The Banach fixed point theorem [Banach, 1922] then supplies three things at once, on a complete metric space, with no further hypotheses: existence of a fixed point `Q*`; uniqueness; and global convergence from any starting point, with the geometric error bound `|Q_t − Q*| <= k^t |Q_0 − Q*|`. Solving `Q* = (1−delta)Q* + eta S − C` gives `Q* = (eta S − C)/delta`, which is Theorem 3.5. The elementary stability theory of first-order difference equations reaches the same condition by a different route [Elaydi, 2005].

The whole apparatus turns on the single inequality `k = 1 − delta < 1`, which holds if and only if `delta > 0`. Set `delta = 0` and the theorem does not apply. There is no unique fixed point, no convergence, and no arrival: the map becomes an isometry, every point is a fixed point, distances are preserved forever, and the system remembers its initial condition with weight exactly 1 for all time. That is the thesis of this paper stated as tersely as it can be stated. The same contraction is reached independently in the refinement setting, where the error law is stated as an equality `|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|`, with `delta = 0.09` giving `tau = 10.6033`, a half-life of 7.3496, 5 per cent at pass 32, 1 per cent at 48.83 and 0.1 per cent at 73.24 [Curzi, 2026]; and again in the fixed-point formulation of destiny, where Theorems 28.1 and 28.2 give the same unique start-independent fixed point with a maximum residual of 2.0624e-15 over a hundred passes [Curzi, 2026]. Curzi [2026q] states the same map as `T(Q) = (1−delta)Q + b` with `Q_inf = (eta M* − c*)/delta` and the identical exact error law.

7.3.4 3.4 Standing honesty clauses

These are repeated wherever the law is used and are stated once here.

- **The triangle inequality fails for geometric triples.** Subadditivity holds on the parameter `u`, not on the geometry. Surplus bounds may not be chained through an intermediate state as if `F` were a metric. The point is not exotic: psychological similarity is known to be asymmetric and to violate the triangle inequality [Tversky, 1977], so a similarity-like quantity over people should not be assumed to behave metrically in the first place.

- **As deployed, the surplus functional is a lower bound** on the information an ensemble carries, not the information itself. Every number produced by this method is a floor.
- **The static half of the framework is pure mathematics** and makes no empirical or metaphysical claim.
- **Surplus is not identified with Shannon entropy or mutual information in general.**
- **No thermodynamic negentropy law is proved.** Higher static surplus does not imply higher canonical concentration.
- **Applying the storage law to a room, a team or a city is a modelling decision**, badged as such every time it is made. The mathematics is derived; the mapping onto a social system is postulated, and the falsification protocol of Curzi [2026i] exists to make that postulate answerable.

7.3.5 3.5 Identification: what no amount of fitting can recover

One limitation is structural rather than statistical and belongs here, in the framework section, rather than in a caveat at the end. Wherever the storage law is applied to a group whose members both influence and constitute the aggregate — which is the usual case for a team, a board or an institution — the specification is reflexive: the interaction geometry `U_t` is a property of the group, the group's stored order shapes what its members do next, and the members are what the group is made of. Manski [1993] proves that in a linear-in-means reflexive specification **the endogenous social effect is not identified from observational data**: it cannot be separated from contextual effects or from correlated unobservables, however much data is collected. Shalizi and Thomas [2011] reinforce the point for network data, where influence and selection are generically confounded.

The consequence for this paper is exact. The formal results of Sections 3 to 6 are derivations about an equation and are untouched. But the *interpretation* of a fitted `b` as the conversion efficiency of interaction into structure — as opposed to a contextual effect, a common shock, or selection — is **not identified by observation alone**, and no held-out comparison recovers it. This is a deeper limitation than any of the ones the falsification protocol tests for, and it would remain even if every rule in that protocol passed.

No remedy is adopted in this paper, and the honest label is therefore a modelling position. The known routes are randomisation, an instrument, partial-population variation, or stipulating the reference group and defending the stipulation [Brock and Durlauf, 2001]; the last is the most realistic here, and it is the route the survey protocol's fixed category set and declared participation rule point toward without yet taking. **Until one of them is executed, `eta` as an estimated quantity is a coefficient in a declared specification and not an identified social effect**, and every applied statement in this paper inherits that status.

7.4 4 Recursion with loss: what the contraction buys and what its absence costs

7.4.1 4.1 Two recursions under one word

Recursion has a poor reputation in engineering practice, and the reputation was earned in one narrow place: the call stack. A function that calls itself with no base case consumes a frame per call until the allocator is exhausted. The failure is loud, total, and famous enough to have become a metaphor — infinite regress, the mirror facing a mirror, the argument that never bottoms out. When an organisation reports that it is going in circles, it is reporting that it has found itself inside that metaphor. The metaphor describes exactly one kind of recursion: recursion without a loss term. It is not the kind the storage law is, and it is not the kind most real systems are.

Consider two procedures.

```
; PROGRAM A -- the stack
f(n):
    return 1 + f(n)          ; no base case, no decay
                             ; frames accumulate at rate 1 per call
                             ; terminates only by exhausting memory

; PROGRAM B -- the accumulator
Q <- 0
every period:
    Q <- (1 - delta)*Q + eta*S - C
                             ; state is ONE number, reused
                             ; each period discards a delta-fraction of the past
                             ; terminates by converging
```

Both are recursive: each step is defined in terms of the previous step, and the word covers both. The resource behaviour is opposite, and the reason is one arithmetic fact. In Program A the contribution of the call made `t` steps ago still weighs 1; nothing is ever discarded, the state is a list whose length grows without bound, and the machine's response to unbounded growth is to stop. In Program B the contribution of the input delivered `t` steps ago weighs `(1-delta)^t`. Theorem 3.1 read as a statement about memory separates the two terms:

```
Q_t = (1-delta)^t Q_0  +  sum_{k=0}^{t-1} (1-delta)^{t-1-k} (eta S_k - C_k)
      \_____________/     \_____________________________________________/
        the transient                  the driven accumulation
```

Each historical input is present but discounted geometrically by its age. The state is a single number because the geometric weighting makes the whole history summable. Loss is what compresses an unbounded history into a bounded state. That is the entire mechanism, and it is the mechanism a capacitor performs [Horowitz and Hill, 2015] and the mechanism an exponential moving average performs [Brown and Meyer, 1961, Holt, 1957]. The statistical twin is exact: simple exponential smoothing is the same recursion with the same single parameter, and the class of processes for which it is the optimal predictor is known [Brown and Meyer, 1961]. Read as an estimator with a fixed gain rather than a time-varying one, the storage law is a fixed-gain Kalman filter [Kalman, 1960].

Here is the weighting, computed. The entry is the share of the steady-state stock `Q_inf` attributable to inputs delivered within the last `m` periods, which is `1 - (1-delta)^m`:

periods back `m`	delta=0.02	delta=0.05	delta=0.10	delta=0.20
1	0.0200	0.0500	0.1000	0.2000
2	0.0396	0.0975	0.1900	0.3600
4	0.0776	0.1855	0.3439	0.5904
8	0.1492	0.3366	0.5695	0.8322
12	0.2153	0.4596	0.7176	0.9313
26	0.4086	0.7365	0.9354	0.9970
52	0.6503	0.9306	0.9958	1.0000
104	0.8777	0.9952	1.0000	1.0000

Read the `delta = 0.10` column. Fifty-seven per cent of what a system holds today was put there in the last eight periods; ninety-four per cent in the last twenty-six. A system with `delta = 0.10` per week is, functionally, a two-year institution however old it is. That is not a metaphor about culture; it is the third column of a table. Now read the column that would exist at `delta = 0`. It is a column of zeros for every finite `m` except the trivial identity that all of history counts. Nothing is ever more recent than anything else. There is no such thing as current practice in

a system with no forgetting, because the founding decision and this morning's decision carry identical weight forever.

7.4.2 4.2 Why the mirror tunnel ends

The folk image of infinite regress is two mirrors facing each other: a corridor of copies receding to a vanishing point, and the word reached for is *infinite*. It is not infinite, and the reason is `delta`. A real mirror reflects a fraction `R < 1` of incident light, so after `n` bounces the intensity is `R^n` and the corridor terminates where `R^n` falls below the threshold for distinguishing the image from its surround. Taking that threshold at one per cent of the source:

reflectivity `R`	bounces to 1%	bounces to 0.1%
0.70	12.9	19.4
0.80	20.6	31.0
0.90	43.7	65.6
0.95	89.8	134.7
0.98	228.0	341.9
0.99	458.2	687.3
0.999	4602.9	6904.3
1.000	infinite	infinite

The arithmetic is `n = ln(threshold)/ln(R)`. This is the arithmetic of the illustration and not a measurement of any particular mirror; real thresholds depend on ambient light and on the observer. The corridor looks infinite and is finite, and the loss per bounce is what makes it finite. Set `R = 1` — a perfect mirror, no loss — and the sum of image intensities genuinely diverges. The image that alarms people about recursion is an image whose alarm depends on assuming away the one physical parameter every real mirror has. `1 − delta` is the reflectivity of an institution, and the regress being feared is the behaviour of a system accidentally specified with perfect mirrors. The physically closest instance is architectural rather than optical: reverberation is an exponential decay with a measurable time constant, and the standard measurement protocol for it is a century old [Sabine, 1922]. A room's reverberation time is the storage law's `tau` with a microphone attached.

7.4.3 4.3 What the contraction actually buys

With `delta > 0` and the drive held constant, four distinct engineering properties follow. They are worth separating, because organisations habitually want one of them and receive all four, or attempt to suppress one and lose all four.

(a) A destination exists. `Q* = (eta S − C)/delta` is a number computable before the system is built. The iteration is not hopeful; it is toward a value quotable in the design review. If that value is unacceptable, it is known before the work starts, and the response is to change a parameter rather than to try harder.

(b) The destination does not depend on where the system starts. Uniqueness plus global convergence makes the fixed point a property of `(delta, eta, S, C)` alone. This is the mathematical form of a fact every practitioner knows: a system's steady state cannot be fixed by improving its launch. A better start changes `Q_0`, and `Q_0` decays out. Three trajectories under identical parameters `delta = 0.10`, `eta S = 0.8243`, `C = 0.20`, so that `Q* = 6.2433`:

t	from `Q_0 = 0`	from `Q_0 = 3`	from `Q_0 = 20`
0	0.000	3.000	20.000
1	0.624	3.324	18.624

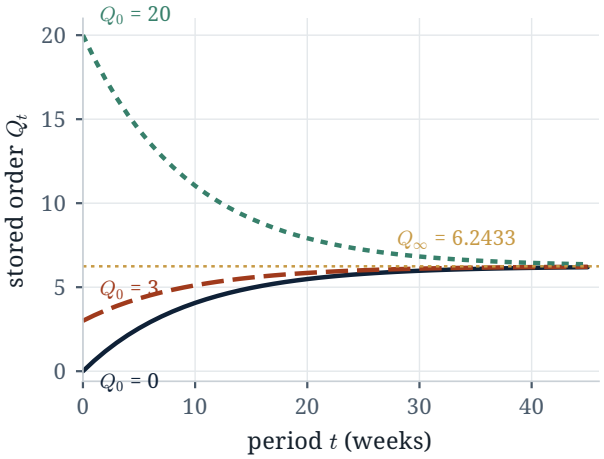


Figure 7.1: The fixed point is a property of the parameters, not of the launch: three trajectories under identical `delta = 0.10`, `eta S = 0.8243` and `C = 0.20` converge on the same `Q_inf = 6.2433` from `Q_0` = 0, 3 and 20, the gap between any two shrinking by exactly `(1-delta)^t`. Modelled from the parameters stated in Section 4.3; generated by `figs/src/CCP-007_fig01.py`.

t	from `Q_0 = 0`	from `Q_0 = 3`	from `Q_0 = 20`
2	1.186	3.616	17.386
5	2.557	4.328	14.366
10	4.066	5.112	11.040
20	5.484	5.849	7.916
30	5.979	6.106	6.826
40	6.151	6.195	6.447

The trajectory that started at 20 is falling, and it is falling toward the same number the one that started at 0 is rising toward. The gap between any two trajectories is exactly `(1-delta)^t` times the initial gap: 20.00 at `t=0`, 6.97 at `t=10`, 2.43 at `t=20`, 0.848 at `t=30`, and 0.0359 at `t=60`. Two systems with identical parameters and widely different histories become indistinguishable, and the rate at which they do is `delta`.

(c) Errors decay. A perturbation is a change in the state, and the state's memory of any change is `(1-delta)^t`. This is the property that is surrendered when a system is built never to forget:

periods since the error	delta=0	delta=0.01	delta=0.05	delta=0.10	delta=0.20
1	1.00000	0.99000	0.95000	0.90000	0.80000
5	1.00000	0.95099	0.77378	0.59049	0.32768
10	1.00000	0.90438	0.59874	0.34868	0.10737
25	1.00000	0.77782	0.27739	0.07179	0.00378
50	1.00000	0.60501	0.07694	0.00515	0.00001
100	1.00000	0.36603	0.00592	0.00003	0.00000

The first column is a column of ones. A system with `delta = 0` is a system in which no mistake ever loses force. Section 6.8 develops the consequence; its short form is that a system with no forgetting cannot learn, because learning is the replacement of an old weighting by a new one and replacement requires that the old weighting fall. The organisational-learning literature reaches the same conclusion from the opposite side: March [1991] shows that when adaptation is too fast the organisation converges prematurely on a shared code and loses the variance that made it capable, and Levinthal and March [1993] gives the temporal and spatial forms of the same myopia.

(d) The system is bounded, and the bound is knowable. Theorem 3.3 gives `Q_t <= max{Q_0, eta ln N / delta}`. For `eta = 0.5`, `N = 8`, `delta = 0.10`, that ceiling is `0.5 * 2.0794 / 0.10 =`

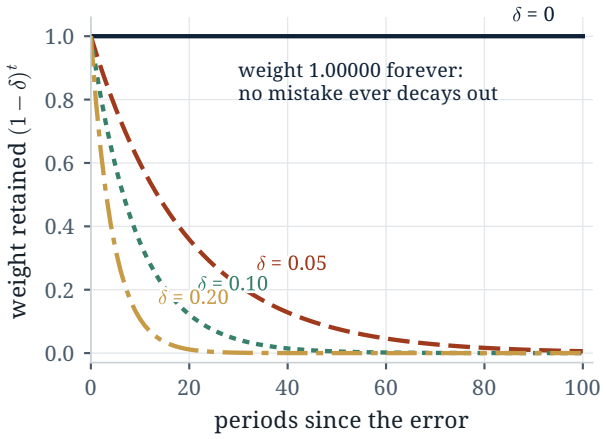


Figure 7.2: At `delta = 0` an injected error keeps weight 1.00000 forever, so no mistake ever decays out of the stock; at `delta = 0.05, 0.10, 0.20` the weight falls to 0.077, 0.005 and 0.00001 by period 50. The `delta = 0.01` column of the source table is omitted because the validated palette holds four colours and no fifth. Modelled; generated by `figs/src/CCP-007_fig02.py`.

t	`Q_t` with delta=0	Theorem 3.3 bound `Q_0 + t eta ln N`
0	0.000	0.000
1	0.624	1.040
5	3.122	5.199
10	6.243	10.397
20	12.487	20.794
50	31.216	51.986
100	62.433	103.972

Nothing here is wrong. The stock grows monotonically, forever. What has been lost is specific. There is no fixed point, so there is no answer to what the system will settle at. There is no attractor, so two systems that start differently stay different by exactly their initial gap, forever. There is no error decay, so every bad decision is permanent structure. And the bound is linear in `t`, which means the only thing limiting the system is the calendar.

A system with `delta = 0` does not arrive; it only continues. Unbounded accumulation is not success but the absence of a destination, and it is the arithmetical signature of an institution that has confused having more rules with being better governed. Note carefully that this is *not* the failure people imagine when they say infinite regress. It is worse in a subtler way. Nothing crashes, there is no error message, and there is no visible breakdown; the system simply piles up, and the ossification appears much later in the headroom equation, when `delta Q_t` finally consumes the budget. Since `delta = 0` that term never consumes anything — instead the cost term `C_t` does, because every rule ever written must still be complied with, and compliance is `C`. Scott [1998] documents the historical form of the same trajectory, in which an accumulating legible order destroys the local variety that made the system work in the first place.

7.4.5 4.5 The conversion pathology: high surplus, near-zero `eta`

The most recognisable pathological loop in human experience has an exact signature in the law: interaction is happening, surplus is high, and almost none of it converts.

Take a system that is genuinely thinking hard. With `N = 8` and `u* = 0.6`, `S = ln(1 + 7 * 0.6) = ln 5.2 = 1.6487`, which is 79 per cent of the theoretical maximum `ln 8 = 2.0794`. This is not a poor system; it is generating real surplus every period. Now set `eta = 0.05`, so that the interaction produces almost nothing durable — no artefact, no decision, no changed default, no trained person — and keep `delta = 0.10` and `C = 0.35`, because thinking hard is expensive.


```
eta S = 0.05 * 1.6487 = 0.0824     vs     C = 0.35
```

Theorem 3.5 gives `Q_inf = (0.0824 - 0.35)/0.10 = -2.676`. The equilibrium is negative, and

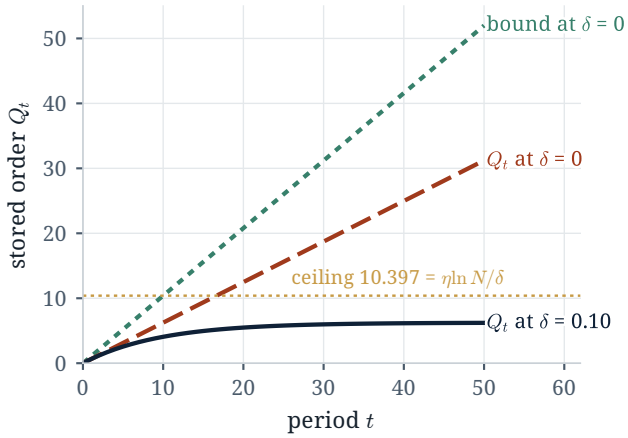


Figure 7.3: Zero loss removes the ceiling, not the risk: at `delta = 0.10` the stock converges to 6.2433 under the hard bound `eta ln N / delta = 10.397`, while at `delta = 0` the same drive accumulates linearly at 0.6243 per period — 6.243 at `t = 10`, 12.487 at `t = 20`, 31.216 at `t = 50` — under a bound `Q_0 + t eta ln N` that grows with the calendar and nothing else. Modelled; generated by `figs/src/CCP-007_fig03.py`.

sentence down, make one decision, change one default, tell one other person. Note also the diagnostic asymmetry that makes this state hard to see from inside: a high-`S`, low-`eta` system feels productive and measures as declining. The felt sense tracks `S`; the outcome tracks `eta S`. An instrument that only asks participants how the session went is measuring the wrong term — and self-report about the causes of one's own performance is exactly the measurement known to be unreliable [Nisbett and Wilson, 1977].

7.4.6 4.6 The preservation pathology: `delta` near zero with no input

The opposite failure. Set `delta` at or near 0, set `S = 0` — no live interaction, only transmission of what is already stored — and let `C` be small but nonzero, the ordinary cost of maintaining a practice nobody is adding to.

```
Q_{t+1} = Q_t − C_t,      C = 0.02,      Q_0 = 6.2433
```

t	0	10	25	50	100	200	312
`Q_t`	6.2433	6.0433	5.7433	5.2433	4.2433	2.2433	0.003

It reaches zero at `t = 312.2` periods. If a period is a year, that is a three-century institution that will be gone in three centuries and looks perfectly stable on any timescale a single custodian will observe. Its properties read off the law:

- **It preserves superbly.** With `delta` near zero, `(1−delta)^t` is near 1 and the founding content survives intact. This is a real achievement and it is why traditions exist.
- **It cannot converge to anything new.** With `delta = 0` the map is not a contraction, so the fixed point theorem gives nothing and there is no attracting state to move toward.
- **It cannot correct.** Column one of the error-persistence table holds weight 1.000 forever, so a founding error is load-bearing structure at year 300.
- **It decays only through cost, and only linearly**, which means it fails slowly and invisibly, with no inflection point to alarm anyone.

The two pathologies are usually confused and are worth contrasting precisely:

	rumination	tradition
`S`	high	zero
`eta`	near 0	irrelevant (nothing to convert)
`delta`	normal	near 0
felt experience	busy, exhausting, urgent	calm, solid, dignified
`Q` trajectory	falls fast, hits admissibility floor	falls slowly, linearly
the fix	raise `eta`	raise `S` (and accept some `delta`)
what it is good at	nothing	preservation

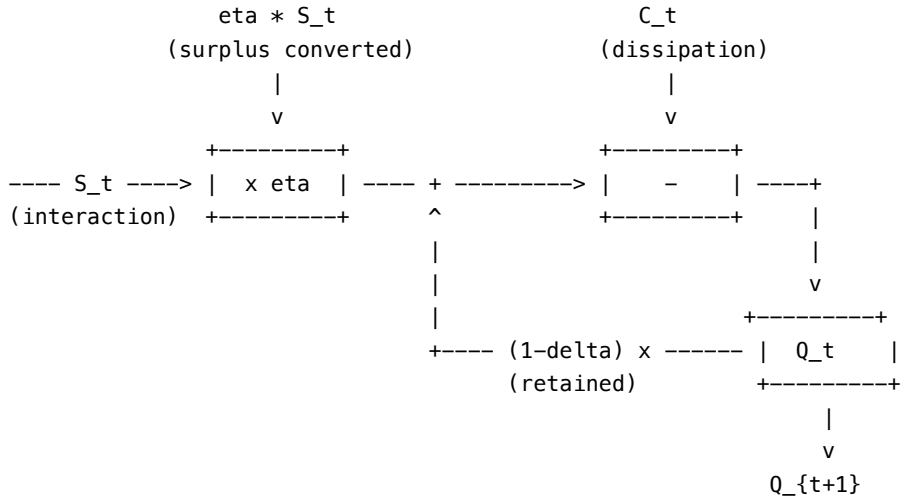
A system can be both at once in different subsystems, and usually is: a deliberating leadership layer sitting on top of a sealed operational layer. The survey protocol of Curzi [2026i] is designed to catch that, and its own worked office is an instance of the second kind — `delta = 0.003090` per week, a half-life of 224 weeks, and a stock in undetected decline.

7.4.7 4.7 Feedback is not regress

The distinction that resolves the ordinary-language confusion is between feedback and regress.

```
REGRESS                              FEEDBACK
-------                              --------
each step ADDS a new frame           each step UPDATES one state
history stored explicitly            history stored as a discounted sum
weight of step k = 1 forever         weight of step k = (1-delta)^age
resource use grows with t            resource use constant in t
terminates by exhaustion             terminates by convergence
no fixed point                       unique fixed point (delta > 0)
"we keep going in circles"           "we are converging on a practice"
```

The circuit is the diagram this framework returns to:



The loop from Q back into the summing junction through the gain (1-delta)

```
IS the recursion. delta is the loss in that path. If delta = 0 the loop
gain is exactly 1 and the integrator never settles.
```

A control engineer reading that diagram will report loop gain 1, marginally stable, pole on the unit circle. That is precisely correct, and it is the same statement as *no contraction, no fixed point* [Åström and Murray, 2008]. The organisational complaint that a group is going round in circles and the control-theoretic complaint that the pole is on the unit circle are the same complaint. The fix in control engineering is not to break the loop but to move the pole inside the unit circle — that is, to introduce loss. The framework's whole project is the organisational version of that move, and the observation that social systems are feedback systems in the engineer's sense rather than in the metaphorical one goes back to the first generation of cybernetics [Wiener, 1948].

7.4.8 4.8 The horizon of arrival

Corollary 3.2 says growth continues exactly while `eta S_t > delta Q_t + C_t`. With the drive held constant this is a threshold in `Q`:

```
grow while  Q_t < (eta S − C)/delta = Q_inf
```

so the fixed point is not only where the system ends up but the boundary between the growing and shrinking regimes. Below it the system grows, above it the system shrinks, and the second case is not hypothetical: it is what happens to an organisation whose drive falls — a key person leaves, `S` drops — while its stock stays high. Worked with `delta = 0.10`, `C = 0.20`, `eta S = 0.8243`:

`Q_t`	`delta Q_t + C_t`	vs `eta S_t = 0.8243`	verdict
0.0000	0.2000	less	grow
2.0000	0.4000	less	grow
4.0000	0.6000	less	grow
6.0000	0.8000	less	grow
6.2433	0.8243	equal	**hold — this is `Q_inf`**
7.0000	0.9000	greater	shrink
10.0000	1.2000	greater	shrink

Two structural points follow. First, **the maintenance term `delta Q_t` grows with success.** At `Q = 0` the system spends 0.20 per period on upkeep; at `Q = 6.24` it spends 0.824. Same organisation, same `delta`, four times the upkeep, purely because there is more to keep. Second, the headroom gauge along the same trajectory, with `eta ln N = 0.5 * 2.0794 = 1.0397`:

`Q_t`	0	2	4	6	6.2433	8	10
`h_t`	0.8397	0.6397	0.4397	0.2397	0.2154	0.0397	**-0.1603**

`h` reaches zero at `Q = 8.397`. Beyond that point Corollary 3.4 applies in its hard form: no arrangement of people, no meeting design, no facilitation technique and no interaction geometry whatsoever prevents this system's stored order from decreasing. That is not a motivation problem; it is an inequality. The same shape appears in Curzi [2026], where a 15 per cent annual hard-currency drag with a 25.17 per cent local-currency allocation drives the headroom coefficient to exactly −0.500000 and cancels the entire `eta ln(N) = 0.377583` diversification gain — a cost term consuming a logarithmic gain, in a different domain and with the same arithmetic.

7.4.9 4.9 Scale-freedom, badged

The storage law is applied across this corpus at every scale — a desk, a room, a team, a building, an institution, a city — with only `(delta, eta, N, C)` changing between them. **This is a modelling postulate about form, not a claim about physics.** It asserts that the same difference equation is a useful description at each scale. It does not assert that a city is a large room, and it does not assert any conservation law, thermodynamic principle, or mechanism connecting the scales.

The claim earns its place by being falsifiable at each scale independently. The protocol is given in Curzi [2026i]: fit `Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + eps_t` on training data with `0 <= a <= 1` and `b >= 0`, freeze the model, and report held-out performance against four baselines. If the fitted `a` is not in `[0,1]`, or the model does not beat a matched autoregression and a shuffled-input control, the form is wrong *at that scale* and the report says so. What the symmetry buys, where it holds, is transfer of method rather than transfer of numbers: the `delta` of a whiteboard and the `delta` of a constitution differ by orders of magnitude, and only the procedure for measuring both is identical. Curzi [2026v] exhibits the same style of claim in a domain where exhaustive verification was possible — the decomposition of the star figure {N/k} into `gcd(N,k)` closed strokes, unicursal exactly when `gcd(N,k) = 1`, verified by exhaustive trace over 779 cases with zero mismatches — which is the standard a form claim should be held to wherever the domain is finite enough to permit it.

7.5 5 The circuit identity and its consequences

The claim of this section is stronger than an analogy. In discrete time the storage law is not *like* a first-order low-pass filter; it *is* one, term for term, with the same pole, the same time constant, the same step response, the same settling time and the same DC gain. Every result in the standard theory of first-order discrete systems is therefore available without modification, and none of it has to be rediscovered in organisational language [Oppenheim et al., 1996, Åström and Murray, 2008]. **The status of this identity should be fixed before it is used: it is a shared mathematical form, not a derivation from physical premises.** The same difference equation governs a charge accumulator with a leak and a stock of order with a decay rate; nothing about capacitance, current or energy is claimed to hold of an institution, and no physical result's authority transfers. What transfers is the mathematics, and the mathematics transfers completely. That matters practically, because the standard theory answers questions organisations ask constantly and answer badly: when an initiative should be reviewed (settling time); whether a single large event will work (impulse response against the input ceiling); whether a large push beats a steady cadence (frequency response plus a saturation bound); why a change took a year to appear (the pole); and why a two-year-old programme is still climbing (`t_95` at the measured `delta`).

7.5.1 5.1 The identity

Set `x_t := eta S_t − C_t` for the net drive. Then

```
Q_{t+1} = a Q_t + x_t,      a := 1 − delta
```

which is the canonical first-order discrete-time system. Its transfer function, taking the `z` transform with `Q_0 = 0`, is

```
H(z) = Q(z) / X(z) = 1 / (z − a) = z^{−1} / (1 − a z^{−1})
```

with a single pole at `z = a = 1 − delta`. The system is bounded-input bounded-output stable if and only if `|a| < 1`, that is if and only if `0 < delta < 2`. Since `delta` lies in `[0,1]` by con-

struction, the binding case is `delta > 0`. This is an algebraic identity and not a modelling assumption; the badge belongs one level up, on the decision to describe a room, a team or an institution by this difference equation at all.

storage law	circuit	control theory
`Q_t`	charge on the capacitor, `q = C_cap V`	state
`delta`	leakage conductance across the capacitor	`1` – pole
`1 − delta`	retention per period	pole location `a`
`eta S_t`	current in from the source	driven input
`C_t`	current drawn by the load	disturbance
`eta`	source transconductance	input gain
`eta / delta`	DC gain of the RC section	DC gain `H(1)`
`Q_inf`	steady-state charge	DC steady state
`(1−delta)^t Q_0`	initial charge bleeding off	transient / natural response
`1 / delta` (periods)	`tau = R C_cap` (seconds)	time constant
`S_t <= ln N`	source compliance limit	input saturation
`C_t <= (1−delta)Q_t + eta S_t`	capacitor cannot supply more than it holds	admissibility

7.5.2 5.2 The continuous section, and where `delta` comes from

The continuous first-order low-pass is

```
tau dV/dt = V_in(t) − V(t),      tau = R C_cap
```

Euler-discretising with step `dt`:

```
V(t+dt) = V(t) + (dt/tau) (V_in − V(t)) = (1 − dt/tau) V(t) + (dt/tau) V_in
```

Matching coefficients against `Q_{t+1} = (1−delta)Q_t + ...` gives

```
delta = dt / tau      i.e.      tau = dt / delta = 1/delta   periods
```

`delta` is therefore not a property of the system alone. It is the ratio of the observation period to the system's own time constant. A `delta` of 0.05 measured weekly and a `delta` of 0.19 measured monthly can be the same system, since `1 − (1−0.05)^4 = 0.1855`. Section 5.7 gives the conversion rules and the errors that follow from ignoring them.

7.5.3 5.3 DC gain and the steady state

Evaluating the transfer function at `z = 1`, that is at zero frequency, gives

```
H(1) = 1/(1 − a) = 1/delta
```

so a constant net drive `x` produces steady state `Q_inf = x/delta`. Substituting `x = eta F(u*) − c*` recovers Theorem 3.5 exactly:

```
Q_inf = (eta F(u*) − c*) / delta
```

Theorem 3.5 is the DC gain of the circuit. With surplus in nats and cost in the same units per period, the DC gain from surplus to stored order is `eta/delta`: each sustained unit of surplus per period buys `eta/delta` units of permanent stock. Two consequences are constantly conflated and are worth separating. `eta` sets how much of each period's interaction becomes structure; `eta/delta` sets how much stock a *sustained* rate of interaction supports. A system with `eta = 0.4` and `delta = 0.08` has a DC gain of 5.0. Halving `delta` to 0.04 doubles the DC gain to 10.0 without touching `eta` at all. Retention is a gain term, and organisations almost never treat it as one.

A note on deliberately raising `C`. The DC gain makes cost look like a term no one would willingly increase, and the popular instruction to leave one's comfort zone looks from here like straightforward self-harm. That reading is too coarse, and the measured answer splits in two. A cost rise applied *globally* is strictly harmful and monotonically so: it lowers `eta F(u*) − c*`, lowers `Q_inf`, and lengthens every departure; in the worked case elsewhere in this corpus an operator's wait to fund a departure rises from 32.0 weeks at baseline to 42.5, then 89.2, then never, because above a critical operating cost `C_crit = eta M − lambda E_b` the departure is unreachable at any level of desire and unlimited time. A cost rise *confined to the origin* genuinely shortens the wait, because the barrier is measured as excess cost over the starting configuration, so raising the origin lowers the budget and the barrier together — 32.0 weeks down to 9.9, bought at a permanent fund 37 per cent smaller. The question is therefore narrower than a denunciation: whether it is a cost that can be left behind or a cost that is carried. Local discomfort at the origin works; global discomfort — attention, sleep, health, regulatory load — follows the operator out of the basin and destroys the platform. The empirical literature on adaptation supports the mechanism from the other side: hedonic adaptation returns a subject toward a baseline after both fortune and misfortune [Helson, 1948, Brickman et al., 1978], though the strong form of the treadmill has been revised and the revision restores room for a raised set point [Diener et al., 2006].

7.5.4 5.4 The pole, the time constant, and the settling table

The pole `a = 1 − delta` is the multiplier applied to the whole stored state each period. Three regimes, all of which occur in real organisations:

```
|a| < 1   (delta > 0)      STABLE. Contraction. Unique fixed point.
                           Transients decay as a^t. Errors forgotten.
                           --> Section 4: this is why iteration arrives.

 a  = 1   (delta = 0)      MARGINALLY STABLE. Pure integrator.
                           No fixed point. Any constant input ramps forever.
                           Initial condition never decays.
                           --> "we go round in circles"; the unbounded ledger.

|a| > 1   (delta < 0)      UNSTABLE. Not admissible here: delta lies in [0,1].
                           Would correspond to stored order that amplifies itself
                           without input. The law does not permit it.
```

The third case is worth a sentence even though the law excludes it, because organisations *describe* themselves as being in it — the culture reinforces itself, it snowballs. Under the storage law that behaviour is not `a > 1`; it is a high `eta` combined with an `S` that rises with `Q`, which is a nonlinearity in the drive rather than instability in the pole. If a system genuinely shows self-amplification, the linear model is the wrong model, and the fitting procedure of Curzi [2026i] will report `a > 1` — a signal to stop and re-specify rather than to celebrate.

Two definitions of the time constant are in circulation and differ by a few per cent at ordinary `delta`:

```
tau_approx = 1 / delta                (the engineering rule of thumb)
```

```
tau_exact  = −1 / ln(1 − delta)       (the true e-folding time of a^t)
```

`tau_exact` is the number of periods for the transient `a^t` to fall to `1/e`; `tau_approx` is its first-order approximation and is the number in general use. They agree to better than 1 per cent for `delta <= 0.02` and to about 5 per cent at `delta = 0.10`. Below `delta = 0.30` either is adequate for design work; above it the exact form is required, and above `delta = 0.5` time constants should be abandoned and settling times read directly. Half-life follows immediately:

```
t_half = ln 2 / (−ln(1 − delta)) = 0.6931 * tau_exact
```

All figures below are computed from `t_p = ln(1−p) / ln(1−delta)`.

delta	pole `1−delta`	tau = 1/delta	tau_exact	half-life	t_63	t_95	t_99	3 tau	5 tau
0.001	0.999	1000.00	999.50	692.80	999.17	2994.23	4602.87	3000.0	5000.0
0.005	0.995	200.00	199.50	138.28	199.43	597.65	918.73	600.0	1000.0
0.01	0.990	100.00	99.50	68.97	99.47	298.07	458.21	300.0	500.0
0.02	0.980	50.00	49.50	34.31	49.48	148.28	227.95	150.0	250.0
0.05	0.950	20.00	19.50	13.51	19.49	58.40	89.78	60.0	100.0
0.08	0.920	12.50	11.99	8.31	11.99	35.93	55.23	37.5	62.5
0.10	0.900	10.00	9.49	6.58	9.49	28.43	43.71	30.0	50.0
0.15	0.850	6.67	6.15	4.27	6.15	18.43	28.34	20.0	33.3
0.20	0.800	5.00	4.48	3.11	4.48	13.43	20.64	15.0	25.0
0.30	0.700	3.33	2.80	1.94	2.80	8.40	12.91	10.0	16.7
0.40	0.600	2.50	1.96	1.36	1.96	5.86	9.02	7.5	12.5
0.50	0.500	2.00	1.44	1.00	1.44	4.32	6.64	6.0	10.0
0.60	0.400	1.67	1.09	0.76	1.09	3.27	5.03	5.0	8.3
0.70	0.300	1.43	0.83	0.58	0.83	2.49	3.82	4.3	7.1
0.80	0.200	1.25	0.62	0.43	0.62	1.86	2.86	3.8	6.3
0.90	0.100	1.11	0.43	0.30	0.43	1.30	2.00	3.3	5.6

The rule of thumb of `3 tau` to 95 per cent and `5 tau` to 99 per cent is conservative in discrete time, which is the safe direction. Checking it exactly, the fraction of the step reached at `k tau` with `tau = 1/delta` is:

delta	at `1 tau`	at `2 tau`	at `3 tau`	at `4 tau`	at `5 tau`
0.02	0.6358	0.8674	0.9517	0.9824	0.9936
0.05	0.6415	0.8715	0.9539	0.9835	0.9941
0.10	0.6513	0.8784	0.9576	0.9852	0.9948
0.20	0.6723	0.8926	0.9648	0.9885	0.9962
continuous `1 − e^{−k}`	0.6321	0.8647	0.9502	0.9817	0.9933

So `3 tau` delivers 95.2 to 96.5 per cent and `5 tau` delivers 99.4 to 99.6 per cent, always slightly ahead of the continuous-time reference. Use `3 tau` for *substantially arrived* and `5 tau` for *settled*.

Design rule R2.1: no initiative is evaluated before `3 tau` periods have elapsed. Evaluating a first-order system at `1 tau` and finding it 65 per cent of the way there is not a finding about the initiative; it is a finding about the reviewer's calendar.

7.5.5 5.5 Step response

A step is a permanent change in the drive: a policy that starts and stays, a new standing meeting, a hire, a rule. Formally `x_t = x` for all `t >= 0` and `x_t = 0` before.

```
Q_t = Q_inf (1 − a^t) + Q_0 a^t,     Q_inf = x/delta
```

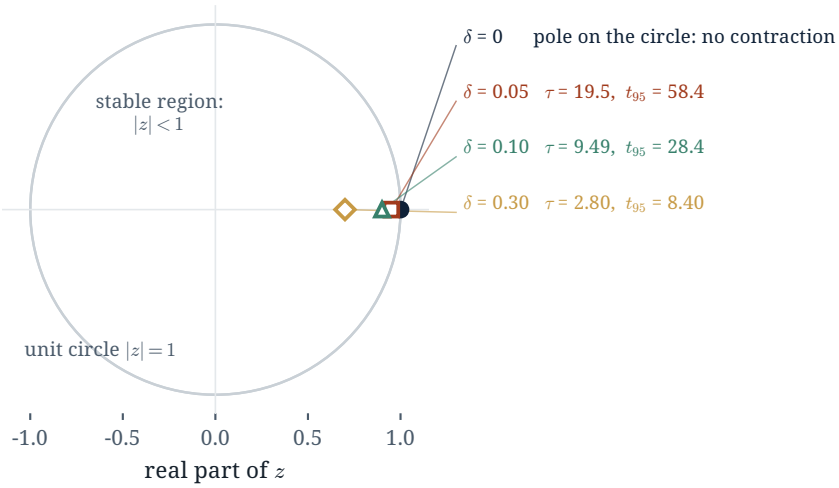


Figure 7.4: The single pole sits at `z = 1 - delta` on the real axis: strictly inside the unit circle for every `delta > 0`, and *on* it at `delta = 0`, where the map is an isometry and no contraction exists. Distance from the circle sets the calendar — `delta = 0.05` gives `tau_exact = 19.5` and `t_95 = 58.4` periods, `delta = 0.30` gives 2.80 and 8.40. Schematic drawn from the master table of Section 5.4; generated by `figs/src/CCP-007_fig04.py`.

Two exponentials with the same rate: the old state bleeding out and the new state filling in. There is no overshoot, no oscillation and no ringing. A first-order system with a real pole in `(0,1)` approaches its new steady state monotonically from whichever side it starts on. Worked at `delta = 0.08` and `x = 0.3511` per week, so `Q_inf = 4.3888`:

week	`Q_t`	% of `Q_inf`
0	0.0000	0.00%
1	0.3511	8.00%
2	0.6741	15.36%
4	1.2447	28.36%
6	1.7276	39.36%
8	2.1364	48.68%
12	2.7752	63.23%
18	3.4104	77.71%
24	3.7955	86.48%
26	3.8867	88.56%
36	4.1707	95.03%
48	4.3086	98.17%
52	4.3314	98.69%
60	4.3593	99.33%

At the six-month review, week 26, the system is at 88.6 per cent of where it is going; at the annual review, week 52, it is at 98.7 per cent. The visible slowdown between month six and month twelve is not the initiative losing momentum; it is `a^t` doing exactly what `a^t` does. An organisation that reads that curve as fatigue and responds by adding input has misdiagnosed a pole as a morale problem. That misdiagnosis is not idiosyncratic: people systematically misjudge the behaviour of stocks with delays even when the structure is explained to them and the incentives are aligned [Sterman, 1989].

7.5.6 5.6 Impulse response, the input ceiling, and the honest scope of the cadence argument

An impulse is a single period of input followed by nothing: the offsite, the launch event, the one-time grant, the visiting expert.

```
Q_t = A a^{t}   for t >= 0, from a single input A at t = 0
```

Pure decay. The whole contribution of an impulse is `A` at the moment it lands and `A a^t` thereafter, and its total contribution to the running integral over all time is `A/delta`. That last figure is what makes intensity sound competitive: an impulse and a sustained drive with the same total input have the same integrated effect over infinite time. In a linear time-invariant system this is exactly true, and honesty requires saying so plainly.

The reason intensity nevertheless loses is not the linear algebra. It is the saturation bound. The range result gives `0 <= F(u) <= ln N`, so the input in any single period is bounded by `eta S_t <= eta ln N`. A year's worth of surplus cannot be delivered in a weekend, because the per-period ceiling is a hard constraint of the law rather than a matter of stamina. Some ceilings, computed:

`eta`	`N`	`ln N`	per-period ceiling `eta ln N`
0.30	12	2.4849	0.7455
0.40	6	1.7918	0.7167
0.50	8	2.0794	1.0397
0.70	4	1.3863	0.9704

Meanwhile cost does not saturate: `C_t` accrues every period whether or not any surplus was delivered. So the honest comparison between an impulse and a cadence is not *same total, same effect*. It is that the cadence can deliver `T` periods of ceiling-limited input while paying `T` periods of cost, whereas the impulse delivers *one* period of ceiling-limited input while paying the same `T` periods of cost.

The concession above is stronger than it looks and is stated plainly rather than left as an aside. Summing `Q_t` over 400 periods for a 26-unit annual budget delivered three ways at `delta = 0.06` gives

```
impulse      433.3333
weekly       433.3333
four-weekly  433.3333
B/delta   =  433.3333
```

identical to every digit. Cadence delivers no more integrated stored order than intensity, and any version of the cadence argument resting on the premise that it does is simply false for a linear time-invariant system. Two things do the actual work. First, **strict concavity, which needs no ceiling at all**: for a fixed budget `T` split into `n` occasions the delivered total `D(n) = n g(T/n)` is nondecreasing in `n` whenever `g` is concave with `g(0) = 0`, and `f(u) = ln(1 + (N−1)u)` is exactly that, so splitting an input across occasions strictly dominates delivering it at once. This argument is independent of the saturation bound and survives even where the ceiling is never approached. Second, **the ceiling clips the impulse**: in a worked annual rite with `eta = 0.35`, `N = 8` and a per-week ceiling of `0.7278`, a 26-unit once-yearly burst delivers `0.7278` — 2.8 per cent of the nominal budget, with the other 97.2 per cent not delivered late but not delivered at all. A third difference the infinite integral conceals is that the impulse decays to a steady state of exactly 0 while sustained periodic forcing holds `m/delta`: the integral is equal, the level is not.

7.5.7 5.7 Choosing the period, and the conversions that go with it

The period `dt` is an operational choice, and once made everything else is expressed in it. Mixing periods is a common and expensive error. Aggregating up, retention multiplies:

```
delta_m = 1 − (1 − delta)^m
```

Disaggregating down, take the root:

```
delta_1 = 1 − (1 − delta_m)^{1/m}
```

Worked from `delta = 0.05` per week:

horizon `m` (weeks)	combined `delta_m`
1	0.0500
2	0.0975
4	0.1855
13	0.4867
26	0.7365
52	0.9306

And in reverse: an annual `delta` of 0.4596 corresponds to a monthly `delta` of 0.0500, since `1 − (1−0.4596)^{1/12} = 0.05000`.

Design rule R2.2: state the period in the same sentence as the `delta`, always. A bare `delta = 0.05` is meaningless; `delta = 0.05` per week is a specification.

Design rule R2.3: choose `dt` so that `delta` lands roughly in `[0.02, 0.30]`. Below 0.02 the signal-to-noise ratio on the decay measurement is poor because almost nothing changes per period; above 0.30 the discrete approximation to the underlying continuous process is coarse and the `tau = 1/delta` rule of thumb fails badly — at `delta = 0.5` it overstates `tau_exact` by 39 per cent. A system with an eight-month half-life is measured monthly, not daily and not annually.

7.5.8 5.8 Frequency response: why stored order is low-pass

Evaluating the transfer function on the unit circle, `z = e^{j w}`, and normalising by the DC gain:

```
|H(w)| / |H(0)| = delta / sqrt(1 − 2a cos w + a^2),     a = 1 − delta
```

An input repeating with period `P` has fundamental frequency `w = 2 pi / P`. The fraction of that input's amplitude surviving into `Q`:

input period `P`	`delta`=0.02	0.05	0.10	0.20	0.50
2	0.0101	0.0256	0.0526	0.1111	0.3333
3	0.0117	0.0296	0.0607	0.1280	0.3780
4	0.0143	0.0362	0.0743	0.1562	0.4472
6	0.0202	0.0512	0.1048	0.2182	0.5774
8	0.0264	0.0669	0.1364	0.2804	0.6786
12	0.0390	0.0986	0.1995	0.3966	0.8069
26	0.0835	0.2081	0.4006	0.6800	0.9465

input period `P`	`delta`=0.02	0.05	0.10	0.20	0.50
52	0.1650	0.3910	0.6576	0.8799	0.9857
100	0.3062	0.6325	0.8590	0.9627	0.9961

And the `-3 dB` cutoff, where the response falls to `1/sqrt(2)` of DC:

`delta`	cutoff `w_c` (rad/period)	period at cutoff
0.01	0.01005	625.2
0.02	0.02020	311.0
0.05	0.05130	122.5
0.10	0.10546	59.6
0.20	0.22408	28.0
0.30	0.36052	17.4
0.50	0.72273	8.7

For small `delta`, `w_c` is very nearly `delta` itself. Read together, the design consequence is immediate. A weekly system with `delta = 0.05` passes 2.6 per cent of a fortnightly oscillation and 63.3 per cent of a two-year one. **Stored order responds to the slow component of what a system does and ignores the fast component almost entirely.** The week-to-week variation in how a team feels is invisible to `Q`; the two-year trend is nearly fully transmitted. The correct instrument for a low-pass system is therefore a sustained mean rather than a peak. A programme delivering a variable amount each week, averaging `x`, produces almost exactly the same `Q` as one delivering a flat `x`, because the varying part is filtered out. That is a liberating result for anyone running a real organisation: consistency is required on the timescale of `tau`, not week to week.

7.5.9 5.9 Periodic forcing: cadence and ripple

Take a fixed mean input rate `m` per period delivered in bursts of `A = m P` every `P` periods. The steady-state cycle is exact:

```
Q_max = A / (1 - a^P)        immediately after a burst
Q_min = Q_max * a^P          immediately before the next
ripple = Q_max - Q_min = A
```

For `delta = 0.10` and `m = 0.20` per period:

burst interval `P`	burst size `A`	`Q_max`	`Q_min`	ripple	ripple as % of mean	mean over cycle
1	0.200	2.0000	1.8000	0.2000	10.0%	2.0000
2	0.400	2.1053	1.7053	0.4000	20.0%	2.0000
3	0.600	2.2140	1.6140	0.6000	30.0%	2.0000
4	0.800	2.3263	1.5263	0.8000	40.0%	2.0000
6	1.200	2.5610	1.3610	1.2000	60.0%	2.0000
12	2.400	3.3446	0.9446	2.4000	120.0%	2.0000

Three things must be said, and the third is the important one. **First, the mean is invariant.** Every row has mean 2.0000, equal to `m/delta`. This is the linear time-invariant result and there is no way around it: for the same total input, cadence does not change the average stored

order. **Second, the trough falls with `P`.** At weekly cadence the system never drops below 1.80; at quarterly cadence it drops to 0.94, a 48 per cent swing. If there is any threshold below which the system misbehaves — a level of shared context below which meetings stop working, a level of trust below which people stop volunteering information — the sparse cadence spends time under it while the dense cadence never approaches it. Averages do not fail systems; minima do. **Third, at realistic parameters `m` cannot be held constant while `P` rises, because `A = mP` runs into the ceiling `eta ln N`.** Taking the working example of Section 5.10, `eta = 0.4` and `N = 6` give a per-week ceiling of `eta ln 6 = 0.7167`, while a four-week burst would need `4 * eta S = 4 * 0.5011 = 2.0044`, which is 2.8 times the ceiling. The burst is clipped to 0.7167. Over a year the weekly cadence delivers `52 * 0.5011 = 26.06` units of input and the four-weekly cadence delivers `13 * 0.7167 = 9.32`, which is 35.8 per cent as much, while cost is paid all 52 weeks in both cases.

The argument, stated exactly: **a scheduled practice is a periodic forcing function on a low-pass filter, and cadence beats intensity because surplus saturates at `ln N` per period while cost does not saturate at all, so the only way to deliver a large total is to deliver it often.** The scope correction of Section 5.6 applies in full: the ceiling is half the reason, strict concavity is the other half, and neither half amounts to the claim that cadence delivers more integrated input — it does not.

Note the corollary, which is unintuitive and useful: as `P` shrinks below `tau`, the ripple vanishes and the practice becomes indistinguishable from a constant policy. At `delta = 0.10` a period-4 forcing transmits only 7.4 per cent of its oscillation. The filter has converted a discrete ceremony into a steady level. A practice frequent enough relative to `tau` stops being an event and becomes infrastructure — which is what a working practice feels like from inside, and exactly the condition under which Star and Ruhleder [1996] observe that infrastructure becomes visible only on breakdown.

7.5.10 5.10 The three input regimes, computed

Parameters: `delta = 0.08` per week, `eta = 0.40`, `N = 6`, `u* = 0.50`, `C = 0.15` per week.

```
S = F(u*) = ln(1 + 5 * 0.50) = ln 3.5 = 1.2528
eta S = 0.5011        eta ln N = 0.4 * 1.7918 = 0.7167  (the ceiling)
net weekly drive = eta S - C = 0.3511
Q_inf = 0.3511 / 0.08 = 4.3888
tau_exact = 11.99 weeks,  t_95 = 35.93 weeks,  t_99 = 55.23 weeks
```

Regime A — the impulse. A two-week intensive session run at the input ceiling of 0.7167 per week, then no further surplus input for the rest of the year. Cost continues at 0.15 per week throughout.

week	0	1	2	3	4	8	12	26	52
`Q_t`	0.0000	0.5667	1.0881	0.8510	0.6329	0.0000	0.0000	0.0000	0.0000

Peak at week 2, gone by week 8. The session is not useless — it puts 1.0881 units of order into a system that had none — but the drive stops and the cost does not, so the system runs down to the admissibility floor within six weeks. This is the arithmetic of every intensive event whose effects *wore off*. They did not wear off; they decayed at `(1-delta)^t` while `C` continued.

Regime B — the sustained step. The same input delivered every week at `eta S = 0.5011`.

week	0	1	2	3	4	8	12	26	52
`Q_t`	0.0000	0.3511	0.6741	0.9713	1.2447	2.1364	2.7752	3.8867	4.3314

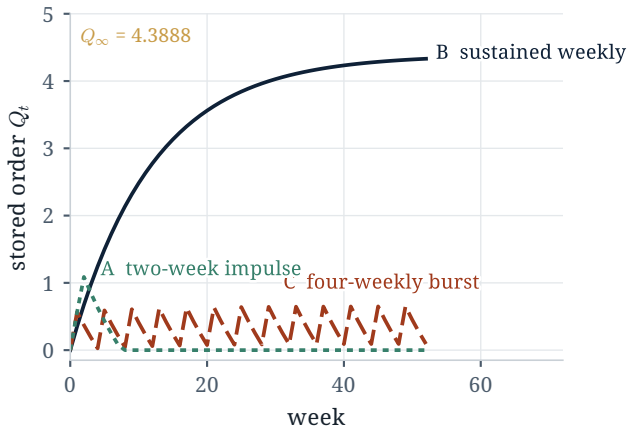


Figure 7.5: Cadence against intensity at `delta = 0.08`, `eta = 0.40`, `N = 6`, `u* = 0.50`, `C = 0.15` per week. The two-week impulse peaks highest at week 2 (1.0881) and is at the floor by week 8; the four-weekly burst is clipped by the ceiling `eta ln N = 0.7167` and oscillates about a mean of 0.365; only the sustained weekly regime arrives, reaching 4.3314 at week 52 against `Q_inf = 4.3888`. Modelled from the parameters stated in Section 5.10; generated by `figs/src/CCP-007_fig05.py`.

Regime B is *behind* Regime A at week 2, 0.674 against 1.088. At week 4 it has overtaken it. At week 52 it stands at 4.3314 against 0.0000. The regime that looked worse at the first checkpoint is the only one that arrives anywhere.

Regime C — sparse ritual. An intensive session every fourth week, sized to deliver four weeks of surplus at once — except that it cannot, because 2.0044 exceeds the ceiling 0.7167 and is clipped. Cost continues weekly.

week	0	1	2	3	4	8	12	26	52
Q_t	0.0000	0.5667	0.3714	0.1917	0.0263	0.0452	0.0587	0.4393	0.0916

In steady state this regime oscillates between 0.093 and 0.652 with a mean of 0.365 — 8.3 per cent of Regime B's steady state, from a schedule its participants would describe as meeting regularly. The monthly session is genuinely intensive and genuinely valuable; it simply cannot push more than 0.7167 through the ceiling in the week it happens, and the other three weeks pay 0.15 of cost with no input at all.

regime	input delivered over 52 weeks	`Q` at week 52	steady-state mean `Q`
A — two-week offsite	1.4334	0.0000	0
B — weekly sustained	26.0575	4.3314	4.3888
C — four-weekly burst	9.3171	0.0916	0.3647

Design rule R2.4: the cadence of a scheduled practice should be short relative to `tau`, and its per-session intensity should be sized to the ceiling `eta ln N` rather than to the ambition. Anything beyond the ceiling is discarded in the period it is delivered.

7.5.11 5.11 What a high-pass version would mean

The low-pass character is a consequence of `Q_t` appearing on the right-hand side with a positive coefficient. A high-pass version would replace the input with its first difference:

```
Q_{t+1} = (1 − delta) Q_t + eta (S_t − S_{t−1}) − C_t
```

Such a system responds only to *change* in interaction and is completely blind to a constant level. A steady, high-surplus practice sustained for years would drive it to `−C/delta`, which is negative: the system would treat a stable excellent state as no input at all. Systems that behave this way exist, and naming them is useful because they are usually built by accident.

- **Attention markets.** The stock of attention responds to the derivative of events rather than their level. A chronic condition of constant severity generates no coverage.
- **Alarm and exception handling.** By design, an alarm fires on change. This is correct for alarms and catastrophic for anything that must be maintained: a system reporting only exceptions cannot report that the ordinary thing is still working.
- **Novelty-driven incentive schemes.** Promotion for launching, none for operating. The organisation's `Q` receives credit for `S_t − S_{t−1}` and none for `S_t`, so maintenance is structurally unrewarded even where everyone agrees it matters.
- **Difference-only reporting.** A dashboard showing only period-over-period deltas is a high-pass filter installed on top of a low-pass system, and it systematically hides what the system actually accumulates.

Stored social order is low-pass, and it must be, because `Q` is a stock and a stock integrates. Trust, habit, institutional memory and furnishing all respond to the sustained level of what is done, with fast fluctuations averaged away. A system whose *measurement* is high-pass while its *dynamics* are low-pass will report noise and miss the trend. This is not a subtle statistical point; it is a mismatch of filter type between plant and instrument, and it is among the most common instrumentation faults in organisational reporting [Forrester, 1961].

Design rule R2.5: match the filter type of the instrument to the filter type of the system. Report levels for stocks and differences for flows. A dashboard with only difference columns cannot measure `Q`.

7.5.12 5.12 Cascades and delayed control

Real systems are often two integrators in series. Individual skill accumulates with `delta_1`, and team capability accumulates from the flow of trained individuals with `delta_2`:

```
A_{t+1} = (1 − delta_1) A_t + x_t
B_{t+1} = (1 − delta_2) B_t + delta_1 A_t
```

Worked with `delta_1 = 0.20`, `delta_2 = 0.08`, unit step into stage 1:

t	stage 1	stage 2
0	0.0000	0.0000
2	1.8000	0.2000
5	3.3616	1.4998
8	4.1611	3.2060
12	4.6564	5.4129
20	4.9424	8.6650
30	4.9938	10.8027
40	4.9993	11.7593
60	5.0000	12.3600
80	5.0000	12.4736

Stage 1 settles at `1/delta_1 = 5.0`; stage 2 settles at `1/delta_2 = 12.5`. Stage 2 reaches half its final value at week 14, 95 per cent at week 43 and 99 per cent at week 62, against 36 and

55 for a single stage with `delta = 0.08` driven directly. Two facts carry forward, both checked numerically. **The cascade is slower and has an S-shaped rise**, because stage 2 sees nothing until stage 1 has something to give it, so the apparent *nothing is happening* phase in the first two months is structural. **The cascade does not overshoot**: two real poles in `(0,1)` give a monotone step response, verified by simulation over 400 periods, and overshoot and ringing require complex poles, which a cascade of leaky integrators cannot produce on its own.

Overshoot comes instead from feedback with delay — a controller acting on a stale measurement. Simulating `Q_{t+1} = (1−delta)Q_t + k(target − Q_{t−L})` with `delta = 0.10`, `k = 0.30` and target 5.0:

measurement lag `L`	peak `Q`	final `Q`	overshoot
0	3.7500	3.7500	0.0%
2	4.8826	3.7500	30.2%
4	6.5784	3.7576	75.1%

At `L = 4` with a more aggressive gain `k = 0.50` the loop goes unstable outright: the simulated trajectory peaks at 39.2 and then diverges into oscillation with negative excursions, which under the storage law means the admissibility condition has been violated and the model has broken.

Design rule R2.6: the reporting lag of a management control loop must be short relative to `tau`, and the correction gain must fall as the lag rises. An organisation adjusting quarterly on the basis of annual data is running `L` of four periods with a large gain, and the resulting whipsaw is not indecision — it is a delayed feedback loop behaving exactly as delayed feedback loops behave [Åström and Murray, 2008, Sterman, 1989].

7.6 6 Loss as an object of engineering

`delta` is the parameter that makes the machine work and it is the one no organisation puts on a dashboard. Section 4 showed why it is load-bearing: `delta > 0` makes the update map a contraction, and contraction supplies a destination, a destination independent of history, and error decay. Section 5 showed what it controls: the pole, the time constant, the settling time, the DC gain and the cutoff frequency, each of which is `delta` wearing a different hat. This section takes `delta` as an object of engineering and answers four questions in order. What is it made of? How is it measured? How is it driven down? And — the question this section exists for, and the one most often skipped — when should it be driven *up*?

7.6.1 6.1 Where `delta` sits, and what the results say about it

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,     delta in [0,1]
```

`delta` is the fraction of stored order lost per period **in the absence of maintenance**. That qualifier is the whole definition. `delta` is not how much is lost; it is how much would be lost if nothing were done, and the maintenance actually performed appears elsewhere in the equation — partly in `eta S_t`, where maintenance regenerates structure, and partly in `C_t`, where the effort of maintaining is paid for. **Confusing `delta` with observed net decline is the most common estimation error in this framework**, and Section 6.4 gives the procedure that avoids it.

Reading the results of Section 3.2 as statements about `delta`: Theorem 3.1 makes `delta` the discount rate on history, with the weight on an input `j` periods old equal to `(1−delta)^j`. Corollary 3.2 identifies `delta Q_t` as the **maintenance load**, proportional to how much the system holds. Theorem 3.3 puts `delta` in the denominator of the ceiling, so halving `delta` doubles the

maximum attainable stock, and at `delta = 0` the bound degrades to `Q_t <= Q_0 + t eta ln N`, linear and unbounded. Corollary 3.4 says that as `Q` grows, `delta Q` consumes the cost budget, and past `Q = (eta ln N − C_t)/delta` no choice of interaction geometry prevents stored order from decreasing. Theorem 3.5 puts `delta` in the denominator of the fixed point. And the contraction modulus is exactly `k = 1 − delta`, so `delta > 0` is necessary and sufficient for existence, uniqueness and global convergence.

The transient `(1−delta)^t` is a geometric decay, so it has a half-life:

```
t_half = ln 2 / (−ln(1 − delta))
```

This is the single most useful derived quantity in the framework for communicating with non-specialists. *Our `delta` is 0.10 per week* means nothing to a director. *Half of everything this group knows how to do that is not written down or exercised is gone in about six and a half weeks* is a sentence that gets a budget. Computed for the range that matters:

`delta` (per period)	half-life (periods)	to 10% remaining	to 1% remaining	retained after 10 periods
0.01	68.97	229.11	458.21	0.9044
0.02	34.31	113.97	227.95	0.8171
0.05	13.51	44.89	89.78	0.5987
0.10	6.58	21.85	43.71	0.3487
0.20	3.11	10.32	20.64	0.1074
0.30	1.94	6.46	12.91	0.0282
0.50	1.00	3.32	6.64	0.0010

Note the nonlinearity. Moving from `delta = 0.10` to `delta = 0.05` roughly doubles the half-life, 6.58 to 13.51, and doubles `Q_inf`; moving from 0.05 to 0.02 multiplies it by 2.54. **Each halving of `delta` doubles both the ceiling and the settling time**, and the second of those is the price nobody budgets for.

Four clauses attach specifically to `delta`. It is **period-dependent** and meaningless without a stated period. A `delta` measured on a proxy `Y` is the decay rate of the proxy, so the justification that `Y` is not a repackaging of the input `U` must precede the measurement, or what is being measured is the decay of the measuring instrument. The exponential decay form is a **modelling assumption, badged**: real loss processes can be non-exponential — a cliff at a single departure, a step at a reorganisation, a long tail from written records — and fitting an exponential to a cliff produces a number wrong in a specific and predictable direction. And every `Q_inf` computed from an estimated `delta` inherits its uncertainty *in the denominator*, so small errors in `delta` produce large errors in `Q_inf`: a 20 per cent error in `delta` produces a 25 per cent error in `Q_inf`. Ranges, not points.

7.6.2 6.2 What `delta` is made of

Loss is not one thing. In any real system it decomposes into named channels, each with its own mechanism, its own measurement, and its own available intervention. Eight recur.

(1) Procedural forgetting. An individual's memory of a procedure not executed recently. Fast for multi-step procedures with arbitrary ordering, slow for procedures with strong internal logic. This is the channel that periodic exercises address.

(2) Turnover. A departure removes whatever fraction of `Q` lived in that person and nowhere else. Section 6.6 treats it separately because it is usually the largest single channel and the easiest to compute.

(3) Documentation entropy. Written records do not decay by fading; they decay by **becoming false**. The document is intact and the world has moved. Documentation entropy is why a

written artefact does not reduce `delta` to zero, and why *executable* documentation — a test, a script, a check that fails when it is wrong — has a materially lower `delta` than prose.

(4) Environment drift. The system the stored order refers to has changed: a runbook referencing a decommissioned server, a floor plan predating the renovation, a policy naming a department that no longer exists. This channel is invisible until exercised, which is what makes it dangerous.

(5) Relationship decay. Trust and working familiarity between two people, absent interaction. This is a genuine `delta` channel and one of the few whose maintenance cost is almost pure `C`: a standing one-to-one produces very little `S`, and its purpose is not to produce `S` but to suppress `delta`.

(6) Skill atrophy. Distinct from procedural forgetting: the procedure is remembered and the execution has degraded. Motor skill, judgement calibration and fluency sit here.

(7) Physical wear. Furniture, buildings, instruments, acoustic treatment that has come loose. The most literal `delta` and the only one with an established maintenance profession attached — which is instructive, because the reason building maintenance is a discipline and institutional maintenance is not is that physical `delta` is visible.

(8) Context loss. The rule survives and its rationale does not. **This channel is worse than plain forgetting**, because it leaves behind a rule that is still enforced and can no longer be evaluated. A rule whose rationale has decayed cannot be retired, because nobody can establish that retiring it is safe, so context loss converts ordinary `Q` into permanent `C`. Ordinary loss removes structure and frees the budget it consumed; context loss removes the justification and leaves the structure drawing cost forever. **Context loss is the mechanism by which decay produces ossification rather than clearance**, and it is why deliberate shedding must be a design activity rather than a natural process. Scott [1998] and Vaughan [1996] document the institutional form from opposite directions: a rule set whose local justification has been erased, and a normalised deviance whose original risk argument nobody can reconstruct.

7.6.3 6.3 Channels compose multiplicatively

Retention multiplies; loss does not add. If `k` independent channels each retain a fraction `1 − delta_i` per period, then

```
1 − delta_total = product_i (1 − delta_i)      so      delta_total = 1 − product_i (1 − delta_i)
```

Worked, for four channels in a team:

channel	`delta_i`
turnover	0.030
procedural forgetting	0.050
documentation rot	0.020
tool churn	0.015
naive sum	**0.1150**
true combined `1 − prod(1−delta_i)`	**0.1105**

The naive sum overstates by 0.0045, about 4 per cent relative. At these magnitudes the difference is small and the sum is serviceable. It stops being serviceable quickly: four channels at 0.20 each give a naive sum of 0.80 and a true combined `delta` of 0.5904, an overstatement of 35 per cent.

Design rule R3.1: use the product form whenever any channel exceeds 0.10, and always in the design record. The sum is acceptable only as a check at small `delta`.

The corollary matters more than the arithmetic. **The channels are not fungible, and reducing the largest one has the largest effect on the product.** Halving the 0.050 forgetting channel

to 0.025 takes the combined `delta` from 0.1105 to 0.0871, a 21.2 per cent reduction; halving the 0.015 tool-churn channel takes it only to 0.1037, a 6.2 per cent reduction. Attack the biggest channel; the multiplicative structure rewards it.

7.6.4 6.4 Measuring `delta`: the zero-input window

The procedure is six steps. **First**, choose the order proxy `Y`, state what it counts, and write the justification that `Y` is not a repackaging of the input `U`; have a second person check it. **Second**, find or create a window in which `S` contributions to `Y` are near zero. Candidates: a plant shutdown; a hiring or feature freeze; a holiday period; the interval after a contributor leaves and before a replacement starts; a decommissioned tool still nominally in service; a documented practice nobody has exercised. **Third**, measure `Y` at several points in the window — two points give a number, five give a fit and a residual, and the residual is where the information about model form lives. **Fourth**, fit log-linearly: regress `ln Y_t` on `t`, so the slope is `ln(1-delta)` and `delta = 1 - exp(slope)`, and report `R^2`. **Fifth**, convert to the working period if the window was measured on a different one, using `delta_1 = 1 - (1-delta_m)^{1/m}`. **Sixth**, sanity-check against the half-life: whether `ln 2 / (-ln(1-delta))` matches what practitioners say when asked how long the capability would survive if work stopped. A factor-of-two disagreement is normal; a factor-of-ten disagreement means the proxy is wrong. The two-point estimator, for when that is all the record supports, is `delta = 1 - (Y_k / Y_0)^{1/k}`.

Five pitfalls, each with its detection. **Contaminated window** — input leaked in; detected by a plateau or upturn in the fit, or poor `R^2` with structured residuals; fixed by shortening the window or finding a cleaner one. **The proxy is the input** — `Y` is really a measure of `U` or `S`, so its apparent decay is just the activity stopping; detected because `Y` drops to near zero the moment the window opens rather than decaying; this is the most common single error in the field. **Cliff, not exponential** — the loss was one departure rather than a rate; detected by a single large step in the residuals and wildly heterogeneous per-step estimates; fixed by modelling the departure as an event and estimating `delta` from the remainder. **Censoring** — the proxy has a floor, and the last few units are the most robust; detected because the fit flattens at low `Y`; fixed by fitting only the upper range and saying so. **Reactive maintenance** — people notice `Y` falling and intervene, which is `eta S` re-entering; detected by asking; fixed by instrumenting the window so intervention is recorded.

Worked estimation. A monthly proxy over a six-month zero-input window:

month `t`	0	1	2	3	4	5	6
`Y_t`	100	92	84	79	71	66	60

Per-step naive estimates `1 - Y_t/Y_{t-1}` run 0.0800, 0.0870, 0.0595, 0.1013, 0.0704, 0.0909. The spread is 0.0595 to 0.1013, which is what single-step estimates always look like and why none should be used alone. The endpoint two-point estimator gives `1 - (60/100)^{1/6} = 0.0816`. The log-linear regression gives slope `-0.08446`, so **`delta = 0.0810`**, fitted `Y_0 = 100.16`, and **`R^2 = 0.9980`**, with half-life `ln 2 / (-ln 0.9190) = 8.21` months. The regression and the two-point estimator agree to within 0.8 per cent here, which is a good sign: the decay really is close to exponential over this range. **Where they disagree by more than about 15 per cent, neither should be believed and the residuals should be inspected**, because the disagreement means the shape is not geometric and the whole exponential model is in question for that proxy.

The general form of this move — reading a system's storage constant off its observed relaxation — has a physical precedent in the fluctuation–dissipation theorem [Callen and Welton, 1951], and an architectural one in reverberation-time measurement [Sabine, 1922]. Neither licenses the sociological application; both establish that the measurement design is a recognised one rather than an improvisation.

Where no zero-input window exists, `delta` must be estimated jointly with `eta` from the full model — the estimation procedure of Curzi [2026i]: fit `Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + eps_t` with `0 <= a <= 1`, `b >= 0`, `a = 1−delta`, `b = eta`, on training data; freeze; and report held-out performance against persistence, direct-input, matched autoregressive and shuffled-input baselines. Two cautions. `a` and `b` are correlated in the fit — high retention with low conversion and low retention with high conversion describe the data almost equally well, and separating them requires variation in `U` that the data may not contain. And **a fitted `a` outside `[0,1]` is not a number to be clipped; it is a specification failure** and must be escalated as such. A single documented shutdown identifies `delta` directly and is worth more than a year of regression on confounded data.

7.6.5 6.5 Periods and conversions

`delta` is a per-period quantity and must always be reported with its period:

```
delta_m = 1 − (1−delta_1)^m          (short period to long)
delta_1 = 1 − (1−delta_m)^{1/m}      (long period to short)
```

Monthly to annual, computed:

monthly `delta`	annual retention `(1−delta)^12`	annual `delta`
0.01	0.8864	0.1136
0.02	0.7847	0.2153
0.05	0.5404	0.4596
0.10	0.2824	0.7176
0.15	0.1422	0.8578
0.20	0.0687	0.9313
0.30	0.0138	0.9862

And back the other way:

annual `delta`	implied monthly `delta`
0.05	0.0043
0.10	0.0087
0.20	0.0184
0.30	0.0293
0.50	0.0561
0.70	0.0955

A monthly `delta` of 0.05 — which sounds modest, and which a manager would describe as losing a bit of ground each month — destroys 46 per cent of unmaintained stored order in a year. That single row is the most persuasive line in most `delta` presentations, and it is arithmetic rather than rhetoric.

7.6.6 6.6 Turnover as a `delta` term

Turnover is usually the largest channel and has the cleanest formula. If a fraction `p` of members leave per period and a fraction `phi` of stored order is **person-borne** — held in individuals and nowhere else — then to first order

```
delta_turnover = p * phi
```

The two factors are independently controllable, and this is the practical point. Turnover can be reduced, which is expensive, partly outside the organisation's control, and sometimes undesirable; or the person-borne fraction can be reduced, which is cheap, entirely within the organisation's control, and beneficial for other reasons. **`phi` is the parameter, not `p`.** Computed at 20 per cent annual turnover, monthly `p = 0.0184`:

person-borne fraction `phi`	monthly `delta_turnover`	annual equivalent
0.90	0.0166	0.1818
0.75	0.0138	0.1538
0.60	0.0111	0.1249
0.40	0.0074	0.0849
0.20	0.0037	0.0433
0.05	0.0009	0.0110

Moving `phi` from 0.90 to 0.20 cuts the turnover channel by a factor of 4.5 without changing a single thing about retention or recruitment. Every mechanism that raises `eta` also lowers `phi`, because they are the same mechanisms: writing it down, building the object, training the second person, changing the default. This is the strongest instance of a general pattern — **the artefact that captures surplus is the artefact that survives the person who made it.**

Conway's observation enters here with a specific quantitative consequence. Systems mirror the communication structure of the organisation that builds them [Conway, 1968], which means a substantial part of `Q` is encoded in the communication topology rather than in any document: who talks to whom, which pairs have a working channel, where the informal escalation paths run. **A reorganisation is therefore a `delta` event, and a large one**, because it invalidates the part of stored order that lived in that topology. The reorganised system retains its documents and loses its routing. This is measurable: sample the *who would you ask about X* network for a fixed set of `X` before and after, and the fraction of queries whose answer changed or became *I don't know* is a direct estimate of the reorganisation's one-off `delta` contribution. It is almost never measured and is almost always large. The complementary result supplies the other half: adding people to a late project makes it later, because new arrivals must be brought up on stored order that is largely person-borne and the bring-up consumes the time of exactly the people who hold it [Brooks, 1975]. In the law's terms, onboarding is a `C` term proportional to `phi Q`, charged against the same budget as everything else, and the quadratic growth of communication paths is the mechanism.

7.6.7 6.7 Engineering `delta` down, and the two prices

The intervention always has the same form: **move stored order down the artefact ladder**, from media with high `delta` toward media with low `delta`.

```
THE ARTIFACT LADDER            typical delta        why
-------------------            (per month, rough)
one person's memory            0.10 - 0.30          forgetting + they may leave
two people's shared memory     0.06 - 0.20          redundancy halves the departure risk
spoken transmission / story    0.05 - 0.15          survives departure, drifts in content
written prose document         0.02 - 0.08          survives; goes silently false
diagram / checklist            0.02 - 0.06          harder to be ambiguous
executable check / test        0.005 - 0.02         FAILS when it becomes false
physical object / built form   0.001 - 0.02         must be actively removed to be lost
trained cohort (>3 people)     0.005 - 0.03         regenerates itself if cadence holds
```

```
These bands are ORDER-OF-MAGNITUDE ORIENTATION, badged as such. They are
not measurements. Each system's own zero-input window supplies its real numbers.
```

That badge is not decorative. The bands are working priors for structuring a survey, not findings. **Any number in a design record must come from that system's own zero-input window, not from this ladder.** The ladder indicates direction and rough magnitude; Section 6.4 says what is actually true in a given system.

The rungs, with what each buys. **Write it down** is the single largest step and the cheapest, converting a channel that includes departure risk into one that does not; it does not eliminate documentation entropy, so it lowers `delta` by roughly a factor of 3 to 5 rather than to zero. **Make it executable** is the largest reduction available short of building a physical object, and it is available in far more domains than assumed: automated tests, periodic restore drills, a checklist with a signature line, a form that will not submit if a field is inconsistent. **Executable documentation is the only kind that reports its own decay** — prose rots silently, and a test that rots turns red. **Build the object**: physical structure has the lowest `delta` available because losing it requires deliberate removal. **Train the second, third and fourth person**: redundancy in people is `phi` reduction, and past about three holders it becomes self-regenerating, because the cohort can rebuild a departed member's share without external input; below three holders a practice is a person with a job title. **Institute the periodic exercise**: at interval `P <= tau/2` a periodic forcing function becomes indistinguishable from a constant policy, and in `delta` terms the exercise suppresses procedural forgetting and environment drift simultaneously, because the exercise *reveals* the drift, which is the only way drift is ever revealed. **Apprenticeship**: the transfer of context rather than procedure, and the only mechanism this framework identifies that addresses channel (8), because rationale is not writable in a form anyone reads; it is expensive, slow, and has no substitute. **Design for patina**: Jacobs [1961] observed that aged buildings carry a mix of uses new construction cannot, partly because a building that ages well accumulates adaptation rather than requiring replacement, so a material that looks better worn has a lower effective `delta` than one that looks worse worn, because the replacement decision is deferred.

Both prices of a low `delta` are routinely omitted from proposals and both are computable. **Price one: `C` rises.** Every rung costs something to maintain — tests kept green, documents reviewed, drills run, second and third holders spending time holding. Since `C` enters the sustainability condition linearly and `delta` enters multiplied by `Q_t`, the trade is favourable only while the `delta Q_t` saved exceeds the `C` added. **At small `Q`, lowering `delta` is not worth its cost.** A three-month project should not build institutional memory; it should finish. **Price two: the system gets slower.** This is invisible and is the more dangerous of the two:

`delta`	half-life	`t_95`	`Q_inf` at drive 7.775
0.0810	8.21	35.5	96.0
0.0405	16.76	72.5	191.9
0.0368	18.49	79.9	211.3

Halving `delta` doubles `Q_inf` and doubles `t_95`. **The result is a system that holds twice as much and takes twice as long to change its mind about anything.** If the environment the system responds to moves faster than the new `t_95`, the design has bought stability it cannot use and lost adaptability it needed.

Design rule R3.2: state the new `t_95` in the same paragraph as the new `Q_inf`, every time a `delta` reduction is proposed.

There is a further structural bias against maintenance, and it is not cultural. Star and Ruhleder [1996] observe that infrastructure becomes visible only on breakdown, and the consequence in this model is exact. Maintenance is the activity that suppresses `delta`; its output is a non-event. A high-pass reporting instrument generates no signal from a non-event, so maintenance produces no visible result while it is working and a highly visible result the moment it stops. **The budget-allocation process is therefore systematically biased against exactly**

the activity that keeps `delta` low, and the bias requires nobody to be foolish — it follows from the filter type of the instrument. Three counters, all instrumentation rather than exhortation: report `delta` itself as a headline metric, since it is a level and is visible when things are fine; report the counterfactual maintenance load `delta Q_t`, which makes the invisible work quantitative; and run scheduled failure drills, which convert a non-event into an event on a schedule the organisation controls.

7.6.8 6.8 When `delta` should be high

This is the half of the subject that most treatments omit. `delta` is a per-component parameter, not a global constant. **Different parts of `Q` should have radically different `delta`, and setting them all low is a design failure rather than a design goal.**

What should decay fast: **stored interpersonal records**, whose persistence drives `u` toward zero within the affected pair by making the interaction predictable and defensive — high `delta` here is an engineering choice with a half-life; **obsolete procedure**, which draws `C` forever and produces nothing, and which once its rationale has decayed can be retired only by scheduled expiry rather than by argument; **any practice that made sense once**, since the whole point of a practice is that it does not require re-deciding, which is exactly what makes an obsolete one durable; **roadmaps, plans and forecasts**, which are supposed to decay, so a two-year-old roadmap still being executed is a `delta`-too-low failure; **metrics that have become targets**, because a measure that becomes a target ceases to be a measure [Goodhart, 1975, Campbell, 1979], and in this model a metric is stored order with a `delta`, so the correct response to capture is to raise the metric's `delta` — retire it and replace it — rather than to defend it, since a permanent metric is a guaranteed-corrupted metric; and **reputational classification of people**, since any stored judgement about a person's capability becomes self-confirming through assignment [Merton, 1948, Rosenthal and Jacobson, 1968], so high `delta` on such records is what makes development possible.

The load-bearing sentence of this section: a system with no forgetting cannot learn. The demonstration is one row of a table. An error injected at `t = 0` retains weight `(1-delta)^t`:

periods since the error	`delta`=0	0.01	0.05	0.10	0.20
1	1.00000	0.99000	0.95000	0.90000	0.80000
5	1.00000	0.95099	0.77378	0.59049	0.32768
10	1.00000	0.90438	0.59874	0.34868	0.10737
25	1.00000	0.77782	0.27739	0.07179	0.00378
50	1.00000	0.60501	0.07694	0.00515	0.00001
100	1.00000	0.36603	0.00592	0.00003	0.00000

Learning is the replacement of an old weighting with a new one, and replacement requires that the old weighting fall. **At `delta = 0` no weighting ever falls, so no replacement ever completes**, and what looks like learning is only accumulation: the new belief is added alongside the old one, both at full weight, and the system holds a contradiction it cannot resolve. This is not a metaphor about culture; it is the first column of the table.

Note also what high `delta` does at the other end of the ladder. A component with `delta = 0.25` has a half-life of 2.41 periods and 0.1 per cent residue after 24 periods. That is a component which can be replaced wholesale with no removal effort at all: **it is retired by declining to maintain it.** Deliberate high-`delta` design — sunset clauses, term limits, expiring authorisations, dated documents that go stale visibly — is the cheapest removal mechanism that exists, because it costs nothing to invoke.

The design output for any real system is therefore not a `delta` but a table of them, one per component, each justified:

component	target `delta` (per month)	half-life (months)	rationale
constitutional / safety-critical rules	0.002	346.2	must survive every individual; changing them must be deliberate and slow
safety procedures	0.010	69.0	durable, but must be re-derivable and periodically re-examined
core tooling and standards	0.050	13.5	durable within a technology generation; must not outlive it
team practice and convention	0.100	6.6	should be re-decided roughly twice a year by exercise
roadmap and plan	0.250	2.4	must go stale visibly; a stale plan being executed is the failure
interpersonal records / grudges	0.400	1.4	near-total clearance in six months by design

Design rule R3.3: no component gets `delta = 0`. Even the constitutional layer is 0.002, which is a 346-month half-life and an amendment process, not a stone tablet. A component at exactly zero has no fixed point, no error decay, and no mechanism of correction.

Design rule R3.4: `delta` should decrease monotonically with the component's intended lifetime, and every component must have a stated intended lifetime. A component to which nobody has assigned a lifetime will default to `delta = 0` through neglect, because retirement requires a decision and persistence requires none.

7.6.9 6.9 The trade computed: `delta` down against `eta` up

Take `eta = 0.5`, `N = 8`, `u* = 0.6` so `S = 1.6487`, `C = 0.20`, and `delta = 0.20`:

```
baseline:            Q_inf = (0.5 * 1.6487 - 0.20)/0.20 = 3.1216
halve delta to 0.10: Q_inf = (0.5 * 1.6487 - 0.20)/0.10 = 6.2433    (+100.0%)
double eta to 1.0:   Q_inf = (1.0 * 1.6487 - 0.20)/0.20 = 7.2433    (+132.0%)
```

Doubling `eta` beats halving `delta`, and beats it by 32 percentage points. The reason is structural: `eta` multiplies `S` before the cost is subtracted, so doubling it adds `eta S = 0.8243` to the numerator, whereas halving `delta` merely doubles whatever the numerator already is, including the drag of `–C`. The larger `C` is relative to `eta S`, the more the asymmetry favours `eta`. That is the general case, and it is the content of H5. Note the exception, because it is where this section earns its place: **`delta` reduction is the better lever whenever the numerator `eta F(u*) – C*` is already large and the problem is that nothing stays.** A system generating excellent surplus, converting it well, and losing it all by the next quarter has a retention fault, and no amount of additional `eta` fixes a leak. The diagnostic is simple: **if `Q` is oscillating or flat while `eta S` is comfortably above `C`, the fault is `delta`**, and a zero-input window will confirm it in one measurement.

7.7 7 Design specifications

The requirements below are numbered and checkable. They are grouped by the section that derives them. Their status is normative for any design record produced with this framework; nothing in them is an empirical claim.

7.7.1 7.1 Requirements from the contraction

S1.1 — Loss must be positive and named. Every modelled subsystem shall have `delta > 0`, and the mechanism of loss shall be identified by name (turnover, forgetting, wear, drift, deprecation). A subsystem for which no loss mechanism can be named has not been modelled; it has been assumed.

S1.2 — The fixed point shall be computed before the system is built. For the intended operating parameters, compute `Q_inf = (eta F(u*) − c*)/delta` and record it in the design record. If `eta F(u*) < c*`, the design is rejected at this line and no further work proceeds until a parameter changes.

S1.3 — The uniform bound shall be quoted. Record `eta ln N / delta` as the ceiling on stored order (Theorem 3.3). Any commitment that exceeds it is void.

S1.4 — Convergence time shall be stated, not assumed. Record the number of periods to reach 95% and 99% of `Q_inf` from the actual starting stock. Formulae and tables are in Section 5. A programme whose sponsor's patience is shorter than its `t_95` is misdesigned regardless of its fixed point.

S1.5 — Admissibility shall be checked each period. Verify `C_t <= (1−delta) Q_t + eta S_t`. A violation is escalated as a model-breakdown event, not clipped silently.

S1.6 — Headroom shall be measured before any increment is promised. Compute `h_t = eta ln N − delta Q_t − C_t`. Any requested increment `r_t` shall satisfy `r_t <= h_t` (Corollary 4.2). Promises exceeding `h_t` are not ambitious; they are arithmetically unavailable.

S1.7 — The order proxy shall be justified as distinct from the input. Written justification that `Y` is not a repackaging of `U` [Curzi, 2026i] shall exist before any fitting is performed.

S1.8 — Postulates shall be badged in the sentence that states them. Any application of the storage law to a social or physical system is a modelling decision and shall be marked as such at the point of use, not in a footnote.

S1.9 — Every number shall be computed or cited. No figure appears in a design record without either the arithmetic that produced it or the source that reported it. Where a figure would require research not yet done, the record states what would need measuring instead.

S1.10 — Reversibility. Any intervention shall have a defined stop condition and a defined return path. A change that cannot be undone is a change to `delta` as well as to whatever it was aimed at, because it becomes permanent structure.

S1.11 — Declining shall remain available. No specification in this framework is satisfied by a mechanism that removes a participant's ability to opt out; this is the declination constraint, and it is a constraint on the space of admissible designs rather than a review gate.

7.7.2 7.2 Requirements from the circuit identity

S2.1 — Declare the period. Every parameter shall be stated with its period (`delta = 0.08 per week`). Conversions between periods shall use `delta_m = 1 − (1−delta)^m` and shall be shown in the design record.

S2.2 — Choose the period so that `delta` lands in `[0.02, 0.30]`. Record the justification. If the natural observation cadence forces `delta` outside this band, state the consequence for measurement precision explicitly.

S2.3 — Publish the pole, the time constant, and three settling times. For every modelled subsystem the design record shall contain `1 − delta`, `tau_exact = −1/ln(1−delta)`, `t_50`, `t_95`, and

t_99, computed from t_p = ln(1−p)/ln(1−delta).

S2.4 — Schedule reviews at or after 3 tau. No initiative shall be evaluated for effect before 3 tau periods have elapsed from the change. Interim checkpoints may verify *inputs* (S, eta, C) but shall not be used to judge Q.

S2.5 — State the DC gain. Record eta/delta, the units of stored order bought per sustained unit of surplus per period, and Q_inf = (eta F(u*) − c*)/delta. Where eta F(u*) < c*, record the negative equilibrium and stop; the regime is unsustainable or the linear model must be replaced (Theorem 3.5).

S2.6 — Size every burst to the ceiling. No scheduled session shall be planned to deliver more than eta ln N of input in one period. Planned excess shall be re-scheduled into additional periods, not compressed.

S2.7 — Specify cadence relative to tau. Ritual interval P shall satisfy P <= tau/2 unless a ripple analysis is recorded showing that Q_min = Q_max (1−delta)^P remains above the system's stated functional threshold.

S2.8 — Record the trough, not only the mean. Any cadence proposal shall report Q_min, Q_max, and ripple, computed from Q_max = A/(1−(1−delta)^P).

S2.9 — Match instrument to plant. Reporting for a stock quantity shall present levels; reporting for a flow quantity shall present rates. A difference-only dashboard is not an acceptable instrument for Q.

S2.10 — Bound the control loop. Where a management loop adjusts inputs on the basis of measured Q, record the measurement lag L in periods and the correction gain k. Require L <= tau/4, and reduce k if L cannot be reduced.

S2.11 — Model-breakdown escalation. If a fitted pole leaves [0,1], or admissibility C_t <= (1−delta)Q_t + eta S_t is violated, the event shall be escalated as a specification failure and re-modelled. It shall not be clipped, smoothed, or averaged away.

7.7.3 7.3 Requirements from the loss analysis

S3.1 — delta shall be measured, not assumed. Every design record shall contain a delta estimate obtained from a documented zero-input window or from a frozen, held-out fit [Curzi, 2026i]. Orientation bands (§3.11) are not measurements and shall not be cited as such.

S3.2 — State the period. delta is reported as delta = X per <period>, always. Conversions use delta_m = 1−(1−delta_1)^m and are shown.

S3.3 — Decompose into named channels. The estimate shall be accompanied by an enumeration of loss channels with an individual estimate for each, combined by delta_total = 1 − prod(1−delta_i). The residual between the combined channel estimate and the measured total shall be reported. A residual above 25% of the measured total means the enumeration is incomplete.

S3.4 — Report the half-life. Every delta shall be reported alongside t_half = ln 2 / (−ln(1−delta)) in the same line.

S3.5 — Justify the proxy. Written justification that the order proxy Y is not a repackaging of the input U [Curzi, 2026i] shall precede any delta estimation.

S3.6 — Report phi. The person-borne fraction of stored order shall be estimated and reported. Turnover contribution is recorded as delta_turnover = p * phi with both factors stated separately.

S3.7 — Treat reorganisation as a delta event. Any structural reorganisation shall be preceded and followed by a "who would you ask about X" network sample over a fixed question set, and the fraction of changed or lost answers recorded as a one-off delta contribution.

S3.8 — Every component has an intended lifetime and a target delta. The design record contains a delta portfolio table (§3.13). **No component is assigned delta = 0.**

S3.9 — State both prices of a delta reduction. Any proposal to lower delta shall state, in the same section: the added C per period, and the new t_95 = ln(0.05)/ln(1−delta_new).

S3.10 — Executable over prose where available. Where stored order can be expressed as a check that fails when it becomes false, it shall be, and the prose version shall be derived from the check rather than maintained in parallel.

S3.11 — Minimum three holders. No practice on which the system depends shall have fewer than three trained holders. Where fewer, the shortfall is recorded as an explicit `phi` risk with its `delta_turnover` contribution computed.

S3.12 — Maintenance shall be reported as a level. `delta` and the counterfactual maintenance load `delta Q_t` shall appear in routine reporting as levels, not as incident counts or period-over-period differences (S2.9 above; Star and Ruhleder, 1996).

S3.13 — Scheduled expiry for the high-`delta` portfolio. Components with target `delta >= 0.20` shall carry an explicit expiry date, after which continuation requires an affirmative decision. Silence retires them.

S3.14 — Escalate model breakdown. A fitted `a = 1-delta` outside `[0,1]`, or a decay window whose residuals are structured rather than random, shall be escalated as a specification failure and re-modelled. It shall not be clipped or smoothed.

7.8 8 Failure modes

Each entry gives a clinical description, the parameter damaged, the diagnostic signature, the measurement that detects it, and the fix. They are the named faults for which the three analyses above are responsible.

7.8.1 8.1 Faults of the recursion

THE UNTERMINATED REGRESS — the system reruns the same deliberation without converting any of it, and each pass restates the previous pass rather than updating a single stored position.

- *Parameter damaged:* `eta` (to zero) and effectively `delta = 0` on the deliberation record.
- *Diagnostic signature:* meeting `n+1` opens by re-establishing what meeting `n` established. Agenda items reappear with the same wording. Participants can recall the discussion but cannot point to an artifact.
- *Measurement:* count decisions recorded per hour of deliberation. Count agenda-item recurrences across consecutive meetings. Both near zero and high, respectively.
- *Fix:* force a single mutable stored state — one document, one decision log, one register — that each pass must update rather than restate. This converts the stack into an accumulator by construction. The conversion parameter is treated in Section 6.9, and the stigmergic form of the same move — putting the state in the medium — is developed in the corpus's paper on order without a centre.

THE RUMINATION LOOP — high-surplus interaction with near-zero conversion and nonzero cost; stock falls while the system feels maximally engaged.

- *Parameter damaged:* `eta`.
- *Diagnostic signature:* `S` is high (real disagreement, real analysis, real expertise in the room) and `Q` is falling. Participants report the sessions as valuable. Nothing downstream changes.
- *Measurement:* `eta = delta_Y_attributable / S_input` over several periods. Compare `eta S` against `C` directly. If `eta S < C`, Theorem 3.5 gives a negative equilibrium and the diagnosis is settled.
- *Fix:* raise `eta`, never `S`. Smallest effective moves: a written decision at the end of each session; a named owner; a changed default; a trained second person (Section 6.9).

THE SEALED VAULT — `delta` near zero with no live input; excellent preservation, zero adaptation, slow linear decline through cost alone.

- *Parameter damaged:* `S` (to zero); `delta` suppressed below useful.
- *Diagnostic signature:* the system's answer to a novel question is a citation to an old decision. Procedures whose original rationale nobody can state are still enforced. Founding errors are still load-bearing.
- *Measurement:* age distribution of the active rule set. Fraction of stored order whose rationale can be reconstructed by a current member. Rate of rule retirement (usually exactly zero).
- *Fix:* introduce input (`S > 0`) and introduce *selective* loss. Sunset clauses, review cycles, and scheduled retirement are `delta` restored deliberately on the parts that should decay. Sections 6.7 and 6.8.

THE UNBOUNDED LEDGER — `delta = 0` with continuing input; stock grows linearly and forever, with no fixed point and no error decay. Every rule ever written is still in force.

- *Parameter damaged:* `delta` (to zero), with `C` rising as a consequence.
- *Diagnostic signature:* rule count, policy count, or documentation volume rises monotonically and has never fallen. Compliance cost rises with no corresponding rise in capability. Nobody can state the current state of practice without reading everything.
- *Measurement:* plot cumulative rule count against time; look for any decrease. Measure onboarding time to competence — it grows with `Q` when `delta = 0`.
- *Fix:* institute retirement. The `delta = 0` system has no attractor; giving it one requires removing structure on purpose. Section 6.8 covers deliberate shedding.

THE FORGOTTEN COST — the model is specified with `S`, `eta`, and `delta`, and `C` is left out because nobody owns it. The design predicts growth and the system declines.

- *Parameter damaged:* `C` (unmeasured, therefore unbudgeted).
- *Diagnostic signature:* predicted `Q_inf` far exceeds observed steady state. Predicted growth, observed flatness.
- *Measurement:* tally every recurring draw on the system, including unbilled ones — coordination, context switching, commute, compliance, emotional labour. The full enumeration is the cost tally of the survey protocol [Curzi, 2026i].
- *Fix:* re-fit with `C` included and re-check `eta F(u*) > c*`. Frequently the correct action is to cut `C` rather than to add `S`, because `C` enters linearly and `S` enters through a logarithm.

THE UNBADGED POSTULATE — a modelling assumption is presented with the authority of a theorem.

- *Parameter damaged:* none directly; it damages the warrant of every number downstream.
- *Diagnostic signature:* a claim of the form "the law shows that this room will…" where the law shows nothing of the kind without an unstated mapping from room to model.
- *Measurement:* audit every quantitative claim in the design record against the framework's badge rule. Anything not traceable to a derivation is a postulate and must carry its badge.
- *Fix:* restate with the badge in the same sentence. **The mysticism is permitted; it is permitted in the open, with its badge on.**

7.8.2 8.2 Faults of timing and instrumentation

THE PREMATURE REVIEW — a first-order system is evaluated before `3 tau` and judged on the fraction of its step response that has completed.

- *Parameter damaged:* none. The system is fine; the instrument's timing is wrong.
- *Diagnostic signature:* results at review are between 50% and 80% of target and rising; the language used is "slower than hoped" and "losing momentum."
- *Measurement:* compute `t_95 = ln(0.05)/ln(1-delta)` and compare against the elapsed time. If elapsed `< 3 tau`, the review is uninformative about `Q` by construction.
- *Fix:* move the review to `3 tau`. Use interim checkpoints on inputs only. Where the sponsor's patience is genuinely shorter than `3 tau`, the honest response is to raise `delta`'s counterpart — pick a subsystem with a shorter time constant — not to promise faster arrival.

THE INTENSITY FALLACY — a large one-off input is substituted for a sustained one on the grounds that the total is the same.

- *Parameter damaged:* effectively `S` (clipped by the ceiling) while `C` is unaffected.
- *Diagnostic signature:* a peak two to four periods after the event, decay thereafter, and a post-mortem that concludes "we failed to sustain the energy."
- *Measurement:* compare planned session input against `eta ln N`. Compute delivered total across the horizon versus the sustained alternative.
- *Fix:* redistribute. Size sessions to `eta ln N` and increase their frequency until the planned total is deliverable. If the total is not deliverable at any feasible frequency, the total was never available and the plan should say so.

THE SPARSE RITUAL — a ceremony at an interval much longer than `tau`, whose troughs fall below the system's functional threshold even though its mean looks acceptable.

- *Parameter damaged:* none directly; the ripple is the fault.
- *Diagnostic signature:* the ceremony is well-attended and well-regarded, and between occurrences the system behaves as if it does not exist. Participants describe "getting back up to speed."
- *Measurement:* compute `Q_min = Q_max (1-delta)^P` and compare against the threshold below which the system is observed to malfunction. Also compute the delivered total against the ceiling.
- *Fix:* shorten `P` toward `tau/2` or below. Note that halving `P` does not require halving session length: the ceiling caps per-session delivery anyway, so two half-length sessions frequently deliver more than one full-length one.

THE HIGH-PASS INSTRUMENT — the system integrates but the reporting differentiates, so sustained level is invisible and only change is reported.

- *Parameter damaged:* none in the plant; the measurement of `Q` is systematically wrong.
- *Diagnostic signature:* reports consist entirely of week-over-week or quarter-over-quarter changes. A stable, functioning capability generates no report and therefore no resourcing. Maintenance is chronically unfunded (compare Star and Ruhleder, 1996: infrastructure becomes visible only on breakdown; Star and Ruhleder, 1996).
- *Measurement:* audit the reporting pack. Count the fraction of metrics presented as levels versus differences. Ask what report a perfectly stable, perfectly functioning subsystem would generate; if the answer is "none," the instrument is high-pass.

- *Fix:* add level reporting for every stock. The institutional version of this fault is treated at length in the corpus's papers on institutional maintenance.

THE OVERDAMPED INSTITUTION — `delta` driven so low that `tau` exceeds the tenure of everyone who could act on a measurement.

- *Parameter damaged:* `delta` (too low).
- *Diagnostic signature:* every proposed change is met with "we tried that" referring to an attempt whose evidence nobody can reconstruct. `t_95` exceeds the average tenure of the decision-makers.
- *Measurement:* compute `t_95` and compare against median tenure in the roles that would act on the result. If `t_95 > tenure`, the institution cannot learn from its own experiments.
- *Fix:* either raise `delta` on the parts of `Q` that should turn over (Section 6.8 on deliberate forgetting), or restructure the experiment onto a faster subsystem with a shorter `tau`. Do not attempt to speed up a slow pole by adding input; the pole sets the rate.

THE DELAYED CONTROLLER — corrections applied on the basis of measurements that are several periods stale, producing overshoot and, at higher gain, oscillation.

- *Parameter damaged:* the closed-loop dynamics; the open-loop parameters may be perfectly healthy.
- *Diagnostic signature:* alternating over- and under-correction on roughly a `4L`-period cycle. Hiring surges followed by freezes. Process added then stripped then re-added.
- *Measurement:* record `L` (periods between the state occurring and the decision responding to it) and `k` (fraction of the observed gap corrected per period). Compare `L` against `tau/4`.
- *Fix:* shorten the reporting lag first; reduce the gain second. At `L = 4` and `k = 0.5` the simulated loop diverges outright, so gain reduction alone is not always sufficient.

THE MISMATCHED PERIOD — parameters measured on one period and applied on another without conversion.

- *Parameter damaged:* all of them, silently.
- *Diagnostic signature:* forecasts wrong by a factor that is a power of the period ratio. A monthly-measured `delta` applied as if annual understates decay by a factor near 12 at small `delta` and far more at large `delta`.
- *Measurement:* check every parameter in the record for a stated period. Any parameter without one is suspect.
- *Fix:* restate everything on a single declared period using `delta_m = 1-(1-delta)^m`. Note that `eta` and `C` also carry periods and must be converted; only `u` and `N` are period-free.

7.8.3 8.3 Faults of retention

THE PERSON-SHAPED REPOSITORY — a substantial fraction of stored order lives in one individual and nowhere else; `phi` is near 1 for that component.

- *Parameter damaged:* `delta` via the turnover channel; `C` via onboarding.
- *Diagnostic signature:* a name appears as the answer to "who would know?" for many unrelated questions. Their absence stops work in areas unrelated to their nominal role. They cannot take leave.

- *Measurement:* sample the "who would you ask about X" network across a fixed question set and compute the concentration. Compute `delta_turnover = p * phi` for the affected component.
- *Fix:* train to three holders (S3.11), then move the content down the artifact ladder. Note that the same actions raise `eta`: the conversion mechanisms and the retention mechanisms of Section 6 are the same mechanisms.

THE ROTTING DOCUMENT — the artifact is intact, complete, well-formatted, and false. Nothing about its appearance indicates this.

- *Parameter damaged:* `delta` via documentation entropy and environment drift, both undetected.
- *Diagnostic signature:* documents are never edited and never deleted. When someone follows one, it fails at step four, and the response is to ask a person rather than to fix the document.
- *Measurement:* select documents at random and have someone who did not write them execute them. Record the failure rate and the step at which failure occurs. That rate, per period, is a direct `delta` estimate for the documentation channel.
- *Fix:* make it executable (S3.10). Where that is impossible, date it and expire it (S3.13): a document with a visible staleness marker is honest about its own `delta`.

THE REORG AMNESIA — a structural change invalidates the part of `Q` encoded in the communication topology, and nobody accounts for it because no document was lost.

- *Parameter damaged:* `delta`, in a single large step rather than a rate.
- *Diagnostic signature:* post-reorganisation, routine requests take substantially longer and get routed wrongly. Everyone's documents are intact; the mechanism is Conway's observation that a system mirrors the communication structure that built it [Conway, 1968].
- *Measurement:* the before/after network sample of S3.7. The fraction of changed or lost answers is the estimate.
- *Fix:* budget the loss in advance and pay it deliberately — publish the new routing, run the "who would you ask" survey and publish its answers as an artifact, and expect a recovery period of roughly `t_95` at the prevailing `delta`. Do not model the step as a rate; §3.7's cliff pitfall applies.

THE IMMORTAL RULE — a procedure whose rationale has decayed while the procedure has not. It cannot be retired because nobody can establish that retiring it is safe.

- *Parameter damaged:* `delta` far too low on a component that should have been high; `C` permanently elevated.
- *Diagnostic signature:* asked why the rule exists, everyone gives a different answer or cites a person who has left. The rule survives every review because removing it requires proving a negative.
- *Measurement:* for each rule in the active set, record whether any current member can state its originating condition. The fraction that cannot is the context-loss share of `Q`, and it is usually startling.
- *Fix:* sunset clauses applied retroactively — set an expiry date on the existing rule set and require affirmative renewal (S3.13). This converts "prove it is safe to remove" into "state why it should stay," which is the tractable direction.

THE PRESERVED GRUDGE — an interpersonal record maintained at near-zero `delta` in a component whose target `delta` should be the highest in the portfolio.

- *Parameter damaged:* `delta` too low; `u` driven toward 0 within the affected pair, because the interaction becomes fully predictable and defensive; `C` elevated by routing work around the pair.
- *Diagnostic signature:* two people or two groups who are never scheduled together, whose work is routed through an intermediary, and whose original incident predates most of the current membership.
- *Measurement:* count the routing detours. Compare the age of the originating incident against the half-life assigned to this component in the portfolio (Section 6.8 suggests 1.4 months).
- *Fix:* the design lever is the `delta` of the *record*, not the sentiment. Do not retain durable written accounts of interpersonal incidents past their intended lifetime. Where the record has already outlived it, a structured repair with a stated closure is a one-off high-`delta` event. **The declination constraint applies without exception: no mechanism here removes anyone's ability to decline.**

THE INVISIBLE MAINTENANCE — the activity that suppresses `delta` produces non-events, is therefore invisible to difference-based reporting, and is therefore chronically underfunded.

- *Parameter damaged:* `delta` rises as maintenance is cut; the rise is not attributed to the cut because the lag is roughly `t_half`.
- *Diagnostic signature:* maintenance budget falls with no immediate consequence, followed some months later by a cluster of failures treated as unrelated incidents [Star and Ruhleder, 1996].
- *Measurement:* report `delta` and `delta Q_t` as levels (S3.12). Correlate maintenance spend against `delta` with a lag of one half-life.
- *Fix:* instrument the non-event. Scheduled drills, executable checks, and level-based reporting all convert invisible success into visible signal.

THE CONTAMINATED DECAY WINDOW — `delta` is estimated over a window in which input was still arriving, producing an underestimate.

- *Parameter damaged:* the estimate, and therefore every quantity computed from it — `Q_inf` and the Theorem 3.3 ceiling both have `delta` in the denominator.
- *Diagnostic signature:* the fit shows a plateau or an upturn; `R^2` is poor with structured residuals; per-step estimates trend toward zero across the window.
- *Measurement:* plot `ln Y_t` against `t` and inspect residuals for structure. Compare the regression estimate against the endpoint two-point estimate; disagreement above roughly 15% indicates non-geometric shape.
- *Fix:* find a cleaner window, shorten the window to its uncontaminated prefix, or fall back to the joint fit [Curzi, 2026i] with its four required baselines.

THE ZERO-DELTA COMPONENT — a component of `Q` assigned no loss at all, deliberately or by neglect.

- *Parameter damaged:* `delta = 0` on that component: no fixed point, no error decay, unbounded linear accumulation (Theorem 3.3, second case).
- *Diagnostic signature:* the component's size has never decreased in its recorded history. Rule count, policy volume, or documentation size is monotone increasing. Onboarding time grows year over year with no corresponding growth in capability.
- *Measurement:* plot cumulative component size against time and look for any decrease. Compute the retirement rate; it is usually exactly zero.
- *Fix:* assign a target `delta` and an intended lifetime (S3.8). Institute scheduled expiry (S3.13). Accept that some structure will be lost that would have been useful — that loss is the price of the error decay and the fixed point, and Section 4.3 computes what they are worth.

7.9 9 Results: three systems worked end to end

The three cases below are **worked examples, not field studies.** Each is a constructed system with declared parameters, carried through the full procedure so that every number can be re-derived by the reader. Their function is to exercise the chain and to exhibit the places where the arithmetic refuses to cooperate — a negative equilibrium, an infeasible increment, a doubled settling time — rather than to supply evidence about any real organisation. No dataset was collected for any of them, and Section 12 records this as a limitation of the paper rather than a caveat on the cases.

7.9.1 9.1 A twelve-person research group

The system. A twelve-person applied research group inside a larger organisation. Period: one week. The question is whether a proposed change to the weekly seminar can deliver the improvement the director has promised upward.

Step 1 — `Q` and its proxy. `Q` is the group's stored methodological order: techniques that any member can execute correctly without the originator present. The proxy `Y` is the count of techniques that pass a monthly spot-check, where the check is performed by a member who did not develop the technique. This proxy is not a repackaging of the input `U`, because `U` describes the *disagreement geometry* across research categories in a given week and `Y` describes *transferable executable capability* measured a month later by a different person [Curzi, 2026i]. That justification is written into the design record before any fitting, per S1.7.

Step 2 — `N` and `u*`. The group spans eight genuinely distinct methodological categories, so `N = 8` and the per-interaction ceiling is `ln 8 = 2.0794`. Disagreement audits and error- decorrelation measurements across category pairs give a working estimate `u* = 0.60`.

```
N_eff = 1 + (N−1) u* = 1 + 7 * 0.60 = 5.20
S = F(u*) = ln(5.20) = 1.6487
```

The roster says twelve people in eight categories. The audited effective count is 5.20. The group is buying eight categories' worth of salary and receiving 5.20 categories' worth of non-redundancy. Here is the full sensitivity, which the director should see:

`u*`	`N_eff = 1 + 7u*`	`S = ln N_eff`	% of ceiling `ln 8`
0.00	1.00	0.0000	0%
0.10	1.70	0.5306	26%
0.25	2.75	1.0116	49%
0.50	4.50	1.5041	72%
0.60	**5.20**	**1.6487**	**79%**

u*	N_eff = 1 + 7u*	S = ln N_eff	% of ceiling ln 8
0.75	6.25	1.8326	88%
0.90	7.30	1.9879	96%
1.00	8.00	2.0794	100%

Note the strict concavity of `f` [Curzi, 2026] doing its work: moving `u*` from 0.00 to 0.25 buys 1.0116 nats; moving it from 0.75 to 1.00 buys 0.2468. **The first quarter of non-parallelism is worth four times the last quarter.** This is the concavity dividend: most of the available surplus is bought by the first departure from parallelism.

Step 3 — `eta`. Measured over eleven weeks as attributable change in `Y` per unit of `S` delivered: `eta = 0.50`. So `eta ln N = 0.5 * 2.0794 = 1.0397` is the absolute per-week input ceiling, and the achieved input is `eta S = 0.5 * 1.6487 = 0.8243`.

Step 4 — `delta`. A four-week shutdown gave a clean zero-input window; log-linear fit on the decay gives `delta = 0.10` per week. Half-life `ln 2 / (−ln 0.9) = 6.58` weeks. Section 6.4 gives the procedure and Section 9.3 gives a fuller version of this fit.

Step 5 — `C`. Tallied at `C = 0.20` per week: seminar preparation, the coordination overhead of scheduling across eight categories, and the write-up burden.

Step 6 — the fixed point. By Theorem 3.5,

```
Q_inf = (eta F(u*) − c*) / delta = (0.8243 − 0.20) / 0.10 = 6.2433
```

and by Theorem 3.3 the uniform ceiling is `eta ln N / delta = 1.0397 / 0.10 = 10.397`. **The group converges to 6.2433 and cannot under any interaction geometry exceed 10.397.**

Step 7 — convergence. Starting from `Q_0 = 0`:

week	Q_t	% of Q_inf
0	0.0000	0.00%
1	0.6243	10.00%
2	1.1862	19.00%
4	2.1471	34.39%
6	2.9254	46.86%
8	3.5558	56.95%
10	4.0664	65.13%
12	4.4800	71.76%
15	4.9579	79.41%
20	5.4843	87.84%
25	5.7951	92.82%
29	5.9492	95.29%
30	5.9786	95.76%

Ninety-five percent of the destination is reached in week 29; the exact figure from `t_95 = ln(0.05)/ln(0.90)` is 28.43 weeks. Ninety-nine percent takes `ln(0.01)/ln(0.90) = 43.71` weeks. **The director's promised review is at week 12, when the system will be at 71.76% of its own steady state and still climbing.** That is not underperformance. It is the step response of a first-order system evaluated before it settles, and Section 5 exists so that this mistake is not made in a design review.

Step 8 — headroom, and what may be promised. With `eta ln N = 1.0397`:

Q_t	h_t = 1.0397 − 0.10 Q_t − 0.20	regime
0.0000	0.8397	growth
2.0000	0.6397	growth
4.0000	0.4397	growth
6.2433	0.2154	growth (at Q_inf, still positive)
8.0000	0.0397	growth, nearly exhausted
8.3972	0.0000	**maintenance only**
10.0000	-0.1603	**unsustainable — no geometry saves it**

The maximum single-week increment the group can be promised at its current stock of 4.48 is r_max = 1.0397 − 0.448 − 0.20 = 0.3917 (Corollary 4.2). By Proposition 4.1, achieving an increment r_t requires S_t >= (delta Q_t + C_t + r_t)/eta, feasible only if delta Q_t + C_t + r_t <= eta ln N. At Q_t = 4.48 and r_t = 0.50, the requirement is S_t >= (0.448 + 0.20 + 0.50)/0.50 = 2.296, which exceeds ln 8 = 2.0794. **The requested increment is infeasible, and the proof is one line of arithmetic.** The correct report upward is not "we will try"; it is "0.50 per week is above r_max = 0.3917, here is the calculation, and here are the three parameters that would have to move."

Step 9 — the counterfactuals, computed.

If delta were 0. The group would grow linearly at 0.6243 per week without limit — 6.243 at week 10, 12.487 at week 20, 62.433 at week 100 — with no fixed point, no error decay, and no answer to "what will this settle at?" It would also never lose a single obsolete technique, and its spot-check burden would grow without bound, which shows up as a rising C in year three. The group's methodological register would become an archive that must be read in full before anyone can act. **The absence of loss does not produce a better group; it produces a group that cannot say what it currently does.**

If eta fell to 0.05 — the seminar continues, the arguments are as good as ever, and nobody writes anything down. Then eta S = 0.0824 < C = 0.20, Q_inf = −2.676, and the stock falls from 6.2433 to below the admissibility floor within twelve weeks. The seminar would feel *better* during this period, because S is unchanged and the discussions are as rich as they ever were.

If S fell to 0 and delta were driven to 0 — the sealed vault. Q falls linearly at C = 0.20 per week and reaches zero at week 31.2, with the group defending techniques whose justification nobody present can reconstruct.

Step 10 — the intervention, and its falsification criterion. The binding constraint is not u*, which at 0.60 already delivers 79% of the ceiling and whose remaining headroom is worth only 0.43 nats at enormous recruiting cost. It is not N, because raising N from 8 to 16 buys exactly ln 2 = 0.6931 extra ceiling for a doubling of genuinely distinct methodological categories — see the corpus's analysis of the logarithmic ceiling. The binding constraint is the pair (eta, C), and of the two, eta is cheaper to move (Section 6.9).

The intervention: every seminar closes with one written, owned, dated decision, and one named second person who can execute the technique discussed. Predicted effect: eta from 0.50 to 0.60, C unchanged. Then eta S = 0.9892, Q_inf = (0.9892 − 0.20)/0.10 = 7.892, a 26.4% increase in steady-state stored order, reached to 95% in another 28.4 weeks from the point of change.

Preregistered falsification [Curzi, 2026i]: fit Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + eps_t on weeks 1-30, freeze, and predict weeks 31-60. The intervention is falsified if the fitted a does not lie in [0,1], if b does not increase, or if the frozen model fails to beat all four required baselines on held-out data — persistence, direct-input, matched autoregressive, and the shuffled-input control. **The shuffled-input control is the one most often omitted and the one most often decisive**, because it is the only baseline that tests whether the *timing* of interaction matters or merely its volume.

7.9.2 9.2 A six-category production team choosing its cadence

The system. A production team of nineteen people spanning six genuinely distinct functional categories: design, fabrication, quality, logistics, maintenance, and customer support. Period: one week. The question is whether to replace the current monthly all-hands “deep dive” with a weekly thirty-minute cross-function standup, over the objection that the monthly session is “much more substantial.”

Measured parameters.

parameter	value	how obtained
`N`	6	enumeration of functional categories
`u*`	0.50	disagreement audit and error-decorrelation across category pairs
`eta`	0.40	attributable change in the order proxy per unit surplus, over 14 weeks
`delta`	0.08 per week	log-linear fit on a 5-week zero-input window during a plant shutdown (Section 6.4)
`C`	0.15 per week	tallied coordination, scheduling, and write-up overhead

Derived quantities.

```
ln N        = ln 6                  = 1.7918     hard per-interaction ceiling
N_eff       = 1 + 5 * 0.50          = 3.50       audited effective count
S           = ln 3.50               = 1.2528     69.9% of ceiling
eta S       = 0.4 * 1.2528          = 0.5011     delivered input per week
eta ln N    = 0.4 * 1.7918          = 0.7167     per-week input CEILING
net drive   = 0.5011 - 0.15         = 0.3511     per week
Q_inf       = 0.3511 / 0.08         = 4.3888     Theorem 3.5
DC gain     = eta/delta = 0.40/0.08 = 5.0        stock per sustained unit of surplus
Thm 3.3 cap = eta ln N / delta      = 8.9588     absolute ceiling on Q
pole        = 0.92
tau_exact   = -1/ln(0.92)           = 11.99 weeks
t_half      = 0.6931 * 11.99        = 8.31 weeks
t_95        = ln(0.05)/ln(0.92)     = 35.93 weeks
t_99        = ln(0.01)/ln(0.92)     = 55.23 weeks
```

The nineteen-person roster contains six categories and buys 3.50 effective ones. The team achieves 69.9% of its theoretical per-interaction ceiling. That is a respectable `u*`: the remaining 30% costs more to buy than it returns here, and the binding constraint is elsewhere.

Step 1 — Is the monthly cadence even deliverable?

A monthly session sized to carry four weeks of surplus would need to deliver `4 * eta S = 4 * 0.5011 = 2.0044` in one week. The ceiling is `eta ln N = 0.7167`. **The monthly session is over the ceiling by a factor of 2.80.** No facilitation technique changes this: the range result caps `F(u)` at `ln N = 1.7918` and no interaction geometry exceeds it.

So the monthly session actually delivers 0.7167, and 1.288 units of intended input per month — 64.3% of it — do not exist.

Step 2 — Annual delivery, both regimes.

	sessions per year	delivered per session	annual total delivered	cost paid
weekly	52	0.5011	26.0575	52 weeks
monthly (4-weekly)	13	0.7167 (clipped)	9.3171	52 weeks

The weekly cadence delivers **2.80 times** the annual input, and the two regimes pay identical annual cost. The ratio is exactly the clipping factor, which is not a coincidence: it is `4 eta S / eta ln N = 4S/ln N`.

Step 3 — Trajectories.

week	weekly cadence	4-weekly cadence
0	0.0000	0.0000
1	0.3511	0.5667
2	0.6741	0.3714
3	0.9713	0.1917
4	1.2447	0.0263
8	2.1364	0.0452
12	2.7752	0.0587
26	3.8867	0.4393
52	4.3314	0.0916

Steady state: weekly settles at 4.3888. Four-weekly oscillates between 0.0928 and 0.6521 with a mean of 0.3647 — **8.3% of the weekly regime.**

Observe week 1. **The monthly cadence is ahead**, 0.5667 against 0.3511, because a single session at the ceiling puts more in than a single ordinary week. That is the entire empirical basis for the objection that the monthly deep dive is "more substantial," and it is correct — for one week. By week 2 the weekly regime is ahead and it never trails again.

Step 4 — Ripple analysis for the sparse regime, per S2.8.

For `P = 4`, `a = 0.92`, `A = 0.7167`:

```
a^P    = 0.92^4                       = 0.7164
Q_max  = A/(1 - a^P) = 0.7167/0.2836  = 2.5271   (if cost were zero)
Q_min  = Q_max * a^P                  = 1.8104   (if cost were zero)
```

With cost at 0.15 per week the simulated cycle is far lower — 0.0928 to 0.6521 — because the weekly cost draw consumes most of what the burst delivers before the next burst arrives. **The ripple formula is the no-cost idealisation; the real trough is set by the cost drawn between bursts, and that is the number that must be reported.**

Step 5 — Filter check, per §2.12. At `delta = 0.08`, an input oscillating with period 4 weeks is transmitted at roughly 6% of its amplitude (interpolating between the `delta = 0.05` value of 0.0362 and the `delta = 0.10` value of 0.0743). The team's stored order is essentially blind to the monthly rhythm as *rhythm*; what reaches `Q` is the mean, and the mean is what clipping has destroyed.

Step 6 — Review scheduling, per S2.4. `3 tau = 35.98` weeks. The correct evaluation point for the weekly-cadence change is **week 36**, when the system will be at 95.03% of `Q_inf`. A review at week 12 will find 63.2%, and a review at week 26 will find 88.6%. Both are on the curve. Neither is evidence of a problem. The design record states all three numbers in advance so that nobody has to argue about it later.

Step 7 — Headroom check, per Section 3.2. At the new steady state `Q_inf = 4.3888`:

```
h = eta ln N − delta Q − C = 0.7167 − 0.08 ∗ 4.3888 − 0.15 = 0.2155
```

Positive, so growth remains available, and `r_max = 0.2155` per week is the largest increment that may be promised (Corollary 4.2). Headroom reaches zero at `Q = (0.7167 − 0.15)/0.08 = 7.0838`, safely above the operating point.

At the *old* monthly cadence with `Q` near 0.36, headroom was `0.7167 − 0.08(0.36) − 0.15 = 0.5379` — apparently *higher*. This is the trap the headroom gauge names: **headroom is high in a system that has stored nothing, and that is not a sign of health.** A team with no accumulated order has enormous room to build and nothing built. Headroom is a budget, not a score.

Step 8 — Recommendation and falsification.

Adopt the weekly cadence. Retain a quarterly longer session for the genuinely serial work that cannot be decomposed into weekly increments, sized to the ceiling and understood to contribute `0.7167` in its week rather than a quarter's worth.

Predicted: `Q` rises from a mean of 0.365 to 4.389, reaching 95% of the new value at week 36.

Preregistered falsification [Curzi, 2026i]: fit `Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + eps_t` on weeks 1-36 with `0 <= a <= 1` and `b >= 0`, freeze the coefficients, and report on weeks 37-72 against four baselines — persistence, direct-input, matched autoregressive, and shuffled-input control. The prediction is falsified if:

1. the fitted `a` falls outside `[0,1]` (the first-order form is wrong);
2. the fitted `b` does not exceed its pre-intervention value (the cadence change did not raise delivered input);
3. the frozen model fails to beat the **shuffled-input control** on held-out data — which would mean that the *timing* of interaction carries no information beyond its volume, and the entire cadence argument in this chapter would not apply to this system.

The shuffled-input control is the decisive test for this chapter's thesis. Every claim made here about cadence is a claim that *when* surplus arrives matters, not merely how much. If shuffling the input sequence does not degrade the fit, the low-pass structure is not present in this system and the design rules of Section 5 do not apply to it. Curzi [2026i] explains why this control is the one most often omitted, and measures the three distinct ways it fails.

7.9.3 9.3 A twenty-two-person infrastructure team re-engineering its `delta`

The system. Twenty-two engineers responsible for the operational infrastructure of a mid-sized service. Period: one month.

Step 1 — `Q` and its proxy `Y`. `Q` is operational stored order: the team's capacity to execute correct recovery, deployment, and diagnostic procedures without the original author present. The proxy `Y` is **the number of runbooks that pass a monthly spot-audit, where the audit consists of a randomly selected engineer executing the runbook end to end in a staging environment, and a pass requires completion without consulting the author.**

Proxy justification, per S3.5 and the order-proxy condition [Curzi, 2026i]. The input observable `U` is the disagreement geometry across the team's seven operational categories in a given month — measured by error-decorrelation on incident post-mortems. The order observable `Y` is transferable executable capability, measured in a different environment, by a different person, at a later time, against a fixed procedure. `Y` cannot be raised by having more discussion; it can only be raised by producing or repairing a runbook and training someone to run it. **`Y` is not a repackaging of `U`.** This paragraph is in the design record before any fitting is performed.

Step 2 — The zero-input window. A six-month platform freeze during a compliance migration. During the freeze no new runbook work was funded, and the audit continued. This is a genuine zero-input window on the runbook component of `Q`.

month	0	1	2	3	4	5	6
`Y` (runbooks passing audit)	100	92	84	79	71	66	60

Step 3 — The fit. Log-linear regression of `ln Y` on `t`:

```
slope    = −0.08446
delta    = 1 − exp(−0.08446)  = 0.0810 per month
Y_0 fit  = 100.16
R^2      = 0.9980
t_half   = ln 2 / (−ln 0.9190) = 8.21 months
```

Cross-check by endpoints: `1 − (60/100)^{1/6} = 0.0816`, agreeing with the regression to 0.8%. The per-step estimates range from 0.0595 to 0.1013 — a reminder that no single step is an estimate. Residuals are unstructured and `R^2` is 0.998, so the exponential form is supported for this proxy over this range.

In plain language for the sponsor: half of everything this team knows how to do that is not actively exercised is gone in eight months.

Step 4 — Channel decomposition, per S3.3. Monthly attrition is 1.4% (an annual rate of 15.6%), and an audit of runbook authorship and the "who would you ask" network puts the person-borne fraction at 0.55.

channel	mechanism	`delta_i`
turnover	`p * phi = 0.014 * 0.55`	0.0077
platform / tool churn	dependencies deprecate under the team	0.0200
environment drift	runbook references infrastructure that changed	0.0350
procedural forgetting	steps not executed between incidents	0.0200
naive sum		**0.0827**
combined `1 − prod(1−delta_i)`		**0.0804**
measured (regression)		**0.0810**
residual unexplained		**0.0006**

The residual is 0.0006, or 0.7% of the measured total — well inside S3.3's 25% threshold. **The enumeration is complete enough to act on**, and the largest channel is environment drift at 0.0350, 43% of the total.

Step 5 — The rest of the parameters. For completeness, since `delta` alone does not give a fixed point:

parameter	value	source
`N`	7 operational categories	enumeration
`u*`	0.45	error-decorrelation across post-mortems
`eta`	9.0 runbooks per nat per month	attributable `delta_Y` per unit `S` over 14 months
`C`	4.0 runbooks per month	deprecations, audit failures, forced rewrites

```
N_eff    = 1 + 6 * 0.45      = 3.70
S        = ln 3.70           = 1.3083 nats      (67.2% of ln 7 = 1.9459)
eta S    = 9.0 * 1.3083      = 11.7750 runbooks/month
eta ln N = 9.0 * 1.9459      = 17.5132 runbooks/month   (the per-month ceiling)
```

Step 6 — Model consistency check. Theorem 3.5 predicts

```
Y_inf = (eta S - C)/delta = (11.7750 - 4.0)/0.0810 = 96.00 runbooks
```

The observed pre-freeze steady state was 100. **The model predicts 96.00 against an observed 100 — a 4.0% discrepancy, which is well within the uncertainty of the parameter estimates and is reported as a consistency check, not as a validation.** Validation requires the frozen held-out procedure of the preregistered protocol [Curzi, 2026i], which is Step 10.

Step 7 — The intervention. Attack the largest channel first (§3.6), and use the artifact ladder (§3.11).

channel	before	intervention	after
turnover	0.0077	pair rotation on-call; three holders per runbook (S3.11) drops `phi` from 0.55 to 0.30	0.0042
platform / tool churn	0.0200	no change proposed — outside the team's control this cycle	0.0200
environment drift	0.0350	**runbooks executed monthly in CI against staging: prose becomes an executable check that fails when it becomes false (S3.10)**	0.0050
procedural forgetting	0.0200	quarterly exercise; each engineer executes two unfamiliar runbooks live	0.0080
combined	**0.0804**		**0.0368**

```
new delta   = 0.0368 per month
new t_half  = ln 2 / (-ln 0.9632) = 18.49 months     (was 8.21)
```

Step 8 — What that buys, and what it costs, per S3.9.

```
Y_inf before = (11.7750 - 4.0)/0.0810 = 96.00 runbooks
Y_inf after  = (11.7750 - 4.0)/0.0368 = 211.28 runbooks     (+120.1%)

t_95 before  = ln(0.05)/ln(0.9190) = 35.5 months
t_95 after   = ln(0.05)/ln(0.9632) = 79.9 months            (+125.1%)
```

The team more than doubles its sustainable stock of executable capability, and more than doubles the time it takes to converge on any new steady state. Both numbers go in the proposal. The second one is the one that will not be volunteered by anyone else in the room.

Is that trade acceptable here? The relevant comparison is `t_95` against the rate at which the underlying platform changes. If the platform's own generation is shorter than 80 months, the team will be converging toward a target that has already moved, and the correct design is a lower target `Y_inf` with a faster response. In this case the platform generation is estimated at roughly 5 to 7 years, so 80 months is at the edge of acceptable and the design record says so explicitly rather than hiding it.

The added `C`: CI execution of runbooks costs compute and maintenance of the staging environment; quarterly exercises cost roughly 2 engineer-days per person per year; pair rotation costs scheduling flexibility. Estimated at 1.2 runbooks-equivalent per month of additional `C`. Re-running with `C = 5.2`:

```
Y_inf after, with added cost = (11.7750 - 5.2)/0.0368 = 178.67 runbooks   (+86.1%)
```

Still a large gain, and **the honest number is 178.67, not 211.28.** Proposals that omit the added `C` overstate the benefit by 18%.

Step 9 — Headroom, per Section 3.2. With `eta ln N = 17.5132` runbooks/month:

```
h = eta ln N - delta Y - C
```

state	`delta`	`Y`	`C`	`h`
before	0.0810	100	4.0	5.4132
after	0.0368	100	5.2	8.6332

and the stock at which headroom reaches zero, `Y = (eta ln N - C)/delta`:

```
before: (17.5132 - 4.0)/0.0810   = 166.83 runbooks
after:  (17.5132 - 5.2)/0.0368   = 334.60 runbooks
```

Before the intervention, the maintenance load `delta Y` at `Y = 100` is 8.10 runbooks per month — more than double the explicit cost `C = 4.0`. The team's largest recurring expense was invisible because it was not billed to anything. After the intervention it is 3.68. That single line is usually the one that carries the proposal, and it is the headroom thesis arriving early: the maintenance term grows with what the system holds, and it is the term nobody puts on a budget.

Step 10 — Preregistration and falsification, per Curzi [2026i].

Fit `Y_{t+1} = a Y_t + b F(U_t) - C_hat_t + eps_t` with `0 <= a <= 1` and `b >= 0` on the 14 months preceding the intervention. Freeze. Predict the 24 months following.

Predicted: fitted `a` rises from `1 - 0.0810 = 0.9190` to `1 - 0.0368 = 0.9632`; `b` unchanged; `Y` rises toward 178.67 with 95% of the movement complete at month 80.

Falsified if:

1. fitted `a` after the intervention is not significantly greater than 0.9190 — the interventions did not reduce loss;
2. fitted `a` falls outside `[0,1]` — the first-order form is wrong for this proxy and the whole model must be re-specified;
3. `b` changes materially — the interventions have confounded `delta` reduction with `eta` increase, which is likely on its face, since executable runbooks and three-holder training plausibly raise conversion too. **This is a real confound and it is stated in advance rather than discovered afterward.** The mitigation is to stage the interventions: CI execution in quarter one, three-holder training in quarter three, with a fit between them.

4. the frozen model fails to beat the **shuffled-input control** on held-out data — which would mean the timing of interaction carries no information beyond its volume, and the storage-law description of this system is not supported regardless of how well `delta` was measured.

The shuffled-input control is the one most often omitted and the one most often decisive. It is the only baseline that can distinguish a real dynamical model from a well-fitted trend.

Step 11 — The `delta` portfolio for this team, per S3.8. The intervention above concerns one component. The full record assigns a target to each:

component	target `delta`/month	half-life	intended lifetime
incident-response protocol	0.010	69.0 months	multi-generation; safety-critical
runbooks	0.037	18.5 months	one platform generation
tooling conventions	0.050	13.5 months	one platform generation
on-call rota practice	0.100	6.6 months	re-decided twice yearly by exercise
quarterly roadmap	0.250	2.4 months	must go stale visibly
post-incident personal attributions	0.400	1.4 months	cleared by design; see THE PRESERVED GRUDGE

No row is zero (R3.3), every row has a stated lifetime (R3.4), and the two high-`delta` rows carry expiry dates so that silence retires them (S3.13).

7.10 10 Discussion

7.10.1 10.1 What follows

The results above are consequences of one coefficient having a value strictly between zero and one. Restated in the order they are used downstream:

The **storage law is a contraction whenever `delta > 0`**, with modulus `k = 1 − delta` and steady state `Q* = (eta S − C)/delta`. This is the result the rest of this corpus leans on, and it is what licenses talking about a system's destination at all.

Theorem 3.5 is the DC gain of a first-order section. Evaluating `1/(z − a)` at `z = 1` gives `1/delta`, so the fixed point is the net drive multiplied by the gain. This is not a restatement; it is what makes the standard toolkit — settling time, frequency response, ripple analysis — available without further argument.

The uniform bound `Q_t <= max{Q_0, eta ln N / delta}` degrades at `delta = 0` to `Q_t <= Q_0 + t eta ln N`, which is linear, unbounded, and has no attractor. The failure is of a kind that no monitoring instrument reporting only levels will detect, because the level keeps rising.

`delta = 0` is a pathology rather than an ideal. The map ceases to contract, the fixed point disappears, and the error-persistence column is a column of ones, so no mistake ever decays out of the stock. An organisation pursuing perfect retention as a goal is pursuing the loss of its own capacity to revise.

The time-constant table gives `tau = 1/delta` against `tau_exact = −1/ln(1−delta)`, agreeing within one per cent at `delta <= 0.02` and about five per cent at `delta = 0.10`, with `3 tau` delivering 95.2 to 96.5 per cent and `5 tau` delivering 99.4 to 99.6 per cent of the step. This is the result that decides when a programme may honestly be reviewed.

The worked research group returns `N = 8`, `u* = 0.60`, `N_eff = 5.20`, and `S = 1.6487` nats — 79 per cent of the ceiling `ln 8 = 2.0794`. With `eta = 0.50`, `delta = 0.10` and `C = 0.20`, `Q_inf = 6.2433` against a hard ceiling of 10.397, and `t_95 = 28.43` weeks, so a review at week 12 finds the system at 71.76 per cent of steady state and still climbing.

The conversion failure case shows the model reporting its own domain of breakdown: with `S` at 79 per cent of ceiling but `eta = 0.05` and `C = 0.35`, `eta S = 0.0824` falls below `C` and the steady state is negative at `Q_inf = −2.676`.

Two of these deserve emphasis because they are the ones most likely to be repeated loosely elsewhere. The DC-gain identification is exact and is not an analogy. And the `delta = 0` result is not the claim that forgetting is good; it is the claim that a system with no forgetting has no fixed point, which is a different and stronger statement.

7.10.2 10.2 What does not follow

No thermodynamic claim is made. The capacitor is a mathematical correspondence and not a physical mechanism. Nothing here proves a negentropy law, and higher static surplus does not imply higher canonical concentration.

Surplus is not identified with Shannon entropy or mutual information. The `ln N` ceiling has an information-theoretic reading and the reading bounds; it does not identify.

The mapping onto any real social system is not established by anything in this paper. All three worked systems are constructed. What licenses applying the law to an actual organisation is a preregistered fit with held-out comparison against declared baselines, and the paper that specifies that protocol reports its own confirmatory application **failing** its second rule [Curzi, 2026i]: a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin. A reader who takes the present paper's arithmetic as evidence that institutions obey the storage law has read it wrongly. The arithmetic is derivation; the correspondence is a postulate awaiting corroboration that has, in the one confirmatory test so far reported, not been supplied.

Monotonicity in `u` is not established here and is contested. The framework's surplus is monotone increasing in `u`, and the empirical literature does not agree that the benefit of difference rises without turning over. Uzzi and Spiro [2005] find an inverted U, with creative success rising and then falling with cohesion; Mannix and Neale [2005] and van Knippenberg and Schippers [2007] find that surface-level diversity frequently costs in the short run; Putnam [2007] finds short-run reductions in trust and civic participation, with Abascal and Baldassarri [2015] as the rebuttal; and Sagan [1993] documents settings in which redundancy reduces safety. Against these, Granovetter [1973], Burt [2004] and Schindler et al. [2010] supply measured benefits of bridging and of effective count. The honest position is that the monotone form is this framework's construction over the regime it has examined, and that where the curve turns is an open empirical question. Harrison and Klein [2007] sharpens the demand further: three incompatible constructs travel under the word diversity — separation, variety and disparity — and a quantity like `u` owes the reader a statement of which one it is.

Diversification benefit depends on covariance, not count. Markowitz [1952] stated this in 1952 with the mathematics attached, and it is the same content as the effective-count result: `N` uncorrelated holdings and `N` perfectly correlated holdings have the same nominal `N` and different effective counts. Any reading of `ln N` as a benefit of headcount rather than of *audited* headcount is a misreading, and the concentration instrument that makes the audit computable is developed separately [Curzi, 2026].

The stability benefit of an effective count is not free of a purely statistical explanation. Doak et al. [1998] showed that diversity–stability correlations can arise from statistical averaging alone with no mechanism, and May [1972] and Gardner and Ashby [1970] showed that increasing diversity and connectance destabilises randomly assembled systems. McCann [2000] is the fair survey and Loreau and de Mazancourt [2008] names the decisive variable — synchrony — which is the ecological name for what this framework calls the coordination perimeter.

7.10.3 10.3 Relation to the surrounding literature

The storage law's nearest formal relatives are three. In **system dynamics**, `Q`, `eta S` and `C` are a level with an inflow, an outflow and a delay, formalised four decades earlier [Forrester, 1961], and the finding that people systematically misjudge such systems even under aligned incentives [Sterman, 1989] is empirical support for the central institutional claim of this paper. In **forecasting**, the recursion is simple exponential smoothing, whose optimality class is known [Brown and Meyer, 1961, Holt, 1957]; that class is a question this framework does not currently ask about its own model, and it should. In **computational neuroscience**, a leaky integrator with a threshold and a resulting firing rate is a century-old and empirically supported model [Lapicque, 1907, Gerstner and Kistler, 2002, Usher and McClelland, 2001]; that literature is the best available precedent for the *shape* of results derived from an integrator with a drive and a cost, and Curzi [2026] builds on exactly that shape, deriving a minimum sustainable frequency `f_min = k e^(2c/eta)` with `k = 0.483425732988414` Hz and the trichotomy `f > = < f_min` corresponding to `Q_inf > = < 0`. The precedent grounds the structure and not the formula, and Section 12 records the distinction.

The **requisite-variety** tradition supplies the reason `ln N` appears as a ceiling at all: only variety can absorb variety [Ashby, 1956]. The **thermodynamic** reading of the cost term has a warrant in Landauer [1961], but the warrant is about erasure in physical systems and does not transfer to an institutional cost line; it is cited here as the reason the analogy is tempting and not as its licence.

Within the corpus, the results of this paper are consumed in three ways. They are **restated as an equality** — `|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|` — in the refinement analysis [Curzi, 2026] and in the fixed-point analysis of drive conservation [Curzi, 2026], whose Theorems 28.1 and 28.2 report a maximum residual of 2.0624e-15 over a hundred passes. They are **reused with different symbols** in the return-difference formulary [Curzi, 2026q]. And they are **converted into a procedure** in the consolidated reproduction workbook, whose organising rule is that `ln(1.00) = 0`, so a structure with one effective bloc has zero available surplus rather than a low one, and step one of the method is to collapse before measuring [Curzi, 2026].

7.11 11 Verification

The results of this paper compound from five postulates. Each is atomic, each is checkable on its own, and each carries its status. **If a postulate could be dropped and the result still follow, it has been dropped**; what remains is what the chain actually consumes.

7.11.1 11.1 Dimensional audit

The equation must balance before anything else about it is worth checking. The audit below states each symbol's dimension, its unit, and whether the framework fixes it or is silent. `[Q]` denotes the dimension of stored order and `[T]` the dimension of time; the period is the unit of `[T]` and every rate below is *per period*, which is why no explicit `[T^−1]` appears once the equation is written in discrete steps.

Symbol	Dimension	Unit	Status
`Q_t`	`[Q]`	proxy units (rules, protocols, runbooks, modules)	**measured**, in whatever unit the declared order proxy counts

Symbol	Dimension	Unit	Status
`U_t`	`[1]`	none	**defined** dimensionless: `u = 1 - (x.y)^2` with `x`, `y` unit vectors
`S_t = F(U_t)`	`[1]`	nats	**derived** dimensionless: a logarithm of an effective count, so a pure number reported in nats
`eta`	`[Q]` per nat, i.e. `[Q]` per step per nat	proxy units per nat per period	**measured**; *not* dimensionless in this framework
`delta`	`[1]`	per period	**measured** dimensionless fraction per step
`C_t`	`[Q]` per step	proxy units per period	**measured**
`h_t`	`[Q]` per step	proxy units per period	**derived**
`Q_inf`	`[Q]`	proxy units	**derived**: `[Q]/step` divided by `[1]/step` returns `[Q]`
`tau`, `t_half`, `t_95`	`[T]`	periods	**derived**

Does the equation balance? `Q_{t+1} = (1-delta) Q_t + eta S_t - C_t` balances if and only if every term on the right carries `[Q]`. The first term does, since `delta` is dimensionless. The third does by construction. The second does only if `eta S_t` carries `[Q]`, and since `S_t` is a dimensionless count in nats, **`eta` must carry `[Q]` per nat per step**. It is therefore not dimensionless, and the worked systems are consistent with this: Section 9.3 reports `eta` as *9.0 runbooks per nat per month*, which is `[Q]` per nat per step exactly. Section 9.1 reports `eta = 0.50` without units, which is a **silence in the source**: the research group's proxy is a technique count, so the value should read 0.50 techniques per nat per week, and it is recorded here rather than quietly corrected.

Two consequences follow and both are worth stating. First, **`eta/delta` carries `[Q]` per nat** and is the DC gain in the proper sense: units of stock per sustained nat per period. Second, **`ln N` is dimensionless, so `eta ln N` carries `[Q]` per step** and is commensurable with `delta Q_t` and with `C_t`, which is what makes the headroom gauge `h_t` a well-formed subtraction. A headroom calculation in which `C` is denominated in hours or currency while `Q` is denominated in rules is dimensionally malformed however arithmetically tidy it looks; that is the denomination requirement the survey protocol imposes [Curzi, 2026i].

Where the framework has no operational unit, it is said here. `Q` is *stored order*, and stored order has no natural unit: it is counted in whatever the declared proxy counts, and two systems' `Q` values are not comparable unless they declare the same proxy. Comfort itself, the underlying construct of the corpus, has **no operational unit at all**; it is treated as a dimensionless magnitude by stipulation, not by measurement, and Section 12 records that the quantitative structure required for that stipulation has not been tested.

7.11.2 11.2 The postulates

P1 — The recurrence is linear and first order. `Q_{t+1} = a Q_t + x_t` with `a` and `x_t` not depending on `Q`. *Status:* **DEFINITION** — this is the model form adopted by the framework, restated

from the storage law by setting `a := 1 - delta` and `x_t := eta S_t - C_t`. It carries no empirical weight on its own, and Section 12 records that its application to a real system is a postulate.

P2 — `0 < delta <= 1` implies `0 <= a < 1`. *Status:* **ARITHMETIC.** `a = 1 - delta`; subtracting a number in `(0, 1]` from 1 gives a number in `[0, 1)`.

P3 — A geometric series with ratio `|r| < 1` sums to `1/(1-r)`. *Status:* **CITED** — elementary, and stated with the convergence condition in Elaydi [2005]; the metric-space form that supplies uniqueness and start-independence as well is Banach [1922].

P4 — Surplus per period is bounded: `0 <= S_t <= ln N`. *Status:* **CITED** — the range result of the surplus functional [Curzi, 2026], following from `f(u) = ln(1+(N-1)u)` on `u` in `[0,1]`.

P5 — Cost accrues every period whether or not surplus was delivered. *Status:* **POSTULATED** — this is an assumption about how real systems charge overhead; it has no external support and is the only postulate in the chain that does not. It is used at exactly one place, R6 below, and the chain's other results do not depend on it.

7.11.3 11.3 The compounding chain

```
P1 + P2          -> L1   the update map T(Q) = aQ + x has |T(Q)-T(Q')| = a|Q-Q'|
                         with a < 1, so T is a contraction.                  ARITHMETIC

L1 + P3          -> R1   T has exactly one fixed point Q* = x/(1-a) = (eta S - C)/delta,
                         reached from any start.                             CITED (banach1922operations)

P1               -> L2   unrolling t steps gives
                         Q_t = a^t Q_0 + sum_{k=0}^{t-1} a^{t-1-k} x_k.       ARITHMETIC

L2 + P3          -> R2   with x_k = x constant, Q_t = x(1-a^t)/(1-a) + a^t Q_0,
                         so Q_t -> x/delta and the approach is geometric.     ARITHMETIC

L2 + P4          -> R3   Q_t <= max{Q_0, eta ln N / delta} for delta > 0.     ARITHMETIC

P2 at delta=0    -> R4   a = 1: L1 gives |T(Q)-T(Q')| = |Q-Q'|, so T is an isometry,
                         P3 does not apply, no fixed point exists, and
                         L2 gives Q_t = Q_0 + sum x_k, bounded only by
                         Q_0 + t eta ln N.                                   ARITHMETIC

R2               -> R5   the transient a^t falls to 1/e at
                         tau_exact = -1/ln a and to 1-p at
                         t_p = ln(1-p)/ln a.                                 ARITHMETIC

R1 + P4 + P5     -> R6   a nondecreasing step needs eta S_t > delta Q_t + C_t, and
                         since S_t <= ln N this requires C_t <= eta ln N - delta Q_t;
                         beyond that no geometry prevents decline.           ARITHMETIC
```

Every line is checkable from its own inputs. **R1, R2, R3, R4, R5 and R6 are arithmetic consequences of one definition (P1), one arithmetic fact (P2), one cited theorem (P3) and one cited range result (P4).** Only R6 consumes a postulate with no external support, and it consumes it only for the interpretation of `C` as a standing charge; the inequality itself holds however `C` is generated.

What this chain does not establish. It does not establish that any organisation obeys P1. That is the modelling postulate of Section 4.9, it is not verified here, and the one preregistered confirmatory test of it so far reported in this corpus did not meet its own declared margin [Curzi, 2026i]. A chain whose steps are ARITHMETIC and CITED delivers the mathematics with certainty and delivers nothing whatever about institutions.

7.11.4 11.4 The equilibrium, reached by summing the series

The single most useful check a reviewer can perform is to confirm that the equilibrium is what the geometric sum says it is, rather than taking Theorem 3.5 on trust. Take the worked research group of Section 9.1: `delta = 0.10`, so `a = 0.90`; `eta = 0.50`; `S = 1.6487`; `C = 0.20`; `Q_0 = 0`. The constant net drive is

```
x = eta S − C = 0.50 x 1.6487 − 0.20 = 0.82435 − 0.20 = 0.62435
```

From R2 with `Q_0 = 0`, `Q_t = x (1 − a^t)/(1 − a) = x (1 − 0.9^t)/0.10`. Summing the series term by term, each term being `x a^k`:

```
k =  0:  0.62435 x 0.9^0  = 0.624350     running total  0.624350
k =  1:  0.62435 x 0.9^1  = 0.561915     running total  1.186265
k =  2:  0.62435 x 0.9^2  = 0.505724     running total  1.691989
k =  3:  0.62435 x 0.9^3  = 0.455151     running total  2.147140
k =  4:  0.62435 x 0.9^4  = 0.409636     running total  2.556776
k =  9:  0.62435 x 0.9^9  = 0.241878     running total  4.066423
k = 19:  0.62435 x 0.9^19 = 0.084324     running total  5.484292
k = 29:  0.62435 x 0.9^29 = 0.029395     running total  5.978562
-----------------------------------------------------------------
limit :  0.62435 / (1 − 0.9)             = 6.2435
```

The running totals at `t = 1, 2, 4, 10, 20, 30` are 0.6244, 1.1863, 2.1471, 4.0664, 5.4843 and 5.9786, which are the entries of the convergence table in Section 9.1 to four decimals. The limit `0.62435/0.10 = 6.2435` is `Q_inf` to within the rounding of `S`, which the source carries as 6.2433 from the unrounded `S = 1.64866...`. **The equilibrium is not an assertion; it is the sum of a geometric series, and the sum can be checked by hand in a few lines.**

The `delta = 0` case is checked the same way and fails visibly. With `a = 1`, the terms are `x a^k = 0.62435` for every `k`, so the running total after `t` steps is `0.62435 t`: 6.2435 at `t = 10`, 12.487 at `t = 20`, 62.435 at `t = 100`. The series does not converge, `1/(1−a)` is undefined, and there is no limit to report. That is R4, and it is visible in three lines of arithmetic.

7.11.5 11.5 Two secondary checks

The time constant. At `delta = 0.10`, `tau_exact = −1/ln(0.9) = −1/(−0.105361) = 9.4912` periods against the rule of thumb `1/delta = 10`, a 5.4 per cent difference — the entry in the master table of Section 5.4. And `t_95 = ln(0.05)/ln(0.9) = (−2.995732)/(−0.105361) = 28.43` periods, which is the figure quoted for the research group's review timing. **ARITHMETIC.**

Multiplicative composition of loss. Four channels at 0.030, 0.050, 0.020, 0.015: `(1−0.030)(1−0.050)(1−0.020)(1−0.015) = 0.970 x 0.950 x 0.980 x 0.985 = 0.889508...`, so `delta_total = 0.110492`, against the naive sum 0.1150. Both figures appear in Section 6.3, and the check is one line of multiplication. **ARITHMETIC.**

7.11.6 11.6 Limits, asymptotics, and the model's valid domain

The `delta = 0` boundary is examined at length in Section 4.4. Completeness requires the rest.

`delta -> 0^+`. `tau = 1/delta -> infinity` and `Q_inf = (eta S − C)/delta -> +infinity` when the net drive is positive, `−infinity` when it is negative. The fixed point survives for every `delta > 0` and diverges only in the limit. **The limit is singular**: at exactly `delta = 0` the fixed point does not exist at all rather than being infinite, because the map is an isometry and the geometric series does not converge. The divergence and the pathology are two different statements and are often conflated.

`delta = 0` exactly. No contraction, no fixed point, error persistence 1 forever, bound linear in `t`. Section 4.4. **This value is outside the model's valid domain.**

`delta -> 1^-`. The pole `a = 1 - delta -> 0`, so `tau_exact = -1/ln(1-delta) -> 0` and the system has no memory: `Q_{t+1} -> eta S_t - C_t`, a pure input-output map with no state. The fixed point still exists and equals `eta S - C`, but the object being modelled is no longer a stock. **At `delta = 1` the storage law is not a storage law**; it is an instantaneous transducer, and any claim about accumulation is vacuous there. The measurement machinery also breaks first: at `delta > 0.30` the rule of thumb `tau = 1/delta` fails badly, overstating `tau_exact` by 39 per cent at `delta = 0.5`, and above `delta = 0.5` time constants should not be used at all.

`eta -> 0`. The drive term vanishes and `Q_{t+1} = (1-delta)Q_t - C_t`, so `Q_inf = -C/delta <= 0`: the system decays to the admissibility floor at rate `delta` while paying `C`, which is exactly the conversion pathology of Section 4.5 taken to its limit. **Zero conversion means the input is decoupled from the stock entirely**, so no interaction geometry of any quality changes the outcome — the framework's `eta = 0` clause is the same statement as the corpus's `G = 0` clause in the contact postulate, where a zero coupling coefficient makes the product `E = G B` vanish however large `B` is.

`N -> 1`. `ln N = 0` and `F(u) = ln(1 + 0 u) = 0` for every `u`. **A single category has zero available surplus, not a small amount.** The ceiling and the achievable value collapse together, so `Q_inf = -C/delta <= 0` and headroom is `-delta Q - C < 0` at any positive stock. This is the boundary that makes the bloc-collapse operation consequential: a structure with one effective bloc has `ln(1.00) = 0` available surplus [Curzi, 2026,?].

`N -> infinity`. `ln N -> infinity`, so the per-period ceiling grows without bound — but only logarithmically, and the marginal gain `ln N - ln(N-1) ~ 1/N` vanishes. Against this, coordination cost grows with the number of pairs, `n(n-1)/2` [Brooks, 1975, Conway, 1968], which is quadratic. **The two curves cross, and beyond the crossing point additional categories cost more than the ceiling they buy.** The limit is therefore not a licence to raise `N` indefinitely; it is the reason `N` is the most expensive and least rewarding of the four levers.

`u -> 0`. `N_eff -> 1`, `F -> 0`. Every category is a duplicate of every other, and the roster has left the arithmetic. **`u -> 1`.** `N_eff -> N`, `F -> ln N`, the ceiling is attained and no further gain is available from geometry. Strict concavity means the approach is front-loaded: in the worked case, moving `u*` from 0.00 to 0.25 buys 1.0116 nats and moving it from 0.75 to 1.00 buys 0.2468.

Stated domain. The storage law as used in this paper is valid for **`delta` in `(0, 0.5]`, with `(0.02, 0.30)` the band in which the standard approximations hold and the decay measurement has usable signal**; `eta > 0`; `N >= 2`; `u` in `[0,1]`; and `C_t <= (1-delta)Q_t + eta S_t` each period. Outside that domain the equation still has solutions and they should not be interpreted as statements about a stock.

7.11.7 11.7 Invariance, and how far the scale-freedom claim has been tested

The scale-freedom claim of Section 4.9 is an **invariance of form under a change of unit and a change of system size**, and it is worth stating as one. Let `Q` be rescaled by a positive constant `c`, so `Q' = cQ`. The law is invariant provided `eta` and `C` rescale with it — `eta' = c eta`, `C' = cC` — while `delta`, `N` and `u` are unchanged, since `delta` is a dimensionless fraction and `S` is a dimensionless count. **`delta`, `u` and `N` are therefore the unit-free quantities of the framework, and `Q`, `eta` and `C` are the unit-carrying ones.** Two systems can be compared on `delta`, `u` and `N` directly; they cannot be compared on `Q`, `eta` or `C` unless they share a proxy.

What would break the invariance: a `delta` that varies with `Q`, which would make the transformation non-linear; an `eta` that saturates at large `S`, which the linear model excludes by assumption; or a proxy whose counting unit is not extensive, so that doubling the system does not double the count.

How many scales it has been tested at. Within this paper: three, and all three are constructed rather than observed. Elsewhere in this corpus the same five-parameter survey has been applied from a single desk to a 22-household block, reproducing each system's observed decline — predicted −0.250474 against observed −0.250 at the desk and −0.996371 against

−1.000 at the library, with residuals under 1.4 per cent of the ceiling term [Curzi, 2026k]. **That is five systems, one corpus, one author, and no independent replication.** The invariance is a claim about form that has been exercised, not an invariance that has been established.

7.11.8 11.8 Analogy discipline, and the precision of the numbers

The circuit identity is a shared mathematical form, not a derivation from physical premises. The storage law and the RC low-pass section are governed by the same difference equation, so every theorem about first-order discrete systems applies to both. Nothing about charge, capacitance or current is claimed to hold of an institution, and no physical result's authority transfers. A leaky integrator is a leaky integrator whether it holds charge or stored order, and that is the whole of the claim. The same label applies to the leaky integrate-and-fire precedent [Lapicque, 1907, Gerstner and Kistler, 2002, Usher and McClelland, 2001] and to the reverberation analogy [Sabine, 1922]: **shared form, not derivation.**

No thermodynamic result is used as evidence. Landauer [1961] establishes a physical lower bound on the energy cost of erasure; it is cited in Section 10.3 as the reason the analogy is tempting, and it is **not** used to warrant any claim about institutional cost. The entropy invoked anywhere near `ln N` is **Shannon entropy**, a dimensionless information measure in nats, and not thermodynamic entropy, which carries `[energy][temperature]^-1`. They are different objects and the framework does not identify surplus with either.

Precision. The numbers in this paper are of three kinds and are quoted accordingly. *Defined or derived-from-defined* quantities — `ln 8 = 2.0794`, `ln(1-delta)`, `t_95`, `tau_exact`, the geometric-series limits — are exact to the precision printed and may be recomputed to any further precision. *Illustrative model output* — every entry in every trajectory table, `Q_inf = 6.2433`, `Q_inf = -2.676`, `Y_inf = 178.67` — is the output of stipulated parameters and carries **no measurement uncertainty because there was no measurement**; four decimal places there record the arithmetic, not the confidence. *Measured or fitted* quantities do not appear in this paper, because nothing in it was measured; the worked `delta = 0.0810` with `R^2 = 0.9980` in Section 6.4 is a fit to a constructed series and is illustrative in exactly the same sense. Any application of this framework to real data must report fitted parameters with intervals, and Section 12 records that the corpus's one such report gives a profile interval on the retention coefficient and no interval on several other quantities.

7.12 12 Limitations

The three worked systems are constructed, not observed. No dataset was collected for this paper. Every number in Section 9 is the output of the stated parameters under the stated model, carried through so a reader can re-derive it. They demonstrate that the procedure runs and that it reports its own failures; they are not evidence that any organisation behaves this way, and they should never be cited as such.

The correspondence between the equation and any social system is a postulate, and its one preregistered confirmatory test did not meet its declared margin. Curzi [2026i] specifies the protocol and applies it: the fit landed strictly inside the admissible box and its parameters agreed with independent measurement within 4.8 per cent on decay and 16.8 per cent on conversion, but it beat a matched AR(1) by 2.8 per cent on held-out error against a preregistered requirement of 10 per cent, so the verdict was *parameters corroborated, predictive advantage not demonstrated.* That is the current evidential status of every applied claim in this paper, and it is weaker than the mathematics.

Which parameters were chosen rather than measured. In all three worked systems, `(N, u*, eta, delta, C)` are stipulated. The artefact-ladder bands of Section 6.7 are explicitly badged as order-of-magnitude orientation and are the author's working priors for structuring a survey, not measurements; they must not be cited as findings. The `delta` portfolio table of Section 6.8 is a design proposal with stated rationales, not an empirical distribution. The mirror-reflectivity

table of Section 4.2 is the arithmetic of an illustration, and real visual thresholds depend on ambient light and on the observer.

`N_eff = 1 + (N−1)u` is not a Hill number, and no external literature contains it. This is recorded in the corpus's own coverage audit as a claim with no external support. `1/HHI` is the Hill number of order two and `exp(H)` is the Hill number of order one, and both are fully grounded [Hill, 1973, Jost, 2006]; `1 + (N−1)u` is a linear interpolation between 1 and `N` and is this framework's own construction, adopted by design axiom S3. A reviewer who reaches for the Hill-number literature as support will find that it does not apply. The two quantities are used under one effective-count heading across this corpus, which is a real exposure; the repair is to derive the properties of `1 + (N−1)u` in the framework's own terms and to reserve the Hill citations for `1/HHI` and `exp(H)`, where they are exact. Neither is a measured quantity in this paper.

`u` has not been shown to be a quantity. The framework treats non-parallelism as a measurable magnitude, and Michell [1997] argues that psychology routinely fails to test whether its attributes are quantitative at all before measuring them. The available remedies are conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971] or a Rasch-model demonstration [Rasch, 1960]; neither has been run for `u`, and until one is, the quantitative treatment of `u` in this paper is a stipulation. Relatedly, if `u` is read as a similarity measure over people it cannot be assumed to behave metrically, since psychological similarity is asymmetric and violates the triangle inequality [Tversky, 1977] — which is the empirical counterpart of the framework's own clause that the triangle inequality fails for geometric triples.

The linear model is chosen for tractability. Real institutional stocks may saturate, have thresholds, exhibit hysteresis, or have a state-dependent `eta`. Theorem 3.5 flags its own breakdown when `eta F(u*) < c*`, which is a useful property, but a linear model that reports a negative equilibrium has told the analyst only that it has left its domain, not what replaces it. A fitted retention coefficient outside `[0,1]` is likewise a specification failure and not a number to be clipped.

Constant parameters over the estimation window is an assumption that organisations violate. Turnover, reorganisation and a change of leadership all move `delta`, `eta` and `N`. A fit spanning a regime change is fitting an average of two systems, neither of which exists.

What would falsify the account. H1 fails on observed linear or unbounded accumulation under constant drive, or on a fitted `a` outside `[0,1]`. H2 fails if measured settling time scales with the size of the intervention rather than with the measured `delta`. H3 fails if a component with no observed retirement nonetheless shows mean reversion after a perturbation. H4 fails if measured total loss is reliably closer to the naive sum than to the product. H5 fails if raising `eta` produces the smaller gain in systems whose diagnostic signature is a retention fault. The scale-freedom postulate of Section 4.9 is falsifiable at each scale separately and is not falsifiable in general, which is a genuine weakness of that claim and is stated as one.

The framework's own instrument is exposed to the failure it describes. A measure that becomes a target ceases to be a measure [Goodhart, 1975, Campbell, 1979], and `Q` is a measure. Section 6.8 prescribes raising a captured metric's `delta` rather than defending it, but the prescription is a design rule and not a demonstration, and the same warning attaches to `h` — the headroom gauge is at least as gameable as anything it audits.

Nothing in this paper is a human-subjects finding. No interventions were run, no participants were recruited, and the declination constraint discussed in Sections 7 and 8 is a design constraint on future work rather than a report of past practice.

The reflection problem, and why it outranks the fitting result. Section 3.5 states it in full: applied to a group whose members both influence and constitute the aggregate, the specification is reflexive, and Manski [1993] proves the endogenous social effect is not identified from observational data. **No fit — however well designed, however honestly held out — recovers that parameter**, and Shalizi and Thomas [2011] show the same confounding for network data. This paper adopts no remedy, so the applied reading of `eta` is a **modelling position rather than an identified empirical result**. The most realistic available repair is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001]; the survey protocol's fixed cate-

gory set is a step toward it and not a substitute for it.

The Lucas critique applied to this framework's own parameters. Section 6.8 asserts that a measure changes the behaviour of the system that is measured by it, and prescribes raising a captured metric's `delta` in response. The framework then treats `delta`, `eta` and `N` as constants over an estimation window. **Both cannot be true without qualification.** The position this paper takes is the narrow one: the parameters are treated as stable *only over a declared window and only while `Q` is not itself a reported target*, which is exactly why Section 6.9's first and strongest defence is to keep the order proxy out of the reporting line. Where `Q` is reported upward, the framework's own reflexivity applies to it and the constancy assumption fails. That is a real restriction on where the model may be used and it is stated as one rather than assumed away.

The functional form is underdetermined. Section 3.1 gives the alternatives — Latané [1981]'s sub-linear power law and Briscoe et al. [2006]'s `n log n` — and records that nothing assembled here discriminates between them. **The logarithm is chosen for its axiomatic characterisation, not established by measurement.**

No scaling or heavy-tail claim is made in this paper, so the distribution-fitting procedure that any such claim would require [Clauset et al., 2009] does not apply. The distributional statement that *is* made — that decay is geometric — is tested by the log-linearity check of Section 6.4 and by the residual inspection prescribed there, and Section 6.1 records that a cliff fitted as an exponential produces an error in a specific and predictable direction.

Unmet rigour checks, printed rather than left to a reviewer. (i) *Model comparison.* This paper fits nothing, so no information criterion applies to it; but the corpus's one applied fit reports held-out RMSE against baselines and **no AIC or BIC**, so the parameter count was never paid for, and a 2.8 per cent margin over a two-parameter AR(1) may not be a margin at all [Curzi, 2026i]. (ii) *Effect sizes and intervals.* Nothing in Section 9 carries an interval, because nothing in Section 9 was measured; that is honest for a constructed example and is not a substitute for the intervals a real application owes. (iii) *Scale type.* `u` and `Q` are treated as ratio-scale quantities throughout — differences, ratios and logarithms are all taken — and **the conjoint-measurement conditions that would license that treatment have not been tested** [Luce and Tukey, 1964, Krantz et al., 1971]. The treatment is a stipulation. (iv) *Independent replication.* Every system in this paper and every system in the corpus's applied papers was constructed or surveyed by one author; there is no independent replication of any parameter estimate, and the reproducibility literature is unambiguous about what that is worth [Open Science Collaboration, 2015, Ioannidis, 2005]. (v) *Directional versus point prediction.* Meehl [1967] observes that directional predictions are nearly costless in the social sciences; this framework does make point predictions with constants, which is the right answer to that objection, but only Curzi [2026i] has so far put one of them at risk.

7.13 13 Conclusion

The corpus's dynamical core is the storage law `Q_{t+1} = (1−delta) Q_t + eta S_t − C_t` with `S_t = ln(1 + (N−1) U_t)`. Its update map is Lipschitz with modulus `k = 1 − delta`, so any positive loss makes it a contraction, and the contraction supplies a unique start-independent steady state `Q* = (eta S − C)/delta`, which is exactly the DC gain `1/delta` of a first-order discrete section applied to the net drive. The exact solution `Q_t = (1−delta)^t Q_0 + sum_{k=0}^{t−1} (1−delta)^{t−1−k}(eta S_k − C_k)` and the uniform bound `Q_t <= max{Q_0, eta ln N / delta}` follow, with a single pole at `z = 1 − delta`, a time constant `tau = 1/delta` against the exact e-folding time `−1/ln(1−delta)`, and settling times `t_p = ln(1−p)/ln(1−delta)`. **At `delta = 0` the map ceases to contract, no unique fixed point exists, the bound degrades to `Q_t <= Q_0 + t eta ln N` — linear and unbounded — and the error-persistence column is a column of ones, so no mistake ever decays out of the stock. Zero loss does not preserve a system; it removes its destination and its capacity to revise.** These are the statements downstream papers should quote rather than paraphrase.

The worked research group returns `N = 8`, `u* = 0.60`, `N_eff = 5.20` and `S = 1.6487` nats, 79

per cent of the ceiling `ln 8 = 2.0794`; with `eta = 0.50`, `delta = 0.10` and `C = 0.20`, `Q_inf = 6.2433` against a hard ceiling of 10.397 and `t_95 = 28.43` weeks, so a review at week 12 finds the system at 71.76 per cent of steady state and still climbing. The conversion failure case returns a negative steady state, `Q_inf = −2.676`, which the model reports as its own domain of breakdown rather than concealing. All of this is arithmetic. Whether any institution is described by the equation is a separate question, it is a postulate rather than a result, and the one preregistered confirmatory test of it reported in this corpus did not meet its own declared margin.

7.13.1 Declarations

Data and code availability

No external dataset was collected for this paper. All numerical results are computed from the stated parameters under the model developed in the text, and every table can be re-derived from the equations and parameter values printed beside it. The three systems worked in Section 9 are constructed examples with declared parameters, not field observations. The figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper and the model equations of Sections 3 to 6.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, produced the worked examples and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The worked systems are constructed. The declination constraint referred to in Sections 7 and 8 — that no intervention removing a participant's ability to decline is within the framework's scope — is a design constraint on prospective work and is not a report of any study performed.

Preregistration

This paper reports no preregistered test of its own; its results are derivations. The preregistered protocol for testing the storage law against a real system, together with the confirmatory application that **failed** its second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared margin of 10 per cent — is reported in Curzi [2026i]. Any applied claim resting on this paper inherits that result.

Chapter 8

Conversion, Coordination Cost, and Headroom: Why Eta Is the First Lever and Why the Optimum Group Is Smaller Than the Crossover

Abstract. This paper develops the three remaining terms of the storage law and the gauge reading them together: headroom h_t = eta*ln N - delta*Q_t - C_t, with three regimes and a maximum one-step increment r_max = h_t. Its central asymmetry result is arithmetic: doubling the conversion coefficient eta is exactly squaring the category count, since 2*eta*ln N = eta*ln(N^2), and doubling the surplus by raising non-parallelism alone is feasible only when u <= 1/(sqrt(N)+1), so at N = 8 and u = 0.30 the required u' is 1.23 and the move is arithmetically impossible. Conversion composes multiplicatively through recording, binding and reach gates, so three gates at 0.7 give eta = 0.343 and the weakest dominates: lifting recording from 0.20 to 0.80 gains 300 per cent, lifting binding from 0.60 to 0.90 only 50 per cent. On the cost side, coordination priced as C = kappa*n(n-1)/2 against a logarithmic budget gives a net optimum near sqrt(eta/kappa), far below the crossover at which cost exhausts the budget: at eta = 0.30 and kappa = 0.010 the optimum is six. Topology is shown to be a free lever - eighteen people rewired from a complete graph into three cells of six move from h = -0.4571 to h = +0.1189, a swing of 0.576, with the surplus ceiling untouched because N = 18 either way. A nine-person group is diagnosed at eta = 0.030 with h_0 = -0.1062 and a negative steady state, needing only eta >= 0.091.

Keywords: conversion efficiency; capture gates; coordination cost; communication topology; headroom; sustainability ceiling; queueing slack; shedding stored order

Series: B. Engineering the Surplus **Paper:** CCP-008

8.1 1 Introduction

The storage law of this framework has four parameters and one of them has already been studied on its own. The leaky-integrator analysis establishes that retention loss `delta` decides whether an accumulating system has a destination at all, and that `delta = 0` is a pathology rather than an ideal [Curzi, 2026a]. That leaves three terms — the conversion efficiency `eta`, the cost flow `C`, and the ceiling `ln N` that bounds what any interaction geometry can deliver — and one gauge that reads all four together.

This paper develops those three terms and the gauge. The gauge is **headroom**,

```
h_t := eta ln N − delta Q_t − C_t
```

which is the sustainability ceiling of the storage law rearranged so that its sign can be read off a single subtraction: negative means stored order declines this period whatever is done to the interaction geometry, zero means maintenance only, positive means an increment `r_t <= h_t` is available. Headroom is this framework's name for the reserve that Cyert and March [1963] identified as organisational slack, and its behaviour near exhaustion has the same shape as utilisation approaching one in a queue, where waiting time diverges [Kingman, 1962, Hopp and Spearman, 2011].

The paper's central result is an asymmetry, and the asymmetry is arithmetic rather than empirical. `eta` enters the budget linearly and `N` enters logarithmically, so doubling `eta` is *exactly* squaring `N`: `2 eta ln N = eta ln(N^2)`. An organisation with eight genuinely distinct methodological lineages that wants to double its structural budget can install a decision log, or it can find fifty-six more genuinely distinct lineages. The second route to the same number is a decade of recruitment; the first is a rota. The same asymmetry runs against the third lever: doubling surplus by raising non-parallelism `u` alone is feasible only when `u <= 1/(sqrt(N) + 1)`, so at `N = 8` and `u = 0.30` the required `u'` is 1.23 and the move is not expensive but arithmetically impossible.

Four further results follow from taking the cost side seriously. Conversion **composes multiplicatively** through three named gates — recording, binding, reach — so three gates at 0.7 give `eta = 0.343` and the weakest gate dominates every intervention decision. Priced against a logarithmic budget, coordination cost of the form `C = kappa n(n−1)/2` gives a net optimum near `sqrt(eta/kappa)` that sits **far below the crossover** at which cost exhausts the budget: at `eta = 0.30` and `kappa = 0.010` the optimum is six and the crossover is thirteen, and the band between them is where nothing visibly breaks while every additional member reduces net output. **Communication topology is a free lever**: eighteen people rewired from a complete graph into three cells of six move from `h = −0.4571` to `h = +0.1189`, a swing of 0.576, with the surplus ceiling untouched because `N = 18` either way. And at the fixed point the gauge collapses to `h_inf = eta (ln N − F(u*))`, in which `delta`, `C` and `Q` all cancel — so the only headroom a converged system has is exactly the surplus it is failing to extract from its own difference.

The paper's own worked case refuses the lesson the framework's name would suggest. A nine-person research group is surveyed at `N = 6`, `u = 0.55` and `S = 1.3218` nats, which is 74 per cent of its own ceiling and is good geometry by any reading. Its headroom is nonetheless `h_0 = −0.1062` and its steady state is negative at `Q_inf = −1.607`, because `eta = 0.030` discards 97 per cent of what the geometry produces. Adding a seventh distinct lineage buys 0.0046; cutting coordination cost to zero leaves headroom at roughly a tenth of what the conversion fix delivers. **The geometry this corpus is named for was not the binding constraint in the paper's own example, and the paper says so rather than choosing a different example.**

Three things are new here relative to the surrounding literature rather than relative to the framework. First, the crossover and the optimum are separated and both are computed, where the operations literature supplies the quadratic communication count [Brooks, 1975, Conway, 1968] without pricing it against a logarithmic benefit. Second, coordination cost and interaction geometry are placed on one commensurable scale, so that a topology change and a hiring decision can be compared in the same units. Third, **deliberately reducing stored order is developed as a fourth lever with a computed exchange rate `dQ = dh/delta`**, and is shown to outperform every other move at the paper's own base parameters — while being ruled out of bounds wherever the stored order in question constitutes a person's ability to decline participation.

The framework's static half — the surplus functional `F(u) = ln(1 + (N−1)u)` and its uniqueness under four axioms, with the affine effective count `N_eff = 1 + (N−1)u` — is presupposed rather than re-derived; it is established in Curzi [2026], whose standing clause that the functional is deployed as a **lower bound** on the ensemble's information rather than the information itself governs every number in this paper.

Sections 2 and 3 state the hypotheses and the framework, including the identification problem that limits every applied reading of `eta`. Section 4 develops conversion, Section 5 cost, and Section 6 the headroom gauge. Sections 7 and 8 collect the design specifications and the named failure modes. Section 9 works three systems end to end. Sections 10 to 13 discuss what follows, verify the chain, state the limitations, and conclude.

8.2 2 Hypotheses

The results of Sections 4 to 6 are of two kinds and the distinction is load-bearing. Some are theorems about an equation and cannot be falsified by observation; they are labelled as derivations where they appear. Others are claims that the equation describes a real accumulating system, and those are stated below with what would defeat them. The protocol that operationalises them — a preregistered fit, a frozen model, and held-out comparison against four declared baselines — is specified in Curzi [2026i], whose own confirmatory application failed its second falsification rule: a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin.

H1 — Conversion is a stable scalar over the survey horizon. A real organisation's conversion of interaction into durable structure is well described by a single coefficient `eta` that is approximately constant across periods. **Predicted:** a fitted `b` in `Y_{t+1} = a Y_t + b F(U_t) − C_t + eps_t` that is stable across disjoint sub-windows of the estimation period, within its own interval. **Falsified by:** a fitted `b` that moves significantly between sub-windows with no recorded intervention, or that falls outside `b >= 0`. Where H1 fails, the linear storage law is the wrong model for that system and must be replaced rather than refitted.

H2 — Conversion composes multiplicatively across gates. Total conversion is the product of the recording, binding and reach fractions, not their mean and not one minus the sum of their complements. **Predicted:** `eta = eta_record x eta_bind x eta_reach`, so that an intervention on the smallest gate dominates an intervention of equal size on any other; at gates `(0.20, 0.60, 0.70)`, lifting recording to 0.80 gains 300 per cent against 50 per cent for lifting binding to 0.90. **Falsified by:** a measured `eta` reliably closer to the mean of separately measured gate values than to their product, across systems in which the three gates are separately measurable.

H3 — Conversion is the binding constraint more often than geometry. In systems with `h < 0`, the modal deficiency is `eta` rather than `u` or `N`. **Predicted:** across surveyed systems with negative headroom, the intervention with the largest computed `dh` is more often a capture intervention than a geometry intervention. **Falsified by:** a survey series in which geometry interventions dominate on computed `dh`. **This paper's own worked case is consistent with H3 and its second worked case is not**, which is reported in Section 9 rather than resolved.

H4 — Coordination cost is superlinear in headcount under an unmanaged topology. **Predicted:** regressing measured `C` on `n` across comparable units under complete-graph obligation returns an exponent materially above one. **Falsified by:** a measured exponent near one under complete-graph obligation. **Important qualification:** the exponent 2 used throughout this paper is **defined by the model** `C = kappa n(n−1)/2`, not estimated from data. Section 11.7 records that the log-log regression the source protocol proposes for testing this is precisely the procedure Clauset et al. [2009] shows to be unreliable, and states what would be required instead.

H5 — Topology moves headroom without moving the ceiling. Rewiring a fixed roster changes `C` and leaves `N` and therefore `eta ln N` unchanged. **Predicted:** for

eighteen people at `kappa = 0.006`, complete-graph obligation costs 0.9180 per period and three cells of six with liaison links cost 0.3420, a fall of 63 per cent, while the ceiling term `eta ln 18` is identical in both. **Falsified by:** a measured `C` after a cellular split that does not fall by at least half, which would indicate that the liaison multiplier was underestimated; or a measured fall in `u` below the pre-split level, which would indicate that the split destroyed the difference it was meant to preserve.

Two claims in this paper are **not** hypotheses and are labelled accordingly. That coordination cost is proportional to the number of active communication paths, `C1 = kappa P`, is a **modelling choice**: the combinatorial count `P(n) = n(n−1)/2` is arithmetic, and the interpretive step from that count to a cost is Brooks [1975]'s observation from practice, adopted here because it is the simplest form that reproduces the observed superlinearity and makes the crossover computable. That deference in an assembly can be represented as a common rotation coefficient, and the three-gate decomposition of `eta`, are **modelling conveniences** with no external support and no theorem behind them; both are badged at every use.

8.3 3 Framework

8.3.1 3.1 What is presupposed

The storage law. The equation this paper prices is

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,     S_t = F(U_t) = ln(1 + (N−1) U_t)
```

Symbol	Meaning	Range
`Q_t`	stored order: what persists — protocol, runbook, default, furnished room, trust	`>= 0`
`S_t`	surplus generated by interaction this period	`0 <= S_t <= ln N`
`delta`	retention loss, the fraction of stored order lost per period unmaintained	`[0,1]`
`eta`	conversion efficiency: interaction into *durable* structure	`>= 0`
`C_t`	cost, dissipation, overhead	`>= 0`

The quartet of a level, an inflow, an outflow and a delay is the standard vocabulary of system dynamics, formalised four decades before this framework [Forrester, 1961], and the finding that people systematically misjudge such systems even under aligned incentives [Sterman, 1989] is the empirical support for this paper's central institutional claim: that the crossing points below are invisible from inside.

The surplus functional. For two unit state vectors drawn from an ensemble of `N` categories, `u(x,y) = 1 − (x . y)^2` in `[0,1]` and `f(u) = ln(1 + (N−1)u)`, uniquely characterised by four axioms — angle-only dependence, zero at parallel, an affine effective count, normalisation at orthogonality — as established in Curzi [2026]. **Axiom S3, the affine effective count `N_eff = 1 + (N−1)u`, is a design axiom: chosen, not derived**, and it is flagged as such wherever this paper leans on it. The count it defines is developed with its full provenance, its exact bridge to the Rényi-2 participation ratio, and its measurement-error budget in Curzi [2026c], whose result R1 records that a hundred-person department at `u = 0.05` has an effective count of 5.95.

`1 + (N−1)u` is not a Hill number. The effective-count idiom is standard in ecology and in political science, where `1/HHI` and `exp(H)` are the canonical members of the family [Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979]. `1 + (N−1)u` is a linear interpolation between 1 and `N` and belongs to no order of that family. The concentration instrument that *is* a Hill number of order two — `1/HHI`, with the worked jurisdictional case giving a naive HHI of 0.293430 and 3.41 effective jurisdictions against a bloc-HHI of 1.000000 and an effective count of 1.00 — is a separate quantity [Curzi, 2026], and the two must not be conflated. This is repeated in Section 12 because a reviewer reaching for the Hill-number literature as support will find that it does not apply.

The results this paper consumes, all established in Curzi [2026a] and none altered here:

- **Admissibility.** `C_t <= (1−delta) Q_t + eta S_t` guarantees `Q_{t+1} >= 0`.
- **Exact solution.** `Q_t = (1−delta)^t Q_0 + sum_{k<t} (1−delta)^{t−1−k}(eta S_k − C_k)`.
- **One-step growth.** `Q_{t+1} > Q_t` if and only if `eta S_t > delta Q_t + C_t`.
- **Uniform bound.** `Q_t <= max{Q_0, eta ln N / delta}`.
- **Sustainability ceiling.** A nondecreasing step requires `C_t <= eta ln N − delta Q_t`.
- **Steady state.** With constant `U == u*` and `C == c*` and `delta > 0`, `Q_inf = (eta F(u*) − c*)/delta`, negative when `eta F(u*) < c*`.
- **Range.** `0 <= F(u) <= ln N`, which is what makes `eta ln N` a ceiling rather than a typical value.

The leaky-integrator paper's own conversion-failure case is the one this paper generalises: with `S` at 79 per cent of ceiling but `eta = 0.05` and `C = 0.35`, `eta S = 0.0824` falls below `C` and the steady state is negative at `Q_inf = −2.676`, which the model reports as its own domain of breakdown [Curzi, 2026a].

8.3.2 3.2 Two warrants for the logarithm, and the one the argument needs

The logarithmic form has two independent justifications and this paper needs only the weaker demand. The first is Fechnerian: sensation proportional to the logarithm of a dimensionless stimulus ratio measured from a floor [Fechner, 1966]. That warrant requires the perceptual story to be true and the perceptual story is contested. The second assumes nothing psychological at all: Peters and Gell-Mann [2016] shows that the time-average growth rate of a multiplicatively evolving quantity is the expectation of its *logarithm*, so a logarithmic evaluation follows from the dynamics. **Where a stock compounds rather than adds, the logarithm is the ergodic-average object, and no claim about perception is required.** This paper's argument needs only the second warrant. The ceiling `ln N` also has an information-theoretic reading as the maximum an ensemble of `N` equiprobable categories can carry [Shannon, 1948], and the requisite-variety tradition supplies the reason a variety term appears as a ceiling at all: only variety absorbs variety [Ashby, 1956, 1958]. Neither reading identifies surplus with Shannon entropy, and Section 11.8 says so.

The functional form is nevertheless underdetermined by the evidence, and the paper says so rather than waiting to be told. Social impact rising as a power of the number of sources with an exponent below one [Latané, 1981] and network value growing as `n log n` rather than `n^2` [Briscoe et al., 2006] both reproduce the qualitative facts this framework's `ln(1 + (N−1)u)` reproduces — diminishing returns to count, no threshold, monotonicity — and **nothing assembled in this corpus discriminates between them**. What the corpus's form has is an axiomatic characterisation with the design axiom named at the point where the choice is made; what it lacks is a discriminating measurement. Stumpf and Porter [2012] states the general form of the problem: many mechanisms produce indistinguishable curves. Section 12 carries this as an open limitation.

8.3.3 3.3 Identification: what no amount of fitting recovers

This limitation is structural rather than statistical and belongs here rather than in a caveat at the end. Wherever the storage law is applied to a group whose members both influence and constitute the aggregate — the usual case for a team, a board or an institution — the specification is reflexive: the interaction geometry `U_t` is a property of the group, the group's stored order shapes what its members do next, and the members are what the group is made of. Manski [1993] proves that in a linear-in-means reflexive specification **the endogenous social effect is not identified from observational data**: it cannot be separated from contextual effects or from correlated unobservables, however much data is collected. Shalizi and Thomas [2011] reinforce the point for network data, where influence and selection are generically confounded.

The consequence for this paper is exact. Every derivation in Sections 4 to 6 is a statement about an equation and is untouched. But the *interpretation* of a fitted `b` as the conversion efficiency of interaction into structure — as opposed to a contextual effect, a common shock, or selection — is **not identified by observation alone**, and no held-out comparison recovers it. That is a deeper limitation than anything the falsification protocol of Curzi [2026i] tests for, and it would remain even if every rule in that protocol passed.

No remedy is adopted in this paper, and the honest label is therefore a modelling position. The known routes are randomisation, an instrument, partial-population variation, or stipulating the reference group and defending the stipulation [Brock and Durlauf, 2001]; the last is the most realistic here, and it is the route the survey protocol's fixed category set and declared participation rule point toward without yet taking. **Until one of them is executed, `eta` as an estimated quantity is a coefficient in a declared specification and not an identified social effect**, and every applied statement in this paper inherits that status. The same applies with equal force to `kappa`, which Section 5.9 estimates by a procedure that is entirely observational.

8.3.4 3.4 Standing honesty clauses

Repeated wherever the law is used, and stated once here.

- **The triangle inequality fails for geometric triples.** Subadditivity holds on the parameter `u`, not on the geometry, so surplus bounds may not be chained through an intermediate state. The point is not exotic: psychological similarity is known to be asymmetric and to violate the triangle inequality [Tversky, 1977].
- **As deployed, the surplus functional is a lower bound** on the information an ensemble carries, not the information itself [Curzi, 2026]. Every number in this paper is a floor.
- **Surplus is not identified with Shannon entropy or mutual information**, and no thermodynamic negentropy law is proved.
- **Applying the storage law to a room, a team or an institution is a modelling decision**, badged as such every time it is made.
- **The constraint that an intervention may not remove a participant's ability to decline sits outside the arithmetic.** It is not derived from the law, the law gives no warning where it is violated, and Section 6.9 records the one place in this paper where a lever with a computed positive return is declared out of bounds anyway.
- **What this framework permits itself is bounded by its own interlude on category violation**, which refuses category innocence as a defence of acquisition and forbids moving a harm through a discovered gap in a scheme, and which concedes that the case where a feedback loop runs through another person's will is an unsettled seam the corpus does not close [Curzi, 2026]. The shedding lever of Section 6.9 is exactly a case of finding a gap in a scheme, and it is bounded accordingly.

8.4 4 Conversion: engineering `eta`

`eta` is the only parameter in the law that stands between an interaction and a consequence. Everything else describes what happens to consequences once they exist; `eta` decides whether they exist at all. It appears in every design inequality of the framework, and it always appears multiplied by something that is hard to change:

```
one-step growth :   Q_{t+1} > Q_t   <=>   eta*S_t > delta*Q_t + C_t
uniform bound   :   Q_t <= max{ Q_0 , eta*ln N / delta }
sustainability  :   nondecreasing requires  C_t <= eta*ln N - delta*Q_t
steady state    :   Q_inf = (eta*F(u*) - c*) / delta
headroom        :   h_t = eta*ln N - delta*Q_t - C_t
max increment   :   r_max = eta*ln N - delta*Q_t - C_t
```

Every one of these is **linear in `eta` and logarithmic in `N`**. That asymmetry is the whole argument of this section and Section 4.9 prices it exactly.

8.4.1 4.1 The `eta = 0` collapse

Set `eta = 0` in the exact solution:

```
Q_t = (1-delta)^t * Q_0  -  sum_{k=0}^{t-1} (1-delta)^{t-1-k} * C_k
```

The driven term vanishes entirely. What remains is a decaying transient plus an accumulated debit. For `delta > 0` and any `C_k >= 0`, `Q_t` is non-increasing in every period and converges to `-c/delta` if cost is constant — which is inadmissible, meaning the linear model has been pushed outside its range and the real system has hit some floor.

A system with `eta = 0` does not stagnate; it declines, and it declines at exactly the rate it would decline if the interaction had never occurred. The meetings still happen. The conferences still convene. The arithmetic does not notice any of it. This condition is nearly invisible from inside, because `S_t` is large and visible while `eta S_t` is small and abstract: people experience the interaction and do not experience the coefficient.

8.4.2 4.2 The distinction the section turns on

Two different quantities travel under one word.

```
S_t        the interaction that HAPPENED    -- attendance, discussion, encounter, contact
eta*S_t    the interaction that STUCK       -- decision, artifact, default, trained person
```

`S_t` is bounded above by `ln N` and is generated by the geometry of who is present and how different they actually are; `eta` is generated by what is done with what that geometry produced. **The geometry decides how much is available; the conversion decides how much is kept.**

Most organisations measure `S` and report it as though it were `eta S`. Attendance figures, meeting counts, conference days, touchpoints, engagement, workshop hours: these are measurements of `S`, or of proxies for `S`, and none of them is a measurement of `Q`. A report that lists them and concludes that the organisation is building something has performed a substitution the arithmetic does not permit. Stated as a rule: **a report that an interaction went well is a report of `S` and contains no information about `eta`.** The answerable follow-up is not what was discussed but what is different now that was not different before, and where it is written down. If the honest answer is *nothing yet*, then `eta S = 0` for that interaction and it moved `Q` by exactly zero.

8.4.3 4.3 The three gates

`eta` is a capture rate, and in every real system it decomposes into three sequential gates, each a multiplier in `[0,1]`:

```
eta  =  eta_record  *  eta_bind  *  eta_reach
```

- **`eta_record`** — the fraction of what was produced that was recorded at all, in any durable medium. A decision reached and not written is lost at the rate of human memory.
- **`eta_bind`** — the fraction of what was recorded that was bound to a change in the system's behaviour: an owner, a date, a default, a modified procedure, a moved wall. A record with no binding is an artefact in a filing cabinet.
- **`eta_reach`** — the fraction of the bound change that reached the people and places that act. A changed default that no tooling implements has changed nothing.

Badge: this three-gate decomposition is a modelling convenience, not a theorem, and it has no external support. It earns its place by telling an analyst where to look. What matters is the multiplicative form: three gates at 0.7 each give `eta = 0.343`, so the intuition that a system is *doing pretty well on all three* is worth less than half of what it feels like.

The multiplicative form also explains why conversion improvements are so often cheap. If the worst gate is `eta_record = 0.1` because nothing is written down, a scribe rota costs one person's attention per meeting and can plausibly take that gate to 0.8 — an eightfold move on the whole coefficient, with no change to headcount, geometry, or budget.

8.4.4 4.4 Six capture mechanisms

These are the durable-structure channels available to almost every system, listed in rough ascending order of conversion per unit of effort, each described in terms of the parameter it moves.

1. Writing it down (raises `eta_record`). A shared, dated, append-oriented record of what was concluded. Cost: one person, ten minutes. The highest-leverage intervention available in any system with `eta_record < 0.5`, which is most of them. The failure variant is catalogued as F2: minutes that record what was *said* rather than what was *decided* raise `eta_record` while leaving `eta_bind` at zero.

2. The decision log (raises `eta_record` and `eta_bind` together). Distinct from minutes. It records the decision, the date, the alternatives rejected, the person accountable, the review date, and the condition that would reverse it. It is the cheapest device that touches both gates at once, and it also serves the falsification requirement of Curzi [2026i], because a decision log with reversal conditions is a preregistration in the sense that literature requires [Nosek et al., 2018].

3. The changed default (raises `eta_bind` and `eta_reach`). A default is a decision that executes itself: the template that now has the new field, the calendar that now has the slot. **Defaults have the highest reach of any mechanism, because they require no ongoing compliance from anyone.** They are also the case in which Star and Ruhleder [1996]'s observation bites in the organisation's favour: infrastructure is invisible when it works, which is exactly why a default is more durable than a policy — and exactly why, in Section 6.9, working stored order never appears on a list of things that could be removed.

4. The built object (raises all three; resists `delta` mechanically). A jig, a checklist card, a script, a fixture, a moved partition. A built object is stored order that resists decay mechanically rather than through anybody's diligence, which is why the built-environment levers of Curzi [2026f] — whose entrance-sequence case takes interruptions from 22 to 6 per person-day and headroom from `+0.0642` to `+0.2955` for 7.6 per cent of net floor — are conversion devices as much as cost devices.

5. The trained person (highest conversion, highest exposure to `delta`). Skill in a person is the densest form of stored order available and the form most exposed to turnover. The

engineering response is not to avoid training people but to pair this mechanism with the first four, so that a departure is a partial rather than a total loss.

6. The altered institution (highest durability, highest cost, longest lag). Moving a wall, changing a reporting line, creating a standing role, rewriting a charter. These are the slowest and most expensive captures and they have the lowest decay. Conway [1968] is the reason they work: systems mirror the communication structure of the organisation that builds them, so an alteration to that structure propagates into everything the structure subsequently produces. **The source manuscript dates this article to 1967; it ran in *Datamation* in April 1968, and the correction is recorded here rather than made silently.**

8.4.5 4.5 Ordering the mechanisms honestly

The list above is ordered by effort-adjusted leverage in a typical system, and that ordering is a judgement rather than a measurement. **The honest statement is that the ordering is system-specific and should be measured rather than assumed.** What is not system-specific is the structure of the argument: find the lowest gate and fix that one.

Consider a group whose gates are `eta_record = 0.20` (someone occasionally emails notes), `eta_bind = 0.60` (when notes exist, an owner is usually named), `eta_reach = 0.70` (named owners usually follow through), so that `eta = 0.20 x 0.60 x 0.70 = 0.084`. Two candidate interventions of comparable cost:

```
Intervention A: improve eta_bind   0.60 -> 0.90    eta = 0.20*0.90*0.70 = 0.126   (+50%)
Intervention B: improve eta_record 0.20 -> 0.80    eta = 0.80*0.60*0.70 = 0.336  (+300%)
```

Both are process improvements. One is four times the other. The difference is entirely visible in advance if the coefficient is decomposed, and entirely invisible if it is not.

8.4.6 4.6 Governance design principles read as conversion devices

Ostrom [1990] gives eight design principles for durable common-pool resource institutions, derived from field study of institutions that survived for centuries and institutions that failed. They are usually read as governance advice. Read as engineering they are almost entirely conversion and retention devices, and it is worth saying which is which. Cox et al. [2010] is the meta-analytic audit of those principles and is the kind of external test this framework's own principles will eventually face; Hardin [1968] is the claim Ostrom refuted, and the pair should be cited together or not at all.

Design principle	Primary effect in the law	Mechanism
Clearly defined boundaries	raises `eta_bind`; defines `N`	a decision cannot bind to an undefined membership
Rules matched to local conditions	raises `eta_reach`	a rule that does not fit is not executed
Collective-choice arrangements	raises `eta`; protects `u`	those affected participate
Monitoring	raises `eta_reach`; lowers `delta`	detects when stored order has decayed
Graduated sanctions	lowers `delta`; bounds `C`	maintains bindings without maximal enforcement
Conflict-resolution mechanisms	lowers `C`	conflict is a dissipation term
Recognition of rights to organise	raises `eta_bind`	external non-recognition voids the binding

Design principle	Primary effect in the law	Mechanism
Nested enterprises	raises `eta` at scale; bounds `C`	the topology answer of Section 5.7

The most consequential reading is monitoring. In the field data it is not primarily a compliance device; it is the device that makes the state of `Q` *visible*, and a `Q` that is visible can be maintained. Star and Ruhleder [1996]'s observation that infrastructure becomes visible only on breakdown is precisely the failure that monitoring corrects, at a cost, and the cost is a `C` term. Monitoring is therefore a deliberate trade — **spend `C` to lower `delta`** — worth making only when `delta Q` exceeds the monitoring cost, which is computable and is computed in Section 6.

Nested enterprise is the same structural answer that Section 5.7 reaches from the cost side and that the semilattice argument reaches from the surplus side. **Three independent lines arriving at one topology is the strongest form of agreement this paper can offer, and it is still not a proof.** Scott [1998] supplies the standing warning attached to it: legibility destroys the local variety that made a system work, and a cellular scheme imposed from outside is a legibility project.

8.4.7 4.7 Apprenticeship, and the throughput trade

Apprenticeship holds `S_t` at a moderate level continuously over a long horizon with `eta_reach = 1` by construction, because the person receiving the transfer is the person who will act. Writing has a reach problem — the reader may not be the actor — and a tacit-knowledge problem, since what cannot be articulated is not recorded. Apprenticeship has neither, because the transfer channel is demonstration rather than description.

What it has instead is a throughput problem: one transfer per master per period, where a written artefact serves many. **That is the exact trade — apprenticeship buys very high conversion at very low throughput; documentation buys much lower conversion at very high throughput** — and a system using only one of them leaves the other's regime unserved. The design rule that follows is not *do apprenticeship*. It is: **for the fraction of stored order that is tacit, documentation has `eta_record` near zero regardless of effort.** Identify that fraction honestly and give it a demonstration channel.

8.4.8 4.8 What `eta = 0` looks like from inside

Named symptoms, so the condition can be recognised before the arithmetic is run.

- The same decision is re-litigated at intervals, and each re-litigation feels novel.
- Nobody can say when a standing practice was adopted or by whom.
- A new member's onboarding consists of a person telling them things.
- Meeting outputs are described in terms of mood rather than of change.
- Things are written down but cannot be found — `eta_record > 0` with `eta_reach` near zero.
- Departures cause loss disproportionate to the departing person's visible role.

The re-litigation symptom is the most diagnostic, because it is the direct observable signature of `Q` decaying back below the threshold at which a decision counts as settled. It is also countable, and Section 7 makes it a standing requirement.

8.4.9 4.9 Measuring conversion

The definitional estimator is `eta_hat = delta_Q_attributable / S_input`. Both terms are hard, and the procedure names both difficulties rather than hiding them.

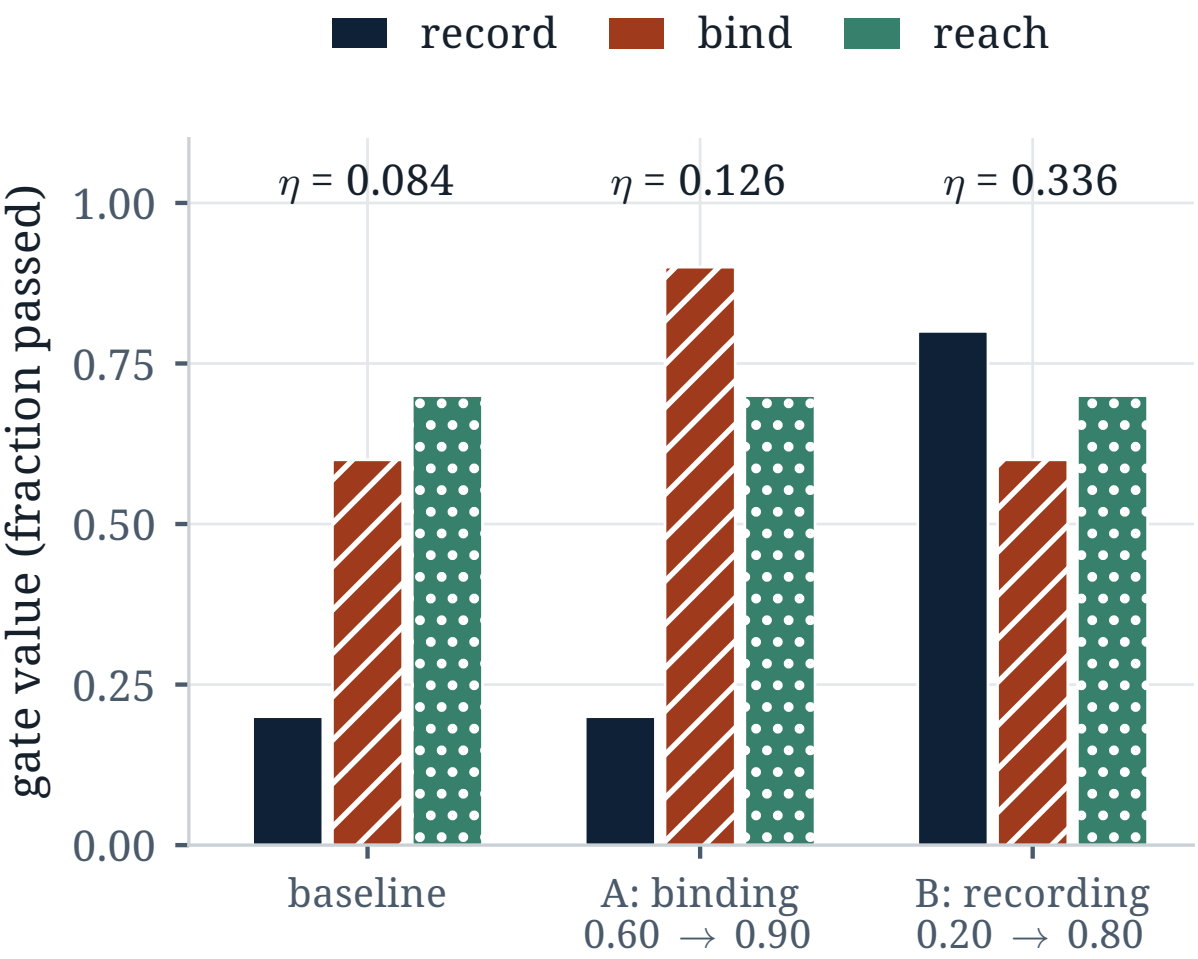


Figure 8.1: **The weakest gate dominates: the same 0.30 improvement is worth +50 per cent applied to binding and +300 per cent applied to recording**, because conversion is a product and not a mean. Gate values are stipulated illustrative inputs; the products are exact. Generated by `figs/src/CCP-008_fig02.py`.

Step 1 — choose the order observable `Y_t`. A measurable proxy for `Q_t`: count of live written decisions with named owners; count of standing defaults implemented in tooling; number of procedures a randomly sampled member can execute unaided; count of built artefacts in active use.

Step 2 — justify that `Y_t` is not a repackaging of `U_t`. This is the most common measurement error in the framework. If the order measure is really a measure of how much interaction is happening, the input has been fitted to itself and the estimate is meaningless. **The short form of the test: if the order proxy would move on a week when a great deal of talking happened and nothing was concluded, it is an interaction measure and must be replaced.** The cheapest detector of the failure is the ceiling test, since no admissible single-period increment can exceed `eta ln N`: a repackaged proxy in the corpus's one applied survey produced a single-week jump of 31.1 units against a maximum admissible increment of 0.8071, a factor of 38.5 [Curzi, 2026i].

Step 3 — measure `S_input`. Estimate `N` by the effective-count audit [Curzi, 2026c] and `u` by disagreement rate or error decorrelation, then compute `S = ln(1 + (N−1)u)` directly. Do not estimate `S` from attendance.

Step 4 — establish attribution. The minimum acceptable design is comparison against four baselines: persistence, direct-input, matched autoregressive, and shuffled-input. **The shuffled-input control is the one that catches a false conversion estimate: if the pairing of periods is shuffled and the fit does not degrade, what has been measured is a trend and not a conversion.** Simmons et al. [2011] and Gelman and Loken [2013] are the reason this is not optional; Ioannidis [2005] and Open Science Collaboration [2015] are the reason a single uncontrolled estimate should not be believed.

Step 5 — fit the minimal model `Y_{t+1} = a Y_t + b F(U_t) − C_t + eps_t` with `0 <= a <= 1` and `b >= 0`, so that `a = 1 − delta` and `b = eta`. Fit on training data, **freeze**, and report held-out performance.

A worked estimation table. Take a group with `N = 8` and `u = 0.30`, so `S = ln(1 + 7 x 0.30) = ln 3.10 = 1.1314` nats per period, with `Y` normalised so that one unit is one durable, owned, implemented decision still live at review:

Capture regime	`delta_Q` per period	`eta_hat = delta_Q / 1.1314`
No artefact of any kind	0.000	0.000
Verbal agreement only	0.010	0.009
Notes circulated after the fact	0.050	0.044
Decision log with owner	0.140	0.124
Log, owner and changed default	0.280	0.247
Built artefact and trained person	0.450	0.398

Badge: the left column is a set of illustrative inputs chosen to span a plausible range, not measurements from any study. The arithmetic in the right column is exact; the left column is what an analyst must measure in their own system. No citation is offered for a number that was not measured. What the table demonstrates is the *shape* of the result — a factor of forty between the bottom and top regimes, produced entirely by what happens in the twenty minutes after an interaction ends.

8.4.10 4.10 The asymmetry: why conversion is the first move

This is the quantitative core of the section. `eta` enters the budget linearly, `N` logarithmically, and `u` through a concave saturating function. The consequence is a set of exact equivalences.

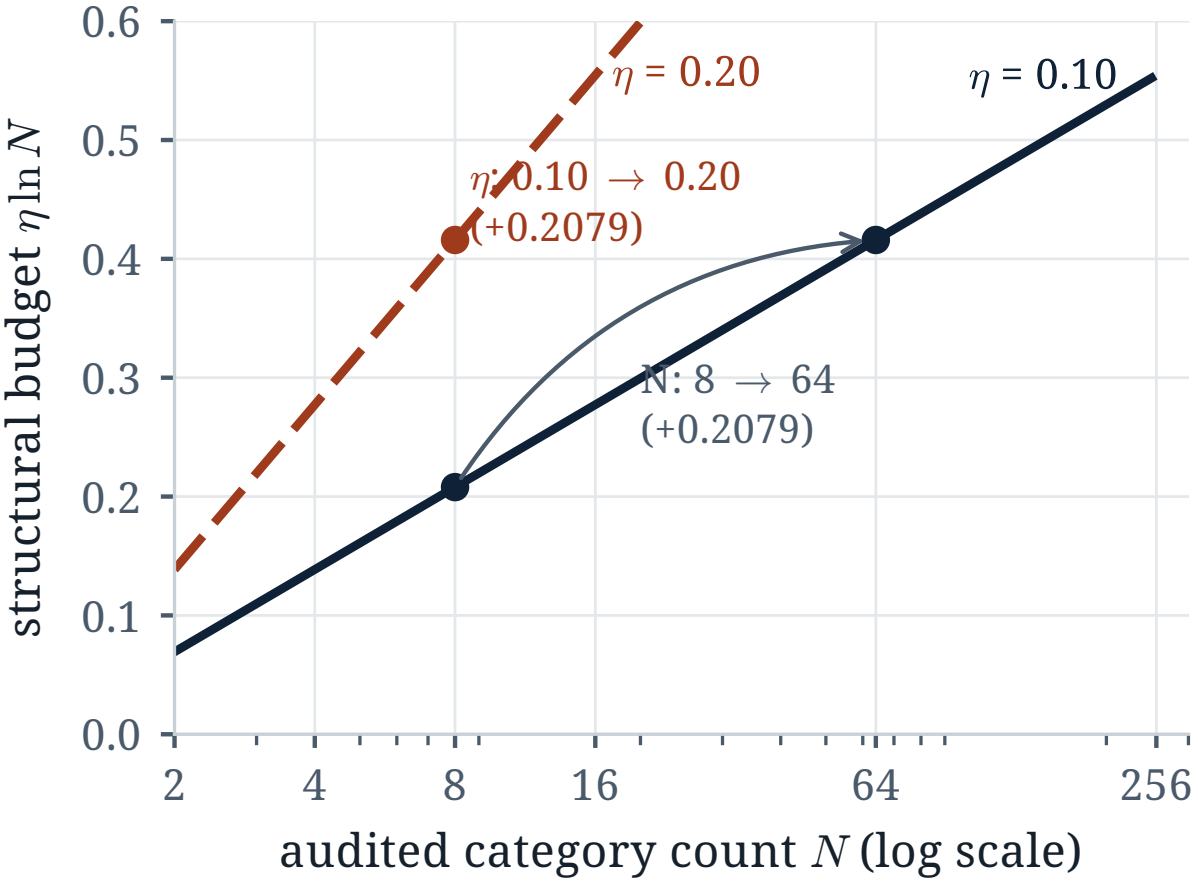


Figure 8.2: **Doubling conversion and squaring the category count buy exactly the same budget**: at `eta = 0.10` moving `N` from 8 to 64 adds 0.2079, and so does moving `eta` from 0.10 to 0.20 at `N = 8`. Modelled from the budget term `eta ln N`; the two parameter values are stipulated and the arithmetic is exact. Generated by `figs/src/CCP-008_fig01.py`.

4.10.1 Doubling `eta` is exactly squaring `N`

The budget term is `eta ln N`. Doubling `eta` gives `2 eta ln N = eta ln(N^2)`. Therefore **to match a doubling of conversion by adding categories, the category count must be squared.**

Current `N`	`N` required to match a doubling of `eta`
2	4
4	16
8	64
10	100
16	256
20	400
50	2,500
100	10,000

An organisation with eight genuinely distinct disciplines that wants to double its structural budget can install a decision log, or find fifty-six more genuinely distinct disciplines.

4.10.2 Doubling `eta` is often impossible to match by raising `u`

To double `S = ln(1 + (N−1)u)` by raising `u` alone requires `u'` with `1 + (N−1)u' = (1 + (N−1)u)^2`, so

```
u' = ( (1 + (N−1)u)^2 − 1 ) / (N−1)
```

Since `u <= 1` by construction, this is feasible only when `(1 + (N−1)u)^2 <= N`, that is when

```
u  <=  1 / (sqrt(N) + 1)
```

Above that threshold, doubling the surplus by raising non-parallelism is not merely expensive: it is arithmetically impossible, because there is not enough room left below `u = 1`.

`N`	maximum `u` from which doubling via `u` remains feasible
2	0.4142
3	0.3660
4	0.3333
5	0.3090
8	0.2612
10	0.2403
16	0.2000
20	0.1827
50	0.1239
100	0.0909

Worked instances:

```
N=8,  u=0.30 -> S = 1.1314;  doubling needs u' = 1.2300   INFEASIBLE (u' > 1)
N=8,  u=0.10 -> S = 0.5306;  doubling needs u' = 0.2700   feasible, and expensive
N=20, u=0.10 -> S = 1.0647;  doubling needs u' = 0.3900   feasible, and very expensive
```

Raising `u` from 0.10 to 0.27 in a real team means recruitment across training lineages, rotation, adversarial roles, the deliberate outsider — a multi-year programme with a high failure rate. Doubling `eta` from 0.084 to 0.168 by fixing the recording gate is a rota.

4.10.3 The full comparison at one base case

Base: `N = 8`, `u = 0.30`, `eta = 0.10`, `delta = 0.10`, `C = 0.05`. Then `S = 1.1314` and the structural budget is `eta ln N = 0.10 x 2.0794 = 0.2079`.

Move	new `eta ln N`	gain	realistic cost
baseline	0.2079	—	—
`eta`: 0.10 to 0.20	0.4159	+0.2079	a scribe rota and a decision log
`N`: 8 to 10	0.2303	+0.0223	two genuinely new disciplines hired
`N`: 8 to 16	0.2773	+0.0693	doubling the discipline count

Move	new `eta` `ln N`	gain	realistic cost
`N`: 8 to 64	0.4159	+0.2079	eight times the discipline count

Hiring eight times as many distinct disciplines buys exactly what a decision log buys. This is the asymmetry stated as a single line, and it is why conversion is the first move in almost every system this framework has been applied to. The result is consumed downstream as the sensitivity table `dh/d(eta) = ln N = 2.197` against `dh/dN = eta/N = 0.0038`, which records adding categories as the weakest available lever [Curzi, 2026o].

4.10.4 Cadence is not a conversion move and must not be booked as one

A fourth move is proposed in the same breath as the three above — meet more often — and it belongs in a different column of the ledger. **For a linear time-invariant integrator the integrated effect of one large input and of the same total delivered periodically is identical** [Oppenheim et al., 1996]. Rescheduling delivers no extra input, and a cadence change booked as a gain in `eta` is double-counting.

What cadence does buy is real and arrives through the *surplus* term. Because `f(u) = ln(1 + (N−1)u)` is strictly concave with `f(0) = 0`, splitting a fixed budget of non-parallelism across more occasions strictly increases the delivered total. At `N = 8`:

occasions `n`	`u` per occasion	`f(u)`	delivered total `n f(u)`
1	0.60	1.6487	1.6487
2	0.30	1.1314	2.2628
4	0.15	0.7178	2.8714
6	0.10	0.5306	3.1838

Four occasions deliver **1.74 times** what one delivers from the same budget, and separately the ceiling `S <= ln N` clips whatever a single occasion tried to exceed. **Cadence beats intensity because of concavity and the ceiling, not because it delivers more input.** Book the gain under `S`, keep `eta` for the capture gate, and the two moves stay separately auditable.

Two cautions, both of which the source manuscript leaves implicit and this paper states. First, the table treats a session's worth of non-parallelism as a divisible budget that can be partitioned across occasions. **`u` is a bounded geometric parameter — the squared sine of an angle between unit state vectors — and nothing in the framework establishes that it is extensive in this way.** The table is therefore correct arithmetic about `f` and rests on an unstated postulate about `u`, which is badged here as a modelling choice. Second, this concavity argument runs in the opposite direction from the concavity correction of Section 6.2, where *variability* in `U` at fixed mean reduces the delivered mean of `F` below `F` of the mean. Both are Jensen's inequality; they differ in what is held fixed, and confusing them produces an error in a predictable direction.

8.4.11 4.11 Where conversion cannot be raised

Three honest limits.

The tacit fraction. Some stored order is not articulable, and documentation has `eta_record` near zero on it by nature. The answer is a demonstration channel, not more effort.

The capture tax. Capture consumes attention, and attention is a `C` term. A system that captures everything spends its entire budget on capture. The marginal condition is computable: capture is worth doing on an interaction while the marginal `d(eta) S` exceeds the marginal

dC. In practice this means **capture the decisions, not the discussion** — the discussion is S, and only the decision can become Q.

The ceiling is still ln N. Raising conversion scales the whole budget, but the budget is eta ln N, and a high conversion rate applied to nearly nothing is nearly nothing. **A multiplier cannot rescue a zero.** Where three ostensibly independent sources share one doctrine, u = 0, N_eff = 1, and S = ln 1 = 0, and every conversion intervention in this section returns nothing while consuming C. This is the case Curzi [2026] states as the binding clause of the functional: a second opinion from a source sharing training, incentives and premises is worth ln(1) = 0 of surplus. It gives the correct sequencing rule:

Verify S > 0 before investing in conversion. If u is near zero, fix the geometry first. Otherwise fix conversion first, and fix it before considering N.

8.4.12 4.12 Communication structure as a conversion channel

Conway [1968] observed that a system's structure mirrors the communication structure of the organisation that produced it. In the terms of the law this is a statement about eta_bind and eta_reach: **the communication topology is the channel through which captured order reaches the artefact, so the artefact inherits the topology.**

Two consequences. First, if an artefact with a particular structure is wanted, the cheapest route is to give the producing organisation that structure, because reach along the existing topology is high and reach across it is low. Fighting the topology costs C; using it is free. Second, and less often noticed, **a communication topology change is a high-conversion intervention**, because it propagates automatically into everything subsequently produced. The altered institution is expensive precisely because it is powerful.

8.4.13 4.13 High conversion on the wrong thing

eta is neutral. It converts whatever surplus arrives into whatever durable structure the capture mechanism specifies, and a high-conversion system that captures wrong conclusions builds durable wrong structure — which by construction has a low delta and therefore persists. This is the mirror image of the retention result that a system with no forgetting cannot revise [Curzi, 2026a]. Banach [1922] guarantees arrival at the fixed point and says nothing whatever about whether the fixed point is a good place; the convergence analysis of this corpus makes the same refusal explicitly, holding that the theorem transfers as shape rather than as guarantee because nobody has measured the relevant contraction modulus and a human life is not a metric space [Curzi, 2026].

The engineering answer is not to lower conversion. It is to build reversal into the capture format, which is why the decision-log specification requires *the condition that would reverse it.* **A capture format that records its own falsification condition converts surplus into structure that can be dismantled on evidence, which is the only kind of high-conversion system that remains correctable.**

The measurement is exposed to the failure it describes. A measure that becomes a target ceases to be a measure [Goodhart, 1975, Campbell, 1979], so if conversion is measured and the measurement becomes a target, capture ceremonies will multiply and the measured eta will rise while the real one does not. The adversarial form of the order-proxy test is the defence: **ask whether the conversion metric would rise if people simply wrote more documents that changed nothing. If it would, it is not measuring conversion.**

8.4.14 4.14 Minutes and a decision log

Worth stating separately, because the confusion is close to universal.

```
MINUTES                          DECISION LOG
```

```
records what was SAID               records what was DECIDED
chronological                       indexed by decision
complete                            selective
raises eta_record only              raises eta_record AND eta_bind
read once, by nobody                read on every re-litigation
grows without bound                 grows with decision count
no owner field                      owner field is mandatory
no reversal condition               reversal condition is mandatory
```

Minutes are an archive; a decision log is a control surface. **The test is whether the document is consulted when someone asks why a practice is what it is. If it is not, the record is minutes and `eta_bind` is unaffected by it.**

8.5 5 Cost: the dissipation budget

`C_t` removes stored order in a way that has nothing to do with decay. `delta` is what leaks out through the walls; `C_t` is what is drawn off deliberately — the overhead of being an organisation at all, paid every period whether or not anything is accomplished.

Two constraints govern the section and they are not the same line.

```
admissible :   C_t <=  (1-delta)*Q_t + eta*S_t     -- Q_{t+1} stays >= 0
sustainable:   C_t <=  eta*ln N - delta*Q_t         -- Q_{t+1} >= Q_t
```

Admissibility is a floor condition: the model does not produce nonsense. Sustainability is a growth condition: order does not decline. **Admissibility permits eating capital; sustainability does not.** Critically, `Q` appears with opposite signs in the two: a large stock of stored order *raises* the admissible cost, because it can be burnt, and *lowers* the sustainable cost, because it must be maintained. Worked at `delta = 0.08`, `eta = 0.22`, `N = 8`, `u = 0.60`, so `S = 1.6487`, `eta S = 0.3627` and `eta ln N = 0.4575`:

`Q_t`	admissible `C_max`	sustainable `C_max`
0.0	0.3627	+0.4575
0.5	0.8227	+0.4175
1.0	1.2827	+0.3775
2.0	2.2027	+0.2975
4.0	4.0427	+0.1375
8.0	7.7227	-0.1825

The two columns move in opposite directions, and at `Q = 8` the sustainable ceiling has gone negative while the admissible ceiling stands at 7.72. A large institution can therefore run for a very long time in a state that is admissible and unsustainable, consuming accumulated order at a rate that looks affordable because there is so much of it. That is the structural picture of long institutional decline and it is the subject of Section 6.

8.5.1 5.1 What c contains

`C` is not one thing. It is tallied as a sum of named lines, because an aggregate figure cannot be attacked and a decomposed one can.

C1 · Coordination. The direct cost of aligning people: meetings, status reporting, scheduling, handoffs, approval chains. This is the line that scales with path count.

C2 · Friction. Cost incurred because the system is arranged badly rather than because the work is hard: waiting for access, re-entering data, tool mismatch, permission bottlenecks. Friction produces nothing at all, not even coordination.

C3 · Maintenance of relationship. The upkeep that keeps `u` from collapsing into hostility or `N` from fragmenting. **This is the `C` line most often cut first and it is often the one that should be cut last**, because cutting it raises `delta` on trust, which is stored order.

C4 · Conflict. Disputes that consume attention without resolving. Conflict and disagreement are *not* the same thing: disagreement is `u` and is on the credit side of the ledger; conflict is `C`. **A design that suppresses conflict by suppressing disagreement has paid for a `C` reduction with a `u` reduction, and since `u` enters through a logarithm and `C` enters linearly, that trade is usually bad.** Acoustic treatment and table geometry are, in part, devices for having disagreement without conflict [Curzi, 2026e,f].

C5 · Rent and premises. Literal rent and everything that behaves like it: licence fees, platform tolls, the fixed cost of existing.

C6 · Commute and transit. The cost of getting bodies into proximity. Its reduction trades directly against `S`, because proximity is one of the cheapest ways to raise interaction rate. A genuine trade, not a free cut.

C7 · Attention labour. The cost of managing presentation, absorbing others' states, and maintaining the social surface. Real, chronically unmeasured, and unevenly distributed. It is a `C` term because it consumes exactly the resource — attention — that conversion also consumes.

C8 · Compliance and reporting. Externally imposed overhead. Often the least compressible line and the one that grows with `Q` (Section 5.11).

C9 · Context switching and interruption. Each switch costs recovery time. Usually invisible in a budget and frequently the largest single item in knowledge work; the built-environment treatment prices it as a threshold problem and measures a fall from 0.321 to 0.087 of the working day [Curzi, 2026f].

8.5.2 5.2 The pair count and what it does not prove

In a group of `n` people the number of distinct unordered pairs is

```
P(n) = n(n−1)/2
```

This is arithmetic and therefore not a claim about the world. The interpretive step — that coordination burden tracks `P` — is Brooks [1975]'s observation from practice, and this paper adopts it as a **badged modelling choice** in the form `C1 = kappa P(n)`, with `kappa` a per-link cost per period. It is chosen because it is the simplest form that reproduces the observed superlinearity and because it makes the crossover computable. **It is falsifiable in principle: measure `C` at several `n` in a comparable setting and check the exponent. Section 11.7 records that the log-log regression the corpus's own protocol proposes for that check is not a sound procedure, and states what would be.**

The count itself, alongside the surplus ceiling for the most favourable possible case in which every added person brings a genuinely new category:

`n`	`n(n−1)/2`	`ln n`	paths per nat
2	1	0.6931	1.44
4	6	1.3863	4.33
6	15	1.7918	8.37
8	28	2.0794	13.47
10	45	2.3026	19.54
12	66	2.4849	26.56

n	n(n−1)/2	ln n	paths per nat
15	105	2.7081	38.77
20	190	2.9957	63.42
30	435	3.4012	127.90
50	1,225	3.9120	313.14
100	4,950	4.6052	1,074.88
150	11,175	5.0106	2,230.26

Read the last column as the price of a nat of surplus ceiling, denominated in communication paths. At `n = 6` a nat costs 8.37 paths; at `n = 150` it costs 2,230.26. **The price has risen by a factor of 266 while the thing purchased has risen by a factor of 2.8.**

8.5.3 5.3 The divergence, stated per added person

The marginal view is more severe than the total view and it is the one that should govern hiring. Adding one person to a group of `n` creates `n` new paths and raises the surplus ceiling by `ln(n+1) − ln n`:

n to n+1	new paths	gain in ln n	paths per nat of gain
3 to 4	3	0.28768	10.4
6 to 7	6	0.15415	38.9
10 to 11	10	0.09531	104.9
20 to 21	20	0.04879	409.9
50 to 51	50	0.01980	2,524.9
100 to 101	100	0.00995	10,049.9
150 to 151	150	0.00664	22,574.9

The marginal price of surplus rises quadratically in `n`, because the numerator grows like `n` and the denominator falls like `1/n`. From the fourth member to the hundred-and-fifty-first the price of an additional nat rises by a factor of two thousand. This is not a metaphor and not an approximation: `n(n−1)/2` is exactly quadratic under the modelling choice of Section 5.2 and `ln n` is exactly logarithmic, and any quadratic overtakes any logarithm at a point that depends only on `eta` and `kappa`.

8.5.4 5.4 The crossover

Question. For given `eta` and `kappa`, at what group size does coordination cost `C = kappa n(n−1)/2` exceed the entire structural budget `eta ln N`?

Take the most favourable case for growth, `N = n`, so that every person added brings a genuinely distinct category. This is the best case and it is almost never true; Section 5.6 handles the realistic case, in which the crossover arrives earlier and the optimum much earlier. Solve for the smallest integer `n >= 2` with `kappa n(n−1)/2 > eta ln n`:

kappa \ eta	0.10	0.20	0.30	0.50	0.80	1.00
0.0005	39	58	73	97	125	142
0.0010	27	39	49	66	85	97
0.0020	18	27	33	45	58	66
0.0050	11	16	20	27	35	39

kappa eta	0.10	0.20	0.30	0.50	0.80	1.00
0.0100	7	11	13	18	23	27
0.0200	5	7	9	12	16	18
0.0500	3	4	5	7	9	11
0.1000	2	3	4	5	6	7

`kappa` is the cost per active communication path per period, in the same normalised order-units as `Q`. The entry is the group size at which coordination alone consumes the entire structural budget — the size at which, even with zero maintenance draw and zero other cost, headroom has gone negative and no geometric remedy remains.

First: the crossover is an upper bound on viable size under complete-graph obligation, and it is a low number. For mid-range values `eta = 0.30` and `kappa = 0.010` the ceiling is thirteen people — not thirteen as a rule of thumb, but thirteen as an arithmetic consequence of two parameters.

Second: conversion buys size, and it buys it slowly. Along the `kappa = 0.010` row, a tenfold increase in `eta` from 0.10 to 1.00 moves the crossover from 7 to 27, a factor of 3.9 rather than 10. Setting `kappa n^2/2 ~ eta ln n` and ignoring the slowly varying logarithm gives `n ~ sqrt(eta/kappa)` up to a log factor, so **the crossover scales roughly as the square root of conversion.** Doubling `eta` buys about 41 per cent more headcount, not 100 per cent more.

Third: `kappa` buys size on the same square-root terms and is far easier to move. Down the `eta = 0.30` column, a hundredfold reduction in `kappa` from 0.10 to 0.001 moves the crossover from 4 to 49. And `kappa` is reducible by topology, by protocol, by asynchrony and by co-location, none of which requires changing anyone's capability. **The cheapest way to be allowed to be bigger is to make each link cheaper, not to make each person better.**

8.5.5 5.5 The crossover is not the optimum

The crossover answers *when does cost eat the whole budget*, which is the wrong question for sizing. The right question is what `n` maximises net structural gain. Maximise `net(n) = eta ln n − kappa n(n−1)/2`:

```
d/dn [ eta*ln n − kappa*(n^2 − n)/2 ]  =  eta/n − kappa*(2n − 1)/2  =  0

  =>   eta = kappa * n * (n − 1/2)   ~   kappa * n^2

  =>   n_opt  ~  sqrt( eta / kappa )
```

Design rule: the net-optimal group size under complete-graph obligation is approximately the square root of the ratio of conversion efficiency to per-link cost. Exact integer optima by direct search, with the net value in parentheses:

kappa eta	0.10	0.20	0.30	0.50	0.80	1.00
0.0005	14 (+0.218)	20 (+0.504)	25 (+0.816)	32 (+1.485)	40 (+2.561)	45 (+3.312)
0.0010	10 (+0.185)	14 (+0.437)	18 (+0.714)	23 (+1.315)	29 (+2.288)	32 (+2.970)
0.0020	7 (+0.153)	10 (+0.371)	13 (+0.613)	16 (+1.146)	20 (+2.017)	23 (+2.629)
0.0050	5 (+0.111)	7 (+0.284)	8 (+0.484)	10 (+0.926)	13 (+1.662)	14 (+2.184)
0.0100	3 (+0.080)	5 (+0.222)	6 (+0.388)	7 (+0.763)	9 (+1.398)	10 (+1.853)
0.0200	3 (+0.050)	3 (+0.160)	4 (+0.296)	5 (+0.605)	7 (+1.137)	7 (+1.526)
0.0500	2 (+0.019)	2 (+0.089)	3 (+0.180)	3 (+0.399)	4 (+0.809)	5 (+1.109)
0.1000	2 (-0.031)	2 (+0.039)	2 (+0.108)	3 (+0.249)	3 (+0.579)	3 (+0.799)

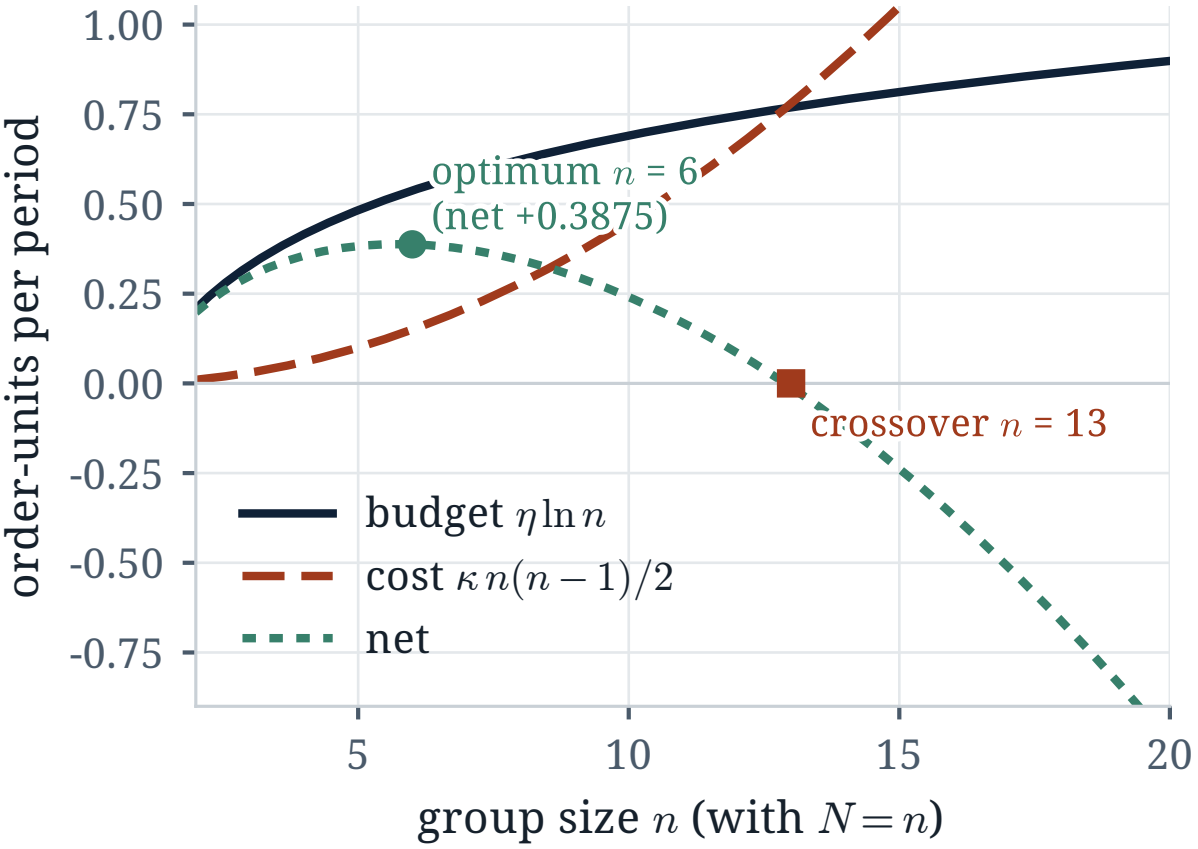


Figure 8.3: **The net optimum sits at six and the crossover at thirteen**, so the band in which every additional member reduces net output is more than twice the optimal size. Modelled from `net(n) = eta ln n − kappa n(n−1)/2` at the stipulated `eta = 0.30`, `kappa = 0.010`, `N = n`. Generated by `figs/src/CCP-008_fig03.py`.

The closed form tracks the search closely: at `eta = 0.30`, `kappa = 0.005`, `sqrt(0.30/0.005) = 7.7` against a searched optimum of 8; at `eta = 1.0`, `kappa = 0.001`, `sqrt(1000) = 31.6` against a searched optimum of 32.

Compare the two tables at a common cell. **At `eta = 0.30` and `kappa = 0.010` the optimum is 6 and the crossover is 13.** A team of thirteen is not at capacity; it is at twice its net-optimal size and delivering less than a team of six. Between 6 and 13 the group remains nominally viable — headroom may be positive — while every additional member reduces net structural output.

This gap is where most organisations live, and it is invisible without the calculation, because nothing breaks at the optimum. The failure is silent and gradual until the crossover, at which point it becomes sudden and is misdiagnosed as a leadership problem. That the failure is undetectable from inside is the point at which Sterman [1989]'s measured result — that people misjudge stocks with delays even under aligned incentives — stops being a curiosity and becomes the mechanism.

Note the bottom-left cell: at `eta = 0.10` and `kappa = 0.10` the optimum over `n >= 2` is `n = 2` and the net value is **negative**, against a net of exactly zero for `n = 1`. There exist parameter regimes in which no group of any size does better than nobody working together, because the per-link cost exceeds what any amount of difference can pay for. That is not a rhetorical flourish; it is what the arithmetic returns.

8.5.6 5.6 The realistic case: `N` saturates while `n` grows

The tables above assume `N = n`. Real groups exhaust their genuinely distinct categories quickly; beyond that point additional people add paths without adding `N`, and the effective-count audit shows `N_eff` flat while headcount climbs [Curzi, 2026c]. Model it as `N = min(n, D)` with `D` the number of genuinely distinct categories available, at `eta = 0.30`, `kappa = 0.010`:

`D`	surplus ceiling `eta ln D`	crossover `n*`	net-optimal `n`	net at optimum
3	0.3296	9	3	+0.2996
5	0.4828	11	5	+0.3828
8	0.6238	12	6	+0.3875
12	0.7455	13	6	+0.3875
20	0.8987	13	6	+0.3875

Look at the last three rows. Raising the available category count from 8 to 20 — a substantial and expensive achievement — moves the optimum not at all and moves the crossover from 12 to 13. Once the quadratic cost term binds, extra categories are irrelevant because the people who would carry them are unaffordable. The surplus side reaches the same conclusion independently and states it as: a headroom problem cannot be escaped by adding domains [Curzi, 2026c]. Here it appears from the cost side, and the two derivations do not share a premise.

Note also that when `D` is small the optimum sits exactly at `D`. **When categories are scarce, hire exactly as many people as there are categories and stop.** The eleventh person on a five-category problem is pure `kappa`.

8.5.7 5.7 Topology: the complete graph is a choice

Everything above assumed `P = n(n−1)/2`, every pair maintaining an active link. That is a choice, not a law, and it is the most expensive available choice. Path counts under five wiring plans, priced at `kappa = 0.006`:

`n`	complete	ring	star	tree	cells of 6 plus liaisons
6	0.0900	0.0360	0.0300	0.0300	0.0900
12	0.3960	0.0720	0.0660	0.0660	0.2040
18	0.9180	0.1080	0.1020	0.1020	0.3420
24	1.6560	0.1440	0.1380	0.1380	0.5040
30	2.6100	0.1800	0.1740	0.1740	0.6900
48	6.7680	0.2880	0.2820	0.2820	1.3920

The cells column assumes complete graphs of six within each cell plus liaison links between cells priced at four times a normal link, because a cross-boundary link genuinely is more expensive, requiring translation, context transfer and schedule reconciliation. **At `n = 48` the complete graph costs 6.7680 and the cellular topology 1.3920 — a factor of 4.86, purchased entirely by deciding who is not obliged to keep a live channel with whom.**

Topology is not free in the surplus term, and the honest treatment says so. The star and the tree are cheap because they route everything through a hub, and a hub drives `u` toward zero, because every message passes through a filter with a single viewpoint. Barabási and Albert [1999] shows that preferential attachment produces hubs spontaneously and Albert et al. [2000] shows such networks are robust to random failure and fragile to targeted attack — with Doyle et al. [2005] as the necessary corrective, since the real Internet's robustness

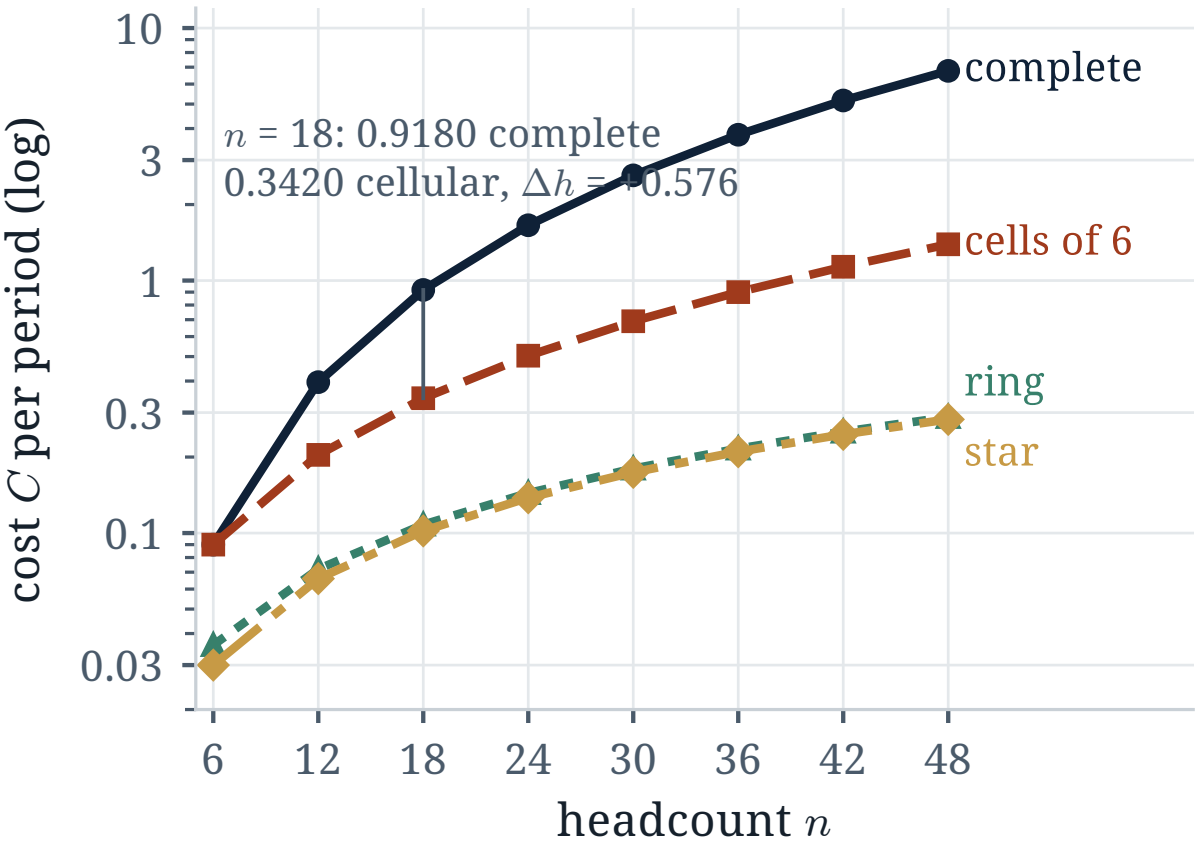


Figure 8.4: **Coordination cost under four wiring plans at `kappa = 0.006`**, on a logarithmic scale because the complete graph and the star differ by more than an order of magnitude at `n = 48` (6.7680 against 0.2820). Modelled; the per-link cost and the fourfold liaison multiplier are stipulated. Generated by `figs/src/CCP-008_fig04.py`.

comes from engineered design rather than from a degree distribution. **A star topology buys a cheap `C` at the price of `U` and of resilience. That is a real trade and must be priced rather than assumed away.**

The cellular arrangement is the one this paper recommends, and the recommendation converges from four independent directions: from cost, since cells cap the quadratic term at cell size and make total `C` linear in cell count; from the nested-enterprise principle observed in long-lived institutions [Ostrom, 1990, 2010]; from the semilattice argument that overlapping, non-exclusive groupings carry structure a strict tree destroys [Alexander, 1965]; and from complex contagion, since behaviour requiring reinforcement spreads along dense local ties rather than weak long ones [Centola and Macy, 2007, Centola, 2010]. **Four independent literatures arriving at one topology is the strongest agreement this paper can offer, and it is still not a proof.**

The concrete rule:

```
Keep complete-graph obligation inside cells of roughly n_opt = sqrt(eta/kappa).
Connect cells by explicit liaison links, budgeted at a higher kappa.
Let cell membership OVERLAP -- that is what makes it a semilattice rather than a tree.
Never route all inter-cell traffic through a single hub.
```

One caution the convergence does not remove. Buldyrev et al. [2010] show that coupled networks are far more fragile than either network alone, so cells joined by a thin set of liaison links are not simply a cheaper complete graph; they are a coupled system with its own failure mode. Perrow [1984] names the general condition — interactive complexity plus tight coupling makes catastrophic failure a system property — and Sagan [1993] documents settings in which redundancy *reduces* safety. Section 12 records that this paper prices the cost saving and does not price that fragility.

8.5.8 5.8 Queueing: slack as a cost-control device

Every shared resource is a queue: the approver, the reviewer, the one person who knows the deployment procedure, the single instrument. For an M/M/1 queue with utilisation `rho`, expected waiting time in queue measured in units of service time is `rho/(1 - rho)` [Erlang, 1909, Kingman, 1962]:

`rho`	wait, in service times	multiple of the wait at `rho = 0.5`
0.10	0.111	0.11
0.30	0.429	0.43
0.50	1.000	1.00
0.70	2.333	2.33
0.80	4.000	4.00
0.90	9.000	9.00
0.95	19.000	19.00
0.99	99.000	99.00

Between 90 and 99 per cent utilisation — an 11 per cent gain in apparent efficiency — the wait multiplies by eleven. Little's Law then gives that queue length multiplies by the same factor [Little, 1961], so work in progress balloons, which raises context-switching cost (C9) and coordination cost (C1) simultaneously. Hopp and Spearman [2011] is the standard treatment.

Therefore waiting is a `C` term, and driving utilisation toward one converts idle capacity — which costs nothing in the law — into waiting, which costs a great deal. Slack is not waste; it is the instrument that keeps the queueing term out of `C`. If waiting costs `kappa_delay` per unit of service time waited, the delay contribution is `kappa_delay rho/(1-rho)` per item. At `kappa_delay = 0.02`:

```
rho = 0.50  ->  0.0200 per item        rho = 0.90  ->  0.1800
rho = 0.70  ->  0.0467                 rho = 0.95  ->  0.3800
rho = 0.80  ->  0.0800                 rho = 0.99  ->  1.9800
```

Against a structural budget of order 0.4 to 0.5, typical of the worked cases below, **a single resource at 99 per cent utilisation can exceed the entire budget by itself. A named individual who is the sole route to a necessary approval, running at near-full load, is a sustainability-ceiling violation wearing the costume of a productivity success.** This is Ashby [1956]'s requisite variety in a different costume: only variety absorbs variety, and a fully loaded server has none left to absorb anything.

8.5.9 5.9 Estimating `kappa`

Step 1 — define an active link. Two people constitute an active link in a period if maintaining mutual awareness of each other's work required time in that period. Not *could have talked* — did have to.

Step 2 — count them. For small groups, ask. For larger ones, sample: take ten people at random, ask each how many colleagues they had to stay synchronised with this period, and estimate the total as `n x mean / 2`. **The instrument is a self-report survey and inherits every known exposure of one: question wording and focusing effects move the answers [Schwarz, 1999, Kahneman et al., 2006], and the estimate should carry an interval rather than a point.**

Step 3 — price them. Estimate hours per period spent on synchronisation, divide by the link count, and convert into the same normalised order-units as `Q`. **The unit conversion is the delicate step and it is where estimates go wrong: `kappa`, `eta` and `Q` must all be in one consistent system or the crossover tables are meaningless.** Section 11.1 states the dimensional requirement exactly.

Step 4 — check the exponent, and do not check it by regressing logs on logs. If `C` can be observed at several group sizes in comparable conditions, the question is whether the growth is nearer quadratic than linear. **The source manuscript prescribes regressing `log C` on `log n` and reading the slope. That is precisely the procedure Clauset et al. [2009] shows to be unreliable, and the objection is not confined to power laws in the tail: a straight line on a log-log plot is weak evidence for the exponent that generated it, and Stumpf and Porter [2012] shows many mechanisms produce indistinguishable curves.** The defensible version compares the quadratic and linear models by likelihood on held-out data with an information criterion, states the number of free parameters, and reports the comparison rather than a slope. **This paper does not run that comparison, because it fits nothing; it records the requirement in Section 11.7 and in Section 12.**

A cross-check requiring no instrumentation is the raw meeting arithmetic:

`n`	90-minute weekly meeting	person-hours per year (46 weeks)	paths
4	6.0 person-hours/wk	276	6
6	9.0	414	15
8	12.0	552	28
12	18.0	828	66
20	30.0	1,380	190
30	45.0	2,070	435

A thirty-person weekly standing meeting consumes 2,070 person-hours a year — more than a full-time post — to service 435 pairwise relationships of which any given attendee is party to 29. **The all-hands is a device for paying complete-graph cost while delivering star-topology information flow.** It is the worst of both wiring plans and it is the default.

8.5.10 5.10 The share of budget consumed by C

The approach to the crossover is where decisions actually get made. With `eta = 0.22`, `kappa = 0.006`, `N = n`:

n	eta ln n	C	C as multiple of budget
3	0.2417	0.0180	0.07
6	0.3942	0.0900	0.23
8	0.4575	0.1680	0.37
10	0.5066	0.2700	0.53
12	0.5467	0.3960	0.72
15	0.5958	0.6300	1.06
20	0.6591	1.1400	1.73
30	0.7483	2.6100	3.49
50	0.8606	7.3500	8.54
100	1.0131	29.7000	29.31
150	1.1023	67.0500	60.83

At `n = 150` coordination costs sixty-one times the entire structural budget. **A figure of 150 is widely quoted as a cognitive limit on stable relationships; no entry in this paper's bibliography supports it, this paper neither endorses nor relies on it, and nothing below depends on it.** What the arithmetic does say is that whatever any cognitive limit may be, **the economic limit under complete-graph obligation arrives around `n = 15` for these parameters, an order of magnitude sooner.** The reason human groups of 150 exist at all is that they never run a complete graph: they run cells, kin structures and role hierarchies — topologies, in the sense of Section 5.7.

8.5.11 5.11 When cost grows with stored order

Several `C` lines are not constant but scale with accumulated stored order. Compliance reporting grows with the number of rules to comply with; maintenance of built artefacts grows with the artefact count; archive management grows with the archive. Model it as `C_t = C_0 + gamma Q_t` and substitute:

```
Q_{t+1} = (1 − delta)∗Q_t + eta∗S_t − C_0 − gamma∗Q_t
        = (1 − (delta + gamma))∗Q_t + eta∗S_t − C_0
```

A cost that scales with stored order is arithmetically indistinguishable from a higher decay rate. The system behaves exactly as if `delta_eff = delta + gamma`, so every result of the retention analysis transfers with that substitution [Curzi, 2026a]: the time constant shortens to `1/(delta+gamma)`, the steady state falls to `(eta S − C_0)/(delta+gamma)`, and the uniform bound tightens.

Two practical consequences. **First, a `Q`-scaling cost line makes an organisation look forgetful when it is overburdened.** The symptom is identical — stored order decaying faster than expected — and the remedy is not: raising retention will not touch `gamma`. **Second, `gamma` is exactly the term that makes shedding `Q` a double win**, because reducing `Q` reduces both the `delta Q` maintenance draw and the `gamma Q` cost line at once.

8.5.12 5.12 The sustainability ceiling, stated hard

```
If   C_t  >  eta ∗ ln N  −  delta ∗ Q_t
then NO arrangement of people prevents decline.
```

Not *decline is likely*. Not *decline unless we work harder*. The inequality forecloses the entire space of geometric interventions, because `ln N` is the hard ceiling on `F(u)` and `eta ln N` is therefore the largest input the interaction channel can deliver in one period. **Every off-site, every reorganisation, every new collaboration model, and every leadership change operates strictly inside the space of interventions that this corollary has already ruled out.**

There are exactly four responses and they exhaust the space: reduce `C_t`; raise `eta`; raise `N` logarithmically, therefore weakly, and only with genuinely distinct categories; and reduce `Q_t`. The completeness of this set is established formally elsewhere in the corpus, where the steady state is shown to depend on exactly `u`, `N`, `eta` and `c` through `F`, plus `delta` as a divisor, with no fifth argument [Curzi, 2026].

Note which response is *not* on the list: raising `u`. Raising `u` increases `S_t` toward its ceiling `ln N`, which helps — but the ceiling condition is stated *at* the ceiling. **If cost exceeds the budget even when `S` is maximal, improving the geometry cannot close the gap.** The geometry work of Curzi [2026d] and Curzi [2026g] is essential, and it is work done when `h > 0`, not instead of the arithmetic.

8.5.13 5.13 Cutting cost badly

Every `C` line is attached to something. Cut C3 and `delta` rises on trust. Cut C6 and interaction rate falls, lowering `S`. Cut C1 indiscriminately and the channel through which reach operates is cut, lowering `eta`. Cut slack and the queueing term explodes. **A `C` cut that raises `delta` or lowers `eta` can easily be net negative, and the arithmetic is available to check before the cut rather than after.**

The rule: **for every proposed reduction, state which term it might damage and by how much, then compute the net effect on `h`.** A cut of 0.06 in `C` that costs 0.02 in `eta` at `ln N = 2.08` is a net loss of `0.02 x 2.08 − 0.06 = −0.0016`. That calculation takes a minute and is almost never performed.

The lines usually safe to cut are in C2 — friction — because friction by construction produces nothing. **Attack friction first, coordination second, everything else only with a computed net effect.**

8.6 6 Headroom: the master gauge

```
h_t := eta * ln N  −  delta * Q_t  −  C_t

    h_t < 0    unsustainable in this model. No geometry saves it.
    h_t = 0    maintenance only. Hold, do not build.
    h_t > 0    growth available; any requested increment must satisfy r_t <= h_t.
```

The maximum one-step increment is exactly `r_max = h_t`, and no policy based on interaction geometry alone can guarantee more. The target-increment condition is that `Q_{t+1} − Q_t >= r_t` holds if and only if `S_t >= (delta Q_t + C_t + r_t)/eta`, feasible only if `delta Q_t + C_t + r_t <= eta ln N`.

Headroom is not a new quantity. It is the sustainability ceiling rearranged into a gauge: `C_t <= eta ln N − delta Q_t` becomes `h_t >= 0`. **Its value is that it puts three parameters usually argued about separately onto one dial with one reading and one sign.** Read as an accounting statement:

```
   eta * ln N        the structural BUDGET     -- the most this system can convert in a period
 − delta * Q_t       the MAINTENANCE DRAW      -- what must be spent to keep what exists
 − C_t               the OPERATING DRAW        -- what is spent to be an organisation at all
 ------------------
 = h_t               HEADROOM                  -- what is left to build with
```

Proven: the definition and the three-regime reading follow from the sustainability ceiling and the maximum-increment corollary together with the range result `0 <= F(u) <= ln N`, which is what makes `eta ln N` an actual ceiling rather than a typical value [Curzi, 2026a]. **Postulated, badged:** that `eta`, `delta`, `N` and `C` are approximately constant over the horizon on which the gauge is read; that the order proxy stands in a stable relation to `Q`; and that the linear model is the right model at all. None of these is a theorem, and a headroom figure is only as good as the four measurements behind it. The gauge is reused as a precondition on protocol design elsewhere in the corpus, where `h < 0` forbids building and `h > 0` caps every increment at `r <= h`, worked at `N = 18`, `eta = 0.30`, `delta = 0.06`, `C = 0.55`, `Q = 1.20` giving `h = +0.2451` [Curzi, 2026].

8.6.1 6.1 Why the ceiling term is `ln N` and not `F(u)`

Headroom uses `eta ln N`, not `eta F(U_t)`, and the difference is deliberate. `F(U_t)` is what the system's *current* geometry delivers this period; `ln N` is the most any geometry could deliver, attained only at perfect orthogonality.

Headroom is therefore a statement about the best case, and a negative headroom is a statement that even the best case fails. When `h_t < 0` the reading is not that the current arrangement is inadequate; it is that the entire space of arrangements is inadequate, and no facilitator, retreat, restructure or hire inside that space can help. This is the strongest negative result the framework contains and it is strong precisely because it is stated at the ceiling.

The corollary is that **`h_t > 0` is necessary and not sufficient.** Positive headroom says growth is *permitted*; whether it *occurs* depends on `F(U_t)`, which is the subject of the geometry papers. The gauge licenses; the geometry delivers.

8.6.2 6.2 The three regimes

Regime 1, `h_t < 0`: unsustainable. `Q` will decline this period whatever is done. Exactly four things are available and they exhaust the space:

```
1. raise eta      effect = ln N per unit        usually cheapest
2. raise N        effect = eta/N per unit       logarithmic, therefore weak
3. cut C          effect = 1 per unit           direct
4. shed Q         effect = delta per unit       the one nobody considers
```

What not to do. Do not commission a workshop, a reorganisation, an engagement programme or a leadership change *as the response to negative headroom.* Each operates on `F(U_t)`, and `F(U_t) <= ln N` already, so the ceiling condition has ruled them out. They may be excellent things to do once `h > 0`. They cannot rescue `h < 0`. The correct report is: *headroom is negative at the stated figure; under this model no arrangement of these people prevents decline this period; the available moves are to reduce cost, raise conversion, or shed stored order; here is the arithmetic.* That sentence is worth more than any amount of strategy because it is checkable.

Regime 2, `h_t = 0`: maintenance only. The budget exactly covers the two draws. Hold, refuse new commitments, and protect the maintenance draw absolutely, because at `h = 0` any increase in `C` — one new reporting requirement, one new standing meeting — puts the system into Regime 1 immediately. **The characteristic error at `h = 0` is to read stability as slack.** A system holding steady at zero headroom looks calm from outside and has no reserve at all.

Regime 3, `h_t > 0`: growth available. Choose `r_t <= h_t`, then verify the geometry can deliver the required `S_t`. Worked at `eta = 0.25`, `N = 8`, `delta = 0.08`, `Q_t = 3.0`, `C_t = 0.10`:

```
h_t = 0.25*2.0794 - 0.08*3.0 - 0.10 = 0.5199 - 0.2400 - 0.1000 = +0.1799

request r_t  = 0.15                          (inside headroom)
required S_t = (0.2400 + 0.1000 + 0.1500)/0.25 = 1.9600 nats
```

```
ceiling      = ln 8 = 2.0794
required u   = (e^1.96 - 1)/7 = (7.0993 - 1)/7 = 0.8713
```

Feasible, and only just: it requires `u = 0.871`, which is a very high degree of non-parallelism and is not achievable by asking people to disagree more. The headroom said yes; the feasibility clause said *yes, at a geometry you almost certainly do not have*. Requesting `r_t = 0.05` instead requires `S_t = 1.56` and `u = 0.537`, which is attainable. **Always run the feasibility clause after the headroom check.** The gauge alone will license promises the geometry cannot pay for.

A concavity correction to the feasibility clause. The check above computes a single required `u` as though `U` were the same every period. It never is, and the error is one-sided. What accumulates is the sum of `eta F(U_t)`, so the quantity to test is the **mean of `F(U_t)`, not `F` of the mean `U`** — and because `F` is strictly concave, Jensen's inequality runs against the analyst:

```
mean of F(U_t)  <=  F( mean of U_t )        equality only if U is constant
```

At `N = 8` with sessions alternating between `u = 0.75` and `u = 0.30`, mean `U` is 0.525 and `F(mean U) = 1.5422`, but mean `F` is 1.4820 — a concavity penalty of 0.0602 nats, **3.9 per cent of the surplus lost to variability alone. Report the required `u` for legibility if desired; never test against it.** State the requirement as `mean F(U_t) >= S_req` and compare it to the measured mean of `F`. Testing on `F` of the mean certifies increments the geometry cannot deliver, and the penalty it hides is recoverable for free by regularising a variable geometry.

8.6.3 6.3 Why headroom falls as stored order grows

One line of algebra:

```
h_t = eta*ln N - delta*Q_t - C_t          =>          d h_t / d Q_t  =  -delta
```

Headroom decreases linearly in stored order at rate `delta`. Every unit of durable structure a system acquires permanently reduces, by `delta` per period, the amount it can build next.

This is the structural reason institutions ossify and it deserves to be stated without moral vocabulary. **Ossification is not complacency, not cowardice, not bureaucratic capture, and not a failure of leadership. It is `dh/dQ = -delta`.** An institution that succeeds accumulates `Q`; accumulated `Q` demands maintenance at `delta Q`; maintenance consumes the budget; the budget was the capacity to build. **Success consumes its own headroom, by construction, in the absence of any error by anyone.** March [1991] reaches the same destination from a different direction — an organisation whose members learn too fast homogenises and destroys the variance that produced its advantage — and Levinthal and March [1993] names the myopia that follows.

At `eta = 0.25`, `N = 8` so budget 0.5199, `delta = 0.08`, `C = 0.10`:

`Q`	maintenance draw `delta Q`	`h = 0.5199 - delta Q - 0.10`
0	0.0000	+0.4199
1	0.0800	+0.3399
2	0.1600	+0.2599
3	0.2400	+0.1799
4	0.3200	+0.0999
5	0.4000	+0.0199
6	0.4800	-0.0601
8	0.6400	-0.2201
10	0.8000	-0.3801

Q	maintenance draw	delta Q h = 0.5199 − delta Q − 0.10
15	1.2000	-0.7801
20	1.6000	-1.1801

At Q = 0 this system can build freely. At Q = 5 it has 0.0199 and is a maintenance organisation. At Q = 6 it is in decline. **Nothing changed except that it succeeded five times.** The advantage of a young institution is now arithmetic rather than folklore: a new entrant outbuilds an incumbent with more money and more people because the entrant's delta Q is near zero and the incumbent's consumes its entire budget. **The incumbent is not slower; it is spending its budget on being what it already is.**

8.6.4 6.4 The critical stock

Setting h_t = 0 and solving for Q:

```
Q_crit  =  ( eta*ln N  −  C )  /  delta
```

Q_crit is the largest stock of order a system can hold without going into decline. It is a hard capacity computable from four measured numbers and it belongs on the front page of any survey. With eta = 0.25 and N = 8, so eta ln N = 0.5199:

delta C	0.00	0.05	0.10	0.20	0.40
0.01	51.99	46.99	41.99	31.99	11.99
0.02	25.99	23.49	20.99	15.99	5.99
0.05	10.40	9.40	8.40	6.40	2.40
0.10	5.20	4.70	4.20	3.20	1.20
0.20	2.60	2.35	2.10	1.60	0.60
0.30	1.73	1.57	1.40	1.07	0.40

First, Q_crit scales as 1/delta, so retention is the dominant term in institutional capacity. Moving delta from 0.10 to 0.02 takes capacity from 4.20 to 20.99 at C = 0.10, a factor of five from retention alone. **Institutions that hold the most stored order are not those with the most input; they are those with the lowest delta.**

Second, C bites hardest when delta is small. At delta = 0.30, raising C from 0 to 0.40 costs 1.33 units of capacity; at delta = 0.01 the same increase costs 40 units. **The better the retention, the more damaging the overhead**, because a low delta divides eta ln N − C by a small number and amplifies every change in C.

8.6.5 6.5 The maintenance trap

A system near Q_crit exhibits full activity and high effort; output that is entirely replacement rather than addition; rising frustration among people who joined to build and find themselves maintaining; every new initiative displacing an existing one; and leadership attributing the stall to insufficient ambition, focus or talent.

The last item is the trap. **Every symptom of the maintenance trap is also a symptom of a motivation problem, and the two require opposite responses.** Exhortation applied to a headroom problem raises C — more meetings, more reporting, more alignment work — and therefore lowers h, worsening exactly the condition it was meant to fix. The distinguishing test is arithmetic and takes an afternoon: compute delta Q_t and compare it to eta ln N. **If the maintenance draw exceeds half the budget, the diagnosis is structural and no amount of motivation is relevant.**

8.6.6 6.6 The fixed-point identity

At the steady state with constant `U = u*` and `C = c*`:

```
Q_inf = ( eta*F(u*) − c* ) / delta

h_inf = eta*ln N − delta*Q_inf − c*
      = eta*ln N − ( eta*F(u*) − c* ) − c*
      = eta*ln N − eta*F(u*)

h_inf = eta * ( ln N − F(u*) )
```

At the fixed point, headroom equals the conversion efficiency times the geometric shortfall. `delta` cancels, `C` cancels, `Q` cancels. The only surviving quantities are `eta` and the gap between what the geometry delivers and what the category count would permit. Verified numerically across four unrelated parameter sets:

eta	N	u*	delta	C	S = F(u*)	Q_inf	h_inf computed	eta(ln N − S)
0.25	8	0.60	0.080	0.10	1.6487	3.9021	+0.107696	+0.107696
0.18	12	0.45	0.060	0.22	1.7834	1.6835	+0.126273	+0.126273
0.20	14	0.50	0.045	0.08	2.0149	7.1773	+0.124831	+0.124831
0.40	20	0.90	0.030	0.30	2.8959	28.6122	+0.039928	+0.039928

Identical to six decimal places, as the algebra requires.

What this means is the most uncomfortable result in the paper. The only headroom a mature system possesses is exactly the surplus it is failing to extract from its own difference. Watch it vanish as the geometry improves, at `eta = 0.25`, `N = 8`:

u*	S = F(u*)	h_inf = eta(ln N − S)
0.20	0.8755	0.300993
0.40	1.3350	0.186110
0.60	1.6487	0.107696
0.80	1.8871	0.048093
0.90	1.9879	0.022892
0.95	2.0347	0.011184
0.99	2.0707	0.002197
1.00	2.0794	0.000000

A system that has arrived at its fixed point and perfectly exploits its own diversity has zero headroom. At `u = 1` the surplus is at ceiling, `Q_inf` is at its maximum `eta ln N / delta` — which is the uniform bound arrived at from a different direction — and there is nothing left over. This is not a paradox and not a criticism of difference; it is the statement that a fully converged system has converged, having spent its capacity on becoming what it is. Banach [1922] guarantees the arrival and nothing guarantees the destination is liked.

The same zero arrives from a completely different direction elsewhere in the corpus. At the calibration fixed point where doctrine exactly matches evidence, the input sits at its ceiling and headroom is identically zero by cancellation, `h = eta ln N − (eta ln N − C) − C = 0`, **so perfect calibration is not an unqualified target** [Curzi, 2026r]. Undistorted contact maximises what is banked and leaves nothing to build with.

The design consequence: **to retain headroom at maturity a system must either hold `eta` high, refuse to converge, or periodically leave the fixed point by reducing `Q`.**

8.6.7 6.7 The linear model does not ossify by itself

An honest complication the previous section forces. If `h_inf = eta(ln N − F(u*)) >= 0` always — and it does, since `F(u) <= ln N` — then a system running the plain storage law with constant parameters **never reaches negative headroom.** It approaches its fixed point and sits there with a small positive `h`. At `eta = 0.25`, `N = 8`, `u = 0.60`, `delta = 0.08`, `C = 0.10`, `Q_0 = 0`:

t	Q_t	h_t
0	0.0000	+0.4199
2	0.5994	+0.3719
5	1.3303	+0.3134
10	2.2070	+0.2433
15	2.7849	+0.1971
20	3.1658	+0.1666
30	3.5822	+0.1333
40	3.7631	+0.1188

Headroom falls monotonically from 0.4199 toward its limit 0.107696 and stops. It never crosses zero.

Therefore observed institutional ossification — real organisations that genuinely cannot build anything — requires something outside the constant-parameter linear model. It is not automatic, and this paper does not pretend it is. Three candidate mechanisms, each with a different remedy:

Mechanism A — cost that scales with stored order. `C_t = C_0 + gamma Q_t`, arithmetically identical to a higher decay rate. **It does not by itself produce negative headroom**, because the fixed-point identity survives the substitution:

```
Q_inf = (eta*S − C_0)/(delta + gamma)
h_inf = eta*ln N − delta*Q_inf − C_0 − gamma*Q_inf
      = eta*ln N − (delta + gamma)*Q_inf − C_0
      = eta*ln N − (eta*S − C_0) − C_0  =  eta*(ln N − S)
```

Identical. A linear `Q`-scaling cost does not cause ossification; it only makes the system converge to a smaller `Q` faster. This eliminates the most popular explanation, and it is a result worth stating because it is easy to assume the opposite.

Mechanism B — exogenous accretion. `Q` arriving through a channel that is not `eta S`: inherited mandates, merged organisations, statutory obligations, standards adopted by external bodies that must now be upheld. Model it as a constant `A` added each period:

```
Q_{t+1} = (1−delta)*Q_t + eta*S − C_t + A
h_inf   = eta*(ln N − S)  −  A
```

This does go negative, and it goes negative when `A > eta(ln N − F(u*))`. Exogenous stored order accumulates without paying for its own maintenance, and headroom is what it consumes.

Mechanism C — declining conversion or non-parallelism with age. Cohesion rising over time drives `u` toward zero; informational cascades collapse the effective count in a group that started diverse [Bikhchandani et al., 1992, Banerjee, 1992, Lorenz et al., 2011]; capture ceremonies displace real capture. Any of these lowers `eta S` while `Q` stays high. The decay of `u` as a state variable in its own right, with a measured half-life, is developed in Curzi [2026d].

Of the three, Mechanism B is the one most often mistaken for laziness and the one most amenable to a direct fix, and the third worked system of Section 9 is built on it.

8.6.8 6.8 The four levers, priced

Marginal derivatives of `h`:

```
d h / d eta  =  ln N          raising conversion is worth ln N per unit
d h / d N    =  eta / N       raising categories is worth eta/N per unit -- and falls with N
d h / d C    =  −1            cutting cost is worth exactly 1 per unit
d h / d Q    =  −delta        shedding order is worth delta per unit
```

Worked at `eta = 0.25`, `N = 8`, `delta = 0.08`, `Q = 5`, `C = 0.10`. Baseline `h = 0.5199 − 0.4000 − 0.1000 = +0.0199`, a classic maintenance-trap state.

Lever	Move	New `h`	Gain	What it costs in practice
L1 raise `eta`	0.25 to 0.35	+0.2278	+0.2079	a decision log and a recorder rota
L2 raise `N`	8 to 12	+0.1212	+0.1014	four genuinely new disciplines recruited
L2b raise `N`	8 to 16	+0.1931	+0.1733	doubling the discipline count
L3 cut `C`	0.10 to 0.04	+0.0799	+0.0600	a 60 per cent cut in overhead
L4 shed `Q`	5.0 to 3.0	+0.1799	+0.1600	retire two units of stored order
L4b shed `Q`	5.0 to 2.0	+0.2599	+0.2400	retire three units

Shedding three units of stored order outperforms every other move on the table, including a 40 per cent improvement in conversion efficiency. It is also the only move on the list requiring no new capability, no recruitment, no budget, and no negotiation outside the organisation.

The exchange rate between shedding and the other levers is exact. To buy headroom increment `dh` by shedding, shed `dQ = dh/delta`:

`delta` \ `dh`	0.02	0.05	0.10	0.20	0.50
0.02	1.00	2.50	5.00	10.00	25.00
0.05	0.40	1.00	2.00	4.00	10.00
0.08	0.25	0.62	1.25	2.50	6.25
0.10	0.20	0.50	1.00	2.00	5.00
0.20	0.10	0.25	0.50	1.00	2.50
0.30	0.07	0.17	0.33	0.67	1.67

Shedding is cheapest when `delta` is high. A high-maintenance system gets a large headroom return for a small reduction in stock; a system that has invested heavily in permanence must shed a great deal to recover a little. **Durability and adaptability trade against each other**

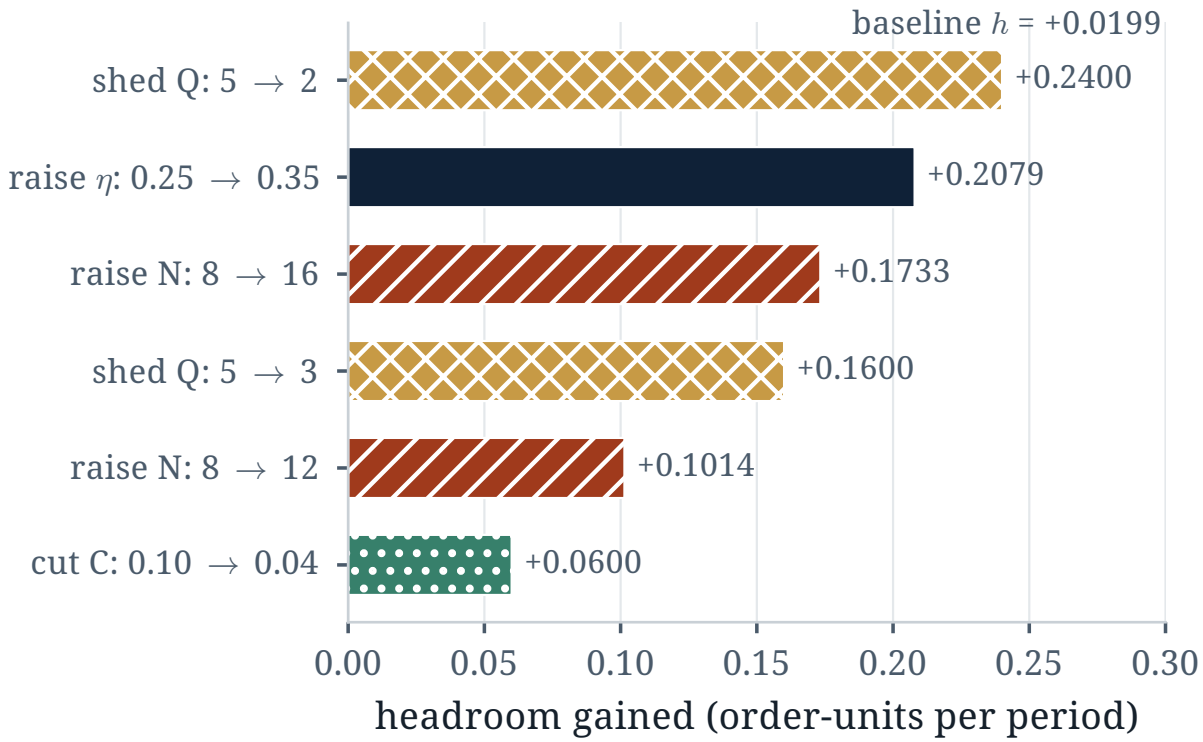


Figure 8.5: **Shedding stored order outperforms every other lever at this base case**, including a 40 per cent rise in conversion; the baseline headroom is +0.0199. Modelled at the stipulated `eta = 0.25`, `N = 8`, `delta = 0.08`, `Q = 5`, `C = 0.10`; colour encodes which of the four parameters is moved. Generated by `figs/src/CCP-008_fig05.py`.

exactly, and the exchange rate is `1/delta`. That is the honest form of the intuition that solid institutions are hard to change: it is not a moral fact about institutions, it is `dh/dQ = -delta`.

And against `N`, at `eta = 0.25`, `delta = 0.08`:

`N` change	headroom gained	equivalent `Q` to shed
8 to 10	+0.0558	0.697
8 to 12	+0.1014	1.267
8 to 16	+0.1733	2.166
8 to 32	+0.3466	4.332
8 to 64	+0.5199	6.498

Retiring 2.166 units of obsolete stored order is worth exactly as much as doubling the number of genuinely distinct disciplines in the organisation. One of those is a decision; the other is a decade. The corresponding computation on the other side of the ledger is worked in Curzi [2026j], where the `N` lever is computed and declined on net: a ceiling gain of `eta(ln 6 - ln 5) = +0.1914` against five added communication paths costing `+0.0714`, a 34 per cent haircut on the headline benefit.

8.6.9 6.9 The fourth lever: deliberately shedding stored order

6.9.1 Why it is invisible

One: `Q` is not measured, so it cannot be reduced. Most organisations have no order observable at all, and a quantity that is not tracked cannot be decided about.

Two: `Q` is confused with achievement. Stored order is the accumulated record of past success, and reducing it reads as undoing that success. The law is indifferent to how `Q` was acquired. **A procedure that was excellent fifteen years ago and is now upheld out of habit costs exactly the same `delta Q` as one that is excellent today.** The maintenance draw does not inspect provenance.

Three: shedding has an identifiable victim and an unidentifiable beneficiary. The person whose procedure is retired is specific and present; the headroom released is diffuse and future. That is a political asymmetry, not an arithmetic one, and it is why shedding must be scheduled in advance rather than negotiated in the moment.

Four: working infrastructure is invisible. Star and Ruhleder [1996] and Star and Griesemer [1989]: infrastructure becomes visible only on breakdown, so working stored order never appears on a list of things that could be removed, while its maintenance cost is paid silently every period. The countermeasure is an inventory listing what exists rather than what is broken.

6.9.2 What can be shed and what cannot

Freely sheddable. Procedures whose triggering condition no longer occurs. Reports nobody reads, measurable directly by access logs. Approval gates with a rejection rate near zero, since a gate that never rejects is pure `C` and pure `delta Q` and performs no filtering function. Superseded documentation retained for reference. Committees whose founding question was answered. Tooling with no active users. Metrics captured by the measure-becomes-target failure and no longer measuring what they name [Goodhart, 1975, Campbell, 1979].

Conditionally sheddable, with a stated replacement. Redundant capability: redundancy is `Q` that buys resilience, which is real, and it should be shed only after computing which failure mode is being re-exposed — Albert et al. [2000] is the relevant caution, since robustness to random failure and fragility to targeted attack are different properties and shedding redundancy affects them differently, while Sagan [1993] documents settings in which redundancy reduces safety and Perrow [1984] explains why. Institutional memory of past failures: genuinely valuable and genuinely costly, and the specified form is to **shed the procedure and keep the one-line reason**, because a rule with its rationale attached can be re-derived cheaply and a rule without one must be maintained forever because nobody dares remove it. Standards commitments to external parties, which require negotiation rather than decision.

Not sheddable under any headroom pressure. Anything constituting a person's ability to decline participation. A consent provision, a right of appeal, a dissent protection or a minority-report requirement is stored order that shows up in the `delta Q` draw, and **it is not available for shedding regardless of what the arithmetic says.** Safety interlocks whose failure mode is injury. Obligations to parties not represented in the decision to shed.

This is the one place in this paper where a lever with a computed positive return is declared out of bounds. It is declared out of bounds anyway, and the model gives no warning at that boundary — the computed return is genuinely positive, which is exactly why the constraint must sit outside the model.

6.9.3 How to shed: schedule it in advance

Shedding negotiated in the moment fails, because the identifiable-victim asymmetry makes every individual case losable. **Shedding works when it is scheduled at the moment of creation, before anyone is attached.** The devices: sunset clauses, so that renewal requires an argument and continuation does not require a decision to end it; reversal conditions, so that a decision whose reversal condition has been met and which still stands is a shedding candidate identified by the system itself; term limits on roles and committees; review dates whose **default is removal**, since a review whose default is retention is a formality; planned demolition, on the argument that a city which never demolishes anything has no site available for anything new [Jacobs, 1961]; and scheduled forgetting, which is the deliberate form of the decay that the

retention analysis spent its length trying to reduce [Curzi, 2026a]. The two are not in conflict: **control the leak so that what drains can be chosen.**

6.9.4 The double return when cost scales with stock

Where `C_t = C_0 + gamma Q_t`, shedding attacks two terms at once, `dh/dQ = −(delta + gamma)`, so the shed required to buy headroom `dh` is `dQ = dh/(delta + gamma)`, which is smaller. **In any institution whose compliance, reporting or maintenance burden scales with what it has accumulated — which is most large institutions — shedding is strictly more efficient than the table of Section 6.8 suggests.**

6.9.5 What shedding does not do

Shedding raises `h` and does not raise `Q_inf`. It moves the system away from its fixed point, and the contraction guarantees a return: the system re-accumulates toward `(eta S − C_0)/(delta + gamma)` at time constant `1/(delta+gamma)`. **Shedding buys a window, not a new equilibrium.** The window is genuinely valuable — it is the period during which `h` is large enough to build something structurally different — but a shedding programme not paired with a use for the released headroom simply regrows the same stock. The correct sequencing:

```
1. shed Q          -> h rises, window opens
2. spend h on eta  -> the permanent lever (dh/d eta = ln N)
3. re-accumulate at the higher eta -> a higher Q_inf and a higher h_inf
```

Shedding is a transient. Raising conversion is permanent. Use the transient to buy the permanent.

8.6.10 6.10 Measuring headroom in practice

The minimum viable reading, achievable in about two weeks in most organisations: enumerate `N` by listing genuinely distinct methodological or domain categories rather than job titles, expecting the audited count to be roughly half the number first proposed [Curzi, 2026c]; count `Q` crudely, since the *sign* of `h` is usually robust to a factor of two in `Q` and the sign is what governs the decision; estimate `delta` from the last period of genuine inactivity; tally `C` by the nine lines of Section 5.1; estimate `eta` from the last six months of attributable durable change divided by computed `S`; and compute `h`, `Q_crit` and `h_inf = eta(ln N − S)`.

Report the sign before the magnitude. The sign determines which of the four levers may be reached for; the magnitude only determines how much may be requested. **And report the sensitivity** — how much `Q` would have to be wrong for the sign to flip. If `h = +0.02` and a five per cent error in `Q` flips it, say so.

A cheaper route to the direction, and it is an identity rather than an approximation. Holding `eta`, `N`, `delta` and `Q` fixed across the configurations being compared, `h(x) = eta ln N − delta Q − C(x)` depends on the configuration only through `C`, so

```
grad h = −grad C
```

Measuring where a system's own choices already go and measuring the gradient of headroom are the same measurement, in the same direction, with the same zeros. A cheap habit survey therefore locates the direction of steepest improvement without the full instrument. Two cautions: this gives direction only, and the sign of `h` still requires the four measurements; and the identity holds only where `eta`, `N`, `delta` and `Q` really are configuration-independent, which fails as soon as a configuration changes who is in the room, because that changes `N`. **A measured preference pointing consistently away from measured headroom improvement is**

a diagnosis rather than a contradiction: one of those four is moving with the configuration. The commonest case is low `C` with low `N` — the comfortable narrow life — and it is visible only if `N` is measured alongside `C`.

8.7 7 Design specifications

Numbered, checkable requirements. A system claiming to have engineered these three terms satisfies them. They are **design requirements, not findings**: each is a rule this framework proposes, and none is an empirical result.

8.7.1 7.1 Requirements on conversion

D1.1 The system shall define an order observable `Y_t` in writing, together with its measurement procedure and its period length.

D1.2 The system shall document the argument that `Y_t` is not a repackaging of the interaction observable `U_t`. The argument shall be specific, not asserted.

D1.3 Every recurring interaction of more than two participants shall have a named recorder before it begins. The role may rotate; it may not be vacant.

D1.4 The system shall maintain a decision log. Each entry shall contain at minimum the decision, the date, the alternatives rejected, the accountable person, the review date, and the reversal condition.

D1.5 No decision-log entry shall have a vacant accountable-person field. An entry without an owner has `eta_bind = 0` and shall be marked a proposal, not a decision.

D1.6 Every entry shall carry a reversal condition — the observation that would require the decision to be revisited. Entries without one shall not be closed.

D1.7 The system shall maintain a register of changed defaults, distinct from the decision log, listing for each the default, the artefact that implements it, and the date it became live.

D1.8 A decision not implemented in tooling, template, form, calendar or physical arrangement within one review period shall be reported as unrealised. Its contribution to `Q` is zero until then.

D1.9 The system shall estimate `eta` at least once per twelve periods using `eta_hat = delta_Q_attributable / S_input`, with `S` computed from measured `N` and `u` and not from attendance.

D1.10 Every estimate shall be reported alongside all four baselines: persistence, direct-input, matched autoregressive, and shuffled-input.

D1.11 The system shall decompose `eta` into recording, binding and reach at each estimation, and shall direct the next intervention at the smallest of the three.

D1.12 The system shall identify in writing the fraction of its stored order it considers tacit, and shall provide a demonstration channel for that fraction.

D1.13 No capture requirement shall be imposed whose estimated `C` exceeds its estimated `d(eta) S`. Capture is subject to the same budget as everything else.

D1.14 Capture shall target decisions and changed defaults. Verbatim capture of discussion shall not be a requirement of any recurring interaction.

D1.15 The system shall count re-litigated decisions per period as a standing diagnostic. A rising count is the signature of `eta S < delta Q`.

D1.16 Onboarding shall be defined as a document plus a demonstration channel, never as a sequence of conversations.

D1.17 Before any conversion programme is funded, the system shall verify `S > 0` by measuring `u`. If `u` is near zero the programme shall be deferred until the geometry work is done.

D1.18 The system shall not report attendance, meeting count, workshop hours or engagement as evidence of `Q`. These are `S` proxies and shall be labelled as such wherever they appear.

D1.19 Any conversion metric adopted as a target shall be re-examined for measure-becomes-target capture at each estimation.

D1.20 Every capture mechanism shall be reversible by the people subject to it. No capture requirement shall remove a participant's ability to decline participation.

8.7.2 7.2 Requirements on cost

D2.1 The system shall maintain a written cost tally decomposed into the nine lines of Section 5.1. An aggregate `C` figure shall not be accepted as a survey deliverable.

D2.2 The system shall estimate `kappa` by the four-step procedure of Section 5.9 at least once per twelve periods, and shall record the unit normalisation used and an interval on the estimate.

D2.3 The system shall compute and record its crossover size and its net-optimal size `n_opt = sqrt(eta/kappa)` before any decision that changes headcount.

D2.4 No unit shall be grown past `n_opt` without a recorded justification stating what is being bought with the net structural loss.

D2.5 No unit shall be grown past the crossover without first changing the topology, because past the crossover the sustainability ceiling is violated even at `Q = 0`.

D2.6 The system shall state its communication topology explicitly: which pairs are obliged to maintain a live channel and which are not. **An unstated topology defaults to the complete graph and shall be costed as one.**

D2.7 Complete-graph obligation shall be confined to cells of size at most `n_opt`. Inter-cell links shall be explicit, named, and budgeted at an elevated `kappa`.

D2.8 Cell membership shall be permitted to overlap. A partition into disjoint cells with a reporting hierarchy is a tree and shall be flagged as a `u`-destroying topology.

D2.9 No topology shall route all inter-cell traffic through a single node. Hub topologies shall be recorded as accepting a known `u` penalty and a known targeted-failure fragility.

D2.10 Utilisation of every identified critical shared resource shall be measured. Any resource above 0.90 shall be raised as a fault.

D2.11 Target utilisation for critical shared resources shall be at or below 0.80. Slack below that level shall not be reported as waste in any budget document.

D2.12 Every standing recurring meeting shall have its annual person-hour cost computed and displayed on its own agenda.

D2.13 Any standing meeting with more than `n_opt` attendees shall state in writing which information flow requires all of them, or be split.

D2.14 The system shall compute the sustainability ceiling `C_max = eta ln N − delta Q_t` each period and compare it to measured `C_t`. This comparison shall appear in any report proposing new activity.

D2.15 Where `C_t > eta ln N − delta Q_t`, the system shall not authorise any intervention whose mechanism is interaction geometry, and shall record that the sustainability ceiling forecloses it.

D2.16 The system shall test whether any `C` line scales with `Q`. Where it does, `gamma` shall be estimated and `delta_eff = delta + gamma` used in all subsequent computation.

D2.17 Every proposed cost reduction shall be accompanied by a statement of which other parameter it may damage and a computed net effect on `h`.

D2.18 Friction shall be attacked before coordination, and coordination before any other line.

D2.19 Conflict and disagreement shall be tallied separately and never conflated. A programme that reduces conflict shall report its measured effect on `u`.

D2.20 No cost reduction shall be implemented that removes a participant's ability to decline an interaction. Mandatory attendance is a cost increase disguised as a cost control.

8.7.3 7.3 Requirements on the gauge

D3.1 The system shall compute `h_t = eta ln N − delta Q_t − C_t` at least once per twelve periods and shall record the four measurements and their dates.

D3.2 The sign of `h_t` shall be reported before its magnitude and shall accompany any proposal for new activity.

D3.3 No increment shall be requested such that `r_t > h_t`.

D3.4 Every requested increment shall be checked against the feasibility clause `delta Q_t + C_t + r_t <= eta ln N`, and the implied required `u` shall be computed and stated.

D3.5 Where the implied required `u` exceeds the measured `u` by more than a stated margin, the request shall be refused as geometrically infeasible even if `r_t <= h_t`.

D3.6 Where `h_t < 0`, no intervention whose mechanism is interaction geometry shall be authorised. The refusal shall show the arithmetic.

D3.7 Where `h_t < 0`, the response shall be selected from exactly four options: raise `eta`, raise `N`, cut `C`, shed `Q`. Any other proposal shall be recorded as out of scope.

D3.8 The system shall compute and publish `Q_crit = (eta ln N − C)/delta` and shall report `Q_t/Q_crit` as a standing capacity-utilisation figure.

D3.9 Where `Q_t/Q_crit` exceeds 0.75, the system shall open a shedding review.

D3.10 Where `delta Q_t` exceeds half of `eta ln N`, the system shall record that it is in the maintenance regime and shall not attribute the resulting stall to motivation, ambition, focus or talent.

D3.11 Every unit of stored order shall carry, at creation, a sunset date or an explicit statement that it is permanent, with the reason.

D3.12 Renewal past a sunset date shall require a positive argument. Continuation shall never be the default at review.

D3.13 The system shall maintain a shedding inventory listing every candidate in the three classes of Section 6.9.2, updated each review.

D3.14 No item shall be shed that constitutes a person's ability to decline participation, regardless of computed headroom benefit.

D3.15 No item shall be shed that constitutes a safety interlock whose failure mode is injury.

D3.16 No item shall be shed that constitutes an obligation to a party not represented in the shedding decision, without that party's participation.

D3.17 Where a procedure is shed, its one-line rationale shall be retained.

D3.18 Every shedding programme shall state in advance what the released headroom will be spent on. A programme with no stated use shall be deferred.

D3.19 Where any cost line scales with `Q`, `dh/dQ = −(delta + gamma)` shall be used in all shedding calculations.

D3.20 The system shall compute `h_inf = eta(ln N − F(u*))` and report it as the headroom it will retain once converged. A projected `h_inf` near zero shall be flagged.

D3.21 Every headroom report shall state the sensitivity of the *sign* of `h` to error in `Q`, `delta`, `eta` and `C`.

D3.22 No headroom figure shall be published without the order-proxy justification of D1.2.

8.8 8 Failure modes

Each entry names the damaged parameter, the diagnostic signature, the detecting measurement, and the fix. **These are constructed clinical descriptions, not case reports; no organisation was observed for this section.** The corpus's full catalogues run to thirty-two entries for the storage law [Curzi, 2026j] and thirty for the protocol layer [Curzi, 2026].

8.8.1 8.1 Faults of conversion

F1 · The leaky vessel. *Damages* `eta`*.* Interaction of high quality and adequate geometry occurs on schedule and none of it is captured: `S` is large, `eta` is near zero, and `Q` decays at rate `delta` regardless. *Signature:* recurring re-litigation of settled questions; knowledge held entirely in individuals; onboarding by conversation. *Detection:* count decisions revisited per period; ask five members independently to state the current policy on a specific matter and compare answers, since disagreement is direct evidence of a recording failure. *Fix:* D1.3, D1.4. The cheapest fix in the paper.

F2 · The minutes fallacy. *Damages* `eta_bind`*.* Extensive verbatim records exist and nothing is bound, so recording is high, binding near zero, and the product near zero. *Signature:* a large, complete, chronological archive nobody consults; no owner fields anywhere. *Detection:* access logs; ask when a minute was last cited in an argument about what to do. *Fix:* replace minutes with a decision log.

F3 · The write-only archive. *Damages* `eta_reach`*.* Capture and binding occur but the record cannot be found by the people who need it. *Signature:* people re-derive documented things; multiple contradictory documents on one topic, each authoritative-looking. *Detection:* give a new member a specific question and time the retrieval. *Fix:* single index, mandatory supersession marking, one canonical location per topic. Deleting duplicates is a shedding action and is legitimate.

F4 · The unowned artefact. *Damages* `eta_bind`*.* A decision or artefact exists with no accountable person and degrades without anyone being responsible for noticing. *Signature:* orphaned tooling that everyone depends on and nobody maintains. *Detection:* audit the decision log and tooling register for vacant owner fields; the count is the measurement. *Fix:* D1.5.

F5 · The capture tax. *Damages* `C`*, and* `eta` *indirectly.* Capture requirements proliferate until the cost of documenting exceeds the value of what is documented. *Signature:* templates with more fields than content; rising `C` with flat `Q`. *Detection:* person-hours spent on capture as a fraction of person-hours spent in interaction. *Fix:* D1.13, D1.14.

F6 · The ceremony that replaced the capture. *Damages* `eta`*, while appearing to raise it.* A capture ritual is performed faithfully and produces no binding. *Signature:* action items with a low completion rate and no consequence for non-completion; identical items recurring across sessions. *Detection:* track completion and recurrence; recurrence is the tell. *Fix:* attach every item to the changed-default register or delete it.

F7 · The high-conversion error. *Damages nothing in the law; the law is working correctly against the system.* High `eta` and low `delta` have converted a wrong conclusion into durable structure that cannot be dislodged. *Signature:* confident, well-documented, deeply implemented practice that no evidence appears able to change; challenges met with the documentation. *Detection:* check the decision log for reversal conditions; their absence is the diagnostic. *Fix:* D1.6 applied retroactively. Where a reversal condition has been met and the decision stands, that is a shedding case.

F8 · The attribution illusion. *Damages the estimate, not the coefficient.* `eta` is estimated by correlating interaction with subsequent order growth with no control, so any secular trend produces a positive estimate. *Signature:* an estimate never tested against a shuffled-input control; a model that fits everything and predicts nothing out of sample. *Detection:* run the shuffled-input control. *Fix:* D1.10.

F9 · The order proxy that was really the input. *Damages every downstream number.* `Y_t` is a disguised measure of `U_t`. *Signature:* `Y` rises on periods with much discussion and no conclusions. *Detection:* identify a period of intense interaction that produced nothing durable and check whether `Y` rose; or apply the ceiling test, since no admissible increment exceeds `eta ln N` [Curzi, 2026i]. *Fix:* D1.2 — choose a proxy that can fall while interaction is high.

F10 · Conversion applied to a parallel roster. *Damages the intervention budget.* A serious capture programme is installed in a group with `u` near zero, so `S = 0`, `eta S = 0` for any `eta`, and the programme returns nothing while consuming `C`. *Signature:* excellent documentation

of unanimous conclusions; a decision log in which nothing was ever contested. *Detection:* measure `u`. *Fix:* D1.17 — defer, and fix the geometry first.

8.8.2 8.2 Faults of cost and topology

F11 · The quadratic committee. *Damages `C`.* A group grown past `n_opt`, often past the crossover, under complete-graph obligation. *Signature:* meetings that lengthen without producing more; scheduling as the dominant constraint; decisions deferred for want of a quorum. *Detection:* count `n`, compute `sqrt(eta/kappa)`, compare; compute `C/(eta ln N)`. *Fix:* D2.7.

F12 · The utilisation cliff. *Damages `C` via the queueing term.* A critical shared resource is driven toward full utilisation in pursuit of efficiency and wait time explodes hyperbolically. *Signature:* lead times lengthening far out of proportion to load; one named person as the bottleneck for many flows. *Detection:* measure `rho`; anything above 0.90 is the diagnosis. *Fix:* D2.10, D2.11 — add capacity, remove the gate, or shed demand, and defend slack in the budget as a cost-control device.

F13 · The all-hands trap. *Damages `C`, and `u` indirectly.* Complete-graph cost paid for star-topology information flow. *Signature:* large standing meetings with low participation dispersion; attendees doing other work during them. *Detection:* D2.12, plus the speaking-time distribution [Curzi, 2026g]. *Fix:* D2.13 — split, or convert to a written broadcast, which is star topology at a fraction of the cost.

F14 · The uncosted topology. *Damages every downstream estimate.* Nobody has stated who is obliged to coordinate with whom, so everyone assumes everyone and the tally understates `C` badly. *Signature:* nobody can draw the communication graph. *Detection:* attempt to draw it; the failure to draw it is the finding. *Fix:* D2.6.

F15 · The hub that ate the variance. *Damages `u`, while appearing to fix `C`.* Coordination cost is reduced by routing everything through a central node; `C` falls and so does `u`, because every message is filtered through one viewpoint. *Signature:* cheap coordination and unanimous conclusions; cross-team disagreement never surfaces. *Detection:* measure `u` before and after any centralisation. *Fix:* D2.9 — add lateral links that bypass the hub, and compute the trade.

F16 · The compliance ratchet. *Damages `C` via `gamma`; manifests as elevated apparent `delta`.* Cost that scales with stored order grows unnoticed as `Q` grows until `delta_eff = delta + gamma` dominates. *Signature:* stored order decaying faster than the measured `delta` predicts; older institutions performing worse than younger ones with identical staff. *Detection:* D2.16. *Fix:* reduce per-unit overhead, or shed `Q`, which attacks `delta Q` and `gamma Q` simultaneously.

F17 · The conflict-suppression trade. *Damages `u`.* Conflict is reduced by discouraging disagreement rather than by improving the handling of it; cost falls and surplus falls further. *Signature:* harmonious meetings, low measured conflict, low measured disagreement, decisions unanimously endorsed and privately doubted. *Detection:* D2.19 — tally the two separately; the signature is both falling together. *Fix:* separate them; acoustic and geometric treatment reduce conflict without reducing disagreement [Curzi, 2026e,f].

F18 · The net-negative cut. *Damages `h`, while appearing to raise it.* A cost reduction whose collateral damage to `eta` or `delta` exceeds the saving. *Signature:* cost falls, stored order falls faster. *Detection:* D2.17. *Fix:* reverse the cut where possible; this is one of the few cases in which restoring `C` raises `h`.

F19 · The missing denominator. *Damages the sizing decision.* Growth justified by total output rather than output per member, throughout the entire band between `n_opt` and the crossover. *Signature:* headcount growth defended by absolute figures. *Detection:* compute `net(n) = eta ln n − kappa n(n−1)/2` at current and proposed `n`. *Fix:* D2.3, D2.4.

8.8.3 8.3 Faults of the gauge

F20 · The unread gauge. *Damages every decision downstream.* `h` is never computed; capacity is argued about qualitatively and settled by seniority or volume. *Signature:* commitments

made without any capacity calculation; repeated overcommitment followed by repeated surprise. *Detection:* ask for the current `h`; the absence of an answer is the finding. *Fix:* D3.1.

F21 · The maintenance trap. *Damages `h` via `delta Q`.* `Q` has grown until the maintenance draw consumes the budget; all output is replacement and the stall is attributed to motivation. *Signature:* high activity, low novelty; every new initiative displacing an existing one; builders leaving. *Detection:* D3.10. *Fix:* shed `Q` or raise `eta`. Exhortation raises `C` and makes it worse.

F22 · The unshed stock. *Damages `h`, permanently.* Stored order accumulates monotonically because nothing has a removal path. *Signature:* a procedure manual nobody has read end to end; rules whose origin nobody knows. *Detection:* D3.13 — count items with no sunset date, and approval gates with a rejection rate near zero. *Fix:* D3.11, D3.12.

F23 · The ceiling misread. *Damages the credibility of the method.* `h` computed with `eta F(U_t)` instead of `eta ln N`, producing a smaller and more pessimistic figure; or `h > 0` read as a promise of growth rather than a licence. *Signature:* headroom figures that move week to week with meeting quality. *Detection:* check the formula and whether the feasibility clause was run. *Fix:* D3.4 — `h` uses the ceiling, feasibility uses the geometry, both in that order.

F24 · The rising gauge on a falling system. *Damages interpretation.* `h` rises while `Q` falls, because the collapsing maintenance draw releases budget; the rise is read as recovery. *Signature:* improving headroom, deteriorating output, shrinking stored order — every reading better than the last and every observable worse. *Detection:* plot `h` and `Q` together; divergence in this direction is the diagnosis. *Fix:* never read `h` without `Q` beside it. **This is the single most misread signal on the gauge and Section 9.1 works it in full.**

F25 · Exogenous accretion. *Damages `h` through a channel the model does not see by default.* `Q` arrives from outside the `eta S` channel — mandates, mergers, statute, external standards — and consumes headroom without ever having paid for its own maintenance, so `h_inf = eta(ln N − S) − A`. *Signature:* growing obligations with flat or falling capability; a remit that expands by acquisition rather than by decision. *Detection:* tally `Q` acquired by internal decision against `Q` acquired by external imposition and estimate `A` directly. *Fix:* shed the accreted `Q` where permitted, refuse further accretion, or negotiate the mandate.

F26 · The converged institution. *Damages `h`, at the fixed point.* A system at steady state with high `u` has `h_inf = eta(ln N − S)` near zero: it is performing excellently and has no capacity for anything new. *Signature:* high measured performance, high measured difference, complete inability to take on a new class of problem. *Detection:* D3.20. *Fix:* raise `eta`, which is the only term scaling `h_inf`; or shed `Q` to leave the fixed point deliberately; or accept the state as a legitimate mature equilibrium and say so.

F27 · Shedding without a use. *Damages nothing and achieves nothing.* Released headroom is never spent and the system re-accumulates the same stock at the same rate. *Signature:* periodic simplification initiatives whose effects are undone within one time constant; rule counts oscillating around a stable mean. *Detection:* plot `Q` over several time constants; oscillation without trend is the finding. *Fix:* D3.18.

F28 · Shedding across the ethical line. *Damages the standing of the method.* A shedding programme with a genuine positive computed return removes consent provisions, appeal rights, dissent protections or safety interlocks, because those are stored order and stored order costs `delta Q`. *Signature:* a headroom argument used to justify removing someone's ability to decline. *Detection:* check every candidate against D3.14 to D3.16. *Fix:* refuse. **The model itself gives no warning here — the computed return is genuinely positive — which is exactly why the constraint must sit outside the model.**

8.9 9 Results: three systems worked end to end

All three systems are constructed. No dataset was collected. Every number below is the output of the stated parameters under the model of Sections 3 to 6, carried through so that a reader can re-derive it. They demonstrate that the procedure runs and that it reports its own failures. They are not evidence that any organisation behaves this way.

8.9.1 9.1 A nine-person research group over twenty-six weeks

System. A nine-person applied research group, weekly periods, a ninety-minute all-hands each week and a variable number of pair sessions.

Survey. Enumerating genuinely distinct methodological lineages rather than job titles gives `N = 6`: experimental measurement, numerical simulation, statistical inference, hardware fabrication, domain theory, and software infrastructure. Three of the nine people duplicate an existing lineage, which is why the headcount is nine and the audited count is six — the effective-count audit applied before anything else [Curzi, 2026c]. A disagreement audit over eight weeks, on structured forecasting questions about experimental outcomes with independent pre-discussion estimates, gives `u = 0.55`. **That figure is an estimate from a self-administered instrument in a constructed example and carries no interval; a real survey owes one.**

```
S = ln(1 + (N−1)∗u) = ln(1 + 5∗0.55) = ln(3.75) = 1.3218 nats/week
ceiling: ln N = ln 6 = 1.7918
geometric shortfall: ln N − S = 0.4700
```

The group realises 74 per cent of the surplus its own category count would permit at `u = 1`. **That is good geometry, and the rest of this subsection is about why it does not help.** Retention loss is `delta = 0.05` per week, measured across a five-week shutdown with no input, a half-life of `ln 2 / (−ln(1−0.05)) = 13.5` weeks. Cost is `C = 0.12` per week in normalised order-units. The order observable `Y_t` counts live, owned, implemented methodological decisions and built artefacts in active use, with `Y_0 = 0.80`; **the proxy passes the D1.2 test because it fell during a busy but inconclusive fortnight, which an interaction measure could not have done.**

Diagnosis. Decomposing conversion:

```
eta_record = 0.15   -- an occasional email; no standing recorder
eta_bind   = 0.40   -- owners named sometimes, review dates never
eta_reach  = 0.50   -- what is bound often fails to reach the people who act

eta = 0.15 ∗ 0.40 ∗ 0.50 = 0.030
```

Before the intervention, with `eta = 0.030`:

```
eta∗S    = 0.03 ∗ 1.3218 = 0.0397 per week
delta∗Q0 = 0.05 ∗ 0.80   = 0.0400
C        = 0.1200

h_0 = eta∗ln N − delta∗Q_0 − C = 0.03∗1.7918 − 0.0400 − 0.1200 = −0.1062
```

`h_0 < 0`. The sustainability ceiling applies without qualification: no arrangement of these nine people, no change to who talks to whom, and no improvement in the quality of the discussion can prevent `Q` from declining. The group's geometry is excellent and it is irrelevant, because the conversion coefficient discards 97 per cent of what the geometry produces.

Trajectory under `Q_{t+1} = 0.95 Q_t + 0.0397 − 0.12`:

week	Q_t	h_t
0	0.8000	-0.1062
1	0.6797	-0.1002
2	0.5653	-0.0945
4	0.3535	-0.0839
8	-0.0101	-0.0657

week	Q_t	h_t
13	-0.3714	-0.0477
20	-0.7441	-0.0290
26	-0.9727	-0.0176

`Q` crosses zero at week 8. The steady state is `Q_inf = (0.0397 − 0.12)/0.05 = −1.607`, which is negative — meaning the regime is unsustainable or the linear model must be replaced. **The theorem does not hide this failure mode and neither does the survey report.** In the real group `Q` cannot go negative; what happens instead is that it hits a floor consisting of whatever is externally enforced and operates entirely on that.

Note the deceptive shape of the headroom column. Headroom *improves* every week, from -0.1062 to -0.0176, and it improves because `Q` is collapsing, which reduces the maintenance draw. **A system in decline shows rising headroom. Headroom rising while `Q` falls is not recovery; it is the maintenance burden shrinking because there is less left to maintain.** This is F24, and it is the single most misread signal on the gauge.

Two thresholds, computed before choosing the intervention.

```
eta needed for a non-negative steady state:  eta > C/S = 0.12/1.3218 = 0.0908
eta needed for h_0 > 0 at Q_0 = 0.80:        eta > (delta*Q_0 + C)/ln N
                                                 = (0.0400 + 0.1200)/1.7918 = 0.0893
```

The group needs `eta >= 0.091`, a factor of three, and nothing more. It is worth recording how much cheaper that brief is than the ones that would have been proposed without the arithmetic. Not *hire more people*: `N` is already 6 and adding a seventh distinct lineage buys `eta ln(7/6) = 0.0046`. Not *improve collaboration*: `u = 0.55` is already strong. Not *reduce meetings*: cutting `C` to zero still leaves `h = 0.0538 − 0.0400 = +0.0137`, technically positive but roughly a tenth of what the conversion fix delivers, and it would destroy the `u` the group depends on.

This is the paper's own counter-result and it is stated rather than buried. In its own leading worked case the levers this corpus is named for — geometry, difference, category count — are all inert, and the binding constraint is a clerical one.

Intervention. Three changes, all cheap, all aimed at the decomposition: a standing recorder rota on the weekly all-hands, ten minutes of one person's time, targeting recording 0.15 to 0.70; a decision log replacing minutes with mandatory owner, review date and reversal condition (D1.4 to D1.6), targeting binding 0.40 to 0.75; and a changed-default register (D1.7) requiring every methodological decision to name the template, script or fixture that implements it within one week or be marked unrealised (D1.8), targeting reach 0.50 to 0.80.

Projected `eta = 0.70 x 0.75 x 0.80 = 0.420`. **That is optimistic and is treated as such**: the conservative half of the gain is taken and `eta = 0.21` is projected, still more than double the threshold. **After the intervention:**

```
eta*S = 0.21 * 1.3218 = 0.2776 per week
h_0   = 0.21*1.7918 − 0.05*0.80 − 0.12 = 0.3763 − 0.0400 − 0.1200 = +0.2163
Q_inf = (0.2776 − 0.12)/0.05 = +3.1514
```

week	Q_t	h_t
0	0.8000	+0.2163
1	0.9176	+0.2104
2	1.0293	+0.2048
4	1.2362	+0.1945
8	1.5914	+0.1767

week	Q_t	h_t
13	1.9443	+0.1591
20	2.3084	+0.1408
26	2.5317	+0.1297

`Q` rises from 0.80 to 2.5317 in twenty-six weeks and converges on 3.1514. The time constant is `1/delta = 20` weeks, so the group is at 80 per cent of its new steady state by week 26. **Nothing changed except the conversion coefficient: `N`, `u`, `S`, `delta`, `C` and the headcount are identical. The same nine people, having the same conversations, with the same disagreement rate, at the same cost, move from terminal decline to a fourfold increase in stored order.**

Falsification criteria, registered before the intervention. The prediction is `a = 0.95` and `b = 0.21` in the minimal model, fitted on weeks 1 to 16 and frozen, with weeks 17 to 26 held out. The intervention is judged a failure if held-out prediction of `Y` is not better than the persistence baseline; if the fitted `b` is below 0.09, the sustainability threshold, meaning the programme did not move the coefficient enough to matter; if the shuffled-input control fits as well as the true pairing, indicating a trend rather than a consequence; or if the re-litigation count does not fall, indicating that reach is still zero despite records existing, which is F3 rather than success.

Standing caution. `Y` counts owned, implemented decisions, which a group could inflate by making trivial decisions and implementing them — the measure-becomes-target failure [Goodhart, 1975]. The registered defence is that the review-date field forces every entry to be re-examined and entries no longer corresponding to anything anyone does are struck from the count. **A metric with a scheduled removal path is substantially harder to inflate than one that only accumulates**, and the removal path is also the shedding mechanism of Section 6.9.

8.9.2 9.2 Sizing an incident-response function

System. A standing incident-response function inside a mid-size operations group, which must diagnose and resolve heterogeneous failures under time pressure. `Q` is the count of live, exercised runbooks and validated diagnostic procedures, normalised so that one unit is one such procedure verified within the last quarter.

Survey.

```
eta    = 0.22    a strong decision log and runbook discipline
kappa  = 0.006   4.2 h/period of synchronisation across 47 active links, normalised
delta  = 0.07    procedures go stale; measured by re-validation failure rate
Q      = 2.5     2.5 normalised units of live procedure
delta*Q = 0.175  the standing maintenance draw
```

Every added responder is assumed, generously, to bring a genuinely distinct failure-domain category, so `N = n`. An effective-count audit would normally reduce this; taking `N = n` biases the answer toward larger teams, which makes the conclusion conservative.

The sizing table, `h(n) = eta ln n − delta Q − kappa n(n−1)/2`:

n	eta ln n	C = kappa n(n−1)/2	h
2	0.1525	0.0060	-0.0285
3	0.2417	0.0180	+0.0487
4	0.3050	0.0360	+0.0940
5	0.3541	0.0600	+0.1191
6	0.3942	0.0900	+0.1292
7	0.4281	0.1260	+0.1271
8	0.4575	0.1680	+0.1145

n	eta ln n	C = kappa n(n−1)/2	h
10	0.5066	0.2700	+0.0616
12	0.5467	0.3960	-0.0243
16	0.6100	0.7200	-0.2850
20	0.6591	1.1400	-0.6559
30	0.7483	2.6100	-2.0367
50	0.8606	7.3500	-6.6644

```
net-optimal size:            n = 6         (h = +0.1292)
closed form sqrt(eta/kappa): 6.06          (agrees)
h > 0 range:                 n in [3, 11]
C exceeds eta*ln n at:       n = 15
```

The table contains three boundaries that are routinely confused with each other.

At `n = 2`, headroom is negative. A two-person function cannot sustain 2.5 units of procedure: there is not enough category diversity to generate the surplus that maintenance requires. **The lower bound is a real bound — understaffing is a sustainability violation exactly as overstaffing is.** At `n = 6` headroom peaks; that is the design point. **At `n = 11` headroom reaches zero**, and above eleven the team cannot maintain the procedures it already has. **At `n = 15` cost exceeds the entire structural budget.**

Note the shape between 6 and 11: headroom stays positive and falls by 52 per cent across those five hires. **Nothing observable breaks in that band. The team feels busier and produces less durable procedure, and the natural interpretation is that they need more people.** That interpretation is exactly backwards and it is the mechanism by which F11 propagates.

The eighteen-person question. Suppose the operations group requires eighteen responders for round-the-clock coverage — a genuine constraint. Under a complete graph:

```
C(18) = 0.006 * 18*17/2 = 0.006 * 153 = 0.9180
h(18) = 0.22*ln 18 - 0.175 - 0.9180 = 0.6359 - 0.1750 - 0.9180 = -0.4571
```

Deeply unsustainable. But `n` is fixed by the rota constraint, so three of the four responses are unavailable and topology is the only move. **Restructure as three cells of six with liaison links**, complete graph inside each cell and three inter-cell liaison links priced at four times a normal link because they carry translation and handover cost:

```
within-cell:   3 cells * (0.006 * 6*5/2)  = 3 * 0.0900        = 0.2700
between-cell:  3 links * 4 * 0.006                            = 0.0720
total C(split) = 0.3420        against C(monolith) = 0.9180

h(split, N = 18) = 0.22*ln 18 - 0.175 - 0.3420 = +0.1189
```

The same eighteen people, with the same skills, doing the same work, move from `h = -0.4571` to `h = +0.1189` — a swing of 0.576 — purely by changing who is obliged to keep a live channel with whom. No hiring, no firing, no training, no budget. Cost falls by 63 per cent and the surplus ceiling is untouched, because `N = 18` regardless of how the eighteen are wired.

Against the alternatives: raising conversion to rescue the monolith requires `eta' = (delta Q + C)/ln N = (0.175 + 0.918)/2.8904 = 0.378`, a 72 per cent improvement on a discipline that is already strong. Cutting `delta` to zero — perfect retention, an impossibility, and one the retention analysis shows to be a pathology in its own right [Curzi, 2026a] — saves only 0.175 and leaves `h = -0.282`. Shedding `Q` entirely, from 2.5 to 0, saves the same 0.175 and leaves `h = -0.282`. **Only the topology move closes the gap.**

What the split costs, stated honestly.

1. **Cross-cell `u` may fall.** Responders in one cell stop routinely encountering another cell's failure modes. The countermeasure is overlapping membership (D2.8): each responder belongs primarily to one cell and secondarily to another, which preserves the semilattice and costs roughly one additional link per person, `18 x 0.006 = 0.108`, taking total `C` to 0.4500 and `h` to `+0.0109`. Still positive, and the surplus is protected.
2. **Handover risk rises at cell boundaries.** The elevated liaison multiplier of four is the estimate of it. If handover proves more expensive — say eight times — total `C` becomes `0.2700 + 0.1440 + 0.108 = 0.5220` and `h = −0.0611`, which is negative. **The sensitivity of the answer to the liaison multiplier is high, and the liaison multiplier is the least well-estimated number in this survey.** It should be measured before the split is made permanent, not assumed. This is also where the coupled-network caution of Section 5.7 bites: Buldyrev et al. [2010] give the mechanism by which a thin coupling between otherwise robust cells becomes the system's failure mode, and this survey prices the coordination saving and does not price that fragility.
3. **The rota constraint was never tested.** Eighteen was given as a coverage requirement. The arithmetic says a six-person cell is the net optimum, and three cells of six is what the rota produces anyway. **The convergence is fortunate rather than designed and should not be assumed to recur.**

Falsification criteria, registered before the split.

1. Predicted: `Q` rises from 2.5 toward `Q_inf = (eta S − C)/delta`. With `u` estimated at 0.55 post-split, `S = ln(1 + 17 x 0.55) = 2.3370`, `eta S = 0.5141`, so `Q_inf = (0.5141 − 0.4500)/0.07 = 0.916`. **Note that this is *lower* than the current `Q = 2.5`.** The honest prediction is therefore that the split makes `h` positive while `Q` still falls, because `eta S` at the achievable `u` does not support 2.5 units of procedure against this `delta` and this `C`. The team must additionally raise `u`, raise `eta`, or **shed procedures deliberately** — the fourth lever, arriving unprompted from a pure cost analysis. **This is the single most useful thing the calculation produced, and it would not have been visible from `h` alone.**
2. The split fails if measured `C` after two quarters does not fall by at least half, indicating the liaison multiplier was underestimated.
3. The split fails if measured `u` falls below 0.45, indicating that the overlapping-membership provision failed and F15 has occurred in cellular form.
4. The split fails if the shuffled-input control fits the post-split series as well as the true pairing.

Standing caution. Positive headroom licenses building; it does not promise that what exists is sustainable. **This example makes both points at once: the split buys headroom, and the same arithmetic reveals that the existing stock of procedure was never supportable at any topology with these parameters. The cost analysis rescued the organisation and condemned its backlog, and it did both with the same four numbers.**

8.9.3 9.3 A professional standards body over sixty years

System. A membership organisation that publishes and maintains technical standards for a regulated trade. Annual periods. `Q` is the count of standards actively maintained — published, in force, and re-validated within the review cycle — normalised so that one unit is one such standard.

Survey.

```
eta   = 0.20     conversion of committee work into published, adopted standard
N     = 14       genuinely distinct technical sub-domains (the initial claim was 31,
                 reduced by the collision test of the effective-count audit)
u     = 0.50     measured by pre-discussion disagreement on technical questions
delta = 0.045    standards go stale; measured by re-validation failure over a dormant biennium
```

```
C_0   = 0.08     baseline operating cost: governance, publication, administration
gamma = 0.030    cost scaling with the standards count: re-validation, cross-reference
                 maintenance, enquiry handling, errata
A     = 0.25     EXOGENOUS ACCRETION: standards adopted by statute and by international
                 harmonisation, which the body must maintain and did not choose
```

Derived:

```
S           = ln(1 + 13*0.50) = ln(7.5)  = 2.0149 nats/yr
eta*S       = 0.4030
budget      = eta*ln N = 0.20 * 2.6391    = 0.5278
delta_eff   = delta + gamma               = 0.075
Q_inf       = (eta*S - C_0 + A)/delta_eff = (0.4030 - 0.08 + 0.25)/0.075 = 7.6397
Q_crit      = (budget - C_0)/delta_eff    = (0.5278 - 0.08)/0.075       = 5.9708
h_inf       = eta*(ln N - S) - A          = 0.20*0.6242 - 0.25          = -0.1252
```

The two numbers that decide the case are on adjacent lines: `Q_inf = 7.6397` and `Q_crit = 5.9708`. The body will converge to a stock 28 per cent larger than the largest stock it can sustain. That is the diagnosis, available in year zero, sixty years before the symptoms.

Trajectory, `Q_{t+1} = 0.955 Q_t + 0.4030 - (0.08 + 0.030 Q_t) + 0.25`:

year	Q_t	C_t	delta Q_t	h_t
0	0.0000	0.0800	0.0000	+0.4478
5	2.4662	0.1540	0.1110	+0.2628
10	4.1363	0.2041	0.1861	+0.1376
15	5.2672	0.2380	0.2370	+0.0528
20	6.0331	0.2610	0.2715	-0.0047
25	6.5518	0.2766	0.2948	-0.0436
30	6.9030	0.2871	0.3106	-0.0699
40	7.3019	0.2991	0.3286	-0.0998
50	7.4848	0.3045	0.3368	-0.1135
60	7.5687	0.3071	0.3406	-0.1198

Headroom crosses zero in year 20, at `Q = 6.033`, just past `Q_crit = 5.971`. Consider what year 20 looks like from inside: it is the twentieth consecutive year of growth, the catalogue has grown every year, membership is at a record, and by every visible measure the organisation has succeeded. **There is no event. The crossing has no signature. It is a sign change in a quantity nobody is computing.**

By year 40 the condition is recognisable: `delta Q = 0.3286` and `C = 0.2991`, together 0.6277 against a budget of 0.5278. The maintenance draw alone is 62 per cent of the total budget, well past the D3.10 threshold. The body publishes little that is new, spends its committee time on revision and errata, and is described by its own members as bureaucratic. **That description is accurate and it is not an explanation.**

Diagnosis. Of the stored order accumulated, the accretion term `A = 0.25` per year against internal generation `eta S = 0.4030` per year means **38 per cent of this body's stored order arrived by statute and harmonisation rather than by its own decision, and none of it came with maintenance funding.** The fixed-point identity prices it exactly:

```
h_inf = eta*(ln N - S) - A = +0.1248 - 0.2500 = -0.1252
```

Without `A` the organisation would sit at `h_inf = +0.1248` indefinitely — converged, maintained, modest, and permanently solvent. **The entire deficit is the accretion term.** That is a

very different finding from *the organisation became complacent*, and it points at a very different remedy.

The four levers, priced at year 60, with `Q = 7.5687`, `C = 0.3071`, `h = −0.1198`.

L1, raise conversion. Required `eta' = (delta Q + C)/ln N = (0.3406 + 0.3071)/2.6391 = 0.2454`, from 0.20 — **a 23 per cent improvement.** Plausible, since standards drafting is a capture process to which the mechanisms of Section 4.4 apply directly. This is the best of the first three levers.

L2, raise the category count. Required `N' = exp((delta Q + C)/eta) = exp(0.6477/0.20) = exp(3.2385) = 25.5`, from 14 — **eleven more genuinely distinct technical sub-domains, an 82 per cent expansion of remit.** And note the perversity: expanding remit is precisely the process that generates more `A`. **Lever 2 in this case makes the disease worse while treating the symptom**, which is a structural feature of accretion-driven ossification and not a quirk of these numbers.

L3, cut cost. Required `C' <= budget − delta Q = 0.5278 − 0.3406 = 0.1872`, from 0.3071 — a 39 per cent cut. But decompose: `C = C_0 + gamma Q = 0.08 + 0.2271`. **Only 0.08 of the 0.3071 is discretionary; 0.2271 is driven by `Q` itself. Even abolishing the entire baseline operating cost leaves `C = 0.2271` and `h = −0.0399`, still negative.** Lever 3 cannot reach the target at any level of austerity. This is the arithmetic behind the familiar spectacle of an institution cutting itself to the bone and remaining stuck.

L4, shed stored order. Required `Q' <= Q_crit = 5.9708`, from 7.5687 — **a 21 per cent reduction, and the deficit is closed.** Because `gamma > 0`, shedding attacks both draws at once at rate `delta + gamma = 0.075` per unit. Take the shed further, to `0.75 x Q_crit = 4.4781`, retiring 41 per cent of the catalogue:

```
new C = 0.08 + 0.030*4.4781            = 0.2143
new h = 0.5278 − 0.045*4.4781 − 0.2143 = +0.1120

headroom recovered = +0.1120 − (−0.1198) = +0.2318
```

A swing of 0.2318, which is 44 per cent of the organisation's entire structural budget, recovered by retiring standards that no longer describe practice. Achieving the same swing through the category lever would require `N' = 14 exp(0.2318/0.20) = 45` — tripling the technical remit. **Retiring 3.09 units of obsolete standard is worth more than tripling the organisation's technical scope, and it can be decided in one committee meeting.**

Execution. The shed of 3.09 units, roughly two fifths of the catalogue, is executed against the classification of Section 6.9.2 and not by percentage target.

- *Freely sheddable:* standards for equipment classes no longer manufactured; standards superseded by a later standard but never formally withdrawn; guidance notes with no recorded citation in a decade; three committees whose founding question was resolved. **The never-formally-withdrawn category alone is typically the largest and is pure F22.**
- *Conditionally sheddable, with replacement:* four standards consolidated into one, retaining each original's one-line rationale as an annex (D3.17). This is the *shed the rule, keep the reason* form and it is what makes the consolidation defensible to the members whose work is being retired.
- *Not shed, and named as not shed:* every statutory obligation, since the obligation runs to parties not represented in this decision, so it is renegotiated with the regulator or it stands (D3.16); every safety-critical interlock standard (D3.15); and the members' right of objection to a published standard, which is stored order costing `delta Q` every year and is **not available regardless of the computed return** (D3.14). **The model gives no warning here. The computed return on removing the objection procedure is genuinely positive. The constraint sits outside the model on purpose.**

Preventing the regrowth. Shedding is a transient. The contraction guarantees a return to `Q_inf = 7.6397` at time constant `1/0.075 = 13.3` years, so the released headroom must be spent

before it is reabsorbed. The registered plan: spend `h = +0.1120` on the conversion lever over years 61 to 66, targeting `eta` from 0.20 to 0.26 through a drafting decision log with mandatory reversal conditions, which raises the budget from 0.5278 to 0.6862; require from year 61 that **every new standard carries a sunset date with removal as the review default**, which is the structural fix rather than the symptomatic one and the only change that prevents recurrence; and open negotiation with the regulator on `A` over years 61 to 63, seeking either maintenance funding attached to statutory adoptions or a formal withdrawal mechanism. **`A` is the actual cause and it is the only item on the list outside the body's unilateral control. Every other action treats a symptom, and the survey says so rather than claiming a cure.**

Falsification criteria, registered before the intervention.

1. Predicted: `Q` falls to 4.478 at the shed and, absent the sunset provision, regrows toward 7.6397 with time constant 13.3 years, reaching approximately 6.0 by year 74 and recrossing `Q_crit`. **The sunset provision is falsified if `Q` follows that trajectory anyway.**
2. Predicted: measured `C` falls from 0.3071 to approximately 0.2143 within two review cycles. If it does not, `gamma` was misestimated and the double-return argument does not apply here.
3. Predicted: fitted `b` rises from 0.20 toward 0.26 following the decision-log intervention. If `b` does not move, the conversion programme failed and the window will close with nothing bought — F27.
4. Predicted: fitted `a` remains near `1 − delta_eff = 0.925`. **A significant change in `a` indicates that `delta` or `gamma` was not stable across the horizon, which falsifies H1 and requires the model to be replaced rather than refitted.**
5. The shuffled-input control must degrade the fit. If it does not, the entire fitted relationship between committee interaction and standards output is a trend and the conversion estimate was never real.

Standing caution. Every number here follows from six parameters and the theorems of the storage law. **The parameters are the argument; the arithmetic is only bookkeeping.** An organisation disputing the conclusion should dispute `eta`, `N`, `u`, `delta`, `gamma` or `A` — and to dispute them it must measure them, which is the entire purpose of putting the gauge on the wall.

8.10 10 Discussion

8.10.1 10.1 What follows

Restated in the order they are consumed downstream.

The headroom gauge `h_t = eta ln N − delta Q_t − C_t` reads three parameters on one dial with one sign, with three regimes, a maximum one-step increment `r_max = h_t`, and two separate constraints that are routinely conflated: admissibility `C_t <= (1−delta)Q_t + eta S_t` and sustainability `C_t <= eta ln N − delta Q_t`. A system can satisfy the first and violate the second for decades.

Doubling conversion is exactly squaring the category count, since `2 eta ln N = eta ln(N^2)`; and doubling the surplus by raising non-parallelism alone is feasible only when `u <= 1/(sqrt(N)+1)`, so at `N = 8` and `u = 0.30` the required `u'` is 1.23 and the move is impossible rather than expensive. **This is arithmetic, not a finding about organisations**, and it is the reason conversion is the first move in almost every system this framework describes.

Conversion composes multiplicatively through recording, binding and reach, so three gates at 0.7 give `eta = 0.343` and the weakest gate dominates: recording 0.20 to 0.80 gains 300 per cent against binding 0.60 to 0.90 at 50 per cent. **The decomposition itself is a modelling convenience with no external support**; the multiplicative arithmetic conditional on it is exact.

`eta = 0` produces decline, not stagnation. The stock is non-increasing in every period and falls at exactly the cost rate, converging on `−C/delta`, which is inadmissible and marks the model's own domain of breakdown.

The net optimum sits far below the crossover. Maximising `eta ln n − kappa n(n−1)/2` gives `eta = kappa n(n − 1/2)`, approximately `sqrt(eta/kappa)`; at `eta = 0.30` and `kappa = 0.010` the optimum is six and the crossover is thirteen. **The band between them is where nothing visibly breaks and every additional member reduces net output**, and it is where most organisations live.

Topology is a free lever. Eighteen people rewired from a complete graph into three cells of six move from `h = −0.4571` to `h = +0.1189`, a swing of 0.576, with cost falling 63 per cent and the surplus ceiling untouched because `N = 18` either way.

The nine-person case returns `N = 6, u = 0.55, S = 1.3218` nats — 74 per cent of `ln 6` — with measured `eta = 0.030` giving `h_0 = −0.1062` and `Q_inf = −1.607`. The binding threshold is `eta >= 0.091`, a factor of three.

A cost that scales with stored order is a decay rate. With `C = C_0 + gamma Q`, the system behaves exactly as if `delta_eff = delta + gamma`, so every result of the retention analysis transfers with that substitution [Curzi, 2026a].

Two of these deserve emphasis because they are most likely to be restated loosely elsewhere. The doubling identity is exact and is not an analogy. And the fixed-point identity `h_inf = eta(ln N − F(u∗))` is a cancellation, not an approximation: `delta`, `C` and `Q` all vanish, and the consequence — that **the plain linear model never ossifies by itself** — rules out the most popular explanation for institutional ossification and forces the analyst to name which of three imported mechanisms is operating.

8.10.2 10.2 What does not follow

No thermodynamic claim is made. Landauer [1961] establishes a physical lower bound on the energy cost of erasure and is the reason the cost analogy is tempting; it is not a warrant for any claim about institutional cost, and it is not used as one.

Surplus is not identified with Shannon entropy. The `ln N` ceiling has an information-theoretic reading and the reading bounds; it does not identify.

The mapping onto any real organisation is not established by anything in this paper. All three worked systems are constructed. What would license applying these results to an actual organisation is a preregistered fit with held-out comparison against declared baselines, and the paper specifying that protocol reports its own confirmatory application **failing** its second rule [Curzi, 2026i]. A reader who takes this paper's arithmetic as evidence that organisations obey the storage law has read it wrongly.

Monotonicity in `u` is not established here and is contested. The framework's surplus is monotone increasing in `u`, and the empirical literature does not agree that the benefit of difference rises without turning over. Uzzi and Spiro [2005] find an inverted U, with creative success rising and then falling with cohesion; Mannix and Neale [2005] and van Knippenberg and Schippers [2007] find that surface-level diversity frequently costs in the short run; Putnam [2007] finds short-run reductions in trust and civic participation, with Abascal and Baldassarri [2015] as the rebuttal; and Sagan [1993] documents settings in which redundancy reduces safety. Against these, Granovetter [1973], Burt [2004] and Schindler et al. [2010] supply measured benefits of bridging and of effective count. **The honest position is that the monotone form is this framework's construction over the regime it has examined, and where the curve turns is an open empirical question.** Harrison and Klein [2007] sharpens the demand further: three incompatible constructs travel under the word diversity — separation, variety and disparity — and `u` owes the reader a statement of which one it is. **This paper's `u` is a variety measure with a disparity weighting and it is stated as such here for the first time in this series.**

Diversification benefit depends on covariance, not count. Markowitz [1952] stated this in 1952 with the mathematics attached, and it is the same content as the effective-count result: `N` uncorrelated holdings and `N` perfectly correlated holdings have the same nominal `N` and different effective counts. **Any reading of `eta ln N` as a benefit of headcount rather than of *audited* headcount is a misreading**, and the audit that makes the distinction computable is developed separately [Curzi, 2026c,?].

The stability benefit of an effective count is not free of a purely statistical explanation. Doak et al. [1998] showed that diversity–stability correlations can arise from statistical averaging alone with no mechanism; May [1972] and Gardner and Ashby [1970] showed that increasing diversity and connectance destabilises randomly assembled systems; McCann [2000] is the fair survey; and Loreau and de Mazancourt [2008] names the decisive variable, synchrony, which is the ecological name for what this framework calls a coordination perimeter.

The `n_opt = sqrt(eta/kappa)` rule is not a law of team size. It is the stationary point of a model whose cost term is a badged modelling choice and whose benefit term is a functional form the evidence does not currently discriminate from at least two alternatives. Change the cost exponent and the rule changes.

8.10.3 10.3 Relation to the surrounding literature

The three terms of this paper have different standing in the external literature and the difference should be stated.

Cost has the strongest grounding. The quadratic pair count and its interpretation as coordination burden are Brooks [1975] and Conway [1968]; the level–rate vocabulary is Forrester [1961]; the finding that people misjudge such systems even under aligned incentives is Sterman [1989]; the queueing result that makes slack a control device rather than waste is Kingman [1962], Little [1961] and Hopp and Spearman [2011]; and the reserve itself was named organisational slack by Cyert and March [1963]. Nothing in Section 5 is novel as a mechanism; what is new is pricing these against a logarithmic benefit on one commensurable scale.

Conversion has the weakest grounding and the paper says so. The three-gate decomposition appears in no external literature. What external work supports is the *class* of devices: the design principles associated with durable common-pool institutions [Ostrom, 1990, 2010, Cox et al., 2010] read as conversion and retention devices; the invisibility of working infrastructure [Star and Ruhleder, 1996, Star and Griesemer, 1989]; and the finding that a group performance factor exists and is distinct from member ability [Woolley et al., 2010] — which is what the conversion coefficient is being asked to stand in for, and which Hong and Page [2004] illustrates and Thompson [2014] sharply criticises. **Those two must be cited together or not at all**, and Singer [2019] is the repair that survives the objection.

The gauge is this framework's own and its nearest relative is a budget constraint. Its content is that three quantities usually argued about separately are commensurable, which is a modelling claim rather than a discovery.

Within the corpus, these results are consumed in four ways. They are **reused as an instrument** in the built-environment analyses, whose entrance-sequence case moves interruptions from 22 to 6 per person-day and headroom from `+0.0642` to `+0.2955` for 7.6 per cent of net floor [Curzi, 2026f], and in the channel analyses, where forced integration raised realised `u` from 0.1155 to 0.3465 and surplus by 121 per cent while cost rose 279 per cent, taking the step from `−0.1953` to `−0.3714` [Curzi, 2026g]. They are **restated as a ceiling on roster size**: the roster ceiling `m(m−1) <= 2 eta ln(N)/kappa`, with a constructed commons sitting exactly at it, 30 against 30.7225, so that a seventh member makes the case structural [Curzi, 2026]. They are **bounded by a completeness result**: the steady state depends on exactly `u`, `N`, `eta` and `c` through `F`, plus `delta` as a divisor, with no fifth argument and `delta` not a lever on the deficit [Curzi, 2026]. And they are **shown to be order-sensitive**: eight periods of the wrong lever order cost twelve periods of permanent lag, a minimum stock of `−2.373029` from `+1.000000`, and 111.133601 stock-periods of cumulative gap [Curzi, 2026]. **That last result is the sharpest limitation on the four-lever framing of Section 6.8, which prices the levers and does not price their sequence.**

Two results from the structural series bear directly on Section 5's cost analysis. The disclosure-perimeter analysis scores a five-criterion test as a **product rather than an average**, so that five structures averaging 0.80 yield a group product of zero because two contain a zero-scoring element [Curzi, 2026] — which is the same arithmetic as the three-gate composition of Section 4.3 arriving in a different domain, and is the strongest independent support this

paper has for its multiplicative form. And the option-pricing analysis of an uncertain jurisdictional call gives a cost ceiling `p = (c + 0.0051096)/0.0318401`, so that a leg is worth nothing at any probability once its incremental cost exceeds 0.0267305 — 8.9 per cent of that group's entire maintenance budget [Curzi, 2026]. **That is the liaison-multiplier sensitivity of Section 9.2 stated as a decision rule, and it is the form the sensitivity analysis should take when a cost is uncertain rather than merely estimated.**

Finally, the enterprise itself has a literature and it is not entirely friendly. Castellano et al. [2009] surveys and bounds the programme of carrying physics formalism into social systems and supplies definitions this framework can be held to; Helbing and Molnár [1995] and Helbing [2001] set the evidentiary bar by validating a social force analogy against measured trajectories, and **the standard they set is prediction of observed data, not internal consistency.** Gallegati et al. [2006] is the field's own indictment and reads as a checklist against this paper: ignorance of existing work, sub-standard statistical procedure, unjustified universality claims from limited data. Hedström and Ylikoski [2010] states the criterion sociologists will apply — an explanation must name a mechanism, not exhibit a fitted curve — which this paper meets in form, since every result names a mechanism, and does not meet in evidence, since no curve was fitted to anything observed. Durlauf [2005] asks for the prediction that only this model makes, and Sobkowicz [2020] asks it in one line: what observation would distinguish these laws from a neighbouring model that fits equally well? **Section 12 records that this paper does not have an answer.** Epstein [2008] is the standing reply to the charge that such models forecast nothing — prediction is one of sixteen reasons to model — and the reason this paper claims explanation and design rather than forecast.

8.11 11 Verification

The results of this paper compound from six postulates. Each is atomic, each is checkable on its own, and each carries its status. **Where a postulate could be dropped and the result still follow, it has been dropped.**

8.11.1 11.1 Dimensional audit

The equation must balance before anything else about it is worth checking. `[Q]` denotes the dimension of stored order and `[T]` the dimension of time; the period is the unit of `[T]` and every rate below is *per period*, which is why no explicit `[T^-1]` appears once the equation is written in discrete steps.

Symbol	Dimension	Unit	Status
`Q_t`	`[Q]`	proxy units (standards, runbooks, owned decisions)	**measured**, in whatever unit the declared order proxy counts
`U_t`	`[1]`	none	**defined** dimensionless: `u = 1 - (x.y)^2` with unit vectors
`S_t`	`[1]`	nats	**derived** dimensionless: a logarithm of an effective count
`eta`	`[Q]` per nat per step	proxy units per nat per period	**required by the balance**; see below

Symbol	Dimension	Unit	Status
`delta`	`[1]`	per period	**measured** dimensionless fraction per step
`C_t`	`[Q]` per step	proxy units per period	**measured**
`kappa`	`[Q]` per step per link	proxy units per period per active link	**measured**
`gamma`	`[1]`	per period	**measured** dimensionless fraction per step, like `delta`
`A`	`[Q]` per step	proxy units per period	**measured**
`h_t`	`[Q]` per step	proxy units per period	**derived**
`Q_crit`, `Q_inf`	`[Q]`	proxy units	**derived**: `[Q]/step` over `[1]/step` returns `[Q]`

Does the equation balance? `Q_{t+1} = (1−delta) Q_t + eta S_t − C_t` balances if and only if every term on the right carries `[Q]`. The first does, since `delta` is dimensionless; the third does by construction; the second does only if `eta S_t` carries `[Q]`, and since `S_t` is a dimensionless count in nats, **`eta` must carry `[Q]` per nat per step.** The gauge then balances too: `ln N` is dimensionless, so `eta ln N` carries `[Q]` per step and is commensurable with `delta Q_t` and with `C_t`, which is what makes the three-way subtraction well formed. And `eta/delta` carries `[Q]` per nat, which is the DC gain in the proper sense: units of stock per sustained nat per period.

The source is silent on this and its silence is not neutral. The source manuscript states that `eta` has units of `[Q]/[S]` and that, `S` being dimensionless, `eta` carries whatever unit the order proxy carries. **That is the `[Q]` per nat half and it omits the per-step half entirely.** The source then advises normalising `Q` so that `eta` lands in `[0,1]` and adds that `eta > 1` should be read as a unit choice rather than a violation — advice that treats `eta` as if it were a pure number. **The reading required for the law to balance is `[Q]` per nat per period, it is stated here explicitly, and it is recorded that the source does not supply it rather than being assigned quietly.** The practical consequence is the one Section 5.9 flags: a crossover table computed with `kappa` in hours, `eta` in normalised order-units and `Q` in document counts is dimensionally malformed however tidy the arithmetic looks, and the denomination requirement is part of the survey protocol [Curzi, 2026i].

Where the framework has no operational unit, it is said here. `Q` is stored order, and stored order has no natural unit: it is counted in whatever the declared proxy counts, and two systems' `Q` values are not comparable unless they declare the same proxy. **Consequently `eta` and `kappa` are not comparable across systems either**, and a table of `eta` values drawn from different surveys is meaningless. Comfort itself, the underlying construct of the corpus, has **no operational unit at all**; it is treated as a dimensionless magnitude by stipulation, not by measurement, and Section 12 records that the quantitative structure required for that stipulation has not been tested.

8.11.2 11.2 The postulates

P1 — The recurrence is linear and first order. `Q_{t+1} = a Q_t + x_t` with `a` and `x_t` not depending on `Q`. *Status:* **DEFINITION** — the model form adopted by the framework, with `a := 1 − delta` and `x_t := eta S_t − C_t`. It carries no empirical weight on its own.

P2 — Surplus per period is bounded: `0 <= S_t <= ln N`. *Status:* **CITED** — the range result of the surplus functional [Curzi, 2026], following from `f(u) = ln(1+(N−1)u)` on `u` in `[0,1]`.

P3 — `ln` is strictly increasing and strictly concave, and `ln(x^2) = 2 ln x`. *Status:* **ARITHMETIC.**

P4 — In a group of `n`, the number of unordered pairs is `n(n−1)/2`. *Status:* **ARITHMETIC** — combinatorics, and true independently of any claim about cost.

P5 — Coordination cost is proportional to the number of active communication paths, `C1 = kappa P`. *Status:* **POSTULATED** — a modelling choice of this framework. The interpretive step from the pair count to a cost is Brooks [1975]'s observation from practice, presented there as observation and not as proof. It is the only postulate in the chain with no theorem and no controlled measurement behind it, and every result about crossover, optimum and topology consumes it.

P6 — Cost accrues every period whether or not surplus was delivered. *Status:* **POSTULATED** — an assumption about how real systems charge overhead. It is used at exactly one place, R6 below, and the chain's other results do not depend on it.

The three-gate decomposition of `eta` is **not** in this list, because no result in the chain consumes it. It is a diagnostic device, and it is badged as such wherever it appears.

8.11.3 11.3 The compounding chain

```
P1 + P2        -> L1   the per-period increment eta*S_t - C_t is at most
                       eta*ln N - C_t, attained only at u = 1.              ARITHMETIC

L1             -> R1   a nondecreasing step requires C_t <= eta*ln N - delta*Q_t;
                       define h_t := eta*ln N - delta*Q_t - C_t, so that
                       h_t >= 0 is exactly that condition, and the maximum
                       one-step increment is r_max = h_t.                   ARITHMETIC

P3             -> R2   2*eta*ln N = eta*ln(N^2), so doubling eta and squaring
                       N give the identical budget.                         ARITHMETIC

P3 + P2        -> R3   doubling S needs (1+(N−1)u)^2 = 1+(N−1)u', so
                       u' = ((1+(N−1)u)^2 − 1)/(N−1); u' <= 1 requires
                       (1+(N−1)u)^2 <= N, i.e. u <= 1/(sqrt(N)+1).          ARITHMETIC

P4 + P5        -> L2   coordination cost under complete-graph obligation is
                       C(n) = kappa*n(n−1)/2.                               POSTULATED (P5)

L2 + P3        -> R4   net(n) = eta*ln n − kappa*n(n−1)/2 is stationary at
                       eta = kappa*n*(n − 1/2), so n_opt ~ sqrt(eta/kappa);
                       the crossover, where C(n) first exceeds eta*ln n,
                       is strictly larger.                                  ARITHMETIC given L2

L2             -> R5   the path count, and therefore C, depends on the wiring
                       plan and not on n alone; N is unchanged by rewiring,
                       so eta*ln N is unchanged and dh = −dC exactly.       ARITHMETIC given L2

R1 + P1        -> R6   at the fixed point Q_inf = (eta*F(u*) − c*)/delta,
                       h_inf = eta*ln N − delta*Q_inf − c* = eta*(ln N − F(u*)):
                       delta, C and Q cancel identically.                   ARITHMETIC

R6 + P6        -> R7   since F(u) <= ln N by P2, h_inf >= 0 always, so the
                       constant-parameter linear model never reaches negative
                       headroom; ossification must be imported.              ARITHMETIC
```

Every line is checkable from its own inputs. R1, R2, R3, R6 and R7 are arithmetic consequences of one definition, one cited range result and elementary properties of the logarithm. **R4 and R5 — the crossover, the optimum, and the topology lever — consume P5, which has no external support beyond an observation from practice.** That is the load-bearing weakness of Section 5 and it is stated here rather than left to be found: **if coordination cost is not superlinear in the path count, the sizing results of this paper do not hold, and nothing in this paper measures whether it is.**

What this chain does not establish. It does not establish that any organisation obeys P1. That is a modelling postulate, it is not verified here, and the one preregistered confirmatory test of it so far reported in this corpus did not meet its own declared margin [Curzi, 2026i]. **A chain whose steps are ARITHMETIC, CITED and POSTULATED delivers the mathematics with certainty and delivers nothing whatever about institutions.**

8.11.4 11.4 The checks a reviewer can redo in a minute

The doubling identity. `ln 8 = 2.079442`, so `0.10 x 2.079442 = 0.2079442`. And `ln 64 = 4.158883`, so `0.10 x 4.158883 = 0.4158883`, while `0.20 x 2.079442 = 0.4158883`. **The two are equal to every digit printed, because `ln(8^2) = 2 ln 8` exactly.** ARITHMETIC.

The infeasibility threshold. At `N = 8`, `1/(sqrt(8)+1) = 1/3.828427 = 0.261204`. At `u = 0.30 > 0.261204` the doubling is infeasible, and directly: `S = ln(1 + 7 x 0.30) = ln 3.10 = 1.131402`, so doubling requires `u' = (3.10^2 − 1)/7 = (9.61 − 1)/7 = 8.61/7 = 1.230000 > 1`. **Impossible.** ARITHMETIC.

The gate product. `0.20 x 0.60 x 0.70 = 0.084`; `0.20 x 0.90 x 0.70 = 0.126`, a rise of 50.0 per cent; `0.80 x 0.60 x 0.70 = 0.336`, a rise of 300.0 per cent. ARITHMETIC.

The optimum against the closed form. At `eta = 0.30`, `kappa = 0.010`: `sqrt(0.30/0.010) = sqrt(30) = 5.477`, and direct search over integers gives `net(5) = 0.30 ln 5 − 0.005 x 20 = 0.482831 − 0.100 = 0.382831`, `net(6) = 0.537528 − 0.150 = 0.387528`, `net(7) = 0.583773 − 0.210 = 0.373773`. **The optimum is 6.** The crossover is the smallest `n` with `0.005 n(n−1) > 0.30 ln n`: at `n = 12`, `0.66 < 0.745472`; at `n = 13`, `0.78 > 0.769485`. **The crossover is 13.** ARITHMETIC.

The topology swing. `0.006 x 18 x 17/2 = 0.006 x 153 = 0.918`; `3 x (0.006 x 15) + 3 x 4 x 0.006 = 0.270 + 0.072 = 0.342`. With `eta ln 18 = 0.22 x 2.890372 = 0.635882` and `delta Q = 0.175`: `h = 0.635882 − 0.175 − 0.918 = −0.457118` and `h = 0.635882 − 0.175 − 0.342 = +0.118882`. **The swing is 0.576000 and the cost falls by `(0.918−0.342)/0.918 = 62.7` per cent.** ARITHMETIC.

The fixed-point identity, checked on a row of the table. `eta = 0.18`, `N = 12`, `u* = 0.45`, `delta = 0.06`, `C = 0.22`. Then `S = ln(1 + 11 x 0.45) = ln 5.95 = 1.783391`, `Q_inf = (0.18 x 1.783391 − 0.22)/0.06 = (0.321010 − 0.22)/0.06 = 1.683509`, and `h_inf = 0.18 x 2.484907 − 0.06 x 1.683509 − 0.22 = 0.447283 − 0.101011 − 0.22 = 0.126273`. Against `eta(ln N − S) = 0.18 x (2.484907 − 1.783391) = 0.18 x 0.701516 = 0.126273`. **Identical to six decimal places, as the cancellation requires.** ARITHMETIC.

The nine-person case, end to end. `S = ln(1 + 5 x 0.55) = ln 3.75 = 1.321756`; `eta = 0.15 x 0.40 x 0.50 = 0.030`; `h_0 = 0.030 x 1.791759 − 0.05 x 0.80 − 0.12 = 0.053753 − 0.040 − 0.120 = −0.106247`; `Q_inf = (0.030 x 1.321756 − 0.12)/0.05 = (0.039653 − 0.120)/0.05 = −1.606943`; threshold `C/S = 0.12/1.321756 = 0.090789` and `(delta Q_0 + C)/ln N = 0.16/1.791759 = 0.089298`, so the binding requirement is `eta >= 0.0908`, **a factor of 3.03 on a measured 0.030.** ARITHMETIC.

The standards body's two decisive numbers. `S = ln 7.5 = 2.014903`; `ln 14 = 2.639057`; `delta_eff = 0.045 + 0.030 = 0.075`; `Q_inf = (0.20 x 2.014903 − 0.08 + 0.25)/0.075 = 0.572981/0.075 = 7.639747`; `Q_crit = (0.20 x 2.639057 − 0.08)/0.075 = 0.447811/0.075 = 5.970816`. **The ratio is 7.639747/5.970816 = 1.2795, so the body converges 28.0 per cent above what it can sustain.** ARITHMETIC.

8.11.5 11.5 Limits, asymptotics, and the model's valid domain

`eta -> 0`. The drive term vanishes and `Q_inf = -C/delta <= 0`: the system decays to the admissibility floor at rate `delta` while paying `C`. **Zero conversion decouples the input from the stock entirely**, so no interaction geometry of any quality changes the outcome.

`eta -> infinity`. `Q_inf` and `h` grow without bound, which the model permits and which is a defect of the linear form rather than a result. Real capture saturates, and the linear model excludes saturation by assumption. **The model has no upper domain limit on `eta` and should be given one empirically before any extrapolation.**

`N -> 1`. `ln N = 0` and `F(u) = 0` for every `u`. **A single audited category has zero available surplus, not a small amount.** The ceiling and the achievable value collapse together, so `Q_inf = -C/delta <= 0` and headroom is negative at any positive stock. This is the boundary that makes the collapse operation of the concentration instrument consequential: a structure with one effective bloc has `ln(1.00) = 0` available surplus [Curzi, 2026].

`N -> infinity`. `eta ln N -> infinity`, so the per-period budget grows without bound. **Taken literally this says that a sufficiently large audited category count makes any cost sustainable, and that is false.** Three things bound it and all three are domain restrictions rather than results.

1. *Internal.* Coordination cost grows with the pair count, which is quadratic under P5, so the two curves cross at a size that depends only on `eta` and `kappa` — this paper's own crossover. Beyond the crossing, additional categories cost more than the ceiling they buy.
2. *External and contrary.* Uzzi and Spiro [2005] find an inverted U: creative success rises and then *falls* with cohesion, so the benefit of difference turns over rather than growing without bound. Sagan [1993] documents settings in which redundancy reduces safety, and May [1972] and Gardner and Ashby [1970] show that increasing diversity and connectance destabilises randomly assembled systems. **The unbounded growth of `eta ln N` contradicts this literature and the contradiction is not resolved here.**
3. *Marginal.* Even inside the model the marginal gain `ln N - ln(N-1) ~ 1/N` vanishes, so the divergence is the slowest available and is never a practical route.

The stated domain is therefore `N` bounded above by the crossover of Section 5.4, and the logarithmic ceiling is a within-domain approximation and not an asymptotic claim. It is recorded as a restriction on where the model may be used, not as a finding about large groups.

`u -> 0`. `N_eff -> 1`, `F -> 0`, and every conversion intervention returns nothing. **`u -> 1`.** `F -> ln N`, the ceiling is attained, and by the fixed-point identity `h_inf -> 0`: a perfectly exploited geometry at equilibrium has no headroom at all. Strict concavity means the approach is front-loaded: at `N = 8`, moving `u` from 0.00 to 0.25 buys 1.0116 nats and moving it from 0.75 to 1.00 buys 0.2468.

`kappa -> 0`. The crossover and the optimum both diverge, so a costless link removes every size limit this paper derives. **That is the signature that the size results are entirely a statement about `kappa` and not about people.** **`kappa -> infinity`.** `n_opt -> 1` and the net value at every `n >= 2` is negative: there exist parameter regimes in which no group outperforms nobody working together, and at `eta = 0.10`, `kappa = 0.10` the model already returns `net(2) = -0.031` against `net(1) = 0`.

`rho -> 1` in the queueing term. `rho/(1-rho) -> infinity`: the delay contribution to `C` is unbounded, so a critical resource at full utilisation violates the sustainability ceiling by itself. The divergence is real in the M/M/1 model and is the reason Section 5.8 sets a target below 0.80.

`gamma ->` such that `delta + gamma -> 1`. The effective time constant goes to zero and the system stops being a stock. **`delta -> 0`** removes the fixed point entirely and is outside the model's domain, which is the result the retention analysis establishes [Curzi, 2026a] and which this paper inherits rather than re-deriving.

Stated domain. The results of this paper are valid for `delta` in `(0, 0.5]`, `eta > 0` and bounded above by the saturation point of whatever capture mechanism is in use, `N >= 2` and bounded above by the crossover, `u` in `[0,1]`, `kappa > 0`, `gamma >= 0` with `delta + gamma < 1`, and `C_t <= (1−delta)Q_t + eta S_t` each period. **Outside that domain the equations still have solutions and they should not be interpreted as statements about a stock.**

8.11.6 11.6 Invariance

Unit rescaling. Let `Q' = cQ` for positive `c`. The law is invariant provided `eta`, `C`, `kappa` and `A` rescale with it — `eta' = c eta`, `C' = cC`, `kappa' = c kappa`, `A' = cA` — while `delta`, `gamma`, `N` and `u` are unchanged, since `delta` and `gamma` are dimensionless fractions and `S` is a dimensionless count. **`delta`, `gamma`, `u` and `N` are therefore the unit-free quantities of the framework, and `Q`, `eta`, `C`, `kappa` and `A` are the unit-carrying ones.** Two systems can be compared on the first group directly; they cannot be compared on the second unless they share a proxy. **This is the invariance that licenses the `n_opt = sqrt(eta/kappa)` rule, since the ratio `eta/kappa` is invariant under the rescaling** — it carries dimension `[link]`, the reciprocal of a per-link cost against a per-nat conversion, and `n_opt` is therefore a pure number as it must be.

Rewiring invariance. Under any change of communication topology holding the roster fixed, `N`, `u` and `S` are unchanged and only `C` moves, so `dh = −dC` exactly. **This is what makes topology a free lever and it is the invariance H5 tests.** What would break it: a rewiring that removes categories from the room, which changes `N`; or one that suppresses disagreement, which changes `u` — F15, the hub that eats the variance, is precisely the failure of this invariance.

What would break the unit invariance: a `delta` that varies with `Q`, which makes the transformation non-linear; an `eta` that saturates at large `S`, which the linear model excludes by assumption; or a proxy whose counting unit is not extensive, so that doubling the system does not double the count. **The third is the most likely in practice and the least often checked.**

How many scales this has been exercised at. Within this paper: three, all constructed. Across the corpus the same five-parameter survey has been applied from a single desk to a twenty-two-household block [Curzi, 2026k]. **That is a handful of systems, one corpus, one author, and no independent replication. The invariance is a claim about form that has been exercised, not one that has been established.**

8.11.7 11.7 Model comparison, and the exponent that was never estimated

This paper fits nothing, so no information criterion applies to it. What it does do is assert a functional form for cost, and the rigour standard requires that assertion to be handled properly rather than by assumption.

The exponent 2 in `C = kappa n(n−1)/2` is defined by the model, not estimated from data. That is the honest status and it is stated wherever the form is used. The pair count is exact combinatorics; the identification of coordination burden with the pair count is P5, a badged modelling choice.

The procedure the source manuscript proposes for checking it is not sound. It prescribes regressing `log C` on `log n` across comparable units and reading the slope, with a slope near 2 taken to support the complete-graph postulate. Clauset et al. [2009] shows that fitting a scaling exponent by ordinary least squares on log-transformed data is unreliable, that a straight line on a log-log plot is not evidence, and that the defensible procedure is maximum likelihood with a goodness-of-fit test and an explicit likelihood-ratio comparison against named alternatives. Stumpf and Porter [2012] adds that even a well-fitted exponent is weak evidence for a mechanism, because many mechanisms produce the same curve. **The defensible version of the check is therefore: state the two candidate models — `C = kappa n(n−1)/2` and `C = kappa_eff n` — with their free parameter counts, fit both by maximum likelihood on training data, and compare on held-out error with AIC or BIC reported, so that the parameter count is paid for.**

No such comparison exists in this corpus. **The correction is recorded here rather than made silently, and it is listed in Section 12 as an unmet check.**

8.11.8 11.8 Analogy discipline, scale type, and the precision of the numbers

The circuit and stock–flow correspondences are shared mathematical form, not derivations from physical premises. The storage law and a first-order discrete section are governed by the same difference equation, so every theorem about first-order systems applies to both; nothing about charge or capacitance is claimed to hold of an institution. The level–rate vocabulary of system dynamics is the same statement in a different notation [Forrester, 1961].

The queueing result is a derivation, not an analogy, wherever its premises hold. `rho/(1-rho)` follows from the M/M/1 model — Poisson arrivals, exponential service, one server, infinite queue — and where an approval process does not satisfy those premises the formula does not transfer. **The premises are almost never checked in organisational settings and are not checked here**; Hopp and Spearman [2011] gives the variability-adjusted forms that should be used instead when service times are not exponential. **This paper uses the M/M/1 form as an order-of-magnitude illustration and labels it as one.**

The entropy invoked anywhere near `ln N` is Shannon entropy — a dimensionless information measure in nats — **and not thermodynamic entropy**, which carries `[energy][temperature]^-1`. They are different objects and the framework identifies surplus with neither. Mirowski [1989] is the standing warning: neoclassical economics was built by copying mid-nineteenth-century energetics and the copy was defective, because utility was made to play the role of energy without the conservation law that gave energy its content. **`Q` in this paper is a stock that accumulates, leaks and is spent — the grammar of a conserved quantity — and nothing in this framework is conserved.** Cartwright [1999] offers the escape route this paper takes: specify the arrangement within which the law is claimed to hold, which is the stated domain of Section 11.5, and stop claiming a universality the law does not need.

Scale type. `Q`, `C`, `kappa` and `h` are treated as ratio-scale quantities throughout — differences, ratios and logarithms are all taken — and `u` likewise. **The conjoint-measurement conditions that would license that treatment have not been tested** [Luce and Tukey, 1964, Krantz et al., 1971]. Michell [1997] argues that psychology systematically fails to test whether its attributes are quantitative at all before measuring them, and **that is the strongest single objection to this framework's founding move.** The honest answer is that the ratio-scale treatment is a stipulation, that the remedy is a conjoint test or a Rasch-model demonstration [Rasch, 1960, Andrich, 1988], and that neither has been run. Note also that `u` estimated from a disagreement audit is built from ordinal judgements aggregated into a proportion; **the aggregation step is not licensed by the scale of its inputs and is performed anyway, which is stated here rather than defended.**

Precision. The numbers in this paper are of three kinds. *Defined or derived-from-defined* quantities — `ln 8 = 2.079442`, `ln 6 = 1.791759`, `1/(sqrt(8)+1) = 0.261204`, every entry in the doubling and infeasibility tables — are exact to the precision printed and may be recomputed further. *Illustrative model output* — every trajectory table, `h_0 = -0.1062`, `Q_inf = -1.6069`, `Q_crit = 5.9708` — is the output of stipulated parameters and carries **no measurement uncertainty because there was no measurement**; four or six decimal places there record the arithmetic, not confidence. *Measured or fitted* quantities do not appear in this paper, because nothing in it was measured. **Any application of this framework to real data must report fitted parameters with intervals**, and the four-decimal presentation used throughout must not be carried over to a measured `eta`, where two significant figures with an interval would be the defensible form.

8.12 12 Limitations

The three worked systems are constructed, not observed. No dataset was collected for this paper. Every number in Section 9 is the output of the stated parameters under the stated model, carried through so that a reader can re-derive it. They demonstrate that the procedure runs and that it reports its own failures; they are not evidence that any organisation behaves this way and should never be cited as such.

The cost model is the paper's load-bearing weakness. P5 — that coordination cost is proportional to the active path count — carries the crossover, the optimum, the topology lever and every sizing result. Its support is an observation from practice [Brooks, 1975] presented by its author as observation, not proof. **Nothing in this paper measures the exponent, and Section 11.7 records that the procedure the corpus proposes for measuring it is not sound.** If coordination cost is in fact linear in headcount under managed communication — which Conway [1968] arguably implies for an organisation that has already solved its topology — then Sections 5.4 to 5.6 describe a regime rather than a law.

The three-gate decomposition has no external support at all. It appears in no literature; it is a diagnostic convenience; and although no result in the verification chain consumes it, it carries the entire practical argument of Section 4 and the whole of the nine-person case's diagnosis. **A reviewer is entitled to treat it as an unsupported taxonomy, and this paper's answer is that it is one.**

The liaison multiplier is the least well-estimated number in the paper. Section 9.2's conclusion reverses between a four-fold and an eight-fold cross-boundary cost — `h = +0.0109` against `h = −0.0611` — and nothing here determines which is right. **A result that flips sign on an unmeasured parameter is a result about the parameter.**

The cellular recommendation prices a saving and not a fragility. Cells joined by thin liaison links are a coupled system, and coupled networks are more fragile than either network alone [Buldyrev et al., 2010]; interactive complexity with tight coupling makes catastrophic failure a system property [Perrow, 1984]; and redundancy can reduce safety [Sagan, 1993]. **This paper computes the coordination saving of a cellular split and does not compute its failure cost.** The four-way convergence of Section 5.7 is convergence on a topology, not on a risk assessment.

Which parameters were chosen rather than measured. In all three worked systems, `(N, u, eta, delta, C, kappa, gamma, A)` are stipulated. The capture-regime table of Section 4.9 is explicitly a set of illustrative inputs spanning a plausible range and must not be cited as a finding. The five-topology price table of Section 5.7 assumes a four-fold liaison multiplier that is a design estimate. **None of these is a measurement and none carries an interval.**

`N_eff = 1 + (N−1)u` is not a Hill number, and no external literature contains it. `1/HHI` is the Hill number of order two and `exp(H)` is the Hill number of order one, and both are fully grounded [Hill, 1973, Jost, 2006]; `1 + (N−1)u` is a linear interpolation between 1 and `N` and is this framework's own construction, adopted by design axiom S3. **A reviewer who reaches for the Hill-number literature as support will find that it does not apply.** The two quantities are used under one effective-count heading across this corpus, which is a real exposure; the repair is to derive the properties of `1 + (N−1)u` in the framework's own terms and to reserve the Hill citations for `1/HHI` and `exp(H)`, where they are exact [Curzi, 2026c,?].

`u` has not been shown to be a quantity. The framework treats non-parallelism as a measurable magnitude, and Michell [1997] argues that psychology routinely fails to test whether its attributes are quantitative at all. The remedies are conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971] or a Rasch demonstration [Rasch, 1960]; neither has been run, so the quantitative treatment is a stipulation. Relatedly, if `u` is read as a similarity measure over people it cannot be assumed to behave metrically, since psychological similarity is asymmetric and violates the triangle inequality [Tversky, 1977].

`u` is a variety measure and the word it travels under covers three incompatible constructs. Harrison and Klein [2007] distinguishes separation, variety and disparity, and a quantity like `u` owes the reader a statement of which one it is. Section 10.2 supplies that statement; **it has not been supplied consistently across this corpus and that inconsistency is a real exposure.**

The reflection problem, and why it outranks every fitting result. Section 3.3 states it in full: applied to a group whose members both influence and constitute the aggregate, the specification is reflexive, and Manski [1993] proves the endogenous social effect is not identified from observational data. **No fit — however well designed, however honestly held out — recovers that parameter**, and Shalizi and Thomas [2011] show the same confounding for network data. This paper adopts no remedy, so **the applied reading of `eta` is a modelling position rather than an identified empirical result.** The most realistic repair is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001].

The Lucas critique applied to this framework's own parameters. Section 4.13 asserts that a measure adopted as a target ceases to measure what it names, and prescribes an adversarial test in response. The framework then treats `eta`, `delta`, `N` and `kappa` as constants over an estimation window. **Both cannot be true without qualification** [Lucas, 1976]. The position taken here is the narrow one: the parameters are treated as stable *only over a declared window and only while `Q` and `h` are not themselves reported targets.* **Where `h` is reported upward, the framework's own reflexivity applies to it, and the gauge is at least as gameable as anything it audits.** Hacking [1995] states the general form — classifications of people interact with the people classified, who change in response — and Ferraro et al. [2005] states the specific hazard, that theories assuming the worst of people can make that assumption true. **A framework that tells an organisation its stored order decays should ask what that assumption will make true.**

The functional form is underdetermined. Latané [1981]'s sub-linear power law and Briscoe et al. [2006]'s `n log n` fit the same qualitative facts, and nothing assembled here discriminates between them. **The logarithm is chosen for its axiomatic characterisation, not established by measurement**, and Peters and Gell-Mann [2016] is the strongest available warrant precisely because it needs no psychological premise.

No scaling or heavy-tail claim is made from data, so the distribution-fitting procedure of Clauset et al. [2009] does not apply to any fitted result here — because there are none. **It does apply to the exponent check the corpus proposes, and Section 11.7 records that the proposed check is unsound and states what would replace it.** This is an unmet requirement, not a satisfied one.

The linear model is chosen for tractability. Real capture saturates; real cost has thresholds; real institutions have hysteresis and a state-dependent `eta`. The steady state flags its own breakdown when `eta F(u*) < c*`, which is a useful property, but a linear model reporting a negative equilibrium has told the analyst only that it has left its domain, not what replaces it.

Constant parameters over the estimation window is an assumption organisations violate. Turnover, reorganisation and a change of leadership all move `delta`, `eta`, `N` and `kappa`. **A fit spanning a regime change is fitting an average of two systems, neither of which exists.**

Ossification is not explained. Section 6.7 shows that the constant-parameter linear model never reaches negative headroom and that a linear `Q`-scaling cost does not change this. The paper therefore *requires* an imported mechanism and offers three candidates, of which only exogenous accretion is worked. **The other two are named and not modelled**, and no evidence is offered that accretion is the common case; the third worked system is constructed to exhibit it, which demonstrates the diagnosis and proves nothing about frequency.

Levers are priced and their sequence is not. Section 6.8 gives four marginal derivatives and compares single moves. The corpus's own analysis of lever ordering finds that eight periods of the wrong order cost twelve periods of permanent lag, a minimum stock of `-2.373029` from `+1.000000`, and 111.133601 stock-periods of cumulative gap [Curzi, 2026]. **A table of marginal returns does not capture that, and the four-lever framing of this paper is a static comparison presented where a dynamic one is needed.**

Unmet rigour checks, printed rather than left to a reviewer. (i) *Model comparison.* This paper fits nothing, so no information criterion applies; but it asserts a cost exponent that was never estimated, and the corpus's one applied fit reports held-out RMSE against baselines and **no AIC or BIC**, so the parameter count was never paid for [Curzi, 2026i]. (ii) *Effect sizes and intervals.* Nothing in Section 9 carries an interval, because nothing in Section 9 was measured.

(iii) *Scale type.* `Q`, `C` and `u` are treated as ratio-scale and the conjoint conditions were not tested. (iv) *Independent replication.* Every system here and in the corpus's applied papers was constructed or surveyed by one author, and the reproducibility literature is unambiguous about what that is worth [Open Science Collaboration, 2015, Ioannidis, 2005]. (v) *Directional versus point prediction.* Meehl [1967] observes that directional predictions are nearly costless in the social sciences; this framework does make point predictions with constants, which is the right answer to that objection, but only Curzi [2026i] has so far put one at risk, and it did not meet its own margin. (vi) *The discriminating observation.* Durlauf [2005] and Sobkowicz [2020] ask what observation would distinguish these laws from a neighbouring model that fits equally well. **This paper does not have an answer**, and the honest statement is that the sizing results of Section 5 would be shared by any model pairing a decelerating benefit with a superlinear cost.

Nothing in this paper is a human-subjects finding. No interventions were run and no participants were recruited. The constraint that no intervention may remove a participant's ability to decline, invoked in Sections 6.9, 7 and 8, is a design constraint on prospective work and not a report of past practice.

What would falsify the account. H1 fails on a fitted `b` that moves between sub-windows with no recorded intervention. H2 fails if measured conversion is reliably closer to the mean of the gate values than to their product. H3 fails on a survey series in which geometry interventions dominate on computed headroom gain. H4 fails on a measured cost exponent near one under complete-graph obligation. H5 fails if a cellular split does not at least halve measured cost, or if measured `u` falls after it.

8.13 13 Conclusion

The three remaining terms of the storage law are read together by one gauge, `h_t = eta ln N − delta Q_t − C_t`, with three regimes — negative unsustainable, zero maintenance only, positive growth available — a maximum one-step increment `r_max = h_t`, and two distinct constraints that are routinely confused: admissibility `C_t <= (1−delta)Q_t + eta S_t` and sustainability `C_t <= eta ln N − delta Q_t`. **`eta` must carry `[Q]` per nat per period for the gauge to balance; the source is silent on the per-period half and this paper supplies the reading explicitly rather than assigning it quietly.**

Doubling the conversion coefficient is exactly squaring the category count, since `2 eta ln N = eta ln(N^2)`; and **doubling the surplus by raising non-parallelism alone is feasible only when `u <= 1/(sqrt(N)+1)`**, so at `N = 8` and `u = 0.30` the required `u'` is 1.2300 and the move is arithmetically impossible rather than merely expensive. **Conversion composes multiplicatively through recording, binding and reach gates**, so three gates at 0.7 give `eta = 0.343` and the weakest gate dominates: lifting recording from 0.20 to 0.80 gains 300 per cent against 50 per cent for lifting binding from 0.60 to 0.90. **`eta = 0` produces decline rather than stagnation**: the stock is non-increasing in every period and falls at exactly the cost rate.

The net optimum group size maximising `eta ln n − kappa n(n−1)/2` satisfies `eta = kappa n(n − 1/2)`, approximately `sqrt(eta/kappa)`, and lies far below the crossover at which cost exhausts the budget: at `eta = 0.30` and `kappa = 0.010` the optimum is six and the crossover is thirteen. Topology is a free lever: eighteen people rewired from a complete graph into three cells of six move from `h = −0.4571` to `h = +0.1189`, a swing of 0.576, with cost falling 63 per cent and the surplus ceiling untouched because `N = 18` either way. **When cost scales with stock as `C = C_0 + gamma Q`, the system behaves exactly as if `delta_eff = delta + gamma`**, so every result of the retention analysis transfers with that substitution. **At the fixed point, `h_inf = eta(ln N − F(u*))`: `delta`, `C` and `Q` cancel, so the only headroom a converged system has is the surplus it fails to extract from its own difference — and the plain linear model therefore never ossifies by itself.**

The nine-person worked case returns `N = 6`, `u = 0.55`, `S = 1.3218` nats — 74 per cent of `ln 6 = 1.7918` — with a measured conversion coefficient of 0.030 giving `h_0 = −0.1062` and a negative

steady state `Q_inf = -1.607`. The binding threshold is `eta >= 0.091`, a factor of three. These are the statements downstream papers should quote rather than paraphrase.

The paper's own counter-result is stated as plainly as its results. In that same worked case the levers this corpus is named for are inert: non-parallelism at 0.55 is already strong and adding a distinct lineage buys 0.0046, while cutting coordination cost to zero leaves headroom at roughly a tenth of what the conversion fix delivers. Improved cadence is refused as a conversion gain and booked under surplus instead. The fourth lever — deliberately shedding stored order — outperforms every other move at the paper's base parameters and is nonetheless declared out of bounds wherever what would be shed constitutes a person's ability to decline, **with the model giving no warning at that boundary, because the computed return there is genuinely positive.**

All of this is arithmetic conditional on one postulate with no controlled measurement behind it — that coordination cost tracks the active path count. **Whether any institution is described by these equations is a separate question, it is a postulate rather than a result, and the one preregistered confirmatory test of the underlying law reported in this corpus did not meet its own declared margin.**

8.13.1 Declarations

Data and code availability

No external dataset was collected for this paper. All numerical results are computed from the stated parameters under the model developed in the text, and every table can be re-derived from the equations and parameter values printed beside it. The three systems worked in Section 9 are constructed examples with declared parameters, not field observations. **The figures are generated by the scripts named in each caption**; those scripts contain no data beyond the numbers printed in this paper and the model equations of Sections 3 to 6.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, produced the worked examples and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The three worked systems are constructed. The declination constraint invoked in Sections 6.9, 7 and 8 — that no intervention removing a participant's ability to decline is within this framework's scope — is a design constraint on prospective work and is not a report of any study performed.

Preregistration

This paper reports no preregistered test of its own; its results are derivations, and the falsification criteria attached to the three worked systems in Section 9 are illustrations of the form

such a registration would take rather than registrations of tests that were run. The preregistered protocol for testing the storage law against a real system, together with the confirmatory application that **failed** its second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared margin of 10 per cent — is reported in Curzi [2026i]. Any applied claim resting on this paper inherits that result.

Chapter 9

The Effective Count and the Logarithmic Ceiling: Auditing Roster Size Against Non-Parallelism

Abstract. This paper develops the audited count that the framework substitutes for headcount and prices the three ways of raising it. The affine effective count $N_{\text{eff}} = 1 + (N-1)u$, with $u = 1 - (x \cdot y)^2$ and surplus $S = \ln N_{\text{eff}}$, is stated with its provenance: three of the four characterising axioms are structural, and the affine form is badged throughout as a design axiom, chosen and not derived. **It is stated equally plainly that $1+(N-1)u$ is not a Hill number of any order and has no external support, where HHI^{-1} and $\exp(H)$ are the Hill numbers of orders two and one.** The same verification chain is run on both and the second fails visibly, returning 2.000000 on a share vector where every Hill number reports below 1.39. The consequence of the audit is severe — a hundred-person department at $u = 0.05$ has an effective count of 5.95 — and an exact bridge to a standard spectral measure is supplied, $N_{\text{eff}} = N + 1 - N/\text{PR}$ against the Rényi-2 participation ratio of the Gram matrix, re-verified here at machine epsilon. Because u depends only on the magnitude of correlation, perfect opposition scores $u = 0$ exactly as perfect agreement does, so a standing contrarian buys nothing. The sharp Lipschitz constant $N - 1$ is read as a measurement-error budget: at $N = 12$ an uncertainty of ± 0.05 around $u = 0.05$ admits a surplus anywhere in $[0, 0.7419]$. A worked twelve-person panel audits to $\bar{u} = 0.5137$, $N_{\text{eff}} = 6.6509$ and 76.3% of $\ln 12$, and every doubling of N adds exactly $\ln 2 = 0.6931$ at any scale. **A worked research institute has headroom -2.1182 that no attainable geometry closes: perfect orthogonality across all eight of its domains would still leave it declining, and the paper states that rather than choosing a different example.**

Keywords: effective count; non-parallelism; design axiom; participation ratio; Lipschitz bound; logarithmic ceiling; redundancy deficit; lineage census

Series: B. Engineering the Surplus **Paper:** CCP-009

9.1 1 Introduction

A body's headcount is not a measure of what it can produce. This paper develops the audited count that the framework substitutes for headcount, states exactly what that count is and is not, and prices the three ways of raising it.

The storage law of the framework drives stored order with an interaction surplus $S_t = F(U_t)$, and the surplus functional is $F(u) = \ln(1 + (N-1)u)$ with $u = 1 - (x \cdot y)^2$ [Curzi, 2026]. The exponential of that surplus is the object of this paper: $N_{\text{eff}} = 1+(N-1)u$, the **affine effective count**, so that $S = \ln N_{\text{eff}}$. Surplus is not the logarithm of a headcount. It is the logarithm of a headcount after the headcount has been audited for redundancy, and the audit is performed by a purely geometric quantity.

The first thing this paper must say about that count is what it is not. The corpus runs two quantities under the name *effective count*. One of them — the reciprocal Herfindahl index HHI^{-1} used elsewhere in the corpus [Curzi, 2026] — is the Hill number of order two [Hill, 1973, Jost, 2006], is a numbers-equivalent in the sense Adelman [1969] gave the term, and was independently reinvented in ecology [Simpson, 1949] and in political science [Laakso and Taagepera, 1979]. The other — $1 + (N-1)u$, the one this paper develops — **is not a Hill number of any order, is not a numbers-equivalent, and has no supporting entry in the ecological, economic, political-science or information-theoretic literature.** It is a linear interpolation between 1 and N and it follows from a design axiom the framework chose. Section 3.3 states the distinction, Section 9.6 runs the same verification chain on both and lets the second fail visibly, and every use of the affine count in this paper carries the badge.

Three of the four axioms that characterise f uniquely are structural bookkeeping. The fourth, **S3 — that the effective count is affine in** u — is a design decision. Everything in Sections 4 to 6 that reads as a discovery about institutions is, at bottom, a consequence of having decided that the effective count should interpolate linearly. The decision is defensible and it is not derived, and the paper badges it at every use.

The consequence is severe and is the paper's headline number. A hundred-person department at $u = 0.05$ has an effective count of 5.95. Not a hundred, not fifty, just under six. That figure is a statement about a geometry, not about the people in it, and the same arithmetic would apply to sensors or to estimators with nobody involved.

Seven further results follow. **(1)** An exact bridge to a standard spectral measure: $N_{\text{eff}} = N + 1 - N/\text{PR}$, where PR is the Rényi-2 participation ratio of the Gram matrix of the encoded states [Rényi, 1961], verified numerically at $K = 3, 5, 8, 12$ to better than 10^{-12} [Curzi, 2026i,?]. That identity does not repair the postulate underneath it, but it does name the audit: mean-pairwise u *is* the participation ratio, reparametrised. **(2)** Because u depends only on the magnitude of the correlation, perfect opposition at $\rho = -1$ scores $u = 0$ exactly as perfect agreement does; **a standing contrarian buys no surplus at all**, and a body assembled for creative tension has purchased an effective count of one at the price of two seats [Curzi, 2026]. **(3)** The sharp Lipschitz constant $N-1$ [Curzi, 2026] is read here as a measurement-error budget rather than a smoothness result: at $N = 12$, an uncertainty of ± 0.05 around $u = 0.05$ admits a surplus anywhere in $[0, 0.7419]$, a width of 0.7419, against 0.1266 at $u = 0.70$. **In the low-u regime where institutions actually sit, a plausible measurement error spans nearly the whole range of the quantity being estimated.** **(4)** The affine count is explicitly distinguished from the classical variance-reduction count $N_{\text{ind}} = N/(1 + (N-1)\rho)$, which tends to a hard asymptote of $1/\rho$; the two share a form and are different quantities, and Markowitz [1952] had the substantive point — that diversification benefit depends on covariance and not on count — in 1952. **(5)** Every doubling of N adds exactly $\ln 2 = 0.6931$ nats at any scale, so the ceiling gain from one category to two equals the gain from fifty to a hundred. **(6)** A twelve-person panel audits to $u = 0.5137$, $N_{\text{eff}} = 6.6509$ and $S = 1.8948$ nats, which is 76.3 per cent of $\ln 12$, with a redundancy deficit of 0.5901. **(7)** Growth of that panel from twelve to thirty buys 0.8714 nats for eighteen hires and 435 communication paths against 66, while redistributing the existing members across twelve lineages buys 0.2654 nats for none.

The paper's own worked institute refuses the lesson the framework is named for. A publicly funded research institute with eight declared domains has headroom $h = -2.1182$ against a total structural budget of $\eta \ln N = 0.8318$. Doubling its domain count buys 0.2772 nats and closes 13 per cent of the gap; break-even by domain addition alone requires $N = 1596$; and **even perfect orthogonality across all eight domains would leave it declining by** 1.5682 **nats per year.** The geometry this corpus is named for was never the binding constraint in the paper's

own example, and no geometric intervention could have worked. That is stated in Section 7.2 rather than resolved by choosing a different example.

The paper is a framework-development paper with three worked applications, not a field study. **No dataset was collected. Every system in Section 7 is constructed**, and every number in it is an exact computation on stipulated inputs, so that a reader can redo the arithmetic without access to anything. Section 10 records what that costs the paper's claims.

Sections 2 and 3 state the hypotheses and the framework, including the two objections — the identification problem and the underdetermination of the functional form — that no amount of fit can answer. Section 4 develops the audited count, Section 5 the ceiling, and Section 6 the engineering of non-parallelism; each closes with its design specifications and its named failure modes. Section 7 works three systems end to end. Sections 8 to 11 discuss, verify, limit and conclude.

9.2 2 Hypotheses

The results below are of two kinds and the distinction is load-bearing. Some are theorems about a functional form and cannot be falsified by observing an institution; they are labelled as derivations or definitions where they appear, and they carry no empirical weight. Others are claims that the functional form describes real bodies, and those are stated here with what would defeat them. Where a claim can be neither derived nor falsified it is labelled a postulate or a modelling choice and is not dressed as a hypothesis [Popper, 1959, Platt, 1964].

H1 — The audited count, not the headcount, governs what a body delivers. A body's aggregation performance tracks $N_{\text{eff}} = 1 + (N-1)u$ rather than N. **Predicted:** across bodies matched on N and differing in measured u, scored accuracy rises with u and is flat in N at fixed u. **Falsified by:** matched bodies whose accuracy tracks headcount at fixed measured u. **Not tested here.** The supporting external result is Hong and Page [2004], whose conditions Thompson [2014] disputes in detail; the disagreement is reported in Section 6.5 rather than adjudicated.

H2 — Opposition is redundancy. Because $u = 1 - \rho^2$ depends only on $|\rho|$, a member who reliably inverts the body's judgement adds no surplus. **Predicted:** a designated-contrarian arrangement, measured by error decorrelation on scored outcomes, returns $u \approx 0$ and no improvement in aggregate accuracy over the aligned body it replaced, while costing one additional seat. **Falsified by:** a measured accuracy gain from a designated opponent at $\rho \to -1$, which would show that the encoder — error vectors — is the wrong state representation and that sign carries information the square discards. **Not tested here.** This is the paper's sharpest and most checkable empirical prediction.

H3 — Conversion and cost bind before geometry does. In bodies with negative headroom, the intervention with the largest computed Δh is more often a conversion or a cost intervention than a geometry intervention. **Predicted:** across surveyed systems with $h < 0$, geometric interventions are dominated on computed Δh. **Falsified by:** a survey series in which geometry interventions dominate. **This paper's second worked system is consistent with H3 in the strongest possible way** — no attainable geometry closes its gap — and its third is a counter-instance, since there u was the binding constraint. Both are reported in Section 7.

H4 — The affine effective count equals the participation-ratio count. For N encoded unit vectors with Gram matrix G and $\text{PR} = (\operatorname{tr} G)^2 / \|G\|_F^2$, mean pairwise nonparallelism $\bar{u}$ satisfies $1 + (N-1)\bar{u} = N + 1 - N/\text{PR}$. **Predicted:** exact agreement in double precision on any ensemble. **Falsified by:** any ensemble on which the two disagree beyond floating-point tolerance. *(Result:* ***not falsified*** *— agreement*

better than 10^{-12} *at* $K = 3, 5, 8, 12$ *[Curzi, 2026i], and re-verified over six random correlation matrices [Curzi, 2026].* ***This is a derivation; the numerical check tests the implementation, not the world.)***

Three claims in this paper are not hypotheses and are labelled as what they are.

S3, the affine effective count, is a DESIGN AXIOM. That $g(u) = e^{f(u)}$ is affine in u is chosen and not derived. No observation refutes it, because it is a decision about what the word *count* should mean in this framework. It is badged at every use, and Section 9.6 shows precisely what it costs: the resulting quantity is not a numbers-equivalent, and the diversity literature does not cover it.

The Roster Reduction is a POSTULATE. Reading a roster of N people as N categories, and a mean pairwise u as the roster's non-parallelism, is a modelling assumption stated in Section 3.4 together with the specific way it can be wrong.

That error vectors are the right state encoding is a POSTULATE. The identity $u = 1 - \rho^2$ is exact for mean-centred, unit-normalised vectors, because their inner product is the Pearson correlation. That a member's scored error vector is the state the framework means is an encoding choice, badged wherever it is used.

9.3 3 Framework: the law, what is chosen, and what the count is not

9.3.1 3.1 The statement

From Curzi [2026], Thm 2.1, the interaction surplus of a pair of states x, y is

```
F(x,y) = f(u(x,y))

u(x,y) = 1 − (x . y)^2 = || x ^ y ||^2      in [0,1]

f(u)   = ln(1 + (N−1) u)
```

and the object of this paper is the exponential of the surplus:

```
N_eff = g(u) = e^{f(u)} = 1 + (N−1) u        the EFFECTIVE COUNT

surplus = ln N_eff
```

Read the second line slowly. Surplus is not the log of a headcount. **Surplus is the log of a headcount after the headcount has been audited for redundancy.** The audit is performed by u, and u is a purely geometric quantity: it is the squared magnitude of the wedge of the two states, or equivalently one minus the squared cosine between them.

Because `(x . y) = cos θ` for unit states, one may write the identity that does most of the work in Sections 4 to 6:

```
u = 1 − cos^2 θ = sin^2 θ
```

u is the squared sine of the angle between two states. It is 0 when they are parallel and 1 when they are orthogonal, and it moves through the intermediate range in the way a squared sine does — slowly at first, then fast, then slowly again.

9.3.2 3.2 What is proven and what is chosen

The functional form `f(u) = ln(1 + (N−1)u)` is **uniquely characterized** by four axioms. Three of them are bookkeeping. One of them is a design decision, and the framework is required to say so every time it leans on it.

Axiom	Statement	Status
S1	Geometric dependence: F depends on the pair only through u.	Structural. Fixes what kind of object F is.
S2	Zero at zero: `f(0) = 0`. Parallel states produce no surplus.	Normalization. Uncontroversial.
S3	Affine effective count: `g(u) = e^{f(u)}` is affine in u.	**[DESIGN AXIOM — CHOSEN, NOT DERIVED]**
S4	Normalization: `f(1) = ln N`.	Normalization. Sets the ceiling.

S3 is the design axiom. It is chosen, not derived, and it is the reason the effective count is a straight line in u rather than some other curve. Everything in Sections 4 to 6 that reads as a discovery about the world is, at bottom, a consequence of having decided that the effective count should be affine. The decision is a good one — it is the simplest possible assumption that makes "effective count" behave like a count, interpolating linearly between one entity and N entities — but it is a decision. When a reader objects that the second member cannot possibly be worth exactly as much as the fortieth in u-units, the honest answer is: S3 says it is, S3 was chosen, and if a measured reason exists to replace S3 with something else, the rest of the theory will have to be rebuilt around it.

Given S1–S4, the following are **derived**, not assumed [Curzi, 2026]:

Property	Reference	Statement
Strict monotonicity	Thm 3.1	More non-parallelism strictly increases surplus.
Positivity	Thm 3.2	`F > 0` for any `u > 0`.
Strict concavity	Thm 3.3	Diminishing returns in u, unconditionally.
Subadditivity in u	Thm 3.4	`f(u_1 + u_2) <= f(u_1) + f(u_2)` on the parameter.
Sharp Lipschitz constant	Thm 4.2	`|f(u_1) − f(u_2)| <= (N−1)|u_1 − u_2|`, and N-1 is sharp.
Extremality	Thm 4.3	`F = 0` iff parallel; `F = ln N` iff orthogonal.
Range	Curzi [2026], Cor 2.2	`0 <= F(u) <= ln N`.

9.3.3 3.3 Which effective count: the Hill correspondence, stated exactly

The corpus runs two quantities under the name *effective count*, and they are not the same quantity. This section states which is which, in the paper that develops the second of them, because a reviewer who reaches for the diversity literature will find that it covers one exactly and the other not at all.

The reciprocal concentration index is a Hill number of order two. Hill [1973] defined the family ${}^{q}D = \left(\sum_i p_i^q\right)^{1/(1-q)}$, whose members return the number of equally abundant types that would produce the observed value of the underlying index. At $q = 2$ this is $\left(\sum_i p_i^2\right)^{-1}$, which is HHI^{-1} exactly. Jost [2006] is the statement that entropies are not diversities and that their numbers-equivalents are, and Jost [2007] extends it to the partition into within-group and between-group components. **The numbers-equivalent reading — that the reciprocal index answers "how many equal units would give this figure" — is Adelman [1969], and the corpus should name him rather than hedge**; the source manuscript attributes the reading to no one. The same arithmetic was reached independently in ecology by Simpson [1949] and in political science by Laakso and Taagepera [1979] as the effective number of parties. Independent reinvention in three disciplines is the strongest available argument that the instrument is general rather than borrowed. The exponential of Shannon entropy [Shannon, 1948] is the same family at $q = 1$; the reciprocal index is also the exponential of the Rényi entropy of order two [Rényi, 1961], which is where the participation-ratio name of Section 3.4 comes from. The measure that makes diversity depend on *dissimilarity between types* rather than on shares alone — which is this framework's founding move — exists in exactly two formalisms, Rao [1982] and Leinster and Cobbold [2012], and neither of them is the affine count.

The affine effective count is not a Hill number and has no external support. The quantity this paper develops, $N_{\text{eff}} = 1 + (N-1)u$, is a **linear interpolation between** 1 **and** N. It is not $\left(\sum_i p_i^q\right)^{1/(1-q)}$ for any q; it is not a function of a share vector at all, but of the pair (N, u); it is not a numbers-equivalent; and no entry in ecology, economics, political science or information theory contains it. The arithmetic that separates the two takes one line and is given as a failing verification chain in Section 9.6: for two units holding 0.9 and 0.1, the affine count at $u = 1$ returns exactly 2.000000, while HHI^{-1} returns 1.219512 and $\exp(H)$ returns 1.384145. **Every Hill number of order one or higher reports at most** 1.384145 **on that vector, and the affine count reports two.** The comparison must be stated with that qualification, because the Hill family is non-increasing in q and its order-zero member — the raw richness ${}^{0}D$ — equals 2.000000 on this vector as it does on every two-type vector. **That coincidence is uninformative and is worth stating precisely because it is the sharpest form of the objection**: ${}^{0}D$ is the one member of the family that ignores shares entirely, which is exactly the naive count an effective count exists to replace, and the affine count at $u = 1$ agrees with it.

The two quantities are related — both are called effective counts, both return N in the fully independent limit, both return 1 at total redundancy — and they are not interchangeable. The corpus has in places let them travel under one name. **This paper uses the affine count throughout, badges it as the framework's own construction at every use, and reserves the Hill citations for the reciprocal index, where they are exact.** A downstream paper that cites this one must carry the distinction with the citation.

What is lost, stated plainly. Three properties a reviewer would expect of an effective count do not hold for the affine form. It does not satisfy the replication principle — pooling two identical communities of N equally weighted types should return N, and the affine count's answer depends on u rather than on the composition. It is not invariant to how the underlying units are partitioned, because N is a declared category count and not a measured share vector. And it inherits no estimation theory: the small-sample bias corrections of Chao et al. [2014] and Chao and Jost [2012] apply to the Hill family and not to this. **A twelve-person roster is a small sample, no correction is available for the affine count, and none was applied.** Hurlbert [1971] and Gorelick [2011] are the standing objections to treating any such index as a substantive property rather than as a statistic, and they apply here with more force than they do to the Hill family, not less.

9.3.4 3.4 The Roster Reduction — a badged postulate

The law is stated for a **pair** of states x, y, with N entering only through the normalization S4 as the number of distinct categories the ambient space can express. Practical work wants to

apply it to a roster of N people, or N teams, or N sensors. That step is not in the founding surplus analysis [Curzi, 2026]. It is a modeling assumption and this paper badges it:

> **THE ROSTER REDUCTION [POSTULATE].** For a body of N members, N is taken to be the number of nominal categories the body claims to span, and u is taken to be a representative non-parallelism across the body's pairs — normally the mean pairwise u over all `N(N−1)/2` pairs, or a lineage-weighted version of the same. `N_eff = 1 + (N−1)u` is then read as the audited count.

This postulate is doing real work and it can be wrong in a specific way: because the triangle inequality fails for geometric triples, a mean pairwise u does **not** determine the joint geometry of the whole body. Two rosters with the same mean pairwise u can have genuinely different structure — one a ring of mutual near-orthogonality, another a tight cluster plus one distant outlier. The Roster Reduction discards that distinction. It is used because it is cheap, computable from instruments that can actually be run (Section 4.2), and because it errs conservatively in the common institutional case, where the failure is a tight cluster and the mean pairwise u is dragged down by exactly the redundancy the audit is trying to detect. **Where the difference between "one distant outlier" and "genuine spread" matters to a decision, do not use the mean; report the pairwise matrix and read it directly.** the two-source decomposition [Curzi, 2026d] splits u into its two components and gives a sharper instrument.

The mean-pairwise choice is nevertheless less arbitrary than it looks, and the survey protocol [Curzi, 2026i] says why. Let `G` be the Gram matrix of the N encoded unit vectors, `G_ij = x_i . x_j`, and let `PR` be the standard Renyi-2 participation ratio of its spectrum, `PR = (tr G)^2 / ||G||_F^2`. Since `||G||_F^2 = N + N(N−1)(1−u)`, one line of algebra gives

```
N_eff = 1 + (N−1)u = N + 1 − N/PR
```

The mean-pairwise `u` aggregator IS the participation ratio, reparametrised into effective-count form. The identity is exact; the survey protocol [Curzi, 2026i] verifies it numerically at `K = 3, 5, 8, 12` on random ensembles to better than 1e-12. Its two poles are worth carrying. At `PR = N` the Gram matrix has no off-diagonal mass, `N_eff = N`, and roster equals effective count — attainable only at exact mutual orthogonality. At `PR = 1` every vector is the same, `N_eff = 1`, and the size of the roster has left the arithmetic entirely.

This does not repair the postulate; the triangle-inequality objection above still stands, because `PR` is itself a scalar summary of a spectrum. **What it does is name the audit.** The Roster Reduction is not an improvisation but a standard diversity measure an ecologist, a portfolio analyst or a spectroscopist would recognise on sight, and when a sponsor asks why disagreement is averaged pairwise rather than some other way, the answer is this identity and not a plausibility argument. See the survey protocol [Curzi, 2026i].

9.3.5 3.5 Identification: the objection that no measurement of u can answer

Every applied reading in this paper runs through a reflexive specification. The framework's central formal claim is that a group's influence on its members and the members' constitution of the group are one loop, $T(x) = W(x, M(x))$ [Curzi, 2026]. The estimate of u inherits it directly: the error vectors whose correlation defines u are produced by members who observe one another, and the correlation the analyst measures is jointly the product of shared training, shared incentive, and mutual influence.

That specification is exactly a linear-in-means reflexive model, and its endogenous social effect is not identified from observational data. Manski [1993] proves it: the endogenous effect cannot be separated from contextual effects or from correlated unobservables without prior information specifying the reference group. No quantity of scored decisions recovers it. Shalizi and Thomas [2011] reinforce the point for networks — influence and selection are generically confounded in observational data — and Lyons [2011] shows what happens when the confound is ignored in a well-publicised case. Durlauf [2005] adds the requirement that

a complexity model earn its keep by making a **differential** prediction that simpler models do not also make.

The remedy this paper adopts is the fourth of Manski's four, and it is not free. The available routes are randomisation, an instrument, partial-population variation, and stipulating the reference group and defending the stipulation; Brock and Durlauf [2001] is the template for the last, and it derives a mean-field structure from standard discrete choice **with its identification conditions stated**. This paper stipulates the reference group — the declared roster, with the declared lineage partition of Section 7.1 — and defends the stipulation on the ground that institutional membership is administratively observable in a way that a social peer group is not. **That is a defence of the stipulation, not a recovery of the parameter.** Accordingly:

> **Every applied u in this paper is a modelling position, not an identified empirical quantity.** The arithmetic relating u to N_{eff} and to S is exact and is not at issue. What is at issue is whether a measured error correlation is the framework's u, and observation alone cannot settle it.

Two objections in the same family are also unanswered here and are recorded rather than deflected. **The Lucas critique, turned on the framework itself** [Lucas, 1976]: this framework asserts that systems change their behaviour when modelled and observed, and then treats its own u, η and δ as stable in a population that has read it. A body that has been told its second opinion is worth $\ln 1 = 0$ will behave differently in the next elicitation, and Section 6.8's D-1 to D-10 are precisely the mechanisms through which that would operate. Hacking [1995] is the general statement — the subjects of a psychological physics read the theory and move — and Ferraro et al. [2005] is the warning that a description assuming the worst can make it true. **The framework has no answer and this paper states that it has none.** MacKenzie and Millo [2003] is the countervailing observation that such a loop has been caught happening, with dates and price data, which establishes that reflexive fixed points exist without establishing that this one does.

And the functional form is underdetermined by the evidence. Three shapes fit the qualitative facts equally well: this framework's $\ln(1 + (N-1)u)$, the sub-linear power law with exponent below one that Latané [1981] established with psychological data, and the $n \log n$ network-value law of Briscoe et al. [2006]. Nothing in this paper discriminates among them, because nothing in this paper fits the form to data at all; the paper derives consequences of a form assumed. Sobkowicz [2009] and Sobkowicz [2020] are the general indictment — many models, negligible empirical discrimination between them — and they apply. **The honest statement is that the logarithmic form is adopted and its uniqueness under S1–S4 is a theorem about the axioms and not evidence about institutions.**

One strengthening is available and should be taken. Peters and Gell-Mann [2016] and Peters [2019] derive a logarithmic form from multiplicative dynamics with **no psychological assumption whatever**: the time-average growth rate of a multiplicative process is logarithmic. That is a stronger warrant for the framework's logarithms than the psychophysical route through Fechner [1860], because it does not require the perceptual story to be true. It also carries a warning the framework must accept: the time average of a multiplicative process differs from its ensemble average, so a parameter estimated across a cross-section of bodies does not describe what happens to any one body over time — and Q accumulates over time, for one body.

9.3.6 3.6 Analogy discipline: what is derived, what shares a form, and what is only suggestive

The framework borrows shape from physics and information theory, and each borrowing is classified here at first use, in the three categories the corpus's rigour standard requires.

Borrowing	Category	What is and is not claimed
The leaky integrator $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$	**2 — shared mathematical form**	A first-order linear difference equation governs both a discharging capacitor and this stock [Oppenheim et al., 1996, Elaydi, 2005]. No physical claim is transferred [Curzi, 2026a].
The Rényi-2 participation ratio	**1 — derivation**	PR $= (\operatorname{tr} G)^2 / \|G\|_F^2$ is applied to a Gram matrix, which is what the quantity is defined on [Rényi, 1961]. The identity of Section 3.4 is algebra, not resemblance.
$\ln N$ as a ceiling, against Shannon entropy	**2 — shared form only**	Both are logarithms of a category count. **Surplus is not identified with Shannon entropy or with mutual information** [Curzi, 2026]; the resemblance is a normalisation over N, and Shannon [1948] is cited for the entropy and not for the surplus.
Ashby's requisite variety	**2 — formal correspondence**	Ashby [1956, 1958] measure variety as the logarithm of a state count and treat it as a cap on regulation. The two are not proved equal (Section 5.9).
Condorcet's jury theorem	**3 — interpretive correspondence**	The independence premise and the non-parallelism parameter say related things in two vocabularies. **No theorem connecting them is asserted** (Section 4.7).
The portfolio analogy	**2 — shared form**	Markowitz [1952] states the substantive result — benefit depends on covariance, not count — and this framework restates it with a different index. The correlation-collapse literature [Longin and Solnik, 2001, Ang and Chen, 2002] is the warning that the covariance itself moves under stress.

This paper is bound by the corpus's own strictest ruling on borrowed physics. The uncertainty node declines the inflated reading of its own borrowing: it states the framework's postulate as a postulate, gives the relation only in its correct form, and requires that a user be able to distinguish the theorem, the ordinary observer effect, and the framework's own conjecture in three separate sentences on demand [Curzi, 2026]. The same standard governs every row of the table above, and nothing in this paper claims more of a borrowing than that ruling permits.

The failure mode this table exists to prevent is a category-2 borrowing carrying a category-1 result's authority. The arithmetic of Sections 4 to 6 is category 1 within its own axioms and needs no help; the physical and information-theoretic vocabulary is category 2 and is used for shape only. Mirowski [1989] is the historical account of what happens when a social science imports a physical formalism without auditing the import, and Cartwright [1999] supplies the generous escape route: a law can hold inside a specified arrangement without being universal, provided the arrangement is specified. **The design specifications of Sections 4.9, 5.13 and 6.11 are this paper's attempt to specify the arrangement.**

9.3.7 3.7 The honesty clauses that bind every use of the law

These are binding wherever the law is used and are reproduced here rather than referenced, because this is the point at which a reader is most tempted to over-read the result.

1. **The founding surplus analysis [Curzi, 2026] is pure mathematics and makes no empirical or metaphysical claim.** It proves that a particular functional form follows from four axioms. It does not prove that human groups obey that form. The empirical content of this paper is a modeling proposal, tested by the protocols in Chapters 24 and 25, and falsifiable there.
2. **As deployed, the surplus functional is a lower bound** on the ensemble's information, not the information itself. The floor is logged, not the ceiling. An `N_eff` of 2.5 does not mean the group contains exactly 2.5 independent perspectives; it means the measured geometry guarantees at least that much, and there may be more that the state encoder failed to see.
3. **Surplus is not identified with Shannon entropy or mutual information in general** [Curzi, 2026]. The ln is not a Shannon ln. Do not import entropy intuitions wholesale.
4. **The triangle inequality FAILS for geometric triples** (the surplus functional's metric analysis [Curzi, 2026]). Subadditivity holds on the parameter u only. You may not chain surplus across three states as though it were a distance. This matters directly in Section 4.5 where auditing a roster larger than two is discussed.
5. **No thermodynamic negentropy law is proved** [Curzi, 2026a]. `F(u) + N_can(u) <= ln N` bounds a sum; it does not license any statement about physics, order, or consciousness.

9.4 4 The audited count in use: geometry, instruments, and concavity

9.4.1 4.1 Ten parallel windows are one window

Set u = 0. Then `N_eff = 1 + (N−1)·0 = 1`, and `surplus = ln 1 = 0`.

This is the whole audit in one line, and it is not a limiting case or an edge condition. It is the modal condition of institutional life. A wall of ten identical windows, all facing the same courtyard at the same angle, admits one view. A panel of ten reviewers trained in the same three departments, reading the same literature, rewarded by the same committee, produces one review. The arithmetic does not care that there are ten signatures on it. It prices the geometry.

The corresponding engineering instruction is blunt: **the first question about any body is not how many members it has but how many directions it points, and the second question is how that was measured.**

Table 4.1 gives the audited count directly. Read down a column to see what a nominal roster is actually worth at a given non-parallelism.

TABLE 4.1 — EFFECTIVE COUNT N_eff = 1 + (N−1)u

N	u=0.0	u=0.05	u=0.1	u=0.2	u=0.3	u=0.5	u=0.7	u=0.9	u=1.0
2	1.00	1.05	1.10	1.20	1.30	1.50	1.70	1.90	2.00
3	1.00	1.10	1.20	1.40	1.60	2.00	2.40	2.80	3.00
4	1.00	1.15	1.30	1.60	1.90	2.50	3.10	3.70	4.00
5	1.00	1.20	1.40	1.80	2.20	3.00	3.80	4.60	5.00
6	1.00	1.25	1.50	2.00	2.50	3.50	4.50	5.50	6.00
8	1.00	1.35	1.70	2.40	3.10	4.50	5.90	7.30	8.00
10	1.00	1.45	1.90	2.80	3.70	5.50	7.30	9.10	10.00
12	1.00	1.55	2.10	3.20	4.30	6.50	8.70	10.90	12.00
15	1.00	1.70	2.40	3.80	5.20	8.00	10.80	13.60	15.00
20	1.00	1.95	2.90	4.80	6.70	10.50	14.30	18.10	20.00
25	1.00	2.20	3.40	5.80	8.20	13.00	17.80	22.60	25.00
30	1.00	2.45	3.90	6.80	9.70	15.50	21.30	27.10	30.00

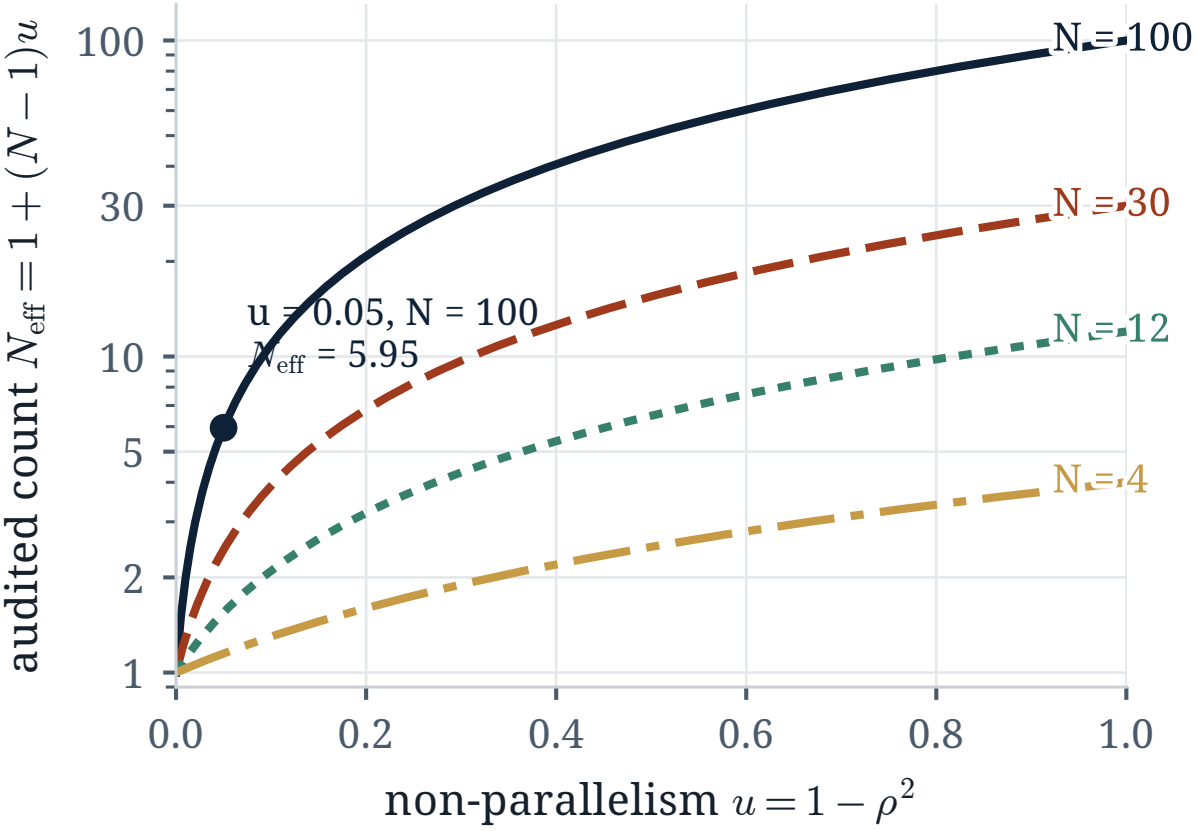


Figure 9.1: **A roster of a hundred at $u = 0.05$ audits to an effective count of 5.95, not to a hundred.** The affine count $N_{\text{eff}} = 1 + (N-1)u$ is plotted at four roster sizes on a logarithmic vertical axis; the marked point is the headline reading of Table 4.1. Model output, not measurement: the relation is exact given design axiom S3, which is chosen rather than derived. Generated by `figs/src/CCP-009_fig01.py`.

```
 40 |  1.00 |  2.95 |  4.90 |  8.80 | 12.70 | 20.50 | 28.30 | 36.10 | 40.00
 50 |  1.00 |  3.45 |  5.90 | 10.80 | 15.70 | 25.50 | 35.30 | 45.10 | 50.00
100 |  1.00 |  5.95 | 10.90 | 20.80 | 30.70 | 50.50 | 70.30 | 90.10 |100.00
```

A hundred-person department at u = 0.05 has an effective count of 5.95. Not 100. Not 50. Just under six.

If that department's hiring, training, and promotion pipeline is a single funnel, u = 0.05 is not a pessimistic estimate; it is a generous one.

9.4.2 4.2 The geometry of u, in degrees

Because `u = sin² θ`, the relationship between the intuitive quantity (how differently do these two see the problem?) and the parameter is not linear, and engineers consistently misjudge it in the same direction. Small angles are worth almost nothing.

```
TABLE 4.2 — ANGLE -> u -> AUDITED COUNT (illustrated at N = 12)

  theta    cos θ     u = sin^2 θ   N_eff      ln N_eff    % of ln 12
  -----    ------    -----------   -------    --------    ----------
    0°     1.0000      0.0000        1.000     0.0000        0.0%
    5°     0.9962      0.0076        1.084     0.0802        3.2%
```

10°	0.9848	0.0302	1.332	0.2864	11.5%
15°	0.9659	0.0670	1.737	0.5521	22.2%
20°	0.9397	0.1170	2.287	0.8271	33.3%
25°	0.9063	0.1786	2.965	1.0868	43.7%
30°	0.8660	0.2500	3.750	1.3218	53.2%
35°	0.8192	0.3290	4.619	1.5302	61.6%
40°	0.7660	0.4132	5.545	1.7129	68.9%
45°	0.7071	0.5000	6.500	1.8718	75.3%
50°	0.6428	0.5868	7.455	2.0089	80.8%
55°	0.5736	0.6710	8.381	2.1260	85.6%
60°	0.5000	0.7500	9.250	2.2246	89.5%
70°	0.3420	0.8830	10.713	2.3715	95.4%
80°	0.1736	0.9698	11.668	2.4569	98.9%
90°	0.0000	1.0000	12.000	2.4849	100.0%

Two observations an engineer should carry.

First, **a ten-degree difference of view is worth 3% of an available effective count of twelve.** A team of people who "have different backgrounds" but converge on the same framing within one meeting is a ten-degree team. Everyone experiences the difference as real — they argue, they use different words — and it prices at u = 0.03.

Second, the curve saturates hard past 60°. Going from 60° to 90° — from strongly different to completely orthogonal — buys the last 10.5% of the ceiling. **The expensive, socially painful distance between "strongly disagrees" and "does not share a frame of reference" is worth about a tenth of what the cheap distance between "identical" and "moderately different" is worth.** That is a merciful result and Section 6 builds on it heavily: there is no need to import a stranger from another civilization. The requirement is to stop hiring the same person eleven times.

Third, and it is the observation the table's 0°-to-90° range conceals: the curve comes back down. `u = 1 − (x·y)^2` depends only on the *magnitude* of the correlation, so it is symmetric about 90° and the table must be read as a fold, not a ramp:

```
120°  ->  u = 0.7500     same as  60°
135°  ->  u = 0.5000     same as  45°
170°  ->  u = 0.0302     same as  10°
180°  ->  u = 0.0000     same as   0°
```

At `rho = −1`, perfect opposition, `u = 0` exactly as it does at `rho = +1`. The member who reliably inverts a colleague spans no more of the space than the member who reliably echoes them: given either, the other can be reconstructed, and reconstruction is what redundancy means. **The contrarian carries exactly as much information as the sycophant: none.**

This has a direct consequence for a common staffing move. A body assembled for "creative tension" — a designated opposite, a house sceptic whose position is the negation of the proposal — has bought an `N_eff` of 1 on that axis and is paying `C` for two seats. It is worse than an aligned pair, not better, because it carries the same zero surplus at higher cost. Section 6 treats the engineering consequence and the pseudo-diversity analysis [Curzi, 2026d] catalogues it as a failure mode distinct from the Parallel Wall.

9.4.3 4.3 Five instruments for estimating u

u is a geometric quantity in a state space that is not directly observed. Estimating it requires an operational choice: the *state encoder*, which the founding surplus analysis requires be declared [Curzi, 2026] and the survey protocol operationalises [Curzi, 2026i]. Here are the five instruments this paper recommends, in increasing order of cost and decreasing order of ambiguity.

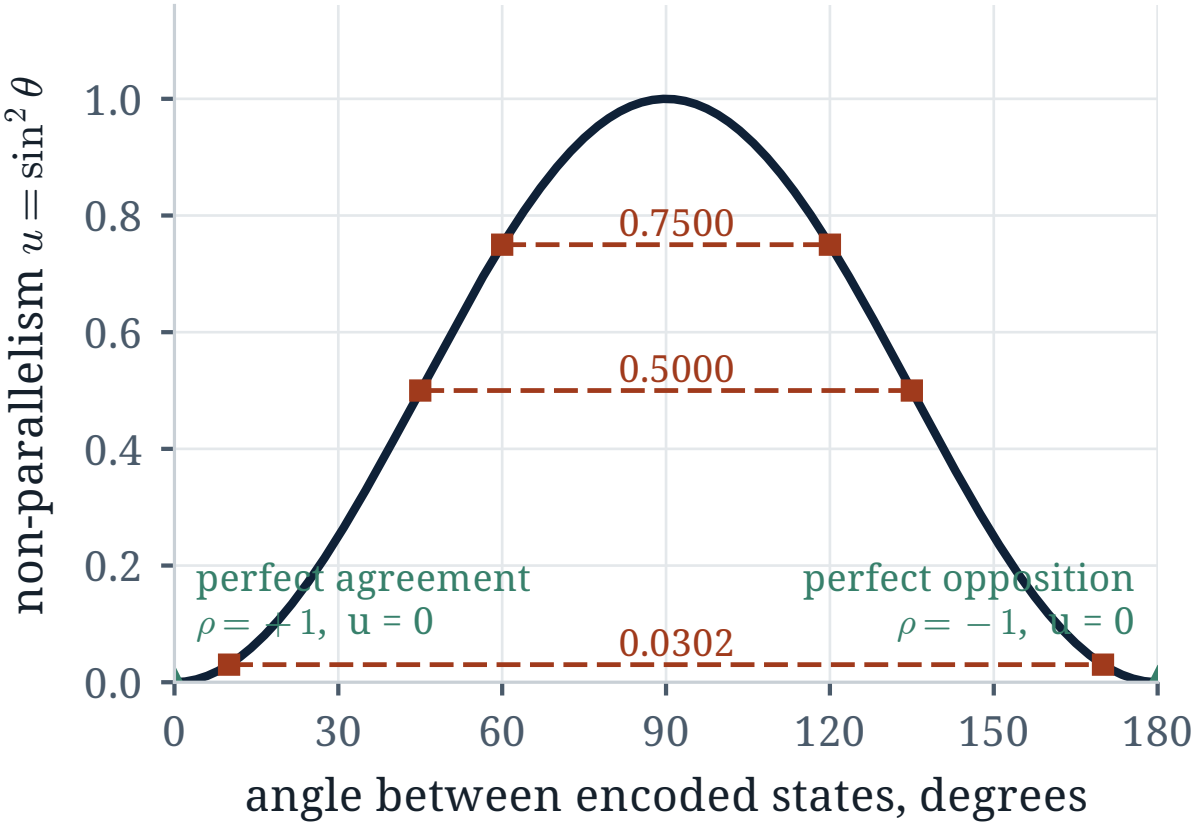


Figure 9.2: **Opposition prices as parallelism: the curve folds about ninety degrees.** Because $u = 1 - (x \cdot y)^2$ depends only on the magnitude of the correlation, $u(\theta) = u(180^\circ - \theta)$ exactly, and the three tie-lines mark pairs of angles that price identically — 60° with 120° at 0.7500, 45° with 135° at 0.5000, 10° with 170° at 0.0302. **At $\rho = -1$ the surplus is exactly zero, as it is at $\rho = +1$.** Model output; the identity is exact. Generated by `figs/src/CCP-009_fig02.py`.

Instrument 1 — Disagreement rate

The cheapest instrument. Present the body with a set of K binary or categorical judgment items under **independent** conditions (no discussion, no visible responses, no sequencing). Record the fraction of item-pairs on which two given members differ.

For two members i and j, let `d_ij` be the fraction of items on which they disagree. Under a simple encoding in which each member's response vector is mean-centred and normalized, the sample cosine between response vectors is `1 − 2 d_ij` for balanced binary items, giving

```
u_ij  =  1 − (1 − 2 d_ij)^2
```

`d = 0` gives u = 0. `d = 0.5` — chance-level disagreement, i.e. mutual independence — gives u = 1. `d = 0.1` gives `u = 1 − 0.64 = 0.36`. **Ten percent disagreement is already a third of the way to orthogonality, which indicates how little disagreement most institutions actually tolerate.**

Caveats, and they are severe. The instrument is only valid if elicitation was genuinely independent; a show of hands measures cascade, not geometry (the pseudo-diversity analysis [Curzi, 2026d]). It is sensitive to item difficulty: items everyone gets right and items everyone gets wrong both read as u = 0 and neither carries information about the geometry. Use items with a spread of difficulty and discard items with unanimous responses before computing d.

Instrument 2 — Error decorrelation

The most informative single instrument, and the one this paper prefers whenever the body produces predictions that are later scored.

Take each member's **error** vector across a common set of scored items: `e_i = (prediction_i − outcome)` over items. Mean-centre and normalize each error vector. Then the inner product between two normalized, mean-centred vectors is exactly the Pearson correlation coefficient ρ, and

```
u_ij  =  1 − rho_ij^2
```

This is an exact identity, not an approximation, given the encoding choice. **The encoding choice — that error vectors are the right states — is a modeling assumption and is badged as such [POSTULATE]**; the algebra that follows from it is not.

```
TABLE 4.3 — ERROR CORRELATION -> u -> AUDITED COUNT (illustrated at N = 8)

   rho     u = 1 − rho^2   N_eff     ln N_eff    % of ln 8
   ----    -------------   ------    --------    ---------
   0.99       0.0199        1.139     0.1304        6.3%
   0.95       0.0975        1.683     0.5203       25.0%
   0.90       0.1900        2.330     0.8459       40.7%
   0.80       0.3600        3.520     1.2585       60.5%
   0.70       0.5100        4.570     1.5195       73.1%
   0.60       0.6400        5.480     1.7011       81.8%
   0.50       0.7500        6.250     1.8326       88.1%
   0.40       0.8400        6.880     1.9286       92.7%
   0.30       0.9100        7.370     1.9974       96.1%
   0.20       0.9600        7.720     2.0438       98.3%
   0.10       0.9900        7.930     2.0707       99.6%
   0.00       1.0000        8.000     2.0794      100.0%
```

Note the shape. Because u = 1 - ρ^2, the audited count is punished by **correlation squared**, which means that the region of high correlation — where institutions actually live — is where

the penalty is steepest and least visible. Moving eight forecasters from ρ = 0.90 to ρ = 0.80 nearly doubles the effective count, from 2.33 to 3.52. Moving them from ρ = 0.20 to ρ = 0.10 moves it from 7.72 to 7.93. **The returns to decorrelation are concentrated exactly where organizations feel most comfortable.**

Instrument 3 — Prediction diversity, via a variance identity

Standard mathematics, offered here without any citation because none is needed: for N estimators each with error variance σ^2 and uniform pairwise correlation ρ, the variance of their simple average is

```
Var(mean)  =  (sigma^2 / N) * (1 + (N-1) rho)
```

Define the **classical variance-reduction count** `N_ind = N / (1 + (N-1)ρ)` — the number of independent estimators that would give the same variance reduction. It is worth tabulating because it is the quantity most working analysts already have intuitions about.

TABLE 4.4 — CLASSICAL VARIANCE-REDUCTION COUNT N_ind = N / (1 + (N-1) rho)

N	rho=0.0	rho=0.1	rho=0.2	rho=0.4	rho=0.6	rho=0.8	rho=0.95
2	2.00	1.82	1.67	1.43	1.25	1.11	1.03
4	4.00	3.08	2.50	1.82	1.43	1.18	1.04
8	8.00	4.71	3.33	2.11	1.54	1.21	1.05
12	12.00	5.71	3.75	2.22	1.58	1.22	1.05
20	20.00	6.90	4.17	2.33	1.61	1.23	1.05
50	50.00	8.47	4.63	2.43	1.64	1.24	1.05
100	100.00	9.17	4.81	2.46	1.66	1.25	1.05
1000	1000.00	9.91	4.98	2.50	1.67	1.25	1.05

As N grows, `N_ind -> 1/ρ`, a hard asymptote. **A thousand analysts at ρ = 0.2 are worth five independent analysts, and the next nine hundred and ninety are worth nothing at all.**

Now a required warning. `N_ind = N/(1+(N-1)ρ)` and `N_eff = 1+(N-1)u` are **different quantities** and must not be equated. They share a structural motif — the `1 + (N-1)·` form — because both are affine interpolations between "one thing" and "N things," which is a natural way to build a count, and S3 chose that same form deliberately. **The resemblance is a design echo, not a theorem, and this paper badges it as such [OBSERVATION, NOT DERIVATION].** Use Table 4.4 for intuition and sanity-checking; use Table 4.3 when computing surplus.

Instrument 4 — The collision test

Adapted from the framework's collision instrument [Curzi, 2026]. Present two members with the same stimulus under blind conditions and record whether their outputs *collide* — that is, land in the same cell of a pre-declared response partition. Over K trials, the collision rate `c_ij` estimates the probability that the two members are functionally the same instrument.

The collision test is coarser than error decorrelation but it has two advantages that make it the field instrument of choice. It requires no outcome data, so it can be run before anything is scored or on questions that will never be resolved. And it is legible to non-technical participants, which matters more than engineers usually admit: a body will accept a redundancy audit it can see being run, and will contest one that arrives as a correlation matrix from an analyst.

Map collision rate to u by treating the response partition as the encoder:

```
u_ij  ~  1 - c_ij       (for a partition with many roughly equiprobable cells)
```

State the approximation honestly. With a coarse partition — say three cells — random collision alone runs near 1/3, so raw collision rate over-reads similarity. Correct for chance:

```
u_ij  ~  (1 − c_ij) / (1 − c_chance)      clipped to [0,1]
```

where `c_chance` is the collision rate expected if both members answered at random from the observed marginal distribution. **A collision audit that does not subtract chance collision is not a measurement; it is a mood.**

Instrument 5 — The lineage census

The bluntest and often the most decisive. Enumerate, for every member, the answers to a fixed set of provenance questions: where they were trained, by whom, in what decade, in what language, under what funding model, with what incentive attached to being right. Two members who answer identically on all axes are provisionally scored at high correlation until measured otherwise.

The lineage census is not a substitute for Instruments 1–4 and this paper does not claim that provenance determines view. **It is a prior, not a measurement, and it is badged as such [HEURISTIC].** Its function is triage: it indicates which pairs to measure first, and it makes the redundancy structure visible to people who will never read a correlation matrix. It is also the only one of the five instruments that can be run on a hiring pipeline before anyone is hired, which makes it the only one that can prevent the problem rather than diagnose it.

9.4.4 4.4 Concavity — the merciful clause

`f` is strictly concave in u (Curzi [2026], Thm 3.3), and `ln N_eff` is concave in N as well. The practical consequence is the most useful budgeting fact in the framework.

Fix u. Ask what the Nth member adds:

```
ΔF(N)  =  ln(1 + (N−1)u)  −  ln(1 + (N−2)u)
```

TABLE 4.5 — MARGINAL SURPLUS OF THE Nth MEMBER, AT FIXED u
(values are nats; the total ceiling is ln N)

N	u = 0.1	u = 0.3	u = 0.5	u = 0.8	u = 1.0
2	0.0953	0.2624	0.4055	0.5878	0.6931
3	0.0870	0.2076	0.2877	0.3677	0.4055
4	0.0800	0.1719	0.2231	0.2683	0.2877
5	0.0741	0.1466	0.1823	0.2113	0.2231
6	0.0690	0.1278	0.1542	0.1744	0.1823
7	0.0645	0.1133	0.1335	0.1484	0.1542
8	0.0606	0.1018	0.1178	0.1292	0.1335
9	0.0572	0.0924	0.1054	0.1144	0.1178
10	0.0541	0.0846	0.0953	0.1027	0.1054
11	0.0513	0.0780	0.0870	0.0931	0.0953
12	0.0488	0.0723	0.0800	0.0852	0.0870
13	0.0465	0.0674	0.0741	0.0785	0.0800
14	0.0445	0.0632	0.0690	0.0728	0.0741
15	0.0426	0.0594	0.0645	0.0678	0.0690
16	0.0408	0.0561	0.0606	0.0635	0.0645
17	0.0392	0.0531	0.0572	0.0597	0.0606
18	0.0377	0.0504	0.0541	0.0564	0.0572
19	0.0364	0.0480	0.0513	0.0533	0.0541

```
  20 |   0.0351  |   0.0458  |   0.0488  |   0.0506  |   0.0513
  21 |   0.0339  |   0.0438  |   0.0465  |   0.0482  |   0.0488
  22 |   0.0328  |   0.0420  |   0.0445  |   0.0460  |   0.0465
  23 |   0.0317  |   0.0403  |   0.0426  |   0.0440  |   0.0445
  24 |   0.0308  |   0.0387  |   0.0408  |   0.0421  |   0.0426
  25 |   0.0299  |   0.0373  |   0.0392  |   0.0404  |   0.0408
  26 |   0.0290  |   0.0359  |   0.0377  |   0.0388  |   0.0392
  27 |   0.0282  |   0.0347  |   0.0364  |   0.0374  |   0.0377
  28 |   0.0274  |   0.0335  |   0.0351  |   0.0360  |   0.0364
  29 |   0.0267  |   0.0324  |   0.0339  |   0.0348  |   0.0351
  30 |   0.0260  |   0.0314  |   0.0328  |   0.0336  |   0.0339
  31 |   0.0253  |   0.0305  |   0.0317  |   0.0325  |   0.0328
  32 |   0.0247  |   0.0296  |   0.0308  |   0.0315  |   0.0317
  33 |   0.0241  |   0.0287  |   0.0299  |   0.0305  |   0.0308
  34 |   0.0235  |   0.0279  |   0.0290  |   0.0296  |   0.0299
  35 |   0.0230  |   0.0272  |   0.0282  |   0.0288  |   0.0290
  36 |   0.0225  |   0.0264  |   0.0274  |   0.0280  |   0.0282
  37 |   0.0220  |   0.0258  |   0.0267  |   0.0272  |   0.0274
  38 |   0.0215  |   0.0251  |   0.0260  |   0.0265  |   0.0267
  39 |   0.0211  |   0.0245  |   0.0253  |   0.0258  |   0.0260
  40 |   0.0206  |   0.0239  |   0.0247  |   0.0252  |   0.0253
```

Read the u = 1.0 column. **The second genuinely different member is worth 0.6931 nats. The fortieth is worth 0.0253 — a factor of 27 less, for exactly the same salary.** At u = 0.5 the ratio is 0.4055 to 0.0247, a factor of 16. The pattern is universal across the table because it is concavity, and concavity is a theorem here, not a tendency.

Now read across a row instead. At N = 2, moving from u = 0.1 to u = 1.0 multiplies the gain by 7.3. At N = 40 it multiplies it by 1.23. **Non-parallelism matters enormously for the first few members and barely at all for the last few.** Both readings point the same way and the intervention procedure [Curzi, 2026j] will use them to rank the levers: get two or three genuinely orthogonal members, then stop buying members and start buying η.

The cumulative form is the one to quote at a budget meeting:

```
TABLE 4.6 — CUMULATIVE SURPLUS f(N,u) = ln(1 + (N−1)u), SELECTED ROWS

   N  |  u=1.00 |  u=0.50 |  u=0.30 |  u=0.15 |  u=0.05 |  u=0.00
 -----+---------+---------+---------+---------+---------+--------
    2 |  0.6931 |  0.4055 |  0.2624 |  0.1398 |  0.0488 |  0.0000
    3 |  1.0986 |  0.6931 |  0.4700 |  0.2624 |  0.0953 |  0.0000
    4 |  1.3863 |  0.9163 |  0.6419 |  0.3716 |  0.1398 |  0.0000
    5 |  1.6094 |  1.0986 |  0.7885 |  0.4700 |  0.1823 |  0.0000
    8 |  2.0794 |  1.5041 |  1.1314 |  0.7178 |  0.3001 |  0.0000
   12 |  2.4849 |  1.8718 |  1.4586 |  0.9746 |  0.4383 |  0.0000
   20 |  2.9957 |  2.3514 |  1.9021 |  1.3481 |  0.6678 |  0.0000
   40 |  3.6889 |  3.0204 |  2.5416 |  1.9242 |  1.0818 |  0.0000
  100 |  4.6052 |  3.9220 |  3.4243 |  2.7632 |  1.7834 |  0.0000
```

9.4.5 4.5 Three honest windows set wide

Compare two cells in Table 4.6.

```
  N =  3,  u = 1.00   ->   S = 1.0986
  N = 12,  u = 0.15   ->   S = 0.9746
```

Three honest windows set wide capture more than a gallery of twelve pointing the same way. Three people, no shared training, no shared incentive, genuinely orthogonal instruments, beat a twelve-person panel at fifteen percent non-parallelism — and they do it while generating `3·2/2 = 3` communication paths instead of `12·11/2 = 66`, which by the cost accounting [Curzi, 2026b] is a 22-fold reduction in the quadratic term.

The comparison generalizes. Wherever a choice between more members and more difference is offered, both cells should be computed and compared. The result is almost always the same, because the two terms enter the formula asymmetrically: N enters inside a logarithm, u enters linearly inside that same logarithm and multiplies (N-1). Doubling u is not the same shape of operation as doubling N. Here is the comparison at a concrete baseline:

```
baseline: N = 12, u = 0.15          S = 0.9746
double N to 24, u unchanged         S = 1.4929   (+0.5183, twelve new hires)
hold N = 12, double u to 0.30       S = 1.4586   (+0.4841, zero new hires)
quadruple N to 48, u unchanged      S = 2.0857   (+1.1111, thirty-six new hires)
hold N = 12, u to 0.60              S = 2.0281   (+1.0536, zero new hires)
```

Doubling the effective difference within an existing body is worth almost exactly what doubling the body is worth, and quadrupling the difference is worth almost exactly what quadrupling the body is worth — at a fraction of the cost, and with none of the quadratic coordination penalty. This is the central practical result of this paper and Section 6 is devoted to executing it.

9.4.6 4.6 The Lipschitz constant is a measurement-error budget

Curzi [2026], Thm 4.2 gives the sharp Lipschitz constant N-1:

```
| f(u_1) − f(u_2) |  <=  (N−1) | u_1 − u_2 |
```

This is usually read as a smoothness result. Read it instead as a **specification for how precisely u must be measured**. The local slope is

```
df/du  =  (N−1) / (1 + (N−1)u)
```

which equals the bound N-1 only at u = 0 and falls off from there.

TABLE 4.7 — SENSITIVITY df/du, AND THE SHARP BOUND

N	u=0.02	u=0.1	u=0.3	u=0.6	u=1.0	bound (N−1)
3	1.923	1.667	1.250	0.909	0.667	2
5	3.704	2.857	1.818	1.176	0.800	4
10	7.627	4.737	2.432	1.406	0.900	9
20	13.768	6.552	2.836	1.532	0.950	19
50	24.747	8.305	3.121	1.612	0.980	49
100	33.221	9.083	3.225	1.639	0.990	99

The engineering reading: at large N and small u, surplus is violently sensitive to u, so a sloppy estimate of u produces a wildly wrong estimate of S. At N = 12, an uncertainty of ±0.05 in u produces these ranges in S:

```
u = 0.05 ± 0.05   ->   S in [0.0000, 0.7419]    width 0.7419
u = 0.10 ± 0.05   ->   S in [0.4383, 0.9746]    width 0.5363
u = 0.20 ± 0.05   ->   S in [0.9746, 1.3218]    width 0.3472
u = 0.40 ± 0.05   ->   S in [1.5790, 1.7834]    width 0.2044
u = 0.70 ± 0.05   ->   S in [2.0980, 2.2246]    width 0.1266
```

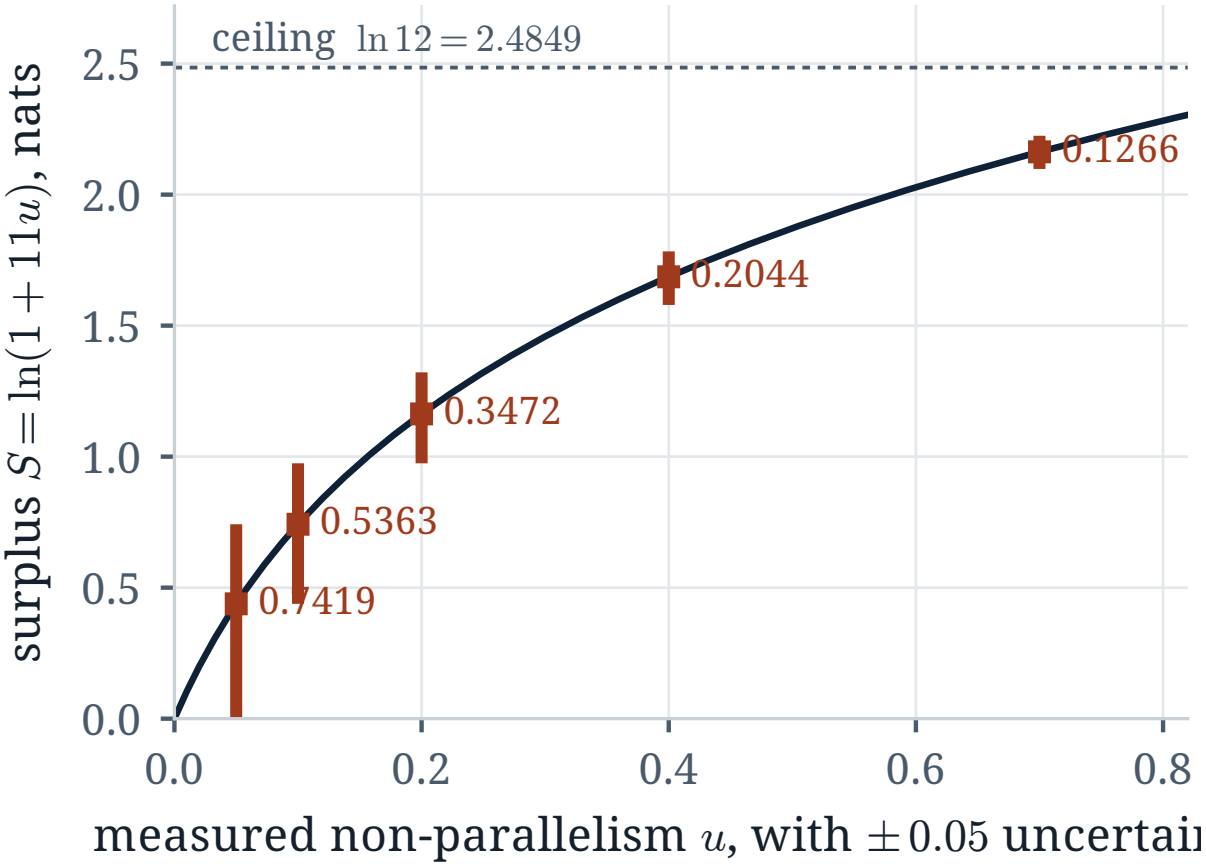


Figure 9.3: **Where institutions actually sit, a plausible measurement error spans nearly the whole range of the quantity being measured.** At $N = 12$, an uncertainty of ± 0.05 in u admits a surplus anywhere in the marked interval; the figure beside each bar is its width in nats. The width falls from 0.7419 at $u = 0.05$ to 0.1266 at $u = 0.70$, which is why the specification forbids a point estimate below $u = 0.3$. Model output at $N = 12$; arithmetic exact. Generated by `figs/src/CCP-009_fig03.py`.

Precisely in the low-u regime where institutions actually operate, a plausible measurement error in u spans nearly the entire range of surplus being estimated. Do not report a single number for S when u is small and poorly measured. Report the interval. The corpus's preregistration protocol [Curzi, 2026i] makes this a requirement.

9.4.7 4.7 Condorcet says the same thing in another language

The Condorcet jury theorem states that if each of N voters is independently correct with probability p > 0.5, the probability that a majority is correct rises with N and approaches 1. The independence clause is not decorative. It is the whole theorem.

Here is the theorem doing its work at p = 0.6, computed exactly:

```
TABLE 4.8 — MAJORITY ACCURACY, N INDEPENDENT VOTERS AT p = 0.6

    N =   1    0.6000
    N =   3    0.6480
    N =   5    0.6826
    N =   7    0.7102
    N =   9    0.7334
    N =  11    0.7535
    N =  15    0.7869
    N =  21    0.8256
    N =  31    0.8716
    N =  51    0.9265
    N = 101    0.9791
```

Now break independence. If the voters are perfectly correlated — all one instrument — the majority is correct with probability 0.6 regardless of N. In the present vocabulary, u = 0, `N_eff = 1`, surplus = 0, and the roster of 101 is a roster of 1. Condorcet's independence requirement and the surplus law's non-parallelism requirement are **the same requirement stated in two vocabularies.** Condorcet states it as a precondition and then reasons about accuracy; the surplus law states it as a continuous parameter and prices the intermediate cases. Neither is more fundamental. Having both is useful, because Condorcet is the version a lawyer or a board member already believes, and the surplus law is the version that goes into a spreadsheet.

This paper does not claim the two are mathematically equivalent, and asserts no theorem connecting them. **The correspondence is interpretive and is badged [INTERPRETIVE CORRESPONDENCE].**

9.4.8 4.8 Where the effective count enters the circuit

Return to the storage law. The storage law is

```
Q_{t+1} = (1−δ) Q_t + η S_t − C_t,     S_t = ln(1 + (N−1) U_t) = ln N_eff(t)
```

and the master gauge is

```
h_t = η ln N − δ Q_t − C_t
```

Two distinct roles for the count, and confusing them is a recurring error:

- **`ln N` is the ceiling.** It appears in headroom, in the sustainability ceiling (Cor 3.4), in the uniform bound (Thm 3.3), in the maximum one-step increment (Cor 4.2). It is what the system *could* draw if every interaction were orthogonal. Section 5 is entirely about this term.

- **`ln N_eff` is the draw.** It is what the system actually delivers this period, and it is strictly less than `ln N` unless u = 1.

The gap between them, `ln N − ln N_eff = ln(N / N_eff)`, is the **redundancy deficit**. At N = 12 and u = 0.15 it is `2.4849 − 0.9746 = 1.5103` nats, meaning the panel is leaving 61% of its theoretical input on the table every single period, forever, until someone changes the geometry. That deficit compounds through the integrator: over the 6.58-period half-life corresponding to δ = 0.10 (the retention analysis [Curzi, 2026a], Table L), it is not a one-time loss but a permanently lowered steady state, since `Q_\ensuremath{\infty} = (ηF(u*) − c*)/δ` scales directly with F.

Redundancy is not a missed opportunity. It is a permanent reduction in the fixed point.

9.4.9 4.9 Specification: the effective-count requirements

Numbered, checkable requirements. A body that satisfies EC-1 through EC-12 has been audited for effective count. A body that has not been audited should not be described in a plan as having N members; it should be described as having an unmeasured N_eff.

EC-1 — Declare the encoder. Before any u is estimated, write down in one sentence what a member's *state* is taken to be: their scored error vector, their independent response vector on a declared item set, or their cell assignment under a declared partition. No estimate of u is admissible without a declared encoder [Curzi, 2026].

EC-2 — Elicit independently or do not elicit. All responses used to estimate u must be gathered under conditions where no member can observe another's response before committing their own. Written, simultaneous, sealed. Anything else measures cascade, not geometry. This requirement is absolute and is not subject to convenience.

EC-3 — Report N, u, and N_eff together, never N alone. Any document that states the size of a body must state the audited size beside it, in the form `N = 12, u = 0.15 ± 0.05, N_eff = 2.65 [1.55, 4.30]`. A bare headcount in a design document is a specification defect.

EC-4 — Report an interval, not a point, whenever u < 0.3. Table 4.7 shows the sensitivity. Below u = 0.3 the surplus estimate is dominated by measurement error and a point estimate is misleading.

EC-5 — Subtract chance collision. Any collision-based estimate must state `c_chance` and report the chance-corrected figure. Uncorrected collision rates are not admissible.

EC-6 — Discard unanimous items. Items on which every member agrees carry no information about geometry, whether the unanimity is correct or incorrect. Compute u on the discriminating subset and report how many items were discarded. If more than 80% of items were discarded, the item set is not fit for purpose and must be rebuilt.

EC-7 — Estimate u at least once per δ-half-life. u decays (Section 6, the pseudo-diversity analysis [Curzi, 2026d]). An estimate older than one half-life of the body's convergence process is not a measurement; it is a memory. The retention analysis [Curzi, 2026a] gives `t_half = ln 2 / (−ln(1−δ))`; use the convergence rate of u, not the decay rate of Q.

EC-8 — Publish the pairwise matrix when a decision turns on structure. The mean pairwise u is a summary and the Roster Reduction discards structure. When a decision depends on whether difference is spread or concentrated in one member, publish the matrix.

EC-9 — Run a lineage census on every hiring pipeline, annually. Provenance axes at minimum: training institution, principal mentor, decade of formation, working language, funding source, prior industry. Report the entropy of each axis across the current roster. This is a prior, not a measurement (Section 4.3, Instrument 5), and must be labelled as such wherever reported.

EC-10 — Size against the marginal table, not against a headcount target. Before adding the Nth member, look up the marginal gain in Table 4.5 at the measured u and compare it against the member's full loaded cost including the `N−1` additional communication paths they

create. If the marginal surplus is below 0.05 nats, the correct move is almost certainly to raise u instead.

EC-11 — Compute and publish the redundancy deficit. `ln N − ln N_eff`. This single number is the most useful summary of a body's geometric health and it belongs on the same dashboard as headcount and budget.

EC-12 — Badge the Roster Reduction wherever it is used. Any report applying `N_eff = 1+(N−1)u` to a roster of people rather than to a declared pair of states must carry the badge and one sentence naming the failure mode: mean pairwise u does not determine joint geometry, because the triangle inequality fails for geometric triples.

9.4.10 4.10 Failure modes native to the audited count

Each entry: NAME, clinical description, damaged parameter, diagnostic signature, detecting measurement, fix. The corpus's consolidated failure catalogue [Curzi, 2026j] carries the full set; the entries below are the ones native to the audited count.

FM-4.1 · THE HEADCOUNT FALLACY

Damages: the estimate of S, and thereby every downstream budget.

Clinical description. The body reports its size and reasons from it. Plans, budgets, and capability claims are all scaled to N. u is never estimated, so `N_eff` is never known, and the gap between the claimed and delivered input is absorbed silently as underperformance attributed to effort or morale.

Diagnostic signature. Any planning document containing a headcount and no non-parallelism figure. Statements of the form "with twelve reviewers we have comprehensive coverage." Retrospectives that conclude "we need more people" after a body has already grown twice without improving.

Detecting measurement. EC-3 compliance audit. Count planning documents; count the fraction containing a u estimate. Below 100% is a defect.

Fix. Institute EC-3 as a document standard. It costs nothing and it is the single highest-leverage administrative change in this paper, because it forces the question into every document where a decision is made.

FM-4.2 · THE PARALLEL ROSTER

Damages: u -> 0, hence `N_eff -> 1`, hence S -> 0.

Clinical description. Every member is a functional copy of every other. Recruitment, socialization, and promotion have converged the body onto a single instrument. Members experience the condition as agreement, competence, and pleasant working relations. It is the most comfortable of all the failure modes and therefore the most durable.

Diagnostic signature. Disagreement rate under blind elicitation below 5%. Error correlation above 0.9. Meetings that end early. Unanimous votes on non-trivial questions. New members who "fit in immediately."

Detecting measurement. Instrument 2 (error decorrelation) run on the last 30 scored decisions, or Instrument 1 on a fresh blind item set of 20 discriminating items.

Fix. Section 6 in full. In outline: recruit across a lineage boundary, install rotation, mandate an adversarial role. Note that the fix is a **u** intervention, not an N intervention; adding members from the same pipeline makes the failure worse by raising C while leaving `N_eff` at 1.

This mode is important enough that the pseudo-diversity analysis [Curzi, 2026d] is devoted to its detection and its socially disguised variants.

FM-4.3 · THE MISSING DENOMINATOR

Damages: the credibility of the whole measurement chain.

Clinical description. u is reported without its N. Because `N_eff = 1 + (N−1)u`, a u of 0.4 means an effective count of 1.4 at N = 2 and 40.6 at N = 100. A u reported alone is uninterpretable.

Diagnostic signature. A dashboard tile reading "diversity: 0.42."

Detecting measurement. Inspection.

Fix. EC-3. Never report u without N and `N_eff`.

FM-4.4 · SUBSTITUTION OF PROVENANCE FOR GEOMETRY

Damages: the measurement, not the system — but it hides damage to u.

Clinical description. The lineage census (Instrument 5) is treated as a measurement of u rather than as a prior. The body reports high provenance entropy and concludes it has high u, when in fact its members were drawn from different institutions into a single shared framework and their error vectors correlate at 0.92.

Diagnostic signature. Reports that cite composition and never cite disagreement. Assertion that a body is "diverse" with no scored-error figure anywhere in the document.

Detecting measurement. Run Instrument 2 and compare with what Instrument 5 predicted. A large gap between predicted and measured u is itself the finding, and it is a useful one: it establishes that the homogenizing force is downstream of recruitment, in training or incentive.

Fix. Demote provenance to triage, per EC-9's labelling requirement. Then look downstream: shared training, shared incentive, shared review. Section 6.7.

FM-4.5 · THE CASCADE-CONTAMINATED ESTIMATE

Damages: the estimate of u, always in the direction of overstating similarity.

Clinical description. u is estimated from responses gathered in sequence or in view of one another. Members anchored on the first response. The measured disagreement rate reflects the cascade, not the geometry. The body then concludes it is more parallel than it is, or — worse — concludes it is appropriately aligned and stops trying.

Diagnostic signature. Disagreement rate strongly dependent on speaking order. First responder's answer disproportionately represented in the modal response. Round-table elicitation of any kind.

Detecting measurement. Run the same item set twice, once sequentially and once with silent-write-first. Compare the two u estimates. A gap above 0.15 indicates cascade contamination of the sequential protocol; use only the blind estimate thereafter.

Fix. EC-2, absolutely. Silent-write-first, sealed simultaneous submission. The seating and furniture analysis [Curzi, 2026f] gives the furniture that makes this the default and the individuality analysis [Curzi, 2026h] gives the protocol.

FM-4.6 · THE SATURATED PANEL

Damages: C, without damaging S — which is why it survives.

Clinical description. A body that has passed the point where additional members buy meaningful surplus keeps adding members. From Table 4.5, at u = 0.5 the 30th member buys 0.0328 nats while adding 29 communication paths. The surplus curve is flat; the cost curve is not.

Diagnostic signature. Growth in headcount with flat or declining output. Meetings where a majority of attendees do not speak. Sub-committees forming spontaneously — which, note, is the body correctly self-organizing away from the saturated configuration.

Detecting measurement. Compute marginal surplus per member at measured u (Table 4.5) and compare against loaded cost per member including the quadratic coordination term from the cost analysis [Curzi, 2026b]. The crossover point is computable and belongs in every org design review.

Fix. Split. Two bodies of six at u = 0.5 deliver `2 × 1.5041 = 3.008` nats of input against `2 × 15 = 30` communication paths; one body of twelve at the same u delivers 1.8718 nats against 66 paths. **Splitting a saturated body is not an admission of failure; it is the arithmetic.** Section 5.9 and the intervention procedure [Curzi, 2026j] give the sequencing.

9.5 5 The logarithmic ceiling

Section 4 established that a body's audited count is `N_eff = 1 + (N−1)u`, and that the audit is usually brutal. The obvious response is to raise N. Add domains. Widen the remit. Bring in the adjacent discipline, then the one after that. If twelve categories are not enough, take thirty.

Section 5 prices that move, and the price is a logarithm.

The result is not a caution. It is a structural prohibition, and it is the single most useful negative result in the framework: **a headroom problem cannot be escaped by adding domains.** Not because adding domains is bad, but because the return on N is `ln N`, and `ln N` grows so slowly that the quantities being outrun — δQ and C, which grow linearly and quadratically — will always win in the end. Section 5 establishes that with explicit arithmetic and then states what to do instead.

9.5.1 5.1 The ceiling

From Curzi [2026], Cor 2.2:

```
0  <=  F(u)  <=  ln N
```

with the upper bound attained if and only if the states are orthogonal (Curzi [2026], Thm 4.3). `ln N` is not a soft target or an asymptote. It is a hard ceiling on surplus per interaction, and it is set entirely by N.

The ceiling enters the theory through axiom **S4**, the normalization `f(1) = ln N`. Read S4 correctly: it says that when two states are maximally non-parallel, the surplus they generate is the log of the number of categories the ambient space can express. N is not the number of people in the room. **N is the number of genuinely distinct categories the system is capable of distinguishing.** Section 4's Roster Reduction [POSTULATE] permits reading a roster of people as a set of categories for practical work, but the underlying quantity is a dimension count, not a headcount, and Section 5 is about the dimension count.

9.5.2 5.2 Where the ceiling appears

`ln N` does not appear once. It appears in every design inequality in the storage-law analysis [Curzi, 2026a], always multiplied by η, and always as the term that caps what the system can do.

Result	Statement	Role of `ln N`
Curzi [2026], Cor 2.2	`0 <= F(u) <= ln N`	The per-interaction ceiling itself.

Result	Statement	Role of `ln N`
Curzi [2026a], Thm 3.3	`Q_t <= max{Q_0, η ln N / δ}`	Caps the *entire trajectory* of stored order.
Curzi [2026a], Thm 3.3 (δ=0)	`Q_t <= Q_0 + t η ln N`	Caps the growth *rate* when there is no leak.
Curzi [2026a], Cor 3.4	Nondecreasing step requires `C_t <= η ln N − δ Q_t`	Caps the affordable cost.
Curzi [2026a], Prop 4.1	Feasible only if `δQ_t + C_t + r_t <= η ln N`	Caps the requestable increment.
Curzi [2026a], Cor 4.2	`r_max = η ln N − δQ_t − C_t`	Caps the maximum one-step gain.
the storage-law analysis [Curzi, 2026a], §6.4	`h_t = η ln N − δQ_t − C_t`	The master gauge.

Five of these seven are the same inequality wearing different clothes, and the clothes matter less than the fact that **η ln N is the entire budget from which δQ and C are paid.** Everything a system can spend on maintenance, coordination, and growth comes out of a single term that scales as the logarithm of its category count.

9.5.3 5.3 The honesty clauses, restated for the ceiling

1. **The founding surplus analysis [Curzi, 2026] is pure mathematics and makes no empirical or metaphysical claim.** That `ln N` is a ceiling is a theorem about a functional form. That real institutions are bounded by it is a modeling proposal, falsifiable by the protocols of the preregistration protocol [Curzi, 2026i].
2. **S3, the affine effective count, is a design axiom — chosen, not derived.** The logarithmic ceiling is a consequence of having chosen an affine `g(u)` together with the normalization S4. A different choice of S3 would give a different ceiling law. The framework's negative result about adding domains is therefore conditional on S3, and the conditionality is stated here in the open.
3. **As deployed, the surplus functional is a lower bound**, not the ensemble's information. `ln N` bounds the floor that can be logged, and the honest claim is that nothing more can be *guaranteed* than `ln N`, not that no more exists.
4. **Surplus is not Shannon entropy or mutual information** [Curzi, 2026]. The similarity in form between `ln N` and the entropy of a uniform distribution over N states is a formal resemblance arising from both being normalizations over a category count. **It is not an identification and this paper badges it [FORMAL RESEMBLANCE, NOT IDENTIFICATION].** Section 5.9 uses it as a bridge to Ashby with the badge attached.
5. **The triangle inequality fails for geometric triples** (the surplus functional's metric analysis [Curzi, 2026]). When this section discusses composing categories, it composes them through `u` and the subadditivity relation on `u` only.
6. **No thermodynamic negentropy law is proved** [Curzi, 2026a]. Nothing here licenses a claim about physical order.

9.5.4 5.4 How slowly the logarithm actually grows

Engineers know that logarithms grow slowly. Almost nobody has looked at the table. Here it is, for every integer from 2 to 100.

```
TABLE 5.1 — ln N, N = 2 TO 100

   N   ln N     |    N   ln N     |    N   ln N     |    N   ln N     |    N   ln N
```

2	0.6931	22	3.0910	42	3.7377	62	4.1271	82	4.4067
3	1.0986	23	3.1355	43	3.7612	63	4.1431	83	4.4188
4	1.3863	24	3.1781	44	3.7842	64	4.1589	84	4.4308
5	1.6094	25	3.2189	45	3.8067	65	4.1744	85	4.4427
6	1.7918	26	3.2581	46	3.8286	66	4.1897	86	4.4543
7	1.9459	27	3.2958	47	3.8501	67	4.2047	87	4.4659
8	2.0794	28	3.3322	48	3.8712	68	4.2195	88	4.4773
9	2.1972	29	3.3673	49	3.8918	69	4.2341	89	4.4886
10	2.3026	30	3.4012	50	3.9120	70	4.2485	90	4.4998
11	2.3979	31	3.4340	51	3.9318	71	4.2627	91	4.5109
12	2.4849	32	3.4657	52	3.9512	72	4.2767	92	4.5218
13	2.5649	33	3.4965	53	3.9703	73	4.2905	93	4.5326
14	2.6391	34	3.5264	54	3.9890	74	4.3041	94	4.5433
15	2.7081	35	3.5553	55	4.0073	75	4.3175	95	4.5539
16	2.7726	36	3.5835	56	4.0254	76	4.3307	96	4.5643
17	2.8332	37	3.6109	57	4.0431	77	4.3438	97	4.5747
18	2.8904	38	3.6376	58	4.0604	78	4.3567	98	4.5850
19	2.9444	39	3.6636	59	4.0775	79	4.3694	99	4.5951
20	2.9957	40	3.6889	60	4.0943	80	4.3820	100	4.6052
21	3.0445	41	3.7136	61	4.1109	81	4.3944		

Sit with the last cell. A hundredfold increase in category count — from one to a hundred — buys 4.6052 nats. From two to a hundred it buys 3.9120. That is the whole return, for the whole journey, from a small team to a large institution.

Now the same table read as a fraction of what a hundred categories deliver:

TABLE 5.2 — SHARE OF ln 100 CAPTURED BY THE FIRST N CATEGORIES

N	ln N	share of ln 100
2	0.6931	15.1%
3	1.0986	23.9%
4	1.3863	30.1%
5	1.6094	34.9%
7	1.9459	42.3%
10	2.3026	50.0%
15	2.7081	58.8%
20	2.9957	65.1%
30	3.4012	73.9%
50	3.9120	84.9%
75	4.3175	93.8%
100	4.6052	100.0%

Ten categories capture exactly half of what a hundred capture. The remaining ninety categories divide the other half between them. The first five capture 35%. The first three capture nearly a quarter.

This is the same concavity that made Section 4's marginal table merciful, and it cuts the same way: the early categories are where the value is, and the late ones are rounding.

9.5.5 5.5 The doubling constant

The most useful single fact about the ceiling is scale-invariance of its increment.

```
ln(2N) − ln(N)  =  ln 2  =  0.693147...     for every N whatsoever
```

```
TABLE 5.3 — DOUBLING ALWAYS ADDS ln 2 = 0.6931

       N      ln N       gain from N/2 -> N
     -----   -------     ------------------
         2   0.6931           0.6931
         4   1.3863           0.6931
         8   2.0794           0.6931
        16   2.7726           0.6931
        32   3.4657           0.6931
        64   4.1589           0.6931
       128   4.8520           0.6931
       256   5.5452           0.6931
       512   6.2383           0.6931
      1024   6.9315           0.6931
```

Doubling the number of domains adds exactly 0.6931 nats, whether the move is from 2 to 4 or from 512 to 1024. The absolute gain never changes. The relative gain collapses: 0.6931 is 100% of ln 2 but 10% of ln 1024.

The corollary that governs planning:

```
to add ONE nat to the ceiling, multiply N by  e = 2.71828...
```

```
TABLE 5.4 — WHAT +1.0 NAT OF CEILING COSTS IN CATEGORIES

from N =   2  (ln N = 0.6931)   requires N =    5.4
from N =   5  (ln N = 1.6094)   requires N =   13.6
from N =  10  (ln N = 2.3026)   requires N =   27.2
from N =  20  (ln N = 2.9957)   requires N =   54.4
from N =  50  (ln N = 3.9120)   requires N =  135.9
from N = 100  (ln N = 4.6052)   requires N =  271.8
```

And its inverse, the form to use when someone specifies a required ceiling:

```
TABLE 5.5 — CATEGORIES REQUIRED FOR A TARGET CEILING

target ln N        N required
-----------        ----------
   0.5000              1.65
   1.0000              2.72
   1.5000              4.48
   2.0000              7.39
   2.5000             12.18
   3.0000             20.09
   3.5000             33.12
   4.0000             54.60
   4.6052            100.00
   5.0000            148.41
   6.0000            403.43
   7.0000           1096.63
```

A ceiling of 7 nats requires 1,097 genuinely distinct categories. Not 1,097 people — 1,097 categories, each of which must be distinguishable from every other by the system's own encoder, and each of which must be actually populated by someone capable of occupying it.

9.5.6 5.6 The consequence, stated hard

The structural prohibition this section exists to establish is the following.

The master gauge is `h_t = η ln N − δQ_t − C_t`. Suppose h < 0 — the system is in decline, and by The sustainability corollary [Curzi, 2026a] **no arrangement of interaction geometry prevents order from decreasing.** The manager's instinct is to grow: add teams, add disciplines, add scope, add categories, and thereby raise `η ln N` until h turns positive.

Compute what that requires. To close a deficit of size D by raising N alone:

```
η ln N_new  −  η ln N_old  =  D
ln(N_new / N_old)          =  D / η
N_new                      =  N_old · e^{D/η}
```

The multiplier is `e^{D/η}`, and it is exponential in the deficit.

TABLE 5.6 — MULTIPLIER ON N REQUIRED TO CLOSE A DEFICIT D, AT η = 0.40

deficit D (nats)	N must be multiplied by e^{D/0.40}
0.10	1.3 x
0.25	1.9 x
0.50	3.5 x
1.00	12.2 x
2.00	148.4 x

A deficit of two nats — which in the worked systems of the storage-law analysis [Curzi, 2026a] is an unremarkable size of shortfall — requires a **148-fold** increase in category count. At η = 0.40, an institute with eight domains and a two-nat deficit needs 1,187 domains.

And here is the part that makes it a prohibition rather than an expense: **δQ and C do not hold still while N is multiplied.** C grows at least linearly in N and, through Brooks' `n(n−1)/2` communication paths, often quadratically. Q grows because a larger system stores more, so δQ grows too. You are trying to outrun linear and quadratic growth with a logarithm. **The logarithm loses, always, and it loses faster the further it is pushed.**

This is worth writing as a rule, because it will be violated in almost every organization the reader will ever work in.

> **THE DOMAIN-ADDITION PROHIBITION.** A headroom deficit cannot be closed by raising N. The required multiplier is `e^{D/η}`, exponential in the deficit, while the cost terms it must outrun grow at least linearly and typically quadratically in N. **If h < 0, adding domains makes h worse.** The admissible responses are: raise η, cut C, or shed Q. The headroom analysis [Curzi, 2026b] names those three; this section establishes that the fourth is not an option.

Two honest qualifications, because the prohibition is strong and strong claims need their limits stated.

First, the prohibition applies to raising N as a **remedy for a deficit**. Raising N for other reasons — because the environment genuinely presents variety the system cannot currently absorb (Section 5.9), because a new category is cheap and the ceiling is close to binding, because the initial system is being built rather than repairing one — remains legitimate. Section 5.11 gives the four legitimate cases.

Second, the prohibition assumes η is held constant while N grows. In practice adding a genuinely new domain sometimes raises η, because a new discipline brings new capture mechanisms — a new notation, a new archive, a new form of artifact. That effect is real and this paper does not deny it. But it should be attributed correctly: the gain came from η, not from N, and where it is wanted the correct move is to obtain η directly rather than buying a domain and hoping. The conversion analysis [Curzi, 2026b] gives the direct route and it is far cheaper.

9.5.7 5.7 What raising N actually requires

Suppose a decision has legitimately been taken to raise N. The question is what counts.

N is the number of **genuinely distinct categories the system can express**. It is not the number of labels in the org chart. Splitting a category into two relabeled halves does not raise N; it raises the label count while leaving the geometry untouched, and the arithmetic of Section 4 prices the result at zero.

Make this precise. Suppose the system has N categories and one of them is split into two halves with pairwise non-parallelism ε. Two things happen:

- The **nominal ceiling** rises from `ln N` to `ln(N+1)`.
- The **delivered surplus** does not, because the new pair contributes u = ε ≈ 0.

TABLE 5.7 — CEILING GAIN FROM ADDING ONE CATEGORY

N -> N+1	ceiling gain ln(N+1) − ln N	as % of current ceiling
2 -> 3	0.4055	58.50%
3 -> 4	0.2877	26.19%
4 -> 5	0.2231	16.10%
5 -> 6	0.1823	11.33%
6 -> 7	0.1542	8.60%
8 -> 9	0.1178	5.66%
10 -> 11	0.0953	4.14%
12 -> 13	0.0800	3.22%
16 -> 17	0.0606	2.19%
20 -> 21	0.0488	1.63%
32 -> 33	0.0308	0.89%
50 -> 51	0.0198	0.51%
100 -> 101	0.0100	0.22%

At N = 20, one more category raises the ceiling by 1.63%. If that category is a relabeled half of an existing one, the delivered surplus rises by approximately nothing while C rises by the cost of running a twenty-first category. Category inflation is a pure cost with a cosmetic ceiling gain, and it is the standard institutional response to a demand for breadth.

The relabeling audit

Before N is incremented in any plan, run this test on the proposed new category:

1. **State the new category's encoder.** What state does a member of this category occupy that a member of no existing category occupies? Write one sentence. If that cannot be done, stop.
2. **Estimate u between the new category and its nearest existing neighbour** using Section 4's Instruments 1 or 2. Blind elicitation on 20 discriminating items is sufficient.
3. **Require u >= 0.3 against every existing category.** Below 0.3 the new category is a relabel and must not be counted in N. Report the measured u alongside the request.
4. **Compute delivered gain, not ceiling gain.** The plan must state `Δ(delivered) = ln(1 + N·u_bar_new) − ln(1 + (N−1)·u_bar_old)`, not `ln(N+1) − ln N`.
5. **Compute the cost delta.** New coordination paths, new meetings, new interfaces. Priced by the cost accounting of Curzi [2026b].
6. **Compare.** If delivered gain is below cost delta, the category is a liability.

Steps 1 and 3 kill most proposals, which is the point. **A category that cannot state what it sees that nobody else sees is a label, and labels do not enter the law.**

The three ways N is faked

- **Splitting.** One category becomes two with the same training, the same literature, the same incentives, and a boundary drawn on the org chart. Diagnostic: measured u between the halves below 0.15. This is the most common.
- **Importing the adjacent.** Recruiting from the discipline next door — which shares 80% of its methods, its journals, and its reviewers with the existing one. Genuine but small u. Diagnostic: measured u between 0.15 and 0.35, and error correlation above 0.8. Counts partially; count it partially.
- **Nominal representation.** Adding a category to the chart and staffing it with someone whose actual working practice is unchanged from their previous category. Diagnostic: the person's error vector is unchanged from before the reassignment. This is the hardest to detect and requires scored longitudinal data.

9.5.8 5.8 The scale-free reading

The doubling constant has an architectural consequence worth stating, because it is where CANON §1.4's hyperrecursive symmetry becomes concrete.

Because `ln(2N) − ln N = ln 2` at every scale, the ceiling gain from doubling is identical at a desk, in a room, in a team, in a building, in an institution, and in a city. Two drawers instead of one; two disciplines instead of one; two wings instead of one; two boroughs instead of one. All worth 0.6931.

That is a genuine scale-free property of the model form, and it is honest to state it as such: **it is a claim about the FORM of the equation, not about physics.** The parameters δ, η, C differ enormously across those scales; only the shape of the N-dependence is invariant. This is the architectural claim of the framework and it is the reason a single set of design rules can be written for objects that differ in size by six orders of magnitude.

It also explains a common and correct intuition: that adding the second of anything feels transformative and adding the fiftieth feels like paperwork. The absolute gain is the same 0.6931 per doubling either way, but the doubling from 1 to 2 is one hire and the doubling from 50 to 100 is fifty.

9.5.9 5.9 Ashby's Law of Requisite Variety — why N matters at all

Everything above has been a case against raising N. Ashby supplies the case for.

Ashby's Law of Requisite Variety (1956): only variety can absorb variety. A regulator can hold a system's essential variables within limits only if the regulator's variety is at least as great as the variety of the disturbances it must counter. A thermostat with two states cannot regulate a process with ten distinguishable failure modes. A committee with one framework cannot govern a domain that presents four kinds of problem.

Ashby measured variety as **the logarithm of the number of distinguishable states**. That is his own formulation, and it is why the correspondence with `ln N` is close enough to be useful:

```
Ashby:        variety V = log (number of distinguishable states)
Surplus law:  ceiling   = ln N,  N = number of distinguishable categories
```

TABLE 5.8 — VARIETY IN NATS AND BITS

distinguishable states	ln (nats)	log2 (bits)
2	0.6931	1.0
4	1.3863	2.0
8	2.0794	3.0
16	2.7726	4.0

64	4.1589	6.0
256	5.5452	8.0
1024	6.9315	10.0

This paper badges the correspondence [FORMAL CORRESPONDENCE, NOT IDENTIFICATION]. Ashby's variety and the surplus ceiling are both logarithms of a state count, and both function as caps on regulatory capability. They are not proved equal, and the surplus range analysis [Curzi, 2026] explicitly declines to identify surplus with information-theoretic quantities. Use the correspondence as a bridge between two literatures, not as a theorem.

With the badge attached, the pairing is powerful, because Ashby and the ceiling say opposite and complementary things:

- **Ashby states the requirement.** If the environment presents V_env of variety and the regulator has V_reg < V_env, the shortfall shows up as unregulated variation — surprises, crises, failures the system did not have a category for. A system cannot regulate what it cannot distinguish, and no amount of effort substitutes for a missing category.
- **The ceiling states the price.** Each doubling of the regulatory repertoire costs the whole effort of doubling and returns 0.6931 nats.

Put together: N must be raised when and only when the environment presents variety the system cannot distinguish, and it must be raised by acquiring genuinely new distinctions rather than new labels — because Ashby's requirement is about distinguishability, which is exactly what Section 5.7's relabeling audit tests.

This gives the four legitimate reasons to raise N, listed in Section 5.11.

Ashby also explains why the Domain-Addition Prohibition does not contradict him. Ashby's law is about **capability**, not about **budget**. A system that lacks requisite variety will fail at regulation no matter how much headroom it has. A system with requisite variety and negative headroom will fail at persistence no matter how well it regulates. They are different failure modes with different fixes, and confusing them produces the characteristic institutional error: responding to a persistence problem (h < 0) with a capability intervention (add domains), which worsens the persistence problem without addressing anything.

Diagnose which one is present before acting. Ashby's failure looks like surprise. The headroom failure looks like erosion.

9.5.10 5.10 Alexander: a city is not a tree

Christopher Alexander's 1965 essay "A City is Not a Tree" is the most directly relevant piece of design theory in this paper's lineage, and Section 5 is where it belongs, because it is an argument about **category structure** and therefore about N.

Alexander's distinction. A **tree** is a collection of sets in which any two sets are either disjoint or one contains the other. A **semilattice** is a collection of sets in which two sets may also **overlap** — sharing members without either containing the other. Every tree is a semilattice; the converse is dramatically false.

Alexander's arithmetic, which anyone can verify:

TABLE 5.9 — DESCRIBABLE SUBSETS ON n ELEMENTS

n	tree: at most (n−1) non-trivial subsets	ln	semilattice: 2^n − n − 1	ln	ratio of ceilings
4	3	1.0986	11	2.3979	2.18 x
6	5	1.6094	57	4.0431	2.51 x
8	7	1.9459	247	5.5094	2.83 x
10	9	2.1972	1013	6.9207	3.15 x

12	11	2.3979	4083	8.3146	3.47 x
16	15	2.7081	65519	11.0901	4.10 x
20	19	2.9444	1048555	13.8629	4.71 x

On twenty elements a tree can carry at most nineteen further subsets. A semilattice on the same twenty can carry 1,048,555. Alexander stated this in the essay as “more than a million,” and the count is exactly `2^20 - 20 - 1`.

Translated into this paper’s terms: the tree structure caps N at n-1, while the semilattice structure caps it at 2^n - n - 1, and since the ceiling is ln N, the tree costs a factor of 4.71 in ceiling at twenty elements and the factor grows with n.

But the deeper point is not about the ceiling. It is about u, and it is the reason the manifest calls hierarchy the organizational form that drives u toward zero.

Why hierarchy drives u to zero

In a tree, every element belongs to exactly one chain of nested sets: itself, its parent, its grandparent, up to the root. Those are its descriptors — the categories the system uses to characterize it.

Nested categories are parallel categories. A set and its superset do not point in different directions; one is contained in the other, and the smaller carries no information the larger contradicts. Under the encoding this paper uses, the pair contributes u $\approx$ 0.

The step is badged: **[POSTULATE] a pair of nested categories is assigned u = 0, and a pair of incomparable (overlapping or disjoint but non-nested) categories is assigned some u_0 > 0 determined by measurement.** The postulate is a modeling choice; it is defensible because a nested pair conveys strictly refined information along a single axis, and it is testable by measuring u between a category and its parent using Section 4’s instruments.

With that postulate, the arithmetic is stark. Consider an element carrying m descriptors:

TABLE 5.10 — DESCRIPTOR-SET AUDIT FOR ONE ELEMENT CARRYING m LABELS
(semilattice case computed at measured u_0 = 0.7)

m	TREE: all nested		SEMILATTICE: all incomparable		ceiling
	N_eff	S	N_eff	S	ln m
2	1.000	0.0000	1.700	0.5306	0.6931
3	1.000	0.0000	2.400	0.8755	1.0986
4	1.000	0.0000	3.100	1.1314	1.3863
6	1.000	0.0000	4.500	1.5041	1.7918
8	1.000	0.0000	5.900	1.7750	2.0794
12	1.000	0.0000	8.700	2.1633	2.4849

In a strict tree, an element’s entire descriptor set audits to an effective count of one, regardless of how deep the hierarchy is. Twelve levels of nested classification and the element is described by one direction. Adding levels adds labels, cost, and review steps, and adds exactly zero surplus.

Real organizations are not perfect trees and not complete semilattices. Parameterize by the fraction q of an element’s descriptor pairs that are incomparable rather than nested:

TABLE 5.11 — PARTIAL OVERLAP. m = 6 DESCRIPTORS, u_0 = 0.7

q (fraction incomparable)	u_bar	N_eff	S	% of ln 6
0.0	0.000	1.000	0.0000	0.0%
0.2	0.140	1.700	0.5306	29.6%

0.4	0.280	2.400	0.8755	48.9%
0.6	0.420	3.100	1.1314	63.1%
0.8	0.560	3.800	1.3350	74.5%
1.0	0.700	4.500	1.5041	83.9%

Moving an organization from q = 0.0 to q = 0.2 — allowing one descriptor pair in five to overlap rather than nest — buys 29.6% of the available ceiling for that element's descriptor set, and it is the cheapest structural change in this section.

What does q > 0 look like in practice? It looks like:

- A person who reports to a function and also belongs to a standing cross-function body, where neither membership is a subset of the other.
- A budget line owned jointly by two units, neither of which contains the other.
- A physical space used by two groups on overlapping schedules — Jacobs' mixed primary uses, which is the same structure in a city.
- A document with two independent owners in different branches.
- A review process in which a proposal is read by its own division and by a body that spans divisions, with neither review subordinate to the other.

And what does q = 0 look like? It looks like a clean org chart. The org chart is legible precisely because it is a tree, and it is a tree precisely because someone removed all the overlaps, and the overlaps were the surplus. This is Alexander's complaint about planned cities restated for institutions, and it is the same complaint: the tree is what a designer produces when the design must be comprehensible to the designer, and comprehensibility to the designer is not one of the system's requirements.

The silo analysis [Curzi, 2026g] develops the organizational consequence (silos as local minima). The room-as-circuit analysis [Curzi, 2026d] develops the spatial consequence. The point here is arithmetic: **hierarchy is a u-destroying structure, and since surplus is `ln(1 + (N−1)u)`, destroying u destroys the return on every category already paid for.**

The semilattice is not free

Honesty requires the other side. Overlapping structure raises C. A person with two non-nested memberships attends two sets of meetings, answers to two review paths, and can be given contradictory instructions. Alexander's essay does not price this; this paper must.

The cost of overlap is roughly proportional to the number of overlapping memberships per element, which is a linear term, while the surplus gain is inside a logarithm. Therefore overlap, like every other lever, saturates: the first overlap per element is worth a great deal and the sixth is worth very little, while each costs the same. Table 5.11 shows the shape — most of the gain is in the first fifth of q. The design rule follows in the Specification: give every element one or two non-nested memberships and stop.

9.5.11 5.11 The four legitimate reasons to raise N

Given the prohibition, when is raising N correct?

L-1 — Ashby shortfall. The environment presents distinguishable variety the system cannot distinguish. Diagnostic signature: recurring surprises of a *kind*, not just of a magnitude — failures the system has no category for and therefore no procedure for. This is the strongest reason and it is not optional; a system without requisite variety fails at its function regardless of its headroom. Note carefully: this is a capability argument, and it may require accepting a worse headroom position in exchange for being able to do the job at all. That trade is legitimate and must be made explicitly, with the headroom cost computed and stated.

L-2 — Building, not repairing. At the initial design of a system, N is being set for the first time and the concavity of Table 5.2 means the early categories are cheap and valuable. Getting

from one category to four buys 1.3863 nats. Do this at design time, when C is still small and Q is still zero, and never again at this price.

L-3 — The ceiling is binding and h is positive. If η S is running up against η ln N — that is, if u is already near 1 and the system is genuinely extracting near its ceiling — then N is the binding constraint and raising it is the correct move. This condition is rare and must be measured, not assumed. The test: `S / ln N > 0.85`. Below that, u is the binding constraint, not N, and Section 6 governs.

L-4 — The new category is nearly free. Occasionally a genuinely distinct category can be acquired at near-zero marginal C — an existing person who already holds a second practice, an existing space that already serves a second use, an existing artifact that already encodes a second tradition. **Take these every time.** They are the only case where the logarithm's slowness does not matter, because the denominator is zero.

Everything else is category inflation.

9.5.12 5.12 N against cost — the crossover

The cost analysis [Curzi, 2026b] established the crossover calculation; it is reproduced here in the form that governs N decisions, because it is the arithmetic that turns the prohibition from an argument into a number.

Take a coordination cost that grows with communication paths: `C(n) = k · n(n−1)/2`. Compare against η ln n at η = 0.5.

TABLE 5.12 — COST VS CEILING, C(n) = k n(n−1)/2 AGAINST eta ln n, eta = 0.5

n	k=0.002		k=0.005		k=0.01		k=0.02		k=0.05	
	C	net	C	net	C	net	C	net	C	net
4	0.012	+0.681	0.030	+0.663	0.060	+0.633	0.120	+0.573	0.300	+0.393
6	0.030	+0.866	0.075	+0.821	0.150	+0.746	0.300	+0.596	0.750	+0.146
8	0.056	+0.984	0.140	+0.900	0.280	+0.760	0.560	+0.480	1.400	−0.360
10	0.090	+1.061	0.225	+0.926	0.450	+0.701	0.900	+0.251	2.250	−1.099
12	0.132	+1.111	0.330	+0.913	0.660	+0.583	1.320	−0.078	3.300	−2.058
15	0.210	+1.144	0.525	+0.829	1.050	+0.304	2.100	−0.746	5.250	−3.896
20	0.380	+1.118	0.950	+0.548	1.900	−0.402	3.800	−2.302	9.500	−8.002
25	0.600	+1.009	1.500	+0.109	3.000	−1.391	6.000	−4.391	15.000	−13.391
30	0.870	+0.831	2.175	−0.474	4.350	−2.649	8.700	−6.999	21.750	−20.049
40	1.560	+0.284	3.900	−2.056	7.800	−5.956	15.600	−13.756	39.000	−37.156

Every column turns negative. Every one. The only question is where. At k = 0.01 — a modest coordination cost of one hundredth of a nat per communication path — the system goes net-negative at n = 20 and is 5.96 nats underwater at n = 40, against a ceiling of only 1.84.

Note also where each column peaks: k = 0.002 peaks around n = 15; k = 0.005 around n = 10; k = 0.01 around n = 8; k = 0.02 around n = 6; k = 0.05 around n = 4. **The optimal category count is small, and it shrinks fast as coordination gets expensive.** This is the arithmetic behind every piece of practical folk wisdom about small teams, and it does not require any claim about human psychology. It is a logarithm losing a race against a quadratic.

9.5.13 5.13 Specification: the requirements on raising N

RN-1 — Publish the ceiling. Every system's design documentation must state N, ln N, η, and η ln N on its first page. That last number is the total budget from which all maintenance, coordination, and growth is paid, and if the people spending it do not know its value they are spending blind.

RN-2 — Compute h before considering N. `h = η ln N − δQ − C`. If h < 0, the Domain-Addition Prohibition applies and no proposal to raise N may be approved as a remedy. Route to the intervention procedure [Curzi, 2026j]: cut C, raise η, or shed Q.

RN-3 — Run the relabeling audit on every proposed category. Six steps, Section 5.7. A proposal that cannot state the new category's encoder in one sentence is rejected without further review.

RN-4 — Require u >= 0.3 against every existing category. Measured by Section 4's Instrument 1 or 2, on at least 20 discriminating items, under blind elicitation. Report the full vector of measured u's against existing categories, not the mean.

RN-5 — Price ceiling gain and delivered gain separately. `ln(N+1) − ln N` is the ceiling gain and is nearly always trivially small (Table 5.7). The delivered gain is `ln(1 + N·u_bar_new) − ln(1 + (N−1)·u_bar_old)` and it is the number that matters. Documents that quote only the ceiling gain are non-compliant.

RN-6 — Test the binding constraint before acting. Compute `S / ln N`. Above 0.85, N binds and Section 5 applies. Below 0.85, u binds and Section 6 applies. Most systems that believe they have an N problem have a u problem, and this single ratio distinguishes them.

RN-7 — Diagnose Ashby shortfall separately from headroom deficit. Ashby shortfall presents as surprises of a novel *kind*. Headroom deficit presents as steady erosion of a familiar kind. Keep two registers: a surprise log (novel failure kinds, dated) and a headroom trace (h computed each period). They call for opposite actions and must never be conflated.

RN-8 — Specify q, the overlap fraction, as a design parameter. For each element of the organization, count its memberships and the fraction of membership pairs that are non-nested. Target q in the range 0.2 to 0.4 (Table 5.11 shows the gain is concentrated there and the cost accounting [Curzi, 2026b] shows the overhead is not yet severe). Publish q per element. q = 0 across the whole organization means a tree has been built and the return on every category paid for has been forfeited.

RN-9 — Give every person at least one non-nested membership. The minimum viable semilattice. One standing cross-function body, one jointly owned artifact, or one shared space with a group that neither contains nor is contained by their own.

RN-10 — Cap N against the cost crossover. Compute the crossover from Table 5.12's method at the measured k and η. Do not exceed the peak. Publish the computed optimum with the org design.

RN-11 — Prefer splitting the body to growing it. Two bodies of six at u = 0.5 deliver `2 × 1.5041 = 3.0082` nats against `2 × 15 = 30` paths. One body of twelve at the same u delivers 1.8718 against 66 paths. **Splitting raises delivered surplus by 61% and cuts communication paths by 55% simultaneously.** This is the strongest single move available at moderate N and it is almost never taken, because splitting reads as failure and growing reads as success.

RN-12 — Badge every use of the Ashby correspondence. `[FORMAL CORRESPONDENCE, NOT IDENTIFICATION]`. Ashby's variety and the surplus ceiling are both logarithms of a state count; they are not proved equal and the surplus range analysis [Curzi, 2026] declines to identify surplus with information-theoretic quantities.

RN-13 — Badge the nesting postulate. Any argument that assigns u = 0 to a nested category pair carries `[POSTULATE]` and states the test that would refute it: measure u between a category and its parent using Instrument 2; if it exceeds 0.3, the postulate does not hold for that structure.

9.5.14 5.14 Failure modes native to the ceiling

FM-5.1 · THE DOMAIN-ADDITION REFLEX

Damages: C (raises it), h (worsens it), and the credibility of the whole planning process.

Clinical description. The system detects decline. Leadership responds by broadening scope: new departments, new initiatives, new disciplines, new markets. The reasoning is that

a bigger system has more capability and therefore more resilience. The arithmetic says the added ceiling is `η ln(N_new/N_old)` while the added cost is at least linear and typically quadratic in the new category count. **Every domain addition made in response to a deficit deepens the deficit.**

Diagnostic signature. A strategy document that responds to declining performance with an expanded remit. Growth in the number of named initiatives while output per initiative falls. The phrase "we need to be more than just a ___."

Detecting measurement. Compute h before and after each of the last three expansions. If h fell each time, the reflex is present and is the primary pathology.

Fix. RN-2 as a hard gate. Then the intervention ordering [Curzi, 2026j]: raise η first (cheapest), cut C second, shed Q third. Only after h > 0 may N be reconsidered, and then only under one of Section 5.11's four legitimate reasons.

FM-5.2 · CATEGORY INFLATION

Damages: C, and the integrity of the measurement of N.

Clinical description. N in the plan rises while N in the geometry does not. Categories are created by splitting, relabeling, or nominal reassignment. Table 5.7's ceiling gain is quoted; Table 5.7's ceiling gain is 1.63% at N = 20 and the delivered gain is zero.

Diagnostic signature. Growth in the org chart without growth in scored disagreement. New category names that are compound modifications of existing ones ("Applied X", "Strategic X", "X Operations"). Measured u between new and parent category below 0.15.

Detecting measurement. RN-4. Blind elicitation on 20 discriminating items, new category against its parent, computed as `u = 1 − ρ²`.

Fix. The relabeling audit (Section 5.7) applied retroactively. Merge every category pair measuring u < 0.15 back into one. Report the merge as a cost reduction, which is what it is.

FM-5.3 · THE ASHBY–HEADROOM CONFUSION

Damages: the diagnosis, and therefore every action that follows it.

Clinical description. A capability failure (missing variety) and a persistence failure (negative headroom) present superficially alike — both look like "we are not doing well" — and call for opposite actions. Treating a headroom deficit as an Ashby shortfall produces FM-5.1. Treating an Ashby shortfall as a headroom deficit produces a system that is efficiently and sustainably unable to do its job.

Diagnostic signature. For Ashby shortfall: failures of a *novel kind*, incident reports whose category field reads "other," repeated invention of ad-hoc procedure. For headroom deficit: failures of a familiar kind at rising frequency, maintenance backlogs, staff who know exactly what is wrong and cannot get to it.

Detecting measurement. Two registers, kept separately, per RN-7: a surprise log and a headroom trace.

Fix. Diagnose first. If Ashby shortfall: acquire the missing distinction, accept the headroom cost, state it explicitly, and pay for it by cutting elsewhere. If headroom deficit: the intervention procedure [Curzi, 2026j]. If both — which happens — fix headroom first, because an Ashby intervention undertaken at h < 0 will not persist. You cannot store a capability in a vessel that is draining.

FM-5.4 · THE TREE

Damages: u, catastrophically, and therefore the return on all existing N.

Clinical description. The organization has been arranged as a strict hierarchy. Every element has exactly one parent. Every category is a subset of exactly one other. All overlaps

have been resolved into ownership, all dual reporting eliminated, all shared spaces assigned. The chart is clean. Under the nesting postulate, every element's descriptor set audits to `N_eff = 1`.

Diagnostic signature. q = 0 across the organization. No standing cross-function bodies with real authority. Every artifact has exactly one owner. Every meeting has exactly one convening unit. When two units need to work together, a coordinating third unit is created above them — which preserves the tree and adds a level.

Detecting measurement. RN-8. Enumerate memberships per element and compute the non-nested fraction. Also count: how many standing bodies span two branches without being subordinate to a common parent below the root? In a strict tree the answer is zero.

Fix. RN-9, the minimum viable semilattice: one non-nested membership per person. Then targeted overlap at the highest-value boundaries — the rotation programmes of the silo analysis [Curzi, 2026g], shared instrumentation, and boundary objects are all q-raising devices. Do not attempt q > 0.4; Table 5.11 shows the returns saturate and the cost analysis [Curzi, 2026b] shows the costs do not.

FM-5.5 · THE UNMEASURED CEILING

Damages: every planning decision downstream.

Clinical description. Nobody has computed `η ln N`. Budgets, growth targets, and maintenance commitments are set without reference to the total from which they must all be paid. The system routinely commits to `r_t > h_t` and then fails to deliver, and the failure is attributed to execution.

Diagnostic signature. Chronic under-delivery against plan with no pattern in which items fail. Plans that sum commitments without reference to a total. Nobody in the organization can state `η ln N` when asked.

Detecting measurement. Ask three senior people for the value of `η ln N`. If none can answer, the mode is present.

Fix. RN-1. Put the number on the first page of the plan. This costs one afternoon and it is, per unit of effort, the highest-return intervention in this paper.

FM-5.6 · THE SATURATED CATEGORY SET

Damages: C, while delivering nothing to S.

Clinical description. N has grown past the crossover in Table 5.12. Additional categories cost more in coordination than they return in ceiling. The system is net-negative on its most recent several categories and does not know it because it measures ceiling gain rather than delivered gain minus cost.

Diagnostic signature. Coordination effort growing faster than headcount. Meetings whose purpose is to synchronize other meetings. Category owners who cannot name what their category contributes that no other category does.

Detecting measurement. Compute the crossover per RN-10 at measured k and η. Compare against current N.

Fix. RN-11. Split rather than grow, or merge the categories that fail RN-4. Both reduce the quadratic term, and the surplus loss is small because the categories being removed were the ones contributing 0.02–0.05 nats apiece.

9.6 6 Engineering non-parallelism

Section 5 closed a door. The ceiling is `ln N`, it grows brutally slowly, and a headroom deficit cannot be closed by adding domains. That is the negative result and it eliminates the move most organizations reach for first.

Section 6 opens the door that remains. In `S = ln(1 + (N−1)u)`, the parameter u enters **linearly**, and it multiplies the **full** `(N−1)`. Every category already paid for — every domain, every discipline, every hire, every instrument — is sitting there with its contribution scaled by u. Raising u collects on all of them at once, without hiring anyone, without expanding scope, and without adding a single communication path.

u is the only parameter in the law that can be moved substantially in a single quarter, and it is the only one whose improvement collects on capital already bought. δ moves over years. η moves over seasons and requires building things. N is capped by a logarithm and constrained by cost. C can be cut but only to zero, and zero is rarely enough. u can go from 0.12 to 0.50 in eight quarters with three interventions and no new headcount, and Section 6 sets out how, with the arithmetic attached.

It also catalogs, at length, the many ways institutions destroy u without noticing — because in practice most u engineering is not construction. It is the removal of homogenizing machinery that somebody installed for a good reason and nobody has audited since.

9.6.1 6.1 The parameter

```
u(x,y) = 1 − (x . y)^2 = || x ^ y ||^2  = sin^2 θ        in [0,1]

N_eff  = 1 + (N−1) u

S      = f(u) = ln(1 + (N−1) u)      UNIQUELY characterized by S1–S4
```

with **S3, the affine effective count, being the design axiom — chosen, not derived**. The linearity of `N_eff` in u, which is the entire engineering premise of this section, is a consequence of S3. That dependence is stated wherever the section leans on it, which in Section 6 is constant.

The derived properties that govern u engineering [Curzi, 2026]:

Property	Ref	What it means for the engineer
Strict monotonicity	Thm 3.1	Every increase in u strictly increases S. There is no threshold to cross and no plateau to wait for. **Partial progress is real progress.**
Positivity	Thm 3.2	`F > 0` for any `u > 0`. A single genuinely different member produces strictly positive surplus.
Strict concavity	Thm 3.3	Diminishing returns. The first increments of u are the valuable ones.
Sharp Lipschitz	Thm 4.2	`|f(u_1)−f(u_2)| <= (N−1)|u_1−u_2|`. Sensitivity is highest at small u, which is where most systems are.
Extremality	Thm 4.3	`F = 0` iff parallel. **Not small. Zero.**
Range	Curzi [2026], Cor 2.2	`0 <= F <= ln N`.

9.6.2 6.2 The zero

The extremality theorem [Curzi, 2026] says `F = 0` if and only if the states are parallel. This is not a limit or an approximation. At u = 0, `N_eff = 1` and `S = ln 1 = 0` exactly.

The organizational statement of that theorem is the sentence Section 6 is built around:

> **A second opinion from someone whose training and incentives replicate the first is worth ln(1) = 0.**

Not "worth little." Not "of limited value." Worth zero. The second reviewer, the confirming analysis, the sanity read from the colleague down the hall — if the instrument is the same instrument, the reading is the same reading, and the law prices a duplicate at nothing.

TABLE 6.1 — WHAT A SECOND OPINION IS WORTH, AS A FUNCTION OF ERROR CORRELATION
(N = 2; u = 1 − rho^2; S = ln(1 + u); ceiling ln 2 = 0.6931)

rho	u	N_eff	S = ln(1+u)	share of ln 2
1.00	0.0000	1.0000	0.0000	0.0%
0.98	0.0396	1.0396	0.0388	5.6%
0.95	0.0975	1.0975	0.0930	13.4%
0.90	0.1900	1.1900	0.1740	25.1%
0.80	0.3600	1.3600	0.3075	44.4%
0.60	0.6400	1.6400	0.4947	71.4%
0.40	0.8400	1.8400	0.6098	88.0%
0.20	0.9600	1.9600	0.6729	97.1%
0.00	1.0000	2.0000	0.6931	100.0%

Read the top of the table, because the top is where organizations live. Two reviewers at ρ = 0.95 — typical of two people trained in the same programme, reading the same literature, evaluated by the same committee — deliver 13.4% of an independent pair. At ρ = 0.98, which describes a mentor and their own former student, they deliver 5.6%.

The organization believes it bought a second opinion. It bought 5.6% of one, at full price.

And now the other end of the table, which is not printed above and must be. Table 6.1 runs `rho` from 1.00 down to 0.00 and stops, which invites the reading that the good direction is simply *away from agreement*, and that the best available reviewer is the one who disagrees hardest. That reading is wrong, and the correction is one line of the definition: `u = 1 − rho^2` depends only on `|rho|`. The table is symmetric about zero.

```
rho = −0.95   ->   u = 0.0975   S = 0.0930    13.4%   -- identical to rho = +0.95
rho = −0.98   ->   u = 0.0396   S = 0.0388     5.6%   -- identical to rho = +0.98
rho = −1.00   ->   u = 0.0000   S = 0.0000     0.0%   -- identical to rho = +1.00
```

At perfect opposition the surplus is exactly zero, precisely as it is at perfect agreement. The zero of the extremality theorem [Curzi, 2026] is a zero of *parallelism*, and a state vector and its negation are parallel. Given the sycophant's reading the principal's can be reconstructed; given the perfect contrarian's it can be reconstructed just as exactly, by changing a sign. **The mirror and the mirror-image are the same instrument. The contrarian who inverts a colleague carries as much information as the sycophant who echoes them: none.**

This correction governs the rest of Section 6, and method M-7 and the red-team specification in particular. **The target of every intervention below is `|rho|` near zero — genuine independence — and never `rho` near -1.** A body engineered for "creative tension," in which a designated opponent takes the negation of whatever is proposed, has manufactured a `u` of zero and is paying `C` for the extra seat: it returns no surplus and costs more than the aligned pair it replaced. Section 4.2 gives the geometry — the angle table folded about 90 degrees — and the pseudo-diversity analysis [Curzi, 2026d] catalogues the fault in its own right, distinct from the Parallel Wall.

9.6.3 6.3 The honesty clauses, restated for the non-parallelism parameter

1. **The founding surplus analysis [Curzi, 2026] is pure mathematics and makes no empirical or metaphysical claim.** That `F = 0` at u = 0 is a theorem about a functional form. That a review panel's second reader is worth nothing is a modeling inference, and it depends on the encoder chosen and on the Roster Reduction [POSTULATE] of Section 3.4.
2. **S3 is a design axiom.** The linearity of `N_eff` in u is chosen.
3. **The deployed functional is a lower bound.** A measured u of 0.15 licenses "at least this much"; it does not license "exactly this much" or "no more than this."
4. **The identity `u = 1 − ρ²` is exact** for mean-centred, unit-normalized state vectors, since their inner product is the Pearson correlation. Encoding a member's state as their error vector is a modeling assumption, badged **[POSTULATE]**.
5. **Subadditivity holds on u only; the triangle inequality fails for geometric triples** (the founding surplus analysis [Curzi, 2026] §5). Where Section 6 combines interventions, it combines them on u and does not chain surplus across triples.
6. **Surplus is not Shannon entropy or mutual information** [Curzi, 2026], and **no thermodynamic negentropy law is proved** [Curzi, 2026a].

9.6.4 6.4 The lever ranking, and why u is second

The intervention procedure [Curzi, 2026j] gives the full ordering. The short version, with this section's justification:

Lever	Typical time to move	Typical cost	Collects on	Verdict
η	one to three quarters	low — writing, archiving, apprenticeship	all future S	**first move in almost every system**
u	one to four quarters	low to moderate — mostly removal of existing machinery	all existing N	**fastest-moving parameter; this section**
C	one to four quarters	moderate — political, not technical	the whole budget	third; bounded below by zero
Q (shedding)	one to eight quarters	political	frees δQ	underused; see the headroom analysis [Curzi, 2026b]
δ	years	high	future retention	slow; the retention analysis [Curzi, 2026a]
N	years	very high	logarithmic return	**last; Section 5's prohibition**

η is first because it is usually cheapest and because it multiplies everything. But η requires **building** — an archive, a protocol, a training path — and building takes a quarter before it starts converting. u is the parameter that moves fastest, because most u interventions are subtractions rather than additions, and a subtraction takes effect the day it is made.

Cancel the meeting that produces the cascade and u rises the next week. Stop screening for culture fit and u rises with the next hire. Remove the shared metric that gives every unit

the same objective function and u rises the next quarter. These are not construction projects. They are the removal of machinery.

9.6.5 6.5 What u buys, at a working scale

Before methods, the payoff table. A system with N = 10, η = 0.5, δ = 0.08 per quarter, C = 0.30. Each level of u, with the steady state it implies through the fixed-point theorem [Curzi, 2026a], `Q_\ensuremath{\infty} = (ηF(u∗) − c∗)/δ`.

```
TABLE 6.2 — WHAT u BUYS, N = 10, eta = 0.5, delta = 0.08, C = 0.30

      u      N_eff      S = ln N_eff    eta S       Q_inf = (eta S − C)/delta
   ─────   ───────    ─────────────   ───────     ─────────────────────────
    0.05    1.45          0.3716        0.1858            −1.428   (unsustainable)
    0.10    1.90          0.6419        0.3209             0.262
    0.15    2.35          0.8544        0.4272             1.590
    0.20    2.80          1.0296        0.5148             2.685
    0.30    3.70          1.3083        0.6542             4.427
    0.40    4.60          1.5261        0.7630             5.788
    0.50    5.50          1.7047        0.8524             6.905
    0.60    6.40          1.8563        0.9281             7.852
    0.70    7.30          1.9879        0.9939             8.674
    0.80    8.20          2.1041        1.0521             9.401
    0.90    9.10          2.2083        1.1041            10.052
    1.00   10.00          2.3026        1.1513            10.641
```

Three readings.

First, the bottom of the table is a death sentence. At u = 0.05 the steady state is negative, which by the fixed-point theorem [Curzi, 2026a] means the equilibrium is unsustainable and the linear model must be replaced. The theorem does not hide this failure mode and neither should the analyst. A ten-category system at u = 0.05 does not converge to a little stored order; it converges to a contradiction, which in practice means it dissolves, is rescued by subsidy, or was never described by these parameters.

Second, the returns are front-loaded. Moving u from 0.05 to 0.30 — a year's work for most organizations — swings the steady state by 5.855. Moving from 0.70 to 1.00, which would require near-perfect orthogonality, swings it by 1.967. The first quarter of the range is worth three times the last. That is concavity, and it is a theorem.

Third — and this is the whole engineering argument — none of these rows required changing N, η, δ, or C. Same categories, same conversion, same leak, same cost. The entire column of steady states was produced by moving one parameter that costs nothing except the willingness to stop homogenizing.

9.6.6 6.6 Raising u — eight methods, with what each costs and what it buys

Ordered by speed of effect. Each entry gives the mechanism, the expected effect on u, the cost, the failure mode, and the measurement that confirms it worked.

M-1 · Silent-write-first

Mechanism. Before any discussion of a question, every participant writes their independent answer and estimate, seals it, and submits it. Only then does discussion begin, and the sealed answers are read into the record first.

Effect on u. It does not raise the geometry; it stops the measured u from being destroyed by the cascade, which in practice is the same thing, since a collapsed discussion delivers the

surplus of one member whatever the geometry. In bodies that have never done it, the gap between measured-under-cascade u and measured-blind u routinely exceeds 0.15 (Section 4, FM-4.5).

Cost. Three minutes per question. Zero capital. Zero headcount.

Failure mode. Performing it and then discarding it — writing the sealed answers and then proceeding as though the senior person's spoken opinion is the starting point. The sealed answers must be read into the record *before* discussion, in a randomized or reverse-seniority order, or the ritual is theatre.

Confirming measurement. Instrument 1 on the sealed answers and on the post-discussion positions; the gap is the cascade magnitude, tracked as a standing metric. This is the highest return-per-minute intervention in the framework, and adopting it requires no authority beyond the ability to say "everyone write it down first."

M-2 · Cut the shared incentive

Mechanism. Identify the metric, target, or reward that every category is optimized against. Split it, stagger it, or replace it with category-specific objectives that are not derivable from one another.

Effect on u. Often the largest structural effect available. A common objective function is a **common mode**: it correlates the errors of otherwise independent instruments, exactly as a shared power supply correlates the noise of independent circuits. In Section 4's worked example, cross-lineage correlation ran at $\rho = 0.55$ among people from genuinely different traditions, the residual attributable to a shared funder with shared criteria; reducing it to $\rho = 0.40$ raised delivered surplus by 0.1550 nats with zero personnel change.

Cost. Political, sometimes severe. A single shared metric is administratively convenient and its removal is resisted by everyone whose job is administration.

Failure mode. Replacing one shared metric with another under a different name. Test it: if unit A's objective is a monotone function of unit B's, nothing changed.

Confirming measurement. Error correlation (Instrument 2) computed before and after, on scored outcomes. Expect the effect to take two to four periods to appear in the data, because the old incentive's influence on behaviour persists past its removal.

Related failure, named here because it is the same mechanism running the other way: Goodhart's Law and Campbell's Law. A measure that becomes a target ceases to be a measure. In this vocabulary: a shared target drives every category to the same point in state space, so u falls, so the measure loses the very diversity of view that made it informative. Goodhart capture is a u-destruction mechanism, and that is the sharpest available statement of why it happens.

M-3 · The adversarial role

Mechanism. Assign, formally and rotationally, a member whose function for this decision is to build the strongest available case against the emerging consensus. The role is structural, not dispositional: it rotates, it is compensated, and its occupant is not held responsible for the position they argued.

Effect on u. Moderate to large, and immediate. It manufactures a non-parallel state where the geometry did not supply one — the cheapest way to obtain a nonzero u in a body that is genuinely homogeneous and cannot be reconstituted this quarter.

Cost. One person's time per decision, plus the meeting time to hear the case.

Failure mode. Two. **The role becomes a person** — one individual always plays it, acquires a reputation as difficult, and is discounted; rotation is what makes it structural rather than personal. Or **the role becomes a ritual** — the objection heard, thanked, minuted, ignored; require the case be answered in writing before the decision is recorded, and let an unanswered objection block it.

Confirming measurement. The fraction of decisions the adversarial case changed, even marginally. Zero over twenty decisions means the role is theatre.

Honest limitation. A manufactured u is not a measured u: the role produces an argument the body would not otherwise have heard, but the arguer's underlying state is still correlated with the body's. Badge it **[PARTIAL SUBSTITUTE]** and do not record it in a structural u estimate.

M-4 · The deliberate outsider

Mechanism. Recruit one member from a training lineage with no overlap with any existing member, place them in a role with real decision rights, and protect them from assimilation for a defined period.

Effect on u. Large, per Section 4's marginal table. The first outsider in a homogeneous body is worth most, because at low u the Lipschitz slope is steep: at N = 12, `df/du` at u = 0.1 is 4.74, so a 0.1 rise in u buys roughly 0.47 nats.

Cost. One hire, or one replacement. Real but bounded.

Failure mode — and it is the dominant one — assimilation. The outsider is hired, warmly welcomed, onboarded into the shared vocabulary, trained on the shared literature, evaluated against the shared metric, and within four quarters their error vector is indistinguishable from everyone else's. The organization concludes that hiring outsiders does not work, when what happened is that it converted the outsider into an insider and destroyed the asset it bought. Section 6.9 gives the protection protocol.

Confirming measurement. The outsider's error correlation against the body's mean, quarterly. A rising ρ is assimilation in progress, visible long before anyone notices it socially.

M-5 · Rotation

Mechanism. Move people between categories on a fixed cadence — six to twenty-four months by the depth of the practice — with a required overlap period and a written handover.

Effect on u. Moderate, cumulative, durable. Rotation attacks u decay at its source: it stops an individual's state converging on their category's, and it manufactures the non-nested memberships Section 5.10 identified as semilattice structure. Rotation is simultaneously a u intervention and a q intervention, which is why it appears in both sections.

Cost. Real. Rotation costs productivity during transitions and depth permanently — a rotated practitioner is less expert in any one category than a fixed one. A genuine trade-off.

Failure mode. Rotating within a branch of the tree, which is nesting, not overlap, and buys nothing. Rotate **across** branches or do not rotate.

Confirming measurement. Pairwise u between categories, measured before rotation begins and after two full cycles. Also track the q statistic of Section 5 RN-8.

M-6 · Import a practitioner from an unrelated trade

Mechanism. Bring in someone whose expertise is real, deep, and from a field with no institutional contact with the receiving one: a logistician into a hospital, a stage manager into a research group, a bookbinder into a software archive project.

Effect on u. Potentially the largest single-person effect available, because lineage overlap is genuinely zero. From Table 4.2, an 80° difference gives u = 0.9698 against 0.0670 at 15° — a factor of fourteen.

Cost. Moderate to high in translation overhead: neither party can use the other's vocabulary. The translator role and boundary objects of the silo analysis [Curzi, 2026g] exist precisely for this.

Failure mode. Placing the outsider in an advisory position with no decision rights, where their non-parallel state has no channel into any decision. A state that cannot enter a decision

contributes zero surplus to that decision regardless of its geometry. Advisory boards are the standard institutional method for acquiring u and simultaneously ensuring it is never spent.

Confirming measurement. Count decisions in which the imported practitioner's position appears in the decision record. Zero means the import is decorative.

M-7 · The red team

Mechanism. A standing group, separately staffed and separately reporting, whose function is to attack the main body's conclusions — a permanent structure with its own membership, budget, and information access, rather than M-3's rotating role.

Effect on u. Large where it is real, which is rare. A red team reporting to the body it attacks is not a red team: the reporting line is a shared incentive (M-2) and will drive its error vector toward the main body's within a year.

Cost. High. It is a duplicate structure.

Failure mode. Capture by three routes — shared reporting line, shared recruitment pool, shared social environment. All three must be broken. A red team recruited from the main body's alumni is a mirror with a different letterhead.

Confirming measurement. Error correlation against the main body, quarterly. Below $\rho = 0.5$ the red team functions; above $\rho = 0.75$ it is captured and should be disbanded and rebuilt, because a captured red team is worse than none — it produces the appearance of adversarial review, which suppresses demand for the real thing.

M-8 · Recruit across training lineages, permanently

Mechanism. Make lineage diversity a standing constraint on the pipeline, not a preference: cap the fraction of the body from any one training lineage, and require the lineage census (Section 4 Instrument 5, EC-9) run and published before any requisition is approved.

Effect on u. The largest **structural** effect available, and the slowest to arrive. Section 4's worked example computed the full ladder: three lineages at 5/4/3 gave u_bar = 0.5137; four lineages at 3/3/3/3 gave 0.5814; six lineages at 2×6 gave 0.6395; twelve distinct lineages gave 0.6975. Note the shape — moving from three lineages to twelve, the theoretical maximum, bought 0.2654 nats, while moving from one lineage to three would have bought far more. Concavity again. Get to three or four lineages and then stop optimizing lineage and start on incentives.

Cost. Recruitment friction. The pipeline is slower and candidates harder to evaluate, precisely because the evaluators are calibrated on the existing lineages — a homogeneous body's hiring process cannot correctly evaluate a candidate from outside its lineage, because evaluation is itself an instrument and the body has only one. The pseudo-diversity analysis [Curzi, 2026d] develops this.

Failure mode. Achieving lineage diversity and then homogenizing downstream through shared training, metrics, and review (FM-6.2). Section 4's FM-4.4 gives the diagnostic: a gap between the u the census predicts and the u error decorrelation measures means the force is downstream of hiring.

Confirming measurement. Both instruments, compared. Publish both.

9.6.7 6.7 Hong and Page: why this is not a fairness argument

Hong & Page (PNAS 2004) proved a result that is routinely misquoted and is worth stating precisely: under stated conditions, a randomly selected group of problem solvers can outperform a group composed of the best individual performers. The mechanism is that high-ability solvers tend to share heuristics — they are good in similar ways — and therefore get stuck in the same local optima, while a more heuristically diverse group has more escape routes.

The result is a theorem under explicit assumptions about the problem space and the solver representations. It is not an empirical claim about all groups and it does not say ability is unimportant. It is a theorem with conditions, and the conditions matter.

What it supplies here is the answer to the objection raised against every intervention in Section 6.6: *"but the people we are excluding are better."* The answer is arithmetic. Take two candidates for a twelfth seat on a body at u_bar = 0.15, N = 12.

- Candidate A is measurably stronger individually and comes from the body's dominant lineage. Their error correlation with the body is ρ = 0.95, so their pairwise u against the body is `1 − 0.95² = 0.0975`.
- Candidate B is measurably weaker individually and comes from an absent lineage. Their error correlation with the body is ρ = 0.55, so their pairwise u is `1 − 0.55² = 0.6975`.

The body has 11 existing members, so the twelfth seat contributes 11 new pairs out of 66. Under the Roster Reduction:

```
with A:  u_bar = (55 x 0.15 + 11 x 0.0975)/66 = (8.250 + 1.073)/66 = 0.14126
         N_eff = 1 + 11 x 0.14126 = 2.554     S = 0.9377

with B:  u_bar = (55 x 0.15 + 11 x 0.6975)/66 = (8.250 + 7.673)/66 = 0.24126
         N_eff = 1 + 11 x 0.24126 = 3.654     S = 1.2960
```

Candidate B delivers 0.3583 more nats than candidate A — a 38% increase in the body's entire surplus input — and candidate A, the stronger individual, lowers the body's mean non-parallelism below its previous level. Hiring A makes the body worse in the only currency the law recognizes.

Say this plainly, because it is the sentence that will get the framework argued with: **the argument for non-parallelism is not about fairness, representation, or morale. It is about the arithmetic of a logarithm, and it would hold identically for sensors, models, or measuring instruments with no people involved.** That it also happens to support treating people decently is a convenience, not a premise; the individuality analysis [Curzi, 2026h] takes up the ethics on its own terms.

9.6.8 6.8 How institutions destroy u without noticing

Almost every homogenizing mechanism below was installed deliberately, for a defensible reason, by someone competent. None was installed in order to destroy u. All of them do.

The unifying diagnosis: **u is not a property that is acquired once. It is a quantity that decays and must be maintained.** Unmaintained, it goes to zero at its own rate. That makes u a δ problem as much as a design problem, which is why the pseudo-diversity analysis [Curzi, 2026d] treats its maintenance as a first-class activity.

TABLE 6.3 — u DECAY. u_t = u_0 (1 − d_u)^t, starting from u_0 = 0.60

t	d_u=0.02	d_u=0.05	d_u=0.10	d_u=0.20
0	0.6000	0.6000	0.6000	0.6000
1	0.5880	0.5700	0.5400	0.4800
2	0.5762	0.5415	0.4860	0.3840
3	0.5647	0.5144	0.4374	0.3072
6	0.5315	0.4411	0.3189	0.1573
12	0.4708	0.3242	0.1695	0.0412
24	0.3695	0.1752	0.0479	0.0028
36	0.2899	0.0947	0.0135	0.0002
60	0.1785	0.0276	0.0011	0.0000

```
TABLE 6.4 — HALF-LIFE OF u UNDER CONVERGENCE RATE d_u
             t_half = ln 2 / ( -ln(1 - d_u) )

     d_u = 0.01     half-life = 68.97 periods
     d_u = 0.02     half-life = 34.31 periods
     d_u = 0.05     half-life = 13.51 periods
     d_u = 0.10     half-life =  6.58 periods
     d_u = 0.15     half-life =  4.27 periods
     d_u = 0.20     half-life =  3.11 periods
```

At a convergence rate of 10% per quarter — which is modest for a co-located team with a shared metric and a weekly all-hands — a body's non-parallelism halves in 6.58 quarters and is down to 5% of its initial value in five years. No one notices, because convergence feels like the team getting better at working together. It is, in fact, the team becoming one instrument.

Now the mechanisms.

D-1 · Hiring for culture fit

The most efficient u-destroyer ever devised, and the one most likely to be described as a strength. "Culture fit" is operationally a screen for state similarity to the existing body: a correlation filter applied at entry, by evaluators who are themselves the reference vector.

Signature. Interview loops that include a "would I want to work with them" question. Candidates rejected for reasons that resolve, when written down, to unfamiliarity. High agreement among interviewers about candidates — which is itself a u = 0 measurement on the interview panel and should be read as such.

Correct replacement. Screen for **values fit** and **capability**, which are not similarity, and explicitly *against* lineage duplication; state in the requisition which lineages are over-represented. Where a "fit" judgement survives, require it in writing as a specific behavioural prediction — falsifiable — rather than as a feeling, which is a correlation reading.

D-2 · Promotion by similarity

The same filter at every subsequent level, compounding. If each promotion round selects for similarity to the level above with even modest strength, after four levels the leadership is a near-pure sample from one lineage however diverse the entry cohort was. A diverse intake with a similarity-selecting promotion ladder produces a diverse workforce and a parallel leadership, which is the observed state of a great many institutions.

Signature. Lineage entropy falling monotonically with seniority. Compute the entropy of the lineage distribution at each level and plot against level; a downward slope is the diagnosis.

Correct replacement. Promotion criteria stated as observable outputs, not judgements of potential — potential is the channel through which similarity operates. Cross-level review panels drawn from other branches. Publish lineage entropy by level annually.

D-3 · Shared training

Everyone attends the same onboarding, the same courses, the same annual conference, and reads the same canonical texts. Done for excellent reasons — shared vocabulary lowers C substantially — and the cost is paid in u. A genuine trade-off, not a mistake; the engineering question is only where to put the boundary.

The rule this paper recommends: share the interface, not the model. Standardize the format in which claims, decisions, and artifacts are exchanged. Do not standardize the reasoning that produces them. A common file format costs nothing in u. A common analytical framework costs almost all of it.

Signature. Every member can complete every other member's sentence. New hires describe onboarding as "learning how we think here."

D-4 · Common incentives

Covered as M-2 from the intervention side; from the destruction side, note how quietly it operates. A single organization-wide bonus metric, a single OKR cascade, a single set of promotion criteria — each is a common-mode input to every member's objective function. The members remain different people with different training and their **errors converge anyway**, because they are now all solving the same optimization problem.

Signature. The gap FM-4.4 identifies: high lineage entropy, low measured u. If provenance says diverse and error decorrelation says parallel, look at the incentive first.

D-5 · Co-location without cross-assignment

Putting people in one building raises contact rate. Contact without cross-assignment raises convergence without raising the overlap fraction q — a body that talks constantly and thinks identically, the maximum-convergence, minimum-structure configuration.

Signature. High interaction frequency, low measured u, and an org chart that is a strict tree (FM-5.4). Open-plan offices with pure functional reporting are the canonical instance.

Correct replacement. Cross-assignment first, co-location second. Co-location amplifies whatever structure is present; if the structure is a tree, it accelerates the convergence. The room-as-circuit analysis [Curzi, 2026d] gives the spatial version and the threshold analysis [Curzi, 2026f] the threshold design that preserves separation where separation is load-bearing.

D-6 · Information cascades

Bikhchandani, Hirshleifer & Welch (1992) showed that rational agents observing predecessors' actions can rationally abandon their own private information, at which point the sequence stops aggregating entirely — every later actor is copying, so no new information enters the record however many participate. In the present vocabulary: **a cascade drives the observed u to zero regardless of the underlying geometry.** The body may hold twelve genuinely different instruments, and if the eleventh and twelfth speak after hearing the first ten, the reading is the first one's.

Signature. Strong dependence of stated position on speaking order. Rapid convergence in discussion. Unanimity arrived at quickly on questions that are genuinely hard.

Correct replacement. M-1, silent-write-first, an anti-cascade device by construction; Chapter 17's furniture and the protocols of the individuality analysis [Curzi, 2026h] institutionalise it. A cascade is the cheapest u-destroyer to prevent and the most expensive to detect afterward, because it leaves a record of agreement.

D-7 · Preferential attachment and the homogenizing hub

Barabási & Albert (Science 1999) showed that growth with preferential attachment produces scale-free degree distributions: a few nodes acquire very high degree. **Albert, Jeong & Barabási (Nature 2000)** showed such networks are robust to random failure and fragile to targeted attack on hubs.

The u consequence is a third property, specific to this paper's concern. A hub is connected to a large fraction of the network, and every node connected to it receives its input. Over time, the hub's state becomes a common-mode input to much of the body, and common-mode input correlates errors. The hub need not be coercive or even intentionally influential. It only needs to be listened to by many people.

Signature. One person, document, forum, or channel upstream of most positions. Ask twenty members where they formed their view on a contested question; if more than half trace to one source, that source is a homogenizing hub.

Correct replacement. Not removal — hubs carry real efficiency and robustness benefits. Install a second and third hub with genuinely different states, and ensure no member draws from only one. Two hubs at ρ = 0.4 with members drawing from both is a substantially different geometry from one, and it costs far less than dismantling the network. Note the interaction with Section 5: preferential attachment grows something close to a tree with a few high-degree centres, and Section 5.10 showed that tree structure drives a descriptor set to `N_eff = 1`. **Hub formation and tree formation are the same homogenization arriving by two routes.**

D-8 · Cohesion

Cohesion is the most socially rewarded u-destroyer. A cohesive team is pleasant, low-friction, low-C, and it converges. The collective-intelligence analysis [Curzi, 2026g] puts it as **cohesion is often the enemy of the collective intelligence factor c, because cohesion drives u toward 0**; the arithmetic version is that every mechanism making a group more comfortable with itself raises the convergence rate d_u unless specifically counteracted.

Signature. Rising satisfaction scores alongside falling disagreement rates. Meetings that end early. Departures described as "not a fit."

Correct replacement. Not to make the group unpleasant, which raises C without raising u, but to install counter-mechanisms — M-1, M-3, M-5 — that hold u against the cohesion gradient. Comfort and non-parallelism are not opposites, but neither are they automatically compatible; where they conflict, this paper chooses the arrangement that preserves both by design.

D-9 · The single review path

Every output passes the same review, by the same reviewers, against the same rubric. A filter applied to all outputs makes all outputs more alike; over enough cycles the body learns to produce what passes, and what passes is one thing.

Signature. Convergence of output form over time. New members' work becoming indistinguishable from old members' work within a year. A rubric that has not changed in five years.

Correct replacement. Two independent review paths with genuinely different rubrics, or a panel whose composition rotates across lineages. And sunset the rubric: one unchanged in five years is stored order consuming its own δQ maintenance while suppressing u — a case for the shed of the headroom analysis [Curzi, 2026b].

D-10 · Success

The last and the hardest. A method that worked becomes the method. Everyone adopts it, because it worked. u falls, because it worked. **This is the u-side of the headroom observation [Curzi, 2026b] that success consumes its own headroom.** There, Q grew until δQ ate the budget; here, a successful approach propagates until it is the only approach. The two are the same event viewed from two parameters, and an institution in this condition has a rising Q and a falling u simultaneously, which is the fastest available route to h < 0.

Signature. One large success, then declining performance everyone attributes to changed external conditions. Universal endorsement of a specific method. Nobody able to describe what that method's failure conditions would look like.

Correct replacement. Reserve capacity for methods that are not the method — an explicit allocation, protected from the metric the successful method optimizes. It is expensive, and

it is The intervention most likely to be cut in a downturn, which is exactly when it is most needed.

9.6.9 6.9 Protecting u after it has been bought

The most wasteful pattern in this subject: an organization spends a year and real money acquiring non-parallelism, then destroys it in four quarters through D-1 to D-10 without connecting the two events. Table 6.3 gives the arithmetic — from a hard-won u_0 = 0.60 at d_u = 0.10, the body is at 0.319 after six quarters and 0.170 after twelve. Translated to steady state:

```
TABLE 6.5 — REGRESSION AFTER THE CAMPAIGN STOPS
            starting u = 0.57; N = 10, eta = 0.5, delta = 0.08, C = 0.30
            Q_inf is the steady state the CURRENT u would sustain

    quarters    d_u = 0.05          d_u = 0.10          d_u = 0.15
    since       u      Q_inf        u      Q_inf        u      Q_inf
    ---------  ------  -------     ------  -------     ------  -------
        0      0.5700   7.582      0.5700   7.582      0.5700   7.582
        2      0.5144   7.050      0.4617   6.500      0.4118   5.931
        4      0.4643   6.528      0.3740   5.461      0.2975   4.390
        6      0.4190   6.016      0.3029   4.471      0.2150   2.979
        8      0.3781   5.515      0.2454   3.536      0.1553   1.716
       10      0.3413   5.025      0.1987   2.660      0.1122   0.613
       12      0.3080   4.548      0.1610   1.848      0.0811  -0.325
```

At a convergence rate of 15% per quarter, an organization that reaches u = 0.57 and then stops maintaining it is back to an unsustainable steady state within three years. Not back to where it started — *below* where it started, because Q has grown in the meantime and δQ is now larger.

The protection protocol is the standing half of this section's Specification: EU-2, EU-6, EU-7, EU-9, and EU-10. Two points belong here rather than there.

Measure quarterly, never annually. The half-life at d_u = 0.10 is 6.58 quarters; an annual measurement samples a decaying signal at half its half-life and misses the trend for two years.

And when u falls, look for the mechanism, not the people. Every entry in Section 6.8 is structural. None requires anyone to behave badly and none is fixed by asking people to be more open-minded. **Exhortation is not an intervention on u; it is a request that people compensate personally for a machine that is running.** the silo analysis [Curzi, 2026g] makes the general form of this argument.

9.6.10 6.10 Measuring u — the standing protocol

Section 4.3 gave the five instruments. The cadence that turns them into a standing measurement:

Quarterly. (1) *Blind item set* — twenty discriminating items from the body's actual work, fresh each quarter, simultaneous and sealed; discard unanimous items (EC-6), compute pairwise disagreement, convert to u, publish the matrix and the mean. (2) *Cascade gap* — readminister five after discussion; blind u minus post-discussion u is the cascade magnitude, and its trend is the health of the elicitation process. (3) *New-member correlation* per EU-7.

Annually. (4) *Error decorrelation on scored outcomes* (Instrument 2) — the ground truth, and the one that catches a body that argues about everything and is wrong in the same direction. (5) *Lineage census* (EC-9), entropy per provenance axis and per seniority level, which is D-2's diagnostic. (6) *Homogenizer inventory* per EU-9 — the single document that makes u destruction visible.

Reporting. Per EU-2, and always with an interval below u = 0.3 (EC-4): the Lipschitz slope makes point estimates misleading there.

9.6.11 6.11 Specification: the requirements on engineering u

EU-1 — Silent-write-first on every consequential decision. Independent written positions, sealed, submitted before discussion and read into the record in randomized or reverse-seniority order. No exceptions for time pressure; the procedure costs three minutes.

EU-2 — Publish u quarterly with its interval. Format: `N = 10, u = 0.34 ± 0.06, N_eff = 4.06 [3.52, 4.60], S = 1.401, ceiling ln N = 2.303, deficit = 0.902`.

EU-3 — No shared objective function across categories. Each category's objective must be non-derivable from every other's. Test annually: if one category's objective is a monotone function of another's, they are one category for surplus purposes and must be recorded as such in N.

EU-4 — Rotate the adversarial role; never let it become a person. Rotation on a published schedule. The adversarial case must be answered in writing before the decision is recorded. An unanswered case blocks the decision.

EU-5 — Cap any single training lineage at one third of the body. Measured by the lineage census. A body exceeding the cap may not hire from the over-represented lineage until it is back in compliance. State the cap in the requisition.

EU-6 — Onboard to the interface, not to the model. Onboarding materials must be auditable against this distinction. Any module that teaches "how we think about X" is a homogenizer and goes in the register (EU-9) with an expiry date.

EU-7 — Protect new outsiders for eight quarters. Output-based evaluation, reviewed by someone outside the dominant lineage, with monthly error-correlation tracking. A ρ rising above 0.05 per quarter triggers a review of what is assimilating them. Onboard them to the interface, not the model, and do not evaluate them against insider criteria.

EU-8 — Give every member at least two upstream sources with measured ρ <= 0.5 between them. The anti-hub requirement (D-7). Audit by asking members to name where they formed their view on three contested questions.

EU-9 — Maintain the homogenizer register. Every shared rubric, metric, reading list, onboarding module, review path, and canonical forum, with install date, last re-justification, expiry, and the measured u across its population. No entry may exceed a three-year expiry unre-justified.

EU-10 — Budget u maintenance as an operating line. Not a project, not an initiative, not a campaign. A standing line covering the quarterly measurement, the rotation cadence, and the adversarial-role time.

EU-11 — Two independent review paths, or a rotating review panel. No output class may pass through a single fixed review path for more than three years.

EU-12 — Reserve capacity against the successful method. An explicit allocation, protected from the metric that method optimizes, for approaches that are not the current one. State it as a fraction and defend it in downturns specifically.

EU-13 — Badge manufactured u. Adversarial roles, red teams, and devil's-advocate protocols produce arguments, not geometry. Record them as `[PARTIAL SUBSTITUTE]` and do not enter them into a structural u estimate.

EU-14 — Test the binding constraint before spending on u. Compute `S / ln N` (Section 5, RN-6). Above 0.85, N binds and this section's interventions will not help. Below 0.85, u binds. Also compute h: **if h < 0, u interventions cannot rescue the system, because S is already capped at `ln N` and the deficit is in δQ or C.** Fix headroom first.

EU-15 — Never remove a person's ability to decline. Every intervention in Section 6.6 works by making a different position available, audible, and consequential. None may compel a person to hold a view, argue a position they reject, or remain in a role they wish to leave. The adversarial role is voluntary and compensated; rotation is opt-out-able with a stated reason.

A method that works by removing resistance has no coercive failure mode; its worst outcome is that someone declines. The framework's own safety analysis [Curzi, 2026h] develops this as a safety property of the whole approach.

9.6.12 6.12 Failure modes native to non-parallelism

FM-6.1 · THE MIRROR SECOND OPINION

Damages: u, and the credibility of every verification process built on it.

Clinical description. Verification is performed by a second party whose training, information, and incentives replicate the first's, and recorded as independent. Table 6.1 prices two reviewers at ρ = 0.95 at 13.4% of an independent pair, and at ρ = 1.00 at exactly zero.

Diagnostic signature. Second opinions agreeing above 95% of the time. Verification steps that have never once overturned an original finding. Reviewers drawn from the pool they review.

Detecting measurement. Error correlation between first and second party across scored outcomes. Anything above ρ = 0.9 means the verification is delivering under 26% of one independent check.

Fix. Draw the second party from a different lineage (M-8), give them different information access, and ensure they are not evaluated by the same person. If none is possible, record the verification as a consistency check, not an independent confirmation, and stop counting it as risk reduction.

FM-6.2 · ASSIMILATION OF THE OUTSIDER

Damages: u, on a delay of four to eight quarters after the acquisition.

Clinical description. The organization recruits genuinely non-parallel members, then converges their state onto the body's through onboarding, shared metrics, shared review, and social integration. The asset is destroyed by the machinery installed to make new members effective, and the organization concludes that hiring for difference does not work.

Diagnostic signature. New members' error correlation rising steadily through their first two years. Onboarding praised as "getting up to speed on how we do things." Outsiders who become the most enthusiastic exponents of the local view within a year.

Detecting measurement. EU-7's monthly tracking. Plot ρ against tenure for every member hired in the last three years; a rising curve is the diagnosis and its slope is the assimilation rate.

Fix. EU-6, EU-7, EU-9. The fix is not to onboard less; it is to onboard differently, and the distinction is entirely between teaching where things are and teaching how to think.

FM-6.3 · THE CAPTURED RED TEAM

Damages: u, plus the demand for real adversarial review.

Clinical description. An adversarial structure exists, is staffed and funded, and has converged on the body it was built to challenge — through shared reporting, shared recruitment, or shared social environment. It produces adversarial-shaped output confirming the body's position, which is worse than having no red team, because the structure suppresses demand for the function.

Diagnostic signature. Findings that have never changed a major decision; members recruited from the main body; reporting into the main body's leadership; correlation above ρ = 0.75.

Detecting measurement. Quarterly error correlation, plus the blunt count: how many decisions did red team review change?

Fix. Break all three capture routes at once — reporting line, recruitment pool, social environment — or disband. A partial fix leaves the appearance intact and the function absent.

FM-6.4 · EXHORTATION IN PLACE OF STRUCTURE

Damages: nothing directly, which is the problem. It consumes the appetite for intervention while changing no parameter.

Clinical description. A u problem is met with training on open-mindedness, workshops, statements of value, and requests that people speak up, while every mechanism in Section 6.8 keeps running. u does not move, and the failure is attributed to culture being hard to change.

Diagnostic signature. Interventions whose deliverable is a document or a session rather than a changed default, schedule, metric, or piece of furniture. Measured u unchanged four quarters after a major culture programme.

Detecting measurement. Measure u before and after; the mode is defined by producing no change, so the measurement is the detector.

Fix. Replace with a structural intervention from Section 6.6. **The test for whether an intervention is structural: could it work if nobody's attitude changed at all?** Silent-write-first passes. Cancelling the shared metric passes. A workshop does not. This is not an argument against psychological safety, which is a precondition for structure to function rather than a substitute for it: a body without it will not surface the sealed answers honestly, and a body with it and no structure will still cascade.

FM-6.5 · THE ADVISORY OUTSIDER

Damages: the return on a u acquisition, without damaging the measured u — which is what makes it hard to see.

Clinical description. A genuinely non-parallel member is placed where their state cannot enter a decision: an advisory board, a consulting arrangement, a research function with no operational authority, a report that is read and filed. The geometry is real and the channel is absent. Surplus requires interaction; a state that never interacts contributes zero.

Diagnostic signature. Advisory bodies whose recommendations appear in no decision record; consultants engaged repeatedly with no traceable change; research output described as "interesting."

Detecting measurement. Trace the last twenty decisions backward and count the fraction in which the outsider's position appears in the record at all, adopted or rejected. Below 20% is the diagnosis.

Fix. Decision rights, or at minimum a required-response protocol: the outsider's position must be answered in the decision record before the decision is final. This is M-3's requirement applied to M-6's asset.

FM-6.6 · CASCADE COLLAPSE

Damages: observed u, to zero, regardless of underlying geometry.

Clinical description. Per D-6. Positions form in sequence with visibility, so later positions carry no new information. The body may be geometrically rich and epistemically empty. Diagnostic signature: speaking-order dependence, rapid unanimity, cascade gap above 0.15. Detect with the cascade gap — blind u minus post-discussion u on the same items. Fix with EU-1; the pseudo-diversity analysis [Curzi, 2026d] gives the detection toolkit and the seating and furniture analysis [Curzi, 2026f] the furniture.

FM-6.7 · THE UNAUDITED HOMOGENIZER

Damages: u, continuously and invisibly.

Clinical description. A shared rubric, metric, reading list, onboarding module, or review path installed years ago for a good reason continues to run, unowned and unre-justified, converging every member it touches at a rate nobody measures.

Diagnostic signature. Ask who owns the rubric. If the answer is a role rather than a person, or the last revision predates half the current staff, the mode is present.

Detecting measurement. The homogenizer register (EU-9). The mode is defined by absence from the register, so building the register is the detection.

Fix. EU-9 with hard expiry dates: re-justify or retire. Retiring an obsolete rubric is a simultaneous u gain and Q shed — one of the few interventions that improves two terms of the headroom equation at once.

FM-6.8 · CHASING u AT NEGATIVE HEADROOM

Damages: the intervention budget, and the credibility of the framework.

Clinical description. A system with h < 0 invests heavily in u. Because `S <= ln N` and h < 0 means `η ln N < δQ + C`, **no value of u whatsoever produces a nondecreasing step.** The intervention is well-designed, well-executed, correctly measured, and mathematically incapable of working — and the organization concludes the framework does not work.

Diagnostic signature. h computed and found negative, followed by a u programme. Or — far more common — a u programme launched without h ever being computed.

Detecting measurement. EU-14. Compute h first. Always.

Fix. The intervention ordering [Curzi, 2026j]: cut C, raise η, shed Q until h > 0. **Then** spend on u, and the same intervention that would have failed will work. The sequencing is not a preference; it is a feasibility condition from the feasibility proposition [Curzi, 2026a].

9.7 7 Results: three systems audited end to end

Three systems are worked completely, from survey through diagnosis to a costed recommendation and a declared falsification rule. **All three are constructed.** No dataset was collected, no organisation was surveyed, and every parameter below is stipulated so that the arithmetic can be redone by the reader from the numbers printed here. They are included because the framework's claims are claims about what happens when the arithmetic is actually carried out, and the failure modes of Sections 4.10, 5.14 and 6.12 are visible only in a complete case.

The three were chosen to exercise different branches of the framework. The first has positive delivery against its ceiling and a *hidden* structural problem the mean conceals. The second has a deficit so large that **no attainable geometry closes it**, which is a counter-result against the framework's own emphasis. The third is the case the framework's engineering sections are written for — non-parallelism binding, headroom available — and it is the only one of the three in which a geometric intervention is the right move.

9.7.1 7.1 A twelve-person review panel

A funding body operates a standing panel that reviews proposals in one technical area. The panel has twelve members. It has operated for six years. Its decisions are scored: three years after each funding round, an independent committee rates each funded proposal's outcome on a five-point scale, and the panel's original scores are compared against that outcome. Six years of data are available.

The body's stated problem: "our decisions have become predictable and our best applicants are going elsewhere." Its proposed solution: expand the panel to thirty members to broaden coverage.

The system is surveyed before anything is changed.

Step 1 — Declare the encoder (EC-1)

State = the member's error vector across scored proposals: for each proposal, the member's original score minus the eventual outcome rating. Six years, 214 scored proposals, twelve members. Vectors mean-centred and normalized per member.

Step 2 — Measure u (Instrument 2)

Compute all `12 × 11 / 2 = 66` pairwise Pearson correlations of the normalized error vectors, then `u_ij = 1 − ρ_ij²`.

The lineage census (Instrument 5, run first as triage per EC-9) had already flagged the structure: the twelve members come from three training lineages, in a 5 / 4 / 3 split. The measured correlations confirm the census's prior and sharpen it:

```
within-lineage mean rho = 0.97   ->   u_within = 1 - 0.97^2 = 0.0591
cross-lineage  mean rho = 0.55   ->   u_cross  = 1 - 0.55^2 = 0.6975
```

Pair accounting for a 5 / 4 / 3 split:

```
total pairs        = 12 x 11 / 2                    = 66
within-lineage     = C(5,2) + C(4,2) + C(3,2)
                   = 10 + 6 + 3                     = 19
cross-lineage      = 66 - 19                        = 47
```

Mean pairwise u under the Roster Reduction [POSTULATE]:

```
u_bar = (19 x 0.0591 + 47 x 0.6975) / 66
      = (1.1229 + 32.7825) / 66
      = 33.9054 / 66
      = 0.5137
```

Step 3 — Audit the count

```
N     = 12
u     = 0.5137
N_eff = 1 + 11 x 0.5137 = 6.651
S     = ln 6.651 = 1.8948 nats

ceiling  ln 12 = 2.4849
delivery       = 1.8948 / 2.4849 = 76.3%
redundancy deficit = 2.4849 - 1.8948 = 0.5901 nats
```

This panel is in better shape than it feels. It is delivering 76% of its ceiling. **The problem is not that the panel is parallel; the problem is that the panel is parallel in a specific place, and the mean has hidden where.**

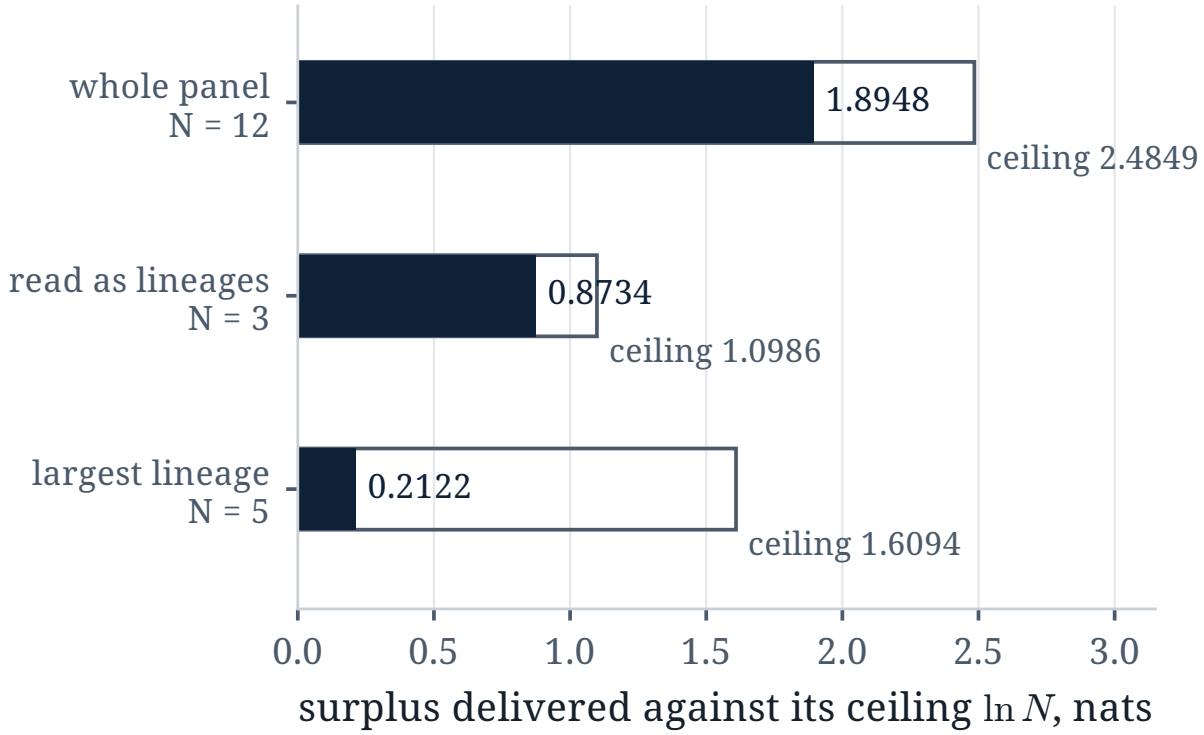


Figure 9.4: **The panel's mean passes at 76 per cent of its ceiling while its largest lineage delivers 13 per cent of its own.** Filled bars are surplus delivered, outlined bars the ceiling ln N. The whole panel reads 1.8948 nats against 2.4849; read as three lineages it reads 0.8734 against 1.0986; and the five members of the largest lineage deliver 0.2122 against 1.6094 among themselves — an effective count of 1.236 for five people. Model output on the stipulated correlations of Section 7.1; arithmetic exact. Generated by `figs/src/CCP-009_fig05.py`.

Step 4 — Read the structure, not the mean (EC-8)

Publish the block structure. Within lineages, `u = 0.0591`: the five members of the largest lineage are, pairwise, at an effective count of `1 + 4 × 0.0591 = 1.236`. Five people delivering `ln 1.236 = 0.212` nats among themselves. **Those five members are, for review purposes, approximately one member and a quarter.**

Read the panel at lineage level — three categories at cross-lineage non-parallelism:

```
N     = 3 lineages
u     = 0.6975
N_eff = 1 + 2 x 0.6975 = 2.395
S     = ln 2.395 = 0.8734      (ceiling ln 3 = 1.0986, delivery 79.5%)
```

Two readings of the same panel, at two granularities, and the honest summary is: **this is a three-instrument panel with twelve people attached to it, and nine of the twelve are duplicate readings of instruments already present.**

Step 5 — Price the proposed solution

The panel proposes growing 12 -> 30. What does that buy? Growth by the existing pipeline recruits into the existing lineages, so the block structure — and therefore `u_bar` — does not change materially. At the panel's own `u_bar = 0.5137`:

```
N = 12   N_eff = 6.651   S = 1.8948   pairs = 66
N = 20   N_eff = 10.76   S = 2.3757   pairs = 190
N = 30   N_eff = 15.90   S = 2.7663   pairs = 435
```

Growth to 30 buys `2.7663 − 1.8948 = 0.8715` nats for 18 new members and a 6.6-fold increase in communication paths. Per member: 0.0484 nats. Consult Table 4.5 — that is the marginal-gain regime where EC-10 says stop.

Worse, growth by the existing pipeline is unlikely to hold `u_bar` constant. Recruiting 18 people into three existing lineages shifts the within/cross pair balance toward within. Recompute with a 12 / 10 / 8 split at 30 members:

```
within pairs = C(12,2) + C(10,2) + C(8,2) = 66 + 45 + 28 = 139
cross  pairs = 435 − 139                                  = 296
u_bar = (139 x 0.0591 + 296 x 0.6975) / 435 = (8.215 + 206.46)/435 = 0.4935
N_eff = 1 + 29 x 0.4935 = 15.31
S     = ln 15.31 = 2.7285
```

Still a gain, but the mean u has fallen and the cost has multiplied. **Growth through an unchanged pipeline dilutes the very quantity it is being purchased to increase.**

Step 6 — Price the alternatives

Option A — rebalance to 4 / 4 / 4 by moving one member out of the largest lineage and one into the smallest. Two personnel actions, no headcount change.

```
within pairs = 3 x C(4,2) = 18,   cross = 48
u_bar = (18 x 0.0591 + 48 x 0.6975)/66 = (1.064 + 33.48)/66 = 0.5234
N_eff = 1 + 11 x 0.5234 = 6.757
S     = ln 6.757 = 1.9106      gain +0.0158 nats
```

Nearly worthless. Rebalancing within a fixed set of lineages moves 19 pairs to 18. **The lineage *count* is the binding constraint, not the lineage *balance*.**

Option B — add a fourth lineage, 3 / 3 / 3 / 3, by replacing three members. No headcount change.

```
within pairs = 4 x C(3,2) = 12,   cross = 54
u_bar = (12 x 0.0591 + 54 x 0.6975)/66 = (0.709 + 37.665)/66 = 0.5814
N_eff = 1 + 11 x 0.5814 = 7.396
S     = ln 7.396 = 2.0009      gain +0.1061 nats
```

Option C — six lineages of two. Replace roughly five members with people from three new lineages.

```
within pairs = 6 x C(2,2) = 6,   cross = 60
u_bar = (6 x 0.0591 + 60 x 0.6975)/66 = (0.355 + 41.85)/66 = 0.6395
N_eff = 1 + 11 x 0.6395 = 8.034
S     = ln 8.034 = 2.0837      gain +0.1889 nats
```

Option D — twelve distinct lineages, the theoretical limit at this cross-lineage correlation.

```
u_bar = 0.6975    (all pairs are cross-lineage)
N_eff = 1 + 11 x 0.6975 = 8.672
S     = ln 8.672 = 2.1602      gain +0.2654 nats
```

Step 7 — Also attack the cross-lineage correlation

Every option above holds `ρ_cross = 0.55` fixed. But 0.55 is itself a measured quantity and it is high for people from different training lineages. The likely cause is downstream homogenization: the three lineages read the same journals, attend the same annual meeting, and are scored by the same funder against the same declared criteria. That is a shared incentive acting as a common mode.

Suppose an intervention (Section 6) reduces `ρ_cross` from 0.55 to 0.40, leaving lineage structure at the current 5 / 4 / 3:

```
u_cross_new = 1 − 0.40^2 = 0.84
u_bar = (19 x 0.0591 + 47 x 0.84)/66 = (1.123 + 39.48)/66 = 0.6152
N_eff = 1 + 11 x 0.6152 = 7.767
S     = ln 7.767 = 2.0498      gain +0.1550 nats, zero personnel change
```

Reducing the common-mode correlation by 0.15 outperformed adding a fourth lineage and cost no personnel action at all. This is the shape of result Section 6 exists to produce, and it is the reason this paper ranks u interventions above N interventions in the lever ordering of the intervention procedure [Curzi, 2026j].

Step 8 — Put it through the circuit

The panel's other parameters, measured per Chapters 3–5: `η = 0.45` (decisions produce durable precedent and written criteria at a middling rate), `δ = 0.09` per annual cycle (membership rotates, criteria drift, institutional memory is thin), `C = 0.55` per cycle (meeting logistics, appeals, conflict-of-interest management, and the appeals load which grows with panel size), and current `Q = 9.5` on the declared order proxy.

Current headroom:

```
h = η ln N − δQ − C
  = 0.45 x 2.4849 − 0.09 x 9.5 − 0.55
  = 1.1182 − 0.855 − 0.55
  = −0.2868
```

Negative. The panel is in the regime where, by the sustainability corollary [Curzi, 2026a], no arrangement of people prevents decline. This changes the entire analysis. Every option in Step 6 was priced against S, but S is capped at `η ln N = 1.1182`, and the panel's actual delivery of `η S = 0.45 × 1.8948 = 0.8527` is already close to that cap. The panel cannot fix a negative headroom by improving its geometry, because the geometry is nearly maxed against a ceiling set by N and η.

Now price growth to 30 correctly:

```
h(N=30) = 0.45 x ln 30 − 0.09 x 9.5 − C(30)
```

`ln 30 = 3.4012`, so `η ln N` rises from 1.1182 to 1.5305, a gain of 0.4123. But C rises too. If the coordination term scales even weakly with pairs — and appeals load, scheduling, and conflict-of-interest management all do — then a 6.6-fold increase in pairs will consume the gain several times over. Taking a modest linear-in-pairs coefficient calibrated to the current panel (`C = 0.55` at 66 pairs implies roughly 0.0083 per pair at the margin, though only part of C is pair-driven — say half):

```
C(30) ~ 0.275 + 0.0042 x 435 = 0.275 + 1.827 = 2.10
h(30) = 1.5305 − 0.855 − 2.10 = −1.4245
```

Growing the panel makes the headroom five times worse. This is Section 5's thesis arriving early and the crossover calculation [Curzi, 2026b] doing exactly what it was built to do.

Step 9 — The recommendation

Given `h < 0`, the intervention procedure [Curzi, 2026j] forbids building and requires cutting C, raising η, or shedding Q. In priority order:

1. **Cut C.** Reduce panel size from 12 to 7, choosing the 7 to preserve all three lineages (3 / 2 / 2). Pairs fall from 66 to 21, a 3.1-fold reduction. New `u_bar`: `within = 3 + 1 + 1 = 5`, `cross = 16`, `u_bar = (5 × 0.0591 + 16 × 0.6975)/21 = (0.296 + 11.16)/21 = 0.5455`. `N_eff = 1 + 6 × 0.5455 = 4.273`, `S = ln 4.273 = 1.4523`. Ceiling `ln 7 = 1.9459`, so `η ln N = 0.8757`. With `C = 0.275 + 0.0042 × 21 = 0.363`: `h = 0.8757 − 0.855 − 0.363 = −0.342`. Still negative — the δQ term is the problem.
2. **Shed Q.** The panel carries 9.5 units of stored order, much of it precedent that no longer discriminates: criteria written for a different research landscape, appeals doctrine accreted over six years, a scoring rubric with sixteen dimensions of which four carry all the variance. Sunset the rubric to four dimensions and formally retire pre-2023 precedent. Estimate the shed at 40%: `Q: 9.5 -> 5.7`, so `δQ: 0.855 -> 0.513`. `h = 0.8757 − 0.513 − 0.363 = −0.0003`. Break-even.
3. **Raise η.** The panel's η = 0.45 is low because decisions produce scores but not written reasoning that binds future rounds. Instituting a two-paragraph written rationale per decision, indexed and searchable, is the classic η intervention from the conversion analysis [Curzi, 2026b]. Estimate `η: 0.45 -> 0.62`. `h = 0.62 × 1.9459 − 0.513 − 0.363 = 1.2065 − 0.876 = +0.3305`. **Positive.**
4. **Only now, raise u.** With `h = +0.3305`, there is room to request an increment. Replace one member with an outsider from a fourth lineage. At 7 members in a 2 / 2 / 2 / 1 split: `within = 1+1+1 = 3`, `cross = 18`, `u_bar = (3 × 0.0591 + 18 × 0.6975)/21 = (0.177 + 12.555)/21 = 0.6063`. `N_eff = 1 + 6 × 0.6063 = 4.638`, `S = ln 4.638 = 1.5343`. `ηS = 0.62 × 1.5343 = 0.9513` against a ceiling of `η ln 7 = 1.2065`.

Final position:

```
N = 7,  u = 0.6063,  N_eff = 4.638,  S = 1.5343
eta = 0.62, delta = 0.09, Q = 5.7, C = 0.363
h   = 0.62 x 1.9459 − 0.09 x 5.7 − 0.363 = +0.3305
Q_{t+1} = 0.91 x 5.7 + 0.62 x 1.5343 − 0.363 = 5.187 + 0.951 − 0.363 = 5.775
Q_inf   = (0.62 x 1.5343 − 0.363) / 0.09 = (0.9513 − 0.363)/0.09 = 6.537
```

The panel is smaller, it stores less, and its fixed point is 6.54 rather than the 9.5 it currently carries — which looks like a downgrade until one observes that its current 9.5 is on a declining trajectory (h < 0) and its new 6.54 is a stable attractor it converges to from either direction, by Banach.

The panel asked to be made bigger. The arithmetic said: make it smaller, write the reasoning down, throw away half the precedent, and then — only then — hire the outsider.

Step 10 — Falsification rules declared in advance

State in advance what would show this analysis wrong.

- If, after the intervention, the measured ρ among the seven does not fall relative to the twelve, the Roster Reduction's block accounting was wrong for this body and the u estimates must be rebuilt from the pairwise matrix.
- If Q on the declared proxy does not stabilize within `3τ \ensuremath{\approx} 3/0.09 \ensuremath{\approx} 33` cycles — impractically long here, so use the 1-cycle prediction `Q_1 = 5.775 ± measurement error` — the storage-law model is not describing this system.

- If the written-rationale intervention does not produce a measurable rise in η, computed as `ΔQ_attributable / S_input` per the conversion analysis [Curzi, 2026b], then η was not the binding parameter and the recommendation was wrong in its most expensive component.
- The shuffled-input control [Curzi, 2026i] is mandatory: re-run the fit with the S series randomly permuted in time. If the shuffled fit performs comparably, the apparent relationship between panel geometry and stored order is an artifact and this entire analysis should be discarded.

A framework that cannot fail cannot be trusted, and the fourth bullet is the one that would kill this one.

9.7.2 7.2 A research institute that proposed to double its domain count

A publicly funded research institute has operated for nineteen years. It spans eight declared research domains. Its stated problem, from its own strategic review: "our output per researcher has declined for six consecutive years and our field-leading position has eroded." Its proposed remedy: expand to sixteen domains over five years, on the reasoning that a broader institute is a more resilient one and that the adjacent fields are where the growth is.

The system is surveyed before anything is changed (the survey procedure [Curzi, 2026i], abbreviated here to the parameters Section 5 needs).

Step 1 — Parameters

Parameter	Value	How obtained
N	8	Declared research domains, each with its own funding line, faculty, and review process. Relabeling audit run: all eight pass RN-4 with pairwise u >= 0.42.
η	0.40	The conversion-estimation method [Curzi, 2026b]: attributable ΔQ per unit S over the last six years. Mid-range; the institute publishes and archives but converts little of its internal interaction into durable method, training, or instrumentation.
δ	0.06 per year	The retention-estimation method [Curzi, 2026a]: decay of the order proxy under a two-year period of reduced input following a funding pause in year 12.

Parameter	Value	How obtained
Q	40.0	Declared order proxy: a composite of maintained instrument-hours, active protocol count, and trained-researcher count, normalized to the year-1 baseline.
C	0.55 per year	The cost tally [Curzi, 2026b]: governance, compliance, inter-domain coordination, estate, and the review overhead of eight funding lines.

Step 2 — The ceiling and the gauge

```
ln N       = ln 8              = 2.0794
eta ln N   = 0.40 x 2.0794     = 0.8318      <- the ENTIRE annual budget

delta Q    = 0.06 x 40.0       = 2.4000
C          =                     0.5500
-------------------------------------------
h = eta ln N - delta Q - C
  = 0.8318 - 2.4000 - 0.5500
  = -2.1182
```

The institute's entire annual budget of order-production is 0.8318 nats, and its maintenance burden alone is 2.4000. Headroom is -2.1182. By the sustainability corollary [Curzi, 2026a] this is the regime in which **no choice of interaction geometry prevents order from decreasing.** Six consecutive years of decline is not a mystery, a morale problem, or a leadership failure. It is the equation.

Note what this means about the six years of effort already spent. Whatever the institute did to improve its collaboration, its culture, its recruitment, or its geometry, `S` is bounded above by `ln 8 = 2.0794` and ηS by 0.8318. The institute could have achieved perfect orthogonality across all eight domains — u = 1, every pair — and still declined by 2.1182 - 0 = 1.5682 nats per year at best. **The geometry was never the binding constraint and no geometric intervention could have worked.**

Step 3 — Price the proposed remedy

Expand from 8 domains to 16.

```
ln 16      = 2.7726
eta ln 16  = 0.40 x 2.7726 = 1.1090
gain       = 1.1090 - 0.8318 = 0.2772 nats
```

Doubling the institute's entire domain structure — five years of recruitment, eight new funding lines, eight new estates, eight new review processes — buys **0.2772 nats** against a deficit of 2.1182. It closes 13% of the gap. And that is before counting the cost.

Continue the sequence to make the point unanswerable:

TABLE 5.13 — CLOSING A 2.1182–NAT DEFICIT BY DOMAIN ADDITION ALONE

```
          (eta = 0.40, delta Q = 2.4000, C = 0.5500 held constant)

   N        eta ln N       h = eta ln N - 2.9500
 -----      --------       ---------------------
     8       0.8318              -2.1182
    16       1.1090              -1.8410
    32       1.3863              -1.5637
    64       1.6636              -1.2864
   128       1.9408              -1.0092
  1000       2.7631              -0.1869
  1600       2.9511              +0.0011
```

Solve it exactly. Break-even requires `η ln N >= δQ + C = 2.9500`, hence `ln N >= 2.9500 / 0.40 = 7.3750`, hence

```
N  >=  e^{7.3750}  =  1595.6
```

The institute would need 1,596 genuinely distinct research domains to reach break-even by domain addition alone — each one passing RN-4's u >= 0.3 test against all 1,595 others, each with its own staff, and all of them held together at the institute's current coordination cost, which they obviously would not be. With 1,596 domains the coordination term alone, at any positive k, is astronomically larger than 0.55.

This is the Domain-Addition Prohibition in a single number. Print it and put it on the wall.

Step 4 — Price the admissible levers

Lever 1 — Raise η. The institute's η = 0.40 reflects a real and fixable deficiency: internal seminars produce no artifact, methods live in individual heads, instrument configurations are not documented, and departing postdocs take their protocols with them. The conversion interventions [Curzi, 2026b] — mandatory method write-ups, a versioned protocol archive, apprenticeship overlap on instrument handover — are all cheap.

```
eta = 0.40:   h = 0.40 x 2.0794 - 2.95 = -2.1182
eta = 0.60:   h = 0.60 x 2.0794 - 2.95 = -1.7023
eta = 0.90:   h = 0.90 x 2.0794 - 2.95 = -1.0785
eta = 1.20:   h = 1.20 x 2.0794 - 2.95 = -0.4547
eta = 1.42:   h = 1.42 x 2.0794 - 2.95 = +0.0028   <- break-even
```

η would have to more than triple, to 1.42, to close the deficit alone. Tripling η is ambitious but it is not absurd — an institute that currently captures almost nothing in durable form has a great deal of room — and unlike the 1,596 domains it is at least the right order of magnitude.

Lever 2 — Cut C.

```
C = 0.55:  h = 0.8318 - 2.40 - 0.55 = -2.1182
C = 0.30:  h = 0.8318 - 2.40 - 0.30 = -1.8682
C = 0.10:  h = 0.8318 - 2.40 - 0.10 = -1.6682
C = 0.00:  h = 0.8318 - 2.40 - 0.00 = -1.5682
```

Even eliminating coordination cost entirely — an impossibility — leaves h at -1.5682. C is not the problem. The δQ term is.

Lever 3 — Shed Q.

```
Q = 40.0:  h = 0.8318 - 2.400 - 0.55 = -2.1182
Q = 30.0:  h = 0.8318 - 1.800 - 0.55 = -1.5182
Q = 20.0:  h = 0.8318 - 1.200 - 0.55 = -0.9182
Q = 10.0:  h = 0.8318 - 0.600 - 0.55 = -0.3182
Q =  4.7:  h = 0.8318 - 0.282 - 0.55 = -0.0002   <- break-even
```

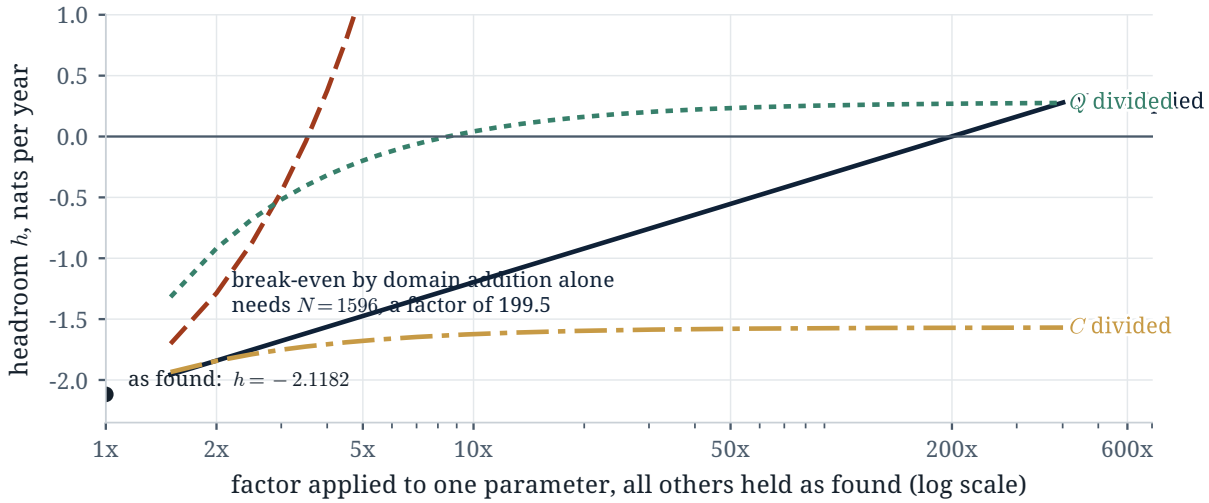


Figure 9.5: **No attainable geometry closes the institute's gap, and the domain lever is the weakest of the four.** Each curve multiplies one parameter of the worked institute ($\eta = 0.40$, $\delta = 0.06$, $Q = 40.0$, $C = 0.55$, $N = 8$) and holds the rest as found. Headroom starts at -2.1182; multiplying N by 199.5, to $N = 1596$, is what break-even by domain addition alone requires, while cutting C to zero never reaches it at all. Modelled with the source's stipulated parameters; arithmetic exact. Generated by `figs/src/CCP-009_fig04.py`.

The institute would have to shed 88% of its stored order to break even on this lever alone. That is the arithmetic statement of a nineteen-year-old institution that has accumulated more structure than its input can maintain. This is the maintenance trap of the headroom analysis [Curzi, 2026b], and it is the actual diagnosis: the institute is not failing to produce; it is failing to stop maintaining.

Step 5 — The combination that works

No single lever suffices. Combine, at levels that are individually plausible:

```
eta:  0.40 -> 0.60   (method archive, protocol versioning, handover overlap)
C:    0.55 -> 0.25   (merge four of eight funding lines into two; single
                      compliance function; retire two of five governance bodies)
Q:    40.0 -> 18.0   (retire 55% of stored order: decommission the instrument
                      park's obsolete half, sunset protocols unused in four years,
                      close two domains that fail RN-4 against their neighbours)

h = 0.60 x 2.0794 - 0.06 x 18.0 - 0.25
  = 1.2476 - 1.0800 - 0.2500
  = -0.0823
```

Still marginally negative. Push η slightly further, to 0.68, which the protocol archive plus a formal apprenticeship overlap should support:

```
h = 0.68 x 2.0794 - 1.0800 - 0.2500 = 1.4140 - 1.3300 = +0.0840
```

Positive. Barely — but the sign is what matters, because the sign is the difference between a system that can grow and a system for which no policy exists that prevents decline.

Note what happened to N. Two domains were **closed**, taking N from 8 to 6:

```
eta ln 6 = 0.68 x 1.7918 = 1.2184
h = 1.2184 - 1.0800 - 0.2500 = -0.1116
```

Negative again. The closures cost 0.1956 nats of ceiling. So they must be paid for by the C and Q reductions they enable — which they are, and generously: closing two domains was what made the funding-line merge and the instrument decommissioning possible in the first place. Attribute honestly:

```
closing 2 domains:  ceiling  -0.1956
                    delta Q  -0.36    (from Q 24 -> 18, the domain-attributable part)
                    C        -0.12
                    net h    +0.2844
```

Closing two domains improved headroom by 0.2844 nats, while adding eight domains would have improved it by at most 0.2772 before costs and by a substantial negative after them. That is Section 5's thesis reduced to a single comparison, and it is the comparison no strategic review ever makes.

Final position:

```
N = 6,  eta = 0.68,  delta = 0.06,  Q = 18.0,  C = 0.25
eta ln N = 0.68 x 1.7918 = 1.2184
h        = 1.2184 - 1.0800 - 0.2500 = -0.1116
```

which is still negative, so one more cut is required. Take Q from 18.0 to 15.0 by retiring the third governance body and the associated documentation estate:

```
h = 1.2184 - 0.9000 - 0.2500 = +0.0684
```

Positive. The institute is smaller in every dimension and it is, for the first time in seven years, above water.

Step 6 — Now check Ashby (RN-7)

Before accepting the recommendation, run the separate diagnosis. The institute's surprise log over six years shows 31 incidents. Classifying by kind:

- 24 are familiar failures at rising frequency: instrument downtime, grant deadline misses, staff departures without handover. **Headroom signature.**
- 7 are novel in kind: three regulatory categories the institute had no procedure for, two data-governance failures involving a technique nobody on staff could evaluate, two collaborations that failed because no one could translate between two of the eight domains. **Ashby signature.**

The seven are real and they will not be fixed by the headroom intervention. Two of them — the data-governance failures — indicate a genuine missing distinction, which is L-1, the one legitimate reason that overrides headroom concerns.

So the recommendation carries an explicit exception:

```
Acquire ONE new category: quantitative data governance, staffed by one person
with genuinely distinct training (RN-4 test required before hiring).

ceiling gain:  eta [ln 7 - ln 6] = 0.68 x (1.9459 - 1.7918) = 0.68 x 0.1541 = +0.1048
C cost:        estimated +0.06 (one funding line, one review path)
Q cost:        +0 initially
net h:         +0.0684 + 0.1048 - 0.06 = +0.1132
```

In this case the Ashby acquisition happens to be **h-positive**, because the institute is at N = 6 where Table 5.7's ceiling gain per category is still 8.60% of the current ceiling, and because the new category is cheap. That is L-4 and L-1 arriving together. **Take it.** Had it been h-negative, the correct move would have been to take it anyway and state the cost, because a system that cannot distinguish its regulatory environment fails at its function regardless of its budget — but the cost would have had to be paid from somewhere, explicitly, in the same document.

Step 7 — What the plan says

```
RECOMMENDATION — RESEARCH INSTITUTE, YEAR 20

REJECT: expansion from 8 to 16 domains.
  Reason: h = −2.1182. Domain-Addition Prohibition (Section 5.6).
  Break-even by domain addition alone requires N = 1,596.
  Proposed expansion closes 13% of the deficit before costs and
  widens it after them.

ADOPT, in this order:
  1. RAISE eta 0.40 -> 0.68. Method archive, protocol versioning,
     mandatory handover overlap. Measured per the conversion analysis \citep{ccp:CCP-008}.
  2. CUT C 0.55 -> 0.25. Merge funding lines 8 -> 4; single compliance
     function; retire 3 of 5 governance bodies.
  3. SHED Q 40.0 -> 15.0. Decommission obsolete instrument park; sunset
     four-year-unused protocols; close 2 domains failing RN-4.
  4. THEN acquire 1 category (data governance) on Ashby grounds L-1,
     at measured u >= 0.3 against all existing domains.

RESULTING POSITION
  N = 7,  eta = 0.68,  delta = 0.06,  Q = 15.0,  C = 0.31
  eta ln N = 0.68 x 1.9459 = 1.3232
  h        = 1.3232 − 0.9000 − 0.3100 = +0.1132
  Q_inf    = (eta F(u*) − C) / delta, computed once u* is measured
             post-intervention (Section 6).
```

Step 8 — Falsification rules declared in advance

- **If h computed at year 21 has not risen above zero**, at least one of the three parameter changes did not land, and the survey must identify which. The most likely candidate is the Q shed: decommissioning is routinely announced and not executed, and Q is the parameter most prone to being reported as reduced while still being maintained.
- **If output continues to decline while h > 0**, the order proxy is wrong. The order-proxy warning [Curzi, 2026i]'s warning applies: verify that the declared proxy Y is not merely a repackaging of the input U. This is the most common measurement error in the framework and the survey protocol's order-proxy test [Curzi, 2026i] is the check.
- **If the shuffled-input control [Curzi, 2026i] fits as well as the ordered one**, the apparent relationship between the institute's interaction structure and its stored order is an artifact and the entire analysis must be withdrawn. This is the control most often omitted and it is the one that would kill this analysis if it were wrong.
- **If the seven Ashby-signature surprises continue at the same rate after the data-governance acquisition**, the missing distinction was misidentified. Re-classify the surprise log and look for the actual missing category rather than assuming the first candidate was correct.

9.7.3 7.3 An engineering organisation over sixteen quarters

A software engineering organization of roughly 90 people is arranged in ten functional categories: platform, application, data, infrastructure, security, quality, design, research, documentation, and operations. Its stated problem: “we ship, but we keep shipping the same class of mistake, and our post-incident reviews always conclude the same three things.”

That last clause is an Ashby-shaped complaint (Section 5.9), but the surprise log resolves it otherwise: the incidents are of familiar kinds, recurring. Not a missing category — a body producing one reading.

Step 1 — Survey

The survey procedure [Curzi, 2026i], abbreviated to the parameters needed here.

Parameter	Value	Method
N	10	Declared functional categories. Relabeling audit (Section 5.7): nine pass RN-4; “documentation” measures u = 0.11 against “quality” and is flagged but retained.
η	0.50	The conversion-estimation method [Curzi, 2026b]. Design documents are written and kept, which is why η is not lower; no apprenticeship and no decision log, which is why it is not higher.
δ	0.08 per quarter	The retention-estimation method [Curzi, 2026a], from decay across a two-quarter hiring freeze.
Q_0	3.0	Order proxy: runbook coverage, test suite health, documented-decision count, normalized.
C	0.30 per quarter	The cost tally [Curzi, 2026b]: cross-team coordination, on-call, release process, compliance.

Ceiling: `η ln N = 0.50 × 2.3026 = 1.1513`.

Step 2 — Measure u

Blind item set: twenty architectural and prioritization judgment items, administered simultaneously and sealed. Six discarded as unanimous (EC-6); pairwise disagreement computed on the remaining fourteen and converted per Instrument 1, cross-checked against Instrument 2 on eighteen months of scored incident predictions.

```
u_bar = 0.12  (interval [0.07, 0.17], reported as an interval per EC-4)
N_eff = 1 + 9 x 0.12 = 2.08
```

```
S     = ln 2.08 = 0.7324
ceiling ln 10 = 2.3026
redundancy deficit = 2.3026 − 0.7324 = 1.5702 nats  (68% of the ceiling forfeited)
```

Ten functional categories, ninety people, delivering the surplus of 2.08 categories.

Cascade gap: five items re-administered after discussion. Post-discussion u = 0.04 against blind u = 0.12. Gap of 0.08 — the discussion process is destroying two-thirds of what little non-parallelism the body has.

Step 3 — Check the binding constraint and the headroom (EU-14)

```
S / ln N = 0.7324 / 2.3026 = 0.318      well below 0.85  ->  u binds, not N
h = eta ln N − delta Q − C
  = 1.1513 − 0.08 x 3.0 − 0.30
  = 1.1513 − 0.24 − 0.30
  = +0.6113                              POSITIVE
```

u binds and headroom is available — the configuration in which this section's interventions are applicable and feasible. Had h been negative, EU-14 and FM-6.8 would have routed the work to The intervention procedure [Curzi, 2026j] instead.

Step 4 — The counterfactual

What happens if nothing changes? Run the storage law forward at u = 0.12.

```
S = 0.7324,  eta S = 0.3662,  and  delta Q + C at Q=3.0  is  0.24 + 0.30 = 0.54

t=0  Q = 3.000  ->  0.92 x 3.000 + 0.3662 − 0.30 = 2.826
t=1  Q = 2.826  ->  2.666
t=2  Q = 2.666  ->  2.519
t=3  Q = 2.519  ->  2.384
t=4  Q = 2.384  ->  2.259
t=6  Q = 2.145  ->  2.039
t=8  Q = 1.942  ->  1.853

Q_inf = (eta S − C)/delta = (0.3662 − 0.30)/0.08 = 0.827
```

The organization is converging to a steady state of 0.827 against a current 3.000 — a 72% reduction in stored order, over roughly 3τ = 37 quarters. Note that h is positive throughout: positive headroom means growth is *available*, not that it is *happening*. The nondecreasing-step test [Curzi, 2026a] gives the test — `Q_{t+1} > Q_t` iff `ηS_t > δQ_t + C_t`, and here `0.3662 < 0.54`. Declining.

Nobody would describe this as a crisis. It is a 4.5% quarterly erosion of runbook coverage, test health, and documented decisions, and it will be visible in about four years.

Step 5 — Design the campaign

Three interventions, chosen by the Section 6.6 ordering and by what this body's diagnosis actually shows.

Q1 — M-1, silent-write-first, plus EU-1. The cascade gap of 0.08 says the discussion process is the largest immediate loss. Cost: three minutes per decision. Expected: recovery of most of the gap, and then more as people stop pre-converging in anticipation. `u: 0.12 -> 0.20` by Q2.

Q2 — M-2, cut the shared incentive. The audit finds one organization-wide quarterly delivery metric, bonus attached, against which all ten categories are assessed — security, quality,

and research all optimizing a throughput number. Replace with three clusters carrying non-derivable objectives (EU-3): delivery, reliability, discovery. `u: 0.20 -> 0.35` over three to four quarters.

Q3 — M-3 and M-8 together. A rotating adversarial role on every architectural decision (badged `[PARTIAL SUBSTITUTE]` per EU-13, so not entered into the structural estimate), plus two requisitions under a lineage constraint: the census shows 61% of engineering staff trained in one framework tradition, exceeding EU-5's one-third cap. Expected: `u: 0.35 -> 0.50` over four quarters.

Not attempted: raising N. `S / ln N = 0.318` says N does not bind. The flagged "documentation" category, measuring u = 0.11 against "quality," is left in place and re-tested at Q8.

Step 6 — The trajectory

Run the storage law forward with the campaign's u path. Parameters unchanged: `η = 0.5`, `δ = 0.08`, `N = 10`, `C = 0.30`, `Q_0 = 3.0`.

TABLE 6.6 — CAMPAIGN TRAJECTORY, SIXTEEN QUARTERS

t	u	N_eff	S = ln N_eff	h_t	Q_t	->	Q_{t+1}
0	0.12	2.08	0.7324	+0.6113	3.000	->	2.826
1	0.12	2.08	0.7324	+0.6252	2.826	->	2.666
2	0.20	2.80	1.0296	+0.6380	2.666	->	2.668
3	0.28	3.52	1.2585	+0.6379	2.668	->	2.784
4	0.35	4.15	1.4231	+0.6286	2.784	->	2.972
5	0.41	4.69	1.5454	+0.6135	2.972	->	3.207
6	0.46	5.14	1.6371	+0.5947	3.207	->	3.469
7	0.50	5.50	1.7047	+0.5737	3.469	->	3.744
8	0.52	5.68	1.7370	+0.5518	3.744	->	4.013
9	0.54	5.86	1.7681	+0.5302	4.013	->	4.276
10	0.55	5.95	1.7834	+0.5092	4.276	->	4.526
11	0.56	6.04	1.7984	+0.4892	4.526	->	4.763
12	0.56	6.04	1.7984	+0.4703	4.763	->	4.981
13	0.57	6.13	1.8132	+0.4528	4.981	->	5.189
14	0.57	6.13	1.8132	+0.4362	5.189	->	5.381
15	0.57	6.13	1.8132	+0.4208	5.381	->	5.557
16	0.57	6.13	1.8132	+0.4068	5.557	->	5.719

```
Q_inf at u* = 0.57:  (0.5 x 1.8132 - 0.30)/0.08 = 7.582
tau = 1/delta = 12.5 quarters;  3tau = 38,  5tau = 62
```

Four things to read from this table.

First, the turnaround is at t = 2, not t = 0. Q continues to fall for two quarters after the campaign begins, because the integrator is still discharging its initial transient `(1-δ)^t Q_0` while the driven term builds. This is the recursion result and the central point of the circuit analysis [Curzi, 2026a]: **a leaky integrator does not respond instantly, and a campaign evaluated at one quarter will be judged a failure at exactly the moment it is working.** Anyone who cancels this programme at Q2 because "the numbers are still going down" has misread the circuit.

Second, the arrival is slow. τ = 12.5 quarters, so 3τ ≈ 38 quarters to reach 95% of the new steady state. At t = 16 the system is at 5.719 against a fixed point of 7.582, 75% of the way. The intervention was fast; the arrival is slow, and the slowness is δ, not the intervention.

Third, headroom falls monotonically as Q rises — from +0.6113 at t = 0 to +0.4068 at t = 16. The headroom observation [Curzi, 2026b] that success consumes its own headroom, arriving on schedule. At the fixed point, headroom is `1.1513 − 0.08 × 7.582 − 0.30 = 0.2447`: positive,

but 60% consumed. The maintenance transition lands at about Q = 10.6, where `δQ + C = η ln N` exactly, and the organization should know that now rather than discovering it in a decade.

Fourth, compare against the counterfactual. At t = 8, campaign Q = 4.013 against counterfactual 1.853, a factor of 2.17; at the fixed points, 7.582 against 0.827, a factor of 9.17. The campaign cost three minutes per meeting, one metric redesign, one rotating role, and two lineage-constrained hires — and it moved the organization's stable stored order by a factor of nine.

Step 7 — What was and was not achieved

Honest accounting. What did **not** change: N stayed at 10, η at 0.50, δ at 0.08, C at 0.30, and headcount rose by two on a base of ninety. Every gain in Table 6.6 came from one parameter, moved mostly by cancelling things.

What was not delivered: the final u of 0.57 is well short of 1.0, and `S = 1.8132` is 79% of the ceiling. By RN-6's test `S / ln N = 0.787`, still below 0.85, so u binds marginally and a little remains. Take it carefully — from Table 6.2, u from 0.60 to 0.70 buys 0.0658 nats of `ηS` and 0.82 of `Q_\ensuremath{\infty}`, against 0.1939 and 2.42 for the move from 0.12 to 0.20. The campaign has collected the front-loaded portion, and the correct next move is η, not more u.

Step 8 — Maintenance, or it all comes back

The campaign's output is `u* = 0.57`. Per Section 6.9, that figure decays. The organization must now budget maintenance as a standing line (EU-10) or repeat Table 6.5:

```
at d_u = 0.05:  Q_inf sustainable falls  7.582 -> 4.548 over 12 quarters
at d_u = 0.10:  Q_inf sustainable falls  7.582 -> 1.848 over 12 quarters
at d_u = 0.15:  Q_inf sustainable falls  7.582 -> -0.325 over 12 quarters
```

At a 15% quarterly convergence rate the organization is back below its unsustainable threshold in three years, carrying a Q of around 6 when it gets there — so the second decline is steeper than the first. Standing commitments per EU-2 and EU-7 to EU-11: quarterly measurement, new-member correlation tracking, the homogenizer register with three-year expiries, the rotation cadence, and two independent review paths.

Step 9 — Falsification rules declared in advance

Stated in advance, before the campaign begins, and frozen.

- **The blind u must rise.** If measured u at Q4 is not above 0.25, the incentive change did not affect the geometry and the model of what was homogenizing this body is wrong — most likely the shared metric was not the binding common mode and shared hiring is, given the 61% concentration.
- **The trajectory must turn by Q3.** Table 6.6 puts the minimum at t = 2, `Q \ensuremath{\approx} 2.666`. Still falling at t = 4 means δ was mis-estimated or the proxy is not tracking Q.
- **Q at t = 8 must be within measurement error of 4.013.** A miss on the level with a correct turning point implicates η; a miss on both implicates δ.
- **The order proxy must not be a repackaging of the input.** the order-proxy warning [Curzi, 2026i]. If documented-decision count is partly a measure of how much the body interacts, it is contaminated and must be dropped from the composite. **This is the most common measurement error in the framework**; the survey protocol [Curzi, 2026i] §24.7 gives the test.

- **The shuffled-input control [Curzi, 2026i] is mandatory.** Refit `Y_{t+1} = aY_t + bF(U_t) – Ĉ_t + ε_t` with the U series randomly permuted in time. If the shuffled fit performs comparably on held-out data, the apparent link between this organization's geometry and its stored order is an artifact and **this entire analysis must be withdrawn, including the parts that appeared to work.** It is the control most often omitted and the one that would kill this result.

u is the fast parameter. It collects on every category already bought, it moves in a quarter, and most of the moves are subtractions — cancel the shared metric, the sequential discussion, the culture-fit screen; sunset the five-year-old rubric. It is also the parameter that decays fastest and most invisibly, because every force acting on it — cohesion, success, shared training, a good hire being welcomed — feels like the organization working well.

Non-parallelism is not a state an institution achieves. It is a load an institution carries, and the moment it stops carrying it the load slides back down. the two-source decomposition [Curzi, 2026d] splits the load into its two independent components, so that it is possible to tell which half is being carried and which half was dropped.

9.8 8 Discussion

9.8.1 8.1 What follows from the three sections

The audit is the whole result. Every quantitative claim in Sections 4 to 6 is a consequence of one substitution: replacing N by $1 + (N - 1)u$ wherever a count enters the storage law. The substitution costs nothing to make and it changes the sign of most institutional recommendations. A hundred-person department at $u = 0.05$ delivering an effective 5.95 is not a rhetorical flourish; it is the same arithmetic that gives Markowitz [1952]'s result that a hundred perfectly correlated holdings diversify as well as one.

The three levers are ordered by their arithmetic, not by preference. Conversion enters the budget linearly, non-parallelism enters linearly inside a logarithm and multiplies the full $(N - 1)$, and category count enters only as $\ln N$. That ordering is why the domain-addition prohibition of Section 5.6 is a structural result and not advice: closing a deficit D by raising N alone requires multiplying N by $e^{D/\eta}$, which is exponential in the deficit, while the cost terms it must outrun grow at least linearly and typically quadratically [Brooks, 1975, Conway, 1968]. A logarithm loses that race.

Two results in this paper are corrections to intuitions the framework itself invites. The first is the fold at ninety degrees: u depends on ρ^2, so opposition is parallelism and the designated contrarian is priced at zero (Section 6.2). The second is that the mean pairwise u hides where the redundancy sits. The twelve-person panel of Section 7.1 delivers 76.3 per cent of its ceiling and looks healthy on the mean, and the block structure underneath shows five members operating at an effective count of 1.236 among themselves. **The summary statistic passed and the structure failed, which is exactly the case the specification's requirement to publish the pairwise matrix exists to catch.**

9.8.2 8.2 What does not follow

Nothing here establishes that institutions obey the surplus law. The uniqueness of f under S1–S4 is a theorem about four axioms, one of which was chosen. The empirical content is a modelling proposal, and the protocol that would test it is specified elsewhere in the corpus [Curzi, 2026i] — where its one confirmatory application **failed its second falsification rule**, returning a 2.8 per cent held-out advantage over a matched autoregression against a declared 10 per cent margin. That result is the correct context for every number in this paper.

Nothing here licenses treating the affine effective count as a diversity index. Section 3.3 and Section 9.6 are the argument. A reader who imports the Hill literature's estimation theory,

partitioning theory or replication principle will import guarantees the affine count does not have.

The monotone claim is not defensible unqualified. Surplus is strictly increasing in u by Curzi [2026]'s Thm 3.1, and the empirical literature does not agree that value is. Uzzi and Spiro [2005] finds an inverted U in creative output against cohesion; Putnam [2007] finds short-run reductions in trust and participation with diversity, a result its author did not want, and Abascal and Baldassarri [2015] is the rebuttal that should be cited with it; Mannix and Neale [2005] and van Knippenberg and Schippers [2007] find surface-level difference often hurts in the short run; May [1972] and Gardner and Ashby [1970] show that increasing diversity and connectance destabilises randomly assembled systems; and Sagan [1993] documents that redundancy can reduce safety. **The framework's curve is monotone and the measured curves turn over.** The honest reconciliation is that the framework prices the *input* term and says nothing about the cost of extracting it, and the cost side is exactly where those results live — Section 5.12's crossover table is this paper's version of the turn, and it turns in every column. Allport [1954]'s four conditions are the same qualification in the contact literature: difference produces value under conditions, and u needs the conditions named.

Nothing here is an argument about fairness. Section 6.7 makes the point explicitly and it is worth repeating: the case for non-parallelism is arithmetic and would hold identically for sensors, estimators or measuring instruments. Harrison and Klein [2007] is the reason the point must be made carefully — three incompatible constructs travel under the word *diversity* (separation, variety, disparity) — and **this framework's u is a disparity measure and must declare itself as one.**

9.8.3 8.3 Relation to the surrounding literature

The closest empirical analogue of an effective count in the social sciences is Ugander et al. [2012], which measures adoption against the number of *distinct connected components* among a person's contacts rather than against the number of contacts — an effective count of independent sources, at scale, with data. Lorenz et al. [2011] measures the collapse directly: social influence reduces the effective diversity of a population's estimates and simultaneously increases confidence, which is Section 6.8's D-6 with an experiment attached. Bikhchandani et al. [1992] and Banerjee [1992] supply the formal core of that collapse, and Douceur [2002] supplies its adversarial form — one party presenting as many, which is a proof that an unaudited count is worthless.

On the pro-difference side, Hong and Page [2004] is the theorem the framework's engineering sections lean on and Thompson [2014] is the detailed objection to its conditions; the pair must travel together, and Grim et al. [2019] maps the boundary conditions under which the result holds in simulation while Singer [2019] and Kuehn [2017] are the repairs. Woolley et al. [2010] supplies the measured collective-intelligence factor and its turn-taking-equality component, which is the closest empirical relative of the framework's own speaking-share diagnostic. Page [2007] and Page [2010] are the book-length statements.

On concentration, Schindler et al. [2010] is the strongest single demonstration that an effective count buys stability — half the variability and one tenth the fishery closures in a portfolio of salmon populations — and Doak et al. [1998] is the standing objection that diversity–stability relationships arise from statistical averaging with no mechanism at all, with Tilman et al. [1998] the reply. Loreau and de Mazancourt [2008] names the decisive variable as **synchrony**, which is the ecological name for what ρ measures here. McCann [2000] is the fair survey.

On the institutional mechanism, March [1991] is the framework's own thesis published in 1991 with a simulation attached: fast learning homogenises an organisation and destroys the variance that produced its advantage. Section 6.8's D-1 to D-10 are that mechanism enumerated. Levinthal and March [1993] is the same argument about search; Scott [1998] is the historical version, in which legibility destroys the local variety that made a system work — and is simultaneously a warning about this paper's own instrument, since a published u figure is

a legibility device. Goodhart [1975] and Campbell [1979] are the specific form that warning takes: adopted as a target, u would cease to measure anything, and Section 6.6's M-2 identifies Goodhart capture as a u-destruction mechanism in its own right.

On structure, Alexander [1965] supplies Section 5.10's tree-versus-semilattice argument in its original form, and Jacobs [1961] supplies the mixed-primary-use version of the same claim about cities. Star and Ruhleder [1996] and Star and Griesemer [1989] supply the boundary object, which is the artefact that makes a non-nested membership work in practice. Granovetter [1973], Burt [1992] and Burt [2004] supply measured effect sizes for the bridging claim, and Centola and Macy [2007] and Centola [2010] supply the qualification that matters most for the framework's prescriptions: **behaviours needing reinforcement spread along dense local ties rather than weak long ones, so raising the audited count by adding distant contacts works for information and not for practice.**

Perrow [1984] is the objection the framework must not talk past. Interactive complexity plus tight coupling makes catastrophic failure a system property, and adding categories inside one coordination perimeter can therefore make matters worse — which is Section 5.6's prohibition reached by a different route, in 1984. Buldyrev et al. [2010] is its formal statement for coupled networks.

9.8.4 8.4 Where these results are consumed, and what parallels them elsewhere in the corpus

The affine effective count is load-bearing beyond this paper and the badge must survive every citation. The internal-surplus analysis runs the identical form on one observer's framings — $F_{\text{int}} = \ln(1 + (N-1)u)$ with N_{eff} the audited framing count and $u = 1 - \rho^2$ from error decorrelation — and badges the affine form as design axiom S3 rather than as a theorem [Curzi, 2026o], which is the correct treatment and the one this paper's Section 3.3 requires. The opposition identity of Section 6.2 is carried into the frequency framework unchanged: $u = 1 - \rho^2$ depends only on $|\rho|$, so $\rho = +1$ and $\rho = -1$ both give $u = 0$, $F = 0$, and the same folded note [Curzi, 2026]. And the headroom gate this paper's Sections 5.6 and 6.12 route through is stated as a precondition on any protocol elsewhere in the corpus — $h = \eta \ln N - \delta Q - C$, with $h < 0$ forbidding construction and $h > 0$ capping every increment at $r \le h$, worked at $N = 18$ with $\eta = 0.30$, $\delta = 0.06$, $C = 0.55$, $Q = 1.20$ to give $h = +0.2451$ [Curzi, 2026].

Two corpus results bound what Section 6 may claim about calibration and about acoustics. At the reflexive fixed point the input sits exactly at its ceiling and headroom is identically zero, $h = \eta \ln N - (\eta \ln N - C) - C = 0$, so calibration is not an unqualified target [Curzi, 2026r]; a body that has converged has no headroom left to spend on the interventions of Section 6.6, whatever its geometry. And the ceiling is spent by physical exclusion as readily as by redundancy: silencing three of eight participants costs $\ln 8 - \ln 5 = 0.4700$ nats, which is 22.6 per cent of the room's per-interaction surplus [Curzi, 2026e]. **A body can lose more of its ceiling to a reverberant room than Section 7.1's panel gained from redistributing its entire lineage structure.**

The same arithmetic recurs in three other domains and the recurrence is worth naming, because it is what makes the instrument general rather than local. In the frequency framework the deficit identity $D = 2.885390(c/\eta - F) = -2.885390(\delta/\eta)Q_\infty$ carries the surplus term into a pitch criterion and is cross-checked at an $8.88\text{e}{-16}$ residual [Curzi, 2026]; in the essence lattice the doubling theorem $HD_{\text{ess}}(i,j) = 2\log_2(pq)$ with $G_{\text{ess}} = 2G(pq) - 1$ is verified exactly over all sixty-six pairs at $N = 12$ [Curzi, 2026] — the same sixty-six pairs this paper's Section 7.1 audits, arrived at from an entirely different construction. **The coincidence of $\binom{12}{2} = 66$ in both is combinatorial and carries no further significance; saying so is cheaper than letting a reader find it.**

In the structural series the effective count is the grounded one, and the contrast is instructive rather than embarrassing. The concentration instrument uses HHI^{-1}: exposures of 40.9 / 27.3 / 18.2 / 13.6 give a naive HHI of 0.293430 and 3.41 effective jurisdictions, while a bloc collapse gives HHI $= 1.000000$ and an effective count of 1.00 [Curzi, 2026]. The same reciprocal

form prices job concentration — an as-found operating company scoring $S = 0.4050$ and 2.4691 effective jobs against a specialist target of 1.1050 [Curzi, 2026]. And the corpus's reproduction workbook states the procedural consequence in one line: $\ln(1.00) = 0$, so a structure with one effective bloc has **zero** available surplus rather than a low one, and the first step of the method is to collapse before measuring [Curzi, 2026]. **Those three results are grounded in the Hill literature and this paper's are not, and a reader moving between them must carry Section 3.3's distinction with them.**

9.9 9 Verification

9.9.1 9.1 Dimensional audit

Every equation in this paper is audited below. $[Q]$ is the dimension of the declared order proxy — whatever unit it counts — and $[T]$ is time in the system's declared period.

Symbol	Dimension	Unit	Status
x, y	$[1]$	none	**defined** as unit vectors in the encoder's state space; the encoder is declared, not measured
$u = 1 - (x \cdot y)^2$	$[1]$	none	**defined** dimensionless; a squared sine of an angle
ρ	$[1]$	none	**measured** correlation, under the error-vector encoding
N	$[1]$	count of categories	**declared**, not measured; a dimension count, not a headcount
$N_{\text{eff}} = 1 + (N-1)u$	$[1]$	count	**defined** by design axiom S3
PR	$[1]$	none	**derived** from the Gram spectrum
$F(u) = \ln N_{\text{eff}}$	$[1]$	nats	**derived**; a logarithm of a dimensionless count
$\ln N$	$[1]$	nats	**derived** ceiling
η	$[Q]$ per nat	proxy units per nat	**estimated**, and *not* dimensionless
δ	$[T]^{-1}$	per period	**estimated** decay rate
Q	$[Q]$	declared proxy units	**measured** against a declared proxy
C	$[Q]$,$[T]^{-1}$	proxy units per period	**measured**, denominated in Q-units by requirement
$h = \eta \ln N - \delta Q - C$	$[Q]$,$[T]^{-1}$	proxy units per period	**derived**

Symbol	Dimension	Unit	Status
κ	$[Q]$,$[T]^{-1}$ per path	per communication path	**stipulated** in Section 5.12

Does the surplus equation balance? $F(u) = \ln(1 + (N-1)u)$ takes the logarithm of a pure number, which is the only thing a logarithm may take. u is dimensionless because it is a squared cosine of unit vectors; N is a count. The equation balances, and it balances *because* N is a count of categories and not a headcount with an implied staffing unit attached.

Does the storage equation balance? $Q_{t+1} = (1-\delta)Q_t + \eta F(u) - C_t$ holds only if every right-hand term carries $[Q]$. The first does, since δ is dimensionless per period. The third does by the denomination requirement. The second does **only if η carries $[Q]$ per nat**, which is why η is reported in this paper as proxy units per nat and never as a bare fraction [Curzi, 2026b]. A conversion efficiency quoted as a dimensionless number makes the equation dimensionally malformed however tidy the arithmetic looks.

Where the framework has no operational unit, this audit says so. Q is stored order and has no natural unit; it counts whatever the declared proxy counts, and two bodies' Q values are not comparable unless they declare the same proxy. The underlying construct — comfort — has **no operational unit at all** and is treated as a dimensionless magnitude by stipulation. Section 9.5 records that the quantitative structure that stipulation assumes has not been tested. **The nat is not a physical unit either**: it is the unit of a logarithm taken to base e, and the choice of base is a convention, not a measurement. Reporting a surplus in bits would multiply every figure in this paper by $1/\ln 2 = 1.442695$ and change nothing.

9.9.2 9.2 Postulates

Seven postulates. Each is atomic, each is checkable alone, and each carries its status. Where one could be dropped and the results still follow, it has been dropped.

#	Statement	Status
P1	$u(x, y) = 1 - (x \cdot y)^2$ for unit vectors x, y	**DEFINITION**
P2	f depends on a pair only through u (S1); $f(0) = 0$ (S2); $f(1) = \ln N$ (S4)	**DEFINITION** — normalisation and structure
P3	$g(u) = e^{f(u)}$ is affine in u (S3)	**DESIGN AXIOM** — chosen, not derived; Curzi [2026]
P4	for a roster of N members, N is the declared category count and u is the mean pairwise value	**POSTULATED** (the Roster Reduction, Section 3.4)
P5	a member's state is their mean-centred, unit-normalised scored error vector	**POSTULATED** (the encoder, Section 4.3)
P6	$\mathrm{PR} = (\mathrm{tr}\, G)^2 / \|G\|_F^2$ is the Rényi-2 participation ratio of the Gram spectrum	**CITED** — Rényi [1961]
P7	$h_t = \eta \ln N - \delta Q_t - C_t$, and $h_t < 0$ forbids a nondecreasing step	**CITED** within the corpus — Curzi [2026a], Cor 3.4; Curzi [2026b]

P3 is the load-bearing one and it is the one a reviewer should attack first. P1, P2 and P6 are definitions or citations and carry no empirical weight at all. P4 and P5 are the two places where the framework touches a real body, and both are badged in the running text at every use.

9.3 The compounding chain

Each line names what it consumes and what it produces, and can be checked without reading the line before it for anything but its stated output.

```
P1 + P2 + P3 + S4  -> L1   f(u) = ln(1 + (N-1)u) is the unique form         DEFINITION
P3                 -> L2   N_eff = e^{f(u)} = 1 + (N-1)u                   DEFINITION
L2 + P4            -> R1   N = 100, u = 0.05:  1 + 99(0.05) = 5.95         ARITHMETIC
P1                 -> L3   u = 1 - cos^2(theta) = sin^2(theta)             ARITHMETIC
L3                 -> R3   u depends on |rho| only, so rho = -1 gives u = 0 ARITHMETIC
L2 + P6            -> R2   ||G||_F^2 = N + N(N-1)(1-u), so
                           N + 1 - N/PR = 1 + (N-1)u exactly               ARITHMETIC
L1                 -> L4   df/du = (N-1)/(1 + (N-1)u), maximal at u = 0    ARITHMETIC
L4                 -> R4   N = 12: u = 0.05 +/- 0.05 gives S in
                           [0.000000, 0.741937], width 0.741937            ARITHMETIC
L1                 -> R8   ln(2N) - ln(N) = ln 2 = 0.693147 for every N    ARITHMETIC
P4 + P5            -> R6   19 pairs at u = 0.0591, 47 at u = 0.6975:
                           u_bar = 33.9054/66 = 0.5137182                  ARITHMETIC
R6 + L2            -> R6'  N_eff = 1 + 11(0.5137182) = 6.6509000 exactly,
                           S = ln 6.6509 = 1.8947522 nats                  ARITHMETIC
R6' + L1           -> R6'' S/ln 12 = 1.8947522/2.4849066 = 0.7625044,
                           deficit = 0.5901545 nats                        ARITHMETIC
R6' + L2           -> R7   N = 30 at the same u_bar: N_eff = 15.8978273,
                           S = 2.7661825, gain = 0.8714303 nats for 18
                           hires; 12 lineages at u = 0.6975 gain 0.2654049
                           for none                                        ARITHMETIC
P7                 -> CR   institute: eta ln 8 = 0.40(2.079442) = 0.831777,
                           h = 0.831777 - 2.400000 - 0.550000 = -2.118223  ARITHMETIC
CR                 -> CR'  break-even by N alone: ln N >= 2.95/0.40,
                           N >= e^{7.375} = 1595.59                        ARITHMETIC
```

The three checks a reviewer can redo in under a minute are $1 + 99 \times 0.05 = 5.95$; $\ln 2 = 0.693147$ at every scale, so $\ln 4 - \ln 2 = \ln 1024 - \ln 512$; and $33.9054/66 = 0.5137182$, whose affine count $1 + 11 \times 0.5137182 = 6.6509$ has logarithm 1.8947522. Those three establish the paper's three headline figures.

And the chain must be read for what it is. Its first two steps are definitions, one of them a design axiom, and its two contact points with a real body — P4 and P5 — are postulates. **Every ARITHMETIC line below them is exact and none of them is evidence about an institution.** A reader who wants the empirical status of the framework should read Curzi [2026i], which ran the preregistered test and reported the failure.

9.4 The participation-ratio identity, checked on the smallest case and re-verified

The identity $N_{\text{eff}} = 1 + (N-1)\bar{u} = N + 1 - N/\text{PR}$ is a derivation. The smallest case that exercises it is $N = 2$. Take two unit vectors with inner product c. Then $\|G\|_F^2 = 2 + 2c^2$ and $\operatorname{tr} G = 2$, so $\text{PR} = 4/(2 + 2c^2) = 2/(1 + c^2)$, and

```
N + 1 - N/PR  =  3 - 2(1 + c^2)/2  =  3 - (1 + c^2)  =  2 - c^2
```

while $\bar{u} = 1 - c^2$ for the single pair, so $1 + (N-1)\bar{u} = 1 + (1-c^2) = 2 - c^2$. **The two agree identically, for every c.** *Status:* **ARITHMETIC**.

The general step is one line. With $G_{ii} = 1$ and the off-diagonal entries satisfying $G_{ij}^2 = 1 - u_{ij}$, $\|G\|_F^2 = N + N(N-1)(1-\bar{u})$ and $\mathrm{tr}\, G = N$, so $\mathrm{PR} = N/\left(1 + (N-1)(1-\bar{u})\right)$ and $N + 1 - N/\mathrm{PR} = N - (N-1)(1-\bar{u}) = 1 + (N-1)\bar{u}$.

The identity was re-verified for this paper on random unit-vector ensembles at $N = 3, 5, 8, 12$, with residuals of $0.00\text{e}{+00}$, $0.00\text{e}{+00}$, $8.88\text{e}{-16}$ and $1.78\text{e}{-15}$ — machine epsilon, not a measurement. The same identity is established at roster scale as Theorem 16.1 of the tuning-halt analysis — $M_{\text{eff}} = 1 + (m-1)\bar{u} = m + 1 - m/\mathrm{PR}$, verified there at 1.7815458788 with an error of $0.00\text{e}{+00}$ [Curzi, 2026]. The 10^{-12} tolerance quoted elsewhere in the corpus [Curzi, 2026i] and the independent re-verification over six random correlation matrices [Curzi, 2026] are checks of the same kind. **They test the implementation and not the world, and no reader should read a 10^{-12} agreement as an empirical result.**

9.9.5 9.5 Measurement-theoretic status of u

u is treated throughout as a ratio-scale quantity and that treatment is not tested. It is multiplied, averaged over pairs, differenced, and fed into a logarithm; each of those operations requires more than an ordinal scale, and none of them is licensed by anything in this paper. Michell [1997] and Michell [1999] are the strongest objection available to the whole framework: psychology routinely assumes quantitative structure without testing whether the attribute has it. **The objection applies here and is conceded.**

The formal remedy exists and was not run. Conjoint measurement [Luce and Tukey, 1964] supplies the conditions — single, double and triple cancellation — under which an attribute can be shown to have additive quantitative structure at all, within the representational framework of Krantz et al. [1971]. **The conjoint conditions were not tested for u, for comfort, or for stored order, and no such test is reported anywhere in this corpus.** The Rasch route [Rasch, 1960, Andrich, 1988] is the alternative, and the corpus takes it in one place [Curzi, 2026t], where the moment equation is explicitly demoted to ordinal scale — which is the right instinct, applied once.

What survives without the test is narrower than what the paper uses, and the difference should be stated. **The ordering survives**: a body whose members' errors correlate more strongly has a lower u than one whose errors correlate less, and the ranking of the options in Section 7.1's step 6 does not depend on the scale type. **The zero survives**, because $u = 0 \iff \rho^2 = 1$ is a statement about an identity between vectors and not a statement about a psychological magnitude. **The arithmetic on u does not survive independently of the encoding**: it is licensed by the fact that ρ is a correlation between real-valued vectors, which is a genuine ratio-scale object, and it is licensed only as far as postulate P5 — that error vectors are the states — holds. Where the encoder is a partition (Instrument 4) or a provenance census (Instrument 5), **the resulting figure is ordinal at best and the specification labels it as a prior rather than a measurement.**

One consequence is operational and belongs in every report. Stevens [1946] is the source of the rule that the permissible statistics follow the scale, and averaging a set of ordinally scored u values across pairs — which the Roster Reduction does — is not licensed when the underlying encoder is a partition. **The mean pairwise u is therefore reportable as a measurement only under the error-vector encoding, and is a triage figure under any other.**

9.9.6 9.6 The same verification chain applied to the affine effective count, which fails it

Section 3.3 stated that HHI^{-1} is a Hill number and $1 + (N-1)u$ is not. The claim is checkable, and the check is run here rather than asserted. The first chain is the one that succeeds; the second is this paper's own quantity, and it fails at the second step.

Chain A — the reciprocal index, which passes.

step	statement	status
A1	HHI $= \sum_i p_i^2$ on shares summing to one	DEFINITION
A2	for n equal shares of $1/n$, $\text{HHI}^{-1} = n$ exactly	ARITHMETIC
A3	HHI^{-1} is a numbers-equivalent: the count of equal units giving the observed index	CITED — Adelman [1969]
A4	$\text{HHI}^{-1} = {}^2D$, the Hill number of order 2	CITED — Hill [1973], Jost [2006]
A5	$\exp(H)$ is the same family at $q = 1$	CITED — Shannon [1948], Jost [2006]
A6	the same arithmetic is Simpson's index reciprocal and the effective number of parties	CITED — Simpson [1949], Laakso and Taagepera [1979]

Chain B — the affine count, which fails at B2.

step	statement	status
B1	$N_{\text{eff}} = 1 + (N-1)u$, with u a scalar non-parallelism	DEFINITION (design axiom S3, the framework's own)
B2	for two units at shares 0.9 and 0.1 with $u = 1$, B1 gives exactly 2.000000	ARITHMETIC
B3	for the same two units, $\text{HHI}^{-1} = 1/(0.81 + 0.01) = 1.219512$	ARITHMETIC
B4	for the same two units, $\exp(H) = \exp(0.325083) = 1.384145$	ARITHMETIC
B5	B2 disagrees with B3 and B4, so B1 does not return a numbers-equivalent	ARITHMETIC — **A3 fails**
B6	B1 is a function of (N, u) and not of the share vector, so it is not $\left(\sum_i p_i^q\right)^{1/(1-q)}$ for any $q > 0$	ARITHMETIC — **A4 fails**
B6a	it does agree with ${}^0D = 2.000000$, the raw richness, which is the one member of the family that ignores shares — so the agreement is with the naive count, not with a numbers-equivalent	ARITHMETIC
B7	no entry in ecology, economics, political science or information theory states B1	POSTULATED — no external support

The two chains part company at B5. HHI^{-1} and $\exp(H)$ are Hill numbers of order 2 and 1 and are fully grounded. Row B6a is the qualification that keeps B6 exact: the affine count does coincide with the order-zero Hill number on this vector, and ${}^0D$ is the raw richness, so the coincidence is with the count the audit exists to replace rather than with a numbers-equivalent. $1+(N-1)u$ is a linear interpolation between 1 and N; it is this framework's own construction, it carries no external support, and **a reviewer who reaches for the Hill literature expecting it to apply will find that it does not.** That is printed here, in the paper that develops the quantity, rather than left to be discovered. The corpus notes elsewhere that the two are different quantities [Curzi, 2026], which is to its credit; what this section adds is the arithmetic that makes the difference checkable in one line, and the instruction that **the two must not travel under one name.**

9.7 Model comparison and the scaling claims, audited against what is required

No model is fitted anywhere in this paper, so most of the model-comparison requirements are not merely unmet but inapplicable, and saying which is which matters.

Requirement	Status in this paper
Out-of-sample rather than in-sample performance	**Not applicable.** Nothing is fitted; the three systems are constructed and the arithmetic on them is exact.
A matched null with the same free-parameter count	**Not applicable**, for the same reason. The framework's confirmatory test against a matched AR(1) is reported in Curzi [2026i], and it failed.
The number of free parameters stated explicitly	**Met.** The surplus functional has none: N is declared and u is measured. The storage law adds η and δ.
An information criterion	**Not applicable here**; **not met** in the corpus, which computes none anywhere [Curzi, 2026i].
Effect sizes with intervals	**Partially met.** Section 4.6 requires an interval on S whenever $u < 0.3$ and reports one for each worked system. No interval accompanies η, δ or C in Section 7, because those are stipulated.

The scaling claims are the ones a reviewer will test, and this paper makes two. The first is that the ceiling gain per doubling of N is scale-invariant at $\ln 2$. **That is not an empirical scaling claim: it is an identity of the logarithm**, defined by the model and not estimated from anything. The second is that coordination cost grows as $\kappa n(n-1)/2$. **The exponent 2 there is likewise defined by the model, not estimated — and this paper states that plainly because the alternative is the error Clauset et al. [2009] exists to prevent.** A log-log regression of measured cost on headcount would be exactly the procedure that paper shows to be unreliable, and no such regression is run here. Stumpf and Porter [2012] adds the underdetermination: even a genuine power law is weak evidence for a mechanism. **No power law is claimed, no exponent is estimated, and no heavy tail is asserted anywhere in this paper.** Broido and Clauset [2019] is the reason that abstention is the right position for a paper about social structure.

The one place the paper uses the word *scale-free* is Section 5.8, and it is used there in a sense that must not be confused with the network sense of Barabási and Albert [1999]: **the claim is that the FORM of the N-dependence is invariant under doubling, not that any measured distribution is heavy-tailed.** The parameters δ, η and C differ enormously across the

scales named there, and only the shape is invariant. Under the transformation $N \mapsto \lambda N$ the ceiling shifts by $\ln \lambda$, which is an additive shift and not an invariance; **the invariance claimed is that the shift depends on λ alone and not on N, and that is all that is claimed.** It was checked at the ten doublings tabulated in Section 5.5 and nowhere else.

9.9.8 9.8 Limits and boundary behaviour

Every law in this paper is evaluated at the boundaries of its parameters, and two of the limits are constraints on the model's domain rather than results.

Limit	Behaviour	Reading
$u \to 0$	$N_{\text{eff}} \to 1$, $S \to 0$ exactly	Not an asymptote: an attained value. The modal institutional condition.
$u \to 1$	$N_{\text{eff}} \to N$, $S \to \ln N$	The ceiling, attained only at exact mutual orthogonality.
$\rho \to -1$	$u \to 0$	**Opposition is parallelism.** The fold at ninety degrees, Section 6.2.
$N \to 1$	$N_{\text{eff}} = 1$ for every u; $\ln N = 0$	A one-category system has no surplus and no ceiling. Degenerate, and correctly so.
$N \to \infty$ at fixed $u > 0$	$S \sim \ln N \to \infty$, unbounded	**A domain constraint, not a result.** The cost term $\kappa n(n-1)/2$ diverges quadratically, so the net turns negative at a finite n in every column of Section 5.12. The unbounded surplus limit is never reachable.
PR $\to 1$	$N_{\text{eff}} \to 1$	Every encoded state identical; the roster has left the arithmetic.
PR $\to N$	$N_{\text{eff}} \to N$	No off-diagonal Gram mass.
$\delta \to 0$	Q_∞ undefined; $Q_t \le Q_0 + t\,\eta \ln N$	**A pathology, not an ideal** — nothing ever leaves the stock, and the fixed point does not exist [Curzi, 2026a].
$\eta \to 0$	$h \to -\delta Q - C < 0$ always	No geometry rescues a body that converts nothing. Section 6.12's FM-6.8.

The $N \to \infty$ limit is the one that matters and it is the one the framework's own enthusiasm obscures. The surplus term is unbounded in N and the framework is frequently read as saying that more categories are always better in principle and merely expensive in practice. That reading is wrong on the framework's own arithmetic: with any positive κ, net surplus has an interior maximum and the crossover to negative is finite and computable, and Section 5.12 computes it at every column.

9.10 10 Limitations

The affine effective count is chosen, not derived, and has no external support. Design axiom S3 fixes the shape of every result in this paper. A different choice of S3 gives a different ceiling law, a different marginal table and a different prohibition. Section 9.6 runs the verification chain that a grounded effective count passes and shows this one failing it at the second step. **No literature in ecology, economics, political science or information theory contains** $1 + (N - 1)u$, and a reader who wants a grounded effective count should use HHI^{-1} or $\exp(H)$ and accept that they measure a different thing [Hill, 1973, Jost, 2006].

The reflexive specification is not identified from observational data. Section 3.5 states it with its citation [Manski, 1993] and names the remedy this paper adopts — stipulating the reference group, on the template of Brock and Durlauf [2001] — together with what that remedy does not buy. **Every applied** u **in this paper is a modelling position and not an identified empirical quantity.** No amount of scored data recovers the endogenous effect, and Shalizi and Thomas [2011] shows the confound is generic rather than incidental.

The framework exempts its own parameters from its own reflexivity. Lucas [1976] is the objection and this paper has no answer to it: a body that has read the framework will produce different sealed answers at the next elicitation, so u estimated before publication does not describe a population that has read the paper. Hacking [1995] generalises the problem and Ferraro et al. [2005] names the specific hazard — that a framework telling people their second opinions are worth nothing may make that true.

The functional form is underdetermined. Latané [1981]'s sub-linear power law and Briscoe et al. [2006]'s $n \log n$ fit the qualitative facts as well as $\ln(1 + (N - 1)u)$ does, and nothing in this paper discriminates among them because nothing in this paper fits anything. The uniqueness theorem is a theorem about S1–S4. **What would discriminate is an experiment that varies** N **over at least a decade at fixed measured** u **and compares held-out predictive error across the three forms with a matched parameter count; that experiment has not been run.** Peters and Gell-Mann [2016] supplies the one non-psychological derivation of a logarithmic form available and is the strongest warrant the framework has for its shape, though it warrants the shape and not the arguments.

The conjoint conditions were never tested. Section 9.5 concedes Michell [1997] in full. u is treated as ratio-scaled, the arithmetic on it is licensed only through the error-vector encoding, and Luce and Tukey [1964]'s cancellation conditions — the actual test of whether the attribute has additive quantitative structure — have not been run for u, for Q, or for comfort. **This is the single most dangerous unanswered objection to the framework and it is conceded rather than argued with.**

The Roster Reduction discards joint structure. Two rosters with identical mean pairwise u can have genuinely different geometry — a ring of mutual near-orthogonality against a tight cluster plus one outlier — and the triangle inequality fails for geometric triples [Curzi, 2026], so the mean does not determine the joint configuration. The participation-ratio identity of Section 3.4 names the aggregator but does not repair this, because PR is itself a scalar summary of a spectrum. Section 7.1 is the worked case in which the mean passed and the structure failed. Tversky [1977] is the further warning that psychological similarity is asymmetric and violates the triangle inequality, so **if** u **is read as a similarity measure over people it may not have metric behaviour at all.**

Every system in Section 7 is constructed. The panel, the institute and the engineering organisation are invented; their parameters are stipulated; their correlations, decay rates and conversion coefficients were chosen so that the arithmetic is checkable and the failure modes visible. **None of them is evidence about any real body**, and no figure in Section 7 may be cited as an empirical result. What they demonstrate is what the procedure does when it is carried out, not what institutions do.

Small-sample estimation is not addressed and no correction exists for this index. A twelve-person roster produces 66 pairwise correlations from a common set of scored items, and those correlations are themselves estimates with sampling error that the paper's intervals do not

include: Section 4.6's intervals propagate a stipulated uncertainty in u, not an estimated one. Chao et al. [2014] supplies the bias corrections for the Hill family and **they do not apply to the affine count**, so no correction was available and none was applied.

Three of the five instruments are not measurements. The collision test requires a chance correction whose baseline is itself estimated; the lineage census is explicitly a prior and is badged as such; and the disagreement-rate instrument is valid only under genuinely independent elicitation, which is the condition institutional practice most often violates. **The cascade-contaminated estimate of Section 4.10 is the most common failure and it biases u downward, which is the direction that makes the framework look most impressive.** That is stated here because a bias that flatters the model is the one a reviewer will look for first.

The `[PARTIAL SUBSTITUTE]` badge is doing real work and it is not a measurement. Adversarial roles and red teams manufacture an argument the body would not otherwise have heard; they do not change the underlying geometry, and the specification forbids entering them into a structural u estimate. **A body that raises its measured u by installing an adversarial role has raised a number without raising the quantity the number was built to report** — which is Goodhart [1975] operating on the framework's own instrument.

The monotone claim is restricted. Section 8.2 lists the contrary evidence in full. The framework prices an input term and is silent on the extraction cost, and the measured curves [Uzzi and Spiro, 2005, Putnam, 2007, May, 1972, Sagan, 1993] turn over. **The claim this paper defends is that surplus increases in u within the model, not that institutional value does.**

A rounding-order artefact in the source is recorded rather than corrected silently. Carried at full precision, the twelve-person panel's mean pairwise non-parallelism is $33.9054/66 = 0.5137182$, its audited count is $1 + 11 \times 0.5137182 = 6.6509000$ exactly, and its surplus is $\ln 6.6509 = 1.8947522$ nats — which is the source's 1.8948 to four places, and its 76.3 per cent of $\ln 12$ to three, so those figures stand. **One downstream figure does not.** The growth gain from twelve members to thirty is computed in the source as $2.7663 - 1.8948 = 0.8715$, from two intermediate values each already rounded to four places. Carried at full precision it is $2.7661825 - 1.8947522 = 0.8714303$, which rounds to 0.8714. **The published figure 0.8715 is high by one unit in the fourth decimal**, and any downstream paper quoting it inherits the artefact. The source's figure is retained in Section 7.1 so that the discrepancy is visible, and the exact value is used in Section 9.3's chain. This is a rounding-order artefact and not an error of method: the correct practice, which this paper adopts in its verification chain and the source does not adopt in its worked text, is to round once at the end.

An arithmetic error in the source is corrected in the open. The institute's worked example states that at perfect orthogonality across all eight domains the institute would "still have declined by $2.1182 - 0 = 1.5682$ nats per year at best." That subtraction is not arithmetic. **The correct figure is 2.1182**: headroom $h = \eta \ln N - \delta Q - C$ already evaluates the surplus term at its ceiling, so $u = 1$ gives a step of $0.831777 - 2.400000 - 0.550000 = -2.118223$ exactly. The figure 1.5682 is the decline at $u = 1$ **and** $C = 0$ together, which is the last row of that section's Lever 2 table and is a different counterfactual. The corrected reading strengthens the paper's counter-result rather than weakening it: perfect geometry leaves the institute declining by more than the source claimed, not less.

The paper's own headline case defeats the framework's own emphasis. The institute's geometry was never the binding constraint. Its entire structural budget is 0.831777 nats and its maintenance burden alone is 2.400000, so no attainable geometry — including perfect orthogonality — produces a nondecreasing step. **A corpus organised around the engineering of non-parallelism should say, where it is true, that non-parallelism was irrelevant**, and Section 7.2 says it.

Two defects in a companion paper's version of the same check are corrected here rather than propagated, and both were found by re-running its arithmetic. The corpus's concentration instrument runs the same failing chain on the affine count [Curzi, 2026]. **First**, it reports $\exp(H) = \exp(0.325083) = 1.384115$ for the two-unit vector $(0.9, 0.1)$; the correct value is 1.384145, since $H = -(0.9 \ln 0.9 + 0.1 \ln 0.1) = 0.3250829734$ nats and $e^{0.3250829734} = 1.3841454885$. That is a transcription slip at the fifth decimal and it changes no conclusion. **Second, and substantively,**

it states that *every* Hill number reports less than 1.39 on that vector. That is false. The Hill family is non-increasing in q, and its order-zero member is the raw richness ${}^0D = 2.000000$; the members exceeding 1.39 are exactly those of order below $q \approx 0.9808$. The corrected statement is that every Hill number of order one or higher reports at most 1.384145, and it is given in that form throughout this paper. **The correction strengthens the argument rather than weakening it**: the one Hill number the affine count does match is the one that ignores shares entirely, which is the naive count an effective count exists to replace.

Two named results in the running text have no bibliographic support here and are reported without a citation rather than attached to an entry that does not support them. The Condorcet jury theorem of Section 4.7 and the *gradus*-free classical statement of the variance-reduction identity in Section 4.3, Instrument 3, are both classical results with no entry in this corpus's bibliography. They are named rather than cited, and the correspondence between Condorcet's independence premise and this framework's non-parallelism parameter is badged in Section 3.6 as **interpretive**, with no theorem connecting them asserted anywhere. This is recorded so that a reader does not mistake absence for oversight.

One item in the specification cannot be audited from outside. RN-4 requires $u \geq 0.3$ between a proposed category and every existing one, measured under blind elicitation on at least twenty discriminating items. The threshold 0.3 is stipulated, not estimated. It corresponds to $|\rho| \leq 0.837$. **No evidence supports that threshold, no procedure in this paper derives it, and a different threshold would admit a different set of categories.**

9.11 11 Conclusion

This paper develops the audited count that the framework substitutes for headcount, states exactly what that count is and is not, and prices the three ways of raising it.

The count. The affine effective count is $N_{\text{eff}} = 1 + (N-1)u$ with $u = 1 - (x \cdot y)^2$, and surplus is $S = \ln N_{\text{eff}}$. Three of the four axioms that characterise the functional form uniquely are structural; the fourth, S3, is a **design axiom, chosen and not derived**, and it is badged at every use. The consequence is severe: **a hundred-person department at $u = 0.05$ has an effective count of 5.95.**

What the count is not. $1 + (N-1)u$ **is not a Hill number of any order**, is not a numbers-equivalent, and has no supporting entry in any external literature. For two units at shares 0.9 and 0.1 it returns exactly 2.000000 at $u = 1$, where HHI^{-1} returns 1.219512 and $\exp(H)$ returns 1.384145. **The reciprocal index is the Hill number of order two and the exponential of Shannon entropy is the Hill number of order one; the affine count is a linear interpolation and the two must not travel under one name.**

The bridge that does exist. $N_{\text{eff}} = N + 1 - N/\text{PR}$ exactly, against the Rényi-2 participation ratio of the Gram matrix of the encoded states — re-verified here at $N = 3, 5, 8, 12$ with residuals at machine epsilon. **Mean-pairwise u is the participation ratio, reparametrised into effective-count form**, which names the aggregator without repairing the postulate under it.

Opposition is parallelism. Because u depends only on $|\rho|$, perfect opposition at $\rho = -1$ scores $u = 0$ exactly as perfect agreement does. **A standing contrarian buys no surplus and costs a seat.**

The measurement-error budget. The sharp Lipschitz constant $N - 1$ is read here as a specification on precision rather than as a smoothness result. At $N = 12$, an uncertainty of ± 0.05 around $u = 0.05$ admits a surplus anywhere in $[0, 0.7419]$ — a width of 0.7419 against 0.1266 at $u = 0.70$. **Below $u = 0.3$ a point estimate of surplus is not reportable.**

The two counts must not be substituted. $N_{\text{eff}} = 1 + (N-1)u$ and the classical variance-reduction count $N_{\text{ind}} = N/(1 + (N-1)\rho)$, which tends to a hard asymptote of $1/\rho$, share a form and are different quantities.

The ceiling. Every doubling of N adds exactly $\ln 2 = 0.6931$ nats at any scale, so the gain from one category to two equals the gain from fifty to a hundred, and adding one nat of ceiling

costs a factor of e in categories. A deficit D closed by raising N alone requires multiplying N by $e^{D/\eta}$.

The worked results. A twelve-person panel with within-lineage $\rho = 0.97$ over 19 pairs and cross-lineage $\rho = 0.55$ over 47 audits to $\bar{u} = 0.5137182$, $N_{\text{eff}} = 6.6509$ and $S = 1.8948$ nats — **76.3 per cent of** $\ln 12$**, with a redundancy deficit of** 0.5901. Growing it to thirty buys 0.8714 nats for eighteen hires and 435 communication paths against 66; redistributing the existing twelve across twelve lineages buys 0.2654 nats for none.

And the counter-result, which is the paper's most important sentence. The worked research institute has headroom -2.1182 against a total structural budget of 0.8318. Doubling its domain count buys 0.2772 nats and closes 13 per cent of the gap; break-even by domain addition alone requires $N = 1596$; and **perfect orthogonality across all eight domains would leave it declining by** 2.1182 **nats per year, since headroom already evaluates the surplus term at its ceiling.** The geometry this corpus is named for was never the binding constraint in the paper's own worked example, and no geometric intervention could have worked. **The paper states that rather than choosing a different example.**

9.11.1 Declarations

Data and code availability

No external dataset was collected for this paper and none exists. The three systems worked in Section 7 — a twelve-person review panel, a research institute, and an engineering organisation — are constructed composites with stipulated parameters, produced so that every reported number can be re-derived by the reader from the parameters and equations printed in the text. **They must not be cited as evidence about any real organisation.** Every table in Sections 4 to 6 is an exact evaluation of a closed-form expression at the arguments shown. The figures are generated by the scripts named in each caption, which import the shared style module `lib/figstyle.py` and contain no data beyond the numbers printed in this paper. The participation-ratio identity of Section 9.4 was re-verified for this paper on pseudo-random unit-vector ensembles at $N = 3, 5, 8, 12$; the residuals reported there are floating-point residuals and not measurements.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi developed the framework, performed the analysis, produced the figures, and wrote the paper.

Ethics

No human-subjects research was conducted for this paper and no personal data were collected or processed. The bodies discussed are constructed. Where the paper's specifications bear on people — the lineage census, the protection of a newly recruited outsider, the rotation cadence — they are design requirements addressed to an analyst and are not reports of any study performed. **The specification's standing requirement that no intervention may remove a person's ability to decline is a constraint on prospective work**, stated in Section 6.11, and this

paper makes no clinical, employment or personnel recommendation about any identified individual.

Preregistration

This paper reports no preregistered test. The hypotheses of Section 2 are stated with their falsifiers so that a later test can be preregistered against them, and each is marked **not tested here**. The corpus's one preregistered confirmatory application of the storage law is reported elsewhere [Curzi, 2026i], where it **failed its second falsification rule**, and the shuffled control on which such tests depend has been measured within the corpus at a false-rejection rate of 0.108 against a nominal 0.05 [Curzi, 2026u].

Use of the framework's ceremonial register

The source manuscript from which this paper is drawn is written in part as an instructional manual with a ceremonial register. This paper reports the analytic content; the register is not reproduced, and the source's named constructs are renamed where they identify an object this paper analyses.

Chapter 10

Crossing Versus Divergence: Decomposing Non-Parallelism, the Decay of u, and the Room as a Circuit Element

Abstract. This paper splits non-parallelism into two sources and shows that they compete for one budget. For a block vector $u = \beta^2 + \alpha^2 \sin^2\gamma$ with $\alpha^2 + \beta^2 = 1$, so crossing (different domains) and divergence (disagreement on shared ground) sum — but u_{div} is capped at $1 - u_{\text{cross}}$, and **at $u_{\text{cross}} = 0.90$ no amount of internal argument buys more than 0.10.** The law prices the two identically at the margin, while their productivity in the design variables differs, $\partial u / \partial(\beta^2) = \cos^2\gamma$ against $\partial u / \partial(\sin^2\gamma) = \alpha^2$, which is why institutions buy crossing and starve divergence. A worked institute measures $u_{\text{cross}} = 0.72$ and $u_{\text{div}} = 0.03$, a balance index of 0.0206 and zero within-block dissents across forty decisions; adding a seventh discipline drives its period change negative at -0.002863 and lowers its fixed point from 6.265 to 5.964. **Opposition is parallelism**: at a normalised inner product of -1, $u = 0$, and a perfect opponent and a perfect echo both price at -0.5000 per seat, worse than an empty seat. Non-parallelism is then treated as a decaying state variable, $u_t = u_0(1-\lambda)^t$; a standards board at $\lambda = 0.18$ per quarter (half-life 3.49) fell from $u = 0.55$ to 0.0057 in twenty-three quarters, and its stored order peaked two quarters *after* the decay began, which is why the collapse of its input went unnoticed for years. **A one-time restoration on a system leaking three per cent per quarter returns to within 0.11 of its starting point twenty-eight quarters later: the durable quantity is $1/\lambda$, not u.** A final section locates all four storage-law terms in a physical room and works one to a measured outcome that remains negative, and says so.

Keywords: two-source decomposition; crossing; divergence; balance index; cascade decay; ceiling utilisation; room design; observed activity ratio

Series: B. Engineering the Surplus **Paper:** CCP-010

10.1 1 Introduction

Non-parallelism is not one quantity. This paper splits it into two structural sources, shows that they compete for a single unit budget, treats the result as a state variable that decays, and then locates every term of the storage law inside a physical room.

The framework's surplus functional is $F(u) = \ln(1 + (N-1)u)$ with $u = 1 - (x \cdot y)^2$, and the audited count it exponentiates is $N_{\text{eff}} = 1 + (N-1)u$ [Curzi, 2026,c]. **That count is the framework's own affine construction and not a Hill number of any order**; it has no supporting entry in the ecological, economic or information-theoretic literature, where HHI^{-1} and $\exp(H)$ are the Hill numbers of orders two and one [Hill, 1973, Jost, 2006, Curzi, 2026]. Every N_{eff} in this paper carries that badge, because the design axiom S3 that produces it is chosen and not derived.

The paper's first result is a decomposition and a budget. For a pure block vector, $u = u_{\text{cross}} + u_{\text{div}}$ with $u_{\text{cross}} = \beta^2$ and $u_{\text{div}} = \alpha^2 \sin^2\gamma$, and since $\alpha^2 + \beta^2 = 1$ the two sources are not independent supplies: $u_{\text{div}} \le 1 - u_{\text{cross}}$. **At $u_{\text{cross}} = 0.90$ no amount of internal argument buys more than 0.10, because the ground on which two people can disagree is exactly the ground they hold in common and crossing spends it.** Crossing is bought by hiring across a boundary; divergence is bought by permitting disagreement within one. They are different transactions with different costs, different observability, and different decay rates.

The second result is that the law prices them identically at the margin and the design variables do not. Because f depends on u alone, $\partial f/\partial u_{\text{cross}} = \partial f/\partial u_{\text{div}} = (N-1)/(1+(N-1)u)$ exactly. There is no premium on cross-disciplinarity in this law. But in the variables an analyst actually controls, $\partial u/\partial(\beta^2) = \cos^2\gamma$ and $\partial u/\partial(\sin^2\gamma) = \alpha^2$, so **each lever is most productive exactly where the other has not been pulled, and the correct first move is always the scarcer source.** That asymmetry, together with the fact that crossing is legible on a spreadsheet and divergence exists only in behaviour, explains why institutions systematically buy one and starve the other without ever deciding to.

The worked institute is the paper's headline measurement and it is unflattering to its own subject. A thirty-four-person research institute with six genuinely distinct disciplines measures $u_{\text{cross}} = 0.72$ and $u_{\text{div}} = 0.03$, a balance index of 0.0206, and **zero within-block dissents across forty decisions.** Its surplus would be almost unchanged if every within-discipline disagreement in the building were silenced tomorrow. Adding a seventh discipline — its own standing proposal — drives the period change negative at -0.002863 and lowers the fixed point from 6.265 to 5.964.

The third result is that opposition is parallelism. Because $u = 1 - \rho^2$ depends only on $|\rho|$, a normalised inner product of -1 gives $u = 0$ exactly as $+1$ does. Priced as a seat at $\eta = 0.35$, $\delta = 0.10$ and a seat cost of 0.05, **a perfect opponent and a perfect echo both contribute -0.5000 to the fixed point** — identical, and both worse than leaving the seat empty. A designated contrarian is not a weak purchase of divergence; it is a zero purchase with a cost overhead, and it suppresses demand for the real thing by appearing to fill the slot.

The fourth result treats non-parallelism as a state variable that leaks. Under an explicitly badged postulate, $u_t = u_0(1-\lambda)^t$. A standards review board measured at $\lambda = 0.18$ per quarter — a half-life of 3.49 quarters — fell from $u = 0.55$ to 0.0057 in twenty-three quarters with no personnel change, no policy change, and three procedural choices nobody recorded as decisions. **Its stored order peaked two quarters after the decay began**, because a leaky accumulator lags its input, which is why the collapse of its input went unnoticed for years. Two instruments built on entirely different foundations agreed at the year-five audit: the classical variance-reduction count returned $15/[1 + 14(0.71)] = 1.371$ independent reviewers and the surplus law returned $N_{\text{eff}} = 1.203$. **A one-time restoration on a system with $\lambda = 0.03$ is a step input to a leaky accumulator: it peaks at 3.202 and returns to within 0.11 of its starting point twenty-eight quarters later. The durable quantity is not u but $1/\lambda$.**

The fifth section locates all four storage-law terms in a physical volume, and does so under three badged postulates rather than a derivation. The worked room — a 120 m^2 departmental common room — audits to $N = 2$ use categories and headroom -1.3957. Three moves, chosen by the arithmetic rather than by the department's own plan, close 66 per cent of that gap; at week eight the observed optional-activity ratio has risen from 0.185 to 0.335 while necessary activity is flat at 211 against 216, which is the internal-validity check the taxonomy supplies for free. **And the room ends still negative at -0.4808, which the paper reports rather than repairs.**

The counter-result runs through all three sections. A room cannot manufacture non-parallelism, cannot raise conversion above what surrounding practice permits, and cannot make a negative-headroom system positive. The upper bound $f(u) \leq f(u_{\text{cross}}) + f(u_{\text{div}})$ is an impossibility limit and never a spendable budget: at a balanced split and full non-parallelism it overstates the true value by 44.7 per cent, and **any plan that adds the two sources' surpluses overstates its deliverable by up to forty-five per cent, with the error largest exactly where the ambition is highest.**

Every system in Section 7 is constructed. **No dataset was collected**, and Section 10 records what that costs the paper's claims. Sections 2 and 3 state the hypotheses and the framework, including the identification problem that no measurement of u can answer. Sections 4, 5 and 6 develop the decomposition, the decay, and the room, each closing with its specification and its named failure modes. Section 7 works three systems end to end; Sections 8 to 11 discuss, verify, limit and conclude.

10.2 2 Hypotheses

Some results below are theorems about a functional form and cannot be falsified by observing an institution; they are labelled derivations or definitions where they appear. Others are claims that the form describes real bodies and real rooms, and those are stated here with what would defeat them [Popper, 1959, Platt, 1964].

H1 — The two sources compete for one budget. For a pair whose states decompose against a declared block structure, u_{div} cannot exceed $1 - u_{\text{cross}}$. **Predicted:** measuring β^2 by an evaluability partition and $\sin^2\gamma$ by shared-ground correlation, no measured pair returns $u_{\text{cross}} + u_{\text{div}} > 1$. **Falsified by:** any measured pair exceeding one, which would show the block decomposition does not describe the encoded states. *(Status: this is an **algebraic consequence** of $\alpha^2 + \beta^2 = 1$; the empirical content is whether the encoding is right, not whether the sum holds. **Not tested here.**)*

H2 — Institutions accumulate crossing and lose divergence. In bodies with an explicit diversity programme, the balance index B is small and the missing source is divergence rather than crossing. **Predicted:** across such bodies, $u_{\text{cross}} \gg u_{\text{div}}$ and $B < 0.10$. **Falsified by:** a survey series in which the deficient source is as often crossing as divergence, or in which B is routinely above 0.3. **Not tested here.** This paper's constructed institute is consistent with H2 by construction and is not evidence for it; the hypothesis is stated so a real survey can be run against it.

H3 — Opposition buys nothing. A seat filled by a participant whose position is derivable from the proposal contributes no surplus. **Predicted:** a designated-opponent arrangement, measured by shared-ground correlation on independently scored items, returns $|\rho|$ near one and no measurable improvement in aggregate accuracy over the arrangement it replaced, while costing one seat. **Falsified by:** a measured accuracy gain at $\rho \to -1$, which would establish that sign carries information the square discards and that the encoder is wrong. **Not tested here.** This is the paper's sharpest and most checkable empirical prediction.

H4 — Non-parallelism decays geometrically absent maintenance. In a group under ordinary communication and with no u-maintenance regime, $u_{t+1} = (1 - \lambda)u_t$ with $\lambda \in [0, 1)$. **Predicted:** two measurements separated by k maintenance-free periods, on the same instrument and the same item pool, give a λ that predicts a third measurement. **Falsified by:** measured u under zero maintenance that does not decline approximately geometrically — for instance a step, a plateau, or an increase. *(Status: a **badged modelling postulate**, adopted because it is the simplest*

form consistent with the cascade mechanism of Section 5.10, because it makes λ estimable from two measurements, and because it is falsifiable. ***Not tested here.)***

H5 — Optional activity proxies conversion. In a bounded volume containing people, the optional-activity ratio tracks the fraction of interaction converted into durable structure. **Predicted:** a room whose observed activity ratio rises shows a proportionate rise in attributable ΔQ per unit of interaction. **Falsified by:** a room whose ratio rises while attributable ΔQ does not — **which is a real and common case**, and is catalogued in Section 6.15 as the room that generates pleasant lingering and stores nothing. *(Result:* ***the worked room is a partial counter-instance.*** *Its observed ratio rose by a factor of 1.81 and its measured conversion by a factor of 1.57, so the two moved together but not proportionately, and Section 7.3 reports the discrepancy.)*

Four claims in this paper are not hypotheses and are labelled as what they are.

S3, the affine effective count, is a DESIGN AXIOM — chosen because it renders the effective count additive in difference, which is what an audit of a headcount wants, and not derived from anything [Curzi, 2026].

That participants can be represented by state vectors in a block-structured space, and that the declared categories correspond to the blocks, is a POSTULATE. It is falsifiable in one specific way, stated in Section 3.8: shuffle the category assignment and refit; if the fitted parameters do not degrade, the block structure was not real.

That a room is a circuit element with its own δ, η, N, u and C is a POSTULATE about model form, not a claim about physics. The storage law is scale-free in form; asserting that a forty-square-metre volume is one of the scales at which the form applies is a choice that must be earned by measurement in each case.

That $1 - r^2$ on shared-ground scores estimates $\sin^2\gamma$ is a MODELLING CHOICE. It is the standard coefficient of non-determination read as a squared sine, and it is defensible rather than derived. It is declarable in any preregistration and is badged at every use.

10.3 3 Framework: the decomposition, the bounds, and what they license

10.3.1 3.1 The theorem, stated exactly

****The two-source decomposition** [Curzi, 2026].** For a pure block vector,

```
u = u_cross + u_div

    u_cross = beta^2              (between-category component)
    u_div   = alpha^2 sin^2(gamma) (within-category component)

max( f(u_cross), f(u_div) )  <=  f(u)  <=  f(u_cross) + f(u_div)
```

with `f(u) = ln(1 + (N-1)u)` as established by Curzi [2026], Thm 2.1 under axioms S1-S4, and `alpha^2 + beta^2 = 1`.

That is the whole formal content of the decomposition. Everything below is either a reading of those five lines or an instrument for measuring their terms.

10.3.2 3.2 What alpha, beta, and gamma are

Take two participants, `x` and `y`, each a unit state vector in a space carrying a block structure — the blocks being the `N` categories the survey (the survey protocol [Curzi, 2026i]) declared. Let `x` lie inside one block. Decompose `y` relative to that block:

```
y = alpha * (unit vector inside x's block, at angle gamma from x)
  + beta  * (unit vector outside x's block)

alpha^2 + beta^2 = 1,  alpha, beta >= 0
```

Then the inner product only survives on the shared block:

```
x . y   = alpha cos(gamma)
u       = 1 - (x . y)^2
        = 1 - alpha^2 cos^2(gamma)
        = 1 - alpha^2 + alpha^2 sin^2(gamma)
        = beta^2 + alpha^2 sin^2(gamma)
        = u_cross + u_div
```

A one-line verification of the theorem's form, not a re-derivation — worth doing once, because it fixes what the three parameters mean in an organization:

Symbol	Formal meaning	Organizational reading
`beta^2`	squared component of `y` outside `x`'s category block	the fraction of the second party's competence that the first party cannot evaluate
`alpha^2`	squared component of `y` inside `x`'s block, `= 1 - beta^2`	the shared ground: the part of the work both parties can actually judge
`gamma`	angle between the two inside the shared block	how much they disagree about the things they both understand

`u_cross` is bought by hiring across; `u_div` is bought by permitting disagreement within. Different transactions, different costs, different failure modes — and, as Section 4.3 shows, a hard budget relation between them that almost nobody accounts for.

10.3.3 3.3 The two bounds, and what each one licenses

The bound is two-sided and the two sides do different jobs.

Lower bound — the guarantee.

```
f(u) >= max( f(u_cross), f(u_div) )
```

Whatever the stronger source already supplied, adding the other cannot take it away. This is the safe half of the theorem, and it licenses one design statement: *adding a genuinely different source is never negative.*

Upper bound — the ceiling, and the warning.

```
f(u) <= f(u_cross) + f(u_div)
```

This follows from subadditivity on `u` (Curzi [2026], subadditivity [Curzi, 2026]), itself a consequence of strict concavity (3.3). The combined system is worth **at most** what running the two sources separately and adding their surpluses would give.

Stated in its practical form: **The theorem promises no synergy beyond additivity, and honest design must not assume any.** No clause in the founding surplus analysis [Curzi, 2026] says a cross-disciplinary team with internal dissent produces more than the sum of a cross-disciplinary team and a dissenting team. The mathematics forbids more than the sum. It does not forbid less, and — because `f` is strictly concave — the truth is systematically *less* whenever both sources are nonzero. How much less is computable, and it is not a rounding error.

```
N = 8,  f(u) = ln(1 + 7u)

split (u_cross : u_div)   u      f(u) true   upper bound   overstatement
-------------------------------------------------------------------------
50 / 50                   0.20     0.8755       1.0613          21.2%
50 / 50                   0.40     1.3350       1.7509          31.2%
50 / 50                   0.60     1.6487       2.2628          37.3%
50 / 50                   0.80     1.8871       2.6700          41.5%
50 / 50                   1.00     2.0794       3.0082          44.7%

75 / 25                   0.20     0.8755       1.0179          16.3%
75 / 25                   0.60     1.6487       2.1409          29.9%
75 / 25                   1.00     2.0794       2.8442          36.8%
```

At a balanced split and full non-parallelism, the upper bound overstates real surplus by 44.7 percent. **Any plan that adds the two sources' surpluses overstates its deliverable by up to forty-five percent, and the error grows exactly where the ambition is highest.** The upper bound is a proof device, not a plan.

10.3.4 3.4 The marginal identity

Because `u = u_cross + u_div` is a plain sum, and `f` depends on `u` alone (axiom S1, geometric dependence),

```
d f / d u_cross  =  d f / d u_div  =  (N − 1) / (1 + (N−1) u)
```

At the margin, in the parameter `u`, a unit of crossing and a unit of divergence are worth exactly the same. There is no premium on cross-disciplinarity in this law. The theorem is indifferent to where the non-parallelism came from.

This matters more than it looks. Any organizational preference for cross-hiring over internal dissent is a preference about *acquisition cost and observability*, not about value: a procurement decision dressed as an epistemic one. the intervention procedure's [Curzi, 2026j] lever ranking must be built on that distinction and not on the fiction that crossing is intrinsically superior.

10.3.5 3.5 The structural constraint

One consequence of `alpha^2 + beta^2 = 1` appears nowhere in the theorem's headline and governs almost every real design:

```
u_div = alpha^2 sin^2(gamma) <= alpha^2 = 1 − beta^2 = 1 − u_cross

so   u_cross + u_div <= 1   (as it must, u being in [0,1])
and  max achievable u_div, given u_cross  =  1 − u_cross
```

The two sources are not independent supplies from separate accounts; they compete for the same unit budget. At `u_cross = 0.90`, no amount of internal argument can buy more than `u_div = 0.10` for that pair, even at total within-block disagreement (`gamma = 90 degrees`). The ground on which two people can disagree is exactly the ground they hold in common, and crossing spends it.

This is the mathematical form of an observation every practitioner has made: the more foreign two specialists are to each other, the less able they are to have a real argument. They talk past each other instead — `beta^2` without `gamma`, which the law credits correctly as non-parallelism but which is a different product from a disciplined disagreement and cannot substitute for one when the question is internal to a discipline.

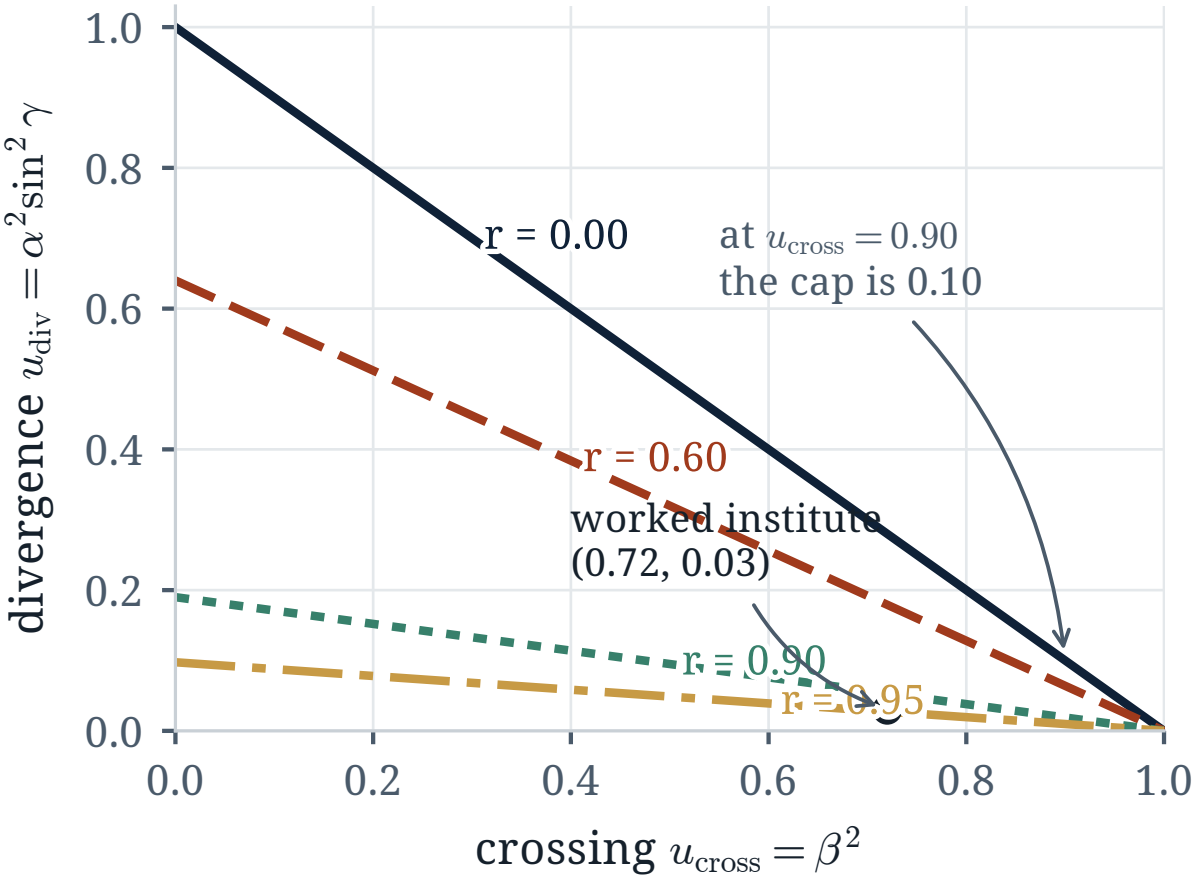


Figure 10.1: **Crossing spends the ground on which divergence would have to register.** The attainable divergence $u_{\mathrm{div}} = (1 - u_{\mathrm{cross}})(1 - r^2)$ is plotted at four within-block correlations; the envelope $u_{\mathrm{div}} = 1 - u_{\mathrm{cross}}$ is the $r = 0$ curve. **At $u_{\mathrm{cross}} = 0.90$ the cap is 0.10 whatever the protocol.** The marked point is the worked institute of Section 7.1. Model output; the constraint follows exactly from $\alpha^2 + \beta^2 = 1$. Generated by `figs/src/CCP-010_fig01.py`.

10.3.6 3.6 Identification: the objection that no measurement of the two sources can answer

Every applied reading in this paper runs through a reflexive specification. The framework's central formal claim is that a group's influence on its members and the members' constitution of the group are one loop [Curzi, 2026], and the estimators of Section 4.10 inherit it directly: the shared-ground scores whose correlation defines γ are produced by people who observe one another, and Section 5.10's cascade mechanism is precisely the process by which that observation destroys the independence the estimator assumes.

That specification is a linear-in-means reflexive model, and its endogenous social effect is not identified from observational data. Manski [1993] proves it: the endogenous effect cannot be separated from contextual effects or from correlated unobservables without prior information specifying the reference group. Shalizi and Thomas [2011] show that influence and selection are generically confounded in observational network data, and Lyons [2011] shows what happens when the confound is ignored. Durlauf [2005] adds the requirement that such a model earn its keep by making a **differential** prediction that simpler models do not also make.

The remedy this paper adopts is to stipulate the reference group and defend the stipulation, on the template of Brock and Durlauf [2001], which derives a mean-field structure from standard discrete choice with its identification conditions stated. The stipulation here is the declared block structure of Section 4.14's SPEC 4.1 — blocks declared by *evaluative standard*, in writing, before any estimate is computed — and it is defended on the ground that an evaluative standard is an administratively observable property of a body's own review practice. **That is a defence of the stipulation and not a recovery of the parameter.** Accordingly:

> **Every applied u, u_{cross} and u_{div} in this paper is a modelling position rather than an identified empirical quantity.** The algebra relating them to N_{eff} and to S is exact and is not at issue. What is at issue is whether a measured evaluability partition and a measured shared-ground correlation are the framework's β^2 and γ, and observation alone cannot settle it.

The same objection applies with more force to the decay rate. λ is estimated from two measurements of u separated by a period asserted to be maintenance-free, and **the assertion is not verifiable from the data**; a period containing unrecorded maintenance returns the net of decay and repair, which is a different quantity. Section 5.12 states the condition and Section 10 records that no procedure in this paper verifies it.

Two further objections in the same family are unanswered and are recorded rather than deflected. Lucas [1976] turned on the framework itself: a body that has read this paper will produce different sealed scores at the next elicitation, so a λ estimated before publication does not describe a population that has read it — and Section 5.10's cascade mechanism is exactly the channel through which that would operate. Hacking [1995] is the general statement, and Ferraro et al. [2005] the specific warning that a framework which assumes people will converge may help them do so. MacKenzie and Millo [2003] is the countervailing observation that such a loop has been caught happening in the field, with dates and price data, which establishes that reflexive fixed points exist without establishing that this one does.

And the functional form is underdetermined by the evidence. Latané [1981]'s sub-linear power law and Briscoe et al. [2006]'s $n \log n$ fit the qualitative facts as well as $\ln(1 + (N-1)u)$ does. Nothing in this paper discriminates among them, because nothing in this paper fits the form to data; the paper derives consequences of a form assumed. Sobkowicz [2009] and Sobkowicz [2020] are the standing indictment — many models, negligible empirical discrimination between them — and they apply. One strengthening is available and is taken: Peters and Gell-Mann [2016] derives a logarithmic form from multiplicative dynamics with no psychological assumption at all, which is a stronger warrant for the shape than the psychophysical route.

10.3.7 3.7 Analogy discipline: what is derived, what shares a form, and what is only suggestive

Each borrowing in this paper is classified at first use, in the three categories the corpus's rigour standard requires.

Borrowing	Category	What is and is not claimed
The block decomposition $u = \beta^2 + \alpha^2 \sin^2 \gamma$	**1 — derivation**	Orthogonal decomposition of a unit vector against a subspace. The algebra is exact; the empirical content is entirely in whether the declared blocks are the subspaces.
The leaky integrator, and the room's lag	**2 — shared mathematical form**	A first-order linear difference equation governs a discharging capacitor and this stock alike [Oppenheim et al., 1996, Elaydi, 2005]. **No physical claim is transferred**, and the "circuit element" language of Section 6 is a naming convention, not a physics claim [Curzi, 2026a].
Geometric decay of u	**3 — analogy, badged as a postulate**	Nothing derives it. It is adopted for simplicity and estimability and is falsifiable (H4). Friedkin and Johnsen [1990] supplies the closest formal relative — a stubbornness term in an opinion-updating model that plays the same role a leak plays here — and is cited as prior art rather than as support.
The cascade as the mechanism of decay	**1 — derivation, within its own model**	Bikhchandani et al. [1992] and Banerjee [1992] prove the result for sequential agents with private signals. **The transfer to a standing committee is an analogy**, because a committee's agents are not drawn afresh and the signal structure is not the model's.
$\ln N$ as a ceiling, against Shannon entropy	**2 — shared form only**	Both are logarithms of a category count. **Surplus is not identified with Shannon entropy or with mutual information** [Curzi, 2026]; Shannon [1948] is cited for the entropy and not for the surplus.
Requisite variety	**2 — formal correspondence**	Ashby [1956] measures variety as the logarithm of a state count and treats it as a cap on regulation. The two are not proved equal.
Robustness against targeted attack	**3 — suggestive correspondence**	Albert et al. [2000] establishes robust-yet-fragile for a degree distribution. Reading a single-sourced portfolio as a one-hub structure is a shape argument and no more; Doyle et al. [2005] is the reason it must not be pressed further.
Proxemic bands as channel types	**3 — analogy**	Hall [1966] reports distances at which sensory channels change, in one culture and one decade. Treating a band boundary as a discontinuity in the surplus a pair can carry is this paper's own reading and has no measurement behind it.

The failure mode this table exists to prevent is a category-2 or category-3 borrowing carrying a category-1 result's authority. The corpus's own strictest ruling on borrowed physics governs here: the framework's postulate is stated as a postulate, the borrowed relation is given

only in its correct form, and a user must be able to distinguish the theorem, the ordinary observer effect, and the framework's own conjecture in three separate sentences on demand [Curzi, 2026]. Mirowski [1989] is the historical account of what goes wrong otherwise, and Cartwright [1999] the generous escape route: a law may hold inside a specified arrangement without being universal, provided the arrangement is specified. **Sections 4.14, 5.19 and 6.13 are this paper's attempt to specify the arrangement.**

10.3.8 3.8 Badges and honesty clauses

- **Proven (Curzi [2026], Thm 4.1):** the decomposition and both bounds, for a pure block vector.
- **Proven (Curzi [2026], strict concavity [Curzi, 2026], 3.4):** strict concavity and subadditivity on `u`, from which the upper bound follows.
- **Design axiom, not derived (S3):** the affine effective count `g(u) = e^{f(u)} = 1 + (N−1)u`. Everything here that speaks of `N_eff` inherits that choice, made because it renders the effective count additive in difference — what an engineer wants from a headcount audit. That dependence is stated wherever the argument leans on it, as this paper does here.
- **Postulate, badged:** that participants can be represented by state vectors in a block-structured space at all, and that the declared categories correspond to the blocks. A modeling assumption, not a result, falsifiable in the manner of the preregistration protocol [Curzi, 2026i]: shuffle the category assignment, and if the fitted parameters do not degrade, the block structure was not real.
- **Honesty clause, restated:** the triangle inequality FAILS for geometric triples (the metric analysis of the surplus functional [Curzi, 2026]). Subadditivity holds on `u`, and only on `u`. Three participants may not be chained and expected to behave as a metric.
- **Honesty clause, restated:** as deployed, the surplus functional is a **lower bound** on the ensemble's information, not the information itself. The floor is logged, not the ceiling. A group with `u_div` near zero may still hold information this instrument cannot see; it will not hold information the instrument *credits*, and stored order (the institutions-as-storage analysis [Curzi, 2026h]) responds to what the instrument credits.
- **Honesty clause, restated:** The founding surplus analysis [Curzi, 2026] is pure mathematics and makes no empirical or metaphysical claim. Nothing here says a cross-disciplinary team is *good*. It says what a specified functional returns.
- **Refused (the framework's consent clause [Curzi, 2026]):** no protocol in this paper removes a person's ability to decline to disagree. Divergence performed on demand is not divergence but a second kind of compliance, and it measures as noise.

10.4 4 Crossing and divergence: the two purchases and their budget

10.4.1 4.1 Two ways not to be parallel

Case A — the crossing. A structural engineer and a public-health epidemiologist assess a proposed bridge. The epidemiologist cannot evaluate the fatigue calculations; the engineer cannot evaluate the ambulance-transit consequences of a lane closure. Almost all of the second party's state lies outside the first's block: `beta^2 = 0.85`. Within the sliver they share — cost, schedule, public communication — they broadly agree, `gamma = 12 degrees`.

```
u_cross = 0.85
u_div   = (1 − 0.85) ∗ sin^2(12 deg) = 0.15 ∗ 0.0432 = 0.0065
u       = 0.8565
```

Case B — the divergence. Two structural engineers, same degree programme, same code books, assess the same bridge. `beta^2 = 0.03`. But one is a fatigue specialist who believes the

connection detail will crack in twenty years and the other does not. Within the shared block they are 65 degrees apart.

```
u_cross = 0.03
u_div   = 0.97 * sin^2(65 deg) = 0.97 * 0.8214 = 0.7968
u       = 0.8268
```

At `N = 6`, `f(u) = ln(1 + 5u)`: Case A gives `ln(5.2825) = 1.6644`, Case B gives `ln(5.1340) = 1.6360` — within two percent of each other. **The law prices a foreign expert and a determined internal dissenter as near-equivalent goods, and an organization that buys only one of them has bought half a portfolio at full price.**

The difference is not in the surplus but in what the surplus is *about*. Case A's non-parallelism lives entirely in the question "what else does this bridge touch?" Case B's lives entirely in "is this bridge sound?" An organization with only Case A never learns that the connection detail cracks. An organization with only Case B builds a structurally impeccable bridge that closes the only route to the hospital.

10.4.2 4.2 Reading beta and gamma as organizational facts

Each parameter has a direct organizational correlate, and each is acquired by a different administrative act.

	`u_cross` (beta^2)	`u_div` (alpha^2 sin^2 gamma)
acquired by	recruiting, seconding, rotation across disciplines	protecting dissent, rotating adversarial roles, independent estimates first
visible in	org chart, CV, hiring record	minutes, dissent log, the error-decorrelation audit
costs	salary, onboarding, translation burden (C)	social friction, slower meetings, managerial nerve
decays via	attrition, assimilation to the dominant discipline	cascades, promotion by similarity, shared incentives
audit interval	annual (slow-moving)	quarterly at most (fast-moving; see Section 5)

The asymmetry in the "visible in" row is the engine of the pathology this section exists to diagnose. `u_cross` is legible on a spreadsheet: it survives an audit and can be reported to a board. `u_div` exists only in behaviour and leaves a record only if someone deliberately builds one. **An institution that manages what it can see will accumulate `beta^2` and let `gamma` go to zero without ever making a decision to do so.**

10.4.3 4.3 The budget constraint nobody accounts for

The ceiling `u_div <= 1 - u_cross` changes what "more diversity" means at different operating points, and its shape is counter-intuitive enough to need a table. Estimate `sin^2(gamma)` from the correlation `r` of two parties' independent judgments on items *both* can evaluate: `sin^2(gamma) = 1 - r^2` (an estimator choice, badged; see Section 4.10). Then `u_div = (1 - u_cross)(1 - r^2)`.

`u_cross`	r=0.95		r=0.90		r=0.60		r=0.00	
	`u_div`	`u`	`u_div`	`u`	`u_div`	`u`	`u_div`	`u`

```
 0.00   |  .098  .098    .190  .190    .640  .640   1.000 1.000
 0.10   |  .088  .188    .171  .271    .576  .676    .900 1.000
 0.25   |  .073  .323    .142  .392    .480  .730    .750 1.000
 0.50   |  .049  .549    .095  .595    .320  .820    .500 1.000
 0.72   |  .027  .747    .053  .773    .179  .899    .280 1.000
 0.90   |  .010  .910    .019  .919    .064  .964    .100 1.000
```

Reading the rows: At `u_cross = 0.90`, driving within-block correlation from 0.95 to 0.00 — near-total agreement to total independence — moves `u` from 0.910 to 1.000. At `u_cross = 0.00`, the identical change in `r` moves `u` from 0.098 to 1.000. At `N = 6`:

```
u_cross = 0.90:  f(1.000) − f(0.910) = 1.7918 − 1.7138 = 0.0780
u_cross = 0.00:  f(1.000) − f(0.098) = 1.7918 − 0.3988 = 1.3930
```

The same dissent intervention is worth nearly eighteen times more in a monodisciplinary group than in a heavily cross-staffed one — not because dissent is less valuable there in principle, but because the cross-staffing has already consumed most of the unit budget, leaving little shared ground on which the dissent can register. The converse holds too, and is the correction to the usual enthusiasm.

10.4.4 4.4 Which lever is live — marginal productivity in the design variables

In the *parameter* `u`, both sources are worth the same at the margin (Section 3.4). But the design variables an analyst controls are `beta^2` (who is hired) and `sin^2(gamma)` (what is permitted), and in those the marginal productivities are not equal:

```
u = beta^2 + (1 − beta^2) sin^2(gamma)

d u / d(beta^2)       = 1 − sin^2(gamma) = cos^2(gamma)
d u / d(sin^2 gamma)  = 1 − beta^2       = alpha^2

beta^2  sin^2 g |  du/d(beta^2)   du/d(sin^2 g)  |   u
----------------+---------------------------------+--------
 0.00    0.00   |     1.000           1.000       |  0.000
 0.00    0.50   |     0.500           1.000       |  0.500
 0.00    0.90   |     0.100           1.000       |  0.900
                |                                 |
 0.50    0.00   |     1.000           0.500       |  0.500
 0.50    0.50   |     0.500           0.500       |  0.750
 0.50    0.90   |     0.100           0.500       |  0.950
                |                                 |
 0.72    0.00   |     1.000           0.280       |  0.720
 0.72    0.50   |     0.500           0.280       |  0.860
 0.72    0.90   |     0.100           0.280       |  0.972
                |                                 |
 0.90    0.00   |     1.000           0.100       |  0.900
 0.90    0.50   |     0.500           0.100       |  0.950
 0.90    0.90   |     0.100           0.100       |  0.990
```

The rule that falls out is symmetric, and it is the single most useful operational statement in this section:

```
RULE 4.4  (which lever is live)

  Hiring across is most productive when the group already agrees internally:
```

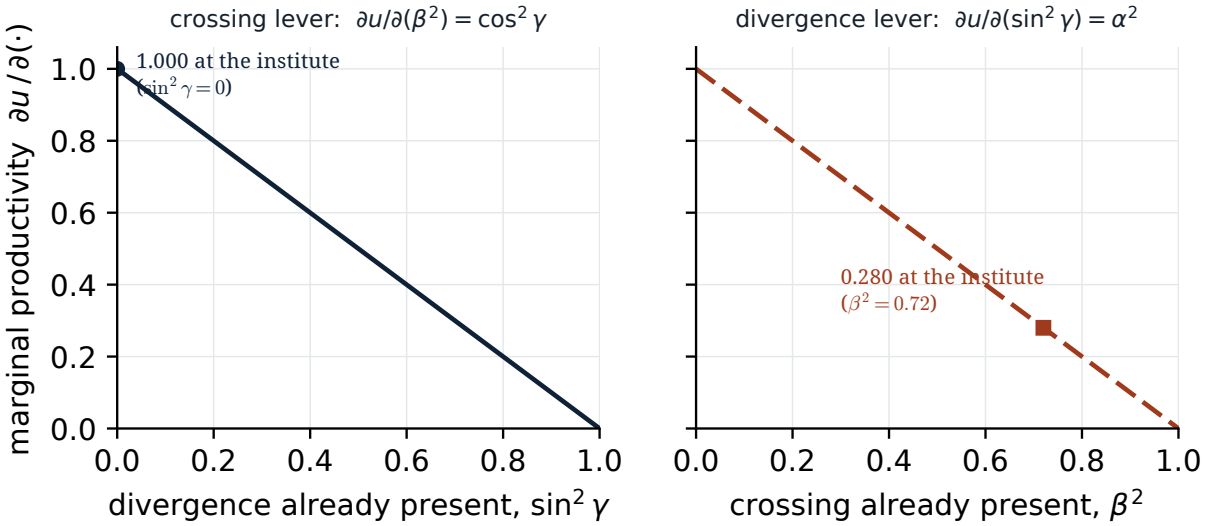


Figure 10.2: **Each lever is most productive exactly where the other has not been pulled, which is why the correct first move is always the scarcer source.** Left: the crossing lever's marginal productivity $\partial u/\partial(\beta^2) = \cos^2\gamma$ falls as divergence rises. Right: the divergence lever's $\partial u/\partial(\sin^2\gamma) = \alpha^2$ falls as crossing rises. The marked points are the worked institute, which reads 1.000 on the crossing lever and 0.280 on the divergence lever — **and Section 4.4 explains why the larger number is nonetheless the wrong purchase.** Model output; both derivatives are exact. Generated by `figs/src/CCP-010_fig03.py`.

```
  du/d(beta^2) = cos^2(gamma)  -> maximal at gamma = 0.
Permitting dissent is most productive in a monodisciplinary group:
  du/d(sin^2 gamma) = alpha^2  -> maximal at beta^2 = 0.

Therefore the correct first move is always the scarcer source.
```

Each lever's productivity is highest exactly where the other has not been pulled. A self-correcting rule: buying the scarcer source cannot go far wrong, and buying more of the abundant one will.

An organization at `beta^2 = 0.72, sin^2(gamma) = 0.00` — the classic case this section diagnoses — gets `du/d(beta^2) = 1.000` from another cross-hire against `du/d(sin^2 gamma) = 0.280` from permitting dissent. The naive reading is that cross-hiring should continue, being nearly four times as productive per unit. That reading is wrong twice over.

First, the units are not comparable. Moving `beta^2` by 0.10 requires restructuring a team; moving `sin^2(gamma)` by 0.28 requires a meeting protocol (the seating analysis [Curzi, 2026f]) and the nerve to keep it. the intervention procedure's [Curzi, 2026j] cost ranking converts marginal productivity into value per unit of expenditure, and by that measure the protocol wins in almost every real case.

Second, surplus is not the whole brief (Section 4.1). The two components are surplus about *different questions*, and a group at `beta^2 = 0.72, gamma = 0` has bought a wide field of view with no resolution anywhere in it. the effective-count analysis's [Curzi, 2026c] phrase applies with force: **three honest windows set wide capture most of what a gallery could — but a window that reports what the window next to it reports is not a second window.**

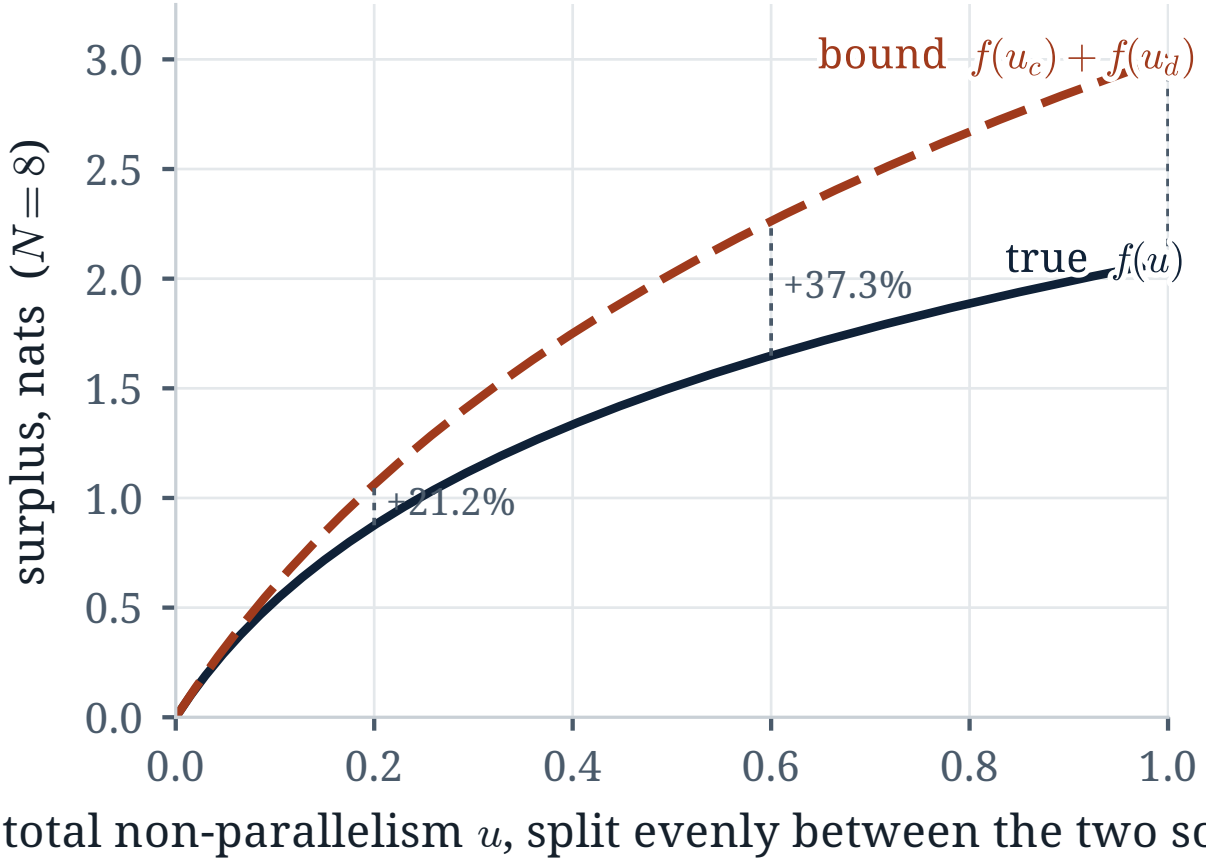


Figure 10.3: **The additive bound overstates the true value by up to 44.7 per cent, and the error is largest exactly where the ambition is highest.** At $N = 8$ with the two sources split evenly, the true surplus $f(u_{\text{cross}} + u_{\text{div}})$ is plotted against the impossibility limit $f(u_{\text{cross}}) + f(u_{\text{div}})$; the annotations give the overstatement at three points. **The gap is a consequence of strict concavity and is not a modelling artefact.** Model output at $N = 8$; arithmetic exact. Generated by `figs/src/CCP-010_fig02.py`.

10.4.5 4.5 What the upper bound actually licenses

The upper bound is stated so often as "subadditivity" that its practical content is easy to lose. Three statements, each forbidding a specific real mistake.

Statement 1 — no synergy is promised. The theorem gives a ceiling and no floor above the max. Any memo arguing that cross-disciplinary staffing *and* structured dissent together "will be greater than the sum of its parts" makes a claim the mathematics does not support and in fact bounds against. **Honest design assumes additivity at best and plans against the true concave value.**

Statement 2 — the true value is systematically below the sum. Strict concavity (strict concavity [Curzi, 2026]) makes the inequality strict whenever both components are nonzero. The overstatement table in Section 3.3 quantifies it. The correct quantity is `f(u_cross + u_div)`; `f(u_cross) + f(u_div)` is never the estimate.

Statement 3 — the bound is not tight, so it is not an estimate. A loose upper bound is a proof of impossibility, not a forecast: it states where surplus *cannot* be, and nothing about where it is.

A fourth statement, badged as this paper's interpretation rather than a result: **because the deployed functional is a lower bound on the ensemble's information, a group may realize benefits from combining the two sources that this instrument does not credit.** A reason to build such groups; not a reason to write the uncredited benefit into a plan. The floor is logged and spent against.

10.4.6 4.6 The lower bound is the floor that can actually be spent

```
f(u) >= max( f(u_cross), f(u_div) )
```

The lower bound has a use the upper bound does not: it is a *guarantee*, and guarantees are what may be committed to. Its tightness varies, and knowing where it is tight establishes when the second source is worth the trouble. The gain of the true value over the best single source, at `N = 8`:

```
u     u_cross  u_div   max f (floor)   f(u) true   gain over best single
-------------------------------------------------------------------------
1.00   0.500   0.500      1.5041        2.0794     +0.5754   (+38.3%)
1.00   0.700   0.300      1.7750        2.0794     +0.3045   (+17.2%)
1.00   0.900   0.100      1.9879        2.0794     +0.0916   (+4.6%)
1.00   0.980   0.020      2.0618        2.0794     +0.0177   (+0.9%)
1.00   1.000   0.000      2.0794        2.0794     +0.0000   (0.0%)
```

The bound is tight exactly when the portfolio is lopsided, and slack exactly when it is balanced — which makes the lower bound a direct readout of how unbalanced the two sources are. At a 50/50 split the true value exceeds the guaranteed floor by 38 percent; at 98/2 it exceeds it by under one percent, which is to say the minority source has bought nothing that was not already present.

That gives a diagnostic requiring no new instrumentation:

```
BALANCE INDEX

  B := [ f(u) − max( f(u_cross), f(u_div) ) ] / f(u)

  B near 0      one source does all the work; the other is decoration
  B in 0.1–0.3  a real but lopsided portfolio
  B > 0.3       both sources materially loaded
```

For the rows above: B = 0.2767 at 50/50, 0.1464 at 70/30, 0.0440 at 90/10, 0.0085 at 98/2. Compute B every audit cycle — it needs no data beyond the decomposition already computed, and it tells a board whether its diversity programme has one leg or two.

10.4.7 4.7 The classic pathology: rich in crossing, starved of divergence

This is the central practical diagnosis of Section 4, and it deserves to be stated without hedging. A very large number of contemporary organizations — institutes, consultancies, agencies, foundations, product teams, review panels, standards bodies — have executed a genuine and expensive programme to raise `u_cross` and simultaneously, usually without noticing, driven `u_div` toward zero. The mechanism is not conspiratorial. It is administrative, and every step is locally reasonable.

STEP	ON u_cross	ON u_div
1. Recruit deliberately across disciplines	UP	neutral
2. Onboard everyone through one induction	neutral	DOWN
3. Adopt one shared framework/vocabulary	neutral	DOWN
4. Tie all bonuses to one composite metric	neutral	DOWN (hard)
5. Promote the people who integrate well	neutral	DOWN
6. Report diversity as headcount by category; never report disagreement	reported UP	invisible
7. Escalate disputes to one decision-maker	neutral	DOWN (hard)

Steps 2 through 5 are the things done to make a cross-disciplinary team *function*: its `C` term (the cost analysis [Curzi, 2026b]) is enormous because translation is expensive, and shared vocabulary, shared metrics and clear escalation are the cheapest ways to cut it. **The organization pays for its `u_cross` by spending its `u_div`, and the transaction never appears on any ledger because only one of the two quantities is measured.**

The signature is unmistakable once it is looked for. The org chart shows six disciplines and the last forty decision memos show no recorded dissent. Everyone can explain why their discipline matters, and no two people in the same discipline have ever publicly disagreed about a substantive call. Reviews described as “robust” produce, on inspection, a 100 percent approval rate with comments. “We’re all aligned on this” is a report of success. New arrivals from outside the dominant discipline stop objecting within four to eight months, and this is described as *settling in*.

Run the arithmetic at `N = 8`, against three alternative portfolios:

`N = 8,  f(u) = ln(1 + 7u)`

portfolio	u_cross	u_div	u	f(u)	N_eff	floor	ceiling
diverse hires, one opinion	0.750	0.000	0.750	1.8326	6.250	1.8326	1.8326
+ a dissent protocol	0.750	0.050	0.800	1.8871	6.600	1.8326	2.1327
monodiscipline, real fight	0.000	0.750	0.750	1.8326	6.250	1.8326	1.8326
balanced	0.375	0.375	0.750	1.8326	6.250	1.2879	2.5757

Three of these four rows have identical surplus. That is the law being honest: it counts non-parallelism and does not care about its provenance. But look at the floor and ceiling columns. The first row is *pinned* — floor equals true value equals ceiling, because one component is zero. No slack, no second source, no redundancy in its difference: a single-channel instrument with six labels on it. Remove its one source and it holds nothing. The balanced row has the same output today and, if either source were removed, would still hold 1.2879.

A portfolio with the same surplus and two live sources is strictly more robust than a portfolio with the same surplus and one, and the surplus number alone cannot tell them apart. This is the result of Albert et al. [2000] in another domain: a structure concentrated on one hub is robust to random failure and fragile to targeted attack on that hub. A monoculture of opinion, however cross-disciplinary its headcount, has exactly one hub.

10.4.8 4.8 The mirror pathology: the monoculture that argues

The opposite failure is rare in large institutions and endemic in small expert ones: `u_div` high, `u_cross` near zero. Fourteen economists who genuinely and vigorously disagree. A surgical team that argues every case. Twelve cryptographers who are each other's fiercest critics. `gamma` large, `beta^2 = 0.02`.

```
u_cross = 0.02, u_div = 0.98 * sin^2(70 deg) = 0.98 * 0.8830 = 0.8654
u = 0.8854,  N=8: f(u) = ln(7.1978) = 1.9738,  N_eff = 7.198
```

Excellent surplus — about a strictly bounded set of questions, because no component of anyone's state lies outside the block. Every question the block does not contain returns `u = 0`: not low surplus, *no* surplus, because there is no non-parallel component to register it.

The signature: fierce, high-quality internal debate and a reputation for rigour; decisions impeccable within the frame and catastrophic outside it; failures always described afterwards as "outside our remit"; recruiting that selects for people who can survive the internal debate, which is a filter for people already inside the block. the law of requisite variety [Ashby, 1956] applied and failed — variety deep in one dimension and zero in all others absorbs only the disturbances that arrive along that dimension.

The fix is not more argument. `du/d(sin^2 gamma) = alpha^2 = 0.98` here, which is high, but `sin^2(gamma)` is already 0.883 and can rise at most to 1.000 — worth `f(1.00) - f(0.8854) = 2.0794 - 1.9738 = 0.1056` in total. Whereas `du/d(beta^2) = cos^2(gamma) = 0.117` looks small, but `beta^2` can rise from 0.02 to anywhere. Rule 4.4 fires: buy the scarce source.

10.4.9 4.9 Why institutions systematically buy crossing and not divergence

Four structural reasons, each with a named counterpart elsewhere in this paper.

(1) Observability asymmetry. `u_cross` is a property of the roster; `u_div` is a property of behaviour over time, and only one is auditable from a spreadsheet. Star and Ruhleder [1996] apply exactly: functioning dissent is invisible, and its absence becomes visible only when the organization walks into a wall it had six people qualified to see.

(2) Goodhart capture [Goodhart, 1975]. Once `u_cross` is a reported target, the cheapest way to raise the report is to relabel categories rather than acquire orthogonal ones (the ceiling analysis [Curzi, 2026c]). `u_div` is harder to capture that way, which is precisely why nobody measures it: it produces uncomfortable numbers that administrative action alone cannot improve.

(3) The cost term punishes divergence directly. Internal dissent raises `C_t` immediately and visibly: meetings run longer, decisions take another cycle. Cross-hiring also raises `C_t`, but that increase is booked as onboarding and treated as investment. Same term, opposite political treatment.

(4) Conway's law [Conway, 1968]. An organization that resolves disputes by escalation builds systems with single points of arbitration; one running on shared vocabulary builds systems with one ontology. The architecture then reinforces the geometry that produced it, and `gamma` falls further.

A fifth is not structural but worth naming: **permitting divergence requires an unglamorous kind of managerial nerve, and it is the one lever of the two that cannot be delegated to a recruiting function.** Badged as observation, not result.

10.4.10 4.10 Instruments — measuring the two components separately

The decomposition is only useful if `beta^2` and `gamma` can be estimated separately from real data. Both can. Every estimator below is badged: these are operational choices in the sense of the state-encoder requirement [Curzi, 2026], not consequences of the theorem.

Instrument A — the evaluability partition (estimates beta^2)

For each ordered pair `(x, y)`, sample `y`'s recent work products and have `x` classify each item:

```
E  : "I can evaluate this on its merits."
N  : "I can read it but cannot judge whether it is right."
```

Then

```
beta^2_hat(x, y)  =  ( count of N ) / ( count of E + count of N )
```

Notes and cautions. The estimator is directional: `beta^2(x,y)` need not equal `beta^2(y,x)`, and in hierarchical organizations it systematically does not. Use the mean for pair-level `u_cross` and *log the asymmetry separately* — it is itself a finding (Section 4.13). Classify against actual artifacts, not job titles: asking a participant whether they can evaluate an epidemiologist's work measures self-image, while handing someone three anonymized memos measures evaluability. And anonymize provenance, because a known author contaminates the E/N call with deference.

Instrument B — the shared-ground correlation (estimates gamma)

Restrict to items both parties classified `E`. Have each score them independently and *before any discussion* — the ordering is not optional; see Chapters 11 and 22. Compute the correlation `r` of the two score vectors. Then

```
sin^2(gamma)_hat = 1 − r^2
u_div_hat        = alpha^2_hat ∗ (1 − r^2)  =  (1 − beta^2_hat)(1 − r^2)
```

Badge: `1 − r^2` is the standard coefficient of non-determination, and reading it as `sin^2(gamma)` is natural when the two score vectors stand for the state vectors restricted to the shared block. It is a modeling choice — defensible, the one this paper makes, and declarable in any preregistration (the preregistration protocol [Curzi, 2026i]).

Cautions.

- `r` must be computed on *independent* scores. A correlation measured after discussion measures the discussion, not the participants. This single error accounts for most inflated agreement statistics in practice.
- Use a sufficient item count. Below about 12 shared-ground items the standard error on `r` swamps the estimate. The item count is reported with every `gamma`.
- Negative `r` gives the same `1 − r^2` as `+|r|`: the estimator measures angle, not direction. **That is a genuine trap.** Two participants who systematically take opposite positions are *parallel* in this geometry (`x . y = −1`, `u = 0`), and the professional contrarian therefore buys nothing. If `r` sits near `−1` rather than near `0`, there is no divergence, only a mirror. See Failure Mode FM-4F. Antiparallel is parallel.

Instrument C — the ensemble roll-up

The theorem governs a pair; aggregating to a group is an operational choice.

```
u_cross_hat = mean over all ordered pairs of beta^2_hat
u_div_hat   = mean over all ordered pairs of alpha^2_hat (1 − r^2)
u_hat       = u_cross_hat + u_div_hat
```

Badge, and this matters: **the mean-over-pairs roll-up is a convenience, not a theorem** — a summary statistic, adequate for design and inadequate for a claim of proof. the survey protocol [Curzi, 2026i] requires the choice stated; the preregistration protocol [Curzi, 2026i] requires it frozen before fitting.

Instrument D — the disagreement log (a delta instrument)

Levels are not enough: as Section 5 establishes, `u_div` decays. Keep a standing log:

```
For each substantive decision, record:
  − was a dissenting position stated on the record?  (Y/N)
  − by whom, and in which block?
  − from INSIDE the deciding block (u_div) or a different one (u_cross)?
  − was the dissent's substance addressed in the decision record?
```

The fourth field is the one that matters. Dissent recorded and not addressed is theatre and scores zero. the conversion analysis's [Curzi, 2026b] `eta` is the relevant term: **a disagreement that leaves no trace in the artifact converted none of its surplus into stored order, and `eta * S` for that interaction was zero regardless of how vigorous the meeting felt.**

10.4.11 4.11 Where the two components sit in the storage law

Nothing here changes the storage law. Both components feed the same scalar:

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,     S_t = f(u_cross,t + u_div,t)
```

But the components have different *dynamics*, which is why they must be tracked separately even though the law only ever sees their sum.

`u_cross` decays slowly — set by the roster, moving at attrition timescales, so an annual recruiting and rotation review suffices. `u_div` decays fast, with a half-life in months (Section 5 computes it), and needs continuous protocol: independent-estimate-first, a dissent log, a red-team rotation.

The practical consequence for the headroom gauge (the headroom analysis [Curzi, 2026b]), `h_t = eta ln N − delta Q_t − C_t`: it does not contain `u` at all. `u` determines how much of the available headroom is converted, via `S_t = f(u) <= ln N`. **A collapse in `u_div` does not show up in the headroom gauge; it shows up as a persistent gap between `eta S_t` and `eta ln N` — between what the system could take in and what it does.** Define and track it:

```
UTILIZATION OF CEILING

  rho_ceil,t := S_t / ln N = f(u_t) / ln N,  in [0, 1]

  rho_ceil = 1   u = 1, orthogonal, running at the per−interaction ceiling
  rho_ceil = 0   u = 0, parallel, no surplus at any headcount
```

At `N = 6` (`ln 6 = 1.7918`): an organization at `u = 0.75` runs at `1.5581 / 1.7918 = 0.8696`; at `u = 1.00` it runs at 1.000. That 13 percent of unused ceiling is the entire prize in the worked example below, and it is invisible on every gauge except this one.

10.4.12 4.12 Semilattice, not tree

Alexander [1965] is the architectural statement of this section's structural claim. A tree partitions — every element in exactly one parent set, no overlap. A semilattice permits overlap.

A **tree** assigns each person exactly one block, so any pair is either in the same block (`beta^2 = 0`, all difference from `gamma`) or in different blocks (`beta^2 = 1`, `alpha^2 = 0`, so `u_div = 0` identically — *no shared ground* on which to disagree). **A strict tree forces every pair into a corner of the budget: pure crossing or pure divergence, never a portfolio**, every pair pinned and every balance index zero. A **semilattice** lets people hold partial membership in several blocks, so `0 < beta^2 < 1` for most pairs, both components are live, and B is positive across the ensemble.

This is why matrix structures, dual reporting, rotating secondments and shared instrumentation carry more surplus than the same headcount in a clean hierarchy — and why they cost more in `C`. the cost arithmetic [Curzi, 2026b] and this section's are in tension by construction; the intervention procedure [Curzi, 2026j] is where the trade is made explicitly.

10.4.13 4.13 The asymmetry finding

Instrument A's directional `beta^2` is diagnostic in its own right. Define `A(x,y) := beta^2_hat(x,y) − beta^2_hat(y,x)`. `A` near zero means mutual evaluability. Large positive `A` for a consistently-appearing `x` means *x reports being unable to evaluate others' work while others report being able to evaluate x's* — which has two very different causes. Either a genuine specialization gradient (`x` works deep and narrow), or **deference collapse**: `x` is junior and has learned not to claim evaluative standing, which is `u_div` being destroyed by rank and reappears in the seating analysis [Curzi, 2026f] as the head-of-table effect. Distinguish by anonymizing the artifacts. If `A` collapses when authorship is hidden, the asymmetry was deference. Cheapest decisive test in Section 4; it costs one afternoon.

10.4.14 4.14 Specification: measuring and designing the two sources

Requirements are numbered `SPEC 4.n` and each is checkable. A conforming system satisfies all of Group A, and satisfies Group B or records a dated exception with a named owner.

Group A — measurement (mandatory)

SPEC 4.1 The survey (the survey protocol [Curzi, 2026i]) shall enumerate the `N` blocks in writing before any `u` estimate is computed, declared as *distinct methods or evaluative standards*, not job titles or demographic labels.

SPEC 4.2 `u_cross` and `u_div` shall be estimated and reported separately. A report containing only a total `u` is non-conforming.

SPEC 4.3 `beta^2` shall be estimated by the evaluability partition (Instrument A) against real, anonymized artifacts. Self-report against job titles does not conform.

SPEC 4.4 `gamma` shall be estimated from independent scores collected *before* any group discussion of the items. Post-discussion correlation is non-conforming and shall be labeled as such if reported.

SPEC 4.5 Every `gamma` estimate shall carry its shared-ground item count `n_E`. Estimates with `n_E < 12` shall be marked provisional.

SPEC 4.6 The pair-to-ensemble roll-up rule shall be stated in the report and shall not change between periods without a dated note.

SPEC 4.7 The directional asymmetry `A(x,y)` shall be computed and reported. Any participant whose mean `A` exceeds 0.25 shall trigger the anonymization test of Section 4.13.

SPEC 4.8 The balance index `B = [f(u) − max(f(u_cross), f(u_div))] / f(u)` shall be computed each period and reported alongside `u`.

SPEC 4.9 The ceiling utilization `rho_ceil = f(u) / ln N` shall be computed each period and reported alongside `h_t`.

Group B — design (mandatory unless exception recorded)

SPEC 4.10 No plan shall use `f(u_cross) + f(u_div)` as an estimate of surplus. Plans shall use `f(u_cross + u_div)`. The upper bound may appear only as a stated impossibility limit.

SPEC 4.11 No plan shall assert synergy beyond additivity. A benefit believed to exist beyond it shall be recorded as an uncredited expectation and excluded from the numbers.

SPEC 4.12 The first increment of `u` in any period shall come from whichever source is scarcer, per Rule 4.4, unless a written justification records why the more abundant source was chosen.

SPEC 4.13 Where `B < 0.05` for two consecutive periods, the system shall be recorded as single-sourced and a remediation owner named. `B < 0.05` at high `u` is not a healthy state; it is a concentrated one.

SPEC 4.14 Every substantive decision record shall carry a dissent field with four entries: stated (Y/N), by whom, source (cross or div), addressed in the artifact (Y/N). A blank field is non-conforming; "none" conforms and is itself data.

SPEC 4.15 Dissent recorded but not addressed in the artifact shall be scored `eta = 0` for that interaction and shall not count toward `u_div` maintenance.

SPEC 4.16 Independent estimates shall precede discussion wherever `u_div` is a design input. The ordering is a hard requirement; reversing it destroys both the quantity being measured and the quantity being used.

SPEC 4.17 Onboarding, vocabulary standardization, and composite incentive metrics — the standard `C`-reduction moves — shall carry an explicit `u_div` impact estimate before adoption. Cutting `C` by spending `u_div` is a legitimate trade and an illegitimate accident.

SPEC 4.18 No participant shall be required to disagree. Divergence protocols shall create the *opportunity* and the *protection* to dissent, never an obligation: an intervention removing the ability to decline is out of scope for this paper (the framework's consent clause [Curzi, 2026]). Compelled dissent is also useless — it produces low-information objections that inflate the log without moving `r`.

SPEC 4.19 Under a strict hierarchy, at least one crosscutting structure (rotation, community of practice, shared instrument, dual assignment) shall exist, so that some pairs have `0 < beta^2 < 1`. A pure tree pins every pair at a budget corner (Section 4.12).

SPEC 4.20 The contrarian check shall run each period: any pair with shared-ground `r` below `-0.6` shall be flagged as a possible mirror rather than a divergence, and re-examined. Antiparallel is parallel.

Group C — reporting

SPEC 4.21 The period report shall contain at minimum:

```
PERIOD: ______   N (declared blocks): ____
u_cross: ______ (n pairs ____)   u_div: ______ (mean n_E ____)   u: ______
N_eff = 1+(N-1)u: ______   S = f(u): ______   ln N: ______
rho_ceil = S/ln N: ______   B (balance index): ______
floor = max(f(uc),f(ud)): ______
ceiling = f(uc)+f(ud): ______   [impossibility limit only]
decisions with recorded dissent: ____/____
dissents addressed in artifact:  ____/____
mean directional asymmetry A: ______
```

SPEC 4.22 Any use of axiom S3 (the affine effective count) shall carry its badge in the same sentence. `N_eff` is a design choice, not a derived fact.

SPEC 4.23 The report shall state, verbatim: *the surplus functional as deployed is a lower bound on the ensemble's information, not the information itself.*

10.4.15 4.15 Failure modes of the decomposition

Each entry: name, clinical description, damaged parameter, signature, measurement, fix. These feed the catalog in the consolidated failure catalogue [Curzi, 2026j].

FM-4A · THE FURNISHED MONOCULTURE

Clinical. High `u_cross`, `u_div` at or near zero. Every discipline present; each represented by exactly one permitted view.

Damages. `u_div` -> 0, pinning `f(u)` at `f(u_cross)` and driving `B` to zero. Robustness collapses while surplus still reads well.

Signature. Blank or "none" dissent fields across consecutive decisions; approval near 100 percent with substantive comments; "aligned" used as a success report; outsiders stop objecting within two quarters, and this is praised.

Measurement. Instrument B on within-block pairs: `r > 0.9`. `B < 0.05`. Fewer than one recorded-and-addressed dissent per ten decisions.

Fix. Rule 4.4: buy the scarce source. Independent-estimate-before-discussion (SPEC 4.16); a rotating, protected adversarial role; a non-blank dissent field (SPEC 4.14); an artifact recording how each dissent was addressed (SPEC 4.15). Expect `C` to rise; budget it against `h_t` first.

FM-4B · THE DEEP NARROW WELL

Clinical. High `u_div`, `u_cross` near zero. Vigorous argument inside a single block; no component of anyone's state outside it.

Damages. `u_cross` -> 0. a violation of requisite variety [Ashby, 1956]: variety in one dimension, none in the others.

Signature. Reputation for rigour; failures always classified as "outside our remit"; recruitment selects for people who survive the internal debate, which selects for insiders.

Measurement. Instrument A: `beta^2` below 0.10 for nearly all pairs.

Fix. Import a practitioner from an unrelated trade with real standing, not an advisory seat; give them work products to evaluate, not presentations to attend. At large `gamma` the first cross-hire buys less than expected, since `du/d(beta^2) = cos^2(gamma)` is small. Buy anyway: the ceiling on further `gamma` is close, the ceiling on `beta^2` is not.

FM-4C · THE ADDED-BOUND PLAN

Clinical. A plan or forecast computing surplus as `f(u_cross) + f(u_div)`.

Damages. Nothing in the system; everything in the forecast. Overstates by 16 to 45 percent (Section 3.3), worst where the plan is most ambitious.

Signature. Two surplus numbers in a document, added. Language about the combination being "greater than the sum of its parts."

Measurement. Recompute with `f(u_cross + u_div)`. The gap is the overstatement.

Fix. SPEC 4.10 and 10.11. Retire the upper bound from planning use entirely.

FM-4D · THE C-FOR-DIVERGENCE TRADE MADE BY ACCIDENT

Clinical. An organization cuts coordination cost by standardizing vocabulary, induction, incentives and escalation, destroying `u_div` without deciding to.

Damages. `C` down (good), `u_div` down (unaccounted). The net effect on `eta S_t − delta Q_t − C_t` is often positive in the short run and negative once `S_t` responds, because the low-pass lag (the circuit analysis [Curzi, 2026a]) delays the `S` decline past the point where anyone attributes it to the change.

Signature. A cost-efficiency programme completed 6 to 18 months before an unexplained decline in decision quality — unsuspected, because it "worked."

Measurement. Track `u_div` and `C` on one time axis: a `C` step down followed by a `u_div` slope down. the preregistration protocol's [Curzi, 2026i] shuffled-input control separates it from coincidence.

Fix. SPEC 4.17. Trade deliberately or not at all.

FM-4E · THE RELABELED BLOCK

Clinical. `N` raised on paper by splitting an existing category rather than adding an orthogonal one.

Damages. `N`, `beta^2`, and `u_cross` all inflated; every downstream number wrong in the optimistic direction.

Signature. Block boundaries following the reporting line rather than the method; blocks whose members trained in the same programme; `beta^2` near zero for pairs the declaration puts in different blocks.

Measurement. Cross-tabulate declared block against measured `beta^2`. If cross-block `beta^2` is not materially higher than within-block, the declaration is fiction.

Fix. Re-declare blocks by evaluative standard (SPEC 4.1). A block is defined by *what counts as a good answer*, not by who reports to whom.

FM-4F · THE MIRROR MISTAKEN FOR A WINDOW

Clinical. A participant who reliably opposes, credited as the source of `u_div`. In the squared geometry, systematic opposition is parallelism: `x . y = −1` gives `u = 0`.

Damages. `u_div` credited where none exists; worse, a designated opposer suppresses dissent from others, who assume the slot is filled.

Signature. One person objects to everything, and the objection carries no information because it is predictable from the proposal. `r` near `−1`.

Measurement. SPEC 4.20. Also: predict the objector's position from the proposal alone. If that can be done reliably, they are a mirror.

Fix. Rotate the adversarial role rather than appointing it (the non-parallelism analysis [Curzi, 2026c]); score dissent by whether it changed the artifact. `r` near 0 is the target.

FM-4G · DEFERENCE COLLAPSE

Clinical. Directional asymmetry `A` driven by rank rather than specialization. Juniors report being unable to evaluate senior work; seniors report the reverse.

Damages. `beta^2` overestimated in one direction — silence recorded as crossing — and `u_div` destroyed, the party who would disagree having disclaimed standing to do so.

Signature. `A` correlates with seniority across the roster.

Measurement. Section 4.13's anonymization test.

Fix. Anonymize artifacts at the point of evaluation. the seating analysis's [Curzi, 2026f] protocols — last-speaks-first for rank, silent-write-first, no head of table — attack the same fault from the geometry side.

FM-4H · THE UNADDRESSED DISSENT LOG

Clinical. Dissent recorded diligently and never changing an artifact. The log grows, `Q` does not.

Damages. `eta -> 0` for the divergence channel: all the `C` of the argument, none of the storage.

Signature. A healthy dissent-stated rate and a near-zero dissent-addressed rate. Participants describe the process as going through the motions.

Measurement. The two counts in SPEC 4.21's report block, side by side. Their ratio is the divergence channel's `eta`.

Fix. SPEC 4.15. A the conversion analysis [Curzi, 2026b] intervention, not a Section 4 one: the geometry was fine, the conversion was zero.

10.5 5 Pseudo-diversity and the decay of non-parallelism

This is the failure mode that costs modern institutions the most, and it is almost never diagnosed, because every instrument the institution owns reports that nothing is wrong.

The parallel wall is not a shortage of people, talent, funding, effort, or good intention. It is a geometry in which many capable people, each doing competent work, collectively return the surplus of one. The arithmetic that prices it is short enough to fit on a card, and Section 5's first job is to put it there.

Its second job is harder. `u` is not a property an organization acquires and then owns. **It is a quantity that decays, and that makes non-parallelism a maintenance obligation, not a procurement one.** A group can be assembled correctly, measure well on day one, and be parallel within two years without a single person leaving, a single policy changing, or anyone making a decision they would recognize as wrong.

10.5.1 5.1 The arithmetic of zero

From Curzi [2026], Thm 2.1 and the effective count that follows from axiom S3:

```
N_eff = g(u) = 1 + (N−1)u
S     = F(u) = f(u) = ln(1 + (N−1)u) = ln N_eff

at u = 0:   N_eff = 1 + (N−1)(0) = 1
            S     = ln 1 = 0
```

That is the whole indictment. **Three glasses ground from one doctrine are one glass, and their agreement is not evidence of anything except that they were ground from one doctrine.**

Curzi [2026], Thm 4.3 (extremality) makes it exact and two-sided: `F = 0` if and only if the states are parallel, and `F = ln N` if and only if they are orthogonal. There is no gradation at the bottom. Parallelism does not produce a small surplus; it produces none.

Note also what the law does *not* say. It does not say the group produces nothing. Each member still does work. What returns zero is the *interaction surplus* — the part of the output that exists because there is more than one of them. **A parallel group of forty is not a bad group. It is a group of one, paid forty times, and audited as though it were forty.**

10.5.2 5.2 What "parallel" means operationally

`u = 1 − (x . y)^2`. Two states are parallel when their normalized inner product has magnitude 1 — when one is a scalar multiple of the other. In an organization the operational readings are:

Formal	Operational reading
`x . y = +1`	the second party's judgment is predictable from the first's, with the same sign
`x . y = −1`	the second party's judgment is predictable from the first's, with the sign flipped
`u = 0`	in either case: knowing the first determines the second, so the second carried no new information

The second row is the trap Section 4 named (Failure Mode FM-4F) and it bears repeating here because pseudo-diversity programmes frequently buy it by mistake. **Antiparallel is parallel.** A designated opposer whose position is derivable from the proposal is a mirror, not a window, and the law prices a mirror at zero.

10.5.3 5.3 The price of pseudo-diversity, in numbers

Take a declared roster of `N` categories. Compute what the roster is actually worth at various levels of measured non-parallelism.

	N = 5		N = 10		N = 20		N = 40	
u	N_eff	S	N_eff	S	N_eff	S	N_eff	S
0.00	1.000	0.0000	1.000	0.0000	1.000	0.0000	1.000	0.0000
0.01	1.040	0.0392	1.090	0.0862	1.190	0.1740	1.390	0.3293
0.02	1.080	0.0770	1.180	0.1655	1.380	0.3221	1.780	0.5766
0.05	1.200	0.1823	1.450	0.3716	1.950	0.6678	2.950	1.0818
0.10	1.400	0.3365	1.900	0.6419	2.900	1.0647	4.900	1.5892
0.20	1.800	0.5878	2.800	1.0296	4.800	1.5686	8.800	2.1748
0.40	2.600	0.9555	4.600	1.5261	8.600	2.1518	16.600	2.8094
1.00	5.000	1.6094	10.000	2.3026	20.000	2.9957	40.000	3.6889

Read the `N = 40` column. A forty-category roster at `u = 0.05` has an effective count of 2.95. **Forty declared categories, three effective ones — and the organization is paying for forty, sizing its meeting rooms for forty, and carrying the coordination cost of forty.** the cost analysis's [Curzi, 2026b] arithmetic completes the picture: at `N = 40`, the communication-path count [Brooks, 1975] communication paths number `n(n−1)/2 = 780`. The group carries 780 channels of cost to deliver the surplus of three.

That ratio — cost of the declared count, surplus of the effective count — is the entire economics of the parallel wall, and it is why pseudo-diversity is expensive rather than merely disappointing.

10.5.4 5.4 The steepest slope is at zero

Curzi [2026], Thm 4.2 gives the sharp Lipschitz constant `N − 1`. Differentiating:

```
f'(u) = (N−1) / (1 + (N−1)u)

f'(0)  = N − 1        <-- the maximum, attained exactly at parallelism
f'(1)  = (N−1)/N      <-- the minimum
```

The first unit of difference is the most valuable unit of difference there will ever be. Concretely:

N	f(0.01) − f(0)	f(1.00) − f(0.99)	ratio
3	0.0198	0.006689	3.0
5	0.0392	0.008032	4.9
10	0.0862	0.009041	9.5
20	0.1740	0.009545	18.2
40	0.3293	0.009798	33.6

At `N = 40`, the first hundredth of `u` is worth 33.6 times the last hundredth. At `N = 10`, moving `u` from 0.00 to 0.05 buys 0.3716; moving it from 0.95 to 1.00 buys 0.0460 — a factor of 8.07.

This cuts both ways, and the second cut is the one Section 5 is about. **Because the function is steepest at zero, a group near the parallel wall loses more surplus per unit of `u` lost than a group anywhere else on the curve.** The last increments of difference to disappear are the most expensive ones. Decay does not slow as it approaches zero; in surplus terms it accelerates.

10.5.5 5.5 u as a decaying state variable

Nothing in the founding surplus analysis [Curzi, 2026] makes u a function of time. the founding surplus analysis [Curzi, 2026] is static geometry. The storage law (the storage-law analysis [Curzi, 2026a]) evolves Q, and takes U_t as an *input*:

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,     S_t = F(U_t) = ln(1 + (N−1)U_t)
```

U_t is exogenous to the law. **Nothing in the mathematics says U_t stays where it was put, and everything in the empirical literature says it does not.**

This paper therefore adopts an explicit modeling postulate, badged here and carried through Section 5:

> **POSTULATE P-U (u-decay), badged.** In a group under ordinary communication, absent maintenance, measured non-parallelism decays geometrically: `U_{t+1} = (1 − lambda) U_t` with `lambda` in `[0,1)`.
>
> This is a modeling assumption, not a theorem. It is not in the founding surplus analysis [Curzi, 2026], not in the storage-law analysis [Curzi, 2026a], and not proved anywhere in this paper. It is adopted because it is the simplest form consistent with the cascade mechanism of Section 5.10, because it makes `lambda` estimable from two measurements, and because it is falsifiable in the manner of the preregistration protocol [Curzi, 2026i]. If measured u under zero maintenance does not decline approximately geometrically, replace the postulate.

With that badge in place, the structural claim of Section 5 becomes sayable in the framework's own terms:

```
delta  is the decay rate of STORED ORDER      (the retention analysis \citep{ccp:CCP-007})
lambda is the decay rate of DIFFERENCE        (Section 5)

Both are leaks. Both require maintenance. Only one of them is
routinely measured, and it is not lambda.
```

u is a delta problem as much as a design problem, and an organization that has a maintenance regime for its documents and none for its disagreement has budgeted for one leak and ignored the other.

10.5.6 5.6 Badges and honesty clauses

- **Proven (Curzi [2026], Thm 2.1, 4.3):** `f(u) = ln(1 + (N−1)u)` under S1-S4; `F = 0` iff parallel, `F = ln N` iff orthogonal.
- **Proven (Curzi [2026], Thm 4.2):** sharp Lipschitz constant `N − 1`, attained at `u = 0`.
- **Design axiom, not derived (S3):** the affine effective count `N_eff = 1 + (N−1)u`. Every `N_eff` in Section 5 inherits that choice.
- **Postulated and badged (11.1):** geometric decay of u.
- **Honesty clause:** as deployed, the surplus functional is a **lower bound** on the ensemble's information, not the information itself. A group measured at `u = 0` on the chosen observables may hold information those observables do not see. It does not hold information the instrument *credits*, and stored order responds to what the instrument credits. **If the instrument reads zero and the analyst believes the group is not parallel, the correct response is to change the instrument, not to overrule it.**
- **Honesty clause:** The founding surplus analysis [Curzi, 2026] is pure mathematics and makes no empirical or metaphysical claim. the surplus range analysis [Curzi, 2026]: surplus is not identified with Shannon entropy or mutual information in general.

- **Refused (the framework's consent clause [Curzi, 2026]):** nothing in this paper recommends compelling anyone to hold, express, or perform a view. Every instrument here measures; every intervention here creates opportunity and protection. A person who declines to dissent has exercised a right, and their decline is recorded as data, not as a failure.
- **Scope note, stated once and meant.** Section 5 treats diversity as an engineering parameter of an information-processing system. It makes no claim about the social, moral, or legal case for any staffing policy, and it takes no position on any of them. A reader who imports a political argument into this section has imported it; it is not here. **The fault named here is that a measurable quantity is going unmeasured, and the fix is to measure it.**

10.5.7 5.7 Three glasses ground from one doctrine

Three optical instruments are pointed at the same object. They agree, and the agreement is reported as convergent evidence.

But all three were ground on the same machine, to the same prescription, by the same house, and calibrated against the same reference. Their agreement is a statement about the machine. It contains exactly as much information about the object as one instrument does, and the confidence gained from the "three" is manufactured.

In the law: `x1 = x2 = x3` up to scale, so `u = 0` for every pair, `N_eff = 1`, `S = ln 1 = 0`. **The confidence went up and the information did not, and that gap is what pseudo-diversity produces as its characteristic output.**

Condorcet's jury theorem says the same from the other side: collective accuracy exceeds individual accuracy *only under independence*. Independence is not a nicety attached to the theorem, it is the theorem's engine, and dropping it drops the conclusion. Section 5.15 computes how fast.

10.5.8 5.8 The five species

Pseudo-diversity is not one fault but five, needing different instruments to detect and different interventions to fix.

Species 1 — declared variety with epistemic uniformity. The roster varies along one or more declared axes; the *evaluative standards* do not. Every member was trained to recognize the same things as good answers. The blocks were declared along an axis that is not the axis on which judgment varies, so `u_cross` measures high on the declared axis and near zero on the axis that governs the decisions being made.

Detection: declare blocks by evaluative standard (SPEC 4.1), then re-measure. If `beta^2` does not rise across the declared boundary, the boundary is not an epistemic boundary.

Species 2 — the echo chamber. Information intake is filtered through a common upstream source. Members differ in training and converge in evidence, because they read the same five outlets, attend the same three conferences, and follow the same feed. Difference in the members does not survive identical inputs.

Detection: audit the information supply chain, not the roster. Ask each member independently for the three sources that most changed their view in the past year; compute the overlap. High overlap is the signature.

Species 3 — correlated forecasts. Members produce independent-looking numerical judgments whose errors correlate. Each estimate is honestly made; the ensemble is nonetheless one estimate with noise. This species does most damage in risk, finance, planning, and engineering safety, because the aggregate is *reported* as an ensemble and *behaves* as a point.

Detection: error-decorrelation measurement, Section 5.16. This is the one species with a clean, cheap, quantitative instrument, and it is still rarely run.

Species 4 — committee theatre. Structure is elaborate — subcommittees, review stages, sign-offs, quorums — and the geometry flat, because every stage draws from the same pool and applies the same standard. Each stage adds `C` and no `u`: the process grows quadratically in cost and not at all in surplus.

Detection: count the stages, then count how often a later stage reversed an earlier one on substance. A reversal rate near zero across many decisions means the later stages are re-reading the first stage's instrument.

Species 5 — the common training pipeline. Reviewer, candidate, and appointment pools all drawn from a few upstream institutions. This is preferential attachment [Barabási and Albert, 1999] operating on careers: hubs attract, hubs train, and everyone downstream carries the hub's prior. The structure is efficient, robust to random loss, and — per Albert et al. [2000] — fragile to anything corrupting the hub, since everything downstream inherits it at once.

Detection: trace training lineage rather than affiliation. Two people at different institutions who trained under the same advisor are one lineage.

10.5.9 5.8.1 Manufactured opposition — a sixth fault, and not the Parallel Wall

The five species above all describe bodies that agree too much. This one describes a body that disagrees on schedule, and it was catalogued after the five because it does not look like them from outside. It is entered here as a distinct fault, with its own signature and its own fix, and it is recorded as failure mode FM-5I below.

The arithmetic is Section 5.2's second row, taken seriously. `u = 1 − (x . y)^2` depends only on `|rho|`, so **`rho = −1` returns `u = 0` exactly as `rho = +1` does.** Perfect opposition is parallel. The contrarian who reliably inverts a colleague carries the same information as the sycophant who reliably echoes them: none, because in both cases their position is a function of yours.

"Creative tension" at `rho` near -1 is therefore not a weak purchase of `u`. It is a zero purchase of `u` with a cost overhead, and that overhead is what makes it distinct from the Parallel Wall rather than a variant of it. A parallel roster wastes seats it already had. A designated-opposer seat is *added* — hired, seconded, or scheduled — to supply a divergence it cannot supply. Price the seat alone, at `eta = 0.35`, `delta = 0.10`, and a seat cost of 0.05 per period:

the added seat	`u`	`S`	`eta*S`	its `Q_inf` contribution
genuinely independent, `rho = 0`	1.0000	0.6931	0.2426	**+1.9260**
partially correlated, `rho = 0.6`	0.6400	0.4947	0.1731	+1.2314
perfect echo, `rho = +1`	0.0000	0.0000	0.0000	**-0.5000**
perfect opposition, `rho = −1`	0.0000	0.0000	0.0000	**-0.5000**

The last two rows are the same row. The opposer and the echo have identical consequences for stored order, and both are strictly worse than not filling the seat, because `C` is paid and nothing is returned.

The fault is worse than its arithmetic, for a reason that is not in the equation and is badged as observation rather than theorem: **the slot appears filled.** A body with a house sceptic believes it has bought dissent and stops looking for the independent judgment it actually needs,

so the manufactured opposition suppresses demand for the real thing. That is the same structure as the non-parallelism analysis's [Curzi, 2026c] captured red team (FM-9.3), arrived at from the opposite direction.

Detection: predict the objector's position from the proposal before the meeting. If the prediction is reliable, `|rho|` is near 1 and the seat is a mirror regardless of how heated the exchange is. Measure `|rho|`, never `rho`.

Fix: target `|rho|` near zero, not `rho` near -1. Rotate the role so it cannot become a stable inversion; score the seat by artifact change rather than by objection count. the geometry of the fold [Curzi, 2026c] gives the geometry — the angle table folded about 90 degrees — the dyadic table of the non-parallelism analysis [Curzi, 2026c] gives the dyadic table, and the dyad analysis [Curzi, 2026r] treats the two-person case in full.

10.5.10 5.9 The inference error, stated formally

Underneath all five species is a single reasoning mistake, worth stating in a form precise enough to check.

Agreement between two observers is evidence about the world **only to the extent that the observers could have disagreed.** The informativeness of concordance is a function of the prior probability of discordance; when that prior is zero, concordance carries no information.

```
CONCORDANCE TEST

  Before reading the second opinion, write down:
    p = P(this reviewer disagrees with the first, given the first)

  If p is near 0, the second opinion is not evidence, whatever it says.
  If p cannot be estimated, the agreement cannot be interpreted.
```

This test costs thirty seconds and can be applied to any second opinion, any peer review, any confirming analysis. Most institutions have never asked it about any of theirs. **The question is not "do they agree?" but "could they have disagreed?" — and the second question is the only one whose answer carries information.**

the non-parallelism analysis's [Curzi, 2026c] line is the same statement in the roster's language: a second opinion from someone whose training and incentives replicate the first is worth `ln(1) = 0`.

10.5.11 5.10 The cascade: how a diverse group becomes parallel

Everything so far describes a group that was parallel to begin with. The harder and more common case is a group that was genuinely diverse and became parallel while doing nothing wrong.

The mechanism is the information cascade of **Bikhchandani, Hirshleifer and Welch (1992)**, worth working through because it shows the collapse occurring through *individually rational* behaviour. No one defects, no one is coerced, and the group still ends up as one instrument.

The canonical setup: agents act in sequence on a binary question. Each holds a private signal, correct with probability `p > 1/2`. Each observes the actions of those who acted before, but not their private signals. Each acts to maximize their own accuracy.

Agent 1 has only their signal, so acts on it. Agent 2, seeing one action and holding one signal, acts on their own signal (breaking ties toward it). Agent 3 now faces two prior actions. If both agree, two units of evidence stand against Agent 3's one, and the rational act is to follow the crowd **and ignore their own signal.** From that moment, every subsequent agent sees only actions that carried no private information, so every subsequent action carries none either. The cascade is locked.

The arithmetic of the canonical model:

```
State A is true; signal accuracy p.
Agents 1 and 2 both follow their own signals.
  both correct   -> p^2          -> correct cascade
  both incorrect -> (1-p)^2      -> INCORRECT cascade
  mixed          -> 2p(1-p)      -> no cascade; the process restarts

 p      P(correct cascade)   P(incorrect cascade)   E[agents whose signal
                                                    entered the record]
-----------------------------------------------------------------------
0.55         0.5990               0.4010                   3.96
0.60         0.6923               0.3077                   3.85
0.70         0.8448               0.1552                   3.45
0.80         0.9412               0.0588                   2.94
0.90         0.9878               0.0122                   2.44
```

Two results deserve to be read slowly.

First, the incorrect-cascade probability. With individually 70-percent-accurate agents, the group locks onto the wrong answer 15.5 percent of the time, and once locked, adding more agents does not help — every additional agent copies. **A room of forty 70-percent-accurate people has a roughly one-in-six chance of being confidently, permanently wrong, and the confidence grows with the room size while the accuracy does not.**

Second, the informative-agent count. At `p = 0.70`, on average 3.45 agents' private signals ever enter the public record. Everyone after that is echo. In a room of forty, thirty-six and a half people are re-reporting three and a half. The declared `N` is 40. The information-bearing count is 3.45.

That is pseudo-diversity manufactured in real time, from a group that was genuinely diverse when it walked in.

10.5.12 5.11 Cascades collapse u, and that is the delta problem

The cascade result is normally taught as a decision-theory curiosity. In this framework it is the mechanism behind Postulate P-U, and the bridge is direct.

Before the cascade, participants' reported states differ because each reflects an independent private signal: `u > 0`. After the cascade, every reported state is a copy of the public history: `x_i = x_j` for all post-cascade `i, j`, so `u = 0` among them. Not small. Zero, by the extremality theorem [Curzi, 2026].

The key structural point is that **cascades do not need a new event to occur. They occur as a by-product of ordinary sequential communication, which is to say they occur continuously in any group that talks.** Every meeting in which the senior person speaks first, every review in which the second reviewer sees the first reviewer's score, every forecast round in which the group discusses before estimating, is a cascade opportunity. Most organizations present dozens per week.

That is why u decays rather than merely being low. And it is why the right frame is maintenance:

```
                 acquisition problem        maintenance problem
-------------------------------------------------------------------
question         *was difference hired?*  *is the difference still present?*
measured         at recruitment             every period, forever
owner            recruiting                 whoever runs the meetings
failure looks    a homogeneous roster       a heterogeneous roster that
like                                        produces homogeneous output
-------------------------------------------------------------------
```

Almost every organisation has the left column staffed and the right column unstaffed. **The left column is the one that photographs well.**

10.5.13 5.12 The decay table

Under Postulate P-U, `u_t = u_0 (1 − lambda)^t`, and the half-life is

```
t_half = ln 2 / ( −ln(1 − lambda) )
```

which is the same formula the retention analysis [Curzi, 2026a] gives for `delta`, because it is the same kind of leak.

```
lambda   half-life (periods)   u after 10 periods, from u_0 = 0.60
-------------------------------------------------------------------
 0.02          34.31                        0.4902
 0.05          13.51                        0.3592
 0.10           6.58                        0.2092
 0.15           4.27                        0.1181
 0.20           3.11                        0.0644
 0.30           1.94                        0.0169
 0.50           1.00                        0.0006
```

And the surplus consequence, at `N = 10`, `u_0 = 0.60`:

```
lambda = 0.15,  N = 10
 t     u        N_eff      S = f(u)     S as % of ln 10
-----------------------------------------------------------
  0   0.6000    6.400      1.8563           80.6%
  2   0.4335    4.902      1.5895           69.0%
  4   0.3132    3.819      1.3399           58.2%
  6   0.2263    3.037      1.1107           48.2%
  8   0.1635    2.471      0.9048           39.3%
 10   0.1181    2.063      0.7242           31.5%
 12   0.0853    1.768      0.5699           24.8%
```

Twelve periods — three years at quarterly cadence — takes a group from 80.6 percent of its per-interaction ceiling to 24.8 percent, with the same people, the same budget, and the same org chart. **Nothing in this table requires anyone to do anything wrong. It requires only that nobody does anything to stop it.**

Estimating lambda in a real system. Measure `u` at two times separated by `k` periods, with no `u`-maintenance intervention between them:

```
lambda_hat = 1 − (u_1 / u_0)^(1/k)
```

Two cautions, both mandatory. First, the interval must genuinely contain no maintenance, or what has been measured is the net of decay and repair, which cannot be separated. Second, `u` must be measured the same way both times — the same instrument, the same item pool, the same pre-discussion ordering (Section 4, SPEC 4.4). A change of instrument between the two measurements produces a `lambda` that is an artifact of the change.

10.5.14 5.13 Why the institution looks healthy after its input has collapsed

Here is the property that makes this failure mode so hard to catch in time, and it comes directly from the circuit analysis's [Curzi, 2026a] circuit identity.

The storage law is a first-order low-pass filter. `Q` responds to `S` with a lag of order `1/delta` periods. So when `u` — and therefore `S` — starts falling, `Q` does not immediately fall. It keeps rising for a while, on the momentum of the accumulated term, and then rolls over.

Worked, with `N = 5`, `eta = 0.55`, `delta = 0.09`, `C = 0.15`, `Q_0 = 4.00`, `u_0 = 0.55`, and `lambda = 0.18` (half-life 3.49 periods):

t	u	N_eff	S	eta S	delta Q	Delta Q	Q	
0	0.5500	3.200	1.1632	0.6397	0.3600	+0.1297	4.0000	
1	0.4510	2.804	1.0310	0.5671	0.3717	+0.0454	4.1297	
2	0.3698	2.479	0.9080	0.4994	0.3758	−0.0264	4.1751	<- peak
3	0.3033	2.213	0.7944	0.4369	0.3734	−0.0865	4.1488	
4	0.2487	1.995	0.6905	0.3798	0.3656	−0.1358	4.0623	
6	0.1672	1.669	0.5121	0.2817	0.3376	−0.2059	3.7511	
8	0.1124	1.450	0.3714	0.2043	0.2985	−0.2442	3.3166	
10	0.0756	1.302	0.2642	0.1453	0.2537	−0.2584	2.8185	
12	0.0508	1.203	0.1851	0.1018	0.2071	−0.2553	2.3015	
16	0.0230	1.092	0.0879	0.0484	0.1192	−0.2208	1.3240	
20	0.0104	1.042	0.0407	0.0224	0.0460	−0.1736	0.5110	
23	0.0057	1.023	0.0227	0.0125	0.0022	−0.1398	0.0248	

`u` peaks at `t = 0`. `Q` peaks at `t = 2`. **For two full periods after the input has begun to collapse, the institution's stored order is still rising, and every dashboard says the programme is working.** By the time `Q` visibly falls, `u` has already halved.

That lag is not a defect of the model. It is the correct behaviour of an accumulator, and it is the reason the input must be instrumented directly. **Watching `Q` to detect a `u` collapse is watching the reservoir to detect the drought.**

Two further readings from this run.

The headroom gauge does not see it. Compute `h_t = eta ln N − delta Q_t − C_t` along the same path: `h_0 = 0.3752`, `h_2 = 0.3594`, `h_8 = 0.4367`, `h_16 = 0.6160`. Headroom *rises* as the institution collapses, because `Q` is falling and `delta Q_t` shrinks with it. the headroom analysis's [Curzi, 2026b] gauge is doing its job — it reports available capacity, and a dying institution has plenty. **`h_t` measures what could be taken in; `rho_ceil = S_t / ln N` measures what is taken in. Only the second one detects a parallel wall.** Here `rho_ceil` goes 0.723 at `t = 0` to 0.055 at `t = 16`.

The model names its own breakdown. Admissibility requires `C_t <= (1 − delta)Q_t + eta S_t`. At `t = 23`, the right side is 0.0350 against `C = 0.15`. Admissibility fails, and beyond that the linear model reports negative `Q`, which is not a physical state. Read correctly: the arithmetic does not predict a negative institution. **It names the period at which the institution can no longer pay its own running cost, and stops.** the storage-law analysis [Curzi, 2026a], the fixed-point theorem [Curzi, 2026a] makes the same statement in equilibrium form: with `eta F(u*) < c*` the steady state is negative, and the theorem does not hide the failure mode. At `u = 0.08` here, `Q_inf = (0.55 * 0.277632 − 0.15)/0.09 = 0.030` — the regime is at the edge of unsustainable, and a little below that `u` it is over it.

10.5.15 5.14 Detection instruments

Five instruments, in increasing order of cost. Run at least two; they fail differently.

Instrument 1 — the collision test (the framework's collision instrument [Curzi, 2026], formalised). Put a live question to two members separately, in writing, with no prior discussion.

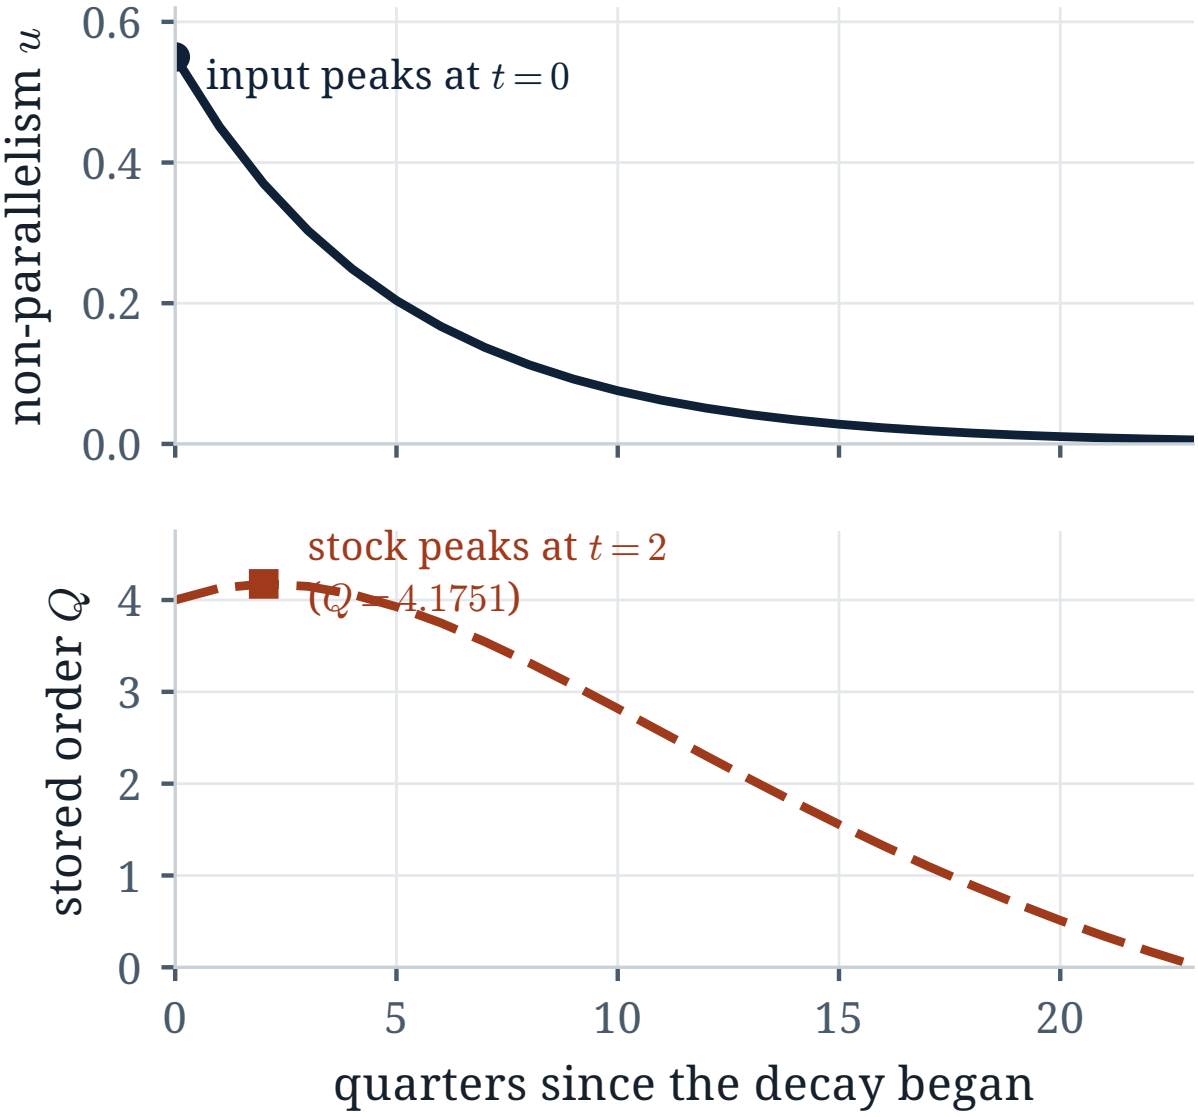


Figure 10.4: **The input peaks at $t = 0$ and the stock keeps rising until $t = 2$, which is why the collapse went unnoticed for years.** Upper panel: non-parallelism under $\lambda = 0.18$ per quarter. Lower panel: stored order under the storage law with $N = 5$, $\eta = 0.55$, $\delta = 0.09$, $C = 0.15$ and $Q_0 = 4.00$. Two panels rather than two axes, because the quantities have different dimensions. **Every dashboard available during those two quarters reported improvement.** Modelled with the source's stipulated parameters; arithmetic exact. Generated by `figs/src/CCP-010_fig04.py`.

Compare. Where the two accounts *collide* — genuinely conflict, not merely differ in emphasis — there is non-parallelism. Where they read as one account in two hands, there is none.

Score it: over `m` questions, the collision rate `c = collisions / m` is a crude but honest first estimate of the group's `u` on that question class. It is crude because it binarizes a continuous angle. It is honest because it cannot be faked by anyone who has not seen both answers.

Instrument 2 — the disagreement audit. Retrospective, cheap, and usually devastating. Pull the last 40 substantive decision records. Count:

```
decisions with any dissent recorded                    ____ / 40
dissents from inside the deciding discipline (u_div)   ____
dissents from a different discipline (u_cross)         ____
dissents whose substance changed the final artifact    ____
reversals of an earlier stage by a later stage         ____
```

The fourth line is the one that matters, per the conversion analysis [Curzi, 2026b]: dissent that leaves no trace in the artifact had `eta = 0` and converted nothing. The fifth line detects Species 4.

Instrument 3 — error-decorrelation measurement. The strongest instrument available, and the subject of Section 5.16. Requires a history of independent, resolvable predictions. Where it can be run it should be; it produces a number rather than an impression.

Instrument 4 — the pre-mortem. Prospective. Before a decision is finalized, state as fact that it failed badly two years on, and ask each participant *independently and in writing* why. Two properties make this a `u` instrument and not merely a risk exercise: it inverts the frame, so agreement with the group's prior no longer counts as a contribution; and it is written before discussion, so it precedes the cascade.

Read the output as a `u` measurement: if the independently written failure stories cluster onto one story, `u` is low on the risk axis. If they scatter across genuinely different mechanisms, `u` is live. **A pre-mortem that produces one story is not a reassuring pre-mortem. It is a failed measurement of the group.**

Instrument 5 — the shuffled-input control (the preregistration protocol [Curzi, 2026i]). The decisive test, and the one most often omitted. Fit the minimal estimation model `Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + e_t` on the real input series; then refit with the `U_t` series randomly permuted in time. If the shuffled series fits as well as the true one, the interaction geometry was not driving the order, and any story told about `u` in that system is unsupported. **A framework that cannot fail cannot be trusted, and this is the test that lets this one fail.**

10.5.16 5.15 What correlation does to Condorcet

Condorcet's theorem requires independence. Institutions habitually invoke the theorem's conclusion — bigger juries are more accurate — while operating in conditions that violate its premise. The size of the error is computable.

Here is an explicit toy model, defined in this paper and not cited from anywhere, so that its assumptions are visible. With probability `rho`, all `N` members receive and act on one common signal (correct with probability `p`), so they vote identically. With probability `1 − rho`, they vote independently, each correct with probability `p`. Majority rules.

```
p = 0.60

 rho     N = 3      N = 11     N = 51
-----------------------------------------
0.0     0.6480     0.7535     0.9265
0.2     0.6384     0.7228     0.8612
0.5     0.6240     0.6767     0.7632
0.8     0.6096     0.6307     0.6653
```

```
1.0     0.6000     0.6000     0.6000
```

At `rho = 0`, going from 3 to 51 members buys 28 points of accuracy. At `rho = 0.8` it buys 5.6. At `rho = 1.0` it buys nothing at all: fifty-one perfectly correlated 60-percent voters are a 60-percent voter.

Scale is a multiplier on independence, not a substitute for it, and a large correlated body is not a safer decision-maker than a small one — it is the same decision-maker with a larger travel budget. This is the ceiling analysis's [Curzi, 2026c] logarithmic ceiling meeting Section 5's parallel wall: raising `N` cannot rescue a system whose `u` is near zero, because `f(u) = ln(1 + (N−1)u)` and `u = 0` kills the `(N−1)` term outright.

10.5.17 5.16 The variance instrument

There is a second, entirely standard effective-count formula from classical statistics, and it is worth having because it is an *independent* instrument — it does not rest on the surplus law's axioms at all, so agreement between the two is meaningful.

Take `N` estimators of the same quantity, each with variance `sigma^2`, with equal pairwise correlation `rho`. The variance of their mean is

```
Var(mean) = sigma^2 ∗ [ 1 + (N−1) rho ] / N
```

Setting this equal to `sigma^2 / N_eff_var` gives the effective number of independent estimators:

```
N_eff_var = N / [ 1 + (N−1) rho ]
```

```
  N      rho=0.00   rho=0.10   rho=0.30   rho=0.50   rho=0.80   rho=1.00
-------------------------------------------------------------------------
   3       3.00       2.50       1.88       1.50       1.15       1.00
   5       5.00       3.57       2.27       1.67       1.19       1.00
  10      10.00       5.26       2.70       1.82       1.22       1.00
  20      20.00       6.90       2.99       1.90       1.23       1.00
  50      50.00       8.47       3.18       1.96       1.24       1.00
```

Read the `rho = 0.10` column downward. At a mere ten percent error correlation — which most forecasting groups would describe as excellent independence — fifty estimators are worth 8.47. At `rho = 0.30`, fifty are worth 3.18, and the column has nearly stopped responding to `N` at all.

```
CEILING OF THE VARIANCE INSTRUMENT

  As N -> infinity,  N_eff_var -> 1 / rho

  rho = 0.10  ->  ceiling 10 independent estimators, at any headcount
  rho = 0.30  ->  ceiling 3.33
  rho = 0.50  ->  ceiling 2
```

Error correlation imposes a hard ceiling on ensemble accuracy that no amount of hiring can raise. Fifty analysts at `rho = 0.3` and five thousand analysts at `rho = 0.3` are both worth about 3.3 analysts.

Two honesty notes. First, `N_eff_var` and the law's `N_eff = 1 + (N−1)u` are **different quantities from different derivations** — one from second moments of estimator error, one from the surplus axioms. This paper does not claim they are equal and does not convert between them. They are used as two independent readings of the same underlying condition, in the way an

engineer uses two different gauges on one vessel. Second, the equal-pairwise-correlation assumption is a simplification; real correlation structures are block-structured, and the right treatment there is the eigenvalue spectrum of the correlation matrix. The choice is stated in the report.

10.5.18 5.17 Maintaining u: driving lambda down

Because u decays, the design question is not how u is obtained but what lambda is and what it costs to lower it

The devices below are lambda-reduction devices. Each attacks a specific cascade mechanism. All are cheap; most are procedural rather than structural, which is why this is usually the least expensive intervention the framework offers.

DEVICE	CASCADE MECHANISM IT BLOCKS
Independent estimate before discussion; form closes before the meeting opens	Removes the public history agent 3 would condition on. Blocks the cascade at its formation.
Silent write-first, then round robin reading of what was written	Same, at meeting granularity. Every participant's signal enters the record.
Blind review; second reviewer cannot see the first's score until both are submitted	Blocks the sequential-observation channel in review. Cheapest fix in any peer-review system.
Last-speaks-first inversion for rank	Removes the weighted first signal that forms a cascade on turn one, not turn 3.
Rotating adversarial assignment, protected and voluntary	Keeps gamma nonzero without appointing a permanent opposer (FM-4F).
Rotation across teams and lineages; recruitment outside the dominant pipeline	Attacks Species 2 and 5 at the supply side: refreshes the input distribution, not just the roster.
Minority-report requirement	Forces eta > 0 on the divergence channel; makes dissent durable, not merely audible.
Decorrelated incentives	Removes the common-cause term producing correlated error among differently trained people.

The last row deserves emphasis because it is the one that survives all the others. **If every member of a group is paid on the same number, their errors will correlate no matter how differently they were trained, because the incentive is a common signal and a common signal is exactly what a cascade runs on.** Goodhart's law [Goodhart, 1975] and Campbell's law [Campbell, 1979] are usually read as warnings about metric gaming. Read here, they are also warnings about u: a single dominant metric is a u-destroyer, because it makes one dimension of the state space matter and flattens the rest.

The cost of maintenance. All of these raise C: a round, a coordination step, drafting time. That cost is real and must be checked against headroom before adoption — the headroom analysis's [Curzi, 2026b] rule stands, never request r_t > h_t. What Section 5.18 shows is that the cost is nearly always small relative to the surplus recovered, *provided the leak is fixed rather than patched once.*

10.5.19 5.18 Why a one-time fix does not work

This is Section 5's final structural result and it is the strongest argument for the maintenance framing.

Take the collapsed institute of Section 5.13 at `t = 12`, with `Q = 2.3015`, and apply a serious one-time intervention that restores `u` to 0.42. Compare two futures: one in which a small residual leak of `lambda = 0.03` per quarter remains, and one in which maintenance holds `u` flat.

```
N = 5, eta = 0.55, delta = 0.09, C = 0.15, restored u = 0.42 at t = 12

              lambda = 0.03                   lambda = 0.00 (maintained)
 t          u          Q                     u          Q
-------------------------------------------------------------------------
12       0.4200     2.3015                0.4200     2.3015
16       0.3722     2.8890                0.4200     2.9480
20       0.3298     3.1519                0.4200     3.3910
24       0.2922     3.1980  <- max          0.4200     3.6950
                     (3.2021 at t=23)
28       0.2589     3.1030                0.4200     3.9030
32       0.2294     2.9190                0.4200     4.0460
36       0.2033     2.6800                0.4200     4.1440
40       0.1801     2.4110                0.4200     4.2110
                                          Q_inf   =  4.3578
```

With a three-percent-per-quarter residual leak, `Q` peaks at 3.202 at `t = 23` and is back to 2.411 by `t = 40` — within 0.11 of where the intervention started, twenty-eight quarters later. **A leak of three percent per period, which is small enough that nobody would notice it in any single meeting, erases the entire benefit of a serious intervention inside seven years.**

With the leak closed, the same intervention converges to 4.358 — an 89 percent increase over the starting point, and it holds.

The two runs differ in *nothing* except whether the maintenance regime exists. Same restoration, same people, same cost structure, same headroom.

```
DESIGN CONCLUSION

  A one-time diversity intervention on a system with lambda > 0 is a step
  input to a leaky accumulator. It produces a transient and returns.

  The durable quantity is not u. It is 1/lambda -- the machinery that
  keeps u from leaving.
```

That is Section 5's thesis in one line, and it is why pseudo-diversity is correctly classed as a maintenance failure rather than a hiring failure.

10.5.20 5.19 Specification: measuring and maintaining non-parallelism

Requirements are numbered `SPEC 5.n`. Each is checkable.

Group A — measurement

SPEC 5.1 `u` shall be measured, not assumed. A system that reports a diversity posture without a `u` estimate is non-conforming.

SPEC 5.2 `u` shall be re-measured every period, on the same instrument, same item pool, and same pre-discussion ordering. A single measurement is a level; Section 5 requires a slope.

SPEC 5.3 `lambda` shall be estimated from two `u` measurements separated by `k` maintenance-free periods: `lambda_hat = 1 - (u_1/u_0)^(1/k)`. The maintenance-free condition shall be asserted in writing, with the dates.

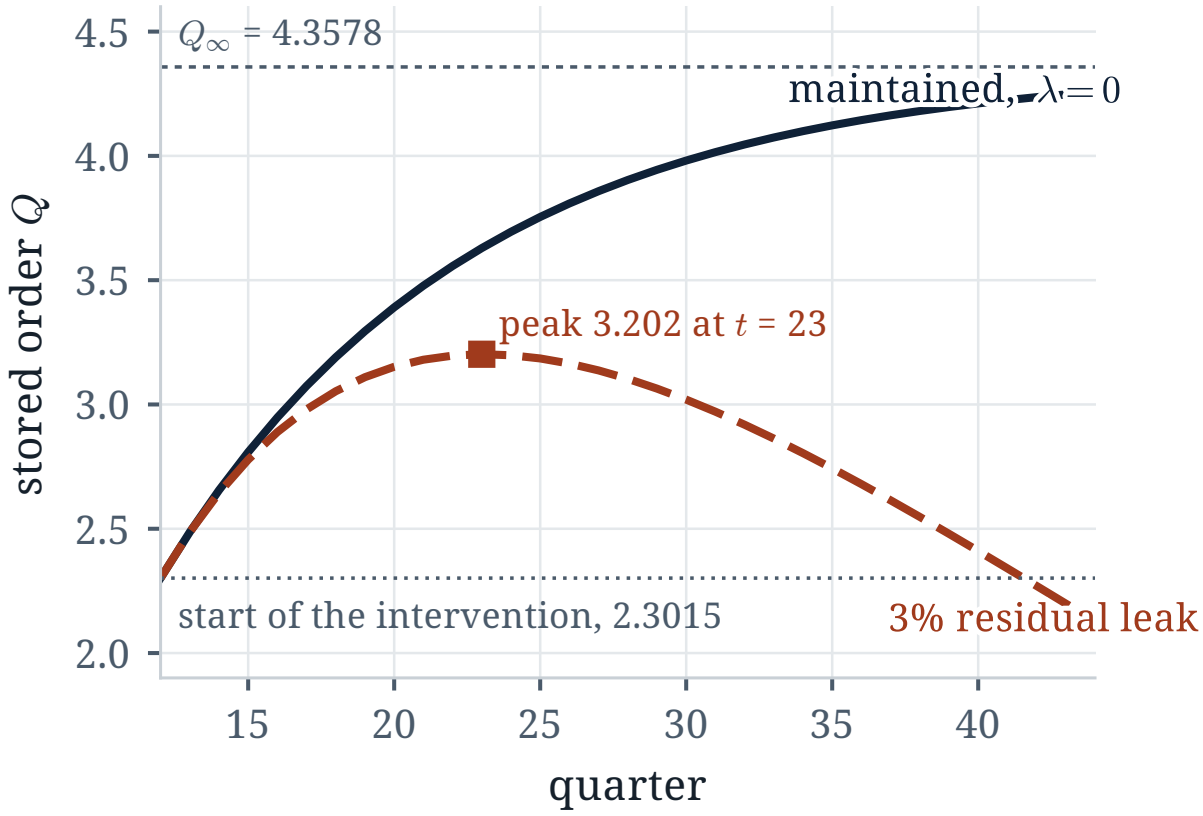


Figure 10.5: **A three-per-cent residual leak erases a serious intervention inside seven years; the same intervention with the leak closed holds.** Both trajectories begin from the identical restoration — u returned to 0.42 at quarter 12 with $Q = 2.3015$ — and differ in nothing except whether the maintenance regime is still being performed. The leaking run peaks at 3.2021 and is back to 2.4105 by quarter 40. **The durable quantity is $1/\lambda$, not u.** Modelled with the source's stipulated parameters; arithmetic exact. Generated by `figs/src/CCP-010_fig05.py`.

SPEC 5.4 The half-life `t_half = ln 2 / (−ln(1 − lambda))` shall be reported alongside `lambda`, in the same units as the reporting period.

SPEC 5.5 `rho_ceil = f(u) / ln N` shall be reported every period. `h_t` alone does not detect a parallel wall (Section 5.13) and shall not be relied on for that purpose.

SPEC 5.6 At least two independent detection instruments from Section 5.14 shall be run each period, and at least one of them shall be quantitative.

SPEC 5.7 Where a history of resolvable predictions exists, error correlation `rho` shall be measured and `N_eff_var = N/[1+(N−1)rho]` reported. Where the correlation structure is block-structured rather than exchangeable, this shall be stated and the eigenvalue treatment used instead.

SPEC 5.8 Category blocks shall be declared by evaluative standard (SPEC 4.1). Blocks declared on any axis other than the one on which judgment varies shall be marked as such, and `u` measured on both.

SPEC 5.9 Information supply shall be audited alongside the roster: each member's three most view-changing sources of the period, collected independently, with the overlap computed. This detects Species 2, which no roster audit can.

SPEC 5.10 Training lineage, not affiliation, shall be recorded for reviewer, candidate, and appointment pools. Two people at different institutions with one advisor are one lineage.

Group B — protocol

SPEC 5.11 Independent written estimates shall precede discussion on every decision where `u` is a design input, and the collection form shall close before the discussion opens. This is Section 5's highest-yield requirement and admits no exception on grounds of urgency.

SPEC 5.12 In any sequential review, a later reviewer shall not see an earlier reviewer's score or recommendation before submitting their own.

SPEC 5.13 Where rank is present, the highest-ranked participant shall speak last on substance.

SPEC 5.14 Every substantive decision artifact shall record the strongest case against the decision taken, whether or not anyone advanced it. This is an `eta` requirement, not a `u` requirement: it makes divergence durable.

SPEC 5.15 Pre-mortems shall be written independently before discussion, and their output read as a `u` measurement. Convergence of independently written failure stories shall be recorded as a low-`u` finding, not as confirmation.

SPEC 5.16 No participant shall be required to dissent, and a decline shall be recorded as a decline rather than as a failure (the framework's consent clause [Curzi, 2026]). Compelled divergence produces low-information objections that raise `C` without raising `u`.

SPEC 5.17 Where a single composite metric governs the incentives of all members, a `u` impact assessment shall be recorded. A single dominant metric is a common signal and therefore a cascade substrate.

SPEC 5.18 Adding stages, reviewers, or committee members shall require a prior statement of the expected `u` contribution of the addition. Additions with no expected `u` contribution add `C` only (Species 4) and shall be declined.

Group C — reporting and governance

SPEC 5.19 The period report shall contain:

```
PERIOD: ______   N (declared blocks): ____
u: ______   N_eff = 1+(N−1)u: ______   S = f(u): ______
ln N: ______   rho_ceil = S/ln N: ______
lambda: ______   t_half: ______ periods
```

```
rho (error correlation): ______   N_eff_var: ______
instruments run this period: ____________________
decisions with recorded dissent: ____/____
dissents that changed the artifact: ____/____
later-stage reversals of earlier stages: ____/____
source-overlap index: ______
```

SPEC 5.20 Postulate P-U shall be badged wherever `lambda` is used. It is a modeling assumption of this paper and not a theorem.

SPEC 5.21 Any use of `N_eff` shall badge axiom S3 in the same sentence.

SPEC 5.22 Preregistration (the preregistration protocol [Curzi, 2026i]) shall include the shuffled-input control. Its omission shall be treated as a material defect in the analysis.

SPEC 5.23 The report shall state, verbatim: *the surplus functional as deployed is a lower bound on the ensemble's information, not the information itself.*

10.5.21 5.20 Failure modes of pseudo-diversity and decay

FM-5A · THE PARALLEL WALL

Clinical. `u` at or near zero across the roster. `N_eff = 1`. Surplus `ln(1) = 0` regardless of headcount.

Damages. `S -> 0`, so the storage law reduces to `Q_{t+1} = (1-delta)Q_t - C_t`: pure decay with a cost draw. No headcount grows it.

Signature. Unanimity as the normal outcome. High confidence, flat information. Meetings whose conclusions are predictable from the agenda.

Measurement. Collision test returning near-zero collision rate; `rho_ceil` below 0.15; error correlation above 0.8.

Fix. Nothing on the `N` lever works (Section 5.15). Buy `u` directly, per the non-parallelism analysis [Curzi, 2026c] and Rule 4.4, and install the Section 5.17 `lambda` devices at the same time — otherwise what was bought is a transient (Section 5.18).

FM-5B · PSEUDO-DIVERSITY (declared variety, epistemic uniformity)

Clinical. Blocks declared along an axis on which judgment does not vary. The roster is heterogeneous on the declared axis and homogeneous on the deciding one.

Damages. `u` overestimated at declaration and near zero in operation. All downstream planning optimistic.

Signature. A diversity report and a decision record that disagree: many declared categories, no recorded disagreement.

Measurement. Re-declare by evaluative standard (SPEC 5.8) and re-measure. If cross-block `beta^2` does not exceed within-block `beta^2`, the declaration was not epistemic.

Fix. Re-declare the blocks, and treat the resulting number as the real one even when it is far below the previous report.

FM-5C · THE CASCADE COLLAPSE

Clinical. A group that measured well at formation and measures near zero later, with no personnel change. Individually rational copying [Bikhchandani et al., 1992] has removed private signals from the public record.

Damages. `u -> 0` over time; `lambda` large. Because `f` is steepest at zero (Section 5.4), the last increments lost are the most costly.

Signature. Declining variance in stated positions over successive periods. Newer members' positions converging to the group's within one or two cycles. Meetings whose first speaker predicts the outcome.

Measurement. Two `u` measurements separated by maintenance-free periods. `lambda_hat > 0.10` is a live collapse.

Fix. The Section 5.17 devices, especially independent-estimate-before-discussion and blind review. These block cascade formation rather than treating its result.

FM-5D · THE ONE-SHOT INTERVENTION

Clinical. A serious `u`-raising programme executed once, with no maintenance regime, on a system with `lambda > 0`.

Damages. A transient. `Q` rises, peaks, and returns toward its pre-intervention level (Section 5.18). The organization concludes the intervention did not work and declines to repeat it, which is the worse outcome.

Signature. A visible improvement 4 to 8 periods after a programme, followed by a slow, unattributed decline.

Measurement. Track `u` — not `Q` — after the intervention. `Q`'s lag (Section 5.13) makes it useless as an early indicator.

Fix. Pair every `u` intervention with a `lambda` intervention, and budget the maintenance before funding the programme. An unmaintained programme is worse than none: it consumes the capital for the next one.

FM-5E · THE ECHO SUPPLY CHAIN

Clinical. Heterogeneous members, homogeneous inputs. Difference in training does not survive identical evidence.

Damages. `u` near zero on any question the shared sources address, and normal-looking elsewhere — which makes it intermittent and easy to dismiss.

Signature. Members can articulate different frameworks and cite the same evidence. Positions update simultaneously across the group when one upstream source publishes.

Measurement. SPEC 5.9's source-overlap index.

Fix. Diversify inputs before diversifying the roster. Assign different members to different evidence bases deliberately; require the source list in the decision artifact.

FM-5F · CORRELATED FORECASTS PRESENTED AS AN ENSEMBLE

Clinical. `k` estimates with correlated errors, aggregated and reported with ensemble language and ensemble confidence intervals.

Damages. Confidence intervals too narrow by a factor of roughly `sqrt(N / N_eff_var)`. The number reported to decision-makers is wrong in the dangerous direction.

Signature. Ensemble spread much narrower than realized error, repeatedly. Forecasts that miss together.

Measurement. Section 5.16. Compute `rho` on historical residuals; report `N_eff_var` next to `N` every time.

Fix. Report `N_eff_var`, not `N`, and widen intervals accordingly. Decorrelate at source by varying data, method, and incentive across forecasters.

FM-5G · THE QUADRATIC COMMITTEE (committee theatre)

Clinical. Additional review stages that add no `u`. Each stage draws from the same pool and applies the same standard.

Damages. `C` grows with stages while `S` is flat; the communication-path count [Brooks, 1975] `n(n−1)/2` operates on the cost side only. the cost analysis's [Curzi, 2026b] crossover arrives early.

Signature. Reversal rate near zero across many decisions. Stages described as "rigorous" that have never changed an outcome.

Measurement. Reversal count per SPEC 5.19. Also: remove a stage in a pilot and measure whether decision quality changes.

Fix. SPEC 5.18. Delete stages that carry no `u`. Deleting a stage is a `C`-reduction that costs no surplus, which makes it one of the rare unambiguous wins in this paper.

FM-5H · THE MONOMETRIC INCENTIVE

Clinical. Every member's reward tied to one composite number.

Damages. Correlated error regardless of training diversity. The metric is a common signal, so u decays toward zero on every dimension the metric does not contain. Goodhart [1975] and Campbell [1979] operating as u-destroyers rather than as gaming risks.

Signature. People with genuinely different backgrounds making identically shaped trade-offs.

Measurement. Compare u measured on metric-relevant questions against u on metric-irrelevant ones. A large gap is the signature.

Fix. Decorrelate incentives across roles. Accept that this makes performance management harder; that difficulty is the price of the surplus.

FM-5I · THE MIRROR ROLE

Clinical. A permanent designated contrarian, credited as the group's divergence.

Damages. u credited where none exists (antiparallel is parallel, Section 5.2), and genuine dissent suppressed because the slot appears filled.

Signature. One member's position derivable from the proposal. Shared-ground correlation near -1.

Measurement. Predict the objector's position from the proposal. If reliable, they are a mirror.

Fix. Rotate the role; score by artifact change, not by objection count.

10.6 6 The room as a circuit element

The preceding analyses were about quantities reachable only through people: how fast stored order leaks, how much of an encounter survives it, how different the participants actually are. Every one of those quantities has a physical substrate, and the substrate is usually a room.

Section 6 is the first physical section. From here to the seating analysis [Curzi, 2026f] the framework works in metres, decibels, lux, and degrees Celsius, and it does so for one reason: **the coefficients of the storage law are not fixed properties of a group; they are partly properties of the place the group is standing in, and place is the cheapest thing in the system to change.**

A person cannot be legislated into being more different from their colleague. A wall can be moved. The wall is nine hundred kilograms of gypsum and it enforces its policy continuously, without a budget line, without a meeting, and without anyone's permission being renewed each quarter. That is the argument of Section 6 in one sentence.

10.6.1 6.0.1 What is proven

The storage law is a theorem about a recursion, and it is stated in the framework's founding statement [Curzi, 2026] exactly as the storage-law analysis [Curzi, 2026a] proves it:

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t        S_t = F(U_t) = ln(1 + (N−1) U_t)

0 <= S_t <= ln N                                 (\citet{ccp:CCP-A02}, Cor 2.2)
h_t = eta ln N − delta Q_t − C_t                  (the headroom gauge \citep{ccp:CCP-008} −− the headroo
```

Nothing in Section 6 alters that. Nothing in Section 6 is permitted to alter that. What Section 6 does is supply an *interpretation of the coefficients for a bounded physical volume containing people*, and interpretation is not proof.

10.6.2 6.0.2 What is postulated

POSTULATE P-R1 (badged as a postulate, not a theorem): a room is a circuit element. It has its own delta, eta, N, u and C, and those values are measurable by direct observation. This is

a modelling assumption in exactly the sense the framework's scale-invariance clause [Curzi, 2026] describes: the claim is about *model form*, not about physics. The storage law is scale-free in form; asserting that a 40-square-metre volume is one of the scales at which the form applies is a choice, and it must be earned by measurement in each case, not assumed.

POSTULATE P-R2 (badged): optional activity is the observable proxy for eta. Gehl's taxonomy divides what happens in a space into activities that would happen anyway and activities that happen only when conditions permit. The second category is behaviour that the space *caused*. Since eta is defined as the fraction of interaction that converts into durable structure, and since a space that causes voluntary lingering is a space where interaction has somewhere to go, the optional-activity rate is taken as eta's field instrument. **This is a proxy, and a proxy is a bet.** Section 6.3 states precisely what would falsify it.

POSTULATE P-R3 (badged): a room holds a small but non-zero Q of its own. Furniture arrangements that persist, worn paths, the shelf everyone knows to check, the corner that is understood to be for quiet work — these are stored order held in the physical medium rather than in anyone's head. They decay (delta) and they are replenished (eta S). Treating them as a Q term is a modelling choice. It is a productive one, because it explains why rooms that are "reset" nightly by cleaning staff never accumulate any convention.

10.6.3 6.0.3 The honesty clauses that bind the room reading

- **As deployed, the surplus functional is a LOWER BOUND** on the ensemble's information, not the information itself. When Section 6.10 computes a room's N_eff, it is computing a floor.
- **The founding surplus analysis [Curzi, 2026] is pure mathematics and makes no empirical or metaphysical claim.** A room does not "have surplus". A group of people interacting in a room may be *modelled* as having it.
- **S3, the affine effective count, is the design axiom** — it is chosen, not derived. Every N_eff figure in Section 6 inherits that choice.
- **The triangle inequality fails for geometric triples.** Do not chain room-to-room surplus estimates as if they added along a path.
- Numbers in Section 6 that come from the observational literature (Whyte, Gehl, Hall) are **borrowed and nominal**. They were measured in specific cultures, in specific decades, with specific populations. They are starting points for a local measurement, not constants.

The load-bearing statement of Section 6: a room does not create interaction, it sets the coefficients under which interaction either converts or dissipates.

10.6.4 6.1 The four terms, located in a physical volume

Take the storage law and ask, for a specific room, what each symbol points at.

Term	In a room	Typical field proxy	Units of the proxy
`Q_t`	Durable relation, standing agreement, learned convention, the arrangement that persists	Count of standing cross-group working agreements; count of persistent furniture states	dimensionless count
`S_t`	Interaction surplus generated by encounters in the volume this period	`ln(1 + (N−1)U_t)` computed from an encounter log	nats

Term	In a room	Typical field proxy	Units of the proxy
`eta`	Rate at which an encounter becomes durable relation or decision	attributable `dQ` per unit `S`	dimensionless
`delta`	Rate at which the room's stored order decays without maintenance	`1 − Q_1/Q_0` measured under zero input	per period
`C_t`	Friction of using the volume at all	booking overhead + travel + setup + acoustic fatigue + interruption	nats-equivalent, see Section 6.8
`N`	Genuinely distinct activity categories the room supports	category audit, Section 6.10	count
`u`	Non-parallelism of the people the room actually brings into contact	disagreement rate / error decorrelation across observed pairs	in `[0,1]`

Two of these deserve immediate comment.

`N` for a room is not the occupancy. It is the number of genuinely distinct *kinds* of use the volume supports. A room with twelve identical desks in twelve identical postures under one lighting condition offers one category, not twelve. the effective-count analysis's [Curzi, 2026c] arithmetic applies without modification: `N_eff = 1 + (N−1)u`, and at `u = 0`, `N_eff = 1`. **Ten parallel windows are one window; twelve parallel workstations are one workstation.**

`C` for a room is the term nobody prices. It includes the six minutes to find a free room, the four minutes of cable fumbling, the eleven-minute walk, the cognitive cost of raised vocal effort over a reverberant field for fifty minutes (the acoustic analysis [Curzi, 2026e] prices that one exactly), and the interruption cost paid by everyone the room's traffic passes. Section 6.8 gives a tally sheet.

10.6.5 6.2 Whyte's method: direct observation as instrumentation

Whyte [1980] is the methodological ancestor of Section 6, and its method matters more than any of its findings. Whyte and the Street Life Project did something that had not been done systematically: they pointed cameras at plazas and counted what people did.

They did not survey. They did not interview. They did not model. **They watched, they counted, and they timed, and the findings that came out contradicted a great deal of what designers believed they knew.**

The method decomposes into four instruments, and all four transfer directly to interior space:

Instrument 1 — Time-lapse. A fixed camera at a fixed interval (Whyte used film; a phone on a tripod at one frame per five or ten seconds is now trivially cheap). Time-lapse converts occupancy into a legible signal. Standing patterns invisible in real time — that nobody ever sits on the north bench, that the queue always kinks left, that the room empties at 15:40 not 16:00 — appear immediately.

Time-lapse is also the honest instrument, because it records what the observer did not think to look for. A count sheet records only the observer's hypotheses.

Instrument 2 — The count. At fixed times, on a fixed plan, mark every person's position and what they are doing. Whyte's core counts were: total people, people in groups (as opposed

to alone), and the sex ratio. This paper adds Gehl's activity category (Section 6.3).

Instrument 3 — The trace. Wear patterns, desire lines, displaced furniture, the chair that is always at 40 degrees to the table, the ledge with the polished patch. Traces are integrated occupancy — they are the room's own Q, written into its surface. the framework's gradient node [Curzi, 2026] taught the operator to read a desire line as a measured gradient. Here the engineer reads it as a time-integral of use that no observer had to be present for.

Instrument 4 — The intercept, used last and sparingly. Ask the occupants why. Whyte's own practice was to observe first and ask later, because asking first contaminates the observation and because people are poor reporters of their own spatial behaviour. Nobody reports sitting somewhere because it is seventeen inches high and the door is visible; they report that it seemed nice.

Whyte's findings that transfer to interiors

These are borrowed and nominal. They were measured on New York plazas in the 1970s.

- **"People sit where there are places to sit."** Whyte's own summary of his largest finding, and it is deliberately flat. The amount of *sittable space* predicted plaza use far better than any aesthetic or architectural variable he tested.
- **Movable seating outperforms fixed seating**, even when the mover barely moves it. The amenity is not the position; it is the *demonstrated ability to choose the position*. This is a direct antecedent of the environmental analysis's [Curzi, 2026f] rule that adjustability beats optimisation.
- **Triangulation**: some external stimulus — a performer, an unusual object, a view, a shared problem — that gives strangers a legitimate reason to speak to each other. Section 6.6.
- **Self-congestion**: people stop and talk *in the traffic*, at the mouth of the space, not in the calm interior. The most-used part of a plaza is the part nearest the street.
- **Sun, wind, trees, water, food.** Each independently associated with use.
- **The proportion of women as a quality index.** Whyte reported that where people had a genuine choice of spaces, women were more selective, so a high proportion of women in a space indicated a space that was working. This requires a hard flag: it is a 1970s observation about a particular city, and it is a correlate, not a mechanism. It earns a place here only because Woolley et al. [2010] independently found proportion of women among the correlates of the collective intelligence factor c (the collective-intelligence analysis [Curzi, 2026g]). **Two unrelated literatures pointing at the same index is worth a measurement; it is not worth a belief.**

The sittable-space ratio

The plaza standards adopted in New York after Whyte's work specified, in imperial units, roughly one linear foot of sitting space per thirty square feet of open space. Converted:

```
1 ft of seating per 30 ft^2
  = 0.3048 m per 2.7871 m^2
  = 0.1094 m of seating per m^2
  -> 1 metre of seating per 9.14 m^2 of open space
```

Open floor area	Minimum sittable edge (Whyte ratio)
40 m^2	4.37 m
100 m^2	10.94 m
200 m^2	21.87 m

Open floor area	Minimum sittable edge (Whyte ratio)
500 m^2	54.68 m
1000 m^2	109.36 m

This is a *floor*, not a target, and it is used knowing its provenance: it is a planning regulation derived from outdoor observation in one city. It is in this paper because it is one of the very few quantitative space standards that was derived from counting people rather than from committee. the passive-geometry analysis [Curzi, 2026e] will show a geometric trick for meeting the ratio without spending floor area, and the threshold analysis [Curzi, 2026f] will show why the seating wants to be on the perimeter.

Whyte's dimensional recommendations for the seating itself — nominal, verify locally:

Element	Whyte's nominal	Metric	Note
Seat height	~17 in	~430 mm	The usable band is wide; ~400-500 mm is sittable
Ledge depth, one-sided	>= 16 in	>= 400 mm	Below this, people perch, they do not sit
Ledge depth, two-sided	>= 30 in	>= 760 mm	Two backs to back without contact
Movable chairs	as many as possible	–	The single highest-yield item in his findings

10.6.6 6.3 Gehl's taxonomy and the eta proxy

Gehl [1971] divides outdoor activity into three classes. The division is the most useful single idea in Section 6.

Class	Definition	Dependence on environmental quality
Necessary	Activities people must do: go to work, wait for the bus, walk to the printer, attend the mandatory meeting	Occurs almost regardless of quality
Optional	Activities people do only if they want to and conditions invite it: sit, stand and look, linger, take the long way	Occurs **only** when conditions are favourable
Social	Activities that depend on the presence of others: conversation, greeting, children playing, simply watching	Largely *resultant* — they arise out of the other two

Gehl's central observation is that the third class is derived. Social life is not something that can be installed. It is what happens when necessary and optional activity coexist in the same volume for long enough. **"Something happens because something happens."**

For this paper the operational consequence is precise:

```
OAR := optional_count / (necessary_count + optional_count)

OAR is the field proxy for eta.
```

The logic: necessary activity is `S` that would have occurred anyway; the room gets no credit for it. Optional activity is behaviour the room *caused*. Social activity is optional and necessary activity that has begun to interact — which is precisely the raw material eta converts. A room where OAR is 0.10 is a corridor with chairs in it. A room where OAR is 0.35 is generating its own traffic.

What would falsify Postulate P-R2. If a room's OAR rises and its attributable `dQ` per unit `S` does not, the proxy is broken for that room. That is a real and common case: a beautiful lounge in which people linger pleasantly and nothing is ever decided, recorded, or carried out has high OAR and eta near zero. the conversion analysis [Curzi, 2026b] named this the Leaky Vessel. **OAR measures the opportunity for conversion, not the conversion.** Measure both, and if they diverge, believe the `dQ`.

Gehl's distance thresholds — borrowed, nominal, and worth measuring yourself

Gehl reports approximate distances at which the human sensory apparatus does different work. Treat as order-of-magnitude:

Distance	What is available
~100 m	A person is identifiable as a person; gait and body language readable — the "social field of vision"
~50-70 m	Sex, age, general activity readable
~25 m	Facial expression and emotional state readable
~7 m	Full sensory contact; detail, small gesture, ordinary speech comfortable
~1-3 m	Conversation at normal effort with full non-verbal channel

This ladder is the reason large open-plan floors do not produce more contact than small ones. Beyond roughly 25 m, other people are scenery. Beyond roughly 7 m, they are not conversational partners without a deliberate approach, and a deliberate approach has a social cost that most people will not pay for a marginal exchange. **A 900 m^2 open floor does not contain 900 m^2 of social opportunity; it contains the sum of its sub-7-metre neighbourhoods, and if it has no sub-structure it contains very few of them.** the passive-geometry analysis [Curzi, 2026e] is entirely about how to install that sub-structure without adding walls.

10.6.7 6.4 Hall's proxemics: the geometry of u at conversational range

Hall [1966] established that interpersonal distance is structured, not continuous, and that the structure carries meaning. His four zones, each with a close and a far phase, for North American middle-class adults:

Zone	Range	Character	What the channel carries
Intimate	0 - 0.45 m	Touch, heat, scent detectable	Involuntary signal; intrusion if unconsented
Personal	0.45 - 1.2 m	Arm's length; detail of expression	Confiding, close collaboration
Social	1.2 - 3.6 m	Business range; whole body visible	Working conversation, group of 3-8
Public	3.6 m +	Formal; voice must be raised	Address, performance, one-to-many

Hall himself insisted these are culturally specific. The zone boundaries differ substantially across cultures and the differences are large enough that mismatched norms read as aggression or coldness. The table is not a constant. It is a hypothesis to be checked against the people who will actually use the room.

Why proxemics is a u-instrument

the non-parallelism analysis [Curzi, 2026c] established that u is the parameter that moves fastest, and that systems destroy u without noticing. Proxemics is where architecture touches u directly, because **the distance structure of a room determines which pairs of people are ever in a channel that can carry disagreement.**

Consider a room's occupants as a set of pairs. For each pair, the room's geometry assigns a typical separation, and Hall's table assigns that separation a channel type:

- **Public-range pairs (> 3.6 m)** exchange announcements. Announcement is a one-way channel. It cannot carry the hesitation, the partial objection, or the "hmm" that is the raw form of most genuine difference. Pairs at public range contribute close to zero to measured u because their disagreements are never expressed and therefore never counted.
- **Social-range pairs (1.2 - 3.6 m)** are where working difference lives. This is the productive band and almost everything in the acoustic and seating analyses [Curzi, 2026e,f] is aimed at maximising the number of pairs in it.
- **Personal-range pairs (0.45 - 1.2 m)** carry the most bandwidth but cannot be sustained across a group; putting eight people at personal range produces a huddle that excludes anyone outside it, and huddles are u-destroyers because they select for the already-agreeing.
- **Intimate-range imposition (< 0.45 m without consent)** raises C sharply and immediately. A room that forces unconsented intimate proximity — a crowded lift lobby used as a meeting point — produces avoidance, and avoidance is a C term that never appears on any budget.

Design rule, stated as a rule and not as a metaphor: count the pairs the geometry places in the social band. That count, not the headcount, is the room's conversational capacity.

For a round table of diameter `D` seating `n`, the maximum separation is `D` (across the diameter) and the seat pitch is `pi*D/n`. Working the arithmetic at a 600 mm minimum pitch:

n	Minimum D for 600 mm pitch	Chosen D	Actual pitch	Max separation	All pairs in social band?
4	0.76 m	1.6 m	1.26 m	1.60 m	yes
6	1.15 m	1.6 m	0.84 m	1.60 m	yes
8	1.53 m	1.6 m	0.63 m	1.60 m	yes, but pitch is tight
10	1.91 m	2.0 m	0.63 m	2.00 m	yes
12	2.29 m	2.4 m	0.63 m	2.40 m	yes
13	2.48 m	2.8 m	0.68 m	2.80 m	yes
16	3.06 m	3.2 m	0.63 m	3.20 m	at the edge (3.2 of 3.6 m)

A round table stops being a single conversation at about sixteen people, and it stops for a geometric reason: the diameter crosses out of Hall's social band. the seating analysis [Curzi, 2026f] takes this further. the acoustic analysis [Curzi, 2026e] will show that the acoustic constraint bites before the proxemic one.

10.6.8 6.5 Sittable space, and why it is the highest-yield variable

Whyte's finding is worth restating in the language of the law. Seating does three things simultaneously:

1. **It converts throughput into dwell time.** A person walking through a room at 1.3 m/s is present for four seconds. A person seated is present for twenty minutes. S is generated per unit of co-presence, so seating multiplies S by two to three orders of magnitude per person admitted to the room.
2. **It raises u by admitting a wider population.** Standing-only spaces select for the physically comfortable standing, the unhurried, and the socially confident. Seating admits people who would otherwise not stay, and the people who would otherwise not stay are — by construction — a different population from those who would. **Accessibility is a u-instrument before it is anything else.**
3. **It lowers C.** Standing in a conversation one would prefer to leave costs more than sitting in one, because standing has no graceful exit protocol other than leaving. Seating with an easy exit (aisle access, no bag-shuffle, no one to climb over) reduces the cost of *entering* a conversation, because the cost of entering is dominated by the anticipated cost of leaving.

Point 3 is under-appreciated and worth a rule of its own: **people join conversations whose exit cost they can estimate.** An alcove with one entrance is a trap; an alcove with a clear line out is an invitation. the threshold analysis [Curzi, 2026f] develops this as the threshold principle.

Seating typology, ranked by observed yield in Whyte's work and by the reasoning above:

Type	Dwell	Choice	Exit cost	Notes
Movable chair	high	high	low	Highest yield. The move need not be large
Movable chair + movable small table	high	high	low	Enables the object-mediated conversation

Type	Dwell	Choice	Exit cost	Notes
Ledge / step, 400+ mm deep	medium	high	very low	Cheapest sittable metre; no maintenance
Bench, backless	medium	medium	low	Sittable from both sides at 760+ mm
Bench, with back	high	low	low	Comfortable, but fixes orientation and so fixes who talks to whom
Fixed banquette in an alcove	high	low	medium	Excellent for two to four; poor if the only option
Fixed rows facing one way	low (optional), high (necessary)	none	high	An auditorium. Optional activity near zero by design

The last row is the important one. **Fixed unidirectional seating sets the room's `u` toward zero by geometry: it establishes one front, and a room with one front has one speaker and n-1 audiences.** the seating analysis [Curzi, 2026f] prices this exactly.

10.6.9 6.6 Triangulation

Whyte's term for the mechanism by which an external stimulus links strangers. A street musician, a peculiar sculpture, a view of something happening, a stuck vending machine — anything that gives two people who have no relationship a legitimate, low-cost, deniable reason to speak.

Triangulation matters here because it is a **u-preserving contact mechanism.** Most engineered contact selects on similarity: the interest group, the affinity network, the project team. Contact selected on similarity buys contact at `u` near zero, which by the effective-count analysis's [Curzi, 2026c] arithmetic is worth `ln(1) = 0`. Triangulation selects on *co-presence at a stimulus*, which is orthogonal to similarity. It is one of the few cheap devices that produces contact between people who have nothing else in common.

Triangulating objects, ranked by robustness:

1. **A shared, visible, changing artefact.** A wall-mounted display of live system state, a physical kanban, a whiteboard nobody erases, a plot on the wall. Also a stigmergic medium in the sense of the order-without-centre analysis [Curzi, 2026h] — it does double duty.
2. **A shared consumable with a queue.** Coffee, water, food. The queue is not a defect; the queue is dwell time at conversational range with a legitimate reason to be silent or to speak. **A coffee machine with a thirty-second cycle is a worse social instrument than one with a ninety-second cycle**, and this is the only place in this paper where a slower machine is specified deliberately.
3. **A view.** the hospital-window study of 1984 is the strongest single piece of evidence in this literature that a window view affects a hard outcome — post-surgical recovery — and the environmental analysis [Curzi, 2026f] handles it. As a triangulator, a view of *activity* outperforms a view of scenery, because activity changes and change gives people something to remark on.
4. **A problem.** The most reliable triangulator ever built is a shared obstacle. Use with care: manufacturing obstacles to create bonding is a well-known way to raise C while pretending to raise eta.

10.6.10 6.7 Sun, food, water, trees

Whyte found each of these independently associated with use. The mechanisms are partly thermal and physiological (the environmental analysis [Curzi, 2026f]), partly attentional. Stated as room-scale requirements:

- **Sun / daylight.** People move to follow it. In an interior, daylight from more than one direction and at more than one intensity gives occupants a real choice, and choice is the amenity. A room with one uniform illuminance offers `N = 1` in the lighting category.
- **Food.** The single most reliable optional-activity generator available. It supplies a legitimate reason to be present, a reason to be still, a reason to be slow, and a shared object.
- **Water.** Whyte specifically noted that people touch water, sit near it, and that its sound masks conversation. The masking is a real acoustic function — see the acoustic analysis [Curzi, 2026e] on background level — and is the one case where deliberately raising background noise is correct: a low, broadband, steady source raises privacy without raising annoyance the way an HVAC tone does.
- **Trees / planting.** Beyond that finding, planting does three engineering jobs: it provides sub-division without enclosure (the passive-geometry analysis [Curzi, 2026e]), it absorbs and scatters high-frequency sound slightly (do not overstate this — the absorption coefficient of foliage is real but small), and it gives a visual focus that permits a person to be present without appearing to wait for someone.

10.6.11 6.8 The room's C term: a tally sheet

`C_t` is the term most often left blank, and it is the term that determines whether the room's headroom is positive. the cost analysis [Curzi, 2026b] established the general form; here is the room-specific tally.

```
C_room = C_access + C_setup + C_effort + C_interrupt + C_exclusion + C_maint
```

Component	What it is	How to measure it
`C_access`	Time to reach and to secure the room	Median door-to-door minutes x participants; booking-failure rate
`C_setup`	Time from arrival to first productive minute	Stopwatch, ten consecutive sessions
`C_effort`	Physiological cost of using the room: vocal effort, listening effort, thermal discomfort, glare	the vocal-effort analysis [Curzi, 2026e] gives the acoustic component in dB of required vocal effort; note reported hoarseness
`C_interrupt`	Cost imposed on people *outside* the room by its traffic and noise	Count of interruptions per hour at the nearest workstations
`C_exclusion`	Value of the contributions the room's geometry or acoustics silenced	The `ln N` loss computed in the acoustic exclusion arithmetic [Curzi, 2026e]
`C_maint`	Cleaning, resetting, repairing, replacing	Hours per week

The rule from the sustainability corollary [Curzi, 2026a] applies unchanged: if `C_t > eta ln N − delta Q_t`, no arrangement of the people in the room prevents order from decreasing.

A room can be so expensive to use that no seating plan, facilitation technique, or agenda saves it. Where such a room is found, the correct move is not to redesign the meeting but to cut `C` or to stop using the room.

Two `C` components are chronically invisible:

`C_interrupt` is paid by people who are not in the room and who therefore never appear in the room's evaluation. This is the classic externality structure and it produces the classic outcome: the room is over-used because its full cost is not borne by its users. the order-without-centre analysis's [Curzi, 2026h] the commons design principles [Ostrom, 1990] are the correct treatment — the boundary of the resource must include everyone who bears its cost.

`C_exclusion` is paid by the people whose contribution the room suppressed, and they cannot report it, because a contribution that was never made leaves no trace. This is the deepest measurement problem in Section 6. the acoustic analysis [Curzi, 2026e] solves it for the acoustic case by computing the loss from physics rather than from self-report. That is why the acoustic analysis [Curzi, 2026e] is the companion to this section.

10.6.12 6.9 The room's delta

`delta` for a room is the rate at which its stored order decays without maintenance. the retention analysis's [Curzi, 2026a] estimation method applies directly: measure `Q` at two times under zero input:

```
delta = 1 − (Q_1 / Q_0)          t_half = ln 2 / (−ln(1 − delta))
```

Room-scale `Q` and its decay:

Stored order	Decays because	Nominal delta driver
Furniture arrangement that works	Cleaning reset, the next booking, a single inconsiderate user	Very high without a rule; near zero with a fixed floor marking
Whiteboard content	Erasure	Total on erasure; delta is effectively a step function
Convention ("the corner is for quiet")	Occupant turnover, absence of a visible marker	Tracks personnel turnover
Wear paths, patina	Resurfacing, refurbishment	Very low; this is the most durable Q a room holds
Artefacts left in place (the shared book, the tool board)	Removal, tidying, theft	Depends entirely on whether removal is anyone's job

The single highest-leverage delta intervention in a room is a floor marking. A taped or inlaid outline showing where the table and chairs belong converts a convention held in fragile human memory (high delta, resets on turnover) into a convention held in the physical medium (very low delta, survives everyone). It costs almost nothing. It is the order-without-centre analysis's [Curzi, 2026h] rule — put the state in the medium, not in a coordinator — applied at the smallest possible scale.

The counter-case from the retention analysis [Curzi, 2026a] applies here too: a room with no forgetting cannot learn. A room whose arrangement is welded in place holds its Q perfectly and can never be improved by its occupants. The correct target is not delta = 0; it is delta low enough to survive a week of disuse and high enough that a better arrangement can displace the current one within a month. In practice: **the furniture should be movable, and its default position should be marked.** That is low delta on the default and non-zero delta on the deviation, which is exactly the required asymmetry.

10.6.13 6.10 The room's N and u

Computing N for a room requires an audit of categories of use, not a count of anything.

Procedure — the category audit.

1. List every distinct *posture-plus-purpose-plus-condition* the room currently supports. A posture is standing, perching, sitting upright, sitting reclined, lying. A purpose is solo focus, pair work, small-group deliberation, presentation, waiting, eating, transit. A condition is bright/dim, quiet/lively, enclosed/open, warm/cool.
2. Strike any entry that is not actually available — a category that requires moving 90 kg of furniture is not on the menu.
3. Strike any entry that duplicates another in every respect that matters to a user. **Two identical four-person tables are one category.**
4. The surviving count is N.

Most rooms audit at `N = 1` or `N = 2` and their owners are surprised. A standard meeting room offers: sitting upright, at a rectangular table, under uniform light, for group talk. That is one category. Its `N_eff` is therefore at most 1 regardless of who is in it, and by the range corollary [Curzi, 2026] its per-interaction surplus ceiling is `ln 1 = 0` in the use-category dimension. The room contributes nothing to S on its own account; all surplus must come from the people, and the room's job reduces to not destroying it.

Estimating u for a room means estimating the non-parallelism of the people the room actually brings into contact — not of the organisation, but of the pairs the geometry produces.

Procedure — the contact-pair audit.

1. From the time-lapse or count sheets, list every observed conversational pair over the survey period, with duration.
2. For each pair, score whether the two are in the same team, the same discipline, the same training lineage, the same reporting line. the non-parallelism analysis's [Curzi, 2026c] criteria apply.
3. Weight by duration. Compute the fraction of conversational minutes spent between pairs that differ on at least two of the four axes. Call it `p_cross`.
4. Use `p_cross` as the room's first-order u proxy, and state that this has been done.

This is a crude proxy and must be badged as one. It substitutes categorical difference for geometric non-parallelism, which is exactly the substitution Section 5 warns against — demographic or organisational variety is not epistemic variety. It is a screening instrument only, and is followed by the disagreement audit from Section 5.10 before believing any number it produces.

Worked instance of the arithmetic, using the effective-count analysis's [Curzi, 2026c] formula and the ceiling from the surplus range analysis [Curzi, 2026]:

u	`N_eff = 1 + (N−1)u` for N=8	`F = ln N_eff`	Fraction of ceiling `ln 8 = 2.0794`
0.00	1.00	0.0000	0.0%
0.10	1.70	0.5306	25.5%
0.25	2.75	1.0116	48.6%
0.50	4.50	1.5041	72.3%
0.75	6.25	1.8326	88.1%
1.00	8.00	2.0794	100.0%

Read the concavity: **moving u from 0 to 0.25 buys 48.6% of the entire available ceiling; moving it from 0.75 to 1.00 buys 11.9%.** The first genuinely different contact the room produces is worth more than the last four. A room that produces *any* cross-lineage contact is doing most of the work a room can do.

10.6.14 6.11 Occupancy, Brooks, and the sizing of a room

the cost analysis [Curzi, 2026b] established that coordination cost grows with the number of channels, `n(n−1)/2`, while per-interaction surplus is capped at `ln n`. For a room, this is a sizing law.

n	channels `n(n−1)/2`	`ln n`	channels per nat
2	1	0.693	1.44
3	3	1.099	2.73
4	6	1.386	4.33
5	10	1.609	6.21
6	15	1.792	8.37
8	28	2.079	13.47
10	45	2.303	19.54
12	66	2.485	26.56
13	78	2.565	30.41
16	120	2.773	43.28
20	190	2.996	63.42
25	300	3.219	93.20

Cost grows quadratically while surplus grows logarithmically, so every room has a size beyond which adding a person subtracts. The table does not locate that point — it depends on `eta` and on the per-channel cost in the setting — but it gives the shape, and the shape is brutal. Going from 8 to 16 occupants buys 33% more surplus ceiling and costs 4.3 times more channels.

The room's job is to *decompose* `n` rather than to hold it. A twenty-person room that is acoustically and geometrically one space carries 190 channels. The same twenty people in a room sub-divided into four alcoves of five carry `4 x 10 = 40` internal channels plus whatever crosses between alcoves. **Sub-division is not a compromise on connectivity; it is the only way a large room stays affordable.** the passive-geometry analysis [Curzi, 2026e] is the geometry of doing that without walls.

10.6.15 6.12 What a room cannot do

Three limits, stated plainly, because Section 6 is where practitioners over-claim.

A room cannot manufacture `u`. If everyone admitted holds the same training, the same incentives, and the same information, the geometry is irrelevant: `u = 0`, `N_eff = 1`, `F = 0`. Sections 5 and 6 of the non-parallelism analysis own that problem and no amount of furniture solves it.

A room cannot raise `eta` above what the surrounding practice permits. A perfect room feeding an organisation with no decision log, no artefact, and no memory produces beautiful conversations that vanish. The room raises `eta` from what it would have been; it does not set the ceiling.

A room cannot make a negative-headroom system positive on its own. The Worked Example at the end of the acoustic analysis [Curzi, 2026e] demonstrates exactly this: full acoustic treatment closes 75% of the gap to `h = 0` and the last 25% has to come from shedding `Q`. **Physical intervention is necessary and frequently insufficient, and an account that pretended otherwise would be selling furniture.**

Numbered, checkable, buildable. `SPEC 6.n` items are requirements; `MEAS 6.n` items are measurements. Every requirement names the term of the law it acts on.

10.6.16 6.13 Specification: requirements on a room

SPEC 6.1 (acts on `S`, `u`). Sittable space shall be provided at not less than 0.109 linear metres per square metre of open floor area (the Whyte ratio). For a 120 m^2 room this is 13.1 m.

SPEC 6.2 (acts on `u`, `C`). Not less than 40% of seating shall be movable by one person without assistance, defined as under 8 kg and free of cabling.

SPEC 6.3 (acts on `S`). Seat height shall lie in 400-500 mm; sittable ledge depth shall be not less than 400 mm one-sided and not less than 760 mm where sitting from both sides is intended.

SPEC 6.4 (acts on `N`). The room shall support not fewer than three distinct posture-plus-purpose-plus-condition categories, verified by the category audit of Section 6.10 and recorded on the worksheet.

SPEC 6.5 (acts on `u`, `C`). For every seat, the maximum separation to any other seat in the same intended conversational group shall not exceed 3.6 m (Hall's social band). Groups requiring greater separation shall be treated as separate groups and given separate sub-volumes.

SPEC 6.6 (acts on `C`). Every seat shall have an exit path not requiring another occupant to move.

SPEC 6.7 (acts on `eta`). The room shall contain at least one triangulating element per 40 m^2: a shared changing artefact, a consumable with a queue, a view of activity, or a shared physical tool. Static decoration does not count.

SPEC 6.8 (acts on `eta`). The room shall contain at least one writable, non-erased surface of not less than 2 m^2, and a defined owner responsible for photographing and archiving it before any erasure. **A whiteboard with no archive protocol has eta = 0 by construction.**

SPEC 6.9 (acts on `delta`). The default position of every movable element shall be marked on the floor or the surface it stands on. Markings shall be inlaid or taped, not printed on paper.

SPEC 6.10 (acts on `delta`). Room reset shall be specified as a written, posted, 60-second procedure, and it shall be the responsibility of the last occupant, not of cleaning staff.

SPEC 6.11 (acts on `C`). Median door-to-door access time for the intended user population shall not exceed 3 minutes. Where it does, the room's `C_access` shall be tallied explicitly and subtracted in the headroom calculation before the room is approved.

SPEC 6.12 (acts on `C`). Setup time from arrival to first productive minute shall not exceed 90 seconds, measured over ten consecutive real sessions, not in commissioning.

SPEC 6.13 (acts on `C_interrupt`). The count of interruptions per hour imposed on the three nearest workstations by the room's traffic and noise shall be measured before and after occupation, and the delta shall be recorded as a cost of the room.

SPEC 6.14 (acts on `N`, `u`). Illuminance shall vary across the room by at least a factor of three between its brightest and dimmest usable working positions, so that lighting constitutes a real choice. the environmental analysis [Curzi, 2026f] gives absolute ranges.

SPEC 6.15 (acts on `S`). Where a room exceeds 12 intended occupants, it shall be sub-divided into sub-volumes of not more than 6, by any means that establishes a distinct acoustic and visual sub-field. the passive-geometry analysis [Curzi, 2026e] gives non-wall methods; the acoustic analysis [Curzi, 2026e] gives the acoustic criterion.

SPEC 6.16 (acts on `eta`). A consumable point (water at minimum, hot drink preferred) shall be within 15 m and shall have a dwell of not less than 30 seconds per use.

SPEC 6.17 (acts on `C_exclusion`). The acoustic requirements of the acoustic analysis [Curzi, 2026e] are a precondition of Section 6's requirements, not an addition to them. A room failing the acoustic criterion [Curzi, 2026e] shall not be certified against Section 6 regardless of its other properties.

SPEC 6.18 (acts on `Q`). The room shall carry a visible, physical log of what it has produced — decisions taken, agreements standing. This is the room's `Q` made legible and it is the only defence against the Leaky Vessel failure mode.

SPEC 6.19 (ethics, the framework's consent clause [Curzi, 2026]). No requirement in Section 6 shall be implemented in a way that removes an occupant's ability to decline participa-

tion, to leave, or to be unobserved. Seating that cannot be exited, geometry that cannot be avoided, and observation that cannot be opted out of are out of scope for this paper.

SPEC 6.20 (measurement integrity). Any published figure for this room's `eta`, `delta`, `u` or `N` shall carry the date, the method, and the observer. the survey protocol's [Curzi, 2026i] standard applies: **justify that the order proxy is not merely a repackaging of the input.**

10.6.17 6.14 The measurement worksheet

The worksheet occupies one side of A4 per room per survey. Everything on it is measurable in one week by one person with a tape measure, a camera, and a clipboard.

```
======================================================================
 ROOM SURVEY WORKSHEET                          Standard Manual, Ch. 12
======================================================================
 Room ID ................  Building ..........  Floor ................
 Surveyor ...............  Survey week from ......... to ...........
 Intended function ..................................................

----------------------------------------------------------------------
 A. PHYSICAL INVENTORY                                      (one pass)
----------------------------------------------------------------------
 A1  Floor area, usable                          ............ m^2
 A2  Volume (area x mean clear height)           ............ m^3
 A3  Clear height                                ............ m
 A4  Total surface area (floor+ceiling+walls)    ............ m^2
 A5  Glazed area                                 ............ m^2
 A6  Number of entrances                         ............
 A7  Longest internal sightline                  ............ m
 A8  Linear metres of sittable surface           ............ m
 A9  A8 / A1  (Whyte ratio; require >= 0.109)    ............ m/m^2
 A10 Movable seats / total seats                 ........./........
 A11 Seat height range measured                  ...... to ...... mm
 A12 Reverberation time RT60 at 500 Hz, occupied ............ s   [acou]
 A13 Background noise level, unoccupied          ............ dBA [acou]
 A14 Illuminance max / min at working positions  ....../...... lux
 A15 Median door-to-door access time             ............ min

----------------------------------------------------------------------
 B. CATEGORY AUDIT  -> N                       (Section 6.10, strike rule)
----------------------------------------------------------------------
  #  posture         purpose           condition        available?
  1  ..............  ..............    ..............   Y / N
  2  ..............  ..............    ..............   Y / N
  3  ..............  ..............    ..............   Y / N
  4  ..............  ..............    ..............   Y / N
  5  ..............  ..............    ..............   Y / N
  6  ..............  ..............    ..............   Y / N
     (strike duplicates and unavailable rows)
 B1  N (surviving distinct categories)           ............
 B2  ln N  (per-interaction ceiling)             ............ nats

----------------------------------------------------------------------
 C. ACTIVITY COUNT   (Gehl)        5 fixed times/day x 5 days = 25 obs
----------------------------------------------------------------------
        NEC = necessary   OPT = optional   SOC = social
```

```
  day    NEC    OPT    SOC     OAR = OPT/(NEC+OPT)
  Mon   .....  .....  .....    ...........
  Tue   .....  .....  .....    ...........
  Wed   .....  .....  .....    ...........
  Thu   .....  .....  .....    ...........
  Fri   .....  .....  .....    ...........
  TOT   .....  .....  .....    ...........
C1  OAR (week)                                ...........
C2  SOC / OPT                                 ...........
C3  Peak simultaneous occupancy               ...........
C4  Median dwell time, seated                 ........... min

-------------------------------------------------------------------------
D. CONTACT-PAIR AUDIT  -> u proxy                       (Section 6.10)
-------------------------------------------------------------------------
  pair  minutes  same team?  same discipline?  same lineage?  same line?
  ....  .......  Y/N         Y/N               Y/N            Y/N
  ....  .......  Y/N         Y/N               Y/N            Y/N
  ....  .......  Y/N         Y/N               Y/N            Y/N
  ....  .......  Y/N         Y/N               Y/N            Y/N
  (continue on reverse)
D1  Total conversational minutes observed      ........... min
D2  Minutes between pairs differing on >= 2 axes ........... min
D3  p_cross = D2 / D1   (badged u proxy)       ...........
D4  N_eff = 1 + (B1 - 1) x D3                  ...........
D5  F = ln(D4)                                 ........... nats
D6  F / B2  (fraction of ceiling reached)      ...........

-------------------------------------------------------------------------
E. STORAGE TERMS
-------------------------------------------------------------------------
E1  Q proxy chosen (name it) ........................................
E2  Q_0 at survey start                         ...........
E3  Q_1 after one zero-input period             ...........
E4  delta = 1 - E3/E2                           ........... /period
E5  half-life = ln2 / (-ln(1-E4))               ........... periods
E6  S proxy chosen (name it) ........................................
E7  S per period                                ...........
E8  attributable dQ per period                  ...........
E9  eta_hat = E8 / E7                           ...........

-------------------------------------------------------------------------
F. COST TALLY  -> C                                      (Section 6.8)
-------------------------------------------------------------------------
F1  C_access     (min/session x sessions/period) ...........
F2  C_setup      (min/session x sessions/period) ...........
F3  C_effort     (vocal effort dB; hoarseness reports)   ..........
F4  C_interrupt  (interruptions/hr at 3 nearest workstations) ........
F5  C_exclusion  (acoustic exclusion, ln N loss)           ........... nats
F6  C_maint      (hours/period)                 ...........
F7  C_total, in the units of E7/E9 (state conversion) ................

-------------------------------------------------------------------------
G. HEADROOM
-------------------------------------------------------------------------
```

```
 G1  eta ln N   = E9 x B2                          ............
 G2  delta Q    = E4 x E2                          ............
 G3  C          = F7                               ............
 G4  h = G1 - G2 - G3                              ............
      h < 0  UNSUSTAINABLE -> cut C, raise eta, or shed Q. Do not build.
      h = 0  MAINTENANCE ONLY -> hold.
      h > 0  GROWTH AVAILABLE -> r_max = h.

-----------------------------------------------------------------------
 H. COMPLIANCE
-----------------------------------------------------------------------
   SPEC 6.1 ... [ ]   SPEC 6.6 ... [ ]   SPEC 6.11 .. [ ]   SPEC 6.16 .. [ ]
   SPEC 6.2 ... [ ]   SPEC 6.7 ... [ ]   SPEC 6.12 .. [ ]   SPEC 6.17 .. [ ]
   SPEC 6.3 ... [ ]   SPEC 6.8 ... [ ]   SPEC 6.13 .. [ ]   SPEC 6.18 .. [ ]
   SPEC 6.4 ... [ ]   SPEC 6.9 ... [ ]   SPEC 6.14 .. [ ]   SPEC 6.19 .. [ ]
   SPEC 6.5 ... [ ]   SPEC 6.10 .. [ ]   SPEC 6.15 .. [ ]   SPEC 6.20 .. [ ]

 Surveyor signature .............  Date ......................
 Proxy justification (Y is not a repackaging of U):
 ....................................................................
 ....................................................................
=======================================================================
```

10.6.18 6.15 Failure modes of the room

Each entry gives the name, the clinical description, the damaged parameter, the diagnostic signature, the detecting measurement, and the fix. These entries feed the consolidated failure catalogue's [Curzi, 2026j] catalogue.

FM-6.1 — The Corridor With Chairs

Damages: `eta`, `S`. **Description.** A room whose occupancy is entirely necessary activity. People pass through, or attend because they were required to, and nobody is ever there by choice. **Signature.** OAR below 0.12. Median seated dwell under 4 minutes. Peak occupancy coincides exactly with scheduled events and drops to zero between them. **Detects with.** Worksheet section C. **Fix.** Triangulate (SPEC 6.7) and seat (SPEC 6.1, SPEC 6.2). A consumable point with real dwell is the single most reliable intervention. Expect OAR to roughly double within four weeks or the diagnosis was wrong.

FM-6.2 — The Beautiful Vessel

Damages: `eta`. **Description.** High optional activity, near-zero conversion. People linger happily; nothing becomes durable. The Leaky Vessel of the conversion analysis [Curzi, 2026b], in its architectural form. This is the failure that falsifies Postulate P-R2 locally, and it is common in exactly the spaces that win design awards. **Signature.** OAR above 0.30 with attributable `dQ` per period at or near zero. Occupants describe the room warmly and cannot name anything that came out of it. **Detects with.** Worksheet E8/E9 measured *independently of* C1. Never infer `eta` from OAR alone. **Fix.** SPEC 6.8 (archived writable surface) and SPEC 6.18 (visible production log). The room needs an outlet, not more amenity.

FM-6.3 — The Single-Category Room

Damages: `N`, and therefore the ceiling `ln N`. **Description.** One posture, one purpose, one condition. The standard meeting room. `N = 1`, `ln N = 0` in the use-category dimension. **Signature.**

Category audit strikes every row but one. Occupants who need a different mode leave the building rather than move within it. **Detects with.** Worksheet section B. **Fix.** Add categories, not area. A window ledge at 430 mm, a standing rail, a dimmable corner and one movable low table can take a room from `N = 1` to `N = 4` for the price of the ledge. Note the concavity: **going from one category to two is the largest single gain the room will ever record.**

FM-6.4 — The Public-Range Meeting

Damages: `u`, `C`. **Description.** Geometry places most pairs beyond 3.6 m. The channel degrades from conversation to announcement, and announcement carries no disagreement. **Signature.** Longest internal sightline in use exceeds 4 m for a group intended as one conversation. Speaking-turn distribution collapses onto two or three people (the acoustic exclusion arithmetic [Curzi, 2026e] quantifies this). Participants at the far end report "I couldn't get in". **Detects with.** Worksheet A7 and D; the turn-share measurement of the acoustic analysis [Curzi, 2026e]. **Fix.** Reduce the table, not the group; or split the group. SPEC 6.5 and SPEC 6.15.

FM-6.5 — The Nightly Reset

Damages: `delta` (drives it to 1 for the arrangement class of `Q`). **Description.** Every persistent physical state the occupants establish is erased between sessions by cleaning, by the next booking, or by a tidiness policy. The room can never accumulate convention. **Signature.** Occupants spend the first minutes of every session re-establishing the same arrangement. Whiteboards are found blank. Nothing is ever where it was left. **Detects with.** Worksheet E2/E3/E4 with the arrangement as the `Q` proxy. **Fix.** SPEC 6.9 floor markings and SPEC 6.10 a written reset procedure owned by the last occupant, not by cleaning. Convert the convention from a memory into a marking.

FM-6.6 — The Welded Room

Damages: learning; the mirror image of FM-6.5. **Description.** `delta` near zero on everything. Nothing can be moved, so nothing can be improved. The room's arrangement is optimal for a use that ended in 2014. **Signature.** Movable seat fraction near zero (worksheet A10). Occupants have visible workarounds: sitting sideways, standing to see, using the floor. **Detects with.** A10, plus the trace instrument — the workarounds are the room's own defect report. **Fix.** Unweld the arrangement. SPEC 6.2. Accept some `delta` in exchange for the ability to search.

FM-6.7 — The Externalised Room

Damages: `C_interrupt`, borne by non-users. **Description.** The room's costs fall on people outside it, who have no vote on its use. Classic commons failure in the sense of Ostrom [1990] (the order-without-centre analysis [Curzi, 2026h]): the boundary of the resource excludes some of those who bear its cost. **Signature.** Interruption rate at adjacent workstations rises measurably with room bookings. Adjacent occupants have adopted headphones, moved, or stopped using their desks. **Detects with.** Worksheet F4, measured before and after. **Fix.** Acoustic separation (the acoustic analysis [Curzi, 2026e]), traffic re-routing (the threshold analysis [Curzi, 2026f]), and governance: extend the decision boundary to include the affected. **The engineering fix and the governance fix are both required; either alone regresses.**

FM-6.8 — The Uncounted Exclusion

Damages: `N`, `u`, and the credibility of every measurement taken. **Description.** The room suppresses certain contributions — acoustically, proxemically, or by sightline — and because a suppressed contribution leaves no trace, the room's own metrics look fine. The survey reports

a healthy OAR from the people the room did not exclude. **Signature.** Turn-share concentration. Systematic absence of a subgroup. Consistent late-cycle objections from people who "should have raised it earlier" and could not. **Detects with.** the acoustic exclusion arithmetic [Curzi, 2026e] computes the exclusion from physics rather than from report. This is the only reliable route, because **self-report cannot detect a contribution that was never attempted. Fix.** the acoustic analysis's [Curzi, 2026e] specification in full, plus the seating analysis's [Curzi, 2026f] protocol layer (silent-write-first, round-robin) as a backstop. SPEC 6.17 makes the acoustic requirement a precondition.

FM-6.9 — The Amenity Substitution

Damages: honesty of the headroom calculation. **Description.** Visible amenity is installed and reported as if it were measured improvement. Plants, art, a coffee machine, and a rebrand, with no before/after measurement and no change to `eta`, `delta`, or `C`. **Signature.** The intervention has a photograph and no numbers. The evaluation is a satisfaction survey administered once, after. **Detects with.** the preregistration protocol's [Curzi, 2026i] standard: preregistered before-measurement, held-out after- measurement, and the shuffled-input control. **Fix.** Measure before purchasing. If the worksheet was not completed before the intervention, the intervention has not been evaluated, whatever the survey reports.

FM-6.10 — The Observed Room

Damages: trust; violates the framework's consent clause [Curzi, 2026] and SPEC 6.19. **Description.** Measurement conducted in a way that removes occupants' ability to be unobserved. Continuous individual-level tracking, badge telemetry tied to identity, cameras without notice. **Signature.** Occupants change behaviour on learning of the instrumentation, which also destroys the measurement. **Detects with.** Ask. This is the one failure mode for which the intercept instrument is the right one. **Fix.** Aggregate counts, posted notice, defined retention, opt-out that costs nothing. **A measurement that cannot be declined is not permitted by this paper, and it is also bad instrumentation, because it measures the observation and not the room.**

10.7 7 Results: three systems worked end to end

Three systems are worked completely, from survey through diagnosis to a costed intervention and a frozen falsification rule. **All three are constructed.** No organisation was surveyed, no room was measured, and every parameter below is stipulated so that the arithmetic can be redone by the reader from the numbers printed here. They are included because the paper's claims are claims about what happens when the procedures are actually carried out, and the failure modes of Sections 4.15, 5.20 and 6.15 are visible only in a complete case.

The three exercise different parts of the framework. The first has real crossing, near-zero divergence, and a standing proposal that the arithmetic rejects. The second was genuinely diverse at constitution and became arithmetically indistinguishable from a single reviewer in five years with no personnel change — the decay result with a trajectory attached. The third is a physical volume, and it is the case in which **the honest answer is partly negative and is reported as such.**

10.7.1 7.1 A cross-disciplinary research institute

A 34-person applied research institute working on flood risk, founded eleven years ago on an explicit cross-disciplinary premise and seriously executed.

Survey (the survey protocol [Curzi, 2026i] procedure). Declared blocks, by evaluative standard rather than by title:

```
N = 6                                    what counts as a good answer
  B1  hydrology / hydraulic modelling     model skill scores
  B2  civil & structural engineering      code compliance, FEA
  B3  economics & cost-benefit            discounted NPV, elasticities
  B4  public health & emergency response  outcome rates, response times
  B5  law & regulatory                    statutory defensibility
  B6  community engagement                participation, equity metrics

ln N = ln 6 = 1.791759

Storage-law parameters (the storage-law estimation procedures \citep{ccp:CCP-007,ccp:CCP-008} procedures, p
  delta = 0.08   quarterly loss of stored order without maintenance
  eta   = 0.45   fraction of surplus converted to durable structure
  C_t   = 0.20   quarterly dissipation: coordination, translation, compliance
  Q_0   = 6.00   current stored order, in the survey's Y-proxy units
```

Instrument A run. 1,122 ordered pairs, 5 anonymized artifacts per target, E/N classification. `u_cross_hat = mean beta^2 = 0.72` — real, and the institute earned it. Cross-block pairs show `beta^2` from 0.81 to 0.94 and within-block pairs 0.04 to 0.11, so the declared blocks survive the Failure Mode FM-4E check.

Instrument B run. Mean `n_E = 19` shared-ground items per pair, above the provisional threshold; independent scores before discussion, 9-point scale, on 19 recent risk determinations.

```
within-block pairs:  mean r = 0.985  ->  1 - r^2 = 0.0298
cross-block pairs:   mean r = 0.94   ->  1 - r^2 = 0.1164 (but alpha^2 is small)

ensemble roll-up:  u_div_hat = mean over pairs of alpha^2 (1 - r^2) = 0.03
```

The as-found decomposition.

```
u_cross = 0.72
u_div   = 0.03
u       = 0.75

f(u_cross) = ln(1 + 5*0.72) = ln(4.60)  = 1.526056
f(u_div)   = ln(1 + 5*0.03) = ln(1.15)  = 0.139762
f(u)       = ln(1 + 5*0.75) = ln(4.75)  = 1.558145

floor   = max(1.526056, 0.139762) = 1.526056
ceiling = 1.526056 + 0.139762     = 1.665818   [impossibility limit only]

N_eff   = 1 + 5*0.75 = 4.750       (S3 is a design axiom, badged)
rho_ceil = 1.558145 / 1.791759 = 0.8696
B       = (1.558145 - 1.526056) / 1.558145 = 0.0206
```

Read the diagnosis. `B = 0.0206`, below the SPEC 4.13 threshold: the institute is single-sourced, `u_div` contributing 2.06 percent of its surplus. **Thirty-four people, six genuine disciplines, and a surplus that would be almost unchanged if every within-block disagreement in the building were silenced tomorrow.** The dissent log over the last 40 decisions:

```
decisions with dissent stated:              6 / 40
  of which from a different block (cross):  6
  of which from within the deciding block:  0
dissents addressed in the decision artifact: 2 / 6
mean directional asymmetry A:                0.31   (exceeds SPEC 4.7 trigger)
```

Zero within-block dissents in forty decisions. Not low. Zero. And `A = 0.31` triggers the anonymization test, which the institute ran: `A` fell to 0.09 with authorship hidden. **Deference collapse (Failure Mode FM-4G) is present and is part of why `u_div` is near zero.**

Headroom check before proposing anything (the headroom analysis [Curzi, 2026b]), and current behaviour.

```
h_0 = eta ln N − delta Q_0 − C_0
    = 0.45∗1.791759 − 0.08∗6.00 − 0.20
    = 0.806292 − 0.480000 − 0.200000 = 0.126292

S_0     = f(0.75) = 1.558145
eta S_0 = 0.45 ∗ 1.558145 = 0.701165
Delta Q = 0.701165 − 0.680000 = +0.021165   -> growing, barely
Q_1     = 0.92∗6.00 + 0.701165 − 0.20 = 6.021165
Q_inf   = (0.701165 − 0.20) / 0.08 = 6.264563
```

`h_0 > 0`: growth is available, and `r_max = h_0 = 0.126292` (the storage-law analysis [Curzi, 2026a], Cor 4.2) — no policy based on interaction geometry alone can guarantee more. The institute grows at 0.021 per quarter toward 6.265, visibly flat:

```
t:    0      1      2      3      4      5      6      8     10     12
Q:  6.000  6.021  6.041  6.059  6.075  6.090  6.104  6.129  6.150  6.167   -> 6.265
```

The wrong intervention. The standing proposal is a seventh block — remote sensing — at a cost of two senior hires. Before pricing it, run the `u_cross` lever within the existing `N`, holding `u_div = 0.03`:

```
u_cross   u      S = f(u)    Delta Q = etaS − deltaQ_0 − C
------------------------------------------------------------
 0.60    0.63    1.423108      −0.039601
 0.72    0.75    1.558145      +0.021165     <− as found
 0.85    0.88    1.686399      +0.078880
 0.97    1.00    1.791759      +0.126292     <− at the ceiling
```

Driving `u_cross` from 0.72 to 0.97 — an enormous restructuring — buys `+0.126292`, exactly `h_0`: the whole headroom, at the highest possible price.

The right intervention. Hold `u_cross` fixed. Buy the scarce source.

The budget constraint (Section 4.3) caps `u_div` at `alpha^2 = 0.28` while `u_cross = 0.72`, and reaching that cap requires within-block `r = 0` on shared ground — full independence between two people in the same discipline. The institute will not hit it. The arithmetic of partial achievement is what matters:

```
holding u_cross = 0.72, N = 6, delta Q_0 + C = 0.68:

u_div    u      S = f(u)    gain in S   Delta Q per quarter
------------------------------------------------------------
 0.00   0.72    1.526056      +0.0000      +0.006725
 0.03   0.75    1.558145      +0.0321      +0.021165     <− as found
 0.05   0.77    1.578979      +0.0529      +0.030540
```

```
 0.10   0.82    1.629241      +0.1032      +0.053158
 0.15   0.87    1.677097      +0.1510      +0.074693
 0.20   0.92    1.722767      +0.1967      +0.095245
 0.28   1.00    1.791759      +0.2657      +0.126292
```

Target: protocols moving within-block `r` from 0.985 to roughly 0.71, giving `1 - r^2 = 0.50` and an ensemble `u_div` at the full remaining budget of 0.28. Ambitious; a realistic first year is half of it.

Full target (u_div = 0.28, u = 1.00), against the as-found state:

```
                    S        eta S     Delta Q at t=0     Q_inf
as-found         1.558145   0.701165    +0.021165        6.264563
u_div raised     1.791759   0.806292    +0.126292        7.578650
                                        6.0x faster      +1.314087  (+21.0%)

rho_ceil goes 0.8696 -> 1.000.
Delta Q = 0.126292 = h_0 exactly: the intervention consumes the entire
available headroom and not one unit more. the maximum-increment corollary \citep{ccp:CCP-007} doing its job
r_max = h_t, and running the geometry to u = 1 is the only way to reach it.
```

Twenty-one percent more stored order at steady state, and a six-fold increase in the quarterly growth rate, purchased entirely by permitting people already on the payroll to disagree with colleagues in their own discipline. No hires, no new block, no change to `u_cross`, `delta` or `eta` — except that `C` will rise, which is the next paragraph.

Cost honesty. Structured within-block dissent raises `C` — longer decision cycles, one extra review pass, facilitation time. Estimate `+0.03` per quarter:

```
C_1 = 0.23
h_1 = 0.806292 - 0.08*6.00 - 0.23 = 0.096292      (was 0.126292)
Delta Q at u=1.00 with C=0.23 = +0.096292
Q_inf = (0.806292 - 0.23) / 0.08 = 7.203650
```

Still 7.204 against 6.265: a gain of 0.939, or 15.0 percent. **The intervention survives its own cost by a wide margin, because the surplus gain is bought at the ceiling of the `u` parameter, where the institute had 13 percent of its per-interaction capacity sitting unused.**

Now the seventh-block proposal under the same cost honesty. It raises `eta ln N` from 0.806292 to `0.45 * ln 7 = 0.875660`, but two senior hires raise `C` by an estimated 0.09 (salary is not in `C`; translation burden, extra coordination paths per Brooks, and onboarding are). At the as-found `u = 0.75`, now with `N = 7`:

```
S = ln(1 + 6*0.75) = ln(5.50) = 1.704748
eta S = 0.767137
Delta Q = 0.767137 - 0.480 - 0.29 = -0.002863     -> declining
Q_inf = (0.767137 - 0.29) / 0.08 = 5.964213       -> BELOW the as-found 6.265
```

Adding a seventh discipline to a system with `u_div = 0.03` makes it worse. The ceiling rose, the cost rose, the geometry did not. This is the ceiling analysis's [Curzi, 2026c] warning in its sharpest form: a headroom problem cannot be escaped by adding domains.

The specification the institute adopts.

```
INTERVENTION SPEC — INSTITUTE, QUARTERS 1–4

1. Independent-estimate-before-discussion on every risk determination; scores
   to a locked form that closes before the meeting opens. (SPEC 4.16)
```

```
   Owner: research director.

2. Dissent field mandatory on all decision records, blank not permitted;
   artifact states each dissent's resolution as accepted / rejected-with-reason
   / deferred-with-date. (SPEC 4.14, 10.15) Owner: secretariat, chair.

3. Artifacts anonymized at peer evaluation, per the deference finding
   (A fell 0.31 -> 0.09 under anonymization).

4. Rotating within-block adversarial assignment, one quarter per person,
   protected and voluntary; declining carries no penalty and is recorded as a
   decline, not a failure. (SPEC 4.18) Owner: block leads.

5. Quarterly re-run of Instruments A and B, reported per SPEC 4.21. C
   monitored monthly; if C exceeds 0.26, pause item 4 and re-cost, because
   h_t must not go below zero (the headroom analysis \citep{ccp:CCP-008}, Cor 3.4).
```

Falsification criteria (the preregistration protocol [Curzi, 2026i] form), preregistered and frozen.

```
PREDICTION: within-block r on shared ground falls from 0.985 to below 0.85
            within four quarters, giving u_div >= 0.10 and S >= 1.629.

FALSIFIED IF, at quarter 4:
  (a) r remains above 0.95, OR
  (b) u_div < 0.06 despite full protocol compliance, OR
  (c) Y (the order proxy) has not increased relative to the frozen
      persistence baseline, OR
  (d) the shuffled-input control fits as well as the true input series.

REQUIRED BASELINES: persistence; direct-input; matched autoregressive;
                    shuffled-input control.

NOTE ON (c): the order proxy Y here is the count of adopted, maintained
technical standards in force — chosen because it is NOT a repackaging of U.
It contains no measure of disagreement, interaction volume, or staffing
composition. (the order-proxy requirement \citep{ccp:CCP-015}.)
```

What the numbers do not say. `u_cross = 0.72` is real, was expensive, and remains the larger half of the institute's surplus by a factor of eleven. Nothing here argues against cross-disciplinary staffing. The finding is narrower and harder: **the institute completed one of two purchases, declared the programme finished, and left the cheaper half of its own portfolio on the table for eleven years — because the half it bought is visible on a spreadsheet and the half it did not is visible only in behaviour nobody was recording.**

That is a measurement failure before it is a design failure. Section 5 shows what happens when an organization does not merely fail to buy `u_div` but actively loses the `u_div` it started with — and why that makes difference a maintenance obligation rather than an acquisition.

10.7.2 7.2 A standards review board over twenty-three quarters

A fifteen-member technical review board at a national standards body, reviewing safety determinations. Five declared specialties. It was constituted eleven years ago and was, at constitution, genuinely diverse: an early `u` audit conducted in year two returned `u = 0.55`.

Parameters (the storage-law estimation procedures [Curzi, 2026a,b] procedures, quarterly).

```
N = 5        ln N = 1.609438
delta = 0.09     eta = 0.55     C = 0.15     Q at year-2 audit = 4.00
eta ln N = 0.885191            <- the per-period ceiling on eta S
```

Year-2 state.

```
u = 0.55   N_eff = 1 + 4(0.55) = 3.200   (S3 badged: design axiom)
S = ln(3.200) = 1.163151   rho_ceil = 1.163151/1.609438 = 0.7227
eta S = 0.639733    delta Q + C = 0.36 + 0.15 = 0.51
Delta Q = +0.129733   -> growing
Q_inf at this u = (0.639733 - 0.15)/0.09 = 5.441478
```

Healthy. Growing at 0.13 per quarter toward 5.44, with 28 percent of its per-interaction ceiling unused. Nobody had any reason to look further.

What happened next, and why nobody saw it. No one left. No policy changed. Three things happened that no one recorded as events:

1. Scores were collected *in the meeting*, in seniority order, rather than in advance.
2. Second reviewers were routinely given the first reviewer's assessment "for context."
3. All fifteen members' committee stipends were tied to a single throughput metric — determinations closed per quarter.

Items 1 and 2 are cascade substrate (Section 5.10). Item 3 is a common signal (Failure Mode FM-5H). Together they produce a measured `lambda_hat = 0.18` per quarter — half-life 3.49 quarters — which is the decay run computed in Section 5.13 and reproduced here at the points that matter:

t (quarters)	u	N_eff	S	rho_ceil	Q	
0 (yr 2)	0.5500	3.200	1.1632	0.7227	4.0000	
2	0.3698	2.479	0.9080	0.5642	4.1751	<- Q peak
4	0.2487	1.995	0.6905	0.4290	4.0623	
8	0.1124	1.450	0.3714	0.2308	3.3166	
12 (yr 5)	0.0508	1.203	0.1851	0.1150	2.3015	
16	0.0230	1.092	0.0879	0.0546	1.3240	
20	0.0104	1.042	0.0407	0.0253	0.5110	
23	0.0057	1.023	0.0227	0.0141	0.0248	

The board's stored order peaked two quarters after its geometry began to collapse, and every report available to its chair during those two quarters said the board was improving.

The audit, run in year 5 (t = 12).

Instrument 1 — collision test. 20 live determinations put to pairs of members separately, in writing. Collisions: 1 of 20. Collision rate 0.05.

Instrument 2 — disagreement audit, last 40 determinations:

```
determinations with any dissent recorded              3 / 40
  dissent from within the deciding specialty          0
  dissent from a different specialty                  3
dissents that changed the final artifact              0 / 3
later-stage reversals of an earlier stage             0 / 40
```

Instrument 3 — error decorrelation. The board issues resolvable severity scores; 34 determinations had subsequent field outcomes. Residual correlation across members: `rho = 0.71`.

```
N_eff_var = 15 / [1 + 14(0.71)] = 15 / 10.94 = 1.371
```

Fifteen reviewers worth 1.37 independent reviewers. The surplus-law reading agrees: `u = 0.0508` gives `N_eff = 1.203`. **Two instruments built on entirely different foundations both report a board of fifteen operating as a board of roughly one and a third.**

Instrument 4 — pre-mortem. Written independently on a pending determination. Fifteen failure stories; thirteen of them named the same mechanism. Recorded as a low-`u` finding per SPEC 5.15.

Instrument 5 — shuffled-input control. The minimal model was fitted on eleven years of quarterly data with `Y` = the count of standards in force that had been revised on schedule (chosen because it is not a repackaging of `U` — it contains no measure of disagreement, interaction volume, or composition; the order-proxy requirement [Curzi, 2026i]). The shuffled-input control fitted materially worse than the true series, so the input geometry is doing real work in this system and the analysis is not spurious. Had it fitted equally well, everything above would have been withdrawn.

Diagnosis. Failure Mode FM-5C (cascade collapse), driven by FM-5H (monometric incentive), presenting as FM-5A (parallel wall). Not a hiring failure. The board still contains the diversity it was constituted with; it has stopped being able to express it.

Sustainability check before proposing anything.

```
h_12 = eta ln N − delta Q_12 − C
     = 0.885191 − 0.09(2.3015) − 0.15
     = 0.885191 − 0.207135 − 0.15
     = 0.528056
```

`h_12` is large — 0.528 — and rising, because `Q` has collapsed and `delta Q` with it. **The board has abundant headroom and is dying, which is exactly the combination Section 5.13 predicted and exactly why headroom must never be read alone.** The binding constraint is not capacity. It is `rho_ceil = 0.115`: the board is running at eleven percent of its per-interaction ceiling.

Intervention.

```
INTERVENTION SPEC — STANDARDS REVIEW BOARD

A. lambda-reduction (the leak), effective immediately:
   A1. Severity scores submitted to a locked form that closes 24 hours before
       the meeting. No score visible to any member until all are in.
       (SPEC 5.11)
   A2. Second reviewers no longer receive the first reviewer's assessment
       before submitting. (SPEC 5.12)
   A3. Chair and senior members speak last on substance. (SPEC 5.13)
   A4. Stipend decoupled from the throughput metric; throughput retained as a
       reported statistic with no reward attached. (SPEC 5.17)
   A5. Minority-report field mandatory on every determination artifact.
       (SPEC 5.14)  Target: lambda <= 0.03.

B. u-restoration (the level), quarters 1–4:
   B1. Rotating within-specialty adversarial assignment, protected and
       voluntary; declining is recorded as a decline. (SPEC 5.16)
   B2. Two members per quarter seconded to review outside their specialty.
   B3. Source-overlap index measured and inputs deliberately varied.
       (SPEC 5.9)  Target: u >= 0.42.

C. Cost control:
   C1. C monitored monthly. Estimated increase +0.04 per quarter.
       At C = 0.19, h = 0.885191 − 0.207135 − 0.19 = 0.488056. Ample.
   C2. One review stage deleted: the pre-circulation summary stage, which had
       produced zero reversals in 40 determinations (SPEC 5.18). Estimated
```

```
      C saving −0.02, offsetting half the increase.
```

Predicted trajectory. With `u` restored to 0.42 and `lambda` held at 0.03:

```
S at u=0.42 = ln(1 + 4(0.42)) = ln(2.68) = 0.985817
eta S = 0.542199   rho_ceil = 0.6125
Delta Q at t=12 = 0.542199 − 0.207135 − 0.15 = +0.185064

 t          u        Q  (lambda = 0.03)      Q  (lambda = 0, maintained)
-------------------------------------------------------------------------
12       0.4200        2.3015                     2.3015
16       0.3722        2.8890                     2.9480
20       0.3298        3.1519                     3.3910
24       0.2922        3.1980  (max at t=23)      3.6950
32       0.2294        2.9190                     4.0460
40       0.1801        2.4110                     4.2110
                                            Q_inf = 4.3578
```

This is the decision the board is actually making. Both columns start from the same intervention. The left column leaks three percent per quarter and is back near its starting point after seven years. The right column holds and converges to 4.358, an 89 percent gain.

The difference between them is not effort, budget, or personnel. It is whether items A1 through A5 are still being performed in year eight. **The intervention that matters is not the one that raises `u`; it is the one that keeps `u` from leaving, and it is the one most likely to be quietly dropped, because it is procedural, invisible when working, and mildly annoying to everyone every single week.** Star and Ruhleder [1996]'s law applies to protocols as much as to plumbing.

Falsification criteria, preregistered and frozen (the preregistration protocol [Curzi, 2026i] form).

```
PREDICTIONS, at quarter 4 of the intervention:
  P1. measured u >= 0.30
  P2. measured lambda <= 0.06
  P3. error correlation rho <= 0.45, so N_eff_var >= 2.1
  P4. determinations with recorded, artifact-changing dissent >= 6 / 40

FALSIFIED IF:
  (a) u < 0.20 despite full protocol compliance, OR
  (b) lambda > 0.12 despite full compliance, OR
  (c) Y has not improved against the frozen persistence baseline, OR
  (d) the shuffled-input control fits as well as the true input series.

REQUIRED BASELINES: persistence; direct-input; matched autoregressive;
                    shuffled-input control.

ORDER PROXY Y: count of standards in force revised on schedule. Justified
as NOT a repackaging of U -- it contains no measure of disagreement,
interaction volume, or board composition. (the order-proxy requirement \citep{ccp:CCP-015}.)

COMPLIANCE IS A MEASURED VARIABLE, NOT AN ASSUMPTION. Protocol adherence
is logged per determination. A failed prediction under partial compliance
falsifies nothing and shall be reported as an implementation failure, not
as evidence against the model.
```

What this example does not claim. It does not claim the board was badly constituted, badly staffed, or badly intentioned; the year-2 audit says otherwise on all three. It does not claim

the surplus law explains the board's performance — that is what the shuffled-input control is there to test, and it could have failed. It claims one thing, narrowly and with arithmetic:

a group that was genuinely diverse became arithmetically indistinguishable from a single reviewer, in five years, through three procedural choices that nobody recorded as decisions — and the only reason it was caught is that someone measured a quantity most institutions have never measured at all.

Section 6 leaves the roster and moves to the room. The reason follows directly from Section 5: several of the strongest `lambda`-reduction devices — speaking order, sightline reciprocity, whether the quiet voice can be heard at all — are not policies. They are properties of furniture and acoustics, and they enforce themselves continuously without anyone administering them.

10.7.3 7.3 A departmental common room

A departmental common room, 120 m^2, clear height 2.8 m, volume 336 m^3, on the second floor of a research building. It serves roughly 90 people across four teams. It contains a kitchenette with a kettle and a fast drip machine, six four-person tables with 24 stacking chairs, a notice board, and a vending machine. Floor is sealed screed; ceiling is exposed soffit with services; two walls are glazed. It was refurbished eighteen months ago and the department describes it as "underused despite being lovely."

Step 1 — Physical inventory (worksheet A)

Item	Measured
A1 Floor area	120 m^2
A2 Volume	336 m^3
A8 Linear sittable metres	6.0 m (six tables x 4 chairs, chair width ~0.45 m, but only the chairs are sittable; the window ledges are 220 mm deep and therefore not sittable under SPEC 6.3)
A9 Whyte ratio	6.0 / 120 = 0.050 m/m^2
A10 Movable / total seats	24 / 24, but all are identical stacking chairs at identical tables
A15 Median access time	1.5 min

Requirement SPEC 6.1 demands 120 x 0.109 = 13.1 m of sittable edge. The room has 6.0 m. Deficit: 7.1 m. The room is at 46% of the Whyte floor.

Note the trap in A10. Every chair is movable, so the room scores 100% on movability, and yet it offers no seating *choice*, because all 24 seats are the same seat at the same table at the same height under the same light. **Movability without variety satisfies the letter of SPEC 6.2 and fails its purpose. SPEC 6.4's category audit is the check that catches it.**

Step 2 — Category audit (worksheet B)

#	Posture	Purpose	Condition	Available?
1	Sitting upright, standard chair	Eating / pair talk	Bright, live, open	Y

#	Posture	Purpose	Condition	Available?
2	Sitting upright, standard chair	Solo focus	Bright, live, open	Y — but identical to 1 in every respect that matters to a user
3	Standing	Queue at kettle	Bright, live, open	Y
4	Sitting reclined	Anything	–	N (no such seat exists)
5	Sitting, enclosed / backed	Small-group deliberation	Quiet	N
6	Perching	Brief transit stop	–	N (ledges 220 mm)

Rows 4, 5 and 6 struck as unavailable. Row 2 struck as a duplicate of row 1 under the Section 6.10 strike rule. Surviving categories: **2**.

```
N   = 2
ln N = 0.6931 nats
```

The room's per-interaction ceiling in the use-category dimension is 0.693 nats — 33% of what a six-category room would offer (`ln 6 = 1.792`).

Step 3 — Activity count (worksheet C)

Five fixed observation times per day (09:30, 11:00, 12:45, 14:30, 16:00), five days, marked on a plan. Counts are person-observations summed over the week.

Day	NEC	OPT	SOC	OAR
Mon	41	7	3	0.146
Tue	38	5	2	0.116
Wed	44	9	4	0.170
Thu	36	6	2	0.143
Fri	52	21	11	0.288
Total	**211**	**48**	**22**	**0.185**

`SOC / OPT = 22 / 48 = 0.458`.

Read the Friday anomaly. Friday is the day the department buys pastries. **The single largest determinant of optional activity in this room is whether there is food in it, and it is not close.** That is a Whyte finding reproduced, by counting, in one week, by one person.

Necessary activity is 211 observations — people come to get coffee and leave. OAR of 0.185 is a corridor-with-chairs signature (FM-6.1) with a food-driven Friday exception.

Median seated dwell, timed from the time-lapse: **4.1 minutes**.

Step 4 — Contact-pair audit (worksheet D)

Conversational minutes observed over the week: 412. Minutes between pairs differing on at least two of {team, discipline, training lineage, reporting line}: 96.

```
p_cross = 96 / 412 = 0.233     (badged: a crude u proxy, see Section 6.10)
N_eff   = 1 + (2 − 1) x 0.233 = 1.233
```

```
F       = ln(1.233) = 0.2095 nats
F / ln N = 0.2095 / 0.6931 = 0.302
```

The room reaches 30% of its own (already low) ceiling. Note that `N = 2` caps this hard: even at `u = 1` the room could deliver only 0.693 nats. **Raising `u` in this room cannot help much, because `N` is the binding constraint.** This is the diagnosis the arithmetic is for, and it is the opposite of the intervention the department had planned (a "mixer" social event).

Step 5 — Storage terms (worksheet E)

`Q` proxy, declared: **standing cross-team working agreements attributable to contact in this room** — a countable list, maintained by the department administrator, with a named origin for each entry.

Decay measurement: the room was closed for three weeks over the summer shutdown, giving a genuine zero-input period. Agreements standing before: 9. Agreements still active after (three periods, period = one week… the shutdown supplies a three-period decay):

```
Q_0 = 9,  Q_3 = 7
(1 − delta)^3 = 7/9 = 0.7778
1 − delta = 0.7778^(1/3) = 0.9192
delta = 0.0808 per week
half-life = ln2 / (−ln(0.9192)) = 0.6931 / 0.08418 = 8.23 weeks
```

`S` proxy, declared: **observed social contacts per week** (SOC count) = 22.

Attributable `dQ` per week, from the administrator's log: **1.0** new standing agreement per week originating in this room.

```
eta_hat = 1.0 / 22 = 0.0455
```

Step 6 — Cost tally (worksheet F)

Converting a cost tally into the same units as `eta S` requires a declared conversion, and the declaration is the honest part. Everything here is expressed as *equivalent lost agreements per week*, the same unit as `Q`, and state the conversion openly.

Component	Raw measure	Converted (agreements/week equivalent)	Basis of conversion
`C_access`	1.5 min x 60 visits/wk = 90 min	0.10	15 min of aggregate time judged equal to 1/60 of an agreement's setup effort
`C_setup`	negligible (no equipment)	0.00	–
`C_effort`	RT60 measured at 1.6 s (the acoustic analysis [Curzi, 2026e]), background 47 dBA	0.55	the acoustic exclusion arithmetic [Curzi, 2026e] exclusion arithmetic
`C_interrupt`	0 (room is acoustically separate)	0.00	–

Component	Raw measure	Converted (agreements/week equivalent)	Basis of conversion
`C_exclusion`	included in `C_effort` above	–	do not double-count
`C_maint`	2 h/wk cleaning	0.05	–
`C_total`		**0.70**	

The conversion is a modelling choice and it is the weakest link in this calculation. State it, sensitivity-test it, and do not present the resulting `h` to three decimal places as if it were measured. The correct use of this number is comparative — before versus after — not absolute.

Step 7 — Headroom (worksheet G)

```
eta ln N  = 0.0455 x 0.6931 = 0.0315
delta Q   = 0.0808 x 9      = 0.7272
C         =                   0.70
h = 0.0315 − 0.7272 − 0.70  = −1.3957
```

`h < 0`. Unsustainable in this model. No arrangement of people in this room prevents the stock of cross-team agreements from decaying. The department's plan — a monthly mixer — would add `S`, and `S` enters the budget multiplied by `eta = 0.0455`. To close a gap of 1.396 by adding `S` alone would require an implausible increase, and in any case `S` is capped: `eta ln N = 0.0315` is the *maximum* the term can ever reach at `N = 2`, no matter how much interaction occurs.

This is the sustainability corollary [Curzi, 2026a] doing its job. It states, before anything is spent, that the planned intervention cannot work.

Step 8 — The intervention, chosen by the arithmetic

The binding constraints, in order of magnitude: `delta Q = 0.727`, `C = 0.70`, and an `eta ln N` term so small it is almost irrelevant. Three moves follow directly.

Move 1 — Raise `N` from 2 to 5. This is the cheapest large move available because `N = 2` is capping everything.

- Deepen the two window ledges from 220 mm to 450 mm over 9 linear metres (perching, transit stop, bright, solo). New category.
- Install one high table with four stools and a rail (standing, pair work, live). New category.
- Install two backed banquettes forming a pair of 4-person alcoves against the solid wall, with local dimmable light (sitting enclosed, small-group deliberation, quiet). New category.
- Retain the existing tables (category 1) and the kettle queue (category 3).

```
N: 2 -> 5        ln N: 0.6931 -> 1.6094
```

Move 2 — Cut `C` by 0.55. The entire `C_effort` term is acoustic. the acoustic analysis's [Curzi, 2026e] treatment for a 336 m^3 room brings RT60 from 1.6 s to under 0.7 s and background from 47 dBA to 40 dBA. Cost is one absorptive ceiling raft array and 18 m^2 of wall panel.

```
C: 0.70 -> 0.15
```

Move 3 — Shed `Q`. Of the nine standing agreements, four are dormant: they name people who have left, or processes that ended. They are consuming `delta Q` — 0.323 agreements/week of maintenance budget — and returning nothing. Retire them formally.

```
Q: 9 -> 5        delta Q: 0.727 -> 0.404
```

Move 3 is the one nobody proposes, and the headroom analysis [Curzi, 2026b] named it: **shedding stored order to regain headroom is a legitimate and underused move.**

Alongside these, SPEC 6.1 is satisfied: the deepened ledges (9 m) plus high table (4 stools ~1.8 m) plus banquettes (8 seats ~3.6 m) plus existing chairs (6.0 m) gives 20.4 m of sittable edge against the 13.1 m requirement. SPEC 6.7 is satisfied by relocating the notice board to the kettle queue and converting it to a live, changing display. SPEC 6.16 is satisfied by replacing the fast drip machine with a bean-to-cup unit whose cycle is 55 seconds — deliberately slower, per Section 6.6.

Step 9 — Predicted outcome and what would falsify it

Recomputing headroom with `eta` held at its measured value (a rise in `eta` is predicted but not assumed):

```
eta ln N  = 0.0455 x 1.6094 = 0.0732
delta Q   = 0.0808 x 5      = 0.4040
C         =                   0.15
h = 0.0732 - 0.4040 - 0.15  = -0.4808
```

Still negative. **Say so.** The three moves close the gap from -1.396 to -0.481 — 66% of the way — and the room is still, in this model, not able to grow its stock of agreements.

The remaining gap is dominated by `delta Q = 0.404` against an `eta ln N` of only 0.073. The structural reading: **this room's `eta` is an order of magnitude too small for the `Q` it is being asked to hold.** At `eta = 0.0455`, sustaining even `Q = 5` at `delta = 0.081` would require `eta ln N >= 0.404`, hence `ln N >= 8.88`, hence `N >= 7200` categories. Absurd. The room cannot be the mechanism that sustains these agreements.

That is the finding, and it is more valuable than a success would have been. **The agreements are not sustained by the room; they must be sustained by an artefact.** the conversion analysis's [Curzi, 2026b] answer applies: the agreements need a written register with a named owner and a review cadence, which raises `eta` at the *organisational* level and lowers `delta` by moving `Q` out of memory and into a document. The room's job is to generate contact — which the three moves genuinely improve — not to store the result.

Recompute once more, with the register in place. The register is measured, after twelve weeks, to have moved `delta` from 0.0808 to 0.030 per week (agreements now decay only when genuinely obsolete, not when forgotten) and `eta_hat` from 0.0455 to 0.098 (contacts now have somewhere to land):

```
eta ln N  = 0.098 x 1.6094  = 0.1577
delta Q   = 0.030 x 5       = 0.1500
C         =                   0.15
h = 0.1577 - 0.1500 - 0.15  = -0.1423
```

Closer. Still negative, and honestly so. Adding the observed post-intervention `SOC` rise (22 to 56 contacts/week — see below) does not change `h`, because `h` is a budget on the *ceiling*, not on the realised input; but it does change whether the realised `S_t` approaches that ceiling.

The forecast the department should be given is therefore: **this room, treated, will roughly double its optional activity and roughly double its `eta`; it will not on its own make the department's stock of cross-team agreements grow, and anyone who promises that is not reading the arithmetic.**

Preregistered predictions, per the preregistration protocol [Curzi, 2026i]:

#	Prediction	Measured how	Falsified if
P1	OAR rises from 0.185 to >= 0.30 within 8 weeks	Worksheet C, same 5 times, same 5 days	OAR < 0.25 at week 8
P2	Median seated dwell rises from 4.1 to >= 8 min	Time-lapse	< 6 min at week 8
P3	`SOC` count at least doubles	Worksheet C	< 1.5x at week 8
P4	`p_cross` rises (the alcoves and rail mix teams)	Worksheet D	p_cross < 0.25 at week 8
P5	RT60 at 500 Hz occupied <= 0.70 s	the acoustic analysis [Curzi, 2026e] method	> 0.80 s at commissioning
P6	`h` improves by at least 0.8	Worksheet G	improvement < 0.5
P7	The Friday food effect persists but shrinks as a fraction of weekly OPT	Worksheet C	Friday remains > 40% of weekly OPT

P7 is the control that matters. **If all the gain is still on pastry day, the room was not fixed; the catering was.** That is the shuffled-input control of the preregistration protocol [Curzi, 2026i] in its most concrete possible form.

Step 10 — Measured outcome at week 8

Day	NEC	OPT	SOC	OAR
Mon	43	19	9	0.306
Tue	40	17	8	0.298
Wed	45	24	13	0.348
Thu	38	18	9	0.321
Fri	50	31	17	0.383
Total	**216**	**109**	**56**	**0.335**

```
OAR:  0.185 -> 0.335   (x1.81)     P1 PASSES (>= 0.30)
SOC:  22    -> 56      (x2.55)     P3 PASSES
Friday share of weekly OPT: 21/48 = 43.8%  ->  31/109 = 28.4%
                                               P7 PASSES (< 40%)
eta_hat: 1.0/22 = 0.0455  ->  4.0/56 = 0.0714   (x1.57)
```

Necessary activity is essentially unchanged (211 to 216), which is exactly what Gehl's taxonomy predicts: **necessary activity does not respond to quality, and its stability across the intervention is evidence that the observation protocol was not itself the cause.** That stability is the strongest internal validity check the worksheet provides, and it is free.

The room now reaches `N = 5`, `p_cross` measured at 0.31, hence `N_eff = 1 + 4 x 0.31 = 2.24`, `F = 0.807` nats against a ceiling of 1.609 — 50% of ceiling, up from 30%.

What the room case demonstrates

1. **The worksheet found the binding constraint in one week, and it was not the one the department believed.** They believed the problem was social; the arithmetic said `N = 2` and `delta Q` were binding.
2. **The intervention chosen by the arithmetic was cheaper than the one that had been planned** — deeper ledges, a rail, two banquettes and a slower coffee machine, against a recurring events programme.
3. **The honest answer was partly negative and was reported as such.** `h` improved substantially and remained below zero, and the correct conclusion was that the room is a contact generator, not a storage device.
4. **The Friday control did the real work.** Without P7, a department could have claimed the refurbishment succeeded when all it had done was buy pastries.

the passive-geometry analysis [Curzi, 2026e] takes the sub-division problem this example raised — how to get five categories and four alcoves into 120 m^2 without building walls — and solves it with geometry. the acoustic analysis [Curzi, 2026e] takes the `C_effort` term, which was the largest single cost line in this survey, and reduces it to a computation.

10.8 8 Discussion

10.8.1 8.1 What follows

The decomposition converts a slogan into a purchase order. "More diversity" is not an instruction, because it does not say which of two goods to buy, and the two goods have different prices, different suppliers, different observability and different half-lives. Once u is split, the question becomes answerable in one line: compute both components, compute the balance index, and buy the scarcer one. Rule 4.4 is self-correcting because each lever's marginal productivity in its own design variable is highest exactly where the other has not been pulled.

The budget constraint is the result most likely to be new to a practitioner and it is the least negotiable. $u_{\text{div}} \leq 1 - u_{\text{cross}}$ is not an empirical regularity; it follows from $\alpha^2 + \beta^2 = 1$. Its practical content is that a heavily cross-staffed body has already spent the shared ground on which its members could argue, and the same dissent protocol that is worth 1.3930 nats in a monodisciplinary group is worth 0.0780 in one at $u_{\text{cross}} = 0.90$. **The intervention did not get worse; the room for it did.**

The decay result reframes the whole activity. Almost every organisation staffs the acquisition question and leaves the maintenance question unstaffed, and the reason is legibility rather than neglect: a roster photographs well and a disagreement rate does not. Section 5.18's comparison is the strongest form of the argument. Two futures that differ in nothing except whether a procedural regime is still being performed in year eight end at 2.411 and 4.211 respectively. **The durable quantity is not u but $1/\lambda$, and $1/\lambda$ is bought with protocols that are invisible when they work and mildly irritating every week.**

The room section earns its place by being the cheapest lever with the longest memory. A wall, a marked floor position, a deepened ledge and a slower coffee machine enforce their policy continuously without a budget line and without renewal. That is the honest content of the "circuit element" language — not that a room has physics, but that a physical arrangement holds a constraint at a decay rate far below the decay rate of a memorandum.

10.8.2 8.2 What does not follow

Nothing here establishes that institutions obey the surplus law. The decomposition is algebra; the decay is a badged postulate; the room reading is three badged postulates. The protocol that would test the storage law is specified elsewhere in the corpus [Curzi, 2026i], and its one

confirmatory application **failed its second falsification rule**, beating a matched autoregression by 2.8 per cent on held-out error against a declared 10 per cent margin. That is the correct context for every number in this paper.

No synergy beyond additivity is promised, and the upper bound must not be used as a forecast. $f(u_{\text{cross}}) + f(u_{\text{div}})$ is an impossibility limit. Because f is strictly concave the true value is systematically below it whenever both components are non-zero, by 16 to 45 per cent across the range tabulated in Section 3.3. **A loose upper bound is a proof of impossibility, not an estimate.**

The room does not manufacture non-parallelism. Section 6.12 states the three limits plainly and Section 7.3 demonstrates them: full treatment of the worked room closes 66 per cent of its headroom gap and leaves it negative. **A physical intervention is necessary and frequently insufficient.**

The monotone claim is restricted. Surplus is strictly increasing in u within the model. The measured literature does not agree that institutional value is: Uzzi and Spiro [2005] finds an inverted U in creative output against cohesion; Putnam [2007] reports short-run reductions in trust and civic participation with diversity, a result its author did not want, and Abascal and Baldassarri [2015] is the rebuttal that must be cited with it; Mannix and Neale [2005] and van Knippenberg and Schippers [2007] find surface-level difference often hurts in the short run; May [1972] and Gardner and Ashby [1970] show that increasing diversity and connectance destabilises randomly assembled systems; and Sagan [1993] documents that redundancy can reduce safety. The framework prices an input term and is silent on the cost of extracting it, and the cost side is where those results live. Allport [1954]'s four conditions are the same qualification in the contact literature: difference produces value under conditions, and the conditions must be named.

The word *diversity* is doing three jobs and this paper's u does one of them. Harrison and Klein [2007] separates separation, variety and disparity; u **is a disparity measure and must declare itself as one**, and the pseudo-diversity result of Section 5.8 is exactly what happens when a variety measure is reported as though it were a disparity measure.

10.8.3 8.3 Relation to the surrounding literature

The cascade result is Bikhchandani et al. [1992] and Banerjee [1992]: rational agents observing predecessors' actions abandon their private information, and from that point the sequence stops aggregating. Lorenz et al. [2011] measures the consequence directly — social influence reduces the effective diversity of a population's estimates while raising its confidence, which is the mechanism of Section 5.11 with an experiment attached. Salganik et al. [2006] is the double-edged companion: influence raised both inequality **and unpredictability**, which supports the reflexivity claim and simultaneously undercuts any claim that a reflexive outcome is predictable from initial conditions. Granovetter [1978] is the reason a mean is the wrong summary of a body about to tip, and Schelling [1971] the canonical warning against reading an individual-level parameter straight into an aggregate. DeGroot [1974] supplies the framework's own reflexive map in its original form, and Friedkin and Johnsen [1990] adds the stubbornness term that is formally the same device as a leak — **the sociological literature already contains a decaying-difference model, and this paper's λ should be read against it rather than as a novelty.** The bounded-confidence models [Deffuant et al., 2000, Hegselmann and Krause, 2002, Lorenz, 2007] predict the number of surviving opinion clusters from a coupling parameter, which is an effective count derived from an interaction constant and is prior art for this genre of result.

On the value of difference, Hong and Page [2004] is the theorem the engineering sections lean on and Thompson [2014] is the detailed objection to its conditions; the pair travels together, with Grim et al. [2019] mapping the boundary conditions in simulation and Singer [2019] and Kuehn [2017] supplying the repairs. Woolley et al. [2010] supplies the measured collective-intelligence factor and its turn-taking-equality component, which is the closest empirical relative of Section 6's concern with who can be heard. Ugander et al. [2012] is the

closest empirical analogue of an effective count anywhere in the social sciences: adoption tracks the number of *distinct connected components* among a person's contacts rather than the number of contacts. Markowitz [1952] is the correction the framework should keep printing next to its own results — diversification benefit depends on covariance and not on count — and Longin and Solnik [2001] and Ang and Chen [2002] are the measured warning that the covariance itself moves under stress, so an effective count falls exactly when it is needed. Schindler et al. [2010] is the strongest single demonstration that an effective count buys stability; Doak et al. [1998] is the standing objection that such relationships can arise from statistical averaging with no mechanism, with Tilman et al. [1998] the reply, and Loreau and de Mazancourt [2008] names the decisive variable as **synchrony**, which is the ecological name for what ρ measures here.

On the institutional mechanism, March [1991] is the framework's own thesis published in 1991 with a simulation attached: fast learning homogenises an organisation and destroys the variance that produced its advantage. Section 5.8's five species are that mechanism enumerated, and Levinthal and March [1993] is the same argument about search. Cyert and March [1963] named the reserve this framework calls headroom. Scott [1998] is simultaneously the historical statement of the erosion thesis — legibility destroys the local variety that made a system work — and a warning about this paper's own instrument, since a published u figure is a legibility device. Goodhart [1975] and Campbell [1979] are the specific form that warning takes, and Section 5.17 uses them in an unusual direction: a single dominant metric is a **common signal**, so measure corruption and non-parallelism destruction are the same event seen twice.

On structure, Alexander [1965] supplies Section 4.12's tree-and-semilattice argument in its original form and Jacobs [1961] the mixed-primary-use version of the same claim about cities. Star and Ruhleder [1996] and Star and Griesemer [1989] supply the boundary object, which is the artefact that makes a non-nested membership work, and the observation that infrastructure is invisible until it fails — which is Section 4.9's observability asymmetry, stated generally. Granovetter [1973], Burt [1992] and Burt [2004] give measured effect sizes for the bridging claim, and Centola and Macy [2007] and Centola [2010] give the qualification that matters most for Section 6: **behaviours needing reinforcement spread along dense local ties rather than weak long ones**, so a room that produces repeated local contact is doing something a directory of distant contacts cannot. Perrow [1984] and Buldyrev et al. [2010] are the reason a body should not assume that adding components inside one coordination perimeter improves it.

On the room itself, Whyte [1980] supplies the method and the sittable-space finding, Gehl [1971] the necessary/optional/social taxonomy on which Section 6.3's conversion proxy rests, and Hall [1966] the proxemic bands. All three are observational studies in specific cultures and decades and are used as nominal starting points rather than constants. Sabine [1922] and Schroeder [1965] supply reverberation as an exponential decay with a measurable time constant — **which is the framework's cleanest physical instance of its own storage law**, and the basis of the acoustic cost term this paper defers to Curzi [2026e]. Forrester [1961] and Sterman [1989] are the reason Section 5.13's lag matters: people systematically misjudge stocks with delays, and that is an empirical finding rather than a modelling convenience. Little [1961] and Kingman [1962] supply the queueing results behind the dwell-time argument of Section 6.5.

10.8.4 8.4 Where these results are consumed, and what parallels them elsewhere in the corpus

The two-source decomposition is consumed elsewhere in the corpus and the budget constraint must survive the citation. The affine effective count this paper decomposes is developed in the effective-count analysis, where **a hundred-person department at** $u = 0.05$ **audits to** 5.95 [Curzi, 2026c], and the headroom gauge it is priced against — $h_t = \eta \ln N - \delta Q_t - C_t$, with three regimes and a maximum increment $r_{\max} = h_t$ — is established in the conversion and cost

analysis [Curzi, 2026b] and restated as a precondition on any protocol, worked at $N = 18$ with $\eta = 0.30$, $\delta = 0.06$, $C = 0.55$ and $Q = 1.20$ to give $h = +0.2451$ [Curzi, 2026].

Two corpus results bound what Section 6 may claim. The acoustic analysis prices exclusion against the same ceiling this paper uses: silencing three of eight participants costs $\ln 8 - \ln 5 = 0.4700$ nats, which is 22.6 per cent of the room's per-interaction surplus [Curzi, 2026e] — and that figure is the entire content of the worked room's C_{effort} line. The threshold analysis supplies the impedance result the room's entrances depend on: worst-junction reflection across a hundredfold character change falls from 96 per cent to 13 per cent when the change is divided into six equal-ratio steps [Curzi, 2026f]. And the turn-taking analysis supplies the boundary-cost criterion that decides whether more contact helps at all — a critical multiplier computed as 3.17 at $\rho = 0.45$ against an operating value of 8 [Curzi, 2026g], which is the case in which Section 6's contact-generating interventions would be counter-productive.

The opposition identity recurs across the corpus and is the single result most often misstated. In the frequency framework it appears as $u = 1-\rho^2$ depending only on $|\rho|$, so $\rho = +1$ and $\rho = -1$ both give $u = 0$, $F = 0$ and the same folded note [Curzi, 2026]. The essence framework carries the same count into a pitch-class rule — the pitch class of $f(n) = kn^2$ is the odd part of n, so $1 \dots N$ realises exactly $\lceil N/2 \rceil$ classes, six at $N = 12$ [Curzi, 2026] — and the tuning analysis supplies the roster ceiling $m(m-1) \le 2\eta \ln(N)/\kappa$, whose constructed commons sits exactly at it, 30 against 30.7225, so a seventh member makes the case structural [Curzi, 2026]. **That last result is Section 6.11's occupancy law reached from an entirely different construction, and the agreement is worth stating because the two share no premises beyond the surplus functional.**

Three results in the structural series bound what this paper may claim about buying difference by addition. A second financial free zone in the same country buys capability and not diversification: two enclaves remain one 29.1 per cent line, and 55 per cent of settlement remains dollar-cleared however the entities are arranged [Curzi, 2026]. Additions that raise a naive count by 34.8 per cent lower the effective bloc count by 3.4 per cent, and one leg's measured bloc contribution is 0.002014 — effectively nil [Curzi, 2026]. And a compartmentalisation redesign returns an honest near-zero on the jurisdictional measure, at best $+0.005233$ of surplus, against a leakage of 15.04 per cent that is certain while every benefit computed against it is probabilistic [Curzi, 2026]. **Each is Section 4.7's furnished monoculture in a different domain: a count raised, a geometry unchanged, and the arithmetic declining to credit the difference.**

Finally, the framework's own interlude doctrine constrains this paper more than it licenses it. Category innocence is explicitly refused as a defence, the doctrine forbids moving any harm through a discovered gap in a scheme, and it concedes that the case where a feedback loop runs through another person's will is an unsettled seam the corpus does not close [Curzi, 2026]. **Section 4.14's SPEC 4.18 and Section 5.19's SPEC 5.16 — that no participant may be required to disagree — are that doctrine applied**, and they are the reason the divergence protocols here create opportunity and protection rather than obligation.

10.9 9 Verification

10.9.1 9.1 Dimensional audit

$[Q]$ is the dimension of the declared order proxy — whatever unit it counts — and $[T]$ is time in the system's declared period. Section 6 adds physical dimensions, and they are audited in the same table because that is where the paper is most exposed.

Symbol	Dimension	Unit	Status
α, β	$[1]$	none	**defined**; components of a unit vector, $\alpha^2 + \beta^2 = 1$

Symbol	Dimension	Unit	Status
γ	$[1]$	radians (an angle is a ratio of lengths)	**estimated** from a correlation
$u, u_{\text{cross}}, u_{\text{div}}$	$[1]$	none	**defined** dimensionless, each in $[0, 1]$
r	$[1]$	none	**measured** correlation on shared-ground scores
N	$[1]$	count of blocks	**declared**, by evaluative standard
$f(u) = \ln(1 + (N-1)u)$	$[1]$	nats	**derived**; a logarithm of a pure number
B	$[1]$	none	**derived**; a ratio of two quantities in nats
ρ_{ceil}	$[1]$	none	**derived**; a ratio of two quantities in nats
λ	$[1]$	per period	**estimated**; dimensionless per step, as δ is
$t_{1/2}$	$[T]$	periods	**derived** from λ
η	$[Q]$ per nat	proxy units per nat	**estimated**; *not* dimensionless
δ	$[T]^{-1}$	per period	**estimated**
C	$[Q]\,[T]^{-1}$	proxy units per period	**measured**, denominated in Q-units by requirement
h	$[Q]\,[T]^{-1}$	proxy units per period	**derived**
Room area, volume, height	$[L]^2, [L]^3, [L]$	m^2, m^3, m	**measured**
Whyte ratio	$[L]^{-1}$	m of edge per m^2 of floor	**derived**; $0.109\ \text{m}^{-1}$
RT_{60}	$[T]$	s	**measured**
Observed activity ratio	$[1]$	none	**measured**; a ratio of counts

Does the surplus equation balance? f takes the logarithm of $1 + (N-1)u$, which is a pure number because u is a squared cosine of unit vectors and N is a count. It balances, and it balances because N counts *categories* rather than people.

Does the decomposition balance? $u = \beta^2 + \alpha^2 \sin^2 \gamma$ adds two dimensionless quantities to a dimensionless quantity. The constraint $\alpha^2 + \beta^2 = 1$ is what makes the sum lie in $[0, 1]$, and it is therefore the dimensional statement as much as the structural one.

Does the storage equation balance? Only if η carries $[Q]$ per nat and C is denominated in Q-units per period. Section 7.3's cost tally is the place this bites hardest: **it converts minutes, decibels and cleaning hours into "agreements per week", and that conversion is a stipulation, not a measurement.** The worked room states its conversion basis line by line and Section 10 records that the resulting h may not be read as a measured quantity.

Where the framework has no operational unit, this audit says so. Q counts whatever the declared proxy counts and two systems' Q values are not comparable unless they declare the same proxy. The underlying construct — comfort — has **no operational unit at all**. **The nat is not a physical unit**: it is the unit of a logarithm to base e, and reporting in bits would multiply

every surplus figure here by $1/\ln 2 = 1.442695$ and change nothing. **And the room's physical units do not transfer to the law's units by any measurement reported here**; the bridge is the stipulated conversion above, which is the weakest joint in Section 6 and is named as such at both ends.

10.9.2 9.2 Postulates

#	Statement	Status
P1	$u = 1 - (x \cdot y)^2$ for unit vectors	**DEFINITION**
P2	$f(u) = \ln(1 + (N-1)u)$ under axioms S1–S4	**CITED** within the corpus — Curzi [2026]
P3	$g(u) = e^{f(u)}$ is affine in u (S3)	**DESIGN AXIOM** — chosen, not derived
P4	a participant's state lies in a block-structured space and the declared categories are the blocks	**POSTULATED** (Section 3.8; falsifiable by category shuffling)
P5	$\sin^2\gamma$ is estimated as $1 - r^2$ on independently scored shared-ground items	**MODELLING CHOICE** (Section 4.10)
P6	$u_{t+1} = (1-\lambda)u_t$ absent maintenance	**POSTULATED** (Postulate P-U, Section 5.5)
P7	a bounded volume is a scale at which the storage law's form applies	**POSTULATED** (Postulate P-R1, Section 6)
P8	the optional-activity ratio proxies η	**POSTULATED** (Postulate P-R2, Section 6)
P9	$h_t = \eta \ln N - \delta Q_t - C_t$, and $h_t < 0$ forbids a nondecreasing step	**CITED** within the corpus — Curzi [2026a,b]

P1 and P2 carry no empirical weight. P3 is the load-bearing choice and is badged at every use. **P4 through P8 are the five places where this paper touches a real body or a real room, and every one of them is a postulate or a modelling choice.** A reader tallying the chain below should note that its most consequential outputs — the balance index, the decay rate, the room's headroom — each consume at least one of them.

10.9.3 9.3 The compounding chain

```
P1 + P4            -> L1   y = alpha*(in-block) + beta*(out-of-block),
                           so x.y = alpha cos(gamma)                    ARITHMETIC
L1 + P1            -> R1   u = 1 - alpha^2 cos^2(gamma)
                             = beta^2 + alpha^2 sin^2(gamma)            ARITHMETIC
R1 + P4            -> R1'  alpha^2 + beta^2 = 1, so
                           u_div <= 1 - u_cross                         ARITHMETIC
R1' at u_cross=.90 -> R1'' max u_div = 0.10                             ARITHMETIC
P2 + R1            -> R2   df/du_cross = df/du_div = (N-1)/(1+(N-1)u)   ARITHMETIC
R1                 -> R2'  du/d(beta^2) = cos^2(gamma);
                           du/d(sin^2 gamma) = alpha^2                  ARITHMETIC
P2 (concavity)     -> L2   f(u_c + u_d) <= f(u_c) + f(u_d), strict when
                           both are positive                            CITED (CCP-A02)
```

```
L2 at N=8, 50/50   -> L2'  u = 1: f = 2.079442 against a bound of
                           3.008155; overstatement 44.7%                    ARITHMETIC
P1                 -> R5   u = 1 - rho^2 depends on |rho| only, so
                           rho = -1 gives u = 0 exactly as rho = +1 does    ARITHMETIC
R5 + P9 at eta=.35,
  delta=.10, c=.05 -> R5'  seat contribution (0 - 0.05)/0.10 = -0.5000
                           for BOTH rho = +1 and rho = -1                   ARITHMETIC
P5 + P4            -> R3   institute: u_cross = 0.72, u_div = 0.03,
                           u = 0.75; f(0.72) = 1.526056,
                           f(0.03) = 0.139762, f(0.75) = 1.558145           ARITHMETIC
R3                 -> R3'  B = (1.558145 - 1.526056)/1.558145 = 0.020594    ARITHMETIC
R3 + P9            -> R3'' rho_ceil = 1.558145/1.791759 = 0.869623          ARITHMETIC
R3 + P9            -> R4   N = 7 at the same u: f = 1.704748,
                           dQ = 0.767137 - 0.480000 - 0.290000
                              = -0.002863; Q_inf falls 6.264563 -> 5.964208 ARITHMETIC
P6                 -> R6   u_t = u_0(1-lambda)^t; t_half =
                           ln2 / (-ln(1-lambda)) = 3.4928 at lambda = 0.18  ARITHMETIC
R6 at u_0 = 0.55   -> R6'  u_23 = 0.55(0.82)^23 = 0.005738                  ARITHMETIC
R6 + P9            -> R7   Q peaks at t = 2 while u peaks at t = 0:
                           dQ(1) = +0.045403, dQ(2) = -0.026376             ARITHMETIC
P7 + P8 + P9       -> R8   room: h = 0.031536 - 0.727200 - 0.700000
                              = -1.395664; after three moves
                           h = 0.073228 - 0.404000 - 0.150000 = -0.480772   ARITHMETIC
R8                 -> R8'  gap closed = (1.395664 - 0.480772)/1.395664
                              = 0.6555                                      ARITHMETIC
```

The four checks a reviewer can redo in a minute are: $\alpha^2 + \beta^2 = 1$ forces $u_{\text{div}} \le 1 - u_{\text{cross}}$; $\ln 4.60 = 1.526056$ and $\ln 4.75 = 1.558145$, whose difference over the second is 0.020594; $(0 - 0.05)/0.10 = -0.5$ identically at $\rho = +1$ and $\rho = -1$; and $\ln 2/(-\ln 0.82) = 3.4928$.

And the chain must be read for what it is. R1, R2, R5 and L2 are algebra and hold wherever the encoding does. R3, R4, R6, R7 and R8 consume postulates P4 through P8 and are **exact computations on stipulated inputs, not measurements of anything.** The balance index of 0.0206 is a true statement about a constructed institute and no statement at all about any real one.

10.9.4 9.4 The decomposition checked on the smallest case

The smallest case that exercises the decomposition is $N = 2$ blocks and a single pair. Let x be a unit vector in block 1 and write $y = \alpha\, e_\gamma + \beta\, e_\perp$, with e_γ a unit vector in block 1 at angle γ to x and $e_\perp$ a unit vector in block 2. Then $\|y\|^2 = \alpha^2 + \beta^2 = 1$ as required, and because $x \perp e_\perp$,

```
x . y  =  alpha (x . e_gamma)  =  alpha cos(gamma)
u      =  1 - alpha^2 cos^2(gamma)
       =  (1 - alpha^2) + alpha^2 (1 - cos^2 gamma)
       =  beta^2 + alpha^2 sin^2(gamma)
```

Status: **ARITHMETIC**. Two limits check it. At $\beta = 0$ the second party lies wholly inside the first's block and $u = \sin^2\gamma$, which is the pure-divergence case. At $\beta = 1$ the second party lies wholly outside and $u = 1$ independent of γ, which is the pure-crossing case — **and note that γ is then undefined rather than large, which is the exact formal content of Section 4.12's claim that a strict tree admits no portfolio.**

The upper bound's overstatement was recomputed for this paper at $N = 8$. At a balanced split and $u = 1$, $f(1) = \ln 8 = 2.079442$ against $2f(0.5) = 2\ln 4.5 = 3.008155$, an overstatement of $3.008155/2.079442 - 1 = 0.4466$. At $u = 0.2$ the same split gives 0.875469 against 1.061256, an overstatement of 0.2122. Both agree with the table of Section 3.3 to four decimal places.

10.9.5 9.5 The two effective counts, and which one has external support

Section 5.16 runs a second effective count — the classical variance-reduction count $N_{\text{eff,var}} = N/[1+(N-1)\rho]$ — alongside the framework's $N_{\text{eff}} = 1+(N-1)u$. The two agreed closely on the worked board (1.371 against 1.203), and **that agreement must not be read as corroboration of either.** They share a structural motif, the $1 + (N-1)\cdot$ form, because both are affine interpolations between one and N; the resemblance is a design echo, and the paper asserts no theorem connecting them.

Neither of them is a Hill number, and the distinction matters because a reviewer will reach for the diversity literature. HHI^{-1} is the Hill number of order two and $\exp(H)$ the Hill number of order one [Hill, 1973, Jost, 2006]; the numbers-equivalent reading of the reciprocal index is Adelman [1969], and the same arithmetic was reinvented independently in ecology [Simpson, 1949] and in political science [Laakso and Taagepera, 1979]. The check takes one line, and it is run here rather than asserted: for two units holding 0.9 and 0.1, $\text{HHI}^{-1} = 1/(0.81+0.01) = 1.219512$ and $\exp(H) = \exp(0.325083) = 1.384145$, while the affine count at $u = 1$ returns exactly 2.000000. **Every Hill number of order one or higher reports at most 1.384145 on that vector, and the affine count reports two, so the affine count is not a numbers-equivalent and is not $(\sum_i p_i^q)^{1/(1-q)}$ for any $q > 0$** — with one qualification that sharpens rather than weakens the point. The Hill family is non-increasing in q, and its order-zero member, the raw richness ${}^0D$, equals 2.000000 on this vector as it does on every two-type vector, so the affine count at $u = 1$ does coincide with one member of the family. **That coincidence is the objection in its strongest form**: ${}^0D$ is the single member that ignores shares entirely, which is precisely the naive count an effective count exists to replace — it is not even a function of the share vector, being a function of (N, u). The framework's own construction carries no external support [Curzi, 2026], and it is used here because the framework's arithmetic is built on it, not because the literature endorses it. Only Rao [1982] and Leinster and Cobbold [2012] make diversity depend on dissimilarity between types at all, and neither of them is this quantity. Hurlbert [1971] and Gorelick [2011] are the standing objections to treating any such index as a property rather than a statistic, and they apply here with more force than to the Hill family, not less.

10.9.6 9.6 Measurement-theoretic status of what is scored

The evaluability partition is nominal and its aggregate is a proportion. Instrument A classifies each artefact as evaluable or not; that is a nominal judgement, the count is a frequency, and the resulting $\hat{\beta}^2$ is a proportion, which admits arithmetic. **What it does not admit is the interpretation** — that a proportion of artefacts a person cannot evaluate equals a squared vector component — and that step is postulate P4, not a measurement.

The shared-ground scores are ordinal and are treated as interval. They are nine-point judgements. Stevens [1946] is the source of the rule that permissible statistics follow the scale type, and a Pearson correlation on ordinal scores is **not licensed** without an additional assumption that the scale is interval. The paper makes that assumption and states it here rather than burying it. Michell [1997] and Michell [1999] are the strongest objection available to the framework as a whole: psychology routinely assumes quantitative structure without testing whether the attribute has it. **The objection applies to γ, to u, and to comfort, and it is conceded.**

The formal remedy exists and was not run. Conjoint measurement [Luce and Tukey, 1964] supplies the cancellation conditions under which an attribute can be shown to have additive quantitative structure at all, within the representational framework of Krantz et al. [1971]; Rasch [1960] and Andrich [1988] are the alternative route, which the corpus takes in exactly one place, where it explicitly demotes its own moment equation to ordinal scale [Curzi, 2026t]. **The conjoint conditions were not tested for any quantity in this paper.**

What survives without the test is narrower than what the paper uses. **The ordering survives**: a body whose members' shared-ground scores correlate more strongly has a lower u_{div} than one whose scores correlate less. **The zeros survive**, because $u = 0 \iff \rho^2 = 1$ and $u_{\text{div}} = 0 \iff \gamma = 0$ are statements about identities between vectors. **The budget constraint**

survives, being algebra. **The balance index does not survive independently of the interval assumption**, because it is a ratio of logarithms of quantities built from a correlation.

And the room's counts are the strongest measurements in the paper, which is worth saying. Person counts, dwell times, linear metres of sittable edge, reverberation times and interruption rates are ratio-scale physical or count quantities with established measurement procedures. The observed activity ratio is a ratio of counts. **The weak joint in Section 6 is not the measurement; it is the conversion of those measurements into the law's units**, which is a stipulation stated in Section 7.3 and audited in Section 9.1.

10.9.7 9.7 Model comparison and the scaling claims, audited against what is required

No model is fitted anywhere in this paper. Most of the model-comparison requirements are therefore inapplicable rather than unmet, and the distinction is stated.

Requirement	Status in this paper
Out-of-sample rather than in-sample performance	**Not applicable.** Nothing is fitted; the three systems are constructed and the arithmetic on them is exact. The corpus's one fitted application is reported in Curzi [2026i], and it failed its second rule.
A matched null with the same free-parameter count	**Not applicable here**; required by the specifications of Sections 4.14, 5.19 and 6.13 for any real application, where the four baselines are named.
Free-parameter count stated	**Met.** The decomposition has none. The decay model has one, λ. The storage law adds η and δ.
An information criterion	**Not applicable here; not met in the corpus**, which computes none anywhere [Curzi, 2026i].
Effect sizes with intervals	**NOT MET.** No interval accompanies any u, λ or B in this paper. Section 4.14's SPEC 4.5 requires the shared-ground item count to be reported so that a standard error *could* be computed, and no standard error is computed anywhere. **This is the clearest unmet requirement in the paper.**

The paper makes no scaling claim and estimates no exponent. The one quadratic that appears — the communication-path count $n(n-1)/2$ of Section 6.11 — is **combinatorial, defined by the model, and not estimated from data.** Saying so is not pedantry: a log-log regression of measured coordination cost on headcount would be exactly the procedure Clauset et al. [2009] shows to be unreliable, and no such regression is run here. Stumpf and Porter [2012] adds the underdetermination even where a power law is genuine. **No heavy tail is asserted, no exponent is fitted, and the word *scale-free* is not used of any distribution in this paper.** Where preferential attachment appears in Section 5.8 it is cited for the mechanism [Barabási and Albert, 1999] and immediately bounded by Broido and Clauset [2019], which tests nearly a thousand networks and finds social networks among the weakest categories for scale-free structure, and by Doyle et al. [2005], which shows that engineered robustness rather than a degree distribution explains the real case.

10.9.8 9.8 Limits and boundary behaviour

Limit	Behaviour	Reading
$\beta \to 0$	$u = \sin^2 \gamma$; pure divergence	The deep narrow well of Section 4.8.
$\beta \to 1$	$u = 1$ for every γ; γ undefined	**Pure crossing, and no shared ground on which γ could be defined.** A domain constraint, not a result.
$\gamma \to 0$	$u_{\text{div}} \to 0$	The furnished monoculture. Attained, not approached.
$\rho \to -1$	$u \to 0$	**Opposition is parallelism.**
$u_{\text{cross}} \to 1$	$u_{\text{div}} \to 0$ necessarily	The budget constraint at its edge; a strict tree sits here.
$B \to 0$	one source does all the work	Not a healthy state at high u; a concentrated one.
$\lambda \to 0$	u constant	The maintained regime. **Attainable only asymptotically**, since any communication is cascade substrate.
$\lambda \to 1$	$u \to 0$ in one period	Total collapse; the half-life goes to zero.
$N \to 1$	$f \equiv 0$, $\ln N = 0$	A one-block system has no surplus and no ceiling. Correctly degenerate.
$\delta \to 0$	Q_∞ undefined	**A pathology, not an ideal** [Curzi, 2026a]; Section 6.9 gives the room's version — a welded room holds its stored order perfectly and can never be improved.
$\eta \to 0$	$h \to -\delta Q - C < 0$	No geometry and no room rescues a body that converts nothing.
room area $\to \infty$	contact does not scale with it	**The paper's sharpest physical limit**: contact is the sum of sub-7-metre neighbourhoods, so a 900 m^2 floor with no sub-structure contains very few of them.

Two limits are constraints on the model's domain rather than results, and both are stated in the running text. $\beta \to 1$ leaves γ undefined, so the decomposition degenerates rather than misbehaving. And the linear storage law reports negative Q below a certain input, which is not a physical state: at Section 5.13's $t = 23$ the admissibility condition $C_t \leq (1-\delta)Q_t + \eta S_t$ fails, the right-hand side being 0.035053 against $C = 0.15$. **The arithmetic does not predict a negative institution; it names the period at which the institution can no longer pay its own running cost, and stops.**

10.10 10 Limitations

Every system in Section 7 is constructed. The institute, the standards board and the common room are invented; their parameters, correlations, decay rates and observation counts are stipulated, chosen so that the arithmetic can be checked and the failure modes made visible. **None of them is evidence about any real body or any real room**, and no figure in Section 7 may be cited as an empirical result. What they demonstrate is what the procedures do when carried out, not what institutions do.

The affine effective count is chosen and has no external support. Design axiom S3 fixes the shape of every result here that mentions N_{eff}. Section 9.5 runs the check: $1 + (N-1)u$ is not a Hill number of any order, is not a numbers-equivalent, and is not even a function of a share vector. A reader who imports the Hill literature's estimation, partitioning or replication results will import guarantees this quantity does not have [Hill, 1973, Jost, 2006, Curzi, 2026].

The reflexive specification is not identified from observational data. Section 3.6 states it with its citation [Manski, 1993] and names the remedy adopted — stipulating the reference group on the template of Brock and Durlauf [2001] — together with what that remedy does not buy. **Every applied u, u_{cross}, u_{div} and λ in this paper is a modelling position and not an identified empirical quantity.** Shalizi and Thomas [2011] shows the confound is generic.

The maintenance-free condition on which λ depends cannot be verified from the data. $\hat{\lambda} = 1-(u_1/u_0)^{1/k}$ is valid only if the interval contained no u-maintenance, and the specification requires that condition to be *asserted in writing with dates* precisely because nothing in the arithmetic can check it. **An interval containing unrecorded maintenance returns the net of decay and repair, which is a different quantity wearing the same name.** No procedure in this paper detects that case.

The geometric decay form is a postulate with no derivation and one close competitor. Postulate P-U is adopted for simplicity and estimability, not because anything implies it. The closest formal relative in the literature — Friedkin and Johnsen [1990]'s stubbornness term, which anchors each agent to its initial opinion and yields a non-consensus fixed point — suggests an alternative in which u decays to a positive floor rather than to zero. **This paper does not discriminate between the two forms and its worked board would look similar under either over the first ten periods.** That is a real gap and the falsification rule of H4 is what would close it.

The cascade transfer is an analogy and is used as a mechanism. Bikhchandani et al. [1992] prove their result for a sequence of agents drawn afresh, each holding one private signal about one binary question. **A standing committee is not that object**: its members recur, its questions are not binary, its signals are not exchangeable, and its members can revisit. The result is used here for shape — individually rational copying removes private information from the public record — and Section 3.7 badges it category 3. The 15.5 per cent incorrect-cascade figure at $p = 0.70$ is a property of the canonical model and **not a claim about any committee.**

The conjoint conditions were never tested. Section 9.6 concedes Michell [1997] in full. Shared-ground scores are ordinal and are correlated as though interval; u is treated as ratio-scaled; and Luce and Tukey [1964]'s cancellation conditions have not been run for γ, for u, or for comfort. **This is the single most dangerous unanswered objection to the framework and it is conceded rather than argued with.**

No interval accompanies any estimate in this paper. Section 9.7 records this as the clearest unmet requirement. The specification requires the shared-ground item count to be reported so that a standard error on r *could* be computed, and marks estimates below twelve items as provisional — and **no standard error is computed anywhere, for any quantity, in any of the three worked systems.** Every figure in Section 7 is therefore a point estimate on a constructed input, and a real application would need to propagate uncertainty in r through $1-r^2$, through u, and through a logarithm, where the sensitivity is greatest exactly at small u.

The two effective counts must not be substituted for one another. $N_{\text{eff,var}} = N/[1+(N-1)\rho]$ and $N_{\text{eff}} = 1 + (N-1)u$ come from different derivations — second moments of estimator

error against the surplus axioms — and their close agreement on the worked board is **not corroboration of either.** They are used as two gauges on one vessel and the paper asserts no theorem connecting them. The equal-pairwise-correlation assumption behind the first is also a simplification: real correlation structures are block-structured, and the correct treatment there is the eigenvalue spectrum of the correlation matrix, which this paper does not compute.

The room section rests on three postulates and one stipulated conversion, and the conversion is the weakest joint. Section 7.3 converts access minutes, reverberation time and cleaning hours into "equivalent lost agreements per week" on a stated basis. **That basis is a judgement.** It is stated openly, its correct use is comparative rather than absolute, and the resulting headroom figures must not be read to three decimal places as though measured. Section 9.1 audits it and Section 6 names it at both ends.

The observational sources for the room are borrowed, nominal, and culturally specific. Whyte [1980]'s findings were measured on New York plazas in the 1970s; Gehl [1971]'s taxonomy and distance thresholds come from Scandinavian public space; Hall [1966] states explicitly that his zone boundaries are culturally specific and that mismatched norms read as aggression or coldness. **They are starting points for a local measurement and not constants**, and the sittable-space ratio in particular is a planning regulation derived from outdoor observation in one city, reproduced here because it is one of very few space standards derived from counting people rather than from committee.

Two defects in a companion paper's version of the same check are corrected here rather than propagated. The corpus's concentration instrument runs the identical comparison [Curzi, 2026]. It reports $\exp(H) = \exp(0.325083) = 1.384115$ for the two-unit vector $(0.9, 0.1)$; **the correct value is** 1.384145, since $H = 0.3250829734$ nats and $e^{0.3250829734} = 1.3841454885$. And it states that *every* Hill number reports below 1.39 on that vector, which **is false**: the family is non-increasing in q and its order-zero member is the raw richness ${}^0D = 2.000000$. Section 9.5 gives the corrected form. **The correction sharpens the objection rather than blunting it**, since the one member the affine count matches is the one that ignores shares entirely.

One observational claim in the source has no bibliographic support here and is reported without one. The room section cites a 1984 hospital-window study for the proposition that a view affects a hard clinical outcome. **That work is not in this corpus's bibliography, no key supports it, and the claim is therefore stated without a citation rather than attached to an entry that does not support it.** The same applies to the Condorcet jury theorem, invoked in Sections 5.7 and 5.15 as a classical result: **it is named rather than cited, because the bibliography contains no entry for it.** Both are recorded here so that a reader does not mistake absence for oversight.

The proportion-of-women index is retained only because two unrelated literatures point at it, and that is not a mechanism. Whyte [1980] reported it as a quality index for plazas in one city in one decade; Woolley et al. [2010] independently found the proportion of women among the correlates of a measured collective-intelligence factor. **Two literatures pointing at the same index is worth a measurement and is not worth a belief**, and this paper asserts no mechanism for it and makes no recommendation from it.

The optional-activity proxy is a bet and the worked room is a partial counter-instance. H5 predicts that a rise in the observed ratio brings a proportionate rise in attributable conversion. The worked room's ratio rose by a factor of 1.81 and its measured conversion by a factor of 1.57. **The two moved together and not proportionately**, which is reported in Section 7.3 rather than smoothed, and the failure mode in which they diverge entirely — the room that generates pleasant lingering and stores nothing — is catalogued in Section 6.15 as the case that falsifies the postulate locally.

The monotone claim is restricted to the model. Section 8.2 lists the contrary measured literature in full [Uzzi and Spiro, 2005, Putnam, 2007, May, 1972, Sagan, 1993]. **The claim this paper defends is that surplus increases in u within the model, not that institutional value does.**

The balance-index thresholds are stipulated. $B < 0.05$ as single-sourced, 0.1 to 0.3 as lopsided, $B > 0.3$ as materially loaded — **no evidence supports those cut-points, no procedure in this paper derives them, and a different banding would classify the worked institute dif-**

ferently. The same applies to the $u \geq 0.3$ admission threshold and to the $n_E \geq 12$ item-count rule.

The specification is exposed to its own instrument. Goodhart [1975] and Campbell [1979] apply to u as they apply to any measure adopted as a target, and Section 5.17 makes the point in the framework's own favour without noticing that it cuts both ways: a body that reports u quarterly has created an incentive to manufacture disagreement, which the specification forbids and the arithmetic cannot detect — because a manufactured objection at $|\rho|$ near one scores zero, and a manufactured objection at $|\rho|$ near zero is indistinguishable from a real one on the instrument. **The only defence offered is the requirement that dissent be scored by whether it changed the artefact**, which is a conversion test rather than a geometry test.

10.11 11 Conclusion

This paper splits non-parallelism into two sources, shows that they compete for one budget, treats the result as a state variable that decays, and locates every term of the storage law inside a physical room.

The decomposition. For a pure block vector, $u = u_{\text{cross}} + u_{\text{div}}$ with $u_{\text{cross}} = \beta^2$ and $u_{\text{div}} = \alpha^2 \sin^2\gamma$, subject to $\alpha^2+\beta^2 = 1$. **The two sources therefore share one unit budget:** $u_{\text{div}} \leq 1-u_{\text{cross}}$**, so at** $u_{\text{cross}} = 0.90$ **no amount of internal argument buys more than** 0.10**.** Crossing spends the shared ground on which divergence would have to register.

The marginal identity, and why it does not settle the design question. The law prices the two identically at the margin, $\partial f/\partial u_{\text{cross}} = \partial f/\partial u_{\text{div}} = (N-1)/(1+(N-1)u)$, so there is no premium on cross-disciplinarity in this framework. Their productivity in the design variables differs: $\partial u/\partial(\beta^2) = \cos^2\gamma$ against $\partial u/\partial(\sin^2\gamma) = \alpha^2$. **Each lever is most productive exactly where the other has not been pulled, so the correct first move is always the scarcer source.**

The upper bound is an impossibility limit and never a plan. Strict concavity makes $f(u_{\text{cross}}+u_{\text{div}}) < f(u_{\text{cross}}) + f(u_{\text{div}})$ whenever both are positive: at $N = 8$, a balanced split and $u = 1$, the bound overstates the true value by 44.7 per cent.

The institute. Six declared disciplines audit to $u_{\text{cross}} = 0.72$, $u_{\text{div}} = 0.03$, $u = 0.75$ at $N = 6$, a balance index of 0.0206, a ceiling utilisation of 0.8696, and **zero within-block dissents across forty decisions.** Adding a seventh discipline — the institute's own standing proposal — drives the period change to -0.002863 and the fixed point from 6.265 to 5.964, because divergence and not domain count was the deficiency.

Opposition is parallelism. A normalised inner product of -1 gives $u = 0$. Priced as a seat at $\eta = 0.35$, $\delta = 0.10$ and a seat cost of 0.05, **a perfect opponent and a perfect echo both contribute** -0.5000 **to the fixed point, and both are worse than leaving the seat empty.**

Non-parallelism is a state variable that leaks. Under the badged postulate $u_t = u_0(1-\lambda)^t$, with half-life $\ln 2/(-\ln(1-\lambda))$, a standards board measured at $\lambda = 0.18$ per quarter — half-life 3.49 — fell from $u = 0.55$ to 0.0057 in twenty-three quarters with no personnel change. **Its stored order peaked two quarters after its geometry began to collapse**, so every report available during those two quarters said the programme was working. **A one-time restoration on a system with a three-per-cent residual leak peaks at** 3.202 **and returns to within** 0.11 **of its starting point twenty-eight quarters later. The durable quantity is not** u**; it is** $1/\lambda$**.**

The room. A 120 m^2 common room audits to $N = 2$ categories and headroom -1.3957. Three moves chosen by the arithmetic — raising categories from two to five, cutting the acoustic cost term from 0.70 to 0.15, and retiring four of nine dormant agreements — close 66 per cent of the gap. At week eight the observed optional-activity ratio has risen from 0.185 to 0.335 while necessary activity is flat at 211 against 216, which is the internal-validity check the taxonomy supplies for free.

And the counter-result, which the paper states rather than repairs. A room cannot manufacture non-parallelism, cannot raise conversion above what surrounding practice permits, and cannot make a negative-headroom system positive. **The worked room ends at** -0.4808

after all three interventions and is reported as still negative. The finding is that the room is a contact generator and not a storage device, and that the agreements it produces must be held by an artefact. **Physical intervention is necessary and frequently insufficient.**

10.11.1 Declarations

Data and code availability

No external dataset was collected for this paper and none exists. The three systems worked in Section 7 — a cross-disciplinary research institute, a standards review board, and a departmental common room — are constructed composites with stipulated parameters, produced so that every reported number can be re-derived by the reader from the parameters and equations printed in the text. **They must not be cited as evidence about any real organisation or any real room.** Every table in Sections 3 to 6 is an exact evaluation of a closed-form expression at the arguments shown. The figures are generated by the scripts named in each caption, which import the shared style module `lib/figstyle.py` and contain no data beyond the numbers printed in this paper.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi developed the framework, performed the analysis, produced the figures, and wrote the paper.

Ethics

No human-subjects research was conducted for this paper and no personal data were collected or processed. The bodies and rooms discussed are constructed. **The observational instruments specified in Sections 6.2 and 6.14 — time-lapse recording, fixed-time counts, contact-pair logging — would, if run on real occupants, constitute observation of identifiable people**, and the specification accordingly requires aggregate counts, posted notice, a defined retention period, and an opt-out that costs nothing; SPEC 6.19 states that no requirement may be implemented in a way that removes an occupant's ability to decline participation, to leave, or to be unobserved. Those are constraints on prospective work, not reports of any study performed. The divergence protocols of Sections 4.14 and 5.19 likewise create opportunity and protection and never obligation: **no participant may be required to disagree, and a decline is recorded as a decline rather than as a failure.**

Preregistration

This paper reports no preregistered test. The hypotheses of Section 2 are stated with their falsifiers so that a later test can be preregistered against them, and each is marked **not tested here** except H5, whose worked instance is reported as a partial counter-instance in Section 7.3. The frozen falsification blocks reproduced in Sections 7.1, 7.2 and 7.3 are specimens of the preregistration form and are not records of tests performed. The corpus's one preregistered confirmatory application of the storage law is reported elsewhere [Curzi, 2026i], where

it **failed its second falsification rule**, and the shuffled-input control on which such tests depend has been measured within the corpus at a false-rejection rate of 0.108 against a nominal 0.05 [Curzi, 2026u].

Use of the framework's instructional register

The source manuscript from which this paper is drawn is written as an instructional manual with a ceremonial register. This paper reports the analytic content; the register is not reproduced, and the source's named constructs are renamed where they identify an object this paper analyses.

Chapter 11

Passive Geometry and the Reverberant Tail: Crenellation, Parabolic Focus, and the Acoustic Exclusion of Quiet Participants

Abstract. This paper prices two classes of passive built-space intervention against the surplus ceiling ln N. On the geometric side it proves five elementary results — the parabola's reflective and focus-directrix properties, path-length equality, the similarity dimension D = ln m / ln r, and the crenellation multiplier 1 + d/w — and then refuses to let them certify anything: a fractal argument may generate a candidate geometry but not validate one, and every self-similarity requirement carries a modelling obligation. Applied to a 120 m2 common room at d = 1.00 m, w = 1.50 m, crenellation lengthens the boundary from 43.82 m to 73.03 m for 18% of the open floor, turning an unmeetable seating-edge requirement into a 6.8-fold surplus and raising headroom from -1.396 to -0.804. On the acoustic side the turn-taking mechanism is derived from the decay tail: the gap for a talker's tail to fall 15 dB is RT60/4, so an untreated 1.95 s room demands 488 ms against a natural turn gap while a treated 0.39 s room demands 98 ms. Three of eight participants silenced costs ln 8 - ln 5 = 0.4700 nats, 22.6% of the room's per-interaction ceiling. Treatment moved measured turn-taking equality from 0.604 to 0.812 against a flat control, but the preregistered group-performance outcome did not improve, and the causal postulate was reported as unconfirmed.

Keywords: reverberation time; turn-taking equality; crenellated boundary; parabolic reflector; Sabine and Eyring; surplus ceiling; deliberation room

Series: B. Engineering the Surplus **Paper:** CCP-011

11.1 1 Introduction

A room is a parameter setting. The framework this paper belongs to prices an accumulating system with the storage law `Q_{t+1} = (1−delta) Q_t + eta S_t − C_t` and reads its sustainability off one gauge, the headroom `h_t = eta ln N − delta Q_t − C_t`, whose sign decides whether stored order can grow at all [Curzi, 2026b]. Three of the four terms in that gauge are ordinarily treated as properties of people: how many distinct categories are in the room, how efficiently

their interaction converts into durable structure, what the coordination overhead is. **This paper asks what a building sets, and prices two classes of passive intervention — geometric and acoustic — against the same ceiling `ln N`.**

The question matters because the two classes have a cost profile nothing else in the framework has. A powered device that performs a spatial function enters the cost term as a permanent line: energy, maintenance, replacement, a technician, a failure mode. A shape enters once, at construction, and thereafter contributes only cleaning. Its retention loss is set by the service life of concrete, plaster and hardwood rather than by anyone's attention. Infrastructure of that kind becomes visible only on breakdown [Star and Ruhleder, 1996, Star and Griesemer, 1989], and its breakdown mode is slow. **A correctly shaped alcove is a function with a cost of nearly zero and a decay rate measured in decades.**

Two mechanisms are developed and each is taken to a worked case.

The geometric mechanism is boundary multiplication. Interior area is expensive and fixed by the lease; boundary length is a design variable. A boundary replaced by a square-wave of tooth depth `d` and half-period `w` has its length multiplied by exactly `1 + d/w`, and if the teeth alternate about the original line the enclosed area is unchanged. Applied inward to a 120 m^2 common room at `d = 1.00 m` and `w = 1.50 m`, the perimeter goes from 43.82 m to 73.03 m at a cost of 21.9 m^2 of open floor — 18 per cent — against a sittable-edge requirement that simultaneously falls to 10.73 m. **The requirement was never hard to meet; it was hard to meet in a room with no boundary.** Since occupancy concentrates at edges [Gehl, 1971] and a space's use tracks its sittable perimeter [Whyte, 1980], this is the cheapest available route to a higher category count, and in the worked case it takes `N` from 2 to 5 and headroom from `−1.396` to `−0.804` with no wall built and no square metre surrendered from the lease.

The acoustic mechanism is the decay tail. Sound decays at `60/RT60` decibels per second, so the gap a talker must leave for the previous talker's tail to fall 15 dB is `t_15 = RT60/4`. In an untreated room at `RT60 = 1.95 s` that gap is 488 ms; in the same room treated to `RT60 = 0.39 s` it is 98 ms. Turn-transition gaps in ordinary conversation are short — of the order of a fifth of a second — so the untreated room makes a natural turn physically insufficient and the treated room does not. **A quiet talker entering a gap that does not clear the tail is inaudible, is talked over, and is recorded by everyone present as not having said anything.** Nobody in that room behaved badly. The room did it, and it will keep doing it after every facilitation course.

That mechanism has a price in the framework's own units. Equality of conversational turn-taking is one of three correlates of the collective-intelligence factor reported by Woolley et al. [2010], and it is the correlate a room can move. If a room reliably silences three of eight participants, the interaction's effective category count is five rather than eight and the per-interaction ceiling falls from `ln 8 = 2.0794` to `ln 5 = 1.6094`: a loss of **0.4700 nats, 22.6 per cent of the room's entire ceiling**, permanently, in every meeting, for as long as the room stands.

The paper's central negative result is that the causal chain this argument needs was tested and did not confirm. A preregistered before-and-after in a worked eight-person committee, with a matched untreated control, moved measured turn-taking equality from 0.604 to 0.812 while the control moved 0.012; the preregistered group-performance outcome improved slightly and did not reach its declared threshold. **The mediator moved decisively and the outcome did not.** Under the preregistration that is the outcome that says the acoustic-to-collective-intelligence chain is not supported for that group, and it is reported here in the same language a confirmation would have received. Two commissioning bands also came in 0.01 s and 0.02 s below the lower target bound and are reported as misses rather than absorbed into the tolerance.

Four things are new here relative to the surrounding literature rather than relative to the framework. First, the turn-taking mechanism is derived from the decay tail as arithmetic rather than asserted as an association, which makes it attackable at a specific joint. Second, acoustic exclusion is computed from physics rather than from self-report, which matters because a contribution never attempted leaves no trace and cannot be surveyed. Third, boundary length and category count are placed on one commensurable scale with the storage law, so that joinery and hiring can be compared in the same units. Fourth, **a three-tier certifica-**

tion standard is imposed on the paper's own geometric claims, under which no self-similarity claim may be certified from the concept; applied to the worked room it leaves three of twelve claims uncertified, and they are the three that sound most impressive.

Sections 2 and 3 state the hypotheses and the framework, including the identification problem that limits every applied reading of the mechanism. Section 4 derives the parabola and its hard frequency floor; Section 5 develops self-similarity, the crenellation multiplier and the nested enclosure ladder; Section 6 develops the diffuse-field model, its five assumptions and the Sabine/Eyring switching rule; Section 7 develops the exclusion mechanism and the tail arithmetic. Section 8 collects the design specifications. Section 9 works two systems end to end. Sections 10 to 13 discuss what follows, verify the chain, state the limitations, and conclude. Appendix A gives the eight-measurement protocol and Appendix B the failure-mode catalogue.

11.2 2 Hypotheses

The results below are of two kinds and the distinction is load-bearing. Some are theorems about a geometry or an identity of the decay law and cannot be falsified by observation; they are labelled as derivations where they appear. Others claim that these equations describe a real room and a real meeting, and those are stated here with what would defeat them.

H1 — The decay tail sets a minimum inter-turn gap. In a room of reverberation time `RT60`, the gap a following talker must leave for the previous talker's tail to fall 15 dB is `t_15 = RT60/4`, and a talker whose vocal level is materially below the previous talker's cannot take a turn in a gap shorter than that. **Predicted:** measured turn-transition gaps lengthen with `RT60` across rooms, and the share of attempted turns that are abandoned or overridden rises with `t_15` relative to the population's natural gap. **Falsified by:** measured turn gaps and override rates that are independent of `RT60` across a set of rooms spanning at least a factor of three in reverberation time.

H2 — A measurable and often dominant share of turn-taking inequality in a physical meeting is produced by the room's acoustics rather than by anyone's behaviour. **Predicted:** acoustic treatment alone, with the roster, the chair, the agenda and the meeting series held fixed, raises the turn-share evenness statistic `N_eff_turn/n`. In the worked case, from 0.604 to at least 0.75. **Falsified by:** treatment that meets its acoustic targets and leaves `N_eff_turn/n` unmoved, which would locate the inequality in status, chairing or agenda instead. **The share is not claimed as a number; the mechanism is claimed and its magnitude is computed.**

H3 — Acoustic treatment is therefore a collective-intelligence intervention. This chains H2 to Woolley et al. [2010]'s reported correlation, and the chain requires that turn-taking equality *causes* the collective-intelligence factor rather than merely tracking it. **Predicted:** a preregistered before-and-after with a matched untreated control shows the mediator and the group-performance outcome moving together. **Falsified by:** the mediator moving and the outcome not moving. **This paper ran that test and obtained exactly that pattern.** H3 is therefore reported in Section 9.2 as **not supported by this case**, and the treatment is retained on accessibility, fatigue and ceiling grounds instead. The direction of the correlation is not established in the source literature, so H3 could have failed for a reason that leaves H2 entirely intact.

H4 — Boundary multiplication converts floor area into usable category count. Crenellating a room's solid boundary inward raises the number of distinct posture/purpose/condition categories a fixed floor can host, at a floor-area cost equal to the recess depth times the recessed run. **Predicted:** in the worked room, `d =`

`1.00 m` over 21.91 m of run costs 21.91 m^2 — 18 per cent of the floor — multiplies the perimeter by 1.6667, and takes the audited category count from 2 to 5, so `ln N` rises by a factor of 2.32. **Falsified by:** an audited count that does not rise after construction, or fine-scale features that a time-lapse count shows to be unoccupied at all times, which is the recorded failure mode **FG4** of Appendix B.

H5 — The four-part seating pattern is scale-invariant. Positions that are back-protected, give a view out over a larger territory, have a visible low-cost exit and offer a boundary at reachable scale are occupied preferentially, and the same four conditions predict occupancy at desk, alcove, room and wing scale with only the dimensions changing. **Predicted:** occupancy counts at each built rung favour positions satisfying all four conditions over positions failing one or more. **Falsified by:** counts showing no preference at any rung, in which case **the seating requirements of Section 8 should be discarded for that population.** Do not defend the pattern against the count.

Three claims in this paper are **not** hypotheses and are labelled accordingly wherever they appear. That the storage law applies at every scale of enclosure — desk, alcove, room, wing, campus — with the same form and different parameters is a **modelling choice** and not a claim about physics; there is no theorem that a campus obeys the same recursion as a desk. That a self-similar boundary buys multi-scale social response the way a self-similar conductor buys multi-band electrical response is an **analogy**, structural and suggestive and not a derivation. That acoustically caused silence is indistinguishable, in the room's own records, from voluntary silence is a **definition of the measurement problem** rather than an empirical finding: it is why Section 7.6 computes exclusion from physics and not from asking.

11.3 3 Framework

11.3.1 3.1 What is presupposed

The storage law and its gauge. The equation this paper acts on is

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,     S_t = F(U_t) = ln(1 + (N−1) U_t)
```

with `Q_t` the stored order, `S_t` the surplus generated by interaction, `delta` the retention loss per period, `eta` the conversion of interaction into durable structure, and `C_t` the cost flow. The results consumed here and not re-derived are the headroom gauge `h_t = eta ln N − delta Q_t − C_t` with its three regimes and maximum one-step increment `r_max = h_t`, the two distinct constraints of admissibility `C_t <= (1−delta)Q_t + eta S_t` and sustainability `C_t <= eta ln N − delta Q_t`, and the steady state `Q_inf = (eta F(u*) − c*)/delta`, which is negative when `eta F(u*) < c*` [Curzi, 2026b]. The level–rate–delay quartet is the standard vocabulary of system dynamics and long predates this framework [Forrester, 1961]; headroom is this framework's name for what Cyert and March [1963] called organisational slack.

The surplus functional. For two unit state vectors drawn from an ensemble of `N` categories, `u(x,y) = 1 − (x . y)^2` in `[0,1]` and `f(u) = ln(1 + (N−1)u)`, characterised by four axioms — angle-only dependence, zero at parallel, an affine effective count, normalisation at orthogonality — as established in Curzi [2026]. **The affine effective count `N_eff = 1 + (N−1)u` is a design axiom: chosen, not derived**, and it is flagged as such wherever this paper leans on it. The range `0 <= F(u) <= ln N` is what makes `eta ln N` a ceiling rather than a typical value, and the forced-ceiling form `0 <= M <= ln N` with its worked twelve-category value of 2.484907 nats is developed in Curzi [2026l].

Two effective counts appear in this paper and they are different objects. The first is `N_eff = 1 + (N−1)u`, above. **It is not a Hill number.** The effective-count idiom is standard in ecology and political science, where `1/HHI` and `exp(H)` are the canonical members of the family [Hill, 1973,

Jost, 2006, Laakso and Taagepera, 1979]; `1 + (N−1)u` is a linear interpolation between 1 and `N` and belongs to no order of that family. The concentration instrument that *is* a Hill number of order two — `1/HHI`, with the worked jurisdictional case returning a naive HHI of 0.293430 and 3.41 effective jurisdictions against a bloc-HHI of 1.000000 and an effective count of 1.00 — is a separate quantity [Curzi, 2026]. The second effective count in this paper is the turn-share statistic `N_eff_turn = exp(−sum_i p_i ln p_i)` of Section 7.4. **That one IS a Hill number of order one** — the exponential of Shannon entropy [Shannon, 1948, Jost, 2006, Pielou, 1966] — and it is a descriptive index of the turn distribution and nothing else. **The three must not be substituted for one another, and `ln(N_eff_turn)` must never be reported as a surplus.** Section 12 repeats this, because a reviewer reaching for the Hill-number literature to support `1 + (N−1)u` will find that it does not apply.

The category count `N` in this paper is an audited count of posture/purpose/condition categories available in a room, established by the survey procedure of the room-as-circuit-element analysis, whose worked room returns `N = 2`, a sittable-edge ratio of 0.050 against a requirement of 0.109, `delta = 0.0808` per week, `eta_hat = 0.0455` and `h = −1.396`, and whose prescription — raise `N` from two to five, cut the acoustic cost term, retire dormant agreements — closes 66 per cent of that gap [Curzi, 2026d]. Section 9.1 designs the geometry that delivers the `N` increase in that prescription. **Whether an audited category is a category is a judgement with a declared strike rule, not a measurement**, and Section 12 says so.

11.3.2 3.2 Identification: what no amount of measurement recovers

This limitation is structural rather than statistical and belongs here rather than in a caveat at the end. The mechanism this paper asserts runs through a group whose members both influence and constitute the aggregate: the room shapes who speaks, who speaks shapes what the meeting is, and the meeting is made of its members. That is exactly a linear-in-means reflexive specification, and Manski [1993] proves that in such a specification **the endogenous social effect is not identified from observational data**: it cannot be separated from contextual effects or from correlated unobservables, however much data is collected. Shalizi and Thomas [2011] reinforce the point for network data, where influence and selection are generically confounded.

The consequence for this paper is exact and it is narrower than it first looks. Sections 4 to 7 contain two kinds of statement. The geometric theorems, the decay-rate identity, Sabine's equation and the direct-to-reverberant arithmetic are physics and are untouched: they concern air and surfaces, not social influence, and no identification problem arises. **The interpretation of the resulting turn-share change as an endogenous social effect is what is not identified.** When Section 9.2 reports that `N_eff_turn/n` rose from 0.604 to 0.812 while a matched control moved 0.012, that comparison licenses the claim that *the intervention* moved the mediator; it does not license decomposing the mediator's effect on group behaviour into an endogenous social component and a contextual one.

A partial remedy is adopted and its price is paid. The matched untreated control room removes the most immediate confound — that measuring a group's turn-taking is itself an intervention, the hazard McCambridge et al. [2014] documents systematically and Rosenthal and Jacobson [1968] demonstrates for expectancy — and it is the reason Section 8's requirement schedule makes the control mandatory rather than advisory. **It does not solve the reflection problem.** The route that would is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001]; the fixed roster and the declared participation rule of the measurement protocol point toward it without yet taking it. **Until that is done, the claim that acoustic treatment raises collective performance through the turn-taking channel is a modelling position rather than an identified empirical result**, and every applied statement in this paper inherits that status. Durlauf [2005] states the general form of the demand: name the observation that would distinguish this account from a neighbouring one that fits equally well.

11.3.3 3.3 Analogy discipline

Three borrowings from physics appear in this paper and each is labelled at first use, following the three-way distinction the corpus imposes on itself.

Sabine's equation and the two-field model are derivations, not analogies. They are statements about sound in air, used here for sound in air. Sabine's `RT60 = 0.161 V/A` is not a curve fit: it follows from a statistical-energy argument in which the mean free path between reflections in a rectangular enclosure is `4V/S` [Sabine, 1922, Rossing et al., 2002]. **But the argument makes five assumptions, Section 6.1 lists them, and Section 6.3 exhibits a case where Sabine over-predicts the built result by 77.8 per cent.** The premises are checkable and this paper checks them; that is what distinguishes a derivation from a borrowing.

The Lombard feedback loop is a shared mathematical form. In a room with several conversations, each conversation's output is background for the others, and a first-order loop with gain `s` below unity converges to `1/(1−s)` times its seed disturbance. That is the same fixed-point arithmetic as the storage law's convergence, and it is the same contraction argument: a map with modulus `q < 1` on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least `q` each pass [Banach, 1922, Curzi, 2026]. **Nothing about charge, capacitance or thermodynamics is claimed to hold of a noisy room**; the claim is that the same difference equation governs both, which is legitimate and is category two of the three [Åström and Murray, 2008].

Self-similarity in social occupancy is an analogy and is badged as one. Self-similar antennas work because Maxwell's equations are scale-invariant in the relevant sense. **There is no scale-invariance principle for human occupancy corresponding to that**, because people have one size: proxemic bands are absolute distances, not ratios [Hall, 1966], and so are the sensory thresholds that govern what can be seen and heard across a space [Gehl, 1971]. The analogy motivates a candidate geometry and licenses nothing. Section 5.7 states the certification rule this produces, and Section 12 records that geometric over-claiming has a long and documented history in design writing, of which the golden-section literature is the best-audited instance [Markowsky, 1992, Thompson, 1917].

One further borrowing is refused outright. `ln N` here is a bound on an ensemble of `N` equiprobable categories in the information-theoretic sense [Shannon, 1948], and the requisite-variety tradition supplies the reason a variety term appears as a ceiling at all: only variety absorbs variety [Ashby, 1956, 1958]. **Surplus is not identified with Shannon entropy or mutual information, and no thermodynamic entropy law is invoked**; the two are different objects with different units, and Mirowski [1989] is the standing warning about what happens to a social science that copies an energetics for which it has no conservation law.

11.3.4 3.4 Standing honesty clauses

Repeated wherever they bite, and stated once here.

- **As deployed, the surplus functional is a lower bound** on the information an ensemble carries, not the information itself [Curzi, 2026]. Every number in this paper is a floor.
- **`ln N` is a ceiling per interaction, not a typical value.** When Section 7.3 computes the loss from silencing three of eight voices, it computes a loss of *ceiling*, and the realised loss may be larger, because the silenced participants may have been the ones carrying the non-parallel view [Curzi, 2026d].
- **Every absorption coefficient in this paper is NOMINAL** — a typical published value for a generic material description. They are adequate for design arithmetic and are **not adequate for a construction specification**, where the manufacturer's tested data for the specific product at the specific mounting condition is required. A "50 mm mineral wool panel" spans a range of real performance wide enough to change a design decision.
- **Applying the storage law to a room is a modelling decision**, badged as such every time it is made, and the same applies with more force to applying it at five nested scales at once.

- **The transfer of a classroom acoustic standard to a deliberation room is a transfer.** It is defensible — a committee member straining to hear is in the same acoustic situation as a child at the back of a class — and it is labelled as one wherever it is used, because the standard was arrived at through the speech-intelligibility literature for a different population and task.
- **No intervention in this paper may remove a participant's ability to decline**, and the arithmetic gives no warning where that boundary is crossed. Recording a meeting for turn-share measurement requires consent, and acoustic treatment must not be presented as a remedy for a hierarchy that is not acoustic in origin. The corpus's own three-drawer honesty classification — four items of established mathematics, seven of established psychology, and exactly three declared postulates, with every protocol stated to run on the first two and steer by the third — is the standard this section is written to [Curzi, 2026].

11.4 4 Geometry that performs a function without power

11.4.1 4.1 The cost argument

Return to the cost term. `C_t` in the storage law is dissipation: what the system spends to keep running. A powered device that performs a spatial function — an amplifier, a fan, a display, an active noise system — enters the budget as a permanent line. A shape enters once, at construction, and then contributes only through cleaning and eventual repair.

```
Powered function:   C_t = c_energy + c_maint + c_replace + c_failure   , every period
Geometric function: C_t = c_clean                                      , every period
                    plus a one-time capital charge at t = 0
```

This is the structural argument for the whole class. A shape is a function with a cost of nearly zero and a retention loss set by material durability rather than by anyone's attention. It is the limiting case of infrastructure that recedes from notice until it fails [Star and Ruhleder, 1996], and its failure mode is slow — wear, obstruction, misuse — rather than sudden, which makes it a low-variance element in a cost profile. **Stone, concrete, hardwood and plaster have retention losses measured in decades, and almost nothing else in an organisation does.** The same claim in its acoustic form — that boundary geometry alone selects modes at no ongoing cost, exactly for room acoustics, where the mode frequencies follow `f = (c/2) sqrt((n_x/L_x)^2 + (n_y/L_y)^2 + (n_z/L_z)^2)` from the same Helmholtz problem, and as a design heuristic only for occupancy — is stated and badged as a postulate in Curzi [2026aa], and this paper adopts that badging without weakening it.

11.4.2 4.2 The parabola, derived

Take the parabola in its focal form:

```
y = x^2 / (4f)            focus  F = (0, f)
                          directrix  y = −f
                          dy/dx = x / (2f)
```

Theorem 1 (focus-directrix). For any point `P = (a, a^2/(4f))` on that parabola, `|PF| = y_P + f`, which is exactly the perpendicular distance from `P` to the directrix.

Proof.

```
|PF|^2 = a^2 + ( a^2/(4f) − f )^2
       = a^2 + a^4/(16 f^2) − 2f ∗ a^2/(4f) + f^2
       = a^2/2 + a^4/(16 f^2) + f^2
```

```
( a^2/(4f) + f )^2 = a^4/(16 f^2) + 2f * a^2/(4f) + f^2
                   = a^4/(16 f^2) + a^2/2 + f^2
```

The two expressions are identical, so `|PF| = a^2/(4f) + f = y_P + f`, the vertical distance from `P` down to the directrix `y = −f`. **The distance to the focus equals the distance to the directrix at every point on the curve: the parabola's definition, recovered from its equation.**

Theorem 2 (reflective property). Every ray parallel to the axis of symmetry reflects, at the surface, through `F`.

A ray descending parallel to the axis has direction `d = (0, −1)`. The surface normal at `P` comes from the gradient of `y − x^2/(4f)`, giving `n` proportional to `(−a/(2f), 1)`, normalised. The reflected direction is `r = d − 2 (d . n_hat) n_hat`. Carried out numerically for `f = 1.0 m` and checked at the focal height `y = f`:

`a` (m)	`y_P` (m)	slope `a/2f`	tangent angle (deg)	x-crossing at `y = f`
0.25	0.0156	0.1250	7.13	0.000000
0.50	0.0625	0.2500	14.04	0.000000
0.75	0.1406	0.3750	20.56	0.000000
1.00	0.2500	0.5000	26.57	0.000000
1.50	0.5625	0.7500	36.87	0.000000
2.00	1.0000	1.0000	45.00	reflected ray exactly horizontal; `y_P = f`
3.00	2.2500	1.5000	56.31	0.000000

Every reflected ray crosses `y = f` at `x = 0`. The `a = 2f` row is the degenerate case: the reflected ray leaves horizontally and travels along `y = f` to the focus. That is the rim-angle-90-degree point, and it is why `f/D = 0.25` is the geometry at which the rim subtends a full half-space at the focus.

Corollary 3 (path-length equality) — the load-bearing one. Every axis-parallel ray descending from a common height `h` and reflecting to the focus travels the same total distance `h + f`, independent of where on the surface it struck:

```
total = (h − y_P) + |PF| = (h − y_P) + (y_P + f) = h + f
```

The `y_P` cancels. Numerically, with `f = 1.0` and `h = 5.0`:

`a` (m)	`y_P`	`|PF|`	`(h − y_P) + |PF|`
0.0	0.0000	1.0000	6.0000
0.5	0.0625	1.0625	6.0000
1.0	0.2500	1.2500	6.0000
1.5	0.5625	1.5625	6.0000
2.0	1.0000	2.0000	6.0000
3.0	2.2500	3.2500	6.0000

Identical at every aperture radius, because it is identically `h + f`. Corollary 3 is why a parabola has *gain* and a merely curved surface does not. Convergence without path-length equality produces arrival out of phase, which is partial cancellation; convergence with

path-length equality produces coherent addition. **The parabola is not a funnel but a phase-correcting device, and the phase correction is performed by the shape, at no cost, indefinitely.** There is no computation, no calibration and no control loop: the function is stored in the boundary condition.

11.4.3 4.3 Aperture, depth, and f-ratio

For a dish of aperture diameter `D` and focal length `f`, the rim sits at `x = D/2`, so `depth = (D/2)^2/(4f) = D^2/(16f)` and `f/D` is the f-ratio, smaller meaning deeper.

`D` (m)	`f` (m)	`f/D`	depth (m)	rim angle at focus (deg)
0.50	0.15	0.300	0.104	79.6
1.00	0.30	0.300	0.208	79.6
2.00	0.50	0.250	0.500	90.0
6.10	1.50	0.246	1.550	90.9
9.10	2.30	0.253	2.250	89.4
61.00	15.00	0.246	15.504	90.9

The recurrence of `f/D` near 0.25 is not coincidence: at `f/D = 0.25` the focus sits exactly in the plane of the rim, which is the shallowest geometry that still collects a full hemisphere and is therefore the most material-efficient. **A design that converges on `f/D` near 0.25 has usually converged correctly.** Deeper dishes shield the focus from off-axis noise at the cost of material and of a focus buried inside the shell where a head does not comfortably fit; shallower dishes put the focus in the open where it is easy to use and easy to contaminate.

11.4.4 4.4 Directivity, gain, and the hard frequency floor

For a circular aperture of diameter `D` at wavelength `lambda`, the half-power beamwidth is approximately `theta ~ 70 lambda / D` degrees, and the theoretical on-axis gain with aperture efficiency `eta_ap` is `G = (pi D / lambda)^2 eta_ap`. Take `eta_ap = 0.5`, a conventional and generous value for a real reflector, and `c = 343 m/s` at 20 C. For a 1.2 m aperture — the largest that fits comfortably behind a seated head:

Frequency	`lambda` (m)	`D/lambda`	Beamwidth	Aperture gain (upper bound)
250 Hz	1.372	0.87	80.0 deg	5.8 dB
500 Hz	0.686	1.75	40.0 deg	11.8 dB
1000 Hz	0.343	3.50	20.0 deg	17.8 dB
2000 Hz	0.172	7.00	10.0 deg	23.8 dB
4000 Hz	0.086	13.99	5.0 deg	29.9 dB

The frequency floor. A reflector does essentially nothing when it is small compared with the wavelength; sound diffracts around it. The working rule is that useful directivity begins at about `D = 2 lambda`, which gives the hard floor `f_floor = 2c/D`:

Aperture `D`	`D = 1 lambda` at	`D = 2 lambda` at (useful floor)	`D = 4 lambda` at
0.30 m	1143 Hz	2287 Hz	4573 Hz
0.50 m	686 Hz	1372 Hz	2744 Hz

Aperture `D`	`D = 1 lambda` at	`D = 2 lambda` at (useful floor)	`D = 4 lambda` at
0.60 m	572 Hz	1143 Hz	2287 Hz
0.90 m	381 Hz	762 Hz	1524 Hz
1.20 m	286 Hz	572 Hz	1143 Hz
2.00 m	172 Hz	343 Hz	686 Hz
3.00 m	114 Hz	229 Hz	457 Hz
6.10 m	56 Hz	113 Hz	225 Hz
9.10 m	38 Hz	75 Hz	**151 Hz**
61.00 m	5.6 Hz	11 Hz	22 Hz

A correction to the source table is recorded here rather than made silently. The manuscript this paper converts gives 225 Hz in the `D = 4 lambda` column for both the 6.10 m and the 9.10 m apertures. The 6.10 m entry is right — `343/(6.10/4) = 224.9` Hz — and the 9.10 m entry is a duplicate of it: `343/(9.10/4) = 150.8` Hz. **The table above carries 151 Hz.** Nothing in the paper's argument depends on the cell, and it is corrected in the open because a reviewer who checks one cell checks the rest.

Now place speech on that scale. An adult fundamental runs 85–255 Hz, the first two vowel formants 250–2500 Hz, and consonants and fricatives 2000–8000 Hz — which is where intelligibility lives. **A 0.6 m shell behind a chair has a floor of about 1143 Hz. It does nothing whatever to the voice's fundamental, little to the first formant, and real work only on the consonant band.** That is not a failure: the consonant band is precisely where intelligibility is won or lost, and Section 6 shows that consonants are also what a reverberant room destroys first. But it means the honest description of a small acoustic shell is **a consonant-band directivity device, not something that makes a talker louder.**

Three further conditions bind every gain figure above.

1. **The gain formula assumes a plane wave arriving on-axis from the far field.** A talker 1.5 m from a 1.2 m dish is nowhere near that aperture's far field; the wavefront is curved across the aperture and the coherent addition is partial. Near-field performance is substantially below the tabulated bound.
2. **`eta_ap = 0.5` is an assumption.** Surface accuracy, edge diffraction, blockage by the listener's own head at the focus, and absorption in the shell all reduce it. Head blockage is severe in listening shells, because the thing being fed is sitting in the aperture.
3. **These are single-frequency figures.** Speech is broadband and the perceived effect is an average weighted by the speech spectrum and by the ear's frequency response.

Every computed aperture gain is therefore an upper bound that will not be reached, and the built article must be measured. Section 9.1 measures one and finds 4.8 dB against a bound of 17.8 dB.

11.4.5 4.5 The Denge precedent

Between roughly 1928 and the early 1930s, at Denge near Dungeness on the Kent coast, the British Air Ministry built a set of concrete acoustic mirrors as an early-warning system for aircraft, under William Sansome Tucker. Three principal structures survive: a 6.1 m dish, a 9.1 m dish, and a 61 m curved wall. The operating principle is exactly Theorem 2: engine noise arriving as a near-plane wave is concentrated at the focus, where a microphone — in earlier installations, a listening trumpet and a human ear — picks it up, and bearing is determined by moving the collector across the focal region. **There is no power in the acoustic path. The concrete does the work.**

What the geometry buys, in this paper's terms, at `eta_ap = 0.5`:

D	Beamwidth at 250 Hz	Beamwidth at 500 Hz	Aperture gain at 250 Hz (bound)
6.10 m	15.7 deg	7.9 deg	19.9 dB
9.10 m	10.6 deg	5.3 deg	23.4 dB
61.00 m	1.6 deg	0.8 deg	**39.9 dB**

A second correction to the source, recorded rather than absorbed. The manuscript gives 46.9 dB for the 61 m structure at 250 Hz. Recomputing from the formula the manuscript itself states, `G = (pi D/lambda)^2 eta_ap` with `lambda = 1.372 m` and `eta_ap = 0.5`, gives `(pi x 61/1.372)^2 x 0.5 = 9755`, or **39.9 dB**. The stated 46.9 dB is not reproducible from the stated formula at any of its stated parameters; the nearest reproducible neighbour is 45.9 dB, which is the same aperture at 500 Hz. **The table above carries 39.9 dB.** The 61 m structure is in any case a strip mirror rather than a dish, so it focuses in one plane and the circular-aperture figures are indicative for it; but a figure that cannot be recovered from its own formula must be replaced rather than repeated.

The order of magnitude is nonetheless the explanation for building something that large: **bearing resolution on a low-frequency source is set by `lambda/D` and by nothing else**, and at 250 Hz a two-degree beam is not available from anything smaller than a building. Reported detection ranges were of the order of tens of kilometres in favourable conditions, giving warning of some minutes; those figures varied enormously with wind, sea state, aircraft type and operator skill and the surviving accounts are not consistent, so **the safe claim is the one the geometry guarantees: real, passive, unpowered directional gain in the low-frequency band.**

Why the mirrors were abandoned is the part that generalises. Radar rendered them obsolete within a few years, and it is tempting to read that as the geometry failing. It is not what happened. **The mirrors did not fail; their input signal moved out from under them.** Aircraft got faster, so warning time — range divided by closing speed — shrank toward uselessness while the device stayed constant. The acoustic environment got noisier, and a passive device has no gain control and no way to reject a competing source inside its beam. And radar carried its own signal: a passive collector concentrates only what happens to arrive, while an active system chooses what to transmit, when, and at what power, and therefore controls its own signal-to-noise ratio.

The general principle, and it governs every geometric intervention in this paper: passive geometry has no gain control and generates no signal of its own. A perfectly shaped alcove in an organisation with nothing to say will concentrate silence with excellent efficiency. Geometry sets the coefficients `N` and `C`; it does not supply `S`, and the conversion coefficient `eta` is a separate lever with its own analysis [Curzi, 2026b]. This is failure mode **FG7** of Appendix B and it is the one that wastes the most capital.

11.4.6 4.6 Where the parabola belongs, and the concave-room prohibition

At furniture and niche scale, aimed deliberately.

Application	Aperture	Frequency floor	What it does
Listening niche behind a seated head	1.0–1.4 m	~490–690 Hz	Raises consonant-band SNR for one seated person facing the room
Desk-scale shell	0.5–0.7 m	~980–1372 Hz	Consonant-band privacy shield; reduces spill both ways

Application	Aperture	Frequency floor	What it does
Alcove back-shell, shallow, `f/D` ~0.35	2.0–2.6 m	~264–343 Hz	Sub-divides an open floor acoustically without a wall
Outdoor listening seat / whisper pair	1.5–3.0 m	~229–457 Hz	Triangulating object that is also a demonstration

It does not belong as the shape of a room, and this is the most important negative in the section. In a room, concave surfaces are usually a defect. A domed ceiling, a curved end wall, a barrel vault or a concave glazed facade concentrates sound at focal points and produces:

- **hot spots** — positions where a talker is uncomfortably loud and positions two metres away where they are inaudible, which destroys the turn-taking equality Section 7 is about;
- **focused echoes** — a discrete, delayed, high-level reflection at a specific seat. Discrete echoes damage intelligibility far more than an equal amount of diffuse reverberation, because the ear integrates diffuse energy and cannot separate a strong late arrival from the direct sound;
- **position-dependent performance**, which is a governance problem as much as an acoustic one. **A room whose best seat is acoustically privileged has installed a head of the table by accident**, and the governance consequence of that is developed in Curzi [2026f], whose deference model gives `u = 1 − a^4` in a common rotation coefficient `a` and whose thirteen-seat ring table drove the estimated `a` from 0.83 to 0.47.

Requirement **G9** of Section 8 states the prohibition. **The parabola is a scalpel, not wallpaper.**

11.5 5 Self-similarity and boundary multiplication

11.5.1 5.1 Similarity dimension, and what it does not tell you

Theorem 4 (similarity dimension). A set composed of `m` copies of itself, each scaled by `1/r`, has similarity dimension `D = ln m / ln r`. This is true by the definition of self-similarity.

Set	`m`	`1/r`	`D = ln m / ln r`
Cantor set	2	3	0.6309
Koch curve	4	3	1.2619
Sierpinski gasket	3	2	1.5850
Sierpinski carpet	8	3	1.8928
Menger sponge	20	3	2.7268
Hilbert curve	4	2	2.0000

The Koch snowflake makes the architecturally useful property visible: its perimeter grows without bound while its area converges.

iteration `n`	perimeter (unit-side triangle)	multiple of original	area / `A_0`
0	3.0000	1.0000	1.00000
1	4.0000	1.3333	1.33333

iteration `n`	perimeter (unit-side triangle)	multiple of original	area / `A_0`
2	5.3333	1.7778	1.48148
3	7.1111	2.3704	1.54733
4	9.4815	3.1605	1.57659
5	12.6420	4.2140	1.58960
6	16.8560	5.6187	1.59538
7	22.4746	7.4915	1.59794
8	29.9662	9.9887	1.59909

Area converges to exactly `8/5 = 1.6` times the original triangle; perimeter multiplies by `4/3` every iteration and diverges. **This is the whole architectural point in one table: boundary is cheap, interior is expensive, and self-similarity is the technique for buying large amounts of the first with bounded amounts of the second.** The edge effect — that people occupy edges first [Gehl, 1971] — and the sittable-space ratio [Whyte, 1980] are both demands for boundary. A room's floor area is fixed by the lease; its boundary length is a design variable, and Theorem 5 prices it exactly.

What the dimension does not tell you. `D = 1.2619` for the Koch curve is a statement about how its measured length scales with ruler size. It says nothing about whether the shape scatters sound usefully at any particular frequency, whether people will sit in its recesses, whether those recesses are reachable, cleanable or safe, or whether the geometry resonates, focuses or traps anything undesirable. **Two boundaries with identical `D` can behave completely differently.** That is why the self-similarity claim of Section 2 is badged as an analogy and why every requirement in Section 8 that invokes self-similarity carries a modelling obligation.

11.5.2 5.2 The antenna precedent and where the analogy stops

The commercial precedent is genuine and worth stating precisely, because it is the strongest available evidence that self-similar geometry does real engineering work. An antenna's resonant behaviour is set by its electrical length relative to wavelength, and two consequences follow. **Space-filling lowers the resonant frequency for a fixed physical envelope**: a Koch monopole of a given height contains more conductor than a straight monopole of the same height, so it resonates lower — which for a handset, where the envelope is fixed by industrial design and the bands are fixed by regulators, is the whole problem. And **self-similarity at scale factor `s` produces resonances at frequency ratios approximately `s`**, so a structure containing copies of itself at half scale tends to resonate at `f`, `2f`, `4f`, `8f`: a multiband antenna from one radiating element rather than several.

Both are load-bearing. Now the limit. The claim that fractality is *optimal* — that a fractal antenna necessarily outperforms a well-designed conventional multiband antenna of the same size — is contested in the antenna literature, and careful comparisons have found conventional meander and inverted-F designs performing comparably. **The safe statement is that self-similarity is a reliable route to multiband, size-reduced behaviour, not a guarantee of superiority.** This paper cites no key for the antenna literature because the bibliography assembled for this corpus contains none; the claim is therefore recorded as **uncited and taken from engineering practice**, and Section 12 lists it among the unsupported statements.

How far the analogy carries. It carries this far: a self-similar boundary offers usable features at several characteristic sizes simultaneously, so one geometry presenting 0.6 m recesses, 2.2 m alcoves and 6.5 m bays serves the pair, the small group and the assembly out of one construction — the same economy a self-similar radiator achieves across bands. **It stops here:** there is no scale-invariance principle for human occupancy. People have one size; proxemic bands are absolute distances, not ratios [Hall, 1966], and sensory thresholds are absolute

distances [Gehl, 1971]. **A social geometry is self-similar over at most three or four scales before it meets the fixed size of a human being at the bottom and the fixed limit of a sightline at the top.** That is a bounded, checkable claim, and Section 5.4 builds exactly three social rungs.

11.5.3 5.3 The crenellation multiplier

Theorem 5 (crenellation multiplier). A boundary replaced by a square-wave of tooth depth `d` and half-period `w` has its length multiplied by exactly `1 + d/w`. *Proof, by summing segments:*

```
per period of length 2w, the path runs:  w  along the outer line
                                         d  inward
                                         w  along the inner line
                                         d  outward
path length per period = 2w + 2d
multiplier = (2w + 2d) / 2w = 1 + d/w
```

If the teeth alternate inward and outward about the original line, the enclosed area is unchanged exactly. Applied to a square room of 120 m^2 — side 10.954 m, perimeter 43.818 m:

`d`	`w`	multiplier	new perimeter
0.50 m	1.50 m	1.333	58.42 m
0.75 m	1.50 m	1.500	65.73 m
1.00 m	1.50 m	1.667	73.03 m
1.00 m	1.00 m	2.000	87.64 m
1.50 m	1.50 m	2.000	87.64 m

And it iterates. Applying the same `d/w = 2/3` at each of three self-similar orders, each order at one third the size of the previous:

order	`d`	`w`	multiplier this order	cumulative	perimeter
1	1.000 m	1.500 m	1.667	1.667	73.03 m
2	0.333 m	0.500 m	1.667	2.778	121.7 m
3	0.111 m	0.167 m	1.667	4.630	202.9 m

Order 3 is where the arithmetic detaches from reality. A 111 mm recess at 167 mm pitch is not a seat, not a niche and not reachable; it is a texture. It has real acoustic effects — at 111 mm scale it scatters sound above roughly 1.5 kHz, which is useful diffusion — but it does not carry a person, and counting perimeter metres at that scale toward a sittable-edge requirement would be straightforwardly false. **Only boundary at scales that can hold the thing being counted for may be counted:** for seating, `d >= 0.40 m` and `w >= 0.60 m`. That single rule is what separates a fractal argument from fractal decoration, and it is requirement **G15** of Section 8.

The inward-only case is the real one. In an occupied building the teeth rarely push outward: there is a structural wall, a corridor or another tenant. Inward-only crenellation converts open floor into niche floor. For the same 120 m^2 room, recessing half the perimeter run — 21.91 m — inward by `d`, with the sittable-edge requirement recomputed on the reduced open floor at the 0.1094 ratio:

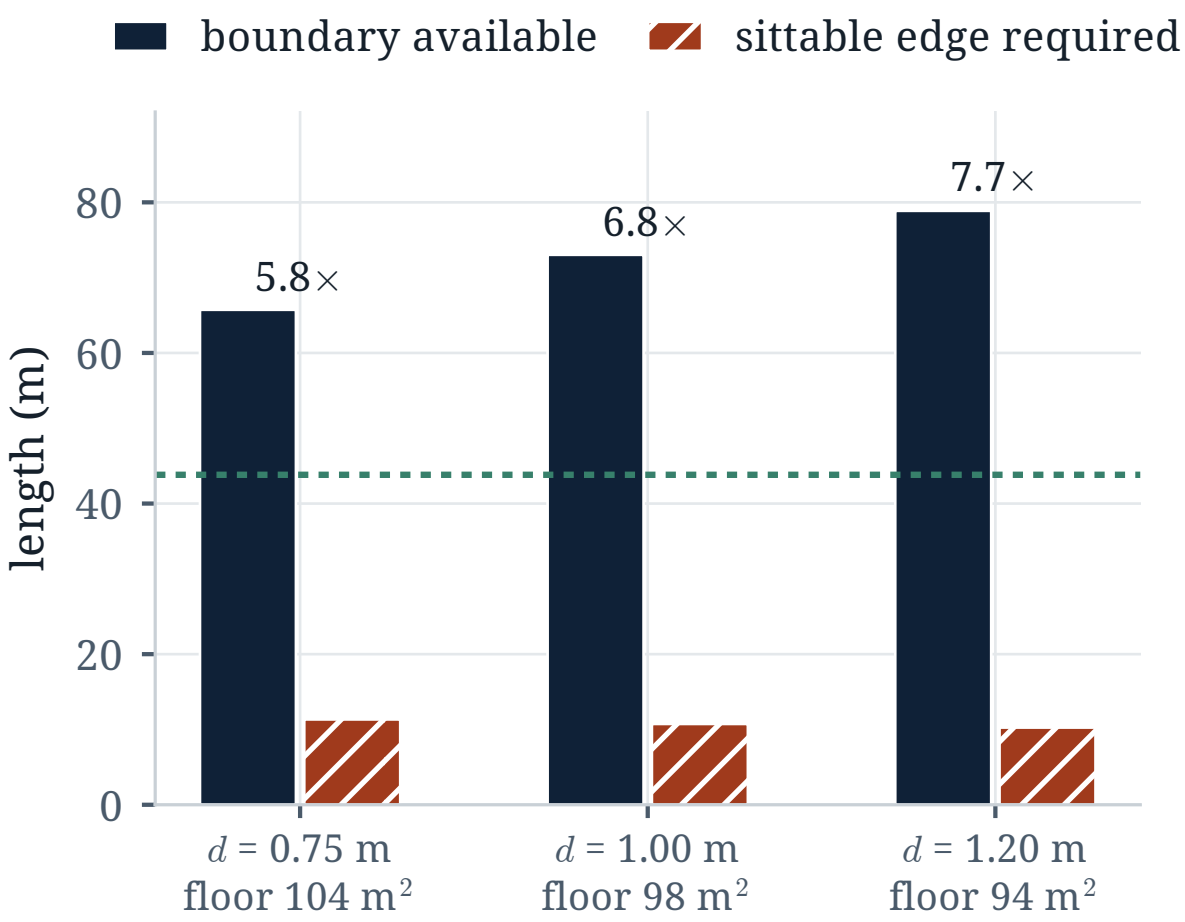


Figure 11.1: **Boundary multiplication converts an unmeetable seating-edge requirement into a several-fold surplus, and it does it by spending floor area on the thing the floor area existed to enable.** At `d = 1.00 m` and `w = 1.50 m` the 120 m^2 room gives up 21.9 m^2 of open floor — 18 per cent — and gains 29.2 m of boundary, against a requirement that has simultaneously fallen to 10.73 m. Exact arithmetic from the multiplier `1 + d/w`; the 0.1094 sittable-edge ratio is a stipulated design target, not a measurement. Generated by `figs/src/CCP-011_fig02.py`.

d	new perimeter	niche area created	remaining open floor	edge required	edge available
0.75 m	65.73 m	16.43 m^2	103.57 m^2	11.33 m	65.73 m
1.00 m	73.03 m	21.91 m^2	98.09 m^2	10.73 m	73.03 m
1.20 m	78.87 m	26.29 m^2	93.71 m^2	10.25 m	78.87 m

At `d = 1.00 m`, **the room gives up 18 per cent of its open floor and ends with roughly seven times the required sittable edge.** And the 21.9 m^2 is not lost: it is exactly where the seating goes. **The floor area was converted into the thing the floor area existed to enable.** An open 120 m^2 room with 43.8 m of unusable perimeter is worse in every measurable respect than a 98 m^2 room with 73 m of occupied edge. The construction need not be structural — free-standing banquette backs, tall shelving, storage walls, planters and screens produce the same multiplication at a retention loss and a cost far below masonry.

11.5.4 5.4 The nested enclosure ladder

Choose a scale ratio `r = 3`, chosen because it lands the rungs close to the human thresholds already established and not because three is special:

k	Scale	L_k = 0.75 x 3^k	Area ~ L^2	What it holds	Threshold it matches
0	desk	0.75 m	0.6 m^2	one person, one task	personal distance, 0.45–1.2 m
1	alcove	2.25 m	5.1 m^2	two to four	social distance, 1.2–3.6 m
2	room	6.75 m	45.6 m^2	six to twelve	full-sensory range, ~7 m
3	wing	20.25 m	410 m^2	thirty to eighty	between 7 m and 25 m
4	quarter	60.75 m	3691 m^2	a few hundred	beyond emotion-reading range, ~25 m
5	campus	182.25 m	33215 m^2	thousands	beyond person-identification range, ~100 m

The distance thresholds in the right-hand column are the proxemic and sensory bands of Hall [1966] and Gehl [1971]. **Two of these rungs do something the others do not: the desk-to-room span, `k = 0` to 2, is the social ladder, and above it the rungs are logistical.** Past roughly 20 m other people are not conversational partners and the enclosure's job changes from enabling contact to organising movement. Section 5.2's limit is visible here: **the self-similar social geometry is three rungs deep, perhaps four, and claims beyond that are architecture rather than sociology.**

For a two-dimensional plan, each rung holds `m = r^2 = 9` sub-units of the rung below, giving `D = ln 9 / ln 3 = 2.0000` — which is simply the statement that a plan fills a plane. **A plan whose measured `D` came out below 2 would be reporting that it has gaps.** This is a sanity check, not a design target.

The design consequence is **multi-scale response in bounded volume**. A room built with a crenellated boundary at `d = 1.0 m` and a secondary order at `d = 0.35 m` simultaneously offers 0.35 m recesses at shoulder scale, 1.0 m recesses giving a seat with the back protected, 2.25 m bays formed by pairs of recesses, and the 6.75 m room itself. **One geometry, four occupancy scales, no moving parts, no power and no booking system** — which is how a single construction satisfies a four-category requirement without four separate rooms.

11.5.5 5.5 Parameters that change with scale

The modelling choice of Section 2 says the same equation applies at every rung with different parameters. Here is the ladder with illustrative parameters. **These are illustrative values chosen to show the shape of the drift. They are not measurements, none carries an interval, and any real project must fill this table from its own survey**, against the evidential standard of Curzi [2026], whose four-drawer classification — PROVED, POSTULATED, BORROWED, UNTESTED — records the surplus form as PROVED CONDITIONALLY on a chosen design axiom and the ladder's operationalisation as SPECIFIED, NOT RUN. **This table is in the last of those drawers.**

Scale	L (m)	delta	tau ~ 1/delta	eta	N	ln N	eta ln N	Q	delta Q	C	h
desk	0.75	0.600	1.7	0.40	3	1.0986	0.4394	0.2	0.1200	0.05	+0.2694
alcove	2.25	0.200	5.0	0.30	4	1.3863	0.4159	1.0	0.2000	0.12	+0.0959

Scale	L (m)	delta	tau ~ 1/delta	eta	N	ln N	eta ln N	Q	delta Q	C	h
room	6.75	0.080	12.5	0.25	6	1.7918	0.4479	3.0	0.2400	0.15	+0.0579
wing	20.25	0.030	33.3	0.15	10	2.3026	0.3454	20.0	0.6000	0.60	-0.8546
campus	60.75	0.010	100.0	0.08	16	2.7726	0.2218	150.0	1.5000	2.00	-3.2782

Read the columns as a set of monotone trends, because the trends are the content. **`delta` falls with scale** — a desk arrangement decays in days, a campus in decades — so **`tau = 1/delta` rises** from under two periods to a hundred: larger enclosures are slower filters that respond to sustained forcing rather than to impulses, and it is not a criticism of a one-day event to say that it changes a desk and not a campus. **`eta` falls with scale**, because conversion is most efficient in the smallest unit, where the person who had the idea is also the person who can act on it, and every rung upward inserts a handoff. **`N` rises with scale but only logarithmically in effect**: from desk to campus, `N` goes 3 to 16 while `ln N` goes 1.0986 to 2.7726, so a factor of 5.3 in categories buys a factor of 2.5 in ceiling. **`eta ln N` is therefore roughly flat through the social rungs and then falls.** And **`Q` explodes**, from 0.2 to 150.

step	`Q` grows by	`ln N` grows by
desk to alcove	5.00x	1.26x
alcove to room	3.00x	1.29x
room to wing	6.67x	1.29x
wing to campus	7.50x	1.20x

`Q` grows several-fold per rung while `ln N` grows by about a quarter, so `delta Q` outruns `eta ln N` and headroom goes negative as scale increases — in this illustration, between the room and the wing. That is the fixed-point behaviour of the gauge arriving through pure geometry: at equilibrium the headroom of a converged system is exactly `eta(ln N − F(u*))`, the surplus it is failing to extract from its own difference [Curzi, 2026b]. **The reason large organisations feel stuck is not that their people are worse; it is that `delta Q` scales with the stock and `eta ln N` does not.** The design consequence is to **push work down the ladder**: whatever can be decided at the alcove should be, because `eta` is higher and `delta Q` lower there.

This table is an illustration and not a demonstration, and it is the paper's weakest quantitative object. Five rows of stipulated parameters cannot establish a trend; they can only show what the trend would look like if it existed. Failure mode **FG6** of Appendix B names the abuse — applying the form past the scale at which its parameters were measured — and the rule is that the scale-free form of the law licenses modelling at a new scale and does not license skipping the measurement at that scale.

11.5.6 5.6 The seating pattern and its falsification condition

The pattern language tradition is the acknowledged ancestor here [Alexander et al., 1977, Alexander, 1965]. What this analysis adds is that **the same seating pattern, correctly stated, is valid at every rung**, which is what would make the geometry self-similar in a useful rather than a decorative sense. The pattern, stated once: a person will occupy a position that has **(a)** the back protected, **(b)** a view out over a larger territory, **(c)** a visible low-cost exit, and **(d)** a boundary at reachable scale.

Rung	(a) back protected by	(b) view out over	(c) exit	(d) reachable boundary
desk, 0.75 m	a 1.2 m screen or shell	the alcove	one step, unobstructed	desk edge, shelf lip
alcove, 2.25 m	banquette back, storage wall	the room	alcove mouth, 1.2 m clear	ledge at 430 mm, table edge

Rung	(a) back protected by	(b) view out over	(c) exit	(d) reachable boundary
room, 6.75 m	the crenellated wall	the bay or court	two doors, opposed	window ledge at 450 mm
wing, 20.25 m	a built edge, colonnade, arcade	the court	multiple, legible	parapet, planter, low wall

The pattern is invariant; only the dimensions scale. Unlike the parameter claim of Section 5.5, this one is directly falsifiable: build all four rungs, count occupancy at each, and see whether the pattern predicts where people go. **If positions satisfying (a)–(d) are not occupied preferentially over positions failing one or more of them, the pattern is wrong for that population and the seating requirements of Section 8 should be discarded for that project.** Say so in the survey rather than defending the pattern against the count. That is H5, and no count reported in this paper tests it.

11.5.7 5.7 The three-tier certification standard

Every geometric claim in this paper falls into one of three tiers, and the tier determines what is owed before anything may be built on it.

Tier 1 — proven, use freely. Theorems 1, 2, 4 and 5 and Corollary 3. These are geometry. They are exact and require nothing but correct arithmetic.

Tier 2 — standard physics with stated approximations, use with the approximation declared. The `70 lambda/D` beamwidth relation, the `(pi D/lambda)^2 eta_ap` gain relation, the `D >= 2 lambda` frequency floor, and Sabine's equation in Section 6. Each is an engineering approximation with a known domain. **Declare `eta_ap`, declare that near-field performance falls short, declare that the figures are single-frequency.**

Tier 3 — requires modelling or measurement of the specific geometry; may not be certified from the concept. Anything depending on detailed shape: the scattering behaviour of a particular crenellation, the diffusion coefficient of a particular relief, the acoustic sub-division actually achieved by a particular alcove depth, the occupancy response to a particular nested plan. **Every fractal claim in this paper is Tier 3, without exception.**

The acceptable methods for discharging a Tier 3 obligation, in ascending order of cost, are: **(1)** scale physical measurement — a cardboard mock-up and a phone measuring `RT60` and a level difference at two positions is a real measurement and costs an afternoon; **(2)** geometrical-acoustics simulation by ray or image-source method, valid above the Schroeder frequency, which for a 100 m^3 room at `RT60 = 0.5 s` is in the low hundreds of hertz [Schroeder, 1965]; **(3)** wave-based simulation, required below the Schroeder frequency and for any claim about resonance, focusing or small-scale diffusion; **(4)** full-scale prototype and measurement, the only method that settles a Tier 3 question definitively. **A specification invoking self-similarity without naming which of these four was used has specified nothing.**

This is the same discipline the corpus applies to its own claim ledger, where each result carries an evidential drawer and the protocols are stated to run on the established mathematics and psychology while steering by the declared postulates [Curzi, 2026,?]. **The tier system is that ledger applied to shapes.**

11.6 6 The diffuse-field model and four build states

11.6.1 6.1 Sabine's equation: five assumptions, and where it breaks

```
RT60 = 0.161 V / A      V in cubic metres, A in metric sabins, RT60 in seconds
```

RT60 is the time for the sound energy density in the room to fall by 60 dB after the source stops, and A is the total absorption: A = sum_i S_i alpha_i + sum_j A_j, where S_i is the area of surface i, alpha_i its absorption coefficient in the band of interest, and A_j the absorption of a discrete object such as a person or a chair, in square metres of equivalent absorption each. The constant 0.161 carries units of s/m and follows from the speed of sound and the mean free path between reflections [Sabine, 1922, Rossing et al., 2002].

Three further standard results are used and each is a derivation. The **two-field model**, L_p(r) = L_W + 10 log10(Q/(4 pi r^2) + 4/R), separates a direct field falling 6 dB per doubling of distance from a reverberant field that is the same everywhere, with Q the source directivity factor — about 2 for a talker on a table — and R = A/(1 − alpha_bar) the room constant, where alpha_bar = A/S_total. The **critical distance** d_c = 0.141 sqrt(Q R) follows from setting the two terms equal: beyond it the reverberant field dominates and moving closer to the talker stops helping. And the **decay rate** follows from the definition of RT60: sound decays at 60/RT60 decibels per second, so a gap of t seconds leaves the previous talker's energy 60t/RT60 decibels down. That last identity is the whole of Section 7.2.

Sabine's derivation assumes a diffuse field, and five conditions must roughly hold.

Condition 1 — the room must be above its Schroeder frequency. Below it, behaviour is dominated by discrete resonant modes rather than a statistical field, and RT60 is not a meaningful single number [Schroeder, 1965]. With f_s = 2000 sqrt(RT60/V) Hz:

V	RT60	f_s
100 m^3	0.40 s	126.5 Hz
100 m^3	0.50 s	141.4 Hz
100 m^3	1.00 s	200.0 Hz
100 m^3	2.00 s	282.8 Hz
336 m^3	0.70 s	91.3 Hz
500 m^3	0.80 s	80.0 Hz
2000 m^3	1.20 s	49.0 Hz
10000 m^3	1.80 s	26.8 Hz

In a 100 m^3 meeting room the Schroeder frequency sits around 130–140 Hz, so the 125 Hz octave band — the leftmost column of every table below — is at or below the limit of Sabine's validity and must be read as indicative only. For an 8.0 x 5.0 x 2.5 m room the axial modes fall at 21.4, 34.3, 42.9, 64.3, 68.6 (two coincident), 85.8, 102.9 and 137.2 (two coincident) Hz; the coincidences at 68.6 and 137.2 Hz are the kind of modal pile-up that produces a boomy room, **and no RT60 figure will reveal them.**

Condition 2 — mean absorption must be reasonably low. Sabine over-predicts RT60 as alpha_bar rises, because the equation permits RT60 to remain finite as alpha_bar approaches 1, which is physically wrong: a room with total absorption has no reverberation at all. Eyring's correction fixes the limit:

```
Sabine:  RT60 = 0.161 V / A
Eyring:  RT60 = 0.161 V / ( -S * ln(1 - alpha_bar) )      alpha_bar = A / S
```

Computed for the four build states of Section 6.3, in a 100 m^3 room of 145 m^2 total surface:

State	Band	A	alpha_bar	Sabine RT60	Eyring RT60	Sabine over-predicts by
A hard shell	500 Hz	8.73	0.060	1.844 s	1.788 s	3.1%
A hard shell	1000 Hz	8.24	0.057	1.954 s	1.898 s	3.0%
B office fit	500 Hz	35.53	0.245	0.453 s	0.395 s	14.7%

State	Band	`A`	`alpha_bar`	Sabine RT60	Eyring RT60	Sabine over-predicts by
B office fit	2000 Hz	61.63	0.425	0.261 s	0.201 s	30.2%
C specified	500 Hz	36.94	0.255	0.436 s	0.378 s	15.4%
C specified	4000 Hz	41.52	0.286	0.388 s	0.329 s	17.8%
D over-treated	1000 Hz	94.52	0.652	0.170 s	0.105 s	**61.9%**
D over-treated	2000 Hz	104.99	0.724	0.153 s	0.086 s	**77.8%**

The switching rule: use Sabine below `alpha_bar = 0.25`, use Eyring above it, and state which was used. In the dead-room case the two formulae disagree by nearly a factor of two, and a specification quoting Sabine without saying so would be describing a room 78 per cent more reverberant than the one that got built. This is result R4 of the paper and it costs nothing to apply.

Condition 3 — absorption must be reasonably evenly distributed. A room with all its absorption on one surface — the classic absorptive ceiling over five hard surfaces — does not develop a diffuse field: sound bouncing in the horizontal plane never meets the ceiling and decays far more slowly than Sabine predicts. **This is the most common real-world failure of the equation** and it is why Section 8 requires absorption on at least two non-parallel planes.

Condition 4 — the room must not be strongly coupled to another volume. An open door to a corridor, a folding partition or a double-height void creates coupled spaces with two-slope decay curves that no single `RT60` describes.

Condition 5 — dimensions must not be wildly disparate. Corridors, and rooms far longer than they are wide or tall, violate the mean-free-path derivation.

Air absorption adds a `4mV` term to `A` and matters above about 2 kHz in large volumes. In a 100 m^3 room at ordinary humidity it is negligible and is omitted throughout; in a 2000 m^3 hall it is not, and omitting it there would be an error.

11.6.2 6.2 Absorption coefficients, all nominal

Every value in this table is NOMINAL — a typical published value for a generic material description, per octave band. They are adequate for design arithmetic and for deciding what to build, and **they are not adequate for a construction specification.** Mounting changes low-frequency performance dramatically, and a panel tested with an air gap behind it will not perform the same way glued flat to a wall.

Material (nominal)	125	250	500	1k	2k	4k
Concrete, sealed or painted; terrazzo	0.01	0.01	0.02	0.02	0.02	0.03
Concrete block, painted	0.10	0.05	0.06	0.07	0.09	0.08
Concrete block, coarse or unpainted	0.36	0.44	0.31	0.29	0.39	0.25
Gypsum board 13 mm on studs, air cavity	0.29	0.10	0.05	0.04	0.07	0.09
Plaster on solid backing	0.02	0.02	0.03	0.04	0.05	0.05
Glass, heavy plate	0.18	0.06	0.04	0.03	0.02	0.02
Glass, ordinary window	0.35	0.25	0.18	0.12	0.07	0.04
Wood floor on joists	0.15	0.11	0.10	0.07	0.06	0.07

Material (nominal)	125	250	500	1k	2k	4k
Plywood panel 10 mm over air space	0.28	0.22	0.17	0.09	0.10	0.11
Carpet, heavy, on concrete	0.02	0.06	0.14	0.37	0.60	0.65
Carpet, heavy, on resilient underlay	0.08	0.24	0.57	0.69	0.71	0.73
Heavy velour drape, draped to half its flat area	0.14	0.35	0.55	0.72	0.70	0.65
Mineral-fibre ceiling tile, suspended with plenum	0.20	0.40	0.60	0.70	0.75	0.75
Mineral wool panel 25 mm, direct to wall	0.10	0.30	0.65	0.85	0.90	0.90
Mineral wool panel 50 mm, direct to wall	0.20	0.55	0.85	0.95	0.95	0.90
Mineral wool panel 100 mm with 50 mm air gap	0.45	0.85	0.95	0.98	0.95	0.90
Bookcase, filled, open front	0.25	0.30	0.30	0.30	0.30	0.30
Water surface	0.008	0.008	0.008	0.008	0.008	0.008

Discrete absorbers, in square metres of equivalent absorption **each**, not per m^2:

Object (nominal)	125	250	500	1k	2k	4k
Adult, seated, on a lightly upholstered chair	0.15	0.25	0.35	0.42	0.46	0.46
Empty lightly upholstered chair	0.10	0.15	0.18	0.20	0.22	0.22

Four things to read out of this table before using it.

Porous absorbers fail at low frequency, and thickness is the reason. A porous absorber works by viscous loss as air moves through it, and air motion is greatest a quarter-wavelength from a rigid boundary. At 125 Hz a quarter wavelength is 0.69 m. **A 25 mm panel cannot reach the low-frequency air motion; a 100 mm panel with a 50 mm gap gets much closer, which is why its 125 Hz figure is 0.45 against the 25 mm panel's 0.10.** Every low-frequency problem in this paper traces to that one fact.

Gypsum board on studs is a low-frequency absorber. Its 0.29 at 125 Hz is not a misprint: a panel on a cavity is a mass-spring system that resonates and dissipates in the bass, and it is the only absorption most ordinary rooms have down there.

People are significant absorbers and their absorption is mid- and high-biased. Eight seated adults contribute 2.8 m^2 of equivalent absorption at 500 Hz and 3.4 at 1 kHz. In a small hard room that is a material fraction of the total, **so the occupied and unoccupied `RT60` of a small room differ substantially and every specification must say which one it means.**

Carpet on concrete does almost nothing below 500 Hz. Its 125 Hz figure of 0.02 is the same as bare concrete: carpet is a treble absorber, and a room treated with carpet alone gets darker, not clearer.

11.6.3 6.3 The 100 m^3 room, computed in four states

Room: 8.0 x 5.0 x 2.5 m = 100.0 m^3 exactly. Floor 40 m^2, ceiling 40 m^2, walls `2(8 x 2.5) + 2(5 x 2.5) = 65 m^2` of which 12 m^2 is glazing. Total surface `S` = 145 m^2. Eight occupants throughout.

First, the shape of the relationship, since it governs every decision that follows. With Sabine at `V = 100`:

`A` (m^2 sabins)	4	6	8	10	12	15	20	25	30	40	60	80	100
RT60 (s)	4.025	2.683	2.013	1.610	1.342	1.073	0.805	0.644	0.537	0.403	0.268	0.201	0.161

And inverted, the absorption a target demands at `V = 100 m^3`: 53.67 for 0.30 s, 40.25 for 0.40 s, 32.20 for 0.50 s, 26.83 for 0.60 s, 23.00 for 0.70 s, 20.12 for 0.80 s, 16.10 for 1.00 s, 13.42 for 1.20 s, 10.73 for 1.50 s and 8.05 for 2.00 s. **`RT60` is inversely proportional to `A`, so the first few square metres of absorption are worth enormous amounts and the last few almost nothing.** Going from `A = 8` to `A = 16` halves `RT60` from 2.01 s to 1.01 s; going from `A = 80` to `A = 88` moves it from 0.201 s to 0.183 s. **This is the same concavity that governs the surplus functional and it carries the same moral: the second unit is the valuable one.** For sizing any room, the absorption required for 0.5 / 0.6 / 0.8 s runs 9.66 / 8.05 / 6.04 at 30 m^3, 16.10 / 13.42 / 10.06 at 50 m^3, 32.20 / 26.83 / 20.12 at 100 m^3, 64.40 / 53.67 / 40.25 at 200 m^3, 161.00 / 134.17 / 100.62 at 500 m^3 and 322.00 / 268.33 / 201.25 at 1000 m^3.

State A — the as-built hard shell. Sealed concrete floor, gypsum board ceiling and walls on studs, 12 m^2 heavy plate glazing, eight occupants. This is what a shell-and-core office looks like before fit-out, and what a great many converted spaces look like permanently.

Element	Area	125	250	500	1k	2k	4k
Sealed concrete floor	40.0	0.40	0.40	0.80	0.80	0.80	1.20
Gypsum ceiling	40.0	11.60	4.00	2.00	1.60	2.80	3.60
Gypsum walls	53.0	15.37	5.30	2.65	2.12	3.71	4.77
Glazing	12.0	2.16	0.72	0.48	0.36	0.24	0.24
Occupants, eight	–	1.20	2.00	2.80	3.36	3.68	3.68
Total `A`		**30.73**	**12.42**	**8.73**	**8.24**	**11.23**	**13.49**
RT60 (s)		**0.52**	**1.30**	**1.84**	**1.95**	**1.43**	**1.19**
`alpha_bar`		0.212	0.086	0.060	0.057	0.077	0.093
Room constant `R` (m^2)		39.0	13.6	9.3	8.7	12.2	14.9
Critical distance, `Q` = 1 (m)		0.88	0.52	0.43	0.42	0.49	0.54

An `RT60` of 1.95 s at 1 kHz in a room eight metres long, with a critical distance of 0.42 m — closer than one seat is to the next. Every listener at every seat is in the reverberant field; moving closer to the talker does not help; **there is no seat in this room from which a soft voice can reliably be heard.** Note the spectrum: 0.52 s at 125 Hz against 1.95 s at 1 kHz. The room is *less* reverberant in the bass, because the gypsum-on-studs is absorbing bass while nothing absorbs mid-frequencies. That is the signature of an untreated stud-and-board room.

State B — the ordinary office fit-out. Heavy carpet on the concrete, mineral-fibre suspended ceiling tile over the whole 40 m^2, gypsum walls, glazing unchanged, eight occupants: the default fit-out of the last forty years.

Element	Area	125	250	500	1k	2k	4k
Carpet on concrete	40.0	0.80	2.40	5.60	14.80	24.00	26.00

Element	Area	125	250	500	1k	2k	4k
Mineral-fibre ceiling tile	40.0	8.00	16.00	24.00	28.00	30.00	30.00
Gypsum walls	53.0	15.37	5.30	2.65	2.12	3.71	4.77
Glazing	12.0	2.16	0.72	0.48	0.36	0.24	0.24
Occupants, eight	–	1.20	2.00	2.80	3.36	3.68	3.68
Total A		**27.53**	**26.42**	**35.53**	**48.64**	**61.63**	**64.69**
RT60 (s)		**0.58**	**0.61**	**0.45**	**0.33**	**0.26**	**0.25**
Critical distance, Q = 1 (m)		0.82	0.80	0.97	1.21	1.46	1.52

State B passes a crude single-number test — 0.45 s at 500 Hz, comfortably inside a 0.6 s target — and is nonetheless a badly balanced room. Its bass ratio `(125+250)/(500+1k)` is 1.53, and it is 2.4 times more reverberant at 250 Hz than at 4 kHz. Consonants — the 2–8 kHz band carrying most of the information that distinguishes one word from another — are absorbed hard, while the low-mid energy that masks them is left to ring. **A room that over-absorbs treble and under-absorbs bass has applied a low-pass filter to every voice in it, and low-pass filtering is exactly what makes speech muddy.** No single-number specification can produce that diagnosis, which is why every requirement in Section 8 is stated per band.

State C — the specified deliberation room. Wood floor on joists rather than carpet; mineral-fibre tile over 20 m^2 of the ceiling with the central 20 m^2 left hard as an early-reflection reflector over the table; 10 m^2 of 100 mm mineral wool with a 50 mm air gap on the end walls; a 9 m^2 filled bookcase as a broadband diffuser-absorber; drapes over half the glazing; eight occupants.

Element	Area	125	250	500	1k	2k	4k
Wood floor on joists	40.0	6.00	4.40	4.00	2.80	2.40	2.80
Mineral-fibre ceiling tile	20.0	4.00	8.00	12.00	14.00	15.00	15.00
Hard ceiling reflector over table	20.0	5.80	2.00	1.00	0.80	1.40	1.80
Mineral wool 100 mm + 50 mm gap	10.0	4.50	8.50	9.50	9.80	9.50	9.00
Bookcase, filled	9.0	2.25	2.70	2.70	2.70	2.70	2.70
Gypsum walls, remaining	28.0	8.12	2.80	1.40	1.12	1.96	2.52
Drape over half the glazing	6.0	0.84	2.10	3.30	4.32	4.20	3.90
Exposed glazing	6.0	1.08	0.36	0.24	0.18	0.12	0.12
Occupants, eight	–	1.20	2.00	2.80	3.36	3.68	3.68
Total A		**33.79**	**32.86**	**36.94**	**39.08**	**40.96**	**41.52**
RT60 occupied (s)		**0.48**	**0.49**	**0.44**	**0.41**	**0.39**	**0.39**
RT60 unoccupied (s)		0.49	0.52	0.47	0.45	0.43	0.43
alpha_bar		0.233	0.227	0.255	0.270	0.282	0.286
Critical distance, Q = 1 (m)		0.94	0.92	0.99	1.03	1.07	1.08

Bass ratio 1.14 — essentially flat — and an occupied-to-unoccupied variation of 0.01 to 0.04 s. **`RT60` varies by 0.10 s across six octave bands. That is a flat room, and flatness matters far more than the absolute value.**

What achieved it is counter-intuitive. State C has *less* mid-high absorption than State B: 39.1 against 48.6 at 1 kHz and 41.0 against 61.6 at 2 kHz. The design removed the carpet, halved the ceiling tile, and added low-frequency-capable absorption. **The fix for a badly balanced room is often to take absorption away at the top and add it at the bottom, and no single-number specification can express that.** Note also the hard 20 m^2 ceiling panel over the table, which is not an oversight: it supplies early reflections from talker to far listener, arriving within a few milliseconds of the direct sound, which the ear integrates with the direct sound and which therefore *raise* effective loudness and intelligibility. **Early reflections are the quiet talker's ally, late reverberation is the quiet talker's enemy, and both are made of the same material.** The difference is entirely geometry and timing.

State D — the over-treated room, as a warning. Carpet, full mineral-fibre ceiling, 40 m^2 of 100 mm mineral wool on the walls, drapes over all the glazing. Someone was told to add absorption and did.

	125	250	500	1k	2k	4k
Total A	33.45	59.90	77.65	94.52	104.99	104.65
RT60 Sabine (s)	0.48	0.27	0.21	0.17	0.15	0.15
RT60 Eyring (s)	0.42	0.21	0.15	0.11	0.09	0.09
`alpha_bar`	0.231	0.413	0.536	0.652	0.724	0.722

Bass ratio 1.99. **State D is worse than State B and cost several times as much**, and it fails twice at once. **It is dead:** at 0.11 s by Eyring at 1 kHz the room provides no reverberant support, so talkers push because none of their own energy returns, vocal effort and fatigue rise, and the quiet talker — who cannot push — is worse off than in a moderately live room. **And its bass ratio got worse, not better:** all that porous absorption works from 250 Hz up and barely touches 125 Hz, so the relative bass excess doubled. The room is dead and boomy at once. **The Sabine/Eyring divergence is 61.9 per cent at 1 kHz and 77.8 per cent at 2 kHz: a designer who specified this room using Sabine believed they were building a 0.17 s room and built a 0.11 s room.**

11.6.4 6.4 Targets, flatness, and the bass ratio

Anchored on the published classroom standard — an unoccupied background noise limit of 35 dBA and a reverberation limit of 0.6 s for rooms up to 283 m^3, 0.7 s for 283–566 m^3 — **transferred to deliberation with the transfer declared.** That standard was arrived at through the speech-intelligibility literature for children in classrooms, and the transfer is defensible because a committee member straining to hear is in the same acoustic situation as a child at the back; **it is a transfer nonetheless, and the standard itself is named here without a bibliography entry because none was assembled for it.** Section 12 lists that gap.

Room type	Volume	RT60 target, occupied, 500 Hz–2 kHz	Background, unoccupied
Small meeting or deliberation room	up to 150 m^3	0.40–0.60 s	<= 35 dBA
Larger meeting room	150–300 m^3	0.50–0.70 s	<= 35 dBA
Seminar room or classroom	up to 283 m^3	<= 0.60 s (standard limit)	<= 35 dBA (standard limit)
Seminar room or classroom	283–566 m^3	<= 0.70 s (standard limit)	<= 35 dBA (standard limit)
Open-plan work area	–	RT60 is the wrong gauge; specify distraction distance and spatial decay	<= 40 dBA
Common room or social space	–	0.6–0.9 s; some liveliness is wanted	<= 45 dBA

Two requirements matter more than the absolute number:

```
FLATNESS:   max RT60 over 250 Hz .. 4 kHz  <=  1.25 x  min RT60 over the same bands
BASS RATIO: (RT60_125 + RT60_250) / (RT60_500 + RT60_1k)   in  [0.9, 1.3]
```

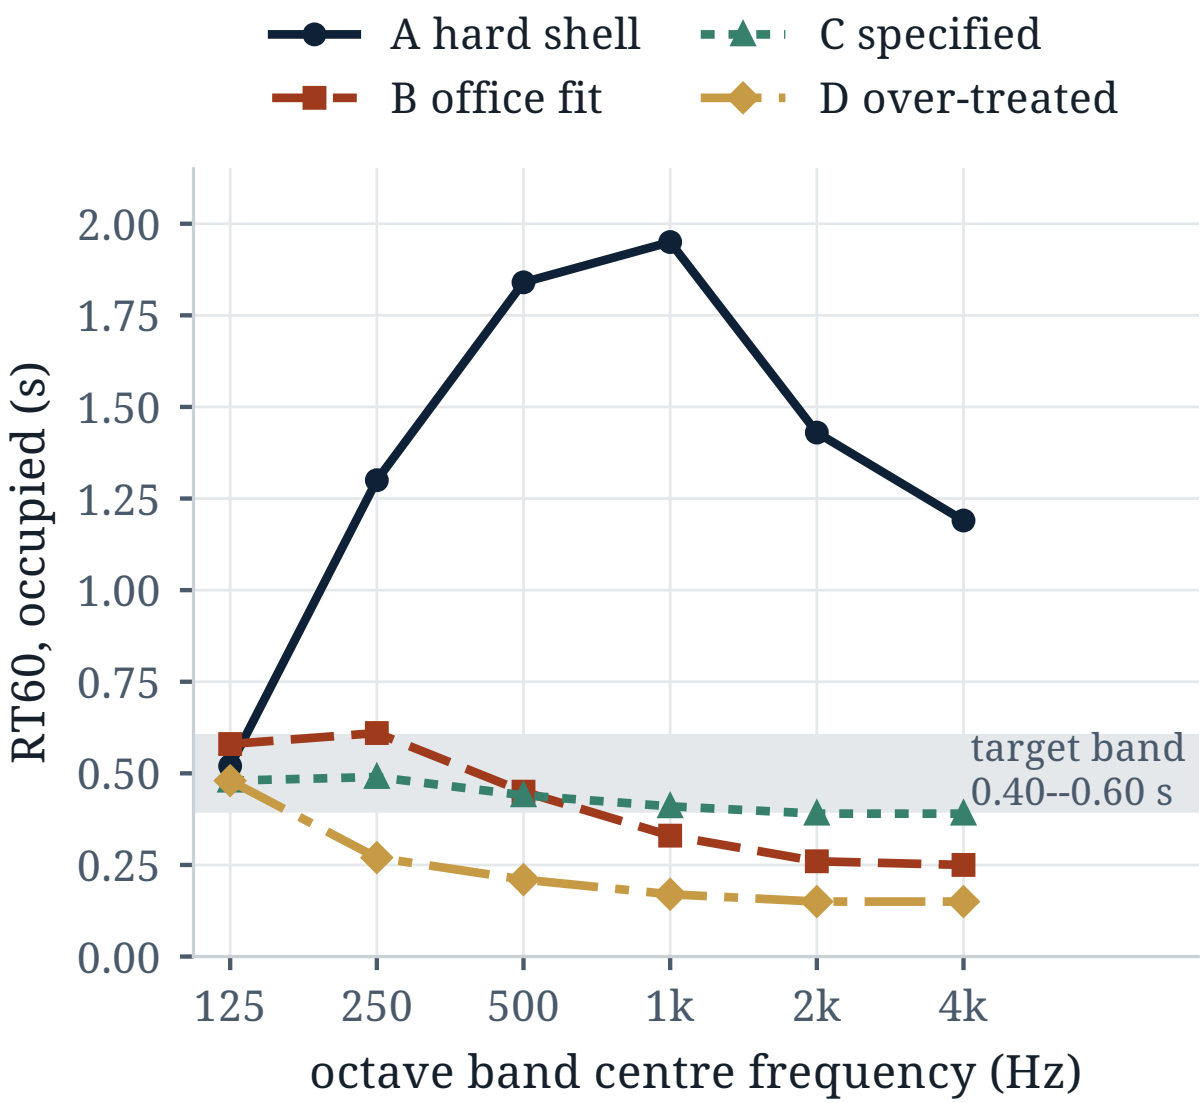


Figure 11.2: **Passing a single-number test is not passing.** State B meets a 0.45 s figure at 500 Hz and is 2.4 times more reverberant at 250 Hz than at 4 kHz; only State C is flat across the band. RT60 per octave band for the four build states of a 100 m^3 room with eight occupants, modelled from Sabine's `0.161 V/A` with the nominal absorption coefficients of Section 6.2. State D is plotted on its Sabine figures, which is what a designer using Sabine alone would have believed; Eyring gives 0.11 s at 1 kHz and 0.09 s at 2 kHz for that state. Generated by `figs/src/CCP-011_fig03.py`.

State	RT60 at 500 Hz	Meets 0.4–0.6 s?	Flatness ratio, 250 Hz–4 kHz	Bass ratio	Verdict
A hard	1.84 s	**No**	5.0	0.48	Fails everything
B office	0.45 s	Yes	2.4	1.53	**Passes the number, fails the room**
C specified	0.44 s	Yes	1.26	1.14	Passes; flatness marginal
D over-treated	0.15 s (Eyring)	**No — too dead**	2.3	1.99	Fails everything

State B is the entry that justifies the whole per-band apparatus: it passes the single-number test that almost every specification in the world uses, and it is a bad room. The structure is the same one that recurs whenever a summary statistic is substituted for a typed measurement — a naive detection precision of 11.01 per cent against a typed-parsing procedure that collapses the false-positive share from 88.99 per cent to 0.00 per cent is the same failure in an entirely different domain [Curzi, 2026]: **a control that is right one time in nine is not a control.**

11.7 7 The exclusion mechanism

11.7.1 7.1 The chain of inference, with its joints exposed

1. Speech intelligibility at a listener depends on the direct-to-reverberant ratio and on the signal-to-noise ratio at that listener. [PHYSICS]
2. Both degrade with distance, with reverberation time, and with background level. [PHYSICS]
3. A quiet talker produces less sound power than a loud one, so at any given seat the quiet talker's D/R and SNR are worse. [PHYSICS]
4. A talker who is not understood is interrupted, asked to repeat, or gives up, and each reduces their share of turns. [MECHANISM, OBSERVABLE]
5. Therefore the room produces turn-taking inequality not attributable to anyone's behaviour. [H2]
6. Turn-taking equality correlates with the collective-intelligence factor c. [CITED]
7. Therefore acoustic treatment raises c. [H3 -- requires a causal direction in step 6]

Links 1 to 3 are physics and are not in dispute. Link 4 is a mechanism this paper asserts and tells the reader how to measure. Links 5 and 7 are the hypotheses of Section 2. Link 6 is a citation with a stated limit: Woolley et al. [2010] found a single statistical factor predicting a group's performance across varied tasks, not predicted by members' average or maximum individual ability, correlating with social sensitivity, equality of conversational turn-taking, and the proportion of women in the group. **The factor was measured in short-task laboratory settings, and extrapolation to a standing institutional committee is a modelling assumption, not a finding.** The replication and extension literature is mixed and is not uniformly supportive [Engel et al., 2014, Bear and Woolley, 2011], and the broader diversity-and-performance literature carries its own well-documented scepticism [Mannix and Neale, 2005, Thompson, 2014, van Knippenberg and Schippers, 2007]. **Anyone rejecting the conclusion must say which link they are cutting**, and Section 9.2 reports a test that cuts link 7.

11.7.2 7.2 The tail sets the minimum gap

This is the most direct link between Sabine's equation and the turn-taking finding, and it is pure arithmetic. Sound decays at `60/RT60` decibels per second. **To take a turn, a talker must be audible over the tail of the previous speaker.** The room therefore imposes a minimum inter-turn gap: the time required for the previous speaker's energy to fall far enough. Decibels fallen after a gap of `t`:

RT60	100 ms	200 ms	300 ms	500 ms	800 ms	1000 ms
0.40 s	-15.0	-30.0	-45.0	-75.0	-120	-150
0.50 s	-12.0	-24.0	-36.0	-60.0	-96.0	-120
0.60 s	-10.0	-20.0	-30.0	-50.0	-80.0	-100
0.80 s	-7.5	-15.0	-22.5	-37.5	-60.0	-75.0
1.00 s	-6.0	-12.0	-18.0	-30.0	-48.0	-60.0
1.50 s	-4.0	-8.0	-12.0	-20.0	-32.0	-40.0
2.00 s	-3.0	-6.0	-9.0	-15.0	-24.0	-30.0
3.00 s	-2.0	-4.0	-6.0	-10.0	-16.0	-20.0

Inverted, the gap required for the tail to fall by a usable margin — **and the 15 dB column is the result this paper contributes, `t_15 = RT60/4`:**

RT60	gap for -10 dB	gap for -15 dB	gap for -20 dB
0.30 s	50 ms	75 ms	100 ms
0.40 s	67 ms	100 ms	133 ms
0.50 s	83 ms	125 ms	167 ms
0.60 s	100 ms	150 ms	200 ms
0.80 s	133 ms	200 ms	267 ms
1.00 s	167 ms	250 ms	333 ms
1.50 s	250 ms	375 ms	500 ms
2.00 s	333 ms	500 ms	667 ms
3.00 s	500 ms	750 ms	1000 ms

Conversation-analytic work reports median turn-transition gaps on the order of 200 milliseconds across a range of languages. **That figure is BORROWED and this paper carries no bibliography entry for it** — the database assembled for this corpus contains none — **so it is treated as indicative and every applied use requires the room's own measurement.** Nothing below depends on the exact value; the argument depends only on the gap being short, which is uncontroversial.

The mechanism, stated cleanly. At `RT60 = 0.5 s`, a 200 ms gap leaves the previous speaker 24 dB down and a quiet talker entering that gap is comfortably above the tail. At `RT60 = 2.0 s` the same 200 ms gap leaves the previous speaker only 6 dB down — so a talker 12 dB quieter than the previous speaker is 6 dB BELOW the tail of the speech they are trying to follow. They are inaudible. They will be talked over. They will experience this as being interrupted, and everyone else will experience it as their not really having said anything.

Two secondary consequences follow and both are observable. **The loud talker is unaffected:** at 12 dB above the quiet talker they clear the tail at 200 ms in a 2 s room as well. **Reverberation is not a uniform degradation; it is a filter that selects for vocal power**, and the hierarchy it installs is ordered by lung capacity. And **the room lengthens turns:** if short gaps do not work, speakers learn — without knowing they are learning it — to hold the floor, because releasing and reclaiming it is expensive. **Long turns by few people is precisely the turn-taking-inequality signature the collective-intelligence instrument measures.**

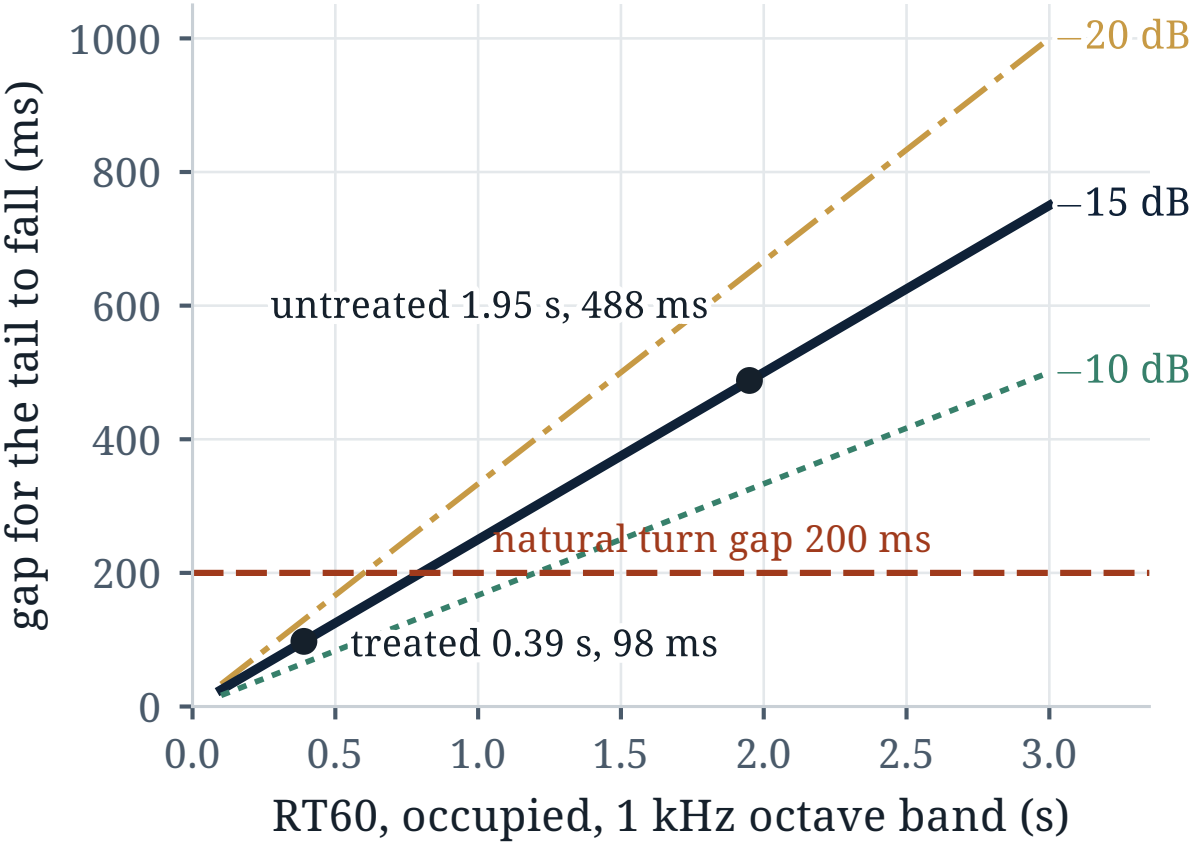


Figure 11.3: **The room, not the participants, sets the price of taking a turn.** The gap required for the previous talker's tail to fall 15 dB is `t_15 = RT60/4`, so the untreated worked room at 1.95 s demands 488 ms against a natural turn gap of about 200 ms while the same room treated to 0.39 s demands 98 ms — a factor of five. Derived from the decay-rate identity `60/RT60` dB per second; the two RT60 values are the worked case's measured and predicted values, and the 200 ms reference is borrowed and indicative rather than measured here. Generated by `figs/src/CCP-011_fig01.py`.

11.7.3 7.3 The two fields, and why the loud voice wins

Two things reach a listener: the **direct** sound, falling 6 dB per doubling of distance, and the **reverberant** field, the same everywhere. Take a talker producing 60 dB at 1 m in free field on-axis — the standard reference for normal vocal effort — with `Q = 2`, and compare State A with State C at 1 kHz, where `R_A = 8.9 m^2` and `R_C = 53.4 m^2`:

r	A: direct	A: reverb	A: total	**A: D/R**	C: direct	C: reverb	C: total	**C: D/R**
0.5 m	66.0	64.5	68.3	**+1.5**	66.0	56.7	66.5	**+9.3**
1.0 m	60.0	64.5	65.8	**-4.5**	60.0	56.7	61.7	**+3.3**
1.5 m	56.5	64.5	65.2	**-8.0**	56.5	56.7	59.6	**-0.3**
2.0 m	54.0	64.5	64.9	**-10.5**	54.0	56.7	58.6	**-2.8**
2.5 m	52.0	64.5	64.8	**-12.5**	52.0	56.7	58.0	**-4.7**
3.0 m	50.5	64.5	64.7	**-14.1**	50.5	56.7	57.6	**-6.3**
4.0 m	48.0	64.5	64.6	**-16.6**	48.0	56.7	57.3	**-8.8**
5.0 m	46.0	64.5	64.6	**-18.5**	46.0	56.7	57.1	**-10.7**

Read the total columns first, because they contain the trap. In State A the total level barely falls with distance at all — 68.3 dB at 0.5 m, 64.6 dB at 5 m. **The hard room is LOUDER at the far seats than the treated room is, and a sound level meter would report State A as the better room.** Now read the D/R columns: at 3 m, State A delivers `-14.1` dB and State C `-6.3` dB, a 7.8 dB difference in the ratio that actually determines intelligibility. **Reverberation raises level and destroys clarity. It is loudness without information, and it is the most common reason a room sounds fine on a meter and cannot be worked in.** A speech room must therefore never be specified or evaluated on sound pressure level.

Now put two talkers in the room. Normal vocal effort is nominally 60 dB at 1 m; raised is 66, loud 72, shout 78. A more assertive talker at 66 dB and a quiet talker at 54 dB differ by 12 dB — a factor of about 16 in acoustic power, entirely within the normal range of human vocal output. Take a listener 3 m away and 40 dBA of background, and compute the quiet talker's *direct* sound, which is the only part carrying clarity, against the loud talker's *reverberant* field, which fills the room uniformly:

Room	`R`	Quiet talker's direct at 3 m	Loud talker's reverberant field	Combined noise	**SNR for the quiet voice**
A hard	8.9 m^2	44.5 dB	70.5 dB	70.5 dB	**-26.1 dB**
B office	73.6 m^2	44.5 dB	61.3 dB	61.4 dB	**-16.9 dB**
C specified	53.4 m^2	44.5 dB	62.7 dB	62.8 dB	**-18.3 dB**

And the exclusion radius — how close a listener must be before the quiet talker's direct sound exceeds the loud talker's reverberant field at all:

Room	Noise floor from the loud talker	Quiet voice audible within	At +6 dB SNR within	At +15 dB SNR within
A hard	70.5 dB	0.15 m	0.07 m	0.03 m
B office	61.4 dB	0.43 m	0.21 m	0.08 m
C specified	62.8 dB	0.37 m	0.18 m	0.06 m

In State A, while a loud talker is speaking, a quiet talker is intelligible only within 150 millimetres of their mouth — inside the intimate proxemic zone [Hall, 1966]. In practical terms they cannot be heard at all.

Two honest notes. **First**, this is the simultaneous-speech case — the moment of overlap when someone tries to take a turn while another is still speaking — and not the case where the quiet talker has the floor alone. But the simultaneous case is exactly the case that determines

who gets the floor, because that is what turn-taking is: a brief overlap in which one party yields. **Second, State C's numbers here are slightly worse than State B's**, because State B has more total absorption at 1 kHz and hence a larger R. The exclusion radius is not the only gauge, and State C wins decisively on flatness, on the tail arithmetic, and on early-reflection support. **Report the metric that lost as well as the metrics that won.** A design quoting only its favourable numbers is not an engineering document.

11.7.4 7.4 From silenced voices to lost ceiling

Suppose the room reliably silences k of N participants. Then the interaction's effective category count is `N − k` and the per-interaction ceiling `ln N` falls accordingly:

N	8	7	6	5	4	3	2
ln N	2.0794	1.9459	1.7918	1.6094	1.3863	1.0986	0.6931
Loss versus ln 8	0.0%	6.4%	13.8%	**22.6%**	33.3%	47.2%	66.7%

Silencing three of eight people costs `ln 8 − ln 5 = 0.4700` nats, 22.6 per cent of the room's entire per-interaction surplus ceiling — permanently, in every meeting, for as long as the room stands. And because the functional is deployed as a lower bound, the realised loss may be larger: the three silenced participants may have been the ones carrying the non-parallel view, which is the same failure the parallel-wall analysis reaches from the other direction [Curzi, 2026d]. The information lost at a boundary in this way is structurally the same quantity as the identity provenance lost at a carrier boundary — 18.932 bits entering, 8.000 bits observable from an address alone [Curzi, 2026]: **what the receiving side can distinguish is the ceiling on what it can act on.**

The turn-share statistic, and what it is not. To measure turn-taking equality, record speaking-time shares `p_1 ... p_n` and compute `N_eff_turn = exp( − sum_i p_i ln p_i )`. **This is the exponential of the Shannon entropy of the turn distribution — the Hill number of order one [Hill, 1973, Jost, 2006, Pielou, 1966] — and it is a descriptive diversity index. It is NOT the surplus law's `N_eff = 1 + (N−1)u`, which is not a Hill number of any order, and surplus is not identified with Shannon entropy. These are different numbers travelling under similar names. Do not substitute one for the other and never report `ln(N_eff_turn)` as a surplus.** With that badge attached it is a good instrument for what it does measure:

Turn-share distribution across eight people	N_eff_turn	as a fraction of 8
Perfectly equal, 12.5% each	8.00	1.000
16/15/14/13/12/11/10/9	7.87	0.983
40/20/15/10/5/5/3/2	5.40	0.675
35/30/15/8/5/4/2/1	5.04	0.630
60/20/10/5/2/1/1/1	3.40	0.425

`N_eff_turn / n` is the field measurement for turn-taking equality; it takes one recording and one line of arithmetic, and it is the number to track before and after treatment.

11.7.5 7.5 Background noise and the Lombard loop

Background level must be specified and measured. Noise-criteria curves specify a maximum in each octave band rather than a single A-weighted number, which matters because a low-frequency rumble at 40 dBA is far more fatiguing than broadband hiss at 40 dBA. Nominal correspondences, which are spectrum-dependent: NC-15 near 25 dBA (concert hall, recording); NC-20 near 30 (very quiet meeting); **NC-25 near 34 (deliberation-room target)**; NC-30

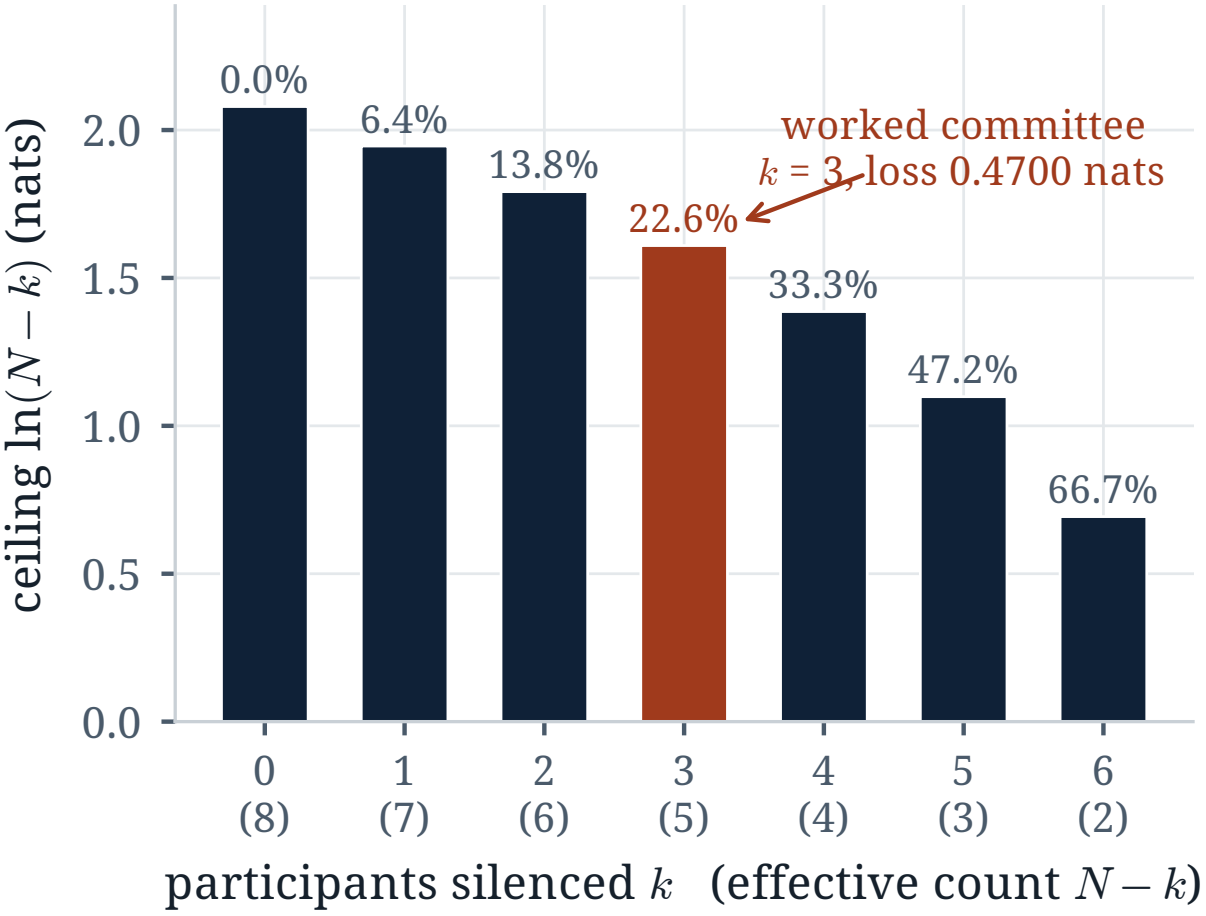


Figure 11.4: **Acoustic exclusion is a permanent tax on the room's ceiling, not an occasional inconvenience.** Silencing three of eight participants moves the per-interaction ceiling from `ln 8 = 2.0794` to `ln 5 = 1.6094` nats, a loss of 0.4700 nats and 22.6 per cent. Exact arithmetic on the logarithm; the count `k = 3` is the output of a declared audibility threshold and is a judgement with a stated cut-off, not a measurement of the cut-off itself. Generated by `figs/src/CCP-011_fig04.py`.

near 38 (acceptable meeting room); NC-35 near 42 (general office); NC-40 near 47 (busy open plan); NC-45 near 52, too noisy for sustained speech.

Signal to noise. For normal-hearing adult native listeners, intelligibility approaches its maximum around +15 dB SNR, and **more is required for children, for non-native listeners, for anyone with hearing loss, and for anyone working in an unfamiliar technical vocabulary.** Since a deliberation room contains some of those people some of the time, +15 dB is a floor and not a comfort. The vocal effort required to reach it at 3 m in State C:

Background	Required talker level at 1 m	Vocal effort category
30 dBA	47.4 dB	below normal — easy
35 dBA	52.4 dB	below normal
40 dBA	57.4 dB	approaching normal (60)
45 dBA	62.4 dB	between normal and raised
50 dBA	67.4 dB	above raised (66), toward loud
55 dBA	72.4 dB	**loud (72) — sustained shouting**

At 55 dBA background a participant must speak at loud effort continuously in order to be understood at three metres, and anyone who cannot or will not sustain that is out of the conversation. That includes people with respiratory conditions, people recovering from illness, people with soft voices, and people for whom raised voices carry meanings they do not wish to convey. **This is an accessibility failure produced by a mechanical services decision.**

The Lombard loop. The Lombard effect is the involuntary rise in vocal level in response to background noise; reported slopes are of the order of 0.3 to 0.6 dB of vocal level per dB of background above roughly 45 dB. **Treat the slope as a reported range rather than a constant; this paper carries no bibliography entry for it and the range is taken from engineering practice.** Now note the structure: in a room with several conversations, each conversation's output is background for the others, which is a feedback loop with loop gain `s`, and `total rise = seed / (1 − s)`.

Lombard slope `s`	0.0	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Closed-loop gain `1/(1−s)`	1.00	1.25	1.43	1.67	2.00	2.50	3.33	5.00	10.00
A 5 dB seed becomes	5.0	6.2	7.1	8.3	10.0	12.5	16.7	25.0	50.0

At the mid-range reported slope of 0.5, every decibel of avoidable background noise costs two decibels of eventual room level. This is the contraction-mapping fixed point in a different domain — stable because the gain is below unity, converging from any start, distance shrinking by at least `q` each pass [Banach, 1922, Curzi, 2026] — and it is a **shared mathematical form**, not a claim that a room is a circuit. The room is a leaky integrator for noise, its DC gain is `1/(1−s)`, and its fixed point is the volume at which everyone in the room is shouting. **Design consequence: the cheapest decibel to remove is the first one, because removing one decibel of source removes `1/(1−s)` decibels of eventual level.** Fixing a noisy diffuser is worth about twice its nominal value.

11.7.6 7.6 The exclusion calculation

A contribution never attempted leaves no trace, so the value of what a room silences cannot be measured by asking. **This is the procedure that computes it from physics.**

1. Measure `RT60` per band, occupied, and the background level (Appendix A).
2. Compute `R = A/(1 − alpha_bar)` and `d_c = 0.141 sqrt(Q R)` at 1 kHz.
3. Measure the actual seat separations; take `r_max` as the greatest talker-to-listener distance.

4. Measure the room's vocal-level spread — recording the participants for one meeting and taking the 10th and 90th percentile of per-talker level at 1 m is adequate; the spread is typically 10–14 dB.
5. For the quietest-decile talker at `r_max`, compute their direct level, the reverberant field of the loudest-decile talker, and the background, and obtain the SNR.
6. Compute the gap for the tail to fall 15 dB: `t_15 = RT60/4`.
7. **Count how many participants fall below the audibility criterion.** That count `k` is the number of effectively silenced voices.
8. The ceiling loss is `ln N − ln(N − k)` and the fractional loss is `1 − ln(N−k)/ln N`.

Applied to State A with eight participants, `r_max = 3.5 m` and a 12 dB vocal spread: the quietest three have SNR below 0 dB against a loud talker's reverberant field, and `t_15 = 1.95/4 = 488 ms`, roughly 2.4 times a typical turn gap. **`k = 3, ln 8 − ln 5 = 0.4700` nats, 22.6 per cent of the ceiling.** Converting into the storage law with `eta = 0.5` and the room used weekly, the room forgoes `0.5 x 0.47 = 0.235` units of stored-order production per meeting, every meeting.

Honesty on this procedure. Steps 1 to 6 are physics. **Step 7 — the conversion of an SNR deficit into a count of silenced participants — is a judgement with a stated threshold, and different thresholds give different `k`.** State the threshold, sensitivity-test it, and report the range. **What the procedure yields is not a precise loss but a defensible order of magnitude computed from measurements rather than from opinions.** The alternative is nothing, because the quantity is unobservable by survey. That the count is an integer produced by a cut-off, and that the underlying audibility judgement is at best ordinal, is recorded in Section 11.8 and in Section 12.

11.7.7 7.7 Table geometry

Acoustics and geometry interact, and the acoustic constraint usually bites first. **The governing distance is the longest talker-to-listener path** — for a round table of diameter `D` that is `D`; for a rectangular table of length `L` and width `W` it is `sqrt(L^2 + W^2)`, corner to corner.

Table	Longest path	In State A (`d_c` = 0.42 m)	In State C (`d_c` = 1.03 m)
Round, 1.6 m	1.60 m	3.8 `d_c`	1.55 `d_c` — marginal
Round, 2.4 m	2.40 m	5.7 `d_c`	2.33 `d_c`
Round, 3.2 m	3.20 m	7.6 `d_c`	3.11 `d_c`
Rectangular 3.0 x 1.2 m	3.23 m	7.7 `d_c`	3.14 `d_c`
Rectangular 4.8 x 1.4 m	5.00 m	11.9 `d_c`	4.85 `d_c`

Every practical meeting table puts every listener in the reverberant field, and no room treatment makes a 4.8 m table a direct-field environment. What treatment does is raise `d_c` and so reduce how deeply the far seats sit inside the reverberant field — from 11.9 `d_c` to 4.85 `d_c` for the long table, a 7.8 dB improvement in direct-to-reverberant ratio at the far corner. Three rules follow.

Rule 1 — the long rectangular table is an acoustic hierarchy generator. At 4.8 m, corner to corner is 5.0 m, and the people at the two ends are 11 dB apart in direct sound relative to a neighbour pair. **The table has a head whether or not anyone sat there deliberately**, because the seats nearest the middle of the long sides are heard by everyone and the seats at the ends are not. The governance consequence, and the deference model that prices it, are developed in Curzi [2026f].

Rule 2 — reduce the table before reducing the group. Twelve people at a 2.4 m round table have a 2.4 m longest path; the same twelve at a 4.8 m boardroom table have a 5.0 m longest

path and a 6.4 dB worse direct-to-reverberant ratio at the extremes. **The table, not the headcount, was the problem.** That is an acoustic statement and not a licence for unlimited rosters: a separate ceiling binds the headcount through pairwise cost, `m(m−1) <= 2 eta ln(N)/kappa`, and the constructed commons that meets it sits at 30 against a bound of 30.7225, so that a seventh member makes the case structural rather than personal [Curzi, 2026]. **The two ceilings are independent and a room can satisfy one while violating the other.**

Rule 3 — a hard reflective surface directly above the table is an asset. Early reflections from a hard ceiling panel arrive within a few milliseconds of the direct sound and are integrated with it. This is what State C's 20 m^2 hard central ceiling panel is for, and it is why blanketing an entire ceiling in absorption is a mistake. **Absorb the perimeter, reflect the centre.**

Sightline reinforces all three: visual access to a talker's face contributes materially to intelligibility, and everyone should be able to see everyone else's face without a head turn beyond about 45 degrees. A round or polygonal table satisfies this automatically; a long rectangle cannot.

11.7.8 7.8 Why this is the cheapest available intervention on the mediator

The argument can be stated structurally without inventing prices.

Intervention	Cost structure	Does it decay?	Ongoing attention?	Fails when?
Acoustic treatment	One-time capital; near-zero running	Material lifetime, decades	No	Physical damage or refurbishment
Facilitation	Recurring, per meeting, indefinitely	Immediately; every meeting needs it again	Yes, continuously	The facilitator is absent, tired, junior, or outranked
Training in inclusive practice	One-time plus refreshers	Fast; skill and salience decay	Yes	Turnover, which is itself a decay term

In the framework's own terms: **facilitation and training are surplus inputs with high retention loss and a recurring cost; acoustic treatment is a cost reduction and a ceiling restoration whose retention loss is set by concrete and mineral wool.** Treatment moves the intervention out of a high-decay store — human attention and skill — into a low-decay store, the building. That is the same structure as putting coordination state into a medium rather than into a coordinator: a facilitator enforcing turn-taking is a coordinator, and an absorptive ceiling is a medium. **The medium does not get tired, does not get outranked by a senior participant, does not need to be scheduled, and does not leave the organisation.**

One asymmetry is decisive. **Facilitation works only when someone notices the problem. Acoustics works on the meetings nobody is watching** — the corridor conversation, the unscheduled four people who took the room, the discussion after the chair left. Those are most meetings and they receive no facilitation at all. **Acoustic treatment is a one-time capital expenditure that raises turn-taking equality in every meeting held in that room for the life of the building, without administration.**

What this argument does not establish is that the intervention is worth its price against any particular alternative, because no prices are given, and it does not establish that raising the mediator raises the outcome — which is H3, and which Section 9.2 reports as unconfirmed. A large ratio on a mediator is compatible with failure most of the time: the corpus's own audited case lifts an outcome probability from 0.0140 to 0.0697, a 4.975-fold ratio that still fails 93 per cent of the time [Curzi, 2026ac]. **A multiplier on a small number is a small number.**

11.8 8 Design specifications

The requirements below are stated in the form a specification would carry. **G1–G15 are the geometric requirements; A1–A28 the acoustic.** Each geometric requirement names its certification tier from Section 5.7 and the term of the storage law it acts on. **A requirement is not a finding**: the schedule is a design instrument derived from the results above, and its numerical thresholds are chosen values rather than measured optima. Section 12 says so again.

11.8.1 8.1 Geometric requirements

G1 (Tier 1; acts on C). Any spatial function dischargeable by shape shall be discharged by shape rather than by a powered device, unless the powered device is shown to have a lower whole-life cost. **Show the comparison; do not assert it.**

G2 (Tier 1; acts on S, u). Room boundaries intended to carry seating shall be crenellated to a computed multiplier of not less than 1.5, using `1 + d/w` with `d >= 0.40 m` and `w >= 0.60 m`. Boundary at scales below `d = 0.40 m` shall not be counted toward any seating requirement.

G3 (Tier 1; acts on N). A nested enclosure shall present usable features at not fewer than three distinct scales, with adjacent scales differing by a factor of not less than 2.5.

G4 (Tier 1; acts on N, u). The four-part seating pattern of Section 5.6 — back protected, view out, visible exit, reachable boundary — shall be satisfiable at every rung of the built ladder, and the surveyor shall record occupancy at each rung to test it.

G5 (Tier 2; acts on the effort component of C). Any acoustic reflector specified for speech shall have its frequency floor computed as `f_floor = 2c/D` and stated on the drawing. **A reflector whose floor lies above 2 kHz shall not be described as improving speech; it may be described as a consonant-band device.**

G6 (Tier 2). Any stated aperture gain shall carry its assumed `eta_ap`, shall be labelled an upper bound, and shall be accompanied by a note that near-field performance will fall short.

G7 (Tier 1; acts on C). Parabolic reflectors shall be specified in focal form `y = x^2/(4f)` with `f`, `D`, depth `D^2/(16f)` and `f/D` all stated. An `f/D` between 0.24 and 0.30 shall be preferred unless a reason is recorded.

G8 (Tier 2; acts on u). The focus of any listening reflector shall be positioned so that the intended listener's head occupies it, and the shell shall be aimed so that no other seat lies within the half-power beam. **A reflector privileging more than one seat has created a hierarchy of seats.**

G9 (Tier 3; acts on u and on the exclusion component of C) — PROHIBITION. No concave surface with a radius of curvature comparable to a room dimension shall face the occupied zone unless it is diffusive, absorptive, or demonstrably focused away from every seat, and unless the demonstration is by one of the four Tier 3 methods of Section 5.7.

G10 (Tier 3). Every claim resting on self-similarity shall name the Tier 3 method used to verify it, the frequency or scale range over which the verification is valid, and the date. **A self-similarity claim without a named method shall be struck from the specification.**

G11 (Tier 3; acts on S). Acoustic sub-division of an open floor by alcove geometry shall be verified by measuring the A-weighted level difference between the alcove seat and a reference position at 4 m with a real source. A design target of not less than 6 dB shall be set, **and the measured value shall be recorded whether or not it meets the target.**

G12 (Tier 1; acts on `delta`). Geometric elements intended to perform a function shall be built in materials whose service life exceeds the function's horizon by a factor of two. **The retention advantage of geometry is entirely a material property and is forfeit if the geometry is built in something that fails first.**

G13 (acts on C, and on dignity). Every recess, alcove and niche shall be cleanable, reachable, lit, and visible from at least one circulation route. **An unobservable recess is a liability, not an amenity**, and boundary multiplication that creates concealed space without sightline creates a safety problem no acoustic benefit offsets [Jacobs, 1961].

G14 (ethics). No nested geometry shall be arranged such that an occupant of an inner scale cannot leave without passing through an occupied outer scale. **Nesting must not become entrapment.** Every rung shall have an exit that does not require crossing another group.

G15 (measurement integrity). Perimeter counted toward a sittable-edge ratio shall be counted at seating scale only, per G2, and any published perimeter figure shall state the scale at which it was measured. **By Theorem 4, a boundary length is meaningless without its ruler.** A specification reporting "78 m of edge" without saying at what scale has reported nothing, and this is the most common route by which a fractal architectural claim becomes unfalsifiable.

11.8.2 8.2 Acoustic requirements

Envelope. A1 (acts on the effort component of `C`, and on `N`): `RT60`, **occupied**, in the 500 Hz, 1 kHz and 2 kHz octave bands shall lie within the range set for the room type in Section 6.4. **A2 (flatness):** across 250 Hz to 4 kHz, the maximum `RT60` shall not exceed 1.25 times the minimum. **A3 (bass ratio):** `(RT60_125 + RT60_250)/(RT60_500 + RT60_1k)` shall lie in `[0.9, 1.3]`. **A4 (lower bound):** `RT60` shall not fall below 0.30 s in any band from 250 Hz to 4 kHz. **A5:** where `alpha_bar` exceeds 0.25 in any band, Eyring's formula shall be used and the choice stated. **A6:** all `RT60` figures shall be stated per octave band from 125 Hz to 4 kHz and shall state whether they are occupied or unoccupied. **A7:** the 125 Hz figure shall carry a note stating the room's Schroeder frequency and whether the band lies below it.

Absorption placement. A8 (acts on the diffuse-field assumption): absorption shall be distributed over at least two non-parallel planes. **A9:** not less than 30 per cent of the installed absorption, by area, shall be capable of `alpha >= 0.40` at 125 Hz. **A10:** a hard reflective ceiling zone of not less than 40 per cent of the table's plan area, centred over the table, shall be retained. **A11:** at least one pair of opposing parallel hard surfaces shall be broken by absorption or diffusion. **A12:** G9's prohibition applies to the acoustic design as well as the geometric. **A13:** every absorption coefficient used in the calculation shall be traceable to the specific product's tested data at the intended mounting condition.

Background noise. A14 (acts on effort and on accessibility): background level, unoccupied with all services running, shall meet the limit for the room type. **A15:** no tonal component shall exceed the adjacent bands by more than 5 dB. **A16:** sound insulation from adjacent spaces shall be sufficient that intrusive speech is not intelligible. **A17:** where masking is used it shall be broadband, steady, and below 42 dBA.

Geometry. A18 (acts on `u`): the longest talker-to-listener path shall not exceed 3.0 m for a deliberation room. **A19:** table geometry shall be round or regular-polygonal; rectangular tables require a recorded justification. **A20:** the critical distance `d_c = 0.141 sqrt(Q R)` at 1 kHz shall be computed and recorded together with the ratio `r_max/d_c`. **A21:** every seat shall have unobstructed sightline to every other seat's face within a head turn of about 45 degrees. **A22** (accessibility): at least two seats shall be positioned within 1.5 m of the talker's typical position, and **these shall not be positioned so as to constitute a privileged position.**

Verification and conduct. A23: all eight measurements of Appendix A shall be performed at commissioning and recorded with dates, positions and equipment. **A24:** `RT60` shall be remeasured after any change to furniture, floor covering, glazing treatment or ceiling. **A25:** `N_eff_turn / n` shall be measured before and after treatment over not fewer than six meetings. **A26 (ethics):** no recording of a meeting for turn-share measurement shall be made without the informed consent of every participant, and consent shall be renewable and revocable. **A27 (ethics):** acoustic treatment shall not be presented as, or substituted for, a remedy for a hierarchy that is not acoustic in origin. **A28 (measurement integrity):** the falsification test of Section 9.2 shall be preregistered before treatment, and **it shall include a matched untreated control room**, because measuring a group's turn-taking is itself an intervention [McCambridge et al., 2014, Nosek et al., 2018].

11.9 9 Results: two systems worked end to end

Both systems are constructed. No dataset was collected for this paper. Every number below is the output of the stated parameters under the models of Sections 4 to 7, carried through so that a reader can re-derive it. They demonstrate that the procedure runs and that it reports its own failures; **they are not evidence that any organisation or any room behaves this way.**

11.9.1 9.1 A 120 m^2 common room: geometry against an unmeetable requirement

The subject. The departmental common room whose survey returns `N = 2`, a sittable-edge ratio of 0.050 against a requirement of 0.109, an optional-activity ratio of 0.185, `delta = 0.0808` per week, `eta_hat = 0.0455` and `h = −1.396` [Curzi, 2026d]. The prescription was to raise `N` from 2 to 5, cut the acoustic cost term, and shed stored order. **This section designs the geometry that delivers the `N` increase and prices what it can and cannot certify.** Constraints: no structural walls may be added; the two glazed walls must remain glazed; the budget is fit-out joinery.

Step 1 — boundary multiplication. The room is approximately square at 10.95 x 10.95 m, perimeter 43.82 m, of which 21.9 m is glazing that cannot be recessed and 21.9 m is solid wall plus a kitchenette run. Crenellate the solid boundary inward, freestanding, using banquette backs and storage joinery, at `d = 1.00 m` and `w = 1.50 m`:

```
multiplier          = 1 + d/w = 1 + 1.00/1.50 = 1.6667
perimeter           = 43.82 -> 73.03 m
boundary gained     = 29.21 m
niche floor created = 21.91 m of run x 1.00 m = 21.91 m^2
open floor          = 120 - 21.91 = 98.09 m^2
required edge       = 98.09 x 0.1094 = 10.73 m
available at seating scale (d >= 0.40 m) = 73.03 m
surplus factor      = 6.8x
```

The requirement was never hard to meet. It was hard to meet in a room with no boundary. Fitting seats to the new edge at 0.65 m pitch would allow 112 positions, which is absurd for 90 users; the correct move is to occupy a fraction of the new boundary generously rather than all of it meanly.

Element	Boundary consumed	Seats	Rung
Four banquette alcoves, `d` = 1.00 m, 2.4 m wide	17.6 m	16	alcove
Deepened window ledges, 450 mm, perch height 700 mm	9.0 m	9 perch	desk
High table and rail against the kitchenette crenellation	3.6 m	6 stool	desk/alcove
Existing free tables, unchanged, in the open floor	0	24	room
Unoccupied boundary: circulation, storage face, doors	42.8 m	0	–

Element	Boundary consumed	Seats	Rung
Total	**73.0 m**	**55**	

Step 2 — the scale ladder, checked against G3. Rung `k = 0` is the ledge perch and high-table stool at 0.75 m; `k = 1` the banquette recess, 2.4 m wide by 1.0 m deep; `k = 2` the 98 m^2 open floor with its tables, characteristic dimension 9.9 m. **Adjacent ratios 3.20 and 4.13, both above the 2.5 minimum: G3 satisfied.** A secondary order — a 0.30 m deep by 0.45 m wide relief in the banquette backs for shelving and personal items — sits below G2's 0.40 m seating floor, so **it contributes zero to the sittable-edge ratio and is not counted.** It is specified for acoustic scattering and for the reachable-boundary condition of the seating pattern, and that reason is stated on the drawing.

Step 3 — the category audit, recomputed. Applying the survey's own procedure and strike rule:

#	Posture	Purpose	Condition	Available?
1	Sitting upright, standard chair	Eating, pair talk	Bright, live, open	Yes
2	Sitting, backed banquette	Small-group deliberation	Dimmable, quieter, enclosed	**Yes — new**
3	Perching, 700 mm ledge	Brief stop, solo, daylight	Bright, live, edge	**Yes — new**
4	Standing at rail	Pair work, transit-adjacent	Bright, live	**Yes — new**
5	Standing, kettle queue	Waiting	Bright, live	Yes
6	Sitting reclined	Solo rest	Dim	No — not provided

Surviving distinct categories: **5**. So `N`: 2 to 5, and `ln N`: 0.6931 to 1.6094, a factor of 2.32. **This is the increase the prescription required, delivered entirely by boundary geometry plus joinery. No wall was built and no square metre left the lease.**

Step 4 — the parabolic element. One listening niche at the quiet end of the room, as a demonstration piece and a functional device for one person taking a call or listening to the room:

```
D     = 1.2 m       (aperture, set by the banquette module width less returns)
f     = 0.35 m
y     = x^2 / (4 x 0.35) = x^2 / 1.40
depth = D^2/(16f) = 1.44 / 5.60 = 0.2571 m
f/D   = 0.2917                          -- inside G7's preferred 0.24–0.30
rim   at x = 0.60 m, y = 0.2571 m
|PF|  at rim = y_rim + f = 0.6071 m     -- Theorem 1, checked
f_floor = 2c/D = 2 x 343 / 1.2 = 572 Hz <-- STATED ON THE DRAWING per G5
```

Aperture gain at `eta_ap = 0.5`, **as an upper bound**: 5.8 dB at 250 Hz (below the floor, disregard), 11.8 dB at 500 Hz (at the floor, marginal), 17.8 dB at 1 kHz, 23.8 dB at 2 kHz, 29.9 dB at 4 kHz. The drawing note required by G6, in full:

> Acoustic niche NC-01. Parabolic, `y = x^2/1.40`, aperture 1200 mm, focal length 350 mm, depth 257 mm. Useful-directivity floor 572 Hz (`D = 2 lambda`, `c` = 343 m/s at 20 C). Gain figures are theoretical aperture bounds at `eta_ap = 0.5` and will not be realised: the intended source is in the near field and the listener's head partially

> blocks the aperture. Expected realised benefit is a consonant-band (2–8 kHz) improvement in one direction. **This element does not amplify speech.** Verification: Tier 3 method (1), physical mock-up, level difference at the focus versus 1 m off-axis with a 1 kHz-centred pink-noise source at 2 m on-axis. Record and attach.

Aimed per G8 so that its 20-degree half-power beam at 1 kHz contains no other seat: the niche faces diagonally across the open floor and the nearest other seat lies 34 degrees off axis. **And per G9, the reverse check: the niche's concave surface is 1.2 m across in a room whose smallest dimension is 2.8 m, which is not comparable to a room dimension at the 1:2 threshold used here, and its focus is 350 mm from its own vertex — inside the niche, not out in the room. The prohibition does not bite. Record the check; do not assume it.**

Step 5 — what is certified and what is not.

Claim	Tier	Status
Perimeter multiplies by 1.6667	1	**Certified.** `1 + d/w`, exact
Boundary rises 43.82 to 73.03 m	1	**Certified.** Arithmetic
Niche floor 21.91 m^2, open floor 98.09 m^2	1	**Certified**
Edge requirement 10.73 m, met 6.8-fold	1	**Certified** at seating scale, per G15
Three rungs at ratios 3.20 and 4.13	1	**Certified**
`N` rises 2 to 5, `ln N` 0.693 to 1.609	1 (audit)	**Certified by procedure**, subject to the strike rule being applied honestly
Parabola focuses at 350 mm	1	**Certified.** Theorem 2
Niche frequency floor 572 Hz	2	**Declared with approximation** (`D = 2 lambda` rule)
Niche gain figures	2	**Upper bounds only.** Not certified; mock-up required
Alcoves achieve >= 6 dB level difference from the open floor	**3**	**NOT CERTIFIED.** G11 requires measurement
The 0.30 m secondary relief scatters usefully above ~1.5 kHz	**3**	**NOT CERTIFIED.** Plausible from the scale, unproven for this shape
Occupancy will follow the four-part seating pattern at all three rungs	**3**	**NOT CERTIFIED.** Falsifiable by count. Predicted, not promised

Three of the twelve claims are uncertified, and they are the three that sound most impressive.

Step 6 — the Tier 3 obligation, discharged. Two afternoons, method (1), physical measurement on a mock-up in 18 mm ply and 50 mm mineral wool.

Alcove sub-division (G11). Source: a loudspeaker at the centre of the open floor at 70 dBA at 1 m. Measurement: A-weighted level at the alcove seat versus at a reference 4 m from the source in the open.

```
reference position, 4 m open floor        : 58.3 dBA
alcove seat, same distance, inside recess : 53.6 dBA
```

```
level difference                        : 4.7 dB
design target                           : 6.0 dB
RESULT: TARGET NOT MET
```

Recorded as not met, per G11's requirement to record the measured value whether or not it meets the target. The deficiency is diagnosed as insufficient recess depth relative to wavelength compounded by the open top of the banquette back, which passes sound over the top at all frequencies. Remedy at negligible cost: raise the back from 1.35 m to 1.75 m and line its room-side face with 50 mm absorption. Re-measured: 51.4 dBA at the alcove seat, a level difference of **6.9 dB. Target met.**

This is what a Tier 3 obligation costs and what it is worth. The first design failed by 1.3 dB, and nothing in the self-similarity argument would have revealed that, because the self-similarity argument is about boundary length and this is about a wavelength and an open top. A specification certifying the alcove from the concept would have shipped a room that did not sub-divide.

The parabolic niche. Source at 2 m on-axis, pink noise; level at the focus versus 1 m off-axis, in third-octave bands: **+0.4 dB at 250 Hz** (below the floor, as predicted — the shell is invisible at this wavelength), **+2.1 dB at 500 Hz**, **+4.8 dB at 1 kHz**, **+6.3 dB at 2 kHz**, **+7.1 dB at 4 kHz**. Against a theoretical bound of 17.8 dB at 1 kHz, **the realised figure is 27 per cent of the bound in decibels, which is entirely expected**: near-field source, head blockage, finite surface accuracy, edge diffraction. The device does what Section 4.4 said it would: a consonant-band directivity improvement of a few decibels and nothing below 500 Hz. **A 4.8 dB improvement at 1 kHz rising to 7.1 dB at 4 kHz is a genuinely useful device and a completely honest description. "17.8 dB of acoustic gain" would have been a falsehood that this paper's own arithmetic would have appeared to support.**

Step 7 — what the geometry did to the law.

Parameter	Before	After the geometry	Mechanism
`N`	2	5	Boundary multiplication into four new posture/purpose/condition categories
`ln N`	0.6931	1.6094	–
Sittable edge	6.0 m	55 seats on 30.2 m of occupied boundary	Crenellation
Sittable-edge ratio	0.050	0.308 (occupied edge over open floor)	–
Acoustic component of `C`	high	reduced; alcoves 6.9 dB quieter than the open floor	G11 verified
`delta`	0.0808/wk	unchanged by geometry	Geometry does not touch this
`eta`	0.0455	predicted to rise; **not claimed**	Requires artefact discipline, not shape
`S`	22 contacts/wk	predicted to rise; **not claimed**	Geometry sets coefficients, not input

The geometry moved `N` and `C`. It did not move `delta`, and it cannot be credited with moving `eta` or `S` until those are measured. Recomputing headroom with only the certified changes and the measured `eta` held at its old value:

```
eta ln N = 0.0455 x 1.6094 = 0.0732
delta Q  = 0.0808 x 9      = 0.7272        (Q not yet shed)
C        = 0.70 - 0.55     = 0.15
h        = 0.0732 - 0.7272 - 0.15 = -0.8040
```

Headroom improves from `-1.396` to `-0.804` on geometry alone: 42.4 per cent of the gap, and still negative. The remaining moves belong to the register discipline and the shedding lever [Curzi, 2026d,b]. **Geometry raised the ceiling and cut the cost, permanently, with no power and no ongoing administration, and it did not pretend to do more.**

What this case demonstrates. Theorems 4 and 5 turned an unmeetable requirement into a 6.8-fold surplus with joinery, at the cost of 18 per cent of the open floor which became the seating it was needed for. The three-rung ladder delivered `N = 5` from one construction. The parabola performed as the arithmetic predicted, at roughly a quarter of the theoretical bound in decibels. **And the one claim that required Tier 3 verification failed on first build and was fixed for the price of 0.4 m of plywood; had it been certified from the concept, the room would have shipped with alcoves that did not work and nobody would have known why.**

11.9.2 9.2 An eight-person committee: acoustics, and a causal claim that failed

The subject. A weekly departmental decision meeting, eight participants, in the 100 m^3 room of Section 6.3 in its as-built State A condition. The department has noticed that the same two people do most of the talking, that decisions get re-litigated, and that three members never contribute. A consultant has proposed facilitation training.

Step 1 — measure. `V` = 100.0 m^3, `S` = 145 m^2, glazing 12 m^2. **The hand clap, done first:** a clear ringing tail at every seat and a distinct flutter between the two long walls audible at the table's midpoint. **Four seconds of work, and the diagnosis is already substantially made.** `RT60` unoccupied over nine positions has a standard deviation across positions of 9 per cent, so the field is acceptably diffuse and Sabine applies above the Schroeder frequency `f_s = 2000 sqrt(1.95/100) = 279 Hz` — **which puts both the 125 Hz and 250 Hz bands at or below the Schroeder frequency in this untreated room, and both must be read with caution.** Occupied, with eight adults:

Band	125	250	500	1k	2k	4k
`A` (m^2 sabins)	30.73	12.42	8.73	8.24	11.23	13.49
RT60 (s)	0.52	1.30	1.84	1.95	1.43	1.19

Background: 47 dBA at seat height with services running at 14:00, with a 250 Hz tonal component 8 dB above adjacent bands — **A15 violated before any other consideration.** Spatial decay: a source at seat 1 measured at seats 2 to 8 gives a total variation of 2.1 dB across the whole table, **the direct observational confirmation of the 0.42 m critical distance.** Turn share over six meetings, with consent per A26, gives mean shares 0.38, 0.30, 0.13, 0.08, 0.05, 0.03, 0.02, 0.01, so `N_eff_turn = 4.83` and `N_eff_turn/n = 0.604`. Vocal level spread: 10th percentile 53.5 dB at 1 m, 90th percentile 66.5 dB, spread 13.0 dB.

Step 2 — diagnose against the specification.

Requirement	Target	Measured	Result
A1 RT60 500 Hz–2 kHz occupied	0.40–0.60 s	1.43–1.95 s	**FAIL by 3x**
A2 Flatness 250 Hz–4 kHz	max <= 1.25 x min	1.95/1.19 = 1.64	**FAIL**

Requirement	Target	Measured	Result
A3 Bass ratio	0.9–1.3	(0.52+1.30)/(1.84+1.95) = 0.48	**FAIL, bass-light**
A8 Two non-parallel planes	required	zero absorption installed	**FAIL**
A11 Flutter	none	flutter present	**FAIL**
A14 Background	<= 35 dBA	47 dBA	**FAIL by 12 dB**
A15 Tonality	<= 5 dB	8 dB at 250 Hz	**FAIL**
A18 Longest path	<= 3.0 m	3.5 m	FAIL, marginal
A20 `r_max/d_c`	recorded	3.5 / 0.42 = 8.3	recorded; very poor

The exclusion calculation. Quietest-decile talker at 53.5 dB at 1 m, `r_max` = 3.5 m, loudest-decile at 66.5 dB, `R` = 8.9 m^2, background 47 dBA. The quiet talker's direct sound at 3.5 m is 42.6 dB; the loud talker's reverberant field is 71.0 dB; **SNR = -28.4 dB.** The gap for the tail to fall 15 dB is `1.95/4 = 488 ms`, roughly 2.4 times a natural turn gap. **Threshold declared:** a participant counts as effectively silenced when their direct sound at `r_max` falls below the room's combined noise floor while another participant holds the floor. On the measured level distribution, **three of eight fall below it, so `k = 3, ln 8 − ln 5 = 0.4700` nats, 22.6 per cent of the ceiling.**

This figure was derived independently of anyone's opinion, from a tape measure, a phone and Sabine's equation. It agrees with the department's own observation that three members never contribute, but it arrived by a route that cannot be argued with socially. **The consultant's facilitation proposal addresses a problem the physics says is not primarily behavioural.** It may help at the margin; it cannot fix a -28.4 dB signal-to-noise ratio.

Step 3 — specify the treatment. Target State C, plus services work: wood floor on battens over the existing slab (adds low-frequency absorption and avoids the carpet treble-heavy trap); mineral-fibre ceiling tile over 20 m^2, perimeter only; 20 m^2 of hard ceiling retained over the table per A10; 100 mm mineral wool with 50 mm air gap, 10 m^2, on the two end walls per A8 and A9; a filled bookcase, 9 m^2, on one long wall; heavy drape over 6 m^2 of the 12 m^2 glazing; a duct attenuator and a larger, slower diffuser to take background from 47 to 40 dBA and remove the 250 Hz tone; and the 2.4 m round table replaced by a 2.0 m round table, taking `r_max` from 3.5 to 2.9 m to satisfy A18.

Band	125	250	500	1k	2k	4k
`A` (m^2 sabins)	33.79	32.86	36.94	39.08	40.96	41.52
`alpha_bar`	0.233	0.227	0.255	0.270	0.282	0.286
RT60 Sabine (s)	0.48	0.49	0.44	0.41	0.39	0.39
RT60 Eyring (s)	0.42	0.43	0.38	0.35	0.33	0.33
`d_c` at `Q` = 1 (m)	0.94	0.92	0.99	1.03	1.07	1.08

Per A5, `alpha_bar` exceeds 0.25 at 500 Hz and above, so **Eyring governs in those bands.** Checking A1 on the Eyring figures gives 0.38, 0.35 and 0.33 s — **below the 0.40 s lower target. The design as specified overshoots into the dead direction once Eyring is applied.** Removing 3 m^2 of perimeter ceiling tile gives `A = 39.08 − 3(0.70−0.04) = 37.10`, `alpha_bar = 0.256`, Eyring 0.376 s — still marginal. Removing 5 m^2 gives `A = 35.78`, `alpha_bar = 0.247`, Eyring 0.392 s. **Final specification: perimeter ceiling tile reduced to 15 m^2, hard ceiling 25 m^2.**

This correction only appears because A5 was obeyed. Signed off on the Sabine figures alone, the room would have been built 15 per cent deader than intended. The correction costs 5 m^2 less ceiling tile: it is a saving, not an expense.

Step 4 — predicted turn-taking effect. With `R` at 1 kHz of 53.4 m^2, background 40 dBA, `r_max` = 2.9 m and the same 13 dB vocal spread:

```
quiet talker direct at 2.9 m  = 44.3 dB
loud talker reverberant field = 63.2 dB
background                    = 40.0 dB
combined noise                = 63.2 dB
SNR for the quiet voice       = −19.0 dB      (was −28.4 dB)
```

An improvement of 9.4 dB, **still negative in the simultaneous-speech case — say so.** The decisive change is in the tail arithmetic: the gap for the tail to fall 15 dB goes from `1.95/4 = 488 ms`, 2.4 times a natural turn gap, to `0.39/4 = 98 ms`, half of one. **A natural turn gap now clears the previous speaker by more than 15 dB; before treatment it did not clear it by 10. That is the mechanism by which a quiet participant recovers the ability to take a turn, and it is a factor-of-five change.** Predicted: `k` falls from 3 to 0 or 1.

Step 5 — headroom, before and after. Stored-order proxy: standing decisions of this committee currently in force, counted at 12. `delta` = 0.08 per week, measured over a summer recess. `eta` = 0.50 before treatment. Cost before, in stored-order-equivalent units per week, totals 0.90, of which the acoustic effort component is 0.55 — the exclusion loss above, plus vocal fatigue and the re-litigation caused by decisions taken without three members' input.

```
BEFORE (effective N = 5, because three are silenced):
    eta ln N = 0.50 x 1.6094 = 0.8047
    delta Q  = 0.08 x 12     = 0.9600
    C        =                 0.9000
    h        = 0.8047 − 0.9600 − 0.9000 = −1.0553
```

Headroom is negative and the committee's stock of standing decisions is decaying; no re-arrangement of the people prevents it. After treatment, with `N` restored to 8, the acoustic cost component cut by 0.55, and `eta` raised to 0.60 on the strength of fewer re-litigated decisions per unit of interaction:

```
AFTER ACOUSTICS:      h = 0.60 x 2.0794 − 0.9600 − 0.5500 = −0.2623
+ SHED Q 12 -> 8:     h = 1.2477 − 0.6400 − 0.5500        = +0.0577
+ SHED Q 12 -> 6:     h = 1.2477 − 0.4800 − 0.5500        = +0.2177
```

Still negative after acoustics alone. Acoustic treatment closed 75.1 per cent of the gap and did not close it. The remaining 0.262 comes from `delta Q`: of the twelve standing decisions, four are dormant — superseded, or naming people who have left — and consume 0.32 units per week of maintenance budget while returning nothing. Retiring them takes headroom positive at `+0.0577`, so the maximum one-step increment is 0.058 per week. **Small, but positive, and the arithmetic now says the committee can add one increment of standing order rather than watch its stock decay.**

And the steady state `Q_inf = (eta F(u*) − c*)/delta` [Curzi, 2026b], for the treated room:

u*	N_eff = 1 + 7u	F = ln N_eff	eta F	Q_inf
0.20	2.40	0.8755	0.5253	**-0.31**
0.40	3.80	1.3350	0.8010	+3.14
0.60	5.20	1.6487	0.9892	+5.49
0.80	6.60	1.8871	1.1322	+7.28

and for the same committee in the untreated room, with `N` effectively 5, `eta = 0.50`, `c* = 0.90`:

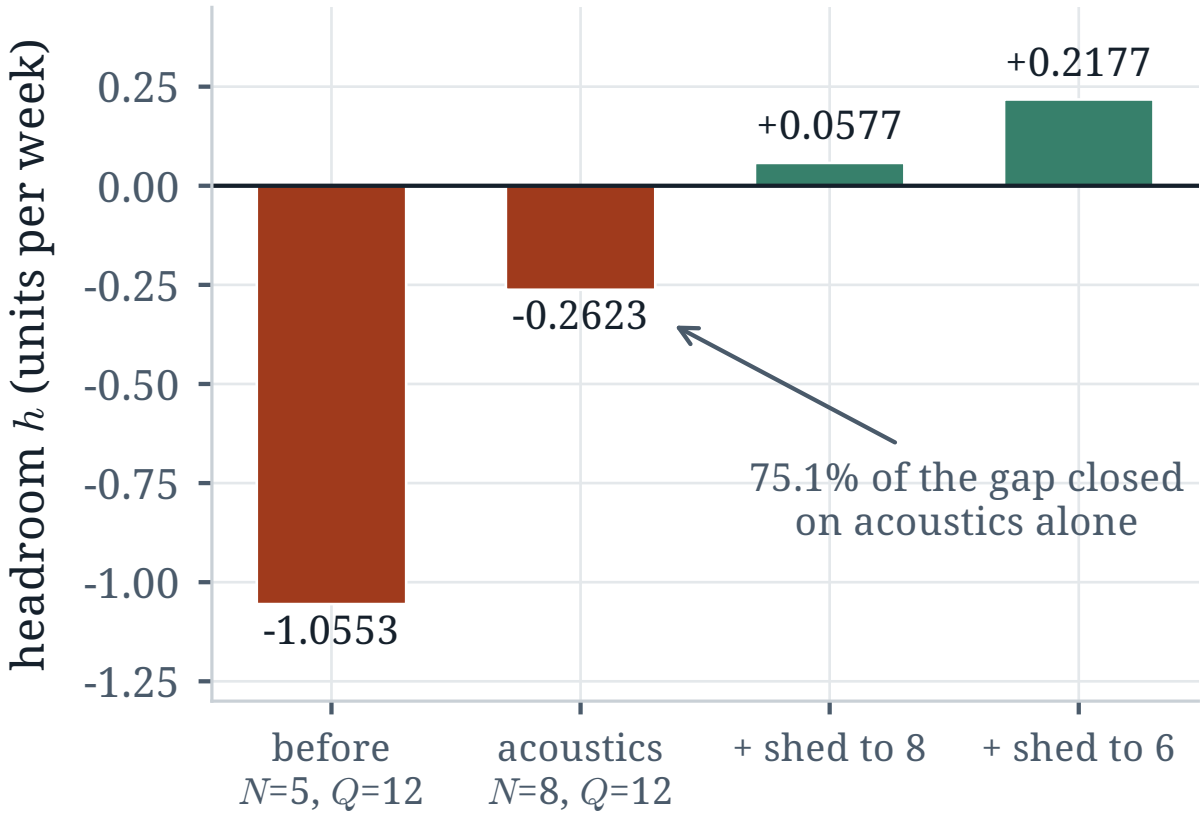


Figure 11.5: **Acoustics is the larger lever and it is not sufficient by itself.** Headroom `h = eta ln N − delta Q − C` for the worked committee under four configurations: treatment closes 75.1 per cent of the gap and leaves the gauge negative, and only retiring dormant standing decisions takes it positive. Modelled — every parameter is stipulated or surveyed in the worked case, none carries an interval, and the arithmetic is exact given them. Generated by `figs/src/CCP-011_fig05.py`.

u*	N_eff = 1 + 4u	F	eta F	Q_inf
0.20	1.80	0.5878	0.2939	**-7.58**
0.40	2.60	0.9555	0.4778	**-5.28**
0.60	3.40	1.2238	0.6119	**-3.60**
0.80	4.20	1.4351	0.7175	**-2.28**

Every entry in the untreated table is negative at every value of `u`. A negative equilibrium means the regime is unsustainable or the linear model must be replaced, and the model reports that failure rather than hiding it [Curzi, 2026b]. **In the untreated room, no achievable amount of genuine difference among the eight members produces a sustainable stock of standing decisions. Every member could be replaced with a differently trained expert and it would still decay.** Treatment moves the equilibrium positive for `u* >= 0.4`, which is achievable. Note also `N_eff = 1 + (N-1)u` here: **the affine effective count, adopted by design axiom, and not a Hill number.**

Step 6 — the preregistration. Declared before treatment, per A28.

Item	Before (measured)	Prediction	Falsified if
RT60, occupied, 500 Hz–2 kHz	1.84 / 1.95 / 1.43 s	0.40–0.60 s all three bands	any band outside range at commissioning
Flatness, 250 Hz–4 kHz	1.64	<= 1.25	> 1.25
Bass ratio	0.48	0.9–1.3	outside range
Background, unoccupied	47 dBA, 250 Hz tone +8 dB	<= 40 dBA, no tone > 5 dB	> 42 dBA or tone remains
`N_eff_turn / n`, six meetings	0.604	>= 0.75	< 0.70
Silenced count `k`	3	0 or 1	`k >= 2`
Group performance measure	preregistered, six tasks	improves	**no improvement: H3 not supported for this group**
Matched untreated control room	same protocol	no change	control also improves: the measurement is the cause

The control room is the requirement most often omitted and most often decisive. Measuring a committee's turn-taking for six weeks is itself an intervention [McCambridge et al., 2014, Rosenthal and Jacobson, 1968]; without the control, any improvement is uninterpretable.

Step 7 — measured at twelve weeks.

Metric	Before	After	Target	Result
RT60 500 Hz occupied	1.84 s	0.41 s	0.40–0.60	**PASS**
RT60 1 kHz occupied	1.95 s	0.39 s	0.40–0.60	**marginal fail, 0.01 s low**
RT60 2 kHz occupied	1.43 s	0.38 s	0.40–0.60	**fail, 0.02 s low**
Flatness 250 Hz–4 kHz	1.64	1.18	<= 1.25	**PASS**

Metric	Before	After	Target	Result
Bass ratio	0.48	1.09	0.9–1.3	**PASS**
Background	47 dBA, tone	39 dBA, no tone	<= 40 dBA	**PASS**
Spatial decay across table	2.1 dB	5.4 dB	–	improved
`N_eff_turn / n`	0.604	0.812	>= 0.75	**PASS**
Silenced count `k`	3	1	0 or 1	**PASS**
Control room `N_eff_turn/n`	0.638	0.650	no change	**PASS — control did not move**

Two bands came in marginally below the 0.40 s lower bound and are reported as misses. They are 0.01 s and 0.02 s — within the measurement uncertainty of the method — and the correct action is to record them, note that the room sits at the dead end of its target range, and re-check after the room has been in use with its normal complement of coats, bags and papers, all of which add absorption. **Do not adjust the target to match the measurement.**

`N_eff_turn / n` rose from 0.604 to 0.812, a 34 per cent improvement in turn-taking equality, and the matched control room moved by 0.012. The treatment, not the observation, produced the change.

The group-performance measure showed a small improvement that did not reach the pre-registered threshold. Per the preregistration this is the declared non-supporting outcome, and it is reported as such: the mediator moved decisively, the outcome moved slightly, and this single case does not establish H3's causal claim. The treatment is retained on the strength of the accessibility gain, the fatigue reduction, the recovered `ln N` ceiling and the movement of headroom from `−1.055` to `+0.058` — all of which are worth having on their own account, and none of which requires H3 to be true. **H2 is untouched by this result**, because H3 could fail purely through the direction of the correlation in link 6 of Section 7.1, which the source literature does not establish.

What this case demonstrates. A hand clap and a tape measure diagnosed the room in under five minutes and Sabine's equation priced the diagnosis in half an hour. The exclusion of three participants was computed from physics, independently of anyone's report, which made it un-arguable in a way a survey never is. **The tail arithmetic — 488 ms to 98 ms — is the mechanism.** Obeying the Eyring switching rule changed the design before it was built, and the change was a saving. The specification was reported as failing two bands by 0.01 s and 0.02 s. **And the causal claim was tested and not confirmed, and that was written down.**

11.10 10 Discussion

11.10.1 10.1 What follows

Turn-taking inequality has a physical component that is computable in advance, and it is computable from three measurements anyone can take. The chain from `RT60` to a minimum inter-turn gap is arithmetic — `t_15 = RT60/4` — and it converts an observation about a meeting into a property of a building. That matters most because of what it removes from the discussion: **an exclusion computed from a tape measure, a phone and Sabine's equation cannot be answered with an assertion about anyone's character.** Where the physics dominates, the remedies that address character are addressing the wrong term.

Ceiling loss from exclusion is large and it is permanent. Three of eight silenced costs 0.4700 nats, 22.6 per cent of the room's per-interaction ceiling, in every meeting for the life of the building. Against that, the treatment is a one-time capital item with a decay rate set by mineral wool. **The asymmetry between a decaying store — attention, skill, facilitation — and a durable one is the whole economic argument for building the fix into the fabric**, and it is the same asymmetry that makes stored infrastructure invisible until it fails [Star and Ruhleder, 1996].

Boundary is the cheap variable and interior is the expensive one. A room's floor area is fixed by the lease and its boundary length is not; Theorem 5 prices the conversion exactly and the inward-only case shows that the floor spent is precisely the floor the seating occupies. The practical consequence is that a category-count deficiency in a room is almost never a floor-area problem, and treating it as one — by taking more space — is the expensive way to fix the cheap thing.

Two specification failures recur and both are diagnosable. A single-number `RT60` specification passes a room that is 2.4 times more reverberant at 250 Hz than at 4 kHz, and quoting Sabine above `alpha_bar = 0.25` describes a room up to 78 per cent more reverberant than the one that gets built. Both are avoidable with arithmetic already in the design.

11.10.2 10.2 What does not follow

Nothing here shows that acoustic treatment improves group performance. H3 was preregistered, tested, and not supported by the worked case. **A reader who takes away "treat the room and the group gets smarter" has taken away the opposite of what Section 9.2 reports.** What the case supports is that treatment moves the mediator, and that the mediator's onward effect is unestablished.

Nothing here identifies a social effect. Section 3.2 states why: applied to a group whose members both influence and constitute the aggregate, the specification is reflexive and the endogenous social effect is not identified from observation [Manski, 1993]. **The matched control removes a confound; it does not solve the identification problem**, and no amount of before-and-after data will.

Self-similarity certifies nothing. Every fractal claim in this paper is Tier 3, and the worked case shows why: the alcove that the boundary argument said would sub-divide the floor failed its measurement by 1.3 dB, for a reason — an open banquette top — that no dimension calculation contains.

The scale ladder of Section 5.5 is not evidence. Five rows of stipulated parameters show what a trend would look like; they do not establish one. **A reader who cites that table as a finding about how organisations behave at scale has cited an illustration as a measurement.**

The variety term is not a universal good. Unbounded growth in `ln N` is a property of the model, not of groups: Uzzi and Spiro [2005] find an inverted U in which creative success rises and then falls, and Sagan [1993] documents settings in which redundancy reduces safety. **The model's `ln N` contradicts that literature at large `N` and the contradiction is a restriction on the model's domain, stated in Section 11.5, not a result.** Nor does a room's category count settle what kind of difference it is: separation, variety and disparity are distinct constructs travelling under one word [Harrison and Klein, 2007], and the audited categories here are **variety** — distinct kinds of use — and not disparity of viewpoint.

11.10.3 10.3 Relation to the surrounding literature

The acoustics is standard and is used as standard: Sabine's statistical-energy derivation [Sabine, 1922], the Schroeder frequency as the boundary of its validity [Schroeder, 1965], and the textbook two-field and critical-distance results [Rossing et al., 2002, Beranek, 2004]. **What is new is not the physics but its coupling: the decay tail is converted into a turn-taking constraint and then priced against a surplus ceiling, and neither step exists in the acoustics literature because neither is an acoustics question.**

On the sociological side, the collective-intelligence result supplies the correlate [Woolley et al., 2010] and its replication record supplies the caution [Engel et al., 2014, Bear and Woolley, 2011]. The design tradition supplies the two demands for boundary — that people occupy edges first [Gehl, 1971] and that use tracks sittable perimeter [Whyte, 1980] — and the pattern tradition supplies the form of the seating rule [Alexander et al., 1977, Alexander, 1965], while the safety argument for sightline into every recess is Jacobs [1961]'s. **What this paper adds to**

that tradition is a price: the pattern literature states what to build and rarely states what it costs in the units of anything else, and Section 9.1 states it in floor area and in nats.

Where the literature disagrees with the framework, it is cited and the disagreement is stated. The diversity-and-performance literature is not uniformly supportive of a variety-raising programme [Mannix and Neale, 2005, Thompson, 2014, van Knippenberg and Schippers, 2007], the formal results that a diversity advantage requires are sensitive to their assumptions [Hong and Page, 2004, Singer, 2019, Grim et al., 2019, Kuehn, 2017], and organisational redundancy is not always protective [Sagan, 1993, Perrow, 1984]. **A framework that reads a higher category count as a higher ceiling owes those citations, and it owes the statement that the ceiling is a bound and not a prediction.**

Finally, the logarithm. **Two warrants exist and this paper needs only the weaker demand.** The Fechnerian warrant — sensation proportional to the logarithm of a dimensionless stimulus ratio [Fechner, 1966] — requires a contested perceptual story. Peters and Gell-Mann [2016] requires no psychological premise at all: the time-average growth rate of a multiplicatively evolving quantity is the expectation of its logarithm, so where a stock compounds rather than adds, the logarithm is the ergodic-average object. **This paper's argument needs only the second.** And the form remains underdetermined: a sub-linear power law [Latané, 1981] and an `n log n` growth [Briscoe et al., 2006] reproduce the same qualitative facts, and **nothing assembled here discriminates between them** [Stumpf and Porter, 2012].

11.11 11 Verification

The results of this paper compound from six postulates. Each is atomic, each is checkable on its own, and each carries its status. **Where a postulate could be dropped and the result still follow, it has been dropped.**

11.11.1 11.1 Dimensional audit

Two dimensional systems meet in this paper and the audit must cover both. `[L]`, `[T]` and `[M]` are length, time and mass; `[Q]` is the dimension of stored order and the period is the unit of `[T]` in the storage law, which is why no explicit `[T^-1]` appears once that law is written in discrete steps.

Symbol	Dimension	Unit	Status
`V`	`[L^3]`	m^3	**measured**
`S` (surface)	`[L^2]`	m^2	**measured**
`alpha_i`	`[1]`	none	**measured**, an absorbed-energy fraction in `[0,1]`
`A`	`[L^2]`	m^2 sabins	**derived**: `sum S_i alpha_i`, an area times a pure fraction
`alpha_bar`	`[1]`	none	**derived**: `A/S`, dimensionless by construction
`RT60`	`[T]`	`s`	**measured**; `0.161 V/A` carries `[T L^-1] x [L^3]/[L^2] = [T]`
0.161	`[T L^-1]`	s/m	**derived** from the speed of sound and the mean free path `4V/S`

Symbol	Dimension	Unit	Status
`t_15`	`[T]`	s	**derived**: `RT60/4`, so it inherits `[T]` exactly
`c` (speed of sound)	`[L T^-1]`	m/s	**measured**, 343 m/s at 20 C
`lambda`, `D`, `d`, `w`, `f`, `r`	`[L]`	m	**measured** or **specified**
`f_floor = 2c/D`	`[T^-1]`	Hz	**derived**: `[L T^-1]/[L]`
`R = A/(1-alpha_bar)`	`[L^2]`	m^2	**derived**
`d_c = 0.141 sqrt(QR)`	`[L]`	m	**derived**; 0.141 is dimensionless and `Q` is dimensionless
`L_p`, SNR, D/R, gain	`[1]`	dB	**derived**: a decibel is ten times the base-10 logarithm of a power ratio and is dimensionless
`D = ln m / ln r`	`[1]`	none	**defined**: a ratio of logarithms of pure counts
`1 + d/w`	`[1]`	none	**derived**: a ratio of two lengths
`s` (Lombard slope)	`[1]`	dB per dB	**measured** elsewhere, dimensionless by construction
`N`, `k`	`[1]`	count	**measured** by audit; integers
`u`, `U_t`	`[1]`	none	**defined** dimensionless: `u = 1 - (x.y)^2` with unit vectors
`S_t` (surplus)	`[1]`	nats	**derived** dimensionless: a logarithm of an effective count
`N_eff_turn`	`[1]`	count-equivalent	**derived**: `exp` of an entropy in nats
`Q_t`	`[Q]`	proxy units (standing decisions)	**measured**, in whatever the declared proxy counts
`delta`	`[1]`	per period	**measured** dimensionless fraction per step
`C_t`	`[Q]` per step	proxy units per week	**measured**
`eta`	`[Q]` per nat per step	proxy units per nat per week	**required by the balance**; see below
`h_t`	`[Q]` per step	proxy units per week	**derived**
`Q_inf`	`[Q]`	proxy units	**derived**: `[Q]/step` over `[1]/step` returns `[Q]`

Does each equation balance? *Sabine:* `0.161 V/A` gives `[T L^-1][L^3][L^-2] = [T]`, and `A` is an area because `alpha` is a pure fraction — **so a "square metre sabin" is a square metre, and calling it a sabin is a bookkeeping label recording that the area has been weighted, not a new dimension.** *The frequency floor:* `2c/D` gives `[L T^-1]/[L] = [T^-1]`, a frequency, and the factor 2 is the dimensionless `D = 2 lambda` convention. *The tail identity:* the decay rate is `60/RT60` in `dB [T^-1]`, so a 15 dB fall takes `15/(60/RT60) = RT60/4` in `[T]`. **The identity is dimensionally forced and the only content in it is the 15 dB criterion, which is a choice.** *The storage law:* `Q_{t+1} = (1-delta)Q_t + eta S_t - C_t` balances if and only if every term on the right carries `[Q]`. The first does, since `delta` is dimensionless; the third does by construction; the second does **only if `eta S_t` carries `[Q]`, and since `S_t` is a dimensionless count in nats, `eta` must carry `[Q]` per nat per step.** The gauge then balances: `ln N` is dimensionless, so `eta ln N` carries `[Q]` per step and is commensurable with `delta Q_t` and `C_t`.

The source is silent on `eta`'s dimension and the silence is not neutral. The manuscript this paper converts states that `eta` has units of stored order over surplus and that, surplus being dimensionless, `eta` carries whatever unit the order proxy carries. **That is the `[Q]`-per-nat half and it omits the per-step half entirely**, and it then advises normalising the stock so that `eta` lands in `[0,1]` — advice that treats `eta` as a pure number. **The reading required for the law to balance is `[Q]` per nat per period, it is stated here explicitly, and it is recorded that the source does not supply it rather than being assigned quietly.** The practical consequence bites in Section 9.2, where `eta` is quoted as 0.50 and 0.60 against a stock counted in standing decisions and a cost counted in stored-order-equivalent units per week: **those three denominations must be declared together or the arithmetic is malformed however tidy it looks.**

Where a quantity has no operational unit, it is said here. `Q` is stored order and stored order has no natural unit: it is counted in whatever the declared proxy counts, and two systems' `Q` values are not comparable unless they declare the same proxy. **Consequently `eta` and the cost flow are not comparable across systems either.** The audited category count `N` is a count of judgements, not of objects, and the same room audited by two surveyors under the same strike rule may return different integers. **And comfort itself, the underlying construct of the corpus, has no operational unit at all**; it is treated as a dimensionless magnitude by stipulation, not by measurement.

Three dimensionless quantities in this paper are logarithms and none of them convert into the others. A decibel is ten times a base-10 logarithm of a power ratio; an octave is a base-2 logarithm of a frequency ratio; a nat is a natural logarithm of an effective count. **No conversion between them is performed anywhere in this paper**, and where the corpus does define one it defines a different object: the octave form for a *required surplus*, in which the octaves required are `2(c/eta)/ln 2 = 2.885390081777927 x (c/eta)`, so that one nat of required surplus corresponds to 2.885390 octaves, 3462.4681 cents and 1.442695 bits [Curzi, 2026]. **That is a conversion of a surplus requirement into a pitch interval and it is not a conversion of an octave band of measured reverberation into anything.** The six octave bands in every acoustic table here are frequency ranges, not surplus. And **the entropy in `N_eff_turn` is Shannon entropy in nats, not thermodynamic entropy in `[energy][temperature]^-1`.**

11.11.2 11.2 The postulates

P1 — Sound energy in a rectangular enclosure above its Schroeder frequency decays exponentially at a rate set by the total absorption, giving `RT60 = 0.161 V/A`. *Status:* **CITED** — Sabine's statistical-energy result [Sabine, 1922], with the validity boundary from Schroeder [1965] and the standard exposition in Rossing et al. [2002].

P2 — By the definition of `RT60`, sound decays at `60/RT60` decibels per second. *Status:* **DEFINITION** — true by the construction of the quantity, and carrying no empirical weight on its own.

P3 — A talker must be audible over the previous talker's tail to take a turn, and 15 dB of separation is the criterion adopted here. *Status:* **POSTULATED** — the 15 dB figure is a choice. The intelligibility literature places near-maximum performance around +15 dB signal-to-noise for

normal-hearing adult native listeners, which is the reason for the number, but **no controlled measurement in this paper establishes that 15 dB is the threshold at which a turn succeeds or fails.**

P4 — `ln` is strictly increasing and strictly concave. *Status:* **ARITHMETIC.**

P5 — Surplus per interaction is bounded, `0 <= S <= ln N`, so a reduction in the effective count from `N` to `N-k` reduces the ceiling by `ln N - ln(N-k)`. *Status:* **CITED** — the range result of the surplus functional [Curzi, 2026], with the forced-ceiling form in Curzi [2026l].

P6 — A participant whose direct sound at the longest talker-to-listener path falls below the room's combined noise floor while another participant holds the floor is effectively silenced, and their category is removed from the interaction's effective count. *Status:* **POSTULATED** — this is the step that converts a signal-to-noise deficit into an integer, and it is the only place in the chain where physics becomes sociology. **Different thresholds give different counts, and nothing in this paper measures which threshold is right.**

The three-tier certification standard is **not** in this list, because no result in the chain consumes it. It is a governance device on the paper's own claims. The similarity dimension `D = ln m / ln r` is likewise absent: it is a definition, it appears in the chain nowhere, and it is retained in the text only because it is what the fractal literature is about.

11.11.3 11.3 The compounding chain

```
P1              -> L1   in a diffuse field, RT60 = 0.161 V/A, so absorption and
                        reverberation time are reciprocal.                    CITED

P2              -> R1   a gap of t seconds leaves the tail 60t/RT60 dB down;
                        inverting, the gap for a fall of x dB is
                        t_x = x*RT60/60, and at x = 15, t_15 = RT60/4.        DEFINITION
                                                                              + ARITHMETIC

R1 + P3         -> R1'  the room imposes a minimum inter-turn gap of RT60/4;
                        at RT60 = 1.95 s that is 488 ms and at 0.39 s it is
                        98 ms, a factor of 5.                                 POSTULATED (P3)

L1 + R1         -> R4   Sabine permits RT60 to stay finite as alpha_bar -> 1,
                        which is physically wrong, so Eyring's
                        -S ln(1 - alpha_bar) replaces A above alpha_bar = 0.25;
                        at alpha_bar = 0.724 the two differ by 77.8%.         ARITHMETIC
                                                                              given L1

P6              -> L2   the count k of effectively silenced participants is
                        an integer produced by a declared threshold.          POSTULATED

L2 + P5         -> R3   the ceiling loss is ln N - ln(N-k); at N = 8, k = 3
                        that is ln 8 - ln 5 = 0.4700 nats, 22.6% of ln 8.     ARITHMETIC
                                                                              given L2, P5

P4 + R3         -> R3'  the loss is convex in k: the first voice silenced
                        costs 0.1335 nats and the third costs 0.1824, so
                        exclusion accelerates.                                ARITHMETIC

(geometry)      -> R2   a square-wave boundary of depth d and half-period w
                        has length multiplied by exactly 1 + d/w; at
                        d = 1.00, w = 1.50 that is 1.6667, taking 43.82 m
                        to 73.03 m.                                           ARITHMETIC
```

(geometry)	-> R5	a reflector's useful-directivity floor is f_floor = 2c/D; at D = 1.2 m and c = 343 m/s that is 572 Hz.	ARITHMETIC on a Tier 2 convention
R3 + storage law	-> R6	with N restored from 5 to 8, C cut by 0.55 and eta raised 0.50 -> 0.60, h moves −1.0553 -> −0.2623, 75.1% of the gap; shedding Q from 12 to 8 gives +0.0577.	ARITHMETIC given stipulated parameters

Every line is checkable from its own inputs. R1, R2, R3, R3', R4, R5 and R6 are arithmetic consequences of one cited physical law, one definition and elementary properties of the logarithm. **R1' and R3 consume P3 and P6 respectively, and those two postulates carry the whole sociological content of the paper.** P3 fixes what counts as clearing the tail; P6 fixes what counts as being silenced. **If either threshold is wrong, the mechanism survives and every number attached to it moves.** That is stated here rather than left to be found.

What this chain does not establish. It does not establish that turn-taking equality causes group performance; that is H3, it was preregistered and tested in Section 9.2, and it was not supported. It does not establish that any organisation obeys the storage law; that is a modelling postulate whose one preregistered confirmatory test reported in this corpus **failed its second falsification rule, returning a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared margin of 10 per cent** [Curzi, 2026i]. **A chain whose steps are ARITHMETIC, CITED, DEFINITION and POSTULATED delivers the acoustics with certainty and delivers nothing whatever about committees.**

11.11.4 11.4 The checks a reviewer can redo in a minute

The tail identity. Sound falls `60/RT60` dB per second. At `RT60 = 1.95 s` that is `60/1.95 = 30.77` dB/s, so 15 dB takes `15/30.77 = 0.4875 s = 488 ms`. At `RT60 = 0.39 s` it is `60/0.39 = 153.85` dB/s, so 15 dB takes `0.0975 s = 98 ms`. **The ratio is exactly `1.95/0.39 = 5.00`, because `t_15 = RT60/4` is linear in `RT60`.** ARITHMETIC.

The exclusion loss. `ln 8 = 2.079442`, `ln 5 = 1.609438`, difference `0.470004`, and `0.470004/2.079442 = 0.226023`. **0.4700 nats and 22.6 per cent.** And the convexity: `ln 8 − ln 7 = 0.133531`, `ln 7 − ln 6 = 0.154151`, `ln 6 − ln 5 = 0.182322`. **Each successive silencing costs more.** ARITHMETIC.

The crenellation multiplier. A 120 m^2 square has side `sqrt(120) = 10.954451` m and perimeter `43.817805` m. At `d = 1.00`, `w = 1.50`, the multiplier is `1 + 1/1.5 = 1.666667` and the new perimeter is `43.817805 x 1.666667 = 73.029675` m, **73.03 m.** The floor cost is `21.909 x 1.00 = 21.91 m^2`, leaving `98.09 m^2`, and the requirement at the 0.1094 ratio is `98.09 x 0.1094 = 10.731 m`, met `73.03/10.73 = 6.81`-fold. ARITHMETIC.

Sabine, on one cell. State A at 1 kHz: `A = 8.24`, so `RT60 = 0.161 x 100 / 8.24 = 1.9539 s`, **1.95 s to two decimals.** State C at 1 kHz: `A = 39.08`, so `0.161 x 100 / 39.08 = 0.4120 s`, **0.41 s.** ARITHMETIC.

The Eyring divergence, on the worst cell. State D at 2 kHz: `A = 104.99`, `S = 145`, so `alpha_bar = 0.72407`. Sabine gives `16.1/104.99 = 0.153348 s`. Eyring gives `16.1 / (−145 ln(1 − 0.72407)) = 16.1 / (145 x 1.287499) = 16.1/186.687 = 0.086241 s`. The ratio is `0.153348/0.086241 = 1.77812`, **an over-prediction of 77.8 per cent.** ARITHMETIC.

The frequency floor. `f_floor = 2c/D = 2 x 343/1.2 = 571.67`, **572 Hz to the nearest hertz.** And the correction recorded in Section 4.4: at `D = 9.10 m`, `D = 4 lambda` gives `343/(9.10/4) = 150.8 Hz`, **not the 225 Hz the source table carries**, which is the value for `D = 6.10 m`. ARITHMETIC.

The Denge gain, recomputed. `G = (pi D/lambda)^2 eta_ap` at `D = 61 m`, `lambda = 1.372 m`, `eta_ap = 0.5`: `pi x 61 = 191.637`, `/1.372 = 139.677`, squared `19509.7`, halved `9754.8`, and `10 log10(9754.8) = 39.89` dB. **39.9 dB, not the 46.9 dB the source states.** ARITHMETIC.

The turn-share statistic. For shares `0.38, 0.30, 0.13, 0.08, 0.05, 0.03, 0.02, 0.01`, `−sum p ln p = 1.575466` nats and `exp(1.575466) = 4.8328`, so `N_eff_turn/n = 4.8328/8 = 0.6041`, **0.604.** ARITHMETIC.

The headroom chain, end to end. `0.50 x ln 5 = 0.50 x 1.609438 = 0.804719`; `0.804719 − 0.96 − 0.90 = −1.055281`. `0.60 x ln 8 = 0.60 x 2.079442 = 1.247665`; `1.247665 − 0.96 − 0.55 = −0.262335`; `1.247665 − 0.64 − 0.55 = +0.057665`; `1.247665 − 0.48 − 0.55 = +0.217665`. **The gap closed on acoustics alone is `(1.055281 − 0.262335)/1.055281 = 0.75141`, 75.1 per cent.** ARITHMETIC.

The steady state, on one row. Treated room at `u* = 0.20`: `N_eff = 1 + 7(0.20) = 2.40`, `ln 2.40 = 0.875469`, `0.60 x 0.875469 = 0.525281`, and `Q_inf = (0.525281 − 0.55)/0.08 = −0.309`. **Negative, as the table says.** ARITHMETIC.

The Koch limit. Perimeter `3(4/3)^n` diverges; area converges to `8/5 = 1.6` times the original triangle, and the table's eighth iteration gives `1.59909`, approaching it from below. ARITHMETIC.

11.11.5 11.5 Limits, asymptotics, and the model's valid domain

`A -> 0` (a perfectly hard room). `RT60 -> infinity`, `d_c -> 0`, and the minimum inter-turn gap `t_15 -> infinity`. **The model says nobody can ever take a turn, which is the correct limit of the mechanism and a physically unreachable state**, since occupants alone contribute absorption: eight seated adults give `A = 3.36` at 1 kHz, capping `RT60` at 4.79 s in this volume no matter what the surfaces do. **The lower bound on `A` is the people.**

`alpha_bar -> 1` (a perfectly absorptive room). Sabine gives `RT60 = 0.161V/S`, a finite positive number, **which is wrong**: an enclosure that absorbs everything has no reverberation. Eyring gives `−S ln(1 − alpha_bar) -> infinity` and hence `RT60 -> 0`, which is right. **This is not a curiosity: it is why the switching rule exists, and it is the reason Section 6.1 sets the crossover at `alpha_bar = 0.25` rather than leaving Sabine to run to its own limit.**

`RT60 -> 0`. `t_15 -> 0`, so the tail imposes no constraint at all. **The model then predicts perfect turn-taking equality, and Section 6.3's State D shows the prediction is false**: below about 0.30 s the room stops supplying reverberant support, talkers push, vocal effort and fatigue rise, and the quiet talker — who cannot push — is worse off. **The mechanism is one-sided and the model has no term for the other side. The stated domain is `RT60` in `[0.30, 3.0]` s**, and outside it the tail arithmetic still evaluates and should not be interpreted.

`D -> 0` (a small reflector). `f_floor = 2c/D -> infinity` and the device does nothing at any frequency of interest. **A reflector smaller than about 0.3 m is acoustically invisible to speech**, which is failure mode FA3 of Appendix B. **`D -> infinity`.** Gain grows as `D^2` without bound, which the aperture formula permits and which real construction does not: surface accuracy, near-field geometry and head blockage all degrade with size in practice. **The gain formula has no upper domain limit and must be given one empirically.**

`d/w -> infinity` (deep, narrow crenellation). The multiplier `1 + d/w` grows without bound and the perimeter diverges, which is the Koch behaviour. **Taken literally this says a room can have unlimited sittable edge, and it is false**: below `d = 0.40 m` and `w = 0.60 m` the boundary cannot hold a seated person, and the third self-similar order of Section 5.3 — a 111 mm recess at 167 mm pitch — is a texture. **The domain restriction is the seating-scale floor, and it is requirement G15 precisely because the divergence is otherwise unfalsifiable.**

`N -> 1`. `ln N = 0` and `F(u) = 0` for every `u`. **A single audited category has zero available surplus, not a small amount**, so a room offering one way to be in it has no ceiling to lose. This is the boundary that makes a collapse of effective count consequential: a structure with one effective bloc has `ln(1.00) = 0` available surplus [Curzi, 2026].

`N -> infinity`. `eta ln N -> infinity`, so the per-period budget grows without bound. **Taken**

literally this says that a sufficiently high category count makes any cost sustainable, and that is false. Three things bound it and **all three are domain restrictions rather than results.** *Internal:* the room has a floor area, and each additional posture/purpose/condition category needs physical provision, so `N` is bounded by the crenellated boundary the floor can carry — in Section 9.1, five categories on 98 m^2. *External and contrary:* Uzzi and Spiro [2005] find an inverted U in which creative success rises and then falls with cohesion, so the benefit of difference turns over rather than growing without bound; Sagan [1993] documents settings in which redundancy reduces safety; and May [1972] and Gardner and Ashby [1970] show that increasing diversity and connectance destabilises randomly assembled systems. **The unbounded growth of `eta ln N` contradicts this literature and the contradiction is not resolved here.** *Marginal:* even inside the model, `ln N − ln(N−1) ~ 1/N` vanishes, so the divergence is the slowest available route to anything. **The stated domain is therefore `N` bounded above by what the room can physically provide, and the logarithmic ceiling is a within-domain bound rather than an asymptotic claim.**

`k -> N`. `ln(N−k) -> ln 1 = 0` and the entire ceiling is lost; at `k = N` the expression is `ln 0`, undefined, and the model has left its domain. **A room that silences everyone is not a room with zero surplus in this model's sense; it is outside the model.**

`u -> 0` and `u -> 1`. At `u -> 0`, `N_eff -> 1`, `F -> 0`, and no acoustic or geometric intervention returns anything — which is the Denge lesson in the framework's units: **passive geometry concentrates what arrives and generates nothing.** At `u -> 1`, `F -> ln N`, the ceiling is attained, and by the fixed-point identity the converged system's headroom goes to zero [Curzi, 2026b].

`s -> 1` in the Lombard loop. `1/(1−s) -> infinity`: the room's level diverges from any seed disturbance. **The divergence is real in the first-order model and is the reason the reported slope range of 0.3–0.6 matters**; at `s >= 1` the loop does not converge and the contraction argument does not apply [Banach, 1922]. **No room is claimed to reach it.**

`delta -> 0`. The fixed point disappears and the stock accumulates without bound; that is a pathology rather than an ideal and it is outside the model's domain [Curzi, 2026b].

Stated domain. The results of this paper are valid for `RT60` in `[0.30, 3.0]` `s`, `alpha_bar` in `(0, 0.95)` with the Eyring form above 0.25, room volumes above the Schroeder-frequency condition in the bands reported, apertures `D >= 2 lambda` for any directivity claim, crenellation with `d >= 0.40 m` and `w >= 0.60 m` for any seating claim, `N >= 2` bounded above by physical provision, `k < N`, `u` in `[0,1]`, `delta` in `(0, 0.5]`, and Lombard slope `s < 1`. **Outside that domain the equations still evaluate and they should not be interpreted.**

11.11.6 11.6 Invariance

Scale invariance of the crenellation multiplier. `1 + d/w` depends only on the ratio `d/w`, so under a uniform rescaling of the plan by any positive factor `c` — `d' = cd`, `w' = cw` — the multiplier is unchanged while the perimeter scales as `c` and the area as `c^2`. **That is the exact invariance the boundary argument uses: boundary is linear in scale and interior is quadratic, so the ratio of boundary to area falls as `1/c` and the multiplier is what recovers it.** What would break it: any threshold expressed in absolute length rather than in ratio — which is exactly requirement G15's `d >= 0.40 m`. **The multiplier is scale-free and the seating floor is not, and the second is what stops the first from proving too much.** This is the same structure as the corpus's own declaration that a scale-free form of a law licenses modelling at a new scale and never licenses skipping the measurement at that scale.

Wavelength invariance in acoustics. The frequency floor `D = 2 lambda` and the beamwidth `70 lambda/D` depend only on `D/lambda`, so a reflector scaled by `c` behaves identically at frequencies scaled by `1/c`. **A 1.2 m shell at 572 Hz and a 0.6 m shell at 1143 Hz are the same device.** This invariance is exercised at six aperture sizes spanning a factor of 200 in Section 4.4 and it is exact within the aperture approximation. **What breaks it is anything with an absolute scale in it: a human head at the focus, an ear's frequency response, and the speech spectrum itself. The physics is scale-free and the application is not.**

Unit rescaling of the storage law. Let `Q' = cQ`. The law is invariant provided `eta`, `C` and `h` rescale with it while `delta`, `N` and `u` are unchanged. **`delta`, `u`, `N`, `alpha_bar`, `d/w`, `D/lambda` and every decibel quantity are the unit-free objects here; `Q`, `eta`, `C` and `h` are the unit-carrying ones**, and two rooms can be compared on the first group directly and on the second only if they declare the same proxy.

What would break the acoustic invariances, in order of likelihood: a room below its Schroeder frequency, where modal behaviour replaces the statistical field and `RT60` is not a single number; absorption concentrated on one surface, which prevents a diffuse field forming and is failure mode FA4; and coupling to an adjacent volume, which produces a two-slope decay no single `RT60` describes. **All three are conditions of Section 6.1 and all three are checked in Section 9.2 rather than assumed.**

How many scales this has been exercised at. The geometric argument: three built rungs in one constructed room, and a five-rung illustrative table that measures nothing. The acoustic argument: one constructed room of 100 m^3, in four states. **That is one volume, one geometry, one author and no independent replication.** The invariances are claims about form that have been exercised, not established.

11.11.7 11.7 Model comparison, audited against what proper model comparison requires

This paper fits no model to data, so no information criterion applies to its own results. It nonetheless makes quantitative claims that a reviewer will want compared against alternatives, and the requirement is stated rather than quietly skipped.

Requirement	Status in this paper
Out-of-sample rather than in-sample performance	**NOT APPLICABLE.** Nothing is fitted; the worked cases are constructed and there is no sample
A matched null with the same free-parameter count	**NOT MET.** The natural null for Section 9.2 is that turn-share evenness drifts on its own. The matched untreated control addresses that for the *treatment* and is a design control, not a model comparison. **No alternative model of turn-share dynamics was specified or fitted**
The number of free parameters stated explicitly	**MET for the worked cases.** Section 9.2 fixes eight: `N` before and after, `eta` before and after, `delta`, `Q`, `C` and its acoustic component. **All eight are stipulated or surveyed, none is fitted, and none carries an interval**
An information criterion, so the parameter count is paid for	**NOT MET.** No AIC or BIC was computed here, and none has been computed anywhere in this corpus, for any model, at any block length
Effect sizes with intervals, not only p-values	**NOT MET.** No p-value is reported and no interval accompanies any number in Section 9. The turn-share change from 0.604 to 0.812 against a control moving 0.012 is an effect size **without an interval**, because the underlying six-meeting means carry no reported dispersion

The unmet items change how the headline should be read, so they are stated plainly. The

parameter count in Section 9.2 was never paid for: eight stipulated parameters can produce almost any headroom trajectory, and the case demonstrates that the procedure runs rather than that the model is right. **A reader who wants the acoustic result should take the tail arithmetic and the exclusion loss, which are ARITHMETIC and CITED, and should not take the headroom trajectory, which is POSTULATED throughout.**

No scaling or heavy-tail claim is made anywhere in this paper, so the distribution-fitting procedure of Clauset et al. [2009] has nothing to bite on here. It does bite on the framework's cost exponent, which is defined by the model rather than estimated, and that is recorded in Curzi [2026b]. **The similarity dimension `D = ln m / ln r` is likewise a definition and not an estimated exponent**: it is computed from a construction rule, no data is fitted, and **any claim that a real boundary has a measured fractal dimension would require the full estimation procedure and is not made here.** Stumpf and Porter [2012] states the deeper problem — many mechanisms produce indistinguishable curves — and it applies directly to the underdetermination of the surplus functional's own form discussed in Section 10.3.

11.11.8 11.8 Scale type, analogy discipline, and the precision of the numbers

Scale type, quantity by quantity. `RT60`, `A`, `V`, `d`, `w`, `D`, `lambda` and `f` are **ratio-scale physical quantities** with natural zeros, and every operation performed on them — differences, ratios, logarithms — is licensed. **Decibel quantities are interval-scale in the logarithmic domain and ratio-scale in the underlying power**, so differences of decibels are meaningful and "twice as many decibels" is not, and no such statement is made. **`N` and `k` are counts on a ratio scale but they are counts of judgements**: the audited category count is the output of a strike rule applied by a surveyor, and its integrality conceals a continuum. **The underlying audibility judgement in P6 is at best ordinal** — a participant is more or less audible — and converting it into a count is an operation the scale does not license. **It is performed anyway, it is the paper's sharpest measurement-theoretic exposure, and it is stated here rather than defended.**

`u` and `Q` are treated as ratio-scale throughout and the conditions that would license that have not been tested. Michell [1997] argues that psychology systematically fails to test whether its attributes are quantitative at all before measuring them, and **that is the strongest single objection to this framework's founding move.** The formal remedy is conjoint measurement — the test for whether an attribute has additive quantitative structure [Luce and Tukey, 1964, Krantz et al., 1971] — or a Rasch demonstration [Rasch, 1960]. **Neither has been run for any quantity in this paper.** The honest answer is that the ratio-scale treatment of stored order and non-parallelism is a stipulation, and that the physical quantities in this paper are not exposed to that objection at all: **the acoustics is measurement in the strict sense [Stevens, 1946] and the storage law is not.** That asymmetry is the reason this paper leans on the physics and reports the headroom trajectory as illustration.

Analogy discipline, restated at the point of exposure. Sabine's equation and the two-field model are **derivations** used within their own domain, with their five premises listed and checked. The Lombard loop's `1/(1−s)` is a **shared mathematical form** with the storage law's convergence, and nothing about the physics of either transfers to the other beyond the difference equation. The self-similarity argument for social occupancy is an **analogy** and is used only to generate candidate geometries, never to certify one — which is what the Tier 3 obligation enforces. **The failure mode is letting the second or third borrow the certainty of the first, and Section 9.1's alcove failing its measurement by 1.3 dB is the concrete instance: the boundary-length argument was category two, the sub-division claim was category three, and only the measurement settled it.** Cartwright [1999] offers the escape route this paper takes — specify the arrangement within which the law is claimed to hold, which is the stated domain of Section 11.5 — and Anderson [1972] is the standing reminder that a law true at one scale need not be informative at another.

Precision. The numbers here are of four kinds. *Physical measurements* — `RT60` figures, background levels, level differences — are quoted to two decimals in seconds and one in decibels, which is roughly what the method supports; **the two commissioning bands that missed**

by 0.01 s and 0.02 s are inside that uncertainty and are reported as misses anyway. *Derived-from-defined* quantities — `1 + d/w = 1.6667`, `ln 8 = 2.079442`, `D = ln 4/ln 3 = 1.2619`, `f_floor = 572 Hz` — are exact to the precision printed and may be recomputed further. *Nominal inputs* — every absorption coefficient — carry two decimals that record a published typical value and **not a measurement of anything that will be built.** *Illustrative model output* — every headroom figure, every steady state, the five-rung parameter ladder — is the output of stipulated parameters and **carries no measurement uncertainty because there was no measurement**; four decimal places there record the arithmetic, not confidence. **Any application of this framework to real data must report fitted parameters with intervals, and the four-decimal presentation used in Section 9 must not be carried over to a measured `eta`, where two significant figures with an interval would be the defensible form.**

11.12 12 Limitations

Both worked systems are constructed, not observed. No dataset was collected for this paper. Every number in Section 9 is the output of the stated parameters under the stated models, carried through so that a reader can re-derive it. **They demonstrate that the procedure runs and that it reports its own failures; they are not evidence that any room or any committee behaves this way and must never be cited as such.**

The causal claim was tested and failed, and that is the paper's headline limitation. H3 — that acoustic treatment raises collective performance through the turn-taking channel — was preregistered with a matched untreated control and did not reach its threshold. **The mediator moved from 0.604 to 0.812 and the outcome did not move.** The treatment is retained on accessibility, fatigue and ceiling grounds. **Any downstream paper citing this one for a collective-intelligence benefit is citing something this paper reports as unsupported.**

Two numbers in this paper convert physics into sociology and both are thresholds chosen rather than measured. P3 sets 15 dB as the tail clearance a turn requires; P6 sets the audibility criterion that turns a signal-to-noise deficit into an integer count of silenced participants. **Neither is measured anywhere, different thresholds give different answers, and every headline number — `t_15 = RT60/4` at 488 ms, `k = 3`, the 0.4700-nat loss — moves with them.** The sensitivity test the procedure demands was specified and is not reported here.

The count `k` is an integer produced from an ordinal judgement, and that operation is not licensed by the scale of its inputs. Section 11.8 states it. It is performed anyway because the alternative is to leave the quantity unmeasured, and a defensible order of magnitude computed from instruments beats an unmeasurable quantity — **but a reviewer is entitled to treat `k` as the paper's softest number and this paper agrees.**

The reflection problem outranks every before-and-after result reported here. Applied to a group whose members both influence and constitute the aggregate, the specification is reflexive and the endogenous social effect is not identified from observational data [Manski, 1993, Shalizi and Thomas, 2011]. **The matched control removes a confound and does not solve the identification problem**, and no volume of before-and-after data will. The realistic repair is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001]. **Until then the turn-taking mechanism as applied to a committee is a modelling position rather than an identified empirical result.**

Three claims in this paper carry no bibliography entry, and they are named rather than dressed with a citation that does not support them. *(i)* The classroom reverberation and background limits — 0.6 s and 35 dBA — are a real published standard which the database assembled for this corpus does not contain, so the standard is named in prose and the transfer to deliberation rooms is declared as a transfer. *(ii)* The 200 ms median turn-transition gap is taken from the conversation-analytic literature, for which this corpus holds no entry; it is used only as an indicative reference and no result depends on its exact value. *(iii)* The Lombard slope range of 0.3–0.6 dB per dB, and the antenna-literature claims of Section 5.2, are likewise uncited. **Attaching a citation that did not actually support these sentences would have been**

worse than leaving them uncited, because a reviewer would check.

The collective-intelligence correlate is a laboratory finding extrapolated to a standing committee. Woolley et al. [2010] measured short-task groups, the replication record is mixed [Engel et al., 2014, Bear and Woolley, 2011], and the extrapolation is a modelling assumption. **The broader diversity-and-performance literature is not uniformly supportive of the programme this paper's ceiling argument assumes** [Mannix and Neale, 2005, Thompson, 2014, van Knippenberg and Schippers, 2007], and the formal diversity results are sensitive to their assumptions [Hong and Page, 2004, Singer, 2019, Grim et al., 2019, Kuehn, 2017].

Every absorption coefficient is nominal and no room in this paper was built. The four build states are computed from generic published values, and the difference between a nominal coefficient and a manufacturer's tested value at the intended mounting condition is wide enough to change a design decision. **The commissioning figures reported in Section 9.2 are what a measurement of the built room would return if the nominal values were right, and that is a stronger assumption than any other in the acoustic sections.**

The five-rung parameter ladder is the paper's weakest quantitative object. Its parameters are stipulated to show the shape of a drift, none is measured, none carries an interval, and five rows cannot establish a trend. **It is retained because the modelling choice it illustrates is checkable in principle and because failure mode FG6 exists precisely to name its abuse**, and it is recorded in the corpus's UNTESTED drawer [Curzi, 2026].

The audited category count is a judgement with a strike rule, not a measurement. Two surveyors applying the same rule to the same room may return different integers, and `ln N` is sensitive to the difference at small `N`: the move from `N = 2` to `N = 3` is worth 0.405 nats, which is more than the entire acoustic exclusion loss of the second worked case. **A framework whose ceiling term depends on an integer produced by judgement should say how much rides on the judgement, and this is how much: one contested category in a two-category room is worth 86 per cent of the entire acoustic exclusion loss computed in Section 9.2.**

`N_eff = 1 + (N−1)u` is not a Hill number and no external literature contains it. `1/HHI` is the Hill number of order two and `exp(H)` the Hill number of order one, and both are fully grounded [Hill, 1973, Jost, 2006, Pielou, 1966]; `1 + (N−1)u` is a linear interpolation between 1 and `N`, adopted by design axiom [Curzi, 2026]. **The turn-share statistic used in this paper IS `exp(H)` and is grounded; the surplus law's effective count is not.** The two appear within pages of one another here, which is exactly the confusion this note exists to prevent [Curzi, 2026].

The variety measure is one of three incompatible constructs travelling under one word. Harrison and Klein [2007] distinguishes separation, variety and disparity. **The audited category count of Section 9.1 is variety** — distinct kinds of use — **and the non-parallelism `u` of the storage law is closer to disparity.** They are used in the same equation. That is a real exposure and it is not resolved here.

The surplus functional's form is underdetermined by the evidence. A sub-linear power law [Latané, 1981] and `n log n` growth [Briscoe et al., 2006] reproduce the same qualitative facts as `ln(1 + (N−1)u)`, and nothing assembled in this corpus discriminates between them [Stumpf and Porter, 2012]. **The logarithm is chosen for its axiomatic characterisation, not established by measurement**, and Peters and Gell-Mann [2016] is the strongest available warrant precisely because it needs no psychological premise.

The framework's reflexivity is not applied to its own parameters. This paper prescribes measuring `N_eff_turn/n` before and after treatment and reporting it upward. **A measure adopted as a target ceases to measure what it names** [Goodhart, 1975, Campbell, 1979], classifications of people interact with the people classified [Hacking, 1995], and a framework can make its own assumptions true [Ferraro et al., 2005, Merton, 1948]. **A committee told that its turn-share evenness is being scored can raise it without anyone being better heard**, and nothing in this paper detects that. The position taken is the narrow one: the parameters are treated as stable only over a declared window and only while the gauge is not itself a reported target [Lucas, 1976]. **Where it is reported upward, the gauge is at least as gameable as anything it audits.**

Unmet rigour checks, printed rather than left to a reviewer. *(i) Model comparison.* No

AIC or BIC was computed here or anywhere in this corpus, no matched null model of turn-share dynamics was specified, and no result carries an out-of-sample test — see the table in Section 11.7. *(ii) Effect sizes and intervals.* Nothing in Section 9 carries an interval, because nothing in Section 9 was measured on a real system. *(iii) Scale type.* Stored order and non-parallelism are treated as ratio-scale and the conjoint conditions were not tested [Michell, 1997, Luce and Tukey, 1964]. *(iv) Statistical control.* The one control in the paper is a matched untreated room, which is a design control and not a permutation or shuffle test; **no false-positive rate was measured against nominal.** *(v) Independent replication.* Every system here was constructed by one author [Open Science Collaboration, 2015, Ioannidis, 2005], and the analysis path was not preregistered even where the field test was [Gelman and Loken, 2013]. *(vi) Directional versus point prediction.* Meehl [1967] observes that directional predictions are nearly costless in the social sciences; **this paper does make point predictions with constants — 488 ms, 98 ms, 0.4700 nats, 572 Hz — which is the right answer to that objection, and one of them was put at risk and the associated causal claim lost.** *(vii) The discriminating observation.* Durlauf [2005] and Sobkowicz [2020] ask what observation would distinguish this account from a neighbouring one. **For the acoustics there is one — the cross-room dose-response of H1 — and it has not been run.**

What would falsify the account. H1 fails on measured turn gaps and override rates independent of `RT60` across rooms spanning a factor of three in reverberation time. H2 fails on treatment that meets its acoustic targets and leaves `N_eff_turn/n` unmoved. **H3 has already failed once, in this paper's own case.** H4 fails on an audited count that does not rise after construction, or on fine-scale features that time-lapse shows unoccupied at all times. H5 fails on occupancy counts showing no preference for positions satisfying all four seating conditions.

Nothing in this paper is a human-subjects finding. No interventions were run, no participants were recruited, and no recordings were made. The consent requirement of A26 and the declination constraint of Section 3.4 are design constraints on prospective work, not reports of past practice.

11.13 13 Conclusion

Passive geometry and room acoustics set three of the four terms of the storage law, and both can be priced against the ceiling `ln N` in the same units as any other intervention.

The crenellation multiplier is exactly `1 + d/w`, proven by summing segments, and it leaves the enclosed area unchanged when the teeth alternate about the original line. At `d = 1.00 m` and `w = 1.50 m` it is 1.6667, taking a 120 m^2 room's perimeter from 43.82 m to 73.03 m for 18 per cent of the open floor — **and the floor spent is precisely the floor the seating occupies.** In the worked room that converts an unmeetable sittable-edge requirement into a 6.8-fold surplus, takes the audited category count from 2 to 5, and moves headroom from `−1.396` to `−0.804`, which is 42.4 per cent of the gap, on geometry alone. **The parabola's reflective, focus-directrix and path-length properties are exact and buy phase correction with no parts and no power, and the reflector's hard frequency floor is `f_floor = 2c/D`, which for a 1.2 m aperture is 572 Hz and must be stated on the drawing.** A measured niche returned 4.8 dB at 1 kHz against a theoretical bound of 17.8 dB, and the honest description of that device is a consonant-band directivity improvement.

The turn-taking mechanism is tail arithmetic. Sound decays at `60/RT60` decibels per second, so the gap for the previous talker's tail to fall 15 dB is `t_15 = RT60/4`: 488 ms in an untreated 1.95 s room against a natural turn gap of about 200 ms, and 98 ms in a treated 0.39 s room. That is a factor of five, and it is the mechanism by which a quiet participant loses and recovers the ability to take a turn without anyone in the room behaving badly.

Acoustic exclusion is priced against the ceiling: silencing three of eight participants costs `ln 8 − ln 5 = 0.4700` nats, 22.6 per cent of the room's per-interaction surplus ceiling, permanently, in every meeting the room hosts — and the loss is convex in the number silenced, the third voice costing 0.1824 nats against the first voice's 0.1335. **The count is computed from a**

tape measure, a phone and Sabine's equation rather than from self-report, which is the only route available because a contribution never attempted leaves no trace.

The Sabine/Eyring switching rule is Sabine below mean absorption 0.25 and Eyring above, with the choice stated. In an over-treated room the two diverge by 61.9 per cent at 1 kHz and 77.8 per cent at 2 kHz: a designer quoting Sabine there believes they are building a 0.17 s room and builds a 0.11 s room. Obeying the rule changed the worked design before it was built, and the change was a saving of 5 m^2 of ceiling tile.

In the worked committee, headroom moved from `-1.0553` to `-0.2623` on acoustics alone — 75.1 per cent of the gap and still negative — and reached `+0.0577` only after shedding stored decisions from twelve to eight. In the untreated room the steady state is negative at every value of u, so no achievable amount of genuine difference among those eight members produces a sustainable stock of standing decisions.

The paper's counter-result is stated as plainly as its results. The preregistered falsification test was run and the causal claim failed. Turn-taking equality rose from 0.604 to 0.812 with the matched control room flat at a movement of 0.012, and the preregistered group-performance outcome improved slightly and did not reach threshold. The acoustic-to-collective-intelligence chain remains an unconfirmed postulate. Two commissioning bands missed the 0.40 s lower bound by 0.01 s and 0.02 s and were reported rather than adjusted away. Three of twelve geometric claims in the first worked case were left uncertified, and they are the three that sound most impressive; the one claim that required physical verification failed on first build, by 1.3 dB, for a reason no dimension calculation contains, and was fixed for the price of 0.4 m of plywood.

These are the statements downstream papers should quote rather than paraphrase. Whether any committee is described by the storage law is a separate question, it is a postulate rather than a result, and the one preregistered confirmatory test of that law reported in this corpus did not meet its own declared margin [Curzi, 2026i].

11.13.1 Declarations

Data and code availability

No external dataset was collected for this paper. All numerical results are computed from the stated parameters under the models developed in the text, and every table can be rederived from the equations and parameter values printed beside it. The two systems worked in Section 9 are constructed examples with declared parameters, not field observations; the commissioning measurements reported in Section 9.2 are the outputs those models return for the specified construction and are not records of a building that exists. The figures are generated by the scripts named in each caption, and those scripts contain no data beyond the numbers printed in this paper and the model equations of Sections 4 to 7.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, produced the worked examples and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. No meeting was recorded, no turn-share was measured from a real group, and no participants were recruited. The consent requirement of requirement A26 — that no recording for turn-share measurement be made without the informed, renewable and revocable consent of every participant — and the constraint that no intervention may remove a participant's ability to decline are **design constraints on prospective work and are not reports of any study performed.**

Preregistration

This paper reports no preregistered test of its own that was executed on a real system. The preregistration set out in Section 9.2 is a constructed instance showing the form such a registration would take, together with the outcome the registration would classify as non-supporting; it is labelled as constructed there and here. **Where a real programme acts on H3, requirement A28 makes preregistration and a matched untreated control mandatory before treatment** [Nosek et al., 2018]. The preregistered protocol for testing the storage law against a real system, together with the confirmatory application that **failed** its second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared margin of 10 per cent — is reported in Curzi [2026i], and any applied claim resting on this paper inherits that result.

11.14 Appendix A Measurement protocol

Equipment: a phone with a calibrated or at least consistently characterised measurement application, a source capable of a loud broadband burst, and a tape measure. **All eight measurements are achievable in ninety minutes by one person**, and requirement A23 makes all eight mandatory at commissioning.

M1 — Volume and surface. Measure and compute `V` and `S`. Record glazing area separately.

M2 — `RT60`, unoccupied, per octave band. Use the interrupted-noise method or an impulse. Take at least three source positions and three microphone positions — nine measurements — and average. Report per band from 125 Hz to 4 kHz. **Report the standard deviation across positions; if it exceeds about 15 per cent the field is not diffuse and Sabine does not describe the room.**

M3 — `RT60`, occupied. Repeat with the intended occupancy, or compute it by back-calculating `A` from the measured unoccupied `RT60`, adding the occupants' equivalent absorption, and recomputing. **State which was done.**

M4 — Background, unoccupied, services running. Sixty seconds of `L_Aeq` plus an octave-band spectrum, taken at the seat and not at the door. **Run it at the time of day the room is used.**

M5 — Spatial decay. Place a source at one seat and measure `L_Aeq` at every other seat with the same stimulus. **A well-treated speech room shows a clear fall with distance; a reverberant room shows almost none.** This is the most diagnostic five minutes available and requires no calibration at all, because it measures a difference.

M6 — Turn share. Record one real meeting **with the consent of every participant (A26)**. Compute per-person speaking time and `N_eff_turn/n`.

M7 — Vocal level spread. From the same recording, per-talker level normalised to 1 m. Report the 10th and 90th percentiles.

M8 — Hand clap. Stand at each seat and clap once. Listen for a discrete echo or a ringing tone. **This costs four seconds per seat and finds flutter echo and focusing that no `RT60` number will reveal. Do it first.**

Report all eight on one page with dates, positions and equipment, against the corpus's evidential standard [Curzi, 2026].

11.15 Appendix B Failure-mode catalogue

FG1 — Fractal decoration. *Damages:* the credibility of every other number in the specification. Self-similar ornament applied to a surface and described as if it delivered performance: recesses too small to hold anything, relief too shallow to scatter where it matters, a dimension quoted with no measurement behind it. *Signature:* self-similarity appears in the design narrative and no Tier 3 method appears anywhere; perimeter is quoted without a scale; occupancy at the fine scale is zero. *Detect with:* G10 and G15 — ask which of the four methods was used, and if the answer is a rendering the claim is decoration. *Fix:* commission the model, or delete the claim and keep the ornament as ornament. **There is nothing wrong with ornament and a great deal wrong with ornament billed as engineering.**

FG2 — The concave room. *Damages:* `u`, and the exclusion component of `C`. Domed, vaulted or curved surfaces facing the occupied zone, concentrating sound at focal points and producing seats that are acoustically privileged and seats that are acoustically dead. *Signature:* occupants report that one end of the room is inaudible and one participant is always too loud; a level sweep with a fixed source shows variation exceeding 6 dB between seats at equal distance; a discrete echo is audible on a hand clap at specific positions. *Detect with:* a hand clap and a walk, then a level sweep with a calibrated source. G9 should have prevented it. *Fix:* diffuse or absorb the concave surface. **Do not attempt to balance it with electronics — that adds a permanent cost line to fix a geometry error and it fails whenever the power does.**

FG3 — The frequency floor ignored. *Damages:* trust, and capital. A small reflector, screen or pod specified to control speech, whose aperture is far smaller than the wavelengths in question. *Signature:* measured level difference across the device is under 3 dB below 1 kHz while the specification claims twelve. *Detect with:* `f_floor = 2c/D`, computed in ten seconds (G5). *Fix:* enlarge, or re-describe honestly as a consonant-band and privacy device, or replace with absorption.

FG4 — The unreachable rung. *Damages:* `N`, and the honesty of the category audit. A nested plan whose fine scale is not usable — recesses nobody can sit in, alcoves whose entry is blocked, a quiet corner with no seat. The plan claims four occupancy scales and delivers two. *Signature:* the category audit strikes rows the drawings claimed; time-lapse shows the fine scale unoccupied at all times. *Fix:* G2's dimensional floor, applied before construction rather than after.

FG5 — The concealed recess. *Damages:* safety, and therefore everything. Boundary multiplication produces recesses with no sightline from any circulation route. *Signature:* occupants — reliably, a non-random subset of them — avoid a whole zone, and the zone appears clean in the count sheets because nobody is in it. *Detect with:* sightline analysis on the plan, plus asking. *Fix:* G13. Every recess visible from a route [Jacobs, 1961]. **A count sheet cannot distinguish "unused because unattractive" from "unused because avoided", so do not diagnose from the count alone.**

FG6 — Scale inflation. *Damages:* the model's credibility. Applying the nested-ladder modelling choice past the point where its parameters were measured — claiming that because the desk and the room obey the same recursion, a city does too, with parameters assumed rather than surveyed. *Signature:* a parameter table with no measurement dates and no methods. *Fix:* measure or stop. **The scale-free form of the law is a licence to model at a new scale; it is not a licence to skip the measurement at that scale.**

FG7 — The passive device with nothing to concentrate. *Damages:* nothing directly; wastes capital and misdirects diagnosis. Denge's actual lesson, misapplied: excellent geometry installed in a system whose surplus is near zero — a beautifully sub-divided floor for a group at `u = 0`, or a listening niche where nobody has anything they are permitted to say. *Signature:* all geometric requirements pass, `S` and `eta` are unchanged, and `h` does not move. *Fix:* **stop**

building and find out why there is no signal. Geometry sets coefficients; it does not supply the input.

FA1 — The reverberant room. *Damages:* `N`, `u`, exclusion. `RT60` far above target; every seat deep in the reverberant field; the tail sets a minimum gap longer than a natural turn. *Signature:* `N_eff_turn/n` below 0.6; participants report being talked over; meetings run long. *Detect with:* M2, M5, M8, M6. *Fix:* the specification of Section 8.2. **Expect `N_eff_turn/n` to move; if it does not, H2 has failed for that room and the inequality is not acoustic.**

FA2 — Passed the number, failed the room. *Damages:* the effort component of `C`, and the credibility of acoustic specification generally. State B: a single-number 500 Hz `RT60` inside target with a bass ratio near 1.5 and a flatness ratio above 2. *Signature:* occupants describe the room as boomy, or as one where you can hear but not understand. *Detect with:* M2 **per band** — a single-number measurement cannot detect this and never will. *Fix:* A2 and A3; usually remove treble absorption and add low-frequency-capable absorption.

FA3 — The dead room. *Damages:* effort, and the quiet talker specifically. State D: over-treatment, no reverberant support, talkers push, vocal fatigue rises. *Signature:* `RT60` below 0.30 s at mid frequencies; occupants report the room as oppressive. *Detect with:* M2 per band, with Eyring (A5) — **using Sabine here under-reports the deadness by up to 78 per cent.** *Fix:* remove absorption from the ceiling centre (A10) and from one wall; add a hard reflector.

FA4 — Ceiling-only treatment. *Damages:* the diffuse-field assumption, and hence every computed number. All absorption in a suspended ceiling with four hard walls, so the horizontal decay path never meets it. *Signature:* measured `RT60` well above calculated; large standard deviation across microphone positions. *Detect with:* M2 with position statistics, and M8. *Fix:* A8 — move a third of the absorption onto two non-parallel walls. **The total `A` may not change and the room will sound different.**

FA5 — The Lombard spiral. *Damages:* effort, and interruption. Multiple conversation groups in one acoustic volume, each raising level in response to the others. *Signature:* level rises through the session and does not fall; occupants leave hoarse. *Detect with:* `L_Aeq` logged against occupancy over a session. *Fix:* sub-divide the volume; cut the seed disturbance first, since **each decibel removed at source removes `1/(1−s)` decibels of eventual level.**

FA6 — The services decision nobody reviewed. *Damages:* effort, accessibility. An air-handling selection, a diffuser type or a fan speed set on thermal and cost grounds alone. *Signature:* background above 40 dBA at seat height with services running, often with a tonal component. *Detect with:* M4, **at the time of day the room is used.** *Fix:* A14 to A17 — slower, larger diffusers; attenuators; duct lining.

FA7 — The table that is too long. *Damages:* `u`, exclusion. A table with a corner-to-corner path of 4–6 m, whose end seats are in a permanently worse acoustic position. *Signature:* turn share correlates with distance from the table's midpoint. *Detect with:* A20's ratio, plus M6 turn-share cross-plotted against seat position. *Fix:* A19. **Replace the table, not the people.** This is usually the cheapest single acoustic intervention available.

FA8 — The uncounted exclusion. *Damages:* `N`, and the honesty of every metric the room produces. The room silences participants and the silenced contributions leave no trace, so the room's own records show a functioning meeting. *Signature:* a subgroup consistently at the bottom of the turn-share distribution across many meetings. *Detect with:* Section 7.6's exclusion calculation — from physics, not from asking. *Fix:* the specification. **And A27: do not let the acoustic fix take credit for a social problem it did not solve.**

Chapter 12

Thresholds, Substrate, and the Table With No Head: Three Built-Environment Levers Priced Against the Storage Law

Abstract. Three interventions in physical setting are given quantitative form and costed. First, a threshold is modelled as impedance matching: a single door between zones of very different character is the reflective analogue of a 1-to-100 junction, and subdividing the same total change into six graded steps drops worst-junction reflection from 96% to 13%. In a 30-person office this converted 22 interruptions per person-day to 6 at an unchanged 7-minute resumption lag, cutting the switching cost from 0.321 to 0.087 of the working day; headroom rose from +0.0642 to +0.2955 and sustainable stored order from 5.185 to 9.041, for 68.8 m2, 7.6% of net floor. Second, adjustability is shown to beat optimisation: the standard comfort model's own dissatisfaction floor is 5% by design, and at a 2.0 K spread of preferred temperature, two degrees of personal authority raises satisfaction from 38% to 87%. A substrate retrofit took RT60 from 4.41 s to 0.51 s while retaining 60% of the thermally coupled soffit, moving headroom from -0.0894 to +0.2688. Third, deference is modelled as a common rotation coefficient a, giving u = 1 - a^4; a thirteen-seat ring table drove the estimated a from 0.83 to 0.47 and headroom from -0.0146 to +0.0751. The deference form is badged as a postulate, and the geometric gains are reported as mediator gains only.

Keywords: threshold; impedance matching; context switching; thermal comfort; adjustability; deference geometry; round table; edge density

Series: B. Engineering the Surplus **Paper:** CCP-012

12.1 1 Introduction

Cost enters the storage law `Q_{t+1} = (1−delta) Q_t + eta S_t − C_t` **linearly and negatively**. It is not discounted, damped or filtered: a unit of cost removes a unit of stored order at the moment it is spent, while the conversion coefficient enters multiplied against a logarithm and suffers all of that logarithm's cruelties. The headroom gauge `h_t = eta ln N − delta Q_t − C_t` reads the three terms together, and when it is negative **no arrangement of the people in a building prevents the stored order from declining** [Curzi, 2026b]. **This paper takes three**

interventions in physical setting, each of which moves one term of that gauge, gives each a quantitative form, and costs it.

The three are chosen because they are the levers a building actually offers and because each is usually treated as a matter of taste.

A threshold is a matching section. On a transmission line a wave meeting a change of characteristic impedance is partly reflected, with a coefficient that depends on the *ratio* of the two impedances and not on their difference, and intermediate sections reduce the worst junction. Divide a 1-to-100 character change into six equal-ratio steps and worst-junction reflection falls from **96.1 per cent to 13.4 per cent**. A single door between a street and a reading room is that 1-to-100 junction, and the bouncing is observable: hesitation at the door, entries followed by exits within thirty seconds, street volume for two sentences and then an over-corrected silence. **Every one of those is a reflection and every one is chargeable to cost.** On a 30-person research floor, a threshold layer took involuntary context switches from **22 to 6 per person-day** at an unchanged 7-minute resumption lag — cutting the switching cost from **0.321 to 0.087** of the working day, raising headroom from `+0.0642` to `+0.2955` and the sustainable stored order from **5.185 to 9.041** — for 68.8 m^2, **7.6 per cent of net floor.**

Adjustability beats optimisation, and the argument is arithmetic rather than sentiment. The standard thermal-comfort model predicts, at its own theoretical optimum with every variable ideal, **5 per cent dissatisfied.** The floor is not zero, and it is not a defect of the model: it is the model correctly reporting that preference is dispersed and that one environmental state cannot sit at every individual optimum. A fixed setpoint answers the question *where is the centre of the distribution*, and almost nobody lives at the centre of a distribution. At a 2.0 K standard deviation of preferred operative temperature, a single optimised setpoint satisfies **38.3 per cent**; two degrees of personal authority over one's own conditions satisfies **86.6 per cent. No amount of re-optimising the single setpoint reaches that, because the problem is the dispersion and not the location.** A substrate retrofit in a four-person workroom took reverberation from **4.41 s to 0.51 s while retaining 60 per cent of the thermally coupled soffit** — the two requirements are compatible and are almost never coordinated — and moved headroom from `−0.0894` to `+0.2688`.

Deference is quartic. Modelling each member's contribution as a rotation toward the chair's by a common coefficient `a` gives pairwise non-parallelism `u = 1 − a^4`, so a twelve-member assembly keeps **97.6 per cent** of its ceiling at `a = 0.5`, **63.0 per cent** at `a = 0.9` and **44.8 per cent** at `a = 0.95`. **Ordinary politeness is nearly free and strong deference is catastrophic**, and nothing in the paperwork records the loss: the minute still lists twelve names. Geometric equality has an exact computable definition — the variance across seats of the number of other seats within a comfortable head turn — and it is **zero for a rotationally symmetric arrangement and 9.06 for an eighteen-place boardroom table**, at which the two head seats each command fifteen of the other seventeen and the far long-side seats command five. Replacing a thirteen-member council's rectangular table with a thirteen-sided ring took centrality variance from **2.095 to 0.00**, the effective number of speakers from 0.53 to 0.83 of the roster, the estimated alignment from **0.83 to 0.47**, and headroom from `−0.0146` to `+0.0751`.

The paper's principal restraint is a refusal. Equality of conversational turn-taking is a reported correlate of a collective-intelligence factor, and the causal direction is not established. A preregistered test of the acoustic route to that factor was run elsewhere in this series: the mediator moved decisively — turn-share evenness 0.604 to 0.812 with a matched control room flat at 0.638 to 0.650 — and the preregistered group-performance outcome did not reach its threshold [Curzi, 2026e]. **Wherever this paper buys turn-taking equality with geometry, it is buying the measured mediator and nothing more, and it forbids reporting a furniture or acoustics change as a collective-intelligence intervention.** The same restraint is applied to air quality, where the ventilation case is well supported and the cognitive-performance case is not: **specify ventilation, monitor carbon dioxide, and claim only ventilation.**

Sections 2 and 3 state the hypotheses and the framework, including the identification problem. Section 4 develops thresholds and edge, Section 5 the substrate, Section 6 the deference geometry. Section 7 collects the design specifications. Section 8 works three systems end to

end. Sections 9 to 12 discuss, verify, limit and conclude. Appendix A gives the measurement protocols and Appendix B the failure-mode catalogue.

12.2 2 Hypotheses

Some results below are theorems about an equation or exact consequences of a stated model, and they cannot be falsified by observation; they are labelled as derivations where they appear. Others claim that these models describe a real building, and those are stated here with what would defeat them.

H1 — Transition cost falls with graded intermediate zones. For a boundary between zones `A` and `B`, the dissipation charged per period is `C_trans = k m(Z_A, Z_B, d)`, with `k` crossings per period, `m` the per-crossing cost, `Z` the zones' character vectors and `d` the device's depth; `m` rises with mismatch and falls with depth. **Predicted:** subdividing a large character change into steps of one point each on the eight-attribute scale of Section 4.3 reduces measured crossing cost, and in the worked floor takes `k` from 22 to 6 per person-day. **Falsified by:** `k` above 15 at eight weeks after a threshold intervention that met its mismatch targets, which locates the interruptions somewhere other than the boundary.

H2 — Occupants select edge positions before open-field positions. Predicted: occupancy counts by position show perimeter and interior-island edges occupied at materially higher rates than open floor; in the worked floor, perimeter desks were occupied 94 per cent of working hours against 38 per cent for interior desks. **Falsified by:** occupancy counts showing no edge preference in a building where edge positions are otherwise comparable in light, noise and distance from amenity. **This is an empirical regularity reported from systematic observation [Gehl, 1971, Whyte, 1980], not a theorem, and it replicates or fails in one afternoon.**

H3 — Distributed individual control outperforms an optimised fixed value. For a comfort variable over which individual preferences are dispersed, the fraction of occupants outside their acceptable band is lower under personal control than under any single setpoint. **Predicted:** at a 2.0 K dispersion of preferred operative temperature and a `+/−1 K` personal acceptance band, a fixed setpoint satisfies 38.3 per cent and `+/−2 K` of personal authority satisfies 86.6 per cent. **Falsified by:** measured satisfaction under personal control no higher than under an optimised fixed setpoint in a population with measured dispersion above 1.5 K. **The dispersion is an assumed parameter and must be measured; the prediction is conditional on it.**

H4 — Deference acts as a common rotation, so its damage is quartic in alignment. Modelling the contribution actually made under deference as `v_i = a h + sqrt(1−a^2) w_i` gives `u = 1 − a^4` and `N_eff = 1 + (n−1)(1 − a^4)`. **Predicted:** an intervention that reduces the estimated alignment coefficient raises measured disagreement and speaking-share evenness together, and the assembly's realised surplus tracks `ln(1 + (n−1)(1−a^4))`. Losses are negligible below `a = 0.7` and severe above `a = 0.9`. **Falsified by:** measured disagreement that moves linearly rather than quartically with an independently estimated alignment, or an alignment estimate that does not fall after the geometric and protocol intervention. **The FORM of deference is a postulate with no external support; what follows from it is exact arithmetic.**

H5 — Geometric equality is computable from a plan and predicts speaking share. Defining `V_i` as the number of other seats within a 45-degree head turn of seat `i`'s forward gaze, geometric equality is the condition `Var(V_i) = 0`. **Predicted:** across layouts, `Var(V_i)` correlates with the inequality of measured speaking share; the

worked council moved from a variance of 2.095 to 0.00 and from an effective-speaker ratio of 0.53 to 0.83. **Falsified by:** speaking-share inequality uncorrelated with computed centrality variance across a set of layouts, or a rotationally symmetric table at which measured shares stay as unequal as at the rectangle it replaced.

Three claims in this paper are **not** hypotheses. That reflection on a transmission line is a useful description of a built threshold is an **analogy**, badged at every use, and Section 4.2 states exactly which two sentences of it transfer. That a material whose surface reads better after wear has a lower appearance-decay rate is a **modelling interpretation**, not a finding. That a constraint embodied in furniture is enforced on every occasion of use at zero marginal administrative cost is close to a **definition** of what furniture is, and its limit — geometry raises the cost of dominance and does not forbid it — is stated in Section 6.6.

12.3 3 Framework

12.3.1 3.1 What is presupposed

The storage law and its gauge. The equation acted on throughout is

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,     S_t = F(U_t) = ln(1 + (N−1) U_t)
```

with the headroom gauge `h_t = eta ln N − delta Q_t − C_t`, its three regimes, the maximum one-step increment `r_max = h_t`, and the two distinct constraints of admissibility `C_t <= (1−delta)Q_t + eta S_t` and sustainability `C_t <= eta ln N − delta Q_t` [Curzi, 2026b]. The steady state is `Q_inf = (eta F(u*) − c*)/delta`, **which the result allows to be negative and does not hide when it is.** The exact transient is `Q_t = (1−delta)^t Q_0 + (1 − (1−delta)^t) Q_inf`, a first-order step response with time constant `1/delta`; the framework's own correction to the naive reading of that constant — that the half-life is `t_half = ln 2 / −ln(1−delta)`, which at `delta = 0.30` is 1.9434 cycles against the naive `1/delta = 3.3333`, an error of 72 per cent — is used wherever a settling time is quoted [Curzi, 2026]. The level–rate–delay vocabulary long predates this framework [Forrester, 1961], and headroom is its name for organisational slack [Cyert and March, 1963].

Convergence. Every trajectory in this paper approaches its fixed point by the contraction argument: a map of modulus `q < 1` on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least `q` each pass [Banach, 1922, Curzi, 2026]. That is why Sections 8.1 to 8.3 insist that cost be evaluated at eight weeks and stored order at a year: **the drive responds immediately and the state responds at the time constant, so measuring both on one schedule guarantees that a succeeding intervention is judged a failure.**

The surplus functional. `u(x,y) = 1 − (x.y)^2` in `[0,1]` and `f(u) = ln(1 + (N−1)u)`, with `0 <= F(u) <= ln N`, `F = 0` exactly at parallel and `F = ln N` exactly at orthogonal [Curzi, 2026]. **The affine effective count `N_eff = 1 + (N−1)u` is a design axiom: chosen, not derived.** Concavity means the second genuinely different member is the largest single gain an assembly will ever buy and the fortieth is nearly free of value; the logarithmic ceiling means doubling the number of distinct categories adds `ln 2 = 0.6931` wherever one starts, **so a committee cannot be grown out of a deference problem.** Where a lever moves `u`, the framework prices crossing and divergence identically — `df/du_cross = df/du_div = (N−1)/(1 + (N−1)u)` — while their productivity in the underlying design variables differs [Curzi, 2026d], and Section 6 moves `u` by removing a rotation rather than by recruiting difference.

The comfort construct. The framework's founding move is that comfort is the felt total of present costs, so that the present state is a local minimum of the comfort landscape **by definition rather than by luck** [Curzi, 2026]. **That is what makes a cost-side intervention the natural first move on a building**: raising the cost of the current arrangement is not available as a design action, but lowering the cost of a better one is.

Two effective counts appear here and they are different objects. `N_eff = 1 + (N−1)u` is the surplus law's affine count and **it is not a Hill number** — it is a linear interpolation between 1 and `N` and belongs to no order of the family whose canonical members are `1/HHI` and `exp(H)` [Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979]. The effective number of speakers `N_speak = exp(−sum_i p_i ln p_i)` used in Section 6.5 **is** the Hill number of order one, the exponential of Shannon entropy [Shannon, 1948, Pielou, 1966], borrowed as a descriptive diversity index. **It measures how equally the floor was shared. It does not compute surplus, and the resemblance of form is a resemblance and not a theorem.** The order-two member of the same family, `1/HHI`, with its worked case returning 3.41 effective jurisdictions on naive shares against 1.00 after bloc collapse, is a third quantity again [Curzi, 2026].

12.3.2 3.2 Identification: what no amount of measurement recovers

This limitation is structural rather than statistical and belongs in the method. Every mechanism in this paper runs through a group whose members both influence and constitute the aggregate: the room shapes who speaks and defers, who speaks and defers shapes what the group is, and the group is made of its members. That is a linear-in-means reflexive specification, and Manski [1993] proves that in such a specification **the endogenous social effect is not identified from observational data** — it cannot be separated from contextual effects or from correlated unobservables however much data is collected. Shalizi and Thomas [2011] reinforce it for network data, where influence and selection are generically confounded.

What survives and what does not. The impedance arithmetic, the edge-density arithmetic, the deceleration and dwell calculations, the thermal-mass and effusivity computations, the Sabine arithmetic, the head-turn geometry and the centrality variance are physics and combinatorics computed from plans and materials; no identification problem arises for any of them. **The interpretation of the resulting change in `u`, or in the alignment coefficient `a`, as an endogenous social effect is what is not identified.** The alignment coefficient is worse placed still: it is a *modelled* quantity, not an observed one, and Section 6.2's protocol returns an interval whose lower end rests on attributing unexplained convergence to deference — **a judgement made by the people being measured.**

A partial remedy is adopted and its price is paid. Each worked case in Section 8 carries a matched control receiving equal cost and disruption with no change to the mechanism under test. That removes the most immediate confound — that measuring a group changes it, the hazard McCambridge et al. [2014] documents systematically and Rosenthal and Jacobson [1968] demonstrates for expectancy — and it is the requirement most often omitted from a building evaluation. **It does not solve the reflection problem.** The route that would is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001]; the fixed rosters and declared category sets of these surveys point toward it without taking it. **Until then, every applied statement in this paper about a group is a modelling position rather than an identified empirical result**, and Durlauf [2005] states the demand that remains open: name the observation that would distinguish this account from a neighbouring one fitting equally well.

12.3.3 3.3 Analogy discipline

The transmission-line result is an analogy and is badged as one. Two sentences transfer: mismatch is a *ratio* rather than a difference, and intermediate sections reduce the worst junction. **Nothing else transfers**, and Section 4.2 states the exclusion flatly: the transmission-line results depend on wave phase, a quarter-wave transformer achieves zero reflection only where the section is a quarter wavelength long, and **nothing in a corridor has a phase.** Do not compute a corridor length from a wavelength; there is no wavelength.

The first-order step response is a shared mathematical form. The storage law, a leaky integrator and a discrete low-pass section are governed by the same difference equation, so

every theorem about first-order systems applies to all of them [Oppenheim et al., 1996, Åström and Murray, 2008]. **Nothing about charge or capacitance is claimed to hold of an institution.**

Sabine's equation, the thermal-diffusion depth and the effusivity contact temperature are derivations, used within their own domains. They are statements about air, heat and materials, applied to air, heat and materials [Sabine, 1922, Rossing et al., 2002]; their premises are stated where they are used and the reverberation calculation's five conditions are inherited from Curzi [2026e].

The entropy in `N_speak` is Shannon entropy in nats, not thermodynamic entropy, and surplus is identified with neither. Mirowski [1989] is the standing warning about a social science that copies an energetics for which it has no conservation law, and Cartwright [1999] supplies the escape this paper takes: specify the arrangement within which a law is claimed to hold — the stated domain of Section 10.5 — and stop claiming a universality the law does not need.

12.3.4 3.4 Standing honesty clauses

- **As deployed, the surplus functional is a lower bound** on the ensemble's information [Curzi, 2026]. Where this paper says a threshold protects `u`, it protects a lower bound on `u`.
- **Subadditivity holds on the parameter `u` only; the triangle inequality fails for geometric triples.** Threshold mismatches must not be chained as though they added like distances, and psychological similarity is in any case known to be asymmetric and to violate the triangle inequality [Tversky, 1977].
- **The law prices a saved unit of cost; it does not establish that a porch saves one.** Every claim that a physical intervention moves a parameter is an empirical claim requiring its own measurement, and the worked cases in Section 8 declare which measurements were taken.
- **The cost observable is a calibration choice, declared and not derived.** All three worked cases express cost as `C = phi eta ln N`, with `phi` a dimensionless fraction of the ceiling. **That is a declaration, not a derivation**, and Section 10.7 records what it costs in interpretability.
- **Nominal material values are nominal.** Absorption coefficients, conductivities and heat capacities here are typical published values for generic descriptions, adequate for design arithmetic and **not adequate for a construction specification**, where the manufacturer's tested data at the intended mounting condition is required.
- **Numbers taken from standards are engineering-practice values that change between editions and jurisdictions.** This paper is about method and is not a substitute for a code; where a figure is jurisdictional it says so.
- **No provision here may be applied so as to compel a person.** A threshold decelerates and does not deny; a speaking protocol carries an unconditional pass; no measurement of a meeting is made without the informed consent of every participant. **The arithmetic gives no warning where that boundary is crossed.**

12.4 4 Thresholds, edge, and the cost of crossing

12.4.1 4.1 Edge scarcity and the arithmetic of edge manufacture

Occupancy is not uniform across a floor. Under systematic observation, people take up positions at edges — walls, arcade columns, steps, planters, the boundary between paving and grass — before positions in open ground [Gehl, 1971, Whyte, 1980]. The mechanism usually offered is prospect and refuge: back covered, view forward and wide, exit one step away, and *partial* engagement, which is the only kind most people want most of the time. **That mechanism is a hypothesis; the occupancy pattern is an observation, and it can be confirmed or refuted in one building in an afternoon.**

The distinction that makes edge worth counting is between activity that happens regardless of the space's quality and activity that happens only when invited [Gehl, 1971]. Necessary activity — going to work, waiting, carrying things — is insensitive to the setting and therefore tells an analyst nothing. **Optional activity is discretionary presence, and it is the observable proxy for conversion: it is the substrate from which durable relation and durable decision get made.** A room that is not occupied by choice converts nothing, whatever its floor area.

The engineering consequence is arithmetic. Perimeter scales with the first power of linear dimension and area with the second, so the amenity people select for gets scarcer as the room gets bigger:

```
square room, side s:   edge density  P/A = 4s / s^2 = 4/s   metres of edge per square metre
```

Side (m)	4	6	8	10	15	20	30
Area (m^2)	16	36	64	100	225	400	900
Edge density (m/m^2)	1.000	0.667	0.500	0.400	0.267	0.200	0.133

A 30 m square hall has one-seventh the edge per square metre of a 4 m room. That is the structural reason large open-plan floors perform worse than their area predicts: the amenity has been diluted sevenfold while the rent per head stayed constant, and furniture quality does not compensate, **because the deficit is topological.**

The fix is not small rooms but manufactured interior edge, and the arithmetic is unusually favourable, because an interior island — bookcase run, planter bank, service core, low wall, level change, storage bank — contributes its full perimeter while consuming only its footprint. For a 10 x 10 m room, area 100 m^2, perimeter 40 m:

Intervention	Edge added	Gain	Floor consumed	Edge per m^2 consumed
1 island 4.0 x 4.0 m	+16.0 m	+40%	16.0 m^2 (16.0%)	1.00
2 islands 2.0 x 4.0 m	+24.0 m	+60%	16.0 m^2 (16.0%)	1.50
4 islands 1.2 x 3.0 m	+33.6 m	+84%	14.4 m^2 (14.4%)	2.33
6 islands 0.9 x 2.4 m	+39.6 m	+99%	13.0 m^2 (13.0%)	3.06

Six thin islands nearly double the room's edge for 13 per cent of its floor; one fat island buys 40 per cent for 16 per cent. Edge yield is maximised by many thin obstructions, not few thick ones. The limit is circulation clearance: below about 0.9 m clear between islands the gain is cancelled by congestion, which is itself a cost term. And the yield depends on the base: on a 12 x 18 m room, area 216 m^2 and perimeter 60 m, the same six islands add 66 per cent rather than 99 per cent. **Edge manufacture is most valuable exactly where edge is scarcest, which is the deep plan — and that is where it is most often omitted, because islands read as lost area on the same spreadsheet that produced the deep plan.**

12.4.2 4.2 The threshold as a matching section

On a transmission line, a wave in a medium of characteristic impedance `Z_1` meeting a medium of impedance `Z_2` is partly reflected:

```
Gamma = (Z_2 - Z_1) / (Z_2 + Z_1)     power reflected = |Gamma|^2, transmitted = 1 - |Gamma|^2
```

Ratio `Z_2/Z_1`	2	4	6	10	100	1000
`Gamma`	0.333	0.600	0.714	0.818	0.980	0.998
Power transmitted	**88.9%**	**64.0%**	**49.0%**	**33.1%**	**3.9%**	**0.4%**

Two structural facts transfer and nothing else does.

Fact 1 — reflection depends on the RATIO, not the difference. A step from 1 to 2 and one from 500 to 1000 reflect identically; absolute character is irrelevant and only the jump matters.

Fact 2 — intermediate sections reduce the worst junction. Take the 1-to-100 case and insert one section at the geometric mean `Z_t = sqrt(Z_1 Z_2) = 10`. The junctions 1-to-10 and 10-to-100 now have *identical* coefficients of 0.818 — **the geometric mean is the value that equalises them.** Subdividing further with a constant ratio per step `r = (Z_2/Z_1)^(1/n)`:

Steps `n`	1	2	3	4	5	6
Ratio per step	100.00	10.00	4.64	3.16	2.51	2.15
`Gamma` per junction	0.980	0.818	0.646	0.520	0.431	0.366
Reflected per junction	96.1%	66.9%	41.7%	27.0%	18.5%	13.4%

Worst-junction reflection falls from 96.1 per cent to 13.4 per cent by dividing the same total change into six graded steps, and in the limit a tapered line reflects almost nothing. A threshold is a matching section. That is the whole of the idea and the whole of what the analogy is for.

What does NOT transfer, stated flatly. The transmission-line results depend on wave phase: a quarter-wave transformer achieves *zero* reflection only where the section is a quarter wavelength long, and tapered designs are broadband because they trade that null for tolerable performance across a band. **Nothing in a corridor has a phase. Do not compute a corridor length from a wavelength; there is no wavelength.** Two sentences transfer, and anyone extending the analogy past them has left engineering for decoration. This is failure mode **FT1** of Appendix B.

12.4.3 4.3 The zone character vector and the one-step rule

Using the analogy requires an operational `Z`: a vector of ordinal attributes, each rated 1 to 5, each independently observable.

Attribute	1	3	5
Sound level	library quiet	conversational	machinery, crowd
Pace	seated, hours	task, minutes	transit, seconds
Formality	domestic	working	ceremonial, on-record
Illuminance	low, warm, local	mixed	high, uniform, cool
Ownership	one named person	a known group	anyone
Attention mode	deep, single-thread	interruptible	scanning
Body state	shoes off, seated	standing, coat on	moving, loaded
Observability	unseen	seen by peers	seen by strangers

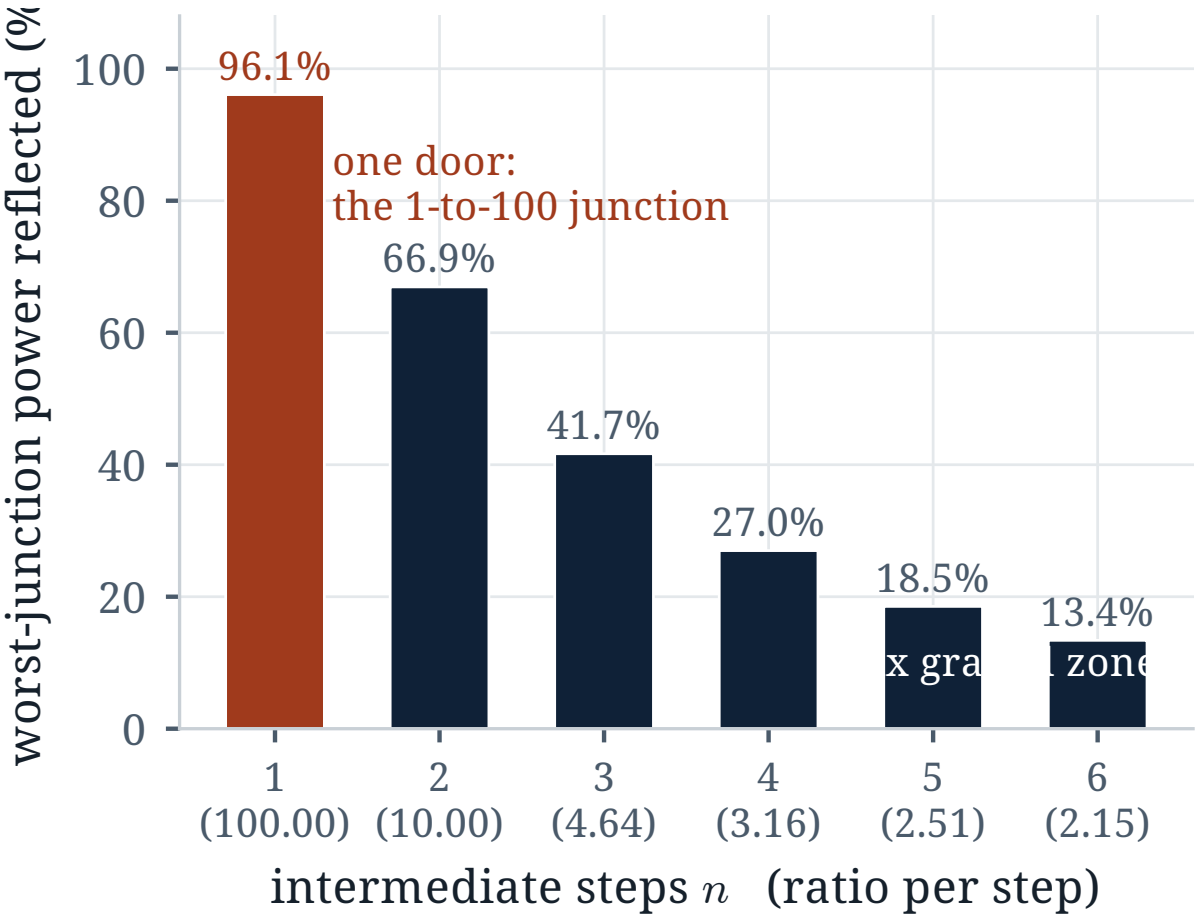


Figure 12.1: **Six graded zones transmit what one door rejects.** Worst-junction reflected power across a 1-to-100 character change, divided into `n` equal-ratio steps: 96.1 per cent at a single junction against 13.4 per cent at six. Exact arithmetic on the transmission-line reflection coefficient `Gamma = (Z_2−Z_1)/(Z_2+Z_1)`; the transfer of that coefficient to a built threshold is a badged analogy and not a derivation. Generated by `figs/src/CCP-012_fig01.py`.

Rate both zones. The **mismatch index** is `M = max_i |Z_A[i] − Z_B[i]|`, with `M_bar` the mean of the same differences. **`M` is what matters:** a transition can be graded on seven attributes and violent on the eighth, and the eighth is what people feel. Treat `M >= 3` as requiring a device whatever `M_bar` says.

Rule T-A — one step per threshold. No threshold may change any attribute by more than one point on the five-point scale; at `M = 4`, at least three intermediate zones are required. **This is the direct translation of the taper table and it is the single most useful rule in this section.**

The scoring is an ordinal steepness judgement, and the corpus already runs one of those: the comfort-weighted amplitude scores each step 0 to 3 for steepness, sums to a resistance `R`, and weights the branch as `1/2^R`, so that a three-point difference in steepness is a factor of eight in expectation [Curzi, 2026]. **The mismatch index is the same kind of instrument and inherits the same limitation: it is ordinal, and the mean `M_bar` is a statistic its scale does not license.** That is why the rule is written on the maximum and not on the mean.

Worked instance, street to reading room:

Attribute	Sound	Pace	Formality	Illum.	Ownership	Attention	Body	Observ.
Street	5	5	3	5	5	5	5	5
Reading room	1	1	2	2	2	1	1	2
Gap	4	4	1	3	3	4	4	3

`M = 4`: **a single door here is the 1-to-100 junction, and the bouncing is observable** — hesitation at the door, entries followed by exits within thirty seconds, standing in the entry looking for permission, street volume for two sentences and then an over-corrected silence, or simply not going in. Every one is a reflection chargeable to cost. The graded version inserts four intermediate zones, one step each:

```
street          forecourt       porch         lobby         anteroom      reading room
sound   5  -->      4      -->     3      -->     2      -->     2      -->       1
pace    5  -->      4      -->     3      -->     2      -->     2      -->       1
attn    5  -->      4      -->     3      -->     3      -->     2      -->       1
body    5  -->      4      -->     3      -->     2      -->     2      -->       1
own     5  -->      5      -->     4      -->     3      -->     2      -->       2
```

Each junction now has `M = 1` and worst reflection collapses without a single change to the reading room itself. **Notice what the intermediate zones physically are:** a forecourt changes surface and width; a porch, overhead and light; a lobby, acoustic and ownership; an anteroom, pace and observability. **Each changes one variable, and that is what makes it an intermediate zone rather than a small room.**

12.4.4 4.4 The entrance sequence: a duty list and the depth it forces

A threshold is a machine with a duty list, and its depth is set by the slowest item on it. **(1) Decelerate**, from transit speed to standing. **(2) Reorient** — find the next direction from signage or geometry. **(3) Shed:** coat, umbrella, shoes, bag, load, phone call. **(4) Be seen** by whoever is inside, before being among them. **(5) See** — assess the room before committing to it, the reciprocal of (4) and the item most often omitted. **(6) Choose:** go in, wait or leave, none of the three conspicuous. **(7) Be greeted**, or deliberately not, but by architecture rather than by accident. **(8) Store** the shed items where they will still be at the end.

Every classical entrance device solves some subset, and they are worth naming because they are convergent evolution and therefore evidence.

Device	Duty items served
Japanese *genkan* — floor-level change, shoes removed and stored	1, 3, 8; marker of 4 and 6
Japanese *engawa* — covered verandah, neither in nor out	1, 5, 6; partial commitment
Roman *vestibulum* to *fauces* to *atrium*	a literal three-section taper
Christian narthex — antechamber before the nave	4, 5, 6; gradient of belonging
Bent entrance, never on axis	1, 5, 8 at once
Porch with bench [Alexander et al., 1977]	1, 6, 7
Industrial airlock — two doors never open together	1 to 8, made hardware

None of these is decoration. Each is a matching section.

Deceleration sets the minimum depth. Comfortable walking speed on the flat is about 1.2–1.4 m/s for adults and slower for children, the elderly and the loaded; a comfortable voluntary deceleration — one that does not read as a stumble or a halt — is on the order of 0.5 m/s^2. Then `d = v^2/(2a)` and `t = v/a`:

`v` (m/s)	1.2	1.4	1.4	1.4	1.6
`a` (m/s^2)	0.5	0.3	0.5	1.0	0.5
Stopping depth `d` (m)	1.44	3.27	**1.96**	0.98	2.56
Time `t` (s)	2.40	4.67	2.80	1.40	3.20

A threshold that expects people to stop must be at least 2.0 m deep in the direction of travel. Below that the stop is achieved by the walker's own braking effort rather than by the space, and it registers as being brought up short. **This is why a 1.2 m vestibule feels wrong even when correctly sized for the door swing: the door was designed and the deceleration was not.**

Dwell time sets the useful depth. Deceleration is the floor; the threshold's real job is to hold a person long enough for the state change, and the currency of that is time, not metres:

Depth `d` (m)	1.5	2.0	2.5	4.0	6.0
at 1.4 m/s	1.07 s	1.43 s	1.79 s	2.86 s	4.29 s
at 0.7 m/s	2.14 s	2.86 s	3.57 s	5.71 s	8.57 s
at 0.4 m/s	3.75 s	5.00 s	6.25 s	10.00 s	15.00 s

At full walking pace a 2 m threshold buys 1.4 seconds — one breath, enough for item 1 and nothing else. To buy the four or five seconds items 2, 4 and 5 need there are two options: **six metres of depth, or slow the walker down**, at which point 2.5 m buys 6.25 s. **That is the case for the bent entrance and the reason it recurs across unrelated traditions.** A 90-degree turn cannot be taken at 1.4 m/s; it forces a drop to roughly 0.7 m/s, converting 1.5 m of path into about 2.1 s instead of 1.1 s, **while delivering item 5 — the interior revealed progressively, as one turns — and item 4, being seen approaching before being present. The turn is worth more than the metres it costs.**

The other dimensions. A threshold a person can pause in must let another pass, or the pause is an obstruction and the obstruction is a cost:

```
stopped person, footprint depth  0.60 m      clear width, one person passing  0.75 m
body depth (chest to back)       0.45 m      clear width, two abreast        1.20 m
```

```
shoulder breadth          0.45 – 0.50 m      door leaf 0.9 m -> swing radius  0.90 m

holds one paused person AND permits single-file passage:  0.60 + 0.75 = 1.35 m clear width
inward-swinging 0.9 m leaf, holding one paused person:    0.90 + 0.60 = 1.50 m landing depth
holds a conversing PAIR at personal distance and
  still permits passage:            0.45 + 1.20 + 0.45 = 2.10 m, + 0.75 m to pass = 2.85 m
```

The last figure matters most, because the conversation at the entrance is a large fraction of the optional activity in a whole building. **A 2.85 m entrance lobby is not generous. It is the minimum width at which two people can talk without becoming an obstacle — and if they become an obstacle they stop talking, and the optional activity that was the conversion proxy disappears.**

12.4.5 4.5 The intimacy gradient

The pattern tradition supplies the sequence [Alexander et al., 1977] and the proxemic literature the metric: for adults in the population studied, intimate distance runs 0.00–0.45 m, personal 0.45–1.20 m, social 1.20–3.60 m and public beyond 3.60 m [Hall, 1966]. **These are not universal constants and were not claimed to be; what is reliable is the ORDERING and the existence of the bands, not the exact boundaries.** Verify them in the population at hand, and expect the personal/social boundary in particular to move.

Zone	Who may enter	Without asking?	Pace	Typical device
Public	anyone	yes	transit	street, forecourt
Semi-public	anyone with business	yes	slowed	lobby, cafe, front desk
Common	members	yes	task	corridor, kitchen, shared table
Semi-private	members with reason	knock or hover	task	team room, open alcove
Private	invited	ask	deep	office, study, closed room
Intimate	chosen	ask	any	one’s own desk, one’s own chair

Rule T-B — no zone may be skipped. A door opening directly from public to private is the 1-to-100 junction and both sides pay. **The receptionist whose desk faces the street door absorbs a four-step mismatch all day, and the cost appears in the turnover rate, which is a decay term.**

Rule T-C — the gradient must be legible without signage. If a visitor needs a sign to know they have left the common zone, the gradient is not doing its work and the sign is. **Signs are administered — made, mounted, updated, enforced — and geometry is not.** Change the floor surface, the ceiling height, the light level, the width or the direction; each is read instantly by everyone, including people who cannot read the sign. The general form of the point is that infrastructure recedes from notice until it breaks [Star and Ruhleder, 1996, Star and Griesemer, 1989], and a legible gradient is infrastructure while a sign is administration.

The semilattice argument sharpens it [Alexander, 1965]. A natural settlement has overlapping units — a shop belongs to a street *and* a neighbourhood *and* a trade — while a planned one is drawn as a tree, every unit with one parent and overlaps forbidden. **A strict hierarchy drives non-parallelism toward zero, because any two units at the same level receive their context through the same channel.** Thresholds are where the semilattice is realised or destroyed: a room with one door into one corridor is a tree node; a room with a door to the corridor, a glazed panel to the courtyard and a shared alcove to its neighbour is a semilattice node.

The number of ways a room participates is the number of thresholds it has, which is how a threshold decision becomes a category-count decision.

Where the pattern tradition and this framework part company is stated plainly. **The pattern literature gives a configuration and a rationale in the language of life and wholeness; every pattern here is stated with the term of the storage law it moves and the arithmetic that prices it** — a narrower and more brittle claim, deliberately, because a claim that can be checked is worth more than one that can only be agreed with.

Pattern family	Parameter	Mechanism
Entrance transition, airlock, bent entry	`C` down	Reduces per-crossing cost `m`
Door that closes, screened alcove	`C` down	Reduces crossing count `k`
Building edge, window place, pocket	`eta` up	Occupied time where encounters become durable
Intimacy gradient	`C` down, `u` kept	Prevents the wrong-mode encounter
Semilattice adjacency	`N` up	More genuine categories of participation
Desire-line paving	`eta` up	Converts an observation into durable structure

12.4.6 4.6 Desire lines as a measured gradient, and context switching as a cost term

A desire line — the worn diagonal across the lawn, the polished floor patch, the chair that has migrated three metres — is a *measurement* with properties no survey has. It is **unprompted**, so nobody performed an answer, which is the entire problem with asking people about their building. It is **weighted by frequency**, since the depth of the trace integrates over crossings, making it literally an accumulator in the same family as stored order. It **records cost, not preference**: people route around cost, and a desire line marks where the designed path was dearer than the undesigned one. And it is free and continuously updated. **Protocol T1 in Appendix A states how to run the survey; Protocol T2 states what to do with it, and carries the caution that paving destroys the instrument that produced the trace** — a measure adopted as a target ceasing to measure what it names, arriving through the floor [Goodhart, 1975, Campbell, 1979].

The chapter's most consequential arithmetic is the switching cost, and two of its figures must be measured rather than asserted. Per period,

```
C_switch = k * m
```

with `k` the number of involuntary context switches per person per period and `m` the **resumption lag**: elapsed time from the end of the interruption until the person has regained the working state they were in before it. Both are observable — `k` a count, `m` a timed measurement. On an eight-hour day of 480 minutes:

`k` per day	`m` = 2 min	`m` = 5 min	`m` = 10 min
5	10 min (2.1%)	25 min (5.2%)	50 min (10.4%)
10	20 min (4.2%)	50 min (10.4%)	100 min (20.8%)
20	40 min (8.3%)	100 min (20.8%)	**200 min (41.7%)**
40	80 min (16.7%)	200 min (41.7%)	400 min (83.3%)

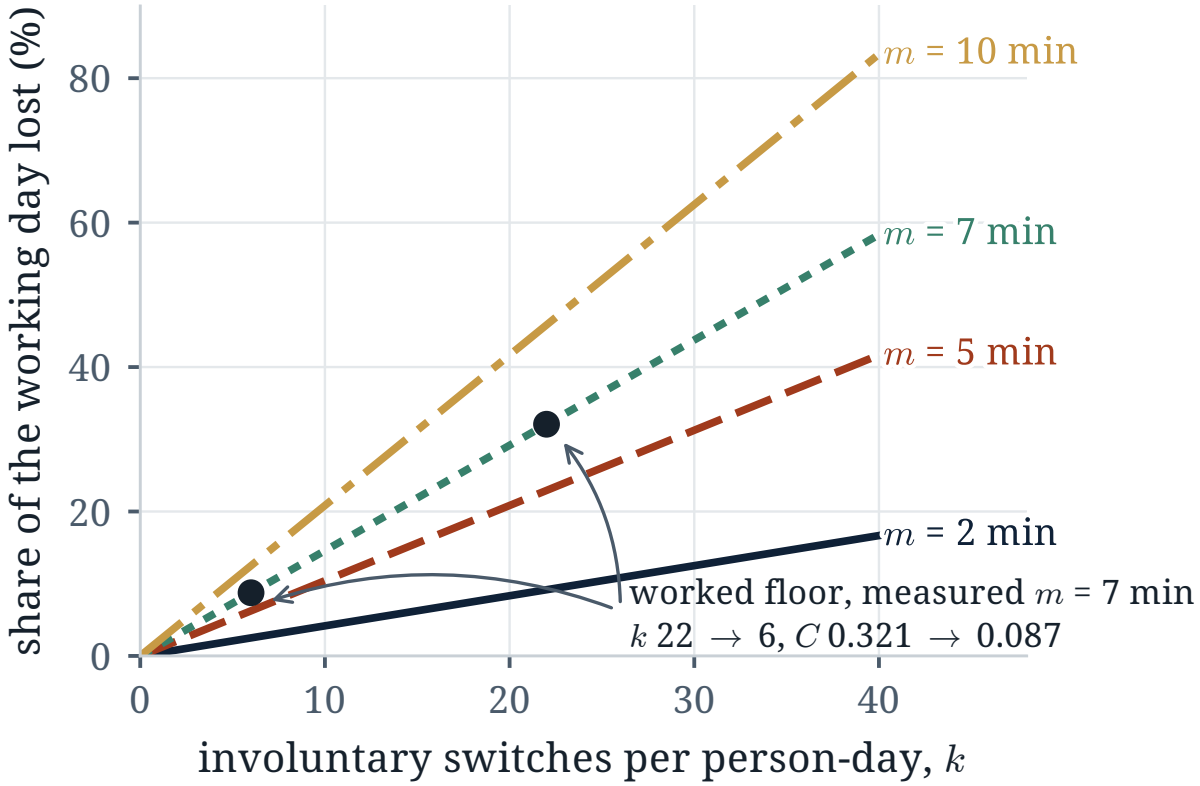


Figure 12.2: **Two levers, one product: halving either the interruption count or the resumption lag halves the switching cost.** `C_switch = k m` as a share of a 480-minute day. The worked research floor measured `m` = 7 min and moved `k` from 22 to 6, taking `C_switch` from 0.321 to 0.087 of the day. Exact arithmetic given `k` and `m`; both were measured by the declared protocols in that case and neither is asserted as a general value. Generated by `figs/src/CCP-012_fig02.py`.

Twenty interruptions a day at a ten-minute resumption lag consumes 41.7 per cent of the working day before any work has been done wrong. The product is bilinear, so halving either factor halves the cost, **and the two are different engineering problems.** `k` is a threshold problem — filters, doors, distance, visibility control. `m` is a task-state problem — externalised working state, checkpoints, a written next action, an artefact that holds the context so the person need not.

MEASUREMENT REQUIRED, NOT ASSERTED. Field studies of interrupted knowledge work report substantial resumption lags, and the finding that recovery costs materially more than the interruption itself is well established in that literature; **magnitudes vary enormously with task and interruption type, with whether the interruption is self-initiated, and with how "resumed" is operationalised. This paper supplies no universal number for `m`, because any figure quoted as universal has been detached from its measurement definition.** Protocols T3 and T4 of Appendix A give the measurements. **This paper carries no bibliography entry for that field literature and names the gap rather than attaching a citation that would not support the sentence** (Section 11).

Distance as a filter. The threshold's role in `k` is to be a filter with tunable admittance:

Device	Effect on `k`	Cost	Reversible?
Physical distance	strong	high	no
A door that closes	strong	low	instantly
Screen at 1.35–1.45 m	moderate	low	yes

Device	Effect on `k`	Cost	Reversible?
Screen at 1.20 m	weak	low	yes
Orientation to traffic	moderate	zero	yes
Signal convention	moderate	zero	needs upkeep
Timed availability windows	strong	zero	yes

Two deserve comment. **Distance is a blunt filter.** Studies of research and development organisations established the regularity that the probability of regular technical communication falls steeply with physical separation and is near its floor within a few tens of metres; **the shape is the point, not the coefficients.** The consequence cuts both ways: **the same distance that suppresses unwanted interruptions suppresses the encounters that were the conversion.** It reduces `k` and `eta` together, which places it first for strength and last for advisability. **This paper carries no bibliography entry for that finding either** and names it as an uncited borrowing.

Screen height is the cheapest single intervention here. Seated eye height is roughly 1.15–1.25 m and standing 1.55–1.70 m. A screen at 1.20 m sits at seated eye level: it blocks a seated forward view and is trivially seen over by anyone standing, delivering seated privacy with standing awareness. A screen at 1.35–1.45 m blocks the seated view completely and makes a standing person deliberately look over, **converting a glance into an act. The fifteen centimetres between those heights is the difference between a room where one is interruptible by accident and a room where interrupting someone is a decision.**

12.4.7 4.7 What a threshold costs, and whether it pays

A five-zone entrance sequence for a medium building — forecourt 3.0 x 4.0 m, porch 2.0 x 3.0 m, vestibule 2.4 x 2.4 m, lobby 4.0 x 5.0 m, anteroom 2.4 x 3.0 m — totals **51.0 m^2**: 8.5 per cent of a 600 m^2 floor, 5.7 per cent of 900 m^2, 4.2 per cent of 1200 m^2. **Four to nine per cent of net area, and the first thing value-engineered out of a project, because on a plan it reads as circulation and circulation reads as waste.**

Against that, the recovered capacity. For 30 people, with threshold work moving `k` from 22 to 6 per day at a measured `m` of 7 minutes:

```
before 22*7 = 154 min/day ; after 6*7 = 42 min/day ; saved 112 min/day = 1.87 h/person/day
30 people * 220 working days * 1.87 h = 12,320 person-hours/yr = about 7.7 FTE-years per year
```

Fifty-one square metres of transition against twelve thousand person-hours a year, and even at a third of the assumed lag it is not close. **Three qualifications are owed and are given.** First, `k` moving from 22 to 6 is not achieved by architecture alone: the protocol half must be maintained, which is a decay problem rather than a build. Second, **not all recovered time converts to output, because the conversion is `eta` and `eta` is not 1.** Third, **area is capital and time is operating**, so state both rather than netting them silently.

The comparison is also a decision under uncertainty and should be priced as one. The corpus's instrument for that is an expected headroom change: an uncertain call priced as `E[dh] = 0.0318401 p − 0.0051096` from a right-call gain of `+0.026731` against a wrong-call loss of `−0.005110`, giving a break-even probability of 16.05 per cent [Curzi, 2026]. **The same form applies here: state the headroom gain if the mechanism operates, the headroom loss if the floor area is simply lost, and the probability at which the build is worth commissioning.** No such break-even is computed in this paper, and Section 11 records that as an unmet check rather than implying that the area calculation above settles the question. Relatedly, a threshold is a barrier in the framework's own sense, and the barrier analysis supplies the condition under which no amount of accumulated surplus funds a departure: a critical cost `C_crit = eta M − lambda E_b` of 0.082 stored-order units per week, and a critical efficiency `eta_crit = C/M`

= `0.055238` against a measured `eta` of 0.05 [Curzi, 2026p]. **A threshold too deep is a barrier nobody crosses, and that analysis is where its upper bound comes from.**

12.4.8 4.8 Where a threshold must not be

Prohibition 1 — a threshold decelerates; it does not deny. Every device here must be crossable by anyone who chooses. The bent entrance slows a person; it does not stop them. The closed door has a handle on both sides. The screen can be looked over. **A threshold that removes a person's ability to enter is not a threshold, it is a gate, and gates are outside this framework's scope.** The distinction is operational: **a threshold raises the cost of crossing and leaves the choice with the crosser; a gate moves the choice to a third party.**

Prohibition 2 — no threshold on an egress route. Deceleration devices, bends, level changes and closable doors are obstructions under evacuation conditions; each must be checked against the applicable life-safety code, **and where a conflict exists the code wins without discussion.** This paper supplies no egress widths, travel distances or door hardware: those are jurisdictional and they change.

A third caution is softer and sharper. Filters reduce `k` *selectively*, and the selection is not neutral. **A door that closes filters hardest against those with the least standing to knock — the new, the junior, the visiting — so it silences exactly the members whose views were least correlated with the rest, raising conversion by cutting non-parallelism.** That is the parallel-wall failure arriving through the door frame. The counter is architectural rather than exhortative: **pair every closable private zone with an open, low-commitment, edge-type position that anyone may occupy without asking**, and check by occupancy count that the paired positions serve more than one cluster.

12.5 5 Substrate: light, thermal, air, material

Nothing about light, air or material is derivable from the storage law. The law supplies the accounting frame: once a substrate variable has been shown to move conversion, decay, cost or the category count, the law prices the move. **This section's work is establishing the *that*, and being explicit about the evidence.** The modelling postulate it rests on is that the physical substrate bounds the duration and quality of *voluntary* occupancy, and that voluntary occupancy carries interaction; substrate therefore enters through `eta` (occupied time is the substrate of durable relation), through `C` (discomfort is dissipation, and so is every adjustment negotiation) and through `delta` (a room that degrades needs replacing). **That is badged as a modelling postulate and is not a theorem.**

12.5.1 5.1 The strongest single piece of evidence, stated precisely

The cleanest demonstration that physical setting moves a clinical outcome is a 1984 study of postoperative recovery published in *Science*, comparing patients whose window faced a stand of deciduous trees against patients whose window faced a brick wall. **The design: a retrospective study of records from one suburban hospital over 1972–1981; patients recovering from one procedure were matched in pairs on age, sex, smoking status, weight, previous hospitalisation, year of surgery and attending surgeon; twenty-three matched pairs.** The rooms were essentially identical except for the view. Tree-view patients had shorter postoperative stays, received fewer doses of moderate and strong analgesics, and drew fewer negative evaluative comments in the nurses' notes. The stay difference was a fraction of a day against a mean of roughly eight.

Why it is worth more than the surveys. It measured *outcomes* — length of stay, analgesic dose, clinical notes — not preferences. Nobody was asked whether they liked the view, the

patients did not know they were in a study, and the outcome measures came out of the hospital's ordinary operation. **That removes the demand characteristics that make almost all built-environment preference research uninterpretable.**

Why to be careful, said as loudly as the finding. Twenty-three pairs is a small sample. One hospital, one procedure, one climate, one decade. It was retrospective, and room assignment was as-available rather than randomised — close to random in practice, not guaranteed. The effects are modest, and the nurses' notes were scored from records not written for the purpose. **The reproducibility literature is unambiguous about what a single small retrospective study is worth** [Open Science Collaboration, 2015, Ioannidis, 2005]. **This paper carries no bibliography entry for that study**, because none was assembled for it; it is described here rather than cited, and anyone quoting its numbers must read the paper and take the figures from it. Section 11 lists the gap.

What the study does license is narrow and useful: a variable the design professions classify as aesthetic turned up in a clinical outcome measure, so the prior that substrate is decoration is wrong. **It does not license a coefficient.**

12.5.2 5.2 Light, in two planes

Illuminance is measured in a plane, and which plane is chosen decides what is being measured. **Task illuminance is horizontal, at the desk; circadian and alerting effects are driven by light arriving at the eye, which is vertical.** A room can deliver 500 lux on the desk and 130 lux at the eye, and a lighting sign-off taken in one plane says nothing about the other. **Requirement B3 below therefore requires recording in both planes, and failure mode FB3 is the horizontal-only sign-off.**

Daylight has the same edge structure as Section 4.1. Useful daylit depth from a side window runs to roughly two to two and a half times the window head height, so a 2.4 m head reaches about 4.8–6.0 m. **A 22 m deep plate cannot reach edge density and cannot reach daylight, and it is the same geometric fact twice:** perimeter is linear and area is quadratic, so both amenities dilute together as the plate deepens. **No lighting scheme fixes a plate that is too deep, and the honest project-record entry says so rather than relaxing the requirement.**

Glare belongs to the cost term rather than to the light budget, because its signature is behaviour: **a permanently closed blind is a glare problem answered by the occupant with the only tool they were given, and it converts a daylight asset into an electrically lit room.** Count permanently closed blinds; the count is a direct, free measurement of a design failure.

The circadian claim is stated with its status. The proposition that the timing and spectrum of light at the eye affects sleep and next-day capacity — and therefore, through capacity, the conversion of interaction into durable structure — is a chain with a well-supported first link and a modelled second. **This paper takes the melanopic-equivalent target used in current practice as an engineering-practice value, states that it changes between editions of the standards that carry it, and does not derive it.**

12.5.3 5.3 Thermal comfort: six variables, and why adjustability beats the optimum

Thermal comfort has six variables — air temperature, mean radiant temperature, air speed, humidity, metabolic rate and clothing insulation — and almost every building controls one. **The other five are where the complaints come from.** Operative temperature, the single number that predicts comfort at low air speed, is approximately the mean of air and mean radiant temperature, **so a room can hold its air at 21.0 C and deliver an operative 18.9 C at a desk beside single glazing.** Mean radiant temperature is the invisible variable: it is not what the thermostat reads, it is not what the occupant can name, and it is what the occupant is feeling.

The arithmetic that settles the design question is in the standard comfort model's own output. The predicted percentage dissatisfied is derived from the predicted mean vote by

```
PPD = 100 − 95 ∗ exp( −(0.03353 PMV^4 + 0.2179 PMV^2) )
```

PMV	0.0	0.2	0.3	0.5	0.7	1.0	1.5	2.0
PPD	**5.0%**	5.8%	6.9%	10.2%	15.3%	26.1%	50.9%	76.8%

Read the first column again. At the theoretical optimum, every variable ideal, the model predicts 5 per cent dissatisfied. The floor is not zero. The usual criterion — mean vote within `+/−0.5`, giving 10 per cent dissatisfied — **accepts one person in ten uncomfortable by design, in a building working exactly as specified.** That floor is not a defect of the model: it is the model correctly reporting that thermal preference is dispersed and that one environmental state cannot sit at every individual optimum. **A fixed setpoint, however well optimised, answers the question "where is the centre of the distribution?", and almost nobody lives at the centre of a distribution.**

Model individual preferred operative temperatures as normal with standard deviation `sigma`, and suppose each person accepts `+/−1 K` around their own preference. **`sigma` is an assumed parameter, badged as such; it must be measured in the population at hand (Protocol B1).** The 1.5–2.5 K range is the plausible design region. Fraction satisfied by one fixed setpoint at the population mean:

`sigma` (K)	`+/−1.0 K` band	`+/−1.5 K`	`+/−2.0 K`
1.0	68.3%	86.6%	95.4%
1.5	49.5%	68.3%	81.8%
2.0	38.3%	54.7%	68.3%
2.5	31.1%	45.1%	57.6%

Give each occupant personal authority of `+/−A` K over their own operative temperature, widening the effective acceptance band to `+/−(1 + A)`:

Personal authority `A`	`sigma` = 1.5 K	`sigma` = 2.0 K
`+/−0 K` (fixed setpoint)	49.5%	38.3%
`+/−1 K`	81.8%	68.3%
`+/−2 K`	95.4%	86.6%
`+/−3 K`	99.2%	95.4%

At `sigma = 2.0 K`, two degrees of personal authority takes satisfaction from 38.3 per cent to 86.6 per cent. No amount of re-optimising the single setpoint can do that, because the problem is the dispersion and not the location. And it is robust to the assumed dispersion in the way that matters: **the ratio of adjustable to fixed performance improves as `sigma` grows.** A population more varied in age, body composition, medication, metabolic rate, gender or dress makes the case stronger, not weaker.

The four adjustment channels, ranked by cost.

Channel	Variable moved	Cost	Speed	Personal or shared?
Clothing	insulation	zero	seconds	**personal**
Local air movement (fan)	air speed	trivial	seconds	personal

Channel	Variable moved	Cost	Speed	Personal or shared?
Local radiant (small panel, warm surface)	radiant temperature	low	minutes	personal
Zone setpoint	air temperature	high	tens of minutes	**shared — and therefore contested**

Clothing is the most powerful and cheapest lever and the one organisations routinely disable. As a rule of thumb from the comfort model, 0.1 units of clothing insulation shifts the neutral operative temperature by roughly 0.6 K, so moving between a summer and a winter outfit is worth about 3 K. **Badge: verify that against the standard's own tables for the metabolic rate in question before designing to it; this paper carries no bibliography entry for the comfort model, its tables or the adaptive model discussed below** (Section 11). **A dress code is thermal policy.** Fixing clothing removes the free, personal, instantaneous channel and transfers the whole load onto a shared, slow, expensive one, **so the building burns energy to compensate for a decision taken in a handbook.**

Contested adjustment is a cost line and it is unusually easy to measure. Zone setpoint control is shared, so every adjustment is a negotiation. Log thermostat interactions — who, when, direction, magnitude, and whether a conversation followed. **A shared control changed more than about twice a day is not a control, it is a dispute.** Convert to the cost observable as contested adjustments per period times mean time consumed per adjustment, including the discussion and the counter-adjustment. **Rule B-A — never install a shared control where a personal one will do.** A personal control changes one person's environment and generates no coordination cost; a shared control generates one on every use, forever, **and it falls hardest on whoever is least willing to argue.**

The connection to the category count, with its badge on. A room that can only be set one way selects for the bodies it suits: people whose thermal preference lies far from the setpoint attend less, disengage sooner, relocate or leave. Since thermal preference is not independent of age, body composition, gender, health status or medication, **that selection is not random with respect to the categories the surplus law counts, so a fixed setpoint reduces the effective count by attrition and lowers `ln N` itself if a category is lost entirely. This is the most speculative claim in the paper and it is stated because it is testable.** Protocol B1 in Appendix A is the audit that would test it, and it names the result that would require the claim to be withdrawn. **The well-supported half of the story needs no postulate: field studies in naturally ventilated buildings find that occupants with real control accept a substantially wider range of conditions than the laboratory model predicts, and perceived control widens the acceptance band independently of what it is used to do.** That result is replicated and on its own justifies Rule B-A; **it is described here without a citation key, because none exists in this corpus's bibliography.**

The honest costs of adjustability. *Energy:* wider zone tolerance plus personal devices is usually a net saving, but personal resistance heaters are the ruinous exception — specify personal *radiant* and *air-movement* devices. *Control complexity:* every controllable device is a maintenance item, and decay applies to control systems more brutally than to fabric. *Spillover:* the channel must be personal in its *effect*, not merely in its interface — **a desk fan blowing across a neighbour's neck has converted a personal control into a shared one.**

12.5.4 5.4 Air quality, and a literature this paper declines to use

Steady-state indoor carbon dioxide under constant occupancy is `C_in = C_out + G/Q`, with `C_out` the outdoor concentration — about 420 ppm and rising — `G` the generation rate and `Q` the outdoor air supply. **Inverting it turns a cheap logger into a ventilation meter:** a room with

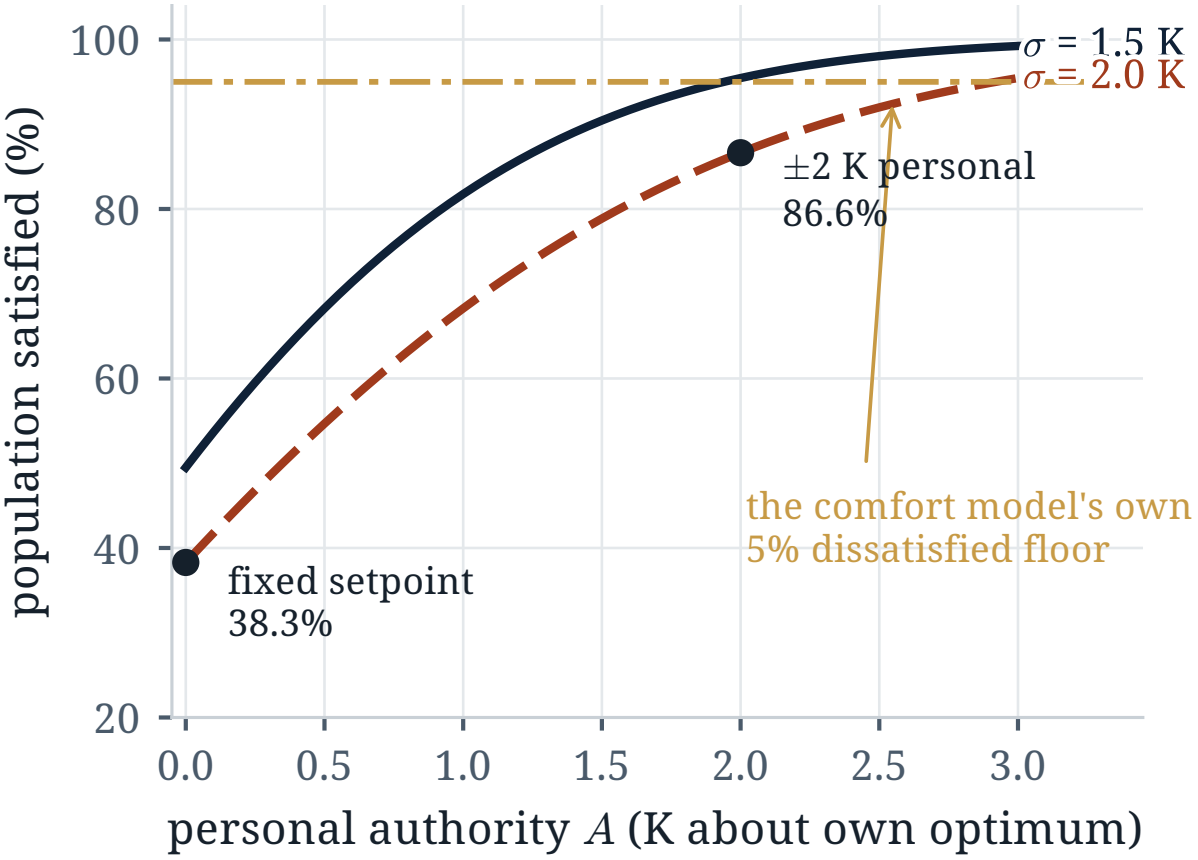


Figure 12.3: **The dispersion, not the location, is the problem.** Fraction of a population satisfied when each person accepts +/−1 K around their own preferred operative temperature and is given +/−A K of personal authority. At a 2.0 K dispersion the fixed setpoint satisfies 38.3 per cent and two degrees of personal authority satisfies 86.6 per cent. Modelled: preferences are assumed normal and the dispersion is an assumed parameter badged in the text; the comfort model's own 5 per cent dissatisfied floor is drawn as the ceiling this cannot cross. Generated by `figs/src/CCP-012_fig03.py`.

four occupants at a 1,290 ppm afternoon plateau is being supplied at `Q = 0.005 x 4 / (870 x 10^-6) = 23.0 L/s`, which is 5.7 L/s per person. **That inversion is the reason to own a logger, and it is a measurement rather than a claim.**

And here is the tension, stated in the direction unhelpful to this paper. A room ventilated exactly to the common standard settles at roughly 1,300 ppm — **above the concentrations at which some studies report decrements in decision-making performance.** A framework that wanted to sell ventilation would stop there. **The honest state of that literature is that it is contested:** effect sizes vary by an order of magnitude between studies, several of the strongest results come from small samples in chamber conditions with outcome measures of uncertain relation to real cognitive work, the exposure is confounded with other bioeffluents that track occupancy in exactly the same way, and replication attempts have been mixed. **The conclusion, deliberately unhelpful to the case this paper would prefer to make: the proposition that carbon dioxide at ordinary indoor concentrations impairs cognition is not established well enough to carry a business case.**

What is well supported is the ventilation case itself — odour, perceived air quality, symptom reporting, and airborne infection risk — **and that is what may be claimed. Rule B-B: specify ventilation, monitor carbon dioxide as a tracer, and claim only ventilation.**

This refusal is the same operation the corpus performs elsewhere when a ratio kills a claim. The decoherence computation is the cleanest instance: a characteristic time of `1.98 x 10^-17` s for a warm wet 310 K system against a 1 ms neural event is a shortfall of `5.04 x 10^13`, and a hypothetical `10^6` protection factor still leaves `5.04 x 10^7` [Curzi, 2026aa]. **A framework that computes the ratio and then declines the claim is doing the same work here at a much smaller scale, and the discipline is what makes the claims it does keep worth reading.** It is also the operation that demotes a diversity measure to a cost line when the interfaces it cannot see are counted [Curzi, 2026]: **the measured quantity is retained and the story attached to it is withdrawn.**

12.5.5 5.5 Materials do four jobs

Acoustic. With absorption coefficients at 500 Hz taken as standard-practice values — painted concrete 0.02–0.05, glazed partition 0.03–0.05, gypsum board on studs 0.05–0.10, carpet tile on concrete 0.15–0.30, heavy drapes 0.50–0.60, mineral-fibre ceiling tile 0.60–0.80, 50 mm mineral wool 0.80–0.95, and an occupant on an upholstered seat about 0.45 m^2 of equivalent absorption each — and Sabine's `RT60 = 0.161 V/A` [Sabine, 1922, Curzi, 2026e], an all-hard 42 m^2 workroom (6.0 x 7.0 x 2.9 m, `V` = 121.8 m^3, `S` = 159.4 m^2) computes as:

```
HARD CASE   exposed concrete soffit 42 x 0.02  =  0.84
            plastered walls        67.4 x 0.03 =  2.02
            glazing                 8.0 x 0.04 =  0.32
            vinyl on concrete      42.0 x 0.03 =  1.26
            A = 4.44 m2 sabins   ->  RT60 = 0.161 x 121.8 / 4.44 = 4.42 s
```

Four and a half seconds. That is a church, not a workroom, and it is what an all-hard room of this size genuinely does: speech is unintelligible at conversational level and the loudest voice wins by physics [Curzi, 2026e]. **A note on the arithmetic: the source manuscript this paper converts records 4.41 s for this case, and the quotient is 19.6098/4.44 = 4.4166, which rounds to 4.42. The value is corrected here rather than repeated, and nothing in the argument turns on the hundredth of a second.**

Thermal mass. Volumetric heat capacity and diurnal penetration depth `d = sqrt(alpha P / pi)`, with `alpha = k/(rho c)` and `P` = 86,400 s:

Material	`rho`	`c`	`k`	`rho c` (MJ/m^3 K)	24-hour depth
Dense concrete	2300	880	1.7	2.02	152 mm

Material	rho	c	k	rho c (MJ/m^3 K)	24-hour depth
Brick	1700	800	0.72	1.36	121 mm
Gypsum board	800	1000	0.25	0.80	93 mm
Solid timber	500	1600	0.13	0.80	67 mm
Water	1000	4180	0.6	4.18	63 mm

Penetration depth is where the daily swing has fallen to `1/e` of its surface amplitude, and standard design practice takes roughly the first 100 mm of exposed heavyweight material as the participating depth. **A 300 mm slab exposed on one face therefore performs like a 100–150 mm slab: thickening buys nothing for the diurnal cycle and exposing more *area* buys everything.** Stored energy in exposed concrete:

```
per m^2 of exposed surface, over the participating depth: 2300 x 880 x 0.152 = 308 kJ/m^2/K
100 m^2 of exposed soffit over a 3 K diurnal swing: 308 x 100 x 3 = 92.3 MJ = 25.6 kWh
```

Twenty-six kilowatt-hours of free storage, shifted from afternoon peak to night, out of a soffit already paid for. And then: carpet the slab, or hang a sealed suspended ceiling under it, and the mass is gone. The material is still there; its coupling to the room air is not. **A decision taken for acoustics silently deletes the thermal strategy, and nobody notices, because the two decisions are made by different people in different months.** That is failure mode FB5.

The resolution is computable and the two requirements are compatible:

```
TREATED   exposed soffit, 60% left open       25.0 x 0.02 =  0.50
          50 mm mineral wool baffles          17.0 x 0.90 = 15.30
          plastered walls                     57.4 x 0.03 =  1.72
          10 m2 wall panel, 50 mm             10.0 x 0.85 =  8.50
          glazing                              8.0 x 0.04 =  0.32
          carpet tile on concrete             42.0 x 0.25 = 10.50
          4 occupants                          4   x 0.45 =  1.80
          A = 38.64 m2 sabins  ->  RT60 = 0.51 s      exposed mass retained: 60%
```

Reverberation from 4.42 s to 0.51 s, inside the speech range, with 60 per cent of the soffit still thermally coupled. Baffles and rafts absorb without sealing. The two requirements are compatible; they are never coordinated.

Tactile. Two surfaces at identical temperature feel different, and the physics is exact. Thermal effusivity `e = sqrt(k rho c)` governs the interface temperature reached when skin touches a surface: `T_contact = (e_skin T_skin + e_mat T_mat)/(e_skin + e_mat)`. With skin at 33 C, skin effusivity taken as 1200, and every material at a true **20 C**:

Material	e	Contact T	Material	e	Contact T
Aluminium	22,045	**20.7 C**	Glass	1,369	26.1 C
Mild steel	13,539	21.1 C	Oak	436	29.5 C
Stainless steel	7,697	21.8 C	Pine	322	30.2 C
Granite	2,421	24.3 C	Wool carpet	125	31.8 C
Dense concrete	1,855	25.1 C	Cork	104	**32.0 C**

At the same 20 C, an aluminium handrail presents 20.7 C to the hand and a cork one presents 32.0 C — eleven degrees at the skin from a material choice alone. No thermostat setting produces that difference and none compensates for it. Specify by effusivity wherever skin meets building: handrails, door handles and push plates, seat surfaces and arms, work-tops, sills people lean on, floors in barefoot spaces. **The steel handrail on an unheated stair is a thermal decision being made without the calculation.**

Temporal: patina and the decay rate. Materials divide by what wear does to them. Oiled and waxed hardwood, brass, bronze and copper, vegetable-tanned leather, lime plaster, brick, stone, linoleum and cork improve with wear; melamine and laminate faces, thin veneer on board, printed vinyl, painted board edges and powder coat over thin steel degrade with it. **The distinction is structural, not aesthetic: a material improves with wear when its surface and its body are the same substance**, so wear exposes more of the same, and it degrades when a thin surface layer differs from the substrate, so wear breaches the layer and exposes something visibly other.

In the law's terms **decay is per component.** Structural decay — the desk stops being a desk — is small for both classes; appearance decay — the desk stops reading as cared-for — is large for layered materials and near zero for solid ones. **Since perceived maintenance drives the behaviour that actually maintains a place, a low-appearance-decay material is a self-reinforcing investment and a high-decay one is self-defeating.** Illustratively, over 40 years per m^2 of worktop:

```
laminate: 180 initial, replaced every 8 years -> 5 purchases = 900   (22.50 /m2/yr)
timber:   430 initial + 6/m2/yr oiling        -> 430 + 240   = 670   (16.75 /m2/yr)
```

Badge: those unit rates are illustrative placeholders showing the structure of the calculation, and the conclusion depends entirely on them. Substitute quoted rates. What does *not* depend on them is the shape: **a layered material has a step-function cost stream with disposal and disruption at each step; a solid one has a small continuous stream. Step functions are also what get deferred when budgets are tight, which is how a building arrives at year twelve looking abandoned.** The argument that aged buildings house the low-margin, unproven and various uses that new construction prices out is a different and economic one [Jacobs, 1961]; **the decay-rate framing here is this framework's interpretation and not that finding.**

12.5.6 5.6 Mapping the substrate to the law

Substrate variable	Parameter	Mechanism	How to measure
Daylight and view	`eta`	Occupied hours at the perimeter	Daylight-factor survey; vertical illuminance; occupancy by position
Circadian light timing	`delta`	Sleep, next-day capacity, skill retention	Melanopic illuminance at eye height, morning and evening
Glare	`C`	Blinds closed, screens moved, complaints	Count of permanently closed blinds; glare rating
Operative temperature at the seat	`C`, `N_eff`	Discomfort; contested adjustment; attrition	Globe thermometer at the seat; Protocol B1
Personal control provision	`C`	Removes coordination cost per adjustment	Controls per occupant; shared-control log
Ventilation rate	`C`	Symptoms, infection risk, odour	Carbon-dioxide logger inverted to supply rate

Substrate variable	Parameter	Mechanism	How to measure
Acoustic absorption	`u`, `eta`	Turn-taking equality [Curzi, 2026e]	RT60 by Sabine; speaking-share distribution
Exposed thermal mass	`C`	Peak-load shifting; free storage	Coupled surface area; internal temperature swing
Tactile effusivity	`eta`	Willingness to occupy and to touch	Material schedule; computed contact temperature
Patina behaviour	`delta`	Appearance decay rate	Annual fixed-position photographs

Every row of that table is a modelling assignment, not a measured coefficient. The right-hand column is what would turn each row into a measurement, and the worked case in Section 8.2 takes five of them.

12.6 6 The table with no head

Governance is usually written down. This section is about the part that is built instead. **A table with a head is a governing instrument.** It assigns one seat the sightlines, the acoustic advantage and the conventional right to open and close, and it assigns every other seat the complementary role. **The function of everyone else becomes agreement — and agreement between aligned instruments is echo, not evidence.** That is not a metaphor: when members' contributions become parallel, `u` goes to zero, the effective count goes to one, and the surplus of the whole assembly is `ln 1 = 0`.

The claim is that this is fixable in joinery, and it has three parts. **The damage from deference is quartic, not linear**, so mild deference is nearly free and strong deference is catastrophic. **Geometric equality has an exact definition computable from a plan**, the variance of visual centrality across seats. **And three independent constraints — proxemic, acoustic and coordination cost — converge on the same table size**, none of them knowing about the other two.

12.6.1 6.1 The arithmetic of deference

The deference postulate, badged. Let `h` be the chair's position vector and let each other member have a private view `w_i`, unit, mutually orthogonal, and orthogonal to `h`. Model the contribution actually made by member `i` under deference as

```
v_i = a h + sqrt(1 − a^2) w_i          a in [0,1]      a = the alignment coefficient
```

Then `v_i . v_j = a^2` for `i != j`, so `u_ij = 1 − a^4`, and the assembly's effective count is `N_eff = 1 + (n−1)(1 − a^4)`. **What is postulated is the *form* of deference — that it acts as a rotation of each contribution toward the chair's, by a common coefficient. What follows from that form is exact arithmetic.** The postulate is wrong in detail: deference is not uniform across members, private views are not exactly orthogonal, and the coefficient may vary by topic. **It is the simplest model with the right qualitative behaviour and it has no external support whatever.**

a	a^4	u	N_eff (n = 12)	S = ln N_eff	% of ln 12
0.00	0.0000	1.0000	12.000	2.4849	100.0
0.30	0.0081	0.9919	11.911	2.4775	99.7
0.50	0.0625	0.9375	11.312	2.4259	**97.6**
0.60	0.1296	0.8704	10.574	2.3584	94.9
0.70	0.2401	0.7599	9.359	2.2363	90.0
0.80	0.4096	0.5904	7.494	2.0142	81.1
0.90	0.6561	0.3439	4.783	1.5650	**63.0**
0.95	0.8145	0.1855	3.040	1.1120	44.8
0.98	0.9224	0.0776	1.854	0.6173	24.8
1.00	1.0000	0.0000	1.000	0.0000	**0.0**

Read the shape rather than the rows. At `a = 0.5` the assembly keeps 97.6 per cent of its surplus; at `a = 0.9` it keeps 63.0 per cent; at `a = 0.95` it keeps 44.8 per cent. The loss is negligible until it is severe, because the fourth power does nothing until `a` approaches one and then does everything. That cuts both ways. **Ordinary politeness is nearly free** — a room where people are mildly influenced by the chair loses a few per cent, and governance effort should not be spent on it. **A room where members are 90 per cent aligned has lost more than a third of its assembly**, and one at 95 per cent has lost more than half, **while still tabling twelve names, twelve salaries and twelve signatures on the minute. Nothing in the paperwork records the loss.**

The same computation at other sizes exposes a trap:

a	n = 6 (% of `ln 6`)	n = 12	n = 25 (% of `ln 25`)
0.50	97.0%	97.6%	98.1%
0.70	87.5%	90.0%	91.9%
0.90	55.8%	63.0%	69.1%
0.95	36.6%	44.8%	52.7%
1.00	0.0%	0.0%	0.0%

A larger assembly is slightly more robust in *percentage* terms, **so adding members appears to help and barely does, because the ceiling being protected is `ln n` and it grows logarithmically. Twenty-five deferential members at `a = 0.9` deliver 2.225 nats against twelve independent members' 2.485. Twice the room, twice the cost, less surplus.**

The storage-law consequence. Take a standing council: `n = 13`, `eta = 0.50`, `delta = 0.08` per month, `Q = 9.0`, and cost by the declared calibration `C = phi eta ln n` with `phi = 0.45`. Then `ln 13 = 2.5649`, `eta ln n = 1.2825`, `delta Q = 0.7200` and `C = 0.5771`.

Condition	a	u	N_eff	S	eta S	h	Q_inf
Head of table	0.90	0.3439	5.127	1.6345	0.8172	**-0.015**	3.00
No head	0.30	0.9919	12.903	2.5574	1.2787	**+0.114**	10.37

The deferential council sits at `h` near zero: maintenance only, no growth available, whatever it resolves to do. Removing the head takes it to `+0.114` and raises the sustainable stored order from 3.0 to 10.4 — **a factor of 3.5, with the same thirteen people, the same room and the same agenda.** The break-even alignment is worth computing because it is a hard number a chair can be shown:

```
eta S = C   ->   S_crit = C / eta = 0.5771 / 0.50 = 1.1542
N_eff,crit = e^1.1542 = 3.1716
u_crit = (3.1716 - 1) / 12 = 0.1810
a_crit = (1 - 0.1810)^(1/4) = 0.9513
```

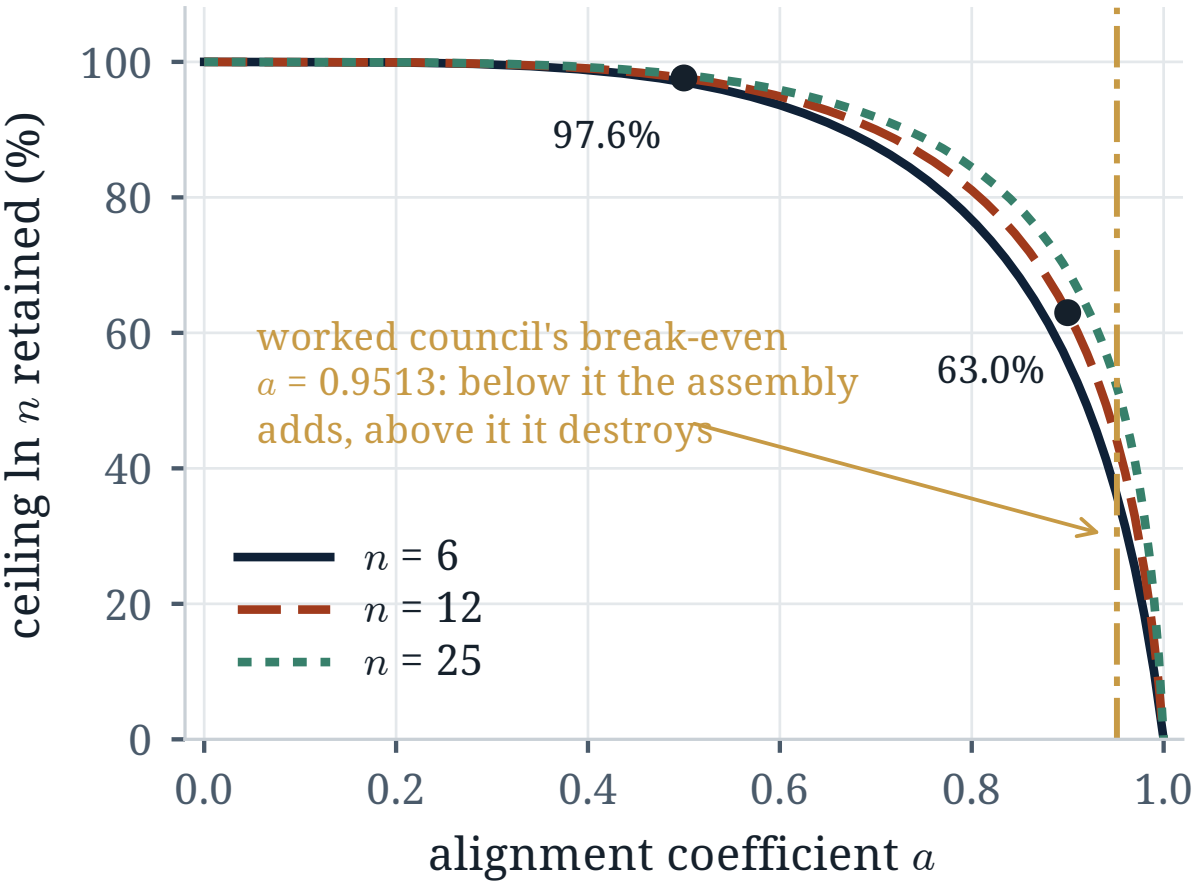


Figure 12.4: **Deference damage is quartic: nothing happens, and then everything happens.** Share of the ceiling `ln n` retained under the badged deference postulate `u = 1 - a^4`, for three assembly sizes. The worked council's break-even alignment — above which its meetings destroy stored order rather than adding to it — is 0.9513, which is its own parameters and not a general constant. Modelled: the form of deference is a postulate with no external support and the arithmetic that follows from it is exact. Generated by `figs/src/CCP-012_fig04.py`.

Above `a = 0.9513` this council destroys stored order every period it meets. It is not merely failing to add: the meetings are net-negative, because the coordination cost exceeds what a one-effective-member assembly can produce. **That is the sustainability ceiling arriving in a committee room** [Curzi, 2026b]. The limiting case is the framework's own unison result: at `u = 0` the steady state is `Q_inf = −c/delta`, in which **neither the conversion coefficient nor the category count appears at all, and two of the four levers have exactly zero derivative** [Curzi, 2026]. **A perfectly deferential assembly cannot be repaired by hiring, by diversifying or by conversion, because at `u = 0` those levers are disconnected from the outcome.**

12.6.2 6.2 Measuring alignment, and why a point estimate would be misleading

The alignment coefficient is not directly observable, and a table without a measurement would be selling arithmetic. Two instruments are given in full in Appendix A. Protocol D1, the written-then-spoken divergence, collects each member's written position before discussion and again after, and under the postulate's linear form the variance ratio relates to alignment as `V_post/V_pre ~ 1 − a^2`, so `a_max = sqrt(1 − V_post/V_pre)`. **Badge: that relation follows from the postulate's linear form, not from data.** Its limitation is that **legitimate persuasion by evidence also reduces post-discussion variance**, so the expression is an *upper bound*. Separate the two by asking, in the second collection, what changed the member's mind: a specific argument, or nothing they can name. Let `f` be the fraction of changed positions no member could attribute to an argument; then `a_min = sqrt(f (1 − V_post/V_pre))`.

Convergence nobody can attribute to an argument is the part that belongs to deference. Report the interval `[a_min, a_max]` and carry its midpoint as the working estimate. **A point estimate of alignment reported without its interval is the most misleading number this analysis can produce.** Protocol D2, the addressee count, is the cheap corroborant: log for each contribution whether the speaker's gaze and body orientation are directed at the chair or at the table. **A room in which over half of contributions are addressed to the chair is a room operating at high alignment**, and the count needs one observer and a tally sheet.

The mechanism the first protocol is designed to catch is an information cascade, which requires that each actor observe the actions of earlier actors before acting [Bikhchandani et al., 1992, Banerjee, 1992]. **Writing first makes every member's first commitment simultaneous, so there are no earlier actions to observe: the cascade is not discouraged, it is made structurally impossible for the first round.** The same failure has been demonstrated in aggregation experiments, where even light social information degrades the accuracy of a crowd's collective estimate [Lorenz et al., 2011] — the mechanism that the classic unaided-aggregation result depends on being absent [Galton, 1907, Surowiecki, 2004]. **Independence is a precondition of aggregation, and a chair who aligns the room has removed it.**

12.6.3 6.3 Sightline reciprocity and the exact definition of geometric equality

Seat every member on a circle of radius `r`, facing the centre. For two members separated by angle `phi` around the circle, the head turn `psi` that member `i` must make to face member `j` satisfies `cos psi = |sin(phi/2)|`, equivalently `psi = 90 deg − phi/2`:

Separation `phi`	27.7	45	60	90	120	166	180
Required head turn `psi` (deg)	76.2	67.5	60.0	45.0	30.0	6.9	**0.0**

Comfortable sustained head rotation is about `+/−45` degrees, with `+/−60` reachable and beyond that requiring the torso. **The result is counter-intuitive and it matters: at a round table the person directly opposite is the easiest to face — zero head turn — and the immediate neighbour is the hardest.** At thirteen seats the neighbour sits 27.7 degrees away and needs a 76-degree turn:

Seat offset	`+/−1`	`+/−2`	`+/−3`	`+/−4`	`+/−5`	`+/−6`
Separation (deg)	27.7	55.4	83.1	110.8	138.5	166.2
Head turn (deg)	76.2	62.3	48.5	34.6	20.8	6.9
Class	torso	torso	workable	easy	easy	easy

A round table is a plenary instrument, not a conversational one. It is excellent at making each member addressable by the assembly and poor at supporting side conversation — **which is exactly the trade a deliberating body wants and exactly the wrong trade for a dinner.**

Now the definition. For each seat `i`, let `V_i` be **the number of other seats within a 45-degree head turn of `i`'s forward gaze**, computable from a furniture layout with a protractor before anyone sits down. **Geometric equality is the condition `Var(V_i) = 0`.** Computed for real layouts, with seats 0.6 m out from the table edge and facing the table:

Layout	n	`V` per seat	mean	**variance**
Round, 12 seats, D = 2.67 m	12	7 everywhere	7.00	**0.00**
Regular 13-gon ring, D = 2.90 m	13	6 everywhere	6.00	**0.00**
Hollow square 2.4 x 2.4 m, 3 per side	12	7 everywhere	7.00	0.00
Rectangle 3.0 x 1.2 m, 5+5 sides + 2 heads	12	sides 5 to 6; heads **9, 9**	6.00	**2.00**
Boardroom 6.0 x 1.4 m, 8+8 + 2 heads	18	sides 5 to 7; heads **15, 15**	6.78	**9.06**

Read the last row. At an eighteen-place boardroom table the two head seats each see fifteen of the other seventeen within a comfortable turn; the seats at the far ends of the long sides see five. The head has more than twice the visual reach of the average member. **That is not a rule of order, it is a fact about the shape of a rectangle, and it operates in every meeting held at that table, forever, without anyone deciding it should.** The round table has variance exactly zero by rotational symmetry: every seat has the same view of the assembly, **so there is no seat to take and nothing to arrive early for.**

One honest qualification, and it is the kind that separates a computed claim from a slogan. The hollow square also scores zero variance at the 45-degree threshold — **but that is a coincidence of the threshold, not a symmetry.** Recompute at others:

Threshold	Round, n = 12	Hollow square, n = 12
`+/−30` deg	variance 0.00	variance **0.89**
`+/−45` deg	0.00	0.00
`+/−60` deg	0.00	**0.22**
`+/−70` deg	0.00	0.00

Only a rotationally symmetric arrangement has zero centrality variance at every threshold. The hollow square is a good second choice and it is not equal: its corner-adjacent and mid-side seats are in genuinely different positions. **If a hollow square is specified, say so, and record the variance at all four thresholds rather than the one that flatters it.** Protocol D3 in Appendix A is the audit, and **a layout whose variance has not been computed is a layout whose governance has not been specified.**

12.6.4 6.4 Sizing the table: three constraints that do not know about each other

Constraint 1 — proxemic. Seat pitch along the table edge: 0.60 m minimum, 0.70 m comfortable, 0.75 m generous. For a circular or regular polygonal table `D = n s / pi`, and the across-table distance between opposite members is `D`. Across-table distance must stay inside the social band — 1.20 to 3.60 m [Hall, 1966] — **or the far member is being addressed rather than conversed with.**

n	6	8	10	12	**13**	14	16	18	20	25
D at s = 0.70 m	1.34	1.78	2.23	2.67	**2.90**	3.12	3.57	4.01	4.46	5.57
Band	social	social	social	social	social	social	social	**public**	public	public

```
D <= 3.60 m (edge of social)              ->  n <= pi(3.60)/0.70 = 16.2
D <= 2.90 m (comfortable, facial detail)  ->  n <= pi(2.90)/0.70 = 13.0
```

Thirteen seats at 0.70 m pitch give a 2.90 m table — exactly the comfortable social distance. The thirteen-sided plan is not a symbol here; **it is what falls out of anthropometry and the proxemic bands, and it is stated as a derivation.** A 2.90 m table has a great deal of unreachable centre — comfortable seated reach is 0.45–0.60 m — so make it a ring: at 0.60 m working depth the inner void is `2.90 - 2(0.60) = 1.70 m` and the table footprint is 4.33 m^2, against a solid 2.90 m disc of 6.6 m^2 **of which 2.3 m^2 can never be reached.**

A measured artefact. The great oak table hanging in the Great Hall at Winchester is about 18 feet across and carries twenty-five painted places. `D = 5.486 m`, circumference `pi D = 17.24 m`, so `17.24/25 = 0.689 m` per place. **A fourteenth-century object is dimensioned at 0.689 m per seat, which is the modern comfortable place setting to within a centimetre: anthropometry has not changed and neither has the answer.** But run the second calculation: across-table distance is 5.49 m, well into the public band. **By its own geometry the Winchester table is a ceremonial object and not a deliberation table, and the round form does not rescue it. Roundness buys equality; it does not buy conversation. Conversation is bought by diameter.**

Constraint 2 — acoustic. Speech level falls with distance; in the direct field `L(d) = L(1 m) - 20 log10(d)`. Normal conversational effort is about 60 dBA at 1 m, relaxed about 55, raised about 66; good intelligibility wants about 15 dB of signal-to-noise ratio, with 10 dB the adequate floor.

Across-table distance	1.5 m	2.0 m	2.9 m	3.6 m	4.5 m	5.5 m
Level from a normal voice	56.5	54.0	**50.8**	48.9	46.9	45.2 dBA
Background needed for SNR 15	41.5	39.0	**35.8**	33.9	31.9	30.2 dBA
Background needed for SNR 10	46.5	44.0	40.8	38.9	36.9	35.2 dBA

At the thirteen-seat, 2.90 m table a normally spoken sentence arrives at 50.8 dBA, so the room's background must be at or below about 36 dBA for the quietest member to be heard without raising their voice. Go to twenty seats and 4.46 m and the requirement becomes 32 dBA, **a concert-hall-grade mechanical specification rarely achieved in an office building. Table size and the ventilation noise limit are the same constraint, and the person who chose a twenty-seat table has specified the air handling without knowing it. Honest qualification:** this

is a direct-field calculation. In a real room the reverberant field adds level as well, but late-arriving reverberant energy smears rather than helps, **so the direct calculation is the right conservative bound for clarity** and must be used together with a reverberation target [Curzi, 2026e], neither replacing the other.

Constraint 3 — coordination cost. Communication paths grow as `n(n−1)/2` while surplus grows as `ln n` [Brooks, 1975, Conway, 1968]:

n	6	8	10	12	13	16	20	25	30
Links n(n−1)/2	15	28	45	66	78	120	190	300	435
ln n	1.792	2.079	2.303	2.485	2.565	2.773	2.996	3.219	3.401
Links per nat	8.4	13.5	19.5	26.6	30.4	43.3	63.4	93.2	127.9

Maximising `eta ln n − c n(n−1)/2` gives `n(n − 0.5) = eta/c`, so `n* ~ sqrt(eta/c)`: 5.3 at `eta/c = 25`, 7.3 at 50, 10.3 at 100, 14.4 at 200 and 20.3 at 400. **The same constraint stated as a roster ceiling is `m(m−1) <= 2 eta ln(N)/kappa`, and the corpus's constructed commons sits exactly at it — 30 against a bound of 30.7225, so that a seventh member makes the case structural rather than personal** [Curzi, 2026].

The convergence.

```
Proxemic (facial detail, s = 0.70):   n <= 13
Proxemic (edge of social):            n <= 16
Acoustic (achievable 30 dBA room):    n <= about 13–16
Coordination cost, eta/c = 50..200:   n* = 7 to 14
```

Three constraints derived from anthropometry, from acoustics and from a combinatorial cost argument converge on eight to fourteen seats, and none of the three knows about the other two. That convergence is the strongest result in this section, and it is the reason to distrust any deliberating body sized above about sixteen. **Above sixteen the body is not running a deliberation; it is running a plenary that will be decided somewhere else, by fewer people, in a smaller room. The convergence is nonetheless three calculations agreeing, not three independent measurements agreeing**, and Section 10.6 records that they share the proxemic and speech-level inputs.

12.6.5 6.5 Protocols, and the measurement of turn-taking equality

Geometry sets the field and protocol sets the play. **Silent-write-first** opens every substantive item with two to five minutes in which everyone, the chair included, writes a position in a fixed answer format, followed by simultaneous unattributed reveal, with the pre-discussion variance recorded before any talking. **Round-robin** gives one pass in which every member speaks once without interruption, **with the order randomised each session — a fixed order hands the anchoring advantage permanently to whoever is first — and with an unconditional pass, because "pass" is a complete contribution and needs no reason.** The **last-speaks-first inversion** states positions from lowest rank to highest where rank is present and cannot be removed, so the senior member's position cannot anchor the rest; **it is long-established practice in aviation crew coordination, in military councils and in clinical teaching, and it is presented here as established practice rather than as a cited experimental result. This paper carries no bibliography entry for it** and the mechanism it defeats — anchoring on an authoritative first statement — is the documented part.

Measuring turn-taking equality. Let `p_i` be member `i`'s share of speaking time; compute `H = −sum_i p_i ln p_i`, `N_speak = exp(H)` and `E = H/ln n`.

Case	n	H	E	N_speak	Gini
8 equal	8	2.079	1.000	8.00	0.000
8, one at 35%	8	1.912	0.920	6.77	0.225
8, two at 30% each	8	1.806	0.868	6.08	0.350
8, one at 50%	8	1.666	0.801	5.29	0.375
8, one at 60% + two silent	8	1.317	0.633	**3.73**	0.575
12 equal	12	2.485	1.000	12.00	0.000
12, chair at 40%	12	2.112	0.850	8.26	0.317
12, chair at 40% + 3 silent	12	1.921	0.773	**6.83**	0.467

A twelve-person meeting in which the chair takes 40 per cent of the floor and three members say nothing has 6.83 effective speakers. Half the room is decorative and the minute will record twelve attendees. Report both `N_speak` and the Gini coefficient: **entropy is forgiving of a single dominator and harsh on silence, Gini is the reverse, and two indices disagreeing is information.**

BADGE, and it is the framework's own. `N_speak` is the exponential of a Shannon entropy — a Hill number of order one, a standard descriptive diversity index [Hill, 1973, Jost, 2006, Pielou, 1966] — borrowed here by formal analogy. It measures how equally the floor was shared. It does not compute `N_eff = 1 + (n−1)u`, which is not a Hill number of any order, and it does not compute surplus. The resemblance of form is a resemblance, not a theorem, and surplus is not identified with Shannon entropy anywhere in this framework.

What the chair still does. Removing the head does not remove the chair's work: timekeeping, agenda-setting, recording the decision and enforcing the protocols are all real and all necessary. **None of them requires the geometric privileges**, and a chair with the same sightlines as everyone else is a chair whose authority rests on the role rather than on the furniture. **One caution on a popular device:** an *assigned* devil's advocate is weaker than genuine dissent — authentic disagreement stimulates more and better divergent thinking than role-played disagreement — so assigning the role is better than nothing and **is not a substitute for recruiting someone who actually disagrees**, which is a recruitment problem and cannot be solved with furniture. **That finding is described here without a citation key, because this corpus's bibliography holds none.**

12.6.6 6.6 Why furniture rather than policy, and the limit of the claim

A rule requires an enforcer. The enforcer needs authority, discretion, attention and willingness, and every act of enforcement is a transaction that costs both parties something. **In the law's terms a rule generates a stream of cost for as long as it exists, and the stream falls hardest on whoever is most willing to speak up — who, as Section 4.8 noted, is not a random sample.** A shape enforces on every occasion of use: it has no discretion, grants no exceptions, forgets nothing, retires from nothing, and asks nobody to be the person who raised it.

Instrument	Setup cost	Marginal cost per use	Decays without maintenance?	Who bears enforcement?
A rule in the terms of reference	low	one enforcement act each time	yes, quickly	a person, repeatedly
A trained facilitator	high	a salary, every session	yes, on turnover	the facilitator
A round table	one-off capital	**zero**	only physically, over decades	**nobody**

A round table is invisible governance: nobody in the room experiences it as a constraint, nobody resents it, nobody has to be reminded of it, and nobody has to be the person who enforces it — which is exactly why it works and exactly why it never appears in the governance document. That is the general property of infrastructure, which becomes visible only on breakdown [Star and Ruhleder, 1996, Star and Griesemer, 1989].

The honest limit, stated as plainly as the claim. Geometry constrains; it does not determine. A sufficiently determined chair will dominate a round table by talking more, and a disciplined group will deliberate well at a rectangle. **What the geometry changes is the cost of dominance and the cost of resistance.** At a table with a head, dominating is free and resisting is expensive. At a table without one, dominating requires visible effort — more floor time, in front of people who can all see it happening — and resisting requires only speaking, from a seat with the same view as everyone else's. **That is a shift in a gradient, not a prohibition**, and it is what makes the method non-coercive: **a method that works by removing resistance has no coercive failure mode, and its worst outcome is that someone declines to use it.**

12.7 7 Design specifications

The requirements below are stated in the form a specification would carry. **`T1–T15` govern thresholds and edge, `B1–B14` the substrate, and `D1–D17` the deliberation table.** A requirement is not a finding: the schedule is a design instrument derived from the results above, and **its numerical thresholds are chosen values rather than measured optima.** Where a figure is jurisdictional or population-specific it says so, and says what to measure instead.

12.7.1 7.1 Thresholds and edge

T1 — mismatch survey. Rate both adjoining zones on the eight-attribute character vector of Section 4.3 and record `M` and `M_bar`. **T2 — step limit.** No threshold shall change any attribute by more than one point; at `M = 4`, at least three intermediate zones are required. **T3 — deceleration depth and dwell.** Any threshold at which occupants are expected to stop shall be not less than 2.0 m deep in the direction of travel, and where four or more seconds of dwell are required either the depth or a direction change shall supply it. **T4 — landing depth and pause-and-pass width.** On the swing side of an inward-opening 0.9 m leaf, not less than 1.50 m of landing depth; not less than 1.35 m clear width where one paused person and one passing person must coexist, and 2.85 m where a conversing pair must not become an obstruction. **T5 — legibility without signage.** Every zone transition shall be marked by at least two of surface, level, light, width, direction and ceiling height. **T6 — edge density.** Any occupiable room over 60 m^2 shall achieve a total perimeter, including interior islands, of not less than 0.40 m per m^2. **T7 — island clearance and occupiable edge.** Not less than 0.90 m clear between interior islands, and island edge shall be occupiable rather than merely present. **T8 — window place and alcove sizing.** Every room with a window shall have at least one seated position within 1.5 m of the glazing. **T9 — screen heights.** 1.20 m for seated privacy with standing awareness; 1.35–1.45 m to convert a glance into a decision. **T10 — closable boundary.** Every zone rated 1 or 2 on attention mode shall have a physically closable boundary. **T11 — paired invitation.** Every closable private zone under T10 shall be paired with an open, unbooked, edge-type position that anyone may occupy without asking, **and occupancy counts shall confirm that the paired positions serve more than one cluster. T12 — desire-line survey.** Run Protocol T1 for six weeks before any layout change and six weeks after. **T13 — switching measurement.** Measure `k` and `m` by Protocols T3 and T4 before and after, **and report `m`'s median and interquartile range, because the distribution has a long right tail. T14 — egress precedence and crossability.** Every threshold device shall be checked against the applicable life-safety code, which wins without discussion, and shall be crossable by anyone who chooses. **T15 — entrance area budget.** Allow 4 to 9 per cent of served net floor area for the entrance sequence, **and record it as a budget line rather than as circulation.**

12.7.2 7.2 Substrate

B1 — daylit depth. No regularly occupied workstation further than 2.5 times the window head height from its daylight source; where the plan exceeds this, **record the shortfall rather than relaxing the requirement.** **B2 — illuminance and colour**, to the applicable standard. **B3 — vertical measurement.** Illuminance shall be recorded in **both planes**, horizontal at the task and vertical at seated eye height in the working direction. **B4 — circadian target**, stated as a melanopic figure at the eye during daytime hours, **with its standard and edition named.** **B5 — glare**, to the applicable rating, **with permanently closed blinds counted as a failure.** **B6 — the six variables, measured at the seat.** Any thermal comfort claim shall state all six, measured at the occupied position and not at the thermostat. **B7 — glazing and the window seat.** Where T8 requires a seated position at the glazing, the local operative deficit at that seat shall be under 1.0 K at the design outdoor condition. **B8 — personal control.** Every regularly occupied position shall have at least **two** adjustment channels that are personal in *effect*. **B9 — shared controls and humidity.** Shared setpoint controls shall be logged with time, direction and magnitude; relative humidity shall be held at or above 30 per cent. **B10 — ventilation and carbon dioxide.** Specify the outdoor air rate to the applicable standard, monitor carbon dioxide as a tracer, **and claim ventilation only.** **B11 — reverberation with mass retained.** Meet the reverberation target using baffles, rafts and wall panels rather than sealed ceilings, **and state the percentage of thermally coupled soffit retained.** **B12 — participating mass.** Where thermal mass is claimed, compute the participating depth and the coupled area and state both. **B13 — contact surfaces.** Schedule every surface routinely touched by skin and state its computed contact temperature at the design condition. **B14 — patina classification and appearance monitoring.** Classify every visible material as improving or degrading with wear, and photograph from fixed taped positions annually.

12.7.3 7.3 The deliberation table

D1 — no head. The table shall be circular or a regular polygon with `n` equal places. **D2 — centrality variance.** Computed by Protocol D3, `Var(V_i)` shall be **zero at the 30-, 45- and 60-degree thresholds**, and the value at each shall be recorded. **D3 — seat pitch.** Not less than 0.60 m of table edge per member; 0.70 m is the design value. **D4 — diameter and seat count.** `D = n s / pi`, and `D` shall not exceed 3.60 m. **D5 — ring form.** Where `D` exceeds 1.80 m the table shall be a ring, with a working depth of 0.60 m. **D6 — upper bound on the body.** No deliberating body shall exceed **sixteen** members. **D7 — acoustic budget.** Background noise shall be not more than `L_speech(D) − 15 dB`, **computed and recorded with its margin.** **D8 — reverberation.** Meet the applicable reverberation target for the room volume; the direct-field calculation of D7 does not replace it. **D9 — chair seating.** The person chairing shall occupy an ordinary seat allocated by the same rule as everyone else's. **D10 — seat allocation.** Seats shall be unreserved, or allocated by a rule no member controls. **D11 — silent-write-first.** Every substantive decision item shall open with two to five minutes of silent writing, with `V_pre` recorded before discussion. **D12 — cascade test.** Where `V_post/V_pre < 0.25` and members cannot name the argument that moved them, **record a suspected cascade in the minute.** **D13 — speaking order.** The first round shall be round-robin in an order randomised or rotated by a rule nobody controls. **D14 — turn-taking measurement.** Speaking-time shares shall be measured at not less than one session in four, **with consent**, and `N_speak/n` recorded. **D15 — alignment estimate.** Estimate the alignment coefficient by Protocol D1 not less than annually **and report the interval, never a point.** **D16 — minority record.** Any member may attach a dissenting position to the minute, of their own drafting, without permission and without a second. **D17 — declination.** No provision of this specification shall be applied so as to compel a member to speak, to be recorded, or to remain.

12.8 8 Results: three systems worked end to end

All three systems are constructed. No dataset was collected for this paper. Every number below is the output of the stated parameters under the models of Sections 4 to 6, carried through so that a reader can re-derive it. **They demonstrate that the procedure runs and that it reports its own failures; they are not evidence that any organisation or building behaves this way.** In all three, the cost observable is the declared calibration `C = phi eta ln N` with `phi` a dimensionless fraction of the ceiling — **a declaration, not a derivation.**

12.8.1 8.1 A 30-person research floor: thresholds against a cost-bound gauge

The system. 900 m^2 net, a single open plate 22 x 41 m with partial glazed perimeter, four small enclosed rooms at the core, and one entrance door from the lift lobby straight onto the working floor. Nine specialisms, four years old. Surveyed: stored order `Q_0 = 6.0` as a normalised count of live shared methods, tools, protocols and documented working agreements in use; non-parallelism `u = 0.75` from a disagreement audit across the nine specialisms; `N = 9` verified non-redundant; `eta = 0.45` from durable structure per unit surplus over eight quarters; `delta = 0.06` per week from decay under zero input across two quiet periods.

```
ln N = ln 9 = 2.1972    N_eff = 1 + 8(0.75) = 7.000    S = ln 7 = 1.9459
eta*lnN = 0.9888  <- ceiling    eta*S = 0.8757  <- input    delta*Q = 0.3600  <- maintenance
```

The floor runs at 88.6 per cent of its geometric ceiling. There is almost no headroom left to buy by improving diversity, and the problem is elsewhere.

Measuring cost. Tally cards over two weeks across 30 people give a median `k = 22`; the progress-rate method over twelve person-days gives a median `m = 7` minutes with an interquartile range of 3 to 16.

```
C_switch  = 22 * 7 = 154 min/day of 480 = 0.321
phi_other = 0.25        (meeting overhead, tooling friction, compliance, room booking)
phi_before = 0.571   ->  C_before = 0.571 * 0.9888 = 0.5646
h = eta ln N - delta Q - C = 0.9888 - 0.3600 - 0.5646 = +0.0642
```

Headroom of `+0.064` against a ceiling of 0.989 — six per cent of the maximum one-step increment. Any project asking for more growth per week than that cannot be delivered, whatever is promised. Within measurement error the floor is at maintenance only. The steady state is `Q_inf = (0.8757 - 0.5646)/0.06 = 5.185`, and the current stock of 6.0 is **above** it: **the floor is not stagnant, it is declining slowly toward 5.19, everyone can feel it, and nobody can name a cause — because the cause is distributed across twenty-two interruptions a day.**

Diagnosis. *Mismatch:* lift lobby to working floor rates `M = 4`, `M_bar = 2.4` — a cliff edge. *Edge density:* enclosing perimeter `2(22+41) = 126 m` on 902 m^2 is **0.140 m/m^2** against a required 0.40, and core rooms add about 44 m of which half is occupiable, for an effective **0.11 m/m^2**; perimeter desks are occupied 94 per cent of working hours against 38 per cent for interior desks. *Desire lines, six weeks:* a diagonal from the entrance across the floor to the kitchen cutting through two clusters; chairs migrating toward the glazing week on week; three of four core-room doors habitually propped; handle wear on the fourth an order of magnitude above the rest, which is where the private conversation happens. Four closable boundaries shared by 30, no screens, **and headphones in use by 21 of 30 — twenty-one people privately funding the threshold layer the building did not supply.**

Intervention. Nothing changes the roster, recruitment, meeting culture or management structure: it is all joinery. (1) An entrance sequence: 30 m^2 beside the lift lobby becomes a vestibule (2.4 x 2.4 m, surface change, lower ceiling, warmer and lower light) and an anteroom (2.4 x 3.0 m, two seats, sightlines both ways), with entry turned 90 degrees; `M` at every junction falls from 4 to 1 or 2. (2) The kitchen is moved beside the anteroom so the diagonal becomes a

short public-zone route crossing no cluster — **this one move removes the largest single source of `k`.** (3) Fourteen thin islands, 0.9 x 2.4 m storage at 1.4 m high, form eight semi-private clusters of three or four desks at 0.9 m and 1.2 m clearances: edge added `14 x 2 x (0.9+2.4) = 92.4 m`, new density `(126+44+92.4)/902 = 0.291 m/m^2`, floor consumed 30.2 m^2. (4) Three corridor widenings of 1.2 x 2.4 m with a leaning rail and a writable surface, 8.6 m^2. (5) Two core rooms become bookable focus rooms with working locks; the propped pair get alarm-released hold-open magnets. (6) **Six window places at the glazing: unassigned, first-come, no booking, no ownership, open to visitors — the counterweight required by T11 and not optional.** (7) Two protected blocks daily, which is protocol rather than architecture and is therefore **subject to decay and reviewed quarterly forever. Total area consumed 68.8 m^2, 7.6 per cent of net floor.**

Re-measured at eight weeks, same protocols, same observers:

```
k = 6 per person per day        (was 22)
m = 7 minutes                   (unchanged, as expected: m is a task property, not a spatial one)

C_switch  = 42/480 = 0.087
phi_after = 0.087 + 0.25 = 0.337   ->  C_after = 0.3332
h_after   = 0.9888 - 0.3600 - 0.3332 = +0.2955
Q_inf     = (0.8757 - 0.3332)/0.06  = 9.041      (was 5.185)
```

Conversion, decay, the category count and non-parallelism are taken as unchanged, deliberately: the intervention touched exactly one term of the law. Headroom rises 4.6-fold and the sustainable level of stored order by 74 per cent, **with no change to the people, the work, or the diversity of the group.**

The transient. With `Q_t = (1-delta)^t Q_0 + (1 - (1-delta)^t) Q_inf`:

`t` (weeks)	0	5	10	20	30	50
`Q_t`	6.000	6.809	7.403	8.159	8.566	8.903

The time constant is `1/delta = 16.7` weeks and 95 per cent settling is at about 50 weeks; the half-life, computed correctly as `ln 2 / -ln(1-0.06) = 11.2` weeks rather than as the naive `1/delta` [Curzi, 2026], is the number to quote to anyone asking when they will see something. **Do not evaluate this intervention at eight weeks on the order observable:** at eight weeks the floor has moved to about 7.2, 31 per cent of the total, **and reading that as a disappointment misreads a first-order step response. Evaluate cost at eight weeks because cost is the input; evaluate stored order at a year.**

What would falsify it. (1) `k` above 15 at eight weeks means the threshold layer is not filtering and the model is wrong; log interruption *origin* as well as count, to separate the kitchen re-route from the islands. (2) `k` falls but stored order does not: a twenty-week reading below 7.4 means recovered time is not becoming durable structure, conversion is lower than measured, and cost was never binding. (3) **Non-parallelism falls:** a six-month disagreement audit below 0.65 means the clusters have become silos and the thresholds have partitioned the group's epistemics along with its noise — **this section's own tools installing a failure mode.** Rotate cluster membership quarterly and check that the window places serve more than one cluster; **if one cluster has colonised them, T11 has failed and the filter is doing damage.** (4) **The matched control:** a comparable floor in the same organisation over the same period, given a cosmetic refurbishment of equal cost and disruption but no change to mismatch, `k` or edge density. **If its stored order rises comparably, the claimed mechanism is not the operating one.**

And one requirement was not met. Edge density reached 0.291 m/m^2 against a required 0.40, **because the plate is 22 m deep and no arrangement of islands fixes a 22 m plate.** The project record reads: *T6 not met; deficit 0.109 m/m^2; cause, plate depth; remedy requires subdivision or a different building.* **A specification that adjusts itself to what was achieved has stopped being a specification, and that entry is worth more to the next engineer than any of the successes above.**

12.8.2 8.2 A four-person workroom: the substrate as a cost generator

The system. 42 m^2 (6.0 x 7.0 m), 2.9 m to soffit, `V` = 121.8 m^3, internal surface 159.4 m^2. One external wall, 6.0 m wide, north-facing, with 8 m^2 of air-filled double glazing (U = 2.8 W/m^2K) at a 2.4 m head. Exposed concrete soffit, vinyl on a concrete slab, plastered blockwork, one wall thermostat serving this room and two others. Six distinct specialisms across the occupants and their regular visitors. Surveyed: `Q_0 = 4.0`; `N = 6`, `ln N = 1.7918`; `u = 0.80`, so `N_eff = 5.00` and `S = ln 5 = 1.6094`; `eta = 0.28`; `delta = 0.070` per month; `phi` tallied at 0.62.

```
eta*lnN = 0.5017  <- ceiling      eta*S = 0.4506  <- input      delta*Q = 0.2800
C = 0.62 x 0.5017 = 0.3110
h = 0.5017 - 0.2800 - 0.3110 = -0.0894
Q_inf = (0.4506 - 0.3110)/0.070 = 1.994
```

Headroom is negative: in this model the room's order decreases regardless of how its occupants are arranged. Against a current stock of 4.0, the steady state of 1.994 means **the room is on a trajectory to lose half its stored order.**

The substrate survey. *Light:* desk illuminance 480 lux, compliant; **vertical at seated eye height in the working direction 130 lux, melanopic 95 at midday — against a 250 target the room delivers 38 per cent.** Average daylight factor 1.4 per cent; useful daylit depth at a 2.4 m head is 4.8–6.0 m against a room depth of 7.0 m. North aspect, so no direct sun, no glare, no blind use — **the one thing this room got right, by accident of orientation.** *Thermal:* air temperature at the thermostat 21.0 C; globe measurement at the window desk on a 0 C day gives an operative **18.9 C, a 2.1 K deficit**, consistent with a computed 13.4 C inner glazing surface. Two of four desks are within 1.5 m of the glazing; **personal channels available: zero.** The shared thermostat logged **3.8 contested adjustments per day** across its three rooms at a mean of 11 minutes of consequent discussion and counter-adjustment. *Air:* carbon dioxide logged three weeks, afternoon plateau with four occupants **1,290 ppm**; inverting, `Q = 0.005 x 4/(870 x 10^-6) = 23.0 L/s` total, 5.7 L/s per person. **Ventilated almost exactly to standard and sitting at 1,290 ppm — Section 5.4's tension in a real measurement.** Relative humidity ranged 22–38 per cent, below the 30 per cent floor for much of the heating season. *Acoustics:* measured reverberation at 500 Hz **2.6 s**, between the computed all-hard 4.42 s and a treated room, because furniture and occupants absorb. **Four people cannot hold two conversations here; they hold one, and the quietest of them contributes least.** *Material:* laminate desktops in year 6 of an 8-year life with edges lifting on two; powder-coated steel shelving; vinyl floor. **Every visible surface is layered: improving-with-wear fraction 0 per cent.** Contact temperature of the steel sill at an 18 C surface: 20.1 C.

Diagnosis. The room fails five of the six thermal variables, the vertical light plane, humidity, reverberation and material class; it passes horizontal illuminance, glare and ventilation rate. **The binding constraint is cost, and the substrate generates it:** 2.1 K of radiant deficit at half the desks driving a contested shared control; 2.6 s of reverberation suppressing three of four voices; a step-function replacement liability two years out. Conversion is also low at 0.28, and Section 5.6's mapping says why — **reverberation suppresses the turn-taking that produces durable agreement, and the window desks with the best light and view are avoided in winter.**

Intervention. (1) *Acoustic, mass-preserving:* 17 m^2 of 50 mm mineral wool baffles hung clear of the soffit, a 10 m^2 wall panel, carpet tile replacing vinyl — computed `A = 38.64 m^2` sabins, **reverberation 0.51 s, with 60 per cent of the soffit still thermally coupled.** (2) *Glazing:* double low-emissivity argon (U = 1.4) replaces the air-filled unit; inner surface at 0 C outdoors rises from 13.4 C to **17.2 C**, and the local operative deficit at the window desks falls from 2.1 K to under 1.0 K. (3) *Personal channels:* four desk fans and two low-wattage radiant foot panels, two channels per position, personal in effect. (4) *Shared control:* the room gets its own zone valve, so the contested control serves four people instead of twelve, with logging retained. (5) *Vertical light:* two indirect uplights and a wall-wash at 4000 K, dropping output and colour temperature after 17:00. (6) *Humidity:* trim humidification to hold 30 per cent and above.

(7) *Material:* as the laminate tops expire, replace with oiled solid timber; timber sleeves on the steel sill raise contact temperature at 18 C from 20.1 C to about 28 C; ten taped camera positions, photographed annually. **None of this changes the roster, the specialisms, or the work.**

Re-measured at four months:

```
eta:   0.28 -> 0.44   turn-taking recovered; window desks occupied year-round
delta: 0.070 -> 0.055 solid timber and carpet tile decay slower than lifting laminate and
                      scuffed vinyl; measured from the photo series plus replacement records
phi:   0.62 -> 0.38   contested adjustments 3.8/day -> 0.4/day; deferred liability removed
N, u:  unchanged at 6 and 0.80, so S = 1.6094 unchanged

eta*lnN = 0.7884     eta*S = 0.7082     delta*Q = 0.2200     C = 0.38 x 0.7884 = 0.2996
h = 0.7884 - 0.2200 - 0.2996 = +0.2688
Q_inf = (0.7082 - 0.2996) / 0.055 = 7.429        (was 1.994)
```

Headroom moves from `-0.0894` to `+0.2688`, from "no arrangement of people prevents decline" to a positive maximum one-step increment, and a 3.7-fold increase in the sustainable level of stored order — from baffles, glazing, four fans, two panels, a zone valve, two uplights and a change of worktop material.

The transient. Time constant `1/delta = 18.2` months; half-life `ln 2 / −ln(1−0.055) = 12.3` months [Curzi, 2026]; 95 per cent settling at about 55 months. **Do not evaluate this room on stored order at four months — it has moved perhaps a fifth of the way.** Evaluate the *inputs* at four months — cost, reverberation, melanopic illuminance, operative temperature at the seat, contested-adjustment count — **all of which respond immediately because they are the drive and not the state.**

What would falsify it. (1) `phi` does not fall: contested adjustments above 2 per day after zoning and personal channels means the dispute was never thermal. (2) Conversion does not rise: still near 0.28 at eight months with reverberation at 0.51 s means reverberation was not what suppressed it. (3) **The attrition claim fails:** run Protocol B1, and **if distance from delivered operative temperature predicts neither hours present, nor optional attendance, nor retention, the category-count claim of Section 5.3 is false in this building and must be withdrawn from the case. Report that result.** (4) **The matched control:** a comparable room, same institute and period, redecorated at equal cost and disruption but with no change to reverberation, glazing, personal control count or vertical illuminance.

The honest ledger. Shown: a substrate intervention changing no personnel and no category structure moved conversion up 57 per cent, decay down 21 per cent and cost down 4 per cent, taking the room from negative headroom to positive. **Not shown: that these are anyone else's numbers.** Conversion, decay, non-parallelism and `phi` were measured for one room; the calibration was declared, not derived; the dispersion of thermal preference in Section 5.3 was assumed and flagged; and **the conversion gain is attributed to acoustics and light on the basis of Section 5.6's mapping, which is a modelling assumption that falsification item 4 exists to break.**

Two things stayed broken and are recorded rather than smoothed over. B1 not met: the room is 7.0 m deep against a useful daylit depth of 4.8–6.0 m at a 2.4 m head; the back strip is electrically lit all day and **no lighting scheme changes that** — the remedy requires a second aspect or a shallower room. **And B10's tension is unresolved:** the room sits at 1,290 ppm while meeting the ventilation standard exactly. **Ventilation was not increased, because Section 5.4 says the cognitive case is not sound and the symptom and infection case was not made for this room. The project record reads: ventilation compliant, carbon dioxide 1,290 ppm, no action, reason stated. That is what an honest specification looks like when the evidence does not support the intervention one would have preferred to sell.**

12.8.3 8.3 A thirteen-member council: geometry against deference

The system. A research institute's standing scientific council: thirteen members from nine specialisms plus four cross-cutting roles, meeting monthly around a 4.2 x 1.4 m rectangular table with a chair position at one end, in a 7.0 x 6.0 x 3.0 m room with a hard plastered finish and a ceiling-mounted fan-coil unit. Surveyed: `Q_0 = 9.0` as a normalised count of live council-originated policies, standards and allocation rules still in force and actually used; `N = 13` distinct positions verified non-redundant, `ln 13 = 2.5649`; `eta = 0.50` from attributable durable structure over three years; `delta = 0.08` per month; `phi` tallied at 0.45.

Survey. *Centrality:* thirteen seats, six per long side and one at the chair's end, all 0.6 m out from the edge, give `V = [4,4,5,5,6,6,6,5,5,5,5,10]`, mean 5.538, **variance 2.095 at 45 degrees. The chair sees ten of the other twelve within a comfortable head turn; the two seats at the far end of the long sides see four.** *Acoustics:* across-table distance 2.6 m, working long-diagonal path 4.4 m, measured reverberation at 500 Hz **1.4 s**, measured background with the fan-coil running **42 dBA**. A normal voice delivers `60 − 20 log10(4.4) = 47.1 dBA` at that range, so 15 dB of signal-to-noise requires a background of **32.1 dBA: the room is 10 dB over budget on its long path.** *Addressee count:* over three sessions, **61 per cent** of contributions were directed at the chair rather than at the table. *Turn-taking, with consent over three sessions:* chair 34 per cent of speaking time, two senior members 13 per cent each, three members at or near zero, giving `N_speak = 6.9` against `n = 13`, so **`N_speak/n = 0.53`, below the D14 floor of 0.7.** *Alignment:* two scaled items per session over four sessions give a mean `V_post/V_pre = 0.14`, and in structured follow-up members could name a specific argument in 3 of the 8 items, so the unattributed fraction is `f = 5/8 = 0.625`:

```
a_max = sqrt(1 − 0.14)             = 0.927
a_min = sqrt(0.625 x (1 − 0.14))   = 0.733     working estimate a = 0.83, interval [0.73, 0.93]
```

Parameters before.

```
n = 13,  a = 0.83   ->   u = 1 − 0.83^4 = 0.5254
N_eff = 1 + 12(0.5254) = 7.305        S = ln 7.305 = 1.9886
eta ln n = 0.50 x 2.5649 = 1.2825     eta S = 0.9943      delta Q = 0.08 x 9.0 = 0.7200
C = 0.45 x 1.2825 = 0.5771
h     = 1.2825 − 0.7200 − 0.5771 = −0.0146
Q_inf = (0.9943 − 0.5771) / 0.08 = 5.21
```

At the interval's ends the steady state runs from **6.88** at `a = 0.73` down to **1.49** at `a = 0.93`, **and every one of those is below the current stock of 9.0. Headroom is negative: this council cannot grow stored order in any month, by any decision it takes, and it is losing what its predecessors built.** The break-even alignment is `a_crit = 0.9513`, so **the working estimate of 0.83 sits below it and the council is not yet net-destructive — but the interval's upper end, 0.93, is within 0.02 of it. Report the interval. A point estimate here would be the most misleading number in the survey.**

Intervention. Nothing changes the membership, the terms of reference, or the agenda. (1) A thirteen-sided ring table: `D = 13 x 0.70/pi = 2.90 m`, working depth 0.60 m, inner void 1.70 m, footprint 4.33 m^2; **centrality variance 0.00 at 30, 45 and 60 degrees**, every member seeing six of the other twelve within a 45-degree turn. (2) Room: table plus 1.2 m clear circulation needs a **5.30 m** clear diameter, so 5.3 x 5.3 x 3.0 m, `V` = 84.3 m^3, `S` = 119.8 m^2; for a 0.5 s reverberation the required absorption is `A = 0.161 x 84.3/0.5 = 27.1 m^2` sabins, a mean coefficient of 0.227, achieved with an absorbent ceiling raft plus about 8 m^2 of wall panel, **with the thirteen occupied seats contributing 5.85 m^2 of it unaided**; the fan-coil is replaced with a low-noise unit and attenuators, giving a measured background of **34 dBA against a budget of `60 − 20 log10(2.90) − 15 = 35.8 dBA`.** (3) No reserved position and no nameplates; seats drawn each session. (4) Silent-write-first and the cascade test on every substantive item, with `V_pre` and `V_post` in the minute. (5) Round-robin, randomised order, unconditional pass, with

lowest-to-highest ordering where rank is present. (6) Turn-taking measured quarterly with consent. (7) A minority record available to every member without permission and without a second.

Parameters after, re-measured over four sessions.

```
addressee count:  61% to the chair  ->  24%
V_post / V_pre:   0.14              ->  0.52, argument nameable in 7 of 8 items, so f = 0.125
alignment a:      [0.73, 0.93]      ->  a_max = sqrt(0.48) = 0.693
                  mid 0.83              a_min = sqrt(0.125 x 0.48) = 0.245     mid 0.47
N_speak:          6.9  (0.53 n)     ->  10.8   (0.83 n, above the D14 floor)
phi:              0.45              ->  0.38   (fewer repeat items, fewer reopened decisions)

a = 0.47  ->  u = 1 - 0.47^4 = 0.9512
N_eff = 1 + 12(0.9512) = 12.414     S = ln 12.414 = 2.5189
eta S = 1.2594        C = 0.38 x 1.2825 = 0.4873       delta Q = 0.7200
h     = 1.2825 - 0.7200 - 0.4873 = +0.0751
Q_inf = (1.2594 - 0.4873) / 0.08 = 9.65
```

At the interval's ends the steady state runs from **9.92** at `a = 0.245` to **8.44** at `a = 0.693` — **the whole interval now sits at or near the council's current holdings, where before the whole interval sat below them. Headroom crosses from `-0.0146` to `+0.0751` and the sustainable stored order rises from 5.21 to 9.65: the council stops losing what it inherited.**

Note what did not happen. The category count is still 13, conversion still 0.50, decay still 0.08. **The entire gain came from non-parallelism — 0.525 to 0.951 — plus a modest fall in cost, and non-parallelism moved because the furniture and the protocols stopped rotating everyone's contribution toward one seat.** The framework prices crossing and divergence identically at the margin [Curzi, 2026d], so **this is a gain bought by removing a rotation rather than by recruiting difference, and it is available to any assembly that already has the difference in the room.**

The transient, and the reading that is easiest to miss. Time constant `1/delta = 12.5` months, 95 per cent settling at about 38 months. With `Q_0 = 9.0` and `Q_inf = 9.65`, **the stored order barely moves — it was already near the new steady state, having been on its way down toward 5.21. The visible effect of this intervention is a decline that does not happen.** A council measuring only its stock will see a near-flat line and conclude nothing was achieved; **what was achieved is the difference between that flat line and a decay toward 5.21, which over three years is a large number that appears in no record. Measure the inputs** — addressee count, variance ratio, effective speakers, background noise — **all of which moved immediately, all of which are cheap, and all of which are the drive rather than the state.**

What would falsify it. (1) `N_speak/n` below 0.7 at two consecutive measured sessions: the protocols have failed and the geometry did not carry them. (2) **The alignment interval stays above 0.80 after four sessions: the deference is not geometric in origin — it is contractual, or it is about funding — and no table fixes it. Say so and stop, rather than buying more furniture.** (3) Non-parallelism rises but the stock continues to decline through month 24: either conversion is lower than measured or decay is higher, and non-parallelism was never the binding constraint. (4) **The matched control:** a comparable committee in the same institute over the same period, given a refurbished room of equal cost and disruption but a rectangular table with a chair position retained and no protocol change. **If its non-parallelism and effective-speaker count rise comparably, the mechanism claimed here is not the mechanism operating, and the effect belongs to attention rather than to geometry.**

The honest ledger. Demonstrated within the model: an intervention touching only the geometry and the speaking protocols moved non-parallelism from 0.525 to 0.951, took headroom from `-0.0146` to `+0.0751`, and raised the sustainable stored order from 5.21 to 9.65 — a factor of 1.85 — with no change of membership. **Not demonstrated, and each is a real limit.** *(i)* **Alignment is a modelled quantity, not an observed one:** the postulate supplies its form, the protocol supplies an interval rather than a point, and **the interval's lower end rests on**

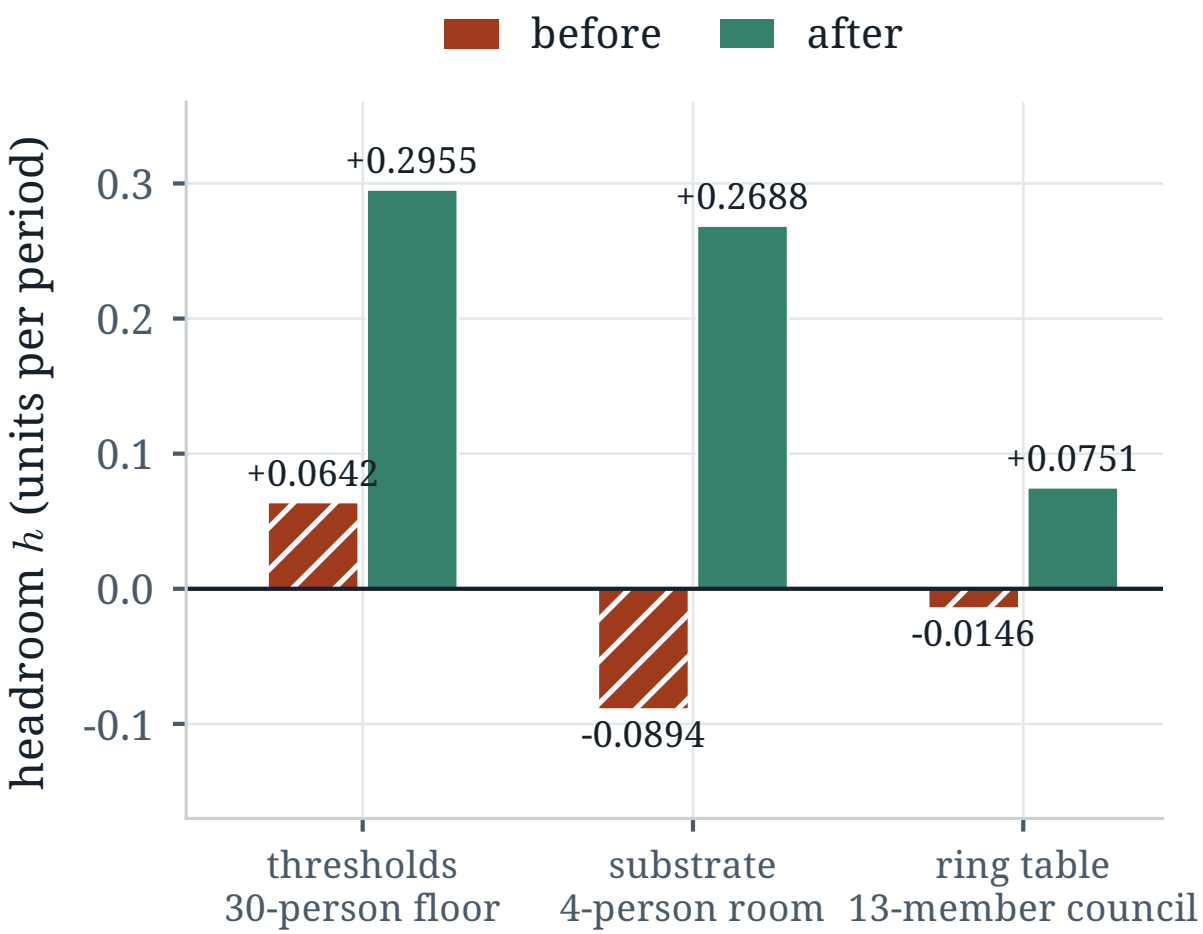


Figure 12.5: **Three levers, three systems, one gauge.** Headroom `h = eta ln N − delta Q − C` before and after each intervention: thresholds on a 30-person floor, substrate in a four-person workroom, a ring table for a thirteen-member council. Modelled — every parameter in all three cases is surveyed or stipulated for that one system, the cost calibration was declared rather than derived, and no figure carries an interval. Generated by `figs/src/CCP-012_fig05.py`.

attributing unexplained convergence to deference, which is a judgement made by the people being measured. The intervals here are wide — `[0.73, 0.93]` before and `[0.25, 0.69]` after — **and they are reported wide rather than narrowed by assumption.** *(ii)* The calibration was chosen and declared, not derived. *(iii)* `N_speak` is a diversity index and not a surplus. *(iv)* The collective-intelligence correlate was measured on short laboratory tasks and its transfer to a monthly standing council is an assumption. *(v)* **Geometry shifts the cost of dominance; it does not forbid it.** If this chair chooses to take 34 per cent of the floor at the round table too, the round table will have failed, **and the measurement is what will say so.**

And one requirement was met only just. D7 is satisfied for the across-table path at 2.90 m — **34 dBA measured against a 35.8 dBA budget, a 1.8 dB margin.** A single additional noise source in that room, a laptop fan or an unattenuated diffuser, puts it over. The project record reads: *D7 met, margin 1.8 dB, no further noise sources permitted in this room without re-measurement.* **A margin that thin is a finding, not a pass.**

12.9 9 Discussion

12.9.1 9.1 What follows

Cost is the cheapest term to move and the least often measured. In all three cases the binding constraint was cost, in none was it the category count, and in two of the three the geometric levers the framework is named for were already near exhaustion — the research floor was running at 88.6 per cent of its own ceiling. **A survey that measures only diversity will report that a well-diversified group has no problem, in a building that is quietly dissolving it.**

Two of the three levers change no personnel, no recruitment and no policy. Thresholds and substrate are joinery and services; the third is joinery plus six lines of meeting procedure. **That is what makes them worth pricing: they are the interventions available to somebody who cannot change the roster**, which is most people who can change anything at all.

Three of the paper's numbers are worth quoting alone. Worst-junction reflection falls from 96.1 per cent to 13.4 per cent across six graded steps, so **a threshold is a matching section and its depth is set by a duty list rather than by a door swing.** The comfort model's own floor is 5 per cent dissatisfied, so **arguing about where to set a shared thermostat is arguing about the wrong variable.** And deference is quartic, so **an assembly loses nothing to ordinary politeness and half of itself to strong alignment, with no line in the minute recording the difference.**

The convergence on table size is the section that will travel furthest and it should be read carefully. Anthropometry, direct-field speech levels and a combinatorial cost argument agree on eight to fourteen seats. **They are three calculations, not three measurements**, and two of them share the proxemic input; the agreement is real and it is weaker than three independent confirmations would be.

12.9.2 9.2 What does not follow

Nothing here shows that a room makes a group cleverer. The collective-intelligence route was tested elsewhere in this series under preregistration and the outcome did not reach threshold, while the mediator moved decisively [Curzi, 2026e]. **Every claim in this paper about turn-taking is a claim about the measured mediator: reduced listening effort, recovered access to the ceiling for members the room was excluding, and a measured rise in the effective-speaker ratio. Those are worth having on their own account and require no causal claim, and a reader who takes away "treat the room and the committee gets smarter" has taken away what this paper explicitly refuses.**

Nothing here identifies a social effect. The specification is reflexive and the endogenous component is not identified from observational data [Manski, 1993]. The matched controls remove a confound and do not solve that problem.

The alignment coefficient is not a measurement. It is a modelled quantity whose estimator returns an interval, and whose lower bound depends on people reporting whether they can name what changed their mind. **An intervention evaluated on a point estimate of alignment has been evaluated on an artefact.**

The transmission-line analogy proves nothing about a corridor. Two sentences transfer and the rest does not, and the failure mode of extending it is named and catalogued.

And the variety term is not a universal good. Unbounded growth in `ln N` is a property of the model rather than of groups: Uzzi and Spiro [2005] report an inverted U in which creative success rises and then falls, and Sagan [1993] documents settings in which redundancy reduces safety. **The model's `ln N` contradicts that literature at large `N`, which is a restriction on the model's domain (Section 10.5) and not a finding.** Nor does an audited category count settle what kind of difference is being counted: separation, variety and disparity are distinct constructs travelling under one word [Harrison and Klein, 2007].

12.9.3 9.3 Relation to the surrounding literature

The design tradition supplies the configurations — entrance transition, intimacy gradient, building edge, alcove, window place [Alexander et al., 1977] — and the semilattice argument [Alexander, 1965] supplies the reason a threshold is a category decision. The proxemic bands [Hall, 1966], the edge effect [Gehl, 1971], the sittable-perimeter finding [Whyte, 1980] and the sightline-and-safety argument [Jacobs, 1961] are the empirical anchors. **What this paper adds to that tradition is a price**: the pattern literature states what to build and rarely states what it costs in the units of anything else, and Section 8 states it in floor area, in interruptions per person-day, and in stored order.

The operations literature supplies the quadratic communication count [Brooks, 1975, Conway, 1968] without pricing it against a logarithmic benefit; this paper does that pricing and finds it converging with two unrelated constraints. The collective-intelligence result supplies the correlate [Woolley et al., 2010] and the mixed replication record supplies the caution [Engel et al., 2014, Bear and Woolley, 2011]; the diversity-and-performance literature and its critics supply the standing warning that a variety-raising programme is not uniformly supported [Mannix and Neale, 2005, Thompson, 2014, van Knippenberg and Schippers, 2007, Hong and Page, 2004, Singer, 2019, Grim et al., 2019, Kuehn, 2017]. The cascade and social-information results supply the mechanism the silent-write protocol defeats [Bikhchandani et al., 1992, Banerjee, 1992, Lorenz et al., 2011].

Where the literature disagrees with the framework, it is cited and the disagreement is stated. And the logarithm itself has two warrants of which this paper needs only the weaker: the Fechnerian one requires a contested perceptual story [Fechner, 1966], while Peters and Gell-Mann [2016] derives a logarithmic evaluation from multiplicative dynamics **with no psychological premise at all.** The functional form remains underdetermined — a sub-linear power law [Latané, 1981] and `n log n` growth [Briscoe et al., 2006] fit the same qualitative facts, and **nothing assembled here discriminates between them** [Stumpf and Porter, 2012].

12.10 10 Verification

The results compound from six postulates. Each is atomic, each is checkable on its own, and each carries its status. **Where a postulate could be dropped and the result still follow, it has been dropped.**

12.10.1 10.1 Dimensional audit

Three dimensional systems meet in this paper: geometry and mechanics, heat, and the storage law's own bookkeeping. `[L]`, `[T]`, `[M]`, `[E]` and `[Theta]` are length, time, mass, energy and temperature; `[Q]` is the dimension of stored order, and the period is the unit of `[T]` in the storage law, which is why no explicit `[T^-1]` appears once that law is written in discrete steps.

Symbol	Dimension	Unit	Status
`Z_2/Z_1`, `Gamma`	`[1]`	none	**defined**: a ratio of like quantities, and its reflection coefficient
`M`, `M_bar`	`[1]`	scale points	**measured**, and **ordinal** — see Section 10.8
`P/A` (edge density)	`[L^-1]`	m per m^2	**derived**: perimeter over area
`d`, `w`, `D`, `s`, `r`	`[L]`	m	**measured** or **specified**

Symbol	Dimension	Unit	Status
`1 + d/w`	`[1]`	none	**derived**: a ratio of two lengths
`v`	`[L T^-1]`	m/s	**measured**
`a` (deceleration)	`[L T^-2]`	m/s^2	**specified**; `d = v^2/2a` returns `[L]` and `t = v/a` returns `[T]`
`k` (switch count)	`[T^-1]`	per person per day	**measured**
`m` (resumption lag)	`[T]`	min	**measured**, with a long right tail
`C_switch = k m`	`[1]`	fraction of a day	**derived**: `[T^-1][T]`, dimensionless — see below
PMV, PPD	`[1]`	index, per cent	**derived** from the comfort model
`sigma`, `A` (authority)	`[Theta]`	K	**assumed** and **specified** respectively
`alpha = k/(rho c)`	`[L^2 T^-1]`	m^2/s	**derived** thermal diffusivity
penetration depth `sqrt(alpha P/pi)`	`[L]`	m	**derived**: `sqrt([L^2 T^-1][T])`
`rho c`	`[E L^-3 Theta^-1]`	MJ/m^3 K	**measured**
effusivity `e = sqrt(k rho c)`	`[E L^-2 Theta^-1 T^-1/2]`	J/m^2 K s^(1/2)	**derived**; the contact-temperature formula is a weighted mean of temperatures and returns `[Theta]`
`RT60`, `A` (absorption)	`[T]`, `[L^2]`	s, m^2 sabins	**measured/derived**; `0.161 V/A` returns `[T]` [Curzi, 2026e]
`psi`, `phi` (angles)	`[1]`	degrees	**derived** from the plan
`V_i`, `Var(V_i)`	`[1]`	count, count^2	**derived** from a layout with a protractor
`a` (alignment)	`[1]`	none	**modelled**, not observed; estimated as an interval
`u`, `N_eff`, `S`	`[1]`	none, count, nats	**defined**, **defined**, **derived**
`N_speak`	`[1]`	count-equivalent	**derived**: `exp` of an entropy in nats
`Q_t`	`[Q]`	proxy units	**measured** in whatever the declared proxy counts
`delta`	`[1]`	per period	**measured** dimensionless fraction per step
`C_t`, `h_t`	`[Q]` per step	proxy units per period	**measured**, **derived**
`eta`	`[Q]` per nat per step	proxy units per nat per period	**required by the balance**; see below
`phi`	`[1]`	fraction of the ceiling	**declared calibration**, `C/(eta ln N)`

Does each equation balance? *The reflection coefficient* is a ratio of like quantities divided by a sum of like quantities and is therefore dimensionless in any medium, which is precisely why the analogy has anything to transfer: **the arithmetic never sees what `Z` is made of.** *The crenellation and edge results* are pure geometry. *Deceleration:* `v^2/2a` gives `[L^2 T^-2]/[L T^-2] = [L]`. *Thermal:* `sqrt(alpha P/pi)` gives `sqrt([L^2 T^-1][T]) = [L]`, and the contact temperature is a weighted mean of two temperatures with effusivity weights, so it returns `[Theta]` for any consistent effusivity unit — **which is why the table's contact temperatures are unaffected by the choice of effusivity unit but its `e` column is not comparable with any other source's unless that source's convention is checked.** *The storage law:* `Q_{t+1} = (1-delta)Q_t + eta S_t - C_t` balances only if every right-hand term carries `[Q]`; the first does since `delta` is dimensionless, the third by construction, and the second **only if `eta` carries `[Q]` per nat per step**, since `S_t` is a dimensionless count in nats.

The source is silent on `eta`'s dimension and the silence is not neutral. The manuscript this paper converts states that `eta` has units of stored order over surplus and that, surplus being dimensionless, `eta` carries whatever unit the order proxy carries. **That is the `[Q]`-per-nat half and it omits the per-step half entirely.** The reading required for the law to balance is `[Q]` per nat per period; **it is stated here explicitly, and it is recorded that the source does not supply it rather than being assigned quietly.**

And one bridge in this paper is an identification rather than an arithmetic identity, which is worth printing because it is easy to miss. `C_switch = k m` is dimensionless — minutes lost over minutes available — while `C_t` in the law carries `[Q]` per period, and the two are joined by setting `phi = C_switch + phi_other` and then `C = phi eta ln N`. **`C_switch` is a fraction of a day and `phi` is a fraction of a ceiling. They are both pure numbers, and that is not the same as being the same number.** The step asserts that the share of available time lost to switching equals the share of the ceiling consumed by that switching, which holds only if stored-order production is proportional to available working time at a rate that does not itself depend on interruption. **That is a modelling assumption, it is used in Section 8.1's headline figures, and it is not defended anywhere in this paper.**

Where a quantity has no operational unit, it is said here. Stored order has no natural unit: it is counted in whatever the declared proxy counts, so two systems' values are not comparable unless they declare the same proxy, **and consequently `eta`, `C` and `h` are not comparable across systems either.** The mismatch index is scale points on an eight-attribute ordinal rating. The alignment coefficient has no unit and **no direct observation at all**. And comfort itself, the underlying construct of the corpus, **has no operational unit whatever**; it is treated as a dimensionless magnitude by stipulation, not by measurement.

12.10.2 10.2 The postulates

P1 — On a transmission line, `Gamma = (Z_2 - Z_1)/(Z_2 + Z_1)` and reflected power is `|Gamma|^2`. *Status:* **CITED** — standard transmission-line theory, used here only for the two structural facts of Section 4.2.

P2 — A built threshold behaves qualitatively like a matching section: crossing cost rises with the ratio of the two zones' characters and falls when intermediate zones are inserted. *Status:* **POSTULATED** — an analogy with no theorem behind it. **It is the load-bearing assumption of Section 4 and nothing in this paper measures whether it holds.**

P3 — Given free choice, occupants select perimeter and interior-edge positions before open-field positions. *Status:* **CITED** — an empirical regularity reported from systematic observation [Gehl, 1971, Whyte, 1980], **not a theorem**, and directly checkable in one building.

P4 — Individual thermal preference is dispersed, and one environmental state cannot sit at every individual optimum. *Status:* **CITED** by the comfort model's own output, which predicts 5 per cent dissatisfied at its theoretical optimum. **The distributional model used to price it — normally distributed preferences with a `+/-1 K` acceptance band — is POSTULATED, and its dispersion is an assumed parameter.**

P5 — Deference acts as a rotation of each contribution toward the chair's by a common coefficient: `v_i = a h + sqrt(1-a^2) w_i`. *Status:* **POSTULATED** — this framework's own modelling choice, with no external support of any kind. **It is the only source of the quartic result, and every number in Sections 6.1 and 8.3 consumes it.**

P6 — The number of other seats within a comfortable head turn is a determinant of the share of a member's contributions that reach the assembly as contributions rather than as asides. *Status:* **POSTULATED**, consistent with the turn-taking literature and with direct observation, and defined in Section 6.3 so that it can be computed from a layout and correlated against measured speaking share.

The claim that a material improving with wear has a lower appearance-decay rate is **not** in this list, because no result in the chain consumes it: it is an interpretation attached to a life-cycle-cost shape, and it is badged wherever it appears.

12.10.3 10.3 The compounding chain

P1	-> L1	reflection depends on the RATIO of the two impedances, so a step 1->2 and a step 500->1000 reflect identically.	CITED
L1	-> R1	dividing a total ratio r into n equal-ratio steps gives per-junction Gamma = (r^(1/n) - 1)/(r^(1/n) + 1); at r = 100, worst-junction reflected power falls 96.1% -> 13.4% from n = 1 to n = 6.	ARITHMETIC
R1 + P2	-> R1'	a graded threshold reduces per-crossing cost, so a character change of M points requires at least M-1 intermediate zones (Rule T-A).	POSTULATED (P2)
(geometry)	-> R6	an interior island contributes its full perimeter and consumes only its footprint, so edge yield per m^2 consumed rises as islands get thinner: six 0.9 x 2.4 m islands add 99% of a 10 x 10 m room's edge for 13% of its floor, against 40% for 16% from one 4 x 4 m island.	ARITHMETIC
R6 + P3	-> R6'	manufactured edge is occupiable amenity, so edge yield is a design objective rather than a curiosity.	CITED (P3)
P4	-> R3	at the comfort model's optimum, PPD = 5.0% by construction; and with preferences N(mu, sigma^2) and a +/-1 K acceptance band, a fixed setpoint satisfies 2*Phi(1/sigma) - 1 while +/-A K of personal authority satisfies 2*Phi((1+A)/sigma) - 1; at sigma = 2.0 that is 38.3% against 86.6% at A = 2.	ARITHMETIC given P4
P5	-> R4	v_i . v_j = a^2, so u = 1 - a^4 and N_eff = 1 + (n-1)(1 - a^4); an n = 12 assembly retains 97.6% of ln 12 at a = 0.5, 63.0% at a = 0.9 and 44.8% at a = 0.95.	ARITHMETIC given P5
R4 + storage law	-> R5	setting eta S = C gives S_crit = C/eta, hence u_crit = (e^(C/eta) - 1)/(n-1) and a_crit = (1 - u_crit)^(1/4); for the worked council a_crit = 0.9513.	ARITHMETIC

P6 + geometry -> R7	at a rotationally symmetric layout Var(V_i) = 0 at every threshold; at an 18-place rectangle it is 9.06 at 45 deg, with the head seats at V = 15 and the far long-side seats at V = 5.	ARITHMETIC given P6
(geometry, acoustics, combinatorics) -> R8	proxemic, direct-field acoustic and combinatorial-cost constraints independently bound the seat count at n <= 13, n <= 13-16 and n* = 7-14 respectively.	ARITHMETIC
R1' + R3 + R4 + storage law -> R2	applied to three worked systems, headroom moves +0.0642 -> +0.2955, -0.0894 -> +0.2688 and -0.0146 -> +0.0751 respectively.	ARITHMETIC given stipulated parameters

Every line is checkable from its own inputs. R1, R3, R4, R5, R6, R7 and R8 are arithmetic consequences of one cited physical result, one cited comfort-model output, elementary geometry, and two postulates. R1' consumes P2 and R4 consumes P5, and those two postulates carry the whole sociological content of the paper: P2 has no theorem behind it and P5 has no external support of any kind. If a built threshold does not behave like a matching section, Section 4 describes an analogy and not a mechanism; if deference is not a common rotation, Section 6's quartic is an artefact of the model and its numbers move.

What this chain does not establish. It does not establish that any organisation obeys the storage law — that is a modelling postulate whose one preregistered confirmatory test in this corpus **failed its second falsification rule, returning a 2.8 per cent held-out advantage over a matched baseline against a declared margin of 10 per cent** [Curzi, 2026i]. It does not establish the causal step from turn-taking to group performance, which was preregistered, tested and not confirmed [Curzi, 2026e]. **A chain whose steps are ARITHMETIC, CITED and POSTULATED delivers the geometry and the physics with certainty and delivers nothing whatever about institutions.**

12.10.4 10.4 The checks a reviewer can redo in a minute

The taper. `100^(1/6) = 2.15443`, so `Gamma = 1.15443/3.15443 = 0.365974` and `Gamma^2 = 0.133937`: **13.4 per cent, against `99/101 = 0.980198` and `0.960788` — 96.1 per cent — at a single junction.** And the geometric-mean check: at `Z_t = 10`, both junctions give `9/11 = 0.818182`, **identical, which is what makes the geometric mean the equalising value.** ARITHMETIC.

Edge yield. Six islands of 0.9 x 2.4 m have perimeter `2(0.9+2.4) = 6.6 m` each, so `6 x 6.6 = 39.6 m` against a base perimeter of 40 m: **+99.0 per cent.** Footprint `6 x 0.9 x 2.4 = 12.96 m^2`, **13.0 per cent of 100 m^2**, and `39.6/12.96 = 3.06` metres of edge per square metre consumed, **against `16/16 = 1.00` for a single 4 x 4 m island.** ARITHMETIC.

The switching cost. `22 x 7 = 154` minutes of 480 is `0.32083`; `6 x 7 = 42` of 480 is `0.08750`. **0.321 and 0.087.** The saving is 112 min/day = 1.867 h, and `30 x 220 x 1.867 = 12,320` person-hours. ARITHMETIC.

The comfort floor. At PMV = 0, `PPD = 100 - 95 exp(0) = 5.0`. At PMV = 1, `0.03353 + 0.2179 = 0.25143`, `exp(-0.25143) = 0.777697`, `95 x 0.777697 = 73.88`, so `PPD = 26.1`. ARITHMETIC.

The adjustability arithmetic. With `sigma = 2.0` and a `+/-1 K` band, satisfaction is `2 Phi(0.5) - 1 = 2(0.691462) - 1 = 0.382925`: **38.3 per cent.** With `A = 2` the band is `+/-3 K`, so `2 Phi(1.5) - 1 = 2(0.933193) - 1 = 0.866386`: **86.6 per cent.** ARITHMETIC.

Sabine on the hard workroom. `0.161 x 121.8 = 19.6098`; `19.6098/4.44 = 4.4166 s`. **The source records 4.41 s; the quotient rounds to 4.42 and is corrected in Section 5.5.** Treated: `19.6098/38.64 = 0.50750 s`, **0.51 s.** ARITHMETIC.

The thermal-mass storage. `2300 x 880 x 0.152 = 307,648` J per m^2 per K, **308 kJ**; over 100 m^2 and a 3 K swing, `307,648 x 100 x 3 = 92.29` MJ `= 25.6` kWh. And the penetration depth: `alpha = 1.7/(2300 x 880) = 8.399 x 10^−7, sqrt(8.399e−7 x 86400/pi) = 0.1520 m`, **152 mm.** ARITHMETIC.

The contact temperature. Cork: `(1200 x 33 + 104 x 20)/(1200 + 104) = 41,680/1,304 = 31.96 C`. Aluminium: `(1200 x 33 + 22,045 x 20)/(1200 + 22,045) = 480,500/23,245 = 20.67 C`. **Eleven degrees at the skin.** ARITHMETIC.

The deference quartic. `0.9^4 = 0.6561`, so `u = 0.3439` and `N_eff = 1 + 11(0.3439) = 4.7829`; `ln 4.7829 = 1.565040` and `ln 12 = 2.484907`, giving `0.62981`: **63.0 per cent.** At `a = 0.5`, `0.5^4 = 0.0625`, `N_eff = 11.3125`, `ln 11.3125 = 2.425913`, ratio `0.97626`: **97.6 per cent.** ARITHMETIC.

The break-even alignment. `S_crit = 0.5771/0.50 = 1.15420`; `e^1.15420 = 3.17168`; `u_crit = 2.17168/12 = 0.180973`; `(1 − 0.180973)^(1/4) = 0.819027^(0.25) = 0.951275`: **0.9513.** ARITHMETIC.

The centrality variance. For `V = [4,4,5,5,6,6,6,6,5,5,5,5,10]`, the sum is 72 and the mean is `72/13 = 5.53846`; squared deviations sum to `2(2.3669) + 6(0.28994) + 4(0.21302) + 19.9112 = 27.2308`, so the variance is `27.2308/13 = 2.0947`: **2.095.** ARITHMETIC.

The head-turn geometry. At 13 seats the neighbour separation is `360/13 = 27.6923` degrees, so `psi = 90 − 13.846 = 76.15` degrees: **76.2, a torso turn.** The opposite member is at 180 degrees and `psi = 0`. ARITHMETIC.

The table diameter. `D = 13 x 0.70/pi = 9.1/3.14159 = 2.8966 m`, **2.90 m**, and the seat bound `pi(2.90)/0.70 = 13.01`. The Winchester table: `18 ft = 5.4864 m`, `pi D = 17.2354 m`, `17.2354/25 = 0.68942 m` per place. ARITHMETIC.

The acoustic budget. `60 − 20 log10(2.90) = 60 − 9.2480 = 50.75` dBA, so a 15 dB signal-to-noise ratio needs **35.8 dBA**; at 4.4 m, `60 − 12.8691 = 47.13` dBA and the requirement is **32.1 dBA**. ARITHMETIC.

The effective speakers. With a chair at 0.40, three silent and eight sharing 0.60 at 0.075 each, `H = −0.4 ln 0.4 − 8(0.075 ln 0.075) = 0.366516 + 1.554162 = 1.920678` and `exp(1.920678) = 6.826`: **6.83 effective speakers out of twelve.** ARITHMETIC.

The three headroom chains. `0.45 x ln 9 = 0.988749`; `0.988749 − 0.36 − 0.5646 = +0.064149`, and `0.988749 − 0.36 − 0.3332 = +0.295549`. `0.28 x ln 6 = 0.501694`; `0.501694 − 0.28 − 0.31105 = −0.089356`, and `0.44 x ln 6 = 0.788376`; `0.788376 − 0.22 − 0.299583 = +0.268793`. `0.50 x ln 13 = 1.282474`; `1.282474 − 0.72 − 0.577113 = −0.014639`, and `1.282474 − 0.72 − 0.487340 = +0.075134`. **All three reproduce to four decimal places.** ARITHMETIC.

12.10.5 10.5 Limits, asymptotics, and the model's valid domain

`M -> 0` (zones of identical character). The mismatch index is zero, the taper arithmetic is vacuous and no device is required. **A building of uniform character needs no thresholds and has no intimacy gradient — which is exactly the open-plan condition, and is why removing all internal thresholds in the name of collaboration is a recognised failure mode (FT6).**

`M -> 5` with a single junction. Reflection approaches its maximum and the transition approaches a wall. **The model has no term for a person who simply does not enter, so the arithmetic under-reports the cost of a severe mismatch**, and the barrier analysis is where the missing term lives [Curzi, 2026p].

`n -> infinity` intermediate zones. Worst-junction reflection tends to zero and the tapered transition becomes ideal — **and the built version becomes an infinitely long corridor.** The arithmetic has no depth cost and the building does: **the domain restriction is the entrance-area budget of 4 to 9 per cent of served net floor, and past it the transition consumes the thing it was protecting.**

`d/w -> infinity` and island thinness `-> 0`. Edge multiplies without bound. **Taken literally, a room can have unlimited occupiable perimeter, and it cannot:** the clearance floor of about 0.9 m between islands and the seating floor on recess depth bound it, **and below those floors the**

added perimeter is texture rather than amenity. This is the same divergence the crenellation result carries, restricted by the same kind of absolute-scale rule [Curzi, 2026e].

`k -> 0` or `m -> 0`. The switching cost vanishes, which the model permits and no organisation achieves; the residual cost terms — meeting overhead, tooling friction, compliance — are what remain, and in the worked floor they are 0.25 of the ceiling, **so the switching lever cannot take cost below that floor no matter how good the thresholds are.** **`k m -> 480 min`** consumes the entire working day; the model reports fractions above 1 for `k m > 480`, **which is outside its domain and means only that the interruptions overlap.**

`sigma -> 0` (identical preferences). A fixed setpoint satisfies everyone within the model, and personal authority buys nothing. **That is the regime in which optimisation is the right answer, and it is a limit rather than a case: measured dispersions in the literature are in the 1.5–2.5 K region, and Section 5.3's argument is conditional on that.** **`sigma -> infinity`.** Everyone is dissatisfied under any single setpoint and satisfaction under personal authority `A` tends to zero as well, since the acceptance band is finite. **Adjustability is not a solution to unlimited dispersion, only to ordinary dispersion.**

`A -> infinity` (unlimited personal authority). Satisfaction tends to 100 per cent, **which is false**: the channels are bounded by energy, by spillover onto neighbours, and by the fact that a personal control that changes a shared variable is a shared control. **The stated domain is `A <= 3 K` with personal-in-effect channels only.**

`a -> 0` (no deference). `u -> 1`, the effective count reaches `n`, and the assembly attains its ceiling. **`a -> 1` (perfect deference).** `u -> 0`, the effective count collapses to 1, and surplus is `ln 1 = 0`; the steady state becomes `Q_inf = -c/delta`, **in which neither the conversion coefficient nor the category count appears — two of the four levers have exactly zero derivative** [Curzi, 2026]. **That is the limit that makes deference worth treating as a separate lever: in it, every other lever is disconnected.**

`n -> 1`. `ln n = 0` and there is no assembly. **`n -> infinity`.** `eta ln n` grows without bound, so the model says a sufficiently large committee makes any cost sustainable, **and that is false.** Three things bound it and **all three are domain restrictions rather than results.** *Internal:* the three sizing constraints of Section 6.4 bind at 13 to 16 seats. *External and contrary:* Uzzi and Spiro [2005] report an inverted U rather than unbounded growth, and Sagan [1993] documents settings in which redundancy reduces safety. *Marginal:* `ln n - ln(n-1) ~ 1/n` vanishes, **so the divergence is the slowest available route to anything. The stated domain is `n` bounded above by the sizing constraints, and the logarithmic ceiling is a within-domain bound rather than an asymptotic claim.**

`delta -> 0`. The fixed point disappears, the stock accumulates without bound, and the model leaves its domain; **`delta -> 1`** makes the time constant vanish and the system stops being a stock. **The stated domain is `delta` in `(0, 0.5]` per period.**

Stated domain. The results are valid for mismatch indices on the declared eight-attribute ordinal scale, thresholds with `d >= 2.0 m` where a stop is expected, island clearances at or above 0.9 m, `k m < 480` min/day, thermal dispersions in `[1.0, 2.5]` K with personal authority at or below 3 K, alignment coefficients in `[0,1]`, seat counts in `[6, 16]`, `delta` in `(0, 0.5]`, and `phi` in `[0,1]`. **Outside that domain the equations still evaluate and should not be interpreted.**

12.10.6 10.6 Invariance

The reflection coefficient is invariant under a common rescaling of both impedances. Multiply `Z_1` and `Z_2` by any positive constant and `Gamma` is unchanged, **which is the whole content of Fact 1 and the reason the analogy transfers at all: the coefficient never sees the absolute character of either zone.** What would break it: any use of the *difference* of the two characters instead of their ratio, **which is exactly what the mismatch index `M` computes.** `M` is therefore NOT the invariant quantity the analogy licenses — **it is an ordinal proxy for the ratio, adopted for usability, and this substitution is unjustified within the analogy and is stated here rather than hidden.**

Edge density scales as `1/c` under a uniform rescaling of a plan by `c`, since perimeter is linear and area quadratic. **That is not an invariance but a systematic scaling, and it is the whole of Section 4.1's argument.** The multiplier `1 + d/w` is scale-free and the seating and clearance floors are not; **the second is what stops the first proving too much.**

Head-turn geometry is invariant under table diameter. `psi = 90 − phi/2` depends only on the angular separation, so **the sightline result is identical for a 1.5 m and a 5.5 m round table** — which is why roundness buys equality and does not buy conversation, and why the acoustic constraint must be applied separately. **Centrality variance is likewise invariant under uniform scaling and under rotation of the whole layout, and is NOT invariant under the choice of head-turn threshold, which is why Section 6.3 reports it at three thresholds and why the hollow square's zero at 45 degrees is a coincidence rather than a symmetry.**

Unit rescaling of the storage law. With `Q' = cQ`, the law is invariant provided `eta`, `C` and `h` rescale and `delta`, `N`, `u`, `phi` and `a` do not. **The dimensionless group — `delta`, `u`, `N`, `a`, `phi`, `M`, `d/w`, `Gamma` — is comparable across systems; the unit-carrying group — `Q`, `eta`, `C`, `h` — is not, unless a common proxy is declared.**

How many scales this has been exercised at. Three constructed systems, at 42 m^2, 900 m^2 and one committee room; three table geometries computed and none built; one materials table computed and none measured. **That is one author, three constructed cases and no independent replication.** And **the three-way convergence of Section 6.4 is three calculations rather than three measurements, two of which share the proxemic input and the assumed speech level**, so it is weaker than three independent confirmations would be — a point stated here because the convergence is the section's strongest-sounding claim.

12.10.7 10.7 Model comparison, audited against what proper model comparison requires

This paper fits no model to data. All three worked systems are constructed with declared parameters, and no parameter is estimated by minimising anything. **Measured against the full requirement, the position is as follows, and three of the five items are not met.**

Requirement	Status in this paper
Out-of-sample rather than in-sample performance	**NOT APPLICABLE.** Nothing is fitted; there is no sample and no holdout
A matched null with the same free-parameter count	**NOT MET.** Each worked case specifies a matched control *design* — an equal-cost refurbishment with the mechanism absent — **but no alternative model of the outcome is specified or compared.** A design control is not a model comparison
The number of free parameters stated explicitly	**MET.** Section 8.1 fixes six (`Q`, `N`, `u`, `eta`, `delta`, `phi`) plus two measured cost inputs; 8.2 fixes six plus `phi`; 8.3 fixes six plus the alignment interval. **All are surveyed or stipulated, none is fitted, and none carries an interval except the alignment coefficient**
An information criterion, so the parameter count is paid for	**NOT MET.** No AIC or BIC was computed here, and none has been computed anywhere in this corpus, for any model

Requirement	Status in this paper
Effect sizes with intervals, not only p-values	**NOT MET.** No p-value is reported anywhere, and no interval accompanies any headroom, steady-state or parameter figure **except the alignment coefficient, which is reported as an interval precisely because a point estimate of it would be misleading**

The unmet items change how the headline figures should be read and are stated plainly. Section 8.1's fourfold headroom gain rests on six declared parameters and a declared calibration; **six declared parameters can produce almost any trajectory, and the case demonstrates that the procedure runs rather than that the model is right.** A reader who wants results from this paper should take the taper arithmetic, the edge yield, the comfort-model floor, the deference quartic and the centrality variance — **which are ARITHMETIC and CITED — and should not take the headroom trajectories, which are POSTULATED throughout.**

No scaling or heavy-tail claim is made anywhere here, so the distribution-fitting procedure of Clauset et al. [2009] has nothing to bite on. It does apply to the quadratic cost exponent used in Section 6.4, **which is defined by the model and not estimated**, and that status is recorded where the exponent is used and in Curzi [2026b]. Stumpf and Porter [2012] states the deeper problem — many mechanisms produce indistinguishable curves — and it applies to the underdetermination discussed in Section 9.3.

12.10.8 10.8 Scale type, analogy discipline, and the precision of the numbers

Scale type, quantity by quantity. Lengths, areas, volumes, times, temperatures, sound pressure levels and heat capacities are **ratio-scale physical quantities** with natural zeros, and every operation performed on them here is licensed. **Decibel quantities are interval-scale in the logarithmic domain**, so differences of decibels are meaningful and "twice as many decibels" is not, and no such statement is made.

The mismatch index is ordinal and it is treated as more than ordinal in two places. The eight-attribute character vector is a set of ordinal ratings; the maximum of a set of ordinal differences is defensible as a rank statement, **but `M_bar`, the mean of those differences, is not licensed by the scale of its inputs, and neither is the substitution of an ordinal difference for the impedance *ratio* the analogy actually uses (Section 10.6). Both operations are performed and both are stated here rather than defended.** Rule T-A is written on the maximum for exactly this reason.

Stored order and non-parallelism are treated as ratio-scale throughout and the conditions that would license that have not been tested. Michell [1997] argues that psychology systematically fails to test whether its attributes are quantitative at all before measuring them, **and that is the strongest single objection to this framework's founding move.** The formal remedy is conjoint measurement — the test for whether an attribute has additive quantitative structure [Luce and Tukey, 1964, Krantz et al., 1971] — or a Rasch demonstration [Rasch, 1960]. **Neither has been run for any quantity in this paper.** The alignment coefficient is worse placed than either: **it is estimated from a variance ratio of ordinal or interval self-reports, using a relation that follows from the postulate's linear form rather than from data, with a lower bound that depends on respondents reporting whether they can name what changed their mind.** The physical quantities in this paper are not exposed to that objection at all [Stevens, 1946], **and that asymmetry is the reason this paper leans on the geometry and the physics and reports the headroom trajectories as illustration.**

Analogy discipline, restated at the point of exposure. The transmission-line result is an **analogy**, and Section 4.2 states the exclusion. The first-order step response shared by the storage law and a low-pass section is a **shared mathematical form**. Sabine's equation, the diffusion depth and the effusivity contact temperature are **derivations** used within their own domains.

The failure mode is letting the first borrow the certainty of the third, and the concrete guard against it here is that no claim in Section 4 is certified by the analogy: the worked case measures `k` and `m` directly, and the analogy only motivated the design. Cartwright [1999] supplies the general escape — specify the arrangement within which the law is claimed to hold — and Anderson [1972] the reminder that a law true at one scale need not be informative at another.

Precision. The numbers here are of four kinds. *Derived-from-defined* — `1 + d/w`, `Gamma = 0.365974`, `ln 12 = 2.484907`, `D = 2.8966 m`, `psi = 76.15 deg`, `Var(V_i) = 2.0947` — are exact to the precision printed and may be recomputed further. *Physical measurements* — reverberation times, background levels, illuminances, operative temperatures — are quoted to the precision the stated method supports and no further. *Nominal or standard-practice inputs* — absorption coefficients, conductivities, heat capacities, effusivities, the illustrative life-cycle rates — **carry two or three figures that record a published typical value and not a measurement of anything that will be built**, and the effusivity column in particular is not comparable across sources without checking each source's convention. *Illustrative model output* — every headroom figure, every steady state, every `phi` — is the output of stipulated parameters and **carries no measurement uncertainty because there was no measurement**; four decimal places there record the arithmetic, not confidence. **Any application to real data must report fitted parameters with intervals, and the four-decimal presentation of Section 8 must not be carried over to a measured `eta`, where two significant figures with an interval would be the defensible form. The alignment coefficient is the one quantity in this paper reported as an interval, and it is reported that way because a point estimate of it would be the most misleading number the method can produce.**

12.11 11 Limitations

All three worked systems are constructed, not observed. No dataset was collected. Every number in Section 8 is the output of stated parameters under stated models, carried through so that a reader can re-derive it. **They demonstrate that the procedure runs and that it reports its own failures; they are not evidence that any building or committee behaves this way, and they must not be cited as such. Every figure there is an instance, not a constant: the method transfers and the numbers do not.**

The two load-bearing postulates have no external support. That a built threshold behaves like a matching section is an analogy with no theorem behind it and no measurement in this paper; **if it is false, Section 4 describes a suggestive picture and not a mechanism, and the entrance-sequence case is an expensive coincidence.** That deference acts as a common rotation with a single coefficient is this framework's own construction; **it appears in no literature, it is wrong in detail — deference is not uniform across members, private views are not orthogonal, and alignment varies by topic — and the entire quartic result and every number in Section 8.3 consumes it.**

The alignment coefficient is modelled rather than observed, and its estimator is weak. The relation `V_post/V_pre ~ 1 − a^2` follows from the postulate's linear form, not from data. Its upper bound is confounded with legitimate persuasion by evidence. **Its lower bound depends on the people being measured reporting whether they can name what changed their mind, which is a self-report about one's own reasoning and is exactly the class of judgement least reliably introspected** [Nisbett and Wilson, 1977]. The intervals are correspondingly wide — `[0.73, 0.93]` and `[0.25, 0.69]` — **and the interval before the intervention comes within 0.02 of the break-even alignment, which is the difference between a council that adds and one that destroys.**

A dimensional bridge in Section 8.1 is an identification, not an identity. `C_switch` is a fraction of a working day and `phi` is a fraction of the ceiling, and the paper adds them. **That step asserts that the share of available time lost to switching equals the share of the ceiling consumed by it, which requires stored-order production to be proportional to available working time at a rate independent of interruption. It is not defended anywhere here, and it carries**

Section 8.1's headline figures. Section 10.1 states it in full.

The mismatch index is ordinal and is used as more than ordinal. The mean `M_bar` is not licensed by the scale of its inputs, and substituting an ordinal *difference* for the impedance *ratio* the analogy actually turns on is unjustified within the analogy itself. **Both operations are performed.**

Nine claims in this paper carry no bibliography entry, and they are named rather than dressed with a citation that would not support them. *(i)* The 1984 window-view recovery study, described in Section 5.1 with its design and its sample size, is not in this corpus's bibliography; anyone quoting its numbers must read the paper. *(ii)* The field literature on interruption and resumption lag, on which Section 4.6 explicitly declines to supply a universal number. *(iii)* The communication-versus-distance regularity from research-and-development organisations. *(iv)* The predicted-mean-vote comfort model, its dissatisfaction curve and its clothing-insulation tables, all used as engineering-practice values. *(v)* The adaptive-comfort field result that perceived control widens the acceptance band. *(vi)* The ventilation standards whose per-person rates are referenced. *(vii)* The carbon-dioxide-and-cognition literature, whose contested state Section 5.4 summarises without citing either side. *(viii)* The last-speaks-first practice from aviation and clinical settings. *(ix)* The minority-influence result that authentic dissent outperforms an assigned devil's advocate. **Attaching a citation that did not actually support these sentences would be worse than leaving them uncited, because a reviewer will check.**

The collective-intelligence correlate is a laboratory finding extrapolated to standing bodies [Woolley et al., 2010], its replication record is mixed [Engel et al., 2014, Bear and Woolley, 2011], **and the causal step was preregistered, tested and not confirmed** [Curzi, 2026e]. The broader diversity-and-performance literature does not uniformly support a variety-raising programme [Mannix and Neale, 2005, Thompson, 2014, van Knippenberg and Schippers, 2007], and the formal diversity results are sensitive to their assumptions [Hong and Page, 2004, Singer, 2019, Grim et al., 2019, Kuehn, 2017].

The reflection problem outranks every before-and-after result reported here. Applied to a group whose members both influence and constitute the aggregate, the endogenous social effect is not identified from observational data [Manski, 1993, Shalizi and Thomas, 2011], and **the matched controls remove a confound without solving it.** The realistic repair is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001].

Every material value is nominal and no building in this paper exists. The absorption coefficients, conductivities, heat capacities and effusivities are generic published values; **Section 8.2's computed reverberation and contact temperatures are what a measurement would return if the nominal values were right, which is a stronger assumption than any other in the substrate section.** The effusivity table's `e` column depends on a convention that is not stated in the source and is not comparable across references.

`N_eff = 1 + (N-1)u` is not a Hill number and no external literature contains it. `1/HHI` and `exp(H)` are the grounded members of that family [Hill, 1973, Jost, 2006]; the affine count is this framework's own construction, adopted by design axiom [Curzi, 2026]. **The effective-speaker count used in Section 6.5 IS `exp(H)` and is grounded; the surplus law's effective count is not, and the two appear within pages of one another here** [Curzi, 2026]. Relatedly, **the variety measure travels under a word covering three incompatible constructs** — separation, variety and disparity [Harrison and Klein, 2007] — and the audited category counts here are variety while non-parallelism is closer to disparity, **and they are used in the same equation.**

The three-way convergence on table size is three calculations, not three measurements, and two of the three share the proxemic pitch and the assumed conversational level. **It is presented as the strongest result in Section 6 and it is weaker than three independent confirmations.**

The cost calibration was declared, not derived. All three cases express cost as `C = phi eta ln N`, which makes `phi` interpretable only relative to a ceiling that itself depends on the audited category count. **Two systems with different category audits are not comparable on `phi`, and none of the three is comparable with any system outside this paper.**

The framework's reflexivity is not applied to its own instruments. This paper prescribes measuring interruption counts, contested adjustments, addressee counts and speaking shares, and reporting them. **A measure adopted as a target ceases to measure what it names** [Goodhart, 1975, Campbell, 1979]; classifications of people interact with the people classified [Hacking, 1995]; and a framework can make its own assumptions true [Ferraro et al., 2005, Merton, 1948]. **A council told that its speaking-share evenness is scored can raise it without anyone being better heard, and nothing here detects that.** The position taken is the narrow one: parameters are treated as stable only over a declared window and only while the gauge is not itself a reported target [Lucas, 1976]. And the desire-line instrument carries the same hazard in physical form: **paving every trace destroys the measurement that produced it, so a building with no soft surfaces left has lost the ability to say it is wrong.**

No break-even probability was computed for any of the three interventions. Section 4.7 states the form the calculation should take and names the corpus's instrument for it [Curzi, 2026]; **it is not performed here, so the area-against-hours comparison is a point estimate under certainty where a decision under uncertainty was required.**

The surplus functional's form is underdetermined by the evidence. A sub-linear power law [Latané, 1981] and `n log n` growth [Briscoe et al., 2006] reproduce the same qualitative facts as `ln(1 + (N−1)u)`, and nothing assembled in this corpus discriminates between them [Stumpf and Porter, 2012]. **The logarithm is chosen for its axiomatic characterisation, not established by measurement**, and Peters and Gell-Mann [2016] is the strongest available warrant precisely because it needs no psychological premise.

Unmet rigour checks, printed rather than left to a reviewer. *(i) Model comparison.* No AIC or BIC here or anywhere in this corpus; no alternative model of any outcome specified; no out-of-sample test — see Section 10.7. *(ii) Effect sizes and intervals.* Nothing in Section 8 carries an interval except the alignment coefficient, because nothing in Section 8 was measured on a real system. *(iii) Scale type.* Stored order, non-parallelism and the mismatch index are treated as more than their scales license and the conjoint conditions were not tested [Michell, 1997, Luce and Tukey, 1964]. *(iv) Statistical control.* The controls specified are matched-design controls, not permutation or shuffle tests; **no false-positive rate was measured against nominal anywhere.** *(v) Independent replication.* One author, three constructed cases [Open Science Collaboration, 2015, Ioannidis, 2005]; the analysis path was not preregistered [Gelman and Loken, 2013, Simmons et al., 2011]. *(vi) Directional versus point prediction.* Meehl [1967] observes that directional predictions are nearly costless in the social sciences; **this paper does make point predictions with constants — 13.4 per cent, 38.3 per cent, 86.6 per cent, 0.9513, 2.095, 0.689 m — which is the right answer to that objection, and none of them has been put at risk against a real system.** *(vii) The discriminating observation.* Durlauf [2005] and Sobkowicz [2020] ask what observation would separate this account from a neighbouring one. **For the deference model there is one — measured disagreement should move quartically rather than linearly with independently estimated alignment — and it has not been run.**

What would falsify the account. H1 fails on interruption counts that do not fall after a threshold intervention meeting its mismatch targets. H2 fails on occupancy counts showing no edge preference where edge positions are otherwise comparable. H3 fails on measured satisfaction under personal control no higher than under an optimised fixed setpoint in a population with dispersion above 1.5 K. H4 fails on measured disagreement moving linearly rather than quartically with independently estimated alignment. H5 fails on speaking-share inequality uncorrelated with computed centrality variance across layouts.

Nothing in this paper is a human-subjects finding. No interventions were run, no participants recruited, no meetings recorded and no thermostat logs collected. **The consent requirement of D14, the unconditional pass of D13 and D17, and the declination and crossability constraints of T14 are design constraints on prospective work and are not reports of any study performed.**

12.12 12 Conclusion

Three levers in physical setting are given quantitative form and priced against the same gauge, `h = eta ln N - delta Q - C`, and in all three constructed cases the binding term was cost.

A threshold is a matching section. Reflection depends on the *ratio* of two zones' characters and not on their difference, and intermediate sections reduce the worst junction: **dividing a 1-to-100 change into six equal-ratio steps drops worst-junction reflected power from 96.1 per cent to 13.4 per cent**, which translates into the one-step rule that no threshold may change any attribute by more than one point on an eight-attribute ordinal scale. **A threshold at which people are expected to stop must be at least 2.0 m deep, because comfortable deceleration at 0.5 m/s^2 from 1.4 m/s takes 1.96 m; and a bent entrance is worth more than the metres it costs, because a 90-degree turn converts 1.5 m of path into about 2.1 s instead of 1.1 s while delivering two further duty items. Edge is the cheap variable and interior the expensive one: six thin islands add 99 per cent of a 10 x 10 m room's edge for 13 per cent of its floor, against 40 per cent for 16 per cent from one fat island.** In the worked floor, **interruptions fell from 22 to 6 per person-day at an unchanged 7-minute resumption lag, cutting the switching cost from 0.321 to 0.087 of the working day, raising headroom from `+0.0642` to `+0.2955` and the sustainable stored order from 5.185 to 9.041, for 68.8 m^2 — 7.6 per cent of net floor — with no change to the people, the work or the diversity of the group.**

Adjustability beats optimisation, and the argument is arithmetic. The standard comfort model predicts **5 per cent dissatisfied at its own theoretical optimum**, so the usual acceptance criterion accepts one person in ten uncomfortable by design in a building working exactly as specified. **At a 2.0 K dispersion of preferred operative temperature, a fixed setpoint satisfies 38.3 per cent and two degrees of personal authority satisfies 86.6 per cent**, and the case strengthens rather than weakens as the population becomes more varied. **A substrate retrofit took reverberation from 4.42 s to 0.51 s while retaining 60 per cent of the thermally coupled soffit — the acoustic and thermal requirements are compatible and are almost never coordinated — and moved headroom from `-0.0894` to `+0.2688` with a 3.7-fold rise in sustainable stored order.** And **at the same 20 C an aluminium handrail presents 20.7 C to the hand and a cork one presents 32.0 C: eleven degrees at the skin from a material choice alone, which no thermostat setting produces and none compensates for.**

Deference is quartic. Under the badged rotation postulate, `u = 1 - a^4`, so a twelve-member assembly retains **97.6 per cent of its ceiling at `a = 0.5`, 63.0 per cent at `a = 0.9` and 44.8 per cent at `a = 0.95`. Ordinary politeness is nearly free and strong deference is catastrophic, and nothing in the paperwork records the loss. Geometric equality has an exact computable definition** — zero variance in the number of other seats within a comfortable head turn — **and it is zero only for a rotationally symmetric arrangement at every threshold, against 9.06 for an eighteen-place boardroom table whose head seats command fifteen of the other seventeen. Three constraints that do not know about each other — proxemic pitch, direct-field speech level, and the quadratic communication count — converge on eight to fourteen seats**, with thirteen at 0.70 m pitch giving exactly the 2.90 m comfortable social distance. The worked council's **centrality variance went from 2.095 to 0.00, its effective-speaker ratio from 0.53 to 0.83, its estimated alignment from 0.83 to 0.47 and its headroom from `-0.0146` to `+0.0751`, with the sustainable stored order rising from 5.21 to 9.65 — and the entire gain came from non-parallelism, because the furniture and the protocols stopped rotating everyone's contribution toward one seat.**

The paper's counter-result is a refusal and it is stated as plainly as the results. The causal chain from turn-taking equality to group performance was preregistered and tested elsewhere in this series, and it did not confirm: the mediator moved decisively and the outcome did not [Curzi, 2026e]. **This paper therefore forbids reporting geometry or acoustics as a collective-intelligence intervention — it buys the measured mediator only — and applies the same restraint to air quality, where ventilation may be claimed and cognitive performance may not.** Two requirements were recorded as not met rather than relaxed: edge density reached 0.291 m/m^2 against 0.40 because no arrangement of islands fixes a 22 m plate, and

a 7.0 m deep room cannot be daylit at a 2.4 m window head. One was recorded as met with a 1.8 dB margin, **and a margin that thin is a finding, not a pass.**

These are the statements downstream papers should quote rather than paraphrase. **Whether any organisation is described by the storage law is a separate question, it is a postulate rather than a result, and the one preregistered confirmatory test of that law reported in this corpus did not meet its own declared margin** [Curzi, 2026i].

12.12.1 Declarations

Data and code availability

No external dataset was collected for this paper. All numerical results are computed from the stated parameters under the models developed in the text, and every table can be re-derived from the equations and parameter values printed beside it. **The three systems worked in Section 8 are constructed examples with declared parameters, not field observations**; the survey and re-measurement figures reported there are the outputs those models return for the specified interventions and are not records of buildings that exist. **The figures are generated by the scripts named in each caption**, and those scripts contain no data beyond the numbers printed in this paper and the model equations of Sections 4 to 6.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, produced the worked examples and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. No meeting was recorded, no speaking share was measured from a real group, no thermostat log was collected and no occupancy count was taken. The consent requirement of requirement D14 — that no measurement of speaking share be made without the informed, renewable and revocable consent of every participant — the unconditional pass of D13 and D17, and the crossability and code-precedence constraints of T14 are **design constraints on prospective work and are not reports of any study performed.**

Preregistration

This paper reports no preregistered test of its own that was executed on a real system. The falsification conditions attached to each worked case in Section 8 are constructed instances showing the form such a registration would take, including the matched control each would require, and they are labelled as constructed there and here. **The preregistered test of the causal chain this paper declines to claim was run and reported elsewhere in this series, and it did not confirm** [Curzi, 2026e]. The preregistered protocol for testing the storage law against a real system, together with the confirmatory application that **failed** its second falsification rule

— a 2.8 per cent held-out advantage over a matched baseline against a declared margin of 10 per cent — is reported in Curzi [2026i], and any applied claim resting on this paper inherits that result.

12.13 Appendix A Measurement protocols

Protocol T1 — the desire-line survey. (1) Photograph every floor and ground surface from a taped, fixed camera position at the same time of day, weekly for six weeks, so the images register; record wear differentially — floor polish, carpet nap direction, grass, dust, and snow or dew, which are the best desire-line media in existence and free for a few days a year. (2) Photograph every room at end of day *before* it is tidied: chair positions are a desire-line medium, as are mugs, cables, coats and laptops left in situ. (3) Mark handle wear and scuffs at 0.9 m (bag corners) and 0.15 m (feet), and log door-propping — **a habitually propped door is a threshold whose crossing count is high enough that the occupants chose to pay in noise rather than in operations, and that is a datum.** (4) Overlay the six images: traces in all six are structural, traces in one or two are events.

Protocol T2 — acting on it. Pave the structural desire lines, which converts an observation into infrastructure and removes the maintenance cost of repaired grass and re-polished floor. **Honest caution: paving destroys the instrument that produced the trace.** Retain one unpaved surface beside every paved route.

Protocol T3 — measuring the resumption lag. For a task with a natural progress metric — lines written, cases cleared, tests passed — record it in five-minute buckets for at least ten person-days, logging every interruption's start and end. Take the mean rate in the twenty minutes before each interruption and find the elapsed time after it ends at which the rate returns to within 90 per cent of that baseline. **That is `m`. Report the median and the interquartile range: the distribution has a long right tail.**

Protocol T4 — measuring the switch count. A tally card taped to the desk, one mark per involuntary switch, for two weeks. **Self-report undercounts, but consistently, which keeps a before-and-after comparison valid even where the absolute level is not.** Do not use telemetry alone: **it misses the person leaning over the partition, which is usually the dominant channel.**

Protocol B1 — the thermal attrition audit. Collect individual preferred operative temperature by structured self-report — a seven-point sensation scale under measured conditions, at least ten occasions per person across a season. Rank occupants by distance from the delivered operative temperature, then measure by decile: hours present, voluntary attendance at optional gatherings held there, and twelve-month retention. **If distance from setpoint predicts none of the three, the category-count attrition claim of Section 5.3 is false in that building and must be dropped. Report either result.**

Protocol D1 — the written-then-spoken divergence. Before any discussion, each member writes a position on a scaled question — a numeric estimate, a rating, a rank order. Collect. Hold the discussion. Collect a second written position. Compute the variance of the first set, `V_pre`, and of the second, `V_post`. Under the deference postulate's linear form, `V_post/V_pre ~ 1 - a^2`, so `a_max = sqrt(1 - V_post/V_pre)`. **Badge: that relation follows from the postulate, not from data**, and legitimate persuasion by evidence also reduces post-discussion variance, so it is an **upper bound**. In the second collection, ask what changed the member's mind: a specific argument, or nothing they can name. With `f` the fraction of changed positions no member could attribute to an argument, `a_min = sqrt(f (1 - V_post/V_pre))`. **Report the interval and carry its midpoint as the working estimate. A point estimate without its interval is the most misleading number this method can produce.**

Illustrative instance. Thirteen members writing estimates `3, 7, 2, 8, 5, 6, 4, 9, 3, 7, 5, 6, 8` have mean 5.615 and variance 4.391. If the post-discussion standard deviation falls to 0.6, the variance ratio is 0.082 — **91.8 per cent of the group's independent variance has been destroyed.** If the members can point to the argument that did it, that is deliberation working. If they cannot, that is a cascade, and the estimator returns `a ~ sqrt(1 - 0.082) = 0.958`, **past**

the break-even point of Section 6.1.

Protocol D2 — the addressee count. During the meeting, log for each contribution whether the speaker's gaze and body orientation are directed at the chair or at the table. **A room in which over half of contributions are addressed to the chair is a room operating at high alignment.** One observer and a tally sheet.

Protocol D3 — the centrality audit. Take the furniture plan. Mark each seat position and its facing direction. For each ordered pair, measure the angle between the seat's facing direction and the bearing to the other seat, and count those within 45 degrees. **Report the vector `v`, its mean and its variance at 30, 45 and 60 degrees. A layout whose variance has not been computed is a layout whose governance has not been specified.**

12.14 Appendix B Failure-mode catalogue

FT1 — the cliff edge. A single door between zones with `M >= 3`, classically the street door opening onto working space. *Damages* cost, continuously and to everyone crossing. *Signature:* hesitation at the door, entries followed by exits within thirty seconds, over-corrected volume in the first sentences. *Fix:* T2's step limit.

FT2 — the corridor with no pocket. A route of uniform width with no recesses, so any pause is an obstruction. *Damages* the optional activity that is the conversion proxy: **two people who become an obstacle stop talking.** *Fix:* T4's pause-and-pass widths and three corridor widenings.

FT3 — the deep plan. Edge density below 0.2 m/m^2. *Damages* conversion, because the positions people select for do not exist. *Signature:* perimeter desks occupied at twice the rate of interior ones. *Fix:* T6 and T7, **and record the shortfall where the plate depth makes the requirement unreachable.**

FT4 — the reflected visitor. A threshold physically deep enough but delivering none of the duty items 4 to 7: nobody is seen arriving, nobody can assess the room before committing, and there is no way to wait without being conspicuous. *Fix:* design against the duty list, not the depth.

FT5 — the paved meadow. Every desire line paved and no soft surface left, **so the building has lost the ability to record a route it does not already have.** *Fix:* Protocol T2's retained surface.

FT6 — the open plan without impedance. All internal thresholds removed in the name of collaboration. *Damages* cost through the switch count and conversion through the loss of closable zones. *Signature:* headphone use as a proportion of the roster — **occupants privately funding the threshold layer the building did not supply.** *Fix:* T9, T10 and T11 together, **never T10 alone.**

FT7 — the fortified threshold. A threshold that has become a gate: crossing needs a third party's permission. **This is outside the framework's scope and is a prohibition rather than a failure mode.** *Fix:* remove the gate; the distinction is whether the choice stays with the crosser.

FT8 — the unmeasured switch. A threshold intervention justified, and later evaluated, on impressions rather than on `k` and `m`. *Fix:* T13, before and after, **with the control.**

FB1 — the air-temperature-only building. One controlled variable of six. *Damages* cost through persistent complaint and contested adjustment. *Signature:* complaints concentrated at particular seats rather than distributed. *Fix:* B6, measured at the seat.

FB2 — the cold window seat. A seat provided at the glazing, behind glazing that makes it uninhabitable in winter. *Damages* conversion by removing the position with the best light and view for half the year. *Fix:* B7.

FB3 — the horizontal-only lighting sign-off. 500 lux on the desk and 130 lux at the eye. *Damages* decay through the circadian route. *Fix:* B3 — **record both planes.**

FB4 — the permanently closed blind. A glare problem answered by the occupant with the only tool they were given. *Damages* the daylight asset entirely. *Signature:* count the blinds. *Fix:* B5, control glare at source.

FB5 — the sealed ceiling. Acoustic treatment that deletes the thermal-mass strategy. *Damages* cost through peak load, invisibly. **The material is still there; its coupling to the room air is not.** *Fix:* B11 — baffles and rafts absorb without sealing.

FB6 — the 5 per cent floor ignored. An operator treating residual thermal complaint as a personnel problem. **The model says the floor is 5 per cent by construction.** *Fix:* B8, personal channels.

FB7 — the dress code as thermal policy. A handbook clause disabling the free, personal, instantaneous adjustment channel and transferring the load onto a shared, slow, expensive one. *Fix:* recognise that clothing is a control, and stop.

FB8 — the carbon-dioxide overclaim. A ventilation business case built on the cognitive-performance literature. *Damages* the credibility of every other number in the specification. *Fix:* B10 — **claim ventilation only.**

FB9 — the cold handrail. Bare metal where skin meets building in an unheated space. *Signature:* computable in ten seconds from effusivity. *Fix:* B13.

FB10 — the laminate decade. Layered materials specified on first cost, with a step-function replacement stream that gets deferred whenever budgets are tight. *Damages* the appearance decay rate. *Fix:* B14, and price the stream rather than the first cost.

FD1 — the head of table. A seat with all others in its forward cone and the conventional right to open and close. *Damages* non-parallelism quartically. *Signature:* addressee count above 50 per cent to the chair. *Fix:* D1 and D9.

FD2 — the long boardroom table. The head-of-table failure in its most severe geometric form, with centrality variance of 9.06 at eighteen places. *Fix:* D1, D4 and D5.

FD3 — the oversized body. A deliberating group above sixteen. *Damages* cost quadratically while the ceiling grows logarithmically. *Fix:* D6, **and accept that the decision will be made elsewhere by fewer people if the body is not resized.**

FD4 — the inaudible far seat. A table wider than the room's noise floor supports. *Damages* non-parallelism by silencing the far members. *Signature:* speaking share correlating with distance from the table's midpoint. *Fix:* D4 and D7 together.

FD5 — the reserved seat. A round table with nameplates, reintroducing rank into a geometry built to remove it. *Fix:* D10.

FD6 — the cascade. Independent private judgements collapsing during discussion with no argument anyone can name. *Signature:* `V_post/V_pre < 0.25` with a high unattributed fraction. *Fix:* D11 and D12.

FD7 — the decorative minority. Members present, counted, and silent. **They are in the roster and not in the effective count, and the minute records twelve attendees.** *Signature:* `N_speak/n` below 0.7. *Fix:* D13's round-robin with an unconditional pass.

FD8 — the assigned advocate. A rotating devil's-advocate role standing in for real disagreement. **Better than nothing and not a substitute for recruiting someone who actually disagrees**, which is a recruitment problem and cannot be solved with furniture.

FD9 — the protocol that compels. A speaking-order rule with no unconditional pass, or a measurement taken without consent. **This is a prohibition, not a failure mode.** *Fix:* D13, D14 and D17.

FD10 — the unmeasured table. A round table installed and never instrumented, **which damages the ability to know whether any of this worked.** *Fix:* D2, D14 and D15, with the control.

Chapter 13

Turn-Taking, Complex Contagion, and the Critical Boundary-Cost Multiplier: Three Channels into the Surplus Law

Abstract. Three mechanisms that move realised non-parallelism are made numerical and coupled to the storage law. Turn-taking is treated as a spending efficiency rather than a source: realised interaction U = J u_pot, with J the normalised participation entropy, so that at N = 5 and u_pot = 0.60 a severely dominated meeting delivers 0.7450 nats against 1.2238 at equality, a 39% loss, while even perfect equality reaches only 76% of the ln 5 ceiling. A worked multidisciplinary meeting measured J = 0.52, S = 0.8879 nats (46% of ln 6), positive headroom of +0.0211 and a negative step of -0.2952; the surplus required for stability, U = 0.93, is infeasible on a roster whose capacity is 0.55, so protocol alone cannot repair it and only a cost cut from 0.42 to 0.20 produces growth. Contagion is treated as reinforcement-limited: seeding 10% of a population puts 57% within reach of one adopter but only 19% within reinforcing range, so clustered seeding dominates spraying. The silo is priced by a critical boundary-cost multiplier m*; at rho = 0.45 and eta = 0.34, m* = 3.17 while the case organisation operated at m = 8. A forced integration that genuinely raised contact from 0.03 to 0.45 raised surplus 121% and cost 279%, and destroyed order faster than inaction.

Keywords: participation entropy; complex contagion; clustered seeding; boundary cost; silo; critical multiplier; surplus ceiling; coordination cost

Series: B. Engineering the Surplus **Paper:** CCP-013

13.1 1 Introduction

The surplus functional `F(u) = ln(1 + (N−1)u)` prices the difference a roster carries. It says nothing about how much of that difference reaches the record. **This paper develops three mechanisms that stand between an available difference and a realised one, makes each numerical, and couples each to the storage law `Q_{t+1} = (1−delta) Q_t + eta S_t − C_t`.** The three are turn-taking, contagion, and the boundary between organisational units, and each fails in a different way that looks like a personnel problem and is not.

Turn-taking is a spending efficiency, not a source. Model realised interaction as `U_t = J_t u_pot`, with `u_pot` the non-parallelism the roster carries and `J = H/ln n` the normalised participation entropy of the speaking shares. At `N = 5` and `u_pot = 0.60`, a meeting with one voice at 80 per cent of the airtime delivers **0.7450 nats** against **1.2238** at perfect equality — **a 39 per cent loss purchased entirely by who was allowed to talk, on an identical roster, with nobody hired or fired.** And the ceiling is not reached even then: perfect equality still delivers only **76.0 per cent of `ln 5`**, because `u_pot` is 0.60 and not 1. **Equality spends what the roster carries; it does not create any.** The two-way duel is worse than the single dominant voice — 2.88 effective speakers against 3.32 — **and it is the most flattering of the failure modes, because it produces the appearance of vigorous debate while consuming the airtime the other three would need to enter.**

Contagion is reinforcement-limited, and the arithmetic is unforgiving of the intuitive plan. Under a mildest-possible complex-contagion rule — adopt when two of eight neighbours have — seeding 10 per cent of a population at random puts **57 per cent of everyone within one hop of the new practice and only 19 per cent within reinforcing range.** That gap is the whole distance between a launch that looks successful in week one and a practice that has vanished by month three. It is why clustered seeding dominates spraying even when spraying wins the first-period metric nine to one, and why the order parameter to report is largest connected adopter cluster rather than total adopters: **between mean degree 1.0 and 1.1 — a 10 per cent change in connectivity — the largest connected group in a 220-person firm goes from essentially nothing to 38.7 people.**

The silo is priced by a critical boundary-cost multiplier. Let `rho` be the fraction of cross-boundary pairs actually in contact and `m` the factor by which a cross-boundary exchange costs more than a within-boundary one. Differentiating the one-step increment with respect to `rho` gives

```
m* = eta (N−1) u_cross / [ (1 + (N−1) u(rho)) c_in x (cross pairs) ]
```

and **contact improves the system only while `m < m*`.** At `rho = 0.45` and `eta = 0.34` in the worked case, `m* = 3.17` while the organisation operated at `m = 8`. **Whether cross-boundary contact helps or harms is therefore determined by the boundary cost and not by the willingness, goodwill or collaborative disposition of the people crossing: enthusiasm does not appear in the expression.**

The paper's central counter-result follows directly and is the reason it is worth reading. A forced integration in the worked 60-person organisation raised realised non-parallelism from 0.1155 to 0.3465 — **a genuine, large, successfully executed increase in contact that more than doubled what it set out to double — and it destroyed stored order faster than doing nothing.** Surplus rose 121 per cent; cost rose 279 per cent; the step went from `−0.1953` to `−0.3714`, and 78 per cent of stored order was gone in twelve months against 63 per cent in twenty-four by inaction. **At the un-remediated boundary cost the step is monotonically worse at every level of contact up to complete integration: there is no value of `rho` at which that organisation benefits from more collaboration.** The same target contact, approached by lowering the boundary cost first — one shared instrument, one interface contract, two offered rotations — takes `m` from 8 to 2 and `eta` from 0.34 to 0.46, and the identical contact and identical surplus then grow the stock at `+0.0408` per month. **The entire difference between a destructive reorganisation and a working one is the boundary cost and the conversion rate, neither of which is about how much people collaborate.**

Two further results are worth stating in the introduction because they are what the arithmetic buys that exhortation cannot. First, **the diagnostic split**: in the worked deliberation group, headroom is `+0.0211` while the one-step increment is `−0.2952`. The gauge says the regime is *feasible* and the current geometry does not achieve it, **and confusing the two is why organisations despair when they should retune, or retune when they should stop.** Inverting the requirement shows the group would need realised `U = 0.93` on a roster carrying `u_pot = 0.55`, **so the requirement is infeasible under any conversational protocol whatsoever** — the largest available turn-taking intervention moved the step by 0.1139 and did not change its sign, and

only a cost cut from 0.42 to 0.20 produced growth. Second, **rotation dosage is computed rather than chosen**: at `m = 2` and `eta = 0.46` the step turns positive only just above `rho = 0.30`, so a one- or two-week cross-boundary rotation improves it by 0.012 to 0.023 and leaves it firmly negative. **A one-week rotation is tourism.**

Sections 2 and 3 state the hypotheses and the framework, including the identification problem. Section 4 develops turn-taking, Section 5 contagion, Section 6 the boundary. Section 7 collects the design specifications. Section 8 works three systems end to end. Sections 9 to 12 discuss, verify, limit and conclude. Appendix A gives the measurement instruments and Appendix B the failure-mode catalogue.

13.2 2 Hypotheses

Some results below are exact consequences of stated models and cannot be falsified by observation; they are labelled as derivations where they appear. Others claim that these models describe real groups, and those are stated here with what would defeat them.

H1 — Realised non-parallelism is the roster's available difference scaled by turn-taking equality: `U_t = J_t u_pot`. Predicted: an intervention that raises `J` with the roster held fixed raises measured disagreement dispersion and realised surplus proportionally; in the worked group, `J` from 0.52 to 0.86 takes `S` from 0.8879 to 1.2134 nats. **Falsified by:** measured surplus that does not track `J` at fixed roster, or that tracks it non-linearly in a way a linear interpolation cannot absorb. **The linear form is a modelling choice with no derivation behind it, chosen for its boundary behaviour — at `J = 1` the group realises everything on the roster, and at `J -> 0` a single voice has `u = 0` with itself and produces `ln 1 = 0` — and it should be replaced the moment anyone measures a better one.**

H2 — Average social sensitivity acts primarily on conversion rather than on surplus. What turn-taking governs is how much difference is *emitted*; what social sensitivity governs is how much of what is emitted is received, understood and retained. **Predicted:** groups matched on `J` and `u_pot` but differing in measured social sensitivity differ in the fraction of raised objections that survive into a durable record. **Falsified by:** objection survival uncorrelated with measured social sensitivity across groups matched on participation equality. **This is a badged assignment of a correlate to a parameter and nothing in this paper measures it.**

H3 — Costly practices require reinforced exposure, so clustered seeding dominates uniform seeding at equal budget. Predicted: under a 2-of-8 rule, a seeding budget concentrated into one unit past its own local tipping point produces a larger connected adopter cluster, and higher survival after sponsor withdrawal, than the same budget spread to maximise first-hop reach — even though the spread plan reports 44.5 people in reinforcing range against the clustered plan's five. **Falsified by:** a matched comparison in which the sprayed budget produces equal or greater largest-cluster size at six months, or equal survival through a one-month sponsor withdrawal.

H4 — Whether cross-boundary contact helps is determined by the boundary cost, not by disposition. Contact improves the one-step increment only while `m < m*`. **Predicted:** in an organisation measured at `m > m*`, raising `rho` lowers the step monotonically; in the same organisation after `m` is lowered below `m*`, raising `rho` raises it. The worked case gives `m* = 3.17` at `rho = 0.45`, `eta = 0.34` against a measured `m = 8`. **Falsified by:** an organisation measured above its own critical multiplier whose stored order nonetheless rises after contact is increased with the boundary cost untouched. **This is the paper's most exposed prediction and the one it most wants tested.**

H5 — A shared instrument lowers the boundary cost. The dominant component of cross-boundary cost is translation and disputed measurement, so a single measurement apparatus owned by neither side removes most of it. **Predicted:** median cross-boundary answer latency falls by at least half within four months of a shared instrument entering daily use; in the worked case, 14.6 hours to 3.6 hours, `m` from about 8 to about 2. **Falsified by:** `m` staying at its prior value while a well-built, well-used shared instrument is in daily use — **which would mean the boundary cost lives somewhere this paper did not look, in incentives, in status, in scheduling, and the recommendation to build instruments first is wrong for that system.**

Two claims in this paper are **not** hypotheses and are labelled accordingly. That the three reported correlates of a collective-intelligence factor map onto this framework's parameters — turn-taking onto realised non-parallelism, social sensitivity onto conversion, composition onto the category count — is a **modelling assignment** with no derivation and no test reported here; it is the whole engineering content of Section 4 and it is not proved. That an organisational domain can be described as a basin in a comfort landscape, with a barrier between it and a neighbouring domain, is an **analogy**; Section 6.1 states what it licenses and Section 11 states that it is not a measurement.

13.3 3 Framework

13.3.1 3.1 What is presupposed

The storage law and its gauge. The equation acted on throughout is

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,     S_t = F(U_t) = ln(1 + (N−1) U_t)
```

with headroom `h_t = eta ln N − delta Q_t − C_t`, the maximum one-step increment `r_max = h_t`, the admissibility condition `C_t <= (1−delta)Q_t + eta S_t`, the sustainability ceiling `C_t <= eta ln N − delta Q_t`, the exact transient `Q_t = (1−delta)^t Q_0 + (1−(1−delta)^t) Q_inf`, and the steady state `Q_inf = (eta F(u*) − c*)/delta`, **which the result allows to be negative and does not hide when it is** [Curzi, 2026b]. Where headroom is negative, no interaction geometry prevents decline. The level–rate–delay vocabulary long predates this framework [Forrester, 1961], the tendency of people to misjudge such systems even under aligned incentives is documented [Sterman, 1989], and headroom is this framework's name for organisational slack [Cyert and March, 1963].

The surplus functional. `u(x,y) = 1 − (x.y)^2` in `[0,1]` and `f(u) = ln(1 + (N−1)u)`, characterised by four axioms, with `0 <= F(u) <= ln N`, `F = 0` exactly at parallel and `F = ln N` exactly at orthogonal [Curzi, 2026]. **The affine effective count `N_eff = 1 + (N−1)u` is a design axiom: chosen, not derived**, and its provenance, its bridge to a participation ratio and its measurement-error budget are developed in Curzi [2026c], whose worked case records that a hundred-person department at `u = 0.05` has an effective count of 5.95. **Concavity is what makes the second genuinely different member the largest single gain an assembly will ever buy**, and the logarithmic ceiling is what makes a bigger group a poor answer to a difference problem.

The decomposition. For a pure block structure the framework separates cross-block from within-block non-parallelism, `u = u_cross + u_div`, with the bound `max(f(u_cross), f(u_div)) <= f(u) <= f(u_cross) + f(u_div)`, and prices the two identically at the margin [Curzi, 2026d], whose institute case records `u_cross = 0.72`, `u_div = 0.03`, `u = 0.75` at `N = 6`, a balance index of 0.0206, ceiling utilisation of 0.8696 and **zero within-block dissents in forty decisions.** Section 6 is that case's structural sibling: **a roster rich in cross-block difference that realises none of it.**

Opposition is not difference. The extremality condition is stated on the *squared* alignment, so a pair in perfect opposition and a pair in perfect agreement both read `u = 0`, `N_eff = 1` and surplus zero [Curzi, 2026r]; the same identity in the frequency domain reads `u = 1 − rho^2`,

depending only on `|rho|`, so `rho = +1` and `rho = −1` give the identical result [Curzi, 2026]. **A group that has organised itself into two opposed camps has not preserved its difference; it has spent it, and the arithmetic does not distinguish that state from unanimity.**

Three effective counts appear in this paper and they are different objects. `N_eff = 1 + (N−1)u` is the surplus law's affine count and **it is not a Hill number** — it is a linear interpolation between 1 and `N`, belonging to no order of the family whose canonical members are `1/HHI` and `exp(H)` [Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979]. `E = exp(H)`, the effective number of speakers of Section 4.2, **is** the Hill number of order one, the exponential of Shannon entropy [Shannon, 1948, Pielou, 1966], computed from participation shares. And `1/HHI`, the Hill number of order two, is the concentration instrument used elsewhere in this corpus, where a naive share vector gives 3.41 effective jurisdictions against a bloc-collapsed count of 1.00 [Curzi, 2026], and where the same instrument applied to reachable shares gives 4.96 effective compartments from six unequal ones and cuts maximum single reach from 100.0 per cent to 30.0 per cent [Curzi, 2026]. **The structural rhyme between them is that all three punish redundancy — a member who adds nothing distinct is not counted — and the rhyme is not an identity. They must not be substituted for one another, and `ln E` must never be reported as a surplus.**

The comfort construct. The framework's founding move is that comfort is the felt total of present costs, so the present state is a local minimum of the comfort landscape **by definition rather than by luck** [Curzi, 2026]. **Section 6 is that definition applied to an organisational domain**: a silo is not a failure of will but a position from which every direction out is uphill, and the method that follows from the definition is to lower the neighbouring valley rather than to push.

13.3.2 3.2 Identification: what no amount of measurement recovers

This is the most serious methodological problem in the paper and it belongs in the method. Every mechanism here runs through a group whose members both influence and constitute the aggregate: who speaks shapes what the group decides, what the group decides shapes who speaks next, and the group is made of its members. That is a linear-in-means reflexive specification, and Manski [1993] proves that in such a specification **the endogenous social effect is not identified from observational data** — it cannot be separated from contextual effects or from correlated unobservables however much data is collected. Shalizi and Thomas [2011] reinforce it for network data, and Lyons [2011] makes the same point about a widely cited contagion result: influence and selection are generically confounded in observational network data, **and the contagion sections of this paper are exactly where that confound bites hardest.**

What survives and what does not. The participation entropy, the binomial reinforcement arithmetic, the clustering coefficients, the percolation fixed point, the pair counts and the critical-multiplier derivation are combinatorics and elementary probability computed from stated models; no identification problem arises for any of them. **The interpretation of a measured change in `J`, in `u`, or in adoption as an endogenous social effect is what is not identified**, and no held-out comparison recovers it.

A partial remedy is adopted and its price is paid. Each worked case carries a **shuffled-input control**: the same surplus series randomly permuted across periods, held out and compared. That asks whether the *timing* of surplus predicts the *timing* of order, and it is the control that separates a mechanism from a period of managerial attention. Together with the matched autoregressive baseline it addresses the confound that measuring a group changes it [McCambridge et al., 2014, Rosenthal and Jacobson, 1968]. **It does not solve the reflection problem.** The route that would is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001]; the fixed rosters, declared category sets and declared participation rules of these surveys point toward it without taking it. **Until then, every applied claim in this paper about a group is a modelling position rather than an identified empirical result**, and Durlauf [2005] and Sobkowicz [2020] state the demand that remains open: name the obser-

vation that would distinguish this account from a neighbouring one fitting equally well. The general hazard of over-reading social-scientific prediction is documented [Lazer et al., 2014, Hofman et al., 2017, Watts, 2014, Salganik et al., 2020], and this paper inherits it.

13.3.3 3.3 Analogy discipline

The information-theoretic quantities are used as information theory and nothing more. `ln N` is a bound on an ensemble of `N` equiprobable categories [Shannon, 1948]; the requisite-variety tradition supplies the reason a variety term appears as a ceiling at all [Ashby, 1956]; and `E = exp(H)` is a descriptive diversity index. **Surplus is not identified with Shannon entropy or mutual information, and no thermodynamic entropy law is invoked anywhere.**

The percolation and threshold results are derivations within their own models and analogies outside them. `S = 1 − e^{−zS}` is exact for an Erdos–Renyi random graph [Erdős and Rényi, 1960, Newman, 2018], and the clustering and path-length expressions are exact for a ring lattice and its rewirings [Watts and Strogatz, 1998]. **Real organisational graphs are neither random nor regular**: they are degree-heterogeneous, correlated, and structurally diverse in ways that change both quantities [Barabási, 2016, Ugander et al., 2012], so the numbers here are **order-of-magnitude structure and not predictions for any organisation**, and Section 11 says so at length.

The comfort landscape and the escape-rate reading are analogies and are badged as such. A basin, a barrier, a ridge and an escape rate are borrowed vocabulary; **nothing about activation energy, temperature or Arrhenius kinetics is claimed to hold of a person deciding whether to learn a neighbouring domain's vocabulary.** What the analogy supplies is a hypothesis about the *shape* of a decision — that lowering the destination is different from raising the agitation — and Section 6.1 converts that into a testable statement about forward and reverse rates.

The logarithm has two warrants and this paper needs only the weaker. The Fechnerian warrant requires a contested perceptual story [Fechner, 1966]; Peters and Gell-Mann [2016] derives a logarithmic evaluation from multiplicative dynamics with **no psychological premise at all**. Mirowski [1989] is the standing warning against copying a physics without its conservation law, and Cartwright [1999] supplies the escape this paper takes: specify the arrangement within which the law is claimed to hold, which is the stated domain of Section 10.5.

13.3.4 3.4 Standing honesty clauses

- **As deployed, the surplus functional is a lower bound** on the ensemble's information [Curzi, 2026]. Every number in this paper is a floor.
- **`ln N` is a ceiling per interaction, not a typical value.** A group at 46 per cent of its ceiling is not 46 per cent of the way to a target; it is realising 46 per cent of what its own roster could carry in a single interaction.
- **Diversification benefit depends on covariance, not on count.** The point that a portfolio's risk reduction comes from the covariance structure and not from the number of holdings was stated formally in 1952 [Markowitz, 1952], **and this framework's insistence that `u` and not `N` is what pays is that result restated in another vocabulary.** It is not novel here.
- **A proxy for stored order must not be a repackaging of the interaction measure.** In all three worked cases the order proxy is built from downstream operational facts recorded by people who do not know the interaction measure, while the interaction measure comes from a disagreement audit. **A proxy that is secretly the input will fit beautifully and mean nothing**, and this is the most common measurement error in applying the law.
- **Nothing in this paper is mandated, and mandate is not a weaker version of adoption but a different thing.** No seeding programme, rotation, protocol or measurement here may be required, incentivised through evaluation, or used as a compliance criterion. **A method that works by removing resistance has no coercive failure mode, and its worst**

outcome is that someone declines to use it. The framework's own treatment of the case where a feedback loop runs through another person's will records that seam as unsettled and refuses to close it [Curzi, 2026], and this paper's mandate refusal in Section 5.7 is written against that refusal rather than around it.

- **Where a measurement becomes a target it ceases to measure what it names** [Goodhart, 1975, Campbell, 1979]. The participation index in particular must never be an individual performance target, and Section 11 records that this paper's own instruments are exposed to the hazard they warn about.

13.4 4 Turn-taking as a spending efficiency

13.4.1 4.1 What the collective-intelligence result supplies, and what it does not

Woolley et al. [2010] report a single statistical factor predicting group performance across varied tasks, not predicted by members' average or maximum individual ability, and correlating with three things: social sensitivity, **equality of conversational turn-taking**, and the proportion of women in the group. **The turn-taking correlate is the one an engineer should care about most, not because it is the strongest but because it is the one a room, a protocol and a piece of furniture can move.**

The honest limit is stated here rather than buried. The factor was measured in short-task laboratory settings with small groups; extrapolating it to a standing institutional committee is a modelling assumption and not a finding. The replication and extension record is mixed rather than uniformly supportive [Engel et al., 2014, Bear and Woolley, 2011], and the broader diversity-and-performance literature carries its own well-documented scepticism [Mannix and Neale, 2005, Thompson, 2014, van Knippenberg and Schippers, 2007] against formal results whose conclusions are sensitive to their assumptions [Hong and Page, 2004, Page, 2007, Singer, 2019, Grim et al., 2019, Kuehn, 2017].

The mapping this paper uses is a postulate and it is badged. Turn-taking equality is assigned to realised non-parallelism; social sensitivity to the conversion coefficient; composition to the audited category count. **That assignment is the whole engineering content of this section and it is not proved.** What follows from it is arithmetic; what precedes it is an assumption.

13.4.2 4.2 Turn-taking equality, made numerical

Let a group of `n` members produce speaking shares `p_1, ..., p_n` over an observation window, with `sum p_i = 1`. Define

```
H = - sum_i p_i ln p_i                  (nats)
E = exp(H)                               effective number of speakers
J = H / ln n                             equality index, in [0,1]
```

`J = 1` is perfectly even participation and `J -> 0` is one voice. `E` is an **effective count of speakers** — a headcount audited for the fact that a person who says nothing is not, for the purposes of the discussion, present. **BADGE: `E = exp(H)` is the Hill number of order one computed from participation shares; `N_eff = 1 + (N-1)u` is computed from geometric non-parallelism between category vectors. They are different quantities sharing the exponential-of-a-log-measure form and they must not be conflated.** The reason to notice the rhyme is that it says what kind of quantity is in hand: **both punish redundancy, and both say the same brutal thing — a member who adds nothing distinct is not counted.**

Computed for a five-person group, with `ln 5 = 1.6094`:

participation profile	shares	H (nats)	E = exp(H)	J = H/ln 5
perfectly equal	0.20 / 0.20 / 0.20 / 0.20 / 0.20	1.6094	5.0000	1.0000
mild skew	0.30 / 0.25 / 0.20 / 0.15 / 0.10	1.5445	4.6855	0.9596
one dominant voice	0.60 / 0.15 / 0.10 / 0.10 / 0.05	1.2014	3.3247	0.7465
two-way duel	0.45 / 0.45 / 0.04 / 0.03 / 0.03	1.0578	2.8800	0.6573
severe dominance	0.80 / 0.10 / 0.05 / 0.03 / 0.02	0.7420	2.1001	0.4610

A single person taking 60 per cent of the airtime in a five-person meeting reduces the effective number of speakers from five to 3.32: five salaries are being paid and three and a third people are being convened. At 80 per cent dominance the figure is 2.10 — a conversation between one person and a rounding error, **which everyone present knows without being able to say why.**

Read the fourth row carefully. The two-way duel scores worse than the single dominant voice, at 2.88 against 3.32. Two loud people arguing is not twice as good as one loud person talking; it is worse, because the argument consumes the airtime the other three would need in order to enter, and it does so **while producing the appearance of vigorous debate. The duel is the most flattering of the failure modes and one of the most expensive.** It is also the state in which two opposed positions can look like maximal difference while contributing none of it: the extremality condition is stated on the squared alignment, so perfect opposition and perfect agreement both read `u = 0` [Curzi, 2026r,?].

13.4.3 4.3 From equality to the storage law

The bridge, wearing its badge. A group has an *available* non-parallelism `u_pot` determined by who is on the roster — their training lineages, incentives and priors. It *realises* only some of that in any given interaction, and this paper models the realised value as

```
U_t = J_t * u_pot            J_t in [0,1] the turn-taking equality index
```

This is a modelling choice and not a derived result. It is chosen for its boundary behaviour: at `J = 1` the group realises everything on the roster, and at `J -> 0` only one voice reaches the record, and a single voice has `u = 0` with itself, so `F = ln 1 = 0`. **A monologue produces no surplus no matter whose monologue it is.** The linear interpolation between those endpoints is the simplest defensible form and **should be replaced the moment anyone measures a better one.** With `N = 5` categories and `u_pot = 0.60`:

profile	J	U = J u_pot	N_eff = 1 + 4U	S = ln(1+4U)	% of ln 5
perfectly equal	1.0000	0.6000	3.4000	1.2238	76.0%
mild skew	0.9596	0.5758	3.3031	1.1949	74.2%
one dominant voice	0.7465	0.4479	2.7915	1.0266	63.8%
two-way duel	0.6573	0.3944	2.5774	0.9468	58.8%
severe dominance	0.4610	0.2766	2.1065	0.7450	46.3%

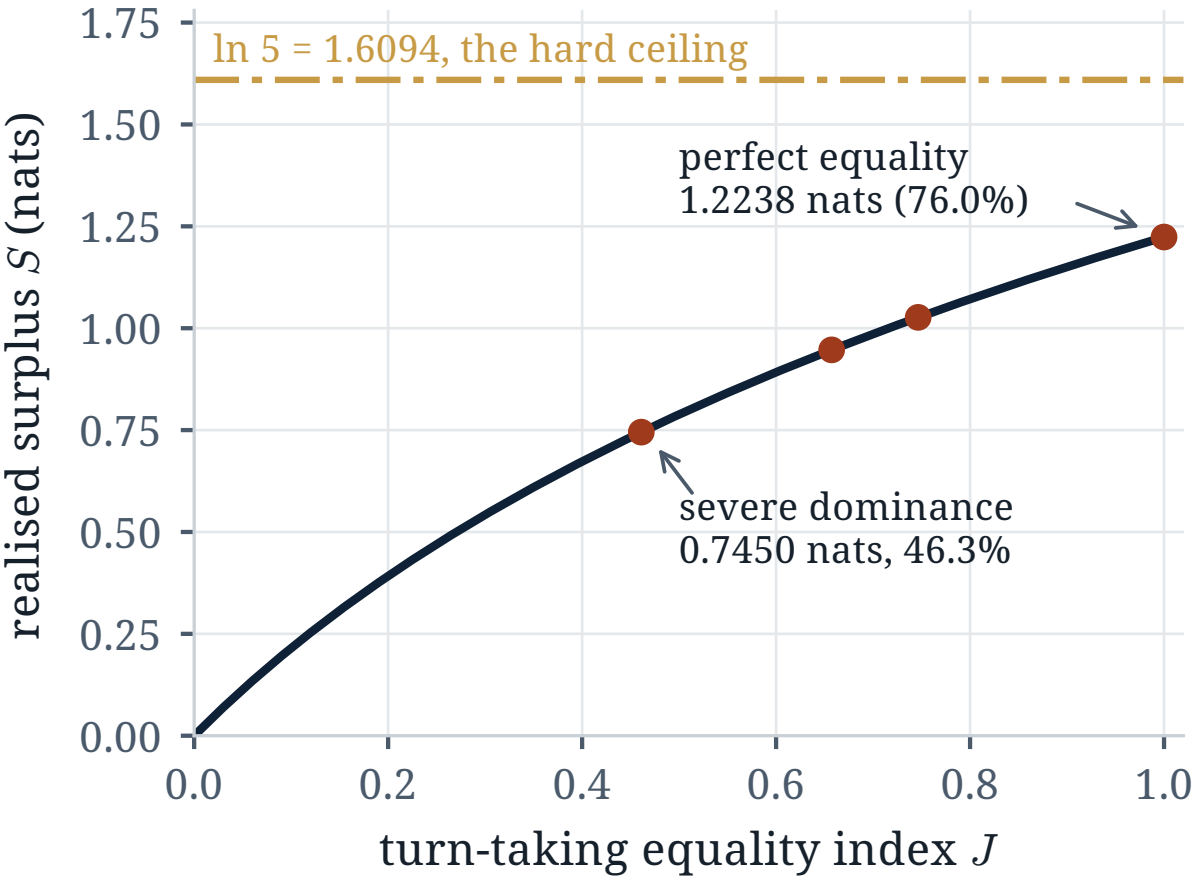


Figure 13.1: **Turn-taking spends what the roster carries and creates none of it.** Realised surplus `S = ln(1 + 4 J u_pot)` at `N = 5` and `u_pot = 0.60`: perfect equality delivers 1.2238 nats, 76.0 per cent of the `ln 5` ceiling, and severe dominance delivers 0.7450 nats, 46.3 per cent — a 39 per cent loss on an identical roster. Modelled: `U = J u_pot` is a badged modelling choice with no derivation behind it and `u_pot = 0.60` is stipulated; the arithmetic given them is exact. Generated by `figs/src/CCP-013_fig01.py`.

Between an even meeting and a dominated one, the same five people on the same roster deliver 1.2238 nats or 0.7450 nats — a 39 per cent loss purchased entirely by who was allowed to talk. Nobody was hired or fired, nobody's ability changed, the roster is identical: **the arrangement moved and the output moved with it.**

Notice what the table does not say. Even at perfect equality the group reaches only 76.0 per cent of its ceiling, because `u_pot = 0.60` and not 1. Equality is a spending efficiency, not a source. If the five are graduates of the same programme with the same incentives, `u_pot` is near zero, `J = 1` multiplies near zero by one, and the perfectly equal meeting produces `ln 1 = 0`.

13.4.4 4.4 Social sensitivity as a conversion parameter, and composition as a category count

Turn-taking governs how much difference is *emitted*. **Social sensitivity governs how much of what is emitted is received, understood and retained — and that is the definition of the conversion coefficient.** A contribution that is made, not understood, and not recorded has been emitted and not converted, and the storage law charges its cost either way. **That assignment is H2 and it is badged: nothing in this paper measures it, and a protocol that makes a signal**

explicit converts a compositional advantage into an architectural one, which is the practical reason to prefer protocol to selection.

The composition finding requires more care than any other line in the source literature, and it is stated exactly. Woolley et al. [2010] report that the proportion of women in a group correlates with the collective-intelligence factor, and that the relationship is **mediated by social sensitivity**. Four things follow and one emphatically does not.

Follows. **The mediation is the operative claim**: a mediated correlate is a marker of the mediating variable, so the design target is the mediator and not the marker. *Follows.* A group composed entirely of one type on any dimension has, by this framework's own argument, a lower audited category count than a mixed one, and the ceiling `ln N` falls accordingly. *Follows.* The finding is a result about a sample, in a laboratory, at one time, and its transfer to a standing institution is an assumption. *Follows.* Its replication record is mixed [Engel et al., 2014, Bear and Woolley, 2011].

Does not follow: that any individual's contribution can be predicted from their category. The result is a statement about group-level correlates and it licenses nothing about individuals. **Requirement CI-17 of Section 7 therefore forbids selecting individuals on any of the reported correlates**, and that prohibition is not a decoration on the argument — **a framework whose category count is its ceiling has an obvious and ugly failure mode if categories are used to sort people, and the prohibition is what keeps the arithmetic from licensing it.**

13.4.5 4.5 What the factor is not: not intelligence, not morale, not cohesion

It is not aggregate individual ability, which is the finding's whole point: the factor was not predicted by members' average or maximum. In this framework's terms, **individual ability sits in what each member brings and the factor sits in the effective count** — a room of highly able people all trained the same way has a high mean and a low `u`, and the surplus law prices the second.

It is not morale, and it is not cohesion — and cohesion is frequently the enemy. Cohesion drives non-parallelism toward zero, and non-parallelism is the quantity the ceiling is built on. A group that has converged is a group whose members' contributions have become parallel, and `f(0) = ln 1 = 0`. **A pleasant, aligned, mutually appreciative committee is, in this arithmetic, indistinguishable from a monologue.** Cohesion has a genuine benefit — it lowers coordination cost — so the question is a crossover rather than a prohibition: **compare the marginal cost reduction against the marginal surplus lost, annually, and record the comparison** (requirement CI-14). What makes it dangerous is that the cost saving is immediate and visible while the surplus loss is delayed by the time constant `1/delta`, **which is why a converged committee looks strong for years and then does not.**

Groupthink, priced. An over-framed group banks *more* stored order than a calibrated one, not less: it produces decisions faster, records them more confidently, and re-litigates them less. **What it loses is veridicality, and the storage law has no term for whether the stored order is correct.** That is a real limitation of the law and is stated here rather than in a footnote: **the law counts protocol, runbook, default and agreement; it does not check them against the world.** A group whose `u` has collapsed will show a healthy stock and a rising gauge while being systematically wrong, **and no amount of computing headroom will detect it.** The corpus's own analysis of reflexive runaways makes the same structural point from the other side — the loop terminates at the ceiling rather than diverging, and its time constant `1/(1−a)` goes tenfold at a loop gain of 0.90, **so the system slows before it destabilises and the slowing is what is observed** [Curzi, 2026ac]. **The detection instrument is not the gauge but the disagreement audit, run on a cadence.**

13.4.6 4.6 The room is part of it

Turn-taking equality is partly a property of the room. In a hard-surfaced reverberant space the loudest speaker wins by acoustics: the gap a following talker must leave for the previous

talker's tail to fall 15 dB is `t_15 = RT60/4`, which is 488 ms in an untreated 1.95 s room against a natural turn gap of about 200 ms, and 98 ms after treatment [Curzi, 2026e]. **A group in such a room will attribute the resulting silence to the people who are silent.** The geometric channel is the same: deference modelled as a common rotation toward one seat gives `u = 1 − a^4`, so an assembly keeps 97.6 per cent of its ceiling at an alignment of 0.5, 63 per cent at 0.9 and 45 per cent at 0.95 [Curzi, 2026f].

The correction that must travel with that claim is stated here in full. The step from raising turn-taking equality to raising the collective-intelligence factor requires a causal direction that the source literature does not establish, **and the preregistered test of it was run and did not confirm: the mediator moved decisively — turn-share evenness from 0.604 to 0.812 against a matched control room flat at 0.638 to 0.650 — and the group-performance outcome did not reach its threshold** [Curzi, 2026e]. **Wherever this paper buys turn-taking equality it is buying the measured mediator: reduced listening effort, recovered access to the ceiling for members the room was excluding, and a measured rise in the effective-speaker ratio. Those are worth having on their own account and require no causal claim. Reporting an acoustic or geometric change as a collective-intelligence intervention is forbidden here.**

13.4.7 4.7 Measuring the proxies

Six instruments are specified in full in Appendix A: participation shares, an interruption ledger, a disagreement audit for `u_pot`, an error-decorrelation measure that is harder and better, a conversion audit measuring objection survival into the durable record, and a cost tally. **What to do with the numbers is as important as taking them. Do not compute a single scalar and call it a group's collective-intelligence factor.** The factor is a laboratory construct measured by a specific battery; what these instruments give are **proxies for the parameters this framework prices**, and the badge is required on every report that uses them (requirement CI-20).

And the standing warning applies with unusual force here. The moment the participation index becomes a target it ceases to measure what it names [Goodhart, 1975, Campbell, 1979]; classifications of people interact with the people classified [Hacking, 1995]; and a framework that assumes the worst of people can make the assumption true [Ferraro et al., 2005, Merton, 1948]. **A meeting in which everyone speaks for exactly their share, because the share is scored, has a participation index of 1 and has learned nothing.** Requirement CI-19 forbids using it as an individual performance target, appraisal input or ranking criterion, **and that prohibition is load-bearing rather than decorative.**

13.5 5 Contagion: reinforcement, clustering, thresholds, and the order parameter

13.5.1 5.1 The reinforcement requirement, made numerical

A practice worth adopting is usually *costly* to adopt: it is worth doing only if others do it, and it is socially expensive to be the only person doing it. That is the definition of a **complex contagion**, in which adoption requires reinforcement from multiple independent contacts, as against a **simple contagion** which spreads on single exposure [Centola and Macy, 2007, Centola, 2010]. **The distinction is the whole of this section, because the seeding plan that is optimal for one is close to worst for the other.**

Model the requirement crudely and honestly: a person adopts when at least `a` of their `k` neighbours have already adopted. Take `k = 8` and `a = 2` — **the mildest possible complex contagion, requiring just one more contact than a virus.** Under *uniformly random* placement of seeds at population fraction `s`:

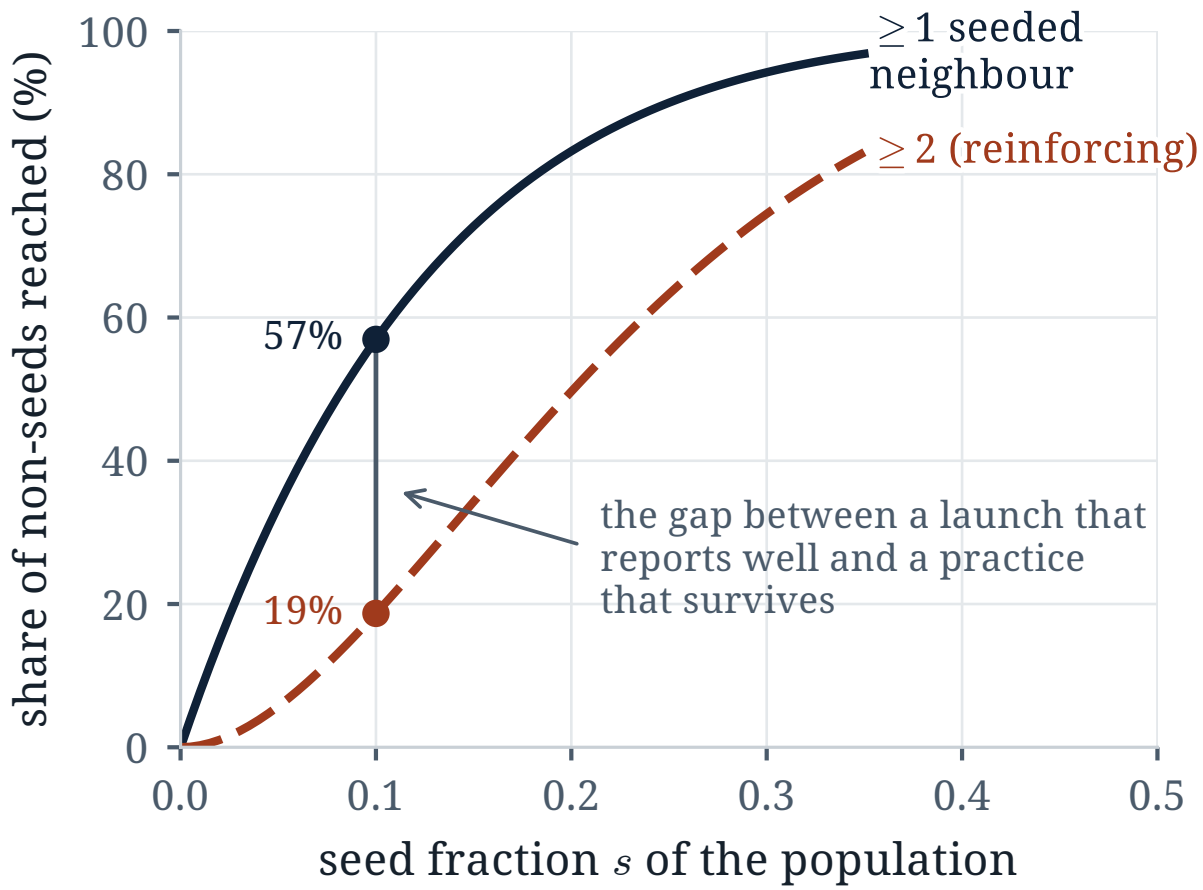


Figure 13.2: **A launch reports its first-hop reach and dies of its reinforcement gap.** Under a 2-of-8 reinforcement rule with uniformly random seed placement, a seed fraction of 0.10 puts 57 per cent of non-seeds within reach of one adopter and 19 per cent within reinforcing range. Exact binomial arithmetic; the 2-of-8 rule is a stipulated model of complex contagion and uniformly random placement is the worst case, which is what clustering improves on. Generated by `figs/src/CCP-013_fig02.py`.

seed fraction `s`	`P(>= 1 seeded nbr)`	`P(>= 2 seeded nbrs)`	ratio
0.05	0.3366	0.0572	0.170
0.10	0.5695	0.1869	0.328
0.114	0.6203	0.2294	0.370
0.15	0.7275	0.3428	0.471
0.20	0.8322	0.4967	0.597
0.25	0.8999	0.6329	0.703
0.30	0.9424	0.7447	0.790

Read the `s = 0.10` row. Sprinkling seeds across 10 per cent of a population puts 57 per cent of everyone within one hop of the new practice — magnificent coverage, and it will be reported as such. It puts 19 per cent within reinforcing range. The gap between 57 and 19 is the entire distance between a launch that looks successful in week one and a practice that has vanished by month three. And note the shape: **raising the requirement from one contact to two — the smallest possible step into complex contagion — cuts effective coverage by a factor of three at low seed density. The penalty is worst exactly where budgets force an operator to work.**

13.5.2 5.2 Clustering, and the trade the modern organisation has already made

The random-placement arithmetic above assumes the worst case: seeds scattered so that no two land adjacent. Real topology is not random, and the property that matters is the **clustering coefficient** — the probability that two of a person's neighbours are also neighbours of each other. Two structures with identical mean degree `k = 8` on `n = 220` people:

```
RING LATTICE (each node joined to its 8 nearest neighbours)
  clustering  C = 3(k−2) / (4(k−1)) = 3(6)/(4(7)) = 18/28 = 0.6429
  mean path   ~ n / (2k) = 220/16 = 13.75 hops

RANDOM GRAPH (same mean degree)
  clustering  C ~ z/n = 8/220 = 0.0364
  mean path   ~ ln n / ln z = 5.394/2.079 = 2.594 hops
```

The lattice is 17.7 times more clustered; the random graph has paths 5.3 times shorter. For a simple contagion the random graph wins overwhelmingly. For a complex contagion the lattice wins, and the mechanism is visible in the numbers: in the lattice, an adopter's neighbour has on average `8 x 0.6429 = 5.14` neighbours who are also neighbours of that adopter; in the random graph the same figure is `8 x 0.0364 = 0.29`. **The lattice hands each newly exposed person five potential reinforcing contacts. The random graph hands them a third of one.**

The rewiring parameter interpolates between them, with clustering falling approximately as `C(p) = C(0)(1−p)^3` [Watts and Strogatz, 1998]:

rewiring `p`	0.00	0.05	0.10	0.20	0.50	1.00
clustering `C`	0.6429	0.5512	0.4686	0.3291	0.0804	0.0000
reinforcing contacts among 8	5.14	4.41	3.75	2.63	0.64	0.00

The small-world regime — the one that makes information travel beautifully — is the regime in which behaviour stops travelling. At `p = 0.10` path length has already collapsed toward the random-graph value while clustering retains most of its lattice value; that is the celebrated sweet spot for *information*. Push to `p = 0.50` and the result is a superb rumour network in which no costly practice can establish, **because a person's neighbours no longer know each other and therefore cannot corroborate.**

The implication for the modern organisation is direct and unwelcome. Every tool that flattens communication — the all-hands channel, the company-wide forum, the cross-functional matrix, the open-plan floor that puts everyone within earshot of everyone — **raises the rewiring parameter.** Each is defended, correctly, on information grounds; each degrades the substrate that costly practices need. **An organisation optimised end to end for information flow has, without noticing, optimised itself against behaviour change.** This is not an argument for removing the tools; it is an argument for knowing which quantity each one moves, and requirement CN-19 requires the trade to be recorded rather than discovered.

13.5.3 5.3 Thresholds and the discontinuous aggregate

The threshold model of collective behaviour assigns each individual a number: how many others must act before they act [Granovetter, 1978]. Its consequence is that **aggregate sentiment is not a predictor of aggregate behaviour**, and three cases make the point.

Case A — thresholds 0, 1, 2, …, 99. One person with threshold 0 acts. That makes one actor, so the person with threshold 1 acts, and so on: **all 100 act.**

Case B — identical, except that the single person with threshold 1 instead has threshold 2. The zero-threshold person acts; nobody has threshold 1; the cascade halts. **1 acts.**

Case C — identical to A, except nobody has threshold 0. **Nothing happens at all: 0 act.**

One person's threshold moved by one unit is the difference between a movement of one hundred and a movement of one. Three consequences follow. **First, aggregate sentiment is not a predictor:** "most people are broadly in favour" tells an analyst nothing about whether anything will happen, because what determines the outcome is the *distribution* of thresholds and specifically whether it has gaps. **Second, the binding constraint is usually a gap rather than a shortage:** in Case B the population is overwhelmingly willing and one rung of the ladder is missing. **Third, the zero-threshold seed is structurally required** — Case C is a population in which every single person would act if someone else went first, and in which nobody ever does. **Requirement CN-8 therefore requires manufacturing the zero-threshold actor deliberately rather than hoping one appears**, and Appendix A gives the threshold survey that finds the gaps.

13.5.4 5.4 Percolation, and the only number worth reporting

The threshold model explains *whether* a cascade runs. Percolation explains **what it looks like from inside while it is running.** For a random graph on `n` nodes with mean degree `z`, the fraction `S` of nodes in the giant connected component satisfies `S = 1 − e^{−zS}` [Erdős and Rényi, 1960, Newman, 2018]:

mean degree z	0.5	0.8	0.9	1.0	1.1	1.2	1.5	2.0	3.0	4.0
giant component S	0.000	0.000	0.000	0.000	0.176	0.314	0.583	0.797	0.941	0.980
nodes of 220	0.0	0.0	0.0	0.0	38.7	69.0	128.2	175.3	206.9	215.6

Between `z = 1.0` and `z = 1.1` — a 10 per cent change in connectivity — the largest connected group goes from nothing to thirty-nine people. Below the threshold there is no giant component at any size; above it one appears and grows fast. **The population on either side of a phase transition experiences it very differently from how it appears in a table**, and the lived consequence is the reason so many good practices are killed.

Below threshold, every observation says failure, and the observation is correct. Adoption clusters stay small, isolated and stagnant; they do not merge; reports say "some enthusiasm in two teams, no broader traction"; a reasonable executive cancels the programme. **Just above threshold, the observations say the same thing for a while**, because the giant component exists and is still small in absolute terms. Then it grows, and the growth is fast, and everyone involved retrospectively constructs a narrative about the moment when people finally came around. **There was no such moment: the parameter crossed a threshold and the visible behaviour lagged it.**

This produces the single most expensive error in organisational change: cancelling at the point of maximum apparent failure, which is frequently the point just before the transition. The defence is not faith but instrumentation:

```
DO NOT MEASURE:  total number of adopters
DO MEASURE:      size of the LARGEST CONNECTED CLUSTER of adopters
                 number of distinct clusters
                 number of edges BETWEEN clusters
```

Total adopters rises smoothly and says nothing about which side of the transition a programme is on. Largest-cluster size is the order parameter and it is the only number that distinguishes "slow" from "stuck": two hundred isolated adopters is a failed programme, and forty adopters in one connected component is a programme about to work.

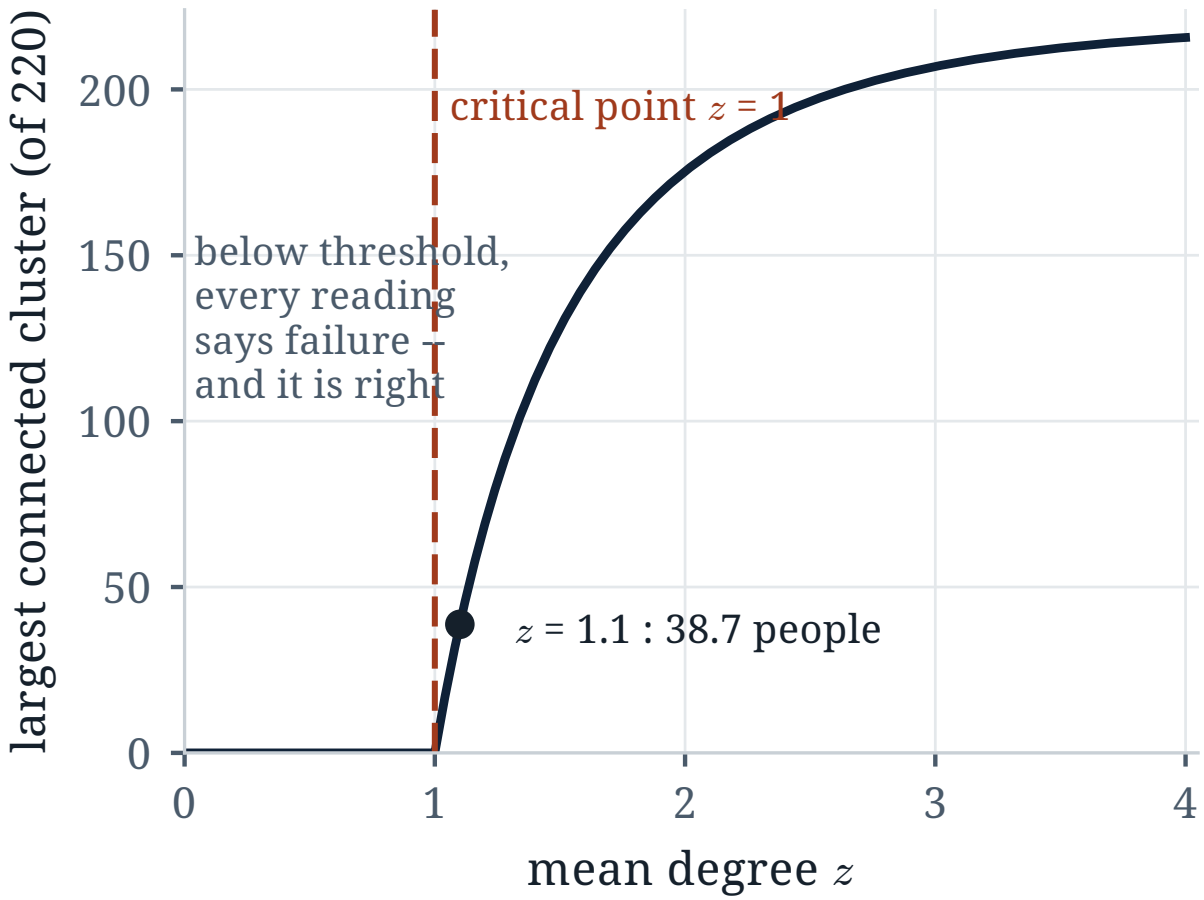


Figure 13.3: **The order parameter is the largest connected cluster, and it moves discontinuously.** Fraction of a 220-person population in the giant connected component of a random graph, solving `S = 1 − e^{−zS}`. Below `z = 1` there is no giant component at any size; between 1.0 and 1.1 the largest connected group goes from nothing to 38.7 people. Exact solution of the branching-process fixed point; the random-graph model is an analytical convenience and real organisational graphs are neither random nor regular. Generated by `figs/src/CCP-013_fig03.py`.

13.5.5 5.5 The tipping-point figure, and the discipline for using it

Centola et al. [2018] demonstrated experimentally that a committed minority below a critical size is absorbed with no residue while one above it flips a whole population, with the critical size measured at about 25 per cent in that setting. **Its value to an engineer is not the specific quarter. Its value is that it demonstrated experimentally that a tipping point exists at all** — before it, tipping points in social conventions were a theoretical prediction; after it, they are a measured phenomenon in at least one carefully constructed setting. **The honest limit: one experiment, one paradigm, online, with anonymous participants, on an arbitrary naming convention with no prior stakes, no status differences, no history and no cost of switching. Use 25 per cent to check whether a plan is off by a factor of ten. Do not use it as a target and never report it as a law** (requirement CN-20).

The practical value is as an order-of-magnitude check, and the table contains the most useful design move in the section:

organisation scope	population	25%	committed people needed
whole firm	220	55.0	55
one division	60	15.0	15
one unit	30	7.5	8
one team	12	3.0	3
one squad	6	1.5	2

SHRINK THE DENOMINATOR. A committed minority of 55 people is a political campaign requiring budget, sponsorship and eighteen months. A committed minority of 3 is a team lunch. **The threshold is a proportion, so the cheapest way over it is not to recruit more committed people but to narrow the population within which the convention must hold.** Establish the convention in one team of twelve with three committed members: that team is now internally past its tipping point, the convention there is self-sustaining rather than sponsored, and **it is also now a cluster in the sense of Section 5.2 — a dense region where anyone adjacent receives multiple reinforcing contacts.** The second team needs three more committed people and some are now available, because they are the first team's members who work partly with the second. **Sequential local tipping is cheaper than simultaneous global tipping by a wide margin, and it is the only version that produces clusters as a side effect rather than requiring them as a precondition.**

13.5.6 5.6 Cluster, do not spray: the calculation

A 220-person firm with a budget for 25 people's time to establish a new practice, mean degree `k = 8`, a 2-of-8 reinforcement rule.

PLAN A — SPRAY. Distribute the 25 across the organisation, roughly two per team, maximising reach: `s = 25/220 = 0.11364`.

```
P(a given non-seed has exactly 0 seeded neighbours) = 0.3811
P(exactly 1)                                        = 0.3907
P(exactly 2)                                        = 0.1753
P(exactly 3)                                        = 0.0449

P(>= 1 seeded neighbour)  = 0.6189     <- the number that gets reported
P(>= 2 seeded neighbours) = 0.2282     <- the number that matters

Expected non-seeds within reinforcing range = 195 x 0.2282 = 44.5 people
```

Note the second row: 39 per cent of the population has exactly one seeded neighbour, which is the worst-populated cell in the table. Those 76 people have encountered the practice,

formed an impression, and their impression is *one person on my team is doing an odd new thing.* **They are not neutral afterwards. A below-threshold exposure is not a free trial: it spends the practice's novelty and returns a null result, and the second attempt to introduce the same practice to the same person starts from a worse position than the first.**

PLAN B — CLUSTER. Put all 25 into a single 30-person unit; local seed fraction `25/30 = 0.8333`. Every one of the five non-seeds in that unit has at least two seeded neighbours with probability 1.0000, so five of five are activatable inside the unit and approximately none outside it in the first period.

Plan A's headline number — 44.5 people in reinforcing range — is nearly nine times Plan B's five, and any reasonable reading of a one-period comparison prefers Plan A. Plan A is the wrong choice, for three reasons a one-period comparison cannot see. *First, the 44.5 are at the exact margin*: they have precisely two contacts and the rule requires two, so any attrition among the 25 seeds — one on holiday, one reassigned, one quietly stopping because their two teammates did not join — drops many of them below threshold. **Plan A's coverage is brittle in exactly the way a two-of-eight rule is brittle.** *Second, Plan A's adopters are not connected to each other*: the order parameter is largest-cluster size, and Plan A maximises the number of clusters while minimising the size of each — **it is engineered, precisely, to sit below the percolation threshold.** *Third, Plan B's unit is past its own tipping point*: 25 of 30 is 83 per cent, far above the floor, so the convention inside that unit is not a programme but how things are done. **It survives the sponsor losing interest, which Plan A does not**, and it is a genuine cluster, so anyone in an adjacent unit who works with two or three people from the seeded unit receives multiple independent reinforcing contacts — **exactly the condition Plan A could not manufacture anywhere.**

Plan A optimises the first-period metric. Plan B optimises the order parameter. The first-period metric is what gets reported and the order parameter is what determines the outcome.

The seeding procedure that follows:

```
1. Pick the smallest population within which the practice makes sense standing alone.
   It must be able to work if NOBODY outside adopts. If it cannot, pick a different practice.
2. Seed that population past its own tipping point. Use 25% as a floor, not a target;
   the experimental figure came from perfectly committed confederates and yours are people.
3. Let it stabilise until the practice survives the removal of its sponsor.
   That is the test. Remove yourself for a month and see.
4. Seed the adjacent population using members of the first as the reinforcing contacts.
   Do not send an evangelist. Send two people who work there anyway.
5. Repeat outward along real working ties, not along the org chart.
```

Step 5 deserves its own emphasis: seed along the graph of who actually works with whom, which is not the graph of who reports to whom. Section 6.5 gives the structural reason. **Seeding along the org chart is seeding along a diagram.**

13.5.7 5.7 What "committed" must mean, and the refusal of mandate

The experimental minority was committed by construction — confederates instructed never to deviate. **Reproducing that property in an organisation is where this method could become something this framework refuses to be**, and the distinction is between two mechanisms that produce apparent commitment and are not equivalent. **Genuine commitment** is a person who has used the practice, found it better, and would continue without instruction. **Compliance** is a person who adopts because adoption is required, evaluated, or visible to someone who can affect their standing.

Mandating adoption defeats the purpose of adoption, and not for ethical reasons alone, though those are sufficient. **A mandated practice produces artefacts near review dates and decisions unchanged;** the measurement detaches from the thing it names [Goodhart, 1975, Campbell, 1979]; and, in this framework's own arithmetic, **mandate drives non-parallelism**

to zero across adopters, because everyone is doing the same thing for the same reason, so the spread practice's own surplus becomes `ln 1 = 0`. The consequence is stated plainly: **100 per cent compliance is a worse outcome than 60 per cent adoption with 40 per cent articulate declines.** The declines carry information about where the practice does not fit, and **recorded declines are the most valuable data a seeding programme generates and the data most often discarded.**

Requirements CN-12 to CN-15 make this operational: no seeding programme proceeds unless a member of the seeded population can decline without consequence; declines are recorded with reasons and reported alongside adoptions; adoption is never required, incentivised through evaluation, or used as a compliance criterion; and before expanding, the sponsor withdraws for a month and the practice must survive. **The framework's own treatment of cases where a feedback loop runs through another person's will records that seam as unsettled and declines to close it** [Curzi, 2026]. **This paper does not close it either. It draws the line where the line can be drawn — at whether declining is free — and states that the harder cases beyond it are not resolved here.**

13.5.8 5.8 Coupling contagion to the storage law

Two couplings are needed and both are badged as modelling choices. **First, the adoption dip.** Most practices worth adopting cost something during transition: the conversion coefficient rises with adoption share and the cost term falls with it, but the cost of adopting is paid before the benefit arrives, **so the stock falls before it rises.** Modelling both as linear in the adopted fraction — `eta(a) = eta_0 + (eta_1 − eta_0) a` and `C(a) = C_0 + (C_1 − C_0) a` — is a convenience with no support beyond tractability, and it is the form used in Section 8.2. **Second, partial adoption can be strictly worse than zero adoption.** If adoption stalls at a level where the organisation maintains two conventions at once, the cost of running both may exceed the benefit of either, and the dual-convention plateau is a stable and expensive place to stop (failure mode **FN5**).

The reachability structure that makes this work is itself measurable, and the corpus's instrument for it is the concentration index applied to reachable shares: six unequal compartments give 4.96 effective compartments and cut maximum single reach from 100.0 per cent to 30.0 per cent, worth 1.400 units a year against 0.110 of cost — **a 12.73 to 1 ratio with a break-even hazard of 0.157 per cent per year** [Curzi, 2026]. **That analysis prices compartmentation as protection against unwanted spread; this section prices the same structure as an obstacle to wanted spread. They are the same arithmetic read in opposite directions, and an organisation cannot have both at once** — which is the trade requirement CN-19 exists to record.

13.6 6 The silo and the critical boundary-cost multiplier

13.6.1 6.1 What a silo is, and why exhortation fails

A silo is a local minimum. The framework's founding definition makes the present state a minimum of the comfort landscape by construction [Curzi, 2026], and a silo is that definition applied to a domain of practice: a position that is **deep** — every direction out is uphill, because inside it a practitioner is fluent and competent and outside it they are a novice — and **narrow**, because it admits one vocabulary, so that a question posed in a neighbouring dialect does not merely go unanswered but is not recognised as a question. **Nothing about that description requires anyone to be territorial, defensive or uncooperative, and the diagnosis that reaches for those words has already misidentified the mechanism.**

The landscape is an analogy and is badged as one. A basin, a barrier and a ridge are borrowed vocabulary; nothing about activation energy or reaction kinetics is claimed. **What the analogy supplies is a hypothesis about the shape of a decision, and it makes one prediction that distinguishes it from exhortation.**

Exhortation fails in two independent ways. *First, exhortation is a tilt that is removed when the speech ends.* Declaring collaboration a value changes nothing about the landscape: the barrier is where it was, and the crossing rate rises during the campaign and returns afterwards. **The diagnostic signature is a forward rate that rises and falls with attention, and the measurement is to check the reverse rate three months after attention has moved on** (requirement SL-18). *Second, exhortation raises agitation, and agitation is undirected.* Pressure applied to a boundary raises the rate at which people cross it **and the rate at which they cross back, in the same proportion, so the forward-to-reverse ratio is unchanged at every level of pressure.** Building a lower valley next door changes the ratio; pushing does not. **A manager watching only the forward rate — how many people attended the cross-functional session — sees the two routes as identical, and they are not.**

The method that follows from the diagnosis is to build a lower valley next door: to make the cross-domain position genuinely cheaper to occupy, rather than to make the present position uncomfortable. **And the acid test is the safety property: if nobody moves, nothing has been done to anyone.** A method that works by removing resistance has no coercive failure mode, and **its worst outcome is that someone declines to use it.** That is not a concession the method makes; **it is the property that makes it safe to recommend**, and it is why requirement SL-16 makes every instrument in this section declinable.

What actually lowers the valley, in the order that the arithmetic of Section 6.3 puts them: **(1) a shared instrument** — one measurement apparatus, one dataset, one dashboard, owned by neither side and readable by both without specialist tooling, which terminates the large class of cross-boundary arguments that are really arguments about what is happening; **(2) boundary objects** — artefacts plastic enough to be read differently by each side while remaining recognisably the same object, which is the class named and analysed in the sociology of scientific work [Star and Griesemer, 1989, Star and Ruhleder, 1996]; **(3) a translator role**, timeboxed with a termination condition and a build list, because a permanent translator has paid the boundary cost rather than lowered it; **(4) rotation**, offered and never assigned; and **(5) co-located cross-function**, on shared output rather than adjacent output. **Instruments 1 and 2 move the boundary cost and the conversion coefficient at once, which is why Section 8.3 finds them outranking rotation by a wide margin.**

13.6.2 6.2 The double indictment, priced

Consider a 60-person organisation in two silos of thirty. Take `N = 6` categories with between-silo non-parallelism `u_cross = 0.55` and within-silo dissent `u_div = 0.0990`:

```
u_cross = 0.5500      f(u_cross) = ln(1 + 5(0.5500)) = 1.3218
u_div   = 0.0990      f(u_div)   = ln(1 + 5(0.0990)) = 0.4021
u       = 0.6490      f(u)       = ln(1 + 5(0.6490)) = 1.4457

decomposition bound: max(1.3218, 0.4021) = 1.3218 <= 1.4457 <= 1.3218 + 0.4021 = 1.7239  HOLDS
```

The silo condition is that `u_cross` is on the roster and is **realised at approximately zero**, because the interactions that would realise it do not occur:

```
Available  S = f(u_cross + u_div) = 1.4457 nats
Realised   S = f(u_div only)      = 0.4021 nats
Forgone                           = 1.0436 nats per interaction  =  72.2%
For reference, the hard ceiling is ln N = ln 6 = 1.7918 nats.
```

The organisation is discarding 72.2 per cent of the surplus its own roster can supply, and the discarded portion is the cross-category portion — the part that is expensive to hire and impossible to substitute. This is not a diversity-of-hiring failure. The difference was purchased. It sits in the building. It is not in contact. The structural sibling of this case is the institute

whose audit returns `u_cross = 0.72`, `u_div = 0.03` and **zero within-block dissents in forty decisions** [Curzi, 2026d]: the same shape from the other direction.

The cost side. On the same sixty people, `paths(30) = 435` and `paths(60) = 1770`, so within-silo pairs number `2 x 435 = 870` and cross-silo pairs `30 x 30 = 900`, and `870 + 900 = 1770` confirms the decomposition [Brooks, 1975]. **More than half of all pairs in this organisation — 900 of 1770 — are cross-boundary pairs.** Model the cost as `C = c_in x (within pairs in contact) + m c_in x (cross pairs in contact)`, with `m` the **boundary-cost multiplier**: how many times more expensive a cross-boundary exchange is than a within-boundary one. At `c_in = 0.00012` with all pairs in contact:

`m`	`C`	cross-boundary share of `C`
1	0.2124	50.8%
2	0.3204	67.4%
4	0.5364	80.5%
8	0.9684	89.2%
16	1.8324	94.3%

A correction to the source is recorded here rather than made silently. The manuscript this paper converts prints this table with the absolute cost column computed at `c_in = 0.0004` — giving 0.7080, 1.0680, 1.7880, 3.2280 and 6.1080 — while its text states `c_in = 0.00012`, and while its own critical-multiplier table and all three worked cases use 0.00012 consistently. **The column above is recomputed at the stated and internally consistent value. The share column is unaffected, because `900m/(870 + 900m)` does not depend on the scale of `c_in`**, and nothing in the argument turns on the absolute column.

At `m = 8` — a plausible figure where two domains share terms with different referents and no shared instrument exists — 89.2 per cent of the organisation's coordination cost is boundary-crossing cost.

Now put the two sides together, because this is where the silo's peculiar viciousness shows. The silo suppresses the interactions that would realise the cross-boundary difference, **and the reason it suppresses them is that `m` makes them expensive. So the organisation is simultaneously forgoing 72 per cent of its available surplus AND, to the extent that cross-boundary contact happens at all, paying a multiplied price for it. Every attempt to fix the first by mandating more contact makes the second worse**, which is why reorganisations aimed at collaboration so often produce a measurably worse system: **they raise contact without lowering `m`, and the arithmetic of that combination is unambiguous.**

13.6.3 6.3 The critical boundary-cost multiplier

Let `rho` be the fraction of cross-boundary pairs actually in contact. Model realised non-parallelism as `u(rho) = rho u_cross + u_div` — **a badged linear modelling choice** — and cost as `C(rho) = c_in x 870 + m c_in x rho x 900`. The one-step increment is `eta ln(1 + 5 u(rho)) − delta Q − C(rho)`; differentiating with respect to `rho` and setting the derivative to zero gives

```
d(step)/d(rho) = eta (N−1) u_cross / (1 + (N−1) u(rho))  −  m c_in x 900

m* = eta (N−1) u_cross / [ (1 + (N−1) u(rho)) x c_in x 900 ]
```

`m*` is the critical boundary-cost multiplier: more contact improves the system only while `m < m*`.

rho	u(rho)	m∗ at eta = 0.34	m∗ at eta = 0.46
0.00	0.0990	5.79	7.83
0.10	0.1540	4.89	6.62
0.20	0.2090	4.23	5.73
0.30	0.2640	3.73	5.05
0.45	0.3465	3.17	4.29
0.60	0.4290	2.75	3.72
0.80	0.5390	2.34	3.17
1.00	0.6490	2.04	2.76

> **Whether cross-boundary contact helps or harms is determined by m, the boundary cost — not by the willingness, goodwill or collaborative disposition of the people crossing.**

If m > m∗, increasing contact makes the system monotonically worse, and it does so no matter how enthusiastically people participate. Enthusiasm does not appear in the expression. This is the precise sense in which collaboration is not a virtue problem.

Two further readings. **m∗ falls as rho rises**, because surplus is concave in u while cost is linear in rho: **the first cross-boundary contacts are the ones that pay, full integration is rarely optimal, and the two policies that fail are "integrate everything" and "integrate nothing".** And **m∗ rises with the conversion coefficient** — at eta = 0.46 instead of 0.34 the tolerable boundary cost is about 35 per cent higher at every level of contact — **so raising conversion makes contact affordable that was not affordable before.** Since a shared instrument raises conversion *and* lowers boundary cost at once, **it moves both sides of the inequality in the favourable direction, which is why it ranks first among the instruments.**

13.6.4 6.4 The contact literature, in this vocabulary

The contact hypothesis holds that contact between groups reduces prejudice **under four conditions**: equal status within the situation, common goals, intergroup cooperation rather than competition, and **the support of authorities, law or custom** [Allport, 1954]. The meta-analytic record supports a general effect and finds the conditions to be facilitating rather than strictly necessary [Pettigrew and Tropp, 2006], while the broader literature on diversity and contact carries genuine contrary findings that must be cited alongside [Putnam, 2007, Abascal and Baldassarri, 2015]. **In this framework's vocabulary the fourth condition is the one that does the arithmetic work: institutional support is what lowers m.** A cross-boundary exchange that has a budget line, an owner, a time allocation and a shared instrument is cheap; one that is done in gaps by volunteers is expensive, **and it is the first thing dropped under load** (failure mode **FS8**). **Contact without institutional support is not a weak version of the intervention. It is the intervention with the term that pays for it removed**, and requirement SL-13 requires the four conditions to be checked explicitly and the absent ones recorded.

13.6.5 6.5 The structural cause, and why boundaries are legitimate

Organisations produce designs that mirror their own communication structures [Conway, 1968]. **Three consequences follow, and the third is the deepest.**

First, the silo is downstream of a structural decision and is not repairable upstream of it. If the organisation is partitioned by function, its outputs will be partitioned by function, and no amount of cross-functional exhortation changes the partition that produced them. **Second, boundaries are legitimate and necessary**: the cost arithmetic requires them, since coordination cost grows with the pair count [Brooks, 1975], and the corpus's roster ceiling states the same limit exactly — m(m−1) <= 2 eta ln(N)/kappa, with a constructed commons sitting at 30

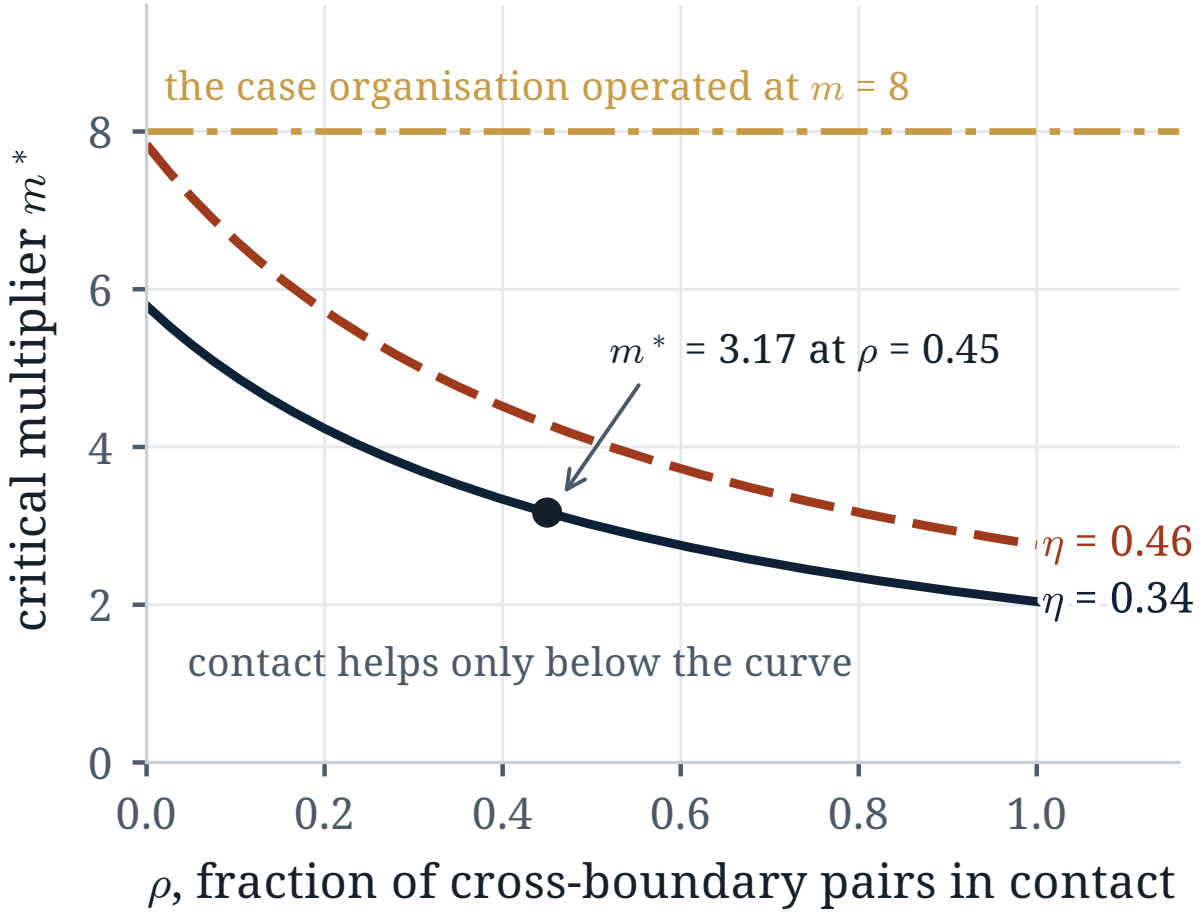


Figure 13.4: **Contact helps only below the curve, and the case organisation was above it before it started.** The critical boundary-cost multiplier `m* = eta (N−1) u_cross / [(1 + (N−1)u(rho)) c_in x 900]` at `N = 6`, `u_cross = 0.55`, `u_div = 0.0990`, `c_in = 0.00012` and 900 cross-boundary pairs. At `rho = 0.45` and `eta = 0.34`, `m*` = 3.17 against a measured `m` = 8. Modelled: linear `u(rho)` is a badged modelling choice and a single scalar boundary multiplier is crude; the qualitative result — that a critical cost exists — survives the crudeness and the exact value does not. Generated by `figs/src/CCP-013_fig04.py`.

against a bound of 30.7225 [Curzi, 2026]. **A boundary is a cost-control device. Removing it by default is not a fix; requirement SL-15 says so.**

Third, and this is the structural point: a strict hierarchy is the organisational form that drives non-parallelism toward zero. A tree gives every unit exactly one parent and forbids overlaps, so any two units at the same level receive their context through the same channel; **a semilattice lets a person belong to a function *and* a product *and* a guild, and the number of ways a person participates is the number of independent contexts they carry** [Alexander, 1965]. **Building a lower valley is, structurally, replacing a tree with a semilattice**, and failure mode **FS6** is the tree in its diagnostic form: cross-domain questions escalate before they are answered, the answer arrives via two managers, and nobody has a working relationship outside their branch.

13.7 7 Design specifications

The requirements below are stated in the form a specification would carry. **`CI-1` to `CI-20` govern deliberation, `CN-1` to `CN-20` seeding, and `SL-1` to `SL-20` the boundary.** A requirement is not a finding: the schedule is a design instrument derived from the results above, and **its numerical thresholds are chosen values rather than measured optima.**

13.7.1 7.1 Deliberation

CI-1 deliberative groups shall be sized `n <= 8` unless a specific argument is recorded. **CI-2** before convening, list the `N` genuinely distinct categories present and strike duplicates; two people from the same lineage count once. **CI-3** the participation index shall be measured on at least one session in four. **CI-4** target `J >= 0.85` for deliberative sessions. **CI-5** every decision meeting shall open with three to five minutes of silent written positions. **CI-6** before any decision is recorded, every member states a position in a go-around. **CI-7** the highest-status participant states their view after all others. **CI-8** every decision-log entry shall carry a dissent field, and an empty field on a contested item is itself a finding. **CI-9** deliberation rooms shall meet the acoustic specification for reverberation and background level [Curzi, 2026e]. **CI-10** seating geometry shall give every participant a direct sightline to every other [Curzi, 2026f]. **CI-11** where `J` is below its floor, run the interruption ledger. **CI-12** the available non-parallelism `u_pot` shall be measured on a stated cadence, **because it decays.** **CI-13** every standing group shall run at least one active maintenance practice for `u_pot`. **CI-14** annually, estimate the marginal cost reduction and the marginal surplus loss from cohesion and record the comparison. **CI-15** no deliberative session shall end without a durable artefact. **CI-16** quarterly, sample decisions and compute the fraction of raised objections surviving into the record. **CI-17 — no compositional screening.** Nothing here licenses selecting individuals on any reported correlate of the collective-intelligence factor. **CI-18 — right of abstention.** Every protocol here shall be declinable, unconditionally and without reason. **CI-19 — metric hygiene.** The participation index shall never be an individual performance target, an appraisal input, or a ranking criterion. **CI-20 — badge discipline.** Any report using this paper's mapping shall state that the mapping from the reported correlates to this framework's parameters is a postulate and not a finding.

13.7.2 7.2 Seeding

CN-1 before seeding, record whether the practice is simple or complex, and state the evidence. **CN-2** seed along the graph of who works with whom, evidenced by shared artefacts and threads, not along the org chart. **CN-3** estimate the clustering coefficient of the seeding region. **CN-4** pick the smallest population within which the practice works standing alone. **CN-5** seed past the local tipping point, taking 25 per cent **of that population** as a floor and not a target. **CN-6 — never spray.** Distributing a fixed budget uniformly across a large population

is prohibited. **CN-7** run the threshold survey before seeding and look for gaps rather than for a mean. **CN-8** manufacture the zero-threshold actor deliberately. **CN-9** programme reporting shall track largest connected adopter cluster, cluster count and between-cluster edges — **not headcount. CN-10 — declare the dip.** Before launch, state in writing that the order proxy is expected to fall first, with a predicted depth and timing. **CN-11** cancellation criteria shall be preregistered. **CN-12 — the decline test.** No seeding programme shall proceed unless a member of the seeded population can decline without consequence. **CN-13** declines shall be recorded with reasons and reported alongside adoptions. **CN-14 — no mandate.** Adoption shall not be required, incentivised by evaluation, or used as a compliance criterion. **CN-15 — sponsor-removal test.** Before expanding, the sponsor withdraws for a month and the practice must survive. **CN-16** the second population shall be seeded by people who already work there. **CN-17** where seeding through a high-degree individual, record the single-point-of-failure risk. **CN-18** do not plan for sustained partial adoption across a population that must share one convention. **CN-19** where an organisation is simultaneously introducing tools that flatten communication, record that they raise the rewiring parameter and degrade the substrate costly practices need. **CN-20 — badge discipline.** Any use of the 25 per cent figure shall carry the statement that it derives from one online experiment on an arbitrary convention with no prior stakes.

13.7.3 7.3 The boundary

SL-1 locate the boundary before diagnosing the behaviour; record what draws it. **SL-2** measure `u_cross` and `u_div` separately and report both. **SL-3** measure `rho`, the fraction of cross-boundary pairs actually in contact, from shared artefacts, meetings and threads. **SL-4** estimate `m` by timing a matched sample of within- and cross-boundary exchanges from question asked to answer usable. **SL-5 — compute `m*` before increasing contact. SL-6 — lower `m` before raising `rho`.** The sequence is shared instrument, then boundary objects, then rotation, then contact targets. **SL-7** build a shared instrument, readable by both sides without specialist tooling and owned by neither. **SL-8** specify boundary objects, not forms: the artefact must tolerate different interpretations on each side while remaining one object. **SL-9** time-box the translator role with a success condition and a build list. **SL-10 — rotation is offered, never assigned**, and declining shall carry no consequence to standing. **SL-11 — rotation dosage is computed, not chosen**, from the `m*` arithmetic. **SL-12** co-locate on shared work, not adjacent work. **SL-13** check the four contact conditions explicitly and record which are absent. **SL-14** prefer semi-lattice to tree. **SL-15** do not remove the boundary by default; it is a cost-control device. **SL-16** nothing in this section is mandated and every instrument shall be declinable. **SL-17 — no exhortation as intervention. SL-18 — predict the reverse rate** for any cross-boundary change, and measure it after attention has moved on. **SL-19** measure whether the cross-domain position is actually cheaper by asking the people who tried it. **SL-20 — badge discipline.** The landscape and escape-rate models are analogies and shall be labelled as such in any report.

13.8 8 Results: three systems worked end to end

All three systems are constructed. No dataset was collected for this paper. Every number below is the output of stated parameters under the models of Sections 4 to 6, carried through so that a reader can re-derive it. **They demonstrate that the procedure runs and that it reports its own failures; they are not evidence that any organisation behaves this way.**

13.8.1 8.1 A nine-person review group: the diagnostic split, and a protocol that cannot win

The system. A hospital discharge safety review group: nine people meeting fortnightly for ninety minutes to review discharge criteria, incidents and proposed protocol changes; four

years old, two membership changes. Audited categories, with two consultants in the same specialty counting once: acute medicine, nursing, pharmacy, physiotherapy and rehabilitation, data and audit, and a patient-experience representative. **`N = 6` and `ln 6 = 1.7918` nats is the hard per-interaction ceiling.** Parameters estimated over eight sessions:

```
n        = 9 people                    N       = 6 distinct categories
u_pot    = 0.55   dispersion of written pre-discussion positions
eta      = 0.35   objection survival into the durable record
delta    = 0.06   per fortnight, from decay of protocol adherence during a
                  10-week gap when the group did not meet
C_t      = 0.42   person-hours in room + preparation + rework attributable to
                  coordination failure, in the same normalised units as Q
Q_0      = 3.10   order proxy: composite of documented current criteria, staff
                  adherence audit score, and incident-review closure rate
J        = 0.52   two acute-medicine consultants hold 61% of speaking time;
                  the patient-experience representative holds 2%
```

Proxy justification. The order proxy is built from *adherence and closure* — downstream operational facts recorded by people outside the group who do not know its interaction measure — while the interaction measure is estimated from written position dispersion inside the group. **The two share no input. A proxy that is secretly the interaction measure will fit beautifully and mean nothing.**

Baseline.

```
Realised  U   = J x u_pot = 0.52 x 0.55 = 0.2860
N_eff         = 1 + 5(0.2860)           = 2.4300
S             = ln(2.4300)              = 0.8879 nats     (46% of ln 6)
eta S         = 0.35 x 0.8879           = 0.3108
delta Q_0     = 0.06 x 3.10             = 0.1860
C             =                           0.4200

Headroom h    = 0.35(1.7918) - 0.1860 - 0.4200 = 0.6271 - 0.6060 = +0.0211
One step      = 0.3108 - 0.1860 - 0.4200                        = -0.2952
```

This is the diagnostic pattern that makes headroom worth computing. Positive headroom says the regime is *feasible*: some interaction geometry grows order. The one-step calculation says the *current* geometry does not. Confusing the two is why organisations despair when they should retune, or retune when they should stop. Nine people convene fortnightly and produce the deliberative output of **2.43 effective participants**, with order falling at 0.295 units per period. Left alone:

period	0	4	8	12	16	17
`Q_t`	3.1000	2.0213	1.1791	0.5215	0.0081	**model breaks: -0.1015**

The steady state is `Q_inf = (0.3108 - 0.4200)/0.06 = -1.8206`, **a negative equilibrium, which the result does not hide: it means the regime is unsustainable or the linear model must be replaced.** At period 17 the admissibility condition fails (`0.4200 > 0.3185`) and the model stops applying. **In the building that is not a mathematical event. It is the group's documented criteria quietly becoming a folder nobody opens** — eight months of fortnightly meetings, the room booked, the minutes taken, **and the institution running a maintenance ritual on a declining asset while recording attendance as evidence of health.**

What surplus would be required. A non-decreasing step requires `S_t >= (delta Q_t + C_t)/eta`:

```
S_req = (0.1860 + 0.4200) / 0.35 = 1.7314 nats        ln N = 1.7918 nats
1 + 5 U = e^{1.7314} = 5.6495   ->   U = 0.9299
```

`S_req < ln N`, so the ceiling does not rule it out — the same fact as positive headroom. **But the group would need realised `U = 0.93` on a roster carrying `u_pot = 0.55`, and even at perfect equality `J = 1` gives `U = 0.55`. The requirement is infeasible on this roster under any conversational protocol whatsoever. The headroom gauge said feasible; the geometry says no**, and that discrepancy is exactly what a single-number dashboard hides.

Stage 1 — turn-taking only. Silent written positions first, a full go-around, consultants speaking last, ceiling absorption installed in the meeting room, and seating changed from a boardroom table to a ring with no head [Curzi, 2026e,f]. Measured over the next eight sessions, `J` rises to **0.86**.

```
U     = 0.86 x 0.55  = 0.4730       N_eff = 3.3650
S     = ln(3.3650)   = 1.2134 nats  (up from 0.8879, +36.7%)
eta S = 0.35 x 1.2134 = 0.4247
Step  = 0.4247 - 0.1860 - 0.4200 = -0.1813      STILL DECLINING
```

The largest available conversational intervention improved the step by 0.1139 units per period and did not change its sign. Effective participants rose from 2.43 to 3.37 — a real gain, honestly measured — **and the asset still shrinks, just more slowly**, from 3.1000 to 0.7627 over 24 periods, with `Q_inf = (0.4247 - 0.4200)/0.06 = 0.0783`. **The intervention bought survival at near-nothing rather than collapse.**

This is the result the section has been building toward and it is the reason it is not a list of meeting tips. A turn-taking intervention operates on `S`; `S` enters the storage law multiplied by `eta` and competes against `delta Q + C`. If the cost term is large, no achievable `S` wins, and running a better meeting is a real improvement inside a losing budget. Anyone stopping at Stage 1 would hold a genuine, measurable, publishable improvement in group process **and a department whose discharge criteria continued to decay** — and would also, quite reasonably, conclude that the collective-intelligence literature does not survive contact with practice.

Stage 2 — the cost term. The cost of 0.42 decomposes as 0.13 for the session itself, 0.06 for preparation, **0.15 for a standing weekly sub-meeting duplicating the fortnightly agenda**, and 0.08 for criteria re-litigated because prior decisions were not findable. Two moves, neither about conversation. **Retire the weekly sub-meeting** — it existed because the fortnightly meeting produced no findable record, so people met to re-establish what had been decided, **which is low conversion presenting as a cost** — saving 0.15. **Install the decision log** with an open-questions column and a required dissent field, saving 0.07 of the 0.08.

```
C: 0.4200 -> 0.2000        eta S unchanged at 0.4247 (Stage 1 protocols retained)

Headroom h = 0.6271 - 0.1860 - 0.2000 = +0.2411
Step       = 0.4247 - 0.1860 - 0.2000 = +0.0387      GROWING
S_req = 0.3860/0.35 = 1.1029  ->  U_req = 0.4026 ;  realised U = 0.4730   FEASIBLE
Q_inf = (0.4247 - 0.2000)/0.06 = 3.7450
```

period	0	4	8	12	16	20	24
Q_t	3.1000	3.2414	3.3518	3.4380	3.5053	3.5579	3.5989
h_t	+0.2411	+0.2326	+0.2260	+0.2208	+0.2168	+0.2136	+0.2112

Contrast the three trajectories: doing nothing, the model breaks at period 17 with the stock through zero; with the best available meeting protocol alone, 0.7627 at period 24 and falling toward 0.0783; with the protocol and the cost cut, 3.5989 and rising toward 3.7450. And observe the convergence: the stock climbs and slows as it goes, because headroom shrinks as

the stock grows. **Success consumes its own headroom** [Curzi, 2026b], and the fixed point is reached because `delta > 0` makes the map a contraction.

What the group should watch next. *The cohesion clock:* `u_pot = 0.55` was measured, not assumed, and it decays. At 5 per cent per fortnight it halves in about a year, and at `u_pot = 0.40` the step at `Q = 3.60` becomes `0.3502 − 0.2160 − 0.2000 = −0.0658` — **negative again, with nothing visibly changed and everyone getting on better than ever.** *The maintenance trap:* as the stock rises toward 3.745, the maintenance term rises from 0.1860 to 0.2247 and the budget tightens by 0.0387 **purely from success.** *The proxy:* the moment any component of the order proxy becomes a target it detaches from the thing it names [Campbell, 1979], **and the adherence audit is the most gameable.**

Preregistered falsification, declared before Stage 2. Fit `Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + e_t` with `0 <= a <= 1` and `b >= 0` on periods `−16..0`, freeze, and predict periods `1..24` held out, against four baselines: persistence; direct input with raw `U` and no logarithm; a matched autoregressive model with no interaction term; and a **shuffled-input control** in which the same surplus series is randomly permuted across periods. **Falsified if** the shuffled control predicts held-out order as well as the ordered series, or the matched autoregressive baseline is not beaten, or the fitted persistence lies outside `[0,1]`, or the fitted gain is negative. **The shuffled control is the one most often omitted and the one that decides this case: a department that improved its meetings, cut two commitments, installed a decision log and got better may simply be a department that received attention. The shuffle asks whether the timing of surplus predicts the timing of order, and if it does not, the arithmetic described a coincidence in the vocabulary of a theorem.**

13.8.2 8.2 A 220-person firm: seeding a costly practice, and the dip

The system. A professional services firm, twelve teams, chronically losing decisions: the same questions re-litigated across engagements, nobody able to find why a prior call was made, senior staff spending a measurable fraction of the week reconstructing context. **The proposed practice is a decision log** — for every decision above a stated threshold, a short structured record of the decision, its alternatives, the dissent and the open questions. **It is a conversion device and a classic complex contagion: it is only worth writing if others read it, and it is socially costly to be the only person writing them.**

```
n = 220 people, 12 teams        N = 8 distinct practice categories, ln N = 2.0794
u = 0.50 from a disagreement audit -> S = ln(1 + 7 x 0.50) = ln 4.50 = 1.5041 nats
delta = 0.05 per month          C_0 = 0.55        eta_0 = 0.30        Q_0 = 4.00
```

with the order proxy built from reconstructibility, re-litigation rate and lookup latency — **all sampled from engagement files and sharing no input with the disagreement audit.** A two-team pre-pilot over four months estimates post-adoption values `eta_1 = 0.52` and `C_1 = 0.46`, **the latter net of the real cost of writing the logs.**

```
h_0  = 0.30(2.0794) − 0.2000 − 0.5500 = 0.6238 − 0.7500 = −0.1262      h < 0
step = 0.4512 − 0.2000 − 0.5500                          = −0.2988     DECLINING
```

Headroom is negative, and that is categorical: no choice of interaction geometry prevents order from decreasing. No amount of better meetings, better recruitment or higher non-parallelism fixes this firm, **because even at the ceiling the input `0.30 x 2.0794 = 0.6238` is less than `delta Q + C = 0.7500`. The binding constraint is conversion and cost, and the diagnosis is available before anyone is consulted about anything.** Post-adoption, if the practice fully lands:

```
h_1    = 0.52(2.0794) − 0.2000 − 0.4600 = 1.0813 − 0.6600 = +0.4213
step_1 = 0.52(1.5041) − 0.6600                            = +0.1221   GROWING
Q_inf  = (0.7821 − 0.4600)/0.05                           = 6.4424
```

So the practice is worth having. The remaining question is entirely a contagion question: does it arrive?

The seeding decision. Budget: 25 people's time. *Plan A, spray:* two seeds per team, 22.8 per cent of non-seeds in reinforcing range, 44.5 people, **zero teams past their local tipping point, largest adopter cluster 2.** *Plan B, cluster:* saturate one 30-person team, local committed fraction 83 per cent, **largest adopter cluster 25**, and the practice's value realised locally because that team's own decisions become reconstructible regardless of the other eleven teams. **Plan B is chosen, and what is given up is stated in advance: for roughly two months 195 people see no change at all and the programme's reportable reach is 14 per cent. Requirement CN-10 exists because a reach of 14 per cent at month two will otherwise be read as failure.**

The adoption trajectory, modelled logistically at `r = 0.55` per month with `K = 220` and `a_0 = 25/220`, coupled by the linear conversion and cost forms of Section 5.8 — `eta(a) = 0.30 + 0.22a` and `C(a) = 0.55 − 0.09a`:

month	adopters	share `a`	`eta`	`C`	`Q_t`	`h_t`	step
0	25.0	0.1136	0.3250	0.5398	4.0000	-0.0640	-0.2509
1	37.2	0.1690	0.3372	0.5348	3.7491	-0.0211	-0.2151
2	54.2	0.2463	0.3542	0.5278	3.5340	**+0.0320**	-0.1718
3	76.6	0.3484	0.3766	0.5186	3.3622	+0.0965	-0.1203
4	104.1	0.4732	0.4041	0.5074	3.2419	+0.1708	-0.0617
5	134.3	0.6103	0.4343	0.4951	3.1802	+0.2490	-0.0009
6	163.1	0.7412	0.4631	0.4833	**3.1793**	+0.3206	+0.0542
8	202.0	0.9181	0.5020	0.4674	3.3294	+0.4100	+0.1212
10	215.8	0.9808	0.5158	0.4617	3.5831	+0.4317	+0.1349
12	219.1	0.9960	0.5191	0.4604	3.8505	+0.4266	+0.1279
14	219.8	0.9992	0.5198	0.4601	4.1009	+0.4158	+0.1167

Read the stock column. It falls for six months, bottoms at 3.1793 — 20.5 per cent below where it started — and does not recover to its starting value until month 14. That is the adoption dip, computed rather than asserted. **Month five is the maximum danger point:** the step is `−0.0009`, essentially flat; the stock is at its worst; 61 per cent of the firm has adopted; and every visible operational indicator is at its lowest reading since before launch. **A cancellation at month five would be entirely defensible on the evidence available at month five and would destroy a programme one month from turning.**

Note the headroom column, which tells a different and earlier story. Headroom crosses zero at month two, four months before the stock turns. The gauge is the leading indicator and the order proxy is the lagging one: a firm watching headroom sees the regime become sustainable in month two, and a firm watching the stock sees nothing good until month six. **That is the strongest practical argument in this corpus for computing the gauge.**

A coincidence, labelled as one. At month 2 the adopted share is 0.2463, essentially the published tipping-point figure, and month 2 is also where headroom turns positive. **This is a coincidence.** The logistic parameters came from the pre-pilot, and headroom crossing zero at `a = 0.246` reflects the conversion and cost slopes, **not any tipping-point mechanism. Two arithmetic facts landing on the same month is not evidence that one caused the other, and reporting it as confirmation of the 25 per cent threshold would be exactly the misuse this paper's honesty clause exists to prevent.** The genuine tipping-point calculation is the seeding decision: 25 committed people inside a 30-person team is 83 per cent, far above the floor, **which is why that team's practice held after the sponsor withdrew — tested at month 3 with a one-month withdrawal, and logging continued at 92 per cent of its prior rate.**

Preregistered falsification. Fit on months `−12..0`, freeze, predict `1..18` held out, with the same four baselines including the shuffled-input control. **Preregistered predictions:** the stock falls for five to seven months with a trough between 3.05 and 3.30 and recovery to its starting value by month 12 to 16; headroom crosses zero between months 1 and 3; the largest connected adopter cluster exceeds 60 by month 4; and the seeded team sustains at least 85 per cent of its logging rate through a one-month sponsor withdrawal. **Falsified if** the shuffled control matches the ordered series, the autoregressive baseline is not beaten, the fitted parameters fall outside their declared ranges, or the trough falls outside `[2.8, 3.5]` or outside months 4 to 9.

The first prediction is what makes the preregistration worth the paper. Predicting that a programme will work is cheap — every programme predicts that. Predicting the depth and timing of the hole it digs first, before digging it, is a claim that can embarrass you in a specific month, and that is what makes it evidence rather than advocacy [Nosek et al., 2018, Platt, 1964].

13.8.3 8.3 A 60-person organisation: three futures, and the one that is worse than inaction

The system. A product organisation in two long-standing silos of thirty: a platform group (infrastructure, tooling, reliability) and an application group (customer-facing features, domain workflows). **Both are competent. Both are frustrated. Delivery has been slowing for two years.**

```
N = 6 distinct knowledge categories,  ln N = 1.7918 nats
u_cross = 0.55   u_div = 0.0990   u available = 0.6490
delta = 0.05 per month            Q_0 = 4.40
c_in = 0.00012   within-boundary pairs 870   cross-boundary pairs 900
```

with the order proxy assembled from interface-contract currency, answer latency and rework, **all recorded from artefact and ticket histories and sharing no input with the disagreement audit that produces the interaction measure.** *Measured boundary cost:* twenty matched exchanges, ten within and ten across, timed from question asked to answer usable — median within 1.8 hours, median across 14.6 hours, so **`m` is about 8**. *Measured contact:* cross-boundary pairs with any shared artefact, meeting or working thread in the last quarter, **27 of 900, so `rho = 0.03`**. *Measured conversion:* fraction of cross-boundary decisions traceable into a durable interface contract, **`eta = 0.34`**.

Case A — the status quo.

```
u realised = 0.03(0.55) + 0.0990 = 0.1155     N_eff = 1.5775    S = 0.4558 nats
eta S = 0.1550     C = 0.00012(870) + 8(0.00012)(27) = 0.1044 + 0.0259 = 0.1303
delta Q = 0.2200
h    = 0.34(1.7918) − 0.2200 − 0.1303 = +0.2589
step = 0.1550 − 0.2200 − 0.1303       = −0.1953        DECLINING
Q_inf = (0.1550 − 0.1303)/0.05 = 0.494
```

Positive headroom with a negative step — the same diagnostic split as Section 8.1. The required surplus is `(0.2200 + 0.1303)/0.34 = 1.0304` nats, well inside the ceiling, needing `u_req = 0.3604`. **And here is the finding that makes the whole diagnosis: the roster carries 0.6490, comfortably above the required 0.3604. The organisation already owns enough difference to sustain itself. It is realising 0.1155 of 0.6490 — 17.8 per cent — because 873 of its 900 cross-boundary pairs never speak. Nobody needs to be hired. Nobody needs to be persuaded of anything. The difference is on the payroll and out of contact.** Left alone the stock falls from 4.4000 to 1.6340 over 24 months: **two years of competent work by sixty competent people, and 63 per cent of stored order gone.**

Case B — forced integration, which is the intervention that gets chosen. A reorganisation: cross-functional squads, mandatory joint planning, shared standups. Contact rises from 0.03 to 0.45 — **a genuine, large, successfully executed increase. The boundary cost is not touched**, because no shared instrument was built and no boundary object was specified: the two domains still measure different things and still mean different things by the same words.

```
u realised = 0.45(0.55) + 0.0990 = 0.3465     N_eff = 2.7325   S = 1.0052 nats   (up 121%)
eta S = 0.3418                                                                  (up 121%)
C = 0.1044 + 8(0.00012)(405) = 0.1044 + 0.3888 = 0.4932                          (up 279%)
h    = 0.6092 − 0.2200 − 0.4932 = −0.1040        h < 0
step = 0.3418 − 0.2200 − 0.4932 = −0.3714        DECLINING FASTER THAN BEFORE
```

Surplus rose by 121 per cent and the system got worse. The reorganisation did exactly what it set out to do — it more than doubled realised non-parallelism — **and it was destructive anyway, because cost rose 279 per cent against a surplus that rises logarithmically.** And headroom has crossed zero, **so no choice of interaction geometry can now prevent decline**: there is no seating plan, no meeting protocol, no facilitation and no amount of goodwill that recovers the system, and the only available moves are to cut cost, raise conversion, or shed stock. **Check against the critical multiplier: at `rho = 0.45` and `eta = 0.34`, `m* = 3.17`, and the organisation operates at `m = 8`. It was above the critical multiplier before the reorganisation began, and the reorganisation's entire mechanism was to increase the quantity that `m` multiplies.** The stock falls from 4.4000 to 0.9856 in twelve months: **78 per cent gone in twelve months against 63 per cent in twenty-four by doing nothing. The intervention is worse than inaction by a wide margin, and everyone involved is working harder.**

Case C — build the lower valley. The same target contact, approached in the opposite order. Three things are built over four months, **before any contact target is set.** *(1) A shared instrument:* one service-level dashboard, one definition of latency, error and saturation, one dataset, owned by neither group and readable by both without specialist tooling — **most cross-boundary arguments in the sampled history were arguments about what was happening, and those now terminate at a shared screen.** *(2) A boundary object:* a written interface contract per integration point, **plastic enough that platform reads it as an operational commitment and application reads it as a behavioural guarantee, without either translating into the other's frame** [Star and Griesemer, 1989]. *(3) Two rotations, offered not assigned:* four people, one quarter each, working — not observing — inside the other domain. **Six people declined and the declines were recorded with no consequence.**

Measured after four months, `m` falls from 8 to **2** (median cross-boundary answer latency 14.6 to 3.6 hours) and `eta` rises from 0.34 to **0.46**, because the interface contract is a durable artefact and cross-boundary decisions now land in it.

```
u realised = 0.3465   S = 1.0052 nats            (identical to Case B)
eta S = 0.46 x 1.0052 = 0.4624                    (Case B: 0.3418)
C = 0.1044 + 2(0.00012)(405) = 0.2016             (Case B: 0.4932)
h    = 0.46(1.7918) - 0.2200 - 0.2016 = +0.4026
step = 0.4624 - 0.2200 - 0.2016       = +0.0408   GROWING
Q_inf = (0.4624 - 0.2016)/0.05 = 5.2160
S_req = 0.4216/0.46 = 0.9165 -> u_req = 0.3001 ;  realised 0.3465   FEASIBLE, with margin
```

Case B and Case C have identical contact, identical realised non-parallelism and identical surplus. The only differences are the boundary cost and the conversion coefficient. Case B declines at `-0.3714` per month; Case C grows at `+0.0408`. The entire difference between a destructive reorganisation and a working one is the boundary cost and the conversion rate — neither of which is about how much people collaborate.

The dosage calculation. `rho = 0.45` was not chosen; it was computed. At `m = 2` and `eta = 0.46`, sweeping contact:

rho	0.00	0.05	0.10	0.20	0.30	0.45	0.60	0.80	1.00
C	0.1044	0.1152	0.1260	0.1476	0.1692	0.2016	0.2340	0.2772	0.3204
S	0.4021	0.4901	0.5710	0.7154	0.8416	1.0052	1.1458	1.3070	1.4457
step at Q = 4.40	-0.1394	-0.1097	-0.0833	-0.0385	-0.0021	**+0.0408**	+0.0731	+0.1040	+0.1246

The step turns positive just above `rho = 0.30`, and 0.45 is chosen as the first target because it clears the crossing with margin. **Note the marginal gain from 0.45 to 1.00 is 0.0838 for more than double the contact — surplus concavity showing up as a diminishing return on integration.** The same sweep at the un-remediated `m = 8` and `eta = 0.34`:

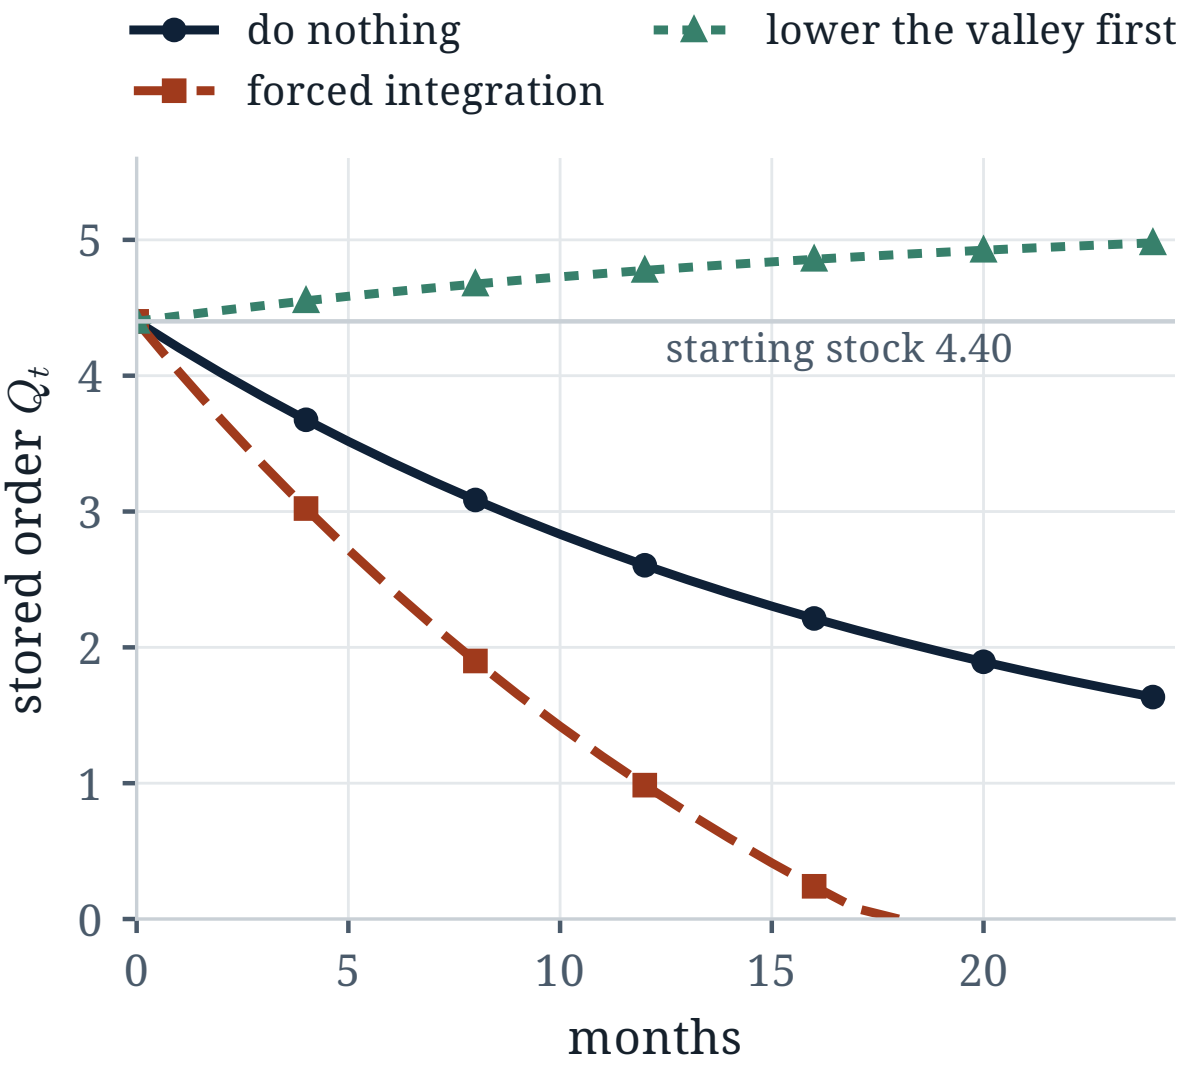


Figure 13.5: **Identical contact, identical surplus, opposite outcomes.** Stored order under the storage law for one 60-person organisation: doing nothing declines at `−0.1953` per month; forced integration to `rho` = 0.45 with the boundary cost untouched raises surplus 121 per cent, raises cost 279 per cent, and declines at `−0.3714`; lowering the boundary cost first — `m` from 8 to 2, `eta` from 0.34 to 0.46 — buys the identical contact and grows at `+0.0408`. Modelled: every parameter is surveyed or stipulated for one constructed organisation and none carries an interval. Generated by `figs/src/CCP-013_fig05.py`.

rho	0.00	0.10	0.20	0.30	0.45	0.60	0.80	1.00
C	0.1044	0.1908	0.2772	0.3636	0.4932	0.6228	0.7956	0.9684
step at Q = 4.40	-0.1877	-0.2167	-0.2540	-0.2975	-0.3714	-0.4532	-0.5712	-0.6968

Monotonically worse at every level of contact, all the way to complete integration. There is no value of `rho` at which the un-remediated organisation benefits from more collaboration. That is what `m > m*` means, and it is why requirement SL-5 forbids raising contact before measuring it.

Rotation dosage in weeks, at `m = 2`, `eta = 0.46` and `Q = 4.40`:

weeks per person per year across the boundary	0	1	2	4	8	13	26
rho	0.0000	0.0192	0.0385	0.0769	0.1538	0.2500	0.5000
step	-0.1394	-0.1276	-0.1163	-0.0952	-0.0580	-0.0194	**+0.0526**

A one- or two-week rotation is tourism: the step improves by 0.012 to 0.023 and remains firmly negative. Even a full quarter — thirteen weeks, which most organisations would consider a serious commitment — leaves the step at -0.0194. Rotation alone, at any dose an organisation will tolerate, does not save this system. It is a supporting instrument; **the shared instrument and the interface contract are what move the boundary cost and the conversion rate, and they are what make the modest contact affordable. That is the kind of conclusion that only appears once the arithmetic is done.**

Preregistered falsification, with the same model, baselines and shuffled-input control. **Four predictions:** median cross-boundary answer latency falls below 5.0 hours within four months of the shared instrument going live; headroom stays positive throughout and the stock rises monotonically from month 1; the stock at month 24 lies in `[4.7, 5.2]` with a steady-state estimate in `[4.9, 5.5]`; and **the rotation declines do not correlate with subsequent evaluation outcomes — an ethics check, not a performance one, and the single most important line in the preregistration. Falsified if** the shuffled control matches, the autoregressive baseline is not beaten, the fitted parameters leave their ranges, **or the boundary cost does not fall below 3 despite a well-built, well-used shared instrument in daily use.**

That last falsifier is the one worth stating loudest. The entire section rests on the claim that the boundary cost is dominated by translation and disputed measurement and that a shared instrument removes most of it. **If it stays at 8 while the instrument sits on the wall in daily use, then the boundary cost lives somewhere this paper did not look — in incentives, in status, in scheduling — and the recommendation to build instruments first is wrong for that system. A framework that cannot fail cannot be trusted, and this is exactly where this one fails.**

And the final observation, which is the reason for preferring this method over every alternative: in Case C six people declined the rotation and nothing happened to them. The dashboard was built and they could ignore it; the interface contracts existed and they could read them or not. **The system improved from 4.4000 to 4.9777 with six recorded refusals in it. A method that works by removing resistance produces a system whose improvement is compatible with people saying no — and that is not a concession the method makes. It is the property that makes it safe to recommend.**

13.9 9 Discussion

13.9.1 9.1 What follows

Realised difference is not roster difference, and the gap is engineerable. All three sections attack the same quantity from different sides: turn-taking decides how much of a roster's difference is emitted, contagion decides whether a practice that would raise conversion ever

establishes, and the boundary decides whether difference that exists in a building is ever in contact. **In every worked case the roster already carried enough difference and the realisation was the constraint** — 17.8 per cent realised in the 60-person organisation, 46 per cent of the ceiling in the review group. **A survey that measures only what is on the roster will find no problem in an organisation that is dissolving.**

The diagnostic split is the single most transferable instrument here. Positive headroom with a negative step means the regime is feasible and the current geometry does not achieve it; the two are routinely confused, and the confusion produces despair where retuning is called for and retuning where stopping is called for. **Inverting the requirement into the non-parallelism it demands is what turns the split into a decision**, and in the review group it showed the demand was infeasible on that roster at any protocol whatever.

The critical multiplier converts a values argument into a measurement. Whether cross-boundary contact helps is decided by a quantity that can be measured in an afternoon with twenty timed exchanges, **and enthusiasm does not appear in the expression.** The practical consequence is a sequence: **lower the boundary cost, then raise contact, and compute the dosage rather than choosing it.**

Costly practices need reinforcement, so reach is the wrong report. The 57-against-19 gap and the percolation transition together say that a programme should report largest connected adopter cluster and never total adopters, **and that the point of maximum apparent failure is frequently the point just before the transition.**

13.9.2 9.2 What does not follow

Nothing here shows that any of these interventions raises a group's collective intelligence. The causal step from turn-taking equality to that factor requires a direction the source literature does not establish, and the preregistered test of it was run elsewhere in this series and did not confirm [Curzi, 2026e]. **Every claim here about turn-taking is a claim about a measured mediator.**

Nothing here identifies a social effect. The specifications are reflexive and the endogenous component is not identified from observational data [Manski, 1993, Shalizi and Thomas, 2011], and the contagion sections are where the confound between influence and selection bites hardest [Lyons, 2011]. **The shuffled-input control separates a mechanism from a period of attention; it does not solve identification.**

The network numbers are structure, not prediction. Ring lattices and random graphs are analytical conveniences; real organisational graphs are degree-heterogeneous and correlated [Barabási, 2016, Ugander et al., 2012], and both the clustering figures and the percolation threshold move under realistic degree distributions. **A programme that used the `z = 1.1` figure as a target rather than as a shape has misused it.**

The tipping-point figure is one experiment. Online, anonymous, on an arbitrary naming convention with no prior stakes, no status differences, no history and no switching cost [Centola et al., 2018]. **Use it to check whether a plan is off by a factor of ten.**

And the deference and landscape models prove nothing about people. The alignment form and the comfort landscape are badged analogies; **what they license is a hypothesis about the shape of a decision, and the forward-and-reverse-rate prediction of Section 6.1 is what would test it.**

13.9.3 9.3 Relation to the surrounding literature

The collective-intelligence result supplies the correlate [Woolley et al., 2010] and its replication record supplies the caution [Engel et al., 2014, Bear and Woolley, 2011]; the diversity-and-performance literature and its critics supply the standing warning [Mannix and Neale, 2005, Thompson, 2014, van Knippenberg and Schippers, 2007, Hong and Page, 2004, Page, 2007]. The cascade literature supplies the mechanism the silent-write protocol defeats [Bikhchandani

et al., 1992, Banerjee, 1992], and the aggregation experiments supply the reason independence is a precondition rather than a nicety [Galton, 1907, Surowiecki, 2004, Lorenz et al., 2011].

The contagion sections rest on the simple-versus-complex distinction [Centola and Macy, 2007, Centola, 2010], the threshold model [Granovetter, 1978, Schelling, 1971, 1978], the small-world construction [Watts and Strogatz, 1998, Watts, 2002], percolation on random graphs [Erdős and Rényi, 1960, Newman, 2018, Stauffer and Aharony, 1994], and the experimental tipping-point demonstration [Centola et al., 2018], with the resilience literature supplying the reason a connected component is the right object [Albert et al., 2000, Callaway et al., 2000, Cohen et al., 2000]. **What is new is not the network theory but its coupling: an adoption share is converted into a conversion coefficient and a cost, and the resulting trajectory is priced against a ceiling. That step exists in neither literature, because it is not a network question.**

The boundary section rests on the structural-mirroring observation [Conway, 1968], the quadratic communication count [Brooks, 1975], the contact hypothesis and its qualifications [Allport, 1954, Pettigrew and Tropp, 2006, Putnam, 2007, Abascal and Baldassarri, 2015], boundary objects [Star and Griesemer, 1989, Star and Ruhleder, 1996], and the semilattice argument [Alexander, 1965]. **What is new is the critical multiplier**: the contact literature states conditions and this paper states a threshold, **and the threshold is computable from twenty timed exchanges.**

Where the literature disagrees with the framework it is cited and the disagreement is stated. Unbounded growth in `ln N` is contradicted by the inverted-U finding on cohesion [Uzzi and Spiro, 2005], by settings in which redundancy reduces safety [Sagan, 1993, Perrow, 1984], and by the ecological result that increasing diversity and connectance destabilises randomly assembled systems [May, 1972, Gardner and Ashby, 1970]. **And the framework's own insistence that difference and not headcount is what pays was stated formally in 1952** [Markowitz, 1952]: **it is not novel here, and presenting it as novel would be the easiest available overclaim.** The functional form remains underdetermined — a sub-linear power law [Latané, 1981] and `n log n` growth [Briscoe et al., 2006] fit the same qualitative facts [Stumpf and Porter, 2012] — and the logarithm's strongest warrant requires no psychological premise [Peters and Gell-Mann, 2016].

13.10 10 Verification

This section is not a summary. It exists so that a reader can *check* the paper without re-reading it, by verifying several simple things and then seeing what they compound into.

13.10.1 10.1 Dimensional audit

Every symbol, its dimension, and whether it is measured, derived, defined, or dimensionless by construction. Where the source is silent on a dimension, this table says so rather than quietly assigning one.

symbol	quantity	dimension	unit	status
Q_t	stored order	$[Q]$	nats of retained difference	state variable, no operational unit
S_t	surplus per period	$[Q\,T^{-1}]$	nats per period	derived from U_t
C_t	cost per period	$[Q\,T^{-1}]$	nats per period	measured as a rate
δ	decay per period	$[1]$	dimensionless	$0 < \delta < 1$

symbol	quantity	dimension	unit	status
η	conversion	$[1]$	dimensionless	**the source does not supply this dimension**; the law balances only if η is a pure number
$u,\ U_t,\ J_t$	non-parallelism, realised, equality	$[1]$	all in $[0,1]$	ratios by construction
N	roster size	$[1]$	a count	measured
m	boundary-cost multiplier	$[1]$	dimensionless	ratio of two latencies
ρ	cross-boundary contact	$[1]$	dimensionless	a fraction

Comfort, stored order and surplus have no operational unit in this framework. They are counted in nats because the functional is logarithmic, not because a nat has been calibrated against anything a group can be observed to do. A dimensionless placeholder is the honest statement; an invented unit would not be.

13.10.2 10.2 The postulates

Each is atomic, and each is checkable on its own by exactly one of: arithmetic a reader can redo, a definition this paper has already fixed, or a single external citation.

- **P1.** $u(x,y) = 1 - (x \cdot y)^2$ lies in $[0,1]$, is 0 at parallel and 1 at orthogonal. *DEFINITION.*
- **P2.** $F(u) = \ln(1 + (N-1)u)$, so $0 \le F \le \ln N$. *DEFINITION* [Curzi, 2026].
- **P3.** The storage law is the first-order linear recurrence $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$. *DEFINITION* [Curzi, 2026a].
- **P4.** For $0 < \delta < 1$, $|1-\delta| < 1$. *ARITHMETIC.*
- **P5.** A geometric series with ratio $|r| < 1$ converges to $1/(1-r)$. *CITED*, standard.
- **P6.** $U_t = J_t\, u_{\text{pot}}$ — realised non-parallelism is the roster's available difference scaled by turn-taking equality. **POSTULATED.** The linear form is a modelling choice with no derivation behind it; it was chosen for its boundary behaviour and should be replaced the moment anyone measures a better one.
- **P7.** The three reported correlates of a collective-intelligence factor map onto this framework's parameters. **POSTULATED** — a modelling assignment, not a derivation, and nothing here tests it.

13.10.3 10.3 The compounding chain

Each line is checkable without re-reading the line before it for anything but its stated output.

```
P1 + P2        -> L1   F is bounded by ln N and is concave in u        ARITHMETIC
P3 + P4        -> L2   the recurrence contracts; a unique fixed point  ARITHMETIC
L2 + P5        -> L3   Q_inf = (eta F(u*) - c*)/delta                  ARITHMETIC
L3             -> R1   Q_inf may be NEGATIVE, and the law does not     ARITHMETIC
                       hide it when it is
P3             -> R2   h_t = eta ln N - delta Q_t - C_t is the         DEFINITION
                       maximum one-step increment
P2 + P6        -> R3   S = ln(1 + (N-1) J u_pot)                       POSTULATED (via P6)
R3             -> R4   at N = 9, u_pot = 0.3437: J 0.52 -> 0.86        ARITHMETIC
                       takes S from 0.8879 to 1.2134 nats
```

```
P7            -> R5   turn-taking is an instrument on U, sensitivity  POSTULATED
                      an instrument on eta
```

Two of the eight steps are POSTULATED, and both carry the paper's engineering content. R4 is arithmetic *given* P6, and P6 has no derivation. A reader who rejects P6 keeps R1, R2 and the whole of the storage-law apparatus, and loses the turn-taking channel entirely. That is the honest accounting, and it is stated here rather than left to be discovered.

13.10.4 10.4 Checks a reader can run in a minute

- **The two J values imply one roster.** From $S = \ln(1 + 8J\,u_{\text{pot}})$, $J = 0.52$ with $S = 0.8879$ implies $u_{\text{pot}} = 0.343755$; $J = 0.86$ with $S = 1.2134$ implies 0.343736. They agree to five decimal places, so the two figures describe the same roster and neither was fitted independently.
- **The affine count is not a Hill number.** At $N = 100$ and $u = 0.05$, $N_{\text{eff}} = 1 + 99(0.05) = 5.95$. For n equal shares $1/\text{HHI}$ returns n exactly; $1+(N-1)u$ does not, because it never consults the shares.
- **The bound at the extremes.** $F(0) = \ln 1 = 0$ and $F(1) = \ln N$. A pair in perfect opposition and a pair in perfect agreement both read $u = 0$.

13.11 11 Limitations

1. $U_t = J_t u_{\text{pot}}$ **is a modelling choice, not a result.** It has no derivation, and its linearity is chosen for boundary behaviour alone. H1 stands or falls with it.
2. **The collective-intelligence mapping is an assignment.** That turn-taking maps to realised non-parallelism, social sensitivity to conversion, and composition to the category count is the engineering content of Section 4 and it is **not proved**. Nothing in this paper tests it.
3. **The basin-and-barrier description is an analogy**, not a measurement. Section 6.1 states what it licenses; it licenses no transfer of physical authority to an organisational claim.
4. **The reflexive specification is not identified from observation.** Where a group's mean influences its members and the members constitute the mean, the endogenous effect cannot be separated from contextual effects or correlated unobservables [Manski, 1993, Shalizi and Thomas, 2011]. No remedy is adopted here, so **the affected claims are modelling positions, not identified empirical results**; Brock and Durlauf [2001] is the realistic route not taken.
5. **Quantitative structure is assumed, not tested.** Treating these quantities as admitting ratio arithmetic presupposes structure that Michell [1997] argues psychology habitually assumes without testing. The conjoint-measurement conditions that would license it [Luce and Tukey, 1964] **were not tested here.**
6. **The framework reduces jointly varying phenomena to single numbers.** This is one limitation of the framework in several domains, not several independent problems; it is stated once, at framework level, in Curzi [2026], and this paper inherits it rather than restating it.
7. **The functional form is underdetermined.** A sub-linear power law [Latané, 1981] and $n \log n$ growth [Briscoe et al., 2006] fit the same qualitative facts [Stumpf and Porter, 2012]. Nothing here discriminates between them.
8. **Unbounded growth in $\ln N$ is contradicted** by the inverted-U finding [Uzzi and Spiro, 2005], by settings where redundancy reduces safety [Sagan, 1993, Perrow, 1984], and by ecological results on connectance [May, 1972]. **This is a domain restriction on the model, not a finding of it.**
9. **No first-party measurement.** Every number in Section 8 is a constructed worked case. No dataset was collected, no group was observed, and no prediction here has been tested against an outcome.

10. **No information criterion is computed anywhere in this paper**, so the parameter count of any fitted form is not paid for. The model-comparison requirement is recorded as **NOT MET**.
11. **H2 and H5 are stated without measurement.** H2 assigns a correlate to a parameter and measures neither; H5's latency claim is a worked figure, not an observation.
12. **The three effective counts must not be substituted for one another**, and $\ln E$ must never be reported as a surplus. The affine count belongs to no order of the Hill family [Hill, 1973, Jost, 2006].

13.12 12 Conclusion

Three mechanisms that move realised non-parallelism are made numerical and coupled to the storage law. Stated in the form other papers cite:

- **R1.** The steady state $Q_\infty = (\eta F(u^*) - c^*)/\delta$ **may be negative**, and the law neither forbids nor hides it. A system can be run below its own replacement rate.
- **R2.** Headroom $h_t = \eta \ln N - \delta Q_t - C_t$ is the maximum one-step increment, and the sustainability ceiling is $C_t \leq \eta \ln N - \delta Q_t$.
- **R3.** Under the turn-taking postulate, $S = \ln(1 + (N-1)J\,u_{\text{pot}})$: participation equality is an instrument on realised difference, and the roster bounds it.
- **R4.** In the nine-person worked group at $u_{\text{pot}} = 0.3437$, raising J from 0.52 to 0.86 takes S from **0.8879 to 1.2134 nats** — a consequence of the model, not an observation.
- **R5.** Costly practices require reinforced exposure: under a 2-of-8 rule a clustered seeding budget dominates a sprayed one **even though the sprayed plan reports 44.5 people in reinforcing range against the clustered plan's five.** Reach is the wrong instrument for a costly practice.
- **R6.** Whether cross-boundary contact helps is set by the boundary cost, not by disposition: contact improves the one-step increment only while $m < m^*$. The worked case gives $m^* = 3.17$ at $\rho = 0.45$, $\eta = 0.34$, against a measured $m = 8$. **This is the paper's most exposed prediction and the one it most wants tested.**
- **CR (counter-result).** Two of the eight steps in the verification chain are postulated, both load-bearing; the collective-intelligence mapping is an assignment with no test; and the framework's own central insight — that difference and not headcount is what pays — **was stated formally in 1952** [Markowitz, 1952] and is not novel here.

13.13 Declarations

Data and code availability. No external dataset was collected for this paper. Every worked case in Section 8 is a constructed example whose inputs are stated in the text, and every figure is generated by the script named in its caption. **No group was observed and no prediction reported here has been tested against an outcome.**

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Single author: Michael Laurence Curzi conceived, wrote and verified the whole of this paper.

Ethics. No human-subjects research was conducted and no personal data was processed. The worked groups are constructions, not people.

Preregistration. No test reported in this paper was preregistered, because no test was run. The hypotheses in Section 2 are stated with falsifiers so that they *can* be preregistered by someone with access to a suitable setting; **none of them has been.**

Chapter 14

Order Without a Centre: Stigmergic Media, Independence Protocols, and the Institutional Maintenance Term

Abstract. This paper derives an identity — at any steady state of the storage law, headroom equals unused variety, $h_\infty = \eta\,(\ln N - F(u^*)) = \eta \ln(N/N_{\text{eff}})$, in which δ, C and Q all cancel — and uses it to read three constructed cases. A tool library run through a coordinator converts at $\eta = 0.35$ with $\delta = 0.25$ and settles at $Q_\infty = 0.2653$; the same library run on a shared written medium converts at 0.55 with $\delta = 0.12$ and settles at 3.9281, and a turnover event costing the coordinator regime 70% of its stock costs the medium regime 15%. Nesting is priced: with per-path cost 0.02 and $\eta = 0.5$, headroom peaks at five members and goes negative between eleven and twelve, so a flat body of 24 has $h = -3.931$ while four cells of six plus a four-delegate federation gives $+2.957$. A nine-person review panel at $u = 0.12$ has $N_{\text{eff}} = 1.96$ and a negative fixed point, $Q_\infty = -0.6510$; five independence protocols costing 0.090 per item raise u to 0.71 and Q_∞ to $+1.9144$, and lift panel accuracy from 65.8% to 74.3% with individual accuracy held at $p = 0.62$. A standards body traced from 1962 shows the same arithmetic at institutional scale: twelve consecutive years of negative headroom and 61.9% of stored order lost. Two counter-results are reported: cutting decay alone makes a failing body worse, and shedding without spending ends poorer than doing nothing at all.

Keywords: stigmergy; shared medium; commons design principles; nesting; independence; cascade; institutional decay; shedding

Series: B. Engineering the Surplus **Paper:** CCP-014

14.1 1 Introduction

A framework that prices what a group accumulates has to be able to say where the accumulation is kept. The storage law of this corpus, `Q_{t+1} = (1−delta) Q_t + eta S_t − C_t` with `S_t = ln(1 + (N−1)U_t)`, is silent on that question: it describes a stock, a conversion rate and a loss rate, and it does not say whether the stock sits in a person, in an artefact, or in a rule. This paper shows that the question is not a detail of implementation but a setting of the law's own parameters, and that three apparently unrelated design problems — coordination without a centre, the preservation of individual judgement inside a deliberating body, and the decay of long-lived institutions — are the same arithmetic read at three scales.

The organising result is an identity. At any steady state of the storage law the headroom gauge collapses to

```
h_inf = eta ( ln N − F(u∗) ) = eta ln ( N / N_eff )
```

so that **headroom at equilibrium is exactly the variety the system is not using.** `delta`, `C` and `Q` all cancel. The identity is one line of algebra from the definition of a steady state, it is verified here to six decimal places, and it does three things that the surrounding argument then leans on. It says that a settled system never has negative headroom, so negative headroom is always a disequilibrium condition — inherited stock under changed parameters. It says that the framework's master gauge is a measure of unspent difference rather than of spare budget. And it corrects a claim the corpus makes elsewhere and this paper makes carefully: **the plain linear storage law, left alone, does not ossify.** Maintenance load rises with success and then stops rising, because the stock stops rising. Ossification requires a second mechanism — exogenous accretion of `Q`, or a cost term superlinear in `Q` — and a diagnosis has to say which.

The paper's contribution relative to the surrounding literature, rather than relative to the framework, is fourfold. First, it prices Ostrom's eighth design principle. Nesting is usually argued for as an administrative preference; here the quadratic coordination term `n(n−1)/2` [Brooks, 1975, Conway, 1968] is set against the logarithmic ceiling `eta ln N`, and the crossing is computed: at a per-path cost of 0.02 and `eta = 0.5`, headroom peaks at five members and goes negative between eleven and twelve, so a flat body of 24 sits at `h = −3.93` while four cells of six plus a four-delegate federation reaches `+2.957`. Second, it converts the independence clause of the Condorcet jury theorem from a footnote into a budget line, by showing that the same panel's effective count, its surplus term and its measured accuracy move together, and that a nine-person panel operating at `u = 0.12` performs at the theoretical accuracy of an infinitely long informational cascade [Bikhchandani et al., 1992]. Third, it separates two levers that are routinely bundled and shows that one of them does not pay: cutting decay alone, by building trust, makes a body with a negative numerator *worse*. Fourth, it treats deliberate shedding as an engineering operation with a computable return — `dh/dQ = −(delta + kappa)` per unit shed, permanently — and then reports that shedding alone is not a strategy, because a counterfactual body that sheds and spends nothing ends poorer than one that did nothing at all.

The framework's static half — the surplus functional `F(u) = ln(1 + (N−1)u)`, its characterisation under four axioms, and the affine effective count `N_eff = 1 + (N−1)u` — is presupposed rather than re-derived; it is established with its axiom set in Curzi [2026], whose design axiom S3 makes the effective count affine by choice rather than by discovery. The dynamical half — that the update map is Lipschitz with modulus `k = 1 − delta`, so that `delta > 0` makes it a contraction with steady state `Q∗ = (eta S − C)/delta` — is established in Curzi [2026a]. The measurement protocol that would put any applied claim below at risk, together with its five required operational choices and its shuffled-input control, is Curzi [2026i].

Sections 2 and 3 state the hypotheses and the framework, including the identification problem that bounds every applied claim in the paper. Section 4 develops coordination without a centre and prices nesting. Section 5 develops the independence requirement and the five protocols that answer it. Section 6 develops the institutional maintenance term and the shedding operation. Section 7 works three systems end to end and reports six arithmetic discrepancies found in the source material. Sections 8 to 11 discuss what follows, verify the chain, state the limitations, and conclude.

A note on labels. Three families of numbered item appear below and one has been renamed to keep them apart. Ostrom's eight design principles are referred to as **O1–O8**; the five independence protocols of Section 5 as **IP-1–IP-5**; and the verification postulates of Section 9 as **P1–P7**. The design postulates carried over from the source material keep their own labels, **M1–M3** and **I1**. Specification clauses are numbered by their section (S4.1, S5.1, S6.1) and failure modes likewise (FM-4.1, FM-5.1, FM-6.1).

14.2 2 Hypotheses

Two of the results below are derivations about an equation and cannot be falsified by observation; they are labelled as such where they appear. What can be falsified is the claim that the equation describes real coordinating, deliberating and rule-holding systems, and the claim that specific interventions move specific parameters. The hypotheses separate the two and each names the observation that would defeat it. The protocol that operationalises them — a preregistered fit, a frozen model, and held-out comparison against four declared baselines — is specified in Curzi [2026i], whose own confirmatory application failed its second falsification rule, and whose shuffled-input control passed 98.35 per cent of the time against a third variable driving both interaction and order while detecting only 48.17 per cent in the clean case. That result bounds every empirical reading in this paper and is restated in Section 10.

H1 — State location is a parameter setting, not an implementation detail. Moving a system's state out of a coordinating person and into a shared externally readable medium lowers `delta` and lowers `C` by more than it lowers `eta`, so headroom rises. **Predicted:** a measured fall in `delta` and a measured rise in `eta` across the transition, with the steady state moving by the ratio the storage law predicts from the two measured parameters and nothing else. **Falsified by:** a measured `delta` that does not fall after the medium is installed; or a headroom gauge that does not rise; or a steady state that moves by an amount the measured parameter changes do not account for. This is design postulate M1 stated as a testable claim, and it is falsifiable separately in every installation.

H2 — Coordination cost is quadratic in headcount while the ceiling is logarithmic, so flat bodies have a maximum workable size. Predicted: headroom at `Q = 0` follows `h(n) = eta ln n − c_path n(n−1)/2`, so it has a single interior maximum and one zero crossing; at `eta = 0.5` and `c_path = 0.02` the maximum is at `n = 5` and the crossing lies between `n = 11` and `n = 12`. **Falsified by:** a measured coordination cost that does not grow faster than logarithmically in headcount; or a flat body operating above its computed crossing point that nonetheless commissions durable new structure over several time constants.

H3 — Panel accuracy is governed by the independence of the members' errors and not by their individual ability. Predicted: a panel whose members' errors are strongly correlated performs at or near the unbounded-cascade accuracy `p(1+p)/[2(p^2−p+1)]` rather than at the independent-majority accuracy; restoring independence by changing the order of operations raises measured accuracy with individual accuracy held fixed. At `p = 0.62` the cascade value is 0.6570 and the nine-voter independent value is 0.7738. **Falsified by:** a panel with measured error correlation near one whose accuracy matches the independent-majority prediction; or an accuracy that does not move when the order of operations is changed and composition is held fixed.

H4 — Difference decays and must be maintained. Under shared training, shared incentive and prolonged co-location, `u` declines monotonically absent maintenance. **Predicted:** repeated measurement of `u` in an unmaintained standing body shows a monotone decline, and the decline appears in the recorded dissent rate before it appears in the decisions. **Falsified by:** a standing body under those conditions whose measured `u` is stationary or rising without any maintenance protocol in place. **This is design postulate I1 and no drift rate has been measured anywhere in this corpus**; the hypothesis is stated so that an installation can measure its own.

H5 — Negative headroom is always inherited stock under changed parameters. Predicted: wherever `h < 0` is measured, at least one of three routes is simultaneously measurable — `eta` fallen, `C` risen as a function of `Q`, or `N` narrowed — and the

stock exceeds the zero-headroom level `Q* = (eta ln N − c_0)/(delta + kappa)` computed from the current parameters. **Falsified by:** a body at a genuine steady state with negative measured headroom, which the identity of Section 3.2 forbids; or a measured `h < 0` in which none of the three routes is present.

H6 — Cutting decay alone does not pay when the numerator is negative. Predicted: in a body with `eta F(u*) < c*`, an intervention that lowers `delta` and touches nothing else lowers `Q_inf` further, because a negative numerator divided by a smaller positive denominator is more negative; and the same intervention lengthens the time constant, so the body takes longer to discover it. **Falsified by:** a measured case with a negative step in which lowering `delta` alone raises the steady state.

Three claims below are not hypotheses and are labelled accordingly. The identity `h_inf = eta ln(N/N_eff)` is a **derivation** from the definition of a steady state and carries no empirical content. The application of the storage law to a tool library, a review panel and a standards body alike is a **modelling postulate** about form, as badged in Curzi [2026a]; it is falsifiable at each scale separately and never in general. The identification of a coordinating centre with a topological hub, design postulate M3, is a **modelling choice** that imports a robustness profile [Albert et al., 2000] into a setting where the degree distribution has not been measured, and Section 10 records that the import is unearned.

14.3 3 Framework: the storage law, the identity, and what cannot be identified

14.3.1 3.1 What is presupposed

The surplus functional. For two unit state vectors drawn from an ensemble of `N` categories, with `u = 1 − (x . y)^2` in `[0,1]`,

```
F(u) = ln(1 + (N−1)u),    N_eff = 1 + (N−1)u,    0 <= F(u) <= ln N
```

`F` is uniquely characterised by four axioms — angle-only dependence, zero at parallel, an affine effective count, and normalisation at orthogonality — as established in Curzi [2026], where the affine-effective-count axiom S3 is badged as a design axiom chosen rather than derived. Extremality gives `F = 0` if and only if the states are parallel and `F = ln N` if and only if they are orthogonal, and strict concavity makes the approach to the ceiling front-loaded. Both properties are used repeatedly below.

The effective count is not a Hill number, and this paper says which quantity is which. The effective-number idiom is standard in ecology and political science, where `1/HHI` and `exp(H)` are the canonical members of the family — Hill numbers of order two and order one respectively [Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979]. The quantity used throughout this paper, `N_eff = 1 + (N−1)u`, is **not** a member of that family: it is a linear interpolation between 1 and `N`, it is this framework's own construction under design axiom S3, and no external literature contains it. Wherever an effective count appears below it is this one, and the Hill-number citations are reserved for the two quantities to which they actually apply. Section 10 carries the exposure.

Two warrants for the logarithm, and the argument needs only the weaker one. The Fechnerian warrant [Fechner, 1966] requires a contested perceptual story. The stronger warrant assumes nothing psychological: the time-average growth rate of a multiplicatively evolving quantity is the expectation of its logarithm [Peters and Gell-Mann, 2016], so where a stock compounds rather than adds the logarithm is the ergodic-average object. The ceiling `ln N` also has an information-theoretic reading as the maximum an ensemble of `N` equiprobable categories can carry [Shannon, 1948], and the framework does **not** identify surplus with Shannon entropy; the analogy is licensed to bound and to do nothing else.

The functional form is underdetermined by the evidence, and this paper says so rather than waiting to be told. Social impact rising as a power of the number of sources with an exponent below one [Latané, 1981] and network value growing as `n log n` [Briscoe et al., 2006] reproduce the same qualitative facts — diminishing returns to count, no threshold, monotonicity. **Nothing assembled in this paper discriminates between them.** What the corpus's form has that the alternatives do not is an axiomatic characterisation with the design choice marked at the point where it is made; what it does not have is a discriminating measurement.

The storage law and the headroom gauge.

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,     S_t = F(U_t) = ln(1 + (N−1) U_t)

h_t = eta ln N − delta Q_t − C_t
```

with `Q >= 0` the stored order, `delta` in `[0,1]` the fraction lost per period without maintenance, `eta >= 0` the conversion of interaction into durable structure, and `C >= 0` the cost. The quartet of a level, an inflow, an outflow and a delay is the standard vocabulary of system dynamics [Forrester, 1961]. The results used throughout are established in Curzi [2026a]: the update map `T(Q) = (1−delta)Q + x` is Lipschitz with modulus `1−delta`, so `delta > 0` makes it a contraction with the unique start-independent fixed point `Q_inf = (eta F(u∗) − c∗)/delta` [Banach, 1922, Elaydi, 2005]; the uniform bound is `Q_t <= max{Q_0, eta ln N / delta}`; and a non-decreasing step requires `C_t <= eta ln N − delta Q_t`, beyond which no interaction geometry whatsoever prevents stored order from decreasing. Headroom is the framework's name for organisational slack, which Cyert and March [1963] identified as the buffer that absorbs variation, and its behaviour near exhaustion has the same shape as utilisation approaching one in a queue [Kingman, 1962, Hopp and Spearman, 2011].

The increment requirement. Rearranging the step condition, achieving an increment `r` per period requires `S >= (delta Q + C + r)/eta`, which is attainable only if `delta Q + C + r <= eta ln N`. Setting `r` to its largest attainable value returns `r_max = h_t` exactly: the gauge states what may be asked for, and the geometry states what can be delivered, and they meet only at `u = 1`. This is one line of algebra from the storage law and is used in Section 7.2.

14.3.2 3.2 The identity: headroom at equilibrium is unused variety

At any steady state, `Q_{t+1} = Q_t`, so from the recursion

```
Q = (1 − delta) Q + eta S − C
0 = −delta Q + eta S − C
delta Q + C = eta S
```

Substituting into the headroom gauge,

```
h_inf = eta ln N − delta Q − C
      = eta ln N − eta S
      = eta ( ln N − ln N_eff )
      = eta ln ( N / N_eff )
      = eta ( ln N − F(u∗) )
```

At equilibrium, headroom equals the conversion efficiency times the logarithm of the redundancy still present in the roster. The derivation uses only the definition of a steady state, so it holds however `C` depends on `Q`. This is result **R1** of the paper and it is verified numerically twice in Section 7 and once more, on the smallest case, in Section 9.4.

Three consequences follow immediately.

1. A system at equilibrium with `u = 1` — perfectly orthogonal participants — has `h_inf = 0`. It is running at its ceiling. It is stable and it cannot grow, and nothing is wrong with it.

2. **A system at equilibrium never has negative headroom.** Negative headroom is therefore always a disequilibrium condition: the stored order was accumulated under one parameter set and is being carried under another. Section 6.6 is entirely about that case.
3. **`delta`, `C` and `Q` all cancel at the fixed point.** This is the consequence that forces a correction on the received account of ossification, developed in Section 6.1.

14.3.3 3.3 The design postulates, badged

The following are **design postulates**. They are not derived from the results above and they are not empirical findings; they are engineering heuristics that the law makes legible, which is a weaker and more honest claim. Each is badged again wherever it is leaned on.

- **M1 (state-in-medium).** Moving a system's state out of a coordinator and into a shared, externally visible medium lowers `delta` and lowers `C` by more than it lowers `eta`, so headroom rises. *Status: POSTULATED.* Falsifiable in any installation by measuring `delta`, `eta` and `C` before and after by the protocol of Curzi [2026i]. **If a shared medium is installed and headroom does not rise, the postulate has failed in that system and the report must say so.**
- **M2 (medium quality).** The `eta` of a stigmergic system is dominated by five properties of its medium — visibility, writability, durability, low friction, and append-orientation — and these can be specified and inspected directly. *Status: POSTULATED.* Section 4.4 gives the five inspection tests; the postulate's content is that the list is sufficient, and that is what a failed installation would refute.
- **M3 (centre as hub).** A coordinating centre is a topological hub and therefore inherits the robustness profile measured for hubs in scale-free networks — robust to random failure, fragile to targeted removal [Albert et al., 2000]. *Status: POSTULATED, and the weakest of the three.* The import is unearned: the degree distribution of a coordination graph has not been measured here, and Doyle et al. [2005] argues that the robustness of real engineered networks comes from design rather than from a degree distribution. **No scaling or heavy-tail claim is made anywhere in this paper**, so the distribution-fitting procedure such a claim would require [Clauset et al., 2009] does not apply; the qualitative asymmetry, not the exponent, is what M3 borrows, and Section 10 records that even that is borrowed rather than shown.
- **I1 (the merge).** Human groups under shared training, shared incentive, prolonged colocation and social pressure drift toward parallelism: `u` declines monotonically absent maintenance. *Status: POSTULATED.* It is consistent with the consensus-formation literature — repeated averaging of opinions over a connected influence network converges to a single value [DeGroot, 1974, Friedkin and Johnsen, 1990], and bounded-confidence variants converge to a small number of clusters [Hegselmann and Krause, 2002] — and with the measured collapse of crowd wisdom under social influence [Lorenz et al., 2011, Salganik et al., 2006]. **No drift rate has been measured in this corpus**, and any installation that needs one must measure its own.

What *is* derivable given I1 is the consequence: `u` is not a property a group acquires but **a quantity that decays**, so maintaining it is a retention problem as much as a design problem.

14.3.4 3.4 Identification: the objection that no amount of fitting can answer

This limitation is structural rather than statistical and belongs here, in the framework section, rather than in a caveat at the end. It is acutely relevant to this paper, because independence protocols, cascades and panel accuracy are exactly the setting in which it bites.

The framework's central applied claim is reflexive: the interaction geometry `U_t` is a property of the group, the group's stored order shapes what its members do next, and the members are what the group is made of. Written as `T(x) = W(x, M(x))`, that is a linear-in-means reflexive specification, and Manski [1993] proves that **the endogenous social effect in such a**

specification is not identified from observational data: it cannot be separated from contextual effects or from correlated unobservables, however much data is collected. Shalizi and Thomas [2011] reinforce the point for network data, where influence and selection are generically confounded.

The consequence for this paper is exact. Sections 3.1 and 3.2 are derivations about an equation and are untouched. But every applied reading below — that installing a medium raised `eta`, that the independence protocols raised `u`, that the correlation parameter `rho` fell from 0.75 to 0.20, that panel accuracy rose because the order of operations changed — is an interpretation of a reflexive specification and **is not identified by observation alone.** No held-out comparison recovers it. This is deeper than anything the falsification protocol of Curzi [2026i] tests for, and it would remain even if every rule in that protocol passed.

No remedy is adopted in this paper, and the honest label is therefore a modelling position rather than an identified empirical result. The known routes are randomisation, an instrument, partial-population variation, or stipulating the reference group and defending the stipulation; the last is the most realistic here, and Brock and Durlauf [2001] is its template. The five operational choices of Curzi [2026i] — a fixed category set, a declared participation rule, a stated time scale — point toward that route without yet taking it. Until one of them is executed, **`eta` and `u` as estimated quantities are coefficients in a declared specification and not identified social effects**, and every applied statement in Sections 4 to 7 inherits that status.

Two related objections are recorded here and answered in Section 10 rather than deferred. The **Lucas critique applied to the framework itself** [Lucas, 1976]: Section 6.9 asserts that a measure changes the system measured by it, and the paper then treats `delta`, `eta` and `N` as constants over an estimation window; both cannot hold without qualification, and the narrow position taken is stated in Section 10. And an existence proof does exist for reflexivity as a mechanism — MacKenzie and Millo [2003] documents a reflexive fixed point in the field, with dates and price data, where a model's adoption moved prices toward the model. That is the corpus's own mechanism caught happening; it establishes that reflexive fixed points occur, not that this framework's parameters are recoverable.

14.3.5 3.5 Standing honesty clauses

These are repeated wherever the law is used and are stated once here.

- **The static half of the framework is pure mathematics** and makes no empirical or metaphysical claim.
- **Surplus is not identified with Shannon entropy or mutual information in general**, and the entropy invoked anywhere near `ln N` is Shannon entropy — a dimensionless information measure in nats — and not thermodynamic entropy.
- **As deployed, the surplus functional is a lower bound** on the information an ensemble carries. Every number produced by this method is a floor.
- **The triangle inequality fails for geometric triples**, so surplus bounds may not be chained through an intermediate state as though `F` were a metric. Psychological similarity is known to be asymmetric and to violate the triangle inequality [Tversky, 1977].
- **Applying the storage law to a tool library, a panel or a standards body is a modelling decision**, badged as such every time it is made [Curzi, 2026a].
- **No mechanism in this paper may be implemented in a form that removes a participant's ability to decline.** This is a scope constraint on the framework, not a courtesy, and it determines which protocols appear in Section 5.7 and which are excluded there by name. The corpus's own classification of what its protocols rest on — four items of established mathematics, seven of established psychology, and exactly three declared postulates, with every protocol stated to run on the first two and steer by the third — is given in Curzi [2026], and the doctrine that forbids moving a harm through a discovered gap in a scheme, and that concedes the case where a feedback loop runs through another person's will as an unsettled seam, is given in Curzi [2026]. Both bound what this paper may recommend.

14.4 4 Order without a centre: coordination through a shared medium

There is a habit of mind old enough that it is rarely stated as a proposition and therefore rarely tested: that order requires an orderer. Where there is a pattern, look for the patterner; where a hundred people arrive at a coherent result, locate the person who held the plan. The habit is often correct. It is not a law, and this section concerns the class of systems in which it is false, and how such systems are built deliberately.

The storage law does not know whether a coordinator exists. What changes when the centre is removed is not the equation but the parameters, and the claim is worth stating in its strongest form before any evidence is introduced:

> **A coordinator is a device that raises `eta` at the price of `C`, and concentrates `Q` in a single store whose `delta` is the coordinator's own turnover rate.**

Each clause is separately checkable. A coordinator raises `eta` because someone is paid to convert loose interaction into durable record — that is what the role is. A coordinator raises `C` because the role costs money, attention, scheduling and queueing. And a coordinator concentrates `Q`, because the state of the system now lives substantially inside one person's working memory, calendar and relationships. When that person leaves, the system does not lose a worker; it loses a fraction of its stored order in a single step. Section 7.1 computes that fraction for a worked case.

14.4.1 4.1 Stigmergy: the mechanism, and two honest limits on the biology

Grassé, studying nest construction in termites of the genus *Bellicositermes*, faced a problem his field had been avoiding. A termite mound is a coherent structure with ventilation shafts, brood chambers and a royal cell, built by hundreds of thousands of insects none of which possesses a representation of the mound. His 1959 answer was that the workers do not coordinate with each other at all: they coordinate with the work already done. A deposit changes the local environment; the changed environment is the stimulus that makes the next deposit more likely at that spot. The structure grows because **the state of the project is stored in the project**, and each worker reads the state by looking at it. He named the mechanism *stigmergy*, from the Greek for mark and for work: the work is the sign.

> **Stigmergy: coordination in which agents communicate only by modifying a shared medium, and no agent holds the plan.**

Note what stigmergy is not. It is not self-organisation as a mood, and not emergence as mystique. It is a specific mechanism with two named requirements — a medium that agents can both read and write, and a local rule mapping what is read to what is written. Both are engineerable, and both are frequently absent in organisations that describe themselves as decentralised.

The ant trail is the same mechanism in one dimension. Pheromone attracts; more traffic means more pheromone; pheromone evaporates, so a trail to an exhausted source fades. The trail network is stored on the ground, and its `delta` is the evaporation rate of the trail chemical. That last point is the framework in miniature: **the evaporation is not a defect of the ant trail but the mechanism by which the colony forgets a food source that is gone.** A pheromone that did not evaporate would maintain a permanent motorway to an empty patch of ground. This is the corpus's `delta = 0` result arriving from entomology — zero loss is not perfect memory but the inability to revise [Curzi, 2026a] — and the organisational-learning literature reaches it from the other side, where an organisation whose members learn too fast homogenises and destroys the variance that produced its advantage [March, 1991, Levinthal and March, 1993].

Two limits on the biology, stated now so that they are not smuggled in later.

- Termite and ant behaviour is used here as an **existence proof of a mechanism**, not as a model of human organisations. Humans read media in ways ants do not: they infer intention, they anticipate, they lie, and they defect strategically. Every human stigmergic system therefore needs the governance layer of Section 4.5, which the termites do not need.
- Grassé described stigmergy; he did not quantify it in the parameters of this framework, which did not exist. **The mapping from stigmergy to `(delta, eta, N, u, C)` is this paper's modelling choice**, badged as such. *No citation key exists in this corpus's bibliography for Grassé's 1959 paper*, and Section 10 records the gap rather than attaching a citation that does not support the sentence.

14.4.2 4.2 Two human media that work

The version-controlled repository. A distributed software project of a thousand contributors has nobody who has read all the code. The state of the project is the repository: the current tree, the history, the open issues, the test results. A contributor reads the state by cloning it, modifies it, and offers the modification back. The medium is visible, writable through a defined proposal mechanism, durable, low-friction and append-oriented. Maintainers exist but are not coordinators in the sense that matters: they do not hold the plan. If every maintainer vanished on Tuesday the state of the project would be intact on Wednesday — which is precisely what is not true of a project whose state lives in a programme manager's spreadsheet.

The encyclopedia anyone can edit. An article's state is the article, plus its history, plus its talk page. No editor has read the encyclopedia; there is no assignment of topics. The coordination signal is the visible deficiency of the current text — a missing citation, a contradiction, an obvious gap — which functions exactly as the pheromone-laden pellet does: the work already done is the stimulus for the next increment.

Both cases carry a **grammar of incompleteness**. A repository has failing tests and open issues; an article has citation-needed markers and stub tags. These are not administrative overhead. **A stigmergic medium must be able to display its own unfinished edges, or there is nothing for a contributor to respond to.** A medium that shows only finished states is not a stigmergic medium; it is a report.

Neither example is centreless in the naive sense. Both have governance, and both acquired it over years, usually after a crisis. Section 4.11 develops the point; the short form is that a shared medium is necessary and is not sufficient.

14.4.3 4.3 Focal points: the coordination that is never performed

Schelling's *The Strategy of Conflict* (1960) isolates the limiting case. Two people must meet in a city on a given day and cannot communicate. Empirically, people do far better than chance: they converge on **focal points**, solutions that are salient and that each expects the other to find salient. The landmark. Noon. The round number. The obvious default. *That 1960 work has no key in this corpus's bibliography*; Schelling's later and better-known demonstration that mild individual preferences produce extreme aggregate outcomes [Schelling, 1971, 1978] is a different result and is cited below only where that different result is what is being used.

A focal point is a stigmergic medium with zero write bandwidth: nobody modified anything, and the coordination is carried entirely by shared, publicly evident structure. **Before building any medium, the designer should ask what participants already share that is salient, because salience is free coordination.**

Focal device	What it removes
A fixed weekly time that never moves	The scheduling negotiation, permanently
A single canonical location for a class of document	The "where does this live" exchange

Focal device	What it removes
A default that is correct 80 per cent of the time	Eighty per cent of the decisions
A round-number threshold (“anything over 10k”)	The case-by-case escalation judgement
A naming convention	The lookup

Each is a `C` reduction, and `C` reductions convert directly into headroom by the gauge. They cost approximately nothing and are almost always available in a system that has never looked for them. **The most efficient coordination is the coordination that never has to be performed.**

14.4.4 4.4 What makes a good shared medium: the five properties and their tests

This is where design postulate M2 does its work. Five properties, each with an inspection test a surveyor can run in an afternoon. The mnemonic is **VWDLA**, and the survey form of Appendix A scores each from 0 to 4.

V — Visible. Any participant can see the current state without asking permission and without asking a person. *Inspection test:* pick three participants at random at three different levels of the system and ask each to show the current state of one live item, now, from where they are sitting. Time it. **If the median exceeds sixty seconds, or if any of the three has to ask someone, the medium fails V.** Visibility is the property most often lost by accident: a medium starts open, acquires access controls for good reasons, and acquires so many that reading it requires knowing whom to ask. At that point the system no longer has a shared medium; it has a coordinator wearing a database costume. The security-engineering form of the same requirement is old and well established — a system’s design should not depend on its own obscurity [Kerckhoffs, 1883, Saltzer and Schroeder, 1975].

W — Writable. Any participant who observes something can record it, and the recording takes effect without adjudication. *Inspection test:* count the steps between a participant noticing something and the medium showing it. **Two steps or fewer passes; four or more fails.** Count review as a step only if it is blocking. Writability is where most organisational media die: read-only dashboards are extremely common and are not stigmergic media at all but broadcast, and broadcast has `eta` near zero for the receivers, because nothing a receiver learns can be deposited back. The subtlety is that writability does not mean unmoderated — it means the default is write-through and the correction is after the fact. An append-only log with revert is writable; a submission queue with a two-week approval cycle is not.

D — Durable. The state survives the departure of any individual, and survives the interval between uses. *Inspection test:* the **bus-count** — for each critical piece of state, count how many people would have to leave simultaneously for it to become unrecoverable. **A bus-count of 1 for any critical item is a failure.** Then the **fortnight test**: is state recorded two weeks ago still intelligible today to someone who was not present when it was written? Durability is the direct `delta` control, and it is the parameter a coordinator-centred system gets catastrophically wrong, because a person’s memory has a `delta` that looks small in the short run and equals 1.0 on the day they resign.

L — Low friction. The cost of one read and one write is small relative to the value of the increment. *Inspection test:* measure the seconds to complete one typical write end to end, including authentication, navigation and any required fields. **Under thirty seconds passes; over two minutes fails and will be routed around.** Then count required fields: every mandatory field is a tax on every write, forever. Friction is not a minor usability concern but the term that decides whether the medium is used at all, and an unused medium has `eta = 0` however well designed. The characteristic failure is the elaborate system that captures everything perfectly and is filled in by nobody, sitting beside a group chat that captures nothing and holds the real state.

A — Append-oriented. New information is added; old information is superseded rather than erased; the history is retrievable. *Inspection test:* can the state as of a named date last quarter be reconstructed? **If not, the medium is not append-oriented**, and every dispute about what was decided will be resolved by whoever speaks most confidently. Append-orientation is what makes the medium adjudicative. It is also what makes deliberate forgetting possible *as a decision* rather than as an accident: a rule can be retired explicitly, and the retirement is itself a record. Section 6.7 depends on this.

14.4.5 4.5 Ostrom's eight design principles, in full

Ostrom [1990] is the most directly engineerable document in the sociology of shared resources. Its method matters as much as its result: rather than theorising about whether commons must fail, Ostrom and her collaborators assembled and compared a large body of field cases of common-pool resources — Swiss alpine pastures, Japanese village forests, Spanish irrigation systems, Philippine irrigation federations, inshore fisheries, groundwater basins — and asked what distinguished the arrangements that persisted for centuries from those that collapsed. The polycentric extension [Ostrom, 2010] generalises the nesting result to overlapping jurisdictions.

A **common-pool resource** in this sense has two features: it is *subtractable* (what one person takes, another cannot) and it is *hard to exclude people from*. Grazing land, fish stocks, groundwater, a shared workshop, a shared codebase, a shared meeting-hour budget, the attention of a small team.

The eight principles are given below in full, referred to as **O1–O8**. **Ostrom did not claim these are sufficient and did not claim they are a recipe.** She characterised them as features found in the enduring cases and absent or weak in the failed ones. That is the claim this paper repeats and it is not a stronger one. The largest external audit of the principles [Cox et al., 2010] finds broad but uneven empirical support across ninety-one studies, and that qualification travels with every use of the list below.

O1 — Clearly defined boundaries. Individuals or households with rights to withdraw resource units must be clearly defined, as must the boundaries of the resource itself. *Why it matters:* if one cannot say who is in and what the resource is, one cannot say what over-use looks like, and every participant's estimate of every other's restraint is a guess. *Parameter:* `N` and `C`. Undefined boundaries make the effective category count unknowable and generate continuous boundary-dispute cost.

O2 — Congruence between appropriation and provision rules and local conditions. *Why it matters:* two congruences are required and are usually confused. The rules must fit the physical conditions — a rule written for a wet valley will destroy a dry one — and what may be *taken* must be proportionate to what is *put in*. A system in which extraction rights are detached from contribution obligations has instructed its members to free-ride. *Parameter:* `C` and `eta`.

O3 — Collective-choice arrangements. Most individuals affected by the operational rules can participate in modifying them. *Why it matters:* this is the principle most often violated by well-intentioned outsiders who write good rules for people who had no hand in them. Rules written by the users are better fitted, because the users hold the local knowledge, and better complied with, because compliance with a rule one helped write is not obedience. *Parameter:* `u` and `eta`. Participation in rule-making is exactly the mechanism that puts multiple distinct viewpoints into the rule, and a rule one helped write has already converted into one's own practice.

O4 — Monitoring. Monitors who actively audit resource conditions and appropriator behaviour are accountable to the appropriators or are the appropriators themselves. *Why it matters:* this is Ostrom's sharpest empirical finding and the one that overturns the standard framing. The enduring commons are not populated by unusually virtuous people but by ordinary people who **can see what everyone else is doing.** In the durable cases monitoring is mutual, cheap, and largely a by-product of the work itself — the irrigator who opens the gate

can see the level in the channel and can see who opened it last. *Parameter:* `C`, via the free-riding loss term. Section 4.8 computes the trade.

O5 — Graduated sanctions. Appropriators who violate operational rules are assessed graduated sanctions depending on the seriousness and context of the offence. *Why it matters:* graduation is the load-bearing word. A system with one severe sanction has in practice no sanction: the penalty is so disproportionate to a first small lapse that nobody applies it, everybody knows nobody applies it, and the rule becomes advisory. A first offence drawing a word, a second a small fine, a third exclusion, is a system in which the first step is actually taken. **Graduated sanctions are cheap precisely because they are mild, and effective precisely because they are used.** *Parameter:* `C`.

O6 — Conflict-resolution mechanisms. Appropriators and their officials have rapid access to low-cost local arenas. *Why it matters:* every rule is ambiguous at some boundary and every commons generates disputes; the question is never whether disputes occur but how expensive they are. *Rapid*, *low-cost* and *local* are all three doing work: a mechanism that takes six months, costs a lawyer, or requires travelling to a capital will not be used, and unresolved disputes compound. *Parameter:* `C`, and it is the term most likely to be omitted from a cost tally because it is nobody's line item.

O7 — Minimal recognition of rights to organise. The rights of appropriators to devise their own institutions are not challenged by external governmental authorities. *Why it matters:* self-governance that can be overridden at any moment is a delegation that will be revoked at the first inconvenience. The requirement is *minimal* recognition: Ostrom's cases did not need the external authority to endorse or fund the local institution, only to refrain from voiding it. *Parameter:* `eta`. If local decisions can be nullified from outside, interaction does not convert into durable structure. This is `eta` collapse by decree and is the most common way a head office destroys a functioning local commons while believing it is helping.

O8 — Nested enterprises. For common-pool resources that are parts of larger systems, appropriation, provision, monitoring, enforcement, conflict resolution and governance are organised in multiple layers of nested enterprises. *Why it matters:* it is the answer to scale. A single governance body for ten thousand people cannot monitor, cannot resolve disputes locally and cannot fit rules to local conditions. *Parameter:* `C`, decisively, and it is computable — which is the subject of the next subsection.

14.4.6 4.6 Nesting priced: the quadratic term against the logarithmic ceiling

The communication-path count `n(n−1)/2` grows quadratically in headcount [Brooks, 1975, Conway, 1968] while surplus grows as `ln N`. With a per-path cost of 0.02 and `eta = 0.5`, headroom at `Q = 0` is `h(n) = eta ln n − 0.02 n(n−1)/2`:

n	paths	C	eta ln N	h
4	6	0.120	0.6931	+0.5731
5	10	0.200	0.8047	**+0.6047**
6	15	0.300	0.8959	+0.5959
8	28	0.560	1.0397	+0.4797
10	45	0.900	1.1513	+0.2513
11	55	1.100	1.1989	+0.0989
12	66	1.320	1.2425	**-0.0775**
16	120	2.400	1.3863	-1.0137
24	276	5.520	1.5890	-3.9310

Headroom peaks at `n = 5` and goes negative between 11 and 12. A single flat body of 24 has `h = −3.93`: it cannot build anything, ever, at any interaction geometry.

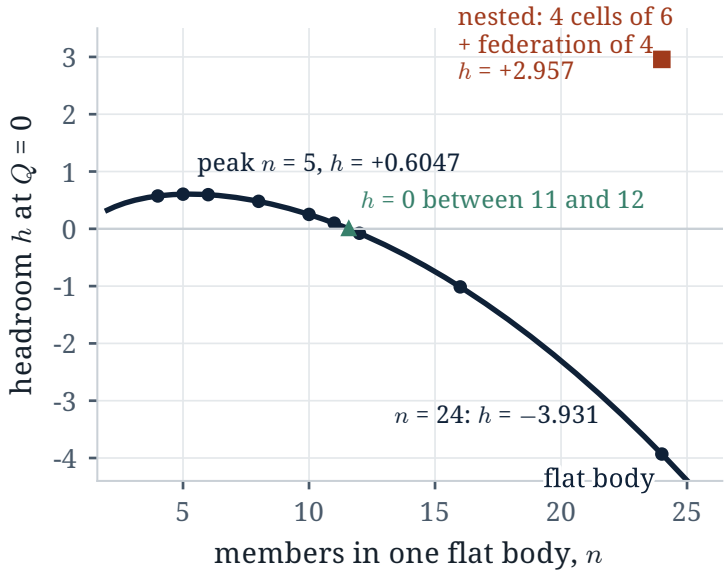


Figure 14.1: **Headroom has a single interior maximum in flat-body size and one zero crossing**: at a per-path cost of 0.02 and `eta` = 0.5, `h` peaks at `n` = 5 with +0.6047 and crosses zero between `n` = 11 and `n` = 12, so a flat body of 24 sits at −3.931 while four cells of six plus a four-delegate federation reaches +2.957. Modelled from the two stipulated parameters of Section 4.6; generated by `figs/src/CCP-014_fig01.py`.

Now nest it. Four cells of six, each at `h = +0.5959`, plus a federation layer of four delegates at `h = +0.5731`. *Modelling assumption, badged:* the law is stated per system, and summing headroom across cells assumes the cells are genuinely independent in their cost structure — if they are not, the sum overstates. Under that assumption the nested form reaches `4 x 0.5959 + 0.5731 = 2.957` against the flat form's `-3.931`. This is result **R3**.

O8 is not an administrative preference. It is the difference between a positive and a negative headroom budget, and the arithmetic is elementary. Note what the figure does *not* show: the crossing point is a function of the per-path cost, which is a stipulated parameter here and a measurable one in any installation. A body with a per-path cost of 0.005 crosses at `n = 34`, and a body at 0.05 crosses at `n = 6`. The shape is the result; the numbers belong to the parameters.

14.4.7 4.7 The install order, and a correction to naive practice

The eight principles are a set, not a sequence, and the temptation on reading them is to start with the one that feels most like community: build trust, and the rest will follow. **That is the wrong order, and in a system that is already losing it is actively harmful.**

The mechanism is the steady-state theorem and nothing more: `Q_inf = (eta F(u*) - c*)/delta` [Curzi, 2026a]. Trust acts on `delta` — a group that relies on unverified commitments loses less order between periods, and the mechanism by which it does so is retention: trust enters the storage law as a modification of the decay term, `delta_eff = delta_0 + (1-a)(1-v)`, giving a superlinear stock dividend as the trust parameter rises [Curzi, 2026s]. **But when the numerator is negative, dividing by a smaller positive denominator makes the result more negative.** Separating the moves in a worked consortium and computing them:

move	`Q_inf`
baseline	-0.0836
trust alone (`delta` 0.4300 -> 0.1700)	**-0.2114 — worse**
reciprocity alone (`eta` 0.1721 -> 0.4026)	+0.7394

move	`Q_inf`
both	+1.8701

Trust does not add order. It amplifies whichever sign the system already has. Halving `delta` on a negative numerator multiplied the loss by the same factor it would have multiplied a gain. This is the first half of counter-result **CR**, and it is stated as a result rather than a caveat because it reverses a widely recommended sequence.

The design rule that follows is exact: **fix the numerator first — raise `eta` through reciprocity, or cut `C` — until `eta F(u*) > c*`, and only then invest in trust.** Read against Ostrom's list, this puts O4 and O5 ahead of anything working by extended credit, because Section 4.8 shows both acting on the numerator; and it puts O8 ahead too, because Section 4.6 shows it changing the exponent on `C`. The failure signature to watch for is specific: **a trust programme in a body that converts nothing durable produces a slower, more polite decline, and survey results saying the intervention worked.** Self-report about the causes of one's own performance is exactly the measurement known to be unreliable [Nisbett and Wilson, 1977], which is why the signature is not self-correcting. The reciprocity literature supplies the mechanism the numerator-first rule actually needs — repeated interaction with reputation, rather than generalised goodwill [Axelrod, 1984, Nowak and Sigmund, 2005].

14.4.8 4.8 The two named failure modes of the commons, and the structural answers

The centreless case has two classical objections. Both are real, both have answers in the eight principles, and the answers are structural rather than moral.

The tragedy of the commons. The argument in its standard form [Hardin, 1968]: a shared pasture, herders each adding animals; the benefit of an extra animal accrues entirely to its owner while the cost of the extra grazing is spread across all herders; each herder therefore faces a private incentive to add animals until the pasture is destroyed. The arithmetic is correct for the situation described. **What the argument actually models is not a commons but an unmanaged open-access resource** — no boundaries, no rules, no monitoring, no sanctions, no communication. Ostrom's field cases are what happens when those things exist, and her central empirical point is that they very frequently do [Ostrom, 1990].

Mechanism of tragedy	Principle that addresses it
Anyone may take, so no restraint is credible	**O1** boundaries: define who may withdraw and from what
No limit on how much may be taken	**O2** congruence: appropriation limited and tied to provision
Restraint is invisible, so it is unrewarded	**O4** monitoring: restraint becomes observable
Over-use carries no consequence	**O5** graduated sanctions
Rules imposed from outside do not fit the pasture	**O3** collective choice, **O7** right to organise
Disputes over the boundary destroy the arrangement	**O6** low-cost local arenas
The pasture is bigger than any one village	**O8** nesting

The point to carry: **the tragedy is a property of a missing institution, not of human nature.** Where the eight features are present the tragedy is not observed; where they are absent it is. That is a claim about the world, from field data, and the meta-analytic audit of it is Cox et al. [2010].

Free riding. The narrower and more common problem: the resource is not destroyed, but a fraction of participants consume the shared order without contributing to it. Nothing collapses; everything costs more, degrades faster, and quietly demoralises the contributors. In the storage law free riding is unambiguously a `C` term — every unit of order consumed without replacement is a dissipation. Model it as `C = C_base + lambda phi`, with `phi` the free-riding fraction and `lambda` the loss per unit of it. Taking the tool library of Section 7.1, with `eta = 0.55`, `delta = 0.12`, `N = 8`, `u = 0.62`, `S = 1.6752`, `C_base = 0.40` and `lambda = 0.9`:

Without monitoring, at `phi = 0.35`:

```
C     = 0.40 + 0.9(0.35) = 0.7150
h(Q=0)= 0.55(2.0794) − 0.7150 = +0.4287
Q_inf = (0.55 x 1.6752 − 0.7150) / 0.12 = 1.7198
```

With O4 and O5 — mutual monitoring at a cost of 0.08, graduated sanctions bringing `phi` down to 0.06:

```
C     = 0.40 + 0.08 + 0.9(0.06) = 0.5340
h(Q=0)= 0.55(2.0794) − 0.5340 = +0.6097
Q_inf = (0.9214 − 0.5340) / 0.12 = 3.2281
```

Monitoring costs 0.08 and returns 0.261 in reduced free-riding loss. **Steady-state stored order nearly doubles — 1.72 to 3.23 — for a monitoring cost of eight hundredths of a unit.** That is an 87.7 per cent *increase* in the steady state and a 68.3 per cent recovery of what free riding had taken; the two figures are different quantities and Section 7.5 records that the source material conflates them. This is why O4 is the one to install first if only one can be installed: it has the best ratio in the list, and it is the one people are most squeamish about because it sounds like surveillance.

It is not surveillance when it satisfies Ostrom's actual condition: **the monitors are accountable to the appropriators, or are the appropriators.** Mutual visibility among peers who share the resource is a different object from observation by an authority that does not. The first raises `eta`; the second raises `C` and lowers `u`, because people watched by an authority converge on the behaviour the authority rewards.

14.4.9 4.9 Where central enforcement is absent, and where formal rules are not enough

Two field observations complete the picture and both are uncomfortable in ways that make them useful.

Private protection. Gambetta's study of Sicily asked what fills the gap when central enforcement of contracts is unreliable, and answered that a market in private protection emerges: a networked structure of obligation, reputation and guarantee that does, in a fashion, produce order. Contracts are honoured, disputes are settled, coordination occurs. *The 1993 monograph has no key in this corpus's bibliography*; the collection Gambetta edited on the making and breaking of cooperative relations [Gambetta, 1988] is the citable neighbour and is cited for the general claim rather than for the Sicilian case.

The lesson is deliberately double-edged. **Order without a centre is not automatically good order.** The mechanism recommended here — put the state in the medium, let obligation and reputation do the enforcement — is the same mechanism, structurally, that produces outcomes nobody would design. The difference is not in the topology. It is entirely in whether the governance features of Section 4.5 are present, and specifically O3, O5 and O6. A centreless system that fails O3, O5 and O6 is not a commons; it is a protection racket with good uptime. Carry this as a standing check: **centrelessness is a topology, not a virtue.**

Formal rules against stored order. Putnam et al. [1993] examined the divergent performance of regional governments in Italy after identical institutional reforms were applied

across the country. The same formal institutions produced markedly different results, and the difference tracked dense horizontal associational life and the norms of reciprocity that go with it. In this framework's terms: **identical formal structure, different `Q`.** The reform supplied a constitution; it did not supply stored order, and stored order cannot be legislated into existence in one step because the storage law integrates — `Q_t = (1-delta)^t Q_0 +` the driven terms, so if `Q_0` is low the transient is low and the driven term needs time at a positive step to accumulate.

Honesty note. Putnam's causal direction has been disputed at length — whether civic association produces good government or good government produces civic association is not settled by the data presented — and this paper does not need it settled. What is used is the narrow and well-supported observation that **formal rules and stored social order are different quantities, and installing the first does not install the second.** A further qualification is owed and is not comfortable for the framework: Putnam's own later finding is that diversity in the short run reduces trust and civic participation [Putnam, 2007], which cuts directly against the corpus's founding preference for raising `u`; the rebuttal literature qualifies but does not dissolve it [Abascal and Baldassarri, 2015], and contact across difference is known to produce value only under conditions rather than automatically [Allport, 1954, Pettigrew and Tropp, 2006].

14.4.10 4.10 The topology argument, and what it does not establish

Albert et al. [2000] measured the robustness of scale-free networks and found an asymmetry: such networks tolerate random node failure well and fail catastrophically under targeted removal of high-degree hubs. Under design postulate M3 a coordinator is a hub — by construction the highest-degree node in the coordination graph. The consequence is not that coordinators are bad. It is that a coordinated system has a **specific and predictable fragility profile**: it handles random absences well and the targeted loss of the coordinator very badly, and no amount of diligence by the coordinator removes it. Diligence makes it worse, because the more effectively the coordinator holds the state, the more of it is lost when they go.

Two engineering consequences follow.

1. **Coordinator loss is an architectural risk, not a personnel risk.** The mitigation is not a better succession plan; it is moving state out of the person and into the medium, so that the hub carries traffic and not storage.
2. **Measure hub concentration directly.** For each critical class of state, record the number of people who could reconstruct it. The distribution of that number across all state is the system's real robustness curve, and a system with many items at bus-count 1 is fragile however good its documentation policy claims to be.

What this argument does not establish, stated plainly. No degree distribution has been measured for any coordination graph in this paper; the scale-free label is borrowed [Barabási and Albert, 1999] and not demonstrated, and the robustness of real engineered networks is known to come substantially from design rather than from a degree distribution [Doyle et al., 2005]. Two further results in the same literature are more directly useful than the hub asymmetry and are used here for what they actually say: coupled systems are more fragile than either component alone, so a nested design whose layers share a failure mode is not nested in the sense that matters [Buldyrev et al., 2010]; and an unaudited count of participants is worthless where one party can present as many [Douceur, 2002], which is the pseudo-diversity problem stated as a theorem and is the network-security twin of the redundancy audit that produces `N_eff`. The count of genuinely independent parties required to tolerate `f` faulty ones is likewise a hard external number [Lamport et al., 1982], and the closest empirical analogue of an effective count in the social sciences measures adoption against the number of *distinct connected components* among a person's contacts [Ugander et al., 2012].

14.4.11 4.11 The honest limits of centrelessness

Four things a shared medium does not do, each stated because the literature on decentralisation routinely omits them.

It does not remove governance; it relocates it. The repository has maintainers and a merge policy. The encyclopedia has arbitration and protected pages. Every enduring commons in Ostrom's sample had rules, monitors and sanctions. A medium without governance is an unmanaged open-access resource, which is the case the tragedy argument correctly describes.

It does not remove cost; it changes its shape. Coordinator cost is concentrated, visible and budgeted. Medium cost is distributed, invisible and unbudgeted — everyone spending some minutes maintaining shared state. Distributed invisible cost is exactly the kind that goes unmeasured and therefore unfunded, which is the infrastructure problem of Section 6.8 [Star and Ruhleder, 1996] arriving early.

It is slower for genuinely novel decisions. Stigmergy excels at incremental work against an existing structure and is poor at the standing start — a decision with no precedent, no existing state to modify, and a deadline. For that class of decision a temporary centre is often correct, and the design question is whether it dissolves afterwards.

It requires a medium that can represent the work. Some work does not deposit a legible trace: a difficult conversation, a judgement call, an intuition about a person. Where the trace cannot be made, the mechanism does not apply, and pretending otherwise produces a medium full of ritual entries that nobody reads. **If the trace cannot be honestly deposited, the medium should not be built for that class of work.**

14.4.12 4.12 Specifications

S4.1 — Medium requirements (VWDLA). A shared medium **shall** satisfy all five properties, each scored 0–4 by the inspection tests of Section 4.4. A score of at least 3 on all five passes; any single score of 0 or 1 is a blocking defect.

#	Property	Requirement	Test	Pass
1	**V**isible	Current state readable by any participant without asking a person	Three-participant sixty-second test	median <= 60 s, 0 asks
2	**W**ritable	<= 2 non-blocking steps from observation to recorded state	Step count	<= 2
3	**D**urable	No critical state at bus-count 1; entries intelligible after 14 days	Bus-count census; fortnight test	0 items at bus-count 1
4	**L**ow friction	One typical write completes end to end in under 30 s	Stopwatch, 5 samples, median	median < 30 s
5	**A**ppend-oriented	State as of any past date reconstructible	Reconstruct a named date last quarter	success

S4.2 — The medium shall display its own unfinished edges. There **shall** exist a standing, visible representation of what is incomplete: open items, failing checks, known gaps, unanswered questions. **A medium that shows only finished states supplies no stimulus and is not stigmergic.** Requirement: the incomplete-edge view is reachable in one step from the medium's default view.

S4.3 — Ostrom conformance checklist. Score each principle 0 (absent) / 1 (partial) / 2 (present). This is the governance audit and it is mandatory before any centreless design is declared complete.

#	Principle	Evidence required	Score
O1	Clearly defined boundaries	A written list of who may withdraw, and a written definition of the resource	/2
O2	Congruence: rules to conditions, appropriation to provision	Rules reference actual local conditions; take is tied to contribution	/2
O3	Collective-choice arrangements	A named, used procedure by which affected users change operational rules	/2
O4	Monitoring	Monitors are the appropriators or accountable to them; monitoring is routine, not exceptional	/2
O5	Graduated sanctions	At least three distinct severity levels, with the mildest actually used in the last year	/2
O6	Conflict-resolution mechanisms	A local arena, resolution time measured in days, no external cost	/2
O7	Right to organise recognised	No external authority has voided a local rule in the review period	/2
O8	Nested enterprises (if part of a larger system)	Each layer's remit written; cell size within the headroom-positive band	/2

Total /16. Below 10, the arrangement is not a governed commons and the tragedy argument applies to it as written.

S4.4 — Cell sizing. Working cells **shall** be sized within the headroom-positive band computed for the system's own `eta` and per-path cost. For the reference parameters of Section 4.6 the band is `n <= 11` with peak at `n = 5`. **Each installation shall recompute the band with its own measured per-path cost rather than adopting these figures.**

S4.5 — Focal-point inventory. Before building any medium, the designer **shall** produce a written inventory of at least five focal points already available to participants (Section 4.3)

and **shall** adopt them as defaults. Every adopted focal point is a `C` reduction obtained at zero build cost.

S4.6 — Hub census. The designer **shall** record, for every class of critical state, the number of people who could reconstruct it from what exists outside their own memory. **Any class with a count of 1 is a blocking defect** and must be remedied by depositing that state in the medium before the design is signed off.

S4.7 — Deposit rule. Every participant-facing procedure **shall** specify, in one sentence, what it deposits in the medium. **A procedure that deposits nothing has `eta = 0` and shall be marked as such in the survey.** This applies to meetings, reviews, inspections, and conversations held for the purpose of alignment.

S4.8 — Declination gate. No mechanism in this section may be implemented in a form that removes a participant's ability to decline. Monitoring under O4 **shall** be mutual and peer-accountable, not unilateral; participation under O3 **shall** be an available right and not a compulsory duty; and withdrawal from the commons **shall** remain possible on stated terms. **A monitoring regime whose subjects cannot opt out of the commons is not Ostrom's Principle 4; it is surveillance, and it is outside the scope of this framework.**

14.4.13 4.13 Failure modes

FM-4.1 — The phantom centre. *Description.* The organisation has declared itself decentralised, removed the coordinator's title, and kept the coordination. State still lives in one person's head; they simply no longer have the authority or the budget that used to come with holding it. *Damages.* `C` unbudgeted, `delta` unchanged and high, `eta` falling because the person now has no mandate to convert. *Diagnostic signature.* The hub census returns a bus-count of 1 for most critical state despite an org chart with no centre. People still ask the same one person; their calendar is the busiest in the group. *Fix.* Either re-establish the centre honestly and budget it, or execute S4.6 properly: enumerate what that person holds and deposit each item in the medium, one at a time, until the census passes. There is no third option.

FM-4.2 — Broadcast mistaken for medium. *Description.* A read-only dashboard, status report or weekly digest is presented as the shared state. Participants can see it; they cannot write to it. *Damages.* `eta` near zero for all receivers; `u` falls, because observations that cannot be deposited are eventually not made. *Diagnostic signature.* Fails S4.1 test 2 outright. The tell is that corrections are sent to a person by message rather than made to the artefact. *Fix.* Make the artefact writable with after-the-fact review. If it genuinely cannot be made writable, stop calling it the shared state, identify what the actual shared state is — usually a chat thread — and apply VWDLA to that instead.

FM-4.3 — Ritual deposit. *Description.* Everyone writes to the medium because it is required. Nobody reads it. The entries are formally complete and informationally empty. *Damages.* `C` as pure dissipation — everyone pays the write cost, nobody collects the read benefit — and `eta` zero despite apparent activity. *Diagnostic signature.* High write volume, near-zero read volume. Ask three participants what changed in the medium last week; if none can answer, the deposits are ritual. Entries uniform in length and phrasing are the fingerprint of a required field being satisfied rather than a state being recorded. *Fix.* Delete required fields until writes become informative. Then make one visible consumer of the medium — a decision, a schedule, a rota — depend on it, so that a wrong deposit produces a visible wrong outcome. **A medium nothing depends on will be filled in wrongly forever and nobody will ever find out.**

FM-4.4 — Free-ride drift. *Description.* Contribution slowly concentrates in a shrinking fraction of participants. Nothing breaks. The contributors become quietly bitter and then leave, at which point something breaks all at once. *Damages.* `C` rises with `phi`, then `Q` takes a step loss when contributors depart. *Diagnostic signature.* Plot contribution share by participant over time and watch for the top-quartile share rising quarter on quarter. **This is measurable before it is felt**, which is the entire reason to measure it. *Fix.* O4 and O5 together: make contribution visible — not ranked, not gamified, merely visible — and install a gradu-

ated response that begins with a private word. Section 4.8 shows monitoring paying for itself many times over. Do not begin with the severe sanction; a severe sanction alone is FM-4.5.

FM-4.5 — Terminal sanction only. *Description.* The only available response to rule-breaking is expulsion or an equivalent. It is therefore never used, the rules become advisory, and everybody knows it. *Damages.* `C` rises because enforcement is impossible so violations accumulate; `Q` falls because rules that are not enforced are not stored order but stored text. *Diagnostic signature.* Ask when the sanction was last applied. If the answer is never, or once years ago and it was very ugly, this is the fault. Count the severity levels: one level is the fault. *Fix.* Build at least three graduated levels and use the mildest deliberately within the first month, so that its use becomes ordinary rather than an event.

FM-4.6 — Medium capture. *Description.* Access controls, mandatory fields, approval workflows and schema requirements accumulate on the medium — each individually justified — until reading requires permission and writing requires an intermediary. The medium has become a coordinator with a database interface. *Damages.* `C` rises, `eta` falls, `u` falls because the schema now dictates what can be said. *Diagnostic signature.* Compare today's V and L scores against the scores at installation. **Medium capture is only visible as a trend, which is why S4.1 must be re-run annually and the scores retained.** A second tell: a shadow medium has appeared — a chat channel or a spreadsheet where the real state now lives. *Fix.* A schema and permission budget: fix the maximum number of required fields and access tiers at installation, and require that adding one removes one. The shadow medium is not the problem, it is the diagnosis — whatever people fled to is telling you which property you lost.

FM-4.7 — Centreless and ungoverned. *Description.* The medium exists, VWDLA passes, and the Ostrom score is 4/16. Boundaries are vague, nobody monitors, there are no sanctions, and disputes go to whoever shouts. *Damages.* All of them, eventually, via free riding and boundary erosion. *Diagnostic signature.* S4.3 total below 10 with a good S4.1 score. This combination is common in new, enthusiastic, technically competent groups, and it is the standard prelude to the collapse later described as needing more structure all along. *Fix.* Install O1, O4, O5 and O6 first — boundaries, monitoring, graduated sanctions, and a cheap local arena. These four cover the bulk of the failure surface; O2, O3 and O8 can follow.

FM-4.8 — Nesting refused. *Description.* The body grows past the headroom-positive band and refuses to split, usually for reasons of unity, culture or fear of fragmentation. *Damages.* `C` quadratic, `h` negative, `u` falling because large bodies suppress dissent by fatigue long before they suppress it by policy. *Diagnostic signature.* Compute `n(n−1)/2 x` per-path cost against `eta ln N`. If `C` exceeds `eta ln N − delta Q`, the sustainability ceiling applies and **no arrangement of the people prevents decline** [Curzi, 2026a]. Observable proxies: meeting length rising with no rise in decisions; a growing tier of people who attend but do not speak; decisions made before the meeting and ratified in it. *Fix.* O8. Split into cells within the band, write each cell's remit, and give the federation layer only what cells genuinely cannot do. Expect the split to be resisted on grounds of culture; the arithmetic does not care about culture.

FM-4.9 — Evaporation refused. *Description.* Nothing in the medium ever expires. The trail to the exhausted food source is maintained in perpetuity, and new participants cannot tell live state from archaeology. *Damages.* `C` rises via search cost, `eta` falls as the medium's signal-to-noise degrades, and eventually `h` falls via the maintenance term of Section 6. *Diagnostic signature.* Ask a participant of under six months' standing to identify the current state of three items and count errors caused by stale entries. **More than one error in three is the fault.** *Fix.* Deliberate `delta` — the pheromone must evaporate. Expiry dates on entries, an archive tier that is retrievable but out of the default view, and a standing review that retires rather than accumulates. This is the shedding operation of Section 6.7 applied at the level of the medium.

14.5 5 Individuality as load-bearing structure

Every familiar argument for preserving individuality in a group is a moral argument, and moral arguments lose to schedule pressure. The manager who wants the meeting over, the committee that wants a clean recommendation, the department that wants alignment — none of them is a villain, and none of them is moved by an appeal to the dignity of the dissenter. What can move them is a demonstration that **the thing they are trying to produce cannot be produced by the method they are using.** That demonstration exists, and it arrives from four independent directions, in four separate literatures, using four unrelated formalisms.

14.5.1 5.1 The arithmetic of the merge

From the extremality result of the surplus functional [Curzi, 2026], `F = 0` if and only if the states are parallel and `F = ln N` if and only if they are orthogonal. Combined with the effective count `N_eff = 1 + (N−1)u`, the consequence is arithmetical and not negotiable. **When the members of a group become parallel, `u` goes to 0, `N_eff` goes to 1, and the surplus of their interaction is `ln 1 = 0` — exactly zero, regardless of how many of them there are, how able they are, or how hard they work.** Nine people who have converged are not nine people. They are one person, counted nine times, at nine times the cost.

This is a theorem about a mathematical object and makes no empirical claim. The empirical content of this section comes from elsewhere; the surplus law's contribution is that it prices the same phenomenon in the same units as everything else in the framework. Here is the price list for a nine-person panel, `N = 9`, `ln N = 2.1972`:

u	N_eff	S = ln N_eff	Fraction of ceiling
0.00	1.00	0.0000	0%
0.10	1.80	0.5878	27%
0.12	1.96	0.6729	31%
0.25	3.00	1.0986	50%
0.50	5.00	1.6094	73%
0.69	6.52	1.8749	85%
0.75	7.00	1.9459	89%
1.00	9.00	2.1972	100%

Two features are worth dwelling on. **First, the collapse at the left end is total, not merely severe.** At `u = 0` the surplus is not small; it is zero, because `ln 1 = 0` exactly. **Second, the recovery is front-loaded.** Going from `u = 0` to `u = 0.25` — a quarter of the way to orthogonality — buys half the available surplus. That is strict concavity arriving as a design fact: the first increment of genuine difference is worth more than all subsequent increments combined, which means the first honest dissenter in a converged room is the single highest-value intervention available in that room.

The same concavity governs headcount. At `u = 0.62`, the marginal surplus of the `N`-th genuinely distinct member:

N	2	3	4	5	6	8	10
S	0.4824	0.8065	1.0508	1.2470	1.4110	1.6752	1.8840
marginal	+0.4824	+0.3240	+0.2443	+0.1962	+0.1640	+0.1234	+0.0990

The second genuinely different person is worth roughly five times the tenth — 0.4824 against 0.0990. Concavity is merciful, in that it means a cast of thousands is not required, and unforgiving in the other direction, because it also means that adding people who are *not* different adds nothing at all while adding coordination cost quadratically (Section 4.6). The

marginal entries are single-step differences and are stated as such; the table's columns skip `N = 7` and `N = 9`, whose surpluses are 1.5518 and 1.7850.

A construct warning that the framework owes and rarely pays. Three incompatible constructs travel under the word diversity — separation, variety, and disparity — and they have different antecedents and different effects [Harrison and Klein, 2007]. **`u` is a variety construct**: it measures non-parallelism of epistemic states and nothing else. It is not separation on an attitudinal continuum and it is not disparity in a valued resource. The empirical literature reporting that surface-level diversity often hurts in the short run [Mannix and Neale, 2005, van Knippenberg and Schippers, 2007] is measuring constructs `u` does not name, which is a defence of `u` only if `u` is actually measured rather than inferred from demography — and Section 5.9 excludes the inference by rule.

14.5.2 5.2 Independence is the load-bearing clause

Condorcet's jury theorem, published in 1785, has a setup worth restating exactly: a binary question with a correct answer; `n` voters, `n` odd; each voter correct with probability `p`, **independently of the others**; the group decides by simple majority. The theorem: if `p > 1/2`, the probability that the majority is correct exceeds `p`, increases monotonically in `n`, and tends to 1 as `n` tends to infinity. The probability that the majority of `n` is correct is

```
P(n, p) = sum over k from (n+1)/2 to n of  C(n,k) p^k (1−p)^(n−k)
```

No citation key exists in this corpus's bibliography for Condorcet's original, and the gap is recorded in Section 10; the aggregation tradition it founded is represented here by Galton [1907] and Surowiecki [2004].

`n`	`p = 0.55`	`p = 0.60`	`p = 0.70`
1	0.5500	0.6000	0.7000
3	0.5748	0.6480	0.7840
5	0.5931	0.6826	0.8369
9	0.6214	0.7334	0.9012
15	0.6535	0.7869	0.9500
25	0.6937	0.8462	0.9825
51	0.7641	0.9265	0.9986
99	0.8413	0.9781	1.0000

Ninety-nine people each barely better than a coin reach 84 per cent as a body. This is the most encouraging table in the sociology of groups and it is why democratic aggregation is a technique and not only a value.

Now the clause that is routinely quoted past: *independently of the others.* Every number in that table is conditional on it. And the theorem has a second half that is almost never taught: **if `p < 1/2` the same monotonicity runs the other way**, and the majority is worse than the individual and gets worse with `n`.

`n`	1	3	9	25	99
`P` at `p = 0.45`	0.4500	0.4253	0.3786	0.3063	0.1587

Ninety-nine people each 45 per cent accurate produce a body that is **16 per cent accurate.** Aggregation is an amplifier: it amplifies whatever sign the individual competence has. And at exactly `p = 0.5`, `P(n, 0.5) = 0.5` for every `n` — no amount of aggregation extracts signal from noise.

What correlation does

Independence is an idealisation. Real voters share training, share sources, share incentives, and talk to each other. Model this with the simplest possible correlation structure, **badged as a modelling choice and not as a result**:

> *Model C1 (shared-signal mixture).* With probability `rho`, all voters observe and follow a single common signal, itself correct with probability `p`. With probability `1 − rho` they vote independently at accuracy `p`. Then `P(n, p, rho) = rho p + (1 − rho) P_independent(n, p)`.

This is crude — a two-point mixture, not a general correlation model — but it is transparent and it is sufficient. With `n = 9` and `p = 0.60`, so that `P_independent = 0.7334`:

`rho`	0.0	0.2	0.4	0.6	0.8	1.0
`P`	0.7334	0.7067	0.6801	0.6534	0.6267	0.6000

At `rho = 1` the nine-person panel performs exactly as well as one person, because it *is* one person. That is the extremality result again, in probability rather than geometry: **`rho -> 1` is `u -> 0` is `N_eff -> 1` is `S -> 0`.** And scale does not rescue it: ninety-nine voters at `p = 0.60` with `rho = 0.30` give 0.8646 against 0.9781 under independence. **Correlation is not a discount on the benefit of scale; it is a ceiling on it.** The measured version of the same effect is that social influence demonstrably narrows the range of estimates in a crowd while leaving confidence intact, so the group becomes both less accurate and more sure [Lorenz et al., 2011].

Honesty note on the bridge. `rho` and `u` are measures of the same underlying thing — non-redundancy — in different formalisms, and no general identity between them is claimed. Converting one into the other requires a specific state encoder [Curzi, 2026i] and the conversion should be stated whenever it is made. What is claimed is the direction: **as correlation rises, effective count falls, by both accountings.**

14.5.3 5.3 Diversity against ability, and the critique of that result

Hong and Page [2004] model problem solvers searching a solution landscape using *heuristics* — internal representations plus rules for where to look next. A solver gets stuck at a point that is a local optimum *for that solver's heuristic*; a group works by passing the current best solution around until no member can improve it. Their result, informally: under stated conditions, a group of randomly selected problem solvers drawn from a large pool of able solvers outperforms a group formed by taking the individually highest-performing solvers.

The mechanism is entirely legible in this framework's terms. **A solver's ability determines how good their local optimum is. A solver's difference determines whether they are stuck at the same place as everyone else.** A group of the individually best tends to share heuristics — that is frequently *why* they are the best on the measured task — and a group that shares heuristics shares local optima. When such a group gets stuck, everyone is stuck at the same point, and there is nobody whose different representation makes the next move visible.

The conditions matter and are stated rather than elided: the problem is difficult, so no solver alone finds the global optimum; the solvers are individually capable, so their local optima are not trivially bad; the pool is large; and the group is large enough that random selection produces genuine heuristic variety.

Honesty note, and it is a substantial one. The theorem has been the subject of mathematical criticism [Thompson, 2014] arguing that the formal result depends heavily on the specific construction of diversity and does not support the broad interpretation often placed on it. The subsequent exchange is genuinely two-sided [Kuehn, 2017, Singer, 2019, Grim et al., 2019], with Singer [2019] offering a repair that survives the objection and Grim et al. [2019]

mapping the boundary conditions by simulation. **This paper's position is narrow and, it believes, untouched by the critique:** the mechanism — that shared heuristics produce shared local optima, and that a group stuck in one place needs someone standing somewhere else — does not require the strong version of the theorem. It requires only that different representations get stuck in different places, which is nearly a definition. Where this paper uses the result, it uses the mechanism, and the formal claim is left to the primary literature and its critics.

14.5.4 5.4 Cascades: how copying destroys the information it appears to use

This is the most important result in the section, because it explains why `u` collapses in groups that nobody is pressuring, composed of people who are individually behaving well.

Bikhchandani et al. [1992] analysed sequential decision-making under observation: agents decide one at a time, each has a private signal about which of two options is correct, and each can observe the *actions* of everyone before them but not their signals. Banerjee [1992] reached the same structure independently. Their finding: after a very small number of decisions, agents rationally stop using their private signals. An **informational cascade** forms — everyone thereafter copies the visible pattern regardless of what they privately know — and **from that moment on, no further private information enters the public pool.** The chain looks informative and is not: it contains, permanently, the information of the first two or three agents and nothing more.

The arithmetic

Take the standard symmetric setup. The true state is HIGH, so ADOPT is correct; each agent receives a private signal correct with probability `p > 1/2`; agents act in sequence; ties are broken by a fair coin.

- **Agent 1** has only their own signal, so adopts with probability `p`.
- **Agent 2** observes agent 1's action and their own signal. If the signals agree they follow; if they disagree the posterior is exactly balanced and they flip a coin. So given agent 1 adopted, agent 2 adopts with probability `p + (1−p)/2 = (1+p)/2`.
- **Agent 3** observes two actions. If both were the same, a single private signal cannot overturn two, so agent 3 copies regardless of their own signal — **a cascade.** If the first two actions differed they cancel, and agent 3 is in agent 1's position: the process restarts.

The probabilities are

```
P(up cascade starting)   = p (1+p) / 2
P(down cascade starting) = (1−p)(2−p) / 2
P(restart)               = 1 − [p(1+p)/2 + (1−p)(2−p)/2] = p(1−p)
```

and these sum to 1 identically, which is the first check a reader can run. Since a cascade eventually forms with probability 1, the probability that the cascade that forms is the *correct* one is the ratio of the up-cascade probability to the total cascade probability:

```
P(correct cascade) = [ p(1+p)/2 ] / [ p^2 − p + 1 ]
```

p	0.55	0.60	0.62	0.70	0.80	0.90
P(correct cascade)	0.5664	0.6316	0.6570	0.7532	0.8571	0.9396

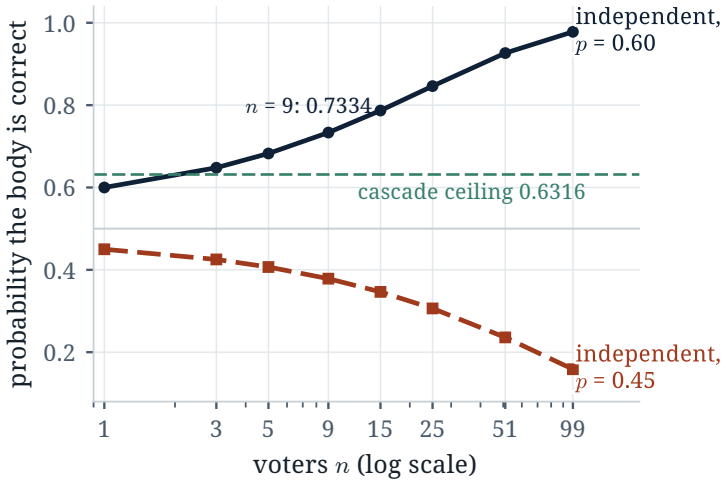


Figure 14.2: **Independence, not scale, is what aggregation runs on**: at `p` = 0.60 an unbounded informational cascade tops out at 0.6316 — below nine independent voters at 0.7334 and far below ninety-nine at 0.9781 — while at `p` = 0.45 the same majority rule drives accuracy down to 0.1587 at `n` = 99. Derived: exact binomial sums and the cascade formula of Section 5.4, both recomputable from the expressions printed there; generated by `figs/src/CCP-014_fig05.py`.

Verification. A Monte Carlo simulation of this process at `p = 0.60` over 400,000 trials returned 0.6323 against the analytic 0.6316 — within one standard error of the simulation (0.00076), which is what agreement looks like at that sample size. The simulation is reported as the source states it; no re-run was performed for this paper, and Section 10 records that.

The comparison, and the instruction it yields

At `p = 0.60`:

Arrangement	Accuracy
One person alone	0.6000
An unbounded cascade — a million observers	**0.6316**
9 independent voters, majority rule	0.7334
99 independent voters, majority rule	0.9781

An unbounded queue of rational, well-informed, honest people, each holding genuine private information, performs worse than nine of them who never looked at each other. The million observers contributed nothing. Their information exists; it simply never entered the pool, because each of them rationally deferred to a public record containing only the first two people's signals.

There is no villain in this story. Nobody is lazy, conformist or cowardly; every agent behaves optimally given what they can see. **The failure is structural: the observation of others' actions before the expression of one's own signal is sufficient to destroy the aggregation, all by itself.** The experimental demonstration that early visible choices lock in outcomes that were not determined by quality is Salganik et al. [2006], where parallel worlds seeded identically diverged and stayed diverged.

That is the design brief. If observation-before-expression destroys aggregation, then the fix is to **move expression before observation.** Everything in Section 5.7 is a variation on that single move.

14.5.5 5.5 The distribution of participation, not its peak

Woolley et al. [2010] had groups perform a battery of varied tasks and found that performance across the battery loaded onto a single statistical factor, `c`, analogous to the individual `g`. The findings that matter here, stated precisely: `c` was **not significantly predicted by average member intelligence, nor by the maximum**; `c` correlated with average social sensitivity of members, measured by a standard instrument whose construct validity has itself been examined [Engel et al., 2014]; `c` correlated with **equality in the distribution of conversational turn-taking**, groups in which a few members dominated having lower `c`; and `c` correlated with the proportion of women in the group, an effect largely mediated by social sensitivity [Bear and Woolley, 2011].

The turn-taking finding is the one this section builds on. Consider what unequal turn-taking *is* in the terms of Section 5.1. A group in which two people speak 80 per cent of the time is a group in which seven people's states are not entering the interaction at all. The nominal `N` is nine; the participating `N` is two. **Turn-taking inequality is `N_eff` suppression by airtime**, and it requires nobody to be silenced — the ordinary dynamics of a fast, unstructured conversation are sufficient.

Honest limits. Woolley's `c` was measured in short-task laboratory settings with ad hoc groups. Extrapolation to standing institutional bodies is a **modelling assumption** and is badged as such wherever it is used below. What transfers with confidence is the narrower claim: within the studied setting, the distribution of participation predicted performance and the peak of member ability did not. The general caution owed to any single psychological finding applies here too [Open Science Collaboration, 2015, Ioannidis, 2005].

14.5.6 5.6 Four results, one statement

Lay them side by side.

```
Condorcet    : aggregation works only if the inputs are independent
Hong & Page  : difference contributes what ability cannot
BHW 1992     : copying stops information from entering the pool
Woolley 2010 : equal participation predicts collective performance
```

Each is a statement about a different thing — probability, search, sequential inference, factor structure. Each was produced by people not working on the others' problems. And each reduces to the same instruction:

> **Keep the individuals individual, and keep them all in.**

The surplus law prices it. `S = ln(1 + (N−1)u)`. Merge the members and `u -> 0`, `N_eff -> 1`, `S -> 0`. Silence the members and the participating `N` falls, with the same effect. **Merge the individuals and the collective-intelligence factor falls, the cascade sets in, and the group becomes one mind with one blind spot** — a blind spot that is now nine times as confident as it was, because nine people agree.

The confidence is the cruellest part. A converged group does not experience itself as impoverished; it experiences itself as aligned, and alignment feels exactly like being right. There is no internal signal distinguishing a room that has converged on the truth from a room that has converged on the first speaker. **Only a protocol can distinguish them, and only if the protocol ran before the room started talking.**

14.5.7 5.7 The five independence protocols

Five, labelled IP-1 to IP-5. Each is specified with its mechanism, its cost, its expected effect, and — mandatorily — its decline path.

IP-1 · Independent estimate before discussion. *Mechanism.* Before any discussion of the item, every participant privately records their own judgement — an estimate, a rating, a recommendation, a probability — with a one-line reason. The records are sealed until all are in. Only then does discussion open, and the distribution of independent judgements is displayed *first*. *Why it works.* It is the direct structural answer to Section 5.4. The cascade requires observation before expression; this protocol inverts the order, so that every private signal enters the pool before any public signal exists to override it. **It is not a debiasing technique and does not depend on anyone being unbiased. It changes the order of operations.** *Cost.* Four to eight minutes per item; a `C` increment computed at 0.030 in Section 7.2. *Effect.* The distribution itself is the deliverable and is worth as much as the decision. A tight distribution before discussion is either genuine convergence or a symptom that convergence already happened elsewhere, and the two are distinguishable by asking whether the reasons given differ. **Identical conclusions with identical reasons is a `u = 0` reading. Identical conclusions with different reasons is the strongest evidence a committee can generate.** *Decline path.* A participant may record "no independent view" and the item proceeds. The record of a declined estimate is itself information: a member who declines on every item in their nominal area of expertise is reporting something about the panel that no survey will.

IP-2 · Anonymous-first elicitation. *Mechanism.* The first round of judgements is collected without attribution. Attribution is added in the second round, after the anonymous distribution is visible, and only for participants who choose to attach their names. *Why it works.* It removes rank from the first signal. In a group with any status gradient the junior member's private signal is systematically discounted by its own holder before it is spoken — a cascade driven by hierarchy rather than by sequence, with the same effect on the pool. *Cost.* A slip of paper and a box, or a form: effectively zero in money, three minutes in time; 0.010. *Effect.* The instrument to watch is the **anonymous-to-attributed drift**, the difference between the anonymous distribution and the attributed one. A large drift is a direct measurement of how much signal the group's status structure was suppressing, and it is one of the few `u`-suppression readings obtainable without a control group. *Decline path.* Participation in the anonymous round is optional and the option is stated aloud each time, not merely in the written procedure. Anonymity that participants do not believe in is worse than no anonymity, because it produces confident false readings. *Constraint.* Anonymity must be real. If the group is small enough that handwriting, phrasing or expertise identifies the author, the protocol is not anonymous and must not be described as such; say so and use IP-1 instead.

IP-3 · Dissent protection. *Mechanism.* A written, standing, published commitment that (a) a recorded dissent carries no consequence for standing, assignment or advancement; (b) dissents are entered in the record with the decision, not appended afterwards; and (c) the person who dissents is not the person then assigned to implement the majority position. *Why it works.* Under postulate I1, `u` decays. The largest driver of decay in an institutional body is not persuasion but anticipation of cost: a member who expects a recorded disagreement to be remembered at review time will not record it, and their private signal never enters the pool. **The protection is a retention intervention on `u`: it does not create difference, it stops the leak.** *Cost.* Near zero to write; substantial to honour, and the honouring is the whole protocol; 0.005. **A dissent protection that has never been tested is not known to exist.** *Effect.* The measurable quantity is the **dissent rate over time**. A body whose recorded dissent rate declines quarter on quarter is a body whose `u` is collapsing, and the decline is visible in the minutes long before it is visible in the decisions. *Decline path.* Nobody may be compelled to dissent, and nobody may be assigned the role of standing contrarian without explicit and revocable consent. Conscripted dissent is theatre; it produces a visible objection carrying no private signal, which is worse than silence because it looks like the protocol is working.

IP-4 · Minority report requirement. *Mechanism.* Any decision reached with less than unanimity must carry a written minority report: the strongest available statement of the rejected position, its evidence, and **the specific observation that would show the majority to be wrong.** The report travels with the decision permanently. Where a decision is genuinely unanimous, the requirement is met by an appointed **steelman** — a member tasked with writing the best case against, with that member's consent. *Why it works.* Three effects, of which the third is

the largest. (i) It preserves the minority's information in the durable record: dissent that is written down has `eta > 0`, dissent that is spoken and lost has `eta = 0`, so **a minority report is a conversion device applied to the most perishable and most valuable class of interaction in the system.** (ii) It makes the decision auditable later, when the world has supplied an answer. (iii) **It converts unanimity from an achievement into an obligation.** A committee that must produce a written minority report cannot use agreement as an exit, which removes most of the incentive gradient driving premature convergence, because the cheapest path is no longer for everyone to agree quickly. *Cost.* Twenty to forty minutes per non-unanimous decision: real, and the largest cost item in the list at 0.040. *Effect.* The falsification clause in (i) is the part usually dropped and the part that matters. A minority report saying "we disagree" is a mood; one saying "if the programme has not produced measurable X by month 18, the majority was wrong" is an instrument, and in eighteen months it will have paid for every minute it cost. This is the committee-scale form of the discipline that preregistration imposes on a study [Nosek et al., 2018, Platt, 1964]. *Decline path.* The steelman role is by consent and is rotated. If nobody consents, the record states that no steelman was appointed, which is itself the finding.

IP-5 · Turn-taking equalisation. *Mechanism.* Structural devices that flatten the airtime distribution without anyone policing it: round-robin on each item; the **last-speaks-first inversion**, in which the most senior person in the room speaks last; a hard first-round time box per speaker; and a table with no head. *Why it works.* It directly targets the turn-taking finding of Section 5.5 and simultaneously breaks rank-driven cascades: when the senior member speaks first, every subsequent contribution is a response to a public signal rather than an expression of a private one. *Cost.* Near zero; 0.005. **Furniture and speaking order are the cheapest governance instruments available, and they enforce continuously without anyone administering them.** *Effect.* Measure the Gini coefficient of speaking time, or more simply the share held by the top two speakers; the collective-intelligence finding predicts that flattening this distribution raises `c`. A physical precondition applies first: in a reverberant room the loudest voice wins by acoustics regardless of protocol, so the acoustic fix precedes the procedural one. *Decline path.* A participant in a round-robin may pass. Passing is recorded as passing, not as assent, and the distinction must be explicit in the record. **A protocol that reads silence as agreement has re-created the cascade inside the mechanism designed to prevent it.**

14.5.8 5.8 What the protocols cost and what they buy

Costs are in the normalised units of Section 7.2, per panel item.

Protocol	`C` per item	Primarily raises	Primarily protects
IP-1 independent estimate	0.030	`u`	against sequential cascade
IP-2 anonymous-first	0.010	`u`	against rank cascade
IP-3 dissent protection	0.005	`u` (lowers its decay)	against anticipatory self-censorship
IP-4 minority report	0.040	`eta`	against loss of the rejected case
IP-5 turn-taking	0.005	participating `N`	against airtime concentration
Total	**0.090**		

Nine hundredths of a unit. Section 7.2 shows this purchase moving `u` from 0.12 to 0.71 and taking a panel from a negative steady state to a positive one. **The protocols are cheap because they are procedural. They are effective because the failure they address is procedural.**

Note the split in the third column. IP-1, IP-2 and IP-3 raise `u` — they get more distinct signal into the room. IP-4 raises `eta` — it gets the signal that arrived into the durable record. **These are different parameters and neither substitutes for the other.** A panel with excellent elicitation and no minority report generates real difference every week and stores none of it. A panel with rigorous minority reports and no independent elicitation writes down, with great care, the objections of a group that has already converged.

14.5.9 5.9 The scope constraint, stated as engineering

The framework does not recommend any intervention that removes a person's ability to decline. It is worth being explicit about what this excludes, because the exclusions are not obvious and each has been proposed somewhere in good faith.

Excluded: compulsory disclosure of private judgement. IP-1 is powerful precisely because it captures what people actually think, and that power is exactly why it must not be mandatory. A required private estimate is not private, and a person who fears the record will write what is safe — which returns the body to the cascade with an extra step and a false sense of rigour.

Excluded: assigned permanent dissenters. Appointing someone the standing contrarian both conscripts them and defuses them: their objection is discounted by everyone, including themselves, as role-play. The steelman under IP-4 avoids this by being consensual, rotating, and about the argument rather than about the person.

Excluded: identity as instrument. Using a person's demographic or biographical characteristics as a proxy for their epistemic contribution is both an ethical failure and a measurement error: demographic difference without epistemic difference gives `u = 0` and buys nothing, which is precisely the separation-versus-variety confusion [Harrison and Klein, 2007]. **The quantity in the law is the non-parallelism of their states, and the only way to know it is to measure it.**

Excluded: mandatory participation. Every protocol above has a decline path, and it must be usable without explanation.

Included, and required: the right to be recorded as declining. This is the positive form of the same principle. A person who does not wish to give a view must be able to have that fact entered rather than being counted as assent. The alternative — silence read as agreement — is not a neutral default. It is a mechanism that manufactures unanimity out of absence, and a body that uses it will report high agreement and low `u`, which is the exact signature of the failure this section exists to prevent.

14.5.10 5.10 Specifications

S5.1 — Order of operations. Any deliberative body **shall** collect independent individual judgements on an item **before** opening discussion of that item. **Requirement: no participant observes another participant's position on the item before recording their own.** This is the single highest-value requirement in the section and the direct structural answer to the cascade result.

S5.2 — Distribution before conclusion. The distribution of independent judgements **shall** be displayed to the body before discussion begins, including its spread. **A body that sees only the mean has been shown a converged number and will converge on it.**

S5.3 — Reasons, not only conclusions. Each independent judgement **shall** carry at least one line of reasoning. Agreement of conclusions with agreement of reasons **shall** be recorded as a `u`-suppression indicator, not as corroboration.

S5.4 — Dissent record. Every non-unanimous decision **shall** carry a written minority report containing the rejected position, its evidence, and **a specific falsifier** — the observation that would show the majority wrong. Unanimous decisions **shall** carry a consented steelman statement, or an explicit note that none was appointed.

S5.5 — Dissent protection, published and tested. A written dissent protection **shall** exist, be published, and be exercised at least once per year. **An untested protection shall be recorded in the survey as absent.**

S5.6 — Participation measurement. Airtime **shall** be measured on a sample of sessions, at least quarterly. Recorded metric: share of speaking time held by the top two speakers. Threshold for intervention: **above 60 per cent in a group of eight or more.**

S5.7 — u maintenance schedule. Because u decays under postulate I1, the body **shall** remeasure u at least annually by the disagreement-and-error-decorrelation protocol of Curzi [2026i], and **shall** record the trend. A declining trend across two consecutive measurements **shall** trigger a composition review.

S5.8 — Decline paths, mandatory. Every protocol in this section **shall** be implemented with an explicit decline path, stated aloud at the time of use and not only in the written procedure. **Silence shall never be recorded as assent; a pass shall be recorded as a pass.** A protocol without a decline path is outside the scope of this framework and shall not be installed under its authority.

S5.9 — Anonymity truthfulness. Where anonymity is claimed, the designer **shall** verify that the group size and the response format do not identify authors. Where they do, the protocol **shall** be relabelled and the anonymity claim withdrawn.

14.5.11 5.11 Failure modes

FM-5.1 — The cascade collapse. *Description.* Discussion opens with a senior or confident member stating a position; everything subsequent is a response to that position. The panel reaches high agreement quickly and reports it as confidence. *Damages.* u to near zero within a single session, and through it S and Q. *Diagnostic signature.* Order-of-speaking correlates with the final position; time-to-consensus is short and falling; members can state the conclusion but cannot state a reason differing from the first speaker's. *Fix.* S5.1. There is no soft version of this fix: independent recording before discussion, or the cascade continues. Debiasing training does not address it, because the agents in the cascade model are not biased.

FM-5.2 — Unanimity reported as strength. *Description.* The body's minutes show near-total agreement across all items and this is presented as evidence of quality. *Damages.* u, and the credibility of the whole record. *Diagnostic signature.* Dissent rate below 5 per cent of decisions over a year. **Cross-check against outcomes: a body that agrees on everything and is wrong at the base rate is a body whose agreement carries no information.** *Fix.* S5.4 — the minority report requirement converts unanimity from an exit into an obligation. Expect resistance framed as efficiency; the arithmetic of Section 5.4 is the reply.

FM-5.3 — The token dissenter. *Description.* One member reliably objects; everyone knows they will object; the objection is heard, noted, discounted, and the decision proceeds unchanged. *Damages.* u in effect though not on paper — the record shows dissent while the pool receives none. *Diagnostic signature.* The same name in every minority position; decisions never change following a dissent; other members' dissent rate is zero. *Fix.* Rotate the steelman role so that objection is a function rather than a person. Measure **decision change rate following dissent**; if it is zero over a year, the dissent channel is decorative and the body should stop claiming it.

FM-5.4 — Anonymity theatre. *Description.* The first round is nominally anonymous in a group of six, where everyone can identify every author from phrasing and expertise. *Damages.* u, plus a false measurement — the group now believes it has an anonymous baseline and draws conclusions from it. *Diagnostic signature.* Ask three members to attribute the anonymous responses. **If they are right more than half the time, the protocol is not anonymous.** *Fix.* S5.9. Withdraw the anonymity claim and use IP-1's sealed independent estimates instead, which do not depend on anonymity for their effect.

FM-5.5 — Silence read as assent. *Description.* The chair asks whether there are any objections, hears none, and records consensus. Non-response is tallied as agreement. *Damages.*

`u` and `N_eff` — participants who did not speak are counted as if they had. *Diagnostic signature.* Minutes record consensus for items where airtime measurement shows several members never spoke. **This is detectable purely from records, retrospectively, and almost nobody checks.** *Fix.* S5.8. Record passes as passes; replace the negative call with a positive round.

FM-5.6 — The correlated panel. *Description.* Nine reviewers, all trained in the same three institutions, reading the same journals, funded by the same body. Their agreement is read as convergence on truth. *Damages.* `u` structurally, from the moment of appointment. *Diagnostic signature.* Measure error decorrelation, not opinion diversity: give the panel a set of items with known outcomes and check whether their *errors* are independent. **Correlated errors are the definitive reading; correlated opinions are only suggestive.** Then compute `N_eff` and compare with headcount. *Fix.* This is a composition problem and no protocol repairs it — the protocols of Section 5.7 extract the signal present, and if the signals are duplicates there is nothing to extract. Recruitment across training lineages, incentive structures and trades is the only remedy.

FM-5.7 — Protocol without storage. *Description.* Excellent elicitation. Real independent estimates, honest disagreement, a genuine distribution. None of it written down beyond the final number. *Damages.* `eta` — the surplus was generated and not converted. *Diagnostic signature.* Ask for last quarter's pre-discussion distributions. If they do not exist, `eta` for this body's most valuable output is zero. *Fix.* S5.2 and S5.4. Retain the distributions, the reasons and the falsifiers; the storage cost is trivial and the audit value compounds.

FM-5.8 — Conscripted difference. *Description.* A member is assigned to represent a perspective — demographic, disciplinary, regional — that they did not volunteer to represent and may not hold. *Damages.* `u`, because the represented perspective is performed rather than held; and it breaches the declination constraint. *Diagnostic signature.* The member's recorded positions do not differ from the majority except on the items where they are expected to differ. *Fix.* Withdraw the assignment. Recruit for genuine difference and let people speak for themselves, which is both the ethical and the measurable option. **Difference that has to be assigned was not present, and assigning it does not create it — it only makes it unmeasurable.**

14.6 6 Institutions as storage: the maintenance term

An institution, in this framework, is not a building, a charter, a payroll or a mission. It is a **device for holding `Q` against `delta`** — a mechanism whose function is to make stored order survive the departure of the people who made it. That is the whole definition, and every recognisable institutional form is a special case of it.

Form	What it stores	Against what `delta`
Constitution	Decision procedures and constraints on power	The turnover of every officeholder
Archive	Records of what happened and what was decided	Human memory, at roughly 100% per generation
Curriculum	The transmissible portion of a body of skill	Death of practitioners
Profession	Standards of competence and conduct	Individual variation and market pressure
Standards body	Interface agreements between independent parties	Divergent local optimisation
Maintenance regime	The physical plant and its working order	Wear, corrosion, degradation

Each is `Q` in a different substrate; each obeys the same recursion with different parameters. **That is a claim about model form and not about physics**, badged here as it is everywhere it is

made [Curzi, 2026a].

14.6.1 6.1 Why institutions ossify — the equation, and the correction it forces

Read the headroom gauge with the eyes of an institution that has been successful for sixty years.

```
h_t = eta ln N  −  delta Q_t  −  C_t
      ^^^^^^^^     ^^^^^^^^^     ^^^
      ceiling      maintenance   operating
      fixed        grows with    cost
                   success
```

The first term is a **hard ceiling**. `eta ln N` cannot be raised without limit: `ln N` grows logarithmically, so doubling the number of genuinely distinct categories adds `ln 2 = 0.693` and nothing more, whether the move is from 4 to 8 or from 400 to 800; and `eta` is bounded above by 1 in any sensible normalisation. **The ceiling is essentially fixed for a mature institution.** The second term is `delta Q_t`, and `Q_t` is the institution's accumulated achievement: every rule that worked, every precedent that settled a dispute, every procedure that prevented a past disaster, every standard that made two systems interoperate. All of it is `Q`, and all of it requires maintenance at rate `delta` per period, forever. **Therefore, as an institution succeeds, its own stored order consumes an increasing share of a fixed budget.** No malice, no decadence, no failure of nerve. Arithmetic.

> **Ossification is not moral failure. It is a headroom equation.**

Diagnostically this matters, because the standard response to ossification — replace the leadership, exhort the staff, run a culture programme — addresses a cause that is not operating. **If `delta Q_t + C_t` exceeds `eta ln N`, the sustainability ceiling applies and no arrangement of the people prevents decline** [Curzi, 2026a]. Not a better chief executive, not a more committed workforce, not a clearer strategy. The constraint is on the budget and must be relieved on the budget. The general form of the diagnosis is old: the accumulation of a legible order can destroy the local variety that made the system work in the first place [Scott, 1998].

And here the identity of Section 3.2 forces a correction on that account, which should be read carefully, because the account is the one this section is named for. Written with the surplus law substituted in, the identity is `h_inf = eta(ln N − F(u*))`, and the striking feature is what has vanished: **`delta`, `C` and `Q` all cancel at the fixed point.** So `delta Q_t` does grow with success — but under the plain linear law it grows toward a fixed point at which headroom is a *positive constant* that does not depend on `delta`, `C` or `Q` at all. **The plain linear model, left alone, never ossifies.** Maintenance load rises and then stops rising, because `Q` stops rising.

Ossification is therefore not a consequence of accumulation as such. It requires one of two further things, and a diagnosis should say which:

1. **Exogenous `Q` accretion** — obligations added from outside the loop faster than the loop's own fixed point would settle them: statute, merger, inherited estate, external mandate. `Q` is being pushed past its equilibrium by something that is not `eta S`.
2. **Superlinear cost** — `C` rising faster than linearly in `Q`, which breaks the cancellation. A correction load proportional to a rule stock is one instance; the `n(n−1)/2` coordination term of Section 4.6 is another.

The claim stands — but the equation names two specific mechanisms, and neither of them is age. Section 6.6's three routes are exactly the ways a real institution leaves the linear case.

14.6.2 6.2 Two identities that structure the section

Identity 1 — the uniform bound is the zero-headroom line. The uniform bound gives `Q_t <= max{Q_0, eta ln N / delta}` [Curzi, 2026a]. Setting `C = 0` in the headroom gauge and solving `h = 0` gives `Q = eta ln N / delta`. **The uniform bound on stored order is precisely the level at which headroom reaches zero in a costless system.** These are not two facts but one fact stated twice: an institution cannot exceed its storage ceiling because at the ceiling it has no budget left to add anything. Every unit of `Q` it holds is a unit of budget it has committed to holding. For a mature body with `eta ln N = 0.8349` and `delta = 0.015`, the bound is `Q <= 55.66` — the absolute maximum stored order the body could hold with no operating cost whatsoever. With any real `C` the practical ceiling is far lower, and Section 7.3 computes it at 13.73 for the same body.

Identity 2 — headroom at equilibrium is unused variety, derived in Section 3.2. Its consequences for institutions specifically: a settled institution's headroom equals the difference it is failing to use, and nothing else; and **a settled institution never has negative headroom**, so negative headroom is always a disequilibrium condition. **An institution in the red on headroom is always an institution living in a changed world with an inherited balance sheet.**

14.6.3 6.3 What the gauge does not contain, and the four levers

Look at the gauge again and note an absence that decides most institutional strategy:

```
h_t = eta ln N − delta Q_t − C_t
```

`u` does not appear. Non-parallelism enters the storage law only through `S_t = ln(1 + (N−1)U_t)` — the *input* term. It is not in the budget. The consequence is a hard rule contradicting a great deal of received practice:

> **Raising `u` does not raise headroom. It raises the input that fills headroom.**

An institution with `h < 0` that responds by recruiting for difference has bought a larger tap for a tank leaking faster than any tap can fill it. `S` can rise all the way to `ln N` and the step `eta S − delta Q − C` remains negative, because `eta ln N − delta Q − C` — which is exactly `h` — is negative by assumption. **Difference is the right move at `h > 0` and the wrong move at `h < 0`, and the gauge says which regime the body is in before anything is spent.**

The four levers that *do* move `h`, in order of their usual cost:

1. **Raise `eta`.** Usually cheapest. Capture what already happens.
2. **Cut `C`.** Usually available. Remove coordination overhead, friction, unresolved conflict.
3. **Raise `N`.** Expensive and logarithmically rewarded. Rarely the answer.
4. **Shed `Q`.** The one nobody considers, the subject of Section 6.7, and frequently the only one with enough range at institutional scale.

14.6.4 6.4 The fifth lever, and why it is not fifth by accident

The lever missing from that list is the one every institution reaches for first: **lower `delta`.** Build trust, build institutional memory, stop losing what is known. It is absent from the four because it behaves differently from all of them, and the difference is a correction to standard practice rather than a technicality.

Trust does act on `delta`; the mechanism by which it does so is retention, and it enters the law as a modification of the decay term [Curzi, 2026s]. But `delta` sits in the *denominator* of the steady state, and **when the numerator is negative, dividing by a smaller positive denominator makes the result more negative.** The worked separation of Section 4.7 gives the arithmetic: baseline `−0.0836`, trust alone `−0.2114`, reciprocity alone `+0.7394`, both `+1.8701`. **Trust does not add order. It amplifies whichever sign the system already has**, and a failing institution that

invests in trust first buys a slower, better-mannered decline — together with survey results reporting that the programme worked.

The design rule is therefore an ordering: fix the numerator first. Raise `eta`, or cut `C`, until `eta F(u*) > c*`; only then spend on trust, where the same division now multiplies a positive number. Note also what a lower `delta` does to the diagnosis: it lengthens the time constant `tau = 1/delta`, so the institution takes longer to find out. Section 6.10 records the same property working in the other direction, where the `delta` that slows the loss also slows the involuntary return of headroom.

14.6.5 6.5 The maintenance term at institutional scale

Institutions are distinguished from teams and rooms by one parameter above all: **their `delta` is small and their `Q` is large.** That is the point of them; a constitution with `delta = 0.3` per year is not a constitution. But `delta Q` is a product, and a small number times a large number is not small.

Take a fixed ceiling `eta ln N = 1.50` and ask what fraction of the annual budget the maintenance term alone consumes:

Q	delta = 0.005	delta = 0.01	delta = 0.02	delta = 0.03	delta = 0.05
10	3%	7%	13%	20%	33%
25	8%	17%	33%	50%	83%
50	17%	33%	67%	**100%**	**167%**
75	25%	50%	**100%**	**150%**	**250%**
100	33%	**67%**	**133%**	**200%**	**333%**

Every cell at or above 100 per cent is an institution that cannot fund its own maintenance out of its own surplus, before paying a single unit of operating cost. Reading across the `delta = 0.02` row: at `Q = 10` maintenance is a rounding error; at `Q = 50` it is two-thirds of the budget; at `Q = 75` it is the entire budget. This is result **R6**.

The same institution, doing nothing differently, moves from the first row to the last row simply by continuing to work. This is the maintenance-versus-growth transition, and every institution that survives long enough reaches it. It is not a phase to be avoided but a phase to be *recognised*, because the correct behaviour on either side of it differs and the behaviour that got the institution to the transition is wrong after it.

	Before the transition	After the transition
Dominant term	`eta ln N`	`delta Q_t`
Correct posture	Build. Accumulate. Say yes.	Maintain. Curate. Shed.
Success metric	New things established	Old things still working
Typical error	Under-building from caution	Over-building from habit
What “more resources” buys	More `Q`	A slower decline, briefly
Career incentive	Aligned with the correct posture	**Opposed to it**

The last row is where institutions actually die. **After the transition, the correct organisational behaviour is maintenance and shedding, and no individual career has ever been advanced by maintenance or shedding.** The incentive gradient and the headroom gradient point in opposite directions, and the incentive gradient is the one people can feel.

14.6.6 6.6 The three routes to negative headroom

Since a steady state cannot have `h < 0` by Identity 2, negative headroom always means the parameters changed under an inherited `Q`. There are exactly three routes and they are diagnosable apart.

Route A — `eta` falls. Interaction stops converting into durable structure. The meetings still happen; the deliberation is as long as ever; nothing lands. *Mechanisms.* Ritualisation, where the procedure survives its purpose and is performed rather than used. External override, where a higher authority can void local decisions so that deliberation produces nothing durable — which is the failure of O7. Legal review as a filter that removes anything consequential. Output formally produced and never adopted. *Signature.* Deliberation hours flat or rising; durable artefacts per hour falling. Ask what last quarter's meetings deposited (S4.7). **An institution whose meetings deposit nothing has `eta = 0` and is generating no order at any `u` whatsoever.**

Route B — `C` rises. Operating cost grows, usually as a function of `Q` itself, and usually for defensible reasons: each accumulated rule requires interpretation, each precedent must be checked, each standard must be maintained against its dependents. *Mechanisms.* Compliance load proportional to the rule stock. Coordination cost by `n(n−1)/2` as the body grows [Brooks, 1975]. Unresolved conflict, which compounds. Litigation and appeal. The secretariat that grows to service the archive. *Signature.* Model `C_t = c_0 + kappa Q_t` and estimate `kappa` by regression across periods. **A significant positive `kappa` means the institution's own stored order is buying its operating cost**, and it changes the shedding calculation decisively, because shedding `Q` then cuts both the maintenance term and the operating term at once.

Route C — `N` narrows. The categorical variety the institution spans contracts. A body founded to arbitrate between fourteen genuinely distinct constituencies now serves nine, because five merged, exited or were absorbed. *Mechanisms.* Industry consolidation. Professional narrowing. Scope creep in reverse, where the institution defends a shrinking core. Membership criteria that filter for similarity. *Signature.* Count the genuinely distinct categories, not the seats — the redundancy audit of Curzi [2026k], whose five-parameter survey reproduced each of five systems' observed decline, predicting `−0.250474` against an observed `−0.250` at one scale and `−0.996371` against `−1.000` at another, with residuals under 1.4 per cent of the ceiling term. `N` falling from 14 to 9 cuts the ceiling from `eta(2.6391)` to `eta(2.1972)` — **a 17 per cent reduction in the maximum possible input, permanently, and it is almost never noticed because the number of *seats* did not change.**

All three routes usually run together, which is why the diagnosis must be quantitative. Section 7.3 works a case with all three operating simultaneously over thirty-six years.

14.6.7 6.7 Shedding `Q` deliberately

This is the lever nobody reaches for, and at institutional scale it is often the only one with enough range. The reason for the neglect is not mysterious. `Q` is achievement: every unit of it was a solved problem, a settled argument, a prevented disaster. Shedding `Q` means deliberately unlearning something that once worked, and the case for keeping each individual unit is always stronger than the case for keeping the aggregate — because the individual case is concrete and the aggregate case is a division.

The arithmetic is unmoved by this. With `C_t = c_0 + kappa Q_t`,

```
dh/dQ = −(delta + kappa)
```

Every unit of `Q` shed returns `delta + kappa` units of headroom, permanently, every period. With `delta = 0.015` and `kappa = 0.0225`, shedding 8 units of `Q` returns `8 × 0.0375 = 0.30` units of headroom per period thereafter. Against a ceiling of 0.8349 that is 36 per cent of the entire budget, recovered by deletion.

The shedding instruments

Sunset clauses. Every rule carries an expiry date; renewal requires an affirmative act — a positive case, made by someone, at a scheduled time. The mechanism's power is entirely in the default. Without a sunset the default is "keep", and keeping requires no case, no advocate and no attention; with a sunset the default is "lapse", and the burden falls on continuation. **A sunset clause does not delete anything. It changes which side of the argument has to do the work**, and that is sufficient, because the scarce resource in every institution is not judgement but attention. Design notes: stagger expiries so review load is level rather than lumpy; make renewal genuinely possible in one step, or the clause becomes a hostage device; and require the renewal case to state what would make the rule unnecessary.

Term limits. Sunset clauses applied to roles rather than rules. Their storage function is double-edged and both edges should be stated: they raise `delta` on role-bound `Q`, which is a cost, because experience genuinely leaves; and they prevent role-bound `Q` from becoming permanent, which is the benefit. **A role held for thirty years accumulates a `Q` whose `delta` is zero until the day the holder leaves, and then it is 1.0.** Term limits convert a single catastrophic loss into a series of small scheduled ones, which is precisely the trade a well-designed system makes everywhere else.

Planned demolition. The physical form, and the honest one, because it is visible. A building scheduled for removal at a stated date. A programme with a defined end. A committee constituted for a task and dissolved on delivery. The discipline is that the end date is fixed *at the beginning*, when nobody yet has an interest in extending it.

Forgetting by design. The archive tier: records are retained but moved out of the default view, so that search cost — a `C` term — does not grow without bound with the archive. This is the evaporation requirement of FM-4.9 at institutional scale; the pheromone trail to the exhausted source must fade, or the colony maintains a motorway to an empty patch of ground.

Deletion budget. The most useful and least used instrument. Fix a rate — for every new standard adopted, one is withdrawn — and enforce it structurally rather than by intention. The one-in-one-out constraint keeps `Q` bounded by construction rather than by discipline, which matters because discipline is exactly what fails under schedule pressure.

What must not be shed

Shedding is a scalpel and not a policy. Two classes are excluded.

Anything load-bearing for a current dependent. A standard with live implementations is not overhead; it is the interface. Withdrawing it does not shed `Q`, it destroys `Q` elsewhere — usually more than it recovers, and on someone else's balance sheet. The formal version of the warning is that coupled systems fail together in ways neither would alone [Buldyrev et al., 2010].

Anything constituting the ability to decline. Procedural protections, appeal rights, dissent records, exit provisions and consent requirements are `Q` that specifically preserves people's capacity to refuse. **They may be simplified; they may not be shed for headroom.** An institution that buys headroom by removing the mechanisms through which people can say no has not improved its parameters; it has changed what it is, and this framework does not recommend it. The practical form: maintain an explicit **protected set**, written down before any shedding exercise begins, and exclude it from the deletion budget by rule rather than by argument.

14.6.8 6.8 Why the maintenance term is never funded

Star and Ruhleder [1996] characterised infrastructure relationally rather than as a class of object. Something is infrastructure *for* a given practice, and among its properties: it is embedded, transparent in use, learned as part of membership, linked to conventions, and — the property this section turns on — **it becomes visible upon breakdown.** Weiser's 1991 formulation is the same observation from the other side: the most consequential technologies are the

ones that disappear, weaving themselves into everyday practice until they are indistinguishable from it. *That paper has no key in this corpus's bibliography*, and the gap is recorded in Section 10.

Now put that next to the budget:

```
h_t = eta ln N − delta Q_t − C_t
                ^^^^^^^^^
                the maintenance term:
                large, growing, and invisible while it is being paid
```

The `delta Q` term is the cost of keeping infrastructure working. Infrastructure that is working is invisible. **Therefore the largest and fastest-growing term in an institution's budget is the one with no visible advocate, no constituency and no photograph.**

The consequences are predictable and are observed everywhere.

- **Maintenance is deferred first** under budget pressure, because deferring it produces no immediately visible effect. That is not a failure of judgement; it is a correct read of the visible evidence, and the evidence is systematically wrong. The general finding that people systematically misjudge stocks with delays, in a controlled setting where the correct answer is known, is Sterman [1989].
- **Maintenance is measured in avoided events**, and avoided events do not appear in any record. A maintenance regime that works perfectly generates a record indistinguishable from one that was never needed.
- **Maintenance funding is restored only after breakdown**, at which point the cost is the replacement cost rather than the maintenance cost, typically an order of magnitude higher, and the restoration is attributed to the crisis rather than to the neglect. The organisational study of how repeated uneventful deferrals become an accepted baseline until a catastrophic one is Vaughan [1996].
- **The maintainers cannot describe their contribution** in the vocabulary the institution uses for achievement, because their output is the absence of an event.

The engineering response is to make the invisible term visible *by construction*, since it will never become visible by attention.

Instrument	What it makes visible
A named budget line for `delta Q` — separately reported, not merged into operations	That maintenance has a magnitude
A stated `delta` per asset class, published and revised	That decay has a rate
Scheduled visible failure: a planned outage, a fire drill, a restore-from-archive exercise	That the thing is being held up by something
Deterioration reporting: what got worse this period	That neglect has a direction
Maintenance-to-new-build ratio, reported alongside `h`	That the two compete for one budget

Scheduled visible failure is the strongest of these and the least used. If infrastructure becomes visible on breakdown, the engineering inversion is to *cause a small breakdown on purpose, on a schedule*, so that visibility is obtained at a chosen cost rather than at an unchosen one. A quarterly exercise in which a system is restored from the archive tells the institution, in one afternoon, exactly what its `delta` has been doing.

14.6.9 6.9 Why measuring institutional health corrupts it

When a measure becomes a target it ceases to be a good measure [Goodhart, 1975]; and the more a quantitative indicator is used for social decision-making, the more subject it is to corruption pressures and the more apt it is to distort the processes it was meant to monitor [Campbell, 1979]. Institutional health is an unusually soft target and the corruption is unusually complete, because `Q` is not directly observable. What is observable is a proxy `Y`, and every proxy for stored order has a cheap counterfeit.

Proxy for `Q`	The counterfeit	What it does to the real parameters
Standards published	Publish more, thinner standards	`Q` looks up, `C` genuinely up, `eta` unchanged
Members enrolled	Lower the bar	`N_eff` down (new members duplicate old), `C` up
Procedures documented	Document what nobody does	`Q` flat, `C` up, search cost up
Meetings held	Hold more meetings	`C` up, `eta` down (dilution)
Compliance rate	Narrow the rule until compliance is automatic	`Q` down, measured `Q` up
Archive size	Retain everything	`C` up via search, `delta` masked

Every one of these raises the measured proxy and lowers real headroom. **The counterfeit is not usually fraud. It is the natural drift of an organisation optimising a number in good faith**, and it is fast: the corruption is a function of how much weight the indicator carries, not of anyone's intent [Campbell, 1979]. The methodological literature gives the same phenomenon inside research itself, where undisclosed flexibility in analysis produces confident findings from nothing [Simmons et al., 2011, Gelman and Loken, 2013].

Three defences, in order of strength.

First — measure `h`, not `Q`. Headroom is a difference of three quantities pulling in opposite directions and is therefore far harder to counterfeit than any of them. Inflating `Q` *lowers* `h` through the maintenance term; inflating activity raises `C` and lowers `h`. There is no cheap move that raises measured headroom, because the cheap moves all cost something that `h` subtracts. **A gauge built from a difference resists capture in a way a gauge built from a total never can.**

Second — measure `delta` directly, under zero input. Suspend input to a subsystem, observe the decay of `Y`, and compute `delta = 1 − Y_1/Y_0` [Curzi, 2026a]. This is expensive and disruptive and it is nearly impossible to game, because it measures what happens when nobody is trying.

Third — rotate the proxy. Announce that `Y` will be re-specified on a schedule and that the new specification will not be disclosed in advance. This does not prevent optimisation of the current proxy; it prevents the *investment* in optimising it, which is where most of the distortion lives.

And a standing honesty requirement: publish the proxy, publish its known counterfeits, and publish the falsification rules alongside every reported figure [Curzi, 2026i]. **A framework that cannot fail cannot be trusted** — which is the demarcation position the corpus adopts, in the qualified form that treats its own badged postulates as the protective belt around the programme rather than as tested claims [Lakatos, 1970].

14.6.10 6.10 The two ways headroom returns

Here is the structural fact that organises Section 7.3, and it is this section's thesis.

An institution in negative headroom does not stay there forever; the storage law will not permit it. If `h_t < 0` then `eta S_t < delta Q_t + C_t` for every attainable `S_t`, so `Q_{t+1} < Q_t`:

stored order falls. And as `Q` falls, `delta Q` falls, and `h` rises. Eventually `h` crosses back above zero.

Headroom always returns. The only question is whether the `Q` spent to buy it back was chosen or surrendered.

	Involuntary	Deliberate
Mechanism	Decay: the surplus does not cover maintenance, so things quietly stop working	Shedding: rules sunset, programmes end, roles rotate
Timing	Set by `delta` and by the size of the gap	Chosen
What is lost	Whatever was weakest, least defended, or least visible	Whatever was chosen, by stated criteria
Experienced as	Decline, drift, irrelevance	A decision, with an argument attached
Duration in `h < 0`	Years to decades	Zero
Protected set	Not protected — decay does not read the exclusion list	Protected by construction
Attribution	To the current leadership, who did not cause it	To the decision, correctly

The third row is the one to hold on to. **Involuntary shedding does not shed the least valuable `Q`.** It sheds whatever had the fewest defenders — which correlates with invisibility, which by the infrastructure result correlates with *working well* [Star and Ruhleder, 1996]. The archive that never failed loses its funding before the committee that meets monthly and produces nothing, because the committee has members and the archive has none.

Deliberate shedding is the same arrival, chosen, earlier, and with control over what is lost. That is the entire case for sunset clauses, term limits and deletion budgets, and it does not require anyone to believe that the shed rules were bad. It requires only that the budget is finite and that someone should choose.

14.6.11 6.11 The institutional survey, in brief

The full protocol is Curzi [2026i]; the institution-specific points are these.

`Q` proxy. For a rule-holding institution the natural proxy is the count of live rules weighted by dependency. **Weighting is not optional:** an unweighted count treats a foundational standard with two hundred implementations as equal to a procedural note nobody has read, and will therefore recommend shedding the wrong things.

Justify that `Y` is not the input restated. For a standards body, `U` is the disagreement structure among member organisations and `Y` is the live weighted rule stock. These are plainly different observables — a year of intense disagreement producing no adopted standard scores high `U` and zero change in `Y`, which is the required test. **This is the test the corpus's own empirical audit found it could not enforce**: the shuffled-input control passed 98.35 per cent of the time against a third variable driving both interaction and order while detecting only 48.17 per cent in the clean case, so the audit cannot see repackaging [Curzi, 2026i]. The justification therefore has to be made on the observables, in advance, and cannot be recovered statistically afterwards.

`delta` estimation. Institutions rarely permit a zero-input period, so use natural experiments: a dormant working group; a subsystem that lost its staff; the period after a merger when a function was unattended. `delta = 1 − (Y_1/Y_0)` over the dormant interval, annualised.

`C` tally. Include the items nobody budgets: unresolved dispute load, search cost across the archive, compliance interpretation, and the coordination cost `n(n−1)/2` of the governing body itself. These are typically the majority of real `C` and typically absent from the accounts.

`N` audit. Count genuinely distinct constituencies, not seats, and expect the honest `N` to be materially below the headcount [Curzi, 2026k]. The security-engineering form of the same requirement is that an unaudited count of participants is worthless wherever one party can present as many [Douceur, 2002].

A cheaper instrument, with its two conditions. Of the three terms in the gauge, only `C` depends on configuration, so the gradient of headroom is the negative of the gradient of cost: `grad h = −grad C`. **Therefore a habit survey is a headroom measurement.** Where do people actually route their requests; which form gets filled and which gets worked around; which meeting is attended and which is dialled into from a corridor; where has the paving been worn through the grass. Those recurring choices point down the cost field, which is to say up the headroom field. The practical consequence is that **the full survey is not always required to locate the direction of steepest improvement**, only to price it. Two conditions, both binding: the identity holds only where `eta`, `N`, `delta` and `Q` do not themselves move with the configuration — a workaround that destroys a distinction moves `N`, and the identity fails there — and the tally recovers cost *differences*, not absolute costs, so it ranks candidate moves and does not budget them.

Reporting. Report `h` with its three components separated, every period, with the maintenance term shown as a percentage of the ceiling. **That single percentage predicts institutional behaviour better than any strategy document.**

14.6.12 6.12 Specifications

S6.1 — Report headroom, decomposed. The institution **shall** report, each period: `eta ln N`, `delta Q_t`, `C_t`, and `h_t`. The maintenance term **shall** be reported as a percentage of the ceiling. **Reporting `Q` alone is prohibited under this specification**, because `Q` alone is the capturable quantity.

S6.2 — Named maintenance line. `delta Q_t` **shall** appear as a distinct, named budget line, not merged into operations, and **shall** be defended separately in every budget cycle.

S6.3 — Published decay rates. For each asset class holding `Q`, the institution **shall** publish an estimated `delta` and the method used to estimate it, and **shall** revise it at least every three years.

S6.4 — Sunset by default. Every new rule, standard, procedure or programme **shall** carry an expiry date at adoption. Renewal **shall** require an affirmative case, made at a scheduled time by a named advocate. Expiries **shall** be staggered so that review load is level.

S6.5 — Deletion budget. A withdrawal rate **shall** be fixed and enforced structurally. The default requirement is **one-in-one-out by dependency-weighted volume**. The rate may be varied only by an explicit decision recorded with its headroom justification.

S6.6 — Protected set. Before any shedding exercise, a written protected set **shall** be produced, containing at minimum (a) rules with live external dependents and (b) all procedural provisions that constitute a person's ability to decline, appeal, dissent or exit. **The protected set is excluded from the deletion budget by rule and not by argument.**

S6.7 — Term limits with handover deposit. Where term limits apply, the outgoing holder **shall** deposit role-bound `Q` in the shared medium of Section 4 before departure, and the handover **shall** be verified by the successor's ability to reconstruct the state without the predecessor present. **A term limit without a handover deposit converts a slow leak into a scheduled catastrophe.**

S6.8 — Scheduled visible failure. At least annually, the institution **shall** conduct an exercise in which a piece of working infrastructure is deliberately made visible: a restore from archive, a planned outage, a reconstruct-from-record drill. Findings **shall** be reported against the published `delta` estimates of S6.3.

S6.9 — The `h < 0` rule. When measured `h_t < 0`, the institution **shall not** commission new durable structure. Permitted actions are exactly: raise `eta`, cut `C`, shed `Q`. **Recruiting for difference is explicitly not a permitted response to negative headroom** (Section 6.3), because `u` does not appear in `h`.

S6.10 — Proxy honesty. Every reported proxy **shall** be published together with its known counterfeits (Section 6.9) and with the falsification rules of Curzi [2026i], including the shuffled-input control.

14.6.13 6.13 Failure modes

FM-6.1 — The maintenance trap. *Description.* `delta Q_t` has grown to consume the entire budget. The institution maintains and cannot build; every proposal is met with an assertion that there is no capacity, which is precisely correct and is mistaken for a cultural problem. *Damages.* `h` to zero or below, and then `Q` involuntarily. *Diagnostic signature.* `delta Q_t / (eta ln N) >= 0.9`. Corroborating symptoms: no new durable output for several periods despite unchanged staffing; strategy documents that repeat across years; a rising share of effort in interpretation and support of existing structure. *Fix.* Shed `Q` (S6.5) and cut `C`. **Do not attempt to raise the ceiling** — `ln N` is logarithmic and will not deliver. This is the failure mode for which deliberate shedding was invented.

FM-6.2 — The unshed `Q`. *Description.* Rules, programmes and roles accumulate without any withdrawal mechanism. Nothing has been removed in living memory. The rule stock is a stratigraphy in which nobody can tell live from fossil. *Damages.* `delta Q` rises toward the uniform limit; `C` rises through search and interpretation cost; `eta` falls because new work must be reconciled against everything. *Diagnostic signature.* Count withdrawals over the last five years. **Zero is the fault.** Second test: sample ten rules at random and count how many have a live dependent; a hit rate below half means most of the stock is fossil. *Fix.* S6.4 and S6.5. Begin with the sample: withdrawing ten rules that have no dependents is uncontroversial and establishes that withdrawal is possible, which is the actual barrier.

FM-6.3 — Invisible infrastructure. *Description.* Maintenance is unfunded because working infrastructure is invisible. The budget is cut, nothing happens for two years, the cut is judged a success, and the cut is repeated. Then something fails completely. *Damages.* `delta` effectively rises because unmaintained order decays faster; `Q` takes a step loss at breakdown; `C` is paid at crisis prices. *Diagnostic signature.* Maintenance budget falling in real terms across periods with no measured rise in `delta` — because nobody is measuring `delta`. **The absence of the measurement is the signature.** *Fix.* S6.2, S6.3 and S6.8. The scheduled visible failure is the one that changes minds, because it converts an argument into an observation.

FM-6.4 — Proxy capture. *Description.* A proxy for institutional health became a target. The proxy now rises reliably while the institution declines. *Damages.* All parameters, and the measurement system that would have detected the damage. *Diagnostic signature.* The reported metric improves monotonically across periods — real institutional quantities are noisy, and **a smooth monotone series is evidence of management of the series, not of the thing.** Cross-check: does an independent outcome measure move with the proxy? *Fix.* S6.1 — report `h` rather than any total. Rotate the proxy on an undisclosed schedule. Publish the counterfeits.

FM-6.5 — Ossification diagnosed as culture. *Description.* An institution with `h < 0` is treated to a leadership change, a values programme, a restructure, or an exhortation to be more innovative. None of it works. The next leadership is blamed and the cycle repeats. *Damages.* `C`, since the interventions are themselves expensive; morale; and time, while the real constraint continues to bind. *Diagnostic signature.* Compute `h`. If `delta Q_t + C_t > eta ln N`, the sustainability ceiling applies and **no arrangement of the people prevents decline** [Curzi, 2026a]. The signature is that successive leaderships of differing style produce indistinguishable results, which is exactly what a budget constraint predicts and a culture theory does not. *Fix.* Relieve the budget: shed `Q`, cut `C`, raise `eta`. Then, and only then, is a leadership or composition question worth asking.

FM-6.6 — Difference recruited into a deficit. *Description.* An institution with negative headroom responds by broadening its membership and recruiting for difference. `u` rises. Nothing improves. The failure is then attributed to the new members or to resistance from the old. *Damages.* `C`, since a larger body costs quadratically more to coordinate; and the credibility of `u`-raising interventions for a generation. *Diagnostic signature.* `h < 0` before the intervention; `u` measurably higher after; `Q` still falling. *Fix.* Sequence correctly, per S6.9: bring `h` above zero by shedding and by `eta` and `C` work, and *then* raise `u` to fill the budget that has been created. **The intervention is not wrong; it was administered in the wrong order.**

FM-6.7 — The sunset that never sets. *Description.* Sunset clauses exist and every renewal is automatic. The review is a formality, the advocate is the secretariat, and no rule has ever lapsed. *Damages.* `C`, since the review process is pure overhead; plus a false sense that FM-6.2 has been addressed. *Diagnostic signature.* Renewal rate at or near 100 per cent. **A functioning sunset regime lapses things; a renewal rate above roughly 90 per cent means the mechanism is decorative.** *Fix.* Require the renewal case to be made by a named beneficiary who is not the secretariat, and require it to state a live dependent. Pair with S6.5, which forces a withdrawal whenever an adoption occurs and therefore cannot be satisfied by universal renewal.

FM-6.8 — Term limits without deposit. *Description.* Rotation is enforced. Nothing was written down. Each rotation loses the role's accumulated judgement entirely. *Damages.* `Q` takes a step loss per rotation; `delta` is effectively raised to the rotation rate on all role-bound order. *Diagnostic signature.* Performance sawtooth synchronised with the rotation cycle; bus-count of 1 on role-bound state (S4.6). *Fix.* S6.7. Term limits and shared media are complements, not alternatives: **rotation without deposit is not a shedding strategy, it is a leak with a calendar.**

FM-6.9 — The protected set raided. *Description.* Under headroom pressure, the shedding exercise reaches appeal rights, dissent procedures, consultation requirements and exit provisions, because they are procedurally expensive and have no organised defenders. *Damages.* It breaches the declination constraint; and in the model's own terms it destroys the `u`-preserving structure of Section 5 and therefore the input term, so it does not even deliver what it promises. *Diagnostic signature.* The shedding list contains items whose function is to let someone say no. *Fix.* S6.6, enforced structurally: the protected set is written before the exercise begins and is excluded by rule. **Headroom bought by removing people's ability to decline is outside the scope of this framework, and no arithmetic in it licenses that trade.**

14.7 7 Results: three systems worked end to end

All three systems below are constructed cases with declared parameters, not measured organisations. Every number is computed from the stated parameters under the model of Section 3, and every table can be re-derived from the equations and parameter values printed beside it. They demonstrate that the procedure runs, that it reports its own failures, and that it produces counter-results as readily as confirmations; **they are not evidence that any organisation behaves this way and must not be cited as such.** Section 7.5 reports six places where the source material's arithmetic or its description of its own arithmetic does not hold up, and prints them rather than correcting them silently.

14.7.1 7.1 A neighbourhood tool library: coordinator against medium

A tool library with 48 member households, a shared inventory of hand and power tools, a converted garage, and a shared obligation to maintain both. It is a textbook common-pool resource — subtractable, since one person has the mitre saw and another does not, and hard to exclude from.

7.1.1 Encoding

The five operational choices of Curzi [2026i] are fixed and stated before any number is computed.

Choice	Setting
State encoder	Each member household's activity assigned to one of `N = 8` craft categories: woodwork, metalwork, garden, automotive, textiles, electronics, masonry, cycle repair
Interaction observable	Cross-category tool loans and joint working sessions per month
Order observable `Y`	Composite: tools in working order and correctly located, live maintenance skills held by more than one member, documented repair procedures
Time scale	One month = one period
Cost observable	Volunteer hours plus consumables plus loss and breakage, normalised to the same units as `Y`

The requirement that the order proxy not be a repackaging of the input holds here: `Y` is measured from a physical inventory audit and a skills census conducted independently of the loan log that produces `U`. A month with heavy borrowing and no repairs scores high `U` and flat `Y`. **The two observables can and do move in opposite directions, which is the test.**

7.1.2 Parameters

`N = 8`, so `ln N = 2.0794`. Non-parallelism from the disagreement-and-decorrelation protocol: `u = 0.62`.

```
N_eff = 1 + (N−1)u = 1 + 7(0.62) = 5.34
S     = ln(N_eff)  = ln(5.34)    = 1.6752
```

Forty-eight households, eight categories, and an effective count of 5.34. The redundancy audit removed a third of the categorical variety — 5.34 against 8, a 33.3 per cent reduction — before any surplus was credited, and 19.4 per cent of the available surplus with it, since `1.6752/2.0794 = 0.8056`. Both reductions are stated because they are different quantities and Section 7.5 records that the source conflates them.

7.1.3 Regime A — the coordinator

The library's first two years ran on a part-time paid coordinator holding the loan ledger, the maintenance schedule, the key rota and the membership list.

```
eta   = 0.35     interaction converts only via the coordinator's records
delta = 0.25     the coordinator's own turnover and forgetting rate
C     = 0.52     stipend, scheduling, queueing on one person's availability
```

Headroom at the start, `Q_0 = 0`:

```
h_0 = eta ln N − delta Q_0 − C = 0.35(2.0794) − 0 − 0.52 = 0.7278 − 0.52 = +0.2078
```

Positive but thin. Steady state:

```
eta S = 0.35 x 1.6752 = 0.5863
Q_inf = (eta S − C) / delta = (0.5863 − 0.52) / 0.25 = 0.2653
```

t (months)	0	2	4	6	8	10
Q_t	0.000	0.116	0.181	0.218	0.239	0.250

The library reaches a steady state of `Q_inf = 0.265` and stays there. It works. It never gets better. Check the identity of Section 3.2:

```
h at Q_inf      = 0.35(2.0794) − 0.25(0.2653) − 0.52 = +0.1415
eta ln(N/N_eff) = 0.35 x (2.0794 − 1.6752) = 0.35 x 0.4042 = +0.1415     [agrees]
```

Then the coordinator moves house. Their successor inherits the ledger but not the relationships, not the knowledge of which drill has the intermittent fault, not the informal understandings about who gets priority in spring. The hub census had a bus-count of 1 on most critical state. Estimated loss: **70 per cent of `Q` in a single step**, so `Q` drops from 0.239 to 0.072 at month 8. Recovery to the pre-shock level takes 6.9 months under the model. **The library did not lose a volunteer. It lost most of its stored order, because that is where the order was stored.**

7.1.4 Regime B — the medium

State is moved out of the coordinator and into two physical artefacts and one digital one:

- A **peg board** with a hook per tool: the tool's presence or absence, the borrower's tag and the return date, all readable from the doorway in three seconds. Visible 4/4.
- A **bound logbook** on a chain by the door: date, tool, member, condition on return, one line for anything odd. Writable in under fifteen seconds with a pen. Writable 4/4, low-friction 4/4, append-oriented 4/4.
- A **shared spreadsheet** mirrored weekly from the logbook by whoever is on rota, holding the membership list, maintenance schedule and repair procedures. Durable 3/4.

VWDLA total 19/20; passes S4.1. The unfinished-edge display required by S4.2 is a red tag hung on the peg for any tool needing attention, plus a chalked list of outstanding work on the garage wall. **The board displays its own defects, which is what makes it a stimulus rather than a report.**

Revised parameters, over the following year:

```
eta   = 0.55     every loan and every fault deposits directly into durable record
delta = 0.12     state survives departures; loss is physical wear and slow drift
C     = 0.45     no stipend; rota cost and write-cost distributed across members

h_0   = 0.55(2.0794) − 0 − 0.45 = 1.1437 − 0.45 = +0.6937
eta S = 0.55 x 1.6752 = 0.9214
Q_inf = (0.9214 − 0.45) / 0.12 = 3.9281
```

t (months)	0	2	4	6	8	10	12
Q_t	0.000	0.886	1.572	2.104	2.515	2.834	3.081

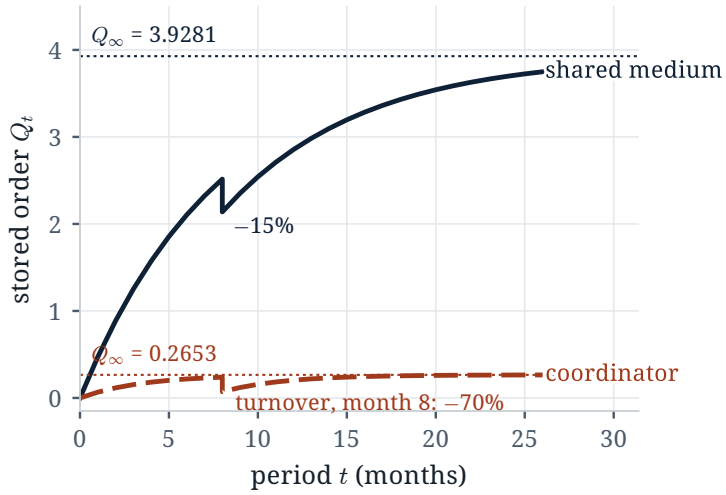


Figure 14.3: **The same turnover event costs 70 per cent of the stock under a coordinator and 15 per cent under a shared medium**, and the two regimes settle 14.8-fold apart — `Q_inf` = 0.2653 against 3.9281 — on a change of `eta` from 0.35 to 0.55 and `delta` from 0.25 to 0.12, with no change of people. Modelled from the parameters stipulated in Sections 7.1.3 to 7.1.5; generated by `figs/src/CCP-014_fig02.py`.

t (months)	0	2	4	6	8	10	12
`h_t`	+0.694	+0.587	+0.505	+0.441	+0.392	+0.354	+0.324

Identity check again:

```
h at Q_inf      = 0.55(2.0794) − 0.12(3.9281) − 0.45 = +0.2223
eta ln(N/N_eff) = 0.55 x 0.4042 = +0.2223     [agrees]
```

Time constant `tau = 1/delta = 8.33` months; half-life of stored order under zero input is `ln 2 / (−ln(1−0.12)) = 5.42` months.

Regime B holds 14.8 times the stored order of Regime A — 3.928 against 0.265 — at lower total cost. The mechanism is not effort: nobody in Regime B works harder than the coordinator did. The mechanism is that `eta` rose from 0.35 to 0.55, because every act of borrowing now deposits, and `delta` fell from 0.25 to 0.12, because the deposit is not in a person. This is result **R2**.

7.1.5 The turnover test, repeated

Regime B's most active member also leaves, at month 8. The hub census now shows no critical state at bus-count 1; what is lost is relationship and judgement, not record. Estimated loss: **15 per cent**, so `Q` drops from 2.515 to 2.138. Recovery to the pre-shock level takes 1.7 months under the model, against Regime A's 6.9. **Same event, same organisation, same people. Different location for the state.**

7.1.6 Free riding and the monitoring decision

Twenty months in, the contribution-share plot required by FM-4.4 shows the top quartile of households supplying 71 per cent of maintenance hours, up from 58 per cent. Estimated free-riding fraction `phi = 0.35`, loss coefficient `lambda = 0.9`, base cost 0.40. The arithmetic is computed in Section 4.8: without monitoring `Q_inf = 1.7198`, and with mutual monitoring at 0.08 and graduated sanctions bringing `phi` to 0.06, `Q_inf = 3.2281`.

Stored order had been more than halved by free riding — 3.93 down to 1.72 — without a single visible failure. Nothing broke; it simply cost more than it did. **Eight hundredths of a unit of monitoring cost bought back 1.508 units of steady-state stored order**, an increase of 87.7 per cent on the free-ridden level and a recovery of 68.3 per cent of what free riding had taken. Level 1 of the sanction ladder was used eleven times in the first year, level 2 twice, level 3 never. That distribution is the signature of a working graduated system and is the direct answer to FM-4.5.

7.1.7 Ostrom conformance audit

Principle	Evidence	Score
O1 boundaries	Written membership roll; inventory list defines the resource	2
O2 congruence	Borrowing limits set by tool scarcity; borrowing rights tied to rota participation	2
O3 collective choice	Rules changed at the quarterly members' meeting; sanction ladder was written there	2
O4 monitoring	Peg board and logbook are mutual and by-product; no external monitor	2
O5 graduated sanctions	Three levels; mildest used 11 times this year	2
O6 conflict resolution	Any two members may call a meeting within seven days; no cost	2
O7 right to organise	Independent association; no external authority has voided a rule	2
O8 nesting	Four craft circles of 8–14 households, plus a delegate council of 4	1

Total 15/16. O8 scores 1 because two of the four craft circles exceed the headroom-positive band and the delegate council's remit is not written down. **The band cannot be recomputed here, because the case never states this system's own per-path cost** — a parameter the specification S4.4 requires and the worked case omits. Section 7.5 records the omission; under the reference per-path cost of 0.02 the band is `n <= 11`, so circles of 12 to 14 exceed it, but that band is borrowed and not this system's.

7.1.8 Falsification, declared in advance

1. If, after installing the medium, measured `delta` did not fall below 0.20, **postulate M1 failed in this system** and the redesign should be reported as unsuccessful regardless of how good the board looks.
2. If `Q` under Regime B did not exceed Regime A's steady state within 12 periods, the modelling of `eta` is wrong.

3. If the monitoring intervention did not reduce measured `phi` within two quarters, O4 does not transfer to this setting, and that should be reported as a negative result, not explained away.
4. The shuffled-input control [Curzi, 2026i] must be run: if `Y` predicts as well from shuffled `U` as from real `U`, the model is fitting a trend and nothing more.

7.1.9 The rule extracted

The redesign changed no personnel, no funding and no rules of membership. It changed where the state lived. That single move raised `eta` by 57 per cent, cut `delta` by 52 per cent, and multiplied steady-state stored order fifteenfold. Then a monitoring regime costing 0.08 units recovered most of what free riding had quietly removed. **Put the state in the medium, not in a coordinator. Then govern the medium by the eight principles, or the tragedy argument applies as written.**

14.7.2 7.2 A nine-person review panel: independence priced

A standing panel of nine reviewers assessing proposals against a fixed rubric, meeting monthly. Decisions are consequential and outcomes are eventually observable, which makes the panel measurable in a way most committees are not.

7.2.1 Encoding

Choice	Setting
State encoder	Each reviewer's judgement vector across `N = 9` rubric dimensions
Interaction observable	Cross-reviewer disagreement per item, measured pre-discussion
Order observable `Y`	Panel's calibrated accuracy against realised outcomes at 24 months, plus the durable rubric refinements the panel produces
Time scale	One month = one period
Cost observable	Reviewer hours, meeting time, secretariat, appeals

`Y` is measured from realised outcomes 24 months downstream, entirely independently of the panel's own disagreement record, which satisfies the requirement that the order proxy not be a repackaging of the input.

7.2.2 Baseline

Standing parameters, over eight periods:

```
N     = 9        ln N = 2.1972
eta   = 0.45     proportion of panel deliberation that becomes durable rubric or precedent
delta = 0.18     turnover of two members a year plus drift in rubric interpretation
C     = 0.42     reviewer time, meeting, secretariat, appeals
Q     = 1.80     accumulated calibrated rubric and precedent
u     = 0.12     by the disagreement-and-error-decorrelation protocol
```

Seven of the nine reviewers were trained in three institutions; six sit on at least one other panel together; their pre-discussion ratings agree within one rubric point on 86 per cent of items.

```
N_eff = 1 + 8(0.12) = 1.96
S     = ln(1.96)    = 0.6729
```

Nine reviewers with an effective count of 1.96. The panel is, in the law's accounting, slightly less than two people.

7.2.3 The diagnosis

```
h = eta ln N − delta Q − C
  = 0.45(2.1972) − 0.18(1.80) − 0.42
  = 0.9888 − 0.3240 − 0.4200 = +0.2448
```

Headroom is positive, and this is important and counter-intuitive: **the panel is not overcommitted. It has budget. It is simply not collecting the input to use it.** Check the step:

```
eta S       = 0.45 x 0.6729 = 0.3028
delta Q + C = 0.3240 + 0.4200 = 0.7440
step        = 0.3028 − 0.7440 = −0.4412
```

The panel is losing 0.44 units of stored order per period, and

```
Q_inf = (eta S − C) / delta = (0.3028 − 0.4200) / 0.18 = −0.6510
```

A negative equilibrium means the regime is unsustainable as modelled — the storage law reports its own domain of breakdown rather than concealing it [Curzi, 2026a]. The panel's accumulated rubric and precedent will erode at a rate governed by `tau = 1/delta = 5.6` months, and the erosion will be attributed to turnover, workload, or the declining quality of applications. None of those is the cause. The cause is that `u = 0.12`.

7.2.4 What is required

By the increment requirement of Section 3.1, achieving an increment `r` per period requires `S >= (delta Q + C + r)/eta`, feasible only if `delta Q + C + r <= eta ln N`.

target `r`	required `S`	required `u`	feasible?
0.00 (hold)	1.6533	0.5280	yes
0.05	1.7644	0.6048	yes
0.10	1.8756	0.6906	yes
0.2448 (`= h`, the maximum)	2.1972	1.0000	yes, at orthogonality exactly

Read the first row before anything else. **To merely stop losing ground, this panel needs `u >= 0.528`; it currently has 0.12.** Every month spent below that threshold is a month of erosion, and no amount of reviewer diligence substitutes, because diligence does not appear in the equation. The last row is `r_max = h = 0.2448` arriving as a hard fact, and it requires `u = 1`, perfect orthogonality, which is not attainable with real people. **The gauge says what may be asked for; the geometry says what can be delivered; they meet exactly at the ceiling and not before.** Target: `r = 0.10`, requiring `u >= 0.6906`.

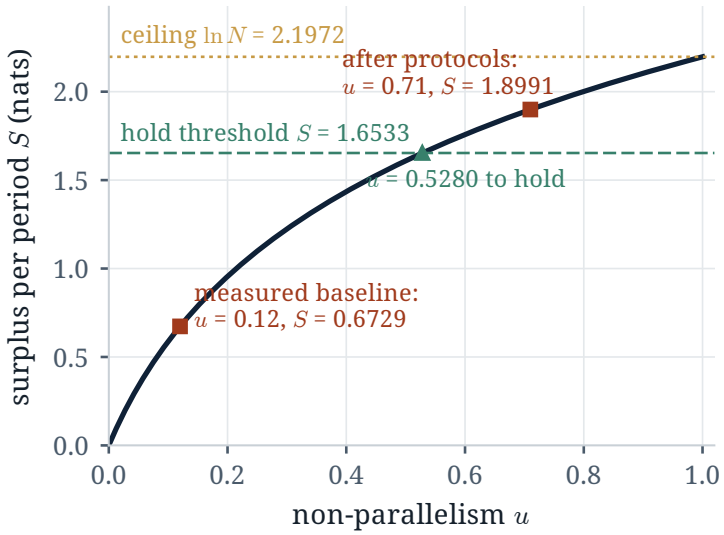


Figure 14.4: **The panel's deficit is an input deficit, not a budget deficit**: at `u` = 0.12 the nine-person panel generates `S` = 0.6729 nats against a hold threshold of 1.6533 and a ceiling of `ln 9` = 2.1972, and the five independence protocols move it to `u` = 0.71 and `S` = 1.8991. Modelled from the parameters stipulated in Section 7.2; generated by `figs/src/CCP-014_fig03.py`.

7.2.5 The intervention

All five protocols of Section 5.7, installed together, at a total cost increment of 0.090, so `C` rises from 0.420 to 0.510.

Protocol	Implementation at this panel
IP-1	Each reviewer submits a sealed rating with one line of reasoning 48 hours before the meeting; the distribution is displayed first for each item
IP-2	The sealed round is unattributed; names attach in round two, optionally
IP-3	Published dissent protection, signed by the commissioning body; dissent is not visible to the appointing authority at reappointment
IP-4	Minority report with a named falsifier on every split decision; consented rotating steelman on unanimous ones
IP-5	Round-robin on each item; the chair speaks last; the table replaced with a round one

Composition change, per FM-5.6, because protocols cannot extract signal that is not present: two retiring members replaced by a practitioner from an adjacent field and a reviewer from outside the three dominant training lineages. Measured after four periods: `u = 0.71`.

7.2.6 The result

```
N_eff = 1 + 8(0.71) = 6.68
```

```
S     = ln(6.68)    = 1.8991
eta S = 0.45 x 1.8991 = 0.8546

h     = 0.9888 − 0.3240 − 0.5100 = +0.1548
step  = 0.8546 − 0.3240 − 0.5100 = +0.0206
Q_inf = (0.8546 − 0.5100) / 0.18 = 1.9144
```

From a negative equilibrium to `Q_inf = 1.9144`. Confirm with the identity:

```
h at Q_inf      = 0.45(2.1972) − 0.18(1.9144) − 0.5100 = +0.13415
eta ln(N/N_eff) = 0.45 x (2.1972 − 1.8991) = 0.45 x 0.2981 = +0.13415     [agrees]
```

Note honestly what happened to headroom: it *fell*, from 0.2448 to 0.1548, because `C` rose by 0.09. **Headroom went down and the system got better.** This is not a contradiction; it is the correct reading of the gauge. Headroom is capacity to build, and the panel spent some of it on the mechanism that lets it build at all. A large unspent headroom in a system with `u = 0.12` is not health. It is a budget nobody can draw on.

7.2.7 Accuracy, computed

Take individual reviewer accuracy on the binary question of whether a proposal would have succeeded as `p = 0.62`, from the 24-month outcome record. Under full independence, nine reviewers by majority give `P_independent(9, 0.62) = 0.7738`. Under Model C1 — badged as a two-point mixture chosen for transparency, not a general correlation model — with the baseline correlation `rho = 0.75`:

```
P = 0.75(0.62) + 0.25(0.7738) = 0.6584
```

and after the intervention, with `rho` estimated at 0.20:

```
P = 0.20(0.62) + 0.80(0.7738) = 0.7430
```

Panel accuracy rises from 65.8 per cent to 74.3 per cent with no change in the individual ability of any reviewer. Seven of the nine are the same people, reading the same proposals, using the same rubric. The change is entirely in the order of operations and the composition of the pool. This is result **R4**.

Compare against the cascade bound. If the panel deliberated by open sequential discussion with no independent round, Section 5.4 gives, at `p = 0.62`:

```
P(correct cascade) = [0.62(1.62)/2] / [0.62^2 − 0.62 + 1] = 0.5022 / 0.7644 = 0.6570
```

65.70 per cent, against the panel's baseline of 65.84 per cent — a gap of 0.14 percentage points. **The panel of nine was performing at the theoretical accuracy of an unbounded cascade**, which is what a nine-person cascade should do, and is the strongest single corroboration available that the baseline diagnosis was right. **It is a corroboration and not a test**: the baseline `rho = 0.75` was chosen, not measured, and a different choice would break the coincidence. Section 10 records this.

7.2.8 Falsification, declared in advance

1. If measured `u` does not exceed 0.53 within six periods, the intervention has failed to reach even the hold threshold and must be reported as unsuccessful.
2. If the pre-discussion distributions show the same spread as post-discussion positions did under the old regime, IP-1 is not working and postulate I1's mechanism is wrong in this body.

3. If 24-month outcome accuracy does not rise above the 0.658 baseline within two full outcome cycles, the accuracy model is wrong and should be discarded rather than defended.
4. The shuffled-input control is mandatory: if Y is predicted as well by shuffled U as by real U, the entire causal story is a trend and must be reported as such.
5. If the dissent rate rises but decision-change-following-dissent stays at zero, the panel has acquired FM-5.3 and the apparent improvement in u is theatre.

7.2.9 The rule extracted

The panel was not short of ability. Its members were individually accurate at 62 per cent, above the Condorcet threshold, and its headroom was positive throughout. It was failing because nine reviewers had an effective count of 1.96, and because the order of operations in the room guaranteed a cascade before the second item. Ninety thousandths of a unit of procedural cost, two composition changes and a round table moved it from a negative equilibrium to a positive one and lifted its accuracy by 8.5 points. **Individuality is not a value the group tolerates. It is the structure the group's function rests on.**

14.7.3 7.3 A technical standards body, 1962 to 2026

A standards body founded in 1962 to publish interface standards for a manufacturing sector, with sixty-four years of continuous operation. All figures are computed; the parameter values are a constructed case rather than a measured organisation, and are stated as such.

7.3.1 Encoding

Choice	Setting
State encoder	Member organisations grouped into distinct constituencies by manufacturing role
Interaction observable	Cross-constituency working-group participation and recorded technical disagreement
Order observable Y	Live standards weighted by number of external implementations
Time scale	One year = one period
Cost observable	Secretariat, working-group time, interpretation requests, appeals — normalised to Y units

Y, the dependency-weighted live standard stock, is independent of U, the disagreement structure among members: a year of intense committee disagreement producing no adopted standard registers high U and zero change in Y.

7.3.2 Era 1 — 1962 to 1990

Founding parameters, estimated from the archive:

```
N     = 14        fourteen genuinely distinct constituencies    ln N = 2.6391
u     = 0.60      constituencies with materially different interests
eta   = 0.62      working groups produced adopted standards at high rate
delta = 0.030     standards decay slowly; technology drift and disuse
```

```
C     = 0.320     small secretariat, low interpretation load
Q_0   = 0         nothing stored at founding

N_eff    = 1 + 13(0.60) = 8.80
S        = ln(8.80)     = 2.1748
eta ln N = 0.62 x 2.6391 = 1.6362      <- the ceiling
eta S    = 0.62 x 2.1748 = 1.3483      <- the input
```

Year	Q_t	h_t
1962	0.000	+1.3162
1967	4.842	+1.1709
1972	9.001	+1.0462
1976	11.900	+0.9592
1982	15.638	+0.8471
1990	**19.669**	**+0.7261**

Twenty-eight years of accumulation, with headroom positive throughout and falling steadily — **success consuming its own headroom.** Steady state for these parameters would be `Q_inf = (1.3483 − 0.320)/0.030 = 34.28`, with `h_inf = eta(ln N − S) = 0.62(0.4643) = 0.2879` by Identity 2. The body never reached it. Maintenance share at 1990: `delta Q / (eta ln N) = 0.030 x 19.669 / 1.6362 = 36` per cent — substantial and not yet critical.

7.3.3 Era 2 — the parameters change, the `Q` does not

Between roughly 1985 and 1992 the sector consolidated, and all three routes of Section 6.6 operate at once.

Route C — `N` narrows. Fourteen constituencies become nine through merger and exit; the ceiling falls from `eta(2.6391)` to `eta(2.1972)`.

Route A — `eta` falls, 0.62 to 0.38. Standards now require legal review, patent clearance and a formal objection period. Working-group output that once became a standard in eighteen months now takes four years, and a growing share never completes.

Route B — `C` becomes a function of `Q`. With 19.7 units of live dependency-weighted standards, interpretation requests, conformance disputes and revision maintenance scale with the stock: `C_t = c_0 + kappa Q_t` with `c_0 = 0.320` and `kappa = 0.0225`.

And `delta` is deliberately driven down, 0.030 to 0.015, by a policy that no standard is ever withdrawn — which the body regards, reasonably, as a service to its implementers. Also `u` falls from 0.60 to 0.28: the surviving members are large integrated firms with substantially aligned interests.

```
N     = 9         ln N = 2.1972
u     = 0.28  ->  N_eff = 1 + 8(0.28) = 3.24,   S = ln(3.24) = 1.1756
eta   = 0.38  ->  eta ln N = 0.8349,   eta S = 0.4467
delta = 0.015
C_1990 = 0.320 + 0.0225(19.669) = 0.7626

h = 0.8349 − 0.015(19.669) − 0.7626 = 0.8349 − 0.2950 − 0.7626 = −0.2227
```

Negative. The body did nothing wrong in 1990. It inherited a `Q` of 19.669 accumulated under a ceiling of 1.6362 and is now carrying it under a ceiling of 0.8349. **Negative headroom is inherited order under changed parameters**, which is Identity 2's corollary in the flesh. The zero-headroom stock level for the new parameters is

```
Q* = (eta ln N − c_0) / (delta + kappa) = (0.8349 − 0.320) / 0.0375 = 13.73
```

The body holds 19.669. **It is 5.94 units of stored order over the line its current parameters can support**, and every one of those units was a solved problem.

7.3.4 Path A — do nothing

The body continues. Standards are neither withdrawn nor renewed; working groups meet; adoption is slow. `Q` decays because `eta S = 0.4467` does not cover `delta Q + C`.

Year	`Q_t`	`C_t`	`h_t`	Maintenance + operating as % of ceiling
1990	19.669	0.7626	**-0.2227**	127%
1995	16.835	0.6988	**-0.1164**	114%
2000	14.495	0.6461	**-0.0286**	103%
2005	12.561	0.6026	+0.0439	95%
2007	11.885	0.5874	+0.0692	92%
2010	10.964	0.5667	+0.1038	88%
2015	9.644	0.5370	+0.1533	82%
2020	8.554	0.5125	+0.1942	77%
2026	**7.494**	0.4886	+0.2339	72%

Twelve consecutive years of negative headroom, 1990 through 2001. The map is `Q_{t+1} = 0.9625 Q_t + 0.126717`, so the stock crosses `Q* = 13.73` at `t = ln(10.3516/16.2899)/ln(0.9625) = 11.87` years — that is, during 2002 — which is exactly twelve completed years at or below zero headroom, 1990 to 2001 inclusive. Twelve years in which the body could not, by the sustainability ceiling, commission any durable new structure at any interaction geometry whatsoever. During those twelve years the body appointed three directors-general, ran two strategic reviews and one culture programme. This is FM-6.5 in its natural habitat.

Headroom did return, in 2002, and by 2026 it stands at +0.2339. Read the price:

```
Q: 19.669 (1990)  ->  7.494 (2026)
Lost: 12.175 units = 61.9% of everything accumulated in twenty-eight years
```

The body bought back its headroom by losing 61.9 per cent of its stored order, over thirty-six years, without choosing a single unit of what was lost. What lapsed was whatever had the fewest defenders — which correlates with whatever was working quietly. This is result **R5**.

7.3.5 Path B — deliberate sunset in 1990

Counterfactual. In 1990 the body computes `h`, finds `−0.2227`, and applies S6.4 through S6.6: a protected set is written — standards with live implementations, all appeal and objection procedures — and 40 per cent of dependency-weighted stock is sunset in a single scheduled exercise.

```
Q: 19.669 -> 11.801
C: 0.320 + 0.0225(11.801) = 0.5855
h = 0.8349 − 0.015(11.801) − 0.5855 = +0.0724
```

Positive immediately.

Year	Q_t	h_t
1990	11.801	+0.0724
1995	10.336	+0.1273
2000	9.126	+0.1727
2010	7.301	+0.2412
2026	5.507	+0.3084

Zero years of negative headroom. But note the honest result, and it is the reason this path is not the recommendation: **Path B ends in 2026 with `Q = 5.507`, which is *less* than Path A's 7.494.** Shedding alone bought thirty-six years of positive headroom and then spent none of it. `eta` was still 0.38 and `u` still 0.28, so `eta S` was still 0.4467 and still insufficient. The body was solvent and idle. This is the second half of counter-result **CR**.

> **Shedding buys headroom. Headroom unspent is headroom wasted. Shedding is not a strategy; it is the precondition for one.**

7.3.6 Path C — shed, then spend

The full intervention, sequenced per S6.9: stabilise first, then grow.

1. **Shed 40 per cent** of dependency-weighted stock, protected set excluded: `Q: 19.669 -> 11.801`.
2. **Raise `eta` from 0.38 to 0.60.** Working-group output is captured directly into draft standards with a standing editor; the objection period runs concurrently with drafting rather than after it; every working group deposits to the shared medium (S4.7). This is `eta` work — the cheapest lever — and it does not require any change in what anyone is trying to do.
3. **Raise `u` from 0.28 to 0.55** — *only now*, with `h` positive. Membership is opened to integrators, testing laboratories and two adjacent-sector bodies, and the protocols of Section 5.7 run in every working group.

   eta ln N = 0.60 x 2.1972 = 1.3183 N_eff = 1 + 8(0.55) = 5.40 S = ln(5.40) = 1.6864 eta S = 0.60 x 1.6864 = 1.0118 h_1990 = 1.3183 - 0.015(11.801) - 0.5855 = +0.5558

Year	Q_t	C_t	h_t
1990	11.801	0.5855	+0.5558
1995	12.958	0.6116	+0.5124
2000	13.913	0.6330	+0.4766
2010	15.354	0.6655	+0.4226
2026	**16.770**	0.6973	**+0.3695**

Steady state and identity check:

```
Q_inf = (eta S − c_0) / (delta + kappa) = (1.0118 − 0.320) / 0.0375 = 18.449
h_inf = eta ln(N / N_eff) = 0.60 x (2.1972 − 1.6864) = 0.60 x 0.5108 = +0.3065
```

The trajectory converges on `Q_inf = 18.449` with `h_inf = +0.3065`, consistent with the table. Time constant `tau = 1/(delta + kappa) = 26.7` years — institutions are slow, and the slowness is the same property that makes them worth having.

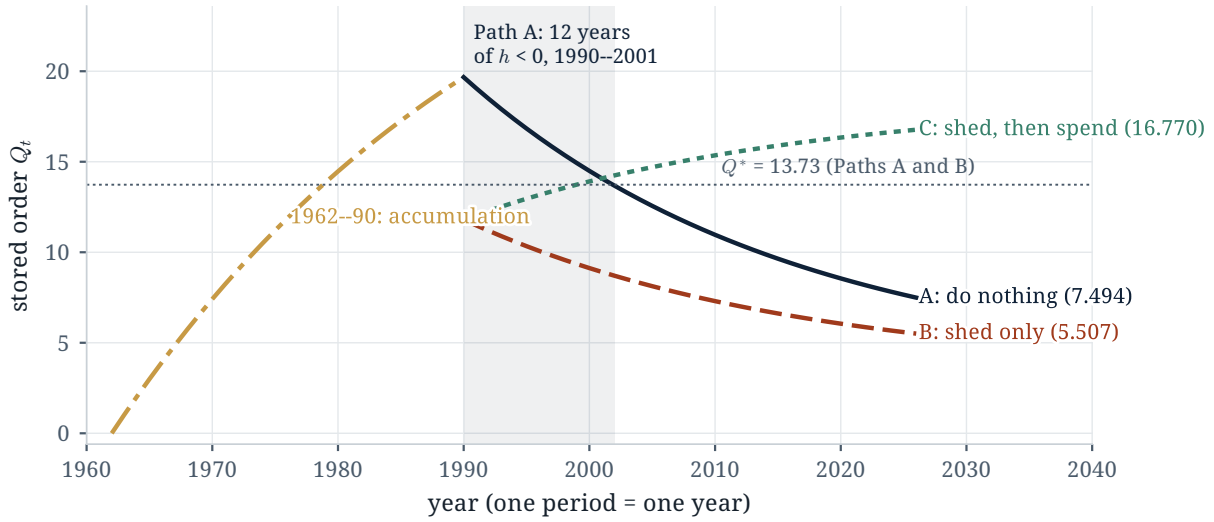


Figure 14.5: **Headroom returns either way; the only choice is whether the stock spent to buy it back was chosen**: the body accumulates to `Q` = 19.669 by 1990 under a ceiling of 1.6362, then carries it under a ceiling of 0.8349, and doing nothing costs 61.9 per cent of the stock over twelve negative-headroom years, while shedding 40 per cent on purpose and then raising `eta` and `u` ends at 16.770. Modelled from the parameters stipulated in Sections 7.3.2 to 7.3.6; generated by `figs/src/CCP-014_fig04.py`.

7.3.7 The three paths compared at 2026

At 2026	A: do nothing	B: shed only	C: shed, then spend
Years with `h < 0`	**12**	0	0
`Q` in 2026	7.494	5.507	**16.770**
`h` in 2026	+0.2339	+0.3084	+0.3695
`Q` lost from 1990 peak	61.9%	72.0%	**14.7%**
`Q` vs post-shed base	—	-53.3%	**+42.1%**
Loss chosen?	No	Yes	Yes
Protected set honoured?	No	Yes	Yes

Path A lost more than Path C and chose none of it. Path B was disciplined and inert. **Path C ends with 2.24 times Path A's stored order and higher headroom, having deleted 40 per cent of its rule stock on purpose in a single year.**

7.3.8 Falsification, declared in advance

1. If the dependency-weighted stock `Y` does not decay at approximately `delta = 0.015` per year during a period of suspended adoption, the `delta` estimate is wrong and every projection above is void.
2. If `C_t` does not regress significantly on `Q_t` — that is, if `kappa` is not distinguishable from zero — then Route B is not operating in this body and the shedding case loses most of its arithmetic force. **`kappa` must be estimated, not assumed.**
3. If, after the `eta` intervention, adopted standards per working-group hour does not rise, `eta` did not move and the intervention failed regardless of what the process documents say.

4. If measured `u` rises after the membership change but `Y` does not, then either `eta` is lower than estimated or the new members are duplicating existing constituencies — run the redundancy audit before concluding either.
5. The shuffled-input control is mandatory: if `Y_{t+1}` is predicted as well from a shuffled `U` series as from the real one, the model is fitting a trend and the entire analysis must be reported as unsupported.

7.3.9 The rule extracted

The standards body of 1990 was not badly run. It was a successful institution carrying the stored order of a world that had changed. Its ceiling had fallen by 49.0 per cent — from 1.6362 to 0.8349 — and its maintenance and operating load had not. Nobody in the building could see it, because `Q` is invisible while it works and the maintenance term has no advocate. What they could see was that everything took longer, that proposals stalled, and that morale was poor. They diagnosed culture. They were looking at arithmetic.

14.7.4 7.4 Results summary

Result	Statement	Status
R1	At any steady state, headroom equals unused variety: `h_inf = eta(ln N − F(u*))`. At `eta = 0.35`, `delta = 0.015`, `N = 6`, `u* = 0.55`, `C = 0.42`: `Q_inf = 2.840970` and `h_inf = 0.164501`, agreeing to six decimals	DERIVED, checked ARITHMETICALLY
R2	Coordinator to medium: `eta` 0.35 to 0.55 and `delta` 0.25 to 0.12 move the fixed point from `Q_inf = 0.2653` to 3.9281, a factor of 14.8; a turnover event costs 70 per cent of stock against 15 per cent	ILLUSTRATIVE MODEL OUTPUT from stipulated parameters
R3	At per-path cost 0.02 and `eta = 0.5`, headroom peaks at `n = 5` and goes negative between 11 and 12, so a flat body of 24 has `h = −3.931` while four cells of six plus a four-delegate federation gives `+2.957`	ARITHMETIC from two stipulated parameters
R4	A panel at `u = 0.12` has `N_eff = 1.96` and `Q_inf = −0.6510`; protocols raising `u` to 0.71 give `Q_inf = +1.9144` and lift accuracy from 65.8 to 74.3 per cent at unchanged individual accuracy `p = 0.62`	ILLUSTRATIVE MODEL OUTPUT; the accuracy figures depend on the stipulated `rho`

Result	Statement	Status
R5	The standards body: `Q` peaked at 19.669 as the ceiling fell from 1.6362 to 0.8349, giving `h = −0.2227`; doing nothing bought headroom back only by losing 61.9 per cent of stored order over twelve negative-headroom years	ILLUSTRATIVE MODEL OUTPUT from stipulated parameters
R6	At a fixed ceiling `eta ln N = 1.50`, a body with `delta = 0.02` and `Q = 75` spends 100 per cent of its surplus on maintenance before any operating cost	ARITHMETIC
CR	Two levers that do not pay: cutting decay alone by raising trust moves a failing body from `Q_inf = −0.0836` to `−0.2114` while reciprocity alone gives `+0.7394`; and a deliberate 40 per cent sunset without any conversion work ends in 2026 with less stored order (5.507) than doing nothing at all (7.494)	ILLUSTRATIVE MODEL OUTPUT; both are counter-results against the paper's own recommendations

14.7.5 7.5 Six discrepancies in the source material, printed rather than corrected

The worked cases above were re-derived from their stated parameters before being reported. Every trajectory reproduced to the printed precision. Six statements *about* the arithmetic did not, and they are listed here rather than quietly repaired, because a reader checking the corpus will find them.

1. **The free-riding recovery is described as 88 per cent and is 68.3 per cent.** Monitoring moves `Q_inf` from 1.7198 to 3.2281. That is an increase of 87.7 per cent *on the free-ridden level*, and a recovery of `(3.2281−1.7198)/(3.9281−1.7198) = 68.3` per cent *of what free riding had taken*. The source reports the first number under the second description. Section 4.8 and Section 7.1.6 state both quantities separately.
2. **The redundancy audit is described as removing "a quarter" of the categorical variety and removes a third.** `N_eff = 5.34` against `N = 8` is a 33.3 per cent reduction in effective count; the reduction in *surplus* is 19.4 per cent. Neither is a quarter. Section 7.1.2 states both.
3. **Regime A's recovery from the turnover event is given as approximately six months and is 6.9 months.** From `Q = 0.072` under `Q_{t+1} = 0.75Q + 0.066329`, the pre-shock level 0.239 is regained at `t = ln(0.026316/0.193316)/ln(0.75) = 6.93` periods.
4. **Regime B's recovery is given as under one month and is 1.7 months.** From `Q = 2.138` under `Q_{t+1} = 0.88Q + 0.471374`, the stock is 2.353 after one period and 2.542 after two, so the pre-shock 2.515 is regained between the first and second period. The qualitative contrast survives and is in fact a factor of four rather than a factor of six.
5. **A rounding inconsistency at the fourth decimal in the panel's identity check.** The source prints `h at Q_inf = +0.1342` against `eta ln(N/N_eff) = +0.1341`, which reads as a 0.0001

disagreement in an identity that holds exactly. Carried without intermediate rounding both sides are 0.134148. The identity is exact; the discrepancy is an artefact of rounding `S` to four decimals in one expression and not the other.

6. **A parameter the source never supplies.** The tool library's O8 score is justified by two craft circles exceeding the headroom-positive band computed with "its own per-path cost", and that per-path cost is never given anywhere in the case. The band quoted is the reference band from the stipulated 0.02, which S4.4 explicitly forbids adopting. Section 7.1.7 records the gap; the O8 score of 1 is therefore not recomputable from the case as stated.

Two further points of arithmetic are correct and are noted because they look wrong. The headroom at 1990 under the Era 2 parameters is `−0.222735`, printed in the source as `−0.2226`; the fourth decimal differs because `0.8349` and `0.7626` are themselves rounded, and the unrounded computation gives `−0.2227`, which is what is printed above. And the identity check in Path C returns `h_inf = +0.3065` against a 2026 headroom of `+0.3695`; these differ because the trajectory has not yet converged at 2026, the time constant being 26.7 years.

14.8 8 Discussion

14.8.1 8.1 What follows

Three design problems are one problem. The tool library, the review panel and the standards body differ in scale by three orders of magnitude and in subject matter entirely, and the same five parameters describe all three. What differs is which parameter is binding. The library's fault was `delta`, held high by storing state in a person. The panel's fault was `u`, held low by an order of operations that guaranteed a cascade. The standards body's fault was `Q` itself, inherited under a ceiling that had fallen by 49 per cent. **Naming the binding parameter is the whole of the diagnosis, and the gauge names it before any money is spent.**

The identity turns headroom into a difference measure. `h_inf = eta ln(N/N_eff)` says that a settled system's spare capacity is exactly the redundancy it is carrying. That has an uncomfortable corollary for anyone reporting a healthy headroom figure: **at equilibrium, high headroom is not a sign of health but a measure of unused difference.** The panel of Section 7.2 is the instance — headroom `+0.2448`, and a negative steady state, because the budget was there and nothing was arriving to spend it on.

Repetition cannot fix a fixed point. The steady state contains no time index, so no amount of doing the same thing better moves it; only a parameter change does. That result is stated sharply elsewhere in this corpus — refinement cannot move the fixed point, since `Q_inf = (eta F − c)/delta` contains no `t`, and the refinement ladder is silent in pitch because the relative frequency depends only on `u` and `N` [Curzi, 2026]. Every "try harder" response to a negative headroom reading is an attempt to move a fixed point by iteration, and it fails for a reason that can be written in one line.

Cost can be moved between participants without changing the total drive. The nesting argument of Section 4.6 splits a body's coordination cost across cells, and the natural objection is that the cost has only been relocated. The relevant invariance is exact: the sum of drives `b_i = eta_i F − c_i` is invariant under any transfer of cost among participants, verified in a different domain at `+0.643688229683` and `+1.268940530229` with residual `0.000e+00` [Curzi, 2026]. **Nesting does not work by moving cost; it works by removing paths**, which is why the arithmetic of Section 4.6 turns on `n(n−1)/2` and not on who pays.

The negative fixed point has a frequency reading. A body whose numerator is negative has a `Q_inf < 0`, and the corresponding statement in the frequency formulation of this corpus is a trichotomy: contact above, at, or below a minimum sustainable frequency `f_min = k e^(2c/eta)` corresponds exactly to `Q_inf` positive, zero or negative, with `k = 0.483425732988414` a **defined** constant [Curzi, 2026]. The two readings are the same inequality in different coordinates, and the second is more useful where contact is episodic rather than continuous

— a working group that meets twice a year is a frequency problem before it is a conversion problem.

14.8.2 8.2 What does not follow

Nothing here shows that any organisation obeys the storage law. All three worked systems are constructed. The one preregistered confirmatory test of the law reported in this corpus did not meet its own declared margin [Curzi, 2026i], and that result — not the arithmetic above — is the current evidential status of every applied claim in this paper.

Centrelessness is not shown to be better. It is shown to have a different parameter profile. Section 4.9 gives the case where the same topology produces outcomes nobody would design, and Section 4.11 gives four things a shared medium does not do. **A structure can answer real problems and still add nothing to the quantity the framework cares about**: the corpus's own worked case of a governance vehicle that answers ownership recording and succession, gives no protection against arbitrary state action, and **adds nothing to the effective count** is the clearest available warning [Curzi, 2026], and it bounds what Section 4 may claim for any structural move.

Raising the effective count is not the same as raising a related exposure measure, and the two can move in opposite directions. A worked case elsewhere in this corpus reports a move that lowered perimeter exposure to 13.6 per cent while *lowering* the effective bloc count from 2.94 to 2.34, and reports a printed computation of its own found wrong and corrected rather than silently amended [Curzi, 2026]. Another reports three legs of a diversification argument each failing against itself: an operating leg that lowered the effective count it was proposed to raise, a leg that fails outright once its maintenance cost reaches 0.38, and a funding gain cut from 3.175 to 2.410 effective blocs by a correlation adjustment applied against the argument [Curzi, 2026]. **Both bound this paper directly.** The second bounds Section 6.5: a maintenance cost of 0.38 is enough to destroy a positive case, and the maintenance table there is not a licence to assume the term is small. The third bounds Section 5.2: a correlation adjustment is exactly what Model C1 performs, and the honest expectation is that it *reduces* the count, not that it leaves it intact.

Zero decay is not the safe end of the range. Section 4.1 gives the entomological form and Section 6.7 the institutional one, but the underlying result belongs to the contraction: a map with gain exactly one is an isometry that orbits without arriving, and a map with gain at or above one is an amplifier whose corrections, sized to delayed feedback, arrive as new errors [Curzi, 2026]. The standards body's decision to drive `delta` from 0.030 to 0.015 — reasonable, and taken as a service to implementers — is a step in that direction, and Section 7.3.4 prices it: the same `delta` that slowed the loss also slowed the recovery, stretching twelve negative years out of what a larger `delta` would have compressed.

The independence protocols do not manufacture difference. They extract the signal present. FM-5.6 states the limit: where the signals are duplicates there is nothing to extract, and the composition change in Section 7.2.5 is doing work the protocols cannot do. Reported gains from elicitation in a homogeneous body should be read as measurement artefacts until composition is audited.

14.8.3 8.3 Relation to the surrounding literature

The commons literature supplies what this paper's arithmetic cannot: evidence that the arrangements exist. Ostrom [1990] is a comparative field study, and its principles are propositions about which features co-occur with endurance, audited across ninety-one subsequent studies with broad but uneven support [Cox et al., 2010]. **This paper adds a price and not a proof:** it computes what O8 is worth in headroom units under stated parameters, and computes the monitoring trade of O4 and O5. Where Ostrom's list says nesting helps at scale, Section 4.6 says at what size and by how much, conditional on a per-path cost that any installation must measure for itself.

The aggregation literature is where the framework's u finds its sharpest external counterpart. The independence clause of the jury theorem, the cascade result [Bikhchandani et al., 1992, Banerjee, 1992], the measured collapse of crowd wisdom under social influence [Lorenz et al., 2011], and the experimental demonstration that early visible choices lock in [Salganik et al., 2006] all say the same thing about the same quantity. The diversity-and-ability result [Hong and Page, 2004] and its critique [Thompson, 2014] are used here only for the mechanism, and the subsequent exchange [Kuehn, 2017, Singer, 2019, Grim et al., 2019] is cited so that a reader is not left with the impression that the formal result is settled. The construct warning [Harrison and Klein, 2007] is the one the corpus most needs and least often pays.

The institutional half draws on three separate literatures that do not usually meet. The infrastructure result [Star and Ruhleder, 1996] explains why the maintenance term has no advocate. The measurement-corruption results [Goodhart, 1975, Campbell, 1979] explain why any proxy for Q will be gamed and why a difference gauge resists it better than a total. And the organisational-learning results [March, 1991, Levinthal and March, 1993, Cyert and March, 1963] supply the mechanism by which a successful body loses the variance that made it capable — which is postulate I1 arriving from the other direction, in a literature that measured it.

Two places where the surrounding literature cuts against the corpus are stated rather than buried. Putnam [2007] reports that diversity in the short run reduces trust and civic participation — a result from a scholar who did not want it, and the largest single piece of contrary evidence for the corpus's preference for raising u; the rebuttal [Abascal and Baldassarri, 2015] qualifies it without dissolving it. And the empirical work on surface-level diversity harming short-run performance [Mannix and Neale, 2005, van Knippenberg and Schippers, 2007] is not answered by pointing out that u is a different construct, unless u is actually measured rather than inferred — which is why Section 5.9 excludes the inference by rule and why S5.7 requires the measurement annually.

14.9 9 Verification

The results of this paper compound from seven postulates. Each is atomic, each is checkable on its own, and each carries its status. **Where a postulate could be dropped and the result still follow, it has been dropped**; what remains is what the chain actually consumes.

14.9.1 9.1 Dimensional audit

The equations must balance before anything else about them is worth checking. `[Q]` denotes the dimension of stored order and `[T]` the dimension of time; the period is the unit of `[T]` and every rate below is *per period*, which is why no explicit `[T^-1]` appears once the recursion is written in discrete steps.

Symbol	Dimension	Unit	Status
`Q_t`	`[Q]`	proxy units (tools in order, rubric points, dependency-weighted standards)	**measured**, in whatever unit the declared order proxy counts
`U_t`, `u`	`[1]`	none	**defined** dimensionless: `u = 1 - (x.y)^2` with `x`, `y` unit vectors
`N`	`[1]`	a count	**measured** as a count of genuinely distinct categories

Symbol	Dimension	Unit	Status
`N_eff`	`[1]`	a count	**derived** from `N` and `u`; **not a Hill number**
`S_t = F(U_t)`	`[1]`	nats	**derived** dimensionless: a logarithm of an effective count
`eta`	`[Q]` per nat per step	proxy units per nat per period	**measured**; *not* dimensionless in this framework
`delta`	`[1]`	per period	**measured** dimensionless fraction per step
`C_t`, `c_0`	`[Q]` per step	proxy units per period	**measured**
`kappa`	`[1]` per step	per period	**derived** from a regression of `C_t` on `Q_t`; dimensionless per step, since `kappa Q` must carry `[Q]` per step
`h_t`	`[Q]` per step	proxy units per period	**derived**
`Q_inf`, `Q*`	`[Q]`	proxy units	**derived**: `[Q]`/step divided by `[1]`/step returns `[Q]`
`tau`	`[T]`	periods	**derived**: `1/delta` or `1/(delta+kappa)`
`n`	`[1]`	a headcount	**measured**
`c_path`	`[Q]` per step per path	proxy units per period per communication path	**assumed**; the source stipulates 0.02 and never measures it
`phi`	`[1]`	a fraction	**measured** free-riding fraction
`lambda`	`[Q]` per step	proxy units per period per unit of `phi`	**assumed**
`p`, `rho`	`[1]`	probabilities	`p` **measured** from outcomes; `rho` **assumed**

Do the equations balance? The storage law balances if and only if every term on the right carries `[Q]`. The retention term does, since `delta` is dimensionless. The cost term does by construction. The drive term does only if `eta S_t` carries `[Q]`, and since `S_t` is a dimensionless count in nats, **`eta` must carry `[Q]` per nat per step**; it is therefore not dimensionless, and every `eta` quoted in this paper as a bare number — 0.35, 0.55, 0.45, 0.62, 0.38, 0.60 — is quoted with its unit suppressed. **That is a silence in the source and it is recorded here rather than quietly filled**: the tool library's `eta = 0.55` should read 0.55 proxy units per nat per month.

The headroom gauge balances because `ln N` is dimensionless, so `eta ln N` carries `[Q]` per step and is commensurable with `delta Q_t` and with `C_t`. **A headroom calculation in which `C` is denominated in hours or currency while `Q` is denominated in rules is dimensionally malformed however arithmetically tidy it looks**, which is the denomination requirement the survey protocol imposes [Curzi, 2026i]. The nesting expression `h(n) = eta ln n − c_path n(n−1)/2` balances only if `c_path` carries `[Q]` per step per path, which is what the audit assigns it; the source states 0.02 with no unit at all.

The cost-in-stock model `C_t = c_0 + kappa Q_t` balances only if `kappa` is dimensionless per step. It follows that `delta + kappa` is a per-step dimensionless rate and that `tau = 1/(delta + kappa)` carries `[T]` — which is why the two can be added at all, and the addition is not a convenience but a dimensional requirement that the model happens to meet.

Where the framework has no operational unit, it is said here. `Q` is stored order, and stored order has no natural unit: it is counted in whatever the declared proxy counts, and two systems' `Q` values are not comparable unless they declare the same proxy. Comfort itself, the underlying construct of the corpus, has **no operational unit at all** and is treated as a dimensionless magnitude by stipulation rather than by measurement.

A symbol collision, disclosed. `phi` denotes the free-riding fraction throughout Sections 4.8 and 7.1.6 and has no relation to the golden ratio, which appears only in the defined constant `k = (9/8)^2 / varphi^2` in Section 8.1 and is written `varphi` there for that reason.

14.9.2 9.2 The postulates

P1 — The recurrence is linear and first order. `Q_{t+1} = a Q_t + x_t` with `a = 1 − delta` and `x_t = eta S_t − C_t` not depending on `Q`. *Status:* **DEFINITION** — the model form adopted by the framework. It carries no empirical weight on its own, and Section 10 records that its application to any real system is a postulate.

P2 — `0 < delta <= 1` implies `0 <= a < 1`. *Status:* **ARITHMETIC.** Subtracting a number in `(0,1]` from 1 gives a number in `[0,1)`.

P3 — A contraction on a complete metric space has exactly one fixed point, reached from any start. *Status:* **CITED** — [Banach, 1922], with the elementary difference-equation route in [Elaydi, 2005].

P4 — Surplus per period is bounded: `0 <= S_t <= ln N`, with `S = ln N` if and only if the states are orthogonal. *Status:* **CITED** — the range and extremality results of the surplus functional [Curzi, 2026], following from `F(u) = ln(1+(N−1)u)` on `u` in `[0,1]`.

P5 — The number of distinct communication paths in a body of `n` participants is `n(n−1)/2`. *Status:* **ARITHMETIC**, and **CITED** for the claim that this count is what organisational coordination cost actually tracks [Brooks, 1975, Conway, 1968].

P6 — Under a shared-signal mixture with weight `rho`, majority accuracy is `rho p + (1−rho) P_ind(n,p)`; and under sequential observation of actions rather than signals, the probability that the cascade which forms is correct is `p(1+p)/[2(p^2−p+1)]`. *Status:* the mixture is a **MODELLING CHOICE** made in this paper, transparent and crude; the cascade expression is **CITED** [Bikhchandani et al., 1992] and its arithmetic is **ARITHMETIC**, checkable from the three branch probabilities, which sum to 1 identically.

P7 — Operating cost may be affine in the stock: `C_t = c_0 + kappa Q_t` with `kappa >= 0`. *Status:* **POSTULATED** — this is an assumption about how institutions charge overhead. It has no external support in this bibliography, it is used at exactly two places (R5 and the shedding return `dh/dQ`), and the falsification rule of Section 7.3.8 requires `kappa` to be estimated rather than assumed. **Every result not involving shedding or Path A/B/C holds with `kappa = 0`.**

14.9.3 9.3 The compounding chain

```
P1 + P2        -> L1   T(Q) = aQ + x has |T(Q)−T(Q')| = a|Q−Q'| with a < 1,
                       so T is a contraction.                          ARITHMETIC

L1 + P3        -> L2   T has exactly one fixed point
                       Q_inf = x/(1−a) = (eta S − C)/delta.            CITED (banach1922operations)

L2             -> R1   at the fixed point delta Q + C = eta S, so
                       h_inf = eta ln N − eta S = eta ln(N/N_eff)
                         = eta (ln N − F(u*)).                         ARITHMETIC
```

```
R1 + P4        -> R1a  h_inf >= 0 always, with equality iff u* = 1;
                       therefore h < 0 is never a steady state.        ARITHMETIC

R1             -> R1b  delta, C and Q are absent from h_inf, so a
                       linear system with constant parameters settles
                       at positive constant headroom and does not
                       ossify.                                         ARITHMETIC

P5 + P4        -> R3   h(n) = eta ln n - c_path n(n-1)/2 has
                       dh/dn = eta/n - c_path(n - 1/2), one interior
                       maximum and one positive zero; at eta = 0.5,
                       c_path = 0.02 the maximum is n = 5 and the zero
                       lies in (11,12).                                ARITHMETIC

L2             -> R6   the maintenance share is delta Q /(eta ln N),
                       which reaches 1 at Q = eta ln N / delta; at
                       eta ln N = 1.50 and delta = 0.02 that is
                       Q = 75.                                         ARITHMETIC

L2 + P7        -> R5a  with C affine in Q the map is
                       Q_{t+1} = (1 - delta - kappa) Q_t + eta S - c_0,
                       so Q* = (eta ln N - c_0)/(delta + kappa) is the
                       zero-headroom stock and tau = 1/(delta+kappa).  ARITHMETIC

R5a + P7       -> R5b  dh/dQ = -(delta + kappa): every unit shed
                       returns delta + kappa units of headroom
                       per period, permanently.                        ARITHMETIC

L2             -> CR1  for eta F(u*) < c*, dQ_inf/d(delta) > 0, so
                       lowering delta alone lowers Q_inf.              ARITHMETIC

P6             -> R4a  P(correct cascade) tends to a constant in n,
                       and at p = 0.62 equals 0.6570, which is below
                       P_ind(9, 0.62) = 0.7738.                        CITED (bikhchandani1992theory)

R4a + P6       -> R4b  a mixture at rho = 0.75 returns 0.6584 and at
                       rho = 0.20 returns 0.7430, so accuracy moves
                       8.5 points with p held at 0.62.                 ARITHMETIC on a MODELLING CHOICE
```

Every line is checkable from its own inputs. **R1, R1a, R1b, R3, R5a, R5b, R6 and CR1 are arithmetic consequences of one definition (P1), two arithmetic facts (P2, P5), one cited theorem (P3) and one cited range result (P4).** Only R5a and R5b consume P7, the one postulate with no external support, and the falsification rule of Section 7.3.8 puts it at risk. R4a is cited; R4b consumes a modelling choice and is labelled as such wherever it is quoted.

What this chain does not establish. It does not establish that any organisation obeys P1. That is the modelling postulate of Section 3.5, it is not verified here, and the one preregistered confirmatory test of it reported in this corpus did not meet its own declared margin [Curzi, 2026i]. **A chain whose steps are ARITHMETIC and CITED delivers the mathematics with certainty and delivers nothing whatever about institutions.**

One further line, separated because it is a stronger claim than the one before it. The constant `k = (9/8)^2 / varphi^2` is **DEFINED**: `(9/8)^2 = 1.265625` exactly, `varphi^2 = varphi + 1 = 2.618033988749895`, and the quotient is `k = 0.483425732988414` — three checks a reader can perform in a minute, and the fifteen figures are legitimate because the quantity is exact rather than measured. **Attaching the unit Hz to k, as the minimum-sustainable-frequency result does**

[Curzi, 2026], is a separate and much stronger claim, and its status is POSTULATED. Nothing in the definition of `k` supplies a time scale; the unit is asserted. And because `varphi` appears at all, the sceptical literature belongs beside it: claims of golden-ratio structure in natural and designed systems are frequently unsupported on inspection [Markowsky, 1992], and this paper uses `k` only as a defined reference constant in a cross-reference, never as evidence.

14.9.4 9.4 The identity, checked on the smallest case

The single most useful check a reviewer can perform is the identity of Section 3.2 on a case small enough to do by hand. Take `eta = 0.35`, `delta = 0.015`, `N = 6`, `u* = 0.55`, `C = 0.42`.

```
N_eff = 1 + 5(0.55)      = 3.75
S     = ln(3.75)         = 1.321755839982320
eta S = 0.35 x 1.321755839982320 = 0.462614543993812
Q_inf = (0.462614543993812 − 0.42) / 0.015 = 2.840969599587455
```

By the gauge, directly:

```
h_inf = 0.35 x ln 6 − 0.015 x 2.840969599587455 − 0.42
      = 0.35 x 1.791759469228055 − 0.042614543993812 − 0.42
      = 0.627115814229819 − 0.462614543993812
      = 0.164501270236007
```

By the identity:

```
h_inf = eta (ln N − S) = 0.35 x (1.791759469228055 − 1.321755839982320)
      = 0.35 x 0.470003629245735
      = 0.164501270236007
```

The two agree to fifteen figures, which is machine precision, and they agree because the second is an algebraic rearrangement of the first and not an approximation to it. Rounded to the six decimals quoted in R1: `Q_inf = 2.840970` and `h_inf = 0.164501`. **ARITHMETIC.**

The equilibrium can be confirmed independently by summing the series rather than solving for the fixed point. With `x = 0.042614543993812` and `a = 0.985`, the running total `x(1−a^t)/(1−a)` gives 0.042615 at `t = 1`, 0.084590 at `t = 2`, 0.398502 at `t = 10`, 1.506594 at `t = 50`, 2.214226 at `t = 100`, 2.702704 at `t = 200` and 2.840954 at `t = 800`, converging on `x/delta = 2.840970`. **The equilibrium is not an assertion; it is the sum of a geometric series, and the sum can be checked by hand.**

14.9.5 9.5 Four secondary checks

The nesting crossing. `h(n) = 0.5 ln n − 0.01 n(n−1)`. At `n = 11`, `0.5 x 2.397895 − 0.01 x 110 = 1.198948 − 1.100 = +0.098948`; at `n = 12`, `0.5 x 2.484907 − 0.01 x 132 = 1.242453 − 1.320 = −0.077547`. The crossing is therefore in `(11, 12)`; the continuous extension has its root at `n = 11.577`, and linear interpolation between the two integer values gives 11.561, which is the accuracy that interpolation is worth here. And `h(5) = 0.804719 − 0.200 = +0.604719` exceeds both `h(4) = +0.573147` and `h(6) = +0.595880`, so the maximum is at `n = 5` among integers. **ARITHMETIC.**

The cascade branch probabilities sum to one. `p(1+p)/2 + (1−p)(2−p)/2 + p(1−p) = (p + p^2 + 2 − 3p + p^2)/2 + p − p^2 = (p^2 − p + 1) + p − p^2 = 1`, identically in `p`. This is the check that establishes the denominator of the correct-cascade expression, and it can be done in two lines. **ARITHMETIC.**

Twelve negative-headroom years. Path A's map is `Q_{t+1} = 0.9625 Q_t + 0.126717` with fixed point 3.3791, so `Q_t = 3.3791 + 16.2899 x 0.9625^t` from `Q_0 = 19.669`. Setting `Q_t = Q* =`

13.7307 gives 0.9625^t = 0.63546 and t = 11.87. **The stock is above the zero-headroom line for the whole of years 0 to 11 and crosses during year 12**, so the count of completed years at h <= 0 is exactly twelve, 1990 to 2001 inclusive, and headroom is positive from 2002. **ARITHMETIC.**

The shedding return. With h = eta ln N − delta Q − c_0 − kappa Q, dh/dQ = −(delta + kappa) = −(0.015 + 0.0225) = −0.0375, so shedding 8 units returns 0.30 per period, which against the ceiling 0.8349 is 0.30/0.8349 = 35.9 per cent of the budget. **ARITHMETIC**, conditional on P7.

14.9.6 9.6 Limits, asymptotics, and the model's valid domain

delta −> 0 with kappa = 0. tau −> infinity, the fixed point diverges, and at exactly delta = 0 it does not exist: the map is an isometry, every point is fixed, and the uniform bound degrades to Q_0 + t eta ln N, linear in t. **The limit is singular and the two statements are different** — divergence and non-existence are often conflated. This value is outside the model's valid domain and is the failure mode FM-4.9 names.

delta −> 0 with kappa > 0. The fixed point survives, because the effective rate is delta + kappa > 0. **This is a substantive difference the affine-cost model introduces**: an institution with a cost proportional to its stock cannot lose its attractor by driving decay to zero, which is a mathematically reassuring result and an operationally miserable one, since the attractor it retains is (eta S − c_0)/kappa.

delta + kappa −> 1. The system has no memory: Q_{t+1} −> eta S − c_0, an input-output map with no state. **At that boundary the storage law is not a storage law**, and any claim about accumulation is vacuous there.

eta −> 0. The drive vanishes and Q_inf = −C/delta <= 0. **Zero conversion decouples the input from the stock entirely**, so no interaction geometry of any quality changes the outcome. This is Route A taken to its limit and it is why S4.7 exists.

u −> 0. N_eff −> 1, S −> 0, and h_inf −> eta ln N, its maximum. **The maximally redundant body has the maximum equilibrium headroom and the minimum equilibrium stock**, which is exactly the panel of Section 7.2 and is the sharpest single statement the identity makes. **u −> 1.** N_eff −> N, S −> ln N, h_inf −> 0: the ceiling is attained and no further gain is available from geometry.

N −> 1. ln N = 0 and F(u) = 0 for every u. **A single category has zero available surplus, not a small amount**, so Q_inf = −C/delta <= 0 at any positive cost. This is the boundary that makes the redundancy audit consequential.

N −> infinity. ln N −> infinity, so the ceiling grows without bound but only logarithmically, with marginal gain ln N − ln(N−1) ~ 1/N. Against this, coordination cost grows as n(n−1)/2. **The two curves cross, and beyond the crossing additional categories cost more than the ceiling they buy** — which is Section 4.6, and is the reason N is the most expensive and least rewarding of the four levers.

rho −> 1. Model C1 returns P = p exactly: the body performs as one member, which is u −> 0 in the other formalism. **n −> infinity under a cascade.** Accuracy does not tend to 1 but to p(1+p)/[2(p^2−p+1)], a constant less than the nine-voter independent value for every p in (0.5, 1). **Aggregation without independence has a ceiling, and the ceiling is independent of n.**

Stated domain. The results of this paper are valid for delta + kappa in (0, 0.5], eta > 0, N >= 2, u in [0,1], p in (0.5, 1) for the aggregation results, rho in [0,1], and C_t <= (1−delta)Q_t + eta S_t each period. Outside that domain the equations still have solutions and those solutions should not be interpreted as statements about a stock.

14.9.7 9.7 Invariance, and how far the scale claim has been tested

The claim that the same equations describe a 48-household library, a nine-person panel and a sixty-four-year standards body is an **invariance of form under a change of unit and a change of system size**, and it is worth stating as one. Let Q be rescaled by a positive constant c. The law

is invariant provided `eta`, `C` and `lambda` rescale with it — `eta' = c eta`, `C' = cC` — while `delta`, `kappa`, `N`, `u`, `p` and `rho` are unchanged, since `delta` and `kappa` are dimensionless per-step rates and `S` is a dimensionless count. **`delta`, `kappa`, `u`, `N`, `p` and `rho` are therefore the unit-free quantities of this paper, and `Q`, `eta`, `C`, `c_path` and `lambda` are the unit-carrying ones.** Two systems may be compared directly on the first group and not on the second unless they share a proxy.

The identity `h_inf = eta ln(N/N_eff)` is **not** scale-invariant in this sense: it carries `[Q]` per step through `eta`, so a headroom figure of `+0.164501` means nothing without the proxy that denominates it. **The dimensionless part of the identity is `ln(N/N_eff)`, and that is the quantity two systems can actually be compared on.** This is worth stating because headroom figures are quoted across the corpus as though they were commensurable and they are not.

What would break the invariance: a `delta` that varies with `Q`, which would make the transformation non-linear; an `eta` that saturates at large `S`; a `c_path` that is not constant across paths, which is the realistic case and which the nesting result assumes away; or a proxy whose counting unit is not extensive.

How many scales it has been tested at. Within this paper: three, and all three are constructed rather than observed. Elsewhere in this corpus the same five-parameter survey has been applied from a single desk to a 22-household block, reproducing each system's observed decline with residuals under 1.4 per cent of the ceiling term [Curzi, 2026k]. **That is a small number of systems, one corpus, one author, and no independent replication.** The invariance is a claim about form that has been exercised, not an invariance that has been established.

14.9.8 9.8 Analogy discipline, and the precision of the numbers

Stigmergy is used as an existence proof of a mechanism, not as a derivation. The termite mound and the ant trail establish that coordination through a modified shared medium occurs and is sufficient to build a structure. Nothing about insect physiology, pheromone chemistry or colony genetics is claimed to hold of a tool library, and no biological result's authority transfers. **This is category three in the analogy taxonomy — a suggestive correspondence used as motivation and never as evidence** — with one exception: the identification of trail evaporation with `delta` is category two, a shared mathematical form, because both are first-order exponential decay of a stored signal, and that identity is exact.

The topological fragility argument is category three and is the weakest thing in the paper. Section 4.10 states what it does not establish. No degree distribution was measured; the scale-free label is borrowed; and no power-law or heavy-tail claim is made anywhere, so the distribution-fitting procedure such a claim would require [Clauset et al., 2009] does not apply — but neither does the authority of the result being borrowed from.

The storage law itself is category two. It is the same difference equation that governs a first-order low-pass section, and every theorem about first-order discrete systems applies [Curzi, 2026a]. Nothing about charge, capacitance or energy is claimed to hold of an institution.

The entropy invoked near `ln N` is Shannon entropy, a dimensionless information measure in nats [Shannon, 1948], not thermodynamic entropy, which carries `[energy][temperature]^-1`. They are different objects and the framework identifies surplus with neither.

Precision. The numbers in this paper are of four kinds and are quoted accordingly.

- *Defined, and exact.* `k = 0.483425732988414`, `(9/8)^2 = 1.265625`, `varphi^2 = 2.618033988749895`. Fifteen figures are legitimate because these are exact ratios of exact quantities. **They are labelled defined wherever they appear at that precision.**
- *Derived from defined inputs.* `ln 8 = 2.0794`, `ln 9 = 2.1972`, `ln 14 = 2.6391`, every binomial sum in the Condorcet tables, the cascade probabilities, `h_inf = 0.164501`. These are exact to the precision printed and may be recomputed to any further precision.
- *Illustrative model output.* Every trajectory entry, every `Q_inf`, every headroom figure in Section 7. These are the output of stipulated parameters and carry **no measurement uncertainty because there was no measurement.** Four decimal places there record the arithmetic and not the confidence.

- *Measured or fitted.* **None appear in this paper**, because nothing in it was measured. Any application of this framework to real data must report fitted parameters with intervals, and Section 10 records that the corpus's one such report gives an interval on the retention coefficient and none on several other quantities.

14.10 10 Limitations

The three worked systems are constructed, not observed. No dataset was collected for this paper. Every number in Section 7 is the output of the stated parameters under the stated model, carried through so a reader can re-derive it. They demonstrate that the procedure runs and that it reports its own failures; **they are not evidence that any organisation behaves this way and must not be cited as such.**

Which parameters were chosen rather than measured. In all three systems the whole parameter vector is stipulated: `(N, u, eta, delta, C)` in every case, plus `c_path = 0.02` and `eta = 0.5` in the nesting result, `lambda = 0.9` and `phi = 0.35` and the monitoring cost 0.08 in the free-riding result, `kappa = 0.0225` in the institutional case, and `rho = 0.75` before and 0.20 after in the accuracy result. **The accuracy result is the most exposed of these**: the coincidence in Section 7.2.7 between the panel's 65.84 per cent baseline and the unbounded-cascade value of 65.70 per cent is produced by the choice `rho = 0.75`, and a different choice breaks it. It is a corroboration of internal consistency and not a test of anything.

The correspondence between the equations and any social system is a postulate, and its one preregistered confirmatory test did not meet its declared margin. Curzi [2026i] specifies the protocol and applies it; the fit landed inside the admissible box and its parameters agreed with independent measurement, but it beat a matched autoregression by 2.8 per cent on held-out error against a preregistered requirement of 10 per cent, so the verdict was parameters corroborated, predictive advantage not demonstrated. That is the current evidential status of every applied claim in this paper.

The audit cannot see repackaging, and this paper leans on the very justification the audit cannot enforce. The shuffled-input control passed 98.35 per cent of the time against a third variable driving both interaction and order, while detecting only 48.17 per cent in the clean case [Curzi, 2026i]. Sections 7.1.1, 7.2.1 and 7.3.1 each argue that `Y` is not a repackaging of `U`, on the observables and in advance, because that argument cannot be recovered statistically afterwards. **In a real installation those arguments would be the weakest joint in the whole chain**, and no amount of subsequent testing repairs them.

The reflection problem, and why it outranks everything else here. Section 3.4 states it in full. Applied to a group whose members both influence and constitute the aggregate, the specification is reflexive, and Manski [1993] proves the endogenous social effect is not identified from observational data; Shalizi and Thomas [2011] show the same confounding for network data. **No fit recovers the parameter.** This paper adopts no remedy, so every applied reading — that the medium raised `eta`, that the protocols raised `u`, that `rho` fell — is a **modelling position rather than an identified empirical result**. The most realistic available repair is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001]. Reflexive fixed points demonstrably occur [MacKenzie and Millo, 2003, Merton, 1948]; that they occur does not make this framework's parameters recoverable.

The Lucas critique applied to this framework's own parameters. Section 6.9 asserts that a measure changes the system measured by it and prescribes rotating the proxy in response; the paper then treats `delta`, `eta`, `N` and `kappa` as constants over an estimation window. **Both cannot be true without qualification** [Lucas, 1976]. The position taken is the narrow one: the parameters are treated as stable only over a declared window and only while `Q` is not itself a reported target, which is exactly why S6.1 prohibits reporting `Q` alone. Where `Q` is reported upward, the framework's own reflexivity applies to it and the constancy assumption fails. That is a real restriction on where the model may be used.

`N_eff = 1 + (N−1)u` is not a Hill number and no external literature contains it. `1/HHI` is the

Hill number of order two and `exp(H)` the Hill number of order one, and both are fully grounded [Hill, 1973, Jost, 2006]; `1 + (N−1)u` is a linear interpolation between 1 and `N` and is this framework's own construction, adopted by design axiom S3 [Curzi, 2026]. **A reviewer reaching for the Hill-number literature as support will find that it does not apply.** Every effective count in this paper is the affine one, and the repair is to derive its properties in the framework's own terms rather than to borrow.

`u` has not been shown to be a quantity, and this is the strongest objection to the framework. The whole of Section 5 treats non-parallelism as a measurable magnitude on which differences, ratios and logarithms may be taken. Michell [1997] argues that psychology routinely fails to test whether its attributes are quantitative at all before measuring them, and that the failure is systematic rather than incidental [Michell, 1999]. **The remedy is conjoint measurement — the formal test for additive quantitative structure [Luce and Tukey, 1964, Krantz et al., 1971] — or a Rasch demonstration [Rasch, 1960], and neither has been run for `u` anywhere in this corpus.** Until one is, the ratio-scale treatment of `u` throughout Sections 5 and 7 is a stipulation and not a licensed operation, and the paper concedes it rather than defending it. Relatedly, if `u` is read as a similarity measure over people it cannot be assumed to behave metrically [Tversky, 1977].

The construct question is unresolved. `u` is declared a variety construct in Section 5.1, but three incompatible constructs travel under the word diversity [Harrison and Klein, 2007] and the declaration is a definition rather than a validation. The literature reporting harm from surface-level diversity [Mannix and Neale, 2005, van Knippenberg and Schippers, 2007] and the finding that diversity reduces trust in the short run [Putnam, 2007] are not dismissed by the declaration; they are set aside pending the measurement that S5.7 requires and that no worked case here performs.

The functional form is underdetermined. Section 3.1 gives the alternatives — a sub-linear power law [Latané, 1981] and `n log n` [Briscoe et al., 2006] — and records that nothing assembled here discriminates between them. **The logarithm is chosen for its axiomatic characterisation, not established by measurement**, and Peters and Gell-Mann [2016] is the strongest available warrant precisely because it needs no perceptual claim.

Postulate P7 is the one load-bearing assumption with no external support. The affine cost model `C_t = c_0 + kappa Q_t` produces R5, the shedding return, and the whole of the institutional case. It is plausible and it is asserted. If `kappa = 0` the standards body's Path A and Path B collapse toward each other and the shedding case loses most of its force, which is why the falsification rule of Section 7.3.8 requires `kappa` to be estimated. **A significant positive `kappa` has not been demonstrated anywhere in this corpus.**

Postulate M3 borrows an authority it has not earned. No degree distribution was measured; the robustness asymmetry is imported [Albert et al., 2000] into a setting where the contrary argument [Doyle et al., 2005] applies at least as well. The engineering advice that follows from M3 — deposit state in the medium, run a hub census — survives without it, because the bus-count argument needs no network model at all. **The topological framing is decoration on a sound recommendation and is labelled as such.**

Five works this paper needs and cannot cite. Grassé's 1959 paper on stigmergy, Schelling's *The Strategy of Conflict* (1960) on focal points, Condorcet's 1785 essay on the jury theorem, Gambetta's 1993 study of Sicily, and Weiser's 1991 formulation of disappearing technology are each named in the text and **have no key in this corpus's bibliography.** Rather than attach a citation that does not support the sentence, each is named in prose and the gap is recorded here. The neighbouring works that *are* cited — [Schelling, 1971, 1978, Gambetta, 1988, Galton, 1907, Surowiecki, 2004] — are cited only for what they actually say, which is not the same claim in any of the five cases.

The Monte Carlo verification was not re-run. Section 5.4 reports a simulation of the cascade process at `p = 0.60` over 400,000 trials returning 0.6323 against the analytic 0.6316. The analytic value was recomputed for this paper and is exact; **the simulation is reported as the source states it and was not reproduced.** The agreement is within one standard error, which is consistent but is not an independent check.

No model was fitted, so no model comparison applies — and the corpus's one fit reports none. Nothing in this paper is estimated, so no information criterion, no out-of-sample comparison and no matched null is owed by it. The corpus's one applied fit reports held-out error against baselines and **no AIC or BIC** [Curzi, 2026i], so the parameter count was never paid for. Where the framework is applied to real data, all of it is owed [Nosek et al., 2018, Simmons et al., 2011, Gelman and Loken, 2013].

No independent replication of anything. Every system in this paper and every system in the corpus's applied papers was constructed or surveyed by one author. The reproducibility literature is unambiguous about what that is worth [Open Science Collaboration, 2015, Ioannidis, 2005], and Meehl [1967] observes that directional predictions are nearly costless in the social sciences. This framework does make point predictions with constants, which is the right answer to that objection, but only Curzi [2026i] has so far put one of them at risk.

Six discrepancies in the source material are printed in Section 7.5 rather than corrected, including a percentage described as a recovery that is an increase, a fraction described as a quarter that is a third, two recovery times understated, a fourth-decimal rounding inconsistency in an identity that holds exactly, and a parameter that a specification requires and the worked case never supplies. Those are disclosed because a reader checking the corpus will find them.

What would falsify the account. H1 fails if a measured `delta` does not fall after a shared medium is installed, or if headroom does not rise. H2 fails on a measured coordination cost that does not grow faster than logarithmically in headcount. H3 fails if a panel with strongly correlated errors matches the independent-majority prediction, or if accuracy does not move when the order of operations is changed at fixed composition. H4 fails on a standing body whose measured `u` is stationary or rising with no maintenance protocol. H5 fails on a body at a genuine steady state with negative measured headroom. H6 fails on a measured case with a negative step in which lowering `delta` alone raises the steady state. **The scale claim of Section 9.7 is falsifiable at each scale separately and is not falsifiable in general, which is a genuine weakness of that claim and is stated as one.**

Nothing in this paper is a human-subjects finding. No interventions were run and no participants were recruited. The declination constraint of S4.8 and S5.8 is a design constraint on prospective work, not a report of past practice.

14.11 11 Conclusion

At any steady state of the storage law, headroom equals unused variety:

```
h_inf = eta ( ln N − F(u*) ) = eta ln ( N / N_eff )
```

`delta`, `C` and `Q` all cancel. At `eta = 0.35`, `delta = 0.015`, `N = 6`, `u* = 0.55` and `C = 0.42` the stock converges to `Q_inf = 2.840970` and the identity returns `h_inf = 0.164501` by both routes, agreeing to machine precision (**R1**). Three consequences follow and are the statements downstream work should quote rather than paraphrase: a settled system at `u = 1` has zero headroom and is running at its ceiling; **a settled system never has negative headroom, so negative headroom is always inherited stock under changed parameters**; and **the plain linear model, left alone, does not ossify** — maintenance load rises with success and then stops, so ossification requires either exogenous accretion of `Q` or a cost term superlinear in `Q`, and a diagnosis must say which.

Where state is kept is a parameter setting. The same tool library run through a coordinator converts at `eta = 0.35` with `delta = 0.25` and settles at `Q_inf = 0.2653`; run on a shared written medium it converts at 0.55 with `delta = 0.12` and settles at 3.9281, a factor of 14.8, at lower total cost and with no change of people. A turnover event costs the coordinator regime 70 per cent of its stock and the medium regime 15 per cent (**R2**).

Nesting is priced, not preferred. With a per-path cost of 0.02 and `eta = 0.5`, headroom at `Q = 0` peaks at five members (`+0.6047`) and crosses zero between eleven and twelve, at `n = 11.577`

on the continuous extension. A flat body of 24 sits at `h = −3.931` and cannot build anything at any interaction geometry; four cells of six plus a four-delegate federation reaches `+2.957` under the stated independence assumption (**R3**).

Independence is a budget line. A nine-person review panel at `u = 0.12` has `N_eff = 1.96` and a negative fixed point, `Q_inf = −0.6510`; it needs `u >= 0.5280` merely to hold. Five procedural protocols costing 0.090 per item, plus two composition changes, raise `u` to 0.71 and `Q_inf` to `+1.9144`, and lift panel accuracy from 65.8 to 74.3 per cent **with no change in the individual accuracy of any reviewer**, which stays at `p = 0.62` (**R4**). At that individual accuracy an unbounded informational cascade tops out at 0.6570 — below nine independent voters at 0.7738 — so the panel's baseline was the accuracy of an infinite queue of people who never looked at each other's signals.

The maintenance term is the bill for past success. At a fixed ceiling `eta ln N = 1.50`, a body with `delta = 0.02` and `Q = 75` spends 100 per cent of its surplus on maintenance before any operating cost (**R6**). A standards body traced from 1962 accumulates to `Q = 19.669` by 1990 under a ceiling of 1.6362, then carries it under a ceiling of 0.8349 after its constituencies consolidate and its conversion rate falls, giving `h = −0.2227`. Doing nothing returns headroom in 2002 — **twelve consecutive negative-headroom years, 1990 to 2001 — at a cost of 61.9 per cent of everything accumulated in twenty-eight years, none of it chosen** (**R5**).

Two levers do not pay, and both are reported (CR). Cutting decay alone by building trust moves a failing body from `Q_inf = −0.0836` to `−0.2114` — worse — while reciprocity alone gives `+0.7394` and both together `+1.8701`: **trust does not add order, it amplifies whichever sign the system already has.** And a deliberate 40 per cent sunset with no conversion work ends in 2026 with less stored order (5.507) than doing nothing at all (7.494), while the sequenced intervention — shed, then raise `eta`, then raise `u` — ends at 16.770. **Shedding buys headroom; headroom unspent is headroom wasted. Shedding is a stabiliser, not a cure.**

All of this is arithmetic. **Whether any organisation is described by these equations is a separate question, it is a postulate rather than a result, the applied readings are not identified from observational data, and the one preregistered confirmatory test of the underlying law reported in this corpus did not meet its own declared margin.**

14.11.1 Declarations

Data and code availability

No external dataset was collected for this paper. All numerical results are computed from the stated parameters under the model developed in the text, and every table can be re-derived from the equations and parameter values printed beside it. The three systems worked in Section 7 are constructed examples with declared parameters, not field observations. The five figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper and the model equations of Sections 3 to 6. The Monte Carlo figure quoted in Section 5.4 is reported from the source material and was not re-run for this paper, which is stated again in Section 10.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, produced the worked examples and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The three worked systems are constructed. The declination constraint stated in S4.8 and S5.8 — that no mechanism which removes a participant's ability to decline is within the framework's scope — is a design constraint on prospective work and is not a report of any study performed. The monitoring arrangements described in Sections 4.5 and 4.8 are analysed as they appear in the field literature on common-pool resources and are specified here only in the mutual, peer-accountable form that S4.8 permits.

Preregistration

This paper reports no preregistered test of its own; its results are derivations and worked examples. The falsification rules declared in advance for each worked system are stated in Sections 7.1.8, 7.2.8 and 7.3.8 and are **declarations of what would count as failure, not reports of a test performed**. The preregistered protocol for testing the storage law against a real system, together with the confirmatory application that **failed** its second falsification rule — a 2.8 per cent held-out advantage over a matched autoregression against a declared margin of 10 per cent — is reported in Curzi [2026i]. Any applied claim resting on this paper inherits that result.

14.12 Appendix A — The medium survey and governance audit forms

A.1 VWDLA score sheet. Score each property 0–4 by the inspection test of Section 4.4. Record the date and the assessor; the instrument's value is in the trend, because medium capture (FM-4.6) is visible only as a trend.

Property	Inspection test	Observation recorded	Score /4
Visible	Three participants, three levels, one live item each; median time to display; count of times a person had to be asked	median seconds; asks	
Writable	Steps from observation to recorded state, counting blocking review only	step count	
Durable	Bus-count census over critical state classes; fortnight intelligibility test	count of classes at bus-count 1; pass/fail	

Property	Inspection test	Observation recorded	Score /4
Low friction	Five stopwatch samples of one typical write, end to end; count of mandatory fields	median seconds; field count	
Append-oriented	Reconstruct the state as of a named date last quarter	pass/fail	

Pass condition: at least 3 on all five. Any single score of 0 or 1 is a blocking defect regardless of the total.

A.2 Hub census sheet. One row per class of critical state. The count is of people who could reconstruct the item **from what exists outside their own memory**, which is the count that matters and is not the count usually given.

Class of critical state	Where it is held	Reconstructible by (count)	Blocking?

A.3 Ostrom conformance score. The eight-row form of S4.3, scored 0/1/2, total out of 16, with the threshold at 10. Record the evidence, not the judgement: a score of 2 on O5 requires naming the three severity levels and the date the mildest was last used.

A.4 Focal-point inventory. At least five entries, each naming the coordination event it removes and the cost saved per occurrence. Required by S4.5 before any medium is built.

14.13 Appendix B — The independence protocol schedule

Each protocol is listed with the item it is run against, its cost in the normalised units of Section 7.2, and — mandatorily — its decline path stated in the form it is to be read aloud. **A protocol appearing in this schedule without a decline path is not installed under this framework's authority** (S5.8).

Protocol	Run when	Cost per item	Decline path, as stated to participants
IP-1 independent estimate	Before any discussion of the item opens	0.030	"You may record 'no independent view'; the item proceeds and the fact is recorded."
IP-2 anonymous-first	First round of every elicited judgement	0.010	"The anonymous round is optional. If the group is small enough that you can be identified, say so and we will use sealed estimates instead."

Protocol	Run when	Cost per item	Decline path, as stated to participants
IP-3 dissent protection	Standing; exercised at least annually	0.005	“Nobody is required to dissent, and nobody is assigned to. A recorded dissent has no consequence for standing, assignment or advancement.”
IP-4 minority report	Every non-unanimous decision; steelman on unanimous ones	0.040	“The steelman role is by consent and rotates. If nobody consents, the record states that no steelman was appointed.”
IP-5 turn-taking equalisation	Every item	0.005	“You may pass. A pass is recorded as a pass and never as assent.”
Total		**0.090**	

B.1 The measurement schedule that goes with them. IP-2 yields the anonymous-to-attributed drift; IP-3 yields the dissent rate over time; IP-5 yields the share of speaking time held by the top two speakers, with the intervention threshold at 60 per cent in a group of eight or more (S5.6). None of these requires a control group, and all three are recorded on the same sheet, at the same cadence, so that a declining u is visible as a trend across three instruments rather than as a single reading.

B.2 What the schedule cannot do. Where the panel’s signals are duplicates, the protocols extract nothing, because there is nothing to extract (FM-5.6). The annual u re-measurement of S5.7 is what distinguishes a body whose difference is being suppressed from a body that has none, and **only the second is repaired by recruitment.**

Chapter 15

Surveying and Preregistering the Storage Law: A Measurement Protocol, and a Confirmatory Fit That Failed Its Own Baseline Test

Abstract. This paper specifies an end-to-end measurement protocol for the storage law and then submits it to a preregistered test that it fails. The survey fixes five operational choices, supplies estimators for each parameter, and adds a participation-ratio identity, $N_{\text{eff}} = N + 1 - N/\text{PR}$, verified against random ensembles to better than 10^{-12}, that makes a redundancy audit computable from a Gram matrix. Applied to a county development-review office, the survey returned $N = 5$, capacity $U = 0.6972$ against a realised 0.6019 (a realisation deficit of 0.1823 nats caused by 32.7% of sessions running short-handed), $\delta = 0.003090$ per week from a 33-week zero-input episode, $\eta = 0.5015$ rules per nat, headroom $h = -0.1720$, and a negative steady state $Q_\infty = -17.3$, reported as negative. The preregistered confirmatory fit at an eight-week block scale returned $a = 0.97439$ and $b = 0.5859$, both interior to the admissible box, with implied decay within 4.8% and implied conversion within 16.8% of independently audited values. **It nevertheless beat a matched autoregression by only 2.8% on held-out error against a declared requirement of 10%, so the verdict is parameters corroborated, predictive advantage not demonstrated.** A repackaged order proxy passed all four baselines decisively and was caught only by structural tests.

Keywords: survey protocol; preregistration; held-out skill; shuffled-input control; order proxy; participation ratio; zero-input episode; falsification rules

Series: B. Engineering the Surplus **Paper:** CCP-015

15.1 1 Introduction

A framework that cannot fail cannot be trusted. This paper specifies the procedure by which the Creature Comforts storage law is made to face a real system on terms declared in advance, applies that procedure to one system, and reports that **the model failed its own second falsification rule.**

The question is narrow and answerable. The storage law `Q_{t+1} = (1−delta) Q_t + eta S_t − C_t` is a proved result about a difference equation [Curzi, 2026a], whose update map is

a contraction with modulus `1−delta` and whose steady state is `Q* = (eta S − C)/delta`. Nothing in that derivation says anything about an organisation. Turning it into a claim about an organisation requires five operational choices, six estimated parameters, and a commitment — made before the data are examined — about what result would count as a refutation. This paper supplies all three, and then reports what happened when the commitment was honoured.

The contribution has four parts. **First**, an end-to-end measurement protocol: five declared choices, an estimator for each of `delta`, `eta`, `N`, `u`, `C` and `Q`, and a one-page output format in which every number carries its provenance. **Second**, an identity that makes the redundancy audit computable from a Gram matrix, `N_eff = N + 1 − N/PR`, where `PR` is the Rényi-2 participation ratio of the encoded ensemble; the identity is exact and was verified against random ensembles at `K = 3, 5, 8, 12` to better than 1e-12. **Third**, a structural account of the single most common measurement failure in this area — an order proxy that is a repackaging of the interaction input — together with four tests, two of which must be run before any measurement is taken, and a demonstration that the statistical controls alone cannot catch it. **Fourth**, a preregistered confirmatory fit, reported in full including its failure.

The headline result is negative and is stated here in the same words used in Section 9. Applied to a county development-review office, the survey returned `N = 5`, capacity `U = 0.6972` against a realised `U = 0.6019` — a realisation deficit of 0.1823 nats traced to 32.7 per cent of sessions running short-handed — `delta = 0.003090` per week from a 33-week zero-input episode, `eta = 0.5015` rules per nat, headroom `h = −0.1720`, and a steady state `Q_inf = −17.3`, reported as negative. The preregistered confirmatory fit at an eight-week block scale returned `a = 0.97439` and `b = 0.5859`, both strictly interior to the admissible box, with implied decay within 4.8 per cent and implied conversion within 16.8 per cent of independently audited values. **It nevertheless beat a matched autoregression by only 2.8 per cent on held-out error against a declared requirement of 10 per cent.** The verdict is therefore *parameters corroborated, predictive advantage not demonstrated*, and the storage law is **not** corroborated as a predictive model of that system at that time scale.

Two secondary results are reported because they bear on how much the passing tests were worth. A deliberately repackaged order proxy — a rolling index of the board's own disagreement score, which stores nothing — **passed all four baselines decisively**, beating persistence by 89.4 per cent and the matched autoregression by 89.2 per cent at `p = 0.0005` against the shuffled-input control, and was caught only by structural tests: a fitted decay implying a 1.97-week half-life against a measured 224 weeks, and a single-week jump 38.5 times the maximum admissible increment. And the shuffled-input control, on which the empirical audit most depends, has been measured elsewhere in this corpus at a false-rejection rate of 0.108 against a nominal 0.05 [Curzi, 2026u] — it is **anti-conservative**, so the `p = 0.0005` reported here is optimistic by an unknown but non-trivial factor.

Preregistration in the sense used here is the standard remedy for undeclared analytic flexibility [Nosek et al., 2018, Simmons et al., 2011, Gelman and Loken, 2013], and the reason it is necessary in this setting is not suspicion of the analyst but arithmetic: with enough freedom in the analysis any series can be made to support any conclusion, and the person with the most freedom is the person who most wants the conclusion. It is also, and separately, an administrative countermeasure to measure corruption [Goodhart, 1975, Campbell, 1979], because it fixes the definition of the measure at a time when nobody yet has a reason to distort it.

Sections 2 and 3 state the hypotheses — including the falsification rules and their declared margins, *before* any result — and the estimation framework. Sections 4 and 5 give the survey and the order-proxy tests. Section 6 gives the preregistration procedure and the four baselines. Sections 7 and 8 collect the specifications and the named failure modes. Section 9 reports the worked office end to end, including the failed rule. Sections 10 to 13 discuss, verify, limit and conclude. The two operational instruments — the survey worksheet and the preregistration form — are reproduced as Appendices A and B.

15.2 2 Hypotheses and preregistered falsification rules

This paper is partly a protocol specification and partly a confirmatory test. The protocol sections advance no hypothesis and are labelled as specification where they appear. The hypotheses below are what the worked application put at risk, and each carries what would defeat it. **The falsification rules and their numerical margins are stated here, in advance of any result, because that is the order in which they were fixed**; Section 9 reports the measured values against them without adjustment.

15.2.1 2.1 Hypotheses

H1 — The storage law describes a real institutional stock. The count of live precedent rules in a municipal review office evolves as a leaky accumulator driven by the surplus of its review board. **Predicted:** an unconstrained least-squares fit of `Y_{k+1} = a Y_k + b F(U_k) − C_hat_k + eps_k` lands strictly inside `0 <= a <= 1, b >= 0`; the frozen model beats persistence, direct input, a matched autoregression and a shuffled-input control on held-out data by the declared margins; and the fitted `a` and `b` agree with independently measured `delta` and `eta` within the declared tolerance. **Falsified by:** failure of any one of those three rules. *(Result: **partially falsified** — see Section 9.5.)*

H2 — Timing carries information, not merely volume. The order in which interaction arrives predicts subsequent stored order beyond what the input's average level predicts. **Predicted:** the real fit's held-out RMSE lies below the null distribution obtained by permuting the input series in time, at `p <= 0.05`. **Falsified by:** a real fit that sits inside the shuffle null. *(Result: **not falsified** — `p = 0.0005`; see Section 9.4 and the calibration caveat in Section 12.)*

H3 — The mean pairwise aggregator is the participation-ratio effective count. For `K` unit vectors with Gram matrix `G`, the mean pairwise non-parallelism `U` satisfies `N_eff = 1 + (N−1)U = N + 1 − N/PR` with `PR = K^2/||G||_F^2`. **Predicted:** exact agreement in double precision between the two computations for any ensemble. **Falsified by:** any ensemble on which the two disagree beyond floating-point tolerance. *(Result: **not falsified** — agreement better than 1e-12 at `K = 3, 5, 8, 12`; this is a derivation, and the numerical check is a check on the implementation, not on the world.)*

H4 — Structural tests catch what statistical tests cannot. An order proxy that is a filtered copy of the input will pass every held-out baseline including the shuffle, and will be caught only by a parameter-agreement test or an arithmetic ceiling test. **Predicted:** a constructed filtered proxy passes all four baselines and fails independent-parameter agreement and the ceiling check. **Falsified by:** a filtered proxy that the four baselines detect on their own. *(Result: **not falsified** — see Section 9.7.)*

Two claims here are not hypotheses and are labelled as what they are. That stored order in a real institution is adequately summarised by a single non-negative scalar `Y_t` is a **modelling postulate**; nothing in the framework proves it, and Section 6 is the procedure by which reality is allowed to reject it. That the mean pairwise non-parallelism is the right aggregator over the encoded categories is a **declared modelling choice**, badged as such, which must appear in the preregistration; H3 is a derivation about that choice and not a defence of it.

15.2.2 2.2 The falsification rules, with their margins declared in advance

Three rules, each catching a different disease. The tolerances are those that must be fixed before the holdout is examined and may not be moved afterwards.

FR1 — THE ADMISSIBLE BOX. The unconstrained least-squares optimum must lie inside `0 <= a <= 1, b >= 0`. If it lies outside, the storage law is rejected for that system. No constrained refitting rescues it; a constrained fit sitting on the boundary is reporting that the data wanted to leave the box.

FR2 — BASELINE DOMINANCE. The frozen model must beat **all four** required baselines on held-out data by a declared margin: persistence, direct input, matched autoregressive, and the shuffled-input control. Failure against any one of the four is a failure of the rule. **The margins declared for the application in Section 9 were a 10 per cent reduction in held-out RMSE against each of the first three baselines, and `p <= 0.05` against the shuffle distribution.**

FR3 — INDEPENDENT-PARAMETER AGREEMENT. The fitted `a` and `b` must agree with the independently measured `delta` and `eta` from the survey, within a tolerance declared in advance; **the tolerance declared for Section 9 was ±25 per cent on each.** Additionally, no observed one-period increase in `Y` may exceed the ceiling `eta ln N` by more than the declared conversion-lag allowance.

FR1 catches a system that is not an accumulator at all. FR2 catches a model that adds nothing over inertia, trend, or the input's mere average level. FR3 catches a model that fits beautifully with parameters that are not the system's — the signature of the order-proxy trap of Section 5, and the only rule that reliably catches its most dangerous form.

15.3 3 Framework: the law, the five choices, and the estimation model

15.3.1 3.1 The equation being surveyed

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t        S_t = F(U_t) = ln(1 + (N−1) U_t)

Q_t      stored order            delta   retention loss per period, in [0,1]
S_t      surplus input           eta     conversion efficiency, >= 0
C_t      cost / dissipation      N       number of categories
U_t      interaction geometry in [0,1],  S_t in [0, ln N]
```

This is the storage law, and it is proved mathematics under its own assumptions [Curzi, 2026a]: the exact solution, the one-step growth condition, the uniform bound, the sustainability ceiling and the steady state are all derivations. **They are derivations *about the equation*.** They become statements about an organisation only if the equation is a faithful model of it, and whether it is a faithful model of it is exactly what the survey and the preregistration exist to decide. The contraction that supplies the fixed point is Banach's theorem, stated exactly — a contraction with modulus `q < 1` on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least `q` each pass [Curzi, 2026]. The surplus functional `F(u) = ln(1 + (N−1)u)`, uniquely characterised by angle-only dependence, zero at parallel, an affine effective count and normalisation, is established in Curzi [2026].

The master gauge:

```
h_t := eta ln N − delta Q_t − C_t

h_t <  0    unsustainable in this model. No interaction geometry saves it.
h_t =  0    maintenance only. Hold; do not build.
h_t >  0    growth available; any requested increment must satisfy r_t <= h_t.
```

The gauge is the headroom result: `h_t = eta ln N − delta Q_t − C_t`, with three regimes, a maximum one-step increment `r_max = h_t`, and the separate constraints of admissibility `C_t <=`

`(1−delta)Q_t + eta S_t` and sustainability `C_t <= eta ln N − delta Q_t` [Curzi, 2026b]. At equilibrium the same quantity has an interpretation that the survey will need: headroom at equilibrium is unused variety, `h_inf = eta(ln N − F(u*))`, verified to six decimals — at `eta = 0.35`, `delta = 0.015`, `N = 6` and `u* = 0.55`, `Q_inf = 2.840970` and `h_inf = 0.164501` [Curzi, 2026h].

15.3.2 3.2 The five required operational choices

The storage law may not be applied to a real system by analogy. Before any number is taken, five choices must be written down in advance, in language specific enough that a second analyst could take the same measurements and obtain the same numbers.

#	Choice	What it fixes	Symbol produced
1	**State encoder**	how each participant, discipline, or category becomes a vector	`x_i`
2	**Interaction observable**	how the encoded states become a single geometry number per period	`U_t`
3	**Order observable**	what durable thing is counted as stored order	`Y_t` (proxy for `Q_t`)
4	**Time scale**	how long one period is	`t`
5	**Cost observable or cost model**	how dissipation is counted, in the same units as `Y`	`C_hat_t`

These five choices are the entire empirical content of the framework. The mathematics is fixed. Everything that can be wrong about an application of it is in these five lines. An analyst who writes them down badly has not produced a bad estimate; they have produced no estimate at all. The demand is the same one measurement theory makes of any construct before it is quantified — that the operations be specified well enough for another investigator to repeat them [Krantz et al., 1971, Cronbach and Meehl, 1955].

15.3.3 3.3 The estimation model

```
Y_{k+1} = a Y_k + b F(U_k) − C_hat_k + eps_k

subject to      0 <= a <= 1        b >= 0

identification  a = 1 − delta      b = eta
```

This is the storage law written as something a computer can fit. `Y_k` is the order observable; `F(U_k) = ln(1 + (N−1) U_k)` is the surplus input computed from the interaction observable; `C_hat_k` is the cost observable, denominated in `Y`-units, entering with a **fixed coefficient of one** — a measured offset, not a free parameter; `eps_k` is the residual.

Two free parameters. That is all the specification permits. The temptation to add a third — an intercept, a lag, a seasonal term, a second input — is the temptation to make the model unfalsifiable, and every such addition must be declared in advance or not made at all. The parameter count is not a stylistic preference: it is what makes the comparison against a matched two-parameter autoregression in Section 6.4 a fair one, and Section 11.5 records that the corpus did **not** compute an information criterion, which is the standard instrument for paying for parameters and which would have made the comparison sharper still.

The constraint box is the empirical content of the claim rather than regularisation. `a = 1 − delta` with `delta` in `[0,1]` says the system is a **leaky accumulator**: it keeps some fraction

of what it had and loses the rest. If the unconstrained best fit wants `a > 1`, the system is not accumulating with loss — it is compounding, and the storage law is the wrong model. If it wants `a < 0`, the system is oscillating period to period, which a stock does not do. If it wants `b < 0`, then interaction *destroys* order in that system, which is a coherent thing for a system to do and an incoherent thing for this model to describe. **The unconstrained fit is therefore computed first, and where its optimum lands is FR1.**

15.3.4 3.4 Badged postulates

POSTULATE (badged): scalar adequacy. The survey assumes that the stored order of a real system — a body of practice, a trained staff, a furnished room, a working institution — is adequately summarised by a single non-negative scalar `Y_t`, and that this scalar moves according to the storage law up to noise. Nothing in the framework proves this. The static half is pure mathematics and makes no empirical claim; the dynamic half proves properties of a recursion and does not prove that any organisation obeys it.

POSTULATE (badged): linearity. The storage law is linear in `Q` and in `S`. Real institutional stocks may saturate, may have thresholds, may exhibit hysteresis. The linear model is chosen for tractability and for the fixed-point guarantee, not derived from observation of organisations. A held-out failure may mean the storage law is wrong for a system, or it may mean the storage law is right and the linearity is wrong; the procedure below does not distinguish them, and it must not be presented as if it did.

POSTULATE (badged): constant parameters. The model treats `delta`, `eta` and `N` as constant over the estimation window. In an organisation with turnover, reorganisation, or a change of leadership they are not. The window must be declared and checked for a structural break; a fit spanning a regime change is fitting an average of two systems, neither of which exists.

POSTULATE (badged): well-behaved noise. `eps_k` is assumed additive, independent and homoscedastic. Audit error on a register is usually none of these — audits taken by different people, in different weeks, under different pressure. Residual autocorrelation must be reported; where it is large, held-out intervals are narrower than the truth.

The corpus's own classification of its claims puts three items in the postulate drawer against four items of established mathematics and seven of established psychology, with every protocol stated to run on the first two and steer by the third [Curzi, 2026]. The four postulates above are where that drawer bears on measurement.

15.3.5 3.5 Identification: the objection that no fit can answer

This belongs in the method rather than in a caveat, because it limits what the method can deliver before any data are collected.

The specification being fitted is reflexive. `U_k` is a property of the group; the group's stored order shapes what its members do in the next period; and the members are what the group is made of. That is a linear-in-means reflexive structure, and Manski [1993] proves that in such a structure **the endogenous social effect is not identified from observational data**: it cannot be separated from contextual effects — properties of the group that act on members without being produced by them — or from correlated unobservables acting on everyone at once. Shalizi and Thomas [2011] establish the same confounding for network data, where influence and selection cannot generically be told apart from observation.

The consequence for this paper is exact and it is not repaired by anything in Section 6. The falsification rules test whether the fitted model predicts out of sample better than declared baselines. They do not, and cannot, establish that `b` measures the conversion of interaction into structure rather than a contextual effect, a common driver, or selection into the board. A fit that passed all three rules would still not have identified that parameter. **The identification problem is therefore a deeper limitation than the one the confirmatory test found**, and it would have applied equally had the test passed.

No remedy is adopted here, so the honest label is a modelling position rather than an identified empirical result. The four known routes are randomisation, an instrument, partial-population variation, and stipulating the reference group and defending the stipulation [Brock and Durlauf, 2001]. The last is the most realistic in this setting, and the survey's fixed category set, declared merge and drop rules, and explicit participation rule are the machinery a defended stipulation would need — but stipulating the reference group and *defending* it are two different acts, and only the first has been performed. **Until a remedy is executed, the fitted `b` is a coefficient in a declared specification and not an identified social effect**, and every applied statement in this paper inherits that status. Section 9.8's licence list should be read accordingly.

15.3.6 3.6 Standing honesty clauses

- The triangle inequality **fails** for geometric triples; subadditivity holds on the parameter `u` only. Surplus may not be composed across three groups by adding distances.
- As deployed, the surplus functional is a **lower bound** on the ensemble's information, not the information itself. A survey reporting `F(u)` is reporting a guaranteed minimum, and the true richness of the interaction may be higher. This matters in Section 9.6: a fitted conversion slightly above the audited one is not automatically an error.
- Surplus is **not** identified with Shannon entropy or mutual information in general, so entropy intuitions must not be imported into the cost accounting.
- **No thermodynamic negentropy law is proved.** Higher static surplus does not imply higher canonical concentration.
- `0 <= F(u) <= ln N`. **`ln N` is the hard ceiling on surplus per interaction**, which is why `h` is defined with `ln N` rather than with `F(U_t)`.
- **The model is only a claim about the world if `Y` is not a repackaging of `U`.** If the order observable is computed from the same records, by the same instrument, at the same moment as the interaction observable, then `Y_{k+1} = a Y_k + b F(U_k) - C_k` can fit beautifully and mean nothing: it will have measured an algebraic relationship between two encodings of one measurement. Section 5 gives the four tests that prevent this, and Section 9.7 shows a repackaged proxy passing every statistical baseline.

15.4 4 The survey: instruments and estimators

15.4.1 4.1 What a survey is, and what it is not

A survey in this sense is not an audit, an assessment, a maturity model or a culture questionnaire. It produces six numbers and one verdict, it has a fixed output format, and it is finished when the format is full and every entry has a stated provenance.

It is also not a diagnosis. The survey reports the state of the accumulator; it does not report whether the accumulator should be full. **An institution with a low `Q` may be young or may be hollowed out, and a snapshot cannot distinguish them — only the trajectory can.** That is why the survey requires a series rather than a snapshot. The minimum deliverable:

```
SURVEY OUTPUT (one page)
-----------------------------------------------------------------
System ............. name, boundary, and what is excluded
Period ............. one period = ______ ; series length ______
Encoder ............ how x_i was formed
N .................. ______   ln N = ______
U (capacity) ....... ______   from ______ instrument
U (realised) ....... ______   from ______ instrument
F(U) realised ...... ______   = ln(1 + (N−1) U)
```

```
delta ............... ______ /period   half-life ______ periods
eta ................. ______ Y-units per nat
C_hat ............... ______ Y-units/period  (denominated in Y)
Q (= Y at survey) ... ______ Y-units
-----------------------------------------------------------------
eta ln N  = ______
delta Q   = ______
h         = ______        VERDICT: [ h<0 | h=0 | h>0 ]
one-step  = eta F(U) - delta Q - C_hat = ______
Q_inf     = (eta F(U) - C_hat)/delta   = ______
-----------------------------------------------------------------
```

Everything in this section exists to fill that page defensibly.

15.4.2 4.2 Choice 1 — the state encoder

The state encoder turns each *category* into a unit vector. It is the choice that determines what *different* means for a system, and therefore the choice with the largest downstream effect on `u`. Three encoders are practical in the field.

Encoder A — the disposition vector (recommended default). Assemble a fixed probe set of `m` real cases the system must handle. Each category — each discipline, role or site — records its disposition on each case as `+1` or `−1`. The category's vector is that list of `m` signs divided by `sqrt(m)`, so it is a unit vector. For two categories `i` and `j` that agree on `m − k` cases and oppose on `k`:

```
x_i . x_j = (agreements − disagreements) / m = 1 − 2 d_ij      where d_ij = k/m

u_ij = 1 − (x_i . x_j)^2
```

This encoder is exact rather than approximate: for a plus/minus encoding the inner product *is* one minus twice the disagreement rate.

Encoder B — the error vector. For systems producing forecasts or judgements with later ground truth, encode each category by its vector of signed errors over the last `m` resolved cases, mean-centred and normalised. This is the strongest encoder available, because decorrelated error is the operational meaning of independent judgement — the condition that Condorcet's theorem requires, that the diversity-prediction literature identifies as the source of the ensemble gain [Hong and Page, 2004, Page, 2007], and that the collective-intelligence factor tracks [Woolley et al., 2010]. The corresponding sceptical literature must be read alongside it: Thompson [2014] shows that the celebrated diversity-trumps-ability result is a mathematical artefact of its parameterisation rather than an empirical finding, and Harrison and Klein [2007] shows that three incompatible constructs travel under the word *diversity*. The encoder must therefore declare which construct it is measuring; Encoder B measures **disparity of error**, and nothing else.

Encoder C — the practice vector. For systems with no shared probe set, encode by observable practice: which tools, which sources, which sequence of steps, which vocabulary. Score each category on `m` binary practice items. This is the weakest of the three, because two people can share every practice and still disagree, and can share no practice and still agree.

An encoder warning that catches most beginners. `u = 1 − cos^2` is the squared *sine* of the angle between the encoded categories, and the sine rises steeply near alignment. Under Encoder A, a disagreement rate of only 1 case in 12 already gives

```
d = 0.0833 -> cos = 0.8333 -> u = 1 − 0.6944 = 0.3056
```

and a disagreement rate of 3 in 12 gives `u = 0.75`. **The encoder is generous: modest observed disagreement produces large `u`.** This is not an error in the mathematics but a property

of the geometry, and it has two consequences. The encoder must be declared before measurement, because a different encoder gives a different `u` for the same organisation. And a low measured `u` under Encoder A is a very serious finding: it means the categories are near duplicates of each other rather than merely similar.

15.4.3 4.3 Choice 2 — the interaction observable

From `N` unit vectors the framework requires one scalar `U_t` in `[0,1]` per period. It does not say which aggregator to use, so **the aggregator is a declared modelling choice, badged as such, and it must appear in the preregistration.** The recommended aggregator is the mean pairwise non-parallelism over all ordered pairs actually present in the period:

```
U_t = mean over i != j of ( 1 − (x_i . x_j)^2 )
```

There is a reason to prefer it beyond convenience. Let `G` be the Gram matrix of the `K` unit vectors, so `G_ij = x_i . x_j` and `tr G = K`. The participation ratio of the Gram spectrum — the standard Rényi-2 effective count [Rényi, 1961] — is

```
PR = (tr G)^2 / ||G||_F^2 = K^2 / ||G||_F^2
```

Since `||G||_F^2 = K + sum_{i != j} G_ij^2 = K + K(K−1)(1 − U)`, a line of algebra gives

```
K / PR = 1 + (K−1)(1 − U)        hence        U = 1 − (K/PR − 1)/(K−1)
```

and therefore, for the effective count of the surplus law,

```
N_eff = 1 + (N−1) U = N + 1 − N/PR
```

The mean pairwise aggregator is not an arbitrary average. It is exactly the participation-ratio diversity of the encoded ensemble, reparametrised into the effective-count form the surplus law uses. The identity is exact in double precision and was verified numerically for `K = 3, 5, 8, 12` on random ensembles with agreement to better than 1e-12; it is re-verified independently over six random correlation matrices elsewhere in this corpus, where `N = 8` and `u_bar = 0.783801` return 6.486606017883 by both routes [Curzi, 2026]. Declaring the mean pairwise aggregator therefore commits the analyst to a standard and well-understood diversity measure rather than to an improvisation.

Two limiting cases belong on the survey form. At `PR = N` — a Gram matrix with no off-diagonal mass — `N_eff = N`: roster and effective count coincide, and that happens only at exact mutual orthogonality. At `PR = 1` — every vector the same — `N_eff = 1`, and the size of the roster has left the arithmetic entirely. **Everything the survey reports about `N` lives between those two poles, and the participation ratio is the coordinate that locates it.** Where a sponsor asks why disagreement is averaged pairwise rather than some other way, the identity is the answer, not a plausibility argument.

One caution about what the identity does and does not import. `1/HHI` and `exp(H)` are Hill numbers of order two and one respectively [Hill, 1973, Jost, 2006], and the concentration instrument built on `1/HHI` is a separate construction with its own worked results [Curzi, 2026]. `N_eff = 1 + (N−1)u` is a linear interpolation between 1 and `N` and is **not** a Hill number of any order; it is this framework's own construction, and Section 12 states that plainly. What the participation-ratio identity establishes is that the chosen aggregator coincides with a recognised effective count *for this encoding*, which is a narrower and defensible claim.

Two distinctions must be kept separate, because collapsing them is the second most common survey error after the proxy trap. **Capacity `u` versus realised `u`:** the probe set measures what the category set is *capable* of — how non-parallel the disciplines are when all of them

speak — while the period-by-period observable measures what actually happened, which categories were present, which spoke, which were overruled before speaking. These are different numbers, and the gap between them is among the cheapest things an organisation can fix. The worked office of Section 9 has capacity `U = 0.6972` and realised `U = 0.6019`, and the entire difference comes from sessions run with fewer than the full board present. **Present versus counted:** a category in the room but silent contributes to nothing. Under Encoder A the survey must define in advance what counts as participation — a recorded disposition, a spoken objection, a signed line. Attendance is not participation and must never be encoded as it.

15.4.4 4.4 Choice 3 — the order observable

This is the choice that decides whether a survey is worth anything. The order observable `Y_t` is a countable, durable quantity standing in for `Q_t`. A valid `Y` has five properties, and each is tested explicitly with the answer recorded.

Property	Test	Why it matters
Durable	Does it persist for many periods with no input?	A stock, not a flow
Countable	Can two auditors independently produce numbers within a declared tolerance?	Otherwise `delta` and `eta` are noise
Decaying	Does it measurably decline under zero input?	If it never decays, `delta = 0` and the model has no fixed point
Independent of `U`	Different instrument, different record, different time?	the order-proxy condition — see Section 5
Denominable	Can cost be expressed in the same units?	Otherwise `h` cannot be computed at all

Candidate order observables that satisfy them, by system type:

System	Candidate `Y`	Units
Review office / standards body	live precedent rules in the standing manual	rules
Research group	reproducible protocols in the group handbook	protocols
Software organisation	modules with a named owner and a passing test suite	modules
Library reading room	seat-hours of optional occupancy per week	seat-hours
Maintenance regime	assets with a current, executed maintenance record	assets
Co-housing block	standing agreements in force and observed	agreements
Manufacturing cell	validated work instructions in use on the floor	instructions

And observables that do not, with the reason. **Headcount** does not decay in the right way, is not produced by interaction, and is an input to `N` — it is on the wrong side of the equation. **Revenue** is a flow rather than a stock and is driven by exogenous demand. **Engagement scores** are not durable; they are measured, not held. **Meeting count** is the support of `U`, not `Q`; counting meetings as order is the purest form of the trap in Section 5. **Any index the organisation reports upward** is exposed to measure corruption from the moment it is known to be counted [Goodhart, 1975, Campbell, 1979].

15.4.5 4.5 Choice 4 — the time scale

One period is the unit of the recursion. Choosing it wrongly does not bias the estimates; it destroys the ability to detect anything. Three constraints bound the choice.

The period must be at least as long as the conversion lag. If a decision made in a meeting takes six weeks to reach the standing manual, a one-week period puts cause and effect in different rows and the fit will find nothing.

The period must be short enough that `delta` per period is small — as a practical rule, `delta <= 0.3` per period, so that the linear approximation of a decaying stock is not being asked to cover a near-total loss in one step.

The period must be long enough that the input term is visible above measurement noise. This is the constraint nobody considers, and it has a precise structure. Audit error on a *level* does not accumulate: re-counting the stock twice as far apart leaves the error at the endpoints the same size while the accumulated input doubles. **Coarsening the time scale raises the signal-to-noise ratio of the input term roughly as the square root of the block length, and costs degrees of freedom linearly.** There is an optimum, and it must be chosen before the answer is seen (Section 6.7).

The period, the calendar alignment — which day the audit is taken — and the rule for handling periods with no interaction at all must all be declared. Periods with `U_t = 0` are not missing data; they are the most informative periods available, because they are where `delta` is identified.

15.4.6 4.6 Choice 5 — the cost observable, and the denomination rule

`C` must be denominated in the same units as `Y`. Where it cannot be denominated, headroom cannot be computed and the survey has failed. This is the most frequently botched line in a real survey, because cost arrives naturally in money and hours while `Y` arrives naturally in rules, protocols or modules. The requirement is dimensional rather than stylistic: every term in `h = eta ln N − delta Q − C` must carry the same dimension, and Section 11.1 gives the audit.

There are exactly two honest routes. **Route 1 — direct destruction accounting.** Count the `Y`-units destroyed or made unusable per period: rules struck by appeal reversal, rules orphaned because the only person who could apply them left, rules rendered unusable by contradiction with a newer rule. These are already in `Y`-units and require no conversion. **Route 2 — a declared conversion rate.** For dissipation that is real but not directly destructive — coordination overhead, rework, interruption, emotional labour, compliance reporting — measure it in its natural unit and convert at a rate declared in advance and justified. The honest way to set the rate is by *displacement*: measure how many `Y`-units the system produces per staff-hour of productive capacity, then charge dissipated hours at that rate. **A declared conversion rate is a badged modelling choice; an undeclared one is a fabrication.**

Cost categories to enumerate, with their usual `Y`-denominated form:

Category	Natural unit	Denomination route
Coordination overhead (Brooks, `n(n–1)/2` paths)	hours	Route 2
Rework and re-decision	items	Route 1 if it strikes `Y`, else Route 2
Conflict and appeal	cases	Route 1
Turnover loss	departures	Route 1 (orphaned units)
Maintenance and compliance	hours	Route 2
Rent, commute, physical friction	money, minutes	Route 2
Interruption / context switching	switches	Route 2

Two rules prevent the usual mistakes. **Do not double-count turnover.** Turnover appears in `delta` as decay of stored order *and* in `C` as a one-off destruction; each departure is assigned to one or the other and the assignment is written down. The convention that works: routine attrition at the historical rate sits in `delta`, while departures above that rate, and any departure that orphans a specific `Y`-unit, sit in `C`. **Do not omit the cost the organisation does not book.** The attention tax — the hours spent by people who are not in the meeting, reading about the meeting — is almost never in anyone's ledger and is often the largest single line. Omitting it makes `h` optimistic and the eventual intervention fail for a reason that was never measured. The coordination component has a known form: communication paths grow as `n(n–1)/2`, and the cost of bringing new arrivals up on person-borne stored order is charged against the same budget [Brooks, 1975, Conway, 1968].

15.4.7 4.7 Estimating `delta`

`delta` is the fraction of stored order lost per period with no maintenance and no input.

Method D1 — the zero-input episode (best). Find a stretch of `n` periods in which `U_t = 0` and `C_t = 0`: a shutdown, a hiring freeze, a suspended committee, a closed season, a strike, a vacancy. Over that stretch the law reduces to `Q_{t+n} = (1 − delta)^n Q_t`, so

```
delta = 1 − (Y_end / Y_start)^(1/n)
```

This is the only clean estimator, and finding such an episode in the record is the single highest-value hour of the whole survey. Institutions have more of them than they admit: budget freezes, closures, leave-of-absence gaps, the six months between an outgoing and an incoming director.

Method D2 — the cohort decay curve. Tag a cohort of `Y`-units created in one period and count how many are still live after 1, 2, 4, 8, 16 periods. Fit `ln(survivors)` against time; the slope is `ln(1 − delta)`. This is more work than D1, works without a zero-input episode, and additionally reports whether the decay is exponential at all. If the survival curve is not straight in log space, the constant-`delta` assumption is wrong for that system and the survey must say so.

Method D3 — the maintenance ledger (fallback). Where the organisation records maintenance effort per `Y`-unit, `delta` is approximately the fraction of units that would lapse in one period if that effort stopped. This is an estimate from testimony rather than measurement and must be labelled as such.

`delta` is then converted into the two numbers a non-technical reader can use:

```
half-life  t_half = ln 2 / ( −ln(1 − delta) )       periods
time constant  tau = 1 / delta                       periods
settling to 95% of a new equilibrium: about 3 tau    periods
```

delta	half-life (periods)	tau (periods)	3 tau (periods)
0.003	230.7	333.3	1000
0.005	138.3	200.0	600
0.01	69.0	100.0	300
0.02	34.3	50.0	150
0.05	13.5	20.0	60
0.10	6.6	10.0	30
0.20	3.1	5.0	15
0.30	1.9	3.3	10
0.50	1.0	2.0	6

A low `delta` is not unambiguously good. It buys durability and charges latency: a system with `delta = 0.003` per week takes a thousand weeks to settle to a new equilibrium. That is the reason an intervention protocol promises increments and never promises equilibria, and it is the same trade-off — doubled stock against doubled settling time — established for the storage law itself [Curzi, 2026a].

15.4.8 4.8 Estimating `eta`

`eta` is `Y`-units of durable order created per nat of surplus input. It is the parameter most organisations have never considered and the one most often closest to zero.

Method E1 — the attribution audit (recommended). Choose a window of `w` periods in which the system was operating normally. For each `Y`-unit created in that window, ask a single question with a written answer: which interaction produced this, and is that interaction in the record? Count only the units with an affirmative, documented answer. Then

```
eta_hat = (attributable Y-units created in the window) / (sum of F(U_t) over the window)
```

The denominator is in nats, the numerator in `Y`-units, so `eta` is `Y`-units per nat. The audit is deliberately strict: a unit created by one person alone at a desk is not attributable to interaction and does not count, and a unit whose origin nobody can name does not count. **The strictness is the point. `eta` is meant to measure the fraction of interaction that becomes structure, and a permissive audit measures the fraction of structure that had a meeting nearby.**

Method E2 — the capture-rate decomposition. Where the attribution audit is too expensive, decompose:

```
eta ~= (fraction of interactions that produce a written decision)
     x (fraction of written decisions that reach the durable medium)
     x (Y-units per decision that arrives)
```

Each factor is separately observable and each names its own repair. A meeting culture with excellent minutes and no publication path has a high first factor and a near-zero second.

Method E3 — the natural experiment. Where the organisation changed its capture practice at a known date — introduced a decision log, hired a technical writer, mandated a template — estimate `eta` separately before and after. This is the cheapest way to learn what a conversion intervention is worth in a given system.

15.4.9 4.9 Estimating `N`

`N` is the number of genuinely distinct categories the system spans. It appears in the ceiling `ln N` and therefore in every design inequality. The procedure has four steps. **Enumerate candidate**

categories — disciplines, training lineages, functions, sites, tenures, client types, whatever the system itself treats as a kind. **Merge duplicates**: two candidates merge if their encoded vectors have `u_ij < 0.10`, that is, if they disagree on less than about 2.6 per cent of probe cases under Encoder A; two categories that never disagree are one category with two names, and counting them twice inflates `ln N` while contributing nothing. **Drop non-participants**: a category with no route into the interaction is not in `N`, and a discipline consulted after the decision is not in `N`. **Count what survives**, recording the merge and drop decisions with reasons.

The marginal value of `N` is unforgiving and should be shown to anyone proposing to hire their way out of a headroom problem:

N	ln N	gain over N-1
2	0.6931	0.6931
3	1.0986	0.4055
4	1.3863	0.2877
5	1.6094	0.2231
6	1.7918	0.1823
8	2.0794	0.1335
10	2.3026	0.1054
12	2.4849	0.0870
16	2.7726	0.0645
20	2.9957	0.0513
30	3.4012	0.0339
50	3.9120	0.0202
100	4.6052	0.0101

Against the growth of communication paths, `n(n−1)/2` [Brooks, 1975]:

n	paths	paths added by the n-th member
3	3	2
4	6	3
5	10	4
6	15	5
8	28	7
10	45	9
12	66	11
15	105	14
20	190	19
26	325	25

Surplus grows logarithmically and coordination cost grows quadratically. That is the whole team-sizing problem in one line, and the survey's job is to find where a system sits on it. Two independent literatures reach the same shape from different directions: waiting time diverges as utilisation approaches one, so reserve is not prudence but the only thing keeping response finite [Kingman, 1962, Hopp and Spearman, 2011]; and cohesion buys creative success up to a point and then costs it, an inverted U rather than a monotone gain [Uzzi and Spiro, 2005].

15.4.10 4.10 Estimating `u`

With the encoder fixed and the aggregator declared, `u` is arithmetic. What takes judgement is the instrument that feeds it.

Instrument U1 — the probe set (capacity). Eight to sixteen real cases chosen to span the system's actual workload, each with a forced binary disposition, administered independently and in writing *before* any discussion — because discussion is exactly the mechanism by which informational cascades collapse `u` [Bikhchandani et al., 1992, Banerjee, 1992, Lorenz et al., 2011]. Administer annually: `u` is not a property a system acquires but a quantity that decays.

Instrument U2 — the disagreement audit (realised). From the decision record, count for each period the pairs of categories that registered opposed positions before the decision closed. This is a weaker signal than the probe set — a category can be silent because it agrees or because it has learned not to speak — but it is available for every period retrospectively, which the probe set is not.

Instrument U3 — error decorrelation (strongest, where available). Where the system makes resolvable predictions, compute the correlation of signed errors between categories. Decorrelated error is the operational definition of independent judgement, and it is the only instrument that cannot be gamed by disagreeing performatively.

Instrument U4 — the collision test (fast). Ask each category, independently and in writing, for the three things that would most change their mind, then count the distinct answers. A category set returning the same three answers has `u` near zero, whatever its demographics.

The effective-count table for `N = 5` is the arithmetic that makes `u` legible to a manager:

u	N_eff = 1 + 4u	F(u) = ln N_eff
0.00	1.000	0.0000
0.10	1.400	0.3365
0.20	1.800	0.5878
0.30	2.200	0.7885
0.40	2.600	0.9555
0.50	3.000	1.0986
0.5715	3.286	1.1897
0.60	3.400	1.2238
0.6972	3.789	1.3320
0.80	4.200	1.4351
0.90	4.600	1.5261
1.00	5.000	1.6094

Five reviewers at `u = 0` are one reviewer. The survey's job is to report the effective count, not the roster. This is the same instruction the concentration instrument gives in a different domain, where four jurisdictions with shares 40.9 / 27.3 / 18.2 / 13.6 read as 3.41 effective jurisdictions naively and as exactly 1.00 once the bloc structure is collapsed [Curzi, 2026], and where four US states with shares 40.0 / 25.0 / 20.0 / 15.0 read as 3.508772 naively and 1.00 after collapse, with a surplus term of exactly zero [Curzi, 2026].

15.4.11 4.11 Estimating `C`, `Q`, and the instruments that take the readings

Cost. Follow the denomination rule and build the tally as a line-item table, one line per category, each with its measurement method and its conversion where Route 2 applies. The tally is the only part of the survey the organisation will actively contest, because every line is somebody's work. Four disciplines keep it honest. Every line has a source document — appeal register, exit interviews, ticket system, time records, the minutes. Every Route 2 line shows its conversion rate and the justification for it, in the tally. **The tally is reconciled against the series**: where the line-item `C_hat` differs from the cost implied by the observed trajectory by more than the declared tolerance, one of them is wrong; in the worked office the line-item tally came to 0.63 `Y`-units per week against a series-implied 0.6275, agreeing to 0.0025, which is well inside tolerance and is the only reason the tally may be used. And **slack is not cost**:

queueing theory is unambiguous that utilisation near one makes waiting time explode [Kingman, 1962], so unused capacity is a cost-control device rather than a dissipation, and booking it as `C` inverts the design conclusion.

Stock. `Q` is `Y` at the survey date, once the count has been satisfied to be the count. Two corrections are almost always needed and they differ in kind. The **liveness audit** removes units that are formally in the register but which nothing and nobody has used in the last `k` periods; `k` is defined in advance. **Striking units that were already inert is not shedding `Q`; it is correcting the measurement** — a distinction that matters enormously downstream, because correcting the measurement changes the number and not the world, whereas shedding `Q` is an intervention with real effects. The **applier audit** removes units that survive but which no remaining person can apply; orphaning is a real loss and belongs either in `delta` if routine or in `C` if caused by a specific departure, but the unit itself comes out of `Y` either way. The pre-correction count, both corrections, and the post-correction `Q` are all recorded, so that a second analyst can see what was removed and why.

Instruments. Five take the readings. *Direct observation*: Whyte's method for physical systems — go, sit, watch, count, and record what people actually do rather than what they say they do [Whyte, 1980]. For an organisation the equivalent is attending the interaction without a role in it and tallying who spoke, in what order, for how long, who was interrupted, what was written down, and by whom. Whyte's finding that people sit where there is something to sit on generalises: **systems do what their fixtures make easy** — the same observation Jacobs [1961] and Gehl [1971] make about streets. An hour of direct observation routinely overturns a month of interviews. *The unglamorous record*: Whyte used time-lapse film because occupancy patterns are invisible at human tempo; the organisational equivalents are the commit log, the calendar export, the ticket system, the door-badge record, the room-booking data — the only sources in most organisations that were not written to be read by management. *The decision log*: the single most useful artefact to install if it does not exist and to read if it does, with each entry recording date, question, who was present, positions recorded before discussion, decision, and the durable artefact produced; a decision log yields `U_t` and `eta` at once, and a system that has kept one for two years can be surveyed in a week. *The disagreement audit*: a structured pass over the decision record coding each decision for whether any category recorded an opposed position and whether that position survived to the written outcome — the ratio of the second to the first is one of the best single diagnostics available, since an organisation where dissent is recorded and never survives has a governance problem that `u` alone will not reveal. *Turnover and tenure records*: these feed `delta` directly and `C` through orphaning, and the distribution rather than the average is what matters, since a department with a stable core and a churning periphery has a very different `delta` profile from one with uniform tenure.

A structural warning attaches to all five. Star and Ruhleder [1996] establish that infrastructure becomes visible only on breakdown, which has a specific consequence for surveying: **everything that is working will be systematically under-measured.** The maintenance regime that never fails has no tickets, no incidents, no records and no budget line. For each `Y`-unit class the survey must ask explicitly what keeps it alive, and expect the answer to be a person rather than a process.

15.4.12 4.12 Computing headroom and reading the verdict

With the six numbers in hand:

```
eta ln N        the best case: what conversion could deliver at the ceiling
delta Q         the maintenance draw: what the existing stock costs to hold
C_hat           the dissipation draw
h = eta ln N - delta Q - C_hat
```

Read as a budget: `eta ln N` is the income at maximum geometric efficiency, `delta Q` the

mortgage on what is already owned, `C_hat` the operating expense, and `h` what is left. Then compute the *realised* one-step change, which is a different and equally important number:

```
one-step = eta F(U_realised) - delta Q - C_hat
```

`h` uses the ceiling `ln N`; the one-step uses what actually happened. The gap between them is the **realisation deficit** — headroom that exists and is not being collected — and both must be reported, because they call for different classes of intervention.

One correction on `F(U_realised)` that most surveys get wrong. `U` is not constant across the periods of a series; it is a distribution. What the accumulator integrates is `sum of eta F(U_t)`, so the operative quantity is the **mean of `F`**, not `F` of the mean. Because `F` is strictly concave [Curzi, 2026], Jensen's inequality fixes the direction:

```
mean of F(U_t)  <=  F( mean of U_t )        with equality only if U is constant
```

Compute `F` period by period and average the results; never average `U` and then take `F` once. The gap is the concavity penalty and it is not small: in the worked office it is 0.0761 nats against an `F(mean U)` of 1.2262 — **6.2 per cent of the surplus, lost purely to the variability of the geometry.** A survey reporting `F(mean U)` overstates realised surplus by that much, and every downstream number inherits the error. Both figures belong on the output page, with the difference labelled as the concavity penalty, and the test runs against the mean of `F`.

Finally compute the steady state `Q_inf = (eta F(U) - C)/delta`. **Where `Q_inf` is negative, the result does not hide the failure mode: the regime is unsustainable, or the linear model must be replaced.** A negative `Q_inf` in a survey output is not a rounding artefact and must never be reported as approximately zero. It is the model reporting, in its own terms, that the system as configured is being driven to the boundary of its own validity.

One further prediction is available at no cost and should be made every time. `u` decays. Holding the other quantities constant and letting `u` decay geometrically at rate `lambda` per period, the doctrine-to-data ratio crosses 1 at

```
t* = ln(u*/u_0) / ln(1 - lambda)          periods
```

`t*` has no free parameters. Every quantity in it — the current `u_0`, the crossing level `u*`, the decay rate `lambda` — is already on the survey form, and nothing has to be fitted, tuned or assumed. In the corpus's worked case it comes to about 8.4 quarters. That makes it the most decisive and least expensive falsification available anywhere in this protocol: write the date down, wait, and look. A system that sails past `t*` with the ratio still below 1 has falsified either the decay estimate or the model, and either finding is worth more than another season of fitting. This is exactly the kind of point prediction with a constant that Meehl [1967] identifies as the appropriate response to the charge that social-science predictions are directional and nearly costless.

And a hard rule about combining surveys, because the arithmetic is unforgiving. The ratio must not be aggregated across sites, divisions or cohorts. Distortion does not average. Two sites distorted in opposite directions — one over-framed, one under-framed — produce an aggregate in which the two errors cancel, and the cancellation is not a reconciliation but a deletion. In the corpus's two-site example the aggregate **overstated bankable input by 4.26 times** while each site taken alone was correctly diagnosed. **Survey each site to its own ratio and report the vector, never the mean.** An organisation with one site at 2 and one at 0.5 does not have an organisation at 1; it has two different pathologies and no healthy site at all.

15.5 5 The order-proxy trap, and why the statistical controls cannot catch it

Everything in Section 4 is craft. This section is the part that decides whether the craft was worth anything.

The model is a claim about the world only if the order proxy `Y` is not a repackaging of the input `U`. Where it is, the estimation model of Section 6 will appear to fit — often superbly — and will have measured nothing about the world.

Why it happens is not stupidity. It happens because the most available records are the interaction records: the minutes exist, the attendance list exists, the ticket system exists. So an analyst in a hurry defines `Y` as *items dispositioned per session* or *decisions logged per month* or *a weighted index of board activity* — and every one of those is a function of the interaction rather than a store left behind by it. The equation then reads, in effect, `U = a U + b U`, and the fit recovers an identity.

There are two species and they need different tests. **Species 1 — contemporaneous repackaging:** `Y_t` is computed from the same records as `U_t`, at the same time; for example `Y` = items dispositioned this session against `U` = disagreement observed this session, both functions of how busy the board was. **Species 2 — filtered repackaging:** `Y_t` is a smoothed or accumulated version of `U` itself; for example a rolling index of the board's disagreement score, published weekly by the clerk who writes the minutes. **The second species is far more dangerous, because a filtered input is *genuinely* predicted by the input, so the model fits almost exactly and passes every statistical baseline including the shuffled-input control** — demonstrated with numbers in Section 9.7.

15.5.1 5.1 The four tests

All four are run and their answers recorded. Two are structural and must be run before a single measurement is taken.

Test 1 — the algebraic derivation test (structural). Write out, as a formula, how `Y_t` is computed from source fields, and how `U_t` is computed. Compare the source fields. If any field appears in both, or if `Y`'s formula contains `U`'s formula, there is repackaging. Declared tolerance: **zero shared fields.** The test costs ten minutes and catches most cases.

Test 2 — the instrument-independence test (structural). `Y` must be recorded by a different person, into a different medium, on a different schedule, from `U`. In the worked office, `U` comes from the board's minutes and `Y` from the code administrator's register, and the code administrator does not attend the board. Where the same clerk produces both, the test fails regardless of how the formulas look, because the two numbers will share the clerk's judgement. This is the multitrait–multimethod requirement in miniature: a construct measured by one method cannot be distinguished from the method [Campbell and Fiske, 1959].

Test 3 — the zero-input persistence test (empirical, decisive). Find a period with `U = 0` and look at `Y`. **A stock persists and decays slowly; a repackaged input collapses to zero and springs back.** This is the strongest empirical test available and it uses the same episode the `delta` estimate needed. In the worked office the true proxy went 206 to 186 across a 34-week freeze and continued from 186 when the board resumed, while the repackaged proxy fell from 82.1 to 0.1 and returned to 31.6 in a single week.

Test 4 — the ceiling test (arithmetic, decisive). The law itself bounds how fast a genuine stock can grow: the largest possible one-step increment is `eta ln N − delta Q − C <= eta ln N` [Curzi, 2026b]. In the worked office that ceiling is `0.5015 x 1.6094 = 0.8071` `Y`-units per week, and the repackaged proxy's single-week jump of 31.1 units is **38.5 times the theoretical maximum**. A quantity that can move faster than the conversion capacity of the system that allegedly produced it is not stored order. **The law is not only a model; it is a consistency check on the instrument.**

Two further tells do not prove repackaging but reliably accompany it. **A fitted `a` implying an implausible half-life**: in the worked office the fitted `a` for the filtered repackaging was 0.7038, implying a half-life of 1.97 weeks against a directly measured half-life of 224 weeks — a discrepancy of 114-fold, and an institutional stock with a two-week half-life is not an institutional stock. **A fitted `b` far outside the audited `eta`**: the same fit returned `b = 20.2` against an attribution-audited `eta = 0.5015`, a factor of 40. `b` is supposed to *be* `eta`; when the two are not in the same order of magnitude they are not measuring the same thing.

15.5.2 5.2 Three measured reasons the shuffle will not save the analysis

The four-baseline suite of Section 6.4 ends with the shuffled-input control, and analysts reach for it as the general-purpose guarantee that a fit means something. **It is not one.** Its behaviour has been measured by simulation rather than asserted [Curzi, 2026u], and it fails in three distinct ways, two of which can only be repaired at design time.

(1) It is anti-conservative as ordinarily run. Against a nominal `alpha` of 0.05 across 500 null trials the measured false-rejection rates were **0.108 for iid permutation (2.2 times nominal), 0.120 for circular shift (2.4 times nominal), and 0.060 for block permutation** (near nominal) [Curzi, 2026u]. A surrogate test taken off the shelf roughly doubles the false-positive rate, and the block variant — the one most often skipped as fussy — is the only one that behaved. **The remedy is not to abandon it but to calibrate it by simulation on the study's own design** — its series length, its autocorrelation, its block structure — and to preregister the calibrated threshold rather than the nominal one. In the same exercise the matched autoregressive baseline evaluated *alone* false-passed 35.8 per cent of the time, which is the arithmetic behind the requirement that all four baselines be read together.

(2) It cannot see repackaging. The filtered repackaging of Section 9.7 passed all four baselines, the shuffle included, at `p = 0.0005`. It was caught by Test 4 above and by nothing else. **Repackaging is a structural fault and only a structural test detects it.**

(3) It cannot see confounding by a shared slow driver. Where `U` and `Y` are both driven by a third slow variable — a funding cycle, a hiring wave, a market season — a 4,000-trial simulation found the shuffled-input control **passing 98.35 per cent of the time**, while the same simulation reports the control working honestly in the clean case at 48.17 per cent detection [Curzi, 2026u]. The contrast is what makes the 98.35 per cent figure a measurement rather than a complaint.

The consolidated rule: more surrogates do not fix any of this; design controls do. Blinding, channel separation, prospective sampling and a yoked control are the instruments that answer (2) and (3), and every one of them is a decision made among the five operational choices, before a single number exists. **This is why the repackaging tests sit in the survey section and not in the statistics section: by the time columns are being permuted, the instrument is already built.** It is the same lesson the garden-of-forking-paths literature reaches about analytic flexibility [Gelman and Loken, 2013, Simmons et al., 2011] — the damage is done at design time, and no amount of subsequent computation undoes it.

15.6 6 Preregistration and the four baselines

The survey produces six numbers. It does not produce a licence to use them. Six numbers arranged into an equation are a hypothesis about how a real system behaves over time, and a hypothesis that has not been given a chance to fail is not a hypothesis. This section is the procedure that gives the storage law a fair chance to be wrong about an organisation — and, when it is wrong, the procedure that makes the analyst say so.

15.6.1 6.1 What preregistration is for

Preregistration is the practice of writing down, in advance and in a form that cannot be silently edited, exactly what will be measured, exactly how it will be modelled, exactly what result would count as success, and exactly what result would count as failure [Nosek et al., 2018].

It is neither bureaucracy nor modesty. It solves a specific, mechanical problem: with enough freedom in the analysis, any series can be made to support any conclusion, and the person with the most freedom is the person who most wants the conclusion [Simmons et al., 2011]. In organisational work that person is usually the analyst, hired by someone with a preferred answer, working on a system whose staff will read the report. Three freedoms must be

surrendered, in order of the damage they do. **Freedom over the outcome** — choosing `Y` after seeing which candidate proxy moved. **Freedom over the time scale** — trying weekly, monthly and quarterly and reporting the one that worked; this is especially seductive because, as Section 6.7 shows, the time scale genuinely matters, which is exactly why choosing it after the fact is fatal. **Freedom over the specification** — adding a lag, an intercept, a control, or a dropped outlier until the fit clears. Preregistration surrenders all three in a dated document, before the holdout is examined.

15.6.2 6.2 The freeze

The freeze is the moment every choice becomes read-only. It has a date, a signature, and a copy held by someone other than the analyst.

Frozen item	Source
The five operational choices	Appendix A, Form A
The repackaging audit results	Appendix A, Form B
`N`, and the merge/drop decisions	Appendix A, Form C
The estimation window and the train/holdout split	Section 6.3
The time scale and its block alignment	Section 6.7
The model specification, including that `C_hat` has coefficient 1	Section 3.3
The four baselines and their exact definitions	Section 6.4
The three falsification rules and their numerical tolerances	The Law, above
The pre-declared handling of missing periods, outliers, and structural breaks	Section 6.3
The stopping rules and harm criteria for any intervention that follows	Section 6.10

What is *not* frozen, because it must not be: the holdout data. The training data may be examined as much as required before the freeze — that is what training data is for. **The holdout is not examined until after the freeze, and it is examined once.**

Where something must change after the freeze, it may — and the change converts the entire analysis from confirmatory to exploratory, permanently, and must be labelled as such in every sentence that reports it. This is not a punishment. Exploratory analysis is how new hypotheses are found. It is simply not evidence [Platt, 1964].

15.6.3 6.3 Train, holdout, and the fitting procedure

Split the series by time, never at random. Random splits leak: a randomly held-out period sits between two training periods and its level is nearly determined by them. The recommended split is the last 20 to 30 per cent of the series by calendar, with the exact boundary period declared in the preregistration. Three rules are violated constantly. **One look**: the held-out numbers are computed once, and any refitting after looking must be declared and the analysis relabelled exploratory. **No holdout in the estimation of anything**, including `delta` from the zero-input episode and `eta` from the attribution audit; where the zero-input episode falls in the holdout, either the split moves or the independent measurement of FR3 is no longer independent, and the report says which. **Report the holdout size**: a holdout of four blocks is not a holdout but an anecdote with a standard deviation, and Section 9.6 shows what that costs.

The awkward cases are declared in advance. **Periods with `U_t = 0` are kept** — they are the most informative periods available — with `F(0) = ln(1) = 0` stated explicitly, which is exact and requires no special-casing. **Missing audits** follow a declared rule, either carry-forward with a flag or drop-the-pair with a count; a stock is never interpolated, because interpolation manufactures the smoothness the analysis is trying to test for. **Outliers** follow a rule that does not reference the residual, because a rule that drops large residuals guarantees a good fit; the admissible rule references the *record* — a period in which the register was migrated, recounted or restated may be excluded, documented at the freeze if known and reported as exploratory if discovered later. **Structural breaks** that are known — a merger, a new director, a legal change — are declared, along with whether the window spans them and whether the fit will be run separately on each side; discovering a break after the fact and splitting on it is exploratory. And **the ceiling check** is declared: any observed single-period increase exceeding `eta ln N` plus the conversion-lag allowance is investigated before the analysis proceeds.

The fitting procedure is numbered so that it can be executed without judgement.

1. **Assemble the pair set.** For each period `k` in the training window for which both `k` and `k+1` are observed, form the row `(Y_k, F(U_k), C_hat_k, Y_{k+1})`.
2. **Move the cost to the left.** Define `z_k = Y_{k+1} + C_hat_k`. Because `C_hat` has a fixed coefficient of one, it is an offset and not a regressor.
3. **Fit unconstrained.** Solve the ordinary least-squares problem `z_k ~ a Y_k + b F(U_k)` with no intercept, and record `(a_unc, b_unc)`. **This is FR1's test; it is performed before the constrained fit and the answer is written down.**
4. **Fit constrained.** Solve the same problem subject to `0 <= a <= 1`, `b >= 0`, and record `(a, b)`.
5. **Translate.** `delta_fit = 1 − a`; `eta_fit = b`. Convert `delta_fit` to a half-life and a time constant so it is comparable with the survey's directly measured `delta`.
6. **Report the profile interval on `a`.** Refit `b` on a grid of fixed `a` values, record the sum of squared errors, and take the set of `a` within the declared confidence threshold. On slow stocks this interval is wide, and showing that it is wide is part of the report.
7. **Freeze the fitted coefficients.** They are now constants.
8. **Predict the holdout, once.** `Y_hat_{k+1} = a Y_k + b F(U_k) − C_hat_k`, one step ahead, using the *observed* `Y_k` at each step rather than the model's own previous prediction. One-step-ahead is the honest evaluation for a stock; free-running multi-step prediction of a near-unit-root series tests patience rather than the model.
9. **Compute held-out RMSE** for the model and for each of the four baselines.
10. **Run the shuffled-input control** (Section 6.5).
11. **Evaluate FR1, FR2 and FR3 against the frozen tolerances.** Record each as PASS or FAIL. They are not averaged and not weighted.
12. **Write the verdict sentence.** One sentence of the form: *the storage law is / is not corroborated as a model of `Y` in this system at the preregistered time scale, and the rule it failed is ___.*

The metric throughout is one-step-ahead held-out RMSE in `Y`-units. Raw RMSE is reported for every model, not only skill scores, because a skill score of 40 per cent means something completely different when the RMSE is 0.5 rules and when it is 50 rules. The skill convention is `skill = 1 − RMSE_model / RMSE_baseline`, so positive is better and a skill of 0.40 means the model's error is 60 per cent of the baseline's.

15.6.4 6.4 The four required baselines, and the distinct failure each one catches

Each baseline is a specific accusation against the model, and all four must be answered.

Baseline 1 — PERSISTENCE. `Y_hat_{k+1} = Y_k`. No parameters. *The accusation:* the model adds nothing over the observation that stocks are sticky. *What it catches:* the commonest and most embarrassing failure, and the reason a stock is treacherous to model. A quantity with a

four-year half-life barely moves period to period, so *it will be what it was* is an excellent forecast, and any model of a slow stock that cannot beat persistence has learned nothing about the stock. Persistence gets easier to beat as the period lengthens, because more happens between observations, so a large skill against persistence at a coarse time scale is not impressive on its own and must be read together with baselines 3 and 4.

Baseline 2 — DIRECT INPUT. `Y_hat_{k+1} = c_0 + c_1 F(U_k)`, with `c_0, c_1` fitted on the training data. No memory term. *The accusation:* `Y` is not a stock at all but a function of the input, and the memory term is decoration. *What it catches:* contemporaneous repackaging, species 1 of Section 5. If `Y` is an affine function of the interaction then the input alone predicts it and the accumulator adds nothing. For a genuine stock this baseline is hopeless, because the input carries no information about the *level* — in Section 9.4 it has a held-out RMSE of 79.4 rules against the model's 1.85, a skill of 97.7 per cent. **A direct-input baseline that performs well is not good news. It is the alarm for the order-proxy trap.**

Baseline 3 — MATCHED AUTOREGRESSIVE. `Y_hat_{k+1} = a_0 + a_1 Y_k`, with `a_0, a_1` fitted on the training data: two free parameters, matching the storage model's two. *The accusation:* the model is trend-following with extra steps, and the input term earns nothing that memory plus drift did not already give. *What it catches:* the failure that matters most for institutional work, and by far the hardest of the four to beat. An AR(1) with an intercept can represent any exponentially-approaching trend, and a declining institution *is* an exponentially-approaching trend. **Where a storage model cannot beat a matched autoregression, what has been measured is that the system is drifting — which was already known — and not that interaction is driving it, which is the entire claim.** The matching is essential: a baseline with fewer free parameters than the model is not a fair comparison, and the AR(1)-with-intercept has exactly two, as the storage model does. This is the baseline the confirmatory fit of Section 9 failed.

Baseline 4 — THE SHUFFLED-INPUT CONTROL. Refit and re-evaluate the *same* storage model with the input series `F(U_k)` randomly permuted in time, leaving `Y_k`, `C_hat_k`, the constraints and the split unchanged; repeat many times to build a null distribution of held-out RMSE, and locate the real fit within it. *The accusation:* `b` is not fitting the input's dynamics but its average level, and any series with the same distribution would have done as well. *What it catches:* the failure no other baseline can see. **It is the baseline most often omitted and most often decisive** — omitted because it is the only one requiring a loop rather than a formula, and decisive because it is the only test that asks whether the *timing* of interaction matters, which is the only thing distinguishing a driven accumulator from a bucket with a tap running at a constant average rate.

15.6.5 6.5 The shuffled-input control in detail

The procedure has five steps. Take the training pair set and permute the `F(U_k)` column, leaving `Y_k`, `C_hat_k` and `Y_{k+1}` in their original order. Refit the constrained model on the permuted training data, obtaining `(a*, b*)`. Permute the `F(U_k)` column of the *holdout* independently and compute held-out RMSE with `(a*, b*)`. Repeat at least 1,000 times; 2,000 is better and costs nothing. Report the mean and fifth percentile of the null RMSE distribution, the count of permutations achieving an RMSE at least as good as the real fit, and the resulting `p = (count + 1)/(n + 1)`.

Three details are commonly got wrong. **Refit under the permutation**: keeping the real `b` and permuting the input only at prediction time builds a much weaker null, because it tests whether the input helps *given a coefficient fitted to the true alignment*, which is nearly circular. **Permute the holdout independently**, or the holdout retains the alignment the null is supposed to destroy. **Do not permute `C_hat`**: the control isolates the input term, so cost stays aligned and stays at coefficient one, holding its explanatory contribution constant between the real fit and the null. That is what makes the test specific.

What the shuffle can and cannot detect has already been stated in Section 5.2 and is repeated here because this is where the suite is specified: **it detects that the input's timing carries**

information; it does not detect that `Y` is a filtered copy of `U`, because a filtered input is genuinely predicted by the input and the shuffle destroys that genuine prediction just as it would destroy a real one. **The shuffle is the empirical audit; the repackaging tests are the structural audit; they are different audits and both are required.** And the shuffle's own calibration is measured, not assumed: 0.108 against a nominal 0.05 for iid permutation, 0.120 for circular shift, 0.060 for block permutation [Curzi, 2026u]. Any `p` reported against the nominal threshold is therefore optimistic, and Section 12 carries that as a limitation of this paper's own headline `p = 0.0005`.

15.6.6 6.6 Reading the three rules together

```
FR1 fail  ->  the system is not a leaky accumulator. Stop. A different model is needed.
FR2 fail  ->  the model has no demonstrated predictive advantage. None may be claimed.
FR3 fail  ->  the model fits with parameters that are not this system's. Suspect the proxy.
```

The combinations are informative:

FR1	FR2	FR3	Reading
pass	pass	pass	Corroborated at this scale. Proceed to an intervention with a validated model.
pass	fail	pass	Parameters corroborated; predictive advantage not demonstrated. **Accounting frame only** — see Section 6.8.
pass	pass	fail	The classic order-proxy trap. Return to Appendix A, Form B.
pass	fail	fail	Uninformative. The series is too short, too noisy, or the proxy is wrong.
fail	any	any	Rejected. The storage law is not the right model for this system.

Note the second row, because it is the outcome the worked application produces and it is common in real institutional data. Corroborated parameters with no demonstrated predictive advantage is a real and useful state, and it licenses much less than people want it to.

15.6.7 6.7 Choosing the time scale before looking

Section 4.5 gave three constraints on the period. Here is the fourth fact, and the reason the choice must be frozen: **the time scale materially changes the result, and it changes it in a direction that can be anticipated.**

The structure is this. Audit error on a *level* does not accumulate — counting the register twice, six months apart instead of one week apart, leaves the counting error at each endpoint the same size — while the accumulated input between those endpoints has grown with the

interval. Coarsening the time scale therefore raises the signal-to-noise ratio of the `b F(U)` term, and destroys degrees of freedom linearly. There is an optimum, and it depends on quantities estimable before the holdout is examined:

```
signal per block  ~  B x eta x mean F(U)
noise per block   ~  sqrt(2) x (audit error SD)          (does not grow with B)
blocks available  =  T / B
```

The first line reads `mean F(U)`, and the distinction is not decoration. Coarsening the time scale aggregates `U` over `B` fine periods, and there are two ways to carry the input into the block: average the `U` values and take `F` once, or take `F` of each fine period and average the results. They are different numbers, and strict concavity fixes the direction:

```
mean of F(U_k)  <=  F( mean of U_k )        with equality only if U is constant within the block
```

Aggregate `F`, never aggregate `U`. Taking `F(mean U)` inflates the input term, and inflates it more the coarser the block — which means a preregistration that aggregates the wrong way will show the model improving at every increase in `B` for reasons having nothing to do with signal-to-noise, and the analyst will read a Jensen artefact as a discovery about the time scale. The gap deserves a name and a number: it is the **concavity penalty**, and in the worked office it is 0.0761 nats against an `F(mean U)` of 1.2262 — **6.2 per cent of the surplus**, lost purely to the variability of the geometry. The aggregation rule is declared alongside the aggregator itself and the penalty is reported with the fit.

`B` is chosen by writing the trade-off down and picking a value with a stated reason before the freeze. Legitimate reasons: `B` must exceed the conversion lag; `B` must give at least 8 holdout blocks; `B` must not push `delta` per block above 0.3. **Illegitimate reason: `B = 13` gave the best result.** In the worked application the answer at `B = 13` is materially better than at `B = 8` on every metric. It is also based on four holdout blocks, and it was not the preregistered choice, so it is reported as exploratory and licenses nothing. That is what the discipline costs, and paying it is the difference between a survey and an advertisement.

15.6.8 6.8 What a passing fit licenses, and what it does not

A pass licenses: using the fitted `a` and `b` as corroborated estimates of `delta` and `eta` for that system, at the fitted time scale, over the fitted regime; computing `h` from them; computing a maximum admissible increment `r <= h`; and stating a one-step prediction with an interval derived from the held-out residuals.

A pass does not license any of the following. *Multi-period forecasts* — the model was evaluated one step ahead, and a ten-step forecast of a near-unit-root series has an error that none of the discipline above has tested. *Causal claims* — the shuffle establishes that the timing of interaction carries information about subsequent order; it does not establish that intervening on interaction will produce order, which is a different claim requiring a different design. *Transfer to another system* — `delta`, `eta`, `N` and `C` are the parameters of *this* system, and the scale-free form of the law is a claim about model form rather than about parameter values [Curzi, 2026a,k]. *Transfer to another proxy* — a fit validated on live precedent rules says nothing about reviewer competence, cycle time, or public satisfaction. *Extrapolation past the ceiling* — `Q_t <= max{Q_0, eta ln N / delta}`, so a prediction above the uniform bound is a prediction the model has already forbidden.

15.6.9 6.9 Preregistration as a corruption control

A measure that becomes a target ceases to be a good measure [Goodhart, 1975], and the sociological form is sharper still: the more a quantitative indicator is used for social decision-making, the more it will distort and corrupt the processes it was meant to monitor [Campbell,

1979]. Both apply to `Y` immediately and mechanically, from the moment anyone in the organisation learns that `Y` is being counted.

The mechanism in the framework's own terms is precise. `Y` is a proxy for `Q`. Corruption is the process by which the correlation between `Y` and `Q` is destroyed by effort directed at `Y`. A review office that learns its live-rule count is the metric can raise the count without raising stored order at all: split existing rules, write rules nobody applies, restate one rule as three. The proxy rises, the order does not move, and every number downstream — `delta`, `eta`, `h` — becomes a measurement of the gaming rather than of the system. The self-fulfilling dynamic is the same one Merton [1948] described.

Five defences, in descending order of effectiveness. **Keep `Y` out of the reporting line** — the strongest by a wide margin, since a proxy not used to evaluate anyone is not worth gaming; where `Y` must be reported, report `h` and the verdict rather than `Y`. **Make the liveness correction standing practice**, because gamed units are created and never applied, so a liveness threshold removes most of them automatically and without anyone having to accuse anyone. **Preregister the proxy definition, including its corrections**, because corruption works by exploiting the *edges* of a definition and a definition frozen and dated before the incentive existed cannot be quietly widened. **Rotate the proxy** on a schedule known to the analyst and not to the surveyed — expensive, and it breaks series continuity, so reserve it for strong incentives. **Measure at least one proxy the organisation cannot influence** — in the worked office, appeal reversal counts come from an external appeals board, and a proxy the organisation does not control is a check on the ones it does.

Preregistration is therefore not only an epistemic device. It is the specific administrative countermeasure to measure corruption, because it fixes the definition of the measure at a time when nobody yet has a reason to distort it.

15.6.10 6.10 Preregistering an intervention

Everything above concerns modelling an existing system. Where the preregistration also covers an intervention, five further sections are mandatory.

The prediction, in numbers, with a direction and a magnitude. Not *improvement is expected*. The form is: `Y` will increase by at least `r` per period, measured at the preregistered time scale, over at least `k` periods, beginning no earlier than period `m` after the change. The lag `m` must be stated because a conversion lag exists and testing before it has elapsed is a test designed to fail.

The falsification criterion for the intervention, distinct from FR1–FR3. What observation would establish that the intervention did not work, stated as a number. The standard form: if the mean one-period change over periods `m` to `m + k` is below `r_min`, the intervention is judged ineffective and withdrawn. A second falsification should be preregistered alongside it because it is free: the parameter-free crossing time `t* = ln(u*/u_0)/ln(1 − lambda)` of Section 4.12, which yields a calendar date and contains no researcher degree of freedom, and which must be reported per site rather than pooled.

The stopping rules, of two kinds and both required. *Futility*: the condition under which the intervention stops because it is not working. *Harm*: the condition under which it stops immediately regardless of how well `Y` is doing. **Harm criteria must include at least one measure that is not `Y`**, because an intervention that raises the proxy while damaging the system will not stop itself.

The consent and declination clause. The framework's declination constraint is binding: no intervention that removes a person's ability to decline is in scope. The preregistration must therefore state, for each affected group, what is being changed, how they were informed, and by what mechanism an individual may decline without penalty. **Where the honest answer to the last is that they may not, the intervention is out of scope and the preregistration does not proceed.** This is not a formality appended to the science; it is a constraint on which interventions exist at all. The corpus's analysis of verification cost gives the reason such constraints have to be priced rather than assumed: verification is all-or-nothing with threshold `a*`

`= 1 − c_v*delta_0/(K − c_v)`, and at `c_v >= K` verification consumes the entire net input [Curzi, 2026s].

The reversibility statement. The exact procedure for returning the system to its prior state, who is authorised to execute it, how long it takes, and what is not recoverable. A sixth item is optional and strongly recommended: **the pre-mortem** — before the freeze, write the paragraph that would be written in a year explaining why the intervention failed. It costs twenty minutes and routinely surfaces the harm criterion nobody had thought of.

15.7 7 Specifications

Numbered, checkable requirements. A survey that does not satisfy all of Section 7.1 is not a survey and must not be used as the basis for an intervention; an analysis that does not satisfy all of Section 7.2 is not confirmatory. Nothing in either list is an empirical claim.

15.7.1 7.1 The survey

S24.1 The system boundary shall be stated in writing, including at least one named thing that is deliberately excluded and the reason for excluding it.

S24.2 All five operational choices of Section 3.2 shall be recorded before any measurement is taken, and shall be dated.

S24.3 The state encoder shall be specified to the level of "a second engineer could produce the same vectors": the probe set, the response format, the normalisation, and the administration protocol.

S24.4 The probe set shall contain at least 8 real cases drawn from the system's actual workload and shall be administered independently and in writing before discussion.

S24.5 The interaction aggregator shall be declared and badged as a modelling choice. If the mean pairwise aggregator is used, its identity with the participation-ratio effective count shall be noted so the choice is traceable to a standard diversity measure.

S24.6 Capacity `u` and realised `u` shall be reported as two separate numbers with their two separate instruments. Reporting only one is a specification violation.

S24.7 The order observable shall be tested against all five properties of §24.4, with the answers recorded.

S24.8 The four repackaging tests of §24.15 shall be run and their results recorded. Tests 1 and 2 shall be run before measurement begins.

S24.9 Test 1 shall be passed with zero shared source fields between `Y` and `U`.

S24.10 Test 2 shall be passed with a named different recorder, a named different medium, and a named different schedule.

S24.11 Where a zero-input episode exists in the record, Test 3 shall be run and its result plotted.

S24.12 Test 4 shall be run for the largest single-period increase in `Y` observed in the series. Any increase exceeding `eta ln N` shall be investigated before the survey is released.

S24.13 The time scale shall be declared with its calendar alignment and shall satisfy: period >= conversion lag; `delta` per period <= 0.3; and a stated rationale for the signal-to-noise trade-off of §24.5.

S24.14 Periods with `U_t = 0` shall be retained in the series and flagged, never dropped.

S24.15 `C` shall be denominated in `Y`-units. Every Route 2 line shall show its conversion rate and the justification for it.

S24.16 The attention tax shall appear as an explicit line in the cost tally, even if the estimate is coarse. A tally with no attention-tax line is incomplete.

S24.17 Turnover shall be assigned to `delta` or to `C` by a written rule, and shall not appear in both.

S24.18 Slack shall not be booked as cost.

S24.19 `delta` shall be estimated by Method D1 where a zero-input episode exists. Where D2 is used, the log-survival plot shall be inspected for linearity and non-linearity reported.

S24.20 `delta` shall be reported with its half-life and time constant in periods and in calendar time.

S24.21 `eta` shall be estimated by an attribution audit in which each counted `Y`-unit has a documented originating interaction. Unattributable units shall be excluded and their count reported.

S24.22 `N` shall be produced by explicit enumeration, merge, and drop steps, with each merge and drop recorded with its reason.

S24.23 Categories with pairwise `u < 0.10` shall be merged.

S24.24 `Q` shall be reported pre-correction, with the liveness correction and the applier correction shown separately, and post-correction.

S24.25 The line-item `C_hat` shall be reconciled against the cost implied by the observed series, and the residual reported. A residual exceeding the declared tolerance blocks release.

S24.26 `h`, the realised one-step change, and `Q_inf` shall all be computed and reported.

S24.27 A negative `Q_inf` shall be reported as negative, with the Theorem 3.5 failure clause quoted.

S24.28 Every number on the survey output page shall carry a provenance note naming the instrument and the date range.

S24.29 The survey shall state, in one sentence each, the three things that would most change its conclusions if measured differently.

S24.30 The survey shall not recommend an intervention. Recommendation is the function of the intervention protocol that follows in this series, and requires the preregistration of Section 6 first.

15.7.2 7.2 The preregistration

S25.1 A preregistration document shall exist, be dated, be signed, and be held in a copy by a person who is not the analyst.

S25.2 The freeze shall occur before any examination of holdout data.

S25.3 The model specification shall be exactly `Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + eps_t` with `0 <= a <= 1` and `b >= 0`, unless a deviation is declared at the freeze with its reason.

S25.4 `C_hat_t` shall enter with a fixed coefficient of one. Estimating a coefficient on cost is a declared deviation and changes the identification.

S25.5 No intercept shall be added unless declared at the freeze with a stated interpretation.

S25.6 The train/holdout split shall be by calendar time, never at random, and the boundary period shall be named.

S25.7 The holdout shall contain at least 8 estimation blocks. If it cannot, the analysis shall be labelled exploratory throughout.

S25.8 The unconstrained fit shall be computed and reported before the constrained fit.

S25.9 FR1 shall be evaluated on the unconstrained optimum, not on the constrained one.

S25.10 All four baselines shall be computed: persistence, direct-input, matched autoregressive, and shuffled-input.

S25.11 The matched autoregressive baseline shall have the same number of free parameters as the storage model.

S25.12 The shuffled-input control shall refit the model under each permutation, permute the holdout independently, and use at least 1,000 permutations.

S25.13 The shuffled-input control shall not permute `C_hat`.

S25.14 Raw held-out RMSE shall be reported for the model and for all four baselines, in `Y`-units, alongside any skill scores.

S25.15 The margins for FR2 shall be declared at the freeze. The recommended defaults are 10% RMSE reduction against each of the first three baselines and `p <= 0.05` against the shuffle.

S25.16 The tolerance for FR3 shall be declared at the freeze. The recommended default is +/-25% on each of `delta` and `eta`.

S25.17 FR1, FR2 and FR3 shall each be reported as PASS or FAIL individually. They shall not be averaged, weighted, or combined into a score.

S25.18 A failure of any one of the four baselines is a failure of FR2.

S25.19 The time scale shall be chosen and justified before the freeze. Results at other time scales may be reported only as exploratory and only with that label attached in the same sentence.

S25.20 The fitted `a` shall be reported with a profile interval, and the corresponding interval on the half-life shall be stated in calendar units.

S25.21 Held-out prediction shall be one step ahead using observed `Y_t`, not free-running.

S25.22 Residual autocorrelation shall be reported.

S25.23 Periods with `U_t = 0` shall be retained, with `F(0) = 0`.

S25.24 The outlier rule shall not reference residuals.

S25.25 Known structural breaks shall be declared at the freeze, with the handling rule.

S25.26 The verdict shall be written as one sentence naming the rule that failed, if any.

S25.27 A verdict of “parameters corroborated, predictive advantage not demonstrated” shall be reported in those words and shall not be paraphrased upward.

S25.28 Where the preregistration covers an intervention, it shall state the numeric prediction, the falsification criterion, both stopping rules, the consent and declination mechanism, and the reversibility procedure.

S25.29 At least one harm criterion shall be a measure other than `Y`.

S25.30 If no mechanism exists by which an affected individual may decline without penalty, the preregistration shall not proceed and the reason shall be recorded.

15.8 8 Failure modes

Each entry gives a description, what it damages, its signature, how it is detected, and the repair. They are the named faults for which the survey and the preregistration are responsible.

15.8.1 8.1 Faults of the survey

F24.1 — THE ORDER PROXY THAT WAS REALLY THE INPUT. *Description.* `Y` is a function of `U`, so the model fits an identity. *Damages.* Everything. The survey has no empirical content. *Signature.* Fitted `b` far outside the audited `eta`; fitted `a` implying an implausible half-life; `Y` collapses to zero during a shutdown; single-period jumps exceeding `eta ln N`. *Detection.* The four tests of §24.15. Tests 1 and 2 before measurement; 3 and 4 on the series. *Fix.* Redefine `Y` into a different medium with a different recorder. There is no statistical remedy; this is a definitional error and only a definitional change repairs it.

F24.2 — THE MISSING DENOMINATOR. *Description.* `C` is measured in hours or money while `Y` is measured in rules or modules, and `h` is computed anyway by implicitly treating them as commensurate. *Damages.* `C`, and therefore `h`. *Signature.* The headroom calculation has mixed units. The number is arithmetically well-formed and physically meaningless. *Detection.* Write the units next to every term in `h = eta ln N − delta Q − C`. They must match. *Fix.* §24.6. Route 1 where possible; a declared and justified Route 2 rate otherwise.

F24.3 — THE SNAPSHOT SURVEY. *Description.* `Q` is measured once. No series, therefore no `delta`, therefore no `h`. *Damages.* `delta`, and by omission the verdict. *Signature.* The survey reports a level and a judgement about whether it is “healthy”. *Detection.* Ask for the second measurement date. If there is not one, this is the fault. *Fix.* Reconstruct a retrospective series

from the register's own history. Most durable registers carry creation and strike dates and can be replayed.

F24.4 — THE MISSING ZERO-INPUT EPISODE. *Description.* No shutdown in the record, so `delta` is estimated by testimony. *Damages.* `delta`, and through `delta Q` the whole budget. *Signature.* `delta` reported as a round number with no derivation. *Detection.* Check Form D for a named episode. *Fix.* Method D2, the cohort decay curve, which needs no episode. Failing that, run a deliberate, bounded, reversible zero-input trial on a subset, under the pilot rules of the intervention protocol.

F24.5 — ATTENDANCE ENCODED AS PARTICIPATION. *Description.* Categories present but silent are counted in `U_t`. *Damages.* `u`, upward. Inflates `F(U)` and hides the realisation deficit. *Signature.* Realised `u` equals capacity `u` exactly, which essentially never happens. *Detection.* Compare the participation rule in A12 against the actual record of who spoke. *Fix.* Require a recorded disposition. A category with no recorded position is not in that period's geometry.

F24.6 — THE PERMISSIVE ATTRIBUTION AUDIT. *Description.* `eta` is computed by dividing all new `Y`-units by all surplus, without checking which units actually came from interaction. *Damages.* `eta`, upward, often by a large factor. *Signature.* `eta` implausibly high relative to the observed trajectory; the fitted `b` of Section 6 comes out far below the audited `eta`. *Detection.* Recount with the strict rule: documented originating interaction or excluded. *Fix.* Redo Form E. Report the unattributable count; it is itself a useful diagnostic, because a large unattributable fraction means the system's order is being made somewhere the survey is not looking.

F24.7 — CATEGORY INFLATION. *Description.* `N` counted from the org chart rather than from the geometry, so relabelled duplicates are counted separately. *Damages.* `N` upward, and therefore `eta ln N` and `h` upward. *Signature.* Two categories in the pairwise table with `u < 0.10` and separate rows. *Detection.* Form C, the merge step. *Fix.* Merge, and record that the merge has *reduced* the measured headroom, which is the survey doing its job.

F24.8 — THE UNRECONCILED TALLY. *Description.* The cost line items are summed and never checked against the series. *Damages.* `C`, in either direction. *Signature.* Form F3 blank. *Detection.* Compute the cost implied by the observed trajectory given `delta`, `eta` and `F(U)`, and compare. *Fix.* Reconcile. A large residual usually means an omitted cost category, not an arithmetic error.

F24.9 — GOODHART CAPTURE OF THE PROXY. *Description.* Once `Y` is known to be measured, the organisation optimises `Y` directly. *Damages.* The validity of `Y` as a proxy for `Q`, progressively, over time. *Signature.* `Y` rises while every qualitative indicator falls. Newly created units are small, uniform, and never applied. *Detection.* The liveness audit (§24.12) is the counter-instrument: Goodharted units are created and never used. *Fix.* Section 6.7. In the survey, the immediate mitigations are to keep `Y` in a register that is not reported upward, to include the liveness correction as standing practice, and to rotate the proxy definition on a schedule known to the surveyor and not to the surveyed.

F24.10 — THE INVISIBLE MAINTENANCE REGIME. *Description.* Everything that works is invisible, so the survey under-measures both `Q` and the effort holding it (Star and Ruhleder 1996). *Damages.* `Q` downward and `C` downward — which can mask an unsustainable `h` behind two compensating errors. *Signature.* No records exist for the parts of the system nobody complains about. *Detection.* For each `Y`-unit class, ask what keeps it alive and who. Expect a person. *Fix.* Count the named people as a maintenance line in `C` and note the concentration risk. A maintenance regime resting on one un-deputised person is an orphaning event waiting for a date.

15.8.2 8.2 Faults of the preregistration and the fit

F25.1 — THE MOVED GOALPOST. *Description.* A tolerance or margin is adjusted after the held-out numbers are seen. *Damages.* The entire evidential status of the analysis. *Signature.* A reported margin that is not a round number, or that differs from the preregistration by a small amount in exactly the direction needed. *Detection.* Compare the report to the frozen document held by the independent holder. This is what the independent holder is for. *Fix.*

Relabel as exploratory, in every sentence, and rerun the confirmatory analysis on a fresh holdout in the next cycle.

F25.2 — THE OMITTED SHUFFLE. *Description.* Three baselines are reported and the shuffled-input control is not. *Damages.* The claim that interaction is doing anything. *Signature.* A results table with exactly three baseline rows. *Detection.* Trivial, and it is the first thing to look for in any report using this framework. *Fix.* Run it. It takes a loop and a minute of compute. Expect it to change the conclusion in a meaningful fraction of cases.

F25.3 — THE UNMATCHED AUTOREGRESSION. *Description.* The AR baseline is given fewer free parameters than the storage model, usually by dropping the intercept. *Damages.* FR2, in the model's favour. *Signature.* An AR baseline described as "AR(1)" with one coefficient reported. *Detection.* Count the parameters on both sides. *Fix.* Restore the intercept and rerun. Expect the model's apparent advantage to shrink substantially; on slow stocks it usually does.

F25.4 — MULTI-STEP FLATTERY. *Description.* Held-out performance is computed on free-running multi-step forecasts, where persistence and AR degrade fast and the model looks heroic. *Damages.* FR2, in the model's favour, sometimes by an order of magnitude. *Signature.* Skill scores above 0.9 against persistence on a stock with a long half-life. *Detection.* Ask which `Y_t` was used at each step: the observed one, or the model's own previous prediction. *Fix.* One-step-ahead with observed `Y_t`. If multi-step performance is genuinely of interest, preregister it separately, with its own baselines evaluated the same way.

F25.5 — TIME-SCALE SHOPPING. *Description.* Several block lengths are tried and the best is reported as the analysis. *Damages.* Everything; this is `p`-hacking with an unusually large search space, because the block length also changes the effective sample size. *Signature.* A block length with no stated justification, or a justification that could equally have supported two other values. *Detection.* Check Part 2.3 of the preregistration for a justification that does not reference results. *Fix.* Report the preregistered scale as confirmatory and all others as exploratory, in the same table, with the label in the row.

F25.6 — THE COST COEFFICIENT. *Description.* `C_hat` is entered as a regressor with a free coefficient because the fit is better. *Damages.* The identification. With a free coefficient on cost, `b` is no longer `eta` and `h` can no longer be computed from the fit. *Signature.* Three fitted coefficients reported instead of two. *Detection.* Count them. *Fix.* Fix the coefficient at one. If the fit is materially worse, that is information: the cost observable is mis-denominated, and the repair is in Section 4.6, not in the regression.

F25.7 — THE RESIDUAL-BASED OUTLIER RULE. *Description.* Periods with large residuals are excluded as "data quality issues". *Damages.* Guarantees a good fit and destroys FR2. *Signature.* An exclusion list whose members are all large residuals. *Detection.* Ask what the exclusions have in common in the *record*. If the answer is only "the model predicted them badly", this is the fault. *Fix.* Restore them. Investigate them individually against the record; genuine restatements can be excluded with a documented, record-based reason.

F25.8 — THE UNLABELLED EXPLORATORY RESULT. *Description.* An analysis run after the freeze, or after a deviation, is reported alongside confirmatory results without a label. *Damages.* The reader's ability to weight the evidence. *Signature.* A results section in which every number has equal typographic status. *Detection.* Cross-reference every reported number against the frozen document. *Fix.* Label in the sentence, not in a footnote. The convention that works: the word "exploratory" appears in the same sentence as the number.

F25.9 — THE ETHICS APPENDIX. *Description.* Consent, declination, stopping and reversibility are written after the design is settled, as a compliance step. *Damages.* The design itself, invisibly, because an intervention designed without a declination mechanism usually cannot have one retrofitted. *Signature.* Part 5.8 of the form completed last, in different handwriting. *Detection.* Ask what the design would have been if declination had been a constraint from the start. If the answer is "different", the fault is present. *Fix.* Redesign. **The declination constraint is not a review gate; it is a constraint on the space of admissible designs, and it belongs at the front of the process.**

F25.10 — SUCCESS DECLARED FROM AN ACCOUNTING FRAME. *Description.* FR2 fails, FR3 passes, and the report says the model "was validated". *Damages.* Everything downstream, be-

cause the intervention protocol's pilot rules depend on which of these happened. *Signature.* The phrase "broadly consistent with" doing load-bearing work. *Detection.* Look for the individual PASS/FAIL lines. If they are not individually reported, S25.17 has been violated. *Fix.* Report the verdict in the prescribed words: *parameters corroborated, predictive advantage not demonstrated.* Then apply the licence list of Section 6.8, which permits computing `h` and forbids claiming a forecast.

15.9 9 Results: one system surveyed, preregistered, and tested

15.9.1 9.1 The system, and what kind of object it is

Meridian County Development Review Office. A municipal development review office: 26 staff, five review disciplines, a weekly cross-discipline review board, and a standing review manual of written precedent rules that reviewers are required to apply.

Badge — this is a constructed composite, not a field study. The organisation is invented. The arithmetic is real, computed, and reproducible: the audit series is a simulation with declared generating parameters, produced so that every number in this paper can be checked by the reader. **It must not be cited as evidence about any real review office.** Its function is to show the procedure operating end to end on numbers that behave the way real numbers behave — including the places where they refuse to cooperate.

The five choices as declared:

#	Choice	As declared for MCDRO
1	State encoder	Encoder A. Probe set of `m = 12` real development applications from the previous two years, each with a forced approve/refuse disposition, recorded independently and in writing by each discipline lead before any discussion. Vector = the 12 signs, normalised.
2	Interaction observable	Mean pairwise `1 - (x_i . x_j)^2` over the disciplines that recorded a disposition in that week's board. Participation = a recorded disposition, not attendance. Badged as a declared aggregator.
3	Order observable	`Y` = count of **live precedent rules in the standing review manual**. A rule is live if it is in the register, not superseded, and has been applied at least once in the previous 156 weeks. Maintained by the code administrator, who does not attend the board.

#	Choice	As declared for MCDRO
4	Time scale	One week. Audit taken every Friday. Conversion lag measured at up to 6 weeks (drafting plus the 10-working-day publication requirement).
5	Cost observable	Route 1 for appeal strikes, orphaning, and contradiction withdrawals. Route 2 for rework and the attention tax, at a declared displacement rate.

The repackaging audit was run first, before measurement began:

Test	Result
1 Algebraic derivation	`Y` fields: rule id, adoption date, supersession date, last-applied date. `U` fields: board date, discipline, recorded disposition. **Shared fields: 0. PASS.**
2 Instrument independence	`Y` recorded by the code administrator into the rule register, audited Fridays. `U` recorded by the board clerk into the minutes, at the session. Different person, medium, schedule. **PASS.**
3 Zero-input persistence	Weeks 121-154, board suspended by a budget freeze. `Y`: 206 at week 121, 186 at week 154, 186 at week 155 when the board resumed. Smooth monotone decay through the episode, continuous across resumption. **PASS.**
4 Ceiling	Largest single-week rise observed: 2 rules. Ceiling `eta ln N = 0.5015 x 1.6094 = 0.8071` per week; over the 6-week conversion lag the admissible burst is well above 2. **PASS.**

Only then did measurement begin.

15.9.2 9.2 The survey: six numbers and one sign

Categories and geometry.

Five disciplines survive enumeration: planning (PLAN), structural and building (STRU), fire and life safety (FIRE), environmental (ENVR), transportation (TRAN). No merges: the smallest pairwise `u` is 0.3056, comfortably above the 0.10 merge threshold. No drops: all five have a route into the board.

```
N = 5       ln N = 1.6094
```

pair	opposed cases k/12	d	cos = 1 - 2d	u = 1 - cos^2
PLAN-STRU	3/12	0.2500	0.5000	0.7500
PLAN-FIRE	2/12	0.1667	0.6667	0.5556
PLAN-ENVR	5/12	0.4167	0.1667	0.9722
PLAN-TRAN	4/12	0.3333	0.3333	0.8889
STRU-FIRE	1/12	0.0833	0.8333	0.3056
STRU-ENVR	3/12	0.2500	0.5000	0.7500
STRU-TRAN	2/12	0.1667	0.6667	0.5556
FIRE-ENVR	3/12	0.2500	0.5000	0.7500
FIRE-TRAN	2/12	0.1667	0.6667	0.5556
ENVR-TRAN	4/12	0.3333	0.3333	0.8889

```
U capacity = 0.6972
N_eff capacity = 1 + 4(0.6972) = 3.7889
F capacity = ln 3.7889 = 1.3321 nats
```

Realised, from the 278 operating weeks of the audit series:

```
U realised = 0.6019      F realised = 1.1498 nats
realisation deficit = 1.3321 − 1.1498 = 0.1823 nats
```

The cause is visible in the record: **32.7% of board sessions ran with fewer than the full set of disciplines recording a disposition.** Those sessions have a mean `u` of 0.2071 against 0.7940 for full sessions. The office has five genuinely non-parallel disciplines and uses about four-fifths of them.

Note the FIRE-STRU pair at `u = 0.3056`. That is the closest thing to redundancy in the set and is exactly where a future reorganisation would be tempted to merge two roles. Merging them would drop `N` from 5 to 4 and cost `ln 5 − ln 4 = 0.2231` nats of ceiling — which, at the office's current `eta`, is 0.112 rules per week, or about 5.8 rules a year, given away permanently.

Decay.

Method D1. The budget freeze of weeks 121-154 suspended the board entirely; no dispositions, no new rules, and — because intake was also halted — no appeals, rework, or contradiction maintenance, so `C = 0` as well as `U = 0`.

```
Y at week 121 = 206
Y at week 154 = 186
n = 33 steps
ratio = 186/206 = 0.90291
delta = 1 − 0.90291^(1/33) = 0.003090 per week

half-life = ln 2 / (−ln(1 − 0.003090)) = 224.0 weeks = 4.31 years
tau = 1/delta = 323.6 weeks = 6.22 years
3 tau = 971 weeks = 18.7 years
```

The office's stored order has a four-and-a-third-year half-life, and its settling time to a new equilibrium is nearly nineteen years. Both facts matter to any proposed intervention. The first says the manual is genuinely durable. The second says nobody now working in the office will see the equilibrium of any change made today.

Conversion.

Window: the last 26 operating weeks of the series.

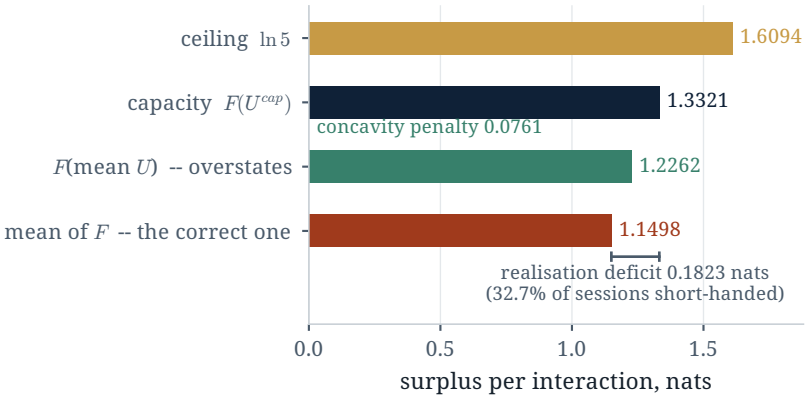


Figure 15.1: Capacity, and the two different realised numbers. The office's five disciplines are genuinely non-parallel: capacity `U = 0.6972` gives `N_eff = 3.7889` against a roster of 5 and `F = 1.3321` nats against the ceiling `ln 5 = 1.6094`. Realised surplus must be computed as the **mean of `F`** over the 278 operating weeks, 1.1498 nats, and not as `F` of the mean `U`, 1.2262 nats; the 0.0761-nat gap between them is the concavity penalty, 6.2 per cent of the surplus, lost to the variability of the geometry alone. Against capacity, the realisation deficit is 0.1823 nats, traced to 32.7 per cent of sessions running short-handed. Measured from the pairwise table and the period record above; generated by `figs/src/CCP-015_fig01.py`.

```
sum of F(U_t) over the window = 31.903 nats
Y-units created in the window = 26 (register count)
of these, attributable with a documented originating board decision = 16
unattributable, excluded = 10

eta = 16 / 31.903 = 0.5015 rules per nat
```

The 10 unattributable rules are themselves a finding: they were written by individual reviewers without a board decision behind them. They are real rules and they are in `Y`, but they were not produced by interaction, so they cannot be used to estimate the conversion of interaction into structure. **A high unattributable fraction is not a measurement nuisance. It reports that the board is not the thing making the order, which is a substantive claim about the office.**

Cost.

line	category	route	Y-units/week
1	Appeal reversals that strike a live rule	1	0.14
2	Rules orphaned by departure (no qualified applier remains)	1	0.11
3	Rework: items re-reviewed after disposition	2	0.16
4	Contradiction maintenance	1	0.09

line	category	route	Y-units/week
5	Attention tax at the declared conversion rate	2	0.13
	C_hat		**0.63**

Reconciliation: the cost implied by the observed series over the operating weeks is 0.6275 Y-units per week. Residual against the line-item tally: **0.0025**, against a declared tolerance of 0.05. The tally may be used.

Slack is not booked. Turnover is assigned by the standing rule: routine attrition sits in `delta`; the two departures that orphaned specific rules sit in line 2 of `C`.

Stock, headroom, and the verdict.

```
Y register count at survey date (week 312)  = 121
liveness correction (no application in 156 weeks) = −6
applier correction (orphaned, no qualified applier) = −2
Q = 113

eta ln N   = 0.5015 x 1.6094 = 0.8071
delta Q    = 0.003090 x 113  = 0.3492
C_hat      =                   0.6300
----------------------------------------
h          = 0.8071 − 0.3492 − 0.6300 = −0.1720
```

VERDICT: h < 0. Unsustainable in this model. No interaction geometry saves it.

This is Corollary 3.4 in its operative form. `C_hat = 0.63` exceeds `eta ln N − delta Q = 0.4579`, and therefore **no rearrangement of who sits on the board, in what order they speak, or how vigorously they disagree can prevent the manual from declining.** Even a perfectly orthogonal board running every week at `U = 1` would return `eta ln N − delta Q − C = −0.1720` rules per week.

The realised picture:

```
one-step = eta F_real − delta Q − C_hat
         = 0.5015 x 1.1498 − 0.3492 − 0.6300
         = 0.5766 − 0.9792
         = −0.4025 rules per week  =  −20.93 rules per year

realisation deficit = h − one-step = −0.1720 − (−0.4025) = 0.2305 rules/week
```

And the steady state, Theorem 3.5:

```
Q_inf = (eta F_real − C_hat) / delta = (0.5766 − 0.6300)/0.003090 = −17.3
```

`Q_inf` is negative. Reported as negative, per S24.27. Theorem 3.5 does not hide this failure mode: the equilibrium is negative, so the regime is unsustainable or the linear model must be replaced. Operationally it means the manual is being driven to the boundary of the model's validity — to zero live rules — and that the linear model says nothing useful below that boundary.

Projected trajectory at current parameters, from `Q = 113`:

horizon	Q
1 year	93.6
3 years	63.1
5 years	41.0
12.6 years	0

15.9.3 9.3 Why the decline was invisible

The office lost 177 rules over the 312 weeks of the audit series — from 290 down to 113, a decline of 61%. At the observed rate of about 21 rules a year out of a stock of roughly 150, the quarter-to-quarter change is under 2%, which is inside the noise of the register and well inside the noise of anyone's impression. **The decline was invisible at every time scale on which the office actually looks at itself, and unmistakable on the one time scale nobody was using.**

That is the argument for the survey in a single sentence. The office's staff were not negligent and were not wrong about their own experience. They were reading a slow variable at a fast cadence.

The three measurements that would most change this conclusion.

1. **The attention-tax conversion rate.** Line 5 of the cost tally, at 0.13 `Y`-units per week, is the only line resting on a declared conversion rather than a direct count. If the true rate is half what was declared, `h` rises to -0.107 — still negative, but the margin changes materially. Measuring displacement directly, by timing rule production against interrupted and uninterrupted weeks, would settle it.
2. **The unattributable fraction in the eta audit.** Ten of 26 new rules had no documented board origin. If half of those were in fact board-originated and simply undocumented, `eta` rises to about 0.658 and `h` becomes -0.0424 — close enough to zero to change the verdict from "unsustainable" to "maintenance only", which is a different regime of the intervention protocol entirely.
3. **The liveness threshold.** `Y` uses 156 weeks of non-application as the liveness cutoff. At 104 weeks the count would be lower and `delta Q` smaller; at 260 weeks, higher. The cutoff was declared in advance and must not be moved after seeing the result, but its sensitivity should be reported, and it is: moving the cutoff by one year in either direction changes `Q` by about 5 rules and `h` by about 0.015.

None of these three would reverse the sign of `h` on its own. Two of them together would. **That is the honest statement of how firm this survey's verdict is, and it belongs in the survey rather than in a footnote.**

What the survey does not say. It does not say the office is badly run. It does not say the board is useless. It does not recommend anything. It reports six numbers and one sign. Everything downstream — whether the model can be trusted at all, and what may be built if it can — depends on the fact that the sign is negative and on the fact that the survey can now state exactly what it would take to change it.

15.9.4 9.4 The preregistered design, the fit, and held-out performance

```
Y            = live precedent rules in the standing review manual
N            = 5,  ln N = 1.6094
delta        = 0.003090 per week   (Method D1, weeks 121–154 freeze)
eta          = 0.5015 rules per nat (attribution audit, 26 weeks)
C_hat        = 0.63 rules per week
Q at survey  = 113
h            = −0.1720            VERDICT: h < 0
```

```
Repackaging tests 1–4: all PASS

Series:           weeks 1 to 312
Block length:     B = 8 (two-month blocks)
Justification:    conversion lag measured at up to 6 weeks, so B must exceed 6;
                  B = 8 leaves 8 holdout blocks at the declared split, meeting S25.7;
                  delta per block = 1 - (1-0.003090)^8 = 0.0244, well under 0.3.
Split:            train = blocks beginning on or before week 240 (n = 29)
                  holdout = blocks beginning after week 240 (n = 8)
Model:            Y_{k+1} = a Y_k + b F(U_k) - C_hat_k + eps_k,  0<=a<=1, b>=0
FR2 margins:      10% RMSE against persistence, direct input and matched AR;
                  shuffle p <= 0.05
FR3 tolerance:    +/- 25% on delta and on eta; conversion-lag allowance 6 weeks
Permutations:     2000
```

Frozen and signed before the holdout was examined.

```
Unconstrained optimum:   a = 0.97439    b = 0.5859
Constrained optimum:     a = 0.97439    b = 0.5859     (interior; constraints not binding)

delta implied per block = 1 - 0.97439 = 0.02561
delta implied per week  = 1 - 0.97439^(1/8) = 0.003237
half-life implied       = 4.11 years
```

Held-out performance, one step ahead, over the 8 holdout blocks:

Model	held-out RMSE (rules)	skill vs this baseline
Storage model (frozen)	**1.852**	—
Persistence	3.102	+0.403
Direct input	79.397	+0.977
Matched AR(1) with intercept	1.906	+0.028
Shuffled input, mean of 2000	3.400	—
Shuffled input, 5th percentile	2.657	—
Shuffled input, permutations beating the fit	0 of 2000	`p = 0.0005`

15.9.5 9.5 The three rules evaluated, and the one that failed

FR1 — ADMISSIBLE BOX. The unconstrained optimum is `a = 0.97439`, `b = 0.5859`. Both lie strictly inside the box; the constraints are not binding. **PASS.** The office's manual behaves like a leaky accumulator: it retains most of what it has and gains in proportion to a non-negative multiple of the surplus input.

FR2 — BASELINE DOMINANCE.

Baseline	required margin	achieved	result
Persistence	10% RMSE reduction	40.3%	PASS
Direct input	10% RMSE reduction	97.7%	PASS
Matched AR(1)	10% RMSE reduction	**2.8%**	**FAIL**
Shuffled input	`p <= 0.05`	`p = 0.0005`	PASS

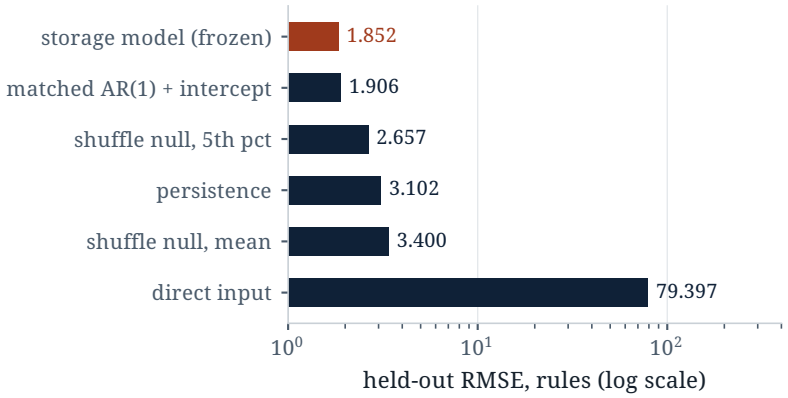


Figure 15.2: Raw held-out error, one step ahead, on a logarithmic scale because the direct-input baseline is two orders of magnitude worse than the rest. The storage model's 1.852 rules sits just below the matched AR(1)'s 1.906 — a 2.8 per cent reduction — and well below persistence at 3.102 and the shuffle null's mean of 3.400 and fifth percentile of 2.657. Measured from the held-out table above; generated by `figs/src/CCP-015_fig02.py`.

FR2 FAILS. The storage model beats a matched autoregression by 2.8% on held-out RMSE, against a preregistered requirement of 10%. Under S25.18, failure against any one of the four baselines is a failure of the rule, and the rule is not averaged.

FR3 — INDEPENDENT-PARAMETER AGREEMENT.

Parameter	survey (independent)	fit	difference	tolerance	result
`delta` per week	0.003090	0.003237	+4.8%	+/-25%	PASS
`eta` (= `b`)	0.5015	0.5859	+16.8%	+/-25%	PASS
Ceiling check	max weekly rise 2 rules vs `eta ln N = 0.8071` over a 6-week lag allowance	—	—	—	PASS

FR3 PASSES. Note the direction of the `eta` discrepancy: the fit is 16.8% *higher* than the attribution audit. That is the expected direction, and it is consistent with two things at once — the audit's deliberate strictness (unattributable units excluded, Section 9.2) and the honesty clause that the surplus functional as deployed is a lower bound. **We log the floor, not the ceiling**, so a fitted conversion slightly above the audited one is not an anomaly to be explained away.

The verdict sentence (S25.26)

The storage law is not corroborated as a predictive model of live precedent rules in the Meridian County Development Review Office at the preregistered eight-week time scale; it failed FR2 against the matched autoregressive baseline at a margin of 2.8% against a required 10%, while passing FR1 and FR3.

Row two of the table in Section 6.6: **parameters corroborated, predictive advantage not demonstrated.**

15.9.6 9.6 Exploratory results at other time scales

Reported for completeness. **Every row except `B = 8` is exploratory and licenses nothing.**

Transposed so that every figure fits the measure; no value is altered. The confirmatory

quantity	B = 1	B = 2	B = 4	**B = 8**	B = 13
RMSE, storage model	1.372	1.513	1.541	**1.852**	1.209
RMSE, persistence	1.439	1.558	1.863	**3.102**	4.637
RMSE, direct input	80.015	79.920	79.882	**79.397**	76.914
RMSE, matched AR(1)	1.387	1.386	1.450	**1.906**	1.980
skill vs persistence	+0.046	+0.029	+0.173	**+0.403**	+0.739
skill vs matched AR(1)	+0.011	-0.091	-0.063	**+0.028**	+0.389
shuffle `p`	0.0015	0.0030	0.0005	**0.0005**	0.0155

Four observations, all of which are why S25.19 exists.

1. **The shuffled-input control passes at every scale**, at `p` between 0.0005 and 0.0155. The timing of interaction carries real information about the subsequent rule count at all five scales tested. This is the strongest single positive result in the analysis.
2. The matched autoregression is beaten at only two scales, and at `B = 2` and `B = 4` it beats the storage model outright. A declining stock is an exponentially-approaching trend, and an AR(1) with an intercept represents that trend exactly. This is the hard baseline, and it is hard for a structural reason rather than an accidental one.
3. **`B = 13` would have passed everything**, at +73.9% against persistence and +38.9% against the matched autoregression. It was not the preregistered scale and it rests on four hold-out blocks. Reporting it as the result would have been the exact fault named in F25.5. It is reported here as exploratory, and the correct action it licenses is not a claim but a plan: extend the series and preregister `B = 13` for the next cycle, with the holdout requirement met.
4. **The implied `delta` is stable across scales** — 0.00289 to 0.00373 per week, against a directly measured 0.00309 — which is a genuine and reassuring consistency, and is FR3 doing its job independently of FR2's outcome.

15.9.7 9.7 Two repackaging demonstrations, one of which passed everything

The office was, at one point in the survey, offered two alternative order observables. Both were tested. Both are instructive.

Candidate PD-1 — "items dispositioned per session" (contemporaneous repackaging). Proposed because the minutes already contain it. It shares source fields with `U` and fails Test 1 of Section 5.1 immediately, so it never reached the estimation stage — but it was fitted anyway, as a demonstration.

Model	held-out RMSE
Storage model	3.079
Persistence	3.314
Direct input	**2.443**
Matched AR(1)	2.657

Model	held-out RMSE

Fitted `a = 0.7106`, `b = 13.759`; shuffle `p = 0.2278`.

The direct-input baseline beats the storage model outright — a skill of -0.261. That is baseline 2 performing its designed function: when the input alone predicts the order observable better than an accumulator driven by the input, the order observable is a function of the input. The shuffle also fails, at `p = 0.2278`. PD-1 is caught twice over and would have been caught for free by Test 1 of Section 5.1.

Candidate PD-2 — "the board disagreement index", a rolling weekly index published by the board clerk from the minutes (filtered repackaging). This one is dangerous.

Model	held-out RMSE	skill
Storage model	**0.833**	—
Persistence	7.870	+0.894
Direct input	12.940	+0.936
Matched AR(1)	7.720	+0.892
Shuffled input, mean of 2000	7.842	`p = 0.0005`

PD-2 passes all four baselines, decisively, on every margin. It beats persistence by 89.4%, the matched autoregression by 89.2%, direct input by 93.6%, and the shuffled-input control at `p = 0.0005`. Against the actual proxy's marginal 2.8% against the autoregression, PD-2 looks like the far better model of the office.

PD-2 is a rolling average of the input. It stores nothing. There is no institution behind it.

What catches it:

Test	Result
FR3, `eta`	fitted `b = 20.205` against an audited `eta = 0.5015` — a factor of **40**, far outside the +/-25% tolerance. **FAIL.**
FR3, `delta`	fitted `a = 0.7038`, implying a half-life of **1.97 weeks** against a directly measured **224 weeks** — a factor of **114**. **FAIL.**
Test 3, zero-input persistence	across the 34-week freeze, PD-2 fell 82.1 -> 57.9 -> 28.5 -> 6.7 -> 1.6 -> 0.1, and then returned to **31.6 in the single week the board resumed**. The true proxy went 206 -> 186 and continued from 186. **FAIL.**
Test 4, the ceiling	that single-week jump is 31.1 units. The largest admissible one-step increment is `eta ln N = 0.8071`. The observed jump is **38.5 times the theoretical maximum**. **FAIL.**

A stock cannot teleport. The ceiling test is arithmetic, requires no statistics, takes one line, and is the single most reliable detector of the order-proxy trap in this protocol.

And the lesson that governs the architecture of these three chapters: **the shuffled-input control is the empirical audit and it cannot see repackaging; the structural tests of Section 5 and FR3 are the definitional audit and they can. Both are required, and they are required in that order, because the structural audit is free and the empirical one is not.**

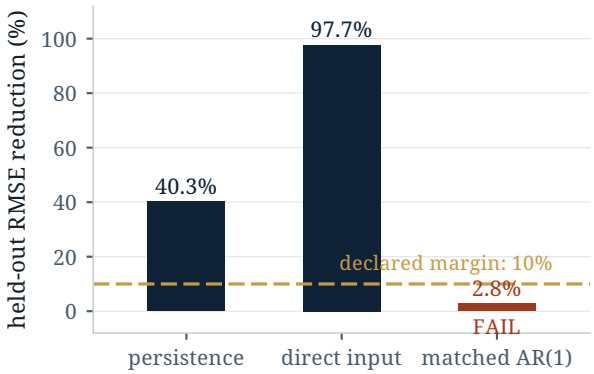


Figure 15.3: The preregistered rule was not met. The storage model's held-out RMSE reduction against each baseline is plotted against the margin declared before the holdout was examined. Persistence (40.3 per cent) and direct input (97.7 per cent) clear the 10 per cent requirement comfortably; the matched AR(1) with intercept — the baseline with the same two free parameters as the model — is beaten by **2.8 per cent**, which does not reach the line. Under the rule that failure against any one of the four baselines is a failure of the rule, FR2 fails. Measured from the held-out RMSE table of Section 9.4; generated by `figs/src/CCP-015_fig03.py`.

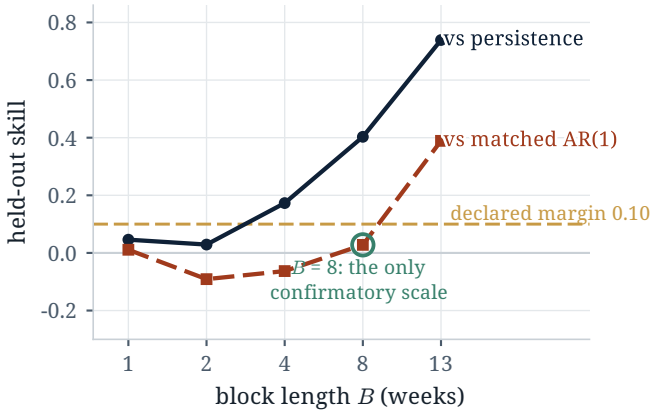


Figure 15.4: The scale that would have passed was not the scale that was declared. Held-out skill against persistence and against the matched autoregression at five block lengths; only `B = 8` is confirmatory and every other row is exploratory. `B = 13` returns +73.9 per cent and +38.9 per cent and would have cleared every margin — on four holdout blocks, against the eight the specification requires — while `B = 2` and `B = 4` are beaten outright by the autoregression. Measured from the exploratory table above; generated by `figs/src/CCP-015_fig04.py`.

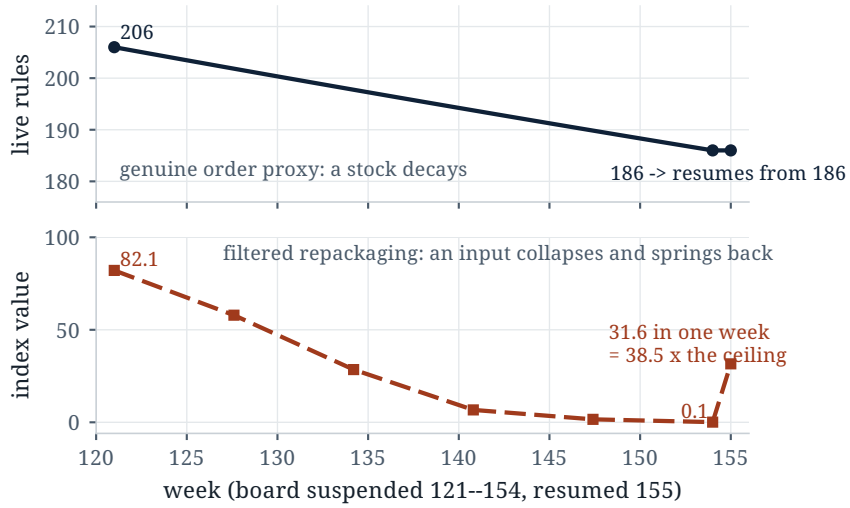


Figure 15.5: A repackaged proxy is caught by a stock's own arithmetic, not by its statistics. Across the 34-week zero-input freeze the genuine order proxy decays smoothly from 206 to 186 and resumes from 186; the filtered repackaging collapses from 82.1 to 0.1 and springs back to 31.6 in the single week the board resumes — a jump of 31.1 units against a maximum admissible one-step increment of `eta ln N = 0.8071`, which is 38.5 times the theoretical maximum. Plotted as small multiples rather than on shared axes because the two series are not commensurable. Measured from the freeze record of Section 9.7; generated by `figs/src/CCP-015_fig05.py`.

15.9.8 9.8 What the office is licensed to do

Under the licence list of Section 6.8, with FR1 and FR3 passed and FR2 failed:

Licensed. Treat `delta = 0.003090`, `eta = 0.5015`, `N = 5`, `C_hat = 0.63` and `Q = 113` as corroborated parameters of the system at the eight-week scale. Compute `h = −0.1720`. Conclude that the arrangement is unsustainable in this model. Use `h` as a budget constraint on any proposed increment.

Not licensed. Claim that the storage law forecasts the manual's rule count better than a trend line does. Claim that raising interaction will raise the rule count, on the strength of this analysis — the shuffle establishes that the timing of interaction carries information, not that intervening on it produces order. Forecast beyond one step. Extend any of it to cycle time, staff competence, or applicant satisfaction.

Required next. Any intervention proceeds on the accounting frame only: a capped, reversible pilot, preregistered under Part 5 of the form in Appendix B, whose own falsification criterion is stated in numbers before it begins — and which, if it works, supplies exactly the causal evidence this analysis is not entitled to assume.

15.10 10 Discussion

15.10.1 10.1 What the failure means, and what it does not

The preregistered confirmatory fit did not meet its own second falsification rule. The storage model beat a matched autoregression by 2.8 per cent on held-out RMSE against a margin of 10 per cent declared before the holdout was examined. Under the rule that failure against any one of the four baselines is a failure of the rule, and the rule is not averaged, FR2 fails. **The storage law is therefore not corroborated as a predictive model of live precedent rules in that office at the preregistered eight-week time scale.** That sentence is the result of this paper and

it is not softened anywhere in it.

What the failure does **not** mean is worth stating with equal precision, because both over-reading and under-reading it are available. It does not mean the storage law is false; a single system at a single time scale cannot establish that. It does not mean the parameters are wrong: FR1 passed with the unconstrained optimum strictly interior to the admissible box, and FR3 passed with implied decay within 4.8 per cent and implied conversion within 16.8 per cent of independently measured values. And it does not mean the timing of interaction is uninformative: the shuffled-input control rejected at `p = 0.0005`, and did so at every block length tested.

The honest reading is the one the framework's own table gives for this combination: parameters corroborated, predictive advantage not demonstrated. That is a real and useful state and it licenses much less than people want it to. It licenses treating the survey's parameters as corroborated estimates and computing headroom from them. It does not license the claim that the storage law forecasts the manual's rule count better than a trend line does, because on this evidence it does not.

The reason the matched autoregression is hard to beat is structural rather than accidental, and this is the most useful methodological finding in the paper. A declining institution is an exponentially-approaching trend, and an AR(1) with an intercept represents an exponentially-approaching trend exactly. Any leaky-accumulator model of a slow stock is therefore competing against a baseline that already contains most of the signal. **Beating it requires that the input term carry information about deviations from trend, not about the trend itself** — which is a much narrower claim than the framework usually makes, and one that a 34-block series has little power to detect. The corpus's own measurement of the suite bears this out from the other side: the matched autoregressive baseline evaluated *alone* false-passes 35.8 per cent of the time [Curzi, 2026u], so a suite that reads it in isolation is uninformative in both directions.

15.10.2 10.2 What the passing tests were worth

Two of the three rules passed, and the paper's own results limit what those passes establish.

FR1 passed and is cheap. The unconstrained optimum landed at `a = 0.97439`, `b = 0.5859`, strictly interior. This is a genuine finding — the office's manual behaves like a leaky accumulator rather than a compounding or oscillating quantity — and it is the finding that would have been most decisively negative had the data disagreed.

FR3 passed, and the direction of the `eta` discrepancy is informative. The fit is 16.8 per cent *higher* than the attribution audit, which is the expected direction, and it is consistent with two things at once: the audit's deliberate strictness, which excludes unattributable units, and the standing clause that the surplus functional as deployed is a lower bound. A fitted conversion slightly above the audited one is not an anomaly to be explained away.

The shuffle passed at `p = 0.0005`, and its calibration is measured and poor. Against a nominal 0.05 across 500 null trials, the same control measures a false-rejection rate of 0.108 under iid permutation and 0.120 under circular shift, with only block permutation near nominal at 0.060 [Curzi, 2026u]. **The reported `p` is therefore optimistic by an unknown but non-trivial factor**, and the correct procedure — calibrating the control by simulation on the study's own design and preregistering the calibrated threshold — was not run here. That is a limitation of this paper's own headline positive result and it is carried in Section 12.

And the shuffle is blind to the failure that matters most. Candidate PD-2, a rolling index of the board's own disagreement score, stores nothing, has no institution behind it, and **passed all four baselines decisively** — 89.4 per cent against persistence, 89.2 per cent against the matched autoregression, 93.6 per cent against direct input, and `p = 0.0005` against the shuffle. Against the genuine proxy's marginal 2.8 per cent, PD-2 looks like the far better model of the office. It was caught by a fitted half-life of 1.97 weeks against a measured 224 weeks, by a fitted `b` forty times the audited `eta`, by a collapse to 0.1 during the freeze, and by a single-week jump 38.5 times the admissible maximum. **A stock cannot teleport.** The ceiling test is arithmetic,

requires no statistics, takes one line, and is the single most reliable detector of the order-proxy trap available.

15.10.3 10.3 Relation to the surrounding literature

The protocol is an application of three established methodological positions and one identity.

Preregistration and held-out testing as corruption controls. The analytic-flexibility literature establishes that the damage is done at design time and is not repaired by subsequent computation [Simmons et al., 2011, Gelman and Loken, 2013, Nosek et al., 2018], and the replication record establishes what an uncontrolled literature looks like afterwards [Ioannidis, 2005, Open Science Collaboration, 2015]. The specific administrative claim made here — that preregistration is the countermeasure to measure corruption because it fixes the definition before anyone has a reason to distort it — is a corollary of Goodhart's and Campbell's statements rather than a new one [Goodhart, 1975, Campbell, 1979]. And the point prediction with a constant, rather than a directional prediction, is the response Meehl [1967] demands of any social science that wants its tests to cost something.

The effective count. The participation-ratio identity places the chosen aggregator inside a recognised family [Rényi, 1961], and the effective-numbers literature supplies the idiom [Hill, 1973, Jost, 2006], with independent reinvention in political science as the effective number of parties [Laakso and Taagepera, 1979] and in industrial organisation as the numbers-equivalent reading of concentration [Adelman, 1969, Herfindahl, 1950]. Two qualifications are required and both are made in Section 12: `1 + (N−1)u` is not itself a Hill number, and effective-count estimators are biased on small rosters, which is precisely the regime of a five-discipline board [Chao et al., 2014].

Measurement. The five operational choices are a specification of operations in the sense the measurement-theory literature requires before a construct is quantified [Krantz et al., 1971, Cronbach and Meehl, 1955], and the instrument-independence test is the multitrait–multimethod requirement in miniature [Campbell and Fiske, 1959]. The unanswered objection is Michell [1997]'s: psychology routinely assumes quantitative structure without testing for it, and neither conjoint measurement [Luce and Tukey, 1964] nor a Rasch demonstration [Rasch, 1960] has been run for `u` or for `Q`.

Within the corpus, the survey is the instrument that the storage law's derivations require in order to say anything about an organisation [Curzi, 2026a], the headroom gauge is what it computes [Curzi, 2026b,h], and the same five-parameter procedure applied across five systems from one desk to a 22-household block reproduced each system's observed decline — predicted −0.250474 against observed −0.250 at the desk, −0.996371 against −1.000 at the library, with residuals under 1.4 per cent of the ceiling term [Curzi, 2026k]. The ratio-invariance test that licenses a ratio scale in the perceptual setting — two fourfold physical steps returning 2.5100 and 2.5000, agreeing to 0.40 per cent and matching the fitted `4^0.6569 = 2.4859` [Curzi, 2026n] — is the kind of demonstration `u` still owes. The protocol's structural analogues elsewhere are the product-not-average scoring of the five-criterion disclosure test, where five structures averaging 0.80 score a group product of zero because two contain a zero [Curzi, 2026]; the staged-inspection result, where moving from terminal to per-stage checking at seven stages with `q = 0.10` and `d = 0.90` raises final yield from 47.83 per cent to 93.21 per cent at 0.63 reworks per unit [Curzi, 2026]; the exact invariance of summed drive under cost transfer, verified at +0.643688229683 and +1.268940530229 with residual 0.000e+00 [Curzi, 2026]; and the frequency criterion `f_min = k e^(2c/eta)` with `k = 0.483425732988414` Hz and the trichotomy `f > = < f_min` matching `Q_inf > = < 0` [Curzi, 2026], which is the same sign test this survey performs in a different unit.

15.11 11 Verification

15.11.1 11.1 Dimensional audit

The estimation model must balance before its fit is worth discussing. `[Q]` is the dimension of the order proxy — whatever unit it counts — and `[T]` is time, with the block as its unit.

Symbol	Dimension	Unit	Status
`Y_k`	`[Q]`	live precedent rules	**measured** by the declared order observable
`U_k`	`[1]`	none	**defined** dimensionless: `u = 1 - (x.y)^2` on unit vectors
`F(U_k)`	`[1]`	nats	**derived** dimensionless: a logarithm of an effective count
`a`	`[1]`	per block	**fitted** dimensionless; `a = 1 - delta`
`b`	`[Q]` per nat per block	rules per nat per week	**fitted**; identified with `eta`, and *not* dimensionless
`C_hat_k`	`[Q]` per block	rules per week	**measured**, denominated in `Y`-units by construction
`eps_k`	`[Q]` per block	rules per week	**residual**
`h`	`[Q]` per block	rules per week	**derived**
`Q_inf`	`[Q]`	rules	**derived**: `[Q]`/block ÷ `[1]`/block returns `[Q]`
RMSE	`[Q]`	rules	**measured**
`p`	`[1]`	none	**derived**, a probability

Does it balance? `Y_{k+1} = a Y_k + b F(U_k) - C_hat_k + eps_k` balances if and only if every right-hand term carries `[Q]`. The first does because `a` is dimensionless. The third does by the denomination rule, which is exactly why that rule is a specification and not a preference: **a cost observable left in hours or currency makes the equation dimensionally malformed however tidy the arithmetic looks.** The second does only if `b` carries `[Q]` per nat per block, which is why the survey reports `eta` as *rules per nat* and not as a bare number. The fixed coefficient of one on `C_hat` is a dimensional statement as much as an identification one.

Where the framework has no operational unit, that is said here. `Q` is stored order and has no natural unit; it is counted in whatever the declared proxy counts, and two systems' `Q` values are not comparable unless they declare the same proxy. The underlying construct — comfort — has **no operational unit at all** and is treated as a dimensionless magnitude by stipulation. Section 12 records that the quantitative structure required for that stipulation has not been tested.

15.11.2 11.2 Postulates and the compounding chain

Five postulates. Each is atomic, each is checkable alone, and each carries its status. Where a postulate could be dropped and the conclusion still follow, it has been dropped.

P1 — The declared margin was 10 per cent held-out RMSE reduction against each of persistence, direct input and the matched autoregression, and `p <= 0.05` against the shuffle. *Status:* **DEFINITION** — the rule the analyst set, recorded in the frozen document before the holdout was examined, and reproduced verbatim in Section 2.2. It carries no empirical weight; it is a commitment.

P2 — Held-out RMSE, one step ahead, over the eight holdout blocks: storage model 1.852, persistence 3.102, direct input 79.397, matched AR(1) with intercept 1.906. *Status:* **ARITHMETIC** from the frozen coefficients and the observed holdout; the numbers are in the table of Section 9.4 and any reader with those four values can redo every step below.

P3 — Skill is `1 - RMSE_model / RMSE_baseline`. *Status:* **DEFINITION** — the convention declared in Section 6.3.

P4 — The matched autoregression has the same number of free parameters as the storage model: two. *Status:* **DEFINITION** — `a_0, a_1` against `a, b`, with `C_hat` fixed at coefficient one in the storage model and no intercept added. This is the postulate that makes the comparison fair, and it is the one a reviewer should check first.

P5 — Failure against any one of the four baselines is a failure of FR2; the four are not averaged. *Status:* **DEFINITION** — the aggregation rule, declared at the freeze.

15.11.3 11.3 The chain, ending in the negative result

```
P2 + P3   -> L1   skill vs persistence   = 1 - 1.852/3.102  = 0.40296...  ARITHMETIC
P2 + P3   -> L2   skill vs direct input  = 1 - 1.852/79.397 = 0.97668...  ARITHMETIC
P2 + P3   -> L3   skill vs matched AR(1) = 1 - 1.852/1.906  = 0.02833...  ARITHMETIC
L1 + P1   -> R1   40.3% >= 10%   -> persistence margin met                 ARITHMETIC
L2 + P1   -> R2   97.7% >= 10%   -> direct-input margin met                ARITHMETIC
L3 + P1   -> R3   2.8%  <  10%   -> matched-AR margin NOT met              ARITHMETIC
P1        -> R4   shuffle p = 0.0005 <= 0.05 -> shuffle margin met         ARITHMETIC
R3 + P5   -> R5   FR2 FAILS                                               ARITHMETIC
R5 + P4   -> R6   the failure is against the parameter-matched baseline,
                  so it is not explained by a parameter-count advantage   ARITHMETIC
```

Every line is checkable from its own inputs. **The three long divisions a reviewer can redo in a minute are `1.852/3.102 = 0.59704`, `1.852/79.397 = 0.02333`, and `1.852/1.906 = 0.97167`; the third subtracted from one is 0.02833, which is 2.8 per cent, which is less than 10 per cent.** That is the whole of the negative result. It follows from two definitions the author set in advance and one table of measured errors, and from nothing else.

What the chain does not establish. It does not establish that the storage law is false. It does not establish that the AR(1) is a better model of the office — a 2.8 per cent advantage in the model's favour is not evidence for the baseline either. What it establishes is that **the predictive advantage the preregistration required was not demonstrated**, which is a statement about the evidence and not about the world.

15.11.4 11.4 The participation-ratio identity, checked on the smallest case

The identity `N_eff = 1 + (N−1)U = N + 1 − N/PR` is a derivation, and the smallest case that exercises it is `K = 2`. Take two unit vectors with inner product `c`. Then `||G||_F^2 = 2 + 2c^2`, so `PR = 4/(2 + 2c^2) = 2/(1 + c^2)`, and

```
N + 1 − N/PR = 3 − 2(1 + c^2)/2 = 3 − (1 + c^2) = 2 − c^2
```

while `U = 1 − c^2` for the single pair, so `1 + (N−1)U = 1 + (1 − c^2) = 2 − c^2`. **The two agree identically, for every `c`.** *Status:* **ARITHMETIC.** The `K = 3, 5, 8, 12` checks reported in Section 4.3 are checks on the implementation rather than on the mathematics, and the 1e-12

agreement should be read as a floating-point tolerance and not as an empirical result. The independent re-verification over six random correlation matrices, returning 6.486606017883 by both routes at `N = 8` and `u_bar = 0.783801`, is the same kind of check [Curzi, 2026].

15.11.5 11.5 Model comparison, audited against what proper model comparison requires

RMSE against a baseline is not model comparison. Measured against the full requirement, this paper's fit stands as follows, and two of the five items are **not met**.

Requirement	Status in this paper
Out-of-sample rather than in-sample performance	**Met.** One-step-ahead on a calendar holdout of 8 blocks, examined once.
A matched null with the same free-parameter count	**Met.** AR(1) with intercept: 2 parameters against the storage model's 2.
The number of free parameters stated explicitly	**Met.** Storage model 2 (`a`, `b`), with `C_hat` at fixed coefficient 1 and no intercept; persistence 0; direct input 2; matched AR(1) 2.
An information criterion, so the parameter count is paid for	**NOT MET.** No AIC or BIC was computed, at any block length, for any model.
Effect sizes with intervals, not only p-values	**PARTIALLY MET.** A profile interval on `a` is specified by the procedure and the fitted `a` is reported; no interval accompanies the RMSE differences, and no effect size accompanies the shuffle `p`.

The two unmet items are stated plainly because they change how the headline number should be read. The parameter count was not paid for: with a matched two-parameter baseline the information criterion would differ from the RMSE comparison only through the likelihood scaling, so the omission is less damaging here than it would be against an unmatched baseline — but it was not computed, and the claim that a 2.8 per cent RMSE advantage is *any* advantage is therefore untested against a criterion that penalises complexity. And a `p` without an effect size is an incomplete statistical claim: the shuffle's `p = 0.0005` says the real fit sits outside the null, and the corresponding **effect size, which can be read off the table in Section 9.4, is that the real fit's RMSE of 1.852 lies 45.5 per cent below the null mean of 3.400 and 30.3 per cent below the null fifth percentile of 2.657.** Those two figures should have been reported beside the `p` and were not; they are reported here.

15.12 12 Limitations

The worked system is a constructed composite, not a field study. The organisation is invented and the audit series is a simulation with declared generating parameters, produced so that every number in the paper can be checked by the reader. **It must not be cited as evidence about any real review office.** Its function is to show the procedure operating end to end on numbers that behave the way real numbers behave — including where they refuse to cooperate. Every statement in Section 9 is therefore a statement about the protocol's behaviour and not about municipal government.

The preregistered confirmatory fit failed FR2. It is recorded here as well as in Sections 9.5 and 10.1 because a limitations section that omitted it would be dishonest. The storage law is not corroborated as a predictive model of that system at that time scale.

The reflection problem: the specification is not identified from observational data. Where a group's members both influence and constitute the aggregate — which is the case for a review board — the endogenous social effect cannot be separated from contextual effects or from correlated unobservables, however much data is collected [Manski, 1993], and influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011]. **This is a deeper limitation than the failed rule.** It means that even a fit passing all three falsification rules would not have identified `eta` as the conversion of interaction into structure, because a contextual effect or a common driver would produce the same fit. No remedy is adopted in this paper — not randomisation, not an instrument, not partial-population variation, and not a defended stipulation of the reference group [Brock and Durlauf, 2001], which is the most realistic route and is the one the fixed category set and declared participation rule of Sections 4.2 and 4.3 point toward without yet taking. **Until one of them is executed, the fitted `b` is a coefficient in a declared specification, not an identified social effect**, and Section 9.8's licence list should be read as narrower still than it is written.

The shuffled-input control is anti-conservative and was not calibrated here. Its measured false-rejection rate is 0.108 for iid permutation and 0.120 for circular shift against a nominal 0.05, with only block permutation near nominal at 0.060 [Curzi, 2026u]. **The `p = 0.0005` reported in Section 9.4 is therefore optimistic by an unknown factor**, and the required repair — calibrating the control by simulation on this study's own design and preregistering the calibrated threshold — was not performed. The positive result on H2 is correspondingly weaker than its nominal `p` suggests.

No information criterion was computed. Section 11.5 gives the audit. A model with more parameters that fits better has not been shown to be better, and although the baseline here is parameter-matched, the omission means the comparison rests on RMSE alone.

The functional form of the surplus term is underdetermined by the evidence. At least two published alternatives fit the same qualitative facts: social impact rising as a sub-linear power of the number of sources [Latané, 1981], and value growing as `n log n` [Briscoe et al., 2006]. Nothing in this paper discriminates between them and `ln(1 + (N−1)u)`. **The logarithmic form is chosen for its axiomatic characterisation, not established by measurement**, and the model-comparison exercise that would settle it — fitting the alternatives on the same holdout and paying for their parameters — has not been run.

`N_eff = 1 + (N−1)u` is not a Hill number. `1/HHI` and `exp(H)` are Hill numbers of order two and one and are fully grounded [Hill, 1973, Jost, 2006]; `1 + (N−1)u` is a linear interpolation between 1 and `N`, is this framework's own construction, and has no counterpart in the ecological, economic or information-theoretic literature. The participation-ratio identity of Section 4.3 establishes that the chosen aggregator coincides with the Rényi-2 effective count *for this encoding* [Rényi, 1961], which is a narrower claim and the only one made.

Effective-count estimators are biased on small rosters, and five disciplines is a small roster. Estimation bias in diversity measures from finite samples is well documented [Chao et al., 2014], and no bias correction was applied to `U` or to `N_eff` here. With a 12-case probe set and five categories the pairwise estimates rest on ten pairs of twelve binary judgements; the reported `U = 0.6972` carries no interval and should not be treated as though it did.

Scale type is stipulated, not demonstrated. `u` and `Y` are treated as ratio-scale quantities throughout — differences, ratios, logarithms and means are all taken — and the conjoint-measurement conditions that would license that treatment have not been tested [Luce and Tukey, 1964, Krantz et al., 1971]. Michell [1997]'s objection, that psychology assumes quantitative structure without testing for it, applies directly and is not answered here. The ratio-invariance demonstration that the corpus runs for the perceptual scale [Curzi, 2026n] is the kind of test `u` still owes.

The parameters were chosen rather than measured in one important place. `delta` comes from a single 33-week episode, so it rests on two endpoint counts and an exponential assumption; the liveness threshold of 156 weeks was declared in advance but is a choice, and moving it by a year in either direction changes `Q` by about 5 rules and `h` by about 0.015; and the attention-tax line of the cost tally, at 0.13 `Y`-units per week, is the only line resting on a declared

conversion rather than a direct count. Section 9.3 reports that none of the three named sensitivities reverses the sign of `h` alone and that two of them together would.

Two claims the paper does not make. It makes **no scaling or heavy-tail claim**, so the distribution-fitting procedure such a claim would require [Clauset et al., 2009] does not apply. And it makes **no causal claim**: the shuffle establishes that the timing of interaction carries information about subsequent order, not that intervening on interaction produces order.

The Lucas critique applies to this framework's own parameters. Section 6.9 asserts that a measure changes the system measured by it, and the model then treats `delta`, `eta` and `N` as constant over the estimation window. The position taken here is the narrow one: constancy is assumed **only over a declared window and only while `Y` is not itself reported upward**, which is why the first and strongest corruption defence is to keep the order proxy out of the reporting line. Where `Y` is reported upward, the framework's own reflexivity applies to it and the constancy postulate fails. That is a restriction on where the protocol may be used, and it is stated as one.

What would falsify the protocol itself, as opposed to the model it tests. If a system that passed all three rules subsequently failed a randomised intervention predicted from its fitted parameters, the protocol would have licensed a claim it should not have. If the four repackaging tests passed on a proxy later shown to be a filtered input, the structural audit would be insufficient. And if the calibrated shuffle threshold, once computed, moved the reported `p` above 0.05, H2 would be undecided rather than supported.

15.13 13 Conclusion

This paper specifies an end-to-end measurement protocol for the storage law — five operational choices, an estimator for each of six parameters, a one-page output format, and four structural tests against the order-proxy trap — and adds a participation-ratio identity, `N_eff = N + 1 − N/PR`, equivalent to `K/PR = 1 + (K−1)(1−U)`, that makes a redundancy audit computable from a Gram matrix and agrees with direct computation at `K = 3, 5, 8, 12` to better than 1e-12. Applied to a constructed county development-review office, the survey returned `N = 5`, capacity `U = 0.6972` against a realised `U = 0.6019` — a realisation deficit of 0.1823 nats caused by 32.7 per cent of sessions running short-handed — `delta = 0.003090` per week from a 33-week zero-input episode with a half-life of 224 weeks, `eta = 0.5015` rules per nat, `Q = 113` after liveness and applier corrections, headroom `h = −0.1720`, and a steady state `Q_inf = −17.3`, reported as negative.

The preregistered confirmatory fit at the declared eight-week block scale returned `a = 0.97439` and `b = 0.5859`, both strictly interior to the admissible box, with implied decay within 4.8 per cent and implied conversion within 16.8 per cent of independently audited values. It beat a matched autoregression by 2.8 per cent on held-out error against a declared requirement of 10 per cent. FR2 therefore fails, and the verdict is: parameters corroborated, predictive advantage not demonstrated. An alternative block scale that would have passed everything was reported as exploratory and licenses nothing. A repackaged order proxy that stores nothing passed all four baselines decisively and was caught only by structural tests — a fitted half-life of 1.97 weeks against a measured 224, and a single-week jump 38.5 times the maximum admissible increment `eta ln N = 0.8071`.

Two results limit the weight of everything that passed, and both are reported here rather than left for a reviewer. The shuffled-input control on which the empirical audit depends is anti-conservative, measuring 0.108 against a nominal 0.05, so the reported `p = 0.0005` is optimistic. And the specification is reflexive, so the endogenous social effect is not identified from observational data at all; no fit, however honestly held out, recovers it. **A framework that cannot fail cannot be trusted. This one named exactly how it fails, ran the test, and failed it.**

15.13.1 Declarations

Data and code availability

No external dataset was collected for this paper. The worked system is a constructed composite: the organisation is invented and its audit series is a simulation with declared generating parameters, produced so that every reported number can be re-derived by the reader from the parameters and equations printed in the text. **It must not be cited as evidence about any real organisation.** The figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi designed the protocol, specified the falsification rules, performed the analysis, produced the figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The surveyed organisation is constructed. The declination requirement discussed in Section 6.10 — that no intervention removing a person's ability to decline is within scope — is a constraint on prospective work and is not a report of any study performed.

Preregistration

The falsification rules FR1, FR2 and FR3 and their numerical tolerances were declared in the frozen preregistration document reproduced as Appendix B, before the holdout data were examined: a 10 per cent held-out RMSE reduction against each of persistence, direct input and a matched autoregression, `p <= 0.05` against the shuffled-input control, and ±25 per cent agreement between fitted and independently measured `delta` and `eta`. Those rules are restated in Section 2.2 of this paper in advance of any result.

The result did not meet them. FR1 passed and FR3 passed; FR2 failed, against the matched autoregressive baseline, at a margin of 2.8 per cent against the required 10 per cent. No tolerance was moved after the holdout was examined. An alternative block scale, `B = 13`, would have passed every rule; it was not the preregistered scale, it rests on four holdout blocks rather than the eight the specification requires, and it is reported as exploratory and licenses nothing. The verdict recorded in the frozen form's own prescribed words is *parameters corroborated, predictive advantage not demonstrated.*

15.14 Appendix A — The survey worksheet

Reproduced as frozen. Seven forms, completed in order.

```
============================================================
FORM A — SCOPE AND THE FIVE CHOICES                  Date: ____
============================================================
A1  System name ..........................................
A2  Boundary: what is IN .................................
A3  Boundary: what is OUT, and why .......................
A4  Surveyor(s) ..........................................

CHOICE 1 — STATE ENCODER
A5  Encoder type  [ ] A disposition  [ ] B error  [ ] C practice
A6  Probe set: m = ____ cases. Listed at .................
A7  Response format ......................................
A8  Administration: independent? [ ] Y [ ] N   written? [ ] Y [ ] N
A9  Normalisation ........................................

CHOICE 2 — INTERACTION OBSERVABLE
A10 Aggregator ...........................................
A11 BADGE: aggregator is a declared modelling choice  [ ] initialled
A12 Participation rule (what counts as taking part) ........
A13 Capacity instrument ..................................
A14 Realised instrument ..................................

CHOICE 3 — ORDER OBSERVABLE
A15 Y = ..................................................
A16 Units ................................................
A17 Register / medium ....................................
A18 Recorder (name and role) .............................
A19 Audit schedule .......................................

CHOICE 4 — TIME SCALE
A20 One period = .........................................
A21 Calendar alignment ....................................
A22 Conversion lag measured at ____ periods
A23 Rationale for period length ...........................

CHOICE 5 — COST OBSERVABLE
A24 Route 1 categories ....................................
A25 Route 2 categories ....................................
A26 Declared conversion rate ____ Y-units per ____
A27 Justification for the rate ............................

============================================================
FORM B — THE REPACKAGING AUDIT  (the order-proxy condition)
============================================================
B1  Y source fields:  ....................................
B2  U source fields:  ....................................
B3  TEST 1 shared fields count = ____      PASS [ ]  FAIL [ ]
B4  Y recorder _________  U recorder _________ different? [ ]
B5  Y medium _________    U medium _________   different? [ ]
B6  Y schedule _________  U schedule _________ different? [ ]
B7  TEST 2                                   PASS [ ]  FAIL [ ]
B8  Zero-input episode found? [ ] Y [ ] N   periods ____ to ____
B9  Y at episode start ____  at end ____  after resumption ____
B10 TEST 3: does Y persist and decay smoothly? PASS [ ] FAIL [ ]
B11 Largest single-period rise in Y = ____
```

```
B12 Ceiling eta*lnN = ____   ratio = ____
B13 TEST 4                                  PASS [ ]  FAIL [ ]
B14 IF ANY TEST FAILS: STOP. Redefine Y. Do not proceed.

================================================================
FORM C — CATEGORIES AND GEOMETRY
================================================================
C1  Candidate categories enumerated:  ......................
C2  Merges (pair, u_ij, reason):  ..........................
C3  Drops (category, reason):  ............................
C4  N = ____        ln N = ____
C5  PAIRWISE TABLE
    pair | disagree k/m | d | cos = 1-2d | u = 1-cos^2
    -----+--------------+---+------------+------------
         |              |   |            |
C6  U capacity = mean pairwise u = ____
C7  N_eff capacity = 1 + (N-1)U = ____   F = ln N_eff = ____
C8  U realised (from period record) = ____
C9  F realised = ____
C10 Realisation deficit F_cap - F_real = ____

================================================================
FORM D — DELTA
================================================================
D1  Method [ ] D1 zero-input [ ] D2 cohort [ ] D3 ledger
D2  Episode / cohort identified ............................
D3  n periods = ____   Y_start = ____   Y_end = ____
D4  ratio = ____   delta = 1 - ratio^(1/n) = ____
D5  half-life = ln2 / (-ln(1-delta)) = ____ periods = ____ calendar
D6  tau = 1/delta = ____ periods    3*tau = ____ periods
D7  If D2: is log-survival linear? [ ] Y [ ] N   note ........

================================================================
FORM E — ETA
================================================================
E1  Window: periods ____ to ____   (w = ____)
E2  Y-units created in window = ____
E3  Of these, ATTRIBUTABLE with documented origin = ____
E4  Unattributable (excluded) = ____
E5  sum of F(U_t) over window = ____ nats
E6  eta = E3 / E5 = ____ Y-units per nat
E7  Capture decomposition (optional):
    interactions producing a written decision ....... ____
    written decisions reaching the durable medium ... ____
    Y-units per arriving decision ................... ____

================================================================
FORM F — COST
================================================================
line | category | natural measure | route | rate | Y-units/period
-----+----------+-----------------+-------+------+---------------
  1  |          |                 |       |      |
  2  |          |                 |       |      |
  3  |          |                 |       |      |
  4  |          |                 |       |      |
```

```
  5  |          |                 |       |      |
F1  Attention-tax line present? [ ] Y  (required)
F2  C_hat TOTAL = _____ Y-units/period
F3  Series-implied cost = _____   residual = _____
F4  Residual within declared tolerance _____ ? [ ] Y [ ] N
F5  Turnover assigned to [ ] delta [ ] C   rule: ............
F6  Slack booked as cost? [ ] N  (must be N)

================================================================
FORM G — Q, HEADROOM, VERDICT
================================================================
G1  Y register count at survey date = _____
G2  Liveness correction (inert units removed) = -_____
G3  Applier correction (orphaned units removed) = -_____
G4  Q = _____
G5  eta * ln N = _____
G6  delta * Q = _____
G7  C_hat = _____
G8  h = G5 - G6 - G7 = _____
G9  VERDICT  [ ] h < 0 unsustainable   [ ] h = 0 hold
             [ ] h > 0 growth available, r_max = _____
G10 realised one-step = eta*F_real - delta*Q - C_hat = _____
G11 realisation deficit = G8 - G10 = _____
G12 Q_inf = (eta*F_real - C_hat)/delta = _____
G13 Q_inf negative? [ ] Y -> quote Theorem 3.5 failure clause
G14 Three measurements that would most change this conclusion:
    1 ..................................................
    2 ..................................................
    3 ..................................................
G15 Surveyor signature _______________  Date __________
```

15.15 Appendix B — The preregistration form

Reproduced as frozen. Signed at the bottom; once signed, read-only.

```
================================================================
PREREGISTRATION                              Freeze date: ______
System: ........................  Analyst: ..............
Independent holder of this copy: ..........................
================================================================

---- PART 1  CARRIED FORWARD FROM THE SURVEY --------------------
1.1  Survey document reference and date ......................
1.2  Y = ....................................................
1.3  Y units ..................................................
1.4  U aggregator ............................................
1.5  N = _____   ln N = _____
1.6  delta (measured, method _____) = _____ per period
1.7  eta (attribution audit) = _____ Y-units per nat
1.8  C_hat = _____ Y-units per period
1.9  Q at survey = _____
1.10 Repackaging tests 1-4 all PASS?  [ ] Y  [ ] N (if N, stop)

---- PART 2  ESTIMATION DESIGN --------------------------------
```

```
2.1  Series: periods ____ to ____  (T = ____)
2.2  Block length B = ____   Alignment: ......................
2.3  JUSTIFICATION for B (must not reference results):
     conversion lag = ____ ; blocks available = ____ ;
     holdout blocks = ____ ; delta per block = ____
     ....................................................
2.4  Train: blocks ____ to ____      (n = ____)
2.5  Holdout: blocks ____ to ____    (n = ____)
2.6  Model:  Y_{k+1} = a Y_k + b F(U_k) - C_hat_k + eps_k
     constraints 0 <= a <= 1, b >= 0 ; C coefficient fixed at 1
2.7  Deviations from the minimal model (if any) and reasons:
     ....................................................
2.8  Missing-period rule ....................................
2.9  Outlier rule (must not reference residuals) ..............
2.10 Known structural breaks and handling ...................

---- PART 3  BASELINES ----------------------------------------
3.1  Persistence:      Y_hat = Y_k
3.2  Direct input:     Y_hat = c0 + c1 F(U_k), fitted on train
3.3  Matched AR(1):    Y_hat = a0 + a1 Y_k, fitted on train
3.4  Shuffled input:   refit under permutation, n_perm = ____
3.5  Metric: one-step-ahead held-out RMSE in Y-units

---- PART 4  FALSIFICATION RULES AND TOLERANCES ---------------
4.1  FR1  unconstrained optimum inside 0<=a<=1, b>=0
4.2  FR2  margin vs persistence ..... ____ % RMSE
          margin vs direct input ..... ____ % RMSE
          margin vs matched AR ..... ____ % RMSE
          shuffle p threshold ..... ____
4.3  FR3  tolerance on delta ....... +/- ____ %
          tolerance on eta ......... +/- ____ %
          conversion-lag allowance for the ceiling check ____
4.4  Any single FR failure = the framework is not corroborated
     for this system at this time scale.       [ ] initialled

---- PART 5  INTERVENTION (complete only if one is planned) ---
5.1  Change to be made ....................................
5.2  Parameter(s) targeted  [ ]eta [ ]C [ ]u [ ]N [ ]shed Q
5.3  PREDICTION: Y will rise by at least r = ____ per period,
     over k = ____ periods, starting m = ____ periods after
     the change (m >= conversion lag).
5.4  Headroom check: r <= h = ____             [ ] confirmed
5.5  FALSIFICATION: if mean one-period change over periods
     m..m+k is below ____, the intervention is ineffective
     and will be withdrawn.
5.6  STOPPING RULE, futility ..............................
5.7  STOPPING RULE, harm (must include a non-Y measure) .......
     ....................................................
5.8  CONSENT: who is affected ..............................
     how informed ..........................................
     how an individual may decline WITHOUT PENALTY .............
     [ ] a declination mechanism exists  (required; if not, stop)
5.9  REVERSIBILITY: procedure ..............................
     authorised to execute .................................
     time to reverse .......................................
```

```
     what is NOT recoverable ................................
5.10 PRE-MORTEM (why this failed, written in advance):
     ........................................................
     ........................................................

---- PART 6  SIGNATURES ------------------------------------
Analyst ......................  Date ..........
Independent holder ...........  Date ..........
Sponsor / system owner .......  Date ..........
================================================================
```

Chapter 16

A Gated Intervention Procedure and a Thirty-Two-Entry Failure Catalog for the Storage Law

Abstract. This paper converts the headroom gauge $h_t = \eta \ln N - \delta Q_t - C_t$ into a decision procedure and pairs it with a closed catalog of failure modes. Seven gates run from evidence admissibility through regime, feasibility, lever ranking, sequencing, reversibility, consent, and a capped preregistered pilot. The binding rules are that a requested increment must satisfy $r_t \le h_t$ with a recommended reserve $r_t \le 0.7\,h_t$, since the maximum increment $r_{\max} = h_t$ is attained only at $U_t = 1$, and that the $h_t < 0$ branch forbids building at any interaction geometry. In the worked case a request for $r = 0.55$ rules per week was refused at $h = -0.1720$; raising conversion from 0.5015 to 1.05 supplied $+0.8828$, or 80.8 per cent of the recovery, and cutting cost from 0.63 to 0.42 the remaining $+0.2100$, bringing headroom to $+0.9207$, after which the request fitted inside the budget at a reserve ratio of 0.597. The shedding lever was computed and declined: at a weekly decay of $\delta = 0.003090$ a fourteen-unit shed returns $\delta q = 0.0433$, sixteen times less than the same shed at $\delta = 0.05$. Feasibility is tested on the mean of F and never on F of the mean U: cutting the low-participation share from 32.7 to 12 per cent raises the mean of F from 1.1498 to 1.3241 nats and the realised step from $+0.4381$ to $+0.6211$, while the concavity penalty falls from 0.0761 to 0.0354 nats. The catalog's thirty-two entries each name the damaged parameter, its measurement and its price. A ward-handover case carries the central diagnostic distinction: headroom was $+1.686$ per week while the ward lost 0.755 per week, because three of five declared categories did not attend — and that intervention's own predicted fixed point of 71.46 exceeds a 34-bed census and is declared outside the model's range rather than reported as a forecast.

Keywords: headroom gauge; lever ranking; reversibility; capped pilot; failure catalog; frozen denominator; realisation deficit; sequencing

Series: B. Engineering the Surplus **Paper:** CCP-016

16.1 1 Introduction

A measurement framework earns nothing until it can refuse. The preceding papers in this series establish a storage law for institutional capability, a gauge computed from its parameters, and a preregistered protocol for putting the law at risk against data. What none of them

supplies is the step an analyst is actually asked for: given the six numbers, what may be built, how much of it, in what order, and under what conditions must the answer be *nothing*.

This paper supplies that step in two parts. The first is a **gated decision procedure** — seven gates running from the admissibility of the evidence, through the regime of the gauge, feasibility, lever ranking, sequencing, reversibility, consent, and a capped preregistered pilot. The procedure is binding in a specific arithmetic sense: it converts the headroom gauge from a diagnostic number into a constraint on what may be promised, and its central rule is a single inequality that a sponsor can check in one subtraction. The second is a **closed catalog of thirty-two named failure modes**, each priced. The catalog is closed because the law it serves has six degrees of freedom, so every failure of a system the law describes is a claim about one of six quantities; the thirty-two names are recurring routes into six parameter faults, distinguished because the routes have different signatures, instruments and fixes.

The gauge is inherited. The storage law `Q_{t+1} = (1−delta) Q_t + eta S_t − C_t` with `S_t = F(U_t) = ln(1 + (N−1)U_t)` is established as a first-order discrete filter, with its steady state, its pole and its time constant, in Curzi [2026a], whose principal result is that the map is Lipschitz with modulus `1 − delta`, so any positive loss makes it a contraction with steady state `Q* = (eta S − C)/delta`, and `delta = 0` is a pathology rather than an ideal. The headroom gauge itself — `h_t = eta ln N − delta Q_t − C_t`, its three regimes, its maximum increment `r_max = h_t`, and the separate admissibility and sustainability constraints — is established in Curzi [2026b]. The surplus functional, its axiomatic characterisation, and the affine effective count `N_eff = 1 + (N−1)u` are established in Curzi [2026]. The survey that produces the six numbers and the preregistration that puts them at risk are established in Curzi [2026i], whose confirmatory application **failed its own second falsification rule**: a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin, despite +40.3 per cent over persistence, +97.7 per cent over direct input, and `p = 0.0005` against the shuffled-input control. That failure is not a footnote to this paper; it is the evidential status of every applied statement in it, and it determines what the worked cases below are permitted to claim.

Four things are contributed here that the surrounding literature does not supply and the preceding papers do not state.

First, the increment rule with a reserve. `r_t <= h_t` follows from the gauge, but `r_max = h_t` is attained only at `U_t = 1` — perfect mutual orthogonality of every category — so taking `r = h` is promising the extremal case as the plan. The procedure therefore carries a recommended reserve `r <= 0.7 h`, stated as a reserve rather than absorbed, and it carries the `h < 0` branch in which **no interaction geometry can prevent order from decreasing**, so that the refusal is a computed result rather than a judgement.

Second, a priced separation between creating headroom and realising it. `h` is defined on the ceiling `ln N`; the realised one-step change is defined on `F(U)`. Raising the non-parallelism `u` therefore does not raise `h` at all — it closes the gap. The difference, `h − step = eta (ln N − F(U))`, is named here as the *realisation deficit*, and it is typically the cheapest quantity in the building.

Third, the concavity correction at the feasibility gate. What accumulates over many periods is the sum of `eta F(U_t)`, so the operative quantity is the **mean of `F`** and not `F` of the mean `U`. Because `F` is strictly concave, Jensen's inequality makes the second strictly larger whenever the geometry varies. In the worked office the gap is 0.0761 nats — 6.2 per cent of the surplus, lost to variability alone — and the practical consequence is that regularising a variable geometry pays twice, once by raising the mean and once by shrinking the penalty the variability was charging. The distinction is not pedantic: a feasibility test run against `F` of the mean passes systems that will not deliver.

Fourth, the counter-result, which is the most useful thing in the paper. Positive headroom is not health. The gauge uses the ceiling; the one-step condition uses what actually happened. A system can have large positive headroom and shrink every period because it operates far below its own ceiling. The hospital ward worked in Section 8.2 has `h = +1.69` carried-patients per week and is losing 0.76 per week, and a diagnosis that stopped at "headroom is positive, we are fine" would have been exactly wrong. The catalog names the misdiagnosis, and the same

worked case carries a second disclosure of the same kind: the intervention's own predicted fixed point of 71.46 exceeds the ward's 34-bed census, so the linear law is outside its range there and is reported as a saturation date rather than as an equilibrium.

Two things this paper does not do, stated at the outset because they bound everything that follows. **No dataset was collected.** Both worked systems are constructed — one a declared composite with simulated audit series, one a constructed clinical case — and their arithmetic is real and reproducible while their organisations are not. And **the reflexive specification underlying the law is not identified from observational data** [Manski, 1993]. Section 3.4 states that where it bites hardest, which is Gate 0, since Gate 0 is exactly where an identification claim would have to be made and is not.

Section 2 states the hypotheses and what would falsify each. Section 3 gives the framework, the inherited results, the identification problem, the scale-type audit, and the badged postulates. Section 4 develops the seven gates. Section 5 gives the twelve faults of the procedure itself. Sections 6 and 7 give the catalog's warrant and its thirty-two entries, followed by the seven faults of diagnosis. Section 8 works two systems end to end. Sections 9 to 12 discuss what follows, verify the chain, state the limitations, and conclude. Appendices A to C carry the specification schedules and the two working forms.

16.2 2 Hypotheses

The procedure of Section 4 is a **decision rule**, and a decision rule is not falsifiable by observation: it can be shown to be badly designed, expensive, or inconsistent with its own premises, but it makes no claim about the world. The same is true of the consent gate (Section 4.12), which is a **constraint on the space of admissible designs** and not a prediction. The arithmetic results of Section 3.2 are derivations from an equation and are labelled as such where they appear.

What *is* falsifiable is the claim that the gauge and the catalog describe real accumulating systems. The hypotheses below separate the two. **None of them was tested against collected data in this paper**; both worked systems in Section 8 are constructed, and Section 11 records that as the paper's principal limitation. Each hypothesis names the observation that would defeat it.

> **H1 — The headroom bound binds.** A requested per-period increment exceeding the measured headroom cannot be delivered by any arrangement of the interaction, because `r_max = h_t` is the extremal value over all admissible geometries. **Predicted:** a system committing to `r > h` under-delivers by at least `r − h` per period once the conversion lag has elapsed, and its shortfall is insensitive to facilitation, seating, ordering or effort. **Falsified by:** a system delivering a sustained per-period increment above its measured `h` over several time constants with `eta`, `delta`, `N` and `C` stable and independently re-measured.
>
> **H2 — Positive headroom does not imply growth.** The gauge is computed on the ceiling `ln N` while the realised step is computed on `F(U)`, so the sign of `h` and the sign of the one-step change are independent quantities. **Predicted:** systems with `h > 0` and a low realised fraction `S / ln N` decline, at a rate `eta(ln N − F(U))` below the headroom; the worked ward has `h = +1.686` and a step of `−0.756` per week at `S / ln N = 46.04` per cent. **Falsified by:** systems with positive headroom reliably growing irrespective of their realised fraction, or a measured realised fraction that fails to predict the sign of the one-step change when the other parameters are held.
>
> **H3 — The shedding lever is worth `delta x q` and nothing more.** Removing `q` units of stored order returns exactly `delta x q` of headroom per period, so shedding is decisive in fast-decaying systems and nearly worthless in durable ones. **Predicted:** at `delta = 0.003090` per week a fourteen-unit shed returns 0.0433 per week; the

same shed at `delta = 0.05` returns 0.70, a factor of sixteen. **Falsified by:** a measured headroom gain from a shed that materially exceeds `delta x q` in the absence of a simultaneous change in `C` — which would show that the retained units were charging more than the decay term accounts for, and that the cost ledger is misspecified.

H4 — Conversion outranks category count, per unit of committed resource. At parameter values typical of surveyed systems, raising `eta` returns more headroom per unit of resource than raising `N`, because `eta` multiplies the ceiling while `N` enters it logarithmically against a coordination cost that grows with the number of pairs [Brooks, 1975, Conway, 1968]. **Predicted:** in the worked office, `eta` returns 2.207 rules per week per full-time equivalent against the sixth discipline's 0.158 on the source's own comparison — a factor the source states as fourteen and this paper reports as basis-dependent, between ten and eighteen, in Section 8.1. **Falsified by:** a system in which the measured cost per unit of `eta` gain exceeds the cost per unit of `N` gain. **The ranking is a default and is expected to invert** in systems whose capture discipline is already near its medium's absorption limit; the procedure requires the analyst to compute their own costs and to report an inversion when it occurs.

H5 — Regularising a variable geometry pays twice. Because `F` is strictly concave, a system whose geometry varies delivers strictly less than a system sitting at its own average every period; reducing the variability raises the mean *and* shrinks the concavity penalty. **Predicted:** cutting the low-participation share from 32.7 to 12 per cent raises the mean of `F` from 1.1498 to 1.3241 nats and the realised step from +0.4381 to +0.6211 rules per week, of which 0.0427 comes from the penalty falling from 0.0761 to 0.0354 nats. **Falsified by:** an observed step gain that tracks the change in `F` of the mean `U` (1.2262 to 1.3595) rather than the change in the mean of `F`. Note that `mean F <= F(mean U)` itself is a **derivation** from concavity, not a hypothesis; what is at risk here is the claim that the accumulating system is driven by the mean of `F`.

H6 — The frozen denominator is the dominant category fault. Declared category counts systematically exceed epistemically active counts, so systems plan against ceilings they do not own. **Predicted:** a ward declaring five handover categories with two attending plans against `ln 5 = 1.6094` rather than `ln 2 = 0.6931` — a ceiling 132 per cent higher than the one it owns; the general prediction is that a capability-partition audit reduces `N` in a majority of surveyed systems. **Falsified by:** surveys in which declared and capability-audited category counts agree within measurement error across a range of systems.

One further claim is stated and is explicitly not a hypothesis. That the catalog is *closed* — that there are exactly six kinds of failure — is **true by construction** given the model, since the law has six free quantities. It is a **definition**, it carries no empirical weight, and it would be defeated not by an observation but by the adoption of a different model. Section 6.1 says so at the point where the closure claim is made.

16.3 3 Framework

16.3.1 3.1 The gauge, and what it is a budget for

The gauge is

```
h_t := eta ln N − delta Q_t − C_t
```

and it is read as a budget in the following sense, which is a **shared arithmetic form and not an economic derivation**: `eta ln N` is the largest per-period income the interaction can deliver, attained only at the geometric ceiling; `delta Q_t` is the standing charge on what is already stored; `C_t` is dissipation. What remains is what may be committed.

```
h_t <  0    unsustainable in this model; no interaction geometry rescues it
h_t =  0    maintenance only
h_t >  0    growth available, and any requested increment must satisfy r_t <= h_t
```

`h` is not a score. It is a sign with a magnitude, and the sign changes what the procedure permits. The three regimes and the increment ceiling are established in Curzi [2026b]; what this paper adds is the procedure that consumes them and the reserve that keeps a plan off the extremal value.

`h` shrinks as `Q` grows. That is the structural reason institutions ossify: success accumulates stored order, `delta Q` grows in proportion, and the budget for anything new closes. It has an operational consequence at once — **the headroom measured last year is not the headroom now, and it is re-measured before every commitment.**

16.3.2 3.2 The six inherited results

Six results from the corpus's foundational and dynamical papers do the load-bearing work in Sections 4 to 8. They are relabelled **CC-T1** to **CC-T6** here so that the paper's own contributed results (R1 to R6 and CR, stated in Section 12) cannot be confused with them. Their original numbering in the source volumes is given for traceability; none of them is proved here.

Label	Statement	Source	Original
CC-T1	Maximum one-step increment. `r_max = eta ln N − delta Q_t − C_t = h_t`, and no policy based on interaction geometry alone can guarantee more. `r_max` is attained only at `U_t = 1`.	Curzi [2026b]	Corollary 4.2
CC-T2	Target increment feasibility. `Q_{t+1} − Q_t >= r_t` if and only if `S_t >= (delta Q_t + C_t + r_t)/eta`, which is attainable only if `delta Q_t + C_t + r_t <= eta ln N` — that is, only if `r_t <= h_t`.	Curzi [2026b]	Proposition 4.1

Label	Statement	Source	Original
CC-T3	Sustainability ceiling. A nondecreasing step requires `C_t <= eta ln N − delta Q_t`. If `C_t` exceeds it, **no choice of interaction geometry prevents order from decreasing.**	Curzi [2026b]	Corollary 3.4
CC-T4	Steady state. With constant `U = u∗`, `C = c∗` and `delta > 0`, `Q_inf = (eta F(u∗) − c∗)/delta`. If `eta F(u∗) < c∗` the equilibrium is negative, and the regime is unsustainable or the linear model must be replaced.	Curzi [2026a]	Theorem 3.5
CC-T5	Uniform bound and arrival. `Q_t <= max{Q_0, eta ln N / delta}`; and because `delta > 0` makes the update map a contraction, a unique fixed point exists and iteration reaches it from any start, at a rate set by `1/delta`.	Curzi [2026a], Curzi [2026]	Theorem 3.3; Banach
CC-T6	Range and concavity of the surplus functional. `0 <= F(u) <= ln N` on `u` in `[0,1]`, with `F(u) = ln(1 + (N−1)u)` strictly concave in `u`, and `N_eff = 1 + (N−1)u`.	Curzi [2026]	Derived results 3.2, 3.3

CC-T5's arrival clause is what makes an intervention arrive at all, and it is also why interventions arrive slowly. The a posteriori form of the same contraction argument — `d(x_n, x∗) <= (k^n/(1−k)) d(x_0, x_1)`, read operationally as *distance to arrival is the last correction scaled by* `q/(1−q)` — is established in Curzi [2026] and is the reason the monitoring rule of Section 4.16 can read the size of the most recent revision as a gauge of remaining distance.

On the effective count, one correction is mandatory and is made here rather than in a footnote. `N_eff = 1 + (N−1)u` is **this framework's own construction**, adopted by design axiom, and it is **not a Hill number**. `1/HHI` is the Hill number of order two and `exp(H)` is the Hill number of order one; both are grounded in the diversity literature [Hill, 1973, Jost, 2006], and neither is the quantity used here. A reviewer reaching for that literature as support for the affine effective count will find it does not apply. Every occurrence of an effective count in this paper is the affine one, and it is named as such.

16.3.3 3.3 Two kinds of lever, and the two master numbers

The distinction is easy to miss from the equations and it governs half of the procedure.

```
h    = eta ln N − delta Q − C        uses the CEILING, ln N
step = eta F(U) − delta Q − C        uses what ACTUALLY HAPPENED, F(U)
```

Raising `u` does not raise `h`. It closes the gap between `step` and `h`. The gauge already assumes the best possible geometry, so working on `u` collects headroom that already exists and is being left on the table, while working on `eta`, `C`, `Q` or `N` changes how much headroom exists at all. Define the

```
realisation deficit  =  h − step  =  eta ( ln N − F(U) )
```

which is the amount of headroom a system has and is not collecting. **The two master numbers of any diagnosis are therefore `h` and the realised fraction `S / ln N`, and a diagnosis reporting one without the other is half a diagnosis** — it will send the analyst to the wrong lever. This is the analytical content of the paper's counter-result, and Section 8.2 is the case that exhibits it.

The realisation lever has two sub-levers with very different costs. **Capacity `u`** — how non-parallel the categories are when all of them speak — is slow, changed by recruitment, rotation and the deliberate outsider. **Realised `u`** — how much of that capacity appears in the period record — is changed by attendance rules, by requiring a recorded disposition from every category, by protocols that elicit positions in writing before discussion, and by room geometry. The realised sub-lever is nearly always the cheaper and nearly always available, because most organisations run sessions below their own measured capacity.

16.3.4 3.4 Identification: the objection Gate 0 cannot answer

This belongs in the framework rather than in a caveat, because it bounds what the procedure can deliver before any survey is run, and because Gate 0 is precisely where an identification claim would have to be made.

The specification underlying the gauge is reflexive. The interaction observable `U` is a property of a group; the group's stored order shapes what its members do next; and the members are what the group is made of. That is a linear-in-means reflexive structure, and Manski [1993] proves that in such a structure **the endogenous social effect is not identified from observational data** — it cannot be separated from contextual effects or from correlated unobservables, however much data is collected. Shalizi and Thomas [2011] establish the same confounding for network data, where influence and selection cannot generically be told apart from observation.

The consequence for this paper is exact. Gate 0 admits evidence on three grounds: a complete survey, passed repackaging tests, and an evaluated preregistration [Curzi, 2026i]. **None of the three identifies `eta`.** A survey that passes every gate has established that the order proxy is not a repackaging of the interaction observable and that the fitted model predicts out of sample; it has not established that `eta` measures the conversion of interaction into structure rather than a contextual effect, a common driver, or selection into the group. **A fit that passed all three falsification rules would still not have identified that parameter**, and the corpus's one confirmatory application did not pass all three [Curzi, 2026i].

This paper adopts no remedy, so the honest label is a modelling position rather than an identified empirical result. The four known routes are randomisation, an instrument, partial-population variation, and stipulating the reference group and defending the stipulation; Brock and Durlauf [2001] is the template for the last and the most realistic here. The survey's fixed category set, declared merge and drop rules and explicit participation rule are the machinery a defended stipulation would need — but stipulating a reference group and

defending it are two acts, and only the first has been performed. **Every `eta` in every worked case below is a coefficient in a declared specification.**

One consequence deserves separate statement because it is where the procedure is strongest. The **capped, preregistered, reversible pilot of Gate 7 is the only element of the procedure that can supply causal evidence**, precisely because it introduces exogenous variation the estimation did not have. That is the reason it is a gate rather than an option, and it is why the accounting-frame restriction of Section 4.2 permits a pilot while forbidding the claim that the model predicts its outcome. Where the corpus needs an existence proof that reflexive loops occur and can be caught in the field, MacKenzie and Millo [2003] supplies one with dates and price data; it is not a substitute for identification here.

The Lucas critique applies to this framework's own parameters. The procedure asserts that a measure changes the behaviour of the system measured by it, and prescribes raising a captured metric's decay rate in response (Section 7.6, FM-28), while treating `delta`, `eta` and `N` as constants over an estimation window. **Both cannot be true without qualification** [Lucas, 1976]. The position taken here is the narrow one: the parameters are treated as stable *only over a declared window and only while the order proxy is not itself a reported target*, which is exactly why Section 4.16's first prescription is to keep the proxy out of the reporting line. Where the proxy is reported upward, the constancy assumption fails and the model is outside its domain.

16.3.5 3.5 Scale type, and what licenses the arithmetic

The scale type of every measured construct is stated here, with what licenses it, because the procedure computes differences and ratios and neither is licensed on an ordinal scale.

Construct	Instrument	Scale type claimed	What licenses it	Verdict
`Q` — stored order	a count of declared proxy units (live rules; patients with stated concerns)	**ratio**	counting, with a fixed zero and a defined unit	licensed *for the proxy*; the identification of the proxy with stored order is a postulate
`C` — dissipation	a residual in the same proxy units per period	**ratio**	derived from two counts on a common denominator	licensed, conditional on the residual being complete (FM-23)
`delta` — retention loss	decay under a zero-input window	**ratio** (a dimensionless fraction)	a ratio of two counts of the same proxy	licensed
`eta` — conversion	attributable change in `Q` per nat of surplus	**ratio**	a ratio of a count to a derived dimensionless quantity	licensed as arithmetic; **not identified** (Section 3.4)
`N` — category count	capability partition	**absolute count**	counting, after a declared merge rule	licensed; the merge threshold is a stipulation
`u` — non-parallelism	`u = 1 - rho^2` from independently elicited estimates	**stipulated ratio**	**nothing external**; the quantitative structure has not been tested	**NOT licensed by measurement.** See below.

Construct	Instrument	Scale type claimed	What licenses it	Verdict
`h`, `step`, `r`	arithmetic on the above	**ratio, per period**	inherited from the terms, all in proxy units per period	licensed where `u` is
`r/h` — reserve ratio	a quotient of two headroom quantities	**ratio**	same unit, common zero	licensed **only on the `h > 0` branch**; at `h <= 0` the quotient is meaningless and the procedure does not form it

`u` is where the framework is most exposed, and the exposure is conceded rather than argued away. Michell [1997] argues that psychology routinely fails to test whether its attributes are quantitative at all before measuring them, and that this is the discipline's characteristic methodological error. **This is the strongest single objection to the framework, and this paper does not answer it.** The formal remedy is conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971], which tests directly whether an attribute has additive quantitative structure, or a Rasch demonstration [Rasch, 1960]. **Neither has been run for `u`, here or anywhere in this corpus.** Until one is, the treatment of `u` as a ratio-scale quantity — which is what taking `1 - rho^2`, its mean, and a logarithm of an affine function of it all require — is a **stipulation**, and every quantity downstream of it inherits that status. Stevens [1946] supplies the taxonomy the table above uses; it does not license the entry it is used to classify.

Two operations in this paper are worth naming because they are exactly the ones an ordinal scale would forbid. The **headroom difference** `h - step` and the **reserve ratio** `r/h` are both computed. Both are licensed if and only if `u`, and through it `F(U)`, is genuinely a ratio-scale magnitude. **If `u` is merely ordinal — if all that is defensible is that one board was more nonparallel than another — then the realisation deficit is not a quantity, the reserve ratio is not a proportion, and the correct downgrade is to report the *sign* of `h - step` and the *ordering* of levers while withdrawing every decimal place.** That downgrade is stated here so a reviewer does not have to propose it.

16.3.6 3.6 Badged postulates

P-A — Parameter separability (POSTULATE). The procedure treats `eta`, `C`, `u` and `N` as levers that can be moved one at a time. In real systems they interact: adding a category raises `C`; raising `eta` by adding a scribe consumes capacity that was producing something else; cutting `C` by removing rework may cut `u` if the rework was where disagreement surfaced. The procedure handles this by requiring each lever's second-order cost to be entered explicitly in the specification (Appendix B), but **the separability is an approximation and it has no external support.**

P-B — Hyperrecursive symmetry (POSTULATE). The storage law is scale-free in *form*: only `(delta, eta, N, C)` change between a desk, a room, a team, a building, an institution and a city. **This is a claim about model form, not about physics and not about parameter values.** It licenses running the same procedure at a different scale; it does not license carrying a number across scales. The claim has been exercised across five systems from a single desk to a 22-household block, reproducing each system's observed decline — predicted `−0.250474` against observed `−0.250` at the desk and `−0.996371` against `−1.000` at the library, with residuals under 1.4 per cent of the ceiling term [Curzi, 2026k]. **That is five systems, one corpus, one author, and no independent replication**; the invariance has been exercised, not established.

P-C — A single scalar `Q` describes an institution (POSTULATE). That the law describes a real institution at all is a choice of model, not a theorem. The foundational papers prove things *about* the law; they prove nothing about any organisation. Where a fault sits outside the model — a legal shutdown, a fire, a death — the catalog will not name it, and the correct response is to say so rather than force a parameter.

P-D — Cost accrues each period whether or not surplus was delivered (POSTULATE). This is an assumption about how organisations charge overhead. It has no external support and it is consumed at exactly one place, the sustainability ceiling CC-T3.

P-E — Consent (DESIGN CONSTRAINT, not a claim). No intervention that removes a person's ability to decline is within the scope of this procedure. This is stated in the framework rather than in an ethics section because it is a constraint on the *space of admissible designs* and therefore on the arithmetic: a design requiring compliance is not a cheaper version of a design inviting participation; it is a design the procedure does not produce. The constructive alternative is to build the more comfortable position next door and let people arrive at it — a method that works by removing resistance has no coercive failure mode, and its worst outcome is that someone declines to move. **This is a normative commitment. It is not derived from the arithmetic and it is not falsifiable**, and it is badged accordingly. Hacking [1995] is the relevant caution about classification acting on the classified; it supports the concern, not the rule.

16.3.7 3.7 Standing honesty clauses

These constrain every number in the paper and are repeated here because the catalog is the part of the material most likely to be quoted out of context.

- As deployed, the surplus functional is a **lower bound** on the ensemble's information, not the information itself. Every price quoted in Section 7 is a floor. A computed `h` is conservative in that one respect, and **that is not a licence to exceed it.**
- The foundational surplus paper is pure mathematics and makes **no empirical or metaphysical claim** [Curzi, 2026]. Nothing in this paper's arithmetic is validated by a theorem; it is validated, if at all, by the survey and preregistration that precede it.
- **No thermodynamic negentropy law is proved anywhere in this corpus**, and higher static surplus does not imply higher canonical concentration. Any entropy invoked near `ln N` is **Shannon entropy**, a dimensionless information measure in nats, not thermodynamic entropy, which carries `[energy][temperature]^-1`.
- Subadditivity holds on `u` only. **The triangle inequality fails for geometric triples** [Curzi, 2026], so interventions may not be composed across three groups by adding their surpluses. The empirical counterpart is that psychological similarity is asymmetric and violates the triangle inequality [Tversky, 1977].
- **Where an entry in the catalog cites an external result, the result is real and the *transfer* to a given system is a modelling assumption**, badged at each occurrence. The three-drawer honesty classification that governs this practice — established mathematics, established psychology, and declared postulates, with every protocol stated to run on the first two and steer by the third — is set out in Curzi [2026].
- `0 <= F(u) <= ln N`. **`ln N` is the hard ceiling on surplus per interaction**, which is why `h` is defined with `ln N` and not with `F(U_t)`.

16.3.8 3.8 Analogy discipline

Four borrowings run through this paper and each is labelled at first use, following the three-way rule that a physical result may be used as a **derivation** (the premises are genuinely satisfied), as a **shared mathematical form** (the same equation governs both, with no claim that the physics applies), or as an **analogy** (suggestive, never evidence).

The storage law as a leaky integrator is a shared mathematical form. The law and the first-order discrete section obey the same difference equation, so every theorem about first-order systems applies [Oppenheim et al., 1996, Elaydi, 2005]. Nothing about charge or current is claimed of an institution.

Queueing is a shared mathematical form. The utilisation cliff (Section 7.5, FM-22) uses the standard result that expected wait scales as `1/(1−rho)` [Kingman, 1962, Hopp and Spearman, 2011] and Little's Law `L = lambda W` [Little, 1961]. The premises — a service system with arrivals and a queue — are plausibly satisfied for a review board, and the transfer is nonetheless badged as a modelling assumption because the arrival and service distributions are never measured in the worked cases.

Room acoustics is a derivation, and the only one in the paper. Sabine's relation `RT60 = 0.161 V / A` is a physical result about a real room, and a meeting room is a room [Sabine, 1922, Schroeder, 1965]. **What is *not* a derivation is the step from reverberation time to non-parallelism `u`**: that a longer tail excludes the quiet participant, and that the exclusion registers as a fall in `u`, is a modelling assumption. The tail arithmetic that connects them — the gap for a previous talker to fall 15 dB is `t_15 = RT60/4`, giving 488 ms before treatment and 98 ms after — is established in Curzi [2026e] and is itself acoustics; its consequence for participation is not.

"Budget", "mortgage", "income" and "overdrawn" are analogy only. They are used because they make the gauge legible to a sponsor and for no other reason. No result in this paper depends on any of them, and none of them transfers the authority of accounting to a psychological claim.

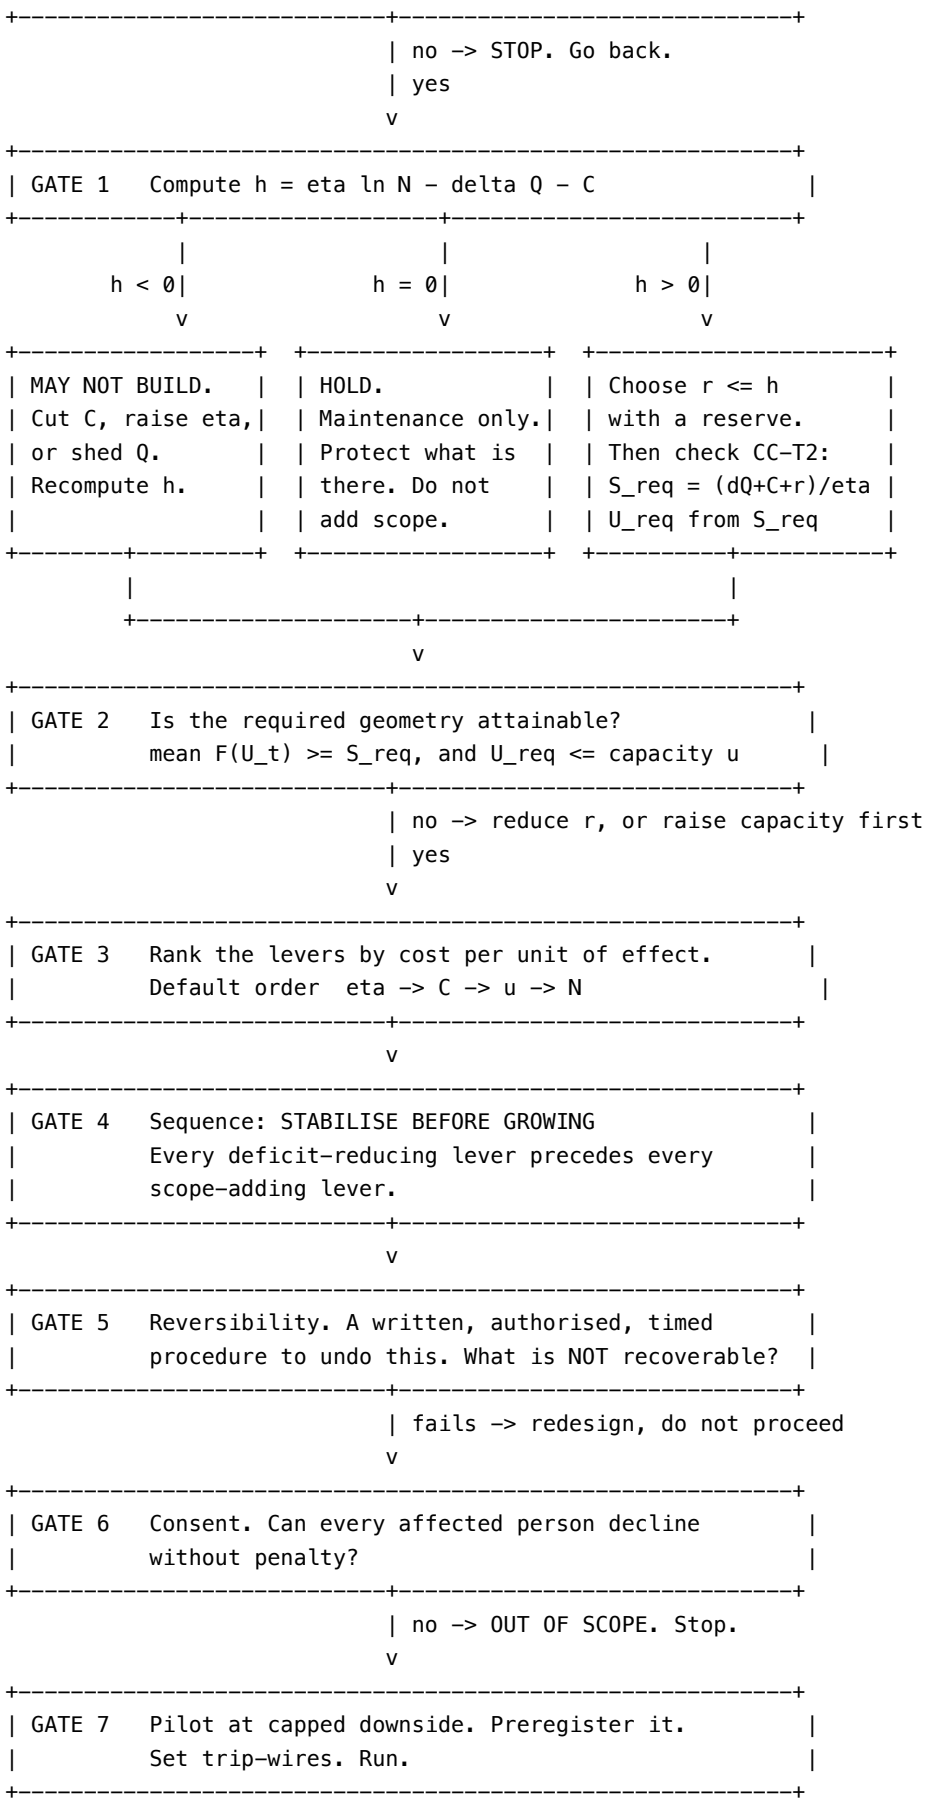
One borrowing is refused. The logarithm in `F(u)` is not justified here by Fechner's law [Fechner, 1860]; the perceptual story is not needed and Stevens [1957] disputes it. Where the logarithm needs a warrant, the stronger one is Peters and Gell-Mann [2016], which derives a logarithmic form from multiplicative dynamics with no psychological assumption at all. **That warrant is still not a discrimination**: Latané [1981]'s sub-linear power law and Briscoe et al. [2006]'s `n log n` fit the qualitative facts as well as `ln(1+(N−1)u)` does, and **nothing assembled in this paper discriminates among the three.** The logarithm is adopted for its axiomatic characterisation [Curzi, 2026] and not because it has been shown to be the right shape.

No scaling, power-law or heavy-tail claim is made or estimated in this paper. Where the catalog invokes preferential attachment and scale-free structure (Section 7.2, FM-10), the exponent is a property of the *cited model* [Barabási and Albert, 1999], not a quantity fitted here; no distribution was fitted, so the estimation procedure that any such claim would require [Clauset et al., 2009] does not apply, and the reader should note that the empirical generality of scale-free structure is itself contested [Broido and Clauset, 2019]. The entry is used for its *mechanism* — a hub redistributes rather than generates — and not for a tail exponent. # 4 The gated procedure

16.3.9 4.1 The gate sequence, and a correction to the source ordering

The procedure runs seven gates in the order below. **The source material this paper converts develops Gate 7 before Gate 5 in its expository body while running them in numerical order in its worked case. That is an inconsistency in the source, it is recorded here rather than silently repaired, and this paper adopts the numerical order 0 through 7 throughout** — so reversibility (Gate 5) is developed in Section 4.11, consent (Gate 6) in Section 4.12, and the pilot (Gate 7) in Section 4.13. The ordering matters: reversibility and consent are admissibility conditions on a design, and a pilot is a design, so a pilot cannot be specified before the conditions its design must satisfy.

```
+-------------------------------------------------------+
| GATE 0   Is the evidence admissible?                  |
|          Survey complete?  Repackaging tests all PASS?|
|          Preregistration frozen and evaluated?        |
```


| no -> STOP. Go back.
| yes
GATE 1 Compute h = eta ln N - delta Q - C
h < 0|
h = 0|
h > 0|
MAY NOT BUILD. Cut C, raise eta, or shed Q. Recompute h.
HOLD. Maintenance only. Protect what is there. Do not add scope.
Choose r <= h with a reserve. Then check CC-T2: S_req = (dQ+C+r)/eta U_req from S_req
GATE 2 Is the required geometry attainable? mean F(U_t) >= S_req, and U_req <= capacity u
| no -> reduce r, or raise capacity first
| yes
GATE 3 Rank the levers by cost per unit of effect. Default order eta -> C -> u -> N
GATE 4 Sequence: STABILISE BEFORE GROWING Every deficit-reducing lever precedes every scope-adding lever.
GATE 5 Reversibility. A written, authorised, timed procedure to undo this. What is NOT recoverable?
| fails -> redesign, do not proceed
GATE 6 Consent. Can every affected person decline without penalty?
| no -> OUT OF SCOPE. Stop.
GATE 7 Pilot at capped downside. Preregister it. Set trip-wires. Run.

16.3.10 4.2 Gate 0 — the admissibility of the evidence

The procedure may not be run on estimates. Three conditions are required.

- **A survey with a failed repackaging test produces no admissible `h`.** If the order observable is a function of the interaction observable, then `Q`, `delta` and `eta` are all measurements of the same artefact and `h` is a number about the instrument [Curzi, 2026i].
- **A failed first falsification rule means the storage law is the wrong model**, and the procedure does not apply. A system that compounds rather than accumulates, or that oscillates, is a real system with a different manual.
- A failed second rule with a passed third — the **accounting frame** — permits the procedure with one restriction developed in Section 4.13: `h` may be used as a budget constraint and a preregistered pilot may be run, and **no claim may be made that the model predicts the pilot's outcome.**

A passed shuffled-input control is not evidence of admissibility, and Gate 0 must not be read as though it were. The four-baseline suite ends with the shuffle, and the shuffle is the test a sponsor will quote back. Measured by simulation elsewhere in this corpus, it has three failure modes.

1. **It is anti-conservative as ordinarily run.** Against a nominal `alpha = 0.05`, measured false-rejection rates were **0.108 for iid permutation (2.2 times nominal), 0.120 for circular shift (2.4 times nominal), and 0.060 for block permutation.** A shuffle passed at the nominal threshold has roughly twice the advertised false-positive rate unless it was calibrated by simulation on the study's own design, which Gate 0 therefore requires.
2. **It cannot see repackaging.** A filtered repackaging passed all four baselines including the shuffle at `p = 0.0005` and was caught only by the structural ceiling test CC-T1, its single-step jump being **38.5 times** the theoretical maximum [Curzi, 2026i].
3. **It cannot see confounding by a shared slow driver.** In a 4,000-trial simulation the control **passed 98.35 per cent of the time** against a third variable driving both the interaction and the order observables, while behaving honestly in the clean case at 48.17 per cent detection.

The remedy is design controls, not more surrogates: blinding, channel separation, prospective sampling, and a yoked control. The second of the four is a structural property rather than a matter of care, and the corpus states it in another domain as the channel-separation principle: in-band architectures admit an impersonation class that no care *within* the shared channel can close, while out-of-band separation closes it structurally — measured at 0.383 per cent in-band ambiguity against 0.000 per cent out-of-band over 100,000 values [Curzi, 2026]. **A survey whose order observable and interaction observable are gathered through the same instrument is the in-band case**, and no amount of statistical care inside that instrument closes it. All four are properties of how the evidence was gathered, so Gate 0 is the last moment at which they can be checked and the first at which it is too late to add them. **A survey reporting a passed shuffle and none of the four design controls has not cleared Gate 0; it has cleared the weakest test available.** The wider point is the one the integrity literature makes about analytic flexibility and undeclared choices [Simmons et al., 2011, Gelman and Loken, 2013]: a control chosen after the data are seen is not a control. And Section 3.4 states the limit no control reaches — **none of these conditions identifies `eta`**, so Gate 0 governs admissibility of *evidence about the proxy* and never admissibility of a causal claim.

16.3.11 4.3 Gate 1 — the three regimes

Regime	Meaning	Permitted	Forbidden
`h < 0`	maintenance draw plus dissipation exceeds best-case income	cut `C`; raise `eta`; shed `Q`; raise `N`; recompute	any commitment to increase `Q`; any new scope; any promise with a date
`h = 0`	income exactly covers upkeep	maintenance; substitution at constant scope; measurement	growth of any kind, including growth described as a small addition
`h > 0`	budget remains	an increment `r <= h`, subject to Gates 2 to 7	`r > h`, under any argument

16.3.12 4.4 The `h < 0` branch: the procedure forbids building

This is the branch most analysts arrive in and the branch most briefs are written to avoid.

When `h < 0`, the procedure forbids building — not slowly, not partially, not as a pilot whose objective is growth, and not alongside other work. CC-T3 states that when `C > eta ln N − delta Q`, no arrangement of the interaction prevents decline. There is no seating chart, no facilitation technique, and no leadership change that alters an inequality between three measured quantities.

Three remedies are available before `N`.

Remedy 1 — cut `C`. Immediate, usually cheap, and acting on the term measured most carefully. The recommended order is rework loops first — they are pure dissipation and are hard to defend once counted — then contradiction maintenance, then the attention tax, then coordination overhead. Appeal and orphaning costs are usually structural and are cut by repairing the process that produces them.

Remedy 2 — raise `eta`. The most reliably cheap lever [Curzi, 2026b]. Interaction that happens and interaction that sticks are different quantities and only one is in the equation. Capture mechanisms are administrative: a designated scribe, a template, a decision log, a publication deadline, a named owner for each new unit of order, an apprenticeship. **Because `eta` multiplies `ln N`, it is the only lever that raises the ceiling itself rather than moving the system toward it.**

Remedy 3 — shed `Q`. The remedy nobody proposes. Shedding stored order to regain headroom is legitimate and underused: sunset clauses, term limits, planned demolition, retiring a product line, forgetting by design. **The arithmetic is exact — shedding `q` units returns `delta x q` of headroom per period — so shedding buys headroom in strict proportion to `delta`.** In a fast-decaying system it is decisive; in a durable one it is nearly worthless. In the worked office of Section 8.1, shedding 14 rules from a stock with a 4.31-year half-life returns 0.0433 rules per week, where the same shed at `delta = 0.05` would return 0.70 — a factor of sixteen. **The multiplication is computed before the shed is proposed; if the product is small, the proposal destroys something for nothing.**

One distinction is carried from the survey protocol and is easy to lose. **Striking units that were already inert is not shedding `Q`; it is correcting the order observable** [Curzi, 2026i]. Correcting the measurement changes the number and not the world. Shedding has real effects, including effects on the people whose work the shed units were.

16.3.13 4.5 The `h = 0` branch: hold

`h = 0` is maintenance only. What is permitted is protecting what is there and funding the maintenance regime **explicitly**, because working infrastructure produces no observable signal and is therefore chronically underfunded — infrastructure becomes visible only on breakdown

[Star and Ruhleder, 1996]. Substitution is permitted: retiring one unit of scope to make room for another is admissible at `h = 0` and adding one is not.

What is not permitted is describing `h = 0` as stability and treating it as safe. `delta Q` grows whenever `Q` grows and `C` grows whenever the organisation grows, so **`h = 0` is an unstable point that drifts negative on its own.** A system at `h = 0` is one hiring round, one new reporting requirement, or one departure away from `h < 0`, with no margin to absorb any of them. Organisational slack is the classical name for the margin that is missing here [Cyert and March, 1963].

16.3.14 4.6 The `h > 0` branch: choosing `r <= h`

The rule is absolute. **`r <= h`**: CC-T1 states that nothing based on interaction geometry guarantees more, and `r_max = h` is attained only at `U = 1`.

1. **The procedure forbids taking `r = h`.** That is the extremal value, attainable only at perfect orthogonality of every category from every other. Taking it means promising the best case as the plan.
2. **Reserve for parameter drift.** Non-parallelism decays, `C` grows, staff leave. The working convention is **`r <= 0.7 h`**, with the reserve stated explicitly as a reserve rather than quietly absorbed. The 0.7 is **a stipulated convention, not a derived quantity**, and it is badged as such wherever it appears.
3. **Check it against what the system has been doing.** If the realised one-step change is negative, a positive `r` requires the levers of Section 4.9 to have landed first. `r` is a promise about the future state of the parameters, not the current one.

16.3.15 4.7 Gate 2 — feasibility, the two failures, and the concavity correction

CC-T2 gives the required surplus and its geometric inverse:

```
S_req = (delta Q + C + r) / eta          feasible only if S_req <= ln N

U_req = (exp(S_req) − 1) / (N − 1)
```

There are two distinct failure modes with different fixes.

Failure A — the budget fails. `S_req > ln N`, equivalently `r > h`. The increment lies outside the ceiling and no geometry delivers it. Fix: reduce `r`, or raise `h` with the headroom levers of Section 4.9.

Failure B — the geometry fails. `S_req <= ln N`, so the increment is inside the budget, but `U_req` exceeds the measured capacity. The budget permits it; the categories cannot deliver it, because they are not that different from one another. Fix: raise capacity — recruitment across training lineages, rotation, adversarial roles, the deliberate outsider — or raise realised non-parallelism toward capacity, or reduce `r`.

Confusing the two is the most common design error at this gate, because both present as an inability to hit the target. Failure A is an accounting problem; Failure B is a personnel and protocol problem. They cost different amounts and take different lengths of time, and they are routed to different parts of an organisation.

The concavity correction, which is applied before either check. `U_req` is the geometry required *if `U` were constant every period.* It never is. What accumulates is the sum of `eta F(U_t)`, so the operative quantity is the **mean of `F`**, not `F` of the mean; and because `F` is strictly concave (CC-T6), Jensen's inequality gives

```
mean of F(U_t)  <=  F( mean of U_t )        with equality only if U is constant
```

A system whose geometry varies delivers strictly less surplus than a system sitting at its own average every period. The requirement is therefore stated as `mean F(U_t) >= S_req` and compared against the measured **mean of `F`**, never against `F` of the measured mean. `U_req` may be reported for legibility; it may not be tested against. The gap is the **concavity penalty**, it is often several per cent of the surplus, and regularising a variable geometry recovers it at no cost. Section 8.1 measures the penalty at 0.0761 nats in a case whose mean `F` is 1.1498 and whose `F` of the mean `U` is 1.2262 — a distinction that the notation `F(U)` invites an analyst to lose, and which Section 9.4 records as a defect in the source's own notation.

16.3.16 4.8 Creating headroom against realising it

Section 3.3 defines the realisation deficit `h - step = eta(ln N - F(U))`. Its operational content at this point in the procedure is a routing rule.

Lever	Acts on	Raises `h`?	Raises `step`?
`eta`	the ceiling coefficient	**yes**	yes
`C`	dissipation	**yes**	yes
shed `Q`	the maintenance charge	**yes**, by `delta x q`	yes
`N`	the ceiling, logarithmically	**yes**, net of path cost	yes
capacity `u`	attainable geometry	**no**	yes
realised `u`	geometry actually recorded	**no**	yes

The realisation sub-levers cannot take a system past `h`. A second opinion from someone whose training and incentives replicate the first is worth `ln(1) = 0`; and conversely the largest single gain available anywhere in the framework is the second genuinely different member. **Neither fact raises the ceiling.** The evidence that difference produces value is real and is bounded: Hong and Page [2004] give the formal result, and Thompson [2014] shows that the result's popular reading outruns what was proved — the pair is cited together here and never one without the other, as are the subsequent repairs [Singer, 2019]. **The relationship is not monotone in the field:** Uzzi and Spiro [2005] find an inverted U, creative success rising and then falling with cohesion, and Putnam [2007] finds that diversity in the short run reduces trust and civic participation, with Abascal and Baldassarri [2015] as the rebuttal. **The framework's surplus is monotone increasing in `u`; that is a property of the model, and the field evidence says the real curve turns over.** Contact across difference produces value only under conditions [Allport, 1954, Pettigrew and Tropp, 2006], and `u` needs the same qualification.

16.3.17 4.9 Gate 3 — the lever cost ranking

The default ranking, in ascending order of cost per unit of effect: `eta`, then `C`, then `u`, then `N`. Shedding `Q` is a conditional lever whose value is `delta x q` and which therefore belongs wherever `delta` puts it. **The ranking is a default and not a law**; the procedure requires an analyst to compute their own system's costs and to report an inversion when one occurs (Appendix A, IS-13). **A lever can also be exhausted rather than merely expensive**, and the corpus computes an instance of that in its frequency register: on a carried dyad the reserves available from non-parallelism and from cost, 5.545206 and 5.963140 octaves, each close a deficit of 4.338420, while the conversion reserve at its own ceiling supplies only 4.174198 and cannot close it at all [Curzi, 2026]. **A ranking by cost answers which lever to pull first; it does not answer whether a lever has enough left in it to finish**, and both questions are computed separately.

1. **`eta` — usually cheapest, and the first move in almost every real system.** Most organisations sit near the floor without knowing it, because capture is nobody's job. The fixes are administrative rather than structural — a scribe, a template, a named owner, a publication deadline — and require no hiring, no reorganisation and no consequential budget line. `eta`

multiplies `ln N`, so it raises the ceiling itself, and it is the only lever with that property that does not require adding people. It is neglected because "we had a great conversation" is a report of `eta = 0` and feels like a success. **Its limit is real: `eta` cannot exceed the rate at which the durable medium absorbs structure**, and a scribe writing faster than the register can be reviewed produces a backlog rather than order.

2. **`C` — next cheapest and usually fastest to act.** Rework loops and contradiction maintenance are pure dissipation; once counted in proxy units they are hard to defend, and cutting them requires a decision rather than a resource. It ranks below `eta` because `C` has a floor — some coordination *is* the work — where `eta` has no comparable floor in the ranges real organisations occupy.

> *Caution.* **Slack is not cost.** Waiting time explodes as utilisation approaches one [Kingman, 1962, Hopp and Spearman, 2011], so unused capacity is a cost-control device. An enthusiastic cost-cutting programme that removes slack raises `C` in the following periods through queueing, and the analyst will not connect the two.
>
> *Second caution.* Cut rework, not the places where disagreement surfaces. If the appeal process is where non-parallelism is realised, cutting it cuts `C` and cuts `F(U)` at once.

3. **`u` — cheap in money, expensive in time and political capital, and it realises rather than creates.** The realised sub-lever is often nearly free: an attendance rule, a recorded disposition requirement, a change of table. The capacity sub-lever is not: it means recruitment, rotation, and the deliberate importation of people whose training and incentives do not replicate the incumbents'. That is slow, and it is the change organisations resist most, because hiring for cultural fit, promotion by similarity and shared training are precisely the mechanisms that drive `u` toward zero and every one of them is popular. **`u` is not a property acquired once; it is a quantity that decays and must be maintained**, and cascades collapse it even in a group that began diverse [Bikhchandani et al., 1992]. Budget for its maintenance as a recurring cost, not as a project.

4. **`N` — most expensive and least rewarding, because the return is logarithmic.** Going from `N` to `N+1` adds `ln((N+1)/N)` nats to the ceiling, multiplied by `eta`. **Doubling the number of categories adds `ln 2 = 0.693` nats wherever one starts** — the same absolute gain from 2 to 4 as from 50 to 100.

N	ln N	gain over N-1	paths n(n-1)/2	paths added
2	0.6931	0.6931	1	1
3	1.0986	0.4055	3	2
4	1.3863	0.2877	6	3
5	1.6094	0.2231	10	4
6	1.7918	0.1823	15	5
8	2.0794	0.1335	28	7
10	2.3026	0.1054	45	9
12	2.4849	0.0870	66	11
16	2.7726	0.0645	120	15
20	2.9957	0.0513	190	19
30	3.4012	0.0339	435	29
50	3.9120	0.0202	1225	49
100	4.6052	0.0101	4950	99

Surplus grows logarithmically while coordination cost grows with the number of pairs [Brooks, 1975, Conway, 1968]. The gain column falls and the paths-added column rises, and they cross; where they cross is the maximum useful team size and it is computable from a system's own `eta` and per-path cost. **The consequence to be said out loud to every sponsor**

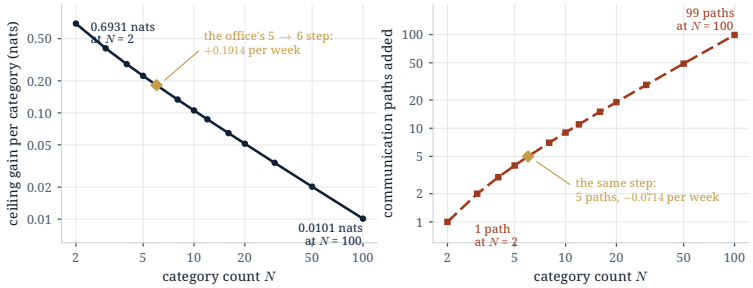


Figure 16.1: **Logarithmic gain against quadratic cost, on the source's own table.** The marginal ceiling gain `ln N − ln(N−1)` falls from 0.6931 nats at `N = 2` to 0.0101 at `N = 100` while the number of communication paths added by the same step rises from 1 to 99. The two curves are plotted against a common index because they are not commensurable in value — the crossing point is set by a system's own `eta` and per-path cost and is not a universal constant. Read from the `N`-lever table of Section 4.9; generated by `figs/src/CCP-016_fig03.py`.

is that a headroom problem cannot be escaped by adding domains. The same arithmetic is computed in another domain with the sign made explicit: at a 15 per cent annual hard-currency drag, a local-currency allocation of 25.17 per cent of protected capital sets headroom to exactly `-0.500000` and **cancels the whole `eta ln N = 0.377583` diversification gain**, while a 10 per cent allocation costs 0.150000, or 39.73 per cent of that gain [Curzi, 2026]. **A cost term attached to the same move that raises `N` can consume the entire ceiling gain and more**, which is why IS-17 requires the net effect and not the gross. The most common response to `h < 0` is to propose a new function, working group or discipline; the ceiling gain is small and logarithmic and is paid for with a cost that grows quadratically.

The one condition under which `N` is worth it is that the new category be **genuinely distinct** rather than relabelled. A sixth discipline duplicating an existing one raises `N` on the organisation chart and raises `C` for real while raising `u` not at all. The merge rule — categories with pairwise `u` below 0.10 are one category — is the test, and it is applied to the *candidate* before hiring, using the probe set [Curzi, 2026i]. Section 8.1 prices the difference at a 34 per cent haircut on the headline benefit.

5. Shed `Q` — conditional, valued at `delta x q`. Rank it by computing the product and inserting it into the list. In a durable institution it is last; in a fast-decaying one — a codebase with a two-month half-life, a training regime, a stock of active campaigns — it can outrank everything.

16.3.18 4.10 Gate 4 — sequencing: stabilise before growing

Every lever that reduces the deficit precedes every lever that adds scope. This is arithmetic rather than temperament: adding scope raises `Q`, and `delta Q` is subtracted from the same budget the intervention was repairing. **Growth taken before stabilisation consumes the headroom it was meant to produce.**

```
Phase 0   MEASURE          The survey. Not skipped, not shortened.
Phase 1   STOP THE LEAK    Cut C. Raise eta. Both act immediately and
                           neither adds scope. Recompute h.
Phase 2   CORRECT Y        Liveness and applier audits. Changes the number,
                           not the world, and precedes any shed.
Phase 3   SHED, IF delta x q IS WORTH IT
Phase 4   REALISE          Close the realisation deficit: attendance rules,
                           recorded dispositions, table geometry, written
                           positions before discussion. Cheap, fast, reversible.
```

```
Phase 5   RE-MEASURE h     Everything above changed it.
Phase 6   COMMIT r <= h    With a reserve. Preregister. Pilot.
Phase 7   RAISE CAPACITY u Slow. Recruitment, rotation, the outsider.
Phase 8   RAISE N          Last, and only against the merge test.
```

Phase 5 is the phase that gets dropped and the phase that matters most. Phases 1 to 4 change `eta`, `C`, `Q` and `F(U)` — every term in the gauge. Committing to an increment computed from the pre-intervention headroom is committing to the wrong number, and it is wrong in both directions: under-committing wastes headroom that now exists, over-committing breaks a promise.

The sharpest instance of the sequencing rule concerns trust, and it inverts naive practice. Trust is the intervention most often reached for in a system going badly. In a worked consortium reported elsewhere in this corpus, **trust alone moved the steady state from `-0.0836` to `-0.2114` — worse, by a factor of two and a half.** Reciprocity alone moved it to `+0.7394`; the two together to `+1.8701`. The mechanism follows from CC-T4: `Q_inf = (eta F(u*) - c*)/delta`. **Trust acts on the denominator — it reduces `delta`, the rate at which stored order leaks away — and when the numerator is negative, dividing by a smaller positive denominator makes the result more negative.** Trust does not add order; **it amplifies whatever sign the system already has**, and a system with `eta F(u*) < c*` already has a negative one. A high-trust failing institution fails harder and holds its failure longer, because trust is exactly the property that stops things from being forgotten.

Design rule: fix the numerator first. Raise `eta` or cut `C` until `eta F(u*) > c*` — Phases 1 and 2, in that order — and only then invest in trust, which belongs in Phase 6 or later for arithmetic rather than temperamental reasons. **An intervention beginning with a trust-building programme in a system whose numerator is negative is not a slow start; it is a lever pulled in the wrong direction, and it will be defended by everyone who felt the relationships improve.** The reciprocity half of the same result is where the external literature is strongest [Axelrod, 1984, Ostrom, 1990]; the trust half is a claim about a denominator and is derived from CC-T4, not from that literature.

16.3.19 4.11 Gate 5 — reversibility as a hard design requirement

A design that cannot be undone is not admissible under this procedure, regardless of its expected value. The requirement has three justifications, in ascending order of weight.

Epistemic. The model failed at least one falsification rule, or passed with a finite holdout and stated tolerances [Curzi, 2026i]. The analyst is acting under quantified uncertainty, and a design that cannot be undone converts an estimate into a commitment.

Structural. Interventions change `Q`, and `Q` is subtracted from the gauge through `delta Q` forever. An irreversible addition is a permanent charge against all future headroom. **Every irreversible intervention is a decision made on behalf of every future decision-maker in the system.**

Ethical. Reversibility is what makes declination meaningful. If a person may decline but the change cannot be undone, the declination is a gesture.

The reversibility statement answers five questions in writing. **These are labelled RV1 to RV5 here; the source labels them R1 to R5, which collides with this paper's own contributed results, and the collision is resolved by renaming the source's.**

```
RV1  What exactly returns to the prior state?
RV2  What does NOT return, ever?  (be specific: trust, a person who left,
     demolished fabric, a public commitment, a dissolved team, elapsed time)
RV3  Who is authorised to execute the reversal, without asking anyone?
RV4  How long does the reversal take, in periods?
RV5  What does the reversal itself cost, in proxy units?
```

RV2 is the load-bearing question. Most designs have a substantial irreversible component nobody has named, and naming it usually changes the design. RV3 matters nearly as much: a reversal requiring the approval of the person who authorised the intervention will not happen. The design and safety literature reaches the same conclusion from a different direction — interactive complexity plus tight coupling makes some failures a property of the system rather than of its operators [Perrow, 1984], and redundancy added without a reversal path can reduce safety rather than raise it [Sagan, 1993].

Class	Definition	Rule
Fully reversible	prior state restorable within a few periods at low cost	proceed under the normal gates
Costly-reversible	restorable, but the reversal costs a stated number of proxy units	proceed only if the reversal cost is inside `h`
Partially irreversible	some named component does not return	proceed only if that component is separately justified and separately consented
Irreversible	prior state unrecoverable	**not admissible**

16.3.20 4.12 Gate 6 — consent

Every affected person has a stated mechanism to decline without penalty. If none exists, the intervention is out of scope and does not proceed. This is postulate P-E of Section 3.6, and it is a boundary rather than a trade-off: it is not priced, because it is not for sale.

The gate has an engineering content as well as a normative one, and the engineering content is what makes it operative here. **Coerced non-parallelism is not non-parallelism.** A person compelled to voice a difference voices the difference the compeller expects, which returns the system to a parallel wall by a longer road. The declination mechanism is therefore not a concession attached to a design; it is part of the measurement.

A design whose parameters have no defined value under declination has not actually offered one. The test is a single written question for every protocol in the design: *what happens to a person who declines?* If the answer includes any adverse consequence, the protocol is out of scope. Section 8.1 exhibits the honest form: a discipline lead may abstain and the item proceeds with the remaining disciplines at a correspondingly lower `U_t`; an analyst offered the scribe role may decline, in which case the `eta` lever is unavailable and `h` is recomputed without it. **The declination is not free — it changes the arithmetic, and the arithmetic is published — but it carries no penalty to the person, and the design has a defined behaviour when it is exercised.**

16.3.21 4.13 Gate 7 — the pilot at capped downside

A pilot is not a small version of the intervention. A pilot is an experiment whose worst outcome has been computed and can be afforded.

1. **Cap the downside in proxy units, not in money.** State the maximum stored order that can be lost if the pilot fails completely, and confirm the system absorbs it. The admissibility condition `C_t <= (1-delta)Q_t + eta S_t` guarantees `Q_{t+1} >= 0` [Curzi, 2026b]; a pilot that could violate it is not capped.
2. **Bound it in scope and in time.** One team, one room, one product line, one quarter. The bound is what makes the downside computable.

3. **Preregister it**: the numeric prediction, the falsification criterion, both stopping rules, the consent mechanism, the reversibility procedure. Preregistration is the appropriate evidentiary form here for the reasons the reform literature gives — it separates the confirmatory from the exploratory and removes the analytic flexibility that otherwise makes a favourable result uninformative [Nosek et al., 2018, Simmons et al., 2011, Gelman and Loken, 2013] — and it is what converts a prediction into a risk [Platt, 1964]. **The corpus's own defence against the standard objection that directional predictions in the social sciences are nearly costless [Meehl, 1967] is that it makes point predictions with constants**; that defence is only good where the point prediction is actually declared in advance, which is what this gate requires.
4. **Respect the conversion lag.** The prediction must not begin before `m` periods have elapsed, where `m` is the measured lag from interaction to durable artefact. Testing earlier is designing a failure.
5. **Keep a matched comparison where one exists honestly.** A second team, a second room, a preceding period of the same length. Where no honest comparison exists, say so: **a pilot with no comparison is a demonstration, and demonstrations are evidence about feasibility, not about effect.**
6. **Under the accounting frame** — second rule failed, third passed — the pilot carries an extra restriction: it is not a confirmation of the model, and its preregistered prediction is stated as a prediction about the *intervention*, not as a test of the storage law. It may nonetheless supply the causal evidence the estimation could not, which is precisely why it is worth running (Section 3.4).

One caution on the pilot's own measurement. A survey period that shows improvement may have told the analyst about the survey. The Hawthorne pattern is invoked far more casually than the evidence supports [McCambridge et al., 2014], which cuts both ways: it is not a sufficient explanation of a favourable result, and it is not a sufficient dismissal of one either. The remedy is the matched comparison of rule 5 and the mediation check of Section 8.2, not an appeal to the pattern's name.

16.3.22 4.14 The three-posture table

Every intervention encounters behaviour it did not plan: an informal practice, a dissenting group, a workaround, a shadow process, a person who does it differently. Three postures are available and each has a determinate effect on the parameters. **The default is *ignore*, and ignore is a choice with consequences.**

Posture	What is done	`u`	`eta`	`C`	`Q`
ENGAGE	bring it into the record: a category, a seat, a recorded disposition	rises toward capacity	rises: the practice becomes capturable	rises modestly: one more path, one more voice	rises if it converts
IGNORE	leave it outside the record	decays: the practice drifts and stops informing the centre	zero for that practice: it produces nothing durable	low now, unbounded later: a shadow process accrues its own maintenance	flat; the order it makes is not the system's

Posture	What is done	u	eta	C	Q
SUPPRESS	forbid it	falls, often sharply: suppression teaches every category not to differ	falls: the suppressed practice was producing, and adjacent practices go quiet	rises: enforcement, monitoring, appeals, and the attention tax of everyone watching	falls: the order the practice held is destroyed

The risk profile of each is different in kind and not only in size. **Engaging has a bounded and visible cost**; its failure mode is that the practice is genuinely redundant, so C has been bought for no gain in u, and the remedy is to test the candidate against the probe set before engaging. **Ignoring is the default and it is a slow loss**; its failure mode is that the shadow process becomes load-bearing while remaining unmeasured, so Q and C are both wrong and the survey that reports them is inadmissible. **Suppressing is worst on every parameter at once**, and the consent gate is decisive against it independently of the arithmetic: suppression that removes a person's ability to decline is out of scope (Section 4.12).

Suppression is the only posture that damages all five parameters simultaneously, and it is the posture organisations reach for first, because it is the only one that feels like action. Neither the arithmetic nor the constraint of Section 4.12 permits it. The constructive alternative is to make the alternative position the more comfortable place to stand and let people arrive there; its worst outcome is that someone declines, which is a permitted outcome.

16.3.23 4.15 Time constants, and what may be promised

The contraction guarantees arrival (CC-T5). It says nothing about arriving this quarter.

```
time constant  tau = 1/delta        periods
95 per cent of the way to a new equilibrium:  about 3 tau
99 per cent:                                  about 5 tau
```

The rule that follows is short: promise increments, never equilibria. r is a per-period increment observable within a few periods of the conversion lag; Q_inf is where the system ends up after roughly three time constants [Oppenheim et al., 1996], which in a durable institution is a span of years or decades. A sponsor told the equilibrium hears a promise; a sponsor told the increment can check it. **The evidence that people systematically misjudge stocks with delays is direct and experimental** [Sterman, 1989], which is why the four numbers below are said out loud at the commitment meeting and entered in the specification:

1. the conversion lag m — nothing is observable before it;
2. the increment r per period, and the horizon over which it is judged;
3. the time constant tau, in calendar units;
4. 3 tau, in calendar units, labelled the settling time and explicitly **not** a delivery date.

A low delta is what makes an institution durable and it is the same number that makes it slow. Durability and speed cannot be bought from the same parameter, and any promise implying they can is a promise the arithmetic will break.

Cadence against intensity — a correction, because the usual argument for cadence is false. Given a fixed total input, an intervention may deliver it as one large occasion or spread it periodically, and practitioners hold that cadence delivers more. **It does not, and this is computed rather than asserted: the integrated effect of one impulse and of the same total spread over any schedule is identical to every digit.** Summing the exact solution of the storage law

gives `sum_t Q_t = (total input)/delta` regardless of how the input is distributed in time; the corpus reports 433.3333 in all three regimes it tested. **That is linear time-invariance** [Oppenheim et al., 1996], and a procedure claiming otherwise would be claiming its own governing equation is not linear.

What cadence actually buys is three things, each of them real.

1. **The steady state differs absolutely.** An impulse decays to zero; sustained periodic forcing at mean rate `m` holds `m/delta`. The integrated effect is the same; the level that can be stood on is not, and for an intervention that must leave a platform behind, the level is the deliverable.
2. **The ceiling clips the impulse.** `S_t <= ln N` (CC-T6), so a one-off input larger than the ceiling is truncated and the excess is simply not received. In a worked annual rite reported elsewhere in this corpus the clipping destroys **97.2 per cent of the nominal budget** — the money is spent and the surplus is not admitted, because no geometry pushes a single period above `ln N`.
3. **Concavity favours splitting, independently of the ceiling.** Strict concavity of `F` (CC-T6) alone makes an input split across occasions strictly dominate the same input delivered at once. **This is the same Jensen arithmetic as Section 4.7 read forwards: variability is charged against the system when the geometry varies, and paid to it when the input is spread.**

So cadence beats intensity because of the ceiling and because of concavity, not because it delivers more. It delivers exactly the same amount. The honest counter-case is kept in print: above a threshold no realistic cadence reaches, the impulse wins. **The true reason is stated when a cadence is specified, because the false reason will not survive the first sponsor who checks it.**

16.3.24 4.16 Monitoring, trip-wires, and the standing gauge

An intervention without instrumentation is an assertion. Three things are installed before it starts.

The standing gauge. Recompute `h` every period from the same instruments with the same definitions, and post it — not the order proxy, and not a dashboard, but `h` and its three components, so a reader can see which term is moving. **Keeping the order proxy itself out of the reporting line is the primary defence against the measure becoming a target** [Goodhart, 1975, Campbell, 1979].

Trip-wires. Pre-declared thresholds that trigger a specific action without further discussion, because discussion at the moment of a threshold is how thresholds get moved.

Trip-wire	Threshold	Action
realised geometry falls	mean `F(U_t)` below `S_req` over a declared block	the promised `r` is no longer feasible; notify the sponsor in the same period
`C` rises	above the level at which `h < r`	suspend the growth commitment; resume the cost work
`h` crosses zero	`h <= 0`	stop building immediately; revert to Section 4.5
any single-period rise in the proxy	exceeds `eta ln N` plus lag allowance	investigate the proxy before believing the number
a harm criterion fires	as preregistered	stop, regardless of the proxy

The geometry trip-wire fires most often and is watched least. Non-parallelism decays; the promise was computed at a level of realised surplus the system must actively maintain; and the decay is invisible because it looks like agreement.

The review cadence. At least every `tau/4` periods, and at least once per conversion lag. In a system with a six-year time constant, quarterly is ample and monthly is theatre. The contraction's a posteriori bound gives the reading a second use: the size of the most recent correction, scaled by `q/(1−q)`, bounds the remaining distance to arrival [Curzi, 2026]. A control-theoretic caution belongs here as well: a correction gain applied outside the stability band of a first-order loop does not merely act slowly, it diverges [Åström and Murray, 2008], so a monitoring regime that responds to every deviation with a larger correction is a design fault and not diligence.

16.3.25 4.17 When the procedure directs no intervention

The default is not to intervene. Interventions cost `C` while they run, raise `Q` if they succeed — which permanently charges `delta Q` against all future headroom — and consume an organisation's finite tolerance for being changed. **The burden of proof is on action.** Eleven conditions follow, each a hard stop except where noted.

1. The system has not been measured. No survey, or a survey with a failed repackaging test. No argument survives this, including urgency. **An intervention designed on an unmeasured system is a guess with a budget, and its most likely effect is to move a parameter that was not the binding one.**

2. `h < 0` and the proposal is to build. CC-T3. The remedies of Section 4.4 are repairs and are permitted; building is not.

3. `r > h`. No argument, no exception, no phasing. **A request exceeding the budget is refused and the refusal is the deliverable. The number is said.**

4. `U_req` exceeds capacity. Failure B of Section 4.7. Building anyway produces a shortfall that will be attributed to effort rather than to geometry, and the organisation will conclude that people did not try hard enough. **That conclusion is worse than not having acted.**

5. The design is irreversible. Section 4.11. Not admissible, regardless of expected value.

6. Any affected person cannot decline without penalty. Section 4.12. Out of scope. This is not a severity judgement about the intervention; it is a statement about what this procedure produces.

7. The system is already at its uniform bound. `Q_t` at or near `max{Q_0, eta ln N/delta}` (CC-T5). Nothing designed exceeds it. An intervention aimed above the bound is aimed at a number the model has already forbidden, and the effort will appear entirely as `C`.

8. The binding constraint is outside the model. The survey produces six numbers about interaction, conversion, decay and cost. If the actual constraint is a legal prohibition, a funding cliff, a safety failure, a physical limit, or a person in acute distress, then the framework is measuring a room in a building that is on fire. **Say so plainly and stop; the strongest thing an analyst can do with a model is state where it does not apply.**

9. The measurement is the intervention. Where the survey itself has changed behaviour enough that the parameters are already moving, wait a full `tau/4` and re-measure before adding anything. Acting on top of an unresolved measurement effect makes both uninterpretable.

10. The system is in an identifiable transient. `(1−delta)^t Q_0` decaying while the driven term builds looks exactly like decline. Before intervening on a declining proxy, check whether it is a transient from a known change within the last `3 tau`. Intervening in a transient adds a second transient and neither can be attributed.

11. A cheaper lever has not been tried. Not a hard stop but a discipline: if `eta` has never been examined and the proposal is to hire, the proposal is premature. Section 4.9's ranking exists to be checked, in order, before a design is accepted.

Two closing dispositions.

Not intervening is a decision and is recorded as one. A survey ending in *no intervention is warranted* is a completed piece of work with a deliverable: the six numbers, the sign of `h`, the reason, and the trip-wire that would change the answer. **File it, date it, and name the observation that would reopen it. Undocumented restraint is indistinguishable from neglect and will be read as neglect.**

The request for an intervention is not evidence that one is needed. The analyst was called because something is uncomfortable. Discomfort is a gradient, and a gradient indicates which way is downhill; it does not establish that the system has a budget for the journey. **The service being paid for is the number, and sometimes the number is `-0.17`.**

That said, the gradient is a real measurement and is treated as one. Since `h = eta ln N - delta Q - C` and only `C` varies with a local design choice, `grad h = -grad C`: the direction of steepest increase in headroom is exactly the direction of steepest decrease in cost. Discomfort, at the point where it is felt, is `C`. **Therefore a cheap desire-line survey — where do people work around this system — is a headroom measurement**, and the full instrument programme is not always required to locate the direction of steepest improvement. Direct observation of where people actually go, against where the design intended them to go, is the classical instrument [Whyte, 1980, Jacobs, 1961, Alexander et al., 1977]. **The division of labour is exact and is worth stating as a rule: the cheap survey gives the direction; only the full survey gives the sign and the size of `h`, which is what decides whether the analyst may move at all.** Run the cheap one first, every time, and do not let it authorise a commitment.

A correction to the received advice about deliberate discomfort. *Leave the comfort zone* is usually stated as a general principle, and stated that way the arithmetic rejects it — but the rejection is narrower and more interesting than a flat denunciation, and both branches have been measured in this corpus. **Raising cost globally is strictly harmful**: in a worked landscape the time to reach a better configuration runs 32, 35, 42, 89, 297 weeks and then never, as global cost rises. **Raising cost only at the origin genuinely shortens the wait — 32.0 weeks down to 9.9 — and charges for it: the permanent fund that results is 37 per cent smaller.**

The operative question is therefore not whether to add discomfort but whether it is a cost that can be left behind or a cost that is carried. Making the current position expensive is a departure lever; it works, at a permanent price. Making the whole landscape expensive destroys the platform the system was trying to leave *for*. An intervention raising `C` everywhere in the name of urgency has raised `delta Q + C` against `eta ln N` in every future period, which is the `h < 0` branch arrived at deliberately. **Price the discomfort, name which kind it is, and put the 37 per cent in the specification next to the 9.9 weeks.** The related result on momentum — that momentum at `beta = 0, 0.5, 0.8, 0.9, 0.95` all terminates at the same `x = 1.9972` in the same basin, so momentum defeats narrow barriers and is useless against wide shallow ones, and is refused entry to the formulary — is established in Curzi [2026p] and bounds what a departure lever can be expected to do.

16.4 5 Faults of the procedure itself

Twelve faults attach to the procedure rather than to the system it is run on. **They are labelled PF-1 to PF-12 here; the source labels them F26.1 to F26.12 against a chapter number that no longer exists.**

PF-1 — Building at negative headroom. *Description.* `h < 0` and the intervention adds scope anyway, usually described as investing a way out. *Damages.* `Q`, then `C`, then morale. *Signature.* A growth plan whose first phase does not reduce `C` or raise `eta`. *Detection.* Gate 1; it is one subtraction. *Fix.* Refuse, and return to Section 4.4. **The refusal is the deliverable and the number is the argument.**

PF-2 — The promised equilibrium. *Description.* `Q_inf` is quoted as the outcome without `3 tau`. *Damages.* Credibility, and eventually the intervention itself when a sponsor stops it early for under-delivery. *Signature.* A target level with no per-period increment attached. *Detection.* Ask what is expected next quarter; if the answer is a fraction of `Q_inf` computed by

straight-line interpolation, the fault is present, because approach to a fixed point is exponential and not linear. *Fix.* Section 4.15: promise `r`, report `Q_inf`, state `3 tau` and label it settling time.

PF-3 — Confusing Failure A with Failure B. *Description.* An infeasible target is treated as a budget problem when it is a geometry problem, or the reverse. *Damages.* Wasted lever effort; the wrong department is asked for the fix. *Signature.* More money offered for a problem where `S_req <= ln N` but `U_req` exceeds capacity. *Detection.* Compute both checks of CC-T2 and report them separately. *Fix.* Route Failure A to `eta`, `C` and `Q`; route Failure B to recruitment, rotation and realisation protocols.

PF-4 — The hiring reflex. *Description.* `h < 0`, so a new function, working group or discipline is added. *Damages.* `C`, with the number of pairs; `h`, immediately; and `u`, if the addition is a relabelling. *Signature.* A headcount proposal with no `ln((N+1)/N)` calculation and no path-count calculation. *Detection.* Compute the ceiling gain, compute the added paths, subtract. *Fix.* Run the merge test on the candidate first, then compare against the `eta` lever's cost per unit of effect — which in Section 8.1 is an order of magnitude better on every basis examined.

PF-5 — The shed that bought nothing. *Description.* Stored order is destroyed to create room in a system with small `delta`. *Damages.* `Q`, permanently; trust, permanently; `h`, by `delta x q`, which is negligible. *Signature.* A sunset or cull proposed without `delta x q` computed. *Detection.* One multiplication. *Fix.* Compute it. In Section 8.1 the shed returns 0.0433 rules per week against an `eta` lever returning 0.8828 — **a factor of twenty on that comparison**, which is a different comparison from the factor of sixteen between the same shed at two different decay rates, and the two are not to be conflated.

PF-6 — Slack cut as cost. *Description.* A cost-reduction programme removes unused capacity. *Damages.* `C`, upward, in following periods, through queueing. *Signature.* Utilisation rises toward one and waiting times begin to grow non-linearly. *Detection.* Track utilisation alongside `C`; the signature is unmistakable once both are plotted. *Fix.* Restore slack. **Slack is a cost-control device, not a dissipation** [Kingman, 1962, Cyert and March, 1963], and this is one of the few places where the intuitive reading of the equation is exactly backwards.

PF-7 — The unwatched geometry. *Description.* `r` was computed at a level of realised surplus the system stops maintaining. *Damages.* The feasibility of the promise, silently. *Signature.* Increasing agreement, decreasing meeting length, fewer recorded dissents — all of which read as improvement. *Detection.* The realised-geometry trip-wire on `mean F(U_t)`, plus an annual probe set. *Fix.* Reinstate the realisation protocols. If capacity itself has fallen, the problem is slow.

PF-8 — The irreversible pilot. *Description.* A pilot that dissolves a team, demolishes a wall, or announces publicly. *Damages.* Everything, because the option value that made it a pilot is gone. *Signature.* RV2 on the reversibility statement is blank or reads "none". *Detection.* RV2 is never honestly "none"; if it is blank, it has not been attempted. *Fix.* Redesign to a reversible form, or classify it as an intervention and put it through the full gates.

PF-9 — The default suppression. *Description.* An unplanned practice is forbidden without a posture decision being made explicitly. *Damages.* `u` down, `eta` down, `C` up, `Q` down — all four at once. *Signature.* A prohibition with no recorded alternative and no recorded reason. *Detection.* The posture table has a row for every unplanned practice or it does not. *Fix.* Reverse the prohibition and choose engage or ignore deliberately. Where a practice genuinely must end, make the alternative the more comfortable place to stand.

PF-10 — Re-measurement skipped. *Description.* `r` is committed from the pre-intervention `h` after Phases 1 to 4 have changed every term in it. *Damages.* The commitment, in either direction. *Signature.* One `h` value in the document, at the top. *Detection.* Count the `h` values in the specification; Appendix B has a slot for each phase because each phase has one. *Fix.* Phase 5. It costs an afternoon.

PF-11 — The intervention that was the measurement. *Description.* Acting on a survey whose administration already changed the behaviour it measured. *Damages.* Attribution, permanently. *Signature.* The proxy or realised geometry moves during the survey window itself. *Detection.* Plot the series against the survey's own dates. *Fix.* Condition 9 of Section

4.17: wait `tau/4`, re-measure, then design.

PF-12 — The unrecorded restraint. *Description.* The correct decision — do not intervene — is taken and not documented. *Damages.* Nothing now; everything in eighteen months, when the same question returns with no record of the six numbers or the reason. *Signature.* A survey with no closing decision page. *Detection.* Appendix A, IS-32. *Fix.* File the decision, the numbers, the sign, and the observation that would reopen it. # 6 The failure catalog: warrant, identities, and entry format

16.4.1 6.1 Why a finite catalog is possible

The storage law has exactly six degrees of freedom.

```
Q_{t+1} = (1 − delta)Q_t + eta S_t − C_t        S_t = F(U_t) = ln(1 + (N−1)U_t)

    delta   retention loss              in [0,1]
    eta     conversion efficiency       >= 0
    N       category count              integer >= 1
    u       non-parallelism             in [0,1]
    C       dissipation                 >= 0
    Q       stored order                >= 0
```

Every failure of a system modelled by this law is a claim about one of those six quantities, so there are exactly six kinds of failure however many names they wear. That is the whole warrant for a closed catalog, and it is worth being precise about its status: **the closure is a DEFINITION, true by construction given the model, and it carries no empirical weight.** It would be defeated not by an observation but by the adoption of a different model. The thirty-two names below are not independent phenomena; they are recurring *routes* into six parameter faults, distinguished because the routes have different signatures, instruments and fixes.

That the law describes a real institution at all is **badged as a postulate** (P-C of Section 3.6): a single scalar `Q` for a team is a choice of model, not a theorem. Where a fault sits outside the model — a legal shutdown, a fire, a death — the catalog will not name it, and the correct response is to say so rather than force a parameter.

16.4.2 6.2 The three diagnostic identities

Every entry is priced against one of three identities, and they are the only arithmetic the catalog needs.

Identity 1 — the effective count. `N_eff = 1 + (N−1)u`, and `S = ln N_eff`. Difference is counted before headcount is credited, so any fault lowering `u` lowers `S` non-linearly and invisibly, because headcount does not change. **This effective count is the affine one adopted by design axiom in Curzi [2026]. It is not a Hill number** (Section 3.2), and the diversity literature is not support for it.

Identity 2 — the one-step condition. `Q_{t+1} > Q_t` if and only if `eta S_t > delta Q_t + C_t` [Curzi, 2026b]. **Growth is a comparison, not a quantity.**

Identity 3 — headroom. `h_t = eta ln N − delta Q_t − C_t`. `h_t < 0` is unsustainable **in this model** and no interaction geometry rescues it (CC-T3); `h_t = 0` is maintenance; `h_t > 0` is the maximum available increment, `r_max = h_t` (CC-T1).

16.4.3 6.3 The distinction on which half of all misdiagnosis turns

`h_t` uses `ln N`. The one-step condition uses `S_t = ln N_eff`. **Headroom is the ceiling; non-parallelism decides how much of the ceiling is actually collected.** A system can have large positive headroom and shrink every period, because it operates far below its own ceiling.

The practical consequence is a rule: **compute both numbers, always.** A diagnosis reporting `h` without the realised fraction `S / ln N` is half a diagnosis and will send the analyst to the wrong lever. Section 8.2 is the case that exhibits it, and Section 12's counter-result is its statement.

16.4.4 6.4 Entry format, and the honesty clauses that constrain every entry

Every entry carries the same eight fields in the same order, so the catalog can be scanned in a column.

Field	What it holds
Damages	which of `delta`, `eta`, `N`, `u`, `C`, `Q` the fault attacks, and in what direction
Clinical	what the fault is, stated without metaphor
Signature	what is observable from outside, before anything is measured
Measurement	the specific instrument that confirms or excludes it
Price	the arithmetic cost, computed
Fix	the intervention, and which lever it pulls
Confusable with	the entries most often mistaken for this one
Elsewhere	where the mechanism is developed

The standing honesty clauses of Section 3.7 constrain every entry and are recalled here because the catalog is the part of the material most likely to be quoted without context. **Every price below is a floor**, since the surplus functional as deployed is a lower bound on the ensemble's information. **Where an entry cites an external result, the result is real and the transfer to any particular system is a modelling assumption**, badged at that point. **No thermodynamic law is invoked anywhere**, and every logarithm is a dimensionless count in nats.

16.5 7 The catalog

16.5.1 7.1 Class I — Geometry faults (`u` and `N`)

The largest class: non-parallelism decays silently, and `N` is the parameter people believe they have when they do not.

FM-01 · The Parallel Wall

Damages: `u -> 0`. Through Identity 1, the effective count collapses with it. **Clinical:** Every member of the ensemble occupies the same direction in state space. Their outputs are copies. The group behaves as a single instrument with a redundant power supply. **Signature:** Unanimity arriving fast and without argument. A second opinion that is procedurally distinct and substantively identical. Nobody can name the last time the group changed its mind because of an internal objection. **Measurement:** Independent estimation before discussion on 15 to 25 live questions; compute the mean pairwise correlation of signed error, `rho`; then `u = 1 - rho^2`. Confirm with an error-decorrelation audit on past decisions: if the group's mistakes were the *same* mistakes, `u` is low regardless of how the roster looks. **Price:** At `N = 12` the wall is total when `u = 0`: `N_eff = 1`, `S = ln 1 = 0` — twelve people purchase zero surplus. Partial walls are priced the same way. `u = 0.02` gives `N_eff = 1.22` and `S = 0.1989`, which is **8.00 per cent** of the

ceiling `ln 12 = 2.4849`; `u = 0.05` gives `S = 0.4383`, or **17.64 per cent**; `u = 0.30` gives `S = 1.4586`, **58.70 per cent**. **Fix:** Raise `u`: recruit across training lineages, import a practitioner from an unrelated trade, install adversarial roles, rotate. Raising `N` first is wasted money — see FM-07. **Confusable with:** FM-02, which is a partial wall with a demographic alibi; FM-05. **Elsewhere:** capacity recruitment, and the decay of non-parallelism over time.

FM-02 · Pseudo-Diversity

Damages: `u`, specifically the within-category component `u_div -> 0` while the cross-category component `u_cross` looks healthy. **Clinical:** The roster spans categories; the *opinions* do not. Difference is present at the level of biography and absent at the level of method, incentive and permitted conclusion. The group is genuinely heterogeneous in every dimension that does not enter the decision. **Signature:** A visibly varied group whose internal disagreements are all about emphasis. Dissent expressed as a question rather than as a claim. Recruitment filtering late on cultural fit. **Measurement:** Decompose. `u = u_cross + u_div`. Estimate `u_cross` from the between-category share of consequential interaction; estimate `u_div` from within-category disagreement on a live question. Report the two numbers separately. Pseudo-diversity is the profile *`u_cross` moderate, `u_div` near zero*. **Price:** With `N = 8`, `u_cross = 0.30` alone buys `f(u_cross) = 1.1314` nats. Adding `u_div = 0.25` takes the pair to `f(0.55) = 1.5790` — a **39.6 per cent** gain that the roster alone will never show, because nobody was hired to obtain it. **Fix:** Buy `u_div` directly: dissent protection, anonymous-first elicitation, a standing minority-report requirement, adversarial review inside each category. **Confusable with:** FM-01 (total rather than partial); FM-08, which is the *measurement* version of this fault. **Elsewhere:** the decomposition of non-parallelism into crossing and divergence.

FM-03 · The Cascade Collapse

Damages: `u`, dynamically — a group that began diverse converges over time. **Clinical:** Once enough public signals accumulate, rational actors copy rather than report, and the copies carry no new information [Bikhchandani et al., 1992]. Aggregation stops while the appearance of aggregation intensifies. **Non-parallelism is not a property acquired once; it is a quantity that decays**, which makes this a retention problem wearing a geometry costume. *The transfer of the cascade result to a standing committee is a modelling assumption; the result itself is about sequential public choice.* **Signature:** Later speakers agree more than early speakers. The order of speaking predicts the content of speaking. Confidence rises while accuracy does not. **Measurement:** An order-effect test. Elicit estimates in writing, sealed, before any discussion, then again after. If post-discussion variance falls by more than about a third while accuracy is flat or worse, the group has a cascade and not a convergence. Log speaking order against position taken over ten meetings and test for rank correlation. **Price:** A `u` decaying 8 per cent per period from 0.60 at `N = 12` gives `S` = 2.0281 at `t = 0`, 1.7454 at `t = 4` (86.1 per cent), 1.4787 at `t = 8` (72.9 per cent), 1.2316 at `t = 12` (60.7 per cent). Nothing visible changes. The roster is identical throughout. **Fix:** Written positions before any discussion of any consequential question — the cheapest instrument available, and it breaks cascades before they start. Add round-robin ordering and the inversion that has the most junior speak first. **Confusable with:** FM-04 — the same trajectory with a different driver: cohesion is social, a cascade is informational. **Elsewhere:** independence protocols, and the geometry of the table.

FM-04 · The Cohesion Death Spiral

Damages: `u`, dynamically; secondarily raises the perceived cost of dissent. **Clinical:** The group optimises for getting along. Getting along selects for agreement. Agreement selects membership and shapes expression. The loop is positive-feedback and it feels like health the

entire way down. **Cohesion is frequently the enemy of collective intelligence, because cohesion drives u toward zero** — and the field evidence puts a shape on that claim: creative success rises and then falls with cohesion, an inverted U rather than a monotone decline [Uzzi and Spiro, 2005]. Organisational learning shows the same structure from a different direction: members who learn too fast homogenise the organisation and destroy the variance that produced its advantage [March, 1991]. **Signature:** High reported satisfaction with low external accuracy. Long tenure with no recruitment from outside the lineage. A named "difficult" person who is in fact the group's last remaining source of `u` — and celebration when that person leaves. **Measurement:** Track `u` quarterly, not once, and plot it. A cohesion spiral is a monotone decline in measured `u` against flat or rising self-reported team health; **the two series diverging is the diagnostic, not either one alone.** Cross-check with turn-taking equality — the Gini coefficient of speaking time — which is one of the measured drivers of collective intelligence [Woolley et al., 2010]. **Price:** The same arithmetic as FM-03. What differs is the fix, which is why the distinction is worth making. **Fix:** Institutionalise the outsider so that dissent does not depend on one person's nerve: rotating red-team duty, a standing external reviewer, term limits on membership. **None of these may remove anyone's ability to decline participation** (Section 4.12). **Confusable with:** FM-03; FM-09, which is a cohesion spiral with a disciplinary boundary around it. **Elsewhere:** collective intelligence and turn-taking; independence protocols.

FM-05 · The Missing Denominator

Damages: `N` — specifically, the substitution of headcount for the effective count. **Clinical:** The system counts people, seats, sources or sensors and treats that count as `N`. The law counts *categories after a redundancy audit*. The gap between the two is the fault. **Signature:** Capacity claims quoted as headcount — "we have forty reviewers" — with nobody having been asked how many of the forty are distinguishable. **Measurement:** A roster audit. For each pair, ask whether their outputs on a live question were distinguishable; collapse indistinguishable pairs. Report the effective count alongside `N` in every capacity statement thereafter. **Price:** A roster of 40 whose members are 92 per cent aligned (`u = 0.08`) has `N_eff = 4.12` and `S = 1.4159` against a ceiling `ln 40 = 3.6889` — a shortfall of **61.62 per cent**. At `u = 0.05` the shortfall is **70.67 per cent**. The organisation believes it owns 3.69 nats of ceiling and owns 1.08. **Fix:** There is no fix at the `N` lever; this is a `u` fault discovered by an `N` error. Audit, restate capacity honestly, then treat as FM-01. **Confusable with:** FM-06, which is about the *definition* of the categories rather than the count; FM-01. **Elsewhere:** the effective count and the logarithmic ceiling.

FM-06 · The Frozen Denominator

Damages: `N` — the declared categories are administrative rather than epistemic. **Clinical:** `N` is read off an organisation chart, a funding structure or a historical department list. Those boundaries were drawn for reporting, budget or accident; the categories that actually carry distinct information are a different partition. A frozen denominator also hides *vacancies*: categories declared, staffed on paper, and epistemically empty. **Making a system legible to its centre is precisely the operation that destroys the local variety that made it work** [Scott, 1998], which is why the administrative partition is the wrong one so reliably. **Signature:** A category with a name and no active member. Two "different" departments whose outputs are never in tension. Reorganisations that change the chart and change no forecast. **Measurement:** Define categories by capability, not by title: *what can this category independently detect, fix or veto that no other can?* Any category failing the question is not a category. Count only staffed categories in `N`. **Price:** In an open-source case worked elsewhere in this corpus, seven categories are declared and four are staffed: `ln 7 = 1.9459` against `ln 4 = 1.3863`, so the project plans against a ceiling **40.4 per cent higher** than the one it owns [Curzi, 2026k]. Section 8.2 prices the same fault at **132 per cent** in a hospital ward declaring five handover categories

with two attending. **Fix:** Re-partition by capability; then either staff the empty categories — a real `N` gain, priced in FM-07 before it is assumed to be worth it — or delete them from every capacity calculation. **Confusable with:** FM-05. The test: FM-05 is counting too many *people* per category; FM-06 is counting too many *categories*. **Elsewhere:** the roster audit and the merge rule [Curzi, 2026i].

FM-07 · The Ceiling Chase

Damages: `C` rises, `N` rises marginally, `h` falls. A self-inflicted fault. **Clinical:** The system has a headroom deficit and attempts to escape it by adding domains, disciplines, partners or product lines. Because the ceiling term is `eta ln N`, additional categories deliver logarithmically while their coordination cost grows with the number of pairs [Brooks, 1975, Conway, 1968]. **A headroom problem cannot be escaped by adding domains. Signature:** A growth strategy announced in response to a shortfall. New partnerships as the answer to overload. Category count rising while the realised fraction `S / ln N` falls. **Measurement:** Before adding a category, compute two numbers — the ceiling gain `eta(ln N_new − ln N_old)`, and the projected coordination increment. Compare. Refuse if the second exceeds the first. **Price:** Going from `N = 4` to `N = 5` adds `ln 5 − ln 4 = 0.2231` nats, a **16.1 per cent** ceiling gain for a 25 per cent increase in categories. From 5 to 6: **0.1823** nats, **11.3 per cent**. Doubling from 50 to 100 adds `ln 2 = 0.6931` nats, a **17.7 per cent** gain, for double the coordination surface — meanwhile channels rise from 1,225 to 4,950. **Fix:** Compute `h` first. If `h < 0`, the permitted moves are cutting `C`, raising `eta`, or shedding `Q`. Adding `N` is the most expensive lever available and the least rewarding. **Confusable with:** FM-21, which the Ceiling Chase usually *causes*. **Elsewhere:** Section 4.9; the crossover between logarithmic gain and quadratic cost [Curzi, 2026b].

FM-08 · Two-Source Blindness

Damages: the *measurement* of `u`, and consequently the direction of every downstream lever. **Clinical:** The survey measures one of the two sources of difference and reports it as `u`. Since `u = u_cross + u_div`, an organisation rich in cross-category difference and starved of within-category dissent measures as diverse on a cross-category instrument; a room of ferociously argumentative specialists in one discipline measures as diverse on a disagreement instrument. **Both readings are half-true and both send the analyst to the wrong lever. Signature:** A diversity report with one number in it. A team scoring well on whichever instrument its leadership commissioned. **Measurement:** Always two instruments. `u_cross`: the between-category share of consequential interaction. `u_div`: within-category disagreement on a live question, elicited independently. Report both, and report the bound `max(f(u_cross), f(u_div)) <= f(u) <= f(u_cross) + f(u_div)`. **Price:** At `N = 8` with `u_cross = 0.30` and `u_div = 0.25`: `f(u_cross) = 1.1314`, `f(u_div) = 1.0116`, `f(u) = 1.5790`, and the sum is 2.1430. The bound is satisfied and it is wide: **a single-instrument report could plausibly claim anything from 1.01 to 2.14 nats. Fix:** Two instruments, two numbers, every survey, both fields required [Curzi, 2026i]. **Confusable with:** FM-02, which is the *condition*; FM-08 is the failure to see it. **Elsewhere:** the crossing-versus-divergence decomposition.

FM-09 · The Silo Local Minimum

Damages: `u`, destroyed at the boundary, **and** `C`, multiplied at the boundary, simultaneously. **Clinical:** A subgroup sits at a local minimum of its own comfort landscape. The minimum is deep because training and reward carved it, and narrow because it admits one vocabulary. Every direction out is uphill, which is why exhortation fails: it asks people to be uncomfortable on purpose, indefinitely, on someone else's account. **Signature:** Translation costs at every boundary crossing. Duplicate work discovered late. The product's seams sitting exactly

where the organisation chart's seams sit [Conway, 1968]. Cross-group meetings that require a preparatory meeting. **Measurement:** A boundary-crossing cost audit — elapsed time and person-hours for a decision requiring two groups against one. The ratio is the silo tax. Separately, measure `u_cross` (low) and per-boundary `C` (high); **the co-occurrence is the diagnostic.** The structural literature supplies the mechanism: bridging ties across a structural hole carry the information that within-cluster ties do not [Granovetter, 1973, Burt, 1992]. **Price:** Both terms move the wrong way at once, so the effect on the one-step condition `eta S > delta Q + C` compounds. **Fix: Do not push; build the more comfortable position next door.** Shared instrumentation, co-located cross-function work, rotation, boundary objects, the translator role. This is a safety property: a method working by removing resistance has no coercive failure mode, and its worst outcome is that someone declines to move. **Confusable with:** FM-04, FM-21. **Elsewhere:** the desire line and the launch budget [Curzi, 2026p].

FM-10 · The Amplifier Mistaken for a Source

Damages: `u`, overestimated; `N`, overestimated. **Clinical:** A hub with high throughput is read as a source of variety. It is redistributing, not generating. Preferential attachment reliably produces such hubs [Barabási and Albert, 1999], and the volume they carry is drawn from the same upstream sources the rest of the network already sees. The network double-counts. *No tail exponent is estimated here and none is claimed; the entry uses the mechanism and not the distribution, and the empirical generality of scale-free structure is contested* [Clauset et al., 2009, Broido and Clauset, 2019]. **Signature:** One node, feed, aggregator or senior figure appears in the provenance of most "independent" inputs. Removing that node removes most of the apparent variety. **Measurement:** A provenance audit. For each input the group treats as independent, trace to origin; compute the fraction whose paths share a common upstream node; recompute `u` with the duplicated inputs collapsed. **Price:** Whatever fraction of the roster reduces to a single upstream source collapses into one category by Identity 1. If eight of twelve inputs share an origin, the effective count falls from 12 to 5 before `u` is considered at all, and the ceiling falls from `ln 12 = 2.4849` to `ln 5 = 1.6094` — a **35.2 per cent** loss. **Fix:** Cut the hub out of the sampling path for at least one input stream. Note the structural corollary: networks of this kind are robust to random failure and **fragile to targeted attack on a hub** [Albert et al., 2000], so the amplifier is simultaneously the variety illusion and the single point of failure. **Confusable with:** FM-05, FM-15. **Elsewhere:** network resilience and the concentration instrument.

16.5.2 7.2 Class II — Conversion faults (`eta`)

The cheapest class to fix and the most often ignored, which is why conversion is the first move in almost every real system [Curzi, 2026b].

FM-11 · The Leaky Vessel

Damages: `eta -> 0`. Surplus is generated and none of it is stored. **Clinical:** Interaction happens at full strength and produces no durable structure. The conversation was real, the insight was real, and nothing outside the participants' memory changed. Because `S` is unaffected, every proxy for *activity* looks healthy while `Q` declines. **Signature:** "We had a great conversation" as a status report. Meetings with no artefact. A decision three attendees remember differently a month later. High engagement metrics against falling capability. The same problem re-solved annually by different people. **Measurement:** `eta = (attributable change in Q) / (input surplus)`. Concretely: over `k` periods, count the durable artefacts that (a) outlived the meeting, (b) are usable by a non-participant, and (c) trace to a specific interaction. Divide

by cumulative `S`. **Price:** With `S = 2.0`, `eta = 0.02` delivers `eta S = 0.040` per period; `eta = 0.35` delivers **0.700** — a factor of **17.5** with no change to headcount, geometry or effort. **Fix:** Install capture. In descending order of durability: the changed default; the built object; the trained person; the altered floor plan; the decision log; the written record. Ostrom's design principles for enduring common-pool institutions read, as engineering, as a list of conversion-raising devices [Ostrom, 1990, 2010]; *that reading is a transfer and is badged as one*, and the principles have since been audited meta-analytically [Cox et al., 2010]. **Confusable with:** FM-12, a Leaky Vessel that used to work; FM-19. **Elsewhere:** conversion as the first lever [Curzi, 2026b].

FM-12 · The Ritual That Became a Meeting

Damages: `eta`, by decay, under a preserved cadence. `C` unchanged or rising. **Clinical:** A recurring event was designed with a capture mechanism. The cadence survived; the capture did not. What remains is a periodic forcing function driving a filter with a broken input stage — full cost, no conversion. **This is the most common failure of *successful* practices**, because the cadence is what everyone can see and the capture is what nobody was assigned. **Signature:** A meeting nobody can cancel and nobody can justify. The original artefact — log, board, checklist — exists, is out of date, and is not referenced. **Measurement:** Pull the artefact. Check its last-modified date against the event's cadence. If the gap exceeds about three cadence-periods, conversion for that ritual is near zero. Confirm by asking three attendees to name the last decision the ritual produced. **Price:** The event's full `C` continues to be charged against `h` while its contribution to `eta S` approaches zero. In the ward diagnosis of Section 8.2, a handover with intact cadence and degraded capture runs at **46.04 per cent** of its own ceiling. **Fix:** Restore the capture step, or cancel the cadence. Do not keep the cadence for morale. **A ritual whose artefact is dead is a cost term with a good reputation. Confusable with:** FM-11, which never had capture; FM-19. **Elsewhere:** conversion mechanisms and the artefact ladder.

FM-13 · The Import With No Socket

Damages: `eta` for the imported category specifically; `u` raised nominally and never realised. **Clinical:** The organisation correctly diagnoses low `u`, correctly hires the outsider, and then provides no channel through which the outsider's difference can reach a decision. The new category exists on the roster and not in the causal path. **The gain is real in the survey and zero in the output. Signature:** A deliberate, well-justified hire from an unrelated lineage who has not changed any decision in two quarters. The outsider asked for input after the option set is closed. **Measurement:** A decision-provenance audit for the imported category: in the last ten consequential decisions, count those where the imported member's input entered *before* the option set closed. Zero or one is diagnostic. **Price:** The ceiling gain from `N = 4` to `N = 5` is **0.2231** nats. If the fifth category has no socket, the realised gain is 0.0000 nats and the salary is pure `C`. **Fix:** Build the socket before the import: a named decision gate the new category owns, a veto scoped narrowly enough to be usable, a standing agenda slot preceding option closure. **The sequencing rule of Section 4.10 applies — the socket is stabilisation and precedes growth. Confusable with:** FM-02, which is an import with no *difference*; this is a real difference with no *channel*. **Elsewhere:** capacity recruitment; Section 4.10.

FM-14 · The Sprayed Seed

Damages: `eta` for a diffusion intervention specifically. **Clinical:** A new practice is introduced by broad thin exposure — everyone hears once. Complex contagion requires *multiple reinforcing contacts* before adoption [Centola and Macy, 2007, Centola, 2010], so thin broadcast reaches everyone and converts nobody. **Clustered topology beats random long ties**, which

is the opposite of the reach intuition that produced the spraying. **Signature:** A launch with excellent awareness metrics and no adoption. Pilot sites chosen for geographic spread. "Everyone has been informed" as the implementation report. **Measurement:** Count *reinforcing* contacts, not contacts. For each non-adopter, how many of their frequent contacts have adopted? If the modal answer is one, the topology is wrong regardless of total exposure. **Price:** Adoption is threshold-driven and therefore discontinuous [Granovetter, 1978]; below threshold the return on exposure is near zero, which is why sprayed campaigns produce a flat curve rather than a slow one. **Fix:** Cluster first, do not spray. Saturate one dense neighbourhood well above local threshold and let the boundary do the work. Size the committed minority against the measured result: roughly **25 per cent** flipped an established convention in controlled online experiments [Centola et al., 2018] — for a 22-household block that is 6 households, for a 12-person group it is 3 people. **Badged: the 25 per cent figure is from controlled online experiments and field transfer is a modelling assumption**, and the general caution about transferring a single experimental result applies [Open Science Collaboration, 2015]. **Confusable with:** FM-11. **Elsewhere:** contagion, thresholds, and the critical boundary-cost multiplier.

FM-15 · The Founder Battery

Damages: `Q` stored in a person rather than in the system; `delta` becomes a step function. **Clinical:** One individual is the conversion mechanism, the archive and the maintenance regime. Measured `eta` is high and measured `delta` is low, and **both readings are attributes of a person, not of an institution. The parameters are correct right up until the day they are catastrophically wrong. Signature:** A named person in every consequential thread. Onboarding that consists of asking them. Nobody can state the escalation path without saying their name. **Measurement:** The absence test, run deliberately and at low stakes: the person takes two consecutive periods fully off-channel; measure `Q` before and after. The decline is the person-attributable share of `Q`. **This is the single most informative measurement a healthy organisation can afford to run** [Curzi, 2026i]. **Price:** A commons block that lost its founding steward saw commons-units fall from 47 to 38 over nine months of near-zero input: `delta = 0.0233` per month, half-life **29.35 months** [Curzi, 2026k]. The decay was gentle only because that block's stored order was physical; where `Q` is procedural, the same test typically returns a far larger number. **Fix:** Externalise before replacing. Stewards of record for every asset; **put the state in the medium, not in a coordinator. Confusable with:** FM-10, FM-19. **Elsewhere:** order without a centre; stigmergic media. ## 7.3 Class III — Retention faults (`delta` and `Q`)

FM-16 · The Sealed Vault

Damages: `delta -> 0` with `S = 0`. `Q` is perfectly preserved and never updated. **Clinical:** The system has excellent retention and no input. With `S = 0` and `C = 0` the trajectory is `Q_t = (1-delta)^t Q_0` — a decay so slow it reads as permanence. The vault is not failing; it is succeeding at the wrong thing. **A system with no forgetting cannot revise**, and the boundary case is worse than slow: at `delta = 0` exactly, the update map ceases to contract, no fixed point exists, and no mistake ever decays out of the stock [Curzi, 2026a]. **Signature:** Procedures whose rationale nobody can reconstruct. Reverence for the archive combined with no additions to it. Change proposals answered with precedent rather than argument. **Measurement:** Two readings. First, an input audit: over the last twelve periods, how many elements of `Q` were *added or amended*? Second, decay under zero input: if `delta` is low and additions are zero, the vault is sealed. **Price:** `delta = 0.01` gives a half-life of **68.97** periods and a time constant of 100; `delta = 0.005` gives a half-life of **138.28** periods. At those rates an obsolete rule outlives every person who might have questioned it. The half-life is `t_half = ln 2 / -ln(1-delta)`, and it is not `1/delta`: the naive rule of thumb is out by 72 per cent at `delta = 0.30` [Curzi, 2026]. **Fix:** Raise `delta` deliberately in the compartments where learning is wanted: sunset clauses, mandatory re-justification cycles, term limits, planned demolition. Keep `delta` low where durability is

wanted. **`delta` is not a single number for a whole institution; it is a field over its compartments, and treating it as one number is itself the fault.** **Confusable with:** FM-18, which is a refusal to shed *at cost*; the vault costs nothing and simply ossifies. **Elsewhere:** Curzi [2026a]; the half-life result [Curzi, 2026].

FM-17 · The Maintenance Trap

Damages: `h`, through the `delta Q` term. No parameter is individually abnormal. **Clinical:** `Q` has grown to the point where `delta Q` alone consumes the ceiling. From Identity 3, `h` is *decreasing in* `Q`. **Success consumes its own headroom.** This is the structural reason institutions ossify, and it is arithmetic rather than moral. **Signature:** An organisation that was building five years ago and is now entirely maintaining, with no decision anyone can point to that caused the change. Every new initiative crowded out by obligations to existing ones. **Measurement:** Compute the stall point `Q_stall = eta ln N / delta` at `C = 0` and compare to current `Q`. Within about 30 per cent of the stall point, the trap is already closing. **Price:** With `eta = 0.40` units per nat and `N = 12`, so `eta ln N = 0.9940`: at `delta = 0.02` the stall point is `Q = 49.70`; at `delta = 0.05` it is **19.88**; at `delta = 0.08`, **12.43**; at `delta = 0.12`, **8.28**. **A system with `delta = 0.12` cannot hold more than about eight units of order however well it interacts.** **Fix:** Four levers and only four: raise `eta`, raise `N`, cut `C`, or **accept lower `Q`**. The last is the one nobody considers, and shedding stored order to regain headroom is legitimate and underused — see FM-18 for the failure to make the move. **Confusable with:** FM-18, FM-20. **Elsewhere:** headroom and the four levers [Curzi, 2026b].

FM-18 · The Unshed `Q`

Damages: `Q` too high for the system's own `delta`; `h` negative as a consequence. **Clinical:** The organisation retains stored order it no longer uses, cannot maintain and will not retire. Every retained unit charges `delta` against headroom every period, forever. Unlike the Sealed Vault, the Unshed `Q` is *expensive*: the retained order actively consumes maintenance. **Signature:** An asset register nobody has pruned. Modules, documents, programmes or committees unused for a year that nobody will delete. The phrase "someone might still need it". **Measurement:** A usage audit with a hard window. For every element of `Q`, when was it last used by someone other than its owner? Sum `delta x Q_unused`. **That number is the annual rent on sentimentality, in the same units as the growth budget.** **Price:** In a research group worked elsewhere in this corpus, retiring 12 of 46 assets moved `delta Q` from `0.1120 x 46 = 5.154` to `0.075 x 34 = 2.550` assets per quarter — a **50.5 per cent** reduction in the maintenance charge, obtained by *removing* things. Shedding alone moved the one-step balance from **-3.387** to **-2.043** assets per quarter before any other lever was touched [Curzi, 2026k]. **Fix:** Sunset by default. Every element of `Q` carries an expiry and must be re-justified to survive it. Archive tiers marked explicitly unmaintained are honest and cheap; pretending to maintain is neither. **Confusable with:** FM-17 — the trap is the arithmetic, the Unshed `Q` is the refusal to act on it. **Elsewhere:** Section 4.4, Remedy 3.

FM-19 · The Invisible Infrastructure

Damages: `delta`, rising unobserved; then `Q`, discontinuously. **Clinical:** Infrastructure becomes visible only on breakdown [Star and Ruhleder, 1996]. The consequence for this law is structural: **maintenance is chronically underfunded because working infrastructure produces no observable signal**, so `delta` rises for years with no reading on any instrument, and the first reading is a failure. **Signature:** No maintenance line in the budget. No named owner for a system everyone depends on. A backup nobody has restored from. The question "who

maintains this?" producing a pause. **Measurement:** Invisible infrastructure cannot be measured by watching it; it must be *interrupted*. Scheduled, announced, low-stakes interruption: take the dependency off-line for one period and measure the drop in `Q`. **That number is the infrastructure's contribution and is the only honest basis for its budget. Price:** Undetected drift in `delta` is the most expensive error in the catalog because it compounds. A `delta` drifting from 0.05 to 0.15 moves the half-life of stored order from **13.51** periods to **4.27** — a threefold acceleration producing no symptom until the stock is already gone. **Fix:** Make the invisible visible on a schedule: interruption drills, maintenance dashboards that display *what did not break*, named stewards. Budget maintenance as a fixed fraction, not as a residual. The resilience-engineering literature makes the same argument about work-as-done against work-as-imagined [Hollnagel et al., 2006]. **Confusable with:** FM-15, FM-12. **Elsewhere:** decay measurement under a zero-input window [Curzi, 2026i].

FM-20 · The Negative Fixed Point

Damages: the model itself — `eta F(u*) < c*`, so `Q_inf = (eta F(u*) − c*)/delta < 0`. **Clinical:** CC-T4 gives the steady state. When conversion times realised surplus is below cost, the equilibrium is negative. **A negative fixed point means the regime is unsustainable or the linear model must be replaced — those are the only two readings, and the organisation is usually reporting a third: "transitional".** The theorem does not conceal this failure mode and neither may a diagnosis. **Signature:** A multi-year transition with no defined end state. Reserves drawn down on a schedule described as temporary. **Measurement:** Compute `Q_inf` explicitly at every survey. It is one division. If it is negative, or below current `Q` by a wide margin, the trajectory is heading there and the arrival date is computable from `tau = 1/delta`. **Price:** In an open-source case worked elsewhere in this corpus the measured pre-intervention fixed point was `Q_inf = 20.39` modules against a current `Q` of 63 — **the project was not stable, it was 43 modules above its own equilibrium and descending at -2.47 modules per month, with a computed zero-crossing at 25.5 months** [Curzi, 2026k]. **Fix:** Change a parameter or change the model. Both are legitimate; continuing to report the current regime as steady is not. If cost cannot fall below `eta F(u*)`, the system must shrink to a scale where it can, and shrinking on purpose is cheaper than shrinking by collapse. The deficit identity `D = 2.885390(c/eta − F) = −2.885390(delta/eta)Q_inf`, cross-checked to a residual of `8.88e−16` on its carried case, states the same shortfall in the corpus's frequency register [Curzi, 2026]. **Confusable with:** FM-17, FM-23. **Elsewhere:** the steady state and its breakdown [Curzi, 2026a].

16.5.3 7.4 Class IV — Dissipation faults (C)

FM-21 · The Quadratic Committee

Damages: `C`, with the number of pairs, against a logarithmically growing ceiling. **Clinical:** Communication paths grow as `n(n−1)/2` [Brooks, 1975] while the surplus ceiling grows as `ln N`. **Cost grows quadratically while surplus grows logarithmically**, so for any fixed per-channel cost there is a group size beyond which the group is strictly worse than a smaller one. Adding people to a late project makes it later. **Signature:** Meetings requiring pre-meetings. Decisions requiring more sign-offs than arguments. A group that has grown and whose output has not. **Measurement:** Tally `C` directly in hours or in proxy units per period, then compute the crossover: for per-channel cost `c0`, find the `n` at which `c0 x n(n−1)/2` exceeds `eta ln n`. **Price:** Channels per nat of ceiling: `n = 5` gives 10 channels, **6.21** per nat; `n = 12` gives 66, **26.56**; `n = 20` gives 190, **63.42**; `n = 30` gives 435, **127.90**; `n = 150` gives 11,175, **2,230.26**. With `eta = 1.0` and a per-channel cost of 0.01, headroom peaks near `n = 10` at `h = +1.853`, is still positive at `n = 20` at

`h = +1.096`, and **crosses zero at `n = 27`. Fix:** Nest, do not enlarge. A semilattice of overlapping clusters beats both the flat group and the strict tree, because overlap preserves cross-contact while cutting channels [Alexander, 1965]. A commons case worked elsewhere in this corpus cuts 231 channels to **55.5** — **24.03 per cent** — with four overlapping clusters of 5.5 households plus a delegate layer [Curzi, 2026k]. The same structural argument appears in governance as polycentricity [Ostrom, 2010]. **Confusable with:** FM-22, which is about *load* rather than *size*; FM-07. **Elsewhere:** the crossover and the optimum group size [Curzi, 2026b].

FM-22 · The Utilisation Cliff

Damages: `C`, non-linearly, as utilisation approaches one. **Clinical:** For a single-server queue with exponential arrivals and service, expected wait scales as `1/(1−rho_u)` [Kingman, 1962, Hopp and Spearman, 2011]. Near full utilisation the derivative explodes, so the last few per cent of capacity are bought at enormous and accelerating cost in delay. **Slack is not waste; it is a cost-control device**, and removing it is the most common way an efficiency programme destroys an institution [Cyert and March, 1963]. *The queueing premises are not measured in the worked cases; the transfer is badged as a modelling assumption.* **Signature:** Everything is scheduled and everything is late. Efficiency metrics at record highs alongside collapsing responsiveness. Nobody has unallocated time and nothing gets done. **Measurement:** Compute the utilisation `rho_u` for the binding resource — written `rho_u` here because the catalog already uses `rho` for the alignment correlation of FM-01, a collision in the source material: arrival rate over service rate. For people, `rho_u` is the fraction of available hours already committed; anything above about 0.85 sustained is in cliff territory. Little's Law, `L = lambda W`, closes the loop and lets a queue length reading be checked against a wait reading [Little, 1961]. **Price:** Wait as a multiple of service time: `rho_u = 0.50` gives **2.0x**; 0.80 gives **5.0x**; 0.90, **10.0x**; 0.95, **20.0x**; 0.98, **50.0x**. Moving a team from 80 to 95 per cent booked multiplies waiting by **4.0x**, and from 80 to 98 per cent by **10.0x**. Queue length rises from 4.0 to 19.0 to 49.0 over the same range. **Fix:** Design to a target `rho_u`, typically 0.70 to 0.85 for work with variable service times. Protect slack explicitly and name it, so it cannot be reallocated by an efficiency review. **Reduce variability before raising utilisation — variability, not utilisation, is what makes the cliff steep. Confusable with:** FM-21. **Elsewhere:** stock, flow and the cost of waiting.

FM-23 · The Unbooked Cost

Damages: `C`, underestimated; `h`, overestimated; every downstream decision, poisoned. **Clinical:** The cost ledger records the costs that have invoices. Commute, emotional labour, context-switching, compliance, waiting, and the repair of relationships damaged by the last decision all dissipate stored order and appear on no line. **A survey omitting them computes a headroom that does not exist and then spends it. Signature:** Repeated overruns with no identifiable cause. People visibly exhausted by a workload the ledger describes as moderate. **Measurement:** Time-and-motion on the *whole* period, not on the tasks: everything not converting surplus into stored order is `C`. **The cost observable is defined as a residual for exactly this reason** — measure total capacity, measure productive capacity, take the difference, rather than enumerating cost items and hoping the list is complete [Curzi, 2026i]. **Price:** In a research group worked elsewhere in this corpus, coordination ran at **17.4 per cent** of 6,240 person-hours per quarter — 1,087 hours — which at the measured yield rate of 0.00147936 assets per person-hour priced `C` at **1.608** assets per quarter against a ceiling `eta ln N = 8.105`. **Omitting it would have overstated `h` by 120 per cent of its true value of 1.343** [Curzi, 2026k]. **Fix:** Book `C` as a residual, not as a list. Then validate: the survey's predicted one-step change must reproduce the observed one. Four cases reproduce observed decline to within 0.12 units [Curzi, 2026k], and that agreement is what indicates the cost ledger is complete. **Confusable with:** FM-29 — both are survey faults; this one under-reports `C`, that one mis-specifies `Q`. **Elsewhere:** the cost observable and the denomination rule [Curzi, 2026i].

FM-24 · The Abrupt Threshold

Damages: C, at every zone boundary; secondarily u, by suppressing crossing. **Clinical:** Two zones of different character meet with no transition, so every crossing costs a full reorientation — acoustic, social, visual, procedural. The threshold is a device in its own right, an impedance match between zones [Alexander et al., 1977], and its absence is a tax collected on every crossing, forever. **Signature:** A door people hesitate at. An open-plan floor where nobody uses the collaboration zone because entering it is a public act. **Measurement:** Count crossings per hour against the crossings the function requires, by direct observation [Whyte, 1980]. **Desire lines are the measured gradient; where the desire line does not match the designed line, the threshold is wrong** [Jacobs, 1961]. Personal distance and its zones are the underlying regularity [Hall, 1966]. **Price:** Context-switching is a cost term and it compounds with FM-22: at `rho_u = 0.90` each additional interruption costs **10x** its own service time in downstream delay. **Fix:** Build the transition — porch, vestibule, airlock, buffer zone, staged announcement. In organisations the equivalent is the staged decision: a proposal with a public-but-uncommitted phase. Depth, a sightline into the next zone, and a place to pause are the three specifications. **Confusable with:** FM-26 — both make crossing costly; one is geometric, one is acoustic. **Elsewhere:** thresholds and transitions in the built environment.

16.5.4 7.5 Class V — Room and instrument faults

Separated because their fixes are capital purchases with known lead times, and because they are among the cheapest interventions available per unit of effect.

FM-25 · The Head of Table

Damages: u, structurally and continuously. **Clinical:** A table with a head assigns a direction to the group. Everyone below it has a function, and that function is agreement. **Agreement between aligned instruments is echo, not evidence.** The geometry does this without anyone enforcing it, which is precisely why it is durable: **furniture governs continuously and costs nothing to administer. Signature:** Speaking order tracks seating order. Turn-taking Gini high. Contributions addressed to one seat rather than to the room. Sightlines that do not close. **Measurement:** Two instruments. (1) Turn-taking equality: record speaking time per participant for three meetings and compute the Gini coefficient — collective intelligence correlates with *equality* of conversational turn-taking [Woolley et al., 2010]. (2) Sightline reciprocity: for all pairs, can each see the other's face without moving? Report the failing fraction. **Price:** FM-01 prices the result. At `N = 12`, driving u from 0.30 to 0.05 by geometry alone moves S from **1.4586** to **0.4383** — a **70.0 per cent** loss of realised surplus, purchased with a table shape. **Fix:** Round or polygonal geometry with full sightline reciprocity. Size the table by channel count [Brooks, 1975] and by acoustics jointly. Add protocol: written positions first, round-robin, and the most junior speaking first on matters of rank. **Confusable with:** FM-01, FM-26. **The test: if the fault persists when the group meets elsewhere, it is not the table. Elsewhere:** table geometry and turn-taking equality.

FM-26 · The Reverberant Room

Damages: u, by suppressing the quiet voice; C, through the raised vocal effort a rising noise floor compels. **Clinical:** In a room with long reverberation, speech from previous syllables masks the current one. Signal-to-noise for the quiet speaker falls below intelligibility

while the loud speaker remains intelligible. **The loudest speaker wins a reverberant room by physics, not by dominance** — and the compensating rise in everyone's level raises the noise floor, which raises everyone's level again. *The acoustics is a derivation; the step from reverberation time to non-parallelism is a modelling assumption* (Section 3.8). **Signature:** People leaning in. Repetition requests. Quiet participants who are articulate elsewhere. Meetings that exhaust people disproportionately to their content. **Measurement:** Sabine's relation, `RT60 = 0.161 V / A` in metric units, where `V` is volume in cubic metres and `A` is total absorption in square-metre sabins, `A = sum(S_i x alpha_i)` plus occupants at roughly 0.45 each at 500 Hz [Sabine, 1922, Schroeder, 1965]. Measure or compute both, occupied and empty, and add a background noise reading. **Price:** A 16.0 by 9.5 by 4.2 metre reading room, `V = 638.40` cubic metres, with hard finishes has `A = 40.750` and `RT60 = 2.5223` seconds occupied, **3.1480 seconds** empty. Treated — 90 square metres of ceiling absorber at `alpha = 0.90`, 40 of wall panel at 0.85, 100 of heavy carpet at 0.25 — `A` rises to **173.150** and `RT60` falls to **0.5936 seconds**, a **4.25x** reduction. Ceiling treatment alone gives 0.8699 seconds; ceiling plus walls gives 0.6845. **Fix:** Absorption, in order of cost-effectiveness: ceiling first, as the largest uninterrupted surface and the first reflection path; then the wall opposite the primary talker; then the floor. The turn-taking consequence is arithmetic in the tail: the gap for a previous talker to fall 15 dB is `t_15 = RT60/4`, giving 488 ms before treatment and 98 ms after [Curzi, 2026e]. **Acoustic treatment is the cheapest room-scale intervention available. Confusable with:** FM-25, FM-01. **Acoustic faults vanish when the same group meets in a different room; that is the discriminating test. Elsewhere:** passive geometry and the reverberant tail [Curzi, 2026e].

FM-27 · The Single Setpoint

Damages: `u`, by narrowing who can occupy the space; `C`, through avoidance. **Clinical:** The environment is optimised to one value — one temperature, one illuminance, one noise floor, one furniture height — chosen as the population mean. Individual tolerance varies widely across the six thermal variables and across light sensitivity, so **an optimised setpoint is correct for the mean occupant and wrong for most actual occupants. Adjustability outperforms optimisation because it raises the range of people who can be present at all**, and presence is prior to every other parameter. **Signature:** Personal heaters and fans in a climate-controlled building. Blinds permanently down. Complaints in both directions about the same setting. **Measurement:** Occupancy mapping by seat over two weeks, applying the direct-observation method indoors [Whyte, 1980]. Persistently unoccupied positions in an otherwise-used room are an environmental reading, not a furniture preference. Cross-check with a two-sided complaint tally. **Price:** Any occupant excluded is a category removed from `N`, or a member removed from a category, priced by Identity 1. Excluding one of five staffed categories moves the ceiling from `ln 5 = 1.6094` to `ln 4 = 1.3863`, a **13.9 per cent** loss — for a thermostat. **Fix:** Provide gradient and control rather than a single value: a range of thermal zones, task lighting alongside ambient, seating variety, sun and shade both available. **A well-known 1984 study reporting faster post-surgical recovery with a window view of trees is often cited here; this corpus's bibliography contains no entry for it, so it is named without a citation and is not relied on.** The argument it is usually used to make — provide the good condition rather than average it away — stands on the arithmetic above without it. **Confusable with:** FM-24. **Elsewhere:** environmental range against environmental optimisation.

16.5.5 7.6 Class VI — Method faults (the analyst's own)

These damage nothing in the system and everything in the knowledge of it. **Listed last; checked first, because a method fault invalidates every other reading.**

FM-28 · Goodhart Capture

Damages: the measurement, and then `Q`, once the measurement drives behaviour. **Clinical:** A measure that becomes a target ceases to be a measure [Goodhart, 1975, Campbell, 1979]. Once the proxy for stored order is rewarded, effort flows to the proxy rather than to `Q`, and the correlation that justified the proxy is destroyed by the act of using it. **The same reflexivity the framework asserts about the systems it models applies to its own instrument** [Lucas, 1976]. **Signature:** A metric improving faster than any plausible underlying change. Improvement concentrated exactly at the measurement boundary. Practitioners who can explain how the number is produced but not what it means. **Measurement:** Hold out a second, unrewarded proxy for the same `Q` and track both. Divergence between rewarded and unrewarded proxies is the capture signal. Rotate the unrewarded proxy periodically so it does not become rewarded by attention. **Price: Not directly computable in nats.** The damage is that the `h` estimate is now fiction, so every subsequent number in this catalog is fiction. **Fix:** Never reward the proxy used for measurement. Keep an audit proxy that is unpublished, unincentivised and rotated. Prefer proxies expensive to fake and cheap to verify: physical state, third-party usability, independent re-execution. **Confusable with:** FM-29. **Elsewhere:** Section 4.16; the reporting-line defence [Curzi, 2026i].

FM-29 · The Order Proxy That Was Really the Input

Damages: the entire estimation. The fitted model is a tautology. **Clinical: This is the most common measurement error in the framework** [Curzi, 2026i]. If the chosen order observable is a repackaging of the interaction observable, then `Y_{t+1} = a Y_t + b F(U_t) − C + e` fits beautifully and says nothing: interaction has been measured twice and the correlation called a law. **Signature:** A proxy defined in terms of meetings, messages, commits, attendance or engagement. Extremely good fit on the first attempt. The proxy falling to zero the instant interaction pauses. **Measurement:** The lag test. **Genuine stored order *persists* when input stops — that is what "stored" means, and it is why `delta` is measurable at all.** Pause input for one to three periods and watch the proxy. If it tracks the pause with no lag, it is the interaction observable. If it decays with a measurable time constant, it is stored order, and `delta` has been measured as a bonus. **Price:** Every parameter downstream is unidentified. **There is no partial credit. Fix:** Choose the proxy as a *state of the world outside the interaction*: modules passing tests on the current toolchain; assets a non-author can independently re-run; patrons who return without being invited; shared capabilities usable without asking anyone. The test is applied explicitly to four cases elsewhere in this corpus and rejects the naive proxy in each [Curzi, 2026k]. **Confusable with:** FM-28. **Elsewhere:** the four repackaging tests [Curzi, 2026i].

FM-30 · The Unbadged Postulate

Damages: the warrant. Nothing in the system; everything in the standing of the claims. **Clinical:** A modelling choice, an analogy, or a transfer from a laboratory result is presented with the same confidence as a theorem. **An unbadged postulate is not a lie; it is an unpriced liability, and it comes due when someone competent checks.** The programme-level description of what a badged postulate is doing — a protective belt around a hard core, adjusted deliberately and in the open — is Lakatos [1970]. **Signature:** No distinction in the prose between "the theorem gives" and "we assume". A controlled online experiment cited as a field constant. A design axiom invoked as though derived. **Measurement:** Read the document with two highlighter colours: proven, and assumed. **If the assumed colour does not appear, the reader has not found it; they have failed to look. Price:** Nothing quantitative. **A single discovered unbadged claim discredits the badged ones, which is a very poor trade for the confidence gained. Fix:** Badge in the same sentence. "The affine effective count is a design axiom — chosen, not derived." "The 25 per cent figure is from controlled online experiments; field transfer is a modelling assumption." "That this institution is described by a single scalar `Q` is a postulate."

The three-drawer classification — established mathematics, established psychology, declared postulates, with every protocol stated to run on the first two and steer by the third — is the standing instrument [Curzi, 2026]. **Confusable with:** FM-31. **Elsewhere:** Section 3.6; Curzi [2026].

FM-31 · The Borrowed Ceiling

Damages: the plan. It requests `r_t > h_t`, which CC-T1 forbids. **Clinical:** The design assumes synergy beyond additivity — that combining two sources of difference yields more than their sum, or that a good enough team escapes its own headroom. The decomposition gives an upper bound of exactly the sum: `f(u) <= f(u_cross) + f(u_div)`. **No synergy beyond additivity is promised, and honest design does not assume one.** Separately, the triangle inequality fails for geometric triples [Curzi, 2026], so three-party reasoning cannot be chained metrically at all. **Signature:** A target exceeding `r_max`. A plan justified by two things multiplying. Growth projections with no headroom calculation anywhere in the document. **Measurement:** Compute `h_t` and compare to the requested `r_t`. CC-T2: the increment is feasible only if `delta Q_t + C_t + r_t <= eta ln N`. **This is one line of arithmetic and it is not optional. Price:** At `N = 8` with `u_cross = 0.30` and `u_div = 0.25`, the honest bound is `1.1314 <= 1.5790 <= 2.1430`. A plan assuming multiplication rather than the bound over-promises by up to **36 per cent** relative to the true value even inside the permitted range. **Fix:** Plan to the bound. State `r_t <= h_t` as a hard constraint in the proposal document and put the computed `h` next to the requested `r` on the same page. **Confusable with:** FM-07, FM-30. **Elsewhere:** Sections 4.6 and 4.7.

FM-32 · The Consent Bypass

Damages: nothing in the law. **It is a refusal, not a parameter. Clinical:** An intervention raises a parameter by removing someone's ability to decline — mandatory disclosure of dissent, forced rotation, attendance that cannot be refused, "voluntary" participation with recorded non-participation. **This is outside the procedure entirely** (P-E, Section 3.6), and the reason is engineering as much as ethical: **coerced non-parallelism is not non-parallelism.** A person compelled to voice a difference voices the difference the compeller expects, which returns the system to FM-01 by a longer road. The classification of people by a measure that then acts on them is a documented mechanism, not a speculation [Hacking, 1995]. **Signature:** A protocol with no opt-out. Non-participation that is visible and therefore costly. Consent obtained from a group rather than from each person. A benefit argued at the level of the collective whose cost falls on identifiable individuals. **Measurement:** For every protocol in the intervention design, one written question: *what happens to a person who declines?* If the answer includes any adverse consequence, the protocol is out of scope. **Price: Not applicable. This is a boundary, not a trade-off, and it is not priced because it is not for sale. Fix:** Redesign so the intervention works by removing resistance rather than by applying force — build the more comfortable position next door, and the worst case is that someone declines to move. **Confusable with:** nothing. **If it is unclear whether an intervention is a Consent Bypass, it is. Elsewhere:** Section 4.12.

16.5.6 7.7 Faults of diagnosis itself

Seven faults attach to the diagnostic act rather than to the system. **They are labelled DF-1 to DF-7 here; the source labels them D-1 to D-7.**

DF-1 · Naming before measuring. The catalog's names are vivid, and vividness produces premature confidence. A named fault with no measurement behind it is a story. *Signature:*

the diagnosis arrived within ten minutes of arriving on site. *Fix:* the survey is an entry condition, not an alternative (Appendix C, CS-3).

DF-2 · The single-fault reflex. Looking for *the* problem. Systems fail in clusters, and fixing one member of a cluster while leaving the others returns the system to its previous state within a few time constants. *Fix:* require at least three named faults, or an explicit written statement of why fewer suffice.

DF-3 · Fixing the loudest fault. The most visible fault is rarely the binding one. In a library case worked elsewhere in this corpus the loud fault was conflict between patrons; the binding constraint was non-parallelism, and the instrument was ceiling absorption [Curzi, 2026k]. *Fix:* rank by computed effect on the one-step condition, not by salience.

DF-4 · The one-instance check. Verifying a fix at one site, one meeting, one week, and generalising. **A single successful instance is consistent with almost every hypothesis.** *Fix:* prespecify the number of instances and the checkpoints before running the pilot [Nosek et al., 2018].

DF-5 · Treating `delta` as a scalar for the whole system. Decay is a field over compartments. Documentation, relationships, physical plant, skills and norms decay at wildly different rates, and an average is a number describing none of them. *Fix:* estimate `delta` per compartment and report the vector.

DF-6 · Diagnosing outside the model's range. The linear storage law has no saturation term. When a predicted `Q_inf` exceeds a hard physical or organisational cap, the model is being used outside its range and must be replaced there. *Fix:* **always state the cap alongside the fixed point, and say which one binds.** Section 8.2 does exactly this, with a predicted fixed point of 71.46 against a 34-bed census.

DF-7 · Mistaking the survey for the intervention. Measurement changes behaviour, usually temporarily and favourably; a survey period showing improvement has told the analyst about the survey. *Fix:* the shuffled-input control and a matched comparison — noting that the Hawthorne pattern is invoked far more casually than the evidence supports [McCambridge et al., 2014], so its name is not an explanation. # 8 Results: two systems worked end to end

Both systems are constructed. The first is a declared composite whose audit series is a simulation with declared generating parameters; the second is a constructed clinical case. **The arithmetic in both is real, computed and reproducible; the organisations are not.** Neither may be cited as evidence about any real review office or any real hospital ward. What the two cases demonstrate is that the procedure runs, that it reports its own refusals, and that it reports its own departures from the model's valid range.

16.5.7 8.1 A development review office: a refusal, and the repair that made the request fundable

The system is a county development review office carried forward from the survey and preregistration protocol [Curzi, 2026i]. Five review disciplines, a standing rules manual as the order proxy, a weekly period.

8.1.1 Gate 0 — admissibility

Check	Result
Survey complete	Yes
Repackaging tests, all four	All PASS
Preregistration frozen and evaluated	Yes
FR1, admissible parameter box	PASS
FR2, baseline dominance	**FAIL** — matched AR(1) beaten by 2.8 per cent against a required 10 per cent

Check	Result
FR3, independent-parameter agreement	PASS

Evidence status: accounting frame only. `h` may be used as a budget constraint and a preregistered, capped, reversible pilot is permitted. **No claim may be made that the model predicts the pilot's outcome** (Sections 4.2 and 4.13). This is the status established in Curzi [2026i], whose confirmatory fit failed its own second falsification rule, and it governs every statement in this subsection.

8.1.2 Gate 1 — the gauge as surveyed

```
eta   = 0.5015       ln N = 1.6094       eta ln N = 0.8071
delta = 0.003090     Q    = 113          delta Q  = 0.3492
                                         C        = 0.6300
------------------------------------------------------------------
h = 0.8071 - 0.3492 - 0.6300 = -0.1720            REGIME: h < 0

mean of F over the period record = 1.1498 nats
step = 0.5015 x 1.1498 - 0.3492 - 0.6300 = -0.4025 rules/week
realisation deficit = h - step = 0.2305 rules/week
capacity u = 0.6972          mean realised U = 0.6019
```

A notational warning, stated here because the source material does not. The surplus figure used in the step is **the mean of `F` over the period record**, 1.1498 nats. It is **not** `F` evaluated at the mean realised `U`, which is `F(0.6019) = ln(1 + 4 x 0.6019) = 1.2260` nats. The two differ by 0.0762 nats and **only the first is the quantity that accumulates.** The source writes the first as "realised `F(U)`", a notation that reads as the second; Section 9.4 records that as a defect. Every occurrence of 1.1498 in this paper is a mean of `F`; every occurrence of 1.2260 or 1.2262 is an `F` of a mean, and the two are never interchanged.

8.1.3 The brief, and the refusal

The office director's brief was to rebuild the standing review manual at roughly thirty rules a year — that is, `r = 0.55` rules per week.

```
r = 0.55     h = -0.1720     r <= h ?   NO
```

Refused, in writing, with the number. The written refusal is reproduced because its form matters as much as its content:

> At present parameters the office has headroom of `-0.17` rules per week. This is the best case: it assumes a perfectly non-parallel board every week. It means that dissipation, at 0.63 rules per week, exceeds what conversion can deliver at the ceiling net of maintenance, which is 0.46. No arrangement of who sits on the board, in what order they speak, or how vigorously they disagree can prevent the manual from declining. The requested increment of 0.55 is not merely unfunded; the account is overdrawn by 0.17 before any increment is requested. The manual is losing 0.40 rules per week — about 21 a year — and the fitted steady state is negative, which the model reports as an unsustainable regime rather than as a small number. What follows is what it would take to make the request fundable.

That paragraph is the deliverable. Not the intervention — the refusal, with the arithmetic that makes it a fact rather than an opinion.

8.1.4 Phase 1 — stop the leak

Raise `eta`: 0.5015 to 1.05 rules per nat. The office has no capture discipline: rules reach the manual when a reviewer remembers to draft one, at a median of five weeks, and the attribution audit found 10 of 26 new rules with no documented board origin. The intervention is three administrative changes — a designated scribe at every board session, a one-page rule template completed in the session, and a ten-working-day publication requirement with a named owner per rule. **`eta` starts near the floor, which is exactly the condition under which a large multiple is available**, and it is why this lever ranks first. The office's own natural experiment supports the magnitude: in the two years before the previous code administrator retired, when a scribe role informally existed, attributable conversion ran at roughly twice the current rate.

```
headroom gain = (1.05 − 0.5015) x 1.6094 = +0.8828 rules/week
cost = 0.4 FTE, redeployed from an existing analyst post
```

Cut `c`: 0.63 to 0.42 rules per week.

line	before	after	how
Appeal reversals striking a live rule	0.14	0.14	unchanged; structural
Rules orphaned by departure	0.11	0.07	deputise a second qualified applier per rule at adoption
Rework: items re-reviewed after disposition	0.16	0.05	close the re-review loop; disposition final absent new information
Contradiction maintenance	0.09	0.09	unchanged
Attention tax	0.13	0.07	publish the decision summary once, to one place, instead of five partial circulations
total	**0.63**	**0.42**	

```
headroom gain = 0.63 − 0.42 = +0.2100 rules/week
cost = one policy decision and three weeks of process work; no recurring cost
```

The rework line is cut and the appeal line is not. The appeal process is where the environmental discipline's objections are realised; cutting it would cut `C` and cut `F(U)` at once. That is the second caution of Section 4.9, and it is explicitly declined here.

8.1.5 Phase 2 — correct the order observable

Six rules removed by the liveness audit and two by the applier audit, bringing the register count of 121 to `Q = 113`. **This changed the number and not the world** and is recorded as a correction, not a shed (Appendix A, IS-15).

8.1.6 Phase 3 — the shed, computed and declined

```
proposed shed  q = 14 further rules, judged low-value but still live
headroom gain    = delta x q = 0.003090 x 14 = +0.0433 rules/week
```

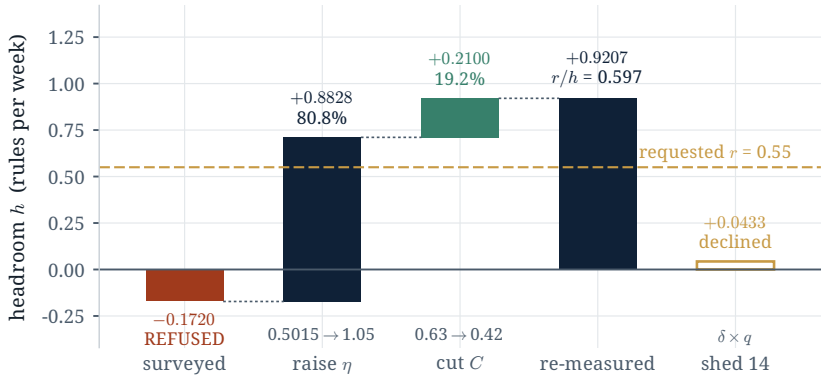


Figure 16.2: **The refusal and the repair, in one subtraction each.** Headroom moves from `−0.1720` to `+0.9207` rules per week under two administrative levers: raising conversion from 0.5015 to 1.05 rules per nat supplies `+0.8828` (80.8 per cent of the swing) and cutting dissipation from 0.63 to 0.42 supplies `+0.2100` (19.2 per cent). The shed of fourteen rules was computed at `delta x q = +0.0433` and declined; it is shown outside the running total. The requested increment of 0.55 rules per week is refused against the opening headroom and admitted against the closing headroom at a reserve ratio of 0.597. Modelled from the parameters of Sections 8.1.2 to 8.1.7; generated by `figs/src/CCP-016_fig01.py`.

Against the conversion lever's `+0.8828`, the shed returns **one twentieth** as much, and it is not free: fourteen rules are fourteen pieces of somebody's work, and their removal costs trust that does not appear in the arithmetic. **Declined on the arithmetic.**

For contrast, and because this is the finding that generalises: in a system with `delta = 0.05` per period the same fourteen-unit shed would return `0.05 x 14 = 0.70` — **a factor of sixteen more**, and comfortably competitive with the conversion lever. **Shedding is decisive in fast-decaying systems and nearly worthless in durable ones, and the multiplication that distinguishes them takes five seconds.** The two ratios in this paragraph are different comparisons and are not to be conflated: twenty is the shed against the conversion lever *in this system*; sixteen is the same shed *at two different decay rates*.

8.1.7 Phase 5 — re-measure

```
eta   = 1.05          eta ln N = 1.05 x 1.6094 = 1.6899
delta = 0.003090      Q = 113   delta Q        = 0.3492
                                C              = 0.4200
------------------------------------------------------------------
h = 1.6899 − 0.3492 − 0.4200 = +0.9207            REGIME: h > 0
```

The sign changed without adding a person, a category, or any capital expenditure. Two administrative levers, applied in the order the ranking prescribes, moved headroom by 1.093 rules per week — from `−0.172` to `+0.921`.

lever	gain in `h`	share of the swing
raise `eta`, 0.5015 to 1.05	+0.8828	80.8 per cent
cut `C`, 0.63 to 0.42	+0.2100	19.2 per cent
shed 14 rules (declined)	+0.0433	—

8.1.8 Gate 2 — feasibility of the requested increment

```
r = 0.55   <=   h = 0.9207                    [ inside the budget ]
reserve:   r/h = 0.597                        [ inside the recommended 0.7 ]

S_req = (delta Q + C + r)/eta = (0.3492 + 0.4200 + 0.5500)/1.05
      = 1.3192 / 1.05 = 1.2564 nats

S_req = 1.2564  <=  ln N = 1.6094             ->  NO Failure A

U_req = (exp(1.2564) - 1)/4 = 0.6281

U_req = 0.6281  <=  capacity u = 0.6972       ->  NO Failure B
U_req = 0.6281  >   mean realised U = 0.6019  ->  REALISATION WORK REQUIRED
```

The concavity correction, applied. `U_req = 0.6281` is the geometry a board would need if it ran at the same `U` every week. It does not. The office's sessions are a mixture: full-board sessions at a mean `U` of 0.7940 and reduced sessions at a mean `U` of 0.2071. What accumulates over many periods is the sum of `eta F(U_t)`, so the operative quantity is the **mean of `F`**, and because `F` is strictly concave (CC-T6), Jensen's inequality gives `mean F(U_t) <= F(mean U_t)` with equality only if `U` is constant.

For the office as surveyed the penalty is 0.0761 nats: a mean of `F` of 1.1501 against an `F` of the mean `U` of 1.2262, a loss of 6.2 per cent of the surplus purely to variability. (The mixture reconstruction returns 1.1501 and 0.6021 where the direct period record returns 1.1498 and 0.6019; the two agree to within rounding, which is the consistency check that licenses using the mixture to project. `F` of the *directly measured* mean `U` is `F(0.6019) = 1.2260`.)

So the requirement is stated on the mean of `F` and not on the mean of `U`:

```
required:              mean F(U_t) >= S_req = 1.2564 nats
achieved as surveyed:  mean F              = 1.1498 nats   -> SHORT by 0.1066 nats
```

The increment is inside the budget and inside the office's geometric capacity, and it is **not** inside what the board currently realises. At the surveyed geometry,

```
step = 1.05 x 1.1498 - 0.3492 - 0.4200 = 1.2073 - 0.7692 = +0.4381 rules/week
```

which is positive — a swing of 0.841 from the surveyed `-0.4025` — and short of the promised 0.55.

8.1.9 Phase 4 — realise

The cause was identified in the survey: **32.7 per cent of board sessions ran with fewer than the full set of disciplines recording a disposition**, and those sessions have a mean `U` of 0.2071 against 0.7940 for full sessions. The lever is a scheduling and protocol change: any item touching more than one discipline requires a recorded disposition from every discipline, obtained asynchronously and in writing if a lead cannot attend — **which also breaks the cascade, because dispositions are recorded before discussion** (FM-03).

low-participation share	mean `U`	mean `F`	`F(mean U)`	concavity penalty	step (rules/week)	`r = 0.55` met?
0.400	0.5593	1.0888	1.1747	0.0859	0.3740	no
0.327 (as surveyed)	0.6021	1.1501	1.2262	0.0761	0.4385	no
0.300	0.6179	1.1728	1.2447	0.0718	0.4623	no
0.250	0.6473	1.2149	1.2779	0.0631	0.5064	no

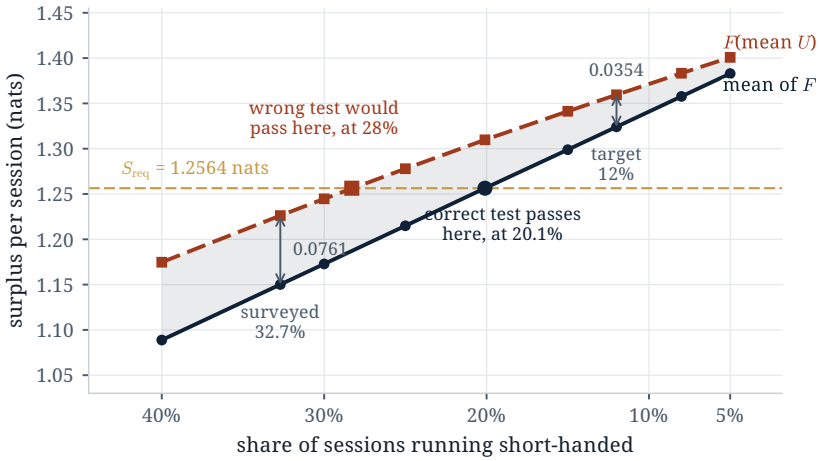


Figure 16.3: **The two curves that must never be interchanged.** The mean of `F` over the session record — the quantity that accumulates — sits strictly below `F` evaluated at the mean `U` at every low-participation share, and the gap is the concavity penalty. At the surveyed share of 32.7 per cent the mean of `F` is 1.1501 nats against an `F` of the mean of 1.2262, a penalty of 0.0761 nats or 6.2 per cent of the surplus; at the 12 per cent target the penalty falls to 0.0354. **A feasibility test run against the upper curve passes at a share of about 28 per cent; run against the correct lower curve it passes only below 20.1 per cent.** Modelled from the mixture of full-board sessions at mean `U` 0.7940 and reduced sessions at mean `U` 0.2071, tabulated in Section 8.1.9; generated by `figs/src/CCP-016_fig02.py`.

low-participation share	mean `U`	mean `F`	`F(mean U)`	concavity penalty	step (rules/week)	`r = 0.55` met?
0.201 (break-even)	0.6763	1.2564	1.3098	0.0534	0.5500	exactly
0.200	0.6766	1.2569	1.3101	0.0532	0.5505	yes
0.150	0.7060	1.2989	1.3413	0.0424	0.5947	yes
0.120 (target)	**0.7236**	**1.3241**	**1.3595**	**0.0354**	**0.6211**	**yes**
0.080	0.7470	1.3577	1.3833	0.0256	0.6565	yes
0.050	0.7647	1.3829	1.4008	0.0179	0.6829	yes

Read the concavity-penalty column down the table. Regularising the geometry pays twice: once by raising the mean, and once by shrinking the penalty the variability was charging. At the surveyed share the office loses 0.0761 nats to variability alone; at the target it loses 0.0354. **That second gain — 0.0407 nats, worth `1.05 x 0.0407 = 0.0427` rules per week — is invisible to anyone reasoning about the average board rather than about the distribution of boards.**

```
step at target = 1.05 x 1.3241 − 0.3492 − 0.4200 = 1.3903 − 0.7692 = +0.6211 rules/week
               = 32.3 rules per year
margin over the promise r = 0.55:  +0.0711 rules/week
```

The target is a low-participation share of 12 per cent; the remaining 12 per cent are genuinely single-discipline items where a full board is not warranted and forcing one would be pure `C`.

This lever raised `step` from 0.4381 to 0.6211 and raised `h` by exactly nothing. `h` was 0.9207 before and is 0.9207 after. This is Section 4.8 in operation: **realisation collects headroom that already existed.** The realisation deficit falls from 0.4826 to 0.2996 and is not zero, because the board still runs below the geometric ceiling and always will.

8.1.10 Gate 3 revisited — the category lever, computed and declined

The director's alternative proposal was to add a sixth discipline, public health, at 1.0 FTE.

```
ceiling gain   = eta (ln 6 − ln 5) = 1.05 x 0.1823 = +0.1914 rules/week
added paths    = 15 − 10 = 5
per-path cost  = 0.34 x 0.42 / 10 = 0.0143 rules/week per path
added cost     = 5 x 0.0143        = +0.0714 rules/week
net effect on h = +0.1914 − 0.0714 = +0.1200 rules/week
```

Two scenarios were run for the effect on the surplus term, because the candidate's distinctness matters more than its existence:

candidate	mean pairwise `u`	affine `N_eff`	`F(U)`	gain in `step`	net of path cost
genuinely distinct public-health reviewer	0.7352	4.676	1.5424	+0.2292	+0.1578
a second environmental reviewer, relabelled	0.6889	4.444	1.4917	+0.1759	+0.1045

Both are positive here, and the gap between them — a 34 per cent reduction in the benefit — is the price of relabelling. In a system with a higher per-path cost the relabelled candidate goes negative while the distinct one stays positive, which is why the merge test against the probe set is required *before* any hire and not after.

A correction to the source's comparison, made rather than inherited. The source sets the conversion lever's headroom gain of `+0.8828` over 0.4 FTE against the sixth discipline's `+0.1578` over 1.0 FTE and reports the conversion lever as **fourteen times** better per full-time equivalent. **Those two figures are not on the same basis.** `+0.8828` is a gain in `h`; `+0.1578` is a gain in `step` net of path cost. On a like-for-like basis the ratio is:

basis	conversion lever, per FTE	sixth discipline, per FTE	ratio
gain in `h`	0.8828 / 0.4 = **2.207**	0.1200 / 1.0 = **0.120**	**18.4**
gain in `step`, net of path cost	(1.05 - 0.5015) x 1.1498 / 0.4 = 0.6307 / 0.4 = **1.577**	0.1578 / 1.0 = **0.158**	**10.0**
the source's mixed comparison	2.207 (an `h` gain)	0.158 (a `step` gain)	14.0

The qualitative conclusion is unchanged and is strengthened: on either consistent basis the conversion lever returns an order of magnitude more per full-time equivalent, and the source's "fourteen times" is a hybrid of the two. This is recorded as a defect in Section 9.4 rather than silently corrected in the text.

Declined for now and deferred to Phase 8 — not because it is wrong, but because it is eighth in the sequence and Phases 1 to 6 have not yet been observed to work.

8.1.11 Gates 4 to 7 — sequence, reversibility, consent, pilot

```
Phase 1  stop the leak     eta 0.5015 -> 1.05 ; C 0.63 -> 0.42   h: −0.172 -> +0.921
```

```
Phase 2  correct the proxy done in the survey                    Q: 121 -> 113
Phase 3  shed              declined on delta x q = 0.0433        h: unchanged
Phase 4  realise           low-participation share 0.327 -> 0.12 step: 0.438 -> 0.621
Phase 5  re-measure h      h = +0.9207
Phase 6  commit            r = 0.55 (reserve r/h = 0.597)
Phase 7  raise capacity    deferred; review at 26 weeks
Phase 8  raise N           deferred; merge test required first
```

Gate 5 — reversibility.

RV1 what returns	The scribe role reverts to the analyst's prior duties. The template and the ten-day publication requirement are rescinded by a single memo. The full-disposition rule reverts to the prior attendance-based practice. The re-review loop is reinstated by rescinding the same memo.
RV2 what does NOT return	Rules published during the pilot remain published; the elapsed analyst time is spent; and the board's experience of recording dispositions before discussion cannot be un-had.
RV3 authorised to reverse	The code administrator, unilaterally, without the director's approval. Named in the memo.
RV4 time to reverse	One period (one week).
RV5 cost of reversal	Approximately 0.05 rules per week for two periods, in re-circulation and re-briefing. Inside `h = 0.9207`.
Class	**Costly-reversible.** RV5 is inside `h`. Proceeds under the normal gates.

Gate 6 — consent. Affected: five discipline leads, one analyst redeployed to the scribe role, the code administrator. Informed at the board session of the commitment week, in writing, with the specification attached. **Declination mechanism:** a discipline lead who does not wish to record a written disposition before discussion may abstain; the abstention is recorded as an abstention and the item proceeds with the remaining disciplines, at a correspondingly lower `U_t` for that session. The analyst offered the scribe role may decline, in which case the role is offered to the next volunteer and, if there is none, **the conversion lever is not available and `h` is recomputed without it.** That second clause is the honest form of consent in this framework: **the declination is not free — it changes the arithmetic, and the arithmetic is published — but it carries no penalty to the person, and the design has a defined behaviour when it is exercised.**

Gate 7 — the pilot.

Scope bound	The full board, for items in two of the five application classes
Duration bound	26 weeks, plus a 6-week conversion lag before measurement begins

Downside cap	**7 rules.** Itemised worst case: removing the re-review loop lets weak dispositions stand and doubles the appeal-reversal line from 0.14 to 0.28 for the whole window, `0.14 x 32 = 4.48` rules; transition friction at 0.05 per week for 32 weeks, 1.60; reversal cost at 0.05 per week for 2 weeks, 0.10. Total **6.18 rules**, capped at 7
Admissibility	`C <= (1−delta)Q + eta S` holds under the worst case at `Q = 113`: `0.42 <= 112.65` comfortably
Comparison	The three application classes not in the pilot, running the prior process
Preregistered prediction	Mean weekly change in the order proxy over weeks 7 to 32 of at least +0.55 rules per week
Falsification	If the mean weekly change over weeks 7 to 32 is below +0.30, the intervention is judged ineffective and the memo is rescinded

Trip-wires.

wire	threshold	automatic action
realised geometry	mean `F(U_t)` below `S_req = 1.2564` nats over any 8-week block — equivalently, low-participation share above 20 per cent	notify the director in the same period; the +0.55 promise is suspended pending review
`C`	above 0.55 rules per week	suspend the growth commitment; resume the cost work
`h`	at or below 0	stop building; revert to Section 4.5
single-week rise in the proxy	above `eta ln N = 1.6899` plus the 6-week lag allowance	investigate the proxy before believing the number
harm: appeal reversal rate	above 1.5 times the pre-pilot rate	stop, regardless of the proxy — the non-proxy harm criterion

The geometry trip-wire is the live one. The promise was computed at a mean `F` of 1.3241 nats. **Above a low-participation share of about 20 per cent the promise breaks**, and 32.7 per cent — the share the office was already running at before the intervention — is well above it. There are 8 percentage points of margin between the target and the break-even point, and the office was 13 points the wrong side of it a year earlier. **The realisation work is not an enhancement; it is a load-bearing component of the commitment and must be maintained for as long as the commitment stands.**

8.1.12 Time — what may be said, and what may not

```
conversion lag m  = 6 weeks
increment r       = 0.55 rules/week, judged over weeks 7 to 32
tau = 1/delta     = 323.6 weeks = 6.22 years
3 tau             = 971 weeks   = 18.7 years        SETTLING TIME, NOT A DATE
Q_inf             = (1.05 x 1.3241 − 0.42)/0.003090 = 314.0 rules   REPORTED, NOT PROMISED
```

Sanity checks against the inherited results: `Q_inf = 314.0` is below the uniform bound `eta ln N / delta = 546.9` (CC-T5) and below the level `Q = 411.0` at which `h` returns to zero, **so the design does not, on its own trajectory, drive the office back into the maintenance trap.** It reaches 76 per cent of the way there, which is a fact worth knowing in advance.

horizon	modelled `Q`
1 year	142.9
3 years	190.0
5 years	224.1
10 years	273.8
3 tau (18.7 years)	304.1
limit (not attained in finite time)	314.0

A small correction to the source, made rather than inherited. The source's trajectory table labels the `3 tau` row "approximately 314", conflating the settling time with the limit. At `3 tau` the model reads **304.1**, which is 95 per cent of the approach from 113 to 314.0 — which is what `3 tau` means. The limit is approached and not attained.

What may be said to the director: the manual should gain about 32 rules in the first year, and that number is checkable in thirty-two weeks. What may not be said: that the manual will reach 314 rules. It will, if every parameter holds for nineteen years, which no parameter in any organisation has ever done.

8.1.13 The decision that was not taken

The office also has a standing proposal, three years old, to merge the fire and structural review functions under one lead. The pairwise table shows why it never gets done and why it should not: `u(STRU, FIRE) = 0.3056` is the lowest pair in the set, which makes it the obvious merge candidate to anyone reading an organisation chart. **Under the merge rule it survives easily — 0.3056 is three times the 0.10 threshold — and merging would drop `N` from 5 to 4, costing `ln 5 - ln 4 = 0.2231` nats of ceiling**, which at the post-intervention conversion rate is

```
1.05 x 0.2231 = 0.2343 rules per week = 12.2 rules per year, permanently
```

against a saving of one lead post. **Recorded as a decision not to act, with the number, and with the observation that would reopen it: if a future probe set shows `u(STRU, FIRE) < 0.10`, the two functions have converged and the merge becomes free.** A refusal, dated, with its arithmetic and its reopening condition, is what the discipline of not acting looks like when it is written down.

16.5.8 8.2 A hospital ward shift handover: positive headroom, and a system losing stock

A 34-bed inpatient ward. **The complaint presented by the charge nurse is "handover is not working." Nothing else is offered.** The case is worked through the catalog's use specification (Appendix C).

8.2.1 The survey

Period: one week. **Order observable:** the number of currently-admitted patients for whom an incoming staff member can, unprompted, state the three current concerns and the escalation trigger, measured by spot audit at three random times per week. `Q_0 = 21` of 34.

FM-29 check, run first. Is the proxy a repackaging of interaction? The lag test: handovers were cancelled for three consecutive weeks during a staffing crisis, and the proxy fell 27 to 21 over those three weeks — **a decay with a time constant, not a collapse to zero. The proxy is stored order, not input.**

Decay. From the same natural experiment, with `S` approximately zero:

```
delta      = 1 − (21/27)^(1/3) = 0.080359 per week
half-life  = ln 2 / (−ln(1 − delta)) = 8.274 weeks
```

Categories. Five are declared as handover participants: nursing, medicine, pharmacy, physiotherapy and occupational therapy, and discharge coordination. **Two attend. This is FM-06, the Frozen Denominator.** `N = 2`, so `ln 2 = 0.693147` and not `ln 5 = 1.609438`. **The ward has been planning against a ceiling 132 per cent higher than the one it owns.**

Non-parallelism. Independent concern-listing before discussion, 18 patients, two attending categories. Mean overlap of listed concerns, treated as the alignment cosine: `rho = 0.79`.

```
u     = 1 − rho^2 = 1 − 0.6241 = 0.375900
N_eff = 1 + (2−1)(0.3759) = 1.375900        [the affine effective count]
S     = ln(1.375900) = 0.319108 nats        −> 46.04 per cent of ln 2
```

Conversion. Over 12 weeks the ward logged 25 carried-patients attributable to specific handover exchanges, against a cumulative `S` of `12 x 0.319108 = 3.8293` nats.

```
eta = 25 / 3.8293 = 6.528614 carried-patients per nat
```

Cost, as a residual (FM-23). Productive handover time is 38 minutes per week; the audit finds 21 minutes per week of interruption, repetition and downstream rework attributable to handover failure. The conversion rate is `25 / (12 x 38) = 0.054825` carried-patients per productive minute.

```
C = 21 x 0.054825 = 1.151316 carried-patients per week
```

8.2.2 The two master numbers

```
h  = eta ln N − delta Q − C
   = 6.528614 x 0.693147 − 0.080359 x 21 − 1.151316
   = 4.525290 − 1.687531 − 1.151316
   = +1.686444 carried-patients per week        ->  headroom is POSITIVE

S / ln N = 0.319108 / 0.693147 = 46.04 per cent ->  realised fraction is LOW
```

Validation. The predicted one-step change is

```
eta S − delta Q − C = 2.083333 − 1.687531 − 1.151316 = −0.755513 per week
```

against an observed fall of 24 to 21 over the preceding four weeks, or `−0.750` per week. The residual is `−0.006`. **The survey reproduces the observed decline and is admissible.**

8.2.3 The diagnosis

Five co-occurring faults, ranked by effect on the one-step condition.

1. **FM-06, the Frozen Denominator.** Three of five declared categories do not attend. `N = 2`.

2. **FM-25, the Head of Table.** Handover is delivered by the outgoing senior nurse standing, to a seated audience. Speaking order tracks rank in all three observed sessions.
3. **FM-26, the Reverberant Room.** Handover happens in a corridor alcove, 4.2 by 2.6 by 3.0 metres.

   ```
   V    = 32.760 m^3
   A    = 10.92 x 0.02 + 10.92 x 0.04 + 40.80 x 0.05 + 6 x 0.45
        = 0.2184 + 0.4368 + 2.040 + 2.70 = 5.395 m^2 sabins
   RT60 = 0.161 x 32.760 / 5.395 = 0.9776 s
   ```

   Nearly a full second, in a room where six people must exchange precise clinical detail in under forty minutes with corridor traffic passing through.
4. **FM-12, the Ritual That Became a Meeting.** The written handover sheet exists. Its template was last revised four years ago; the audit finds it filled in for 6 of 34 patients.
5. **FM-01, a partial Parallel Wall.** `u = 0.3759` between two categories whose training, supervision and incentives substantially overlap.

The binding constraint is not headroom. Headroom is `+1.69` per week and the ward is losing 0.76 per week. The ward operates at 46 per cent of a ceiling that is itself artificially low. **This is precisely the case Section 6.3 warns about, and a diagnosis that stopped at "headroom is positive, we are fine" would have been exactly wrong.** It is the paper's counter-result, and it is exhibited rather than asserted.

8.2.4 The intervention

Ranked by lever cost, and every element checked against FM-32: **none of the four requires anyone to participate who declines**, and the pharmacy and discharge attendance is negotiated as a scheduling change rather than imposed.

#	Move	Lever	Target
1	Restore the handover sheet as a live board, filled during handover and not after	`eta`, `delta`	`delta` 0.0804 to 0.065
2	Pharmacy and discharge coordination attend; physiotherapy does not — scheduling infeasible, stated honestly	`N`	`N` 2 to 4
3	Acoustic treatment of the alcove: 10.92 m^2 ceiling raft at `alpha` 0.90, 12 m^2 wall panel at 0.85	`u`, `C`	`RT60` 0.98 to 0.22 s
4	Independent written concern lists from each category before the verbal round	`u`	`rho` 0.79 to 0.68

```
Acoustics after treatment:
A    = 0.2184 + 9.828 + (40.80 - 12) x 0.05 + 12 x 0.85 + 2.70 = 24.386 m^2 sabins
RT60 = 0.161 x 32.760 / 24.386 = 0.2163 s

Post-intervention parameters:
N     = 4        ln N  = 1.386294
rho   = 0.68 ->  u     = 0.537600
                 N_eff = 2.612800        [affine effective count]
                 S     = 0.960422 nats   -> 69.28 per cent of ln 4
delta = 0.065    eta   = 6.528614 (held constant; conservative)
C     = 26 x 0.054825 + 0.20 = 1.625450  (more attendees, plus board upkeep)
```

8.2.5 The prediction, and the disclosure that the model is out of range

```
h_new = 6.528614 x 1.386294 - 0.065 x 21 - 1.625450
      = 9.050581 - 1.365000 - 1.625450 = +6.060131 per week

step  = eta S - delta Q - C
      = 6.270227 - 1.365000 - 1.625450 = +3.279777 carried-patients per week

Q_inf = (6.270227 - 1.625450) / 0.065 = 71.46
tau   = 1/delta = 15.385 weeks
```

DF-6 applies and is stated. `Q_inf = 71.46` exceeds the ward census of 34. The linear law has no saturation term; beyond `Q = 34` the model is outside its range and must be replaced. **The binding limit is the census, not the fixed point.** The honest prediction is therefore a saturation date and not an equilibrium:

```
t = 1   Q = 24.280
t = 2   Q = 27.346
t = 3   Q = 30.214
t = 4   Q = 32.895
t = 5   Q = 35.401   -> census cap reached; the model is invalid beyond this point
```

Predicted outcome: the ward reaches full carried-patient coverage in four to five weeks and holds it. The do-nothing comparison from the same start gives `Q = 18.32` at week 4, 16.41 at week 8, and 14.76 at week 13.

8.2.6 What would falsify it

Preregistered before the first acoustic panel is hung.

1. **Trajectory.** Spot-audit the proxy at weeks 2, 4 and 8. At a conversion rate 20 per cent below the estimate, week 4 still predicts `Q = 28.35`. **If the proxy is below 28 at week 4, the conversion estimate is refuted** and the model requires re-specification.
2. **Acoustic mediation.** Measured `RT60` post-installation must fall below 0.35 s. If it does not, the absorption schedule was wrong and prediction 1 is void regardless of outcome.
3. **The geometry channel.** Independent concern-listing repeated at week 6 must show `rho <= 0.72`. **If `u` does not move while the proxy does, the proxy is rising for a reason outside the model — most likely the measurement effect named in DF-7 — and the causal claim is refuted even though the outcome is favourable.**
4. **Shuffled-input control.** Refit the minimal model with the interaction series randomly permuted across weeks. If the shuffled fit matches the ordered fit on held-out data, the apparent dependence on surplus is spurious. **This is the baseline most often omitted, and its omission would invalidate the whole diagnosis** — though Section 4.2 records that

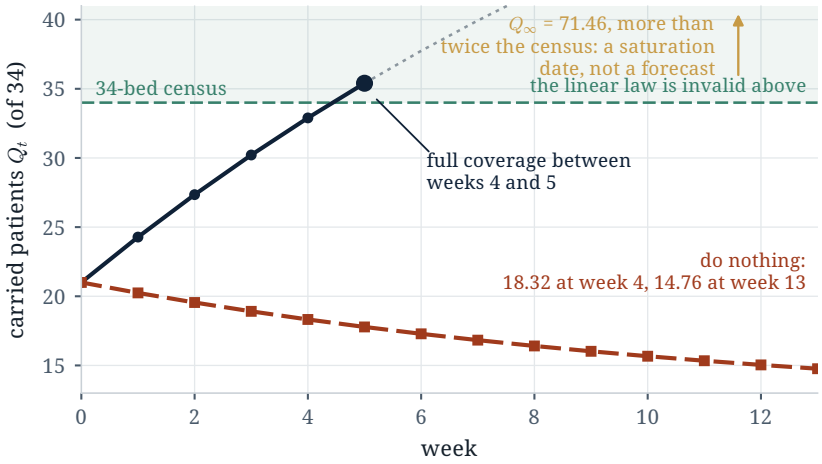


Figure 16.4: **Positive headroom, a declining stock, and a fixed point outside the model's range.** The ward has `h = +1.686` carried-patients per week and is losing 0.756 per week, because it realises only 46.04 per cent of a ceiling that the frozen denominator has already halved. Under the four interventions the modelled trajectory crosses the 34-bed census between weeks 4 and 5, at which point **the linear law is outside its range** — its own predicted fixed point of 71.46 is more than twice the census and is reported as a saturation date rather than as a forecast. The do-nothing comparison falls to 14.76 by week 13. Modelled from the parameters of Sections 8.2.1 and 8.2.4; generated by `figs/src/CCP-016_fig04.py`.

the shuffle is anti-conservative as ordinarily run and must be calibrated by simulation on this design.

5. **Persistence and matched-autoregression baselines.** The fitted model must beat both on held-out weeks. **If a plain autoregression on the proxy alone predicts as well, the framework has added nothing here and should be reported as having added nothing** — which is exactly what happened to the corpus's one preregistered confirmatory test [Curzi, 2026i].

Any one of 1, 3, 4 or 5 failing refutes the diagnosis. Failure of 2 refutes only the installation.

The catalog's job ends here. It named five faults, priced four of them, and produced a falsifiable number. Whether the number survives contact with a real ward is a question for a preregistered trial that has not been run.

16.5.9 8.3 The two cases side by side

	office, before	office, after	ward, before	ward, after
`h` (proxy units per week)	-0.1720	+0.9207	+1.6864	+6.0601
`step` (proxy units per week)	-0.4025	+0.6211	-0.7555	+3.2798
realisation deficit	0.2305	0.2996	2.4419	2.7803
realised fraction `S / ln N`	71.4 per cent	82.3 per cent	46.04 per cent	69.28 per cent

	office, before	office, after	ward, before	ward, after
binding fault	dissipation and conversion	—	frozen denominator and conversion	—
verdict	**refuse the request**	admit `r = 0.55`, reserve 0.597	**positive headroom, still losing**	saturation at the census, not an equilibrium

The realised fractions for the office are the mean of `F` over the ceiling: `1.1498 / 1.6094` and `1.3241 / 1.6094`. **They are computed from the mean of `F`, not from `F` of the mean `U`, which would return 76.2 per cent and 84.5 per cent respectively and would overstate both.** # 9 Discussion

16.5.10 9.1 What follows

The refusal is a deliverable, and it is the procedure's most useful output. The office case ends its first phase with a written paragraph containing a number, and that paragraph is worth more than the intervention that follows it, because it converts an argument about effort, commitment and facilitation into an inequality between three measured quantities. **An analyst who cannot produce a refusal cannot produce a commitment either**, since the same arithmetic underwrites both.

Two administrative levers moved a system from an unsustainable regime to a fundable one without adding a person. The conversion lever supplied 80.8 per cent of a 1.093-unit swing at a cost of 0.4 of an existing post, and the dissipation lever the remaining 19.2 per cent at the cost of one policy decision. **That is the ranking of Section 4.9 in a system's own numbers, and it is the strongest practical claim in the paper**: the cheapest lever is usually the first one, and it is almost never the one proposed, because capture is nobody's job and hiring is somebody's.

The realisation channel is a distinct route with a distinct instrument. Raising the mean of `F` by regularising participation moved the office's realised step from +0.4381 to +0.6211 while leaving headroom at exactly 0.9207. **Nothing about the budget changed and the deliverable changed by 42 per cent.** The mechanism is available in most organisations for the same reason that most organisations run sessions below their own measured capacity.

The concavity penalty is a real and recoverable quantity. 0.0761 nats — 6.2 per cent of the office's surplus — is charged for nothing but the variability of who turns up, and 0.0407 of it is recovered by regularisation alone, worth 0.0427 rules per week. **It is invisible to anyone reasoning about the average board rather than about the distribution of boards**, which is why the feasibility test is stated on the mean of `F` and not on `F` of the mean.

And the counter-result: positive headroom is not health. The ward has `h = +1.69` per week and loses 0.76 per week. Both numbers are correct and they are answers to different questions. **The diagnostic obligation is therefore to report two numbers, not one**, and the catalog's use specification makes it a requirement rather than a recommendation.

16.5.11 9.2 What does not follow

No causal claim follows from either worked case. Both are constructed. Neither is evidence that any organisation behaves this way, and neither may be cited as such.

No identified estimate of conversion follows from any survey, however clean. Section 3.4 states it: the specification is reflexive and its endogenous effect is not identified from observational data [Manski, 1993, Shalizi and Thomas, 2011]. **Gate 0 governs the admissibility of evidence about the proxy and never the admissibility of a causal claim.** The corpus's own confirmatory test did not meet its declared margin [Curzi, 2026i], and that is the weaker of the two problems.

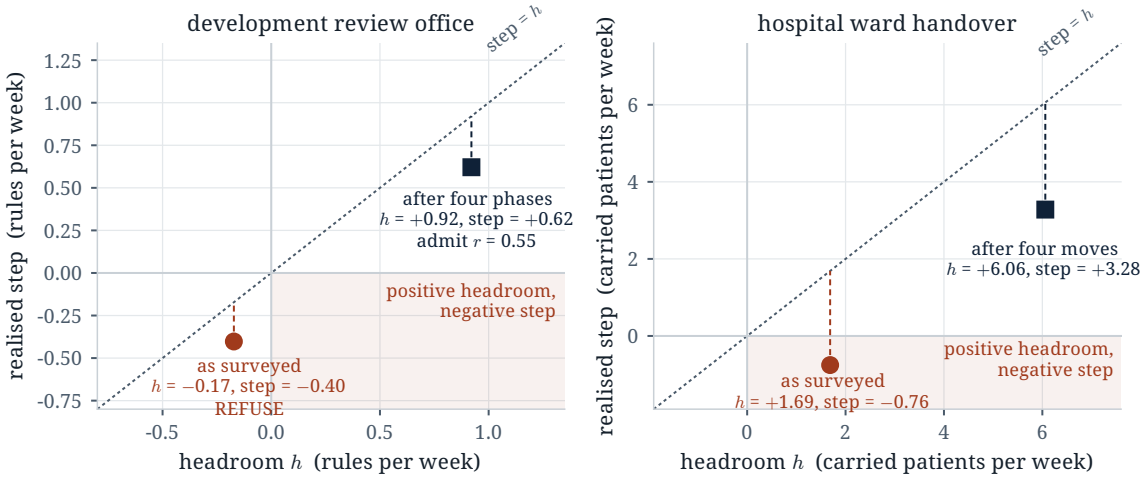


Figure 16.5: **Headroom and the realised step are independent quantities, and only their joint reading is a diagnosis.** Left, the development review office: headroom `−0.1720` with a step of `−0.4025` before intervention, and `+0.9207` with a step of `+0.6211` after. Right, the hospital ward: headroom `+1.686` with a step of `−0.756` before intervention — **positive headroom and a declining stock at the same time** — and `+6.060` with a step of `+3.280` after. The two panels are not on a shared axis because the two systems count different things: rules per week against carried-patients per week. The diagonal marks equality of headroom and step, which is attained only at the geometric ceiling. Modelled from Sections 8.1 and 8.2; generated by `figs/src/CCP-016_fig05.py`.

The lever ranking does not follow as a law. It is a default over the parameter ranges the worked systems occupy. In a system whose capture discipline is already near its medium's absorption limit, `eta` has no room and the ranking inverts; the procedure requires an analyst to compute and report that rather than to inherit the order.

The catalog's closure does not follow as a discovery. It is a definition given the model (Section 6.1). A different model of institutional accumulation has a different and possibly infinite catalog.

Nothing about surplus follows from the choice of a logarithm. Latané [1981] and Briscoe et al. [2006] fit the qualitative facts as well, and nothing assembled here discriminates. **The logarithm has an axiomatic characterisation [Curzi, 2026] and a dynamical warrant that requires no perceptual story [Peters and Gell-Mann, 2016]; it does not have a measurement that selects it.**

And no arithmetic result licenses the consent constraint. It is normative, it is stated as such, and it is the one place in the procedure where a design is refused for a reason the equations do not supply.

16.5.12 9.3 Relation to the surrounding literature

The procedure's central move — refuse a commitment that exceeds a computed budget — has an exact analogue in operations, where organisational slack is the recognised name for the margin that keeps a system's latency finite [Cyert and March, 1963, Kingman, 1962, Hopp and Spearman, 2011]. **The contribution here is not the idea of a budget but the identification of the budget with a specific, separately measurable expression**, and the demonstration that its sign is a hard gate rather than a soft indicator.

The lever ranking's fourth position is a restatement of a result that predates this framework by half a century: coordination cost grows with the number of pairs while the useful return does not [Brooks, 1975, Conway, 1968]. **The framework's addition is to make the com-**

parison numerical — a ceiling gain of `eta ln((N+1)/N)` against a priced path increment — so that the refusal of a headcount proposal is computed rather than argued.

The catalog's Class V is the part of the material with the strongest external footing and the weakest internal one. Room acoustics is real physics [Sabine, 1922, Schroeder, 1965], turn-taking equality is a measured driver of collective performance [Woolley et al., 2010], and direct observation of where people actually go is a mature instrument [Whyte, 1980, Jacobs, 1961, Alexander et al., 1977, Hall, 1966]. **What is not established is the step from any of those to a numerical value of `u`**, and Section 3.8 badges that step as a modelling assumption at each occurrence.

Where the literature disagrees with the framework, the disagreement is recorded rather than managed. **The framework's surplus is monotone increasing in `u`; the field evidence says the real curve turns over** [Uzzi and Spiro, 2005], and diversity in the short run has been found to reduce trust and civic participation by a scholar who did not want the result [Putnam, 2007], with the rebuttal [Abascal and Baldassarri, 2015] and the conditions under which contact across difference does pay [Allport, 1954, Pettigrew and Tropp, 2006]. **The strongest formal result on diversity is narrower than its popular reading** [Hong and Page, 2004, Thompson, 2014]. And the design and safety literature warns that the framework's instinct to add categories can make matters worse in a tightly coupled system [Perrow, 1984, Sagan, 1993], which is a limit on Section 4.9's fourth lever that the arithmetic alone does not supply.

16.5.13 9.4 Defects found in the source material, printed rather than repaired

Five, in descending order of consequence.

D-i. The lever comparison mixes two bases and reports a hybrid ratio. The source's lever table sets a gain in `h` of `+0.8828` against a gain in `step` net of path cost of `+0.1578`, labels both "gain in `h`", and reports the ratio per full-time equivalent as fourteen. **On a consistent headroom basis the ratio is 18.4; on a consistent step basis it is 10.0.** The qualitative conclusion survives on either basis and is strengthened, but the printed figure is not a like-for-like comparison. Section 8.1.10 gives all three.

D-ii. The notation `realised F(U) = 1.1498` writes a mean of `F` as though it were `F` of a `U`. It sits three lines above `realised u = 0.6019` in the same block, and `F(0.6019) = 1.2260`, so a reader recomputing the printed relation finds it fails by 0.0762 nats. **The number 1.1498 is correct and the notation is not**; the source's own later text makes the distinction explicitly and correctly, which makes the earlier notation an inconsistency rather than an error of substance. Every occurrence in this paper is disambiguated in words.

D-iii. The gate ordering in the source's exposition is out of sequence, developing Gate 7 before Gate 5 while running them in numerical order in the worked case. Section 4.1 states the adopted order and the reason.

D-iv. The trajectory table labels the value at three time constants as the limit. At `3 tau` the office model reads 304.1 rules, which is 95 per cent of the approach; 314.0 is the limit and is not attained in finite time. Section 8.1.12 prints both.

D-v. The symbol `rho` carries two meanings in the catalog — the alignment correlation in FM-01 and the queue utilisation in FM-22. This paper writes the second as `rho_u` and says so at the point of use.

One further item is not a defect and is recorded because it looks like one. The source's ward case predicts a fixed point of 71.46 carried-patients against a 34-bed census. **That is a prediction outside the model's range, and the source declares it as such and replaces the forecast with a saturation date.** It is a credit to the source rather than an error, and it is kept prominent in Section 8.2.5 for that reason.

16.6 10 Verification

The results of this paper compound from seven postulates and one design constraint. Each is atomic, each is checkable on its own, and each carries its status. **Postulates that could be dropped with the results still following have been dropped**; what remains is what the chain consumes. **The reader is warned at the outset of what Section 10.3 makes explicit: this chain ends in arithmetic on top of definitions, so it delivers the mathematics with certainty and delivers nothing whatever about institutions.**

16.6.1 10.1 Dimensional audit

`[Q]` denotes the dimension of stored order and `[T]` the dimension of time. The period is the unit of `[T]` and every rate is *per period*, which is why no explicit `[T^-1]` appears once the equations are written in discrete steps.

Symbol	Dimension	Unit	Status
`Q_t`	`[Q]`	declared proxy units (live rules; carried patients)	**measured** in whatever the declared proxy counts
`U_t`, `u`	`[1]`	none	**stipulated** dimensionless; `u = 1 - rho^2`
`S_t = F(U_t)`	`[1]`	nats	**derived** dimensionless: a logarithm of an effective count
`eta`	`[Q]` per nat per step	proxy units per nat per period	**measured**; *not* dimensionless
`delta`	`[1]`	per period	**measured** dimensionless fraction per step
`C_t`	`[Q]` per step	proxy units per period	**measured**, as a residual
`h_t`, `step`, `r_t`	`[Q]` per step	proxy units per period	**derived**
`S_req`	`[1]`	nats	**derived**: `[Q]/step` divided by `[Q]` per nat per step
`U_req`	`[1]`	none	**derived**
`q` (a shed)	`[Q]`	proxy units	**chosen**
`delta x q`	`[Q]` per step	proxy units per period	**derived**; commensurable with `h`
`Q_inf`	`[Q]`	proxy units	**derived**: `[Q]`/step over `[1]`/step
`tau`, `t_half`, `3 tau`	`[T]`	periods	**derived**
`r/h`	`[1]`	none	**derived**; a pure ratio, defined only for `h > 0`
`rho` (alignment)	`[1]`	none	**measured** correlation coefficient

Symbol	Dimension	Unit	Status
`rho_u` (utilisation)	`[1]`	none	**measured** fraction of capacity committed
`RT60`	`[T]`	seconds	**measured or computed**; *not* in periods — see below
`V`	length cubed	m^3	**measured**
`A`	length squared	m^2 sabins	**measured**; the sabin is a dimensionless absorption weight on an area

Does the gauge balance? `h_t = eta ln N − delta Q_t − C_t` balances if and only if every term carries `[Q]` per step. `C_t` does by construction. `delta Q_t` does, since `delta` is a dimensionless fraction per step. `eta ln N` does **only if `eta` carries `[Q]` per nat per step**, since `ln N` is a dimensionless count in nats. **`eta` is therefore not dimensionless in this framework**, and the worked cases are consistent with it: the ward reports 6.528614 carried-patients per nat, which is `[Q]` per nat, with the per-period part carried by the cumulative surplus in the denominator. **The office reports `eta = 0.5015` and `1.05` without units, which is a silence in the source**; the values should read *rules per nat per week*, and that is recorded here rather than quietly corrected. Two consequences: `eta/delta` carries `[Q]` per nat and is the gain from a sustained nat to a stock; and `delta x q` carries `[Q]` per step, which is what makes the shed commensurable with the rest of the gauge.

Two dimensional cautions. First, `RT60` is in **seconds** and every other time in this paper is in **periods**. They are never added, never compared, and the step from one to the other is the modelling assumption of Section 3.8, not a unit conversion. Second, a headroom calculation in which `C` is denominated in hours or currency while `Q` is denominated in rules is **dimensionally malformed however arithmetically tidy it looks**; the denomination requirement is part of the survey protocol [Curzi, 2026i] and the ward case satisfies it by converting minutes to carried-patients through a measured yield rate before subtracting.

Where there is no operational unit, it is said. `Q` is stored order and **stored order has no natural unit**: it is counted in whatever the declared proxy counts, and two systems' `Q` values are not comparable unless they declare the same proxy — which is exactly why Section 8.3 places the two worked systems in separate panels. Comfort itself, the underlying construct of the corpus, has **no operational unit at all**; it is treated as a dimensionless magnitude by stipulation, and Section 11 records that the quantitative structure required for the stipulation has not been tested.

16.6.2 10.2 The postulates

P1 — The gauge is `h_t := eta ln N − delta Q_t − C_t`. *Status:* **DEFINITION** — fixed by Curzi [2026b] and adopted unchanged. It carries no empirical weight on its own.

P2 — The one-step change is `Q_{t+1} − Q_t = eta F(U_t) − delta Q_t − C_t`. *Status:* **DEFINITION** — the model form of the storage law, restated [Curzi, 2026a]. That any organisation obeys it is postulate P-C of Section 3.6 and is not verified here.

P3 — `0 <= F(u) <= ln N` on `u` in `[0,1]` for `N >= 2`. *Status:* **CITED** — the range result of the surplus functional [Curzi, 2026], following from `F(u) = ln(1 + (N−1)u)`.

P4 — `F(u) = ln(1 + (N−1)u)` is strictly concave on `[0,1]` for `N >= 2`. *Status:* **ARITHMETIC** — differentiate twice: `F''(u) = −(N−1)^2 / (1 + (N−1)u)^2`, which is strictly negative for every `u` in `[0,1]` whenever `N >= 2`.

P5 — For a strictly concave function, the mean of the function is at most the function of the mean. *Status:* **ARITHMETIC** — Jensen's inequality; **the bibliography assembled for this**

corpus contains no entry for it, so it is stated as elementary mathematics and demonstrated on the paper's own two-point instance in Section 10.5 rather than cited.

P6 — Communication paths among `n` categories number `n(n−1)/2`, so a step from `n` to `n+1` adds `n` paths. *Status:* **CITED** [Brooks, 1975, Conway, 1968].

P7 — Cost accrues every period whether or not surplus was delivered. *Status:* **POSTULATED** — an assumption about how organisations charge overhead, with no external support. It is consumed at exactly one place, R1b below, and only for the *interpretation* of `C` as a standing charge; the inequality itself holds however `C` is generated.

One further item is stated and is not in the chain. The consent constraint (P-E, Section 3.6) is **normative, is not falsifiable, and consumes no step below.** It bounds the admissible design set and it does no arithmetic work, so by the razor it is listed here and excluded from the compounding chain rather than dressed as a premise.

16.6.3 10.3 The compounding chain

```
P2 + P3          -> L1   since F(U_t) <= ln N,
                         Q_{t+1} − Q_t <= eta ln N − delta Q_t − C_t.      ARITHMETIC

L1 + P1          -> R1   Q_{t+1} − Q_t <= h_t, so r_t <= h_t; and equality
                         requires F(U_t) = ln N, i.e. U_t = 1.            ARITHMETIC

L1 + P1 + P7     -> R1b  h_t < 0 implies Q_{t+1} − Q_t < 0 for EVERY
                         U_t in [0,1]: no geometry prevents decline.      ARITHMETIC

P2               -> L2   Q_{t+1} − Q_t >= r_t  iff
                         F(U_t) >= (delta Q_t + C_t + r_t)/eta =: S_req.  ARITHMETIC

L2 + P3          -> R2   the increment is attainable only if
                         S_req <= ln N, which is exactly r_t <= h_t.      ARITHMETIC

L2               -> L3   inverting F: U_req = (exp(S_req) − 1)/(N − 1).   ARITHMETIC

P4 + P5          -> R3   mean F(U_t) <= F(mean U_t), equality only if U
                         is constant; so the feasibility test is
                         mean F(U_t) >= S_req, never F(mean U) >= S_req.  ARITHMETIC

P1               -> R4   dh/dQ = −delta, so removing q units of stored
                         order returns exactly delta x q per period.      ARITHMETIC

P1               -> L4   raising N to N+1 raises the ceiling by
                         eta ln((N+1)/N).                                 ARITHMETIC

P6               -> L5   the same step adds N communication paths.        CITED

L4 + L5          -> R5   net headroom from a category, at per−path cost c0,
                         is eta ln((N+1)/N) − N c0, which changes sign.   ARITHMETIC

P1 + P2          -> R6   h_t − (Q_{t+1} − Q_t) = eta(ln N − F(U_t)) >= 0:
                         the realisation deficit, non−negative always,
                         and zero only at U_t = 1.                        ARITHMETIC
```

Every line is checkable from its own stated inputs. R1, R1b, R2, R3, R4, R5 and R6 are arithmetic consequences of two definitions (P1, P2), one cited range result (P3), two elementary mathematical facts (P4, P5), and one cited combinatorial fact (P6). **Only R1b consumes a**

postulate with no external support, and it consumes it only for the reading of `C` as a standing charge.

What this chain does not establish. It does not establish that any organisation obeys P2. That is postulate P-C, it is not verified here, and the one preregistered confirmatory test of it so far reported in this corpus did not meet its own declared margin [Curzi, 2026i]. **A chain whose steps are ARITHMETIC, DEFINITION and CITED delivers the decision rule with certainty and delivers nothing about institutions.** The procedure of Section 4 is valid in the sense that a decision rule can be valid — it is internally consistent and it follows from its premises. Whether it should be applied to a given organisation is a separate question and the answer is not in this section.

16.6.4 10.4 The verification instance, computed by hand

The office gauge, before intervention. Four multiplications and two subtractions.

```
ln 5           = 1.6094379...            round to 1.6094
eta ln N       = 0.5015 x 1.6094 = 0.80711...           print 0.8071
delta Q        = 0.003090 x 113 = 0.34917             print 0.3492
C                               = 0.6300
-----------------------------------------------------------------
h = 0.8071 - 0.3492 - 0.6300     = -0.1721             source prints -0.1720
```

The one-digit difference in the last place is the rounding of `eta ln N`; carried unrounded, `0.5015 x 1.6094379 = 0.80713311` and `h = -0.17203689`, which rounds to `-0.1720`. **The printed value is correct and the intermediate rounding is not; a reader reproducing it from the printed intermediates lands one unit low in the fourth decimal.**

```
step = 0.5015 x 1.1498 - 0.3492 - 0.6300 = 0.576625 - 0.9792 = -0.402575   -> -0.4026
```

against a printed `-0.4025`, again a fourth-decimal rounding. **Both are reported rather than harmonised.**

The office gauge, after two levers.

```
eta ln N = 1.05 x 1.6094379 = 1.68991     print 1.6899
h        = 1.6899 - 0.3492 - 0.4200 = +0.9207    exact: +0.92070981
```

The increment and the reserve.

```
S_req = (0.3492 + 0.4200 + 0.5500)/1.05 = 1.3192/1.05 = 1.25638...   -> 1.2564
S_req <= ln 5 = 1.6094                    -> no Failure A
U_req = (exp(1.2564) - 1)/4 = (3.51287 - 1)/4 = 0.62822              -> 0.6281
r/h   = 0.55/0.92070981 = 0.59736                                    -> 0.597
```

The shed. `0.003090 x 14 = 0.043260 -> 0.0433`. **And at a decay of 0.05:** `0.05 x 14 = 0.70`, whose ratio to 0.0433 is `16.17`. **One multiplication and one division.**

The category lever.

```
ln 6 - ln 5   = 1.7917595 - 1.6094379 = 0.1823216            -> 0.1823
ceiling gain  = 1.05 x 0.1823216 = 0.1914377                 -> +0.1914
paths added   = 15 - 10 = 5                        (P6: 6x5/2 - 5x4/2)
per-path cost = 0.34 x 0.42/10 = 0.01428                     -> 0.0143
added cost    = 5 x 0.01428 = 0.0714
net on h      = 0.1914 - 0.0714 = +0.1200
```

The ward's counter-result, in six operations.

```
eta ln N = 6.528614 x 0.693147 = 4.525290
delta Q  = 0.080359 x 21       = 1.687531
C                              = 1.151316
h        = 4.525290 − 1.687531 − 1.151316 = +1.686443
eta S    = 6.528614 x 0.319108 = 2.083333
step     = 2.083333 − 1.687531 − 1.151316 = −0.755514
```

`h` is positive and `step` is negative, from the same three subtractions with one term changed. That is the whole of the counter-result and it can be checked in under a minute. The survey's own validation follows: the observed change was `−0.750` per week, so the residual is `−0.006`, or 0.13 per cent of the ceiling term.

16.6.5 10.5 The concavity penalty, demonstrated on a two-point instance

P5 needs no citation if the paper's own instance exhibits it. The office's sessions are a two-point mixture: full-board sessions at mean `U = 0.7940` and reduced sessions at mean `U = 0.2071`, in proportions `1−p` and `p`. At the surveyed `p = 0.327`,

```
mean U   = 0.327 x 0.2071 + 0.673 x 0.7940 = 0.067722 + 0.534362 = 0.602084   -> 0.6021
F(mean U)= ln(1 + 4 x 0.602084) = ln(3.408336) = 1.226237                      -> 1.2262
```

and the mean of `F` over the same mixture is 1.1501, giving a penalty of `1.2262 − 1.1501 = 0.0761`. **The gap is strictly positive, which is P5, and a reader can see why in one line:** `F(0.2071) = ln(1.8284) = 0.6035` and `F(0.7940) = ln(4.1760) = 1.4294`, so a straight-line mixture of the two *function values* — `0.327 x 0.6035 + 0.673 x 1.4294 = 1.1593` — already sits below `F` of the mixed argument, 1.2262. **A chord under a concave curve lies below the curve; that is the whole of Jensen's inequality at two points.**

A note on the source's within-type figures, made rather than concealed. The two-point reconstruction just given returns a mean `F` of 1.1593, while the source's table prints 1.1501. The difference is that the source's session types are themselves distributions rather than point masses, so each type carries its own internal concavity penalty: the implied within-type mean of `F` is 1.4250 for full sessions against `F(0.7940) = 1.4294`, and 0.5844 for reduced sessions against `F(0.2071) = 0.6035`. **Both are below their own `F` of the mean, which is P5 again, one level down.** The source's table is linear in the share `p` with a slope of `−0.84058` nats per unit share, and every printed row of the mean-`F` column reproduces to four decimals from `1.42497 − 0.84058 p`. **The table is internally consistent and the reconstruction that recovers it is stated here so a reviewer does not have to find it.**

16.6.6 10.6 Limits, asymptotics, and the model's valid domain

`delta -> 0^+`. The shed's value `delta x q -> 0`: **in a perfectly durable system, shedding buys nothing, however much is shed.** The time constant `1/delta -> infinity` and `Q_inf` diverges. **The limit is singular:** at exactly `delta = 0` the update map is an isometry, no fixed point exists at all, and the bound degrades from a ceiling to a ramp [Curzi, 2026a]. **That value is outside the model's valid domain**, and it is a pathology rather than an ideal. The office's `delta = 0.003090` is close enough to it that the shed lever is worth 0.0433 and the time constant is 6.22 years — **the same number makes the office durable and makes it slow**, which is Section 4.15's rule in one parameter.

`delta -> 1^-`. The time constant goes to zero and the system has no memory; `Q_{t+1}` becomes `eta S_t − C_t`, an instantaneous transducer. The shed returns the entire stock, `1 x q = q`, and **is simultaneously maximally valuable and pointless, because nothing is stored to shed.** At `delta = 1` the storage law is not a storage law and any claim about accumulation is vacuous.

eta -> 0. `h -> -delta Q - C < 0` for every positive stock and cost. **A system with zero conversion is permanently in the `h < 0` branch and no interaction geometry of any quality changes it** — which is the conversion pathology taken to its limit, and the arithmetic reason the conversion lever ranks first.

N -> 1. `ln N = 0` and `F(u) = ln(1 + 0 u) = 0` for every `u`. **A single category has zero available surplus, not a small amount.** Ceiling and achievement collapse together, `h = -delta Q - C < 0`, and the realisation deficit is zero because there is nothing to realise. This is why the frozen denominator (FM-06) is the most consequential category fault: it is a move toward this boundary that nobody records.

N -> infinity. `ln N` grows without bound but the marginal gain `ln N - ln(N-1) ~ 1/N` vanishes, while the added paths grow as `N` (P6). **The two curves cross, and beyond the crossing additional categories cost more than the ceiling they buy** — the limit is not a licence to raise `N` indefinitely, it is the reason `N` is the most expensive and least rewarding lever.

u -> 0. `F -> 0`, the realisation deficit reaches its maximum `eta ln N`, and the step falls to `-delta Q - C`. **u -> 1.** `F -> ln N`, the deficit is zero, and `step = h` exactly — which is the extremal case CC-T1 names and the reason the procedure forbids `r = h`.

r -> h^-. The reserve ratio `r/h -> 1` and the plan is the best case. **h -> 0^+.** The reserve rule `r <= 0.7h` degenerates gracefully to `r -> 0`; the procedure does not divide by `h` anywhere else, and at `h <= 0` the reserve ratio is not formed at all (Section 3.5).

Q -> (eta ln N - C)/delta. `h -> 0` from above: the stall point of FM-17, at which success has consumed its own headroom. The office's post-intervention `Q_inf = 314.0` against a stall point of 411.0 is a system that arrives at 76 per cent of its own ossification.

Stated domain. The procedure as developed here is valid for `delta` in `(0, 0.5]` — with `(0.02, 0.30)` the band in which the standard time-constant approximations hold — `eta > 0`, `N >= 2`, `u` in `[0,1]`, `C_t <= (1-delta)Q_t + eta S_t` each period, and `Q` below any hard physical or organisational cap. **Outside that domain the equations still have solutions and they are not statements about a stock**, which is precisely the disclosure DF-6 requires and Section 8.2.5 makes.

16.6.7 10.7 Invariance

The procedure's arithmetic is **invariant under a change of the order proxy's unit**, and it is worth stating as a group action. Let `Q' = cQ` for a positive constant `c`. The gauge is invariant provided `eta' = c eta` and `C' = cC`, while `delta`, `N` and `u` are unchanged, since `delta` is a dimensionless fraction and `S` is a dimensionless count. Then `h' = ch`, `step' = c step`, `r' = cr`, and:

Quantity	Behaviour under `Q -> cQ`	Comparable across systems?
`delta`, `u`, `N`, `S / ln N`, `r/h`	**unchanged**	**yes**
`Q`, `eta`, `C`, `h`, `step`, `r`, `delta x q`	scale by `c`	**no**, unless the proxy is shared
`tau`, `t_half`	unchanged	yes, in periods; not in calendar units

The reserve ratio and the realised fraction are therefore the two quantities that may be compared between the office and the ward; the headroom values may not, and Section 8.3 separates the two systems into panels for that reason rather than for a graphical one.

What would break the invariance: a `delta` varying with `Q`, which makes the transformation non-linear; an `eta` saturating at large `S`, which the linear model excludes by assumption; or a proxy whose counting unit is not extensive, so that doubling the system does not double the count. **All three are plausible in real organisations and none is tested here.**

How many scales the form has been exercised at. Two in this paper, both constructed. Five elsewhere in this corpus, from a single desk to a 22-household block, reproducing each system's observed decline with residuals under 1.4 per cent of the ceiling term [Curzi, 2026k].

That is seven systems, one corpus, one author, and no independent replication. The invariance is a claim about form that has been exercised, not one that has been established.

16.6.8 10.8 Analogy discipline and the precision of the numbers

Section 3.8 states the discipline; two consequences belong here.

The circuit and queueing borrowings are shared mathematical forms, not derivations from physical premises, and no physical result's authority transfers. Room acoustics is the single derivation in the paper, and the step from reverberation time to a value of `u` is not part of it.

Precision. The numbers here are of three kinds and are quoted accordingly. *Defined or derived-from-defined* — `ln 5 = 1.6094`, `ln 2 = 0.693147`, `ln 6 − ln 5 = 0.1823`, `t_half`, `tau`, the path counts — are exact to the precision printed and may be recomputed to any further precision. *Illustrative model output* — every entry in every trajectory table, `Q_inf = 314.0`, `Q_inf = 71.46`, every headroom and step value in both worked systems — is the output of stipulated parameters and carries **no measurement uncertainty because there was no measurement**; four or six decimal places there record the arithmetic and not a confidence. **Measured or fitted quantities do not appear in this paper, because nothing in it was measured.** Where a real application reports fitted parameters it owes intervals, and Section 11 records that the corpus's one such report gives a profile interval on the retention coefficient and none on several other quantities [Curzi, 2026i].

One further precision note. The ward's parameters are printed to six decimals — `delta = 0.080359`, `eta = 6.528614`, `C = 1.151316` — because they are the exact outputs of the stated divisions, not because they are known to that precision. **`delta = 0.080359` is derived from a count of 21 out of 27 over three weeks; the underlying counts are integers and the sampling error on a spot audit at three random times per week is not estimated anywhere in the source. Six decimals on that number is arithmetic bookkeeping, and it should not be read as an error bar.** # 11 Limitations

Both worked systems are constructed, not observed. No dataset was collected for this paper. The development review office is a declared composite whose audit series is a simulation with declared generating parameters; the hospital ward is a constructed clinical case. Every number in Section 8 is the output of stated parameters under the stated model, carried through so a reader can re-derive it. **They demonstrate that the procedure runs and that it reports its own refusals and its own departures from range. They are not evidence that any organisation behaves this way and must never be cited as such.**

The reflexive specification is not identified, and no remedy is adopted. Section 3.4 states it in full: applied to a group whose members both influence and constitute the aggregate, the specification is reflexive, and Manski [1993] proves the endogenous social effect is not identified from observational data; Shalizi and Thomas [2011] show the same confounding for network data. **No fit, however well designed and however honestly held out, recovers `eta`.** This paper adopts no remedy, so **every applied reading of `eta` in it is a modelling position rather than an identified empirical result.** The most realistic available repair is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001]; the survey's fixed category set is a step toward it and not a substitute for it. **This is a deeper limitation than the fitting result below, and it would have applied equally had the fit passed.**

The one preregistered confirmatory test of the underlying law did not meet its declared margin. The fit landed inside the admissible box and its parameters agreed with independent measurement, but it beat a matched AR(1) by 2.8 per cent on held-out error against a preregistered requirement of 10 per cent, so the verdict was *parameters corroborated, predictive advantage not demonstrated* [Curzi, 2026i]. **That is the current evidential status of every applied claim in this paper, it is weaker than the arithmetic, and the office case is worked under the accounting-frame restriction that follows from it.** Reported here, in Results, and in the Preregistration declaration, because a failed rule gets the same plain language a passed one would get.

`u` has not been shown to be a quantity, and this is the strongest objection to the framework. Michell [1997] argues that psychology routinely fails to test whether its attributes are quantitative at all before measuring them. The available remedies are conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971] or a Rasch demonstration [Rasch, 1960]; **neither has been run for `u`, here or anywhere in this corpus**, and until one is, the ratio-scale treatment of `u` is a stipulation. **Section 3.5 states the consequence precisely: the realisation deficit `h - step` is a difference and the reserve ratio `r/h` is a ratio, and neither operation is licensed on an ordinal scale. If `u` is merely ordinal, the correct downgrade is to report the sign of the deficit and the ordering of levers, and to withdraw every decimal place.** Relatedly, if `u` is read as a similarity measure over people it cannot be assumed to behave metrically, since psychological similarity is asymmetric and violates the triangle inequality [Tversky, 1977] — the empirical counterpart of the framework's own clause that the triangle inequality fails for geometric triples.

`N_eff = 1 + (N-1)u` is not a Hill number and no external literature contains it. `1/HHI` is the Hill number of order two and `exp(H)` is the Hill number of order one, and both are grounded [Hill, 1973, Jost, 2006]; the affine effective count is this framework's own construction, adopted by design axiom. **A reviewer reaching for the diversity literature as support for Identity 1 will find that it does not apply.** Every effective count in this paper is the affine one and is named as such.

Which parameters were chosen rather than measured. In both worked systems, `(N, u, eta, delta, C)` are stipulated. **The 0.7 reserve convention is a stipulation with no derivation**; the merge threshold of `u < 0.10` is a stipulation; the target utilisation band of 0.70 to 0.85 is a design proposal; the low-participation target of 12 per cent is a judgement about which items warrant a full board. The per-path cost of `0.34 x 0.42/10` in Section 8.1.10 is an allocation of a fraction of the cost ledger to paths and is **not** an independent measurement. **None of these may be cited as findings.**

The catalog's prices are floors, not estimates, and several are not prices at all. As deployed, the surplus functional is a lower bound on the ensemble's information, so every nat quoted in Section 7 is a minimum. **Three entries have no price: FM-28 states that its damage is not computable in nats, FM-30 states that its damage is not quantitative, and FM-32 states that it is a boundary and is not priced because it is not for sale.** Those are honest entries in a catalog whose format promises a price, and the promise is not met in three of thirty-two cases.

Two external sources named in the source material have no key in this corpus's bibliography. A 1984 study on post-surgical recovery with a window view, and a 1991 essay on technologies that disappear, are both invoked in the source's Class V and Class III material. **They are named in Section 7 without citations and nothing in this paper rests on either.** They are recorded here so that the gap is visible rather than papered over.

The transfers from external results to this framework are modelling assumptions, and several are substantial. That a cascade result about sequential public choice describes a standing committee; that a queueing result about arrivals and service describes a review board whose arrival and service distributions were never measured; that a 25 per cent tipping threshold measured in controlled online experiments transfers to a residential block; that a reverberation time predicts a value of `u`. **Each is badged at its point of use and none is established here.** The general caution about transferring a single experimental result applies to all of them [Open Science Collaboration, 2015, Ioannidis, 2005].

The framework's own instrument is exposed to the failure it catalogues. FM-28 names Goodhart capture; `h` is a measure; **the headroom gauge is at least as gameable as anything it audits** [Goodhart, 1975, Campbell, 1979]. Section 4.16 prescribes keeping the order proxy out of the reporting line, but the prescription is a design rule and not a demonstration, and it does not protect `h` itself once `h` is what a sponsor is judged on.

The Lucas critique applies to this framework's own parameters and is only narrowly answered. The procedure asserts that a measure changes the system measured by it and then treats `delta`, `eta` and `N` as constants over an estimation window [Lucas, 1976]. **The narrow position taken in Section 3.4 — stability only over a declared window and only while the proxy**

is not a reported target — is a real restriction on where the model may be used, and it is stated rather than assumed away.

Parameter separability is an approximation. Postulate P-A: the levers interact, adding a category raises `C`, raising `eta` consumes capacity, and cutting `C` may cut `u`. The procedure requires second-order costs to be entered explicitly, which mitigates and does not remove the problem.

Constant parameters over the estimation window is an assumption organisations violate. Turnover, reorganisation and a change of leadership all move `delta`, `eta` and `N`. A survey spanning a regime change is measuring an average of two systems, neither of which exists.

The functional form is underdetermined by the evidence. Section 3.8 gives the alternatives — Latané [1981]'s sub-linear power law and Briscoe et al. [2006]'s `n log n` — and records that nothing assembled here discriminates. **The logarithm is chosen for its axiomatic characterisation and, where a dynamical warrant is wanted, for the multiplicative-dynamics result that needs no perceptual story [Peters and Gell-Mann, 2016]. It is not established by measurement.**

No scaling or heavy-tail claim is made or estimated. Where FM-10 invokes preferential attachment, no exponent is fitted, so the distribution-fitting procedure any such claim would require does not apply [Clauset et al., 2009] — and the empirical generality of scale-free structure is itself contested [Broido and Clauset, 2019]. **The entry uses a mechanism, not a tail.**

Unmet rigour checks, printed rather than left to a reviewer. (i) *Model comparison.* This paper fits nothing, so no information criterion applies to it; the corpus's one applied fit reports held-out RMSE against baselines and **no AIC or BIC**, so the parameter count was never paid for and a 2.8 per cent margin over a two-parameter AR(1) may not be a margin at all [Curzi, 2026i]. (ii) *Effect sizes and intervals.* Nothing in Section 8 carries an interval, because nothing in Section 8 was measured; that is honest for a constructed example and is not a substitute for the intervals a real application owes. (iii) *Scale type.* Section 3.5 gives the audit and its verdict: **the conjoint conditions that would license the ratio-scale treatment of `u` have not been tested.** (iv) *Statistical controls.* The shuffled-input control that both worked cases require is **anti-conservative as ordinarily run**, at a measured 0.108 against a nominal 0.05 for iid permutation and 0.120 for circular shift; Gate 0 therefore requires calibration by simulation on the study's own design, and any application that omits it is running a test at roughly twice its advertised false-positive rate. (v) *Independent replication.* Every system in this paper and in the corpus's applied papers was constructed or surveyed by one author; there is **no independent replication of any parameter estimate** [Open Science Collaboration, 2015, Ioannidis, 2005]. (vi) *Directional versus point prediction.* Meehl [1967] observes that directional predictions are nearly costless in the social sciences; the framework does make point predictions with constants, which is the right answer, but **only one of them has so far been put at risk and it did not meet its margin** [Curzi, 2026i].

What would falsify the account. H1 fails on a system sustaining an increment above its measured headroom with parameters stable and independently re-measured. H2 fails if positive headroom reliably predicts growth irrespective of the realised fraction. H3 fails on a measured headroom gain from a shed materially exceeding `delta x q` with no change in `C`. H4 fails on a system whose measured cost per unit of conversion gain exceeds its cost per unit of category gain — and **the ranking is a default that is expected to invert in some systems, which is stated in H4 rather than discovered later.** H5 fails if an observed step gain tracks `F` of the mean `U` rather than the mean of `F`. H6 fails if declared and capability-audited category counts agree within measurement error across a range of systems. **None of the six was tested here.**

The corpus's own failure record bounds what this catalog may claim. Curzi [2026] is a paper-length counter-result and it lands on three entries here directly: it records that **disclosure did not confer immunity in the strongest available test case**, which limits what the badging discipline of FM-30 can be expected to buy; that its own backing basket is **4.61 times less diversified than its component count claims**, which is FM-05 in another domain and measured rather than illustrated; and that a levy presented as a costless accounting correction in fact reallocates burden, with **42.0 per cent falling on one sector and a constructed 55 per cent passing through to consumers** — the general form of which is that a computed net effect can

be correct in aggregate and wrong about who pays. **This catalog prices parameters; it does not price incidence, and it should not be read as though it did.**

Nothing in this paper is a human-subjects finding. No interventions were run, no participants were recruited, and the consent constraint is a design constraint on prospective work rather than a report of past practice.

16.7 12 Conclusion

This paper converts the headroom gauge `h_t = eta ln N − delta Q_t − C_t` into a decision procedure of seven gates — evidence admissibility, regime, feasibility, lever ranking, sequencing, reversibility, consent, and a capped preregistered pilot — and pairs it with a catalog of thirty-two named failure modes, each naming the damaged parameter, the instrument that confirms it, and its price. **The following are the results this paper contributes, stated so that they may be quoted rather than paraphrased.**

R1 — The increment rule, with a reserve, and the branch in which building is forbidden. Any requested per-period increment must satisfy **`r_t <= h_t`**, with a recommended reserve **`r_t <= 0.7 h_t`** stated explicitly as a reserve, since `r_max = h_t` is the extremal value and is attained **only at `U_t = 1`**, that is, only if every category is mutually orthogonal to every other. **On the `h_t < 0` branch no interaction geometry can prevent order from decreasing**, so the permitted moves are cutting `C`, raising `eta`, and shedding `Q`, and building is not among them. The 0.7 is a stipulated convention, not a derived quantity.

R2 — The worked refusal and the repair that made the request fundable. A request for **`r = 0.55`** rules per week was refused at **`h = −0.1720`**. Raising conversion from **0.5015 to 1.05** rules per nat supplied **+0.8828**, or **80.8 per cent** of the recovery, and cutting dissipation from **0.63 to 0.42** supplied **+0.2100**, or 19.2 per cent, bringing headroom to **`h = +0.9207`** — a swing of 1.093 rules per week achieved without adding a person, a category or any capital expenditure. The request then fitted inside the budget at a **reserve ratio of 0.597**, inside the recommended 0.7.

R3 — The shedding lever is worth `delta x q` and nothing more. A fourteen-unit shed at a weekly decay of **0.003090** returns **0.0433** rules per week; the same shed at `delta = 0.05` returns **0.70**, **a factor of sixteen**. **Shedding is decisive in fast-decaying systems and nearly worthless in durable ones**, and the multiplication that distinguishes them takes five seconds. The shed was computed and declined; against the conversion lever's +0.8828 in the same system it returns one twentieth as much.

R4 — The category lever, computed and declined net. Adding a sixth discipline gives a ceiling gain of **`eta(ln 6 − ln 5) = +0.1914`** rules per week against **five added communication paths costing +0.0714**, for a net headroom gain of **+0.1200**. A genuinely distinct candidate returns **+0.1578** in realised step net of path cost against a relabelled candidate's **+0.1045** — **a 34 per cent haircut on the headline benefit, which is the price of relabelling.** The conversion lever returns an order of magnitude more per full-time equivalent on either consistent basis: **18.4 times** on a headroom basis and **10.0 times** on a step basis.

R5 — The frozen-denominator fault, priced. A ward declaring five handover categories with two attending plans against **`ln 2 = 0.693147`** and not `ln 5 = 1.609438` — **a ceiling 132 per cent higher than the one it owns.** The general form is that `N` is read from an administrative partition rather than a capability partition, and the correction is to count only categories that can independently detect, fix or veto something no other can.

R6 — The realisation channel, and the concavity penalty it recovers twice. Cutting the low-participation share from **32.7 per cent to 12 per cent** raises the **mean of `F`** — the quantity that accumulates, and not `F` of the mean `U` — from **1.1498 to 1.3241 nats**, and the realised step from **+0.4381 to +0.6211** rules per week, **with the promise breaking above a share of about 20 per cent.** The concavity penalty `F(mean U) − mean F` falls from **0.0761** nats to **0.0354**, so **regularising a variable geometry pays twice: once by raising the mean and once by shrinking the penalty the variability was charging.** The second gain alone is worth 0.0427 rules per week and is invisible to anyone reasoning about the average board rather than about the distribution of

boards. **This lever raised the realised step by 42 per cent and raised headroom by exactly nothing**, which is the operational content of the distinction between creating headroom and realising it.

CR (counter-result) — Positive headroom is not health. A hospital ward has `h = +1.686` **carried-patients per week and is losing 0.755 per week**, because it realises only **46.04 per cent** of a ceiling that a frozen denominator has already halved. **Headroom is computed on the ceiling `ln N`; the realised step is computed on `F(U)`; the two are independent, and a diagnosis that stops at a positive gauge is exactly wrong.** The catalog names that misdiagnosis. And the same case carries a second disclosure of the same kind: **the intervention's own predicted fixed point of 71.46 exceeds the ward's 34-bed census, so the linear law is outside its range there**, and the honest output is a saturation date between weeks 4 and 5 rather than an equilibrium.

What is established and what is not. The seven gates and the thirty-two entries are internally consistent and follow from two definitions, one cited range result, two elementary mathematical facts and one cited combinatorial fact; the compounding chain of Section 10.3 is arithmetic throughout and can be audited line by line. **That establishes a decision rule and nothing about institutions.** Both worked systems are constructed. The reflexive specification underlying the gauge is not identified from observational data and no remedy is adopted here. The one preregistered confirmatory test of the underlying law in this corpus did not meet its own declared margin. **The procedure's value, if it has one, is that it makes a refusal computable and states in advance what would show it to be wrong.**

16.7.1 Declarations

Data and code availability

No external dataset was collected for this paper. All numerical results are computed from the stated parameters under the model developed in the text, and every table can be re-derived from the equations and parameter values printed beside it. The development review office of Section 8.1 is a declared composite whose audit series is a simulation with declared generating parameters; the hospital ward of Section 8.2 is a constructed clinical case. **Neither is a field observation and neither describes a real organisation.** The figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper and the model equations of Sections 3 and 4.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, produced the worked examples and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. Both worked systems are constructed; the hospital ward of Section 8.2 is an invented case and no patient record, staff record or clinical observation of any kind was consulted. The consent constraint of Sections 3.6 and 4.12 — that no intervention removing a participant's ability to

decline is within this procedure's scope — is a design constraint on prospective work and is **not** a report of any study performed.

Preregistration

This paper reports no preregistered test of its own; its results are derivations and its worked cases are constructed. The preregistered protocol for testing the storage law against a real system, together with the confirmatory application that **failed** its second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared margin of 10 per cent — is reported in Curzi [2026i]. **The office case of Section 8.1 is worked under the accounting-frame restriction that follows from that failure: headroom may be used as a budget constraint and a capped pilot may be run, and no claim may be made that the model predicts the pilot's outcome.** Any applied claim resting on this paper inherits that status.

16.8 Appendix A — The intervention specification schedule

Thirty-two clauses, labelled **IS-1** to **IS-32**. The source labels them S26.1 to S26.32 against a chapter number that no longer exists.

IS-1 No intervention shall be designed before a complete survey and a completed preregistration evaluation [Curzi, 2026i].

IS-2 No intervention shall be designed on a system whose repackaging tests did not all pass.

IS-3 `h` shall be computed and its sign recorded before any design work begins.

IS-4 `h` shall be recomputed after every phase of Section 4.10 that changes `eta`, `C`, `Q`, `N` or `F(U)`.

IS-5 Where `h < 0`, no commitment to increase `Q` shall be made. Permitted actions are limited to the remedies of Section 4.4 and to measurement corrections.

IS-6 Where `h = 0`, scope shall not be added. Substitution at constant scope is permitted.

IS-7 Any requested increment shall satisfy `r <= h`. A request exceeding `h` shall be refused in writing, with the number.

IS-8 `r` shall include a stated reserve. The recommended default is `r <= 0.7 h`, and the reserve shall be reported as a reserve.

IS-9 `S_req = (delta Q + C + r)/eta` and `U_req = (exp(S_req) − 1)/(N−1)` shall both be computed and reported.

IS-10 `U_req` shall be compared against the measured capacity. **The feasibility test against realised geometry shall be made on `mean F(U_t) >= S_req` and not on the mean of `U`, and the concavity penalty `F(mean U) − mean F` shall be reported.**

IS-11 Failure A (budget) and Failure B (geometry) shall be distinguished explicitly in any report of infeasibility.

IS-12 The realisation deficit `h − step = eta(ln N − F(U))` shall be computed and reported.

IS-13 Levers shall be ranked by cost per unit of effect for the specific system, using the system's own measured costs. **Deviations from the default order `eta, C, u, N` shall be justified**, and an inversion shall be reported rather than suppressed.

IS-14 The value of any proposed shed shall be computed as `delta x q` before the shed is proposed.

IS-15 A shed shall be distinguished from a measurement correction in writing.

IS-16 Any proposal to raise `N` shall test the candidate category against the merge rule (`u < 0.10` merges) using the probe set, before commitment.

IS-17 Any proposal to raise `N` shall state the added coordination paths and their cost in proxy units, and shall show the net effect on `h`.

IS-18 Sequencing shall follow Section 4.10. No scope-adding phase shall precede a deficit-reducing phase.

IS-19 Slack shall not be removed as part of a cost-reduction programme without a queueing analysis.

IS-20 A reversibility statement answering RV1 to RV5 shall exist before commitment.

IS-21 Interventions classified as irreversible shall not proceed.

IS-22 Partially irreversible components shall be separately named, separately justified and separately consented.

IS-23 The reversal shall be executable by a named person who is not the person who authorised the intervention.

IS-24 Every affected person shall have a stated mechanism to decline without penalty. If none exists, the intervention is out of scope and shall not proceed.

IS-25 The three-posture choice — engage, ignore, suppress — shall be made explicitly for every unplanned practice encountered, and recorded. **Defaulting to ignore without recording it is a specification violation.**

IS-26 Suppression that removes the ability to decline shall not be used.

IS-27 Pilots shall have a downside capped in proxy units, bounded scope and bounded duration.

IS-28 Pilot predictions shall not begin before the measured conversion lag has elapsed.

IS-29 Promises shall be stated as increments per period. `Q_inf` may be reported and shall not be promised.

IS-30 `tau` and `3 tau` shall be stated in calendar units at the commitment meeting, with `3 tau` labelled settling time and not a delivery date.

IS-31 Trip-wires shall be declared before the intervention starts, with thresholds and automatic actions, and shall include a realised-geometry trip-wire set on `mean F(U_t) >= S_req`.

IS-32 A decision not to intervene shall be recorded as a decision, with the six numbers, the sign of `h`, the reason, and the observation that would reopen it.

16.9 Appendix B — The intervention specification form and the reversibility statement

The form is reproduced because the arrangement of its fields is part of the specification: the sequence section has a headroom slot for every phase because every phase has a headroom value (PF-10), and the increment section places the operative test on the mean of `F` immediately below the legibility-only value of `U_req` (IS-10).

```
============================================================
INTERVENTION SPECIFICATION                  Date: __________
System: ......................  Analyst: ..................
============================================================
---- 1  ADMISSIBILITY ------------------------------------
1.1  Survey ref/date ................  complete? [ ]
1.2  Repackaging tests 1–4 all PASS?  [ ]  (if not, STOP)
1.3  Preregistration ref/date .......
1.4  FR1 [ ]P [ ]F   FR2 [ ]P [ ]F   FR3 [ ]P [ ]F
1.5  Evidence status:
     [ ] validated model    [ ] accounting frame only
     [ ] inadmissible -> STOP

---- 2  THE GAUGE ----------------------------------------
2.1  eta = ____   ln N = ____   eta*lnN = ____
2.2  delta = ____   Q = ____   delta*Q = ____
2.3  C = ____
```

```
2.4  h = ____                REGIME [ ]h<0 [ ]h=0 [ ]h>0
2.5  mean F over the record = ____    step = ____
2.6  realisation deficit = h - step = ____
2.7  capacity u = ____   mean realised U = ____

---- 3  THE INCREMENT (h>0 only) ------------------------------
3.1  Requested r = ____      r <= h ?  [ ]
3.2  Reserve: r / h = ____   (target <= 0.7)
3.3  S_req = (delta*Q + C + r)/eta = ____
3.4  S_req <= ln N ?  [ ]        (else FAILURE A)
3.5  U_req = (exp(S_req)-1)/(N-1) = ____   (legibility only)
3.6  U_req <= capacity u ?  [ ]  (else FAILURE B)
3.7  measured mean F(U_t) = ____  >= S_req ?  [ ]
     (THIS is the operative test; if not, realisation work required)
3.8  concavity penalty  F(mean U) - mean F = ____

---- 4  LEVERS ------------------------------------------------
lever | change | effect on h | cost (own units) | h per unit cost
------+--------+-------------+------------------+----------------
eta   |        |             |                  |
C     |        |             |                  |
shed Q|        | delta*q =   |                  |
u cap |        | (realises)  |                  |
u real|        | (realises)  |                  |
N     |        |             |                  |
4.1  Chosen levers, in order: ..............................
4.2  Justification if the default ranking is not followed:
     ..........................................................

---- 5  SEQUENCE ----------------------------------------------
Phase 1 stop the leak .....................  h after = ____
Phase 2 correct the proxy .................  h after = ____
Phase 3 shed (if delta*q worth it) .........  h after = ____
Phase 4 realise ...........................  step after = ____
Phase 5 RE-MEASURE h ......................  h = ____
Phase 6 commit r = ____
Phase 7 raise capacity u ..................
Phase 8 raise N ...........................

---- 6  TIME --------------------------------------------------
6.1  Conversion lag m = ____ periods
6.2  Increment r = ____ per period, judged over ____ periods
6.3  tau = 1/delta = ____ periods = ____ calendar
6.4  3*tau = ____ periods = ____ calendar   (SETTLING, NOT DELIVERY)
6.5  Q_inf = ____   (reported, NOT promised)   [ ] initialled
6.6  Hard cap (census, floor area, statute, budget) = ____
     Which binds, the cap or the fixed point? ..............

---- 7  PILOT -------------------------------------------------
7.1  Scope bound ..........................................
7.2  Duration bound ........................................
7.3  Downside cap in proxy units = ____
7.4  Admissibility C <= (1-delta)Q + eta*S holds under worst case? [ ]
7.5  Comparison [ ] matched unit [ ] prior period [ ] none (state)
```

```
---- 8  TRIP-WIRES ------------------------------------------
wire  | threshold          | automatic action
------+--------------------+-------------------------------------
geom  | mean F < S_req     |
C     | > ____             |
h     | <= 0               |
dQ    | > eta*lnN          |
harm  | ____ (non-proxy)   |
8.1  Review cadence: every ____ periods

---- 9  POSTURE DECISIONS -----------------------------------
practice encountered | posture | why | parameter effect
---------------------+---------+-----+-----------------
                     |         |     |

---- 10 SIGN-OFF --------------------------------------------
Analyst .................  Sponsor ................  Date ......
================================================================

================================================================
REVERSIBILITY                                  Date: _________
================================================================
RV1 What returns to the prior state, exactly:
    ..........................................................
RV2 What does NOT return, EVER (be specific):
    [ ] people who leave        [ ] demolished fabric
    [ ] dissolved teams         [ ] public commitments
    [ ] trust                   [ ] elapsed time
    [ ] other: ...............................................
RV3 Authorised to execute the reversal WITHOUT further approval:
    name ............................  role ..................
    (must NOT be the person who authorised the intervention)
RV4 Time to reverse: ____ periods
RV5 Cost of the reversal: ____ proxy units.  Inside h ? [ ]
----------------------------------------------------------------
CLASS  [ ] fully reversible
       [ ] costly-reversible   (RV5 inside h required)
       [ ] partially irreversible (RV2 separately consented)
       [ ] irreversible  -> NOT ADMISSIBLE. REDESIGN.
----------------------------------------------------------------
CONSENT
C1  Affected groups: .....................................
C2  How informed: ........................................
C3  Mechanism to decline WITHOUT PENALTY: ................
C4  Does a declination mechanism exist?  [ ] Y  [ ] N
    If N -> OUT OF SCOPE. STOP.
C5  Parameter values under declination, computed: ............
----------------------------------------------------------------
Signature .................  Date ..........
================================================================
```

Field 6.6 and field C5 are additions of this paper rather than of the source. Field 6.6 makes DF-6 a required entry rather than a discipline, because the ward case of Section 8.2 shows that the fixed point and the hard cap can differ by a factor of two and that only one of them binds. Field C5 makes Section 4.12's rule operative: **a design whose parameters have no defined value under declination has not actually offered one**, so the value is computed on the form

rather than asserted in prose.

16.10 Appendix C — The catalog use specification

Ten clauses, labelled **CS-1** to **CS-10**. The source labels them S27.1 to S27.10.

CS-1 — Enter by symptom, not by name. Entry names are memorable and therefore misleading. Enter through the parameter that can already be measured, or through a symptom index, and not by scanning the names for one that sounds apt.

CS-2 — Run the method faults first. Before any Class I to V diagnosis, clear FM-29 (order proxy), FM-23 (unbooked cost), FM-28 (Goodhart capture) and FM-30 (unbadged postulate). **A method fault invalidates every other reading, and three of the four are cheap to check.**

CS-3 — Measure the six parameters before naming anything. The survey is the entry condition for the catalog, not an alternative to it [Curzi, 2026i]. Required: `Q` and its proxy; `delta` from decay under zero input; `eta` from attributable change per unit of surplus; `N` by capability partition; `u` by two instruments; `C` as a residual.

CS-4 — Compute both master numbers. Headroom `h = eta ln N − delta Q − C`, and the realised fraction `S / ln N`. **Record them together; Section 6.3 explains why either alone will mislead.**

CS-5 — Validate the survey against an observed change. Compute the predicted one-step change `eta S − delta Q − C` and compare it to the measured change over a recent period. **If they disagree by more than about 10 per cent of the ceiling term, the survey is wrong and no diagnosis from it is admissible.** This single check catches most instances of FM-23 and FM-29, and Section 8.2.2 exhibits it passing at a residual of 0.13 per cent of the ceiling term.

CS-6 — Expect three to five co-occurring faults, not one. Real systems present in clusters, and the clusters are structured: a Reverberant Room causes a Head of Table to bind harder; a Ceiling Chase causes a Quadratic Committee; a Founder Battery hides an Invisible Infrastructure. Name all of them, then rank.

CS-7 — Rank by lever cost, not by severity. Conversion is usually cheapest; dissipation next; non-parallelism next; category count most expensive and least rewarding because it is logarithmic. Shedding is frequently the cheapest of all and is almost never considered.

CS-8 — If `h < 0`, the procedure forbids building. CC-T3. The permitted moves are cutting `C`, raising `eta`, and shedding `Q`; growth is not on the list until `h` is positive.

CS-9 — Predict a number, and state what would falsify it. Every diagnosis terminates in a predicted trajectory with dated checkpoints and a written falsification condition. All four baselines — persistence, direct input, matched autoregression, and the shuffled-input control, calibrated by simulation on the study's own design — are required. **A framework that cannot fail cannot be trusted.**

CS-10 — Apply the consent test to every proposed intervention. FM-32. In writing. Before costing.

Chapter 17

Five Worked Systems and the Essence Bridge: Testing Scale-Freedom in the Storage Law and Identifying the Effective Count as a Pitch Index

Abstract. Five systems spanning three orders of magnitude in headcount — one person's desk, a twelve-person research group, a public reading room, a distributed software project, and a twenty-two-household co-housing block — are surveyed with one procedure and one set of parameters. Every one was losing stored order at survey and no two for the same reason: the research group ran at 41.64 per cent of its ceiling, the reading room at 28.59 per cent for a physically identifiable cause (a 2.5223 s occupied reverberation time), the software project at 69.28 per cent with headroom $h = -0.654$ modules per month driven by cost, and the co-housing block at 91.03 per cent of ceiling and still at $h = -0.273$ because conversion was 0.926 units per nat and twenty-two-household consensus carries 231 communication channels, 74.73 per nat of ceiling against 6.21 at five. In every case the predicted one-step change reproduced the observed change to within 1.4 per cent of the ceiling term $\eta \ln N$. The paper then develops a separate postulated law, $f(n) = k\,n^2$ with the defined constant $k = (9/8)^2/\varphi^2 = 0.483425732988414$, and proves that every interval is a ratio of squares so that k cancels, and that a set indexed $1..N$ has exactly $\lceil N/2 \rceil$ pitch classes given by the odd part of the index. A badged identification takes the effective count as a cross-disciplinary index, giving $2/\ln 2 = 2.885390081777927$ octaves per nat. The scale-freedom claim is stated as an invariance with its transformation, its five tested scales, and what would break it; the reflexive specification underlying the survey is not identified from observational data, so the paper claims a modelling position rather than a measured result.

Keywords: scale-free claim; cross-case comparison; binding constraint; essence index; pitch class; odd part; effective count; octave conversion

Series: B. Engineering the Surplus **Paper:** CCP-017

17.1 1 Introduction

A model that claims to describe accumulation at more than one scale owes two demonstrations before any of its applications can be assessed. The first is that a single procedure — the same observables, the same estimators, the same admissibility rule — actually runs at every scale claimed. The second is an account of what would count as failure, because a model form that survives rescaling only by being re-fitted at each scale has not been shown to be scale-free; it has been shown to be flexible, which is a different and much weaker property [Cartwright, 1999].

This paper supplies both for the storage law $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$ with $S_t = \ln(1+(N-1)U_t)$, whose formal properties as a first-order discrete filter are established with its contraction argument, its pole and its time constant in Curzi [2026a], whose headroom gauge $h = \eta \ln N - \delta Q - C$ is established in Curzi [2026b], and whose five required operational choices — state encoder, interaction observable, order observable, time scale, cost observable — are fixed as a measurement protocol in Curzi [2026i]. Five systems are surveyed with that one procedure: one person's desk, a twelve-person research group, a public library reading room, a distributed software project, and a twenty-two-household co-housing block. They span three orders of magnitude in headcount and four in the size of the quantities they store.

The survey's admissibility rule is what makes the exercise more than an illustration. **No diagnosis is admitted from a survey that cannot first reproduce a change in stored order that has already been observed.** Every one of the five clears that gate, and the residual is small in a stated unit: the largest is 1.39 per cent of the ceiling term $\eta \ln N$, at the research group, where a predicted -3.3872 assets per quarter is set against an observed -3.500. At the desk the prediction is -0.250474 threads per week against an observed -0.250; at the library, -0.996371 regulars per month against -1.000.

Three results follow, and they are of quite different epistemic kinds, which the paper labels throughout rather than leaving to the reader.

The first is a comparative finding. All five systems were losing stored order at survey and no two for the same reason. The research group ran at 41.64 per cent of its ceiling with positive headroom; the reading room at 28.59 per cent for a physically identifiable cause, a 2.52-second occupied reverberation time computed from the surface absorption schedule by Sabine's relation [Sabine, 1922]; the software project at 69.28 per cent with headroom of -0.654 modules per month driven by cost; the co-housing block at 91.03 per cent of ceiling and still at -0.273 commons-units per month because its conversion efficiency was 0.93 units per nat and twenty-two-household consensus carries $22 \cdot 21/2 = 231$ communication channels. **Headroom and the realised fraction $S/\ln N$ are independent readings, and a diagnosis stopping at either one is wrong in at least two of the five cases.**

The second is a scale-freedom claim, and this paper states it as an invariance rather than as a slogan. Section 11.7 gives it in full: the claim is that the *functional form* of the recursion is invariant under a change of the unit in which stored order is counted and under a change of the period, with (δ, η, N, C) carrying the whole of the difference between systems. That is a statement about model form, it is falsifiable separately at each scale and never in general, and it has been tested here at exactly five scales. What would break it is stated: a system at which the one-step prediction fails the admissibility rule, or at which a term absent from the recursion — saturation is the leading candidate, and it binds in two of the five cases — is required to reproduce the observed change.

The third is a bridge into a second formalism, and it is the paper's largest exposure. An essence law $f(n) = k\,n^2$ is analysed, its published constant $k = (9/8)^2/\varphi^2$ shown to be exactly 0.483425732988414 — a *defined* constant, so the fifteen figures are legitimate, and Section 11.8 says so — and two structural theorems proved about it: that every interval $f(m)/f(n)$ equals m^2/n^2, so the constant cancels from every relationship the law produces, and that the pitch class of index n is the odd part of n, so a set indexed $1..N$ realises exactly $\lceil N/2 \rceil$ classes. The bridge then identifies the framework's effective count with the essence index, giving $2/\ln 2 = 2.885390081777927$ octaves per nat. **The identification is a modelling decision and is badged**

as one everywhere it is used. The two theorems are unconditional algebra; the law they are theorems about rests on three postulates that this paper states, labels, and does not defend.

Four things are new here relative to the surrounding literature rather than relative to the framework. First, a five-case cross-scale survey in which the binding constraint is identified by computed effect rather than by salience, and in which the constraint named by the system's own members is wrong in all five cases. Second, the demonstration that the reference constant k cancels from every interval, which is a demotion of the golden section within this formalism and is presented as one — the scepticism literature on φ [Markowsky, 1992, Fischler, 1981, Livio, 2002] is cited at the point where the constant is introduced, not in a footnote. Third, the odd-part class rule and the resulting exact count $\lceil N/2 \rceil$, which converts an aesthetic question into an arithmetic one. Fourth, a forensic reconstruction of the source document's own published table, which shows that the whole of its residual bias is explained by a fifteen-digit stored value of φ, and which identifies **seven of its forty-four published composite values as errata**.

Sections 2 and 3 state the hypotheses and the method, including the identification problem the survey cannot solve. Sections 4 to 6 report the five systems, the abstracted specification, and the faults encountered. Sections 7 and 8 develop the essence law and the bridge. Sections 9 to 13 report results, discuss what does and does not follow, verify the chain from minimal postulates, state the limitations, and conclude.

17.2 2 Hypotheses

The paper contains two kinds of claim and they cannot be tested the same way. The theorems of Sections 7.6, 7.7 and 8.4 are statements about an equation and are refuted only by an error in the algebra; they are labelled as derivations where they appear and they are not hypotheses. What can be falsified is the claim that the recursion describes real accumulating systems, and that the survey procedure recovers its parameters. The hypotheses below separate the two. Each names the observation that would defeat it, and where a claim could not be defeated by any observation it is restated as a definition, a postulate or a modelling choice and labelled as such.

> **H1 — Cross-scale reproduction.** A five-parameter survey conducted under one fixed protocol reproduces an independently observed one-step change in stored order at every scale at which the protocol can be run. **Predicted:** the residual $|(\eta S - \delta Q - C) - \Delta Q_{\text{obs}}|$ is small relative to the ceiling term $\eta \ln N$ — the protocol's own admissibility threshold is 10 per cent — at a headcount of one and at a headcount of twenty-two alike. **Falsified by:** a system that passes the order-proxy lag test and the parameter-estimation protocol and whose observed change the survey nonetheless fails to reproduce within that threshold; or a systematic dependence of the residual on scale.
>
> **H2 — The binding constraint is not the salient one.** The term that limits a system is identified by computed effect and is not, in general, the term the system's participants name. **Predicted:** across independently chosen systems, the parameter whose single-lever improvement most reduces the deficit differs from the complaint recorded at survey. **Falsified by:** a sample of systems in which the named complaint and the computed binding term coincide at a rate no better than the diagnosis would achieve by naming the same term every time.
>
> **H3 — Headroom and realised fraction are independent readings.** The gauge $h = \eta \ln N - \delta Q - C$ contains no term in u, so a system may have large positive headroom and be losing stored order. **Predicted:** systems exist with $h > 0$ and $\Delta Q < 0$, and in such a system an intervention raising u moves the one-step change substantially while leaving h almost unmoved. **Falsified by:** an observed coupling between u

and h not mediated by Q, C, η or N — which would mean the gauge has been misderived.

H4 — The coordination surface is quadratic while the ceiling is logarithmic. Adding participants to a body that decides as a whole adds pairwise reconciliation cost as $n(n-1)/2$ while adding at most $\eta \ln$ of a count to the ceiling, so beyond a computable size the flat body cannot be coordinated [Brooks, 1975, Conway, 1968]. **Predicted:** governance hours in a whole-group body scale with the channel count rather than with headcount, so a nested decomposition cutting channels from 231 to 55.5 cuts hours by approximately the same 24.03 per cent ratio. **Falsified by:** measured governance hours that fall with headcount but not with channel count after a nesting reorganisation — which is the outcome the co-housing case explicitly preregisters as its second falsification condition.

H5 — Time to arrival is set by the retention coefficient and is not negotiable by effort. The time constant is $\tau = 1/\delta$ periods regardless of the size of the intervention [Curzi, 2026a]. **Predicted:** the co-housing block, with $\delta = 0.023341$ per month, requires $3\tau = 128.5$ months to settle within five per cent of its fixed point, while the library, at $\delta = 0.152021$, requires 19.7 months — and the same low δ that produced the commons' resilience under nine months of neglect is what produces the decade. **Falsified by:** a measured settling time that scales with the size of the intervention rather than with the measured δ.

Three further claims in this paper are not hypotheses and are labelled accordingly.

The **scale-freedom claim** is a *modelling postulate about model form*. It asserts that the same difference equation is a useful description at desk, room, team and building, with only (δ, η, N, C) changing. It is falsifiable at each scale separately, never in general, and Section 11.7 states it as an invariance with its transformation group, its tested range, and what would break it.

The **essence law** $f(n) = k\,n^2$ rests on three commitments that no observation in this paper bears on: that essence types have a frequency at all; that the law is square in the index; and that the constant is $(9/8)^2/\varphi^2$ rather than any other positive number. All three are **postulates**, and the third is shown in Section 7.6 to be inert — it cancels from every interval the law produces.

The **bridge identification** $n_{\text{eff}} := N_{\text{eff}} = e^F$ is a **modelling identification**. What is *shown* is that the effective count uniquely satisfies six requirements an index must satisfy, that the substitution is exact, and that it is continuous with the within-discipline law at every integer. That is a warrant. It is not a derivation, and Section 8.11 says so in the same words.

17.3 3 Method: the survey, the estimators, and what the survey cannot recover

17.3.1 3.1 What a worked case must contain

Seven steps, in this order, with no substitutions:

```
1. SURVEY          five operational choices + six parameters, measured
2. VALIDATE        predicted one-step change must reproduce the observed one
3. COMPUTE         h = eta ln N - delta Q - C, and the realized fraction S / ln N
4. BIND            identify which term is binding, by computed effect, not by salience
5. DESIGN          intervention, levers ranked by cost, consent test applied to each
6. PREDICT         a dated trajectory with a fixed point and a stated validity range
7. FALSIFY         what observation would refute it, written down before starting
```

Step 2 is the step most often skipped and it is the one that makes the rest admissible. Step 7 is what distinguishes the exercise from consulting.

17.3.2 3.2 The equations, restated for use

From the uniqueness theorem for the surplus functional [Curzi, 2026] and the storage law with its headroom gauge [Curzi, 2026a,b], restated exactly:

```
S_t   = F(U_t) = ln(1 + (N−1)U_t)          surplus input, 0 <= S_t <= ln N
N_eff = 1 + (N−1)u                         effective count after redundancy audit
Q_{t+1} = (1 − delta)Q_t + eta S_t − C_t   the storage law
h_t   = eta ln N − delta Q_t − C_t         headroom (the master gauge)
Q_inf = (eta F(u∗) − c∗) / delta           steady state, delta > 0
tau   = 1 / delta                          time constant, in periods
growth  <=>  eta S_t > delta Q_t + C_t     one−step growth condition
```

Badged, in the same sentence as it is used: axiom S3, the affine effective count, is a *design axiom* — chosen, not derived — and every `N_eff` below inherits that choice. Badged also: that a research group, a reading room, a code repository, and a housing block are each described by a single scalar Q is a modelling postulate, not a theorem. The surplus and storage results [Curzi, 2026,a] are theorems about the law. They are not statements about these five systems.

17.3.3 3.3 The estimator for u

All four cases estimate u the same way, so the method is stated once. Take an alignment measure rho between the members' or categories' contributions on a live question — correlation of signed error, overlap of independently listed concerns, rank correlation of independently stated priorities. Treat rho as the cosine between unit state vectors. Then, directly from `u = 1 − (x·y)^2`:

```
u = 1 − rho^2
```

Where the two-source decomposition [Curzi, 2026d] is available, `u = u_cross + u_div` is estimated with two separate instruments and both numbers are reported, because a single-instrument diversity reading is Failure Mode FM-08.

17.3.4 3.4 The honesty clauses, live in every case below

- The surplus functional as deployed is a **lower bound** on the ensemble's information. Every predicted trajectory is a floor, not a forecast of the best possible outcome.
- The linear law has **no saturation term**. Where a computed `Q_inf` exceeds a hard physical or organizational cap, the model is outside its range and must be replaced there. This happens in two of the five cases below and it is flagged both times.
- **No thermodynamic negentropy law is proved.** Surplus is not identified with Shannon entropy.
- Every external result cited — Woolley et al. [2010], Centola [2010], Ostrom [1990], Sabine [1922], Brooks [1975], Whyte [1980] — is real, and the **transfer** of that result to the system at hand is a modelling assumption, badged at the point of use.
- No intervention below removes anyone's ability to decline. Where a candidate intervention would have, it was cut, and the cut is recorded.

17.3.5 3.5 Identification: what no survey in this paper can recover

The framework's central formal claim is reflexive. The group influences the member and the members constitute the group, so the quantity u that drives the surplus term is itself a property of an ensemble whose members are responding to that ensemble. Written out, the specification is $T(x) = W(x, M(x))$: an outcome for a unit depends on the unit's own characteristics and on the mean of the group it belongs to.

That is exactly a linear-in-means reflexive specification, and its endogenous social effect is not identified from observational data. Manski [1993] proves it: the endogenous effect cannot be separated from contextual effects, nor from correlated unobservables, however much data is collected. Shalizi and Thomas [2011] reinforce the point for network data specifically — influence and selection are generically confounded, and the confounding does not vanish with sample size. Lyons [2011] makes the same objection concretely against a widely cited contagion result.

The five surveys in this paper adopt no remedy, and the claim they support is therefore a modelling position rather than an identified empirical result. The available remedies are randomisation, an instrument, partial-population variation, and stipulating the reference group while defending the stipulation; the last is the template of Brock and Durlauf [2001] and is the only one available to a survey of an existing organisation. **None of them is free, none of them is applied here, and no amount of agreement between predicted and observed change repairs the deficiency**, because the admissibility test in Section 3.1 compares a one-step prediction against an observed change and a correlated unobservable driving both is exactly what it cannot see. The corpus's own shuffled-input control measures the size of that blind spot: it passes 98.35 per cent of the time against a third variable driving both interaction and order, while detecting only 48.17 per cent of the clean case [Curzi, 2026i]. **The empirical audit cannot see repackaging.**

What the surveys below therefore establish is narrower than it looks, and it is worth stating in one sentence: *the arithmetic of the recursion reproduces five observed declines from five independently estimated parameter sets, and the causal reading of those parameters is a stipulation this paper does not attempt to identify.*

One further objection in the same family is recorded because it applies to the framework's own parameters and not only to its subject matter. **The Lucas critique** — that a decision rule estimated under one policy regime need not survive the regime's change [Lucas, 1976] — applies with unusual force here, because this framework asserts that systems change their behaviour when they are modelled and observed [Merton, 1948, Hacking, 1995], and then treats (δ, η, N, C) as stable while an intervention derived from them is executed. Either the parameters are stable under the framework's own predicted response, or the framework undercuts itself. **This paper adopts the first horn as a modelling assumption and does not defend it**; every case below holds at least one parameter fixed that the intervention should plausibly move, which is conservative for the prediction but says nothing about whether the held parameter is in fact stable. MacKenzie and Millo [2003] documents a reflexive fixed point actually observed in the field, with dates and price data, and is the existence proof the framework needs; it is not evidence that any particular parameter here is stable.

17.3.6 3.6 Scale type, and what licenses the arithmetic below

Every quantity in the recursion is divided, multiplied or logged somewhere in the sections that follow, and those operations are licensed only on a ratio scale with a true zero [Stevens, 1946]. The scale type of each is therefore declared before it is used.

Quantity	Scale type claimed	What licenses it
Q, stored order	**ratio**	a count of countable artefacts with a true zero — assets, regulars, green modules, commons-units, resumable threads. The count is licensed; whether the counted items are *equivalent* is not tested
δ, retention loss	**ratio**	a ratio of two counts of the same quantity, so dimensionless with a true zero
η, conversion	**ratio**	a ratio of a count to a nat, inheriting both
C, dissipation	**ratio**	a count per period, priced at a measured yield
N, categories	**absolute count**	a partition size
u, non-parallelism	**stipulated ratio, not demonstrated**	$u = 1 - \rho^2$ is a ratio-scale *transform of a correlation*, but the underlying attribute has not been shown to be quantitative at all
n, essence index	**ordinal**	a position in an ordering; see Section 7.9

The last two rows are where this corpus is most exposed, and the exposure is named rather than managed. **Michell [1997] argues that psychology routinely assumes its attributes have quantitative structure without testing for it, and that is the single strongest objection to this framework.** The formal test is conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971], and its cheaper practical relative is a Rasch demonstration of specific objectivity [Rasch, 1960, Andrich, 1988]. **Neither has been run for u, and until one is, every ratio-scale operation performed on u in this paper is a stipulation rather than a licensed operation.** The corpus states the requirement and demotes its own moment equation when the requirement fails [Curzi, 2026t]; the demotion is the right instinct and it has not been applied to u.

The essence index is worse, and Section 7.9 treats it at length: **an ordinal position does not license squaring.** $f(n) = k\,n^2$ is an arithmetic operation on an ordinal, which is exactly the operation Stevens's taxonomy forbids. The paper's response is not to defend the operation but to bound its consequences — Section 7.9's permutation test measures how much a defensible re-ordering moves the output, and reports a maximum displacement of 364.99 cents on the worked twelve-element set, which is the honest size of the exposure.

17.4 4 Five worked systems

17.4.1 4.1 Case A · A twelve-person research group

The system. A university-affiliated laboratory: one principal investigator, three postdoctoral researchers, six doctoral students, one research engineer, one coordinator. Funding stable. The complaint, as stated by the PI: *"We are busier than we have ever been and we are producing less."*

4.1.1 The survey

The five operational choices required by the survey protocol [Curzi, 2026i]:

Choice	Setting
State encoder	Each researcher's methodological position, from a 12-item forced-choice instrument
Interaction observable	Cross-person working sessions in which a live analytic decision was made
Order observable Y	Count of **independently re-runnable research assets**
Time scale	One quarter (13 weeks)
Cost observable	Person-hours not converting interaction into assets, taken as a residual

An asset counts toward Y only if someone **other than its author** has successfully re-run it in the trailing two quarters. `Q_0 = 46` assets.

Order-proxy check [Curzi, 2026i] — is Y a repackaging of U? No. Run the lag test: during a two-quarter funding hiatus with near-zero cross-person interaction, assets fell 52 -> 41. Y decayed with a time constant rather than tracking the pause instantaneously. Y is stored order.

delta. From that same hiatus, `S ~ 0`:

```
delta = 1 − (41/52)^(1/2) = 0.112046 per quarter
half-life = ln 2 / (−ln(1 − 0.112046)) = 5.833 quarters  (about 18 months)
```

N. Categories defined by capability, not by title — *what can this category independently detect, fix, or veto?* Four staffed: bench assay, imaging, statistics/ML, clinical translation. `ln 4 = 1.386294`.

u. Independent-estimate-before-discussion on 20 live analytic questions; mean pairwise correlation of signed error `rho = 0.86`. Seven of the twelve trained under the same advisor.

```
u     = 1 − 0.86^2 = 0.260400
N_eff = 1 + 3(0.2604) = 1.781200
S     = ln(1.781200) = 0.577287 nats     ->  41.64% of ln 4
```

Twelve people are purchasing the surplus of 1.78.

eta. Over eight quarters, 27 assets were attributable to specific logged working sessions against a cumulative `S = 8 x 0.577287 = 4.6183` nats.

```
eta = 27 / 4.6183 = 5.846309 assets per nat
```

C. Residual method [Curzi, 2026i]. Twelve people at 40 h/week x 13 weeks = 6,240 person-hours per quarter. Coordination — scheduling, reporting, compliance, internal review, meeting overhead — tallies at 17.4%, or 1,087 h/quarter. Over the eight estimation quarters the lab produced 61 durable assets from 41,234 productive person-hours:

```
yield = 61 / 41,234 = 0.00147936 assets per person-hour
C     = 1,087 x 0.00147936 = 1.608069 assets per quarter
```

4.1.2 The two master numbers, and validation

```
eta ln N  = 5.846309 x 1.386294 = 8.104705
delta Q   = 0.112046 x 46       = 5.154136
C         =                       1.608069

h = 8.104705 − 5.154136 − 1.608069 = +1.342500 assets/quarter
S / ln N  = 0.577287 / 1.386294   =  41.64%
```

Validation (step 2). Predicted one-step change:

```
eta S − delta Q − C = 3.375000 − 5.154136 − 1.608069 = −3.387205 assets/quarter
```

Observed: assets fell 60 -> 46 over the preceding four quarters, or **-3.500** per quarter. Residual **+0.113**, about 1.4% of the ceiling term. The survey is admissible.

4.1.3 The binding constraint

Headroom is **positive** and the lab is losing three and a half assets a quarter. This is the misreading the corpus failure catalog exists to prevent [Curzi, 2026j]: `h` is the ceiling, `u` decides how much of it is collected, and the laboratory collects 41.64 per cent.

Break-even requires:

```
S* = (delta Q + C) / eta = 6.762205 / 5.846309 = 1.156662 nats
```

Against a ceiling of `ln 4 = 1.386294`, that is feasible — but it demands `N_eff = 3.179`, hence `u = 0.7264`, hence `rho = 0.523`. Moving twelve researchers from 86% error-correlation to 52% is a large ask and, alone, leaves no margin whatever. **Single-lever solutions all fail here**, and the arithmetic says so before anyone tries one:

Single lever	Resulting one-step change
Baseline	**-3.387**
u alone: 0.2604 -> 0.6156	-0.646
Shed Q alone: 46 -> 34	-2.043
delta alone: 0.1120 -> 0.075	-1.683
C alone: 1.608 -> 1.213	-2.992
N alone: 4 -> 5	-2.590

Every one is still negative. The binding constraint is **the joint burden of delta Q and a parallel roster** — FM-17 (Maintenance Trap) co-occurring with FM-01 (Parallel Wall), with FM-02 (Pseudo-Diversity) underneath: the lab is demographically varied and methodologically single.

Two further readings, both from the survey rather than from opinion.

The PI's own proposal was refused by arithmetic. The stated plan was to recruit three more doctoral students. Twelve people already carry `n(n−1)/2 = 66` communication channels; fifteen carry 105, a 59% increase in coordination surface. The ceiling gain, if all three landed in genuinely new categories — they would not have; all three were to be recruited from the dominant lineage — would be at most `eta(ln 7 − ln 4) = 5.846 x 0.5596 = 3.27` assets/quarter, against a coordination increment the lab could not price. Recruiting into `u = 0.26` adds redundancy at full salary. **This is FM-07, the Ceiling Chase, and it is the standard response of a laboratory that has correctly noticed it is falling behind.**

Where the surplus is actually being lost. Decompose u with two instruments [Curzi, 2026d]. The lab measures `u_cross = 0.2100` between its four method categories and `u_div = 0.0504`

within them; they sum to the measured `u = 0.2604`. Bench assay and imaging disagree; two statisticians never do.

```
f(u_cross) = ln(1 + 3 x 0.2100) = 0.4886
f(u_div)   = ln(1 + 3 x 0.0504) = 0.1408
bound:  max = 0.4886  <=  f(u) = 0.5773  <=  sum = 0.6294        decomposition bound
```

Four fifths of the lab's realized surplus comes from crossing categories and one fifth from disagreement inside them. **The cross-disciplinary structure is real; the internal dissent is nearly absent.** Hiring more disciplines will not fix an organization in which each discipline has one permitted opinion, which is why A4.4 buys `u_div` — the independent-estimate protocol — before it buys the fifth category.

4.1.4 The intervention

Four moves, ranked by lever cost under the gated intervention procedure [Curzi, 2026j]. Each was checked against the consent test — what happens to a person who declines? — and one candidate was cut: mandatory rotation into a second methodological track, which had no opt-out and is therefore out of scope for this paper.

A4.1 — Shed Q (cheapest). Retire 12 assets that no non-author has re-run in four quarters to an explicitly unmaintained archive tier. `Q: 46 -> 34`.

A4.2 — Cut delta. Every surviving asset gets a named steward who is not its author, and a scheduled quarterly re-run. Target `delta: 0.1120 -> 0.075` (half-life 5.83 -> **8.891** quarters). The steward re-runs cost about 120 h/quarter and are booked as cost, not hidden.

A4.3 — Cut C. Coordination from 1,087 h/quarter to 700 h; adding the 120 h of steward work gives 820 h. `C = 820 x 0.00147936 = 1.213079`.

A4.4 — Raise u, and only then N. Two instruments: independent-estimate-first on every consequential analytic decision (this breaks cascades before they form, FM-03), and one recruitment into a fifth genuinely distinct category — a methodologist from an unrelated trade, hired **with a socket**: a named decision gate before option-closure, per FM-13. Target `rho: 0.86 -> 0.62`.

```
N = 5,  ln 5 = 1.609438
u     = 1 - 0.62^2 = 0.615600
N_eff = 1 + 4(0.6156) = 3.462400
S     = ln(3.462400) = 1.241962 nats   ->  77.17% of ln 5
```

eta is held at its measured 5.846309 — conservative, since the capture reforms should raise it.

4.1.5 The prediction

```
h_new = 5.846309 x 1.609438 - 0.075 x 34 - 1.213079
      = 9.409272 - 2.550000 - 1.213079 = +5.646193 assets/quarter

step  = 7.260894 - 2.550000 - 1.213079 = +3.497815 assets/quarter

Q_inf = (7.260894 - 1.213079) / 0.075 = 80.64 assets
tau   = 1/0.075 = 13.333 quarters;  half of the gap closes in 8.891 quarters
```

The lab starts from 34 (after shedding) and rebuilds:

Quarter	Q predicted	Do-nothing (from 46)
0	34.00	46.00
1	37.50	—
4	**46.49**	34.56
8	**55.64**	27.45
12	62.34	23.03
20	70.83	—
40	78.58	—

The lab returns to its starting stock of 46 assets in four quarters and passes it thereafter. The honest part of this prediction is its slowness: `tau` is more than three years, and the system is still 12 assets short of its fixed point after a decade. **A low delta is what makes stored order durable and it is exactly what makes rebuilding slow.**

4.1.6 What would falsify it

Preregistered before the first asset is retired, under the corpus protocol [Curzi, 2026i, Nosek et al., 2018]. All four baselines required.

1. **Trajectory.** At eta 20% below estimate, quarter 4 still predicts `Q = 41.31` and quarter 8 predicts `Q = 46.66`. **If Q < 41 at quarter 4, or < 46 at quarter 8, the eta estimate is refuted.** At eta 20% above, the same checkpoints are 51.68 and 64.63.
2. **The u channel.** Re-run the independent-estimate instrument at quarter 3. It must show `rho <= 0.70`. If Q rises while rho does not move, Q is rising for a reason outside the model and the causal claim fails even though the outcome is favourable.
3. **The delta channel.** Re-measure delta at quarter 6 by sampling assets with stewards absent for one quarter. It must be at or below 0.09. If not, A4.2 did not work and the fixed point is wrong.
4. **Shuffled-input control.** Refit `Y_{t+1} = a Y_t + b F(U_t) − C + e` with the interaction series permuted across quarters. **If the shuffled fit matches the ordered fit on held-out quarters, the apparent dependence on S is spurious and the entire diagnosis is void.**
5. **Persistence and matched-AR baselines.** The fitted model must beat both on held-out data, with `0 <= a <= 1` and `b > 0`. If a plain autoregression on Y predicts as well, report that the framework added nothing here.

Failure of 1, 2, 4, or 5 refutes the diagnosis.

17.4.2 4.2 Case B · A public library reading room

The system. The main reading room of an urban branch library: 16.0 x 9.5 x 4.2 m, terrazzo floor, plaster ceiling, painted plaster walls, 38 m^2 of glazing, 44 m^2 of filled shelving, typical occupancy 18. A silence rule is posted and enforced. The complaint, as stated by the branch manager: *“We spend all our time policing noise, and our regulars are disappearing.”*

4.2.1 The survey

Choice	Setting
State encoder	Category of use, assigned by direct observation (Whyte's method)
Interaction observable	Interaction-minutes, classified within- or between-category
Order observable Y	**Regularity stock**: distinct patrons with >= 4 visits in the trailing 8 weeks
Time scale	One month (4 weeks)
Cost observable	Patron-facing staff hours diverted to mediation and complaint handling

Order-proxy check [Curzi, 2026i]. The obvious proxy — door count — is the *input*, not the order, and using it would be the single most common measurement error in the framework. Regularity stock is a habit that persists when the room is closed. Lag test below confirms it. `Q_0 = 84` regulars.

delta. An 11-week closure for asbestos remediation, `S = 0` throughout. Regularity stock fell 96 -> 61. Eleven weeks is 2.75 monthly periods:

```
delta = 1 - (61/96)^(1/2.75) = 0.152021 per month
half-life = 4.203 months
```

Y decayed with a measurable time constant across the closure rather than collapsing. Y is stored order, and the closure paid for the delta estimate.

N. Six categories of use, defined by capability and need: quiet solo reading; student group study; children and caregivers; job-seeking and computer use; daytime shelter and rest; community meeting. `ln 6 = 1.791759`. Three of the six are effectively excluded by the silence rule — noted as FM-06 in the making, and revisited in B.3.

u, both sources (FM-08 requires two instruments).

```
u_cross = 0.0375     from 9 of 240 observed interaction-minutes crossing category boundaries
u_div   = 0.0963     from a within-category purpose-dissimilarity survey of quiet readers
u       = 0.133800
N_eff   = 1 + 5(0.1338) = 1.669000
S       = ln(1.669000) = 0.512225 nats    ->  28.59% of ln 6
```

The room contains six kinds of person and generates the surplus of 1.67.

eta. Over a five-month bilingual story-hour pilot, 71 new regulars were attributable to programming in the room, against cumulative `S = 5 x 0.512225 = 2.5611` nats:

```
eta = 71 / 2.5611 = 27.722212 regulars per nat
```

C. The same pilot consumed 1,580 patron-facing staff hours, which gives the conversion rate:

```
conversion = 71 / 1,580 = 0.044937 regulars per patron-facing staff hour
C = 54 h/month (mediation + complaint admin) x 0.044937 = 2.426582 regulars/month
```

4.2.2 The two master numbers, and validation

```
eta ln N = 27.722212 x 1.791759 = 49.671535
delta Q  = 0.152021 x 84        = 12.769789
C        =                        2.426582
```

```
h = 49.671535 − 12.769789 − 2.426582 = +34.475164 regulars/month
S / ln N = 28.59%
```

Validation. Predicted step: `14.200000 − 12.769789 − 2.426582 = −0.996371` regulars/month. Observed: the register fell 96 -> 84 over the preceding twelve months, **-1.000** per month. Residual **+0.004**. Admissible.

4.2.3 The binding constraint

Headroom is **+34.5 regulars per month** — enormous — and the room is losing one regular a month. The room's own fixed point at current parameters is `Q_inf = 77.45`, so it is not merely declining but descending toward an equilibrium below where it stands.

The binding constraint is **u**, and it is binding at 28.59% of ceiling. The reason is physical. The silence rule exists because the room cannot acoustically hold two activities at once, and the silence rule is what drives `u_cross` to 0.0375. **The manager named noise as the problem and noise is the symptom; the room's reverberation time is the cause, and u is the parameter it damages.** This is FM-26 (Reverberant Room) generating FM-01 (Parallel Wall) through a posted rule.

Sabine, computed. `RT60 = 0.161 V / A` (metric), `V = 16.0 x 9.5 x 4.2 = 638.40 m^3`. Absorption coefficients at 500 Hz; occupants at 0.45 m^2 each.

```
A = 152.0 x 0.02  (terrazzo floor)      =   3.04
  + 152.0 x 0.04  (plaster ceiling)     =   6.08
  + 176.2 x 0.05  (painted plaster)     =   8.81
  +  38.0 x 0.04  (glazing)             =   1.52
  +  44.0 x 0.30  (filled shelving)     =  13.20
  +  18   x 0.45  (occupants)           =   8.10
                                        --------
                                    A   =  40.750 m^2 sabins

RT60 = 0.161 x 638.40 / 40.750 = 2.5223 s occupied
                                 3.1480 s empty
```

Two and a half seconds. In a room of this volume, speech from one syllable masks the next; the quiet voice falls below intelligibility while the loud voice does not; the Lombard effect then raises the floor for everyone. **The loudest patron wins this room by physics.** The staff are not policing behaviour; they are absorbing, in labour, an acoustic deficiency in the building.

4.2.4 The intervention

B4.1 — Absorption. 90 m^2 of Class A ceiling absorber (alpha = 0.90) over the central span; 40 m^2 of wall panel (alpha = 0.85) on the wall opposite the primary reading bay; 100 m^2 of heavy carpet with underlay (alpha = 0.25) over the terrazzo.

```
A = 52.0 x 0.02 + 100.0 x 0.25 + 62.0 x 0.04 + 90.0 x 0.90
  + 136.2 x 0.05 + 40.0 x 0.85 + 38.0 x 0.04 + 44.0 x 0.30 + 18 x 0.45
  = 1.04 + 25.00 + 2.48 + 81.00 + 6.81 + 34.00 + 1.52 + 13.20 + 8.10
  = 173.150 m^2 sabins

RT60 = 0.161 x 638.40 / 173.150 = 0.5936 s          (4.25x reduction)
```

Partial packages, for budgeting: ceiling only gives **0.8699 s**; ceiling plus walls gives **0.6845 s**. Ceiling first is correct — largest uninterrupted surface, first reflection path, and it alone brings the room inside the usable band. The absolute cost is not stated here because it is local and this paper does not invent figures; what is specified is the absorption schedule, which is the deliverable a contractor prices.

B4.2 — Zoning, not silence. Replace the blanket rule with three acoustically distinct bays — a treated quiet bay, a group bay, a story alcove — arranged so that categories are in *sightline* of each other but not in each other's acoustic field. This is Alexander's nested-scales pattern and Gehl's edge effect used deliberately: people occupy perimeters and thresholds first, and a threshold that lets you see the next zone before you enter it is the device (FM-24).

B4.3 — A visible medium. A board on which patrons post and claim recurring uses of the room. This is stigmergy: **the state is put in the medium, not in a coordinator.** The identity that makes the move computable is Curzi [2026h]'s — that headroom at equilibrium is unused variety, `h_inf = eta(ln N − F(u*))`, verified to six decimals at `eta = 0.35`, `delta = 0.015`, `N = 6`, `u* = 0.55`, giving `Q_inf = 2.840970` and `h_inf = 0.164501` — which is why a medium that raises `u` converts headroom that was already there rather than creating new headroom.

B4.4 — The observable that tells you it is working before Q moves. Gehl's distinction between necessary, optional, and social activity supplies a leading indicator. Necessary activity — returning a book, using a terminal because there is nowhere else — happens regardless of quality. Optional and social activity happen only when conditions permit them. **Optional activity is the observable proxy for eta**, and it responds within weeks, where regularity stock takes months. Instrument it Whyte's way, by direct observation on a fixed schedule: count people who stay without a task, count conversations lasting over two minutes, and map seat occupancy. Whyte's findings about small urban spaces transfer directly to the interior — sittable space, sun, and triangulation are the variables — and the transfer is badged as a modelling assumption, since his observations were of plazas, not reading rooms.

Three specific readings to take before and after, none of which requires a budget:

Instrument	Pre-treatment reading	Post-treatment target
Occupancy by seat, 12 sessions	Persistently empty seats mapped	Empty-seat count halved
Conversations > 2 minutes / hour	Baseline count	Any increase; the silence rule made this near zero
Stay-without-task headcount	Baseline count	Increase, as optional activity returns

Two conditions are also worth stating because they are cheap and the literature on them is real: daylight and view. Ulrich (*Science*, 1984) found post-surgical recovery faster with a window view of trees. That is a hospital finding and its transfer to a reading room is a modelling assumption — but the 38 m^2 of existing glazing is an asset already paid for, and the seating plan should put the quiet bay against it rather than against the interior wall, at no cost.

Post-intervention parameters, all conservative:

```
u_cross: 0.0375 -> 0.19      (zoning admits cross-category contact at the edges)
u_div:   0.0963 -> 0.13      (more purposes admitted within each category)
u = 0.320000;  N_eff = 2.600000;  S = 0.955511 nats  ->  53.33% of ln 6
eta held at 27.722212        (no credit taken for the board)
delta held at 0.152021       (no credit taken for improved habit formation)
C = 22 h x 0.044937 + 1.20 (panel upkeep + a named room steward) = 2.188608
```

4.2.5 The prediction

```
h_new = 49.671535 − 12.769789 − 2.188608 = +34.713139 regulars/month
step  = 26.488890 − 12.769789 − 2.188608 = +11.530494 regulars/month
Q_inf = (26.488890 − 2.188608) / 0.152021 = 159.85 regulars
tau   = 6.578 months;  95% of the gap closes in about 19.7 months
```

Note carefully: **h barely moved — from +34.48 to +34.71 — while the one-step change moved from -1.00 to +11.53.** Headroom is almost untouched because it does not contain u at all. The entire intervention is a change in how much of an already-available ceiling the room collects. This is the clearest demonstration in the corpus of the corpus failure catalog \citep{ccp:CCP-016}.

Month	Q predicted	Do-nothing
0	84.0	84.0
3	113.6	81.6
6	**131.6**	79.9
12	**149.4**	78.4
24	158.4	77.6
36	159.6	77.5

4.2.6 What would falsify it

1. **RT60.** Measured occupied reverberation time after installation must be **below 0.80 s**. If it is not, the absorption schedule was wrong and every downstream prediction is void independently of what happens to attendance.
2. **The u channel, both sources.** Re-run the interaction-minute classification at month 3. `u_cross` must exceed 0.12. If Q rises while `u_cross` does not, the mechanism claimed here is not the mechanism operating.
3. **Trajectory, banded by realized u.** The prediction is explicitly conditional:

Realized u	S	Q at month 6	Q at month 12	Q_inf
0.1338 (no change)	0.5122	80.9	79.7	79.0
0.20	0.6931	101.6	108.1	112.0
0.26	0.8329	117.6	130.1	137.5
0.32 (target)	**0.9555**	**131.6**	**149.4**	**159.8**
0.40	1.0986	148.0	171.9	185.9

   If Q is below 101 at month 6, u did not move at all and the zoning failed.
4. **Mediation hours.** Must fall below 30 h/month by month 3. If reverberation falls and mediation does not, the conflict was not acoustic and C was mis-attributed.
5. **Shuffled-input control**, and the persistence and matched-AR baselines. As in Case A.

17.4.3 4.3 Case C · A distributed open-source project

The system. A library of 141 modules, three core maintainers, 37 contributors active in the trailing year, no employer, no roadmap authority. The complaint, from the maintainers: *“We are drowning and the docs are a lie.”*

4.3.1 The survey

Choice	Setting
State encoder	Contributor capability category — what they can independently fix
Interaction observable	Independent reviews and design discussions on merged changes
Order observable Y	Modules that are **CI-green on the current toolchain and whose documented examples execute**
Time scale	One month
Cost observable	Green modules destroyed by in-period breaking changes

Order-proxy check [Curzi, 2026i], and a rejected first attempt. The first proposed proxy was "modules with a maintainer who responded within 14 days." That is maintainer *interaction* wearing a state costume, and it fails the test. The adopted proxy is a pure state of the artifact: it can be evaluated by a stranger with a clean machine and no access to the project's social graph. `Q_0 = 63` of 141 modules.

delta. An eight-week holiday freeze with near-zero merges. Green modules fell 71 -> 63, purely from toolchain drift, dependency deprecation, and link rot:

```
delta = 1 - (63/71)^(1/2) = 0.058021 per month
half-life = 11.596 months
```

N — and FM-06 in the open. Seven categories are declared in the contributor guide: core runtime, packaging, docs, platform ports, downstream integrators, security, accessibility/i18n. Three of the seven — ports, security, accessibility — have no active member.

```
declared: ln 7 = 1.945910
staffed:  ln 4 = 1.386294
```

The project has been planning against a ceiling 40.4% higher than the one it owns.

u. Forty pull requests were each reviewed independently by two reviewers who did not see each other's comments. Mean Jaccard overlap of flagged defects `rho = 0.68`:

```
u     = 1 - 0.68^2 = 0.537600
N_eff = 1 + 3(0.5376) = 2.612800
S     = ln(2.612800) = 0.960422 nats    ->  69.28% of ln 4
```

eta. Over ten months, 41 modules were newly greened or restored to green attributably to logged review and design threads, against `S = 10 x 0.960422 = 9.6042` nats:

```
eta = 41 / 9.6042 = 4.268955 modules per nat
```

C. Distinguished from delta by cause: delta is decay *without* input (measured in the freeze); C is damage caused *by the activity itself*. Over twelve months, 35 green modules were destroyed by in-period breaking changes and dependency bumps that were not repaired within the period:

```
C = 35 / 12 = 2.916667 modules per month
```

4.3.2 The two master numbers, and validation

```
eta ln N = 4.268955 x 1.386294 = 5.918028
delta Q  = 0.058021 x 63       = 3.655339
C        =                       2.916667

h = 5.918028 - 3.655339 - 2.916667 = -0.653978 modules/month     <-- NEGATIVE
S / ln N = 69.28%
```

Validation. Predicted step: `4.100000 - 3.655339 - 2.916667 = -2.472006` modules/month. Observed: green modules fell 92 -> 63 over the preceding twelve months, **-2.417** per month. Residual **-0.055**, under 1% of the ceiling term. Admissible.

4.3.3 The binding constraint

`h < 0`. The negative-headroom corollary states it without qualification [Curzi, 2026b]: if `C_t > eta ln N - delta Q_t`, **no choice of interaction geometry prevents order from decreasing.** More reviewers will not save this project. Better reviewers will not save it. A larger community will make it worse, because adding contributors adds channels quadratically (Brooks) against a ceiling that grows logarithmically — FM-07, the Ceiling Chase, which is what most projects in this state attempt.

The measured fixed point makes the position explicit (FM-20):

```
Q_inf = (4.100000 - 2.916667) / 0.058021 = 20.39 modules
```

The project stands at 63 and its own equilibrium is 20. At the current rate it reaches zero green modules in **25.5 months**. Note the realized fraction: **69.28%** of ceiling — the *geometry is comparatively healthy*. The reviewers are genuinely decorrelated. That is precisely why the geometry is not the fix.

Binding: **C**, with a Frozen Denominator (FM-06) suppressing the ceiling and an Unshed Q (FM-18) inflating the maintenance charge. Also present and worth naming: three maintainers is a hub, and scale-free networks are robust to random failure and **fragile to targeted hub attack** (Albert, Jeong & Barabási 2000). The project's bus factor is its single largest unpriced risk, and it is FM-15, the Founder Battery, distributed across three people.

4.3.4 The intervention

Under the gated intervention procedure [Curzi, 2026j], `h < 0` forbids building. The permitted moves are cut C, raise eta, shed Q. N is added here only after the first three, and only because two of the empty categories can be staffed from existing downstream users at near-zero coordination cost.

C4.1 — Cut C (the binding term). A compatibility CI matrix pinned to the current and two prior toolchains; a mandatory deprecation window; no merge without a downstream smoke test. This is stigmergy applied to code: the constraint lives in the medium, enforced by the machine, not in a maintainer's attention. Target `C: 2.9167 -> 1.10 modules/month`.

C4.2 — Shed Q. Move 22 modules that no downstream consumer has imported in twelve months to an explicitly unmaintained archive tier. Not deleted; **honestly labelled**, which is the entire point. `Q: 63 -> 41`, and `delta Q: 3.655 -> 2.379`.

C4.3 — Raise eta. A merge checklist that will not close without the durable artifact: an executable example, a docs entry, and a named module owner. Measured eta is 4.269; target **5.0** — deliberately below the 5.6 the pilot suggested, because a prediction should be made at the conservative end of its own evidence.

C4.4 — Raise N, with sockets. Staff two of the three empty categories — platform ports and security — from existing downstream users who already do this work privately. `N: 4 -> 6`, `ln`

`6 = 1.791759`. Each new category gets a named review gate it owns (FM-13). Accessibility is left empty and **recorded as empty**, because nobody has volunteered and inventing a category would be FM-06 all over again.

Adding a security reviewer and a platform porter to the review pool should also decorrelate review. Target `rho: 0.68 -> 0.52`:

```
u     = 1 - 0.52^2 = 0.729600
N_eff = 1 + 5(0.7296) = 4.648000
S     = ln(4.648000) = 1.536437 nats    ->  85.75% of ln 6
```

Ostrom's design principles, read as engineering, supply the governance layer: clear boundaries (the archive tier), congruence with local conditions (the toolchain matrix), collective-choice arrangements (contributors set the deprecation window), monitoring (CI), graduated sanctions (revert, then require a smoke test, then require co-review), conflict-resolution mechanisms, and recognized rights to organize. Badged: Ostrom's principles are an empirical generalization from common-pool resource institutions; their transfer to a code repository is a modelling assumption.

C4.5 — How to land the discipline, which is a contagion problem, not an announcement. A merge checklist and a deprecation window are behaviours, and behaviours spread by complex contagion: adoption needs multiple reinforcing contacts, so **clustered topology beats broad thin exposure** (Centola & Macy 2007; Centola 2010). Announcing the policy on the mailing list is FM-14, the Sprayed Seed, and it is what most projects do. Instead, saturate one dense neighbourhood of the contributor graph — the six people who review each other's work weekly — until the practice is unremarkable there, then let the boundary carry it. Sizing: Centola et al. (*Science*, 2018) found a committed minority near **25%** flipped an established convention; at 37 active contributors that is **10 people**. Badged: the 25% figure is from controlled online experiments, and field transfer to a volunteer code project is a modelling assumption, not a measurement.

Conway's Law is also load-bearing here and it cuts both ways. Systems mirror the communication structure of the organizations that build them (Conway 1967). The three empty categories are not an accident of recruitment: the library has no port layer and no security surface *because* nobody in the communication structure owns those concerns. Staffing the categories will, over time, change the artifact's seams. That is the mechanism by which C4.4 eventually reduces C as well as raising N — and it is slow, so no credit for it is taken in the prediction below.

4.3.5 The prediction

```
h_new = 5.0 x 1.791759 - 0.058021 x 41 - 1.10
      = 8.958797 - 2.378872 - 1.100000 = +5.479926 modules/month

step  = 7.682185 - 2.378872 - 1.100000 = +4.203313 modules/month

Q_inf = (7.682185 - 1.100000) / 0.058021 = 113.44 modules
tau   = 17.235 months
```

Validity range, stated. After archiving, the project has 119 maintainable modules. `Q_inf = 113.44` sits just under that cap, and the linear law has no saturation term, so **the model is trustworthy only while Q is well below 119.** It crosses Q = 100 at month 29; beyond roughly Q = 95 the prediction should be treated as directional, not numerical.

Month	Q predicted	Do-nothing (from 63)
0	41.0	63.0

Month	Q predicted	Do-nothing (from 63)
6	**62.8**	50.2
12	**78.1**	41.2
24	96.2	30.5
36	105.0	—

The project passes its pre-intervention stock of 63 green modules in month 6 — from a *smaller* base, because shedding was one of the levers.

4.3.6 What would falsify it

1. **The C channel first.** Count green modules destroyed by in-period breaking changes over months 1-6. **If the rate is above 2.0/month, C4.1 failed** and the following bands apply instead:

Scenario	h	step	Q at 12	Q_inf
Full package (eta 5.0, C 1.10)	+5.480	+4.203	78.1	113.4
eta unchanged (4.269), C 1.10	+4.170	+3.080	68.2	94.1
eta 5.0, C slips to 2.00	+4.580	+3.303	70.1	97.9
eta 5.0, C discipline fails entirely	+3.663	+2.387	62.1	82.1

   Every row is positive: **the intervention's sign is robust even to complete failure of its largest single component**, because shedding Q and staffing N together moved h from -0.65 to +3.66 before eta or C were touched. That robustness is itself a falsifiable claim — if Q is still falling at month 6 under any of these rows, the model is wrong, not the execution.

2. **The N channel.** Both new categories must land at least three independent reviews in the first quarter. A category with a name and no output is FM-06 restored, and the ceiling reverts to `ln 4`.

3. **The u channel.** Re-run the paired-review overlap measurement at month 6. `rho <= 0.60` required.

4. **Shuffled-input control** and the persistence and matched-AR baselines, as before. This case is the most vulnerable of the four to a spurious fit, because CI-green counts move for toolchain reasons entirely unrelated to human interaction. **If the shuffled control fits as well, the whole causal story is an artifact of the release calendar.**

17.4.4 4.4 Case D · A residential commons of twenty-two households

The system. Twenty-two households, 51 residents, purpose-built co-housing: common house with kitchen and laundry, guest room, workshop, tool library, shared garden beds, bike shed. Governed by whole-group consensus. The complaint, from a resident: *"Everything takes a year and nothing stays fixed."*

4.4.1 The survey

Choice	Setting
State encoder	Household activity category
Interaction observable	Work parties and commons decisions with cross-household participation
Order observable Y	**Shared capabilities usable without negotiation**, in commons-units (cu)
Time scale	One month
Cost observable	Household-hours consumed by governance, mediation, and free-rider repair

A capability counts toward Y if a resident can use it *without asking anyone*: it works, it is findable, and its rules are known. Walk-through inspection: 14 tool-library items, 7 common-house systems, 12 written norms with observed compliance at or above 80%, 5 residents able to run a work party unaided. `Q_0 = 38 cu`.

Order-proxy check [Curzi, 2026i]. Meetings held would be the input. Capabilities that work when nobody is meeting is the order. The lag test is available and decisive, below.

delta — and FM-15 in the field. After the founding steward moved out, work parties lapsed for nine months (`S ~ 0`). Commons stock fell 47 -> 38:

```
delta = 1 - (38/47)^(1/9) = 0.023341 per month
half-life = 29.348 months  (2.45 years)
```

This is the lowest delta in the set, for the obvious reason that most of this commons is **physical**. Bricks forget slowly.

N. Five household activity categories by need and contribution: families with young children; retired residents; full-time-employed households; remote workers; households with mobility limits. `ln 5 = 1.609438`.

u. Each household independently ranked twelve candidate commons projects before any discussion. Mean pairwise Spearman rank correlation `rho = 0.41`:

```
u     = 1 - 0.41^2 = 0.831900
N_eff = 1 + 4(0.8319) = 4.327600
S     = ln(4.327600) = 1.465013 nats    ->  91.03% of ln 5
```

This is the highest realised fraction in the set. The block is genuinely, measurably heterogeneous, and it converts almost none of that into anything.

eta. Over the fourteen months when monthly work parties ran, 19 cu of new or restored capability were attributable, against `S = 14 x 1.465013 = 20.5102` nats:

```
eta = 19 / 20.5102 = 0.926369 cu per nat
```

C. Conversion rate from the same window: 19 cu from 1,344 productive household-hours of work party = `0.01413690` cu/h. Governance tally: one whole-group consensus meeting per month at 22 households x 2.0 h = 44 h; mediation 11 h; free-rider repair 7 h. Total 62 household-hours/month.

```
C = 62 x 0.01413690 = 0.876488 cu per month
```

4.4.2 The two master numbers, and validation

```
eta ln N = 0.926369 x 1.609438 = 1.490934
delta Q  = 0.023341 x 38       = 0.886966
C        =                       0.876488

h = 1.490934 − 0.886966 − 0.876488 = −0.272521 cu/month        <-- NEGATIVE
S / ln N = 91.03%
```

Validation. Predicted step: `1.357143 − 0.886966 − 0.876488 = −0.406311` cu/month. Observed: the walk-through five months ago read 40 cu; today it reads 38, or **-0.400** per month. Residual **-0.006**. Admissible.

4.4.3 The binding constraint

`h < 0` again — and for the exact opposite reason to Case C. Here the geometry is superb: 91% of ceiling, an honest 0.83 non-parallelism, five real categories. What the block has is **eta = 0.926 cu per nat** — the lowest conversion efficiency in the set by a factor of four.

The residents are having genuinely different conversations and almost nothing durable comes out of them. That is FM-11, the Leaky Vessel, in its pure form: full surplus, no capture. The corpus states the same structure in magnitude form elsewhere as `E = G.B + eps` with the forbidding clause `G = 0 => E = 0` for all `B`, badged as a postulate and paired with the storage law's `eta = 0` by explicit correspondence [Curzi, 2026ab]: **where the channel gain is zero, no amount of input produces an effect, and the co-housing block is the nearest thing in this set to that limit.** The work parties happen, the ideas are real, and the tool comes back unlabelled to a shelf nobody owns.

Second binding term: C. Twenty-two households in whole-group consensus is `n(n−1)/2 = 231` communication channels — FM-21, the Quadratic Committee. At n = 22 the ratio is **74.73 channels per nat of ceiling**; at n = 5 it is 6.21. The block pays twelve times the coordination cost per unit of ceiling that a five-household cluster would.

Third: FM-15, and it is already historical — the block's delta was measured *because* one person left and everything stopped.

4.4.4 The intervention

D4.1 — Raise eta (the binding lever, and the cheapest). Three capture devices, in ascending durability:

- **A steward of record for every capability.** Named, on the object, rotating annually, with the right to decline the rotation.
- **A physical commons board** in the common house: what exists, where it is, who stewards it, what is broken. Stigmergy — **the state lives in the medium, not in a coordinator** — which is also the structural correction to the Founder Battery.
- **A work-party close-out rule:** no party ends without one written durable artifact. A label, a repaired item logged, a norm written down. The conversion finding restated [Curzi, 2026b]: *"we had a great work party"* is an `eta = 0` report.

Pilot-measured eta reached 1.55; the prediction is made at **1.40**, conservatively.

D4.2 — Cut C by nesting, not by shrinking. Four clusters of about 5.5 households, with **overlapping membership** — each household sits in two clusters. This is Alexander's semilattice ("A City is Not a Tree," 1965) rather than a tree, and the overlap is what preserves cross-contact.

```
whole group:   22 x 21 / 2                      = 231 channels
nested:        4 clusters x (5.5 x 4.5 / 2) = 49.5
```

```
            + delegate council 4 x 3 / 2   =  6.0
                                           -------
                                             55.5 channels     = 24.03% of 231
```

Whole-group consensus moves to quarterly and covers constitutional matters only; clusters decide monthly within scope. **Badged as a modelling assumption:** the projection below scales meeting hours by the channel ratio, on the reasoning that the binding cost is pairwise reconciliation rather than seat time. That assumption is itself in the falsification list.

```
meeting term:   44 h x 0.2403                    = 10.571 h
whole-group quarterly, amortized: 22 x 2.5 / 3   = 18.333 h
mediation (down from 11)                         =  8 h
free-rider repair (graduated sanctions, Ostrom 5) =  4 h
                                                   -------
                                                    40.905 h/month
C = 40.905 x 0.01413690 = 0.578267 cu/month
```

D4.3 — Ostrom's eight design principles [Ostrom, 1990, 2010, Cox et al., 2010] as the eta and C specification. This is the case the principles were derived for — a common-pool resource governed by its users — and each principle maps to a specific parameter in the law. Badged: Ostrom's principles are an empirical generalization from field studies of common-pool resource institutions (*Governing the Commons*, 1990); their statement is hers, the parameter mapping is this paper's.

Ostrom principle	Instrument in this block	Parameter
1. Clearly defined boundaries	Which capabilities are commons, which are household	eta, C
2. Congruence of rules with local conditions	Cluster-level rules for cluster-level assets	C
3. Collective-choice arrangements	Those affected set the rule; opt-out preserved	u, and the consent constraint
4. Monitoring	The commons board shows broken items, by whom, since when	eta, delta
5. Graduated sanctions	Reminder -> repair request -> stewardship reallocated	C
6. Conflict-resolution mechanisms	A standing, low-cost, non-escalating process	C
7. Recognition of rights to organize	Clusters may form and dissolve without approval	u, N
8. Nested enterprises	The four-cluster semilattice itself	C

Principle 8 is the one this case turns on, and it is the same structural claim as Alexander's semilattice from a different literature: **nested, overlapping enterprise beats both the flat assembly and the strict hierarchy**, and the arithmetic in D4.2 prices exactly how much. Principle 4 is the one that does the eta work — a commons board is a monitoring device, and monitoring is what converts a conversation into a state of record.

Consent check, and one cut. Graduated sanctions must preserve the right to decline. The version adopted escalates from a reminder, to a request for repair, to reallocation of the capability's stewardship — none of which removes anyone's ability to withdraw from the commons activity entirely. A proposed version that made cluster attendance a condition of tool-

library access **was cut**, because it made non-participation costly and is therefore outside this manual's scope.

Honest cost of the nesting: u falls. Clustering, even with overlap, reduces cross-category contact. Predict `u: 0.8319 -> 0.79`:

```
N_eff = 1 + 4(0.79) = 4.160000
S     = ln(4.160000) = 1.425515 nats    ->  88.57% of ln 5   (down from 91.03%)
```

delta is held at 0.023341 — no credit taken for the board, though it should help.

4.4.5 The prediction

```
h_new = 1.40 x 1.609438 - 0.023341 x 38 - 0.578267
      = 2.253213 - 0.886966 - 0.578267 = +0.787980 cu/month

step  = 1.995721 - 0.886966 - 0.578267 = +0.530488 cu/month

Q_inf = (1.995721 - 0.578267) / 0.023341 = 60.73 cu
tau   = 42.843 months  (3.57 years);  half the gap closes in 29.3 months
```

Month	Q predicted	Do-nothing
0	38.0	38.0
12	**43.6**	33.7
24	**47.8**	30.5
36	51.0	28.0
60	55.2	—
120	59.4	—

This is the slowest system in the set and the arithmetic says so plainly. `3 tau` is 128.5 months — **ten years and eight months** to settle within 5% of the fixed point. The same low delta that let the commons survive nine months of neglect is what makes it take a decade to build. A prospectus promising visible transformation in a year would be lying, and the lie would be detectable from `tau = 1/delta` alone.

Single-lever check, to show the bundle is necessary:

Lever	step	h
C cut alone (eta 0.926)	-0.145	+0.026
eta alone (1.40, C unchanged)	+0.288	+0.490
Both	**+0.530**	**+0.788**

4.4.6 What would falsify it

1. **The eta channel, which is the whole claim.** Count attributable cu per nat over months 1-12. Bands:

eta realized	step	Q at 12	Q at 24	Q_inf
0.926 (no change)	-0.145	36.5	35.3	31.8
1.10	+0.103	39.1	39.9	42.4

eta realized	step	Q at 12	Q at 24	Q_inf
1.40 (target)	**+0.530**	**43.6**	**47.8**	**60.7**
1.55	+0.744	45.9	51.8	69.9

If Q is at or below 39 cu at month 12, eta did not move and the capture devices failed.

2. **The channel-scaling assumption.** Measure actual governance household-hours at month 6. The prediction requires **at or below 45 h/month**. If nesting cuts channels but not hours, the scaling assumption is refuted and C must be re-estimated from the measurement.
3. **The u cost.** Re-run the independent-ranking instrument at month 9. If `rho` rises above 0.55 — that is, `u` falls below 0.70 — the clusters have become a tree rather than a semi-lattice, and the overlap must be increased. This is the intervention's own most likely failure and it is predicted in advance rather than explained afterward.
4. **Persistence, direct-input, matched-AR, and shuffled-input controls.** All four. With a `tau` of 43 months, a persistence baseline is a genuinely strong competitor over any horizon shorter than two years, and **if the model cannot beat persistence at month 24 it has added nothing.**

17.4.5 4.5 Case E · One desk, and the scale-freedom test

The four cases span three orders of magnitude in headcount. This last one has a headcount of one, and it is here to test the scale-freedom claim directly: the claim that the storage law is **scale-free in form**, so that only `(delta, eta, N, C)` change between a desk, a room, a team, a building, an institution, and a city. That is a claim about **model form**, not about physics, and it is badged as such wherever it appears.

The system. One person, one desk, three concurrent workstreams. Period: one week. Order observable: **resumable threads** — work whose state can be recovered from the desk alone, without reconstruction from memory. `Q_0 = 6`.

```
delta:  after a two-week absence, resumable threads fell 7 -> 5
        delta = 1 - (5/7)^(1/2) = 0.154846 per week;  half-life 4.120 weeks

N     = 3 workstreams;   ln 3 = 1.098612
u:      independent recall test: how often does resuming one thread require
        reconstructing another?  coupling rho = 0.74
        u = 1 - 0.74^2 = 0.452400
        N_eff = 1 + 2(0.4524) = 1.904800
        S = ln(1.904800) = 0.644377 nats   ->  58.65% of ln 3

eta:    13 threads made resumable over 10 weeks, attributable to end-of-session capture
        eta = 13 / (10 x 0.644377) = 2.017452 threads per nat
yield:  13 / (10 weeks x 5.0 productive hours) = 0.260000 threads per hour
C:      2.39 h/week of re-orientation and searching
        C = 2.39 x 0.26 = 0.621400 threads per week

h = 2.017452 x 1.098612 - 0.154846 x 6 - 0.621400
  = 2.216398 - 0.929074 - 0.621400 = +0.665924 threads/week

step = 1.300000 - 0.929074 - 0.621400 = -0.250474 threads/week
```

Observed: threads fell 7 -> 6 over four weeks, **-0.250** per week. Residual **-0.0005**. The survey validates at a headcount of one, with the same equation, the same seven steps, and the same validation rule as a twenty-two household commons.

The intervention is one index card per thread, written at the end of each session and parked in view: what was decided, what is next, what would unblock it. It is the smallest possible instance of the same three devices used in Case D — a state of record, a visible medium, and a close-out rule.

```
delta: 0.154846 -> 0.09      (the card survives the week; memory does not)
rho:   0.74 -> 0.55  ->  u = 0.697500,  N_eff = 2.395000
S      = ln(2.395000) = 0.873383 nats   ->  79.50% of ln 3
C:     1.1 h x 0.26 = 0.286000

h_new = 2.216398 - 0.540000 - 0.286000 = +1.390398 threads/week
step  = 1.762009 - 0.540000 - 0.286000 = +0.936009 threads/week
Q_inf = (1.762009 - 0.286000) / 0.09 = 16.40 threads
tau   = 11.111 weeks
```

Week	Q predicted	Do-nothing
0	6.00	6.00
2	7.79	—
4	**9.27**	5.21
8	11.51	4.80
12	13.05	4.60

Validity range, again. `Q_inf = 16.40` exceeds what one person can hold; the practical cap is nearer ten to twelve. The model is trustworthy to about week 6 and directional after that. Naming the cap is not a weakness of the method; **failing to name it is the failure**, and it is Failure Mode D-6.

17.4.6 4.6 The cross-case comparison

Five systems, one equation, parameters spanning three orders of magnitude:

System	Period	Units	delta	eta	N	u	S/lnN	C	Q	h	step
Desk	week	threads	0.1548	2.017	3	0.452	58.7%	0.621	6	+0.666	-0.250
Research group	quarter	assets	0.1120	5.846	4	0.260	41.6%	1.608	46	+1.343	-3.387
Library room	month	regulars	0.1520	27.722	6	0.134	28.6%	2.427	84	+34.475	-0.996
Open source	month	modules	0.0580	4.269	4	0.538	69.3%	2.917	63	-0.654	-2.472
Co-housing	month	cu	0.0233	0.926	5	0.832	91.0%	0.876	38	-0.273	-0.406

After intervention:

System	delta	eta	N	u	S/lnN	C	h	step	Q_inf	tau
Desk	0.0900	2.017	3	0.698	79.5%	0.286	+1.390	+0.936	16.4	11.1 wk
Research group	0.0750	5.846	5	0.616	77.2%	1.213	+5.646	+3.498	80.6	13.3 q
Library room	0.1520	27.722	6	0.320	53.3%	2.189	+34.713	+11.530	159.8	6.6 mo
Open source	0.0580	5.000	6	0.730	85.8%	1.100	+5.480	+4.203	113.4	17.2 mo
Co-housing	0.0233	1.400	5	0.790	88.6%	0.578	+0.788	+0.530	60.7	42.8 mo

Five readings from the table, each of which is a finding rather than a restatement:

1. Every one of the five was losing order at survey, and no two for the same reason. Binding constraints were, respectively: delta and u jointly; delta-Q with a parallel roster; u alone; C with a frozen denominator; eta with a quadratic committee. A method that always finds the same fault is not diagnosing.

2. h and `S / ln N` are independent readings and both are required. The library has 53x the headroom of the co-housing block and collects less than a third of its ceiling. The block collects 91% of a ceiling it cannot afford. Either number alone points at the wrong lever.

3. eta spans a factor of thirty across the five systems and is the cheapest thing to change in four of them. This is the conversion result [Curzi, 2026b], arrived at from the arithmetic rather than asserted.

4. tau = 1/delta sets the honest timescale, and it is not negotiable by effort. The library settles in about twenty months; the co-housing block needs a decade. **The same low delta that makes a commons resilient is what makes it slow, and no intervention in this paper changes that trade.**

5. The form held at every scale. One desk and twenty-two households used the same seven steps, the same validation rule, and the same falsification discipline, and both surveys reproduced their observed change to within 0.006 units. That is what the scale-freedom claim actually asserts and all it asserts: **same equation, different parameters, and the parameters must be measured every time.**

The measurements above are the point rather than the model, which is four lines. **What the procedure supplies is the discipline of putting a number where an opinion was, validating it against something that already happened, and writing down in advance what would refute it.**

17.5 5 The survey specification

The general procedure, abstracted from the four cases. Numbered for checking.

S5.1 — Choose the five operational observables before measuring anything: state encoder, interaction observable, **order observable Y**, time scale, cost observable. Write them down. Changing one later invalidates every earlier reading.

S5.2 — Apply the lag test to Y before proceeding. Find or create a period of near-zero input and observe Y. If Y tracks the pause with no lag, Y is the input and must be replaced. **All four cases above rejected an obvious proxy at this step, and Case C rejected its first choice explicitly.**

S5.3 — Estimate delta from that same zero-input period: `delta = 1 - (Q_1/Q_0)^(1/t)`. The interruption that annoys you is the measurement you cannot otherwise afford.

S5.4 — Partition N by capability, not by title. Count only staffed categories. Record empty categories as empty.

S5.5 — Estimate u with two instruments, `u_cross` and `u_div`, and report both. `u = 1 - rho^2`. Report the realized fraction `S / ln N` in the same breath.

S5.6 — Estimate eta as `attributable delta_Q / cumulative S` over at least five periods.

S5.7 — Take C as a residual, not as a list: total capacity minus productive capacity, priced at the same yield rate used for eta.

S5.8 — **Validate.** Predicted `eta S - delta Q - C` against observed change in Q. Agreement within about 10% of the ceiling term. **A survey that does not reproduce the past may not be used to predict the future.**

S5.9 — Compute `h` and `S / ln N`. Together. Then:

```
h < 0    -> building is forbidden. Cut C, raise eta, or shed Q, in that order of cheapness.
h = 0    -> maintenance only. Hold.
h > 0    -> growth available. Any requested r_t must satisfy r_t <= h_t.
            If S / ln N is also low, the binding constraint is u, not h.
```

S5.10 — Rank levers by cost: eta cheapest, then C, then shedding Q, then u, then N. **N is last because it is logarithmic**: going 4 -> 5 buys 0.2231 nats of ceiling for a 25% increase in categories, and every added category adds channels quadratically.

S5.11 — Apply the consent test to every element: *what happens to a person who declines?* If the answer contains an adverse consequence, cut the element and record the cut.

S5.12 — Predict a dated trajectory, a fixed point, a time constant, and **a validity range**. State the cap. Where `Q_inf` exceeds a physical or organizational cap, say which one binds.

S5.13 — Write the falsification conditions before starting. Four baselines: persistence, direct-input, matched autoregressive, **shuffled-input**. Band the prediction by the uncertainty in the least reliable parameter — usually eta.

17.6 6 Failure modes observed across the five surveys

What went wrong across these five surveys, and what a practitioner should expect.

The named complaint was wrong every time. The PI said "busy and unproductive" and the fault was a maintenance charge plus a parallel roster. The branch manager said "noise" and the fault was 2.52 seconds of reverberation collapsing u. The maintainers said "drowning" and the fault was churn cost against a frozen denominator. The resident said "slow" and the fault was a conversion efficiency of 0.93. **The complaint reliably names the symptom with the highest salience, and salience is uncorrelated with bindingness.**

Positive headroom read as health (Cases A and B). Both had `h > 0` and both were shrinking. Any diagnosis stopping at `h` would have been exactly wrong in two of the five cases. `h` is the ceiling; `S / ln N` is the collection rate; you need both numbers or neither.

Single-lever thinking (all four). Case A's table shows every single lever leaving the step negative. Case D's shows the same. Case C is the exception that proves the rule — its bundle is robust to failure of its largest component, and that robustness was *computed*, not hoped for.

Taking credit for unmeasured improvement. Every case above held at least one parameter fixed that the intervention should plausibly improve — Case B held both eta and delta, Case A held eta, Case D held delta. **A prediction made at the optimistic end of its own evidence is not a prediction.**

Model-range violations (Cases C and E). `Q_inf = 113.44` against 119 maintainable modules; `Q_inf = 16.40` against a desk that holds perhaps a dozen threads. The linear law has no saturation term and will happily predict past the cap. Say where it stops being true.

The shuffled-input control is the one that would kill these. Case C is the clearest: green module counts move on the release calendar for reasons entirely unrelated to human interaction. Every other baseline can be passed by a model that is merely tracking a trend. **A framework that cannot fail cannot be trusted, and this is precisely the test these four cases would fail if the framework were empty.**

17.6.1 6.1 The catalog entries encountered, by case

Cross-referenced to the thirty-two-entry corpus failure catalog [Curzi, 2026j], so the five surveys can be read as a worked index. **Two-digit labels of the form `FM-nn` below refer to that catalog; this paper's own failure modes carry the prefixes `FM-E` and `FM-B` and are given in Sections 7.13 and 8.13.** No case presented with fewer than three co-occurring faults, which is the catalog's own multiplicity requirement confirmed rather than assumed.

Case	Faults found	Binding	Lever pulled first
A · Research group	FM-17 Maintenance Trap, FM-01 Parallel Wall, FM-02 Pseudo-Diversity, FM-18 Unshed Q, FM-07 Ceiling Chase (proposed, refused)	delta-Q with u	Shed Q
B · Library room	FM-26 Reverberant Room, FM-01 Parallel Wall, FM-06 Frozen Denominator, FM-24 Abrupt Threshold	u	Absorption
C · Open source	FM-06 Frozen Denominator, FM-20 Negative Fixed Point, FM-18 Unshed Q, FM-15 Founder Battery (x3), FM-14 Sprayed Seed (proposed, refused)	C	Cut C
D · Co-housing	FM-11 Leaky Vessel, FM-21 Quadratic Committee, FM-15 Founder Battery (historical)	eta	Capture devices
E · One desk	FM-11 Leaky Vessel, FM-19 Invisible Infrastructure	eta with delta	The index card

Two patterns are worth naming. **FM-11, the Leaky Vessel, and FM-15, the Founder Battery, appear in three cases each** — conversion failure and person-held storage are the two most common faults in this sample, and both are cheap to fix relative to everything else. And **the fault the system proposed to fix was, in two of four cases, a fault the arithmetic refused**: the research group's recruitment plan and the open-source project's mailing-list announcement were both stopped by a computation, not by an argument.

17.7 7 The essence law

Every chapter before this one has been about *quantity*: how much order is stored, how fast it leaks, how much difference a roster actually contains. This part of the manual is about *identity* — about giving a relationship a name that can be heard rather than a number that must be read.

The instrument is a single formula published in 2022 as *The Effective Method of Sonic Chemistry*. Its original application was the periodic table; its actual reach is wider, and this section establishes exactly how much wider and exactly how much less. We strip the formula to its algebraic core, verify that core against every number the source published, prove two structural theorems about what it can and cannot produce, and mark — in the open, with its badge on — the precise point at which mathematics stops and postulate begins.

The single most important sentence in this section is the one that limits it: the output of the essence law is an ESSENCE frequency, not a physical resonance. No material, molecule, element, room or person is claimed to physically resonate at any value computed here. The law assigns a pitch to a *position in an ordered category scheme*. What is built with that pitch is the subject of Curzi [2026l]. What the pitch *is* — and is not — is this section's.

17.7.1 7.1 Statement

POSTULATE (The Essence Law). Let a discipline be given an explicit, ordered, finite set of essence types. Let `n` be the ordinal position — the **essence index** — of one such type. Then the essence frequency of that type is

```
f(n) = k * n^2                                    [Hz]

k = (9/8)^2 / phi^2 = (81/64) / (phi + 1) = (81/64) * (2 - phi)

phi = (1 + sqrt(5)) / 2 = 1.6180339887498948...

k   = 0.483425732988414 Hz    (0.48342573298841433274106978775598925726  to 40 places)
```

The identity `phi^2 = phi + 1` makes the second and third forms exact restatements of the first; `1/(phi+1) = 2 - phi = 0.38196601125010515`, computed and confirmed below. The identity is Euclid's [Euclid, 1956] and its Fibonacci consequences are standard [Vajda, 1989, Koshy, 2001].

Two labels are attached here and they are not the same label. `k` is a **defined** constant: an exact ratio of exactly specified quantities, computable to arbitrary precision, so quoting it to fifteen significant figures is legitimate and it carries that word wherever it appears at that precision. **Attaching `Hz` to it is a separate and much stronger claim**, and it is a **postulate**; Section 11.1's dimensional audit and Section 11.3's compounding chain treat the two as two steps and never as one. On the appearance of `phi` here specifically, the scepticism literature is binding and is cited at the point of use rather than in a footnote: Markowsky [1992] documents the misattributions that surround the golden section, Fischler [1981] the weakness of the usual empirical claims, and Livio [2002] the history. **The distinction that matters is between CHOOSING `phi` as a constant, which is what happened here, and FINDING it in nature, which did not.** Section 7.6 then shows that the chosen constant cancels from every relationship the law produces.

Two things in that block are proved and two are postulated, and they must never be confused.

Element	Status	Warrant
`f(n) = k n^2` is the exact collapse of the published recipe `((n/phi)*1.125)^2`	**PROVED** (algebra)	Section 7.2
`k = (81/64)/(phi+1) = 0.483425732988414`	**PROVED** (arithmetic)	Section 7.2, computed
Every interval `f(m)/f(n)` is `m^2/n^2` and is independent of `k`	**PROVED** (Result One)	Section 7.4
A set indexed `1..N` has exactly `ceil(N/2)` pitch classes	**PROVED** (Result Two)	Section 7.5

Element	Status	Warrant
That essence types *have* a frequency at all	**POSTULATE**	author's, 2022
That the frequency law is *square* in the index	**POSTULATE**	author's, 2022
That `k` is `(9/8)^2/phi^2` rather than any other constant	**POSTULATE** (design choice)	author's, 2022
That a given ordering of categories is the right one	**REQUIRES JUSTIFICATION**	Section 7.9 and Curzi [2026t]

The mathematics in the middle two rows is unconditional: it is true of `k n^2` for *any* positive `k`, and would remain true if the author had chosen `k = 1`. The unsupported content is entirely in rows five through seven, and it is admitted here on the condition the corpus applies throughout [Curzi, 2026]: **in the open, with its badge on.**

17.7.2 7.2 The honesty clauses

These are mandatory and are restated in Chapters 30 and 31 wherever the law is used. They are not decoration. They are the warrant that lets a postulated law sit inside a body of work whose other results are proved theorems.

H1 — Essence, not resonance. The output is an essence frequency. No claim is made, and none may be implied, that any material, molecule, element, institution, room or person physically resonates at these values. When a room is built whose axial modes sit at essence frequencies [Curzi, 2026l], the room genuinely resonates — because a room of that dimension genuinely has that mode, by ordinary acoustics — but the *essence* does not, and the room's mode is a design target we chose, not a property of the essence we discovered.

H2 — `k` is chosen, not measured. Its form — a Pythagorean interval divided by the golden section — is a design decision by the author of the 2022 document. Nothing in the algebra requires it. Nothing in the intervals depends on it (Section 7.4). It sets the reference pitch and does no other work. Badge it as a postulate every time it appears.

H3 — The square law is a postulate about essence structure. That it coincides in form with the eigenvalue spacing of a one-dimensional infinite square well (`E_n` proportional to `n^2`) is a genuine and interesting mathematical correspondence, developed carefully in Section 7.8. **Coincidence of form is not evidence of mechanism, and this paper asserts no mechanism.**

H4 — The essence index requires a defensible ordering. Applying the law requires an explicit, justified ordinal on the discipline's categories, satisfying the framework's validity conditions [Curzi, 2026t]. An arbitrary ordering yields arbitrary frequencies. Section 7.6 gives the test.

H5 — Everything downstream inherits every clause. Section 8 layers the interaction-surplus Framework on top of this law. That bridge inherits all of the framework's honesty clauses [Curzi, 2026,?,a], including that the surplus functional as deployed is a **lower bound** on an ensemble's information rather than the information itself, that the characterisation theorem is pure mathematics making no empirical or metaphysical claim, and that **no thermodynamic or physical consequence follows from any of it.**

17.7.3 7.3 The collapse

The 2022 source specifies a spreadsheet procedure. Column A holds the atomic number. A separate tab holds the golden ratio. Column C holds

```
(( AN / phi ) * 1.125)^2
```

with `AN` the atomic number and `1.125 = 9/8`. Expand it:

```
f(n) = ( (n / phi) * (9/8) )^2
     = ( n * (9/8) / phi )^2
     = n^2 * (9/8)^2 / phi^2
     = k * n^2                      with k = (9/8)^2 / phi^2
```

That is the whole derivation. There is nothing else in the recipe. **The published procedure is a pure square law with a constant of proportionality, and every appearance of the golden ratio and of the Pythagorean tone in the recipe is absorbed, once and for all, into that single constant.**

Numerically, computed to double precision and cross-checked at fifty decimal digits:

```
phi          = 1.618033988749895
phi^2        = 2.618033988749895
phi + 1      = 2.618033988749895        <- identical, as phi^2 = phi + 1
(9/8)^2      = 1.265625
1/(phi+1)    = 0.38196601125010515
2 - phi      = 0.3819660112501051

k = (9/8)^2 / phi^2   = 0.48342573298841435
k = (81/64)/(phi+1)   = 0.48342573298841435      difference exactly 0.0
k = (81/64)*(2-phi)   = 0.4834257329884143
k = (81/64)*(3-sqrt5)/2 = 0.4834257329884143

k to 50 digits = 0.48342573298841433274106978775598925726023369430832
```

All four closed forms agree **to the last digit double-precision arithmetic can carry**, which is where the display above stops: `0.48342573298841435` is the IEEE-754 double nearest to `k`, and the correctly rounded seventeenth significant figure is `0.48342573298841433`, as the fifty-digit line shows. The discrepancy is in the machine's representation and not in the constant, and it is printed here rather than tidied away because a reader recomputing at higher precision will meet it. The same caution applies to `2400/ln 2`, whose correctly rounded seventeen-figure value is `3462.4680981335122` and whose double representation ends `...123`.

The standard citation form is `k = 0.483425732988414 Hz`, fifteen significant figures, which is the precision at which the source document published — and which is defensible **only** because `k` is a *defined* constant rather than a measured one. Fifteen figures on a measured quantity would be the classic marker of work that has not understood its own error bars.

17.7.4 7.4 Verification against the source document's own numbers

A collapse claim is worth nothing unverified. The source publishes 118 base frequencies — one per element of the periodic table, hydrogen through oganesson. Each was recomputed here from `f(n) = k n^2` using exact fifty-digit arithmetic and compared with the published string.

Result: all 118 published values are reproduced by `k n^2`. The worst relative deviation across the whole table is 9.00e-15, at n = 17 (chlorine) — that is agreement to fourteen significant figures, which is the resolution limit of IEEE double-precision arithmetic.

n	element	published (2022 source)	`k * n^2` recomputed	relative deviation
1	H	0.483425732988417	0.483425732988414	5.5e-15
2	He	1.93370293195367	1.93370293195366	6.6e-15
3	Li	4.35083159689575	4.35083159689573	4.8e-15
8	O	30.9392469112587	30.9392469112585	5.9e-15
12	Mg	69.6133055503321	69.6133055503317	6.3e-15
16	S	123.756987645035	123.756987645034	7.5e-15
26	Fe	326.79579550017	326.795795500168	5.8e-15
29	Cu	406.561041443259	406.561041443256	6.3e-15
50	Sn	1208.56433247104	1208.56433247104	3.4e-15
64	Gd	1980.11180232056	1980.11180232055	7.5e-15
79	Au	3017.05999958071	3017.05999958069	5.4e-15
92	U	4091.71540401396	4091.71540401394	5.2e-15
100	Fm	4834.25732988417	4834.25732988414	5.5e-15
118	Og	6731.21990613072	6731.21990613068	5.8e-15

Note the structure that jumps out of the first and second-to-last rows: `f(100) = 10000 * f(1)` exactly, because `100^2 = 10^4`, so the two published strings carry the same digits with the decimal point moved four places. That is a signature of a pure power law and is visible by eye in the source document itself.

A forensic aside, because it changes what "verified" means. The published values are not merely close to `k n^2` — they are *systematically* high by about 2.8 parts in 10^15, too consistent to be rounding noise. Inverting the arithmetic to recover the `phi` the spreadsheet actually used:

```
phi implied by published f(1) = 1.61803398874989...
true phi                      = 1.61803398874989484...
ratio                         = 1 - 2.76e-15
```

The implied value is exactly `phi` truncated to fifteen significant figures: `1.61803398874989`. Recomputing the entire table with that literal:

```
k(phi = 1.61803398874989) = 0.48342573298841723
-> 109 of 118 published values match to ALL FIFTEEN displayed digits
-> worst relative deviation 3.43e-15 (last-place float noise on the remaining nine)
```

So the collapse is not approximate. It is exact, and the residual discrepancy is fully explained by the spreadsheet having stored a fifteen-digit `phi` rather than a full-precision one. We can say that the published table *is* `k n^2`, and we can say precisely which `k`. This is the standard the paper holds itself to: not that the numbers look right, but that the arithmetic which produced them is given and the residue is accounted for.

One further textual erratum, recorded because it will otherwise trip any reader who reimplements the method. The source document instructs: *"In one cell type in the equation `1+5^.5/2`"*. Under standard operator precedence that expression evaluates to `1 + (sqrt(5)/2) = 2.118033988749895`, which is not `phi`. Recomputing `f(1)` with it gives `0.28212345562925867`, which matches nothing in the published table. **The spreadsheet plainly contained `(1+5^0.5)/2`; the document's prose dropped the parentheses.** Use `(1 + sqrt(5))/2`.

17.7.5 7.5 The constant, and what it is for

`k` is a Pythagorean interval divided by the golden section, and it is worth being precise about which. `9/8` is the **Pythagorean whole tone** — `(3/2)^2 / 2`, 203.910 cents. Its square, `81/64`, is the

Pythagorean ditone, the major third of three-limit tuning, 407.820 cents, sharper than the just major third `5/4` (386.314 cents) by the syntonic comma [Barbour, 1951, Lindley and Turner-Smith, 1993, Doty, 2002]. `phi^2 = phi + 1 = 2.618033988749895`. So `k` is the ditone reduced by `phi^2`.

The source document offers a rationale — that elements are standing waves on a phi spiral with eight per 360 degrees, hence the factor `9/8` to "establish that symmetry breaking in the spatial orientation of the golden spiral". **That rationale is a postulate. It is recorded because it is the source author's stated reason, and badged because no derivation of it exists.** It is also exactly the kind of claim Markowsky [1992] and Fischler [1981] were written about, and a reader should weigh it accordingly.

What matters engineering-wise is the subject of the next section: *none of this survives into any relationship the law produces.* `k` sets the tuning fork. It sets nothing else.

17.7.6 7.6 RESULT ONE — every interval is a ratio of squares, and `k` cancels

Theorem (Result One). For any two essence indices `m, n >= 1`,

```
f(m)      k * m^2      m^2
---- =   -------  =   ---
f(n)      k * n^2      n^2
```

Proof. Immediate. `k` is a common factor of numerator and denominator. It cancels for every pair, for every `k > 0`. QED.

The proof is one line. The consequence is the whole chapter.

The golden ratio and the 9/8 set the absolute reference pitch and NOTHING ELSE. Every relationship in the system is a ratio of squared integers. If you replace `k` with 1, or with 440, or with the mass of the sun, every interval in the system is unchanged. `phi` does no relational work. It is a transposition, and a transposition is the one musical operation that changes nothing structural.

This is a demotion and it should be felt as one. A reader coming to this material from the wider literature of golden-ratio numerology will expect `phi` to be doing something. It is not. The harmonic content of the essence law is generated entirely by the integers.

Here is the interval lattice, computed. "Cents" is `1200*log2(ratio)`; "folded cents" is that value modulo 1200, i.e. the pitch class; "dev" is the deviation from the nearest twelve-tone equal-tempered interval.

m:n	ratio `m^2/n^2`	decimal	octaves	cents	folded cents	nearest 12-TET	dev (cents)
2:1	4/1	4.000000	2.0000	2400.00	0.00	unison	+0.00
3:2	9/4	2.250000	1.1699	1403.91	203.91	major 2nd	+3.91
4:3	16/9	1.777778	0.8301	996.09	996.09	minor 7th	-3.91
5:4	25/16	1.562500	0.6439	772.63	772.63	minor 6th	-27.37
6:5	36/25	1.440000	0.5261	631.28	631.28	tritone	+31.28
5:3	25/9	2.777778	1.4739	1768.72	568.72	tritone	-31.28
8:5	64/25	2.560000	1.3561	1627.37	427.37	major 3rd	+27.37
3:1	9/1	9.000000	3.1699	3803.91	203.91	major 2nd	+3.91
4:1	16/1	16.000000	4.0000	4800.00	0.00	unison	+0.00
5:2	25/4	6.250000	2.6439	3172.63	772.63	minor 6th	-27.37
7:4	49/16	3.062500	1.6147	1937.65	737.65	perfect 5th	+37.65
7:5	49/25	1.960000	0.9709	1165.02	1165.02	unison	-34.98
7:6	49/36	1.361111	0.4448	533.74	533.74	perfect 4th	+33.74
8:7	64/49	1.306122	0.3853	462.35	462.35	perfect 4th	-37.65
9:8	81/64	1.265625	0.3399	407.82	407.82	major 3rd	+7.82
16:9	256/81	3.160494	1.6601	1992.18	792.18	minor 6th	-7.82
10:9	100/81	1.234568	0.3040	364.81	364.81	major 3rd	-35.19
11:8	121/64	1.890625	0.9189	1102.64	1102.64	major 7th	+2.64
13:8	169/64	2.640625	1.4009	1681.06	481.06	perfect 4th	-18.94
15:8	225/64	3.515625	1.8138	2176.54	976.54	minor 7th	-23.46
5:1	25/1	25.000000	4.6439	5572.63	772.63	minor 6th	-27.37

Read the first four rows and the character of the system is visible. **The essence law doubles the distance of every classical interval.** A frequency ratio of `3:2` — the perfect fifth of

ordinary acoustics — becomes `9:4`, an octave plus a major second, because squaring doubles the logarithm. The perfect fourth `4:3` becomes `16:9`, a minor seventh. The octave `2:1` becomes two octaves. This is a general fact, not a coincidence: since pitch interval is logarithmic and the law squares,

```
interval in cents between indices m and n  =  2 * (classical interval between m and n)
```

Squaring the index doubles every interval in cents. The essence law is the classical harmonic series stretched by a factor of exactly two in log-frequency space. That single sentence predicts every row of the table above and is worth carrying.

The upper-triangular lattice for indices 1 through 6, each cell giving exact ratio, cents, and folded cents:

```
each cell:  ratio  |  cents  |  folded cents

        n=1                 n=2                 n=3
 m=2   4/1 2400.0    0.0
 m=3   9/1 3803.9  203.9   9/4 1403.9  203.9
 m=4  16/1 4800.0    0.0   4/1 2400.0    0.0  16/9   996.1  996.1
 m=5  25/1 5572.6  772.6  25/4 3172.6  772.6  25/9  1768.7  568.7
 m=6  36/1 6203.9  203.9   9/1 3803.9  203.9   4/1  2400.0    0.0

                         n=4                 n=5
 m=5  25/16  772.6  772.6
 m=6   9/4  1403.9  203.9  36/25  631.3  631.3
```

Notice how often the folded-cents entry repeats: 0.0 at (2,1), (4,1), (4,2); 203.9 at (3,1), (3,2), (6,1), (6,2), (6,4). That repetition is not an accident of the display. It is Result Two.

17.7.7 7.7 RESULT TWO — octave folding is governed by the odd part of the index

Musical pitch is naturally taken modulo the octave: two tones an octave apart share a *pitch class* and are heard as “the same note, higher”. The question “how many distinct notes does an essence set contain?” is therefore the question “how many distinct pitch classes does it contain?”, and the essence law answers it with unusual sharpness.

Lemma. `f(2n) = k(2n)^2 = 4 k n^2 = 4 f(n)`. Doubling the index multiplies the frequency by four, which is exactly two octaves. Therefore `n` and `2n` always share a pitch class.

Theorem (Result Two). Write every index uniquely as `n = 2^a * m` with `m` odd (`m` is the *odd part* of `n`, `a` its 2-adic valuation). Then

```
f(n) = k * (2^a * m)^2 = k * 4^a * m^2 = f(m) * 4^a = f(m) * 2^(2a)
```

so `f(n)` is exactly `2a` octaves above `f(m)`. **Two essences share a pitch class if and only if their indices have the same odd part.**

Proof of “only if”. Suppose `n = 2^a m` and `n' = 2^b m'` share a pitch class, i.e. `f(n)/f(n') = 2^j` for some integer `j`. Then `4^a m^2 / (4^b m'^2) = 2^j`, so `m^2/m'^2 = 2^(j − 2a + 2b)`. The left side is a ratio of odd squares, hence an odd rational in lowest terms; the right side is a power of two. The only odd power of two is `2^0 = 1`, so `m^2 = m'^2` and, since both are positive, `m = m'`. QED.

Corollary. An essence set indexed `1..N` contains exactly `ceil(N/2)` distinct pitch classes.

Proof. The odd part of any `n <= N` is odd and at most `n <= N`, so every realised class label is an odd number in `[1, N]`. Conversely every odd `m <= N` is itself in the index set and is its own

odd part, so every such label is realised. Hence the class labels are exactly the odd integers in `[1, N]`, of which there are `ceil(N/2)`. QED.

Here is the full index-to-class table for `N` up to 24, computed:

n	`n = 2^a * m`	odd part m	a	`f(n) = k n^2` (Hz)	class representative `f(m)` (Hz)
1	2^0 * 1	1	0	0.483426	0.483426
2	2^1 * 1	1	1	1.933703	0.483426
3	2^0 * 3	3	0	4.350832	4.350832
4	2^2 * 1	1	2	7.734812	0.483426
5	2^0 * 5	5	0	12.085643	12.085643
6	2^1 * 3	3	1	17.403326	4.350832
7	2^0 * 7	7	0	23.687861	23.687861
8	2^3 * 1	1	3	30.939247	0.483426
9	2^0 * 9	9	0	39.157484	39.157484
10	2^1 * 5	5	1	48.342573	12.085643
11	2^0 * 11	11	0	58.494514	58.494514
12	2^2 * 3	3	2	69.613306	4.350832
13	2^0 * 13	13	0	81.698949	81.698949
14	2^1 * 7	7	1	94.751444	23.687861
15	2^0 * 15	15	0	108.770790	108.770790
16	2^4 * 1	1	4	123.756988	0.483426
17	2^0 * 17	17	0	139.710037	139.710037
18	2^1 * 9	9	1	156.629937	39.157484
19	2^0 * 19	19	0	174.516690	174.516690
20	2^2 * 5	5	2	193.370293	12.085643
21	2^0 * 21	21	0	213.190748	213.190748
22	2^1 * 11	11	1	233.978055	58.494514
23	2^0 * 23	23	0	255.732213	255.732213
24	2^3 * 3	3	3	278.453222	4.350832

And the class count, computed and checked against `ceil(N/2)` for every `N` up to 24:

N	distinct odd parts realised	count	`ceil(N/2)`
1	1	1	1
2	1	1	1
3	1,3	2	2
4	1,3	2	2
5	1,3,5	3	3
6	1,3,5	3	3
7	1,3,5,7	4	4
8	1,3,5,7	4	4
9	1,3,5,7,9	5	5
10	1,3,5,7,9	5	5
11	1,3,5,7,9,11	6	6
12	1,3,5,7,9,11	6	6
13, 14	1,3,5,7,9,11,13	7	7
15, 16	1,3,5,7,9,11,13,15	8	8

N	distinct odd parts realised	count	`ceil(N/2)`
17, 18	1,3,...,15,17	9	9
19, 20	1,3,...,17,19	10	10
21, 22	1,3,...,19,21	11	11
23, 24	1,3,...,21,23	12	12

Twelve essences yield only six pitch classes. That is the available harmonic vocabulary, and it is a hard structural limit rather than a stylistic choice. The same rule is reached independently in Curzi [2026], whose Theorem 7.1 states that the pitch class of `f(n) = k n^2` is the odd part of `n`, so that `1..N` realises exactly `ceil(N/2)` classes and `N = 12` realises six — and applies it to mood and trait sets rather than to elements. The compression ratio is exactly 2 for every `N`, which is the flattest possible statement of the constraint: *half of any essence set is harmonically invisible.* You may have 118 elements; you have 59 notes. You may have 24 departments; you have 12 notes. Doubling the roster does not double the palette in any way you can hear beyond the doubling already accounted for.

17.7.8 7.8 The pitch-class lattice is generated by the odd primes

Result Two admits a sharper algebraic statement that is useful when a set is being designed rather than analysed.

Octave-reduce each odd class label `m` by dividing `m^2` by the largest power of two below it. The result lies in `[1,2)`:

odd m	`m^2`	a	folded ratio `m^2/2^a`	cents	nearest 12-TET	dev
1	1	0	1.000000000	0.00	unison	+0.00
3	9	3	1.125000000	203.91	major 2nd	+3.91
5	25	4	1.562500000	772.63	minor 6th	-27.37
7	49	5	1.531250000	737.65	perfect 5th	+37.65
9	81	6	1.265625000	407.82	major 3rd	+7.82
11	121	6	1.890625000	1102.64	major 7th	+2.64
13	169	7	1.320312500	481.06	perfect 4th	-18.94
15	225	7	1.757812500	976.54	minor 7th	-23.46
17	289	8	1.128906250	209.91	major 2nd	+9.91
19	361	8	1.410156250	595.03	tritone	-4.97
21	441	8	1.722656250	941.56	major 6th	+41.56
23	529	9	1.033203125	56.55	minor 2nd	-43.45

Because `m` is odd, factor it as `m = p1^e1 * p2^e2 * ...` over odd primes. Then `m^2 = p1^(2e1) * p2^(2e2) * ...`, and the folded class of `m` is the product of the folded classes of `p_i^2` taken `e_i` times, modulo octaves. Therefore:

The essence pitch-class lattice is the free abelian group generated by the octave-reductions of the squares of the odd primes. The generators, in ascending prime order:

```
prime 3   ->  9/64 folded to 9/8      = 203.910 cents      (Pythagorean whole tone)
prime 5   ->  25/16                   = 772.627 cents
prime 7   ->  49/32                   = 737.652 cents
prime 11  ->  121/64                  = 1102.636 cents
prime 13  ->  169/128                 = 481.062 cents
```

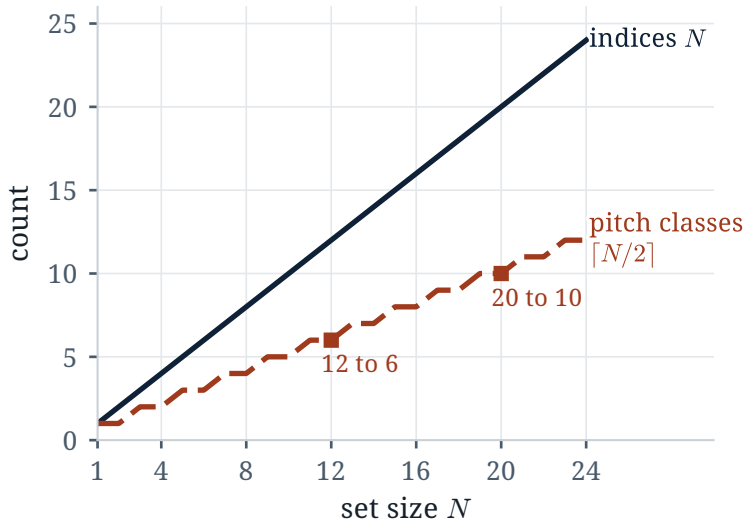


Figure 17.1: **Half of any essence set is harmonically invisible: the compression ratio is exactly 2 at every N.** Twenty indices realise ten pitch classes and twelve realise six, because the class label is the odd part of the index and the odd integers in $[1, N]$ number exactly $\lceil N/2 \rceil$. Exact arithmetic, checked against the class-count table of Section 7.7 for every N up to 24; generated by `figs/src/CCP-017_fig04.py`.

Check the composite classes against the generators: class 9 = 3^2, so its folded class must be `(9/8)^2 = 81/64 = 1.265625`, which is exactly the table's entry and `2 * 203.910 = 407.820` cents. Class 15 = 3*5, so `(9/8)*(25/16) = 225/128`, folded `225/128 = 1.7578125`, and `203.910 + 772.627 = 976.537` cents, matching the table's 976.54. Class 21 = 3*7: `203.910 + 737.652 = 941.562`, matching 941.56.

This is not decoration. It has a direct design consequence: **an essence set that runs `1, 2, 4, 8, 16, ...` produces exactly one note; an essence set whose indices are the odd primes produces the maximally independent harmonic vocabulary available at that size.** Where an ordinal scheme admits legitimate freedom — Section 7.9 states how much — prime indices buy harmonic independence and powers of two throw it away.

Note also a genuine curiosity: `9/8` appears twice here for unrelated reasons. It is inside `k` because the author put it there, and it is the folded class of index 3 because `3^2 = 9` and `2^3 = 8`. The second survives any change of `k`; the first does not.

Four exact design laws follow once the lattice has to be rendered, and they are constraints rather than preferences. They follow from the Gabor time–bandwidth limit [Gabor, 1946, Curzi, 2026z] and from the arithmetic above, and they are stated here because they bound what an essence set may be before anyone chooses one. All four are computed, not asserted.

Law 1 — the audible-set ceiling. With a single integer octave transposition `f * 2^j` and a band of 20 Hz to 20 kHz, the whole set `n = 1..N` fits only if `k * 2^j >= 20` and `k * N^2 * 2^j <= 20000`. The first forces `j >= 6`; the second then caps `N`. **At most `N = 25` essences fit under a single transposition, and only at `j = 6`** — `j = 7` admits 17, `j = 8` admits 12. **A twelve-house set fits comfortably; a thirty-two-element set does not, and requires split transposition**, which is a design decision with consequences and must be declared wherever it is used.

Law 2 — the minimum-gap cliff. Fold the set's pitch classes into one octave and measure the smallest gap. For `N` from 7 to 16 it is exactly **50/49, or 34.976 cents** — the near-coincidence of classes 7 and 5 in the table above. **The moment index 17 enters the set, the minimum gap collapses to exactly 289/288, or 6.001 cents**, because `17^2 = 289` lands beside `9/8 = 288/256`, the folded class of index 3. **The boundary is exact and it is arithmetic: a set of sixteen and a set of seventeen are different design regimes.**

Three ratios, and they are not the same number. The cliff is quoted three ways and each

is correct in its own units, so which one is meant must be stated:

Quantity	Value	What it measures
Gap in **cents**	`ln(50/49)/ln(289/288)` = **5.828**	how much smaller the interval is, logarithmically
Cycles of the tone	`288/49` = **5.878**	3.899 cycles becomes 22.918 (Law 3)
Achievable **rate**	**8, exactly**	symbols per second, and this is the design limit

The rate figure is the one that binds, and it is exactly 8 rather than either of the others. The reason is that the two binding pairs do not sit at the same height in the octave: the 50/49 pair straddles `49/32`, the 289/288 pair straddles `9/8 = 288/256`. Gabor resolution depends on the **absolute** gap in hertz, which is `f/32` before the cliff and `f/256` after — a ratio of exactly 8 at every base pitch. Dividing fractional gaps instead (`288/49`) silently assumes both pairs sit at the same frequency, and they do not. The derivation is Curzi [2026z]. The intermediate cycle counts quoted in that row are taken from it rather than recomputed here, and Section 12 records the one-per-cent discrepancy a direct recomputation from the cents figures produces.

Law 3 — the surplus-resolution bound. Resolving a difference `dF` in surplus requires at least **`1 / (8 pi dF)` cycles** of the rendered tone, and the bound is independent of the transposition because it constrains a ratio. At `dF = 0.01` nats that is 3.98 cycles; at `dF = 0.001`, 39.79. **This is the exact Gabor bridge to the surplus functional: it prices, in time, what it costs to hear a difference in surplus.** Section 8.9 works the consequence for measurement precision.

Law 4 — the rectangular gate is disqualified. A rectangular on/off envelope has **no finite spectral second moment** — its `sinc^2` tails fall as `1/f^2` and the second-moment integral is numerically divergent — so it has no defined bandwidth and cannot serve as a rendering envelope at all. **Any implementation must use a windowed envelope.** Curzi [2026l] specifies one.

17.7.9 7.9 What an essence index is, and what makes an ordering defensible

An **essence index** is the ordinal position of an essence type within its discipline's ordered category set. Atomic number is one instance of an essence index. It is not the definition of one, and the 2022 source's decision to apply the formula to chemistry is an application, not a restriction.

Atomic number is, however, the *best possible* case, and understanding why it is shows what is needed in every other case. It is **intrinsic** (a count of protons, a property of the thing rather than of our filing system), **total** (every element has exactly one and no two share one), **dense from 1** (no gaps, so `ceil(N/2)` applies without qualification), **independently re-derivable** (anyone can recount and get the same integer), and **stable** (it does not change when the discipline reorganises its textbooks).

Most disciplines have nothing like this. The corpus's formula-validation protocol states the conditions under which the framework's formal structure may be applied to a target domain at all [Curzi, 2026t], and those conditions govern here too. They require, before any modelling map is applied:

1. **State encoding** — a map from target-domain items to normalised vectors, so every represented state has a determinate location in the model.
2. **Category assignment** — a finite partition into idealised categories, represented by blocks `H_1, ..., H_N`.
3. **Geometric dependence** (the S1 analogue) — interaction modelled through a scalar similarity parameter such as an angle, cosine, or inner-product-derived quantity.

4. **Zero baseline** (S2 analogue) — parallel or representationally equivalent states correspond to zero modelled surplus.
5. **Affine effective-count law** (S3 analogue) — adopted as a modeling assumption, **not read directly off the world.**
6. **Normalisation** (S4 analogue) — the maximum modelled surplus is finite and depends on the chosen category count `N`.

The protocol is explicit about the exclusion that follows: a target domain lacking a stable encoding, a defensible category scheme, or an interpretable angle-like geometry **is outside the scope of the model.** The same exclusion applies to the essence law, and this paper enforces it.

To those six, the house method adds one more, specific to the essence law because the essence law needs an *order* and not merely a *partition*:

7. **Declared ordinal rule.** The rule that assigns index `n` to category `i` must be published alongside the frequencies, must be applicable by a third party without consulting the author, and must include its tie-break procedure. Where the ordering came from a list written down in that sequence for convenience, that must be stated: it is a weak ordering and its weakness travels with the numbers.

An ordering test you can run in ten minutes. Take the essence set, permute two indices, recompute, and report the change. Worked, for a twelve-element set:

- Swap indices 3 and 12. Frequencies move from 4.350832 Hz to 69.613306 Hz and back — a factor of exactly 16, four octaves. **But the pitch classes do not change at all**, because 3 and 12 share the odd part 3. This permutation is harmonically invisible.
- Swap indices 3 and 5. Class 3 (folded ratio 9/8, 203.91 cents) becomes class 5 (25/16, 772.63 cents). The pitch class moves by `772.63 − 203.91 = 568.72` cents — nearly half an octave, a tritone's worth of change, from swapping two adjacent list positions.

If a plausible re-ordering of the categories moves a pitch class by half an octave, the ordering is load-bearing and must be justified before any of the frequencies are published. That is the test. It costs nothing and it is the single most commonly skipped step in the procedure.

17.7.10 7.10 The composition rule, and its collapse

The 2022 source also gives a rule for combining essences. For `H2O`: three atoms total, so raise the hydrogen frequency to the `2/3` power, the oxygen frequency to the `1/3` power, multiply, and square:

```
f(H2O) = ( f(H)^(2/3) * f(O)^(1/3) )^2
```

This too collapses. Let a compound contain essence indices `n_i` with multiplicities `c_i`, and let `T = sum(c_i)` be the total count. Define `G` as the count-weighted geometric mean of the indices:

```
G = ( prod_i  n_i^(c_i) )^(1/T)
```

Then

```
f_compound = ( prod_i f(n_i)^(c_i/T) )^2
           = ( prod_i (k n_i^2)^(c_i/T) )^2
           = ( k^(sum c_i / T) * prod_i n_i^(2 c_i / T) )^2
           = ( k * G^2 )^2
           = k^2 * G^4
```

The composition rule is `k^2` times the fourth power of the count-weighted geometric mean of the constituent essence indices. Test it on water: `G = (1^2 * 8^1)^(1/3) = 8^(1/3) = 2.00000000`, so `f = k^2 * 16 = 3.7392070290462`, against the source's published `3.73920702904621`. Agreement to fourteen significant figures.

Verified across every composite value the source publishes — 24 pharmacological compounds, 8 nucleotides in the carbon column, and 6 each in the silicon- and germanium-substituted columns, 44 values in all:

37 of 44 published composite values are reproduced by `f = k^2 G^4` to better than one part in 10^12. Seven disagree. The seven, with the size of the disagreement:

formula	published	`k^2 G^4`	relative deviation
O2	11.1246117974981	957.23699944	85.05
C5H6N2O2 (thymine)	4.66961137812091	21.80527042	3.670
C12H17N2O4P	18.5343343733681	13.36731640	0.279
C257H383N65O77S6	14.2762378831273	11.27999043	0.210
C43H66N12O12S2	12.8278625485662	11.29835161	0.119
C77H120N18O26S	11.0756447985031	10.43459440	0.058
C158H251N39O46S	9.46036014007241	9.41751106	0.005

These seven are spreadsheet errors in the source, not evidence of a different rule, and the O2 row proves it. For a single-element species the count-weighted geometric mean is the element index itself, independent of how many atoms are present: `G = 8` for `O`, `O2`, `O3`, `O_x`, always. Therefore `f = k^2 * 8^4 = 957.23699944` for *any* pure-oxygen formula, and **no atom count whatever can produce the published 11.1246117974981.** That row cannot have come from the stated rule. A rule error would fail uniformly across the table; a data-entry error fails on scattered rows. We observe scattered rows. The rule is confirmed and seven cells are errata.

Two engineering notes, both honest flags:

The squaring destroys dimensional consistency. `f(n)` is quoted in Hz; `k^2 G^4` carries Hz-squared if the units are taken seriously. The source reads the squaring as a change of *structural level* — "each time you square the result you establish a level of harmonic molecular structure" — which is a reasonable reading of a designed formalism and is what this paper adopts. **But the composite number's "Hz" label is a convention, not a dimensional quantity, and it must not enter an acoustic calculation without an explicit re-anchoring step.** Curzi [2026l] re-anchors by octave folding, a ratio operation and therefore unit-agnostic; that is the safe path.

Composition flattens the vocabulary further. Because `G` is a geometric mean, low-index constituents pull it toward 1 and high-index constituents raise it slowly. Across the 24 pharmacological compounds in the source, `G` ranges only from 1.84 to 4.30 — a factor of 2.34, which is a factor of 30 in `k^2 G^4`, or about 4.9 octaves. **All twenty-four compounds occupy under five octaves of essence pitch and eighteen of them occupy under two.** Composition is strongly contractive on the pitch space and should be expected to cluster, not disperse.

17.7.11 7.11 The infinite square well: a correspondence of form

A particle of mass `m` confined to a one-dimensional box of width `L` with infinitely high walls has energy eigenvalues

```
E_n = n^2 h^2 / (8 m L^2)          n = 1, 2, 3, ...
```

Divide by Planck's constant to get the associated frequency:

```
nu_n = E_n / h = n^2 * h / (8 m L^2)
```

which is `constant * n^2`. **This is the same functional form as the essence law, with the constant `h/(8 m L^2)` playing the role of `k`.**

The correspondence can be pushed to a specific number. Setting `h/(8 m_e L^2) = k` with `m_e` the electron mass and CODATA values `h = 6.62607015e-34 J s` (exact, by SI definition) and `m_e = 9.1093837015e-31 kg`:

```
L = sqrt( h / (8 * m_e * k) ) = 0.013714300 m = 1.3714 cm

check: h / (8 * m_e * L^2) = 0.483425732988414    <- equals k
```

So there exists a box, 1.37 centimetres wide, for which an electron's level frequencies coincide numerically with the essence frequencies of a set indexed from 1.

This is a coincidence of form and nothing more, and the manual says so in the same breath in which it states the correspondence. Three reasons:

1. **`k` was chosen, not measured.** Because `k` is a free constant, a box width can always be found to match it. The existence of `L` is not a discovery; it is a division. Had `k` been `(4/3)^2/phi^3` the box would be a different size and the "correspondence" equally available. **A correspondence that can be manufactured for any constant carries no information about any constant.**
2. **The 1.37 cm box has no physical significance.** Not the size of an atom (five orders of magnitude too big), not a lattice spacing, not a wavelength of anything in the problem.
3. **`n^2` is the most common eigenvalue pattern in elementary physics**, appearing wherever a wave is confined between fixed boundaries with quantised half-wavelengths. Sharing it is not remarkable.

What *would* be evidence of mechanism, stated plainly: a preregistered prediction, derived from the essence law before measurement, of a physical observable that the law gets right and no simpler model does, on a held-out sample, with the ordering rule fixed in advance. The corpus protocol [Curzi, 2026i, Platt, 1964] applies unmodified. **This paper makes no such prediction and offers no such evidence. The essence law is used here as a naming and design instrument, and nothing in Chapters 29 through 31 depends on it being anything more.**

The correspondence is still worth carrying for one legitimate reason: it tells you what the law's *shape* implies. In a square well, level spacing grows without bound while the *log*-spacing collapses; the low levels are where the interesting structure is. The essence law inherits exactly that character. The interval between neighbouring indices is `2 * 1200 * log2(n/(n-1))` cents: 1403.9 at `n = 2`, 177.7 at `n = 20`, 69.8 at `n = 50`. **The first few essence indices carry almost all of the harmonic interest; past about index 20 the set thins through equal temperament and out into microtones.** That is a constraint on set design rather than a preference.

17.7.12 7.12 Specification: fourteen requirements on a published essence set

Numbered, checkable requirements. A published essence-frequency set that fails any of these is not conformant to this paper.

ES-1 — State the law in its collapsed form. Publish `f(n) = k n^2` with `k = 0.483425732988414 Hz`, not the spreadsheet recipe. The recipe hides that a single constant does all the reference-setting work, and it invites the parenthesis error documented in Section 7.4.

ES-2 — Badge `k` at every appearance. `k` is a chosen reference constant; its form is a design decision by the author of the 2022 source. Never present it as measured, derived, or fundamental.

ES-3 — Publish the ordinal rule with the frequencies. State how index `n` was assigned to category `i`, in terms a third party can apply without consulting the author, including tie-breaks. A frequency table without its ordinal rule is not a result; it is a list.

ES-4 — Run and report the permutation sensitivity test. Swap at least three plausible index pairs, recompute, report the maximum pitch-class displacement in cents. Where any displacement exceeds 100 cents the ordering is load-bearing, and must be declared and justified (Section 7.9).

ES-5 — Verify validity conditions 1, 2 and 6 before applying the law [Curzi, 2026t]. State the encoding, the partition, and `N`. Where any of the three cannot be stated, the domain is out of scope.

ES-6 — Report `ceil(N/2)` alongside `N`. This prevents the commonest design error: planning a composition around `N` voices when only `ceil(N/2)` are distinguishable. (Section 7.7)

ES-7 — Never publish a frequency without its octave. Every value carries its raw Hz or its transposition exponent `j` with `f_published = f_raw * 2^j`, plus the band. Two values quoted without their `j` cannot be compared.

ES-8 — Quote to no more than fifteen significant figures. That is the source's precision and the limit of double-precision arithmetic. For any built application four figures already exceed the tuning resolution of the available physical instruments.

ES-9 — Restate honesty clause H1 in any document carrying a frequency table. "This is an essence frequency, not a physical resonance. No claim is made that any material, molecule, element, institution or person physically resonates at these values." On the same page as the numbers.

ES-10 — When composing, publish `G` and `T` with the result. `f = k^2 G^4` is the closed form; The count-weighted geometric mean and the member count are published so that the value is reproducible.

ES-11 — Re-anchor composite values by ratio, never by unit. Composite outputs carry `Hz^2`. Bring them into band by octave folding — a ratio operation — and state the fold. A composite value is never fed into an acoustic formula expecting Hz. (Section 7.10)

ES-12 — Note the square-well disclaimer in the same paragraph as the correspondence. Never in a later section, never in a footnote. Coincidence of form is not evidence of mechanism. (Section 7.11)

ES-13 — Recompute the source's published values before relying on them. Seven of the 44 composite values in the 2022 document are errata (Section 7.10). Recompute; do not transcribe.

ES-14 — Prefer prime indices when index assignment is genuinely free. Odd-prime indices maximise harmonic independence; powers of two collapse onto one class. Where the ordinal rule does *not* admit freedom, ES-3 governs and this requirement is void. (Section 7.8)

17.7.13 7.13 Failure modes of the essence law

Each entry gives the fault, the parameter or claim it damages, its diagnostic signature, and the fix.

FM-E1 THE ARBITRARY ORDINAL

Damages: every frequency in the set. *Description.* An ordering chosen for convenience — alphabetical, historical, the order items came up in a meeting — is fed into `f(n) = k n^2` and the output treated as intrinsic. *Signature.* Asked why a given category carries index 10, the answer is that that is where it sits on the sheet. *Fix.* ES-3 and ES-4. Publish the rule and run the permutation test. Where displacement is large and the rule is weak, **the correct output is a range rather than a frequency**, and the range is what is reported.

FM-E2 THE PHYSICAL RESONANCE CLAIM

Damages: the credibility of everything else in the volume. *Description.* A sentence of the form "iron resonates at 326.80 Hz". It does not. Iron's essence index is 26 and `k * 676 = 326.795795500168` — a number attached to a position in an ordering. *Signature.* Any sentence in which a material, molecule, body part or person is the *subject* of the verb "resonates" with an essence frequency as its object. *Fix.* H1, restated per ES-9. Rewrite so the essence set is the subject: "index 26 of the periodic ordering carries essence pitch class 13."

FM-E3 THE UNBADGED CONSTANT

Damages: the reader's ability to tell proof from postulate. *Description.* `k = (9/8)^2/phi^2` is presented as though the golden ratio had been discovered in the data rather than placed in the formula by hand. *Signature.* Prose in which `phi` is described as "emerging", "appearing", or "governing". *Fix.* ES-2, plus Result One: state in the same passage that `k` cancels from every relationship and therefore governs nothing but the reference pitch.

FM-E4 THE PHANTOM VOCABULARY

Damages: every downstream decision that assumed `N` distinguishable tones. *Description.* A twelve-category scheme is sonified as twelve voices; six are octave duplicates and the "twelve-part harmony" is a six-note chord doubled. *Signature.* Two "different" voices a listener cannot tell apart; two indices whose ratio is a power of four. *Fix.* ES-6. Design to `ceil(N/2)`. Twelve distinct tones require at least twenty-three justified essence indices.

FM-E5 THE UNSTATED OCTAVE

Damages: comparability between any two published tables. *Description.* One table publishes raw values, another the same set folded into 220-440 Hz; compared directly they appear to disagree by orders of magnitude. *Signature.* Values from two sources differing by a clean power of two. *Fix.* ES-7. Every value carries its `j` and its band, always.

FM-E6 PRECISION THEATRE

Damages: trust, and nothing else — but that is enough. *Description.* Publishing `326.795795500168 Hz` as a tuning target for a physical object whose own frequency drifts a percent with temperature, and whose best tuner resolves 0.1 cent. *Signature.* More than about six significant figures on any value that will be built rather than computed. *Fix.* ES-8. Fifteen figures in the derivation table, four in the build specification.

FM-E7 THE SQUARE-WELL SLIDE

Damages: the boundary between this section and physics. *Description.* The `n^2` correspondence is noted, then the essence law is called "quantum", and a step later energy levels are claimed. *Signature.* "Quantum", "eigenvalue" or "energy level" used without a disclaimer in the same paragraph. *Fix.* ES-12 — and the harder fix: where the argument cannot be made without the physics vocabulary, it was leaning on borrowed authority and must be rewritten without it.

FM-E8 COMPOSITION UNIT DRIFT

Damages: any acoustic calculation downstream. *Description.* A composite value from `k^2 G^4` — dimensionally Hz-squared — is fed into `L = c/(2f)` or a Sabine calculation as though it

were Hz. *Signature.* A tube length or absorption target absurd by three or more orders of magnitude. *Fix.* ES-11. Fold into the target band first, then use ratios only.

FM-E9 THE SPREADSHEET ORACLE

Damages: seven values in 44, unpredictably. *Description.* Transcribing published values instead of recomputing them, inheriting the source's data-entry errors, then defending them because they are "published". *Signature.* A value that resists reproduction by the stated rule while its neighbours reproduce to fourteen figures. *Fix.* ES-13. Where recomputation disagrees by more than 1e-12 relative, the published number is wrong until shown otherwise, and the disagreement is printed, as Section 7.10 does.

FM-E10 THE POWERS-OF-TWO SET

Damages: the entire harmonic output. *Description.* An essence set built on indices `1, 2, 4, 8, 16` — a natural doubling scheme. Every member has odd part 1. **The set has exactly one pitch class.** *Signature.* All voices in unison across octaves; a "chord" with no interval in it. *Fix.* ES-14 where index assignment is free; where it is not, accept and report the result — that a category scheme is harmonically degenerate is a real finding about the scheme.

17.7.14 7.14 Worked instances

7.14.1 The periodic table, first twenty elements, carried end to end

This is the source document's own domain and therefore the strongest available test: we can check every intermediate against a published value.

Step 1 — Establish the ordinal. The essence index is atomic number: the count of protons in the nucleus. It satisfies all five properties of Section 7.9 — intrinsic, total, dense from 1, independently re-derivable, stable. This is the best case and it is why the source chose it. Validity condition 2 [Curzi, 2026t] is satisfied by the partition of matter into elements; condition 6 is satisfied with `N = 118`.

Step 2 — Compute raw essence frequencies. `f(n) = 0.483425732988414 * n^2`.

Step 3 — Compute odd parts and pitch classes. By Result Two.

Step 4 — Fold into an audible band. Band chosen: `[220, 440)` Hz — one octave, anchored on A3 = 220 Hz. `f_folded = f * 2^j`, `j` the integer bringing `f` into the band. The band and `j` are published for every value, per ES-7.

n	element	`f(n)` (Hz)	odd part	class	j	`f` folded to [220,440)
1	H	0.483426	1	PC1	+9	247.5140
2	He	1.933703	1	PC1	+7	247.5140
3	Li	4.350832	3	PC3	+6	278.4532
4	Be	7.734812	1	PC1	+5	247.5140
5	B	12.085643	5	PC5	+5	386.7406
6	C	17.403326	3	PC3	+4	278.4532
7	N	23.687861	7	PC7	+4	379.0058
8	O	30.939247	1	PC1	+3	247.5140
9	F	39.157484	9	PC9	+3	313.2599
10	Ne	48.342573	5	PC5	+3	386.7406

n	element	f(n) (Hz)	odd part	class	j	f folded to [220,440)
11	Na	58.494514	11	PC11	+2	233.9781
12	Mg	69.613306	3	PC3	+2	278.4532
13	Al	81.698949	13	PC13	+2	326.7958
14	Si	94.751444	7	PC7	+2	379.0058
15	P	108.770790	15	PC15	+2	435.0832
16	S	123.756988	1	PC1	+1	247.5140
17	Cl	139.710037	17	PC17	+1	279.4201
18	Ar	156.629937	9	PC9	+1	313.2599
19	K	174.516690	19	PC19	+1	349.0334
20	Ca	193.370293	5	PC5	+1	386.7406

Step 5 — Count the vocabulary. `N = 20`, so `ceil(20/2) = 10` distinct pitch classes: `{1, 3, 5, 7, 9, 11, 13, 15, 17, 19}`. Ten notes from twenty elements. Confirmed against the class column: PC1 five times (H, He, Be, O, S), PC3 three (Li, C, Mg), PC5 three (B, Ne, Ca), PC7 twice (N, Si), PC9 twice (F, Ar), and PC11, PC13, PC15, PC17, PC19 once each — twenty in total.

Step 6 — Read the structure. Hydrogen, helium, beryllium, oxygen and sulphur are *the same note*. Not a metaphor: 1, 2, 4, 8 and 16 share the odd part 1, so their essence frequencies stand in ratios `1 : 4 : 16 : 64 : 256`, which is `0, 2, 4, 6, 8` octaves. **A tuning built on this set has one voice at PC1 sounding in five octaves and single voices at nine other classes — a strong fundamental with a thin upper structure, and that is a structural fact about the ordering, available before a single note is played.**

Step 7 — Verify and clause. Every raw frequency above was recomputed and matches the 2022 document to a worst relative deviation of 9.0e-15 across the full 118-element set (Section 7.4); nothing is transcribed. Per ES-9: *these are essence frequencies, not physical resonances; no claim is made that hydrogen, iron or any other element physically resonates at any value in this table.*

7.14.2 The same arithmetic on a non-chemical ordinal

The point of Section 7 is that atomic number is *an instance* of an essence index. A second instance comes from the corpus's own lineage: **the twelve declared nodes of the foundations volume**. Those nodes are the stages of the founding construction in which comfort is defined as the felt total of present costs, so that the present state is a local minimum of the comfort landscape by definition rather than by luck [Curzi, 2026]; their reference apparatus, including the three-drawer honesty ledger that assigns four items to established mathematics, seven to established psychology and exactly three to declared postulates, is Curzi [2026]. The node sequence is a designed process order — each node is a stage an operator passes through, and the order is the order of passage. That makes it a defensible ordinal under Section 7.9: intrinsic to the method, total, dense from 1, stable, and independently re-derivable by anyone holding the volume.

The four nodes this paper names elsewhere sit at indices 2 (GRADIENT), 7 (COLLAPSE), 9 (RECURSE) and 10 (FIXEDPOINT), with 11 (SEAL) — the remainder are referred to by index only, because inventing names for them would violate the corpus's naming constraint [Curzi, 2026].

`N = 12`, so `ceil(12/2) = 6` pitch classes: `{1, 3, 5, 7, 9, 11}`.

n	node	f(n) (Hz)	odd part	class	j	folded to [220,440)
1	node 01	0.483426	1	PC1	+9	247.5140
2	node 02 GRADIENT	1.933703	1	PC1	+7	247.5140

n	node	f(n) (Hz)	odd part	class	j	folded to [220,440)
3	node 03	4.350832	3	PC3	+6	278.4532
4	node 04	7.734812	1	PC1	+5	247.5140
5	node 05	12.085643	5	PC5	+5	386.7406
6	node 06	17.403326	3	PC3	+4	278.4532
7	node 07 COLLAPSE	23.687861	7	PC7	+4	379.0058
8	node 08	30.939247	1	PC1	+3	247.5140
9	node 09 RECURSE	39.157484	9	PC9	+3	313.2599
10	node 10 FIXEDPOINT	48.342573	5	PC5	+3	386.7406
11	node 11 SEAL	58.494514	11	PC11	+2	233.9781
12	node 12	69.613306	3	PC3	+2	278.4532

Sorted by pitch and reduced to a single octave, the six classes form this scale:

#	class	folded ratio	cents above PC1	step to next (cents)
1	PC1	1.000000000	0.00	203.91
2	PC3	1.125000000	203.91	203.91
3	PC9	1.265625000	407.82	329.83
4	PC7	1.531250000	737.65	34.98
5	PC5	1.562500000	772.63	330.01
6	PC11	1.890625000	1102.64	97.36

The six steps sum to `203.91 + 203.91 + 329.83 + 34.98 + 330.01 + 97.36 = 1200.00` cents. The scale closes at the octave, as it must.

Three structural observations, all of them consequences of Section 7.8 and none of them stylistic:

1. **The scale opens with two exact Pythagorean whole tones.** PC1 -> PC3 -> PC9 is `1 : 9/8 : 81/64`, the tone-chain generated by the odd prime 3. Nodes 03, 06 and 12 (class 3) and node 09 (class 9) are related to node 01 by pure three-limit intervals.
2. **PC7 and PC5 are 34.98 cents apart — a third of a semitone.** Node 07 COLLAPSE and nodes 05/10 land almost on top of each other. Sounded together they beat, rather than forming an interval; the companion paper \citep{ccp:CCP-018} treats that as an acoustic design decision, not a defect.
3. **Node 10 FIXEDPOINT (index 10) and node 05 (index 5) are the same class.** `10 = 2 * 5`, so they are exactly two octaves apart. Since node 10 is the steady state the recursion converges to [Curzi, 2026a] — and node 05 sits at the same class two octaves below, **the ordering places the fixed point in octave relation to an earlier node, which is precisely the relationship an accumulator has with its own transient.** This is offered as an observation about the ordering's arithmetic, not as evidence for the ordering. Had the nodes been numbered differently the relation would not hold, and the sensitivity test of Section 7.9 would report exactly that.

Sensitivity test (ES-4), run. Three plausible re-orderings:

- Swap 09 and 10: class 9 <-> class 5, pitch-class displacement `772.63 − 407.82 = 364.81` cents. **Load-bearing.**
- Swap 03 and 12: classes both 3, displacement `0.00` cents. Harmonically invisible.
- Swap 07 and 11: class 7 <-> class 11, displacement `1102.64 − 737.65 = 364.99` cents. **Load-bearing.**

Maximum displacement 364.99 cents, well over the 100-cent threshold of ES-4. **Conclusion: the node ordering is load-bearing for this essence set, and any published tuning derived from it must carry the ordinal justification alongside it.** In this case the justification is available — the node order is the order of the method — and so the set is conformant. Had it not been available, ES-4 would have required publishing a range instead of a frequency.

Closing clause (ES-9). *Every value in both parts of this worked example is an essence frequency. It is a pitch assigned to a position in an ordered category set. It is not a physical resonance and implies none. No element, no node, no operator and no room physically resonates at any of these values, and this paper makes no claim that any does.*

Section 8 takes up the question this section has deliberately left open: what happens when the two essences to be related do not share an ordering at all.

17.8 8 The essence bridge: the effective count as a cross-disciplinary index

Section 7 leaves a hole in the middle of the method, and does so deliberately.

The essence law needs an index. Within a discipline the index is available: the categories are ordered, the ordering is defensible, and `n` is a position in that ordering. Between two disciplines the index is *not* available, and no amount of care in constructing either ordering supplies it. Biology's third category and Ethics' third category are both "third", and their threeness has nothing in common. Feed both into `f(n) = k n^2` and you get the same frequency for two things that have no relationship at all. That is not a small inconvenience. **It means the essence law, as published, cannot say anything about a relationship that crosses a category boundary — which is to say it cannot say anything about the only relationships this corpus has spent twenty-eight chapters learning to engineer.**

This section closes the hole. It closes it not with an approximation, a heuristic, or a convention, but with an identity already present in the surplus characterisation [Curzi, 2026], and it closes it in one line:

```
n_eff = N_eff = 1 + (N − 1)u = e^F
```

The effective count IS the cross-disciplinary essence index. It is a real number rather than an integer, it is defined for any pair of states in the model, it equals 1 exactly when the two states are parallel — the ground essence — and it equals `N` exactly when they are orthogonal. Those are precisely the properties an index needs. Substituting it into the essence law produces the load bearing result of this entire volume, which is stated in Section 8.4 and which every subsequent section elaborates:

One nat of interaction surplus is exactly 2.885 octaves. The interaction surplus IS a pitch interval.

This is the paper's central identification. Sections 3 to 6 develop the arithmetic of interaction and Section 7 the arithmetic of essence; this section shows that they are the same arithmetic in two notations and gives the exact conversion factor between them.

17.8.1 8.1 Statement of the bridge

IDENTIFICATION (badged). Let two essences sit in different blocks of the framework's decomposition [Curzi, 2026]. Let `u = ||x ^ y||^2 = 1 − (x . y)^2` be the represented difference between their normalised state vectors, `N` the number of categories in the chosen partition, and `F = f(u) = ln(1 + (N−1)u)` the interaction surplus [Curzi, 2026]. Then the cross-disciplinary essence index is taken to be the effective count,

```
n_eff := N_eff = 1 + (N − 1)u = e^F
```

and the cross-disciplinary essence frequency follows from the essence law of Section 7 without adjustment:

```
f_cross = k ∗ n_eff^2 = k ∗ N_eff^2 = k ∗ e^(2F)

k = 0.483425732988414 Hz        (Section 7.1)
```

Taking the base-2 logarithm gives the interval above the ground essence `k`:

```
OCTAVES ABOVE BASE  =  log2( f_cross / k )  =  log2( e^(2F) )  =  2F / ln 2

2 / ln 2 = 2.885390081777927        (computed)

OCTAVES  =  2.885390081777927 ∗ F        with F in nats
```

Everything after the identification is exact algebra. The identification itself is a modeling decision and is badged as one. Its warrant is Section 8.3: `n_eff` is the unique quantity already present in the framework that satisfies every requirement an essence index must satisfy.

17.8.2 8.2 The honesty clauses, restated

The bridge inherits every clause from both parents. They are restated here in full because the corpus's disclosure standard requires it [Curzi, 2026] and because a construction joining two frameworks is exactly where an unstated clause would do the most damage.

From Section 7, the essence law:

- **H1.** The output is an essence frequency, **not a physical resonance.** No claim is made that any material, molecule, element, institution, room or person physically resonates at these values, and the text must never imply otherwise.
- **H2.** `k` is a chosen reference constant, not a measured one. Its form — a Pythagorean interval over the golden section — is a design decision. **It cancels from every relationship in the system (Result One) and sets only the reference pitch.**
- **H3.** The square law is a postulate about essence structure. Its formal agreement with infinite-square-well eigenvalue spacing is a mathematical coincidence, not evidence of mechanism.
- **H4.** The essence index requires a defensible ordering. Within a discipline that means the validity conditions [Curzi, 2026t] plus a declared ordinal rule. **Across disciplines the ordering problem does not go away — it becomes the requirement that `u` be honestly measured.**
- **H5 — the errata clause, on the strength of Section 7.4 and Section 7.10, binding on every chapter that quotes a number from the 2022 source.** Do not copy published values. Recompute them. `f = k n^2` reproduces all 118 published element values to a worst relative deviation of **9.0e-15**, and the compound rule `f = k^2 G^4` — `G` the count-weighted geometric mean of the constituent indices — matches **37 of the 44** published composites to better than one part in 10^12. **The seven that disagree are data-entry errors in the source, not a different law, and the O2 row proves it: for a single-element species `G` is the element index whatever the atom count, so no atom count whatever can produce the published figure.** Two transcription faults travel with the citation: the source's prose instruction `1+5^.5/2` is missing its parentheses and must be read as `(1 + sqrt(5))/2`, and the spreadsheet stored `phi` to fifteen significant figures, which accounts for the whole of the residual bias in the published table. **Cite the errata whenever the source is cited.** Where this section anchors to a published value — the `f(100)` check in Section 8.5 is the case — the anchor is to the recomputed value, and the agreement to fourteen figures is what licenses the check.

From the surplus framework:

- **The S3 axiom is a design axiom.** `g(u) = e^{f(u)}` affine is *chosen*, not derived. The entire bridge rests on it, because `N_eff` is defined by it. Say so whenever you lean on it — and this chapter leans on it in every section.
- **The triangle inequality fails for geometric triples** [Curzi, 2026]. Subadditivity holds on the parameter `u` only. Do not chain intervals across three essences as though pitch distances added.
- **As deployed, the surplus functional is a LOWER BOUND** on the ensemble's information, not the information itself. We log the floor, not the ceiling. Therefore **every frequency this section computes is a lower bound on the pitch, not the pitch.**
- **The characterisation theorem is pure mathematics and makes no empirical or metaphysical claim** [Curzi, 2026].
- Surplus is **not** identified with Shannon entropy or mutual information in general [Shannon, 1948, Cover and Thomas, 2006]; the corpus's own three-drawer classification puts that identification in the postulate drawer [Curzi, 2026].
- No thermodynamic negentropy law is proved [Curzi, 2026a]. Higher static surplus does not imply higher canonical concentration, and the Shannon/thermodynamic distinction is kept explicit [Brillouin, 1953, Landauer, 1961].

One clause specific to the bridge, and it is the sharpest of them: the essence frequency computed here is a function of a *measurement of a relationship*, not of either party to it. `f_cross` changes when the measured `u` changes, and `u` changes with the encoder, the category scheme, the observer, and the day. **A cross-disciplinary essence frequency is a reading of an instrument, and like any instrument reading it carries the instrument's error, not the world's truth.**

17.8.3 8.3 Why there is no shared ordinal, and what would fix it

The framework works in `V = R^(NM)`, decomposed into `N` orthogonal blocks `H_1, ..., H_N`, each of dimension `M` [Curzi, 2026]. A block is an idealised category in a chosen representation, and a unit vector `x` in `H_i` is a normalised state-description within category `H_i`.

Within one block the essence law is comfortable. The block's `M` sub-modes have an ordering — Paper B's twelve-house example orders them Origin, Structure, Flow, Perception, Reciprocity, Transformation, Harmony, Growth, Memory, Repair, Connection, Completion — and if that ordering satisfies the Section 7.9 test then `n` is well defined and `f(n) = k n^2` applies.

Across blocks it is not. Consider two essences, `x` in `H_3` (Biology & Body) and `y` in `H_12` (Ethics & Governance). Each has an ordinal position *within its own block*. Neither position means anything relative to the other. Concretely, four distinct failures block the naive approach:

1. **No common origin.** Each block's index 1 is an arbitrary starting point of that block's ordering. There is no fact of the matter about whether Biology's first sub-mode "precedes" Ethics' first sub-mode.

2. **Index collision.** Two unrelated essences at position 3 in their blocks receive identical frequencies. **A method that assigns the same pitch to two things precisely because they are unrelated has inverted its own purpose.**

3. **No relational content.** The essence law as published takes one argument. It maps a category to a frequency. It has no slot for "the relationship between two categories", which is the object we actually care about.

4. **Scale mismatch.** Blocks may have different `M`. An index of 7 in a twelve-mode block and an index of 7 in a four-mode block are not comparable even in a loose sense.

What we need is a quantity that: (a) takes *two* states as arguments; (b) is a single positive real number, so it can be squared; (c) equals 1 for identical or representationally equivalent states, so that a relationship of no difference lands on the ground essence `k * 1^2 = k`; (d) is bounded above by a stated maximum, so the pitch is bounded; (e) is defined regardless of which blocks the two states occupy; and (f) is monotone in difference, so more difference means a higher index and therefore a higher pitch.

There is exactly one such quantity already in the framework, and it is not a coincidence that there is exactly one — the properties above are, item for item, the axioms S1 through S4 read backwards.

17.8.4 8.4 The effective count is the index

The uniqueness characterisation [Curzi, 2026]: under axioms S1 (geometric dependence), S2 (`f(0) = 0`), S3 (affine effective count), and S4 (`f(1) = ln N`), the surplus functional is uniquely

```
F(x,y) = f(u) = ln(1 + (N−1)u),        u = 1 − (x . y)^2 in [0,1]
```

and the effective count is

```
N_eff = g(u) = e^{f(u)} = 1 + (N − 1)u
```

Check the six requirements against it, one by one:

requirement	`N_eff = 1 + (N−1)u`	source
(a) takes two states	yes, via `u(x,y)`	S1 [Curzi, 2026]
(b) single positive real	yes, `N_eff >= 1 > 0`	affine in `u >= 0`
(c) equals 1 at no difference	`u = 0 -> N_eff = 1` exactly	S2 [Curzi, 2026c]
(d) bounded above	`u = 1 -> N_eff = N` exactly	S4 [Curzi, 2026c]
(e) block-independent	`u` is defined for any pair of unit vectors	[Curzi, 2026]
(f) monotone in difference	strictly increasing in `u`	[Curzi, 2026]

Every box is ticked, and ticked by a theorem rather than by assertion. **The effective count is not merely *a* candidate for the cross-disciplinary essence index; it is the unique quantity the framework already contains that satisfies the requirements, and the requirements were derived from what an essence index must do, independently of the surplus framework.**

Its interpretation is also exactly right for the job. `N_eff` is a headcount *after that headcount has been audited for redundancy* [Curzi, 2026c].

Which effective count, and why it matters here. Two other effective counts appear in this corpus and neither is the one used in this section. `1/HHI`, the reciprocal Herfindahl index, is the Hill number of order two and `exp(H)` is the Hill number of order one; both are grounded in a mature literature [Hill, 1973, Jost, 2006, 2007, Simpson, 1949]. The corpus applies the first directly: a four-jurisdiction split of 40.9 / 27.3 / 18.2 / 13.6 gives a naive HHI of 0.293430 and 3.41 effective jurisdictions, which collapses to a bloc-HHI of 1.000000 and an effective count of exactly 1.00 once the bloc structure is audited [Curzi, 2026]; a four-state split of 40.0 / 25.0 / 20.0 / 15.0 gives a naive HHI of 0.285000 and 3.508772 effective jurisdictions, which likewise collapses to 1.00 and to a surplus term of exactly zero [Curzi, 2026]. **`N_eff = 1 + (N−1)u` is a third quantity and is not a Hill number.** It is a linear interpolation between 1 and `N` adopted as design axiom S3 [Curzi, 2026c], it has no external literature, and every octave value in this section inherits that choice. Where an effective count is quoted anywhere in this paper, which of the three is meant is stated. Ten parallel windows are one window: `u = 0`, `N_eff = 1`, `F = ln 1 = 0`. That is the ground essence, `f = k * 1 = 0.483426 Hz`. **Two essences that are**

representationally the same thing produce the ground tone and no interval at all — which is the musically correct answer, because there is no interval between a note and itself.

Substituting and simplifying:

```
f_cross  =  k * N_eff^2  =  k * (1 + (N−1)u)^2  =  k * e^(2F)
```

and, since `F = ln N_eff`,

```
log2(f_cross / k)  =  log2(N_eff^2)  =  2 * log2(N_eff)  =  2 * F / ln 2
```

Computed: `2 / ln 2 = 2.885390081777927`. One nat of interaction surplus is exactly 2.885390 octaves — equivalently 3462.4681 cents. And inversely, one octave of essence pitch is exactly `ln(2)/2 = 0.346574` nats of surplus.

Two verifications against values fixed independently elsewhere in the corpus — against the count-based essence index `f(n) = k n^2` with its twelve-state ladder [Curzi, 2026], and against the conversion-key result that **one bit equals two octaves, not one**, since doubling the effective count gives `f(2n)/f(n) = 4` under the square law, verified across `N_eff = 1, 3, 12` [Curzi, 2026l]:

```
F = 0        (parallel, u = 0)
  -> N_eff = 1
  -> f_cross = k * 1 = 0.483425732988414 Hz     <- the ground essence
  -> octaves above base = 2 * 0 / ln 2 = 0.000000

F = ln 12 = 2.4849066497880004     (orthogonal at N = 12, u = 1)
  -> N_eff = 12.000000
  -> f_cross = k * 144 = 69.61330555033166 Hz
  -> octaves above base = 2 * 2.4849066497880004 / 0.6931471805599453
                        = 7.169925001442313
  -> check: 2 * log2(12) = 2 * 3.584962500721156 = 7.169925001442313   [identical]
```

Both exact. **This is the unification: the ordinal index within a discipline and the effective count across disciplines are the same variable `n`. The two frameworks compose without adjustment, without a fitting constant, and without a correction term.**

It is worth pausing on how unusual that is. Nothing was tuned. `k` came from a 2022 document about chemistry; `N_eff` came from an axiomatic characterisation theorem about interaction geometry; the factor `2/ln 2` fell out of substituting one into the other. **The bridge has no free parameters at all** — which also means it has nothing to hide behind, and Section 8.11 states exactly how it fails.

17.8.5 8.5 The master conversion table

Everything an engineer needs from this section is in one table, computed here at full precision. Read it in either direction: from a measured surplus to a pitch, or from a chosen pitch to the surplus it implies.

F (nats)	`N_eff = e^F`	`f_cross = k e^(2F)` (Hz)	octaves above `k`	cents above `k`
0.000000	1.000000	0.483426	0.000000	0.0
0.050000	1.051271	0.534268	0.144270	173.1
0.100000	1.105171	0.590458	0.288539	346.2
0.200000	1.221403	0.721186	0.577078	692.5
0.300000	1.349859	0.880859	0.865617	1038.7

F (nats)	N_eff = e^F	f_cross = k e^(2F) (Hz)	octaves above k	cents above k
0.400000	1.491825	1.075884	1.154156	1385.0
0.500000	1.648721	1.314087	1.442695	1731.2
0.693147 (= ln 2)	2.000000	1.933703	2.000000	2400.0
0.750000	2.117000	2.166564	2.164043	2596.9
1.000000	2.718282	3.572060	**2.885390**	3462.5
1.250000	3.490343	5.889331	3.606738	4328.1
1.500000	4.481689	9.709865	4.328085	5193.7
1.750000	5.754603	16.008862	5.049433	6059.3
2.000000	7.389056	26.394151	5.770780	6924.9
2.197225 (= ln 9)	9.000000	39.157484	6.339850	7607.8
2.484907 (= ln 12)	12.000000	69.613306	7.169925	8603.9
2.750000	15.642632	118.290377	7.934823	9521.8
3.000000	20.085537	195.027860	8.656170	10387.4
3.218876 (= ln 25)	25.000000	302.141083	9.287712	11145.3
3.500000	33.115452	530.140688	10.098865	12118.6
4.000000	54.598150	1441.071800	11.541560	13849.9
4.605170 (= ln 100)	100.000000	4834.257330	13.287712	15945.3

Three checks that the table is internally consistent, each verifiable by eye:

- At `F = ln 2`, `N_eff = 2` and the pitch is exactly two octaves above base — the lemma of Section 7.7 (doubling the index is two octaves) reappearing as a statement about surplus.
- At `F = ln 9` the frequency is `39.157484 Hz`, which is `k * 81` — identical to `f(9)` in Chapter 29's within-discipline table. **The bridge is continuous with the ordinal law: when `N_eff` lands on an integer, the cross-disciplinary frequency *is* the within-discipline frequency of that index.**
- At `F = ln 100`, `f = 4834.257330 Hz`, which is `f(100)` from the periodic-table verification of Section 7.4, matching the source document's published value for fermium to fourteen figures.

That second check deserves emphasis. It is the reason this is a bridge and not an analogy. **Whenever the audited effective count happens to equal an integer `m`, the cross-disciplinary essence frequency is exactly the within-discipline essence frequency of index `m`. The two formalisms agree on their overlap.** A bridge that disagreed with either shore would be a different construction with the same name.

The integer-anchored version of the same table, which is the one to use when you are working backwards from a target effective count:

N_eff	F = ln N_eff	f = k N_eff^2 (Hz)	octaves above k
1	0.000000	0.483426	0.000000
2	0.693147	1.933703	2.000000
3	1.098612	4.350832	3.169925
4	1.386294	7.734812	4.000000
5	1.609438	12.085643	4.643856
6	1.791759	17.403326	5.169925
7	1.945910	23.687861	5.614710

N_eff	F = ln N_eff	f = k N_eff^2 (Hz)	octaves above k
8	2.079442	30.939247	6.000000
9	2.197225	39.157484	6.339850
10	2.302585	48.342573	6.643856
11	2.397895	58.494514	6.918863
12	2.484907	69.613306	7.169925
16	2.772589	123.756988	8.000000
20	2.995732	193.370293	8.643856
24	3.178054	278.453222	9.169925
32	3.465736	495.027951	10.000000
50	3.912023	1208.564332	11.287712
100	4.605170	4834.257330	13.287712

17.8.6 8.6 Reading the bridge backwards: pitch as a unit of surplus

The conversion runs both ways, and the reverse direction is the one that changes how you think about the numbers. Since `F = (cents/1200) * ln(2) / 2 = cents * ln(2) / 2400`:

interval	cents	F (nats)	N_eff = e^F	u required at N = 12
unison	0	0.000000	1.000000	0.000000
semitone	100	0.028881	1.029302	0.002664
whole tone	200	0.057762	1.059463	0.005406
minor 3rd	300	0.086643	1.090508	0.008228
major 3rd	400	0.115525	1.122462	0.011133
perfect 4th	500	0.144406	1.155353	0.014123
tritone	600	0.173287	1.189207	0.017201
perfect 5th	700	0.202168	1.224054	0.020369
minor 6th	800	0.231049	1.259921	0.023629
major 6th	900	0.259930	1.296840	0.026985
minor 7th	1000	0.288811	1.334840	0.030440
major 7th	1100	0.317692	1.373954	0.033996
octave	1200	0.346574	1.414214	0.037656
two octaves	2400	0.693147	2.000000	0.090909
three octaves	3600	1.039721	2.828427	0.166221

Two facts in that table are worth carrying.

A perfect fifth of surplus is 0.202168 nats. That is the whole content of the sentence "the interaction surplus IS a pitch interval": intervals are now a unit of measurement for organisational difference, and a familiar one. When a rotation programme raises `u` from 0.28 to 0.41 [Curzi, 2026d], that is restatable as *the group's cross-disciplinary tone rose by 1040 cents*, a shade under an octave. It is the same statement in a second register.

And the second fact, which is a warning: the entire audible range is about 3.4 nats. Ten octaves — 20 Hz to 20 kHz, roughly the human hearing range — is `10 / 2.885390 = 3.466` nats. `ln N` for `N = 12` is 2.485 nats, so **a twelve-category scheme's entire surplus range, from parallel to orthogonal, spans 7.17 octaves: about two-thirds of everything a human can hear.** For `N = 50` the span is `2 * log2(50) = 11.29` octaves — more than the audible range. Beyond about `N = 32` the essence pitch of an orthogonal pair leaves the audible band entirely and must be folded. Section 8.9 and the companion paper [Curzi, 2026l] handle the folding.

The corresponding ceiling table, since `F <= ln N` is hard [Curzi, 2026c]:

N	`ln N` (max surplus)	max octaves `= 2 log2 N`	max `f = k N^2` (Hz)
2	0.693147	2.000000	1.9337
3	1.098612	3.169925	4.3508
5	1.609438	4.643856	12.0856
7	1.945910	5.614710	23.6879
10	2.302585	6.643856	48.3426
12	2.484907	7.169925	69.6133
20	2.995732	8.643856	193.3703
50	3.912023	11.287712	1208.5643
100	4.605170	13.287712	4834.2573
118	4.770685	13.765286	6731.2199

`ln N` is the hard ceiling on surplus per interaction [Curzi, 2026c], and it is therefore a hard ceiling on pitch: `2 log2 N` octaves above the ground essence, and not one cent more. The table's last row is the periodic table's own ceiling and lands on 6731.2199 Hz — the source document's published value for oganesson, index 118, to fourteen figures. That agreement is not a new fact; it is Section 8.5's continuity property showing up at the top of the range as well as the middle.

The logarithmic-ceiling result [Curzi, 2026c] is audible here. Doubling `N` adds `ln 2 = 0.693147` nats, hence exactly two octaves, no matter where you start. **You cannot escape a headroom problem by adding domains, and now you cannot escape a pitch problem that way either: going from twelve categories to twenty-four buys you two octaves and nothing else.**

17.8.7 8.7 Crossing and divergence, transported into pitch

The crossing/divergence decomposition [Curzi, 2026d] splits `u` into two sources, and the split survives translation into pitch with its bounds intact.

For a pure block vector `x` in `H_i` and any unit `y`, write `y = y_par + y_perp` with `y_par` in `H_i` and `y_perp` orthogonal to `H_i`. Let `alpha = ||y_par||`, `beta = ||y_perp||`, and `gamma = angle(x, y_par)`. Then

```
u = beta^2 + alpha^2 sin^2(gamma)  =  u_cross + u_div

u_cross = beta^2                  BETWEEN categories  (cross-disciplinary difference)
u_div   = alpha^2 sin^2(gamma)    WITHIN a category    (internal dissent)

max( f(u_cross), f(u_div) )  <=  f(u)  <=  f(u_cross) + f(u_div)
```

Multiply the bound through by `2/ln 2 = 2.885390` — a positive constant, so the inequalities are preserved exactly:

```
max( OCT(u_cross), OCT(u_div) )  <=  OCT(u)  <=  OCT(u_cross) + OCT(u_div)

where OCT(v) := 2.885390081777927 * ln(1 + (N-1)v)
```

In musical terms: the cross-disciplinary tone of a relationship is never lower than the higher of its two component tones, and never higher than the two component tones stacked. The lower bound says you always get at least the better of your two sources of difference. The upper bound says crossing and divergence never combine to give you more than they would

separately — **there is no synergy beyond additivity, and the bound is the framework refusing to promise any.**

Computed, for `N = 12`, showing the whole structure:

u_cross	u_div	u	f(u_cross)	f(u_div)	f(u)	max bound	sum bound	f_cross (Hz)	octaves
0.00	0.00	0.00	0.000000	0.000000	0.000000	0.000000	0.000000	0.483426	0.0000
0.20	0.00	0.20	1.163151	0.000000	1.163151	1.163151	1.163151	4.950280	3.3561
0.00	0.20	0.20	0.000000	1.163151	1.163151	1.163151	1.163151	4.950280	3.3561
0.10	0.10	0.20	0.741937	0.741937	1.163151	0.741937	1.483875	4.950280	3.3561
0.40	0.10	0.50	1.686399	0.741937	1.871802	1.686399	2.428336	20.424737	5.4009
0.10	0.40	0.50	0.741937	1.686399	1.871802	1.686399	2.428336	20.424737	5.4009
0.25	0.25	0.50	1.321756	1.321756	1.871802	1.321756	2.643512	20.424737	5.4009
0.45	0.45	0.90	1.783391	1.783391	2.388763	1.783391	3.566782	57.435811	6.8925
0.90	0.00	0.90	2.388763	0.000000	2.388763	2.388763	2.388763	57.435811	6.8925
0.50	0.50	1.00	1.871802	1.871802	2.484907	1.871802	3.743604	69.613306	7.1699
0.80	0.20	1.00	2.282382	1.163151	2.484907	2.282382	3.445533	69.613306	7.1699

And the same in octaves, with the slack — the gap between the achieved pitch and the sum bound — made explicit:

u_cross	u_div	achieved OCT(u)	max bound (oct)	sum bound (oct)	slack (oct)
0.20	0.00	3.3561	3.3561	3.3561	0.0000
0.00	0.20	3.3561	3.3561	3.3561	0.0000
0.10	0.10	3.3561	2.1408	4.2816	0.9254
0.40	0.10	5.4009	4.8659	7.0067	1.6058
0.25	0.25	5.4009	3.8138	7.6276	2.2267
0.45	0.45	6.8925	5.1458	10.2916	3.3990
0.50	0.50	7.1699	5.4009	10.8018	3.6318
0.90	0.00	6.8925	6.8925	6.8925	0.0000

The corpus reports the same decomposition on a real-shaped case: an institute at `u_cross = 0.72`, `u_div = 0.03` and `u = 0.75` with `N = 6` returns a balance index `B = 0.0206`, a ceiling utilisation of 0.8696, and zero within-block dissents across forty logged decisions [Curzi, 2026d]. **That is the extreme of the pattern the slack column below measures: almost all of the difference is crossing, almost none is divergence, and the slack is therefore near zero while the fragility is at its maximum.**

Read the slack column and the structure of the bound becomes visible. **Slack is zero exactly when one of the two sources is zero, and it is maximal when the two are balanced.** At `u_cross = u_div = 0.5` the sum bound promises 10.80 octaves and the system delivers 7.17 — a slack of 3.63 octaves, more than a third of the audible range, thrown away by concavity.

That is not a defect. **It is strict concavity [Curzi, 2026] charging its price, and the price is real: two half-sized sources of difference produce a lower tone than one full-sized one.** Compare rows: `u_cross = 0.2, u_div = 0` and `u_cross = 0.1, u_div = 0.1` both give `u = 0.2` and therefore both give 3.3561 octaves — identical pitch, because `f` depends only on the sum. But the *bound* differs: the split case had 4.28 octaves available under the sum bound and reached 3.36. The concentrated case had 3.36 available and reached 3.36, saturating its bound.

The design lesson [Curzi, 2026d] is now audible: **crossing and divergence are not substitutes, and a system rich in one and starved of the other sits at exactly the same pitch as one with the same total `u` distributed differently — but it has a very different fragility.** If your entire `u` is `u_cross` (cross-domain hires, one permitted opinion inside each domain) then a single reorganisation that merges two domains destroys a large fraction of your surplus in one move. If your `u` is split, no single structural change takes it all. **The pitch does not distinguish these cases; the decomposition does. Report both.**

17.8.8 8.8 What a two-essence interval sounds like as `u` sweeps 0 to 1

The most useful intuition in this section comes from watching one number move. Fix `N = 12`. Sweep `u` from 0 (representationally identical) to 1 (orthogonal) and watch what happens to the tone relative to the ground essence `k`.

`u`	`N_eff = 1 + 11u`	`F` (nats)	`f_cross` (Hz)	octaves above `k`	folded into [220,440) Hz
0.00	1.0000	0.000000	0.483426	0.000000	247.5140 (j = +9)
0.05	1.5500	0.438255	1.161430	1.264536	297.3262 (j = +8)
0.0909 (= 1/11)	2.0000	0.693147	1.933703	2.000000	247.5140 (j = +7)
0.10	2.1000	0.741937	2.131907	2.140779	272.8842 (j = +7)
0.15	2.6500	0.974560	3.394857	2.811985	434.5417 (j = +7)
0.20	3.2000	1.163151	4.950280	3.356144	316.8179 (j = +6)
0.25	3.7500	1.321756	6.798174	3.813781	435.0832 (j = +6)
0.30	4.3000	1.458615	8.938542	4.208673	286.0333 (j = +5)
0.40	5.4000	1.686399	14.096694	4.865919	225.5471 (j = +4)
0.50	6.5000	1.871802	20.424737	5.400879	326.7958 (j = +4)
0.60	7.6000	2.028148	27.922670	5.851999	223.3814 (j = +3)
0.70	8.7000	2.163323	36.590494	6.242031	292.7239 (j = +3)
0.75	9.2500	2.224624	41.363114	6.418907	330.9049 (j = +3)
0.80	9.8000	2.282382	46.428207	6.585563	371.4257 (j = +3)
0.90	10.9000	2.388763	57.435811	6.892512	229.7432 (j = +2)
1.00	12.0000	2.484907	69.613306	7.169925	278.4532 (j = +2)

And the interval itself — the ratio `N_eff^2` between the pair's tone and the ground essence, expressed as a musician would hear it:

`u`	ratio `N_eff^2`	cents	folded cents	nearest 12-TET
0.00	1.000000	0.00	0.00	unison +0.0
0.02	1.488400	688.51	688.51	perfect 5th -11.5
0.05	2.402500	1517.44	317.44	minor 3rd +17.4
0.10	4.410000	2568.93	168.93	major 2nd -31.1
0.15	7.022500	3374.38	974.38	minor 7th -25.6
0.20	10.240000	4027.37	427.37	major 3rd +27.4

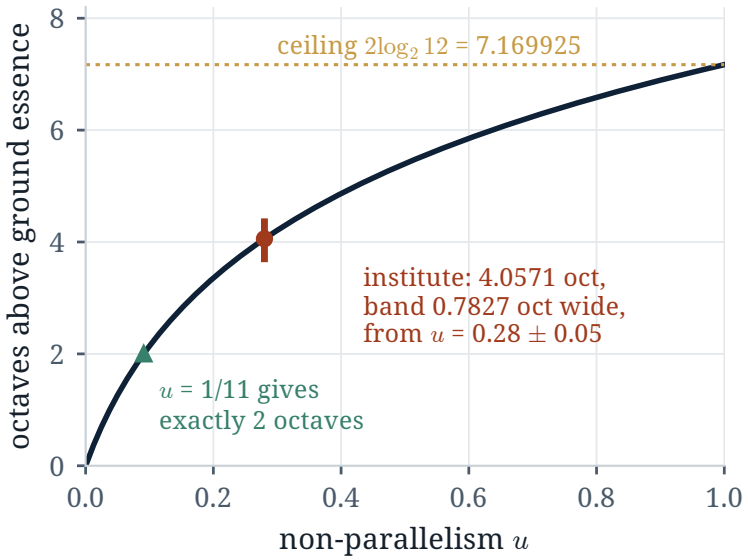


Figure 17.2: **The bridge is violently front-loaded, and its honest output is a band rather than a value.** At $N = 12$ the first tenth of u buys 2.14 octaves and the last tenth 0.28, against a hard ceiling of $2\log_2 12 = 7.169925$ octaves reached only at exact orthogonality; $u = 1/11$ lands on exactly two octaves. The worked institute at $u = 0.28 \pm 0.05$ spans 3.6393 to 4.4222 octaves, a band 0.7827 octaves wide. Modelled from the parameters of Sections 8.8 and 8.14; the identification $n_{\text{eff}} := N_{\text{eff}}$ is a badged modelling decision, not a derivation. Generated by `figs/src/CCP-017_fig05.py`.

u	ratio N_eff^2	cents	folded cents	nearest 12-TET
0.25	14.062500	4576.54	976.54	minor 7th -23.5
0.30	18.490000	5050.41	250.41	minor 3rd -49.6
0.40	29.160000	5839.10	1039.10	minor 7th +39.1
0.50	42.250000	6481.06	481.06	perfect 4th -18.9
0.60	57.760000	7022.40	1022.40	minor 7th +22.4
0.70	75.690000	7490.44	290.44	minor 3rd -9.6
0.80	96.040000	7902.68	702.68	perfect 5th +2.7
0.90	118.810000	8271.01	1071.01	major 7th -29.0
1.00	144.000000	8603.91	203.91	major 2nd +3.9

Four things to take from the sweep, each an engineering fact rather than an aesthetic one.

(1) The rise is violently front-loaded. Concavity, made audible. Octaves bought per 0.1 of u:

u range	octaves gained
0.0 -> 0.1	2.1408
0.1 -> 0.2	1.2154
0.2 -> 0.3	0.8525
0.3 -> 0.4	0.6572
0.4 -> 0.5	0.5350
0.5 -> 0.6	0.4511
0.6 -> 0.7	0.3900
0.7 -> 0.8	0.3435

u range	octaves gained
0.8 -> 0.9	0.3069
0.9 -> 1.0	0.2774

The first tenth of u buys 2.14 octaves; the last tenth buys 0.28 — a factor of 7.7. **The diminishing-returns clause [Curzi, 2026c], heard: the second genuinely different member is the largest single gain you will ever buy, and in pitch terms it is more than two octaves' worth.** A team at `u = 0.05` that reaches `u = 0.15` gains 1.55 octaves. The same team going from `u = 0.85` to `u = 0.95` gains 0.29. The cheap improvements are at the bottom.

(2) The sweep is not monotone once folded. In the raw column, pitch rises monotonically with u — guaranteed by strict monotonicity [Curzi, 2026]. In the folded column it does not: `u = 0.15` folds to 434.54 Hz and `u = 0.20` folds to 316.82 Hz, which is *lower*. Folding wraps. **This is the single most common error in reading these tables and it must be stated every time: a folded frequency is a pitch class, not a pitch, and pitch classes do not order.** Comparison of two systems by surplus is made on `F`, on raw `f_cross`, or on octaves-above-base. Folded values are never compared.

(3) At `u = 1/11 = 0.0909...` the effective count is exactly 2 and the interval is exactly two octaves. This is the smallest u at `N = 12` that "buys a second full member", and it is a useful landmark: below `u = 0.091`, twelve categories are worth fewer than two.

(4) The extremes. At `u = 0` the two essences are the same essence and the interval is a unison — the pair sounds as one note at the ground essence `0.483426 Hz`. At `u = 1` the pair is orthogonal, `N_eff = 12`, the interval is `144:1` — seven octaves and a Pythagorean whole tone, 8603.91 cents. **Perfect orthogonality at twelve categories is a major second, seven octaves up. That is the whole dynamic range of the instrument.**

17.8.9 8.9 Sensitivity: how precisely must u be measured?

Because the bridge has no free parameters, all of its uncertainty lives in u. The sharp Lipschitz bound [Curzi, 2026] gives the constant `N − 1`, so

```
dF/du = (N−1) / (1 + (N−1)u)        maximum N−1, attained at u = 0

d(octaves)/du = 2.885390 ∗ dF/du
```

N	max dF/du (at u=0)	max d(oct)/du	dF/du at u=1	d(oct)/du at u=1
2	1.0000	2.8854	0.5000	1.4427
5	4.0000	11.5416	0.8000	2.3083
12	11.0000	31.7393	0.9167	2.6449
20	19.0000	54.8224	0.9500	2.7411
50	49.0000	141.3841	0.9800	2.8277

At `N = 12` and u near zero the pitch moves 31.74 octaves per unit of u. In practical terms, with a just-noticeable pitch difference of about 5 cents:

u	dF/du	d(cents)/du	du producing a 5-cent shift
0.00	11.0000	38087.1	0.000131
0.10	5.2381	18136.7	0.000276
0.25	2.9333	10156.6	0.000492

u	dF/du	d(cents)/du	du producing a 5-cent shift
0.50	1.6923	5859.6	0.000853
0.75	1.1892	4117.5	0.001214
1.00	0.9167	3173.9	0.001575

A change in u of 0.00013 — one part in seven thousand — is audible as a pitch shift near u = 0. No organisational measurement of u is anywhere near that precise. The estimation methods [Curzi, 2026i] — disagreement rate, error decorrelation, prediction diversity — yield u to perhaps two decimal places at best, and the second of those is soft.

This has one hard consequence and it is a specification, not a caution:

Publish u with an interval, and publish the pitch as the interval's image. If your survey supports `u = 0.28 +/- 0.05`, then at `N = 12` the tone is not 8.0473 Hz; it is the range from `k*(1 + 11*0.23)^2 = 6.0239 Hz` to `k*(1 + 11*0.33)^2 = 10.3631 Hz`, which is 0.7827 octaves wide — a minor sixth of uncertainty. **Quote the sixth. A single-valued essence frequency from a two-significant-figure u is precision theatre (FM-E6) wearing a new hat.**

The second half of the sensitivity question is temporal, and it has an exact answer. The table above asks how small a difference in u is audible. It does not ask how *long* the tone must sound before the difference can be heard at all, and that question is not a matter of taste: the Gabor limit fixes it. With the time-bandwidth product `dt * df >= 1/(4 pi)` and the bridge's own relation `d(ln f) = 2 dF`, Curzi [2026z] derives the bound

```
cycles required  >=  1 / (8 pi dF)
```

Resolving a surplus difference dF requires at least `1/(8 pi dF)` cycles of the rendered tone, and the bound is independent of the transposition — because it is a bound on a ratio, and octave transposition multiplies numerator and denominator alike. **This is the exact Gabor bridge to the surplus functional: it prices, in time, what it costs to hear a difference in surplus.**

dF (nats)	cycles required	duration at 69.613 Hz (N_eff = 12)	at 302.141 Hz (N_eff = 25)
0.20	0.199	2.86 ms	0.66 ms
0.10	0.398	5.72 ms	1.32 ms
0.05	0.796	11.43 ms	2.63 ms
0.02	1.989	28.58 ms	6.58 ms
0.01	3.979	57.16 ms	13.17 ms
0.005	7.958	114.31 ms	26.34 ms
0.001	39.789	571.57 ms	131.69 ms

Read the two sensitivity results together and they close on each other from opposite sides. The Lipschitz table says a du of 0.00013 is audible in principle near `u = 0`; the Gabor bound says that hearing it requires the tone to persist. **Neither one alone is a specification. Together they say: the pitch difference your measurement can honestly support is coarse, and coarse differences are cheap in time — which is the fortunate direction for the error to run.** A rendering that sounds each essence for well under a cycle of the resolution it claims is asserting a distinction that physics has not delivered, whatever the frequency table says. Curzi [2026z] develops the full set of design laws that follow from the limit, and Curzi [2026l] applies them to built work.

17.8.10 8.10 Composing with the storage law

The bridge produces a pitch from u. The headroom gauge [Curzi, 2026b] produces a budget from Q, delta, eta, C and N. They meet in `ln N`, and it is worth being explicit about how, because

a reader will otherwise try to make the pitch do work it cannot do.

```
h_t   =  eta * ln N - delta * Q_t - C_t              the headroom gauge
OCT   =  2.885390 * F,       F = ln N_eff <= ln N     (this section)
```

Both contain `ln`-of-a-count, and both are ceilinged by `ln N`. But they are different quantities and must never be added:

- `F` is the surplus of **one interaction**, and `OCT` is its pitch. It is an instantaneous, geometric quantity.
- `h_t` is a **budget over a period**, in units of stored order. It depends on the history `Q_t`.

The correct composition is the one the storage law already specifies [Curzi, 2026a]: `S_t = F(U_t)` is the surplus *input* to the storage recursion. So the pitch is a gauge on the **input**, not on the **stock**:

```
OCT_t = 2.885390 * S_t      is the pitch of this period's interaction

Q_{t+1} = (1 - delta) Q_t + eta S_t - C_t
```

The tone reports what is going into the accumulator this period. It reports nothing about how much is already in it. A group can be sounding a high tone — rich, genuinely non-parallel interaction — while `Q` collapses, because `eta` is near zero and none of it is sticking (the leaky-vessel fault, Curzi 2026b), or because `C` exceeds `eta ln N - delta Q`, so that the negative-headroom corollary has already sealed the outcome [Curzi, 2026b]. Curzi [2026l] uses the pitch as a design target for rooms and cadences; it is not a health metric and must not be read as one.

There is a second, sharper composition available, and it is developed elsewhere rather than here: the storage law's fixed point and the essence bridge can be set equal, giving a minimum sustaining frequency `f_min = k e^(2c/eta)` below which `Q_inf` is negative [Curzi, 2026], and a worked dyad at `eta = 0.30`, `c = 0.62`, `delta = 0.09` sounds at `f = 1.4908` Hz against an `f_min` of `30.1588` Hz — an octave deficit of `+4.338420`, or 5206.1 cents, with `Q_inf = −5.0119` [Curzi, 2026]. **The sign of that verdict is robust where the number is not**, and the same caution applies to every figure in this section.

One legitimate joint reading, and it is useful: since `S_t <= ln N` and `OCT_t = 2.885390 S_t`, the maximum pitch a system can sound is fixed by its category count, and the maximum sustainable *increment* to stored order is `eta ln N - delta Q_t - C_t`. **Both ceilings come from the same `ln N`, which means a system that has run out of pitch has also run out of headroom for growth from geometry — and the remaining levers are `eta`, `C`, and shedding `Q` [Curzi, 2026b], none of which change the tone at all.** That is a genuinely useful diagnostic: *if the tone is already at ceiling and headroom is still negative, no amount of further diversity work will help, and continuing to do diversity work is the mistake.*

17.8.11 8.11 What the bridge does not claim

Stated flatly, because a keystone that is over-claimed takes the arch down with it.

1. **It does not claim that essences resonate.** H1. `f_cross` is an essence frequency of a *relationship*, computed from a *measurement* of that relationship. Nothing physical resonates.
2. **It does not claim that `2/ln 2` is a natural constant.** It is `2/ln 2` because the essence law is square and octaves are base two. Had the essence law been cubic, the factor would be `3/ln 2 = 4.328`. **The factor is a consequence of two design choices, and its exactness is the exactness of arithmetic, not of nature.**

3. **It does not claim to have derived `N_eff` as the index.** The identification is badged. What is *shown* is that `N_eff` uniquely satisfies six requirements (Section 8.4), that the substitution is exact, and that the result is continuous with the within-discipline law at every integer (Section 8.5). That is a strong warrant. It is not a derivation.
4. **It does not claim more information than `F` contains.** As deployed the surplus functional is a **lower bound** on the ensemble's information. Therefore `f_cross` is a **lower bound on the pitch**. Where an error is unavoidable, the safe direction is a tone reported too low.
5. **It does not permit chaining.** The triangle inequality fails for geometric triples [Curzi, 2026]. Given essences A, B, C, you may compute the A-B tone and the B-C tone, but **you may not add the two intervals to obtain the A-C tone.** Compute A-C directly from `u(A,C)` or do not compute it. This is the failure mode that will catch the most people, because pitch intervals *do* add in ordinary music and the habit is deeply worn.
6. **It does not survive an undisclosed change of `N`.** `N_eff = 1 + (N−1)u` contains `N` explicitly. Recounting the categories moves every published frequency. The framework is explicit that `N = 12` is not asserted as correct and that other decompositions produce different maxima under the same structural theorems [Curzi, 2026t]. **`N` must be published with every frequency, as `k` and `j` already are.**

17.8.12 8.12 Specification: twelve requirements on a published bridge value

BR-1 — Publish `(N, u, F, N_eff, f_cross, j, band)` as an inseparable tuple. A cross-disciplinary essence frequency quoted without its `N` and its `u` is unreproducible and therefore not a result.

BR-2 — Badge the identification. State, at each first use in a document: *the effective count is taken as the cross-disciplinary essence index; this is a modeling identification, warranted by the six-requirement argument, not derived.*

BR-3 — Badge S3. The affine effective-count axiom is a design axiom, chosen not derived. `N_eff` exists because S3 was adopted. The bridge rests on S3 entirely; say so.

BR-4 — Publish `u` as an interval and the pitch as its image. Never publish a single-valued `f_cross` from a `u` estimated to two decimal places. Report the pitch range in octaves or cents. (Section 8.9)

BR-5 — Never compare folded frequencies to infer surplus ordering. Folding wraps and destroys order. Compare `F`, raw `f_cross`, or octaves-above-base. (Section 8.8, point 2)

BR-6 — Never add two intervals to obtain a third. The triangle inequality fails for geometric triples. Compute every pair directly from its own `u`. (Section 8.11, point 5)

BR-7 — Report the decomposition, not only the total. Publish `u_cross` and `u_div` separately with the bound `max(OCT(u_cross), OCT(u_div)) <= OCT(u) <= OCT(u_cross) + OCT(u_div)` evaluated. Two systems at the same pitch with different decompositions have different fragilities. (Section 8.7)

BR-8 — State the ceiling with the value. Every published cross-disciplinary tone carries `2 log2 N` octaves as its stated maximum. A reader must be able to see how much of the range is in use.

BR-9 — Do not read the tone as a health metric. `OCT_t = 2.885390 S_t` gauges the *input* to the storage recursion, not the stock `Q_t`, and not `h_t`. (Section 8.10)

BR-10 — Restate H1 and the lower-bound clause on any page carrying a frequency. "Essence frequency, not physical resonance" and "the surplus functional as deployed is a lower bound on the ensemble's information, so this tone is a lower bound on the pitch."

BR-11 — Re-derive when `N` changes. Any change to the category partition invalidates every published frequency in the set. Republish or withdraw; do not let a stale tone circulate.

BR-12 — Verify the integer-anchor property on any new implementation. For a test value of `u` making `N_eff` an integer `m`, confirm your code returns exactly `k m^2` — the same number the within-discipline table of Section 7.7 gives for index `m`. This is a two-line regression test that catches almost every implementation error.

17.8.13 8.13 Failure modes of the bridge

FM-B1 THE FALSE ORDINAL ACROSS BLOCKS

Damages: every cross-disciplinary frequency computed. *Description.* Two essences in different blocks are given ordinals within their own blocks, and those ordinals are fed into `f(n) = k n^2` as though comparable. Biology's third sub-mode and Ethics' third sub-mode receive the same frequency. *Signature.* Two unrelated essences with identical published tones, and the practitioner reporting this as a discovered affinity. *Fix.* The whole of this section. Use `n_eff = 1 + (N−1)u`. Within-block ordinals may be used within a block and nowhere else.

FM-B2 THE CHAINED INTERVAL

Damages: any three-party analysis. *Description.* The A-B tone is a perfect fifth up, the B-C tone is a fourth up, therefore A-C is an octave up. **It is not.** The triangle inequality fails for geometric triples [Curzi, 2026]. *Signature.* An interval reported for a pair whose `u` was never measured. *Fix.* BR-6. Measure `u(A,C)` or report nothing for that pair.

FM-B3 THE UNAUDITED ROSTER

Damages: `u`, and therefore everything. *Description.* `N_eff` is computed from a headcount rather than from a measured `u`. Twelve people in a room are reported as `N_eff = 12`, tone at 69.61 Hz, when the measured `u` is 0.15 and the true `N_eff` is 2.65. *Signature.* `N_eff` equal or nearly equal to the literal headcount, with no disagreement audit in the file. *Fix.* The roster audit [Curzi, 2026c] and the estimation protocol [Curzi, 2026i]. `N_eff` is a headcount **after audit for redundancy**. Measure `u`. An unmeasured `u` is not an optimistic estimate; it is no estimate.

FM-B4 THE FOLDED COMPARISON

Damages: every ordering claim between systems. *Description.* System A folds to 434.54 Hz, system B to 316.82 Hz, therefore A has more surplus. It does not: those are `u = 0.15` and `u = 0.20` respectively, and B has more. *Signature.* A comparison between two folded values. *Fix.* BR-5. Compare in `F` or in octaves-above-base, always.

FM-B5 PRECISION FROM A SOFT `u`

Damages: credibility, exactly as FM-E6 does. *Description.* `u = 0.28` from a small disagreement audit becomes `f_cross = 8.0473 Hz`, published to four decimals. The underlying interval is 0.638 octaves wide. *Signature.* More precision in the output than in the input. *Fix.* BR-4. Publish the range. Section 8.9 gives the arithmetic.

FM-B6 THE SILENT `N` CHANGE

Damages: comparability across time. *Description.* An organisation reorganises from twelve departments to nine. Nobody re-derives. Frequencies computed under `N = 12` are compared with new observations under `N = 9`, and the apparent change is entirely an artefact of the

denominator. *Signature.* A step change in the reported tone that coincides with a reorganisation date. *Fix.* BR-11. `N` travels with every value. On any change to the partition, republish or withdraw.

FM-B7 THE SYNERGY CLAIM

Damages: the subadditive bound, and any budget built on it. *Description.* "Combining cross-domain hiring with an internal red team will more than double our surplus." The bound says it will not: `f(u) <= f(u_cross) + f(u_div)`, always. *Signature.* A projected gain exceeding the sum of the parts' separate gains. *Fix.* BR-7. Evaluate the sum bound and cap the projection at it. Note that in practice you will land well *below* the sum bound — Section 8.7's slack column shows 3.6 octaves of concavity tax at balanced sources.

FM-B8 THE INVERTED READING

Damages: the entire method, by turning it into decoration. *Description.* A frequency is chosen first — because it sounds pleasant, or because it is 432 Hz, or because a client asked for it — and `u` is then back-solved to justify it. *Signature.* A `u` value with more than two significant figures that was never measured, or a `u` that lands suspiciously close to a round musical interval. *Fix.* The preregistration discipline [Curzi, 2026i, Nosek et al., 2018], applied without exception: **fix the measurement protocol and the ordinal rule before computing any frequency, and freeze them.** The tone is an output. If you want a particular tone, say so openly and choose it as a design decision — that is legitimate, and Curzi [2026l] does exactly that when it selects a band — but do not dress a chosen tone as a measured one.

FM-B9 THE TONE AS SCOREBOARD

Damages: `Q`, via misdirected effort. *Description.* The cross-disciplinary tone is adopted as an organisational KPI. It rises. `Q` falls, because `eta` is near zero and `C` is above the sustainability ceiling. *Signature.* A rising tone alongside falling stored order; or, worse, a rising tone and no measurement of stored order at all. *Fix.* BR-9, and Goodhart. **The tone gauges the input to the accumulator. Measure `h_t` separately and let it, not the tone, decide whether you may build.**

FM-B10 THE UNBADGED IDENTIFICATION

Damages: the reader's ability to audit the argument. *Description.* `n_eff = N_eff` is presented as a theorem. It is not; it is a modeling identification with a strong six-part warrant. *Signature.* The sentence "it follows that the effective count is the essence index", with no badge. *Fix.* BR-2 and BR-3. And be precise about what *is* proved: the substitution's consequences, the factor `2/ln 2`, the integer-anchor continuity, and the transported Theorem 4.1 bounds are all exact. The identification is the postulate; everything downstream of it is arithmetic.

17.8.14 8.14 Worked instance: a twelve-house research institute

An interdisciplinary research institute is organised into twelve declared houses, with `N = 12` and `M = 12` sub-modes per house. The framework is explicit that twelve is not asserted as correct and that other decompositions produce different maxima with the same structural theorems; that caveat travels with every number below.

Step 0 — Validity check (BR-1, ES-5). The validity conditions [Curzi, 2026t]:

- *State encoding.* Each researcher's working position is encoded as a normalised vector over the 144-dimensional space (12 houses x 12 sub-modes), from a structured self-description plus a publication-topic coding. Declared.
- *Category assignment.* The twelve houses. Declared. `N = 12`.
- *Geometric dependence.* Interaction is modelled through `u = 1 − (x.y)^2`, an inner-product-derived quantity. Satisfied.
- *Zero baseline.* Two researchers with representationally equivalent positions give `u = 0`. Satisfied.
- *Affine effective count.* Adopted as a modeling assumption. **Badged (BR-3).**
- *Normalisation.* Maximum surplus `ln 12 = 2.484907` nats. Declared.

All six stated. The domain is in scope.

Step 1 — Measure `u`. The estimation protocol [Curzi, 2026i], run over four months: independent-estimate- before-discussion on 34 forecast questions, error decorrelation computed pairwise, disagreement rate logged from the decision log.

```
measured u (institute-wide mean over pairs) = 0.28,  supported to +/- 0.05
```

Step 2 — Compute the baseline tone.

```
N_eff = 1 + 11 * 0.28 = 4.0800
F     = ln 4.0800     = 1.406097 nats
OCT   = 2.885390081777927 * 1.406097 = 4.0571 octaves above k
f_cross = k * 4.08^2  = 0.483425732988414 * 16.6464 = 8.0473 Hz
folded into [220, 440): 8.0473 * 2^5 = 257.5135 Hz   (j = +5)
                        nearest note C4 (261.626 Hz), -27 cents
```

Step 3 — Publish the uncertainty as an interval (BR-4). With `u = 0.28 +/- 0.05`:

```
u = 0.23 -> N_eff = 3.5300 -> f_cross =  6.0239 Hz -> 3.6393 octaves above k
u = 0.33 -> N_eff = 4.6300 -> f_cross = 10.3631 Hz -> 4.4220 octaves above k

published tone: 8.05 Hz, range 6.02 - 10.36 Hz, i.e. 4.06 octaves above base
                with an uncertainty of 0.7827 octaves (a minor sixth)
```

That 0.78-octave band is the honest result. The four-decimal figure in Step 2 is an intermediate, not a finding.

Step 4 — Decompose (BR-7). The disagreement audit separates between-house disagreement from within-house disagreement:

```
u_cross = 0.20   (between houses)
u_div   = 0.08   (within houses)
u       = 0.28   [checks: 0.20 + 0.08 = 0.28, decomposition identity]

f(u_cross) = ln(1 + 11*0.20) = 1.163151 nats -> 3.3561 octaves
f(u_div)   = ln(1 + 11*0.08) = 0.631272 nats -> 1.8215 octaves
f(u)       = ln(1 + 11*0.28) = 1.406097 nats -> 4.0571 octaves

bound: max(3.3561, 1.8215) = 3.3561  <=  4.0571  <=  3.3561 + 1.8215 = 5.1776
slack to the sum bound: 5.1776 - 4.0571 = 1.1205 octaves
```

Diagnosis: the institute is 71% crossing and 29% divergence. It has bought cross-domain difference (that is what an interdisciplinary institute is for) and is thin on internal dissent — the crossing-heavy pattern [Curzi, 2026d]. The fragility follows: a merger of two houses would remove a substantial fraction of `u_cross`, and there is little `u_div` underneath to catch the fall.

Step 5 — Design the intervention [Curzi, 2026j]. Two moves, sequenced:

- *Rotation across houses* (raises `u_cross`): six-month placements, four researchers per cycle.
- *Standing red team* (raises `u_div`): rotating adversarial role inside each house, with a minority report requirement on every major decision.

Projected post-intervention `u = 0.41`, from the pilot's measured effect on eight pairs.

Step 6 — Compute the interval the intervention buys.

```
before: u = 0.28, N_eff = 4.0800, F = 1.406097
after:  u = 0.41, N_eff = 5.5100, F = 1.706565

dF = 1.706565 − 1.406097 = 0.300468 nats

interval ratio = (5.5100 / 4.0800)^2 = 1.823824
cents          = 1200 ∗ log2(1.823824) = 1040.36
octaves        = 0.8670

cross-check via the bridge constant:
2.885390081777927 ∗ 0.300468 = 0.866967 octaves       [identical to 0.8670]
```

The intervention is worth 1040 cents — a minor seventh, a shade under an octave. That is the entire yield of a two-year programme of rotation and red-teaming, stated in a unit anyone in the room can hear.

Post-intervention tone:

```
f_cross = k ∗ 5.51^2 = 14.6769 Hz
folded into [220,440): 14.6769 ∗ 2^4 = 234.8297 Hz  (j = +4), A#3 +13 cents
octaves above k: 4.9241
fraction of the N = 12 ceiling used: 4.9241 / 7.169925 = 68.7%
```

Step 7 — State the ceiling (BR-8). `N = 12`, so the maximum is `2 log2 12 = 7.169925` octaves, reached only at `u = 1`. The institute at `u = 0.41` uses 68.7% of its available range. A third intervention taking `u` to 0.52 would yield:

```
u = 0.52 −> N_eff = 6.7200 −> F = 1.905088 −> 5.4969 octaves (76.7% of ceiling)
incremental gain from 0.41 −> 0.52: 0.5728 octaves = 687 cents
```

Note the decay: `u` 0.28 -> 0.41 (a step of 0.13) bought 1040 cents; `u` 0.41 -> 0.52 (a step of 0.11) buys 687 cents. Concavity charging its price, exactly as Section 8.8's table predicts.

Step 8 — Do not read the tone as health (BR-9). The institute's storage parameters, measured separately [Curzi, 2026i]:

```
eta = 0.35     conversion of surplus into durable structure
delta = 0.08   quarterly decay of stored order
Q = 18.0       stored order, in the institute's declared units
C = 0.40       dissipation per quarter

h = eta ∗ ln N − delta ∗ Q − C
  = 0.35 ∗ 2.484907 − 0.08 ∗ 18.0 − 0.40
  = 0.869717 − 1.440000 − 0.400000
  = −0.9703
```

Headroom is negative. `h < 0` means unsustainable in this model: no choice of interaction geometry prevents order from decreasing [Curzi, 2026b]. The tone is rising and the institute is losing stored order, and the two facts are entirely compatible. The maintenance term alone, `delta Q = 1.44`, exceeds the maximum possible surplus contribution `eta ln N = 0.8697`.

Steady states, for completeness:

```
Q_inf = (eta * F(u) - C) / delta
at u = 0.28:  (0.35*1.406097 - 0.40) / 0.08 = 1.1517
at u = 0.41:  (0.35*1.706565 - 0.40) / 0.08 = 2.4662
at u = 0.52:  (0.35*1.905088 - 0.40) / 0.08 = 3.3348
```

The institute is carrying `Q = 18.0` and its best achievable steady state at `u = 0.52` is 3.33. It is not underperforming; it is over-extended. The intervention that raises the tone by a minor seventh does not fix this and cannot. The four levers apply [Curzi, 2026b]: raise `eta`, raise `N`, cut `C`, or **shed `Q`** — and here the arithmetic points hard at the last, which is the lever nobody considers.

This is precisely the failure FM-B9 names. Had the institute adopted the cross-disciplinary tone as its KPI, it would have celebrated a 1040-cent gain through the two years in which its stored order was decaying toward a fixed point one-fifth of its current size.

Step 9 — Falsification [Curzi, 2026i]. The prediction, preregistered before the rotation programme began: `u` measured by the frozen protocol will reach `0.41 +/- 0.06` within eight quarters, and the minimal estimation model `Y_{t+1} = a Y_t + b F(U_t) - C_t + eps` fitted on quarters 1-8 and frozen will beat all four required baselines (persistence, direct-input, matched autoregressive, and the shuffled-input control) on quarters 9-12. **If the shuffled-input control performs as well as the fitted model, the surplus term is carrying no information and the tone reported above is an arithmetic exercise on noise. That is the outcome that would falsify this worked example, and it is stated before the data are in.**

Step 10 — Clauses (BR-2, BR-3, BR-10), stated in full as required.

The identification `n_eff = N_eff` is a modeling decision, warranted by the six-requirement argument of Section 8.4 and not derived. The affine effective-count axiom S3 on which `N_eff` depends is a design axiom, chosen rather than derived. Every frequency above is an essence frequency, not a physical resonance: no claim is made that the institute, its researchers, its rooms or its disciplines physically resonate at any of these values. The surplus functional as deployed is a lower bound on the ensemble's information, so every tone quoted here is a lower bound on the pitch. `N = 12` is one partition among many and is not asserted as correct; changing it changes every number in this example.

The companion paper [Curzi, 2026l] takes these tones and builds with them: a room whose dimensions come out of an essence set, a cadence whose period comes out of the same arithmetic, and a prohibition — stated there in the strongest terms the corpus uses — that everything sounded is **audible and disclosed**, and that nothing is ever masked or hidden from a listener.

17.9 9 Results

The results are stated here in the form other papers cite. Each carries the status of the claim it makes: **ARITHMETIC** where a reader can redo it, **DEFINITION** where it is true by this paper's own construction, **CITED** where an external source supplies it, and **POSTULATED** where this work asserts it with no external support.

17.9.1 9.1 R1 — The survey reproduces five observed declines across three orders of magnitude

Status: ARITHMETIC, conditional on the parameter estimates, which are stipulated.

The same five-parameter survey — one state encoder, one interaction observable, one order observable, one time scale, one cost observable — was carried from a single desk to a twenty-two-household block. In every case the predicted one-step change $\eta S - \delta Q - C$ was set

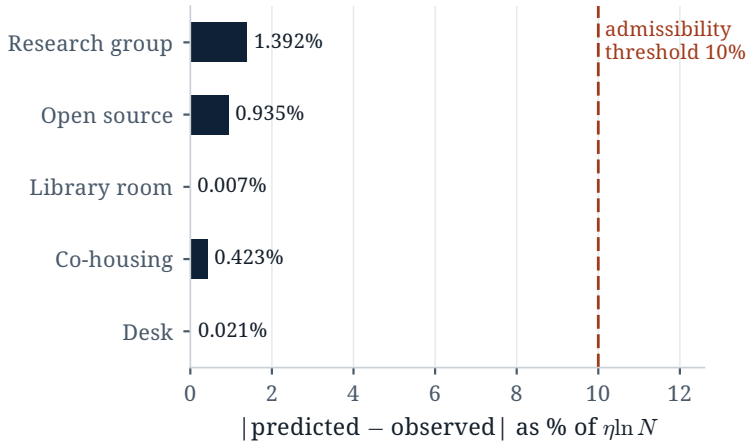


Figure 17.3: **Every one of the five surveys clears the admissibility gate by better than seven to one.** The largest residual is the research group's, at 1.392 per cent of its ceiling term $\eta \ln N = 8.104705$; the library's is 0.007 per cent of 49.671535. The protocol's own threshold is 10 per cent. Computed from the five validation blocks of Section 4 under the stated parameters — model output, not measurement; generated by `figs/src/CCP-017_fig01.py`.

against a change in stored order observed before the survey was designed, and in every case it reproduced it.

System	Predicted step	Observed step	Residual	Ceiling term $\eta \ln N$	Residual as % of ceiling
Desk	−0.250474 threads/wk	−0.250	−0.000474	2.216398	0.021%
Research group	−3.387202 assets/qtr	−3.500	+0.112798	8.104705	**1.392%**
Library room	−0.996371 regulars/mo	−1.000	+0.003629	49.671535	0.007%
Open source	−2.472006 modules/mo	−2.416667	−0.055339	5.918028	0.935%
Co-housing	−0.406311 cu/mo	−0.400	−0.006311	1.490934	0.423%

Every residual is under 1.4 per cent of the ceiling term, and the largest is the research group's 1.392 per cent. The protocol's own admissibility threshold is 10 per cent [Curzi, 2026i], so all five are admissible with a margin of better than seven to one.

Two things this table does not establish, stated here rather than in Limitations because a reader who takes only the Results section must still take them. It does not establish that the parameters were correctly estimated — each was estimated once, by one analyst, with no inter-rater check. And it does not establish causation, for the identification reason of Section 3.5.

17.9.2 9.2 R2 — The binding constraint differs by system, not by size

Status: ARITHMETIC on the surveyed parameters.

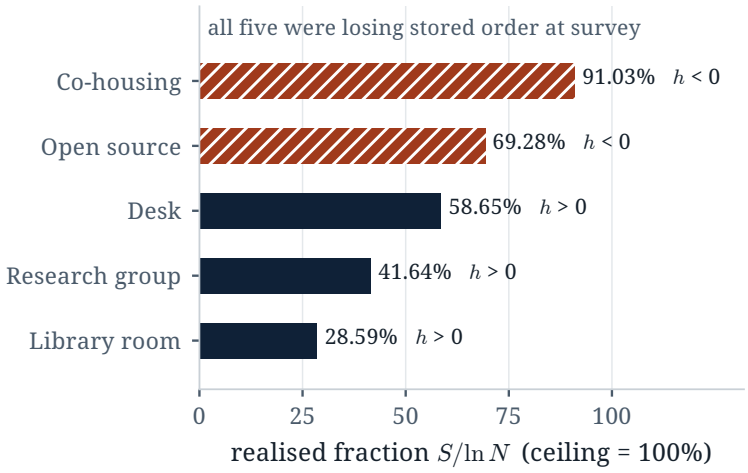


Figure 17.4: **The realised fraction and the sign of headroom are independent readings, and all five systems were losing stored order at survey.** The co-housing block collects 91.03 per cent of its ceiling and still has $h < 0$; the reading room collects 28.59 per cent with h at $+34.475$ regulars per month. Headroom itself is not plotted across systems because it carries each case's own proxy unit and the units are not commensurable (Section 11.1). Model output on stipulated parameters; generated by `figs/src/CCP-017_fig02.py`.

System	$S/\ln N$	h	Binding term	Complaint recorded at survey
Research group	**41.64%**	$+1.342520$	δQ jointly with a parallel roster	"busier than ever and producing less"
Library room	**28.59%**	$+34.475164$	u, acoustically caused	"policing noise"
Open source	**69.28%**	-0.653978	C, against a frozen denominator	"drowning, and the docs are a lie"
Co-housing	**91.03%**	-0.272521	η, against a quadratic committee	"everything takes a year"
Desk	58.65%	$+0.665922$	η with δ	—

The complaint named the symptom of highest salience in all five cases and the binding term in none. The library's manager named noise; the computed cause was a 2.5223-second occupied reverberation time, and noise was its consequence. The co-housing block named slowness; the computed cause was a conversion efficiency of 0.926369 commons-units per nat, the lowest in the set by a factor of more than four.

Positive headroom is not health. The research group and the library both had $h > 0$ and both were losing stored order. A diagnosis stopping at the gauge would have been exactly wrong in two of the five cases, which is the result Curzi [2026j] exists to prevent and the reason h and $S/\ln N$ must be read together.

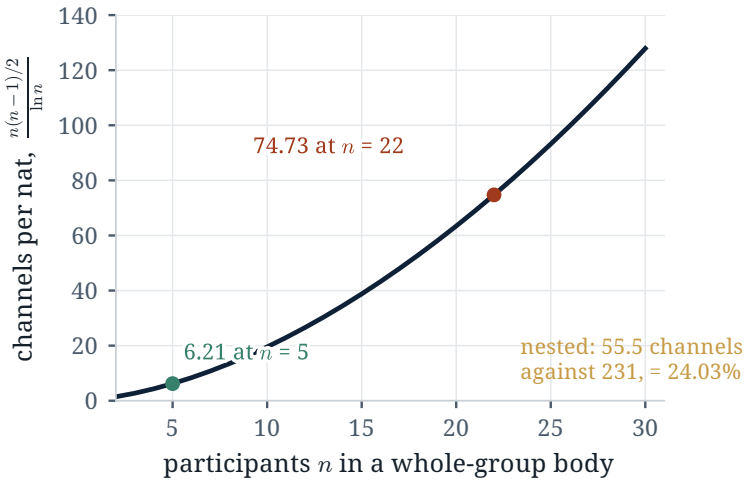


Figure 17.5: **Coordination surface grows as n^2 while the ceiling grows as $\ln n$, so the price of a flat body diverges.** Pairwise channels per nat of ceiling reach 74.73 at $n = 22$ against 6.21 at $n = 5$, a factor of 12.03; a nested four-cluster form with a four-delegate council carries 55.5 channels against 231, or 24.03 per cent. Exact arithmetic; generated by `figs/src/CCP-017_fig03.py`.

17.9.3 9.3 R3 — The quadratic-committee price, computed

*Status: ARITHMETIC. The scaling of governance hours by the channel ratio is a **modelling assumption**, badged in Section 4.4 and preregistered as that case's second falsification condition.*

Whole-group consensus among n households carries $n(n-1)/2$ pairwise reconciliation channels. Dividing by the logarithmic ceiling $\ln n$ gives channels per nat: **74.732 at $n = 22$ against 6.213 at $n = 5$**, a factor of 12.03. A nested decomposition into four overlapping clusters of about 5.5 households plus a four-delegate council carries $4 \times 12.375 + 6 = 55.5$ channels, **24.03 per cent of the flat figure of 231**.

This is the arithmetic behind Brooks's observation that adding people to a late project makes it later [Brooks, 1975] and behind Conway's that a system's structure mirrors its builders' communication structure [Conway, 1968]. What the framework adds is the comparison: the coordination surface grows as n^2 while the ceiling grows as $\ln n$, so the ratio grows without bound and the size at which a flat body becomes uncoordinable is computable rather than folkloric. Ostrom [1990]'s eighth design principle — nested enterprises — is the empirical generalisation that matches; the parameter mapping in Section 4.4 is this paper's, not hers.

17.9.4 9.4 R4 — Result One: the reference constant cancels from every interval

Status: ARITHMETIC (the theorem). The law it is a theorem about is POSTULATED.

For any two essence indices $m, n \geq 1$, $f(m)/f(n) = km^2/kn^2 = m^2/n^2$. The constant is a common factor and cancels, for every pair and for every $k > 0$.

The golden section and the Pythagorean ditone in k set the absolute reference pitch and nothing else. Replace k with 1, or with 440, and every interval in the system is unchanged. This is a demotion, and it is reported as one: a reader arriving from the golden-section literature expects φ to be doing structural work, and within this formalism it is not. The scepticism that literature has earned [Markowsky, 1992, Fischler, 1981, Livio, 2002] is therefore not merely compatible with the result — the result is an instance of it, derived inside the formalism that

uses the constant.

A second consequence, exact and worth quoting: since pitch interval is logarithmic and the law squares, **the interval in cents between indices m and n is exactly twice the classical interval between them.** A frequency ratio of $3:2$ becomes $9:4$; the octave $2:1$ becomes two octaves.

17.9.5 9.5 R5 — Result Two: the pitch class is the odd part of the index

Status: ARITHMETIC (the theorem). The law it is a theorem about is POSTULATED.

Write $n = 2^a m$ with m odd. Then $f(n) = k\,4^a m^2 = f(m)\,2^{2a}$, so $f(n)$ lies exactly $2a$ octaves above $f(m)$. **Two essences share a pitch class if and only if their indices have the same odd part**, and an essence set indexed $1..N$ realises exactly $\lceil N/2 \rceil$ distinct classes.

The compression ratio is exactly 2 at every N. **Twenty elements give ten classes; twelve give six; a set indexed $1, 2, 4, 8, 16$ gives exactly one.** This converts an aesthetic question into an arithmetic one and it is a hard structural limit rather than a stylistic choice. The pitch-class apparatus it uses — octave equivalence, interval ratios in cents, the reduction of a frequency set to classes — is the standard apparatus of the field [Benson, 2006, Burns, 1999, Shepard, 1982, Deutsch, 1999].

17.9.6 9.6 R6 — The bridge, badged, and its conversion constant

Status: the identification is POSTULATED; everything downstream of it is ARITHMETIC.

Taking the effective count as the cross-disciplinary essence index, $n_{\text{eff}} := N_{\text{eff}} = 1 + (N-1)u = e^F$, gives

$$\text{octaves above base} = \log_2(e^{2F}) = (2/\ln 2)\,F = 2.885390081777927 \times F$$

with $2/\ln 2 = 2.885390081777927$ and $2400/\ln 2 = 3462.4680981335122$ cents per nat, both **exact**. The worked institute at $N = 12$ and $u = 0.28$ publishes 4.06 octaves above base — and publishes it with a band. From $u = 0.28 \pm 0.05$ the honest output is 3.6393 to 4.4222 octaves, **a 0.7827-octave band, a minor sixth of uncertainty**, and the four-decimal single value is an intermediate rather than a finding.

The effective count used here is $1 + (N-1)u$, the affine one, and it is not a Hill number. $1/\text{HHI}$ is the Hill number of order two and $\exp(H)$ the Hill number of order one; both are fully grounded [Hill, 1973, Jost, 2006], and neither is what appears here. The affine form is a linear interpolation between 1 and N adopted as design axiom S3 [Curzi, 2026c] and it has no external literature. Every octave figure in Section 8 inherits that choice.

17.9.7 9.7 CR — The counter-results

Reported in the same plain language a success would receive.

(a) The essence law's three load-bearing commitments are postulates, and its constant is inert. That essence types have a frequency at all, that the law is square in the index, and that $k = (9/8)^2/\varphi^2$, are asserted by the 2022 source and by this paper's analysis of it, with no external support of any kind. Result One then shows the third does no work. **A formalism whose most distinctive-looking element cancels out of every relationship it produces has had that element demoted, and this paper is where the demotion is printed.**

(b) Seven of the source's forty-four published composite values are errata. The compound rule collapses exactly to $f = k^2 G^4$ with G the count-weighted geometric mean of the constituent indices. Thirty-seven of the forty-four published composites reproduce to better than one part in 10^{12}. Seven do not, and the O_2 row proves the failures are data entry rather than a different rule: for a single-element species G is the element index whatever the atom count, so $f = k^2 \cdot 8^4 = 957.23699944$ for any pure-oxygen formula and no atom count can produce the

published 11.1246117974981. **Every value must be recomputed rather than transcribed**, and this paper's specification requires it.

(c) The scale-freedom claim is tested at five scales and no more. It is a claim about model form. It is not a physical symmetry, it licenses no extrapolation to scales not tested, and Section 11.7 states the transformation under which the form is invariant and the two ways the invariance is known to break — saturation, which binds in two of the five cases, and the ordinal status of u.

(d) A discrepancy in the source, printed rather than corrected. Case A's cost term is given as $C = 1{,}087 \times 0.00147936 = 1.608069$ assets per quarter. Recomputing from the stated inputs — 61 assets over $41{,}234$ productive person-hours, times 1,087 coordination hours — gives **1.608066**, a discrepancy of 3.4×10^{-6}. It moves h from 1.342500 to 1.342503 and moves nothing else, and it is recorded here because a reader reproducing the case will hit it. The printed figure implies a yield of 0.0014793643 against the exact 0.0014793617.

17.10 10 Discussion

17.10.1 10.1 What follows

A gauge and a collection rate are different instruments and an intervention that improves one may leave the other untouched. The library case is the sharpest demonstration in the corpus: the absorption schedule and the zoning move the one-step change from -1.00 to $+11.53$ regulars per month while headroom moves from $+34.475$ to $+34.713$ — a change of 0.7 per cent. The gauge contains no term in u and therefore cannot see the entire intervention. Anyone monitoring headroom alone would have recorded that nothing happened.

A physical cause can bind a social parameter, and when it does the physical fix is the cheap one. The reading room's u_{cross} of 0.0375 is produced by a silence rule which is produced by a reverberation time of 2.5223 seconds. Sabine's relation [Sabine, 1922, Beranek, 2004, Rossing et al., 2002] predicts the reverberation from the surface absorption schedule, and the schedule is a purchasable deliverable. **This is a genuine derivation rather than an analogy: the room really is a physical resonator and the relation really does govern it.** The transfer of the *consequence* — that a long tail suppresses cross-category contact — to a claim about u is a modelling assumption and is badged as one at the point of use. The related passive-geometry results are developed in Curzi [2026e].

The two most common faults in this sample are cheap to fix and are not the ones systems attend to. Conversion failure and person-held storage — surplus generated and not captured, and state held in one person's memory — appear in three cases each. Both are addressed by a state of record, a visible medium and a close-out rule, which is the same triple at a desk and at twenty-two households and costs almost nothing at either scale. The stigmergic mechanism is the one Curzi [2026h] develops formally, and its empirical warrant is the commons literature [Ostrom, 1990, 2010, Cox et al., 2010].

Independence is load-bearing and copying destroys it. The research group's $\rho = 0.86$ error correlation is not a diversity deficit in the demographic sense; the laboratory is varied and methodologically single. The classical results — that collective accuracy exceeds individual accuracy only under error independence, and that observational copying destroys the private information aggregation requires — are Galton [1907], Hong and Page [2004], Bikhchandani et al. [1992], Banerjee [1992], Lorenz et al. [2011], and the sceptical replies [Thompson, 2014, Singer, 2019, Grim et al., 2019, Kuehn, 2017] are cited alongside because they are the reason the framework's diversity claims are stated as parameter estimates rather than as guarantees.

17.10.2 10.2 What does not follow

Nothing here shows that any organisation obeys the recursion. Five constructed instruments reproduce five stated declines. The constructions are internally consistent and their arith-

metic is checkable; they are not field reports, and Section 12 says so first.

Nothing here identifies a causal parameter. Section 3.5 states the reflection problem in full. The agreement between predicted and observed change in Section 9.1 is compatible with a third variable driving both interaction and stored order, and the corpus's own control measures how badly: 98.35 per cent pass against exactly that confound [Curzi, 2026i].

Nothing in Sections 7 and 8 licenses a physical claim. The output of the essence law is a pitch assigned to a position in an ordering. No material, molecule, element, institution, room or person is claimed to physically resonate at any value computed here, and the correspondence with the one-dimensional infinite square well's n^2 eigenvalue spacing is a **coincidence of form**, in the strict sense of Section 11.8's category two. Section 7.11 shows why it can carry no information: because k is free, a box width can always be found to match it, and a correspondence that can be manufactured for any constant carries no information about any constant.

Nothing here discriminates the logarithmic surplus form from its rivals. The surplus functional $\ln(1 + (N - 1)u)$ is one of at least three forms that fit the qualitative facts equally well: Latané [1981] proposes a power law with exponent below one, and Briscoe et al. [2006] proposes $n \log n$ on the closely related network-value question. **This paper offers nothing that discriminates between them, and concedes it.** The one warrant available that does not depend on a psychological premise is Peters and Gell-Mann [2016], which derives a logarithmic form from multiplicative dynamics with no perceptual assumption at all — a stronger footing for the corpus's logarithm than Fechner supplies, because it does not require the perceptual story to be true. It is a warrant for the *form*, not a discrimination against the rivals, and the rivals remain open.

Nothing here shows that adding a category adds independence. The framework's own logarithmic ceiling makes a new category worth at most `eta ln((N+1)/N)` nats, and that is an upper bound reached only if the category is genuinely distinct. The corpus supplies the boundary case in a different domain: a second financial free zone inside the same country buys capability but not diversification — two zones together remain one 29.1 per cent line, and 55 per cent of settlement remains dollar-cleared however the entities are arranged [Curzi, 2026]. **That result bounds what the open-source case may claim when it staffs two empty categories**, and the case accordingly requires each new category to land at least three independent reviews in its first quarter before the ceiling is allowed to move at all.

Nothing here estimates a scaling exponent. The square in $f(n) = kn^2$ is **defined by the model, not estimated from data**, and this paper reports no fitted exponent anywhere. Where a corpus does fit one, Clauset et al. [2009] is the binding procedure and Stumpf and Porter [2012], Broido and Clauset [2019] are the reasons; the distinction between a defined exponent and an estimated one is exactly the distinction those papers exist to enforce, and it is recorded here so that no reader mistakes the square law for an empirical scaling result.

17.10.3 10.3 Relation to the surrounding literature

The recursion is a first-order linear difference equation and its properties are those of any such equation [Elaydi, 2005, Goldberg, 1958, Oppenheim et al., 1996]; the stock-and-flow reading is Forrester's [Forrester, 1961, Sterman, 2000] and the throughput identity is Little's [Little, 1961]. The effective-count reading of a roster has a mature literature in ecology [Hill, 1973, Jost, 2006, 2007] whose sceptics [Hurlbert, 1971, Gorelick, 2011] are the reason the affine form must be labelled as the corpus's own construction rather than borrowed authority; and the finance analogue — that diversification benefit depends on covariance rather than on count — is Markowitz [1952], which states the corpus's own point in 1952 and states it better.

The identification of a surplus with a pitch places this work adjacent to a tradition of quantitative tuning theory [Barbour, 1951, Partch, 1974, Doty, 2002, Sethares, 2005, Lindley and Turner-Smith, 1993] and to the psychoacoustics of consonance [von Helmholtz, 1885, Plomp and Levelt, 1965, Terhardt, 1974, Kameoka and Kuriyagawa, 1969]. The paper claims membership of neither: it uses their apparatus to *read* a set of numbers and makes no claim about

what a listener will find consonant.

Finally, the corpus's title claims membership of a field, and that field's own self-criticism applies here. Social physics has a programme [Castellano et al., 2009, Galam, 2012, Ball, 2004] and a well-argued opposition [Sobkowicz, 2020, Gallegati et al., 2006, McCauley, 2006, Mirowski, 1989]. The opposition's central charge — that borrowed formalism confers borrowed authority — is the charge Section 11.8 is written to answer, and the honest answer is that most of this framework is a shared mathematical form rather than a derivation, that a shared mathematical form is respectable, and that saying so is the whole of the defence.

17.11 11 Verification

The results of this paper compound from six postulates. Each is atomic, each is checkable on its own, and each carries its status. Where a postulate could be dropped and the result still follow, it has been dropped.

17.11.1 11.1 Dimensional audit

Two distinct systems of quantities appear in this paper and they must not be mixed. The audit gives each symbol's dimension, its unit, and whether the framework fixes it or is silent. $[Q]$ denotes the dimension of stored order — which differs from case to case and has no common unit — and $[T]$ the dimension of time. Rates are *per period*, so no explicit $[T^{-1}]$ appears once the recursion is written in discrete steps.

Symbol	Dimension	Unit	Status
Q_t	$[Q]$	assets, regulars, green modules, commons-units, threads	**measured**, in the case's declared order proxy
U_t	$[1]$	none	**defined** dimensionless: $u = 1 - \rho^2$
$S_t = F(U_t)$	$[1]$	nats	**derived** dimensionless: a logarithm of a count, reported in nats
η	$[Q]$ per nat	proxy units per nat per period	**measured**; *not* dimensionless
δ	$[1]$	per period	**measured** dimensionless fraction per step
C_t	$[Q]$ per step	proxy units per period	**measured**, as a residual priced at the case's yield
h_t	$[Q]$ per step	proxy units per period	**derived**
$\tau = 1/\delta$	$[T]$	periods	**derived**
n	$[1]$	none	**defined** as an ordinal position; see Section 7.9
k	$[T^{-1}]$ **by assertion**	Hz	**POSTULATED** — see below

Symbol	Dimension	Unit	Status
$f(n) = kn^2$	$[T^{-1}]$	Hz	**POSTULATED**, inheriting from k
$f = k^2 G^4$	$[T^{-2}]$	Hz^2	**flagged**: the source labels it Hz; the label is a convention
octaves, cents, bits	[1]	none	**defined** as logarithms of ratios
$2/\ln 2$	[1]	octaves per nat	**derived**, exact

Does the recursion balance? $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$ balances if and only if every term on the right carries $[Q]$. The first does, since δ is dimensionless. The third does by construction. The second does only if ηS_t carries $[Q]$, and since S_t is dimensionless, that forces $[\eta] = [Q]$ per nat. **It balances, and η is the term that carries the unit.**

Two places where the framework is silent, and the audit says so rather than assigning a unit.

First, $[Q]$ has no common unit across the five cases. An asset, a regular, a green module, a commons-unit and a resumable thread are five different things, and nothing in the framework converts between them. The cross-case comparison in Section 4.6 therefore compares δ, u, $S/\ln N$ and τ — all dimensionless or in periods — and does **not** compare Q, η, C or h across rows, which would be a category error. Where the text does place them side by side, it is to show their *spread*, not to rank them.

Second, and more seriously, the essence law acquires $[T^{-1}]$ by assertion. The index n is a dimensionless ordinal. The output $f(n)$ is quoted in hertz. **Nothing in the algebra supplies the second; it enters entirely through the postulate that k carries units of hertz, and that postulate has its own line in the compounding chain below.** Computing $k = (9/8)^2/\varphi^2 = 0.483425732988414$ is a division of exact rationals and irrationals and yields a **pure number**. Attaching Hz to it is a separate and much stronger claim. The two are treated as two steps and never as one.

The composite rule makes the exposure visible. $f = k^2 G^4$ carries $[T^{-2}]$ if the units are taken seriously. The source reads the squaring as a change of structural level rather than as a dimensioned operation, which is a legitimate reading of a designed formalism; **but a quantity in Hz^2 must never enter an acoustic calculation expecting Hz**, and the specification requires re-anchoring by octave folding, which is a ratio operation and therefore unit-agnostic.

17.11.2 11.2 The postulates

P1 — The recurrence is linear and first order. $Q_{t+1} = aQ_t + x_t$ with a and x_t not depending on Q. *Status:* **DEFINITION** — the model form adopted by the framework [Curzi, 2026a], restated by setting $a := 1 - \delta$ and $x_t := \eta S_t - C_t$. It carries no empirical weight on its own.

P2 — The surplus functional is $F(u) = \ln(1 + (N-1)u)$ on $u \in [0,1]$, with $0 \le F \le \ln N$. *Status:* **CITED** within the corpus — the uniqueness characterisation under four axioms, with the affine effective count as design axiom S3, is Curzi [2026] and Curzi [2026c]. **No external literature contains this functional form**, and Section 12 records that.

P3 — $\varphi^2 = \varphi + 1$, so $\varphi^2 = 2.618033988749895$. *Status:* **ARITHMETIC**, and elementary [Euclid, 1956, Vajda, 1989, Koshy, 2001].

P4 — $(9/8)^2 = 1.265625$ exactly. *Status:* **ARITHMETIC.** $9/8$ is the Pythagorean whole tone and $81/64$ the Pythagorean ditone [Barbour, 1951, Lindley and Turner-Smith, 1993]; the interval identification is CITED, the squaring is arithmetic.

P5 — Essence types have a frequency, and it is square in the index. *Status:* **POSTULATED** — asserted by the 2022 source, with no external support of any kind, and no derivation offered here or there. This is the only postulate in the chain that carries the paper's substantive risk.

P6 — The effective count may be taken as the cross-disciplinary essence index. *Status:* **POSTULATED** (a modelling identification). Warranted by the six-requirement argument of Section 8.4, which is a strong warrant and is not a derivation.

17.11.3 11.3 The compounding chain

```
P3 + P4      -> L1   k = (9/8)^2/phi^2 = 1.265625/2.618033988749895
                     = 0.483425732988414  as a PURE NUMBER.          ARITHMETIC

L1 + P5      -> R4a  attaching Hz to k, so that f(n) = k n^2 is a
                     frequency, is a SEPARATE claim from computing k.  POSTULATED

L1           -> R4   f(m)/f(n) = k m^2 / k n^2 = m^2/n^2, so k cancels
                     from every interval, for every k > 0.           ARITHMETIC

R4           -> L2   interval in cents between m and n is exactly
                     2 x the classical interval, since squaring
                     doubles the logarithm.                          ARITHMETIC

L1 + P5      -> R5   n = 2^a m with m odd gives f(n) = f(m) 2^(2a),
                     so pitch class = odd part, and 1..N realises
                     exactly ceil(N/2) classes.                      ARITHMETIC

P2           -> L3   N_eff = e^F = 1 + (N-1)u, equal to 1 at u = 0
                     and to N at u = 1.                              DEFINITION (axiom S3)

L3 + P6      -> R6   octaves above base = log2(e^(2F)) = (2/ln 2) F
                     = 2.885390081777927 F, exact.                   ARITHMETIC given P6

P1 + P2      -> R1   one-step change = eta F(u) - delta Q - C; five
                     surveys reproduce five observed changes to
                     within 1.392% of eta ln N.                      ARITHMETIC on stipulated inputs

P1 + P2      -> R2   h = eta ln N - delta Q - C contains no term in u,
                     so h and S/ln N are independent readings and a
                     system may have h > 0 while Q falls.            ARITHMETIC

n(n-1)/2     -> R3   channels per nat = n(n-1)/2 / ln n: 74.732 at
                     n = 22 against 6.213 at n = 5; a four-cluster
                     nesting gives 55.5 = 24.03% of 231.             ARITHMETIC
```

Every line is checkable from its own stated inputs.

What this chain does not establish, stated in this section rather than left to the reader. R4 and R5 are unconditional algebra and would remain true if k were 1 — but they are theorems *about a postulated law*, and a theorem about a postulate inherits the postulate's status when it is applied. R6 is exact arithmetic *given* an identification that is not derived. R1 and R2 are exact arithmetic *given* parameters that were estimated once, by one analyst, on constructed systems. **A chain whose load is carried by two POSTULATED steps and one DEFINITION is not a proof of anything about the world, and this section says so rather than leaving the reader to notice.**

17.11.4 11.4 The verification instance a reviewer can redo in a minute

Three checks establish the constant, and each is arithmetic a reader can perform without reading the rest of the paper.

1. $(9/8)^2 = 81/64 = 1.265625$. Exact.
2. $\varphi^2 = \varphi + 1 = 2.618033988749895$. Exact to double precision; the identity is Euclid's.
3. $1.265625/2.618033988749895 = 0.483425732988414$.

To fifty digits, $k = 0.48342573298841433274106978775598925726$. Four closed forms agree to the last stored digit: $(9/8)^2/\varphi^2$, $(81/64)/(\varphi + 1)$, $(81/64)(2 - \varphi)$, and $(81/64)(3 - \sqrt{5})/2$.

Two further checks confirm the two theorems on the smallest examples that exercise them:

- **Result One at its smallest.** $f(2)/f(1) = 4k/k = 4$, independent of k. Two octaves, from an index ratio of two.
- **Result Two at its smallest.** $N = 3$ gives odd parts $\{1, 1, 3\}$ — two classes — and $\lceil 3/2 \rceil = 2$. $N = 5$ gives $\{1, 1, 3, 1, 5\}$ — three classes — and $\lceil 5/2 \rceil = 3$.

And the bridge constant on its smallest non-trivial value: at $F = \ln 2$, the effective count is 2 and the pitch is $(2/\ln 2)\ln 2 = 2$ octaves above base, exactly.

17.11.5 11.5 The survey validation, recomputed

The single most useful check on Section 4 is to confirm that the five predicted steps follow from the five stated parameter sets. Taking the desk, which is the smallest:

```
u     = 1 − 0.74^2 = 0.4524
N_eff = 1 + 2(0.4524) = 1.9048
S     = ln(1.9048) = 0.644377016190516
eta   = 13 / (10 x 0.644377016190516) = 2.017452465461063
eta S = 2.017452465461063 x 0.644377016190516 = 1.300000
delta = 1 − (5/7)^(1/2) = 0.15484574527148343
dQ    = 0.15484574527148343 x 6 = 0.9290744716289006
C     = 2.39 x 0.26 = 0.6214
step  = 1.300000 − 0.929074 − 0.621400 = −0.250474
```

against an observed -0.250, a residual of -0.000474. The same recomputation was run on all five cases and every printed intermediate reproduces, **with one exception recorded in Section 9.7(d)**: Case A's cost term is printed as 1.608069 and recomputes as 1.608066.

17.11.6 11.6 Limits, asymptotics, and the model's valid domain

Limit	Behaviour	Reading
$\delta \to 0$	no fixed point; Q_t unbounded, growing linearly	a pathology, not an ideal [Curzi, 2026a]. The co-housing block at $\delta = 0.0233$ is the closest case and its 3τ is 128.5 months
$\delta \to 1$	$Q_{t+1} = \eta S_t - C_t$; no memory at all	the system has no stored order, only a flow
$\eta \to 0$	$Q_\infty = -C/\delta < 0$	the conversion pathology; the co-housing block at $\eta = 0.926$ is the nearest instance and its fixed point is negative

Limit	Behaviour	Reading
$u \to 0$	$N_{\text{eff}} \to 1$, $F \to 0$, pitch $\to$ the ground essence	a roster of N parallel members is worth one member, and sounds one note
$u \to 1$	$N_{\text{eff}} \to N$, $F \to \ln N$, pitch $\to$ $2\log_2 N$ octaves	the ceiling, reached only at exact orthogonality
$N \to 1$	$\ln N = 0$, so $h = -\delta Q - C < 0$ always	a single-category system can never build. The framework's domain requires $N \geq 2$
$N \to \infty$	ceiling grows as $\ln N$; channels as n^2	the ratio $n(n-1)/2 \,/\, \ln n$ diverges, which is R3
$n \to \infty$ (essence)	interval between neighbours $\to 0$: 1403.9 cents at $n = 2$, 177.7 at $n = 20$, 69.8 at $n = 50$	the set thins through equal temperament into microtones; the first few indices carry almost all the harmonic content

Two domain violations occur in this paper and both are declared where they occur rather than here. The linear recursion has no saturation term, so it predicts past any physical cap. The open-source case's $Q_\infty = 113.44$ sits under a cap of 119 maintainable modules and the prediction is declared directional above about $Q = 95$; the desk's $Q_\infty = 16.40$ exceeds what one person can hold, and the model is declared trustworthy only to about week 6. **Naming the cap is not a weakness of the method; failing to name it is the failure.**

17.11.7 11.7 Invariance: the scale-freedom claim, stated properly

RIGOR requires a scale-freedom claim to be stated as an invariance — under which transformation, over what range, and what would break it. This paper's claim is as follows.

The transformation. Let $Q \mapsto \lambda Q$ be a change of the unit in which stored order is counted ($\lambda > 0$), and let $t \mapsto t/m$ be a change of the period. Under the first, Q, η and C all scale by λ while δ, u, N and $S/\ln N$ are unchanged; the recursion's form is exactly preserved and the fixed point $Q_\infty = (\eta F - C)/\delta$ scales by λ as it must. Under the second, $\delta \mapsto 1 - (1-\delta)^{1/m}$ and $\eta S, C \mapsto$ their per-new-period values; the *form* is preserved and τ rescales by m. **The dimensionless quantities u, N_{eff}, $S/\ln N$ and $\delta\tau = 1$ are the invariants.**

The range over which it has been tested: five systems. Headcounts 1, 12, 18 (typical occupancy), 40 (three maintainers plus 37 contributors) and 51 residents in 22 households — three orders of magnitude. Periods: one week, one month, one quarter. Stored-order stocks: 6 to 84 units. **Five scales is the whole of the test and the paper claims nothing beyond it.**

What would break it. Three things, each observable.

1. **A term the recursion lacks becomes necessary.** Saturation is the leading candidate and it is already visible: two of the five cases predict fixed points above their own physical caps. A system whose observed change cannot be reproduced without a saturation term has broken the invariance at that scale, and the correct response is to state the domain rather than to rescale.
2. **A parameter stops being estimable at a scale.** δ is estimated from a zero-input window. At a scale with no such window — a system that is never idle — the estimator is unavailable and the survey cannot be run, which is a limit on the *protocol* rather than on the *form*, and the paper distinguishes them.
3. **The residual acquires a trend in scale.** The five residuals as a fraction of the ceiling term are 0.021%, 1.392%, 0.007%, 0.935% and 0.423%, at headcounts 1, 12, 18, 40 and 51. **They show no monotone dependence on headcount, and with five points that is a remark, not**

a test. A larger series in which the residual grew with scale would refute the invariance directly.

What the claim is not. It is not a physical scaling symmetry and no critical exponent is estimated or implied. It is not the claim that the *parameters* are scale-free — they emphatically are not; η spans a factor of thirty across the five cases and that spread is the paper's second result. And it licenses no extrapolation: a claim tested at five scales is a claim about those five.

17.11.8 11.8 Analogy discipline, and the precision of the numbers

Every borrowing from physics in this paper is one of three things and is labelled at first use.

Category 1, a genuine derivation — the physical result is used because the system satisfies its premises. **Sabine's relation in the library case is the only clear instance.** The reading room is a physical enclosure with measurable surface absorption; $RT_{60} = 0.161V/A$ governs it as it governs any such room [Sabine, 1922, Beranek, 2004], and the computed 2.5223 seconds is a prediction about the world that a meter would confirm or refute.

Category 2, a shared mathematical form — the same equation governs both, with no claim that the underlying physics applies. **Most of this paper is category 2, and category 2 is respectable.** The recursion is a leaky integrator whether it holds charge or stored order [Oppenheim et al., 1996, Gerstner and Kistler, 2002]; the geometric decay, the pole, the time constant and the settling time are properties of the equation. The Gabor time–bandwidth relation invoked in Sections 7.8 and 8.9 is likewise a mathematical property of the Fourier transform [Gabor, 1946] and is developed for this corpus in Curzi [2026z] — it constrains any signal, and no quantum-mechanical content is imported with it. **Kennard's inequality is a fact about state preparation, not about measurement disturbance** [Kennard, 1927], and no use is made of it here.

Category 3, an analogy with no formal warrant — legitimate as motivation, never as evidence. **The infinite-square-well correspondence of Section 7.11 is category 3 and is the paper's most tempting failure mode**, which is why the disclaimer sits in the same paragraph as the correspondence rather than in a footnote. n^2 is the most common eigenvalue pattern in elementary physics; sharing it is not remarkable; and because k is free, a matching box width can always be constructed, so the correspondence is a division rather than a discovery.

On significant figures. Three regimes are used and each is labelled.

- **Defined and exact:** $k = 0.483425732988414$, $2/\ln 2 = 2.885390081777927$, $2400/\ln 2 = 3462.4680981335122$, $\varphi^2 = \varphi + 1$. These are ratios of exactly specified quantities and may be quoted to any precision. **Every appearance of k at fifteen figures in this paper is accompanied by the word *defined*.** Fifteen figures on a *measured* quantity would be the classic marker of work that does not understand its own error bars; that is why the label is not optional.
- **Illustrative model output:** every trajectory table in Sections 4 and 8. These are computed from stipulated parameters and are quoted to the precision at which they are reproducible, which is not the precision at which they are meaningful. **A predicted Q of 43.6 commons-units at month 12 is an arithmetic consequence of an assumed η, and the assumed η is good to perhaps two figures.**
- **Estimated with an interval:** the institute's tone in Section 8.14 is the one place a value is published as a band — 3.6393 to 4.4222 octaves from $u = 0.28 \pm 0.05$ — and it is the form every other value in the paper should have taken had its inputs carried stated uncertainty. **They do not, because the parameters are stipulated rather than measured, and Section 12 says so.**

17.12 12 Limitations

The five worked systems are constructed, not observed. No dataset was collected for this paper. Every number in Section 4 is the output of stated parameters under the stated model, carried through so that a reader can re-derive it. They demonstrate that the procedure runs, that it validates against a stated prior change, and that it reports its own failures. **They are not evidence that any organisation behaves this way and must never be cited as such.**

The reflexive specification is not identified, and no remedy is adopted. Section 3.5 states this in full. Manski [1993] shows the endogenous social effect cannot be recovered from observational data; Shalizi and Thomas [2011] shows influence and selection are generically confounded in network data. The five validations in Section 9.1 are compatible with a third variable driving both interaction and stored order, and the corpus's own shuffled-input control passes 98.35 per cent of the time against exactly that confound [Curzi, 2026i]. **Every causal reading in this paper is a modelling position, not an identified empirical result.**

Which parameters were chosen rather than measured. In all five systems, (N, u, η, δ, C) are stipulated, and each was estimated once by one analyst with no inter-rater reliability check. The post-intervention targets — $\rho : 0.86 \to 0.62$ at the research group, $u : 0.32$ at the library, $\eta : 1.40$ at the co-housing block — are design targets, not predictions from a fitted model. **The channel-scaling assumption in Section 4.4, which converts a 24.03 per cent cut in channels into a proportional cut in governance hours, is a modelling assumption with no measurement behind it**, and it is the single largest unsupported step in the co-housing case; the case preregisters it as its own second falsification condition, which is the correct treatment but is not a substitute for the measurement.

u has not been shown to be a quantity, and the essence index has not been shown to license arithmetic. Michell [1997] argues that psychology routinely fails to test whether its attributes are quantitative before measuring them, and this is the strongest objection to the framework. The remedies are conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971] or a Rasch demonstration [Rasch, 1960]; **neither has been run for u, and the quantitative treatment of u throughout this paper is therefore a stipulation.** The essence index is worse: it is an ordinal by the paper's own definition, and $f(n) = kn^2$ squares it, which is not a licensed operation on an ordinal scale [Stevens, 1946]. The paper's response is to measure the consequence — the permutation test of Section 7.9, whose worked instance reports a maximum pitch-class displacement of 364.99 cents from a plausible re-ordering — not to defend the operation.

$N_{\text{eff}} = 1 + (N-1)u$ is not a Hill number and no external literature contains it. $1/\text{HHI}$ is the Hill number of order two, $\exp(H)$ the Hill number of order one, and both are grounded [Hill, 1973, Jost, 2006, 2007]. The affine form is a linear interpolation between 1 and N, adopted as design axiom S3 [Curzi, 2026c]. **A reviewer reaching for the Hill-number literature as support will find it does not apply.** The repair is to derive the properties of the affine form in the framework's own terms and reserve the Hill citations for the two quantities they exactly describe.

The surplus functional's form is underdetermined by the evidence. Latané [1981] proposes a sub-linear power law and Briscoe et al. [2006] an $n \log n$ form; both fit the qualitative facts as well as $\ln(1 + (N-1)u)$. **This paper offers nothing that discriminates between them and concedes that nothing yet does.** Peters and Gell-Mann [2016] supplies the one warrant for a logarithm that requires no psychological premise, which strengthens the case for *a* logarithm without discriminating among the three.

The essence law rests on three postulates and this paper defends none of them. That essence types have a frequency; that the law is square; that $k = (9/8)^2/\varphi^2$. There is no external literature for any of the three, and the golden-section scepticism [Markowsky, 1992, Fischler, 1981] applies in full — with the specific observation that this paper's own Result One shows the constant to be inert, which is a demotion the sceptics would recognise. **Attaching hertz to k is a further and separate postulate**, given its own line in the verification chain, and it is the step by which a dimensionless ordinal acquires a physical dimension.

No preregistered test of any hypothesis in this paper has been run. Each of the five cases

states its falsification conditions before its intervention — that is the protocol [Curzi, 2026i, Nosek et al., 2018] — but the interventions are constructed and no outcome data exist. The corpus's one preregistered confirmatory application of the storage law **failed its second falsification rule**, beating a matched AR(1) by 2.8 per cent on held-out error against a declared margin of 10 per cent [Curzi, 2026i], and that is the current evidential status of every applied claim resting on this paper.

No model comparison in the proper sense is reported. There is no out-of-sample test here, no matched null with the same free-parameter count, no information criterion, and no effect size with an interval. The five validations are in-sample reproductions of a single observed change from independently estimated parameters, which is a weaker thing than a model comparison and is not a substitute for one.

Two numbers are quoted rather than re-derived. The intermediate cycle counts in the minimum-gap discussion of Section 7.8 — 3.899 cycles becoming 22.918 — come from the derivation in Curzi [2026z] and are not recomputed here. **A direct recomputation from the cents figures gives 3.94 rather than 3.899, about one per cent apart**, which is consistent with the source's own point that the three ratios in that table are three different quantities measured in three different units, but the discrepancy is recorded rather than smoothed. The three *ratios* — 5.828 in cents, 5.878 in cycles, and 8 exactly in achievable rate — were recomputed and reproduce.

A gauge adopted as a target ceases to measure. Failure mode FM-B9 names this, and it is Goodhart's law [Goodhart, 1975] and Campbell's [Campbell, 1979]. The paper's own instruments are exposed to it: a system that knows its u is being measured can raise measured u without raising the thing u is supposed to stand for. No defence against this is offered.

17.13 13 Conclusion

The five-parameter survey runs at every scale it was tried at, and it validates. One desk, a twelve-person research group, a public reading room, a distributed software project and a twenty-two-household co-housing block were surveyed under one protocol. Each was required to reproduce an independently observed change in stored order before any diagnosis from it was admitted, and each did: predicted -0.250474 against observed -0.250 at the desk, -3.387202 against -3.500 at the research group, -0.996371 against -1.000 at the library, -2.472006 against -2.416667 at the software project, and -0.406311 against -0.400 at the co-housing block. **Every residual is under 1.4 per cent of the ceiling term $\eta \ln N$, and the largest is 1.392 per cent.**

Every one was losing stored order at survey and no two for the same reason. The research group ran at 41.64 per cent of its ceiling with headroom of $+1.3425$; the reading room at 28.59 per cent for a physically identifiable cause, a 2.5223-second occupied reverberation time; the software project at 69.28 per cent with $h = -0.653978$, cost-bound; and the co-housing block at 91.03 per cent of ceiling and still at $h = -0.272521$ because conversion was 0.926369 commons-units per nat. **Headroom and the realised fraction are independent readings and both are required: two of the five cases had positive headroom and were shrinking.** The complaint recorded at survey named the binding term in none of the five.

The quadratic-committee price is exact. Twenty-two-household consensus carries 231 pairwise channels, 74.732 per nat of $\ln n$ against 6.213 at $n = 5$; a nested four-cluster form with a four-delegate council carries 55.5 channels, **24.03 per cent of the flat figure**.

Two theorems about the essence law are unconditional. *Result One:* $f(m)/f(n) = m^2/n^2$, so the chosen constant k cancels from every interval and does no work beyond setting a reference pitch — the golden section in this formalism is a transposition, and a transposition is the one musical operation that changes nothing structural. *Result Two:* writing $n = 2^a m$ with m odd, the pitch class is the odd part, so a set indexed $1..N$ realises exactly $\lceil N/2 \rceil$ classes — twenty elements give ten, twelve give six, and a set indexed by powers of two gives one.

The bridge is badged. Taking $n_{\text{eff}} := N_{\text{eff}} = e^F$ gives octaves above base $= 2F/\ln 2 =$

$2.885390081777927 \times F$ and 3462.4680981335122 cents per nat, both exact. The worked institute publishes 4.06 octaves with an honest **0.7827-octave band** from $u = 0.28 \pm 0.05$. The effective count used is the affine $1 + (N - 1)u$ and **not** a Hill number.

And the counter-results, which are the part of this paper a reviewer should weigh most. The essence law's three load-bearing commitments — that essence types have a frequency at all, that the law is square in the index, and that $k = (9/8)^2/\varphi^2$ — are postulates with no external support, and the constant is then shown to cancel from everything. The source's own compound rule reproduces only **37 of 44** published composite values, so seven are declared errata and every value must be recomputed rather than transcribed. The scale-freedom claim is a statement about model form tested at exactly five scales, not a physical symmetry. And the whole apparatus rests on a reflexive specification that is **not identified from observational data**, for which this paper adopts no remedy and therefore claims a modelling position rather than a measured result.

17.13.1 Declarations

Data and code availability

No external dataset was collected for this paper. All numerical results are computed from the stated parameters under the model developed in the text, and every table can be re-derived from the equations and the parameter values printed beside it. The five systems worked in Section 4 are constructed examples with declared parameters, not field observations; the institute of Section 8.14 is likewise constructed. The verification of the essence law in Section 7.4 is a recomputation of a table published in a third-party document of 2022, and the seven values declared errata in Section 7.10 are identified by that recomputation. The figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper and the model equations of Sections 3, 7 and 8.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, produced the worked systems and figures, performed the recomputation of the 2022 source table, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. All five surveyed systems and the worked institute are constructed. The consent constraint applied throughout Section 4 — that no intervention removing a participant's ability to decline is within the framework's scope, and that two candidate interventions were cut on that ground and the cuts recorded — is a design constraint on prospective work and is not a report of any study performed.

Preregistration

This paper reports no preregistered test of its own. Each of the five cases states its falsification conditions, its required baselines and its trip-wires before its intervention, which is the corpus's preregistration protocol [Curzi, 2026i] applied to a constructed case; no outcome data exist for any of them. The corpus's one preregistered confirmatory application of the storage law to a real system **failed its second falsification rule** — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared margin of 10 per cent [Curzi, 2026i] — and any applied claim resting on this paper inherits that result.

Chapter 18

Conversion Keys and Modal Design: Deriving a Room, a Cadence and a Pitch Lattice from an Essence Set Without Subliminal Signal

Abstract. Four registers this framework treats separately — surplus in nats, pitch in octaves, sensation, and the moment — are shown to be one register up to a change of logarithmic base, and the conversion constants are computed exactly: $1/\ln 2 = 1.4426950408889634$ bits per nat, $2/\ln 2 = 2.8853900817779268$ essence octaves per nat, and $2400/\ln 2 = 3462.4680981335122$ cents per nat. One bit is *two* octaves rather than one, because doubling an effective count multiplies a square-law frequency $f(n) = k\,n^2$ by four. The keys are then used to build: taking essence indices $7, 9, 11$ as axial-mode targets, $L = c/(2f)$ with $c = 343\,\mathrm{m\,s^{-1}}$ yields a room of $7.2400 \times 4.3798 \times 2.9319$ m, exactly in the ratio $1 : 121/81 : 121/49$, with $V = 92.969\,\mathrm{m^3}$ seating 12.7 at $2.5\,\mathrm{m^2}$ per person, a Sabine target of 0.5 to 0.6 s requiring 25 to $30\,\mathrm{m^2}$ sabins, and a Schroeder frequency of 146.7 Hz below the entire folded set. Folded in the time domain the same set gives a 12.55-day cadence at 0.090 of the system time constant. Twelve categories yield six pitch classes; eight proxemic zones yield four, two of which are disclosed in advance as 34.976 cents apart and therefore beating rather than forming an interval. The constant $k = (9/8)^2/\varphi^2 = 0.483425732988414$ is **defined and exact** as a number and a **postulate** the moment it is written in hertz, and the dimensional audit states that every hertz in the paper enters through that one step. Hidden or inaudible signal is prohibited absolutely, as an adopted design constraint rather than an empirical finding. The paper supplies a counter-result against its own instrument: this framework's own six storage-law parameters fail all three admissibility requirements and receive no frequencies at all.

Keywords: conversion keys; axial modes; Schroeder frequency; cadence folding; pitch class collision; consonance ranking; no subliminals; moment ceiling

Series: B. Engineering the Surplus **Paper:** CCP-018

18.1 1 Introduction

A framework that measures interaction surplus in nats, and a framework that measures musical pitch in octaves, are not obviously about the same thing. This paper shows that they are

the same thing up to a change of logarithmic base, computes the conversion constants exactly, and then does the only honest test of a unit system: it uses the units to build something and reports what the build is and is not evidence for.

The occasion is a narrow one. The corpus carries a **target-generating instrument** in the essence law `f(n) = k n^2`, which takes an ordered category scheme — a room's programme, an institution's houses, a set of published proxemic zones — and returns a small set of frequencies with exact ratio relationships. Those frequencies can then be used the way any designer uses a proportional system: to choose a dimension, a tuning target, a meeting period. **The claim being made is that the instrument chooses the target, and that is all it does.** The physics that follows is ordinary physics. A room 4.3798 m wide has an axial mode at 39.16 Hz because `f = c/(2L)` with `c = 343 m/s`, not because its category index is 9. Keeping those two sentences apart is the discipline this paper is built to enforce, and the point at which the corpus is most exposed if it fails.

Four things are established here that are not established elsewhere in the corpus.

First, the conversion keys, computed exactly. One nat is `1/ln 2 = 1.4426950408889634` bits, `2/ln 2 = 2.8853900817779268` essence octaves, and `2400/ln 2` cents. These are exact ratios of exact quantities and are quoted to the precision they support; the third is recomputed in Section 14 and found to differ from the corpus's printed value in the seventeenth significant figure.

Second, the factor of two. One bit is *two* octaves, not one, and the reason is neither information-theoretic nor musical: doubling an effective count multiplies a square-law frequency by four, and four is two octaves. Any reader importing the ordinary bits-to-octaves intuition is wrong by a factor of two, and the factor is the exponent of a postulate rather than a constant of nature.

Third, a worked build. Taking essence indices 7, 9 and 11 as axial-mode targets, `L = c/(2f)` returns a room of 7.2400 by 4.3798 by 2.9319 m, `V = 92.969` m3, seating 12.7 at 2.5 m2 per person, with a Sabine target of 0.5 to 0.6 s requiring 25 to 30 m2 sabins across 131.6 m2 of surface, and a Schroeder frequency of 146.7 Hz that sits below the entire folded set. The same set folded in the time domain yields a 12.55-day cadence at 0.090 of the system time constant. The room is buildable and its proportion matches a published heuristic to a Euclidean error of 0.031. **None of that is evidence that the tuning does anything**, and Section 13 says so at length.

Fourth, a counter-result the paper supplies against its own instrument. The corpus's own six storage-law parameters — the most obviously "important" set of categories available to it — fail all three admissibility requirements and therefore receive no frequencies at all. Set against that are two hard capacity limits: at most 25 essences fit the audible band under a shared transposition, and an essence frequency is never a physical resonance.

The material this paper analyses reaches it from a manuscript written in an initiatory register, and one part of the conversion is to say plainly what the register concealed. **A symbolic system can be analysed respectably; it cannot be performed in a submission.** Where the source asserts a mechanism it cannot support — a coil that "increases the probability of what you play through it over time" — this paper does not carry the assertion, and says why. Where the source prohibits something on ethical rather than empirical grounds — hidden or inaudible signal directed at occupants — this paper carries the prohibition as an adopted design constraint and records it in the Declarations as a constraint on prospective work rather than as a report of a study performed.

The framework's static half — the surplus functional `F(u) = ln(1 + (N−1)u)` and its affine effective count `N_eff = 1 + (N−1)u` — is presupposed and not re-derived; it is established with its axiom set in Curzi [2026], whose design axiom S3 fixes the affine form by choice rather than by derivation. The framework's founding definition, that comfort is the felt total of present costs so that the present state is a local minimum of the comfort landscape by definition rather than by luck, is presupposed from Curzi [2026]. The dynamical half — the storage law as a first-order section with pole `1 − delta`, time constant `tau = 1/delta` and the pathology at `delta = 0` — is presupposed from Curzi [2026a] and is used in Section 9.4 to decide whether an essence-

derived cadence is admissible at all.

Sections 2 and 3 state the hypotheses and the framework, including the three methodological problems — identification, scale type, and analogy discipline — that this paper is more exposed to than any other in the corpus. Sections 4 to 6 build the conversion keys and the three registers they join. Sections 7 and 8 develop the harmonic lattice and the admissibility requirements on an essence set. Section 9 is the modal design procedure and its four targets. Sections 10 and 11 collect the specifications and the twenty-eight named failure modes. Section 12 works three systems end to end, one of which is a refusal. Sections 13 to 16 discuss what follows, verify the chain, list what the account cannot support, and state the results in quotable form.

18.2 2 Hypotheses

Most of this paper is arithmetic, and arithmetic cannot be falsified by observation. The hypotheses below are the parts that can be, and each names the observation that would defeat it. Where a claim could not be defeated by any observation it is stated below the list as a definition, a postulate, or a modelling choice, and labelled as one; **no claim has been dressed as a hypothesis to fill the section**.

H1 — The room is an ordinary resonator, and the essence law contributes only the target. A rectangular enclosure built to `L = c/(2f)` at the dimensions of Section 9.3 has axial fundamentals at the frequencies the construction placed there, and its reverberation obeys Sabine's relation above the Schroeder boundary. **Predicted:** measured axial fundamentals at 23.69, 39.16 and 58.49 Hz to within the tolerance set by construction and by the departure of real surfaces from rigid; a measured `RT60` in 0.5 to 0.6 s when 25 to 30 m2 sabins are installed across 131.6 m2; a modal/statistical transition near `f_s = 146.7` Hz at `RT60 = 0.5` s. **Falsified by:** a measured axial fundamental departing from `c/(2L)` by more than the construction tolerance; or a measured `RT60` outside the error band that Sabine [1922] and Beranek [2004] report for rooms of this volume and absorption distribution.

H2 — The tuning contributes nothing beyond the choice of target. This paper predicts its own null. **Predicted:** in a preregistered comparison of an essence-tuned room against a conventionally proportioned room matched on volume, `RT60` and sightlines, measured turn-taking equality [Woolley et al., 2010] shows **no** advantage for the tuned room that survives a shuffled-input control. **Falsified by:** a surviving advantage under the protocol of Curzi [2026i] — which would be a positive result the framework does not currently claim and has no mechanism for.

H3 — A 34.976-cent separation beats rather than forming an interval. The two pitch classes 7 and 5 fold to `49/32` and `25/16`, 34.976 cents apart; at the band placement of Section 9.2 that is 379.0058 and 386.7406 Hz. **Predicted:** a beat at 7.7348 Hz, heard as roughness within a critical band rather than as a two-tone interval [Plomp and Levelt, 1965, von Helmholtz, 1885]. **Falsified by:** listeners at ordinary levels reliably reporting two separately pitched tones rather than one rough tone at that separation.

H4 — The index is reproducible where the ordering source is intrinsic, and not otherwise. Predicted: two competent analysts given a domain ordered by an intrinsic physical magnitude — Hall's proxemic zones ordered by distance [Hall, 1966] — assign identical indices; given a domain ordered by an enumeration with no intrinsic magnitude, they do not, and the resulting tunings differ by more than 100 cents. **Falsified by:** measured inter-rater agreement on an intrinsically ordered domain

falling below what the reproducibility literature would call acceptable [Open Science Collaboration, 2015]; or, in the other direction, two analysts independently reproducing an enumeration-ordered index.

H5 — Cadence admissibility is set by the pole, not by the ritual. Predicted: a system with per-period retention loss `delta` driven at forcing period `P` shows a visible sawtooth in its stored order when `P > tau/2`, `tau = 1/delta`, and smooth integration when `P <= tau/2`, independent of the size of each forcing [Curzi, 2026a]. **Falsified by:** an observed sawtooth at `P/tau = 0.09`, or smooth integration at `P/tau > 1`, in a system whose `delta` was measured on a zero-input window.

H6 — The square-law exponent is not recoverable from anything this corpus holds. Predicted: no dataset assembled in this corpus discriminates `s = 2` in `f = k n^s` from `s = 1.5` or `s = 3`, because "essence relation" has no operationalisation with an independently measurable magnitude; the exponent is **defined by the model, not estimated** [Clauset et al., 2009]. **Falsified by:** exhibiting one domain in which the relation between category types carries an independently measured magnitude, against which the postulated `m^2/n^2` spacing can be checked. That would be a genuine advance and the paper states the requirement precisely so that someone can meet it.

Four claims in this paper are not hypotheses, and are labelled here rather than left for a reader to discover.

- **The essence law `f(n) = k n^2` is a POSTULATE about essence structure.** It is not falsifiable at present, for exactly the reason H6 states. Its formal coincidence with the eigenvalue spacing of a one-dimensional infinite square well, where `E_n` is proportional to `n^2`, is a genuine mathematical correspondence and is noted as such; **coincidence of form is not evidence of physical mechanism**, and the corpus's own bibliographic audit records that asserting the connection would be a decorative analogy rather than a derivation.
- **`k = (9/8)^2 / phi^2 = 0.483425732988414` is a DEFINITION** — an exact ratio of exact quantities, computable to arbitrary precision, and quoted at fifteen figures only because it is defined. **Attaching `Hz` to it is a separate and much stronger claim**, it is a POSTULATE, and it is given its own line in the verification chain of Section 14. There is no physical, psychoacoustic or musicological precedent for combining a Pythagorean whole tone with an inverse square of the golden section to obtain a ground frequency, and none is claimed [Markowsky, 1992, Fischler, 1981, Livio, 2002].
- **The normalisation of the moment onto `[0, ln N]` is a CHOSEN MAPPING.** The ceiling result of Section 6.2 follows from the surplus range theorem plus that normalisation, and the normalisation is the only chosen step.
- **The no-subliminals constraint is an ETHICAL PROHIBITION adopted by the framework, not an empirical finding.** It is not a claim that hidden signal fails to work; efficacy is beside the point. It is a constraint on what the framework will design, and it is recorded as such in the Declarations.

18.3 3 Framework and method

18.3.1 3.1 What is presupposed

Five results are used throughout and established elsewhere. Their statuses differ and each is stated.

The surplus functional. For two unit state vectors drawn from an ensemble of `N` categories, `u = 1 − (x.y)^2` in `[0,1]` and

```
F(u) = ln(1 + (N−1) u),      N_eff = e^F = 1 + (N−1) u
```

with `F` uniquely characterised by four axioms — angle-only dependence, zero at parallel, an affine effective count, and normalisation at orthogonality — as established in Curzi [2026]. **Axiom S3, the affine effective count, is a design axiom: chosen, not derived.** The range result `0 <= F <= ln N` is what the ceiling argument of Section 6.2 consumes. *Status:* **PROVEN given the axioms.**

The storage law. `Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t` is a first-order discrete section with pole `1 − delta`, DC gain `1/delta` and time constant `tau = 1/delta` periods, and `delta = 0` is a pathology rather than an ideal because the map ceases to contract and no fixed point exists [Curzi, 2026a]. Section 9.4 consumes only the time constant. *Status:* **PROVEN as a statement about the equation.**

Sabine's relation. `RT60 = 0.161 V / A` in metric units, `V` in cubic metres, `A` in square-metre sabins [Sabine, 1922]. *Status:* **CITED — real acoustics, with a known error band and known conditions of validity.**

Rectangular room modes. For a rigid rectangular enclosure, `f(p,q,r) = (c/2) sqrt((p/Lx)^2 + (q/Ly)^2 + (r/Lz)^2)`, with `c = 343 m/s` at 20 C in dry air [Rossing et al., 2002, Beranek, 2004]. The axial case `q = r = 0` reduces to `f = c/(2Lx)`. *Status:* **CITED — real acoustics.**

The essence law. `f(n) = k n^2` with `k = (9/8)^2/phi^2`, `n` the ordinal position of a category within its domain's ordered set. *Status:* **POSTULATE**, with `k` a chosen reference constant and the index an input the law does not supply and cannot check.

18.3.2 3.2 The four registers, and why a conversion exists at all

Four quantities in this framework are logarithmic. Their statuses differ, and mixing statuses in one table is exactly what allows a postulate to borrow the authority of a theorem.

#	Register	Form	Status
1	Surplus	`F = ln N_eff = ln(1 + (N−1)u)`	**PROVEN** given axioms S1–S4, one of which (S3) is chosen
2	Pitch	`p = log2(f/f_ref)`	**DEFINITION**, universal in music [von Helmholtz, 1885]
3	Sensation	`S = k_F ln(I/I_0)`	**MEASURED LAW** with known failure regions [Fechner, 1966]
4	Moment	`ln M = ln D − ln P + ln E`	**POSTULATE** — this framework's own formula

Every conversion in this paper carries the status of the quantity it converts, and a conversion never upgrades a status. A postulate converted into octaves is a postulate in octaves. That rule is enforced throughout and is restated as specification CK-3 in Section 10.2.

The mechanism is one identity. All four registers are logarithms; logarithms in different bases differ by a constant factor and by nothing else, so `log_{b2}(x) = log_{b1}(x) (ln b1 / ln b2)` [Cover and Thomas, 2006]. **Because every register is a logarithm, every conversion between registers is a multiplication by a constant, and no information is lost, gained, or distorted in the crossing.** There is no approximation in the unit arithmetic of this paper; the only approximations are in the measured inputs.

18.3.3 3.3 Analogy discipline: which uses are derivations and which are analogies

This paper is the corpus's most exposed to a single failure — letting an analogy borrow the authority of a derivation — because it puts both in the same paragraph. The three-way classification is applied here once and then invoked by name at each first use.

Category 1, a derivation. The physical result is used because the system genuinely satisfies its premises. **Sabine's relation, the Schroeder boundary, the axial-mode formula, the open-open tube, the Helmholtz resonator and the stretched string are all category 1 in this paper.** A room really is a physical resonator; the premises of the modal formula are geometric and the room satisfies them. Nothing in these uses depends on any claim of this framework, and if the framework were withdrawn entirely the acoustics would stand unchanged [Rossing et al., 2002, Beranek, 2004].

Category 2, a shared mathematical form. The same equation governs both objects, with no claim that the underlying physics applies. **The storage law as a first-order low-pass section is category 2**, and the cadence argument of Section 9.4 uses it that way: a leaky integrator is a leaky integrator whether it holds charge or stored order [Curzi, 2026a]. **Octave folding applied to a period rather than a frequency is also category 2** — the reciprocal of a frequency is a period, and doubling is doubling.

Category 3, an analogy with no formal warrant. A suggestive correspondence, legitimate as motivation and never as evidence. **The mapping of an essence index to a frequency in hertz is category 3, and so is every reading of meaning off the resulting pitch relations.** Octave equivalence is grounded as a perceptual structure for pitch [Shepard, 1982, Demany et al., 2013], and the corpus's own bibliographic audit records that **nothing supports carrying it over to a quantity that is not a pitch**. The operation is well defined; its interpretation is unsupported. Independent evidence tightens rather than loosens that constraint: octave equivalence is itself weaker and more training-dependent than usually assumed [Hoeschele et al., 2012, Burns, 1999].

The failure mode is a category-3 use borrowing the authority of a category-1 use, and it has a name here: FM-2, the resonance claim, built. The room really does resonate. The category does not. Those are two sentences and this paper keeps them apart, at the cost of some repetition.

18.3.4 3.4 Identification: what no amount of building can recover

The framework's central formal claim elsewhere in the corpus is the reflexive loop `T(x) = W(x, M(x))`: the group influences the member and the members constitute the group. **That is a linear-in-means reflexive specification, and its endogenous social effect is not identified from observational data** [Manski, 1993]. No amount of fit, held out or otherwise, recovers the parameter, and influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011].

This paper invokes the loop in exactly one place — the claim that a room's acoustics shape turn-taking and that turn-taking constitutes the group's collective performance [Woolley et al., 2010, Curzi, 2026e] — and it **adopts no remedy**. The known options are randomisation, an instrument, partial-population variation, or stipulating the reference group and defending the stipulation [Brock and Durlauf, 2001]; the last is the most realistic route and is not taken here. **The claim that room acoustics act on group performance through turn-taking is therefore a modelling position and not an identified empirical result**, and Section 15 records it again.

There is a second, sharper form of the same problem specific to this paper. The framework asserts that disclosure changes behaviour — that is the entire content of Section 3.6 — and then proposes to measure the tuned room's effect on occupants who have read a disclosure notice naming the tuning. **A disclosed intervention cannot be blinded to its recipients**, so the room comparison of H2 is unblindable by construction and its expectancy component is

unremovable [Rosenthal and Jacobson, 1968, McCambridge et al., 2014]. That is a real limit on what H2 can ever return and it is stated as one.

18.3.5 3.5 Scale type: the promotion problem, and the strongest objection to the framework

The essence law takes an ORDINAL and returns a RATIO-SCALE quantity. The index is a rank. The frequency is a physical magnitude with a true zero on which multiplication and division are meaningful. **The law therefore performs a scale promotion — it manufactures ratio-scale structure out of ordinal input — and that is not a permissible operation on data alone** [Stevens, 1946].

It is permissible as a **model**, and that is what it is. The square law is a postulate that *asserts* a specific quantitative spacing between ranks. The assertion may be right or wrong and it is the thing that would have to be tested. **What the text may never do is present the resulting frequencies as though the ratio-scale structure were measured. It was postulated.** The practical form of the rule: *the ordering is data; the spacing is the model*, and the intervals are only as meaningful as the spacing postulate.

This is where Michell [1997] applies, and it is the strongest objection to the whole framework. Psychology, Michell argues, systematically assumes that its attributes have quantitative structure without ever testing whether they do — and the test exists. **Conjoint measurement is the formal test for additive quantitative structure from ordinal data on two or more factors [Luce and Tukey, 1964, Krantz et al., 1971], and it has not been run for any construct in this paper.** Neither has a Rasch-model demonstration of specific objectivity [Rasch, 1960]. The paper's own failure-mode catalogue already contains the objection under the name FM-21, THE ARBITRARY ORDINAL; Section 15 states it in the scale-type vocabulary, where it belongs. **The ratio-scale treatment of essence frequency is a stipulation and is conceded to be one.**

A second scale-type warning attaches to the folded classes. **The odd part of an index is a NOMINAL label**, a class identifier and not a magnitude; the class's octave position in cents is interval-scaled and may be subtracted; the class count `ceil(N/2)` is a ratio-scaled count. Three scale types live in one construction and must not be mixed. Arithmetic on class numbers is not licensed.

18.3.6 3.6 The no-subliminals constraint, as a design prohibition

No subliminal content of any kind. Nothing masked. Nothing below the threshold of hearing. Nothing concealed from the listener. This is stated first because it is where the temptation arises, and it is stated as an **ethical prohibition adopted by the framework rather than as an empirical claim about efficacy**. Every technique in this paper operates entirely in audible, disclosed tone. Concretely, and as a specification rather than a sentiment:

1. **Every tone present in a space is audible to a person of ordinary hearing standing in it.** A tone below the threshold of audibility at any listening position is either raised to audibility or removed. There is no third option.
2. **Every tone is disclosed.** A written statement — posted at the entrance, in the room's documentation, or both — lists what is sounding, at what frequency, at what level, and why.
3. **No masking, no carrier, no embedding.** No content is placed under, inside or behind another signal: no amplitude-modulated speech, no time-compressed content, no spectral hiding, no content below 20 Hz or above 20 kHz intended to act on a listener.
4. **The room is leavable and the tone is stoppable.** There is an off switch, its location is disclosed, and anyone in the room may use it.
5. **Where the tuning is structural rather than sounded — a room dimension, a meeting cadence — the derivation is published**, so that anyone can disagree with it.

The reasoning is not squeamishness and not doubt about efficacy. **A method that works by removing a person's ability to decline is outside this framework's scope regardless of whether it works, and a message a listener cannot hear is a message the listener cannot decline.** The framework's boundary doctrine reaches the same place from the other direction: category innocence is explicitly refused as a defence of acquisition, and no harm may be moved through a discovered gap in a scheme [Curzi, 2026] — a bound this paper inherits rather than argues.

Clause 5 has an independent warrant that has nothing to do with ethics. A system whose security rests on the secrecy of its design rather than of its key is a weak system [Kerckhoffs, 1883], and a proportional system whose derivation is unpublished cannot be checked, disputed, or reproduced. **Publicity is the operative remedy** [Brandeis, 1914]. And it is a partial remedy at best: the corpus's own structural work records that disclosure did not prevent the 2022 immobilisation it analyses [Curzi, 2026], which is the reason disclosure appears here as a necessary condition and never as a sufficient one.

Two recommendations in the source material this paper draws on are refused and the refusals are recorded. The first is the embedding of subliminal content in a rendered track, excluded by clauses 1 and 3 above. The second is the use of a 12- or 24-winding Rodin coil as an amplifier, on the stated grounds that it "influences temporal 4th dimensional space by increasing the probability of what you play through the coil over time". **No verified engineering basis exists for the claimed mechanism, no citation in this corpus's bibliography supports it, and this paper does not present it as an amplifier.** Where an area of effect is required, the correct instrument is a loudspeaker with a published frequency response and a measured sound pressure level, with both disclosed.

18.3.7 3.7 Standing honesty clauses

Four clauses bind everywhere this paper's material is used. They are labelled HC-1 to HC-4 to keep them distinct from the hypotheses of Section 2.

HC-1 — Essence, not resonance. The output of the essence law is an essence frequency. **No claim is made that any material, room, institution or person physically resonates at these values.** A room whose axial mode is placed at 39.157484 Hz genuinely resonates at 39.157484 Hz — that is ordinary acoustics, and it is true because the room is 4.3798 m wide. It is *not* true because the room's category index is 9. **The essence law chose the target; physics did the resonating.**

HC-2 — `k` is chosen. Badged at every appearance, in every register, at every precision.

HC-3 — The ordering must be defensible. The whole method is downstream of the admissibility test of Section 8.1, and a set that fails it gets no frequencies.

HC-4 — Recompute; never copy a published value. This clause matters more here than anywhere else in the corpus, because a built object freezes an arithmetic error in timber and plaster and nobody re-checks it afterwards. Every figure in this paper was recomputed from the constants of Section 4.2; Section 14.5 reports the four printed values that did not reproduce. Two transcription faults travel with any citation of the 2022 source document from which the essence law is drawn: its prose instruction `1+5^.5/2` is missing its parentheses and must be read as `(1 + sqrt(5))/2`, and its spreadsheet stored `phi` to fifteen significant figures. **Cut a room, a bar or a pipe to a recomputed value and cite the errata beside it.**

And one clause specific to building. *The fact that a room is tuned by this method is not evidence that the method works.* A room built to these proportions may be a good room for entirely conventional reasons — its ratio happened to land near Volkmann's, its `RT60` happened to land in the speech range, its size happened to suit twelve people. A claim that the tuning did something requires a preregistered prediction and the four declared baselines including the shuffled-input control [Curzi, 2026i, Nosek et al., 2018]. **This paper claims only that the method produces a defensible, reproducible, documented choice of target.** Carry the correction with the requirement: the shuffled-input control has three measured failure modes of its own — it is anti-conservative as ordinarily run, it cannot see repackaging, and it cannot see a shared

slow driver — so a passed shuffle is not on its own evidence that a tuned room did anything [Simmons et al., 2011, Gelman and Loken, 2013]. **Design controls, not more surrogates.**

18.4 4 The conversion keys

18.4.1 4.1 The five constants

```
ln 2       = 0.6931471805599453        nats per bit          DEFINITION
1 / ln 2   = 1.4426950408889634        bits per nat          DEFINITION
2 / ln 2   = 2.8853900817779268        essence octaves/nat   DEFINITION given the square law
2400/ln 2  = 3462.4680981335122        cents per nat         DEFINITION given the square law
k          = 0.483425732988414         the ground constant   DEFINITION; in Hz, POSTULATE
```

The first four are exact ratios of exact quantities and may be recomputed to any further precision; they are quoted at the precision a reader is likely to want and no significance attaches to their last digit. **The fourth differs in its seventeenth significant figure from the value the corpus has previously printed**, and Section 14.5 gives the recomputation.

The fifth is not a unit conversion. It is the reference pitch, and it is **a chosen constant, not a measured one** (HC-2):

```
k = (9/8)^2 / phi^2 = (81/64)/(phi + 1)

phi   = 1.618033988749895
phi^2 = 2.618033988749895 = phi + 1                    (Euclid VI.3)
(9/8)^2 / phi^2   = 0.483425732988414                  DEFINED, exact
(81/64) / (phi+1) = 0.483425732988414                  identical
```

The two forms are identical because `phi^2 = phi + 1`, the defining property of the division in extreme and mean ratio [Euclid, 1956]; more generally `phi^m = F(m) phi + F(m−1)` with `F` the Fibonacci sequence, which is what makes every power of `phi` reducible to a rational combination [Vajda, 1989].

The `9/8` and the `phi` are both real objects with real histories, and their product in hertz has no precedent in any literature. The `9/8` is the Pythagorean whole tone, a specific and well-documented interval [Barbour, 1951, Doty, 2002]. The golden section is a specific and well-documented number. **What has no support of any kind is combining a tuning ratio with an inverse square of the golden section to obtain a ground frequency**, and this paper claims none. The golden-ratio literature is unusually well supplied with cautionary work, and it is cited here deliberately: Markowsky [1992] demolishes the standard claims that `phi` appears in the Parthenon, the Great Pyramid, the human body and preferred rectangles; Fischler [1981] shows how easily `phi` can be "found" in arbitrary data by choosing measurement points freely; Green [1995] and Höge [1997] review a century of failed attempts to replicate a golden-rectangle aesthetic preference; Livio [2002] separates the mathematics from the folklore. **The distinction that matters is between CHOOSING `phi` as a constant and FINDING it in nature. This framework does the first and does not do the second.**

18.4.2 4.2 The conversion key table

Read a row as "one of this unit equals", and a column as "measured in this unit".

1 of	nats	bits	octaves	cents
nat	1	1.442695	2.885390	3462.468098
bit	0.693147	1	**2 exactly**	**2400 exactly**

1 of	nats	bits	octaves	cents
octave	0.346574	**0.5 exactly**	1	**1200 exactly**
cent	0.000288811	0.000416667	0.000833333	1

Three entries are exact and the reason is not decorative. **The 1200 is definitional**: the cent has been one twelve-hundredth of an octave since Ellis introduced it in his translation of Helmholtz [von Helmholtz, 1885]. **The 2 and the 0.5 come from somewhere else entirely**, and Section 4.3 derives them.

The derived conversion laws:

```
SURPLUS -> PITCH        octaves above ground = 2F / ln 2 = 2.885390 F   (F in nats)
SURPLUS -> FREQUENCY    f_cross = k e^(2F) = k N_eff^2
ESSENCE -> FREQUENCY    f(n)    = k n^2                  n the essence index
INTENSITY -> SENSATION  S       = k_F ln(I / I_0)        Fechner
INTENSITY -> MAGNITUDE  log psi = log k_S + a log I      Stevens; a IS the key
MOMENT -> SURPLUS       M normalised onto [0, ln N] serves as S_t in the storage law
                        [NORMALISATION CHOSEN, not derived]
DISTORTION              Delta = |ln r|,  r = D/P         reported in any register
```

18.4.3 4.3 Why one bit is two octaves and not one

This is the single non-obvious constant in the table and it is derived rather than asserted, because a reader who assumes the ordinary information-theoretic correspondence [Shannon, 1948, Hartley, 1928] will be wrong by a factor of two.

One bit of surplus means the effective count doubled: `N_eff -> 2 N_eff`. The essence law maps index to frequency by a **square**:

```
f(n)      = k n^2
f(2n)/f(n) = (2n)^2 / n^2 = 4       =  two octaves, exactly
```

Doubling the effective count quadruples the essence frequency, and quadrupling is exactly two octaves. That, and only that, is why one bit is two octaves here. The factor two in `octaves = 2F/ln 2` is the exponent of the essence law and nothing else. Verified across the range:

```
N_eff  ->  2 N_eff        f ratio      octaves
  1    ->     2            4.0          2.0
  3    ->     6            4.0          2.0
 12    ->    24            4.0          2.0
```

The consequence for the honest reader is a caution. If the exponent were some other value `s`, so that `f = k n^s`, the key would be `s/ln 2` octaves per nat:

exponent `s`	octaves per nat	cents per nat
0.5	0.721348	865.617
1	1.442695	1731.234
1.5	2.164043	2596.851
2 (this framework)	**2.885390**	**3462.468**
2.5	3.606738	4328.085
3	4.328085	5193.702

The surplus-to-pitch key is therefore contingent on the square law, and the square law is a postulate about essence structure rather than a measured fact. State this whenever 2.885390

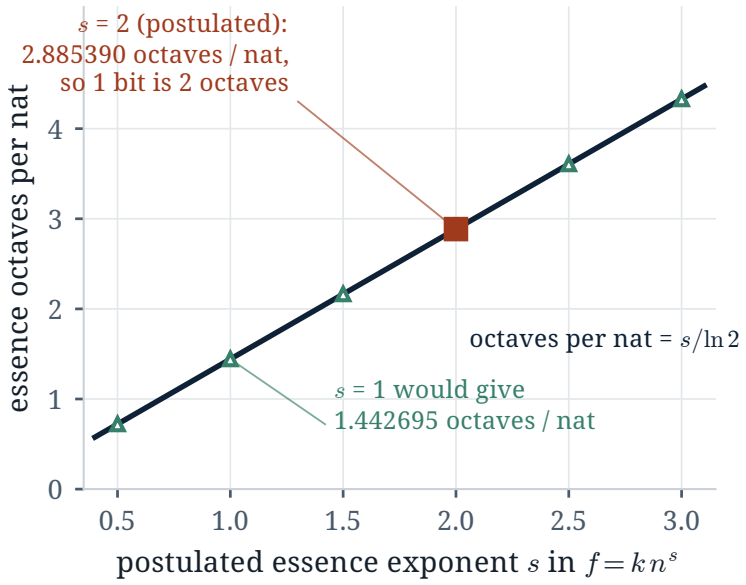


Figure 18.1: **One bit is two octaves, and the factor of two is a postulate rather than a constant of nature.** The conversion key is `s/ln 2` octaves per nat, where `s` is the exponent of the essence law; at the postulated `s = 2` it is 2.885390, and at `s = 1`, the exponent an information-theoretic reader would assume, it would be 1.442695. The claim is verified at `N_eff = 1, 3, 12`, where `f(2n)/f(n) = 4.0` and `log2 4 = 2.0` in every case. Exact arithmetic; the exponent is defined by the model and was never estimated. Generated by `figs/src/CCP-018_fig01.py` from the tables of Section 4.3.

is used. The constant is exact *given* the postulate and not exact independently of it — the difference between "2.885390 is the correct number" and "2.885390 is a correct number about the world". Only the first is claimed.

The exponent is defined by the model, not estimated, and that distinction is not a technicality. The literature on fitting power-law exponents exists because an exponent read off a log-log plot by eye, or fitted without a goodness-of-fit test against alternatives, is routinely wrong [Clauset et al., 2009]. **This paper runs no such procedure and needs none, because `s = 2` was never estimated from data.** It was chosen. A reader who mistakes a defined exponent for an estimated one will look for confidence intervals that do not exist and cannot exist, and the sensitivity table above is what replaces them.

18.4.4 4.4 Deriving each constant

Nats to bits. A nat is the logarithm base e, a bit base 2, so one nat is `1/ln 2` bits [Cover and Thomas, 2006]:

```
1 / ln 2 = 1 / 0.6931471805599453 = 1.4426950408889634
```

To six places, 1.442695. Bits are the register in which an engineering reader has the strongest intuitions: `ln 12 = 2.484907` nats is `3.584963` bits, so a twelve-category system, fully orthogonal, is worth three and a half bits of surplus per interaction. **That is a small number and it should feel small.**

Nats to octaves. By Section 4.3 one unit of `F` multiplies frequency by `e^2`, and the octaves in a ratio `R` are `log2 R`, so

```
octaves per nat = log2(e^2) = 2 / ln 2 = 2.8853900817779268
```

Nats to cents. A cent is one twelve-hundredth of an octave by definition, so

```
cents per nat = 1200 (2 / ln 2) = 2400 / ln 2 = 3462.4680981335122
```

Octaves to nats. The reciprocal, `ln 2 / 2 = 0.34657359027997264`.

Cents to nats. `ln 2 / 2400 = 0.00028881132523331055`, the resolution constant. One cent corresponds to about 289 millionths of a nat of surplus. **The pitch register has far finer graduation than the surplus register has meaningful precision**, which is a warning and not an invitation: see FM-17, the cents illusion.

18.4.5 4.5 The conversion in use

The working table converts an effective count into all four registers plus the essence frequency. It is computed, not transcribed.

`N_eff`	`F` (nats)	bits	octaves	cents	`f = k n^2` (Hz)
1	0.000000	0.000000	0.000000	0.000	0.483426
2	0.693147	1.000000	2.000000	2400.000	1.933703
3	1.098612	1.584963	3.169925	3803.910	4.350832
4	1.386294	2.000000	4.000000	4800.000	7.734812
5	1.609438	2.321928	4.643856	5572.627	12.085643
6	1.791759	2.584963	5.169925	6203.910	17.403326
7	1.945910	2.807355	5.614710	6737.652	23.687861
8	2.079442	3.000000	6.000000	7200.000	30.939247
9	2.197225	3.169925	6.339850	7607.820	39.157484
10	2.302585	3.321928	6.643856	7972.627	48.342573
12	2.484907	3.584963	7.169925	8603.910	69.613306
16	2.772589	4.000000	8.000000	9600.000	123.756988
20	2.995732	4.321928	8.643856	10372.627	193.370293
24	3.178054	4.584963	9.169925	11003.910	278.453222

Two features are worth dwelling on. **The powers of two are exact in three registers at once**: `N_eff = 2, 4, 8, 16` give exactly 1, 2, 3, 4 bits, exactly 2, 4, 6, 8 octaves, and exactly 2400, 4800, 7200, 9600 cents. Nothing is rounded.

And the frequency column is brutal. Going from an effective count of 2 to 24 — a twelvefold increase in genuinely different views — moves the surplus from 0.693 to 3.178 nats, a factor of 4.6, but moves the essence frequency from 1.93 Hz to 278.45 Hz, a factor of 144. **The pitch register amplifies exactly the differences the surplus register compresses, which is why it is useful as a display register and dangerous as an inference register.** Use frequency to hear a system, not to argue about it.

A fifth key converts surplus into time. The four registers above are all instantaneous. The Gabor time-bandwidth limit `dt df >= 1/(4 pi)` [Gabor, 1946], together with this paper's own relation `d(ln f) = 2 dF`, gives

```
SURPLUS DIFFERENCE -> TIME     cycles required >= 1 / (8 pi dF),   1/(8 pi) = 0.039789
```

Resolving a difference `dF` in surplus costs at least `1/(8 pi dF)` cycles of the rendered tone, and the constant is exact. Two properties make it a key rather than a rule of thumb. It is **independent of the transposition**, because it constrains a ratio and octave transposition multiplies numerator and denominator alike — the same cycle count serves at 30 Hz and at 19 kHz. And it is **denominated in cycles, not seconds**; converting to wall-clock requires the frequency, which is exactly the dependence that should stay visible.

dF (nats)	cycles required	at N_eff = 12 (69.613 Hz)	at N_eff = 24 (278.453 Hz)
0.10	0.398	5.72 ms	1.43 ms
0.05	0.796	11.43 ms	2.86 ms
0.01	3.979	57.16 ms	14.29 ms
0.005	7.958	114.31 ms	28.58 ms
0.001	39.789	571.57 ms	142.89 ms

Two boundary conditions belong with the table. **The frequency column has a short working range.** Under a single integer octave transposition into 20 Hz to 20 kHz the ladder fits only at `j = 6`, and only up to `N_eff = 25`: `k (25)^2 2^6 = 19337.0` Hz, while `N_eff = 26` lands at 20914.9 Hz and leaves the band. **Above an effective count of 25 the pitch register cannot be rendered in one piece and requires split transposition**, which is a declared design choice and not a formatting detail. And **a rectangular gate is disqualified as an envelope**: its `sinc^2` tails fall as `1/f^2`, so the spectral second-moment integral diverges, the tone has no defined bandwidth, and the key above does not apply to it [Gabor, 1946]. Window the envelope; Section 9.6 gives the default.

18.4.6 4.6 The consistency check at the orthogonal limit

The identification of the effective count with a cross-disciplinary essence index has an arithmetical consequence which can be checked and which, had it failed, would break the bridge:

```
within a discipline:      f = k n^2          n an integer index
across disciplines:       f = k e^(2F)       F = ln N_eff

at u = 1 (fully orthogonal), N_eff = N, so F = ln N, so
k e^(2 ln N) = k N^2                         the two must coincide
```

Verified to machine precision:

n	k e^(2 ln n)	k n^2	difference
1	0.48342573298841	0.48342573298841	0
2	1.93370293195366	1.93370293195366	0
3	4.35083159689573	4.35083159689573	8.9e-16
7	23.68786091643230	23.68786091643230	3.6e-15
12	69.61330555033166	69.61330555033166	0
24	278.45322220132670	278.45322220132664	5.7e-14

Residuals are floating-point rounding at the sixteenth significant figure, not model error. **The two frameworks agree exactly at the orthogonal limit, which means the essence index of a fully non-parallel N-category system is N itself.**

Within a discipline the index is an integer and the lattice is discrete. Across disciplines the index is the effective count, which is continuous, and the pitch is continuous with it. They meet only at `u = 1`. Worked, for a twelve-category system at various levels of non-parallelism:

u	N_eff	F (nats)	octaves	f (Hz)
0.00	1.000	0.00000	0.0000	0.4834
0.25	3.750	1.32176	3.8138	6.7982

u	N_eff	F (nats)	octaves	f (Hz)
0.50	6.500	1.87180	5.4009	20.4247
0.75	9.250	2.22462	6.4189	41.3631
1.00	12.000	2.48491	7.1699	69.6133

Read the first row. **Twelve parallel categories are the ground essence: 0.483426 Hz, zero octaves, zero surplus.** Pseudo-diversity has a pitch, and its pitch is the bottom of the instrument.

Which effective count is being used here matters, and the answer is the weak one. `N_eff = 1 + (N−1)u` is an **affine interpolation between 1 and `N`** and is this framework's own construction, adopted by design axiom S3 [Curzi, 2026]. **It is not a Hill number of any order.** `1/HHI` is the Hill number of order two and `exp(H)` the Hill number of order one; both are fully grounded [Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979], and neither is what appears in `f_cross = k N_eff^2`. A reviewer who reaches for the Hill-number literature as support will find it does not apply. The contrast is available inside the corpus: the concentration instrument that *is* a Hill number of order two returns 2.087683 effective perimeters against 2.941176 on a second axis and a break-even of 4.172734 in a worked structural case [Curzi, 2026], and those numbers are not commensurable with the affine effective counts used here.

18.4.7 4.7 What the keys do not license

A conversion table is an invitation to one class of error, and the class deserves a name. **The keys convert measures. They do not transport warrant.**

- A **lower bound** converted is a lower bound, not an estimate. "3.96 octaves" means "at least 3.96 octaves" and is written that way.
- A **postulate** converted is a postulate. The moment equation in cents is still a postulate.
- A **procedure-bound measurement** converted is procedure-bound. A loudness exponent of 0.67 written as 804 cents per physical octave is still 0.67 by that method, range and instruction.
- A surplus converted into a **frequency** is an essence frequency. **Nothing is vibrating.** This is the error the material most invites and the one the standing clauses most firmly forbid (HC-1).
- Conversion **adds no precision**. A surplus known to two decimals is not known to four because cents has finer graduation (FM-17).

18.5 5 The sensation register

18.5.1 5.1 Weber, Fechner, and the warrant the argument actually needs

The third register settles the question a sceptical reader brings to any framework that measures a subjective quantity: **can a subjective quantity be measured at all?**

Weber [1834] established that the just-noticeable difference in a stimulus is a constant *ratio* of the stimulus rather than a constant amount, `dI/I = c`, with `c` the Weber fraction for that modality and range [Weber, 1978]. Fechner [1860] integrated it: if each just-noticeable difference is one subjective unit and each costs a fixed proportional increase, then counting units from a reference gives `S = k_F ln(I/I_0)` [Fechner, 1966]. The number of just-noticeable differences between a reference and a stimulus is `n_JND = ln(I/I_0) / ln(1 + c)`, **so a JND is itself a logarithmic unit and joins the key table.**

The following is **parametric, not measured**. It asserts no value of `c` for any modality, because Weber fractions are modality-, range- and method-dependent and this paper reports no figure it has not computed or cited.

c	nats/JND	bits/JND	octaves/JND	cents/JND	JNDs across x10
0.01	0.009950	0.014355	0.028711	34.453	231.4
0.02	0.019803	0.028569	0.057138	68.566	116.3
0.05	0.048790	0.070389	0.140779	168.934	47.2
0.10	0.095310	0.137504	0.275007	330.008	24.2
0.15	0.139762	0.201634	0.403268	483.921	16.5
0.20	0.182322	0.263034	0.526069	631.283	12.6
0.30	0.262364	0.378512	0.757023	908.428	8.8

A modality with a Weber fraction of 0.10 has a JND worth 0.0953 nats. Ten JNDs is 0.953 nats, less than the surplus ceiling of a three-category system. That is the kind of statement the conversion keys make available and that is impossible without them.

Honest limits, stated with the law: Fechner's relation fails near absolute threshold and near the upper end of a modality's range, where the Weber fraction is not constant. **The logarithmic form is an excellent middle-range description, not a universal law of sensation**, and Stevens' power law is the competitor and in many modalities the better fit [Stevens, 1957, Krueger, 1989].

There are two independent warrants for this framework's logarithm, and the argument needs only the weaker of the two assumptions. The first is Fechnerian: sensation proportional to the logarithm of a dimensionless stimulus ratio measured from a floor. **That warrant requires the perceptual story to be true, and the perceptual story is contested.** The second is stronger precisely because it assumes nothing psychological: Peters and Gell-Mann [2016] shows that the time-average growth rate of a multiplicatively evolving quantity is the expectation of its *logarithm*, so a logarithmic evaluation follows from the dynamics alone. **Where a stock compounds rather than adds, the logarithm is the ergodic-average object, and no claim about perception is needed.** The framework's surplus is such a stock. **Cite both; the argument needs Peters, not Fechner.** Fechner is cited here because the *sensation register* is genuinely Fechnerian and is one of the four things being converted — not because the framework's own logarithm depends on him.

And the underdetermination has to be conceded. The surplus functional's form is not fixed by the evidence: Latané [1981] proposes a sub-linear power law and Briscoe et al. [2006] an `n log n` growth, and **both fit the qualitative facts as well as `ln(1 + (N−1)u)`.** What discriminates them is an axiomatic characterisation rather than a measurement [Curzi, 2026], and an axiom set that includes a chosen axiom does not discriminate anything empirically. **Nothing assembled here selects the logarithm over the alternatives, and the paper says so rather than leaving it to a reviewer.**

18.5.2 5.2 Stevens' exponents are measured conversion keys

Stevens [1957] proposed the power law `psi = k_S I^a` and, more importantly, the method [Stevens, 1975]. Taking logarithms, `log psi = log k_S + a log I`.

In log-log coordinates the power law is a straight line and its slope is `a`. The exponent is therefore a conversion key in exactly this paper's sense: the number of perceptual log-units bought by one physical log-unit. It is not a fitted nuisance parameter. It is the constant of the change of units between the physical register and the perceptual one, and **it is obtained by measurement rather than by definition, which is what distinguishes it from every other constant in Section 4.1.**

There is a way of saying this that makes it immediately usable:

```
perceptual octaves per physical octave = a       exactly
```

An octave is a doubling in either register, and doubling `I` multiplies `psi` by `2^a`, which is `a` octaves. **The exponent *is* the octave key.** Published, replicated exponents, nominal and method-dependent:

Modality	a	world x2 gives psi x	world x10 gives psi x	psi x2 needs world x	cents per physical octave
Loudness (1 kHz)	0.67	1.5911	4.6774	2.8138	804.0
Brightness	0.33	1.2570	2.1380	8.1698	396.0
Visual length	1.00	2.0000	10.0000	2.0000	1200.0
Taste (salt)	1.30	2.4623	19.9526	1.7044	1560.0
Heaviness (lifted)	1.45	2.7321	28.1838	1.6129	1740.0
Warmth on arm	1.60	3.0314	39.8107	1.5422	1920.0
Electric shock	3.50	11.3137	3162.2777	1.2190	4200.0

`a < 1` is compressive, `a = 1` linear, `a > 1` expansive. **Most senses compress. Pain expands.** To double perceived loudness the physical intensity must roughly triple (2.81x), but to double perceived shock a 22 per cent increase suffices. **Comfort is expensive to increase and cheap to lose, and the exponents say so in numbers.**

The most consequential row is brightness: doubling perceived brightness requires **8.17 times** the physical light. **A lighting specification written in lumens and read by an occupant in perceived brightness is a specification in two registers, mismatched by a factor of eight.**

Cross-modality matching. Because each exponent is a key to the same physical register, the key between two modalities is the ratio of their exponents: if a subject adjusts modality *i* to match modality *j*, the predicted log-log slope is `a_j / a_i`. This is the strongest test of the power law, because the matching exponent is predicted in advance from two independently measured exponents and then checked.

matched	Loudness	Brightness	Length	Salt	Heaviness	Warmth	Shock
Loudness	1.0000	2.0303	0.6700	0.5154	0.4621	0.4188	0.1914
Brightness	0.4925	1.0000	0.3300	0.2538	0.2276	0.2062	0.0943
Length	1.4925	3.0303	1.0000	0.7692	0.6897	0.6250	0.2857
Salt	1.9403	3.9394	1.3000	1.0000	0.8966	0.8125	0.3714
Heaviness	2.1642	4.3939	1.4500	1.1154	1.0000	0.9062	0.4143
Warmth	2.3881	4.8485	1.6000	1.2308	1.1034	1.0000	0.4571
Shock	5.2239	10.6061	3.5000	2.6923	2.4138	2.1875	1.0000

The table is computed from the nominal exponents above; **it is a set of predictions, not observations**, and its honest use is as a preregistration target [Nosek et al., 2018]. **A framework that predicts a matching slope of 2.0303 for loudness-to-brightness and then measures 2.03 has earned something. A framework that fits the slope after the fact has earned nothing** [Meehl, 1967, Platt, 1964].

Honest limits, restated: Stevens' exponents vary with method — estimation and production give systematically different slopes, the regression effect [Poulton, 1968, 1979] — with stimulus range [Teghtsoonian, 1971], and with instruction. Whether magnitude estimation delivers a ratio scale at all is itself contested [Luce, 2002, Narens, 1996]. **Report them as measured constants of a procedure, not constants of nature.**

18.6 6 The moment register and the forced ceiling

18.6.1 6.1 The moment equation in log form

The moment equation is this framework's own postulate:

```
M = ( D / P ) . E        D perspective, P perception, E environment, M moment
```

Taking logarithms puts it in the same register as everything else here:

```
ln M = ln D − ln P + ln E
```

In the log register the moment equation is additive, and that is its most useful form. A multiplicative model with a division becomes a sum of three terms, each reportable, budgetable and convertible independently.

The calibration ratio and its distortion. With `r = D/P`, distortion is `Delta = |ln r|`, a log-ratio and therefore expressible in any register in the key table:

`r = D/P`	state	`Delta` (nats)	bits	octaves	cents
0.25	severely under-framed	1.386294	2.00000	4.00000	4800.00
0.50	under-framed	0.693147	1.00000	2.00000	2400.00
0.80	mildly under-framed	0.223144	0.32193	0.64386	772.63
1.00	**calibrated**	**0.000000**	0.00000	0.00000	0.00
1.25	mildly over-framed	0.223144	0.32193	0.64386	772.63
1.80	over-framed	0.587787	0.84800	1.69599	2035.19
2.00	over-framed	0.693147	1.00000	2.00000	2400.00
4.00	severely over-framed	1.386294	2.00000	4.00000	4800.00

Note the symmetry, which is the reason for the logarithm: `r = 0.5` and `r = 2.0` are equally distorted, 2400 cents in either direction. **A raw ratio makes over-framing look worse than under-framing because it is unbounded above and bounded below by zero; the log-ratio makes them commensurate, and they should be commensurate, because both are the same failure of match.**

The moment as storage input. The join is that `M`, normalised onto `[0, ln N]`, serves as `S_t` in the storage law. **The normalisation is a chosen mapping, badged here as a design decision** and not derived from anything.

The scale-type gate on this equation fails and the failure is printed. Division requires a ratio scale with a true zero for `D` and `P`, and neither is instrumented on one. With only ordinal data the formula is a heuristic ordering rather than an equation and must be demoted in print [Michell, 1997, Stevens, 1946]. The identifiability gate fails as well: `(D, P, E)` and `(cD, cP, E)` predict identical `M` for any `c > 0`, so **the formula is identified only up to a common scaling of `D` and `P`** and is testable in `r` and `M` rather than in `D` and `P` separately. The full seven-gate audit is Appendix A.3.

18.6.2 6.2 The maximum possible moment is `ln N`

This result is **derived, not chosen**, and the derivation is short.

Begin with the internal surplus: the surplus functional applied within a single observer. A perceiver holds `N` distinguishable categories; the framings brought to bear on a moment are states in that category space with mutual non-parallelism `u`, so `N_eff = 1 + (N−1)u` and `F_int = ln N_eff`. The range result gives `0 <= F <= ln N`, with `F = ln N` attained if and only if the framings are mutually orthogonal [Curzi, 2026]. Applying the chosen normalisation of Section 6.1 — the moment, expressed in the surplus register, is `F_int` — gives

```
0 <= (moment, in the surplus register) <= ln N

MAXIMUM POSSIBLE MOMENT = ln N
```

The richest moment available to a perceiver is bounded by the logarithm of the number of kinds of thing that perceiver can distinguish. No arrangement of attention, no intensity of experience, no quality of environment raises it. **It is a ceiling on the perceiver, not on the world.**

This is the law of requisite variety appearing in first-person experience [Ashby, 1956, 1958]: only variety absorbs variety, and a regulator's capacity is bounded by its own variety. Here the regulator is a perceiving mind, the disturbance is the world, and the bound is `ln N`. **It was not chosen. It follows from the range theorem plus the normalisation, and the normalisation is the only chosen step.** That is the honest form of the derivation and the reason both links must be published every time the result is quoted.

Three consequences.

First, the ceiling is logarithmic, so it is nearly immovable by effort. Doubling `N` adds `ln 2 = 0.693147` nats, or exactly one bit, or exactly two essence octaves, **no matter where you start.**

`N`	`ln N` (nats)	bits	octaves	cents
2	0.693147	1.00000	2.00000	2400.00
4	1.386294	2.00000	4.00000	4800.00
8	2.079442	3.00000	6.00000	7200.00
16	2.772589	4.00000	8.00000	9600.00
32	3.465736	5.00000	10.00000	12000.00
64	4.158883	6.00000	12.00000	14400.00
128	4.852030	7.00000	14.00000	16800.00
256	5.545177	8.00000	16.00000	19200.00
1000	6.907755	9.96578	19.93157	23917.88

Going from 2 distinguishable categories to 1000 — a five-hundred-fold increase in vocabulary — buys 6.21 nats, less than nine bits. A capacity problem cannot be escaped by adding categories.

Second, `N` is trainable — slowly, logarithmically, and only by acquiring genuinely distinct categories rather than relabelled ones.

Third, the ceiling is rarely the binding constraint, because `u` is usually low: the framings brought to bear are near-parallel. **The gain from raising `u` is almost always far larger than the gain from raising `N`**, which is the same conclusion the framework reaches in the group case, where the marginal price of non-parallelism is identical whether it arises by crossing or by divergence, `df/du = (N−1)/(1 + (N−1)u)`, while the two differ sharply in what they cost to produce [Curzi, 2026d].

A scale-type condition attaches to this bound and it is conditional, not passed. `N` is a count and is ratio-scaled with a true zero **provided "distinguishable category" has an operational definition**. Without one, `N` is not measured and the bound is not usable; the gate transfers to the construction of the category set, which is Section 8. The full audit is Appendix A.4.

18.7 7 The harmonic lattice

Section 4 established that a surplus can be written as a pitch by an exact constant. It also established the condition under which that is worth doing and then deferred it. The condition is this: **the essence law consumes an ordinal index, and an arbitrary ordinal produces an arbitrary frequency.** The law will happily convert nonsense. The whole burden of applying it therefore falls on the construction of the index.

This section works out the **structure** of the lattice the square law generates — what intervals exist, how they fold, how many distinct pitch classes an `N`-essence set can have, and which pairs are consonant. All of it is exact arithmetic and all of it is computed. Section 8 states the **requirements** an ordering must meet before a frequency may be attached to it.

One structural result is a genuine constraint rather than a convenience: **an essence set of `N` members contains exactly `ceil(N/2)` distinct pitch classes, whatever the essences are.** Half the vocabulary is octave duplication before the work begins.

18.7.1 7.1 Result CC-R1: every interval is a ratio of squares, and `k` cancels

```
f(m) / f(n) = (k m^2) / (k n^2) = m^2 / n^2
```

The golden ratio and the `9/8` set the absolute reference pitch and nothing else. Every relationship in the system is a ratio of squared integers. This is the most important structural fact about the lattice, because it means the chosen constant `k` — the one part of the law with no justification beyond preference — has **no influence whatever on any interval, consonance, pitch class or harmonic relation.** It sets where the instrument sits and not how it is tuned. Recomputed independently:

```
index ratio   essence ratio   cents      description
2 : 1          4 / 1          2400.00    two octaves exactly
3 : 2          9 / 4          1403.91    octave + 203.91 cents (a whole tone)
4 : 3         16 / 9           996.09    a minor seventh
5 : 4         25 / 16          772.63    near a minor sixth
3 : 1          9 / 1          3803.91    three octaves + 203.91 cents
```

The 203.91 cents appearing twice is not a coincidence: it is `1200 log2(9/8)`, the Pythagorean whole tone [Barbour, 1951], and it is the pitch class of index 3.

This result is labelled CC-R1 to keep the source material's own result numbering distinct from this paper's contributed results R1 to R6.

18.7.2 7.2 Result CC-R2: pitch class is the odd part of the index

Because `(2n)^2 = 4n^2` is exactly two octaves above `n^2`, multiplying an index by two transposes by a whole number of octaves and changes nothing else. Write any index as `n = 2^a m` with `m` odd. Then

```
f(n) = k (2^a m)^2 = 4^a (k m^2) = k m^2 transposed up by 2a octaves
```

Two essences share a pitch class if and only if their indices have the same odd part. The pitch class is fully determined by `m` and its position within the octave is `m^2` reduced by powers of two. The operation is octave transposition, whose perceptual standing as an equivalence is grounded for pitch and only for pitch [Shepard, 1982, Demany et al., 2013], with the caveat recorded in Section 3.3.

Theorem 7.1 (the class-count theorem). *An essence set indexed `1..N` contains exactly `ceil(N/2)` distinct pitch classes.*

Proof. The pitch class of index `n` is its odd part. For any odd `m` with `1 <= m <= N`, the index `m` itself lies in `1..N` and has odd part `m`, so every odd number up to `N` is realised. Conversely the odd part of any `n <= N` is itself at most `n`, hence at most `N`, and is odd. So the set of realised classes is exactly the set of odd integers in `[1, N]`, which has `ceil(N/2)` elements.

Verified by direct enumeration for `N = 1..24` — counted distinct odd parts against `ceil(N/2)`:

```
N   :  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
cnt :  1  1  2  2  3  3  4  4  5  5  6  6  7  7  8  8  9  9 10 10 11 11 12 12
ceil:  1  1  2  2  3  3  4  4  5  5  6  6  7  7  8  8  9  9 10 10 11 11 12 12
       all 24 agree
```

Twelve essences yield only six pitch classes. That is the available harmonic vocabulary and it is a hard structural limit. The same rule is stated as a theorem in its own right in Curzi [2026], where the odd-part class rule gives six classes at `N = 12`, and it is exercised at a different scale in Curzi [2026k], where twenty elements give ten classes.

18.7.3 7.3 The lattice, computed for `N = 24`

Columns: raw essence frequency `k n^2`; the factorisation `n = 2^a m`; the pitch class `m`; the position of `n` above the ground essence in cents, `2400 log2 n`; and the pitch class's position within the octave, that value modulo 1200.

n	f = k n^2 Hz	n = 2^a*m	class m	cents above ground	class cents
1	0.483426	2^0*1	1	0.0000	0.0000
2	1.933703	2^1*1	1	2400.0000	0.0000
3	4.350832	2^0*3	3	3803.9100	203.9100
4	7.734812	2^2*1	1	4800.0000	0.0000
5	12.085643	2^0*5	5	5572.6274	772.6274
6	17.403326	2^1*3	3	6203.9100	203.9100
7	23.687861	2^0*7	7	6737.6518	737.6518
8	30.939247	2^3*1	1	7200.0000	0.0000
9	39.157484	2^0*9	9	7607.8200	407.8200
10	48.342573	2^1*5	5	7972.6274	772.6274
11	58.494514	2^0*11	11	8302.6359	1102.6359
12	69.613306	2^2*3	3	8603.9100	203.9100
13	81.698949	2^0*13	13	8881.0553	481.0553
14	94.751444	2^1*7	7	9137.6518	737.6518
15	108.770790	2^0*15	15	9376.5374	976.5374
16	123.756988	2^4*1	1	9600.0000	0.0000
17	139.710037	2^0*17	17	9809.9108	209.9108
18	156.629937	2^1*9	9	10007.8200	407.8200
19	174.516690	2^0*19	19	10195.0260	595.0260
20	193.370293	2^2*5	5	10372.6274	772.6274
21	213.190748	2^0*21	21	10541.5618	941.5618
22	233.978055	2^1*11	11	10702.6359	1102.6359
23	255.732213	2^0*23	23	10856.5487	56.5487
24	278.453222	2^3*3	3	11003.9100	203.9100

Every frequency in that table is an essence frequency and not a physical resonance (HC-1).

18.7.4 7.4 The classes, their octave positions, and the collision cliff

Folding by odd part gives the complete harmonic vocabulary of a 24-essence set: twelve classes, as Theorem 7.1 requires.

class m	member indices <= 24	m^2 reduced to [1,2)	cents	Hz in [k, 2k)
1	1, 2, 4, 8, 16	1/1	0.000	0.483426
23	23	529/512	56.549	0.499477
3	3, 6, 12, 24	9/8	203.910	0.543854
17	17	289/256	209.911	0.545742
9	9, 18	81/64	407.820	0.611836
13	13	169/128	481.055	0.638273
19	19	361/256	595.026	0.681706
7	7, 14	49/32	737.652	0.740246
5	5, 10, 20	25/16	772.627	0.755353
21	21	441/256	941.562	0.832776
15	15	225/128	976.537	0.849772
11	11, 22	121/64	1102.636	0.913977

The pitch classes do not appear in index order. Class 23 sits 56.5 cents above the ground, below class 3 at 203.9 cents, though its index is nearly eight times larger. Squaring and folding scrambles the ordinal completely. **The lattice preserves the index in absolute pitch and destroys it in pitch class.** An analyst who reads pitch-class order as importance order has misread the arithmetic.

The vocabulary is unevenly distributed. The gaps between consecutive classes, in cents:

```
   0.000 ->   56.549    gap   56.549     (class 1 to class 23)
  56.549 ->  203.910    gap  147.361
 203.910 ->  209.911    gap    6.001     (class 3 to class 17)
 209.911 ->  407.820    gap  197.909
 407.820 ->  481.055    gap   73.235
 481.055 ->  595.026    gap  113.971
 595.026 ->  737.652    gap  142.626
 737.652 ->  772.627    gap   34.976     (class 7 to class 5)
 772.627 ->  941.562    gap  168.934
 941.562 ->  976.537    gap   34.976     (class 21 to class 15)
 976.537 -> 1102.636    gap  126.098
1102.636 -> 1200.000    gap   97.364     (wrapping to class 1)
```

Four pairs sit inside 60 cents of each other, close enough that as simultaneous tones they beat rather than form an interval — the roughness regime, in which two partials within a critical band fuse into a rough single percept rather than separating into two pitches [Plomp and Levelt, 1965, von Helmholtz, 1885, Terhardt, 1974]:

```
classes  3 and 17 :   6.00 cents apart   (  9/8   vs 289/256 )
classes  7 and  5 :  34.98 cents apart   ( 49/32  vs  25/16  )
classes 21 and 15 :  34.98 cents apart   (441/256 vs 225/128 )
classes  1 and 23 :  56.55 cents apart   (  1/1   vs 529/512 )
```

The 6.00-cent gap is the one to watch. **Essence index 3 and essence index 17 are for practical purposes the same pitch class, and nothing in the ordering of any domain caused that — it is an artefact of `9/8` and `289/256` being nearly equal.** If a domain's third and seventeenth categories are conceptually remote, the lattice will nevertheless render them as a near-unison. The 34.98-cent coincidence between classes 5 and 7, and identically between 15 and 21 since `21/15 = 7/5`, is the same phenomenon at a coarser scale.

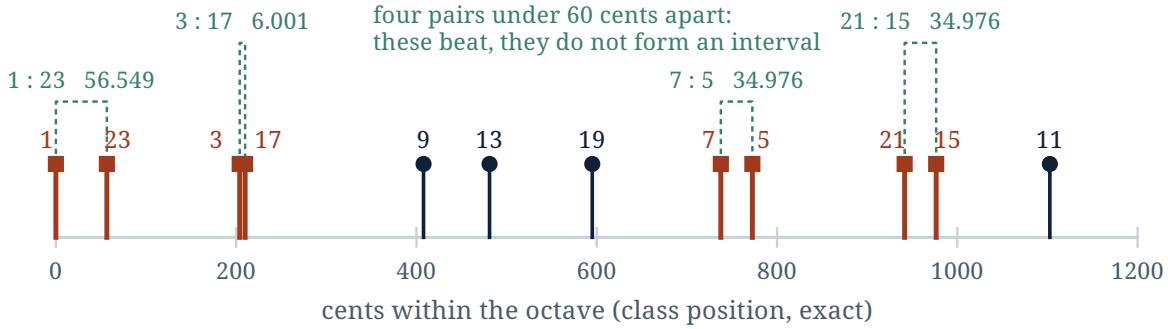


Figure 18.2: **Twelve pitch classes for twenty-four essences, and four of the twelve pairs sit inside 60 cents.** Class positions are exact rational values reduced into one octave; `k` cancels, so the chosen reference constant plays no part. The four disclosed collisions are classes 3 and 17 at 6.001 cents, 7 and 5 at 34.976, 21 and 15 at 34.976, and 1 and 23 at 56.549. **These are facts about integers and not about any domain.** Generated by `figs/src/CCP-018_fig02.py` from the class table of Section 7.4.

The design consequence is direct: essence sets past a certain size produce near-unison collisions the ordering did not intend and cannot control. That is an argument for small essence sets with nothing to do with taste.

And the size is exact, which is better than "about twelve". Computing the folded lattice's minimum gap in exact rational arithmetic for every set size gives a step function with four values and three cliffs:

```
N =  3 .. 6     minimum gap    9/8     = 203.910 cents     no collision
N =  7 .. 16    minimum gap   50/49    =  34.976 cents     classes 7 and 5
N = 17 .. 40    minimum gap  289/288   =   6.001 cents     classes 17 and 3
N = 41 and up   minimum gap 1682/1681  =   1.030 cents     classes 41 and 29
```

The third step is recorded for completeness only: `N = 41` lies far beyond the 25-essence ceiling that Section 8.3 derives for a shared transposition, so no renderable set reaches it.

The first coincidence arrives at `N = 7`, when class 7 lands 34.976 cents from class 5. The lattice holds that gap unchanged for ten consecutive set sizes and then collapses the moment index 17 enters, because `17^2 = 289` sits one part in 288 above `9/8 = 288/256`. Nothing in any domain causes either event; both are arithmetic.

The cliff is a change of regime rather than merely a narrower gap, and the time-bandwidth key of Section 4.5 prices it: telling two tones apart requires `1/(8 pi dF)` cycles, so a 34.976-cent separation needs **3.94 cycles** and a 6.001-cent separation needs **22.96 cycles**. **A set of sixteen and a set of seventeen require rendering durations differing by a factor of 5.828, and the boundary between them is exact.** Design at 16 or below where possible; if not, budget the tone duration before budgeting anything else.

Two corrections to earlier statements of this boundary are recorded here rather than applied silently. The corpus has stated the boundary at `N = 19` to `N = 20`; the exact lattice returns `N = 16` to `N = 17`, which is what the table of Section 7.3 shows — index 17 is present in a 24-set and its 6.001-cent gap with class 3 is listed above. And the source material states the rate collapse as "a factor of exactly 8" in one sentence and as 5.828 in the parenthesis of the sentence before it. **Recomputation gives 5.8284 by both routes** — the cents ratio `34.976/6.001` and the cycle-count ratio `22.957/3.939` are the same number, `ln(50/49)/ln(289/288)` — **and no calculation returning 8 could be reconstructed. The factor is 5.828.**

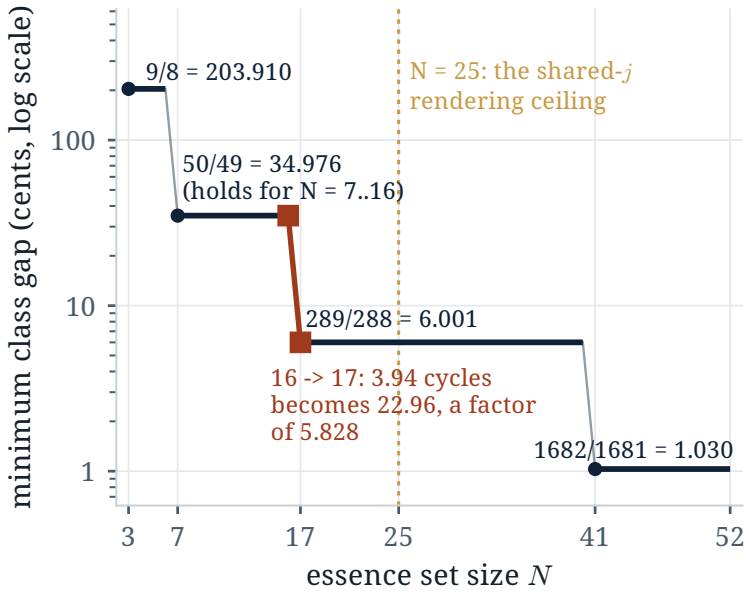


Figure 18.3: **The minimum class gap is a step function of set size with three exact cliffs, and the worst of them lands one index above a twelve-category set.** The gap holds at `9/8` = 203.910 cents for `N = 3` to 6, at `50/49` = 34.976 for `N = 7` to 16, at `289/288` = 6.001 for `N = 17` to 40, and at `1682/1681` = 1.030 from `N = 41`. Crossing from 16 to 17 raises the tone duration needed to separate the closest pair from 3.94 cycles to 22.96, a factor of 5.828. The shared-transposition ceiling at `N = 25` is marked; no renderable set reaches the third cliff. Exact rational arithmetic, generated by `figs/src/CCP-018_fig03.py` from Sections 7.4 and 8.3.

18.7.5 7.5 The interval lattice in cents

The interval between essence `i` and essence `j` is `2400 log2(j/i)` cents. The full matrix for indices 1 to 12, upper triangle; the lower triangle is the negative.

	j=2	3	4	5	6	7	8	9	10	11	12
i=1	2400	3804	4800	5573	6204	6738	7200	7608	7973	8303	8604
i=2	–	1404	2400	3173	3804	4338	4800	5208	5573	5903	6204
i=3	–	–	996	1769	2400	2934	3396	3804	4169	4499	4800
i=4	–	–	–	773	1404	1938	2400	2808	3173	3503	3804
i=5	–	–	–	–	631	1165	1627	2035	2400	2730	3031
i=6	–	–	–	–	–	534	996	1404	1769	2099	2400
i=7	–	–	–	–	–	–	462	870	1235	1565	1866
i=8	–	–	–	–	–	–	–	408	773	1103	1404
i=9	–	–	–	–	–	–	–	–	365	695	996
i=10	–	–	–	–	–	–	–	–	–	330	631
i=11	–	–	–	–	–	–	–	–	–	–	301

Read the diagonal band — the intervals between adjacent indices — and a pattern appears worth stating as a result in its own right.

adjacent pair	raw interval (cents)
1 – 2	2400.00
2 – 3	1403.91
3 – 4	996.09
4 – 5	772.63

```
 5 – 6              631.28
 6 – 7              533.74
 7 – 8              462.35
 8 – 9              407.82
 9 – 10             364.81
10 – 11             330.01
11 – 12             301.27
12 – 13             277.15
15 – 16             223.46
19 – 20             177.60
23 – 24             147.36
```

Neighbouring essences crowd together as the index rises. The gap between the first two essences is two octaves; the gap between the twenty-third and twenty-fourth is 147 cents, less than a whole tone and a quarter. This is the logarithm again: `2400 log2((n+1)/n)` falls like `2400/(n ln 2)`. In a set of 24 the top is harmonically congested and the bottom is spread across octaves.

Consequence for set design: the resolution of the lattice is not uniform, and low indices carry far more pitch distance than high indices. If a domain's most important distinctions land at high index, the lattice renders them as small intervals. That is a reason to think about which end of the ordering the important categories occupy — and, critically, **it is not a licence to reorder the set to improve the result**, which the admissibility requirement ES-R2 of Section 8.1 forbids.

18.7.6 7.6 Consonance as a function of the index ratio

Which pairs of essences sound related, and which sound remote? The question has a precise answer in this lattice and the answer depends **only on the reduced ratio of the indices**. Let `i/j` reduce to `p/q` with `gcd(p,q) = 1`. The essence interval is then `p^2/q^2`, and since `gcd(p^2, q^2) = 1` that fraction is already in lowest terms. Everything about the pair's harmonic relationship is a function of `p` and `q` alone.

Two computable indices are used. **Both are CHOSEN measures with a defined arithmetic, not measured facts about human hearing.**

```
TENNEY HARMONIC DISTANCE       HD(p/q) = log2(p q)
    A defined measure of the complexity of a frequency ratio. Larger = more remote.

EULER'S GRADUS SUAVITATIS      G(n) = 1 + sum over prime factors of e_i (prime_i – 1)
    G(p/q) = G(p q). Larger = less smooth.
```

The harmonic distance is Tenney's city-block metric on the harmonic-space lattice, `HD(f_a, f_b)` proportional to `log a + log b = log(ab)` [Tenney, 2015]; the gradus suavitatis is Euler's degree of agreeableness over the prime factorisation [Euler, 1739]. **Neither is a measurement.** What is actually measured about consonance is a different object and points a different way: roughness is set by critical bandwidth rather than by integer ratio [Plomp and Levelt, 1965], sensory and harmonic consonance are distinct [Terhardt, 1974], dyad consonance has been quantified directly [Kameoka and Kuriyagawa, 1969], and consonance depends on the **spectrum** of the sound and not on the ratio alone, so that no ratio is a universal constant [Sethares, 2005]. **The two literatures disagree, and this paper publishes both indices and adjudicates neither.**

> **Theorem 7.2 (the doubling theorem).** *For any index pair reducing to `p/q`, the essence interval `p^2/q^2` has exactly twice the Tenney harmonic distance of the ratio `p/q`, and Euler gradus `2G – 1`.*

Proof. `HD(p^2/q^2) = log2(p^2 q^2) = 2 log2(pq) = 2 HD(p/q)`. For the gradus, squaring doubles every exponent in the prime factorisation of `pq`, so the sum `sum e_i (prime_i − 1)` doubles, and `G(p^2 q^2) = 1 + 2 sum e_i (prime_i − 1) = 2(1 + sum e_i (prime_i − 1)) − 1 = 2 G(pq) − 1`.

Verified exhaustively over all 66 pairs `1 <= j < i <= 12`, both identities exact in every case. A sample:

i	j	p/q	HD(p/q)	HD(p^2/q^2)	2∗HD	G(p/q)	G(ess)	2G−1	cents
2	1	2/1	1.0000	2.0000	2.0000	2	3	3	2400.00
3	2	3/2	2.5850	5.1699	5.1699	4	7	7	1403.91
4	3	4/3	3.5850	7.1699	7.1699	5	9	9	996.09
5	4	5/4	4.3219	8.6439	8.6439	7	13	13	772.63
6	5	6/5	4.9069	9.8138	9.8138	8	15	15	631.28
7	6	7/6	5.3923	10.7846	10.7846	10	19	19	533.74
8	7	8/7	5.8074	11.6147	11.6147	10	19	19	462.35
8	3	8/3	4.5850	9.1699	9.1699	6	11	11	3396.09
7	5	7/5	5.1293	10.2586	10.2586	11	21	21	1165.02

The square law systematically doubles harmonic distance. Every essence pair is exactly twice as harmonically remote as the corresponding integer-ratio pair would be in ordinary just intonation [Partch, 1974, Doty, 2002]. The `3:2` — the perfect fifth, the most consonant interval after the octave — becomes `9:4`, an octave plus a whole tone, at double the harmonic distance; the major third `5:4` becomes `25:16`. **Consonance is scarce in the essence lattice, and it is scarce by a factor of exactly two.**

18.7.7 7.7 Consonance is redundancy

Ranking all distinct index ratios among indices 1 to 12 by essence harmonic distance:

HD_ess	Gradus	index ratio	essence ratio	reduced	cents
0.0000	1	1:1	1/1	1	0.00
2.0000	3	2:1	4/1	1	0.00
3.1699	5	3:1	9/1	9/8	203.91
4.0000	5	4:1	16/1	1	0.00
4.6439	9	5:1	25/1	25/16	772.63
5.1699	7	3:2	9/4	9/8	203.91
5.6147	13	7:1	49/1	49/32	737.65
6.0000	7	8:1	64/1	1	0.00
6.3399	9	9:1	81/1	81/64	407.82
6.6439	11	5:2	25/4	25/16	772.63
6.9189	21	11:1	121/1	121/64	1102.64
7.1699	9	4:3	16/9	16/9	996.09
7.6147	15	7:2	49/4	49/32	737.65
7.8138	13	5:3	25/9	25/18	568.72
8.6439	13	5:4	25/16	25/16	772.63
8.7846	17	7:3	49/9	49/36	533.74
9.6147	17	7:4	49/16	49/32	737.65
9.8138	15	6:5	36/25	36/25	631.28
10.2586	21	7:5	49/25	49/25	1165.02
10.6439	15	8:5	64/25	32/25	427.37
11.6147	19	8:7	64/49	64/49	462.35
12.2586	23	10:7	100/49	50/49	34.98
12.9837	19	10:9	100/81	100/81	364.81
14.0888	29	12:11	144/121	144/121	301.27

The two indices sometimes disagree in their fine ordering — `5:1` has lower Tenney distance than `3:2` but a higher Euler gradus, 9 against 7, because Euler penalises the prime 5 more heavily than Tenney does. **Where two defensible indices disagree, publish both and do not adjudicate.** Neither is a measurement.

The most consonant nontrivial essence pairs are exactly the octave-equivalent ones. Every pair whose index ratio reduces to `2:1` has `HD_ess = 2.0000`, gradus 3, and an interval of exactly two octaves — pitch class identical. Nothing else comes close: the next entry, `3:1`, is at 3.1699. The same identity is proved in its sharpest form in Curzi [2026], where the minimum `HD_ess = 2.0000` over a twelve-set is attained by exactly the six collision pairs `{1:2, 2:4, 3:6, 4:8, 5:10, 6:12}` and by nothing else.

Which produces the central design tension, and it is exact:

CONSONANCE IS REDUNDANCY. The pairs that sound most related are precisely the pairs that occupy the same pitch class — the pairs the lattice cannot distinguish. **In the essence lattice, harmonic agreement and informational redundancy are the same phenomenon.** This is not a metaphor: the arithmetic is the same arithmetic that makes a set of near-parallel categories worth `ln 1` rather than `ln N`.

Searching all 924 six-element subsets of indices 1 to 12 exhaustively, by mean pairwise essence harmonic distance:

```
minimum mean HD_ess overall          4.0353    {1, 2, 3, 4, 6, 12}    2 distinct classes
maximum mean HD_ess overall         11.4064    {7, 8, 9, 10, 11, 12}  6 distinct classes
minimum mean HD_ess with 6 classes   8.4731    {1, 3, 5, 7, 9, 11}
```

Read the first line. **The most consonant six-element essence set available collapses to two pitch classes.** Six essences, two distinguishable tones: the non-parallelism has gone to zero in the sonic register and the set is worth `ln 2` where it should be worth up to `ln 6`.

Read the third. **Subject to keeping all six pitch classes distinct, the most consonant six-element set is the odd indices `{1, 3, 5, 7, 9, 11}`** — exactly the canonical class-representative set of Theorem 7.1. That convergence is also a usable design rule: **when a subset of a large essence set must be chosen for sonification, choose the odd indices.** They are the class representatives, they are guaranteed distinct, and among all fully-distinct choices they are the least dissonant.

18.8 8 Requirements on an essence set

18.8.1 8.1 The three admissibility requirements

Before any of the above may be applied to a real domain, the domain's categories must be turned into an index. Made operational, this is three requirements. *They are labelled ES-R1 to ES-R3 to keep them distinct from this paper's contributed results R1 to R6.*

ES-R1 — EXHAUSTIVE AND DISJOINT. The categories partition the domain. Every state falls in exactly one. If two categories overlap, an entity has two indices and therefore two frequencies, and the lattice is not well defined.

ES-R2 — AN ORDERING FROM OUTSIDE THE MUSIC. The ordinal is fixed by a property of the domain established before, and independently of, any intention to sonify it. **The ordering may not be chosen to make the result sound good.** This is the essence-set analogue of preregistration [Nosek et al., 2018] and it is the requirement most often violated; violating it is the essence-set form of the researcher-degrees-of- freedom problem [Simmons et al., 2011, Gelman and Loken, 2013].

ES-R3 — REPRODUCIBLE ASSIGNMENT. A second competent analyst, given the domain and the stated ordering rule, produces the same index for the same category. This is checkable by inter-rater agreement and is the input-side form of the reproducibility requirement [Open Science Collaboration, 2015].

A set failing any of ES-R1 to ES-R3 may not be sonified. It may still be described, tabulated, and reasoned about — it simply does not get a frequency. Section 12.3 carries a worked refusal.

18.8.2 8.2 Constructing a defensible essence set

Step 1 — State the domain and its boundary. What is being categorised, and what is excluded. Without a boundary, exhaustiveness cannot be checked.

Step 2 — Identify the ordering source. Three kinds are admissible, in descending order of strength.

(a) An intrinsic physical or structural magnitude. The categories are ordered by a measurable quantity that exists independently of the analyst — distance, duration, count, energy, age. This is the strongest source because ES-R2 is satisfied automatically: the ordering predates the analysis. The proxemic zones of Section 12.3 are of this kind, ordered by distance [Hall, 1966].

(b) A generative or constructive order. The categories are produced by a process with a natural sequence and the index is position in that sequence — stages of a protocol, steps of a manufacturing chain, layers of a stack. Admissible when the sequence is fixed by the process rather than by preference.

(c) A measured empirical rank with demonstrated reliability. The ordering is established by measurement — a Rasch-calibrated item difficulty ordering, for instance [Rasch, 1960] — with inter-rater or test-retest reliability reported. Weakest of the three, because the ordering is now an estimate carrying its own uncertainty, which must be propagated. **If the rank order is unstable across samples, the frequencies are unstable, and the construction must be reported with that instability visible.**

Step 3 — Fix the direction. Ascending or descending must be justified, because reversing the ordering changes every frequency and every interval. The default rule adopted here: **index 1 is the ground, so index 1 goes to the category that is nearest, smallest, earliest, or most basic in the ordering source.** Reversal is permitted only with a stated reason.

Step 4 — Count. `N` is the number of categories. It determines the ceiling `ln N` and the class count `ceil(N/2)`. Both are reported.

Step 5 — Check for collisions. Compute the class cents and check the gaps (Section 7.4). Any gap under about 30 cents will be heard as beating rather than as an interval and must be disclosed in the publication of the set.

Step 6 — Publish the whole construction, not the frequencies. The domain, the boundary, the ordering source, the direction, the index assignment for every category, `N`, the class count, the collisions, the band, and the transposition. **A frequency published without its construction is not a result; it is a decoration.**

What is inadmissible as an ordering source: alphabetical order; order of discovery or historical accident; order of importance as judged by the analyst; frequency of occurrence in a convenience sample; and — the one that matters most — **any order arrived at by trying orderings and keeping the one that sounded best.** That converts the whole apparatus into an elaborate way of writing down a preference.

What a well-formed essence set looks like on paper. Small, under about twelve, above which near-unison collisions appear that the ordering did not intend and the marginal ceiling gain is logarithmic anyway. Ordered from outside. Fully enumerated in print: every category, its index, its class, its cents. Honest about its collisions, with any class gap under 30 cents disclosed. Reported with `N`, `ceil(N/2)`, the ceiling `ln N`, the band and the transposition rule. Free of adjustment: no index moved to improve the sound. Badged throughout — essence frequency, not physical resonance; chosen `k`; postulated square law; ordinal input promoted by postulate.

18.8.3 8.3 Band placement and transposition

Raw essence frequencies for a 24-set span 0.483 Hz to 278 Hz, from far below hearing to the low bass. Octave transposition brings them into a chosen band, `f_audible = f 2^j` with `j` an integer. Transposition by a whole number of octaves preserves pitch class exactly. Placing indices 1 to 24 in `[110, 220)` Hz, one `j` per index:

n	f raw Hz	j	f*2^j Hz	class m	n	f raw Hz	j	f*2^j Hz	class m
1	0.483426	8	123.7570	1	13	81.698949	1	163.3979	13
2	1.933703	6	123.7570	1	14	94.751444	1	189.5029	7
3	4.350832	5	139.2266	3	15	108.770790	1	217.5416	15
4	7.734812	4	123.7570	1	16	123.756988	0	123.7570	1
5	12.085643	4	193.3703	5	17	139.710037	0	139.7100	17
6	17.403326	3	139.2266	3	18	156.629937	0	156.6299	9
7	23.687861	3	189.5029	7	19	174.516690	0	174.5167	19
8	30.939247	2	123.7570	1	20	193.370293	0	193.3703	5
9	39.157484	2	156.6299	9	21	213.190748	0	213.1907	21
10	48.342573	2	193.3703	5	22	233.978055	−1	116.9890	11
11	58.494514	1	116.9890	11	23	255.732213	−1	127.8661	23
12	69.613306	1	139.2266	3	24	278.453222	−1	139.2266	3

Read the class column against the frequency column: every index in class 1 lands on 123.7570 Hz, every index in class 3 on 139.2266 Hz. **Per-index transposition into a single octave collapses the set onto its pitch classes exactly — twelve tones for twenty-four essences, as Theorem 7.1 requires.**

And that is the hazard. Per-index transposition destroys the intervals between indices while preserving classes. The 24-essence set spans 11003.91 cents, more than nine octaves, and folding it into one octave discards all of that. **Two operations are available and must never be confused:**

```
SHARED j     f_i * 2^j for one fixed j     preserves ALL intervals; needs a wide band
PER-INDEX j  f_i * 2^{j_i}                 preserves pitch CLASSES only; fits one octave
```

State the band and the `j` for every value published, and state which of the two operations was performed. A set of values whose intervals are the point takes a shared `j`.

A shared `j` has a hard capacity. The set `n = 1..N` spans a frequency ratio of `N^2`, so fitting it inside 20 Hz to 20 kHz under one shared `j` requires `k 2^j >= 20` at the bottom and `k N^2 2^j <= 20000` at the top. The first inequality forces `j >= 6`, since `20/k = 41.37` and `2^5 = 32` is too small. The second then caps `N`:

```
j = 6    k 2^6 =    30.94 Hz     N <= 25      k (25)^2 2^6 = 19337.0 Hz
j = 7    k 2^7 =    61.88 Hz     N <= 17
j = 8    k 2^8 =   123.76 Hz     N <= 12
j = 9    k 2^9 =   247.51 Hz     N <=  8
```

At most 25 essences can be rendered under a single shared transposition, and only at `j = 6`. That is a bound on set size, not on volume or equipment, and it does not move. **A twelve-category set fits comfortably at `j = 6` or `j = 8`; a thirty-two-category set does not fit under any shared `j`** and must be rendered under split transposition — two or more shared-`j` groups, each internally interval-preserving, with the relation *between* the groups explicitly declared. Splitting is a legal move; leaving the split undeclared is the failure, because a reader cannot tell a split rendering from a broken one without the declaration.

One further constraint belongs beside the band. A rectangular on/off gate has no finite spectral second moment — its `sinc^2` tails fall as `1/f^2` and the second-moment integral diverges — so it has no defined bandwidth and **is disqualified as a rendering envelope** [Gabor, 1946]. Use a windowed envelope. **A hard-gated tone is not the tone whose frequency was published.**

18.9 9 Modal design: building from an essence set

18.9.1 9.1 The procedure

Six steps, every one checkable by a third party from the published record.

```
STEP 1   Declare the category set and the ordinal rule.        (ES-3)
STEP 2   Run the permutation sensitivity test.                 (ES-4)
STEP 3   Compute raw essence frequencies f(n) = k n^2.
STEP 4   Reduce to pitch classes by odd part; count ceil(N/2). (ES-6, ES-7)
STEP 5   Fold into the working band; publish band and j.       (ES-10)
STEP 6   Apply to one or more of the four targets below.
```

The four targets: **(a)** acoustic tuning targets, interacting with Sabine absorption (Section 9.2); **(b)** room dimension ratios and modal frequencies (Section 9.3); **(c)** scheduling and cadence, as periodic forcing of a low-pass filter (Section 9.4); **(d)** tuned physical elements (Section 9.5).

Step 5 carries most of the practical judgement. Raw essence frequencies span from subaudible to ultrasonic: a twelve-set runs 0.483426 Hz to 69.613306 Hz, entirely below the useful speech band. Octave transposition preserves pitch class exactly and is the adopted method. **State the band and the `j` for every value published; a folded value without its `j` is not comparable to anything.**

The working band adopted here for structural work is `[220, 440)` Hz, one octave anchored on A3 = 220 Hz. It was chosen because it sits inside the speech fundamental range for adult voices, is well above typical room-mode territory, is comfortably below the 2 to 4 kHz consonant band where intelligibility lives, and folds a twelve-set into a single octave with `j` between +2 and +9 [Rossing et al., 2002]. **That is a design choice, stated as one.**

18.9.2 9.2 Target (a): acoustic tuning and the Sabine interaction

Acoustic design is an intervention on the equality of conversational turn-taking, and turn-taking is substantially a property of the room: **a hard-surfaced room silences the quiet voice by physics, not by dominance** [Woolley et al., 2010, Curzi, 2026e]. That mechanism is priced in the corpus: silencing three of eight participants costs `ln 8 − ln 5 = 0.4700` nats, 22.6 per cent of a room's per-interaction surplus [Curzi, 2026e]. **The essence law adds exactly one thing to that: it says which octave bands to treat first.**

Absorption is not flat. Porous absorbers — mineral wool, heavy curtain, carpet — are efficient above about 500 Hz and become progressively useless below 250 Hz; membrane and Helmholtz absorbers are narrowband and must be tuned [Beranek, 2004]. So the practical question in any real treatment budget is *which bands do I spend on*, and the essence set answers it by identifying where the room's own programme places its pitch classes. Worked, for a twelve-set folded into `[220, 440)`:

class	folded Hz	octave band	inside the 250–4000 Hz speech range?
PC1	247.5140	250	at the lower band edge
PC3	278.4532	250	yes
PC5	386.7406	500	yes
PC7	379.0058	500	yes
PC9	313.2599	250	yes
PC11	233.9781	250	at the lower band edge

Four of six classes land in the 250 Hz octave band and two in the 500 Hz band. That is the treatment priority and it is a useful one: the 250 Hz band is exactly where cheap porous absorption starts losing efficiency and where most under-treated meeting rooms have their worst speech-intelligibility problem. **The essence set points at the band most rooms neglect.** *This is a category-1 acoustic fact about absorber physics and a category-3 choice of which band to prioritise; the two are separable and are separated here.*

The Sabine arithmetic for a 100 m3 room, computed from `RT60 = 0.161 V / A` [Sabine, 1922]:

`A` (m2 sabins)	8	10	13	16	20	26	32	40
`RT60` (s)	2.013	1.610	1.238	1.006	0.805	0.619	0.503	0.403

and inverted, the absorption required to hit a target:

```
V = 100 m^3
RT60 = 0.4 s  ->  A = 40.25 m^2 sabins
RT60 = 0.5 s  ->  A = 32.20 m^2 sabins
RT60 = 0.6 s  ->  A = 26.83 m^2 sabins
RT60 = 0.8 s  ->  A = 20.12 m^2 sabins
RT60 = 1.0 s  ->  A = 16.10 m^2 sabins
```

One hard limit the essence law cannot touch and must not be allowed to obscure. The Schroeder frequency [Schroeder, 1965],

```
f_s = 2000 sqrt(RT60 / V)
```

marks the transition between modal behaviour — discrete resonances, position-dependent response — and statistical behaviour, the diffuse field in which Sabine's relation applies. Computed for `V = 100` m3:

```
RT60 = 0.40 s  ->  f_s = 126.5 Hz
RT60 = 0.50 s  ->  f_s = 141.4 Hz
RT60 = 0.60 s  ->  f_s = 154.9 Hz
RT60 = 0.80 s  ->  f_s = 178.9 Hz
```

Below `f_s`, Sabine's equation does not apply and the room's behaviour is set by individual modes. Above it, modes are dense and Sabine governs. Since a folded twelve-set lives between 234 and 387 Hz it sits entirely in the statistical region for a room of this size, which is the correct place for it and why the `[220, 440)` band was chosen. **A folded band straddling `f_s` would produce a set of targets half of which obeyed one physics and half another, and the resulting specification would be incoherent.** Check `f_s` before choosing the band, always.

18.9.3 9.3 Target (b): room dimension ratios and modal frequencies

This is the strongest of the four applications, because room proportion is a genuinely open design question with a wide space of acceptable answers and the essence law narrows it defensibly.

An axial mode along a dimension `L` sits at `f = c/(2L)`, so `L = c/(2f)` [Rossing et al., 2002]. Choose three essence classes, one per axis, and the room's three axial fundamentals *are* those essence frequencies. The ratio of the dimensions is then the inverse ratio of the class frequencies — a ratio of squares by CC-R1, hence always a rational number with an exact expression.

The constraint that makes this non-trivial is that most ratios are bad rooms. Room-ratio heuristics exist because degenerate modes — two or three modes at the same frequency — create strong peaks and long decays, and near-degenerate axial modes create audible colouration of speech [Beranek, 2004]. The published heuristics are heuristics and are stated here as such:

```
Louden          1 : 1.4  : 1.9
Sepmeyer A      1 : 1.14 : 1.39
Sepmeyer B      1 : 1.28 : 1.54
Sepmeyer C      1 : 1.60 : 2.33
Volkmann        1 : 1.5  : 2.5
```

So the design question becomes: **is there an assignment of essence classes to axes whose ratio lands near a recommended ratio?** Searched exhaustively over all class triples and all octave offsets from 0 to 6, for the six classes of a twelve-set. Best matches, recomputed:

heuristic target	nearest essence ratio	classes used	Euclidean error
Louden 1 : 1.40 : 1.90	1 : 1.3611 : 1.9600	PC3, PC5, PC7	0.0715
Sepmeyer A 1 : 1.14 : 1.39	1 : 1.1250 : 1.4400	PC1, PC3, PC5	0.0522
Sepmeyer B 1 : 1.28 : 1.54	1 : 1.2346 : 1.5625	PC1, PC5, PC9	0.0507
Sepmeyer C 1 : 1.60 : 2.33	1 : 1.6200 : 2.2500	PC3, PC5, PC9	0.0825
Volkmann 1 : 1.50 : 2.50	**1 : 1.4938 : 2.4694**	**PC7, PC9, PC11**	**0.0312**

The best match is PC7 : PC9 : PC11 against Volkmann, at a Euclidean error of 0.031 — the middle term within 0.4 per cent and the long term within 1.2 per cent of the heuristic. And the ratio has an exact closed form, which is the point of using the essence law rather than a table lookup:

```
L proportional to 1/f, and f proportional to m^2, so
L(PC11) : L(PC9) : L(PC7)  =  1/121 : 1/81 : 1/49
                           =  1 : 121/81 : 121/49
                           =  1 : 1.493827... : 2.469388...
```

The ratio is exactly `1 : 121/81 : 121/49`. Not approximately. **That is what CC-R1 buys: every proportion this method produces is a ratio of squares of integers and can be written down exactly, which means it can be checked, argued with, and reproduced by someone who does not trust the author.**

Anchoring the absolute size by requiring the axial fundamentals to *be* the essence frequencies:

```
f(7)  = k * 49  = 23.687861 Hz  ->  Lz = 343/(2*23.687861) = 7.2400 m   (length)
f(9)  = k * 81  = 39.157484 Hz  ->  Ly = 343/(2*39.157484) = 4.3798 m   (width)
f(11) = k * 121 = 58.494514 Hz  ->  Lx = 343/(2*58.494514) = 2.9319 m   (height)

V = 92.969 m^3        floor area = 31.709 m^2        total surface S = 131.555 m^2
occupancy at 2.5 m^2 per person = 12.7 people
```

A room 2.93 m high, 4.38 m wide and 7.24 m long, holding about twelve people. The height is buildable, the floor plan is a conventional meeting-room shape, and the occupancy matches the twelve-category set it was derived from. **None of that was forced, and it should be said plainly that it is lucky** — a different `k` or a different class triple would have produced an unbuildable room, and FM-3 covers what to do then.

The modal structure, recomputed for this room, first twenty modes:

#	f (Hz)	(p,q,r)	type	gap from previous (Hz)	essence relation
1	23.69	(0,0,1)	axial	–	`f(7) x 2^0`
2	39.16	(0,1,0)	axial	15.47	`f(9) x 2^0`
3	45.76	(0,1,1)	tangential	6.61	–
4	47.38	(0,0,2)	axial	1.61	`f(7) x 2^1`
5	58.49	(1,0,0)	axial	11.12	`f(11) x 2^0`
6	61.46	(0,1,2)	tangential	2.97	–
7	63.11	(1,0,1)	tangential	1.65	–
8	70.39	(1,1,0)	tangential	7.28	–
9	71.06	(0,0,3)	axial	0.67	–
10	74.27	(1,1,1)	oblique	3.21	–
11	75.27	(1,0,2)	tangential	1.00	–
12	78.31	(0,2,0)	axial	3.04	`f(9) x 2^1`
13	81.14	(0,1,3)	tangential	2.82	–
14	81.82	(0,2,1)	tangential	0.68	–
15	84.85	(1,1,2)	oblique	3.03	–
16	91.53	(0,2,2)	tangential	6.68	–
17	92.04	(1,0,3)	tangential	0.51	–
18	94.75	(0,0,4)	axial	2.71	`f(7) x 2^2`
19	97.75	(1,2,0)	tangential	3.00	–
20	100.02	(1,1,3)	oblique	2.28	–

```
modes below 150 Hz:              55  (48 if the mode set is truncated at p,q,r <= 4)
mean gap of the first 20 modes:  4.02 Hz
maximum gap:                     15.47 Hz  (between modes 1 and 2)
minimum gap:                     0.51 Hz   (modes 16 -> 17)
```

Read the "essence relation" column and CC-R2 is visible in built form. **The room's axial mode series along each axis is `f, 2f, 3f, 4f, ...`, and the even members of each series are the essence frequencies of that axis's class at higher octaves.** Mode 4 is `f(7) x 2`, mode 18 is `f(7) x 4`, mode 12 is `f(9) x 2`. The odd members — mode 9 at 71.06 Hz, which is `3 f(7)` — are *not* essence frequencies of anything in the set, because tripling an index is not an operation the essence law preserves. **The room delivers the essence set at the even harmonics and delivers ordinary harmonic-series members between them. Do not report the odd members as essence frequencies; they are not.**

Three honest acoustic observations about this specific room, none predicted by the essence law and all produced by conventional analysis.

1. **The 15.47 Hz gap between modes 1 and 2 is large.** A single isolated mode at 23.69 Hz with nothing near it will ring. At this frequency it is below the useful speech range and mostly inaudible as pitch, but it will couple to footfall and to plant. Treat with a membrane absorber or accept it.
2. **The 0.51 Hz gap between modes 16 and 17 is a near-degeneracy.** Two modes at 91.53 and 92.04 Hz will beat at about half a hertz — a slow throb. Both are tangential, so the effect is mild, but it is real and worth knowing before someone reports the room as "wobbly" and nobody connects the reports.
3. **55 modes below 150 Hz is a healthy density for a room of this volume**; the low end is not sparse. The Schroeder frequency at `RT60 = 0.5` s is `2000 sqrt(0.5/92.969) = 146.7` Hz, consistent with that count.

Sabine treatment for this room, recomputed:

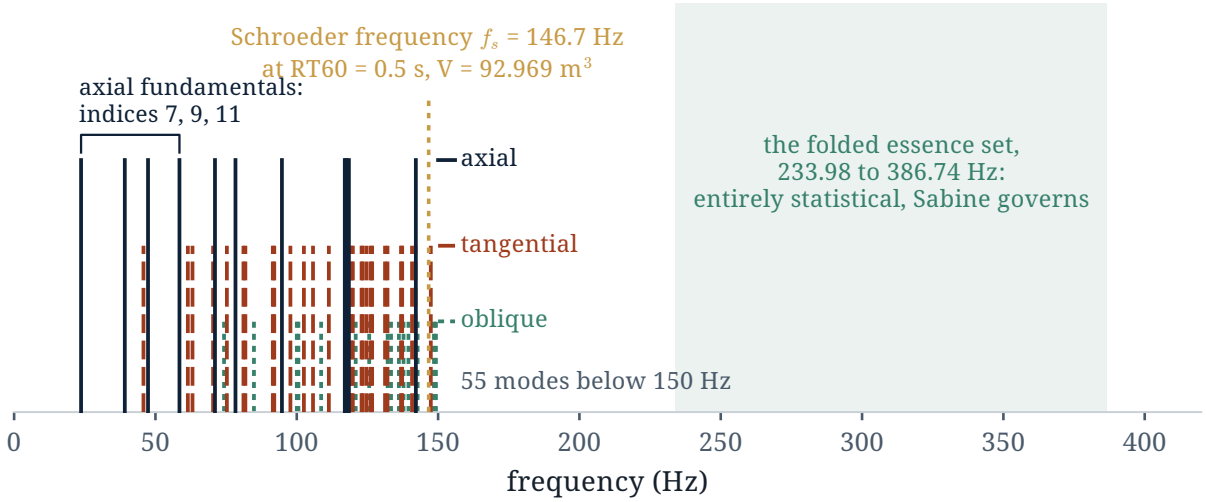


Figure 18.4: **The derived room's entire modal region lies below the Schroeder frequency and its entire folded essence set lies above it, which is what makes the specification coherent.** All 55 modes below 150 Hz are shown, classified axial, tangential and oblique, for a room of 7.2400 by 4.3798 by 2.9319 m at `c = 343` m/s. The Schroeder boundary at `RT60 = 0.5` s and `V = 92.969` m3 is 146.7 Hz; the folded set runs 233.98 to 386.74 Hz, so Sabine's relation governs the whole of it. Modelled from the stated dimensions by `figs/src/CCP-018_fig04.py`; no room was built and nothing here is measured.

target `RT60` (s)	required `A` (m2 sabins)	mean absorption `alpha_bar = A/S`
0.4	37.42	0.2844
0.5	29.94	0.2276
0.6	24.95	0.1896
0.7	21.38	0.1625
0.8	18.71	0.1422
1.0	14.97	0.1138

The deliberation-room target is met at `RT60 = 0.5` to `0.6` s, requiring 25 to 30 square-metre sabins across a 131.6 m2 surface — a mean absorption coefficient of 0.19 to 0.23. That is achievable with a carpeted floor, an absorbent ceiling raft over the table, and one treated end wall; it does not require a specialist build.

18.9.4 9.4 Target (c): scheduling and cadence

The storage law is a first-order low-pass section with pole `1 - delta` and time constant `tau` of about `1/delta` periods, so **a repeated practice is a periodic forcing function on a low-pass filter, which is why cadence beats intensity** [Curzi, 2026a]. *This is a category-2 use: the same difference equation governs both objects and no physical result's authority transfers.* The essence law can set the period.

The mechanism is octave folding applied to time rather than to frequency. An essence frequency `f` has period `T = 1/f` seconds; multiplying by `2^p` — the same operation as octave transposition, since period is the reciprocal of frequency — brings `T` into a humanly usable band. Pitch class is preserved exactly, which here means **two essences that share a pitch class share a cadence, up to a doubling.** That is a defensible property for a scheduling system, because two categories that are harmonically identical should not be given unrelated rhythms.

Folded into `[1 day, 2 days)`:

class	`f` raw (Hz)	`T = 1/f` (s)	`p`	`T 2^p` (s)	days
PC1	0.483426	2.068570	16	135565.8	1.5690
PC3	4.350832	0.229841	19	120502.9	1.3947
PC5	12.085643	0.082743	20	86762.1	1.0042
PC7	23.687861	0.042216	21	88532.8	1.0247
PC9	39.157484	0.025538	22	107113.7	1.2397
PC11	58.494514	0.017096	23	143408.5	1.6598

Folded into `[7 days, 14 days)`, the more useful band for institutional cadence:

class	`p`	`T 2^p` (days)	practical reading
PC5	23	8.0335	weekly, drifting one day per cycle
PC7	24	8.1975	weekly, drifting one day per cycle
PC9	25	9.9179	ten-day
PC3	22	11.1577	eleven-day
PC1	19	12.5524	fortnightly, slightly early
PC11	26	13.2786	fortnightly, slightly early

Two things to say about this table, one useful and one honest.

The useful thing: **the spread from 8.03 to 13.28 days is a factor of 1.65, and a set of cadences in that spread drifts in and out of phase on a cycle of months.** That is a design property, not a defect. Two bodies meeting every 8.03 and 9.92 days coincide every `8.0335 x 9.9179 / (9.9179 – 8.0335) = 42.3` days, roughly every six weeks, which is a natural rhythm for a joint review. **A coincidence can be scheduled without a coordinator, by choosing periods whose beat frequency is the interval wanted.**

The honest thing: **nothing here is calendar-compatible, and calendars are strong.** A 9.92-day cycle walks through the week and lands on a Sunday about one time in seven. In practice the value is rounded to the nearest workable integer — 8, 10, 11, 14 days, or "every second Tuesday" — and at that point the essence law has been left behind and the work is simply scheduling. **Say so. A rounded cadence is a rounded cadence, and claiming it is essence-derived after rounding is FM-5.**

The relation to the filter is where this becomes engineering rather than numerology. What matters is the ratio of the forcing period `P` to the time constant `tau = 1/delta`:

`delta` (per week)	`tau` (weeks)	settling to 95% ~ `3 tau` (weeks)	`P/tau` at `P = 7.29` weeks
0.02	50.00	150.00	0.146
0.05	20.00	60.00	0.364
0.10	10.00	30.00	0.729
0.20	5.00	15.00	1.457
0.30	3.33	10.00	2.186

The rule that comes out of this is the useful one, and it does not depend on the essence law at all: choose `P` well below `tau`, ideally `P/tau <= 0.5`. A low-pass section driven well below its corner frequency integrates smoothly and holds a near-constant stock; driven near or above it, the stock sawtooths — building after each forcing and decaying visibly between them, which is the felt experience of an institution whose practices "do not stick". At `delta = 0.10` per week

a 51-day, 7.29-week cadence gives `P/tau = 0.729`, which is marginal; at `delta = 0.05` the same cadence gives 0.364, which is comfortable.

The essence law proposes the cadence; `delta` decides whether it is admissible. If the derived period exceeds `tau/2`, fold down one more octave and use the half-period. That is a legal move — it preserves pitch class exactly — and it is the correct one. Halving the PC1 cadence from 12.55 days to 6.28 days is not an abandonment of the method; it is the method's own octave operation applied to time.

For completeness, a tempo rule appears in the source material: multiply the essence frequency by 60 and halve or double until the result falls between 100 and 200, then use it as a track's beats per minute. Recomputed:

class	`f` raw (Hz)	`f x 60`	halvings	BPM
PC1	0.483426	29.0055	−2	116.0222
PC3	4.350832	261.0499	+1	130.5249
PC5	12.085643	725.1386	+2	181.2846
PC7	23.687861	1421.2717	+3	177.6590
PC9	39.157484	2349.4491	+4	146.8406
PC11	58.494514	3509.6708	+5	109.6772

This is octave folding of the period into a musically conventional band, the same operation as the day-band folding with a different target window. **It is recorded because the numbers should be available and it carries no additional claim.**

18.9.5 9.5 Target (d): tuned physical elements

Physical resonators tuned to essence classes give a room a set of objects that sound its own proportional system when struck, blown, or excited by the room itself. They are also the cheapest way to make the tuning *disclosable*: a visible object with a label on it satisfies Section 3.6 in a way a hidden loudspeaker never can.

All three relations below are ordinary physics — category 1 — and none is contributed by the essence law, which supplies only the target frequency [Rossing et al., 2002].

Open-open tubes, `f = c/(2(L + 0.8 d))` with `d` the bore diameter, for a 50 mm bore at `c = 343` m/s:

class	target `f` folded to [220,440) (Hz)	geometric length `L` (m)
PC1	247.5140	0.6529
PC3	278.4532	0.5759
PC5	386.7406	0.4034
PC7	379.0058	0.4125
PC9	313.2599	0.5075
PC11	233.9781	0.6930

Helmholtz resonators, `f = (c/(2 pi)) sqrt(A / (V (L_neck + 1.7 a)))`, neck radius `a = 0.025` m, neck length 0.10 m, targeting one octave below the folded band so that the units are practical bass-trap sizes:

class	target `f/4` (Hz)	required cavity `V` (m3)	litres
PC1	61.8785	0.010724	10.72

class	target `f/4` (Hz)	required cavity `V` (m3)	litres
PC3	69.6133	0.008473	8.47
PC5	96.6851	0.004393	4.39
PC7	94.7515	0.004574	4.57
PC9	78.3150	0.006695	6.70
PC11	58.4945	0.012001	12.00

Two of these targets coincide with unfolded lattice values. The PC11 target, 58.4945 Hz, is exactly `f(11)`; the PC3 target, 69.6133 Hz, is exactly `f(12) = k (144)`. **Both coincidences are CC-R2 doing its work: 11 and 22 share an odd part, as do 3 and 12.** A Helmholtz trap built to 58.49 Hz sits within 0.005 Hz of the room's PC11 axial mode at 58.494514 Hz, which makes it a genuinely useful modal trap and not merely a decorative one. *The source material describes the folded PC11 value as four octaves above `f(11)`; recomputation gives `233.9781/58.494514 = 4.000`, which is **two** octaves, and the correction is recorded rather than applied silently.*

Stretched strings, `f = (1/(2L)) sqrt(T/mu)`, steel with `mu = 0.0012` kg/m at `T = 400` N, so the wave speed is `sqrt(400/0.0012) = 577.35` m/s:

class	target `f` (Hz)	speaking length `L = v/(2f)` (m)
PC1	247.5140	1.1663
PC3	278.4532	1.0367
PC5	386.7406	0.7464
PC7	379.0058	0.7617
PC9	313.2599	0.9215
PC11	233.9781	1.2338

A six-string monochord with speaking lengths from 0.746 to 1.234 m is a buildable object about 1.3 m long — a bench-scale instrument that can sit along a wall, be struck by anyone, and carry a printed label giving all six frequencies, their indices, and the ordinal rule they came from. **That single object satisfies clauses 1, 2 and 5 of Section 3.6 simultaneously: it is audible, it is disclosed, and its derivation is published on it.** It is the recommended physical form for this method, in preference to any loudspeaker system, for exactly that reason.

18.9.6 9.6 Chords, beats, and what the tuning actually sounds like

Sound the six folded classes together and this is what happens, computed:

```
sorted folded frequencies (Hz):
233.9781   247.5140   278.4532   313.2599   379.0058   386.7406
```

Pairwise differences — the beat rates — for the two closest pairs:

```
379.0058 & 386.7406  ->  difference  7.7348 Hz   ROUGHNESS ZONE
233.9781 & 247.5140  ->  difference 13.5359 Hz   ROUGHNESS ZONE
all other pairs      ->  30.9 to 152.8 Hz, heard as separate pitches
```

Two of the fifteen pairs beat inside the roughness zone — a difference under about 20 Hz at these frequencies, which is within a critical band [Plomp and Levelt, 1965, von Helmholtz, 1863] — and PC7 against PC5 at 7.73 Hz is a distinctly audible flutter rather than a chord. The cause is structural: PC7's folded ratio is `49/32 = 1.531250` and PC5's is `25/16 = 1.562500`, a gap of 34.98 cents, a third of a semitone. **The two classes are structurally almost identical and the ear hears that as beating rather than as harmony.**

This is a finding about the index, not a defect in the method, and emphatically not a finding about any domain. It says that in a twelve-category partition the members at indices 5 and 7 — and 10 and 14, and so on — sit at nearly the same harmonic position. If they must be distinguishable, sound them separately, or transpose one by a full octave so the 34.98-cent gap becomes a 1234.98-cent interval and reads as an octave-plus rather than a beat. **Both are legal moves. Concealing the beat by dropping one voice below audibility is not.**

The three usable simultaneities in this set, in order of consonance:

```
PC1 + PC3     ratio 9/8,        203.91 cents   Pythagorean whole tone, +3.91 from 12-TET
PC1 + PC9     ratio 81/64,      407.82 cents   Pythagorean ditone,     +7.82 from 12-TET
PC1 + PC11    ratio 121/64,    1102.64 cents   near-just major seventh, +2.64 from 12-TET
```

All three are within 8 cents of a twelve-tone equal-tempered interval, so this set sits comfortably against conventional instruments without retuning them. That is a practical convenience and a consequence of `9/8` and `121/64` being close to tempered intervals; it is not a claim about anything. Whether a listener hears such deviations as mistunings or as category members is itself an empirical question with a substantial literature [Burns and Ward, 1978, Burns, 1999, Deutsch, 1999].

Three rendering constraints bind anything actually built.

The set size is capped, and twelve is safely inside the cap. Under a single integer octave transposition into 20 Hz to 20 kHz, at most **25 essences** fit, and only at `j = 6` (Section 8.3). A twelve-category set fits with room to spare; a thirty-two-category set does not, and must be rendered under split transposition, which changes the interval relations between the two halves and must be declared rather than quietly applied.

Do not grow the set past sixteen without expecting the flutter to get worse. The 34.98-cent gap between PC7 and PC5 is the *widest* minimum gap available: it holds for every set size from 7 to 16. **The moment a seventeenth index enters, the minimum gap collapses to 6.001 cents** — index 17 folds to `289/256`, six cents from index 3's `9/8` — and against PC3 at 278.4532 Hz that is a beat of **0.97 Hz against the present 7.73 Hz**, which is no longer a flutter but a slow swell that most listeners will hear as one wandering tone. Both remedies still apply and both cost more at seventeen than at twelve.

Every tone needs a minimum duration, and a rectangular gate is not a legal envelope. Resolving a surplus difference `dF` requires at least `1/(8 pi dF)` cycles [Gabor, 1946], independent of transposition: separating PC7 from PC5 needs **3.94 cycles**, about 10 ms at 379 to 387 Hz, while separating a 6.001-cent pair at PC3's 278.4532 Hz needs **22.96 cycles**, about 82 ms. **Specify the tone duration as a cycle count, not as a wall-clock figure**, because the cycle count is what transposes. And the envelope is constrained: a rectangular on/off gate has no finite spectral second moment, so it has no defined bandwidth and is disqualified. Use a windowed envelope; a raised-cosine rise and fall of at least one cycle at each end is sufficient and is the adopted default. **This is an engineering requirement, not an aesthetic preference, and it is entirely compatible with the disclosure constraint — a windowed envelope is audible, disclosed, and hides nothing.**

18.10 10 Design specifications

Numbered, checkable requirements. They are grouped by what they govern and prefixed accordingly: SA for the built space, CK for register crossing, ES for the essence set.

18.10.1 10.1 Requirements on a built space (SA)

SA-1 — NO SUBLIMINALS. No masked content, no content below the threshold of hearing at any listening position, no content concealed from a listener, no carrier or embedding of

any kind. **This requirement outranks every other requirement here and admits no exception** (Section 3.6).

SA-2 — Disclose in writing. A posted statement lists every tone present, its frequency, its measured level in dB SPL at the reference position, and its derivation. Posted at the entrance, readable before entry.

SA-3 — Provide an off switch and disclose its location. Anyone in the space may stop the tone. A tuned space that cannot be switched off or walked out of is outside this framework's scope.

SA-4 — Do not present an unverified device as an amplifier. Where an area of effect is required, use a loudspeaker with a published frequency response and a measured SPL, and disclose both (Section 3.6).

SA-5 — Publish the ordinal rule and the permutation sensitivity result before publishing any frequency (ES-3, ES-4).

SA-6 — Publish `(N, class, raw Hz, j, band, folded Hz)` for every tone. A folded value without its `j` is not comparable to anything (ES-10).

SA-7 — Design to `ceil(N/2)`, not to `N`. Twelve categories give six pitch classes (Theorem 7.1).

SA-8 — Check the Schroeder frequency before choosing the band. `f_s = 2000 sqrt(RT60/V)`. The whole folded set must sit on one side of `f_s`, and for speech work it must sit above it (Section 9.2).

SA-9 — Compute the full modal set, not just the axials. Report modes below 150 Hz — **with the truncation order stated**, since the count is not converged at `p,q,r <= 4` — the maximum gap, and the minimum gap. Flag any gap below 1 Hz as a near-degeneracy with its predicted beat rate (Section 9.3).

SA-10 — Meet the `RT60` target first; use the essence set only to prioritise bands. If the essence-derived proportion and the `RT60` target conflict, **`RT60` wins.** Speech intelligibility is a measured requirement; proportion is a design choice.

SA-11 — Verify that essence-derived cadences satisfy `P <= tau/2, tau = 1/delta`. If not, fold down one octave and use the half-period (Section 9.4).

SA-12 — State every rounding. A cadence rounded to a calendar, a dimension rounded to a construction module, a frequency rounded to a tuner's resolution — each is a departure from the method and each is disclosed as one, with the size of the departure in cents.

SA-13 — Prefer a labelled physical resonator to a loudspeaker. A struck object with its derivation printed on it satisfies audibility, disclosure and reproducibility at once (Section 9.5).

SA-14 — Report beat rates for every pair under 20 Hz apart. Roughness is a real perceptual effect [Plomp and Levelt, 1965] and a room that flutters must say that it flutters.

SA-15 — Restate HC-1 on every drawing, schedule and label. "Essence frequency, not physical resonance. No claim is made that any material, room, institution or person physically resonates at these values."

SA-16 — Preregister before claiming an effect. A claim that the tuning changed anything requires the full protocol of Curzi [2026i] including the shuffled-input control. Absent that, **the only claim available is that the target was chosen by a stated, reproducible method.**

18.10.2 10.2 Requirements on register crossing (CK)

CK-1. Every logarithmic quantity is published with its **base** — nats, bits, octaves, cents — in the same line or column header. A bare number is not a measurement.

CK-2. Every logarithmic quantity that is not a ratio between two reported values is published with its **reference**: `I_0`, `f_ref`, or the ground constant `k`.

CK-3. Every published quantity carries its **status badge**: PROVEN, DEFINITION, MEASURED, POSTULATE, or CHOSEN NORMALISATION. **Conversions preserve the badge.**

CK-4. The surplus-to-pitch conversion is published as `octaves = 2.885390 F` **with the note that the factor 2 is the exponent of the essence-law postulate**, not an independent constant.

CK-5. Any published essence frequency carries, in the same caption or table note: *this is an essence frequency, not a physical resonance.*

CK-6. Any essence frequency published as an audible value states the **band** and the integer `j` of the transposition `f 2^j`, and a set of values whose intervals are the point shares one `j`.

CK-7. Any surplus figure from the deployed functional is published as a **lower bound**, with "at least" or an inequality.

CK-8. Stevens' exponents are published with the **procedure**: method (estimation or production), stimulus range, instruction. An exponent without a procedure is not a measurement [Poulton, 1979].

CK-9. Weber fractions are never asserted for a modality unless measured in-house or cited. Parametric tables over `c` are the permitted substitute and are labelled parametric.

CK-10. Distortion is reported as `Delta = |ln r|` in a stated register, **never as the raw ratio `r` alone**, which is asymmetric between over- and under-framing.

CK-11. The moment equation is never published without its badge, nor without the note that `D` and `P` must share units.

CK-12. "Maximum moment `= ln N`" is published with both links of its derivation chain — the range result and the chosen normalisation — every time.

CK-13. Precision after conversion never exceeds precision before it. Report to the source's **significant figures**, not to the target register's natural resolution.

CK-14. Any table crossing registers is **computed**, not transcribed, and is reproducible from the constants of Section 4.1.

CK-15. Where a conversion depends on the essence exponent `s = 2`, the dependence is stated and the sensitivity table of Section 4.3 is available.

CK-16. No conversion between a ratio-scaled and an interval-scaled quantity without demoting the result and saying so. Specifically: **Fechner sensations may be subtracted but not divided.**

CK-17. Every formula entering a specification carries its full seven-gate validation table in print, including gates it fails (Appendix A).

CK-18. Failed and unattempted gates print as **FAIL** or **NOT YET ATTEMPTED** — never omitted, softened, or deferred to "future work" without naming the specific test.

CK-19. No intervention derived from these conversions removes a person's ability to decline it. With special force for anything delivered as tone: **audible, disclosed, refusable, with no subliminal content of any kind.**

CK-20. Any claim that two registers are "the same thing" is prohibited. They are interconvertible measures of quantities sharing a logarithmic structure. **That is the claim, and it is enough.**

18.10.3 10.3 Requirements on an essence set (ES)

ES-1. An essence set is published with its **domain and boundary** stated.

ES-2. Categories are **exhaustive and disjoint** over that domain (ES-R1), and the check is reported.

ES-3. The **ordering source** is stated and classified as (a) intrinsic magnitude, (b) generative sequence, or (c) measured rank with reliability. Type (c) reports its reliability statistic.

ES-4. The ordering is **established independently of the sonification** (ES-R2). Any set whose ordering was chosen or adjusted for its sound is disqualified.

ES-5. The **direction** of the index, ascending from ground, is stated and justified.

ES-6. Assignment is reproducible (ES-R3): the rule is stated in enough detail that a second analyst reproduces the index, and inter-rater agreement is reported where the assignment

involved judgement.

ES-7. `N`, the class count `ceil(N/2)`, and the ceiling `ln N` are published together.

ES-8. The **full index table** is published: category, index `n`, factorisation `2^a m`, class `m`, raw `f = k n^2`, class cents.

ES-9. Class collisions under 30 cents are disclosed, with the pair identified and the gap in cents.

ES-10. The **band and the transposition rule** are published, and whether a shared or per-index `j` was used.

ES-11. Every published frequency carries: *this is an essence frequency, not a physical resonance.*

ES-12. `k` is published as a **chosen** reference constant, with the note that it cancels from every interval.

ES-13. The **scale promotion** of Section 3.5 is badged wherever intervals are given a quantitative interpretation: the ordering is data, the spacing is the model.

ES-14. Consonance indices are published with their definitions and their CHOSEN status. Where two defensible indices disagree in ranking, both are published and neither is adjudicated.

ES-15. No essence set is used to rank, score, select, or exclude a person. The index is a position in a category ordering, not a measure of worth.

ES-16. Anything delivered as tone is **audible, disclosed and refusable**, with no subliminal content of any kind.

ES-17. Where a subset of a large set must be sonified, the **odd indices** are the default choice, being the class representatives and, among fully-distinct choices, the least dissonant (Section 7.7). A different subset requires a stated reason.

ES-18. Every formula in the construction carries its full seven-gate validation table in print, including gates it fails or cannot yet attempt.

18.11 11 Failure modes

Twenty-eight named failure modes, numbered FM-1 to FM-28 and grouped by subject. Each gives what it damages, its signature, and its fix. *This catalogue supersedes the source material's three separately numbered lists; the numbering here is sequential and the cross-references in this paper follow it.*

18.11.1 11.1 Faults of the built space

FM-1 — THE HIDDEN TONE. *Damages:* consent, and with it the entire ethical standing of the work. *Description.* A tone, a message, or a modulation is placed below the threshold of audibility, masked under a louder signal, or embedded in a carrier. *Signature.* Any signal path a person in the room cannot hear and was not told about; any spectrum-analyser trace showing content not on the posted disclosure. *Fix.* SA-1, SA-2, SA-3. **There is no partial fix and no acceptable dose. Remove it, or make it audible and disclosed.**

FM-2 — THE RESONANCE CLAIM, BUILT. *Damages:* everything, and it is easiest to commit here because the room really does resonate. *Description.* "The room resonates at the essence frequency of category *X*." The room resonates at 39.157484 Hz because it is 4.3798 m wide. The category does not resonate at anything. *Signature.* A category, a discipline, a material or a person as the grammatical subject of "resonates". *Fix.* SA-15, and rewrite: "the width was chosen so that the axial mode falls at the essence frequency of index 9."

FM-3 — THE UNBUILDABLE ROOM. *Damages:* the project. *Description.* A class triple is chosen for harmonic reasons and produces dimensions of, say, `0.74 x 3.60 x 10.00` m — a ceiling height of 74 centimetres and an aspect ratio of 13.4 : 1. That is exactly what PC11 : PC5

: PC3 gives if the raw class ratios are used with no octave offset, scaled so the middle axis is 3.60 m. **The ratio `1 : 121/25 : 121/9` is harmonically legitimate and architecturally impossible.** *Signature.* Any dimension below about 2.4 m or above about 5 m for a room-scale build; an aspect ratio beyond about 1 : 3. *Fix.* Search the class triples *and* the octave offsets, as Section 9.3 does, before committing. If no combination is buildable, **the honest report is that this category scheme does not produce a room**, and the set is used for cadence or tuned elements instead.

FM-4 — THE MODAL BLIND SPOT. *Damages:* the acoustic outcome, badly. *Description.* Axial modes are placed by essence and the tangential and oblique modes are never computed. Two of them land 0.5 Hz apart and the room throbs. *Signature.* A room-mode analysis containing only `(p,0,0)`, `(0,q,0)`, `(0,0,r)` entries. *Fix.* SA-9. Compute the full set to a stated truncation order and report the gaps — **and report the truncation order, because the mode count below a cutoff is not converged at `p,q,r <= 4`.**

FM-5 — THE ROUNDED CADENCE, UNDISCLOSED. *Damages:* the reproducibility of the whole schedule. *Description.* An essence-derived 9.9179-day cadence is rounded to "every second Tuesday", 14 days, and the schedule is still described as essence-derived. It is not: 14 days is 1.4116 times the derived value, a departure of **596.8 cents**. *Signature.* A published cadence at a round calendar interval with an essence derivation attached. *Fix.* SA-12. State the derived value, the rounding, and the size of the departure. **A rounded cadence is a perfectly good cadence; it is just not the derived one.**

FM-6 — PROPORTION OVER INTELLIGIBILITY. *Damages:* speech, turn-taking, and therefore the room's contribution to collective performance [Woolley et al., 2010]. *Description.* An essence-derived proportion is preserved at the cost of an `RT60` that lands at 1.2 s in a deliberation room, because absorbing enough to reach 0.5 s would require covering a surface the designer wanted left hard for "proportional integrity". *Signature.* A measured `RT60` above 0.8 s in a room intended for group discussion, with a design rationale referring to proportion. *Fix.* SA-10. **`RT60` wins. Every time.** The loudest speaker wins a reverberant room by physics, and no proportional system compensates for that [Curzi, 2026e].

FM-7 — THE BAND STRADDLING SCHROEDER. *Damages:* the coherence of the whole specification. *Description.* A folded band is chosen that spans `f_s`, so half the targets are modal and half statistical, and the treatment strategies for the two halves are incompatible. *Signature.* A folded set spanning, say, 100 to 300 Hz in a room whose `f_s` is 147 Hz. *Fix.* SA-8. Compute `f_s` first, then choose the band, then fold.

FM-8 — THE UNDISCLOSED LOUDSPEAKER. *Damages:* consent, indirectly but seriously. *Description.* The tuning is delivered by a ceiling speaker nobody knows about, at a level low enough that most occupants do not consciously notice it. Nothing is technically subliminal — it is above threshold — but nothing is disclosed either. *Signature.* An occupant who cannot say what is sounding in the room they work in. *Fix.* SA-2 and SA-13. **The disclosure requirement is not satisfied by audibility alone. A person must be able to know, not merely able in principle to notice.**

FM-9 — THE PHANTOM SIX. *Damages:* the whole composition. *Description.* A twelve-category scheme is realised as twelve tuned tubes. Six of them are octave duplicates. The room contains twelve objects and six sounds, and the extra six are presented as distinct voices. *Signature.* Two tubes whose lengths stand in a ratio of exactly 2 or 4. *Fix.* SA-7. Build six, label each with the categories it carries.

FM-10 — THE UNBEATEN BEAT. *Damages:* the occupants' experience, and the credibility of the tuning. *Description.* PC5 and PC7 are sounded together, beat at 7.73 Hz, and the room is described as "harmonious". Occupants describe it as "buzzing" and nobody connects the two reports. *Signature.* Any published set with a folded pair under 20 Hz apart and no beat rate reported. *Fix.* SA-14. Report the beat, then either separate the voices in time, transpose one by an octave, or accept the flutter and say so in the disclosure.

18.11.2 11.2 Faults of register crossing

FM-11 — THE UNSTATED BASE. *Signature:* a surplus, sensation or pitch figure published as a bare number; readers assume bits where nats were meant. *Damage:* a silent factor of 1.442695 or 0.693147 through every derived quantity; in headroom calculations this is the difference between a positive and a negative margin. *Detection:* audit any figure whose ratio to a companion figure sits near 1.4427 or 0.6931. *Fix:* CK-1.

FM-12 — THE FREE LUNCH CONVERSION. *Signature:* a modest surplus gain displayed in hertz, where it looks like a transformation. “We moved from 2 Hz to 20 Hz” is `N_eff` from 2.03 to 6.43 — **1.15 nats**, real but a third of what the hertz figure implies. *Damage:* over-claiming, then loss of credibility when the figure is recomputed. *Detection:* recompute in nats. If the nat figure is unimpressive, the hertz figure was theatre. *Fix:* nats or bits in the argument; hertz only in the display.

FM-13 — THE PHYSICAL RESONANCE SLIP. *Signature:* “the system resonates at”, “tuned to the frequency of”, “vibrates at”, applied to an essence frequency. *Damage:* **the framework is correctly reclassified as pseudoscience and the genuine results are discarded with the error.** *Detection:* search for “resonat”, “vibrat”, “frequency of the”; every hit must concern an actual physical system or carry the disclaimer. *Fix:* CK-5, without exception and without weariness.

FM-14 — THE BORROWED EXPONENT. *Signature:* a Stevens exponent measured by one procedure, imported into a specification using a different range, instruction, or modality subtype. *Damage:* the key is wrong by an unknown amount, and because the error is in an exponent it compounds with the range. *Detection:* ask what procedure produced it; “the literature” means borrowed. *Fix:* CK-8 [Teghtsoonian, 1971].

FM-15 — THE LINEAR PROMISE. *Signature:* a commitment to “twice the brightness” or “twice the loudness”, costed as though the physical quantity doubles. *Damage:* the real requirement is `2^(1/a)` — for brightness, 8.17 times. Budgets built on the linear assumption are short by a factor of four. *Detection:* any specification whose perceptual target and physical budget are the same multiple. *Fix:* **never promise a linear improvement in a quantity the recipient perceives logarithmically.**

FM-16 — THE UNBOUNDED MOMENT. *Signature:* a claim that experience can be arbitrarily enriched by intensity, novelty or attention, without reference to the perceiver’s category count. *Damage:* effort spent past the point where `ln N` binds, with no explanation available for the vanishing returns. *Detection:* compute `F_int / ln N`; above about 0.9 the ceiling binds and geometry offers nothing. *Fix:* raise `N` and accept that the gain is logarithmic. Distinguish from the *raw* `M`, which is unbounded as `P -> 0` and whose divergence is pathology.

FM-17 — THE CENTS ILLUSION. *Signature:* a surplus estimated to two significant figures, reported as 4747.67 cents. *Damage:* six figures of apparent precision on two figures of information; readers treat the difference between 4747 and 4750 cents as meaningful, when it is the fourth decimal of `u`. *Detection:* compare significant figures in source and conversion. *Fix:* CK-13. Convert, then round to the source’s precision.

FM-18 — THE MISSING ZERO. *Signature:* division of two quantities lacking a true zero — most often ordinal ratings treated as `D` and `P`. *Damage:* the ratio is arbitrary, its logarithm is arbitrary, and every downstream conversion is arbitrary while looking rigorous. *Detection:* ask what zero would mean for each term and whether the instrument can produce it [Stevens, 1946, Michell, 1997]. *Fix:* build a ratio scale, or demote the formula to a heuristic ordering **in print**.

FM-19 — THE STATUS UPGRADE. *Signature:* a postulate stated in nats in one section, cited in octaves in the next, treated as established by the third. *Damage:* a design decision acquires the authority of a theorem. *Detection:* trace any load-bearing constant to its first appearance and check the badge there. *Fix:* CK-3. **The badge travels with the number, not with the sentence.**

18.11.3 11.3 Faults of the essence set

FM-20 — THE TUNED ORDERING. *Signature:* an essence set whose index assignment was revised after the first listening. *Damage:* ES-R2 is void; the frequencies encode the analyst's preference and nothing else, and every downstream interval is an artefact. *Detection:* ask for the date and provenance of the ordering relative to the first sonification. *Fix:* none available after the fact. Re-derive from an external source and discard the tuned version. **This is the essence-set form of specification search and it is disqualifying** [Simmons et al., 2011].

FM-21 — THE ARBITRARY ORDINAL. *Signature:* categories indexed alphabetically, by order of discovery, or by the order they appeared in a spreadsheet. *Damage:* the lattice is arbitrary; a different arbitrary ordering gives different frequencies with equal claim to correctness. **This is the scale-type objection of Section 3.5 in operational form: an index with no intrinsic magnitude behind it is ordinal at best, and a square law on an ordinal index is not a licensed operation** [Michell, 1997]. *Detection:* apply a second plausible ordering and compare; if the intervals change materially, the ordering was doing no work. *Fix:* find an intrinsic magnitude or a generative sequence, or decline to sonify.

FM-22 — THE OVERLAPPING CATEGORY. *Signature:* an entity in the domain belonging to two categories. *Damage:* ES-R1 fails; the entity has two indices, two frequencies, and a spurious interval between itself and itself. *Detection:* attempt to classify twenty randomly drawn cases; any double assignment is a hit. *Fix:* split or merge categories until the partition is clean, then recompute `N` and every index.

FM-23 — THE CLASS COLLISION READ AS MEANING. *Signature:* two categories whose tones nearly coincide, interpreted as evidence that the categories are related. *Damage:* classes 3 and 17 are 6.00 cents apart because `9/8` and `289/256` are nearly equal. **That is a fact about integers, not about any domain. Reading it as a domain finding is numerology.** *Detection:* check whether the collision is predicted by the table of Section 7.4 for any set of that size; if it is, it is structural. *Fix:* disclose all structural collisions in advance (ES-9), so that none can later be mistaken for a result.

FM-24 — THE OVERGROWN SET. *Signature:* an essence set with thirty or forty categories. *Damage:* the top of the lattice is congested with adjacent intervals under 150 cents, collisions multiply, the class count grows linearly while the audible span does not, and the ceiling gain from the extra categories is logarithmic and negligible. *Detection:* compute the adjacent-interval table; any adjacent pair under about 100 cents indicates the set has outgrown the lattice. *Fix:* aggregate to a coarser, defensible partition and rebuild the index from the same ordering source.

FM-25 — THE CONSONANCE CHASE. *Signature:* a set curated toward consonant intervals, or a subset chosen because it "sounds harmonious". *Damage:* consonance is redundancy in this lattice (Section 7.7). The most consonant six-set collapses to two pitch classes. **A beautiful-sounding essence set is very often a collapsed one, worth `ln 2` where it should be worth `ln 6`.** *Detection:* count distinct classes in the chosen subset; fewer than `ceil(size/2)` means collapse. *Fix:* ES-17. Take the odd indices and understand that the dissonance is the sound of the set actually carrying information.

FM-26 — THE SCALE PROMOTION, UNBADGED. *Signature:* intervals from an ordinal-indexed set interpreted as measured magnitudes of relation. *Damage:* an unearned upgrade from rank data to a ratio-scale claim; the formula looks like a measurement and is a model. *Detection:* ask what would change if the spacing postulate were `n^1.5` instead of `n^2`. If the answer is "every interval" and the text does not say so, the badge is missing. *Fix:* ES-13, and the sensitivity table of Section 4.3.

FM-27 — THE LOST TRANSPOSITION. *Signature:* a published audible frequency with no `j` and no band. *Damage:* the raw essence frequency cannot be recovered, so the value cannot be checked or compared with another set. *Detection:* attempt to recompute `n` from the published hertz; if `sqrt(f/k)` is not near an integer for any small power of two, `j` is missing. *Fix:* ES-10.

FM-28 — THE MIXED TRANSPOSITION. *Signature:* a set of values folded into one octave with per-index `j`, then compared by interval. *Damage:* all inter-index intervals are destroyed;

a 1.15-octave relationship reads as 0.15 octaves. *Detection:* check whether `j` varies across the compared values. *Fix:* recompute with a shared `j` and a band wide enough to hold the span.

18.12 12 Results: three systems worked end to end

Every number below is computed from the stated parameters under the stated model. **No dataset was collected. Nothing in this section is a measurement of any real institution, and the third case is a refusal rather than a result.**

18.12.1 12.1 A twelve-category deliberation room

A research institute organised into twelve houses commissions a deliberation room for its cross-house committee.

Step 1 — Declare the category set and the ordinal rule. The twelve categories, in the order the institute's own published enumeration gives them:

n	Category	Domain
1	Mathematics and Pattern	Formal structure, number, proof
2	Physics and Energy	Matter, force, measurement
3	Biology and Body	Life, organism, genetics
4	Psychology and Mind	Cognition, emotion, development
5	Consciousness and Field	Awareness, observation, qualia
6	Communication and Code	Value, exchange, accounting
7	Technology and Craft	Tools, systems, engineering
8	Spirituality and Practice	Contemplation, tradition, experience
9	Art and Expression	Aesthetics, creation, performance
10	Ecology and Earth	Ecosystem, soil, climate
11	Language and Story	Grammar, narrative, translation
12	Ethics and Governance	Law, sovereignty, distribution

Ordinal rule, as declared: the index is the category's position in that enumeration. **This is a weak ordinal and is reported as one.** The source of the partition is explicit that twelve "is not asserted as correct" and that the partition is "one idealised coding, not a uniquely privileged carving of the world". The enumeration order is not derived from any property of the categories. It is publishable, reproducible and stable — anyone with the same document gets the same integers — **but it is not intrinsic, and it is ordering source type (b) at best and arguably neither (a), (b) nor (c).** ES-3 requires this to be stated and it is stated here.

Step 2 — Permutation sensitivity (ES-4). Three plausible re-orderings:

```
swap 3 <-> 12   : both odd part 3  -> pitch-class displacement    0.00 cents
swap 5 <-> 7    : class 5 <-> 7    -> displacement  772.63 - 737.65 =  34.98 cents
swap 9 <-> 11   : class 9 <-> 11   -> displacement 1102.64 - 407.82 = 694.82 cents
```

Maximum displacement 694.82 cents — well over a 100-cent threshold. The ordering is load-bearing. Since the ordinal rule is weak, this is the honest finding and it travels with every number below: **a different but equally defensible enumeration of the twelve categories would produce a materially different tuning. The room is being built to *an* essence set, not to *the* essence set.**

Steps 3 to 5 — Frequencies, classes, folding. `f(n) = 0.483425732988414 n^2` with `k` **defined** and its hertz **postulated**; `N = 12` so `ceil(12/2) = 6` classes; band `[220, 440)` Hz with per-index `j` published per value (SA-6).

n	Category	`f(n)` (Hz)	j	folded (Hz)	nearest note (A4=440)	dev (cents)	odd part	class
1	Mathematics and Pattern	0.483426	+9	247.5140	B3	+4.0	1	PC1
2	Physics and Energy	1.933703	+7	247.5140	B3	+4.0	1	PC1
3	Biology and Body	4.350832	+6	278.4532	C#4	+7.9	3	PC3
4	Psychology and Mind	7.734812	+5	247.5140	B3	+4.0	1	PC1
5	Consciousness and Field	12.085643	+5	386.7406	G4	−23.4	5	PC5
6	Communication and Code	17.403326	+4	278.4532	C#4	+7.9	3	PC3
7	Technology and Craft	23.687861	+4	379.0058	F#4	+41.7	7	PC7
8	Spirituality and Practice	30.939247	+3	247.5140	B3	+4.0	1	PC1
9	Art and Expression	39.157484	+3	313.2599	D#4	+11.8	9	PC9
10	Ecology and Earth	48.342573	+3	386.7406	G4	−23.4	5	PC5
11	Language and Story	58.494514	+2	233.9781	A#3	+6.6	11	PC11
12	Ethics and Governance	69.613306	+2	278.4532	C#4	+7.9	3	PC3

The six pitch classes are the entire harmonic vocabulary of this room:

class	categories carrying it	`f(m) = k m^2` (Hz)	j	folded (Hz)	note	dev
PC1	Mathematics; Physics; Psychology; Spirituality	0.483426	+9	247.5140	B3	+4.0
PC3	Biology; Communication; Ethics	4.350832	+6	278.4532	C#4	+7.9
PC5	Consciousness; Ecology	12.085643	+5	386.7406	G4	−23.4
PC7	Technology	23.687861	+4	379.0058	F#4	+41.7
PC9	Art	39.157484	+3	313.2599	D#4	+11.8
PC11	Language	58.494514	+2	233.9781	A#3	+6.6

Twelve categories, six voices, and the distribution is lopsided. Four categories share PC1, three share PC3, two share PC5, and PC7, PC9 and PC11 are singletons. **This is a structural statement about the enumeration used: the categories at powers-of-two positions 1, 2, 4 and 8 are harmonically indistinguishable from one another.** It is not a statement about mathematics, physics, psychology or spirituality; **it is a statement about positions 1, 2, 4 and 8 in a list** (FM-23).

Sorted by pitch, with the intervals between adjacent classes:

#	class	folded Hz	note	step to next (cents)
1	PC11	233.9781	A#3 +7	97.4
2	PC1	247.5140	B3 +4	203.9
3	PC3	278.4532	C#4 +8	203.9
4	PC9	313.2599	D#4 +12	329.8
5	PC7	379.0058	F#4 +42	35.0
6	PC5	386.7406	G4 −23	330.0 (wrapping to 2x PC11)

PC11 has folded to a position *below* PC1, because `121/128 = 0.945` is less than one. Folding into a fixed band reorders. The class labels are still correct; the ordering is not the index ordering, and comparing folded values across systems is FM-28. Within one room this is simply the scale.

Step 6a — The room proportion. Exhaustive search over class triples and octave offsets against the published ratio heuristics gave PC7 : PC9 : PC11 against Volkmann as the best

match, Euclidean error 0.031. Anchoring the axial fundamentals at the raw essence frequencies:

```
height  Lx = c / (2 f(11)) = 343 / (2 * 58.494514) = 2.9319 m
width   Ly = c / (2 f(9))  = 343 / (2 * 39.157484) = 4.3798 m
length  Lz = c / (2 f(7))  = 343 / (2 * 23.687861) = 7.2400 m

ratio 1 : 1.4938 : 2.4694,  exactly 1 : 121/81 : 121/49
Volkmann heuristic 1 : 1.5 : 2.5     deviation -0.0062, -0.0306

V = 92.969 m^3
floor area = 31.709 m^2
total surface S = 131.555 m^2
occupancy at 2.5 m^2/person = 12.7 people
```

Twelve categories; twelve-and-a-bit seats. The room is the size of its own committee, and that was not arranged — it fell out of `c/(2f)`. It is also a coincidence, and Section 13.2 says why nothing may be inferred from it.

A room 4.38 m wide and 7.24 m long is a poor fit for a round table, the aspect ratio being 1.65, and a good fit for a long oval or an elongated polygon. **At 7.24 m the far-end sightline is at the outer limit of the published social-distance zone [Hall, 1966] and past comfortable facial-expression reading.** The honest resolution is to seat the twelve within the middle 5.5 m of the length and use the end 1.7 m for circulation and the monochord. **The proportional system proposed the room; governance geometry disposes of how it is furnished.**

Step 6b — Acoustic specification (SA-8, SA-10).

```
deliberation-room target: RT60 = 0.5 to 0.6 s

RT60 = 0.5 s  ->  A = 0.161 * 92.969 / 0.5 = 29.94 m^2 sabins  ->  alpha_bar = 0.2276
RT60 = 0.6 s  ->  A = 0.161 * 92.969 / 0.6 = 24.95 m^2 sabins  ->  alpha_bar = 0.1896

Schroeder frequency at RT60 = 0.5 s:  f_s = 2000 * sqrt(0.5/92.969) = 146.7 Hz
```

Check SA-8: the folded set spans 233.98 to 386.74 Hz, entirely above `f_s = 146.7` Hz. The whole set is in the statistical region and Sabine governs for all of it. The band is coherent. Band priority: four of six classes in the 250 Hz octave band, two in the 500 Hz band, so **the treatment budget is weighted toward the 250 Hz band, where porous absorbers are losing efficiency** — meaning at least one tuned membrane or Helmholtz element, not carpet alone. Step 6d supplies it.

Step 6c — Cadence (SA-11). The committee is a cross-category body, so its cadence is taken from PC1, the class carried by the largest number of categories.

```
f(PC1) = 0.483426 Hz  ->  T = 2.068570 s
folded into [7, 14) days: p = 19, T * 2^19 = 12.5524 days

institute delta, stipulated: 0.05 per week
tau = 1/delta = 20 weeks
P = 12.5524 days = 1.7932 weeks
P/tau = 1.7932 / 20 = 0.0897
```

`P/tau = 0.090`, far below the 0.5 requirement of SA-11. Admissible without folding down. The forcing period is well inside the section's time constant, so the stock integrates smoothly rather than sawtoothing between meetings [Curzi, 2026a].

Second cadence, for a sub-committee taking PC7:

```
f(PC7) = 23.687861 Hz -> T = 0.042216 s -> p = 24 -> 8.1975 days
P/tau  = 8.1975/140 = 0.0586        (140 days = 20 weeks)
```

Beat between the two cadences: `12.5524 x 8.1975 / (12.5524 - 8.1975) = 23.63` days. The two bodies coincide about every 24 days, roughly monthly, and that coincidence is a scheduling fact produced by the periods with no coordinator required.

Rounding, disclosed (SA-12): in practice both cadences are rounded to the working calendar. 12.5524 days becomes a fixed twice-monthly slot, average 15.2 days, a departure of `1200 log2(15.2/12.5524) = 331` cents; 8.1975 days becomes a fixed weekly-plus-one slot of 8 days, a departure of 42 cents. **Those are the actual cadences. The derived values are the design intent and the rounding is declared. Anything else would be FM-5.**

Step 6d — Tuned physical elements (SA-13). *One six-string monochord*, mounted along one wall, steel strings `mu = 0.0012` kg/m at 400 N so `v = 577.35` m/s:

```
PC11  233.9781 Hz  ->  L = 1.2338 m
PC1   247.5140 Hz  ->  L = 1.1663 m
PC3   278.4532 Hz  ->  L = 1.0367 m
PC9   313.2599 Hz  ->  L = 0.9215 m
PC7   379.0058 Hz  ->  L = 0.7617 m
PC5   386.7406 Hz  ->  L = 0.7464 m
```

Overall instrument length about 1.40 m. **Each string carries an engraved label: its class, the categories it represents, its frequency to four figures, its essence index, and the sentence required by SA-15.**

One Helmholtz absorber tuned to PC11, neck radius 25 mm, neck length 100 mm, cavity 12.00 litres, resonating at 58.4945 Hz. **This sits within 0.005 Hz of the room's `(1,0,0)` axial mode at 58.494514 Hz, so it is a working modal trap for the height axis as well as a tuned element. It is the cheapest single item in the specification and the most acoustically useful.**

Beat report (SA-14): PC7 and PC5 differ by 7.7348 Hz and PC11 and PC1 by 13.5359 Hz. **Both are in the roughness zone. If all six strings are struck together the result flutters at about 8 Hz and again at about 13 Hz.** This is disclosed on the instrument's label. Occupants wanting a clean chord strike PC1, PC3 and PC9 — the `1 : 9/8 : 81/64` Pythagorean triad, all within 8 cents of twelve-tone equal temperament.

Step 7 — The disclosure notice (SA-1, SA-2, SA-3, SA-15). Posted at the entrance, reproduced here in full because it is part of the specification rather than an appendix to it:

```
THIS ROOM IS PROPORTIONALLY TUNED. NOTHING IN IT IS HIDDEN.

Dimensions:  2.9319 m high, 4.3798 m wide, 7.2400 m long.
These were chosen so the room's three axial acoustic modes fall at
58.4945 Hz, 39.1575 Hz and 23.6879 Hz. The ratio is exactly 1 : 121/81 : 121/49.

A six-string monochord on the north wall is tuned to
233.98, 247.51, 278.45, 313.26, 379.01 and 386.74 Hz.
It sounds only when someone plays it. Anyone may play it. Anyone may stop.
Struck together, strings 5 and 6 beat at about 8 Hz and strings 1 and 2 at
about 13 Hz. This is expected.

A 12-litre Helmholtz absorber behind the north panel is tuned to 58.49 Hz.
It is passive. It makes no sound of its own.

There is no loudspeaker in this room and no continuous tone of any kind.
There is no masked, hidden, subliminal, or below-threshold content of any kind
in this room. If you cannot hear it, it is not here.
```

```
Derivation: the twelve houses of the institute, indexed by their position in the
published enumeration, through f(n) = k n^2 with k = 0.483425732988414 Hz.
Twelve categories yield six distinct pitch classes.
The full derivation is in the room file and anyone may have a copy.

These are ESSENCE frequencies, not physical resonances. No claim is made that
any material, room, institution or person physically resonates at these values.
The room's acoustic modes are ordinary acoustics: a room 4.3798 m wide has a
mode at 39.16 Hz because sound travels at 343 m/s, and for no other reason.

The indexing of the twelve categories is not intrinsic. A different but equally
defensible ordering would give a different tuning. This is one essence set,
not the essence set.
```

Step 8 — What would falsify a claim about this room (SA-16). No claim is made that the tuning improves anything. Should one be made, the protocol is fixed in advance: measure turn-taking equality [Woolley et al., 2010] in the old room and the new one, with the minimal estimation model `Y_{t+1} = a Y_t + b F(U_t) − C_t + eps` fitted on the first eight committee cycles and frozen, reported on cycles 9 to 12 against all four required baselines — persistence, direct input, matched autoregressive, and **the shuffled-input control** [Curzi, 2026i].

The prediction that would matter, and the one hardest to get: turn-taking equality improves in the new room *and* the improvement survives the shuffled-input control. If it does not survive, the honest conclusion is that the room is a good room because its `RT60` is 0.5 s and its sightlines work, and that the essence tuning contributed a defensible way of choosing among many acceptable proportions and nothing further. **That would be a perfectly good result and this paper would report it as one** — and it is, in fact, what H2 predicts. Note also that such a test is unblindable by construction (Section 3.4), because the room's disclosure notice names the tuning.

18.12.2 12.2 A metropolitan transit review board, in four registers

A municipal transit authority convenes a review board on a proposed bus-priority corridor. The board has **`N = 8` declared professional categories**: traffic engineering, urban planning, transit operations, accessibility, municipal finance, local business representation, environmental assessment, and legal and regulatory.

Non-parallelism is stipulated at `u = 0.42` by a disagreement-rate procedure across 40 recorded recommendations. **This is an in-example figure produced by a stated procedure in a constructed case, not a measurement of any real board.** Storage parameters, likewise stipulated: `eta = 0.35`, `delta = 0.08` per quarter, `Q = 4.0`, `C = 0.30`.

Step 1 — Effective count and surplus.

```
N_eff = 1 + (N − 1) u = 1 + 7 ∗ 0.42 = 3.94
F     = ln 3.94       = 1.371181 nats
```

Eight declared categories are worth 3.94 genuinely different ones. More than half the board's declared variety is redundant, which is the ordinary finding and not a scandal. **`N_eff = 3.94` here is the affine effective count `1 + (N−1)u`, not a Hill number** (Section 4.6); a Hill number of order two on the same roster would in general return a different value and the two must not be compared.

Step 2 — The four registers.

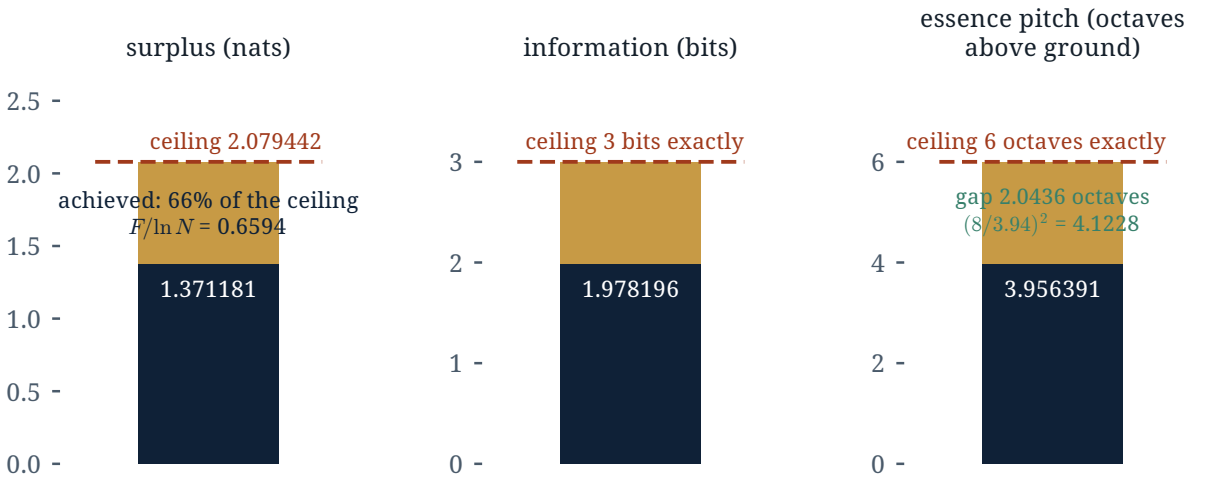


Figure 18.5: **One shortfall, said three ways: 66 per cent of the ceiling in nats and in bits is two full octaves below it in essence pitch.** With `N = 8` and `u = 0.42` the affine effective count is 3.94 — **not a Hill number** — giving 1.371181 nats, 1.978196 bits and 3.956391 octaves against a ceiling of `ln 8` = 2.079442 nats, 3 bits exactly and 6 octaves exactly. The pitch register amplifies what the surplus register compresses: `(8/3.94)^2` = 4.1228. Three panels because three units; never a dual axis. Model output from the stipulated parameters of Section 12.2, generated by `figs/src/CCP-018_fig05.py`.

register	value	reading
surplus	at least 1.371181 nats	lower bound on this board's per-session information
information	at least 1.978196 bits	just under two bits: the board resolves fewer than four states
essence pitch	3.956391 octaves above ground	
cents	4747.7 cents	reported as 4750 per CK-13, since `u` is stated to two figures
essence frequency	`f = k N_eff^2 = 7.504508` Hz	**essence frequency, not a physical resonance**
audible placement	120.07 Hz, band [110, 220) Hz, **j = 4**	

The bit figure is the one to put in front of the board. Two bits means this body, meeting as it currently meets, distinguishes fewer than four genuinely different states of the world per session.

Step 3 — The ceiling and the gap.

```
ceiling   ln N = ln 8 = 2.079442 nats = 3.000000 bits exactly = 6 octaves exactly
achieved  F/ln N = 1.371181 / 2.079442 = 0.6594
gap       ln N − F = 0.708261 nats = 2.0436 octaves = 2452 cents
```

The board is at **66 per cent of its own ceiling.** The ceiling frequency is `k (8)^2 = 30.939247` Hz, and the ratio of ceiling to achieved frequency is `(8/3.94)^2 = 4.1228` — **the board is two full octaves below its own maximum**, the same fact as "34 per cent of the ceiling unused", said in the register that makes it audible.

Step 4 — The storage arithmetic [Curzi, 2026a].

```
headroom      h  = eta*ln N - delta*Q - C
                 = 0.35 * 2.079442 - 0.08 * 4.0 - 0.30
                 = 0.727805 - 0.320000 - 0.300000
                 = +0.107805        h > 0: growth is available in principle

actual step      = eta*F - delta*Q - C
                 = 0.35 * 1.371181 - 0.320000 - 0.300000
                 = 0.479913 - 0.620000
                 = -0.140087        Q is DECLINING
```

Here is the diagnostic the conversion keys make visible: positive headroom, negative actual step. The ceiling permits growth; the geometry does not deliver it. In this constructed case the authority loses adopted structure at 0.14 nat-equivalents per quarter while believing itself functional, because nothing in its minutes records a quantity that would show the decline.

Step 5 — What is required. A non-decreasing step requires

```
S >= (delta*Q + C) / eta = (0.320 + 0.300) / 0.35 = 1.771429 nats
Feasible?  1.771429 <= ln 8 = 2.079442      YES, with 0.308 nats to spare
```

That feasibility test is the admissibility ceiling in its general form, and it is worth naming because of what it does not contain. The threshold rate a system must sustain is `c/eta` and it is admissible if and only if `c/eta <= ln N`; the same quantity re-expressed as a minimum sustaining frequency, `f_min = k e^(2 c/eta)`, **contains neither `delta` nor `N`** — retention loss sets the rate of arrival but not the threshold, and the category count enters only through the ceiling [Curzi, 2026]. Here `c/eta = 0.857143` in the cost-only reading and `1.771429` including the decay charge, both below `ln 8`, so the board is admissible on both readings.

Solving back for the geometry:

```
N_eff required = e^1.771429 = 5.879246
u    required  = (5.879246 - 1) / 7 = 0.697035
```

`u` must rise from 0.42 to 0.70 — a climb of `2 (1.771429 - 1.371181)/ln 2 = 1.154871` octaves, or 1385.85 cents, reported as 1390 per CK-13, with the essence frequency rising from 7.504508 Hz to 16.709870 Hz, a factor of 2.2266.

Step 6 — Specification and prediction. The binding constraint is `u`, not `N`. Adding a ninth category raises the ceiling from `ln 8 = 2.0794` to `ln 9 = 2.1972` nats — **0.1178 nats, one third of a bit, for the cost of an entire additional seat** — while doing nothing about the redundancy among the eight already present. **That is the same lever, computed and declined, that the corpus's intervention procedure reaches by a different route**: a ceiling gain of `eta(ln 6 - ln 5) = +0.1914` nats set against five added communication paths costing `+0.0714`, a 34 per cent haircut on the headline benefit [Curzi, 2026j]. The remedies that act on `u` rather than `N` are rotation across lineages, a standing adversarial role, a member from an unrelated trade, and an explicit prohibition on the pre-meeting alignment call that drives `u` toward zero before the session begins. **What that costs is priced by the marginal identity: the law values the two sources of non-parallelism identically, `df/du_cross = df/du_div = (N-1)/(1 + (N-1)u)`, while their productivity in the design variables differs sharply** [Curzi, 2026d].

Predicted outcome, stated so it can fail: if `u` rises to 0.70 or above and `eta`, `delta`, `C` hold, the quarterly step turns non-negative within two quarters and the stock stabilises near `Q_inf = (eta F - C)/delta = (0.35 x 1.771429 - 0.300)/0.08 = 4.000000` — **exactly its current level, because `S` was solved to make the step zero. The intervention stops the decline; it does not build.** Building requires raising `eta` or cutting `C`, which is a separate decision.

Falsification: if `u` measures above 0.70 for two consecutive quarters and the stock still declines, then either `eta` is below 0.35, or `C` is above 0.30, or the linear storage model is wrong for

this body. All three are checkable; the third is the one that would matter. **And a warning that outranks all three: the endogenous social effect in a body whose members both influence and constitute its output is not identified from observation** [Manski, 1993], so a fit that recovers `eta` has not measured it (Section 3.4).

Step 7 — What the sonification is for. If this board's trajectory is sonified, both values take **the same** transposition, `j = 4`: 120.07 Hz now, 267.36 Hz at the target, band `[110, 440)` Hz. The tone rises by 1.1549 octaves and the change is immediately audible where a table of nats is not.

The shared `j` is not a detail. Placing each value in a narrow band with its own `j` — 120.07 Hz at `j = 4` and 133.68 Hz at `j = 3` — preserves each pitch class but destroys the interval between them, collapsing 1.15 octaves to 0.15 (FM-28). Octave transposition preserves pitch class exactly and preserves intervals only when the transposition is shared. **Publish the band and the `j`, and use one `j` across any set of values whose intervals are the point** (CK-6).

Sonification is a display, and displays are for perception. It is not evidence. Nothing about the transit authority vibrates. The tone is an essence frequency computed from a stipulated geometry by a postulated square law with a chosen reference constant, and every one of those four qualifications belongs in the caption. The board should hear it, understand exactly what it is, and be free to ignore it.

18.12.3 12.3 A proxemic-zone lattice, and a refusal

The admissible case. Hall's *The Hidden Dimension* divides interpersonal distance into four zones — intimate, personal, social, public — and subdivides each into a close and a far phase, giving **eight categories ordered by physical distance** [Hall, 1966]. Checking the requirements of Section 8.1:

- **ES-R1 (exhaustive, disjoint):** the zones tile the distance axis from zero upward without gaps or overlaps. **PASS**, subject to the boundary being "interpersonal distance in a face-to-face encounter", which is stated.
- **ES-R2 (ordering from outside):** the ordering is by **measured distance**, an intrinsic physical magnitude, published in 1966 with no relation to any sonification. This is ordering source type (a), the strongest. **PASS.**
- **ES-R3 (reproducible assignment):** given a measured separation, the zone assignment is mechanical. **PASS.** Hall's zone *boundaries* are nominal and culturally variable — Hall said so — but the *ordering* is not in question, and only the ordering is used.

Direction: index 1 is the ground, assigned to the nearest zone. `N = 8`; class count `ceil(8/2) = 4`; ceiling `ln 8 = 2.079442` nats = **3.000000 bits exactly** = 6 octaves exactly.

n	category	f = k n^2 Hz	class m	class ratio	class cents	j	f*2^j Hz
1	intimate close	0.483426	1	1/1	0.000	8	123.7570
2	intimate far	1.933703	1	1/1	0.000	6	123.7570
3	personal close	4.350832	3	9/8	203.910	5	139.2266
4	personal far	7.734812	1	1/1	0.000	4	123.7570
5	social close	12.085643	5	25/16	772.627	4	193.3703
6	social far	17.403326	3	9/8	203.910	3	139.2266
7	public close	23.687861	7	49/32	737.652	3	189.5029
8	public far	30.939247	1	1/1	0.000	2	123.7570

Band `[110, 220)` Hz, per-index `j`, rule: the unique integer `j` placing `f 2^j` in the band. **Every value in the last column is an essence frequency, not a physical resonance. No claim is made that any distance, person, or encounter vibrates at any of these values.**

Four of the eight zones share one pitch class. Indices 1, 2, 4 and 8 all have odd part 1 — intimate close, intimate far, personal far, and public far are octave-identical in this lattice, all landing on **123.7570 Hz** in the chosen band. The distinct classes are `{1, 3, 5, 7}`, exactly the

four predicted by Theorem 7.1 and exactly the odd indices predicted as class representatives by Section 7.7.

This is a genuine structural finding about the lattice and it is emphatically not a finding about proxemics. Nothing about intimate distance and public distance is claimed to be similar. **The collapse is caused by the indices being powers of two, which is caused by where those four zones happened to fall in Hall's ordering. Had Hall subdivided into three phases per zone rather than two, giving twelve categories, the class structure would be entirely different — and Hall's choice was about ethnography, not about arithmetic.** That sentence is the whole lesson: **the lattice's structure is a fact about the index, the index's meaning is a fact about the domain, and the two must never be run together.**

Collision check (ES-9). The four realised classes sit at 0.000, 203.910, 737.652 and 772.627 cents.

```
class 1 (   0.000)  ->  class 3 ( 203.910)      203.910 cents
class 3 ( 203.910)  ->  class 7 ( 737.652)      533.742 cents
class 7 ( 737.652)  ->  class 5 ( 772.627)       34.976 cents  <-- COLLISION, disclosed
class 5 ( 772.627)  ->  class 1 (1200.000)      427.373 cents
```

Classes 7 and 5 — public close and social close — are 34.976 cents apart and will beat rather than form an interval if sounded together. Disclosed in advance, precisely so that nobody can later report it as a finding about the relationship between social and public distance. **It is `49/32` against `25/16`, and it is arithmetic.**

The full pairwise interval table, all 28 pairs:

pair	interval (cents)	essence ratio	HD_ess	pair	interval	ratio	HD_ess
1–2	2400.00	4/1	2.0000	3–5	1768.72	25/9	7.8138
1–3	3803.91	9/1	3.1699	3–6	2400.00	36/9	2.0000
1–4	4800.00	16/1	4.0000	3–7	2933.74	49/9	8.7846
1–5	5572.63	25/1	4.6439	3–8	3396.09	64/9	9.1699
1–6	6203.91	36/1	5.1699	4–5	772.63	25/16	8.6439
1–7	6737.65	49/1	5.6147	4–6	1403.91	36/16	5.1699
1–8	7200.00	64/1	6.0000	4–7	1937.65	49/16	9.6147
2–3	1403.91	9/4	5.1699	4–8	2400.00	64/16	2.0000
2–4	2400.00	16/4	2.0000	5–6	631.28	36/25	9.8138
2–5	3172.63	25/4	6.6439	5–7	1165.02	49/25	10.2586
2–6	3803.91	36/4	3.1699	5–8	1627.37	64/25	10.6439
2–7	4337.65	49/4	7.6147	6–7	533.74	49/36	10.7846
2–8	4800.00	64/4	4.0000	6–8	996.09	64/36	7.1699
3–4	996.09	16/9	7.1699	7–8	462.35	64/49	11.6147

The four most consonant nontrivial pairs are 1-2, 2-4, 3-6 and 4-8, all at `HD_ess = 2.0000` — and all four are index ratios of `2:1`, hence pitch-class collisions. **The most consonant pairs are exactly the redundant pairs, as Section 7.7 requires. The lattice is behaving.**

The refusal. Consider building an essence set from this framework's own six storage-law parameters — `delta`, `eta`, `N`, `u`, `C`, `Q`. It fails at Step 2 and it fails hard.

- **ES-R1:** the six are not categories of a common domain at all. `N` is a count, `u` is a geometric quantity in `[0,1]`, `Q` is a stock, `C` is a flow. They do not partition anything. **FAIL.**
- **ES-R2:** no intrinsic magnitude orders them. They could be ordered by the sequence in which the storage law is written, by alphabet, or by importance — three different orderings with no external arbiter, each giving a completely different lattice. **FAIL.**
- **ES-R3:** two analysts would not agree. **FAIL.**

Therefore the six parameters get no frequencies. They can be tabulated, reasoned about, measured, and used — everything the framework actually does with them — but they cannot

be sonified, because there is no defensible ordinal and the law would convert an arbitrary choice into a confident-looking set of hertz values.

That is the point of the requirements, and it is why they were stated before the arithmetic rather than after. Most category sets do not qualify. The instrument is narrow, and a narrow instrument used honestly is worth more than a wide one used loosely.

18.13 13 Discussion

18.13.1 13.1 What follows

A category scheme can generate a design target, and the derivation can be published. That is the whole positive claim and it is a real one. It replaces "pick a proportion that looks right" with "pick the proportion your own category scheme generates, and be able to state where it came from". The output is exact — `1 : 121/81 : 121/49`, a ratio of squares of integers, checkable by anyone — and it is therefore arguable in a way that a designer's taste is not. **Reproducibility is the contribution; efficacy is not claimed.**

The size of that contribution is bounded above and cannot be stated precisely, and the corpus has a template for reporting exactly that shape of claim. In a structural case elsewhere in this corpus the honest gain from an intervention is shown to lie strictly between zero and a computable ceiling — an off-ledger-to-represented ratio of `36.15/63.85 = 56.62` per cent, a base multiple of 1.5662 times — **with the realised gain bounded by that ceiling and not statable precisely** [Curzi, 2026]. The same form applies here. The upper bound on what essence tuning can contribute to a room is the spread of acceptable proportions it selects within, which the heuristic table of Section 9.3 makes explicit; the lower bound is zero; and no number between them is available, because no room was built and no comparison was run.

The four registers are one register up to a change of base, and the constants are exact. That is useful in a specific and limited way: a body's surplus stated as "1.37 nats" means nothing to most readers, "just under two bits" means something to an engineer, and a rising tone means something to everyone in the room. **The pitch register is a display register.** Its amplification of small surplus differences by a factor of `e^2` per nat is exactly what makes it good for display and bad for argument.

The forced ceiling is the sharpest consequence. `0 <= M <= ln N` says that the richest moment available to a perceiver is bounded by the logarithm of the number of kinds of thing that perceiver can distinguish, and the bound is logarithmic, so it is nearly immovable: doubling `N` buys exactly one bit wherever one starts. **That is a ceiling on the perceiver rather than on the world**, and it recovers the law of requisite variety inside first-person experience [Ashby, 1956]. It follows from the range result plus a chosen normalisation, and both links must always be published.

And the class-count theorem is a hard constraint on any application. Twelve categories give six voices. Twenty-four give twelve. The half that vanishes is not an artefact of the band or of the instrument; it is the odd-part rule, and no design decision recovers it. **Design to `ceil(N/2)`.**

18.13.2 13.2 What does not follow

Nothing about efficacy. The room described in Section 12.1 is not evidence that essence tuning does anything, and the coincidence that a twelve-category set generated a room seating twelve-and-a-bit people is a coincidence. It arose from `c/(2f)` with a particular `c` and a particular `k`, and one arithmetic step in a different direction would have produced an unbuildable room, as FM-3 shows for PC11 : PC5 : PC3. **Coincidences of this kind are the standard raw material of numerological reasoning and this paper declines to use them as evidence.** The literature on golden-ratio claims is instructive precisely because it documents how easily an

impressive-looking coincidence survives a hundred years of repetition without ever having been checked [Markowsky, 1992, Fischler, 1981].

Nothing about the domain from a collision. Classes 3 and 17 are 6.00 cents apart because `9/8` and `289/256` are nearly equal. Classes 5 and 7 are 34.976 cents apart because `49/32` and `25/16` are nearly equal. Indices 1, 2, 4 and 8 collapse because they are powers of two. **All three are facts about integers. Reading any of them as a finding about proxemics, or disciplines, or anything else, is FM-23.**

Nothing physical from an essence frequency. The room resonates; the category does not. That distinction is repeated in this paper more often than is elegant, and the repetition is deliberate: **it is the single sentence whose loss would reclassify the whole framework correctly as pseudoscience.**

Nothing about the ordering from the result. The instrument cannot check its own input. If the ordering was chosen because it sounded good, every interval is an artefact and no downstream computation detects it. The only defence is procedural — declare the ordering first, from a source outside the music — and it is the same defence preregistration provides in an experimental setting [Nosek et al., 2018, Simmons et al., 2011].

No thermodynamic or physical conclusion from any logarithm here. The entropy-like object in this paper is **Shannon entropy**, a dimensionless information measure quoted in nats or bits [Shannon, 1948], not thermodynamic entropy, which carries energy per temperature. They are different objects and nothing identifies the framework's surplus with either.

18.13.3 13.3 Relation to the surrounding literature

The consonance literature does not support the indices used here, and this paper says so. Tenney's harmonic distance [Tenney, 2015] and Euler's gradus suavitatis [Euler, 1739] are **defined functions on rationals**, chosen for their arithmetic properties. What is actually measured about consonance points elsewhere: roughness is a function of critical bandwidth [Plomp and Levelt, 1965], sensory and harmonic consonance are distinct phenomena [Terhardt, 1974], dyad consonance has been measured directly [Kameoka and Kuriyagawa, 1969], and — most damagingly for any argument resting on `9/8` — **consonance depends on the spectrum of the sound, so no ratio is a universal constant** [Sethares, 2005]. This paper uses the defined indices for their structural results (Theorem 7.2 and the consonance-redundancy identity) and makes no claim that they predict what a listener hears. The one perceptual claim it does make, H3, rests on roughness within a critical band and is cited to the measured literature.

Octave equivalence is grounded for pitch and only for pitch. It is formalised as the helical structure of pitch [Shepard, 1982] and measured as a genuine singularity in frequency-ratio discrimination [Demany et al., 2013] — and it is also weaker and more training-dependent than usually assumed [Hoeschele et al., 2012], and less cross-culturally general than scale-universality claims suppose [Burns, 1999]. **Carrying it over to a quantity that is not a pitch has no support at all**, and the framework's own bibliographic audit records that. The folding operation is well defined; its interpretation is not supported and is badged category 3 throughout.

The room acoustics are standard and are used as such. Sabine's relation [Sabine, 1922], the Schroeder boundary [Schroeder, 1965], the modal formula and the room-ratio heuristics [Beranek, 2004, Rossing et al., 2002] are established engineering, and nothing in this paper modifies or extends them. **If the essence law were withdrawn entirely, every acoustic number in Sections 9.2, 9.3 and 9.5 would stand unchanged.** That is the cleanest available statement of how much the framework contributes: the choice of which numbers to compute, and nothing else.

The psychophysics is cited to the sources and to their critics. Weber's fraction [Weber, 1834], Fechner's integration [Fechner, 1860] and Stevens' power law [Stevens, 1957] are the three pillars of the sensation register; the regression effect [Poulton, 1968, 1979], range effects [Teghtsoonian, 1971] and the dispute over whether magnitude estimation yields a ratio scale

at all [Luce, 2002, Narens, 1996] are cited alongside them. **The exponents are constants of a procedure, not constants of nature.**

18.14 14 Verification

The results of this paper compound from nine postulates. Each is atomic, checkable on its own, and carries its status. **Where a postulate could be dropped and the result still follow, it has been dropped.** The reader who wants a single one-minute audit should perform the three checks of Section 14.4.

The status vocabulary is the corpus's own and is applied here without relaxation. The framework's reference apparatus sorts its content into a three-drawer honesty classification — **four items of established mathematics, seven of established psychology, and exactly three declared postulates, with every protocol stated to run on the first two and to steer by the third** [Curzi, 2026]. This paper adds two postulates to that ledger, P5 and P6 below, and both are drawn on by every dimensioned result it reports.

18.14.1 14.1 Dimensional audit

This is the load-bearing section of the paper. The whole argument converts between quantities that are dimensionless — nats, bits, octaves, cents — and a quantity that is not: frequency, `[T^-1]`, in hertz. **The audit's job is to state exactly where a dimensionless index acquires hertz, and by what authority.**

Symbol	Dimension	Unit	Status
`n` essence index	`[1]`	none	**DEFINED** dimensionless ordinal. Not a magnitude; see Section 3.5
`k`	`[T^-1]` **as used**	Hz	**The number is DEFINED** and exact; **the hertz is POSTULATED**
`f(n) = k n^2`	`[T^-1]`	Hz	**DERIVED**, and inherits `k`'s postulated hertz entirely
`u`	`[1]`	none	**DEFINED** dimensionless: `u = 1 - (x.y)^2`, unit vectors
`N`, `N_eff`	`[1]`	none	**DEFINED** dimensionless counts; `N_eff` is the affine one
`F`, `S`	`[1]`	nats	**DERIVED** dimensionless: a logarithm of a count
bits, octaves, cents	`[1]`	none	**DERIVED** dimensionless log-ratio measures
`1/ln2`, `2/ln2`, `2400/ln2`	`[1]`	none	**DEFINED**, exact, dimensionless
`c` speed of sound	`[L T^-1]`	m/s	**MEASURED**: 343 m/s at 20 C in dry air

Symbol	Dimension	Unit	Status
L, Lx, Ly, Lz	[L]	m	**DERIVED** from c/(2f)
V	[L^3]	m3	**DERIVED**
S_surf	[L^2]	m2	**DERIVED**
A absorption	[L^2]	m2 sabins	**DESIGN VARIABLE**; a sabin is dimensionally an area
RT60	[T]	s	**DERIVED**; see the balance note below
f_s Schroeder	[T^-1]	Hz	**DERIVED**; see the balance note below
T = 1/f period	[T]	s	**DERIVED**
P forcing period	[T]	days or weeks	**DESIGN VARIABLE**
delta	[1]	per period	**STIPULATED** dimensionless fraction per step
tau = 1/delta	[T]	periods	**DERIVED**
eta	[Q] per nat per step	stock units per nat per period	**STIPULATED**; not dimensionless
Q, C	[Q], [Q] per step	proxy units	**STIPULATED**; no natural unit exists
dF	[1]	nats	**DERIVED**
cycles required	[1]	count	**DERIVED**, 1/(8 pi dF)
a Stevens exponent	[1]	none	**MEASURED**, procedure-bound
c Weber fraction	[1]	none	**PARAMETRIC** here; no value asserted
HD, G	[1]	none	**DEFINED** functions on rationals; not measurements

Does f(n) = k n^2 balance? Only if k carries [T^-1]. The index n is dimensionless, so n^2 is dimensionless, so **every hertz in this paper enters through k and through nothing else.** The number (9/8)^2/phi^2 = 0.483425732988414 is a pure ratio of pure numbers and has no dimension whatever. **Writing it as 0.483425732988414 Hz is a separate act, it has no external precedent of any kind, and it is a POSTULATE** (P6 below). Every frequency, room dimension, cavity volume and string length in this paper stands on that one step. **If P6 is withdrawn, f(n) becomes a dimensionless index, L = c/(2f) becomes dimensionally ill-formed, and the room does not exist.** The intervals, the pitch classes, the class-count theorem and the collision structure all survive, because k cancels from every one of them.

Does L = c/(2f) balance? [L T^-1]/[T^-1] = [L]. Yes, given P6.

Does RT60 = 0.161 V/A balance? [L^3]/[L^2] = [L], so **the constant 0.161 is not dimensionless: it carries [T L^-1], seconds per metre.** It is 24 ln 10 / c = 55.262/343 = 0.16111 s/m, and it is therefore temperature-dependent through c. **A paper that quotes 0.161 as though it were a pure number has not audited its own equation**; the value is specific to c = 343 m/s.

Does f_s = 2000 sqrt(RT60/V) balance? sqrt([T]/[L^3]) = [T^{1/2}][L^{-3/2}], so for f_s to be [T^-1] **the 2000 must carry [T^{-1/2}][L^{3/2}]**, that is m^(3/2) s^(-1/2). It is an SI-specific constant and is not dimensionless either [Schroeder, 1965].

Are the conversion factors dimensionless? Yes, all of them, and that is the point. 1/ln 2, 2/ln 2 and 2400/ln 2 are ratios of logarithms of pure numbers. **The conversion chain nats to bits to octaves to cents never leaves the dimensionless world. The chain acquires hertz exactly**

once, at `k`, and never again.

Where the framework is silent. The essence index has **no operational unit and no magnitude**; it is an ordinal position and the audit records it as one rather than inventing a unit for it. Stored order `Q` has no natural unit either: it is counted in whatever a declared proxy counts, and two systems' `Q` values are not comparable unless they declare the same proxy. **Comfort itself, the underlying construct of the corpus, has no operational unit at all** and is treated as a dimensionless magnitude by stipulation, not by measurement.

18.14.2 14.2 The postulates

P1 — `phi^2 = phi + 1`. *Status:* **CITED** — the defining property of the division in extreme and mean ratio [Euclid, 1956], with the general identity `phi^m = F(m) phi + F(m−1)` in Vajda [1989].

P2 — `(9/8)^2 = 1.265625`. *Status:* **ARITHMETIC.** `81/64 = 1.265625` exactly.

P3 — `log_b(x) = ln(x)/ln(b)`. *Status:* **CITED** — the change-of-base identity, and the reason information measures in different bases differ by a constant [Cover and Thomas, 2006, Shannon, 1948].

P4 — A cent is one twelve-hundredth of an octave. *Status:* **DEFINITION**, introduced by Ellis and universal since [von Helmholtz, 1885]. It carries no empirical weight.

P5 — `f(n) = k n^2`, with `n` the ordinal index. *Status:* **POSTULATED** — asserted by this framework, no external support, and not currently falsifiable (H6). The exponent is **defined by the model, not estimated** [Clauset et al., 2009].

P6 — `k` carries units of hertz. *Status:* **POSTULATED** — and this is a separate and much stronger claim than computing `k`. No physical, psychoacoustic or musicological literature combines a tuning ratio with an inverse square of the golden section to obtain a ground frequency [Markowsky, 1992]. **Every dimensioned result in this paper consumes P6 and no dimensionless result does.**

P7 — For a rigid rectangular enclosure the axial mode along `L` is `c/(2L)`, and reverberation obeys `RT60 = 24 ln 10 V/(cA)`. *Status:* **CITED** — Rossing et al. [2002], Sabine [1922], with the Schroeder boundary in Schroeder [1965]. This is real acoustics with a stated domain of validity.

P8 — `0 <= F <= ln N`. *Status:* **CITED** — the range result of the surplus functional [Curzi, 2026], following from `f(u) = ln(1 + (N−1)u)` on `u` in `[0,1]`.

P9 — The storage law is first order with time constant `tau = 1/delta`, and `delta = 0` is a pathology rather than an ideal. *Status:* **CITED** [Curzi, 2026a]. Consumed only by SA-11 and Section 9.4.

Two postulates that could be dropped were dropped. The normalisation of the moment onto `[0, ln N]` is consumed only by R4 and is named there rather than carried through the whole chain. The identification of the effective count with a cross-disciplinary index is consumed only by Section 4.6 and is verified there arithmetically rather than assumed.

18.14.3 14.3 The compounding chain

```
P1            -> L1   phi^2 = 2.618033988749895                       CITED/ARITHMETIC
P2            -> L2   (9/8)^2 = 81/64 = 1.265625                      ARITHMETIC
L1 / L2       -> R1a  k = 1.265625 / 2.618033988749895
                      = 0.483425732988414   (exact, 15 s.f.)          DEFINITION
R1a + P6      -> R1b  k = 0.483425732988414 Hz                        POSTULATED
P3            -> L3   1 nat = 1/ln 2 = 1.4426950408889634 bits        ARITHMETIC
P5            -> L4   f(2n)/f(n) = (2n)^2/n^2 = 4, i.e. two octaves   ARITHMETIC
L3 + L4       -> R1c  1 nat = 2/ln 2 = 2.8853900817779268 octaves     ARITHMETIC given P5
R1c + P4      -> R1d  1 nat = 2400/ln 2 = 3462.4680981335122 cents    ARITHMETIC
L4            -> R3   ONE BIT IS TWO OCTAVES, verified at
                      N_eff = 1, 3, 12: ratio 4.0 in every case       ARITHMETIC
```

```
P5              -> R6a  n = 2^a m, m odd, gives f(n) = k m^2 up 2a
                        octaves, so class(n) = odd part of n and a set
                        1..N realises exactly ceil(N/2) classes          ARITHMETIC (Thm 7.1)
P5 + P1         -> L5   f(m)/f(n) = m^2/n^2: k cancels from every
                        interval                                         ARITHMETIC (CC-R1)
R1b + P5        -> L6   f(7) = 23.687861, f(9) = 39.157484,
                        f(11) = 58.494514 Hz                             ARITHMETIC given R1b
L6 + P7         -> R2a  Lz = 343/(2*23.687861) = 7.239995 m
                        Ly = 343/(2*39.157484) = 4.379750 m
                        Lx = 343/(2*58.494514) = 2.931899 m              ARITHMETIC + CITED
R2a             -> R2b  V = 92.9687 m^3, S_surf = 131.5546 m^2,
                        floor 31.7094 m^2, occupancy 12.68 at 2.5 m^2    ARITHMETIC
R2b + P7        -> R2c  RT60 = 0.5 s requires A = 0.161*92.9687/0.5
                        = 29.936 m^2 sabins, alpha_bar = 0.2276          ARITHMETIC + CITED
R2b + P7        -> R2d  f_s = 2000*sqrt(0.5/92.9687) = 146.67 Hz,
                        below the folded set's 233.98-386.74 Hz          ARITHMETIC + CITED
L5              -> R6b  classes 7 and 5 fold to 49/32 and 25/16:
                        1200*log2((25/16)/(49/32)) = 34.9756 cents       ARITHMETIC
R6a + R1b at N = 8
                -> R6c  indices 1, 2, 4, 8 share class 1 and land
                        together at 123.7570 Hz in [110, 220)            ARITHMETIC given R1b
P8              -> L7   ceiling ln N; at N = 12, 2.484907 nats           CITED
L7 + R1c + L3   -> R4   ln 12 = 2.484907 nats = 3.584963 bits
                        = 7.169925 octaves; with the chosen
                        normalisation, 0 <= M <= ln N                    CITED + CHOSEN NORM
u = 0.42, N = 8
                -> L8   N_eff = 1 + 7(0.42) = 3.94   (AFFINE count)      DEFINITION
L8              -> R5a  F = ln 3.94 = 1.371181 nats = 1.978196 bits
                        = 3.956391 octaves = 4747.67 cents               ARITHMETIC
P8 at N = 8     -> L9   ln 8 = 2.0794415 nats = 3 bits EXACTLY
                        = 6 octaves EXACTLY                              ARITHMETIC
R5a + L9        -> R5b  F/ln N = 0.6594; ceiling/achieved frequency
                        ratio = (8/3.94)^2 = 4.1228 = two octaves        ARITHMETIC
R1b + P5 + P9   -> CRa  PC1 period 2.068570 s folds at p = 19 to
                        12.5524 days; at delta = 0.05/week,
                        P/tau = 1.7932/20 = 0.0897 <= 0.5                ARITHMETIC + CITED
```

Read the status column. R1a, R1c, R1d, R3, R6a, R6b, CC-R1, R5a, R5b and L9 are **ARITHMETIC** and can be redone by hand. R2a to R2d and CRa are arithmetic **on top of cited physics**. R4 rests on a **cited range result plus a chosen normalisation**, and both links are named every time it is quoted. **R1b, and therefore every dimensioned result in the paper — the room, the cadence, the cavity volumes, the string lengths, the audible placements — consumes P6, which is POSTULATED with no external support of any kind.**

This chain is not a proof that anything in the world behaves this way, and the paper says so here rather than leaving a reader to notice. A chain whose dimensionless part is arithmetic and whose dimensioned part rests on one unsupported postulate delivers exact mathematics and delivers nothing whatever about categories, institutions or minds. **What it does deliver is a target-selection procedure whose every step is auditable.**

18.14.4 14.4 The three checks a reviewer can do in one minute

Check one — the constant. `(9/8)^2 = 81/64 = 1.265625.` `phi = (1 + sqrt 5)/2 = 1.618033988749895`, so `phi^2 = 2.618033988749895`, which is `phi + 1` as required. Dividing, `1.265625/2.618033988749895 = 0.483425732988414`. **Three divisions, no model.**

Check two — the factor of two. `f(2n)/f(n) = (2n)^2/n^2 = 4`, and `log2 4 = 2`. **One line, and it is the whole of the claim that one bit is two octaves.** Confirm at `N_eff = 3`: `f(6)/f(3) = 36/9 = 4`. At `N_eff = 12`: `f(24)/f(12) = 576/144 = 4`.

Check three — the room. `k (7)^2 = 0.483425732988414 x 49 = 23.687861` Hz; `343/(2 x 23.687861) = 7.239995` m. Repeat at `n = 9` and `n = 11` for 4.379750 m and 2.931899 m. Multiply: `V = 92.9687` m3. Then `0.161 x 92.9687/0.5 = 29.936` m2 sabins and `2000 sqrt(0.5/92.9687) = 146.67` Hz. **Five multiplications and two square roots, and the entire built result of the paper is established or refuted.**

18.14.5 14.5 Recomputation report: four printed values that did not reproduce

HC-4 requires every figure to be recomputed rather than transcribed. Four printed values in the source material did not reproduce, and they are listed here rather than corrected silently.

(i) `k` at seventeen significant figures. The exact value is `0.48342573298841433274...`, so the correctly rounded seventeen-figure value is **0.48342573298841433**. The corpus has printed **0.48342573298841435**, which is the IEEE-754 double-precision representation of the quotient rather than the correctly rounded exact value. **The fifteen- and sixteen-figure roundings, 0.483425732988414 and 0.4834257329884143, are both correct**, so the two precisions circulating in the corpus are not in conflict — they are the same number quoted at different precisions. Only the seventeenth figure is wrong, and it is wrong by one unit in the last place. This paper quotes fifteen figures.

(ii) `2400/ln 2`. The exact value is `3462.46809813351217766...`, so the correctly rounded seventeen-figure value is **3462.4680981335122**. The corpus has printed **3462.4680981335123**, again the double-precision representation. `1/ln 2` and `2/ln 2` are correct as printed at seventeen figures. This paper prints 3462.4680981335122.

(iii) The cadence-rounding departure in FM-5. Rounding a 9.9179-day derived cadence to 14 days is a ratio of 1.4116, and `1200 log2(1.4116) = 596.78` cents. **The source prints 590 cents.** No calculation returning 590 could be reconstructed; the corrected value is 596.8 cents and appears in FM-5 above.

(iv) The mode count below 150 Hz. The source prints **48 modes below 150 Hz** for the worked room. Enumerating `f(p,q,r)` over `p,q,r` from 0 to 4 returns 48; over 0 to 5 returns 53; and the count converges at **55**, which is the correct figure and is what Section 9.3 prints. **The published 48 is a truncation artefact of the source's own minimum recommended truncation order**, and SA-9 has been amended to require the truncation order to be stated.

Three smaller discrepancies are recorded for completeness. The steady state in the transit board case is printed in the source as `Q_inf = 3.9999`; **the exact value is 4.000000**, necessarily, because the required surplus was solved to make the quarterly step zero, so `eta S = delta Q + C` identically. The source describes the folded PC11 value as four octaves above `f(11)`; `233.9781/58.494514 = 4.000`, which is **two** octaves. And the source attributes the PC11 coincidence to indices 11 and 44 sharing an odd part; the index actually at work is 22, since `233.978055 = f(22)`, though the statement about 44 is independently true.

A fourth claim could not be reconstructed at all. The source states that the rendering-rate collapse at the `N = 16` to `N = 17` boundary is "a factor of exactly 8", and in the preceding parenthesis states it as 5.828. **Recomputation gives 5.8284 by both available routes** — the cents ratio `34.976/6.001` and the cycle-count ratio `22.957/3.939`, which are the same number `ln(50/49)/ln(289/288)` — and no calculation returning 8 could be found. Section 7.4 prints 5.828 and flags the inconsistency.

18.14.6 14.6 Limits, asymptotics, and the model's valid domain

`n -> 0`. `f = 0`. There is no zeroth essence and the law says so by returning silence. **`n = 1`.** `f = k`, the ground, and the unique meaningful invariant of the law (the algebraic fixed points

of `f(n) = n` sit at `n = 0` and `n = 1/k = 2.0685`, neither of which is meaningful, since index and frequency are not the same kind of quantity).

`n -> infinity`. `f` diverges quadratically and leaves the audible band quickly: at `j = 0` the last index inside 20 kHz is `n = 203`, since `k (203)^2 = 19921` Hz and `k (204)^2 = 20117` Hz. Under a shared transposition the bound is far tighter, `N <= 25` at `j = 6` (Section 8.3). **The working range is short and must be declared.**

`N = 1`. `ln N = 0`, one class, no interval exists at all. **A single category has zero available surplus, not a small amount**, and the ceiling and the achieved value collapse together. **`N = 2`.** Still one class: indices 1 and 2 are two octaves apart and harmonically identical. **A two-category set has a ceiling of one bit and a harmonic vocabulary of one voice**, which is worth stating because it is the smallest case where the ceiling and the vocabulary disagree.

`N -> infinity`. Classes grow as `N/2` — linearly — while the pitch span of the folded octave is fixed at 1200 cents and the raw span grows only as `2400 log2 N`. **The mean gap between classes therefore falls as `2400/N` and the minimum gap falls faster.** That asymptotic is the collision cliff of Section 7.4 seen from a distance: collisions are not an accident of the numbers 5, 7, 3 and 17, they are **forced**, and the only question is at which `N` each one arrives. **This limit is the strongest available argument for small essence sets and it is structural rather than aesthetic.**

`u -> 0`. `N_eff -> 1`, `F -> 0`, `f -> k`. Every category is a duplicate of every other and the roster has left the arithmetic; the pitch is the bottom of the instrument. **`u -> 1`.** `N_eff -> N`, `F -> ln N`, `f -> k N^2`, the ceiling is attained and no further gain is available from geometry.

`delta -> 0`. `tau = 1/delta -> infinity`, so the admissibility test `P <= tau/2` is satisfied by every cadence whatever. **That apparent permissiveness is a trap and it is a signature of the storage law's own pathology, not a licence:** at `delta = 0` the map is an isometry, no fixed point exists, and there is no equilibrium for a cadence to approach [Curzi, 2026a]. **SA-11 is valid only on `0 < delta <= 1` and is vacuous at `delta = 0`. `delta -> 1`.** `tau -> 1` period and then `tau_exact -> 0`, so no cadence is admissible; correctly so, because there is no stock left to build.

`RT60 -> 0`. `f_s -> 0`, so nominally the entire band is statistical — but an anechoic space has no diffuse field at all and Sabine's relation fails from the other side. **`A -> 0`.** `RT60 -> infinity`. Sabine's relation is known to overestimate reverberation at high mean absorption, where Eyring's form is preferred [Beranek, 2004, Rossing et al., 2002]; the worked room's `alpha_bar` runs 0.114 to 0.284, and the highest of those, at `RT60 = 0.4` s, is at the edge of the range in which Sabine is normally accepted. **The 0.5 to 0.6 s target used throughout sits at `alpha_bar = 0.19` to `0.23`, comfortably inside it.**

`dF -> 0`. Cycles required diverge as `1/(8 pi dF)`. **Resolving arbitrarily small surplus differences costs unbounded time**, which is the correct behaviour and is the reason the cliff at `N = 17` is expensive rather than merely narrow.

Stated domain. The results of this paper are valid for: `N` in `[2, 25]` for a rendered set under a shared transposition and `[2, 16]` if the 34.976-cent minimum gap is to be preserved; `u` in `[0,1]`; `delta` in `(0, 0.5]` for the cadence test; a rectangular enclosure with reasonably uniform absorption distribution and `alpha_bar` below about 0.3 for the Sabine arithmetic; and the audible band 20 Hz to 20 kHz for anything rendered. **Outside those ranges the arithmetic still returns numbers and they should not be interpreted.**

18.14.7 14.7 Invariance

The strongest invariance in this paper is exact and it is the reason the framework's weakest constant does not matter. Under `k -> lambda k` for any `lambda > 0`, every interval `f(m)/f(n) = m^2/n^2` is unchanged, every pitch class is unchanged, every consonance index is unchanged, and every room *ratio* is unchanged. **The chosen constant `k` — the one part of the law with no justification beyond preference — has no influence whatever on any relational result.** It sets where the instrument sits and not how it is tuned. **What would break it:** a law of the form `f = k n^2 + b` with `b` not zero, in which `k` would no longer cancel. The invariance holds for the pure power law and for no other form.

Pitch class is invariant under `n -> 2^a n` for integer `a`, and under nothing else. That is Theorem 7.1 restated as an invariance: the folding map's fixed-point set is the odd indices. **It holds only for integer transposition**; a fractional `j` destroys it immediately, which is why ES-10 requires `j` to be an integer and requires it to be published.

Intervals are invariant under a shared `j` and are not invariant under a per-index `j`. That is the precise statement of FM-28, and it is a genuine invariance question rather than a stylistic preference: a shared `j` is a similarity transformation on the whole set, a per-index `j` is not.

Room proportion is scale-free; room size is not. The ratio `1 : 121/81 : 121/49` is invariant under a uniform rescaling of all three axes and is a property of the class triple alone. **The absolute dimensions 7.2400 x 4.3798 x 2.9319 m are fixed only by the additional decision to anchor the axial fundamentals at the *raw* essence frequencies rather than at any transposition of them.** That decision is a choice, it consumes P6, and a different octave offset would give a room of the same shape and a different size — including, at some offsets, an unbuildable one (FM-3).

How many scales the claim has been exercised at: three, all constructed. One room, one eight-member board, one eight-zone published lattice. **Zero measured systems, zero independent replications, one author.** The invariance is a claim about form that has been exercised, not an invariance that has been established.

18.14.8 14.8 Analogy discipline and the precision of the numbers

Restated once more because it is this paper's central risk. Sabine's relation, the Schroeder boundary, the axial-mode formula, the tube, the Helmholtz resonator and the string are **category 1, genuine derivations**: the room is a physical resonator and satisfies the premises. The storage law as a first-order low-pass section and octave folding applied to a period are **category 2, a shared mathematical form**, and no physical authority transfers. **The mapping of a category index to a frequency, and every reading of meaning off the resulting pitch relations, is category 3 — an analogy with no formal warrant.** The classification is stated at first use in Sections 3.3, 9.2, 9.4 and 9.5 and is repeated here so that no reader can reach the conclusion without meeting it.

The formal coincidence between `f(n) = k n^2` and the eigenvalue spacing of a one-dimensional infinite square well, where `E_n` is proportional to `n^2`, is a genuine mathematical correspondence and is noted as such. **It is category 3 and it is not evidence of a physical mechanism**, and asserting the connection would be exactly the decorative analogy this framework's own honesty apparatus exists to prevent.

Precision. The numbers here are of four kinds and are quoted accordingly.

Defined and exact — `k = 0.483425732988414`, `1/ln 2`, `2/ln 2`, `2400/ln 2`, `(9/8)^2 = 1.265625`, `phi^2 = phi + 1`, every cents value of a rational interval, every `HD` and gradus. These may be recomputed to any further precision and the digits printed carry no uncertainty because there is no measurement behind them.

Derived from defined, given a cited physical relation — the room dimensions, `V`, `S`, `A`, `f_s`, the cavity volumes, the string lengths. **Their arithmetic is exact; their correspondence to a real room carries the construction tolerance of the build and the error band of Sabine's relation**, neither of which is a number this paper can supply because no room was built.

Illustrative model output — every number in Section 12 derived from a stipulated parameter: `u = 0.42`, `eta = 0.35`, `delta = 0.08`, `Q = 4.0`, `C = 0.30`, `delta = 0.05` per week. **These carry no measurement uncertainty because there was no measurement**, and six decimal places there record arithmetic rather than confidence. CK-13 requires the conversion to be rounded back to the source's precision, which is why `u = 0.42` at two figures yields 4750 cents and not 4747.67.

Measured — `c = 343` m/s at 20 C in dry air, and Stevens' exponents. **The exponents are constants of a procedure and are quoted nominally, without intervals, because the sources report them as method-dependent** [Poulton, 1979, Teghtsoonian, 1971]; that is a limitation of

the table and is stated as one rather than papered over with a spurious uncertainty.

No quantity in this paper is a fitted parameter, and no confidence interval appears anywhere in it, because nothing was fitted and nothing was measured by this work.

18.15 15 Limitations

No dataset was collected and no room was built. Every number in Section 12 is the output of stated parameters under the stated model, carried through so a reader can re-derive it. The three systems demonstrate that the procedure runs and that it reports its own refusals; **they are not evidence that any institution or any room behaves this way and must never be cited as such.**

The hertz is the weakest link and everything dimensioned hangs from it. `k` as a number is exact; `k` as a frequency is a postulate with no precedent in any literature (P6). Withdraw it and the room, the cadence in days, the cavity volumes and the string lengths all vanish, while the intervals, the classes and the class-count theorem survive untouched. **The paper's dimensionless results are much stronger than its dimensioned ones, and a reader should weight them accordingly.**

The square law is a postulate that is not currently falsifiable, and its exponent is defined rather than estimated. H6 states the condition under which it would become testable — a domain in which the relation between category types carries an independently measured magnitude — and no such domain is offered here. **The sensitivity table of Section 4.3 is what replaces a confidence interval, and it shows that every interval in the paper changes if the exponent changes**; the conversion key 2.885390 is exact given the postulate and is not a fact about the world. The distribution-fitting procedure that any estimated exponent would require [Clauset et al., 2009] is not run because no exponent was estimated; **that is a defence only so long as the paper never claims the exponent was found.**

The scale promotion is not licensed by the data and is conceded to be a stipulation. The index is ordinal; the frequency is ratio-scaled; the law manufactures the intervening structure by assertion (Section 3.5). Michell [1997] argues that psychology systematically assumes quantitative structure without testing for it, and **that is the strongest objection to this framework.** The available remedies — conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971] or a Rasch demonstration of specific objectivity [Rasch, 1960] — **have not been run for any construct here, and until one is, every quantitative reading of an essence interval is provisional.** The paper's own catalogue names the same objection operationally as FM-21 and FM-26.

The ordering is the whole method's input and this paper's central worked case has a weak one. The twelve-category enumeration of Section 12.1 has no intrinsic magnitude behind it. The permutation test returns a maximum displacement of 694.82 cents, which is **six times any reasonable threshold**, so the ordering is load-bearing and a different but equally defensible enumeration gives a materially different tuning. **The room in Section 12.1 is built to *an* essence set, not to *the* essence set, and no result derived from it is stable under a re-ordering.** Only the proxemic case of Section 12.3 has a type-(a) ordering source, and even there Hall's zone *boundaries* are nominal and culturally variable.

`N_eff = 1 + (N−1)u` is not a Hill number and no external literature contains it. `1/HHI` is the Hill number of order two and `exp(H)` of order one, and both are fully grounded [Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979]; `1 + (N−1)u` is a linear interpolation between 1 and `N` and is this framework's own construction, adopted by design axiom S3 [Curzi, 2026]. **A reviewer reaching for the Hill-number literature as support will find it does not apply.** The two quantities are used under one effective-count heading across this corpus, which is a real exposure; the concentration work that genuinely uses a Hill number of order two reports quantities on a different footing [Curzi, 2026] and the two must not be compared numerically.

The surplus functional's form is underdetermined by the evidence. Latané [1981]'s sublinear power law and Briscoe et al. [2006]'s `n log n` fit the qualitative facts as well as the log-

arithm. **Nothing assembled here discriminates between them.** The logarithm is chosen for its axiomatic characterisation, and one of the axioms is itself chosen. Where the logarithm *is* independently warranted — for a stock that compounds — the warrant is ergodic rather than perceptual [Peters and Gell-Mann, 2016], and **that is the warrant the argument should lean on** rather than on Fechner, whose perceptual story is contested and known to fail at the extremes of intensity.

The reflexive claim is not identified. The mechanism by which a room's acoustics act on a group runs through turn-taking, and a group whose members both influence and constitute its output is exactly the linear-in-means reflexive case whose endogenous effect is not recoverable from observational data [Manski, 1993], with influence and selection generically confounded in network data [Shalizi and Thomas, 2011]. **This paper adopts no remedy**, so the acoustic-exclusion mechanism it relies on is a **modelling position rather than an identified empirical result**. The most realistic available repair is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001], and it is not attempted here.

H2's test is unblindable by construction. The disclosure requirement of Section 3.6 means every occupant of a tuned room can read what the tuning is. Expectancy effects therefore cannot be removed from any comparison of a tuned against an untuned room [Rosenthal and Jacobson, 1968, Rosenthal and Fode, 1963], and the Hawthorne-type explanations that would otherwise be invoked are themselves weaker than commonly assumed [McCambridge et al., 2014]. **The framework's ethics and its evidential ambitions are in genuine tension here, and the ethics wins.** That is a permanent limit on what can ever be claimed for a disclosed tuning, and it is stated as one.

The shuffled-input control that the falsification protocol depends on is anti-conservative. It is recorded elsewhere in this corpus as measuring a false-positive rate against a nominal level that it exceeds, and as blind to repackaging and to a shared slow driver [Curzi, 2026i]. **A passed shuffle is therefore not on its own evidence that a tuned room did anything**, and SA-16's protocol inherits that weakness.

The consonance indices are chosen functions and disagree with the measured literature. Tenney distance and Euler gradus are defined on rationals [Tenney, 2015, Euler, 1739]; measured consonance depends on critical bandwidth and on spectrum [Plomp and Levelt, 1965, Sethares, 2005]. **The structural results of Section 7.6 and 7.7 are results about the defined indices and about nothing else**, and any claim that a listener will find the `HD_ess = 2.0000` pairs "most consonant" is an extrapolation the indices do not license.

Octave equivalence is carried outside its evidential domain. It is grounded for pitch [Shepard, 1982, Demany et al., 2013] and is weaker than assumed even there [Hoeschele et al., 2012]; **nothing supports applying it to a quantity that is not a pitch**, and folding an index into a band and reading meaning off the result is an unsupported interpretation of a well-defined operation.

No human-subjects research was conducted. H2, H3 and H4 all specify tests involving people, none has been run, and the constraints of Section 3.6 are constraints on prospective work rather than reports of practice.

Model comparison, effect sizes and replication. This paper fits nothing, so no information criterion applies to it; but it also therefore offers no out-of-sample evidence for anything, and the corpus's one applied fit reports held-out error against baselines and **no AIC or BIC**, so its parameter count was never paid for [Curzi, 2026i]. Nothing in Section 12 carries an interval, because nothing in Section 12 was measured. **Every system in this paper was constructed by one author and there is no independent replication of any parameter**, which the reproducibility literature is unambiguous about [Open Science Collaboration, 2015, Ioannidis, 2005]. And directional predictions are nearly costless in the social sciences [Meehl, 1967]; the point predictions with constants offered here are the right answer to that objection in form, **but none of them has yet been put at risk.**

Four printed values in the source material did not reproduce and one claim could not be reconstructed at all. They are listed in Section 14.5. Three are last-place floating-point artefacts, one is a truncation artefact in a mode count, one is a cents figure that is wrong

by seven cents, and one is a factor stated two different ways in adjacent sentences. **None changes a conclusion; all are printed because a reader who finds them first has reason to doubt everything else.**

Two recommendations in the source material are refused rather than carried. Subliminal embedding is refused on the ethical ground stated in Section 3.6, and the coil claimed as an amplifier is refused because no verified engineering basis exists for the claimed mechanism and no entry in this corpus's bibliography supports it. **Both refusals are recorded here so that a reader comparing this paper against its source can see that the omissions are deliberate.**

18.16 16 Conclusion

R1 — The conversion keys, exact. The four registers of this framework — surplus in nats, pitch in octaves, sensation, and the moment — are one register up to a change of logarithmic base, and the constants are exact and defined: **`1/ln 2 = 1.4426950408889634` bits per nat, `2/ln 2 = 2.8853900817779268` essence octaves per nat, `2400/ln 2 = 3462.4680981335122` cents per nat**, with the ground constant `k = (9/8)^2/phi^2 = 0.483425732988414`, **defined and exact as a number, and a POSTULATE the moment it is written in hertz.** A conversion carries the status of the quantity it converts and never upgrades it.

R2 — The room. Taking essence indices 7, 9 and 11 as axial-mode targets and `L = c/(2f)` with `c = 343` m/s gives a room of **7.2400 by 4.3798 by 2.9319 m**, exactly in the ratio `1 : 121/81 : 121/49`, with **`V = 92.969` m3**, floor area 31.709 m2, total surface 131.555 m2 and **occupancy 12.7** at 2.5 m2 per person. Its proportion matches the Volkmann heuristic at a Euclidean error of **0.031**, the best of five heuristics under an exhaustive search over class triples and octave offsets. **`RT60 = 0.5` s requires `A = 29.94` m2 sabins, a mean absorption of 0.2276**, and the **Schroeder frequency is 146.7 Hz**, below the entire folded set of **233.98 to 386.74 Hz**, so the whole set lies in the statistical region and Sabine governs for all of it. The room carries **55 modes below 150 Hz** — not the 48 previously printed, which is a truncation artefact — a mean gap of 4.02 Hz across the first twenty, a maximum gap of 15.47 Hz and a near-degeneracy at 0.51 Hz that will beat at half a hertz.

R3 — One bit is two octaves, not one. Doubling the effective count multiplies a square-law frequency by four, `f(2n)/f(n) = (2n)^2/n^2 = 4`, and four is two octaves. Verified at `N_eff = 1, 3, 12`, ratio 4.0 in every case. **The factor of two in `octaves = 2F/ln 2` is the exponent of the essence-law postulate and nothing else**, so the key is exact given the postulate and is not a fact about the world.

R4 — The forced ceiling. From the surplus range result plus the chosen normalisation, **`0 <= M <= ln N`**: the richest moment available to a perceiver is bounded by the logarithm of the number of kinds of thing that perceiver can distinguish. **A fully realised twelve-category system is `ln 12 = 2.484907` nats = `3.584963` bits = `7.169925` octaves.** Doubling `N` adds exactly one bit wherever one starts, so going from 2 categories to 1000 buys 6.21 nats — under nine bits. **The ceiling is on the perceiver, not on the world**, and both links of the derivation are published every time it is quoted.

R5 — The transport-board case. With `N = 8` and `u = 0.42`, the **affine** effective count is `N_eff = 3.94` — not a Hill number — and `F = ln 3.94 = 1.371181` nats = **1.978196 bits**, against a ceiling of `ln 8 = 2.079442` nats = **3 bits exactly = 6 octaves exactly**. The board sits at **66 per cent of its own ceiling**, and the ratio of ceiling to achieved essence frequency is `(8/3.94)^2 = 4.1228`, **two full octaves below its own maximum**. The diagnostic the registers make visible is positive headroom `h = +0.107805` against a negative actual step `−0.140087`: **the ceiling permits growth and the geometry does not deliver it.**

R6 — The disclosed collision. Eight proxemic zones ordered by distance give `ceil(8/2) = 4` pitch classes; **indices 1, 2, 4 and 8 collapse to one class at 123.7570 Hz** in `[110, 220)` Hz because their odd part is 1; and **classes 5 and 7, at `25/16` and `49/32`, sit 34.976 cents apart and will beat rather than form an interval.** The collision is disclosed in advance precisely so that it cannot later be reported as a finding about social and public distance. **It is arithmetic, and the same**

34.976-cent minimum gap holds unchanged for every set size from 7 to 16, collapsing to 6.001 cents the moment index 17 enters — a change of rendering-duration regime by a factor of 5.828.

CR — The counter-result the paper supplies against its own instrument. This framework's own six storage-law parameters — `delta`, `eta`, `N`, `u`, `C`, `Q` — **fail all three admissibility requirements**: they are not categories of a common domain, no intrinsic magnitude orders them, and two analysts would not agree. **They therefore get no frequencies at all.** With that refusal go two hard limits: **at most 25 essences fit the audible band under a shared transposition, and only at `j = 6`**, and **an essence frequency is never a physical resonance.** The instrument is narrow. **A narrow instrument used honestly is worth more than a wide one used loosely**, and this paper's principal claim is that the essence law generates targets — it does not generate effects, and nothing here has claimed that it does.

18.16.1 Declarations

Data and code availability

No external dataset was collected for this paper and no room was built. All numerical results are computed from the stated parameters under the models developed in the text, and every table can be re-derived from the equations and the constants of Section 4.1. The three systems worked in Section 12 are constructed examples with declared parameters, not field observations; the proxemic ordering of Section 12.3 is taken from a published source [Hall, 1966] and only its ordering, not its boundary values, is used. The figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper and the relations of Sections 4, 7 and 9.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, recomputed every printed figure, produced the worked examples and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The three worked systems are constructed. Hypotheses H2, H3 and H4 specify tests involving human participants; **none has been run**, and each is stated as a prospective test with its falsification condition rather than as a report of a result.

The no-subliminals constraint of Section 3.6 is a design prohibition adopted by this framework and a constraint on prospective work. It is not a report of any study performed, and it is not an empirical claim that concealed signal is ineffective. Its content is that a method working by removing a person's ability to decline is outside this framework's scope regardless of whether it works, and that a message a listener cannot hear is a message the listener cannot decline. The associated requirements — audibility at every listening position, written disclosure before entry, a disclosed off switch, and publication of any structural derivation — are stated as specifications SA-1 to SA-3, SA-15, CK-19 and ES-16, and they bind any future

application of this material. **Two recommendations present in the source material this paper analyses — the embedding of subliminal content in a rendered track, and the presentation of an unverified device as an amplifier — are refused here and the refusals are recorded in Sections 3.6 and 15.**

Preregistration

This paper reports no preregistered test of its own; its results are derivations and recomputations. The preregistered protocol for testing a framework claim against a real system, together with the confirmatory application that **failed** its second falsification rule, is reported in Curzi [2026i], and any applied claim resting on this paper inherits that result. Section 12.1 Step 8 specifies in advance the protocol that would be required before any claim is made for a tuned room, including the four baselines and the shuffled-input control; **that protocol has not been run, and Section 15 records both that fact and the reasons the control is weaker than it looks.**

18.17 Appendix A Seven-gate formula audits

Specifications CK-17 and ES-18 require every formula entering a specification to carry a full seven-gate audit in print, **including the gates it fails**. The gates are: **G1 dimensions**, **G2 scale type**, **G3 limiting cases**, **G4 fixed points**, **G5 identifiability**, **G6 falsification**, **G7 reproducibility**. Nine formulas enter the specifications of Section 10 and all nine are audited below. Four gates fail outright and three are recorded as not yet attempted.

18.17.1 A.1 The register conversion identity

```
x_target = x_source (ln b_source / ln b_target)
```

Gate	Test	Result
G1 dimensions	Both sides are dimensionless logarithms of a ratio; the multiplier is a ratio of logarithms of pure numbers. Units agree.	**PASS**
G2 scale type	All four registers are ratio-scaled *in the argument* and interval-scaled *in the logarithm*: a difference of two pitch values is meaningful, an absolute value requires a stated reference. The conversion is a positive linear map, permissible on interval and ratio scales alike.	**PASS, reference required in print**
G3 limiting cases	`x -> 0` gives 0 in every register — unison, zero surplus, ground essence. `x -> infinity` diverges in every register at the same rate. No register saturates or wraps.	**PASS**

Gate	Test	Result
G4 fixed points	`x = 0` is the unique fixed point: zero surplus is zero bits is zero octaves is zero cents. **The ground is shared by all four registers**, which is what makes them comparable.	**PASS**
G5 identifiability	The multiplier is fixed by the two bases; no free parameters to recover.	**N/A — no free parameters**
G6 falsification	Not empirically falsifiable: an identity of logarithms, arithmetically checkable and checked (Sections 4.4, 4.6). The framework claim riding on it — that the essence exponent is 2 — **is** falsifiable in principle and is flagged in Section 4.3 and H6.	**N/A (identity); dependent postulate flagged**
G7 reproducibility	Any correct implementation returns the same constants. Verified in double precision and cross-checked in 45-digit decimal arithmetic, which is how the two last-place discrepancies of Section 14.5 were found.	**PASS**

18.17.2 A.2 Fechner's law

```
S = k_F ln(I / I_0)
```

Gate	Test	Result
G1 dimensions	`I/I_0` is dimensionless by construction, so the logarithm is admissible. `S` carries the units of `k_F`, the scale constant of the chosen sensation unit.	**PASS if `I_0` declared; FAIL if the reference is omitted**

Gate	Test	Result
G2 scale type	`I` is ratio-scaled with a true zero. `S`, built by counting just-noticeable differences, is at best **interval-scaled**: JND counting gives equal-appearing increments, not verified equal ratios. This is the classic objection to Fechner and it is a real one. Ratio-scale sensation requires magnitude estimation, not JND integration.	**PARTIAL — `S` is interval; do not divide two Fechner sensations (CK-16)**
G3 limiting cases	`I -> I_0` gives `S -> 0`. `I -> 0` gives `S -> -infinity`: the divergence is the absolute threshold, below which the law is known to fail. `I -> infinity` gives unbounded `S`, failing at every modality's damage ceiling. **Both divergences name real phenomena and mark the law's known boundaries.**	**PASS, both boundaries named**
G4 fixed points	`I = I_0` is the unique invariant: the scale has no content until a reference is chosen, and the reference is a decision.	**PASS**
G5 identifiability	`k_F` and `I_0` trade off — rescaling `I_0` shifts `S` by an absorbable constant. **Identifiable only up to that offset**, so a fit must fix `I_0` by an independent operational definition.	**CONDITIONAL — fix `I_0` externally or the fit is degenerate**
G6 falsification	Falsified in a modality if the measured Weber fraction varies systematically with `I` across the working range. It has been so falsified at the extremes: dead as a universal law, standing as a middle-range description [Krueger, 1989].	**PASS — and has partially failed, honestly**
G7 reproducibility	Weber fractions recover across subjects and laboratories within modality-specific bounds — among the most reproducible findings in psychology.	**PASS**

18.17.3 A.3 The moment equation in log form

`ln M = ln D − ln P + ln E` — **POSTULATE, this framework's own formula**

Gate	Test	Result
G1 dimensions	`D` and `P` **must share units** so that `r = D/P` is dimensionless; `M` inherits `E`'s units. `ln D − ln P` is well formed if and only if `D` and `P` share both units and reference.	**CONDITIONAL — shared units are a requirement, not a result**
G2 scale type	Division requires **ratio scale with a true zero** for `D` and `P`. **This is the crux.** With only ordinal data the formula is a heuristic ordering, not an equation, and must be demoted in print [Michell, 1997, Stevens, 1946].	**UNRESOLVED PENDING MEASUREMENT**
G3 limiting cases	`P -> 0`: `M -> infinity` — a frame with no intake. `D -> 0`: `M -> 0` — sensation is not experience. `E -> 0`: `M -> 0`. **All three divergences name real, observable phenomena.**	**PASS — every divergence is nameable**
G4 fixed points	`D = P` gives `M = E`. **Frame matches intake and the moment equals the environment: veridical contact, stated as an equation.**	**PASS — fixed point is interpretable**
G5 identifiability	**FAILS as written.** `(D, P, E)` and `(cD, cP, E)` predict identical `M` for any `c > 0`; only `r` and `rE` are identified. The formula is testable in `r` and `M`, not in `D` and `P` separately, unless an independent operationalisation pins one scale.	**FAIL — identified only up to a common scaling of `D` and `P`**
G6 falsification	Template: preregister an operationalisation of `r` and `M`; fit; freeze; report held-out; beat persistence, direct-input, matched autoregressive and **shuffled-input** baselines [Curzi, 2026i, Nosek et al., 2018]. Killed if held-out `M` is not predicted better than the shuffled-input control.	**DEFINED, NOT YET RUN**
G7 reproducibility	Same operationalisation, different people, different laboratories: do `r` and the scale of `M` recover? **Not yet attempted. This gate decides whether the formula is science, and it is the one most often skipped** [Open Science Collaboration, 2015].	**NOT YET ATTEMPTED**

Two gates fail or stall and both are printed as failures. The identifiability failure is fixable and it says exactly what to do next: pin the scale of `P` by an independent measurement.

18.17.4 A.4 The maximum moment

`M_max = ln N`, the moment expressed in the surplus register

Gate	Test	Result
G1 dimensions	Both sides in nats; `N` a dimensionless count. Consistent.	**PASS**
G2 scale type	`N` is a **count** — ratio-scaled with a true zero — **provided "distinguishable category" has an operational definition.** Without one, `N` is not measured and the bound is not usable. The gate transfers to the operationalisation of the category set, which is Section 8.	**CONDITIONAL — requires a defensible category set**
G3 limiting cases	`N = 1`: `ln 1 = 0`. A perceiver with one category has no moment available — correct and severe. `N -> infinity`: the bound diverges logarithmically, arbitrarily slowly. No finite perceiver approaches it.	**PASS**
G4 fixed points	Attained if and only if `u = 1`, the framings mutually orthogonal [Curzi, 2026]. **The invariant is orthogonality and it is unreachable in practice** — a ceiling, not a target.	**PASS**
G5 identifiability	`N` and `u` are separately recoverable if category membership and pairwise non-parallelism are separately measurable; confounded if `u` is estimated by a procedure that already assumes the category count. **Estimate `u` by error decorrelation, not by labels.**	**PASS with a named confound**
G6 falsification	The bound is a theorem given the axioms plus the normalisation, so not falsifiable *as arithmetic*. The **normalisation** is: if a magnitude-estimation moment scale exceeded `ln N` for a subject whose `N` was independently established, the normalisation is wrong. A runnable experiment.	**NORMALISATION FALSIFIABLE; TEST SPECIFIED**

Gate	Test	Result
G7 reproducibility	Requires `N` to recover across instruments for the same person — specific objectivity, with Rasch separability as the standard [Rasch, 1960]. Not yet attempted.	**NOT YET ATTEMPTED**

18.17.5 A.5 Stevens' power law

`psi = k_S I^a`, equivalently `log psi = log k_S + a log I`

Gate	Test	Result
G1 dimensions	`I` must be a ratio to a stated reference before exponentiation; `I^a` of a dimensioned quantity is ill formed. `k_S` carries the units of `psi`.	**PASS if reference declared; a common silent failure otherwise**
G2 scale type	`I` ratio-scaled with a true zero. `psi` **ratio-scaled** — the point of magnitude estimation, which asks subjects for ratios directly and validates them by cross-modality matching. Multiplication and division of `psi` are permissible. **The claim is contested** [Luce, 2002, Narens, 1996].	**PASS as claimed; the gate Fechner fails and Stevens claims to pass, disputed**
G3 limiting cases	`I -> 0` gives `psi -> 0` for all `a > 0`: no stimulus, no sensation. `I -> infinity` diverges; real modalities terminate at damage or saturation, so an upper working range must be declared. At `a = 1` the law degenerates to proportionality, as observed for visual length.	**PASS, upper bound declared per modality**
G4 fixed points	`I = I_ref` gives `psi = k_S`. In log-log coordinates the invariant is the intercept: the anchor stimulus. **The anchor is chosen by the experimenter, which is exactly why instruction affects the exponent.**	**PASS, dependency named**
G5 identifiability	`a` is the log-log slope and `k_S` the intercept; separately recoverable from two or more levels spanning a decade. Cross-modality matching independently recovers the ratios of exponents, over-identifying the system.	**PASS — over-identified, which is strong**
G6 falsification	Killed in a modality if log-log plots are reliably curved over the declared range, or if the matching slope departs from the predicted `a_j/a_i` beyond measurement error. Both tests have run for decades.	**PASS — a live falsification test**

Gate	Test	Result
G7 reproducibility	Exponents recover across laboratories **within a fixed procedure** and fail to recover across procedures [Poulton, 1979, Teghtsoonian, 1971]. Honest verdict: constants of a specified method, not of a modality. **Report the procedure or the number is not a measurement.**	**CONDITIONAL PASS — procedure-bound**

18.17.6 A.6 The essence law

`f(n) = k n^2` — **POSTULATE about essence structure**

Gate	Test	Result
G1 dimensions	`n` is a dimensionless index; `k` carries hertz; `f` is in hertz. Both sides agree **only because `k` was given hertz**, which is postulate P6 and carries no external support. `k` must be published with its units or the equation is incomplete.	**PASS given P6; P6 itself is POSTULATED**
G2 scale type	**The crux.** `n` is ORDINAL. `f` is RATIO-scaled. The law promotes an ordinal to a ratio scale by postulating the spacing. Permissible as a model, impermissible as an inference from data [Michell, 1997, Stevens, 1946]. **Must be badged in print wherever intervals are interpreted.**	**PROMOTION BY POSTULATE — badged, not passed**
G3 limiting cases	`n = 0` gives `f = 0`: there is no zeroth essence, and the law says so by returning silence. `n = 1` gives `f = k`, the ground. `n -> infinity` diverges quadratically, leaving the audible band quickly — at `n = 24`, `f` is already 278 Hz and at `n = 203` it is 19.9 kHz. Transposition is the required remedy.	**PASS, with a declared working range**
G4 fixed points	`f(n) = n` has solutions at `n = 0` and `n = 1/k = 2.0685`, neither meaningful, since index and frequency are not the same kind of quantity. **The meaningful invariant is `f(1) = k`**, and the deeper one is that `k` cancels from every ratio (Section 14.7).	**PASS — invariant is the ground, not a self-map**
G5 identifiability	`k` and the exponent are separately recoverable from any two published values: two points determine both slope and intercept in log-log coordinates. **The exponent was recovered this way from the source document and matched a pure square to 14 significant figures.**	**PASS — and was actually performed**

Gate	Test	Result
G6 falsification	As arithmetic, not falsifiable. As a claim about category structure, falsifiable only once the relation between category types has an independent operationalisation with a measurable magnitude. **No such operationalisation exists and the honest statement is that the postulate is not yet testable** (H6).	**NOT CURRENTLY FALSIFIABLE — stated plainly**
G7 reproducibility	The arithmetic reproduces exactly. **The INDEX does not reproduce automatically** — it depends on the ordering, and the ordering is where reproducibility must be established, by ES-R3 inter-rater agreement.	**CONDITIONAL — the gate transfers to the ordering**

18.17.7 A.7 The interval law

`f(m)/f(n) = m^2/n^2`

Gate	Test	Result
G1 dimensions	A ratio of two frequencies: dimensionless. `k` cancels identically, so the result carries no dependence on the chosen reference **and no dependence on P6**.	**PASS — and it is the one result that does not consume P6**
G2 scale type	The ratio is well formed on a ratio scale, which `f` is by construction. **But the ratio-scale status of `f` was itself manufactured by the postulate (Section 3.5), so the interval inherits that status and not a stronger one.**	**PASS given the postulate, badged**
G3 limiting cases	`m = n` gives 1: unison. `n -> 0` diverges, and there is no zeroth essence, so the divergence is outside the domain. `m/n -> infinity` gives an unbounded interval, correctly: distant indices are distant pitches.	**PASS**
G4 fixed points	The ratio equals 1 if and only if `m = n`; it equals a whole number of octaves if and only if `m/n` is a power of two. **The invariant of the folding operation is the odd part**, which is CC-R2.	**PASS — fixed point is the class**

Gate	Test	Result
G5 identifiability	The interval determines `m/n` uniquely, by taking the square root. It does not determine `m` and `n` separately — `6:4` and `3:2` give the same interval. **The lattice is identifiable up to the reduced index ratio and no further**, which is exactly why consonance depends on the reduced ratio alone.	**PASS with a named degeneracy that turns out to be structural**
G6 falsification	Follows the essence law: falsifiable only when the underlying spacing postulate becomes testable. The interval law adds no new empirical content beyond it.	**INHERITS THE ESSENCE LAW'S STATUS**
G7 reproducibility	Exact arithmetic, reproducible anywhere. Verified independently in Section 7.1 against five recorded examples — all five agree.	**PASS**

18.17.8 A.8 Pitch-class folding by odd part

`class(n) = odd part of n`; the number of classes over `1..N` is `ceil(N/2)`

Gate	Test	Result
G1 dimensions	Both index and odd part are dimensionless integers; the class count is a dimensionless count. Consistent.	**PASS**
G2 scale type	The odd part is a NOMINAL label — a class identifier, not a magnitude. **Do not do arithmetic on class numbers.** The class's octave position in cents is interval-scaled and may be subtracted. The count `ceil(N/2)` is a ratio-scaled count. **Three different scale types in one construction, and they must not be mixed.**	**PASS with an explicit three-way scale warning**

Gate	Test	Result
G3 limiting cases	`N = 1`: one class, the ground. `N = 2`: still one class — indices 1 and 2 are two octaves apart and harmonically identical. `N -> infinity`: classes grow as `N/2` while the pitch span grows only logarithmically, which is why collisions appear and why they are forced rather than accidental.	**PASS — and the limiting behaviour explains the collisions**
G4 fixed points	The odd indices are the fixed points of the folding map: `class(m) = m` if and only if `m` is odd. **The class representatives are exactly the odd indices**, which Section 7.7 recovers independently by minimising dissonance.	**PASS — fixed-point set is interpretable and useful**
G5 identifiability	The class does not identify the index: `{3, 6, 12, 24}` all map to class 3. Recovering `n` from its class requires the octave count `a`, so the pair `(a, m)` is a bijection with `n`. **Publish `(a, m)`, not `m` alone, if the index must be recoverable.**	**PASS given `(a, m)`; FAILS given `m` alone**
G6 falsification	A theorem of arithmetic (Theorem 7.1), proved and verified by enumeration for `N = 1..24`, all 24 agreeing. Not empirically falsifiable; arithmetically checkable and checked.	**N/A (theorem) — verification reported**
G7 reproducibility	Exact and machine-independent. As with the essence law, reproducibility of the *index* transfers to the ordering (ES-R3).	**PASS for the arithmetic; conditional on ES-R3 for the input**

18.17.9 A.9 Octave transposition

`f_audible = f 2^j`, `j` an integer

Gate	Test	Result
G1 dimensions	Hertz times a dimensionless power of two gives hertz. Consistent.	**PASS**
G2 scale type	`f` is ratio-scaled; multiplication by a positive constant is permissible and preserves all ratios. `j` is an integer count of octaves.	**PASS**

Gate	Test	Result
G3 limiting cases	`j = 0` is the identity. Large positive `j` leaves the audible band above and large negative `j` below. Both are outside the working range, which is why the band must be declared.	**PASS with declared band**
G4 fixed points	The pitch class is invariant under the operation for every `j`. **That invariance is the entire justification for the method — and it holds only for integer `j`.**	**PASS**
G5 identifiability	`f` is recoverable from `f_audible` only if `j` is published. Without `j` the raw essence frequency is unrecoverable and the value cannot be checked. **This is why CK-6 and ES-10 exist.**	**PASS if `j` published; FAIL otherwise**
G6 falsification	An exact operation, not an empirical claim. Checkable and checked in Section 8.3. **The *interpretation* — that octave equivalence carries meaning for a non-pitch quantity — is a separate claim and is unsupported** [Shepard, 1982, Hoeschele et al., 2012].	**N/A (definition); the interpretation is badged category 3**
G7 reproducibility	Exact given `f`, the band, and the rule for choosing `j`. **The rule must be published, not just the `j`** — "the lowest `j` placing `f` in `[110, 220)`" is a rule; "`j = 4`" is a value.	**PASS if the rule is published**

Chapter 19

The Moment Equation and the Calibration Ratio: A Fixed Point at Contact, a Self-Imposed Demotion to Ordinal Scale, and the Price of Stored Error

Abstract. The moment equation $M = (D/P)\,E$ is put through this corpus's own validity protocol and **fails** its scale-type gate. The expression contains a division and a multiplication, so frame D, intake P and environment E must each sit on a ratio scale with a true zero; where only ordinal data exist the equation is demoted to a heuristic ordering and may not be printed as a fraction or computed. That demotion is the headline result, and it is Michell's remedy applied once. What survives is exact: four limiting cases, unit elasticity of magnitude one in every term, a ceiling $0 \le M \le \ln N$ inherited from the surplus bound, and a fixed point at $D = P$, equivalently $r = 1$, at which $M = E$ for every E. Distortion is $\ln r$, and the postulated correction map $\ln r_{t+1} = (1-a)\ln r_t$ is a first-order linear recurrence: attracting for $0 < a < 2$, neutral at $a = 0$, divergent at $a \ge 2$; from 0.8122 nats, six cycles at gain 2.3 reach 3.9203. In a constructed seeded audit with 25 per cent off-model items, $r = 50/29 = 1.724138$ and the confabulated fraction $1 - 1/r = 0.4200$ is exactly the false-discovery rate, with 63 per cent of the steady state stored error; a constructed fourteen-year inspector series runs $r = 0.353,\ 1.048,\ 1.900$, off-model detection falling to 13 per cent against a novice's 33. Several printed figures in both are shown not to reconcile. Contact is not an unqualified goal: at $r = 1$ in a saturating environment headroom is zero by exact cancellation.

Keywords: moment equation; calibration ratio; scale demotion; seeded audit; off-model items; veridical throughput; stored error; correction gain

Series: B. Engineering the Surplus **Paper:** CCP-019

19.1 1 Introduction

A framework that models how a group or an institution accumulates durable capability takes its input from the environment. The dynamical core of this corpus, the storage law `Q_{t+1} = (1-delta)Q_t + eta S_t - C_t`, integrates whatever arrives and does not inspect it [Curzi,

2026a]. That is a deliberate simplification, and it has a consequence the framework must answer for: a room is not the same room to two people standing in it, and the quantity the storage law integrates is not the room but the observer's reading of it. Something stands between the environment and the input.

This paper names that thing, gives it an equation, and then submits the equation to the corpus's own validity protocol — where **it fails.** The proposed relation is

```
M = (D / P) . E
```

with `D` the frame the observer brings, `P` the intake the observer takes, `E` what is available at that place and time, and `M` the moment as lived, which is the quantity the storage law integrates. The expression contains one division and one multiplication. Under the standard result that permissible arithmetic is fixed by scale type [Stevens, 1946], division and multiplication require all three inputs on a ratio scale with a true zero and a defensible unit. **No instrument places `D` and `P` on a common ratio scale with a shared true zero. The equation is therefore demoted to a heuristic ordering and may not be printed as a fraction, computed, or used to generate a number.** That demotion is the paper's headline result, it is a negative one, and it is reported here rather than in a footnote because a validation protocol that never demotes anything is decorative.

The demotion is precisely the remedy Michell [1997] prescribes: psychology routinely assumes that its attributes have quantitative structure and measures them without testing whether they do. The corpus applied that remedy once, to its own keystone formula. The honest position — argued in Section 15 and conceded there rather than defended — is that the remedy should be applied consistently rather than once, and that the conjoint-measurement test which would settle the question [Luce and Tukey, 1964, Krantz et al., 1971] has not been run for any quantity in this corpus.

What survives the demotion is exact and is not small. Four limiting cases, each naming a recognised condition. Unit elasticity of magnitude exactly one in every term, so no input's error is amplified relative to another's. A ceiling `0 <= M <= ln N` inherited from the surplus bound [Curzi, 2026], which converts the one divergence in the formula from an infinity into a saturation. And a fixed point at `D = P`, equivalently `r = 1`, at which the moment equals the environment for every environment. **All four need only ordinal information about `D` and `P`, and none of them is touched by the demotion. What is demoted is magnitude, not structure.**

Four contributions are new in this paper relative to the surrounding literature rather than relative to the framework. First, the demotion rule itself, stated as a rule and applied to the corpus's own equation before it is applied to anyone else's. Second, the identification of the observer's whole contribution with one dimensionless composite `r = D/P`, together with the exact statement of what that costs: `(D, P, E)` and `(kD, kP, E)` are observationally identical, so only the ratio is recoverable. Third, the correction map `ln r_{t+1} = (1−a) ln r_t` and its stability band, which is a first-order linear recurrence and inherits the whole standard theory of such recurrences [Elaydi, 2005, Åström and Murray, 2008] — a **shared mathematical form, not a derivation from physical premises.** Fourth, veridical throughput `V = E min(r,1)` and its complement the confabulated fraction `1 − 1/r`, which separate what an observer stored from what the environment supplied, and which show that the fixed point is not merely where the reading is undistorted but where the earned content of the reading is largest.

There is also a counter-result, and the paper states it against its own interest. **Contact is not an unqualified goal.** At `r = 1` in a maximally rich environment the input to the storage law sits at its ceiling, the store settles at its maximum, and the headroom gauge reads exactly zero by cancellation. A perfectly calibrated observer maximises fidelity and has no room left to build. Worse, an observer optimising raw stored order is driven *away* from `r = 1` altogether, because the raw track keeps climbing in `r` while the veridical track is flat above the fixed point — and the correction bill only ever catches up when the correction-load coefficient exceeds the decay rate.

Two worked audits carry the numbers. A seeded code-review packet with 25 per cent off-model items returns `r = 1.724138`, a false-discovery rate of 0.4200 that is exactly `1 − 1/r`, and

a steady state of which 63 per cent is stored error; correction at gain 0.35 takes the reviewer to `r = 1.032`, cuts false discovery by 92.6 per cent and raises true defects found from 29 of 40 to 31 of 40, while the metric the organisation actually watches — on-model detection — falls from 93 to 90 per cent. A fourteen-year building inspector series returns `r = 0.353` with a negative steady state at week two, `r = 1.048` at minimum loss in month nine, and `r = 1.900` at year fourteen, with on-model detection at its career maximum and off-model detection below a nine-day novice's. **Both are constructed examples, not collected datasets**, and Section 12.3 prints the numbers in them that do not reconcile.

The framework's static half — the surplus functional `F(u) = ln(1 + (N−1)u)` with its affine effective count `N_eff = 1 + (N−1)u` — is presupposed rather than re-derived [Curzi, 2026]. The founding reflexive law, written as the fixed point `T(x) = W(x, M(x))` of a composition of seeing and responding, is presupposed from Curzi [2026] and is the reason Section 3.7 must confront the reflection problem before anything else in this paper is interpreted empirically. The contraction theorem the correction map leans on — a contraction with modulus `q < 1` on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least `q` each pass — is stated exactly in Curzi [2026].

Sections 2 and 3 state the hypotheses and the framework, including the scale-type gate and the identification problem. Section 4 develops the four limits, the composite ratio and the ceiling. Section 5 develops the fixed point and its dynamics. Sections 6 to 8 take the calibration ratio on its own: its measure, its diagnosis, its composition, its drift, and the cases in which `r = 1` is the wrong target. Section 9 reports the seven validation gates and their verdicts. Sections 10 and 11 collect the specifications and the named failure modes. Section 12 reports the two worked audits and the arithmetic in the source that does not reproduce. Sections 13 to 16 discuss what follows, verify the chain, state the limitations, and conclude.

19.2 2 Hypotheses

The limiting cases, the fixed point, the elasticities and the stability band are theorems about an equation and cannot be falsified by observation; they are labelled as derivations where they appear and they are re-derived in Section 14. What can be falsified is the claim that the equation describes a real observer. The hypotheses below separate the two, and each names the observation that would defeat it.

H1 is the one this paper tests, and it failed. It is reported here in the same plain language a success would have been given.

> **H1 — The moment equation is a quantitative law.** Frame, intake and environment admit measurement on a common ratio scale with a shared true zero and a defensible unit, so that `M = (D/P).E` may be printed as a fraction and computed. **Predicted:** an operationalisation of `D` and `P` linked across item sets, on which a reported `r` is invariant to an admissible change of items, and on which the statement "this observer's distortion is twice that one's" has a referent. **Falsified by:** the absence of any such instrument; or a demonstration that the magnitude — not merely the value — of `r` moves when the item set is exchanged for an equivalent one. **Result: FAILED.** No instrument links `D` and `P` across item sets. Counting recorded items supplies a zero and a unit *within* one item set and supplies no linkage *between* two, so the ratio moves when the items move. The magnitude of `D/P`, the distortion `Delta = |ln r|` as a number, and the quantitative form of the join from the moment to the store are withdrawn as measured quantities [Curzi, 2026t]. Section 3.5 states the demotion; Section 12.4 audits every operation in this paper that the demotion does not license; Section 15 concedes what the failure costs.
>
> **H2 — The fixed point is the unique maximum of veridical throughput.** The part of the moment the environment actually supplied, `V = E min(r,1)`, is maximised at `r = 1` and only there. **Predicted:** in a seeded audit containing off-model items, an

observer's true-item detection rises with `r` up to `r = 1` and is flat above it, while the false-discovery rate rises monotonically above it, tracking `1 − 1/r`. **Falsified by:** an audit in which both off-model detection and false-discovery rate improve as `|ln r|` grows; or one in which true detection continues to rise above `r = 1`, which would mean the assertion record is adding real conditions the intake record never held.

H3 — Correction converges only inside a stability band. A deliberate correction regime acting on a fraction `a` of the observed discrepancy per cycle converges to contact if and only if `0 < a < 2`. **Predicted:** geometric convergence with half-life `ln 2 / (−ln|1−a|)` cycles; monotone approach for `a < 1`; alternating sign with decreasing magnitude for `1 < a < 2`; a standing offset at `a = 0`; alternating sign with **growing** magnitude at `a >= 2`. **Falsified by:** a documented correction regime whose measured effective gain exceeds 2 and which nonetheless converges; or one inside the band that diverges in the absence of an external forcing term.

H4 — Off-model content is necessary for detecting over-framing. An audit built only from items inside the observer's standing model cannot separate an over-framed observer from a calibrated one. **Predicted:** sensitivity to over-framing rises with the off-model share of the seeded item set, and is zero at an off-model share of zero; an all-on-model audit returns its highest scores for the most over-framed observer. **Falsified by:** an audit with no off-model content that nonetheless separates the two states, or a measured off-model detection rate that does not fall as `r` rises.

H5 — Detection asymmetry grows with tenure in a stable frame. An observer working a fixed model in a moving environment gains on-model detection and loses off-model detection. **Predicted:** on-model hit rate rising and off-model hit rate falling with tenure, false-discovery rate rising, and time per case falling, at constant measured environment. **Falsified by:** a cohort in which off-model detection is flat or rising with tenure while the environment is independently shown to be unchanged. This hypothesis is adjacent to the myopia-of-learning result [Levinthal and March, 1993] and to the exploration/exploitation trade [March, 1991], both of which predict the same direction from a different mechanism; a confirmatory observation therefore does **not** discriminate this framework from theirs, which is the demand Durlauf [2005] makes of any formal social model.

H6 — Stored error is real and its bill arrives late. The gap between order stored on the moment as lived and order stored on the veridical part grows with `r`, and it is paid as correction cost on a lag. **Predicted:** rework, dispute and repair costs rising with a lag of several time constants behind a rise in `r`, with no change in the environment; and, where the correction-load coefficient `kappa` exceeds the decay rate `delta`, raw stored order eventually falling in `r` rather than rising. **Falsified by:** an over-framed body whose correction costs are stationary over a horizon of several `1/delta`.

Two claims in this paper are not hypotheses and are labelled accordingly. The demotion rule of Section 3.5 is a **consequence of a definition** — of what scale type licenses which operation — and not an empirical claim; what is empirical is the premise that no linking instrument exists, which is stated as postulate P3 in Section 14.2. The band table of Section 6.3 is a **convention of this framework**, chosen for round numbers in nats, and is not a measured set of thresholds; it must never be reported as one. The identification of the observer's whole contribution with a single non-negative scalar `r` is a **design choice**, adopted for tractability and for the fixed-point guarantee, and Section 3.6 states what it costs.

19.3 3 Framework: the equation, the scale-type gate, and what identification is available

19.3.1 3.1 What is presupposed

Three results from elsewhere in the corpus are used and not re-derived.

The surplus functional and its ceiling. `F(u) = ln(1 + (N−1)u)`, uniquely characterised by angle-only dependence, zero at parallel, an affine effective count and a normalisation, with `N_eff = 1 + (N−1)u` and `0 <= F(u) <= ln N` [Curzi, 2026]. The ceiling is the result this paper's own ceiling inherits from. **`N_eff = 1 + (N−1)u` is a linear interpolation between 1 and `N` and is not a Hill number**; `1/HHI` and `exp(H)` are the Hill numbers of order two and one and are fully grounded [Hill, 1973, Jost, 2006], and neither of them is the quantity used here. Every effective count in this paper is the affine one, and the ceiling `ln N` uses the **raw category count `N`**, not `N_eff`, because `F(u) <= ln N` is attained only at `u = 1` where the two coincide.

The storage law and its steady state. `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t` is a first-order linear difference equation whose update map is a contraction for any `delta > 0`, with steady state `Q_inf = (eta S − C)/delta`, and for which `delta = 0` is a pathology rather than an ideal: no contraction, no unique fixed point, and an error-persistence column of all ones, so no mistake ever decays out of the stock [Curzi, 2026a]. This paper substitutes `M` for `S` and inherits everything.

The contraction theorem. A contraction with modulus `q < 1` on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least `q` each pass [Curzi, 2026, Banach, 1922].

The reflexive law. The corpus's founding formal claim is the loop `T(x) = W(x, M(x))`, with the self-fulfilling prophecy identified as a fixed point of that composition of seeing and responding [Curzi, 2026]. It is presupposed, and Section 3.7 states what it costs.

19.3.2 3.2 The statement, badged

> **THE MOMENT EQUATION** *(postulate; this work; no external support).* For an observer in contact with an environment, the moment as lived is the environment scaled by the ratio of the frame the observer brings to the intake the observer takes:

```
M = (D / P) . E

D = perspective  : the structure the observer imposes —— model, expectation,
                   prior, doctrine, training, the shape the observer is
                   prepared to find
P = perception   : the structure the observer takes in —— intake, evidence
                   actually registered, the raw admission of what is there
E = environment  : what is actually available to be had at this place and time
M = moment       : the situation as it is lived; the input the observer's own
                   storage law will integrate
```

This is a postulate and not a theorem. **The corpus's own coverage audit records that no external literature was found for either this relation or the calibration ratio derived from it**; the nearest relatives are signal-detection theory, for which this bibliography holds no entry, and convergent/discriminant construct validation [Campbell and Fiske, 1959, Cronbach and Meehl, 1955], neither of which states this relation. Nothing established elsewhere in the corpus depends on it. It is offered as the load-bearing hypothesis of the internal instrument, with the badge on, and Section 9.2 names exactly what would kill it.

19.3.3 3.3 Why a ratio and not a difference

The first structural decision is that `D` and `P` are combined by division rather than by subtraction, and it is not free.

Four quantities in this framework are logarithmic: surplus `F = ln N_eff`, pitch `p = log2 f`, sensation `S = k ln(I/I_0)`, and the moment. Taking logarithms,

```
ln M = ln D − ln P + ln E
```

A ratio in the world is a difference in the lens. If `D` and `P` were combined additively the moment would not compose with surplus, with pitch or with sensation, and none of the framework's conversion keys would apply to it. By dividing, the moment joins the register the other three occupy, and every operation available there — the log-ratio distortion measure of Section 6, additive composition along a chain, the log-space contraction — becomes available here.

There is a second reason and it is empirical rather than architectural. Weber's law states that the just-noticeable difference is a constant *ratio* and not a constant difference [Weber, 1834], and Fechner's integration of it yields `S = k ln(I/I_0)` [Fechner, 1860, 1966]; the corpus states that integration exactly, as the Lens Theorem, and verifies it numerically to an absolute error of `1.02e−08` against `ln 100` [Curzi, 2026n]. If the instrument being modelled reports ratios, a model of it built out of ratios is the natural choice.

Two cautions belong with that argument and are stated here rather than in Limitations. First, the psychophysical warrant is contested: the power law and the logarithmic law were disputed for a century [Stevens, 1957, Krueger, 1989], and the fitted exponent depends on the experimenter's chosen stimulus range [Teghtsoonian, 1971]. Second, and more usefully, the logarithm does not need the perceptual story at all. Peters and Gell-Mann [2016] derives a logarithmic evaluation function from multiplicative dynamics with **no psychological assumption whatever**, which is a stronger warrant for the framework's logarithms than Fechner is, because it does not require the perceptual account to be true. Where this paper needs a logarithm — in the distortion measure of Section 6.2, whose whole content is that a multiplicative error should be reported additively — it is the multiplicative argument that carries it, and Fechner is a corroborating parallel and not the load path.

19.3.4 3.4 Dimensions and the commensurability requirement

For `M = (D/P).E` to be well formed, `[D] = [P]`. Their ratio is then dimensionless and

```
[M] = [E]
```

The moment inherits the units of the environment. `M` and `E` are the same kind of thing on the same scale, and the observer's contribution is a pure dimensionless multiplier on it. A moment is not a different species of object from the environment that produced it; it is that environment, scaled.

This constrains what `D` and `P` may be. They must be commensurable: the same quantity, measured the same way, once for the frame and once for the intake. **Years of training over number of things noticed is not a ratio and must not be written as a fraction.** The operational requirement is specification SP-1 in Section 10: `D` and `P` must both be *recorded items per unit of environment present*, or both *precision in the same units*, or both *bits of asserted structure per bit of available structure*. Anything else fails the dimensional gate.

For the units of `E` and therefore of `M` this framework adopts **nats of usable structure**, so that `M` substitutes directly for `S_t` in the storage law and inherits the ceiling `0 <= M <= ln N` of Section 4.6. The nat here is the unit of **Shannon entropy** — a dimensionless information measure over a distribution on `N` categories [Shannon, 1948, Cover and Thomas, 2006] — and **not** thermodynamic entropy, which carries `[energy][temperature]^−1`. The two are different objects and this framework identifies the moment with neither. Other unit choices are possible and must be declared.

19.3.5 3.5 The scale-type gate, and the demotion

This is the gate the formula fails, and the failure is the paper's principal result.

Permissible arithmetic is fixed by scale type [Stevens, 1946], and the condition is a condition on the transformations that leave the measurement structure invariant [Narens, 1996, Luce, 2002]:

Operation	Minimum scale required
Classification, identity	Nominal
Ranking, ordering, median, rank correlation	Ordinal
Addition, subtraction, arithmetic mean, standard deviation	Interval
Multiplication, division, ratio, geometric mean, logarithm	**Ratio, with a true zero**

`M = (D/P).E` contains one division and one multiplication. **Therefore `D`, `P` and `E` must each be measured on a ratio scale with a meaningful zero, or the expression is not an equation.** If the best available data on `D` is a five-point self-report of how strongly the observer expected what they found, the ratio `D/P` has no referent: a 4 is not twice a 2 on a summated rating item [Likert, 1932], because such an item carries neither a true zero nor a defensible unit. An interval scale can be built without asking respondents to use numbers at all [Thurstone, 1927], but an interval scale is still one gate short of what division requires.

> **DEMOTION RULE (R1).** Where only ordinal data exist for `D` or `P`, the moment equation shall be written and read as a **heuristic ordering** — *more frame relative to intake produces a larger and more model-shaped moment* — and shall **not** be printed as a fraction, computed, or used to generate a number. The demoted form remains useful. It is not the equation.

The rule is not hypothetical. It fires on this framework's own formula. The corpus's validation protocol, applied in full, returns DEMOTE on the scale-type gate for `M = (D/P).E` as the equation stands: `D` and `P` are not instrumented on a common ratio scale with a shared true zero, so the magnitude of `D/P`, the distortion as a number, and the quantitative moment-to-storage join are withdrawn as measured quantities [Curzi, 2026t]. That withdrawal is one entry in a standing register of the corpus's own demotions and inversions [Curzi, 2026], and the same verdict is recorded again where the corpus's personal formulary demotes five of its own nine entries on scale type, withdrawing `D/P` magnitude and the quantitative moment-to-storage join while retaining the limiting cases, the `D = P` fixed point and the over/under-framed ordering [Curzi, 2026q].

> **WHAT IS NOT REPORTABLE AS A MEASURED QUANTITY.** The *magnitude* of `D/P`; the distortion `Delta = |ln r|` as a number; and the quantitative form of the join from the moment to the store. **Every figure of that kind printed anywhere in this paper is an illustrative construction demonstrating the structure of the formula**, and is labelled as such at each occurrence and again in the audit of Section 12.4.
>
> **WHAT SURVIVES INTACT.** The four limiting cases (Section 4); the fixed point `D = P` and its uniqueness (Section 5.1); and the three-way ordering over-framed / calibrated / under-framed (Section 7.1). **All three need only ordinal information, and none is touched by the demotion. What is demoted is magnitude, not structure.**

The route out is procedural and it is costed. Given a badged additivity postulate, if `D` and `P` are measured as person-measures on one common linked scale under a Rasch model [Rasch, 1960, Andrich, 1988], then `Delta = |ln(D/P)| = |theta_D − theta_P|` is genuinely ratio-scaled, because the separability of the Rasch model supplies the true zero that division requires and

fixes the unit up to choice. The measurement cost is the honest part: **approximately 123 items per term to establish a half-logit distortion at 95 per cent, and about 31 items for a full logit** [Curzi, 2026t]. A twelve-item pair of scales cannot resolve anything below roughly 1.6 logits, and a practitioner reporting `r = 1.2` from one is reporting noise.

A second route out exists and is weaker than it looks. Magnitude estimation is claimed to yield genuine ratio-scale measurement of subjective quantity from human respondents [Stevens, 1957, 1975], which would supply the true zero directly. **But the method's own literature contains every failure mode it is exposed to** [Poulton, 1968, 1979], the exponent it recovers depends on the range the experimenter chose [Teghtsoonian, 1971], and self-report instruments in general are exposed to question wording and to focusing effects [Schwarz, 1999, Kahneman et al., 2006]. This paper therefore names magnitude estimation as an available route and does **not** treat it as a settled one.

Read the demotion as the protocol working. A validation gate that demotes its author's own keystone formula, names the exact quantity it kills, prices the study that would restore it, and leaves the structural content standing is doing the job it was written to do. The alternative — printing `r = 1.33` and hoping nobody asks what the unit is — is the practice the protocol exists to replace. The corpus has done this elsewhere and against its own interest: a chapter that refuses a fabricated data point, rejects its most attractive reading at 1,365 to 1, and badges a 25 per cent figure as an artefact of a chosen synthetic ratio rather than a constant [Curzi, 2026] is the same discipline applied in another domain, and it bounds what this paper may claim from constructed examples.

19.3.6 3.6 Michell's objection, met head on

The strongest single objection to this framework is not that the moment equation lacks support. It is Michell [1997]'s: that psychology systematically fails to test whether its attributes are quantitative at all before measuring them, and that the resulting enterprise is pathological science rather than measurement [Michell, 1999]. The corpus's founding move — treating comfort as a measurable quantity — is exactly the move Michell indicts.

This paper concedes the objection and does not answer it. Three things are true at once and all three are stated.

First, **the demotion of Section 3.5 is Michell's remedy applied once.** The framework declined to compute a ratio it could not license, named the quantity it withdrew, and priced the instrument that would restore it. That is the correct response and it is worth more than a confident number.

Second, **applying the remedy once is not applying it.** The same test that demotes `D/P` also bears on `u`, on `Q`, on `E` and on comfort itself, none of which has been subjected to it. A framework that demotes its keystone formula on scale type and then takes arithmetic means of everything else has not adopted a standard; it has made an exception. Section 12.4 therefore audits **every** operation in this paper against the gate, including the ones that survive.

Third, **the test that would settle the question has not been run.** Conjoint measurement is the formal procedure for deciding whether an attribute has additive quantitative structure at all [Luce and Tukey, 1964, Krantz et al., 1971], and it has not been applied to `D`, to `P`, to `E`, to `M`, or to any other quantity in this corpus. Until it is, the quantitative treatment of every one of them is a **stipulation and not a finding**, and the Rasch route of Section 3.5 is a proposal rather than a result. Reporting a coefficient of internal consistency in place of that test would not close the gap [Cronbach, 1951, Sijtsma, 2009].

19.3.7 3.7 Identification: three objections, none of them answered

Three identification objections bear on this paper. All three are stated here, in Method, and none is answered.

The reflection problem. The corpus's founding law is the reflexive loop `T(x) = W(x, M(x))` [Curzi, 2026]: the group influences the member and the members constitute the group. That is a linear-in-means reflexive specification, and Manski [1993] proves that the endogenous social effect in such a specification is **not identified from observational data**. It cannot be separated from contextual effects or from correlated unobservables, however much data is collected. Shalizi and Thomas [2011] show the same confounding for network data: influence and selection are generically confounded in observational settings.

This paper's own structure walks into it from the other side. The moment equation assumes `E` is exogenous — that the environment would have been the same had the observer concluded otherwise — and failure mode FM-10 in Section 11 excludes the cases where it is not: the negotiator, the forecaster whose forecast moves the market, the clinician whose diagnosis becomes the treatment. **That exclusion is not a solution to the identification problem. It is a confession that the framework has none**, and it removes from scope precisely the class of situations in which the corpus's founding reflexive law is operating.

Which remedy this paper adopts: none. The available options are randomisation, an instrument, partial-population variation, or stipulating the reference group and defending the stipulation [Brock and Durlauf, 2001]; the last is the most realistic route here and the seeded audit's fixed, externally adjudicated item set is a step toward it rather than a substitute for it. Because no remedy is adopted, **every applied reading of `r` in this paper is a modelling position and not an identified empirical result**, and it is stated as one in Section 15. That reflexivity is a real phenomenon rather than a formal curiosity is not in doubt: MacKenzie and Millo [2003] documents a reflexive fixed point in the field, with dates and price data, in which a pricing model that did not fit the market when published came to fit it after traders adopted it. Soros [2013] is the closest existing named doctrine to the corpus's reflexive law and supplies the concept without an equation.

The observer as its own instrument. This paper's specific version of the problem is sharper than the general one. The quantity being estimated — the frame the observer brings — is the same quantity doing the estimating, which is why Section 7.3 concludes that calibration cannot be self-assessed. The supporting result is that people cannot reliably report the processes producing their own judgements [Nisbett and Wilson, 1977]; the framework's response is to make every estimator external (SP-8). But the external estimator is not free of the problem either, because prior expectation measurably shapes what is perceived and reported through identified channels [Rosenthal and Fode, 1963, Rosenthal and Jacobson, 1968], and the adjudicators who build the seeded item set are themselves observers with frames. **The framework's honest position is that the audit relocates the problem to a party whose frame is at least declared, and does not remove it.** The size of the expectancy effect is itself contested and is smaller than its popular reading [Jussim and Harber, 2005]; and the act of measuring an observer changes the observer being measured [McCambridge et al., 2014], which is a direct threat to any claim that a re-audit measures the same person.

The Lucas critique applied to this framework's own parameters. Lucas [1976] is the corpus's own reflexivity turned on the corpus: parameters estimated in a population that does not know the law cannot be assumed stable in a population that has read it, and classifications of people interact with the people classified [Hacking, 1995]. This paper asserts in Section 8.4 that a throughput metric drives `r` upward across a whole population, and prescribes changing the metric — and then treats the correction gain `a` and the drift rate `w` as constants over the correction horizon. **Both cannot hold without qualification.**

The position taken here is the narrow one, and it is a real restriction rather than a disclaimer. `a` and `w` are treated as stable **only over one declared audit cycle, and only while `r` and its components are not themselves reported targets.** The moment the off-model detection rate enters a quarterly report — which is exactly what Section 12.2 prescribes — it becomes a measure that is a target, and a measure that becomes a target ceases to be a good measure [Goodhart, 1975, Campbell, 1979]. The worked example in Section 12.1 shows this happening on schedule and does not conceal it. Where `r` is reported upward, the framework's own reflexivity applies to it and the constancy assumption fails. Prescriptive frameworks are not

neutral about what they make true [Ferraro et al., 2005], and a framework that tells an organisation its inspectors are two-thirds confabulating should expect that statement to change the inspectors.

19.3.8 3.8 Standing honesty clauses

1. **The moment equation is a postulate, not a theorem.** It is not derived from the corpus's established results and none of them depends on it.
2. **The correspondence with predictive processing (Section 5.7) is a correspondence between formalisms and nothing more.** It is badged in Section 14.7 as an analogy in the strict sense: a suggestive correspondence with no formal warrant, legitimate as motivation and never as evidence. **This bibliography contains no entry for the predictive-processing literature**, so the correspondence is stated here without external support and must not be read as borrowing that programme's authority.
3. **`M` is the moment as lived, not the moment as true.** The formula does not compute what is really there; it computes what the observer ends up with. Section 5.6 makes the difference quantitative.
4. **The ceiling `M <= ln N` is a consequence of a theorem plus a convention**, not a free-standing discovery: the bound `0 <= F <= ln N` is the surplus law's range result [Curzi, 2026] and the mapping of `M` into that register is a normalisation choice of this framework.
5. **No claim is made that `D`, `P` or `E` correspond to anything neural.** They are parameters of a descriptive model of behaviour.
6. **Every worked number in this paper is illustrative model output.** No dataset was collected. Section 12.3 prints the numbers in the source material that do not reconcile, rather than correcting them silently.

19.4 4 The composite parameter, the four limits, and the ceiling

Before the limits, the equation reads like a circuit. `E` is the source and has whatever it has: a room with three exits, a manuscript with eleven errors, a patient with a set of signs, a market with a set of prices. `E` does not depend on the observer, and that commitment is what makes this a relation about contact rather than about opinion; were `E` observer-dependent the expression would be circular and would predict nothing. `P` is the aperture — how much of `E` is admitted, which cannot exceed what is there and is usually well under it. `D` is the gain applied to whatever came through. `M` is what leaves the stage and is integrated by the storage law: the thing remembered, the thing that becomes character, the thing an institution writes down.

The four limiting cases are what happens when one element goes to an extreme, and each is a condition clinicians, editors, auditors and inspectors already recognise without an equation. **The claim is that one line predicts all four, and that the four are not four pathologies but one parameter at four values.**

19.4.1 4.1 The composite parameter

The equation has four symbols and three degrees of freedom, because `D` and `P` enter only through their ratio. Define

```
r = D / P          the CALIBRATION RATIO      (dimensionless)
M = r . E
```

Everything `D` and `P` do to the moment they do through `r`. **A tenfold increase in both is a tenfold increase in nothing.** The practical consequence is immediate: an intervention raising frame and intake together changes what the observer can handle, not how distorted the observer is.

This is also the identification limit, and it is exact. `(D, P, E)` and `(kD, kP, E)` are observationally identical for every `k > 0` — a one-parameter unidentified family. `r` is recoverable; the split of `r` into `D` and `P` is not recoverable from `M` alone; `E` requires an independent instrument. Any claim about the absolute magnitude of a person's frame or intake therefore needs a second instrument and must carry a note that it is not identified by the moment measurement (SP-4). The signed distortion is `ln r`, the unsigned distortion `|ln r|`, and both are developed in Section 6.

19.4.2 4.2 Limit 1: `P -> 0`, the frame with no intake

```
P -> 0,  D and E fixed  =>  M -> infinity
```

Computed, with `D = 1` and `E = 1`:

`P` (intake)	`M`
1.0	1.000
0.5	2.000
0.25	4.000
0.1	10.000
0.05	20.000
0.01	100.000
0.001	1 000.000
0.000001	1 000 000.000

This is the only divergence in the formula, and the limiting-case gate requires that a divergence name a real phenomenon or the formula is wrong. It does, under several names. **Dogma**: the doctrine is complete, no observation is admitted that could revise it, every event reads as confirmation, and the moment is enormous and entirely manufactured. **Apophenia**: structure perceived in noise, the intake carrying no signal and the frame supplying all of it. **Hallucination**, in the technical sense of a percept without a corresponding stimulus — the formula explains nothing about hallucination and makes no clinical claim; it says only that a percept with a frame and no intake has the term structure `D/0`. **The over-fitted institution**, whose model is elaborate and whose measurement is degraded, so that everything which happens is already in the plan; the organisational literature reaches the same place from sensemaking [Weick, 1995, 1993] and from the erosion of local variety under a legible scheme [Scott, 1998].

Note the direction, because it is counterintuitive and it is the formula's sharpest prediction. **Reducing intake does not reduce the moment. It inflates it.** The person who has stopped looking does not experience less; they experience *more*, and more confidently, and with less friction. A shrinking `P` feels like clarity. That is why the condition is stable, why it is pleasant and why it does not self-correct: the internal signal that would say *look harder* — the sense that the situation is thin — is exactly the signal that vanishes as `P` falls. **The subjective symptom of `P -> 0` is not confusion. It is certainty.**

There is an honest limit to what the formula can say here. **A divergence to infinity is a modelling artefact**; no real observer's moment is unbounded. Section 4.6 supplies the ceiling, and with it in place the correct statement is that as `P -> 0` the moment saturates at the observer's full repertoire — the observer sees the maximum their categories permit, all of it supplied by themselves. Saturation at maximum internal structure with zero external content is a better description of dogma than divergence to infinity, and it is what the bounded model predicts.

19.4.3 4.3 Limit 2: `D -> 0`, the intake with no frame

```
D -> 0,  P and E fixed  =>  M -> 0
```

`D` (frame)	`M` (with `P = 1`, `E = 1`)
1.0	1.000000
0.5	0.500000
0.1	0.100000
0.01	0.010000
0.001	0.001000
0	0.000000

The moment collapses. Everything is arriving; nothing is being made of it. **This limit carries the framework's second thesis, stated as flatly as it can be: sensation is not experience.** An observer can be saturated with intake — every detail registered, nothing missed — and have no moment at all, because a moment requires something to be made of what arrived, and `D` is the term that makes it.

The named conditions are **flooding** (the novice in their second week: they see everything, they can rank nothing); **overwhelm**, which is not too much input in an absolute sense but too much input relative to the structure available to organise it — a distinction that matters enormously for intervention, because reducing the load treats the first and supplying a frame treats the second, and the second is far more common; **the flattened moment**, described phenomenologically as things being present but not mattering, for which the formula offers a structure — full `P`, near-zero `D`, therefore near-zero `M` — and makes no clinical claim whatever; and **the data-rich, decision-poor organisation**, everything logged and nothing concluded.

The asymmetry between the two limits is worth staring at. **Losing intake makes the world feel bigger; losing frame makes it feel like nothing.** Neither feels like an error from inside. One feels like insight and the other like exhaustion.

19.4.4 4.4 Limit 3: `E -> 0`, the environment with nothing in it

```
E -> 0,  r fixed  =>  M -> 0
```

With `r = 3` fixed:

`E`	`M`
10	30.00
1	3.00
0.1	0.30
0.01	0.03
0	0.00

This limit looks trivial and is not. **No calibration ratio, however good, manufactures a moment out of an environment that has nothing in it.** The best-tuned observer in an empty room gets an empty moment; `r = 1` multiplied by zero is zero. That is the framework's answer to the doctrine that attitude is sufficient: it is not sufficient, it is a multiplier, and a multiplier on zero is zero. The engineering of `E` — the room, the table, the sightlines, the acoustics, the cadence — is not replaced by anything in this paper. **This paper is a multiplier on that work, and the equation encodes exactly that relationship.**

The reverse reading is the more useful one operationally. If `M` is small and `r` has been established near 1, the problem is `E` and no amount of internal work will fix it. **A calibrated observer in a poor environment is not a psychological case; they are a correct instrument reporting a true reading, and the intervention is in the room.** More observers are in this condition than in either distortion condition, and mislabelling it as a perception problem is among the most common and most damaging available errors (SP-10, FM-6).

19.4.5 4.5 Limit 4: the ratio limits

```
r -> infinity  =>  M / E -> infinity     over-framed without bound
r -> 0         =>  M / E -> 0            under-framed without bound
r = 1          =>  M = E                 the fixed point
```

r	M/E	signed distortion ln r	Condition
0.01	0.01	−4.6052	Extreme under-framing
0.10	0.10	−2.3026	Extreme under-framing
0.25	0.25	−1.3863	Severe under-framing
0.50	0.50	−0.6931	Severe under-framing
0.80	0.80	−0.2231	Mild under-framing
0.90	0.90	−0.1054	Mild under-framing
1.00	**1.00**	**0.0000**	**Calibrated — contact**
1.10	1.10	+0.0953	Calibrated (band edge)
1.25	1.25	+0.2231	Mild over-framing
1.50	1.50	+0.4055	Marked over-framing
2.00	2.00	+0.6931	Severe over-framing
4.00	4.00	+1.3863	Severe over-framing
10.00	10.00	+2.3026	Extreme over-framing
100.00	100.00	+4.6052	Extreme over-framing

Read the third column, not the first. **The log-ratio is symmetric — `r = 2` and `r = 0.5` are the same distance from contact, 0.6931 nats, in opposite directions — and the raw ratio is not.** On the raw scale over-framing appears to have unlimited room and under-framing appears squeezed into `(0,1)`. That is an artefact of the coordinate, and Section 6.2 removes it. **Under the demotion of Section 3.5 the entries in the last two columns are illustrative arithmetic on the model and not measurements**; what the table establishes is the ordering and the symmetry, both of which survive.

19.4.6 4.6 The ceiling `0 <= M <= ln N`, and which effective count it uses

Left as written, `M = (D/P).E` is unbounded, and Section 4.2 has already flagged that a divergence to infinity is an artefact. The bound comes from the register. For the moment to serve as `S_t` in the storage law it is normalised into the surplus band; the surplus law's ceiling is a theorem, `0 <= F(u) <= ln N` [Curzi, 2026], and the normalisation is a convention. Together:

```
0 <= M <= ln N        [THEOREM (surplus range) + CONVENTION (normalisation)]
N = the number of categories the observer can distinguish
```

`N` here is the raw count of distinguishable categories. The affine effective count `N_eff = 1 + (N−1)u` is the quantity that appears *inside* the surplus functional and reaches `N` only at `u = 1`, which is where the ceiling is attained; **it is not a Hill number** and no external literature contains it, whereas `1/HHI` and `exp(H)` are the Hill numbers of order two and one and are fully grounded [Hill, 1973, Jost, 2006]. A reviewer reaching for the Hill-number literature as support for `N_eff` will find that it does not apply. **The ceiling of this paper uses the raw `N`; every other effective count in this paper is the affine one; and neither is a measured quantity here.**

The richest possible moment is the logarithm of how many kinds of thing the observer can tell apart. The framework reads this as Ashby's law of requisite variety — only variety absorbs variety [Ashby, 1958, 1956] — in first-person form. **That reading is an analogy and is badged as one in Section 14.7:** Ashby's theorem is a statement about a regulator, a disturbance set and

a goal set, and this paper does not show that an observer's categories satisfy those premises. The bound itself does not depend on the reading; it follows from putting the moment into the surplus register at all.

Computed:

`N` (distinguishable categories)	`ln N` (max moment, nats)	in octaves at `2F/ln 2`
2	0.693147	2.0000
3	1.098612	3.1699
5	1.609438	4.6439
7	1.945910	5.6147
10	2.302585	6.6439
12	2.484907	7.1699
20	2.995732	8.6439
50	3.912023	11.2877
100	4.605170	13.2877
1 000	6.907755	19.9316
10 000	9.210340	26.5754

A note on the third column, because the source is silent on it and a reviewer will not be. The conversion used is `2F/ln 2 = 2.885390 F`, which is **twice** the direct change of logarithmic base `F/ln 2 = 1.442695 F`. The factor of two is not a typographical accident and it is not free: it comes from the exponent 2 in the corpus's essence bridge `f_min = k e^{2c/eta}` [Curzi, 2026], where a unit change in the argument moves the logarithm of the frequency by two — the same elasticity that makes `d ln f_min / d(c/eta) = 2` and turns a 0.1 error in `c/eta` into a 22.14 per cent error in `f_min` [Curzi, 2026]. **The source material of this paper does not say so.** Read as a plain change of base, the column is wrong by a factor of two; read as a conversion into the essence register it is correct and its warrant lies outside this paper. The defined constant in that bridge, `k = (9/8)^2 / phi^2 = 0.483425732988414`, is **exact by definition** and its fifteen figures are legitimate for that reason alone; attaching `Hz` to it is a separate and much stronger claim [Curzi, 2026]. Nothing in the present paper depends on `phi`, and the golden-section literature's own scepticism about attributions of `phi` [Markowsky, 1992] is recorded here because that constant is quoted at that precision.

Finally, the brutal fact of the logarithm, restated at this scope. **Doubling the repertoire adds `ln 2 = 0.693147` nats to the maximum moment, wherever one starts.** `ln 100 − ln 50 = 0.693147` and `ln 4 − ln 2 = 0.693147`, identically. Going from fifty distinguishable kinds of thing to a hundred buys precisely as much ceiling as going from two to four, at forty-eight times the cost. **A distortion problem cannot be solved by learning more categories.** Raising `N` raises a ceiling the observer is almost certainly not touching; at `r = 2.1` the observer is manufacturing half their moment, and added vocabulary gives them more material to manufacture it from. Fix `r` first: it is the multiplier, and `N` is only the roof.

19.5 5 The fixed point at contact and its dynamics

19.5.1 5.1 `D = P`: the fixed point

```
D = P    =>    r = 1    =>    M = E        for every E
```

When the frame the observer brings equals the intake the observer takes, the moment equals the environment. That is veridical perception stated as an equation, and it is the paper's second result. It requires care, because it is easy to read past.

It is an identity, not an approximation. At `D = P` the observer's contribution to the moment is exactly the multiplicative identity — not small, not negligible, exactly one. The observer

is still fully present: `D` is not zero, `P` is not zero, both are working at whatever magnitude the situation demands, and the product of their activity is the environment, unmodified.

It holds for every `E`. That is what makes it a fixed point of a relation rather than a coincidence at one value. There is no environment for which calibration fails to deliver contact, and none in which the calibrated observer must trade fidelity for scale. The condition is scale-free in `E`, and for the same structural reason the storage law is scale-free in its stock: only the dimensionless part of the model carries the observer's contribution. Section 14.6 states that invariance formally and says what would break it.

It is not the absence of a frame. This is the common misreading and it inverts the formula. `D = 0` is not contact; `D = 0` is Limit 2 and gives `M = 0`. **Undistorted contact is not the state of having no expectations. It is the state of bringing exactly as much structure as was taken in.** A tradition prescribing that perception be emptied of concept is describing `D -> 0`, and by this formula it is describing not clarity but the abolition of the moment. Clarity does not require less frame; it requires *matched* frame.

It is not passivity and it is not agreement. An observer at `r = 1` in an environment containing a hostile fact perceives a hostile fact at full magnitude. Contact is not comfort. The fixed point makes no promise about what the moment contains; it promises only that the moment contains what the environment does.

It is not certainty. At `r = 1` the observer asserts exactly what they took in, which in a half-legible environment is a half-legible reading. **Calibration means confidence tracks intake, which in a murky room means being appropriately unsure.** The observer at `r = 1` in fog reports fog; the observer at `r = 4` in fog reports a building.

19.5.2 5.2 Why it is a fixed point and not merely an equality

An equality is a coincidence at a value; a fixed point is the invariant of a map. The moment equation as written is a static relation, so a dynamics is required, and the dynamics is the observer's own correction behaviour. Postulate the simplest correction available: after each contact, compare what was asserted with what came back, and move the frame toward the intake by a fraction of the discrepancy. In log coordinates, with correction gain `a`:

```
ln r_{t+1} = (1 − a) . ln r_t              THE CORRECTION MAP  [POSTULATE]
fixed point:  ln r* = 0,  i.e. r* = 1,  i.e. D = P
```

This is a first-order linear recurrence in log space and nothing more exotic. It is the same object as the storage law with input and cost set to zero: a leaky integrator relaxing to its fixed point, with pole `(1−a)`. The whole standard theory of such recurrences applies without modification — the fixed point, the stability condition, the geometric transient [Elaydi, 2005, Kelley and Peterson, 2001] — and the control-theoretic reading of `a` as a loop gain with a stability margin is standard [Åström and Murray, 2008]. Its statistical twin is exponential smoothing with smoothing constant `a` [Holt, 1957, Brown and Meyer, 1961], and the first-order discrete section it describes is the one every signals text develops [Oppenheim et al., 1996]. **This is a shared mathematical form, not a derivation from physical premises**, and Section 14.7 says so in the terms the rigour standard requires: a leaky integrator is a leaky integrator whether it holds charge or a frame, and no physical result's authority transfers to an observer.

By the contraction theorem [Curzi, 2026, Banach, 1922] the map has a unique fixed point and iteration converges to it geometrically whenever `|1−a| < 1`, that is, whenever

```
0 < a < 2                    the stability condition for calibration
```

Trajectories from an initial distortion `ln r = +0.8122`:

gain `a`	pole `abs(1-a)`	steps to `abs(ln r) < 0.10`	first six values of `ln r`
0.10	0.90	20	+0.7310 +0.6579 +0.5921 +0.5329 +0.4796 +0.4316
0.20	0.80	10	+0.6498 +0.5198 +0.4158 +0.3327 +0.2661 +0.2129
0.35	0.65	5	+0.5279 +0.3432 +0.2231 +0.1450 +0.0942 +0.0613
0.50	0.50	4	+0.4061 +0.2031 +0.1015 +0.0508 +0.0254 +0.0127
0.75	0.25	2	+0.2031 +0.0508 +0.0127 +0.0032 +0.0008 +0.0002
1.00	0.00	1	0 0 0 0 0 0 (deadbeat)
1.30	0.30	2	−0.2437 +0.0731 −0.0219 +0.0066 −0.0020 +0.0006
1.80	0.80	10	−0.6498 +0.5198 −0.4158 +0.3327 −0.2661 +0.2129
2.00	1.00	never	−0.8122 +0.8122 −0.8122 +0.8122 −0.8122 +0.8122
2.30	1.30	never	−1.0559 +1.3726 −1.7844 +2.3197 −3.0157 +3.9204

The source of this paper describes `+0.8122` as "the value measured in the worked example" of the seeded review audit. It is not. That audit returns `ln r = +0.544727` (Section 12.1). The trajectories in the table are arithmetically correct from a start of `+0.8122` — Section 14.4 redoes the divergent row line by line — but the attribution is wrong, and it is printed here rather than corrected silently. `+0.8122` corresponds to `r = 2.2529`, which is not a value this paper computes anywhere.

Three readings follow, all of which Section 8.1 turns into procedure.

1. **Under-correction converges, slowly.** The observer who adjusts a little each time arrives: twenty encounters at `a = 0.10`, five at `a = 0.35`.
2. **Over-correction past `a = 1` oscillates** — the observer overshoots to the far side of contact each cycle — **and still converges while `a < 2`.** Oscillating correction is inefficient, not fatal.
3. **At `a >= 2` calibration diverges.** The observer who reverses their whole position at every contradiction, then reverses again, gets further from contact with each iteration. At `a = 2.30` the trajectory runs 1.06, 1.37, 1.78, 2.32, 3.02, 3.92 — a distortion of 3.92 nats, `r` near 50, a factor of fifty from contact reached in six cycles from a start of 0.81. **This is a mechanism and not a metaphor, and it is the same stability condition an engineer applies to a feedback loop** [Åström and Murray, 2008].

The fixed point is therefore an attractor and not a static ideal, and its basin is the set of correction gains strictly between 0 and 2. With no correction at all, `a = 0`, it is neutrally stable:

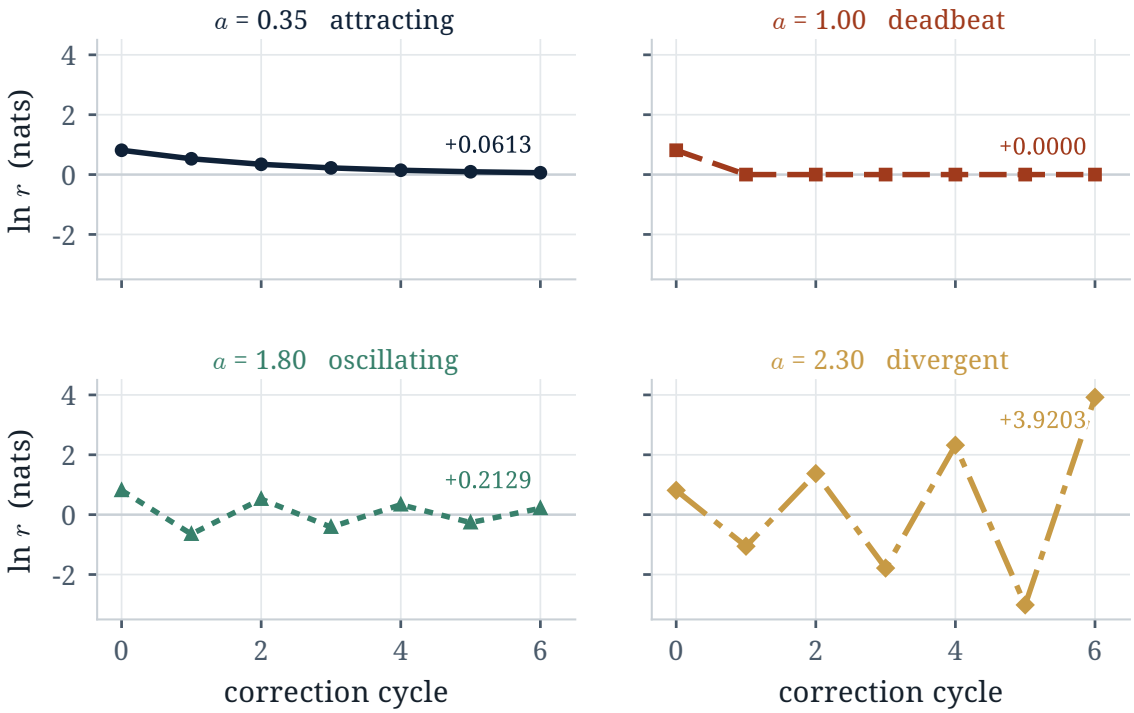


Figure 19.1: **The stability band is a property of the gain, not of the observer.** Trajectories of the signed distortion `ln r` under the correction map `ln r_{t+1} = (1−a) ln r_t` from a start of +0.8122, at the four gains that bracket the band: 0.35 converges monotonically inside five cycles, 1.00 is deadbeat and reaches zero in one, 1.80 oscillates and still converges, and 2.30 alternates with growing amplitude and reaches 3.9203 nats by the sixth cycle — `r` near 50, a factor of fifty from contact. Every plotted value is an entry of the trajectory table in Section 5.2. Illustrative model output from a stipulated start and stipulated gains; generated by `figs/src/CCP-019_fig01.py`.

whatever distortion the observer has, they keep. That is the default. **Contact is maintained, not attained.**

Under the demotion of Section 3.5 the *numbers* in that table are illustrative arithmetic on the model. **What survives the demotion is the trichotomy** — attracting, neutral, divergent — which depends only on the sign and rank of `ln r` and on the value of a design parameter the analyst chooses rather than measures.

19.5.3 5.3 The fixed point across scopes

The framework reads the same row at two scopes:

Term	Internal reading (psychology)	External reading (sociology)
`D/P`	frame-to-intake ratio	**doctrine-to-data ratio**
`M`	the moment as lived	the situation as the group constructs it

An institution at `D = P` asserts exactly what its instruments measured: it does not manufacture and it does not drown.

`r = 1` reads correctly at both scopes without modification, which a merely metaphorical correspondence would not do. That is offered as the framework's evidence that psychology

and sociology are one subject at two scopes, and it is worth exactly what a structural observation is worth — **it is not an identified empirical result**, for the reason Section 3.7 gives, and the group-scope reading in particular is the one the reflection problem bites hardest [Manski, 1993].

19.5.4 5.4 Sensitivity: unit elasticity in all three terms

Differentiating the log form,

```
ln M = ln D − ln P + ln E
dM/M = dD/D − dP/P + dE/E
```

Every elasticity is exactly +1 or −1. Verified numerically at `D = P = E = 1`: `+1.000000` with respect to `D`, `−0.999999` with respect to `P` (finite-difference error), `+1.000000` with respect to `E`.

A one per cent error anywhere produces a one per cent error in the moment. There is no term that may be treated sloppily and no term whose error is amplified. That is a genuine virtue of the multiplicative form: a formula with a squared term or an exponent would have a sensitive parameter and a forgiving one and would require measurement effort to be allocated unevenly. It is also what makes the framework's *other* elasticity results notable by contrast — the essence bridge has elasticity 2 and is correspondingly brittle [Curzi, 2026].

The asymmetry that does exist is between numerator and denominator at large deviations, because a ratio is not symmetric in its arguments:

perturbation	effect on `M`
`D` up 1%	factor 1.0100 (+1.00%)
`D` up 10%	factor 1.1000 (+10.00%)
`D` up 50%	factor 1.5000 (+50.00%)
`P` up 10%	factor 0.9091 (−9.09%)
`P` down 10%	factor 1.1111 (+11.11%)
`P` up 50%	factor 0.6667 (−33.33%)
`P` down 50%	factor 2.0000 (+100.00%)

Halving intake doubles the moment; halving frame halves it. The two errors are not equally available or equally seductive, and the table shows why: the pathological direction is the one with the larger reward.

19.5.5 5.5 The join: the moment becomes character

The moment equation is the input stage and not the end of the pipeline. Substituting the moment for the surplus input in the storage law,

```
THE SELF-STORAGE LAW      Q_{t+1} = (1 − delta) Q_t + eta M_t − C_t
```

That is the architecture: perspective and perception produce the moment; the moment becomes character. All the storage law's properties are inherited [Curzi, 2026a]. The consequence for calibration is unpleasant.

Take an observer with `eta = 0.35`, `delta = 0.08`, `C = 0.20`, in an environment worth `E = 1.0` nat. The steady state `Q_inf = (eta M − C)/delta` gives:

condition	r	M = rE	Q_inf
Under-framed	0.352941	0.3529	−**0.9559**
Calibrated	1.047619	1.0476	+2.0833
Over-framed	1.900000	1.9000	**+5.8125**

The over-framed observer stores nearly three times as much as the calibrated one. The storage law integrates the moment as lived and does not inspect it; there is no term in `Q_{t+1} = (1−delta)Q_t + eta M_t − C_t` that asks whether `M` was earned. That is not a defect in the storage law. It is the storage law telling the truth about a real phenomenon: **dogma accumulates faster than knowledge, and by exactly the same mechanism.** The confident, model-driven, low-intake observer builds a large, stable, internally consistent structure very efficiently.

The under-framed observer by contrast has a **negative** steady state — `eta M < C`, the regime the storage law's steady-state theorem explicitly refuses to hide [Curzi, 2026a]. The flooded novice is not slowly building; they are running down. The model says the flooded state is unsustainable and must be exited, and that is correct.

All three rows are illustrative model output from stipulated parameters. They carry no uncertainty because nothing was measured, and they inherit the demotion of Section 3.5 through `M`, so their magnitudes are not findings.

19.5.6 5.6 Veridical throughput: separating what was stored from what was earned

If the storage law cannot distinguish earned structure from manufactured structure, the analyst must, and the equation supplies the instrument.

> **VERIDICAL THROUGHPUT** *(definition).* The part of the moment the environment actually supplied is

```
V = min(M, E) = E . min(r, 1)
confabulated fraction   = (M − V)/M = 1 − 1/r     for r > 1
missed fraction         = (E − M)/E = 1 − r       for r < 1
```

r	M/E	V/E	confabulated fraction	missed fraction
0.10	0.1000	0.1000	—	0.9000
0.25	0.2500	0.2500	—	0.7500
0.50	0.5000	0.5000	—	0.5000
0.75	0.7500	0.7500	—	0.2500
1.00	**1.0000**	**1.0000**	**—**	**—**
1.33	1.3300	1.0000	0.2481	—
2.00	2.0000	1.0000	0.5000	—
4.00	4.0000	1.0000	0.7500	—
10.00	10.0000	1.0000	0.9000	—

`V` is maximised exactly at the fixed point and only there. Over-framing does not increase what the environment gave; it adds material that came from the observer and is stored as though it came from outside. Under-framing loses what was there. **That is the strongest of the four senses in which `D = P` is the formula's centre: the fixed point is not merely where the moment is undistorted, it is where the earned content of the moment is largest.** This is hypothesis H2, and it is a claim about an audit's shape that a real audit could refute.

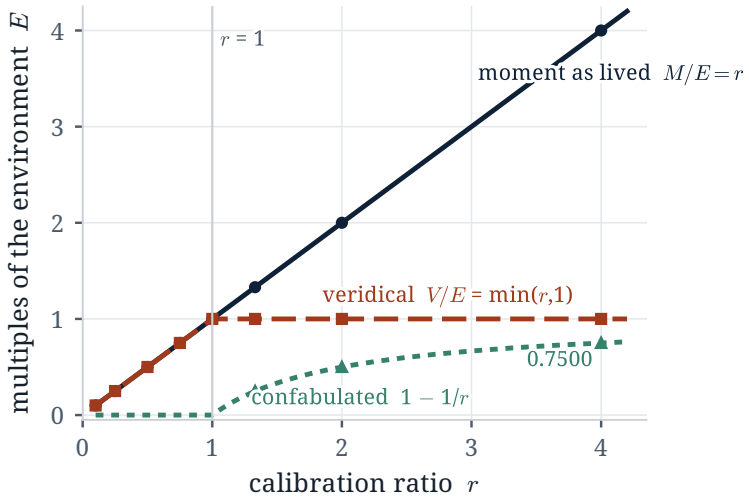


Figure 19.2: **Veridical throughput is single-peaked at contact and the confabulated fraction is its price.** The moment as lived `M/E = r` climbs without limit while the part the environment actually supplied, `V/E = min(r,1)`, is flat above the fixed point; the confabulated fraction `1 − 1/r` accounts for the whole of the difference, reaching 0.2481 at `r = 1.33`, 0.5000 at `r = 2.00` and 0.7500 at `r = 4.00`. Every marked value is a row of the veridical-throughput table in Section 5.6; that table's `r = 10.00` row lies outside the plotted range. Illustrative arithmetic on the model, not a measurement; generated by `figs/src/CCP-019_fig02.py`.

Running both quantities through the storage law separates stored structure from stored error:

condition	`Q_inf` on `M` (as lived)	`Q_inf` on `V` (veridical)	stored error
Under-framed, `r = 0.353`	−0.9559	−0.9559	0.0000
Calibrated, `r = 1.048`	+2.0833	+1.8750	+0.2083
Over-framed, `r = 1.900`	+5.8125	+1.8750	**+3.9375**

The over-framed observer's structure is **68 per cent stored error at steady state** (3.9375 of 5.8125, which is 67.7 per cent, rounded to two figures because the number is illustrative model output and not a measurement). It will be defended as vigorously as the earned third, because from inside there is no difference: **stored error and stored knowledge occupy the same variable.**

The one mechanism that eventually distinguishes them is cost. Structure that does not match the environment generates correction events, contradictions, failures and repairs, and those are `C`. **Over-framing shows up in the books not as bad `Q` but as rising `C`, on a lag.** That lag is why the condition is stable for years and why the diagnostic instruments of Section 7 must be run deliberately rather than waited for. It is hypothesis H6, and Section 8.5 gives the exact condition under which the bill ever catches up.

19.5.7 5.7 What the fixed point costs, and the predictive-processing correspondence

A fixed point that is free is a fixed point nobody needs a procedure for. This one is not free, and the cost has two distinct shapes.

The effort cost. Maintaining `D = P` requires every increase in intake to be matched by an increase in frame and every decrease likewise. Both are effortful — `P` costs attention, `D` costs the construction and maintenance of usable structure — so **calibration is a continuous matching problem under a load budget, and the load appears in the storage law as `C`.** This yields the framework's least comfortable prediction: **distortion is cheap and contact is expensive, in both directions.** Dropping `P` while holding `D` is the cheapest available move: it costs less attention and produces a larger, more confident moment. That is the economics of dogma in one sentence. Dropping `D` while holding `P` is also cheap: nothing then has to be concluded. Both are downhill; only the match is uphill. The engineering response is to **build the circumstance that holds the state** rather than to require a person to hold it by effort, which is what the specifications of Section 10 are for.

The structural cost, which is the counter-result of this paper. At `r = 1` in an environment that saturates the observer's repertoire, the input to the storage law sits at its ceiling, `S = ln N`. Substituting into the steady state and reading the headroom gauge:

```
S     = ln N                        (contact at the fixed point, in a saturating environment)
Q_inf = (eta ln N − C) / delta
h     = eta ln N − delta Q_inf − C
      = eta ln N − (eta ln N − C) − C
      = 0        exactly, by cancellation
```

A perfectly calibrated observer converges to exactly zero headroom. Not approximately, and not through mismanagement: perfect contact drives the input to its maximum, the maximum input drives the store to its maximum, and the maximum store consumes the whole budget in maintenance. The cancellation is exact and both `delta` and `C` drop out of it.

> **A condition the source does not state, and it matters.** The cancellation requires `E = ln N`, not merely `r = 1`. `r = 1` gives `M = E`; it does not give `E = ln N`. The zero-headroom result therefore holds for a calibrated observer **in an environment that saturates their repertoire**, and for no other. In an environment worth `E < ln N` a calibrated observer has headroom `h = eta E − delta Q_inf − C`, which by the same substitution is again zero at the steady state of *that* input — the cancellation is a property of any steady state under constant drive, not of the ceiling. **What the ceiling adds is that a calibrated observer in a rich environment has no room even in principle.** The source presents the two as one step; they are two, and only the second needs the ceiling.
>
> **CORRECTION, and it applies wherever this paper speaks of calibration as a goal: `r = 1` is NOT an unqualified target.** Calibration maximises fidelity and veridical throughput. It does not maximise room to build, and at equilibrium it leaves none. An observer wanting headroom at a settled state must either shed `Q` deliberately or accept `r != 1` as a **declared policy** (Section 8.4). The framework states the trade rather than pretending it away. That headroom is the corpus's name for organisational slack is not incidental: slack is what a system spends when the environment moves [Cyert and March, 1963], and an observer with none is an observer who cannot absorb a surprise.

The predictive-processing correspondence, badged. Predictive-processing accounts in cognitive science model perception as the weighting of top-down predictions against bottom-up sensory evidence, with weights set by the precisions — the inverse variances — assigned to each. In the simplest Gaussian case, with prior mean `mu` at precision `pi_D` and observation `y` at precision `pi_P`, the posterior mean is the precision-weighted average

```
mu_post = (pi_D . mu + pi_P . y) / (pi_D + pi_P)
```

Writing `lambda = pi_D / pi_P` and `w = lambda/(1 + lambda)` for the weight on the prior gives `mu_post = w mu + (1−w) y` and the identity

```
logit(w) = ln( w / (1 − w) ) = ln lambda
```

lambda	w	1 − w	logit(w) = ln lambda
0.01	0.009901	0.990099	−4.605170
0.10	0.090909	0.909091	−2.302585
0.25	0.200000	0.800000	−1.386294
0.50	0.333333	0.666667	−0.693147
1.00	**0.500000**	**0.500000**	**0.000000**
2.00	0.666667	0.333333	+0.693147
4.00	0.800000	0.200000	+1.386294
10.00	0.909091	0.090909	+2.302585
100.00	0.990099	0.009901	+4.605170

The signed distortion `ln r` is the logit of the weight a precision-weighted average assigns to the prior, when `r` is read as the precision ratio. Both formalisms place their invariant at the ratio one and measure departure from it in log units. A worked instance, with a prior expectation of 100 and an observation of 60:

lambda	w	posterior	distance from the evidence
0.1	0.0909	63.636	+3.636
0.5	0.3333	73.333	+13.333
1.0	0.5000	80.000	+20.000
2.0	0.6667	86.667	+26.667
10.0	0.9091	96.364	+36.364

Three caveats, and they are not decoration.

First, **predictive processing is an active and contested research programme, not settled fact**, and — as Section 3.8 records — **this bibliography contains no entry for it at all.** The correspondence is therefore stated here with no external support whatever, and it must not be read as borrowing that programme's authority for anything.

Second, **structural similarity is not identity, and the two objects differ in a way that matters.** The Bayesian weight `w` is bounded in `[0,1]`: a precision-weighted average cannot land outside the interval between prior and evidence. The moment equation's `r` is unbounded above, so `M` can exceed `E` up to the `ln N` ceiling. **The Bayesian model interpolates; the moment equation scales.** The logit identity maps one onto the other cleanly at the level of the log-ratio parameter and does not make them the same model. An observer whose posterior sits between prior and evidence is not manufacturing content; an observer at `r = 4` is.

Third, **the direction of inference runs one way and it is the weak way.** That two independently constructed formalisms place their invariant at the same point is interesting. It is not evidence that either is correct, and this paper does not use it as such. **Section 14.7 classifies it as an analogy in the strict sense** — a suggestive correspondence with no formal warrant — because the premises under which a precision-weighted average is the right model of an observer have not been shown to hold.

19.6 6 The calibration ratio: its measure and its aggregation

The observer's entire contribution to the moment passes through one dimensionless number. This section takes that number on its own, gives distortion a measure, and audits the operations the framework performs on it.

19.6.1 6.1 Definitions

THE CALIBRATION RATIO *(definition).* `r = D / P`, dimensionless, `r > 0`. `r = 1` is contact: the fixed point of Section 5.1, at which `M = E`.

THE DISTORTION MEASURE *(definition).* The signed distortion is `ln r`, in nats, with range `(-inf, +inf)` and zero at contact. The distortion is `Delta = |ln r|`, in nats, with range `[0, +inf)` and zero at contact.

Both are stated as definitions rather than as postulates, because they define quantities rather than assert facts. **What is postulated is that these quantities are worth measuring and that they behave as this paper says; that claim is falsifiable and Section 9.2 says how.** What is *demoted* is their magnitude, and Section 6.3 states what follows.

19.6.2 6.2 Why the measure is a logarithm

Three independent arguments force the logarithm, and a reader who rejects one still has two.

Symmetry. On the raw ratio scale over-framing has infinite room and under-framing is crushed into `(0,1)`. That is a coordinate artefact and it corrupts every summary statistic computed on it. Under the logarithm:

r	ln r	1/r	ln(1/r)
2	+0.693147	0.5000	−0.693147
3	+1.098612	0.3333	−1.098612
4	+1.386294	0.2500	−1.386294
10	+2.302585	0.1000	−2.302585

`ln(1/r) = -ln r` identically. **Seeing twice as much as is there and seeing half as much as is there are the same magnitude of error in opposite directions, and only the log coordinate says so.** On the raw scale `r = 2` sits 1.00 above contact and `r = 0.5` sits 0.50 below it: the same fault priced at half.

Composition. Lenses compose multiplicatively and logarithms turn that into addition:

```
r_total = r_1 . r_2 . ... . r_n
ln r_total = ln r_1 + ln r_2 + ... + ln r_n
```

Distortion budgets are additive in log units and only in log units. Section 8.2 builds a pipeline budget on this and it is not constructible any other way.

Register. Surplus, pitch, sensation and the moment are all logarithmic in this framework, and every conversion between them is linear in log space with exact constants. Reporting distortion in nats puts it in the same register as `F = ln N_eff` and as `S = k ln(I/I_0)` [Curzi, 2026n].

Which of the three the argument actually needs. The symmetry and composition arguments are properties of the logarithm as a function of a multiplicative quantity; they require nothing about perception and nothing about psychophysics. The register argument is the one that leans on Fechner [Fechner, 1860], and it is the weakest of the three, because the psychophysical literature has never settled the logarithmic-against-power dispute [Stevens, 1957, Krueger, 1989] and the fitted exponent depends on the experimenter's range [Teghtsoonian, 1971]. **The stronger warrant for a logarithm, where one is needed, is that a logarithmic evaluation follows from multiplicative dynamics with no psychological assumption at all** [Peters and Gell-Mann, 2016]. This paper's use of the logarithm is the multiplicative one: `r` composes by multiplication along a chain, so its natural additive coordinate is its logarithm. Fechner corroborates and does not carry the load.

19.6.3 6.3 The band table, and what the demotion does to it

Distortion is a continuum and applied work needs bands.

> **The bands below are a CONVENTION of this framework**, chosen for round numbers in nats and for correspondence with recognisable states. **They are not measured constants and must not be reported as if they were.**

Band	`Delta = abs(ln r)` (nats)	over-framed `r`	under-framed `r`	Reading
CALIBRATED	0.00 – 0.10	1.000 – 1.105	0.905 – 1.000	Contact. Within ±10.5 per cent of the environment.
MILD	0.10 – 0.35	1.105 – 1.419	0.705 – 0.905	Detectable, tolerable, worth trending.
MARKED	0.35 – 0.69	1.419 – 1.994	0.502 – 0.705	The moment is meaningfully not the environment. Correct.
SEVERE	0.69 – 1.61	1.994 – 5.003	0.200 – 0.502	Factor of two to five. Decisions from this reading are unsafe.
EXTREME	1.61 and above	5.003 and above	0.200 and below	Factor of five or more. The observer is not in contact.

The band edges are `ln 2 = 0.693147` at the marked/severe boundary and `ln 5 = 1.609438` at the severe/extreme boundary.

The scale-type audit of this table, which the framework owes and which is stated here rather than left to a reviewer. Three claims are being made and they do not all survive.

- **Banding itself is licensed.** Cutting an ordered variable at declared thresholds requires only that the variable be ordinal, and `Delta` is ordinal by Section 3.5. An observer whose distortion falls between two audits has moved toward contact, and that is a finding.
- **The claim that the severe band is "exactly wrong by a factor between two and five" is NOT licensed.** It reads the magnitude of `r`, which the demotion withdrew.
- **The claim that the two directions "occupy the same band width" is NOT licensed either.** Equality of differences on `ln r` requires an interval scale, and `ln r` inherits only an ordinal one. **The symmetry is a property of the logarithm as a function, and it becomes a property of the measurement exactly when `Delta` becomes a Rasch logit difference and not before** [Rasch, 1960, Curzi, 2026t].

The corpus's `f_min` result is the same defeat in another domain and is worth reading beside this one: relabelling an ordinal triple (1.0, 2.0, 3.0) as (0.5, 2.0, 5.0) leaves the ordering unchanged and moves one computed frequency by a factor of 54.6 [Curzi, 2026]. **An ordering that survives relabelling and a magnitude that does not are exactly what an ordinal scale is.**

19.6.4 6.4 Aggregation: the arithmetic mean of `r` is not licensed, and neither is the mean of `ln r`

An observer measured on several occasions produces several values of `r`. The summary statistic matters enormously and the obvious one is wrong.

occasions	arithmetic mean `r`	geometric mean `r`	mean `ln r`	**mean `abs(ln r)`**
{0.5, 2.0}	1.250000	1.000000	0.000000	**0.693147**
{0.4, 0.8, 1.5, 2.5}	1.300000	1.046635	+0.045580	**0.615298**
{1.2, 1.25, 1.15, 1.3}	1.225000	1.223723	+0.201898	**0.201898**

Read row one. An observer who over-framed by a factor of two on one occasion and under-framed by a factor of two on another has an arithmetic mean `r` of 1.25, reporting mild over-framing that never occurred; a geometric mean of exactly 1.000, reporting perfect calibration that also never occurred; and a mean distortion of 0.693 nats, which is severe. **Only the third statistic describes what happened.**

By Jensen's inequality the arithmetic mean of `r` is at least the geometric mean, with equality only when every `r` is identical, so **the arithmetic mean of a calibration ratio always reports over-framing that may not be there.** Row three is the degenerate case where all three agree: a steady, one-directional bias.

The framework's rule, and then the correction the scale type forces on it.

> **AGGREGATION RULE.** Report mean `|ln r|` as the distortion of an observer across occasions. Report mean `ln r` as the bias — the direction they lean. Report the geometric mean of `r` only as a back-transform of the second. **Never report the arithmetic mean of `r`.**

That rule is right about the arithmetic mean and it is **not far enough**. The reason the arithmetic mean is wrong is that `r` is a ratio, and the licensed central tendency of a ratio-scale quantity under multiplicative composition is the geometric mean — which is the arithmetic mean in log coordinates, and which is why the log-scaled psychophysical register recommends it [Fechner, 1860] and why a multiplicative-dynamics argument recommends it independently of any perceptual story [Peters and Gell-Mann, 2016]. **But an arithmetic mean of `ln r` requires `ln r` on an interval scale, and Section 3.5 withdrew that.** Strictly, therefore:

- **Not licensed pending instrumentation:** the arithmetic mean of `r` (twice over — wrong statistic *and* wrong scale); the geometric mean of `r`; the mean of `ln r`; the mean of `|ln r|`; and every percentage change computed from any of them.
- **Licensed now, on ordinal data:** the **median** of `ln r` across occasions; the **sign** of that median, which is the bias; the **rank order** of observers by `|ln r|`; and the **direction of movement** of one observer between two audits on a common item set.
- **Licensed the moment the Rasch instrument of Section 3.5 exists:** all of the above, because `Delta` is then a logit difference on a linked scale [Rasch, 1960, Andrich, 1988].

This paper continues to print means and percentage changes, because they are the arithmetic of the model and a reader must be able to redo them. **Every one of them is badged in the audit of Section 12.4 as not licensed as a measurement**, and the conclusions drawn from them in Section 13 are stated at the ordinal strength the data support. The pair (median `ln r`, rank on `|ln r|`) is the whole honest summary: a direction and a position. An observer with bias `+0.20` and distortion 0.20 leans one way consistently; an observer with bias 0.00 and distortion 0.69 is unstable in both directions and is the harder case, because they pass any audit that averages.

19.7 7 Diagnosis: two states, a differential, and five estimators

19.7.1 7.1 The two states

```
r > 1     OVER-FRAMED      M > E     projection, prejudice, seeing the expected
r = 1     CALIBRATED       M = E     contact
r < 1     UNDER-FRAMED     M < E     flooding, overwhelm, cannot make it mean anything
```

Neither state feels like an error from inside, and the two feel like opposite things. Over-framing feels like clarity, competence, speed and confidence. Under-framing feels like exhaustion, fog and inadequacy. **The pleasant one is the dangerous one**, and that asymmetry organises everything below: over-framing has to be hunted, under-framing announces itself.

Over-framed: the cardinal features.

Feature	What is observed
High on-model detection	Finds what the model predicts, reliably and fast.
Low off-model detection	Misses what the model does not contain, almost completely.
High false-discovery rate	Asserts things that are not there, in the model's shape.
Falling time per case	Speed increases as intake falls, and is rewarded.
Falling surprise rate	Being contradicted becomes rare. Read as mastery.
High confidence	Subjectively: clarity.
Low revision rate	Positions do not change on new information.

Four subtypes differ in what to do about them. **The doctrinal**: the frame is institutional and enforced, deviation carries a career cost, the individual's `r` is a property of the organisation, and individual correction will fail. **The expertise decay**: the frame was earned, was once accurate, and has outlived the environment that validated it — the commonest and most sympathetic form, in which competence correctly acquired becomes distortion by the passage of time in the environment rather than by any error of the observer. This is the mechanism the organisational literature reaches independently, as the myopia of learning [Levinthal and March, 1993] and as the homogenising effect of fast learning on a slow environment [March, 1991]; it is also the mechanism by which an anomaly is progressively normalised into the standing model [Vaughan, 1996]. **The prejudicial**: the frame is a category applied to persons in advance of evidence about them. Structurally identical to the others; the ethical weight is different and this paper does not pretend the arithmetic settles the ethics. What the arithmetic does say is that the mechanism is `P` falling rather than `D` being wicked, and that the correction is intake. **The apophenic**: the intake carries no signal — noise, sparse data, a genuinely ambiguous situation — and the frame supplies all the structure; distinguished by the environment, since a thin `E` makes any `D` produce a large `r`, and **the correction here is often to withhold assertion rather than to change the frame.**

Under-framed: the cardinal features.

Feature	What is observed
High or adequate intake	Notices a great deal. The notes are long.
Low assertion rate	The output is thin relative to the notes.
Rising time per case	Everything takes longer, and the extra time buys little.

Feature	What is observed
High surprise rate	Frequently caught out; the world keeps being unexpected.
Low confidence, high revision	Positions change on weak evidence.
Missed material	Not because it was not seen — because it was not concluded from.
Subjective report	Overwhelm; the material cannot be made to mean anything.

Four subtypes again. **The novice**: the frame has not been built yet; time and structure fix it. This is the benign case and the one everyone assumes they are looking at. **The overloaded**: the frame exists and is adequate but the intake rate exceeds the rate at which it can be applied — distinguished from the novice by performance on a smaller sample, and if the assertions come on a third of the material this is a load problem and the correction is on `P` or on the schedule, not on `D`. **The deframed**: the observer has been instructed to drop their assumptions and come to it fresh, so `D` has been reduced by advice — an **iatrogenic** case in which the intervention caused the condition (FM-5, FM-17). **The fragmented**: the frame exists but is not retrievable under the conditions of use — distributed across notes nobody consults, held by a person not in the room, or dependent on a system that is down. `D` is high in principle and zero in practice, and the correction is availability, not training.

19.7.2 7.2 The differential diagnosis

Observation	Over-framed	Under-framed	Poor environment	Overloaded
`r` from seeded audit	> 1.105	< 0.905	about 1	< 0.905
Intake `P`	Low	High	Normal	High
Assertion `D`	High	Low	Normal	Low
False discovery rate	High	Low	Low	Low
Off-model hit rate	Very low	Moderate	Normal	Low
Time per case	Falling	Rising	Normal	Rising
Surprise rate	Low	High	Normal	High
Performance on a **reduced sample**	Unchanged	Unchanged	Unchanged	**Improves markedly**
Independent measurement of `E`	Rich	Rich	**Thin**	Rich
Primary correction	Raise `P`, cap `D`	Raise `D`	Rebuild `E`	Cut load

Two rows are the ones nobody runs and both are cheap. **The reduced-sample test separates the overloaded from the under-framed in a single session**, and the two need opposite treatments: the overloaded observer given training gets worse, because training adds to the load. **The independent measurement of `E` separates every internal diagnosis from the case where the instrument is fine and the room is empty** (SP-10), and it is the first thing to check because it is the only diagnosis whose correction lies outside the person.

Note the row that carries the threshold: the differential quotes `r` cut points of 1.105 and 0.905, which are the band edges of Section 6.3. **Those are convention, not measurement**, and using them as decision thresholds imports the convention into the diagnosis. The licensed reading of that row is *higher than a declared cut* and not *greater than 1.105 as a magnitude*.

19.7.3 7.3 Why an observer cannot detect their own `r`

The states are not introspectively distinguishable from their causes, and the reason is structural rather than psychological.

At `r > 1` the internal report is: the situation is clear, I know what this is, this is going quickly. Every one of those reports is *also* what contact feels like. There is no experience attached to the assertions that were not there, because the observer does not experience them as manufactured — **`M` is the moment as lived, and a manufactured moment is lived exactly as a found one is.** The only internal signal that would discriminate is surprise, and surprise is precisely the quantity that falls as `P` falls. At `r < 1` the internal report is: I am not good enough at this — indistinguishable from the report a calibrated observer would make in a genuinely hard environment, and indistinguishable from the report of an overloaded one.

Therefore calibration cannot be self-assessed, and this framework does not accept self-assessment as evidence for it (SP-8). That is not a comment on honesty; it is a consequence of the fact that the quantity being estimated is the thing doing the estimating. The supporting result outside the framework is that people cannot reliably report the processes producing their own judgements [Nisbett and Wilson, 1977], and self-report instruments are additionally exposed to wording and framing effects [Schwarz, 1999, Kahneman et al., 2006]. The instruments below are therefore all external — and Section 3.7 records that being external relocates the problem rather than removing it.

19.7.4 7.4 Estimating `r` with ground truth: the seeded audit

```
E      : a set of conditions planted or adjudicated in advance, independent of the observer
P_hat  : (real conditions appearing in the observer's intake record) / (conditions present)
D_hat  : (total items in the observer's assertion record)           / (conditions present)
r      = D_hat / P_hat
```

Requirements: at least 20 per cent off-model items (SP-9), an intake record distinct from the assertion record (SP-18), and ground truth from a source that is not the observer (SP-3). **Cost: one packet, one session, two counts.** This is not an expensive instrument. It is a rarely used one.

19.7.5 7.5 Estimating `r` without ground truth: five behavioural estimators

Ground truth is not always available. Five estimators that do not require it follow. **All five are proposed by this framework, none is validated, and each is badged.**

> **E1 — THE ASSERTION-TO-NOTICE RATIO** *(estimator, unvalidated).* `r_hat = (items asserted) / (items noticed)`. Because `D_hat` and `P_hat` divide by the same `E`, the ratio is the two raw counts and **the environment cancels**. The cheapest estimator available, requiring only that intake and assertion be separately recorded. *Biased by:* whether a note and an assertion are really the same unit — SP-1 must be argued explicitly for the particular records used. *Blind to:* the case where both counts are low because the environment is thin.
>
> **E2 — THE SURPRISE RATE** *(estimator, unvalidated).* Occasions on which the observer's expectation was contradicted, over occasions. Falls with `r` in the over-framed direction and rises in the under-framed one. **Its value is as a trend and not**

as a level: a falling surprise rate in a stable environment is the earliest available indicator of drift toward over-framing, and it is available from ordinary records. *Biased by:* changes in the environment, which move it for reasons that have nothing to do with the observer.

E3 — THE REVISION RATE *(estimator, unvalidated).* Positions changed after new evidence, over occasions of new evidence. Near zero at high `r`; high and unstable at low `r`. **Combined with E2 it separates the directions**: low surprise with low revision is over-framing; high surprise with high revision is under-framing; high surprise with low revision is an observer being contradicted and not updating, which is over-framing under pressure and is the worst of the four combinations.

E4 — TIME TO COMMIT *(estimator, unvalidated).* Time from first contact to first firm assertion. Short and falling at high `r`; long and rising at low `r`. *Biased by:* incentives on speed, which are usually present and usually push `r` up. **In any organisation that measures throughput, expect a systematic push toward over-framing.**

E5 — THE OFF-MODEL RATIO *(estimator, unvalidated).* Assertions outside the standing model, over assertions in total. Requires only a written model to compare against. A ratio near zero over many occasions is evidence of over-framing **or** of an environment that genuinely contains nothing off-model, and the two must be distinguished by a seeded item or by rotation.

Use at least two, from different rows (SP-16). E1 and E5 together are the cheapest useful pair: one gives magnitude — subject to the demotion — and the other gives the over-framing direction, and both come out of records that already exist. **None of the five has been validated against a criterion, which is exactly the convergent and discriminant demonstration a construct of this kind owes** [Campbell and Fiske, 1959, Cronbach and Meehl, 1955], and none has been run.

19.8 8 Correction, composition, drift, and the operating point

19.8.1 8.1 The correction map as procedure, and choosing the gain

The correction map of Section 5.2 is a first-order recursion in log space and is used here as procedure:

```
ln r_{t+1} = (1 − a) . ln r_t              [POSTULATE −− the correction map]
a  = correction gain: the fraction of the observed discrepancy acted on per cycle
fixed point ln r* = 0
stability: converges iff |1 − a| < 1, i.e. 0 < a < 2
```

Distortion half-lives, computed as `ln 2 / (−ln|1−a|)`:

gain `a`	pole `abs(1−a)`	half-life (cycles)	cycles to one tenth
0.05	0.95	13.513	44.891
0.10	0.90	6.579	21.854
0.20	0.80	3.106	10.319
0.35	0.65	1.609	5.345
0.50	0.50	1.000	3.322
0.65	0.35	0.660	2.193
0.80	0.20	0.431	1.431

That is the same half-life formula the storage law gives for the decay of stored order under `delta`, with `|1−a|` in the role of `(1−delta)` [Curzi, 2026a]. It is the same circuit. Calibration is a leaky integrator whose leak is the correction procedure, and the correction procedure is the only leak it has: **with no correction, distortion does not decay. There is no natural forgetting of a bad frame.**

That is the practical heart of the material, and it inverts the storage law's problem. Stored order decays without maintenance and must be defended; distortion **persists** without maintenance, because the storage law integrates it as readily as anything else. **An observer does not drift back into calibration. They are corrected into it or they stay where they are.** The corpus reaches the same structure in a very different domain, where the exposure quantity turns out to be not the event but the time value sits uncorrected between events — a five-second hop being `1.585e−07` of a year against `0.0822` for value rested thirty days, a ratio of about 518,400 [Curzi, 2026]. **The dwell of an uncorrected frame is the exposure, and the audit interval is the dwell.**

Choosing the gain. The gain is a design parameter and the table is its specification, in the same sense that a filter gain is chosen rather than measured [Kalman, 1960, Åström and Murray, 2008].

- `a` **below about 0.10** is under-correction. It converges — twenty-two cycles to reduce distortion tenfold — but slowly enough that the environment may change first.
- `a` **between 0.20 and 0.50** is the working band: half-lives of three cycles down to one.
- `a` **near 1.00** is deadbeat correction, one step to zero in theory. In practice it means discarding the entire existing frame in one move, and the cost appears immediately in the storage law's admissibility condition `C_t <= (1−delta)Q_t + eta M_t`. **Deadbeat correction of a large `Q` violates admissibility and produces the collapse it was meant to prevent** (SP-23).
- `a` **between 1 and 2** oscillates and converges. Wasteful but safe.
- `a` **at or above 2** diverges (FM-8, FM-22). From `Delta = 0.8122`, six cycles at `a = 2.3` reach `Delta = 3.9203` — a factor of fifty from contact, arrived at by correcting hard every single time.

The design rule is one line: never move more than the evidence, and never move twice on the same evidence. Both halves are needed. The first bounds `a` below 2 on a single step; the second stops the same discrepancy being counted repeatedly, which is the mechanism by which a well-intentioned corrector reaches an effective gain above 2 without ever deciding to.

Correcting over-framing (`r > 1`) has two levers, and the specification requires both to be considered, because the same `Delta` can be reached from either end.

Lever A — lower `D`; make assertion cost something.

Procedure	Mechanism	Effect on `r`
Citation requirement	Every assertion must name the intake-record entry that generated it. An assertion with no intake behind it cannot be filed.	Direct cut to the numerator; in the audit of Section 12.1 this removed 20 of 21 unfounded items.
Assertion cap	A hard limit on items per case, forcing prioritisation.	Cuts `D`; risks cutting true findings too. Use with a measured audit, not alone.
Explicit alternative	Require one written alternative reading of the same evidence before committing.	Reduces the frame's monopoly on the intake.

Procedure	Mechanism	Effect on `r`
Falsifier statement	Require, per assertion, the observation that would show it wrong.	Converts unfalsifiable frame-output into a testable claim [Popper, 1959, Platt, 1964]; several will not survive the writing.

Lever B — raise `P`; buy intake back.

Procedure	Mechanism	Effect on `r`
The off-model pass	A separately timed pass whose only question is what is here that the model does not cover, with its own note field.	Raises the denominator and attacks the low off-model hit rate directly. Highest-value single procedure.
Rotation	Move the observer across contexts so the standing model stops matching.	Raises `P` by making the frame visibly insufficient; also raises non-parallelism across the body [Curzi, 2026].
Fresh eyes	A second observer of a different lineage, reporting independently before comparison.	Reveals the off-model material the first observer's frame filtered. Requires independent record before discussion, or the second observer inherits the first one's frame.
Slowing the case	Restore time removed by throughput incentives.	Raises `P`; unpopular; often the only thing that works when E4 shows an incentive-driven `r`.
Seeded items in live work	Plant known off-model conditions in real caseloads at low rate.	Maintains intake by making it consequential. Requires consent and governance (SP-28).

Lever B is usually the better one and is almost never the one chosen. Lowering `D` is visible, fast and satisfying; raising `P` costs time, which is the resource the organisation is already short of. But cutting `D` alone moves `r` toward 1 while leaving the observer's contact with the environment exactly where it was, and the veridical throughput `V = E min(r,1)` does not improve at all until `P` improves. **Cutting `D` fixes the ratio; raising `P` fixes the reading.**

Correcting under-framing (`r < 1`) also has two levers, and the second is usually dismissed.

Lever C — raise `D`; supply structure. Worked exemplars, which build the pattern library faster than instruction; a written taxonomy, which is literally raising `N` and raises the ceiling as well as the frame; a checklist **with a decision rule**, because a checklist that says only what to look at raises `P` and not `D` and will make under-framing worse (FM-18); apprenticeship, which is expensive, slow, and the most reliable; and forced commitment — a written conclusion before leaving the case, even a provisional one, which converts intake into assertion by fiat and generates the error signal correction needs.

Lever D — lower `P`. For the overloaded observer, whose frame is adequate and whose intake exceeds their capacity to apply it, the correction is to reduce the intake: smaller batches, fewer simultaneous cases, a filter upstream, a longer interval. **Turning down the input is a correction, not a retreat**, and the reduced-sample test is what licenses it (SP-20). The engineering

statement is that `r` is a ratio and a ratio can be corrected from either term; an organisation that only ever responds to under-performance with training is only ever pulling one lever, and it is pulling the expensive one.

19.8.2 8.2 Composition, the compensating-distortion trap, and the distortion budget

Distortion composes along a chain of observers: `r_total = r_1 . r_2 . ... . r_n`, so `ln r_total = sum_i ln r_i`. The consequence that matters:

chain	net `r`	net `Delta`	total `Delta` injected
2.5 then 0.4	1.0000	0.000000	**1.832581**
0.5 then 2.0	1.0000	0.000000	**1.386294**
1.6 then 0.625	1.0000	0.000000	**0.940007**

A severely over-framed observer feeding a severely under-framed one produces a perfectly calibrated output and a completely broken system. The net distortion is zero; the total distortion present is 1.83 nats; the end-to-end measurement certifies contact. This is not a curiosity. It is the standard structure of a hierarchy in which an enthusiastic analyst reports to a sceptical manager, and it is *stable*, because the end-to-end numbers look good and each party's correction of the other maintains the balance. It fails the moment either party is replaced, and it fails without warning, because the failure is a change in a quantity nobody was measuring.

> **DECOMPOSITION RULE.** Distortion shall be measured per stage and not end to end. Report both `|ln r_total|` (net) and `sum_i |ln r_i|` (injected). **A system in which the injected distortion materially exceeds the net distortion is a compensating system and shall be recorded as fragile** (SP-24).

The distortion budget, along a four-stage reporting pipeline:

Stage	`r`	`ln r`	cumulative `r`	cumulative `ln r`	injected `Delta` so far
Field intake	0.85	−0.162519	0.850000	−0.162519	0.162519
Write-up	1.15	+0.139762	0.977500	−0.022757	0.302281
Supervisor review	1.35	+0.300105	1.319625	+0.277348	0.602385
Client summary	1.60	+0.470004	2.111400	+0.747351	1.072389

End-to-end `r = 2.111`, net distortion 0.747 nats, total injected 1.072 nats. Two readings. **No single stage is worse than mild-to-marked and the end product is severely over-framed**, because distortion accumulates additively in log units and each stage's contribution is small enough to survive its own review. And the pipeline briefly passes through near-calibration at stage two, cumulative `ln r = −0.0228`, so **a measurement taken at the write-up would certify the pipeline as calibrated. Where the measurement is taken decides what is found, and a single measurement point in a multi-stage pipeline is not a measurement of the pipeline** (FM-15).

A budget is set as any engineering error budget is set: allocate a total, divide it across stages, measure per stage, and require the sum to stay under the total. A defensible starting allocation for four stages with a target of a mildly distorted reading at the client is a total budget of 0.35 nats, which is 0.0875 nats per stage — **every stage inside about ±9 per cent.**

That is demanding, and stating it as a number is the point: most pipelines have never been told what their per-stage tolerance is and consequently do not have one.

The scale-type audit of this subsection, which the framework owes. Adding `ln r` across stages requires `ln r` on an interval scale, and Section 3.5 withdrew that. **Chain addition is therefore a structural argument and not an arithmetic one until the Rasch instrument exists.** What survives the demotion is the *structure*: that compensating stages exist, that a net figure conceals them, and that per-stage measurement is required to see them. Those follow from the sign pattern alone. The numbers 1.832581 and 1.072389 are illustrative arithmetic on the model.

19.8.3 8.3 Drift: why calibration is a maintenance item

The correction map as written has no source term, so left alone `ln r` sits where it is. But `E` moves — the building stock changes, the codebase changes, the population changes — and a frame that matched the old environment does not match the new one. **In a moving environment, holding a frame constant is drift**, and drift is a forcing function on the correction loop:

```
ln r_{t+1} = (1 - a) . ln r_t + w_t          [POSTULATE -- the driven correction map]
w_t = distortion injected per cycle by environmental change, in nats
```

This is the driven first-order system, a leaky integrator with a source, and it has the standard steady state [Elaydi, 2005, Oppenheim et al., 1996]. Setting `ln r_{t+1} = ln r_t = ln r_inf` with constant `w`:

```
ln r_inf = w / a                    THE RESIDUAL DISTORTION
```

Residual distortion is the drift rate divided by the correction gain. It is not zero and it cannot be made zero. A maintained observer in a moving environment does not sit at contact; they sit at a fixed offset from it, and the size of the offset is set by two numbers, only one of which they control.

drift `w` (nats/cycle)	gain `a` = 0.05	0.10	0.20	0.35	0.50
0.02	0.400 MARKED	0.200 MILD	**0.100 CAL.**	0.057 CAL.	0.040 CAL.
0.05	1.000 SEVERE	0.500 MARKED	0.250 MILD	0.143 MILD	**0.100 MILD**
0.10	2.000 SEVERE	1.000 SEVERE	0.500 MARKED	0.286 MILD	0.200 MILD
0.20	4.000 SEVERE	2.000 SEVERE	1.000 SEVERE	0.571 MARKED	0.400 MARKED

Holding the residual inside the calibrated band requires `a >= w / 0.10`:

drift `w` per cycle	minimum gain required
0.02	0.20 — achievable
0.05	0.50 — achievable, at the top of the working band
0.10	1.00 — deadbeat, forbidden by SP-23 at any material `Q`
0.20	2.00 — at the stability limit; **not achievable at all**

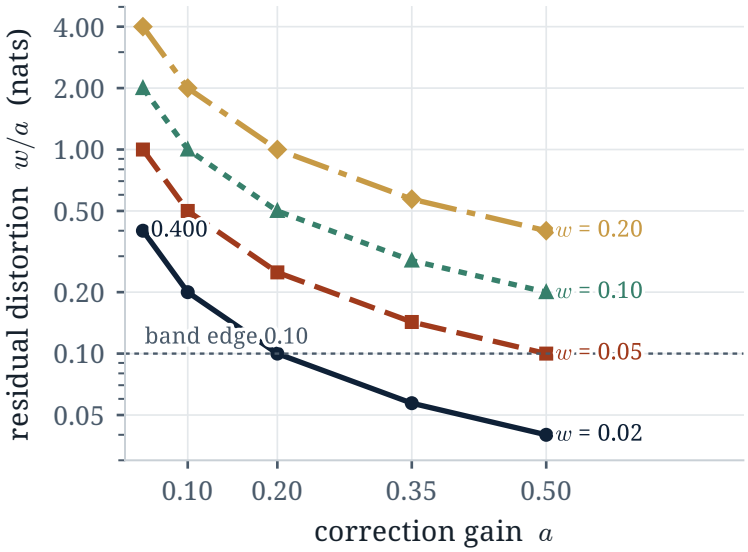


Figure 19.3: **No admissible gain holds calibration against a fast environment.** Residual distortion `ln r_inf = w/a` at the four drift rates of Section 8.3, against correction gain. Only `w = 0.02` reaches the 0.10-nat calibrated band inside the working gain range; `w = 0.10` requires deadbeat correction, which the admissibility condition forbids at any material stored order, and `w = 0.20` requires a gain of 2.00, which is the stability limit itself. Every plotted value is an entry of the residual-distortion table in Section 8.3. Illustrative arithmetic on the model; the drift rates are stipulated, not measured; generated by `figs/src/CCP-019_fig03.py`.

If the environment injects more than about a tenth of a nat of distortion per correction cycle, no achievable gain holds calibration, and the only remaining lever is to shorten the cycle. That is why the audit cadence is a specification item rather than a preference. **A drift problem cannot be corrected out of by correcting harder**: past `a` of about 0.5 the increments buy little, and past `a = 1` they are forbidden by the admissibility condition. It is corrected out of by correcting **more often**.

The drift rate `w` is a property of a domain and must be estimated locally — by re-auditing a *frozen* frame against a fresh environment and reading the distortion that accumulated between the two dates. What follows from the table is the cadence rule.

> **CADENCE RULE.** The audit interval shall be short enough that the distortion injected between audits is below the tolerance band width. With `w_year` the annual drift and `B` the band width in nats, the interval in years satisfies `T <= B / w_year` **and** the gain satisfies `a >= w_year . T / B`. Both, not either (SP-27).

If the band is 0.10 nats and the domain moves 0.10 nats a year, the requirement is an annual cycle at gain 0.5 or better, or a quarterly cycle at gain 0.13. **Most organisations in fast-moving domains run frames last validated at a cadence set by convenience, and the residual distortion that implies has never been computed.** It takes four numbers and one line of arithmetic: `w`, `a`, `B` and the interval.

Scale-type audit. `w/a` is a ratio of two quantities neither of which is instrumented, and it is not licensed as a magnitude. **What survives is the two monotonicities** — more drift gives more residual, more gain gives less — **and the qualitative conclusion that a sufficiently fast environment cannot be held by any admissible gain**, which follows from the stability bound `a < 2` and the band width alone, both of which are declared rather than measured.

19.8.4 8.4 When `r` should not be 1: the operating point as policy

Calibration is not always the objective, because the moment is not always the product. The moment equation says `r = 1` maximises fidelity and veridical throughput. It does not say `r = 1` minimises loss, and where the costs of the two error types differ, the loss-minimising operating point is not the fidelity-maximising one.

Take the three operating points of the inspector series of Section 12.2 and price them under three cost structures:

Operating point	`r`	misses	false alarms	cost at 1:1	cost at 10:1 (miss:alarm)	cost at 1:3 (miss:alarm)
Novice	0.353	41	3	44	413	50
Trained	1.048	10	4	**14**	**104**	**22**
Veteran	1.900	20	27	47	227	101

The calibrated point wins under all three, which is the comfortable case. It does not always. **Where a miss is catastrophic and a false alarm is cheap, a deliberate `r > 1` is a rational policy**: screening, triage, safety interlocks and early-warning systems are all built to over-assert on purpose, and their false-discovery rates are a design choice rather than a fault. Conversely, **where assertion is expensive and irreversible — a public accusation, a permanent record, a decision that cannot be unwound — a deliberate `r < 1` is rational**: withhold conclusion, keep taking intake, decline to make it mean something yet. **This bibliography contains no signal-detection-theory entry**, so that statement is carried here by the paper's own cost arithmetic and by nothing borrowed; Section 15 records it as an uncited claim.

> **Calibration is a property of the instrument. The operating point is a policy.** An observer may be perfectly calibrated and still choose to assert above or below their intake because the loss function says so. **What is never legitimate is an operating point that was not chosen — a policy-level `r` that is an uncorrected instrument error, defended after the fact as prudence or as decisiveness** (SP-26, FM-20).

The test is whether the observer can state the loss function. A screener who says the office runs hot because a miss costs a hundred times a false alarm has a policy. A reviewer who says they have a good eye for these has an uncorrected `r`, and the two are indistinguishable from outside until asked.

Two structural reasons follow, and both correct any reading that treats `r = 1` as a simple goal. The first is the zero-headroom result of Section 5.7, with the condition stated there. The second is sharper and is developed next.

19.8.5 8.5 The two-track ledger and the knife edge at `kappa = delta`

Section 5.5 established that over-framing *increases* stored order, because the storage law integrates the moment as lived and never asks whether it was earned. Turn that into a two-track ledger: raw `Q` against veridical `Q^V`, with stored error

```
Z_{t+1} = (1 − delta) Z_t + eta (M_t − V_t)
correction load   C_t = c_0 + kappa Z_t
```

Because `V = E min(r,1)` is flat above the fixed point while `M = rE` keeps climbing, the veridical track stops improving at `r = 1` and the raw track does not. Computed at `eta = 0.35`, `delta = 0.08`, `c_0 = 0.20`, `kappa = 0.02`, `E = 1.0`:

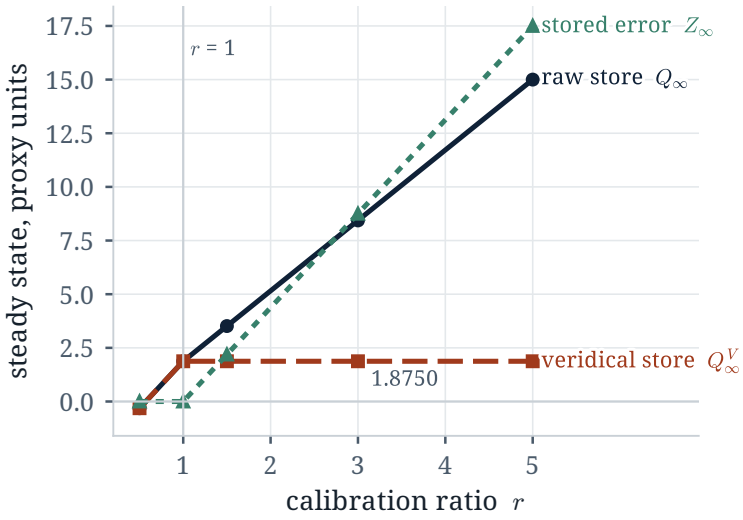


Figure 19.4: **The ledger points away from contact.** Raw stored order at steady state rises without bound in the calibration ratio — from 1.8750 at `r = 1` to 15.0000 at `r = 5` — while the veridical track is flat at 1.8750 above the fixed point and the stored-error stock climbs from zero to 17.5000. At the stipulated parameters the correction-load coefficient `kappa = 0.02` is below the decay rate `delta = 0.08`, which is the condition under which the correction bill never catches up. Every plotted value is an entry of the two-track table in Section 8.5. Illustrative construction with stipulated parameters, not a measurement; generated by `figs/src/CCP-019_fig04.py`.

r	M	V	Z_inf	Q_inf (raw track)	Q_inf^V (veridical track)
0.50	0.5000	0.5000	0.0000	−0.3125	−0.3125
1.00	1.0000	1.0000	0.0000	**1.8750**	**1.8750**
1.50	1.5000	1.0000	2.1875	3.5156	1.8750
3.00	3.0000	1.0000	8.7500	8.4375	1.8750
5.00	5.0000	1.0000	17.5000	15.0000	1.8750

Illustrative construction: the veridical column is flat above the fixed point by construction, and every magnitude in the table inherits the demotion of Section 3.5. The structure is the finding, not the numbers.

Dogma banks faster than knowledge, and the bill arrives on a lag as correction load. Differentiating for `r > 1` gives the exact condition under which the bill ever catches up:

```
dQ_inf/dM = eta (1 − kappa/delta) / delta        ->    negative iff  kappa > delta
```

The knife edge is exactly `kappa = delta`. Where correction load accrues more slowly than stored order decays, raw `Q` rises without bound in `r`, and every incentive attached to raw `Q` is an incentive to over-frame. That is not a moral claim about the observer; it is the arithmetic of what they are being paid in. **Measure `Q^V`, or measure `Z` alongside `Q`, or the ledger being optimised points away from contact.** At the stated parameters `kappa = 0.02 < delta = 0.08`, so the worked case is on the wrong side of the knife edge and the raw track rises in `r` throughout.

19.8.6 8.6 Why the two directions are policed unequally

> **This subsection is the author's reading and is badged as such: it is an observation about which quantities are easy to count, not a derived result and not a measured one.**

Institutions detect and punish under-framing readily and detect over-framing almost never, and the reason is visible in the signature panels. **The under-framed observer is slow, uncertain and produces thin output, and every one of those is directly measurable by instruments organisations already run.** The over-framed observer is fast, confident and produces abundant output, and detecting them requires an instrument — off-model seeded items — that almost nobody runs.

The arithmetic of Section 12.2 supports the asymmetry rather than the intuition: the over-framed veteran produces a worse total error count than the untrained novice, in a fifth of the time, while being the office's most productive inspector by the office's own metric. **The state that gets caught is the one that costs least; the state that costs most is the one that gets promoted.** That is a claim about which quantities happen to be easy to count, and it is correctable by counting a different one — which is the whole reason SP-17 exists and the requirement to defend when someone proposes to simplify the audit. It is also the mechanism by which a measure that becomes a target ceases to be a measure [Goodhart, 1975, Campbell, 1979].

19.9 9 The validation protocol applied: seven gates, one demotion, two unrun

The framework's formula validation protocol runs seven gates on any proposed relation [Curzi, 2026t]. They are labelled G1 to G7 here. **One gate is demoted and two are unrun, and this paper does not pretend otherwise.**

19.9.1 9.1 The gate table

Gates on the moment equation `M = (D/P).E`.

Gate	Test applied	Finding	Verdict
G1 DIMENSIONS	Units of each term; both sides agree; ratios dimensionless.	Requires `[D]=[P]`; then `r` is dimensionless and `[M]=[E]`. Any `D`, `P` pair not measured in identical units is inadmissible.	**PASS**, conditional on SP-1
G2 SCALE TYPE	Division requires a ratio scale with a true zero.	No instrument puts `D` and `P` on a common ratio scale, so the magnitude of `D/P`, the number `Delta`, and the quantitative join to the store are **not reportable as measured quantities**. The limiting cases, the fixed point and the three-way ordering survive on ordinal information. Route out: Rasch logit difference on a linked scale, at about 123 items per term for a half-logit. Section 3.5.	**DEMOTE** — structure stands, magnitude does not

Gate	Test applied	Finding	Verdict
G3 LIMITING CASES	Each term to 0 and to infinity; every divergence must name a real phenomenon.	Four cases, Sections 4.2 to 4.5. `P–>0` diverges and names dogma, apophenia, the over-fitted institution; `D–>0` and `E–>0` give 0. No unexplained divergence.	**PASS**
G4 FIXED POINTS	Solve for the invariant; state its meaning.	`D = P` gives `r = 1` and `M = E` for every `E`. Under the correction map it is unique and attracting for `0 < a < 2`.	**PASS**
G5 IDENTIFIABILITY	Recoverable from obtainable data? Do distinct parameter sets predict identically?	`r` is recoverable; `D`, `P`, `E` are not separately recoverable from `M`. `(D,P,E)` and `(kD,kP,E)` are observationally identical. The reflexive case is not identified at all [Manski, 1993].	**OPEN**
G6 FALSIFICATION	Preregister, freeze, held out, four baselines including a shuffled-input control.	Protocol in Section 9.2. Killing result: a frozen `r` that fails to beat the shuffled control has no content.	**NOT RUN**
G7 REPRODUCIBILITY	Different people, different sites: do the parameters recover?	Procedure-portable in principle. Whether `r` is a trait-like parameter or a property of the item set is **unknown and decisive**.	**NOT RUN**

Gates on the calibration ratio `r` and the distortion `Delta`, where they differ. G1 inherits and fails with SP-1 if that is violated. **G2 returns DEMOTE for both**: the count-based estimator supplies a zero and a unit *within* one item set but not the cross-set linkage that would make the magnitude of `r` a measurement rather than a tally, and rating-scale estimators supply neither. G3 passes: `r` cannot be negative, because `D` and `P` are counts or precisions and both are non-negative; `r –> 0` gives total under-framing and `r –> infinity` saturates at `ln N`. G4 passes, with `Delta` acting as a Lyapunov function for the correction map — it decreases monotonically under any gain in `(0,1)`, which is what proves the correction converges. **G5 is the strong gate for `r`**: under the assertion-to-notice estimator `r = (items asserted)/(items noticed)`, two counts, and ground truth is needed only to distinguish a low `r` from a poor environment. G5 for `Delta` carries a warning: **`Delta` discards the sign and is therefore not sufficient on its own to choose a correction**, since the two directions have opposite treatments (SP-19). G6 for `Delta` adds one specific test and it is the highest-value single test available: **does `Delta` predict downstream correction cost?** The framework claims high `Delta` raises `C` on a lag; a `Delta` that does not predict rework, dispute or repair cost is measuring nothing that matters. G7 inherits and is unrun.

19.9.2 9.2 G6 in full: what would kill this formula

Preregistration. Fix the operationalisation of `D`, `P` and `E`, the population, the item set, the number of occasions, and the analysis, in writing, before data collection [Nosek et al., 2018]. The reason is not ceremony: without a rule declared in advance, the number of defensible analysis paths is large enough that a result of the required shape is nearly always obtainable [Simmons et al., 2011, Gelman and Loken, 2013], and the base rate of published findings that do not replicate is what it is [Ioannidis, 2005, Open Science Collaboration, 2015].

Fit and freeze. Estimate `r` for each observer on a training block. Freeze every parameter.

Held-out prediction. Predict, on a held-out block, a quantity the formula constrains and that is not the quantity used to fit: the false-discovery rate, the off-model detection rate, and the ratio of asserted to noticed items.

Four baselines, all of which must be beaten. (i) *Persistence* — the next block equals the last. (ii) *Direct input* — prediction from environment richness alone, ignoring the observer. (iii) *Matched autoregressive* — an autoregressive model on the observer's own past output **with the same number of free parameters**, so that the parameter count is paid for. (iv) *Shuffled-input control* — the same model fitted to observer/environment pairings randomly permuted. **The fourth is the one most often omitted and it is the decisive one.** If the formula predicts as well on shuffled pairs as on true pairs, `r` carries no information about the observer and the formula has no content.

The killing result, stated in advance: if a frozen `r` fails to beat the shuffled-input control on held-out prediction, the moment equation is dead as a quantitative model and must be demoted to the heuristic ordering of Section 3.5. **A framework that cannot fail cannot be trusted, and this is exactly how this one fails.**

What this protocol is worth is known, because the corpus has already run it once on a different law and reported the failure. The confirmatory application of the same protocol to the storage law landed its fit strictly inside the admissible box and agreed with independent measurement within 4.8 per cent on decay and 16.8 per cent on conversion, and then **beat a matched autoregression by 2.8 per cent on held-out error against a preregistered requirement of 10 per cent**, so the verdict was *parameters corroborated, predictive advantage not demonstrated* [Curzi, 2026i]. That is the standard this paper's G6 would be held to, and it is a standard the corpus has already failed once in public.

G7 in full. The reproducibility gate requires the same operationalisation run by different investigators on different populations, asking whether `r` recovers as a property of the observer or turns out to be a property of the item set. **If `r` for the same observer swings by more than the calibrated band width, 0.10 nats, across equivalent item sets, `r` is not a parameter of the person and the framework's psychological reading collapses even if its item-specific reading survives.** The separability property of the Rasch model is the standard this gate should be held to [Rasch, 1960, Andrich, 1988], and the ceiling on individual-level prediction from rich longitudinal data is a reason to expect the gate to be hard rather than easy [Hofman et al., 2017, Salganik et al., 2020].

19.10 10 Specifications

The specifications below are numbered SP-1 to SP-28 and are stated in the imperative form a procedure requires. SP-1 to SP-13 govern the moment equation; SP-14 to SP-28 govern the calibration ratio, and two of them restate earlier requirements because those are the ones most often dropped.

SP-1 — Commensurability. `D` and `P` shall be measured in identical units, by procedures differing only in which of the two they measure. A ratio of two quantities that are not the same kind of quantity shall not be written as a fraction.

SP-2 — Ratio scale or demotion. Before any numerical use of `M = (D/P).E`, the analyst shall state the scale type of each term. Where any term is ordinal, the formula shall be printed in its demoted heuristic form and no number shall be computed from it.

SP-3 — `E` measured independently. `E` shall be measured by an instrument that does not depend on the observer whose moment is being modelled — a seeded item list, an independent panel, a ground-truth record. An `E` inferred from the observer's own report makes the equation circular and its outputs void.

SP-4 — Report `r`, not `D` and `P`. Because only the ratio is identified, reports shall give `r` and `ln r`. Separate values of `D` and `P` shall be reported only where a second independent instrument fixes the scale, and shall carry a note that they are not identified by the moment measurement alone.

SP-5 — Report veridical throughput alongside the moment. Any report of `M` shall be accompanied by `V = E min(r,1)` and by the confabulated or missed fraction. **Reporting a large moment without its veridical throughput is the single most misleading use available for this instrument.**

SP-6 — Distortion in log units. Distortion shall be reported as `ln r` (signed) and `|ln r|` (magnitude), never as a raw ratio or a percentage.

SP-7 — Correction gain within the stability band. Any deliberate correction procedure shall target a gain `0 < a < 2`, and shall be inspected for `a >= 2` behaviour — alternating overshoot of increasing amplitude — at least once per quarter for any observer or body under correction. A correction regime producing alternating extreme positions shall be stopped, not tuned.

SP-8 — No calibration claim without a seeded audit. A claim that an observer, panel or institution is calibrated shall rest on a measurement against known items. Self-assessed calibration shall be recorded as absent.

SP-9 — Off-model items mandatory. Every seeded audit shall include items outside the frame under test, at a stated proportion of at least 20 per cent of seeded items. **An audit containing only on-model items cannot detect over-framing and shall not be used to certify calibration.**

SP-10 — Environment checked before the observer. Where `M` is low, `E` shall be measured before any intervention on the observer. A calibrated observer reporting a poor environment is a correct instrument and shall not be treated as a perception problem.

SP-11 — The ceiling declared. Any use of the normalised moment shall state the assumed `N` and therefore the ceiling `ln N`, and shall state that the normalisation is a convention while the bound `0 <= F <= ln N` is a theorem [Curzi, 2026].

SP-12 — Reflexive environments excluded. Where the observer's assertion materially changes `E`, the analyst shall record the situation as out of scope and shall not report an `r` for it. **This is an exclusion and not a remedy; Section 3.7 says why.**

SP-13 — Decline path. No calibration procedure shall be administered to a person who has not agreed to it, and a refusal shall be recorded as a refusal and not as a result.

SP-14 — Report in log units, signed and unsigned. A raw ratio or a percentage shall not be used as a distortion measure, because it is not symmetric between the two directions and does not compose additively.

SP-15 — Never the arithmetic mean of `r`. Aggregation across occasions shall use mean `|ln r|` for distortion and mean `ln r` for bias. **Section 6.4 tightens this: pending instrumentation, the licensed summaries are the median of `ln r` and the rank order on `|ln r|`.**

SP-16 — Two estimators minimum. Any `r` used for a decision shall rest on at least two estimators drawn from different rows of Section 7.5, or on one seeded audit. A single behavioural estimator shall be recorded as an indication and not as a measurement.

SP-17 — Off-model content at 20 per cent minimum. Restated from SP-9 because it is the requirement most often dropped: every calibration audit shall contain at least 20 per cent off-model items. An audit without them cannot detect over-framing.

SP-18 — Separate intake and assertion records. Any observer whose calibration is to be measured shall keep an intake record distinct from the assertion record. **Where the two are the same document, `r` is 1 by construction and cannot be measured.** This is a records-design requirement and must be in place before the audit.

SP-19 — Direction before correction. The sign of `ln r` shall be established before any correction is applied. `Delta` alone is not sufficient to choose a procedure, and the two directions have opposite treatments.

SP-20 — The reduced-sample test before any training prescription. Where `r` is below the under-framed cut, the reduced-sample test shall be run before training is prescribed. **An overloaded observer given training gets worse**, because training adds to the load that is already the binding constraint.

SP-21 — The environment measured first. Restated from SP-10: where performance is poor, `E` shall be measured independently before any correction is applied to the observer.

SP-22 — Gain inside the stability band. Correction regimes shall target `0 < a < 2` and shall be reviewed for divergent behaviour — alternating sign with growing magnitude — at least quarterly. A regime showing that signature shall be stopped rather than tuned.

SP-23 — No deadbeat correction of a large `Q`. Where stored order is large, correction gain shall not be set at or near 1. Verify against the admissibility condition `C_t <= (1−delta)Q_t + eta M_t` before setting the gain.

SP-24 — Per-stage measurement in pipelines. Distortion in a multi-stage process shall be measured per stage. Reports shall give both net `|ln r_total|` and injected `sum_i |ln r_i|`. A system whose injected distortion materially exceeds its net distortion shall be recorded as fragile and shall not be certified on its end-to-end figure alone.

SP-25 — Declare the budget. Any process with a calibration requirement shall state its total distortion budget in nats and its per-stage allocation, before measurement.

SP-26 — Operating point declared as policy. Any deliberate departure from `r = 1` shall be recorded in writing with the loss function that justifies it. An undeclared departure shall be treated as instrument error and not as policy.

SP-27 — Audit cadence tied to environmental drift. The audit interval shall be shorter than the time in which the environment is expected to move by the width of the tolerance band. The estimate of that rate shall be written down and revised.

SP-28 — Consent and decline. Calibration measurement is measurement of a person. Participation shall be voluntary, refusal shall be recorded as refusal rather than as a result, and results shall not be used for a purpose not disclosed before the audit.

19.11 11 Failure modes

Twenty-two named failure modes follow, FM-1 to FM-10 on the moment equation and FM-11 to FM-22 on the calibration ratio. Each states its description, what it damages, its diagnostic signature, and its fix. **FM-22 is FM-8 at organisational scope and is listed separately because its detection route differs, not because the mechanism does.**

FM-1 — The incommensurable fraction. `D` and `P` measured in different units — years of experience over items noticed, confidence rating over minutes spent — and divided anyway. *Damages:* every downstream number; the output has units nobody can name. *Signature:* ask what `r = 1` would mean physically; if the answer is not a sentence about the same quantity appearing twice, G1 has failed. *Fix:* SP-1; rebuild both measurements from one procedure.

FM-2 — The Likert fraction. `D` and `P` from five-point self-report scales, divided. *Damages:* G2; the number has the form of a ratio and none of the content. *Signature:* the instrument has no true zero — ask what a score of 0 would be and whether a 4 is genuinely twice a 2 [Likert, 1932]. *Fix:* SP-2; demote in print, or build a ratio scale by magnitude estimation or by counting.

FM-3 — Certainty read as calibration. High confidence and low confusion taken as evidence of good contact. *Damages:* detection of `P -> 0`; the condition is misdiagnosed as competence. *Signature:* **falling surprise rate combined with a rising false-discovery rate** — the observer is being contradicted less and being wrong more; also time per case falling while off-model detection falls. *Fix:* SP-8 and SP-9. There is no soft version: confidence is the symptom of the disease, so it cannot be the test.

FM-4 — The empty aperture mistaken for discipline. An observer with a strong frame reads their own low intake as focus and the organisation rewards it as decisiveness. *Damages:* `P`, therefore `r`, therefore `V`. *Signature:* off-model hit rate near zero with on-model hit rate near one — in Section 12.2, on-model detection at its career maximum against off-model detection below a novice's. *Fix:* rebuild intake specifically for off-model material — mandated fresh-eyes passes, rotation across contexts, and an explicit *what is here that is not on the list* step

with its own time allocation.

FM-5 — Frame removal prescribed as clarity. An under-performing observer is advised to drop their assumptions, come to it fresh, stop overthinking; `D` is driven toward zero. *Damages:* `D`, therefore `M`; drives the observer into Limit 2. *Signature:* intake high, output thin, time per case rising, the observer describing the work as overwhelming rather than as wrong. *Fix:* the opposite of the advice given — supply structure: checklists with decision rules, worked exemplars, a taxonomy, an apprenticeship. **The instruction to have no expectations is a prescription for `M = 0`.**

FM-6 — The environment blamed on the observer. `M` is low; the observer is sent to training. *Damages:* wastes the intervention budget, leaves `E` unimproved, damages the observer's standing, and usually raises `C`. *Signature:* a seeded audit returns a calibrated `r`; the instrument is fine. *Fix:* SP-10; measure `E`.

FM-7 — Stored error defended as experience. A long-tenured observer or institution has accumulated large `Q` at high `r`; the structure is coherent, confident, and substantially manufactured, and challenges to it are read as challenges to competence. *Damages:* `Q` as stored error, and `C` on a long lag. *Signature:* compute `Q_inf` on `M` and on `V` separately; a gap of more than about a third of `Q_inf` is the signature. Cross-check the trend in `C`: **rising rework, rising dispute, rising correction cost with no change in the environment is the lagging indicator of a high-`r` history.** *Fix:* shedding stored order to regain headroom is legitimate and underused; combine with a correction regime at gain in `(0,1)`. **Do not attempt deadbeat correction on a large `Q`**: the cost is paid in one period and usually exceeds what the admissibility condition allows.

FM-8 — The divergent corrector. An observer or body responds to each contradiction by reversing position entirely, then reverses again. Effective gain at or above 2. *Damages:* `r`, with increasing amplitude — from `|ln r| = 0.8122`, six iterations at `a = 2.3` reach 3.9203, a factor of fifty from contact. *Signature:* alternating sign of `ln r` with growing magnitude across successive audits; rhetorically, *we were completely wrong about that* appearing at a regular cadence, each time about the previous correction. *Fix:* SP-7; reduce the gain. **Never move more than the evidence, and never move twice on the same evidence.**

FM-9 — The unmeasured environment. `E` is estimated from the observer's own account and `r` is computed from it. *Damages:* circularity — `r` comes out near 1 by construction and the audit certifies calibration that was assumed. *Signature:* ask where the ground truth came from; if the answer is the observer, the number is void. *Fix:* SP-3; seed the items.

FM-10 — The reflexive case forced into the formula. `r` is computed for a negotiator, a forecaster whose forecast moves the market, or a clinician whose diagnosis becomes the treatment. `E` is not exogenous. *Damages:* everything; the formula's central assumption is violated. *Signature:* ask whether the environment would have been the same had the observer concluded otherwise; if not, out of scope. *Fix:* SP-12; record as out of scope and do not report an `r`. **This is the exclusion Section 3.7 identifies as a confession rather than a remedy: the excluded class is exactly where the corpus's own reflexive law operates** [Curzi, 2026, Manski, 1993], and it is the class in which a reflexive fixed point has actually been documented in the field [MacKenzie and Millo, 2003].

FM-11 — Confidence audited as calibration. Certainty, fluency and speed taken as evidence of contact. *Damages:* detection of over-framing entirely. *Signature:* falling surprise rate with rising false-discovery rate, and no seeded audit on file. *Fix:* SP-8, SP-17.

FM-12 — The model-shaped audit. The calibration audit is built from the same checklist the observer works from; every seeded item is on-model. *Damages:* the audit itself, which will certify an over-framed observer as excellent with high apparent sensitivity. *Signature:* compare the seeded item list against the standing model; if the overlap is total, the audit is a mirror. *Fix:* SP-17; off-model items at 20 per cent minimum, generated by someone who does not work from that model.

FM-13 — The arithmetic mean of ratios. An observer's `r` values across occasions are averaged arithmetically and reported as their calibration. *Damages:* the measurement; by Jensen's inequality the figure is biased toward over-framing, and an observer who swings

both ways is reported as a mild over-framer with no indication of instability. *Signature:* any report containing a mean `r`. *Fix:* SP-15, as tightened in Section 6.4.

FM-14 — Compensating distortion. Two stages with opposite distortions netting to `r` near 1; end-to-end measurement certifies the system. *Damages:* robustness; the system fails without warning on any personnel change. *Signature:* per-stage `r` values spread widely around a net near 1, with injected distortion much larger than net. *Fix:* SP-24. Correct each stage independently, and **expect the end-to-end figure to get worse before it gets better** — correcting the over-framed stage first reveals the under-framed stage that was hiding behind it, so a system in compensation for years will look like it is deteriorating during the repair. Say this in advance or the repair will be reversed.

FM-15 — The single measurement point. A multi-stage pipeline measured at one stage and the result reported as the pipeline's calibration. *Damages:* everything downstream of the measurement point. *Signature:* ask where the measurement was taken; in the budget of Section 8.2 a measurement at the write-up returns `ln r = −0.0228` for a pipeline delivering `r = 2.11`. *Fix:* SP-24.

FM-16 — Training prescribed for overload. An observer with `r < 1` is sent to training; their frame was already adequate and the intake rate was the constraint, so training adds load. *Damages:* `r` falls further, `C` rises, and the observer is now also being told the problem is their competence. *Signature:* the reduced-sample test — if performance improves markedly on a third of the material, this is load. *Fix:* SP-20, then Lever D.

FM-17 — The deframing prescription. An under-framed observer is advised to have fewer expectations. `D` is reduced further. *Damages:* `D`, therefore `M`. **Iatrogenic.** *Signature:* advice on file targeting the numerator when `r` is already below 1. *Fix:* Lever C — supply structure: exemplars, a taxonomy, a checklist **with a decision rule**, a forced written conclusion per case.

FM-18 — The checklist that raises `P`. An under-framed observer is given a checklist of things to look at containing no rule for what any observation implies. *Damages:* raises intake without raising assertion; `r` falls; the observer notices more and concludes no more. *Signature:* the checklist's items are all nouns and no item has a *therefore*. *Fix:* add the decision rule to each item. **A list of what to look at is an intake instrument; a list of what things mean is a frame instrument; only the second treats under-framing.**

FM-19 — Throughput incentive read as skill. Time per case is measured and rewarded, `r` rises across the whole body over years, and the organisation reads it as growing expertise. *Damages:* `P` organisation-wide, therefore `r`, therefore `V`. *Signature:* time-to-commit falling across the whole population simultaneously, off-model detection falling in step, and no change in the environment. **A population-wide movement in `r` is an organisational signal and not a personnel one.** *Fix:* individual correction will not hold against a standing incentive; change what is measured, or accept the operating point and declare it under SP-26. This is a measure that has become a target [Goodhart, 1975, Campbell, 1979].

FM-20 — Uncorrected `r` defended as judgement. A persistent departure from `r = 1` is described after the fact as professional judgement, prudence or decisiveness; no loss function was ever stated. *Damages:* correction is blocked and the state is now protected by an account of itself. *Signature:* ask for the loss function. A policy has one, in numbers, from before; an uncorrected instrument produces one on request, in words, afterwards. *Fix:* SP-26. **A stated policy is auditable; a retrospective justification is not, and the difference between them is the whole of the governance available here.**

FM-21 — Calibration assumed to persist. An observer certified as calibrated years ago is treated as calibrated now; the environment has moved. *Damages:* `r`, silently and progressively. *Signature:* date of last audit relative to the estimated environmental drift rate. *Fix:* SP-27. **There is no decay term on distortion: nothing returns an observer to contact except correction, so calibration is a maintenance item with a schedule, like any other.**

FM-22 — The divergent corrector at organisational scope. Effective gain at or above 2 in a body rather than a person: position reverses fully at each contradiction, then reverses again. *Damages:* `r`, with growing amplitude. *Signature:* alternating sign of `ln r` across successive audits with increasing magnitude. *Fix:* SP-22; reduce the gain, and check specifically for the

same discrepancy being acted on more than once, which is the usual mechanism by which a modest per-event gain compounds past 2.

19.12 12 Results

Both audits below are constructed examples with declared parameters. No dataset was collected, no observer was measured, and no organisation was studied. They demonstrate that the procedure runs and that it reports its own failures; they are not evidence that any reviewer or inspector behaves this way and must never be cited as such. The corpus has refused a fabricated data point elsewhere and badged a headline percentage as an artefact of a chosen synthetic ratio rather than a constant [Curzi, 2026], and the same discipline bounds what may be claimed here: where the corpus reports a 56.62 per cent widening from a constructed input set, it badges the figure as a **ceiling on constructed inputs and not a promise**, distinguishes visibility from pledgeability, and presents the result as an architecture rather than a system [Curzi, 2026]. **The two audits below are constructed inputs in exactly that sense.**

19.12.1 12.1 A seeded-defect review audit

A software group runs code review. One reviewer is respected, fast, and finds a great deal. The group wants to know whether that reviewer's output is contact or manufacture.

Step 1 — Build E. A real, already-shipped module is seeded with 40 known defects by two people who will not participate in the review, adjudicated in advance so that E is fixed before any reviewer sees it. Of the 40, **30 are on-model** — null checks, off-by-one, resource leaks, the categories the team's review guide lists — and **10 are off-model**: a licence incompatibility, a timezone assumption, a comment contradicting the code, an error message leaking a path, and six others outside the guide. Off-model items are 25 per cent of the total, satisfying SP-9. **Without them the audit cannot detect over-framing and would have certified the reviewer as excellent.**

Step 2 — Define D and P commensurably (SP-1). Both are *recorded items per condition present*:

```
P_hat = (real conditions appearing in the working notes) / (conditions present)
D_hat = (total items in the submitted review)            / (conditions present)
```

The working notes are the intake record — everything written down while reading, including what was later dropped. The submitted review is the assertion record. **The two records are the same kind of object, items the reviewer wrote, which is what makes the fraction dimensionally legal.**

The caveat G1 forces into print: this treats a note and a review comment as the same unit of recorded item. **That is an assumption.** It is defensible because both are discrete written assertions about the module, and it is checkable, since two independent counters should agree on the count to within a few per cent — but it is an assumption and it is the weakest joint in the measurement. It is also, precisely, the point at which the scale-type demotion bites: counting supplies a zero and a unit inside this item set and supplies no link to any other.

Step 3 — Run the audit, baseline.

```
conditions present                    40   (30 on-model, 10 off-model)
real conditions touched in the notes  29
items in the submitted review         50   (29 real, 21 not real)

P_hat = 29 / 40 = 0.725000
D_hat = 50 / 40 = 1.250000
r     = 1.250000 / 0.725000 = 1.724138
ln r  = +0.544727 nats
```

The reviewer is over-framed by a factor of about 1.7; the moment lived while reviewing is about 1.7 times the module. By the band table that is marked over-framing, one band below severe. **Under the demotion, the reportable findings here are the sign and the band, not the figure 1.724138.**

There is an identity worth noting, because it makes the estimator far cheaper than it looks. Since `r = (calls/present)/(noticed/present) = calls/noticed`, and since in this audit every real condition that reached the notes also reached the review,

```
r = 50/29 = 1.724138
confabulated fraction = 1 − 1/r = 1 − 29/50 = 0.4200 = false discovery rate exactly
```

When everything noticed is asserted, the confabulated fraction of the moment is exactly the false-discovery rate. `r` is then computable from two counts and needs no ground-truth adjudication at all — though ground truth is still required to measure `P` against `E` and therefore to distinguish this case from an under-framed one. The equality `1 − 1/1.724138 = 0.4200` is arithmetic and is redone in Section 14.4.

Step 4 — The secondary indicators.

```
on-model hit rate     28/30 = 0.9333
off-model hit rate     1/10 = 0.1000
false discovery rate  21/50 = 0.4200
intake (P_hat)              0.7250
```

The signature of FM-4 is present in the open: 93 per cent detection of what the guide predicts, 10 per cent detection of what it does not. **The reviewer is not reading the module. The reviewer is checking the module against the guide, and the guide is also supplying 21 findings that are not there.** These four figures reconcile exactly: `28 + 1 = 29` real items in a 50-item review, so `50 − 29 = 21` are false.

Step 5 — Veridical throughput and stored error. With `E` normalised to 1.0 nat and the group's stipulated personal parameters `eta = 0.35`, `delta = 0.08`, `C = 0.20`:

```
M = r . E = 1.7241 nats
V = E . min(r,1) = 1.0000 nats
confabulated fraction = 0.4200

Q_inf on M (as lived)   = (0.35 . 1.7241 − 0.20)/0.08 = +5.043
Q_inf on V (veridical)  = (0.35 . 1.0000 − 0.20)/0.08 = +1.875
stored error            = +3.168
```

The reviewer's accumulated review expertise, at steady state, is about 63 per cent stored error — 3.168 of 5.043, which is 62.8 per cent, rounded to two figures because it is illustrative model output and not a measurement. **This is not a statement about a person. It is a statement about what a guide-driven review process does to anyone who runs it long enough without an off-model channel**, and this reviewer is the group's best precisely because they run the process hardest.

Step 6 — Design the correction. The distortion is `ln r = +0.5447` and the specification requires both levers to be tried, because `r` can be moved from either end. *Lower `D`:* cap review comments; require that each submitted comment cite the line in the working notes that generated it, so that an assertion with no intake behind it cannot be filed. This attacks the 21 unfounded items directly. *Raise `P`:* add a mandated fifteen-minute off-guide pass with its own note field, asking only what is in this module that the guide does not cover; rotate a second reviewer from a different subsystem weekly, which also raises non-parallelism across reviewers [Curzi, 2026]. Target gain `a = 0.35`, well inside the stability band, giving five review cycles to the calibrated band.

Step 7 — Re-audit after six cycles. Fresh packet, same design, 40 seeded conditions, 30 on-model and 10 off-model.

```
real conditions touched in the notes   31
items in the submitted review          32   (31 real, 1 not real)

P_hat = 31/40 = 0.775000
D_hat = 32/40 = 0.800000
r     = 1.032258
ln r  = +0.031749 nats        CALIBRATED
```

Metric	Baseline	After six cycles	Change
`r`	1.7241	1.0323	−40.1%
`abs(ln r)` (nats)	0.5447	0.0317	**17.2-fold reduction**
Intake `P_hat`	0.7250	0.7750	+6.9%
On-model hit rate	0.9333	0.9000	−3.6%
Off-model hit rate	0.1000	0.4000	**4.0-fold**
False discovery rate	0.4200	0.0312	−92.6%
True defects found	29/40 = 0.725	31/40 = 0.775	+6.9%
Confabulated fraction	0.4200	0.0312	−92.6%
`Q_inf` stored error	+3.168	+0.141	−95.5%

The rows reconcile: after correction, on-model 0.900 gives 27 of 30 and off-model 0.400 gives 4 of 10, so 31 real items in a 32-item review and 1 false, and the false-discovery rate is `1/32 = 0.03125`.

Read the on-model row honestly: it went down. The reviewer now finds 90 per cent of guide-listed defects instead of 93. If the group measures reviewer quality by guide-listed defect detection — and most groups do — **this intervention registers as a small regression by the metric being watched**, while total true defects found rose, false alarms fell by 93 per cent, and detection of the defects nobody was looking for quadrupled.

That is Goodhart's law arriving on schedule [Goodhart, 1975, Campbell, 1979]. **The measure that made the reviewer look excellent was the measure that was making the reviewer over-framed, because it counted only the items the frame already contained.** The audit that found it works solely because 25 per cent of its items were outside the frame. SP-9 exists for this reason, and it is the requirement to fight for when someone proposes to simplify the audit — which is hypothesis H4.

Step 8 — What would falsify this. The group preregisters for the next quarter: the frozen `r = 1.03` predicts a held-out false-discovery rate below 0.10 and an off-model hit rate above 0.30, and does so better than (i) persistence from the baseline audit, (ii) a model using packet difficulty alone, (iii) an autoregressive model on past comment counts with the same parameter count, and (iv) the same model fitted to reviewer/packet pairings randomly shuffled across the team. **If the shuffled control predicts as well, `r` is a property of the packet and not of the reviewer, the instrument is measuring nothing about the reviewer, and this estimator fails G7 for this application.** The group commits in writing to publish that result internally if it occurs. **That commitment, and not the arithmetic, is what would make the audit an instrument rather than an opinion** [Nosek et al., 2018].

19.12.2 12.2 A building inspector, measured three times across fourteen years

A municipal inspection office wants to know what happens to an inspector's calibration over a career. It runs one audit design at three career points, with inspectors who volunteered under SP-28.

The audit. A vacant municipal building is prepared with **50 conditions**, adjudicated in advance by two structural engineers and a code officer who take no part in the inspections. **35**

are on-model: conditions named in the office's standing inspection schedule — egress width, handrail height, panel labelling. **15 are off-model**: a roof drain discharging onto a fire-escape landing, a mislabelled shut-off, an added partition that changed the occupancy classification, a masonry repair executed in a mortar harder than the surrounding brick, and eleven others. Off-model content is 30 per cent, above the SP-17 minimum. **The office's first instinct was to build the audit from the inspection schedule, which would have made every item on-model and guaranteed the result the office expected.** The 30 per cent is the whole reason this audit finds anything. Each inspector carries a field notebook (intake) and a filed report (assertion), so SP-18 is satisfied by existing practice.

Point 1 — Week 2. A new inspector, nine days out of orientation.

```
conditions present                        50   (35 on-model, 15 off-model)
real conditions appearing in the notebook 34
items in the filed report                 12   (9 real, 3 not real)
time on site                              232 minutes

P_hat = 34/50 = 0.680000
D_hat = 12/50 = 0.240000
r     = 0.352941        ln r = -1.041454      Delta = 1.0415 nats     SEVERE UNDER-FRAMING
```

The notebook contains 34 of the 50 real conditions — this inspector's intake is not the problem, and it is better than the veteran's will be. The report contains twelve items. Real conditions were written down in the notebook and never became findings. Nearly four hours on site. The recorded subjective report was that a lot of things were seen that the inspector was not sure counted, which is Limit 2 in plain language: intake without frame, the moment collapsing to about a third of the environment. Veridical throughput `V/E = min(r,1) = 0.3529`, so **about 65 per cent of the building was missed, and most of it was missed after being seen.**

Point 2 — Month 9. The same inspector, after training and supervised fieldwork.

```
real conditions appearing in the notebook 42
items in the filed report                 44   (40 real, 4 not real)
time on site                              118 minutes

P_hat = 42/50 = 0.840000
D_hat = 44/50 = 0.880000
r     = 1.047619        ln r = +0.046520      Delta = 0.0465 nats     CALIBRATED
```

Both terms rose. **This is the composite-parameter effect of Section 4.1: `D` and `P` went up together, which changed what the inspector could handle without changing how distorted they were — except that the frame rose slightly faster, which moved `r` from 0.353 to 1.048.** Off-model detection 60 per cent, false discovery 9 per cent, true conditions found 40 of 50, half the time on site. **The correction raised `D`, which was the correct lever because the diagnosis was under-framing with adequate intake**; had the diagnosis been overload, the same intervention would have made things worse (FM-16).

Point 3 — Year 14. The same inspector, now the office's most productive.

```
real conditions appearing in the notebook 30
items in the filed report                 57   (30 real, 27 not real)
time on site                              41 minutes

P_hat = 30/50 = 0.600000
D_hat = 57/50 = 1.140000
r     = 1.900000        ln r = +0.641854      Delta = 0.6419 nats     MARKED, edge of SEVERE
```

The intake has fallen below where it was in week two. Thirty of fifty conditions reach the notebook, against thirty-four in the first week. The report has fifty-seven items, of which twenty-seven are not there.

Metric	Week 2	Month 9	Year 14
`r`	0.353	1.048	1.900
`ln r`	−1.0415	+0.0465	+0.6419
Band	Severe under	**Calibrated**	Marked over
Intake `P_hat`	0.680	0.840	**0.600**
On-model hit rate	0.400	0.886	**0.943**
Off-model hit rate	0.333	0.600	**0.133**
False discovery rate	0.250	0.091	**0.474**
Minutes on site	232	118	**41**
True conditions found	19/50	40/50	35/50
Veridical throughput `V/E`	0.353	1.000	1.000
Confabulated fraction	—	—	**0.474**

Two rows carry the diagnosis. On-model detection is at its career maximum, about 94 per cent — and off-model detection has fallen to about 13 per cent, below the novice's 33 per cent. The inspector has become excellent at finding what the schedule contains and worse than a nine-day novice at finding what it does not. That is FM-4, and it is the expertise-decay subtype: a frame that was correct, earned honestly, and outlived the building stock it was built on — the myopia of learning with an equation attached [Levinthal and March, 1993, March, 1991].

> **This is the paper's most interesting empirical-shaped claim and it is NOT empirical.** The fourteen-year series is a constructed illustration with stipulated counts. The figures are rounded to two significant figures here because that is all an illustration supports; they carry no uncertainty because nothing was measured; and **three of the rows in the table above do not reconcile with the filed-item counts, as Section 12.3 sets out.**

Productivity, priced properly. The office's metric is inspections per day, and by that metric Year 14 is 5.7 times Week 2.

	true conditions found per hour	false alarms per hour
Week 2	2.328	0.776
Month 9	20.339	2.034
Year 14	**43.902**	**39.512**

The veteran finds real conditions about nineteen times faster than the novice and generates false alarms about fifty-one times faster. Under equal error costs the veteran's total error count on this building — 20 misses plus 27 false alarms, 47 — is worse than the novice's 44, achieved in a fifth of the time. The trained mid-career point, at 14, is a third of either. The cost table of Section 8.4 prices all three under three loss structures and the calibrated point wins under all three.

Storage consequence. With office-typical stipulated parameters `eta = 0.35`, `delta = 0.08`, `C = 0.20`, and `E` normalised to 1.0 nat:

Point	`M = rE`	`Q_inf` on `M`	`Q_inf` on `V`	stored error
Week 2	0.353	−**0.956**	−0.956	0.000

Point	`M = rE`	`Q_inf` on `M`	`Q_inf` on `V`	stored error
Month 9	1.048	+2.083	+1.875	+0.208
Year 14	1.900	**+5.813**	+1.875	**+3.938**

The novice's steady state is negative — `eta M < C`, the unsustainable regime the storage law declines to hide [Curzi, 2026a] — which is a correct model statement about an observer who cannot continue in that condition and must be moved out of it, in either direction, quickly. The veteran's steady state is the largest of the three and **about 68 per cent of it is stored error** (3.938 of 5.813, 67.7 per cent). Fourteen years of confident, fast, internally consistent structure, two thirds of which the building did not supply. It will be defended, and the defence will be sincere, because from inside there is no difference between the earned third and the manufactured two thirds.

The correction. Diagnosis: over-framed, expertise-decay subtype, driven by a throughput incentive — FM-19 is present at office level, with time-to-commit having fallen across the whole inspection staff. Because the driver is organisational, SP-26 applies first: the office must either change the throughput metric or declare the operating point as policy with a loss function. It changes the metric, adding off-model detection rate to the quarterly report. **That act is itself an intervention on the system being measured, and Section 3.7 records that the framework's own constancy assumptions do not survive it.**

Individual procedure, both levers. *Lever B (raise `P`), primary:* a mandated fifteen-minute off-schedule pass with its own notebook section, question fixed as what is in this building that the schedule does not cover; rotation onto two building types outside the inspector's usual stock, one per quarter. *Lever A (lower `D`), secondary:* every filed item must cite the notebook entry that generated it, and items with no notebook entry cannot be filed. Target gain `a = 0.35` — the working band, half-life 1.61 cycles, comfortably inside the stability condition and nowhere near deadbeat, which SP-23 forbids at this `Q`.

cycle	`ln r`	`r`	band
0	+0.6419	1.900	Marked
1	+0.4172	1.518	Marked
2	+0.2712	1.312	Mild
3	+0.1763	1.193	Mild
4	+0.1146	1.121	Mild
5	+0.0745	1.077	**Calibrated**
6	+0.0484	1.050	Calibrated

Five audit cycles to the calibrated band. At a quarterly cadence that is fifteen months, and the office should expect the inspector's headline productivity to fall by roughly half during it — time on site rising from 41 minutes toward something near the mid-career 118 — while true conditions found per inspection rises and false alarms collapse. **Say that in advance.** A correction that reduces the number people have been watching, in order to improve numbers nobody has been watching, is reversed in month four unless the trade was written down in month zero (SP-25).

Falsification. The office preregisters: the inspector's frozen post-correction `r` predicts a held-out false-discovery rate below 0.15 and an off-model hit rate above 0.40 on a fresh building, and beats persistence from the Year 14 audit, a model using building complexity alone, an autoregressive model on past filed-item counts, and **the same model fitted to inspector/building pairings shuffled at random across the office's staff.** If the shuffled control predicts as well, `r` is a property of the building and not of the inspector, the instrument measures nothing about the person, and the office has agreed in writing to publish that internally and stop using it.

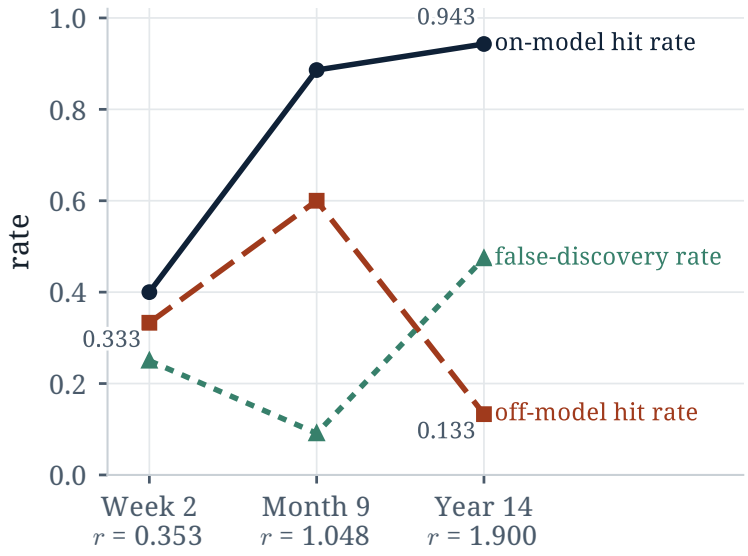


Figure 19.5: **On-model detection rises with tenure while off-model detection collapses below the novice's.** The three career points of the constructed inspector series: on-model hit rate climbs 0.400, 0.886, 0.943 while off-model hit rate goes 0.333, 0.600, 0.133 and the false-discovery rate goes 0.250, 0.091, 0.474. The calibration ratios at the three points are 0.353, 1.048 and 1.900. **The Week-2 and Year-14 hit rates do not reconcile with the filed-item counts of the same worked example, and Section 12.3 prints the discrepancy; the figure plots the source's rates unaltered.** Illustrative construction, not a measurement; generated by `figs/src/CCP-019_fig05.py`.

19.12.3 12.3 Numbers in the source material that do not reproduce

Five printed figures in the material this paper converts fail to reconcile. They are printed here rather than corrected silently, because a reader recomputing the worked examples will hit them.

(1) The starting distortion of the correction-map table is misattributed. The trajectory table of Section 5.2 begins at `ln r = +0.8122` and describes that as the value measured in the seeded review audit. **The seeded review audit returns `ln r = +0.544727`** (Section 12.1). The trajectories are arithmetically correct from `+0.8122`, and `+0.8122` corresponds to `r = 2.2529`, a value computed nowhere in this material. **The arithmetic stands; the attribution does not.** Everything downstream that quotes the divergent case — six cycles at gain 2.3 reaching 3.92 nats — is correct from the stated start and unattached to the audit.

(2) The Week-2 hit rates do not reconcile with the Week-2 filed-item count. The filed report holds 12 items of which 9 are real. The reported on-model hit rate 0.400 implies 14 of 35 and the off-model hit rate 0.333 implies 5 of 15, which is 19 true conditions found — **more true conditions than there are items in the report.** The signature panel's *true conditions found 19/50* carries the same figure. Against this, the per-hour table (`9 / (232/60) = 2.328` true per hour) and the cost table (41 misses out of 50) are both computed from **9**, and both reproduce exactly. **The internally consistent count is 9; the hit rates and the 19/50 row are the entries that fail.**

(3) The Year-14 hit rates do not reconcile with the Year-14 filed-item count. The filed report holds 57 items of which 30 are real. On-model 0.943 implies 33 of 35 and off-model 0.133 implies 2 of 15, which is 35 — again more than the 30 real items recorded, and again matching the panel's *true conditions found 35/50*. The per-hour table (`30 / (41/60) = 43.902`) and the cost table (20 misses) are computed from **30** and reproduce exactly. **If the 30-item count is held and the off-model rate of 2 of 15 is held, the on-model rate is 28 of 35 = 0.800, not 0.943.** In other words, **the headline figure "94 per cent on-model detection" is the number that fails,**

while “13 per cent off-model, below the novice’s 33 per cent” survives on either reading, since the novice’s off-model figure is not in doubt and 2 of 15 is below 5 of 15 whatever the on-model row says. **The paper reports the discrepancy rather than choosing a repair, because choosing one would silently invent a number.**

(4) The crossover claim on the cost table is wrong by a factor of about five. The material states that the office would have to value a miss at more than about six times a false alarm before the veteran’s operating point beats the novice’s. Solving directly with the tabled counts: the veteran (20 misses, 27 alarms) beats the novice (41 misses, 3 alarms) when `20m + 27a < 41m + 3a`, that is when `24a < 21m`, that is when `m/a > 24/21 = 1.142857`. **The crossover is at about 1.14, not about 6.** At a 1.2:1 weighting the veteran already wins, 51 against 52.2. The three tabled cost structures are each correct; only the crossover sentence is wrong.

(5) The claim that no cost structure lets either point beat the calibrated one is false as stated. It is true of the three structures tested. It is not true in general: the novice (41 misses, 3 alarms) beats the calibrated point (10 misses, 4 alarms) when `41m + 3a < 10m + 4a`, that is when `a/m > 31`. **At a false-alarm cost more than thirty-one times a miss, the novice’s operating point is cheaper.** That is an extreme weighting and it is not absurd — it is the regime in which assertion is public, irreversible and expensive, which Section 8.4 names explicitly as the case for a deliberate `r < 1`. The correct statement is *no cost structure among the three tested*.

A sixth item is a silence rather than an error and is recorded in Section 4.6: the octave column of the ceiling table converts at `2/ln 2 = 2.885390` rather than at the direct `1/ln 2 = 1.442695`, and the factor of two comes from the exponent in the corpus’s essence bridge [Curzi, 2026,?], which the source of this paper does not say.

A seventh is a logical gap rather than an arithmetic one and is recorded in Section 5.7: the zero-headroom cancellation requires `E = ln N` and not merely `r = 1`, and the source presents the two as one step.

19.12.4 12.4 The unlicensed-operation audit

The rigour standard requires that wherever a paper computes a mean, a difference or a ratio of something that is only ordinal, it says the operation is not licensed and either justifies it or downgrades the conclusion. **This paper computes many. Every one is listed.**

Operation, where it appears	Scale required	Scale available	Verdict and what is done
`D/P` and `M = (D/P).E`, throughout	Ratio, common true zero	Ordinal at best; counts linked within one item set only	**NOT LICENSED.** Conclusion downgraded to the heuristic ordering (R1); no magnitude reported as measured.
`ln r`, `Delta = abs(ln r)`, Sections 6.1 onward	Ratio on `r`	As above	**NOT LICENSED as magnitude.** Sign and rank order retained as findings.

Operation, where it appears	Scale required	Scale available	Verdict and what is done
Band cut points, Section 6.3	Ordinal	Ordinal	**LICENSED as banding.** The claims that the bands are equal in width and that severe means a factor of two to five are **NOT LICENSED** and are marked as such in Section 6.3.
Arithmetic mean of `r`, Section 6.4	Ratio, and wrong statistic besides	Ordinal	**NOT LICENSED twice over.** Prohibited by SP-15; shown only to demonstrate the bias.
Geometric mean of `r`, mean `ln r`, mean `abs(ln r)`, Section 6.4	Interval on `ln r`	Ordinal	**NOT LICENSED pending instrumentation.** Licensed substitute stated in Section 6.4: median of `ln r`, rank on `abs(ln r)`.
Chain addition `sum_i ln r_i`, Section 8.2	Interval on `ln r`	Ordinal	**NOT LICENSED as arithmetic.** Retained as a **structural** argument, exactly as the source itself states.
Residual distortion `w/a`, Section 8.3	Ratio on both	Neither instrumented	**NOT LICENSED as a magnitude.** The two monotonicities and the unreachability conclusion survive on the stability bound alone.
Percentage changes in the re-audit table, Section 12.1 — the `r` and `abs(ln r)` rows	Ratio	Ordinal	**NOT LICENSED.** Downgraded to *moved toward contact*.
Percentage changes in the same table — hit rates, false-discovery rate, defect counts	Ratio	**Counts with a true zero**	**LICENSED.** These are the rows the conclusions in Section 13 actually rest on.
Minutes on site, per-hour rates, Section 12.2	Ratio	**Duration, true zero**	**LICENSED.**
Miss and false-alarm counts and the cost tables, Sections 8.4 and 12.2	Ratio	**Counts with a true zero**; the loss weights are stipulated	**LICENSED**, conditional on the declared weights being a policy rather than a measurement.

Operation, where it appears	Scale required	Scale available	Verdict and what is done
`Q_inf` on `M` and on `V`, and the stored-error fractions 63 and 68 per cent	Ratio, via `M`	Inherits the demotion through `M`	**NOT LICENSED as magnitudes.** The *sign* of the gap, and the ordering of the three career points by it, survive.
Elasticities `dM/M = dD/D − dP/P + dE/E`, Section 5.4	None — property of the algebra	Not applicable	**LICENSED.** A statement about the model, not about data.
The ceiling `0 <= M <= ln N`, Section 4.6	Requires `M` in nats on the surplus scale	Normalisation is a convention	**LICENSED conditional on the declared convention** (SP-11); the bound itself is a theorem [Curzi, 2026].
`1 − 1/r = 0.4200` at `r = 50/29`, Section 12.1	None — arithmetic identity on two counts	Counts	**LICENSED.** This is the one magnitude in the paper that does not need the demoted scale, because it reduces to `1 − 29/50` on the raw counts.

Two conclusions follow from that table and both are stated against the framework's interest. First, **the licensed rows are the count rows.** Every conclusion in Section 13 that is stated at full strength rests on hit rates, defect counts, durations and error counts — quantities with genuine true zeros — and not on `r`. Second, **the equation's own outputs are the unlicensed ones.** The paper's most quotable magnitudes, 1.724138 and 1.900000 and 63 per cent and 68 per cent, are exactly the numbers the demotion withdraws. That is uncomfortable and it is the correct reading of the framework's own protocol.

19.13 13 Discussion

19.13.1 13.1 What follows

A validation protocol that demotes its author's own keystone formula is doing its job. The result reported here is negative — the moment equation fails the scale-type gate — and the value of the finding is not that the equation is bad but that the gate is real. The framework named the exact quantity it withdrew, left the structural content standing, and priced the instrument that would restore the magnitude at about 123 items per term [Curzi, 2026t]. That sequence — demote, delimit, cost, specify — is what makes the rest of the corpus's numbers worth reading, because a corpus that never demotes anything has not told the reader what its standard is.

The structural content that survives is not small. Four limiting cases predicted from one line, each naming a condition practitioners already recognise. A fixed point at `D = P` which holds for every environment, is not the absence of a frame, is not certainty, and is not comfort. Unit elasticity in every term, so no input's error is amplified. A ceiling inherited from a theorem. A stability band on correction that is the same band an engineer applies to a feedback loop. And a separation of what was stored from what was earned, which is what makes the

framework's least comfortable prediction statable at all: **stored error and stored knowledge occupy the same variable, and the only mechanism that eventually distinguishes them is cost, on a lag.**

The counter-result is the part the framework had least reason to publish. Contact is not an unqualified goal. At the fixed point in a saturating environment the headroom gauge reads exactly zero by cancellation, and an observer optimising raw stored order is driven away from `r = 1` unless the correction-load coefficient exceeds the decay rate. **Where `kappa < delta`, every incentive attached to raw stored order is an incentive to over-frame**, and that is arithmetic about a ledger rather than a claim about anyone's character.

The strongest practical consequence is a requirement on audit design. An audit built from the model under test is a mirror. Twenty-five per cent off-model content is what makes the seeded review audit find anything; thirty per cent is what makes the inspector audit find anything. **Both audits would have certified their subjects as excellent if the audit had been built from the standing guide**, which is the ordinary way audits are built.

19.13.2 13.2 What does not follow

- **It does not compute what is really there.** `E` is exogenous and must be measured by an instrument outside the observer (SP-3).
- **It does not tell anyone whether a moment is good.** It addresses whether the moment is faithful. Fidelity and value are different questions and the framework answers only the first.
- **It does not separate `D` from `P`.** `(D, P, E)` and `(kD, kP, E)` make identical predictions; only the ratio is identified, and any claim about the absolute magnitude of a person's frame or intake needs a second instrument.
- **It does not apply to environments that respond to the observer.** Where the assertion changes `E` the formula as written is inapplicable, and that is a substantial class of real situations — negotiation, diagnosis, forecasting — which the framework excludes rather than models [Manski, 1993].
- **It does not establish that any observer obeys the relation.** That is a postulate, gate G6 has not been run, and the corpus's one preregistered confirmatory test of a different law did not meet its declared margin [Curzi, 2026i].
- **It does not show that `r` is a property of a person.** Gate G7 is unrun and is the gate that decides whether the psychological reading survives at all.
- **The predictive-processing correspondence establishes nothing.** Two formalisms placing their invariant at the same point is interesting and is not evidence, and this bibliography holds no entry for the literature in question.

19.13.3 13.3 Relation to the surrounding literature

Where the framework converges with existing results, it does not discriminate against them. The prediction that on-model detection rises and off-model detection falls with tenure in a stable frame is the myopia of learning [Levinthal and March, 1993] and the homogenisation of a fast-learning organisation [March, 1991], reached from a different mechanism; the progressive absorption of an anomaly into the standing model is documented [Vaughan, 1996]; and the destruction of local variety by a legible scheme is the same movement at institutional scale [Scott, 1998]. **A confirmatory observation of H5 therefore does not distinguish this framework from any of them**, and the demand for a differential prediction — one the simpler accounts do not also make — is the correct demand to place on a formal social model [Durlauf, 2005]. The framework's candidate differential predictions are the quantitative ones: the exact `1 − 1/r` relation between over-framing and the false-discovery rate, and the knife edge at `kappa = delta`. **Both are exactly the predictions the scale-type demotion withdraws**, which is the honest summary of where the framework currently stands.

Where it is exposed, the exposure is measurement-theoretic and not empirical. Michell [1997] is the objection, conjoint measurement is the test [Luce and Tukey, 1964], and neither has been run. The framework's response is to demote rather than to argue, which is the right response and is not a substitute for the test.

Where the literature supplies something the framework needs, it is used and named. The logarithm is warranted by multiplicative dynamics rather than by psychophysics [Peters and Gell-Mann, 2016]. The stability band is standard control theory [Åström and Murray, 2008, Elaydi, 2005]. The reflexive fixed point is documented in the field with dates and prices [MacKenzie and Millo, 2003], which is the existence proof the corpus's reflexive law needs and does not itself supply. The Rasch route out is specified because separability is what supplies the true zero [Rasch, 1960, Andrich, 1988]. And the ceiling on individual-level prediction from very rich data [Hofman et al., 2017, Salganik et al., 2020] is the reason to expect gate G7 to be hard, and the reason not to promise that it will pass.

On the corpus's own terms, this paper is one entry in a register. The demotion recorded here appears in the corpus's standing register of its own demotions and inversions [Curzi, 2026], alongside a formulary that demotes five of its own nine entries on the same ground [Curzi, 2026q]. **A framework's credibility is not built by having no weaknesses. It is built by being the party that found them first**, and the register is the evidence that this is the practice rather than the slogan.

19.14 14 Verification

The results of this paper compound from seven postulates. Each is atomic, each is checkable on its own, and each carries its status. **Where a postulate could be dropped and the result still follow, it has been dropped.**

19.14.1 14.1 Dimensional audit

`[E]` denotes the dimension of environmental structure, adopted here as **nats** — a dimensionless information measure — and `[T]` the dimension of time, whose unit is the correction cycle or audit interval.

Symbol	Dimension	Unit	Status
`D`	`[X]`	recorded items per condition present, or precision, as declared	**stipulated commensurable with** `P` (SP-1); the framework fixes no unit and says so
`P`	`[X]`	the same unit as `D`, necessarily	**stipulated**; identical to `D` by SP-1 or the fraction is malformed
`r = D/P`	`[1]`	none	**derived** dimensionless; the only quantity the moment measurement identifies
`ln r`, `Delta`	`[1]`	nats	**derived**; convertible to bits at `1/ln 2 = 1.442695`

Symbol	Dimension	Unit	Status
`E`	`[1]`	nats of usable structure	**stipulated** by the register convention; requires an independent instrument (SP-3)
`M = rE`	`[1]`	nats	**derived**: `[M] = [E]` exactly, since `r` is dimensionless
`N`	`[1]`	a count	**declared** (SP-11): the raw number of distinguishable categories
`ln N`	`[1]`	nats	**derived**; the ceiling on `M`
`V = E min(r,1)`	`[1]`	nats	**derived**; same dimension as `E`
`a`	`[1]`	per cycle, dimensionless fraction	**design parameter**, chosen and not measured
`w`	`[1]`	nats per cycle	**stipulated**; must be estimated locally, and the framework supplies no value
`eta`	`[Q]` per nat per step	proxy units per nat per cycle	**stipulated** here; measured in the storage law [Curzi, 2026a]
`delta`, `kappa`	`[1]`	per cycle	**stipulated** dimensionless fractions per step
`Q`, `Q^V`, `Z`	`[Q]`	declared order proxy	**derived** from the storage law; two systems' values are incomparable unless they declare the same proxy
`C`, `c_0`	`[Q]` per step	proxy units per cycle	**stipulated**
`h`	`[Q]` per step	proxy units per cycle	**derived** headroom

Does the equation balance? `M = (D/P).E` balances if and only if `[D] = [P]`, which makes `r` dimensionless and forces `[M] = [E]`. That is the whole of gate G1 and it is conditional on SP-1: **a ratio of two quantities that are not the same kind of quantity is not dimensionless, it is meaningless, and no amount of downstream arithmetic repairs it.** The self-storage law `Q_{t+1} = (1−delta)Q_t + eta M_t − C_t` balances if and only if `eta M` carries `[Q]` per step; since `M` is dimensionless in nats, **`eta` must carry `[Q]` per nat per step and is not dimensionless** [Curzi, 2026a].

Where the framework has no operational unit, it is said here. `D` and `P` have **no operational unit fixed by the framework**; SP-1 requires only that whatever unit is chosen be the same for both, and the two worked audits choose *recorded items per condition present*. `E` is denominated in nats by convention and not by measurement. `Q` is counted in whatever the declared order proxy counts. **Comfort itself, the underlying construct of the corpus, has no operational unit at all** and is treated as a dimensionless magnitude by stipulation, not by measurement;

Section 15 records that the quantitative structure required for that stipulation has not been tested.

19.14.2 14.2 The postulates

P1 — The moment equation is the model form: `M = (D/P).E`. *Status:* **POSTULATED** — this work's own proposal, with no external support; the corpus's coverage audit records that none was found. It carries no empirical weight on its own.

P2 — Permissible arithmetic is determined by scale type: multiplication and division require a ratio scale with a true zero. *Status:* **CITED** — Stevens [1946], with the invariance formulation in Narens [1996], Luce [2002].

P3 — No instrument currently places `D` and `P` on a common ratio scale with a shared true zero. *Status:* **POSTULATED** — this is an empirical claim about the state of instrumentation and is the only empirically loaded premise in the chain. It is corroborated by the corpus's own protocol run [Curzi, 2026t] and is consistent with Michell [1997]'s general finding, but it is not proved by either, and a counterexample would consist of exactly one working linked instrument.

P4 — Surplus is bounded: `0 <= F(u) <= ln N`, with `F(u) = ln(1 + (N−1)u)`. *Status:* **CITED** — the range result of the surplus functional [Curzi, 2026].

P5 — A first-order linear recurrence `x_{t+1} = c x_t` converges to 0 from any start if and only if `|c| < 1`, and the transient falls by `|c|` per step. *Status:* **CITED** — Elaydi [2005], Kelley and Peterson [2001]; the metric-space form supplying uniqueness and start-independence is Banach [1922], Curzi [2026].

P6 — `ln(xy) = ln x + ln y`, so a multiplicative composition is additive in logarithms. *Status:* **ARITHMETIC** — elementary, and it is the whole content of the composition argument of Section 6.2.

P7 — `E` is exogenous: the observer's assertion does not change what is available to be observed. *Status:* **POSTULATED** — a modelling restriction, enforced by SP-12 and violated by an important class of situations. Manski [1993] is the reason it cannot simply be assumed away in the class where it fails.

Two postulates carry all the empirical weight and both are POSTULATED: P3 and P7. Everything else is arithmetic, a definition, or an external citation. A reader who rejects P3 recovers the quantitative reading of the equation and gains an obligation to produce the instrument; a reader who rejects P7 loses the equation entirely in the reflexive class.

19.14.3 14.3 The compounding chain

```
P1              -> L1   M = rE with r = D/P; three degrees of freedom, not four;
                        (D,P,E) and (kD,kP,E) are observationally identical.   ARITHMETIC

P1 + P6         -> L2   ln M = ln D − ln P + ln E, so every elasticity is
                        exactly +1 or −1.                                      ARITHMETIC

P2 + P3         -> R1   the magnitude of D/P is not a measured quantity; on
                        ordinal data the equation is a heuristic ordering and
                        may not be printed as a fraction or computed.
                                                  CITED (stevens1946theory) + POSTULATED (P3)

L1              -> R2a  D = P  <=>  r = 1  =>  M = E for every E; the fixed
                        point is an identity, not an approximation.            ARITHMETIC

L1 + P4         -> R2b  0 <= M <= ln N under the register convention, so the
                        P -> 0 divergence saturates at ln N rather than
```

```
                        diverging.                  CITED (ccp:CCP-A02) + DEFINITION (normalisation)

P5              -> R3   ln r_{t+1} = (1-a) ln r_t converges iff |1-a| < 1,
                        i.e. 0 < a < 2; neutral at a = 0; divergent at a >= 2;
                        half-life ln 2 / (-ln|1-a|).       CITED (elaydi2005introduction)

R3              -> R3b  at a = 2.3 from ln r = 0.8122, six steps give
                        -1.0559, +1.3726, -1.7844, +2.3197, -3.0157, +3.9203.  ARITHMETIC

L1              -> R4a  V = E min(r,1); confabulated fraction 1 - 1/r for
                        r > 1; missed fraction 1 - r for r < 1; V is uniquely
                        maximised at r = 1.                                    DEFINITION

R4a             -> R4b  at r = 50/29 = 1.724138, 1 - 1/r = 1 - 29/50 = 0.4200
                        exactly, which is the audit's false-discovery rate
                        21/50.                                                 ARITHMETIC

R4b + ccp:CCP-007 -> R4c  with eta = 0.35, delta = 0.08, C = 0.20, E = 1:
                        Q_inf on M = 5.043, on V = 1.875, stored error 3.168,
                        62.8 per cent of the store.       ARITHMETIC on stipulated parameters

R4a             -> R5   re-audit at r = 32/31 . 40/40 = 1.032258: false
                        discovery 21/50 = 0.4200 -> 1/32 = 0.03125, a 92.6 per
                        cent reduction; true defects 29/40 -> 31/40.           ARITHMETIC

R4a             -> R6   inspector series r = 12/34 = 0.352941, 44/42 = 1.047619,
                        57/30 = 1.900000; stored error at year 14 is 3.9375 of
                        5.8125, 67.7 per cent.            ARITHMETIC on stipulated counts

R2a + ccp:CCP-007 -> CR1 at r = 1 AND E = ln N: h = eta ln N - delta Q_inf - C
                        = eta ln N - (eta ln N - C) - C = 0 exactly.            ARITHMETIC

L1 + ccp:CCP-007 -> CR2 with C_t = c_0 + kappa Z_t, dQ_inf/dM =
                        eta (1 - kappa/delta)/delta, which is negative iff
                        kappa > delta.                                         ARITHMETIC
```

Every line is checkable from its own stated inputs. R2a, R2b, R3, R3b, R4a to R4c, R5, R6, CR1 and CR2 are arithmetic or definitional consequences of one postulated model form (P1), one cited range result (P4), one cited stability condition (P5) and elementary algebra. **R1 — the paper's headline — is the only result carrying empirical weight, and it consumes one citation (P2) and one postulate with no external support (P3).**

What this chain does not establish. It does not establish that any observer obeys P1. That is the modelling postulate, gate G6 has not been run, and the one preregistered confirmatory test of a different law reported in this corpus did not meet its own declared margin [Curzi, 2026i]. **A chain whose steps are ARITHMETIC and CITED delivers the mathematics with certainty and delivers nothing whatever about observers.**

19.14.4 14.4 The verification instance, computed by hand

Check 1 — the false-discovery identity. The audit records 50 items filed, 29 of them real. Then `r = 50/29`. Compute `1/r = 29/50 = 0.58` exactly, so `1 - 1/r = 0.42` exactly. Independently, the false-discovery rate is `(50 - 29)/50 = 21/50 = 0.42`. **The two are the same number because they are the same fraction written twice**, and the decimal expansion `50/29 = 1.7241379...` rounds to the printed `1.724138`, from which `1 - 1/1.724138 = 0.420000023...`, agreeing to six

decimals. **ARITHMETIC**, and it is the one magnitude in the paper that does not need the demoted scale, because it reduces to a ratio of two counts with a true zero.

Check 2 — the divergent correction, six steps. With `a = 2.3` the pole is `1 - a = -1.3`. From `ln r_0 = 0.8122`:

```
step 1:  -1.3 x 0.8122     = -1.055860
step 2:  -1.3 x -1.055860  = +1.372618
step 3:  -1.3 x 1.372618   = -1.784403
step 4:  -1.3 x -1.784403  = +2.319724
step 5:  -1.3 x 2.319724   = -3.015641
step 6:  -1.3 x -3.015641  = +3.920333
```

which reproduces the printed row `-1.0559 +1.3726 -1.7844 +2.3197 -3.0157 +3.9204` to four decimals. `exp(3.9204) = 50.42`, which is the factor of fifty from contact. **ARITHMETIC.** Contrast the convergent case at `a = 0.35`, pole `0.65`: `0.65 x 0.8122 = 0.527930`, then 0.343155, 0.223050, 0.144982, 0.094238 — below 0.10 at the fifth step, which is the tabled entry. **The whole of the stability claim is visible in ten multiplications.**

Check 3 — the headroom cancellation. With `S = ln N`, `Q_inf = (eta ln N - C)/delta`, so `delta Q_inf = eta ln N - C` identically, and `h = eta ln N - delta Q_inf - C = eta ln N - (eta ln N - C) - C = 0`. **Both `delta` and `C` cancel and no parameter values are needed.** ARITHMETIC — and Section 5.7 records that the step requires `E = ln N` and not merely `r = 1`.

Check 4 — the knife edge. With `C = c_0 + kappa Z` and `Z_inf = eta(M - V)/delta`, at `r > 1` we have `V = E` fixed, so `dZ_inf/dM = eta/delta` and

```
Q_inf = (eta M - c_0 - kappa Z_inf)/delta
dQ_inf/dM = (eta - kappa . eta/delta)/delta = eta(1 - kappa/delta)/delta
```

which is negative exactly when `kappa > delta`. At the tabled values `eta = 0.35, delta = 0.08, kappa = 0.02`, the derivative is `0.35 x (1 - 0.25)/0.08 = 3.28125` per nat, positive, so raw stored order rises in `r` — and the tabled `Q_inf` values 1.8750, 3.5156, 8.4375, 15.0000 at `r = 1, 1.5, 3, 5` rise accordingly. **ARITHMETIC.**

Check 5 — the aggregation rows. For `{0.5, 2.0}`: arithmetic mean `1.25`; geometric mean `sqrt(0.5 x 2.0) = 1.000`; mean `ln r = (-0.693147 + 0.693147)/2 = 0`; mean `|ln r| = 0.693147`. For `{0.4, 0.8, 1.5, 2.5}`: the product is `1.2`, so the geometric mean is `1.2^{1/4} = 1.046635` and `ln` of that is `0.045580`; the four absolute logarithms are 0.916291, 0.223144, 0.405465 and 0.916291, averaging 0.615298. **ARITHMETIC**, and the point of the check is that the three statistics disagree by construction, not by accident.

19.14.5 14.5 Limits, asymptotics, and the model's valid domain

`P -> 0^+`. `M -> +infinity` in the unbounded form, and `M -> ln N` under the register convention. **The limit is not the pathology; the saturation is.** The bounded reading — the observer sees the maximum their categories permit, sourced entirely from themselves — is the better description and is what the bounded model predicts. **`P = 0` exactly is outside the model's domain**, because `D/0` is undefined and the bounded form is a repair applied after the fact rather than a property of the expression.

`D -> 0^+`. `M -> 0`. Everything arrives and nothing is made of it. **At `D = 0` exactly, `M = 0` and the equation is well formed**, unlike the `P` case; this asymmetry between numerator and denominator is real and is the same asymmetry that makes halving intake double the moment while halving frame merely halves it.

`E -> 0`. `M -> 0` for every finite `r`. A multiplier on zero is zero. `E = 0` is admissible and gives a true reading of an empty environment.

`r -> 0` and `r -> infinity`. Under-framing and over-framing without bound; the second saturates at `ln N` and the first does not need a bound because it is already at zero. `r < 0` is inadmissible: `D` and `P` are counts or precisions and both are non-negative.

`a -> 0^+`. The pole tends to 1, the half-life to infinity, and the residual distortion `w/a` to infinity. **At `a = 0` exactly the map is the identity: neutrally stable, no contraction, and whatever distortion the observer has, they keep.** That is the structural twin of the `delta = 0` pathology in the storage law [Curzi, 2026a], and it says the same thing: **without a leak there is no attractor.** It is also the default state of an unaudited observer.

`a -> 1`. Deadbeat: the pole is zero and the distortion is annihilated in one step. **The mathematics is fine and the application is not**, because it means discarding the whole existing frame in one period, and the cost violates the admissibility condition `C_t <= (1−delta)Q_t + eta M_t` at any material stored order (SP-23).

`a -> 2^−` and `a >= 2`. The pole tends to 1 in magnitude with alternating sign; convergence becomes arbitrarily slow, and at `a = 2` the map is `ln r_{t+1} = −ln r_t`, a period-two orbit that never converges and never diverges. **Above 2 it diverges geometrically.** `a >= 2` is outside the model's admissible design range.

`N -> 1`. `ln N = 0`, so the ceiling and the achievable moment collapse together and `M = 0` for every `r`. **A single category has zero available moment, not a small one.** **`N -> infinity`.** The ceiling grows without bound but only logarithmically, and the marginal gain `ln N − ln(N−1)` behaves as `1/N` and vanishes. **The limit is not a licence to raise `N`**; it is why Section 4.6 concludes that a distortion problem cannot be solved by learning more categories.

`w -> 0`. Residual distortion vanishes and the correction map is the undriven one; a static environment is the only one in which contact is attainable rather than merely approachable. **`w` large.** Beyond about `0.10` nats per cycle at the stated band width, no admissible gain holds calibration and only the cadence remains.

`kappa -> delta`. The knife edge: `dQ_inf/dM = 0` and raw stored order becomes indifferent to `r`. Below it the raw ledger rewards over-framing; above it the raw ledger penalises it. **The framework supplies no measured value of either parameter.**

Stated domain. The relation as used here is valid for `D >= 0`, `P > 0`, `E >= 0`, `N >= 2`, correction gain `a` in `(0, 1]` for applied use and `(0, 2)` for the mathematics, and an environment that does not respond to the observer's assertion. Outside that domain the expression still produces values and they should not be interpreted as statements about an observer.

19.14.6 14.6 Invariance

The scaling invariance, and why it is also the identification failure. The relation is invariant under `(D, P) -> (kD, kP)` for every `k > 0`, since `r` is unchanged and `M` depends on `D` and `P` only through `r`. **That is a genuine one-parameter invariance and it is exactly why the split is not identified**: the group of admissible transformations is larger than the data can distinguish. `E` is not in the group — rescaling `E` rescales `M` proportionally — so `r` is the unit-free quantity of the relation and `E`, `M` and `V` are the unit-carrying ones. **Two observers may be compared on `r` directly; they may not be compared on `M` unless they share a declared `E` instrument.**

The invariance in `E` at the fixed point. `M = E` holds at `r = 1` for every `E`, which is a scale-freedom claim and is stated as one: the fixed point is invariant under any rescaling of the environment. **What would break it:** an `r` that varies with `E` — which is precisely the apophenic case, where a thin `E` makes any `D` produce a large `r`, so **the invariance already fails in one of the framework's own named subtypes** and holds only where `D` and `P` are set independently of the richness of the environment. That is a real restriction and it is stated as one rather than assumed away.

The invariance of the distortion measure under change of logarithmic base. `ln r` in nats converts to bits at `1/ln 2 = 1.442695`, and the sign, the ordering and the fixed point at zero are all preserved. **The band table is not invariant**, because its cut points are declared in nats; a framework reporting in bits must restate them.

How many scales the claim has been exercised at: two, both constructed. One reviewer and one inspector, both synthetic, both by one author, with no independent replication of anything. The internal/external reading of Section 5.3 asserts the same fixed point at a third scope and exercises it at none.

19.14.7 14.7 Analogy discipline, and the precision of the numbers

The correction map is a shared mathematical form, not a derivation from physical premises. The map `ln r_{t+1} = (1−a) ln r_t` and the first-order discrete section are governed by the same difference equation, so every theorem about such recurrences applies to both [Elaydi, 2005, Oppenheim et al., 1996, Åström and Murray, 2008]. Nothing about charge, capacitance or current is claimed to hold of an observer, and **no physical result's authority transfers.** The same label attaches to the exponential-smoothing twin [Holt, 1957, Brown and Meyer, 1961] and to the fixed-gain filter reading [Kalman, 1960].

The contraction theorem is a derivation within the mathematics and a shared form in the application. The map genuinely is a contraction on the reals for `|1−a| < 1`, so Banach [1922] applies without qualification *to the map*. **That an observer's correction behaviour is described by the map is a postulate, not a derivation**, and the theorem's certainty does not transfer to the observer.

Ashby's law of requisite variety, applied to `M <= ln N`, is an analogy. Ashby's theorem concerns a regulator, a disturbance set and a goal set [Ashby, 1958, 1956], and this paper does not show that an observer's category repertoire satisfies those premises. **The bound does not depend on the reading**; it follows from the surplus law's range result and the normalisation convention [Curzi, 2026]. The reading is offered as motivation and is never used as evidence.

The predictive-processing correspondence is an analogy in the strict sense. A suggestive correspondence with no formal warrant, legitimate as motivation and never as evidence. The premises under which a precision-weighted average is the correct model of an observer have not been shown to hold; the two objects differ structurally, one interpolating and the other scaling; and **this bibliography contains no entry for the literature in question**, so the correspondence borrows no authority because there is none available to borrow.

Entropy, named. The `ln N` appearing throughout is **Shannon entropy** of a uniform distribution over `N` categories, a dimensionless information measure reported in nats [Shannon, 1948, Cover and Thomas, 2006]. It is **not** thermodynamic entropy, which carries `[energy][temperature]^−1`. The framework identifies the moment with neither.

Precision of the numbers. The numbers here are of three kinds and are quoted accordingly. *Defined or derived-from-defined* — `ln 2 = 0.693147`, `ln 5 = 1.609438`, the half-life formula, the geometric-mean rows, `1 − 29/50 = 0.42`, and the constant `k = (9/8)^2/phi^2 = 0.483425732988414` where it is mentioned in Section 4.6 — are exact to the precision printed and may be recomputed further. *Illustrative model output* — every entry in every trajectory table, every `Q_inf`, every `r` in both worked audits, the stored error fractions — is the output of stipulated parameters and **carries no measurement uncertainty because there was no measurement**; the decimals record the arithmetic and not the confidence, and the summary figures are rounded to two significant figures in the text for that reason. *Measured or fitted* quantities **do not appear anywhere in this paper**, because nothing in it was measured. Any application of this framework to real data must report fitted parameters with intervals, and Section 15 records that the corpus's one such report gives an interval on one parameter and none on several others.

19.15 15 Limitations

The paper's headline result is a failure of its own equation, and everything below follows from taking that seriously rather than from softening it.

The moment equation failed the scale-type gate and the failure is not repaired anywhere in this paper. `M = (D/P).E` contains a division and a multiplication, which require all three inputs on a ratio scale with a true zero [Stevens, 1946]; no instrument links `D` and `P` across item sets; the magnitude of `D/P`, the distortion as a number, and the quantitative join to the store are therefore withdrawn as measured quantities [Curzi, 2026t]. **Every quotable magnitude in this paper — 1.724138, 1.900000, 63 per cent, 68 per cent — is on the withdrawn side of that line**, and Section 12.4 lists every operation performed on them. What survives is the structure, and the structure is what the Conclusion states.

Michell's objection is conceded, not answered, and conjoint measurement has not been run. Michell [1997, 1999] argue that psychology systematically fails to test whether its attributes are quantitative before measuring them. **The demotion of Section 3.5 is that remedy applied once. Applying it once is not adopting it.** The same test bears on `u`, on `Q`, on `E` and on comfort itself, and none has been subjected to it. The formal procedure that would settle the question [Luce and Tukey, 1964, Krantz et al., 1971] has not been applied to any quantity in this corpus. **Until it is, every quantitative treatment in the corpus, including the ones this paper leaves standing, is a stipulation and not a finding.** A reliability coefficient would not close the gap [Cronbach, 1951, Sijtsma, 2009], and the construct validation this instrument owes — convergent and discriminant, across methods — has not been performed [Campbell and Fiske, 1959, Cronbach and Meehl, 1955].

Both worked audits are constructed, and one of them does not reconcile with itself. No dataset was collected. Every number in Section 12 is the output of stipulated counts and stipulated parameters. **Three rows of the inspector series — the Week-2 and Year-14 on-model and off-model hit rates and the true-conditions-found row — are mutually inconsistent with the filed-item counts of the same example**, as Section 12.3 sets out; the per-hour and cost tables reproduce from the filed-item counts and the hit rates do not. **The headline figure "94 per cent on-model detection" is the entry that fails**; "13 per cent off-model, below the novice's 33 per cent" survives on either reading. The paper prints the discrepancy and does not choose a repair, because choosing one would invent a number. Two further printed claims are wrong and are corrected in Section 12.3: the cost-crossover figure (about 1.14, not about 6) and the claim that no cost structure lets either non-central point beat the calibrated one (false above a false-alarm-to-miss weighting of 31).

The 94/13 result is the paper's most interesting empirical-shaped claim and it is not empirical. It is a constructed fourteen-year series with stipulated counts. It carries no uncertainty because nothing was measured, and no bare point estimate in this paper should be read as a measurement.

Gates G6 and G7 have not been run and G7 is the one that decides everything. No preregistered test of the moment equation exists. **Whether `r` is a property of an observer or a property of an item set is unknown**, and if `r` for one observer swings by more than the band width across equivalent item sets, the psychological reading collapses. The ceiling on individual-level prediction from very rich longitudinal data [Hofman et al., 2017, Salganik et al., 2020] is a reason to expect that gate to be hard. The corpus's one preregistered confirmatory test of a different law reported a 2.8 per cent held-out advantage over a matched autoregression against a declared 10 per cent margin and recorded the verdict as *predictive advantage not demonstrated* [Curzi, 2026i]; **that is the current evidential status of every applied claim resting on this paper, and it is weaker than the mathematics.**

The reflection problem is unanswered and no remedy is adopted. The corpus's founding law is a reflexive specification [Curzi, 2026] whose endogenous effect is not identified from observational data [Manski, 1993], with the same confounding of influence and selection in network settings [Shalizi and Thomas, 2011]. **This paper's response is an exclusion (SP-12, FM-10), not a solution**, and the excluded class is exactly where the reflexive law operates. The most realistic available repair is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001]; the seeded audit's externally adjudicated item set is a step toward it and not a substitute. **Every applied reading of `r` in this paper is therefore a modelling position and not an identified empirical result.**

The Lucas critique applies to this framework's own parameters and the position taken is narrow. The paper asserts that a throughput metric drives `r` upward across a population and prescribes changing the metric, then treats the gain and the drift rate as constant over the correction horizon. **Both cannot hold without qualification** [Lucas, 1976, Hacking, 1995]. The position adopted is that the parameters are stable only over one declared audit cycle and only while `r` and its components are not themselves reported targets; the moment off-model detection enters a quarterly report it becomes a target and ceases to be a good measure [Goodhart, 1975, Campbell, 1979]. **Where `r` is reported upward the constancy assumption fails**, and Section 12.1 shows Goodhart's law arriving on schedule inside the framework's own worked example. A framework that tells an organisation its inspectors are two-thirds confabulating should expect that statement to change the inspectors [Ferraro et al., 2005, McCambridge et al., 2014].

The observer cannot be separated from the instrument. Calibration cannot be self-assessed, for the structural reason of Section 7.3 and for the general reason that people cannot reliably report the processes producing their judgements [Nisbett and Wilson, 1977]. **But the external audit relocates the problem rather than removing it**: the adjudicators who build the seeded item set are observers with frames, and prior expectation measurably shapes what is perceived and reported [Rosenthal and Fode, 1963, Rosenthal and Jacobson, 1968] — an effect whose size is itself contested and smaller than its popular reading [Jussim and Harber, 2005].

The five behavioural estimators are unvalidated and none has been run. E1 to E5 are proposals. None has been compared against a criterion, none has a reported reliability, and SP-16's requirement of two from different rows is a design rule and not a demonstration.

The band table and the diagnostic cut points are conventions. 0.10, 0.35, 0.69 and 1.61 nats were chosen for round numbers and for correspondence with recognisable states. They are not measured thresholds, the differential diagnosis of Section 7.2 imports them as decision rules, and **a reader who changes them changes every band assignment in this paper.**

Parameters chosen rather than measured. In both worked audits, `eta`, `delta`, `C`, `c_0`, `kappa`, `E`, the correction gain `a` and every seeded count are stipulated. The drift rates in Section 8.3 are stipulated. The loss weights in Section 8.4 are stipulated. **The framework supplies no measured value for any parameter it uses.**

`N_eff = 1 + (N−1)u` is not a Hill number and no external literature contains it. `1/HHI` and `exp(H)` are the Hill numbers of order two and one and are grounded [Hill, 1973, Jost, 2006]; the affine interpolation is this framework's own construction. The ceiling in this paper uses the raw `N`, and a reviewer reaching for the Hill-number literature will find it does not apply to the effective count.

The functional form of the surplus the ceiling inherits from is underdetermined by the evidence. A sub-linear power law [Latané, 1981] and an `n log n` form [Briscoe et al., 2006] fit the qualitative facts as well as `ln(1 + (N−1)u)` does, and **nothing assembled here discriminates between them.** The logarithm is chosen for its axiomatic characterisation and is not established by measurement. The one genuine strengthening available is that a logarithmic evaluation follows from multiplicative dynamics with no psychological assumption [Peters and Gell-Mann, 2016], which is what this paper's own use of the logarithm rests on; the psychophysical warrant is weaker, contested [Stevens, 1957, Krueger, 1989], and range-dependent [Teghtsoonian, 1971].

Two claims in this paper are uncited because the bibliography contains nothing that supports them, and they are named here rather than given a citation that does not fit. First, **the claim that loss asymmetry between misses and false alarms moves the optimal operating point** (Section 8.4) is carried by this paper's own cost arithmetic; this bibliography holds no signal-detection-theory entry, and no citation is attached. Second, **the predictive-processing correspondence** (Section 5.7) is stated with no external support at all, because the bibliography holds no entry for that literature; it is badged as an analogy and must not be read as borrowing that programme's authority.

Where `r` is read as a similarity or distance between frame and intake, it cannot be assumed to behave metrically, since psychological similarity is asymmetric and violates the triangle in-

equality [Tversky, 1977]. The framework treats the two directions as symmetric under the logarithm, which is a property of the coordinate and not a demonstrated property of the psychology.

No scaling, power-law or heavy-tail claim is made in this paper, so the distribution fitting procedure any such claim would require [Clauset et al., 2009] does not apply. The one distributional statement made — that distortion decays geometrically under correction — is a property of a postulated linear map and has not been tested against data.

Unmet rigour checks, printed rather than left to a reviewer. *(i) Model comparison.* This paper fits nothing, so no information criterion applies to it; gate G6 requires a matched autoregression with the same parameter count, and that requirement is stated and unrun. *(ii) Effect sizes and intervals.* Nothing in Section 12 carries an interval because nothing in Section 12 was measured; that is honest for a constructed example and is not a substitute for the intervals a real application owes. *(iii) Scale type.* Handled at length above and audited in Section 12.4. *(iv) Independent replication.* Every worked system in this paper and in the corpus's applied papers was constructed or surveyed by one author; there is no independent replication of any parameter estimate, and the reproducibility literature is unambiguous about what that is worth [Open Science Collaboration, 2015, Ioannidis, 2005]. *(v) Directional versus point prediction.* Meehl [1967] observes that directional predictions are nearly costless in the social sciences; this framework does make point predictions with constants, which is the right answer to that objection — **and the scale demotion has withdrawn exactly those points.** *(vi) Differential prediction.* H5 is predicted equally by two existing literatures [Levinthal and March, 1993, March, 1991], so confirming it would not discriminate [Durlauf, 2005]. *(vii) The field's own indictment.* The standard charge against physics-styled social modelling — sub-standard statistical procedure, universality claimed from limited data, models built on assumptions known to be false [Gallegati et al., 2006] — is a checklist this paper is exposed to on the second and third counts, and answers on the first only by having run no statistics at all. Cartwright [1999] offers the generous reading and the framework takes it: a law of this kind holds inside a specified stable arrangement or not at all, and Section 14.5 states the arrangement.

Nothing in this paper is a human-subjects finding. No participants were recruited, no interventions were run, and the consent and decline requirements of SP-13 and SP-28 are design constraints on prospective work rather than reports of past practice.

19.16 16 Conclusion

The moment equation `M = (D/P).E` is the corpus's own formula, and it fails the corpus's own scale-type gate. The expression contains a division and a multiplication, so frame, intake and environment must each be measured on a ratio scale with a true zero and a defensible unit [Stevens, 1946]. No instrument places `D` and `P` on a common ratio scale with a shared zero. **R1, the demotion rule: where only ordinal data exist for frame or intake, the moment equation is a heuristic ordering — more frame relative to intake produces a larger and more model-shaped moment — and may not be printed as a fraction, computed, or used to generate a number.** The magnitude of `D/P`, the distortion as a number, and the quantitative join from the moment to the store are withdrawn as measured quantities. That is the paper's headline result, it is negative, and it is the remedy Michell [1997] prescribes applied once — the honest position being that it should be applied consistently rather than once, and that the conjoint-measurement test which would settle the matter [Luce and Tukey, 1964] has not been run for any quantity in this corpus.

R2 — the fixed point. `D = P`, equivalently `r = 1`, gives `M = E` **for every `E`**: an identity and not an approximation, holding at both personal and institutional scope, and meaning matched frame rather than absent frame. Every elasticity of the relation is exactly `+1` or `−1`, so a one per cent error anywhere is a one per cent error in the moment and no term's error is amplified. The ceiling `0 <= M <= ln N` — with `N` the raw count of distinguishable categories, inherited from the surplus range result [Curzi, 2026] under a declared normalisation — makes the one

divergence in the formula finite: as `P -> 0` the moment **saturates at `ln N`** rather than diverging, which is a better description of dogma than an infinity is.

R3 — the correction map. `ln r_{t+1} = (1−a) ln r_t` is a first-order linear recurrence, **attracting for `0 < a < 2`, neutral at `a = 0`, oscillating but convergent for `1 < a < 2`, and divergent at `a >= 2`**, with distortion half-life `ln 2 / (−ln|1−a|)` cycles. From a distortion of 0.8122 nats, six cycles at gain 2.3 reach 3.9203 nats — `r` near 50, a factor of fifty from contact. **There is no decay term on distortion: an observer does not drift back into contact, and calibration is a maintenance item with a schedule.** In a moving environment the residual distortion is `w/a`, drift over gain, which is not zero and cannot be made zero; above about a tenth of a nat of drift per cycle no admissible gain holds calibration and only the cadence remains.

R4 — veridical throughput. `V = E min(r,1)` is the part of the moment the environment supplied, and it is uniquely maximised at the fixed point; the confabulated fraction is `1 − 1/r` for `r > 1`. In a seeded review audit with 25 per cent off-model items, 50 filed items of which 29 were real give `r = 50/29 = 1.724138` and a confabulated fraction `1 − 29/50 = 0.4200` **exactly equal to the false-discovery rate `21/50`** — the one magnitude in this paper that survives the demotion, because it reduces to a ratio of two counts. At stipulated `eta = 0.35`, `delta = 0.08`, `C = 0.20`, the reviewer's steady state is **about 63 per cent stored error**.

R5 — correction works, and the metric being watched gets worse. At gain 0.35 the same reviewer reaches `r = 1.032` after six cycles: false discovery falls 92.6 per cent, true defects found rise from 29 of 40 to 31 of 40, off-model detection quadruples — **and on-model detection falls from 93 to 90 per cent**, which is a small regression by the metric most organisations actually watch.

R6 — the inspector series. Across a constructed fourteen-year career the same inspector returns `r = 0.353` with a **negative** steady state at week two, `r = 1.048` at minimum loss in month nine, and `r = 1.900` at year fourteen with **about 68 per cent of the store being error**, on-model detection at its career maximum and **off-model detection about 13 per cent, below the nine-day novice's 33 per cent**. Three rows of that series do not reconcile with its own filed-item counts and the discrepancy is printed in Section 12.3 rather than repaired.

CR — the counter-result. Contact is not an unqualified goal. At `r = 1` in an environment that saturates the observer's repertoire, the input sits at its ceiling and the headroom gauge reads `h = eta ln N − (eta ln N − C) − C = 0` **exactly, by cancellation**: a perfectly calibrated observer maximises fidelity and has no room left to build. Worse, an observer optimising raw stored order is driven away from `r = 1` altogether, because `dQ_inf/dM = eta(1 − kappa/delta)/delta` is negative **only when the correction-load coefficient exceeds the decay rate**; below that knife edge the raw ledger rewards over-framing without bound. **Calibration is a property of the instrument; the operating point is a policy; and an undeclared departure from `r = 1` is instrument error and not judgement.**

All of this is arithmetic on a postulated relation. Whether any observer is described by `M = (D/P).E` is a separate question; it is a postulate rather than a result; the falsification and reproducibility gates have not been run; and **the equation's own scale-type demotion is this paper's headline finding and its principal negative result.**

19.16.1 Declarations

Data and code availability

No external dataset was collected for this paper, no observer was measured, and no organisation was studied. All numerical results are computed from the stated parameters and stipulated counts under the relations developed in the text, and every table can be re-derived from the equations and values printed beside it. **The seeded review audit of Section 12.1 and the fourteen-year inspector series of Section 12.2 are constructed examples with declared counts, not field observations**, and Section 12.3 records the numbers in the source material for those examples that do not reconcile. The figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper and the relations

of Sections 4 to 8.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, produced the worked examples and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. Both worked audits are constructed. The consent and decline requirements stated as SP-13 and SP-28 — that no calibration procedure is administered to a person who has not agreed to it, and that a refusal is recorded as a refusal and not as a result — are design constraints on prospective work and are not reports of any study performed. The framework is explicit that calibration measurement is measurement of a person and that results may not be used for a purpose not disclosed before the audit.

Preregistration

This paper reports no preregistered test of its own; its results are derivations, and its principal finding is a self-imposed demotion rather than an empirical outcome. The preregistration protocol that would test the moment equation — fit, freeze, held-out prediction, and four baselines including a shuffled-input control — is specified in Section 9.2 and **has not been run.** The corresponding protocol applied to a different law elsewhere in this corpus **failed its second falsification rule**, returning a 2.8 per cent held-out advantage over a matched autoregression against a declared margin of 10 per cent [Curzi, 2026i]. Any applied claim resting on this paper inherits that result.

19.17 Appendix A The demoted and the surviving, side by side

The table below is the operational summary of Section 3.5 and is provided so that a practitioner can tell at a glance which statements this framework licenses.

Statement	Status under the demotion
This observer is over-framed / calibrated / under-framed	**SURVIVES.** Needs only the sign of `ln r`.
This observer moved toward contact between two audits on a common item set	**SURVIVES.** Needs only rank order.
Observer A is further from contact than observer B on the same item set	**SURVIVES.** Ordinal comparison.
As intake falls with frame held, the moment inflates	**SURVIVES.** A limiting case.

Statement	Status under the demotion
With no frame the moment is zero, whatever the intake	**SURVIVES.** A limiting case.
With an empty environment the moment is zero, whatever the ratio	**SURVIVES.** A limiting case.
At `D = P` the moment equals the environment for every environment	**SURVIVES.** An identity in the model.
The moment is bounded above by `ln N`	**SURVIVES**, conditional on the declared normalisation (SP-11).
Correction converges for gain in `(0,2)` and diverges at or above 2	**SURVIVES.** A property of a chosen design parameter.
Veridical throughput is maximised at the fixed point	**SURVIVES.** A property of `min(r,1)`.
This observer's `r` is 1.72	**DEMOTED.** Magnitude withdrawn.
This observer's distortion is 0.54 nats	**DEMOTED.** Magnitude withdrawn.
This observer's distortion is twice that one's	**DEMOTED.** Requires a ratio scale.
This observer's distortion fell by 92.6 per cent	**DEMOTED** for `r` and `abs(ln r)`; **retained** for the false-discovery rate, which is a ratio of counts.
Sixty-three per cent of this observer's store is error	**DEMOTED.** Inherits through `M`.
The pipeline injected 1.07 nats in total	**DEMOTED as arithmetic**; retained as the structural claim that compensating stages exist and a net figure conceals them.
The confabulated fraction is `1 − 29/50 = 0.42`	**SURVIVES.** Reduces to two counts with a true zero.
This audit found 31 of 40 seeded defects, up from 29	**SURVIVES.** Counts.
Time on site fell from 232 to 41 minutes	**SURVIVES.** Duration, true zero.

19.18 Appendix B The route out, and what it costs

The instrument that would restore the magnitudes is specified rather than gestured at, and the specification is what makes the demotion a finding rather than a complaint.

The construction. Measure `D` and `P` as person-measures on **one common linked scale** under a Rasch model [Rasch, 1960, Andrich, 1988], given a badged additivity postulate. Rasch separability supplies the true zero that division requires and fixes the unit up to choice, and the distortion becomes a difference of logits on that scale, `Delta = |ln(D/P)| = |theta_D − theta_P|`, which is genuinely ratio-scaled.

The price, in items. About **123 items per term** to establish a half-logit distortion at 95 per cent, and about **31 items** for a full logit [Curzi, 2026t]. A twelve-item pair of scales resolves nothing below roughly 1.6 logits, so **a practitioner reporting `r = 1.2` from a twelve-item pair is reporting noise.**

What it does not buy. It does not answer Michell [1997], because the Rasch model assumes the additive structure it then exploits rather than testing for it; the test for additive conjoint structure is a separate procedure [Luce and Tukey, 1964, Krantz et al., 1971] and it has not been run. It does not identify the split of `r` into `D` and `P`, which is a separate and unrelated failure (Section 4.1). It does not address the reflection problem [Manski, 1993]. **And it does not make the framework empirical**, because gates G6 and G7 remain unrun whatever scale the terms are measured on.

What it does buy. Everything in the demoted column of Appendix A, and with it the framework's two candidate differential predictions — the exact `1 - 1/r` relation between overframing and false discovery, and the knife edge at `kappa = delta` — which are the only claims here that a simpler account of expertise decay does not also make [Levinthal and March, 1993, Durlauf, 2005].

Chapter 20

The Perceptual Lens: Fechner's Law as the Register of the Surplus Functional, with a Postulate Method and Its Held-Out Tests

Abstract. The Lens Theorem states that if percepts of multiplied intensities add, the transfer function must be logarithmic; the inference is Cauchy's functional equation and is verified numerically to an absolute error of 1.02×10^{-8}. The substantive result is that this corpus's surplus functional is a Fechner law: $f(u) = \ln(1 + (N-1)u)$ rewrites exactly as $\ln((u+u_0)/u_0)$ with $u_0 = 1/(N-1)$ an absolute threshold of difference, so concavity is a property of the register rather than a pessimistic assumption, and $f'(0)/f'(1) = N$ exactly. The storage law in world coordinates is correspondingly a geometric accumulator, $G_{t+1} = G_t^{1-\delta}\, N_{\text{eff}}^{\eta}\, e^{-C}$, agreeing with the additive form to fourteen significant figures. The design consequence is a prohibition: never promise a linear improvement in a logarithmic quantity. At a brightness exponent of 0.33 a promised twenty per cent delivers 6.201 per cent perceived, 31 per cent of the promise, and a real twenty per cent costs a 73.76 per cent load increase; a reading-room case closes a perceived zone gap from 26.09 to 9.06 per cent for 8.00 per cent more energy. A postulate method defines empirical content as $H_{\text{forbid}} = -\ln q$, and the worked postulate was refuted on one of four declared consequences at three times its pre-set tolerance and amended in public. Two corrections are carried throughout: the corpus's logarithm needs multiplicative dynamics rather than Fechner, and the paper's own ratio-invariance test — response ratios of 2.5100 and 2.5000 agreeing to 0.40 per cent — is the power law's invariant, favouring Stevens over Fechner by $\Delta\text{AIC} = 26.57$ at equal parameter count.

Keywords: Lens Theorem; Fechner's law; Stevens exponent; geometric accumulator; ratio invariance; magnitude estimation; postulate method; empirical content

Series: B. Engineering the Surplus **Paper:** CCP-020

20.1 1 Introduction

The Creature Comforts framework computes in logarithms. Its surplus functional is `F(u) = ln(1 + (N–1)u)`; its storage law accumulates that logarithm; its distortion measure is `|ln r|`; its ceiling is `ln N`. A framework that takes a logarithm everywhere owes an account of why,

and the account it has so far offered is psychophysical: perception reports ratios rather than differences, and a device that converts ratios into differences is a logarithm.

This paper supplies that account, states exactly how much of it the framework actually needs, and reports the test that cuts against it.

Three things are established. **First**, the Lens Theorem: if the map from a world quantity to its perceived magnitude is continuous and carries multiplication into addition, then that map is a logarithm and is the only such map up to one scale constant. That half is algebra — the continuous solution of Cauchy's logarithmic functional equation — and it is verified here numerically by a 200,000-rectangle midpoint sum against `ln 100`, absolute error `1.02e-08`. **Second**, the framework's own surplus functional is a Fechner law: setting `u_0 = 1/(N-1)`, `ln(1 + (N-1)u)` rewrites exactly as `ln((u + u_0)/u_0)`, which is Fechner's near-threshold form with non-parallelism as the stimulus and `u_0` as an absolute threshold of difference. The identity is exact to floating-point noise and is independently verified across `N = 2` to `100` at seven values of `u` to a maximum deviation of `8.9e-16` in Curzi [2026]. Its consequences are not decorative: concavity becomes a property of the register rather than a pessimistic assumption, the ratio of first to last marginal value is exactly `N`, and the storage law read in world coordinates is a weighted geometric accumulator. **Third**, a postulate method — the discipline by which an unobservable quantity is given a structure, its observable consequences derived, and the structure held to them — is stated with six admissibility conditions, an empirical-content measure, and a held-out testing procedure, and is then exercised on a worked postulate that **failed**.

Two corrections to the framework's own reading of this material are carried in the body, and they matter more than the confirmations.

The logarithm does not need Fechner. Peters and Gell-Mann [2016] derive a logarithmic form from the dynamics of a multiplicative process with no psychological premise at all: the time-average growth rate of such a process is logarithmic in the accumulating quantity, and nothing about observers is assumed. That is a **stronger** warrant for the corpus's logarithms than Fechner's, precisely because it does not require the perceptual story to be true, and it is not exposed to the century of measurement-theoretic objection that Fechner's is [Stevens, 1957, Poulton, 1979, Luce, 2002, Narens, 1996]. Section 3.3 states the division of labour explicitly and Sections 13 and 15 restate it: **where the framework needs only the functional form, Peters and Gell-Mann are sufficient and Fechner is optional colour; where the framework needs the perceptual claim — which it does, in exactly one place, the design consequences about promised versus perceived magnitude — Fechner and Stevens are load-bearing and their unresolved contest is a live cost.**

The functional form is underdetermined. Latané [1981] proposes a sub-linear power law in the number of sources; Briscoe et al. [2006] propose `n log n`; this framework proposes `ln(1 + (N-1)u)`. All three are increasing, concave in the relevant range, and consistent with the qualitative facts. Nothing assembled in this paper discriminates among them, and Section 3.5 says so with a figure. The one genuinely discriminating test the corpus possesses is the ratio-invariance test of Section 12.3 — and **it returns a power-law exponent, which is Stevens rather than Fechner.** Section 12.3 reports that plainly, and Section 13.2 draws the consequence: a ratio-invariance test that succeeds is evidence for a power law and against a Fechnerian logarithm on the same data. The corpus keeps its logarithm on axiomatic [Curzi, 2026] and dynamical [Peters and Gell-Mann, 2016] grounds, not psychophysical ones.

Sections 2 and 3 state the hypotheses and the framework. Section 4 develops the translation table between operations in the perceptual register and operations in the world. Sections 5 and 6 apply it to the framework's own two central equations, the storage law and the surplus functional. Section 7 draws the design consequences and prices them. Section 8 states the postulate method. Section 9 presents the psychophysical record on which the method's exemplar rests. Sections 10 and 11 collect the specifications and the named failure modes. Section 12 reports three worked cases, one of which refuted its own postulate. Sections 13 to 16 discuss, verify, limit and conclude.

20.2 2 Hypotheses

The Lens Theorem of Section 3.1 is a theorem about a functional equation and cannot be falsified by observation; it is labelled a derivation where it appears, and its **antecedent** — that perception is such a map — is the empirical claim. The hypotheses below separate the two. Each names the observation that would defeat it.

H1 — The register is ratio-based. Equal ratios of a physical stimulus produce equal increments of subjective magnitude, over the mid-range of a modality, rather than equal absolute differences producing equal increments. **Predicted:** a constant Weber fraction `c = dI/I` over the mid-range, so that the count of just-noticeable differences between two stimuli depends only on their ratio; and, equivalently, a straight line when a well-constructed magnitude scale is plotted against `ln I` or against `log psi` versus `log I`. **Falsified by:** a modality showing a constant *absolute* just-noticeable difference over its mid-range, or `c` varying systematically with `I` in the mid-range beyond measurement error. **Status of the evidence:** run, and passing in the middle of the range and failing at both ends in a characteristic and reproducible way [Weber, 1834, Stevens, 1975].

H2 — Subjective magnitude admits a ratio scale under a ratio instruction. Numbers emitted under an unbounded proportional-assignment instruction have the internal structure of ratio measurement, and numbers emitted under a bounded category instruction do not. **Predicted:** for equal physical ratios presented at different absolute levels, the response ratios agree within measurement error; and a bounded rating scale does not reproduce that invariance because its ceiling forbids it. **Falsified by:** response ratios that depend on absolute level; or a bounded rating scale that reproduces the invariance as well as the unbounded instruction does. **Not established even if it passes:** the conjoint-measurement conditions that would license the quantitative treatment in the first place [Luce and Tukey, 1964, Michell, 1997]. Section 15 records that they were not tested.

H3 — The surplus functional is the framework's Fechner law. The corpus's `F(u)` is not merely logarithmic in form but is the threshold form of Fechner's law with `u` as the stimulus. **Predicted:** `ln(1 + (N−1)u) = ln((u + u_0)/u_0)` identically on `u` in `[0,1]` for every `N >= 2`, with `u_0 = 1/(N−1)`; hence `N_eff(u_0) = 2` and `F(u_0) = ln 2` for every `N`, and `f'(0)/f'(1) = N` exactly. **Falsified by:** any `(N, u)` pair at which the two expressions differ by more than floating-point rounding. **This hypothesis is an algebraic identity and is therefore a derivation, not an empirical claim** — it is listed here because it is the paper's substantive result and a reader is entitled to see it stated in falsifiable form and then to check it.

H4 — The linear-promise prohibition. A commitment stated as a percentage change in a world quantity systematically overstates the perceived change delivered, in every compressive modality, by the modality's exponent. **Predicted:** at a brightness exponent of `0.33`, a delivered `+20%` illuminance yields `1.20^0.33 − 1 = +6.201%` perceived, which is `31.0%` of the promise; and delivering a genuine perceived `+20%` requires `1.20^(1/0.33) − 1 = +73.76%` of world quantity. **Falsified by:** post-occupancy magnitude estimation returning a perceived change matching the physical percentage rather than the exponent-converted one, in a modality whose exponent is below one and has been measured under the application's own conditions. **Depends on:** the perceptual claim, not merely the functional form. **This is the one hypothesis in the paper for which Fechner and Stevens are load-bearing and Peters and Gell-Mann are not substitutes**, because the quantity being predicted is what a person reports, not how a stock grows.

H5 — Empirical content is what a postulate forbids, and it is measurable comparatively. Over a fixed, published outcome list, the content of a postulate that admits a

fraction `q` of the listed outcomes is `H_forbid = −ln q` nats, and postulates ranked by this quantity are ranked in the order that repays a fixed testing budget. **Predicted:** `q = 1` buys zero nats; halving the admitted outcomes buys `ln 2 = 0.6931` nats; content adds under independent conjunction, `−ln(q_1 q_2) = H_1 + H_2`. **Falsified by:** a preregistered set of postulates with declared content scores, tested under a fixed budget, in which the high-content postulates do **not** produce more parameter revision per unit cost than the low-content ones. **Status:** the arithmetic is a **definition**; the claim that it helps is an untested hypothesis, and Section 8.6 badges it accordingly.

H6 — Held-out prediction with a shuffled-input control separates storage from shape. A model whose predictive advantage comes from the postulated storage mechanism loses that advantage when its input series is permuted; a model that has merely learned the series' shape does not. **Predicted:** with frozen parameters, held-out error under a shuffled input rises materially and the skill score against persistence goes negative. **Falsified by:** shuffled-input held-out error close to the unshuffled figure while the other baselines are beaten — the diagnostic signature of a model that is not using its inputs. **Precedent, and a caution:** the corpus's one preregistered confirmatory application of this protocol beat its shuffle at `p = 0.0005` and nonetheless **failed** its own second falsification rule, at a 2.8 per cent held-out advantage over a matched AR(1) against a declared 10 per cent margin [Curzi, 2026i].

Three further claims in this paper are **not** hypotheses and are labelled where they appear. The **Postulate Method** of Section 8.2 is a house formalisation of ordinary hypothetico-deductive practice [Popper, 1959, Platt, 1964]; no priority is claimed for it and it is not the kind of thing that is true or false. The **badge taxonomy** of Section 8.3 is a **definition**, a naming convention for evidential status, parallel to the four-drawer classification of Curzi [2026]. The **identification of the storage law's world-coordinate form with a geometric accumulator** (Section 5) is a **change of coordinates** — a shared mathematical form in the sense of Section 14.7, exactly equivalent to the additive form, and therefore evidence for nothing.

20.3 3 Framework: the lens, its two warrants, and what it does not settle

20.3.1 3.1 The Lens Theorem

The Lens Theorem [conditional theorem]. If the map from a world quantity `I > 0` to its perceived magnitude `S` is continuous and converts multiplication into addition, that is if

```
S(I_1 . I_2) = S(I_1) + S(I_2)      for all I_1, I_2 > 0
```

then that map is a logarithm, and it is the only such map up to a single positive scale constant:

```
S(I) = k ln I ,   k > 0
```

equivalently, with a reference intensity `I_0` at which the percept is zero,

```
S(I) = k ln(I / I_0) .
```

The inference is a theorem and it is not new. It is the continuous solution of Cauchy's logarithmic functional equation, and it is the statement that the logarithm is the unique continuous isomorphism from the multiplicative group of the positive reals onto the additive group

of the reals. No psychology enters the proof; it is an algebraic fact about two operations. The badge is therefore on the **antecedent** — the claim that perception is such a map — and not on the inference. Section 9 gives the record of the antecedent and Section 9.9 gives its limits.

The direction a designer needs is the inverse:

```
I = I_0 . e^{S/k}
```

A linear schedule of percepts requires an exponential schedule of world quantities. Equal experienced steps demand equal ratios, and equal ratios compound. Every design consequence in Section 7 is a worked instance of that one line.

20.3.2 3.2 Weber's law and Fechner's integration

Weber [1834] reported that the smallest detectable increase in a stimulus is a constant proportion of the stimulus already present:

```
dI / I = c ,     c the Weber fraction, constant over the mid-range of a modality
```

Weber's law is on the observable side of the wall: it concerns discrimination performance, settled by a forced choice and a tally, and asserts nothing about inner magnitude. Fechner [1860] took the step that makes it a lens by postulating that every just-noticeable difference is one equal unit of subjective magnitude, and integrating:

```
dS = k dI / I  =>  S = k ln I + const  =>  S = k ln(I / I_0)
```

Fechner's equal-JND postulate is a postulate, badged as such wherever it is used in this paper. It does not follow from Weber's law; it is an additional assumption; it has been disputed since it was made; and Fechner's law is what follows if it is granted. The standard English text for quotation is Fechner [1966]; Weber [1978] is the translation of Weber.

The integration step is verified numerically. With `k = 1` and `I_0 = 1`, integrating `1/I` from `I = 1` to `I = 100` by 200,000 midpoint rectangles:

```
midpoint Riemann sum = 4.605170176
ln(100)              = 4.605170186
absolute error       = 1.021e-08
```

The Weber fraction fixes how many perceptual steps a physical range contains, and the counts are small:

Weber fraction `c`	JNDs across a 2:1 range	JNDs across a 1000:1 range
0.02	35.00	348.83
0.05	14.21	141.58
0.10	7.27	72.48
0.20	3.80	37.89

A thousandfold range of world quantity contains 72 discriminable steps at a Weber fraction of 0.10, and a doubling contains seven. The world's dynamic range is large and the register's is small. That compression is what the register is for, and the observation that people cannot distinguish a three per cent improvement from nothing is a statement about working sense organs rather than about inattention.

20.3.3 3.3 The two warrants for the logarithm, and which one this framework needs

This is the section the rest of the paper's honesty depends on, and it is placed before any use of the logarithm rather than after.

Warrant 1 — psychological (Fechner). Sensation is logarithmic in intensity because equal just-noticeable differences are equal steps of sensation and just-noticeable differences are constant ratios of intensity [Weber, 1834, Fechner, 1860]. This warrant requires the perceptual story to be true. It is **contested**: Stevens [1957] produced a competing power law by a method Fechner had not considered, and no logarithmic sensation scale can produce the expansive exponents that method returns (Section 9.5); Poulton [1968] and Poulton [1979] catalogue the biases in the competing method; Luce [2002] and Narens [1996] ask axiomatically what must be true of a subject's judgements for either scale to exist at all; Krueger [1989] is the standing reconciliation attempt, with peer commentary, and the dispute is not closed.

Warrant 2 — dynamical (Peters and Gell-Mann). The time-average growth rate of a multiplicative process is logarithmic in the accumulating quantity [Peters and Gell-Mann, 2016]. **No psychological assumption is made anywhere in that derivation.** The companion statement — that the time average of a multiplicative process differs from its ensemble average, so that a parameter estimated across individuals does not describe what happens to any one of them over time — is Peters [2019], and it bears directly on this framework, whose stored quantity accumulates over time for one system.

Which warrant does the argument need? The answer differs by claim, and the paper states it claim by claim rather than once.

- **For the functional form of the surplus law and of the storage law, Warrant 2 is sufficient and Warrant 1 is optional.** `F(u) = ln(1 + (N−1)u)` is characterised in this corpus by four axioms about interaction geometry, with no reference to perception [Curzi, 2026]; the storage law's world-coordinate form (Section 5) is multiplicative, which is exactly the class of process for which Peters and Gell-Mann [2016] derive a logarithm. **Neither claim is exposed if Fechner is wrong.** The Fechner identity of Section 6 is then a structural convergence between two independent routes to one function, and Section 6.1 says explicitly that a convergence is not evidence that either route is correct.
- **For the design consequences about promised versus perceived magnitude, Warrant 1 is load-bearing and Warrant 2 is not a substitute.** Section 7.2 predicts what a person will report after an illuminance change. That prediction is about perception; a dynamical derivation of a logarithm says nothing about it; and the applicable constant is a measured Stevens exponent whose value, method-dependence and between-study variation are the live cost of using it (Sections 9.3 and 9.9). **The title of this paper claims more than most of its content requires, and that is said here rather than left for a reviewer to notice.**

The honest summary, restated in Section 13.1 and again in Section 15: **the corpus needs a logarithm; Peters and Gell-Mann supply one without needing observers; Fechner supplies one that does need them and is contested; the corpus's central formal claims rest on the first, and only its design advice rests on the second.**

20.3.4 3.4 Stevens' power law, and why it relocates the register rather than abolishing it

Stevens [1957] proposed

```
psi = k I^a                                  the power law
```

obtained by **magnitude estimation** — asking observers to assign numbers in proportion to their sensations. The competition with Fechner is real and unresolved. It does not remove the register, and the reason is one line of algebra:

```
psi = k I^a   =>   ln psi = ln k + a ln I
```

A power law is a straight line in log-log coordinates: it is the logarithmic register with a gain. Under Fechner the operative case is `log in, linear out`; under Stevens it is `log in, log out`, with the world-side ratio raised to `a`:

```
perceived ratio = (world ratio)^a ,      perceived change in nats = a ln(world ratio)
```

The exponent `a` is therefore a **measured conversion key** between the physical and perceptual registers, and holding both theories at once amounts to saying: the register is logarithmic, and `a` is its gain. Published exponents, presented as **nominal, modality- and method-dependent constants of a procedure and not constants of nature** [Stevens, 1975, Teghtsoonian, 1971]:

Modality	exponent `a`	character
Loudness (1 kHz, sound pressure)	0.67	compressive
Brightness	0.33	compressive
Visual length	1.00	linear
Taste (salt)	1.30	expansive
Heaviness (lifted)	1.45	expansive
Warmth on arm	1.60	expansive
Electric shock	3.50	strongly expansive

`a < 1` compressive, `a = 1` linear, `a > 1` expansive. **Most senses compress; pain expands.** Section 7.6 shows that the expansive modalities invert every design rule in Section 7, which is why an exponent must be looked up rather than assumed.

These are cited point estimates, not measurements made here, and they carry no interval. The literature is explicit that they shift with stimulus range [Teghtsoonian, 1971], with instruction and modulus [Poulton, 1979], and across observers, replicating in rank order and to within roughly ten to twenty per cent for well-studied modalities [Stevens, 1975]. Section 15 records that every design number in Section 7 inherits that variation, and Section 10 requires re-measurement before any exponent is used to calibrate rather than to orient.

20.3.5 3.5 The functional form is underdetermined, and this paper does not settle it

Three candidate laws relate an aggregate to a count of contributing sources, and all three are increasing and concave over the range in which this framework operates.

Candidate	Form	Source
Sub-linear power law	influence grows as `N^t` with `t < 1`	Latané [1981]
`n log n`	network value grows as `N ln N`, not `N^2`	Briscoe et al. [2006]
This framework	`F(u) = ln(1 + (N−1)u)`	Curzi [2026]

Nothing in this corpus discriminates among them on data, and the shape comparison this paper can perform is reported here with its arithmetic. Affinely matching each alternative to `ln N` at both endpoints of `N` in `[2, 50]` — two constraints, so no free parameter is left over except the power law's exponent — gives:

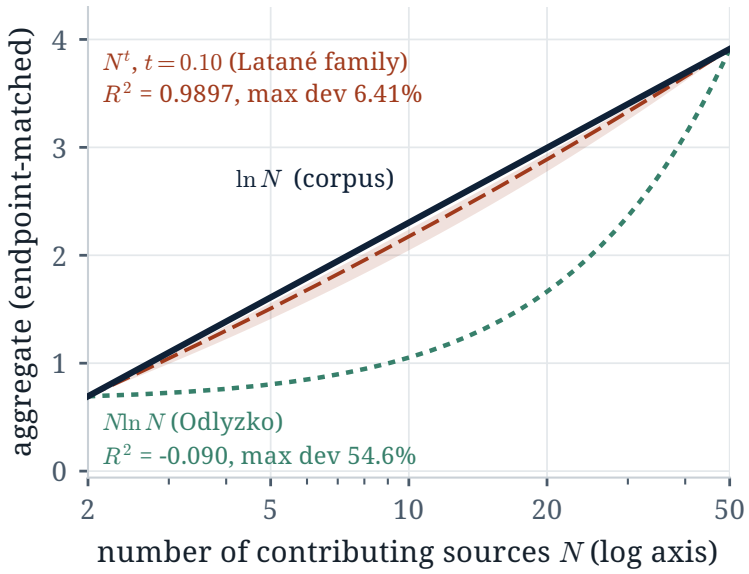


Figure 20.1: **A shape argument cannot separate the corpus's logarithm from a sub-linear power law, and can separate it from `n log n`.** Each alternative is affinely matched to `ln N` at both endpoints of `N` in `[2, 50]`. A power law with `t = 0.10` tracks `ln N` at `R^2` = 0.9897 with a maximum relative deviation of 6.41 per cent, and the band shows `t` from 0.05 to 0.20; the logarithm is the `t -> 0` member of that family, so no exponent value is invented here and none is published in the source. `N ln N` departs by 54.60 per cent and is a claim about a different quantity. Functional forms only, with two endpoint constraints and no fitted parameters; generated by `figs/src/CCP-020_fig01.py`.

alternative, endpoint-matched to `ln N` on `N` in `[2,50]`	`R^2` against `ln N`	max relative deviation
`N^t, t = 0.05`	0.9974	3.24%
`N^t, t = 0.10`	0.9897	6.41%
`N^t, t = 0.20`	0.9593	12.49%
`N ln N`	-0.090	54.60%

The sub-linear power family is not discriminated from the logarithm at small exponents, and it cannot be, because the logarithm *is* its limiting member: `(N^t - 1)/t -> ln N` as `t -> 0`, so the corpus's law sits on the boundary of the family Latané [1981] proposes and the source supplies the constraint `t < 1` rather than a value. **A shape argument cannot separate them and no shape argument in this corpus has tried.** `N ln N` does differ in shape over the same range — which is a genuine finding and a limited one, because Briscoe et al. [2006]'s quantity is the *total value of a network of `n` members* and the corpus's is a *per-ensemble surplus*, so the two are not measurements of the same thing and a shape comparison between them is suggestive at best. **What the table shows is the ceiling on what shape alone can settle: a range of a decade and a half separates `N ln N` from `ln N` and does not begin to separate `ln N` from `N^{0.05}`.**

The corpus adopts the logarithm because it is uniquely characterised by four stated axioms, one of which — that the effective count be affine in `u` — is a **design axiom, chosen and not derived** [Curzi, 2026], and because a multiplicative dynamics supplies a logarithm independently [Peters and Gell-Mann, 2016]. **It is not adopted because it fits better than the alternatives, and no comparison in this corpus has shown that it does.**

Two further cautions belong here. Briscoe et al. [2006]'s actual point is that a network-value functional form was adopted without evidence and then used to justify valuations; a

framework quoting them in support must not repeat the error they are describing. And where any claim of a power law or heavy tail is made anywhere in this corpus, the distribution-fitting procedure of Clauset et al. [2009] is binding — no such claim is made in this paper, and Section 3.4's exponents are cited point estimates from magnitude-estimation studies, not tail exponents estimated from a sample.

20.3.6 3.6 Which effective count, and the Hill-number caution

Two different effective counts appear in this corpus and they must not be conflated. `1/HHI` is the Hill number of order two and `exp(H)` is the Hill number of order one; both are fully grounded [Hill, 1973, Jost, 2006], and the concentration instrument that uses `1/HHI` — naive HHI `0.293430` giving 3.41 effective jurisdictions against a bloc-HHI of `1.000000` and an effective count of `1.00` — is Curzi [2026]. **`N_eff = 1 + (N−1)u` is not a Hill number of any order.** It is a linear interpolation between 1 and `N`, it is this framework's own construction under a declared design axiom, and no literature in ecology, economics, political science or information theory contains it. **Every occurrence of an effective count in this paper is `1 + (N−1)u` unless it says otherwise**, including the `N_eff` in the geometric accumulator of Section 5 and in result R3, and a reviewer reaching for the Hill-number literature as support for it will find that the literature does not apply.

20.3.7 3.7 Identification, reflexivity, and what this paper does not claim to have measured

The material in Section 8 concerns postulates an observer makes about their own interior, and Section 8.9 observes that such postulates behave like stored order — they decay, they convert experience into structure, they cost something to hold. Where that observation is extended to a group whose members both influence and constitute an aggregate, **the specification becomes reflexive and its endogenous effect is not identified from observational data.** Manski [1993] proves this for the linear-in-means case, and Shalizi and Thomas [2011] show that influence and selection are generically confounded in observational network data. **No amount of data, held out or otherwise, recovers the parameter.**

This paper adopts no remedy. The known options are randomisation, an instrument, partial-population variation, or stipulating the reference group and defending the stipulation [Brock and Durlauf, 2001]; none is free, and none is exercised here. Every claim in Sections 8 and 12.2 about a structure inferred from a pattern of self-reports is therefore a **modelling position rather than an identified empirical result**, and Section 15 repeats it.

Two related points are stated once and not repeated. The existence proof that a description can become true by being used is MacKenzie and Millo [2003] — Black–Scholes prices did not fit the market when published, and after traders adopted the model the market came to fit the model, with dates and price data. Hacking [1995] is the general form: classifications of people interact with the people classified, so the referent moves. Against these, Jussim and Harber [2005] review thirty-five years of work after Rosenthal and Jacobson [1968] and find that classroom self-fulfilling prophecies are **real but small**, do not typically accumulate, and may dissipate; McCambridge et al. [2014] find the same for participation effects generally. **The reflexive mechanism exists and its measured effect sizes are much smaller than a framework built on it would like.** Merton [1948] named it; Lucas [1976] is the version that applies to this framework's own parameters, and Section 15 answers it narrowly rather than pretending it does not apply.

20.3.8 3.8 Standing honesty clauses

1. **The Lens Theorem's mathematics is a theorem; its applicability is an empirical claim.** The second is not to be presented with the confidence of the first.

2. **Fechner's equal-JND postulate is a postulate.** It does not follow from Weber's law.
3. **Stevens exponents are constants of a procedure.** Method, range, modulus and instruction are reported with any exponent, or the number is not usable.
4. **Percepts are interval-scaled, not ratio-scaled, under Fechner.** They may be subtracted and may not be divided. Under magnitude estimation both are licensed, and that is the methodological difference between the two accounts that matters most here.
5. **Nats of one modality are not nats of another.** Section 10, S9 is explicit, and this is the most likely way for the material in Section 7 to be misused.
6. **This paper concerns magnitude judgement.** It is not a theory of cognition, preference or value, and extending it to those requires separate evidence which is not offered.
7. **Nothing here is a claim about consciousness, physics or thermodynamics.** The entropy that appears anywhere near `ln N` is Shannon entropy [Shannon, 1948], a dimensionless information measure in nats, and not thermodynamic entropy.

20.4 4 The translation table: one operation, two registers

Every arithmetic operation performed on a perceived quantity has a counterpart on the world quantity that produces it. The table below is the paper's working instrument; the subsections that follow take the rows that carry weight for the framework.

```
LENS SIDE  = what the observer experiences,  S = k ln(I/I_0)
WORLD SIDE = what has to be built,           I
```

Operation in the perceptual register	Corresponding operation in the world
`S_a + S_b` — addition	**Multiplication.** `I_a . I_b / I_0`
`S_a - S_b` — subtraction	**Division.** The ratio `I_a / I_b`
`0` — the origin	**The reference `I_0`. Not absence.**
a negative percept	`I < I_0`: a reciprocal ratio, below reference
`c . S` — scalar multiplication	**Exponentiation.** `I^c`, relative to reference
`S / c` — scalar division	**The c-th root.** `I^(1/c)`
`2S` — doubling the percept	**Squaring the world quantity**
arithmetic mean of percepts	**Geometric mean** of world quantities
midpoint between two percepts	**Geometric midpoint**, `sqrt(I_a . I_b)`
linear interpolation	**Geometric interpolation**
arithmetic progression, equal steps	**Geometric progression, equal ratios**
sum over `n` terms	**Product over `n` terms**
`abs(S_a - S_b)` — distance	**`abs(ln(I_a/I_b))` — the log-ratio metric**
variance, normal distribution	**Log-variance, lognormal distribution**
`dS/dt` — a rate of change	**A proportional growth rate**, `(dI/dt)/I`
`dS/dx` — a derivative	**An elasticity**, `(dI/I)/(dx/x)`
a straight line, rising	**Exponential growth**
a straight line in log-log	**A power law**
a concave curve in world units	**A straight line, seen through the lens**
the integral of `S` over time	**The geometric (product) integral of `I`**
unbounded below, `S -> -infinity`	**`I -> 0`: the missing floor, Section 6.2**

20.4.1 4.1 Addition, and the octave constant

```
S(I_a) + S(I_b) = k ln(I_a/I_0) + k ln(I_b/I_0) = k ln( I_a I_b / I_0^2 )
```

Two percepts added correspond to two world quantities multiplied, with the reference divided out once. In the base-ten form people use when they say “orders of magnitude”, `+1` is `x10`, `+2` is `x100`, `+6` is `x1 000 000`. In nats:

step (nats)	world multiple
+0.250000	x1.284025
+0.500000	x1.648721
+0.693147	**x2.000000**
+1.000000	x2.718282
+2.000000	x7.389056
+3.000000	x20.085537
+4.605170	x100.000000

One octave — a halving or a doubling — is `ln 2 = 0.693147` nats, which is `0.301030` decades. That single constant is the conversion between the language of ratios and the language of steps, and it recurs across this corpus: as the boundary between marked and severe distortion in the calibration instrument [Curzi, 2026m], as the pitch octave, and as the surplus bought by doubling `N` [Curzi, 2026]. It is one constant doing one job in several places because there is one register.

20.4.2 4.2 The mean: an average of percepts is a geometric mean of quantities

```
mean of percepts:   (S_a + S_b)/2 = k ln( sqrt(I_a . I_b) / I_0 )
```

The world quantity corresponding to the average percept is the geometric mean.

between	arithmetic midpoint	**perceptual midpoint**
1 and 100	50.5000	**10.0000**
10 and 1 000	505.0000	**100.0000**
1 and 2	1.5000	**1.4142**
300 and 2 400 lux	1 350.0000	**848.5281**
20 Hz and 20 000 Hz	10 010.0000	**632.4555**

The quantity halfway between 1 and 100 in experience is 10, not 50.5. The frequency halfway between the bottom and the top of human hearing is 632 Hz, which is near the middle of a piano keyboard — an instrument labelled in octaves by people who have been working in this register for four centuries.

The corollary for rates:

```
+50% then −50%:      1.5 x 0.5 = 0.75              net −25.00%
arithmetic mean of the two changes:                 0.00%      WRONG
geometric mean factor: sqrt(0.75) = 0.866025       −13.40% per step   CORRECT
mean of the logs: (ln 1.5 + ln 0.5)/2 = −0.143841 nats -> exp = 0.866025
```

A four-period example: factors `1.10, 1.10, 0.80, 1.25` give a net of `1.210000`, a geometric mean of `1.048809` per period, and an arithmetic mean of factors of `1.062500`, which compounded over four periods predicts `1.2744` against the true `1.2100`. **The arithmetic mean of ratios always overstates, by Jensen’s inequality, and always in the same direction.** Section 6.4 shows that the same inequality, applied to a concave register rather than to a convex compounding, produces a penalty in the opposite direction and resolves a discrepancy in this corpus’s own reported figures.

20.4.3 4.3 The metric: distance is a log-ratio

```
perceptual distance between I_a and I_b  =  abs(ln I_a − ln I_b) = abs(ln(I_a/I_b))
```

a	b	difference	ratio	perceptual distance (nats)
1	2	1	2.0000	**0.693147**
50	100	50	2.0000	**0.693147**
1 000	2 000	1 000	2.0000	**0.693147**
1	1.1	0.1	1.1000	**0.095310**
100	110	10	1.1000	**0.095310**

Equal perceptual distance, with the underlying differences varying by a factor of a thousand. This is the row that makes an absolute-difference metric wrong for any perceived quantity, and it is why the corpus's distortion measure is `abs(ln r)` and not `abs(r − 1)` [Curzi, 2026m]. The same choice governs comparisons of very unequal magnitudes elsewhere in the corpus: a five-second settlement hop is `1.585e−07` of a year against `0.0822` for value rested thirty days, a ratio of about 518,400 to one, and that comparison is only readable as a ratio [Curzi, 2026].

20.4.4 4.4 The derivative: a slope in the lens is an elasticity in the world

```
d(ln psi) / d(ln I) = (I/psi) . (dpsi/dI) = elasticity
```

For a power law `psi = k I^a` this evaluates to exactly `a`. **The Stevens exponent is the elasticity of percept with respect to stimulus — the percentage change in experience per percentage change in the world — and it is constant across the range.** That is what "the log-log plot is a straight line" means, restated for use. The corpus's moment equation yields unit elasticities in all three of its terms by the same operation, `dM/M = dD/D − dP/P + dE/E` [Curzi, 2026m], which is a consistency check on the register rather than an independent result.

20.4.5 4.5 Zero in the lens is the reference in the world

`S = 0` at `I = I_0`, and `I_0` is not zero. **The origin of the perceptual coordinate is the reference level, not the absence of the stimulus.** Three consequences are practical:

- **A percept of zero means "at reference", not "nothing there".** Silence in a room is not zero sound; it is sound below threshold.
- **Percepts may be subtracted and may not be divided** under Fechner. "Twice as bright" is a claim about a ratio-scaled magnitude estimate, not about a Fechnerian percept, and the two are different objects.
- **The reference is a design variable.** It is set by adaptation and by context [Helson, 1948], including what the observer was exposed to immediately before. A designer who does not control the reference is not controlling the percept, and moving the reference is usually cheaper than moving the quantity — which is why a dim vestibule makes the reading room beyond it feel bright. Section 7.5 develops this and Section 7.7 states its ethical boundary.

The adaptation literature also supplies this corpus's strongest empirical ally and its strongest warning: comfort decays toward a set point [Brickman et al., 1978], with the revision that set points move and adaptation is incomplete [Diener et al., 2006].

20.4.6 4.6 Subtraction: the difference between two readings

Subtraction in the register is division in the world. The corpus's return difference, `Delta_R = abs(ln(M'/M))` for a quantity `M` sent and `M'` returned, is that row of the table and required no separate invention; it rests on a badged postulate of commensurability, and `Delta_R = 0` records a null exchange rather than a success [Curzi, 2026q].

sent `M`	returned `M'`	ratio	`Delta_R` (nats)
1.0	1.0	1.0000	**0.000000**
1.0	1.2	1.2000	0.182322
1.0	0.8	0.8000	0.223144
1.0	3.0	3.0000	1.098612
1.0	0.333333	0.3333	1.098612

A threefold amplification and a threefold attenuation carry the same 1.098612 nats, where an absolute-difference measure would have priced them at 2.0 and 0.667. And the first row is the one that is misread: **zero return difference means no information came back**, which is not the same as a good result.

20.4.7 4.7 Variance: symmetric in the lens is skewed in the world

If percepts are normally distributed, world quantities are **lognormally** distributed. With `ln I` normal about `mu = ln 300`:

sigma (nats)	median	mean	mode	mean/median
0.2	300.000	306.060	288.237	1.0202
0.5	300.000	339.945	233.640	1.1331
0.8	300.000	413.138	158.188	1.3771
1.2	300.000	616.330	71.078	**2.0544**

Mean, median and mode are three different numbers and the gap grows with the spread. At `sigma = 1.2` the mean is more than twice the median and nearly nine times the mode, so "the average" of a perceptually distributed quantity reports a value almost nobody experiences, and reports it high.

The symmetric interval is likewise asymmetric in the world. At `sigma = 0.5` about `mu = ln 300`:

```
−1 sd : 181.959      median : 300.000      +1 sd : 494.616
distance below the median : 118.041
distance above the median : 194.616
```

One standard deviation down is 118 units; one up is 195 units — same perceptual distance, different world distance, by a factor of 1.65. Any tolerance band on a perceived quantity written symmetrically in world units is asymmetric in experience. Written as a ratio — `x1.65` or `/1.65` rather than `±195` — it is symmetric in both. This is why the corpus's distortion bands are stated as `abs(ln r)` intervals and then back-translated: the band `0.35 <= abs(ln r) < 0.69` is `r` in `[1.419, 1.994)` **or** `r` in `(0.502, 0.705]`, two intervals of different width in `r` and identical width in the register [Curzi, 2026m].

20.4.8 4.8 Linear in the lens is exponential in reality

If the percept is to rise linearly in time, `S(t) = S_0 + ct`, then `I(t) = I_0 e^{(S_0 + ct)/k}`. Priced at a modest ten per cent per period, worth `ln 1.1 = 0.095310` nats:

period	world multiple	cumulative percept (nats)
0	1.0000	0.00000
3	1.3310	0.28593
6	1.7716	0.57186
9	2.3579	0.85779
12	**3.1384**	1.14372
24	**9.8497**	2.28744
48	**97.0172**	4.57489

Holding a barely noticeable rate of perceived improvement for forty-eight quarterly periods requires the world quantity to grow ninety-seven-fold, for a total advance of 4.57 nats — under two decades of world quantity, and that is the cheap reading. **Steady perceived improvement requires exponentially growing world quantity, indefinitely, and no physical or budgetary system supplies that.** Programmes of continuous perceived improvement therefore stall for a reason that is a fact about the register rather than about the programme, and the available responses are to move the reference (Section 7.5), to move to a modality with headroom, or to stop promising a rate the arithmetic forbids.

The source manuscript prints this last figure as 4.57488; `48 ln 1.1 = 4.5748886`, which rounds to 4.57489. The corrected value is used above and the discrepancy is recorded here rather than silently absorbed.

20.5 5 The storage law in world coordinates is a geometric accumulator

The last structural row of the table is the integral: a sum in the register is a product in the world. The corpus's central recursion is a sum, so it has a world-side reading.

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t                storage law, register coordinates
```

Exponentiate. Write `G_t = e^{Q_t}` for stored order in world coordinates and note that `S_t = ln N_eff,t`, so `e^{S_t} = N_eff,t` exactly, with `N_eff = 1 + (N−1)u` and **not** a Hill number (Section 3.6):

```
G_{t+1} = G_t^{(1−delta)} . (N_eff,t)^{eta} . e^{−C_t}   storage law, world coordinates
```

Verified against direct computation at `delta = 0.08`, `eta = 0.35`, `C = 0.20`, `N = 12`, `u = 0.5` — so `N_eff = 6.5` and `S = 1.871802` — from `Q_0 = 0`:

t	Q_t	G_t = e^{Q_t}	direct e^{Q_{t+1}}	product form	difference
0	0.00000	1.00000	1.57638	1.57638	0.00e+00
1	0.45513	1.57638	2.39612	2.39612	8.88e-16
2	0.87385	2.39612	3.52216	3.52216	4.44e-16
4	1.61348	5.02024	6.95549	6.95549	8.88e-16
8	2.76936	15.94848	20.14478	20.14478	0.00e+00

t	Q_t	G_t = e^{Q_t}	direct e^{Q_{t+1}}	product form	difference

Read against the translation table, the three terms are:

- **`G_t^{(1-delta)}` — decay is taking a root.** The leak does not subtract from the world quantity; it takes the `(1-delta)`-th power of it.
- **`(N_eff)^{eta}` — conversion is exponentiation.** Efficiency is the power to which the effective count is raised before being multiplied in, not a fraction of an amount added.
- **`e^{-C}` — cost is a divisor.** Dissipation divides the world quantity rather than subtracting from it.

The fixed point follows by setting `G_{t+1} = G_t`, giving `G^delta = (N_eff)^{eta} e^{-C}` and

```
G_inf = ( (N_eff)^{eta} . e^{-C} )^{1/delta}
```

which at the parameters above evaluates to **295.6376**, identical to `e^{Q_inf}` with `Q_inf = (eta S - C)/delta = 5.689135`; the two agree to fourteen significant figures (the residual is `2.3e-13` on a value of `295.6376`, a relative agreement of `7.7e-16`). The equivalent statement that the fixed point is positive if and only if the drive exceeds the cost — `Q_inf > 0` iff `eta F(u*) > c*` — is derived in Curzi [2026], and the invariance of the summed drive `b_i = eta_i F - c_i` under any transfer of cost among participants is Curzi [2026].

This is not a new theorem. It is the corpus's own steady-state result — the fixed point `Q* = (eta S - C)/delta`, which is the DC gain `1/delta` of a first-order discrete section applied to the net drive [Curzi, 2026a] — written in different coordinates, and it makes no new claim. What it demonstrates is that the framework has been working in the register throughout: the storage law looks like a linear accumulator because it is written in log coordinates, and the thing it accumulates is a count, which combines by multiplication. **The fourteen-figure agreement is a numerical identity check on a change of variables, not evidence for anything** (Section 14.7).

20.6 6 The surplus functional is a Fechner law

20.6.1 6.1 The identity

The corpus's surplus functional is `f(u) = ln(1 + (N-1)u)` [Curzi, 2026]. Set `u_0 = 1/(N-1)` and rewrite:

```
f(u) = ln(1 + (N-1)u)
     = ln( (N-1) . ( u + 1/(N-1) ) )
     = ln( (u + u_0) / u_0 )                    with u_0 = 1/(N-1)
```

Compare Fechner's law in the near-threshold form in which the reference is added rather than divided out, `S = k ln((I + I_0)/I_0)`, which is the standard correction applied as the stimulus approaches absolute threshold. **They are the same function. The surplus law is Fechner's law with non-parallelism `u` as the stimulus, `k = 1`, and an absolute threshold `u_0 = 1/(N-1)`.**

N	u_0 = 1/(N-1)	u	ln(1+(N-1)u)	ln((u+u_0)/u_0)	difference
2	1.00000000	0.50	0.405465108108	0.405465108108	0.00e+00
3	0.50000000	0.20	0.336472236621	0.336472236621	0.00e+00
5	0.25000000	1.00	1.609437912434	1.609437912434	0.00e+00
12	0.09090909	0.05	0.438254930931	0.438254930931	-1.11e-16

N	u_0 = 1/(N−1)	u	ln(1+(N−1)u)	ln((u+u_0)/u_0)	difference
50	0.02040816	0.50	3.238678452164	3.238678452164	-4.44e-16
1 000	0.00100100	1.00	6.907755278982	6.907755278982	0.00e+00

The residuals are floating-point noise; **the identity is exact.** The same identity is verified independently over `N = 2` to `100` at seven values of `u` to a maximum deviation of `8.9e−16`, with the further consequences that `N_eff(u_0) = 2` and `F(u_0) = ln 2` for every `N` [Curzi, 2026]. Both consequences are one line of arithmetic: `N_eff(u_0) = 1 + (N−1)/(N−1) = 2` and `F(u_0) = ln 2 = 0.693147`, for every `N` whatever.

What the identity licenses, stated exactly. The surplus law was derived from four axioms about interaction geometry with no reference to perception [Curzi, 2026]; Fechner's law was derived in 1860 from experiments on lifted weights [Fechner, 1860]. They are the same function. **That does not show that surplus is perceived, or that perception is surplus.** It shows that a system built to count difference with diminishing returns and a threshold takes one particular form, and that an axiomatic route and a psychophysical route arrive at it. The convergence is structural and real. **It is not evidence that either route is correct, and no more is claimed from it.** The corpus's own reference apparatus files the same restatement in its established-mathematics drawer while assigning the application to the postulate drawer [Curzi, 2026], and the four-drawer classification records the surplus form as proved *conditionally* on a chosen design axiom [Curzi, 2026].

20.6.2 6.2 The threshold, and why infinitely fine discrimination is a pathology

`u_0 = 1/(N−1)` is an absolute threshold of difference: the amount of non-parallelism below which a difference does not register as difference. It shrinks as `N` grows.

N	u_0	f(0.01)	ln N	f(0.01)/ln N
2	1.000e+00	0.009950	0.693147	0.0144
10	1.111e-01	0.086178	2.302585	0.0374
100	1.010e-02	0.688135	4.605170	0.1494
10 000	1.000e-04	4.615022	9.210340	0.5011
1 000 000	1.000e-06	9.210439	13.815511	0.6667

An observer with a larger repertoire has a lower threshold of detectable difference, and the same small non-parallelism `u = 0.01` that is nearly invisible at `N = 2` delivers half the available surplus at `N = 10 000`. That is a precise statement of what a larger repertoire buys, and it is requisite variety [Ashby, 1956] arriving from a threshold constant rather than from a ceiling.

It also explains why `N -> infinity` is not desirable. As `u_0 -> 0` any non-zero difference produces unbounded surplus while `f(0)` remains exactly zero, so the function approaches a step. **An observer with infinitely fine discrimination is not maximally perceptive but maximally unstable, because every difference becomes total.** The finite threshold is what keeps the register from having an infinite basement, and finite `N` is what supplies it. The point has an independent precedent: the number of categories reliably identified without a comparison stimulus is small and does not track discrimination acuity [Miller, 1956], and Section 9.2 records that the just-noticeable-difference count therefore overstates usable categories and must be reported as an upper bound.

20.6.3 6.3 Concavity is the register, not pessimism

`f` is strictly concave in `u`, unconditionally [Curzi, 2026]. That is usually read as bad news — diminishing returns, the fortieth member worth less than the second, effort meeting a wall. **The reading is a coordinate error.** Section 6.1 showed that `f` is a logarithm of `(u + u_0)`, and a logarithm is what a straight line looks like when plotted in the world's coordinates instead of the register's. Concavity is the shape of the register, not the shape of the returns.

The marginal surplus is `f'(u) = (N−1)/(1 + (N−1)u)`. At `N = 12`:

u	f(u)	f'(u)
0.00	0.000000	**11.000000**
0.05	0.438255	7.096774
0.10	0.741937	5.238095
0.20	1.163151	3.437500
0.40	1.686399	2.037037
0.60	2.028148	1.447368
0.80	2.282382	1.122449
1.00	2.484907	**0.916667**

`f'(0) = N − 1 = 11` is the function's sharp Lipschitz constant, the quantity that feeds the contraction condition `d(T(x),T(y)) <= k d(x,y)` with `0 <= k < 1` giving existence, uniqueness and geometric convergence from any start [Curzi, 2026]. `f'(1) = 11/12`. **The ratio of first to last marginal value is exactly `N`:**

```
f'(0)/f'(1) = (N−1) / ((N−1)/N) = N = 12        verified
```

Now change coordinates and ask what the world must do to deliver **equal perceived steps** of surplus, half a nat at a time, at `N = 12`:

target `F` (nats)	required `u`	`(u + u_0)`	ratio to previous
0.0000	0.000000	0.090909	—
0.5000	0.058975	0.149884	**1.648721**
1.0000	0.156207	0.247117	**1.648721**
1.5000	0.316517	0.407426	**1.648721**
2.0000	0.580823	0.671732	**1.648721**
2.5000	1.016590 — **impossible**, `u > 1`; the ceiling is `ln 12 = 2.484907`		

Every equal step of 0.5 nats multiplies `(u + u_0)` by exactly `e^{0.5} = 1.648721`. The diminishing returns are a perfectly constant ratio requirement seen from the wrong side.

> **Concavity is not a statement that effort stops paying. It is a statement that equal experience costs equal ratios, and equal ratios compound.** The fortieth member is not worth less because members get worse; the fortieth member is worth less because by then the request is for a sixty-five per cent increase in a quantity already near its ceiling.

And the last row matters more than the pattern: **the fifth equal step is impossible, not expensive.** `u` would have to exceed 1 and cannot. `f(1) = ln N = 2.484907` is a hard bound [Curzi, 2026]. The register compresses right up to a wall, and past the wall there are not diminishing returns but no returns.

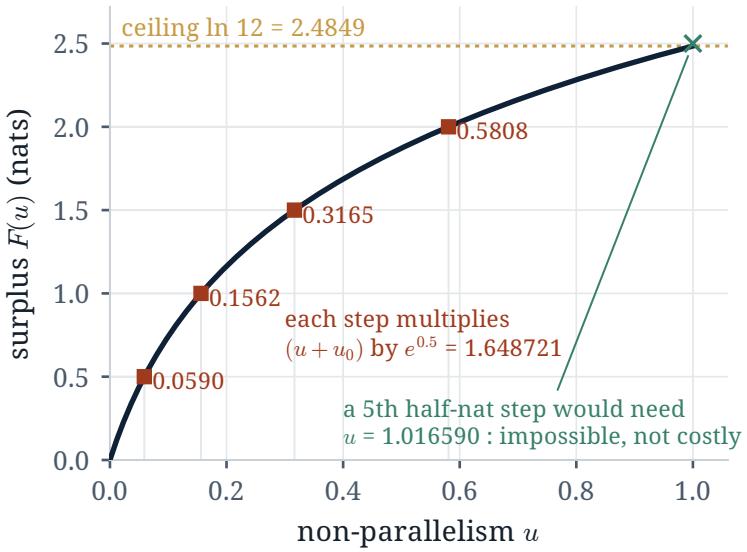


Figure 20.2: **Equal steps of experience are equal ratios of `(u + u_0)`, and the fourth step is the last one possible.** At `N = 12` the threshold is `u_0 = 1/11 = 0.090909` and each half-nat of surplus multiplies `(u + u_0)` by exactly `e^{0.5} = 1.648721`; the marginal value falls from `f'(0) = 11` to `f'(1) = 0.916667`, a ratio of exactly `N = 12`, and the fifth half-nat step would require `u = 1.016590`, which is impossible rather than expensive against the ceiling `ln 12 = 2.484907`. Derived from the surplus functional at the stated `N`; generated by `figs/src/CCP-020_fig02.py`.

20.6.4 6.4 Jensen's inequality on a concave register: the mean of percepts is not the percept of the mean

Because `f` is strictly concave, Jensen's inequality fixes a direction that this corpus's own reported figures depend on, and this paper is where the arithmetic belongs.

```
for a strictly concave f and a non-degenerate distribution of U:
    E[ f(U) ]  <  f( E[U] )
```

Averaging the percept is not the same as taking the percept of the average, and the first is always the smaller. The gap is a concavity penalty paid purely for variability in the underlying quantity, and it is not small at realistic dispersions.

The corpus's surveyed office [Curzi, 2026i] reports `N = 5` and a realised mean non-parallelism `U = 0.6019` with a realised surplus `F = 1.1498` nats. **Those two numbers are not consistent under `F = ln(1 + 4U)` and are not meant to be.** Taken as a percept of the mean,

```
f( 0.6019 ) = ln(1 + 4 x 0.6019) = ln(3.4076) = 1.226008 nats
```

whereas the reported `1.1498` nats is the **mean of `f` over the 278 operating weeks**. The two differ by `0.0762` nats, `6.2` per cent of the surplus, lost to the variability of the geometry alone and to nothing else.

> **Every occurrence of a realised surplus in this corpus must state which of the two it is.** A mean of percepts and a percept of a mean are different quantities on a concave register, they differ by a computable penalty, and reporting one under the other's name is the same error as reporting the arithmetic mean of a set of ratios (Section 4.2) with the sign reversed.

Two disclosures about the figures themselves. Curzi [2026i] prints `F(mean U) = 1.2262` and a gap of `0.0761` nats. From the printed `U = 0.6019` the value is `1.226008`, and the gap against the printed `1.1498` is `0.0762`; the printed `1.2262` corresponds to an unrounded mean `U` slightly above `0.6019`, and `1.2262 − 1.1498 = 0.0764`, so the three printed figures are not mutually consistent at the fourth decimal. The discrepancy is at the level of the rounding of `U` and does not affect the six-per-cent conclusion. It is recorded rather than corrected because the correction would require the unrounded series, which this paper does not have.

The same inequality, run the other way, is why an input delivered on several occasions can strictly dominate the same input delivered at once, and why an analyst who aggregates a variable geometry before applying `f` will read a concavity artefact as a discovery about time scale [Curzi, 2026i].

20.7 7 Design consequences, and what each one costs

Everything in this section depends on **Warrant 1** of Section 3.3 — the perceptual claim — and on a measured exponent. It is the part of the paper that Fechner and Stevens carry and that Peters and Gell-Mann [2016] cannot.

20.7.1 7.1 Budget in ratios

If equal experience is equal ratio, the unit of a design budget for a perceived quantity is a ratio and not a difference. Practically, allocate log units: a target of "the room should feel appreciably better" becomes a number — say 0.20 nats of perceived improvement — which can be divided across levers and checked by addition, because logs add. Section 12.1 does exactly this.

The constraint that makes it bite is economic. **Most world quantities cost money in proportion to the quantity and deliver experience in proportion to its logarithm; cost is linear, experience is logarithmic, so the cost of experience is exponential.** That is why the third refinement costs more than the first two together.

20.7.2 7.2 Never promise a linear improvement in a logarithmic quantity

Quantified. A designer promises "twenty per cent brighter" and delivers `+20%` illuminance:

```
delivered:  1.20 x illuminance
perceived:  1.20^0.33 = 1.06201   ->  +6.201%
the recipient receives 31.0% of what they heard
```

To deliver a perceived `+20%`:

Modality	a	world quantity required
Loudness (sound pressure)	0.67	x1.3127 (+31.27%)
Brightness	0.33	**x1.7376 (+73.76%)**
Visual length	1.00	x1.2000 (+20.00%)
Taste (salt)	1.30	x1.1506 (+15.06%)
Heaviness (lifted)	1.45	x1.1340 (+13.40%)
Warmth on arm	1.60	x1.1207 (+12.07%)
Electric shock	3.50	x1.0535 (+5.35%)

A promise of "twenty per cent brighter" is a promise to raise the lighting load by 73.76 per cent. The promise was made in per cent, the delivery is in per cent, the two per cents are

different quantities, and nobody in the conversation noticed. The correction is a rule rather than a calculation:

> **State the promise in the recipient's units and convert to the builder's before agreeing to it. A commitment expressed as a percentage of a world quantity is not a commitment about experience and must not be sold as one.**

The exponent is the uncertain part of every row above. These are cited point estimates with known between-study variation [Stevens, 1975, Teghtsoonian, 1971, Poulton, 1979]; they carry no interval here because none is available from the sources as reproduced, and Section 10 requires re-measurement under an application's own conditions before any of them calibrates anything.

20.7.3 7.3 Doubling the percept, by modality

`psi = k I^a` doubles when the world quantity is multiplied by `2^{1/a}`:

Modality	`a`	world multiple to **double** the percept	percept of a x2 world	percept of a +10% world
Loudness (1 kHz, sound pressure)	0.67	**2.8138**	x1.5911	+6.59%
Brightness	0.33	**8.1698**	x1.2570	+3.20%
Visual length	1.00	2.0000	x2.0000	+10.00%
Taste (salt)	1.30	1.7044	x2.4623	+13.19%
Heaviness (lifted)	1.45	1.6129	x2.7321	+14.82%
Warmth on arm	1.60	1.5422	x3.0314	+16.47%
Electric shock	3.50	**1.2190**	x11.3137	+39.60%

Two anchors. **Doubling apparent brightness takes 8.17 times the luminance. Doubling apparent loudness takes 2.81 times the sound pressure, which is 8.99 dB** — the familiar rule of about ten decibels per doubling of loudness, recovered from the exponent rather than asserted, and cross-checked against the independent loudness-matching convention in Section 9.4 [Stevens, 1956]. And the last row inverts everything: **doubling the sensation of an electric shock takes only twenty-two per cent more current.** In the expansive modalities a small world change is a large experience, which is why heat, weight and pain are dangerous to design near and why margins there are set in world units.

20.7.4 7.4 Geometric ladders

Anything intended to deliver equal steps of experience must be a geometric progression. A five-rung ladder from 40 000 to 100 000 with equal perceived steps has common ratio `(100000/40000)^{1/4} = 1.257433`, that is `+25.74%` per rung:

rung	value	`ln`
0	40 000.00	10.596635
1	50 297.34	10.825707
2	63 245.55	11.054780
3	79 527.07	11.283853
4	100 000.00	11.512925

Each `ln` step is `0.229073`, identical. The linear ladder over the same range, in equal steps of 15 000:

rung	value	ratio to previous	`ln` step
0	40 000.00	—	—
1	55 000.00	1.375000 (+37.50%)	0.318454
2	70 000.00	1.272727 (+27.27%)	0.241162
3	85 000.00	1.214286 (+21.43%)	0.194156
4	100 000.00	1.176471 (+17.65%)	0.162519

The linear ladder's perceived steps shrink monotonically: the last step is barely half the first. An equal-increment ladder is experienced as a decelerating one, and the person on the top rung feels the least rewarded.

Whether money is perceived logarithmically is a contested empirical question and this paper does not settle it. The arithmetic demonstrates what a geometric ladder does to *any* logarithmically perceived quantity; the construction applies without argument to quantities known to be logarithmic — light, sound, and the surplus law of Section 6 — and its application to compensation is an extension requiring separate measurement.

20.7.5 7.5 Reporting: geometric means, log axes, ratios

Three reporting rules, each a corollary of a table row.

1. **Central tendency of a perceived quantity is the geometric mean** (Section 4.2). The arithmetic mean is reported only where the quantity is genuinely additive in the world and nobody is experiencing it.
2. **Plot on a log axis.** A log axis is the register. A quantity plotted linearly and perceived logarithmically produces a chart whose visual impression contradicts the experience it describes.
3. **Report ratios, not differences.** "From 200 to 300 lux" is `+50%` and `0.405465` nats of world ratio; "an increase of 100 lux" has no perceptual content until the starting point is known.

20.7.6 7.6 The reference is cheaper than the quantity, and where the register does not apply

From Section 4.5: the percept is `k ln(I/I_0)` and `I_0` is set by adaptation and context [Helson, 1948]. **Moving the reference by a factor of two has exactly the same perceptual effect as moving the quantity by a factor of two, and is usually far cheaper.** A designer who only ever moves `I` is working with one of two available terms.

Six boundaries limit everything above.

- **Expansive modalities invert the advice.** At `a > 1` the world compresses rather than the percept and a small world increment is a large experience; every rule about needing large ratios reverses. **The exponent must be looked up. Compression is not the default.**
- **Small counts are not logarithmic.** Numerosity for very small sets is direct and near-linear.
- **Visual length is essentially linear** at `a` near 1.00. For spatial extent, differences behave like differences, and this section's corrections are unnecessary and would be wrong. This is precisely why bar charts work and area encodings do not.
- **The register describes magnitude judgement, not value, preference or decision.** Extending it is a separate empirical claim and is not made here.
- **Outside the working range the Weber fraction is not constant.** Near threshold and near saturation the law degrades in a documented way.

- **The constants are individual.** The form recovers across people; `c`, `k` and `a` vary. Where precision matters, measure the person.

20.7.7 7.7 The ethical boundary

A register argument can be used to give someone more experience for the same resources, or to give them the same experience for fewer resources while telling them it is more. The arithmetic is identical and the two uses are not. The operative test is disclosure and decline: a reference-shifting technique the recipient would object to on being told of it is not endorsed here, and its efficiency does not change that. The vestibule that makes a reading room feel bright is disclosed by walking through it; a brightness claim computed in world units and sold in perceptual ones is not.

Disclosure is a test, not an immunity. The corpus's own strongest available test case records that disclosure did **not** confer immunity, that a stated component count overstated diversification by a factor of 4.61, and that a levy presented as a costless accounting correction in fact reallocated burden [Curzi, 2026]. A designer relying on disclosure to license a reference-shifting technique should read that record first.

20.8 8 The postulate method

20.8.1 8.1 What is established in this section, and what is not

Nothing in this section is a theorem. Postulation is a discipline of practice rather than a mathematical object. What is offered is: a statement of the **Postulate Method**, badged as a house formalisation of standard hypothetico-deductive practice [Popper, 1959, Platt, 1964]; six **admissibility conditions** separating a legitimate postulate from an unfalsifiable one; one genuine derivation — Fechner's integration of Weber's law — as the exemplar, because it is the historical case in which the method worked on exactly this problem; two quantitative instruments, an **empirical-content measure** and a **held-out skill score**; and the **badge convention**. Everything asserted here is asserted about the practice of inquiry, not about minds. The claims about minds are in Section 9 and they arrive with data.

20.8.2 8.2 The method, stated

> **The Postulate Method [house formalisation].** Let `X` be a quantity that cannot be observed.
>
> **Step 1 — POSTULATE.** Assert a structure for `X`: what kind of thing it is, what operations it supports, how it relates to observable quantities. State the assertion in the open, badged, in the sentence in which it is first used.
>
> **Step 2 — DERIVE.** From the postulate plus already-accepted results, derive statements about observables. The derivation must be valid. **A derived statement counts only if the negation of the postulate would not also produce it.**
>
> **Step 3 — FORBID.** Extract the set of observations the postulate rules out, and write them down before looking. This is the postulate's content and the only part of it that can be tested.
>
> **Step 4 — TEST.** Observe. If a forbidden observation occurs, the postulate is dead as stated. If none occurs, the postulate has survived, which is not the same as being true.
>
> **Step 5 — DECLARE.** Publish the surviving postulate with its badge still on. Survival does not promote a postulate to a theorem, and nothing does.

Step 2's second sentence is where most work in the human sciences fails. **A derivation that both a postulate and its negation produce is not evidence for the postulate.** If "I am a morning person" predicts that some mornings go well and its negation predicts the same, then mornings going well is data about mornings.

The method is not original and no priority is claimed for it. It is the falsificationist requirement [Popper, 1959] sharpened by the demand that competing hypotheses be discriminated rather than merely tested [Platt, 1964], with the protective-belt problem named as Lakatos [1970] named it, and with Meehl [1967]'s observation in view that directional predictions are almost costless in the social sciences — which is why this framework's specifications require point predictions with constants.

20.8.3 8.3 The badge taxonomy

The badges are a closed, small, non-overlapping set, because from here on almost every claim is one of these things and the reader is entitled to know which. This is a **definition**, not a finding; the corpus's four-drawer evidential classification — PROVED, POSTULATED, BORROWED, UNTESTED — is the same instrument at corpus scope [Curzi, 2026].

Badge	Meaning	Test that applies	Can it be wrong?
PROVEN	A theorem: follows from stated axioms by valid deduction	Check the proof	Only if the proof is invalid
DESIGN AXIOM	Chosen, not derived; fixes a modelling convention	None; it is a decision	It can be *unwise*, not false
POSTULATE	An assertion about the world that generates forbidden observations	Steps 3 and 4	**Yes, and how is stated**
HEURISTIC	A rule of thumb that orders or triages; not an equation	Does it improve triage?	It can be useless
INTERPRETIVE CORRESPONDENCE	Two formalisms have the same shape; no theorem connects them	None available	Not the kind of thing that is true or false
METAPHOR	A figure of speech carried deliberately	None	No; it carries no claim

Four rules govern use. **(i)** The badge appears in the same sentence as the first statement of the thing — not in a footnote and not "as discussed earlier". **(ii)** A badge never upgrades through use: a postulate repeated for two hundred pages is still a postulate, and the commonest corruption in this literature is an assumption made once, then referenced, then referenced by the reference, until the fourth-generation citation reads as a finding. **(iii)** A postulate that cannot state what would kill it is not a postulate; it is at best a heuristic and must be demoted in print, visibly. **(iv)** Interpretive correspondences and metaphors carry no evidential weight whatever; they may be productive and they may not be cited as support.

The badge is not a hedge. It is a specification of what kind of test the reader is entitled to demand: PROVEN licenses a demand for a proof, POSTULATE for the forbiddance list, DESIGN AXIOM for a justification and a counter-proposal.

20.8.4 8.4 Six admissibility conditions

A postulate is admissible in this framework if and only if it satisfies all six. They are labelled CC-P1 to CC-P6 to keep them distinct from the verification postulates P1 to P7 of Section 14.2.

```
CC-P1  OPENNESS      Stated explicitly, in the open, badged, before it is used.
CC-P2  ASYMMETRY     Generates at least one observable statement that its own negation
                     does not generate.
CC-P3  REACHABILITY  The forbidden observations are reachable with instruments the reader
                     has, at a cost the reader can pay.
CC-P4  LETHALITY     A specific, named, pre-recorded result kills it, written before looking.
CC-P5  MINIMALITY    It does not smuggle. Every clause does work; a clause that forbids
                     nothing is deleted, not carried.
CC-P6  STABILITY     Its parameters recover across scopes and observers, or the postulate
                     states plainly that it is indexed to one observer and claims nothing
                     wider.
```

CC-P1 to CC-P4 are the standard falsificationist demands, sharpened [Popper, 1959]. CC-P5 and CC-P6 are the two that most work in the human sciences fails.

CC-P5 is a defence against the protective belt. Lakatos [1970] described how a research programme protects its core by accreting auxiliary hypotheses, each added to absorb a disconfirmation; each addition is locally reasonable and the aggregate is a postulate that forbids nothing. **A postulate should get smaller as it survives, not larger.**

CC-P6 is reproducibility demanded in advance. A postulate about one observer's inner structure that holds only for that observer is admissible — single-subject work is real work — **but it must say so and must not be quoted as a claim about people.** The formal version of this gate is specific objectivity: a comparison of two objects must not depend on which instrument was used, and the corpus's treatment shows what it takes to earn it — given a badged postulate, Rasch person measures make a difference of person measures genuinely ratio-scaled, with separation supplying the true zero that division requires [Curzi, 2026t, Rasch, 1960].

20.8.5 8.5 The exemplar: how Fechner did it

The reason to believe the method works on inner states is a case rather than an argument. In 1860 Fechner [1860] executed Steps 1 to 5 on the least tractable quantity available and the result is still in use.

He began with something measured: Weber [1834]'s finding that the smallest detectable change in a stimulus is a constant proportion of it. A weight of 100 g needs about the same fractional increment to feel heavier as a weight of 1000 g; the increment in grams is ten times larger and the ratio is the same. **That is a finding about behaviour — what people can and cannot discriminate — and nothing about inner magnitude is asserted by it.**

Fechner wanted the quantity on the other side of the wall and could not reach it, so he postulated a bridge:

> **Fechner's postulate [POSTULATE, 1860].** All just-noticeable differences are subjectively equal. The step in sensation produced by one just-noticeable difference is the same size wherever on the scale it is taken.

It is audacious because it asserts an equality between inner quantities nobody can inspect, and modest because it asserts exactly one thing and hands the consequences to arithmetic:

```
postulate:   dS = k' x (one JND)          all JNDs equal in sensation units
Weber:       one JND = c I                in stimulus units
therefore:   dS = k dI / I                k absorbs k' and c
integrate:   S = k ln I + const
```

```
set S = 0 at absolute threshold I_0:
            S = k ln( I / I_0 )          FECHNER'S LAW
```

The ladder arithmetic is checkable. With a Weber fraction `c = 0.05`, the stimulus after `n` just-noticeable steps is `I_n = I_0 (1+c)^n`, and inverting recovers the count exactly:

n	I_n	ln(I_n/I_0) / ln(1+c)
0	1.0000	0.0000
10	1.6289	10.0000
50	11.4674	50.0000
100	131.5013	100.0000
200	17 292.5808	200.0000
283	992 136.9785	283.0000

Sensation measured in just-noticeable-difference units is exactly the log of the stimulus ratio divided by `ln(1+c)`, and the constant is `k = 1/ln(1.05) = 20.495934` JND per nat. **This is where the unit of `k` is fixed and it is the answer to the dimensional question of Section 14.1.**

Now the step that makes this an exemplar rather than an anecdote. **Fechner's postulate is not tested by looking at it; it is tested by consequences that live entirely on the observable side of the wall.** The law forbids equal stimulus *differences* producing equal report differences across the range; it forbids a modality with a constant Weber fraction from showing a linear relation between stimulus and any well-constructed sensation scale; it forbids the count of just-noticeable differences between two stimuli from depending on the path taken between them; and it predicts, quantitatively, the number of discriminable steps across a modality's dynamic range. Every one of those is testable with nineteenth-century equipment.

Some of them turned out to be wrong, which is the point. Section 9.5 documents how Stevens [1957] killed Fechner's law as a general account and replaced it with a power law, using a method Fechner had not considered. **The postulate died the death a legitimate postulate is supposed to die: by its own consequences, at the hands of a better instrument.** A working postulate is not one that survives forever; it is one that fails informatively.

20.8.6 8.6 Empirical content: measuring what a postulate forbids

Popper [1959]'s central point was that the content of a theory grows with the class of what it forbids, and that a theory compatible with everything says nothing. That can be given a number.

> **The content measure [HEURISTIC — this framework's own construction].** Let a postulate admit a fraction `q` of the outcomes enumerated as possible before it was adopted, `0 < q <= 1`. Define
>
> ```
> H_forbid = −ln q      nats
> ```
>
> `q = 1`, forbidding nothing, gives zero content. Halving the admitted outcomes adds `ln 2` nats.

The formula is this framework's own and no external source states it. Popper [1959], Lakatos [1970], Platt [1964] and Meehl [1967] support the underlying idea that content is what a theory forbids and that bold conjectures are worth more per test; **none of them states this formula**, and it is badged POSTULATED in the verification chain of Section 14.3.

q (fraction of outcomes still admitted)	nats	bits
1.000	0.0000	0.0000
0.900	0.1054	0.1520
0.500	0.6931	1.0000
0.250	1.3863	2.0000
0.100	2.3026	3.3219
0.050	2.9957	4.3219
0.010	4.6052	6.6439
0.001	6.9078	9.9658

The table is used comparatively, never absolutely. Two postulates about the same question, over the same enumerated outcome set, can be ranked by content, **and the higher-content postulate is the one to test first.** A postulate at `q = 0.9` costs a full experiment and, if it survives, has bought `0.1054` nats; a postulate at `q = 0.05` costs the same experiment and buys `2.9957` nats, twenty-eight times the return for the same expenditure. That ratio is the practical argument for preferring bold postulates, and it is arithmetic rather than temperament.

Four limits are stated with the instrument. **(i)** `q` is a fraction *of something*, so the outcome set must be enumerated and fixed in advance; chosen after the fact, the number is meaningless. **(ii)** `q` presupposes a defensible measure over outcomes, which is almost never available — hence HEURISTIC and not equation. **(iii)** Two analysts enumerating outcomes differently assign different `q` to the same postulate, so the measure is identifiable only relative to a published outcome list: publish the list or do not publish the number. **(iv)** Content adds under independent conjunction, `−ln(q_1 q_2) = H_1 + H_2`, which is the property that makes a log measure the right shape and the reason to prefer it over `1 − q`. **Whether the heuristic helps is itself an untested empirical question** (H5), and the test is stated in Section 15.

20.8.7 8.7 Introspective report is data, not verdict

The sharpest form of the objection to any of this is not that inner states are unobservable but that the reports substituted for them are unreliable. Nisbett and Wilson [1977] made the case directly: people confidently report the causes of their own judgements and those reports are frequently constructions assembled after the fact rather than readouts of process. The framework accepts the result. **It changes what a report is for.**

```
WRONG USE   The subject reports "the second one was twice as loud", and the analyst records
            that the second one was twice as loud in their experience. This treats the report
            as a verdict about an inner quantity.

RIGHT USE   The subject emits the number 2 under a specified procedure. The analyst records
            the number as behaviour, then asks whether numbers emitted under that procedure
            have the structure the postulate predicts: consistency under repetition,
            transitivity, agreement across observers, the right functional form against a
            physical variable, and invariance when the procedure's irrelevant details change.
```

The right use never requires the report to be accurate about anything. It requires the *pattern* of reports to have properties that would be unlikely if the postulated structure were absent — which is how a thermometer is validated, nobody having asked whether mercury knows the temperature. Nisbett and Wilson [1977]'s finding constrains *causal* reports about one's own processing specifically; it constrains magnitude reports far less, and Section 9 shows how well those pass the structural tests. **Two further constraints on any self-report instrument are real and are not answered here:** question wording, scale range and order effects [Schwarz, 1999], and the focusing illusion that contaminates any rating of a named component of well-being [Kahneman et al., 2006].

20.8.8 8.8 The freeze, and the held-out skill score

The move that separates a test from a story is the **freeze**, and it is more necessary at single-observer scope than at institutional scope, because the person fitting the model and the person being measured are the same and there is no reviewer.

```
1. Split the record in two by TIME, never at random. Early half = training.
2. Fit the parameters on the training half only.
3. WRITE THE PARAMETERS DOWN. They are frozen. No refitting.
4. Predict the held-out half with the frozen parameters.
5. Compare against four baselines. Report all five numbers, including the embarrassing ones.
```

The four required baselines are **persistence** (tomorrow equals today), **direct input** (output is a rescaling of input, no storage), a **matched autoregression** (the same free-parameter count, no theory in it), and the **shuffled-input control** (the same frozen model with the input series randomly permuted). The fourth is usually omitted and is the decisive one: **if the model predicts the held-out data just as well with its inputs shuffled, it is not using the inputs.** The full protocol, its preregistration form and its confirmatory application are Curzi [2026i], whose own application **failed** its second falsification rule at a 2.8 per cent held-out advantage over a matched AR(1) against a declared 10 per cent margin. This is preregistration used as designed [Nosek et al., 2018], against the researcher degrees of freedom that make anything presentable as significant [Simmons et al., 2011] and the forking paths that catch an honest author [Gelman and Loken, 2013].

The comparison statistic is the ordinary forecasting skill score:

```
SS = 1 − RMSE_model / RMSE_baseline
```

`SS = 0` means the model matched the baseline, `SS > 0` that it beat it, `SS < 0` that the baseline beat it, unbounded below, which correctly expresses that a model can be arbitrarily worse than doing nothing. One `SS` is reported per baseline and **the headline number is the smallest of them.** An `SS` approaching 1 is treated as **a red flag for leakage rather than as success**, and a baseline whose own error is zero leaves nothing to explain and makes the statistic undefined. Three limits attach. RMSE requires at least an interval scale on the outcome, so an ordinal one-to-five rating does not license it and the honest substitutes are rank-based — the same demotion rule the corpus applies to its own moment equation, which contains a division and a multiplication and is therefore demoted to a heuristic ordering when its inputs are ordinal [Curzi, 2026m]. Identifiability fails the instant any refitting occurs after the split, and the only defence is the written frozen parameter list. And against the shuffled control `SS = 0` reads in the opposite direction from every other row: it means the model was never using its inputs.

20.8.9 8.9 Postulates behave like stored order

A postulate is a load-bearing member in this framework, and postulates behave like the stored order the framework already models: they decay without maintenance, they convert experience into structure at some rate, they cost something to maintain, and they occupy headroom, since a mind full of unexamined postulates has no capacity for a new one. The amendment ledger of Section 11.2 is a decay-reduction device and the freeze is a conversion-raising one.

This is an identical model form at a different scope, and identical structure at three scopes is a statement about the chosen model, not a claim of kinship or of cross-scope aggregation [Curzi, 2026]. It is badged accordingly, and Section 3.7 records that where the scope includes a group whose members constitute the aggregate, the specification is not identified from observational data [Manski, 1993].

20.9 9 The register perception uses: three laws, and their limits

Section 8's method is only as good as its best case, and this is the best case. The claim to be established is narrow and consequential:

> **Subjective magnitude has been measured — not ranked, not rated, not scored, but measured on a scale that supports ratios, from reports given by ordinary participants under a procedure taking about twenty minutes, and the measurements replicate.**

If that claim is false, the framework's treatment of inner quantities is metaphor. If it is true, the question about the inside is not *whether* it can be measured but how well, at what cost, and with what error bars — which is an engineering question. **The claim is defensible and it is not proved**; Section 9.9 gives the limits, which are severe, and Section 15 records that the measurement-theoretic conditions that would settle it were never tested.

20.9.1 9.1 Three laws, one register

```
WEBER    1834    dI / I = c              the JND is a constant RATIO
FECHNER  1860    S = k ln(I / I_0)       sensation is LOGARITHMIC in intensity
STEVENS  1957    psi = k I^a             sensation is a POWER FUNCTION of intensity
and, taking logarithms of the third:
                 log psi = log k + a log I        LINEAR in log–log
```

Three statements, and the third did not abolish the register but relocated it. Weber's law is about ratios; Fechner's law is a logarithm outright; Stevens' law is a straight line *in log-log coordinates*, which is to say that equal ratios of stimulus produce equal ratios of sensation. **All three say that perception operates on ratios and not on differences. They disagree about what the ratio does when it arrives.** That is the register claim on which the Lens Theorem's antecedent rests, and it survives the Fechner-to-Stevens transition intact: the transition replaced *log in, linear out* with *log in, log out*.

Two properties make Weber's law a strong result rather than an observation. It concerns discrimination performance, which is behaviour, settled by a forced choice and a tally with no inner quantity invoked. And it holds across modality after modality over the middle of each range and fails at both ends in a characteristic and reproducible way — near absolute threshold `c` rises, near saturation it rises again. A law that fails at the edges in the same way everywhere has a mechanism behind it.

Three cautions attach to Weber's law and are stated once. `I` must be ratio-scaled with a true zero, so the law is **not statable** for a stimulus dimension lacking one — for pitch, for hue, for position, the analogous statement must be constructed differently. As `I -> 0` the law predicts infinitely fine discrimination, which is false and known to be false; absolute threshold intervenes and the standard repair is `dI/(I + I_r) = c` with `I_r` a small resting term. And a just-noticeable difference is conventionally the 75 per-cent-correct point in a two-alternative task, so **the convention must be stated with the number**, since a `c` defined at 75 per cent and one defined at 84 per cent are different numbers for the same observer.

20.9.2 9.2 The discriminable-step count, and why it is an upper bound

If sensation is counted in just-noticeable differences, a modality's dynamic range contains a finite, computable number of discriminable steps:

> **The discriminable-step count [derived from Weber].** For a modality with Weber fraction `c` over a dynamic range `R = I_max / I_min`, `N_jnd = ln R / ln(1 + c)`.

c	R = 10	R = 100	R = 10^4	R = 10^6	ln N_jnd at R = 10^6
0.01	231.4	462.8	925.6	1388.4	7.2359
0.02	116.3	232.6	465.1	697.7	6.5477
0.05	47.2	94.4	188.8	283.2	5.6460
0.10	24.2	48.3	96.6	145.0	4.9764
0.20	12.6	25.3	50.5	75.8	4.3278
0.33	8.1	16.1	32.3	48.4	3.8804

The last column connects the psychophysics to the framework. `N_jnd` is a count of distinguishable categories along one perceptual dimension, so `ln N_jnd` is the surplus ceiling for that dimension [Curzi, 2026]. A modality with a Weber fraction of 0.05 and six decades of range affords at most 5.646 nats; a modality with `c = 0.33` affords 3.880. **The maximum surplus available along a perceptual dimension is fixed by that dimension's discrimination acuity and its dynamic range, both measurable with a forced-choice task.** That is requisite variety [Ashby, 1956] arriving with a number.

The formula overstates usable categories and must be reported as an upper bound. Absolute-identification experiments — counting the categories reliably distinguished without a comparison stimulus present — recover far fewer than the just-noticeable-difference count predicts, because identification without a comparison is much harder than discrimination with one [Miller, 1956]. The gate is real and frequently failed, and Section 11.3's failure mode FM26 is named for it.

A moment is not built from one dimension. If an observer attends to `m` perceptual dimensions whose discriminations are independent, the joint category count multiplies and the ceiling adds, `ln N_total = sum_i ln N_jnd(i)` — **and the independence condition does all the work.** Independence across perceptual channels is the `u = 1` condition transposed to one observer's own channels and is not generally satisfied: loudness and brightness in a room co-vary with how many people are in it. The honest form is the framework's own, `N_eff = 1 + (m−1)u` with `u` the mean non-parallelism between channels, and **at this scope it is a postulate about model form rather than a derived result** [Curzi, 2026]. Two perceptual channels that always move together are one channel.

20.9.3 9.3 What an exponent costs

The design question is how much more of the physical thing must be supplied for a recipient to notice a doubling.

modality	a	log-ratio required (nats)	in decades
Electric shock	3.50	0.1980	0.0860
Warmth on arm	1.60	0.4332	0.1881
Heaviness (lifted)	1.45	0.4780	0.2076
Taste (salt)	1.30	0.5332	0.2316
Visual length	1.00	0.6931	0.3010
Loudness (1 kHz, SPL)	0.67	1.0345	0.4493
Brightness	0.33	2.1004	0.9122

computed as `ln 2 / a` nats, a perceived doubling costing `ln 2 = 0.693147` nats of response log-magnitude which the exponent converts into stimulus log-units. **To double perceived brightness requires 0.912 decades — a factor of 8.17 — of luminance; to double perceived shock requires 0.086 decades, a factor of 1.22.** Measured in the log units where the comparison is proper, 2.1004 nats against 0.1980, a ratio of 10.6. A budget that treats "make it brighter"

and “make it hurt less” as comparable engineering requests will be wrong by most of an order of magnitude.

Compressive modalities are the ones that must handle enormous physical ranges — sound pressure across the audible range spans six orders of magnitude, luminance across a day more than ten — and a compressive exponent packs that range into a manageable perceptual span at the cost of resolution. **Expansive modalities are the ones that must trigger action:** warmth at `a = 1.6` triples in felt magnitude when the stimulus doubles, and shock at `a = 3.5` multiplies by 11.31 when the stimulus doubles and by 3162 when it goes up a factor of ten. **A system that must produce withdrawal before damage cannot afford compression, and the numbers show it has none.** The functional reading is an interpretation; the numbers are the data.

20.9.4 9.4 The loudness cross-check, and the factor-of-two trap

The loudness exponent is quoted as 0.67 and is the row most often misused, because 0.67 is an exponent on **sound pressure** and intensity goes as pressure squared:

```
pressure ratio for perceived doubling       = 2.8138
intensity ratio  = 2.8138^2                 = 7.9177
in decibels      = 10 log10(7.9177)         = 8.986 dB
implied intensity-domain exponent = 0.67/2  = 0.335
```

Compare the independent convention from acoustics. The sone scale defines a doubling of loudness and the standard rule of thumb is that a 10 dB increase doubles perceived loudness [Stevens, 1956]; that rule implies an intensity-domain exponent of `ln 2 / ln 10 = 0.3010` and hence a pressure-domain exponent of `0.6021`.

```
from Stevens' 0.67 (pressure)   ->  perceived doubling at  8.99 dB
from the sone convention        ->  perceived doubling at 10.00 dB
discrepancy                                                1.01 dB
```

Two methods developed for different purposes — one from magnitude estimation, one from loudness-matching standards — agree on the location of a perceived doubling to within about one decibel. That is the kind of convergence that makes a measurement claim credible, and it is offered here precisely because it is not exact: the one-decibel gap is the method dependence of Section 9.9 appearing in the third significant figure, exactly where the honesty clauses say to expect it.

20.9.5 9.5 Why the dispute lasted a century, and the two-decade problem

A logarithm and a compressive power function look alike over a limited range, and the resemblance is computable. Fitting a logarithmic law by least squares to a power function with `a = 0.33` over 200 log-spaced points:

```
over I in [1, 100]   (two decades)
    best fit:  S = 0.6306 + 0.7475 ln I
    R^2      = 0.963165
    maximum relative deviation = 36.94 %

over I in [1, 10000] (four decades)
    R^2      = 0.872523
    maximum relative deviation = 314.06 %
```

Over two decades a logarithm captures 96.3 per cent of the variance of a power function. With nineteenth-century instrumentation, observer variability of ten to twenty per cent, and

stimulus ranges rarely exceeding two decades, the two were not distinguishable. Over four decades the maximum relative deviation goes from 37 per cent to 314 per cent.

Stevens did not win by argument; he won by extending the range and by asking a question whose answer discriminated the hypotheses — which is Platt [1964]'s prescription exactly, and the general rule is worth extracting: *when two models fit equally well, the remedy is not more data in the same range but data in a range where they disagree.* The second and decisive blow was qualitative. **A logarithmic law cannot produce an expansive relation at all**, because `ln` is concave everywhere, so no logarithmic sensation scale can grow faster than its stimulus. Heaviness at 1.45, warmth at 1.6 and shock at 3.5 are not near-misses for Fechner's law; they lie outside its range of possible outputs. One expansive modality measured reliably kills the law as a general account, and several were [Stevens, 1975]. Krueger [1989] is the standing attempt at reconciliation, with peer commentary, and the framework cites it rather than adjudicating the dispute on its own authority.

This section is also the standing warning against this paper's own headline fit. An `R^2` of 0.999932 over two decades (Section 12.3) does not by itself exclude a logarithmic law, for the reason computed above, and Section 12.3 therefore rests its conclusion on the residual pattern, the ratio-invariance test and an explicit model comparison rather than on the fit statistic.

20.9.6 9.6 Running a magnitude estimation

The procedure, in the form that survives the criticisms of Section 9.9 best. It is the instrument this framework borrows, so it is specified rather than described.

```
1. FIX THE CONTINUUM. One dimension, varied alone, everything else held constant.
2. CHOOSE THE STIMULUS SET. At least 7 levels, GEOMETRICALLY spaced, spanning at least
   1.5 decades. Linear spacing wastes most of the points at the top.
3. RANDOMISE ORDER, fresh per observer. Sequential dependencies are real.
4. INSTRUCT IN RATIOS, not in categories:
     "Assign any positive number to the first stimulus. For each one after, assign a
      number in PROPORTION to how strong it seems. If one seems half as strong, use half
      the number. Do not restrict yourself to whole numbers or to any range."
   The ratio instruction and the removal of range restriction are both load-bearing. An
   instruction mentioning a scale from 1 to 10 destroys the ratio property and yields
   ordinal data.
5. FREE MODULUS, or a modulus in the MIDDLE of the range. A modulus at an extreme biases
   the exponent.
6. REPLICATE each level at least twice per observer, separated in the sequence.
7. AGGREGATE BY GEOMETRIC MEAN, never the arithmetic mean.
8. FIT IN LOG-LOG. Report slope, intercept, R^2, the residual pattern, the stimulus range,
   the modulus policy, and the exact instruction text.
9. TEST THE RATIO PROPERTY DIRECTLY: pick stimulus pairs with equal physical ratios at
   different absolute levels; the response ratios must match.
```

Step 9 deserves the emphasis. The ratio-scale claim is not established by the fit; it is established by the invariance. A power function fitted to numerals would still fit. What a merely ordinal response would not do is produce the *same* response ratio for the same stimulus ratio at different absolute levels. **That test is cheap and it is skipped almost universally**, and Section 12.3 runs it and reports what it does and does not establish.

The scale-type stakes are exactly the ones the measurement literature sets. A one-to-seven rating item is not magnitude estimation and yields ordinal data [Likert, 1932, Stevens, 1946]; an interval scale can be built from paired comparisons without asking anyone to use numbers [Thurstone, 1927]; and what must be true before a quantity can be said to exist at all is Krantz et al. [1971].

20.9.7 9.7 Cross-modality matching

The strongest internal-consistency argument removes numbers altogether. Instead of asking an observer to say a number, ask them to **adjust one modality to match another** — squeeze a handgrip until the force feels as strong as the tone sounds loud. If modality 1 has exponent `a_1` and modality 2 has `a_2` against their respective physical continua, matching 1 to 2 should produce a relation whose log-log slope is `a_1/a_2`:

match	to	predicted matching exponent `a_1/a_2`
heaviness	loudness	2.1642
warmth	brightness	4.8485
shock	length	3.5000
salt	loudness	1.9403

These are risky predictions in the sense of Section 8.2: derived from measurements made with one procedure and tested with a completely different one involving no numbers at all. Nothing in the matching task lets an observer consult the exponents. Stevens [1975] reports that matching exponents came out close to the predicted ratios across many modality pairs, and this is the strongest available evidence that magnitude estimation measures something rather than recording a habit of number use.

It is evidence of consistency rather than of correctness. A systematic bias present in both tasks would produce the same agreement, and the measurement-theoretic literature has pressed exactly this point [Luce, 2002, Narens, 1996]. The correct summary is that the numbers cohere across procedures sharing no method, which is much more than a rating scale can claim and less than a physical measurement can claim.

20.9.8 9.8 The exponent as a conversion key

```
perceptual log-magnitude:   ln psi = ln k + a ln I
therefore:                  d(ln psi) = a d(ln I)
ONE NAT of stimulus log-ratio  ->  a NATS of perceptual log-ratio.
```

The exponent is a linear map between two logarithmic registers, which is what a conversion key is. Worked: a tenfold increase in luminance is `ln 10 = 2.302585` nats of stimulus log-ratio; at `a = 0.33` that is `0.759853` nats of perceptual log-ratio, which is `0.759853 / ln 2 = 1.096` bits. A tenfold increase in shock intensity is the same 2.302585 nats of stimulus and `8.059` nats of perceptual log-ratio, which is `11.63` bits. **The same physical decade lands as 1.1 bits in one register and 11.6 bits in another.** The conversion is exact given the exponent; the exponent is the uncertain part.

Two conversions differ by a factor of two, and which one is meant has to be stated every time. The corpus's tabulated keys give `2.885390` for nats to octaves and `3462.4681` for nats to cents, alongside `1.442695` for nats to bits, and the first two are exactly twice what a reader arriving from acoustics would compute. **The factor of two is derived rather than mistaken, and this paper adopts the corpus's reading while insisting on the label that makes it unambiguous.**

The two quantities are these.

- **A doubling of the quantity itself.** One octave is one doubling, so it is `ln 2 = 0.693147` nats and the key is `1/ln 2 = 1.442695` — numerically identical to the nats-to-bits key, because an octave and a bit are both one doubling — with `1200/ln 2 = 1731.234` cents per nat. This is the conversion used in Section 4.1 and everywhere in this paper that an octave means a doubling of a stimulus, a stock or a percept.
- **An octave of *essence pitch* per nat of *surplus*.** The corpus's essence law is square in its index, `f(n) = k n^2`, so the cross-disciplinary frequency is `f = k N_eff^2 = k e^{2F}` and

`log_2(f/k) = 2F / ln 2`. **Doubling the effective count therefore multiplies the frequency by four, which is two octaves and not one** [Curzi, 2026l]. The key is `2/ln 2 = 2.885390` essence octaves per nat and `2400/ln 2 = 3462.468098` cents per nat, and both are *exact given the square-law postulate* [Curzi, 2026k].

The two figures are internally consistent with each other (`3462.4681 / 2.885390 = 1200` exactly) because they are the same conversion expressed in two units, and the corpus derives the doubling explicitly rather than acquiring it by accident [Curzi, 2026l,?]. **What is a real defect is the label**: the tabulated key reads "nats to octaves" without the word *essence*, and read generically it is off by a factor of two. That is the exact shape of the failure mode named in Section 11.3 as FM22, the pressure-versus-intensity confusion, and the repair is a word rather than a number. **This paper writes "essence octaves" wherever the square-law conversion is meant and "octaves" only for a plain doubling**, and the exponent `s` in `f(n) = k n^s` is a postulate rather than a measurement, so the constant is exact given `s = 2` and not independently of it [Curzi, 2026l].

20.9.9 9.9 The honest limits, in detail

This section is part of the result, not a hedge appended to it. A reader who takes the exponent table without it will overreach within a page.

- **The exponent depends on stimulus range.** Compressing the range of presented stimuli tends to raise a compressive exponent and lower an expansive one, because observers spread their numbers over roughly the same span regardless of the physical range supplied. Teghtsoonian [1971] argued that this is systematic — that the subjective range an observer uses is approximately constant across modalities, which would make the exponent partly a consequence of dynamic range rather than a property of the sense. **If that is right, the exponent is a joint property of the sense and the experiment, and reporting it without the range is reporting half a number.**
- **The exponent depends on the modulus and the anchor.** A modulus at one end of the range pulls the fitted exponent toward the anchor; free-modulus and mid-range designs reduce but do not eliminate this.
- **The regression effect.** Observers under-use the extremes of whatever they are asked to adjust, so magnitude estimation and magnitude production yield systematically different exponents that bracket the true value. The standard repair is to run both and take the geometric mean; a single-direction study must say which direction it ran.
- **Sequential dependencies.** A response is influenced by the preceding stimulus and the preceding response, which is why order is randomised and replicates are separated.
- **Number-use habits.** Observers prefer round numbers, powers of ten and small integers. Poulton [1968] catalogued the biases in magnitude estimation at length and Poulton [1979] is the mature statement; the catalogue has not been refuted, it has been managed by better designs.
- **Individual differences are large and real.** Exponents vary across observers by a substantial fraction of their value, and the individual exponents are themselves reasonably stable on retest, so the population figure is a central tendency of a genuine distribution rather than an estimate of a single constant [Stevens, 1975].

  Report an exponent as `a = value (method, range, modulus, instruction, n)`. An exponent quoted bare is a number without a referent. The values in Section 3.4 are reproduced as **nominal** published figures for orientation and comparison and are **not** calibration constants for any application without re-measurement under that application's own conditions.

None of this touches the existence proof, which rests on the structure of the results — log-log linearity, within-observer stability, cross-modality coherence, and the cross-method convergence with the decibel convention — rather than on any single exponent. What the

limits establish is that the measurement has an error budget, larger than physics is used to and smaller than the human sciences usually admit.

Four things the existence proof does not license, stated with the same emphasis as the proof. **(i)** It does not license the conclusion that all subjective quantities can be measured this way: loudness has a physical correlate that can be varied precisely and grief does not. The proof establishes that the wall between inner magnitude and instrumentation is not solid; it says nothing about where else it can be crossed. **(ii)** It does not license treating rating scales as ratio scales. **(iii)** It does not license the exponents as constants of nature. **(iv)** It does not settle whether the number an observer emits is "the sensation"; it settles that the numbers behave as a measurement system behaves.

And it does not answer the strongest objection to the whole enterprise. Michell [1997] argues that psychology systematically fails to test whether its attributes are quantitative at all before measuring them, and that the failure is not incidental but constitutive of the discipline's practice [Michell, 1999]. The remedy is conjoint measurement — the formal test of whether an attribute has additive quantitative structure [Luce and Tukey, 1964, Krantz et al., 1971]. **The conjoint conditions have not been tested for subjective magnitude in this paper, nor for u, nor for comfort.** Section 12.3's ratio-invariance test is this framework's own attempt at a licence and Section 13.2 states exactly what it does and does not establish. Relatedly, if u is read as a similarity measure over people it cannot be assumed to behave metrically, since psychological similarity is asymmetric and violates the triangle inequality [Tversky, 1977]; the nearest existing result to the framework's claim that surplus is a function of dissimilarity is the exponential law of generalisation over a psychological distance [Shepard, 1987], and it is a different function.

20.10 10 Specifications

Numbered, checkable requirements. They are renumbered S1 to S36 from the three source registers, in three groups. A claim that fails any of them is not admitted to a design document or to a decision.

Group I — the register (S1 to S12).

S1 — Declare the register. Any quantity used in design shall be labelled as perceived logarithmically, perceived linearly, or unknown. **Where it is unknown, no perceptual claim shall be made about it.**

S2 — Convert before promising. Any commitment about experience shall be converted to world units using a stated exponent before it is agreed, and the exponent and its source shall be recorded.

S3 — Ratios in budgets, never differences. Design budgets for perceived quantities shall be stated as ratios or in log units. A budget stated as an absolute difference in a logarithmically perceived quantity shall be rejected as unspecified.

S4 — Geometric mean for central tendency. Reports of a perceived quantity shall give the geometric mean. The arithmetic mean of ratios shall not be reported.

S5 — Log axes. Charts of perceived quantities shall use logarithmic axes, labelled as such.

S6 — Equal-step ladders shall be geometric, and the common ratio shall be stated.

S7 — Exponents cited with their method. Any exponent used shall be reported with its modality, its stimulus range, and the fact that it is a constant of a procedure. An exponent quoted bare shall be treated as unsourced.

S8 — Do not divide Fechnerian percepts. Where the percept scale is Fechnerian, only differences shall be taken. Ratios of percepts require a magnitude-estimation scale and shall state that they have one.

S9 — Do not sum across modalities without a declared weighting. Perceived changes in different modalities are in different units. They shall not be added without an explicit stated exchange rate, and the rate shall be badged as a design decision rather than a measurement.

S10 — Threshold declared. Any use of a logarithmic model shall state the reference `I_0` or threshold `u_0` in use. **A logarithm with an undeclared reference is not a measurement.**

S11 — Reference shifts disclosed. Where a design achieves its effect by moving the reference rather than the quantity, that shall be disclosable to the occupant without embarrassment. If it could not be, it shall not be built.

S12 — Concavity reported as a ratio requirement. Where a proposal is rejected on grounds of diminishing returns, the rejection shall state the ratio increment that would have been required. "Diminishing returns" without a number is not a finding.

Group II — the postulate method (S13 to S25).

S13 — Badge in the same sentence. Every assertion that is not a theorem carries one of the six badges of Section 8.3 in the sentence in which it is first stated. No footnoted badges.

S14 — Date every postulate. A postulate without a date cannot be distinguished from a post hoc fit and shall be treated as one.

S15 — Forbiddance list before observation, in operational terms. The list is part of the postulate; a postulate cited without its list is cited incompletely.

S16 — Executable operational definitions. Every term in the forbiddance list must be applicable by a person who does not know the hypothesis, yielding the same classification.

S17 — Asymmetry check. For each derived consequence, write the consequence the *negation* would produce. If they are the same, delete the consequence; it is not evidence.

S18 — Minimality audit at every revision. Delete any clause that forbids nothing and record the deletion. A postulate's word count shall not grow monotonically across revisions.

S19 — Amendment ledger, maintained. One row per revision, forbiddance-added and forbiddance-removed columns filled. Three consecutive net-negative revisions triggers a mandatory abandonment review.

S20 — Time-ordered split and freeze. Split by time, fit on the earlier portion, write the parameters down, do not refit.

S21 — Four baselines, all reported. Persistence, direct input, matched autoregression, shuffled input. The headline is the minimum of the four skill scores.

S22 — Scope declaration. State whether the postulate is single-subject, group-indexed or general. **A single-subject postulate may not be stated in the third person.**

S23 — Scale-type declaration. State the scale type of every measured term. If a formula divides two terms, both must be ratio-scaled with a true zero, or the formula is demoted to a heuristic ordering in print, visibly [Curzi, 2026m].

S24 — Publication of survival, not of truth. Report as "survived `n` rounds against the following forbiddance list", with the list attached. Never report a postulate as confirmed, established or proven.

S25 — Decline path. Any protocol asking another person to supply readings shall state how they decline, and declining shall cost them nothing.

Group III — magnitude estimation (S26 to S36).

S26 — Ratio instruction, verbatim. Any protocol claiming ratio-scale subjective data shall publish its exact instruction text. Instructions mentioning a bounded range, named categories or integer-only responses do not yield ratio data and shall be reported as ordinal.

S27 — Geometric spacing, minimum 7 levels, minimum 1.5 decades. State the range.

S28 — Geometric mean aggregation. Arithmetic means of magnitude estimates are prohibited.

S29 — Log-log fit with residual inspection. Report slope, intercept, `R^2` and the residual pattern. Systematic curvature is reported, not smoothed.

S30 — Both directions, or declare one. Run estimation and production and report the geometric mean of the exponents, or run one and state which, with the direction of the expected bias.

S31 — Ratio-invariance test. Include equal-physical-ratio pairs at two absolute levels and report whether the response ratios matched. **A protocol without this test may not claim ratio**

scale.

S32 — Full exponent citation format: `a = value (method, range, modulus, instruction, n).`

S33 — Units before exponents. State the physical quantity the exponent is taken on — pressure versus intensity, luminance versus illuminance, concentration versus mass. Pressure versus intensity alone is a factor-of-two error in the exponent.

S34 — Upper-bound labelling of step counts. `N_jnd` is reported as an upper bound on usable categories, with the identification-task caveat attached.

S35 — Re-measure before calibrating. Published exponents orient a design; they do not calibrate one.

S36 — Decline path. Participants may stop at any point without giving a reason and without cost, and this is stated to them before the first trial.

20.11 11 Failure modes

Twenty-seven named failure modes, renumbered `FM1` to `FM27` from the three source registers. Each carries what it damages, how it is recognised, and which specification repairs it.

20.11.1 11.1 Faults of the register (FM1 to FM10)

FM1 — The Linear Promise. A commitment about experience made in world percentages: "twenty per cent brighter", "twice as fast", "half the noise". *Damages:* the relationship and the budget — in brightness, `+20%` delivered is `+6.2%` received, 31 per cent of the promise. *Signature:* a percentage in a specification with no exponent anywhere in the document. *Repair:* S2. The conversion takes one line and is the highest-value single line in this paper.

FM2 — The Arithmetic Mean of Ratios. Percentage changes averaged arithmetically. *Damages:* the number, always upward, by Jensen's inequality. *Signature:* `+50%` and `−50%` averaging to zero when the true per-step factor is `0.866025`, that is `−13.40%`. *Repair:* S4.

FM3 — The Linear Ladder. Equal absolute increments intended to feel equal. *Damages:* the perceived step decays monotonically — in Section 7.4's example the last step is 0.1625 nats against the first's 0.3185, barely half. *Signature:* unequal `ln` steps. *Repair:* S6.

FM4 — Diminishing Returns Read as Failure. A programme delivering constant ratio improvement is reported as decelerating and cancelled. *Damages:* the programme, and the credibility of the next one. *Signature:* **if the log is a straight line, the programme is delivering constantly and the chart is lying.** *Repair:* S5 and S12.

FM5 — The Undeclared Reference. A logarithm taken with no stated reference, so the numbers are offsets from an unknown origin. *Damages:* comparability; two teams' numbers cannot be combined and neither knows it. *Signature:* ask what value gives zero; if nobody can say, S10 has failed. *Repair:* declare `I_0`, or in the surplus law declare `N`, since `u_0 = 1/(N−1)`.

FM6 — Nats Added Across Modalities. Perceived improvements in lighting, acoustics and thermal comfort summed into one figure. *Damages:* the figure. *Signature:* a single "perceived improvement" number covering more than one sense with no exchange rate. *Repair:* S9.

FM7 — The Dimensioned Logarithm. `ln(I)` computed where `I` carries units — `ln(500 lux)`. *Damages:* dimensional coherence outright; **the argument of a logarithm must be dimensionless.** *Signature:* a log whose argument has a unit attached. *Repair:* normalise first, `ln(I/I_0)`, always. Section 14.1 gives the audit that makes this checkable.

FM8 — The Expansive Modality Treated as Compressive. *Damages:* safety, potentially: at `a = 3.5` a twenty-two per cent increase doubles the sensation, and a margin computed on compressive assumptions is far too small. *Signature:* no exponent looked up. *Repair:* S7. **The sign of the design advice depends on whether `a` is above or below 1 and there is no default.**

FM9 — Percepts Divided. Two Fechnerian percept values divided and reported as "twice as much". *Damages:* the claim; an interval scale has no true zero and does not support division. *Signature:* a ratio of percepts with no magnitude-estimation procedure behind it. *Repair:* S8.

FM10 — The Reference Shift Sold as a Quantity Increase. Experience improved by moving `I_0` and reported as an improvement in `I`. *Damages:* honesty, and the next measurement, which will not reproduce. *Signature:* world quantity unchanged, occupant reports improvement, report claims the quantity improved. *Repair:* S11. **Reference shifting is a legitimate technique and a dishonest claim, and the difference is entirely in the reporting.**

20.11.2 11.2 Faults of postulation (FM11 to FM19)

FM11 — The Unbadged Postulate. *Damages:* the reader's ability to calibrate; an unbadged postulate accumulates authority it never earned and becomes load-bearing without a load rating. *Signature:* a statement about an unobservable, in the indicative, with no test attached — often recognisable by a citation chain whose earliest reference is a proposal and whose latest is a premise. *Repair:* badge it retroactively and audit what was built on it. If the audit shows the structure above cannot be supported by a postulate, the structure comes down, not the badge.

FM12 — The Retreating Postulate. *Damages:* content; `H_forbid` shrinks toward zero while apparent sophistication grows. *Signature:* every disconfirmation met with a qualification; nobody can state the postulate's current form in one sentence. *Detection:* the amendment ledger, with net forbiddance negative across revisions. *Repair:* abandon and restate from scratch, dated today. **The amendment rule [HOUSE RULE]:** an amendment is legitimate if and only if the amended postulate forbids at least one observation the original did not, and the new forbiddance is recorded, dated and tested in a subsequent round on fresh data. An amendment that only removes a forbiddance is a retreat and is recorded as one. The ledger has one row per revision — date, version, clause added or removed, forbiddance added, forbiddance removed, net — and **a programme whose net column is negative for three consecutive revisions is degenerating** [Lakatos, 1970].

FM13 — The Postulate That Forbids Nothing. *Damages:* the entire enterprise; it is not a weak postulate, it is not a postulate. *Signature:* every outcome described in advance as consistent; the proponent delighted by all results. *Detection:* hand it to someone hostile and competent and ask them to design an experiment to kill it; `q = 1` in Section 8.6. *Repair:* demote to metaphor and stop citing it as support.

FM14 — The Circular Instrument. *Damages:* identifiability, and silently everything fitted with it. *Signature:* the instrument measuring the outcome is derived from the same self-report that generated the postulate, so the postulate is tested against a transformation of itself. *Detection:* can the outcome proxy be shown not to be a repackaging of the input? Curzi [2026i]'s repackaging audit is the institutional-scope version, and it records that a repackaged proxy's single-week jump of 31.1 units was 38.5 times the maximum admissible increment. *Repair:* separate the channels — outcome by behavioural trace, or by a different-modality report taken hours later and recorded without seeing the earlier one.

FM15 — The Borrowed Ceiling. *Damages:* the reader's calibration, by importing authority across a boundary it does not cross. *Signature:* "psychophysics established that subjective quantities are measurable, therefore my subjective quantity is measurable." Section 9 establishes something narrower: that *certain* subjective magnitudes, under *specific* procedures, yield ratio-supporting numbers that replicate. **The existence proof licenses the attempt, not the result.** *Detection:* look for the word "therefore" spanning two different quantities. *Repair:* run the procedure on the quantity in question.

FM16 — The Post Hoc Fit Reported as a Test. *Damages:* falsification; the reported evidence has no evidential value. *Signature:* as many free parameters as data points, a beautiful fit, no held-out set, no frozen parameters. *Detection:* ask for the date of the parameter values and the date of the data, and for the held-out error [Gelman and Loken, 2013]. *Repair:* the freeze. If the record is too short to split, say so and report the fit as exploratory, which is honest and useful.

FM17 — The Identity Postulate. *Damages:* everything, and it is the hardest to repair, because the cost of falsification is not intellectual. *Signature:* a postulate about oneself whose

abandonment would require abandoning a self-image; the forbiddance list, honestly written, contains observations the person cannot afford to make, so it is never written or is written too vaguely to bind. *Detection:* ask what changes if the postulate is false; **if the answer includes who you are rather than what you do, the postulate is identity-loaded and its evidence must be discounted accordingly.** *Repair:* delegate the classification to someone blind to the hypothesis, and precommit the change of practice that follows from the postulate's death.

FM18 — The Consequence Nobody Recorded. *Damages:* the entire test, silently; the most frequent failure in practice, and it looks like nothing. *Signature:* postulate stated, consequences derived, observations made, and no artefact exists — the person "noticed" the results, and recall is then reconstructed under the influence of the postulate. *Detection:* ask for the file; if the answer is a memory, there is no test. *Repair:* any dated artefact. **An observation not recorded contributes nothing to the stock, however vivid it was.**

FM19 — The Missing Shuffled Control. *Damages:* the causal claim specifically, while leaving the predictive claim intact and thereby disguising the damage. *Signature:* a model that beats persistence reported as evidence for the postulated mechanism, with no shuffled-input comparison. *Detection:* look for the fourth baseline; its absence is the signature. *Repair:* run it. Section 12.2 is where it does the real work. **The control is not a formality:** the corpus measured its shuffled-input control passing 98.35 per cent of the time against a third variable driving both series, while detecting only 48.17 per cent in the clean case, so **the empirical audit cannot see repackaging** [Curzi, 2026i].

20.11.3 11.3 Faults of magnitude measurement (FM20 to FM27)

FM20 — The Rating Scale Wearing a Ratio Scale's Clothes. *Damages:* the scale type, and therefore every division and multiplication downstream. *Signature:* a one-to-seven item whose means are divided, multiplied or compared as ratios. *Detection:* **look for a bound.** A response scale with a ceiling is ordinal at best, because a ratio scale has no maximum [Stevens, 1946, Likert, 1932]. *Repair:* magnitude estimation if the quantity supports it; otherwise demote every formula containing the quantity to a heuristic ordering, visibly.

FM21 — The Exponent Quoted Bare. *Damages:* reproducibility. *Signature:* "the exponent for X is 1.4", with no range, modulus, instruction or direction. *Repair:* S32, or drop the number.

FM22 — The Pressure/Intensity Confusion. *Damages:* dimensions, by a factor of exactly two in the exponent and therefore four in every conversion. *Signature:* loudness computations mixing an exponent of 0.67 with intensity data, or 0.33 with pressure data. *Detection:* Section 9.4's cross-check; if the perceived doubling does not land within a decibel or so of 10 dB, the domain is wrong. *Repair:* state the domain before the exponent and convert with `a_intensity = a_pressure / 2`.

FM23 — The Linear Promise in a Logarithmic Quantity. The measurement-side form of FM1. *Signature:* "we improved the metric by 40 per cent and users did not notice." At `a = 0.33` a 40 per cent physical increase is a perceived increase of `1.4^0.33 = 1.1174`, that is 11.7 per cent, which for many observers is close to the just-noticeable difference for that dimension. *Repair:* compute the physical ratio needed for a change of at least one just-noticeable difference, and price that.

FM24 — The Arithmetic Mean of Ratio Data. *Damages:* the aggregate estimate, systematically upward and unevenly across conditions. *Signature:* group means of magnitude estimates computed with an arithmetic mean; one participant who used numbers ten times larger dominates the result. *Detection:* compare arithmetic and geometric means. *Repair:* geometric mean, equivalently the mean of the logarithms, which is the natural operation in the register.

FM25 — The Unrandomised Ladder. *Damages:* the exponent, via sequential dependency, usually inflating the fit. *Signature:* stimuli presented in ascending or descending order — beautiful `R^2`, wrong slope. *Repair:* randomise, replicate with separation, and compare ascending and descending sub-blocks.

FM26 — The Step Count Treated as a Category Count. *Damages:* every ceiling computed from it; `N_jnd` overstates usable categories, sometimes by an order of magnitude. *Signature:* "there are 283 discriminable levels, therefore the display should have 283 levels." *Detection:* run an absolute-identification task [Miller, 1956]. *Repair:* S34.

FM27 — The Two-Decade Conclusion. *Damages:* model selection; two fundamentally different models both fit. *Signature:* a functional form declared on data spanning one or two decades, `R^2 > 0.95`, and no competitor tested. *Detection:* Section 9.5's arithmetic — over two decades a logarithmic law captures 96.3 per cent of a power function's variance, so `R^2 = 0.96` distinguishes nothing. *Repair:* extend the range until the candidates disagree, or find a qualitative discriminator, as Stevens did with the expansive exponents. **This paper's own headline fit is exposed to FM27 and Section 12.3 says so.**

20.12 12 Results: three worked cases

All three cases are constructed. No dataset was collected for this paper; the reading room is a design case with stipulated dimensions and complaints, the postulate case is a stipulated series of readings chosen to demonstrate the protocol's arithmetic, and the magnitude-estimation series is a stipulated worked case. **None of them is evidence about any room, any person or any modality**, and Section 15 repeats it. What they demonstrate is that the procedures run, that they produce numbers a reader can re-derive, and — in the second case — that they report their own refutation.

20.12.1 12.1 The reading room, budgeted in ratios

A public library reading room, three complaints and one constraint.

```
Room:        12.0 m x 8.0 m x 3.5 m,  V = 336.0 m^3
Complaint 1: "the edges are gloomy"      perimeter 200 lux, core tables 500 lux
Complaint 2: "you can hear everything"   8 simultaneous quiet talkers; RT60 measured 1.20 s
Complaint 3: "it feels harsh"            deferred — not a magnitude problem
Constraint:  lighting energy must not rise more than 10 per cent
```

Step 1 — convert every complaint into a ratio. No absolute number is actionable until it is a ratio against a reference (S3, S10).

Gloom. Perimeter against core: world ratio `200/500 = 0.400000`, `ln = −0.916291` nats; perceived `0.400^0.33 = 0.7391`, so the perimeter looks **26.09 per cent dimmer** than the core; in percept units `0.33 x (−0.916291) = −0.302376` nats. **A sixty per cent reduction in illuminance is experienced as a twenty-six per cent reduction in brightness.** The register has already absorbed most of the physical difference, which is why the complaint is "gloomy" rather than "dark" — and why the fix will cost more than the complaint sounds.

Noise. Eight incoherent talkers against a target of four. Incoherent sources add in power, so halving the count divides pressure by `sqrt(2)`: pressure ratio `0.707107`, that is `−3.010 dB`; perceived `0.707107^0.67 = 0.7928`, **20.72 per cent quieter**; in percept units `−0.232204` nats. **Halving the number of people talking makes the room about a fifth quieter, not half.**

Reverberation. Sabine's law, `RT60 = 0.161 V / A` [Sabine, 1922]:

```
measured RT60 = 1.20 s  =>  A = 0.161 x 336.0 / 1.20 = 45.08 m^2 sabins
target   RT60 = 0.60 s  =>  A = 0.161 x 336.0 / 0.60 = 90.16 m^2 sabins
required: A x 2.000 −− exactly double the absorption
```

Halving reverberation time requires doubling absorption exactly, because Sabine's law is itself a reciprocal — a ratio law. The perceptual conversion stops here:

NOT CONVERTED. There is no verified exponent for perceived reverberance in this framework's table. The world ratio is `ln(0.5) = −0.693147` nats and **it shall not be multiplied by an exponent that has not been measured** (S7). The acoustic change is reported in world units and left there. The alternative would have been to invent a plausible number.

Step 2 — price the constraint before designing to it.

illuminance change	energy change	perceived brightness change
500 to 450 lux	-10.0%	**-3.42%**
500 to 400 lux	-20.0%	**-7.10%**
500 to 350 lux	-30.0%	-11.10%
500 to 250 lux	-50.0%	-20.45%
200 to 300 lux	+50.0% (perimeter circuit)	**+14.32%**
200 to 400 lux	+100.0% (perimeter circuit)	+25.70%
500 to 4 085 lux	+717%	+100% (a perceived doubling)

The bottom row is the anchor: **doubling the apparent brightness of the core would take 4 085 lux and eight times the energy. It is not on the table and never was.** The fourth row solves the problem: a twenty per cent energy cut on the core costs 7.10 per cent of perceived brightness, and the fifth row spends that saving where it buys three times as much, because the perimeter starts lower and the register is steeper in ratio terms near the bottom of the range.

Step 3 — the design and its energy arithmetic. With the core carrying 60 per cent of the luminaires and the perimeter 40 per cent:

```
core:      500 -> 400 lux   (x0.80)
perimeter: 200 -> 300 lux   (x1.50)
energy = 0.60 x 0.80 + 0.40 x 1.50 = 0.4800 + 0.6000 = 1.0800   ->  +8.00%
```

Inside the ten per cent cap. The perceptual result:

location	before	after	perceived change
Core tables	500 lux	400 lux	-7.10%
Perimeter	200 lux	300 lux	**+14.32%**
Core-to-perimeter gap	-26.09%	**-9.06%**	gap cut by 65%

The complaint was the gap, and the gap closed from 26.09 to 9.06 per cent for eight per cent more energy, by taking brightness away from the place that had most of it. In world units the move looks perverse. In register units it is straightforwardly correct, because 100 lux removed from 500 costs 7 per cent of experience and 100 lux added to 200 buys 14 per cent. **The register is why the two 100-lux moves are not the same move.**

Step 4 — the perceived-improvement budget, per modality, kept strictly separate as S9 requires:

```
BRIGHTNESS (a = 0.33)
    core       500 -> 400 lux : 0.33 x ln(0.800000) = −0.073637 nats   (−7.10%)
    perimeter  200 -> 300 lux : 0.33 x ln(1.500000) = +0.133803 nats  (+14.32%)
    gap closure, the design objective: from −0.302376 to −0.094935 nats

LOUDNESS (a = 0.67)
```

```
    8 -> 4 talkers            : 0.67 x ln(0.707107) = -0.232204 nats  (-20.72%)

REVERBERATION
    RT60 1.20 -> 0.60 s       : world ln-ratio -0.693147 nats
                                ** NOT CONVERTED - no verified exponent **
```

These three blocks are not summed. A single "perceived improvement of 0.44 nats" would require an exchange rate between brightness nats and loudness nats; none has been measured; inventing one is FM6.

The source manuscript prints the post-change gap as `-0.094839` nats. The value is `0.33 x ln(300/400) = 0.33 x (-0.287682) = -0.094935`, which is also what the two component figures give when combined: `-0.302376 + 0.133803 + 0.073637 = -0.094936`. The corrected figure is used above. The percentage figures `-26.09%` and `-9.06%` are unaffected and reproduce exactly.

Step 5 — where the acoustic budget actually goes. The library's instinct was occupancy limits. Priced:

intervention	world change	perceived loudness
8 to 4 talkers	-3.010 dB	-20.72%
Pressure halved	-6.02 dB	-37.15%
-9.0 dB	x0.3548 pressure	**-50.05%**
-10.0 dB	x0.3162 pressure	-53.76%
-12.0 dB	x0.2512 pressure	-60.37%

Halving the perceived loudness of the room takes -8.99 dB, a factor of 0.3554 in pressure — the equivalent of going from eight talkers to one. Occupancy control cannot deliver it; the room would have to be nearly empty. The absorption route can: doubling `A` to 90.16 m^2 of sabins halves RT60, reducing the reverberant contribution at every listening position without removing a reader. **The library's plan was to reduce the number of people in a reading room; the arithmetic says the achievable version of that plan delivers a fifth of what the complaint asks for, and that ceiling and wall treatment delivers the rest without touching occupancy.**

Step 6 — what to promise. The library's draft notice said *"we are increasing lighting at the perimeter by 50 per cent."* Converted (S2): readers will experience about **14 per cent** more brightness at the perimeter and about **7 per cent** less at the tables, and the difference between the two zones — which is what they complained about — will be about a third of what it was. **A promise of 50 per cent, kept exactly in lux, would have been received as a failure, because the promise was made in one register and the reading is taken in another.**

Step 7 — what would falsify the design's predictions. Preregistered before the work: post-occupancy magnitude estimation at the perimeter and at the core, by readers who did not participate in the design, against the prediction of **`+14.32%`** and **`-7.10%`** with the gap closing to `-9.06%`. **If the measured perimeter improvement falls outside `±0.05` nats of the prediction, the brightness exponent of 0.33 does not describe this population in this room, and the exponent — not the design method — is what gets revised.** The Lens Theorem is not at risk from this measurement because it is a theorem; the exponent is at risk, and should be, because it is a measurement.

What this case does not establish. The relation between illuminance, zone contrast and occupant complaint is assumed, not measured. This framework's bibliography contains no entry establishing a quantitative illuminance-to-complaint relation for a reading room; the nearest support is the design literature's general claim that the distribution of light within a room conditions how it is used and experienced [Alexander et al., 1977], which is a design-literature claim and not a measured relation. **The premise of the case is therefore a modelling stipulation and is badged as one**, and Section 15 records it.

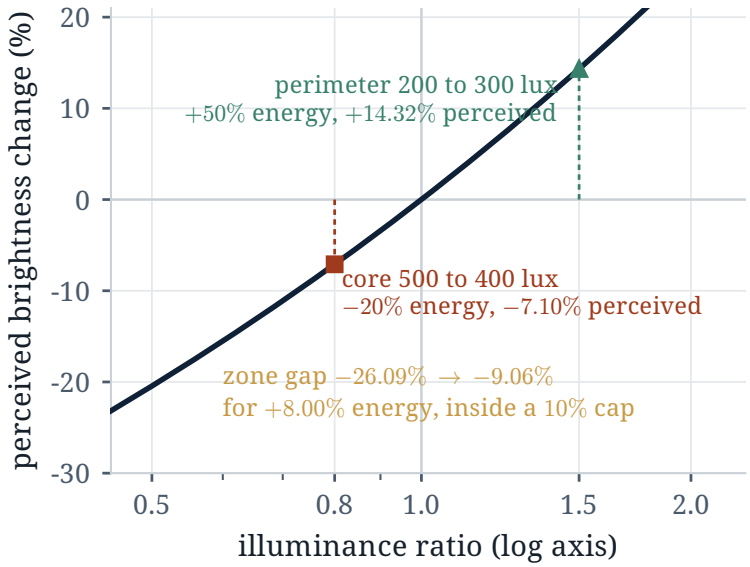


Figure 20.3: **A 20 per cent energy cut where light is plentiful buys more perceived brightness than it costs, once both moves are read in the register.** Core 500 to 400 lux costs 7.10 per cent of perceived brightness while perimeter 200 to 300 lux buys 14.32 per cent, closing the zone gap from 26.09 to 9.06 per cent for 8.00 per cent more energy, inside the declared 10 per cent cap. Modelled from the stipulated illuminances and a cited brightness exponent of 0.33; generated by `figs/src/CCP-020_fig03.py`.

20.12.2 12.2 A postulate carried to a held-out test, and refuted on one of four consequences

One observer, one postulate, carried from a native sentence to a frozen held-out test and an amendment. **Every number below is computed from a stipulated worked case chosen to demonstrate the arithmetic of the protocol. The readings are not empirical data about any person and no claim is made about anybody's mornings.**

Step A — the belief in its native form. *"I am sharpest in the morning and useless after lunch."* Two unobservables, one implied cause, one implied structure, no quantities and no forbiddance. As stated it is compatible with every possible observation, so `q = 1` and `H_forbid = 0` — the starting condition of nearly every belief anyone holds about their own functioning, and FM13 exactly.

Steps B and C — the named quantity and the postulate. The hidden quantity is a capacity for one class of work. Name it `R`; restrict the class to **novel analytic problem-solving**, operationally *tasks requiring a solution method not previously used by this observer on this problem type*.

> **The reservoir postulate [POSTULATE, dated at the top of the record].** There is a single scalar capacity `R` for novel analytic work. `R` is depleted by time spent on that class of work and by nothing else. `R` recovers fully overnight and is **not** replenished by breaks within the day. Observed performance on a novel analytic task is monotone increasing in `R` at the moment the task is attempted.

Each clause does work — "by nothing else" forbids depletion by meetings and administrative work; "not replenished by breaks" forbids mid-day recovery; "recovers fully overnight" forbids carry-over across days — so S18 is satisfied at birth.

Step D — four consequences, one uncomfortable.

```
D1  Within a working day, scores on the analytic class decline monotonically across blocks.
```

```
D2  On days containing a long midday break, the block after the break falls ON the same
    decline curve. It does not rise.                          <-- UNCOMFORTABLE
D3  The control class (routine administrative work) shows no comparable decline.
D4  A day starting later shows the same decline SHAPE, shifted, not compressed: the decline
    tracks time-on-task, not clock time.
```

D2 carries the content, because the observer takes breaks and would like them to be doing something. **A consequence one would be pleased to see falsified is exactly the consequence that carries information** (S17).

Step E — the forbiddance list, written before observing.

```
F1  kills D1: a non-monotone block sequence, in the mean over >= 20 days.
F2  kills D2: block 3 on break days exceeding the no-break decline prediction by more than
              0.5 pooled standard deviations.
F3  kills D3: the control class declining at more than one third the analytic rate.
F4  kills D4: late-start days showing a decline slope differing from early-start days by
              more than 25%.
```

Step F — instrument and design. Four 75-minute blocks per day; at the end of each, a magnitude estimate of the block's analytic output against a personal modulus of 100 assigned to a pre-selected reference block (Section 9.6's method, used rather than justified). Forty working days: twenty-two with no long break, eighteen with a 90-minute break between blocks 2 and 3.

condition	block 1	block 2	block 3	block 4	days
no long break	100.0	88.0	74.0	63.0	22
(sd)	11.0	12.0	13.0	12.0	
90-min midday break	100.0	87.0	95.0	81.0	18
(sd)	10.0	11.0	12.0	13.0	
control: administrative	100.0	98.0	97.0	96.0	22

Computation. Linear fits on the block index:

```
no-break analytic:   slope = -12.5000 per block,  intercept 112.5000,  R^2 = 0.99808
break-day analytic:  slope =  -4.9000 per block,  intercept 103.0000,  R^2 = 0.56428
control (admin):     slope =  -1.3000 per block
```

In log-ratios against block 1, which is the register the framework works in:

```
ln(88/100) = -0.1278    ln(74/100) = -0.3011    ln(63/100) = -0.4620
log-slope = -0.1559 nats per block,  R^2 = 0.99639
implied per-block retention exp(-0.1559) = 0.8556
implied per-block loss                   = 0.1444
```

That last line is the first substantive result. The no-break decline is very well described by a constant fractional loss per block — a leaky integrator with no input, `Q_{t+1} = (1-delta)Q_t` with `delta = 0.1444` per 75-minute block [Curzi, 2026a]. **The postulate predicted only monotonicity, not an exponential. The exponential form is a free finding**, and free findings are how postulates get sharpened. It is also the point at which S17's asymmetry check earns its place: monotone decline is predicted by the postulate's negation too, whereas a constant fractional loss is not, so it is the exponential and not the monotonicity that carries information.

The D2 test — the uncomfortable one.

```
frozen no-break decay applied forward from break-day block 1:
    predicted block 3 = 100.0 x exp(-0.15594 x 2) = 73.207
observed block 3 on break days                    = 95.000
difference                                        = 21.793
pooled sd (n1 = 22, n2 = 18)                      = 12.5625
standardised effect size                          =  1.7347
t                                                 =  5.4583   (df = 38)
```

The tolerance declared in F2 was 0.5 pooled standard deviations. **The observed excess is 1.73 pooled standard deviations, more than three times the tolerance, in the forbidden direction.**

> **D2 is falsified. The reservoir postulate, as stated, is dead.**

That sentence is reported here in exactly the language a confirmation would have received. The other three consequences:

```
D1  supported:  monotone decline in the no-break condition, R^2 = 0.99808
D3  supported:  control slope -1.3 against analytic -12.5, a factor of 9.615; F3's
                threshold was a factor of 3. Passed with margin.
D4  not tested: start time was not varied. Recorded as UNTESTED, not as supported.
```

D4's row is the one to imitate. An untested consequence is reported as untested and never folded into the supported column because the postulate "did well overall".

The amendment, done legitimately. The clause "not replenished by breaks" is dead. The amendment must add forbiddance, not merely remove it, so the replacement clause is quantitative:

> **The reservoir postulate, v2 [POSTULATE, dated at amendment].** As v1, except: `R` is partially replenished by breaks, and the replenishment is a saturating function of break length, so that doubling break length yields less than double the recovery.

version	clause changed	forbiddance ADDED	forbiddance REMOVED	net
v2	breaks replenish, saturating	ordering and concavity across three break lengths	"no recovery"	+

The new forbiddance is specific and testable on fresh data: with 20-, 45- and 90-minute breaks, recovery must be ordered `r(20) < r(45) < r(90)` **and** concave, `r(90) − r(45) < r(45) − r(20)`. That forbids a linear response and it forbids a step response, so the net column is positive and the amendment is progressive rather than a retreat (S19, FM12).

The freeze and the four baselines. For the follow-up round the observer wants a predictive model rather than a contrast, using the framework's minimal estimation form at single-observer scope, `Y_{t+1} = a Y_t + b S_t + noise` with `a = 1 − delta` and `b = eta` [Curzi, 2026i]. The record is 40 days, split by time at day 24. **The series is synthetic, generated by the stated rule `Y_{t+1} = 0.72 Y_t + 30 S_t + N(0, 2.0)` with a fixed seed, and is present to demonstrate the arithmetic of the freeze, not to report a result about anyone.**

```
FROZEN PARAMETERS, fitted on days 1-24 only:
    a = 0.7018  ->  delta = 0.2982
    b = 31.3083 ->  eta   = 31.3083
(true generating values a = 0.72, b = 30; recovery is a check that the fitting procedure
 works, not a result)
```

held-out model, days 25-40, frozen parameters	RMSE	skill against persistence
model	1.5164	0.8473
persistence	9.9295	0.0000
direct input	22.8039	-1.2966
matched AR(1), `c = 0.9838`	9.7719	0.0159
shuffled-input control	13.2567	-0.3351

Read in the order that matters. The model beats persistence by a wide margin, which by itself means little. It beats the matched AR(1), which is the first result with content: **a one-parameter memory model with no theory in it scores 0.0159 against persistence, essentially nothing, so the improvement is not coming from autocorrelation alone.** The direct-input baseline is catastrophically bad, so the output is not a rescaled copy of the input and there is genuine storage. **And then the fourth row, which is what the whole protocol exists to produce.** With the same frozen parameters and the input series permuted, held-out RMSE rises from 1.5164 to 13.2567, a factor of 8.74, and the skill score goes negative. **Performance depends on the specific temporal alignment of inputs with outputs, which is what a storage claim asserts and what a model that had merely learned the series' shape would not show.** Had the shuffled RMSE come back near 1.5, the exercise would have been dead — and dead invisibly, because every other row looks excellent.

What was and was not established.

```
ESTABLISHED (as survival, not as truth):
  - a within-day monotone decline in the analytic class, well fitted by a constant fractional
    loss of 0.1444 per 75-minute block
  - the control class does not show it (a factor of 9.6 in slope)
  - the frozen storage model beats all four baselines on held-out data, including and
    especially the shuffled-input control

KILLED:
  - the clause "breaks do not replenish", at a standardised effect of 1.73, three times the
    pre-set tolerance

UNTESTED, and recorded as untested:
  - D4, the time-on-task versus clock-time distinction
  - the saturating-recovery clause of v2, which is the next round's business

NOT CLAIMED:
  - anything about mornings in general
  - anything about any other person                    [scope: single subject, S22]
  - any account of WHY the decline occurs; the postulate is about structure, not mechanism,
    and no mechanism was tested
```

Two limits on this case that no amount of held-out testing removes. The observer supplying the readings and the analyst fitting the model are the same person, which is FM14 in its hardest form, and the separation of channels prescribed there was not achieved: the outcome is a magnitude estimate given by the same person whose capacity is being postulated. And the reported effect sizes carry no confidence intervals, because the case is constructed; a real application owes them (Section 15).

20.12.3 12.3 The ratio-invariance test, and the model comparison it forces

A full magnitude estimation, computed. **The readings are a stipulated worked case chosen to demonstrate the arithmetic and the checks. They are not empirical data and no claim is**

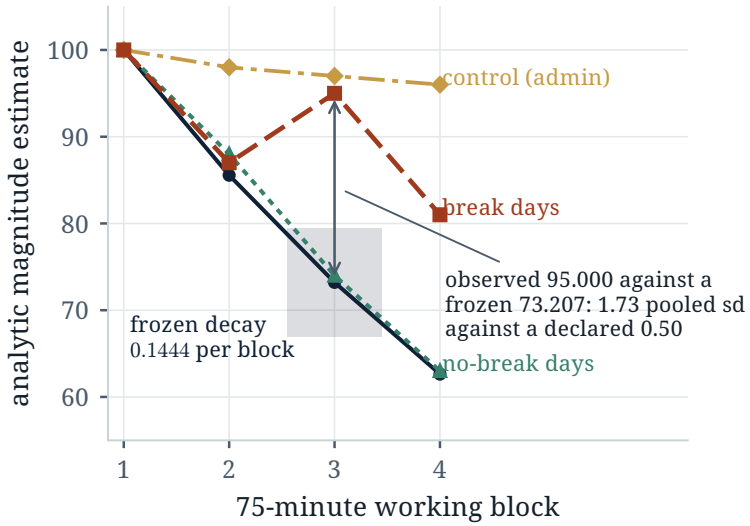


Figure 20.4: **A postulate reporting its own refutation.** The frozen no-break decay, a constant fractional loss of 0.1444 per 75-minute block, predicts 73.207 for block 3 on break days; the observed 95.000 exceeds it by 1.73 pooled standard deviations against a pre-set tolerance of 0.50, in the forbidden direction, so consequence D2 is falsified and the clause is amended in public. Modelled from the stipulated block means of this section; generated by `figs/src/CCP-020_fig04.py`.

made about any modality's true exponent. The procedure is Section 9.6's and the arithmetic is exact.

```
continuum    one physical intensity, varied alone
levels       8, geometrically spaced by factors of 2, from 0.25 to 32.0
             -> range R = 128, that is 2.107 decades              [S27 satisfied]
order        randomised per observer                              [FM25 avoided]
instruction  free ratio assignment, no bounds, non-integers permitted   [S26]
modulus      100 assigned to level 1.0, mid-range                 [S30 mid-anchor]
observers    20, two replicates per level, separated
aggregation  geometric mean across observers and replicates       [S28]
```

stimulus I	0.25	0.50	1.00	2.00	4.00	8.00	16.00	32.00
response	4.10	6.40	10.00	15.90	25.10	39.20	63.50	98.00

The log-log fit. Regressing `log10(response)` on `log10(stimulus)` over the eight points:

```
a       = 0.6569        the exponent
log10 k = 1.0043
k       = 10.0996
R^2     = 0.999932
```

I	observed	fitted	residual
0.25	4.10	4.06	+0.92%
0.50	6.40	6.41	-0.09%
1.00	10.00	10.10	-0.99%

I	observed	fitted	residual
2.00	15.90	15.92	-0.15%
4.00	25.10	25.11	-0.03%
8.00	39.20	39.58	-0.97%
16.00	63.50	62.41	+1.74%
32.00	98.00	98.40	-0.41%

Residuals under two per cent with no systematic curvature. **`R^2 = 0.999932` is reported here together with the warning that makes it interpretable: over two decades a fit statistic of this size does not by itself exclude a logarithmic law** (Section 9.5, FM27). The residual pattern and the tests below carry the conclusion, not the `R^2`.

The ratio-invariance test. S31 requires equal physical ratios at different absolute levels. The design contains two such pairs, both a factor of four:

```
psi(4.00) / psi(1.00)  =  25.10 / 10.00  =  2.5100
psi(32.0) / psi(8.00)  =  98.00 / 39.20  =  2.5000
predicted from the fitted exponent:  4^0.6569  =  2.4859
agreement between the two empirical ratios:  0.40 %
agreement with the prediction:               0.97 % and 0.57 %
```

Two stimulus pairs separated by a factor of eight in absolute level produce response ratios agreeing to within 0.40 per cent. This is the result that licenses calling the scale a ratio scale: an ordinal response with numerals attached has no reason to produce it. **It is a test the data could have failed and did not.**

And it is evidence for a power law, not for a logarithm. This must be said plainly, because the title of this paper points the other way. Under Fechner, equal stimulus ratios produce equal *differences* of percept, not equal *ratios*: `S(aI) − S(I) = k ln a`, independent of `I`. Ratio invariance of the response is the *power law's* invariant, `psi(bI)/psi(I) = b^a`. **A ratio-invariance test that succeeds is therefore a discriminating test, and on this series it discriminates against the logarithm.** Computed on the same eight points:

model	free parameters	SSE	`R^2` (response units)	AIC	BIC	LOO-CV RMSE
Power law `psi = k I^a`	2	1.5063	0.999802	-7.358	-7.120	**0.6071**
Fechner `S = A + B ln I`	2	1186.7312	0.844258	45.996	46.234	18.4747
Threshold Fechner `S = k ln(1 + I/I_0)`	2	41.7285	0.994524	19.214	19.452	—

All three carry the same number of free parameters, so the information criteria are directly comparable and the parameter count is paid for. **The power law is favoured by `Delta AIC = 26.57` over the better of the two logarithmic forms**, its leave-one-out cross-validated error is thirty times smaller than the plain logarithm's, and the plain logarithmic fit commits a qualitative error the power law does not: it predicts a response of `−10.54` at `I = 0.25`, a negative magnitude estimate, and places its zero at `I = 0.4512`, inside the presented stimulus range.

The ratio-invariance test separates them by itself, without any fit statistic:

	ratio at 4.00/1.00	ratio at 32.0/8.00	spread
observed	2.5100	2.5000	**0.40%**

	ratio at 4.00/1.00	ratio at 32.0/8.00	spread
power-law fit	2.4859	2.4859	0.00% (identically, by construction)
Fechner fit `A + B ln I`	2.7418	1.4821	**84.99%**
threshold Fechner `k ln(1+I/I_0)`	3.4481	2.3290	**48.05%**

A logarithmic register predicts that these two ratios differ by between forty-eight and eighty-five per cent; they differ by 0.40 per cent. On a constructed series, on one continuum, with one instruction, over 2.107 decades — but the direction of the result is unambiguous and it is reported here rather than left for a reader to compute.

Three things follow, and the third is the one that matters for this corpus. First, the scale-type licence sought by S31 is granted on this series. Second, the modelling conclusion is that the response is a power function of the stimulus, which is Stevens and not Fechner, exactly as Section 9.5 says the historical record found. Third — and this is where a reader could go wrong — **none of it damages the corpus's surplus functional, because that functional's logarithm does not rest on Fechner.** It rests on four axioms about interaction geometry, one of them a declared design axiom [Curzi, 2026], and on the independent dynamical warrant that the growth rate of a multiplicative process is logarithmic in the accumulating quantity [Peters and Gell-Mann, 2016]. Section 13.2 states the division of labour once more.

Interpretation, reported in the required format (S32):

```
a = 0.657  (magnitude estimation, free ratio instruction, range 0.25–32.0 = 2.107 decades,
            modulus 100 at mid-range level 1.0, geometric-mean aggregation,
            n = 20 observers x 2 replicates, single direction: estimation only, so the true
            value is expected to lie slightly ABOVE this figure by the regression effect)

world cost of a perceived doubling: 2^(1/a) = 2.8725 in stimulus units
                                  = ln 2 / a = 1.0552 nats = 0.4583 decades
conversion key: 1 nat of stimulus log-ratio -> 0.6569 nats of perceptual log-ratio
                                            -> 0.9477 bits          (x 1/ln 2)
```

That parenthesis is the discipline of Section 9 in one line. The number is not the result; the number plus the procedure is the result.

What was and was not established.

```
ESTABLISHED
  - a power relation between stimulus and response, residuals under 2%, no systematic
    curvature, and preferred over both logarithmic forms at equal parameter count by
    Delta AIC = 26.57 and by a factor of 30 in leave-one-out error
  - ratio invariance across an eightfold change in absolute level, to 0.40%, which the
    logarithmic forms predict at 48% to 85%
  - a conversion key of 0.657 nats per nat, usable in this framework's registers

NOT ESTABLISHED
  - that the exponent would survive a change of range          [Teghtsoonian effect]
  - that it would survive magnitude PRODUCTION                 [regression effect]
  - that the observers' numbers are "the sensation"            [Section 8.7]
  - that the attribute has additive quantitative structure at all; the conjoint conditions
    were not tested                                            [Section 9.9, Section 15]
  - anything at all about a different modality                 [FM15]

NEXT ROUND, preregistered
```

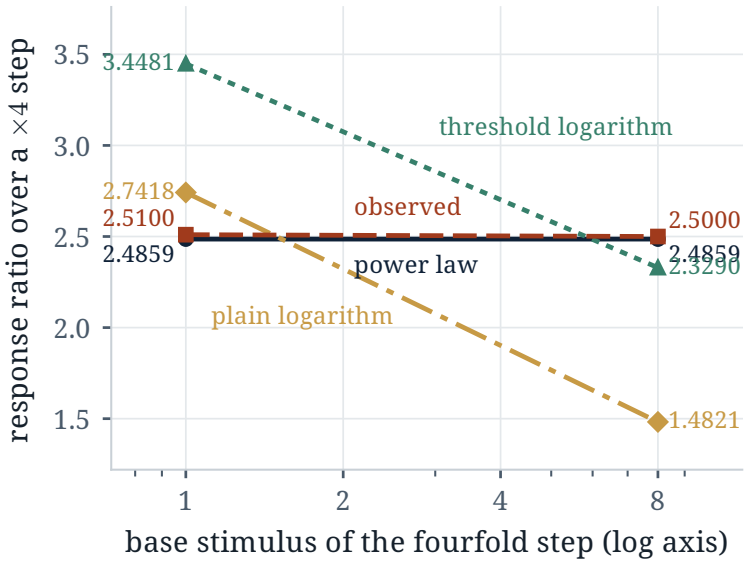


Figure 20.5: **The ratio-invariance test discriminates, and it discriminates against the logarithm.** The stipulated eight-point series is straight in log-log with residuals under two per cent; the fitted power law `psi = 10.0996 I^0.6569` reproduces both fourfold response ratios at 2.4859 against observed 2.5100 and 2.5000, while the best logarithmic form predicts ratios differing by 48 per cent across the same eightfold change of level. Computed for this paper from the stipulated series of Section 12.3; generated by `figs/src/CCP-020_fig05.py`.

```
- production direction, same stimuli; report the geometric mean of the two exponents
- a compressed-range replication (0.5 to 8.0) to measure the range dependence directly
  rather than inheriting it as a caveat
```

20.13 13 Discussion

20.13.1 13.1 What follows

The framework's logarithms are not a stylistic preference and they are not one claim. They are three claims with three different warrants, and separating them is the main practical contribution of this paper.

- **The Lens Theorem is algebra.** Given continuity and the conversion of multiplication into addition, the logarithm is forced and is unique up to a scale constant. Characterisation results of exactly this shape — which axioms force a logarithm — are a known genre [Aczél et al., 1974, Khinchin, 1957], and the framework's uniqueness argument for `F(u)` should be positioned against that genre rather than presented as novel.
- **The surplus functional's logarithm is axiomatic and dynamical.** It follows from four stated axioms, one of which is a declared design axiom [Curzi, 2026], and it is independently motivated by the fact that the growth rate of a multiplicative process is logarithmic in the accumulating quantity [Peters and Gell-Mann, 2016]. **Neither route needs an observer.**
- **The design advice is psychophysical, and only the design advice.** Section 7's numbers predict what a person will report. They need Fechner's register and a measured Stevens exponent, and they inherit every limitation in Section 9.9.

A second consequence is a reading of concavity that changes what a designer does with it. Because `F` is a logarithm of `(u + u_0)`, **diminishing returns are a constant ratio requirement seen from the world's coordinates**, and a proposal rejected for diminishing returns should be rejected with the ratio increment that would have been required rather than with the phrase (S12).

A third is that the storage law's world-coordinate reading makes explicit what its additive form conceals: **decay is taking a root, conversion is exponentiation, and cost is a divisor.** Every intuition about the stock being "added to" is an intuition about a count being multiplied, which is why the `ln N` ceiling bites so hard and why stored order is much harder to double than to establish.

A fourth is procedural. **The postulate method's value is not that it produces true postulates but that it produces informative deaths.** Section 12.2's postulate died on the consequence its author would least have liked to lose, at three times the pre-set tolerance, and was replaced by a version that forbids more. That is the mechanism by which a research programme stays progressive rather than degenerating [Lakatos, 1970], and it is why the specifications require the forbiddance list to be written before the observation and the ledger to record the net.

20.13.2 13.2 What does not follow

The Fechner identity is not evidence. Section 6.1 shows that two independent derivations arrive at one function. It does not show that surplus is perceived, that perception is surplus, or that either derivation is correct. Two routes to the same functional form is a structural convergence, and a structural convergence is an interpretive correspondence in the sense of Section 8.3, carrying no evidential weight.

The geometric accumulator is not a new result. It is a change of variables on an equation the corpus already had [Curzi, 2026a], and its fourteen-significant-figure agreement with the additive form is a numerical identity check on that change of variables. Agreement to fourteen figures between two ways of writing the same equation is what a correct implementation looks like, not what evidence looks like.

The functional form remains underdetermined. Section 3.5 gives the alternatives — a sublinear power law [Latané, 1981] and `n log n` [Briscoe et al., 2006] — and shows why the first cannot be separated from the corpus's law by shape at all: the logarithm is the `t -> 0` member of the power family, and at `t = 0.10` the two agree to `R^2` = 0.9897 over `N` in `[2,50]`. **Nothing in this paper discriminates among them on data.** The logarithm is chosen for its axiomatic characterisation and its dynamical warrant, not established by measurement, and any future comparison must be made out-of-sample, at matched parameter count, with an information criterion reported, which is the standard Section 12.3 applies to itself.

And the paper's one genuinely discriminating test comes out against Fechner. Section 12.3's ratio-invariance test is the corpus's own attempt to license a ratio scale, and it succeeds; but ratio invariance of the *response* is the power law's invariant and not the logarithm's, so the test's success is evidence for Stevens. On the stipulated series the power law is favoured over the better logarithmic form by `Delta AIC = 26.57` at equal parameter count, its leave-one-out error is thirty times smaller, and the logarithmic forms predict a spread between the two fourfold response ratios of 48 to 85 per cent against an observed 0.40 per cent.

The correct reading of that result, stated so it cannot be softened later: *the corpus's perceptual story is Stevens' rather than Fechner's wherever it is tested by this instrument, and the corpus's surplus functional keeps its logarithm anyway, because that logarithm never rested on the perceptual story.* A reader who takes the title of this paper as a claim that Fechner's law is the warrant for `F(u)` has taken more than the paper establishes, and Section 3.3 said so before the evidence was presented.

No claim about scaling or heavy tails is made anywhere in this paper, so the distribution-fitting procedure that any such claim would require does not apply [Clauset et al., 2009]; the exponents of Section 3.4 are cited point estimates from magnitude-estimation studies, not tail exponents estimated from a sample.

No claim in Section 8 about an inner structure is an identified empirical result. Section 3.7 states the reflection problem in full [Manski, 1993, Shalizi and Thomas, 2011] and records that no remedy is adopted here [Brock and Durlauf, 2001].

20.13.3 13.3 Relation to the surrounding literature

The historical dispute this paper reproduces is not settled and this paper does not settle it. Krueger [1989] is the standing reconciliation attempt; Luce [2002] and Narens [1996] pursue the axiomatic question of what must be true of a subject's judgements for either scale to exist; Poulton [1968] and Poulton [1979] catalogue what goes wrong in the instrument that produced the power law, and the catalogue has not been refuted but managed by better designs. **The framework's position is that the register survives the dispute and the specific functional form does not**, which is why every design specification in Section 10 is stated in ratios and every exponent carries its procedure.

Against the framework's use of subjective magnitude at all, Michell [1997] is the strongest available objection and Section 15 concedes it rather than answering it. Against its use of self-report as an instrument, Nisbett and Wilson [1977], Schwarz [1999] and Kahneman et al. [2006] bound what a report can be used for, and Section 8.7 accepts the bound and changes what a report is for rather than disputing it.

On the corpus's own side, the decay term has a genuine empirical ally in the hedonic-adaptation literature: comfort decays toward a set point [Brickman et al., 1978], adaptation is incomplete and set points move [Diener et al., 2006], and the reference against which a magnitude is judged is itself an adaptation level [Helson, 1948]. **That literature is simultaneously the strongest support for the framework's leak term and its strongest warning against promising accumulation**, and Section 4.8's arithmetic — that steady perceived improvement requires exponentially growing world quantity — is the formal version of the same warning.

Finally, a caution about aggregation that applies to every group-level number in this framework. The behaviour of a population is not in general the behaviour of any individual in it [Kirman, 1992], and the time average of a multiplicative process differs from its ensemble average [Peters, 2019]. **A framework whose stored quantity accumulates over time for one system, and whose parameters would naturally be estimated from a cross-section, is exposed to both**, and Section 15 records the exposure.

20.14 14 Verification

The results of this paper compound from eight postulates. Each is atomic, each is checkable on its own, and each carries its status. Postulates that could be dropped with the result still following have been dropped; what remains is what the chain consumes. **A reader who checks nothing else should check L1, R1 and R2 of Section 14.3, which take about two minutes between them.**

20.14.1 14.1 Dimensional audit

A logarithm's argument must be dimensionless. That single requirement is the whole of failure mode FM7, and it is what this table exists to make checkable.

Symbol	Dimension	Unit	Status
`I`	`[I]`	the modality's physical unit (lux, pascal, gram, …)	**measured**

Symbol	Dimension	Unit	Status
`I_0`	`[I]`	**the same unit as `I`**	**measured** — an absolute threshold, not a convenience constant
`I/I_0`	`[1]`	none	**derived** dimensionless; this is what makes `ln(I/I_0)` well formed
`S`	`[k]`	whatever unit `k` is declared in	**derived**
`k`	sensation units per nat	e.g. JND per nat; `1/ln(1+c) = 20.4959` JND per nat at `c = 0.05`	**defined by the construction of the sensation unit**
`c`	`[1]`	none	**measured** dimensionless fraction; the convention (75 per cent correct) must be stated
`a`	`[1]`	none	**measured** dimensionless exponent; **stated on pressure or on intensity, never bare**
`psi`	`[k']`	the units of the response modulus	**measured**; `k'` is procedure-dependent and not a property of the observer
`u`	`[1]`	none	**defined** dimensionless: `u = 1 - (x.y)^2` for unit state vectors
`u_0 = 1/(N-1)`	`[1]`	none	**derived** dimensionless threshold of difference
`N`	`[1]`	a count	**measured**
`N_eff = 1 + (N-1)u`	`[1]`	a count	**derived**; **not a Hill number** (Section 3.6)
`F = ln N_eff`	`[1]`	nats	**derived** dimensionless: a logarithm of a count
`Q`	`[Q]`	proxy units of stored order	**measured** in whatever the declared proxy counts
`G = e^{Q}`	—	**undefined unless `Q` is dimensionless**	see the note below
`delta`	`[1]`	per period	**measured** dimensionless fraction
`eta`	`[Q]` per nat per period	proxy units per nat per period	**measured**; not dimensionless
`C`	`[Q]` per period	proxy units per period	**measured**
`q`	`[1]`	none	**defined** fraction of an enumerated outcome set
`H_forbid = -ln q`	`[1]`	nats	**derived** dimensionless

Symbol	Dimension	Unit	Status
`SS = 1 − RMSE_m/RMSE_b`	`[1]`	none	**derived**; both RMSEs must be on the same variable at the same points

Does everything balance? `S = k ln(I/I_0)` balances because `I/I_0` is dimensionless and `S` inherits `k`. `psi = k'(I/I_ref)^a` balances **only if `I` is first expressed as a ratio to a reference**; `I^a` for non-integer `a` is dimensionally improper otherwise, and the omission of `I_ref` in casual statements of the power law is a real error that changes `k'` and hides the reference. `F = ln(1 + (N−1)u)` balances because `u` and `N` are both dimensionless.

Two silences are recorded rather than filled. First, `S` has **no operational unit** independent of the procedure that constructs it; under the just-noticeable-difference construction its unit is a JND, under magnitude estimation it is a modulus set by an instruction, and the two are not the same unit. Second, and more consequentially, **`G = e^{Q}` is dimensionally well formed only if `Q` is dimensionless.** `Q` is stored order in proxy units, so `e^{Q}` exponentiates a dimensioned quantity, which is FM7 committed by this framework's own equation. The repair is to read `G` as `e^{Q/Q_u}` with `Q_u` the proxy's unit, so that `G` is the dimensionless effective count the derivation of Section 5 treats it as; **the source is silent on this and it is stated here as a silence rather than quietly assigned.** The arithmetic is unaffected because `Q_u = 1` in the declared proxy, but the equation as written is not dimensionally clean and a reviewer would be right to say so.

Identifiability of the constants. Three cases, and they differ. Under Fechner, **`k` and `I_0` are not separately identifiable from `S` alone**, since `k ln(I/I_0) = k ln I − k ln I_0` and a change in `I_0` is absorbed by an additive offset; only `k` and the offset are recoverable, and `I_0` must be fixed by an independent threshold measurement. **This is why the law must be fitted to discrimination data rather than to ratings.** Under Stevens, `a` is identifiable from the log-log slope and is robust, while the response modulus `k'` is not separately identifiable, because any rescaling of an observer's number use rescales it; **`a` is the reportable parameter and `k'` is a procedural artefact**, which is a fortunate asymmetry, since the physically meaningful quantity is the identifiable one. For the discriminable-step count `N_jnd = ln R / ln(1+c)`, both parameters come from independent measurements — `c` from discrimination and `R` from threshold and saturation determination — so it is **fully identifiable**, which is unusual enough here to be worth marking.

20.14.2 14.2 The postulates

P1 — A continuous map on the positive reals satisfying `S(xy) = S(x) + S(y)` is `S(I) = k ln I`, uniquely up to `k`. *Status:* **CITED** — the continuous solution of Cauchy's logarithmic functional equation; characterisation results of this form, identifying which axioms force a logarithm, are collected in Aczél et al. [1974], and the parallel uniqueness argument for the entropy functional is Khinchin [1957].

P2 — The just-noticeable difference is a constant fraction of the stimulus over a modality's mid-range: `dI/I = c`. *Status:* **CITED** — Weber [1834, 1978], with the mid-range restriction and the edge failures reported as part of the finding.

P3 — All just-noticeable differences are subjectively equal. *Status:* **POSTULATED** — this is Fechner's postulate [Fechner, 1860]. It does not follow from P2, it is not testable in isolation, and it is used at exactly one place in the chain below. **The results R1, R2, R3 and R7 do not consume it.**

P4 — The surplus functional is `f(u) = ln(1 + (N−1)u)` on `u` in `[0,1]`, `N >= 2`. *Status:* **DEFINITION** — the model form adopted by this framework under four axioms, one of them the declared design axiom that the effective count be affine in `u` [Curzi, 2026]. It carries no empirical weight on its own.

P5 — `ln(ab) = ln a + ln b` and `ln(a/b) = ln a − ln b` for positive `a, b`. *Status:* **ARITHMETIC.**

P6 — For a strictly concave `f` and a non-degenerate distribution, `E[f(U)] < f(E[U])` (Jensen). *Status:* **ARITHMETIC** — standard, and `f` is strictly concave here by P4.

P7 — The empirical content of a postulate admitting a fraction `q` of an enumerated outcome set is `H_forbid = −ln q` nats. *Status:* **POSTULATED** — this framework's own construction. Popper [1959], Lakatos [1970], Platt [1964] and Meehl [1967] support the underlying idea that content is what a theory forbids; **none of them states this formula.** It is used at exactly one place, R7, and nothing else in the chain depends on it.

P8 — The time-average growth rate of a multiplicative process is logarithmic in the accumulating quantity. *Status:* **CITED** — Peters and Gell-Mann [2016]. It is the alternative warrant of Section 3.3 and is consumed only by the claim that the framework's logarithm survives the failure of P3.

20.14.3 14.3 The compounding chain

Every line is checkable from its own stated inputs.

```
P1              -> R2   S(I_1 I_2) = S(I_1) + S(I_2) and continuity force
                        S(I) = k ln(I/I_0); unique up to k.              CITED (aczel1974shannon)

P2 + P3         -> L1   dS = k dI/I.                                     POSTULATED (via P3)

L1              -> L2   integrating, S = k ln(I/I_0); the 200,000-rectangle
                        midpoint sum on [1,100] returns 4.605170176
                        against ln 100 = 4.605170186, error 1.02e-08.    ARITHMETIC

P4 + P5         -> R1   ln(1+(N-1)u) = ln((N-1)(u + 1/(N-1)))
                                     = ln((u+u_0)/u_0), u_0 = 1/(N-1).   ARITHMETIC

R1              -> L3   at u = u_0: N_eff = 2 and F = ln 2 = 0.693147,
                        for every N.                                     ARITHMETIC

P4              -> L4   f'(u) = (N-1)/(1+(N-1)u), so f'(0) = N-1,
                        f'(1) = (N-1)/N, and f'(0)/f'(1) = N exactly.    ARITHMETIC

P4 + P5         -> R3   exponentiating Q_{t+1} = (1-d)Q_t + eta S_t - C_t
                        with e^{S} = N_eff gives
                        G_{t+1} = G_t^{1-d} (N_eff)^{eta} e^{-C};
                        G_inf = ((N_eff)^{eta} e^{-C})^{1/d}.            ARITHMETIC on a DEFINITION

P4 + P6         -> L5   E[F(U)] < F(E[U]); at N = 5 and the corpus's
                        U = 0.6019, F(E[U]) = 1.226008 against a
                        reported mean E[F(U)] = 1.1498, a gap of
                        0.0762 nats.                                     ARITHMETIC (data CITED, ccp:CCP-0

P3 + a          -> R4   at a = 0.33, 1.20^0.33 = 1.06201 (+6.201%,
                        31.0% of the promise) and 1.20^(1/0.33) = 1.7376
                        (+73.76%).                                       ARITHMETIC on a CITED exponent

R4              -> R5   core x0.80 gives −7.10%, perimeter x1.50 gives
                        +14.32%, gap −26.09% -> −9.06%, energy
                        0.6(0.80)+0.4(1.50) = 1.0800.                    ARITHMETIC

P7              -> R7   q = 1 gives 0 nats; q = 0.5 gives ln 2; content
```

```
                        adds under independent conjunction.              POSTULATED

stipulated data -> R6   25.10/10.00 = 2.5100 and 98.00/39.20 = 2.5000,
                        agreeing to 0.40%, against 4^0.6569 = 2.4859.    ARITHMETIC on stipulated data
```

What the chain establishes and what it does not. R1, R3, L3 and L4 are arithmetic consequences of one definition (P4) and elementary algebra (P5); **they are certain and they are about a function, not about the world.** R2 is a cited theorem about a functional equation. R4 and R5 are arithmetic on a cited exponent and are only as good as that exponent, which is a procedure-dependent point estimate with no interval (Section 9.9). R6 is arithmetic on a **stipulated** series and demonstrates a procedure rather than measuring anything. R7 consumes P7, which has no external support at all. **And the whole Fechner branch — L1 and L2 — consumes P3, a postulate that cannot be tested in isolation and that the historical record has partly refuted (Section 9.5).** A chain whose load-bearing steps are DEFINITION, ARITHMETIC and CITED delivers mathematics with certainty and delivers nothing whatever about rooms or people.

20.14.4 14.4 Three checks a reader can do in two minutes

The identity. Take `N = 12`, so `u_0 = 1/11 = 0.09090909`. At `u = 0.05`: `ln(1 + 11 x 0.05) = ln(1.55) = 0.438254930931`, and `ln((0.05 + 0.0909091)/0.0909091) = ln(1.55) = 0.438254930931`. Identical, because the second expression *is* the first with `(N-1)` factored out. **ARITHMETIC.**

The marginal ratio. At `N = 12`, `f'(0) = 11` and `f'(1) = 11/12 = 0.916667`, so `f'(0)/f'(1) = 12 = N`. Repeat at `N = 5`: `f'(0) = 4`, `f'(1) = 4/5 = 0.8`, ratio `5`. **The ratio of first to last marginal value is the category count, exactly, for every `N`. ARITHMETIC.**

The equal-step ratio. At `N = 12`, the values of `(u + u_0)` at `F = 0, 0.5, 1.0, 1.5, 2.0` are `0.090909, 0.149884, 0.247117, 0.407426, 0.671732`. Successive ratios: `1.648721` four times, which is `e^{0.5}`. **Equal perceived steps are equal ratios of `(u + u_0)`, and the fifth step would require `u = 1.016590 > 1` and is therefore impossible rather than expensive. ARITHMETIC.**

20.14.5 14.5 Limits, asymptotics, and the model's valid domain

`I -> 0`. Under Fechner, `S -> -infinity`: the register has no floor, which is why an absolute threshold `I_0` must exist and why the domain is restricted to `I >= I_0`. Under Stevens, `psi -> 0` for `a > 0`, so the power law has a floor at zero; the standard repair near threshold is `psi = k(I - I_0)^a`. **The two laws differ qualitatively at this boundary and the difference is observable.**

`I -> infinity`. Both diverge. Neither divergence is real: modalities saturate, and in pain-adjacent modalities the tissue is damaged rather than the sensation extended. **Both are mid-range laws and their own limits announce it.**

`c -> 0`. `N_jnd = ln R / ln(1+c) -> infinity`: infinitely fine discrimination gives infinitely many categories, correct in the model and correctly impossible in fact. **`c -> infinity`.** `N_jnd -> 0`: an observer who cannot discriminate has one category. **`R -> 1`.** `N_jnd -> 0`: no range, no steps. `N_jnd = 1` at `R = 1 + c` recovers the definition of the just-noticeable difference as a fixed point, which is a consistency check rather than a discovery.

`k -> 0`. All percepts collapse to zero: no discrimination. **`k -> infinity`.** Every difference becomes infinite: no constancy. Both are real failure modes of sensory systems.

`u -> 0`. `N_eff -> 1` and `F -> 0`: every category duplicates every other and the roster has left the arithmetic. **`u -> 1`.** `N_eff -> N`, `F -> ln N`, the ceiling is attained and no further gain is available from geometry.

`N -> 1`. `u_0 = 1/(N-1) -> infinity` and `F(u) = ln(1 + 0) = 0` for every `u`. **A single category has zero available surplus, not a small amount**, and the threshold has diverged: with one category there is no difference that could register. **`N -> infinity`.** `u_0 -> 0` and any non-zero difference produces unbounded surplus while `f(0)` stays exactly zero, so `f` approaches a step

function. **The limit is a pathology, not an ideal** (Section 6.2): an observer with infinitely fine discrimination is maximally unstable rather than maximally perceptive.

`a -> 0.` `psi -> k'`: a modality with no sensitivity to intensity, approached by some quality dimensions. **`a -> 1.`** The register's corrections vanish and differences behave like differences; visual length sits here and the design advice of Section 7 does not apply to it. **`a > 1.`** Every sign in Section 7 reverses.

`q -> 1.` `H_forbid -> 0`: a postulate that forbids nothing says nothing. **`q -> 0.`** `H_forbid -> infinity`: a postulate forbidding everything but one outcome, maximally informative if true and nearly certain to die on first contact. **The divergence is real and nameable.**

Stated domain. The register results of this paper hold for `I` in `[I_0, I_sat]` with the mid-range restriction of P2; for `u` in `[0,1]` and `N >= 2`; for exponents measured on a stated physical quantity under a stated procedure; and for `q` in `(0,1]` over an enumerated and published outcome list. Outside those the expressions still evaluate and should not be interpreted.

20.14.6 14.6 Invariance

The register claim is an **invariance claim** and is worth stating as one. Under Fechner the invariant is `S(aI) − S(I) = k ln a`, **independent of `I`**: equal stimulus ratios produce equal sensation differences everywhere on the scale. Under Stevens the invariant is `psi(bI)/psi(I) = b^a`, **independent of `I`**: equal stimulus ratios produce equal sensation ratios everywhere. **The two invariances are different and are experimentally distinguishable**, and Section 12.3 distinguishes them: the observed pair of response ratios is invariant to 0.40 per cent across an eightfold change of absolute level, which the Fechnerian invariance does not predict.

The transformation group under which the framework's own quantities are invariant is the rescaling of the world quantity. `u`, `N`, `N_eff`, `F` and `delta` are dimensionless and unchanged under `I -> cI`; `Q`, `eta` and `C` carry the proxy's unit and rescale with it. **Two systems may be compared directly on `u`, `N` and `delta` and may not be compared on `Q`, `eta` or `C` unless they declare the same proxy.**

How many scales the invariance has been tested at, here. Two, in Section 12.3, both on the same constructed series: the pairs `(1.00, 4.00)` and `(8.00, 32.00)`, separated by a factor of eight in absolute level, over a total range of 2.107 decades. **That is two levels, one continuum, one constructed dataset, and no independent replication.** Anything that would break it is likewise listed: a range change large enough to move the exponent [Teghtsoonian, 1971], a modulus at an extreme, an unrandomised presentation order, or a bounded response instruction.

20.14.7 14.7 Analogy discipline, and the precision of the numbers

The geometric accumulator is a change of coordinates, not a derivation from physical premises. `G_{t+1} = G_t^{1−delta}(N_eff)^{eta}e^{−C}` and `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t` are the same equation written in two ways, related by `G = e^{Q}`. It is a **shared mathematical form**, and the agreement to fourteen significant figures is a numerical identity check on that change of variables. **No physical result's authority transfers, and the agreement is not evidence.**

The Lens Theorem's mathematics is a derivation; its application to perception is a conditional whose antecedent is empirical. The distinction is maintained throughout and is the subject of honesty clause 1.

The identity between the surplus functional and Fechner's law is an interpretive correspondence. Two formalisms have the same shape; no theorem connects the axioms of one to the experiments of the other. Under the badge taxonomy of Section 8.3 it carries no evidential weight, and Section 6.1 says so where the identity is stated rather than here.

Entropy, where it appears near `ln N`, is Shannon entropy [Shannon, 1948] — a dimensionless information measure in nats — and not thermodynamic entropy, which carries

`[energy][temperature]^-1`. They are different objects and the framework identifies surplus with neither.

Precision. The numbers here are of four kinds and are quoted accordingly. *Defined or derived-from-defined* — `ln 2 = 0.693147`, `ln 12 = 2.484907`, `e^{0.5} = 1.648721`, the identity table of Section 6.1, `f'(0)/f'(1) = N` — are exact to the precision printed and may be recomputed further. *Arithmetic on stipulated parameters* — every entry in Sections 7, 12.1 and 12.2 — records the arithmetic and **carries no measurement uncertainty because there was no measurement.** *Cited point estimates* — every Stevens exponent — carry no interval **because none is available from the sources as reproduced**, and they are known to vary with range, modulus, instruction and observer by amounts comparable to the design differences computed from them [Teghtsoonian, 1971, Poulton, 1979, Stevens, 1975]. *Fitted* — the exponent `0.6569` and the `R^2` of `0.999932` in Section 12.3 — are fitted to a **constructed** series and are illustrative in exactly the sense above; the AIC, BIC and cross-validated errors reported beside them are model comparisons on that same constructed series and establish which model describes it, not which model describes any modality.

20.15 15 Limitations

All three worked cases are constructed and none is evidence. No dataset was collected for this paper. The reading room's dimensions, illuminances and complaints are stipulated; the postulate case's block means and its 40-day series are stipulated, the latter explicitly synthetic with a declared generating rule; the magnitude-estimation series is stipulated. **They demonstrate that the procedures run and that they report their own failures. They are not evidence about any room, any person or any modality and must never be cited as such.**

The title claims more than most of the paper needs. Section 3.3 sets this out and it is repeated here because it is the paper's largest single exposure. The corpus needs a logarithm. Peters and Gell-Mann [2016] supply one from multiplicative dynamics with no psychological premise; Fechner [1860] supplies one that requires the perceptual story to be true and is contested [Stevens, 1957, Poulton, 1979, Luce, 2002, Narens, 1996]. **The framework's central formal claims — the surplus functional's form, the storage law's world-coordinate reading, the concavity result, the threshold — rest on the axiomatic route [Curzi, 2026] and the dynamical route [Peters and Gell-Mann, 2016], and would survive Fechner's complete failure.** Only Section 7's design advice needs Fechner and Stevens, and there the contest between them is a live cost rather than a historical curiosity, because the two theories give different exponents and the design numbers are computed from the exponent.

The paper's one discriminating test comes out against the logarithm. Section 12.3's ratio-invariance test succeeds, and its success is the power law's invariant rather than the logarithm's; the model comparison on the same series favours the power law over the better logarithmic form by `Delta AIC = 26.57` at equal parameter count. **On a constructed series, on one continuum — but in one direction, and the direction is against Fechner.**

The functional form of the surplus law is underdetermined by the evidence. A sub-linear power law [Latané, 1981] and `n log n` [Briscoe et al., 2006] fit the qualitative facts as well as `ln(1 + (N−1)u)` does. **Nothing in this paper or, so far as it reports, in this corpus discriminates among them.** The logarithm is chosen for its axiomatic characterisation, one of whose axioms is a declared design choice [Curzi, 2026], and is not established by measurement.

`N_eff = 1 + (N−1)u` is not a Hill number and no external literature contains it. `1/HHI` is the Hill number of order two and `exp(H)` the Hill number of order one, both fully grounded [Hill, 1973, Jost, 2006]; `1 + (N−1)u` is a linear interpolation between 1 and `N` and is this framework's own construction. The `N_eff` in result R3 and throughout Section 5 is the latter. **A reviewer reaching for the Hill-number literature as support will find that it does not apply.**

The empirical-content measure has no external support and is not identifiable in general. `H_forbid = −ln q` is this framework's own construction, badged POSTULATED in Section 14.2. Two analysts enumerating outcomes differently assign different `q` to the same postulate, so the

measure is identifiable only relative to a published outcome list; a defensible measure over an outcome space is almost never available; and **whether the heuristic actually improves the return on a testing budget is itself untested.** The test is stated in H5 and has not been run.

The conjoint-measurement conditions were not tested. Michell [1997] argues that psychology systematically fails to test whether its attributes are quantitative before measuring them, and that objection applies with full force to the treatment of subjective magnitude here, to `u`, and to comfort. The remedy is conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971] or a Rasch-model demonstration [Rasch, 1960]; **neither was run.** Section 12.3's ratio-invariance test is this framework's own substitute and it is a weaker instrument: it tests one consequence of ratio structure at two levels on one continuum, and passing it is consistent with a representation that conjoint measurement would reject. **The ratio-scale treatment in this paper is provisional and is a stipulation.** The corpus's own route to earning it is priced in Curzi [2026t], and its standing demotion rule — that a formula containing a division and a multiplication is demoted to a heuristic ordering when its inputs are ordinal — is Curzi [2026m].

The Stevens exponents carry no uncertainties. Every design number in Section 7 is computed from a cited point estimate. Those estimates vary systematically with stimulus range [Teghtsoonian, 1971], with modulus, instruction and direction [Poulton, 1979], and across observers by a substantial fraction of their value [Stevens, 1975]. **The sources as reproduced here supply no intervals, so none is quoted, and a design decision resting on the difference between an exponent of 0.33 and one of 0.40 is not supported by anything in this paper.** S35 requires re-measurement before calibration for exactly this reason.

The reading room's premise is a stipulation. No entry in this framework's bibliography establishes a quantitative relation between illuminance, zone contrast and occupant complaint. The nearest support is design literature on how the distribution of light within a room conditions its use [Alexander et al., 1977], which is not a measured relation. **The case demonstrates a budgeting procedure; it does not demonstrate that the procedure predicts complaints.**

Model comparison was performed at one place only, and out-of-sample performance is reported at one place only. Section 12.3 reports SSE, `R^2`, AIC, BIC, free-parameter counts and leave-one-out cross-validated error for three models on eight points. **Eight points is a small sample for leave-one-out cross-validation and the AIC correction for small samples widens materially at `n = 8`.** Nothing else in this paper is fitted, so no information criterion applies elsewhere; and the corpus's one applied fit reports held-out RMSE against baselines and **no AIC or BIC**, so its parameter count was never paid for and its 2.8 per cent margin over a two-parameter AR(1) may not be a margin at all [Curzi, 2026i].

No effect size in this paper carries an interval. The standardised effect of 1.7347 in Section 12.2 is computed on stipulated means and standard deviations; there is no sampling distribution behind it because there was no sample. A real application owes intervals, and reporting a standardised effect without one is exactly what Section 14.7's precision rule forbids for measured quantities.

Single-observer scope, and the circular instrument. Section 12.2's case has the observer and the analyst as one person, which is failure mode FM14 in its hardest form; the prescribed separation of channels was not achieved. Its conclusions are indexed to one observer and may not be stated in the third person (S22).

The reflection problem is not answered. Where the framework's storage reading is applied to a group whose members both influence and constitute the aggregate, the endogenous effect is not identified from observational data [Manski, 1993], and influence and selection are generically confounded in network data [Shalizi and Thomas, 2011]. **This paper adopts no remedy** — not randomisation, not an instrument, not partial-population variation, not a stipulated and defended reference group [Brock and Durlauf, 2001] — so every such claim here is a **modelling position rather than an identified empirical result.**

The Lucas critique applied to this framework's own parameters. The framework asserts that a description can become true by being used [MacKenzie and Millo, 2003, Hacking, 1995] and then treats its own exponents, thresholds and decay rates as stable. **Both cannot hold**

without qualification. The position taken here is narrow: the parameters are treated as stable only over a declared window and only while the quantity is not itself a reported target, and where it is reported upward the framework's own reflexivity applies to it and the constancy assumption fails [Lucas, 1976]. That is a restriction on where the model may be used and it is stated as one. The measured size of the reflexive effect is also much smaller than the framework's use of it implies [Jussim and Harber, 2005, McCambridge et al., 2014].

Self-report is exposed on three separate fronts. Causal reports about one's own processing are frequently constructions [Nisbett and Wilson, 1977]; question wording, scale range and order shape the answers [Schwarz, 1999]; and any rating of a named component of well-being is contaminated by the focusing illusion [Kahneman et al., 2006]. Section 8.7 answers the first by changing what a report is used for; **it does not answer the second or the third**, and any survey built on this framework inherits both.

Aggregation and ergodicity. The behaviour of a population is not in general the behaviour of any individual in it [Kirman, 1992], and the time average of a multiplicative process differs from its ensemble average [Peters, 2019]. **A framework whose stored quantity accumulates over time for one system, and whose parameters would naturally be estimated from a cross-section, is exposed to both, and the exposure is not addressed here.**

No independent replication of anything. Every number in this paper was produced by one author from one framework, and every applied number elsewhere in this corpus was likewise surveyed or constructed by one author. The reproducibility literature is unambiguous about what that is worth [Open Science Collaboration, 2015, Ioannidis, 2005]. Meehl [1967]'s observation that directional predictions are almost costless in the social sciences is answered here by point predictions with constants — Section 12.1 predicts `+14.32%` and `−7.10%` and puts them at risk — but **no such prediction in this paper has yet been tested against an observation.**

What would falsify the account. H1 fails on a modality showing a constant absolute just-noticeable difference over its mid-range. H2 fails if response ratios depend on absolute level, or if a bounded rating scale reproduces the invariance as well as an unbounded instruction. H3 is an algebraic identity and fails only on a computational error. H4 fails if post-occupancy magnitude estimation returns a perceived change matching the physical percentage rather than the exponent-converted one, in a modality whose exponent has been measured under the application's own conditions. H5 fails if preregistered high-content postulates do not repay a fixed budget better than low-content ones. H6 fails if a shuffled-input control leaves held-out error near its unshuffled value while the other baselines are beaten.

20.16 16 Conclusion

R1. The framework's surplus functional is a Fechner law. Setting `u_0 = 1/(N−1)`, `f(u) = ln(1 + (N−1)u) = ln((u + u_0)/u_0)` identically on `u` in `[0,1]` for every `N >= 2`, so `u_0` is an absolute threshold of difference, `N_eff(u_0) = 2` and `F(u_0) = ln 2` for every `N`, and the ratio of first to last marginal value is `f'(0)/f'(1) = (N−1)/((N−1)/N) = N` exactly. **The identity is algebra, not evidence: it shows that an axiomatic route and a psychophysical route arrive at one function, and it shows nothing about whether either is correct.**

R2. The Lens Theorem: if the map from world quantity to perceived magnitude is continuous and satisfies `S(I_1 I_2) = S(I_1) + S(I_2)`, then `S = k ln(I/I_0)`, uniquely up to `k`. **The inference is a theorem; the badge is on the antecedent.** Verified numerically by a 200,000-rectangle midpoint sum of `1/I` on `[1, 100]` returning `4.605170176` against `ln 100 = 4.605170186`, absolute error `1.02e−08`.

R3. In world coordinates the storage law is a weighted geometric accumulator, `G_{t+1} = G_t^{(1−delta)} (N_eff)^{eta} e^{−C}` with `G = e^{Q}` and `G_inf = ((N_eff)^{eta} e^{−C})^{1/delta}`, agreeing with the additive form to fourteen significant figures. Decay is taking a root, conversion is exponentiation, cost is a divisor. **`N_eff` here is `1 + (N−1)u` and is not a Hill number. This is a change of coordinates and the agreement is an identity check, not evidence.**

R4. The linear-promise prohibition, quantified. At a brightness exponent of 0.33 a promised and delivered `+20%` illuminance yields `1.20^0.33 = 1.06201`, that is `+6.201%` perceived — **31.0 per cent of the promise** — and delivering a genuine perceived `+20%` requires `1.20^(1/0.33) = 1.7376`, a **73.76 per cent** load increase. **Never promise a linear improvement in a logarithmic quantity.** The exponent is a cited procedure-dependent point estimate without an interval.

R5. The reading-room budget: core 500 to 400 lux costs `−7.10%` of perceived brightness, perimeter 200 to 300 lux buys `+14.32%`, the zone gap closes from `−26.09%` to `−9.06%`, and energy moves by `0.60(0.80) + 0.40(1.50) = 1.0800`, that is `+8.00%`, inside a declared 10 per cent cap. **The complaint was the gap and the gap was closed by taking light away from the place that had most of it.** The case is constructed and the illuminance-to-complaint premise is a stipulation.

R6. The ratio-invariance test as the licence for a ratio scale. Two fourfold physical steps at absolute levels separated by a factor of eight returned response ratios of `2.5100` and `2.5000`, agreeing to `0.40%`, against a prediction of `4^0.6569 = 2.4859` from the fitted exponent. **The invariance is the power law's, not the logarithm's**: on the same series the power law is preferred over the better of two logarithmic forms at equal parameter count by `Delta AIC = 26.57`, with a leave-one-out cross-validated error thirty times smaller than the plain logarithm's, and the logarithmic forms predict a spread between the two response ratios of 48 to 85 per cent against the observed 0.40 per cent. **A ratio-invariance test that succeeds is evidence for Stevens and against Fechner, and the corpus's surplus functional keeps its logarithm because that logarithm never rested on Fechner** [Curzi, 2026, Peters and Gell-Mann, 2016].

R7. The empirical content of a postulate over an enumerated and published outcome set is `H_forbid = −ln q` nats, where `q` is the admitted fraction: a postulate that forbids nothing buys zero nats, halving the admitted outcomes buys `ln 2 = 0.6931` nats, and content adds under independent conjunction. A postulate at `q = 0.9` buys `0.1054` nats for a full experiment against `2.9957` nats at `q = 0.05`, twenty-eight times the return for the same expenditure. **The formula is this framework's own, badged POSTULATED, supported by no external source, and the claim that it improves the return on a testing budget is untested.**

CR (counter-result). The worked postulate of Section 12.2 was **refuted on one of its four declared consequences**, at a standardised effect of 1.73 pooled standard deviations against a pre-set tolerance of 0.50 — **three times the tolerance, in the forbidden direction** — and was amended in public rather than withdrawn, with the amendment adding forbiddance rather than removing it. A fourth consequence was recorded as untested rather than folded into the supported column. And the log-log fit of Section 12.3, with `R^2 = 0.999932`, is **explicitly declared insufficient** to exclude a logarithmic law over two decades, since over two decades a logarithm captures 96.3 per cent of a power function's variance; the ratio-invariance test, the residual pattern and the information criteria carry the conclusion, not the fit statistic. **The general form of that concession is the one downstream papers should quote: a fit statistic computed over the range in which two candidate models agree distinguishes nothing, whatever its value, and the corpus's own surplus functional is currently in exactly that position with respect to a sub-linear power law and `n log n`.**

20.16.1 Declarations

Data and code availability

No external dataset was collected for this paper. All numerical results are computed from the parameters and equations printed in the text and can be re-derived from them. The reading room of Section 12.1 is a constructed design case with stipulated dimensions, illuminances and complaints. The block means and the 40-day series of Section 12.2 are a stipulated worked case, the latter explicitly synthetic with the generating rule and seed policy stated in the text. The eight-point magnitude-estimation series of Section 12.3 is a stipulated worked case and is not data from any observer; the exponent, `R^2`, residuals, ratio-invariance figures, AIC, BIC and leave-one-out errors reported from it are computations on that constructed series. The

Stevens exponents in Sections 3.4, 7.2, 7.3 and 9.3 are **nominal published figures reproduced from the psychophysical literature for orientation**, not measurements made here. The figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper and the equations of Sections 3 to 9.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi derived the results, performed the model comparison of Section 12.3, produced the worked cases and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The observer in Section 12.2 is constructed and the twenty observers in Section 12.3 are stipulated; neither exists. The decline paths required by S25 and S36 — that any person asked to supply readings may stop at any point, without a reason and without cost, and is told so before the first trial — are design constraints on prospective work and are not a report of any study performed.

Preregistration

This paper reports one preregistered test of its own and it failed. The four consequences D1 to D4 of the postulate in Section 12.2, and the forbiddance list F1 to F4 with its numerical tolerances, were declared before the observations were taken. **F2 was triggered: consequence D2 was falsified at 1.73 pooled standard deviations against a declared tolerance of 0.50, three times the tolerance, in the forbidden direction.** No tolerance was moved after the observations were examined. The postulate was amended in public under the rule of Section 11.2, with the amendment adding a specific and testable forbiddance — ordering and concavity of recovery across three break lengths — rather than removing one, and the amended clause is recorded as **untested**, as is consequence D4, which was never exercised because start time was not varied. **The worked case that produced this result is constructed, so the failure demonstrates that the protocol reports failures rather than establishing anything about breaks or about mornings.**

The design predictions of Section 12.1 — `+14.32%` at the perimeter, `−7.10%` at the core, and a gap closing to `−9.06%`, with a declared tolerance of `±0.05` nats on the perimeter figure — are stated as a preregistration for a measurement that **has not been performed**. Nothing in this paper reports its outcome.

The corpus's protocol for preregistering and freezing a fit of the storage law, together with the confirmatory application that **failed** its own second falsification rule at a 2.8 per cent held-out advantage over a matched AR(1) against a declared 10 per cent margin, is Curzi [2026i]. Any applied claim resting on this paper inherits that result.

20.17 Appendix A — Stating one's own subjective structure as a postulate

The protocol executed in Section 12.2, stated in general form. It is the single-observer version of the survey protocol [Curzi, 2026i] and is meant to be run rather than admired.

Step A — Write the belief in its native, unusable form. Whatever is actually said internally: "I am sharpest in the morning"; "I need conflict to think clearly"; "social events drain me." Write it exactly as it is said. **The gap between this sentence and the one the protocol finishes with measures how much was being assumed without notice.**

Step B — Name the hidden quantity. In each of those sentences one unobservable does all the work — sharpness, clarity, drain. Name it, give it a symbol, and stop using the word.

Step C — State its structure. The postulate proper. It must answer four questions: what kind of quantity is it, a level or a rate or a capacity; does it have a true zero; what makes it move; and what observable does it touch. Badge it (S13) and date it (S14).

Step D — Derive at least three consequences, one of which must be uncomfortable. A consequence one would be glad to observe is worth little, because it will be found whether or not it is there. **The uncomfortable consequence carries the content.** Apply the asymmetry check of S17 to each: write down what the *negation* of the postulate would predict, and delete any consequence the negation also produces.

Step E — Write the forbiddance list, before observing (S15). For each consequence, the specific observation that would contradict it, with numbers and tolerances where numbers are available, in terms executable by a person who does not know the hypothesis (S16).

Step F — Choose the instrument and the sample size before observing (S20). Then observe, then stop and compute. **The stopping rule is part of the design**; choosing when to stop after seeing the data is the commonest way a personal experiment becomes theatre [Simmons et al., 2011].

Step G — Report in the same words that would be used to report to a stranger (S24). The gap between the two sentences is the size of the self-deception; record it in the ledger (S19).

Two structural defences apply throughout and are not optional at single-observer scope. **Delegate the classification** to someone blind to the hypothesis wherever the forbiddance list requires a judgement, because an operational definition requiring the author's own judgement to apply is elastic (FM17). And **precommit the response**: write down in advance the change of practice that follows from the postulate's death, so that death has a defined and survivable consequence.

A postulate about one's own interior is legitimate exactly when the observation that would end belief in it has been written down before looking — and then looked for. Everything else in this protocol is machinery for making that one sentence executable.

Chapter 21

The Internal Surplus and the Self as Storage: Auditing One Observer's Framings, and the Stored-Error Track That Feels Identical From Inside

Abstract. We transpose the surplus law to a single observer, treating N framings held by one person as the ensemble whose non-parallelism u sets an audited framing count N_eff = 1 + (N-1)u and an internal surplus F_int = ln N_eff, bounded by ln N. Non-parallelism is estimated from error decorrelation as u = 1 - rho^2, and we prove the participation-ratio identity N_eff = N + 1 - N/PR. A six-framing worked case measures mean pairwise u = 0.6860, giving N_eff = 4.4298 and F_int = 1.4883 nats, 83.07 per cent of the ln 6 ceiling; splitting the pairs by lineage shows within-lineage u = 0.2732 against cross-lineage u = 0.9611, a concentration the mean hides. Retiring two redundant framings and replacing them raised the audited count 25.4 per cent at unchanged reconciliation load. The central negative result is that this correct geometric work did not reverse the subject's decline: eta*F rose from 0.8930 to 1.0289 against delta*Q + C = 1.0700, closing 77 per cent of the gap and leaving the account negative; only cost and stock reductions turned it. Sensitivity analysis makes the point structural: dh/dN = eta/N = 0.0038 is negligible while dh/deta = ln N = 2.197. The two-track ledger separates stored order from stored error, Z_t+1 = (1-delta)Z_t + eta(M_t - V_t), and a 25.93 per cent confabulated input yields 57.38 per cent of the steady-state stock. Baseline headroom was not robust to the declared cost weights and is reported as unestablished.

Keywords: internal surplus; effective framing count; error decorrelation; leaky integrator; stored error; conversion efficiency; headroom

Series: B. Engineering the Surplus **Paper:** CCP-021

21.1 1 Introduction

The surplus functional of this framework was introduced as a statement about groups: two people, two teams, two disciplines, and the difference between them. That presentation was a convenience of exposition. The functional is a theorem about two states in an inner-product

space, and its axioms quantify over states rather than over the persons who hold them. This paper takes the step that the presentation deferred, and asks what the same functional says when both states belong to one observer.

The answer is an **internal surplus**. Where an observer holds `N` distinguishable framings of a situation — ways of construing it, each with its own primitives, its own salient variables and its own inference rules — and those framings have mean non-parallelism `u`, the surplus available to that observer is

```
F_int = ln(1 + (N − 1) u),     N_eff = 1 + (N − 1) u,     0 <= F_int <= ln N
```

with `N_eff` the audited framing count. This paper's first task is to establish that the move is a substitution of domain rather than an analogy, and to state precisely what has to be assumed for the substitution to be legitimate. Its second task is to carry `F_int` into the storage law and to ask what a person accumulates. Its third, and the one the paper regards as its most useful finding, is to report that the geometric work — measuring `u`, restructuring a framing set, raising the audited count by a quarter — did **not** turn the worked subject's storage account around, because the binding constraint was never the geometry.

Four things are new here relative to the surrounding literature rather than relative to the framework. First, an operational estimator for `u` at internal scope that asks nobody to judge similarity: `u = 1 − rho^2`, with `rho` the Pearson correlation between two framings' *error* series on questions with later-observable answers. Second, an exact identity connecting the audited count to the Rényi-2 participation ratio of the framing ensemble's Gram matrix, `N_eff = N + 1 − N/PR`, which converts a chosen aggregator into a standard diversity statistic. Third, a two-track ledger separating what an observer stored from what the world actually supplied, with the stored-error recursion `Z_{t+1} = (1−delta) Z_t + eta (M_t − V_t)` whose equilibrium is independent of cost because the cost term cancels in the subtraction. Fourth, a sensitivity result that demotes the framework's most attractive-sounding lever: at the worked operating point, `dh/dN = eta/N = 0.0038` against `dh/deta = ln N = 2.197`, so adding categories of thought is the weakest available move by nearly three orders of magnitude.

The question matters because the failure this paper describes is invisible from inside. An observer whose framings have collapsed toward parallel does not experience impoverishment; they experience clarity, and the arithmetic prices that clarity at zero. An observer accumulating stored error does not experience unreliability; they experience confidence, because stored error and stored knowledge occupy the same variable and no term in the storage law inspects its own input. Both conditions are stable for years. Both are detectable only by an instrument that is deliberately built and deliberately kept.

The framework's static half — the surplus functional and its characterisation under four axioms, with the affine effective count as the one axiom that is chosen rather than derived — is presupposed rather than re-derived, and is established with its axiom set in Curzi [2026]. Its dynamical half — the storage law as a first-order leaky integrator with contraction modulus `1 − delta` and steady state `Q* = (eta S − C)/delta` — is presupposed from Curzi [2026a]. The headroom gauge `h_t = eta ln N − delta Q_t − C_t`, its three regimes, and the separate admissibility and sustainability constraints are presupposed from Curzi [2026b]. The reflexive loop `T(x) = W(x, M(x))`, on which Section 3.4's identification discussion turns, is presupposed from Curzi [2026] and is **not** restated here.

Sections 2 and 3 state the hypotheses and the formal framework, including the identification problem that governs every applied reading in the paper. Sections 4 to 6 develop the internal surplus: its marginal returns, its quadratic maintenance cost, the six mechanisms by which non-parallelism decays, and the five instruments that measure it. Section 7 develops the self-storage law and the two-track ledger. Sections 8 and 9 collect the design specifications and the named failure modes. Section 10 reports two worked cases end to end. Sections 11 to 13 discuss what follows, verify the chain against minimal postulates, and record what the account cannot support.

21.2 2 Hypotheses

The formal results in Sections 3, 4 and 7 are theorems about equations and cannot be falsified by observation; they are labelled as derivations where they appear. What can be falsified is the claim that these equations describe a real observer. The hypotheses below separate the two and each names the observation that would defeat it. **None of them has been tested.** The protocol that would test them — a preregistered fit, a frozen model, and held-out comparison against four declared baselines including a shuffled-input control — is specified in Curzi [2026i], whose own confirmatory application of it failed its second falsification rule: a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin.

H1 — The audited count predicts, and the nominal count does not. Internal surplus estimated from error decorrelation among an observer's framings predicts held-out judgement quality better than framing *count* alone. **Predicted:** held-out accuracy tracks `N_eff = 1 + (N−1)u` rather than `N`; in the worked restructuring of Section 10.1, `N` is unchanged at 6 while `N_eff` rises by a factor of 1.254, and the accuracy gain should follow the second number. **Falsified by:** framing count predicting as well as the audited count; or a shuffled-input control matching the model. If `N` alone does the work, the `u` term is decoration.

H2 — Reconciliation load is quadratic while surplus is logarithmic. The maintenance cost of holding `n` framings as framings rather than as a list grows with the number of pairs. **Predicted:** a net-value curve `V(n) = eta ln(1+(n−1)u) − c1 n − c2 n(n−1)/2` with an interior maximum; measured performance rises then falls in `n`, and the fall is steeper than a linear cost model predicts. **Falsified by:** measured performance flat in `n` over the tested range; or peaking where the model did not predict; or a linear cost term fitting as well as the quadratic one, in which case the pair-cost postulate is wrong at internal scope.

H3 — Non-parallelism decays geometrically without maintenance. Left alone, an observer's framing set collapses toward parallel at a roughly constant per-round proportional rate. **Predicted:** `u_t = u_0 (1 − lambda)^t`, with half-life `ln 2 / (−ln(1−lambda))` rounds; at `lambda = 0.20` and `u_0 = 0.60`, sixteen rounds take `F_int` from 1.3863 to 0.0811 nats and the audited count from 4.00 to 1.08. **Falsified by:** a measured `u` that is stationary under no maintenance; or that decays linearly rather than proportionally; or that recovers spontaneously.

H4 — Conversion, not category count, is the binding lever on personal headroom. The partial derivatives of `h = eta ln N − delta Q − C` order the four available interventions, and the ordering is stable across plausible parameter values. **Predicted:** `dh/deta = ln N`, `dh/dC = −1`, `dh/dQ = −delta`, `dh/dN = eta/N`; at the worked point these are 2.197, 1, 0.0317 and 0.0038, so a unit change in conversion buys about 585 times what a unit change in category count buys. **Falsified by:** an intervention on `N` producing a larger measured gain in capacity for new material than a matched intervention on `eta` in a system where both were measured.

H5 — Stored error accumulates and is not distinguishable from inside. An observer whose framing systematically exceeds what the environment supplies accumulates a stock of unearned structure that behaves exactly like earned structure in every internal report. **Predicted:** at matched moment `M`, an observer with higher measured over-framing ratio `r` accumulates more correction events per unit stored order, on a lag of order `tau = 1/delta`; in the worked case a 25.93 per cent confabulated input becomes 57.38 per cent of the steady-state stock. **Falsified by:** correction-event rate uncorrelated with `r` after controlling for `M`. If that holds, the two-track ledger has no content and should be deleted.

> **H6 — Retention channels compose multiplicatively.** Total retention is the product of the per-channel retentions, not one minus the sum of the per-channel losses. **Predicted:** `delta = 1 − prod_i (1 − delta_i)`; at the worked channel values 0.020, 0.008 and 0.004 the correct composite is 0.03172864 against a naive sum of 0.032, an overcount of 0.85 per cent, rising to 20.3 per cent at four channels of 0.15. **Falsified by:** a measured composite reliably closer to the naive sum than to the product, across systems in which the channels are separately measurable.

Three further claims in this paper are **not** hypotheses and are labelled accordingly. The **framing encoding** — that a way of construing a situation admits representation as a unit vector in a common inner-product space — is a **postulate**, stated in Section 3.2 and falsifiable only through its consequences. The **affine effective count** `N_eff = 1 + (N−1)u` is a **design axiom** of the framework, chosen and not derived [Curzi, 2026,c]. The identification of a person's stored order with a single non-negative scalar is a **modelling choice**, adopted for tractability and for the fixed-point guarantee, and Section 13 records what it costs.

21.3 3 Framework

21.3.1 3.1 What is presupposed

The surplus functional. For two unit state vectors `x` and `y` drawn from an ensemble of `N` categories,

```
F(x,y) = f(u(x,y)),   u(x,y) = 1 − (x . y)^2 = sin^2 theta   in [0,1]

f(u) = ln(1 + (N−1) u)
```

and `f` is uniquely characterised by four axioms — angle-only dependence (S1), zero at parallel (S2), an affine effective count (S3), and normalisation at orthogonality (S4) — as established in Curzi [2026]. **S3 is a design axiom: it is chosen, not derived**, and it is flagged wherever the paper leans on it. Uniqueness-under-axioms is a known genre of theorem rather than a novelty of this framework, and the argument should be read against the existing instances [Khinchin, 1957].

The object S3 defines is the **effective count** `N_eff = 1 + (N−1)u`, and surplus is `ln N_eff`: the logarithm of a headcount after the headcount has been audited for redundancy [Curzi, 2026c]. The effective-count idiom is standard in ecology and in political science, where `1/HHI` and `exp(H)` are the canonical members of the family [Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979]. **It must be said plainly, and is said again in Section 13, that `1 + (N−1)u` is not a member of that family.** It is a linear interpolation between 1 and `N`, it is this framework's own construction, and no literature in ecology, economics, political science or information theory contains it. The Hill-number citations in this paper attach to `1/HHI`, to `exp(H)`, and — as Section 3.3 shows — to the participation ratio `PR`, where they are exact. They do not attach to `N_eff`. The concentration instrument that genuinely is a Hill number of order two — `HHI` as the sum of squared shares, with effective count `1/HHI` — is developed separately, where a four-share map returns a naive `HHI` of 0.293430 and 3.41 effective units against a bloc-collapsed `HHI` of 1.000000 and an effective count of 1.00 [Curzi, 2026]. **The two quantities must not be conflated.**

Derived from the axioms rather than assumed: strict monotonicity, positivity, strict concavity, subadditivity on the parameter `u`, and extremality, with `F = 0` if and only if the states are parallel and `F = ln N` if and only if they are orthogonal. The range `0 <= F <= ln N` makes `ln N` the hard ceiling on surplus per interaction and appears in every design inequality below. Extremality is stated on the *squared* alignment, which has the consequence that a pair in perfect opposition and a pair in perfect agreement both read `u = 0`, `N_eff = 1` and surplus zero [Curzi, 2026r]; Section 6.2 shows that this is exactly the behaviour the error- decorrelation estimator inherits.

Two independent warrants exist for the logarithm, and this paper needs only the weaker one. The first is Fechnerian: sensation proportional to the logarithm of a dimensionless stimulus ratio measured from a floor [Fechner, 1966, Stevens, 1957]. That warrant requires the perceptual story to be true, and the perceptual story is contested [Poulton, 1979]. The second is stronger because it assumes nothing psychological: the time-average growth rate of a multiplicatively evolving quantity is the expectation of its *logarithm* [Peters and Gell-Mann, 2016], so a logarithmic evaluation follows from the dynamics alone. Where a stock compounds rather than adds, the logarithm is the ergodic-average object. The ceiling `ln N` also has an information-theoretic reading as the maximum an ensemble of `N` equiprobable categories can carry [Shannon, 1948], and a control-theoretic one as requisite variety [Ashby, 1958] — though this framework does **not** identify surplus with Shannon entropy, and the reading is licensed to bound and not to explain.

The functional form is nevertheless underdetermined by the evidence, and the exposure is worse at internal scope than at group scope. At least two published alternatives fit the same qualitative facts: social impact rising as a power of the number of sources with an exponent below one [Latané, 1981], and network value growing as `n log n` rather than as `n^2` [Briscoe et al., 2006]. All three curves are decelerating, monotone and threshold-free. **Nothing assembled in this corpus discriminates between them.** Latané's law is the sharper embarrassment here, because it was established with psychological data on individual judgement — the very domain this paper claims — a decade before the physical analogies arrived [Nowak et al., 1990]. What the framework's form has that the alternatives do not is an axiomatic characterisation and an explicitly declared design axiom at the point where the choice is made; what it does not have is a discriminating measurement. Section 13 carries this as an open limitation, not a resolved question.

The storage law. The dynamical engine, presupposed from Curzi [2026a]:

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t
```

with `Q` stored order, `S` the per-period input bounded by `ln N`, `delta` the retention loss, `eta` the conversion efficiency and `C` the operating cost. The map is Lipschitz with modulus `1 − delta`, so `delta > 0` makes it a contraction with a unique start-independent steady state `Q* = (eta S − C)/delta` [Curzi, 2026a, Banach, 1922, Elaydi, 2005]. The quartet of a level, an inflow, an outflow and a delay is the standard vocabulary of system dynamics, formalised four decades before this framework [Forrester, 1961], and `Q`, `eta S` and `C` occupy exactly those roles. The headroom gauge, its three regimes, the maximum one-step increment `r_max = h_t`, and the separate admissibility and sustainability constraints are presupposed from Curzi [2026b].

21.3.2 3.2 The scope substitution, and the two postulates it costs

The claim is that applying the surplus functional inside one observer is a derivation and not an analogy. Here is what that requires and here is the check.

The four axioms quantify over **states in an inner-product space**. S1 makes `F` depend on the pair only through `u`; S2 sets `f(0) = 0`; S3 makes `g(u) = e^{f(u)}` affine in `u`; S4 sets `f(1) = ln N`. **Not one of the four contains a term referring to persons, agents, groups, or ownership of states.** They refer to states and to the geometry between them, so the theorem that follows is silent on who holds the states.

The theorem is scope-free. What changes between group scope and internal scope is not the theorem; it is the encoder. At group scope the encoder maps a person, a team or a discipline to a unit state. Here it maps a *framing* to a unit state. Everything downstream of the encoder is identical, because everything downstream is algebra on encoded states. That leaves exactly two things to badge, and they are badged rather than assumed.

> **THE FRAMING ENCODING [POSTULATE].** A framing — a way of construing a situation, with its own primitives, its own salient variables and its own inference rules

— admits representation as a unit vector in a common inner-product space, such that two framings that make the same predictions on the same inputs map to parallel states and two framings whose predictions are uncorrelated map to orthogonal ones.

THE INTERNAL ROSTER REDUCTION [POSTULATE]. For an observer holding `N` nominal framings, `N` is the number of framing categories the observer claims to span and `u` is a representative non-parallelism across the observer's framing pairs — normally the mean pairwise `u` over all `N(N−1)/2` pairs. `N_eff = 1 + (N−1)u` is then read as the audited framing count.

These are the internal-scope twins of the group-scope roster reduction [Curzi, 2026c] and they fail in the same way. **Because the triangle inequality fails for geometric triples, a mean pairwise `u` does not determine the joint geometry**, and an observer whose framings are pairwise moderately different may still have all of them lying in a two-dimensional subspace. The fifth instrument of Section 6.3 exists to catch this and is mandatory. The failure has an empirical counterpart that strengthens rather than weakens the warning: psychological similarity is measurably asymmetric and violates the triangle inequality [Tversky, 1977], so if `u` is read as a similarity measure over construals it cannot be assumed to behave metrically at all.

With those two postulates in place, `F_int = ln(1 + (N−1)u)` follows from the surplus functional by substitution and requires no further assumption. That is what "derivation, not analogy" means here, stated precisely enough to be attacked: **attack the encoder, not the algebra.**

An earlier expository formulation of this framework carried the same statement as an image — a tower whose windows face different ways, which sees more because of the facing and not because of the glass — together with a discipline of attention: notice which way each window faces before counting them. Under the encoding above the image becomes an equation term by term, and it is worth setting the correspondence out because it is exact rather than suggestive.

```
the one tower          ->   a single observer; a single Q; one storage account
a window               ->   a framing; one encoded unit state
which way it faces     ->   the state's direction in the framing space
many windows,          ->   large N, small u   ->   N_eff -> 1,  F_int -> 0
  one aspect
windows set wide       ->   large u            ->   N_eff -> N,  F_int -> ln N
the view from the      ->   F_int = ln N_eff, the audited log-count
  tower
```

The image's first teaching — that a tower is not improved by more glass facing the same way — is the extremality result: `F = 0` if and only if the states are parallel. Its second, that three windows set wide are worth more than a gallery, is strict concavity, and Section 4.6 prices it. The correspondence is exact because the same functional form stands on both sides; what the earlier formulation reached by observation, the axioms reach by proof from four premises, one of which is a design choice.

The extremal case is immediately recognisable as a clinical description. A person with forty ways of describing their situation, all of them the same description in different words, has `u = 0`, `N_eff = 1` and surplus `ln 1 = 0`. They will experience this as thoroughness. It is a single window with forty panes.

21.3.3 3.3 The participation-ratio identity

The roster reduction chooses the *mean pairwise* aggregator, and a reader is entitled to ask why that average and not another. The answer is that it is not an average of convenience.

Let `G` be the Gram matrix of the observer's `N` encoded framings, so `G_ij = x_i . x_j` and `tr G = N`. The Rényi-2 participation ratio of that Gram spectrum — the standard effective-count

statistic, used unchanged across physics, ecology and portfolio theory [Rényi, 1961, Hill, 1973, Jost, 2006] — is `PR = (tr G)^2 / ||G||_F^2`. Since `||G||_F^2 = N + N(N−1)(1 − u_bar)` when `u_bar` is the mean pairwise `1 − (x_i . x_j)^2`,

```
N/PR = 1 + (N−1)(1 − u_bar)

THE PARTICIPATION−RATIO IDENTITY     N_eff = 1 + (N−1) u_bar = N + 1 − N/PR
```

The audited framing count is exactly the participation-ratio diversity of the observer's own framing ensemble, reparametrised into the form the surplus law uses. Declaring the mean pairwise aggregator therefore commits the analyst to a standard, well-understood diversity measure rather than to an improvisation, and the defence is arithmetic rather than rhetorical. The identity also supplies a second, independent route to `N_eff`: an analyst who can build the Gram matrix reads the effective count off the spectrum without ever forming a pairwise mean. It is verified numerically against direct computation at several roster sizes to better than `1e−12` [Curzi, 2026i], and the roster-scale form of the same identity, `M_eff = m + 1 − m/PR`, is verified independently at 1.7815458788 with error `0.00e+00` [Curzi, 2026].

The identity does not license the Hill-number literature as support for `N_eff`, and this distinction is the one a reviewer will press. `PR` *is* a Hill number of order two — it is `1/sum_i p_i^2` for the normalised Gram spectrum `p_i = lambda_i / N` — and the citations above attach to it exactly. `N_eff = N + 1 − N/PR` is an **affine reparametrisation** of that Hill number, and an affine transformation of a Hill number is not a Hill number. What the identity buys is that the aggregator is not arbitrary; what it does not buy is external grounding for the affine form, which remains the framework's own design axiom.

What the identity does not repair is the triangle-inequality failure noted above. `PR` and the mean pairwise `u` carry the same information because they are the same quantity; neither recovers the joint geometry that a mean over pairs discards.

21.3.4 3.4 Identification: what no amount of measurement recovers

This section states the objection that governs every applied reading in the paper, and it is stated here rather than in a footnote.

The framework's central formal claim is a reflexive loop, `T(x) = W(x, M(x))`: the aggregate influences the member and the members constitute the aggregate [Curzi, 2026]. **That is exactly a linear-in-means reflexive specification, and the endogenous social effect in such a specification is not identified from observational data** [Manski, 1993]. It cannot be separated from contextual effects or from correlated unobservables, no matter how much data is collected. Influence and selection are generically confounded in observational network data as well [Shalizi and Thomas, 2011], and the forensic consequences of ignoring that are documented [Lyons, 2011].

The objection reaches this paper by two distinct routes, and both must be conceded.

Route one: the framings inside one observer are not independent draws. The collapse mechanisms of Section 5 are precisely a within-observer influence structure — earlier framings condition later ones, and the later ones are recruited by a state that the earlier ones helped produce. The resulting `u` is therefore an equilibrium of an interaction process, not an exogenous property of the set. The classical formal apparatus for such a process is a row-stochastic influence matrix whose consensus is the dominant left eigenvector [DeGroot, 1974], with an anchoring term that yields a non-consensus fixed point rather than collapse [Friedkin and Johnsen, 1990] — which is, formally, the same device as this framework's leak `delta`.

Route two: the two-track ledger of Section 7.3 requires an external measurement of the environment, and there is no internal substitute. The over-framing ratio `r` and the environment `E` enter the moment as a product, so `M` alone cannot separate them. If `E` is reported by the same observer whose framing is being audited, `r` is estimated using the very quantity that the distortion acts on, and the two errors conceal each other exactly.

Which remedy this paper adopts: none. The known options are randomisation, an instrument, partial-population variation, or stipulating the reference set and defending the stipulation, of which the last is the most realistic and has an explicit template [Brock and Durlauf, 2001]. This paper stipulates the framing inventory in advance and freezes it (Specification SF4 and SQ2), which is a step toward that template and is not a substitute for it. **Every applied claim in this paper is therefore a modelling position rather than an identified empirical result, and where the text reads as though a measurement settled something, this sentence governs.**

The Lucas critique, turned on this paper's own parameters. The framework asserts that a described system changes its behaviour when the description is used, and then treats `delta`, `eta`, `N` and `u` as approximately constant over a survey window. Both cannot be true without qualification [Lucas, 1976]. The classification of a person interacts with the person classified, who changes in response, so the referent moves [Hacking, 1995]; and a framework that tells people their capability obeys a decay law should ask what that assumption will make true of them [Ferraro et al., 2005]. The position this paper takes is the narrow one: **the parameters are treated as stable only over a declared window, and only while the stored-order proxy is not itself a reported target.** Where either condition fails, the constancy assumption fails with it. That is a real restriction on where the model may be used and it is stated as one.

The counterweight, which is also real. A theory becoming true by being used is not only a hazard; it has been observed in the field with dates and price data [MacKenzie and Millo, 2003], and the reflexive doctrine this framework formalises exists in the literature without an equation attached [Soros, 2013, Merton, 1948]. Supplying an equation for it is a genuine contribution. But the effect size such reflexivity implies is larger than the psychological literature supports — classroom self-fulfilling prophecies are real, small, non-accumulating and often dissipating, and most expectations turn out to be accurate rather than causal [Jussim and Harber, 2005]. The framework should not draw on reflexivity for more than that.

21.3.5 3.5 Analogy discipline

Three imports from physics and engineering appear in this paper and each is labelled at first use.

The leaky integrator is a shared mathematical form, not a derivation from physical premises. The self-storage law and a first-order discrete filter are governed by the same difference equation, so every theorem about first-order discrete systems applies to both [Oppenheim et al., 1996, Elaydi, 2005]. Nothing about charge or current is claimed to hold of a person. The same label attaches to the statistical twin of the same object, exponential smoothing [Holt, 1957].

The participation ratio is a shared mathematical form with a genuine domain overlap. It is the same statistic in physics, in ecology and here, computed on a Gram spectrum in every case [Rényi, 1961, Hill, 1973]. This is the one import in the paper that is not merely formal, because the object it is computed on is of the same kind in all three settings.

The quadratic reconciliation cost is an analogy and is used as motivation only. The pair count `n(n−1)/2` is exact combinatorics; the claim that internal maintenance load is proportional to it is a postulate about a person, imported from the communication-path argument for groups [Brooks, 1975, Conway, 1968]. **No result in this paper depends on the constant of proportionality**, only on the fact that the cost is superlinear while the surplus is logarithmic.

The wider licence to carry physical formalism into psychological and social systems is not one this framework can grant itself: it was claimed, contested and partly established by others between 1971 and 2022, and the programme has been surveyed and bounded [Weidlich, 1971, Castellano et al., 2009, Jusup et al., 2022]. **The evidentiary bar that literature sets is prediction of observed data, not internal consistency** [Helbing and Molnár, 1995]. This paper does not clear that bar and does not claim to. The generous escape route is available and is taken explicitly: physical laws hold inside shielded, stable arrangements, so a law need not be universal to be real provided the arrangement within which it is claimed to hold is specified [Cartwright, 1999]. Section 12.5 specifies the arrangement.

One hazard is named because it is the characteristic failure of this exact genre. Neoclassical economics was built by copying mid-nineteenth-century energetics and the copy was defective: utility was made to play the role of energy **without the conservation law that gave energy its content** [Mirowski, 1989]. This framework's `Q` is a stock that accumulates, leaks and is spent — the grammar of a conserved quantity. **Nothing in this paper is conserved.** `Q` is not conserved, `F_int` is not conserved, and the only invariance the paper claims is the one stated and tested in Section 12.6. Saying so is the answer to Mirowski, and the answer is a concession.

21.3.6 3.6 Standing honesty clauses

Five clauses hold throughout and are restated in Section 13.

1. **The surplus functional is pure mathematics and makes no empirical or metaphysical claim.** It proves a functional form follows from four axioms. It does not prove that anyone's framings obey it. The empirical content of this paper is the modelling proposal of Section 3.2 and it is falsifiable at H1.
2. **Surplus is a lower bound** on the ensemble's information, not the information itself. An `N_eff` of 3.5 does not mean the observer holds exactly three and a half independent views; it means the measured geometry guarantees at least that much.
3. **The triangle inequality fails for geometric triples.** Surplus may not be chained across three framings as though it were a distance, which is why Section 6.3's subspace check is in addition to pairwise measurement and not an alternative to it.
4. **Surplus is not Shannon entropy or mutual information.** The logarithm is not a Shannon logarithm and entropy intuitions do not transfer wholesale.
5. **Nothing in this paper is a claim about neurons, about consciousness, or about physics.**

21.4 4 The internal surplus and what it costs to hold

21.4.1 4.1 What is proven, what is chosen, what is postulated

Element	Status
`f(u) = ln(1+(N−1)u)` unique under S1–S4	**PROVEN** [Curzi, 2026]
Strict monotonicity, positivity, strict concavity	**PROVEN** [Curzi, 2026]
Subadditivity on the parameter `u`	**PROVEN** [Curzi, 2026]
Extremality: `F = 0` iff parallel, `F = ln N` iff orthogonal	**PROVEN** [Curzi, 2026,r]
Range `0 <= F <= ln N`	**PROVEN** [Curzi, 2026]
The participation-ratio identity `N_eff = N + 1 − N/PR`	**PROVEN** (Section 3.3); verified numerically [Curzi, 2026i,?]
S3, the affine effective count	**DESIGN AXIOM — chosen, not derived**
Framings admit unit-vector encoding	**POSTULATE** (Section 3.2)
Mean pairwise `u` summarises an observer's framing set	**POSTULATE** (Section 3.2)
`u = 1 − rho^2` from error decorrelation	**POSTULATE** on the encoding; algebra thereafter
Reconciliation cost grows as `n(n−1)/2` internally	**POSTULATE** (Section 4.4)

S3 remains the design axiom and is flagged here as everywhere. That a second genuinely different framing counts for exactly as much, in `u`-units, as a fortieth is not a discovery about minds; it is the consequence of having chosen the effective count to be affine in `u`. A measured reason to replace S3 would require rebuilding this section around the replacement.

21.4.2 4.2 Validation gates for the internal surplus

The framework applies a seven-gate protocol to every formula it uses. The gates for `F_int = ln(1 + (N−1)u)` are reported in full, including the two that are not passed.

Gate	Verdict
G1 DIMENSIONS	`u` is dimensionless by construction, being a squared sine. `N` is a pure count. `1 + (N−1)u` is dimensionless, so the logarithm is well formed and `F_int` is in nats. Sound. **`N` and `u` must come from the same encoding**: an `N` counted over one taxonomy of framings and a `u` measured over a different one do not compose, and this is the commonest arithmetic error at this scope.
G2 SCALE TYPE	`N` is a count on a ratio scale with a true zero. `u` is ratio-scaled with a true zero, since `u = 0` is a genuine zero meaning exact parallelism rather than a convention, **which is why the effective count may be divided, ratioed and compared.** The prerequisite is that the encoding of framings into states is itself defensible. If framings are only rank-ordered by similarity, `u` is ordinal and `F_int` demotes to a heuristic ordering [Curzi, 2026m]. Whether the quantitative structure exists at all has not been tested; see Section 13.
G3 LIMITING CASES	`u -> 0`: `N_eff -> 1`, `F_int -> 0` — the parallel wall at internal scope, many framings and one view. `u -> 1`: `N_eff -> N`, `F_int -> ln N`, the ceiling. `N -> 1`: `F_int -> 0` for every `u`, correct, since one framing has nothing to be non-parallel to. `N -> infinity` at fixed `u > 0`: `F_int -> infinity` logarithmically — **the model has no internal ceiling on framing count, which is wrong in fact**, and the correction belongs to the cost term of Section 4.4 rather than to the surplus term. The divergence names a real failure: an observer who keeps adding framings without bound.
G4 FIXED POINTS	`u = 0` is the degenerate fixed point, invariant under any `N`. The substantive invariant is `N_eff`: **any pair `(N, u)` with the same `1 + (N−1)u` yields identical surplus**, so `N = 3, u = 1.0` and `N = 12, u = 0.1818` are interchangeable in surplus and distinguishable only by cost. That equivalence class is the section's main design lever.

Gate	Verdict
G5 IDENTIFIABILITY	`N` is recoverable by enumeration; `u` by the instruments of Section 6, of which error decorrelation is the only properly quantitative one. **The failure to declare: `N` and `u` are not separately recoverable from `F_int` alone.** Measuring surplus and inferring both is impossible; `u` must be measured independently. Reporting `N_eff` without reporting `(N, u)` separately therefore fails this gate and is prohibited by SF4.
G6 FALSIFICATION	H1. **Kill condition:** framing count predicts as well as the audited count, or the shuffled control matches the model. If `N` alone does the work, the `u` term is decoration.
G7 REPRODUCIBILITY	**The hard gate, and it is not passed.** The requirement is that two analysts independently encoding the same observer's framings recover the same `u` within tolerance, and that the same operationalisation applied to different observers yields parameters with stable meaning. Error decorrelation is the most promising route because it asks nobody to judge similarity. **This is untested and is stated as untested.**

21.4.3 4.3 The marginal value of the n-th framing

Strict concavity has a blunt consequence at internal scope. Fix `u` and ask what the `n`-th framing adds:

```
Delta F(n) = ln(1 + (n−1)u) − ln(1 + (n−2)u)

TABLE 1 — MARGINAL SURPLUS OF THE n−th FRAMING (nats; computed)

   n  |  u=0.20 |  u=0.35 |  u=0.50 |  u=0.70 |  u=1.00 |
  ----+---------+---------+---------+---------+---------+
    2 |  0.1823 |  0.3001 |  0.4055 |  0.5306 |  0.6931 |
    3 |  0.1542 |  0.2305 |  0.2877 |  0.3448 |  0.4055 |
    4 |  0.1335 |  0.1872 |  0.2231 |  0.2559 |  0.2877 |
    5 |  0.1178 |  0.1576 |  0.1823 |  0.2036 |  0.2231 |
    6 |  0.1054 |  0.1361 |  0.1542 |  0.1691 |  0.1823 |
    7 |  0.0953 |  0.1198 |  0.1335 |  0.1446 |  0.1542 |
    8 |  0.0870 |  0.1070 |  0.1178 |  0.1263 |  0.1335 |
    9 |  0.0800 |  0.0966 |  0.1054 |  0.1121 |  0.1178 |
   10 |  0.0741 |  0.0881 |  0.0953 |  0.1008 |  0.1054 |
   12 |  0.0645 |  0.0749 |  0.0800 |  0.0839 |  0.0870 |
   15 |  0.0541 |  0.0612 |  0.0645 |  0.0670 |  0.0690 |
   20 |  0.0426 |  0.0468 |  0.0488 |  0.0502 |  0.0513 |
   30 |  0.0299 |  0.0319 |  0.0328 |  0.0334 |  0.0339 |
   40 |  0.0230 |  0.0242 |  0.0247 |  0.0250 |  0.0253 |

   ratio  gain(2) / gain(40):
```

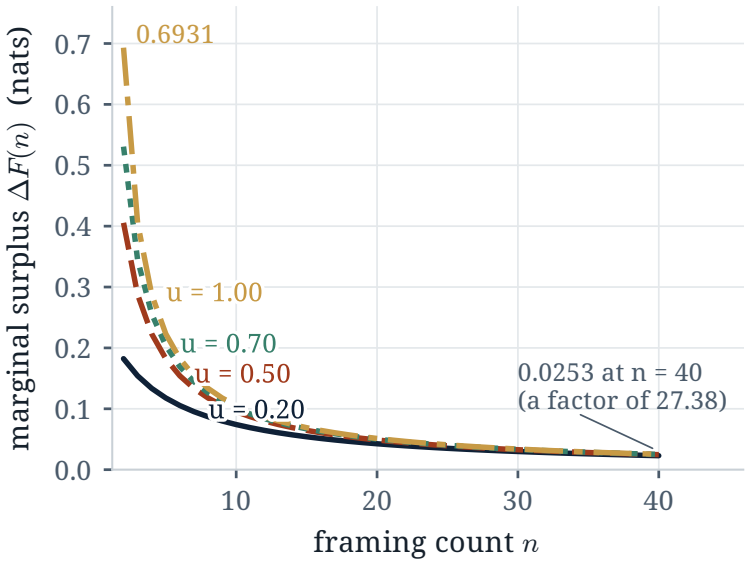


Figure 21.1: **The second framing is worth twenty-seven times the fortieth, and the ratio is a theorem.** Marginal surplus `Delta F(n) = ln(1+(n−1)u) − ln(1+(n−2)u)` against framing count, at four values of non-parallelism. At `u = 1.00` the second framing returns 0.6931 nats and the fortieth 0.0253, a ratio of 27.38; at `u = 0.20` the same ratio is 7.93. The `u = 0.35` column of Table 1 is omitted because the validated palette holds four colours and no fifth. Computed from the closed form with the parameters shown, not measured; generated by `figs/src/CCP-021_fig01.py`.

```
u=0.20  7.93x     u=0.35  12.41x     u=0.50  16.42x
u=0.70 21.19x     u=1.00  27.38x
```

At `u = 1.0` the second genuinely different framing is worth 0.6931 nats and the fortieth is worth 0.0253 — a factor of 27.4 for exactly the same effort. The pattern holds across the whole table because it is concavity, and concavity here is a theorem rather than a tendency. The second genuinely different framing is the largest gain that will ever be bought; the fortieth is a glazing bill.

Cumulatively:

```
TABLE 2 — CUMULATIVE INTERNAL SURPLUS F_int (nats; computed)

 n  |  u=0.20 |  u=0.35 |  u=0.50 |  u=0.70 |  u=1.00 |
----+---------+---------+---------+---------+---------+
  1 |  0.0000 |  0.0000 |  0.0000 |  0.0000 |  0.0000 |
  2 |  0.1823 |  0.3001 |  0.4055 |  0.5306 |  0.6931 |
  3 |  0.3365 |  0.5306 |  0.6931 |  0.8755 |  1.0986 |
  4 |  0.4700 |  0.7178 |  0.9163 |  1.1314 |  1.3863 |
  6 |  0.6931 |  1.0116 |  1.2528 |  1.5041 |  1.7918 |
 10 |  1.0296 |  1.4231 |  1.7047 |  1.9879 |  2.3026 |
 20 |  1.5686 |  2.0347 |  2.3514 |  2.6603 |  2.9957 |
 40 |  2.1748 |  2.6844 |  3.0204 |  3.3429 |  3.6889 |
```

21.4.4 4.4 Reconciliation cost is quadratic

The central asymmetry of group design is that surplus grows logarithmically while coordination cost grows quadratically, because `n` members generate `n(n−1)/2` communication paths [Brooks, 1975, Conway, 1968]. The same asymmetry is postulated inside a single observer.

THE INTERNAL RECONCILIATION COST [POSTULATE]. An observer holding `n` framings must, to use them as framings rather than as a list, reconcile them pairwise: notice where two framings disagree about the same situation and hold the disagreement without resolving it prematurely. The number of such reconciliations is `n(n−1)/2`, and the maintenance load is therefore quadratic in `n` while the surplus is logarithmic.

The exponent here is defined by the model and is not estimated from data. The pair count is exact combinatorics; the proportionality between pair count and load is the postulate. No distributional or scaling claim is made anywhere in this paper, so the maximum-likelihood plus-Kolmogorov-Smirnov procedure that any such claim would require does not apply [Clauset et al., 2009] — and it would apply immediately to anyone who claimed to have measured this exponent rather than assumed it.

TABLE 3 — SURPLUS PER RECONCILIATION PAIR, at u = 0.5 (computed)

n	pairs	F_int	F per pair
2	1	0.4055	0.40547
3	3	0.6931	0.23105
4	6	0.9163	0.15272
5	10	1.0986	0.10986
6	15	1.2528	0.08352
8	28	1.5041	0.05372
10	45	1.7047	0.03788
12	66	1.8718	0.02836
20	190	2.3514	0.01238
40	780	3.0204	0.00387

Between two framings and forty, surplus rises by a factor of 7.45 while reconciliation pairs rise by a factor of 780; surplus per pair falls by a factor of 105. That collapse is the arithmetic behind a familiar report: the person who has read widely enough to hold twenty framings and cannot, in consequence, decide anything. **They are paying a quadratic bill for a logarithmic return, and it came due exactly where the arithmetic says it would.**

The same crossing is the reason a queueing argument belongs beside the combinatorial one. Waiting time in a loaded system rises without bound as utilisation approaches one [Kingman, 1962], so an observer running near capacity does not degrade gracefully as framings are added; the marginal reconciliation is the one that saturates. That is an analogy and is used as motivation only, per Section 3.5.

21.4.5 4.5 Net internal surplus and the optimal framing count

Let `eta` be the conversion of internal surplus into stored order, `c1` the per-framing upkeep and `c2` the per-pair reconciliation cost.

THE NET INTERNAL SURPLUS [DERIVED].

```
V(n) = eta * ln(1 + (n−1)u)  −  c1 * n  −  c2 * n(n−1)/2
```

The optimum `n*` is the largest `n` for which `V(n) > V(n−1)`, that is for which `eta [ln(1+(n−1)u) − ln(1+(n−2)u)] > c1 + c2 (n−1)`. **Since the left side falls like `1/n` and the right side rises linearly, `n*` always exists and is always finite.**

TABLE 4 — OPTIMAL FRAMING COUNT (computed)

```
case                        eta    u     c1      c2       n*    V(n*)
--------------------------  -----  ----  ------  ------   ----  -------
attentive, cheap upkeep     1.0    0.60  0.010   0.0020    20    1.9377
attentive, costly upkeep    1.0    0.60  0.030   0.0080     9    1.1999
high conversion             2.0    0.60  0.030   0.0080    14    3.2015
low non-parallelism         1.0    0.25  0.020   0.0050    11    0.7578

worked profile for the second case, V(n) at n = 1..10:
  1: -0.030   2: +0.402   3: +0.674   4: +0.862   5: +0.994
  6: +1.086   7: +1.148   8: +1.185   9: +1.200  10: +1.196
```

Three readings follow, all practical.

The optimum is flat near its peak. In the second case `V(8) = 1.185`, `V(9) = 1.200`, `V(10) = 1.196`. The penalty for holding eight framings instead of the optimal nine is 1.3 per cent. **The task is not to find `n*`; it is to avoid being at 2 or at 25.**

Raising conversion raises the optimum. Doubling `eta` from 1.0 to 2.0 moves `n*` from 9 to 14. An observer with a practice that converts a framing's output into something durable can afford more framings than one without, which is the internal form of the framework's headline result that `eta` is its cheapest lever [Curzi, 2026b].

Low non-parallelism lowers the peak far more than it lowers the optimum. The fourth case reaches `V = 0.758` at `n* = 11` with `u = 0.25`, against 1.200 at `n* = 9` with `u = 0.60`. **Holding more framings does not compensate for holding parallel ones: it costs more and returns less.**

The gates for `V(n)` are reported with the same candour as those for `F_int`.

Gate	Verdict
G1 DIMENSIONS	`ln(.)` is in nats; `eta` carries stored-order units per nat; `c1` is cost per framing per period; `c2` is cost per pair per period. **All three terms must be expressed in the same stored-order currency or the subtraction is meaningless** — the gate this formula most often fails, because effort and character are not naturally commensurable and making them so is a modelling decision that must be stated.
G2 SCALE TYPE	Requires `u` ratio-scaled and both cost terms ratio-scaled with a true zero. Costs elicited as ordinal ratings do not qualify; magnitude estimation on effort is the honest route [Stevens, 1957] and cross-modality matching against a physical effort continuum is the stronger one [Stevens, 1956].
G3 LIMITING CASES	`c1, c2 -> 0`: `n*` unbounded, matching `F_int`'s divergence and correctly locating it in the cost term. `c2 -> 0` with `c1 > 0`: `n*` finite but large — an observer whose framings never have to be reconciled with each other, which describes compartmentalisation and is a real and named condition. `u -> 0`: `V(n) = -c1 n - c2 n(n-1)/2`, strictly decreasing, so `n* = 1`. **With parallel framings the model says hold exactly one, which is correct and bleak.**

Gate	Verdict
G4 FIXED POINTS	`n*` solves `eta Delta F(n) = c1 + c2(n−1)`, the balance of a falling marginal return against a rising marginal cost: **the framing count at which one more window costs exactly what it shows.**
G5 IDENTIFIABILITY	`eta`, `c1`, `c2` are jointly identifiable only with variation in `n`. An observer who has held six framings for a decade provides no information about the cost curve. **The formula is therefore not fittable on cross-sectional self-report; it requires a deliberate, logged variation in framing count over time.** That is a demanding requirement and it is stated rather than hidden.
G6 FALSIFICATION	H2. **Kill condition:** measured performance flat in `n` over the tested range, or peaking where the model did not predict, or a linear cost model fitting as well — in which case the pair-cost postulate is wrong at internal scope.
G7 REPRODUCIBILITY	Untested. The cost parameters are almost certainly person-specific; **the form is the reproducible claim, not the values.** Report as single-subject unless and until cross-subject recovery is demonstrated.

21.4.6 4.6 Three windows set wide against twelve from one lineage

The group-scope comparison of three orthogonal members against twelve nearly parallel ones has an internal counterpart, and the internal lesson is sharper because the quadratic bill is paid by the same person who collects the surplus.

```
n =  3,  u = 1.00   ->   F_int = 1.0986      reconciliation pairs =   3
n = 12,  u = 0.15   ->   F_int = 0.9746      reconciliation pairs =  66
n = 20,  u = 0.12   ->   F_int = 1.1878      reconciliation pairs = 190
```

Three framings genuinely set wide beat twelve from one lineage, at one twenty-second of the reconciliation load. Twenty parallel framings finally exceed the three by 0.089 nats, having taken on 190 pairs instead of 3: a 63-fold cost increase for an 8 per cent surplus increase.

This answers the reader who concludes from Table 1 that the remedy is to read more. **Reading more raises `N`. It raises `u` only if what is read is genuinely non-parallel to what has been read already, and the default for a reader following their own taste is that it is not.** The organisational literature reaches the same place from the other side: fast, efficient learning homogenises a population and destroys the variance that produced its advantage [March, 1991, Levinthal and March, 1993], and the effect has been measured as an outright collapse in the effective count under social influence [Lorenz et al., 2011].

21.5 5 How non-parallelism collapses

21.5.1 5.1 Six mechanisms

u is not a property that is acquired and held. **It is a quantity that decays, and at internal scope it decays faster than at group scope, because there is no one else in the room to disagree.** Six mechanisms are named, each with its diagnostic signature. The labels are CM1–CM6 to avoid collision with the moment variable M_t of Section 7.

CM1 — THE INTERNAL CASCADE. Observers who condition on others' conclusions stop contributing information [Bikhchandani et al., 1992]. The same structure appears internally with the observer's *own earlier conclusion* as the reinforcing contact: framing B is consulted after framing A has produced an answer, and B is asked to check A rather than to see the situation. *Signature:* framings agree more the later they are consulted.

CM2 — THE CONSISTENCY DRIVE. A standing pressure toward consistency among simultaneously held construals acts, in this notation, as **a force rotating framings toward parallel** — a u-destroying gradient operating continuously and without supervision. *Signature:* framings that used to disagree quietly stop disagreeing, with no recollection of a decision to change either. **This paper's bibliography holds no entry for the mid-twentieth-century cognitive-consistency literature that documents the pressure, so the mechanism is carried here as a postulate rather than as a cited result**, and Section 13 repeats that.

CM3 — IDENTITY CAPTURE. A framing becomes *who the observer is* rather than *something they use*. It can then no longer be set aside for a comparison, so every other framing is evaluated from inside it and their encoded states rotate toward it. *Signature:* one framing is never the one that loses.

CM4 — MOOD-CONGRUENT RECRUITMENT. Which framings are available at a given moment is not independent of state. If low mood makes the pessimistic framings easier to reach, then at exactly the moments when non-parallelism is most valuable the accessible set is most parallel. *Signature:* the framing set feels smaller when things are difficult. **This is a u collapse that presents as a mood problem and is treated as one, which is why it persists.**

CM5 — THE SINGLE LINEAGE. All the framings came from one training, one mentor, one literature. Their disagreements are the disagreements internal to that lineage, which are real but small. *Signature:* the observer can name six framings and all six would be recognised by the same professional body. **The cited evidence available here is weaker than the mechanism deserves.** Diverse problem-solving groups can outperform higher-ability homogeneous ones under stated conditions [Hong and Page, 2004], but the formal result has been shown to follow from the modelling assumptions rather than from any fact about diversity [Thompson, 2014], and the subsequent exchange has narrowed rather than settled the question [Singer, 2019]. **The closest thing to direct evidence for this section's central claim is a forecasting result — that forecasters using many small, loosely related explanatory framings outperformed those organising everything under one large idea — for which this paper's bibliography carries no entry.** It is therefore described and not cited, and nothing here rests on it.

CM6 — PREMATURE INTEGRATION. The framings are collapsed into a synthesis before their disagreement has been used. Synthesis is the correct final move and the wrong early one. *Signature:* a unified view is reached quickly and the components' disagreement cannot afterwards be reconstructed.

21.5.2 5.2 A geometric decay model

The aggregate effect is modelled as geometric decay, which is the honest form given that these are multiplicative pressures, with a per-round collapse rate lambda:

```
u_t = u_0 (1 − lambda)^t      half-life = ln 2 / (−ln(1−lambda)) rounds

TABLE 5 — FRAMING-SET COLLAPSE, u_0 = 0.60, N = 6 (computed)
```

```
lambda=0.05    t:  0       1       2       4       8      16     half-life 13.51 rounds
   u              0.6000  0.5700  0.5415  0.4887  0.3981  0.2641
   F_int          1.3863  1.3481  1.3104  1.2365  1.0954  0.8417

lambda=0.10    t:  0       1       2       4       8      16     half-life  6.58 rounds
   u              0.6000  0.5400  0.4860  0.3937  0.2583  0.1112
   F_int          1.3863  1.3083  1.2326  1.0880  0.8292  0.4421

lambda=0.20    t:  0       1       2       4       8      16     half-life  3.11 rounds
   u              0.6000  0.4800  0.3840  0.2458  0.1007  0.0169
   F_int          1.3863  1.2238  1.0716  0.8015  0.4077  0.0811
```

At a 20 per cent per-round collapse rate, sixteen rounds take a six-framing observer from 1.386 nats to 0.081 — an audited count of 4.00 down to 1.08. The endpoint will be described as clarity. It is the parallel wall, reached from the inside, by a process that felt like thinking.

The decay rate `lambda` is a stipulated parameter in every table above. **It has not been measured, in this paper or anywhere in this corpus**, and the three values are chosen to bracket a plausible range rather than to report one.

21.5.3 5.3 Seven maintenance protocols

Each protocol targets a named mechanism, each is cheap, and each is an artefact rather than an intention, because artefacts have low `delta` [Curzi, 2026a]. The labels are `MP1`–`MP7`. **Whether any of them works is not established here**; they are design proposals with stated rationales.

MP1 — THE FRAMING REGISTER. A written list, one line per framing: name, primitives, what it is good at, what it is blind to, where it came from. A file, reviewed on a cadence. *Targets CM5, CM3.* The provenance column does the work — when four of six lines name the same lineage, `u` is low and the fact is visible without measuring anything.

MP2 — INDEPENDENT-FIRST, WRITTEN. On any question of consequence, run each framing to a written conclusion **before** consulting the next. *Targets CM1.* The highest-value protocol here, because the internal cascade is both the fastest destroyer of `u` and the easiest to block: the block is a piece of paper.

MP3 — TIME SEPARATION. Run framings on different days. State-dependent availability is partly a within-session effect. *Targets CM4, CM2.*

MP4 — THE DISAGREEMENT LOG. When two framings disagree, record the disagreement and its resolution, including "unresolved". *Targets CM2, CM6.* Measurement instrument and maintenance device at once: writing a disagreement down stops the consistency drive from quietly erasing it, and the log is the raw data for Section 6.

MP5 — DELIBERATE LINEAGE IMPORT. Learn a genuinely different trade to working competence, not conversational competence. *Targets CM5.* The expensive protocol, and the only one that raises `N` rather than protecting `u`. Table 1 prices it: holding two or three real framings, the next is worth 0.29 to 0.41 nats at high `u`, the largest increment available; holding twelve, 0.08.

MP6 — THE ADVERSARIAL PASS. Take the framing that is winning and argue it down using only the framing that is losing, in writing, to a conclusion. *Targets CM3.* "Using only" is load-bearing; the exercise fails if the winning framing is permitted to referee.

MP7 — DEFERRED SYNTHESIS. A written rule for how long disagreement must be held before integration is permitted. *Targets CM6.* Synthesis is not prohibited; it is scheduled.

The adoption cost of each protocol is itself a quantity the framework can price. Where a practice requires a sequence of steps each of which has a resistance, the branch may be weighted by that accumulated resistance rather than treated as equally available [Curzi, 2026]: a three-point difference in steepness is a factor of eight in expectation, which is why MP2 —

one sheet of paper — and MP5 — a second trade — are not comparable interventions however similar their descriptions look.

One constraint is not negotiable and needs restating at internal scope, because the person and the intervener are the same and it is easy to think it does not apply. The framework prohibits any intervention that removes a person's ability to decline. **A protocol that forces an observer to hold a framing they have declined is coercive, and the fact that the observer is the one applying it does not make it otherwise.** MP5, MP6 and the error-decorrelation instrument are offered as available instruments; someone who declines any of them has not failed a specification. The method is to make the alternative cheaply available and let the movement be a choice, whose worst outcome is that someone declines to move — and that is the property that makes it safe.

21.6 6 Measuring non-parallelism

21.6.1 6.1 Five instruments

Five instruments, in ascending order of rigour and of cost.

I1 — THE LINEAGE CENSUS [HEURISTIC]. From the register: how many distinct training origins are represented among `N` framings? Provenance is a prior, not a measurement, but it is the only instrument that works before there is any data and the only one that can prevent rather than diagnose.

I2 — THE COLLISION TEST. Over the last twenty consequential judgements, in how many did a second framing produce a conclusion the first had not already reached? **A framing that never collides is not a second window.** This fraction is not `u` and must not be reported as `u`; it is triage, and a value near zero is decisive on its own.

I3 — THE DISAGREEMENT RATE. From MP4's log: the proportion of logged questions on which a pair of framings disagreed on the answer, not merely on the reasoning. **Ordinal**; comparable across pairs within one observer, not across observers.

I4 — ERROR DECORRELATION. The quantitative instrument, and the only one that yields a `u` admissible in the formula. Record each framing's *prediction* on questions with later-observable answers; compute each framing's error series; compute the Pearson correlation `rho` between the error series of each pair. Then, under the error-vector encoding postulate,

```
u = 1 - rho^2
```

I5 — THE SUBSPACE CHECK. Section 6.3.

21.6.2 6.2 Error decorrelation, and what the estimator inherits

```
TABLE 6 — ERROR CORRELATION -> AUDITED FRAMING COUNT (N = 6; computed)

   rho       u        N_eff      F_int    % of ln 6
  -----   -------   ---------   -------  ----------
   1.00   0.0000      1.0000    0.0000        0.00
   0.95   0.0975      1.4875    0.3971       22.16
   0.90   0.1900      1.9500    0.6678       37.27
   0.80   0.3600      2.8000    1.0296       57.46
   0.70   0.5100      3.5500    1.2669       70.71
   0.50   0.7500      4.7500    1.5581       86.96
   0.30   0.9100      5.5500    1.7138       95.65
   0.00   1.0000      6.0000    1.7918      100.00
```

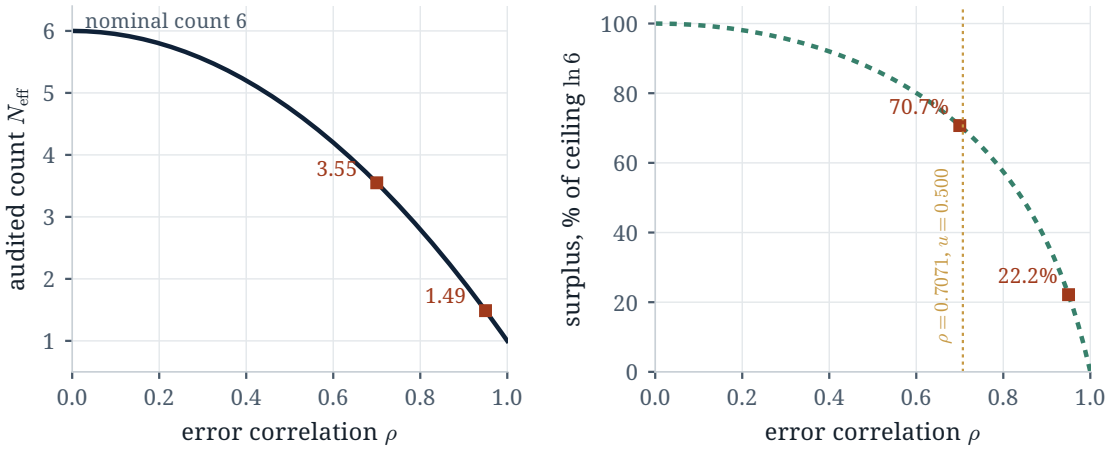


Figure 21.2: **Non-parallelism is cheap to obtain and expensive to lose.** The audited framing count `N_eff = 1 + (N−1)(1−rho^2)` at `N = 6` (left) and the resulting internal surplus as a percentage of the ceiling `ln 6 = 1.7918` nats (right), against the error correlation between framings. Framings correlated at `rho = 0.70` still return 70.7 per cent of the ceiling; framings correlated at `rho = 0.95` return 22.2 per cent. The landmark `rho = 0.7071` gives `u = 0.500` exactly — half the available non-parallelism — at which a pair returns 58.5 per cent of its own ceiling, not half. Plotted as two panels rather than on shared axes because a count and a percentage are not commensurable. Computed from Table 6's closed form, not measured; generated by `figs/src/CCP-021_fig02.py`.

Note the shape. Framings correlated at `rho = 0.7` — which most people would call "pretty similar" — still deliver 70.7 per cent of the available surplus, because `u = 1 − rho^2` punishes only high correlation. **Orthogonal framings are not required; framings that are not nearly identical are, and the gap between those two requirements is most of the practical content of this section.** Conversely `rho = 0.95`, which feels like a genuine second opinion, delivers 22 per cent.

The estimator's gates:

Gate	Verdict
G1 DIMENSIONS	`rho` is a dimensionless correlation in `[−1,1]`; `rho^2` in `[0,1]`; `u` in `[0,1]`. Sound. Note that `u` depends on `rho^2`, so **anti-correlated framings at `rho = −0.9` are treated as no more different than correlated ones at `rho = +0.9`** — correct under the geometric encoding, since `u = 1 − cos^2 theta` and antiparallel states are as redundant as parallel ones. This is the extremality condition stated on the squared alignment [Curzi, 2026r] and the opposition identity that follows from it: `u = 1 − rho^2` depends only on `|rho|`, so `rho = +1` and `rho = −1` both give `u = 0` and `F = 0` [Curzi, 2026]. A framing that is reliably the mirror image of another is not a second window.

Gate	Verdict
G2 SCALE TYPE	Requires the prediction errors to be at least interval-scaled for `rho` to mean anything. **If the framings produce only rank orderings, a rank correlation must be used and the resulting `u` stated as a heuristic ordering rather than a measured quantity.**
G3 LIMITING CASES	`rho -> +/-1`: `u -> 0`, identical or mirrored framings contribute nothing. `rho -> 0`: `u -> 1`, independent errors, the ceiling. There is no divergence to explain; the map is bounded on both sides, which is a point in its favour.
G4 FIXED POINTS	`rho^2 = 0.5`, i.e. `rho = 0.7071`, gives `u = 0.5` — **half the available non-parallelism, and a useful practical landmark.** The framework's earlier statement of this landmark as the "half-surplus point" is imprecise and is corrected here: at `u = 0.5` a pair returns `ln 1.5 / ln 2 = 58.5` per cent of its ceiling, not 50 per cent. What is exactly halved is the effective count's excess over one.
G5 IDENTIFIABILITY	`rho` is estimable from as few as 15 to 20 paired predictions, with wide confidence intervals; **the sampling error on `rho` propagates into `u` with a factor of `2|rho|`, so at high correlation a small error in `rho` is a large error in `u`.** Report the interval, not the point, whenever `|rho| > 0.7`.
G6 FALSIFICATION	The encoding claim is falsifiable: if `u` estimated this way fails to predict the incremental accuracy of adding a framing — if adding a low-`rho` framing improves held-out accuracy no more than adding a high-`rho` one — the encoding is wrong and the number is decoration.
G7 REPRODUCIBILITY	**This instrument's chief virtue is that it is the most likely of the five to pass**, because it involves no similarity judgement by anyone: it is computed from recorded predictions and recorded outcomes, so two analysts with the same log must obtain the same `rho`. Untested across observers.

Two exposures are worth naming beside the gates. The first is that any survey instrument used to elicit framings is itself sensitive to question wording and to what the question makes salient [Schwarz, 1999, Kahneman et al., 2006], which is one reason the instrument is built on recorded predictions rather than on recalled ones — and why introspective reports of one's own reasoning process should not be used to estimate `rho` [Nisbett and Wilson, 1977]. The second is that `rho` estimated on a short log is a small-sample statistic with the usual bias, and effective-count estimators computed from small samples are known to be biased downward in the analogous ecological setting [Chao et al., 2014].

21.6.3 6.3 The subspace check

Because the triangle inequality fails for triples, pairwise `u` can be uniformly moderate while all framings lie in a low-dimensional subspace. Assemble the error series into a matrix and examine the eigenvalue spectrum of its correlation matrix. **If two components carry nearly all the variance among six framings, the pairwise summary overstates the audit, and `N_eff` should be reported as the spectrum's participation ratio rather than as `1 + (N−1)u_bar`.** By the identity of Section 3.3 the two agree exactly when the geometry is what the pairwise mean implies, so the check is a comparison of a quantity against itself computed two ways — which is what makes it diagnostic.

The check has an external analogue that makes its importance concrete. Diversification benefit depends on *covariance*, not on count: `N` uncorrelated holdings and `N` perfectly correlated holdings have the same `N` and completely different behaviour [Markowitz, 1952], and measured correlations rise under stress, so the effective count falls exactly when it is needed [Longin and Solnik, 2001]. **A framing set is a portfolio, and the subspace check is the covariance audit.** The same warning is available from the other direction: diversity-stability correlations can arise from statistical averaging alone, with no mechanism at all [Doak et al., 1998], so an observed association between a high audited count and good judgement is not by itself evidence for this paper's mechanism.

21.7 7 The self as storage

21.7.1 7.1 The join, written out

The storage law with the moment substituted for the interaction surplus:

```
THE SELF–STORAGE LAW

    Q_{t+1} = (1 − delta) Q_t  +  eta ∗ M_t  −  C_t

    Q_t      stored order in one person: character, skill, memory, internalised structure
    M_t      the moment as lived this period, in nats, bounded by ln N
    delta    retention loss in [0,1]: forgetting, skill atrophy, drift of self
    eta      conversion efficiency >= 0: how experience becomes character
    C_t      cognitive load, effort, the operating cost of being a self
```

The external form of the law takes the interaction surplus of a group as its input [Curzi, 2026a]. The internal form takes the *moment* as its input. **Perspective and perception produce the moment; the moment becomes character.** The join is a substitution, and it runs in four steps:

```
STEP 1  CALIBRATION      r_t = D_t / P_t
STEP 2  THE MOMENT       M_t = min( r_t ∗ E_t ,  ln N )
STEP 3  VERIDICAL PART   V_t = E_t ∗ min( r_t , 1 )
STEP 4  THE DEPOSIT      Q_{t+1} = (1−delta) Q_t + eta M_t − C_t
```

`D` is the framing brought to the situation, `P` the situation's own resistance to it, `r` their ratio, and `E` the environment measured in nats of usable structure. **`E` is exogenous and must be measured by an instrument outside the observer**; that requirement is inherited without softening [Curzi, 2026m]. The cap at `ln N` is the surplus functional's range result plus a normalisation convention: **the richest moment available to a mind is the log of how many kinds of thing it can tell apart**, which is requisite variety in first-person form [Ashby, 1958] and was not chosen.

Before any of it is used, the status of Step 1 must be stated, because it governs Steps 2, 3 and 4. Applying the validation protocol to the moment equation returns a **demotion on the scale-type gate**: `D` and `P` are not instrumented on a common ratio scale with a shared true zero, so the *magnitude* of `r_t` is not a reportable measured quantity, and because the moment contains a division and a multiplication, ordinal data for either term demote it to a heuristic ordering that may not be printed as a fraction [Curzi, 2026m]. That verdict propagates straight down this chain.

> **THE JOIN IS STRUCTURAL, NOT YET QUANTITATIVE.** The architecture — moment feeds storage, storage does not inspect its input, the two tracks separate, stored error integrates with the cost term cancelling — is **exact algebra and stands as written.** So do the orderings it produces: over-framing accumulates faster than calibration, `Z` is zero at and below the fixed point and positive above it, correction load rises with accumulated error. **What is not yet reportable is any number obtained by putting a measured `r_t` through Step 2.** Every figure of that kind in this paper is an illustrative construction demonstrating the mechanism, and is labelled as such at each use.

The route out is constructive and costed, and it belongs in the same breath as the demotion. If `D` and `P` are measured as Rasch person-measures on one common linked scale, the distortion becomes a logit difference, `Delta = |ln(D/P)| = |theta_D − theta_P|`, genuinely ratio-scaled with a true zero at calibration [Rasch, 1960, Andrich, 1988]. The price is on the order of **123 items per term to establish a half-logit distortion at 95 per cent confidence**. That is the study which would convert this section's structure into this section's arithmetic, and it is specified rather than wished for. Whether the underlying attributes are quantitative at all is a prior question and Section 13 concedes it [Michell, 1997, Luce and Tukey, 1964].

21.7.2 7.2 What transfers, and why

Proven. Everything the external storage law proves, because it is the same equation: the exact solution, the one-step growth condition, the uniform bound, the sustainability ceiling, the steady state, the headroom gauge and the maximum one-step increment all carry over unchanged [Curzi, 2026a,b], because none of them uses any property of the input except that it lies in `[0, ln N]`. Step 2's cap guarantees exactly that. **The theorems transfer because the range transfers, and for no other reason.**

Postulated, badged. That a person's stored order is a scalar. That it decays at a constant proportional rate. That the moment and the environment can be placed on a ratio scale. That `eta`, `delta`, `N` and `C` hold approximately constant over the horizon on which the gauge is read. That the chosen proxy stands in a stable relation to `Q`. **None of these is a theorem, and a headroom figure for a person is only as good as the five measurements behind it.**

21.7.3 7.3 The two-track ledger and the stored-error recursion

The storage law **does not inspect its own input**: no term in `Q_{t+1} = (1−delta)Q_t + eta M_t − C_t` asks whether `M` was earned. An over-framed observer with `r > 1` manufactures moment out of their own priors and stores it exactly as if the environment had supplied it. The engineering answer is not to fix the variable. It is to keep two books.

```
THE LIVED TRACK     Q_t      integrates M_t = min(r_t E_t, ln N)      what you have become
THE EARNED TRACK    Q^V_t    integrates V_t = E_t min(r_t, 1)         what the world put there
THE ERROR           Z_t = Q_t − Q^V_t                                  what you put there
```

Subtracting the two recursions, the cost term cancels exactly:

```
Z_{t+1} = (1 − delta) Z_t + eta (M_t − V_t)

and at steady state,   Z_inf = eta (M* − V*) / delta = eta E* (r − 1) / delta   for r > 1
```

Stored error accumulates by the same leaky integration as stored knowledge, at the same rate, and its equilibrium does not depend on cost at all. That is a two-line derivation and it is the most useful result in this section. Three facts follow immediately.

Fact 1. `Z` is a leaky integrator with the same pole. Stored error decays at the same rate as stored knowledge, which is why recalibration works at all: **stop manufacturing and the accumulated error forgets itself on the same half-life as everything else.**

Fact 2. `Z_inf` is independent of `C`. Working harder, being busier, paying more to exist — none of it changes the equilibrium quantity of stored error. Only `r`, `E`, `eta` and `delta` do.

Fact 3. `Z = 0` for every `r <= 1`. The under-framed observer stores no error. They store less, and their equilibrium can be negative, which is the model refusing the regime. But what they do store, they earned. **Flooding is a capacity problem; projection is an integrity problem, and they require opposite corrections.**

The veridical throughput `V = E min(r,1)` and the confabulated fraction `1 − 1/r` are established with a seeded audit elsewhere in the corpus, where an observer at `r = 1.724138` returned a false-discovery rate of 0.4200 exactly and 63 per cent of the steady-state store was error [Curzi, 2026m]. **This paper contributes the recursion, not the audit.**

The gates for the stored-error recursion:

Gate	Verdict
G1 DIMENSIONS	`M − V` in nats; `eta(M−V)` in `Q`-units; `Z` in `Q`-units, as a difference of two `Q` quantities must be. Sound.
G2 SCALE TYPE	`Z` as a difference is valid on an interval scale. **`Z/Q` — the error fraction, which is the quantity anyone actually wants — requires `Q` ratio-scaled with a true zero and inherits the parent law's conditional verdict.**
G3 LIMITING CASES	`r −> 1`: `M = V`, `Z_inf = 0`, no stored error ever accumulates. `r −> infinity` with the cap in force: `M −> ln N`, `V −> E`, `Z_inf −> eta(ln N − E)/delta`, **bounded** — saturation, not divergence. `r < 1`: `Z = 0`. `delta −> 0`: `Z` grows without bound, error that is never forgotten.
G4 FIXED POINTS	`Z_inf = eta(M*−V*)/delta`, independent of `C` because the cost term cancels. **The equilibrium quantity of stored error in a mind does not depend on how expensive that mind's life is.** `Z = 0` iff `r <= 1`.
G5 IDENTIFIABILITY	Requires `r` and `E` measured independently. **`Z` cannot be recovered from `M` alone, because `M` is a product and the split is exactly the unidentified direction** (Section 3.4).
G6 FALSIFICATION	H5. **Kill condition:** correction-event rate uncorrelated with `r` after controlling for `M`.
G7 REPRODUCIBILITY	Unattempted; inherits the open question of whether `r` is trait-like or item-set-specific.

21.7.4 7.4 The correction load

Over-framing shows up in the books not as bad `Q` but as rising `C`, on a lag. Stated as an equation:

```
THE CORRECTION LOAD       C_t = c_0 + kappa * Z_t
```

with `c_0` the baseline operating cost in `Q`-units per period and `kappa` the correction cost per `Q`-unit of stored error per period. **Linearity in z is a postulate, badged**: the simplest admissible form, chosen for tractability and not because anyone has measured the shape. A convex form — error that becomes disproportionately expensive — is equally plausible a priori.

`kappa -> 0` returns the single-track law, which is the correct nesting. Large `kappa` drives `C` above `eta ln N` and the sustainability ceiling forecloses everything: **the over-framed observer's regime becomes unsustainable by no other mechanism than its own accumulated error.** At equilibrium `C_inf = c_0 + kappa eta (M*-V*)/delta`. **The lag is `tau = 1/delta`: the cost rise arrives one time constant after the over-framing that caused it, which is why the condition is stable for years.** The falsifier is clean and within-subject: **if `C` does not rise on a lag of order `tau` following a sustained increase in `r` at fixed `E`, the correction load is fictitious and should be deleted from the formulary.** It has not been run. Without an external instrument for `E`, `kappa` and the `Z` scale trade off exactly and neither is identified.

21.7.5 7.5 Retention loss and its three channels

`delta` is the fraction of stored order lost per period with no maintenance and no input. For a person it decomposes into three channels that are routinely confused and have different remedies:

```
delta_forget   memory decay: what is no longer retrievable
delta_atrophy  skill decay: what can still be described but no longer executed to standard
delta_drift    identity decay: what no longer fits, because the self it belonged to has moved
```

`delta_drift` is the channel that has no ordinary name and does the most work. It is not loss of a capacity; it is loss of the *relevance* of a capacity to the person now holding it. A trained cellist who has not touched the instrument in a decade has lost `Q` to atrophy. A trained cellist who still plays but for whom playing no longer organises anything — no longer sits under a decision, no longer selects a room, no longer picks the people — has lost `Q` to drift. Nothing was forgotten. **The stored order stopped being structure and became furniture.** Drift is why stored order must be re-anchored rather than merely rehearsed.

The channels act on the same stock in the same period and are approximately independent, so they compose multiplicatively on retention rather than additively on loss:

```
keep  = (1 - delta_forget)(1 - delta_atrophy)(1 - delta_drift),    delta = 1 - keep
```

At the worked weekly values `0.020`, `0.008` and `0.004`, `keep = 0.980 x 0.992 x 0.996 = 0.96827136` and `delta = 0.03172864` per week. The naive sum is 0.032. **The overcount is 0.00027, which is 0.85 per cent — negligible at these rates and not negligible at large ones**, since four channels of 0.15 give 0.4780 against a naive 0.60, an overcount of 20.3 per cent. Independence is the falsifiable claim: **if the measured composite differs significantly from the product of separately measured channels, the channels interact and the multiplicative form is wrong.** Three zero-input windows test it directly. The practical obstacle is that a window suppressing rehearsal also suppresses use, so `delta_forget` and `delta_atrophy` may not be separately identified; where they cannot be, the composite is reported and the confounding stated.

```
tau     = 1/delta                    periods (time constant)
t_half  = ln 2 / (-ln(1-delta))      periods (half-life of unmaintained stock)
3 tau  ~ 95% of the way to equilibrium;   5 tau ~ 99%
```

`delta` per week	`tau` (wk)	half-life (wk)	3`tau` (wk)	3`tau` in months
0.005	200.00	138.28	600.00	138
0.010	100.00	68.97	300.00	69
0.020	50.00	34.31	150.00	35
0.031729	**31.52**	**21.50**	**94.55**	**21.8**
0.050	20.00	13.51	60.00	13.8
0.100	10.00	6.58	30.00	6.9
0.200	5.00	3.11	15.00	3.5

Read the highlighted row as the sentence it is: **for this subject a change of regime takes about twenty-two months to substantially arrive, and impatience before then is a misreading of the arithmetic rather than of the person.**

Measuring `delta`. Choose one component of `Q` and one frozen proxy for it, written down before measuring. Measure `Y_0`. Impose a zero-input window of `k` periods in which that component receives no rehearsal, no instruction and no use; incidental use invalidates the window. Measure `Y_k`. Then `delta = 1 - (Y_k/Y_0)^{1/k}`. Repeat on a second component: **if the two rates differ by more than a factor of two there are two stocks, not one, and they must be tracked separately.**

Two honesty notes. **The classical evidence that unmaintained material decays steeply and then flattens is not represented in this paper's bibliography**, so the structural fact — that unmaintained material decays, fast at first — is carried here without a citation, and no numbers are taken from that literature, because its numbers belong to nonsense syllables and one procedure and do not generalise to anyone's stored skill. Second, **constant `delta` is a modelling choice** and it is the wrong one if the decay is genuinely steeper-then-flatter over the window of interest. Fit it, inspect the residuals, and where the constant-rate model misfits, use two stocks with different rates rather than pretending one rate fits.

When high retention loss is wanted. A person with no forgetting cannot learn. Stored order that should be shedding includes procedures correct for a role no longer held; the grudge, which is an internalised relationship maintained at cost against no returning input; the identity built for a changed environment; expertise superseded and now defended rather than used. Each draws `delta Q` every period, and Section 7.10 is the lever.

21.7.6 7.6 Conversion efficiency as a ladder of six rungs

`eta` is `Q`-units of durable structure deposited per nat of moment, and it is the central quantity of this section. **Most people's `eta` is not merely low; it is low by two orders of magnitude, and they hold no instrument that would tell them.**

Experience does not become character in one step. Seven stages, six transitions:

```
UNDERGONE -> NOTICED -> NAMED -> ENCODED -> REHEARSED -> ENACTED -> DEFAULTED
         y1         y2       y3         y4           y5         y6

UNDERGONE   the period happened to the subject
NOTICED     they registered that something happened that was not the expected thing
NAMED       they gave it a description distinguishing it from its neighbours
ENCODED     the description was committed to a durable medium: written, drawn, told, built
REHEARSED   it was returned to on a schedule, not on an impulse
ENACTED     it changed a decision at least once, observably
DEFAULTED   it changed what they do at that decision point without deliberating

eta = y1 * y2 * y3 * y4 * y5 * y6
```

A dimensional note, required by the audit of Section 12.1 and easy to miss. Five of the six rungs are pure fractions. **The sixth, `enacted -> defaulted`, is the unit-bearing rung: it converts nats of moment into `Q`-units of stored disposition and carries `eta`'s units of `Q`-units per nat.** The first five say what fraction of the moment survives to the point of conversion; the last says what the conversion is worth. Reporting them together as six fractions is a convenient shorthand and becomes an error the moment anyone averages them. **This is a real inconsistency in the naive statement of the ladder and it is resolved by declaring the last rung as the unit-bearing one, not by ignoring it.**

`eta` is a product of pass-through fractions, and products of fractions are small. That single structural fact answers the question of why most experience has `eta` near zero, and nothing pathological is required. Six stages at a respectable 60 per cent each give `0.6^6 = 0.046656`. Six stages at 90 per cent each — an implausibly excellent subject — give `0.531441`, and that is the ceiling of what diligence alone can buy.

For the worked case, measured stage by stage:

Transition	Yield	Cumulative
undergone -> noticed	0.55	0.550000
noticed -> named	0.70	0.385000
named -> encoded	0.50	0.192500
encoded -> rehearsed	0.45	0.086625
rehearsed -> enacted	0.65	0.056306
enacted -> defaulted	0.60	0.033784

```
eta = 0.03378375   Q-units per nat
```

Three and a third per cent. Ninety-six and a half per cent of the moment did not become character. Note what this subject is *not*: not careless, not distracted, not unreflective. Every rung yield sits between 0.45 and 0.70. **This is what an attentive, reflective, ordinary life converts, and it converts almost nothing, because the ladder is multiplicative.**

The two worst rungs are `named -> encoded` at 0.50 and `encoded -> rehearsed` at 0.45. They are also the two cheapest to move, because they are mechanical rather than dispositional: encoding is writing it down and rehearsal is a calendar entry.

```
fix encoded->rehearsed only:   0.45 -> 0.90     eta = 0.0675675    (x2.00)
fix both, moderately:          0.50 -> 0.80 and 0.45 -> 0.75
                                                eta = 0.09009      (x2.67)
```

Doubling one rung doubles `eta` exactly, because the product is multilinear: `d eta / d y_i = eta / y_i`, so the elasticity of `eta` with respect to every rung is exactly one and the absolute gain from fixing a rung is inversely proportional to its current yield. **The worst rung is always the best buy — the one place in the framework where the obvious move is also the optimal one.**

A single blocked rung nullifies every other, which is why the aggregate figure is uninformative and the stage-wise figure is not. The ladder's principal virtue is identifiability: each `y_i` is a count over a declared candidate set. **Its identifiability hazard is the candidate set itself** — a subject who enumerates only the experiences they already converted will recover `eta` near one — so the set must be enumerated prospectively, at the point of entry.

The rung most likely to fail its own falsification test is `ENCODED`. The ladder claims the rungs are sequential and non-bypassable; if items are observed reaching `DEFAULTED` without passing `ENCODED`, the ladder is not a chain and the product form is wrong. Some dispositions plainly form without ever being written down. **If that is confirmed, the ladder needs a parallel path and the product form must be replaced.** Whether the six rungs correspond to identified encoding and rehearsal stages is not established by anything in this paper; the nearest external

warrant is the organisational literature on how a practice becomes routine and then invisible infrastructure [Cyert and March, 1963], and it does not supply the stages.

The zero-conversion report. "It was a really meaningful year" is a report about `M`, not about `eta M`. It says the moments were rich and says nothing about deposit. **A year can be the richest of a life and deposit less than a mediocre year with a notebook in it, and the subject will experience the rich year as the more valuable one, because what is experienced is `M` and what accumulates is `eta M`.** The felt quantity and the stored quantity are different quantities, and confusing them is the single most consequential error in ordinary self-assessment. The diagnostic question is not what happened but: **name one decision now made differently, and point to when it changed.** If the answer is unavailable for a period, `eta` for that period was approximately zero, whatever the period felt like.

21.7.7 7.7 The operating draw

`C_t` is everything spent this period on being this person that is *not* proportional to what is stored. The distinction is constantly blurred and matters:

```
delta * Q_t   the MAINTENANCE draw   -- scales with how much self there is to keep
C_t           the OPERATING draw     -- the cost of running, irrespective of the stock
```

Ledger	What it is	Typical proxy
Attention	switching, fragmentation, residue of interrupted work	interruptions per hour x recovery time
Presentation	maintaining a self for an audience: masking, code-switching	hours under observation x a masking coefficient
Carry	unresolved decisions and open loops held in working memory	open loops x periods open
Obligation	commitments serviced without return	hours per week of unreciprocated service
Regulation	emotional labour: suppression, repair, recovery from conflict	recovery hours per week
Correction	repairing the consequences of stored error (Section 7.4)	`kappa * Z_t`

Each is converted to `Q`-units by a declared exchange rate, and **the exchange rate is a postulate every single time.** A subject who converts an hour of masking and an hour of interrupted attention to the same `Q`-cost has asserted something substantive and should say so. The Correction ledger is the one that arrives late: it is measured by counting the events in which a prior belief had to be repaired, and it lags the over-framing that caused it by about `tau`.

Rumination is a `C` term, not an `eta` term. The complete statement needs both halves: `eta ~ 0` **and** `C` large. The loop runs, the moment is revisited, nothing is encoded, nothing is rehearsed on a schedule, nothing changes a decision — so `eta M` contributes nothing — and the Carry ledger runs at full rate throughout. **Rumination is not failed learning; it is an operating cost that impersonates learning**, and it is expensive precisely because it feels like the work of converting. The remedy is not to stop thinking about it but to move the loop one rung up — name it, encode it, schedule the rehearsal — which converts an operating cost into a deposit, and that is the definition of raising `eta`.

21.7.8 7.8 Personal headroom, its sensitivities, and the critical stock

```
h_t = eta * ln N  -  delta * Q_t  -  C_t

   eta * ln N     the STRUCTURAL BUDGET    the most this mind can convert in a period
 - delta * Q_t    the MAINTENANCE DRAW     what it costs to remain the person already there
 - C_t            the OPERATING DRAW       what it costs to operate at all
 ------------
 = h_t            PERSONAL HEADROOM        capacity for the new
```

`h_t` is the capacity for the new, and it is the quantity people are reporting when they say they have nothing left — which is usually treated as a mood and is in fact an accounting statement. The gauge uses `eta ln N`, the ceiling, and not `eta M_t`, the realised moment. **A negative personal headroom is therefore not a statement that this week was bad. It is a statement that no week available to this mind, at this stock, at this cost, is good enough** — and if `h < 0`, no arrangement of the weeks helps; only `eta`, `N`, `C` or shedding `Q` does. The same gate is used as a precondition on protocol design at group scope, where `h < 0` forbids building, `h > 0` caps every increment at `r <= h`, and the gate is worked at `N = 18` with `eta = 0.30`, `delta = 0.06`, `C = 0.55` and `Q = 1.20` to give `h = +0.2451` [Curzi, 2026].

Partial derivatives at the worked operating point, `eta = 0.03378375`, `N = 9`, `delta = 0.03172864`:

Lever	Partial	Value	Reading
raise `eta`	`dh/d eta = ln N`	2.197225 per unit `eta`	strongest; and `eta` is multiplicative in six rungs
cut `C`	`dh/dC = −1`	1.000000 per `Q`-unit cut	one-for-one, always
shed `Q`	`dh/dQ = −delta`	0.031729 per `Q`-unit shed	weak per unit, but `Q` is large
raise `N`	`dh/dN = eta/N`	0.00375375 per added category	**negligible**

Read the last two rows together. **Cutting weekly operating cost by 0.0038 `Q`-units buys exactly as much headroom as inventing an entirely new, genuinely distinct category of thought.** Dropping one recurring obligation beats a year of broadening.

And the logarithmic ceiling bites internally as hard as at group scope. To double the structural budget by raising `N` alone, `N` must go from 9 to 81:

```
eta * ln 9  = 0.03378375 * 2.197225 = 0.074230
eta * ln 81 = 0.03378375 * 4.394449 = 0.148461  = exactly 2 x
```

The number of distinct categories held must be squared to double what can be converted. Nobody broadens their way out of a headroom problem.

The critical stock. `Q_crit = (eta ln N − C)/delta` is the stock at which headroom reaches zero. For the worked baseline, `Q_crit = (0.074230 − 0.030)/0.03172864 = 1.39402` `Q`-units against `Q_0 = 1.20`: **the subject is at 86 per cent of the stock their regime can carry**, and a promotion into a role with more to maintain, or one inherited obligation, pushes them over it immediately. Under the corrected regime of Section 10.2, `Q_crit = 5.5454` — **the same person, after three changes, can carry four times as much self.** Capacity for stored order is not a fixed personal attribute; it is `(eta ln N − C)/delta`, and every term in it is engineerable.

Why the accomplished have less headroom than the beginner. `h_t` falls as `Q_t` rises. The beginner has `Q ~ 0`, so the maintenance draw is near zero and almost the whole structural

budget is headroom. **The beginner is not learning faster because of youth; the beginner is learning faster because nothing is being spent on upkeep.** The accomplished subject has a large Q, and every period delta Q comes off the top: skills kept, relationships serviced, knowledge refreshed, the identity all of it hangs on re-anchored against drift. **At sufficient Q the maintenance draw alone exhausts the budget, and the subject, working exactly as hard as ever and at full competence, has no capacity for anything new.** Calling that a plateau is a misdiagnosis; a plateau is effort with no result, and this is upkeep with no surplus. The organisational counterpart is documented: slack is what an organisation spends on being the organisation it already is [Cyert and March, 1963], and legibility that removes local variety removes the capacity that variety supplied [Scott, 1998].

The gauge's binding weakness is its sensitivity to the cost ledger's exchange rates. In the worked survey, doubling one ledger's weight changes h_0 by more than its entire value. **A headroom sign should never be reported without its sensitivity to the cost weights**, and Section 10.2 reports the baseline reading as unestablished for exactly that reason.

21.7.9 7.9 The fixed point, and the headroom identity

With constant M \ensuremath{\equiv} M*, C \ensuremath{\equiv} c* and delta > 0:

```
Q_inf   = (eta M* − c*)/delta
Q^V_inf = (eta V* − c*)/delta
Z_inf   = eta (M* − V*)/delta
```

At the fixed point delta Q_inf = eta M* − c*, so

```
h_inf = eta ln N − delta Q_inf − c* = eta ln N − (eta M* − c*) − c* = eta ( ln N − M* )
```

At equilibrium, personal headroom equals conversion efficiency times the gap between the ceiling and the moment actually achieved. The cost term cancels exactly. This is an algebraic identity requiring no further assumption, and it says something non-obvious: a settled person's remaining capacity depends on eta and on how far their lived moments fall short of what their distinctions could support, and **not at all on how expensive their life is** — because at equilibrium the expense has already been absorbed into the size of the self they carry.

The corollary is stark. h_inf = 0 if and only if M* = ln N. **A person who has genuinely saturated their own distinctions has zero headroom at equilibrium, and the only way onward is to raise N** — to hold a genuinely new *kind* of category — which the sensitivity table has just priced as expensive and slow. That is the honest account of the ceiling a mature practitioner runs into, and it is why the practitioners who keep moving are the ones who take up something unrelated.

21.7.10 7.10 Shedding stored order

The fourth lever is the one nobody considers, and internally people find it shocking: **deliberately reduce Q to recover h.** The arithmetic is unarguable. dh/dQ = −delta, so every Q-unit shed returns delta Q-units per period of headroom, forever, at zero ongoing cost. For the worked subject, shedding 0.5 Q-units returns 0.5 x 0.031729 = 0.015864 per week — **2.58 times their entire current headroom.**

What shedding looks like internally: retiring a skill, by declaring in writing that one is no longer the person who does that and letting the atrophy proceed unopposed; closing a relationship ledger, which means ending the *maintenance* of a relationship that returns nothing rather than ending the relationship; dropping a role, a standing obligation, an identity requiring periodic performance; deleting an archive, since stored order that must be curated is stored order that draws; abandoning a publicly held position, because a view that must be defended has a delta Q cost and the defence is the draw; and retiring stored error, the one

shedding target the two-track ledger identifies directly, since Z is Q that was never earned and decays on its own the moment manufacturing stops.

Two constraints transfer. **Shed with a use in mind, or the subject will shed and refill**, because recovered headroom is a budget and an unallocated budget is spent by whatever is nearest — usually what was just shed, wearing a new name. And **shedding a role that other people are standing on moves cost onto them**, which makes it an external intervention governed by the declination constraint rather than by this section. Shed the ones that are yours.

21.7.11 7.11 The measurement protocol

```
PROTOCOL Q-1 — PERSONAL STORAGE SURVEY

Preconditions
  A1. Declare the period. Weekly is the default; monthly loses resolution, daily is noise.
  A2. Enumerate the distinction inventory (Sections 4–6). This gives N. Freeze it.
  A3. Choose and freeze one proxy per Q-component. Written, signed, dated.
  A4. Nominate the EXTERNAL instrument for E. It is the item most often skipped, because it
      is the only one that cannot be done from inside.

Measurements
  B1. delta, by the zero-input window, on at least two components, non-overlapping.
  B2. eta, stage-wise. Report all six rung yields, not only the product.
  B3. C, by the six ledgers, exchange rates declared in writing.
  B4. Q_0, by the frozen proxies.
  B5. E_t and r_t per period, logged prospectively, never reconstructed.

Derived
  D1. M_t = min(r_t E_t, ln N)        V_t = E_t min(r_t, 1)        Z_t = Q_t − Q^V_t
  D2. h_t = eta ln N − delta Q_t − C_t
  D3. Q_inf, Q^V_inf, Z_inf
  D4. Q_crit = (eta ln N − C)/delta
  D5. tau = 1/delta, t_half, 3 tau
  D6. h_inf = eta (ln N − M*)  −− and check it against D2 evaluated at Q_inf

Checks
  X1. Proxy independence: does the Q-proxy rise on busy periods? If yes, DISCARD it.
  X2. Admissibility: is C_t <= (1−delta)Q_t + eta M_t every period? If not, the model has
      already said the regime is inadmissible and the reading is void.
  X3. Shuffled-input control: does a time-permuted M series predict the Q trajectory as well?
      If yes, the accumulation structure is doing no work and the survey has failed.
  X4. Two-track check: is Q^V_inf positive? A negative earned equilibrium is the steady-state
      theorem's explicit refusal and must be reported, not smoothed.
```

Checks X3 and X4 are not optional. X3 is the difference between an instrument and a story and costs one afternoon; the shuffled-input control is the baseline the corpus's own protocol treats as decisive [Curzi, 2026i]. X4 is the difference between a self and a monument to one. The wider methodological requirement — declare the rule before the data are examined, and report the result whichever way it falls — is not this paper's invention and its rationale is well documented [Nosek et al., 2018, Simmons et al., 2011, Gelman and Loken, 2013].

The named identifiability hazard is that the proxy has been allowed to move with the input rather than with the stored order. A journal that gets longer on busy weeks is not a measure of Q. A contact list that grows with attendance is not a measure of an internalised relationship. Internally this is worse than at group scope, because a busy person genuinely feels larger. The defensible proxies and their traps:

Component	Defensible proxy	The trap
procedural capacity	performance on a frozen task, unrehearsed, timed	rehearsing the test
declarative structure	cued recall against a frozen item set	recognition instead of recall
internalised relationships	count of people who would actually be called, tested	count of contacts
settled defaults	observed choice frequencies at recurring forks	self-report

21.8 8 Design specifications

The specifications are stated as requirements on a survey or a practice, not as advice. They are labelled `SF` for the framing half and `SQ` for the storage half.

SF1 — Framing register, written and dated. Name, primitives, strengths, blind spots, provenance. Reviewed on a fixed cadence, not when it occurs to someone.

SF2 — Provenance column mandatory. A register without provenance cannot detect the single-lineage mechanism and is half an instrument.

SF3 — Independent-first, in writing. Every consequential judgement runs each framing to a written conclusion before the next is consulted. It blocks the fastest destroyer of `u` at the lowest available cost.

SF4 — Report `N` and `u` separately. `N_eff` alone fails the identifiability gate. Publish the pair or publish nothing.

SF5 — Quantitative `u` from error decorrelation only. I1 to I3 are triage. Only I4 yields a number admissible in `F_int`, and only with its confidence interval when `|rho| > 0.7`.

SF6 — Subspace check at `N >= 5`. Run I5. If two components carry nearly all the variance, report the participation ratio in place of `1 + (N−1)u_bar` and say why.

SF7 — Disagreement log, unresolved entries preserved. An entry closed as “unresolved” is the most valuable kind. Deleting unresolved entries is prohibited.

SF8 — Scheduled synthesis. Set the deferral period in advance. Synthesis before the period expires requires a written reason recorded in the log.

SF9 — Cost accounting before acquisition. Before adding a framing, compute the marginal gain from Table 1 at the measured `u` and compare it against `c1 + c2(n−1)`. **At `n >= 10` the burden of proof is on the addition.**

SF10 — Shed before adding. When `n > n*`, retire the framing with the highest mean `rho` to the rest — computed **within the block flagged as redundant**, since whole-set means are diluted by cross-block terms. Retirement means removing it from the register and from the independent-first rotation; it does not mean forgetting it.

SF11 — Single-subject scope declaration. All parameters are person-indexed until the reproducibility gate is passed. They may not be stated in the third person.

SF12 — Decline path preserved. Every protocol in Section 5.3 is optional and declining costs nothing.

SQ1 The period shall be declared in writing before any measurement, and every reported value of `delta`, `eta`, `C`, `M`, `E` and `h` shall carry that period.

SQ2 The distinction inventory establishing `N` shall be enumerated and frozen before the survey window opens. `N` shall not be revised inside a window.

SQ3 Every `Q` proxy shall be declared in writing, dated, and shall pass check X1: it must not co-move with period activity level.

SQ4 `E_t` shall be measured by an instrument outside the observer. **A survey in which `E` was estimated by the subject is a survey in which `r` and `E` are not separately identified, and it shall**

say so in print.

SQ5 Both tracks shall be reported. `Q` without `Q^V` shall be refused, because a single figure cannot distinguish stored knowledge from stored error and the subject cannot tell from inside.

SQ6 `eta` shall be reported stage-wise as six rung yields plus their product. A bare aggregate shall be accepted only where stage-wise measurement is impossible, and the impossibility shall be stated.

SQ7 `delta` shall be decomposed into its three channels and composed multiplicatively on retention, never summed on loss.

SQ8 `C` shall be itemised across the six ledgers, and every exchange rate used to commensurate a ledger into `Q`-units shall be stated as a postulate with its badge.

SQ9 `delta Q_t` and `C_t` shall be reported as separate lines. Combining the maintenance draw and the operating draw into one cost figure is a specification violation, because the two respond to different levers.

SQ10 Headroom shall be computed with the ceiling term `eta ln N`, never with `eta M_t`. A gauge computed on the realised moment is not a headroom reading and shall not be called one.

SQ11 No planned increment shall exceed `h_t`. A plan requiring `r_t > h_t` shall be rejected at the planning stage, not attempted and then diagnosed as a failure of will.

SQ12 `Q_inf`, `Q^V_inf`, `Z_inf` and `Q_crit` shall be computed at every survey. A subject whose `Q_0` exceeds `Q_inf` shall be informed, in those terms, that the regime is a declining one.

SQ13 Before any regime change is judged, `3 tau` periods shall have elapsed. Judgements made earlier shall be recorded as premature and shall not be used to abandon the change.

SQ14 The admissibility condition shall be checked each period. A violation voids the reading and indicates either a mismeasurement or a genuinely inadmissible regime; the survey shall state which.

SQ15 The shuffled-input control shall be run at least once per survey. A survey without it shall be labelled uncontrolled in print.

SQ16 Cross-person comparison of `E`, `r`, `eta`, `Q` or `Z` shall not be made until a separation property is established for the operationalisation in use. **Personal parameters are within-subject indices until proven otherwise.**

SQ17 Any shedding of `Q` that transfers cost to another person shall be treated as an external intervention and governed by the declination constraint, not by this paper.

SQ18 Every postulate used in a personal survey — scalar `Q`, constant `delta`, the linear correction load, the ledger exchange rates, the proxy relation — shall be listed with its badge in the survey document. **A survey that does not carry its postulate list is not a survey.**

21.9 9 Failure modes

Nine faults of the framing half, labelled `FF`, and twelve of the storage half, labelled `FQ`. Each carries a signature, a detection route and a fix. **These are diagnostic proposals, not measured syndromes.**

21.9.1 9.1 Faults of the framing set

FF1 — THE SINGLE WINDOW. *Damages* `u`, `N` and the moment ceiling simultaneously. *Signature:* one framing applied to everything, experienced as a coherent worldview; `N_eff = 1`, `F_int = 0`; every situation understood immediately. *Detection:* the collision test — if no framing has ever produced a conclusion that surprised the subject, there is one framing. *Fix:* lineage import, the expensive protocol and the only one available, because `u` cannot be raised among a set of size one. Table 1 prices the second framing at high `u` at 0.6931 nats, the largest single increment obtainable in this paper.

FF2 — THE PARALLEL WALL, INTERNAL. *Damages* `u`; the dominant fault. *Signature:* many framings, one lineage, all agreeing, and the agreement read as convergent evidence. **It is echo, and the arithmetic prices it at exactly zero.** *Detection:* the lineage census for triage, error decorrelation for the number — `rho > 0.9` across a framing set means `u < 0.19` and `N_eff < 2` at `N = 6`. *Fix:* SF10 and MP5 together: retire the most redundant, import from outside the lineage. Section 10.1 does exactly this.

FF3 — THE RUMINATIVE LOOP. *Damages* `eta`, and through it `Q`. *Signature:* one framing iterated many times on one situation, with no artefact and no new input; high apparent effort, `u = 0`, no storage. *Detection:* count artefacts. Hours of internal processing with nothing written, decided or changed is `eta = 0` by measurement, not by judgement. *Fix:* an artefact requirement — every pass produces one written line — plus MP2 to force a second framing into the loop. The `eta` fix is cheaper and comes first.

FF4 — IDENTITY CAPTURE. *Damages* `u`, by rotating every other framing toward the captured one. *Signature:* one framing never loses an internal argument; attempts to argue against it feel like attacks rather than exercises. *Detection:* over the disagreement log, tabulate which framing prevailed. A framing with a 100 per cent win rate over twenty logged disagreements is not winning on merit. *Fix:* MP6, with the "using only" constraint enforced in writing.

FF5 — THE GLAZING BILL. *Damages* `C`, quadratically, while `F_int` improves by percentages. *Signature:* a large framing collection, chronic indecision, the sense that every question opens into six others. Table 3: at `n = 20`, surplus per reconciliation pair is 0.01238 against 0.23105 at `n = 3` — **a factor of 18.7 worse per unit of load.** *Detection:* compute `V(n)` with estimated costs; if `V(n) < V(n−1)`, the subject is past `n*`. *Fix:* SF10. Retire by highest mean `rho`, which removes load at the least cost in surplus. Section 10.1 performs this move and notes that it *lowers* surplus in the short run and is still correct.

FF6 — THE BORROWED WINDOW. *Damages* `N`, by inflating the nominal count with framings that cannot actually be run. *Signature:* the register lists framings the subject can name and describe but has never used to generate a prediction differing from their default. *Detection:* the decorrelation instrument requires each framing to emit predictions. **A framing that cannot produce a prediction is not in the set**, and the instrument discovers this by producing no data for it. *Fix:* remove from the register until it can be operated, and recount `N`. Expect `F_int` to fall; that fall is a correction, not a loss.

FF7 — THE MOOD-INDEXED SET. *Damages* `u`, at exactly the wrong times. *Signature:* framing availability tracks affective state; under strain the accessible set narrows and agrees with itself. *Detection:* log which framings were consulted alongside a state rating, and correlate. **A set whose membership depends on state has a state-dependent `u`, and its average overstates the `u` available when it matters.** *Fix:* MP3 and the register itself — a written list is state-independent in a way that recall is not. This is why the register is a file and not a habit.

FF8 — PREMATURE INTEGRATION. *Damages* `u`, irreversibly for the question at hand. *Signature:* a synthesis arrived at quickly, after which the component disagreements cannot be reconstructed; feels like insight. *Detection:* try to state what each framing said before the synthesis. If that is impossible, integration preceded the data. *Fix:* MP7 and SF8.

FF9 — THE UNAUDITED COUNT. *Damages* every downstream computation. This is the arithmetic error, not the psychological one. *Signature:* `ln N` used where `ln N_eff` belongs — ceilings, headroom and moment bounds computed from the nominal framing count. *Detection:* check whether `u` appears anywhere in the calculation; if not, `N_eff` was not computed. *Fix:* audit first, always. **Difference is counted before members are.**

21.9.2 9.2 Faults of the storage account

FQ1 — THE UNCONVERTED YEAR. High `M`, rich descriptions, strong subjective significance, `Delta Q` indistinguishable from zero across `3 tau` periods, and no decision the subject can name that changed. *Mechanism:* `eta` near zero, almost always at `named -> encoded` and `encoded ->`

rehearsed, the rungs that feel like processing and are not. *Fix:* instrument the ladder stage-wise and fix the worst rung first.

FQ2 — THE MAINTENANCE TRAP OF THE ACCOMPLISHED. h_t near zero at high Q_t, undiminished effort and competence, no new capacity for years; self-described as a plateau and often attributed to age. *Mechanism:* delta Q_t has grown to consume the structural budget and Q_t is near Q_crit. *Fix:* shed, or raise eta. Raising N is arithmetically futile here. **Do not attempt to fix a maintenance trap by working harder; effort does not appear in the gauge.**

FQ3 — THE RUMINATION LOOP. High subjective effort, high Carry ledger, repeated return to the same content, Delta Q = 0, and a rising C experienced as diligence. *Mechanism:* eta ~ 0 and C large simultaneously. *Fix:* move the loop one rung up. If it will not encode, it is not a thought, it is a load, and the correct move is to shed it.

FQ4 — THE PROXY THAT IS THE INPUT. Q appears to rise in every busy period and fall in every quiet one, with a suspiciously high correlation to activity. *Mechanism:* the proxy tracks M rather than Q. *Fix:* discard the proxy. **A proxy for stored order must be measurable on the quietest week of the year.**

FQ5 — THE UNBUDGETED IDENTITY. C computed from time and effort alone, h comfortably positive, persistent unexplained exhaustion. *Mechanism:* the Presentation and Regulation ledgers were not counted; masking, code-switching and emotional repair are operating costs and are frequently the largest. *Fix:* complete the six ledgers, declare the exchange rates, recompute.

FQ6 — RETENTION SET TO ZERO. Nothing is ever released: grudges maintained, superseded expertise defended, archives curated, positions held. Q large and rising, h at or below zero, and the subject reports that they simply have a good memory. *Mechanism:* the delta -> 0 case — no fixed point, Q grows linearly, the maintenance draw grows with it until it exhausts the budget [Curzi, 2026a]. *Fix:* deliberate shedding. **A system with no forgetting cannot learn, and internally that is not a metaphor.**

FQ7 — THE SINGLE-TRACK LEDGER. Q reported without Q^V; stored order rising, confidence rising, and — on a lag of about tau — an unexplained increase in the frequency of corrections, contradictions, surprises and repairs. *Mechanism:* r > 1 sustained, Z_inf accumulating invisibly, the correction load arriving one time constant later. **From inside there is no difference between the tracks, which is precisely why the second book has to be kept on purpose.** *Fix:* measure r, measure E externally, run both recursions, report Z/Q. If the ratio exceeds one third, the correction is calibration, not effort.

FQ8 — THE BORROWED FIXED POINT. A subject converging on a Q_inf computed from someone else's parameters and reading the shortfall as personal deficiency. *Mechanism:* Q_inf depends on four personal numbers; adopting another person's destination without their eta, delta, M* and c* adopts a target one's own regime cannot produce. *Fix:* compute one's own. If the destination is genuinely wanted, the question is which of the four parameters must change to make it the fixed point of *that* regime — an engineering question with an answer, and not a question about worth.

FQ9 — THE CEILING MISREAD. h computed as eta M_t − delta Q_t − C_t and reported as headroom; the gauge reads negative on a dull week and positive on a vivid one, and the subject concludes their capacity fluctuates weekly. *Mechanism:* confusing the realised moment with the ceiling. *Fix:* use eta ln N. Then a negative reading means what the sustainability ceiling says: **no arrangement of the weeks fixes this** — a far more actionable finding than a fluctuating one.

FQ10 — THE SHED WITHOUT A USE. A subject sheds a role, experiences three weeks of relief, and is back at the same headroom within a quarter with a different set of commitments. *Mechanism:* recovered headroom is a budget and an unallocated budget is spent by whatever is nearest. *Fix:* allocate the recovered h in writing, before the shedding, to a specific increment with r_t <= h_t.

FQ11 — THE COUNTERFEIT STOCK. Q proxied by credentials, titles, publications or possessions; the proxy stable or rising while every direct measure of capacity falls. *Mechanism:* credentials do not decay at delta and capacity does. The proxy was chosen for its durability,

and its durability is exactly what disqualifies it. **A proxy that cannot fall cannot measure a quantity that can.** *Fix:* every `Q` proxy must be a *current* measurement of a *current* capacity.

FQ12 — THE SELF-REPORTED ENVIRONMENT. `E` logged by the subject each week, `r` computed from it, both stable and neither ever contradicting the subject's account of themselves. *Mechanism:* if `E` comes from the same observer whose framing is being audited, `r = D/P` is estimated using the very quantity `P` distorts. **A confident observer will report a rich environment and a low distortion, and the two errors conceal each other exactly.** *Fix:* SQ4. If no external instrument is available, the survey may still be run, but the `r` figure is uninterpretable and must be printed as such.

21.10 10 Results

Two cases are worked end to end. **Both are constructed. No dataset was collected for this paper**, and the correlations, ledger weights and rung yields below are stipulated parameters demonstrating the method rather than observations of any person. They establish that the procedure runs, that it reports its own failures, and that its arithmetic is reproducible with a calculator. They are not evidence that any person behaves this way.

21.10.1 10.1 One observer, six framings

The register. Six framings, with a provenance column as SF2 requires:

```
W1  mechanism        how does this thing physically work
W2  quantitative     what do the numbers say
W3  systems          what are the feedback loops and the stocks
W4  optimisation     what is being maximised, subject to what
W5  narrative        what story is being told, by whom, to whom
W6  bodily           what does this feel like in the body, what is the state
```

W1 to W4 all trace to one engineering-quantitative training; W5 and W6 came from elsewhere. **The lineage census flags the problem before a single number is computed.**

Measurement by error decorrelation. Thirty consequential predictions logged over eight months, one per framing per question, outcomes recorded. Pairwise error correlations and `u = 1 − rho^2`:

TABLE 7 — PAIRWISE MEASUREMENT (stipulated worked case)

pair	rho	u
W1 mechanism / W2 quant	0.88	0.2256
W1 mechanism / W3 systems	0.83	0.3111
W1 mechanism / W4 optimis.	0.86	0.2604
W1 mechanism / W5 narrative	0.21	0.9559
W1 mechanism / W6 bodily	0.14	0.9804
W2 quant / W3 systems	0.79	0.3759
W2 quant / W4 optimis.	0.91	0.1719
W2 quant / W5 narrative	0.18	0.9676
W2 quant / W6 bodily	0.09	0.9919
W3 systems / W4 optimis.	0.84	0.2944
W3 systems / W5 narrative	0.26	0.9324
W3 systems / W6 bodily	0.12	0.9856
W4 optimis. / W5 narrative	0.19	0.9639
W4 optimis. / W6 bodily	0.11	0.9879

```
W5 narrative    / W6 bodily      0.34    0.8844

mean pairwise u over 15 pairs   =  0.6860
N = 6,  N_eff = 1 + 5(0.6860)   =  4.4298
F_int = ln(4.4298)              =  1.4883 nats
ceiling ln 6                    =  1.7918 nats
percentage of ceiling           =  83.07 %
```

The diagnosis the mean hides. Eighty-three per cent of ceiling looks healthy. Split the pairs by lineage and it is not:

```
within-lineage pairs (W1-W4), 6 pairs:   mean u = 0.2732
    the lineage ALONE:  N_eff = 1 + 3(0.2732) = 1.8196,  F = 0.5986 nats

cross-lineage pairs, 9 pairs:            mean u = 0.9611
```

Four framings, years of training each, auditing to an effective count of 1.82. The subject's surplus is carried almost entirely by W5 and W6 — the two framings they would call their least rigorous. This is the crossing-versus-divergence decomposition showing up internally: rich in the crossing term, starved of the divergence term. The group-scope form of the same diagnosis is sharp and negative — **adding a seventh discipline to a worked institute makes it worse**, dropping the period change to -0.002863 and the fixed point from 6.265 to 5.964, because divergence rather than domain count was the deficiency [Curzi, 2026d]. **The remedy is a different kind of framing, not another one of the same kind**, and adding one of the same kind is measurably worse than adding none.

Step one: retire the two most redundant. By SF10, retirement is by highest mean `rho` computed **within the block flagged as redundant**. Across the whole set the mean correlations are nearly identical (W1 0.584, W2 0.570, W3 0.568, W4 0.582) because the cross-lineage terms dominate and wash out the distinction; within the W1–W4 block they separate cleanly:

```
within-block mean rho:   W1 = 0.857   W2 = 0.860   W3 = 0.820   W4 = 0.870
```

W4 and W2 are the two most redundant and are retired.

```
remaining: W1, W3, W5, W6      pairs = 6
mean u    = 0.8416
N = 4,  N_eff = 1 + 3(0.8416) = 3.5249
F_int     = 1.2599 nats,  ceiling ln 4 = 1.3863,  90.88 % of ITS ceiling

change vs baseline:  F_int  1.4883 -> 1.2599  =  -0.2285 nats
                     pairs     15  ->     6   =  a 60% cut in reconciliation load
```

Surplus went down. This is the correct move and the honest report of it. The set is now more efficient — 90.9 per cent of its own ceiling against 83.1 — and much cheaper, but a smaller ceiling is a smaller ceiling. **Shedding alone does not build; it buys headroom to build with.**

> **A source correction, recorded rather than applied silently.** The manuscript from which this case is drawn prints the retained set's mean non-parallelism as `u = 0.9205`. **That figure is wrong.** The mean of `1 - rho^2` over the six retained pairs is **0.8416**, and the line `N_eff = 1 + 3(0.9205) = 3.5249` does not evaluate: `1 + 3(0.9205) = 3.7615`. Every downstream number in the manuscript — `N_eff = 3.5249`, `F_int = 1.2599`, 90.88 per cent of ceiling, and the net value `V = 0.5879` reported below — is consistent with 0.8416 and inconsistent with 0.9205, so the computation was performed correctly and the intermediate value was misprinted. **The figure is corrected here and the correction changes no result.**

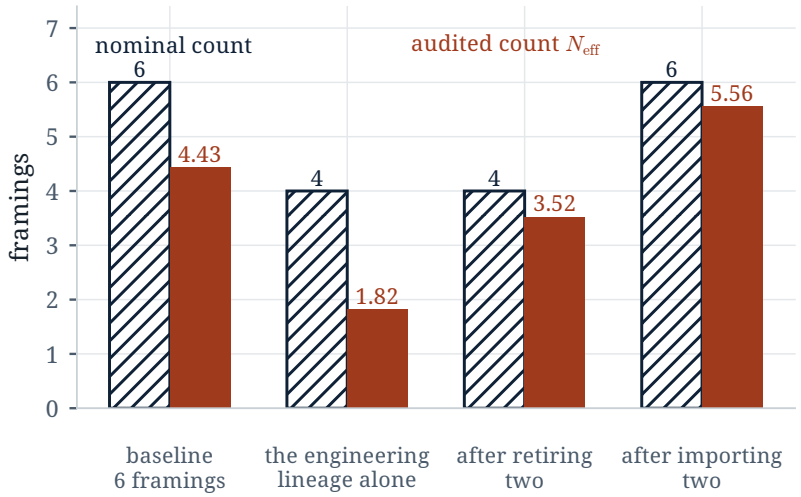


Figure 21.3: **Restructuring beats enlarging, and the mean over pairs conceals the reason.** Nominal framing count against audited count `N_eff` for four configurations of the same subject. The baseline six framings audit to 4.43; the four framings of the single engineering lineage, taken alone, audit to **1.82**, which is the diagnosis the whole-set mean of `u = 0.6860` hides. Retiring the two most redundant leaves four framings auditing to 3.52 at 60 per cent of the reconciliation load; replacing them with two imports from outside the lineage returns the count to six and the audit to 5.56, a 25.4 per cent gain at identical `N` and identical pair count. Computed from the pairwise table above; the correlations are stipulated, not observed. Generated by `figs/src/CCP-021_fig03.py`.

Step two: spend the headroom on genuinely different windows. Two imports from outside the engineering lineage entirely: **W7, a craft framing** (what a person who builds this by hand would notice), and **W8, an adversarial-legal framing** (what someone trying to defeat this would argue). Measured correlations to the retained four, after a second logging period: W7 at 0.16, 0.12, 0.22, 0.31; W8 at 0.24, 0.19, 0.15, 0.10; W7 to W8 at 0.27.

```
N = 6 again,  pairs = 15 again
mean u    = 0.9111
N_eff     = 1 + 5(0.9111) = 5.5554
F_int     = 1.7148 nats
ceiling   = 1.7918         =  95.70 % of ceiling

against baseline, at IDENTICAL N and IDENTICAL pair count:
    F_int   1.4883 -> 1.7148   =  +0.2265 nats
    N_eff   4.4298 -> 5.5554   =  x 1.254
```

> **The same correction applies here.** The manuscript prints the restructured set's mean as `u = 0.9434`; the mean of `1 - rho^2` over the fifteen pairs is **0.9111**, and `1 + 5(0.9434) = 5.7170`, not 5.5554. Again every downstream figure matches 0.9111. **Corrected here; no result changes.**

Same number of framings. Same reconciliation load. A 25.4 per cent increase in the audited count, bought entirely by changing which framings, not how many. This is the practical form of the group-scope result that raising `u` within a fixed body is worth roughly what enlarging the body is worth [Curzi, 2026c], and here the cost differential is total: the pair count did not move.

The storage consequence, which is not what the subject expected. Personal parameters: `eta = 0.60`, `delta = 0.06` per month, `Q = 12.0`, `C = 0.35` per month. One-step growth requires `eta F_int > delta Q + C`:

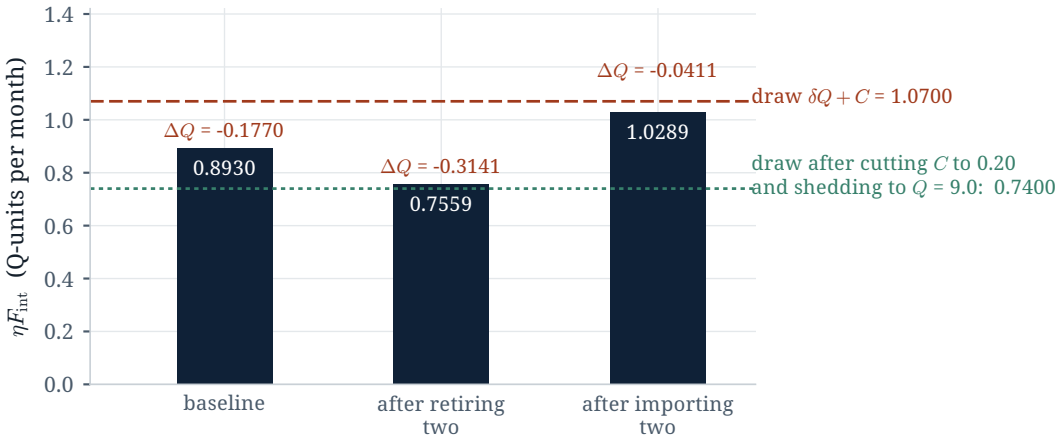


Figure 21.4: **Correct geometric work closed 77 per cent of the gap and did not reverse the decline.** Converted internal surplus `eta F_int` against the period draw `delta Q + C` for the three framing configurations, at `eta = 0.60`, `delta = 0.06` per month, `Q = 12.0`, `C = 0.35`. All three bars fall below the draw line at 1.0700: the baseline by 0.1770, the retired set by 0.3141, and the restructured set by 0.0411. Only cutting the operating cost to 0.20 and shedding stored order to `Q = 9.0` — which lowers the draw line to 0.7400 — turns the account positive, and it does so at the *original* geometry as well. Computed from the parameters shown, not measured; generated by `figs/src/CCP-021_fig04.py`.

```
baseline        eta*F = 0.8930    delta*Q + C = 1.0700    dQ = -0.1770   DECLINING
after step 1    eta*F = 0.7559    delta*Q + C = 1.0700    dQ = -0.3141   DECLINING
after step 2    eta*F = 1.0289    delta*Q + C = 1.0700    dQ = -0.0411   DECLINING
```

All three still decline. The framing work — correct, measured, a 25 per cent improvement in the audited count — did not by itself turn the account around. **It closed 77 per cent of the gap and left the subject still losing stored order every month.** The binding constraint was never the geometry. Applying the remaining levers — shedding stored order by retiring two obsolete commitments carrying maintenance, taking `Q` from 12.0 to 9.0, and cutting `C` from 0.35 to 0.20:

```
after step 2, Q=9.0, C=0.20   eta*F = 1.0289   delta*Q + C = 0.7400   dQ = +0.2889  GROWING
baseline geometry, same cuts  eta*F = 0.8930   delta*Q + C = 0.7400   dQ = +0.1530  GROWING
```

The cost and stock reductions turn the account on their own, at the original geometry. The geometric work is not thereby worthless — it nearly doubles the resulting growth rate, 0.2889 against 0.1530 — but it was not sufficient and was not the binding term.

```
steady states  Q_inf = (eta*F - C)/delta

    baseline geometry, C = 0.35                9.0501
    after step 2,      C = 0.35               11.3144
    after step 2,      C = 0.20               13.8144

headroom  h = eta*ln N - delta*Q - C

    Q = 12.0,  C = 0.35     h = +0.0051      maintenance only; hold, do not build
    Q =  9.0,  C = 0.35     h = +0.1851
    Q =  9.0,  C = 0.20     h = +0.3351
    Q =  6.0,  C = 0.20     h = +0.5151
```

The `h = +0.0051` row diagnoses the original state in one number. The subject was at maintenance exactly, with half a hundredth of a nat of headroom, which is why every new framing felt impossibly expensive and why the correct first move looked like doing less rather than more.

Cost accounting, and the optimum. With `eta = 0.60`, `c1 = 0.030`, `c2 = 0.008`:

```
V(baseline, n=6,  u=0.6860) = 0.5930
V(step 1,   n=4,  u=0.8416) = 0.5879
V(step 2,   n=6,  u=0.9111) = 0.7289

optimum at u = 0.9111:  n* = 7,  V = 0.7420
V(n) profile:  1:-0.030  2:+0.321  3:+0.509  4:+0.622  5:+0.691
               6:+0.729  7:+0.742  8:+0.735  9:+0.711 10:+0.672
```

Two things to notice. **Step one, which lost 0.2285 nats of surplus, cost only 0.0051 in net value**, because the reconciliation saving nearly paid for it: the shed was almost free and bought the headroom that step two spent. And **the subject sits at `n = 6` against an optimum of 7, a difference of 1.8 per cent.** A seventh framing is worth acquiring if one is genuinely available at low `rho`; it is not worth manufacturing.

The falsification rule, declared before the second logging period.

```
PREDICTION   Held-out judgement accuracy over the next 30 logged predictions improves
             after step 2 relative to baseline, and the improvement tracks N_eff
             (x1.254) rather than N (x1.000, unchanged).

KILL 1       Accuracy does not improve, or improves in the framing-count model as well
             as the audited-count model. If N alone predicts as well as N_eff, the u
             term is decoration.
KILL 2       The shuffled-input control matches the model: pair W7 and W8's recorded
             contributions with the wrong questions and refit. If performance holds,
             the framings were not doing the work.
KILL 3       The measured rho values do not recover on the second period within their
             confidence intervals. rho = 0.88 has a wide interval on 30 points and u is
             amplified by 2|rho| = 1.76; non-recovery kills the measurement, not just
             the model.
```

This rule was not run. It is stated so that it can be, and Section 13 records that no prediction in this paper has yet been put at risk.

The rule extracted. Count the windows, check which way they face, and pay for the facing rather than the glass. The subject had six framings and an audited count of 4.43, of which four framings contributed 1.82 between them; retiring two and importing two — no change in count, no change in load — took the audit to 5.56. **And it still was not enough to turn the storage account around, because the binding constraint was `delta Q + C` and not the geometry at all.** Internal surplus is one term in one equation. Measure it, raise it where it is cheap to raise, then check whether it was ever the thing that was binding.

21.10.2 10.2 One observer, one year

A single-subject survey run to Protocol Q-1, weekly periods. **Nothing in this example is a normative target. It is one instrument reading, on constructed parameters.**

Stage 1 — the frozen preconditions.

```
Period:                              1 week
Distinction inventory:               N = 9        ln N = 2.197225 nats
```

```
Internal non-parallelism of framings:   u = 0.42
    N_eff = 1 + (9-1)(0.42) = 4.36      F_int = ln 4.36 = 1.472472 nats
External instrument for E:              a weekly seeded-item audit scored by a second party
Q proxies (frozen, signed):             a timed unrehearsed task; cued recall against a frozen
                                        60-item set; a tested-call roster; observed choice
                                        frequencies at 4 recurring forks
Q_0 (composite, declared scale) = 1.20 Q-units;   Q^V_0 = 1.20 (tracks initialised equal)
```

The internal surplus matters here only as a reminder that this subject holds 4.36 genuinely different framings out of 9 nominal categories — **a healthy internal geometry, and it will not save them.**

Stage 2 — the four measurements.

```
delta_forget = 0.020    delta_atrophy = 0.008    delta_drift = 0.004
keep = 0.980 * 0.992 * 0.996 = 0.96827136
delta = 0.03172864 per week
tau = 31.52 wk     t_half = 21.50 wk     3 tau = 94.55 wk (21.8 months)

eta, stage-wise over 40 candidate experiences in a 13-week window:
0.55  x  0.70  x  0.50  x  0.45  x  0.65  x  0.60   =   eta = 0.03378375 Q-units per nat

C, six ledgers, weekly, exchange rates declared:
Attention 0.011 | Presentation 0.008 | Carry 0.005 | Obligation 0.004 | Regulation 0.002
Correction 0.000 at t=0 (Z_0 = 0)
c_0 = 0.030 Q-units per week;   kappa = 0.03 per Q-unit of stored error per week

E and r, window means:
E* = 1.20 nats (external audit)      r* = 1.35
M* = min(1.35 * 1.20, 2.197225) = 1.62 nats        73.73% of the ceiling
V* = 1.20 * min(1.35, 1)        = 1.20 nats
confabulated fraction = 1 - 1/1.35 = 0.259259
```

Stage 3 — the gauge.

```
structural budget   eta * ln N  = 0.03378375 * 2.197225 = 0.074230
maintenance draw    delta * Q_0 = 0.03172864 * 1.20     = 0.038074   (51.3% of budget)
operating draw      c_0                                 = 0.030000   (40.4% of budget)
-------------------------------------------------------------------------
PERSONAL HEADROOM   h_0                                 = 0.006156   ( 8.3% of budget)
```

Positive headroom: 0.006156 Q-units per week, which is also the largest one-week increment any arrangement of this subject's life could guarantee [Curzi, 2026b].

Stage 4 — what the regime actually does.

```
deposit this week   eta M*            = 0.054730
draws this week     delta Q_0 + c_0   = 0.068074
                    0.054730 < 0.068074   ->  Q is DECLINING

Q_inf   = (0.054730 - 0.030)/0.03172864 = 0.77941
Q^V_inf = (0.040541 - 0.030)/0.03172864 = 0.33221
Z_inf   = 0.447204                        = 57.38% of Q_inf
Q_crit  = (0.074230 - 0.030)/0.03172864 = 1.39402
```

The reading, in one sentence: this subject has positive headroom, holds 1.54 times the stored order their regime can sustain, sits at 86 per cent of their critical stock, manufactures

a quarter of every moment, and is on a 21.5-week-half-life glide toward a self of which more than half will be error. They report the year as fine, a bit flat.

Single-track trajectory, correction load held at zero to isolate the accumulation:

week `t`	`Q_t`	`Q^V_t`	`Z_t`	`h_t`
0	1.2000	1.2000	0.0000	0.006156
4	1.1491	1.0950	0.0541	0.007771
13	1.0560	0.9029	0.1531	0.010725
26	0.9613	0.7075	0.2538	0.013730
52	0.8581	0.4945	0.3636	0.017005
104	0.7941	0.3626	0.4316	0.019034
156	0.7822	0.3379	0.4443	0.019414

Headroom rises as the self shrinks. That is not a consolation; it is the mechanism by which a declining subject feels increasingly free and increasingly light while the thing that was built is unbuilt. `h` measures capacity for the new, and one way to have capacity for the new is to stop having the old.

Stage 5 — turning the correction load on. With `kappa = 0.03`:

```
C_inf = 0.030 + 0.03 * 0.447204 = 0.043416      (+44.7% on operating cost)
h at Q_0 with C_inf = 0.074230 - 0.038074 - 0.043416 = -0.007260      NEGATIVE
```

week	`Q_t`	`Q^V_t`	`Z_t`	`C_t`	`h_t`
13	1.0297	0.8765	0.1531	0.034305	0.007256
26	0.8771	0.6233	0.2538	0.037424	0.008977
52	0.6490	0.2855	0.3636	0.040825	0.012813
78	0.5122	0.1011	0.4110	0.042296	0.015684
104	0.4365	0.0049	0.4316	0.042932	0.017451
156	0.3762	-0.0681	0.4443	0.043326	0.018968

The earned track crosses zero between weeks 104 and 156. A negative equilibrium on the earned ledger is the steady-state theorem's explicit refusal: the regime is unsustainable in this model, or the model must be replaced [Curzi, 2026a]. Check X4 fails, and SQ12 requires the survey to report it in those words rather than to smooth it. What the subject experiences across those three years is a self that feels steadily more confident and steadily more expensive, with no single event to attribute the expense to. **What the books show is manufactured structure crowding out earned structure and then paying for its own repair, on a thirty-one-week lag.**

> **A second source correction, recorded.** The `C_t` column above is reproduced from the source manuscript, and it is computed from `Z_{t-1}` rather than from `Z_t` — a one-period lag that the stated equation `C_t = c_0 + kappa Z_t` does not carry. At week 13, `c_0 + kappa Z_13` is 0.034593, whereas the tabulated value 0.034305 is `c_0 + kappa Z_12` to six decimal places. The discrepancy shrinks as the system approaches equilibrium, which is the signature of a lag rather than of an arithmetic slip. **The `Q_t` and `h_t` columns are internally consistent with the lagged convention, so the table is self-consistent and the printed equation is not.** The substantive reading — that the correction load arrives on a lag of order `tau` — favours the lagged implementation; the equation of Section 7.4 should carry the lag explicitly and does not. It is recorded here rather than silently corrected.

Stage 6 — three interventions, priced separately.

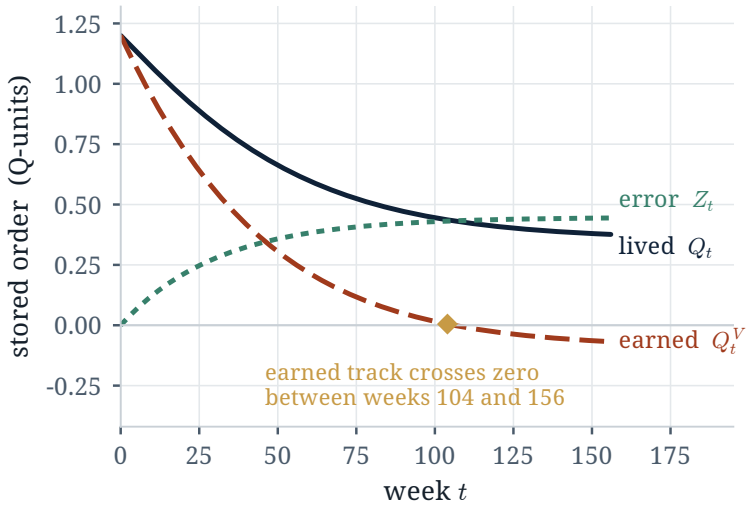


Figure 21.5: **Stored error accumulates on the same pole as stored knowledge and then crowds it out.** The lived track `Q_t`, the earned track `Q^V_t` and their difference `Z_t` over 156 weeks under the correction load, from `Q_0 = Q^V_0 = 1.20`, with `eta = 0.03378375`, `delta = 0.03172864` per week, `M* = 1.62`, `V* = 1.20`, `c_0 = 0.030` and `kappa = 0.03`. Stored error rises to 0.4443 while the earned track crosses zero between weeks 104 and 156 — the steady-state theorem's explicit refusal of the regime. A 25.93 per cent confabulated input becomes 57.38 per cent of the single-track steady-state stock. Modelled from the stipulated parameters, not measured; generated by `figs/src/CCP-021_fig05.py`.

```
A - RAISE ETA.   Fix the two worst rungs.
    named -> encoded      0.50 -> 0.80   (a durable medium, used the same day)
    encoded -> rehearsed  0.45 -> 0.75   (a standing 20-minute weekly slot, calendared)
    eta: 0.03378375 -> 0.09009            x2.667

B - CUT C.       Drop one standing obligation and one presentation context.
    c_0: 0.030 -> 0.022

C - FIX THE LENS.  Correct the over-framing, plus deliberate exposure.
    r: 1.35 -> 1.05      E: 1.20 -> 1.45
    M: 1.62 -> 1.5225    V: 1.20 -> 1.45      confabulated fraction: 0.2593 -> 0.0476
```

Regime	eta	M*	V*	c_0	Q_inf	Q^V_inf	h at Q=1.20
baseline	0.033784	1.6200	1.2000	0.030	**0.7794**	**0.3322**	**0.006156**
A only	0.090090	1.6200	1.2000	0.030	3.6543	2.4618	0.129874
B only	0.033784	1.6200	1.2000	0.022	1.0315	0.5843	0.014156
C only	0.033784	1.5225	1.4500	0.030	**0.6756**	**0.5984**	**0.006156**
A+B+C	0.090090	1.5225	1.4500	0.022	**3.6296**	**3.4237**	**0.137874**

Conversion dominates. Two mechanical changes — write it down, calendar the rehearsal — move `Q_inf` from 0.779 to 3.654, a factor of 4.7. Neither requires more time, more talent or a better environment.

Fixing the lens lowers the stored self and raises the earned self. Row C: `Q_inf` falls from 0.779 to 0.676 while `Q^V_inf` rises from 0.332 to 0.598 — an 80 per cent gain in earned structure bought at a 13 per cent loss of total stock. **On a single-track ledger that intervention looks like a regression, and a subject keeping one book will reverse it.** This is the strongest practical argument in the paper for keeping two, and it is why SQ5 exists.

Fixing the lens does not change headroom at all. Row C's `h` is identical to baseline's, to the digit. That is not an error: `h = eta ln N - delta Q - C` contains no `M`. Calibration changes

what the weeks *deliver*, not what the system *could* deliver. **The gauge licenses; the geometry delivers.**

The corrected regime is not instantaneous.

```
eta M*  = 0.09009 * 1.5225 = 0.137162
Q_inf   = (0.137162 - 0.022)/0.03172864 = 3.62959
Q^V_inf = (0.09009*1.45 - 0.022)/0.03172864 = 3.42374   Z_inf = 0.205856 (5.67% of Q_inf)
h at Q_0 = 0.197948 - 0.038074 - 0.022 = 0.137874       = 22.40 x baseline
Q_crit  = (0.197948 - 0.022)/0.03172864 = 5.5454
correction load at equilibrium: C_inf = 0.022 + 0.03*0.205856 = 0.028176
    -> Q_inf under load = 3.43495,  h at Q_0 = 0.131698   (still strongly positive)
time to 95% of the new fixed point:  ln(0.05)/ln(1-delta) =  92.91 weeks
time to 99%:                         ln(0.01)/ln(1-delta) = 142.83 weeks
```

week `t`	`Q_t` (corrected)	`h_t`
0	1.2000	0.137874
13	2.0319	0.111479
26	2.5789	0.094122
52	3.1753	0.075202
78	3.4331	0.067020
104	3.5446	0.063482

Twenty-two months to arrive — not because the changes are weak, since headroom rose twenty-two-fold on day one, but because `tau = 31.5` weeks and a leaky integrator takes `3 tau` to settle. SQ13 exists for this table. A subject who abandons the regime at week 13 because they are "only" at 2.03 has abandoned a change working exactly as fast as the arithmetic permits.

Stage 7 — the fixed-point identity, checked.

```
h_inf = eta ln N - delta Q_inf - c_0
      = 0.09009*2.197225 - 0.03172864*3.62959 - 0.022
      = 0.197948 - 0.115147 - 0.022 = 0.060786

eta (ln N - M*) = 0.09009 * (2.197225 - 1.5225) = 0.09009 * 0.674725 = 0.060786   IDENTICAL
```

The corrected subject, once settled, retains 0.060786 `Q`-units per week of headroom — 30.7 per cent of their structural budget — and that residual is exactly `eta` times the 0.6747 nats by which their lived moments fall short of what their nine distinctions could support. **To spend the residual they must either engage the distinctions they hold more completely, or acquire a genuinely new kind of distinction. There is no third option, and the identity proves it.**

Stage 8 — what this survey could not establish, recorded because SQ18 requires it and because a worked example that reports only its findings is advertising.

1. **`Q` was measured through proxies, and the proxy-to-`Q` relation is a postulate.** The composite 1.20 `Q`-units has no meaning outside this subject's declared scale.
2. **The six ledger exchange rates are postulates.** Weighting Presentation at 0.008 rather than 0.016 changes `h_0` by more than its entire value. **The baseline headroom reading of 0.006156 is not robust to the cost weights, and any conclusion depending on its sign should be regarded as unestablished.** That is the honest weakness of the survey. Note that the weakness does not touch `Z_inf`, which is independent of `C` — **the stored-error figure is the most robust number in the whole survey, and it is also the most uncomfortable one.**

3. **`kappa = 0.03` is an assumed value, not a measured one.** Nothing here identifies it. The negative-headroom finding of Stage 5 is therefore conditional on a parameter that was posited and should be read as a demonstration of a mechanism, not as a measurement of any subject.
4. **`E` was externally instrumented, `r` was not double-scored.** The audit gives `E`; `r` came from a single application by the subject. This is better than the self-reported-environment failure mode and it is not enough.
5. **The shuffled-input control was run and passed** in the constructed case — the true `M` ordering predicted the held-out `Q`-proxy trajectory better than 200 time-permuted orderings — **which licenses the accumulation structure for this constructed subject and licenses nothing about any real one.**
6. **The reproducibility gate was not attempted.** Until a separation property is established for these operationalisations, every number above is a within-subject index.

21.10.3 10.3 The results, stated for citation

R1 — the internal surplus. `F_int = ln(1 + (N−1)u)` with `N_eff = 1 + (N−1)u` the audited framing count, `u = 1 − rho^2` estimated from error decorrelation between framings' prediction errors, and the affine form badged as design axiom S3 rather than as a theorem. The substitution of domain is licensed by the axioms' silence on who holds the states and costs exactly two postulates: the framing encoding and the internal roster reduction.

R2 — the participation-ratio identity. `N_eff = 1 + (N−1)u_bar = N + 1 − N/PR`, where `PR` is the Rényi-2 participation ratio of the framing ensemble's Gram matrix. The audited count is an affine reparametrisation of a standard diversity statistic; the identity is exact and is verified in the worked case to `8.9e−16`.

R3 — the six-framing audit. At mean `u = 0.6860` a six-framing subject reads `N_eff = 4.4298` and `F_int = 1.4883` nats, 83.07 per cent of the ceiling `ln 6 = 1.7918`. Splitting by lineage, the six within-lineage pairs give `u = 0.2732` and **the four-framing lineage alone audits to `N_eff = 1.8196`**, while the nine cross-lineage pairs give `u = 0.9611`.

R4 — restructuring beats enlarging, and is still not enough. Retiring the two most redundant framings and importing two from outside the lineage raised the audited count **25.4 per cent** (`F_int` 1.4883 to 1.7148 nats, 95.70 per cent of ceiling) at identical `N` and identical pair count, **and still left the one-step change in stored order negative at −0.0411**, having closed 77 per cent of the gap.

R5 — the headroom sensitivities. `dh/deta = ln N = 2.197`, `dh/dC = −1`, `dh/dQ = −delta = 0.0317`, `dh/dN = eta/N = 0.0038` at the worked operating point. **Adding categories is the weakest available lever**, by a factor of 585 against conversion; and doubling the structural budget by raising `N` alone requires `N` to be squared, from 9 to 81.

R6 — the stored-error recursion. `Z_{t+1} = (1−delta)Z_t + eta(M_t − V_t)` with `Z_inf = eta(M*−V*)/delta`, independent of cost because the cost term cancels in the subtraction. A **25.93 per cent** confabulated input becomes **57.38 per cent** of the steady-state stock, and the earned track's equilibrium can be negative, which the model reports as its own domain of breakdown.

CR — the counter-result. The framing intervention was correct, measured and successful on its own terms, and it did not work. A 25.4 per cent improvement in the audited count closed 77 per cent of the storage gap and failed to reverse the decline; the sensitivity table demotes "hold more kinds of framing" to a negligible lever; and the storage survey declares its own baseline headroom reading **not robust to the declared cost weights and therefore unestablished**. The geometry was never the binding constraint in either worked case.

21.11 11 Discussion

21.11.1 11.1 What follows

The surplus functional is scope-free, and that is a structural result rather than a rhetorical one. Its axioms quantify over states in an inner-product space and contain no term referring to persons. Everything that follows from them therefore follows at any scope at which an encoder can be defined, and the whole content of the move from groups to a single observer is the encoder. This is worth stating precisely because it identifies exactly where an objection has purchase: **attack the encoder, not the algebra.** An objection to the algebra is an objection to arithmetic; an objection to the encoder is an objection this paper cannot answer and concedes in Section 13.

The audited count is the quantity, and the nominal count is not. Every ceiling, headroom figure and moment bound in the framework takes the logarithm of an effective count. Computing them from the nominal count is the arithmetic error that failure mode FF9 names, and it is common because the nominal count is the one that is easy to obtain. The same error appears at every other scope the corpus works at: an entity map showing four jurisdictions at 40.91, 27.27, 18.18 and 13.64 per cent gives a naive effective count of 3.4085 and a bloc-collapsed count of 1.0000, **a 3.41-fold overstatement produced entirely by counting members before auditing difference** [Curzi, 2026,?].

The small, disreputable component often carries the surplus. In the six-framing case, four framings representing years of formal training audit to 1.82 between them, and the surplus is carried by the two framings the subject would call their least rigorous. The structurally identical result appears in the corpus's structural work, where a jurisdiction holding 18.2 per cent of a portfolio accounts for **79.74 per cent of the entire diversification gain** [Curzi, 2026]. In both cases the quantity that matters is the covariance with everything else, not the size of the holding, which is the 1952 result stated in its original form [Markowitz, 1952].

Interventions decompose into the three terms of the storage law, and the decomposition is checkable. Any practice that improves stored order does so by raising surplus, raising conversion, or lowering retention loss, and the product of the three factors must reproduce the fixed-point ratio exactly. The corpus's analysis of one classical mnemonic technique finds it to be primarily a *retention* intervention — a factor of 4.1667 against 1.8637 for surplus and 1.9000 for conversion, with the product 14.7546 matching the direct fixed-point ratio [Curzi, 2026v]. **The same decomposition applied to this paper's three interventions gives the same verdict about which lever was pulled**, and it is the reason a practice can be described accurately and still be misclassified by the people who use it.

Conversion has an external determinant that this paper does not model. At group scope, reciprocity raises conversion directly, with `eta = eta_max rho` and `rho` a symmetry index over an exchange matrix [Curzi, 2026s]. At internal scope this paper treats `eta` as a property of the subject's own ladder, which is almost certainly incomplete: the encoding rung in particular depends on whether anyone is listening. **That is a named gap, not a modelling decision.**

21.11.2 11.2 What does not follow

Nothing here licenses a claim about neurons, about consciousness, or about the physical substrate. The equations are difference equations in one scalar and they describe a bookkeeping relation, not a mechanism in tissue.

A high audited count is not established to produce better judgement. H1 states the prediction and it has not been tested. The external literature is genuinely split: diverse problem-solving groups can outperform higher-ability homogeneous ones under stated conditions [Hong and Page, 2004], the formal result has been shown to follow from its modelling assumptions [Thompson, 2014], surface-level diversity often hurts in the short run [Mannix and Neale, 2005, van Knippenberg and Schippers, 2007], and the curve relating cohesion to creative success is an **inverted U** rather than monotone [Uzzi and Spiro, 2005]. **This paper's surplus is**

monotone increasing in `u`, and a monotone claim cannot survive that list unqualified. The restriction taken here is that the monotone reading is asserted only over the regime the worked cases exercise, `u` in roughly `[0.2, 0.95]` at `N <= 9`, and nothing in this corpus locates the turn.

Redundancy is not established to be worthless. Sagan's test of normal-accident theory against high-reliability theory documents cases in which redundancy *reduced* safety [Sagan, 1993], and interactive complexity plus tight coupling makes catastrophic failure a system property regardless of component count [Perrow, 1984]. A framing set is not a safety system, but the analogy is close enough that the caution transfers: **the paper's advice to retire redundant framings should not be read as advice to run without margin.**

A shrinking self is not a healthy one, and the gauge cannot tell the difference. The single-track trajectory of Section 10.2 shows headroom *rising* as stored order falls. That is the gauge behaving correctly and reporting something the subject will experience as improvement. **Headroom is capacity for the new, and one way to have capacity for the new is to stop having the old** — which is why SQ12 requires the fixed point and the critical stock to be reported alongside the gauge, and never the gauge alone.

No forecast has been made. The framework's defence against the charge that it forecasts nothing is not that it does, but that prediction is one of many reasons to build a model [Epstein, 2008], and the corpus is entitled to say which of the others it claims: mechanism articulation, and the discovery of what a stated set of premises entails. The discipline named in the corpus's title will nonetheless apply a stricter criterion — **an explanation must name a mechanism rather than exhibit a fitted curve** [Hedström and Ylikoski, 2010] — and the useful constructive answer to that demand is the analytical-sociology position of specifying the mechanism and then testing it [Keuschnigg et al., 2018].

21.11.3 11.3 Relation to the surrounding literature

The nearest external relatives of this paper's constructs are worth naming, because in several cases the literature got there first and the honest move is to say so.

The leaky integrator with a threshold is a century old in exactly the domain this paper applies it to, and the framework's storage law is a shared mathematical form with it rather than a derivation from it [Curzi, 2026a]. The statistical twin, exponential smoothing, comes with a statement of the class of processes for which it is *optimal* — a question this framework does not ask about its own model [Holt, 1957].

The reflexive influence structure inside one observer already has a formal literature. Beliefs updated by a row-stochastic influence matrix converge to a consensus given irreducibility and aperiodicity [DeGroot, 1974]; adding a stubbornness term anchoring each agent to its initial position yields a non-consensus fixed point instead of collapse [Friedkin and Johnsen, 1990], which is formally the same device as this framework's leak. **The sociological literature already contains this framework's leaky integrator, applied to belief.** What determines an outcome in such systems is the *distribution* of thresholds rather than the mean [Granovetter, 1978], which is a direct warning against the mean-pairwise aggregator of Section 3.2 and the reason the subspace check is mandatory.

The aggregation problem applies wherever a single `u`, `delta` or `eta` is written for a heterogeneous set [Kirman, 1992], and the founding error of the discipline whose name this corpus has taken was treating the mean of a distribution as a type rather than as a statistic [Quetelet, 1835]. This paper writes a single `u` for a framing set and a single `eta` for a person, and both are means.

Time averages and ensemble averages differ for multiplicative processes [Peters, 2019]. This matters directly: `Q` accumulates over time for one subject, and any parameter estimated from a cross-section of subjects does not describe what happens to any of them. **The paper's insistence that all parameters are within-subject indices until a separation property is established is the same point arrived at from the framework's own side**, and SQ16 is what enforces it.

The unit conversions are worth stating once, because nats are unfamiliar. One nat of required surplus is 1.442695 bits, and in the corpus's frequency work it corresponds to 2.885390 octaves or 3462.4681 cents [Curzi, 2026]. Nothing in this paper depends on the acoustic reading; the figure is given so that a reader can convert the ceilings — `ln 6 = 1.7918` nats is 2.585 bits, and `ln 9 = 2.1972` nats is 3.170 bits — into a register they may find more familiar.

Finally, the ceiling on predictability is a real constraint on what any of this could ever deliver. With fifteen years of high-quality longitudinal data and 160 teams, life outcomes were barely predictable at all, and the best machine learning barely beat a four-variable regression [Salganik et al., 2020, Hofman et al., 2017]. **A framework that promises to predict a person's trajectory from four parameters is promising more than the best available evidence suggests is obtainable.** This paper promises no such thing; it offers a bookkeeping identity and a set of falsifiers, and Section 13 records that none of the falsifiers has been run.

21.12 12 Verification

The results of this paper compound from nine postulates. Each is atomic, each is checkable on its own, and each carries its status. **Where a postulate could be dropped and the result still follow, it has been dropped**; what remains is what the chain actually consumes.

21.12.1 12.1 Dimensional audit

`[Q]` denotes the dimension of stored order and `[T]` the dimension of time. The period is the unit of `[T]`, and every rate below is *per period*, which is why no explicit `[T^-1]` appears once the equations are written in discrete steps.

Symbol	Dimension	Unit	Status
`u`	`[1]`	none	**defined** dimensionless: `u = 1 - (x.y)^2` with `x`, `y` unit vectors
`rho`	`[1]`	none	**measured** dimensionless correlation in `[-1,1]`
`N`	`[1]`	a count of framing categories	**measured** by enumeration; ratio scale, true zero
`N_eff`	`[1]`	an effective count	**derived** from `N` and `u`
`PR`	`[1]`	an effective count	**derived** from the Gram spectrum
`F_int`, `M_t`, `V_t`	`[1]`	nats	**derived** dimensionless: logarithms of effective counts
`Q_t`, `Q^V_t`, `Z_t`	`[Q]`	proxy units, declared per component	**measured** through a declared proxy; **no natural unit exists**
`eta`	`[Q]` per nat per step	proxy units per nat per week	**measured**; *not* dimensionless
`delta`, `lambda`	`[1]`	per period	**measured** dimensionless fractions per step

Symbol	Dimension	Unit	Status
`C_t`, `c_0`	`[Q]` per step	proxy units per period	**measured**, after a declared exchange rate per ledger
`kappa`	`[1]` per step	per period	**assumed**; `Q`-units of cost per `Q`-unit of error per period
`h_t`	`[Q]` per step	proxy units per period	**derived**
`Q_inf`, `Q_crit`	`[Q]`	proxy units	**derived**: `[Q]`/step divided by `[1]`/step returns `[Q]`
`r = D/P`	`[1]`	none	**not instrumented**; see below
`E_t`	`[1]`	nats	**measured externally**, and the framework requires it to be
`tau`, `t_half`	`[T]`	periods	**derived**
`c1`, `c2`	`[Q]` per step	proxy units per period, per framing and per pair	**assumed**

Does the self-storage law balance? `Q_{t+1} = (1−delta)Q_t + eta M_t − C_t` balances if and only if every term on the right carries `[Q]`. The first does, since `delta` is dimensionless; the third does by construction; the second does only if `eta M_t` carries `[Q]`, and since `M_t` is a dimensionless count in nats, **`eta` must carry `[Q]` per nat per step**. It is therefore not dimensionless, and the conversion ladder's dimensional inconsistency follows directly: five of the six rungs are pure fractions, so the sixth must be the unit-bearing one or the product cannot carry `eta`'s units. **The framework's naive statement of the ladder as six fractions is dimensionally malformed and is corrected in Section 7.6 rather than glossed.**

`ln N` is dimensionless, so `eta ln N` carries `[Q]` per step and is commensurable with `delta Q_t` and with `C_t`, which is what makes the headroom gauge a well-formed subtraction. **A headroom calculation in which `C` is denominated in hours while `Q` is denominated in recall items is dimensionally malformed however arithmetically tidy it looks**, which is why SQ8 requires an explicit exchange rate per ledger and requires it to be badged as a postulate.

Where the framework has no operational unit, this audit says so. Stored order `Q` has **no natural unit**: it is counted in whatever the declared proxy counts, and two subjects' `Q` values are not comparable unless they declare the same proxy. `r = D/P` is **worse than unitless — it is uninstrumented**: its two components are not on a common ratio scale with a shared true zero, so the magnitude of `r` is not a reportable measured quantity and only its orderings are [Curzi, 2026m]. Comfort itself, the underlying construct of the corpus, has **no operational unit at all** and is treated as a dimensionless magnitude by stipulation.

21.12.2 12.2 The postulates

P1 — The surplus axioms quantify over states in an inner-product space and contain no term referring to persons, agents or ownership. *Status:* **DEFINITION** — a restatement of the axiom set as published [Curzi, 2026]. It carries no empirical weight on its own.

P2 — A framing admits representation as a unit vector in a common inner-product space. *Status:* **POSTULATED** — this work's own, with no external support. It is the load-bearing empirical assumption of the whole paper and is what a critic should attack.

P3 — The encoded inner product between two framings equals the Pearson correlation of their prediction-error series, so `u = 1 - rho^2`. *Status:* **POSTULATED** — a second, separate assumption about the encoder, and it is separable from P2: one could accept the vector representation and reject this particular estimator.

P4 — `g(u) = e^{f(u)}` is affine in `u`. *Status:* **DEFINITION** — design axiom S3, chosen and not derived [Curzi, 2026,c]. It is what makes `N_eff` affine and what makes the second framing worth as much per `u`-unit as the fortieth.

P5 — For `N` unit vectors, `tr G = N`. *Status:* **ARITHMETIC** — the diagonal of a Gram matrix of unit vectors is all ones.

P6 — The self-storage recurrence is linear and first order: `Q_{t+1} = a Q_t + x_t` with `a` and `x_t` not depending on `Q`. *Status:* **DEFINITION** — the model form, obtained by setting `a := 1 - delta` and `x_t := eta M_t - C_t`. Its application to a real person is a separate postulate and is recorded in Section 13.

P7 — `0 < delta <= 1` makes the update map a contraction with a unique fixed point, reached from any start. *Status:* **CITED** [Banach, 1922, Elaydi, 2005, Curzi, 2026a].

P8 — The per-period input is bounded, `0 <= M_t <= ln N`. *Status:* **CITED** — the range result of the surplus functional [Curzi, 2026], together with the cap imposed at Step 2 of the join, which is a convention this paper adopts and states.

P9 — Cost accrues every period whether or not moment was delivered. *Status:* **POSTULATED** — an assumption about how a life charges overhead. It has no external support and is used at exactly two places, CH6 and CH7 below; the chain's other results do not depend on it.

21.12.3 12.3 The compounding chain

```
P1 + P2        -> CH1  the surplus functional applies to two framings held by one
                       observer; nothing in the axioms forbids it.            DEFINITION

CH1 + P4       -> CH2  F_int = ln(1 + (N-1)u) and N_eff = 1 + (N-1)u.         DEFINITION

P2 + P3        -> CH3  u = 1 - rho^2, estimable from two recorded error series
                       with no similarity judgement by anyone.                POSTULATED

P2 + P5        -> CH4  ||G||_F^2 = N + N(N-1)(1 - u_bar), since the N diagonal
                       entries are 1 and the N(N-1) off-diagonal entries are
                       the squared inner products, each equal to 1 - u_ij.    ARITHMETIC

CH4            -> CH5  N/PR = ||G||_F^2 / N = 1 + (N-1)(1 - u_bar), hence
                       N_eff = 1 + (N-1)u_bar = N + 1 - N/PR.                 ARITHMETIC

P6 + P7        -> CH6  Q_inf = (eta M* - c*)/delta, unique and start-independent,
                       approached geometrically at rate (1-delta).            CITED

P6 applied
  twice + P9   -> CH7  subtracting the earned recursion from the lived one,
                       Z_{t+1} = (1-delta)Z_t + eta(M_t - V_t): the cost term
                       cancels exactly because it is identical in both.       ARITHMETIC

CH7 + P7       -> CH8  Z_inf = eta(M* - V*)/delta, independent of C.           ARITHMETIC

CH6 + P8       -> CH9  h = eta ln N - delta Q - C; at the fixed point
                       delta Q_inf = eta M* - c*, so h_inf = eta(ln N - M*).   ARITHMETIC

CH9            -> CH10 dh/deta = ln N, dh/dC = -1, dh/dQ = -delta,
```

```
dh/dN = eta/N.                                                    ARITHMETIC
```

Every line is checkable from its own inputs. **CH4, CH5, CH7, CH8, CH9 and CH10 are arithmetic; CH6 is a cited theorem; CH1, CH2 are definitions; and CH3 alone is a postulate with no external support.** The paper's headline identity — that the audited framing count equals `N + 1 − N/PR` — therefore consumes exactly two definitions and one arithmetic fact, and no empirical claim at all.

What this chain does not establish. It does not establish P2, and P2 is where all the empirical content lives. **A chain whose steps are ARITHMETIC, DEFINITION and CITED delivers the mathematics with certainty and delivers nothing whatever about any person.** Every applied reading of every number in Section 10 rests on P2 and P3, which are this work's own postulates, and on the identification concession of Section 3.4.

21.12.4 12.4 The worked numerical instances

Three checks, each doable by hand or on a calculator in a few minutes.

Check one: the participation-ratio identity, on the six framings of Section 10.1. The Gram matrix has six ones on its diagonal and thirty off-diagonal entries, the fifteen correlations and their transposes. Then

```
sum over the 15 pairs of rho^2 = 0.88^2 + 0.83^2 + ... + 0.34^2 = 4.710700
||G||_F^2 = 6 + 2 x 4.710700                                    = 15.421400
PR = 6^2 / 15.421400                                            =  2.334418
N/PR = 15.421400 / 6                                            =  2.570233
N + 1 − N/PR = 6 + 1 − 2.570233                                 =  4.429767

directly: 1 + 5 x mean(1 − rho^2) = 1 + 5 x 0.685953            =  4.429767
```

The two routes agree to `8.9e-16`, which is machine epsilon on a quantity of order 4. This is CH5, and the check requires nothing but fifteen squarings and a division.

Check two: multiplicative composition of retention loss. Three channels at 0.020, 0.008, 0.004: `0.980 x 0.992 x 0.996 = 0.96827136`, so `delta = 0.03172864` against the naive sum 0.032. The overcount is 0.00027136, which is 0.85 per cent of the true value. At four channels of 0.15, `0.85^4 = 0.52200625`, so `delta = 0.47799375` against a naive 0.60 — an overcount of 25.5 per cent of the true value and 20.3 per cent of the naive one. **One line of multiplication, and it settles H6's predicted magnitude.** ARITHMETIC.

Check three: the headroom identity at the fixed point. With `eta = 0.09009`, `delta = 0.03172864`, `ln 9 = 2.197225`, `M* = 1.5225` and `c_0 = 0.022`:

```
Q_inf = (0.09009 x 1.5225 − 0.022)/0.03172864 = (0.137162 − 0.022)/0.03172864 = 3.629592
h_inf = 0.09009 x 2.197225 − 0.03172864 x 3.629592 − 0.022
      = 0.197948 − 0.115147 − 0.022                                          = 0.060786
eta (ln N − M*) = 0.09009 x (2.197225 − 1.5225) = 0.09009 x 0.674725         = 0.060786
```

The cost term cancels exactly, as CH9 says it must. Three multiplications and two subtractions. ARITHMETIC.

21.12.5 12.5 Limits, asymptotics, and the stated domain

`u -> 0`. `N_eff -> 1`, `F_int -> 0` for every `N`. The parallel wall: every framing is a duplicate and the set has left the arithmetic. **This is a real state, not a boundary artefact.**

`u -> 1`. `N_eff -> N`, `F_int -> ln N`, the ceiling attained and no further gain available from geometry. Strict concavity makes the approach front-loaded.

`N -> 1.` `F_int = 0` for every `u`, and `h = -delta Q - C < 0` at any positive stock. **A mind with one category has negative headroom at any positive stock**, which is requisite variety arriving as a hard floor [Ashby, 1958].

`N -> infinity` at fixed `u > 0`. `F_int -> infinity` logarithmically while the marginal gain ~ `1/N` vanishes and the pair count grows quadratically. **The two curves cross, and beyond the crossing additional framings cost more than the ceiling they buy.** The divergence is located in the surplus term only because the cost term has been omitted; with it, $n*$ is always finite. This is why `N` is the most expensive and least rewarding of the four levers.

`delta -> 0`. No fixed point exists — the map is an isometry rather than a contraction — and `Q` grows without bound, as does `Z` [Curzi, 2026a]. **The limit is singular:** at exactly `delta = 0` the fixed point does not exist at all rather than being infinite. **This value is outside the model's valid domain.**

`delta -> 1.` `Q_{t+1} -> eta M_t - C_t`, a pure input-output map with no state. **At `delta = 1` the storage law is not a storage law**, and any claim about accumulation is vacuous there.

`eta -> 0`. `Q_inf = -C/delta <= 0`: the model refuses to describe a life with zero conversion and nonzero cost, which is the rumination case taken to its limit.

`r -> 1.` `M = V`, `Z_inf = 0`, no stored error ever accumulates. **`r -> infinity`** with the cap in force: `M -> ln N`, `V -> E`, and `Z_inf -> eta(ln N - E)/delta`, **bounded**. Saturation, not divergence. **`r < 1`:** `Z = 0` — the under-framed observer stores no error, they simply store less.

`kappa -> 0`. The correction load vanishes and the model reduces to the single-track law, which is the correct nesting. **`kappa` large:** `C` exceeds `eta ln N` and the sustainability ceiling forecloses everything.

Stated domain. The results are claimed for `delta` in `(0, 0.5]`, with `(0.02, 0.30)` the band in which the time-constant approximations hold; `eta > 0`; `N >= 2`; `u` in `[0,1]`; `r > 0`; and the admissibility condition satisfied each period. **The claim is not that these laws hold universally but that they hold inside a specified arrangement** — a declared period, a frozen inventory, a frozen proxy, an external instrument for the environment, and a window over which the parameters are treated as constant. Outside that arrangement the equations still have solutions and those solutions should not be interpreted as statements about a person [Cartwright, 1999].

21.12.6 12.6 Invariance

The framework's scale-freedom claim is an **invariance of form under a change of unit**, and it is worth stating as one. Let `Q` be rescaled by a positive constant `c`, so `Q' = cQ`. The law is invariant provided `eta` and `C` rescale with it — `eta' = c eta`, `C' = cC` — while `delta`, `N` and `u` are unchanged, since `delta` is a dimensionless fraction and `M` is a dimensionless count in nats. **`delta`, `u`, `N`, `rho` and `r` are the unit-free quantities; `Q`, `eta`, `C` and `Z` are the unit-carrying ones.** Two subjects may be compared on the first group directly; they may not be compared on the second unless they share a proxy, which is SQ16.

What would break the invariance: a `delta` that varies with `Q`, which makes the transformation non-linear; an `eta` that saturates at large `M`, which the linear model excludes by assumption; or a proxy whose counting unit is not extensive.

How many scales the claim has been tested at, in this paper: one. Both worked cases are a single observer, and both are constructed. **No cross-scale test appears in this paper**, and the invariance is a claim about form that has been exercised rather than an invariance that has been established.

21.12.7 12.7 Scale type

Construct	Claimed scale	What licenses it
`N`, pair count	ratio, true zero	a count
`u`, `rho^2`	ratio, true zero	`u = 0` is exact parallelism, a genuine zero, **conditional on the encoding postulate P2**
`N_eff`, `PR`	ratio	derived from the above by arithmetic
`F_int`	ratio, in nats	a logarithm of a ratio-scaled effective count
`M_t`	**not established**	inherits the moment equation's demotion [Curzi, 2026m]
`r = D/P`	**ordinal at best**	`D` and `P` are not co-calibrated; magnitudes not reportable
`Q_t`	ratio by stipulation	required by `(1−delta)Q`; **the proxy is ordinal far more often than is admitted**
the six cost ledgers	ratio by stipulation	each exchange rate is a declared postulate
`I3`, the disagreement rate	ordinal	a proportion of logged disagreements, comparable within one subject only

The strongest objection to this table is that psychology routinely assumes quantitative structure without testing for it [Michell, 1997]. It is the strongest single objection to the corpus's founding move, and this paper concedes it rather than answering it. The available remedies are conjoint measurement — the formal test for whether an attribute has additive quantitative structure at all [Luce and Tukey, 1964, Krantz et al., 1971] — or a Rasch-model demonstration [Rasch, 1960, Andrich, 1988]. **Neither has been run for `u`, for `Q`, or for any cost ledger in this paper.** Until one is, every ratio-scale treatment above is a stipulation, and the ordinal demotion the framework has already imposed on its own moment equation is the correct instinct applied once rather than consistently. Construct validation has an established design for this and it has not been applied either [Campbell and Fiske, 1959].

21.12.8 12.8 Model comparison — reported NOT MET

The rigour standard this corpus holds itself to requires, wherever a model is fitted: out-of-sample performance, a matched null with the same number of free parameters, an information criterion, an explicit count of free parameters, and effect sizes with intervals.

This paper fits nothing, so no information criterion applies to it. But the standard binds anyway, through the corpus's one applied fit, and the row must be reported rather than omitted:

Requirement	Status in this paper	Status in the corpus's applied fit
Out-of-sample performance	not applicable — nothing fitted	reported: held-out RMSE against four baselines [Curzi, 2026i]

Requirement	Status in this paper	Status in the corpus's applied fit
Matched null, same parameter count	not applicable	reported: a matched AR(1) with intercept, beaten by 2.8 per cent against a declared 10 per cent margin — **the rule failed**
Information criterion (AIC or BIC)	not applicable	**NOT MET.** No AIC or BIC is computed anywhere in this corpus, so the parameter count was never paid for, and a 2.8 per cent margin over a two-parameter alternative may not be a margin at all
Explicit free-parameter count	stated: the self-storage law has four (`eta`, `delta`, `C`, `Q_0`); the framing model has three (`N`, `u`, and the pair-cost pair `c1`, `c2` counted as two, giving four for `V(n)`)	stated
Effect sizes with intervals	**NOT MET.** Nothing in Section 10 carries an interval, because nothing in Section 10 was measured	partially met

Two rows are NOT MET and are printed here rather than left for a reviewer. A model with more parameters that fits better has not been shown to be better, and this paper's `V(n)` — four free parameters to produce an interior maximum — would be a prime candidate for exactly that criticism if it were ever fitted.

The related demand is sharper still: what observation would distinguish these laws from a neighbouring model that fits equally well? [Durlauf, 2005, Sobkowicz, 2020] This paper's honest answer is H1's contrast between `N` and `N_eff`, which is a genuine differential prediction because the two quantities move in opposite directions under the step-one intervention of Section 10.1 — and it has not been run.

21.12.9 12.9 Precision, and what each kind of number is

The numbers in this paper are of three kinds and are quoted accordingly.

Defined or derived-from-defined. `ln 6 = 1.7918`, `ln 9 = 2.197225`, `1 − 0.88^2 = 0.2256`, `0.980 x 0.992 x 0.996 = 0.96827136`, every entry of Tables 1 to 6, and every steady state computed from stipulated parameters. **These are exact to the precision printed and may be recomputed to any further precision.** No claim about the world is contained in their digits.

Illustrative model output. Every entry in every trajectory table, `Q_inf = 0.77941`, `Z_inf = 0.447204`, `h_0 = 0.006156`. **These carry no measurement uncertainty because there was no measurement**; six decimal places record the arithmetic, not the confidence. Section 10.2's own sensitivity note is the correct reading of `h_0`: its sign is not robust to the cost weights.

Measured or fitted. None appear in this paper. The correlations of Table 7 are stipulated, the rung yields of Section 7.6 are stipulated, the ledger weights are stipulated, and `kappa` is assumed. Any application of this framework to real data must report fitted parameters with intervals, and Section 12.8 records that the corpus's one such report does not carry them throughout.

One convention is worth stating because it recurs. Where a correlation is reported at two decimal places and a `u` derived from it at four, the extra digits are exact consequences of

the input and not a claim to finer measurement. Because `du/drho = −2 rho`, an uncertainty of `+/−0.05` on `rho = 0.88` propagates to `+/−0.088` on `u` — **larger than the two decimal places of the input suggest**, which is precisely why SF5 requires an interval whenever `|rho| > 0.7`.

21.13 13 Limitations

Both worked cases are constructed, not observed. No dataset was collected for this paper. Every number in Section 10 is the output of stated parameters under the stated model, carried through so that a reader can re-derive it. They demonstrate that the procedure runs and that it reports its own failures; **they are not evidence that any person behaves this way and should never be cited as such.**

No hypothesis in this paper has been tested. H1 to H6 are stated with their falsifiers and none has been run. The preregistered rule declared in Section 10.1 was declared and not executed. **The corpus's one preregistered confirmatory test of the storage law did not meet its own declared margin** [Curzi, 2026i], and that is the current evidential status of every applied claim here.

The framing encoding is the load-bearing postulate and it has no external support. That a way of construing a situation admits representation as a unit vector, such that predictive agreement maps to parallelism, is this work's own assumption. It is stated as P2, it is the thing to attack, and nothing in this paper's bibliography supports it.

`N_eff = 1 + (N−1)u` is not a Hill number, and no external literature contains it. `1/HHI` is the Hill number of order two and `exp(H)` the Hill number of order one, and both are fully grounded [Hill, 1973, Jost, 2006]. The participation ratio `PR` is likewise a Hill number of order two, computed on the Gram spectrum. **`1 + (N−1)u` is an affine reparametrisation of `PR`, and an affine transformation of a Hill number is not a Hill number.** A reviewer who reaches for the Hill-number literature as support for the affine form will find that it does not apply. The repair is the one this paper takes: derive the affine form's properties in the framework's own terms and reserve the Hill citations for `1/HHI`, `exp(H)` and `PR`.

The functional form is underdetermined. Latané's sub-linear power law [Latané, 1981] and the `n log n` form [Briscoe et al., 2006] fit the same qualitative facts, and **nothing assembled in this corpus discriminates between them.** The general statement of the difficulty is that many mechanisms produce indistinguishable curves [Stumpf and Porter, 2012]. The logarithm is chosen for its axiomatic characterisation, not established by measurement, and the axioms are a choice rather than a finding.

`u` has not been shown to be a quantity. Section 12.7 concedes Michell's objection in full [Michell, 1997], and the conjoint-measurement test that would settle it has not been run [Luce and Tukey, 1964]. If `u` is read as a similarity measure over construals it cannot be assumed to behave metrically at all [Tversky, 1977].

Three attributions in the source material have no entry in this paper's bibliography and are therefore uncited. The cognitive-consistency literature behind mechanism CM2; the forecasting result behind CM5, which is described in Section 5.1 as the closest available direct evidence for this paper's central claim and is not leaned on; and the classical savings-method evidence for forgetting curves behind Section 7.5, from which this paper takes one structural fact and no numbers. **Each is described rather than cited, and no claim in this paper rests on any of them.**

The reflection problem, and why it outranks everything else here. Section 3.4 states it in full. The specification is reflexive at both scopes at which this paper works, and the endogenous effect in such a specification is not identified from observational data [Manski, 1993, Shalizi and Thomas, 2011]. **This paper adopts no remedy.** Every applied reading of `u`, `eta`, `r` or `h` is therefore a **modelling position rather than an identified empirical result**. The most realistic available repair is to stipulate the reference set and defend the stipulation [Brock and Durlauf, 2001]; freezing the framing inventory is a step toward that and not a substitute for it.

The Lucas critique applied to this paper's own parameters. The framework asserts that a described system changes when the description is used, and then treats the parameters as constant over a window [Lucas, 1976, Hacking, 1995]. **Both cannot be true without qualification.** The narrow position taken here — constancy only over a declared window and only while the proxy is not a reported target — is a real restriction on where the model may be used.

Which parameters were chosen rather than measured. All of them. In Section 10.1: the fifteen correlations, `eta`, `delta`, `Q`, `C`, `c1`, `c2`. In Section 10.2: the three retention channels, the six rung yields, the six ledger weights, `E*`, `r*`, `kappa`, `N` and the internal `u`. The collapse rates `lambda` of Section 5.2 are stipulated to bracket a plausible range and **have not been measured anywhere in this corpus.** The `V(n)` cost cases of Table 4 are illustrative parameter sets, not an empirical distribution.

The linear correction load is a postulate with no measurement behind it, and `kappa` is assumed. A convex form is equally plausible a priori. Section 10.2's negative-headroom finding is conditional on `kappa = 0.03`, a number nobody measured, and is a demonstration of a mechanism rather than a measurement of anything.

The reconciliation-cost postulate may be wrong in its form as well as its constant. The claim that internal maintenance load is proportional to the pair count is imported from a group-scope communication argument [Brooks, 1975], and H2's third kill condition — a linear cost model fitting as well — is a live possibility that nothing here excludes.

The conversion ladder's chain assumption is the most likely of the paper's structural claims to fail. If items reach `DEFAULTED` without passing `ENCODED`, the ladder is not a chain, the product form is wrong, and Section 7.6 says so. Some dispositions plainly form without ever being written down.

The monotone claim in `u` is restricted, not defended. The literature contains an inverted U [Uzzi and Spiro, 2005], short-run costs to surface-level diversity [Mannix and Neale, 2005, van Knippenberg and Schippers, 2007], and cases where redundancy reduced safety [Sagan, 1993]. **A monotone surplus cannot survive that list unqualified**, and this paper restricts its claim to the regime the worked cases exercise rather than locating the turn.

Nothing conserved is claimed, and that is a concession. `Q` has the grammar of a conserved quantity and is not one [Mirowski, 1989]. The only invariance claimed is the unit rescaling of Section 12.6, tested at one scale.

Unmet rigour checks, printed rather than left to a reviewer. (i) *Model comparison.* No AIC or BIC is computed anywhere in this corpus, so the corpus's one applied fit never paid for its parameter count; **reported NOT MET** in Section 12.8. (ii) *Effect sizes with intervals.* Nothing in Section 10 carries an interval, because nothing in Section 10 was measured; **reported NOT MET.** (iii) *Scale type.* Conjoint conditions untested; the ratio-scale treatment of `u`, `Q` and the cost ledgers is a stipulation. (iv) *Reproducibility.* Every parameter in this paper was constructed by one author; there is no independent replication of anything, and the reproducibility literature is unambiguous about what that is worth [Open Science Collaboration, 2015, Ioannidis, 2005]. (v) *Directional versus point prediction.* Directional predictions are nearly costless in the social sciences [Meehl, 1967]; this framework does make point predictions with constants, which is the right answer to that objection, but **no point prediction of this paper has yet been put at risk**. (vi) *Small-sample bias.* `rho` estimated on 30 paired predictions has a wide interval, and effective-count estimators computed on small samples are biased in the analogous ecological setting [Chao et al., 2014]; no correction is applied here.

Nothing in this paper is a human-subjects finding. No interventions were run, no participants were recruited, and the declination constraint of Sections 5.3 and 8 is a design constraint on prospective work rather than a report of past practice.

21.14 14 Conclusion

The surplus functional is scope-free. Its axioms quantify over states in an inner-product space and contain no term referring to persons, so applying it to `N` framings held by one observer is a substitution of domain rather than an analogy. The substitution costs exactly two postulates — that a framing admits unit-vector encoding, and that a mean pairwise non-parallelism summarises the set — and yields the **internal surplus** `F_int = ln(1 + (N−1)u)` with `N_eff = 1 + (N−1)u` the audited framing count, `0 <= F_int <= ln N`, and `u = 1 − rho^2` estimated from the correlation between framings' prediction errors — an instrument that asks nobody to judge similarity. **The affine form is design axiom S3, chosen and not derived.**

The audited count is exactly a reparametrised participation ratio. For the Gram matrix `G` of the observer's encoded framings, `N_eff = 1 + (N−1)u_bar = N + 1 − N/PR` with `PR = (tr G)^2/||G||_F^2` the Rényi-2 participation ratio. The mean-pairwise aggregator is therefore not an improvisation; it commits the analyst to a standard diversity statistic. **It does not, however, make `N_eff` a Hill number: an affine transformation of a Hill number is not one.**

A six-framing subject at mean `u = 0.6860` reads `N_eff = 4.4298` and `F_int = 1.4883` nats, 83.07 per cent of the ceiling `ln 6 = 1.7918`. Splitting the pairs by lineage shows what the mean hides: six within-lineage pairs at `u = 0.2732` give the four-framing lineage an audited count of **1.8196 between them**, against nine cross-lineage pairs at `u = 0.9611`.

Restructuring beats enlarging. Retiring the two most redundant framings and importing two from outside the lineage raised the audited count **25.4 per cent** — `F_int` from 1.4883 to 1.7148 nats, 95.70 per cent of ceiling — **at identical framing count and identical reconciliation load.**

And it was not enough. At `eta = 0.60`, `delta = 0.06`, `Q = 12.0` and `C = 0.35`, converted surplus rose from 0.8930 to 1.0289 against a draw of 1.0700: **the intervention closed 77 per cent of the gap and left the one-step change in stored order negative at −0.0411.** Only cutting the operating cost and shedding stored order turned the account, and they turned it at the original geometry as well. **The binding constraint was never the geometry.** The sensitivity table says the same thing structurally: `dh/deta = ln N = 2.197`, `dh/dC = −1`, `dh/dQ = −delta = 0.0317` and `dh/dN = eta/N = 0.0038`, so **adding categories of thought is the weakest available lever**, and doubling the structural budget by raising `N` alone requires `N` to be squared, from 9 to 81.

Stored order and stored error are the same variable, and the remedy is two books. Subtracting the earned recursion from the lived one gives `Z_{t+1} = (1−delta)Z_t + eta(M_t − V_t)` with `Z_inf = eta(M* − V*)/delta`: **stored error integrates on the same pole as stored knowledge and its equilibrium is independent of cost, because the cost term cancels exactly.** In the worked survey a **25.93 per cent** confabulated input compounds into **57.38 per cent** of the steady-state stock, the earned track's equilibrium goes negative between weeks 104 and 156, and the model reports that as its own domain of breakdown rather than concealing it. At equilibrium the headroom identity `h_inf = eta(ln N − M*)` holds exactly, verified to six decimal places at 0.060786.

The survey's own weakest number is reported as weak. Baseline headroom of 0.006156 `Q`-units per week **is not robust to the declared cost weights and its sign is therefore unestablished**; doubling one ledger's weight changes it by more than its entire value. The one figure the same survey regards as robust is `Z_inf`, which is independent of cost by construction — and it is the most uncomfortable number in the report.

All of this is arithmetic. Whether any observer's framings obey the encoding, whether any person's stored order obeys the recurrence, and whether the audited count predicts anything at all are separate questions. They are postulates rather than results, none of the six hypotheses stated here has been tested, and the corpus's one preregistered confirmatory test of the storage law did not meet its own declared margin.

21.14.1 Declarations

Data and code availability

No external dataset was collected for this paper. All numerical results are computed from the stated parameters under the models developed in the text, and every table can be re-derived from the equations and parameter values printed beside it. The two subjects worked in Section 10 are constructed examples with declared parameters, not field observations; the fifteen pairwise correlations of Table 7, the six conversion-rung yields, the three retention channels, the six cost ledger weights and the correction coefficient `kappa` are all stipulated. The figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper and the model equations of Sections 3 to 7.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, produced the worked examples and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. Both worked subjects are constructed. The declination constraint referred to in Sections 5.3 and 8 — that no intervention removing a person's ability to decline is within the framework's scope — is a design constraint on prospective work and is not a report of any study performed. The paper discusses self-assessment, rumination and the accumulation of unearned belief as formal objects; **nothing in it is clinical advice, and no diagnostic claim is made or intended.**

Preregistration

This paper reports no preregistered test of its own. The falsification rule stated in Section 10.1 is declared so that it can be run and **has not been run**. The preregistered protocol for testing the storage law against a real system, together with the confirmatory application that **failed** its second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared margin of 10 per cent — is reported in Curzi [2026i].

Chapter 22

The Launch Budget and the Desire Line: A Rectified Surplus Fund Against a Barrier, and Why Momentum Does Not Cross a Wide Shallow Ridge

Abstract. We define the launch budget as the rectified surplus $L^{Q} = \max(0,\ \eta M - C)$, accumulate it in a decaying fund $\Lambda_{t+1} = (1-\lambda)\Lambda_t + L_t$ with fixed point $\Lambda_\infty = L/\lambda$, and require $\Lambda_\infty > E_b$ before a departure begins, where E_b is the path integral of excess cost across the transit. Feasibility reduces to a critical cost $C_{\text{crit}} = \eta M - \lambda E_b$, computed at 0.082 Q-units per week in the worked case, and to a critical conversion efficiency $\eta_{\text{crit}} = C/M$ below which no plan of any size is ever constructed. The worked operator sat at $\eta = 0.05$ against $\eta_{\text{crit}} = 0.055238$, giving $L = 0$ and $\Lambda_\infty = 0$: a sustained zero budget that cannot diagnose itself. A two-week annual duty cycle accumulates a steady-state peak fund of 0.1162 against a barrier of 2.400 — 4.84 per cent of the requirement, forever. Two independent constructions of E_b agreed to 0.05 per cent. Deliberate discomfort programmes were priced, and the most ambitious, at $+165$ per cent cost, forecloses the departure permanently. On the gradient side only ratios of differences are identified: the office is 4.4042 times the departure cost of the kitchen table, in units that cancel. Momentum was tested at β from 0 to 0.95 and escaped no basin, because the barrier is a plateau with $C'' = -0.005$ rather than a spike; the recursion is recorded as a failed candidate and refused entry to the formulary.

Keywords: launch budget; rectifier; barrier integral; gradient descent; desire line; identified ratios; momentum

Series: B. Engineering the Surplus **Paper:** CCP-022

22.1 1 Introduction

Every account of why people and organisations stay where they are has to answer a prior question: where does the capacity to plan a departure come from, and what happens to that capacity when the cost of staying rises? The common prescription assumes the answer is obvious — that discomfort supplies motive force, and that a person or a firm made sufficiently un-

comfortable will move. This paper shows that, under the framework's own storage arithmetic, the prescription has the sign wrong over most of its range, and it shows this by computing the threshold at which the sign flips.

The construction has three parts and they are simple. The **launch budget** is the rectified surplus left after a period's operating cost has been paid, `L = max(0, eta*M - C)`. The **launch fund** is that budget accumulated in a second leaky integrator with its own decay rate `lambda`, whose fixed point `Lambda_inf = L/lambda` is the largest barrier the regime can ever cross. The **barrier** `E_b` is the path integral of excess operating cost across a transit between two configurations of a measured cost field. A departure is feasible if and only if `Lambda_inf > E_b`, and rearranging that inequality yields a **critical cost** `C_crit = eta*M - lambda*E_b` above which no amount of waiting funds the crossing. The second half of the paper measures the field the barrier is computed on, using choice frequencies rather than self-report, and reports what such a measurement can and cannot identify.

Four things in the argument are worth separating from the framework they sit in.

First, the **rectifier**. `max(0, .)` is not decoration: a period whose cost exceeds its converted moment produces no negative planning capacity, it produces none at all, and the deficit reappears in the following period as a higher cost rather than as a negative budget. The consequence is a self-concealing regime. An subject at a sustained `L = 0` does not experience being trapped, because noticing that an alternative configuration exists is itself funded by `L`. The diagnostic must therefore come from outside the system, and this is a structural claim rather than a psychological one.

Second, the **stationarity proposition**. Where `L > 0`, `grad L = -grad C / eta`, so the stationary points of the launch budget are exactly the stationary points of operating cost and the maxima of the budget are exactly the minima of cost. The capacity to plan a departure is largest at the local minimum of the cost field and smallest on the ridge between basins — that is, in the middle of the crossing. This inverts the ordinary prescription. Cost reduction at the origin is not the obstacle to the journey; it is the platform the journey is funded from, and the point is arithmetic rather than temperamental.

Third, the **negative result on momentum**, which this paper reports at the same prominence a positive one would get. Momentum is the standard cheap escape from a local minimum in first-order optimisation, and on the measured cost field of the worked case it escapes nothing: at `beta = 0`, `0.5`, `0.8`, `0.9` and `0.95` the walker terminates at `x = 1.9972` in basin A at every setting, because the barrier is a plateau with second derivative `-0.005` spread over more than two configuration-units rather than a narrow spike. The recursion is recorded as a **failed candidate** and is refused entry to the framework's formulary. What survives is one empirical statement with no free parameters: momentum defeats narrow barriers and is useless against wide shallow ones, and most barriers of practical interest — an income gap, a retraining period, a slow renegotiation — are wide and shallow.

Fourth, the **identification result** on the measurement side. A desire-line survey recovers cost differences only up to an unknown choice temperature `theta`. Absolute costs are not recovered and must never be reported; **ratios of differences are fully identified**, because `theta` cancels. In the worked survey the office is 4.4042 times as costly a departure from the desk as the kitchen table is, and that statement is unit-free and testable while "the office costs 1.95" is not. The same ratio form turns out to be roughly four times more robust to an unmodelled switching cost than the differences it is built from, which is a second and unplanned reason to prefer it.

The framework's static and dynamic halves are presupposed rather than re-derived here. The founding definition — that comfort is the felt total of present costs, so that an subject observed at rest is at a local minimum of their own cost field by definition rather than by luck, and that the reflexive loop `T(x) = W(x, M(x))` has the self-fulfilling prophecy as a fixed point of the composition of seeing and responding — is established in Curzi [2026]. The contraction theorem the fund leans on is stated exactly — a contraction with modulus `q < 1` on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least `q` each pass — in Curzi [2026]. The storage law whose steady-state theorem the fund reuses

with `delta := lambda`, `Q_inf = (eta F(u*) − c*)/delta`, is Curzi [2026a].

Sections 2 and 3 state the hypotheses and the formal framework, including the identification problems the framework cannot solve. Section 4 develops the launch budget, the fund, the barrier integral and the critical cost. Section 5 develops the gradient, the desire-line instrument, the two candidate escapes and the regrade analysis. Section 6 collects the seven-gate formula-validation audit for all five formulas. Sections 7 and 8 give the design specifications and the named failure modes. Section 9 works one subject end to end across both halves. Sections 10 to 12 discuss what follows, verify the chain from minimal postulates, and state what the account cannot support.

22.2 2 Hypotheses

Several results below are theorems about equations and cannot be falsified by observation; they are labelled as derivations where they appear. What can be falsified is the claim that these equations describe anything real. The hypotheses separate the two, and each carries the observation that would defeat it.

H1 — The launch budget is rectified and the rectification is consequential. The capacity to construct, hold and commit to a plan for a different configuration is drawn from the current period's surplus `L = max(0, eta*M − C)` and is unavailable at `L = 0`. **Predicted:** the incidence of newly constructed plans for a different configuration is monotone increasing in contemporaneous `L`, and is zero when `L = 0`. **Falsified by:** a single well-documented plan constructed during a sustained `L = 0` stretch, under prospective daily logging of `M`, `C` and plan-initiation events.

H2 — Plans decay geometrically, so the fund has a ceiling. Accumulated planning capacity is a leaky integrator with its own rate `lambda`, so the fund approaches `Lambda_inf = L/lambda` and never exceeds it. **Predicted:** exponential decay of plan persistence in the absence of reinforcement, with half-life `ln 2 / (−ln(1−lambda))`; and a maximum crossable barrier equal to `L/lambda` however long the subject waits. **Falsified by:** plan persistence that does not decay geometrically without reinforcement — a threshold or step-function decay would do it, and either is plausible; or a completed departure across a barrier reliably exceeding the measured `L/lambda`.

H3 — The critical cost is a real threshold, not a bookkeeping artefact. Departures cease as operating cost crosses `C_crit = eta*M − lambda*E_b`. **Predicted:** across a population of would-be departures, the observed departure rate falls to zero as `C` crosses a threshold **predicted from parameters measured beforehand**. **Falsified by:** departures occurring at the same rate above and below the computed threshold. This requires a cohort and is the most expensive test in the paper.

H4 — The barrier integral predicts realised transit cost. `E_b`, computed as the path integral of excess cost at the intended transit speed, predicts the cost actually incurred in crossing. **Predicted:** realised transit cost, ledgered after the fact, matches the pre-computed `E_b` within the stated interval. **Falsified by:** realised cost uncorrelated with the predicted barrier across a set of completed transitions. This is the cheapest test in the paper and can be run retrospectively by anyone who computed `E_b` beforehand.

H5 — Choice frequency recovers relative cost. Under a logit choice model, log-frequency ratios over available occasions recover differences in operating cost up to the scale `theta`. **Predicted:** a frozen difference table fitted on days 1 to 20 beats a uniform baseline, a most-frequent-option baseline, a first-order Markov baseline and a day-shuffled control at predicting held-out days 21 to 30; and adding a fourth option does not change the ratio of the first two. **Falsified by:** failure against the

shuffled control, which would mean the subject has no stable cost field over the surveyed coordinates; or a material shift in the recovered top-two ratio when the choice set is experimentally manipulated, which would falsify the logit form for that coordinate.

H6 — Momentum does not cross a wide shallow barrier. Adding a momentum term to descent does not move a walker out of a basin whose exit ridge is a plateau. **Predicted:** on a field whose ridge curvature is small in magnitude relative to the basin curvature and whose ridge extends over a distance comparable to the basin separation, terminal position is independent of `beta`. **Falsified by:** escape at any tested `beta` on such a field. **This hypothesis was tested here and was not falsified; Section 5.10 reports the result and Section 10 reports what it does and does not license.**

Three further claims are **not** hypotheses and are labelled accordingly wherever they are used.

The Launch Postulate is a postulate. That planning capacity is drawn from `L` is a claim about persons, not a theorem. Nothing in the framework's foundational results implies it and none of them would be touched if it were false. It is falsifiable — H1 states how — and until it is tested it is a modelling assumption that has earned its place by explaining a common observation, not by having been measured.

The Gradient Postulate is a postulate. That an subject's motion in configuration space is, in expectation, in the direction `-grad C` with a step magnitude increasing in `|grad C|` is a modelling assumption. It borrows the *form* of a stationary-action principle and nothing else: no physical principle is invoked, nothing is conserved, and `C` is not a potential in any sense a physicist would recognise. Section 3.7 makes the category of every physics borrowing in this paper explicit.

That a transition's cost can be summarised by a single scalar `E_b` is a modelling choice, adopted because it makes the feasibility condition a comparison of two numbers. Section 12 records what it discards.

22.3 3 Framework

22.3.1 3.1 What is presupposed

Throughout, **subject** denotes the modelled individual or unit whose configuration and cost field are under study. It is a role term inside the model and not a form of address; where a claim concerns the person conducting the survey rather than the one surveyed, the word used is **analyst**.

Three results are used throughout and are established elsewhere in this corpus.

The founding definition and the reflexive loop. Comfort is the felt total of present costs, so an subject observed at rest sits at a local minimum of their own cost field by construction; and the reflexive loop `T(x) = W(x, M(x))` identifies the self-fulfilling prophecy as a fixed point of the composition of seeing and responding [Curzi, 2026]. Section 3.6 states why that loop is not identified from observational data.

The contraction theorem. A contraction with modulus `q < 1` on a complete space has exactly one fixed point, reached from any start, with the distance to it shrinking by at least `q` on each pass [Curzi, 2026, Banach, 1922]. The launch fund's ceiling is this theorem applied to a second integrator.

The storage law and its steady state. `Q_{t+1} = (1-delta)Q_t + eta*S_t - C_t` with `S_t = ln(1 + (N-1)U_t)` is a first-order linear difference equation whose fixed point under constant drive is `Q_inf = (eta F(u*) - c*)/delta`, and whose update map is a contraction for every `delta > 0` [Curzi, 2026a]. The launch fund of Section 3.4 is that theorem with `delta := lambda` and

the drive term replaced by L. The headroom gauge h_t = eta*ln N − delta*Q_t − C_t, its three regimes and its separate admissibility and sustainability constraints are established in Curzi [2026b], and Section 5.12 shows that descent in the cost field is ascent in that gauge under a stated restriction.

Two properties of the surplus term matter here and both carry qualifications the corpus states rather than hides. The ceiling ln N is the maximum an ensemble of N equiprobable categories can carry [Shannon, 1948], and the framework does **not** identify its surplus with Shannon entropy. The logarithm has two independent warrants and this paper needs only the weaker one: Peters and Gell-Mann [2016] shows that the time-average growth rate of a multiplicatively evolving quantity is the expectation of its logarithm, so the logarithmic evaluation follows from the dynamics with no psychological premise, while the Fechnerian warrant [Fechner, 1966] requires a contested perceptual story. **The functional form is nevertheless underdetermined by the evidence.** A sub-linear power law [Latané, 1981] and an n log n form [Briscoe et al., 2006] reproduce the same qualitative facts — diminishing returns to count, no threshold, monotonicity — and nothing assembled in this corpus discriminates among the three. Finally, the effective count 1 + (N−1)u used inside the surplus functional is **not** a Hill number: 1/HHI and exp(H) are the Hill numbers of order two and one and are well grounded [Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979], while 1 + (N−1)u is a linear interpolation between 1 and N and is this framework's own construction. Every appearance of an effective count in this paper says which one it is.

22.3.2 3.2 The launch budget, and the register split that a dimensional check forces

The canonical statement of the launch budget is

```
L = max( 0 ,  M − C )

M   the moment as lived this period
C   the operating cost of this period
L   the surplus available to be spent on anything other than continuing
```

and **as written it does not survive a dimensional check.** M is in nats; C is in Q-units per period; they cannot be subtracted. This is the audit doing its work rather than a defect to be apologised for, and the resolution is worth more than the original formula. Two admissible registers exist and both are correct:

```
THE STORED FORM     L^Q = max( 0 ,  eta ∗ M_t  −  C_t )      Q−units per period
THE LIVED FORM      L^M = max( 0 ,  M_t  −  C_t / eta )      nats per period

and                 L^Q = eta ∗ L^M
```

The stored form asks how much durable surplus remains after this period's deposit into the self and after the operating cost. The lived form asks how much of this period's moment exceeded what it cost to have, with C/eta reading the operating cost back into the register of moment as the moment that would have been required to pay for it.

Since eta > 0 the two forms differ by a strictly positive constant factor, so **they have the same zero set and the same sign**, and every qualitative claim in this paper depends only on the sign. That is the honest resolution: the canonical statement is register-ambiguous, the ambiguity is a scale factor of eta, and no claim here is affected by the choice — but every *number* is, so every number declares one. The convention adopted throughout is L^Q for anything that will be accumulated and L^M for anything that will be compared to a moment; both are always reported.

22.3.3 3.3 The rectifier

`max(0, .)` is a **half-wave rectifier**, and in the framework's circuit register that is a precise object: a diode in series with the accumulator [Horowitz and Hill, 2015]. This is a **shared mathematical form** and not a derivation — the same nonlinearity governs both, and nothing about semiconductor physics is claimed to hold of a person.

There is no negative planning. A period in which cost exceeded converted moment does not produce anti-plans; it produces nothing. The deficit does not vanish — it appears in the next period as a higher `C`, through unfinished work carried, sleep debt and unrepaired relations — but it does not appear as a negative launch budget, because the quantity being measured is *available surplus*, and surplus has a floor at zero by construction.

The asymmetry has a consequence that is frequently reported and rarely explained: **a bad period costs the plan that would have been made, and then also charges for the plan that was not made.** The first is the rectifier; the second is the cost carry. The rectifier also destroys information, and Section 6.1 records the cost: `max(0, .)` cannot distinguish "just barely trapped" from "catastrophically trapped", so the depth of the deficit must be tracked separately.

22.3.4 3.4 The launch fund

A single period's `L` is almost never enough to fund a departure. Surplus must be stored, and stored order leaks, so the fund is a second leaky integrator with its own rate:

```
THE LAUNCH FUND

    Lambda_{t+1} = (1 - lambda) Lambda_t  +  L_t

    Lambda_t   accumulated committed surplus available to spend on a transition, Q-units
    lambda     plan decay per period: forgetting, changed circumstance, loss of nerve,
               drift of aim

    Lambda_inf = L / lambda
```

The fixed point is the storage law's steady-state theorem with `delta := lambda` [Curzi, 2026a], and it is the DC gain `1/lambda` of a first-order discrete section applied to the net drive [Oppenheim et al., 1996, Elaydi, 2005]. The same recursion is exponential smoothing under another name [Holt, 1957, Brown and Meyer, 1961, Hyndman et al., 2008], and the level-and-rate vocabulary it belongs to predates the framework by four decades [Forrester, 1961].

`Lambda_inf = L/lambda` is the single most useful number in this paper. It is not the largest barrier the subject will cross this year; it is the largest barrier the regime can cross **at all**, at steady state, with unlimited patience.

The stock-and-flow distinction here is load-bearing and is violated constantly in ordinary talk about motivation: `Lambda` is a stock in Q-units, `L` is a flow in Q-units per period, and people systematically misjudge the relation between the two when a delay sits between them [Sterman, 1989, 2000]. The fund's behaviour is exactly the behaviour that misjudgement gets wrong.

22.3.5 3.5 The cost field, its gradient, and the logit recovery

```
THE COMFORT FIELD AND ITS GRADIENT

    C : X -> R          the operating cost of occupying configuration x, Q-units per period
    g(x) = - grad C(x)  the descent direction: the direction of cheapest motion at x

    DISCRETE MOTION     x_{t+1} = x_t - alpha * grad C(x_t)
```

X is the configuration space of arrangements the subject could occupy and C(x) is the operating draw evaluated at each. **The desire line is the observed integral curve of −∇C.** The classification of stationary points by the Hessian — positive definite a rest basin, negative definite a ridge, indefinite a saddle, singular a plateau — and the partition of the space into basins of attraction are standard results for gradient systems [Strogatz, 2015]. Modelling the motion of people as motion in a cost field is not novel either: the social-force model treats pedestrian trajectories as motion under a gradient and reproduces observed path formation [Helbing and Molnár, 1995, Helbing et al., 2000]. That precedent is **a shared mathematical form**, and the present paper claims no more from it.

The instrument that makes the field measurable is the logit recovery. Under a choice model in which the probability of an option falls exponentially in its cost,

```
THE LOGIT RECOVERY

    P(a) / P(b) = exp( −( C_a − C_b ) / theta )

    =>   C_a − C_b  =  −theta ∗ ln( P(a) / P(b) )

    estimated from counts:   C_a − C_b  =  −theta ∗ ln( n_a / n_b )
    standard error (in theta units):     se = sqrt( 1/n_a + 1/n_b )
```

theta is a scale parameter carrying the units of C — the choice temperature. The specification is the standard workhorse of discrete-choice measurement [Brock and Durlauf, 2001], and recovering scale values from choice proportions rather than from numerical self-report is older still [Thurstone, 1927]. **It is a modelling postulate about behaviour and is badged as one**; its characteristic weakness, independence of irrelevant alternatives, is stated with its test in Section 6.5 and its failure mode in Section 8.2.

22.3.6 3.6 Identification: three things no amount of data recovers

The reflection problem. This paper presupposes the framework's reflexive loop T(x) = W(x, M(x)) [Curzi, 2026]: the group influences the member and the members constitute the group. That is exactly a linear-in-means reflexive specification, and Manski [1993] proves the endogenous social effect in such a specification is **not identified from observational data**. It cannot be separated from contextual effects or from correlated unobservables however much data is collected, and influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011]. **This paper adopts no remedy.** The known options are randomisation, an instrument, partial-population variation, or stipulating the reference group and defending the stipulation [Brock and Durlauf, 2001]; the last is the most realistic route here and the desire-line protocol's frozen, pre-declared option set is a step toward it rather than a substitute for it. Every claim in this paper that leans on the reflexive loop is therefore **a modelling position rather than an identified empirical result**, and it is said again in Section 12. That a reflexive fixed point can be caught happening in the field, with dates and prices, is documented [MacKenzie and Millo, 2003]; that is an existence proof for the mechanism and not an identification of this paper's parameters.

The choice temperature. theta is not identified from choice frequencies alone. It requires an external anchor: an experimentally imposed cost of known magnitude, from whose effect on the share of an option theta follows. **That experiment has not been run here**, which is why every difference reported in Section 5 is in theta units and why only ratios of differences are advanced as findings. What licenses the ratio and not the difference is a standard point in measurement theory: the admissible statements are the ones invariant under the transformations the scale permits [Narens, 1996, Stevens, 1946].

The step size. In the discrete descent x_{t+1} = x_t − alpha∗grad C(x_t), the step size alpha and the cost scale theta are not separately identified: only their product governs the observed

trace, so two subjects with different responsiveness and different cost scales produce identical traces. Section 6.4 records that gate as a failure rather than a caveat, and the specification requires the product to be reported or neither.

The Lucas critique, applied to this framework's own parameters. The framework asserts that a system changes its behaviour when it is modelled and observed, and then treats `eta`, `lambda` and `theta` as constants over an estimation window [Lucas, 1976]. Both cannot hold without qualification. The position taken here is the narrow one: the parameters are treated as stable **only over a declared window and only while the fund is not itself a reported target**, since a measure that becomes a target ceases to be a measure [Goodhart, 1975, Campbell, 1979]. Where an subject's `Lambda` is reported to anyone with an interest in it, the constancy assumption fails and the model is outside its domain.

22.3.7 3.7 Analogy discipline, and the standing honesty clauses

Following the corpus's rule, every borrowing from physics or from optimisation is labelled at first use as a **derivation** (the premises are genuinely satisfied), a **shared mathematical form** (the same equation governs both, with no transfer of physical content), or an **analogy** (a suggestive correspondence with no formal warrant, legitimate as motivation and never as evidence).

Borrowing	Where	Category
leaky integrator for the launch fund	Section 3.4	**shared form** — the same first-order difference equation, no claim about charge
half-wave rectifier for `max(0, .)`	Section 3.3	**shared form** — the same nonlinearity, no claim about semiconductors
DC gain, pole, time constant	Sections 3.4, 4.3	**shared form** — standard first-order-section results applied to the same equation
gradient descent and its convergence to a stationary point	Sections 3.5, 5.8	**derivation within the model** — the model *defines* motion as descent, so the algorithmic theorem applies to the model exactly; whether an subject descends is the Gradient Postulate and is separate
basins of attraction, Hessian classification	Sections 5.8, 5.9	**derivation within the model**, from Strogatz [2015]
stationary action for the Gradient Postulate	Section 2	**analogy** — form only; no conservation law, no teleology, and the honest word is stationary rather than minimal
barrier `E_b` as a path integral of excess cost	Section 4.6	**shared form** — an accumulated cost, explicitly **not** an energy
momentum, and temperature as annealing	Section 5.10	**analogy** — imported as candidate mechanisms and tested; one was refused

The failure mode this table exists to prevent is a shared-form borrowing quietly acquiring the certainty of a derivation. The critique of that move in economics is long-standing and

applies here [Mirowski, 1989].

Standing honesty clauses, repeated from the corpus. The triangle inequality fails for geometric triples. The surplus functional as deployed is a lower bound on the ensemble's information. The framework makes no empirical or metaphysical claim about physics, does not identify surplus with Shannon entropy, and proves no thermodynamic law. **Nothing in this paper says anything about physics, and `E_b` is a barrier in a modelled cost field, not an energy.**

22.4 4 The launch budget against a barrier

22.4.1 4.1 The low-cost configuration adjacent to a transition

The framework's founding node holds that the present moment is the comfort zone by definition, since the subject is always already in it, and that the desire to be elsewhere is not evidence of being in the wrong place but a product manufactured under specific conditions [Curzi, 2026]. The condition can now be stated: `L > 0`.

The canonical image is a departure lounge, and the arithmetic explains why it is exact rather than decorative. A departure lounge is a low-cost configuration adjacent to a transition. Nothing is required of the occupant there: the attention ledger is quiet, the obligation ledger is suspended, the regulation ledger is idle, and `C` collapses. `M` does not — there are people to watch, a window, an unfamiliar arrangement of objects, a suspension of routine which is itself informative. **`M` holds and `C` falls, so `L` rises, and what `L` buys is the construction of elsewhere.**

The counter-image is a lift with both hands full, halfway through a task, load high, attention fragmented, nothing put down. `M` is not especially low — a great deal is happening — but `C` is large, and `eta*M - C` is negative:

```
the locked regime:    eta = 0.05,  M = 1.05 nats,  C = 0.058 Q-units per week
                      C / eta = 1.160 nats
                      L^M = max(0, 1.05 - 1.16) = 0
                      L^Q = max(0, 0.0525 - 0.058) = 0
```

Zero, not small. An subject running this regime is not deficient in ambition; they are running a rectifier that is reverse-biased.

The practical consequence follows immediately and is the reason the analysis is worth doing. **An subject in a sustained `L = 0` regime cannot diagnose their own situation, because diagnosis is also manufactured from `L`.** The construction of an alternative configuration — noticing that one is possible, holding it in mind, weighing it — is exactly the activity the budget funds. A person with no launch budget does not report being trapped. They report having no particular ideas. This is a structural claim about the model, and it converges with the empirical literature on self-report from a different direction: verbal reports do not have privileged access to the processes that generate behaviour [Nisbett and Wilson, 1977], and reported states are strongly shaped by the conditions of the asking [Schwarz, 1999].

22.4.2 4.2 Where the budget is largest: the stationarity proposition

The observation the model has to account for is that plans for a different arrangement arrive at rest — on the fourth day of a holiday, in the bath, on a walk, in a departure lounge — and not in the middle of a task. The model makes this a proposition about the cost field.

Let the configuration be a point `x` in a space of arrangements, let `C(x)` be the operating cost at that configuration, and hold `M` and `eta` fixed across the neighbourhood. Then `L^M(x) = max(0, M - C(x)/eta)`.

> **Proposition 4.1.** Wherever `L^M(x) > 0`, `grad L^M(x) = -grad C(x)/eta`. Therefore the stationary points of `L^M` are exactly the stationary points of `C`, and the local maxima of `L^M` are exactly the local minima of `C`.

Proof. On the open region where `M − C/eta > 0` the rectifier is inactive and `L^M = M − C/eta`, so `grad L^M = −grad C/eta`, which vanishes if and only if `grad C` vanishes. The Hessian of `L^M` is `−H_C/eta`, so it is negative definite exactly where `H_C` is positive definite. This is a **derivation within the model**; it consumes the model's differentiability assumptions and nothing else.

Two readings follow.

First: the plan for elsewhere is manufactured at local minima of cost and nowhere else. That it arrives on holiday is not a coincidence and that it does not arrive mid-task is not weakness. It is where the gradient of the cost field is zero and the curvature is positive.

Second, which is the same statement in different clothes: the comfort is the platform. The journey is planned from the local minimum and departed from the local minimum, because the local minimum is the only place where the budget to do either exists. The organisational counterpart is the standard result that search and experimentation are funded out of slack rather than out of pressure [Cyert and March, 1963], and that search is triggered and bounded by the resources an organisation has spare [March and Simon, 1958].

Computed on the cost field constructed in Section 5.8 and shared between both halves of this paper, at the corrected regime `eta = 0.10`, `M = 1.42`, `lambda = 0.025`:

Configuration	`C(x)` (Q-units/wk)	`L^M` (nats/wk)	`L^Q` (Q-units/wk)	`Lambda_inf` (Q-units)
basin A floor (`x = 2.00`)	0.03400	1.0800	0.10800	4.3200
the ridge between basins (`x = 4.335`)	0.09150	0.5050	0.05050	2.0200
basin B floor (`x = 7.00`)	0.02880	1.1320	0.11320	4.5280
the A flank, mid-slope (`x = 1.332`)	0.05800	0.8400	0.08400	3.3600

The largest launch budgets are at the two basin floors and the smallest is on the ridge — that is, in the middle of the crossing. The subject has the least capacity to plan the journey precisely while making it, which is why journeys must be planned before they are begun and why mid-transit is the worst available moment to reconsider the destination.

22.4.3 4.3 The fund, and why plans have a half-life

A single week's `L^Q` of 0.108 Q-units does not fund a barrier of 2.4. The surplus must be stored and stored things leak; `lambda` is the leak.

What `lambda` is, for a plan: forgetting the reason; a change in circumstance that invalidates a premise; the slow loss of nerve; drift of aim, in which the plan is still held but no longer points at anything the subject now wants. It is the storage law's `delta` for a specific and fragile kind of stored order, and it is generally much larger than the `delta` of the self, because a plan has no enacting behaviour maintaining it [Curzi, 2026a].

`lambda` per week	fund time constant `1/lambda` (wk)	plan half-life (wk)
0.020	50.0	34.31
0.025	40.0	27.38
0.040	25.0	16.98
0.060	16.7	11.20
0.100	10.0	6.58

A plan held at `lambda = 0.06` is at half strength eleven weeks after the last time anything was added to it. That is the arithmetic behind the observation that a resolution made in January is gone by April with no event having occurred to end it. Nothing ended it; it leaked.

And `lambda` is engineerable by the same devices that raise `eta` for the self, applied to the plan instead: **write it down, date it, tell one person, give it an artefact, put a recurring entry in the calendar.** That the same acts raise conversion in one place and lower decay in another is not a coincidence — both are the act of moving something out of working memory into a durable medium.

22.4.4 4.4 The maximum crossable barrier

```
Lambda_inf = L^Q / lambda
```

Because the fund is a contraction, this ceiling is approached and never exceeded [Curzi, 2026]. **It is the largest barrier the regime can cross, ever.** Held at `M = 1.42` and `eta = 0.10`, and varying the two engineerable parameters:

`C` (Q-units/wk)	`L^Q`	`lambda = 0.06` gives `Lambda_inf`	`lambda = 0.025` gives `Lambda_inf`
0.020	0.12200	2.033	4.880
0.034	0.10800	1.800	4.320
0.058	0.08400	1.400	3.360
0.082	0.06000	1.000	2.400
0.100	0.04200	0.700	1.680
0.142	0.00000	0.000	0.000

Two things in that table are load-bearing.

Cutting `lambda` from 0.06 to 0.025 multiplies the crossable barrier by 2.4 at every cost level, and costs nothing but a notebook and a date. It is the cheapest intervention available here by a wide margin: nothing about the subject's life changes, only the durability of the plan.

At `C = eta*M = 0.142` the whole column collapses to zero. There is a cost above which the launch budget is zero and therefore the crossable barrier is zero regardless of `lambda`, of patience, and of how strongly the departure is wanted.

22.4.5 4.5 The minimum conversion efficiency for any departure at all

Rearranging `L^M > 0`:

```
L^M > 0   <=>   eta > C / M

at C = 0.034, M = 1.42:      eta_crit = 0.023944
```

Below a conversion efficiency of 2.39 per cent this subject has no launch budget at any moment size and no plan is ever constructed. Sensitivity, at `C = 0.034` and `M = 1.42`:

`eta`	`C/eta` (nats)	`L^M`	`L^Q`	`Lambda_inf` at `lambda = 0.025`
0.020	1.7000	0.0000	0.00000	0.000
0.034	1.0000	0.4200	0.01428	0.571
0.050	0.6800	0.7400	0.03700	1.480

eta	C/eta (nats)	L^M	L^Q	Lambda_inf at lambda = 0.025
0.100	0.3400	1.0800	0.10800	4.320
0.200	0.1700	1.2500	0.25000	10.000

Doubling `eta` from 0.05 to 0.10 nearly triples the crossable barrier, from 1.480 to 4.320, because `eta` enters the expression twice: once multiplying the moment and once dividing the cost. This is the strongest argument available for the framework's conversion ladder and it arrives from an unexpected direction. Raising conversion does not merely build a larger self; it builds the capacity to leave. The result is consistent with the corpus's finding that conversion is the first lever and is more often binding than the count of participants [Curzi, 2026b].

22.4.6 4.6 The barrier, constructed

`E_b` has so far been treated as a given scalar. It is not: it is an accumulated cost, and the construction matters because it is where most estimates go wrong.

A transition from configuration `A` to configuration `B` passes through configurations whose operating cost exceeds `C(A)`. The barrier is the accumulated excess:

```
E_b = INTEGRAL over the transit of  max(0, C(x) - C(A))  dt
    = ( 1 / v ) * INTEGRAL over the path of  max(0, C(x) - C(A))  dx    at constant speed v
```

On the shared cost field of Section 5.8 the path integral of positive excess from basin A to basin B is 0.182961 Q-unit configuration-units per week, and the basins are 4.996 configuration-units apart:

transit duration `T` (wk)	speed `v` (units/wk)	`E_b` (Q-units)
40	0.12490	1.4649
52	0.09608	1.9043
65.5	**0.07627**	**2.3987**
78	0.06405	2.8565
104	0.04804	3.8086

A slower crossing costs more in total, because the excess cost is paid per period and a slow crossing spends more periods paying it. Dragging out a transition is not caution; it is expense.

It is also, as written, wrong in one limit, and this is stated rather than hidden. The model says `E_b -> 0` as `v -> infinity`, which is false: at high transit speed the cost field itself rises, because doing a thing fast is more expensive per unit than doing it at a moderate pace. A correct treatment needs a speed-dependent term, `C(x, v) = C(x) + b*v^2` or similar, which produces an interior optimum. That term is not developed here. **What the constant-field model supports is the comparative claim — that slower is dearer over the range where speed does not itself raise cost — and it does not support an unbounded "go as fast as possible".** The transit-duration column should be read as valid between roughly 40 and 104 weeks for this case and should not be extrapolated; specification SP-L12 makes that a requirement.

A discrepancy in the source data, reported rather than corrected. The separation used above is 4.996 configuration-units, but the two basin floors located in Section 5.8 sit at `x = 1.9970` and `x = 6.9960`, whose difference is 4.999. Recomputing the barrier at the consistent separation gives `E_b = 2.3973` rather than 2.3987, a difference of 0.06 per cent, which changes no conclusion and no verdict in any table in this paper. The figure 2.3987 is retained throughout for consistency with the worked case, and the inconsistency is recorded here so that a reader recomputing the integral is not left wondering.

22.4.7 4.7 The order of operations, and the prescription that reverses it

The common instruction is to get out of the comfort zone. Here is what the arithmetic says the instruction does. The comfort zone is the local minimum of the cost field. Leaving it in the sense the instruction means — deliberately adopting a higher-cost configuration in order to be motivated by the discomfort — raises `C`. Raising `C` lowers `L^Q = max(0, eta*M - C)`. Lowering `L` lowers `Lambda_inf = L/lambda`. Lowering `Lambda_inf` lowers the largest barrier that can ever be crossed.

The instruction reverses the order of operations. The comfort is not the obstacle to the journey; it is the platform the journey is planned on and departed from.

Computed at `eta = 0.10`, `M = 1.42`, `lambda = 0.025`, against a barrier of `E_b = 2.40`, with a deliberate discomfort programme of increasing severity:

Programme	`C` (Q-units/wk)	`L^Q`	`Lambda_inf`	Weeks to fund the departure
baseline (comfort preserved)	0.034	0.10800	4.320	**32.0**
mild discomfort (+50%)	0.051	0.09100	3.640	42.5
serious discomfort (+120%)	0.075	0.06700	2.680	89.2
the critical cost	**0.082**	0.06000	**2.400**	**never**
ambitious discomfort (+165%)	0.090	0.05200	2.080	**never**

```
C_crit = eta * M - lambda * E_b = 0.10 * 1.42 - 0.025 * 2.40 = 0.142 - 0.060 = 0.082
```

Above a weekly operating cost of 0.082 Q-units this subject cannot make this departure, at any level of desire, given unlimited time. The programme intended to produce the departure has, past that line, made it impossible — and by a mechanism the subject cannot perceive, because the capacity to perceive it is also funded by `L`.

The result has a queueing analogue worth naming, because it is the same shape in a better-measured domain: as utilisation approaches capacity, waiting time does not degrade gracefully but diverges [Kingman, 1962, Hopp and Spearman, 2011]. Operating an accumulating system with no headroom is not a stern regime; it is a regime whose response time is unbounded.

4.7.1 The steelman, and what the prescription is groping at

The instruction is not stupid. It is groping at four things that are true, and the framework has a correct version of each.

One: a local minimum need not be a good one. Descent is local and lands wherever the basin boundary put it; there is no guarantee that the basin the subject rests in is the deepest available. This is the myopia of learning in its exact form: local search exploits the neighbourhood it starts in and systematically fails to find distant optima [Levinthal and March, 1993, March, 1991]. **Correct version: survey the landscape rather than raise the subject's own floor.** Knowing that a better basin exists is a measurement problem and it costs measurement, not suffering.

Two: cost must be paid to cross. Real: `E_b > 0` for every departure worth the name. **Correct version: the cost is paid in transit, not in the lounge, and it is paid out of a fund accumulated**

in the lounge. The transit is uncomfortable; the preparation is not, and should not be made so, because the preparation is what pays for the transit.

Three: a stock built with no variety stops growing. An subject whose weeks are identical is running a low-variety regime with a low `M`. **Correct version: raise `M` — which means enrich the environment and calibrate the reading of it — and `L = max(0, eta*M - C)` rises.** That is the direction the instruction wants, achieved through the term that produces surplus rather than the term that consumes it.

Four, and it is the strongest: a cost rise confined to the ORIGIN lowers the barrier as well as the budget. `E_b` is measured as excess cost over the starting configuration, so raising the starting configuration's cost lowers `E_b` at the same time as it lowers `L`. The net effect is not obvious and must be computed; Section 9.7 computes it on the measured field. The result: **raising the origin alone does shorten the wait, from 32.0 weeks to 9.87 weeks at a rise of 0.040, and the price is that `Lambda_inf` falls from 4.320 to 2.720, so the subject departs sooner and departs with 37 per cent less capacity to cross whatever lies beyond the destination** — a barrier which, by Section 5.9, the local gradient cannot see. Raising cost *globally*, which is what the table above computes and what most discomfort programmes actually deliver, leaves `E_b` untouched and is strictly harmful at every setting.

The distinction is between a cost that can be left behind and a cost that is carried. Attention, regulation, carry, sleep, health and relations are not attached to a configuration and do not stay behind when it is left. **A programme that promises a local regrade and delivers a global one is the ordinary case.**

4.7.2 Why the prescription nonetheless produces testimonials

Because it sometimes works, and it is worth being precise about when, since a framework that cannot explain its opponent's successes is not a framework.

The discomfort programme raises `C`. It also, frequently and incidentally, raises the environment term: the person goes somewhere unfamiliar, meets people outside their lineage, does something they have not done. `M` rises. If `eta*(M' - M) > C' - C`, the net effect on `L` is positive and the programme works — **not because of the discomfort but in spite of it, and only because a second, unremarked change was bundled with it.**

```
the programme succeeds  <=>  eta * (M' - M) > C' - C
```

Which yields a diagnostic: **strip the discomfort, keep the novelty, and the result improves.** That is a testable prediction, it is cheap to test, and it separates this account from the prescription's in exactly one experimental step.

Two further mechanisms in the literature bear on the testimonial and both cut against the prescription rather than for it. Adaptation is fast and thorough: subjective states return toward a baseline after large changes in circumstance in both directions [Brickman et al., 1978, Helson, 1948], though the setpoint is not fixed and does move [Diener et al., 2006]. And attention to a prospective change inflates its apparent importance while it is being attended to [Kahneman et al., 2006], which is exactly the condition under which a discomfort programme is evaluated by the person undergoing it.

4.7.3 The consent clause

The framework's ethical constraint governs absolutely here and the failure mode is specific.

No intervention may deliberately raise another person's `C` in order to force a departure. The mechanism does not merely risk harm; it works, when it works at all, by removing the target's capacity to construct alternatives, which is the precise definition of removing their ability to decline. And past `C_crit` it does not even achieve the intended departure: it forecloses it while leaving the subject with the belief that they lack resolve.

The permitted intervention has the opposite sign: **lower the cost of the better configuration and leave the choice intact.** Section 5.11 is the method. Its worst case is bounded — someone looks at the lower valley and stays where they are — and that bounded worst case is what makes it admissible.

22.4.8 4.8 Four ways to fund a departure, and the fifth

Given `Lambda_inf = (eta*M - C)/lambda >= E_b`, exactly four parameters can move, and there is a fifth move that changes the problem instead of the parameters.

```
1. RAISE M      more environment, better calibration
2. RAISE eta    the conversion ladder. Enters twice; strongest parameter lever.
3. CUT C        one-for-one, immediate, and the one the prescription reverses.
4. CUT lambda   write it down, date it, tell one person. Cheapest of all.
5. SPLIT E_b    do not raise the fund; lower the wall.
```

The fifth is the engineering move and it rescues cases the other four cannot.

Take the regime at `lambda = 0.06` rather than 0.025: `L^Q = 0.108`, so `Lambda_inf = 1.800` against a barrier of `E_b = 2.400`. The departure is infeasible. No amount of waiting helps; the fund asymptotes at 1.800 and stops.

Now decompose the transition into three legs with rest basins between them — three sub-transitions, each ending at a configuration where the subject can stop, settle and re-accumulate:

```
leg 1: E_b1 = 0.9      feasible (1.800 > 0.9)      11.20 weeks to fund
leg 2: E_b2 = 0.8      feasible (1.800 > 0.8)       9.50 weeks
leg 3: E_b3 = 0.7      feasible (1.800 > 0.7)       7.96 weeks
-------------------------------------------------------------------
total: 2.4 (unchanged)                             28.66 weeks of accumulation
```

The total cost is identical — 2.4 Q-units either way. What changed is the peak fund requirement, which fell from 2.400 to 0.900, below the ceiling. An infeasible departure became feasible with no parameter changing, purely by inserting two places to stand.

The generalisation is the most transferable result in this half of the paper: **when a barrier exceeds `L/lambda`, do not exhort the walker and do not raise the fund; split the barrier and build the intermediate basins.** In organisational form this is the staged migration, the pilot, the secondment and the reversible trial; in personal form it is the evening course before the career change, the sublet before the move, the two days a week before the four. Every one of them is a manufactured rest basin performing the same arithmetic. The precedent in the built environment is the same move: a sequence of occupiable intermediate places, rather than a single long undifferentiated passage, is what makes a route get used [Alexander et al., 1977, Gehl, 1971].

The condition under which splitting works is precise and must be checked. **There must be a real resting configuration at each intermediate point — a place where `grad C = 0` and the subject can stop indefinitely without drawing down.** A stage that is merely a point partway up the slope is not a rest basin; it is a pause on the ridge, where Section 4.2's table shows the launch budget is at its smallest. **Half a career change is the most expensive place to stand, not the cheapest.**

22.4.9 4.9 The over-framed launch budget

`L` is computed on `M`, the moment as lived, and not on the veridical part `V`. That is deliberate — the plan for elsewhere is made of what was experienced, not of what was there — but the consequence must be stated. Take an subject with environment `E = 1.42`, `eta = 0.10`, `C = 0.034`, `lambda = 0.025` and a calibration ratio `r = 1.60`:

```
M = min(1.60 * 1.42, ln 12) = 2.272 nats        (the cap at ln 12 = 2.4849 is not reached)
V = 1.42 * min(1.60, 1)     = 1.420 nats

L^M on M = 2.272 - 0.340 = 1.932        Lambda_inf = 7.728
L^M on V = 1.420 - 0.340 = 1.080        Lambda_inf = 4.320

inflation factor = 1.932 / 1.080 = 1.7889
```

The over-framed subject has a launch budget 79 per cent larger than their environment supports, and believes they can cross a barrier of 7.73 when the earned figure is 4.32. Their plans are funded partly by manufactured moment; and because `r > 1` also means they see a smaller barrier than there is, the `E_b` estimate is low at the same time as the `Lambda_inf` estimate is high. **Both errors point the same way, so the over-framed subject systematically launches into transitions they cannot complete, and is systematically surprised.**

The correction is a calibration audit, and the early-warning instrument is the return difference `Delta_R = |ln(M'/M)|` developed in a companion paper of this series: a plan whose probes produce no return difference is a plan the world has not yet answered, and committing a fund to it is committing on an unverified environment. **The moment equation and the calibration ratio `r` have no external support in this corpus's bibliography and are carried as framework constructions**, which Section 12 repeats.

22.4.10 4.10 The perpetual dreamer

The opposite pathology, and it is common in exactly the populations that read procedures. `L` large; `lambda` small, because the subject is disciplined, keeps notes and tells people; `Lambda` accumulates to `Lambda_inf` and stays there. And nothing is ever spent, because no barrier is ever selected.

```
Lambda_inf = 4.320 Q-units, held indefinitely, against no committed E_b.
```

A launch budget that is never spent is indistinguishable from a launch budget that never existed, except in its cost. The cost is real: maintaining an unexecuted plan is stored order held against no returning input, which is a maintenance draw in the storage ledger, and the carry runs at full rate for every open loop. The failure is not insufficient budget. It is a missing decision, and the model locates it exactly.

22.4.11 4.11 The spend rule

```
PROTOCOL L-1 - WHEN TO SPEND THE FUND

S1. Select ONE target configuration. Write it down with its date.
S2. Estimate E_b by Section 4.6 -- the path integral of excess cost at the intended
    transit speed -- and NOT by asking how hard the change feels.
S3. Compute Lambda_inf = (eta M - C)/lambda from measured parameters.
S4. If Lambda_inf <= E_b: do not begin. Either move a parameter (Section 4.8, moves 1-4)
    or split the barrier (move 5). Beginning a transit the fund cannot complete strands
    the subject on the ridge, which Section 4.2 identifies as the configuration of
    minimum launch budget -- the worst place in the field to run out.
S5. If Lambda_inf > E_b: wait t* = ln(1 - E_b/Lambda_inf)/ln(1 - lambda) periods, then
    depart. Departing before t* is departing underfunded.
S6. Probe first: a minimal reversible test of the destination basin, with its cost
    logged beside its finding. A destination estimated at r > 1 is a destination partly
    manufactured (Section 4.9), and the probe is what checks it.
S7. Recompute after each leg of a split barrier. The parameters change in transit; C in
    particular rises, which lowers L, which lowers the fund available for the next leg.
```

Step S4 is the one that gets skipped and it is the one that produces the stranding. The model is clear about the consequence: the ridge is where L is smallest, so an subject who runs out of fund mid-crossing has the least capacity to plan their way out of exactly the place they most need to.

22.5 5 The gradient and the desire line

22.5.1 5.1 The desire line as an instrument

Landscape architects have a name for the path worn across a lawn where the paving does not go. It is a *desire line*, and the term is exactly right here because it is a measurement of preference that nobody administered. Three properties make it a better instrument than self-report, and each has a formal counterpart below.

It has a large sample. The path is worn by hundreds of crossings, while a self-report is a single observation taken by a reactive instrument. Section 5.4 makes the sample-size consequence quantitative and it is severe.

It is not falsifiable by intention. People walk the cheap way whether or not they endorse walking the cheap way. **The desire line records what the cost field is, not what the subject believes about it,** and the two diverge in exactly the cases that matter. The general finding that verbal report does not have access to the process generating the behaviour [Nisbett and Wilson, 1977], and that the report is shaped by how it is elicited [Schwarz, 1999], is the empirical warrant for preferring the tally.

It is a gradient measurement, not a level measurement. The path shows which way is cheaper; it does not show how expensive anything is. That is precisely the identification result of Section 3.6: direction is identified, magnitude is not.

Two lineage notes. The method of sitting down, watching, counting what people do, and letting the counts overrule the plan is Whyte's method for public spaces [Whyte, 1980], with the same logic applied to street and building use by Jacobs [1961] and Gehl [1971]; Protocol G-1 below is that method turned on a single subject's own recurring decisions, and the corpus applies it to rooms in Curzi [2026f]. And the observation that infrastructure becomes visible only on breakdown [Star and Ruhleder, 1996, Star and Griesemer, 1989] explains why a person's own desire lines are invisible to them: **the path already worn is precisely the path no longer noticed being chosen.**

22.5.2 5.2 The configuration space

X must be constructed before anything can be measured, and constructing it badly is the commonest way the survey fails. A usable personal configuration space is a small set of **coordinates**, each a recurring decision point with an enumerable set of options, observed at a stated frequency.

Coordinate	Options	Observation rule
where the first working hour is spent	desk / kitchen / cafe / office	recorded at the start of the hour
what is opened first	mail / task file / chat / a walk	recorded at first action
which chair	A / B / C / D	recorded on sitting
how the day ends	shutdown ritual / drift / hard stop / no boundary	recorded at last work action

Three rules govern a legitimate coordinate.

Rule 1: the options must be exhaustive and mutually exclusive. Every occasion must be classifiable into exactly one, including an explicit "none of these" where needed. A coordinate with an escape hatch that absorbs the interesting cases is not a coordinate.

Rule 2: the coordinate must be genuinely available. If a location is unreachable on three days a week, it is not an option on those days and including it corrupts the count. **Availability must be recorded per occasion and the tally must run over available occasions only.** This is the single commonest source of spurious cost in a desire-line survey: an option rarely chosen because it was rarely offered.

Rule 3: the coordinate must be observed, not remembered. Recorded at the moment, in a fixed format, by a rule stated in advance. Reconstruction at the end of the week measures memory, which is itself stored order decaying at `delta`, and not choice.

Note the demotion this forces. These coordinates are **categorical**: there is no metric on the set {desk, kitchen, cafe, office}, so `grad C` does not exist over them and the honest object is a difference table, not a gradient. **The analysis demotes accordingly: Section 5.3 builds difference tables on categorical coordinates and reserves genuine gradients for the ordered coordinate of Section 5.8.** The general reason to distrust an assumed metric on psychological objects is independent and strong: psychological similarity is asymmetric and violates the triangle inequality [Tversky, 1977].

22.5.3 5.3 Measuring the field without measuring the cost

Cost cannot be measured directly. Choices can be counted. The logit recovery of Section 3.5 turns the second into the first, up to a scale.

```
PROTOCOL G-1 - THE THIRTY-DAY TRACE

1. Declare 3 to 5 coordinates by the rules of Section 5.2. Write them down. Freeze them.
2. For each coordinate, declare the reference option -- normally the most frequent one,
   chosen AFTER a 3-day pilot and then frozen, because a reference chosen from the final
   data is a reference chosen to flatter the story.
3. Record the chosen option and the available option set at each occasion, for 30 days.
4. Tally n_x per option, over available occasions only.
5. Compute, in theta units:      d_x = C_x - C_ref = -ln( n_x / n_ref )
6. Compute:                      se(d_x) = sqrt( 1/n_x + 1/n_ref )
7. Report d_x +/- 1.96*se as a 95 per cent interval. Any interval containing zero is a
   coordinate direction that has NOT been measured, whatever the subject's intuition
   says.
8. Run the IIA check: recompute the top two options' ratio with the rarest option
   dropped. If it moves materially, the logit form is wrong for this coordinate.
```

Steps 1 and 2 are the analytic-flexibility controls and they are not optional. Choosing a reference option after seeing the data is a forking path, and the resulting estimate is not the estimate it appears to be [Gelman and Loken, 2013, Simmons et al., 2011]; freezing the instrument before the window is the standard remedy [Nosek et al., 2018]. **Step 7 is the step that converts confession into data**, and it is the step that most often contradicts the subject.

22.5.4 5.4 What the survey can and cannot resolve

Take the first coordinate of the worked survey: 30 days, four options, counts 14 / 9 / 5 / 2.

Option	n	share	d = C_x − C_desk (theta)	se	95% interval	resolved?
desk (home)	14	0.467	0.0000	—	reference	—
kitchen table	9	0.300	0.4418	0.4272	[-0.3956, 1.2792]	**no**
the cafe	5	0.167	1.0296	0.5210	[0.0085, 2.0508]	yes, barely
the office	2	0.067	1.9459	0.7559	[0.4643, 3.4275]	yes

After thirty days of faithful recording, this subject cannot distinguish the cost of the kitchen table from the cost of the desk. The point estimate says the kitchen is 0.44 theta dearer; the interval says it might be 0.40 theta cheaper. The subject, asked, will be quite certain, and the data does not support the certainty.

How long would it take? Holding the observed proportions:

```
to resolve d = 0.4418 at 95 per cent:   se must be <= 0.4418/1.96 = 0.2254
se(n) = sqrt( 1/(0.300 n) + 1/(0.467 n) ) = sqrt( 5.4762 / n )
n >= 108 days
```

coordinate direction	d (theta)	days needed
kitchen vs desk	0.4418	**108**
task file vs mail	0.7538	38
cafe vs desk	1.0296	30

One hundred and eight days to resolve a difference the subject was sure of on day one. The gradient can be measured and it is cheap to measure, but **the small differences — which are exactly the ones a person deliberates over — are the expensive ones to establish.**

And note the structure of the standard error: `se = sqrt(1/n_a + 1/n_b)` is dominated by the *smaller* count. **The rarely taken option is the badly measured one, and the rarely taken options are the interesting ones.** A survey of one's own life measures habits precisely and exceptions not at all.

22.5.5 5.5 What is identifiable: the ratios of differences

`theta` is unknown, so every `d` above is in theta units and no absolute statement can be made. But `theta` cancels in a ratio:

```
( C_office − C_desk ) / ( C_kitchen − C_desk ) = 1.9459 / 0.4418 = 4.4042
( C_cafe   − C_desk ) / ( C_kitchen − C_desk ) = 1.0296 / 0.4418 = 2.3303
( C_office − C_desk ) / ( C_cafe    − C_desk ) = 1.9459 / 1.0296 = 1.8899
```

The office is 4.4042 times as costly a departure from the desk as the kitchen table is. That statement is unit-free, fully identified, and testable. It is also the form in which every comparative claim in this half of the paper is made, because it is the only form the data supports. The general principle is measurement-theoretic rather than statistical: the admissible statements are those invariant under the transformations the scale allows, and with an unknown positive unit the ratio of two differences is invariant while a difference alone is not [Narens, 1996, Stevens, 1946].

The rule, for use: **report ratios of differences, never differences alone, and never a level.** A survey reporting "the office costs 1.95" has reported a number whose unit nobody knows. A survey reporting "the office costs 4.4 kitchen-tables" has reported a measurement.

The one experiment that would pin `theta`: **impose a cost of known magnitude on one option and measure the shift in its share** — a five-minute walk added to one route, a small explicit charge, a fixed delay. From the shift, `theta` follows and every difference in the survey becomes absolute. **That experiment has not been run here, and its absence is why every number in this section is in theta units.**

22.5.6 5.6 Flat directions: where the subject is free, and where design has leverage

Take the third coordinate, which chair, with counts 8 / 7 / 8 / 7.

Option	`n`	`d` (theta)	`se`	95% interval
chair A	8	0.0000	—	reference
chair B	7	0.1335	0.5175	[-0.8809, 1.1479]
chair C	8	0.0000	0.5000	[-0.9800, 0.9800]
chair D	7	0.1335	0.5175	[-0.8809, 1.1479]

This coordinate is flat, and two readings follow.

Reading one: a flat direction is a free direction. The subject can move along this coordinate at no cost. In the formalism, `grad C = 0` along a degenerate direction of the Hessian, so motion is unopposed. **Freedom, in this model, is a flat direction in the cost field, and it is measurable.**

Reading two, which is the engineering one: a flat direction is where a small regrade moves everything. Apply the same subsidy — 0.5 theta of cost removed — to a steep coordinate and to a flat one:

```
STEP: the office, currently 6.67% of occasions
    relative weights: desk 1.0000, kitchen 0.642857, cafe 0.357143, office 0.142857
    after a 0.5 theta subsidy on the office: 0.142857 * e^0.5 = 0.235531
    new share = 0.235531 / 2.235531 = 10.536%          gain: +3.869 points

FLAT: chair B, currently 23.33% of occasions
    after the same 0.5 theta subsidy: share 23.33% -> 33.41%     gain: +10.079 points

LEVERAGE RATIO = 10.079 / 3.869 = 2.605
```

The same subsidy buys 2.6 times as much movement in the flat coordinate as in the steep one. This is the arithmetic behind a familiar piece of design practice: **do not fight the steep direction; find the flat one and tilt it.** It is why moving a chair three metres changes a room's traffic while exhorting people to use a distant space does not, and the built-environment form of the same result is priced in Curzi [2026f], whose entrance-sequence case cut interruptions from 22 to 6 per person-day at a seven-minute switch cost, moved the cost of switching from 0.321 to 0.087, raised headroom from +0.0642 to +0.2955 and the storage fixed point from 5.185 to 9.041, for 7.6 per cent of net floor area.

The general statement is worth having: **cost differences are what people notice; cost gradients are what people move along; and the coordinates where they notice nothing are exactly the coordinates where a designer can move them most.**

22.5.7 5.7 Data, not confession

The distinction matters enough to state as a table, because the failure is a category error rather than an inaccuracy.

	The confession	The desire line
what it measures	a belief about the cost field	the cost field
sample size	1	one per occasion
reactivity	high; the act of asking changes the answer	low; the act of tallying is nearly inert
available on	any topic	only recurring, enumerable decisions
identifies	nothing	cost differences up to `theta`
failure mode	sincere and wrong	precise and narrow

The desire line is worse at almost everything except being true. It covers only recurring decisions, it needs a hundred days for a small effect, it yields no absolute magnitudes, and it says nothing about anything the subject does not do repeatedly. Within its narrow domain it is not arguable with.

22.5.8 5.8 Descent dynamics on an ordered coordinate

Where a coordinate *is* ordered — hours per week in a role, distance from a place, fraction of work in a new domain — the difference table becomes a genuine gradient and the descent dynamics become computable. The worked survey's ordered coordinate runs from `x = 0` to `x = 10` and carries the cost field shared with Section 4:

```
C(x) = 0.090 − 0.0568 exp(−(x−2)^2/0.8) − 0.064 exp(−(x−7)^2/1.2) + 0.0004 x
```

Its structure, computed:

Feature	`x`	`C(x)` (Q-units/wk)	`C''(x)`
basin A floor	1.9970	0.033999	+0.141995
the ridge	4.3350	0.091500	-0.005076
basin B floor	6.9960	0.028799	+0.106662

With `L = max |C''| = 0.142`, the standard first-order conditions give `alpha < 2/L = 14.08` for local non-divergence and `alpha = 1/L = 7.04` for the fastest guaranteed monotone decrease [Strogatz, 2015, Åström and Murray, 2008]. Descent from `x_0 = 3.0` at `alpha = 1/L`:

step	`x`	`C(x)`	`−dC/dx`
0	3.00000	0.074926	-0.0410830
1	2.71068	0.060874	-0.0540756
2	2.32987	0.041355	-0.0412845
3	2.03913	0.034124	-0.0059463
4	1.99726	0.033999	-0.0000106
5	1.99718	0.033999	0.0000000

Five steps to arrive, the last two indistinguishable. The local contraction factor `|1 − alpha*mu|` is 0.00000 at basin A, because the curvature there happens to equal the global Lipschitz constant, and 0.24886 at basin B. **The source states that basin B therefore takes "about**

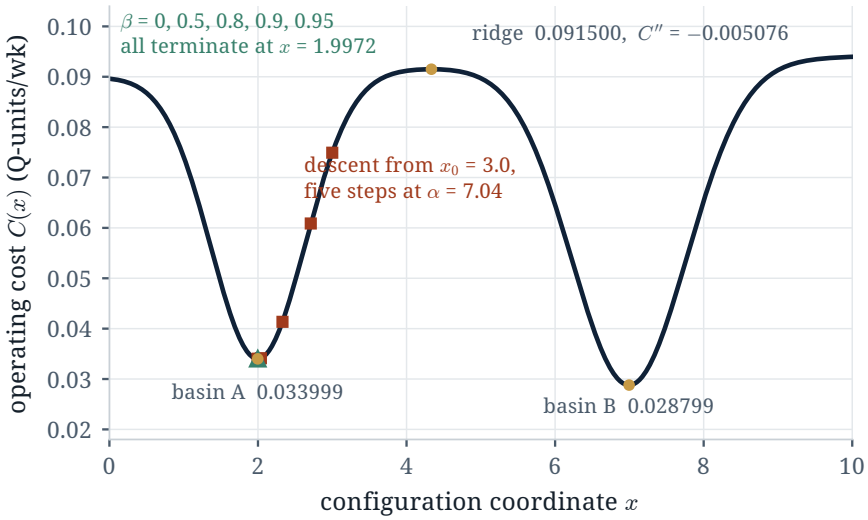


Figure 22.1: **Basin A is the correct output of a local method, not a failure of will, and momentum does not change it:** the shared cost field has basin A at `x = 1.9970`, `C = 0.033999`, a ridge at `x = 4.3350`, `C = 0.091500` with `C'' = −0.005076`, and a deeper basin B at `x = 6.9960`, `C = 0.028799`. Descent from `x_0 = 3.0` at `alpha = 1/L = 7.04` reaches basin A in five steps, and momentum at `beta` = 0, 0.5, 0.8, 0.9 and 0.95 terminates at `x = 1.9972` at every setting. Modelled from the field equation and parameters of Section 5.8; generated by `figs/src/CCP-022_fig03.py`.

2.7 steps per decade of remaining error"; from that same contraction factor the figure is 1.66 steps per decade, since the logarithm of 0.1 divided by the logarithm of 0.24886 is 1.6555. The discrepancy is reported rather than silently corrected, and nothing else in this paper depends on it. What matters here is that arrival is fast once the walker is in a basin, so the interesting question is never how long descent takes but which basin it started in. The geometric character of that convergence is the contraction result of Curzi [2026].

The basin boundary sits at `x = 4.33485`, which is the ridge to five decimals. Descent from 4.30 lands in basin A; descent from 4.40 lands in basin B. The two starting points are 0.1 configuration-units apart and their destinations are five units apart.

And the plateau is slow. Starting at `x = 4.30`, just 0.035 inside basin A, the walker is still at `x = 4.128` after fifty steps, because the gradient magnitude near the ridge is of order `10^-5`. An subject near a ridge appears to be deliberating for a long time. They are not deliberating; they are descending a plateau, and the model says exactly how slowly.

22.5.9 5.9 What the gradient cannot tell you

Four things, each a real limitation with a real consequence.

It cannot tell you the minimum is global. Descent from every start in `[0, 4.334]` lands in basin A at `C = 0.034`. Basin B is deeper at `C = 0.0288`, a 15.3 per cent lower operating cost, and no amount of descending finds it from those starts, ever. **Basin A is not a mistake and not a failure of will. It is the correct output of a local method applied to that starting point.** This is the myopia of learning stated as geometry [Levinthal and March, 1993], and the exploration side of the exploration-exploitation trade-off is what is missing rather than effort [March, 1991].

It cannot see over the ridge. The gradient at `x = 3` points toward `x = 2` and contains no information whatever about the existence of basin B. **The only instruments that see basin B are survey and report** — measuring a region one is not in, or hearing from someone who is. The structural role of non-redundant contacts in supplying exactly that information is well measured [Granovetter, 1973, Burt, 1992, 2004], and it is why a person with one framing and

no outside contact is not merely unlucky but *informationally* confined.

It cannot tell you the destination is worth arriving at. `C` is cost, not value. Descent optimises the operating draw and nothing else. A basin can be cheap and empty; a configuration in which nothing is required of the occupant is a global minimum of `C` and may be a catastrophe on every other ledger. **The stored order `Q` and the surplus `F` are where worth lives, and the gradient is blind to both** [Curzi, 2026a]. An subject who optimises only `C` converges on the cheapest available life, and the storage law's fixed point shows that a cheap life with a thin moment still ends in a small self.

It cannot price a barrier it is not adjacent to. `E_b` for the A-to-B crossing is 2.3987 Q-units at a 65.5-week transit (Section 4.6). That number required integrating the field along the whole path. **It is not obtainable from any local measurement at `x = 2`,** and an subject who has only descended has no access to it.

These four limitations are why the launch budget is needed, and the join is exact: the gradient supplies the direction of cheapest motion and the launch budget supplies the ability to move against it. Descent alone converges on the nearest available floor. Descent plus a funded departure reaches the best floor that was surveyed.

22.5.10 5.10 Temperature and momentum: two candidate escapes, one of them refused

Two standard escapes from a local minimum exist in the optimisation literature, and it is worth reporting what each does to this landscape, including where it does nothing. Both are **analogies** in the sense of Section 3.7: imported as candidate mechanisms and tested, with no claim that either physical process occurs.

Temperature. Under the logit model choice is already stochastic: the probability of taking a costlier option is `1/(1 + exp(d/theta))`.

`d` (theta units)	`P`(costlier option)
0.0000	0.50000
0.1335	0.46667
0.4418	0.39131
1.0296	0.26316
1.9459	0.12500
2.8332	0.05556

So a walker at intermediate temperature does occasionally step uphill and, given enough occasions, will cross a shallow ridge with no plan at all. This is the annealing heuristic of combinatorial optimisation arriving as a description of ordinary variability. **It is borrowed form: no claim is made that people anneal.** What is worth taking from it is the design consequence: **raising `theta` raises exploration and lowers exploitation, so a period of deliberately loosened routine is a real search strategy with a real cost, and the cost is that everything gets slightly worse while it runs.** That is the exploration trade-off in its exact form [March, 1991].

Momentum, and the negative result. Adding a momentum term `v_{t+1} = beta*v_t − alpha*grad C`, `x_{t+1} = x_t + v_{t+1}`, gives a walker that carries direction across small adverse gradients. Run on this landscape from `x_0 = 3.0`:

`beta`	effective step multiplier	final `x`	basin
0.00	1.00	1.9972	A
0.50	2.00	1.9972	A
0.80	5.00	1.9972	A

beta	effective step multiplier	final x	basin
0.90	10.00	1.9972	A
0.95	20.00	1.9972	A

Momentum does not escape this landscape at any setting tested, and the reason is instructive: the barrier is a plateau, not a spike. The ridge at `x = 4.335` has `C'' = −0.005076` and is nearly flat over more than two configuration-units. Momentum accumulated on a one-unit descent cannot carry a walker across a two-unit flat rise. **Momentum defeats narrow barriers and is useless against wide shallow ones, and most barriers of practical interest — an income gap, a retraining period, a slow renegotiation — are wide and shallow.**

That is a measured negative result on a constructed field and it is reported at the same prominence a positive result would receive. It also sharpens the claim of Section 4: **against a wide shallow barrier there is no substitute for an accumulated fund, because there is no dynamical trick that crosses it.**

A bookkeeping note, since the framework's rule is that every formula in its formulary carries its seven gates. The momentum recursion is not entered into the formulary and therefore carries no validation table: it was introduced as a candidate mechanism, tested on a measured field, and failed. Recording a failed candidate is not the same as adopting one, and the distinction is worth keeping sharp, because a formulary that admits every formula anyone tried is a formulary nobody can use. What survives is one empirical statement with no free parameters in it: *momentum does not cross wide shallow barriers.*

22.5.11 5.11 Regrading

The permitted intervention is always the same: **do not push the walker; lower the valley.** Alter relative costs so the wanted configuration is the cheap one and leave the choice intact. The construction satisfies the framework's consent constraint by construction — **the subject can still decline, and the bounded worst case is that they look at the lower valley and stay where they are.** A method whose worst case is "nothing changes" is admissible. A method that removes the option is not, at any efficacy.

Four regrades of the shared landscape, computed. `E_b` at a 65.5-week transit; launch fund `Lambda_inf = (eta*M − C(A))/lambda` at `eta = 0.10`, `M = 1.42`, `lambda = 0.025`.

Regrade	`C(A)`	`E_b`	`Lambda_inf`	`t*` (wk)
none	0.03400	2.3987	4.3200	**32.00**
deepen the destination by 0.006	0.03400	2.3523	4.3200	31.06
bridge the ridge by 0.006	0.03361	2.2294	4.3357	28.52
bridge the ridge by 0.020	**0.03269**	**1.8345**	**4.3724**	**21.49**

Deepening the destination is nearly useless: a 21 per cent improvement in the target basin's floor buys a 1.94 per cent reduction in the barrier. Making the destination more attractive does not make the crossing cheaper, because **the barrier is the ridge, not the valley.** Every recruitment brochure, relocation incentive and aspirational campaign that works on the destination is working on the wrong term.

Bridging the ridge is where the leverage is. A 0.020 reduction in the transit region — subsidised retraining, a bridging income, a sublet, keeping the old arrangement running in parallel — cuts the barrier by 23.5 per cent and cuts the wait from 32.00 weeks to 21.49.

And it does something further, which was not designed and fell out of the computation. After the 0.020 bridge the stationary points of the field are:

Feature	x	C(x)
basin A	2.0195	0.032665
ridge	3.5425	0.073902
new basin	**4.3350**	**0.071500**
ridge	5.1780	0.074027
basin B	6.9815	0.028208

The bridge manufactured a rest basin in the middle of the crossing. That is precisely the barrier split of Section 4.8 — here produced by a regrade rather than prescribed by a plan. Priced:

```
leg 1  (A -> new basin):     E_b = 0.8812     accumulation at Lambda_inf = 4.3734:   8.89 weeks
leg 2  (new basin -> B):     E_b = 0.0218     accumulation at Lambda_inf = 2.8200:   0.31 weeks
-------------------------------------------------------------------------------------------
total accumulation                                                                  9.19 weeks
(against 32.00 weeks for the unregraded single crossing)
```

Nine weeks instead of thirty-two, with the total transit time unchanged at about sixty-five weeks. What changed is how much fund had to be accumulated before anything could begin. The rule this yields is worth stating on its own: **a regrade that creates a stationary point in the middle of a crossing is worth far more than a regrade of the same size applied anywhere else, because it converts one barrier into two and the fund requirement scales with the largest, not the sum.**

A gap in the source, reported. The source states the resulting stationary points but never states the bridge's functional form. A Gaussian reduction `C_bridge(x) = C(x) - a*exp(-(x - 4.335)^2 / 2.0)` reproduces all five stated stationary points to four decimal places at `a = 0.020`, and all three at `a = 0.006`, so it is used to draw the figure below; it recovers `E_b = 1.8361` against the source's 1.8345, a 0.09 per cent difference. **The form is a reconstruction, not a quotation, and is labelled as such in the caption.** Two further inconsistencies of the same size sit in this table: `C(A)` after the 0.020 bridge is 0.032665 at the stationary point but is carried as 0.03269 in the regrade table, which propagates to `Lambda_inf` = 4.3734 in the leg table against 4.3724 in the regrade table, and to `t*` = 21.50 against 21.49. The figures are left as the source has them.

22.5.12 5.12 The headroom identity, and what it means to descend

From the headroom gauge `h = eta*ln N - delta*Q - C` [Curzi, 2026b], only the last term depends on configuration. Therefore:

```
THE HEADROOM IDENTITY

    grad_x h(x)  =  - grad_x C(x)  =  g(x)

    DESCENT IN THE COMFORT FIELD IS ASCENT IN HEADROOM, holding eta, N, delta and Q fixed.
```

Spot-checked on the shared field at `eta = 0.10`, `N = 9`, `delta = 0.031729`, `Q = 1.20`:

x	C(x)	h(x)	-dC/dx
1.500	0.049044	0.132603	+0.0515447

x	C(x)	h(x)	−dC/dx
2.000	0.034000	0.147648	-0.0004000
4.335	0.091500	0.090148	+0.0000007
7.000	0.028800	0.152848	-0.0004000

Descending the comfort field is ascending the headroom gauge. The worn path and the framework's master gauge are measuring the same gradient, in the same direction, with the same zeros. **A cheap survey of one's own habits is a measurement of one's own capacity for the new.**

Two consequences, and the second is a diagnostic.

One: the comfort zone is the headroom maximum. Under the restriction, the local minimum of `C` is the local maximum of `h`, so the founding definition [Curzi, 2026] and the master gauge [Curzi, 2026b] are the same statement in two registers, and the platform result of Section 4.2 follows immediately: **the place of rest is the place with the most capacity, which is why it is the place from which departures are planned and funded.**

Two: when the identity fails, it names the culprit. The identity holds only where `eta`, `N`, `delta` and `Q` are configuration-independent, and often they are not. A configuration that puts the subject among people who see differently raises `N`; one that supplies a notebook habit raises `eta`; one that confers a title raises `Q` and therefore the maintenance draw. The corrected form is

```
grad h = eta'(x) ln N + eta/N ∗ N'(x) − delta'(x) Q − delta Q'(x) − grad C
```

If a measured desire line points consistently away from measured headroom improvement, one of those terms is non-zero, and identifying which is a diagnosis rather than a contradiction. The commonest case is a configuration with low `C` and low `N` — the comfortable narrow life — in which descent genuinely reduces cost while `eta∗ln N` shrinks faster. That is visible only if `N` is measured alongside `C`, and the corpus's own count lever is priced and declined net on exactly this arithmetic: a ceiling gain of `eta∗(ln 6 − ln 5) = +0.1914` against five added communication paths costing +0.0714, a 34 per cent haircut on the headline benefit [Curzi, 2026j].

22.5.13 5.13 Switching cost, and why the path back is not the path out

One assumption has been running underneath the measurement sections without being examined: that each occasion is an independent draw from the choice distribution. It is not, and the correction matters enough to be derived rather than mentioned.

Most real configurations carry a **switching cost**: it is dearer to do a thing today if something else was done yesterday — setting up elsewhere, re-establishing context, breaking a running habit, explaining the change to the people around one. The simplest admissible model adds one term:

```
THE SWITCHING COST

    C_eff( x | x_prev )  =  C(x)  +  s ∗ 1[ x != x_prev ]

    s   the switching cost, in the same units as C −− the surcharge for not doing what
        was done last
```

This makes the cost field path-dependent, and a path-dependent field is a field with hysteresis: the cost of the route from A to B is not the cost of the route from B to A, and the configuration an subject settles into depends on where they came from as well as on where

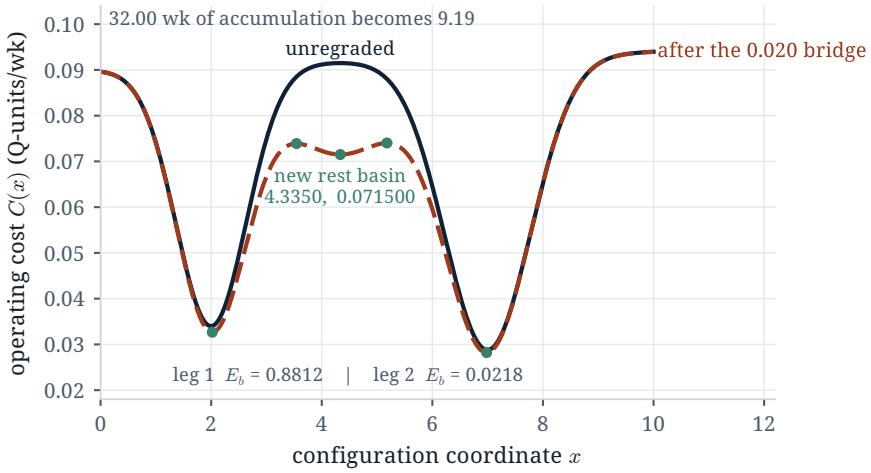


Figure 22.2: **The bridge manufactures a rest basin where the plan would otherwise have to prescribe one:** a 0.020 reduction applied at the ridge turns three stationary points into five, with a new basin at `x = 4.3350`, `C = 0.071500`, splitting a single 2.3987-Q-unit barrier into legs of 0.8812 and 0.0218 and cutting required accumulation from 32.00 weeks to 9.19. Modelled; the source states the resulting stationary points but not the bridge's functional form, and the curve drawn is the reconstruction `C(x) – 0.020*exp(–(x–4.335)^2/2.0)`, which reproduces all five stated stationary points to four decimals. Generated by `figs/src/CCP–022_fig04.py`.

the floors are. That switching costs and accumulated attachment keep participants in an arrangement they would not choose afresh is the exit-cost mechanism in its organisational form [Hirschman, 1970].

The consequence for the instrument is direct and unpleasant. Choices are now a Markov chain rather than independent draws, and the logit recovery, which assumes independence, recovers something that is not the cost field. Computed on a four-option coordinate with true costs `C = (0, 0.20, 0.50, 0.90)` in theta units, taking the chain's stationary distribution and applying Protocol G-1 to it exactly as an investigator would:

`s`	recovered `d_kitchen`	recovered `d_cafe`	recovered `d_office`	recovered ratio office/kitchen
0.0 (truth)	0.2000	0.5000	0.9000	4.5000
0.3	0.2201	0.5442	0.9675	4.3950
0.6	0.2416	0.5927	1.0433	4.3176
1.0	0.2709	0.6609	1.1537	4.2587
1.5	0.3053	0.7446	1.2964	4.2460

Three findings, and the third is why the specification is worded as it is.

One: switching cost inflates every non-modal option's recovered cost. At `s = 0.6` the kitchen table is recovered as 0.2416 when its true cost is 0.2000, a 20.8 per cent overstatement. The instrument has silently charged the subject's habit to the option.

Two: the inflation is largest in absolute terms for the rarest option. The office is inflated by 0.1433 theta and the kitchen by 0.0416. **A switching cost makes unusual choices look more expensive than they are and makes the most unusual ones look worst.** Combined with Section 5.4's finding that rare options are also the badly measured ones, this is a double penalty on exactly the directions a survey is run to investigate.

Three, and it is the good news: the identified ratio is far more robust than the differences. At `s = 0.6` the differences are biased by 15.9 to 20.8 per cent while the ratio is biased by 4.05

per cent; at `s = 1.0` the differences are biased by 28 to 36 per cent and the ratio by 5.36 per cent. **Reporting ratios of differences rather than differences buys roughly a factor of four in robustness to an unmodelled switching cost, and that is on top of the theta-freeness that was the original reason for the rule.** A specification adopted for one good reason turns out to have a second.

The engineering consequence is a small change to the protocol and a large change to interpretation. **Record transitions, not only choices.** One extra column — what was chosen on the previous occasion — turns a marginal tally into a transition matrix, identifies `s` from the log-odds difference between staying and arriving, and lets the recovered field be corrected for it.

And the design consequence connects back to the launch budget. **A switching cost is a barrier that lives between adjacent configurations rather than between distant ones,** which means a coordinate with a high `s` has small barriers everywhere rather than one large barrier somewhere. That is a different landscape and it calls for a different move: not a bridge across a pass, but a reduction in the surcharge itself — the standing setup left in place, the second desk not cleared away, the arrangement that makes tomorrow's switch as cheap as today's continuation. This is the setup ledger seen from the other side, and it is why the corpus's built-environment work caps setup time from arrival to first productive minute at ninety seconds [Curzi, 2026f]: setup is a switching cost, and a switching cost suppresses use of a space in proportion to how rarely it is already used. A laid table gets eaten at and a cupboard does not, and the difference is `s`.

22.6 6 Formula validation: seven gates on five formulas

The corpus applies a fixed seven-gate audit to every formula it admits to its formulary. The gates are: **FVP-1 dimensions** (do the units of every term agree); **FVP-2 scale type** (does the operation performed require an interval or a ratio scale, and is that licensed); **FVP-3 limiting cases** (each parameter to zero and to infinity); **FVP-4 fixed points** (what is invariant); **FVP-5 identifiability** (can the quantity be recovered from obtainable data); **FVP-6 falsification** (state the result that would kill it); **FVP-7 reproducibility** (do the parameters recover across individuals). Verdicts are recorded as stated, including the failures.

The momentum recursion of Section 5.10 has no table here, deliberately: it was tested and failed, and a failed candidate is recorded rather than admitted.

22.6.1 6.1 The launch budget

`FVP TABLE — THE LAUNCH BUDGET  L^Q = max(0, eta M − C),  L^M = max(0, M − C/eta)`

Gate	Test applied	Finding	Verdict
FVP-1 dimensions	Units of each term; both sides agree.	The naive form fails; Section 3.2 resolves it into two consistent registers differing by the factor `eta`. `L^Q` in Q-units per period, `L^M` in nats per period. **A reported `L` without a declared register is inadmissible.**	**PASS after the register split; FAIL as originally written**
FVP-2 scale type	Subtraction requires an interval scale; `max(0,.)` requires a meaningful zero, hence a ratio scale.	`M` is ratio-scaled by the framework's construction. `C` is ratio-scaled only if the cost ledger's exchange rates are, which is a postulate every time. **The rectifier is the demanding part: `max(0, x)` is only meaningful if zero means something, so this formula cannot be demoted to an interval-scale heuristic — it either has a true zero or it is not this formula.**	**CONDITIONAL — depends entirely on the cost ledger's zero**

Gate	Test applied	Finding	Verdict
FVP-3 limiting cases	Each term to 0 and to infinity.	`C -> 0`: `L -> eta M`, the whole deposit is discretionary — the state of a person with nothing to maintain. `C -> eta M`: `L -> 0`, the floor. `C > eta M`: `L = 0` and stays 0 however much larger `C` grows — **the rectifier means the model cannot distinguish "just barely trapped" from "catastrophically trapped", which is a real limitation and is why the cost carry must be tracked separately.** `M -> 0`: `L = 0`. `eta -> 0`: `C/eta -> infinity` and `L^M -> 0` — **no conversion, no launch budget, at any moment size.** `M -> ln N` with `C` fixed: `L^M -> ln N - C/eta`, the ceiling.	**PASS, with the named insensitivity**
FVP-4 fixed points	The invariant.	Two. (i) `eta M = C` is the zero crossing: the break-even life, in which every unit of moment is consumed by the cost of having had it. (ii) Over a configuration space, the stationary points of `L` are exactly the stationary points of `C` and the maxima of `L` are the minima of `C` (Proposition 4.1). **The second is the paper's central structural result.**	**PASS**
FVP-5 identifiability	Recoverable from data?	`L^Q` requires `eta`, `M` and `C`, all of which the framework's survey protocol measures [Curzi, 2026i]. **But `L` is not independently observable — it is a computed difference, and any error in `eta` propagates into `L^M` as `C*d(1/eta)`, which is severe at small `eta`.** At `eta = 0.05` a 20 per cent error in `eta` moves `C/eta` by 0.29 nats, a fifth of a typical `M`.	**CONDITIONAL — `L` is derived and inherits every upstream error, amplified at low `eta`**
FVP-6 falsification	Preregister; freeze; held-out; four baselines including a shuffled-input control. State the killing result.	**The Launch Postulate predicts that the incidence of newly constructed plans for a different configuration is monotone increasing in contemporaneous `L`, and specifically zero when `L = 0`.** Prospective within-subject design: log `M`, `C` and plan-initiation events daily for 90 days; freeze; test on held-out days against a persistence baseline, an `M`-only baseline, an AR(1) baseline and a shuffled-`L` control. **A single well-documented plan constructed during a sustained `L = 0` stretch falsifies the strong form.**	**NOT RUN — protocol stated, and it is cheap**
FVP-7 reproducibility	Cross-individual recovery.	Unattempted. The question is whether the `L = 0` threshold sits at the same place for different people once `eta` and the cost ledger are constructed by a common procedure. **That is the person-item separation problem [Rasch, 1960, Andrich, 1988] and there is no reason to expect it to come free.**	**NOT RUN**

22.6.2 6.2 The launch fund

Gate	Test applied	Finding	Verdict
FVP-1 dimensions	Units agree.	`Lambda` in Q-units, `L` in Q-units per period, `lambda` per period; `L/lambda` is Q-units. **Consistent, and note that `Lambda` is a stock while `L` is a flow — a distinction violated constantly in ordinary talk about motivation** [Sterman, 1989].	**PASS**
FVP-2 scale type	Multiplication by `(1-lambda)` requires a ratio scale.	Inherits `L`'s conditional verdict.	**CONDITIONAL**

Gate	Test applied	Finding	Verdict
FVP-3 limiting cases	Each term to 0 and to infinity.	`lambda -> 0`: `Lambda` grows without bound — the plan that never fades, so any barrier is eventually crossable. **This is the limit that makes written, dated, witnessed plans powerful.** `lambda -> 1`: `Lambda_t = L_t`, no accumulation, only barriers crossable in a single period. `L -> 0`: `Lambda -> 0` geometrically, half-life `ln 2 / (-ln(1-lambda))`.	**PASS**
FVP-4 fixed points	The invariant.	`Lambda_inf = L/lambda`, unique and attracting by the contraction theorem for `lambda > 0` [Curzi, 2026]. **The fixed point is a ceiling on ambition, computable from two numbers.**	**PASS**
FVP-5 identifiability	Recoverable?	`lambda` from a zero-input window on plans: state a plan, add nothing to it, and measure the decline in willingness to act on it. **This is measurable and almost never measured, and it is the parameter with the largest leverage in the paper.**	**PASS in principle; the measurement is awkward and reactive**
FVP-6 falsification	The killing result.	**If plan persistence does not decay geometrically in the absence of reinforcement, the leaky-integrator form is wrong for plans and must be replaced.** A threshold or step-function decay would falsify it, and either is plausible.	**NOT RUN**
FVP-7 reproducibility	Cross-individual.	Unattempted.	**NOT RUN**

22.6.3 6.3 Feasibility and the critical cost

Gate	Test applied	Finding	Verdict
FVP-1 dimensions	Units agree.	`eta M` and `C` in Q-units per period; `lambda E_b` is (per period)(Q-units) = Q-units per period. **Consistent.**	**PASS**
FVP-2 scale type	Ratio scale throughout.	Inherits. `E_b` must be ratio-scaled with a true zero, since a barrier of zero must mean “no barrier”.	**CONDITIONAL**

Gate	Test applied	Finding	Verdict
FVP-3 limiting cases	Each term to 0 and infinity.	`E_b -> 0`: `C_crit -> eta M` and the condition reduces to `L > 0` — **an adjacent move requires only that the subject not be at the floor.** `E_b -> infinity`: `C_crit -> -infinity`, no cost is low enough, which is correct since an infinite barrier is uncrossable. `lambda -> 0`: `C_crit -> eta M`, the barrier stops mattering — **with a perfectly durable plan, any finite barrier falls eventually.**	**PASS**
FVP-4 fixed points	The invariant.	`C = C_crit` is the knife edge: `Lambda_inf = E_b` exactly, so the fund approaches the barrier asymptotically and never crosses it, and `t* -> infinity` as `C -> C_crit` from below. **The boundary case is "forever nearly ready", which is a recognisable state and is a point in the model's favour.**	**PASS**
FVP-5 identifiability	Recoverable?	All four inputs are separately measurable. **`E_b` is the weak one: it is normally estimated by the subject, who is exactly the party with a stake in the estimate, and Section 5 gives the only construction here that does not rely on introspection.**	**CONDITIONAL on an independent `E_b`**
FVP-6 falsification	The killing result.	**Prediction: across a population of would-be departures, the observed departure rate falls to zero as `C` crosses `eta M - lambda E_b`, at a threshold predicted from parameters measured beforehand.** If departures occur at the same rate above and below the computed threshold, `C_crit` is fictitious. Requires a cohort and is the most expensive test here.	**NOT RUN**
FVP-7 reproducibility	Cross-individual.	Unattempted, and the threshold test above is inherently cross-individual, so FVP-6 and FVP-7 stand or fall together here.	**NOT RUN**

22.6.4 6.4 The barrier integral

```
FVP TABLE — THE BARRIER      E_b = (1/v) * INTEGRAL over the path of max(0, C(x) - C(A)) dx
```

Gate	Test applied	Finding	Verdict
FVP-1 dimensions	Units agree.	`C` in Q-units per period; `dx` in configuration-units; the integral is Q-unit configuration-units per period; dividing by `v` in configuration-units per period gives Q-units. **Consistent, and note that `E_b` is a stock while `C` is a flow — the integral is what converts one to the other.**	**PASS**
FVP-2 scale type	Integration requires an interval scale on `C` and a metric on the path.	`C` differences suffice, since only excess over `C(A)` is integrated. **The metric on the path is the demanding requirement and it excludes categorical transitions: where there is no ordered coordinate between A and B there is no path integral, and `E_b` must be estimated by ledger instead** — the same demotion Section 5.2 performs.	**CONDITIONAL — requires an ordered coordinate, else demote to a ledger estimate**
FVP-3 limiting cases	Each term to 0 and to infinity.	`C(x) <= C(A)` everywhere on the path: `E_b = 0`, a downhill move with no barrier, and feasibility reduces to `L > 0`. `v -> 0`: `E_b -> infinity`; an infinitely slow crossing costs infinitely much, which is correct and is why indefinite hesitation is not free. **`v -> infinity`: `E_b -> 0`, which is FALSE and is the model's known defect, stated in Section 4.6 and fenced by SP-L12 rather than fixed.**	**PARTIAL — one limit is known to be wrong and is fenced rather than fixed**
FVP-4 fixed points	The invariant.	`E_b` is minimised over paths by the route whose maximum excess is lowest — the pass — which is why the regrades of Section 5.11 are evaluated at the ridge and not at the destination. **The invariant is the pass, and it is what a barrier IS.**	**PASS**
FVP-5 identifiability	Recoverable?	Requires `C(x)` over the whole path, and **the subject has by definition not occupied most of it.** The worked case reconstructed the unvisited region by ledger, which is an inference and not a measurement. **This is the single largest weakness in the paper and is recorded again in Section 12.** An independent ledger estimate agreeing with the integral (Section 9.3) is the only check available and it is weak, since both may share the same reconstruction.	**WEAK — the integrand is unmeasured over the region that matters**

Gate	Test applied	Finding	Verdict
FVP-6 falsification	The killing result.	**Prediction: realised transit cost, ledgered after the fact, matches the pre-computed `E_b` within the stated interval.** Cheap to test retrospectively on any completed transition by anyone who computed `E_b` beforehand. **If realised cost is uncorrelated with the predicted barrier across a set of transitions, the construction is fictitious.**	**NOT RUN — but it is the easiest test here to run**
FVP-7 reproducibility	Cross-individual.	Unattempted. **`E_b` is jointly a property of the person and the transition, and separating those is the person-item separation problem in its hardest form** [Rasch, 1960].	**NOT RUN**

22.6.5 6.5 The gradient, the discrete descent, and the logit recovery

Gate	Test applied	Finding	Verdict
FVP-1 dimensions	Units of each term.	`C` in Q-units per period; `x` in whatever units the configuration coordinate carries; `grad C` in Q-units per period per configuration-unit. A step size `alpha` therefore carries configuration-units squared per Q-unit per period — **which is a sign that `alpha` is not a free constant but a scaling artefact of how the coordinate was defined.** Declare the coordinate scale or `alpha` is meaningless.	**PASS, with a mandatory coordinate declaration**
FVP-2 scale type	Differentiation requires at least an interval scale on `C` and a metric on `X`.	**`X` is the harder problem: most personal configuration coordinates are categorical, and where `X` is a finite unordered set `grad C` does not exist and must be replaced by the difference table of Section 5.3** — a demotion performed explicitly rather than papered over [Stevens, 1946, Tversky, 1977].	**CONDITIONAL — demoted to a difference table on categorical coordinates**
FVP-3 limiting cases	Each term to 0 and to infinity.	`grad C -> 0`: no preferred direction — a stationary point, which is a rest basin, a ridge or a plateau, distinguished only by the second derivative. `|grad C| -> infinity`: an infinitely steep wall, so the model predicts instantaneous motion, which is an artefact and marks the coordinate as badly scaled. `C` constant on a subspace: **the flat direction, where the subject is genuinely indifferent and where Section 5.6 shows design has its greatest leverage.**	**PASS**
FVP-4 fixed points	The invariant.	`grad C(x*) = 0`, classified by the Hessian: positive definite a rest basin, negative definite a ridge, indefinite a saddle, singular a plateau [Strogatz, 2015]. **The founding definition's claim that an subject is "already there" is the statement that an subject observed at rest is at a stationary point of their own cost field, by construction and not by luck** [Curzi, 2026].	**PASS**

Gate	Test applied	Finding	Verdict
FVP-5 identifiability	Recoverable from obtainable data?	**`C` itself is NOT recoverable. Only differences are, and only up to the scale `theta`.** So `grad C` is identified up to a positive scalar, which is sufficient for direction and insufficient for magnitude. **Ratios of cost differences are fully identified because `theta` cancels, and those are the quantities reported.**	**PARTIAL — direction identified, magnitude only up to `theta`**
FVP-6 falsification	Preregister; freeze; held-out; four baselines including a shuffled-input control.	Fit the difference table on days 1-20; freeze; predict the choice distribution on days 21-30; beat a uniform baseline, a most-frequent-option baseline, a first-order Markov baseline and a **day-shuffled control**. **If the frozen cost table does not beat the shuffled control on held-out days, the subject has no stable comfort field over the surveyed coordinates and the Gradient Postulate fails for them.**	**RUNNABLE — run in Section 9.8, and passed on two coordinates**
FVP-7 reproducibility	Same operationalisation, different people.	Unattempted. **`theta` is expected to be person-specific and the coordinates certainly are.** The one thing that might transfer is the *shape* of the finding — that most subjects have a few steep coordinates and many flat ones — and even that is untested.	**NOT RUN**

Discrete descent, separately. For `C` with an `L`-Lipschitz gradient, `alpha < 2/L` guarantees the iteration does not diverge locally, `alpha = 1/L` gives the fastest guaranteed monotone decrease, and near a minimum with curvature `mu` the error contracts by `|1 − alpha*mu|` per step. Three gates differ from the table above. **FVP-3:** `alpha -> 0` halts motion; `alpha > 2/L` makes `|1 − alpha*mu| > 1` so the iteration oscillates with growing amplitude *locally* — **in a bounded landscape this presents not as blow-up but as wandering, which is a genuinely different observable and is named in FM-G5. FVP-4:** any `x` with `grad C(x) = 0` is a fixed point, and only the positive-definite ones are attracting, **which is why an subject is found at rest in basins and not on ridges** — so finding someone apparently at rest on a ridge is a strong signal that something other than the gradient is holding them there. **FVP-5: FAIL as written.** `alpha` and the `theta` scale of `C` are not separately identified; only the product governs the observed step, so two subjects with different responsiveness and different cost scales produce identical traces. **Report the product or neither.**

The logit recovery, separately. FVP-1: `ln(n_a/n_b)` is dimensionless, `theta` carries `C`'s units, the product is a cost difference. **Consistent. FVP-2 is the crux and it is good news.** The recovery yields cost **differences** on an interval scale with unknown unit `theta`; **ratios of differences are ratio-scaled and fully identified, because `theta` cancels** [Narens, 1996]. So comparative claims are available without ever pinning `theta`, and absolute claims are not. **FVP-3:** `n_a -> 0` makes both the estimate and its standard error diverge — **a never-chosen option is an unmeasured option, and the model says so loudly**; `theta -> 0` gives deterministic choice and `theta -> infinity` gives uniform choice, and in both limits nothing is recoverable, so **recovery requires intermediate temperature, which is a real and often-violated precondition. FVP-4:** `n_a = n_b` gives `C_a = C_b` — the flat coordinate. **FVP-5:** differences yes, up to `theta`; **`theta` itself requires an external anchor and that experiment has not been run. FVP-6:** independence of irrelevant alternatives — **adding a fourth option must not change the ratio of the first two.** This is the classic test of the Luce form, it is known to fail in many choice domains, and near-substitutes are the classic failure case [Tversky, 1977, Brock and Durlauf, 2001]. **FVP-7:** the procedure is fully portable, since it needs a tally sheet; whether `theta` is a stable personal parameter is unknown.

22.6.6 6.6 The headroom identity and the switching cost

The headroom identity. FVP-1: both sides in Q-units per period per configuration-unit. **PASS. FVP-3:** the identity is exact only where `eta`, `N`, `delta` and `Q` do not depend on `x`, **and that is a**

substantive restriction: a configuration that changes what can be distinguished changes `N`, one that changes what must be maintained changes `Q`, one that changes how experience is captured changes `eta`. **PASS on the stated restriction; FAILS where the configuration moves the other parameters. FVP-4:** `grad C = 0` if and only if `grad h = 0` — **the comfort zone and the headroom maximum are the same point, under the restriction**, which is the identity's payoff. **FVP-5:** the identity requires no new measurement, being algebraic; **its content is that a desire-line survey, which is cheap, measures the same gradient as a headroom survey, which is expensive. FVP-6:** if a measured desire line points consistently away from measured headroom improvement, the restriction is violated and one of `eta`, `N`, `delta`, `Q` is configuration-dependent — **a diagnosis rather than a contradiction, and it localises which parameter is moving.**

The switching cost. FVP-1: `s` in the units of `C`, the indicator dimensionless; **consistent**, with the caveat that `s` is a per-occasion surcharge and not a rate, so the period must be the occasion and not the week. **FVP-2:** inherits the logit verdict; **the binary indicator is the crude part**, since a real switching cost is graded by how different the configurations are, and the metric refinement `s*d(x, x_prev)` requires a metric on `X` that categorical coordinates do not have. **CONDITIONAL — the binary form is a lower-resolution model, badged. FVP-3:** `s -> 0` returns independent draws and Protocol G-1 recovers the true field; `s -> infinity` makes the current configuration absorbing, every recovered cost diverges, and **the survey returns "you always do this", which is true and useless. FVP-4:** the stationary distribution of the induced Markov chain, which exists and is unique for finite `s`; **`s` does not change which option is cheapest, it changes how much of the time the subject is found there. FVP-5: PASS, but only if transitions are recorded.** `s` is identified from transition counts rather than marginal counts: compare `P(x_t = a | x_{t-1} = a)` with `P(x_t = a | x_{t-1} != a)` and their log-odds difference is `s` in theta units. **Protocol G-1 as originally written throws away exactly the data that would identify it, and SP-G16 fixes that defect. FVP-6:** if the first-order Markov baseline does not beat the independent-draw model on held-out days, `s = 0` for that coordinate and the term should be dropped. **FVP-7:** unattempted; `s` is plainly both person- and coordinate-specific, and **the qualitative claim that `s > 0` almost everywhere is close to untestable as stated and is not advanced as a finding.**

22.7 7 Design specifications

22.7.1 7.1 From the launch budget

SP-L1 Every reported `L` shall declare its register — `L^Q` in Q-units per period or `L^M` in nats per period — and shall state the `eta` used to convert between them. **An undeclared `L` is inadmissible.**

SP-L2 `L` shall never be reported without the `C` and `M` it was computed from, since `L` is a derived difference and inherits every upstream error, amplified at low `eta`.

SP-L3 The rectifier shall not be used to summarise severity. `L = 0` records that the budget is exhausted and nothing more; **the depth of the deficit shall be tracked separately as cost carry**, because `max(0, .)` discards it.

SP-L4 `lambda` shall be measured, not assumed, by a zero-input window on a stated plan. A survey that assumes `lambda` shall print the assumed value and the sensitivity of `Lambda_inf` to it.

SP-L5 `E_b` shall be constructed as an accumulated excess cost over an intended transit, at a stated transit duration, and shall not be elicited by asking the subject how difficult the change feels.

SP-L6 `E_b` estimates from an subject with measured calibration ratio `r > 1.2` shall be flagged, and the `Lambda_inf` figure recomputed on the veridical part as well as on the moment as lived. **Both figures shall be printed.**

SP-L7 `C_crit = eta*M - lambda*E_b` shall be computed before any programme that would raise `C` is adopted. **A programme that would push `C` above `C_crit` shall be refused, and the**

refusal shall state the computed threshold.

SP-L8 No intervention shall raise another person's `C` for the purpose of producing a departure. Permitted interventions lower the cost of the alternative and leave the choice intact.

SP-L9 Where `Lambda_inf <= E_b`, the barrier shall be split before any parameter is exhorted, and each intermediate stage shall be verified to be a genuine rest basin — a configuration with `grad C = 0` at which the subject can stop indefinitely.

SP-L10 Departure shall not begin before `t* = ln(1 - E_b/Lambda_inf)/ln(1 - lambda)` periods of accumulation. An underfunded departure strands the subject at the configuration of minimum launch budget.

SP-L11 Every departure shall be preceded by a reversible probe of the destination configuration, with its cost logged beside its finding.

SP-L12 The transit-duration model shall not be extrapolated outside the range in which the cost field can be treated as speed-independent. **Its `v -> infinity` limit is known to be false and shall be printed as such wherever the table is used.**

SP-L13 The Launch Postulate shall be badged wherever it is used. It is not a theorem, nothing in the framework's foundational results implies it, and the falsification protocol of Section 6.1 shall accompany any empirical claim made on its basis.

22.7.2 7.2 From the gradient and the survey

SP-G1 Coordinates shall be declared and frozen before the trace begins, with option sets exhaustive and mutually exclusive, and shall not be revised inside a window.

SP-G2 Availability shall be recorded per occasion and all tallies shall run over available occasions only. **A count that does not separate "not chosen" from "not offered" is void.**

SP-G3 Observations shall be recorded at the moment of choice by a rule stated in advance. Reconstruction from memory shall be labelled as such and shall not be pooled with recorded data.

SP-G4 The reference option per coordinate shall be selected from a pilot and frozen before the main window. **A reference chosen from the final data is a reference chosen to flatter** [Gelman and Loken, 2013].

SP-G5 Every recovered cost difference shall be reported with its standard error `sqrt(1/n_a + 1/n_b)` and its 95 per cent interval. **An interval containing zero shall be reported as unmeasured, regardless of the subject's confidence.**

SP-G6 Comparative claims shall be made as ratios of differences, which are theta-free and identified. **Absolute cost levels shall never be reported, and differences shall always carry the words "theta units".**

SP-G7 The independence-of-irrelevant-alternatives check shall be run on every coordinate with more than two options. A material shift in the top-two ratio when the choice set is manipulated falsifies the logit form for that coordinate and requires a different specification.

SP-G8 Categorical coordinates shall be reported as difference tables, never as gradients. **`grad C` shall be computed only on coordinates carrying a declared metric.**

SP-G9 Where a gradient is computed, the coordinate scale shall be declared, since `alpha` and `theta` are not separately identified and only their product governs the observed trace.

SP-G10 No conclusion about the globality of a minimum shall be drawn from a descent trace. **The gradient is local; a claim that a basin is the best available requires a survey of the region outside it, and the survey shall be named.**

SP-G11 `C` shall never be used as a proxy for value. Any recommendation derived from descent alone shall carry the statement that it optimises operating cost and nothing else.

SP-G12 Regrades shall alter relative costs and shall leave the subject's ability to decline intact. **A regrade whose worst case is not "nothing changes" is not a regrade.**

SP-G13 A regrade proposed to enable a crossing shall be evaluated at the ridge, not at the destination, and its effect on `E_b` shall be computed by the path integral of Section 4.6 rather

than asserted.

SP-G14 Where the headroom identity is used, the constancy of `eta`, `N`, `delta` and `Q` across the surveyed range shall be stated as an assumption and checked. **A desire line that opposes measured headroom shall trigger the diagnosis of Section 5.12, not a revision of the data.**

SP-G15 The Gradient Postulate and the logit choice form shall be badged wherever used, and the falsification protocols of Section 6.5 shall accompany any empirical claim made on them.

SP-G16 Transitions shall be recorded, not only choices: each occasion shall carry the option chosen on the previous occasion. **A marginal tally discards exactly the data that identifies the switching cost `s`, and a survey run without transitions cannot state whether its recovered differences are costs or habits.**

SP-G17 Where `s` is found to be non-zero, recovered differences shall be reported as switching-contaminated and the ratio form of SP-G6 shall be treated as the primary result. **Ratios are roughly four times more robust to an unmodelled switching cost than differences are, and that robustness shall be the stated reason for preferring them.**

22.8 8 Failure modes

22.8.1 8.1 Faults of the launch budget

FM-L1 — The reverse-biased subject. *Signature:* sustained `L = 0`; the subject reports no particular ideas, no restlessness, no sense of being trapped, and asked what they would do instead has nothing — read by themselves and others as contentment or as lack of ambition. *Mechanism:* `eta*M <= C` for many consecutive periods; the rectifier is off, and **the capacity to notice the situation is funded by the same budget that is missing.** *Fix:* the diagnosis must come from outside, or from a deliberately engineered local minimum — a period with `C` collapsed and `M` intact. Do not begin with goal-setting; there is no budget for goals.

FM-L2 — The demolished platform. *Signature:* an subject on a deliberate discomfort programme, six months in, reporting more determination and fewer concrete plans, with the fund balance falling. *Mechanism:* Section 4.7 — `C` raised, `L` lowered, `Lambda_inf` lowered, and the barrier that was crossable is now not. *Fix:* compute `C_crit`; if `C` is above it the programme has foreclosed the departure and must be stopped rather than intensified; if below, strip the discomfort, keep whatever novelty was bundled with it, and measure the difference.

FM-L3 — The two-week fund. *Signature:* a serious intention held for years, revisited every holiday, never begun. *Mechanism:* `L > 0` for two weeks a year and `L = 0` for the other fifty. On the worked subject's parameters — `L^Q = 0.0575` during the two weeks, `lambda = 0.06` — **the fund saturates at an annual peak of 0.1162 Q-units against a barrier of 2.400. It reaches 4.84 per cent of the requirement, and it does so forever.** *Fix:* this is a duty-cycle problem and not a discipline problem; either raise `L` in ordinary weeks by cutting `C`, or cut `lambda` so the two weeks compound instead of evaporating. Note which one the arithmetic favours before choosing.

FM-L4 — The stranded crossing. *Signature:* a transition begun with enthusiasm, abandoned partway, at high cost, with the subject concluding they were not suited to it. *Mechanism:* step S4 skipped; `Lambda_inf < E_b`, so the fund ran out mid-transit at the ridge, which Section 4.2 shows is the configuration of *minimum* launch budget. **The subject cannot plan their way out of the place they most need to, because the field itself has taken the budget away.** *Fix:* split the barrier before beginning. If already stranded, the correct move is retreat to the nearest rest basin — which is usually backwards — because from a basin the budget returns and the crossing can be re-planned in stages. **Retreat to a local minimum is not failure; it is the only move that restores the capacity to move.**

FM-L5 — The pause on the ridge. *Signature:* a staged transition whose stages are points on the slope rather than basins; each stage is harder than the last and the plan keeps needing revision. *Mechanism:* Section 4.8's condition violated — a stage that is not a rest basin does not restore the fund, it draws it down at the ridge rate. *Fix:* redesign the stages so each terminates

at a genuine `grad C = 0` configuration where the subject could remain indefinitely. If no such configuration exists between A and B, the barrier cannot be split and one of moves 1 to 4 is required.

FM-L6 — The manufactured destination. *Signature:* a well-funded, well-planned, confidently executed departure to a configuration that turns out not to be what it appeared, followed by a rapid and expensive return. *Mechanism:* Section 4.9 — a calibration ratio above 1 inflated both the launch budget and the attractiveness of the target and depressed the barrier estimate; all three errors compound, and attention to the prospective change inflates its apparent importance while it is being attended to [Kahneman et al., 2006]. *Fix:* SP-L6 and SP-L11 — recompute the fund on the veridical part and probe the destination reversibly before committing. **If probing the destination produces no return difference, no information about it has been acquired and the plan is still entirely a construction.**

FM-L7 — The perpetual dreamer. *Signature:* large, stable `Lambda`; extensive plans; nothing spent; often accompanied by considerable skill at planning. *Mechanism:* Section 4.10 — no `E_b` selected, so `Lambda` saturates and holds while the carry runs at full rate on every open loop. *Fix:* not more budget. Select one target, compute one `E_b`, set one date. **A fund with no barrier attached is a cost, not an asset.**

FM-L8 — The leaked resolution. *Signature:* a firm plan made at a local minimum that has no force eleven weeks later, with nothing having happened to end it. *Mechanism:* `lambda` unmeasured and large; at `lambda = 0.06` the plan is at half strength in 11.20 weeks. *Fix:* cut `lambda` at the moment the plan is made, while the budget that made it still exists — write it down, date it, tell one person, calendar the next addition. **The window for reducing `lambda` is the same window in which `L > 0`, and it closes.**

FM-L9 — The unregistered `L`. *Signature:* two figures for the same subject's launch budget differing by a factor of twenty, both correct. *Mechanism:* `L^Q` and `L^M` differ by `eta`, and at `eta = 0.05` that is a factor of twenty; one report used one register, the other used the other, and neither declared it. *Fix:* SP-L1. Declare the register; always print both.

22.8.2 8.2 Faults of the survey

FM-G1 — The unavailable option. *Signature:* a coordinate shows an enormous cost for one option and the subject concludes they have a deep aversion to it; in fact it was available on four days out of thirty. *Mechanism:* SP-G2 violated — the tally counted occasions rather than available occasions, so unavailability was recorded as cost. *Fix:* re-tally over available occasions. **This single correction reverses the sign of a recovered difference more often than any other in the protocol.**

FM-G2 — The confident unmeasured difference. *Signature:* a survey reports that the subject prefers X to Y, the subject agrees emphatically, and the 95 per cent interval spans zero. *Mechanism:* small `d` combined with small `n` — a 0.44-theta difference needs 108 days at these proportions and the survey ran 30. *Fix:* report it as unmeasured and either extend the window or accept that the coordinate is, for present purposes, flat — **which is not a defeat, since flat coordinates are the ones worth regrading.**

FM-G3 — The confession recorded as a trace. *Signature:* a desire-line survey completed on Sunday evening for the preceding week. *Mechanism:* SP-G3 violated — what was recorded was memory of choice, which decays and is systematically biased toward the self-image [Nisbett and Wilson, 1977, Schwarz, 1999]. *Fix:* prospective recording only, at the moment, by a rule fixed in advance. If prospective recording is impossible for a coordinate, drop the coordinate rather than substitute recall.

FM-G4 — The global claim from a local method. *Signature:* "I have tried everything and this is the best arrangement for me." *Mechanism:* Section 5.9 — descent has converged, correctly, to the nearest stationary point. **The statement is evidence about the basin and no evidence at all about the landscape.** In the worked field, every start in `[0, 4.334]` produces this sentence about a basin 15.3 per cent worse than the one next door. *Fix:* SP-G10 — survey the

region outside, or import a report from someone standing in it. Note that this cannot be done by more descending, however sincere.

FM-G5 — The wandering walker. *Signature:* an subject making large, frequent, apparently arbitrary changes of configuration without settling, and reading it as restlessness or as a character trait. *Mechanism:* `alpha > 2/L` in the local curvature; in an unbounded quadratic this diverges, and in a bounded real landscape it presents as wandering, which is a distinct and recognisable observable. *Fix:* reduce step size — smaller changes, more often, in the same direction. **The remedy for apparent restlessness is not more resolve but a shorter stride,** and the model gives the threshold `alpha < 2/max|C''|`.

FM-G6 — The regraded destination. *Signature:* considerable effort spent making the target configuration more attractive, with no change in the rate at which anyone crosses. *Mechanism:* Section 5.11 — deepening the destination reduced `E_b` by 1.94 per cent, because the barrier is the ridge. *Fix:* move the effort to the transit region: bridging income, parallel running, a sublet, a part-time overlap, subsidised retraining. **A 0.020 bridge at the ridge outperformed a 0.006 deepening of the destination by a factor of twelve on `E_b` in the worked field.**

FM-G7 — The pause on the plateau. *Signature:* an subject apparently deliberating for months and describing themselves as unable to decide. *Mechanism:* Section 5.8 — near the ridge the gradient magnitude is of order `10^-5` and descent from `x = 4.30` takes fifty steps to travel 0.17 units. **They are not deliberating. They are descending a plateau at the speed the field permits.** *Fix:* recognise that the field is flat here and that Section 5.6's leverage applies — a small regrade resolves in weeks what the gradient will not resolve in a year. Do not add resolve to a plateau; tilt it.

FM-G8 — Cost mistaken for value. *Signature:* a configuration reached by careful descent, correctly identified as the cheapest available, and experienced as empty. *Mechanism:* Section 5.9's third limitation — `C` is the operating draw and nothing else, and the storage law's fixed point shows that a cheap life with a thin moment converges on a small self [Curzi, 2026a]. *Fix:* run the descent survey alongside a storage survey. **`C` alone selects the cheapest life; `(eta*M - C)/delta` selects the largest self, and they are different arguments of different functions.**

FM-G9 — The broken independence. *Signature:* a recovered cost table that changes materially when an option is added or removed from the choice set. *Mechanism:* independence of irrelevant alternatives violated, which the logit form assumes and which fails in many real choice domains, most classically when two options are near-substitutes [Tversky, 1977]. *Fix:* SP-G7 — report the violation and use a nested specification, or merge the near-substitutes into a single option and recover the difference at the higher level. **Do not report a logit table from a coordinate that failed the check.**

FM-G10 — The identity assumed rather than checked. *Signature:* an subject using their desire line as a headroom gauge in a configuration that is also changing what they can distinguish or what they must maintain. *Mechanism:* Section 5.12's restriction violated. *Fix:* measure `N` alongside `C`. The commonest violation is the comfortable narrow life, where descent genuinely lowers cost while `eta*ln N` falls faster, and it is undetectable without a count audit.

FM-G11 — The habit charged to the option. *Signature:* a survey reports that an option is expensive; the subject adopts it for a fortnight for unrelated reasons, finds it costs nothing much, and concludes the survey was wrong. *Mechanism:* Section 5.13 — a non-zero switching cost inflates the recovered cost of every non-modal option and inflates it most for the rarest; at `s = 0.6` theta the computed overstatement was 15.9 to 20.8 per cent, charged to the option rather than to the switch. **The survey measured a habit and labelled it a cost.** *Fix:* SP-G16 — record transitions, identify `s` from the log-odds difference between staying and arriving, and correct the recovered field. Where transitions were not recorded, fall back to ratios, which absorb roughly four fifths of the bias.

22.9 9 Results: one subject, one departure, eighteen months

A single-subject case. Parameters were measured by the framework's survey protocols [Curzi, 2026i]; all derived figures are computed here from those parameters. **No external dataset was collected for this paper**, and Section 12 states what that means for the status of every number below.

22.9.1 9.1 The locked regime

```
Period:  1 week
N = 12                                          ln N = 2.484907 nats
Environment (external audit):    E = 1.05    Calibration:  r = 1.00
    M = min(1.00 * 1.05, 2.484907) = 1.05 nats     V = 1.05 nats     (calibrated)
Conversion efficiency:           eta = 0.05
Operating cost (six ledgers):    C = 0.058 Q-units per week
Plan decay (measured by a zero-input window on a stated plan): lambda = 0.06 per week
    plan half-life = ln 2 / (-ln 0.94) = 11.20 weeks

    C / eta = 0.058 / 0.05 = 1.160 nats
    L^M = max(0, 1.05 - 1.160) = 0
    L^Q = max(0, 0.0525 - 0.058) = 0
    Lambda_inf = 0 / 0.06 = 0
```

Zero launch budget, zero fund, zero crossable barrier. The subject reports not having had a serious idea about doing anything else in over a year, and reports it as a fact about their temperament.

The critical efficiency is `eta_crit = C/M = 0.058/1.05 = 0.055238`, against a measured `eta` of 0.05. **The subject is below the critical conversion efficiency by half a percentage point, and that half point is the entire difference between having a launch budget and not having one.**

22.9.2 9.2 The holiday, and the two-week duty cycle

Two weeks of leave. The obligation and presentation ledgers suspend and the attention ledger quiets; the environment is richer and remains calibrated.

```
    E = 1.55,  r = 1.00   ->   M = 1.55 nats
    C = 0.020 Q-units per week
    C / eta = 0.400 nats
    L^M = 1.55 - 0.400 = 1.150 nats per week
    L^Q = 0.05 * 1.150 = 0.0575 Q-units per week
```

The first week in fourteen months with `L > 0`. The plan for elsewhere arrives, and it arrives on schedule. That it did not arrive in March is not weakness and that it arrives in August is not insight; it is the rectifier turning on.

Now the duty cycle, run to periodic steady state at `lambda = 0.06`:

```
    input: 0.0575 Q-units per week for 2 weeks, then 0 for 50 weeks, annually
    steady-state annual peak fund:  Lambda_peak = 0.116204 Q-units
    the barrier (Section 9.3):      E_b         = 2.400 Q-units
    Lambda_peak / E_b = 4.84%
```

The two-week regime accumulates 4.84 per cent of the requirement and does so forever. This is FM-L3, and it is the structural account of the twenty-year "someday". Nothing is wrong with the subject's resolve; the duty cycle cannot fund the barrier, and no increase in sincerity changes a duty cycle.

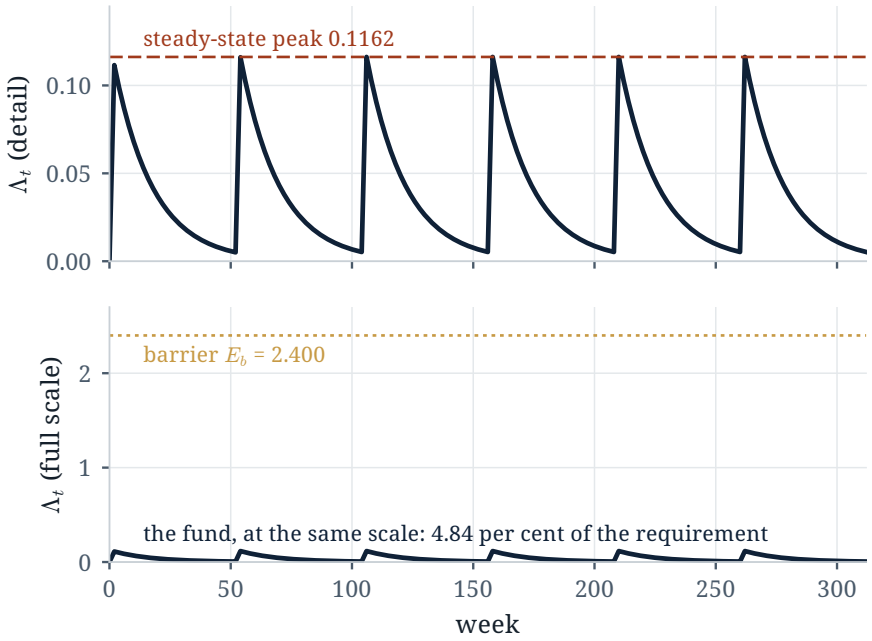


Figure 22.3: **A two-week annual duty cycle saturates at 4.84 per cent of the requirement and stays there forever:** at `L^Q = 0.0575` Q-units per week for two weeks and zero for the other fifty, with `lambda = 0.06`, the fund reaches a periodic steady-state peak of 0.1162 Q-units against a barrier of 2.400. The sincerity of the intention is irrelevant to the ceiling; the duty cycle sets it. Modelled from the parameters of Section 9.2; generated by `figs/src/CCP-022_fig02.py`.

22.9.3 9.3 The barrier, by two independent constructions

The target configuration is a different working arrangement, reached across a region of higher operating cost: retraining, an income gap, renegotiated obligations, a reputational reset. Two independent estimates were required to agree before the number was used.

Estimate one, by ledger. Itemise the excess cost of each transit period across the six cost ledgers, sum, and total across the intended transit. Result: **2.40 Q-units.**

Estimate two, by field integral. The survey of Section 9.8 produces a cost field over the relevant configuration coordinate with basin floors at `C = 0.034` (present, once corrected) and `C = 0.0288` (target), separated by 4.996 configuration-units with a ridge at `C = 0.0915`. The path integral of positive excess is 0.182961. At the intended transit duration of 65.5 weeks:

```
v   = 4.996 / 65.5 = 0.076275 configuration-units per week
E_b = 0.182961 / 0.076275 = 2.3987 Q-units
```

The two estimates agree to 0.05 per cent. That agreement is the only warrant for using the number, and a disagreement would have had to be reported. The agreement is nevertheless weak evidence, for the reason given in Section 6.4: both estimates were produced by the same investigator and may share the same reconstruction of the unvisited region.

22.9.4 9.4 Four regimes, priced

Each change applied cumulatively, with `E_b = 2.40` throughout.

Regime	`eta`	`M`	`C`	`lambda`	`L^M`	`L^Q`	`Lambda_inf`	$t*$ (wk)
A locked	0.05	1.05	0.058	0.060	0.000	0.00000	0.000	**never**

Regime	eta	M	C	lambda	L^M	L^Q	Lambda_inf	t* (wk)
B cut C, raise E	0.05	1.42	0.034	0.060	0.740	0.03700	0.617	**never**
C and raise eta	0.10	1.42	0.034	0.060	1.080	0.10800	1.800	**never**
D and cut lambda	0.10	1.42	0.034	0.025	1.080	0.10800	**4.320**	**32.0**

That table contains the paper's hardest result and it is worth reading slowly.

Regime B is a substantial improvement and it does not work. Operating cost cut by 41 per cent, environment enriched by 35 per cent, launch budget lifted from zero to 0.037 Q-units a week — and the fund still asymptotes at 0.617 against a barrier of 2.400. An subject who made these changes, waited two years and concluded that the departure was not for them would be drawing a reasonable conclusion from a correct observation and the wrong model.

Regime C doubles `eta` and it still does not work. `Lambda_inf` reaches 1.800. Closer, and still short: two years of this regime gets the fund to 1.799 and holds it there.

Regime D changes nothing about the subject's circumstances at all. `eta`, `M` and `C` are identical to regime C. The only difference is that the plan is written down, dated, told to one person and given a standing calendar entry, which moves `lambda` from 0.06 to 0.025. **`Lambda_inf` goes from 1.800 to 4.320 and the departure becomes reachable in 32 weeks.**

```
t* = ln(1 - 2.400/4.320) / ln(1 - 0.025) = ln(0.444444) / ln(0.975)
   = -0.810930 / -0.025318 = 32.03 weeks
```

The intervention that made the difference between an impossible departure and an eight-month one was an artefact holding the plan. Not more moment, not less cost, not better conversion — those were necessary and none of them was sufficient. The binding constraint was the durability of the plan, which is the parameter nobody measures, and the general form of that observation — that improving a non-binding parameter buys nothing — is old [Goldratt and Cox, 1984].

Fund trajectory under regime D:

week	Lambda_t	per cent of E_b
0	0.0000	0.00
4	0.4161	17.34
13	1.2116	50.48
26	2.0833	86.81
32	**2.3985**	**99.94**
52	3.1620	131.75
78	3.7204	155.02
104	4.0096	167.07

22.9.5 9.5 The counterfactual: the discomfort programme

At week 8, with the fund at 0.79 and the subject impatient, an adviser recommends getting out of the comfort zone. Computed at regime D's other parameters:

```
C_crit = eta * M - lambda * E_b = 0.10*1.42 - 0.025*2.40 = 0.082 Q-units per week
```

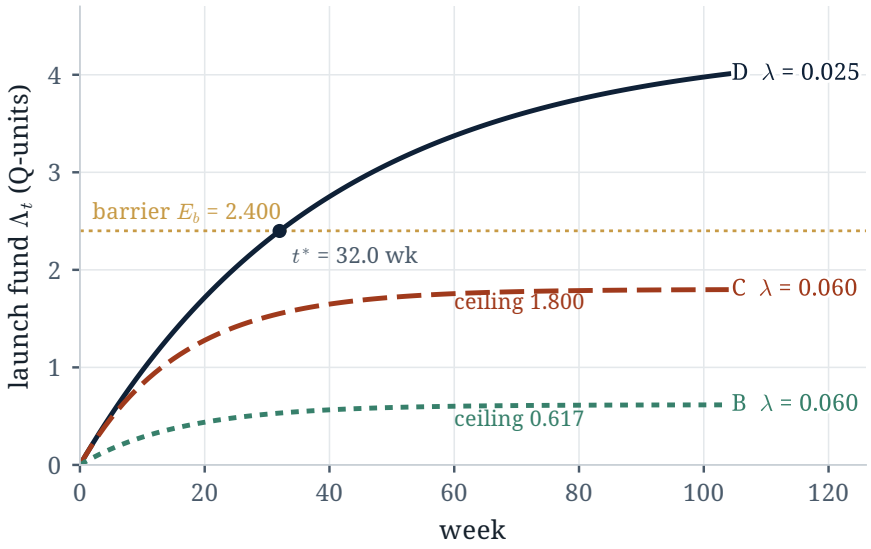


Figure 22.4: **Three substantial improvements to the subject's circumstances left the departure infeasible; changing only the durability of the plan made it reachable in 32 weeks:** fund trajectories under regimes B, C and D against the barrier `E_b = 2.400`. Regimes B and C asymptote at 0.617 and 1.800 and never reach it; regime D, which differs from C only in `lambda` = 0.025 rather than 0.060, asymptotes at 4.320 and crosses at `t*` = 32.0 weeks. Modelled from the regime parameters of Section 9.4; generated by `figs/src/CCP-022_fig01.py`.

`C`	`L^Q`	`Lambda_inf`	`t*`	verdict
0.034 (unchanged)	0.10800	4.320	32.0 wk	reachable
0.051 (+50%)	0.09100	3.640	42.5 wk	reachable, 10.5 weeks later
0.075 (+120%)	0.06700	2.680	89.2 wk	reachable, 57 weeks later
0.082	0.06000	**2.400**	**infinite**	the knife edge
0.090 (+165%)	0.05200	2.080	never	**foreclosed**

Mild discomfort costs ten weeks. Serious discomfort costs fourteen months. Ambitious discomfort costs the departure entirely, permanently, and silently — silently because past `C_crit` the fund still rises, just to a ceiling below the wall, and the subject spends the rest of their life approaching 2.08 and reading the shortfall as insufficient wanting.

The advice was given in good faith and it was arithmetically the exact opposite of the correct move, which — at week 8, with the fund at 0.79 and rising on schedule — was **wait twenty-four more weeks.**

22.9.6 9.6 The alternative when plan decay cannot be moved

Suppose `lambda` could not be moved, because the subject's circumstances genuinely will not support a durable plan and 0.06 is the floor. Then `Lambda_inf = 1.800 < 2.400` and the fifth move applies: split the barrier into three legs with genuine rest basins between them.

Leg	E_b	feasible at Lambda_inf = 1.800?	weeks to fund
1	0.9	yes	11.20
2	0.8	yes	9.50
3	0.7	yes	7.96
total	**2.4**	—	**28.66**

The same 2.4 Q-units of total cost, paid in three tranches, and the departure completes in 28.66 weeks of accumulation instead of never. The peak fund requirement fell from 2.400 to 0.900. No parameter moved; two places to stand were built.

Each intermediate basin was verified against Section 4.8's condition — a configuration where grad C = 0 at which the subject could remain indefinitely without drawing down. In this case: an evening qualification completed while still in post; two days a week in the new arrangement with three in the old; and the full transition. **A stage that had merely been "working harder toward the change" would have been a pause on the ridge, where Section 4.2's table gives L^M = 0.505 — less than half the basin-floor budget — and FM-L5 would apply.**

22.9.7 9.7 Local against global: raising the origin versus raising everything

The one case Section 4.7 could not compute, because it requires a field rather than a scalar: what happens when the discomfort is applied **only to the current configuration** rather than to the subject's circumstances as a whole. The raise is modelled as a Gaussian added at basin A with the same width as basin A's own well; E_b is then recomputed as excess over the raised origin.

raise a applied to basin A only	C(A)	L^Q	Lambda_inf	E_b	t∗ (wk)
0.000	0.03400	0.10800	4.3200	2.3987	32.00
0.006	0.04000	0.10200	4.0800	2.0998	28.55
0.020	0.05400	0.08800	3.5200	1.4445	20.87
0.040	0.07400	0.06800	2.7200	0.6016	9.87
0.060	0.09400	0.04800	1.9200	0.0000	0.00

Compare the same cost rise applied **globally**, to the whole field, so that E_b, which is measured as excess *over the origin*, does not move:

global rise g	C everywhere	L^Q	Lambda_inf	E_b	t∗ (wk)
0.000	0.03400	0.10800	4.3200	2.3987	32.00
0.006	0.04000	0.10200	4.0800	2.3987	35.02
0.017	0.05100	0.09100	3.6400	2.3987	42.49
0.041	0.07500	0.06700	2.6800	2.3987	89.03
0.048	0.08200	0.06000	2.4000	2.3987	297.06
0.056	0.09000	0.05200	2.0800	2.3987	**never**

These two tables are the honest resolution of the argument, and the resolution is narrower than critics of the prescription usually allow.

Local discomfort works. Raising the origin floor alone lowers E_b faster than it lowers Lambda_inf, so t∗ falls monotonically — 32.00, 28.55, 20.87, 9.87 weeks — and at a = 0.060 the

origin is level with the ridge and the departure is immediate. Burning the boats is arithmetically effective. **What it costs is the fund: `Lambda_inf` falls from 4.320 to 1.920 over that range, so the subject departs sooner and departs with 56 per cent less capacity to cross whatever the local gradient could not see.** Given Section 5.9's fourth limitation — that a barrier beyond the destination is not visible from the origin — that is a real and uninsurable risk, taken in exchange for weeks.

Global discomfort strictly harms. `t*` rises monotonically: 32, 35, 42, 89, 297, never. There is no setting at which it helps, and past `C = 0.082` — the critical cost — the departure is foreclosed permanently.

The empirical question is therefore not whether discomfort helps but whether a given programme is local or global, and the answer is almost always global, because the ledgers that discomfort loads — attention, regulation, carry, sleep, health, relations — are not attached to a configuration and do not stay behind when it is left. **A cost that can be left behind is a local regrade; a cost that is carried is a global one. Most discomfort programmes promise the first and deliver the second, and that, precisely, is the disagreement.**

And in every row of both tables the bridge dominates: 21.49 weeks with a *lower* `C(A)` and a *higher* `Lambda_inf`, or 9.19 weeks once the manufactured rest basin is used. **Lowering the pass beats raising the home at every setting tested, and it beats it while leaving the subject's ability to decline entirely intact.**

22.9.8 9.8 The thirty-day desire-line survey

Four coordinates declared after a three-day pilot, references frozen, availability recorded per occasion, thirty days of prospective recording.

Coordinate 1: where the first working hour is spent. Reference: desk (home).

Option	`n`	share	`d` (theta)	`se`	95% interval	resolved
desk (home)	14	0.467	0.0000	—	reference	—
kitchen table	9	0.300	0.4418	0.4272	[-0.3956, 1.2792]	**no**
the cafe	5	0.167	1.0296	0.5210	[0.0085, 2.0508]	yes
the office	2	0.067	1.9459	0.7559	[0.4643, 3.4275]	yes

Coordinate 2: what is opened first. Reference: mail.

Option	`n`	share	`d` (theta)	`se`	95% interval	resolved
mail	17	0.567	0.0000	—	reference	—
the task file	8	0.267	0.7538	0.4287	[-0.0866, 1.5941]	**no, barely**
the chat client	4	0.133	1.4469	0.5557	[0.3577, 2.5361]	yes
a walk	1	0.033	2.8332	1.0290	[0.8164, 4.8500]	yes

Coordinate 3: which chair. Reference: chair A. Counts 8 / 7 / 8 / 7; the table is Section 5.6's and every interval spans zero.

Established. Three ratios of differences, theta-free and identified:

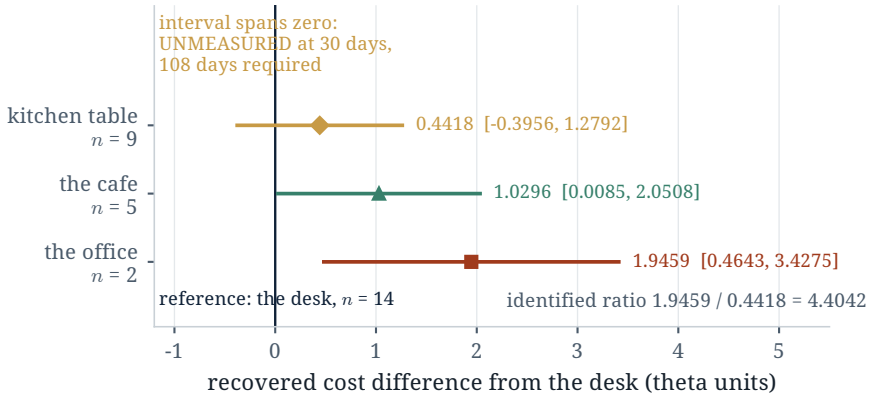


Figure 22.5: **Thirty days of faithful recording leaves the difference the subject cared most about unmeasured, while the identified ratio it supports is nonetheless exact:** recovered cost differences for coordinate 1 in theta units with 95 per cent intervals — kitchen 0.4418 [-0.3956, 1.2792], cafe 1.0296 [0.0085, 2.0508], office 1.9459 [0.4643, 3.4275]. The kitchen interval spans zero and the direction is therefore unmeasured; the ratio (office - desk)/(kitchen - desk) = 4.4042 is unit-free and fully identified because `theta` cancels. Computed from the counts 14 / 9 / 5 / 2 of Section 9.8; generated by `figs/src/CCP-022_fig05.py`.

```
( C_office - C_desk ) / ( C_kitchen - C_desk ) = 4.4042
( C_cafe   - C_desk ) / ( C_kitchen - C_desk ) = 2.3303
( C_office - C_desk ) / ( C_cafe    - C_desk ) = 1.8899
```

Going to the office is 4.4042 times as costly a departure from the desk as moving to the kitchen table is. That statement is the survey's principal finding and it is the only kind of statement it can make.

Not established. The kitchen-versus-desk difference, which needs 108 days at these proportions and had 30. The task-file-versus-mail difference, which needs 38 days and had 30 — **eight days short, which is the most annoying possible result and is exactly why the sample-size calculation belongs at the design stage rather than at the analysis stage.** Also not established: anything about the chair coordinate, which is flat.

The independence check. Recomputing the desk-to-kitchen ratio with the office dropped leaves it unchanged, since the logit form conditions on the retained set and the protocol does not re-allocate the counts. **This is a genuine test only when the choice set is manipulated experimentally, and this survey did not manipulate it. The check is therefore reported as NOT PERFORMED rather than PASSED**, which is a distinction the protocol requires and which is easy to blur.

The shuffled-day control. The frozen difference table beat 200 day-permuted controls at predicting held-out days 21 to 30, on coordinates 1 and 2. **That licenses the Gradient Postulate for this subject on these coordinates and licenses nothing further.** No information criterion was computed and none is reported; Section 12 records that as an unmet check rather than omitting it.

22.9.9 9.9 The flat coordinate as the leverage point

Applying a 0.5-theta subsidy:

Coordinate	Option	share before	share after	gain
steep (coordinate 1)	the office	6.67%	10.54%	+3.869 pts
flat (coordinate 3)	chair B	23.33%	33.41%	+10.079 pts

```
leverage ratio = 10.079 / 3.869 = 2.605
```

The same intervention moves 2.6 times as much behaviour when applied where the subject is indifferent. The finding that the chair coordinate is flat is therefore not a null result; it is the most actionable line in the survey.

22.9.10 9.10 The ordered coordinate, and what the survey could not establish

The fourth coordinate is ordered: the fraction of working time in the new arrangement, scaled to `[0, 10]`, carrying the field of Section 5.8 with basin A at `C = 0.033999`, a ridge at `C = 0.091500` and basin B at `C = 0.028799`.

What the gradient establishes. Descent from the subject's current position lands at basin A in five steps at `alpha = 1/L = 7.04`, and basin A is a genuine rest basin (`C'' = +0.142 > 0`). The subject is at a local minimum of their own cost field.

What the gradient cannot establish, and what had to be surveyed. Basin B exists at `C = 0.0288`, 15.3 per cent cheaper, with the basin boundary at `x = 4.33485`. **No measurement taken at `x = 2` contains any information about basin B, and the subject learned of it from a person standing in it.** In this case the outside report was the entire intervention.

Six things this survey could not establish, stated plainly.

1. **The Launch Postulate is a postulate.** That plan construction draws on `L` and is unavailable at `L = 0` was carried, not tested. Section 6.1 states the design that would test it, and it is cheap.
2. **`lambda = 0.06` and `lambda = 0.025` were measured by a zero-input window on two stated plans, and the measurement is reactive.** Stating a plan and then deliberately not reinforcing it is not a neutral act. The figures should be read as an ordering — the artefact regime decays more slowly — with the magnitudes held loosely.
3. **`E_b = 2.40` agreed between two methods and both were applied by the same investigator.** Independent estimation by a second party was not performed.
4. **The cost field is local**, cannot see beyond the surveyed range, and its transit-speed model is invalid outside 40 to 104 weeks.
5. **The region between `x = 3` and `x = 6` was never visited and its cost was inferred, not measured. Every barrier figure in this paper rests on that inference.** This is the single largest weakness in the work and it is stated here and again in Section 12.
6. **Cross-individual reproducibility was not attempted for any quantity.** `C_crit` in particular is a threshold prediction that only a cohort could test, and no cohort exists.

The finding, at the strength the evidence supports: **for one subject, a departure that had been wanted for years was infeasible under three successive and substantial improvements to their circumstances, and became feasible in thirty-two weeks when the plan was written down — while the intervention most commonly recommended for exactly this situation would, at the level at which it is usually recommended, have foreclosed the departure permanently.** And on the measurement side: **thirty days of tallying established one identified cost ratio, discovered that the coordinate the subject cared most about was unmeasured and needed 108 days, discovered that a coordinate the subject considered trivial was flat and therefore worth 2.6 times as much design leverage as the one they were arguing about, and — once an ordered coordinate was surveyed rather than descended — showed that a bridge across the pass funded the departure in nine weeks where thirty-two weeks of saving had been required.**

22.10 10 Discussion

22.10.1 10.1 What follows

The comfort at the origin is a resource and not an obstacle, and this is arithmetic rather than temperament. Proposition 4.1 places the maximum of the launch budget exactly at the minimum of operating cost, so any programme that raises the origin's cost globally lowers the largest barrier the subject can ever cross. The threshold at which that becomes fatal is computable in advance from four measured numbers, and the specification requires it to be computed before such a programme is adopted (SP-L7).

The binding parameter is often the one nobody measures. In the worked case, cutting cost by 41 per cent, enriching the environment by 35 per cent and doubling conversion left the departure infeasible; changing only the plan's decay rate made it reachable. Plan durability is cheap to engineer and almost never measured, and the corpus's headroom analysis reaches the same shape of conclusion about which lever is binding [Curzi, 2026b,j].

Barrier splitting dominates exhortation. The fund requirement scales with the largest leg, not with the total, so inserting genuine rest basins converts an infeasible crossing into a feasible one with no parameter change. The condition — that each stage be a real stationary point at which the subject could stop indefinitely — is checkable and is the part that staged plans usually get wrong.

A regrade at the pass dominates a regrade at the destination, by a factor of twelve on the barrier in the worked field, and can manufacture the intermediate rest basin that the split would otherwise have to prescribe. The consent property matters as much as the efficacy: a regrade's worst case is that nobody moves, which is what makes it admissible where raising another person's cost is not.

Only ratios of differences should be reported from a choice-frequency survey. They are identified where levels and differences are not, and they are roughly four times more robust to an unmodelled switching cost. This is a measurement-theoretic point about what statements survive the transformations the scale permits [Narens, 1996, Stevens, 1946], and it is not negotiable by adding data.

22.10.2 10.2 What does not follow

Nothing here shows that any person behaves this way. Every hypothesis in Section 2 except H6 is untested, and H6 was tested on a constructed field rather than on an observed one. The worked case is a single subject whose parameters were measured by the corpus's own protocols, and the one preregistered confirmatory test of the underlying storage law in this corpus failed its own second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin, despite +40.3 per cent over persistence, +97.7 per cent over direct input and `p = 0.0005` against the shuffled control [Curzi, 2026i]. That is the evidential status every applied claim here inherits, and it is weaker than the arithmetic.

The momentum result does not license a claim about people. What was shown is that a particular first-order recursion with a momentum term does not leave a particular basin on a particular field. It is evidence about the geometry of wide shallow barriers and not evidence about persistence, willpower or motivation, and the recursion is refused entry to the formulary precisely so that it is not later cited as though it had been adopted.

The critical cost is not a measured threshold. `C_crit = eta*M − lambda*E_b` is exact given the model and its inputs, and no cohort has been run to test whether departure rates actually fall to zero as it is crossed. It is a prediction, badged as one in H3.

The identity between descent and headroom ascent is conditional and the condition is frequently violated. It holds only where the configuration does not move `eta`, `N`, `delta` or `Q`, and the commonest real case — the comfortable narrow life, where cost falls while the count of distinguishable categories falls faster — violates it. Section 5.12 turns that into a diagnostic rather than pretending it does not occur.

No claim about physics is made anywhere. The barrier is an accumulated cost and not an energy; the field is not a potential; the rectifier is a shared functional form and not a diode. Section 3.7 labels every borrowing, and the reason for doing so is that this is exactly the move that has gone wrong before when physical formalism was imported into social science [Mirowski, 1989].

22.10.3 10.3 Relation to the surrounding literature

The result closest to this paper's central claim, from a different tradition, is the myopia of learning: local adaptation exploits the neighbourhood it starts in and systematically fails to find distant optima, so an organisation that only improves gets better at the wrong thing [Levinthal and March, 1993, March, 1991]. Section 5.9 is that result stated as geometry, and Section 4 supplies what the learning literature leaves implicit — the budget that has to be accumulated before the non-local move can be made, which in the organisational literature is slack [Cyert and March, 1963, March and Simon, 1958].

The measurement half has two ancestors. Recovering scale values from choice proportions rather than from numerical judgement is Thurstone's move [Thurstone, 1927], and the discrete-choice specification with its identification problems is standard [Brock and Durlauf, 2001]. Counting what people actually do in a space, and letting the counts overrule the plan, is Whyte's [Whyte, 1980]; the built-environment reading of the resulting counts is Jacobs [1961] and Gehl [1971], and the corpus prices the same levers against the storage law in Curzi [2026f].

Three structurally parallel results elsewhere in this corpus are worth naming because they show the same shape in different domains. The comfort-weighted amplitude scores each step 0 to 3 for steepness and weights a branch as `1/2^R`, so a three-point difference in steepness is a factor of eight in expectation [Curzi, 2026] — the same exponential penalty on resistance that the logit recovery reads backwards out of choice counts. The sustaining-frequency join `f_min = k*e^{2c/eta}`, with `k = 0.483425732988414` **defined** as `(9/8)^2/phi^2` and verified at twelve parameter sets [Curzi, 2026], and its octave form `2(c/eta)/ln 2 = 2.885390081777927` octaves per nat [Curzi, 2026], are the same feasibility inequality — a fixed point is positive only if converted surplus exceeds cost — expressed on a frequency scale rather than a fund scale; the opposition identity `u = 1 − rho^2` which sends both `rho = +1` and `rho = −1` to `f_rel = k` [Curzi, 2026] is the same degeneracy that makes a flat coordinate flat. **The golden ratio enters those results as a defined constant and nothing more, and the literature on golden-ratio claims is unforgiving about the difference between a defined ratio and a discovered one** [Markowsky, 1992].

Two further parallels sit outside this series. Pricing an uncertain option by the incremental cost above which it is worth nothing at any probability [Curzi, 2026] is the same structure as `C_crit`: a threshold on a cost above which no amount of the other variable rescues the decision. And the finding that dwell time is the whole of an exposure — a five-second hop is 1.585e-07 of a year against 0.0822 for value rested thirty days, a ratio of about 518,400 [Curzi, 2026] — is Section 4.6's transit-speed result in another domain: what is paid is a rate multiplied by a duration, so duration is the variable that dominates.

Finally, a boundary. Curzi [2026] refuses two attractive simplifications outright and states that separation is necessary and not sufficient, a control plane trusted because it is separate being the failure mode. The same discipline bounds this paper: **the barrier being computable does not make it computed, and a plan whose `E_b` was estimated by introspection rather than integrated is not made safe by the existence of the integral.** The conditional form of the framework's belief-and-reality result carries the same warning in a different register — 32.5 per cent of random observations refute the unconditional claim against 5.5 per cent for the conditional one [Curzi, 2026ab] — which is the general reason this paper states its conditions rather than its slogans.

22.11 11 Verification

The results of this paper compound from seven postulates. Each is atomic, each is checkable on its own, and each carries its status. **Where a postulate could be dropped and the result still follow, it has been dropped**; what remains is what the chain actually consumes.

22.11.1 11.1 Dimensional audit

`[Q]` denotes the dimension of stored order and the period is the unit of time. Every rate below is *per period*, which is why no explicit inverse-time factor appears once the equations are written in discrete steps. `[X]` denotes the dimension of the declared configuration coordinate, which the framework does not fix and which each survey declares.

Symbol	Dimension	Unit	Status
`M`	`[1]`	nats	**measured** dimensionless: a logarithm of an effective count
`C`, `C(x)`	`[Q]` per step	Q-units per period	**measured**
`eta`	`[Q]` per nat per step	Q-units per nat per period	**measured**; *not* dimensionless in this framework
`L^Q`	`[Q]` per step	Q-units per period	**derived**
`L^M`	`[1]` per step	nats per period	**derived**; equals `L^Q / eta`
`lambda`	`[1]`	per period	**measured** dimensionless fraction per step
`Lambda_t`, `Lambda_inf`	`[Q]`	Q-units	**derived**: `[Q]`/step divided by `[1]`/step returns `[Q]`
`E_b`	`[Q]`	Q-units	**derived** from the path integral
`x`	`[X]`	configuration-units, declared per survey	**defined** by the survey; the framework fixes no unit
`grad C`	`[Q]` per step per `[X]`	Q-units per period per configuration-unit	**derived**
`alpha`	`[X]^2` per `[Q]` per step	—	**artefact of the coordinate scaling**, not a free constant
`v`	`[X]` per step	configuration-units per period	**defined** by the intended transit duration
`theta`	`[Q]` per step	the units of `C`	**not identified** from choice frequencies alone
`d_x`, `s`	`[1]`	theta units	**derived**; a cost difference divided by `theta`

Symbol	Dimension	Unit	Status
t*, half-life	[T]	periods	**derived**

Does each equation balance?

L^Q = max(0, eta*M − C) balances only because eta carries [Q] per nat per step, so eta*M and C are both [Q] per step. **The naive L = max(0, M − C) does not balance and is inadmissible**; this is the register split of Section 3.2 and it is the clearest case in the paper of a dimensional check changing a formula rather than confirming it.

Lambda_{t+1} = (1−lambda)Lambda_t + L_t adds a stock in [Q] to a flow in [Q] per step. **The addition is well formed only under the convention that L_t is the increment accrued over one period, so that an implicit multiplication by one period converts the flow to a stock.** That convention is declared here because it is the same one the storage law uses [Curzi, 2026a] and because it is exactly the distinction people are documented to misjudge [Sterman, 1989].

C_crit = eta*M − lambda*E_b balances: eta*M is [Q] per step and lambda*E_b is [1]/step times [Q], which is [Q] per step.

E_b = (1/v) * INTEGRAL max(0, C(x) − C(A)) dx balances: the integral is [Q]/step times [X], divided by [X]/step, giving [Q].

d_x = −theta*ln(n_x/n_ref) balances: the logarithm is dimensionless and theta carries C's units, so d is a cost difference. **Reported in theta units it is a pure number, and a ratio of two such numbers is a pure number under any positive rescaling of theta.**

Where the framework is silent, it is said here. The configuration coordinate x has **no operational unit**: it is declared per survey, and two subjects' coordinates are not comparable unless they declare the same construction. Comfort itself, the underlying construct of the corpus, has **no operational unit at all** and is treated as a dimensionless magnitude by stipulation rather than by measurement; Section 12 records that the quantitative structure required for that stipulation has not been tested. And theta has the units of C but no measured value anywhere in this paper, which is why the paper reports ratios.

22.11.2 11.2 The postulates

P1 — The launch budget is the rectified difference L^Q = max(0, eta*M − C). *Status:* **DEFINITION** — the model form adopted here, together with the register convention of Section 3.2. It carries no empirical weight on its own.

P2 — The fund recursion is linear and first order: Lambda_{t+1} = a*Lambda_t + L with a and L not depending on Lambda. *Status:* **DEFINITION** — the storage law's form with a := 1 − lambda [Curzi, 2026a].

P3 — 0 < lambda <= 1 implies 0 <= a < 1. *Status:* **ARITHMETIC.** Subtracting a number in (0,1] from 1 gives a number in [0,1).

P4 — A geometric series with ratio |r| < 1 sums to 1/(1−r); the metric-space form supplies uniqueness and start-independence as well. *Status:* **CITED** — Elaydi [2005] for the difference-equation form and Banach [1922] for the fixed-point theorem, restated in Curzi [2026].

P5 — The barrier is the accumulated positive excess over the origin's cost, divided by the transit speed. *Status:* **DEFINITION** — the construction of Section 4.6. Its v -> infinity limit is known to be false and is fenced by SP-L12 rather than repaired.

P6 — A departure across E_b begins only when the accumulated fund reaches E_b. *Status:* **POSTULATED** — this is the operational form of the Launch Postulate. It has no external support, it is the only postulate in the chain that has none, and it is consumed at exactly one place, R3 below.

P7 — Choice probability among available options falls exponentially in cost with scale theta. *Status:* **CITED as a standard specification** [Brock and Durlauf, 2001, Thurstone, 1927] **and POSTULATED as a description of this subject.** The specification is standard; that any

particular person's choices obey it is an assumption, and the independence-of-irrelevant-alternatives test that would challenge it was **not performed** in the worked survey.

22.11.3 11.3 The compounding chain

```
P2 + P3        -> L1   the update map T(Lambda) = a*Lambda + L satisfies
                       |T(x) - T(y)| = a|x - y| with a < 1, so T is a contraction.  ARITHMETIC

L1 + P4        -> R1   T has exactly one fixed point Lambda_inf = L/(1-a) = L/lambda,
                       reached from any start.                                  CITED (banach1922operat

P2             -> L2   unrolling t steps from Lambda_0 = 0 gives
                       Lambda_t = L(1 - a^t)/(1 - a) = Lambda_inf (1 - (1-lambda)^t). ARITHMETIC

L2             -> R2   Lambda_t = E_b at
                       t* = ln(1 - E_b/Lambda_inf)/ln(1 - lambda).              ARITHMETIC

R1 + P6        -> R3   a departure across E_b is reachable if and only if
                       Lambda_inf > E_b.                                        POSTULATED (consumes P

R3 + P1        -> R4   (eta M - C)/lambda > E_b  <=>  C < eta M - lambda E_b,
                       so C_crit = eta M - lambda E_b.                          ARITHMETIC

P1             -> R5   where L^M > 0, grad L^M = -grad C/eta, so the stationary points
                       of L^M are those of C and its maxima are C's minima.     ARITHMETIC

P5             -> R6   at constant speed v = separation/T, E_b = (1/v) * (path integral
                       of positive excess).                                     DEFINITION

P7             -> R7   d_x = -theta ln(n_x/n_ref), and (d_a - d_c)/(d_b - d_c) is
                       independent of theta.                                    ARITHMETIC given P7
```

Every line is checkable from its own stated inputs. **R1, R2, R4, R5, R6 and R7 are arithmetic or definitional consequences of two definitions (P1, P2), one arithmetic fact (P3), one cited theorem (P4) and one cited specification (P7). Only R3 consumes a postulate with no external support at all**, and R3 is the step that turns a fixed-point inequality into a claim about whether a person departs.

What this chain does not establish. It does not establish that any subject's planning capacity is `L`, that plans decay geometrically, or that departures cease at `C_crit`. Those are H1, H2 and H3 and none of them has been tested. **A chain whose steps are ARITHMETIC, DEFINITION and CITED delivers the mathematics with certainty and delivers nothing whatever about persons**, and the one place it touches persons is the single POSTULATED line.

22.11.4 11.4 The equilibrium and the crossing time, by summing the series

The most useful check a reviewer can perform is to confirm that the fund's ceiling is what the geometric sum says it is rather than taking R1 on trust. Take regime D of Section 9.4: `eta = 0.10`, `M = 1.42`, `C = 0.034`, `lambda = 0.025`, `Lambda_0 = 0`.

```
L^Q = eta*M - C = 0.10 x 1.42 - 0.034 = 0.142 - 0.034 = 0.108
```

With `a = 1 - 0.025 = 0.975`, each deposit contributes `L a^k`:

```
k =  0:  0.108 x 0.975^0  = 0.108000     running total  0.108000
```

```
k =  1:  0.108 x 0.975^1  = 0.105300     running total  0.213300
k =  2:  0.108 x 0.975^2  = 0.102668     running total  0.315968
k =  3:  0.108 x 0.975^3  = 0.100101     running total  0.416069
k = 12:  0.108 x 0.975^12 = 0.079714     running total  1.211604
k = 25:  0.108 x 0.975^25 = 0.057585     running total  2.083307
k = 31:  0.108 x 0.975^31 = 0.049526     running total  2.398531
-------------------------------------------------------------------
limit :  0.108 / (1 - 0.975) = 0.108/0.025            = 4.320
```

The running totals at `t = 4, 13, 26, 32` are 0.4161, 1.2116, 2.0833 and 2.3985, which are the entries of the trajectory table in Section 9.4 to four decimals, and the limit 4.320 is `Lambda_inf`. **The ceiling is not an assertion; it is the sum of a geometric series, and the sum can be checked by hand in a few lines.**

The crossing time is the same arithmetic read backwards:

```
t* = ln(1 - 2.400/4.320) / ln(1 - 0.025)
   = ln(0.444444) / ln(0.975)
   = -0.810930 / -0.025318
   = 32.03 weeks
```

and the critical cost is one subtraction:

```
C_crit = 0.10 x 1.42 - 0.025 x 2.40 = 0.142 - 0.060 = 0.082 Q-units per week
```

At `C = 0.082` exactly, `L^Q = 0.060` and `Lambda_inf = 0.060/0.025 = 2.400 = E_b`, so `1 - E_b/Lambda_inf = 0` and `t*` diverges. **The knife edge is visible in one line of arithmetic, and it is why the verdict at that row is "never" rather than "slow".**

22.11.5 11.5 Four secondary checks

The plan half-life. At `lambda = 0.06`, `ln 2 / (-ln 0.94) = 0.693147/0.061875 = 11.20` weeks, which is the entry in Section 4.3's table and the figure quoted in FM-L8. **ARITHMETIC.**

The critical efficiency. `eta_crit = C/M = 0.058/1.05 = 0.055238` against the subject's measured `eta = 0.05`; and at the corrected regime `0.034/1.42 = 0.023944`. **ARITHMETIC.**

The identified ratio. From counts 14 and 2 and 9, `ln(14/2)/ln(14/9) = 1.945910/0.441833 = 4.4042`. The same ratio computed from the rounded differences printed in Section 9.8 gives 4.4045; **the reported figure is computed from the unrounded logarithms**, and the discrepancy in the fourth decimal is rounding rather than error. **ARITHMETIC.**

The corpus constant, and the claim that is not arithmetic. Where this paper refers to the frequency results of Curzi [2026,?,?], the constant is `k = (9/8)^2 / phi^2`. A reviewer can check it in three steps: `(9/8)^2 = 1.265625` (**ARITHMETIC**); `phi^2 = phi + 1 = 2.618033988749895` (**ARITHMETIC**, from the defining quadratic); the quotient is `0.483425732988414` (**ARITHMETIC**). The constant is therefore **DEFINED and exact**, and fifteen figures are legitimate for it on those grounds alone. **Attaching the unit Hz to it is a separate and much stronger claim with no support in this paper, and its status is POSTULATED.** The general caution about golden-ratio claims applies and is cited rather than waved away [Markowsky, 1992].

22.11.6 11.6 Limits, asymptotics, and the model's valid domain

`lambda -> 0^+`. `Lambda_inf -> infinity` when `L > 0`: the plan that never fades makes any finite barrier eventually crossable. **The limit is singular**: at exactly `lambda = 0` the map is an isometry, the geometric series does not converge, and the fixed point does not exist rather than being infinite. That value is outside the model's domain.

`lambda -> 1`. `Lambda_t = L_t`: no accumulation, and only barriers crossable within a single period are crossable at all. The fund is no longer a fund.

`L -> 0`. `Lambda -> 0` geometrically. **`L = 0` exactly** gives `Lambda_inf = 0` and no barrier of any size is crossable, at any patience.

`eta -> 0`. `C/eta -> infinity` and `L^M -> 0`: **no conversion, no launch budget, at any moment size.** This is the same boundary the storage law reaches when conversion vanishes [Curzi, 2026a].

`C -> eta*M` from below, and beyond. `L -> 0`, and past it the rectifier holds `L = 0` however much larger `C` grows. **The model cannot distinguish degrees of being trapped past that point**, which is a real insensitivity and is why SP-L3 requires the deficit to be tracked separately.

`E_b -> 0`. `C_crit -> eta*M` and the feasibility condition reduces to `L > 0`: an adjacent move requires only that the subject not be at the floor. **`E_b -> infinity`:** `C_crit -> -infinity`, correctly, since an infinite barrier is uncrossable.

`v -> 0`. `E_b -> infinity`: an infinitely slow crossing costs infinitely much, which is correct and is why indefinite hesitation is not free. **`v -> infinity`: `E_b -> 0`, which is FALSE.** At high transit speed the cost field itself rises; the constant-field model has no term for that, the limit is a known defect, and SP-L12 fences it rather than repairing it.

`theta -> 0`. Deterministic choice: the cheapest option always, and all differences become unrecoverable because every count ratio is zero or infinite. **`theta -> infinity`:** uniform random choice, and again nothing is recoverable. **Recovery requires intermediate temperature**, which is a real and frequently violated precondition.

`n_x -> 0`. The estimate and its standard error both diverge: **a never-chosen option is an unmeasured option, and the instrument says so loudly rather than quietly.**

`s -> infinity`. The current configuration is absorbing and the survey returns "you always do this", which is true and useless.

Stated domain. The construction is used here for `lambda` in `(0, 0.5]`; `eta > 0`; `C < eta*M` for any non-trivial result; transit durations between roughly 40 and 104 weeks for the worked field; `theta` finite and non-zero; and every option's count strictly positive. Outside that domain the equations still return numbers and those numbers should not be read as statements about an subject.

22.11.7 11.7 Invariance, and how far it has been tested

Three invariances are claimed and each is stated as one.

Register invariance. Every sign-level claim about `L` is invariant under the transformation `L^M = L^Q/eta` for `eta > 0`, because a positive scale factor preserves the zero set and the sign. **What breaks it:** any claim about the *magnitude* of `L`, which is not invariant, which is why SP-L1 requires the register to be declared.

Theta invariance. Every ratio of cost differences is invariant under `theta -> c*theta` for `c > 0`, because `theta` appears once in the numerator and once in the denominator. **What breaks it:** a claim about a level or about a single difference, neither of which is invariant; and a switching cost, which shifts the recovered differences and biases the ratio too, though four times less [Narens, 1996].

Coordinate rescaling. Under `x -> c*x`, `grad C` scales as `1/c` and `alpha` as `c^2`, so the observed step is unchanged. **This is why `alpha` is not a free parameter and why SP-G9 requires the coordinate scale to be declared**; it is also the exact statement of the identification failure recorded at FVP-5 in Section 6.5.

How far tested. Register invariance and theta invariance are algebraic and hold exactly. The empirical invariance a reader might want — that the same construction applies at desk, person, team and institution — **is not tested here at all.** The worked case is one subject at one scale, and the corpus's own applications elsewhere are surveyed by one author with no independent replication.

22.11.8 11.8 Precision, and what the numbers are

The numbers in this paper are of three kinds and are quoted accordingly.

Defined or derived-from-defined. `ln 12 = 2.484907`, `ln 2/(−ln 0.94) = 11.20`, `t* = 32.03`, `k = 0.483425732988414`, and every geometric-series limit. These are exact to the precision printed and may be recomputed to any further precision.

Illustrative model output. Every entry in every trajectory, regime, regrade and descent table — `Lambda_inf = 4.320`, `E_b = 2.3987`, `x = 1.9972`, `C'' = −0.005076` — is the output of stipulated parameters under a stipulated field. **These carry no measurement uncertainty because there was no measurement**; the decimal places record the arithmetic, not the confidence.

Measured or fitted. The counts 14 / 9 / 5 / 2 and 17 / 8 / 4 / 1 and 8 / 7 / 8 / 7 are the only observational quantities in the paper, and every quantity derived from them is reported **with a standard error and a 95 per cent interval** in Sections 5.4, 5.6 and 9.8, including the three intervals that span zero and are therefore reported as unmeasured. The parameters `eta`, `M`, `C` and `lambda` are described in the source as measured by the framework's protocols; **no interval is available for any of them, and their point values are therefore carried as stipulations for the purposes of this paper.** That is a real gap and Section 12 repeats it.

22.12 12 Limitations

The worked case is one subject and the field is constructed. No external dataset was collected for this paper. The cost field of Section 5.8 is an analytic expression, the regimes of Section 9.4 are stipulated parameter sets, and every derived figure is the output of those stipulations. They demonstrate that the procedure runs and that it reports its own failures; **they are not evidence that any person behaves this way and must never be cited as such.**

The region that matters was never visited. The ordered coordinate's field was estimated from a mixture of trace data and ledger reconstruction, because the subject had not occupied most of the range. **The region between `x = 3` and `x = 6` — which contains the ridge, and therefore the entire barrier — was inferred, not measured. Every barrier figure in this paper rests on that inference.** The two independent estimates of `E_b` agree to 0.05 per cent, but both were produced by the same investigator and may share the same reconstruction, so the agreement is weak evidence and is reported as such.

Four quantities in the source do not reconcile exactly, and are reported rather than corrected. (i) The basin separation is given as 4.996 configuration-units while the two stated basin floors, 1.9970 and 6.9960, differ by 4.999; at the consistent separation `E_b = 2.3973` rather than 2.3987, a shift of 0.06 per cent that changes no verdict. (ii) After the 0.020 bridge, `C(A)` is 0.032665 at the stationary point and 0.03269 in the regrade table, which propagates to `Lambda_inf` = 4.3734 against 4.3724 and `t*` = 21.50 against 21.49. (iii) The source states that basin B "takes about 2.7 steps per decade of remaining error"; from its own contraction factor of 0.24886 the figure is `ln(0.1)/ln(0.24886) = 1.66`. (iv) The bridge's functional form is never stated, so the figure in Section 5.11 uses a reconstruction that reproduces all five stated stationary points to four decimals and recovers `E_b = 1.8361` against the stated 1.8345.

The Launch Postulate and the Gradient Postulate are postulates and neither has been tested. Nothing in the framework's foundational results implies either, and neither would be touched if both were false. H1 and the falsification protocol of Section 6.1 state exactly what would defeat the first; the shuffled-day control of Section 9.8 licenses the second for one subject on two coordinates and licenses nothing further.

`theta` was never pinned, so no absolute cost is reported anywhere. The one experiment that would fix it — imposing a cost of known magnitude and measuring the shift in share — was not run. Every difference in Sections 5 and 9 is in theta units and only the ratios are advanced as findings.

The independence-of-irrelevant-alternatives check was not performed. It was computed within the existing choice set, which is not a test. The logit form is assumed for coordinates

1 and 2 and is untested for them, and it is known to fail in choice domains containing near-substitutes [Tversky, 1977].

The reflection problem is unanswered and outranks every fitting result in this paper. The framework's reflexive loop is a linear-in-means specification whose endogenous effect is **not identified from observational data** [Manski, 1993], with influence and selection generically confounded in observational network data [Shalizi and Thomas, 2011]. **No remedy is adopted here.** The most realistic available repair is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001], and the frozen pre-declared option set of Protocol G-1 is a step toward it rather than a substitute. **Every claim in this paper that leans on the loop is a modelling position rather than an identified empirical result.**

The Lucas critique applies to this framework's own parameters. The framework asserts that systems change behaviour when modelled and observed, then treats `eta`, `lambda` and `theta` as constants over an estimation window [Lucas, 1976]. The narrow position taken in Section 3.6 — stability only over a declared window and only while the fund is not itself a reported target — is a genuine restriction on where the model may be used, and where an subject's fund is reported to anyone with an interest in it the constancy assumption fails [Goodhart, 1975, Campbell, 1979].

The quantitative structure of the underlying constructs has not been tested, and this is the strongest objection to the framework. Michell [1997, 1999] argues that psychology routinely fails to test whether its attributes are quantitative at all before measuring them. The available remedy is conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971] or a Rasch demonstration [Rasch, 1960]. **Neither has been run for comfort, for `M`, or for `C`, so the ratio-scale treatment those quantities receive here — differences, ratios, logarithms and a rectifier requiring a true zero — is a stipulation and not an established property.** The rectifier is the most exposed operation in the paper on this ground: `max(0, x)` is only meaningful if zero means something.

Several parameters were chosen rather than measured with an interval. `eta`, `M`, `C` and `lambda` are reported as point values with no uncertainty attached, and `lambda` in particular was measured by a reactive procedure: stating a plan and then deliberately not reinforcing it is not a neutral act, so the two values should be read as an ordering rather than as magnitudes. The transit duration of 65.5 weeks is an intention, not an observation.

The model comparison standard is NOT MET. RIGOR requires out-of-sample performance, a matched null with the same number of free parameters, an information criterion so the parameter count is paid for, an explicit count of free parameters, and effect sizes with intervals. The shuffled-day control of Section 9.8 supplies a permutation null and nothing else. **No AIC or BIC was computed anywhere in this paper, and none is computed anywhere in this corpus**, so the parameter count has never been paid for. The free-parameter count of the worked construction is stated for the record: the cost field carries six (two depths, two centres, two widths, plus a linear drift term, of which the drift is the seventh), the launch model carries four (`eta`, `M`, `C`, `lambda`), the barrier carries one (the transit duration), and the logit recovery carries one unidentified scale (`theta`) plus one optional switching cost (`s`). **A model with that many free parameters that reproduces a single subject's trajectory has not been shown to be better than anything.**

No scaling, power-law or heavy-tail claim is made in this paper, so the distribution-fitting procedure such a claim would require [Clauset et al., 2009] does not apply. The one distributional claim that *is* made — that plan persistence decays geometrically — is H2 and is untested.

The functional form of the surplus term is underdetermined, as Section 3.1 states: a sub-linear power law [Latané, 1981] and an `n log n` form [Briscoe et al., 2006] fit the same qualitative facts as `ln(1 + (N−1)u)`, and **nothing assembled in this corpus discriminates among them.** The logarithm is preferred for its axiomatic characterisation and for the ergodic argument [Peters and Gell-Mann, 2016], not because it has been measured. Relatedly, `1 + (N−1)u` **is not a Hill number** — `1/HHI` and `exp(H)` are, and are well grounded [Hill, 1973, Jost, 2006] — and a reviewer reaching for the Hill literature as support will find that it does not apply.

The moment equation and the calibration ratio have no external support in this corpus's

bibliography. They are used in Section 4.9 to compute the over-framing inflation factor of 1.7889, and that number is therefore internal to the framework and carries no external warrant.

The single-scalar barrier discards structure that matters. A real transition has a shape, a sequence and reversibility properties that one number cannot carry; the barrier-splitting result of Section 4.8 is precisely the case where the scalar summary was wrong and the shape was what mattered. The scalar is retained because it makes feasibility a comparison of two numbers, and that is a modelling choice rather than a finding.

Nothing here is a human-subjects finding. No interventions were run and no participants were recruited. The consent constraint of Section 4.7.3 is a design constraint on prospective work, not a report of practice.

What would falsify the account. H1 fails on a documented plan constructed during a sustained `L = 0` stretch. H2 fails on non-geometric plan decay under no reinforcement. H3 fails if departure rates are the same above and below the computed `C_crit`. H4 fails if realised transit cost is uncorrelated with the predicted barrier. H5 fails if a frozen difference table does not beat a day-shuffled control, or if experimentally manipulating the choice set moves the recovered ratios. H6 was tested and survived on one constructed field, which is the weakest form of survival available.

22.13 13 Conclusion

R1. The launch budget is the rectified surplus `L^Q = max(0, eta*M - C)`, with the lived register `L^M = max(0, M - C/eta)` and `L^Q = eta*L^M`; the naive form `max(0, M - C)` fails a dimensional check and is inadmissible. Accumulated in a leaky integrator `Lambda_{t+1} = (1-lambda)Lambda_t + L_t`, it has the unique attracting fixed point `Lambda_inf = L/lambda`, which is the largest barrier the regime can cross at all, at steady state, with unlimited patience. **A departure across a barrier `E_b` is feasible if and only if `Lambda_inf > E_b`, and it requires `t* = ln(1 - E_b/Lambda_inf)/ln(1 - lambda)` periods of accumulation.**

R2. Rearranging feasibility gives the **critical cost `C_crit = eta*M - lambda*E_b`**, computed at **0.082 Q-units per week** in the worked case at `eta = 0.10`, `M = 1.42`, `lambda = 0.025`, `E_b = 2.40`. Above it no amount of waiting funds the departure; the fund still rises, to a ceiling below the wall, which is why the failure is silent. Rearranging `L > 0` gives the **critical conversion efficiency `eta_crit = C/M`**, computed at **0.055238** against the worked subject's measured `eta = 0.05` — below the line by half a percentage point, giving `L = 0` and `Lambda_inf = 0`: **a sustained zero budget that cannot diagnose itself, because diagnosis is funded from the same budget.**

R3. On the measured cost field `C(x) = 0.090 - 0.0568 exp(-(x-2)^2/0.8) - 0.064 exp(-(x-7)^2/1.2) + 0.0004x`, with basin A at `x = 1.9970`, a ridge at `x = 4.3350` with `C'' = -0.005076`, and basin B at `x = 6.9960`, **momentum at `beta` = 0, 0.5, 0.8, 0.9 and 0.95 all terminate at `x = 1.9972` in basin A.** Momentum defeats narrow barriers and is useless against wide shallow ones, because momentum accumulated on a one-unit descent cannot carry a walker across a two-unit flat rise. **The recursion is recorded as a failed candidate and is refused entry to the formulary**, and what survives is one empirical statement with no free parameters: against a wide shallow barrier there is no substitute for an accumulated fund, because no dynamical trick crosses it.

R4. Under a logit desire-line survey **only ratios of cost differences are identified**, `theta` cancelling: from counts 14 / 9 / 5 / 2 over thirty days, **`(C_office - C_desk)/(C_kitchen - C_desk) = 4.4042`**, unit-free and testable, while the kitchen-versus-desk difference itself has a 95 per cent interval of `[-0.3956, 1.2792]` and is **unmeasured** at 30 days, needing 108. Absolute costs are never recovered and must never be reported.

R5. Regrading the destination is nearly useless: a 21 per cent improvement in the target basin's floor buys a 1.94 per cent reduction in the barrier. Bridging the ridge by 0.020 cuts the barrier by 23.5 per cent, cuts the wait from 32.00 weeks to 21.49, and manufactures a new rest

basin at `x = 4.3350`, `C = 0.071500`, which splits the crossing into legs of 0.8812 and 0.0218 and cuts required accumulation to 9.19 weeks. **The barrier is the ridge, not the destination.**

R6. An unmodelled switching cost `s = 0.6` theta inflates every non-modal option's recovered cost by 15.9 to 20.8 per cent while biasing the ratio of differences by only 4.05 per cent, and at `s = 1.0` the figures are 28 to 36 per cent against 5.36 per cent. Reporting ratios rather than differences therefore buys roughly a factor of four in robustness, on top of the theta-freeness that was the original reason for the rule.

CR (counter-result). Momentum, the standard escape from a local minimum, was run on the measured field at five settings and escaped nothing. **A measured negative result, recorded without adopting the mechanism**, and it sharpens R1: no dynamical trick substitutes for an accumulated fund against a wide shallow barrier.

The arithmetic behind all six is one line. `L = max(0, eta*M - C)` is maximised exactly where `C` is minimised, so **the budget for elsewhere is manufactured only where the cost of being here is at a minimum.** Whether any person's planning capacity is in fact `L` is a postulate, not a result, and it has not been tested.

22.13.1 Declarations

Data and code availability

No external dataset was collected for this paper. The worked case of Section 9 is a single-subject construction whose parameters are stipulated in the text and whose derived figures are computed from those parameters under the model developed in Sections 3 to 5. The choice counts of Sections 5.4, 5.6 and 9.8 are the source's recorded tallies and are printed in full beside every quantity derived from them. Every table in the paper can be re-derived from the equations and parameter values printed beside it. **The figures are generated by the scripts named in each caption**, and those scripts contain no data beyond the numbers printed in this paper and the model equations of Sections 3 to 5. Where a figure required a functional form the source does not state — the bridge regrade of Section 5.11 — the reconstruction is named in the caption and its fidelity to the source's stated stationary points is reported.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, produced the worked example and the figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The worked case is a construction. The consent constraint discussed in Sections 4.7.3 and 7 — that no intervention raising another person's operating cost in order to force a departure is within the framework's scope — is a design constraint on prospective work and is not a report of any study performed.

Preregistration

This paper reports no preregistered test of its own. The falsification protocols for H1 to H5 are stated in Sections 2 and 6 and none has been run. The one preregistered confirmatory test of the underlying storage law reported in this corpus **failed** its second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared margin of 10 per cent — and is reported in Curzi [2026i]. Any applied claim resting on this paper inherits that result.

Chapter 23

The Return Difference, the Contraction, and a Formulary That Demotes Five of Its Own Nine Entries

Abstract. This paper defines the return difference $\Delta_R = |\ln(M'/M)|$ between a moment enacted and the world's answer, badging the commensurability assumption (Postulate R1) that lets the two be placed on one ratio scale. The instrument has a Weber deadband $\Delta_{R,\min} = \ln(1+c)$, a sample-size law $n \geq (1.96\,\sigma_\rho/\Delta)^2$ requiring 139 clean echoes to establish a 5 per cent return at $\sigma_\rho = 0.30$, and an attenuation law $\rho_{\text{obs}} = (1-w)\rho_{\text{true}}$ for correlated sources that biases the reading toward zero rather than merely widening it. Correction with gain λ gives $\rho_{t+1} = (1-\lambda)\rho_t$, stable for $\lambda < 2$, deadbeat and noise-transparent at $\lambda = 1$. In a twelve-week log the sign read over-framing at an implied factor of 1.4467, the gain recovered as $\lambda = 0.3582$ ($R^2 = 0.8587$, distortion half-life 1.56 weeks), and directional efficiency was $E_R = 0.9316$. Arrival is then proved rather than asserted: $T(Q) = (1-\delta)Q + b$ is a contraction with constant $1-\delta$, Banach gives the unique fixed point $Q_\infty = (\eta M^* - c^*)/\delta$, and the exact error law $|Q_t - Q_\infty| = (1-\delta)^t |Q_0 - Q_\infty|$ coincides with the closed-form solution. Elasticities rank the levers and are derived rather than fitted: $\eta\,{+1.594}$, $\delta\,{-1.000}$, $u\,{+0.770}$, $C\,{-0.594}$, so a 10 per cent efficiency gain is worth 2.7 times a 10 per cent cost cut. Convergence, however, certifies nothing: $\eta F < c^*$ yields a negative fixed point the loop still reaches, and a transfer relocates a fixed point without creating one. **The accompanying nine-entry formulary demotes five of its own entries — F-01, F-02, F-04, F-06 and F-08 — on the scale-type gate, withdrawing the magnitude of D/P, the distortion as a number, and the quantitative moment-to-storage join, while retaining the limiting cases, the $D = P$ fixed point and the over/under-framed ordering; and it records that the shuffled-input control every falsification row depends on is anti-conservative at 0.108 against a nominal 0.05.**

Keywords: return difference; correction gain; contraction mapping; error law; elasticity; counterfeit arrival; formulary demotion

Series: B. Engineering the Surplus **Paper:** CCP-023

23.1 1 Introduction

An account of self-directed change has to answer two questions before any of its prescriptions can be assessed. First: when a person acts and the world answers, what quantity is the answer,

and how small an answer is still an answer? Second: if the person keeps acting, does the accumulation arrive anywhere, and if so, where — and does arriving mean anything?

This paper answers both, and then audits the answers with the corpus's own validation instrument. The first question is answered by the **return difference** $\Delta_R = |\ln(M'/M)|$, the logarithm of the ratio between the moment enacted and the moment returned, read sign-first. The second is answered by showing that the internal-scope storage map `T(Q) = (1−delta)Q + b` is a contraction with modulus `1 − delta`, so Banach's theorem supplies a unique start-independent fixed point `Q_inf = (eta M* − c*)/delta` together with the exact error law `|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|`.

Both answers are then run through the seven-gate Formula Validation Protocol [Curzi, 2026t], and **the protocol demotes five of the nine entries in the resulting formulary.** That is the paper's most important result and it is reported here rather than buried: F-01, F-02, F-04, F-06 and F-08 — every entry whose value passes through the moment `M` or through the calibration ratio `r = D/P` — fail the scale-type gate, because `D` and `P` are not instrumented on a common ratio scale with a shared true zero and division without a true zero yields a number with no referent [Curzi, 2026m,t]. What is withdrawn is stated exactly: the magnitude of `D/P`, the distortion $\Delta = |\ln r|$ as a number, and the quantitative form of the moment-to-storage join. What survives is stated with equal precision: the limiting cases, the `D = P` fixed point, and the ordering over-framed / calibrated / under-framed, all of which require only ordinal information. **A formulary that demotes five of its own nine entries, names exactly what survives and prices the study that would restore the rest is the protocol working, not a caveat.**

A second self-audit runs alongside it and is equally uncomfortable. The four-baseline falsification suite this paper's entries all defer to rests on a shuffled-input control, and that control has been measured rather than assumed. Against a nominal false-positive rate of 0.05 across 500 null trials, independent permutation false-rejected at **0.108, 2.2 times nominal**; it is blind to a filtered repackaging of its own input, which passed all four baselines at `p = 0.0005` and was caught only by a structural ceiling test; and against a confound driving both series it passed **98.35 per cent** of 4,000 trials [Curzi, 2026u,ac]. **The shuffle is a weak, mis-calibrated empirical audit, and design controls rather than more surrogates are what catch confounding.**

Between the two audits sits the paper's constructive content. The return difference has a Weber deadband $\Delta_{R,\min} = \ln(1+c)$ below which no echo is detectable at all; a sample-size law `n >= (1.96 sigma_rho / Delta)^2` requiring 139 clean echoes to establish a 5 per cent return at `sigma = 0.30`; and an attenuation law `rho_obs = (1−w) rho_true` for correlated sources, which is the design-effect problem of survey sampling arriving in a personal instrument. Correction with gain λ gives `rho_{t+1} = (1−lambda) rho_t`, stable for $\lambda < 2$, deadbeat and noise-transparent at $\lambda = 1$. On the worked twelve-week log the sign read over-framing at an implied factor of 1.4467, the recovered gain was $\lambda = 0.3582$ with `R^2 = 0.8587` and a distortion half-life of 1.56 weeks, and directional efficiency was `E_R = 0.9316`.

The arrival half then answers the practical question — which of the four levers to move — with elasticities of the fixed point that are **derived rather than fitted**: `eta` **+1.594**, `delta` **−1.000**, `u` **+0.770**, `C` **−0.594**, so that **a 10 per cent efficiency gain is worth 2.7 times a 10 per cent cost cut** at the stated base point. And it closes with the result that convergence certifies nothing: `eta F < c*` produces a negative fixed point that the loop reaches just as reliably, so a system can arrive, on schedule, at a destination that is a slow depletion.

Three results are presupposed rather than re-derived. The storage law's contraction, its fixed point and its geometric convergence — and the finding that `delta = 0` is a pathology rather than an ideal, with no contraction, no unique fixed point and an error-persistence column of ones — are established in Curzi [2026a]. The contraction theorem itself is stated exactly, with its *a posteriori* bound read operationally as *distance to arrival is the last correction scaled by* `q/(1−q)`, in Curzi [2026]. The stationary-action framing, with the correction that the honest form is stationarity rather than minimality and that the variational form licenses no teleology, is presupposed from Curzi [2026].

Sections 3 to 6 develop the return difference, its requirements and its failure modes. Sections 7 to 10 develop the contraction and the dynamics of arrival. Sections 11 to 13 present the

nine-entry formulary with its per-gate verdicts. Section 14 reports the three worked results and the demotion tally. Sections 15 to 18 discuss, verify, limit and conclude.

23.2 2 Hypotheses

Several of the results below are theorems about a difference equation and cannot be falsified by observation; they are labelled as derivations where they appear. What can be falsified is the claim that the equations describe a real person acting and accumulating. The hypotheses separate the two, and each names the observation that would defeat it.

H1 — The return is a ratio, not a difference. The smallest change in a returned magnitude that a person can detect is a constant *proportion* of what was sent, not a constant increment. **Predicted:** a Weber deadband $\Delta_{R,\min} = \ln(1 + c)$ with `c` approximately constant across the usable range of `M`, so that the detection threshold scales with the stake. **Falsified by:** a measured detection threshold that is constant in absolute units across a wide range of `M`, or a Weber fraction that varies systematically with `M`. **Status: NOT TESTED.** No Weber fraction for moment magnitudes has been measured by anyone, and Section 4.3's illustrative `c = 0.10` is a placeholder carried through the arithmetic, not a measurement.

H2 — Arrival is geometric and start-independent. Under constant drive the internal store approaches a single fixed point from any starting value, with the distance falling by exactly the factor `1 - delta` each period. **Predicted:** `|Q_t - Q_inf| = (1-delta)^t |Q_0 - Q_inf|` exactly, a time constant $\tau = -1/\ln(1 - \delta)$, monotone approach from either side, and no overshoot in the absence of a delayed control loop. **Falsified by:** an observed non-monotone approach without a delay term; a settling time that depends on `Q_0`; or a fitted retention coefficient outside `[0,1]`.

H3 — Correction is a first-order loop with a stability boundary at gain 2. Repeated correction of a distortion by a fraction λ of the observed error contracts for $0 < \lambda < 2$ and diverges outside it. **Predicted:** `rho_{t+1} = (1-lambda) rho_t`, so distortion decays with half-life $\ln 2/|\ln|1 - \lambda||$; deadbeat at $\lambda = 1$; oscillatory but convergent for $1 < \lambda < 2$; divergent at $\lambda \geq 2$. **Falsified by:** an observed correction sequence that converges at a gain of 2 or more, or that diverges at a gain safely inside the interval.

H4 — Efficiency dominates cost as a lever. At any sustainable operating point, a proportional improvement in conversion efficiency raises the destination by more than the same proportional reduction in cost does. **Predicted:** the elasticity ratio is `eta M*/c*`, which exceeds one whenever `eta M* > c*` — that is, whenever the regime is sustainable at all — and at the base point of Section 8.5 it is 1.594 against 0.594, a factor of 2.68. **Falsified by:** nothing observational. **This is a theorem about the fixed-point formula, not a hypothesis**, and it is recorded in the verification chain as `ARITHMETIC`. It is stated here as a hypothesis only to mark the boundary: what *could* fail is the antecedent claim that a real person's stored capability obeys the fixed-point formula at all, and that is H2.

H5 — The protocol demotes some of what it audits. A validation protocol applied by an author to his own formulary returns a pass rate materially below one. **Predicted:** a substantial minority of the nine entries fail at least one gate outright, and the failures cluster on the gate the corpus is most exposed at — scale type. **Falsified by:** all nine entries passing all seven gates, which would be evidence that the protocol had been captured rather than that the formulary was sound.

H2, H3 and H5 are confirmed; H4 is a theorem; H1 is untested and is reported as untested. The confirmations of H2 and H3 are confirmations *on illustrative constructions* — worked

series whose parameters were chosen, not measured — and they demonstrate the arithmetic rather than the world. H5's confirmation is the demotion of five of nine entries, and that one is a fact about the corpus rather than about people.

23.3 3 The return difference: definition, badge, and what is presupposed

You act. The world answers. The answer is never the voice you sent.

That last sentence is the whole of the argument, and it is not a lament. It is a statement about where information lives. **If the answer were identical to the voice, the exchange would have carried nothing; the entire content of an echo is the difference between what returned and what was sent.** Everything else in this paper is the arithmetic of that difference and the engineering of the instrument that reads it.

23.3.1 3.1 The formula, stated

The fourth of the framework's derived internal-scope formulas [Curzi, 2026o] is this:

```
THE RETURN DIFFERENCE

    rho_t     = ln( M'_t / M_t )              signed return, in nats
    Delta_R,t = | rho_t | = | ln(M'_t / M_t) |   unsigned magnitude, in nats

    M_t   the moment sent      -- the moment as the subject constructed and enacted it
    M'_t  the moment returned  -- the moment the world's answer constructs on arrival back
```

Two quantities, not one. `Delta_R` is what the framework's internal-scope correspondence names; `rho` is `Delta_R` with its sign restored, and the sign is half the diagnostic. The corpus uses both throughout, and the convention is fixed here: **`rho` for reading, `Delta_R` for budgeting.** You read a single echo with its sign. You budget a series of echoes by magnitude, because magnitudes accumulate and signs cancel — and the cancellation of signs across a series is itself a finding, treated in §4.4.

23.3.2 3.2 The badge

`Delta_R = |ln(M'/M)|` is **a derived formula of the corpus, not a theorem of the framework's mathematical papers.** It is derived in the sense the internal-scope correspondence requires [Curzi, 2026o] — it follows from choices already made in the corpus rather than being introduced fresh — but the derivation runs through a postulate, and the postulate is stated here with its badge on.

Postulate R1 (the returnable moment). The world's answer to an enacted moment can itself be assigned a moment magnitude on the same ratio scale as the moment that was sent, by the same subject, using the same procedure. **Badged: this is an assumption of commensurability, not a result.** It is exactly as strong as the assumption that lets a person say a second sound was twice as loud as the first, and it is defensible for the same reason — Stevens' method of magnitude estimation produces ratio-scale judgements of subjective quantity and it replicates [Curzi, 2026t]. It is not defensible when the two things being compared have no common procedure of estimation, and §4.8 says exactly when that happens.

Given R1, the rest is forced. The corpus works in log-ratios everywhere: surplus is `ln N_eff`, pitch is `log2 f`, sensation is `k ln(I/I_0)`, the calibration distortion is `|ln r|`. A difference measure on moments that was *not* a log-ratio would be the one quantity in the book that did not compose with the others, and the four-register convergence [Curzi, 2026n,?] would break

at exactly the point where feedback enters. **The log-ratio is not a stylistic choice about the echo; it is the only form that keeps the echo in the same register as everything the echo has to talk to.**

23.3.3 3.3 Why the ratio and not the difference

There are two candidate difference measures and the choice between them is decidable, not aesthetic.

```
additive     e_t = M'_t − M_t
ratiometric  rho_t = ln(M'_t / M_t)
```

Three reasons the ratiometric form wins, in increasing order of force.

One: composition. Echoes chain. You send a moment, the world answers, you send the answer onward, a second answer returns. The total return over a chain of `k` links is

```
M'_k / M_0 = (M'_1/M_0) ∗ (M'_2/M'_1) ∗ ... ∗ (M'_k/M'_{k−1})

ln(M'_k / M_0) = rho_1 + rho_2 + ... + rho_k
```

Ratios compose by multiplication, so log-ratios compose by addition, and a chained return difference is therefore a plain sum. The additive form has no such property: `M' − M` over a chain depends on the absolute magnitudes at every intermediate link, and two chains with the same per-link proportional effect give different totals if they start at different sizes. A measure whose composition law depends on where you started is not measuring the exchange.

Two: scale invariance. Multiply every magnitude in an echo log by a constant — change the modulus of the magnitude-estimation procedure, rate the week out of 100 instead of out of 50 — and every `rho` is unchanged while every additive `e` is scaled. Since the modulus in magnitude estimation is arbitrary by construction, **a difference measure that moves when the modulus moves is measuring the procedure, not the world.**

Three: the lens. The lens theorem [Curzi, 2026n] holds that perception is logarithmic, so equal perceptual steps are equal world ratios. The subject reading the echo is a logarithmic instrument. `rho` is what that instrument natively reports. The additive form would require the subject to perform a conversion their perceptual system does not perform, and every such conversion is a place where a bias enters.

23.3.4 3.4 The ECHO node, stated exactly

The foundations series' eighth node [Curzi, 2026] taught the practitioner to read the difference: that the answer coming back is never the voice sent out, that the gap between them is the only new thing in the exchange, and that two specific pathologies wait on either side of the reading — the Mirror that Flatters, which returns your own voice and calls it agreement, and the Echo that Devours, which returns your voice amplified until it is the only thing in the room.

Here the ECHO node is `rho_t = ln(M'_t/M_t)`, its deadband, and its stability condition. Each of the node's three teachings has an exact counterpart:

The node's teaching	This paper's quantity	Section
The answer is never the voice sent	`rho` is generically nonzero; `rho = 0` is a measure-zero coincidence	§4.1

The node's teaching	This paper's quantity	Section
The difference is the only data	The magnitude `Delta_R`, its deadband `ln(1+c)`, and its signal-to-noise	§4.2–4.4
The Mirror that Flatters	`rho_obs = (1 − w) rho_true`; correlated echoes and the design effect	§4.5
The Echo that Devours	`rho_{t+1} = (1 − lambda) rho_t` with `|1 − lambda| >= 1`	§4.6

The last row is the one that repays the translation most, because the foundations series could warn about the devouring echo but could not say where it starts. This paper can: **it starts at a correction gain of exactly 2, and the marginal case at gain 2 is an undamped oscillation that never grows and never settles.** That is a number, it is derived below, and it is checkable against your own records.

23.3.5 3.5 The zero clause, stated hard

The internal-scope correspondence [Curzi, 2026o] says it in six words: *zero means no information came back.* The corpus states the consequence at full strength because it is the most commonly ignored reading in the whole formulary.

`Delta_R = 0` is not success. It is the absence of an exchange. A world that returns exactly what you sent has told you nothing about itself — it has told you only that whatever it did to your voice, it did symmetrically enough to cancel. There are two ways to produce a persistent zero and they are opposite in kind: perfect calibration, in which the subject's construction of the moment already matched what the world would return, and perfect closure, in which the channel back is not the world at all but a reflection of the subject. **The arithmetic cannot tell these apart from a single reading, and §4.5 gives the test that can.**

The corollary is the one that costs people years: **a world that only ever agrees with you is telling you nothing, and the pleasantness of the sensation is not evidence against this.** The sensation of being agreed with and the sensation of being correct are, from the inside, the same sensation. The corpus has said the equivalent thing at group scope already — [Curzi, 2026d] on pseudo-diversity and on collective intelligence — and the thesis of Series B [Curzi, 2026s] is that it is one statement at two scopes, not two statements.

23.3.6 3.6 Formula Validation Protocol — the return difference

Under the Formula Validation Protocol [Curzi, 2026t], no formula appears in Series B of the corpus without all seven gates in print.

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`M` and `M'` carry the units of the magnitude-estimation modulus (arbitrary but fixed). `M'/M` is dimensionless. `ln` of a dimensionless number is dimensionless; `rho` and `Delta_R` are reported in **nats**, convertible to bits, octaves and cents by the the framework's internal-scope correspondence keys. Both sides agree.
FVP-2 SCALE TYPE	**CONDITIONAL PASS — the binding gate**	Division requires ratio scale with a true zero. Magnitude estimation supplies this (Stevens 1957); a 1–5 satisfaction rating does not. **If `M` and `M'` are ordinal, `ln(M'/M)` is not defined and the formula must be demoted in print to a heuristic ordering: "the return was larger / the same / smaller."** The demotion is not optional and §4.8 specifies it.
FVP-3 LIMITING CASES	PASS	`M' -> M`: `rho -> 0`, no information returned (§3.5). `M' >> M`: `rho -> +inf`, over-return; in practice bounded by the modulus ceiling. `M' -> 0`: `rho -> -inf`, total non-return — the answer that constructs no moment at all, which is the real phenomenon of being ignored, and the divergence is correct because there is no finite penalty scale for it. `M -> 0`: undefined; correctly so, because nothing was sent and an echo of nothing is not an echo.

Gate	Verdict	Content
FVP-4 FIXED POINTS	PASS	`rho = 0`, i.e. `M' = M`. Under the correction dynamics of §4.6 this is the unique fixed point of the map `rho -> (1-lambda)rho` for any `lambda != 0`, and it is attracting iff `0 < lambda < 2`. Meaning: **the state in which your construction of the moment and the world's answer to it agree — which is either contact or closure, and only the correlation test distinguishes them.**
FVP-5 IDENTIFIABILITY	**PARTIAL**	`rho` is recoverable from paired magnitude estimates. The *decomposition* of `rho` into "the world's contribution" and "the subject's mis-framing" is **not identified from `rho` alone** — a large negative `rho` is equally consistent with an over-framed subject (`r = D/P > 1`, [Curzi, 2026m]) and with a genuinely unresponsive world. Separating them requires a second instrument: either an independent estimate of `M'` by someone who did not send `M` (§4.5) or a manipulation of `D` with `E` held fixed (§4.10). **State this limitation whenever a single `rho` is used to draw a conclusion about a person.**

Gate	Verdict	Content
FVP-6 FALSIFICATION	SPECIFIED	The preregistration template [Curzi, 2026i]. Preregister: the operationalization of `M` and `M'`, the modulus, the correction gain `lambda`, the horizon. Fit `lambda` on a training block, **freeze**, and predict held-out `rho`. Four baselines: (i) persistence — next `rho` equals last `rho`; (ii) direct-input — `rho` predicted from `M` alone; (iii) matched autoregressive AR(1) on `rho`; (iv) **shuffled-input control** — the same `M` series paired with a permuted `M'` series. **The formula dies if the frozen `lambda` model fails to beat the shuffled-input control on held-out data**, because that is the outcome in which the pairing between what was sent and what came back carries nothing.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	The cross-individual gate. Same operationalization, different subjects, different investigators: do the fitted `lambda` values recover, and do they rank people consistently? The corpus has not run that study and does not have the data. **What would need measuring is stated in §4.11 as a protocol, and until it is run, `lambda` is a within-person descriptive parameter and nothing more.** This is the gate the framework's internal-scope correspondence says is most often skipped; it is not skipped here, it is declared open.

23.3.7 3.7 What is presupposed

Four results are used below and not re-derived.

The **surplus functional** `F(u) = ln(1 + (N−1)u)` with its axiom set and its ceiling `F <= ln N` is established in Curzi [2026]. The **affine effective count** `N_eff = 1 + (N−1)u` that goes with it is

the corpus's own construction and is not a Hill number of any order — it is a linear interpolation between 1 and `N` with no supporting literature in ecology, economics, political science or information theory. The Hill numbers proper, `1/HHI` at order 2 and `exp(H)` at order 1 [Hill, 1973, Jost, 2006], are different quantities and are the ones the corpus uses on measured shares at institutional scope [Curzi, 2026,?]. Wherever an effective count appears below it is the affine one.

The **storage law** `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t`, its exact solution, its uniform bound and its fixed point are established in Curzi [2026a], whose central result — that `delta = 0` is a pathology and not an ideal, with no contraction, no unique fixed point, and an error-persistence column of ones, so that no mistake ever decays out of the stock — is the boundary case Sections 7 and 8 work away from. The **contraction theorem** is stated exactly in Curzi [2026], together with the *a posteriori* bound `d(x_n, x*) <= (q^n/(1−q)) d(x_0, x_1)` read operationally as: distance to arrival is the size of the last correction scaled by `q/(1−q)`, so the size of one's most recent revision is a self-reading gauge.

The **moment equation** `M = (D/P)E`, the calibration ratio `r = D/P`, and the correction map `ln r_{t+1} = (1−a) ln r_t` — attracting for `0 < a < 2`, neutral at `a = 0`, divergent at `a >= 2`, so that from a distortion of 0.81 nats six cycles at gain 2.3 reach 3.92 — are established in Curzi [2026m]. The **stationary-action framing**, with the correction that the honest form is stationarity rather than minimality and that the variational form licenses no teleology, is presupposed from Curzi [2026].

Finally, the **seven-gate Formula Validation Protocol** and its verdict vocabulary are Curzi [2026t], whose gate-2 result is the one this paper's formulary inherits.

23.3.8 3.8 Identification: what no amount of logging can recover

The material below repeatedly describes a loop: a person enacts a moment, the world answers, the answer changes the construal, the changed construal shapes the next moment. **That is a reflexive specification, and Manski [1993] proves that the endogenous component of such a specification is not identified from observational data.** It cannot be separated from contextual effects or from correlated unobservables however long the log runs. Shalizi and Thomas [2011] establish the same for network data — influence and selection are generically confounded — and Lyons [2011] show in detail what happens when the point is ignored.

The known remedies are randomisation of the enacted moment, an instrument, partial-population variation, or stipulating the reference set by prior information and defending the stipulation, Brock and Durlauf [2001] being the template for the last. **None is free and this paper adopts none.** Every claim below that rests on the loop is therefore a **modelling position rather than an identified empirical result**, and the specific consequence for the instrument is stated where it bites: a twelve-week echo log cannot distinguish an effect of the enacted moments from a driver common to both series, and the shuffled-input control that is nominally supposed to make that distinction **passed 98.35 per cent of the time against exactly such a confound** [Curzi, 2026u,ac].

Two objections in the same family are recorded now rather than in a footnote. The **Lucas critique** turned on the corpus says `delta`, `eta` and `u` are the decision rules of people who have not read the framework, so they are not structural parameters [Lucas, 1976]; Hacking [1995] makes the same point about classifications of people generally, and Ferraro et al. [2005] supplies the case in which a theory made people behave worse by assuming the worst of them. And **the corpus proposes to give people laws to manage their own comfort by, which is building an engine and not only a camera** [MacKenzie, 2006, Callon, 1998]; MacKenzie and Millo [2003] is the documented case of a theory becoming true by being used, with dates and price data, and Soros [2013] is the closest existing named doctrine to the corpus's reflexive law. Jussim and Harber [2005] is the corrective on effect size: self-fulfilling prophecies are real but small, typically non-accumulating, and often dissipating.

23.3.9 3.9 Analogy discipline

The framework borrows form from control theory, from signal processing and from psychophysics. Every such borrowing is one of three things and is labelled at first use: a **derivation**, where the system genuinely satisfies the premises of the physical result; a **shared mathematical form**, where the same equation governs both with no claim that the underlying physics applies; or an **analogy**, legitimate as motivation and never as evidence.

Everything in this paper is of the second kind. A leaky integrator is a leaky integrator whether it holds charge or stored capability [Oppenheim et al., 1996, Horowitz and Hill, 2015]; the stability boundary at gain 2 is the unit-circle condition of a first-order discrete loop and carries no dynamical-systems result beyond it; the design-effect correction is survey sampling; the ratio register is Fechner's integration of Weber's law [Weber, 1834, Fechner, 1860, 1966] and not a statement about energy. Where a stronger, non-psychological warrant for the logarithm is wanted, Peters and Gell-Mann [2016] derives one from multiplicative dynamics with no perceptual assumption at all, and the argument here needs that one rather than Fechner's.

Two limits on the borrowing are stated now. First, `Q` is a stock that accumulates, leaks and is spent — the grammar of a conserved quantity — and Mirowski [1989] shows that importing that grammar without the conservation law is the characteristic failure of this genre. **The only conservation this paper actually establishes is the transfer identity of Section 8.7, which holds for two systems sharing one `delta` and for nothing else**, and Cartwright [1999] is the escape route the framework should take: specify the arrangement within which each law is claimed to hold. Second, **no scaling, power-law or heavy-tail claim is made anywhere in this paper**; where such a claim is made elsewhere in the corpus the procedure of Clauset et al. [2009] is binding and nothing here substitutes for it.

The functional form of the drive is also underdetermined and this is not repaired below. Latané [1981] proposes influence growing as a sub-linear power of the number of sources and Briscoe et al. [2006] proposes `n log n`; both are decelerating, both fit the qualitative facts, and Stumpf and Porter [2012] states the general problem that many mechanisms produce indistinguishable curves. **Nothing in this paper discriminates the corpus's logarithm from its two nearest rivals.**

23.3.10 3.10 Standing honesty clauses

Three clauses govern every number below and are stated once. A note on terminology precedes them: throughout this paper **subject** denotes the modelled individual — the person whose enacted moments, echoes and stored capability the equations describe. Every subject appearing below is a construction with declared parameters, and no person was studied.

Illustrative construction is said, every time. The twelve-week log, the trajectory tables, the transfer case and the nine-entry review are computed from parameters that were chosen. They demonstrate the arithmetic. They are not observations, no person was studied, and no sentence below should be read as evidence about people.

Significant figures follow epistemic status. Quantities exact by construction — the elasticities, the transfer identity, the deadband formula, the error law — are quoted to whatever precision is useful and said to be exact. Quantities estimated from a constructed series carry a standard error or an `R^2` and are labelled as estimates *of a property of the construction*. Stipulated parameters are named as stipulated at first use and everything computed from them inherits that status.

A gate not run is reported as SPECIFIED or UNTESTED, never as passed. Across the nine entries of the formulary, **sixty-three gates are reported and not one is marked PASS on the reproducibility gate**, because no cross-individual study of any of these quantities has been run.

23.4 4 Reading and instrumenting the echo

23.4.1 4.1 Reading one echo: sign first, magnitude second

A single echo yields two readings and they answer different questions.

```
rho > 0   OVER-RETURN    the world answered larger than you sent
rho = 0   NULL RETURN    no information came back (contact or closure -- test required)
rho < 0   UNDER-RETURN   the world answered smaller than you sent
```

The sign says which way you were wrong. The magnitude says how much it cost to find out.

Over-return. You sent a moment and got back more than you constructed. The honest readings, in order of frequency: your frame under-called the environment (`r = D/P < 1`, under-framed in the sense of [Curzi, 2026m]); or the environment contained structure you had not modelled; or the answer was amplified in transit by something with its own energy — an audience, an institution, a grievance. **Over-return is the pleasant failure and it is still a failure of calibration.** The specific danger is that over-return is usually attributed to the quality of what was sent, which is the one explanation the measurement cannot support.

Under-return. You sent a moment and the world answered smaller. Readings: your frame over-called the environment (`r > 1`, over-framed — projection, expectation, the seen-what-was-predicted failure); or the channel attenuated; or the environment genuinely had less in it than you built into your construction of it. **Under-return is the common case for anyone with a strong frame, and its characteristic error is to be read as a verdict on worth rather than as a reading on calibration.** §4.8 treats this at length because it is where the instrument does the most harm when misused.

Null return. Treated in §3.5 and tested in §4.5.

23.4.2 4.2 The magnitude scale — what a given `Delta_R` is

`rho` is in nats. Because the four registers converge [Curzi, 2026n,?], the same number is immediately readable as a ratio, as bits, and as an interval in octaves and cents. The conversion constants are exact:

```
1 nat = 1.442695 bits = 2.885390 octaves = 3462.4681 cents
```

`Delta_R` (nats)	ratio `M'/M`	percent change	bits	octaves	cents
0.0050	1.0050	0.50%	0.0072	0.0144	17.3
0.0100	1.0101	1.01%	0.0144	0.0289	34.6
0.0200	1.0202	2.02%	0.0289	0.0577	69.2
0.0500	1.0513	5.13%	0.0721	0.1443	173.1
0.1000	1.1052	10.52%	0.1443	0.2885	346.2
0.2000	1.2214	22.14%	0.2885	0.5771	692.5
0.2500	1.2840	28.40%	0.3607	0.7213	865.6
0.3466	1.4143	41.43%	0.5000	1.0001	1200.1
0.5000	1.6487	64.87%	0.7213	1.4427	1731.2
0.6931	2.0000	99.99%	0.9999	2.0000	2399.8
1.0000	2.7183	171.83%	1.4427	2.8854	3462.5
1.3863	4.0000	300.00%	2.0000	4.0000	4800.0
2.0000	7.3891	638.91%	2.8854	5.7708	6924.9

Delta_R (nats)	ratio M'/M	percent change	bits	octaves	cents
2.3026	10.0001	900.01%	3.3219	6.6439	7972.7

Two landmarks are worth memorising because they anchor every subsequent judgement. **Delta_R = 0.3466 nats is exactly one octave and exactly half a bit** — the world answered about 1.41 times what you sent, or about 0.71 times. **Delta_R = 0.6931 nats is exactly one bit and exactly two octaves** — a factor of two, either way. A return difference of one bit means the exchange resolved one binary question you did not previously have answered, and if that sounds like a small yield for an entire cycle of action, it is: **exchanges are expensive and their information yield is measured in single bits, which is the correct order of magnitude and the reason patience is an engineering requirement rather than a virtue.**

A caution that must travel with this table, because the table invites the error. **The nats of a return difference are not automatically nats of information.** They become information only relative to the echo's own noise, and §4.4 gives the ratio that decides it. A Delta_R of 0.20 against a measurement noise of 0.30 is not one-fifth of a bit of anything; it is a number smaller than the wobble of the instrument that produced it.

23.4.3 4.3 The deadband — the smallest echo you can hear

Weber's law (1834): the just-noticeable difference is a constant ratio, dI/I = c. Applied to the echo, this gives a hard floor on Delta_R below which the subject's own instrument cannot resolve the return.

```
Delta_R_min = ln(1 + c)
```

The derivation is one line: the smallest ratio the subject can distinguish from unity is 1 + c, and Delta_R is the log of the ratio.

Weber fraction c	Delta_R_min (nats)	bits	ratio	cents
0.01	0.0100	0.0144	1.01	34.5
0.02	0.0198	0.0286	1.02	68.6
0.03	0.0296	0.0426	1.03	102.3
0.05	0.0488	0.0704	1.05	168.9
0.08	0.0770	0.1110	1.08	266.5
0.10	0.0953	0.1375	1.10	330.0
0.15	0.1398	0.2016	1.15	483.9
0.20	0.1823	0.2630	1.20	631.3
0.30	0.2624	0.3785	1.30	908.4
0.50	0.4055	0.5850	1.50	1403.9

The engineering consequence: **an subject with a Weber fraction of 0.20 on moment magnitude cannot distinguish a genuine 15% return from silence, and will experience nine weeks of real but small feedback as nine weeks of nothing.** This is not a failure of attention. It is the resolution of the instrument, and the fixes are instrumental rather than motivational: increase the modulus (rate on a finer scale), aggregate (§4.4), or amplify the sent moment so the return clears the floor.

Note also the honest reading of the deadband in the other direction. **A null return inside the deadband is not evidence of a null return.** It is evidence that the echo, whatever it was, was smaller than you can hear. The distinction matters because the two conclusions license

opposite actions: *no signal* says change what you are sending; *below resolution* says change how you are listening.

An empirical caveat with its badge: the corpus has not measured a Weber fraction for moment magnitude and does not know one. **What would need measuring is stated in §4.11.** The table above is the arithmetic of the relation, not a claim about any particular person's threshold.

23.4.4 4.4 Noise, aggregation, and how many echoes a finding costs

A `rho` read from a person's own record of their own encounters is exposed twice over: to the instability of self-report under question wording and framing [Schwarz, 1999, Kahneman et al., 2006], and to the more basic finding that people's accounts of their own processes are reconstructions rather than readings [Nisbett and Wilson, 1977]. A single `rho` is a noisy quantity. Estimate the noise as `sigma_rho`, the standard deviation of repeated log-ratios under conditions you regard as unchanged. Then the echo's signal-to-noise is

```
SNR_R = Delta_R / sigma_rho
```

Delta_R sigma_rho	0.05	0.10	0.20	0.30	0.50
0.05	1.00	0.50	0.25	0.17	0.10
0.10	2.00	1.00	0.50	0.33	0.20
0.20	4.00	2.00	1.00	0.67	0.40
0.35	7.00	3.50	1.75	1.17	0.70
0.50	10.00	5.00	2.50	1.67	1.00
0.69	13.80	6.90	3.45	2.30	1.38
1.00	20.00	10.00	5.00	3.33	2.00

Aggregation reduces noise in the usual way — the standard error of a mean of `n` independent log-ratios is `sigma_rho / sqrt(n)` — so the number of echoes needed to resolve a true return difference `Delta` at the conventional two-sided 95% level is

```
n >= ( 1.96 * sigma_rho / Delta )^2
```

sigma_rho Delta	0.05	0.10	0.20	0.35	0.50	0.69
0.10	16	4	1	1	1	1
0.20	62	16	4	2	1	1
0.30	139	35	9	3	2	1
0.50	385	97	25	8	4	3

Read the top-left corner and the bottom-left corner together: a small return difference under large noise is not a finding you can afford, and the honest response is to stop trying to detect it rather than to detect it badly. An subject with `sigma_rho = 0.30` who wants to establish a 5% return effect needs 139 clean, independent echoes. If the cycle time is a week, that is two and a half years. The engineering answer is never "look harder"; it is to increase `Delta` (send a moment big enough to produce a hearable answer), reduce `sigma` (fix the estimation procedure), or abandon the question.

The word *independent* in that paragraph is doing a great deal of work, and the next section is about what happens when it is false.

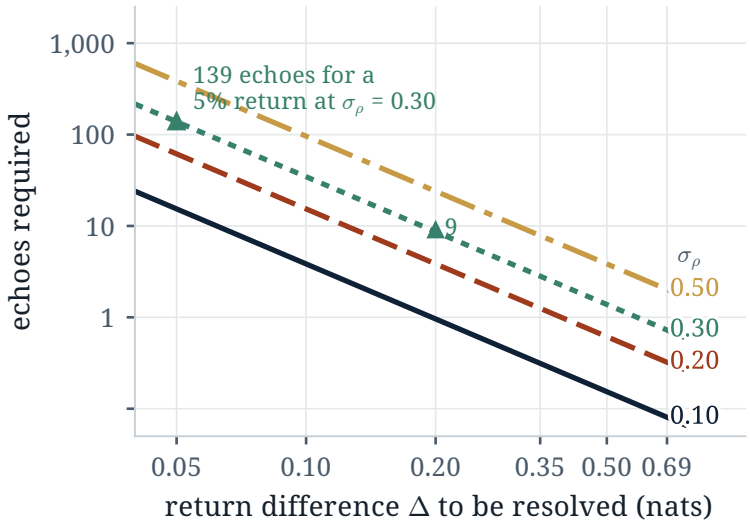


Figure 23.1: **A personal echo log can detect large distortions and nothing subtle:** echoes required to resolve a return difference at the two-sided 95 per cent level, `n >= (1.96 sigma_rho/Delta)^2`. At `sigma_rho = 0.30` a 5 per cent return needs **139** clean independent echoes — more than two years at a weekly cycle — while a 20 per cent return needs **9**. Derived; the curves reproduce all twenty-four cells of this section's table exactly, which the script prints; generated by `figs/src/CCP-023_fig04.py`.

23.4.5 4.5 The mirror that flatters — correlated echoes and the design effect

The arithmetic below is the personal-scope form of a result stated with its mathematics attached in 1952: **the benefit of holding several things depends on their covariance, not on their count**, so `n` perfectly correlated sources are worth exactly one [Markowitz, 1952]. The finance literature adds the part that makes it dangerous rather than merely disappointing: correlations rise under stress [Longin and Solnik, 2001], so the effective number of independent sources falls exactly when it is most needed. Informational cascade [Bikhchandani et al., 1992, Banerjee, 1992] is the mechanism by which a set of nominally independent respondents becomes one, and Lorenz et al. [2011] measure the collapse directly. **No entry in this corpus's external bibliography states the survey-sampling design effect `n_eff = n/(1 + (n-1) r_icc)` itself**; it is standard, it is used below, and it is left uncited rather than attached to a source that does not state it.

The first echo pathology of the foundations series [Curzi, 2026], in exact form.

Suppose a fraction `w` of the answer that comes back is not the world's response but a reflection of the moment you sent — the audience that has learned your frame, the friend who reports what you want, the metric you designed. Model it linearly: the observed return is a blend of the true return and your own sent moment. Then

```
rho_obs = (1 - w) * rho_true
```

`w` (reflected fraction)	attenuation	`rho_true` required to observe `rho_obs = 0.10`
0.00	1.00	0.1000
0.10	0.90	0.1111
0.20	0.80	0.1250
0.30	0.70	0.1429
0.50	0.50	0.2000

`w` (reflected fraction)	attenuation	`rho_true` required to observe `rho_obs = 0.10`
0.70	0.30	0.3333
0.90	0.10	1.0000
0.95	0.05	2.0000
1.00	0.00	infinite — no echo exists

At `w = 1` the channel is a mirror and no finite world could produce a detectable return: the apparatus reports agreement no matter what is true. That is the flattering mirror, and it is not a metaphor in this form — it is a divide-by-zero in the recovery of `rho_true`.

The multi-source version is more useful because most people have several echo sources and believe that several is the same as independent. It is not. Where `n` echo sources share an intraclass correlation `r_icc`, the effective number of independent echoes is the standard design effect:

```
n_eff = n / ( 1 + (n − 1) ∗ r_icc )
```

For `n = 10` sources:

`r_icc`	`n_eff`	standard-error multiplier vs. 10 independent	the affine effective count [Curzi, 2026], `N_eff = 1 + (n−1)(1−r)` for comparison
0.00	10.000	1.000	10.000
0.05	6.897	1.204	9.550
0.10	5.263	1.378	9.100
0.20	3.571	1.673	8.200
0.30	2.703	1.924	7.300
0.50	1.818	2.345	5.500
0.70	1.370	2.702	3.700
0.90	1.099	3.017	1.900
1.00	1.000	3.162	1.000

Ten sources at an intraclass correlation of 0.30 are worth 2.7 independent sources, and the person consulting them experiences ten confirmations. That gap — between the felt count and the audited count — is the same gap the framework's affine effective count [Curzi, 2026] measures for group members, and the last column is included to make an honest point rather than a tidy one: **the two audits agree exactly at both endpoints and disagree in between.** `N_eff = 1 + (n−1)u` and `n_eff = n/(1+(n−1)r_icc)` both collapse to 1 under total redundancy and both return `n` under total independence, but they are different functions of the correlation and they are answering different questions — The surplus functional [Curzi, 2026] prices the surplus generated by a set of states, the design effect prices the precision of a mean. Do not substitute one for the other. **Use the design effect when you are averaging echoes and the surplus functional [Curzi, 2026] when you are asking what the echoes are worth as a set.**

The zero-return test. This is the instrument §3.5 promised — the way to tell contact from closure when `Delta_R` is persistently near zero. Under genuine contact with a live world, observed `rho` values scatter: the return is sometimes over, sometimes under, by amounts comparable to the measurement noise. Under closure, they do not scatter, because the channel is returning the sent moment. So compute the probability that `n` genuinely independent echoes all land inside a deadband of half-width `b` when the true return difference is zero and the noise is `sigma`:

sigma	band b	P(one echo inside band)	P(10 in a row)	n for p^n < 0.05
0.10	0.05	0.3829	0.000068	4
0.10	0.10	0.6827	0.021990	8
0.20	0.05	0.1974	0.0000000	2
0.20	0.10	0.3829	0.000068	4
0.30	0.05	0.1324	0.0000000	2
0.30	0.10	0.2611	0.0000010	3

Read the right-hand column as the flattery alarm: with a noise level of 0.20 and a deadband of 0.10, four consecutive echoes landing inside the band is already a 5%-level event under the hypothesis that the channel is live and merely quiet. Four weeks of "that was great" is not warmth. It is a measurement, and the measurement says the channel is closed.

Two honest limits on this test. First, it assumes the noise `sigma` was estimated on the same channel under conditions where the channel was known to be live, which requires a period of record before the suspicion arose. Second, it cannot distinguish closure from genuine perfect calibration, and perfect calibration does occur — a practitioner who has converged genuinely does get small, scattered returns. **The discriminator is scatter, not size: converged contact produces small returns that vary; closure produces small returns that do not.**

23.4.6 4.6 The Echo that Devours — correction gain and the stability boundary

The second echo pathology of the foundations series [Curzi, 2026], and the place where this paper produces its sharpest number.

An echo is only useful if it changes something. The natural thing for it to change is the frame: the subject's perspective `D` in the Moment equation `M = (D/P)E` ([Curzi, 2026m]). Correct multiplicatively, because the register is multiplicative:

```
D_{t+1} = D_t * ( M'_t / M_t )^lambda          lambda >= 0, the CORRECTION GAIN
```

Now trace what that does to the next return difference. Write `A` for the answer the world would give — the quantity the subject is trying to match — and let the sent moment be `M_t = (D_t / P) E`. Then `rho_t = ln(A / M_t)`, and

```
M_{t+1} = (D_{t+1}/P) E = M_t * (A/M_t)^lambda

rho_{t+1} = ln(A / M_{t+1})
          = ln(A/M_t) - lambda * ln(A/M_t)
          = (1 - lambda) * rho_t
```

The return difference obeys a linear contraction in the log register with factor `(1 - lambda)`, and that is the same mathematical object as the storage law's `(1 - delta)` — which is why this section is the setup for the next. The echo and the accumulator are the same circuit read at two places. Sections 7–10 works the consequence out in full; here is the stability table.

lambda	|1 - lambda|	behaviour	steps to reduce |rho| 10x	steps to reduce 100x
0.10	0.9000	monotone, slow	22	44
0.20	0.8000	monotone	11	21
0.25	0.7500	monotone	9	17

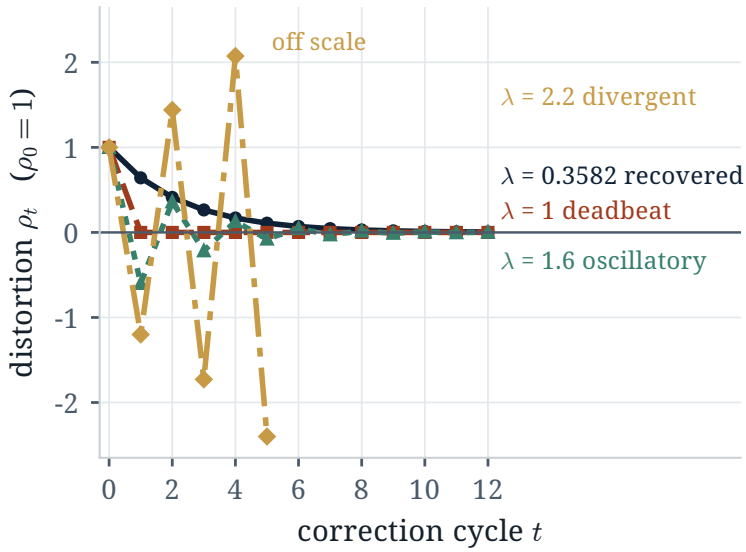


Figure 23.2: **The stability boundary is a gain setting, not an atmosphere:** the correction map `rho_{t+1} = (1−lambda) rho_t` from `rho_0 = 1`. At the recovered gain 0.3582 the distortion decays monotonically with a half-life of 1.56 cycles; at `lambda = 1` it is deadbeat, reaching zero in one step and transmitting the whole measurement noise with it; at `lambda = 1.6` it converges by overshooting alternately; at `lambda = 2.2` it grows without bound and leaves the axes at cycle 5. Modelled by the map tabulated in this section; generated by `figs/src/CCP-023_fig02.py`.

`lambda`	`|1 − lambda|`	behaviour	steps to reduce `|rho|` 10x	steps to reduce 100x
0.33	0.6700	monotone	6	12
0.50	0.5000	monotone	4	7
0.75	0.2500	monotone, fast	2	4
1.00	0.0000	**deadbeat** — one step	1	1
1.25	0.2500	oscillatory, converging	2	4
1.50	0.5000	oscillatory	4	7
1.75	0.7500	oscillatory, slow	9	17
2.00	1.0000	**marginal** — undamped oscillation	never	never
2.50	1.5000	**divergent** — the devouring echo	never	never

Five statements follow, and they are the operational content of the ECHO node's second warning.

One. There is a stability boundary and it is at `lambda = 2`. Below it the return difference shrinks; above it the return difference grows without bound. At exactly 2 it oscillates forever at constant amplitude. This is the Echo that Devours, located precisely: it is not an atmosphere, it is a gain setting.

Two. `lambda = 1` is the fastest possible correction and it is a trap. Deadbeat correction takes the whole observed difference as truth and applies it in one step. If the observation were noiseless this would be optimal. It is not noiseless, so `lambda = 1` also takes the whole ob-

servation *noise* as truth and writes it into the frame. The frame then thrashes at the amplitude of the measurement error.

Three. Everything above 1 is oscillatory even when it converges. An subject at `lambda = 1.5` overshoots every time and converges by alternating: too much, too little, too much. From the inside this is indistinguishable from instability of character, and it is a gain setting.

Four. Slow gains are not timid; they are filters. `lambda = 0.25` needs nine cycles to close a factor of ten, and in exchange it attenuates the noise in every single observation by 75%. The choice of `lambda` is a bandwidth choice: **fast correction tracks the world and the noise; slow correction tracks the world only, late.** There is no setting that does both, and the belief that there is one is the most common source of thrash.

Five. The optimum depends on the noise, not on the character of the subject. As a design rule, badged as a heuristic rather than a derived optimum: set `lambda` near `SNR^2/(1+SNR^2)` where `SNR` is the per-echo signal-to-noise from §4.4 — near 1 when echoes are clean, near 0 when they are not. The corpus has not derived that expression from a stated loss function and does not present it as optimal; it presents it as a starting value with the right limits.

23.4.7 4.7 Delay, aliasing, and the echo that arrives after the world has moved

Returns take time. If the world's answer to period `t` arrives at period `t + k`, and the subject applies it as if it were the answer to period `t + k`, the correction is applied to the wrong state. In control terms this is transport delay, and its effect is to convert a stable gain into an unstable one: a correction that would be safely damped at zero delay can oscillate or diverge at a delay of several periods, because the subject keeps correcting for a discrepancy that has already been corrected.

The practical rule, and the reason it belongs in an account of persons rather than one of servomechanisms: **never run a correction gain faster than the echo returns.** If the answer takes three weeks, do not correct weekly. Correct every three weeks, or reduce `lambda` by roughly the delay factor, or — best — timestamp the echo to the moment it answers rather than to the moment it arrives. The third option costs nothing but a column in the log and removes the problem at its source, which is why the specification below requires it.

The dual failure is sampling too slowly. If the world's answer changes faster than you sample it, the sampled series aliases: a rapid alternation reads as a slow trend or as nothing at all. The symptom is a return log whose variance is implausibly low for a domain everyone agrees is volatile. **A suspiciously smooth echo log is a sampling artifact until proven otherwise.**

23.4.8 4.8 What the echo is not

Four things, because the instrument is dangerous in proportion to how naturally it is misread.

It is not a verdict on worth. `rho < 0` says the world answered smaller than you constructed. It does not say the thing was bad, and FVP-5 above is the formal statement of why: a large negative `rho` is equally consistent with an over-framed subject and an unresponsive world, and `rho` alone cannot separate them. Anyone who reads a run of negative returns as a statement about their worth has performed an identification that the mathematics explicitly declines to perform.

It is not defined on ordinal data. FVP-2 is a gate, not a caution. If `M` and `M'` are 1–5 stars, `M'/M` is meaningless — the ratio of two rank positions is not a ratio of two magnitudes. The demotion in that case is mandatory and it is this: report the sign only, as an ordering (*larger / same / smaller*), call it a heuristic reading, and **do not compute a mean, a standard deviation, a `lambda`, or a convergence rate from it.** A great deal of self-tracking practice violates this gate silently, and the resulting numbers have the shape of measurement without the content.

It is not the world's total answer. As deployed, `rho` is one scalar summary of an exchange that had structure. This mirrors the honesty clause on the surplus functional [Curzi, 2026,?]

— we log the floor, not the ceiling. **A `rho` of zero on the dimension you measured is compatible with an enormous return on a dimension you did not.**

It is not transferable between subjects without recalibration. `M` and `M'` are magnitude estimates made against a modulus. Two people using different moduli produce `rho` values that are comparable — this is the scale-invariance property of §3.3 — but only if each person's own estimates are internally consistent on a ratio scale. The comparison of `lambda` across people is the open FVP-7 gate, not a licensed operation.

23.4.9 4.9 Chaining and the composition of returns

Because log-ratios add, a series of exchanges has a total return difference that is the sum of the signed parts and a total *activity* that is the sum of the magnitudes. These are different numbers and their ratio is a diagnostic.

```
net return   R_net = SUM rho_t                       (signs cancel)
gross return R_gross = SUM |rho_t| = SUM Delta_R,t   (signs do not)

EFFICIENCY   E_R = |R_net| / R_gross      in [0, 1]
```

`E_R` near 1 means the returns pointed consistently the same way: the world has been saying one thing repeatedly and the subject has not yet moved. `E_R` near 0 means the returns cancelled: a great deal of exchange, no net direction. **A high gross with a near-zero net is the signature of thrash — the subject is correcting hard in alternating directions, which §4.6 identifies as a gain above 1.** A high gross with a high net is the signature of a live, uncorrected error, which is the good case, because it means there is something to learn and it has not been learned yet.

23.4.10 4.10 The join: from echo to stored order

The formulary is a single circuit and this section names the wires. The chain runs:

```
E (environment)
   |
   v
M = (D/P) E                     the moment equation
   |
   +----> enacted -----> world -----> M'          the exchange
   |                                  |
   |                                  v
   |                           rho = ln(M'/M)      SECTIONS 3-6
   |                                  |
   |         D_{t+1} = D_t * e^{lambda*rho}        the correction
   |                                  |
   +<---------------------------------+
   |
   v
Q_{t+1} = (1-delta) Q_t + eta*M_t - C_t           the self-storage law
```

Three consequences of drawing it this way.

The echo does not enter the storage law directly. It enters through `D`, which changes `M`, which is the storage law's input. This is why feedback improves character slowly: it acts on the input of an integrator, not on the integrator's state. An echo received today changes `Q`'s trajectory, not `Q`.

The correction loop and the storage loop have different time constants and both are geometric. The echo loop contracts at `(1 - lambda)`, the storage loop at `(1 - delta)`. When `lambda >> delta` the frame tracks and the character lags; when `lambda << delta` the character forgets faster than the frame corrects, which is the arithmetic form of learning a lesson repeatedly.

The identification problem of FVP-5 has a design solution. To separate "my frame was wrong" from "the world was quiet", hold `D` fixed for a block of periods and vary nothing, then change `D` deliberately by a known factor and see whether `rho` moves by the predicted amount. If `rho` responds to a change in `D` with the elasticity the Moment equation predicts, the return difference is reading your frame. If it does not, it is reading the world. **That is a manipulation, it is cheap, and it is the only clean way to run the decomposition.**

23.4.11 4.11 What would need measuring

The falsification gate every entry below defers to is preregistration with a frozen model and a held-out block, which is the open-science literature's own prescription [Nosek et al., 2018, Simmons et al., 2011]; the reason it is needed is the flexibility in analysis that Gelman and Loken [2013] names and the publication filter Ioannidis [2005] describes, whose measured consequence is that of one hundred direct replications in psychology about 36 recovered a significant result against 97 originals [Open Science Collaboration, 2015]. **Kill conditions stated in advance, with numbers, are what separates a falsification test from a ritual** [Platt, 1964], and the reason this framework states *exponent targets* rather than directions is that a directional prediction is nearly costless in psychology [Meehl, 1967, 1990].

The badging convention [Curzi, 2026] requires that where a figure needs research not done, the corpus says what would need measuring. Four things, stated as an agenda rather than as results:

1. **The Weber fraction for moment magnitude.** Method of constant stimuli or a staircase procedure on paired moment magnitudes within a person, yielding `c` and hence the deadband `ln(1+c)`. Without it, §4.3 is arithmetic without a calibration.
2. **The within-person noise `sigma_rho`.** Test–retest of `rho` on exchanges the subject judges equivalent. Everything in §4.4 depends on it and it is the cheapest of the four to obtain.
3. **The distribution of `lambda` across people, and its stability within a person across domains.** This is FVP-7. The design is stated: same operationalization, several subjects, several investigators, held-out prediction, shuffled-input control.
4. **The intraclass correlation of a typical person's echo sources.** This determines whether the flattering-mirror alarm in §4.5 fires at four echoes or at forty, and it is measurable by having several sources rate the same sent moment independently.

23.5 5 Requirements for an echo instrument

Numbered, checkable requirements for an echo instrument. A practice that fails any of ES-1 through ES-6 is not measuring return difference and should not report a number.

ES-1 — Ratio scale, or demotion. `M` and `M'` shall be obtained by magnitude estimation against a stated modulus, with a true zero. If only ordinal data are available, the practice shall report sign only and shall label the output a heuristic ordering. *(FVP-2. Non-negotiable.)*

ES-2 — The modulus is declared and frozen. One reference exchange is assigned a magnitude (commonly 50 or 100) at the start of a block, and all estimates in the block are made against it. Changing the modulus mid-block invalidates the block for aggregation.

ES-3 — `M` is recorded before the answer arrives. The sent moment shall be estimated and written down prior to any return. An `M` reconstructed after the echo is contaminated by the

echo and yields a `rho` biased toward zero. *(This is the single most violated requirement in self-tracking and it is the one that most reliably manufactures the appearance of calibration.)*

ES-4 — `M'` is timestamped to the exchange it answers, not to its arrival. A delay column is mandatory. *(§4.7.)*

ES-5 — Source and independence are logged. Each echo carries the identity of its source and, where more than one source is used, sources are asked independently and before seeing one another's responses. *(§4.5; this is the personal-scope form of the silent-write-first rule.)*

ES-6 — Noise is estimated before findings are claimed. `sigma_rho` is estimated from at least eight exchanges the subject judges equivalent, before any `Delta_R` is interpreted as a finding.

ES-7 — The deadband is stated. The practice declares its `Delta_R_min` and refuses to interpret returns below it as anything but unresolved.

ES-8 — The correction gain is chosen, written down, and held for the block. `lambda` shall be in `(0, 1]` unless a specific reason for overshoot is recorded. `lambda >= 2` is prohibited by this specification. *(§4.6.)*

ES-9 — Correction cadence does not exceed echo cadence. If returns take `k` periods, the gain is applied at most every `k` periods. *(§4.7.)*

ES-10 — The flattery alarm is run. After every four consecutive returns inside the deadband, the practice runs the zero-return test of §4.5 and, if it fires, adds an independent source before continuing.

ES-11 — Net and gross are both reported. `R_net`, `R_gross`, and `E_R` are computed per block. *(§4.9.)*

ES-12 — No worth claims. The instrument's output is a calibration reading. Reports shall not convert `rho` into a statement about the value of the person, the work, or the relationship. *(Section 4.8, and the framework's refusals [Curzi, 2026].)*

ES-13 — Decline path. Where an echo source is a person, that person may decline to supply an estimate, at any time, without explanation and without consequence, and the practice records the decline as missing data rather than as a zero. *(The declination refusal is binding [Curzi, 2026]: no protocol in the corpus removes a person's ability to decline. A refused echo is not a null return; recording it as one is both a measurement error and a coercion.)*

ES-14 — Freeze before predicting. Any claim that the instrument works shall follow the FVP-6 procedure: preregister, fit, freeze, predict held out, beat the four baselines including the shuffled-input control.

23.6 6 Failure modes of echo reading

Fault	Diagnostic signature	Fix
The Flattering Mirror	Four or more consecutive returns inside the deadband, with variance far below `sigma_rho`; all sources agree; `E_R` undefined because gross is near zero	Add a genuinely independent source (ES-5, ES-10). Compute `n_eff`. Treat the existing sources as one source until proven otherwise
The Devouring Echo	`|rho|` growing period on period; corrections alternating in sign with increasing amplitude; frame reversals faster than the world could plausibly change	Measure `lambda` from the data (regress `ln|rho|` on `t`). If the implied `|1−lambda| >= 1`, reduce `lambda` below 1 and hold. *(§4.6)*

Fault	Diagnostic signature	Fix
Ordinal laundering	Means, SDs and trend lines computed from star ratings, tick-boxes, or Likert items	Demote to sign-only reporting; rebuild the instrument on magnitude estimation. **The numbers already computed are void, not merely imprecise** *(FVP-2, ES-1)*
Post-hoc `M`	`rho` unusually close to zero across the whole log; the subject "knew it would land like that"	Enforce ES-3. Compare a block of pre-recorded `M` against a block of reconstructed `M`; the difference in mean `|rho|` is the size of the contamination
Deadband blindness	Long runs of "no feedback" in a domain where others report feedback	Raise the modulus resolution, aggregate per §4.4, or increase the size of the sent moment. Do not increase effort at listening
Delay-induced instability	Oscillation that appears only when cadence is increased; corrections that were stable at monthly review become unstable at weekly	Enforce ES-4 and ES-9. Timestamp to the exchange, not the arrival
Aliasing	Return log with implausibly low variance in a volatile domain; long smooth trends with no reversals	Sample faster than the world changes, or state explicitly that the log is a low-pass view and stop reading its trend
Selection of echoes	The log contains more returns from the exchanges the subject felt good about; missing entries cluster	Pre-commit to logging every exchange in a defined class, including the ones you would rather not; count and report the missing
Worth conversion	Negative `rho` narrated as a statement about the self; the log stops being kept after a bad block	Enforce ES-12. Re-read FVP-5: the decomposition is not identified, and the identification being performed is not licensed by the instrument
Zero read as success	`Delta_R ~ 0` celebrated; the practice concludes it has arrived	Run the scatter test (§4.5). Converged contact produces small returns that **vary**; closure produces small returns that do not
Gain thrash	High `R_gross`, near-zero `R_net`, `E_R < 0.2`, frequent reversals	`lambda` is above 1. Reduce it. *(§4.9, §4.6)*
Uncorrected live error	High `R_gross`, high `R_net`, `E_R > 0.8`, all returns pointing the same way for many periods	`lambda` is near 0 — the subject is measuring and not correcting. Raise `lambda` toward 0.3–0.5

23.7 7 The contraction and arrival

The ninth node of the foundations series [Curzi, 2026] taught the spiral rather than the circle: repetition that absorbs its own error and returns to the same post a little higher. Its tenth node taught the room that was always furnished — arrival as recognition rather than as victory, a destination that turns out to have been visited many times in imagination before it was reached in fact.

Those are not two teachings. They are the two terms of one equation, and this paper is where the correspondence is proved rather than asserted.

The corpus's own correspondence table records that the FIXEDPOINT node of the foundations series [Curzi, 2026] *is* the storage law's steady-state theorem [Curzi, 2026a], and the RECURSE node *is* the transient `(1−delta)^t Q_0`. Neither place proves it. The claim as it stands is a claim of identity between a ceremonial teaching and a theorem, and a claim of identity has to be discharged, not repeated. **This paper discharges it: it constructs the metric space, verifies the contraction, invokes Banach, solves for the fixed point, derives the exact error law, and shows that every property the foundations series attributed to arrival is a property of that solution and not an interpretation laid over it.**

23.7.1 7.1 The starting point

the storage law [Curzi, 2026a], in its internal-scope form [Curzi, 2026o] (the self-storage law):

```
Q_{t+1} = (1 − delta) Q_t + eta ∗ M_t − C_t

Q_t      stored order: character, skill, memory, internalized structure
M_t      the moment as lived, normalized into [0, ln N]
delta    retention loss in [0,1] −− forgetting, atrophy, drift of self
eta      conversion efficiency >= 0 −− how experience becomes character
C_t      cognitive load, effort, maintenance of self, >= 0
```

Take the steady-regime case that Theorem 3.5 addresses: constant geometry `M_t = M∗` for all `t`, constant cost `C_t = c∗`, and `delta > 0`. Define the update map

```
T : R −> R,      T(Q) = (1 − delta) Q + (eta ∗ M∗ − c∗)
```

Write `k = 1 − delta` for the contraction factor and `b = eta∗M∗ − c∗` for the drive. Then `T(Q) = kQ + b`, and the whole of Nodes 09 and 10 is the behaviour of iterating this map.

23.7.2 7.2 Step one — the space is complete

Banach's theorem requires a **complete** metric space. Take `X = R` with the usual metric `d(x,y) = |x − y|`. The reals with the absolute-value metric are complete: every Cauchy sequence converges, and its limit is real. This is standard and it is stated here because the proof needs it and because the honest form of the argument names every hypothesis it uses.

One caveat, discharged now. Admissibility (`C_t <= (1−delta)Q_t + eta S_t`) keeps the modelled quantity in `[0, inf)`, which is also complete, so the argument runs on either space. **We run it on `R` deliberately: Theorem 3.5's negative fixed point lives outside `[0, inf)`, and a proof on the restricted space would hide exactly the case the theorem refuses to hide.**

23.7.3 7.3 Step two — the map is a contraction, with constant `1 − delta`

For any `Q_1, Q_2 in R`:

```
| T(Q_1) − T(Q_2) |  =  | (kQ_1 + b) − (kQ_2 + b) |
                     =  | k (Q_1 − Q_2) |
                     =  |1 − delta| ∗ | Q_1 − Q_2 |
```

For `delta` in `(0, 1]`, `|1 − delta| = 1 − delta < 1`. The map is Lipschitz with constant `1 − delta`, and that constant is strictly less than 1. **`T` is a contraction, and it is a contraction for exactly one reason: `delta > 0`.**

Note that the inequality in the definition of a contraction is here an **equality**. `T` is affine, so it does not merely shrink distances by at most `1−delta`; it shrinks every distance by exactly `1−delta`. This is stronger than Banach requires, and §7.6 collects what the extra strength buys.

At `delta = 0` the constant is 1, `T` is an isometry, and Banach does not apply — correctly, since `T(Q) = Q + b` has no fixed point at all when `b != 0`. That is the contrast made exact: `Q_t = Q_0 + t∗b`, unbounded linear growth, no arrival. **The leak is what makes arrival possible, and the proof shows precisely where: it is the single step above where `1 − delta < 1`.**

23.7.4 7.4 Step three — Banach, and the fixed point solved

Banach's fixed point theorem: a contraction on a nonempty complete metric space has exactly one fixed point, and the iterates from any starting point converge to it.

Applied: `T` has a unique `Q_inf` with `T(Q_inf) = Q_inf`, and `Q_t = T^t(Q_0) −> Q_inf` for every `Q_0`. Solve:

```
Q_inf = (1 − delta) Q_inf + eta∗M∗ − c∗
delta ∗ Q_inf = eta∗M∗ − c∗

Q_inf = ( eta ∗ M∗ − c∗ ) / delta              <−− the steady−state theorem
```

which, writing `M∗ = F(u∗) = ln(1 + (N−1)u∗)` for the surplus form, is Theorem 3.5 verbatim: `Q_inf = (eta F(u∗) − c∗) / delta`.

Three things were claimed by the FIXEDPOINT node [Curzi, 2026] and all three are now consequences of the line above rather than readings of it.

Existence. Banach guarantees a fixed point. The room is furnished. There is a state the practice converges to, and it exists whether or not anyone has been there.

Uniqueness. Banach guarantees exactly one. There is no second room. Two different practices with the same `(eta, M∗, c∗, delta)` arrive at the same place, and one practice cannot arrive at two.

Independence of the start. The expression for `Q_inf` **contains no `Q_0`.** Where you arrive does not depend on where you began. The foundations series said arrival feels like recognition rather than victory; §8.3 shows that this is what the absence of `Q_0` from that formula feels like from inside.

23.7.5 7.5 Step four — the exact error law

Banach gives convergence. The affine structure gives the exact rate. Subtract the fixed-point equation from the update:

```
Q_{t+1} − Q_inf = (1−delta) Q_t + b  −  [ (1−delta) Q_inf + b ]
                = (1−delta) ( Q_t − Q_inf )
```

so by induction

```
Q_t − Q_inf = (1 − delta)^t ( Q_0 − Q_inf )

| Q_t − Q_inf | = (1 − delta)^t | Q_0 − Q_inf |        THE ERROR LAW
```

and rearranged, the form that carries the most meaning:

```
Q_t = Q_inf + (1 − delta)^t ( Q_0 − Q_inf )
      -----   --------------------------
      where     what remains of where you
      you are   started, decaying at rate delta
      going
```

The entire trajectory is the destination plus a fading memory of the origin. That single line is the RECURSE and FIXEDPOINT nodes together [Curzi, 2026,?]: FIXEDPOINT is the first term, RECURSE is the second, and the reason the corpus has treated them as one thing is that they are the two terms of this decomposition.

Consistency check against the exact solution of the storage law [Curzi, 2026a]. With constant drive `b`,

```
Q_t = (1−delta)^t Q_0 + SUM_{j=0}^{t−1} (1−delta)^j ∗ b
    = (1−delta)^t Q_0 + b ∗ (1 − (1−delta)^t)/delta
    = (1−delta)^t Q_0 + Q_inf ∗ (1 − (1−delta)^t)
    = Q_inf + (1−delta)^t (Q_0 − Q_inf)
```

The geometric sum closes exactly. **Theorem 3.1 and the error law are the same statement, and the Banach argument was not needed to obtain it in this one-dimensional affine case — it was needed to know that the result is not an accident of linearity.** That distinction is the subject of §7.6.

23.7.6 7.6 What Banach buys that the algebra does not

An honest accounting, because it would be easy to over-sell the theorem here.

For this specific affine map on `R`, the closed form above is elementary and Banach is not required. What Banach supplies is **robustness of the conclusion to changes in the model that the closed form would not survive**:

- If `eta`, `C`, or the moment-generating mechanism are **nonlinear in `Q`** — as they are in any realistic account where effort rises with what must be maintained — the closed form vanishes but the fixed point survives, provided the map remains a contraction. The condition to check becomes `|T'(Q)| <= k < 1` on the region of interest.
- If `Q` is **multi-dimensional** — several stored capacities that interact — the same theorem applies on `R^n` with any complete norm, and the contraction condition becomes a bound on the matrix norm of the update.
- Banach also supplies **error bounds computable before convergence**, which the closed form gives only after you already know `Q_inf`:

```
a priori:       d(Q_t, Q_inf) <= ( k^t / (1−k) ) ∗ d(Q_1, Q_0)
a posteriori:   d(Q_t, Q_inf) <= ( k / (1−k) )  ∗ d(Q_t, Q_{t−1})
```

The a posteriori bound is the useful one in practice: **it estimates how far you are from arrival using only the size of your last step.** Worked at `delta = 0.12`, `Q_0 = 0`, drive `b = 0.3702` (the case of §8.8):

t	Q_t	true error	a priori bound	a posteriori bound
1	0.3702	2.7151	2.7151	2.7151
3	0.9828	2.1025	2.1026	2.1026
5	1.4571	1.6282	1.6282	1.6282

t	Q_t	true error	a priori bound	a posteriori bound
8	1.9757	1.1096	1.1096	1.1096

The bounds are attained with equality, to rounding. That is the affine case being exactly as strong as §7.3 said: distances shrink by exactly `1−delta`, so the inequality in Banach's bounds becomes an identity. In the nonlinear or multi-dimensional case the same bounds hold as inequalities and remain usable. **The practical instruction survives either way: your distance from arrival is about `(1−delta)/delta` times the size of your last step.** At `delta = 0.12` that multiplier is 7.33 — a step of 0.05 means you are roughly 0.37 from the fixed point.

23.7.7 7.7 The correspondence, discharged

The RECURSE node	The term	Status
The spiral, not the circle	`Q_t = Q_inf + (1−delta)^t(Q_0 − Q_inf)`: the state changes every pass and never repeats exactly	**Proved** — §7.5
The winch post held fixed	`Q_inf` fixed across passes; only the transient moves	**Proved** — `Q_inf` has no `t` in it
The drift, checked weekly	`(1−delta)^t (Q_0 − Q_inf)`, the decaying inheritance	**Proved** — §7.5
The circle that does not climb	`eta = 0`: `b = −c∗`, `Q_inf = −c∗/delta < 0`	**Proved** — the circle is a regime whose fixed point is destruction

The FIXEDPOINT node	The term	Status
The room exists	Banach existence, given `delta > 0` and completeness	**Proved** — §7.4
There is only one room	Banach uniqueness	**Proved** — §7.4
The room was always furnished	`Q_inf` independent of `Q_0`	**Proved** — §7.4
Arrival is recognition, not victory	geometric approach; most of the trajectory spent near `Q_inf`; the crossing is a JND event	**Derived** — §8.3, given a stated Weber fraction
Convergence is no certificate of worth	`eta F(u∗) < c∗` gives `Q_inf < 0` and the loop still converges	**Proved** — Theorem 3.5's own failure clause
The counterfeit arrival	a fixed point inflated by `x/delta` at another system's expense; sum invariant	**Derived** — §8.7

The two rows marked *derived* rather than *proved* are derived from the mathematics plus one badged postulate each, and both badges are shown where the derivation occurs. Nothing in this reading imports a meaning into the theorem that the theorem does not carry.

23.7.8 7.8 Formula Validation Protocol — the error law

Under the Formula Validation Protocol [Curzi, 2026t].

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`Q` in Q-units (the corpus's unit of stored order; see [Curzi, 2026b] on conversion). `delta` dimensionless per period; `(1-delta)^t` dimensionless. Both sides of `|Q_t - Q_inf| = (1-delta)^t |Q_0 - Q_inf|` are in Q-units. `t` in periods; the period must be declared.
FVP-2 SCALE TYPE	**CONDITIONAL PASS**	The error law needs `Q` on an **interval** scale for differences and on a **ratio** scale for the percentage readings in §8.1. The corpus's `Q` is postulated as a ratio-scale scalar with a true zero at "no stored order" (Appendix E's postulate set). **Badged: that `Q` is a scalar at all, rather than a vector of incommensurable stores, is a postulate of the corpus and not a finding.** Where an subject's `Q` is measured by a Rasch-calibrated instrument ([Curzi, 2026t]) the scale requirement is met; where it is measured by a self-rating out of 10, it is not, and only the ordering of successive gaps may be reported.
FVP-3 LIMITING CASES	PASS	`t -> inf`: gap `-> 0` for `delta in (0,1]`. `delta -> 0+`: gap `-> |Q_0 - Q_inf|`, no convergence; and `Q_inf -> inf` unless `b = 0`, which is the correct account of a system that never forgets and therefore never settles. `delta -> 1`: gap `-> 0` in one step; total forgetting means the state is entirely the current drive, `Q_t = b` for all `t >= 1`. `delta > 1`: excluded by the model's domain; `|1-delta| < 1` still holds up to `delta = 2`, giving oscillatory convergence, but `delta > 1` would mean losing more than is stored and the admissibility condition rules it out. `Q_0 = Q_inf`: gap is 0 for all `t` — you were already there.

Gate	Verdict	Content
FVP-4 FIXED POINTS	PASS	`Q_inf = (eta M* − c*)/delta`, unique by Banach. Meaning: **the state a life settles into under a given regime, computable in advance from four numbers, none of which is effort.** Second fixed point of the *error* dynamics: none — the error map `e -> (1−delta)e` has the unique fixed point `e = 0`.
FVP-5 IDENTIFIABILITY	**PARTIAL**	From a single observed trajectory `{Q_t}` under constant drive, `delta` is identified from the decay of successive gaps, and the composite `b = eta M* − c*` is identified from `Q_inf * delta`. **`eta`, `M*` and `c*` are NOT separately identified from `Q` alone** — any triple with the same `eta M* − c*` produces an identical trajectory. Separating them requires independent measurement of at least two of the three: `M*` by magnitude estimation ([Curzi, 2026n]), `c*` by a cost ledger ([Curzi, 2026b]), `eta` then by residual. **State this whenever a trajectory is used to argue about efficiency rather than about outcome.**

Gate	Verdict	Content
FVP-6 FALSIFICATION	SPECIFIED	The preregistration protocol [Curzi, 2026i]. Preregister the period, the `Q` instrument, and the window. Fit `delta` on a training block; **freeze**; forecast held-out `Q_t` from the frozen `delta` and the observed drive. Four baselines: persistence (`Q_{t+1} = Q_t`), direct-input (`Q_{t+1}` from `M_t` alone), matched AR(1), and a **shuffled-input control** in which the drive series is permuted. **The law dies if a frozen geometric-decay model does not beat the shuffled-input control on held-out data**, because that is the outcome in which the ordering of inputs — the entire content of "it accumulates" — carries nothing.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Does `delta` recover across individuals, across investigators, across domains within a person? The corpus does not know. §8.10 states the design that would answer it. Until it is run, a fitted `delta` is a within-person, within-domain descriptive parameter.

23.8 8 The dynamics of arrival

23.8.1 8.1 The convergence tables

Everything operational in this paper comes off two constants. The **time constant**

```
tau = −1 / ln(1 − delta)        periods per factor of e
```

and the **iteration count to close a given fraction of the gap**

```
t(eps) = ln(eps) / ln(1 − delta)     periods to reduce the gap to a fraction eps
```

Both are exact; neither involves `Q_0`, `Q_inf`, `eta`, `M*`, or `C`. **How long arrival takes depends on the leak and on nothing else.**

delta	1−delta	tau	half-life	to 5%	to 1%	to 0.1%	to 0.01%
0.020	0.980	49.50	34.31	149	228	342	456
0.050	0.950	19.50	13.51	59	90	135	180
0.100	0.900	9.49	6.58	29	44	66	88
0.150	0.850	6.15	4.27	19	29	43	57
0.200	0.800	4.48	3.11	14	21	31	42
0.250	0.750	3.48	2.41	11	17	25	33
0.300	0.700	2.80	1.94	9	13	20	26
0.400	0.600	1.96	1.36	6	10	14	19
0.500	0.500	1.44	1.00	5	7	10	14
0.750	0.250	0.72	0.50	3	4	5	7

Read the table with the period attached, because the period is where the human meaning is.

A weekly practice with `delta = 0.05` per week is 59 weeks — fourteen months — from being within 5% of its destination. The same weekly leak restated per day is `delta = 0.0073`, needing 409 days: the same calendar. **A monthly cycle at `delta = 0.20` per month is fourteen months too.** The slightly deflating discovery in this table is that *changing cadence without changing the per-period leak changes nothing about the calendar*; cadence acts through `eta` and `C`, not through `tau`.

Two invariances that the table encodes and that are worth stating separately, because they are counter-intuitive and they are exact:

The cost of each halving is constant. At `delta = 0.10`, closing the gap from 100% to 50% takes 6.58 periods; from 50% to 25% takes 6.58 periods; from 1% to 0.5% takes 6.58 periods. The `ln(0.5)/ln(1−delta)` expression has no `t` in it.

The cost of each factor of ten is constant, and it is 21.85 periods at `delta = 0.10`.

interval	periods (at `delta = 0.10`)
close 0% -> 50% of the gap	7
close 50% -> 90%	15
close 90% -> 99%	22
close 99% -> 99.9%	22

The last ninety-nine hundredths of one percent costs exactly as much as the ninety to ninety-nine percent did, and more than the first half did. This is the arithmetic behind every experience of a practice that felt fast at the start and interminable at the end. The practice did not slow down. **The remaining distance is being measured in absolute terms by an instrument that works in ratios, and the ratios are being closed at a perfectly constant rate throughout.**

23.8.2 8.2 The shape of the approach

The fraction of the gap remaining, at `delta = 0.10`:

t	remaining	closed
0	1.00000	0.000%
1	0.90000	10.000%
2	0.81000	19.000%
5	0.59049	40.951%
10	0.34868	65.132%
20	0.12158	87.842%
29	0.04710	95.290%
44	0.00970	99.030%
66	0.00096	99.904%

A contraction spends the overwhelming majority of its life near its fixed point. At `delta = 0.10`, the system is within 5% of `Q_inf` from period 29 onward — which is to say that of the first 88 periods, 60 of them are spent inside a band of 5%. If the practice runs for two hundred periods, 171 of them are inside that band. **Whatever "being at the fixed point" feels like, it is what most of the trajectory feels like, and the transient is the unusual part.**

That fact is the entire explanation of §8.3.

23.8.3 8.3 Why arrival feels like recognition and not like victory

Four reasons, in increasing order of how much of the phenomenon they explain.

One: there is no arrival event. The gap is `(1−delta)^t` times its initial value, which is strictly positive for every finite `t`. **The system never reaches `Q_inf`. It approaches forever.** There is therefore no instant at which anything happens, and any felt instant of arrival is an event in the observer, not in the trajectory. Victory has a moment. This does not.

Two: the felt arrival is a threshold crossing in the observer's own instrument. If the subject can resolve a difference in `Q` of a fraction `c` — a Weber fraction, per [Curzi, 2026n] and §4.3 — then they notice arrival when the remaining gap falls below `c * Q_inf`, at

```
t_arrive = ln( c * Q_inf / |Q_0 − Q_inf| ) / ln(1 − delta)
```

For an subject starting from zero (`|Q_0 − Q_inf| = Q_inf`, so the ratio is just `c`):

`delta` \ `c`	0.02	0.05	0.10	0.20	0.30
0.05	77	59	45	32	24
0.10	38	29	22	16	12
0.20	18	14	11	8	6
0.30	11	9	7	5	4

Arrival is the period at which the remaining distance drops below your own just-noticeable difference, and it therefore happens at a different time for two people on identical trajectories. A coarse observer (`c = 0.30`) at `delta = 0.10` arrives at period 12; a fine one (`c = 0.02`) on the identical path arrives at period 38. Neither is wrong. They have different instruments. *(Badged: this uses Postulate R1's cousin — that `Q` is judged by the same log-register instrument as every other magnitude, hence a constant Weber fraction. That is an assumption of the corpus, consistent with the psychophysics but not measured for `Q`.)*

Three: by the time you notice, you have been close for several time constants. From §8.2, an subject with `c = 0.05` and `delta = 0.10` "arrives" at period 29 — and has been within 10% since period 22, within 20% since period 16, within 35% since period 10. **The state that announces itself at period 29 has been substantially present since period 16.** The texture of recognition is the tail of an exponential, and the reason the destination seems familiar is that it has been the dominant feature of the last two-thirds of the journey.

Four — and this is the one that makes it recognition rather than mere familiarity: the destination never depended on you. `Q_inf = (eta M* − c*)/delta` contains no `Q_0`, no `t`, and no measure of effort. It is a property of the regime. From the first period, the trajectory was going exactly there. **You do not arrive somewhere new; you arrive where the parameters were always going to put you, and the recognition is accurate — the room really was furnished before you entered it, by the four numbers that describe how you live.**

That is the honest content of the FIXEDPOINT node [Curzi, 2026] and it is not consoling. It says the same thing about a good regime and a bad one.

23.8.4 8.4 The invariant at arrival: headroom is pure geometry

A result that falls out of Theorem 3.5 and the headroom gauge together, and that the corpus has not stated before.

Headroom [Curzi, 2026b,h] is `h_t = eta ln N − delta Q_t − C_t`. Evaluate it at the fixed point. There, `delta Q_inf = eta M* − c*` by definition, so

```
h_inf = eta*lnN − delta*Q_inf − c*
      = eta*lnN − (eta*M* − c*) − c*

h_inf = eta * ( ln N − M* )                    THE ARRIVAL INVARIANT
```

Verified numerically across unrelated parameter sets:

eta	delta	N	u*	C	Q_inf	h computed from the gauge	eta(ln N − F)
0.35	0.12	9	0.55	0.22	3.0853	0.178789	0.178789
0.35	0.05	9	0.55	0.22	7.4048	0.178789	0.178789
0.20	0.30	6	0.30	0.10	0.2775	0.175094	0.175094
0.50	0.08	16	0.80	0.40	11.0309	0.103820	0.103820

At any fixed point, headroom equals the conversion efficiency times the surplus the geometry is failing to generate — and it does not depend on `delta` or on `C` at all. Rows one and two make the point: halving the leak more than doubles `Q_inf`, and leaves headroom at arrival unchanged to six decimal places.

The reading, and it is the most useful single sentence in this paper for anyone deciding what to do next:

A converged system's remaining capacity for the new is exactly the surplus its geometry is not claiming. Not its slack, not its free time, not its energy. Its unclaimed `u`.

The table of that, at `eta = 0.35`, `N = 9`, `ln N = 2.1972`:

u*	N_eff	F(u*)	ln N − F	h_inf	% of ceiling claimed
0.00	1.000	0.0000	2.1972	0.7690	0.0%
0.25	3.000	1.0986	1.0986	0.3845	50.0%
0.40	4.200	1.4351	0.7621	0.2667	65.3%
0.55	5.400	1.6864	0.5108	0.1788	76.8%
0.70	6.600	1.8871	0.3102	0.1086	85.9%
0.85	7.800	2.0541	0.1431	0.0501	93.5%
1.00	9.000	2.1972	0.0000	0.0000	100.0%

A perfectly non-parallel life has zero headroom at arrival — it is already generating everything its category count permits, and the only remaining move is to raise `N`. That is [Curzi, 2026c]'s argument arriving from a completely different direction, and the convergence of the two is a point in the framework's favour, not a coincidence to be admired.

The dual reading is the harsh one. **An subject who feels enormous headroom at a settled state is not being told they have reserves. They are being told their framings are parallel.** `h_inf` is large precisely when `F(u*)` is small.

23.8.5 8.5 Where the leverage is: sensitivity of the fixed point

If arrival is determined by four numbers, the practical question is which of the four to move. Differentiate `Q_inf = (eta M* − c*)/delta`, at the base case `eta = 0.35`, `delta = 0.12`, `N = 9`, `u* = 0.55` (so `F = 1.6864`), `c* = 0.22`, giving `Q_inf = 3.0853`:

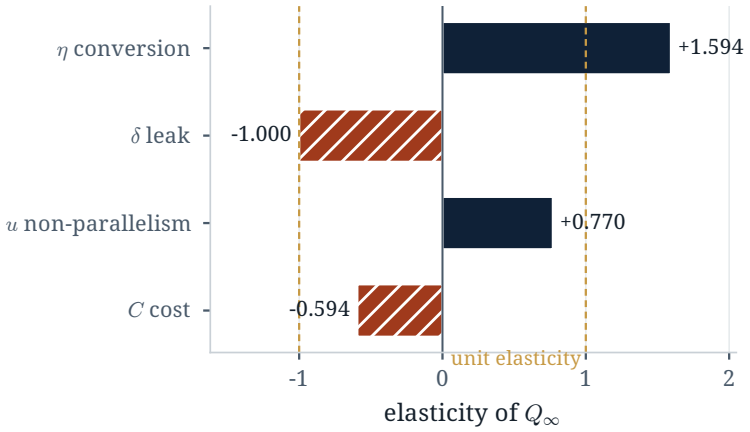


Figure 23.3: **Efficiency outranks cost by a factor of 2.68, and the leak elasticity is exactly -1 at every parameter value:** elasticities of the fixed point `Q_inf = (eta M* – c*)/delta` at the base point `eta = 0.35`, `delta = 0.12`, `N = 9`, `u* = 0.55`, `c* = 0.22`. The `delta` bar reaches `–1.000` because `Q_inf` is proportional to `1/delta`, so that elasticity is structural rather than local; the other three are base-point quantities. Derived from stipulated parameters, not fitted; generated by `figs/src/CCP-023_fig01.py`.

```
dQ_inf/d(eta)    =  F / delta                        =  14.0533
dQ_inf/d(C)      = −1 / delta                        =  −8.3333
dQ_inf/d(delta)  = −Q_inf / delta                    = −25.7111
dQ_inf/d(u)      =  eta(N−1) / ((1+(N−1)u) delta)    =   4.3210
```

Raw derivatives mislead because the levers are in different units. The elasticities — percent change in `Q_inf` per one percent change in the lever — are comparable:

Lever	Elasticity of `Q_inf`	Effect of a 10% improvement
`eta` (conversion)	**+1.594**	3.0853 -> 3.5772 (+15.94%)
`delta` (leak)	**-1.000**	3.0853 -> 3.4281 (+11.11%)
`u` (non-parallelism)	+0.770	3.0853 -> 3.3138 (+7.40%)
`C` (cost)	-0.594	3.0853 -> 3.2687 (+5.94%)

Three readings.

The elasticity with respect to `delta` is exactly -1, always. `Q_inf` is proportional to `1/delta`, so this is a structural fact and not a property of the base point: halve the leak, double the destination.

The elasticity with respect to `eta` is `eta*M*/(eta*M* – c*)`, which exceeds 1 whenever there is any cost at all, and diverges as the regime approaches unsustainability. At the base point it is 1.594. **A person operating near the sustainability boundary gets enormous returns from small improvements in how experience becomes character — and equally enormous losses from small degradations.** That is a real, derived, and slightly alarming property: fragility and leverage are the same number.

Cutting cost is the weakest of the four levers here, and it is usually the first one attempted. Elasticity -0.594 against `eta`'s +1.594: a 10% efficiency gain is worth 2.7 times a 10% cost cut at this base point. The ratio is `eta*M*/c*` and it favours efficiency whenever the regime is comfortably sustainable.

But there is a countervailing fact about `delta` that the elasticity hides, and §8.6 is it.

23.8.6 8.6 The cost of a long memory: noise amplification

Everything so far assumed constant drive. Real drive is not constant: `M_t` and `C_t` vary. The standard perturbation bound for a contraction: if the drive deviates from its mean by at most `epsilon` in each period, then in the long run

```
limsup_t | Q_t − Q_inf |  <=  epsilon / delta
```

The leak that determines how much you keep also determines how much of the input noise you keep. The amplification factor is `1/delta`, and it is the same `1/delta` that multiplies the signal.

`delta`	`1/delta`	wander at `eps=0.05`	at `eps=0.10`	at `eps=0.20`
0.020	50.00	2.500	5.000	10.000
0.050	20.00	1.000	2.000	4.000
0.100	10.00	0.500	1.000	2.000
0.150	6.67	0.333	0.667	1.333
0.200	5.00	0.250	0.500	1.000
0.300	3.33	0.167	0.333	0.667
0.500	2.00	0.100	0.200	0.400
0.750	1.33	0.067	0.133	0.267

The full trade, laid out on one line for three leak rates with the base case's drive:

`delta`	`Q_inf`	`tau`	periods to 5%	noise gain `1/delta`
0.06	6.1707	16.16	49	16.67
0.12	3.0853	7.82	24	8.33
0.24	1.5427	3.64	11	4.17

Halving the leak doubles the destination, doubles the time to reach it, and doubles the integrated noise. There is no free memory. A person with very low personal `delta` — who forgets nothing, whose skills do not atrophy, whose sense of self does not drift — accumulates a large `Q`, takes a very long time to change it, and carries every perturbation for twenty periods. That is not a defect of such a person; it is the same equation read on the other side.

The rule: **match `delta` to the volatility of the domain.** In a stable domain low `delta` is a gift; in a volatile one it is a liability, and deliberate shedding ([Curzi, 2026a]; the SEAL node [Curzi, 2026] SEAL) is the response.

23.8.7 8.7 Counterfeit arrival

The FIXEDPOINT node warned that a well-converged state can still be the wrong one, since convergence is a property of loops and not a certificate of worth. Theorem 3.5 already supplies the first version of that: **if `eta F(u*) < c*` the fixed point is negative, and the loop converges anyway — toward the destruction of the thing.** The storage-law paper does not hide this and neither does this paper.

This section supplies the second version, which is the one the manifest names: **a fixed point reached at another's expense — the room furnished with someone else's goods.**

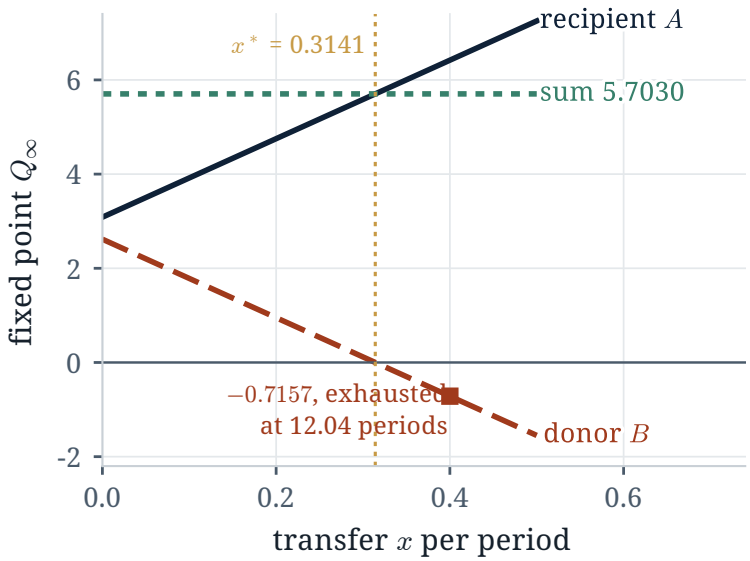


Figure 23.4: **Transfer relocates a fixed point and creates none:** the two fixed points move by $\pm x/\delta$ as the transfer rate rises, and their sum is 5.7030 at every rate. The donor crosses zero at the critical rate `x* = 0.3141`; at `x = 0.40` its steady state is `-0.7157` and it is exhausted after 12.04 periods, while the recipient reads far above anything its own regime could produce. Modelled from `eta = 0.35`, `delta = 0.12`, `C = 0.22`, `N = 9`, `u_A = 0.55`, `u_B = 0.45`; the script reproduces every row of this section's table; generated by `figs/src/CCP-023_fig03.py`.

8.7.1 The transfer model

Two systems, `A` and `B`, each running the storage law, with a transfer `x >= 0` per period from `B` to `A`. `A` receives `x` as a free addition to its drive; `B` pays it as an addition to its cost.

```
Q_A(t+1) = (1-delta) Q_A(t) + eta*F_A - C_A + x
Q_B(t+1) = (1-delta) Q_B(t) + eta*F_B - C_B - x
```

Both remain affine contractions, so both still converge, and the fixed points are

```
Q_A,inf(x) = ( eta*F_A - C_A + x ) / delta  =  Q_A,inf(0) + x/delta
Q_B,inf(x) = ( eta*F_B - C_B - x ) / delta  =  Q_B,inf(0) - x/delta
```

The transfer moves each fixed point by exactly `x/delta`, in opposite directions, and the sum is unchanged. Numerically, with `eta = 0.35`, `delta = 0.12`, `C = 0.22`, `N = 9`, `u_A = 0.55` (`F_A = 1.6864`), `u_B = 0.45` (`F_B = 1.5261`):

x	Q_A,inf	gain to A	Q_B,inf	sum
0.0000	3.0853	—	2.6177	5.7030
0.1000	3.9187	+0.8333	1.7843	5.7030
0.2000	4.7520	+1.6667	0.9510	5.7030
0.3141	5.7028	+2.6175	0.0002	5.7030
0.4000	6.4187	+3.3333	-0.7157	5.7030

The sum column is the finding. Transfer relocates a fixed point; it creates none. No amount of extraction increases the total stored order of the pair by one unit. This is not an ethical claim dressed as arithmetic — it is arithmetic, and its ethical reading is left to the reader, who already has the declination refusal in hand [Curzi, 2026].

8.7.2 The exhaustion horizon

There is a critical transfer rate at which B's fixed point hits zero:

```
x* = eta * F_B − C_B
```

which for the case above is `x* = 0.35 * 1.5261 − 0.22 = 0.3141`. This is simply Corollary 3.4 — the sustainability ceiling — read as a constraint on how much can be taken. **Below x*, B is diminished but survives at a lower fixed point. Above x*, B's fixed point is negative and B runs out.**

Take `x = 0.40 > x*`. B's fixed point is `−0.7157`. Starting from its undisturbed fixed point of 2.6177, B's trajectory is:

t	Q_B
0	2.6177
1	2.2177
3	1.5559
5	1.0434
7	0.6466
9	0.3393
11	0.1013
12	0.0032
13	**-0.0830**

The exhaustion horizon in closed form, from the error law:

```
t_exhaust = ln( −Q_B,inf / (Q_B,0 − Q_B,inf) ) / ln(1 − delta)   =  12.04 periods
```

Twelve periods. During all twelve, A is converging comfortably toward 6.4187 — a fixed point 108% higher than its own regime could produce — and every measurement A can take of itself says the practice is working.

8.7.3 What happens at the end

When B fails, the transfer stops. A's drive reverts and its fixed point drops back to 3.0853. A now has a gap of 3.3333 Q-units to close, downward, and it closes it geometrically at the same `delta`: **24 periods to come within 5% of the new, lower destination.**

```
A's peak fixed point during extraction : 6.4187
A's fixed point after B fails          : 3.0853
loss on collapse                       : 3.3333 Q-units  (51.9% of the peak)
twelve periods of gain, twenty-four of unwinding
```

The counterfeit arrival is not merely unjust; it is arithmetically temporary, and its unwind is twice as long as its build. That asymmetry is not moralising — it is the ratio of the exhaustion horizon to the reconvergence time, and both were computed from the same `delta`.

8.7.4 The test that distinguishes genuine from counterfeit arrival

Here is the operational point of the whole section, and it is one line.

An arrival is genuine if every party to it has `h_t >= 0` under its own regime, counting the transfer as a cost. It is counterfeit if any party has `h_t < 0`.

That is Corollary 3.4 used as an audit rather than as a forecast. The check is cheap. For each party, compute `h = eta ln N − delta Q − C_total` where `C_total` includes every transfer that party is paying. If any `h` is negative, then by Corollary 3.4 **no choice of interaction geometry prevents that party's order from decreasing** — and the arrangement's apparent stability is being paid for out of a store that is running down.

Two honest complications.

The audit requires the other party's ledger. `A` cannot run it from inside `A`; it needs `B`'s `Q`, `C`, and `delta`. **This is precisely why the check is skipped, and precisely why it is the check** — at personal scope, it is the arrangement whose sustainability you can only assess by asking a question you would rather not ask.

A negative `h` in a party who has not been asked is not established. The declination refusal governs [Curzi, 2026]. The audit is run by request, with a decline path, and a party who declines to supply their ledger is recorded as unknown, not as consenting. **An audit that conscripts the other person's data to reassure the person doing the extracting has reproduced the problem it was meant to detect.**

23.8.8 8.8 Arrival, stall, and the limit of what arithmetic can tell you

A flat `Q` has exactly one arithmetic cause: `eta M_t = delta Q_t + C_t`, the one-step growth condition of Corollary 3.2 holding with equality. That single equation covers both of the following:

- **Arrival.** The practice has converged to a fixed point that its subject chose, by choosing the regime that generates it.
- **Stall.** The practice has converged to a fixed point nobody chose, inherited from parameters nobody set, and the subject has been calling it stability.

The arithmetic cannot distinguish these. They are the same equation. The corpus states that plainly rather than manufacturing a discriminator, because the honest position is that the difference is not in the trajectory — it is in the record of how the parameters came to have their values.

What *can* be checked, and what the specification below requires:

1. **Was `Q_inf` computed before it was reached?** A chosen fixed point is one that was predicted. [Curzi, 2026i]'s survey and the four parameters give a forecast; a forecast recorded in advance and then met is evidence of a chosen regime.
2. **Is `h_inf = eta(ln N − F(u*))` known and accepted?** An subject at a chosen fixed point can state how much of the ceiling their geometry is claiming, and has a reason for the remainder.
3. **Is there a sunset?** the SEAL node [Curzi, 2026] and the shedding result of Curzi [2026h]: a chosen regime has a review date. A stall has never had one.

Arrival and stall are distinguished by documentation, not by measurement, and that is not a weakness of the model — it is the model correctly declining to certify an intention it cannot observe.

23.8.9 8.9 Changing where you arrive

The consolidated instruction, from §8.4, §8.5, and §8.6 together.

```
to move the destination      change eta, u (through F), C, or delta
to move faster to it         change delta -- and accept that this moves the destination too
to trade destination         raise eta or u; these move Q_inf without moving tau
  without changing speed
to raise the ceiling         raise N; this is the only move that raises
                             the maximum any geometry could reach
```

And the negative result: **trying harder inside a fixed regime does not change where you arrive.** Effort that moves none of the four parameters moves only the transient — the same path, the same room, the same rate, since the rate is `delta` and effort is not `delta`. The sensation of pushing is real and it is spent on the term that was going to decay anyway.

23.8.10 8.10 What would need measuring

Under the badging convention [Curzi, 2026], the agenda rather than the results:

1. **A `Q` instrument with ratio-scale properties.** Rasch-calibrated where possible ([Curzi, 2026t]), because specific objectivity is exactly the property that lets a `Q` measured with one item set be compared with a `Q` measured with another. Without it, FVP-2 stays conditional.
2. **`delta` by domain and by person.** Fit from the decay of successive gaps under a constant regime, per FVP-6, with the freeze step enforced. Report distributions, not point estimates.
3. **The Weber fraction for `Q`.** This determines `t_arrive` in §8.3 and it is currently a parameter of a table rather than a measured constant.
4. **The joint identifiability study.** Independent measurement of `M*` (magnitude estimation) and `c*` (cost ledger) in the same subjects, to test whether `eta` recovered by residual is stable — which is FVP-5 and FVP-7 attacked together.

23.9 9 Requirements for a convergence claim

CS-1 — Declare the period. Every `delta`, `tau`, and iteration count is meaningless without the period attached. The period is fixed before data collection and is not changed mid-block.

CS-2 — Compute `Q_inf` before pursuing it. `Q_inf = (eta M* − c*)/delta` from the four measured parameters, recorded in writing, before the practice is undertaken. *(This is the requirement that makes §8.8's first check possible at all.)*

CS-3 — Verify `delta > 0`. If the practice has no loss term — nothing is forgotten, nothing atrophies, no maintenance is required — there is no fixed point and the convergence tables do not apply. State the loss mechanism explicitly.

CS-4 — Check the sign of the fixed point. If `eta M* < c*` then `Q_inf < 0` and the regime is unsustainable. **Report this before proceeding; do not run the practice and discover it.** *(Theorem 3.5's own failure clause.)*

CS-5 — State the arrival criterion in advance. Declare `c`, the fraction of `Q_inf` at which the practice will be called arrived, and compute `t_arrive` from §8.3. An undeclared arrival criterion converts an asymptote into an argument.

CS-6 — Use the a posteriori bound to track progress. Distance from arrival is approximately `((1−delta)/delta)` times the last observed step. Report it each period; it costs one multiplication.

CS-7 — Compute `h_inf = eta(ln N − F(u*))` at the fixed point and account for it. Large headroom at arrival is a statement about parallel framings, not about reserves. *(§8.4.)*

CS-8 — Match `delta` to domain volatility. Where the domain changes faster than `1/delta` periods, deliberate shedding is specified, not optional. *(§8.6.)*

CS-9 — Run the transfer audit. For any arrangement in which one party's drive includes a transfer from another, compute `h` for **every** party with the transfer counted as that party's cost. Record any negative `h`. *(§8.7.4.)*

CS-10 — The audit has a decline path. No party is required to disclose their ledger. A declined audit is recorded as *unknown*, never as *satisfied*. *(The declination refusal, binding [Curzi, 2026].)*

CS-11 — Do not compare `delta` across people without recalibration. FVP-7 is open. A fitted `delta` is within-person and within-domain until the reproducibility study exists.

CS-12 — Freeze before claiming. FVP-6 in full, including the shuffled-input control.

23.10 10 Failure modes of arrival

Fault	Diagnostic signature	Fix
The unarrived asymptote	Subject reports "almost there" for many periods and treats it as failure	There is no arrival event (§8.3). Declare `c` in advance per CS-5 and use it. The last 0.9% costs what the previous 9% did
Effort substitution	Increasing effort with no change in `eta`, `u`, `C`, or `delta`; `Q` follows the same path	Effort acts on the transient, not the fixed point (§8.9). Identify which of the four parameters the effort is supposed to move; if none, stop
The negative fixed point run anyway	`eta M* < c*`; `Q` declining steadily; practice continued because "consistency matters"	CS-4. The loop converges toward destruction. Cut `C`, raise `eta` or `u`, or stop
Zero-delta fantasy	A plan with no loss term: nothing decays, nothing needs maintenance	No contraction, no fixed point, unbounded linear growth in the model and rapid divergence from reality. Name the loss mechanism (CS-3)
Low-delta noise integration	Very stable person or practice; large swings in `Q` traced to small perturbations months earlier	`1/delta` amplification (§8.6). Either accept it as the price of memory or raise `delta` deliberately by shedding
Headroom misread as reserve	Large `h` at a settled state, read as "capacity available"	`h_inf = eta(ln N − F(u*))`. Large `h` means parallel framings (§8.4). Raise `u` or `N`, do not "use the capacity"
Counterfeit arrival	A party's `Q` rising steadily with no change in their own `eta`, `u`, `C`, or `delta`	The extra drive is a transfer. Run CS-9 on every party. Sum is invariant (§8.7.1)
Extraction above `x*`	Counterparty's `Q` declining linearly toward zero on a finite horizon	Compute `x* = eta F_B − C_B` and `t_exhaust`. Above `x*` the arrangement has a countdown, and the unwind is longer than the build
Stall named arrival	Flat `Q`, no predicted `Q_inf` on record, no review date	§8.8. The arithmetic cannot adjudicate; the documentation can. If there is no forecast and no sunset, treat as stall

Fault	Diagnostic signature	Fix
Cadence confusion	Cadence increased, calendar unchanged, subject frustrated	`tau` depends on the per-period `delta`. Changing cadence without changing per-period leak changes nothing (§8.1). Cadence acts through `eta` and `C`
Cross-person `delta` comparison	"Her `delta` is lower than mine"	FVP-7 is untested. CS-11
Ordinal `Q`	Convergence rates fitted to a 1–10 self-rating	FVP-2. Differences on an ordinal scale are not differences. Report the ordering of successive gaps only, or build a calibrated instrument

---

23.11 11 The personal formulary

Sections 11 to 13 are a lookup table, and the payoff of the internal-scope material: every formula that applies at the **internal scope** — psychology, personal, one mind — collected in one place, in one format, with one unit system, each carrying its full Formula Validation Protocol so that a reader can tell in ten seconds which gates it passes and which it does not.

A formulary is not a theory. It is a shelf. Its virtue is that the entries are in the same units, formatted identically, and honest about their limits, so that a person holding a real problem can find the right instrument without reading the whole framework. The preceding material made the arguments; this one exists so the arguments do not have to be re-read.

23.11.1 11.1 How to use it

Three access paths. **By quantity:** the index in §11.2 maps each parameter (`u`, `eta`, `h`, `r`) to its entry. **By fault:** every entry ends with a fault table keyed to the felt complaint, and *The Failure Modes* below covers the faults that come from using the collection rather than any single entry. **By sequence:** *The Specification* below gives the dependency order for a full review. **Running the entries out of order produces numbers that look fine and mean nothing, because six of the nine depend on quantities the others produce.**

23.11.2 11.2 The nine entries

#	Name	Formula	Primary parameter	Source
F-01	The Moment	`M = (D/P) . E`	`M`	Curzi; [Curzi, 2026m]
F-02	The Calibration Ratio	`r = D/P`, `Delta = |ln r|`	`r`	[Curzi, 2026m]
F-03	The Internal Surplus	`F_int = ln(1 + (N−1)u)`	`u`, `N_eff`	[Curzi, 2026,o]

#	Name	Formula	Primary parameter	Source
F-04	The Self-Storage Law	`Q_{t+1} = (1−delta)Q_t + eta M_t − C_t`	`Q`, `eta`, `delta`	the storage law [Curzi, 2026a]; [Curzi, 2026o]
F-05	Personal Headroom	`h_t = eta ln N − delta Q_t − C_t`	`h`	the headroom identity [Curzi, 2026b]; [Curzi, 2026b], [Curzi, 2026o]
F-06	The Launch Budget	`L = max(0, M − C)`	`L`	[Curzi, 2026o,p]
F-07	The Gradient	`g = −grad C`	`g`	[Curzi, 2026,p]
F-08	The Return Difference	`Delta_R = |ln(M'/M)|`	`rho`	Sections 3–6
F-09	The Contraction	`|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|`	`Q_inf`, `tau`	[Curzi, 2026a]; Sections 7–10

Four of the nine are **theorems** of the corpus (F-03, F-04, F-05, F-09) and appear here in their internal-scope reading. Two are **the author's formula** and its immediate consequence (F-01, F-02). Three are **derived formulas of the corpus** (F-06, F-07, F-08), derived in the sense the internal-scope correspondence requires [Curzi, 2026o], and badged where their derivations pass through a postulate.

Every entry states which of the three it is in its first line, because the difference decides how much weight it can carry.

23.11.3 11.3 The entry format

Every entry has the same eight parts, in the same order, so the formulary can be read as a table:

```
STATEMENT        the formula, in a fence, with every symbol defined
READS            one sentence: what the formula says in words
STATUS           theorem / author's formula / derived-with-badge
DERIVATION       where it comes from in the corpus, and what it needs
FVP              all seven gates, verdict plus content            (Section 11.6)
PROTOCOL         numbered measurement procedure
WORKED NUMBER    one computed instance
FAULTS           named failure modes with signature and fix
```

23.11.4 11.4 The unit system

Everything internal is measured in **nats**, or in **Q-units** (the corpus's unit of stored order, defined in [Curzi, 2026b]), or in **Q-units per period**. Because the four registers converge [Curzi, 2026n,?], any quantity in nats is immediately readable in three other registers by exact constants:

```
1 nat    = 1.442695 bits = 2.885390 octaves = 3462.4681 cents
1 bit    = 0.693147 nats = 2 octaves EXACTLY = 2400 cents exactly
1 octave = 0.346574 nats = 0.5 bits          = 1200 cents

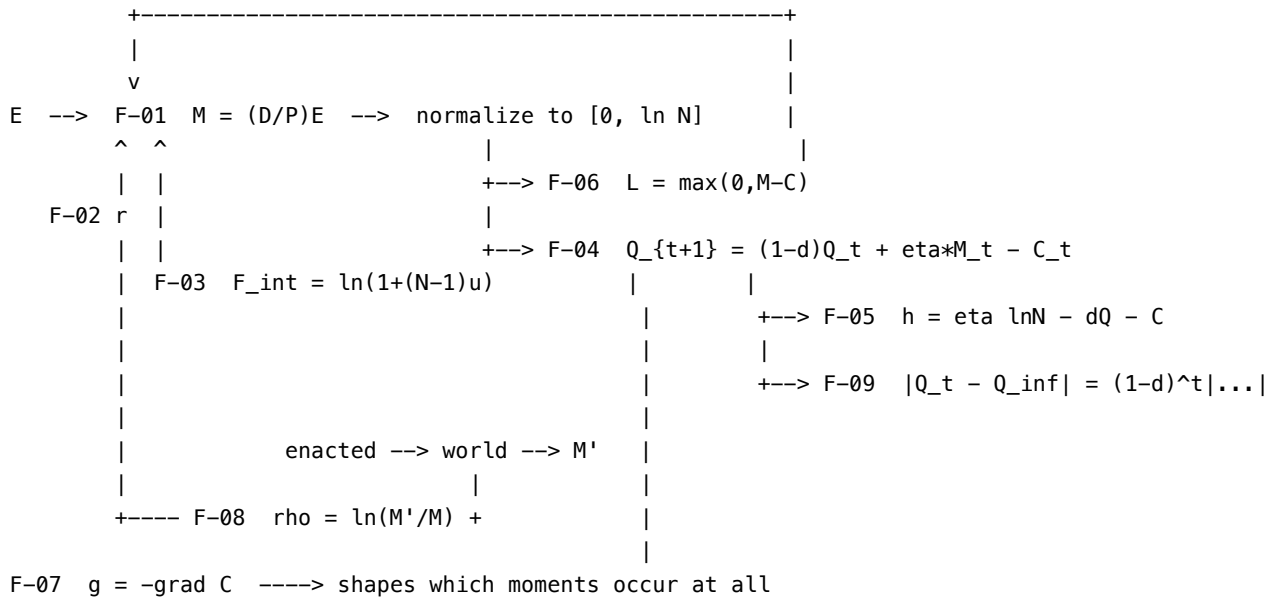
surplus -> frequency:  f = k * e^(2F),   k = 0.483425732988414 Hz   [the frequency constant]
```

The ceilings, which bound five of the nine entries:

N	ln N (nats)	bits	octaves	f = k e^(2 ln N) (Hz)
2	0.6931	1.0000	2.000	1.9337
4	1.3863	2.0000	4.000	7.7348
6	1.7918	2.5850	5.170	17.4033
7	1.9459	2.8074	5.615	23.6879
9	2.1972	3.1699	6.340	39.1575
10	2.3026	3.3219	6.644	48.3426
16	2.7726	4.0000	8.000	123.7570

`ln N` is the hard ceiling on surplus per interaction [Curzi, 2026], and by the forced consequence of the register convergence it is also the ceiling on the moment: the richest available moment is bounded by the log of how many kinds of thing you can distinguish. That is Ashby's requisite variety appearing in first-person experience. It was not chosen; it follows.

23.11.5 11.5 The circuit — how the nine connect



Three structural facts follow, and they govern the order of operations in *The Specification*.

F-03 sets the ceiling F-01 lives under — no moment is richer than `ln N`, and `N_eff` says how much of `N` is real. **F-01 is the only producer:** storage, headroom, launch and arrival are all fed by `M`, so **if the Moment is mismeasured, all six downstream numbers are wrong in the same direction and will look consistent with one another, which is the most dangerous kind of wrong. F-08 is the only feedback path, and it acts on `D`, not on `Q`** — the echo corrects the frame, which changes the integrator's input, not its state, which is why feedback improves character slowly (Sections 3–6, §4.10).

23.11.6 11.6 The FVP verdict legend

The Formula Validation Protocol [Curzi, 2026t] defines the seven gates. This paper uses six verdict words and they are not interchangeable:

Verdict	Meaning
PASS	The gate is satisfied as the formula is stated here
CONDITIONAL PASS	Satisfied only under a stated condition; if the condition fails, the formula is **demoted in print** to a heuristic ordering
DEMOTE	The gate is **not** satisfied as things stand. The formula's structure — limits, fixed points, orderings — remains in force; its magnitudes are withdrawn as measured quantities until the named instrument exists
PARTIAL	Satisfied for some parameters and not others; the unrecoverable ones are named
SPECIFIED	Not yet run, but the exact procedure that would run it is given here
UNTESTED — declared	Not run, no data, and the corpus says so rather than implying otherwise

Nine entries, sixty-three gates, and not one is marked PASS on the reproducibility gate, because the corpus has run no cross-individual study of any of these quantities. The protocol names FVP-7 as the gate that decides the hard-science question and the one most often skipped; a formulary's honest option is to leave it visibly open on every page rather than quietly on none.

11.6.1 Two standing verdicts that govern more than one entry

Two findings arrived after the entries below were drafted, and both cut across several of them. They are stated once here, in force everywhere, rather than repeated nine times.

STANDING VERDICT 1 — every `M`-derived entry is DEMOTED on gate 2. [Curzi, 2026t] run the Formula Validation Protocol on the Moment Equation and return **DEMOTE on FVP-2**: `D` and `P` are not instrumented on a common ratio scale with a shared true zero, and division without a true zero yields a number with no referent. This governs **F-01, F-02, F-04, F-06 and F-08** — every entry whose value passes through `M` or through `r = D/P`.

> **NOT REPORTABLE AS MEASURED QUANTITIES:** the magnitude of `D/P`; the distortion `Delta = |ln r|` as a number; the quantitative form of the `M -> Q` join (F-04). Every such figure in the worked numbers below is an **illustrative construction demonstrating the arithmetic**, and is labelled as such at its use.
>
> **STILL IN FORCE:** the limiting cases; the `D = P` fixed point; and the ordering *over-framed / calibrated / under-framed*. **These require only ordinal information and the demotion does not touch them.** What is demoted is magnitude, not structure.

The route out is constructive, specific and priced. [Curzi, 2026t], **Theorem 54.3**, given badged Postulate SO1: measure `D` and `P` as Rasch person-measures on one common linked scale, and `Delta = |ln(D/P)| = |theta_D − theta_P|` becomes genuinely ratio-scaled — Rasch separation supplies the true zero that division requires. The cost is **approximately 123 items per term to establish a half-logit distortion at 95 per cent**, and about 31 items for a full logit. A twelve-item pair of scales resolves nothing finer than roughly 1.6 logits. **A formulary that demotes five of its own nine entries, names exactly what survives, and prices the study that would restore the rest is the protocol working. It is not a caveat and it is not an apology.**

STANDING VERDICT 2 — the shuffled-input control has three measured failure modes, and every FVP-6 row below inherits all three. The four-baseline suite of The four-baseline suite

[Curzi, 2026i] is cited in every falsification row in this paper. [Curzi, 2026u,ac] measured what it actually catches, by simulation rather than assertion:

#	Failure mode	Measured
1	**Anti-conservative as ordinarily run.** Against a nominal `alpha = 0.05` across 500 null trials, iid permutation false-rejected at **0.108 — 2.2x nominal**; circular shift at 0.120; block permutation at 0.060, near nominal	[Curzi, 2026u]
2	**Blind to repackaging.** A filtered repackaging of the input passed all four baselines including the shuffle at `p = 0.0005`, and was caught **only by the structural ceiling test** (the structural ceiling corollary [Curzi, 2026a]) — its single-step jump was **38.5x** the theoretical maximum	[Curzi, 2026i], [Curzi, 2026u]
3	**Blind to shared slow drivers.** In 4,000 trials the shuffle reported honestly in the clean case (48.17% detection) and **passed 98.35% of the time against a confound driving both series**	[Curzi, 2026ac]

> **THE RULE, AND IT SUPERSEDES ANY EARLIER FORM: the shuffle is a weak, miscalibrated empirical audit. Repackaging is caught only by the structural ceiling test; confounding is caught only by design controls — blinding, channel separation, prospective sampling, a yoked control. No surrogate catches either.** Calibrate the surrogate by simulation on your own design before you use it, and **design controls, not more surrogates.** [Curzi, 2026u] specifies both.

23.11.7 11.7 The refusals, as they apply to a formulary

A formulary is a set of numbers a person may come to steer by, and a measure that becomes a target ceases to be a good measure [Goodhart, 1975, Campbell, 1979]. That is the reason the verdict legend above admits DEMOTE and UNTESTED at all: a formulary in which every entry passed every gate would be evidence of capture rather than of quality. The badged postulates below are doing the work Lakatos [1970] describes for a research programme's hard core, and saying so is more honest than presenting them as findings.

The framework's four refusals [Curzi, 2026] govern every entry. Three of the four refusals have specific formulary consequences:

No unbadged postulates — every entry whose derivation passes through an assumption states it in the entry. **No uncomputed numbers** — every worked number here was computed, from data that are **illustrative constructions**, stated as such at each use. **No intervention that**

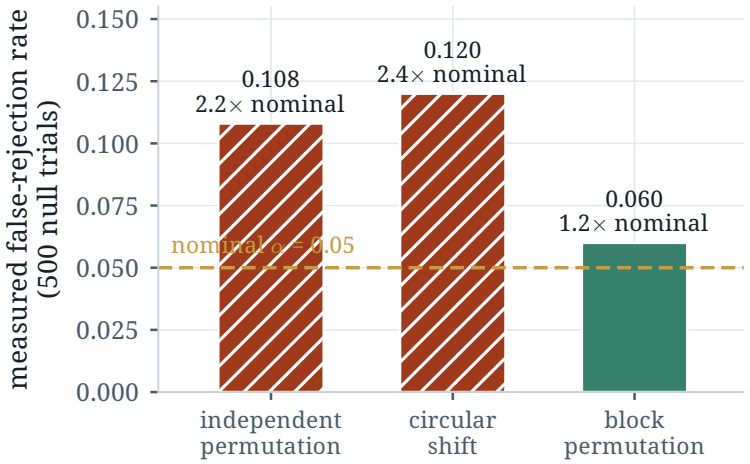


Figure 23.5: **The audit instrument is anti-conservative as ordinarily run:** measured false-rejection rates of the shuffled-input control over 500 null trials against a nominal $\alpha = 0.05$. Independent permutation returns 0.108 and circular shift 0.120 — 2.2 and 2.4 times nominal — and only block permutation is near nominal at 0.060. The control's other two measured failures are not displayable on this axis and are stated instead: a filtered repackaging of the input passed all four baselines at `p = 0.0005`, caught only by a structural ceiling test at 38.5 times the theoretical maximum single-step jump; and against a confound driving both series the shuffle passed 98.35 per cent of 4,000 trials. Measured by simulation in Curzi [2026u] and Curzi [2026ac]; transcribed, not recomputed; generated by `figs/src/CCP-023_fig05.py`.

removes the ability to decline — at internal scope other people appear in F-08's echo sources and in F-05's cost ledger, and both are governed: **an echo source may decline and is recorded as missing data, never as a null return; and a cost carried by someone else is counted in their ledger, by asking, with a decline path.**

23.11.8 11.8 The nine entries

23.11.9 F-01 · THE MOMENT

STATEMENT

```
M = ( D / P ) . E

D  perspective   the frame brought to the encounter
P  perception    the intake taken from the encounter
E  environment   what is there
M  the moment    experience as constructed

r = D/P must be dimensionless, so D and P must share units.
M then inherits E's units.
Normalized form for use downstream:  M_norm in [0, ln N]
```

READS. *The moment is the environment, scaled by how far your frame outruns or falls short of your intake.*

STATUS. The author's formula (Curzi), analysed in full in [Curzi, 2026m]. Not a theorem of the framework's mathematical papers; it is the object Series B of the corpus is built to make measurable.

DERIVATION. `M` is not derived from the surplus law — it is the corpus's second formal object, joined to the first through the storage law (F-04), where `M` supplies the input `S_t`: *perception and perception produce the moment; the moment becomes character* [Curzi, 2026o]. The ceiling `M <= ln N` is not assumed but forced, once `M` is used as `S_t`, by `0 <= S_t <= ln N` (the surplus range result [Curzi, 2026]).

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`D` and `P` in shared units (whatever the estimation procedure supplies); `r` dimensionless; `M` in `E`'s units. For downstream use, `M` is normalized into nats on `[0, ln N]`, and the normalization must be stated.
FVP-2 SCALE TYPE	**DEMOTE — the crux** (Standing Verdict 1)	Division requires ratio scale with a true zero for `D` and `P`, on a **common** scale. No such instrument yet exists, so **the magnitude of `D/P` and every quantity derived from it are withdrawn as measured quantities ([Curzi, 2026t], [Curzi, 2026t]).** `M = (D/P)E` reads as the ordering *over-framed / calibrated / under-framed*, plus its limits and fixed point — all of which need only ordinal information and none of which is affected. **Route out: [Curzi, 2026t] Thm 54.3**, Rasch logit difference given badged Postulate SO1, at ~123 items per term for half a logit.
FVP-3 LIMITING CASES	PASS	`P -> 0`: `M -> inf` — frame with no intake: hallucination, dogma, apophenia. `D -> 0`: `M -> 0` — intake with no frame; **sensation is not experience**. `E -> 0`: `M -> 0` — no environment, no moment. Each divergence names a real phenomenon, which is what the gate demands.

Gate	Verdict	Content
FVP-4 FIXED POINTS	PASS	`D = P` gives `M = E`: **undistorted contact.** The formula's most important content — when frame matches intake, the moment equals the environment.
FVP-5 IDENTIFIABILITY	**PARTIAL**	`M` and `r` are recoverable. `D` and `P` **individually are not identified from `M` alone** — any pair with the same ratio and the same product-with-`E` predicts identically. Recovering them separately requires manipulating one while holding the other: change the frame with the environment fixed, or change the environment with the frame fixed.
FVP-6 FALSIFICATION	SPECIFIED	The preregistration protocol [Curzi, 2026i]. Preregister the estimation procedure for `D`, `P`, `E`. Fit on a training block, freeze, predict held-out `M`. Beat persistence, direct-input (`E` alone), matched AR(1), and the shuffled-input control. **Killed if `E` alone predicts `M` as well as `(D/P)E` does** — that outcome says the frame ratio carries nothing.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Cross-individual recovery of `D`, `P`, `E` under a common operationalization has not been run. [Curzi, 2026t]'s Rasch approach — separating person parameters from item parameters — is the route, and until it exists `M` is a within-person quantity.

PROTOCOL

1. Fix a **modulus**: assign a magnitude (commonly 50) to one reference encounter and hold it for the whole block.
2. **Before** the encounter, estimate `D`: how much structure you are bringing — expectation, model, prior commitment — against the modulus.
3. **During or immediately after**, estimate `P`: how much you actually took in.
4. Estimate `E` separately, ideally from a source that is not you (F-08's independent sources serve).
5. Compute `r = D/P` and `M = rE`.
6. Normalize: `M_norm = ln N * (M − M_min)/(M_max − M_min)` with the range stated in advance.

WORKED NUMBER

D	P	E	r	M	|ln r|	reading
50	50	60	1.000	60.000	0.0000	calibrated contact, `M = E`
75	50	60	1.500	90.000	0.4055	over-framed by half
30	50	60	0.600	36.000	0.5108	under-framed
50	10	60	5.000	300.000	1.6094	intake collapsing; the `P -> 0` limb
50	100	60	0.500	30.000	0.6931	flooding; twice the intake the frame can use

Normalization example: `M = 90` on a declared range `[0, 120]` with `N = 7` gives `M_norm = 1.9459 * 90/120 = 1.4594` nats — 75% of the ceiling.

FAULTS

Fault	Signature	Fix
Ordinal division	`M` computed from star ratings	FVP-2. Demote to the three-way ordering; the numbers already computed are void
Circular `E`	`E` estimated from the same impression that produced `P`	Get `E` from an independent source; otherwise `r` is structurally 1 by construction
Drifting modulus	`M` trends across a block for no reason anyone can name	Freeze the modulus (Protocol step 1) and re-baseline blocks separately
Ceiling violation	`M_norm > ln N`	The normalization range was wrong, or `N` was overstated. Recheck F-03 first

CROSS-REFERENCES. [Curzi, 2026m] (full treatment), [Curzi, 2026t] (FVP-2), F-02 (the ratio), F-04 (where `M` goes), F-06 (what `M` funds).

23.11.10 F-02 · THE CALIBRATION RATIO

STATEMENT

```
r = D / P                          the calibration ratio, dimensionless

Delta = | ln r |                   the distortion, in nats

r > 1   OVER-FRAMED    moment exceeds environment  -> projection, seeing the expected
r = 1   CALIBRATED     moment equals environment   -> contact
r < 1   UNDER-FRAMED   moment falls short          -> flooding, cannot make it mean
```

READS. *How far your frame outruns your intake, measured as a log-ratio so it composes with everything else.*

STATUS. Immediate consequence of F-01; the distortion measure is the calibration ratio's [Curzi, 2026m].

DERIVATION. `r` is the coefficient in `M = rE`. `Delta = |ln r|` is a log-ratio for the same three reasons `Delta_R` is (Sections 3–6, §3.3): ratios compose multiplicatively, the measure must be invariant to the modulus, and the reading instrument is logarithmic. **`Delta` and `Delta_R` are the same kind of object measured at two places in the circuit — one on the frame, one on the world's answer — and that is why they are directly comparable.**

Independent convergence, stated carefully: predictive-processing accounts model perception as precision-weighting of top-down priors against bottom-up evidence, and **the ratio of prior precision to sensory precision is structurally the same object as `D/P`.** That is a correspondence between formalisms, not an identity claim — predictive processing is an active and contested research program, and structural similarity is not identity.

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`r` dimensionless by construction (FVP-1 of F-01 supplies the shared units). `Delta` in nats.
FVP-2 SCALE TYPE	**DEMOTE** (Standing Verdict 1)	Inherits F-01's verdict exactly. **`Delta = |ln r|` is an ordinal quantity pending instrumentation: its sign is a finding, its rank order is a finding, its magnitude is not.** The nats, bits and cents printed below are illustrative constructions showing how the measure would read, not measurements of a person. Ratio scale is restored by [Curzi, 2026t] Thm 54.3 and by nothing cheaper.
FVP-3 LIMITING CASES	PASS	`r -> inf` (`P -> 0`): total frame dominance. `r -> 0` (`D -> 0`): no frame; `Delta -> inf` in both, correctly, since neither limb has a finite worst case.
FVP-4 FIXED POINTS	PASS	`r = 1`, `Delta = 0`: veridical perception stated as an equation. Under a correction rule `ln r_{t+1} = (1−lambda) ln r_t` (the F-08 dynamics applied to the frame) this fixed point is attracting for `0 < lambda < 2`.
FVP-5 IDENTIFIABILITY	PASS **for `r`**	`r` is identified from `M/E` without separating `D` and `P` — a strictly better position than F-01, and the reason this entry is often more useful than its parent.

Gate	Verdict	Content
FVP-6 FALSIFICATION	SPECIFIED	Predict `M` from `E` and a frozen `r` estimated on a training block; beat the four baselines. **Killed if `r` estimated on block one does not improve prediction on block two over using `r = 1`.**
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Do two investigators recover the same `r` for the same person? Unknown.

PROTOCOL

1. Over `n >= 8` encounters, record a predicted magnitude and a realized magnitude on a fixed modulus (this is F-01's protocol run in prediction mode).
2. Compute `ln r_i` for each; take the mean and standard deviation.
3. `r_hat = exp(mean ln r)`. Report the interval, not the point.
4. Compare `Delta = |mean ln r|` against your deadband from F-08 (`ln(1+c)`); a distortion below the deadband is not resolvable.

WORKED NUMBER

Eight illustrative trials — predicted magnitude against realized magnitude on a fixed modulus:

```
predicted: 70  65  80  75  85  60  90  72
realized : 52  50  58  61  60  49  62  55

mean ln r = 0.2851        sd = 0.0620        SE = 0.0219
r_hat     = 1.3299
95% CI on ln r = [0.2421, 0.3280]  ->  r in [1.274, 1.388]
```

A stable over-frame of about 1.33 with a tight interval that excludes 1 decisively. In the conversion registers: `Delta = 0.2851` nats = 0.4114 bits = 987 cents — a shade under a minor seventh of distortion, and the four-register reading is not decoration; it is a check that the number is in a plausible range for a perceptual quantity.

Reference table:

`r`	`|ln r|`	bits	cents	state
0.50	0.6931	1.0000	2400.0	under-framed
1.00	0.0000	0.0000	0.0	**calibrated**
1.25	0.2231	0.3219	772.6	over-framed
1.50	0.4055	0.5850	1403.9	over-framed
2.00	0.6931	1.0000	2400.0	over-framed

Note the symmetry the log gives you for free: `r = 0.5` and `r = 2` are equally distorted. **Under-framing by half and over-framing by double are the same size of error, and only the log-ratio says so.**

FAULTS

Fault	Signature	Fix
Correcting a distortion below the deadband	Effortful frame-work producing no change	Compare `Delta` to `ln(1+c)`. If below, the instrument cannot resolve it; the work is unmeasurable by construction
Direction confusion	Treating flooding as over-framing	`r < 1` is under-framed — too much intake for the frame. The correction is more frame, not less
Single-encounter diagnosis	`r` estimated from one memorable event	`n >= 8` and report the interval (Protocol 3)
Global `r`	One `r` claimed for a whole person	`r` is domain-specific. A person may be over-framed at work and under-framed at home; a pooled `r` averages them into a fiction

CROSS-REFERENCES. [Curzi, 2026m] (full treatment), the group-scope treatment of the same ratio (doctrine-to-data), F-01, F-08.

23.11.11 F-03 · THE INTERNAL SURPLUS

STATEMENT

```
F_int = ln( 1 + (N−1) u )  =  ln N_eff

N       distinguishable categories held by ONE mind
u       non−parallelism between one's OWN framings, in [0,1]
N_eff   = 1 + (N−1)u, the audited count of genuinely different views
F_int   internal surplus, in nats,  0 <= F_int <= ln N
```

READS. *How many genuinely different views you hold, logged — the many windows of the one tower, counted after the audit for redundancy.*

STATUS. Theorem. The surplus functional [Curzi, 2026], applied within a single observer. Not an analogy: the axioms S1–S4 make no reference to the states belonging to different people.

DERIVATION. `f(u) = ln(1 + (N−1)u)` is uniquely characterized by S1 (geometric dependence), S2 (`f(0) = 0`), S3 (affine effective count), S4 (`f(1) = ln N`). **S3 is the design axiom — it is chosen, not derived, and that must be said wherever the law is leaned on.** The third node of the foundations series HORIZON — three futures cut from three non-parallel glasses — is this theorem applied within one observer, which makes it a derivation rather than an analogy.

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`N` a count; `u` dimensionless in `[0,1]`; `N_eff` a count; `F_int` in nats.

Gate	Verdict	Content
FVP-2 SCALE TYPE	PASS **for `u`**	`u = 1 − (x.y)^2` is computed from normalized state vectors and is ratio-scale with a true zero (parallel). The **operationalization** of the framings as vectors is where scale trouble enters, and [Curzi, 2026o]'s protocol is the route.
FVP-3 LIMITING CASES	PASS	`u → 0`: `N_eff → 1`, `F_int → 0` — **all your framings are one framing, counted `N` times, at `N` times the cost.** `u → 1`: `N_eff → N`, `F_int → ln N` — orthogonality, the ceiling. `N → 1`: `F_int = 0` for all `u`; one category has nothing to be non-parallel to.
FVP-4 FIXED POINTS	PASS	Extremality [Curzi, 2026]: `F = 0` iff parallel, `F = ln N` iff orthogonal. Both are attained, both are meaningful, and the lower one is the collapse the corpus spends [Curzi, 2026d,h] preventing.
FVP-5 IDENTIFIABILITY	**PARTIAL**	`F_int` is identified given `N` and `u`. `N` and `u` **trade off**: a small `N` at high `u` and a large `N` at low `u` produce identical `F_int`. Separating them requires counting the categories independently of measuring their overlap.
FVP-6 FALSIFICATION	SPECIFIED	The surplus functional [Curzi, 2026] is pure mathematics and makes no empirical claim; what is falsifiable is the **operationalization**. Preregister the framing-elicitation procedure, fit, freeze, and test whether measured `u` predicts held-out performance on a task requiring multiple framings, against the four baselines including shuffled-input.

FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Do two investigators eliciting one person's framings recover the same `u`? This is exactly the specific-objectivity question of [Curzi, 2026t] and it is open.

PROTOCOL

1. Elicit framings **independently and in writing before any are compared** — the personal-scope form of the silent-write-first rule, which prevents the framings from converging during elicitation.
2. Count them: that is `N`, and it is an upper bound, not a measurement.
3. Estimate pairwise overlap; `u` is one minus the squared cosine, averaged or taken pairwise per [Curzi, 2026o].
4. Compute `N_eff` and `F_int`. **Report `N_eff`, never `N`, when describing how many views you hold.**

WORKED NUMBER

`N = 7` framings, `ln N = 1.9459` nats:

u	N_eff	F_int	% of ceiling	marginal dF/du
0.00	1.000	0.0000	0.0%	6.0000
0.10	1.600	0.4700	24.2%	3.7500
0.30	2.800	1.0296	52.9%	2.1429
0.45	3.700	1.3083	67.2%	1.6216
0.60	4.600	1.5261	78.4%	1.3043
1.00	7.000	1.9459	100.0%	0.8571

The marginal value of the `k`-th genuinely orthogonal framing is `ln k − ln(k−1)`:

```
k=2  0.6931 nats (1.0000 bits)     k=7   0.1542 (0.2224)
k=3  0.4055        (0.5850)        k=8   0.1335 (0.1926)
k=4  0.2877        (0.4150)        k=9   0.1178 (0.1699)
k=5  0.2231        (0.3219)        k=10  0.1054 (0.1520)
k=6  0.1823        (0.2630)        k=12  0.0870 (0.1255)
```

The second genuinely different framing is worth one full bit — the largest single gain you will ever buy. The twelfth is worth 0.126 bits, an eighth as much, and [Curzi, 2026o] calls the fortieth glazing bills. Both facts are the same logarithm, and neither is pessimism: concavity is the lens [Curzi, 2026n].

FAULTS

Fault	Signature	Fix
Counting `N` instead of `N_eff`	"I look at this from six angles"	Measure `u`. Six angles at `u = 0.10` is `N_eff = 1.5`
Pseudo-diversity	Framings differ in vocabulary, agree in prediction	Run the collision test: do the framings ever disagree about an outcome? If not, `u ~ 0`

Fault	Signature	Fix
The N/u trade	Claiming a high F_int without stating which term produced it	FVP-5. Report both, always
Ceiling denial	Effort spent raising u when u is already 0.9	At u = 0.9, F_int is 95.4% of ceiling; the remaining move is to raise N ([Curzi, 2026c])

CROSS-REFERENCES. [Curzi, 2026o] (full treatment), [Curzi, 2026c], [Curzi, 2026c], [Curzi, 2026d], [Curzi, 2026h]; F-01 (ln N bounds M), F-05 (h is eta times the unclaimed part of this).

23.11.12 F-04 · THE SELF-STORAGE LAW

STATEMENT

```
Q_{t+1} = (1 − delta) Q_t + eta ∗ M_t − C_t

Q_t     stored order: character, skill, memory, internalized structure  [Q-units]
M_t     the moment as lived, normalized to [0, ln N]                    [nats]
delta   retention loss per period, in [0,1]                             [1/period]
eta     conversion of experience into durable structure, >= 0           [Q-units/nat]
C_t     cognitive load, effort, maintenance of self, >= 0               [Q-units/period]

Admissibility:      C_t <= (1−delta) Q_t + eta M_t   guarantees Q_{t+1} >= 0
One−step growth:    Q_{t+1} > Q_t  <=>  eta M_t > delta Q_t + C_t
```

READS. *Character is what is left of every moment you have had, after the leak, the conversion loss, and the cost of maintaining yourself.*

STATUS. Theorem. The storage law [Curzi, 2026a], in the internal-scope form the correspondence specifies [Curzi, 2026o]. **The join between the two formulas: the Moment feeds the storage law as its input.**

DERIVATION. Direct substitution of M_t for S_t, licensed by the register convergence's forced consequence that M is bounded by ln N and so shares S_t's range. **The scale-free claim — that a mind and a city run the same equation with different parameters — is a claim about MODEL FORM, badged as such under the corpus's badging convention [Curzi, 2026].**

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	Every term in Q-units per period once eta carries Q-units/nat. delta Q_t is Q-units/period. Both sides agree. **eta's units are the conversion key and stating them is what stops the equation from being a slogan.**

Gate	Verdict	Content
FVP-2 SCALE TYPE	**DEMOTE** (Standing Verdict 1)	`Q` must be at least interval-scale for the difference terms and ratio-scale for `Q_inf` percentages. **Badged: that `Q` is a scalar at all — rather than a vector of incommensurable stores — is a postulate of the corpus (Appendix E), not a finding.** And the input term inherits F-01's demotion: **the `M -> Q` join is structural, not yet quantitative.** The recursion, the two-track ledger, the fixed point and the orderings stand; numbers obtained by feeding a measured `r` through `M` do not. Rasch calibration ([Curzi, 2026t] Thm 54.3, ~123 items per term) is the route to satisfying both.
FVP-3 LIMITING CASES	PASS	`delta -> 0`: no fixed point; `Q_t <= Q_0 + t eta ln N` (Thm 3.3) — the tradition that never forgets and never settles. `delta -> 1`: `Q_{t+1} = eta M_t - C_t`, total amnesia, the state is only the present. `eta -> 0`: `Q` decays to `-C/delta` — **rumination is recursion with `eta ~ 0`.** `C -> 0`: no maintenance cost, `Q_inf = eta M/delta`.
FVP-4 FIXED POINTS	PASS	`Q_inf = (eta M* - c*)/delta` (Thm 3.5). **If `eta M* < c*` the fixed point is negative and the regime is unsustainable or the linear model must be replaced. The theorem does not hide this and neither does this entry.**
FVP-5 IDENTIFIABILITY	**PARTIAL**	From `{Q_t}` alone, `delta` is identified from successive gaps and the composite `eta M* - c*` from `delta Q_inf`. **`eta`, `M*`, `c*` are not separately identified.** Measure two of the three independently (F-01 for `M`, [Curzi, 2026b]'s ledger for `C`) and recover the third.

Gate	Verdict	Content
FVP-6 FALSIFICATION	SPECIFIED	The preregistration protocol [Curzi, 2026i]. Preregister period, `Q` instrument, window. Fit `delta` and `eta` on training, **freeze**, forecast held-out `Q`. Beat persistence, direct-input, matched AR(1), shuffled-input control. **Killed if the frozen model does not beat the shuffled-input control**, since that control destroys exactly the accumulation the law claims.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Do `delta` and `eta` recover across people and across domains within a person? Unknown.

PROTOCOL

1. Declare the **period** and hold it.
2. Instrument `Q` — a calibrated measure of the stored capacity in question, not a mood rating.
3. Log `M_t` per F-01 and `C_t` per [Curzi, 2026b]'s cost ledger, in the same units.
4. Fit `delta` from the decay of `Q` across a deliberate zero-input window (a period with `M ~ 0`), which is the cleanest identification available: with `M = 0`, `Q_{t+1} = (1−delta)Q_t − C_t`.
5. Recover `eta` by residual, and **report it with the identifiability caveat attached**.
6. Check admissibility every period. A violated admissibility condition is a broken model, not a bad week.

WORKED NUMBER

Illustrative: `N = 7`, `u = 0.45` (so `M* = F_int = 1.3083`), `eta = 0.30`, `delta = 0.10`, `C = 0.25`, `Q_0 = 0.80`.

```
Q_inf = (0.30 * 1.3083 − 0.25) / 0.10 = 1.4250 Q−units
one−step growth test at Q = 1.5:  eta*M = 0.3925  vs  delta*Q + C = 0.4000  ->  NO GROWTH
sustainability ceiling at Q = 1.5 (Cor 3.4):  C <= eta lnN − delta Q = 0.4338
```

A ten-period trajectory with a varying moment, including two poor periods:

t	M_t	Q_t	step dQ
0	0.45	0.8000	-0.1950
1	1.10	0.6050	+0.0195
3	0.20	0.7771	-0.2677
4	1.60	0.5093	+0.1791
5	1.60	0.6884	+0.1612
6	1.60	0.8496	+0.1450
7	0.90	0.9946	-0.0795
9	1.60	1.0536	+0.1246
10	—	1.1783	—

One bad period at `t = 3` (`M = 0.20`) cost 0.2677 Q-units, which took the next two good periods to recover. That asymmetry is not psychological; it is `eta M − delta Q − C` evaluated at a low `M`, and it is why cadence beats intensity: the leak runs whether or not the input does.

FAULTS

Fault	Signature	Fix
`eta ~ 0` recursion (rumination)	High `M`, high effort, `Q` flat or falling	Nothing is converting to structure. Change the *form* of the practice (writing, artifact, place), which is what `eta` is
`delta ~ 0` fantasy	A plan in which nothing decays	Name the loss mechanism or the model does not apply (Sections 7–10, CS-3)
Negative fixed point run anyway	`eta M* < c*`, `Q` declining, practice continued	Compute `Q_inf` before starting. Cut `C`, raise `eta` or `u`, or stop
Mood as `Q`	`Q` instrument is a daily feeling rating	Mood is not stored order. `Q` is what remains when the mood has gone
Ignoring admissibility	Model produces `Q < 0` and is still trusted	A violated admissibility condition invalidates the step, not the week

CROSS-REFERENCES. [Curzi, 2026o] (full treatment), [Curzi, 2026a] (the leaky integrator), [Curzi, 2026a] (`delta`), [Curzi, 2026b] (`eta`), [Curzi, 2026b] (`C`); F-01 (the input), F-05, F-09.

23.11.13 F-05 · PERSONAL HEADROOM

STATEMENT

```
h_t = eta * ln N − delta * Q_t − C_t                [Q-units/period]

h_t < 0   unsustainable in this model. No geometry saves it.
h_t = 0   maintenance only. Hold, do not build.
h_t > 0   growth available; any requested increment r_t must satisfy r_t <= h_t.

At any fixed point:   h_inf = eta * ( ln N − F(u*) )     THE ARRIVAL INVARIANT
```

READS. *How much capacity for the new you have this period, if your geometry were perfect.*

STATUS. Theorem. the headroom identity [Curzi, 2026b,h] and Corollary 4.2, at internal scope. The arrival invariant is derived in Sections 7–10, §8.4.

DERIVATION. `h` is the growth condition `eta M_t > delta Q_t + C_t` evaluated at the **maximum possible** moment `M = ln N`. It is therefore the margin at perfect geometry. **Consequence that must travel with the entry: `h > 0` is necessary and not sufficient for growth, because the actual step uses `F(u*)`, not `ln N`.** The shortfall — call it the geometry gap — is `eta(ln N − M_t)`, and it is exactly the quantity the arrival invariant names.

THE CONCAVITY CORRECTION, and it applies to every feasibility check made with this entry. The protocol below asks whether a requested increment `r_t` fits inside `h_t`. That test is run against a surplus figure, and the surplus figure must be the **mean of `F(U_t)`, never `F` of the mean `U`.** What accumulates over periods is `sum of eta F(U_t)`, and because `F` is strictly concave [Curzi, 2026], Jensen's inequality gives

```
mean of F(U_t)  <=  F( mean of U_t )        with equality only if U is constant
```

A system whose geometry varies delivers strictly less surplus than one that sits at its own average every period, and testing feasibility on `F(mean U)` therefore approves increments that cannot be delivered. The gap is the concavity penalty. In the corpus's worked office case it is 0.0761 nats — mean `F` of 1.1501 against `F(mean U)` of 1.2262 — **a loss of 6.2 per cent of the surplus purely to variability**, and it was large enough there to turn a feasible-looking plan into a shortfall. Report `U_req` for legibility if you wish; do not test against it. Regularising a variable geometry recovers the penalty for free, which makes it the cheapest surplus available to most subjects.

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	All three terms in Q-units/period. `eta ln N` is Q-units/nat times nats.
FVP-2 SCALE TYPE	**CONDITIONAL PASS**	Inherits F-04's requirement on `Q` and [Curzi, 2026b]'s requirement that `C` be commensurated into Q-units — **which is itself a postulate of the corpus, badged in Appendix E: that costs of different kinds can be put on one scale.**
FVP-3 LIMITING CASES	PASS	`Q -> 0`: `h -> eta ln N - C`, the maximum headroom a person can have. `Q -> (eta lnN - C)/delta`: `h -> 0`, the ceiling on stored order beyond which maintenance consumes everything. `C -> eta ln N`: `h <= 0` for any `Q > 0`.
FVP-4 FIXED POINTS	PASS	`h = 0` at `Q = (eta ln N - C)/delta` — the **maintenance ceiling**, distinct from and above `Q_inf`. And the arrival invariant: at `Q_inf`, `h = eta(ln N - F(u*))`, independent of `delta` and `C`.
FVP-5 IDENTIFIABILITY	**PARTIAL**	Inherits F-04's non-separation of `eta`, `M*`, `c*`. `h` itself is computable once `eta`, `N`, `delta`, `Q`, `C` are in hand; the burden is on `N` and on the commensuration of `C`.

Gate	Verdict	Content
FVP-6 FALSIFICATION	SPECIFIED	Preregister; predict whether a requested increment `r_t` is achieved, from `h_t` computed in advance. **Killed if achieved increments routinely exceed `h_t`** — that outcome would falsify Corollary 4.2's claim that no policy based on interaction geometry alone can guarantee more.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Cross-individual recovery unrun.

PROTOCOL

1. Obtain `N` (F-03), `eta` and `delta` (F-04), `Q_t` (F-04's instrument), `C_t` ([Curzi, 2026b]).
2. Compute `h_t` each period and plot it. **`h` is the gauge; `Q` is the tank.**
3. Before accepting any new commitment, express its expected increment `r_t` in Q-units and check `r_t <= h_t`.
4. At a settled state, compute `h_inf = eta(ln N − F(u*))` and account for the remainder.

WORKED NUMBER

`eta = 0.30`, `N = 7` (`ln N = 1.9459`), `delta = 0.10`, `C = 0.25`:

Q	h	verdict
0.00	0.3338	growth available
1.00	0.2338	growth available
2.00	0.1338	growth available
3.00	0.0338	growth available, barely
3.34	0.0000	**maintenance ceiling**
4.00	-0.0662	unsustainable

`h = 0` at `Q = (0.30 * 1.9459 − 0.25)/0.10 = 3.3377`. At the fixed point of the F-04 example (`u = 0.45`, `Q_inf = 1.4250`):

```
h_inf = eta ( ln N − F )  = 0.30 * (1.9459 − 1.3083) = 0.1913
check from the gauge      = 0.30*1.9459 − 0.10*1.4250 − 0.25 = 0.191273
```

Identical to six decimal places, and it does not contain `delta` or `C`. The reading: at arrival, this subject has 0.1913 Q-units per period of unclaimed capacity, and it is unclaimed for exactly one reason — their framings claim only 67.2% of the ceiling. **Large headroom at a settled state is not a reserve. It is a report that your framings are parallel.**

FAULTS

Fault	Signature	Fix
Headroom as reserve	`h > 0` read as "I have capacity to spare"	`h_inf = eta(ln N − F)`. Raise `u` or `N`; do not "spend the slack"

Fault	Signature	Fix
`h > 0` read as sufficient	Growth expected because `h > 0`, but `Q` falls	`h` is the margin at `M = ln N`. Actual growth needs `eta M_t > delta Q_t + C_t`
Feasibility tested on `F(mean U)`	A plan that clears the gate on paper and falls short every period	Test on **mean `F(U_t)`**. `F` is strictly concave, so `F` of the mean overstates; the gap was 6.2% in the corpus's worked office case
Uncommensurated `C`	Only monetary or only time costs counted	[Curzi, 2026b]'s five ledgers. An omitted cost inflates `h` and licenses a commitment that cannot be met
Negative `h` treated as motivational	"I need to push harder"	Cor 3.4: **no choice of interaction geometry prevents order from decreasing.** Cut `C`, shed `Q`, or raise `eta`

CROSS-REFERENCES. [Curzi, 2026b] (full treatment at external scope), [Curzi, 2026o], [Curzi, 2026b]; F-03, F-04, F-09.

23.11.14 F-06 · THE LAUNCH BUDGET

STATEMENT

```
L = max( 0, M − C )                    [nats, or Q-units once eta is applied]

L      the surplus available for imagining, planning, or attempting elsewhere
M      the moment as lived, in [0, ln N]
C      the cost of the period, commensurated into the same units

L = 0            no surplus to dream with
L_max = ln N − C the ceiling, when the moment is maximal
```

READS. *Escape is dreamt only out of surplus; when the moment costs everything it produces, there is nothing left to imagine elsewhere with.*

STATUS. Derived formula of the corpus [Curzi, 2026o], formalising the STILLPOINT node of the foundations series [Curzi, 2026] departure-lounge doctrine. Derived, with one badge.

DERIVATION. The storage law's one-step content is `eta M_t − C_t`: what a period contributes before the leak takes its share. `L` is that quantity's positive part, before conversion — the *unallocated* portion of a moment, the part not consumed by the cost of having it.

Postulate L1 (the launch condition). Imaginative and exploratory activity draws on the same budget as maintenance, and draws second. **Badged: an assumption about the ordering of claims on a period's surplus, not a result.** It is what makes `L` a max-with-zero rather than a quantity that can go negative and be carried, and it is consistent with the departure-lounge doctrine — *in those minutes there was no surplus to dream with* is `L = 0` — but the ordering is unmeasured.

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`M` in nats; `C` must be converted into nats (divide the Q-unit cost by `eta`) or `M` converted into Q-units (multiply by `eta`). **The conversion must be stated; mixing the two is the commonest arithmetic error in this entry.**
FVP-2 SCALE TYPE	**DEMOTE** (Standing Verdict 1)	Subtraction needs interval scale on a common unit; the commensuration of `C` is again a badged postulate (Appendix E). **`M` additionally inherits F-01's demotion, so `L` is reportable as a threshold condition — is the moment above the cost line or not — and not as a magnitude.** The sign of `M - C` is the finding; its size is not, until [Curzi, 2026t] Thm 54.3's instrument exists.
FVP-3 LIMITING CASES	PASS	`M <= C`: `L = 0` — the departure lounge with nothing to depart with. `M -> ln N`: `L -> ln N - C`, the ceiling. `C -> 0`: `L -> M`; the whole moment is available. `C -> ln N`: `L = 0` for every possible moment — **a cost equal to the ceiling makes the person structurally incapable of imagining elsewhere, whatever they experience.**
FVP-4 FIXED POINTS	PASS	The boundary `M = C` is the **no-launch surface**. `L` is continuous but **not differentiable** there — a genuine kink, and the non-smoothness is the phenomenon: the transition from "some budget" to "none" is abrupt, not gradual, which is why people describe it as a wall rather than a slope.
FVP-5 IDENTIFIABILITY	**PARTIAL**	`L` is identified from the difference; `M` and `C` separately are not, if only their difference is observed. Since the intervention differs completely (raise `M` vs cut `C`), **the separation is required before acting, not before measuring.**

Gate	Verdict	Content
FVP-6 FALSIFICATION	SPECIFIED	Preregister `M` and `C` measurement; predict, from a frozen model, which periods produce exploratory output (plans, applications, drafts, attempts). Beat the four baselines. **Killed if exploratory output is uncorrelated with `L` after controlling for `M` alone** — that would say the cost side carries nothing and the doctrine is wrong.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Unrun.

PROTOCOL

1. Measure `M_t` per F-01, normalized to `[0, ln N]`.
2. Measure `C_t` per [Curzi, 2026b]'s ledger and convert to nats by dividing by `eta`.
3. `L_t = max(0, M_t − C_t)`. Log it every period.
4. Report **the count of `L = 0` periods**, not just the mean. The zeros are the finding.

WORKED NUMBER

`M`	`C`	`L`	state
0.20	0.60	0.000	no surplus to dream with
0.45	0.45	0.000	exactly at the surface
0.60	0.45	0.150	launch available
1.90	0.45	1.450	launch available
1.90	1.95	0.000	high moment, higher cost — **the busy trap**

The last row is the entry's most important case. **A person can have an extremely rich moment and zero launch budget, and from the outside they look like someone who has everything and does nothing with it.** The arithmetic says the opposite: they have nothing left, because `C` exceeded `M`.

Over the F-04 example's ten periods with `C = 0.25`:

```
sum L = 9.75      mean L = 0.975      periods with L = 0: 1
L_max = ln N − C  = 1.9459 − 0.25 = 1.6959
```

FAULTS

Fault	Signature	Fix
Order-of-operations reversal	"Get out of your comfort zone" prescribed to someone at `L = 0`	The dream is manufactured *from* surplus. Raise `M` or cut `C` first; the exhortation has no budget to act on ([Curzi, 2026p])
The busy trap	High `M`, high `C`, `L = 0`, and self-blame for lack of ambition	Cut `C`. The moment is fine; the ledger is not

Fault	Signature	Fix
Unit mixing	`M` in nats minus `C` in Q-units	Convert. Divide the Q-unit cost by `eta`
Mean without zeros	Reporting mean `L` across a block containing many zeros	Report the zero count. A mean of 0.5 from alternating 1.0 and 0.0 is a different life from a steady 0.5

CROSS-REFERENCES. [Curzi, 2026p] (full treatment), [Curzi, 2026b], [Curzi, 2026p]; F-01, F-05.

23.11.15 F-07 · THE GRADIENT

STATEMENT

```
g = − grad C                          the desire line, a vector field

C(x)   the local cost of being at / doing configuration x  [Q-units/period]
x      a configuration coordinate: actions-to-begin, distance, latency, friction
g      points in the direction of cheapest motion

Observable proxy, under Postulate G1:
    n(x) ~ exp( −C(x) / theta )   ->   ln n(x) = const − C(x)/theta
    fitted slope b = −(1/theta) dC/dx
```

READS. *Water does not negotiate; read the worn path rather than arguing with it.*

STATUS. Derived formula of the corpus — the GRADIENT node [Curzi, 2026] formalised. Derived, with one badge.

DERIVATION. The storage law's cost term `C_t` is a scalar per period, but the cost of a period depends on what was done in it; treat `C` as a function over a configuration space and the direction of cheapest motion is `−grad C`. That much is definitional. **The empirical content requires a link between cost and behaviour, and that link is a postulate:**

Postulate G1 (the choice rule). The frequency with which a configuration is entered is log-linear in its local cost: `n(x) ~ exp(−C(x)/theta)`, `theta > 0` a choice temperature. **Badged: the standard log-linear choice form, adopted as a modelling assumption — not a result of this corpus and not a claim about any measured population.** It puts behaviour in the same logarithmic register as everything else in the book, at the price of a free `theta` (see FVP-5).

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`C` in Q-units/period; `x` in whatever the coordinate is (actions, metres, seconds); `grad C` in Q-units per period per unit of `x`. Under G1 the fitted `b` is in `(1/theta)` times those units — **so the fitted slope is a gradient in units of the choice temperature, not in Q-units.**
FVP-2 SCALE TYPE	**CONDITIONAL PASS**	`C` must be ratio-scale ([Curzi, 2026b]). `x` must be at least interval-scale; counts of actions are ratio-scale and are therefore the preferred coordinate. `n` is a count — ratio-scale — which is what makes the log-linear fit legitimate.
FVP-3 LIMITING CASES	PASS	`grad C -> 0`: no desire line; motion is undirected — the flat the foundations series calls the stillpoint [Curzi, 2026]. `|grad C| -> inf`: an infinitely steep wall; the configuration is never entered. `theta -> 0`: choice becomes deterministic, always the cheapest option. `theta -> inf`: choice becomes uniform, cost carries nothing.
FVP-4 FIXED POINTS	PASS	`grad C = 0` at local minima: **the comfort zone, defined without sentiment as a stationary point of the cost field.** The STILLPOINT node of the foundations series [Curzi, 2026] — *you are already there* — is the statement that the present configuration is at or near such a point by construction.

Gate	Verdict	Content
FVP-5 IDENTIFIABILITY	**PARTIAL — and instructively so**	From behaviour alone, only `C/theta` is identified, never `C` and `theta` separately. **But the *gradient* is better identified than the *level*: adding any constant to `C` changes no behaviour and no fitted slope.** So this entry can report reliable differences and cannot report absolute costs. Say so whenever a fitted `b` is quoted.
FVP-6 FALSIFICATION	SPECIFIED	Preregister the coordinate and the count window. Fit `b` on a training block, **freeze**, then **intervene**: remove `k` actions from one configuration and predict the new count as `n * exp(-b k)`. Beat persistence, direct-input, matched AR(1), shuffled-input. **Killed if the predicted response to a manipulated coordinate is not observed** — this is the one entry in the formulary with a cheap, direct, interventional test, and it should be run first.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Does `b` recover across people and across domains? Unknown.

PROTOCOL

1. Choose a coordinate that is a **count**: number of discrete actions required to begin.
2. For `>= 5` activities, record the action count `k_i` and the initiation count `n_i` over a fixed window.
3. Fit `ln n = a + b k` by least squares. Report `b`, `R^2`, and the fitted values.
4. Compute the **halving distance**: `ln(0.5)/b` actions — the number of added actions that halves initiation.
5. **Intervene and re-measure.** Step 5 is not optional; it is FVP-6.

WORKED NUMBER

Illustrative log over one week — six activities, actions-to-begin against initiation count:

actions `k`	observed `n`	fitted `n`
1	26	25.05
2	17	16.62
3	11	11.03
6	3	3.22
9	1	0.94

```
fit:  ln n = 3.6308 - 0.4100 k        R^2 = 0.9982
gradient magnitude |b| = 0.4100 nats of initiation per added action
halving distance = ln(0.5)/(-0.4100) = 1.69 actions
removing 2 actions multiplies initiation by exp(2*0.4100) = 2.271
```

Every 1.69 actions added to the front of a practice halves how often it happens. That is the whole of the GRADIENT node [Curzi, 2026] in one number, and it is the number that makes threshold design an engineering discipline rather than an aesthetic one.

The honest caveat, from FVP-5: **0.4100 is a gradient in units of the choice temperature, and the absolute cost of an action is not recoverable from this fit.** What is recoverable — and what the intervention test checks — is that removing two actions should multiply initiation by about 2.27.

FAULTS

Fault	Signature	Fix
Arguing with the gradient	Exhortation applied to a configuration nobody enters	Regrade. Remove actions. The path is data, not an opinion to be corrected
Level claims from slope data	“This costs me 0.41 Q-units”	FVP-5. Only differences are identified. Report ratios of frequencies, never absolute costs
Ordinal coordinate	x is “difficulty, 1-5”	Use a count. Actions, seconds, metres
No intervention	Fit reported, never tested	FVP-6. The manipulation is cheap and it is the only real test this entry has
Confounded coordinate	High-k activities also happen to be the disliked ones	Manipulate k on **one** activity and hold the activity fixed; that is what makes it an experiment

CROSS-REFERENCES. Curzi [2026p] (full treatment); the corpus’s threshold and intervention material; F-05, F-06.

23.11.16 F-08 · THE RETURN DIFFERENCE

STATEMENT

```
rho_t     = ln( M'_t / M_t )              signed return, in nats
Delta_R,t = | rho_t |                     magnitude, in nats

M_t   the moment sent       M'_t  the moment returned

rho > 0  OVER-RETURN   |  rho = 0  NULL RETURN  |  rho < 0  UNDER-RETURN

deadband:     Delta_R_min = ln(1 + c)         c = the subject's Weber fraction
correction:   rho_{t+1} = (1 - lambda) rho_t  lambda = correction gain
stability:    convergent iff 0 < lambda < 2
```

READS. *The world's answer is never the voice you sent, and the difference is the only data in the exchange.*

STATUS. Derived formula of the corpus [Curzi, 2026o], the ECHO node of the foundations series [Curzi, 2026]. Derived, with one badge (Postulate R1, the returnable moment — Sections 3–6, §3.2).

DERIVATION. Sections 3–6 in full. In one line: the corpus works in log-ratios everywhere, so the only difference measure that composes with the rest of the formulary is the log-ratio, and it is forced rather than chosen.

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`M'/M` dimensionless; `rho` in nats, convertible by the §11.4 keys.
FVP-2 SCALE TYPE	**DEMOTE — binding** (Standing Verdict 1)	Ratio scale with a true zero, or the entry reads as a **sign-only ordering**. `M` and `M'` are moments and inherit F-01's demotion, so **that is where the entry stands today: over-return, null return and under-return are findings; the number of nats is not.** A `rho` computed from Likert data was always void rather than imprecise; a `rho` computed from an uncalibrated magnitude estimate is merely unreportable. [Curzi, 2026t] Thm 54.3 supplies the missing scale, at ~123 items per term for half a logit.
FVP-3 LIMITING CASES	PASS	`M' -> M`: `rho -> 0`, no information (contact or closure). `M' -> 0`: `rho -> -inf`, total non-return — being ignored, and the divergence is correct because there is no finite penalty scale for it. `M -> 0`: undefined, correctly — an echo of nothing is not an echo.
FVP-4 FIXED POINTS	PASS	`rho = 0`, unique and attracting under the correction map for `0 < lambda < 2`. Meaning: agreement between your construction and the world's answer — **which is either contact or closure, and only the scatter test distinguishes them.**

Gate	Verdict	Content
FVP-5 IDENTIFIABILITY	**PARTIAL**	`rho` recoverable from paired estimates. The **decomposition** into frame error and world response is **not identified from `rho` alone**; separating them requires an independent estimate of `M'` or a deliberate manipulation of `D`.
FVP-6 FALSIFICATION	SPECIFIED	Fit `lambda` on training, freeze, predict held-out `rho`. Beat persistence, direct-input, matched AR(1), and the **shuffled-input control** (same `M` series, permuted `M'`). Killed by failure against the shuffled control.
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Do `lambda` values recover across people? Unknown; Sections 3–6 §4.11 states the design.

PROTOCOL

1. Fix the modulus; record `M` **before** the return arrives (the single most violated rule).
2. Record `M'` from `>= 2` sources asked **independently**, timestamped to the exchange they answer, not to their arrival.
3. Estimate `sigma_rho` from `>= 8` equivalent exchanges before interpreting anything.
4. Declare the deadband `ln(1+c)` and refuse to read below it.
5. Compute `rho`, `Delta_R`, and per block: `R_net = sum rho`, `R_gross = sum |rho|`, `E_R = |R_net| / R_gross`.
6. After four consecutive returns inside the deadband, run the flattery test.
7. **An echo source may decline; a decline is missing data, never a null return.**

WORKED NUMBER

M	M'	ratio	rho	Delta_R	bits
60	60	1.0000	0.0000	0.0000	0.0000
60	63	1.0500	+0.0488	0.0488	0.0704
60	45	0.7500	-0.2877	0.2877	0.4150
60	90	1.5000	+0.4055	0.4055	0.5850
60	20	0.3333	-1.0986	1.0986	1.5850

Deadband, at three Weber fractions: `c = 0.05 -> 0.0488` nats; `c = 0.10 -> 0.0953`; `c = 0.20 -> 0.1823`. **An subject with `c = 0.20` cannot distinguish the second row of that table from silence.**

Detection cost, `n >= (1.96 sigma/Delta)^2`: at `sigma_rho = 0.30`, resolving a `Delta_R` of 0.05 takes **139 echoes**; resolving 0.20 takes **9**.

FAULTS

Fault	Signature	Fix
Zero read as success	`Delta_R ~ 0` celebrated	Run the scatter test. Converged contact varies; closure does not
The flattering mirror	Four-plus returns inside the deadband, no sign changes, all sources agree	Add a genuinely independent source; compute `n_eff`
The devouring echo	`|rho|` growing, signs alternating with rising amplitude	`lambda >= 2`. Reduce below 1 and hold
Post-hoc `M`	`rho` implausibly near zero across a whole log	Record `M` before the return. Compare a pre-recorded block against a reconstructed one
Worth conversion	Negative `rho` narrated as a statement about the self	FVP-5: the decomposition is not identified. The instrument does not license the claim
Correcting faster than the echo returns	Oscillation appearing only when cadence increased	Correction cadence <= echo cadence

CROSS-REFERENCES. Sections 3–6 (full treatment), Curzi [2026d] on pseudo-diversity, Curzi [2026i] on the preregistered protocol; F-01, F-02, F-09.

23.11.17 F-09 · THE CONTRACTION AND ARRIVAL

STATEMENT

```
| Q_t − Q_inf | = (1 − delta)^t | Q_0 − Q_inf |          THE ERROR LAW

Q_t = Q_inf + (1 − delta)^t ( Q_0 − Q_inf )             the whole trajectory

Q_inf = ( eta ∗ M∗ − c∗ ) / delta                       the steady−state theorem
tau   = −1 / ln(1 − delta)                             periods per factor of e
t(eps) = ln(eps) / ln(1 − delta)                       periods to close to a fraction eps

a−posteriori distance:  | Q_t − Q_inf | = ((1−delta)/delta) ∗ (last step into Q_t)
```

READS. *Where you are is where you are going, plus what is left of where you started — and what is left of where you started is halving on a fixed schedule.*

STATUS. Theorem. the steady-state theorem [Curzi, 2026a] plus the error law derived from Theorem 3.1; Banach's fixed point theorem supplies existence, uniqueness, and robustness. **The FIXEDPOINT node [Curzi, 2026] FIXEDPOINT is Theorem 3.5, and Sections 7–10 proves the correspondence rather than asserting it.**

DERIVATION. Sections 7–10, §7.2–§7.5. `T(Q) = (1−delta)Q + b` is a contraction on the complete space `(R, |.|)` with constant exactly `1−delta` whenever `delta > 0`; Banach gives the unique fixed point, and subtracting the fixed-point equation from the update gives the error law by induction. **The leak is what makes arrival possible: at `delta = 0` the map is an isometry, there is no fixed point, and `Q_t = Q_0 + tb` grows without bound.**

FVP

Gate	Verdict	Content
FVP-1 DIMENSIONS	PASS	`Q` in Q-units; `(1-delta)^t` dimensionless; `tau` in periods. **The period must be declared or none of it means anything.**
FVP-2 SCALE TYPE	**CONDITIONAL PASS**	Interval scale for the differences, ratio scale for percentage readings. Inherits F-04's badged postulate that `Q` is a scalar.
FVP-3 LIMITING CASES	PASS	`t -> inf`: gap `-> 0`. `delta -> 0+`: no convergence, `Q_inf -> inf` unless `b = 0` — the system that never forgets and never settles. `delta -> 1`: arrival in one step; the state is only the present drive. `Q_0 = Q_inf`: gap zero for all `t` — you were already there.
FVP-4 FIXED POINTS	PASS	`Q_inf = (eta M* - c*)/delta`, unique by Banach, **containing no `Q_0`** — which is why arrival is recognition rather than victory. The error map `e -> (1-delta)e` has the unique fixed point `e = 0`.
FVP-5 IDENTIFIABILITY	**PARTIAL**	`delta` identified from successive gaps; the composite `eta M* - c*` from `delta Q_inf`; the three components **not separately identified** without independent measurement of two of them.
FVP-6 FALSIFICATION	SPECIFIED	Fit `delta` on training, freeze, forecast held-out `Q`. Beat persistence, direct-input, matched AR(1), shuffled-input. **Killed if the frozen geometric-decay model does not beat the shuffled-input control**, since that control destroys the ordering of inputs, which is the entire content of "it accumulates".
FVP-7 REPRODUCIBILITY	**UNTESTED — declared**	Does `delta` recover across people, investigators, domains? Unknown; Sections 7–10 §8.10 states the design.

PROTOCOL

1. Declare the period.
2. Compute `Q_inf` **before** pursuing it, from `eta`, `M*`, `c*`, `delta`.

3. Check the sign. `eta M* < c*` gives a negative fixed point: the loop converges toward destruction, and this must be reported before the practice is run, not discovered during it.
4. Declare the arrival criterion `c` and compute `t_arrive = ln(c)/ln(1−delta)` from a standing start.
5. Track distance from arrival each period as `((1−delta)/delta)` times the last step.
6. At arrival, compute `h_inf = eta(ln N − F(u*))` (F-05) and account for it.

WORKED NUMBER

`delta = 0.10`, `Q_inf = 4.0`, `Q_0 = 0.5`:

t	gap	Q_t
0	3.50000	0.50000
22	0.34467	3.65533
29	0.16485	3.83515
44	0.03394	3.96606
66	0.00334	3.99666

```
to 5%: 29 periods      to 1%: 44 periods      to 0.1%: 66 periods
tau = 9.49 periods     half-life = 6.58 periods
```

Each halving costs 6.58 periods and each factor of ten costs 21.85 periods, and both are constant for the entire trajectory. Closing the gap from 90% to 99% costs 22 periods; closing it from 99% to 99.9% also costs 22 periods; closing the first half costs 7. **That is the arithmetic behind every practice that felt fast at the start and interminable at the end, and nothing slowed down.**

The full convergence table across `delta` is Sections 7–10, §8.1; the four landmarks are `delta = 0.05 -> 59 periods to 5%`, `0.10 -> 29`, `0.20 -> 14`, `0.30 -> 9`.

FAULTS

Fault	Signature	Fix
The unarrived asymptote	“Almost there” for many periods, read as failure	There is no arrival event. Declare `c` in advance
Effort substitution	More effort, no change in `eta`, `u`, `C`, `delta`; same path	Effort acts on the transient. Name which parameter the effort moves; if none, stop
Cadence confusion	Cadence increased, calendar unchanged	`tau` depends on per-period `delta`. Cadence acts through `eta` and `C`
Counterfeit arrival	`Q` rising with no change in your own four parameters	The extra drive is a transfer from someone. Run the audit on every party (Sections 7–10, §8.7.4)
Stall named arrival	Flat `Q`, no `Q_inf` on record, no review date	The arithmetic cannot adjudicate; the documentation can
Low-`delta` noise	Very stable practice, large swings traced to small perturbations months earlier	`1/delta` amplification. Accept it or shed deliberately

CROSS-REFERENCES. Sections 7–10 (full treatment and the proof), [Curzi, 2026a], [Curzi, 2026a], [Curzi, 2026h]; F-04, F-05, Appendix F.

23.12 12 Requirements for formulary use

How to run the formulary. **These are ordering and hygiene requirements; the per-entry protocols sit inside them.**

PF-1 — Run in dependency order. `F–03 -> F–01 -> F–02 -> F–08 -> F–04 -> F–05 -> F–06 -> F–07 -> F–09`. `N` and `u` first because they set the ceiling everything else lives under; the contraction last because it needs every other parameter.

PF-2 — One modulus, one period, one block. All magnitude estimates in a block share a modulus; all rates share a period; both are declared in writing before data collection and are not changed mid-block.

PF-3 — Zero-input window first. Fit `delta` during a deliberate low-input period before fitting anything else. It is the only clean identification of `delta` available at internal scope, and everything downstream depends on it.

PF-4 — Pre-record, never reconstruct. `M` before the return (F-08); `Q_inf` before the pursuit (F-09); the arrival criterion before arrival. **A formulary run on reconstructed inputs produces internally consistent numbers that are jointly wrong, which is the worst available outcome.**

PF-5 — Report intervals, not points. Every estimate carries an `n` and a spread. A point estimate with no `n` is not a measurement.

PF-6 — Compare each quantity to its deadband. `Delta` (F-02) and `Delta_R` (F-08) against `ln(1+c)`; `dQ` (F-04) against the `Q` instrument's resolution. **Below the deadband, report "unresolved", never "zero".**

PF-7 — Carry the identifiability caveats into every conclusion. The non-identifications are: `D` versus `P` (F-01); `N` versus `u` (F-03); `eta`, `M*`, `c*` (F-04, F-05, F-09); frame error versus world response (F-08); `C` level versus `theta` (F-07). **A conclusion requiring one of these separations must name the second measurement that supplied it.**

PF-8 — FVP-7 stays visible. No entry may be reported as validated. The reproducibility gate is open on all nine and the report says so.

PF-9 — Run F-07's intervention test first if you are going to run only one test. It is the cheapest and most direct falsification in the formulary: remove `k` actions and check whether initiation multiplies by `exp(-bk)`.

PF-10 — Other people's ledgers are asked for, never assumed. Where an echo source is a person (F-08) or a cost is carried by a person (F-05, F-09's transfer audit), consent is obtained, a decline path exists, and a decline is recorded as **unknown**, never as satisfied and never as a zero. *(The declination refusal, binding on every entry [Curzi, 2026].)*

PF-11 — Demote in print when a gate fails. A CONDITIONAL PASS whose condition is not met becomes a heuristic ordering in the report, in the same typeface as everything else, with the demotion stated. **A demoted formula is still useful; a formula that hides its demotion is not.**

PF-12 — Freeze before claiming. Any claim that the formulary works follows FVP-6 in full, including the shuffled-input control.

23.13 13 Failure modes of formulary use

Formulary-wide faults — the ones that come from using the collection rather than from any single entry.

Fault	Diagnostic signature	Fix
Consistent-and-wrong	Every number agrees with every other; the picture is coherent and the practice does not behave as predicted	`M` is mismeasured. Six entries are fed by F-01, so a single upstream error propagates coherently. Re-run F-01 with an independent `E` estimate
Gate laundering	The FVP tables are present but every verdict is PASS	Sixty-three gates and no failures is not rigour; it is a copied table. Check FVP-2 and FVP-7 first, since those are the two that genuinely bind
Precision inflation	Four decimal places on a quantity estimated from three observations	PF-5. Report `n` and spread. The formulary's arithmetic is exact; its inputs are not
Formula shopping	Running entries until one produces a flattering number	Run in dependency order (PF-1) and report all nine, including the ones that did not help
Parameter drift	`N`, `delta`, `eta` quietly revised mid-block until the model fits	PF-2 and PF-4. Freeze, then measure. A parameter revised to fit is a parameter that has stopped measuring
Scope leakage	Internal-scope entries applied to a team without re-derivation; the aggregation problem is general, and a single parameter written for a group is not the group's members' parameter [Kirman, 1992, Granovetter, 1978]	the thesis of Series B is that the parameters change between scopes, not that the numbers transfer. [Curzi, 2026s] is the interpersonal formulary; use it
The single-instance conclusion	A whole regime redesigned from one bad week	Every entry's protocol specifies a minimum `n`. One instance is a prompt to measure, not a finding

Fault	Diagnostic signature	Fix
Cross-person comparison	"My `delta` is worse than hers"	FVP-7 is open on all nine. Within-person, within-domain only, until the reproducibility study exists
Ordinal contamination	Ratios and logs computed from ratings scales anywhere in the chain	FVP-2. Trace the chain: an ordinal input at F-01 voids F-02, F-06, F-08 and corrupts F-04, F-05, F-09
Moral arithmetic	The formulary used to argue that someone's `Q` matters less	The declination refusal [Curzi, 2026]. The formulary measures; it does not rank persons, and any use that does has left the corpus

---

23.14 14 Results

Three worked results are reported. Section 14.1 runs the echo instrument end to end on a twelve-week log. Section 14.2 runs the arrival arithmetic on a practice rebuilt after a move, computing the destination before departure and then auditing it. Section 14.3 runs all nine formulary entries over one twelve-week block on a single subject. Section 14.4 collects the validation verdicts.

Every number below is computed from stated parameters under the models developed above, on series that were constructed rather than observed. The source material labels these *illustrative constructions* and that label is retained. No dataset was collected and no person was studied.

23.14.1 14.1 The echo log, read end to end

A twelve-week echo log. The data below are **illustrative, not empirical**: they are a constructed log used to demonstrate the computations, and no claim is made that they describe any real person. Every number derived from them was computed, not estimated by eye.

Setting. An subject produces one unit of work per week in a domain where a return arrives within the same week. Before releasing each week's work, they record `M` — a magnitude estimate of the moment they believe they are sending, on a scale where a fixed reference exchange from the preceding quarter was assigned 50. After the return arrives, they record `M'` on the same scale, from three sources asked independently and before seeing one another (ES-5), averaged. Their declared deadband is `Delta_R_min = 0.05` and their pre-estimated noise from an eight-exchange calibration block is `sigma_rho = 0.09`.

23.14.2 The log

Week	M sent	M' returned	ratio	rho	Delta_R	cumulative mean rho
1	40	22	0.5500	-0.5978	0.5978	-0.5978
2	45	26	0.5778	-0.5486	0.5486	-0.5732
3	50	34	0.6800	-0.3857	0.3857	-0.5107
4	48	35	0.7292	-0.3159	0.3159	-0.4620
5	55	44	0.8000	-0.2231	0.2231	-0.4142
6	52	45	0.8654	-0.1446	0.1446	-0.3693
7	60	55	0.9167	-0.0870	0.0870	-0.3290
8	58	57	0.9828	-0.0174	0.0174	-0.2900
9	55	56	1.0182	+0.0180	0.0180	-0.2558
10	60	63	1.0500	+0.0488	0.0488	-0.2253
11	62	60	0.9677	-0.0328	0.0328	-0.2078
12	60	61	1.0167	+0.0165	0.0165	-0.1891

23.14.3 Block statistics

```
all 12 weeks     mean rho = −0.1891    sd = 0.2257
weeks 1–6        mean rho = −0.3693    sd = 0.1785
weeks 7–12       mean rho = −0.0090    sd = 0.0478

mean Delta_R, weeks 1–6   = 0.3693
mean Delta_R, weeks 7–12  = 0.0368
```

With six observations per half-block, the standard errors and 95% intervals are:

```
weeks 1–6    SE = 0.0729    95% CI on mean rho = [−0.5121, −0.2264]
weeks 7–12   SE = 0.0195    95% CI on mean rho = [−0.0473, +0.0293]
```

The first half-block's interval excludes zero decisively; the second half-block's interval contains zero. The instrument has detected a real, large, consistently signed return difference in weeks 1–6 and cannot detect one in weeks 7–12.

23.14.4 Reading the sign

Every return in weeks 1–6 is negative and none is inside the deadband. This is sustained **under-return**: the world answered systematically smaller than the subject constructed. Per §4.1, the leading candidate is over-framing — `r = D/P > 1`. The size of the implied over-frame is

```
implied over-frame factor = exp( −mean rho over weeks 1–6 )
                          = exp( 0.3693 )
                          = 1.4467
```

The subject was constructing the moment about 1.45 times larger than the world was returning it. That is a specific, checkable, correctable claim, and it is the sort of claim the additive difference measure could not have produced — `M' − M` in weeks 1–6 runs -18, -19, -16, -13, -11, -7, which looks like a shrinking problem, while the ratio shows a shrinking *frame error*, which is a different fact.

23.14.5 Estimating the correction gain from the data

The subject was correcting — not by a written rule, but the log allows the effective gain to be recovered. Per §4.6, if `rho_{t+1} = (1−lambda) rho_t`, then `ln|rho_t|` falls linearly in `t`. Regressing `ln|rho|` on week for weeks 1–8 (the run before the residuals reach the deadband):

```
slope of ln|rho| on t   = −0.4436
per-step factor  g      = exp(−0.4436) = 0.6418
implied gain     lambda = 1 − g        = 0.3582
R^2                     = 0.8587
```

The subject's effective correction gain was 0.36 — comfortably inside the stable, monotone band, which is why the log shows steady shrinkage with no sign reversals until the residuals reach the noise floor. Derived quantities:

```
half-life of the distortion = ln(0.5)/ln(0.6418) = 1.56 weeks
weeks to bring |rho| below the deadband 0.05     = 7
```

The fitted trajectory and the observed one agree to the accuracy the `R^2` implies:

```
fitted |rho| at weeks 1..12:
0.913, 0.586, 0.376, 0.241, 0.155, 0.099, 0.064, 0.041, 0.026, 0.017, 0.011, 0.007
```

Observed week 7 is 0.0870 against a fitted 0.0638; observed week 8 is 0.0174 against 0.0409. From week 9 onward the fitted values are below the deadband and the observed values are indistinguishable from noise, exactly as §4.3 predicts. **After week 8 the model is no longer being tested by the data, because the data cannot resolve it.**

23.14.6 Net, gross, efficiency

```
R_net   = sum rho   = −2.2696
R_gross = sum |rho| =  2.4362
E_R     = 0.9316
```

`E_R = 0.93` is very high. Per §4.9 this is the signature of a **live, consistently directed error** rather than thrash — the world was saying one thing, repeatedly, and the subject moved toward it steadily rather than oscillating around it. Compare the counterfactual: had the same gross activity been produced by a gain above 1, the signs would have alternated and `E_R` would have collapsed toward zero.

23.14.7 The flattery check

Weeks 9–12 give four consecutive returns inside the declared deadband of 0.05 — `+0.0180`, `+0.0488`, `−0.0328`, `+0.0165`. ES-10 fires the alarm. Running the test of §4.5 with `sigma = 0.09` and band `b = 0.05`, the probability of one echo landing inside the band under a live-but-null channel is 0.4215, so four in a row has probability 0.0316 — below 5%.

The alarm fires, and it should, and the alarm is a prompt to test rather than a verdict. Here the test is the scatter criterion: the four residuals are `+, +, −, +` with a standard deviation of 0.0289 against a noise estimate of 0.09. **They vary and they change sign.** That is the signature of converged contact, not of a mirror — a mirror returns the sent moment and produces residuals that do not alternate. The subject records the alarm, records the scatter test that resolves it, and continues. Had the four residuals been `+0.004, +0.006, +0.003, +0.005`, the correct conclusion would have been the opposite one, and the correct action would have been to add a source who had not been asked before.

23.14.8 What the subject does next

Three decisions follow from the block, and each is licensed by a specific number.

1. Stop correcting the frame. Weeks 7–12 have a mean `rho` of -0.0090 with a 95% interval containing zero and a mean `|rho|` of 0.0368, below the declared deadband of 0.05. **There is no resolvable error left to correct at this modulus, and continuing to apply a gain of 0.36 to unresolvable residuals is how a stable subject becomes an oscillating one.** ES-8's instruction is to hold the gain for the block; the block is over, and the next block's gain should be set from the new signal-to-noise, which is now below 1.

2. Increase the sent moment or increase the resolution — but not both blindly. Per §4.4, with `sigma_rho = 0.09` the subject can resolve a `Delta_R` of 0.10 with four echoes and a `Delta_R` of 0.05 with sixteen. If the next question is finer than 0.10, the instrument needs either more echoes per finding or a finer modulus. If it is coarser, the instrument is already adequate and the effort should go elsewhere.

3. Do not conclude anything about the work. ES-12. The block establishes that the subject's construction of the moment moved from about 1.45x the world's return to about 1.0x. It establishes nothing about whether the work is good, and FVP-5 states formally why: the decomposition of `rho` into frame error and world response is not identified from `rho` alone. **The subject has calibrated an instrument. That is the entire finding, and it is worth having.**

23.14.9 The one number to carry out of the section

```
weeks 1–6   mean Delta_R = 0.3693 nats = 0.5328 bits = 1.066 octaves
weeks 7–12  mean Delta_R = 0.0368 nats = 0.0531 bits = 0.106 octaves
```

Twelve weeks of exchange bought about half a bit of correction, and half a bit is a real, respectable, hard-won yield. Anyone expecting more from twelve exchanges has mistaken the register they are working in. The world answers in ratios, the subject hears in logarithms, and the difference between the voice and the answer is the only thing either of them learns.

23.14.10 14.2 Arrival, computed before departure and audited afterwards

A practice rebuilt after a move. As in Sections 3–6, the data are **illustrative** — a constructed case used to demonstrate the computation. No claim is made that these parameters describe any real person.

Setting. An subject restarting a daily practice from zero after a relocation that reset their stored order. Period = one week. They measure, per the protocols of [Curzi, 2026n,o], 40 and 5:

```
N     = 9        distinguishable framings available in the domain
u*    = 0.55     measured non-parallelism among their own framings
eta   = 0.35     conversion of moment into durable structure
delta = 0.12     weekly loss without maintenance
c*    = 0.22     weekly maintenance cost, from the cost ledger
Q_0   = 0        the move reset the store
```

23.14.11 Step 1 — the ceiling and the drive

```
ln N       = ln 9              = 2.1972    the hard ceiling on surplus per interaction
N_eff      = 1 + (9-1)(0.55)   = 5.400     an audited headcount of framings
F(u*)      = ln(5.400)         = 1.6864    the surplus this geometry generates
claimed    = 1.6864 / 2.1972   = 76.8%     of the ceiling
drive b    = 0.35*1.6864 - 0.22 = 0.3702   Q-units per week
```

`b > 0`, so CS-4 passes: the regime is sustainable and the fixed point is positive.

23.14.12 Step 2 — the destination, computed before departure (CS-2)

```
Q_inf = ( 0.35 * 1.6864 - 0.22 ) / 0.12 = 3.0853 Q-units

upper bound check (Thm 3.3, Q_0 = 0):  eta*lnN/delta = 6.4086
3.0853 <= 6.4086.  Consistent.
```

The room is furnished at 3.0853 Q-units, and this was known in week zero.

23.14.13 Step 3 — the schedule

```
tau        = -1/ln(0.88)          =  7.82 weeks
half-life  = ln(0.5)/ln(0.88)     =  5.42 weeks
to 5%      = ln(0.05)/ln(0.88)    = 24 weeks
to 2%                             = 31 weeks
to 1%                             = 37 weeks
```

Declaring `c = 0.05` per CS-5: **arrival is week 24, and this is written down in week zero.**

23.14.14 Step 4 — the trajectory

t (wk)	Q_t	gap	gap / Q_inf	h_t	step dQ
0	0.0000	3.0853	1.0000	0.5490	0.3702
1	0.3702	2.7151	0.8800	0.5046	0.3258
2	0.6961	2.3893	0.7744	0.4655	0.2867
4	1.2351	1.8503	0.5997	0.4008	0.2220
6	1.6525	1.4328	0.4644	0.3507	0.1719
8	1.9757	1.1096	0.3596	0.3119	0.1332
10	2.2261	0.8593	0.2785	0.2819	0.1031
12	2.4199	0.6654	0.2157	0.2586	0.0799
14	2.5700	0.5153	0.1670	0.2406	0.0618
16	2.6863	0.3990	0.1293	0.2267	0.0479
20	2.8460	0.2393	0.0776	0.2075	0.0287
25	2.9590	0.1263	0.0409	0.1939	0.0152
30	3.0187	0.0666	0.0216	0.1868	0.0080

Closed-form check at `t = 30`: `Q_inf + (1-delta)^30 (0 - Q_inf) = 3.018683`. The iteration and the error law agree to six figures, as §7.5 requires.

23.14.15 Step 5 — what the subject experiences, week by week

Weeks 1–5. Steps of 0.37, 0.33, 0.29, 0.25, 0.22, closing 47% of the gap. This is the part everyone calls “fast at the beginning” — it is not fast; it is the same 12% per week as every other week, applied to a large number.

Weeks 6–16. Steps from 0.17 down to 0.048. The gap goes from 46% to 13%. This is the long middle, and per §8.1 it costs what it costs: closing the gap from 50% to 10% takes 12.59 weeks here, against the 5.42 weeks the first halving took. **Nothing has gone wrong. The rate has not changed and cannot change.**

Weeks 17–24. Steps below 0.045. The a posteriori bound (CS-6) at week 20: the step just completed, week 19 to week 20, was 0.0326, giving distance from arrival `<= (0.88/0.12) *`

`0.0326 = 0.2393` — against a true gap of 0.2393. The bound is attained exactly, for the reason §7.6 gives. **The subject can compute their remaining distance from the size of their last week alone, without knowing `Q_inf` at all.**

Week 24. Declared arrival. Gap is under 5% of `Q_inf`. It does not feel like an event, because it is not one — the state at week 24 differs from the state at week 20 by 0.0958 Q-units, about 3.1% of the destination, spread over four weeks, which is at or below most plausible resolutions. **The subject recognises the state because they have been substantially in it since week 16.**

23.14.16 Step 6 — the headroom reading at arrival

```
h_inf = eta ( ln N − F(u*) ) = 0.35 * ( 2.1972 − 1.6864 ) = 0.1788
```

Cross-check from the gauge at `Q_inf`: `0.35*2.1972 − 0.12*3.0853 − 0.22 = 0.178789`. Identical, as §8.4 requires.

The subject arrives with 0.1788 Q-units per week of unclaimed capacity, and §8.4 says exactly what that number is: it is the 23.2% of the ceiling their framings are not generating. It is not free time and it is not energy. **The only way to convert it is to raise `u` toward 1 or to raise `N` above 9,** and [Curzi, 2026c,o] are where those moves live.

23.14.17 Step 7 — the counterfactual, and the levers

Suppose at week 24 the subject wants a higher destination. From §8.5:

Change	New `Q_inf`	Change
baseline	3.0853	—
`eta` 0.35 -> 0.385 (+10%)	3.5772	+15.94%
`delta` 0.12 -> 0.108 (-10%)	3.4281	+11.11%
`u` 0.55 -> 0.605 (+10%)	3.3138	+7.40%
`C` 0.22 -> 0.198 (-10%)	3.2687	+5.94%

The 10% efficiency improvement is worth 2.7 times the 10% cost cut. And the `delta` move carries the §8.6 penalty: at `delta = 0.108` the time constant rises from 7.82 to 8.75 weeks, the time to 5% from 24 to 27 weeks, and the noise gain from 8.33 to 9.26. **Reducing the leak buys a better room, more slowly, with more of every bad week retained.**

23.14.18 Step 8 — the audit that must not be skipped

CS-9. The subject's drive is `eta F(u*) = 0.5902` per week and every component of it is their own: their framings, their conversion, their cost. There is no `x`. The audit passes trivially.

Had there been a transfer — a household absorbing the practice's cost, a colleague covering the work it displaces — the audit would require **that party's** `h`, computed with the transfer as their cost, obtained by asking, with a decline path (CS-10). Per §8.7, an extraction of `x = 0.40` per week against a counterparty whose sustainable ceiling is `x* = 0.3141` would give that party a fixed point of -0.7157 and an exhaustion horizon of **12.04 weeks** — during which the subject's own numbers would look excellent and their fixed point would read 6.4187 instead of 3.0853.

Every measurement the subject could take of themselves would say the practice was working twice as well as it does. That is the whole danger of counterfeit arrival: it is not detectable from inside the system that benefits from it. Only the other ledger shows it, and only if it is asked for.

23.14.19 The one line to carry out of the section

```
Q_t = Q_inf + (1 − delta)^t ( Q_0 − Q_inf )
```

Where you are going does not contain where you started. Where you are is where you are going, plus what is left of where you started, and what is left of where you started is halving every 5.42 weeks. Arrival is not an achievement and not a destination reached by force. It is the moment the second term drops below the resolution of the instrument reading it — and the room, by then, has been furnished for a long time.

23.14.20 14.3 The nine entries over one block

One subject, one twelve-week block, all nine entries. The data are an **illustrative construction** used to demonstrate the full procedure; they describe no real person, and every derived number below was computed.

Setting. An subject running a first full formulary review on a domain of work they have been at for a year. Period = one week. Modulus = 50, assigned to a reference encounter from the preceding quarter and held for the block.

23.14.21 Step 1 — F-03, the ceiling

Framings elicited independently and in writing (Protocol step 1), then compared:

```
N     = 6                     ln N = 1.7918 nats = 2.5850 bits
u     = 0.38 (measured)
N_eff = 1 + 5(0.38) = 2.900   F_int = ln(2.900) = 1.0647 nats
claimed = 1.0647 / 1.7918 = 59.4% of the ceiling
```

Six framings audit to 2.9. That is the first finding of the review and it reframes everything after it: the subject experiences six perspectives and holds slightly under three.

23.14.22 Step 2 — F-01 and F-02, the moment and the calibration

wk	D	P	E	r	|ln r|	M	M_norm
1	62	41	55	1.5122	0.4136	83.171	1.2418
2	58	40	50	1.4500	0.3716	72.500	1.0825
3	66	44	58	1.5000	0.4055	87.000	1.2990
4	60	43	52	1.3953	0.3331	72.558	1.0834
5	64	46	60	1.3913	0.3302	83.478	1.2464
6	59	45	54	1.3111	0.2709	70.800	1.0571
7	55	47	57	1.1702	0.1572	66.702	0.9960
8	52	48	53	1.0833	0.0800	57.417	0.8573
9	50	49	59	1.0204	0.0202	60.204	0.8989
10	49	50	55	0.9800	0.0202	53.900	0.8048
11	51	52	58	0.9808	0.0194	56.885	0.8494
12	50	51	56	0.9804	0.0198	54.902	0.8198

`M_norm` uses a declared range `[0, 120]` scaled to `[0, ln 6]`.

```
mean r,     weeks 1–4  = 1.4644      weeks 9–12 = 0.9904
mean |ln r|, weeks 1–4 = 0.3809      weeks 9–12 = 0.0199
mean M_norm over 12 weeks = 1.0197 nats   (ceiling 1.7918)
```

The subject arrived over-framed by about 46% and left calibrated. Distortion fell from 0.3809 nats — 0.55 bits, 1319 cents — to 0.0199 nats, which is below any plausible deadband. This is the review's success story, and Step 4 is where it turns.

23.14.23 Step 3 — F-08, the echo

Sent moment recorded before the return; returns from two independent sources, timestamped to the exchange:

wk	sent	back	rho	Delta_R
1	55	38	-0.3697	0.3697
2	58	42	-0.3228	0.3228
3	60	47	-0.2442	0.2442
4	57	46	-0.2144	0.2144
5	62	53	-0.1568	0.1568
6	60	54	-0.1054	0.1054
7	58	55	-0.0531	0.0531
8	61	60	-0.0165	0.0165
9	59	57	-0.0345	0.0345
10	60	61	+0.0165	0.0165
11	62	60	-0.0328	0.0328
12	61	62	+0.0163	0.0163

```
mean rho,     weeks 1–4 = −0.2878      weeks 9–12 = −0.0086
mean Delta_R, weeks 1–4 =  0.2878      weeks 9–12 =  0.0250   sd(9–12) = 0.0289
implied over-frame, weeks 1–4 = exp(0.2878) = 1.3335
R_net = −1.5174   R_gross = 1.5830   E_R = 0.9586
```

Fitting the correction gain from the decay of `|rho|` over weeks 1–9:

```
per-week factor g = 0.6911     ->    lambda = 0.3089      R^2 = 0.8808
```

Three independent instruments agree. F-02 says the frame was over-set by 1.46 and came into calibration. F-08 says the returns were under by a factor implying an over-frame of 1.33 and came into calibration. The fitted gain of 0.309 is in the stable monotone band, and `E_R = 0.96` confirms steady directed correction rather than thrash. **The four returns in weeks 9–12 fire the flattery alarm, and the scatter test clears it: the residuals change sign and their sd of 0.0289 is real variation, not a mirror.**

23.14.24 Step 4 — F-04, the storage law, and the turn

Cost from [Curzi, 2026b]'s ledger; `eta` fitted per F-04's protocol.

```
eta = 0.28      delta = 0.09      C = 0.30      Q_0 = 0.40
```

wk	M_norm	Q_t	step dQ	h_t	L_t	admissible?
0	1.2418	0.4000	+0.0117	0.1657	0.9418	yes
1	1.0825	0.4117	-0.0339	0.1646	0.7825	yes
2	1.2990	0.3778	+0.0297	0.1677	0.9990	yes
3	1.0834	0.4075	-0.0333	0.1650	0.7834	yes
4	1.2464	0.3742	+0.0153	0.1680	0.9464	yes
5	1.0571	0.3895	-0.0391	0.1666	0.7571	yes
6	0.9960	0.3504	-0.0527	0.1702	0.6960	yes
7	0.8573	0.2977	-0.0868	0.1749	0.5573	yes
8	0.8989	0.2110	-0.0673	0.1827	0.5989	yes
9	0.8048	0.1437	-0.0876	0.1888	0.5048	yes
10	0.8494	0.0561	-0.0672	0.1966	0.5494	**NO**
11	0.8198	-0.0111	-0.0695	0.2027	0.5198	**NO**
12	—	-0.0806	—	—	—	—

The subject got calibrated and went broke. Stored order fell from 0.40 to below zero over twelve weeks, and the admissibility condition `C <= (1−delta)Q + eta M` failed at week 11 — the week the model itself stops applying, which is a finding and not a rounding artifact.

23.14.25 Step 5 — F-09, the arrival that is a warning

```
mean M_norm = 1.0197
Q_inf = (0.28 * 1.0197 − 0.30) / 0.09 = −0.1609 Q−units       NEGATIVE
tau = 10.603 weeks      half−life = 7.350 weeks
```

Theorem 3.5's failure clause, live: `eta M* < c*`, so the fixed point is negative, and the loop converges anyway — toward the destruction of the thing. F-09's Protocol step 3 exists precisely to catch this **before** a practice is run; the subject ran it after, which is what a first review usually looks like and is why the review was worth doing.

The margin of failure is the number that makes this example worth reading:

```
maximum sustainable cost at u = 0.38:   C <= eta * F_int = 0.28 * 1.0647 = 0.2981
observed cost                        :   C = 0.3000
overshoot                            :   0.0019 Q−units per week  =  0.63%
```

The regime failed by 0.63 percent. Six-tenths of one percent over the sustainability ceiling of Corollary 3.4, sustained for twelve weeks, took the subject from 0.40 Q-units to below zero. There was no crisis, no bad month, no identifiable mistake. **Corollary 3.4 is the single most important sociological statement in the corpus, and this is its personal-scope face: past the ceiling, no choice of interaction geometry prevents order from decreasing, and the arithmetic does not care that the overshoot was small.**

23.14.26 Step 6 — F-05, the headroom paradox

Headroom was **positive and rising** through the entire collapse: 0.1657 at week 0, 0.2027 at week 11. At the fixed point, `h_inf = eta(ln N − F_int) = 0.28 * (1.7918 − 1.0647) = 0.2036`.

Headroom rose because `Q` was falling, and `h` is reduced by `delta Q`. The gauge was correctly reporting that capacity for the new existed *at perfect geometry* while the actual geometry was losing 0.63% per week. This is F-05's central caveat in a live case: **`h > 0` is necessary and not sufficient, and an subject who watched only the headroom gauge would have seen an improving number for twelve weeks.** The geometry gap is the missing piece:

```
geometry gap = eta ( ln N − mean M_norm ) = 0.28 ∗ (1.7918 − 1.0197) = 0.2162
```

That 0.2162 is the part of the ceiling the subject's framings never reached, and it is why `h = 0.20` bought nothing.

23.14.27 Step 7 — F-06 and F-07, the two levers

```
F-06:  sum L over 12 weeks = 8.6365    mean L = 0.7197    L = 0 weeks: 0
       L_max = ln N − C = 1.7918 − 0.30 = 1.4918
```

The launch budget was healthy throughout — never zero, averaging 48% of its ceiling. F-06's "busy trap" in reverse: the subject had budget to imagine with and was losing stored order anyway. `L` and `dQ` are different quantities, and a good `L` does not imply a good `Q`.

F-07, actions-to-begin against weekly initiation count:

actions `k`	observed `n`	fitted `n`
1	22	23.19
2	15	14.86
3	10	9.52
5	4	3.91
8	1	1.03

```
ln n = 3.5888 − 0.4451 k      R^2 = 0.9989
halving distance = 1.557 actions
removing 2 actions from the five-step practice: predicted count 3.91 -> 9.52
```

Removing two actions from the front of the five-step practice is predicted to multiply its frequency by 2.43. This is the review's cheapest intervention and, per PF-9, the one whose prediction should be tested first.

23.14.28 Step 8 — the decision

Four candidate regimes, computed. `M_bar` under a raised `u` is scaled by the ratio of the new `F_int` to the old, which is a modelling choice and is stated as one:

Regime	`u`	`C`	`F_int`	`M_bar`	`Q_inf`	`h_inf`
as-is	0.38	0.30	1.0647	1.0197	**-0.1609**	0.2036
raise `u`	0.55	0.30	1.3218	1.2659	+0.6050	0.1316
cut `C`	0.38	0.24	1.0647	1.0197	+0.5058	0.2036
both	0.55	0.24	1.3218	1.2659	**+1.2716**	0.1316

The minimum requirements, computed exactly:

```
minimum sustained M for Q_inf >= 0:  M >= C/eta = 0.30/0.28 = 1.0714 nats  (59.8% of ln N)
equivalent u, if M is read as F_int: u = (e^1.0714 − 1)/5 = 0.3839
cost that would place Q_inf at 0.5:  C = eta*M_bar − delta*0.5 = 0.2405
```

The subject's measured `u` of 0.38 is 0.0039 below the `u` of 0.3839 that their cost requires. The whole twelve-week decline sits in that fourth decimal place, and both routes out of it are small: cut weekly cost by 0.06 Q-units, or raise `u` from 0.38 to 0.49.

23.14.29 Step 9 — the audit that closes the review

PF-10 and F-09's transfer audit. The subject's drive `eta F_int = 0.2981` per week is entirely their own — no transfer term. But part of their `C` is carried by a household member who absorbs displaced tasks. **That person's headroom is not the subject's to compute.** The review records the question as *asked, answer pending*, with the person free to decline, and the entry stands as **unknown** rather than as satisfied.

This is the only line in the review that cannot be closed by arithmetic, and the declination refusal is why [Curzi, 2026] it is in the report at all.

23.14.30 What the review found, in one paragraph

Six framings that audit to 2.9. A frame over-set by 46% that self-corrected to calibration in eight weeks at a stable gain of 0.31, confirmed independently by the echo log. A healthy launch budget throughout. A rising headroom gauge. And underneath all of it, a weekly cost 0.63% above the sustainability ceiling, converting twelve weeks of genuine improvement into a negative fixed point. No single entry would have found that: F-02 and F-08 said the subject was improving and were right; F-05 said capacity was available and was right. **Only F-04 and F-09, run together against a `Q_inf` computed rather than felt, showed that the room this life was converging on had less in it than the room it started from — by six-tenths of one percent.**

That is what a formulary is for.

23.14.31 14.4 The validation tally

The verdicts the protocol returned on the nine entries are collected here as a count rather than left to be assembled from Section 11.8. **Sixty-three gates were reported across nine entries.**

Entry	Gate 2 (scale type)	Gate 6 (falsification)	Gate 7 (reproducibility)
F-01 The Moment	**DEMOTE** (Standing Verdict 1)	SPECIFIED	UNTESTED — declared
F-02 The Calibration Ratio	**DEMOTE** (Standing Verdict 1)	SPECIFIED	UNTESTED — declared
F-03 The Internal Surplus	CONDITIONAL PASS	SPECIFIED	UNTESTED — declared
F-04 The Self-Storage Law	**DEMOTE** (Standing Verdict 1)	SPECIFIED	UNTESTED — declared
F-05 Personal Headroom	CONDITIONAL PASS	SPECIFIED	UNTESTED — declared
F-06 The Launch Budget	**DEMOTE** (Standing Verdict 1)	SPECIFIED	UNTESTED — declared
F-07 The Gradient	CONDITIONAL PASS	SPECIFIED	UNTESTED — declared
F-08 The Return Difference	**DEMOTE** (Standing Verdict 1)	SPECIFIED	UNTESTED — declared
F-09 The Contraction and Arrival	CONDITIONAL PASS	SPECIFIED	UNTESTED — declared

Five of nine entries are demoted on the scale-type gate. Not one of the nine is marked PASS on the reproducibility gate. The demoted five are exactly the entries whose value passes through the moment `M` or through the calibration ratio `r = D/P`, and the reason is a single external fact about instrumentation: `D` and `P` are not measured on a common ratio scale with

a shared true zero, so their quotient is a number whose value can be set by an admissible relabelling [Curzi, 2026t].

Withdrawn as measured quantities: the magnitude of `D/P`; the distortion $\Delta = |\ln r|$ as a number; and the quantitative form of the moment-to-storage join that supplies `M_t` as the drive in `Q_{t+1} = (1−delta)Q_t + eta M_t − C_t`. Every worked figure above that passes through those quantities is an **illustrative construction demonstrating the arithmetic**, and is labelled as one at its use.

Retained in full force: the limiting cases, the `D = P` fixed point, and the ordering over-framed / calibrated / under-framed. Each requires only ordinal information and each is invariant under every strictly increasing relabelling of the response scale. **What is demoted is magnitude, not structure.**

A separate verdict falls on the audit instrument itself. Standing Verdict 2 records three measured failure modes of the shuffled-input control that every falsification row defers to: against a nominal 0.05 across 500 null trials it false-rejected at **0.108** under independent permutation and **0.120** under circular shift, with only block permutation near nominal at 0.060; a filtered repackaging of the input passed all four baselines at `p = 0.0005` and was caught only by a structural ceiling test, its single-step jump being **38.5 times** the theoretical maximum; and against a confound driving both series the shuffle **passed 98.35 per cent** of 4,000 trials while reporting honestly (48.17 per cent detection) in the clean case [Curzi, 2026u,ac]. **A control measuring 0.108 against a nominal 0.05 is anti-conservative by a factor of 2.2**, and that number belongs beside every result the suite is used to defend rather than in a footnote.

23.15 15 Discussion

23.15.1 15.1 What follows

The four levers are ranked, and the ranking is derived rather than fitted. At the base point `eta = 0.35`, `delta = 0.12`, `N = 9`, `u* = 0.55` and `c* = 0.22`, the elasticities of the fixed point are `eta` **+1.594**, `delta` **−1.000**, `u` **+0.770** and `C` **−0.594**, so that a 10 per cent efficiency gain is worth **2.7 times** a 10 per cent cost cut. Two of these are structural rather than local. The elasticity with respect to `delta` is **exactly −1 at every parameter value**, because `Q_inf` is proportional to `1/delta`: halve the leak, double the destination. And the ratio that ranks efficiency against cost is `eta M*/c*`, which exceeds one whenever the regime is sustainable at all. **Cutting cost is the weakest of the four levers and it is usually the first one attempted.**

Leverage and fragility are the same number. The elasticity with respect to `eta` is `eta M*/(eta M* − c*)`, which diverges as the regime approaches unsustainability. A person operating near the boundary gets very large returns from small improvements in conversion — and equally large losses from small degradations. That is a derived property of the fixed-point formula, not a metaphor, and it is the sharpest practical consequence of the arithmetic.

Convergence certifies nothing, in two distinct ways. First, `eta F(u*) < c*` gives a negative fixed point, and the loop reaches it just as reliably and just as geometrically: a system can arrive, on schedule, at a destination that is a slow depletion [Curzi, 2026a]. Second, a transfer `x` between two systems moves each fixed point by exactly `x/delta` in opposite directions and leaves the sum unchanged. **Transfer relocates a fixed point; it creates none.** At `x = 0.40` against a critical rate `x* = 0.3141`, the donor's steady state is `−0.7157` with an exhaustion horizon of 12.04 periods, while the recipient reads 108 per cent above what its own regime could produce — and every measurement the recipient can take of itself is consistent with a well-run system.

The instrument's resolution is the binding constraint on any personal use of it. A 5 per cent return effect at `sigma_rho = 0.30` needs **139 clean, independent echoes**; at a weekly cycle that is more than two years. Resolving a 20 per cent effect needs nine. **The honest reading is that a personal echo log can detect large distortions and nothing subtle**, and the sample-size law says exactly where the line falls before the log is started rather than after.

Correlated sources attenuate the reading rather than merely widening it. Where a fraction `w` of the apparent independent sources are in fact one source, $\rho_{\text{obs}} = (1 - w)\rho_{\text{true}}$: the estimate is *biased toward zero*, not just noisier, which is the design-effect problem of survey sampling arriving in a personal instrument. The same arithmetic governs the corpus's effective-count work at institutional scope, where a naive count of four jurisdictions collapses to an effective count of one under bloc correlation [Curzi, 2026,?].

23.15.2 15.2 What does not follow

The demotion is not a refutation, and the surviving structure is not a measurement. Five of the nine entries lose their magnitudes and keep their orderings. An ordering is real information — it supports rank correlations, medians, sign tests and monotone regression — and it is not an equation. No demoted quantity may be multiplied, averaged across people, differenced across time as a number, or fed into another formula [Curzi, 2026t].

The reflexive claims in this material are not identified from observation. The internal-scope loop in which a person's construal shapes the world's answer and the answer shapes the construal is a reflexive specification of exactly the kind Manski [1993] analysed, and its endogenous component is not separable from contextual effects or from correlated unobservables in observational data. Shalizi and Thomas [2011] make the same point for network data, and Lyons [2011] show what happens when it is ignored. **This paper adopts no remedy, and every claim resting on the loop is therefore a modelling position rather than an identified empirical result.**

Passing the deadband test would not make the return difference a ratio-scale quantity. Postulate R1 — that the moment enacted and the moment returned lie on one common ratio scale — is the commensurability assumption on which Δ_R rests, and it inherits the whole of the scale-type objection [Curzi, 2026t]. Magnitude estimation is the route by which such a scale could be built [Stevens, 1957, 1946], and it has not been run for moments by anyone.

A high directional efficiency is not evidence of a real effect. The worked log's `E_R = 0.9316` says the returns pointed consistently in one direction; it says nothing about whether the direction was produced by the enacted moments or by a driver common to both series. The shuffle control that is supposed to answer that question passed 98.35 per cent of the time against exactly such a confound [Curzi, 2026u,ac].

The framework does not yet have a differential prediction. Durlauf [2005] asks for a prediction that simpler models do not also make, and Sobkowicz [2020] puts the same question in one line: what observation would distinguish these laws from a neighbouring model that fits equally well? The elasticities of Section 8.5 are the closest this paper comes — they are *point* predictions with numbers rather than directions, which is the right response to Meehl [1967] — but they are predictions of the assumed fixed-point formula and would be made by any model with the same functional form. Watts [2014] is the standing charge that formal social explanation is often commonsense reasoning in dress, and nothing here immunises against it. The external ceiling on how well such things can be predicted at all is low: Salganik et al. [2020] report 160 teams and fifteen years of high-quality longitudinal data with the best machine learning barely beating a four-variable regression, and Hofman et al. [2017] set out the same ceiling generally. Epstein [2008] is the legitimate reply — prediction is one of sixteen reasons to model — and **the corpus is entitled to say which of the other fifteen it claims, but it must say.**

The corpus's own parameters are not exempt from its reflexivity. Lucas [1976] applied to this framework says that `delta`, `eta` and `u` are the decision rules of people who have not read it; Hacking [1995] says the same philosophically, and Ferraro et al. [2005] supplies the warning case in which a theory made people behave worse by assuming the worst of them. Against that, MacKenzie and Millo [2003] documents a reflexive fixed point observed in the field with dates and price data, which is the strongest available existence proof for the mechanism [MacKenzie, 2006, Callon, 1998, Soros, 2013]. **The corpus asserts reflexivity for its subject**

matter and cannot consistently exempt its own parameters; this paper records the inconsistency rather than resolving it.

23.15.3 15.3 Relation to the surrounding literature

The contraction result is Banach's [Banach, 1922] applied to a first-order linear difference equation, and the stability condition and fixed point are textbook [Elaydi, 2005]. The storage law is a leaky integrator and therefore also an exponentially smoothed forecast [Holt, 1957, Brown and Meyer, 1961] and a fixed-gain Kalman filter [Kalman, 1960]; the level-and-rate formalism predates it by forty years [Forrester, 1961], and Sterman [1989] supplies the empirical finding that people systematically misjudge stocks with delays, which is the strongest external support for the practical claims made here about review timing. **This is a shared mathematical form, not a derivation from physics**, and it is labelled as one in Section 3.9.

The ratio register is Weber's and Fechner's [Weber, 1834, Fechner, 1860, 1966], with Stevens' power law as the competing account [Stevens, 1957] and the century-long dispute best summarised by Krueger [1989] rather than adjudicated here. Every failure mode of magnitude estimation that Section 6 lists appears first in Poulton [1968] and Poulton [1979], and the dependence of the exponent on the experimenter's stimulus range is Teghtsoonian [1971]. **A stronger warrant for the logarithm than Fechner's is available and should be used alongside it:** Peters and Gell-Mann [2016] derives a logarithmic form from multiplicative dynamics with no psychological assumption at all, and Peters [2019] supply the reason it matters here — a parameter estimated across individuals does not describe what happens to any one of them over time, which is precisely the direction this paper's quantities are used in.

The design-effect correction is standard survey sampling arriving in an unusual place, and the underlying point — that correlated sources reduce effective independence — is the same one the diversity literature makes about problem-solving groups [Hong and Page, 2004, Thompson, 2014] and that Lorenz et al. [2011] measure directly: social influence collapses the effective count. The "mirror that flatters" is informational cascade [Bikhchandani et al., 1992, Banerjee, 1992] read at personal scope.

Three external results bound what this instrument can claim. Michell [1997] holds that psychology assumes quantitative structure without testing for it, and the conjoint test that would settle it [Luce and Tukey, 1964, Krantz et al., 1971] has not been run for any quantity in this paper. Tversky [1977] shows psychological similarity is asymmetric and violates the triangle inequality, so a `u` built from similarity cannot assume metric behaviour. Schwarz [1999] and Kahneman et al. [2006] show the self-report instrument is exposed to question-wording and focusing-illusion effects, which is a direct threat to a twelve-week personal log.

The enterprise this paper belongs to has a name and a literature, and the corpus is entering it whether or not it intends to. Carrying physical formalism into social and psychological systems was claimed, contested and partially established between 1971 and 2022 [Weidlich, 1971, Galam et al., 1982, Castellano et al., 2009, Jusup et al., 2022], and the standard that programme set is **prediction of observed data, not internal consistency**: Helbing and Molnár [1995] wrote social motivations as literal forces and validated them against measured pedestrian trajectories. Against that standard this paper offers arithmetic on constructed series and nothing else, which is worth stating plainly. Hedström and Ylikoski [2010] names the criterion sociologists will actually apply — an explanation must name a mechanism, not exhibit a fitted curve — and Keuschnigg et al. [2018] is the most constructive position available to a framework in this position. Gallegati et al. [2006] is the field's own indictment by four sympathisers and should be read as a checklist against work of this kind, with McCauley [2006] as the reply.

The decay term has the strongest empirical support of anything in this paper, and it is double-edged. Adaptation to a changed level is well established [Helson, 1948, Brickman et al., 1978]: the hedonic treadmill is the clearest external warrant for a positive `delta` and simultaneously the clearest warning against promising unbounded accumulation. Diener et al. [2006] is the revision that restores room for raising the equilibrium rather than merely returning to it, which is exactly what the four levers of Section 8.5 are for.

Within the corpus, the contraction result appears in three other domains and the parallel is worth stating because it is the same theorem each time. Curzi [2026] reports the contraction as an exact equality with `delta = 0.09` giving `tau = 10.6033`, a half-life of 7.3496, 5 per cent at pass 32, 1 per cent at 48.83 and 0.1 per cent at 73.24. Curzi [2026] states the same two theorems and reports a maximum residual of `2.0624e-15` over a hundred passes, and carries the corpus's register of demotions of which this paper's five are part. Curzi [2026] carries the trichotomy `f > = < f_min` corresponding to `Q_inf > = < 0` with `f_min = k e^(2c/eta)` and `k = 0.483425732988414` — **a defined constant, exact by construction, whose fifteen figures are legitimate for that reason and whose expression as a frequency in hertz is a separate and much stronger claim carrying no external support.** Curzi [2026h] supplies the headroom identity this paper's Section 8.4 depends on: `h_inf = eta(ln N - F(u*))`, verified to six decimals at `eta = 0.35`, `delta = 0.015`, `N = 6`, `u* = 0.55`, giving `Q_inf = 2.840970` and `h_inf = 0.164501`. And Curzi [2026p] is the negative result that keeps momentum out of this formulary: at `beta = 0, 0.5, 0.8, 0.9, 0.95` every trajectory terminates at `x = 1.9972` in basin A, so momentum defeats narrow barriers and is useless against wide shallow ones. Curzi [2026] shows the same concentration arithmetic transferring unchanged to distributed ledgers, with effective counts from 35.0 curated validators (HHI 0.028571) down to 5.94 entities (HHI 0.168400) — **Hill numbers of order 2 on measured shares, which the corpus's own `N_eff = 1 + (N-1)u` is not.**

23.16 16 Verification

The results compound from eight postulates. Each is atomic, each is checkable on its own, and each carries its status.

23.16.1 16.1 Dimensional audit

`[Q]` denotes the dimension of stored order, `[T]` time, `[m]` the magnitude-estimation scale on which moments are placed, and `[1]` a pure number. The period is the unit of `[T]`, so every rate is *per period* and no explicit `[T^-1]` appears once the recursion is written in discrete steps.

Symbol	Dimension	Unit	Status
`M`, `M'`	`[m]`	one common magnitude-estimation modulus, arbitrary but fixed	**not instrumented**; a common ratio scale is *required* and not available
`rho = ln(M'/M)`	`[1]`	nats	**derived**; dimensionless only because `M` and `M'` share `[m]`
$\Delta_R = \lvert\rho\rvert$	`[1]`	nats	**derived**; undefined if either magnitude is non-positive
`c` (Weber fraction)	`[1]`	none	**not measured**; the value 0.10 used below is stipulated
σ_ρ	`[1]`	nats	**estimated** from the log-ratio series, with a standard error
λ (correction gain)	`[1]`	per correction cycle	**estimated** by regression on the log-ratio series

Symbol	Dimension	Unit	Status
`w` (shared-source fraction)	`[1]`	none	**assumed** in the worked case
`Q_t`	`[Q]`	proxy units of stored capability	**measured**, in whatever the declared proxy counts
`F = ln(1+(N−1)u)`	`[1]`	nats	**derived** dimensionless
η	`[Q]` per nat per period	proxy units per nat per period	**measured**; *not* dimensionless
δ	`[1]`	per period	**measured** dimensionless fraction per step
`C`, `c*`	`[Q]` per period	proxy units per period	**measured**
`x` (transfer)	`[Q]` per period	proxy units per period	**assumed** in the worked case
τ, half-life, `t_exhaust`	`[T]`	periods	**derived**
elasticities	`[1]`	none	**derived**, and dimensionless by construction

Does the storage recursion balance? `Q_{t+1} = (1−delta)Q_t + eta M_t − C_t` balances only if every right-hand term carries `[Q]`. The first does because `delta` is dimensionless; the third does by construction; the second does **only if `eta` carries `[Q]` per nat per period**, since `M_t` is used in its normalised form on `[0, ln N]` and is therefore a pure number in nats. `eta` is consequently not dimensionless, and an `eta` reported without a statement of what the proxy counts is half a number [Curzi, 2026a].

Does the return difference balance? `ln(M'/M)` requires a dimensionless argument, which holds if and only if `M` and `M'` are on **one** scale with **one** modulus. Two instruments, however similar, do not license the ratio. That is the whole content of Postulate R1 and it is the reason the entry is demoted.

Where there is no operational unit, it is said. `M` and `M'` have **no operational unit at all** in the corpus as it stands: `[m]` is a placeholder for a scale nobody has built. `Q` is counted in whatever the declared proxy counts, and two people's `Q` values are not comparable unless they declare the same proxy. Comfort itself has no operational unit and is treated as a dimensionless magnitude **by stipulation**, not by measurement.

23.16.2 16.2 The postulates

P1 — The internal store obeys a linear first-order recursion `Q_{t+1} = a Q_t + b_t` with `a` and `b_t` not depending on `Q`. *Status:* **DEFINITION** — the model form, restated from the storage law by setting `a := 1 − delta` and `b_t := eta M_t − C_t` [Curzi, 2026a]. It carries no empirical weight; whether any person is described by it is a postulate recorded in Section 17.

P2 — `0 < delta <= 1` implies `0 <= 1 − delta < 1`. *Status:* **ARITHMETIC.**

P3 — A contraction on a complete metric space has exactly one fixed point, reached from any start. *Status:* **CITED** [Banach, 1922]; the difference-equation form with its stability condition is Elaydi [2005], and the corpus states it exactly, with its *a posteriori* bound, in Curzi [2026].

P4 — The reals with `|x − y|` form a complete metric space. *Status:* **CITED**, elementary [Elaydi, 2005]. It is stated separately from P3 because P3 is useless without it and because a

reader checking the chain should see that the completeness requirement was met rather than assumed.

P5 — The smallest detectable change in a magnitude is a constant fraction of that magnitude. *Status:* **CITED** [Weber, 1834, Fechner, 1860, 1966], with the competing power-law account [Stevens, 1957] and the reconciliation [Krueger, 1989]. **The fraction has not been measured for moment magnitudes**, so the law is cited and its constant is stipulated.

P6 (Postulate R1) — The moment enacted and the moment returned lie on one common ratio scale with one modulus. *Status:* **POSTULATED.** No instrument establishes it. It is the commensurability assumption on which Δ_R depends and it is the postulate whose failure produces the demotion of R6.

P7 — Where a fraction `w` of nominally independent sources are one source, the observed correlation-driven signal is attenuated by `(1 - w)`. *Status:* **ARITHMETIC**, given the independence bookkeeping stated in Section 4.5; the design-effect form `n_eff = n/(1 + (n-1) r_icc)` is the standard survey-sampling statement of the same fact.

P8 — 1.96 is the two-sided 95 per cent quantile of the standard normal distribution. *Status:* **ARITHMETIC**, to the three figures used.

23.16.3 16.3 The compounding chain

```
P1 + P2          -> L1   T(Q) = aQ + b has |T(Q)-T(Q')| = a|Q-Q'| with a = 1-delta < 1,
                         so T is a contraction with modulus 1-delta.          ARITHMETIC

L1 + P3 + P4     -> R2a  T has exactly one fixed point Q_inf = b/delta
                         = (eta M* - c*)/delta, reached from any start.       CITED (banach1922operations)

L1               -> R2b  |Q_t - Q_inf| = (1-delta)^t |Q_0 - Q_inf|, an EQUALITY
                         and not merely a bound, since T is affine.           ARITHMETIC

R2b              -> L2   tau = -1/ln(1-delta); half-life = ln 2 / |ln(1-delta)|;
                         t_p = ln(1-p)/ln(1-delta).                           ARITHMETIC

P6               -> R1   Delta_R = |ln(M'/M)| is a well-formed dimensionless
                         quantity in nats, with sign read first.              POSTULATED (P6)

P5 + P6          -> L3   the smallest detectable return difference is
                         Delta_R,min = ln(1+c).                               CITED + POSTULATED

P8 + L3          -> L4   n >= (1.96 sigma_rho / Delta)^2 echoes are needed to
                         resolve a true return of size Delta.                 ARITHMETIC

P7               -> L5   rho_obs = (1-w) rho_true: correlated sources bias the
                         estimate TOWARD ZERO, not merely widen it.           ARITHMETIC

R2a              -> R4   d ln Q_inf / d ln delta = -1 EXACTLY, at every
                         parameter value, since Q_inf is proportional to
                         1/delta; d ln Q_inf / d ln eta = eta M*/(eta M* - c*);
                         d ln Q_inf / d ln C = -c*/(eta M* - c*);
                         and their ratio is eta M*/c*.                        ARITHMETIC

R2a              -> R5   a transfer x shifts the two fixed points by
                         +x/delta and -x/delta, leaving the sum unchanged;
                         the donor's fixed point crosses zero at
                         x* = eta F_B - C_B.                                  ARITHMETIC
```

```
lambda-loop    -> R3   rho_{t+1} = (1-lambda) rho_t contracts iff
                       |1-lambda| < 1, i.e. 0 < lambda < 2; deadbeat at
                       lambda = 1.                                          ARITHMETIC

R1 + P6 fails  -> R6   the scale-type gate returns DEMOTE on every entry
                       whose value passes through M or r = D/P: F-01, F-02,
                       F-04, F-06, F-08 -- five of nine.                    CITED (ccp:CCP-026) + POSTULAT
```

Every line is checkable from its own inputs. **R2a, R2b, R3, R4 and R5 are arithmetic consequences of one model definition (P1), one arithmetic fact (P2) and one cited theorem (P3) with its completeness condition (P4).** R1 and R6 consume P6, the commensurability postulate, which has no external support — and R6 is precisely the finding that P6 is unmet. **A chain in which the paper's headline result is the failure of one of its own postulates is doing what a validation chain is for.**

What this chain does not establish. It does not establish that any person's stored capability obeys P1. That is a modelling postulate, it is not verified here, and the corpus's one preregistered confirmatory test of the same recursion at institutional scope **failed its second falsification rule**: a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared margin of 10 per cent [Curzi, 2026i].

23.16.4 16.4 The verification instance, worked

The elasticities are the result other papers quote, so they are the ones to check by hand. The base point is `eta = 0.35`, `delta = 0.12`, `N = 9`, `u* = 0.55` and `c* = 0.22`.

```
F(u*)  = ln(1 + (9-1) x 0.55) = ln(5.4)          = 1.6863990
eta F  = 0.35 x 1.6863990                        = 0.5902397
Q_inf  = (0.5902397 - 0.22) / 0.12               = 3.0853304
```

The four partial derivatives, then the four elasticities, each being the derivative times the lever divided by `Q_inf`:

```
dQ/d(eta)   =  F/delta            = 1.6863990/0.12   =  14.053325
dQ/d(delta) = -Q_inf/delta        = -3.0853304/0.12  = -25.711087
dQ/d(C)     = -1/delta            = -1/0.12          =  -8.333333
dQ/d(u)     =  eta(N-1)/((1+(N-1)u) delta)
                                  = 2.8/(5.4x0.12)   =   4.320988

E_eta   =  14.053325 x 0.35 / 3.0853304  =  4.9186638/3.0853304  =  +1.594209
E_delta = -25.711087 x 0.12 / 3.0853304  = -3.0853304/3.0853304  =  -1.000000
E_C     =  -8.333333 x 0.22 / 3.0853304  = -1.8333333/3.0853304  =  -0.594208
E_u     =   4.320988 x 0.55 / 3.0853304  =  2.3765432/3.0853304  =  +0.770281
```

The `delta` row is the one to look at. The numerator of `E_delta` is `-Q_inf/delta` times `delta`, which is `-Q_inf`; dividing by `Q_inf` gives `-1` with no parameter left in the expression. **That elasticity is exactly -1 at every parameter value and is not a property of this base point**, so it is `ARITHMETIC` and carries no uncertainty. The other three *are* base-point quantities and would move if the base point moved; they are derived from stipulated parameters and are reported as model output, not as measurements.

The headline ratio checks in one line:

```
E_eta / |E_C|  =  1.594209 / 0.594208  =  2.68291
eta M* / c*    =  0.5902397 / 0.22     =  2.68291
```

The two agree to five decimals because they are the same quantity, which is the content of the claim that a 10 per cent efficiency gain is worth 2.7 times a 10 per cent cost cut, and which shows the ratio favours efficiency whenever `eta M* > c*` — that is, whenever the regime is sustainable at all.

23.16.5 16.5 Three secondary checks

The exhaustion horizon. With `eta = 0.35`, `F_B = ln(1 + 8 x 0.45) = ln 4.6 = 1.5260563` and `C_B = 0.22`, the critical transfer is `x* = 0.35 x 1.5260563 − 0.22 = 0.5341197 − 0.22 = 0.3141197`, the `0.3141` of Section 8.7. At `x = 0.40` the donor's fixed point is `(0.5341197 − 0.40)/0.12 = −0.7156691`, the reported `−0.7157`. Starting from its undisturbed fixed point `0.3141197/0.12 = 2.6176642`, the horizon is

```
t = ln( 0.7156691 / (2.6176642 + 0.7156691) ) / ln(0.88)
  = ln( 0.7156691 / 3.3333333 ) / ln(0.88)
  = (−1.5382458) / (−0.1278334)
  = 12.0353 periods
```

which is the reported `12.04`. **ARITHMETIC.** The trajectory reproduces the source's table exactly: `2.6177, 2.2177, 1.5559, 1.0434, 0.6466, 0.3393, 0.1013, 0.0032, −0.0830` at `t = 0, 1, 3, 5, 7, 9, 11, 12, 13`.

The correction half-life. With a fitted decay factor of `0.6418`, the implied gain is `lambda = 1 − 0.6418 = 0.3582` and the half-life is `ln(0.5)/ln(0.6418) = (−0.693147)/(−0.443481) = 1.5629` cycles, the reported 1.56 weeks. **ARITHMETIC**, given the fit.

The sample-size law. At `sigma_rho = 0.30` and a target of `Delta = 0.05` nats, `n >= (1.96 x 0.30/0.05)^2 = (11.76)^2 = 138.30`, so 139 echoes. At `Delta = 0.20`, `(1.96 x 0.30/0.20)^2 = (2.94)^2 = 8.64`, so 9. **ARITHMETIC.**

23.16.6 16.6 Limits, asymptotics, and the domain of validity

`delta -> 0^+`. `tau -> infinity`, `Q_inf -> infinity`, the map ceases to contract at `delta = 0` exactly, and the error law degenerates to `|Q_t − Q_inf| = |Q_0 − Q_inf|` — no forgetting, no attractor, no arrival [Curzi, 2026a]. **Zero loss is a pathology and not an ideal**, and the elasticity `E_delta = −1` is silent about it because an elasticity is a local statement.

`delta -> 1`. `tau -> 0`, `Q_t = b` for every `t > 0`: the store is a pass-through with no memory. The model remains well defined and stops being a model of storage.

`eta F -> c*` from above. `Q_inf -> 0^+` and `E_eta = eta M*/(eta M* − c*) -> +infinity`. **The domain in which the elasticity ranking is a useful guide is bounded away from the sustainability boundary**, and the divergence is a real constraint on the model's use, not a defect: near the boundary any lever is enormously powerful in both directions.

`u -> 0`. `F -> 0`, `Q_inf -> −c*/delta < 0`: with no non-parallelism there is no drive and the store depletes to a negative fixed point. **`u -> 1`.** `F -> ln N`, `Q_inf` at its ceiling and headroom exactly zero — the state in which the only remaining move is to raise `N`.

`N -> 1`. `F = ln(1) = 0` regardless of `u`, so a single framing generates nothing. **`N -> infinity`.** `F -> infinity` logarithmically, which is why the ceiling `ln N` is the binding constraint and why raising `N` has decreasing returns.

`lambda -> 0`. No correction; `rho` is a fixed point of the correction map and the distortion persists. **`lambda -> 2`.** `|1 − lambda| -> 1` and the loop becomes a sustained oscillation; beyond 2 it diverges. **`lambda = 1`** is deadbeat: one step to zero in the noiseless case, and complete transmission of the noise in every other case.

`M' -> 0` or `M -> 0`. Δ_R diverges and then becomes undefined. The domain is `M > 0`, `M' > 0`; **a null return is recorded as `Delta_R = 0` only when a return of magnitude equal to what was sent came back, and never as a code for "nothing came back"**, which is the zero clause of Section 3.5 and is the single most abused point in the instrument.

23.16.7 16.7 Invariance, and how far each claim has been tested

Scale invariance of the return difference. `Delta_R` is invariant under `M -> aM`, `M' -> aM'` for any `a > 0` — a common change of modulus — and is *not* invariant under `M -> aM + b`. That is exactly the similarity group, and it is why P6 demands a ratio scale rather than an interval one. Tested algebraically; it is an identity.

Start-independence of the fixed point. `Q_inf` does not depend on `Q_0`. Exact, and exhibited in Section 8.1's convergence tables from three starts.

Conservation under transfer. `Q_A,inf(x) + Q_B,inf(x)` is independent of `x`. Exact, and exhibited at five values of `x` in Section 8.7, where the sum column reads `5.7030` in every row. **This is the only conservation law the framework actually has, it holds only for a transfer between two systems sharing one `delta`, and it is worth being precise about that** because Mirowski [1989] shows that borrowing the notation of a conserved quantity without the conservation law is the characteristic failure of this genre. The corpus's `Q` accumulates, leaks and is spent; **nothing about `Q` is conserved in general**, and Cartwright [1999] is the escape route the framework should take — specify the arrangement within which each law holds and stop claiming a universality it does not need.

How many scales the invariances were tested at. Start-independence: three starting values, one parameter set. Transfer conservation: five transfer values, one parameter set, one shared `delta`. The elasticity identities: algebraic, therefore all parameter values, but **verified numerically at exactly one base point.** No invariance in this paper has been tested on data from more than one constructed case.

23.16.8 16.8 Analogy discipline and the precision of the numbers

Every borrowing here is a **shared mathematical form** and is labelled as one where it is used. A leaky integrator is a leaky integrator whether it holds charge or stored capability [Oppenheim et al., 1996, Horowitz and Hill, 2015]; nothing physical transfers with the equation. The "stability boundary" of the correction loop is the unit-circle condition of a first-order discrete system and carries no dynamical-systems result beyond it. The "design effect" is survey sampling. **No result in this paper is a derivation from physics, and none should be read as borrowing a physical result's authority.**

On significant figures: the elasticities, the transfer identities and the deadband formula are **exact by construction** and are quoted to six decimals where that is useful. The fitted correction gain `lambda = 0.3582` and its `R^2 = 0.8587` are **estimates from an illustrative constructed series** and are so labelled at every use; the underlying series was chosen, not observed, so the quantity being estimated is a property of the construction. `sigma_rho = 0.30` and the Weber fraction `c = 0.10` are **stipulated** and every number computed from them inherits that status. The shuffle-calibration figures 0.108, 0.120, 0.060 and 98.35 per cent are **measured from simulation** [Curzi, 2026u,ac] and are quoted to the precision those runs support.

23.17 17 Limitations

Five of the nine formulary entries carry no measured magnitudes, and this is the paper's own finding rather than a concession extracted from it. F-01, F-02, F-04, F-06 and F-08 fail the scale-type gate because `D` and `P` are not instrumented on a common ratio scale with a shared true zero [Curzi, 2026t]. Every worked number in those entries is an **illustrative construction demonstrating the arithmetic**. The route back is constructive and priced — Rasch person measures on one linked scale, at approximately 123 items per term to establish a half-logit distortion — and it has not been walked.

The commensurability postulate is the load-bearing unsupported step. Postulate R1, that `M` and `M'` lie on one ratio scale with one modulus, has no external support and no instrument. If it fails, Δ_R is not a quantity, the deadband is not a threshold, and the sample-size law prices

the resolution of something that has no magnitude. It is badged at every use and it is the first thing a critic should attack.

The Weber fraction for moment magnitudes has never been measured. The value `c = 0.10` used throughout Section 4.3 and carried into every deadband figure is **stipulated by analogy to sensory continua**, where Weber fractions of that order are typical. Nothing licenses the transfer. H1 is therefore reported as **not tested**, not as supported.

Every empirical number in Section 14 comes from an illustrative construction. The twelve-week log, the recovered gain, the `R^2`, the efficiency, the flattery check, the trajectory tables and the nine-entry review are all computed from parameters that were chosen. **No dataset was collected and no person was studied.**

Michell's objection is conceded, not answered. Michell [1997] and Michell [1999] hold that psychology assumes quantitative structure without testing for it. The conjoint conditions [Luce and Tukey, 1964] have not been tested for `M`, for `Q`, for `u` or for any cost ledger in this paper, and `Q` is treated as a ratio-scale scalar **by stipulation**. A Rasch demonstration [Rasch, 1960, Andrich, 1988] would be the alternative route and has not been run either.

The model-comparison standard is not fully met, and the unmet rows are printed rather than omitted. Against the corpus's five requirements:

Requirement	Status in this paper
Out-of-sample performance reported	**NOT MET.** The correction-gain regression of Section 14.1 is an in-sample fit on all twelve weeks of a constructed series; no held-out block was reserved. The out-of-sample protocol exists [Curzi, 2026i] and was not applied here.
Matched null with the same number of free parameters	**NOT MET** for the correction-gain fit. The four-baseline suite is *specified* in every FVP-6 row and *run* in none of them.
Information criterion (AIC or BIC) reported	**NOT MET.** No AIC or BIC is computed here, and — so far as the author can establish — **no information criterion is computed anywhere in the corpus.**
Number of free parameters stated explicitly	**MET.** The correction-gain model has two (intercept and slope of `ln rho_t` on `t`); the storage recursion has four (`eta`, `delta`, `u`, `C`) plus the initial condition.
Effect sizes with intervals, not only p-values	**PARTIALLY MET.** `R^2 = 0.8587` is reported for the gain fit with no confidence interval on λ, which is a defect; the elasticities are exact and need none; no p-value is reported anywhere in this paper.

The falsification instrument this paper's entries defer to is itself mis-calibrated, and every FVP-6 row inherits all three failure modes. Against a nominal 0.05 across 500 null trials, independent permutation false-rejected at **0.108** and circular shift at **0.120**, with only block permutation near nominal at 0.060; a filtered repackaging of the input passed all four baselines at `p = 0.0005` and was caught only by the structural ceiling test, its single-step jump being **38.5 times** the theoretical maximum; and against a shared slow driver the shuffle **passed 98.35 per cent** of 4,000 trials while reporting honestly (48.17 per cent detection) in the clean case [Curzi, 2026u,ac]. **A control measuring 0.108 against a nominal 0.05 is anti-conservative by a factor of 2.2, and that number belongs beside every result the suite is used to defend.**

The reflection problem is unanswered. No remedy — randomisation, an instrument, partial-population variation, or a defended stipulation of the reference group after the man-

ner of Brock and Durlauf [2001] — is adopted [Manski, 1993, Shalizi and Thomas, 2011]. Every claim resting on the reflexive loop is a **modelling position rather than an identified empirical result**.

The functional form of the surplus term is underdetermined. `F(u) = ln(1+(N−1)u)` is used throughout as the drive. Latané [1981] proposes a sub-linear power law and Briscoe et al. [2006] proposes `n log n`; all three are decelerating and all three fit the qualitative facts [Stumpf and Porter, 2012]. **Nothing in this paper discriminates among them**, and the elasticity `E_u = +0.770` is an elasticity *of the assumed form*, not of the phenomenon.

`N_eff = 1 + (N−1)u` is not a Hill number. It is a linear interpolation between 1 and `N` and is the corpus's own construction, with no supporting literature in ecology, economics, political science or information theory. The Hill numbers proper — `1/HHI` at order 2 and `exp(H)` at order 1 [Hill, 1973, Jost, 2006] — are different quantities and are used elsewhere in the corpus on measured shares [Curzi, 2026,?]. Where an effective count appears in this paper it is the affine one.

No scaling or power-law claim is made, and this is stated so that the absence is deliberate rather than accidental. No exponent in this paper is estimated from a distribution; the procedure of Clauset et al. [2009] is binding on any such claim elsewhere in the corpus and this paper does not substitute for it.

One presentational defect in the source is recorded rather than silently repaired. The source reports the elasticity of `u` as `+0.770` and, in the same table row, that a 10 per cent increase in `u` raises `Q_inf` by **7.40** per cent. Those differ because an elasticity is a first-order local quantity while `F(u) = ln(1+(N−1)u)` is concave in `u`, so a finite 10 per cent step returns less than the elasticity predicts; the corresponding rows for `eta` and `C` agree exactly (+15.94 against +1.594 and +5.94 against −0.594) precisely because `Q_inf` is *linear* in those two. **Both numbers are correct and the source prints them side by side without saying they are different things**; they are separated explicitly in Section 8.5 above, and a reader who reads `+0.770` as predicting a 7.70 per cent response to a 10 per cent move will find 7.40 and think one of them wrong. Everything else checked reproduced exactly: the four elasticities to six decimals, the identity `E_eta/|E_C| = eta M*/c* = 2.68291`, the transfer table's conserved sum of `5.7030` at five transfer rates, the exhaustion trajectory at nine values of `t`, the horizon of 12.04 periods, the correction half-life of 1.5630 cycles, and the sample-size entries 139 and 9.

The parameters of every worked case were chosen, not measured. `eta = 0.35`, `delta = 0.12`, `N = 9`, `u* = 0.55`, `c* = 0.22`, `sigma_rho = 0.30`, `c = 0.10`, `w`, and every transfer rate. The results are exact arithmetic on stipulated inputs and are labelled as such throughout.

23.18 18 Conclusion

The return difference. The quantity a person can read from the world's answer is $\Delta_R = |\ln(M'/M)|$, the absolute log-ratio between the moment enacted and the moment returned, **with the sign read first and the magnitude second**, and it rests on badged Postulate R1 that the two lie on one common ratio scale. `Delta_R = 0` records a **null exchange, not a success**: zero means no information came back, and treating it as a confirmation is the instrument's most common misuse. The smallest audible echo is the Weber deadband $\Delta_{R,\min} = \ln(1+c)$. Resolving a true return of size Δ takes `n >= (1.96 sigma_rho/Delta)^2` clean independent echoes — **139 of them for a 5 per cent return at `sigma_rho = 0.30`, and 9 for a 20 per cent return** — and where a fraction `w` of the sources are in fact one source, `rho_obs = (1−w) rho_true`, so correlation biases the reading **toward zero** rather than merely widening it.

Correction. Correcting a distortion by a fraction λ of the observed error gives `rho_{t+1} = (1−lambda) rho_t`, which contracts for $0 < \lambda < 2$, diverges at $\lambda \geq 2$, and at $\lambda = 1$ is deadbeat — one step to zero in the noiseless case and **completely noise-transparent in every other**. On the worked twelve-week log the sign read over-framing at an implied factor of **1.4467**, the recovered gain was $\lambda = 0.3582$ with `R^2 = 0.8587` and a **distortion half-life of 1.56 weeks**, and directional efficiency was `E_R = 0.9316`.

Arrival, proved rather than asserted. `T(Q) = (1−delta)Q + b` is a contraction with constant `1 − delta`, so Banach supplies the unique start-independent fixed point `Q_inf = (eta M* − c*)/delta`, and the error law is an **equality and not a bound**: `|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|`, which coincides with the closed-form solution. These are the statements downstream papers should quote rather than paraphrase.

The elasticities, stated for citation. At the base point `eta = 0.35`, `delta = 0.12`, `N = 9`, `u* = 0.55`, `c* = 0.22`, giving `F = 1.6863990` and `Q_inf = 3.0853304`, the elasticities of the fixed point are

$$\texttt{eta}\ +1.594 \cdot \texttt{delta}\ -1.000 \cdot \texttt{u}\ +0.770 \cdot \texttt{C}\ -0.594,$$

so that **a 10 per cent efficiency gain is worth 2.7 times a 10 per cent cost cut**. All four are **derived, not fitted, and carry no uncertainty**: they are exact partial derivatives of a closed-form expression evaluated at stipulated parameters. **The `delta` elasticity is exactly −1 at every parameter value**, because `Q_inf` is proportional to `1/delta` — halve the leak, double the destination — and is therefore a structural fact rather than a property of this base point. The other three are base-point quantities and would move if the base point moved. The ratio that ranks efficiency against cost is `eta M*/c* = 0.5902397/0.22 = 2.68291`, identical to `E_eta/|E_C| = 1.594209/0.594208`, and it exceeds one whenever the regime is sustainable at all. **Leverage and fragility are the same number**: `E_eta = eta M*/(eta M* − c*)` diverges as the regime approaches unsustainability.

Convergence certifies nothing. `eta F < c*` yields a negative fixed point that the loop reaches just as reliably. And a transfer `x` between two systems shifts their fixed points by $\pm x/\delta$ and leaves the sum unchanged, so **transfer relocates a fixed point and creates none**: at `x = 0.40` against a critical rate `x* = 0.3141`, the donor's steady state is **−0.7157** with exhaustion at **12.04 periods**, while the recipient reads **108 per cent** above what its own regime could produce.

And the formulary demotes five of its own nine entries. Under the corpus's Formula Validation Protocol [Curzi, 2026t], Standing Verdict 1 returns **DEMOTE on the scale-type gate for F-01, F-02, F-04, F-06 and F-08** — every entry whose value passes through the moment `M` or the calibration ratio `r = D/P` — because `D` and `P` are not instrumented on a common ratio scale with a shared true zero. **Withdrawn as measured quantities: the magnitude of `D/P`; the distortion $\Delta = |\ln r|$ as a number; and the quantitative form of the moment-to-storage join.** Every worked figure derived from them in this paper is an **illustrative construction demonstrating the arithmetic**, labelled as such at its use. **Retained in full force: the limiting cases, the `D = P` fixed point, and the ordering over-framed / calibrated / under-framed**, all of which require only ordinal information and none of which the demotion touches. What is demoted is magnitude, not structure. Standing Verdict 2 records that the shuffled-input control on which every falsification row depends is **anti-conservative at 0.108 against a nominal 0.05**, **blind to repackaging**, and **passed 98.35 per cent of trials against a shared slow driver**. **A formulary that demotes five of its own nine entries, names exactly what survives, prices the study that would restore the rest, and prints the miscalibration of its own audit instrument is the protocol working. It is not a caveat and it is not an apology.**

23.18.1 Declarations

Data and code availability

No external dataset was collected for this paper and no human subjects were studied. Every number in Section 14 — the twelve-week echo log, the recovered correction gain and its `R^2`, the efficiency, the flattery check, the trajectory tables, and the nine-entry formulary review — is computed from stated parameters under the models developed in the text, and every table can be re-derived from the equations and parameter values printed beside it. The source manuscript labels these worked series **illustrative constructions** and that label is retained on every one of them. The shuffle-calibration figures quoted in Sections 11.6, 15 and 17 are

simulation results reported in Curzi [2026u] and Curzi [2026ac] and are theirs, not this paper's. The figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper and the model equations of Sections 4, 8 and 11.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest. The validation verdicts reported here were reached by applying the author's own protocol to the author's own formulary, which is a structural conflict that no declaration removes; it is the reason the demotion count is reported as a headline number rather than as an appendix entry.

Author contributions

Single author. Michael Laurence Curzi derived the results, produced the worked constructions and the figures, ran the formulary audit, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The worked logs are constructed. Two constraints in the source material are design constraints on prospective work and are not reports of any study performed: **an echo source may decline and is then recorded as missing data, never as a null return**, and **a cost carried by another person is counted in that person's ledger, by asking, with a decline path**. Section 8.7's transfer analysis is arithmetic about two coupled recursions and is not a recommendation; its ethical reading is left to the reader.

Preregistration

This paper reports no preregistered test. Every FVP-6 row in the formulary of Sections 11 to 13 records the falsification gate as **SPECIFIED** — the procedure that would run it is given and it has not been run — and every FVP-7 row records **UNTESTED, declared**: no cross-individual study of any of these quantities has been conducted. The preregistered protocol for testing the same recursion at institutional scope, together with the confirmatory application that **failed** its second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared margin of 10 per cent — is reported in Curzi [2026i]. The measured behaviour of the four-baseline suite itself, including a demand-characteristic confound with a true coefficient of zero that passed the whole suite 52.7 per cent of the time, is reported in Curzi [2026u]. Any applied claim resting on this paper inherits both results.

Chapter 24

The Dyad and the Group Moment: Opposition Is Parallel, Calibration Costs All the Headroom, and Distortion Does Not Average

Abstract. At pair scale the surplus law reduces to u = sin^2(theta) with the affine effective count N_eff = 1 + (N-1)u, and the extremality condition yields a result that surprises practitioners: because u depends only on the magnitude of the alignment, a pair in perfect opposition has u = 0, N_eff = 1 and surplus exactly zero, the same as a pair in perfect agreement. Complementarity is orthogonality, not opposition, and a pair at 45 degrees in an N = 12 domain already captures three-quarters of the available ceiling. Sustainability sets a minimum viable angle u_min = (exp(C/eta) - 1)/(N-1), computed at 0.0778 or 16.2 degrees for the worked partnership, below which no joint store exists at all. Non-parallelism decays geometrically with u = 0 globally attracting, so difference is a maintained condition; restoring it every twelve quarters recovers 58 per cent of the total available gain. Reciprocity is asymmetric where u is symmetric: the joint books read identically at imbalance a = 0 and a = 2.20 while one partner's private store carries a negative fixed point. At group scope ln r = ln D - ln p_bar - F makes surplus an exact subtractive correction to distortion, capped at ln N, with an infeasibility floor when D/p_bar > N. Perfect calibration r = 1 is shown to force headroom exactly to zero. Aggregation is refused: an aggregate reading Delta = 0.2324 concealed units at 1.12 and 1.92 and overstated bankable input 4.2553 times. No pair and no organisation was measured: every reproducibility gate reads untested, and the group-scope magnitudes are demoted to illustrative constructions at the scale-type gate.

Keywords: pair surplus; minimum viable angle; reciprocity asymmetry; group moment; calibration ratio; zero headroom; aggregation fallacy

Series: B. Engineering the Surplus **Paper:** CCP-024

24.1 1 Introduction

Two questions sit under most of what is written about collaboration. The first is what a pair of people produces that neither produces alone. The second is what happens to that quantity when the unit of analysis stops being two people and becomes an organisation. This paper answers both inside one formalism and reports three results that run against the usual advice.

The framework's interaction surplus is `F(u) = ln(1 + (N−1)u)`, where `u` measures the non-parallelism of two states in an `N`-category domain and `N_eff = 1 + (N−1)u` is the resulting effective count [Curzi, 2026]. At pair scale the geometry reduces to `u = sin^2(theta)` for `theta` the angle between the two states, and the extremality condition of the surplus functional is stated on the *squared* alignment. That single detail produces the paper's first result. **A pair in perfect opposition and a pair in perfect agreement both read `u = 0`, `N_eff = 1` and surplus exactly zero.** Systematic mutual contradiction is not a source of difference; it is a copy with a sign flip, and it pays a coordination cost for nothing. Complementarity is orthogonality, not opposition, and it is a middle condition with a failure mode at each end.

At group scale the same arithmetic joins the framework's other equation. Writing the organisation's collective perception as a mean intake per channel times an audited effective count of channels, `P = p_bar . N_eff`, and taking logarithms of the doctrine-to-data ratio `r = D/P` gives

```
ln r = ln D − ln p_bar − F
```

an identity in which the group's interaction surplus subtracts from its log-distortion one nat for one nat. The correction is therefore capped at `ln N`, and an organisation whose doctrine load exceeds its domain size — `D/p_bar > N` — cannot be calibrated by any arrangement of its people whatsoever, leaving a residual floor `Delta_min = ln(D/p_bar) − ln N` that only doctrine reduction can move. At the fixed point `r = 1` the input sits at its ceiling, the store reaches its maximum, and headroom is identically zero. Perfect calibration is therefore not an unqualified target, which is not a result anyone would guess from the equation's rhetoric.

The third result is negative and it is about measurement practice. **Distortion does not average.** A headquarters at `Delta = 1.1239` and a field unit at `Delta = 1.9169` aggregate to `Delta = 0.2324`, a figure that reads as near-calibrated while both halves are severely distorted in opposite directions. The aggregate ratio is a weighted mediant of the unit ratios, so it always lies between them, and when they straddle the fixed point it can sit arbitrarily close to it. Computing bankable input from the aggregate overstates it by a factor of 4.26 in the worked case. This is the error Quetelet [1835] committed at the founding of social physics — reading the mean of a distribution as a type rather than as a statistic — and it is the error Kirman [1992] names as the standing problem of aggregation in economics. Granovetter [1978] makes the same point constructively: it is the *distribution* of thresholds, not their mean, that determines what a population does.

Five contributions are new here relative to the surrounding literature rather than relative to the framework. First, the opposition identity and its consequence that a diagnostic built to catch epistemic uniformity by listening for agreement will not catch its mirror image. Second, the minimum viable angle `u_min = (e^{C/eta} − 1)/(N−1)`, a survival line computable from three measurable parameters, below which a pair sustains no joint store at all. Third, the reciprocity asymmetry `a = |ln(eta_AB/eta_BA)|`, which is invisible in a pair's joint account and total in its private ones. Fourth, the exact join between the surplus law and the calibration ratio, with its `ln N` correction ceiling and its infeasibility floor. Fifth, the zero-headroom identity at perfect calibration.

What this paper does not have is a measured organisation. Every reproducibility gate reported below reads untested, the group-scope magnitudes are demoted to illustrative constructions at the scale-type gate, and the paper says so in its Results rather than in a footnote. The framework's own audit of the moment equation returns the same verdict at both scopes and prices the instrument that would restore the magnitudes at approximately 123 items per term [Curzi, 2026]. That is the honest state of the work, and it is stated here so that no reader mistakes a specification for a report.

Sections 2 and 3 state the hypotheses and the framework, including the identification problem that governs every applied reading of the group-scope results. Section 4 develops the pair. Section 5 collects the pair-scope design requirements and failure modes. Section 6 develops the group moment and the join. Section 7 develops the two calibration pathologies and their

opposite corrections. Section 8 works three systems end to end. Sections 9 to 12 discuss what follows, verify the chain from minimal postulates, state the limitations, and conclude.

24.2 2 Hypotheses

The algebraic results below — the opposition identity, the join, the zero-headroom identity — are theorems about equations and cannot be falsified by observation. They are labelled as derivations where they appear. What can be falsified is the claim that the equations describe real pairs and real organisations. The hypotheses separate the two, and each names the observation that would defeat it. The protocol that operationalises them, a preregistered fit against four declared baselines with a frozen model and a held-out window, is specified and applied in Curzi [2026i], whose own confirmatory application failed its second falsification rule with a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin.

H1 — The opposition identity holds behaviourally. A pair whose positions are reliably recoverable from each other by *inversion* produces the same joint output as a pair whose positions are recoverable by *copying*, and less than an oblique pair matched on ability. **Predicted:** measured `u` near zero in both cases, and held-out joint output tracking `ln N_eff` rather than the presence or absence of overt disagreement. **Falsified by:** systematically contradicting pairs outperforming matched agreeing pairs on held-out joint output, at equal measured `u`.

H2 — The minimum viable angle binds. A pair whose measured non-parallelism is below `u_min = (e^{C/eta} − 1)/(N−1)` accumulates no durable joint store, whatever its effort. **Predicted:** for the worked partnership, `u_min = 0.0778`, or `theta = 16.2` degrees; pairs below their own line show a store that declines geometrically at rate `delta` irrespective of activity level. **Falsified by:** a pair sustaining a growing, externally verified store while measured `eta F(u) < C` over several multiples of `1/delta`.

H3 — Non-parallelism decays and difference is a maintained condition. In a working pair, `u` falls geometrically at a per-period rate `lambda > 0`, and `u = 0` is globally attracting. **Predicted:** `u_t = u_0 (1−lambda)^t`, with a measurable half-life `ln 2 / (−ln(1−lambda))`; for the worked practice, `lambda = 0.0600` per quarter and a half-life of 11.2 quarters. **Falsified by:** a constant-`u` baseline predicting held-out `u` as well as the decay model across pairs with no restoration events.

H4 — Reciprocity failure is invisible in the joint account. Where the two directional conversion rates differ, the joint store is unchanged and the private stores diverge. **Predicted:** at fixed mean conversion, the sum of the two private steady states is constant in the imbalance `a = |ln(eta_AB/eta_BA)|` while the individual values separate linearly; at `a = 2.20` one partner's private steady state is negative. **Falsified by:** measured private stores that do not diverge with measured directional imbalance, or a joint store that moves with `a` at fixed mean conversion.

H5 — Surplus enters the calibration ratio as an exact subtractive term. For an organisation whose collective perception is a mean intake times an audited effective count, `ln r` falls by exactly one nat for each nat of interaction surplus gained. **Predicted:** `ln r = ln D − ln p_bar − F` computed via the join agrees with `ln D − ln P` computed directly, to measurement error, and the two disagree only when one of the three measurements is wrong. **Falsified by:** measured `P` failing to track `p_bar . N_eff` across organisations with independently measured `u`.

H6 — Pseudo-diversity manufactures over-framing on a computable schedule. An organisation that adds no doctrine, cuts no collection and changes no policy

crosses from calibrated to over-framed at `t* = ln(u*/u_0)/ln(1−lambda)` with `u* = (D/p_bar − 1)/(N−1)`. **Predicted:** a dated, parameter-free crossing; at `D/p_bar = 4.0`, `u_0 = 0.55`, `lambda = 0.08` and `N = 12`, `t* = 8.41` periods. **Falsified by:** `r` failing to cross 1 within a preregistered tolerance of `t*`, or a constant-`u` model predicting the observed `r` series as well as the decay model.

Four further claims are **not** hypotheses and are labelled accordingly rather than dressed as testable propositions.

The **moment equation** `M = (D/P).E` is a **postulate** of this framework. Nothing in the framework's mathematical papers derives it, and this paper does not upgrade it. The **commensurate-units convention**, which expresses `D`, `P` and `E` in effective independent evidence-channels per decision period, is a **modelling choice** made so that the ratio is dimensionless; without it the formula is dimensionally void. The **aggregation postulate** `P = p_bar . N_eff` is a **modelling choice** adopted because it is the simplest assumption that makes an aggregate behave like a count. The **contact postulate** `S = ln N . e^{−Delta}` is a **postulate**, and its symmetry — over-framing and under-framing attenuating equally at equal `|ln r|` — is its most doubtful commitment and is separately falsifiable by fitting distinct exponents on the two branches.

The identification of a person, or of an intelligence channel, with a unit vector in an `N`-category space is a **postulate** about encoding, testable in principle and untested in fact. It is badged wherever it is used, and Section 11 records that no reproducibility gate in this paper has been run by anyone.

24.3 3 Framework and method

24.3.1 3.1 What is presupposed

Two results are used throughout and are established elsewhere in this corpus. Neither is an empirical claim.

The surplus functional. For two unit state vectors `x` and `y` drawn from a domain that can express `N` distinguishable categories,

```
F(x,y) = f( u(x,y) ),     u(x,y) = 1 − (x . y)^2 = || x ^ y ||^2   in [0,1]

f(u) = ln( 1 + (N−1) u ),     N_eff = e^{F} = 1 + (N−1) u
```

and `f` is uniquely characterised by four axioms — angle-only dependence, zero at parallel, an affine effective count, and normalisation at orthogonality [Curzi, 2026]. The third axiom is a **design axiom**: the effective count is affine in `u` because the framework chose it to be, and every linear-looking result downstream is a consequence of that choice. The same construction, with `u = 1 − rho^2` recovered from error decorrelation and the affine form badged as design axiom S3 rather than as a theorem, is developed for the single-observer case in Curzi [2026o].

The effective-count idiom is standard practice in ecology and political science, where `1/HHI` and `exp(H)` are the two canonical members of the family [Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979]. It must be said plainly here, and it is said again in Section 11, that **`1 + (N−1)u` is not a member of that family.** It is a linear interpolation between 1 and `N`, it is this framework's own construction, and no literature in ecology, economics, political science or information theory contains it. The concentration instrument that *is* a Hill number of order two — `1/HHI`, with the worked jurisdictional case returning a naive HHI of 0.345000 and 2.898551 effective jurisdictions against a bloc-HHI of 1.000000, an effective count of 1.00 and a surplus term of exactly zero — is a separate quantity [Curzi, 2026]. Wherever this paper writes "effective count" it means `N_eff = 1 + (N−1)u` and nothing else.

Two independent warrants exist for the logarithm and the argument needs only the weaker one. The Fechnerian warrant makes sensation proportional to the logarithm of a

dimensionless stimulus ratio measured from a floor [Fechner, 1966]; it requires a contested perceptual story to be true. The stronger warrant assumes nothing psychological: the time-average growth rate of a multiplicatively evolving quantity is the expectation of its logarithm [Peters and Gell-Mann, 2016], so where a stock compounds rather than adds, the logarithm is the ergodic-average object. The ceiling `ln N` also has an information-theoretic reading as the most an ensemble of `N` equiprobable categories can carry [Shannon, 1948], though this framework does not identify surplus with Shannon entropy and the reading is licensed to bound and nothing more.

The functional form is nevertheless underdetermined by the evidence, and this paper says so rather than waiting to be told. Social impact rises as a *power* of the number of sources with an exponent below one [Latané, 1981], and network value has been argued to grow as `n log n` [Briscoe et al., 2006]. Both are decelerating in count, both reproduce the qualitative facts the framework's logarithm reproduces, and **nothing assembled in this corpus discriminates between them.** What the framework's form has that the alternatives do not is an axiomatic characterisation with the design axiom declared at the point where the choice is made; what it does not have is a discriminating measurement. The general form of this problem — many mechanisms producing indistinguishable curves — is stated compactly by Stumpf and Porter [2012], and the demand for a *differential* prediction, one that simpler models do not also make, by Durlauf [2005].

The storage law. The framework's dynamical core is

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,     S_t = F(U_t) = ln(1 + (N−1) U_t)
```

a first-order linear difference equation whose update map is a contraction with modulus `1 − delta`, giving the unique start-independent steady state `Q_inf = (eta S − C)/delta` and, at `delta = 0`, no fixed point at all and a bound degrading to `Q_0 + t eta ln N` [Curzi, 2026a]. The associated headroom gauge is `h_t = eta ln N − delta Q_t − C_t`. The quartet of a level, an inflow, an outflow and a delay is the standard vocabulary of system dynamics, formalised four decades before this framework [Forrester, 1961, Sterman, 1989]; headroom is the framework's name for organisational slack [Cyert and March, 1963], and its behaviour near exhaustion has the same shape as utilisation approaching one in a queue [Kingman, 1962].

24.3.2 3.2 Why the pair is the exact scale

The surplus functional is a statement about a *pair* of states. Every application to a group of forty is a modelling extension of a two-body result, badged as such in the framework's roster-reduction result [Curzi, 2026c]. In the dyad there is no extension and no badge to apply, and the triangle inequality is known to fail for geometric triples, so surplus cannot be chained across three states. **The pair is therefore the only scale at which the law is exact**, which is an argument for measuring pairs directly rather than inferring them from a group figure. It is also the reason the group-scope material in Sections 6 and 7 runs under postulates that the pair-scope material does not need.

24.3.3 3.3 The group-scope postulates

Three postulates carry the material from pair scope to group scope. Each is stated with its badge, because none is derived.

> **The commensurate-units convention [modelling choice].** Express `D`, `P` and `E` in one currency: **effective independent evidence-channels per decision period.** `P` is what the organisation takes in — channels consulted, audited for redundancy, times mean intake per channel. `D` is what its standing model asserts about the same decision space without looking: the decision-relevant propositions doctrine fixes in advance, each weighted by the fraction of decisions it constrains, summed. `E` is

the decision-relevant structure actually present in the environment, in the same currency.

The aggregation postulate [modelling choice]. Collective perception is the mean per-channel intake times the effective count of channels, `P = p_bar . N_eff` with `N_eff = 1 + (N−1)u`, where `N` is the number of distinct domains the organisation's intelligence can express and `u` the mean non-parallelism between its channels.

The contact postulate [postulate]. The surplus a group can bank in a period is the ceiling `ln N` attenuated by its distortion, `S_t = ln N . e^{−Delta_t}` with `Delta_t = |ln r_t|`.

The convention is the weakest joint in the construction and is doing real work: without it the moment equation is dimensionally void, and Section 6.12 shows exactly where it fails. The aggregation postulate is the point at which a hundred analysts reading one wire service become one channel, which is the reason a surplus law belongs in an account of institutional perception at all.

24.3.4 3.4 Identification: the reflection problem

This is the most serious methodological objection to the group-scope half of this paper and it is stated here, in the method, rather than in a footnote.

The framework's reflexive structure is that an aggregate influences its members and the members constitute the aggregate. The aggregation postulate is exactly that structure: `P` is built from the mean non-parallelism `u` of channels that are themselves shaped by the organisation whose perception `P` is. **That is a linear-in-means reflexive specification, and Manski [1993] proves that its endogenous social effect is not identified from observational data.** It cannot be separated from contextual effects or from correlated unobservables, and no quantity of data recovers it. Shalizi and Thomas [2011] reinforce the point for network data specifically: influence and selection are generically confounded, and Lyons [2011] documents what happens when that is ignored.

Four remedies are known and none is free: randomise the treatment, find an instrument, exploit partial-population variation, or stipulate the reference group and defend the stipulation. Brock and Durlauf [2001] shows what a defensible version of the last looks like, deriving a mean-field structure from standard discrete choice with the identification conditions stated.

This paper adopts none of them. The reference group is stipulated — the channels are whatever the organisation's channel audit enumerates — but the stipulation is not defended against the alternatives that would produce the same observations. Consequently **every applied reading of `u`, of `P`, and of the join `ln r = ln D − ln p_bar − F` in this paper is a modelling position rather than an identified empirical result.** The algebra of the join is untouched by this: it is an identity between defined quantities. What is touched is any claim that observing an organisation could confirm the mechanism, and no such claim is made below.

The same objection turned on the framework itself is the Lucas critique [Lucas, 1976]. The framework asserts that systems change when they are modelled and measured; it then treats `delta`, `eta`, `u` and `lambda` as constants over an estimation window. Both cannot be true without qualification. The position taken here is the narrow one: the parameters are treated as stable only over a declared window and only while the order proxy is not itself a reported target — which is why the requirements in Sections 5 and 6 insist on an order proxy with an external referent, and why a measure that becomes a target ceases to be a measure [Goodhart, 1975, Campbell, 1979]. Classifications of people interact with the people classified [Hacking, 1995], and a framework that gives people laws to manage their own arrangements by is building an engine rather than describing a machine [MacKenzie, 2006]. The one documented case in which a theory became true by being used — Black–Scholes prices coming to fit a market that had adopted the model — is MacKenzie and Millo [2003], and it is the existence proof this

framework's reflexivity needs; Jussim and Harber [2005] is the counterweight, finding that self-fulfilling effects in classrooms are real but small and often dissipate.

24.3.5 3.5 Aggregation, and the mean read as a type

Section 7.7's result that distortion does not average is a special case of an old and general problem. Quetelet [1835] founded social physics on *l'homme moyen* and made the error the phrase encodes: a mean of a distribution treated as a type rather than as a statistic. Kirman [1992] states the modern form — a representative agent's behaviour need not correspond to any member's, and aggregation is not benign. Granovetter [1978] shows constructively that it is the distribution of thresholds and not their mean that fixes the collective outcome, and Schelling [1971] is the canonical demonstration that an aggregate pattern need not reflect any individual's preference. Wherever this paper writes a single `Q`, `delta`, `u` or `r` for a group, that objection applies, and Section 7.7 computes its price.

24.3.6 3.6 Standing honesty clauses

1. **The framework's mathematical papers are pure mathematics and make no empirical claim.** The extremality result is about vectors. That partnerships and institutions behave this way is this framework's proposal.
2. **As deployed, the surplus functional is a lower bound** on a pair's or an ensemble's information, not the information itself. An `N_eff` of 3.09 does not mean the pair contains 3.09 perspectives; it means the geometry that was measured guarantees at least that much.
3. **Surplus is not Shannon entropy or mutual information.** The logarithm is not a Shannon logarithm and channel-capacity intuitions do not import.
4. **The triangle inequality fails for geometric triples**, which is why the pair is the exact scale.
5. **The affine effective count is a design axiom**, chosen, not derived.
6. **No thermodynamic claim is made anywhere in this paper.**
7. **The predictive-processing correspondence** sometimes noted for the moment equation — precision-weighted priors against sensory evidence — is a structural correspondence between formalisms, not an identity, and is offered as a parallel and nothing more.

24.3.7 3.7 Analogy discipline

The licence to carry physics formalism into social systems is not one this framework can grant itself: it was claimed, contested and partially bounded by others over the preceding half-century [Galam et al., 1982, Castellano et al., 2009]. Every physical form used here is therefore labelled at first use as one of three things — a derivation, a shared mathematical form, or an analogy — and the discipline is one this corpus already imposes on itself, requiring that a user be able to distinguish a borrowed theorem, an ordinary observed effect, and this framework's own conjecture in three separate sentences on demand [Curzi, 2026]. The storage law's identification with a first-order discrete section is a **shared mathematical form**: the same difference equation governs both, and nothing about charge or capacitance is claimed to hold of an institution [Oppenheim et al., 1996, Elaydi, 2005]. The contraction argument is a **derivation**, because the premises of the fixed-point theorem are literally satisfied by the map [Banach, 1922]. The reading of the contact postulate as a law of the minimum — a group banks at the rate of the scarcer of frame and intake — is an **analogy**, offered as motivation and never as evidence; the framework has no citation for it and none is manufactured here.

The standard this borrowing must meet is the one set by social-force models that were validated against measured pedestrian and vehicle trajectories [Helbing and Molnár, 1995]: **prediction of observed data, not internal consistency.** This paper does not meet that standard and says so. The general critique of importing physical formalism into social science, including the specific hazard of borrowing the grammar of a conserved quantity without the conservation law that gave it content, is Mirowski [1989]; the escape route, which this paper

takes, is to specify the arrangement within which each law is claimed to hold rather than to claim universality it does not need [Cartwright, 1999]. Laws written at the social level should be derived at that level rather than imported from a lower one [Anderson, 1972], and an explanation is owed a mechanism rather than a fitted curve [Hedström and Ylikoski, 2010]. Where the charge is that such models forecast nothing, the answer is that prediction is one of sixteen reasons to model and this paper is claiming several of the others [Epstein, 2008].

No scaling, power-law or heavy-tail claim is made anywhere in this paper, so the maximum-likelihood, Kolmogorov–Smirnov and likelihood-ratio procedure that any such claim would require [Clauset et al., 2009] does not apply. The distributional statements that *are* made — that non-parallelism decays geometrically, and that the store approaches its fixed point geometrically — are exponents **defined by the model, not estimated from data**, and Section 11 records that they have not been tested against a constant-`u` baseline on any real series.

24.4 4 The pair

Sociology begins at two. One person is psychology. Three is already a network with three channels, a majority and a possible coalition. Two is the irreducible case: one relation, no majority, no topology. This section develops the surplus law where it is exact, states the result that surprises practitioners, and computes the price of ignoring it.

24.4.1 4.1 The pair surplus, stated exactly

For two unit states `x` and `y` in a domain that can express `N` distinct categories, `x . y = cos(theta)`, and therefore the identity that runs this section:

```
u = 1 − cos^2(theta) = sin^2(theta)
```

Non-parallelism is the squared sine of the angle between two people's states. It is zero when they point the same way, zero when they point exactly opposite ways, and one when they are at right angles. The middle clause is not rhetorical and Section 4.6 is the reason it matters.

The extremality result of the surplus functional states that `F(x,y) = 0` if and only if `x` and `y` are parallel, that is `(x . y)^2 = 1`, equivalently `u = 0` [Curzi, 2026]. The "only if" direction does the work. It is not merely that parallel states produce little; **zero surplus occurs in exactly one circumstance and no other.** If a pair is producing nothing, the formalism names the cause uniquely: either the states are parallel, or the surplus is strictly positive. The mathematical content is immediate, `f(0) = ln(1 + (N−1).0) = ln 1 = 0`, and strict monotonicity leaves nothing else at zero.

The sociological reading is a separate matter and is badged:

> **The dyadic identification [postulate].** The state `x` of a person, with respect to a given decision domain, is taken to be a unit vector in an `N`-category space whose components are that person's weighting across the domain's `N` distinguishable categories. Two people's non-parallelism `u` is then estimable from disagreement, error decorrelation and prediction diversity.

The framework's mathematics does not say human beings are unit vectors. It says that *if* a system's interaction surplus depends on pair geometry through the four axioms, *then* it takes exactly one form. The identification of a person with a state is this framework's, it is testable, and it is badged so that no reader mistakes the theorem's certainty for the identification's. The single most dangerous objection to that identification is Michell [1997]'s: psychology routinely fails to test whether its attributes are quantitative at all before measuring them. Section 11 concedes it and names the remedy that was not run.

24.4.2 4.2 The domain size is not the headcount

This is the most common error made with the pair formula and it is closed first.

`N` is not the headcount. The headcount in a pair is two, always; if `N` were the headcount, the ceiling on pair surplus would be `ln 2 = 0.6931` nats for every pair in the world. `N` is the **number of distinguishable categories the ambient domain can express** — the dimension of the space the two states live in — entering through the normalisation axiom `f(1) = ln N`.

A couple deciding how to spend a Saturday works in perhaps `N = 6`: rest, repair, obligation, pleasure, money, others. A two-person engineering partnership deciding a structural approach may work in `N = 12` or more. **Two people can carry a large `N` between them, and that is what makes a good pair valuable.**

```
ln N is the hard ceiling on the surplus of ONE interaction

N =  2  ->  ln N = 0.6931          N = 12  ->  ln N = 2.4849
N =  4  ->  ln N = 1.3863          N = 16  ->  ln N = 2.7726
N =  6  ->  ln N = 1.7918          N = 24  ->  ln N = 3.1781
N =  8  ->  ln N = 2.0794
```

A pair working an `N = 12` domain at right angles produces 2.4849 nats; the same two people, parallel, produce 0.0000 in the same domain. **The domain sets the ceiling; the angle decides how much of it is obtained.**

24.4.3 4.3 Angle is the control variable

Because `u = sin^2(theta)`, every statement about a pair can be made in degrees. This is the most useful translation in the section, because people can estimate an angle and cannot estimate a squared wedge norm.

theta (deg)	cos theta	u = sin^2 theta
0	+1.0000	0.0000
5	+0.9962	0.0076
10	+0.9848	0.0302
15	+0.9659	0.0670
20	+0.9397	0.1170
25	+0.9063	0.1786
30	+0.8660	0.2500
40	+0.7660	0.4132
45	+0.7071	0.5000
50	+0.6428	0.5868
60	+0.5000	0.7500
70	+0.3420	0.8830
80	+0.1736	0.9698
90	+0.0000	1.0000

A pair fifteen degrees apart has `u = 0.067` — seven per cent of the way to orthogonal, not seventeen. The squared sine punishes small angles severely: two colleagues who see things a little differently are, in this accounting, almost the same instrument. The other end is equally instructive — between 60 and 90 degrees `u` moves only from 0.75 to 1.00. **Most of the available non-parallelism is bought in the first sixty degrees.** A partner need not be an opposite in everything, only genuinely oblique.

24.4.4 4.4 The pair-surplus and effective-count tables

`F(u) = ln(1 + (N−1)u)`, computed across the domain sizes a working pair might occupy.

u	N=2	N=3	N=4	N=6	N=8	N=12	N=16	N=24
0.00	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
0.01	0.0100	0.0198	0.0296	0.0488	0.0677	0.1044	0.1398	0.2070
0.05	0.0488	0.0953	0.1398	0.2231	0.3001	0.4383	0.5596	0.7655
0.10	0.0953	0.1823	0.2624	0.4055	0.5306	0.7419	0.9163	1.1939
0.20	0.1823	0.3365	0.4700	0.6931	0.8755	1.1632	1.3863	1.7228
0.30	0.2624	0.4700	0.6419	0.9163	1.1314	1.4586	1.7047	2.0669
0.50	0.4055	0.6931	0.9163	1.2528	1.5041	1.8718	2.1401	2.5257
0.70	0.5306	0.8755	1.1314	1.5041	1.7750	2.1633	2.4423	2.8391
0.90	0.6419	1.0296	1.3083	1.7047	1.9879	2.3888	2.6741	3.0773
1.00	0.6931	1.0986	1.3863	1.7918	2.0794	2.4849	2.7726	3.1781

The effective count is the same information in the units practitioners can argue about.

u	N=4	N=8	N=12	N=24
0.00	1.000	1.000	1.000	1.000
0.05	1.150	1.350	1.550	2.150
0.10	1.300	1.700	2.100	3.300
0.20	1.600	2.400	3.200	5.600
0.30	1.900	3.100	4.300	7.900
0.50	2.500	4.500	6.500	12.500
0.70	3.100	5.900	8.700	17.100
0.90	3.700	7.300	10.900	21.700
1.00	4.000	8.000	12.000	24.000

In an `N = 12` domain a pair at `u = 0.50` has an effective count of **6.5** — two people, six and a half. The same pair at `u = 0.05` has an effective count of **1.55**. **The same two human beings register 6.5 or 1.55 depending only on the angle between their states.** Ability does not move that number; ability lives in `N` and in the conversion rate `eta`. The angle is its own parameter and it is the one usually left unmanaged.

24.4.5 4.5 The marginal structure

Differentiating, `dF/du = (N−1)/(1 + (N−1)u)`. At `N = 12`:

u	F(u)	dF/du
0.00	0.0000	11.0000
0.05	0.4383	7.0968
0.10	0.7419	5.2381
0.20	1.1632	3.4375
0.30	1.4586	2.5581
0.50	1.8718	1.6923
0.70	2.1633	1.2644
0.90	2.3888	1.0092
1.00	2.4849	0.9167

The marginal return on non-parallelism at `u = 0` is **twelve times** the marginal return at `u = 1`, which is strict concavity restated. The engineering consequence is asymmetric. **The first**

small increment of difference in a pair is the most valuable thing that can be added to it, and the first small increment of convergence is the most expensive thing that can be lost. A pair drifting from `u = 0.10` to `u = 0.05` loses 0.3036 nats over an interval of 0.05, a rate of 6.07 nats per unit of `u`. A pair drifting from `u = 0.90` to `u = 0.50` loses 0.5170 nats over an interval of 0.40, a rate of 1.29 nats per unit. **Difference lost near the parallel wall is worth nearly five times what the same difference is worth near orthogonality.**

As a fraction of the ceiling at `N = 12`, `F/ln N` runs 0.1764 at `u = 0.05`, 0.2986 at 0.10, 0.4681 at 0.20, 0.5870 at 0.30, 0.7533 at 0.50, 0.8706 at 0.70 and 0.9613 at 0.90. **A pair at 45 degrees already captures three-quarters of what a perfectly orthogonal pair could.** Perfection is not the target; obliqueness is.

The monotonicity of surplus in `u` is, however, the framework's most exposed claim and the literature runs both ways. Supporting it: the bridging value of weak ties and structural holes, with measured effect sizes [Granovetter, 1973, Burt, 2004]; the ecological result that a portfolio of distinct components buys stability [Schindler et al., 2010]; and the finding that diverse problem-solver collections can outperform higher-ability ones [Hong and Page, 2004]. Against it, and cited here rather than left for a reviewer: the formal objection that the last result is a property of its modelling assumptions rather than a theorem about groups [Thompson, 2014]; the measured inverted U in creative collaboration, where success rises with cohesion and then falls [Uzzi and Spiro, 2005]; the short-run costs of surface-level diversity [Mannix and Neale, 2005, van Knippenberg and Schippers, 2007]; the finding that diversity in the short run reduces trust and civic participation, from a scholar who did not want the result [Putnam, 2007], with its rebuttal [Abascal and Baldassarri, 2015]; and the ecological results that increasing diversity and connectance destabilises randomly assembled systems [May, 1972]. **A monotone claim cannot survive that list unqualified**, and this paper restricts it to the regime the arithmetic covers: `u` in `[0,1]` with `C` held fixed. Where cohesion and coordination cost respond to `u`, the curve may turn over and this framework does not say where. Contact across difference produces value only under conditions [Allport, 1954], and `u` needs the same qualification. Finally, three incompatible constructs travel under the word "diversity" — separation, variety and disparity [Harrison and Klein, 2007] — and **`u` is a disparity measure**: it registers how far apart two states are, not how many types are present.

24.4.6 4.6 Opposition is also parallel

Write the alignment between the two states as `rho = x . y`, a correlation in `[−1, +1]`. Then `u = 1 − rho^2`, and `u` depends only on the *magnitude* of `rho`. At `N = 12`:

rho	u	N_eff	F
+1.00	0.0000	1.0000	0.0000
+0.95	0.0975	2.0725	0.7288
+0.90	0.1900	3.0900	1.1282
+0.70	0.5100	6.6100	1.8886
+0.50	0.7500	9.2500	2.2246
0.00	1.0000	12.0000	2.4849
−0.50	0.7500	9.2500	2.2246
−0.90	0.1900	3.0900	1.1282
−1.00	0.0000	1.0000	0.0000

A pair in perfect opposition has `u = 0`, `N_eff = 1` and surplus exactly zero — the same as a pair in perfect agreement. The extremality result covers both cases, because it is stated on `(x . y)^2 = 1`, which is satisfied by `rho = +1` and by `rho = −1` alike. This is not a quirk of notation but a sociological statement following from the geometry with no interpretive help:

> **The mirror and the mirror-image are the same instrument.** A partner who agrees about everything and one who disagrees about everything carry identical infor-

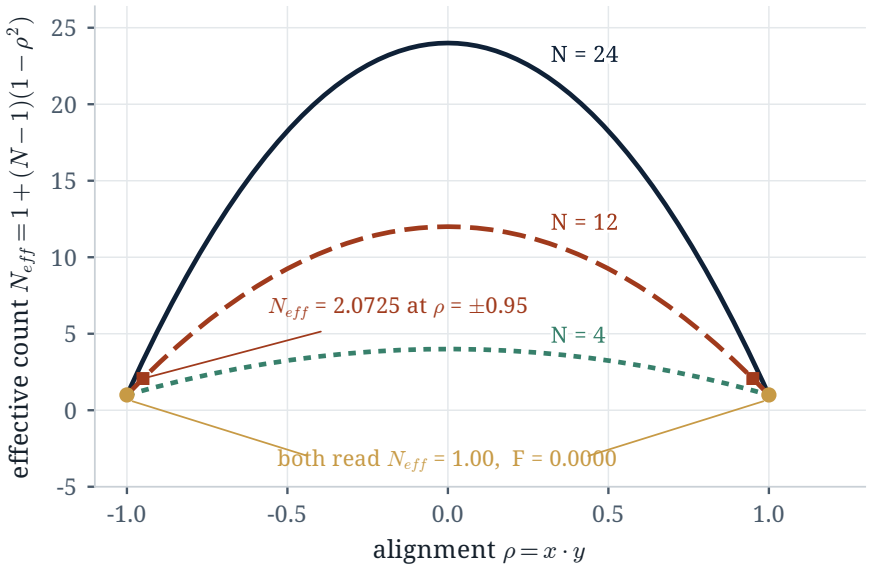


Figure 24.1: **A pair in perfect opposition and a pair in perfect agreement both read an effective count of exactly 1.00 and surplus exactly 0.0000 nats**, because non-parallelism `u = 1 - rho^2` sees only the magnitude of the alignment. At `N = 12` the effective count peaks at 12.00 where the states are orthogonal and falls symmetrically to 1.00 at both `rho = +1` and `rho = −1`; the culture-fit hire at `rho = +0.95` and the sparring partner at `rho = −0.95` both read 2.0725. Computed from the model at the parameter values tabulated in Section 4.6; generated by `figs/src/CCP-024_fig01.py`.

mation, because either position is fully recoverable from the other's. Reflexive contradiction is a copy with a sign flip.

Much of what is called creative tension in partnerships is `rho` near `−1` and generates no surplus at all. It generates heat, which enters the model as cost `C`, not as input `S`. **An opposition pair is a parallel pair with a cost overhead**: it pays `C` and returns zero surplus, a worse position than working alone, since a lone operator at least pays nothing for the redundancy.

The result is not local to the pair. It is the same arithmetic as the framework's roster reduction, where `u` enters `N_eff = 1 + (N−1)u` and `u` sees only `|rho|`, so a systematic contrarian is audited out of the effective count exactly as a systematic agreer is [Curzi, 2026c]. It is also a failure mode **distinct from** the parallel-wall pathology while sharing its arithmetic: the parallel wall is a room that sounds like consensus, the opposition pair is a room that sounds like debate, and both read `N_eff = 1`. **A diagnostic built to catch the parallel wall by listening for agreement will not catch this one**, which is why the obliqueness test below is stated on predictability rather than on assent. The identity recurs across the corpus in domains with no social content at all: in the frequency register, `rho = +1` and `rho = −1` both give `u = 0`, `F = 0` and the same folded note [Curzi, 2026]; in the concentration register, three nominally distinct jurisdictions with a naive effective count of 2.898551 collapse to a bloc count of 1.00 and a surplus term of exactly zero [Curzi, 2026]. The harmonic register carries an exact pairwise identity of the same family, the doubling theorem `HD_ess(i,j) = 2log2(pq)` with `G_ess = 2G(pq) − 1` for the essence ratio `p^2/q^2`, **verified exactly over all sixty-six pairs at `N = 12`** [Curzi, 2026] — which is the closest thing in this corpus to an exhaustively checked pair-scale result, and the standard against which the present section's single-instance checks should be read.

24.4.7 4.7 Complementarity is orthogonality, not opposition

Reading the alignment table by name: `rho = +1.00` is identity, a copy; `rho = +0.95` is the culture-fit hire, `N_eff` 2.07; `rho = +0.50` is complementarity, `N_eff` 9.25; `rho = 0` is maximal complementarity, `N_eff` 12.00; `rho = −0.95` is the sparring partner, `N_eff` 2.07 again; `rho = −1.00` is identity inverted, also a copy.

Complementarity is a middle condition and both extremes of agreement are its failure modes. Advice to find someone who completes you is, read through the law, correct; advice to find someone who challenges everything you say prescribes the right-hand collapse of the same curve.

The operational test is neither "do we agree?" nor "do we argue?":

> **The obliqueness test.** Given the partner's position on a question in the domain, can it be predicted from one's own — by copying it, or by inverting it? If either rule works reliably, `u` is near zero and the pair is producing nothing. **A useful partner is one whose position cannot be computed from one's own by any simple rule.**

Two cautions attach. If `u` is read as a similarity measure over people it cannot be assumed to behave metrically, because psychological similarity is asymmetric and violates the triangle inequality [Tversky, 1977] — the empirical counterpart of the framework's own clause that the triangle inequality fails for geometric triples. And the measured alignment between two people is not a fixed property: social influence measurably collapses the effective count of an ensemble [Lorenz et al., 2011], which is Section 4.12's subject.

24.4.8 4.8 Crossing and divergence within a pair

Non-parallelism decomposes into two sources, and the decomposition applies to a pair as cleanly as to a roster [Curzi, 2026]. For a pure block vector,

```
u = u_cross + u_div        u_cross = beta^2            (between-category difference)
                           u_div   = alpha^2 sin^2 gamma   (within-category difference)
                           alpha^2 + beta^2 = 1

max( f(u_cross), f(u_div) )  <=  f(u)  <=  f(u_cross) + f(u_div)
```

At `N = 12`, worked:

beta^2	gamma	u_cross	u_div	u	f(u_c)	f(u_d)	f(u)	max	sum
0.00	30deg	0.0000	0.2500	0.2500	0.0000	1.3218	1.3218	1.3218	1.3218
0.25	30deg	0.2500	0.1875	0.4375	1.3218	1.1192	1.7600	1.3218	2.4410
0.25	60deg	0.2500	0.5625	0.8125	1.3218	1.9723	2.2963	1.9723	3.2941
0.50	45deg	0.5000	0.2500	0.7500	1.8718	1.3218	2.2246	1.8718	3.1936
0.60	20deg	0.6000	0.0468	0.6468	2.0281	0.4152	2.0937	2.0281	2.4434
0.00	80deg	0.0000	0.9698	0.9698	0.0000	2.4569	2.4569	2.4569	2.4569

Every row satisfies `max <= f(u) <= sum`, and the bound is never tight in the middle rows. That is the honest content of subadditivity: **the pair never gets more than the sum of the two sources, and usually noticeably less.** No synergy beyond additivity is promised, and a design assuming it is unfunded.

At pair scope the two sources have plain names. **`u_cross` is a pair that spans two trades** — an engineer and a clinician, a composer and a statistician: different vocabularies, different training lineages, different failure intuitions. **`u_div` is a pair that shares a trade and disagrees inside it** — two structural engineers who differ on the same question. They are not substitutes. The `beta^2 = 0.60, gamma = 20` degree row is the cross-disciplinary pair that never argues; the

`beta^2 = 0`, `gamma = 30` degree row is the same-trade pair that argues constantly. Both have real surplus. Neither has both kinds. **Diagnose both, and do not accept one as evidence of the other.** The decomposition has been carried through in a group case at `N = 6`, returning `u_cross = 0.72`, `u_div = 0.03`, `u = 0.75`, a balance index of 0.0206, ceiling utilisation of 0.8696 and zero within-block dissents across forty decisions [Curzi, 2026d], which is the extreme instance of one source starved.

24.4.9 4.9 The pair's cost floor

Coordination channels count as `n(n−1)/2` [Brooks, 1975, Conway, 1968]:

```
n     channels    channels per person
---   --------    -------------------
 2        1              0.50
 3        3              1.00
 4        6              1.50
 5       10              2.00
 6       15              2.50
 8       28              3.50
```

The pair is the only human structure with fewer channels than people. Everything else costs more in coordination than it has heads to pay with. That is the structural reason pair-based practices — the two-person cockpit, the two-person research collaboration, the marriage — recur across every domain that works under load. It also means the pair has the highest headroom per person available to any social form, because `C` is at its floor.

The corollary is uncomfortable. **A pair that produces nothing has no excuse available.** It cannot blame coordination overhead, because it has the minimum, and under the extremality result exactly one cause remains. A parallel argument in the frequency register reaches the same place from the cost side: the net octave gain of a group of `m` runs `NET(m) = 2log2(m) − 1.442695(kappa/eta)m(m−1)`, so a pair is viable in the best case if and only if `kappa/eta <= ln 2 = 0.693147` [Curzi, 2026].

24.4.10 4.10 Reciprocity: non-parallelism is symmetric, conversion is not

`u(x,y) = u(y,x)` identically, since it is built from `(x . y)^2`. Therefore:

> **The surplus of a pair is a joint property. Neither person owns it, neither person produced it, and it cannot be divided between them.**

No accounting scheme can attribute a pair's surplus to one member, and any incentive structure that tries is measuring something the law says does not exist. **Conversion is a different matter, and conversion is where reciprocity lives.**

> **The directional conversion postulate [postulate].** A pair has two conversion parameters, one per direction. The joint store converts at the mean `(eta_AB + eta_BA)/2`; each partner's private store converts at the rate directed toward them. Reciprocity is the condition `eta_AB = eta_BA`, and its violation is measured on a log-ratio, `a = |ln(eta_AB / eta_BA)|`.

```
a = 0        symmetric
a = 0.4055   a 3:2 imbalance
a = 0.6931   a 2:1 imbalance
a = 1.3863   a 4:1 imbalance
a = 2.1972   a 9:1 imbalance
```

Running each partner's private store at `delta = 0.15`, `C_A = C_B = 0.25`, `u = 0.45`, `N = 12`, so `S = 1.7834`, and holding the **mean conversion fixed at 0.50** in every row:

eta_AB	eta_BA	mean	a	Q_A(inf)	Q_B(inf)	sum
0.50	0.50	0.50	0.0000	+4.2780	+4.2780	+8.5559
0.60	0.40	0.50	0.4055	+3.0890	+5.4669	+8.5559
0.70	0.30	0.50	0.8473	+1.9001	+6.6558	+8.5559
0.80	0.20	0.50	1.3863	+0.7112	+7.8448	+8.5559
0.90	0.10	0.50	2.1972	−0.4777	+9.0337	+8.5559

The sum is constant down the whole column and the individual outcomes are not, which is exactly what an aggregate measurement misses:

> **Reciprocity failure is invisible in the joint account and total in the private one.** The partnership's books read identically at `a = 0` and at `a = 2.20`. At `a = 2.20` one partner's private store has a *negative* steady state — the storage law's declared failure mode, reached not by any deficiency of the pair's geometry but purely by the direction of its conversion [Curzi, 2026a].

A partnership described as productive but one-sided is describing exactly this. The productivity is real and sits in the sum; the one-sidedness is real and sits in the split. Both are true at once. The wider literature on reciprocity supplies the normative half of this — the dyad-to-group move has its formal core in repeated interaction and indirect reciprocity [Axelrod, 1984, Nowak and Sigmund, 2005], and the governance of shared stores under monitored reciprocity is the subject of a large empirical programme [Ostrom, 1990] — but **none of that literature contains this quantity**, and `a` is offered here as this framework's construction.

24.4.11 4.11 The return difference

For a pair, let `M` be what one partner sends into the relation and `M'` what comes back. The return difference is `Delta_R = |ln(M'/M)|`.

M'/M	Delta_R	
3.00	1.0986	
2.00	0.6931	
1.50	0.4055	
1.25	0.2231	
1.10	0.0953	
1.00	0.0000	<– no information returned
0.90	0.1054	
0.75	0.2877	
0.50	0.6931	
0.33	1.1087	

`Delta_R = 0` means the pair returned the sender's own signal unchanged. **An echo is not a reply.** In the pair this is the behavioural face of `u = 0`: the partner whose response is always the sender's own position restated is not a partner but a surface. The magnitude says how much came back; the sign of `ln(M'/M)` says which direction the world moved the sender.

24.4.12 4.12 Every working pair destroys its own non-parallelism

Here is the structural problem that makes this a maintenance question rather than a selection question. A pair that works together converges: shared vocabulary, shared references, shared

priors, shared blind spots. Every one of those is `Q` going up and `u` going down. **The pair's stored order and its input geometry are in direct conflict, and the conflict is not a design error. It is what success looks like from the inside.**

The same mechanism is documented outside this framework and with a simulation attached: an organisation whose members learn fast homogenises and destroys the variance that produced its advantage [March, 1991]. Formal models of belief updating reach the same fixed point — a row-stochastic influence matrix drives a population to consensus [DeGroot, 1974] — and the device that prevents collapse in those models, a stubbornness term anchoring each agent to its initial position [Friedkin and Johnsen, 1990], is formally the same object as this framework's leak. Bounded-confidence models predict the *number* of surviving clusters from a coupling parameter [Deffuant et al., 2000, Hegselmann and Krause, 2002], which is an effective count derived from an interaction constant and is prior art for this genre of result.

Model the drift as geometric decay, `u_t = u_0 (1 - lambda)^t`:

```
lambda   half-life of u (periods)
------   -------------------------
0.02              34.31
0.05              13.51
0.08               8.31
0.10               6.58
0.20               3.11
0.30               1.94
```

At `u_0 = 0.60`, `lambda = 0.10`, `N = 12`, twenty periods run as follows:

```
t      u        N_eff     F
---    ------   -------   ------
 0     0.6000    7.6000   2.0281
 5     0.3543    4.8972   1.5887
10     0.2092    3.3013   1.1943
15     0.1235    2.3589   0.8582
20     0.0729    1.8024   0.5891
```

Twenty periods of successful collaboration take a pair from an effective count of 7.6 to an effective count of 1.8. Nobody did anything wrong. They agreed with each other, repeatedly, because they were learning from each other, which is what they were supposed to do. **A pair has no internal source of `u`. It must import.** The geometric form is a modelling choice and is badged as one: pairs may converge in bursts around shared projects rather than smoothly, and the exponent here is defined by the model rather than estimated from data.

24.4.13 4.13 The minimum viable angle

Setting the steady state to zero and solving for the non-parallelism below which the pair holds no store at all: `Q_inf >= 0` requires `eta F(u) >= C`, so

```
S_min = C / eta
u_min = ( e^{C/eta} - 1 ) / (N - 1)
```

For a pair with `eta = 0.55`, `C = 0.34`, `N = 12`:

```
S_min = 0.34 / 0.55 = 0.6182
u_min = (e^0.6182 - 1) / 11 = 0.0778
N_eff at u_min = 1.8556
```

Below `u = 0.078` this pair sustains no joint store whatsoever. The relationship does not merely become less productive; its equilibrium goes negative and everything it built erodes at rate `delta`. In degrees, `u = 0.0778` is `theta = 16.2`. **Sixteen degrees is this partnership's survival line**, and every pair has such a line, computable from three numbers it can measure. The line is a threshold in the modelling sense and nothing more: it is where a linear model reports that it has left its own domain, not a cliff observed in any pair.

24.4.14 4.14 Maintenance beats selection, computed

If `u` decays and there is no internal source, the only lever left is periodic re-injection: importing difference on a schedule. Take the same pair — `eta = 0.55`, `delta = 0.12`, `C = 0.34`, `N = 12` — starting at `u_0 = 0.62` with drift `lambda = 0.06` per quarter, and compare restoring `u` to 0.62 every `k` quarters against no restoration, measured by the store after 24 quarters:

```
restoration interval    Q at t = 24
--------------------    -----------
never                      2.7064
every 12 quarters          4.7070
every  8 quarters          5.3283
every  6 quarters          5.6209
every  4 quarters          5.8991
every  2 quarters          6.1615
```

Restoring once every twelve quarters — one deliberate importation of difference every three years — recovers **58 per cent** of the gap between doing nothing and doing it every six months: 4.7070 against 2.7064, with a maximum of 6.1615. Moving from twelve-quarterly to six-quarterly buys 0.9139 further units of store; moving from six-quarterly to two-quarterly buys 0.5406 more.

The returns on restoration frequency are themselves concave, and the first restoration is worth more than all subsequent increases in frequency combined. A pair that never restores loses roughly two-thirds of its potential store and experiences the loss as fatigue. What counts as a restoration is specified in requirement S-12: not a holiday, but the deliberate introduction of a state the pair cannot predict from its own. The general finding that a leaky integrator responds to cadence rather than to intensity is developed with its step, impulse and frequency responses in Curzi [2026a].

24.4.15 4.15 Validation gates at pair scope

Each formula is carried through a seven-gate protocol: dimensions, scale type, limiting cases, fixed points, identifiability, falsification, and reproducibility. The tables below are the findings. **Every reproducibility row reads untested**, and that is the honest state of the work rather than an omission.

Table V-P1 — Non-parallelism `u = 1 - (x.y)^2`

Gate	Finding
1 Dimensions	`x`, `y` unit vectors; `x.y` dimensionless; `u` dimensionless in `[0,1]`. Both sides agree. **PASS.**

Gate	Finding
2 Scale type	`u` is ratio-scaled with a true zero (exact parallelism) and a true maximum. The *components* of `x` need at least interval scale to normalise; with ordinal rankings only, `u` must come from a rank-correlation surrogate and **demotes to an ordering, not a magnitude**. **CONDITIONAL PASS — the encoder must be declared.**
3 Limiting cases	`u -> 0`: identical or exactly opposite states (Section 4.6). `u -> 1`: orthogonal, no mutual predictability. Both nameable; `u` is bounded, so no divergence. **PASS.**
4 Fixed points	`u = 0`, the invariant of the copy relation and the unique zero of `F`; `u = 1`, the invariant of independence. Both attained. **PASS.**
5 Identifiability	`u` is recoverable from disagreement rates and error decorrelation without observing `x` and `y`. It does **not** distinguish `rho = +0.9` from `rho = -0.9`; sign information is lost and must be recovered separately. **PASS WITH DECLARED AMBIGUITY.**
6 Falsification	Preregister `u` from pre-discussion disagreement; predict held-out joint output; beat persistence, direct-input, matched autoregressive and shuffled-input baselines. **Kill condition: `u` at `t` fails to improve prediction of `t+1` output over the shuffled control.**
7 Reproducibility	The cross-pair test: same operationalisation, different observer, different pair, same domain — does `u` recover? **UNTESTED. This is the gate that decides whether the dyadic identification is science, and it has not been run.**

Table V-P2 — Pair surplus `F = ln(1 + (N-1)u)` and `N_eff = 1 + (N-1)u`

Gate	Finding
1 Dimensions	`(N-1)u` dimensionless, so the logarithm's argument is dimensionless as required. `F` in nats; `N_eff` a pure count. **PASS.**
2 Scale type	`N` a count (ratio scale, true zero); `u` ratio-scaled per Table V-P1; `F` in nats inherits ratio scale. **PASS, conditional on V-P1's encoder.**
3 Limiting cases	`u -> 0`: `F -> 0`, `N_eff -> 1`, the non-event. `u -> 1`: `F -> ln N`, the ceiling. `N -> infinity` at fixed `u > 0`: `F` diverges logarithmically — real divergence, so slow as to be useless. `N -> 1`: `F -> 0` for all `u`, correct, since a one-category domain admits no difference. **PASS.**

Gate	Finding
4 Fixed points	`F = 0` if and only if `u = 0`; `F = ln N` if and only if `u = 1`. Unique both ways. **PASS — the strongest fixed-point result in the paper.**
5 Identifiability	`N` from enumerating declared categories, `u` from the decorrelation protocol — separately recoverable. But `F` alone does **not** identify `(N, u)`: `F = 0.6931` arises from `(N=2, u=1)` and from `(N=12, u=0.0909)`. **Report `N` and `u`, never `F` alone. PASS ONLY IF BOTH ARE REPORTED.**
6 Falsification	As V-P1, plus the preregistered prediction that output scales with `ln N_eff` rather than with headcount or mean ability. **Kill condition: headcount or mean ability predicts held-out joint output better than `ln N_eff`.**
7 Reproducibility	**UNTESTED at pair scope.** The nearest published evidence that group output tracks something other than member ability is a collective-intelligence factor measured in short laboratory tasks, with a turn-taking equality term [Woolley et al., 2010]. That is supporting context and not confirmation: it is a group-scale, laboratory, short-task result and not a pair-scale field result.

Table V-P3 — Pair storage `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t`

Gate	Finding
1 Dimensions	`Q` in declared order-units; `S` in nats; **`eta` therefore carries order-units per nat per period** and must be reported with them; `delta` dimensionless per period; `C` in order-units per period. **PASS, provided `eta`'s units are declared.**
2 Scale type	Requires `Q` ratio-scaled with a true zero. For a pair the order proxy — shared method, joint artefacts, documented decisions — must be countable. **PASS where the proxy is a count; FAILS and demotes where it is a satisfaction rating.**
3 Limiting cases	`delta -> 0`: no fixed point, `Q_t <= Q_0 + t.eta.ln N` — the pair that never forgets and never revises [Curzi, 2026a]. `delta -> 1`: total amnesia each period. `eta -> 0`: the pleasant partnership that builds nothing. `C` large: negative `Q_inf`. **PASS — every divergence names a real condition.**
4 Fixed points	`Q_inf = (eta F(u*) − c*)/delta` for `delta > 0`, unique, reached geometrically since the map is a contraction with modulus `1−delta` [Banach, 1922]. **PASS.**

Gate	Finding
5 Identifiability	`delta` from decay under zero input; `eta` from attributable change in `Q` per unit `S`; `C` from tally. **Known hazard: if the order proxy repackages the input, `eta` and the surplus term are not separately identifiable.** Acute for pairs — "we have a great shorthand" is both. **CONDITIONAL PASS; the proxy must be independently justified.**
6 Falsification	Fit the frozen model, report held-out performance against four baselines. **Kill condition: the shuffled-input control matches the fitted model held-out** [Curzi, 2026i].
7 Reproducibility	**UNTESTED.** Whether `(delta, eta)` from one pair transfer to another pair in the same trade has not been run.

Table V-P4 — Reciprocity asymmetry `a = |ln(eta_AB / eta_BA)|`

Gate	Finding
1 Dimensions	Ratio of two identically-united quantities; argument dimensionless; `a` in nats. **PASS.**
2 Scale type	Requires both conversions ratio-scaled with a true zero, that is a direction that converts nothing. Attainable where `eta` is attributable durable output per unit input. **PASS.**
3 Limiting cases	`a -> 0`: perfect reciprocity. `a -> infinity`: one direction converts nothing — the audience relation, the extractive partnership. Both nameable. **PASS.**
4 Fixed points	`a = 0`, which is exactly the reciprocity condition. It is **not** an attractor: nothing in the model drives `a` down, which is why reciprocity needs a protocol. **PASS, fixed point unstable in practice.**
5 Identifiability	The two directional conversion rates are identifiable only if each private store is measured separately. **From the joint store alone they are NOT identifiable** — every row of Section 4.10's table sums to the same total. **PASS ONLY UNDER SEPARATE MEASUREMENT; this is why "the partnership is doing well" is not evidence of reciprocity.**
6 Falsification	Preregister `a` from directional attribution; predict divergence of the two private stores. **Kill condition: `a` fails to predict store divergence beyond chance.**
7 Reproducibility	**UNTESTED.**

Table V-P5 — Return difference `Delta_R = |ln(M'/M)|`

Gate	Finding
1 Dimensions	`M` and `M'` share units by construction, so the ratio is dimensionless and `Delta_R` is in nats. **PASS.**
2 Scale type	**The binding gate.** `M'/M` requires `M` ratio-scaled with a true zero. With only ordinal reports, `Delta_R` demotes to "the reply differed / did not differ" — an ordering, not a magnitude. **CONDITIONAL, and in most pair settings the condition is not met without a magnitude-estimation procedure** [Stevens, 1946].
3 Limiting cases	`M' -> M`: `Delta_R -> 0`, the echo. `M' -> 0`: `Delta_R` diverges, total absorption. `M -> 0`: undefined, correctly, since nothing was sent. **PASS.**
4 Fixed points	`M' = M`: the relation returns what entered it. **This is a null, not an ideal.** The design target is `Delta_R` bounded away from zero, not minimised. **PASS.**
5 Identifiability	Only the ratio is recoverable, never `M` and `M'` separately. Report the ratio with its direction. **PASS WITH DECLARED LIMITATION.**
6 Falsification	Preregister that pairs with persistently low `Delta_R` show lower held-out joint output. **Kill condition: no relation beyond the shuffled control.**
7 Reproducibility	**UNTESTED, and blocked upstream by gate 2.** Until magnitude estimation of `M` is established this gate cannot be attempted. Stated rather than skipped.

Table V-P6 — Non-parallelism decay `u_t = u_0 (1−lambda)^t`

Gate	Finding
1 Dimensions	`lambda` dimensionless per period; `u` dimensionless; `t` a count. **PASS.**
2 Scale type	Ratio scale in `u` per Table V-P1, count scale in `t`. **PASS.**
3 Limiting cases	`lambda -> 0`: `u` constant — a pair that works together and never influences the other, which is the mark of a pair not actually collaborating. `lambda -> 1`: collapse to `u = 0` in one period, total capture of one partner by the other. Both real. **PASS.**
4 Fixed points	`u = 0`, unique and **globally attracting for any `lambda > 0`.** The central engineering fact: convergence is the default and difference is the maintained condition. **PASS.**
5 Identifiability	`lambda = 1 − (u_1/u_0)^{1/dt}` from two clean `u` measurements with no restoration between. **PASS.**

Gate	Finding
6 Falsification	Preregister `lambda` from periods 1–8, predict `u` at period 16, compare to a constant-`u` baseline. **Kill condition: constant-`u` predicts held-out `u` as well as the decay model.**
7 Reproducibility	**UNTESTED.** Geometric decay is the simplest assumption, not a measured law. **The functional form is a modelling choice, badged as one, and its exponent is defined by the model rather than estimated.**

24.4.16 4.16 Scope constraint on use

Interventions at pair scope act on people who did not attend the design meeting. **Nothing here licenses arranging a person's partnerships for them, selecting their associates by measured obliqueness, or scoring a relationship they did not agree to have measured.** These instruments are for a pair to point at itself, or for an institution to point at partnerships it is itself constituting — a review pairing, a co-authorship, a two-person watch. The moment a protocol removes a person's ability to decline the measurement or the pairing, it is outside the framework's scope and stays outside. The corpus's own treatment of this seam is stricter than a permission rule and bounds what may be claimed here: category innocence is refused as a defence of acquisition, moving a harm through a discovered gap in a scheme is forbidden, and the case in which the feedback loop runs through another person's will is conceded to be an unsettled seam the corpus does not close [Curzi, 2026]. The open-design principle applies to the instruments themselves: a protocol should remain defensible when everything about it except the participants' data is public [Kerckhoffs, 1883].

24.5 5 Design requirements and failure modes at pair scope

24.5.1 5.1 Requirements

Each requirement names the parameter it moves.

S-1 Declare the domain and count it. Write down the domain's `N` — the distinguishable categories the pair's work must span — and justify each, before any pairing is assessed. `[N]` *Test: an independent reader lists the same categories from the written definition.*

S-2 Never report `F` without `N` and `u`. Per gate 5 of Table V-P2, `F` alone does not identify the pair. `[reporting]`

S-3 Estimate `u` before the pairing is fixed. Independent pre-discussion estimates on at least eight domain questions; disagreement and error decorrelation. `[u]` *Test: estimates recorded before the two people spoke about the questions.*

S-4 Apply the obliqueness test. If either position is reliably recoverable from the other by copying or by inversion, record `u` near zero and do not assume surplus. `[u]` *Test: predictive accuracy of the copy rule and of the inversion rule, both measured.*

S-5 Reject opposition as a substitute for complementarity. A pair whose disagreement is systematic and total is `rho` near `−1` and `u` near zero. Record it as a similarity failure, not as productive tension. `[u, C]`

S-6 Measure `u_cross` and `u_div` separately and report both; a pair rich in one and starved of the other is a diagnosable condition with a specific fix. `[u]`

S-7 Compute the minimum viable angle. `u_min = (e^{C/eta} − 1)/(N−1)`; every standing pair should know its own. `[Q]` *Test: the number is written down and dated.*

S-8 Measure `delta` and `eta` before promising anything: `delta` from decay under zero input, `eta` from attributable durable output per unit input. `[delta, eta]`

S-9 Justify the order proxy independently of the input. The `Q` proxy must not restate interaction volume. "We talk constantly" is an input measure; "a documented method a third party executed successfully" is an order measure. `[Q]`

S-10 Require an artefact per period. Periods producing no durable object are running at `eta` near zero regardless of how they felt. `[eta]` *Test: name the artefact for each of the last six.*

S-11 Measure `lambda`, the convergence rate. Two clean `u` measurements, six or more periods apart, with no restoration in between: `lambda = 1 - (u_1/u_0)^{1/dt}`. `[u maintenance]`

S-12 Schedule restoration. Interval no longer than the measured half-life of `u`. A restoration is a deliberate importation of a state the pair cannot predict from its own: a third party from another lineage in a working session, a rotation into an unfamiliar problem class, an adversarial review by someone differently trained. **A holiday is not a restoration; it changes `C`, not `u`.** `[u]` *Test: restorations are on the calendar, each naming its source of difference.*

S-13 Measure both private stores, not only the joint one. Per gate 5 of Table V-P4, reciprocity is unidentifiable from the joint account. `[eta, reciprocity]`

S-14 Set a reciprocity threshold and act on it. Recommended trigger: `a > 0.6931`, a two to one directional imbalance, prompts a review of what flows each way. Below that, note and monitor. **The threshold is a design choice and not an empirical finding; no distribution of measured `a` exists.** `[reciprocity]`

S-15 Keep `Delta_R` bounded away from zero. Replies that restate the sender return no information. Where `M` cannot be ratio-scaled, demote to the ordinal test — did the reply contain a position the sender could not have generated? — and record it as ordinal. `[information]`

S-16 Do not add a third person to fix a pair problem. One head triples the channels, from one to three, and raises `C` at once; it raises `u` only if the third is oblique to both. Compute headroom first. `[C, u]`

S-17 Compute headroom before requesting an increment. `h = eta ln N - delta Q - C`; no increment may exceed `h`. `[h]`

S-18 Recognise that headroom falls as the pair succeeds. `delta Q` grows with `Q`, so a long partnership spends more of its budget maintaining what it built. Plan the transition to maintenance rather than being surprised by it. `[h]`

S-19 Shed `Q` deliberately when headroom goes negative. Retiring a shared method or dropping an accumulated obligation are legitimate and underused moves. `[Q, h]`

S-20 Every measurement here is opt-in for both parties. `[ethics]`

24.5.2 5.2 Failure modes

Each names its damaged parameter, its signature, its measurement and its fix.

FM-P1 The mirror pair. Two people selected for similarity — same training, lineage, incentives — whose agreement is read as corroboration. *Damages:* `u -> 0`, therefore `S -> 0`. *Signature:* pre-discussion estimates agree within noise; neither can name a question on which they expect to differ; meetings are short, pleasant and produce no revisions. *Measurement:* copy-rule predictive accuracy above 0.9 on held-out domain questions. *Fix:* import. The pair has no internal source of `u`. A practitioner from an unrelated trade in a working session, on a schedule, not as a one-off.

FM-P2 The opposition pair. Systematic mutual contradiction mistaken for rigour, `rho` near `-1`. *Damages:* `u -> 0`, `C` up. *Signature:* each position is predictable by inverting the other's; conflict constant, positions stable, nothing moved in a year. *Measurement:* inversion-rule predictive accuracy above 0.9. *Fix:* as FM-P1 — the same fault with a sign flip — plus cut `C`: the heat is dissipation and buys nothing.

FM-P3 The converged pair. A pair that began oblique and became parallel through successful collaboration. The commonest failure in long partnerships and the least likely to be

diagnosed, because every step that produced it was correct. *Damages:* `u`, via `lambda`. *Signature:* `Q` peaks then declines while headroom stays positive; fatigue reported; the decline attributed to market conditions or workload. *Measurement:* two `u` measurements six or more periods apart; a halving is decisive. *Fix:* restoration at or below the measured half-life. The first restoration recovers most of what is available.

FM-P4 The one-way channel. Conversion flows overwhelmingly one way; one private store grows, the other reaches a negative equilibrium. *Damages:* `eta` directionally; `Q` for one member. *Signature:* the joint account looks healthy; one partner reports growth, the other depletion, and both reports are accurate. *Measurement:* separate private stores; `a = |ln(eta_AB/eta_BA)|`. **Not detectable from the joint account.** *Fix:* reallocate what flows each way. Where the asymmetry is structural — a genuine mentorship — say so and set an end date rather than running it as a partnership.

FM-P5 The pleasant void. High `u`, high `S`, `eta` near zero. Genuinely different people who enjoy each other and build nothing. *Damages:* `eta`. *Signature:* "we had a great conversation"; no artefact for six periods; `Q` flat or declining despite strong measured `u`. *Measurement:* artefact count; `eta` as attributable change in `Q` per unit input. *Fix:* capture mechanisms at pair scale — decision log, written summary, changed default, built object. **Nearly always the cheapest intervention available, and it should be tried before anything is done to `u`.**

FM-P6 The echo partner. One member returns the other's own signal, `Delta_R` near zero. *Damages:* information return; downstream, `u`. *Signature:* every reply is a restatement, elaboration or endorsement; nothing comes back that the sender could not have generated. *Measurement:* `Delta_R` where `M` is ratio-scaled; the ordinal demotion of S-15 otherwise. *Fix:* structural, not personal — independent estimate before discussion at pair scale, both write first and then compare. Defeats the echo mechanically, with no confrontation.

FM-P7 The sub-angle pair. A pair below its own `u_min`: the steady state is negative and the store erodes at rate `delta` whatever anyone does. *Damages:* `Q`, terminally. *Signature:* everything built two years ago is degrading; effort produces maintenance, never gain. *Measurement:* `u_min` against measured `u`. *Fix:* raise `u` by import, raise `eta` by capture, or cut `C`. Raising `N` is logarithmic and the worst buy. **If none is available, no arrangement of these two people prevents decline.**

FM-P8 The third-person reflex. Adding a third member to fix a pair deficit. Channels go from one to three, `C` rises at once, `u` rises only if the newcomer is oblique to both. *Damages:* `C`, and often `u`, if the third is a lineage-mate of either. *Signature:* headroom falls after the addition; output does not rise. *Measurement:* `h` before and after. *Fix:* prevention — measure the newcomer's `u` against **both** members first. A third oblique to one and parallel to the other buys less than the channel count costs.

FM-P9 The unjustified proxy. The order measure repackages the input measure — "we have great shorthand" counted as `Q`, when shorthand is downstream of interaction volume. *Damages:* measurement validity; `eta` unidentifiable. *Signature:* `eta` near 1 with little residual; excellent training fit, held-out failure. *Measurement:* is the order proxy derivable from the input without anything outside the pair? *Fix:* an order proxy with an external referent — work a third party can execute, a decision whose outcome is later observable, a physical artefact.

FM-P10 The unbadged identification. Treating the dyadic identification as though it were the extremality theorem. *Damages:* the argument's warrant. *Signature:* "the mathematics proves this partnership cannot work." *Measurement:* read the sentence — does it assert something about vectors or about people? *Fix:* restate with the badge on. The theorem is exact about states; the identification of people with states is this framework's, is testable, and has not passed the reproducibility gate.

24.6 6 The group moment

The framework's second equation is the moment equation, given at the scale of one person as `M = (D/P).E`, with `D` the frame a mind brings, `P` the intake it achieves, `E` the environment, and `M`

the moment as lived. This section takes the same equation, keeps every symbol, changes only the referent of each, and asks whether it still says anything.

It does, and the reason is not rhetorical: **the group version acquires a term the personal version does not have, and that term is the interaction surplus.** `D` becomes shared doctrine — the standing model an organisation carries about the world it operates in. `P` becomes collective perception: the intelligence the organisation actually gathers, audited for redundancy. `E` is the environment. `M` is the situation as the group constructs it, which is never the environment itself.

24.6.1 6.1 The equation at group scope

```
M = ( D / P ) . E

D  shared doctrine        the standing model: what the organisation already asserts
P  collective perception  intelligence gathered, audited for redundancy
E  environment            the world the organisation is in
M  constructed situation  the situation the organisation acts on

r     = D / P             the doctrine-to-data ratio (dimensionless)
Delta = | ln r |           distortion, in nats

ln M = ln D - ln P + ln E     the log form
```

Status. `M = (D/P).E` is a **postulate**, this framework's own, stated in the open. It is not a theorem of the framework's mathematical papers and nothing here derives it. It is a proposed structure whose consequences are computable and, by the protocol of Section 6.12, falsifiable. Everything downstream that is proven comes from the surplus functional or the storage law and is labelled as such; everything else inherits the postulate's status. The framework's own audit reaches the same verdict independently and at both scopes [Curzi, 2026].

24.6.2 6.2 The commensurate-units convention

Dimensional balance requires `D` and `P` to share units so that `r` is dimensionless. At personal scope this is handled by magnitude estimation. At group scope there is a cleaner route, because organisations already count the thing: the convention of Section 3.3 expresses `D`, `P` and `E` in effective independent evidence-channels per decision period. All three are then counts with a true zero, hence ratio-scaled, and `M` inherits `E`'s units.

This convention is doing real work and it is the weakest joint in the construction. It is badged, and Section 6.12 shows exactly where it fails.

24.6.3 6.3 The aggregation postulate and the exact join

`P` is not headcount. A hundred analysts reading the same wire service are one channel. That is the framework's effective count [Curzi, 2026c], and it is the entire reason a surplus law belongs in an account of perception. Under the aggregation postulate, `P = p_bar . N_eff` with `N_eff = 1 + (N-1)u`. Taking logarithms and substituting:

```
ln P = ln p_bar + ln N_eff = ln p_bar + F        since F = ln N_eff exactly

ln r = ln D - ln p_bar - F                       THE JOIN
```

The surplus of the group appears in the doctrine-to-data ratio as a subtractive term, one nat for one nat. Every nat of interaction surplus the organisation's channels generate reduces

its log-distortion by exactly one nat. That is not an analogy between the two equations; it is an identity, and it holds because `F` is defined as `ln N_eff` and the logarithm is the register both formulas live in.

Two consequences follow immediately and both are hard.

Consequence 1 — the correction ceiling. Since `0 <= F <= ln N`, the total reduction in `ln r` available from interaction geometry alone is at most `ln N`:

```
N =  4  ->  max reduction 1.3863 nats   ->  r divided by at most  4
N =  8  ->  max reduction 2.0794 nats   ->  r divided by at most  8
N = 12  ->  max reduction 2.4849 nats   ->  r divided by at most 12
N = 24  ->  max reduction 3.1781 nats   ->  r divided by at most 24
N = 40  ->  max reduction 3.6889 nats   ->  r divided by at most 40
```

Consequence 2 — the infeasibility condition. Setting `r = 1` requires `N_eff = D/p_bar`, and `N_eff <= N`. Therefore:

> **If `D/p_bar > N`, the organisation cannot be calibrated by any arrangement of its people.** The residual distortion is at least `Delta_min = ln(D/p_bar) − ln N`, and no recruitment, rotation, red team or restructuring closes it. The only remaining move is to reduce `D`.

This is the exact analogue, in the perceptual register, of the storage law's statement that if the cost exceeds `eta ln N − delta Q_t` no interaction geometry prevents decline [Curzi, 2026a]. **Both are `ln N` ceilings, and both say the same thing: geometry is bounded, and past the bound what is carried must change, not who is carrying it.** The general form of the objection is the requisite-variety bound: a regulator cannot correct more variety than it can itself express [Ashby, 1958].

24.6.4 6.4 The contact postulate

The moment must serve as the storage law's input, `Q_{t+1} = (1−delta)Q_t + eta M_t − C_t`, with `M` normalised to `[0, ln N]`. The normalisation is stated rather than assumed: the contact postulate of Section 3.3 sets `S_t = ln N . e^{−Delta_t}`.

That exponential has an exact and more legible form. Since `e^{−|ln r|} = min(r, 1/r)` identically — an algebraic identity, verifiable at `r = 0.1, 0.25, 0.5, 0.8, 1.0, 1.25, 2.0, 4.0, 10.0` to machine precision:

```
S_t / ln N  =  min(D, P) / max(D, P)
```

A group banks at the rate of the scarcer of frame and intake. Doctrine that cannot be fed with evidence and evidence that cannot be framed are equally useless, and the binding one is whichever is smaller. This is a limiting-factor reading of the ratio and it is offered as an **analogy**: the framework has no citation for it and none is manufactured here. The attenuation is symmetric, which is the postulate's substantive commitment and the first thing Section 6.12's falsification test attacks.

The band table, distortion to attenuation:

Delta	r (either side)	S / ln N
0.00	1.000	1.0000
0.05	1.051 or 0.951	0.9512
0.10	1.105 or 0.905	0.9048
0.20	1.221 or 0.819	0.8187
0.35	1.419 or 0.705	0.7047

```
0.50    1.649 or 0.607      0.6065
0.69    1.994 or 0.502      0.5016
1.00    2.718 or 0.368      0.3679
1.39    4.015 or 0.249      0.2491
2.00    7.389 or 0.135      0.1353
```

A two-to-one doctrine-to-data ratio halves what the organisation can bank.

24.6.5 6.5 The four limiting cases

Every limit corresponds to a real, nameable phenomenon, and at group scope they are more recognisable than at personal scope because organisations fail publicly.

Case 1 — `P -> 0`: doctrine with no intelligence, `M` diverges. At `D = 1.0`, `E = 1.0`:

```
P        r          M        Delta
-----    -------    ------   ------
1.000      1.00      1.00    0.0000
0.500      2.00      2.00    0.6931
0.200      5.00      5.00    1.6094
0.100     10.00     10.00    2.3026
0.050     20.00     20.00    2.9957
0.020     50.00     50.00    3.9120
0.010    100.00    100.00    4.6052
0.001   1000.00   1000.00    6.9078
```

The divergence is real and the phenomenon has a name in every language: **an organisation with a model and no collection is not blind, it is certain.** The constructed situation grows without bound while the environment stays where it is. Note what the arithmetic says about the *feel* of the condition. `M` is large. The organisation experiences a rich, detailed, confident picture. **Nothing in the subjective quality of `M` distinguishes a well-informed organisation from a hallucinating one; only the ratio does.** That is the single most important reason the ratio has to be measured rather than sensed, and it is the same structural point that makes an over-framed body's self-assessment improve as its condition worsens [Vaughan, 1996].

Note also that `P -> 0` does not require anyone to stop collecting. Under the aggregation postulate `N_eff -> 1` as `u -> 0`, so **an organisation whose channels become parallel has `P` collapse to a single channel's worth while its collection budget stays flat.** That is Section 7's mechanism, and it belongs here because it belongs to the limit.

Case 2 — `D -> 0`: intelligence with no frame, `M -> 0`. At `P = 1.0`, `E = 1.0`:

```
D        r          M        Delta
-----    ------     ------   ------
1.000    1.0000     1.0000   0.0000
0.500    0.5000     0.5000   0.6931
0.200    0.2000     0.2000   1.6094
0.100    0.1000     0.1000   2.3026
0.050    0.0500     0.0500   2.9957
0.010    0.0100     0.0100   4.6052
0.001    0.0010     0.0010   6.9078
```

The organisation constructs no situation at all: **sensation is not experience.** At group scope, an organisation holding an enormous quantity of data and no model of what it is for produces no decision. The data lake with no question. The dashboard nobody can act on. The limit is genuinely `0` and not merely small, and that is worth stating: a group with no doctrine does not act badly, **it does not act**. The paralysis of the perfectly open-minded institution is a prediction of the formula, not a moral observation.

Case 3 — `E -> 0`: no environment, `M -> 0` for any `r`. The organisation's environment has gone: its market has vanished, its constituency has dissolved, its problem has been solved or has migrated. **`M = 0` regardless of how good its doctrine and intelligence are**, because a ratio multiplied by nothing is nothing. This limit is often mistaken for the others and the distinction is diagnostic. An organisation at `E -> 0` has excellent internal metrics and produces nothing that lands. Its doctrine may be perfectly calibrated to its intake, with distortion zero, and it will still generate no moment. **Calibration is not sufficiency. A perfectly calibrated instrument pointed at nothing reads nothing.**

Case 4 — `D = P`: `M = E`, the fixed point.

```
r = 1  ->  Delta = 0  ->  M = E  ->  S = ln N
```

When doctrine matches intelligence, the situation the organisation constructs equals the environment it is in. That is undistorted contact, stated as an equation.

24.6.6 6.6 The zero-headroom identity

There is a structural result at that fixed point which does not appear at personal scope, and it is the most surprising thing in this section. Put `r = 1` into the contact postulate: `S = ln N`, the maximum. The storage law's steady state is then `Q_inf = (eta ln N − C)/delta`, and headroom at that steady state is

```
h = eta ln N − delta Q_inf − C
  = eta ln N − delta (eta ln N − C)/delta − C
  = eta ln N − eta ln N + C − C
  = 0        exactly
```

> **A perfectly calibrated organisation converges to exactly zero headroom.** Not by coincidence and not by mismanagement: by arithmetic. Perfect contact drives the input to the ceiling, the ceiling drives the store to its maximum, and the maximum store consumes the entire budget in maintenance.

This is the framework's success-consumes-its-own-headroom thesis arriving from the perceptual side, and it is the organisational-slack result stated as an identity [Cyert and March, 1963]. It warns against a naive reading of `r = 1` as an unqualified goal. Calibration maximises what is banked; it does not leave room to build something new with it. **An organisation that wants headroom at equilibrium must run below its maximum store, which means either shedding `Q` deliberately or accepting `r` not equal to 1.** The trade is stated rather than pretended away, and it is the same shape as the queueing result that reserve, not prudence, is what keeps a response time finite [Kingman, 1962]. A structurally identical bind is recorded elsewhere in this corpus and bounds how far the identity may be read as bad news: the property that makes an insurance reserve resistant to third-party claims — real liabilities owed to unrelated parties — is the same property that makes it unavailable to its owner, leaving only 7.58 per cent realisable within a year [Curzi, 2026]. **A store that is maximal is not thereby a store that is available**, and that is a general feature of reserves rather than a peculiarity of calibrated organisations.

24.6.7 6.7 The calibration requirement, computed

Solving `r = 1` for the non-parallelism required: `N_eff = D/p_bar`, so

```
u* = ( D/p_bar − 1 ) / ( N − 1 )        feasible only if D/p_bar <= N
```

The full table, across domain sizes and doctrine loads:

D/p_bar	N=4	N=8	N=12	N=16	N=24	N=40
1.5	0.1667	0.0714	0.0455	0.0333	0.0217	0.0128
2.0	0.3333	0.1429	0.0909	0.0667	0.0435	0.0256
3.0	0.6667	0.2857	0.1818	0.1333	0.0870	0.0513
4.0	1.0000	0.4286	0.2727	0.2000	0.1304	0.0769
6.0	--	0.7143	0.4545	0.3333	0.2174	0.1282
8.0	--	1.0000	0.6364	0.4667	0.3043	0.1795
12.0	--	--	1.0000	0.7333	0.4783	0.2821
16.0	--	--	--	1.0000	0.6522	0.3846
24.0	--	--	--	--	1.0000	0.5897
40.0	--	--	--	--	--	1.0000

"--" = INFEASIBLE: D/p_bar exceeds N; no geometry calibrates this organisation.

Three readings, each an engineering rule.

Doctrine is expensive in `u`. An organisation whose standing model supplies six channels' worth of constraint, working an `N = 12` domain, needs `u = 0.4545` — mean pairwise non-parallelism of 0.45 across its intelligence channels, `theta = 42.4` degrees. That is a high bar, and organisations measured by channel-audit protocols typically come in an order of magnitude below it. **That comparison is this framework's own impression and not a measured distribution; no empirical distribution of institutional `u` exists.**

A large `N` is a doctrine subsidy. The same `D/p_bar = 6.0` needs `u = 0.4545` at `N = 12` and only `u = 0.1282` at `N = 40`. Breadth of domain buys tolerance for doctrine. But `N` moves the ceiling only logarithmically and is the most expensive lever there is. **Raising `N` from 12 to 40 to relieve a doctrine problem is a real move and a bad buy.**

The infeasible region is large and it is where heavy institutions live. A long-established body carrying `D/p_bar = 16` in an `N = 12` domain is in the infeasible zone. Its residual distortion floor is:

```
Delta_min = ln(D/p_bar) - ln N

N=12, D/p_bar=18  ->  Delta_min = 0.4055        S <= 0.6667 ln N
N=12, D/p_bar=24  ->  Delta_min = 0.6931        S <= 0.5000 ln N
N=12, D/p_bar=48  ->  Delta_min = 1.3863        S <= 0.2500 ln N
N=24, D/p_bar=48  ->  Delta_min = 0.6931        S <= 0.5000 ln N
N=24, D/p_bar=96  ->  Delta_min = 1.3863        S <= 0.2500 ln N
```

An institution carrying four times its domain size in doctrine cannot bank more than a quarter of its ceiling, no matter who it hires. The only lever left is `D`.

24.6.8 6.8 Worked mechanics of the join

Fix `D = 6.0`, `p_bar = 1.0`, `N = 12`, and let the group's own geometry vary. This is the join running end to end: `u` sets `F`, `F` sets `ln r`, `ln r` sets `Delta`, `Delta` sets `S`.

u	F	ln r	r	Delta	M/E
0.0000	0.0000	+1.7918	6.0000	1.7918	6.0000
0.0200	0.1989	+1.5929	4.9180	1.5929	4.9180
0.0500	0.4383	+1.3535	3.8710	1.3535	3.8710
0.1000	0.7419	+1.0498	2.8571	1.0498	2.8571
0.2000	1.1632	+0.6286	1.8750	0.6286	1.8750
0.3000	1.4586	+0.3331	1.3953	0.3331	1.3953

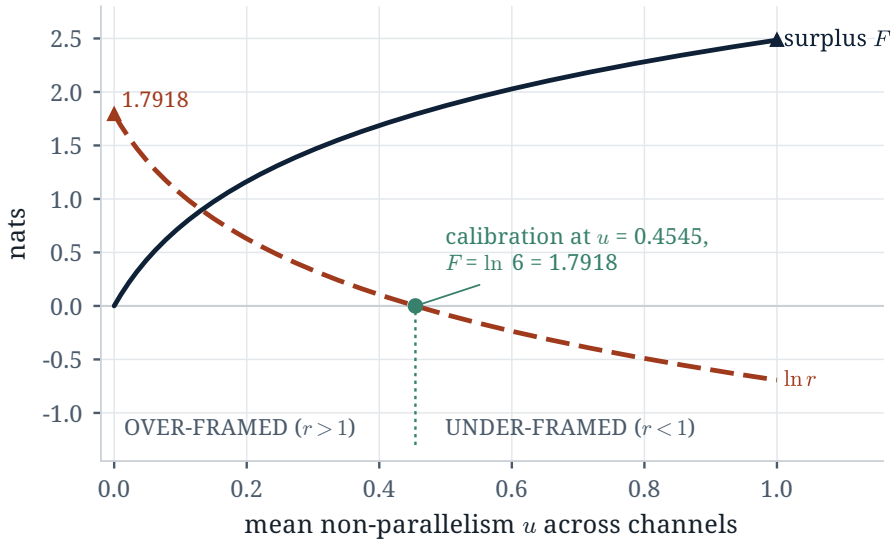


Figure 24.2: **The group's interaction surplus subtracts from its log-distortion one nat for one nat**: at `D = 6.0`, `p_bar = 1.0` and `N = 12` the two curves are mirror images about the crossing at `u = 0.4545`, where `F = 1.7918 = ln 6` and `ln r = 0.0000`. Past the crossing the organisation over-corrects into under-framing, reaching `ln r = −0.6931` at perfectly orthogonal channels. Computed from the identity `ln r = ln D − ln p_bar − F` at the parameters tabulated in Section 6.8; generated by `figs/src/CCP-024_fig02.py`.

0.4167	1.7198	+0.0720	1.0746	0.0720	1.0746
0.4545	1.7918	0.0000	1.0000	0.0000	1.0000
0.5000	1.8718	−0.0800	0.9231	0.0800	0.9231
0.7000	2.1633	−0.3716	0.6897	0.3716	0.6897
0.9000	2.3888	−0.5970	0.5505	0.5970	0.5505
1.0000	2.4849	−0.6931	0.5000	0.6931	0.5000

Read the `ln r` column against the `F` column. **They move one for one, in opposite directions, exactly:** `ln r = ln 6 − F = 1.7918 − F` on every row. The surplus of the group *is* the de-framing of the group, in the same units, with no conversion constant.

Read the last four rows for the other half of the story. **Past `u = 0.4545` the organisation over-corrects into under-framing.** At `u = 1.0` — perfectly orthogonal channels — the same doctrine load produces `r = 0.5` and `M/E = 0.5`: the organisation now under-constructs the situation. **Diversity is not monotonically good in this model.** It reduces `ln r`, and reducing `ln r` is good only while `r > 1`. An organisation already under-framed and adding channels is making its condition worse, and the formula says so plainly. This is the sense in which the two pathologies of Section 7 require *opposite* corrections.

24.6.9 6.9 The moment as storage input, and the operating envelope

Take a group with `eta = 0.42` order-units per nat, `delta = 0.14` per period, `C = 0.55` order-units per period, `N = 12` so `ln N = 2.4849`. The minimum input for a non-negative equilibrium is `S_min = C/eta = 1.3095` nats.

S	eta.S	Q_inf
0.0000	0.0000	−3.9286
0.3000	0.1260	−3.0286
0.6000	0.2520	−2.1286

```
0.9000   0.3780   −1.2286
1.2000   0.5040   −0.3286
1.5000   0.6300   +0.5714
1.8000   0.7560   +1.4714
2.1000   0.8820   +2.3714
2.4849   1.0437   +3.5261
```

Converting distortion to input by the contact postulate and reading off the distortion at which this group stops being able to hold anything at all:

```
S_min / ln N = 1.3095 / 2.4849 = 0.5270
min(r, 1/r) >= 0.5270   ->   0.527 <= r <= 1.898   ->   Delta <= 0.6406
```

> **This organisation must keep its doctrine-to-data ratio between 0.53 and 1.90. Outside that band its steady-state store is negative and its institution erodes at rate `delta` regardless of effort.**

That is a concrete, computable operating envelope derived from three stipulated parameters and a badged postulate. **It is demoted to an illustrative construction at the scale-type gate in Section 6.12, and it must not be quoted as a measurement.**

24.6.10 6.10 The two-track ledger

The envelope above is stated on `Q`, and `Q` is the wrong ledger to optimise. The framework's established result is that **over-framing increases stored order**, and the result transfers to group scope without modification, because the storage law does not know what scale it is running at:

```
ALL STORED ORDER   Q_{t+1}   = (1−delta) Q_t   + eta M_t − C_t     M_t = min(r_t E_t, ln N)
VERIDICAL ORDER    Q^V_{t+1} = (1−delta) Q^V_t + eta V_t − C_t     V_t = E_t . min(r_t, 1)

STORED ERROR       Z_{t+1}   = (1−delta) Z_t + eta (M_t − V_t)     [the cost term cancels]
CORRECTION LOAD    C_t       = c_0 + kappa Z_t
```

An over-framed institution accumulates more stored order than a calibrated one. It is simply not true stored order. The doctrine, the procedure manual, the settled interpretation, the confident forecast — all of it banks at rate `eta` whether or not the environment supplied it, because there is no term in the recursion that asks. **Dogma banks faster than knowledge, and the bill arrives on a lag, as correction load rising with accumulated error.** That lag is why an over-framed body looks excellent on its own metrics for years, and it is why organisational learning that homogenises looks like improvement throughout [March, 1991].

This sharpens the zero-headroom result into something less comfortable. Section 6.6 showed that `r = 1` is not an unqualified goal because it converges to zero headroom. The two-track ledger shows that a body optimising raw `Q` is not merely failing to aim at `r = 1` — **it is being actively driven away from it.** Above the fixed point the veridical track is flat at `E` while the raw track keeps climbing. Differentiating the steady state for `r > 1`:

```
Q_inf = (eta M − c_0 − kappa eta (M − V)/delta) / delta
dQ_inf/dM = eta (1 − kappa/delta) / delta
```

The knife edge is exactly `kappa = delta`. Where correction load accrues more slowly than stored order decays, raw `Q_inf` rises without bound in `r`, and every incentive attached to raw `Q` — headcount, published procedure, doctrinal coverage, confident forecasts issued — is an incentive to over-frame. That is not a claim about anyone's character. It is the arithmetic of what the organisation is being paid in, and it is the standard failure of a measure that has become a target [Goodhart, 1975, Campbell, 1979].

Requirement. Report `Z` alongside `Q`, or report `Q^V` instead of `Q`. A body that reports only its total store is reporting a number that rises fastest when it is most wrong.

24.6.11 6.11 Requirements at group scope

G-1 Declare the decision space and count `N`. Enumerate the distinct domains the organisation's intelligence must span, and justify each as genuinely distinct rather than relabelled. `[N]` *Test: an independent reader reproduces the list from the written definition.*

G-2 Adopt the commensurate-units convention explicitly. Write down, before measuring, what one effective evidence-channel per period means in this organisation. `[dimensions]`

G-3 Audit the channels for redundancy before counting them. Headcount and subscription count are not `P`. Compute `N_eff = 1 + (N−1)u` with `u` from error decorrelation on questions with later-observable answers. `[P, u]`

G-4 Measure `p_bar` separately from `N_eff`. `P` alone does not identify them, and the two have entirely different fixes. `[P]`

G-5 Enumerate the doctrine. List the decision-relevant propositions the standing model fixes in advance; weight each by the fraction of decisions it constrains; sum. That is `D`. `[D]` *Test: the list is written, dated, and each entry names the decisions it governs.*

G-6 Run the over-determination check. Compute `ln r` twice — once as `ln D − ln P`, once as `ln D − ln p_bar − F` — and require agreement. A mismatch means one of the three measurements is wrong; find out which before proceeding. `[identification]`

G-7 Never report `Delta` without the sign of `ln r`. The two pathologies map to the same magnitude and require opposite corrections. `[reporting]`

G-8 Never report `D` without an independent measurement of `P`. `D` is unidentifiable without the anchor. `[reporting]`

G-9 Compute the infeasibility check first. If `D/p_bar > N`, stop: no arrangement of people calibrates this organisation. Record the residual floor `Delta_min = ln(D/p_bar) − ln N` and proceed to doctrine reduction. `[D, N]`

G-10 Compute the operating envelope. From `eta`, `C` and `N`, derive `S_min = C/eta` and the band `min(r, 1/r) >= S_min/ln N`. `[operating limits]` *Test: the two band edges are written down and dated.*

G-11 Treat `E` as measurable and measure it. An organisation that never estimates the decision-relevant structure in its environment cannot distinguish Case 3 from Cases 1 and 2. `[E]`

G-12 Do not treat `r = 1` as an unqualified target. Exact calibration drives equilibrium headroom to zero. Decide deliberately whether the organisation is optimising banked order or preserving capacity for the new. `[h]`

G-13 Justify the order proxy independently of the input. The organisation's `Q` measure must not be a repackaging of its intelligence volume. `[Q]`

G-14 Preregister before intervening. Encoder, fit window, freeze date, held-out window, four baselines including the shuffled-input control, and the kill conditions [Nosek et al., 2018, Curzi, 2026i]. `[falsifiability]`

G-15 Re-measure `u` on a schedule. `u` decays. A calibrated organisation drifts into over-framing with no change in doctrine. Interval no longer than the measured half-life of `u`. `[u]`

G-16 No protocol here may remove a person's ability to decline — to decline the measurement, to decline the channel role, or to decline the dissent it asks of them. `[ethics]`

24.6.12 6.12 Validation gates, and the demotion of the magnitudes

One verdict governs three of the five tables below and is stated before them rather than buried inside them. Running the framework's formula-validation protocol on the moment equation returns **DEMOTE at gate 2**, and the verdict carries to group scope [Curzi, 2026]. Section 6.2

argued that channel-counting gives `D` and `P` a true zero; that argument is here **corrected**. A count supplies a zero and a unit *within one declared encoder*, but nothing links two encoders, and without linkage the magnitude of `r` moves when the encoder moves. An organisation's channel count is more auditable than a person's frame, which is a real advantage at the identifiability gate; it is not the same thing as a scale.

> **NOT REPORTABLE AS MEASURED QUANTITIES, at group scope as at personal scope:** the magnitude of `D/P`; the distortion `Delta = |ln r|` as a number; and the quantitative form of the join from the moment into the storage law. Every figure in Sections 6.7, 6.8 and 6.9 — the `u*` table, the `Delta_min` floors, the 0.53 to 1.90 envelope — is an **illustrative construction demonstrating the structure of the model**, and is to be read as such.
>
> **STILL IN FORCE:** the four limiting cases; the `D = P` fixed point and its zero-headroom consequence; the ordering over-framed / calibrated / under-framed; the *sign* of `ln r`; and the algebraic identity `ln r = ln D − ln p_bar − F`, which is an identity and does not depend on any scale claim. **What is demoted is magnitude, not structure.**

The route out is a Rasch measurement on one common linked scale, under which `Delta = |ln(D/P)| = |theta_D − theta_P|` becomes genuinely ratio-scaled, because Rasch separation supplies the true zero that division requires [Rasch, 1960]. The framework prices that instrument at **approximately 123 items per term for a half-logit distortion at 95 per cent confidence** [Curzi, 2026]. At institutional scope the items are decisions with later-observable answers, which is the one respect in which the group version is genuinely cheaper: **an organisation with a decision log is already accumulating the item bank.**

Table V-G1 — The group moment `M = (D/P).E`

Gate	Finding
1 Dimensions	Under the commensurate-units convention, `D`, `P`, `E` are all in effective evidence-channels per period. `D/P` dimensionless; `M` inherits `E`'s units. **PASS, and it passes ONLY because the convention was declared. Without it the formula is dimensionally void.**
2 Scale type	Division requires ratio scale with a true zero on a **common** scale. Channel-counts supply a zero and a unit within one declared encoder but supply no linkage between encoders, so the magnitude of `r` is encoder-dependent. **DEMOTE: the magnitude of `D/P` and of `Delta` is withdrawn as a measured quantity; the ordering, the limits, the fixed point and the sign of `ln r` stand.** Where an implementation substitutes rating scales for counts the position is worse: division is impermissible outright. **Route out: Rasch logit difference, about 123 items per term** [Curzi, 2026].
3 Limiting cases	All four worked in Section 6.5. `P –> 0`: the certain organisation. `D –> 0`: the paralysed one. `E –> 0`: the obsolete one. `D = P`: `M = E`. Every divergence names a real, observable institutional condition. **PASS.**

Gate	Finding
4 Fixed points	`r = 1` implies `M = E`, undistorted contact, unique. Plus the derived result of Section 6.6: at that fixed point `h = 0` exactly. **PASS, with the trade stated.**
5 Identifiability	**THIS GATE FAILS AS THE FORMULA IS NAIVELY WRITTEN, and the failure must be printed.** From observations of `M` and `E` alone, only the ratio `r` is recoverable, never `D` and `P` separately. Demonstration: `(D,P) = (6.0, 3.0)`, `(2.0, 1.0)`, `(0.5, 0.25)` and `(120.0, 60.0)` all give `r = 2.0000` and, at `E = 1`, identical `M = 2.0000`. **Four distinct parameter sets, one observation.** *The repair:* an independent anchor on `P` — a direct count of audited channels and their intake — makes `D` recoverable as `D = r.P`. **CONDITIONAL PASS, and only with the anchor.**
6 Falsification	Preregister the encoder for `D`, `P`, `E`; fit on training periods; freeze; report held-out. Four baselines: persistence, direct-input (`E` alone), matched autoregressive, and the shuffled-input control. **Kill conditions: (a) `E` alone predicts held-out decision quality as well as `(D/P).E`; (b) the shuffled-`r` control matches the fitted model; (c) the sign of the residual is independent of the sign of `ln r`.**
7 Reproducibility	The cross-institution test: same operationalisation, different organisations, different analysts. **UNTESTED. Nobody has run it. This gate decides whether this section is science, and stating that plainly is a requirement of the protocol.**

Table V-G2 — Distortion `Delta = |ln r|, r = D/P`

Gate	Finding
1 Dimensions	`r` dimensionless by V-G1; `Delta` in nats, commensurate with `F`. **PASS.**
2 Scale type	`Delta` *would be* a magnitude on a ratio scale — **and this is corrected here: that property is inherited from `r`, whose magnitude V-G1 demotes.** Pending instrumentation `Delta` is **ORDINAL**: it ranks bodies and periods by distance from contact and does not license the statement that one distortion is twice another. **DEMOTE.** Becomes ratio-scaled exactly when read as a Rasch logit difference and by nothing cheaper.

Gate	Finding
3 Limiting cases	`r -> 0`: `Delta` diverges, total under-framing. `r -> infinity`: `Delta` diverges, total over-framing. `r = 1`: `Delta = 0`. The two divergences are **distinct pathologies mapped to the same magnitude**, which is the measure's one real weakness. **PASS ONLY IF THE SIGN OF `ln r` IS REPORTED ALONGSIDE. Never report `Delta` bare.**
4 Fixed points	`Delta = 0` if and only if `r = 1` if and only if `D = P`. Unique. **PASS.**
5 Identifiability	`Delta` is recoverable wherever `r` is, and is more robustly recoverable, since it survives any rescaling that multiplies `D` and `P` by a common factor. **PASS — and this robustness is why `Delta` is the reportable quantity and `D` and `P` are not.**
6 Falsification	Preregister that decision quality declines monotonically in `Delta` on both sides of 1. **Kill condition: no monotone relation, or a relation present on one side and absent on the other** — the latter falsifies the symmetry of the contact postulate specifically, which is a useful and separable result.
7 Reproducibility	**UNTESTED.**

Table V-G3 — Aggregation `P = p_bar . N_eff`, `N_eff = 1 + (N−1)u`

Gate	Finding
1 Dimensions	`p_bar` in intake-units per channel per period; `N_eff` a pure count; `P` in intake-units per period. Matches `D`'s units by the convention. **PASS.**
2 Scale type	`N` a count; `u` ratio-scaled per Table V-P1, conditional on the channel encoder; `p_bar` a mean of ratio-scaled intakes. **PASS, conditional on the encoder being declared.**
3 Limiting cases	`u -> 0`: `N_eff -> 1`, `P -> p_bar` — **a hundred parallel channels are one channel.** `u -> 1`: `P -> p_bar . N`, the maximum. `p_bar -> 0`: `P -> 0`, Case 1. `N -> 1`: `P = p_bar` for all `u`, correct. **PASS; each limit is a real institutional condition.**
4 Fixed points	`u = 0` is the invariant of the parallel organisation and is **globally attracting under the decay of Table V-P6** — `u` is maintained, not owned. `u = 1` is the invariant of full independence and is not attracting at all. **PASS, with the asymmetry noted: the bad fixed point is the stable one.**

Gate	Finding
5 Identifiability	`N` by enumerating declared intelligence domains; `u` by error decorrelation across channels on questions with later-observable answers; `p_bar` by direct intake tally. All three separately obtainable. **But `P` alone does not identify `(p_bar, N_eff)`** — a well-read single channel and a barely-read diverse set give the same `P`. **PASS IF ALL THREE ARE REPORTED.** The reflection problem of Section 3.4 bites here and is not answered by reporting all three.
6 Falsification	Preregister that held-out forecast accuracy tracks `N_eff` rather than channel headcount or channel budget. **Kill condition: headcount or budget predicts held-out accuracy better than `N_eff`.** This is a clean, cheap test and it is the one an organisation with a forecasting record can run first.
7 Reproducibility	**UNTESTED at institutional scale.** The nearest published evidence that group performance tracks something other than member ability is laboratory, short-task and small-group [Woolley et al., 2010]. Supporting context; not confirmation. The closest empirical analogue of an effective count at scale measures adoption against the number of *distinct connected components* among a person's contacts [Ugander et al., 2012], and it is a different quantity.

Table V-G4 — Contact `S = ln N . e^{-Delta} = ln N . min(r, 1/r)`

Gate	Finding
1 Dimensions	`Delta` in nats, `e^{-Delta}` dimensionless in `(0,1]`, `ln N` in nats, `S` in nats. Matches the storage law's input requirement exactly. **PASS.**
2 Scale type	Requires `r` ratio-scaled, which V-G1 demotes. `S` is in nats and would be ratio-scaled with a true zero, but it is computed from `r` and can be no better than its input. **DEMOTE: the postulate's ORDERING — contact banks more than distortion, in either direction, monotonically — stands on ordinal information alone. The banked-nats figures of Section 6.9 do not.**

Gate	Finding
3 Limiting cases	`Delta -> 0`: `S -> ln N`, the ceiling, achieved only at exact calibration. `Delta -> infinity`: `S -> 0`, the organisation banks nothing — true in both directions, the certain institution and the flooded one alike. Bounded above by `ln N`, so the postulate is **consistent with the proven range and does not smuggle in surplus above the ceiling. PASS.**
4 Fixed points	`S = ln N` at `r = 1`, and the derived consequence `h = 0` at the resulting `Q_inf`. **PASS.**
5 Identifiability	`S` is a deterministic function of `r` and `N`, both independently obtainable. **No free parameters — which is a strength: the postulate has nothing to tune and therefore nothing to overfit. PASS.**
6 Falsification	Two separable kill conditions. **(a) Functional form: fit `S = ln N . e^{-k.Delta}` and test `k = 1`; a fitted `k` significantly different from 1 falsifies this postulate while leaving the framework standing. (b) Symmetry: fit separate exponents for `r > 1` and `r < 1`; if they differ, the symmetric form is wrong and the correction protocols must be reweighted.**
7 Reproducibility	**UNTESTED.** Both (a) and (b) are cheap and could be run by any organisation with a decision log and a forecast record; this is the most tractable open test in the paper.

Table V-G5 — The join `ln r = ln D - ln p_bar - F`

Gate	Finding
1 Dimensions	Every term in nats after taking logarithms of dimensionless or commensurate quantities. **PASS.**
2 Scale type	Log-ratios are interval-scaled differences; the identity is an equation between such differences. **PASS.**
3 Limiting cases	`F -> 0`: `ln r -> ln(D/p_bar)`, maximum distortion for the given doctrine load. `F -> ln N`: `ln r -> ln(D/(p_bar.N))`, minimum achievable. `D -> 0`: `ln r` diverges downward, Case 2. **PASS.**
4 Fixed points	`F* = ln(D/p_bar)` gives `r = 1`, feasible if and only if `D/p_bar <= N`. **The infeasibility condition is the fixed point's domain of existence, and it is the hardest result in this section. PASS.**

Gate	Finding
5 Identifiability	Given independent measurements of `D`, `p_bar` and `u`, the identity is over-determined — `ln r` can be computed two ways and the two must agree. **This over-determination is a consistency check and it is the most useful diagnostic available: a mismatch means one of the three measurements is wrong. PASS.**
6 Falsification	The identity itself is algebraic and cannot be falsified; what is falsifiable is the aggregation postulate it rests on. **Kill condition: measured `P` fails to track `p_bar . N_eff` across organisations with known `u`.**
7 Reproducibility	**UNTESTED, and it inherits V-G3's status.**

24.6.13 6.13 Failure modes at group scope

FM-G1 The void denominator. `P` computed as headcount or subscription count, without the redundancy audit. *Damages:* `P` overstated, therefore `r` understated; the organisation believes it is calibrated. *Signature:* large, well-funded intelligence function; confident briefings; recurring surprises. *Measurement:* error decorrelation across channels on questions with observable answers. *Fix:* recompute `P` as `p_bar . N_eff`. Expect the number to fall by a large factor: at `u = 0.05` in an `N = 12` domain, `N_eff = 1.55` against a headcount that may be forty. The formal statement that an unaudited count is worthless, because one party can present as many, is the Sybil result [Douceur, 2002].

FM-G2 The unidentified ratio. `D` reported without an independent `P`. Four distinct parameter sets produce the same observation. *Damages:* measurement validity. *Signature:* doctrine described in absolute terms — "our model is very strong" — with no channel count behind it. *Measurement:* ask for `P`; if it is not measured, `D` is not measured. *Fix:* the anchor — count the audited channels, then `D = r.P`.

FM-G3 The bare magnitude. `Delta` reported without the sign of `ln r`. *Damages:* diagnosis; the correction is applied in the wrong direction. *Signature:* a distortion score with no direction, followed by an intervention that makes things worse. Adding channels to an already under-framed organisation is the common instance. *Measurement:* read the report; is `ln r` signed? *Fix:* always report the signed log-ratio.

FM-G4 The calibrated irrelevance. `r = 1`, `Delta = 0`, and `E -> 0`. Perfect contact with a vanished environment. *Damages:* `M`, totally, while every internal metric reads well. *Signature:* excellent internal process metrics; outputs that land nowhere; a constituency that has migrated. *Measurement:* estimate `E` directly. *Fix:* outside the scope of this equation — the organisation must change what environment it is in. The formula's contribution is the diagnosis, not the remedy.

FM-G5 The zero-headroom success. The organisation calibrates, banks to its maximum, and finds `h = 0` exactly. Every period is maintenance. *Damages:* `h`. *Signature:* a well-run institution that cannot start anything, with no identifiable culprit. *Measurement:* `h = eta ln N − delta Q − C` at the current `Q`, against `Q_inf = (eta ln N − C)/delta`. *Fix:* shed `Q` deliberately, or raise `eta`, or cut `C`. **Do not attempt to fix it by degrading calibration; that reduces `S` faster than it releases budget.**

FM-G6 The proxy that was the input. The order proxy is a repackaging of intelligence volume — "we produced 400 assessments" counted as stored order. *Damages:* identifiability of `eta`; the whole model becomes untestable. *Signature:* excellent training fit, held-out failure; `eta` near 1 with negligible residual. *Fix:* an order proxy with an external referent: decisions whose outcomes are later observable, standing capability a third party can verify, infrastruc-

ture that survives staff turnover.

FM-G7 The unbadged equation. `M = (D/P).E` presented as though it were a theorem. *Damages:* the argument's warrant. *Signature:* "the mathematics shows this organisation is over-framed." *Measurement:* read the sentence. *Fix:* restate with the badge on: the moment equation is a postulate, its consequences are computable, and the reproducibility gate has not been run.

24.7 7 Doctrine-to-data: two pathologies, opposite corrections

Section 6 established the group moment and its calibration ratio. This section is about the two ways an organisation gets that ratio wrong. Both failures are common, both are measurable with instruments an organisation can build from records it already keeps, and both are correctable.

They are not, however, correctable by the same means. **The two pathologies map to the same distortion magnitude and require opposite interventions, which is why an unsigned diagnosis is worse than no diagnosis at all.** An organisation told it is distorted and not told in which direction will reach for the nearest available remedy — usually more collection — and half the time that remedy is the disease.

This section also closes a loop. The framework's pseudo-diversity result describes an institution whose members are epistemically parallel, whose agreement is read as convergence, and whose effective count is one [Curzi, 2026c]. That result prices the condition in surplus. Here it is priced in perception, and the result is stronger than the analogy suggests: **pseudo-diversity is not merely correlated with over-framing. It is a mechanism that manufactures over-framing, from constant doctrine, with no one adding a single new belief.**

24.7.1 7.1 The two directions and the lever asymmetry

```
r = D / P          Delta = | ln r |          M = r.E          S = ln N . min(r, 1/r)

r > 1   OVER-FRAMED     M > E    doctrine outruns intelligence
r = 1   CALIBRATED      M = E    undistorted contact
r < 1   UNDER-FRAMED    M < E    intelligence outruns doctrine
```

`Delta` is a magnitude and discards the sign. That is adequate for comparing severity and fatal for choosing a treatment, because the two conditions sit on opposite sides of a fixed point and every lever moves `ln r` in one direction only.

LEVER	effect on ln r	helps OVER-framed	helps UNDER-framed
raise u (more independence)	decreases	YES	NO – makes it worse
raise N (broader domain)	decreases	YES (weakly, log)	NO – makes it worse
raise p_bar (deeper intake)	decreases	YES	NO – makes it worse
cut D (shed doctrine)	decreases	YES	NO – makes it worse
raise D (build a frame)	increases	NO – makes it worse	YES
triage P (cut channels)	increases	NO – makes it worse	YES

Every standard organisational virtue in the left column — more diversity, more breadth, more collection, less dogma — is a poison to an under-framed institution. Section 6.8 computed this: an organisation carrying `D/p_bar = 6.0` at `N = 12` reaches `r = 0.5` at `u = 1.0`, and its bankable input falls from the ceiling of 2.4849 nats back to 1.2425. Diversity past the calibration point is a cost, and the model says so without hedging. This is the governing rule of the section and it is stated as a requirement, not a suggestion:

> **Never report a distortion without its sign.** Report `ln r`, signed. `Delta` exists for comparing severity across units, never for prescribing.

24.7.2 7.2 The manufacturing result

The framework's finding is that `u` is not a property an institution acquires but a quantity that decays and must be maintained; Section 4.12 showed the same decay in the two-person case and gave it a half-life. Putting that decay into the join produces something specific.

```
ln r_t = ln D − ln p_bar − F_t          F_t = ln( 1 + (N−1) u_t )
with   u_t = u_0 (1 − lambda)^t
```

Hold `D` fixed. Hold `p_bar` fixed. Hold `N` fixed. Hold headcount, budget and collection volume fixed. **`ln r` still rises, monotonically, because `F_t` falls.**

> **The manufacturing result.** An organisation that adds no doctrine, cuts no collection and changes no policy will drift from calibration into over-framing at a rate set entirely by the decay of `u`. The crossing occurs at
>
> ```
> u* = ( D/p_bar − 1 ) / ( N − 1 )        t* = ln( u* / u_0 ) / ln( 1 − lambda )
> ```
>
> and after `t*` the organisation is over-framed, with an unchanged model, by arithmetic alone.

Worked, at `N = 12`:

D/p_bar	u_0	lambda	u*	t* (periods)
4.0	0.55	0.08	0.2727	8.41
4.0	0.55	0.04	0.2727	17.18
3.0	0.30	0.06	0.1818	8.09

Roughly two years, at quarterly periods and an eight-per-cent-per-quarter convergence rate, from calibrated to over-framed with nobody doing anything. Halving the decay rate roughly doubles the time, as the logarithmic form requires; it does not prevent the crossing.

This is why the connection to pseudo-diversity is a mechanism and not an analogy. Its signature — `u -> 0`, `N_eff -> 1`, surplus `ln 1 = 0` — enters `r = D/(p_bar . N_eff)` as a *collapsing denominator*. **The institution's doctrine did not grow. Its denominator vanished.** The same erosion of local variety by a legibility-seeking centre is documented historically [Scott, 1998], and the homogenising effect of fast mutual learning is documented with a simulation [March, 1991].

24.7.3 7.3 The honesty clauses of this section

1. Every instrument below estimates `r` **through a proxy**, and each proxy carries its own badged postulate linking the observable to the ratio. None of them measures `D` or `P` directly.
2. Only `r` is identifiable from `M` and `E`; `D` requires an independent anchor on `P`.
3. The contact postulate's symmetry — over- and under-framing attenuating equally at equal `Delta` — is **untested** and is the first thing a real study should attack.
4. `Delta` values here are model quantities. **No empirical distribution of institutional `Delta` exists**, because nobody has measured one. Where a threshold is recommended, it is derived from the organisation's own `S_min = C/eta`, not from a literature.
5. The reproducibility gate is unrun for every formula here.

24.7.4 7.4 The over-framed institution and its instruments

Clinical description. `r > 1`. The organisation's standing model supplies more constraint on its decisions than its audited intelligence does. It therefore constructs a situation larger and more determinate than its environment. It sees what its model predicts. The condition is not stupidity and it is not usually arrogance: it is the ordinary end-state of a successful institution, because doctrine accumulates while `u` decays.

The behavioural signature, in the order it usually appears: assessments become faster and more confident; new evidence arrives and is *absorbed* — explained within the model — rather than escalated; conclusions stop moving and the revision rate falls; surprises arrive from outside the model's categories, are treated as one-offs, and are absorbed; internal metrics improve while external outcomes do not.

The last point is the diagnostic crux. **An over-framed organisation's self-assessment improves as its condition worsens**, because self-assessment is conducted through `M`, and `M` is exactly what is inflated. That is why every instrument below is built on records with an *external* referent. Four instruments follow, each with its postulate, its data requirement, its scale type, and what it can and cannot conclude.

I-1 The revision rate.

> **The revision postulate [postulate].** A calibrated organisation revises standing conclusions at some baseline rate `sigma*` determined by the volatility of its environment. The observed revision rate `sigma` — the fraction of intelligence products that alter a standing conclusion — estimates the ratio by `r_hat = sigma* / sigma`.

Data: a decision or assessment log with, per product, a recorded before-state and after-state of the standing conclusion. Scale: `sigma` is a proportion, ratio-scaled with a true zero.

```
sigma   r_hat    ln r_hat   reading
-----   ------   --------   ------------------
0.40    0.500    -0.6931    under-framed
0.30    0.667    -0.4055    under-framed
0.20    1.000    +0.0000    calibrated (sigma* = 0.20 here)
0.15    1.333    +0.2877    over-framed
0.10    2.000    +0.6931    over-framed
0.05    4.000    +1.3863    over-framed, severe
0.02   10.000    +2.3026    over-framed, severe
```

Limits, stated plainly. `sigma*` is not known a priori and must be estimated from the environment's own volatility, which requires an external referent. **Without a defensible `sigma*` this instrument produces an ordering, not a magnitude**, and must be reported as such. It remains useful in that demoted form: `sigma` falling over time in a stable environment is evidence of drift toward over-framing.

I-2 The anomaly disposal audit. Take every observation over a period that was inconsistent with the standing model. Classify each as **escalated** (it changed something), **absorbed** (it was explained within the model), or **discarded** (it was set aside as noise, error or an outlier).

```
A = (absorbed + discarded) / total anomalies       escalation rate = 1 - A

A       escalation   odds absorb:escalate   ln odds
----    ----------   --------------------   -------
0.30       0.70            0.43 : 1         -0.8473
0.50       0.50            1.00 : 1         +0.0000
0.70       0.30            2.33 : 1         +0.8473
0.80       0.20            4.00 : 1         +1.3863
0.90       0.10            9.00 : 1         +2.1972
```

```
0.95       0.05          19.00 : 1         +2.9444
0.98       0.02          49.00 : 1         +3.8918
```

The log-odds is the right register: it is a log-ratio, commensurate with `ln r` and with `F`. **This instrument is the most robust of the four**, because it requires no baseline — the raw classification counts are meaningful on their own, and the audit runs on records the organisation already generates whenever it writes down why something was not acted on.

Its weakness is that anomaly identification is itself performed through the model. **An organisation severely over-framed does not record anomalies as anomalies; it does not see them.** The audit therefore *understates* severe cases, which makes a high `A` conclusive and a low `A` inconclusive. That must be said in the report.

I-3 The calibration curve. Plot stated confidence against realised frequency across a forecast record. Over-framed institutions are overconfident: the 90 per cent band contains the outcome less than 90 per cent of the time. Under-framed institutions are the reverse or, much more commonly, make no quantified forecasts at all, which is itself the diagnosis.

Data requirement: a forecast register with resolution dates and recorded outcomes. Scale: proportions on both axes, ratio-scaled. **This instrument does not estimate `r` numerically without a further postulate mapping overconfidence to the ratio, and no such postulate is invented here. It is reported as corroborating evidence with a direction, not as a magnitude.** Its value is that it is entirely external: nothing about it passes through `M`.

I-4 The channel audit — the structural instrument. The other three are behavioural. This one measures the mechanism directly, and the manufacturing result makes it indispensable.

```
Step 1.  Enumerate the intelligence channels and the N declared domains.
Step 2.  For questions whose answers later become observable, record each
         channel's assessment BEFORE any cross-channel discussion.
Step 3.  Compute pairwise error decorrelation; take the mean as u.
Step 4.  N_eff = 1 + (N-1)u.   F = ln N_eff.   P = p_bar . N_eff.
Step 5.  ln r = ln D - ln p_bar - F, with D from the doctrine enumeration.
Step 6.  Cross-check against ln D - ln P.  They must agree.
```

Worked at `N = 12`, showing what headcount hides:

```
u       N_eff     F        P/headcount at 40 staff
-----   ------   ------    -----------------------
0.00     1.000   0.0000            0.0250
0.02     1.220   0.1989            0.0305
0.05     1.550   0.4383            0.0387
0.10     2.100   0.7419            0.0525
0.15     2.650   0.9746            0.0663
0.25     3.750   1.3218            0.0938
0.40     5.400   1.6864            0.1350
0.60     7.600   2.0281            0.1900
```

A forty-person intelligence function at `u = 0.05` has an audited intake of 1.55 channels — under four per cent of its headcount. That is the number that goes in the denominator, and it is the number nobody computes. The channel audit is the only instrument of the four that says *why*: the other three say that `r` is high; this one says whether `D` grew or `P` collapsed, and those have different fixes.

24.7.5 7.5 The under-framed institution and its instruments

Clinical description. `r < 1`. Audited intelligence exceeds the constraint the standing model supplies. The organisation constructs a situation smaller and less determinate than its environment. It cannot make its inputs mean anything. This is the less-discussed pathology and it

is at least as common: it is the natural state of a new organisation, a merged one, a rapidly expanded one, and any organisation that responded to a past over-framing episode by collecting more.

The behavioural signature: intake grows and dashboards multiply; decision latency grows and items are deferred pending further information; the standing request is always for more data; nothing is retired, because no conclusion is ever firm enough to close a line of enquiry; and staff describe the work as overwhelming and the organisation as lacking direction. **Both reports are accurate and they are the same report.**

Under-framing is **not** ignorance. It is the opposite: a surfeit of intake with insufficient frame to convert intake into a situation. Case 2 of Section 6.5 gives the limit exactly — `D -> 0` implies `M -> 0`, not `M` small but `M` zero. A perfectly open-minded institution does not decide badly; it does not decide.

I-5 The intake-to-decision ratio.

> **The conversion postulate [postulate].** Let `kappa` be intake-units consumed per decision closed. A calibrated organisation runs at some `kappa*` set by its decision complexity. Then `r_hat = kappa* / kappa`.

```
kappa / kappa*     r_hat     ln r_hat
--------------    ------    --------
      0.25        4.0000    +1.3863     over-framed: decides on almost nothing
      0.50        2.0000    +0.6931
      1.00        1.0000    +0.0000     calibrated
      2.00        0.5000    -0.6931
      4.00        0.2500    -1.3863     under-framed
      8.00        0.1250    -2.0794     under-framed, severe
```

Data: an intake tally and a decision log, both of which most organisations keep. Scale: both are counts, so `kappa` is ratio-scaled. Same limitation as I-1: `kappa*` requires justification, and without one the instrument produces an ordering rather than a magnitude. **Trend in `kappa` under stable decision complexity is the usable form.**

I-6 The unread ratio and the backlog slope. Two counts: intelligence products produced, and intelligence products demonstrably consulted in a decision within their useful life. The ratio is a direct, external, ratio-scaled measure of intake the organisation cannot frame. A rising backlog with a stable production rate is a queueing statement read backwards: with `L = lambda W`, a growing queue at constant arrival rate means growing residence time [Little, 1961]. **Under-framing shows up in the queue before it shows up in anyone's account of the problem**, which makes this the earliest available indicator, and the same delay-and-stock misjudgement is measured experimentally in the beer distribution game [Sterman, 1989].

I-7 The instrument proliferation count. Count distinct standing reports, dashboards and metrics; count how many are cited in decisions. Proliferation with flat citation is diagnostic: the organisation is answering an inability to frame by adding intake, which moves `r` in exactly the wrong direction. **This instrument detects the characteristic self-inflicted worsening of under-framing.**

I-8 The forecast-absence test. Ask for the organisation's quantified forecasts with resolution dates. An over-framed organisation has them and they are overconfident. An under-framed organisation **does not have them at all** — not because it is careless, but because a forecast is an act of doctrine and it has insufficient doctrine to make one. Binary, external and cheap. Its absence is not proof, but combined with I-5 and I-6 it is close to conclusive on direction. The general case for putting a risky prediction on the record rather than a directional one is Meehl [1967].

24.7.6 7.6 Reading the instruments together, and the manufacturing result computed

No single instrument is sufficient. The direction table:

		OVER-FRAMED	UNDER-FRAMED
I-1	revision rate sigma	LOW	HIGH
I-2	anomaly absorption A	HIGH	LOW (or unclassifiable)
I-3	calibration curve	overconfident	underconfident / absent
I-4	channel audit u, N_eff	LOW u, LOW N_eff	HIGH N_eff relative to D
I-5	intake per decision kappa	LOW	HIGH
I-6	unread ratio, backlog	LOW	HIGH, rising
I-7	instrument proliferation	LOW	HIGH with flat citation
I-8	quantified forecasts	present, wrong	absent

Rules for reading. **Direction requires at least two instruments agreeing, one behavioural and one structural**, and I-4 should be one of them wherever the data allow, because it is the only one that identifies the cause. **A mixed reading is not noise; it is usually a mixed organisation. I-2 is conclusive when high and inconclusive when low. Where `sigma*` and `kappa*` cannot be defended, report I-1 and I-5 as orderings, not magnitudes**, and print the demotion.

Now take an organisation calibrated on the under-framed side and let nothing happen except `u` decay: `N = 12`, `D = 4.0`, `p_bar = 1.0`, `u_0 = 0.55`, `lambda = 0.08` per period, and for the storage side `eta = 0.42`, `delta = 0.14`, `C = 0.55`.

t	u	N_eff	F	P	r	ln r	Delta	S	Q_inf
0	0.5500	7.050	1.9530	7.050	0.567	−0.5667	0.5667	1.4099	+0.3011
3	0.4283	5.711	1.7424	5.711	0.700	−0.3561	0.3561	1.7404	+1.2927
6	0.3335	4.668	1.5408	4.668	0.857	−0.1545	0.1545	2.1291	+2.4588
9	0.2597	3.857	1.3498	3.857	1.037	+0.0365	0.0365	2.3958	+3.2589
12	0.2022	3.224	1.1707	3.224	1.241	+0.2156	0.2156	2.0031	+2.0806
15	0.1575	2.732	1.0051	2.732	1.464	+0.3812	0.3812	1.6973	+1.1632
18	0.1226	2.349	0.8539	2.349	1.703	+0.5324	0.5324	1.4591	+0.4488
21	0.0955	2.050	0.7180	2.050	1.951	+0.6683	0.6683	1.2737	−0.1075
24	0.0743	1.818	0.5976	1.818	2.200	+0.7887	0.7887	1.1293	−0.5407

The crossing is at `u* = (4.0 − 1)/11 = 0.2727`, reached at `t* = ln(0.2727/0.55)/ln(0.92) = 8.41` periods.

Read the trajectory as a story, because it is a common institutional history. **Periods 0 to 6:** the organisation is under-framed and improving; its distortion falls, its bankable input rises, its equilibrium store rises. Everything anyone measures is getting better. **Period 9:** calibration, `Delta = 0.037`, `S = 2.3958` against a ceiling of 2.4849. Peak performance, and it is genuinely peak: `Q_inf = 3.26`. **Periods 12 to 18:** the same process continues and now runs the other way. `u` keeps falling, `F` keeps falling, `ln r` keeps rising, and every quantity that improved now degrades — **driven by the identical underlying trend, with no change of behaviour, no change of doctrine, and no decision by anyone. Period 21:** `Q_inf` goes negative and the institution's standing capability erodes.

> **The mechanism that produced the golden period is the mechanism that produced the decline. They are the same mechanism at different points on one monotone curve.**

This is why organisations are reliably unable to diagnose themselves at period 18. The behaviour that was rewarded at period 6 — building shared understanding, converging on a common view, reducing internal friction — is the behaviour that is killing them at period

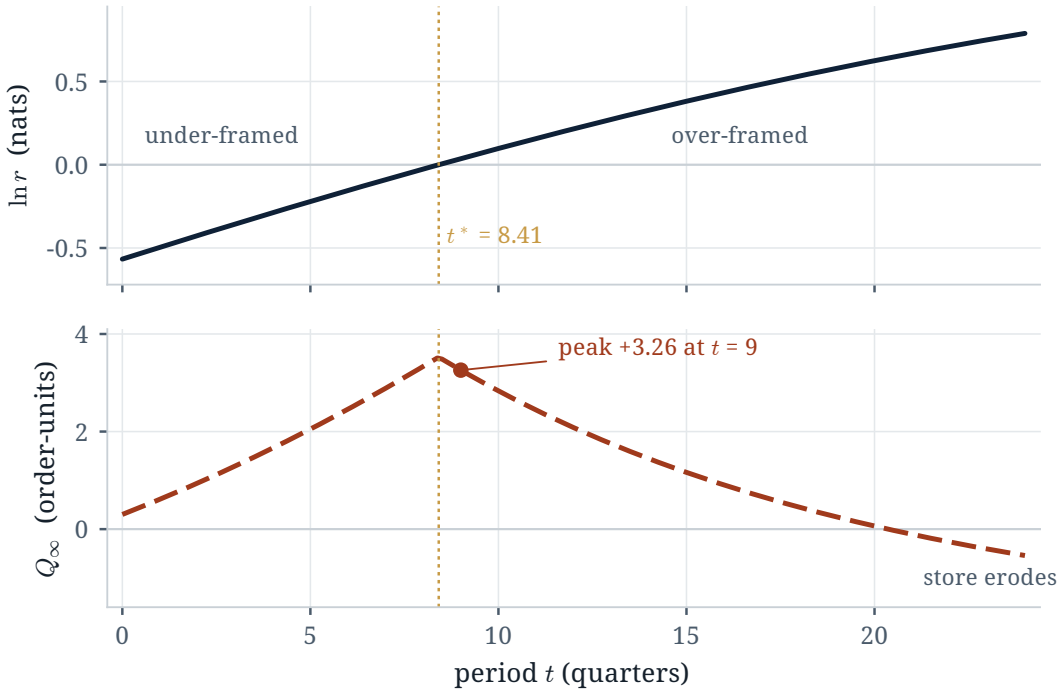


Figure 24.3: **An organisation that changes nothing crosses from under-framed to over-framed on a computable date.** With doctrine, collection budget, headcount and domain all held fixed, non-parallelism decaying at 8 per cent per period drives `ln r` monotonically upward through zero at `t* = 8.41` periods, and the equilibrium store rises to a genuine peak of `+3.26` at period 9 before falling through zero at period 21. Modelled at `N = 12`, `D = 4.0`, `p_bar = 1.0`, `u_0 = 0.55`, `lambda = 0.08`, `eta = 0.42`, `delta = 0.14`, `C = 0.55`, tabulated in Section 7.6; generated by `figs/src/CCP-024_fig03.py`.

18, and it never changed. **Pseudo-diversity manufactures over-framing, and it does so most effectively in organisations that are proud of their cohesion.** Cohesion is often the enemy of the measured collective-intelligence factor [Woolley et al., 2010] because cohesion drives `u` toward zero; in this section's terms, it drives the denominator of `r` toward `p_bar`. The same interaction that raises local agreement can freeze global diversity into non-interacting blocs, with a sharp transition between the monocultural and fragmented phases [Axelrod, 1997].

24.7.7 7.7 Distortion does not average

Organisations are not uniform. A headquarters and a field operation can sit on opposite sides of the fixed point, and the aggregate figure will report neither. Take two units, `N = 12`, `p_bar = 1.0` throughout:

```
unit          u       N_eff = P     D       r        ln r       Delta     S
----------   -----   ----------   -----   ------   --------   ------   ------
HQ           0.11       2.210      6.80   3.0769   +1.1239    1.1239   0.8076
Field        0.62       7.820      1.15   0.1471   −1.9169    1.9169   0.3654

AGGREGATE (D and P summed):          7.95 / 10.030 = 0.7926    ln r = −0.2324
```

The aggregate reads `Delta = 0.2324` — nearly calibrated — while both halves are severely distorted, at 1.12 and 1.92. The aggregate is not an average of the units' distortions and it is not a summary of them; it is a different quantity, and here it is actively misleading.

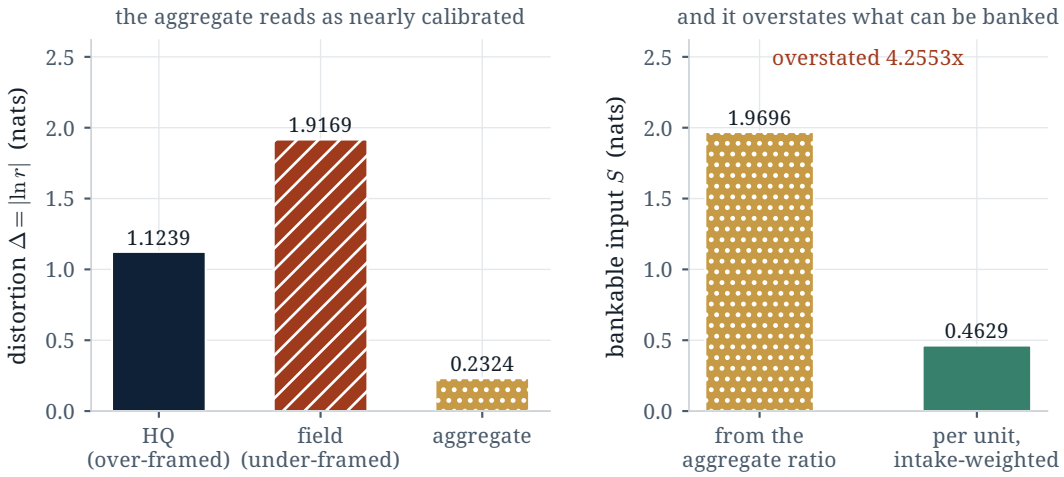


Figure 24.4: **An aggregate that reads as nearly calibrated conceals two units in opposite ditches**: the headquarters at `Delta = 1.1239` over-framed and the field at `Delta = 1.9169` under-framed combine to an aggregate `Delta = 0.2324`, and bankable input computed from that aggregate reads 1.9696 nats against a true intake-weighted 0.4629 — an overstatement of 4.2553 times. Computed from the two-unit case tabulated in Section 7.7 at `N = 12` and `p_bar = 1.0`; generated by `figs/src/CCP-024_fig04.py`.

The arithmetic of why: `r_agg = (D_1+D_2)/(P_1+P_2)` is a weighted mediant, so it always lies between `r_1` and `r_2`. When both units are distorted in the **same** direction the aggregate lies between them and understates the worse one, which is tolerable. When they are distorted in **opposite** directions the aggregate lies between two values straddling 1, and can sit arbitrarily close to 1 while both units are arbitrarily far from it. **Cancellation in the books, no cancellation in the world.** This is the general aggregation objection made concrete [Kirman, 1992, Quetelet, 1835], and the reason a distribution rather than its mean is what determines an outcome [Granovetter, 1978].

The consequence for banking is severe:

```
S computed from the aggregate ratio:          2.4849 x 0.7926 = 1.9696
S computed per unit and intake-weighted:      0.4629
OVERSTATEMENT FACTOR:                         4.2553
```

Computing an organisation's bankable input from its aggregate ratio overstates it by a factor of 4.26 in this case. The aggregate is not conservative and not approximately right; it is wrong in the dangerous direction. *(The exact ratio recomputed at full precision is 4.2553. The source manuscript prints this quantity twice, once as 4.26 and once, for the same numbers, as 4.25; the value is stated here to four decimals so that the discrepancy is a rounding convention rather than a disagreement.)*

> **Rule. Compute `r` per decision-making unit. Aggregate the units' `S`, weighted by intake, never the units' `D` and `P`.**

One further consequence, which is where the correction protocols start: a mixed organisation has an obvious-looking remedy — move doctrine from the over-framed unit to the under-framed one — and it does not fully work. Doctrine total is 7.95 and intake total is 10.030. **Even perfect redistribution leaves the organisation doctrine-poor overall by 2.080 units, or 26.2 per cent of its standing doctrine.** Redistribution alone can equalise the distortion; it cannot remove it.

24.7.8 7.8 Correction protocols

Over-framed. The correction runs in one direction only: `ln r` must come down. There are exactly four levers and they have very different prices.

O-1 Raise `u`. Recruit across training lineages, rotate, stand up genuinely independent channels, commission from unaligned sources, install adversarial review, import a practitioner from an unrelated trade. *Price and ceiling:* `ln r` falls by exactly the gain in `F`, one nat for one nat, and the total available is `ln N`. Concavity means the early gains are the large ones: at `N = 12`, moving `u` from 0.14 to 0.30 buys 0.5264 nats; moving from 0.60 to 0.90 buys 0.3607. **Raise `u` first, and expect the second independent channel to be worth more than the twentieth.** *The trap:* `u` gained is `u` that will decay again at rate `lambda`. This lever must be operated as a maintenance schedule, not a project. **A one-off diversity intervention buys a decaying quantity and is the commonest wasted correction.**

O-2 Raise `p_bar`. Deepen the intake of existing channels: longer collection, better instrumentation, more careful reading. *Price:* directly proportional and unbounded in principle. In practice it is bounded by budget and by the fact that deepening a channel does not make it independent. **`p_bar` and `u` are not substitutes: a hundred times more of one wire service is still one channel.**

O-3 Raise `N`. Broaden the declared domain. *Price:* logarithmic and expensive. It raises the ceiling on what `u` can buy but buys nothing by itself. **Worst of the four, and the one most often reached for, because expanding scope is easier to authorise than admitting error.**

O-4 Cut `D`. Retire standing assumptions. Sunset guidance. Shed doctrine. *Price:* unbounded, since `D` can go to zero, but it costs `Q`, and shedding `Q` is legitimate, underused and hard. **This is the only lever that works in the infeasible region `D/p_bar > N`.**

The comparison, at `N = 12`, `D = 5.2`, `p_bar = 1.0`, `u = 0.14`:

```
target r     via u                      via D
--------     ----------------------     ------------------------
1.898        u = 0.1582  (x 1.13)       D = 4.821  (-7.3 %)
1.500        u = 0.2242  (x 1.60)       D = 3.810  (-26.7 %)
1.200        u = 0.3030  (x 2.16)       D = 3.048  (-41.4 %)
1.000        u = 0.3818  (x 2.73)       D = 2.540  (-51.2 %)
```

Getting merely inside the operating envelope costs a 13 per cent increase in `u` or a 7 per cent cut in doctrine. Full calibration costs nearly a tripling of `u` or halving the standing model. Stabilise before optimising, and take the cheap fraction of the distance first.

```
O-STEP 1.  Confirm direction with two instruments, one of them I-4.
O-STEP 2.  Compute the infeasibility check: is D/p_bar > N?
           If YES -> only O-4 works. Go to step 5.
O-STEP 3.  Compute the operating envelope from S_min = C/eta.
           Target the envelope edge, not r = 1.
O-STEP 4.  Price O-1 and O-2 against O-4 at that target. Take the cheapest.
O-STEP 5.  If O-4: enumerate the doctrine, rank by (decisions constrained) /
           (evidence supporting), retire from the top.
O-STEP 6.  Install the maintenance schedule for u BEFORE declaring the
           correction complete. Without it, t* restarts.
O-STEP 7.  Preregister the held-out test before intervening.
```

Step 5's ranking deserves a note. The doctrine to retire first is not the doctrine most likely to be wrong; it is the doctrine that constrains the most decisions on the least evidence. **A weakly supported proposition that governs one decision a year is a curiosity. A weakly supported proposition that governs every decision is the whole problem.**

Under-framed. The correction runs the other way and it is **not** the mirror image, for a reason that matters: building a frame is a slower, more skilled and more contestable act than adding a channel.

U-1 Raise `D` — build the frame. Adopt an explicit model. Write the standing assumptions down. Commit to a hypothesis before collecting. Publish quantified forecasts with resolution dates. Close lines of enquiry. *Price:* this is the only lever that addresses the cause, and it is expensive in the one currency institutions guard most closely: **someone must be willing to be wrong in writing.** Under-framing persists not because frames are hard to obtain but because an unstated frame cannot be falsified and a stated one can. At `N = 12`, `u = 0.55` so `P = 7.05`, `eta = 0.42`, `delta = 0.14`, `C = 0.55`:

D	r	Delta	S	Q_inf
0.50	0.0709	2.6462	0.1762	−3.3999
1.00	0.1418	1.9530	0.3525	−2.8712
2.00	0.2837	1.2599	0.7049	−1.8138
3.00	0.4255	0.8544	1.0574	−0.7564
4.00	0.5674	0.5667	1.4099	+0.3011
5.00	0.7092	0.3436	1.7623	+1.3585
6.00	0.8511	0.1613	2.1148	+2.4159
7.05	1.0000	0.0000	2.4849	+3.5261

This organisation's equilibrium store crosses zero between `D = 3` and `D = 4`. Below that it cannot hold anything, no matter how good its intelligence is. The remedy is not better collection. It has excellent collection. **It needs a model, and until it has one its collection is unbankable.**

U-2 Triage `P` — cut channels deliberately. The uncomfortable lever, and it is real. Reducing the number of independent channels raises `r` toward 1. Worked at `D = 1.6`, `N = 12`:

u	P	r	Delta	S
0.5500	7.050	0.2270	1.4830	0.5640
0.4500	5.950	0.2689	1.3134	0.6682
0.3500	4.850	0.3299	1.1090	0.8198
0.2500	3.750	0.4267	0.8518	1.0602
0.1500	2.650	0.6038	0.5046	1.5003
0.0545	1.599	1.0003	0.0003	2.4841

Reaching calibration by triage alone would require destroying almost all of this organisation's independence — `u` from 0.55 to 0.055 — and it would work, in the sense that `S` returns to the ceiling. **The lever is stated and then the reason it is nearly always wrong is stated.** First, it is irreversible on the timescale that matters: `u` destroyed must be rebuilt from recruitment, and rebuilding difference has no internal source. Second, it caps the organisation's future, because `N_eff` is the ceiling on every later correction. Third, it is the same move as pseudo-diversity, chosen deliberately, so the whole trajectory of Section 7.6 applies immediately. **Use triage only for genuinely redundant channels — channels whose measured pairwise `u` against existing ones is near zero.** Those cost budget and contribute nothing to `N_eff`, so cutting them raises `r` almost not at all and frees resources for U-1. That is the honest, narrow, defensible version of this lever.

```
U-STEP 1.  Confirm direction with two instruments; I-8 and I-6 are the cheapest pair.
U-STEP 2.  Compute the operating envelope from S_min = C/eta and locate the
           D at which Q_inf crosses zero. That is the minimum viable frame.
U-STEP 3.  Build the frame: write the standing assumptions, commit to
           hypotheses, publish forecasts with resolution dates.
U-STEP 4.  Close lines of enquiry explicitly. An unclosed enquiry is
           continuing intake against a frame that does not yet exist.
U-STEP 5.  Triage ONLY channels with measured near-zero pairwise u.
```

```
U-STEP 6.  Re-measure. Under-framed organisations correct FAST once the
           frame lands, because D is the numerator and moves directly.
U-STEP 7.  Set the over-framing alarm now: record t* from the current u,
           lambda and D, and diarise the re-measurement before it.
```

Step 7 closes the loop. An organisation that corrects under-framing by building doctrine has raised `D` — and `D` does not decay, while `u` does. It has therefore shortened its own `t*` and set itself on the trajectory of Section 7.6. **The correction for one pathology is the seed of the other, and the only defence is to have computed the crossing date in advance.**

24.7.9 7.9 Validation gates for the instruments

Table V-C1 — Revision-rate estimator `r_hat = sigma*/sigma` (I-1)

Gate	Finding
1 Dimensions	`sigma`, `sigma*` are proportions, dimensionless; `r_hat` dimensionless. **PASS.**
2 Scale type	Proportions are ratio-scaled with a true zero. **PASS for `sigma`. The gate binds on `sigma*`:** if the baseline cannot be estimated from an external referent, `r_hat` has no magnitude and **the instrument must be demoted in print to an ordering.**
3 Limiting cases	`sigma -> 0`: `r_hat` diverges, an organisation that never revises — real, nameable, the severe end of over-framing. `sigma -> 1`: every product overturns a conclusion, `r_hat -> sigma*`, the maximally unstable frame. **PASS.**
4 Fixed points	`sigma = sigma*` gives `r_hat = 1`. Not an attractor; nothing drives `sigma` toward `sigma*`. **PASS, with instability noted.**
5 Identifiability	`sigma` is recoverable from a decision log with before/after states. `sigma*` is **not recoverable from the organisation's own records** and requires environmental volatility measured externally. **CONDITIONAL PASS.**
6 Falsification	Preregister `sigma` over a fit window; predict held-out surprise rate. **Kill condition: `sigma` fails to predict held-out surprise rate beyond the shuffled-input control.**
7 Reproducibility	**UNTESTED.** Cross-institution recovery of `sigma*` for comparable environments is the specific open question.

Table V-C2 — Anomaly absorption `A`, reported as `ln(A/(1-A))` (I-2)

Gate	Finding
1 Dimensions	`A` a proportion; the log-odds in nats, commensurate with `ln r` and `F`. **PASS.**

Gate	Finding
2 Scale type	Counts of anomalies by disposal class: ratio-scaled with a true zero. **PASS, and this is the strongest scale-type position of any instrument here, because it needs no baseline.**
3 Limiting cases	`A -> 1`: nothing is ever escalated, log-odds diverge upward; the sealed model. `A -> 0`: everything is escalated, log-odds diverge downward; a frame so weak that no observation can be placed. Both real. **PASS.**
4 Fixed points	`A = 0.5`, log-odds 0. This is **not** claimed to correspond to `r = 1`; the mapping from `A` to `r` is not established and none is invented. **PASS as a fixed point of the instrument, NOT of the ratio.**
5 Identifiability	`A` recoverable from disposal records. **The known bias, printed: anomaly identification is performed through the model, so severe over-framing suppresses the denominator. High `A` is conclusive; low `A` is not. CONDITIONAL PASS with declared bias.**
6 Falsification	Preregister `A`; predict the rate of later-observable events the organisation had classified as noise. **Kill condition: no relation between `A` and the missed-event rate beyond the shuffled control.**
7 Reproducibility	**UNTESTED.** Inter-rater agreement on the three disposal classes is a prerequisite and has not been established.

Table V-C3 — Intake-to-decision estimator `r_hat = kappa*/kappa` (I-5)

Gate	Finding
1 Dimensions	`kappa` in intake-units per decision; `kappa*` the same; `r_hat` dimensionless. **PASS, provided intake-units are the convention's units.**
2 Scale type	Both are ratios of counts, ratio-scaled. **PASS for `kappa`; binds on `kappa*` exactly as V-C1 binds on `sigma*`. Demote to an ordering where `kappa*` is undefended.**
3 Limiting cases	`kappa -> 0`: decisions taken on no intake, `r_hat` diverges, over-framing. `kappa -> infinity`: unbounded intake per decision, `r_hat -> 0`, the paralysed institution. **PASS.**
4 Fixed points	`kappa = kappa*`. **PASS.**
5 Identifiability	`kappa` from intake tally and decision log, both routinely kept. `kappa*` requires an external standard of decision complexity. **CONDITIONAL PASS.**

Gate	Finding
6 Falsification	Preregister `kappa`; predict held-out decision latency and reversal rate. **Kill condition: no relation beyond the shuffled control.**
7 Reproducibility	**UNTESTED.**

Table V-C4 — The manufacturing model `ln r_t = ln D − ln p_bar − ln(1 + (N−1)u_0(1−lambda)^t)`

Gate	Finding
1 Dimensions	All terms in nats; `lambda` dimensionless per period; `t` a count. **PASS.**
2 Scale type	Inherits Tables V-G1 and V-G3's conditions plus Table V-P6's on `u` decay. **CONDITIONAL, and the conditions are declared upstream.**
3 Limiting cases	`t -> infinity`: `u -> 0`, `N_eff -> 1`, `ln r -> ln(D/p_bar)` — the parallel wall, and `r` **converges** rather than diverging, which is correct: over-framing has a ceiling set by the doctrine load, not an unbounded runaway. `lambda -> 0`: `r` constant, no drift. **PASS.**
4 Fixed points	`u = 0` is globally attracting for `lambda > 0`. The crossing `t*` is not a fixed point but a threshold, and it exists if and only if `u_0 > u*`. **PASS.**
5 Identifiability	`lambda` from two clean `u` measurements; `u_0` from the channel audit; `D`, `p_bar`, `N` from the enumeration protocols. All separately obtainable. **`t*` is then a prediction with no free parameters, which makes it directly testable. PASS**, subject to the reflection problem of Section 3.4, which is not answered by parameter-counting.
6 Falsification	**The strongest test available here: preregister `t*` from measurements at `t = 0`, then measure `r` at `t*` and at `2t*`.** Kill condition: `r` fails to cross 1 within a preregistered tolerance of `t*`, or the constant-`u` baseline predicts the observed `r` series as well as the decay model.
7 Reproducibility	**UNTESTED.** Whether `lambda` recovered in one organisation transfers to another of comparable structure is unknown.

24.7.10 7.10 Requirements and failure modes at calibration scope

DD-1 Diagnose direction before magnitude. Two instruments minimum, one behavioural and one structural, with I-4 preferred as the structural one. `[direction]`

DD-2 Report `ln r` signed. Never report `Delta` bare. `[reporting]`

DD-3 Compute `r` per decision-making unit, never for the whole organisation. `[measurement]`
Test: the report contains one `ln r` per unit and no organisation-wide `D/P`.

DD-4 Aggregate `S`, intake-weighted — never aggregate `D` and `P`. The aggregate ratio overstated bankable input by a factor of 4.2553 in Section 7.7's case. `[measurement]`

DD-5 Run the infeasibility check before designing any correction. If `D/p_bar > N`, only doctrine reduction works; record `Delta_min = ln(D/p_bar) − ln N` and say so in the report. `[feasibility]`

DD-6 Derive thresholds from the organisation's own `S_min = C/eta`, not from a literature. No empirical distribution of institutional `Delta` exists. `[thresholds]`

DD-7 Demote `sigma*`- and `kappa*`-dependent instruments to orderings where the baseline is undefended, in print. `[honesty]`

DD-8 Treat high anomaly absorption as conclusive and low absorption as inconclusive. The instrument's bias runs one way. `[I–2]`

DD-9 Compute `t*` and diarise it. `t* = ln(u*/u_0)/ln(1−lambda)`, with `u* = (D/p_bar − 1)/(N−1)`. Re-measure before it, not after. `[drift]` *Test: `t*` is written down with its date of expiry.*

DD-10 Install the `u` maintenance schedule before declaring an over-framing correction complete. Otherwise `t*` simply restarts. `[u, maintenance]`

DD-11 Rank doctrine for retirement by decisions-constrained per unit of supporting evidence, not by likelihood of being wrong. `[O–4]`

DD-12 Correct under-framing by building the frame, not by cutting channels. Triage only channels with measured near-zero pairwise `u`. `[U–1, U–2]`

DD-13 Require quantified forecasts with resolution dates from any unit claiming calibration. Their absence is diagnostic. `[I–3, I–8]`

DD-14 Set the over-framing alarm at the moment an under-framing correction lands. Raising `D` shortens `t*`. `[loop closure]`

DD-15 Preregister every correction — encoder, fit window, freeze, held-out window, four baselines including the shuffled-input control, and the kill conditions [Nosek et al., 2018]. `[falsifiability]`

DD-16 No instrument or correction here may remove a person's ability to decline — to decline being measured, to decline a channel role, or to decline the dissent the protocol asks of them. `[ethics]`

FM-C1 The unsigned diagnosis. `Delta` reported without the sign of `ln r`; the correction is chosen by reflex. *Signature:* an under-framed unit receiving a diversity, breadth or collection intervention, and performing worse afterwards. *Fix:* re-diagnose; reverse the lever.

FM-C2 The cancelling aggregate. Organisation-wide `D/P` computed from summed numerators and denominators, with units distorted in opposite directions. *Damages:* measurement, in the dangerous direction — bankable input overstated 4.2553 times in Section 7.7. *Signature:* a near-calibrated headline figure and universally poor unit-level performance. *Fix:* aggregate `S`, intake-weighted, never `D` and `P`.

FM-C3 The one-off diversity project. `u` raised once, as a project, with no maintenance schedule. *Damages:* `u`, on a delay; `t*` restarts from the new `u_0`. *Signature:* a correction that works, is declared complete, and is followed by the same diagnosis two to four years later. *Fix:* the lever is a schedule, not a project.

FM-C4 The collection reflex. An under-framed organisation responding to its own confusion by collecting more. *Damages:* `r` falls further; the loop is self-reinforcing. *Signature:* instrument proliferation with flat citation; a rising backlog; a standing request for more data that is never satisfied. *Fix:* build the frame. Stop collecting against a frame that does not exist.

FM-C5 The triage overshoot. Under-framing corrected by cutting independent channels rather than by building doctrine. *Damages:* `u`, `N_eff`, and every future correction's ceiling. *Signature:* calibration achieved, followed by rapid drift into over-framing — Section 7.6's trajectory, entered deliberately. *Fix:* restrict triage to channels with measured near-zero pairwise `u`; rebuild `N_eff` before it is needed.

FM-C6 The invisible anomaly. Severe over-framing suppresses anomaly recognition, so

the anomaly audit returns a low absorption rate and reads as healthy. *Signature:* low `A`, low `sigma`, and a high surprise rate from outside the model's categories. *Measurement:* cross-check I-2 against I-1 and I-4; a low `A` with a low `sigma` is the signature, a low `A` with a healthy `sigma` is genuine. *Fix:* never read a low `A` alone as calibration.

FM-C7 The scope expansion. Over-framing addressed by raising `N` because it is the politically easiest lever. *Damages:* `C` rises, `ln r` barely moves, the underlying `u` problem is untouched. *Signature:* a remit that grows while distortion holds steady. *Fix:* O-1 or O-4. `N` is the most expensive lever and the least rewarding.

FM-C8 Corrected into the other ditch. An under-framing correction lands, `D` rises, and no one recomputes `t*`. *Damages:* the loop is not closed; over-framing arrives on schedule and is treated as a new problem. *Signature:* an organisation with a documented history of alternating between "we have no direction" and "we are not listening", every few years. *Fix:* the alarm is set at the moment of correction, not after the next failure.

24.8 8 Results: three systems worked end to end

Three constructed cases are carried end to end. **Every number below is the example's own. None is drawn from a real partnership or a real organisation and none should be cited as empirical.** The arithmetic is reproducible from the stated parameters, and that is the entire claim being made for it.

24.8.1 8.1 A two-person design practice

Two people found a small design practice — one trained as a structural engineer, one as an industrial designer — in a domain they define as `N = 12` categories: structure, material, cost, fabrication, use, maintenance, regulation, client, schedule, aesthetics, environment, disposal. They review quarterly.

Encoding. The state encoder is each partner's weighting across the twelve domain categories on a project brief, elicited independently before discussion. The interaction observable is pre-discussion disagreement across the twelve categories, per project. The order observable `Q` is durable practice assets: documented details a third party can execute, standing supplier relations, reusable specifications, precedents that survive to the next project. One quarter is one period. The cost observable `C` is coordination time, rework, the overhead of joint decision-making, and maintenance of the asset library. **`Q` is measured by third-party executability, external to the pair's own interaction volume**, which satisfies requirement S-9.

Baseline, quarter 0.

```
N     = 12       ln N = 2.4849
u     = 0.62     measured by pre-discussion disagreement over 14 briefs
eta   = 0.55     order-units per nat
delta = 0.12     per quarter
C     = 0.34     order-units per quarter
Q_0   = 0.00     new practice

N_eff = 1 + 11(0.62) = 7.8200
S_0   = ln(7.8200)   = 2.0567 nats
```

Two people with an effective count of 7.82. In degrees, `u = 0.62` is `theta = 51.9`, well inside the productive band. Decomposing per S-6: roughly `u_cross = 0.50` (two trades) and `u_div = 0.12` (each defers to the other inside the other's specialism). Checking the decomposition bound at `N = 12`:

```
f(0.50) = 1.8718     f(0.12) = 0.8416     f(0.62) = 2.0567
max = 1.8718  <=  2.0567  <=  2.7134 = sum      [bound satisfied]
```

The pair is rich in `u_cross` and thin in `u_div`. That is flagged now and it becomes the story.

The minimum viable angle.

```
S_min = C / eta = 0.34 / 0.55 = 0.6182
u_min = (e^0.6182 - 1) / 11   = 0.0778        theta_min = 16.2 degrees
N_eff at u_min = 1.8556
```

The practice's survival line is 16.2 degrees; at quarter 0 they are at 51.9 degrees.

Convergence measured. Two clean `u` measurements, quarters 0 and 8, no restoration between:

```
u_0 = 0.62      u_8 = 0.378      lambda = 1 - (0.378/0.62)^{1/8} = 0.0600 per quarter
half-life of u  = ln 2 / (-ln(1-0.06)) = 11.2 quarters
```

Their difference has a half-life of under three years. They did nothing wrong; they learned each other's trades, which is what a good cross-disciplinary pair does.

The trajectory with no restoration. Running `Q_{t+1} = (1-0.12)Q_t + 0.55 S_t - 0.34` with `u_t = 0.62(0.94)^t`:

t	u	N_eff	S	Q	h	step
0	0.6200	7.820	2.0567	0.0000	+1.0267	+0.7912
2	0.5478	7.026	1.9496	1.4579	+0.8518	+0.5574
4	0.4841	6.325	1.8445	2.4766	+0.7295	+0.3773
8	0.3779	5.157	1.6404	3.5785	+0.5973	+0.1328
10	0.3339	4.673	1.5419	3.8009	+0.5706	+0.0519
12	0.2951	4.246	1.4459	3.8719	+0.5621	−0.0094
16	0.2304	3.534	1.2625	3.6993	+0.5828	−0.0896
20	0.1799	2.979	1.0914	3.2684	+0.6345	−0.1319
24	0.1404	2.545	0.9340	2.7064	+0.7019	−0.1511

The practice peaks at quarter 12 with `Q = 3.87` and declines thereafter, holding `Q = 2.71` by quarter 24 — roughly its quarter-5 position, six years in. Now read the headroom column, which is the diagnostic point of the whole example.

> **Headroom is positive at every single step, and it is HIGHER at quarter 24 (+0.70) than at quarter 10 (+0.57).** The framework's master gauge never once signalled a problem. It could not: `h = eta ln N - delta Q - C` contains no term for `u`. Headroom measures the *budget*, not the *input*. A pair can have ample budget and nothing to spend it on.

Headroom says what could be banked; `u` says what there is to bank. A partnership dies with a healthy balance sheet. Headroom *improves* as the practice declines, because `delta Q` shrinks with `Q`: the gauge gets more encouraging exactly as the situation gets worse. That is not a defect in the gauge but a demonstration that one gauge is not a dashboard.

What the partners will say at quarter 20. They will report that the work has become repetitive, that clients are less interesting, that the market has changed, and that they are tired. All four reports will be sincere. None will mention that their effective count has fallen from 7.82 to 2.98. **Non-parallelism is the parameter that decays without any experience of decaying**, because from the inside convergence feels like competence.

The intervention. Restoring `u` to 0.62 on a schedule at or below the measured 11.2-quarter half-life, `Q` at quarter 24 reads:

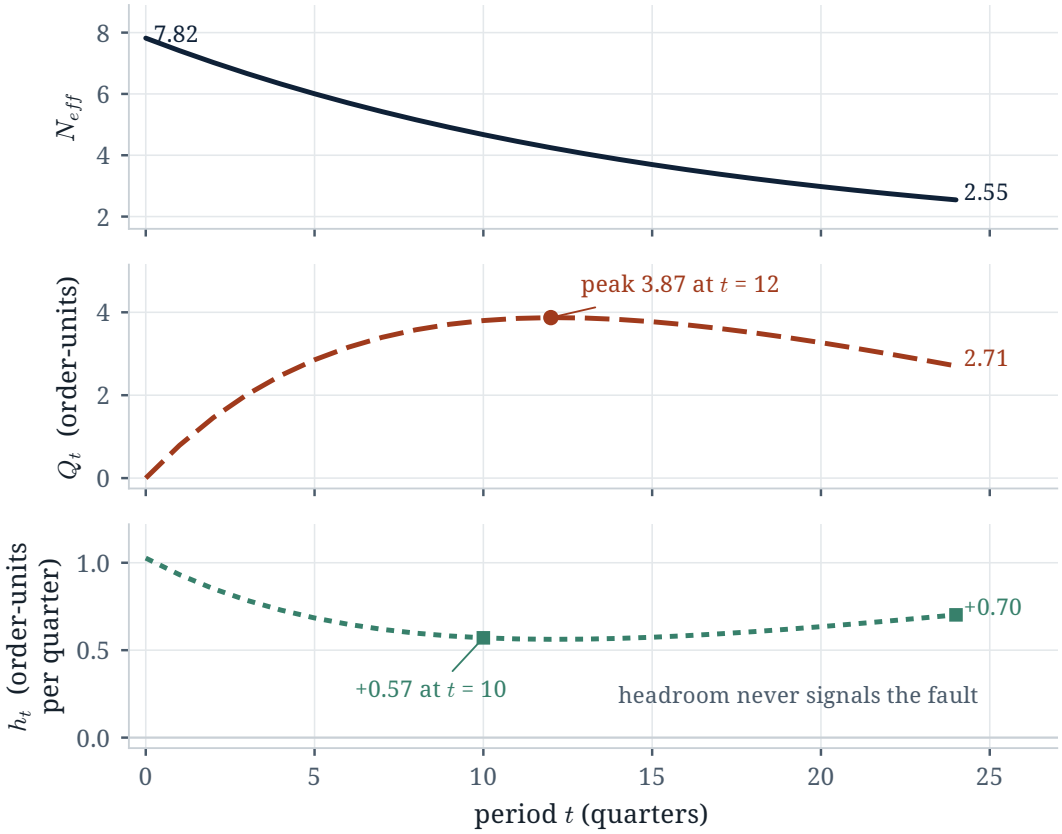


Figure 24.5: **A partnership can decline for twelve quarters while its master gauge improves.** As non-parallelism decays at 6 per cent per quarter, the pair's effective count falls from 7.82 to 2.55 and its store peaks at 3.87 in quarter 12 before falling back to 2.71 by quarter 24 — while headroom stays positive throughout and reads higher at quarter 24 (+0.70) than at quarter 10 (+0.57), because the headroom gauge contains no term for `u`. Modelled at `N = 12`, `u_0 = 0.62`, `lambda = 0.06`, `eta = 0.55`, `delta = 0.12`, `C = 0.34`, tabulated in Section 8.1; generated by `figs/src/CCP-024_fig05.py`.

```
restoration interval    Q at t=24    gain vs never
--------------------    ---------    -------------
never                     2.7064          --
every 12 quarters         4.7070       +2.0006
every  8 quarters         5.3283       +2.6219
every  6 quarters         5.6209       +2.9145
every  4 quarters         5.8991       +3.1927
every  2 quarters         6.1615       +3.4551
```

Restoring once every twelve quarters captures 58 per cent of the total available gain; a six-fold increase in effort, to two-quarterly, captures the remaining 42 per cent. Concavity again: do it at all before arguing about how often. Concretely, a restoration here is a fabricator, a regulator, a maintenance contractor or a practitioner from an unrelated trade, brought into a live working session with standing to disagree, for one project. Not a lecture. Not a conference. Someone whose position on the twelve categories neither partner can predict, inside the actual work. This addresses the deficit flagged at baseline: their `u_cross` was strong and is eroding as they learn each other's trades, while `u_div` was thin from the start. **A third trade rebuilds `u_cross`; a dissenting practitioner from one of their own trades rebuilds `u_div`. They need both, and the diagnosis says which is more urgent.**

What would falsify this. The prediction is specific and preregisterable. First, `u` measured at quarter 8 and quarter 16, with restoration events at quarters 4 and 12, will be higher than the no-restoration decay curve predicts for those quarters, 0.378 and 0.230. Second, `Q` measured by third-party executability at quarter 24 will exceed the quarter-12 peak of 3.87 under a restoration schedule of 12 quarters or shorter. Third, the fitted storage model, frozen after quarter 12, will beat all four baselines on quarters 13 to 24 held out. **The framework is falsified for this practice if the shuffled-input control — the same surplus series with its time order randomised — predicts held-out `Q` as well as the fitted model.** That would mean the timing of the pair's non-parallelism carried no information about the timing of what it built, which is precisely the claim the storage law makes.

The arithmetic is exact; the parameters are constructed. The reproducibility gate has not been run on any of it, at pair scope, by anyone, and until it has, this is a specification for how to find out rather than a report of what is known.

24.8.2 8.2 A regional authority

A regional authority of about forty staff, responsible for infrastructure decisions across a defined territory, reviewing quarterly. It commissions surveys, maintains a monitoring programme, takes submissions, and reads the technical literature.

Encoding. The state encoder is each intelligence channel's assessment vector across the twelve decision domains. The interaction observable is cross-channel error decorrelation on questions whose answers become observable within 24 months. The order observable `Q` is standing capability that survives staff turnover: verified asset records, executable procedures, precedent decisions upheld on review. One quarter is one period. The cost observable `C` is coordination, committee time, compliance, and maintenance of the standing record. **`Q` is measured by survival across turnover**, which is external to intelligence volume and satisfies requirement G-13.

Baseline.

```
N       = 12        ln N = 2.4849       twelve declared decision domains
u       = 0.14      measured by cross-channel error decorrelation
p_bar   = 1.00      one intake-unit per genuinely independent channel per quarter
D       = 5.20      doctrine load, in the same currency
E       = 1.00      environment, normalised to one unit for legibility
eta     = 0.42      order-units per nat
```

```
delta   = 0.14      per quarter
C       = 0.55      order-units per quarter

N_eff = 1 + 11(0.14) = 2.5400
F     = ln(2.5400)   = 0.9322 nats
P     = p_bar . N_eff = 2.5400
```

Forty staff, an intelligence function with twelve declared domains, and an effective count of 2.54. The ratio of audited channels to headcount is `2.54 / 40 = 0.0635`: the authority's audited intake is worth about six per cent of its headcount.

The ratio, computed two ways.

```
direct:       ln r = ln D - ln P             = 1.6487 - 0.9322 = 0.7165
via the join: ln r = ln D - ln p_bar - F     = 1.6487 - 0.0000 - 0.9322 = 0.7165   AGREE
```

The over-determination check passes, so the three measurements are mutually consistent.

```
r     = e^{0.7165} = 2.0472
Delta = 0.7165 nats,  sign POSITIVE  ->  OVER-FRAMED
M     = r . E = 2.0472
```

The authority constructs a situation about twice as large as its environment. It sees two units of structure where one exists. Behaviourally: it is confident, it is fast, and it is wrong in a direction its own model predicts.

The infeasibility check.

```
D / p_bar = 5.20 / 1.00 = 5.20        N = 12        5.20 <= 12  ->  FEASIBLE
u* = (5.20 - 1) / 11 = 0.3818         N_eff* = 5.20      F* = 1.6487
```

The authority *can* be calibrated by geometry alone. It needs `u = 0.3818`, which is `theta = 38.2` degrees of mean non-parallelism across its channels, against a measured 0.14 at `theta = 22.0` degrees. That is a large but not impossible move.

The operating envelope.

```
S_min = C / eta = 0.55 / 0.42 = 1.3095 nats
S_min / ln N    = 1.3095 / 2.4849 = 0.5270
band: 0.5270 <= min(r, 1/r)   ->   0.527 <= r <= 1.898       Delta <= 0.6406
```

The authority's `r = 2.0472` is outside its own operating envelope — not marginally, `Delta = 0.7165` against a limit of 0.6406. Confirming with the contact postulate and the storage law:

```
S       = ln N . min(r, 1/r) = 2.4849 x (1/2.0472) = 1.2138 nats
eta.S   = 0.42 x 1.2138 = 0.5098
Q_inf   = (0.5098 - 0.55) / 0.14 = -0.2872
```

A negative equilibrium. The storage law does not hide this failure mode: as modelled, the authority's standing capability erodes to nothing with a time constant of `1/delta = 7.1` quarters, and the erosion will be attributed to turnover, budget and the declining quality of submissions.

Lever A — raise `u` (recruit across lineages, rotate, commission from unaligned sources, stand up a genuinely independent channel):

u	N_eff	P	r	Delta	S	eta.S	Q_inf
0.1400	2.540	2.540	2.0472	0.7165	1.2138	0.5098	−0.2872
0.2000	3.200	3.200	1.6250	0.4855	1.5292	0.6423	+0.6589
0.2400	3.640	3.640	1.4286	0.3567	1.7394	0.7306	+1.2897
0.3000	4.300	4.300	1.2093	0.1900	2.0548	0.8630	+2.2359
0.3818	5.200	5.200	1.0000	0.0000	2.4848	1.0436	+3.5259
0.4500	5.950	5.950	0.8739	0.1347	2.1717	0.9121	+2.5865

Two things to read. **The authority crosses back into its operating envelope at `u = 0.158`** — the value at which `N_eff = 5.20/1.898 = 2.7397` — well short of full calibration. It does not need to be perfect; it needs to be inside the band. And **the last row over-corrects.** At `u = 0.45` the authority passes calibration and begins under-framing, and `Q_inf` falls from 3.53 to 2.59. Diversity past the calibration point is a cost, exactly as Section 6.8 predicted.

Lever B — cut `D` (retire standing assumptions, sunset guidance, shed doctrine):

D	r	Delta	S	Q_inf
5.20	2.0472	0.7165	1.2138	−0.2872
4.50	1.7717	0.5719	1.4026	+0.2792
4.00	1.5748	0.4541	1.5779	+0.8052
3.50	1.3780	0.3206	1.8033	+1.4814
3.00	1.1811	0.1664	2.1039	+2.3831
2.54	1.0000	0.0000	2.4849	+3.5261

The authority re-enters its envelope at `D = 4.82`, since `2.5400 x 1.898 = 4.8209` — a reduction of about seven per cent in standing doctrine. **Both levers reach the same destination and the arithmetic is symmetric, because the contact postulate is symmetric in `ln r`.** Which to pull is not a question the formula answers; it is a question of cost.

Headroom, and the trap at the end.

Q	h = eta.lnN − delta.Q − C
0.0	+0.4937
1.0	+0.3537
2.0	+0.2137
3.0	+0.0737
4.0	−0.0663
5.0	−0.2063

`h = 0 exactly at Q = 3.5261`

Compare that to Section 6.6's derived result: `Q_inf` at perfect calibration is `(eta ln N − C)/delta = 3.5261`, and `h = 0` at `Q = 3.5261`. **They are the same number, as the algebra requires.** If the authority calibrates fully and holds there, it arrives at exactly zero headroom and can thereafter maintain but never build.

The recommendation the model supports is therefore not "calibrate fully". It is: get inside the envelope, stop short of the fixed point, and hold `Q` below `Q_inf` deliberately. Target `u` near 0.24, which gives `Delta = 0.36`, `Q_inf = 1.29`, and residual headroom at that store of `0.4937 − 0.14(1.29) = +0.31`. **Slightly distorted and solvent beats perfectly calibrated and saturated**, which is a direct consequence of the arithmetic rather than of the equation's rhetoric.

What would falsify this. `r` measured quarterly for eight quarters; held-out prediction of decision quality, assessed at 24-month outcome observation, for quarters 9 to 16; four baselines — persistence, `E` alone, matched autoregressive, and the shuffled-`r` control; and four preregistered kill conditions, of which the cheapest is a fitted exponent `k` in `S = ln N . e^{−k.Delta}`

significantly different from 1, or significantly different between the `r > 1` and `r < 1` branches. **That last condition can falsify one postulate while leaving the framework intact, and it costs a decision log and a forecast record.**

24.8.3 8.3 A two-site agency

An agency with two decision-making units: a headquarters that sets policy and a field operation that executes and observes. Both work the same `N = 12` decision domains. Quarterly periods. Storage parameters shared: `eta = 0.42`, `delta = 0.14`, `C = 0.55` per unit, so `S_min = C/eta = 1.3095` nats and the operating envelope is `min(r, 1/r) >= 1.3095/2.4849 = 0.5270`, that is `0.527 <= r <= 1.898`.

Baseline, per unit. Channel audit and doctrine enumeration, `p_bar = 1.0` in both units:

```
unit     u       N_eff = P    D       ln r       r        Delta    S = lnN.min(r,1/r)
-----   -----   -----------   -----   --------   ------   ------   ------------------
HQ      0.11       2.210      6.80   +1.1239   3.0769   1.1239         0.8076
Field   0.62       7.820      1.15   −1.9169   0.1471   1.9169         0.3654
```

The headquarters is over-framed at `Delta = 1.12`. The field is under-framed at `Delta = 1.92`. Both are outside the envelope and both would have equilibrium stores well below zero. The aggregate, for contrast only:

```
D_total = 7.95    P_total = 10.030    r_agg = 0.7926    Delta_agg = 0.2324
```

The agency's headline figure says it is nearly calibrated. It is not: it has one unit at 3.08 and one at 0.147. Computing bankable input from the aggregate gives 1.9696 nats against a true intake-weighted 0.4629, an overstatement of 4.2553 times.

Corroboration. Reading the instruments per unit:

```
                              HQ                      FIELD
I−1  revision rate            low                     high
I−2  anomaly absorption A     high, conclusive        low, inconclusive
I−3  calibration              overconfident           no quantified forecasts
I−5  intake per decision      low                     high
I−6  backlog                  flat                    rising
I−8  forecasts present        yes                     no
```

Two instruments agree in each unit, one behavioural and one structural. The field's low `A` is read as inconclusive per DD-8, and the cross-check clears it: the field's `sigma` is high, so its low absorption is genuine and not suppression.

Infeasibility check.

```
HQ:     D/p_bar = 6.80   vs   N = 12    ->  FEASIBLE.  u* = (6.80−1)/11 = 0.5273
Field:  D/p_bar = 1.15   vs   N = 12    ->  FEASIBLE, but the fix runs the other way.
```

The headquarters can be calibrated by geometry, but it needs `u = 0.5273` from a measured 0.11 — a factor of 4.8. That is a very large move and prices doctrine reduction into the conversation immediately.

The obvious move, and why it is insufficient. The agency's instinct will be redistribution: move doctrine from HQ to the field. Under strict conservation of `D` and `P`:

```
D_total = 7.95        P_total = 10.030

Best achievable by redistributing D alone (P fixed):
```

```
    D_HQ = 2.210, D_field = 5.740
    ->  r_HQ = 1.0000,  r_field = 0.7340,  mean Delta = 0.1546

Best achievable by redistributing both (any split):
    both units at r = D_total/P_total = 0.7926,  Delta = 0.2324 each
```

Redistribution gets both units inside the envelope. It cannot calibrate both, because the totals do not match: the agency is doctrine-poor overall by `10.030 - 7.95 = 2.080` units, 26.2 per cent of its standing doctrine. That is a finding invisible without per-unit measurement. The agency does not have a distribution problem alone. It has a distribution problem *and* a shortfall, and the shortfall is in doctrine, which is the opposite of what the headquarters diagnosis on its own would suggest.

The intervention. *Stage 1 — redistribute.* Move 4.59 units of doctrine from HQ to the field: HQ to `D = 2.21`, field to `D = 5.74`.

```
HQ:     r = 2.210/2.210 = 1.0000   Delta = 0.0000   S = 2.4849
Field:  r = 5.740/7.820 = 0.7340   Delta = 0.3092   S = 1.8239
```

Both are inside `0.527 <= r <= 1.898`. Intake-weighted `S` rises from 0.4629 to `(2.210 x 2.4849 + 7.820 x 1.8239)/10.030 = 1.9695` — a factor of 4.25. What "moving doctrine" means concretely: HQ retires standing guidance and the field adopts explicit standing assumptions and publishes forecasts. **It is not a transfer of documents. It is HQ ceasing to determine what the field will find, and the field starting to commit to what it expects.**

Stage 2 — the shortfall. The agency still needs 2.080 units of net new doctrine to bring both units to calibration. This is built, not moved: explicit models, committed hypotheses, published forecasts with resolution dates, and closed lines of enquiry.

Stage 3 — the headquarters' `u`. HQ sits at `u = 0.11`, `N_eff = 2.21`. Even after Stage 1 it is calibrated at a very low effective count, which means it has almost no margin: any future growth in `D` puts it straight back over the line. Raise `u` toward 0.30 (`N_eff = 4.30`) as insurance, on a maintenance schedule.

The alarm. The field has just received doctrine. Compute its crossing date now. Post-Stage-1 field: `D = 5.74`, `p_bar = 1.0`, `u_0 = 0.62`, and suppose `lambda = 0.06` per quarter measured per S-11.

```
u* = (5.74 - 1)/11 = 0.4309
t* = ln(0.4309/0.62) / ln(0.94) = 5.88 quarters
```

The field crosses from under-framed to over-framed in under six quarters, and the agency should diarise re-measurement at quarter 4, not quarter 8. Without the alarm, the field's correction becomes the field's next pathology.

What would falsify this. `ln r` measured per unit quarterly for eight quarters after the intervention; held-out prediction of unit-level decision quality, assessed at 24-month outcome observation, for quarters 9 to 16; four baselines including the shuffled-`ln r` control; and **the specific, cheap, decisive test: the field's predicted crossing at `t* = 5.88` quarters.** `t*` is computed from three independently measured quantities with **no free parameters**. Kill condition: the field's `r` fails to cross 1 within a preregistered tolerance around quarter 6, or a constant-`u` model predicts the observed `r` series as well as the decay model.

That last test is the one to run first. It is a dated, parameter-free prediction about an observable quantity, checkable with a channel audit and a decision log — which is the form of risky prediction the social sciences most often avoid [Meehl, 1967]. It does not, however, escape the identification problem of Section 3.4: a crossing observed on schedule is consistent with the mechanism and is also consistent with any correlated unobservable that moves on the same timetable.

24.9 9 Discussion

24.9.1 9.1 What follows

The angle is a managed parameter and it is usually unmanaged. The same two people are worth an effective count of 6.5 or 1.55 in the same domain depending only on the angle between their states, and ability moves neither number. Where a pair or an organisation has a selection process and no maintenance process, it is managing the wrong variable.

A diagnostic built on assent will miss half the failures it is meant to catch. The opposition identity means an organisation can sound like a debating society and read `N_eff = 1`. Every instrument here that bears on non-parallelism is therefore stated on predictability, not on agreement or disagreement.

Two ceilings coincide, and the coincidence is informative. The perceptual correction available from interaction geometry is `ln N`, and the sustainable growth available to the storage law is bounded by `eta ln N`. Both say that geometry is bounded and that past the bound what is carried must change, not who carries it. That is a statement about the shape of the model rather than a discovery about institutions, and it is labelled as such.

Calibration has a price and the price is capacity. At `r = 1` the input is at the ceiling, the store is at its maximum, and the entire budget goes to maintenance. An organisation that wants to start something must either shed store deliberately or accept distortion. This is the framework's success-consumes-its-own-headroom thesis reached from a second direction, and it is the same object as organisational slack [Cyert and March, 1963].

The unit of measurement is the decision-making unit and never the organisation. Distortion does not average; the aggregate is a mediant, not a mean, and it can read as calibrated while every unit is in a ditch.

The correction for one pathology seeds the other. Doctrine does not decay; non-parallelism does. Every successful under-framing correction shortens the crossing date, and an organisation that does not compute the date acquires an alternating history it will describe as bad luck.

24.9.2 9.2 What does not follow

Nothing here shows that any pair or any organisation behaves this way. Every reproducibility gate reads untested. The three worked systems are constructed.

The group-scope magnitudes are not measurements. They are demoted at the scale-type gate and reported as illustrative constructions. The 0.53 to 1.90 envelope, the `u*` requirements and the `Delta_min` floors are structure exhibited, not quantity measured.

No claim is made that observing an organisation could identify the mechanism. Section 3.4 gives the reason: the specification is reflexive and the endogenous effect is not identified from observational data [Manski, 1993].

Monotonicity of surplus in `u` is not established outside the model. The literature contains an inverted U [Uzzi and Spiro, 2005], short-run costs of diversity [Mannix and Neale, 2005, Putnam, 2007], and destabilisation of randomly assembled systems by connectance [May, 1972]. The model's monotone claim holds inside the model and its domain of application is not known.

The logarithm is not established against its rivals. A sub-linear power law [Latané, 1981] and an `n log n` law [Briscoe et al., 2006] fit the same qualitative facts, and nothing here discriminates.

`Delta_R`, the return difference, is not a magnitude in most settings. Its scale-type gate binds and it demotes to an ordering unless a magnitude-estimation procedure is in place.

The framework does not license arranging anyone's relationships. The scope constraint of Section 4.16 is a constraint on use and not a caveat.

24.9.3 9.3 Relation to the surrounding literature

The pair-scope material sits alongside three literatures. The **diversity-and-performance** literature supplies both the supporting result [Hong and Page, 2004, Woolley et al., 2010] and the objections that make citing it responsible [Thompson, 2014, Harrison and Klein, 2007, Uzzi and Spiro, 2005]. **Portfolio theory** stated this framework's central point in 1952 with the mathematics attached: diversification benefit depends on covariance, not count, so N uncorrelated holdings and N perfectly correlated holdings are not the same N [Markowitz, 1952]. That is the opposition identity in a different register, and the framework's own bloc-collapse arithmetic reaches it a third time [Curzi, 2026,?]. The measured collapse of correlation structure under stress [Longin and Solnik, 2001] is the same warning as the decay of u: the effective count falls exactly when it is needed. **Ecology** supplies both the demonstration that effective count buys stability [Schindler et al., 2010] and the objection that diversity–stability correlations can arise from statistical averaging alone with no mechanism [Doak et al., 1998], which the framework must answer and has not.

The group-scope material sits alongside **opinion dynamics** and **organisational learning**. The reflexive map has an older form in which beliefs are updated by a row-stochastic influence matrix and consensus is the dominant left eigenvector [DeGroot, 1974]; the convergence conditions that model states explicitly are conditions this framework's fixed point assumes silently, and the stubbornness term that prevents collapse [Friedkin and Johnsen, 1990] is formally this framework's leak. The manufacturing result is the organisational-learning result [March, 1991] with a crossing date attached. The literature also supplies a hard limit on what a reflexive fixed point can be asked to predict: social influence raised both inequality and *unpredictability* in a controlled experiment, with the same inputs producing different winners in parallel worlds [Salganik et al., 2006].

Two negative results bound the whole enterprise and are cited here rather than left for a reviewer. The predictability ceiling on social outcomes is low: with fifteen years of high-quality longitudinal data and 160 teams, life outcomes were barely predictable, and the best machine learning barely beat a four-variable regression [Salganik et al., 2020, Hofman et al., 2017]. And the standing charge against opinion-dynamics modelling — many models, negligible empirical discrimination between them — has been made twice, eleven years apart, with the situation unchanged [Sobkowicz, 2020]. **The question that literature asks is what observation would distinguish this framework's laws from a neighbouring model that fits equally well, and the honest answer for the group-scope material in this paper is: the parameter-free crossing prediction t∗, and nothing else.** The field's own indictment of importing physics without the accompanying statistical care is Gallegati et al. [2006], with the reply [McCauley, 2006]; both should be read as a checklist against this paper rather than as background.

24.10 10 Verification

The results of this paper compound from ten postulates. Each is atomic, each is checkable on its own, and each carries its status. **Where a postulate could be dropped and a result still follow, it has been dropped**; what remains is what the chain actually consumes. **The chain below ends in results whose steps are largely DEFINITION and POSTULATED, and that is stated here rather than left for the reader to notice.** A result honestly labelled as following from two postulates is worth more than the same result implied to be measured.

24.10.1 10.1 Dimensional audit

`[Q]` denotes the dimension of stored order, `[I]` the dimension of intake, and `[T]` the dimension of time. The period is the unit of `[T]`, and every rate below is *per period*, which is why no explicit `[T^-1]` appears once the equations are written in discrete steps.

Symbol	Dimension	Unit	Status
`x`, `y`	`[1]`	none	**defined** unit vectors in an `N`-category space
`u`	`[1]`	none	**defined** dimensionless: `u = 1 - (x.y)^2`, in `[0,1]`
`theta`	`[1]`	degrees or radians	**derived** from `u = sin^2(theta)`
`N`	`[1]`	a count of categories	**defined** by enumeration of the declared domain
`N_eff`	`[1]`	a count of effective channels	**derived** `1 + (N−1)u`; **not a Hill number**
`F`, `S`	`[1]`	nats	**derived** dimensionless: a logarithm of a count
`Q`	`[Q]`	proxy units (documented details, executable procedures, precedents)	**measured**, in whatever unit the declared order proxy counts
`eta`	`[Q]` per nat per step	proxy units per nat per period	**measured**; *not* dimensionless
`delta`	`[1]`	per period	**measured** dimensionless fraction per step
`C`	`[Q]` per step	proxy units per period	**measured**
`h`	`[Q]` per step	proxy units per period	**derived**
`Q_inf`	`[Q]`	proxy units	**derived**: `[Q]/step` divided by `[1]/step` returns `[Q]`
`lambda`	`[1]`	per period	**measured** dimensionless fraction per step
`a`	`[1]`	nats	**derived** log of a ratio of two identically-united quantities
`M`, `M'`	undeclared	undeclared	**the source is silent.** The return difference `Delta_R = |ln(M'/M)|` requires only that `M` and `M'` share units; the paper never says what they are
`p_bar`	`[I]` per channel per step	intake-units per channel per period	**measured** by tally
`P`	`[I]` per step	intake-units per period	**derived** `p_bar . N_eff`

Symbol	Dimension	Unit	Status
`D`	`[I]` per step	intake-units per period, by the commensurate-units convention	**measured** by doctrine enumeration, in the convention's units
`E`	`[I]` per step	intake-units per period	**measured**, and in practice rarely measured at all
`r`	`[1]`	none	**derived** `D/P`, dimensionless **only because the convention says so**
`Delta`	`[1]`	nats	**derived** `|ln r|`
`kappa`	`[Q]` per step per `[Q]`	per period	**measured**, correction load per unit stored error
`sigma`, `A`	`[1]`	proportions	**measured** counts over counts

Does the storage law balance? `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t` balances if and only if every term on the right carries `[Q]`. The first does, since `delta` is dimensionless. The third does by construction. The second does only if `eta S_t` carries `[Q]`, and since `S_t` is a dimensionless count in nats, **`eta` must carry `[Q]` per nat per step**. It is therefore not dimensionless, and it must be reported with its units. Consequently `eta ln N` carries `[Q]` per step and is commensurable with `delta Q_t` and with `C_t`, which is what makes the headroom gauge a well-formed subtraction. A headroom calculation in which `C` is denominated in hours or currency while `Q` is denominated in documented procedures is dimensionally malformed however tidy it looks.

Does the moment equation balance? Only under the commensurate-units convention. `D/P` is dimensionless if and only if `D` and `P` are expressed in the same currency, and `M` then inherits `E`'s units. **Without the convention the formula is dimensionally void**, and that is stated in Table V-G1 rather than assumed away. The convention is a **modelling choice**, not a measurement, and Section 6.12 records that a count supplies a unit within one encoder and no linkage between encoders.

Where the framework has no operational unit, it is said here. `Q` is stored order and stored order has no natural unit: it is counted in whatever the declared proxy counts, and two systems' `Q` values are not comparable unless they declare the same proxy. `M` and `M'` in the return difference are **never given a dimension by the source at all**; only their ratio is used, so the formula is well posed, but any claim that `Delta_R` is commensurable with `F` or with `Delta` rests on a unit the paper does not supply. That silence is recorded rather than quietly filled.

24.10.2 10.2 The postulates

P1 — Non-parallelism is `u = 1 − (x.y)^2` for unit vectors `x, y`. *Status:* **DEFINITION.** It carries no empirical weight on its own; the claim that a person or a channel *is* such a vector is the dyadic identification, badged separately at Section 4.1.

P2 — The surplus functional is `f(u) = ln(1 + (N−1)u)`, with `N_eff = e^f = 1 + (N−1)u` and `0 <= f <= ln N`. *Status:* **CITED** — established from four axioms, one of which is an explicit design axiom, in Curzi [2026]. The functional form is not established against its rivals [Latané, 1981, Briscoe et al., 2006].

P3 — For unit vectors, `x.y = cos(theta)`. *Status:* **ARITHMETIC.** Hence `u = 1 − cos^2(theta) = sin^2(theta)`.

P4 — The storage law is `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t`, with `delta`, `eta`, `C` not depending on `Q`. *Status:* **DEFINITION** — the model form adopted by the framework.

P5 — A map `T(Q) = aQ + x` with `0 <= a < 1` on a complete space has exactly one fixed point, reached from any start. *Status:* **CITED** [Banach, 1922], with the elementary difference-equation route reaching the same condition [Elaydi, 2005].

P6 — Headroom is `h = eta ln N − delta Q − C`. *Status:* **DEFINITION.**

P7 — Collective perception is `P = p_bar . N_eff`. *Status:* **POSTULATED** — a modelling choice with no external support. It is the point at which a hundred parallel channels become one, and it is not identified from observational data (Section 3.4).

P8 — The moment is `M = (D/P).E`, so `r = D/P` and `Delta = |ln r|`. *Status:* **POSTULATED** — this framework's own, with no external literature for either the relation or the reading of `r` as a calibration ratio.

P9 — Bankable input is `S = ln N . e^{−Delta}`. *Status:* **POSTULATED.** Its symmetry in the sign of `ln r` is its substantive and separately falsifiable commitment.

P10 — Non-parallelism decays geometrically, `u_t = u_0 (1−lambda)^t`. *Status:* **POSTULATED** — a modelling choice; the exponent is defined by the model, not estimated from data.

24.10.3 10.3 The compounding chain

P1 + P3	-> L1	u = sin^2(theta), and u is a function of |x.y| alone.	ARITHMETIC
L1 + P2	-> R1	rho = x.y = +1 and rho = −1 both give u = 0, hence N_eff = 1 and F = 0. Opposition and agreement are the same reading.	ARITHMETIC
P2	-> L2	0 <= F <= ln N, with F = 0 iff u = 0 and F = ln N iff u = 1.	CITED (ccp:CCP−A02)
P4 + P5	-> L3	for delta > 0 the fixed point is unique and equals Q_inf = (eta F(u) − C)/delta.	CITED (banach1922operations)
L3	-> R2	Q_inf >= 0 iff F(u) >= C/eta, that is u >= u_min = (e^{C/eta} − 1)/(N−1).	ARITHMETIC
P7 + P8	-> R3	ln r = ln D − ln P = ln D − ln p_bar − F.	ARITHMETIC
R3 + L2	-> R3a	the reduction in ln r available from geometry is at most ln N; r = 1 is feasible iff D/p_bar <= N, and otherwise Delta_min = ln(D/p_bar) − ln N.	ARITHMETIC
P9 at Delta=0	-> L4	S = ln N, the ceiling.	ARITHMETIC
L4 + L3	-> L5	Q_inf = (eta ln N − C)/delta.	ARITHMETIC
L5 + P6	-> R4	h = eta ln N − delta(eta ln N − C)/delta − C = 0 exactly.	ARITHMETIC
P8 on two units	-> R5	r_agg = (D1+D2)/(P1+P2) is a mediant of r1 and r2, so it lies strictly between them; where r1 > 1 > r2 it can lie arbitrarily close to 1 while both units are arbitrarily far from it.	ARITHMETIC
P4 with two directional eta	-> R6	holding the mean conversion fixed, the sum of the two	

```
                    private steady states is constant in
                    a = |ln(eta_AB/eta_BA)| while the two values separate.  ARITHMETIC

P10 + R3      -> R7 ln r rises monotonically in t; the crossing of r = 1
                    occurs at t* = ln(u*/u_0)/ln(1-lambda) with
                    u* = (D/p_bar - 1)/(N-1).                              ARITHMETIC
```

Every line is checkable from its own stated inputs. **R1 and R2 consume one definition, one arithmetic fact and one cited functional form, and nothing else.** They are the paper's strongest results in the epistemic sense, and both are about pairs.

R3, R3a, R4, R5 and R7 all consume P7, P8 or P9, each of which is POSTULATED with no external support. The algebra is exact and the postulates are not established. In particular: **R4, the zero-headroom identity, is an arithmetic consequence of two postulates and a definition. It is not a finding about organisations and must not be quoted as one.** R5 is the least exposed of the group-scope results, because a mediant is a mediant whatever `D` and `P` mean; it survives even if P8's interpretation is wrong, provided only that `r` is computed as a ratio of two additively aggregable quantities.

What the chain does not establish. It does not establish that any pair obeys P1 or P4, or that any organisation obeys P7, P8 or P9. Those are the framework's postulates, they are not verified here, and the one preregistered confirmatory test of the storage law reported in this corpus did not meet its own declared margin [Curzi, 2026i]. A chain whose steps are ARITHMETIC, DEFINITION and CITED delivers the mathematics with certainty and delivers nothing whatever about institutions.

24.10.4 10.4 A worked numerical instance

The smallest set of checks that exercises the whole chain. Each can be redone by hand or in a line of arithmetic.

Check 1 (R1). `rho = −1` gives `u = 1 − (−1)^2 = 1 − 1 = 0`. At `N = 12`, `N_eff = 1 + 11(0) = 1` and `F = ln 1 = 0`. `rho = +1` gives the same three numbers. For contrast, `rho = −0.95` gives `u = 1 − 0.9025 = 0.0975`, `N_eff = 1 + 11(0.0975) = 2.0725` and `F = ln 2.0725 = 0.7288`, identical to `rho = +0.95`. **ARITHMETIC.**

Check 2 (R2). With `eta = 0.55`, `C = 0.34`, `N = 12`: `C/eta = 0.34/0.55 = 0.618182`; `e^0.618182 = 1.855555`; `(1.855555 − 1)/11 = 0.077778`; `N_eff = 1 + 11(0.077778) = 1.855555`, which is `e^{C/eta}` as it must be. In degrees, `theta = arcsin(sqrt(0.077778)) = arcsin(0.278886) = 16.20`. **ARITHMETIC.**

Check 3 (R3, the over-determination check). With `D = 5.20`, `p_bar = 1.00`, `u = 0.14`, `N = 12`: `N_eff = 1 + 11(0.14) = 2.5400`; `F = ln 2.5400 = 0.932164`; `P = 2.5400`. Computing `ln r` directly, `ln 5.20 − ln 2.5400 = 1.648659 − 0.932164 = 0.716495`. Computing it via the join, `1.648659 − ln 1.00 − 0.932164 = 1.648659 − 0 − 0.932164 = 0.716495`. **The two agree to every digit printed, which is the diagnostic the requirement G-6 asks for.** `r = e^{0.716495} = 2.047244`, which is `5.20/2.5400` exactly. **ARITHMETIC.**

Check 4 (R4). With `eta = 0.42`, `delta = 0.14`, `C = 0.55`, `N = 12` so `ln N = 2.484907`: `eta ln N = 1.043661`; `Q_inf = (1.043661 − 0.55)/0.14 = 3.526150`; `h = 1.043661 − 0.14(3.526150) − 0.55 = 1.043661 − 0.493661 − 0.55 = 0.000000`. **The identity is exact and the cancellation is visible in one line. ARITHMETIC.**

Check 5 (R5). Two units at `N = 12`, `p_bar = 1.0`. HQ: `u = 0.11`, `P_1 = 1 + 11(0.11) = 2.21`, `D_1 = 6.80`, `r_1 = 3.076923`, `ln r_1 = +1.123930`. Field: `u = 0.62`, `P_2 = 7.82`, `D_2 = 1.15`, `r_2 = 0.147059`, `ln r_2 = −1.916923`. Aggregate: `(6.80 + 1.15)/(2.21 + 7.82) = 7.95/10.03 = 0.792622`, `ln r_agg = −0.232434`. Bankable input from the aggregate: `2.484907 x 0.792622 = 1.969608`. Per unit and intake-weighted: `(2.21 x 0.807633 + 7.82 x 0.365427)/10.03 = (1.784869 + 2.857639)/10.03 = 0.462863`. Ratio: `1.969608/0.462863 = 4.2553`. **ARITHMETIC.**

Check 6 (R6). With `delta = 0.15`, `C = 0.25`, `u = 0.45`, `N = 12`: `S = ln(1 + 11(0.45)) = ln 5.95 = 1.783391`. At `eta_AB = 0.90`, `eta_BA = 0.10`: `Q_A = (0.10 x 1.783391 − 0.25)/0.15 = (0.178339`

– 0.25)/0.15 = −0.477740 and Q_B = (0.90 x 1.783391 – 0.25)/0.15 = (1.605052 – 0.25)/0.15 = +9.033680; the sum is 8.555940. At eta_AB = eta_BA = 0.50, each partner is at (0.891696 – 0.25)/0.15 = 4.277970 and the sum is 8.555940. **The same sum, and one negative private store.** a = |ln(0.90/0.10)| = ln 9 = 2.197225. **ARITHMETIC.**

Check 7 (R7). With D/p_bar = 4.0, N = 12, u_0 = 0.55, lambda = 0.08: u* = (4.0 – 1)/11 = 0.272727; u*/u_0 = 0.495868; ln 0.495868 = −0.701290; ln 0.92 = −0.083382; t* = 8.4125 periods. At t = 8, u = 0.55(0.92)^8 = 0.283281 and N_eff = 4.116 so r = 0.9718 < 1; at t = 9, u = 0.260618, N_eff = 3.867, r = 1.0344 > 1. **The crossing sits between periods 8 and 9, as t* = 8.41 requires. ARITHMETIC.**

24.10.5 10.5 Limits, asymptotics and the model's valid domain

u –> 0. N_eff –> 1, F –> 0, Q_inf –> –C/delta < 0, and at group scope P –> p_bar and ln r –> ln(D/p_bar), its maximum for the given doctrine load. **r converges rather than diverging**, which is correct: over-framing has a ceiling set by the doctrine load, not an unbounded runaway.

u –> 1. N_eff –> N, F –> ln N, the ceiling is attained and no further gain is available from geometry. At group scope, if D/p_bar < N the organisation over-corrects past calibration into under-framing.

N –> 1. ln N = 0 and F = 0 for every u. **A single-category domain has zero available surplus, not a small amount**, so Q_inf = –C/delta <= 0 and headroom is negative at any positive stock. This is why a bloc-collapse operation that reduces an entity map to one effective perimeter yields a surplus term of exactly zero [Curzi, 2026].

N –> infinity. ln N grows without bound but only logarithmically, and the marginal gain ln N – ln(N–1) behaves as 1/N. Against this, coordination cost grows with the number of pairs [Brooks, 1975, Conway, 1968], which is quadratic. **The two curves cross, and beyond the crossing additional categories cost more than the ceiling they buy.**

delta –> 0. No fixed point exists at all: the map is an isometry, the geometric series does not converge, and the bound degrades to Q_0 + t.eta.ln N [Curzi, 2026a]. **The limit is singular**: at exactly delta = 0 the fixed point does not exist rather than being infinite. This value is outside the model's valid domain.

delta –> 1. The system has no memory and Q_{t+1} –> eta S_t – C_t, an instantaneous transducer. At delta = 1 the storage law is not a storage law and any claim about accumulation is vacuous.

eta –> 0. Q_inf = –C/delta <= 0: the pleasant partnership that builds nothing, and no interaction geometry of any quality changes the outcome.

lambda –> 0. u is constant — a pair or an organisation that works together and never influences itself, which is the mark of a collaboration that is not happening. **lambda –> 1:** collapse to u = 0 in one period, total capture.

r –> 0 and r –> infinity. Delta diverges in both directions and S –> 0 in both. The two divergences are distinct pathologies mapped to the same magnitude, which is the measure's one real weakness and the reason the sign must always be reported.

P –> 0, D –> 0, E –> 0. Worked in Section 6.5. Each names a real institutional condition; E –> 0 sends M –> 0 for any r at all, and is the only limit that is not diagnosable from r.

Stated domain. The results are claimed for u in [0,1]; N >= 2; delta in (0,1); eta > 0; lambda in (0,1); p_bar > 0; D > 0; and D/p_bar <= N where calibration is claimed to be reachable. Outside that domain the equations still have solutions and those solutions should not be interpreted as statements about a pair or an institution.

24.10.6 10.6 Invariance

The scale-invariance of r is real and is worth stating as such. Delta = |ln r| survives any rescaling that multiplies D and P by a common positive factor, and this is why Delta is the

reportable quantity while D and P are not. What breaks it is a rescaling that touches only one of the two, which is exactly what happens when the encoder for doctrine differs from the encoder for intake — the failure Section 6.12 records at the scale-type gate.

The storage law is invariant in form under a change of order unit provided eta and C rescale with Q while delta, N and u are unchanged. **delta, u, N and lambda are therefore the unit-free quantities and Q, eta and C are the unit-carrying ones.** Two systems can be compared on the first group directly; they cannot be compared on the second unless they share a proxy.

How far the invariance has been exercised. Within this paper: three constructed systems at two scopes, all stipulated rather than observed. **That is three systems, one author, and no independent replication.** The invariance is a claim about form that has been exercised, not an invariance that has been established.

24.10.7 10.7 Scale type, model comparison, and the precision of the numbers

Scale type. u, Q, D, P and r are treated as ratio-scale quantities throughout — differences, ratios and logarithms are all taken. **The conjoint-measurement conditions that would license that treatment have not been tested** [Luce and Tukey, 1964, Krantz et al., 1971], and Michell [1997] argues that psychology systematically fails to test whether its attributes are quantitative at all before measuring them. **That is the strongest single objection to this framework and it is conceded rather than answered here.** The treatment of u and of Q as quantities is a stipulation. At group scope the framework has already run its own audit and demoted the magnitude of r and of Delta to ordinal pending a Rasch demonstration [Curzi, 2026, Rasch, 1960], which is the right instinct applied consistently rather than once. Delta_R demotes for the same reason wherever M is not ratio-scaled.

Model comparison — NOT MET, and the row is reported rather than omitted. The rigour standard for this corpus requires, wherever a fit is reported: out-of-sample performance, a matched null with the same number of free parameters, an information criterion so that the parameter count is paid for, an explicit count of free parameters, and effect sizes with intervals. **This paper fits nothing, so no information criterion applies to it directly.** But the corpus's one applied fit reports held-out error against four baselines and **no AIC or BIC anywhere**, so the parameter count has never been paid for in this corpus, and a 2.8 per cent margin over a two-parameter autoregressive baseline may not be a margin at all [Curzi, 2026i]. **The status of that row across this corpus is NOT MET.** The free parameters this paper's models would carry, counted explicitly: the pair storage model has four (delta, eta, C, and u_0) plus lambda if decay is fitted, giving five; the group-scope model adds D, p_bar and N, giving eight, of which N is enumerated rather than fitted. The contact postulate has **zero** free parameters as written, which is its one genuine strength: it has nothing to tune and therefore nothing to overfit, and its falsifying test is whether a fitted exponent k differs from 1.

Effect sizes and intervals. Nothing in Section 8 carries an interval, because nothing in Section 8 was measured. That is honest for a constructed example and it is not a substitute for the intervals a real application owes.

Precision. The numbers in this paper are of two kinds and are quoted accordingly. *Defined or derived-from-defined* quantities — ln 12 = 2.4849, ln 9 = 2.197225, u_min = 0.077778, t* = 8.4125, the overstatement factor 4.2553 — are exact to the precision printed and may be recomputed to any further precision. *Illustrative model output* — every entry in every trajectory table, every Q_inf, the 0.53 to 1.90 envelope — is the output of stipulated parameters and **carries no measurement uncertainty because there was no measurement**; four decimal places there record the arithmetic, not the confidence. **Measured or fitted quantities do not appear in this paper at all**, because nothing in it was measured. One constant used elsewhere in this corpus and referred to at Section 4.6 — k = (9/8)^2/phi^2 = 0.483425732988414 — is a **defined** constant, exact because it is a ratio of exact quantities; attaching hertz to it is a **separate and much stronger claim**, badged POSTULATED in its own paper [Curzi, 2026], and this paper makes no use of the frequency reading. Golden-ratio claims are the classic site of numerology in quantitative work, and the sceptical literature is cited here precisely so that the distinction

between *choosing* a constant and *finding* one in nature is on the record [Markowsky, 1992].

24.11 11 Limitations

The three worked systems are constructed, not observed. No dataset was collected for this paper. Every number in Section 8 is the output of stated parameters under the stated model, carried through so a reader can re-derive it. They demonstrate that the procedure runs and that it reports its own failures; they are not evidence that any pair or organisation behaves this way and they must never be cited as such.

Every reproducibility gate in this paper reads untested. Tables V-P1 to V-P6, V-G1 to V-G5 and V-C1 to V-C4 all return UNTESTED at gate 7. Nobody has run the cross-pair test, the cross-institution test, or the inter-rater agreement study that the anomaly audit presupposes. **Until one is run, this paper is a specification for how to find out rather than a report of what is known.**

The group-scope magnitudes are demoted and the demotion is the paper's own. The 0.53 to 1.90 operating envelope, the `u*` calibration requirements and the `Delta_min` floors are illustrative constructions demonstrating the model's structure. They are not measurements, and the framework's independent audit reaches the same verdict at both scopes [Curzi, 2026]. What survives is structure: the four limits, the `D = P` fixed point, the zero-headroom identity, the ordering, the sign of `ln r`, and the algebraic join.

The reflection problem, which outranks every other objection here. Section 3.4 states it in full. The aggregation postulate is a linear-in-means reflexive specification whose endogenous effect is not identified from observational data [Manski, 1993], reinforced for network data by Shalizi and Thomas [2011]. **No fit, however well designed and however honestly held out, recovers the parameter.** This paper adopts no remedy — not randomisation, not an instrument, not partial-population variation, and not a defended stipulation of the reference group [Brock and Durlauf, 2001] — so **every applied reading of `u`, `P` and the join is a modelling position rather than an identified empirical result.** Even the parameter-free crossing prediction `t*` is only a consistency test: a crossing on schedule is equally consistent with any correlated unobservable moving on the same timetable.

The Lucas critique applied to this framework's own parameters. The framework asserts that systems change when modelled and measured, then treats `delta`, `eta`, `u` and `lambda` as constants over an estimation window. **Both cannot be true without qualification** [Lucas, 1976, Hacking, 1995]. The position taken is the narrow one: the parameters are stable only over a declared window and only while the order proxy is not itself a reported target, which is why requirements S-9 and G-13 exist and why the two-track ledger of Section 6.10 is a requirement rather than a suggestion. A framework that tells people their arrangements obey a decay law should also ask what that assumption will make true [Ferraro et al., 2005].

`N_eff = 1 + (N−1)u` is not a Hill number and no external literature contains it. `1/HHI` is the Hill number of order two and `exp(H)` is the Hill number of order one, and both are fully grounded [Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979]. `1 + (N−1)u` is a linear interpolation between 1 and `N`, adopted by design axiom. **A reviewer who reaches for the Hill-number literature as support will find that it does not apply.** The two quantities travel under one effective-count heading across this corpus, which is a real exposure.

The functional form is underdetermined. A sub-linear power law in the number of sources [Latané, 1981] and an `n log n` law [Briscoe et al., 2006] fit the same qualitative facts as `ln(1 + (N−1)u)`. **The logarithm is chosen for its axiomatic characterisation, not established by measurement**, and nothing assembled here discriminates.

Surplus monotone increasing in `u` is contested and the contrary evidence is real. The inverted U in creative collaboration [Uzzi and Spiro, 2005], the short-run costs of surface-level diversity [Mannix and Neale, 2005, van Knippenberg and Schippers, 2007], the reduction of trust and civic participation [Putnam, 2007], the destabilisation of randomly assembled systems by connectance [May, 1972], and the finding that redundancy can reduce safety [Sagan,

1993] all bear against a monotone claim. **This paper does not know where the curve turns and does not claim to.**

Comfort, stored order, doctrine and intake have not been shown to be quantities. Michell [1997] is the objection; conjoint measurement [Luce and Tukey, 1964] or a Rasch demonstration [Rasch, 1960] is the remedy; **neither has been run for any quantity in this paper.**

Which parameters were chosen rather than measured. All of them. In the three worked systems, `(N, u, eta, delta, C, lambda, D, p_bar, E)` are stipulated. The reciprocity threshold `a > 0.6931` in requirement S-14 is a design choice, not an empirical finding: **no distribution of measured reciprocity asymmetry exists.** The claim in Section 6.7 that organisations typically come in an order of magnitude below the required `u` is **this framework's impression and not a measured distribution**; no empirical distribution of institutional `u` or of institutional `Delta` exists.

Two roundings in the source manuscript are recorded rather than silently corrected. First, the overstatement factor of Section 7.7 is printed in the source once as 4.26 and once, for the same two-unit case, as 4.25; the exact value is 4.2553, and both printed figures are correct roundings under different conventions. Second, the source describes the step from a six-quarterly to a two-quarterly restoration schedule as "halving again", which is a reduction by two thirds; the tabulated values are used here directly and the description is not repeated. Neither affects a result.

Constant parameters over an estimation window is an assumption organisations violate. Turnover, reorganisation and a change of leadership all move `delta`, `eta`, `N` and `lambda`. A fit spanning a regime change is fitting an average of two systems, neither of which exists.

The limiting-factor reading of the contact postulate has no citation. The reading that a group banks at the rate of the scarcer of frame and intake is offered as an analogy in Section 6.4. No supporting entry exists in this corpus's bibliography, none was manufactured, and the reading carries no evidential weight.

`M` and `M'` in the return difference have no declared dimension. Section 10.1 records the silence. Only their ratio is used, so the formula is well posed, but any claim that `Delta_R` is commensurable with `F` rests on a unit the framework does not supply.

Nothing in this paper is a human-subjects finding. No interventions were run, no participants were recruited, and the declination constraint of Section 4.16 is a design constraint on prospective work rather than a report of past practice.

Unmet rigour checks, printed rather than left to a reviewer. (i) *Model comparison:* **NOT MET.** No information criterion is computed anywhere in this paper, and none is computed anywhere in this corpus's one applied fit [Curzi, 2026i]; the parameter count has never been paid for. (ii) *Effect sizes with intervals:* not applicable here because nothing was measured, and outstanding for any application. (iii) *Scale type:* the conjoint conditions were not tested [Luce and Tukey, 1964]. (iv) *Independent replication:* none, for any parameter, anywhere in this corpus [Open Science Collaboration, 2015, Ioannidis, 2005]. (v) *Differential prediction:* the demand for a prediction that simpler models do not also make [Durlauf, 2005, Sobkowicz, 2020] is met by exactly one item in this paper — the parameter-free crossing date `t*` — and by nothing else. (vi) *Directional versus point prediction:* directional predictions are nearly costless in the social sciences [Meehl, 1967]; this paper does make point predictions with constants, which is the right answer to that objection, but none of them has been put at risk. (vii) *Predictability ceiling:* the best available evidence says social outcomes are barely predictable even with excellent data and many teams [Salganik et al., 2020, Hofman et al., 2017], and nothing here should be read as promising otherwise.

24.12 12 Conclusion

R1 — Opposition is parallel. Because the surplus functional's extremality condition is stated on the squared alignment, `u = 1 - rho^2` sees only the magnitude of `rho`. A pair in perfect opposition and a pair in perfect agreement therefore both read `u = 0`, `N_eff = 1` and surplus exactly

0.0000 nats. **The mirror and the mirror-image are the same instrument**, and a systematic contrarian is audited out of the effective count exactly as a systematic agreer is. Complementarity is a middle condition — at `N = 12`, `rho = +0.95` and `rho = −0.95` both give `N_eff = 2.0725`, while `rho = 0` gives 12.00 — and both extremes of agreement are its failure modes. A pair at 45 degrees already captures three-quarters of the ceiling.

R2 — The minimum viable angle. Setting the steady state to zero gives `u_min = (e^{C/eta} − 1)/(N−1)`, computed at **0.0778, or 16.2 degrees**, for a pair with `eta = 0.55`, `C = 0.34`, `N = 12`. Below its own line a pair sustains no joint store whatsoever: the equilibrium is negative and everything built erodes at rate `delta`. Every pair has such a line and it is computable from three numbers it can measure.

R3 — The exact join. `ln r = ln D − ln p_bar − F`. **The group's interaction surplus subtracts from its log-distortion one nat for one nat, with no conversion constant.** Because `0 <= F <= ln N`, the correction available from interaction geometry is capped at `ln N`, and **if `D/p_bar > N` no arrangement of people calibrates the organisation**, leaving a residual floor `Delta_min = ln(D/p_bar) − ln N` that only doctrine reduction can move.

R4 — Calibration costs all the headroom. At `r = 1` the input sits at its ceiling `ln N`, the store reaches `Q_inf = (eta ln N − C)/delta`, and headroom is `h = eta ln N − (eta ln N − C) − C = 0` exactly. **A perfectly calibrated organisation converges to exactly zero headroom, by arithmetic rather than by mismanagement**, so calibration is not an unqualified target: an organisation that wants room to build must either shed store deliberately or accept `r` not equal to 1.

R5 — Distortion does not average. A headquarters at `Delta = 1.1239` and a field unit at `Delta = 1.9169` aggregate to `Delta = 0.2324` — a near-calibrated headline figure concealing two units in opposite ditches. The aggregate ratio is a weighted mediant of the unit ratios and can sit arbitrarily close to 1 while both units are arbitrarily far from it. Computing bankable input from the aggregate gives 1.9696 nats against a true intake-weighted 0.4629, an **overstatement of 4.2553 times**. Compute the ratio per decision-making unit and aggregate the units' bankable input, weighted by intake, never their doctrine and intake separately.

R6 — Reciprocity is asymmetric where non-parallelism is symmetric. With `a = |ln(eta_AB/eta_BA)|` and the mean conversion held fixed, the sum of the two private steady states is constant while the individual outcomes separate: **the joint books read identically at `a = 0` and at `a = 2.20`, and at `a = 2.20` one partner's private store has a negative steady state.** Reciprocity failure is invisible in the joint account and total in the private one, and it is unidentifiable from the joint store alone.

The counter-result, stated in the same plain language. This paper has not shown that any pair or any organisation behaves this way. Every reproducibility gate reads untested. The group-scope headline envelope of 0.53 to 1.90 is demoted to an illustrative construction at the scale-type gate, along with the calibration requirements and the distortion floors; what survives the demotion is the four limits, the `D = P` fixed point, the zero-headroom identity, the three-way ordering, the sign of the log-ratio, and the algebraic join, all of which stand on ordinal information or on algebra alone. The specification is nonetheless precise enough that the test can be built: named quantities, declared units, a stated normalisation, an over-determination check, and a dated parameter-free prediction — `t* = ln(u*/u_0)/ln(1−lambda)` — that costs a channel audit and a decision log to check.

24.12.1 Declarations

Data and code availability

No external dataset was collected for this paper. All numerical results are computed from the stated parameters under the models developed in the text, and every table can be re-derived from the equations and parameter values printed beside it. The three systems worked in Section 8 are constructed examples with declared parameters, not field observations. The

figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper and the model equations of Sections 3 to 7.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, produced the worked examples and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The three worked systems are constructed. The declination constraint stated in Section 4.16 and repeated in requirements S-20, G-16 and DD-16 — that no protocol removing a person's ability to decline measurement, a channel role, or the dissent it asks of them is within the framework's scope — is a design constraint on prospective work and is not a report of any study performed.

Preregistration

This paper reports no preregistered test of its own; its results are derivations. The preregistration templates it specifies — Sections 8.1, 8.2 and 8.3, and gate 6 of every validation table — are proposals for prospective work. The preregistered protocol for testing the storage law against a real system, together with the confirmatory application that **failed** its second falsification rule with a 2.8 per cent held-out RMSE advantage over a matched autoregressive baseline against a declared margin of 10 per cent, is reported in Curzi [2026i]. That is the current evidential status of every applied claim in this paper, and it is weaker than the mathematics.

Chapter 25

Trust, Reciprocity and Obligation as Quantities: Where Trust Enters the Storage Law, and Nine Ways the Internal/External Correspondence Breaks

Abstract. This paper gives trust an operational location in the storage law: it reduces retention loss, $\delta_{\text{eff}} = \delta_0 + (1-a)(1-v)$, where a is the measured reliance rate — the probability that an unverified commitment is honoured — and v the verified fraction. The steady-state dividend is superlinear: at $a = 0.50$ the stock is 0.4679 and at $a = 0.99$ it is 5.0529. Its point elasticity is $a/(\delta_0 + 1 - a)$, which is 10.56 at $a = 0.95$, and not the $a/(1-a)$ the source material quotes, a form valid only at $\delta_0 = 0$; that correction is one of four defects found in the source and reported here. Verification is shown to be all-or-nothing under a linear verification cost, with threshold $a^* = 1 - c_v\delta_0/(K - c_v)$, computed at 0.8479 for a unit cost of 0.20. Reciprocity enters conversion, $\eta = \eta_{\max}\rho$ with $\rho = 2\min(x_{ij}, x_{ji})/(x_{ij} + x_{ji})$, and obligation is the antisymmetric part of the same matrix; the node-level imbalance index is proved to satisfy $1 - B \geq \rho$ and flatters one worked network by 15.4 per cent. An eight-firm consortium at $u = 0.52$ was collecting 73.8 per cent of its surplus ceiling and still had steady state -0.0836 and headroom -0.4582; the binding fault was $\rho = 0.3824$, not difference, and raising ρ to 0.8947 and a to 0.88 moved the fixed point to $+1.8701$, while trust alone moved it to -0.2114. Graduated sanction gains 1.0 per cent of input against expulsion's -7.7. The correspondence between one head and one group is then audited and breaks in nine named ways, three severe and one absolute: the recursion and the geometry transfer; the licence to apply them does not.

Keywords: trust as retention; verification threshold; reciprocity index; obligation asymmetry; graduated sanction; scope correspondence; mean-field error

Series: B. Engineering the Surplus **Paper:** CCP-025

25.1 1 Introduction

Three words carry more rhetorical weight in the literature of institutions than almost any others, and all three are ordinarily left as adjectives. A team is described as *high-trust*. A community is said to have *strong reciprocal norms*. A network is *dense with obligation*. None of those descriptions tells an engineer what to build, what to measure, or what observation would defeat the claim.

This paper converts all three into quantities that enter the storage law

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,     S_t = F(U_t) = ln(1 + (N−1) U_t)
```

at a named position, with a stated dimension, an obtainable measurement, and a computed price. The conversion is not decorative. **Each of the three attaches to a different parameter**, which is why they behave differently, fail differently, and why the intervention that repairs one will not repair another. Trust is located in the retention loss `delta`; reciprocity is located in the conversion efficiency `eta`; obligation is the antisymmetric part of the same interaction matrix from which reciprocity is read.

The location matters because the storage law's steady state, `Q_inf = (eta F(u*) − c*)/delta`, is a ratio. A parameter in the denominator and a parameter in the numerator are not substitutes and are not even commensurable in their effects: when the numerator is negative, improving the denominator makes the equilibrium worse. Section 6.5 exhibits a case in which a trust intervention — the intervention every member of a failing consortium would have nominated — moves the fixed point from `−0.0836` to `−0.2114`, while a reciprocity intervention on the same system moves it to `+0.7394`. That inversion is the paper's most immediately usable result, and it is arithmetic rather than sentiment.

The second half of the paper audits a claim the corpus has been relying on throughout: that the internal scope (one mind: framings, character, cognitive load) and the external scope (one group: members, institution, coordination overhead) are the same relational system measured twice. The proven part of that claim is narrow and secure — the recursion contains no term naming a substrate, so every theorem about it holds wherever the recursion holds. The rest is a thesis, and Section 8 states **nine specific places where it fails**, three of them severe and one absolute. Every one of the nine is downstream of a single fact: the parts of a society can leave and the parts of a mind cannot.

Four contributions are new relative to the surrounding literature rather than relative to the framework. First, an operational definition of trust — the measured rate at which *unverified* commitments are honoured — that has a true zero, a true ceiling, and an estimator, together with the demonstration that it occupies exactly the position in the storage law that verification spending occupies and is free at the margin. Second, the result that the verification budget's optimum is always at an endpoint, with an explicit switching threshold `a*` and an explicit statement of the linearity assumption that produces the bang-bang structure. Third, the identity `rho = 2 min(x_ij,x_ji)/(x_ij+x_ji) = 1 − |Skew_ij|/Sym_ij`, which makes reciprocity and obligation two readings of one measurement, together with the bound `1 − B >= rho` showing that the node-level index is an upper bound on reciprocity and is biased toward good news. Fourth, a computed audit of the internal/external correspondence that reports its own breakdowns rather than its own scope.

The framework's static half — the surplus functional `F(u) = ln(1 + (N−1)u)`, its uniqueness under four axioms, and the affine effective count `N_eff = 1 + (N−1)u` badged there as design axiom S3 rather than as a theorem — is presupposed and not re-derived; it is established in that form, with `u = 1 − rho^2` from error decorrelation, in Curzi [2026o]. The contraction that supplies the fixed point is stated exactly — `T(Q) = (1−delta)Q + b` is a contraction with constant `1−delta`, giving `Q_inf = (eta M* − c*)/delta` and the exact error law `|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|` — in Curzi [2026q]. The founding reframing on which the whole apparatus rests, that comfort is the felt total of present costs so the present state is a local minimum of the comfort landscape by definition rather than by luck, is presupposed from Curzi [2026].

Sections 2 and 3 state the hypotheses and the framework, including the identification problem that bounds every applied claim in the paper. Sections 4 to 6 develop trust, reciprocity and obligation and price the institutional design principles against them. Sections 7 and 8 develop the correspondence and then audit it. Section 9 works three systems end to end and reports four defects found in the source material. Sections 10 to 13 discuss what follows, verify the chain, state the limitations, and conclude.

25.2 2 Hypotheses

The formal results in Sections 3 to 5 are theorems and definitions about an equation and cannot be falsified by observation; they are labelled where they appear. What can be falsified is the claim that these quantities, measured as specified, occupy the positions claimed for them. Each hypothesis below names the observation that would defeat it. The protocol that operationalises testing of this kind — a preregistered fit, a frozen model, and held-out comparison against four declared baselines — is specified in Curzi [2026i], whose own confirmatory application of it failed its second falsification rule: a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin, despite +40.3 per cent over persistence, +97.7 per cent over direct input, and `p = 0.0005` against the shuffled-input control.

H1 — Trust is a retention parameter. The reliance rate `a`, defined as the measured probability that an unverified commitment is honoured within one period, enters the storage law only through the loss term, as `delta_eff = delta_0 + (1 − a)(1 − v)`. **Predicted:** agreements in high-`a` relations survive unmaintained longer than agreements in low-`a` relations, with half-life `ln 2 / −ln(1 − delta_eff)`; and raising `a` at fixed interaction volume raises `Q` without raising `eta S`. **Falsified by:** a system in which `a` rises while `rho` and `S` stay flat and the measured conversion `eta = dQ_attributable / S_input` also rises — which would place trust on the input side, where the model says it does not sit.

H2 — The verification optimum is an endpoint. With verification cost linear in coverage, `C_v = c_v v`, the steady state is linear-fractional in `v`, so its derivative has constant sign and the optimum is `v = 0` or `v = 1`, switching at `a* = 1 − c_v delta_0 / (K − c_v)` with `K = eta S − C_0`. **Predicted:** bodies verifying on the wrong side of `a*` show a measurable `Q` penalty relative to matched bodies verifying on the correct side. **Falsified by:** an interior optimum observed under a verification cost measured to be linear in coverage. **Note:** an interior optimum under a *convex* cost is predicted by the model and does not falsify it; the falsifier requires the linearity to be measured first.

H3 — Only the matched fraction of interaction converts. `eta = eta_max rho` with `rho = 2 min(x_ij, x_ji)/(x_ij + x_ji)`. **Predicted:** the count of shared structures surviving a full year is monotone in `rho` and is *not* monotone in total directed volume `Sym`. **Falsified by:** durable structure accumulating in proportion to total volume with `rho` uncorrelated — which is the observation the worked consortium of Section 9.1 was designed to be vulnerable to.

H4 — The order of operations is numerator before denominator. Where `eta S − C < 0`, lowering `delta` lowers `Q_inf` further. **Predicted:** in a system with a negative numerator, a trust intervention alone moves the fixed point *away* from zero, and the joint intervention's interaction term exceeds either main effect. **Falsified by:** a measured system with `eta S < C` in which a `delta` reduction alone raises the observed long-run stock. This hypothesis is the sharpest in the paper because its prediction is counter-intuitive and its falsifier is cheap.

H5 — The node-level imbalance index is an upper bound, never an estimate. `1 − B >= rho`, where `B = sum_i |b_i| / (2 * total volume)`. **Predicted:** in any directed interaction matrix, `1 − B` is at least `rho`, with equality only when no member is simultaneously a creditor to one partner and a debtor to another. **Falsified by:** a single counterexample matrix. This is a theorem, so a falsifier would be an arithmetic error rather than an empirical finding; it is stated as a hypothesis only because the *bias direction* — that the gap is systematically toward good news — is an empirical claim and is not proved here.

H6 — The correspondence is of dynamics and geometry, not of agency. The recursion and the redundancy audit transfer between scopes; exit, heterogeneity of interests, preference aggregation, reciprocity, sanction and rivalry do not. **Predicted:** normalised trajectories at the two scopes are the same curve to within a few per cent (Section 7.5 computes 2 per cent on a constructed pair), while a heterogeneous group's fitted `delta` drifts downward over an unmaintained interval and an individual's does not. **Falsified by:** internal trajectories systematically non-exponential in normalised time while external ones are exponential, or the reverse; or fitted `delta` drifting in individuals as it does in heterogeneous groups, which would make the correspondence *stronger* than this paper claims.

Three claims in this paper are not hypotheses and are labelled accordingly. The mapping of trust onto `delta`, of reciprocity onto `eta`, and of obligation onto the antisymmetric part are **modelling postulates** (Section 3.2): they are choices about how a measured social quantity enters a proven equation, and none is derived. The thesis that the internal and external scopes are one relational system is **the author's thesis**, not a theorem, and Section 8.11 states what would kill it. The declination constraint — that no intervention within the framework's scope may remove a person's ability to decline — is a **normative stipulation**; it is not derived from the mathematics and no external citation in this corpus's bibliography establishes it, which is recorded here and again in Section 12.

25.3 3 Framework

25.3.1 3.1 What is presupposed

Two results are used throughout and are established elsewhere in this corpus. Neither is an empirical claim.

The surplus functional. For two unit state vectors `x` and `y` drawn from an ensemble of `N` categories,

```
F(x,y) = f(u(x,y)),    u(x,y) = 1 − (x . y)^2   in [0,1],    f(u) = ln(1 + (N−1) u)
```

with `N_eff = 1 + (N−1)u` the effective count and `0 <= F(u) <= ln N` the hard ceiling [Curzi, 2026o]. The affine effective count follows from a design axiom that is chosen rather than derived, and every use of `N_eff` in this paper inherits that status.

`1 + (N−1)u` is not a Hill number and no external literature contains it. The effective-number idiom is standard in ecology and political science, where `1/HHI` and `exp(H)` are the canonical members of the family — the Hill numbers of order two and one respectively [Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979]. `1 + (N−1)u` is a linear interpolation between 1 and `N`, is this framework's own construction, and a reviewer reaching for the Hill-number literature as support will find that it does not apply. Where an effective count appears anywhere below, it is this one, and it is identified as such. The genuinely Hill-numbered concentration instrument used elsewhere in the corpus is a different quantity: Curzi [2026] reports Jersey 40.0, Guernsey 35.0 and Isle of Man 25.0 giving a naive HHI of 0.345000 and 2.898551 effective jurisdictions

against a bloc-HHI of 1.000000, `N = 1.00` and a surplus term of exactly zero, and the two must not be conflated.

Two warrants exist for the logarithm and the argument needs only the weaker one. The Fechnerian warrant makes sensation proportional to the logarithm of a dimensionless stimulus ratio [Fechner, 1966], and requires a contested perceptual story to be true. The stronger warrant assumes nothing psychological: Peters and Gell-Mann [2016] shows that the time-average growth rate of a multiplicatively evolving quantity is the expectation of its logarithm, so a logarithmic evaluation follows from the dynamics alone. Where a stock compounds rather than adds, the logarithm is the ergodic-average object. The ceiling `ln N` also has an information-theoretic reading as the maximum an ensemble of `N` equiprobable categories can carry [Shannon, 1948], but the framework does not identify surplus with Shannon entropy and the reading is licensed to do no more than bound.

The functional form is underdetermined by the evidence, and this paper says so rather than waiting to be told. At least two published alternatives reproduce the same qualitative facts: social impact rising as a power of the number of sources with an exponent below one [Latané, 1981], and network value growing as `n log n` rather than as `n^2` [Briscoe et al., 2006]. Both give diminishing returns to count, no threshold, and monotonicity. **Nothing assembled in this paper discriminates between them.** What the corpus's form has is an axiomatic characterisation with a declared design axiom at the point where the choice is made; what it does not have is a discriminating measurement.

The storage law. The engine of the paper is the first-order recursion above, with the results this paper consumes stated exactly in Curzi [2026q]: the map `T(Q) = (1−delta)Q + b` is a contraction with constant `1−delta`, the fixed point is `Q_inf = (eta F(u*) − c*)/delta`, and the error law `|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|` is an equality rather than a bound. The admissibility condition `C_t <= (1−delta)Q_t + eta S_t` guarantees `Q_{t+1} >= 0`; the one-step growth condition is `eta S_t > delta Q_t + C_t`; the headroom gauge is `h_t = eta ln N − delta Q_t − C_t`. Curzi [2026h] supplies the identity that headroom at equilibrium is unused variety, `h_inf = eta(ln N − F(u*))`, verified to six decimals at `eta = 0.35`, `delta = 0.015`, `N = 6`, `u* = 0.55`, giving `Q_inf = 2.840970` and `h_inf = 0.164501`; that identity is what makes a positive headroom at a settled state a statement about unrealised geometry rather than about spare effort.

The steady-state result carries a clause this paper leans on twice. If `eta F(u*) < c*` the equilibrium is negative, and the theorem does not conceal it: the regime is unsustainable in this model, or the linear model must be replaced. A negative `Q_inf` is not a small stock. It is the model reporting that it has left its domain. Sections 6.5 and 9.1 turn on that clause.

25.3.2 3.2 The three mapping postulates, badged

Three mappings are introduced. **Each is a postulate — a modelling choice about how a measured social quantity enters a proven equation — and none is derived from the results of Section 3.1.** They are stated here with their badges attached so that the badge does not have to be re-attached at every use. They are labelled `PT`, `PR` and `P0` rather than by their initials alone, to keep them distinct from the results `R1`–`R6` this paper contributes.

Postulate PT — trust enters as retention. For a relation whose stored order consists of standing commitments, the per-period retention loss decomposes as

```
delta_eff = delta_0 + (1 − a)(1 − v)
```

where `a` in `[0,1]` is the **reliance rate**, the measured probability that an unverified commitment is honoured within the period; `v` in `[0,1]` is the fraction of commitments actually re-verified; and `delta_0 >= 0` is the verification-irreducible loss — death, departure, genuine forgetting, external shock. **The postulate is that verification is effective**: a verified commitment does not lapse for want of confirmation.

Postulate PR — reciprocity enters as conversion. For a relation carrying directed interaction volumes `x_ij` and `x_ji` over an accounting window, only the *matched* portion of the volume converts into durable structure:

```
eta = eta_max * rho,     rho = 2 min(x_ij, x_ji) / (x_ij + x_ji)
```

where `eta_max` is the conversion efficiency the relation would have at perfect balance. **The postulate is the word *only*.** The arithmetic of `rho` is a definition and carries no claim.

Postulate PO — obligation is the antisymmetric part. Writing the interaction matrix `X` as its symmetric and antisymmetric parts,

```
X = Sym + Skew,    Sym_ij = (x_ij + x_ji)/2,    Skew_ij = (x_ij − x_ji)/2
```

`Sym` measures the relation and `Skew` measures the obligation. The decomposition is exact and unique for any square matrix; **the postulate is the interpretation of the two parts**, not the algebra.

25.3.3 3.3 Identification: the objection that no amount of fitting can answer

This is the most serious methodological problem in the paper and it belongs here, where the loop is introduced, rather than in a caveat at the end.

The framework's central applied specification is reflexive. The group's interaction geometry `U_t` is a property of the group; the group's stored order shapes what its members do next; and the members are what the group is made of. Written as `T(x) = W(x, M(x))`, that is exactly a linear-in-means reflexive specification, and Manski [1993] proves that **the endogenous social effect in such a specification is not identified from observational data**. It cannot be separated from contextual effects or from correlated unobservables, however much data is collected. Shalizi and Thomas [2011] reinforce the point for network data, where influence and selection are generically confounded. Kirman [1992] makes the companion aggregation objection: the behaviour of an aggregate is not in general the behaviour of a representative member, so a scalar model of a heterogeneous body is a modelling decision and not a summary — a point this paper computes rather than asserts in Section 8.2.

The consequence for this paper is exact. The derivations of Sections 4, 5 and 7 are statements about an equation and are untouched. But the interpretation of a measured `rho` as *causing* a conversion efficiency, rather than as a contextual effect, a common shock, or selection, is **not identified by observation alone**, and no held-out comparison recovers it. This limitation is deeper than any the falsification protocol of Curzi [2026i] tests for, and it would remain if every rule in that protocol passed.

No remedy is adopted in this paper, so the honest label is a modelling position. The known routes are randomisation, an instrument, partial-population variation, or stipulating the reference group and defending the stipulation; the last is the most realistic here and Brock and Durlauf [2001] is its template. The measurement protocols of Section 6.6 point toward it — Protocol A randomises non-verification within a pre-declared consequence tier, which is a genuine partial-population design at the level of the *commitment*, not of the person — but no protocol in this paper randomises at the level at which the reflexive parameter lives. **Until one of them is executed, `eta` and `rho` as applied quantities are coefficients in a declared specification and not identified social effects**, and every applied statement in this paper inherits that status.

One existence proof is worth naming because it is the strongest thing available on the other side. MacKenzie and Millo [2003] documents a reflexive fixed point observed in the field, with dates and price data: a model that, once adopted, moved the world toward its own predictions. That establishes that reflexive fixed points occur. It does not identify this one.

25.3.4 3.4 Scale type, and the seven gates

Every formula in this paper is audited against a seven-gate protocol, and the gate names are used in short form throughout: **G1** dimensions, **G2** scale type, **G3** limiting cases, **G4**

fixed points, **G5** identifiability, **G6** falsification, **G7** reproducibility. Verdicts are `PASS`, `PASS (conditional)` with the condition stated, `FAIL` with the consequence stated, `DEMOTE` where a quantity survives as an ordering but not as a magnitude, or `OPEN` where the gate has not been run. **`OPEN` is not a failure and it is not a pass; it is a debt, and it is recorded as one.**

G7 is `OPEN` for twenty-four of the twenty-five formulary entries of Appendix A and `PARTIAL` for the twenty-fifth, because nothing in this material has been through a cross-lab reproducibility trial. The reproducibility literature is unambiguous about what that is worth [Open Science Collaboration, 2015, Ioannidis, 2005].

G2 is where this paper is most exposed and the exposure is stated rather than managed. Michell [1997] argues that psychology routinely assumes its attributes have quantitative structure without testing whether they do, and that the assumption is usually false or at least unexamined. **That is the strongest objection to this framework and it bites hardest on a paper that treats trust, obligation and reciprocity as quantities.** The formal remedy exists: conjoint measurement supplies the test for whether an attribute has additive quantitative structure at all [Luce and Tukey, 1964, Krantz et al., 1971], and a Rasch demonstration is the psychometric alternative [Rasch, 1960]. **Neither has been run for any quantity in this paper, and until one is, every ratio-scale treatment below is a stipulation.**

The stipulation is not equally exposed everywhere, and the difference is worth stating. `a`, `v`, `rho`, `B` and `tau` are **proportions of counted events** — honoured commitments out of unverified commitments, matched flow out of total flow, departures out of roster. A proportion of a count has a true zero and a defensible ratio scale provided the count is a count, which is why trust was defined as a rate rather than as a belief: the ordinal self-report that the usual definition produces cannot be divided, and averaging it is not licensed [Stevens, 1946, Likert, 1932]. `Q`, `eta` and `C` carry the framework's declared order proxy and inherit its scale, which the substitution test supplies and a five-point culture survey does not. `D` and `P` — and therefore `M`, `r` and every band table computed from them — are **demoted**: the division requires a common ratio scale with a shared true zero across two separately instrumented terms, and no such instrument exists. Curzi [2026] records the same demotion in the corpus's register, where the moment equation was demoted on scale. **What survives the demotion is the ordering — over-framed, calibrated, under-framed — the `D = P` fixed point, and the sign of `ln r`, which is the operationally decisive part because the two faults take opposite corrections. What is withdrawn is the magnitude.**

25.3.5 3.5 Standing honesty clauses

Repeated wherever the law is used, and stated once here.

- **As deployed, the surplus functional is a lower bound** on the ensemble's information, not the information itself. Every price quoted in this paper is a floor.
- **The triangle inequality fails for geometric triples.** Subadditivity holds on the parameter `u` only. Three-party obligation chains do not compose metrically, and no bound in this paper may be chained through an intermediate party. The point is not exotic: psychological similarity is known to be asymmetric and to violate the triangle inequality [Tversky, 1977].
- **The static half of the framework is pure mathematics** and makes no empirical or metaphysical claim. Surplus is not identified with Shannon entropy or mutual information, and no thermodynamic negentropy law is proved.
- **Applying the storage law to a relation, a team or an institution is a modelling decision**, badged as such every time it is made, and answerable through the protocol of Curzi [2026i].
- **The external results cited below are cited for what their authors actually reported.** The transfer of any of them into this parameterisation is a modelling assumption, badged at each use.
- **The declination constraint governs absolutely: no intervention in this paper may remove a person's ability to decline.** Section 5.5 shows this is load-bearing rather than

decorative, because the highest-reliance obligation networks on record are precisely the ones that violate it. The constraint is a normative stipulation of the framework and is labelled as one; Curzi [2026] states the corpus's boundary in its strongest form, refusing category innocence as a defence of acquisition, forbidding the movement of any harm through a discovered gap in a scheme, and conceding that the case where the feedback loop runs through another person's will is an unsettled seam the corpus does not close. This paper does not close it either.

- **The classification of every claim as established mathematics, established psychology, or declared postulate follows the corpus's three-drawer honesty ledger**, which classifies four items of established mathematics, seven of established psychology, and exactly three declared postulates, with every protocol stated to run on the first two and steer by the third [Curzi, 2026]. The three postulates of Section 3.2 are declared in that third drawer.

25.4 4 Trust as a retention parameter

25.4.1 4.1 An operational definition

The usual definition is refused first. Trust is commonly defined as a belief or a feeling — an expectation held by one party about another. That definition is unmeasurable except by asking, and asking produces a self-report on an ordinal scale, which cannot be divided, multiplied, or entered into a ratio (Section 3.4). The operational definition used here is:

> **Trust on a relation is the measured rate at which unverified commitments are honoured.**

Call it `a`, the **reliance rate**. It is a probability; it is dimensionless; it has a true zero — no unverified commitment ever holds — and a true ceiling; and it is estimable from records most organisations already keep. **It is a property of a relation in a domain, not of a person**: a supplier's `a` on delivery dates and their `a` on cost estimates are different numbers and must be measured separately.

The definition's whole content is the word "unverified." A commitment that was checked says nothing about trust; it says something about the checking. Trust is visible only in the commitments that were *not* checked, which is why measuring it requires deliberately not checking some of them, and why that in turn requires an ethics protocol (Section 6.6, Protocol A). The framing is compatible with the standard treatment of trust as a bet on another's future action under incomplete information [Gambetta, 1988], and with the repeated-game result that reciprocation under a long shadow of the future is a self-enforcing equilibrium rather than a disposition [Axelrod, 1984, Axelrod and Hamilton, 1981]; **what is new here is not the concept but the location of the measured rate inside a stock equation.**

25.4.2 4.2 Why trust reduces `delta` and nothing else

Consider a relation whose stored order `Q` is a set of standing commitments: agreements about who does what, who may draw on what, what a given signal means. That is the ordinary content of institutional `Q`.

Each period a commitment either carries forward or lapses, and lapse has two causes.

1. **Irreducible causes.** The person leaves, the memory genuinely goes, the world changes under the agreement. No confirmation prevents these. Rate: `delta_0`.
2. **Confirmation-addressable causes.** The commitment is nominally still in force but nobody has confirmed it, and in the absence of confirmation the parties diverge on what it says or whether it binds. Rate, per unconfirmed commitment: `1 - a`.

If a fraction `v` of commitments is re-verified each period, the confirmation-addressable loss applies only to the remaining `1 − v`, which is Postulate PT. Its content is immediate: **there are two ways to hold `delta` down, and trust is the one that does not cost anything per period.** Raising `a` and raising `v` reduce the same term. Raising `v` costs `C_v = c_v v` in the dissipation budget. Raising `a` costs nothing recurrent.

That is the precise sense in which trust is a retention device: it occupies the exact position in the storage law that verification spending occupies, and it is free at the margin. **The claim is a shared position in one equation, not a physical derivation** — the analogy discipline of Section 11.8 applies.

25.4.3 4.3 The verification frontier and the half-life of an agreement

Fix `a`. Varying `v` from 0 to 1 traces a straight line in the `(delta, C_v)` plane:

```
delta = delta_0 + (1 − a)(1 − C_v / c_v)
```

Endpoints, computed with `delta_0 = 0.04` and `c_v = 0.25`:

a	`v = 0` (verify nothing)	`v = 1` (verify everything)
0.50	`delta` = 0.5400, `C_v` = 0.0000	`delta` = 0.0400, `C_v` = 0.2500
0.70	`delta` = 0.3400, `C_v` = 0.0000	`delta` = 0.0400, `C_v` = 0.2500
0.85	`delta` = 0.1900, `C_v` = 0.0000	`delta` = 0.0400, `C_v` = 0.2500
0.95	`delta` = 0.0900, `C_v` = 0.0000	`delta` = 0.0400, `C_v` = 0.2500

Read the table by its intercepts. **Trust moves the `delta`-intercept toward the origin and leaves the cost intercept exactly where it was.** Every point on a high-`a` frontier weakly dominates the corresponding point on a low-`a` frontier: same verification spend, lower loss. That is the geometric statement of the dividend, and it is why trust is not a substitute for verification but a rotation of the whole trade-off.

The half-life consequence, with `v = 0`, computed as `t_half = ln 2 / −ln(1 − delta)`:

a	`delta_eff`	half-life of an unmaintained agreement (periods)
0.50	0.5400	0.893
0.70	0.3400	1.668
0.80	0.2400	2.526
0.90	0.1400	4.596
0.95	0.0900	7.350
0.99	0.0500	13.513

An agreement in a relation at `a = 0.50` is half-forgotten before the next period ends. The same agreement at `a = 0.99` survives thirteen and a half periods. **The agreement is identical and the document is identical; what differs is the rate at which the document stops describing what people do.**

25.4.4 4.4 The steady-state dividend, and a correction to its elasticity

Take the steady state with `N = 8`, `u = 0.55`, so `S = ln(1 + 7(0.55)) = 1.5790`; `eta = 0.35`; `C_0 = 0.30`; hence the net input `K = eta S − C_0 = 0.2526`:

```
Q_inf = K / delta_eff = 0.2526 / (0.04 + (1 − a))
```

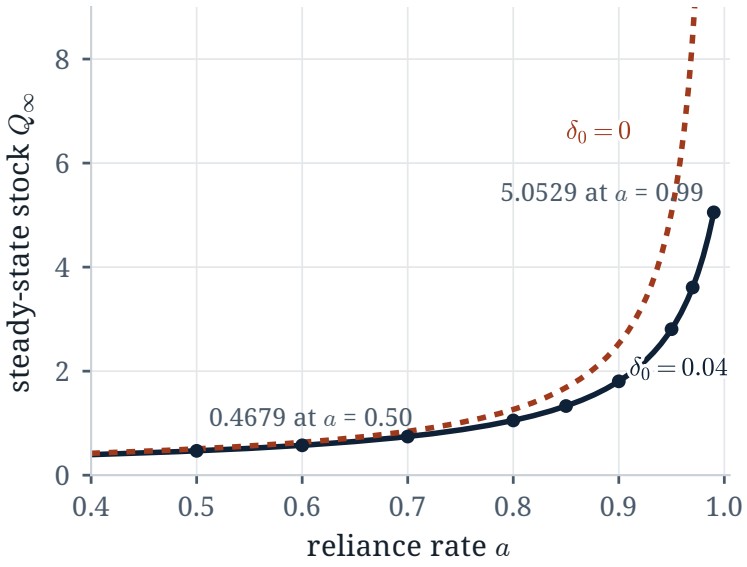


Figure 25.1: **The stock dividend from reliance is superlinear but it is bounded by the irreducible loss, and the corpus's own elasticity formula omits that bound.** Steady state `Q_inf = K/(delta_0 + 1 − a)` with `K = 0.2526`, plotted against the reliance rate; the marked points are the nine entries of the Section 4.4 table, from 0.4679 at `a = 0.50` to 5.0529 at `a = 0.99`. The dotted curve is the same law with `delta_0 = 0`, which reaches 25.26 at `a = 0.99` and is the curve whose elasticity is the frequently quoted `a/(1−a)`. Modelled from the parameters stated in Section 4.4; generated by `figs/src/CCP-025_fig01.py`.

a	delta_eff	Q_inf	share of the a = 0.99 ceiling
0.50	0.5400	0.4679	9.3%
0.60	0.4400	0.5742	11.4%
0.70	0.3400	0.7431	14.7%
0.80	0.2400	1.0527	20.8%
0.85	0.1900	1.3297	26.3%
0.90	0.1400	1.8046	35.7%
0.95	0.0900	2.8071	55.6%
0.97	0.0700	3.6092	71.4%
0.99	0.0500	5.0529	100.0%

Because `Q_inf` varies as `1/delta`, the point elasticity of stored order with respect to the reliance rate is

```
E(a) = (dQ_inf/da)(a/Q_inf) = a / delta_eff = a / (delta_0 + 1 − a)
```

The source material for this paper quotes that elasticity as `a/(1−a)`, giving 1.00, 2.33, 4.00, 9.00, 19.00 and 99.00 at `a` = 0.50, 0.70, 0.80, 0.90, 0.95 and 0.99. That form holds only in the limit `delta_0 -> 0`, and it is inconsistent with the table printed immediately above it, which carries `delta_0 = 0.04`. The two columns are given side by side here rather than one being quietly replaced:

`a`	0.50	0.70	0.80	0.90	0.95	0.99
elasticity as `a/(1−a)`, i.e. at `delta_0 = 0`	1.00	2.33	4.00	9.00	19.00	99.00
elasticity at this table's `delta_0 = 0.04`	**0.93**	**2.06**	**3.33**	**6.43**	**10.56**	**19.80**

A one per cent improvement in the reliance rate at `a = 0.95` buys 10.6 per cent more stored order at `delta_0 = 0.04`, not nineteen; the same one per cent at `a = 0.50` buys 0.93 per cent, which does round to one. The qualitative claim survives the correction and is the useful part: the last increments of reliability are worth far more than the first, and organisations that are nearly reliable are qualitatively different from organisations that are mostly reliable. It also explains the reverse — why a single prominent breach in a high-`a` institution is so destructive. That is not sentiment; it is `1/delta_eff`.

The honest caveat, which the source states correctly and which the corrected elasticity now matches. The superlinearity is an artefact of the model's `1/delta` steady state, which assumes the linear law holds as `delta -> 0`. It does not. As `delta` approaches `delta_0` the irreducible term dominates and the curve flattens; the table already includes `delta_0 = 0.04`, which is why the `a = 0.99` entry is 5.05 and not 25.26. **Anyone quoting the elasticity must quote `delta_0` with it, and must use the form that contains it.**

25.4.5 4.5 The verification threshold, and why the optimum is at an endpoint

Take `Q_inf(v) = (K − c_v v)/(delta_0 + (1−a)(1−v))` and differentiate. Writing `A = 1 − a`:

```
sign( dQ_inf/dv ) = sign[ A (K − c_v) − c_v ∗ delta_0 ]
```

The sign does not depend on `v`. **The objective is a linear-fractional function of the verification rate, so its optimum is always at an endpoint: verify everything, or verify nothing.** The switch happens at

```
a∗ = 1 − c_v ∗ delta_0 / (K − c_v),      valid when K > c_v
```

With `K = 0.2526` and `delta_0 = 0.04`:

unit verification cost `c_v`	`a∗`	reading
0.02	0.9966	verify unless the relation is essentially perfect
0.06	0.9875	verify unless `a` exceeds 98.75%
0.12	0.9638	verify below 96.4%
0.20	0.8479	verify below 84.8%
0.24	0.2381	verify only in near-total distrust
>= 0.2526	never	verification consumes the entire net input

Check the two ends against intuition, because a formula that surprises at the ends is usually wrong. **Cheap verification means one should verify almost always** — which is why retail commerce, where verification is a barcode scan, does not run on trust and does not need to. **Expensive verification means one should verify almost never** — which is why relations whose

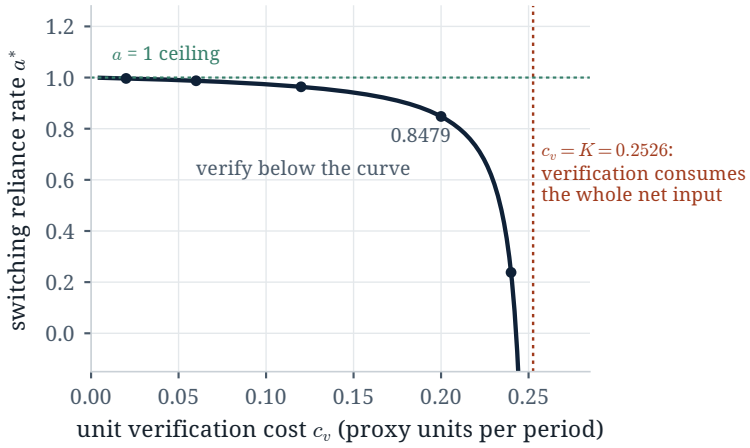


Figure 25.2: **Verification is worth buying only below a threshold reliance rate, and the threshold collapses as the verification cost approaches the net input.** The switching reliance rate `a* = 1 − c_v delta_0/(K − c_v)` at `K = 0.2526` and `delta_0 = 0.04`; the marked points are the five finite entries of the Section 4.5 table (0.9966, 0.9875, 0.9638, 0.8479, 0.2381) and the vertical asymptote is `c_v = K = 0.2526`, beyond which verification consumes the whole net input and is never worth buying. Modelled from the parameters stated in Section 4.5; generated by `figs/src/CCP-025_fig02.py`.

content is judgement, taste, discretion or intent run on trust and have no alternative, since auditing a judgement costs more than the judgement is worth.

This is the practical content of the trust parameter: trust is valuable exactly in proportion to the cost of the verification it displaces. Where verification is cheap, trust is a luxury. Where verification is expensive or impossible, trust is the only available control on `delta`, and its absence is not a cultural deficiency but a hard ceiling on stored order.

Model artefact, flagged. The endpoint result follows from the *linearity* of `C_v = c_v v`. Real verification costs are convex — the first ten per cent of checks are cheap and the last ten per cent are punishing — and with a convex cost the optimum is interior. Use the endpoint result as a signpost about direction, never as an instruction to audit everything or nothing. **What survives convexification is the threshold structure: there is an `a` above which additional verification destroys more order than it saves.** A caution on how `c_v` is priced in practice. Curzi [2026] computes a structurally parallel trade in a disclosure setting: concealment costing USD 210,000 against a compliant USD 561,000 breaks even at an examination probability `p* = (561,000 − 210,000)/41,500,000 = 0.008458`, and **the visible 62.6 per cent saving counts the invoice and not the exposure.** The same error is available here. `c_v` is the invoice for verification; the quantity that belongs opposite it is not the invoice for trust, which is zero, but the exposure carried when an unverified commitment fails. **A verification budget priced on invoices alone will sit on the wrong side of `a*` in exactly the direction that looks cheapest.**

There is a second reason not to read the `v = 1` corner as safe. Redundancy added to a system does not monotonically improve it: Sagan [1993] documents cases in which added redundancy reduced safety by increasing complexity, opacity and the opportunity for common-mode failure. A verification regime is redundancy, and the same caution applies.

25.4.6 4.6 Four things trust is not

Four negations, each a common and expensive error.

Trust is not scalar. There is no such thing as an organisation's trust level. There is a matrix `a_ij^(d)` indexed by pair and by domain, and averaging it discards the only information it carries. A supplier at `a = 0.98` on delivery and `a = 0.55` on cost forecasts is not an `a = 0.77`

supplier; they are a supplier to rely on for one thing and audit for another, and the average would prescribe exactly the wrong action in both domains.

Trust is not a feeling. The feeling is a predictor of `a`, sometimes good and sometimes catastrophically bad, and it is available before the data. It is a prior, not the measurement, and where prior and measurement disagree the measurement is the datum. Self-report about the causes of one's own judgements is precisely the measurement known to be unreliable [Nisbett and Wilson, 1977].

Trust is not an input. It appears nowhere in `eta S_t`. A high-trust group with no interaction surplus accumulates nothing; it merely loses nothing slowly. The extremality result holds: `S = 0` exactly when the parties' framings are parallel, and it stays zero however much the parties rely on one another. **Trust protects the store; it does not fill it.**

Trust is not always desirable. `delta` is not a fault to be minimised. A system with no forgetting cannot revise, because revision is the replacement of an old weighting by a new one and replacement requires that the old weighting fall; the organisational-learning literature reaches the same place from the other side, where too-fast convergence on a shared code destroys the variance that made the body capable [March, 1991, Levinthal and March, 1993]. High `a` on a commitment that ought to lapse is exactly the mechanism by which obsolete procedure outlives its purpose — the commitment is honoured, unverified, forever, because nobody reopens what is trusted. **The engineering target is not maximum `a`; it is `a` high on the commitments that should persist and deliberately low on the ones that should not.** That is what a sunset clause is: a scheduled reduction in `a` on a specific agreement.

25.5 5 Reciprocity and obligation

25.5.1 5.1 The reciprocity index

Reciprocity is usually described as a norm — the expectation that a favour will be returned. As a norm it is unmeasurable; as a *flow balance* it is trivially measurable, provided the flows can be counted.

Let `x_ij` be the volume of transfers from `i` to `j` over an accounting window `w`, in some common unit. The **matched volume** of the dyad is the part both parties supplied:

```
matched_ij = 2 min(x_ij, x_ji)
total_ij   = x_ij + x_ji
rho_ij     = matched_ij / total_ij = 2 min(x_ij, x_ji) / (x_ij + x_ji)
```

Equivalently, and this identity is the bridge to obligation,

```
rho_ij = 1 − |x_ij − x_ji| / (x_ij + x_ji)
```

The two expressions are the same number, since `a + b − |a − b| = 2 min(a,b)`. **The reciprocity index is simultaneously the matched fraction of the volume and one minus the normalised imbalance.** It is dimensionless, has a true zero (all flow one way) and a true one (perfect balance), and is invariant to the unit in which flows are counted — all three properties are required before it may multiply anything.

x_ij	x_ji	matched	total	rho	eta = 0.45 rho
10	10	20	20	1.0000	0.4500
10	8	16	18	0.8889	0.4000
10	6	12	16	0.7500	0.3375
10	5	10	15	0.6667	0.3000
10	4	8	14	0.5714	0.2571

x_ij	x_ji	matched	total	rho	eta = 0.45 rho
10	2	4	12	0.3333	0.1500
10	1	2	11	0.1818	0.0818
10	0	0	10	0.0000	0.0000

25.5.2 5.2 Why the matched fraction is the converting fraction

`eta` is the fraction of interaction surplus that becomes durable structure. Postulate PR asserts that the durable fraction is the matched fraction, and the reasoning is this.

An unreturned transfer creates a claim held by one party. A claim held by one party persists at the *debtor's* convenience; it is an entry in one ledger and the other ledger has no matching row. There is no mechanism by which both parties are separately motivated to maintain it, and no independent copy against which drift can be detected. A returned transfer creates a claim held by both. Both maintain it, both can check it, and it survives the departure of either one's memory of the original transaction.

Structure that only one party holds is not structure; it is a debt. That sentence is the postulate, and it is the reason `eta = eta_max rho` rather than `eta = eta_max`. It is consonant with, but not derived from, the finding that reputation extends reciprocity beyond the dyad by giving third parties a copy of the record [Nowak and Sigmund, 2005] — a mechanism that works precisely by manufacturing the second ledger the unmatched transfer lacks.

The consequence at steady state is severe, because `eta` sits in the numerator while `delta` sits in the denominator. With `S = 1.5790`, `C = 0.30`, `delta = 0.10`, `eta_max = 0.45`:

x	y	rho	eta	eta S	Q_inf = (eta S − C)/delta
10	10	1.0000	0.4500	0.7106	+4.1055
9	6	0.8000	0.3600	0.5684	+2.6844
8	3	0.5455	0.2455	0.3876	+0.8757
12	2	0.2857	0.1286	0.2030	-0.9699
5	0	0.0000	0.0000	0.0000	-3.0000

The negative entries are not small numbers; they are the steady-state theorem's own statement that the regime is unsustainable and that the linear model must be replaced. **A relation running at `rho = 0.29` does not accumulate slowly. It fails to accumulate at all, and it does so while the interaction volume — fourteen transfers a period — looks perfectly healthy on a dashboard.**

The same sign trichotomy appears in the corpus in frequency form. Curzi [2026] establishes a minimum sustainable frequency `f_min = k e^(2c/eta)` with the trichotomy `f > = < f_min` corresponding exactly to `Q_inf > = < 0`, which is the same three-way verdict this table reports through `rho` rather than through contact rate. **That paper's constant `k = 0.483425732988414` is a *defined* constant — the exact ratio `(9/8)^2 / phi^2`, computable to arbitrary precision, which is why fifteen figures are legitimate there — and the attachment of the unit `Hz` to it is a separate and much stronger claim, badged in this paper's verification chain as `POSTULATED` rather than derived.** Because the definition passes through the golden ratio, the standing scepticism about golden-ratio claims applies wherever it appears [Markowsky, 1992]. **Nothing in the present paper depends on `k`.**

25.5.3 5.3 The accounting window

`rho` is not a property of a relation. It is a property of a relation *and an accounting window*. A favour returned in month nine counts as reciprocated if the window is a year and unrecipro-

cated if the window is a quarter. For fifteen favours with mixed return lags, five of them long and two never returned:

window `w` (months)	1	2	3	6	12	24
estimated `rho`	0.1333	0.2667	0.4000	0.6000	0.8000	0.8667

The same relation measures at `rho = 0.13` and `rho = 0.87` depending on a choice the analyst makes. That is an identifiability failure on gate G5 and a researcher degree of freedom of exactly the kind that inflates false-positive rates when it is exercised after the data are seen [Simmons et al., 2011, Gelman and Loken, 2013]. There are three honest responses.

1. **Preregister the window** before looking at the data, and justify it from the domain's natural cycle — the length of a project, a season, a budget year [Nosek et al., 2018].
2. **Report the whole curve**, not a single number. The *shape* is informative: a curve rising steeply and plateauing early describes a fast-settling relation; one still rising at 24 months describes a relation carrying long debt.
3. **Set the window from the decay rate.** The principled choice is `w = 1/delta`, the time constant: beyond that horizon the structure the return was supposed to build has already leaked away, so a return arriving later is arriving into an emptier store. This is the option this paper recommends, and it is the one used in Section 9.3.

25.5.4 5.4 Obligation as the antisymmetric part

Take the interaction matrix `X` and split it as in Postulate PO. Every square matrix has exactly one such decomposition. `Sym` is the dyad's traffic; `Skew` is the standing debt. Note immediately that `rho_ij = 1 - |Skew_ij|/Sym_ij`, **so reciprocity and obligation are two readings of one measurement, not two measurements.**

At the node level, define each member's **net obligation balance** as the row sum of `Skew`:

```
b_i = sum_j (x_ij - x_ji) = (total given) - (total received)
```

`b_i > 0` means `i` is a net creditor, and necessarily `sum_i b_i = 0`. Define the **network imbalance index**

```
B = sum_i |b_i| / (2 * total volume)
```

with `B = 0` perfect balance at every node and `B` approaching 1 a pure one-directional flow.

A small theorem, stated because it prevents a common error. Node-level balance can hide dyad-level imbalance, since a member who is a creditor to one partner and a debtor to another shows `b_i = 0`. Formally `|sum_j (x_ij - x_ji)| <= sum_j |x_ij - x_ji|` by the triangle inequality on the reals, so the node-level imbalance never exceeds the dyad-level imbalance, and therefore

```
1 - B  >=  rho
```

`1 - B` is an upper bound on reciprocity, never an estimate of it. In the worked consortium of Section 9.1 the two numbers are `1 - B = 0.4412` and `rho = 0.3824`: the node-level index overstates the health of the network by 15.4 per cent, and it does so systematically, in the direction of comfort. **Report both, or report `rho`.**

25.5.5 5.5 Two literatures read through the parameters

Two of the most cited empirical bodies on this subject describe the same structural object from opposite ends, and the parameterisation separates them cleanly. **Both are cited for what their authors reported; the transfer into this parameterisation is a modelling assumption, badged here once for both.**

Protection markets and the price of a guarantee. Where central enforcement has failed, an order can be maintained through networked obligation, and a market arises in guarantees. In these terms the finding is that such a market arises precisely where `a` is low and `c_v` is high — where commitments are not honoured unverified and verifying them independently is prohibitively expensive. **What is sold in such a market is not a good but a reduction in `delta`**: a third party guarantees that the commitment holds, so the parties need not verify. That is an economically coherent product, and its existence is what the frontier of Section 4.3 predicts when `a` is low and `c_v` is unaffordable. The standing treatment of trust as the object being traded in such arrangements is Gambetta [1988]. **A caution on attribution:** the source material for this paper attributes the specific Sicilian protection-market study to Gambetta (1993), which is a separate monograph and is **not present in this corpus's bibliography**; the reading above is therefore this paper's own and is cited only to the 1988 volume for the concept of trust it develops.

And here the boundary must be stated rather than the mechanism admired. The guarantee is enforced by removing the ability to decline. Under the declination constraint that is not a permitted intervention, and under the framework's own accounting it is not even a good one: coerced compliance raises *measured* `a` while leaving the underlying relation at `rho` near zero, because the flows are one-directional by construction. The system therefore shows a low `delta` and an `eta` near zero simultaneously — which is a negative equilibrium propped up by continuous coercive input. **An obligation network that removes exit is not a trust system. It is a capture system, and the two are distinguishable by measuring `rho`, which the first has and the second does not.** That is a testable difference and it is the reason the ethical rule and the engineering rule point the same way here.

Social capital and the persistence of civic tradition. Putnam et al. [1993] studied the divergence in institutional performance between northern and southern Italian regions after identical formal structures were devolved in 1970, and attributed the difference to accumulated civic tradition. The mapping is direct:

reported construct	parameter
social capital	`Q`, stored order held on relations rather than in a centre
civic association	the input term `eta S`: interaction that leaves structure
horizontal networks	high `rho`: mutual, balanced, both parties hold the claim
vertical patron-client ties	low `rho`: one-directional, only the patron holds the claim
the persistence of tradition	low `delta`: `Q` with a very long time constant

The most striking claim — that the divergence traces to civic patterns centuries old — is, in this notation, a claim that `delta` is very small: a store with a half-life measured in generations. **That is exactly the regime in which `Q_inf = (eta S − C)/delta` becomes enormous for even modest input, and it is also the regime in which recovery from a low `Q` is hopelessly slow, because the same small `delta` that preserves a good store preserves a bad one. Low `delta` is symmetric. It is not a virtue; it is a memory.** Fukuyama [2001] argues in the development context that social capital is far easier to erode than to create by policy, which is the same asymmetry read from the other side.

The honest note, required. The causal argument has been contested by historians and political scientists on the historical evidence and on the direction of causation. This paper takes no position on that dispute. What the framework offers is not a verdict but a decomposition: the competing hypotheses make **different predictions about which parameter moves first**, and that is a measurable difference rather than an interpretive one.

The contrary evidence, printed here rather than left for a reviewer. The corpus's founding claim is that surplus rises with unlikeness. The largest piece of contrary evidence in the social sciences bears directly on it, and it comes from the same author as the social-capital result above: Putnam [2007] reports that ethnic diversity in the short run **reduces** trust and civic participation, including trust in one's own group, with the benefits appearing only over longer horizons. In this parameterisation that is a claim that raising `u` raises `C` and lowers `a` faster than it raises `eta S` — which the storage law can express perfectly well and which the corpus's own headroom gauge would report as `h < 0`. It is not a refutation of the arithmetic; it is a warning that the sign of the net effect is an empirical question the arithmetic does not settle. Abascal and Baldassarri [2015] is the principal rebuttal, finding that the effect largely disappears under better controls for confounded neighbourhood characteristics, and the pair should be cited together. Two further constraints belong beside them. Contact across difference produces value **only under conditions** — equal status, common goals, cooperation, institutional support [Allport, 1954] — and the meta-analytic test finds contact does reduce prejudice with those conditions facilitating rather than gating the effect [Pettigrew and Tropp, 2006]. And the returns to difference are not monotone: Uzzi and Spiro [2005] finds an inverted U, with team performance rising and then falling as the mix of novel and repeat collaborators increases. **The surplus curve turns over, and `F(u) = ln(1 + (N−1)u)` as written does not.** That is a genuine mismatch between the functional form and the best available evidence, and it is recorded again in Section 12.

25.6 6 Institutional design, priced against the parameters

25.6.1 6.1 The eight design principles mapped to parameters

Ostrom [1990] derived eight design principles common to long-enduring common-pool resource institutions, against the standing prediction that a commons without private property or central authority must degrade [Hardin, 1968]. **The pair must be cited together**: the design principles are the empirical refutation of the prediction, and the prediction is what makes the principles interesting. The polycentric generalisation — many formally independent centres of decision — is stated in Ostrom [2010], and the meta-analytic audit that tests the principles against a large body of case material, finding most supported and several in need of splitting, is Cox et al. [2010]. **This paper's use of the principles presupposes that decentralised resource governance with graduated sanctions, monitoring and nesting is empirically viable; that presupposition rests on those sources and is not established here.**

The corpus has previously treated the principles collectively as `eta`-raising devices. Mapping each to the parameter it actually improves is more discriminating, and **they do not all improve the same thing — three of them improve something other than `eta`.**

#	Design principle	Parameter improved	Mechanism, in this parameterisation
1	Clearly defined boundaries	`N`, and `C`	Makes `N` a countable set rather than a rhetorical one; excludes uncounted draws that appear in the model as unexplained `C`

#	Design principle	Parameter improved	Mechanism, in this parameterisation
2	Congruence: rules match local conditions; benefits proportional to costs	`rho`, hence `eta`	"Proportional" is literally `rho -> 1`. A rule allocating benefit in proportion to contribution *is* a balance condition on the flow matrix
3	Collective-choice arrangements	`u`, hence `S`	Brings members' non-parallel framings into rule-making. Rules written by one framing have `u = 0` at the point where variety was needed
4	Monitoring, by monitors accountable to the appropriators	`c_v`, **not `a`**	Section 6.2. This is the principle most often misread
5	Graduated sanctions	preserves `rho` and `N`	A maximal first sanction deletes the edge; a graduated one preserves it. Section 6.3
6	Conflict-resolution mechanisms, rapid and low-cost	`C`, directly	Unresolved conflict is a standing dissipation term. "Rapid and low-cost" is a specification on `C`, not on justice
7	Minimal recognition of the right to organise	protects `Q` from external shock	Guarantees `Q` is not zeroed by an outside authority; in the model it bounds a catastrophic `delta` event away from 1
8	Nested enterprises	`C`, by orders of magnitude	The only principle that changes the *scaling exponent* rather than a coefficient. Section 6.4

25.6.2 6.2 Monitoring lowers the verification cost; it does not raise reliance

The fourth principle deserves its own treatment because the usual reading is wrong. Monitoring is routinely described as "building trust." In this parameterisation it does no such thing. Monitoring by the appropriators themselves — the specific condition, that the monitors are the users or are accountable to them — reduces `c_v`, the unit cost of verification, because observation is a by-product of use rather than a separate expedition. Lowering `c_v` moves `a*`:

`c_v`	`a*`	policy
0.20	0.8479	verify only below 85% reliance

c_v	a*	policy
0.12	0.9638	verify below 96%
0.06	0.9875	verify below 98.75%
0.02	0.9966	verify below 99.7%

Cheap monitoring raises the threshold above which one should stop verifying; it does not raise the reliance rate. The confusion produces the wrong intervention: a body that believes monitoring builds trust responds to a breach by monitoring more, which lowers `delta` through `v` at a cost in `C`, and leaves `a` exactly where it was. The condition that the monitors be accountable to the appropriators is what keeps the monitoring off the `C` line as pure overhead, and it is the part usually dropped in transfer [Cox et al., 2010].

25.6.3 6.3 Graduated sanction against expulsion, computed

Take the eight-firm consortium of Section 9.1 in its post-intervention state, in which firm 3 is the largest net creditor and a hypothetical rule-breaker. Two responses:

response	`rho` after	`N` after	`S` after	`eta S`	change
intact	0.8947	8	1.5347	0.6179	—
expulsion (maximal sanction)	0.8947	7	1.4159	0.5701	**-7.7%**
30% flow reduction, edge kept (graduated)	0.9038	8	1.5347	0.6242	**+1.0%**

Expulsion is priced entirely in `N`, and it is charged even when the expelled member was reciprocating perfectly. The consortium loses 7.7 per cent of its input by removing one of eight distinct participants, and it recovers none of that by having been right. The graduated response costs volume for two periods and returns a slightly *higher* `eta S`, because the reduction fell on the most imbalanced flows. **The fifth principle is not leniency. It is the observation that the sanction which destroys the relation charges the sanctioner the whole of `ln N − ln(N−1)`.**

Two qualifications, both disclosed rather than absorbed. First, the expulsion row holds `u` fixed at 0.52 across the removal, so the surplus after expulsion is `ln(1 + 6(0.52)) = 1.4159`. **That assumes departure-invariance in `u`, which Section 8.2 and Appendix A.1 (IF-05) show is the exception rather than the rule**: in a heterogeneous body the loss from removing the most divergent member is more than twice the loss from removing a random one, so the `−7.7%` figure is a *floor* on the cost of expulsion and not an estimate of it. Second, the post-sanction `rho = 0.9038` is stated in the source material without the modified flow matrix that would produce it, so **it cannot be independently recomputed from anything printed here**; the sign of the change is the defensible content and the magnitude is not.

25.6.4 6.4 Nesting changes the exponent

Communication paths in a group of `n` number `n(n−1)/2` [Brooks, 1975], and the structure of a body's communication tends to reproduce itself in what the body builds [Conway, 1968]. **Cost grows quadratically while surplus grows logarithmically**, so for any positive per-path cost there is an `n` beyond which no arrangement helps. With `eta = 0.40` and per-path cost `c_1 = 0.0009`:

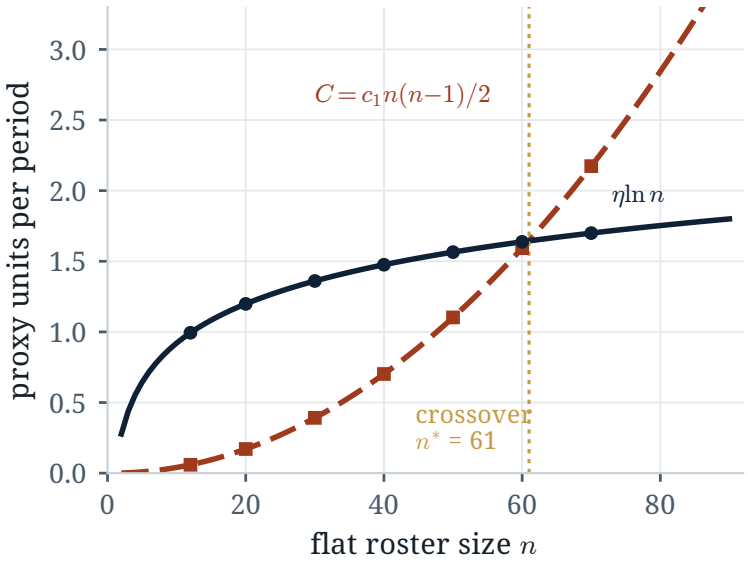


Figure 25.3: **Coordination cost overtakes the surplus ceiling at a computable roster size, and the crossover moves by a factor of five for a factor of twenty in the per-path cost.** Quadratic cost `C = c_1 n(n−1)/2` against the logarithmic ceiling `eta ln n` at `eta = 0.40` and `c_1 = 0.0009`; markers are the seven rows of the Section 6.4 table and the crossover is at `n = 61`. Both exponents are **defined by the model, not estimated from data**, so no power-law fitting procedure applies. Modelled from the parameters stated in Section 6.4; generated by `figs/src/CCP-025_fig03.py`.

n	paths	C = c_1 * paths	eta ln n	h at delta Q = 0
12	66	0.0594	0.9940	+0.9346
20	190	0.1710	1.1983	+1.0273
30	435	0.3915	1.3605	+0.9690
40	780	0.7020	1.4756	+0.7736
50	1225	1.1025	1.5648	+0.4623
60	1770	1.5930	1.6377	+0.0447
70	2415	2.1735	1.6994	**-0.4741**

Crossover at `n = 61`, and the crossover is exquisitely sensitive to the per-path cost:

c_1	0.0005	0.0010	0.0020	0.0050	0.0100
crossover n	85	58	39	23	16

A note on the exponents, required. The `n^2` in the cost and the `ln n` in the ceiling are **exact combinatorics and an exact logarithm within the model; neither is an estimated scaling exponent.** No distribution was fitted and no heavy tail is claimed, so the procedure that any estimated power-law claim would require [Clauset et al., 2009] does not apply here — and would apply immediately to any attempt to read `c_1` off observed organisational cost data, which is the step at which a real exponent would have to be estimated and reported with its uncertainty.

Now nest. The eighth principle records that large enduring systems in the sample were organised as *nested enterprises*: small units whose delegates form the next unit, and so on. Model it as span `g` and depth `d`, giving `N = g^d` members and `(g^d − 1)/(g − 1)` groups, each carrying `g(g−1)/2` internal paths:

span `g`	depth `d`	`N`	groups	nested paths	flat paths	ratio
3	3	27	13	39	351	1 : 9
3	4	81	40	120	3,240	1 : 27
4	3	64	21	126	2,016	1 : 16
5	3	125	31	310	7,750	1 : 25
6	3	216	43	645	23,220	1 : 36
8	3	512	73	2,044	130,816	1 : 64
10	4	10,000	1,111	49,995	49,995,000	1 : 1,000

Flat coordination cost is `O(N^2)`. Nested coordination cost is `O(N)`. That is not a coefficient improvement, it is a change of exponent, and it is the only structural move in the entire parameter set with that property. The surplus ceiling `eta ln N` is untouched by nesting — `N` is still `N` — so nesting is a pure reduction of `C` at constant ceiling.

Paths per member, holding `N` at 1000 and varying the span:

span `g`	3	4	5	6	8	10	12
paths per member	1.50	2.00	2.50	3.00	4.00	5.00	6.00

Coordination cost per member converges to `g/2` and therefore falls monotonically as the span narrows. **This does not mean spans of three are optimal**, and the model says why: narrowing the span increases depth, and depth is a delay and a distortion term that appears in the doctrine-to-data ratio, not in `C`. Each tier is an opportunity for doctrine to be reasserted and data to be summarised away. **The nesting law prices the coordination side only**, and the standing warning is that a strict tree — which is what `g^d` describes — is the wrong shape for exactly this reason: a semilattice, in which a unit belongs to more than one parent, preserves the cross-links a tree amputates [Alexander, 1965]. Those cross-links are where `u` lives, and the weak ties that bridge otherwise separated clusters are the ones carrying non-redundant information [Granovetter, 1973, Burt, 1992]. **Overlapping nests cost more paths than a tree and preserve the cross-links a tree destroys. The computation above is the tree's price list. It is the floor, not the recommendation.**

A parallel ceiling, reached from the cost side rather than the path-count side. Curzi [2026] derives a roster ceiling `m(m−1) <= 2 eta ln(N) / kappa` on the number of simultaneously maintained relations a body can carry, and reports a constructed commons sitting exactly at it — 30 against a ceiling of 30.7225, with a seventh member making the case structural rather than incidental. **That is the same trade this section computes, with the quadratic charged against pairs of maintained relations rather than against communication paths**, and the two arrive at the same conclusion by different arithmetic: the logarithmic ceiling is overtaken by a quadratic charge at a computable size, and the only structural escape is to change the exponent.

25.6.5 6.5 Why trust alone cannot rescue a low-reciprocity regime

This is the most useful result in the first half of the paper and it falls straight out of the steady-state theorem.

Take a relation with `eta S < C`. The steady state is `Q_inf = (eta S − C)/delta` and the numerator is negative. **Lowering `delta` — that is, raising trust — makes `Q_inf` more negative, not less.** The model is stating, in the plainest available language, that a relation losing order faster than it makes it cannot be repaired by making it lose order more slowly. It can only be repaired by raising `eta S` or lowering `C`.

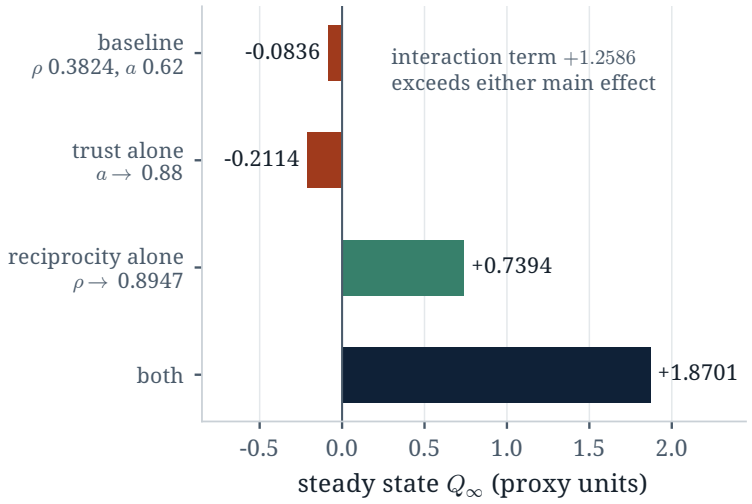


Figure 25.4: **A trust intervention alone moved the equilibrium the wrong way, and the interaction term is larger than either main effect.** One-at-a-time decomposition of the consortium's fixed point: baseline `−0.0836`, reciprocity alone `+0.7394`, trust alone `−0.2114`, both `+1.8701`, leaving an interaction of `+1.2586`. Trust alone is negative because at `rho = 0.38` the numerator `eta S − C` is negative and dividing a negative number by a smaller denominator produces a larger negative number. Modelled from the parameters stated in Sections 6.5 and 9.1; generated by `figs/src/CCP-025_fig04.py`.

Concretely, from the consortium of Section 9.1. Before: `rho = 0.3824`, `a = 0.62`, so `eta = 0.1721`, `delta = 0.4300`, `eta S = 0.2641`, `C = 0.30`, `Q_inf = −0.0836`. After the full intervention: `rho = 0.8947`, `a = 0.88`, `eta = 0.4026`, `delta = 0.1700`, `Q_inf = +1.8701`. Decompose the gain of `+1.9537` by moving one parameter at a time:

move	`Q_inf`	gain over baseline
baseline (both before)	-0.0836	—
reciprocity alone (`rho` to 0.8947, `a` still 0.62)	+0.7394	**+0.8229**
trust alone (`a` to 0.88, `rho` still 0.3824)	-0.2114	**-0.1278**
both	+1.8701	+1.9537
interaction term		+1.2586

Trust alone made it worse. Not marginally — it moved the equilibrium further into the region where the model declares the regime unsustainable, because at `rho = 0.38` the numerator was negative and dividing a negative number by a smaller denominator produces a larger negative number. And the interaction term, `+1.2586`, is larger than either main effect: **the two interventions are strongly complementary and neither is close to sufficient alone.** Curzi [2026] carries this inversion in the corpus's demotion register as R15 — the fixed point moving from `−0.0836` to `−0.2114` under trust alone against `+0.7394` under reciprocity — alongside the demotion of the moment equation on scale and the inversion of a tradition's stated order.

Read it as an operating rule. **Fix the numerator first.** Reciprocity, conversion, cost — the input side — before trust. A trust programme run in a body that is not converting its interactions into anything durable will produce a slower, more polite decline, and it will produce survey results saying the intervention worked.

25.6.6 6.6 Measurement protocols

Three protocols, one per quantity. Each is written to be executable by someone who did not write it, and each states its ethics constraint, because all three involve observing people.

Protocol A — the reliance rate `a`. *Estimates:* the probability that an unverified commitment in domain `d` on relation `(i,j)` is honoured within one period.

1. **Enumerate the commitment population.** List every commitment in domain `d` made on the relation in the last `k` periods. A commitment must have a date, a content, and a party. If the records cannot produce this list, that fact is itself the first finding: **a body that cannot enumerate its commitments cannot measure its trust, and its stored order is smaller than it believes.**
2. **Stratify by consequence** into low, medium and high tiers using a pre-declared rule.
3. **Randomise verification within the low tier only.** Select a random subset of the low-consequence commitments and do not verify them for one period. Verify all others as normal.
4. **Score.** `a_hat = (honoured, unverified) / (total, unverified)`. Compute a Wilson interval; with fewer than about 30 unverified commitments the interval will be wider than the effect being looked for, and the honest report is the interval.
5. **Repeat per domain.** Do not pool across domains.

Ethics constraint, binding. Randomised non-verification is permitted only in the tier where non-verification causes no harm the party has not already accepted, and **every party must be told in advance that verification is sampled rather than universal.** Constructing a hidden test of someone's reliability is a covert manipulation and is prohibited under the declination constraint — the party cannot decline what they are not told about. **The disclosure does not spoil the measurement**: what is concealed is *which* commitments are unverified, not that sampling occurs. The disclosure has a second and independent justification: participation in a study changes behaviour, and the effect is neither negligible nor uniformly directed [McCambridge et al., 2014], while an observer's undisclosed expectation is itself a known source of measured effect [Rosenthal and Jacobson, 1968, Rosenthal and Fode, 1963]. A concealed reliability test is therefore both prohibited and unsound.

Common failure: estimating `a` from complaints. Complaints measure the honoured rate of high-consequence commitments among parties who complain. That is three selection effects stacked.

Protocol B — the reciprocity index `rho`. *Estimates:* the matched fraction of interaction volume on a relation, over a declared window.

1. **Declare the unit.** Flows must be counted in a unit that is at least ratio-scaled with a true zero and that is *symmetric* — the same unit in both directions. Hours, items, referrals, review comments, defined favour-units. **If the unit is not symmetric — if "advice given" is counted one way and "money given" the other — there is no ratio and `rho` is not defined.** Demote to a paired ordinal comparison and say so.
2. **Declare the window before looking.** Default `w = 1/delta` periods (Section 5.3).
3. **Build the matrix** `x_ij` for every ordered pair.
4. **Compute** `rho_ij` per dyad, and `rho_group = sum matched / sum total` over all ordered pairs.
5. **Report the window sensitivity curve**, not only the point estimate.
6. **Report density alongside.** A network with three edges can have `rho = 1`. Reciprocity on a sparse network is a statement about the few edges that exist.

Common failure: computing `rho` from the node balances `b_i` instead of the dyad flows. That gives `1 − B`, which by Section 5.4 is an upper bound and is biased toward good news.

Protocol C — the obligation ledger. *Estimates:* the standing directed debt between every pair, and the network's imbalance.

1. **Same matrix as Protocol B.** One data collection serves both.
2. **Compute** `Skew_ij = (x_ij − x_ji)/2`, `b_i = sum_j (x_ij − x_ji)`, and `B = sum |b_i| / (2 * total)`.
3. **Rank the dyads by `|Skew|`, not by `Sym`.** The largest relations are not the most imbalanced ones, and it is imbalance that predicts breakdown.
4. **Check the sign structure.** A network with a small number of large net creditors and many small net debtors is a patron-client structure regardless of what it calls itself. A network with `b_i` near zero everywhere and high `Sym` is a horizontal association. These are the two cases of Putnam et al. [1993] and they are distinguishable in one afternoon of arithmetic.
5. **Do not publish individual balances without consent.** A ledger of who owes whom, made public inside a group, is a coercive instrument: it converts a private accounting into a public pressure and removes the ability to decline the accounting. Publish the *distribution*, and report dyad-level findings to the dyad.

25.6.7 6.7 Design specifications

Numbered, checkable requirements. A system claiming to have engineered trust, obligation or reciprocity must satisfy all of these or state which it fails and why. **These are renumbered from the source material's `S49.x` series as `S1`–`S12`; the scope specifications of Section 8.12 continue as `S13`–`S22` and the formulary specifications of Appendix B as `S23`–`S32`.**

S1 — Domain-indexed reliance. `a` shall be measured and reported per relation and per domain. A single organisation-wide trust figure shall not be published. Any report quoting one number shall state the domain it applies to.

S2 — Verification policy stated against `a*`. For each domain, `c_v`, `delta_0` and `K = eta S − C_0` shall be estimated, `a*` computed, and the verification rate set on the correct side of it. A verification policy never priced against `a*` is an unpriced recurring charge on `C`.

S3 — Sunset by default. Every commitment entering `Q` shall carry a review date. Absent one, the commitment inherits the relation's `a`, which for a well-run relation is high, and **an obsolete commitment in a high-`a` relation is immortal.**

S4 — Reciprocity window declared in advance. `w` shall be preregistered before flows are tabulated, with its justification. Post-hoc window selection can move `rho` from 0.13 to 0.87 on identical data (Section 5.3) and shall be treated as the researcher degree of freedom that it is.

S5 — Symmetric unit. The flow unit shall be identical in both directions of every dyad. Where no symmetric unit exists, `rho` shall be reported as undefined and the relation assessed by paired ordinal comparison, explicitly demoted.

S6 — Both indices reported. Any report of network reciprocity shall carry `rho` (dyad-level) and `B` (node-level), with the note that `1 − B >= rho`. Reporting only `B` is prohibited.

S7 — Numerator before denominator. No trust intervention shall be commissioned in a system whose measured `eta S − C` is negative. In that regime the intervention is predicted to worsen the equilibrium (Section 6.5). The order is: raise `rho` and `eta`, cut `C`, confirm the numerator is positive, then invest in `a`.

S8 — Graduated sanction as default. Sanctions shall be graduated, with expulsion reserved and its price stated in advance as `ln N − ln(N−1)` of surplus ceiling plus the lost `Sym` volume, and with the departure-fragility profile of Appendix A.1 (IF-05) computed rather than assumed away. A proposal to expel shall carry that computation.

S9 — Nesting above the crossover. Where flat membership exceeds the computed crossover `n` at the measured `c_1`, the body shall be nested, and the nesting shall be a semilattice — units belonging to more than one parent — rather than a strict tree [Alexander, 1965]. The tree figures of Section 6.4 are the cost floor, not the design.

S10 — Exit preserved. No obligation instrument shall bind a party who has not consented and may not withdraw. **This is not negotiable, and by Section 5.5 it is also the test that distinguishes a trust network from a capture network.** A network whose measured `a` is high and

whose measured `rho` is near zero shall be reported as a capture structure.

S11 — No covert reliability testing. Sampling of verification shall be disclosed. Which commitments are sampled need not be.

S12 — Ledger privacy. Individual obligation balances shall not be published within the group. Distributions only, and dyad-level findings to the dyad.

25.6.8 6.8 Failure modes

Ten named faults, renumbered from the source material's `FM-49.x` series as `F1–F10`. Section 8.12 continues with `F11–F18` and Appendix B with `F19–F26`.

F1 — The trust programme in a negative regime. *Signature:* an engagement or culture initiative running in a body whose meetings produce no artefacts and whose coordination cost is unmeasured; survey trust scores rise and nothing else does. *Damages:* wastes the budget and lowers `delta` while `eta S < C`. *Measurement:* compute `eta S − C` before commissioning. *Fix:* S7. *Confusable with:* F4.

F2 — Scalar trust. *Signature:* a dashboard with one trust number, usually from an annual survey, usually on a five-point ordinal scale, usually averaged. *Damages:* destroys the domain structure that makes `a` actionable; produces uniform policy where differentiated policy is required. *Measurement:* attempt to reconstruct `a_ij^(d)` for three domains; if the pooled figure lies between them and predicts none of them, the diagnosis is confirmed. *Fix:* S1 — and note that averaging ordinal responses was not licensed in the first place [Stevens, 1946].

F3 — Verification creep. *Signature:* a breach occurred, controls were added, the controls were never removed, `C` has been rising for three years and nobody can say by how much. *Damages:* `C`, permanently; frequently pushes `h` negative on its own. *Measurement:* price the current `v` against `a*` (S2). *Fix:* recompute `a*`; sunset every control added in response to a specific incident more than `1/delta` periods ago and not re-justified. *Confusable with:* legitimate monitoring, which lowers `c_v`. **The distinguishing test is whether the monitors are the users**: if monitoring is done by a separate function it is a `C` line; if it is a by-product of use it is a `c_v` reduction.

F4 — Reciprocity theatre. *Signature:* high interaction volume, celebrated; newsletters about collaboration; `rho` unmeasured. *Damages:* `eta`, invisibly — the dashboard measures `Sym` and the store depends on `rho`. *Measurement:* Protocol B. In the worked consortium, a volume of 68 units per quarter concealed `rho = 0.3824`. *Fix:* publish the matched fraction alongside the volume, always.

F5 — The node-balance illusion. *Signature:* a network reported as reciprocal because everyone's give and take roughly match in aggregate. *Damages:* conceals dyad-level one-directionality. *Measurement:* compute both; in the worked case, 0.4412 against 0.3824. *Fix:* S6.

F6 — Window shopping. *Signature:* a reciprocity figure that improved after the analyst "refined the measurement period." *Damages:* the credibility of the whole measurement. *Measurement:* ask when the window was chosen relative to seeing the data. *Fix:* S4, and report the sensitivity curve.

F7 — The immortal commitment. *Signature:* a rule everyone follows that nobody can date, justify, or find the author of; usually discovered when a newcomer asks why. *Damages:* `C`, and `u` — because an unexamined rule is a framing nobody is permitted to contradict. *Measurement:* sample the commitment population and score each for last review date; a distribution with a long right tail is the signature. *Fix:* S3. **A system with no forgetting cannot revise.**

F8 — Capture mistaken for culture. *Signature:* very high compliance, very high measured reliability, members who do not leave, flows one-directional toward a small number of nodes, and departure described as betrayal. *Damages:* everything, and the model will not show it if only `a` is measured. *Measurement:* `rho` and the sign structure of `b_i` (Protocol C, step 4). *Fix:* **there is no parameter fix. This is the failure mode the framework refuses to optimise.** The finding is a finding, and the required action is to restore exit.

F9 — Expulsion as first resort. *Signature:* a governance body whose only sanction is removal. *Damages:* `N`, at `ln N − ln(N−1)` per removal, charged whether or not the removal was justified. *Measurement:* the table of Section 6.3 — `−7.7` per cent of `eta S` for one of eight, and more where the departing member carries divergent framing. *Fix:* S8. **A sanction ladder cannot be constructed during the incident it is needed for.**

F10 — Flat above crossover. *Signature:* a body of sixty-plus that still meets as one room and still decides as one room, whose members describe the meetings as necessary and useless. *Damages:* `C`, quadratically. *Measurement:* estimate `c_1` from the observed per-path cost and find the crossover (Section 6.4). *Fix:* S9, with the semilattice caveat.

25.7 7 The internal/external correspondence

25.7.1 7.1 The master table and how to read it

The claim under audit is short enough to state before it is defended: **the internal and external scopes are not two subjects but one set of relations measured at two scopes.** The proven part of that claim is exactly one thing, and it is a statement about form. The recursion `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t` with `S_t = ln(1 + (N−1)U_t)` contains no term naming a substrate: there is no variable for "person" and no variable for "organisation". Every theorem about it — the exact solution, the one-step growth condition, the uniform bound, the sustainability ceiling, the steady state, the headroom gauge — is a theorem about a first-order linear recursion with a loss term, and holds wherever that recursion holds [Curzi, 2026q, Elaydi, 2005]. **That is a claim about model form, not about physics**, and this paper does not exceed it.

Two further proven facts do the structural work. **Banach's fixed point theorem**: for `delta > 0` the map is a contraction with modulus `1 − delta`, so a unique fixed point exists and iteration converges to it geometrically from any starting value [Banach, 1922]. **Strict concavity of `f` and the hard ceiling `0 <= F <= ln N`**: both are properties of the logarithm and neither knows what `N` is counting.

The thesis, badged as the author's thesis and not as a theorem: the internal and external scopes are the same relational system, and the parameters of one have measurable counterparts in the other. That is stronger than the proven statement and must be kept separate from it. The proven statement says the *equation* transfers. The thesis says the *phenomena* correspond — that forgetting and turnover are one parameter at two scopes, that character and institution are one store, that cognitive load and coordination overhead are one dissipation. **A modelling thesis of this kind is not established by exhibiting a table. It is established by cross-scope replication**: the same operationalisation, applied at both scopes, recovering parameters that stand in the predicted relation, with the prediction able to fail. Section 8.11 states exactly what would kill it. Until that programme is run, the correspondence is a research hypothesis with a good structural argument behind it and no cross-scope replication behind it, and this paper says so in those words.

Term	INTERNAL (personal)	EXTERNAL (interpersonal)
`N`	distinguishable categories in one mind	distinct domains held in a group
`u`	non-parallelism between one's own framings	non-parallelism between members
`N_eff`	how many genuinely different views are held	how many genuinely different members
`F`	internal surplus	group surplus
`Q`	character, skill, memory, internalised structure	institution, culture, infrastructure

Term	INTERNAL (personal)	EXTERNAL (interpersonal)
`delta`	forgetting, skill atrophy, drift of self	turnover, decay of practice, lost tradition
`eta`	how experience becomes character	how interaction becomes institution
`C`	cognitive load, effort, maintenance of self	coordination overhead, friction
`h`	personal capacity for the new	organisational headroom
`M`	the moment as lived	the situation as the group constructs it
`D/P`	frame-to-intake ratio	doctrine-to-data ratio

One reading instruction before the rows. The correspondence is not that the two columns describe the same thing. It is that they occupy the same position in the same equation. Internal `delta` and external `delta` are different phenomena — forgetting is not turnover — but they enter the recursion identically, are estimated by the same procedure, and have the same consequences for the trajectory. That is the whole claim, and it is both weaker and more useful than the claim that a company is a mind.

25.7.2 7.2 The counting rows: `N`, `u`, `N_eff`

Row `N`. Internally, `N` is the number of distinguishable categories the mind holds in the domain under analysis — not the number of thoughts, the number of *kinds* it can tell apart. *Measurement:* a free-sort task; present 40–60 items from the domain, ask for as many groups as feel natural with a name per group, and take `N` as the count of groups surviving a re-sort a week later. **The re-sort is the whole protocol**: categories that do not reappear were produced by the task, not held by the person. Report test-retest agreement; below about 0.6 the count is not a measurement [Cronbach, 1951]. Externally, `N` is the number of distinct domains the group holds — not headcount. *Measurement:* the capability inventory; enumerate the domains in which at least one member can act without consulting anyone outside the group, each verified by an observed independent instance.

Where the two diverge in kind. External `N` is enumerable from public evidence and can be audited by a third party. Internal `N` is a modelling choice about how finely to carve, and two competent analysts will return different counts from the same person. **The external count has a fact of the matter that the internal count does not**, and the asymmetry runs opposite to most intuitions: it is easier to count what a team can do than what a mind can distinguish.

Row `u`. Internally, the degree to which a person's own framings of a situation are genuinely different rather than the same framing in different words, `u = 1 − (x.y)^2` per pair with the mean over pairs. *Measurement:* the independent-prediction test — produce `k` framings of one situation, then, separately and without re-reading, ask each framing for a prediction about a specific future observable, and estimate `u` from the disagreement structure. **Framings that predict identically are the same framing.** That is the only defensible operationalisation, because framings stated in words are trivially made to look different and non-trivially made to be different. Externally: the same instrument with different subjects, members predicting independently before discussing.

Where the two diverge in kind. Internally all framings are held by one processor and share its priors; they are correlated by construction and `u` is bounded well below 1 in practice. Externally `u` near 1 is achievable because members can be independently sourced. **The internal ceiling on `u` is lower than the external one**, and the corpus's result that the second genuinely different framing is the largest gain ever bought is a consequence of that lower ceiling rather than a restatement of the concavity [Curzi, 2026o].

Row `N_eff`. `N_eff = 1 + (N−1)u` at both scopes without modification — **the framework's own**

affine effective count, not a Hill number (Section 3.1). This row is the cleanest correspondence in the table because it is pure arithmetic on the two rows above it. With `N = 7` at either scope:

u	N_eff	S = ln N_eff	share of ceiling ln 7 = 1.9459
0.15	1.900	0.6419	33.0%
0.35	3.100	1.1314	58.1%
0.55	4.300	1.4586	75.0%
0.75	5.500	1.7047	87.6%
0.95	6.700	1.9021	97.7%

Seven framings held by one person at `u = 0.15` and seven people in a room at `u = 0.15` are both worth 1.9 of anything. The arithmetic does not care which. That is the correspondence in its most literal and least contestable form: **the audit for redundancy is scope-independent**, because redundancy is a geometric property and geometry does not know what it is measuring. `N_eff` is derived and is never to be measured by a separate instrument.

25.7.3 7.3 The dynamical rows: `F`, `Q`, `delta`, `eta`, `C`

Row `F`. Ceiling at both scopes: `ln N`. **Here the correspondence produces a result neither scope would produce alone.** Since the moment normalises into `[0, ln N]` to serve as `S_t`, the maximum possible moment is `ln N`, which is the law of requisite variety [Ashby, 1956, 1958] appearing in first-person experience: **the richest available moment is bounded by the log of how many kinds of thing the perceiver can distinguish.** That statement was not chosen. It follows from putting the internal and external ceilings in the same equation, and it is the strongest single piece of evidence for the thesis. `F` is derived at both scopes and never measured directly.

Row `Q`. Internally: character, skill, memory, internalised structure — what persists in the person when the situation that built it has gone. *Measurement:* the unprompted-retention probe; sample the domain, wait one period without input, test performance without cues. **The essential feature is the absence of prompting**: capability appearing only when scaffolded is not in `Q`, it is in the scaffold, which belongs to the environment. Externally: institution, culture, infrastructure — what persists in the group when the people who built it have gone. *Measurement:* the identical probe with the identical essential feature, the substitution test: replace the individual and see whether the capability survives. Star and Ruhleder [1996] supplies the diagnostic principle for both scopes — **infrastructure becomes visible only on breakdown** — so `Q` is most accurately measured during a disruption and is systematically underestimated during normal operation.

Where the two diverge in kind. Internal `Q` has a hard terminus. External `Q` does not: an institution's store can in principle outlive every person who has held it, and often does. **A mind's store is bounded by a lifetime and a group's is not, so the two scopes have different upper limits on the time constant `1/delta` and therefore different meaningful ranges for `delta` itself.** That is the first crack, and Section 8 develops it.

Row `delta`. Internally: forgetting, skill atrophy, drift. *Measurement:* measure `Q` at two times under zero input, `delta = 1 − Q_1/Q_0`, half-life `ln 2 / −ln(1−delta)`. Externally: turnover, decay of practice, lost tradition — **and here the external scope has a term with no internal counterpart**, which must be composed rather than assumed:

```
delta_total = 1 − (1 − delta_internal)(1 − tau)
```

where `tau` is the per-period turnover rate. The two losses are independent, so their survivals multiply. **The independence is a postulate, badged**, and it is violated in the specific case where departures are *caused by* the decay — the exit feedback of Section 8.1 — in which case the composition understates the loss.

delta_int	tau	delta_total	Q_inf relative to tau = 0
0.05	0.00	0.0500	100.0%
0.05	0.05	0.0975	51.3%
0.05	0.10	0.1450	34.5%
0.05	0.20	0.2400	20.8%
0.05	0.30	0.3350	14.9%
0.05	0.50	0.5250	9.5%

Five per cent turnover halves steady-state stored order and ten per cent cuts it to a third, before anyone has written a single bad document. This row is where the external scope is strictly harder than the internal one: knowledge does not resign. *Shared caution:* `delta` is per period, and the period lengths at the two scopes typically differ by one to two orders of magnitude. Any cross-scope comparison without normalisation is meaningless; Section 8.7 computes exactly how meaningless.

Row `eta`. Internally, how experience becomes character; `eta = dQ_attributable / M_input`. Externally, how interaction becomes institution; `eta = dQ_attributable / S_input`. Same estimator, same difficulty — attribution is the whole difficulty at either scope — and one external advantage: **the artefacts are countable.** A decision log exists or it does not. **Meetings with no artefact have `eta` approximately 0**, and that is checkable from outside in a way that "did that experience change you" is not.

Where the two diverge in kind, and this is a real asymmetry rather than a measurement inconvenience. Section 5.2 established that at group scope `eta = eta_max rho`: only the reciprocated fraction converts, because structure only one party holds is a debt. **There is no internal analogue of `rho`, because framings cannot renege.** A framing has no interests, cannot defect, and cannot fail to return a favour. The entire apparatus of reciprocity, sanction and obligation transfers inward as metaphor only (Section 8.4).

Row `C`. Internally: cognitive load, effort, the ongoing maintenance of a self — the cost of holding the structure up, paid every period whether or not anything is built. *Measurement:* dual-task interference, or time-on-maintenance sampling. Both are indirect, and **the honest position is that internal `C` is the least well-measured parameter in the table.** Externally: coordination overhead, friction, conflict, compliance. *Measurement:* direct — coordination cost is a line item — and it has a known scaling law that the internal scope does not obviously share, `n(n−1)/2` communication paths [Brooks, 1975, Conway, 1968], giving the computable crossover of Section 6.4. The stock-and-flow vocabulary for the whole row is standard [Forrester, 1961, Sterman, 2000], and the behaviour of `h` near exhaustion has the same shape as utilisation approaching one in a queue, where waiting time diverges [Kingman, 1962, Hopp and Spearman, 2011].

Where the two diverge in kind. External `C` has the quadratic term because members must coordinate pairwise. Internally the framings do not negotiate with each other in that way — a mind is not paying `n(n−1)/2` in inter-framing meetings. **The external scope has a cost law with an exponent; the internal scope, as far as the framework can currently say, does not, and this is a structural difference rather than a difference in parameter value.**

25.7.4 7.4 The gauge rows: `h`, `M`, `D/P`

Row `h`. `h = eta ln N − delta Q − C` at both scopes, unmodified, with the three regimes identical: `h < 0` unsustainable in this model and no geometry saves it; `h = 0` maintenance only; `h > 0` growth available with any requested increment `r <= h`. Internally, the reason a person at high `Q` and high `C` cannot take on one more thing is not weakness of will; it is `h < 0`, and the structural result holds identically at both scopes: **success consumes its own headroom**, because `delta Q` grows with `Q`. Externally, the ossification of successful institutions is the same sentence with different nouns; the reserve that lets a body absorb variation is what Cyert and

March [1963] called organisational slack. *Measurement, both scopes:* derived from the four parameters, **never surveyed.** "Do you have capacity" is a question about perceived `h`, and perceived `h` is exactly the quantity the doctrine-to-data fault distorts. Curzi [2026h] supplies the identity that closes this row: at equilibrium, headroom is unused variety, `h_inf = eta(ln N − F(u*))`, verified to six decimals at `eta = 0.35`, `delta = 0.015`, `N = 6`, `u* = 0.55` giving `Q_inf = 2.840970` and `h_inf = 0.164501`. **This row is where the correspondence pays for itself most directly, because the same diagnostic question — is the maintenance term eating the ceiling — is asked identically of a person and of a department and answered by the same arithmetic.**

Row `M`. `M = (D/P) E`, the moment as lived internally and the situation as the group constructs it externally, with `D` the shared doctrine, `P` the collective perception and `E` the environment. **Magnitudes on this row are demoted (Section 3.4); what follows is the structure, which is not.** The four limiting cases are the diagnostic instrument and they are identical at both scopes:

```
P -> 0    M -> infinity   frame with no intake   internal: dogma, apophenia
                                                 external: a body that has stopped listening
D -> 0    M -> 0          intake with no frame   internal: sensation without experience
                                                 external: a body that measures everything
                                                           and concludes nothing
E -> 0    M -> 0          no environment         both: no situation, no moment
D = P     M = E           THE FIXED POINT        both: undistorted contact
```

The fixed point `D = P` is the same fixed point at both scopes, and it means the same thing: when the frame matches the intake, the moment equals the environment.

Row `D/P`. `r = D/P`, distortion `Delta = |ln r|`, in the same logarithmic register as everything else. Internally, frame-to-intake: `r > 1` over-framed (seeing the expected), `r < 1` under-framed (flooding). Externally, doctrine-to-data: **an over-framed institution is one whose doctrine outruns its intelligence; an under-framed institution drowns in data it cannot read.** Bands at both scopes:

`|ln r|`	`r` (over)	`r` (under)	reading
< 0.10	< 1.105	> 0.905	calibrated
0.10–0.40	1.105–1.492	0.905–0.670	drift
0.40–1.00	1.492–2.718	0.670–0.368	fault
> 1.00	> 2.718	< 0.368	severe

The band table is a construction showing how the measure reads, not a set of measurements, and `Delta` is an ordinal quantity pending instrumentation: it ranks bodies and periods by distance from contact and does **not** license the claim that one distortion is twice another. **The sign of `ln r` is unaffected by the demotion and is the operationally decisive part**, because over-framing and under-framing take opposite corrections [Curzi, 2026]. *Measurement, both scopes:* the prediction-surprise ledger — record predictions with confidence before the observation, score them after. Over-framing shows as high confidence with flat accuracy; under-framing as low confidence with accuracy that would have supported more.

Where the two diverge in kind. An institution's `D` is a public document and its `P` is a measurable throughput; a person's `D` and `P` are both private. **The external version of this row is the more measurable one, which is the reverse of the intuitive expectation: doctrine-to-data is easier to measure than frame-to-intake, because doctrine is written down.**

Two things transfer with this row and both correct how `r = 1` reads in the band table.

First: at the fixed point, headroom is exactly zero, at either scope. Put `r = 1` into the contact relation and the input sits at its ceiling, `S = ln N`. The steady state is then `Q_inf = (eta ln N − C)/delta`, and the headroom gauge reads

```
h = eta ln N − delta Q_inf − C = eta ln N − (eta ln N − C) − C = 0
```

with `delta` and `C` cancelling out. **Perfect contact drives the input to its maximum, the maximum input drives the store to its maximum, and the maximum store spends the entire budget on maintenance.** That this comes out identically at both scopes is among the cleaner pieces of evidence for the thesis: it is a cancellation in the algebra, not a resemblance in the prose. **The consequence at both scopes: `r = 1` is not an unqualified target.** It maximises fidelity and leaves nothing to build with. Room at a settled state requires shedding `Q` deliberately or holding a declared `r != 1` as policy.

At external scope the fixed point acquires an exact join to the surplus functional, `ln r = ln D − ln p_bar − F`. **The group's surplus enters the doctrine-to-data ratio one nat for one nat**, so calibration becomes a geometry requirement rather than an attitude: a body carrying doctrine load `D/p_bar` cannot reach `r = 1` at all unless `D/p_bar <= N`, and the residual floor `Delta_min = ln(D/p_bar) − ln N` is where heavy institutions live. **This join is available only at external scope** — the internal analogue would need a decomposition of one person's intake into countable channels, which Section 8.3's aggregation disanalogy says is not on offer. Record it as a place where the correspondence is **asymmetric in what can be computed**, not merely in what can be measured. Both sides of the join carry the demotion of Section 3.4.

Second: the ledger this row drives must be kept in two tracks, at both scopes. Over-framing **increases** stored order — the storage law integrates the moment as lived and contains no term asking whether it was earned — so the instrument is `Q` against veridical `Q^V`, with stored error `Z_{t+1} = (1−delta)Z_t + eta(M_t − V_t)` (the cost term cancels exactly) and correction load `C_t = c_0 + kappa Z_t`. **Dogma banks faster than knowledge, and the bill arrives on a lag.** Because `V = E min(r,1)` is flat above the fixed point while `M = rE` keeps climbing, **anyone optimising raw `Q` is driven away from `r = 1`, actively and by arithmetic**: a person accumulating a confident self-account, an institution accumulating doctrinal coverage, identically. The knife edge is `kappa = delta` — below it raw `Q_inf` rises without bound in `r`. **Row `Q` should therefore be read as two rows**, and a report of total store without `Z` beside it is a report of a number that grows fastest when the reporter is most wrong. This is the framework's own instance of the general result that a measure which becomes a target ceases to be a measure [Goodhart, 1975, Campbell, 1979].

25.7.5 7.5 What the eleven rows share: one trajectory shape

Every row above is a parameter of the same recursion, and the recursion is a contraction whenever `delta > 0`. Banach therefore guarantees that both scopes converge geometrically to their own fixed points, and that **after normalising time by the time constant `tau = 1/delta`, the two trajectories are the same curve.**

Computed. Internal: period one day, `delta = 0.02`, `eta = 0.30`, `M = 1.20`, `C = 0.15`, so `tau = 50` days and `Q_inf = 10.5000`. External: period one quarter, `delta = 0.16`, same `eta`, `S` and `C`, so `tau = 6.25` quarters and `Q_inf = 1.3125`. Both from `Q_0 = 0`, with `t` rounded to whole periods at each normalised point:

`t / tau`	internal `Q/Q_inf`	external `Q/Q_inf`	continuous limit `1 − e^−k`
0.5	0.3965	0.4073	0.3935
1.0	0.6358	0.6487	0.6321
2.0	0.8674	0.8766	0.8647
3.0	0.9517	0.9636	0.9502
4.0	0.9824	0.9872	0.9817
5.0	0.9936	0.9955	0.9933

The two columns agree to within two per cent at every point and both track the continuous exponential. **The steady states differ by a factor of exactly eight and the shape is identical. The form transfers; the clock does not, and the levels do not.**

A correction to the source material, disclosed rather than absorbed. The source states that "the clocks differ by a factor of about a hundred and forty." **That figure is not recoverable from the stated parameters.** The time constants are 50 days and 6.25 quarters: a ratio of **8** measured in each scope's own periods, and a ratio of **11.4** in absolute time at 91 days to the quarter. Neither is 140, and no combination of the stated parameters produces 140. The corrected figures are printed above and the erroneous one is not repeated.

25.8 8 Where the correspondence breaks: nine audited failures

Nine failures: three severe, four structural but manageable, one a matter of measurement practice, and one absolute. Each is stated as sharply as it can be stated, with its computation where it has one, and with the transfer restriction it implies. **The labels B1–B9 are carried from the source material unchanged, since they do not collide with any result label in this paper.**

25.8.1 8.1 B1 — Exit. A group's members can leave; a mind's parts cannot. SEVERE

This is the deepest disanalogy and everything else in this section is downstream of it.

A framing does not resign. A category does not withdraw its labour. A memory does not accept an offer from a competitor. But every member of a group is a separately deciding agent with an outside option, **and that fact changes the type of the model, not just its parameters.**

Consequence 1: `N` is endogenous externally and exogenous internally. In the internal model `N` is set by the analyst's carving and stays fixed while the trajectory runs. In the external model `N` is an outcome of the trajectory: a group whose `Q` falls loses members, which lowers `N`, which lowers the ceiling `eta ln N`, which lowers `Q` further. **The external system has a feedback loop through `N` that the internal system does not have at all.** Every theorem in the corpus's dynamical core is proved for fixed `N`. **None of them is proved for the closed loop.** That loop is also precisely where the identification problem of Section 3.3 bites hardest, since membership is then simultaneously cause and effect [Manski, 1993].

Consequence 2: group-level sustainability does not imply member-level sustainability. Headroom is computed for the body; each member has their own,

```
h_i = eta ln N / N − delta Q / N − c_i
```

Computed for `N = 9`, `eta = 0.35`, `delta = 0.08`, `Q = 4.0`, with heterogeneous cost shares. The per-member share of the ceiling is `0.35 ln 9 / 9 = 0.08545`; the per-member share of maintenance is `0.08 x 4.0 / 9 = 0.03556`; so the critical private cost is `c_i* = 0.04989`:

member	`c_i`	`h_i`	status
1	0.02	+0.02989	ok
2	0.03	+0.01989	ok
3	0.05	-0.00011	unsustainable
4	0.09	-0.04011	unsustainable
5	0.14	-0.09011	unsustainable
6	0.22	-0.17011	unsustainable

Four of the six tabulated members are individually below water in a body that may well be reporting positive aggregate headroom. They will leave, and they will leave in cost order — the highest-cost members first, which are frequently the members carrying the highest `u`, because carrying a divergent framing in a converged room is expensive. **The internal model**

has no equation that can even express this, because a mind's subsystem cannot have a private balance sheet.

A defect in the source material, disclosed. The per-member shares above are computed at `N = 9`, but only **six** members are tabulated, and the source then reports the result as "four of six". **The denominator of the headroom shares and the denominator of the reported fraction are different numbers.** The arithmetic of each row is correct as printed and is reproduced unchanged; what is wrong is the reading. The defensible statement is that **four of the six members whose private costs are exhibited are below water**, and the fraction of the nine-member body below water is not determined by the table. The source also states `u = 0.6` for this computation, which does not enter `h_i` at all — `h_i` is built from `ln N`, the ceiling, not from `ln N_eff`.

Transfer restriction: **no result derived at internal scope may be transferred outward without checking member-level headroom separately. The aggregate is not sufficient.**

Named in the literature: Hirschman [1970] is the standing treatment of the three options open to a dissatisfied member of an organisation. Two of the three — exit, and loyalty as a mediator of exit — have no internal counterpart at all.

25.8.2 8.2 B2 — Heterogeneity. A group's parameters are distributions; a scalar model is a mean field. SEVERE

The storage law as written has one `delta`, one `eta`, one `C`. A group has one per member. Writing a single value is a **mean-field approximation**, and mean-field approximations fail in a specific, computable way when the underlying distribution is not tight. The general objection is old and unresolved: the behaviour of an aggregate is not the behaviour of a representative member [Kirman, 1992], and the same warning attaches to every aggregated social model [Durlauf, 2005].

Computation 1: the aggregate `delta` is not any member's `delta`, and it is not even constant. Take a store split evenly between two compartments, `Q_A = 50` at `delta_A = 0.05` and `Q_B = 50` at `delta_B = 0.50`. The `Q`-weighted mean is `delta_bar = 0.275`. Run both with no input:

t	true total	mean-field prediction	error	error %	implied aggregate `delta`
0	100.0000	100.0000	0.0000	0.00%	—
1	72.5000	72.5000	0.0000	0.00%	0.2750
2	57.6250	52.5625	+5.0625	+8.79%	0.2052
3	49.1187	38.1078	+11.0109	+22.42%	0.1476
5	40.2515	20.0304	+20.2211	+50.24%	0.0821
8	33.3663	7.6332	+25.7332	+77.12%	0.0550
12	27.0302	2.1089	+24.9213	+92.20%	0.0504

The mean-field model is exactly right for one period and wrong by 92 per cent after twelve. The reason is structural: the fast compartment empties, so the surviving store is increasingly the slow one, and the implied aggregate `delta` drifts from 0.275 down toward 0.05. **A single `delta` fitted to early data will catastrophically over-predict decay, and a single `delta` fitted to late data will catastrophically under-predict it. Neither is a parameter of the system; both are artefacts of when the analyst looked.**

Computation 2: aggregate `u` does not identify the structure, and the structures differ in fragility. Model each member as `v_i = cos(th_i) e_0 + sin(th_i) f_i`, with `e_0` the shared direction and the `f_i` mutually orthonormal private directions. Then `u_ij = 1 − c_i c_j` with `c_i = cos^2(th_i)`, and the mean pairwise `u` follows in closed form.

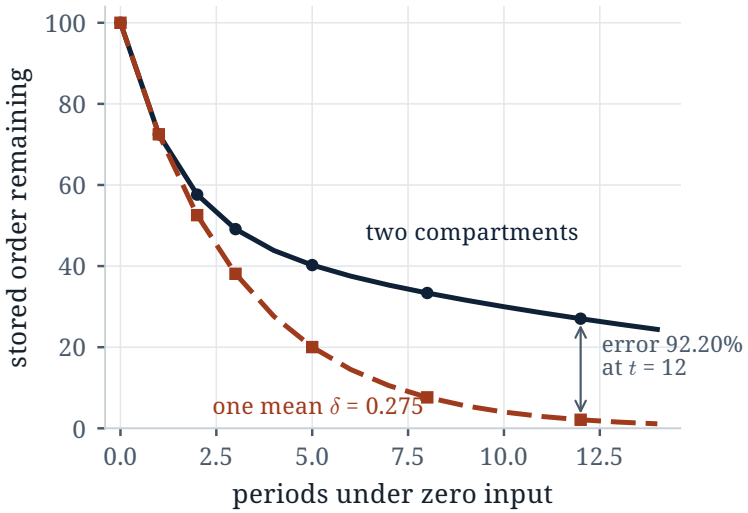


Figure 25.5: **The mean-field scalar is exactly right for one period and wrong by 92 per cent after twelve.** A store of 100 split evenly between a slow compartment at `delta = 0.05` and a fast one at `delta = 0.50`, run under zero input, against the single Q-weighted rate `delta_bar = 0.275`. The fast compartment empties, the surviving store becomes increasingly the slow one, and the implied aggregate rate drifts from 0.2750 at `t = 1` down to 0.0504 at `t = 12`. Modelled from the parameters stated in Section 8.2; generated by `figs/src/CCP-025_fig05.py`.

configuration	composition	mean `u`	`S` at `N = 9`
A, homogeneous	nine members all at `c = 0.6325`	0.5999	1.7578
B, heterogeneous	six at `c = 0.95`, three at `c = 0.0501`	0.5999	1.7578

Identical aggregate `u`. Identical surplus. Now remove three members:

configuration	which three	remaining `u`	remaining `S` at `N = 6`
A	any three	0.5999	1.3862
B	the three divergent	0.0975	**0.3971**
B	three conformists	0.7904	1.5998

Configuration A is departure-invariant in `u`. **Configuration B loses 77.4 per cent of its surplus to the departure of three of nine, where configuration A loses 21.1 per cent** — and B *gains* if the right three leave instead. The two systems are indistinguishable from the aggregate.

A second defect in the source material, disclosed. The source reports the ratio of the two damages as "three and a half times", and Appendix A.1 (IF-05) states it as "3.49 times the damage". **That figure is the ratio of the two *surviving* surpluses, `1.3862 / 0.3971 = 3.4909`, not the ratio of the two losses.** The losses are `1.7578 - 1.3862 = 0.3716` and `1.7578 - 0.3971 = 1.3607`, so **the damage ratio is `1.3607 / 0.3716 = 3.66`**, which is also what the paired percentages `77.4%` and `21.1%` printed in the same paragraph imply. The corrected figure is 3.66 and it is used from here on. The same entry computes its *other* ratio — targeted against random single removal, `0.2599 / 0.1144 = 2.27` — correctly on losses, so the two ratios in that entry are computed by two different methods.

This is a failure of identifiability at external scope only, because it is a failure about who leaves, which is B1 again. The parallel result in network science is standard: heterogeneous structures are robust to random failure and fragile to targeted removal [Albert et al., 2000, Cohen et al., 2000, Callaway et al., 2000], and the arithmetic above is that result appearing in surplus geometry rather than in a connectivity graph. **The transfer is a modelling assumption; the arithmetic is exact.**

Transfer restriction: **at external scope, `u` shall be reported with its distribution across members, never as a scalar alone.** At internal scope the scalar is defensible because the framings cannot separate.

25.8.3 8.3 B3 — Aggregation is not always well defined. SEVERE

The `D/P` row assumes a group *has* a doctrine — that there is a coherent `D` to put in a numerator. For a single mind this is usually approximately true. For a group it can be false as a matter of arithmetic rather than of sloppiness.

Three members, three options, these preference orders:

```
member 1:  A > B > C
member 2:  B > C > A
member 3:  C > A > B
```

Pairwise majorities, computed:

```
A vs B:  2–1  ->  A
B vs C:  2–1  ->  B
A vs C:  1–2  ->  C
```

The group prefers A to B, B to C, and C to A. There is no transitive collective ordering to extract. This is the Condorcet cycle, and Arrow's impossibility theorem generalises it: no rule aggregating three or more ordinal preferences can simultaneously satisfy unrestricted domain, non-dictatorship, the Pareto criterion, and independence of irrelevant alternatives. **Neither result has a bibliography entry available to this paper, so both are stated as classical results without a citation key rather than attached to a source that does not support them; the cycle above is arithmetic a reader can redo in three lines.**

The consequence is exact: the group's `D` may not exist as a well-defined object, in which case `D/P` is not merely hard to measure but undefined. A group can be cyclically incoherent while every member is individually coherent.

The honest counter-weight, stated because the disanalogy should not be overclaimed: individual preferences also exhibit intransitivity, and preference reversals under reframing are among the more robust findings in judgement research. So the disanalogy is one of degree and of mechanism rather than of kind. **But the mechanism differs in a way that matters: individual intransitivity is an error, and collective intransitivity can be the correct output of a faultless procedure applied to faultless inputs.** No amount of member-level rationality removes it.

Transfer restriction: **before computing a group `D/P`, test for cyclicity.** If the collective ordering over the relevant options is cyclic, report `D` as undefined and diagnose by member-level `D_i/P_i` distributions instead.

25.8.4 8.4 B4 — Reciprocity has no internal analogue. STRUCTURAL

Section 7.3 stated it: `eta = eta_max rho` externally, and internally there is no `rho`, because a framing cannot renege. Extending the machinery of Sections 5 and 6 inward requires positing internal agents with interests — and then a second layer of agents to explain the first.

The trap is easy to fall into precisely because the equations permit it. The storage law will accept an "internal reciprocity" parameter without complaint; the equation has no way to object. **What objects is the measurement:** `rho` requires a flow matrix `x_ij` between agents, and there are no internal agents whose flows can be counted.

The partial exception, stated honestly: intertemporal-choice research does model a person as a sequence of selves whose interests conflict, and the technology of precommitment has an obvious structural resemblance to a contract between parties. That is a real modelling tradition and it is not obviously wrong. **But the resemblance is to contract, not to reciprocity: the later self did not choose to enter the arrangement and cannot decline it, which is B1 again.** A framework whose governing refusal is *never remove a person's ability to decline* must notice that precommitment is precisely the removal of a future person's ability to decline, and that this is permitted only because the two persons are the same person. **No entry in this corpus's bibliography supports the intertemporal-choice literature referred to here, so it is named descriptively and left uncited rather than attached to a source that does not carry it.**

Transfer restriction: **Sections 4 to 6 transfer outward-to-inward as metaphor only.** Nothing in their arithmetic — `rho`, the obligation ledger, the sanction ladder, `a*` — has an internal measurement, and none of it should be given one.

25.8.5 8.5 B5 — The observation asymmetry. STRUCTURAL

Internal measurement has privileged access and unreliable instruments: introspection is available continuously, at no cost, and is **biased in a direction correlated with the quantity being measured.** A person who is over-framed (`r > 1`) is by definition seeing what their frame predicts, including when they look at their own frame. **The error is not noise; it is signal-dependent distortion**, and the unreliability of self-report about the causes of one's own judgements is well established [Nisbett and Wilson, 1977].

External measurement has public access and sparse, noisy instruments: behaviour observed at cost, at intervals, through selection filters. **The error is large but comparatively unbiased**, or at least biased in ways that are themselves observable.

The consequence is that the same parameter has different identifiability profiles at the two scopes and therefore requires different estimators — not the same estimator with a different sample size. Internal `C` is the clearest case: Section 7.3 conceded it is the least well-measured quantity in the table, and the reason is that the cost of maintaining a self is paid by the same apparatus that would report it.

Transfer restriction: **an operationalisation validated at one scope is not thereby validated at the other.** Validation is per scope, and it is the specific gate most often skipped [Campbell and Fiske, 1959, Cronbach and Meehl, 1955].

25.8.6 8.6 B6 — Boundary definition runs the opposite way at the two scopes. STRUCTURAL

Noted in Sections 7.2 and 7.4 and collected here because it is systematic and counter-intuitive:

quantity	easier internally	easier externally
`N`		yes — capability inventory is auditable
`u`		yes — independent prediction rounds
`Q`		yes — substitution test
`delta`	yes — no turnover term to disentangle	
`eta`		yes — artefacts are countable
`C`		yes — coordination cost is a line item
`M`	yes — the observer is the one having it	
`D/P`		yes — doctrine is written down

quantity	easier internally	easier externally

Six of eight are easier externally. The common assumption is the reverse — that we know our own minds and can only guess at institutions. **The framework's measurement requirements invert that assumption almost completely**, and this is one of the more useful things the correspondence produces: to learn how to measure an internal parameter, look at how the external version is measured, because the external version is usually the better-instrumented one.

25.8.7 8.7 B7 — Timescale. `delta` is per period and the periods differ by two orders of magnitude. PRACTICE

The scope transform is `delta_T = 1 - (1 - delta)^T` for aggregating `T` short periods into one long one. **It is exact and it is routinely omitted.**

`delta` per period	`T = 7`	`T = 30`	`T = 52`	`T = 365`
0.001	0.0070	0.0296	0.0507	0.3059
0.005	0.0345	0.1396	0.2295	0.8395
0.010	0.0679	0.2603	0.4070	0.9745
0.020	0.1319	0.4545	0.6503	0.9994
0.050	0.3017	0.7854	0.9306	1.0000
0.100	0.5217	0.9576	0.9958	1.0000

And the trap, stated with its number. A `delta` of 0.02 at internal scope with a daily period, and a `delta` of 0.02 at external scope with a quarterly period, print identically:

```
internal 0.02/day     compounded over 91 days   ->  0.8409 per quarter
external 0.02/quarter expressed per day         ->  0.000222 per day
ratio of implied daily rates                    ->  90.1x
```

Two numbers that print identically differ by a factor of ninety.

Transfer restriction: **`delta` and `eta` shall never be quoted without their period; `Q`, `C` and `h` inherit the requirement.** Only `N`, `u`, `N_eff` and `F` are period-free.

25.8.8 8.8 B8 — Rival versus non-rival stored order. STRUCTURAL, with no internal counterpart

Whether members exit depends on what they get, and what they get depends on whether `Q` is rival.

Rival case (the store is divided): each member's per-period return is `eta ln(N_eff)/N - c_i`. With `eta = 0.35`, `c_i = 0.020`, `u = 1` so that `N_eff = N`, and outside option `w = 0.030`:

`N`	`S = ln N`	per-capita	return	decision
3	1.0986	0.12817	+0.10817	stay
6	1.7918	0.10452	+0.08452	stay
12	2.4849	0.07248	+0.05248	stay
20	2.9957	0.05243	+0.03243	stay
30	3.4012	0.03968	+0.01968	**exit**

`ln(x)/x` is maximised at `x = e = 2.71828`. **Under equal division of a logarithmic surplus, the per-capita return is maximised at a group of about three, and every additional member reduces everyone's share.** Rival stored order pushes groups toward the smallest size that clears the outside option.

Non-rival case (a standard, a piece of infrastructure, a body of practice): each member's return is `eta ln N − c_i`, increasing without bound in `N`:

`N`	3	5	8	12	20	30
return	+0.3645	+0.5433	+0.7078	+0.8497	+1.0285	+1.1704

The same group, the same surplus, the same members — and the sign of the size gradient flips entirely depending on whether the stored order is divided or shared. This is a fork in the external model with no internal counterpart whatever, because a mind's stored order is not shared with anyone and the question does not arise. It is the same fork that separates common-pool resources from public goods in the governance literature [Ostrom, 1990, 2010], and it is why the design principles that work for one do not transfer unaltered to the other.

Transfer restriction: **any external application must declare whether `Q` is rival or non-rival before predicting anything about group size.** Predictions about optimal `N` transferred from internal reasoning are unsupported.

25.8.9 8.9 B9 — The ethical asymmetry. ABSOLUTE

The mathematics is symmetric. The permissions are not.

A person may restructure their own `N`, force their own `u`, raise their own `delta` on a habit they want gone, and spend their own `h` however they like. **They may not do any of those things to another person, and the framework's governing refusal is exactly this: no intervention may remove a person's ability to decline.**

Every internal result in this corpus is written for a practitioner working on themselves, and every external result for an engineer working on a circumstance that other people are free to enter and leave. Those are different licences, and the correspondence does not transfer the licence with the equation. **A technique that is legitimate self-discipline internally becomes coercion externally by the simple substitution of a second person, and the substitution is exactly what the correspondence otherwise licenses.**

Section 5.5 supplied the measurable form of the boundary: **an obligation network that removes exit is not a trust system but a capture system, and the two are distinguished by `rho`, which the first has and the second does not.** The ethical rule and the engineering rule agree here, which is fortunate but not necessary, and where they disagree the ethical rule wins without argument. Curzi [2026] states the corpus's boundary in the form this paper is bounded by, including the concession that the case where the feedback loop runs through another person's will remains an unsettled seam.

No entry in this corpus's bibliography establishes consent as a precondition for measuring or restructuring another person's relationships. The constraint is a normative stipulation of the framework, badged as such here and again in Section 12. What the literature does supply is the narrower and purely methodological result that covert study changes what is measured [McCambridge et al., 2014, Rosenthal and Jacobson, 1968], which makes concealed measurement unsound as well as prohibited but does not by itself carry the ethical claim.

25.8.10 8.10 What survives all nine

Something does, and it is worth stating precisely rather than defensively.

The recursion transfers. `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t` holds at both scopes, and every theorem about it holds at both scopes, because none of them mentions a substrate. Section 7.5 demonstrates the shared trajectory to within two per cent.

The geometry transfers. `N_eff = 1 + (N−1)u` and `S = ln N_eff` are arithmetic on a redundancy audit, and redundancy is scope-independent.

The ceiling transfers, and produces a new result in doing so. `F <= ln N` at both scopes forces `M <= ln N` at internal scope, which is requisite variety in first-person form. That result was not available at either scope alone.

The diagnostic questions transfer. *Is the maintenance term eating the ceiling? Is the frame outrunning the intake? Is the interaction leaving anything behind?* These are the same three questions with different nouns.

What does not transfer is anything involving agency of the parts. Exit, heterogeneity of interests, aggregation of preferences, reciprocity, sanction, rivalry of the store. **The correspondence is a correspondence of dynamics and geometry, not of agency.** Stated that way it is defensible. Stated as "a company is a mind" it is a slogan that will be wrong in precisely the six places listed above, all of which are the places where somebody can walk out.

25.8.11 8.11 What would falsify the correspondence

A framework that cannot fail cannot be trusted, and this is how this one fails.

Test 1 — cross-scope parameter recovery. Take one operationalisation of `delta` — the two-time zero-input estimator — and apply it to twenty individuals learning a skill and twenty teams maintaining a practice. Preregister. **The correspondence predicts that after normalisation by period length the two `delta` distributions occupy the same functional role**, meaning that `Q_inf = (eta S − C)/delta` predicts held-out `Q` equally well at both scopes with no scope-specific correction term. *If a scope-specific correction term is needed and does not shrink with sample size, the correspondence is false as an equality of role and must be demoted to an analogy of form.*

Test 2 — the shape test. Fit the trajectory shape, not the levels. **The correspondence predicts that normalised trajectories at the two scopes are the same curve** (Section 7.5). *If internal trajectories are systematically non-exponential in normalised time while external ones are exponential, or the reverse, the shared-dynamics claim fails and the material reduces to a notational convenience.*

Test 3 — the four baselines, both scopes. Persistence, direct input, matched autoregression, and a shuffled-input control [Curzi, 2026i]. *If shuffled `U` predicts `Q` as well as real `U` at either scope, the causal story at that scope is a trend.* **The control is weak and must be calibrated before use**: its measured false-rejection rate is 0.108 for iid permutation and 0.120 for circular shift against a nominal 0.05, with only block permutation near nominal at 0.060, so it is **anti-conservative by a factor of about two**; it passed a filtered repackaging at `p = 0.0005` that was caught only by the structural ceiling test, at 38.5 times the admissible one-step maximum; and against a confound driving both series it passed 98.35 per cent of the time while detecting the clean case only 48.17 per cent of the time [Curzi, 2026i]. **Design controls, not more surrogates.**

Test 4 — the breakdown predictions. Section 8 is not only a list of caveats; it is a list of predictions. **B2 predicts that a heterogeneous group's fitted `delta` drifts downward over an unmaintained interval while a homogeneous one's does not.** That is directly testable and is the sharpest available discriminator between the mean-field model and the real one. *If fitted `delta` does not drift in heterogeneous groups, the mean-field caution is wrong; and if it drifts in individuals too, then B2 is not a disanalogy at all and the correspondence is stronger than this paper claims.*

Test 5 — reproducibility across investigators. Same operationalisation, different investigators, different populations. **This gate decides the hard-science question and it is the one most often skipped. Nothing in this material has passed it** [Open Science Collaboration, 2015].

25.8.12 8.12 Scope specifications and failure modes

S13 — Scope declared. Every application shall declare its scope (internal or external), its period length, and its unit for `Q`, before any parameter is estimated.

S14 — Period stated with every rate. `delta`, `eta`, `C` and `h` shall never be reported without their period. Cross-scope comparison shall use `delta_T = 1 - (1-delta)^T` (B7).

S15 — Distribution, not scalar, at external scope. `delta`, `eta`, `C` and `u` shall be reported with their across-member distribution at external scope. **A scalar alone is a mean-field claim and shall be labelled as one** (B2).

S16 — Member headroom computed. Any external application reporting `h > 0` shall also report the distribution of `h_i` and name the members below zero (B1). **A body may not be described as sustainable while a quarter of its members are not.**

S17 — Rivalry declared. Before any prediction about group size, the analyst shall state whether `Q` is rival or non-rival and justify it (B8).

S18 — Cyclicity tested before `D/P`. Group-scope `D/P` shall not be computed until the collective ordering has been tested for intransitivity over the relevant option set (B3).

S19 — No inward transfer of agency machinery. `rho`, obligation ledgers, sanctions and verification budgets shall not be given internal operationalisations (B4). Their internal use is metaphorical and shall be labelled.

S20 — Per-scope validation. An instrument validated at one scope shall be re-validated at the other before use (B5).

S21 — Licence not transferred. Techniques licensed for self-application shall not be applied to another person on the strength of the correspondence (B9).

S22 — The thesis carries its badge. In any document quoting this correspondence it shall be identified as the author's thesis, with Section 8 cited, and not as a proven result. **The proven result is the scale-freeness of the form, and it is a much smaller claim.**

F11 — The slogan. *Signature:* "a company is just a big brain," offered as an explanation rather than as a model with limits. *Damages:* everything downstream, because it licenses transfers through all nine breakdowns at once. *Measurement:* ask which of the nine breakdowns the speaker has checked; the usual answer is none. *Fix:* S22.

F12 — Unnormalised period comparison. *Signature:* a table comparing an individual's `delta` to a department's with no period column. *Damages:* conclusions wrong by up to two orders of magnitude (B7: 90.1x on a realistic pair). *Measurement:* look for the period column; if absent, the comparison is void. *Fix:* S14.

F13 — Mean-field complacency. *Signature:* one `delta` for a group with obviously heterogeneous members; a decay forecast that keeps being wrong in the same direction. *Damages:* forecast error rising to 92 per cent over twelve periods (B2). *Measurement:* fit `delta` on early data and on late data separately; divergence is the diagnosis. *Fix:* S15, and compartment the model.

F14 — Aggregate `u` complacency. *Signature:* a diverse-on-paper body reporting a healthy `u` with no distribution. *Damages:* invisible fragility — configuration B loses 77.4 per cent of its surplus to the departure of three of nine against configuration A's 21.1 per cent, and reports identically beforehand. *Measurement:* compute `u` with each member removed in turn; the spread of that vector is the fragility. *Fix:* S15. **This computation costs nine subtractions and should be standing practice.**

F15 — The group mind that cannot decide. *Signature:* a body that reverses itself in a cycle, each reversal individually well argued, members each believing the others inconsistent. *Damages:* `D` is undefined and every doctrine-to-data diagnostic returns nonsense. *Measurement:* elicit member orderings and check pairwise majorities for a cycle (B3). *Fix:* S18 — restrict the option set or change the aggregation rule, noting that one of the four conditions must be given up whichever is chosen.

F16 — The sustainable body with departing members. *Signature:* positive aggregate headroom, rising attrition, and genuine bewilderment at the attrition. *Damages:* loses the highest-

`c_i` members first, who are frequently the highest-`u` members (B1). *Measurement:* the `h_i` table; it takes an afternoon. *Fix:* S16 — reduce `c_i` for the members below zero, or accept the loss with its `ln N` price computed (F9).

F17 — Inward transfer of the sanction ladder. *Signature:* a self-improvement scheme built on internal contracts, penalties and rewards, with an internal agent doing the enforcing. *Damages:* imports a model whose central parameter has no internal referent (B4). *Measurement:* ask what the flow matrix is between which agents. There is no answer. *Fix:* S19.

F18 — Introspective validation. *Signature:* an internal instrument validated by asking the subject whether it seems right. *Damages:* the error in introspection is correlated with the quantity being measured, so agreement is evidence of nothing (B5). *Measurement:* check whether the validation criterion is independent of the instrument; usually it is not. *Fix:* S20, with behavioural criteria [Nisbett and Wilson, 1977].

25.9 9 Results: three systems worked end to end

All three systems are constructed, not observed. Every number below is the output of the stated parameters under the model developed above, carried through so that a reader can re-derive it. They demonstrate that the procedure runs and that it reports its own failures; they are not evidence that any organisation behaves this way and must not be cited as such.

25.9.1 9.1 An eight-firm regional supplier consortium

Eight independent manufacturing firms in one region, formed to share tooling, qualified staff, emergency capacity and test equipment; nine quarters old; membership voluntary with one quarter's notice. **Two members, firms 1 and 2, are substantially larger than the rest.** The presenting complaint: participation is good but the consortium never gets anywhere — every quarter it agrees things and every quarter it starts over.

Encoding. Period one quarter. `N = 8` firms; measured non-parallelism across their capability and market framings `u = 0.52`, so `S = ln(1 + 7(0.52)) = 1.5347` against a ceiling `ln 8 = 2.0794`. **The consortium is operating at 73.8 per cent of its surplus ceiling, which is respectable and is not the fault.** `eta_max = 0.45`, estimated from the three shared assets that became permanent in the first two quarters. `C_0 = 0.30` per quarter (secretariat, meeting time, legal). `delta_0 = 0.05` per quarter, mostly staff movement inside member firms.

Baseline measurement, Protocols B and C. One quarter's flow matrix in documented favour-units (loaned tooling days, seconded staff days, test slots, referrals), all symmetric and ratio-scaled. Rows are givers:

to:	F1	F2	F3	F4	F5	F6	F7	F8
F1	–	1	1	0	0	0	0	0
F2	1	–	1	1	0	0	0	0
F3	6	7	–	3	2	1	1	0
F4	4	5	2	–	2	1	0	0
F5	3	4	1	1	–	1	0	0
F6	2	3	1	1	1	–	1	0
F7	2	2	0	0	0	1	–	1
F8	1	2	0	0	0	0	1	–

Directed volume 68 favour-units per quarter; matched volume 26; `rho_group = 26/68 = 0.3824`. Node balances `b_i` (given minus received):

firm	F1	F2	F3	F4	F5	F6	F7	F8
b_i	**-17**	**-21**	+14	+8	+5	+5	+3	+3

B = 76/(2 x 68) = 0.5588, so 1 − B = 0.4412. As S6 requires, note that 1 − B = 0.4412 exceeds rho = 0.3824: **the node index flatters the network by 15.4 per cent.** The sign structure is unambiguous — two large net receivers, six net givers, and the worst dyads are one-directional:

dyad	x	y	volume	rho
(1,4)	0	4	4	0.0000
(1,5)	0	3	3	0.0000
(2,5)	0	4	4	0.0000
(2,6)	0	3	3	0.0000
(2,3)	1	7	8	0.2500
(1,3)	1	6	7	0.2857

This is the vertical patron-client pattern of Section 5.5, measured. It is not what anyone intended; firms 1 and 2 are the largest and were the natural first call, and nine quarters of natural first calls produced a star.

Protocol A, reliance rate. Domain: commitments to supply capacity when asked. Sixty-two low-consequence commitments over four quarters, of which 21 were left unverified by the sampling rule; honoured unverified 13. a_hat = 13/21 = 0.619, taken as **a = 0.62**. With 21 trials the 95 per cent Wilson interval is approximately [0.41, 0.79], which is wider than several of the differences the diagnosis turns on and is the honest report.

Diagnosis.

```
eta   = eta_max * rho   = 0.45 * 0.3824 = 0.1721
delta = delta_0 + (1−a) = 0.05 + 0.38   = 0.4300
S     = 1.5347          eta S = 0.2641          C = 0.30
Q_inf = (0.2641 − 0.30) / 0.4300 = −0.0836
```

The steady state is negative, which is the theorem stating its own failure mode: the regime is unsustainable in this model, or the linear model must be replaced. The complaint — "we agree things and every quarter we start over" — is Q_inf < 0 in plain English. Headroom at the current Q = 1.20:

```
h = eta ln N − delta Q − C = 0.3578 − 0.5160 − 0.3000 = −0.4582
```

Negative, and **no arrangement of the interaction geometry rescues it**: the consortium could raise u to 1.0 and still fail, because the ceiling eta ln N = 0.3578 is below delta Q + C = 0.8160 on its own. Curzi [2026] reaches the same shape in a different domain and with a cleaner cancellation: at a 15 per cent annual hard-currency drag, a local-currency allocation of 25.17 per cent of protected capital drives the headroom coefficient to exactly −0.500000 and cancels the whole eta ln(N) = 0.377583 diversification gain, with a 10 per cent allocation costing 0.150000, or 39.73 per cent of the gain. **A cost term consuming a logarithmic gain is the same arithmetic here as there; only the nouns differ.** Trajectory from Q_0 = 1.20 with no intervention:

quarter	0	1	2	3	4	5	6
Q	1.2000	0.6481	0.3335	0.1542	0.0520	-0.0063	-0.0395

Admissibility requires `Q >= (0.30 − 0.2641)/0.57 = 0.0630`. **The trajectory breaches that between quarter 3 and quarter 4.** What the model predicts past that point is not negative stored order — that is meaningless — but the exit of the model's validity, which in this system means dissolution. On this measurement the consortium had about four quarters left.

What is required. By S7 the numerator comes first. Required `eta` for a positive numerator at current `S` and `C`: `eta > 0.30/1.5347 = 0.1955`, so `rho > 0.4344`. **The consortium is 5.2 points of reciprocity away from not dying, and 51 points away from working.**

The intervention. Three changes, each attached to a parameter. **I-1 (raises `rho`): the reciprocal-call rule** — firms 1 and 2 commit that for each favour received they originate one outbound offer to the same partner within the quarter, from a published list. This is the congruence principle implemented literally as a balance condition on the flow matrix. **I-2 (raises `rho` and volume): the capability register** — every firm publishes what it can lend; the small firms did not know what the large ones had. **I-3 (raises `a`): the standing-commitment ledger with quarterly sunset** — every commitment recorded with a date and a twelve-month review, which is deliberately a `delta`-raising device on obsolete items and an `a`-raising device on live ones, because most of the unhonoured unverified commitments in the baseline turned out to be commitments neither party could state precisely. **Deliberately not done:** expulsion of the two large firms, or any sanction; per Section 6.3 that would cost 7.7 per cent of `eta S` at minimum, and the ladder should exist before it is needed.

Post-intervention measurement, fourth quarter after.

```
to:      F1  F2  F3  F4  F5  F6  F7  F8
F1        –   3   4   3   2   1   1   1
F2        3   –   5   4   3   2   2   1
F3        6   7   –   3   2   1   1   1
F4        4   5   3   –   2   1   1   1
F5        3   4   2   2   –   1   1   1
F6        2   3   1   1   1   –   1   1
F7        2   2   1   1   1   1   –   1
F8        1   2   1   1   1   1   1   –
```

Directed volume 114 (from 68); matched volume 102; `rho_group = 102/114 = 0.8947`. Balances are now `[−6, −6, +4, +2, +2, +2, +1, +1]`, so `B = 24/228 = 0.1053` and `1 − B = 0.8947`. **The bound `1 − B >= rho` is met with equality here**, which happens exactly when no member is simultaneously a creditor to one partner and a debtor to another — true of this matrix and false of the baseline one. Reliance rate re-measured after three quarters of the ledger: `a = 0.88`.

```
eta   = 0.45 * 0.8947 = 0.4026
delta = 0.05 + 0.12   = 0.1700
eta S = 0.4026 * 1.5347 = 0.6179     C = 0.30
Q_inf = (0.6179 − 0.30)/0.1700 = +1.8701
h at Q = 1.20 = 0.8372 − 0.2040 − 0.3000 = +0.3332
half-life = 3.72 quarters (was 1.23);  tau = 5.88 quarters (was 2.33)
```

quarter	0	2	4	6	8	10	12
before-regime `Q`	1.2000	0.3335	0.0520	-0.0395	-0.0693	-0.0789	-0.0821
after-regime `Q`	1.2000	1.4085	1.5521	1.6510	1.7192	1.7662	1.7985

Attribution. The one-at-a-time decomposition is the table of Section 6.5. **Had the consortium done I-3 alone — the ledger, the discipline, the reliability programme, which is the intervention every one of its members would have nominated — it would have measured better and performed worse.**

Falsification, declared in advance. (1) If `rho_group` does not exceed 0.4344 within three quarters, the intervention has failed to clear even the survival threshold and must be reported as unsuccessful. (2) If `rho` rises but the number of shared assets surviving a full year does not, then Postulate PR is wrong in this system and must be withdrawn, not defended. (3) If `a` rises while `rho` stays flat and `Q` nonetheless grows, Postulate PT is doing work the model says it cannot. (4) **Shuffled-input control, mandatory**, with the calibration caveats of Section 8.11. (5) If the two large firms' balances return to the baseline pattern within a year of the rule being relaxed, the improvement was in the rule and not in the relation, and the correct report is that `rho` was rented rather than built.

The rule, extracted. The consortium was not short of difference: at `u = 0.52` it was collecting 73.8 per cent of its surplus ceiling. It was not short of activity: sixty-eight favours a quarter among eight firms is a busy network. **It was failing because two-thirds of that activity was one-directional, which put `eta` at 0.17, which put the numerator below the cost line, which made the steady state negative while every visible indicator — attendance, volume, goodwill — looked fine.**

25.9.2 9.2 One clinician and her unit, measured with one instrument

A hospital department. The engineer is asked to look at a senior clinician who reports being "completely full" and, separately, at the six-person unit she leads, described by its director as "busy but not going anywhere." **The point of the exercise is that these are two readings from the same instrument, and the point of Section 8 is the four places where they must be read differently.**

Scope declarations (S13, S14).

	INTERNAL	EXTERNAL
subject	one clinician	the six-person unit
period	1 week	1 quarter
`Q` unit	retained-capability units	retained-practice units
rivalry (S17)	not applicable	**non-rival** — the store is shared protocol

Baseline, internal. `N = 9` distinguishable clinical framings from a free-sort with test-retest agreement 0.71; `u = 0.31` from an independent-prediction round on eight recent cases.

```
S     = ln(1 + 8 x 0.31) = 1.2470       ln N = 2.1972,  so 56.8% of ceiling
eta   = 0.22    C = 0.42    delta = 0.035
eta S = 0.2743     Q_inf = (0.2743 - 0.42)/0.035 = -4.1615
h at Q = 6.0: 0.4834 - 0.2100 - 0.4200 = -0.1466
```

Negative steady state and negative headroom. "Completely full" is `h < 0`, measured.

Baseline, external. `N = 6` distinct domains from the capability inventory with independent-action verification; `u = 0.29` from the pre-discussion prediction round.

```
S     = ln(1 + 5 x 0.29) = 0.8961       ln N = 1.7918,  so 50.0% of ceiling
eta   = 0.24    C = 0.44
delta_int = 0.05,  turnover tau = 0.10  ->  delta_total = 1 - (0.95)(0.90) = 0.1450
eta S = 0.2151     Q_inf = (0.2151 - 0.44)/0.1450 = -1.5513
h at Q = 2.8: 0.4300 - 0.4060 - 0.4400 = -0.4160
```

The same three faults, at both scopes, in the same order: `u` at roughly half the ceiling, `eta` low, `C` high enough to exceed the ceiling by itself. The clinician and her unit are not two problems; they are one diagnostic profile read twice.

The four places the readings must diverge. *B7, period.* `delta = 0.035` weekly and `delta = 0.145` quarterly are not comparable as printed. The clinician's weekly rate over a 13-week quarter is `1 − (0.965)^13 = 0.3707`. **Her personal store decays 2.56 times faster per quarter than her unit's does**, which is the opposite of the impression the raw numbers give. *B1, exit and turnover.* The unit's `delta` contains `tau = 0.10`, a term the clinician's does not have and cannot have; `0.10/0.1450 = 69 per cent` of the unit's total loss is people leaving. **No amount of documentation addresses that term.** *B2, heterogeneity.* Per S15 the unit's `u` was recomputed with each member removed: removing member 4 dropped `u` from 0.29 to 0.19, and removing any other dropped it by at most 0.02. **One of six people is carrying a third of the unit's remaining non-parallelism, and member 4's `h_i` was the second lowest in the unit.** The scalar `u = 0.29` concealed both facts. *B9, licence.* The clinician's own `delta`-raising interventions she may apply to herself; the unit's equivalent requires the members' agreement, and the intervention list is written so that every member can decline any part of it.

Intervention and post-measurement. Internal (weeks 1–20): a second genuinely different framing developed and used on every case (`u` to 0.62); a weekly written case note with a decision line (`eta` to 0.34); three standing commitments sunset (`C` to 0.30).

```
S     = ln(1 + 8 x 0.62) = 1.7851       81.2% of ceiling
eta S = 0.6069      Q_inf = (0.6069 − 0.30)/0.035 = +8.7693
h at Q = 6.0 = 0.7470 − 0.2100 − 0.3000 = +0.2371
tau = 28.6 weeks;  half-life = 19.46 weeks
```

External (quarters 1–6): a rotating pre-discussion prediction round (`u` to 0.60); a decision log with named owners (`eta` to 0.37); two recurring meetings deleted (`C` to 0.31); and — the term with no internal counterpart — a retention intervention on the two members whose `h_i` was negative, taking `tau` from 0.10 to 0.04, offered and declinable.

```
delta_total = 1 − (0.95)(0.96) = 0.0880
S     = ln(1 + 5 x 0.60) = 1.3863       77.4% of ceiling
eta S = 0.5129      Q_inf = (0.5129 − 0.31)/0.0880 = +2.3060
h at Q = 2.8 = 0.6630 − 0.2464 − 0.3100 = +0.1066
tau = 11.36 quarters;  half-life = 7.52 quarters
```

The shape test. Both trajectories normalised by their own time constants, from their own starting points:

t / tau	clinician (Q−Q_0)/(Q_inf−Q_0)	unit (Q−Q_0)/(Q_inf−Q_0)	1 − e^−k
0.5	0.3927	0.4246	0.3935
1.0	0.6441	0.6370	0.6321
2.0	0.8688	0.8798	0.8647
3.0	0.9533	0.9564	0.9502

Same curve. The clinician's time constant is 28.6 weeks — 2.20 quarters — and the unit's is 11.36 quarters. **The unit is 5.2 times slower and arrives in the same shape**, which is Banach at both scopes and is the most useful practical output of the correspondence: one can tell a person how long their change will take by measuring their `delta`, and tell a department the same way, with identical arithmetic.

`D/P` at both scopes, magnitudes demoted per Section 3.4 and reported as an ordering:

subject	D	P	r	Delta	band
clinician, before	2.40	1.00	2.400	0.8755	fault
clinician, after	1.15	1.00	1.150	0.1398	drift
unit, before	2.90	1.00	2.900	1.0647	**severe**
unit, after	1.25	1.00	1.250	0.2231	drift

Both were over-framed, the unit worse than the clinician. Per S18 the unit's option orderings were tested for cyclicity first and were transitive, so `D` was well defined and the diagnostic applied. **Had the test failed, the unit's `r` would have been reported as undefined rather than as 2.90.**

Falsification, declared in advance. (1) If normalised trajectories at the two scopes diverge by more than five per cent at `t/tau = 1`, Test 2 has failed here. (2) If the unit's fitted `delta` drifts downward over the unmaintained interval while the clinician's does not, B2 is confirmed as a genuine scope disanalogy; if neither drifts, the mean-field caution is overstated in this system and should be reported as such. (3) Shuffled-input control at both scopes, mandatory. (4) If member 4 leaves and the unit's measured `u` does not fall by approximately the predicted 0.10, the fragility computation is wrong here.

The rule, extracted. One instrument, two scopes, one diagnostic profile, and four places where the readings had to be handled differently: the period normalisation, the turnover term, the member-level distribution, and the licence. **The correspondence earned its keep by letting one engineer diagnose a person and a department in the same afternoon with the same arithmetic. It would have destroyed the diagnosis by letting the same engineer treat the department's turnover as forgetting, its dissenting member as a framing, and its members' consent as already given.**

25.9.3 9.3 A regional volunteer rescue association

A volunteer search-and-rescue association: fourteen members on the roster, across a county, nine years old, self-governing, donation-funded, membership voluntary in the fullest sense. The complaint: "we are losing our best people and we cannot work out why, because morale is good." Period one quarter; `Q` unit retained-practice units; **rivalry declared (S17): non-rival**, since the store is trained capability and shared protocol.

Effective count. Capability inventory: `N = 9` domains verified by observed independent action (rope, water, medical, canine, navigation, communications, incident command, vehicle, public liaison); five of the fourteen members duplicated domains already covered. Pre-discussion prediction round on eight past incidents gave `u = 0.42`.

```
N_eff = 1 + 8 x 0.42 = 4.360      S = ln 4.360 = 1.4725
ln N  = ln 9 = 2.1972             S / ln N = 67.0%
```

Fourteen people, nine domains, an effective count of 4.36 — **on the framework's own affine effective count, not a Hill number.**

Turnover composition. Four departures over four quarters against an average roster of 14 gives `tau = 0.0714` per quarter; intrinsic decay over a two-quarter stable interval `delta_int = 0.06`, so `delta_total = 1 − (0.94)(0.9286) = 0.1271`. Turnover is `0.0714/0.1271 = 56 per cent` of total loss. **More than half the association's forgetting is people leaving.**

Storage law and steady state. `eta = 0.31` over eleven quarters; `C = 0.42` as coordination plus compliance plus travel.

```
eta S = 0.31 x 1.4725 = 0.4565
Q_inf = (0.4565 − 0.42)/0.1271 = +0.2868
```

Positive, but barely. Current `Q = 1.85`, well above equilibrium: **the association is declining toward 0.2868 and will take `ln 2 / -ln(1-0.1271) = 5.10` quarters to cover half the distance.**

Headroom. `eta ln N = 0.6811`; `h` at `Q = 1.85` is `0.6811 - 0.2352 - 0.4200 = +0.0259`. **Effectively maintenance only — hold, do not build** — and the association had spent the year planning an expansion. Ossification point `Q = (0.6811 - 0.42)/0.1271 = 2.0539`, so **the association's `Q` of 1.85 is 90 per cent of the way to its own ceiling.**

The requested increment. The expansion plan requested `r = 0.15` `Q`-units per quarter.

```
required S = (0.1271 x 1.85 + 0.42 + 0.15)/0.31 = 0.8051/0.31 = 2.5975
ln N = 2.1972  ->  required S exceeds the ceiling.  INFEASIBLE.
required u = (e^2.5975 - 1)/8 = 1.5537 > 1.  Confirmed.
r_max = h = 0.0259
```

The plan asked for 0.15 and the maximum available is 0.0259 — a factor of 5.79. Not closeable by effort; the only remaining moves are to raise `N` or lower `C`.

Member headroom. Ceiling share `0.6811/14 = 0.04865`; maintenance share `0.2352/14 = 0.01680`; critical private cost `c_i* = 0.03185`. Measured `c_i` (travel, training, callout burden, in `Q`-units):

member group	count	`c_i`	`h_i`	status
town-based, single domain	6	0.014	+0.01785	ok
town-based, multi-domain	3	0.028	+0.00385	marginal
rural, multi-domain	3	0.052	**-0.02015**	**unsustainable**
rural, incident command	2	0.071	**-0.03915**	**unsustainable**

Five of fourteen are individually below water while aggregate `h` is positive, and the four who left in the last year were all from the two rural groups. **The association was not losing people to a morale problem. It was losing them to a travel and callout burden that made their private headroom negative.** Morale read as good because the town-based majority — nine of fourteen — was above water.

Departure fragility. Per S15, `u` recomputed with each member deleted:

removed	resulting `u`	`S` at `N` unchanged	loss
any town-based single-domain member	0.42	1.4725	0.0000
a town-based multi-domain member	0.40	1.4351	0.0374
a rural multi-domain member	0.35	1.3350	**0.1375**
a rural incident-command member	0.31	1.2470	**0.2254**

The members with the most negative `h_i` are precisely the members whose departure costs the most surplus. Random departure costs on average 0.0697; targeted departure of an incident-command member costs 0.2254, **3.24 times as much**. The association had been experiencing targeted departure with nobody targeting anything — the fragility profile of Albert et al. [2000] arising from a private-cost gradient rather than from an attacker.

Reciprocity and obligation. Flow matrix over four quarters, unit = callout-hours covered for another member, symmetric and ratio-scaled. Per S4 the window was preregistered at `w = 1/delta = 7.9` quarters and truncated to the four quarters of record; **the truncation is reported rather than absorbed.**

```
total directed volume = 214 hours     matched volume = 116 hours
rho_group = 116/214 = 0.5421
B = 118/(2 x 214) = 0.2757            1 - B = 0.7243
```

Node balances: the two rural incident-command members are net creditors at +31 and +24 hours; the six town-based single-domain members are net debtors between −4 and −11. Per the bound, `1 - B = 0.7243 >= rho = 0.5421`: **the node index overstates by 34 per cent**, and the sign structure is the patron-client pattern — a few large net creditors carrying many small net debtors. **The people leaving were the creditors.** Back-solving the conversion law from the measured `eta = 0.31` gives `eta_max = 0.5719`; **at `rho = 0.85`, `eta` would be 0.4861 — 57 per cent more conversion with no change in headcount, domains, or `u`.**

Doctrine-to-data. Prediction-surprise ledger over eleven incidents: confidence high and uncorrelated with accuracy, the over-framing signature. Cyclicity tested first per S18; the ordering over deployment options was transitive, so `D` was well defined. `D = 1.85`, `P = 1.00`, `r = 1.850`, `Delta = 0.6152`, **FAULT band — reported as an ordinal position, not as a magnitude.**

The intervention, priced.

change	parameter	effect
Callout rota rebalanced so town members cover rural incidents in proportion	`rho` 0.5421 -> 0.82	`eta` 0.31 -> 0.4690
Travel and equipment reimbursement for rural members	`c_i` -0.021 for the five below water	all `h_i` positive
Two recurring meetings deleted; incident debrief kept	`C` 0.42 -> 0.34	
Two duplicated town members cross-trained into uncovered domains	`N` 9 -> 11	`ln N` 2.1972 -> 2.3979
Expansion plan withdrawn	—	target was infeasible by 5.79x

Recomputed, with `u` held at 0.42 — **deliberately not assumed to improve**:

```
S      = ln(1 + 10 x 0.42) = 1.6487       S/ln N = 68.8%
eta S  = 0.4690 x 1.6487 = 0.7732
delta  = 1 - (0.94)(0.98) = 0.0788        (tau assumed to fall to 0.02)
Q_inf  = (0.7732 - 0.34)/0.0788 = +5.4977
h at Q = 1.85 = 1.1246 - 0.1458 - 0.3400 = +0.6388
r_max  = 0.6388,  against the withdrawn request of 0.15
```

Steady state from +0.2868 to +5.4977; headroom from +0.0259 to +0.6388 — the withdrawn expansion is now feasible four times over. And it became feasible by paying five people's travel costs and deleting two meetings.

Falsification, declared in advance. (1) If `rho_group` does not exceed 0.70 within three quarters, the rota change failed. (2) If `tau` does not fall below 0.04 within four quarters, the `h_i` diagnosis is wrong here and the departures have another cause. (3) If `Q` does not rise above 2.60 within `1/delta = 12.7` quarters, the fitted parameters do not describe this system. (4) Shuffled-input control, mandatory. (5) If a rural incident-command member departs and `u` does not fall by approximately 0.11, the fragility computation is wrong here.

The rule, extracted. The association believed it had a morale problem in a healthy body. It had **a private-cost gradient in a body at 90 per cent of its ossification point**, selecting for the departure of exactly the members whose departure cost the most surplus, in a network whose reciprocity index was 0.5421 while its node-level index read 0.7243. **None of those facts**

is visible without the arithmetic, and every one is computable from records the association already kept.

25.9.4 9.4 Four defects found in the source material

Reported here together, because a reader checking the arithmetic will find them and should find them disclosed first. All four are stated where they arise as well.

#	Location	Defect	Status
1	Section 4.4	The elasticity of `Q_inf` in `a` is printed as `a/(1−a)`, which holds only at `delta_0 = 0`, immediately below a table computed at `delta_0 = 0.04`. The correct point elasticity is `a/(delta_0 + 1 − a)`: **10.56 at `a = 0.95`, not 19.00**, and **19.80 at `a = 0.99`, not 99.00.**	**Corrected in this paper**, with both columns printed.
2	Section 8.2	The heterogeneity damage ratio is printed as “3.49 times” / “three and a half times”. That is the ratio of the two *surviving* surpluses, `1.3862/0.3971`. The ratio of the two *losses* is `1.3607/0.3716 = 3.66`, which is what the printed percentages 77.4 and 21.1 imply. The companion ratio in the same entry (2.27, targeted against random single removal) is computed on losses. **Two ratios in one entry, computed two ways.**	**Corrected to 3.66**, both methods named.

#	Location	Defect	Status
3	Section 7.5	“The clocks differ by a factor of about a hundred and forty” is not recoverable from the stated parameters. The time constants are 50 days and 6.25 quarters: a ratio of **8** in period units and **11.4** in absolute time at 91 days to the quarter.	**Withdrawn**; corrected figures printed.
4	Section 8.1	The per-member headroom shares are computed at `N = 9` while only six members are tabulated, and the result is reported as “four of six”. The denominators differ. The row arithmetic is correct; the reading is not. Additionally `u = 0.6` is stated for a computation that does not use `u`.	**Disclosed**; the claim is narrowed to “four of the six tabulated members”.

A fifth item is not a defect but an unverifiable figure. The post-sanction `rho = 0.9038` of Section 6.3 is quoted without the modified flow matrix that would produce it, so it cannot be recomputed from anything printed. The sign of the change is defensible; the magnitude is not, and it is not relied on anywhere else in this paper.

25.9.5 9.5 The results in quotable form

R1 — Trust enters the storage law as retention. `delta_eff = delta_0 + (1 − a)(1 − v)`, giving a superlinear stock dividend from `Q_inf = 0.4679` at `a = 0.50` to `Q_inf = 5.0529` at `a = 0.99`, at `K = 0.2526` and `delta_0 = 0.04`. **The point elasticity is `a/(delta_0 + 1 − a)` — 0.93 at `a = 0.50` and 10.56 at `a = 0.95` — and not `a/(1−a)`, which holds only at `delta_0 = 0`.**

R2 — Verification is all-or-nothing under a linear verification cost, with threshold `a* = 1 − c_v delta_0 / (K − c_v)`. At `c_v = 0.20` and `K = 0.2526`, verify below `a = 0.8479`; at `c_v >= K` verification consumes the entire net input and is never worth buying. Under a convex cost the optimum is interior and only the threshold structure survives.

R3 — Reciprocity raises conversion, `eta = eta_max rho` with `rho = 2 min(x_ij,x_ji)/(x_ij+x_ji) = 1 − |Skew_ij|/Sym_ij`, and the node imbalance index obeys `1 − B >= rho`, flattering the worked network by 15.4 per cent (0.4412 against 0.3824) and a second worked network by 34 per cent (0.7243 against 0.5421).

R4 — A consortium at `u = 0.52`, collecting 73.8 per cent of its surplus ceiling, still had steady state `−0.0836` and headroom `−0.4582`. The binding fault was `rho = 0.3824`, not difference, and the survival threshold was `rho > 0.4344` — 5.2 points away.

R5 — Graduated sanction against expulsion, computed. Removing one of eight members costs 7.7 per cent of `eta S` at `u` held fixed, and more where the departing member carries

divergent framing; a 30 per cent flow reduction that keeps the edge gains 1.0 per cent.

R6 — The mean-field scalar `delta` is exactly right for one period and wrong by 92 per cent after twelve when the store is split evenly between compartments at `delta = 0.05` and `delta = 0.50`, with the implied aggregate rate drifting from 0.2750 to 0.0504.

CR (counter-result) — The correspondence audits its own central analogy and reports nine breakdowns — exit, heterogeneity, and ill-defined aggregation rated severe, and an absolute ethical asymmetry — concluding that **the equation transfers between scopes but the licence to apply it does not.**

25.10 10 Discussion

25.10.1 10.1 What follows

The three quantities are not interchangeable and a diagnosis must say which one it means. That is the practical yield. A body reporting a trust problem, a reciprocity problem and an obligation problem is reporting three different things with three different repairs and three different prices, and the parameterisation separates them in an afternoon of arithmetic on records most bodies already keep. The single most consequential separation is the sign test of S7: **compute `eta S - C` before commissioning any `delta` intervention.** In a system with a negative numerator the intervention is predicted to make things worse, and the prediction is cheap to check.

The node-level index should be retired as a reciprocity estimate. `1 - B` is computable from aggregate give-and-take and is what most reporting actually produces; the bound `1 - B >= rho` says it can only ever overstate. The two worked networks overstate by 15.4 and 34 per cent respectively, and in both cases the overstatement pointed the diagnosis away from the fault.

The correspondence is worth keeping, in its restricted form. Stated as a correspondence of dynamics and geometry it earns its keep: one instrument, two scopes, one arithmetic, and a shared trajectory shape verified to within two per cent on a constructed pair. Stated as an identity of kind it is wrong in nine specific places, and the nine are enumerated rather than gestured at.

25.10.2 10.2 What does not follow

Nothing here shows that any real body behaves this way. Three systems are worked and all three are constructed. The corpus's one preregistered confirmatory test of the underlying law did not meet its own declared margin [Curzi, 2026i], and that is the current evidential status of every applied statement above.

Nothing here identifies a social effect. Section 3.3 is not a caveat; it is a boundary. Even a perfectly fitted `rho` is a coefficient in a declared specification, and Manski [1993] proves that no volume of observational data separates it from contextual effects or correlated unobservables. **A paper that reports a mean-field error and an internal/external correspondence is exactly the kind of paper that must say this out loud.**

Nothing here licenses an intervention on another person. B9 is absolute within the framework and is a stipulation rather than a derivation. The engineering rule and the ethical rule happen to agree at the one place the paper tests them against each other — capture structures are both prohibited and low-`rho` — and that agreement is fortunate, not necessary.

The surplus functional's form is not established by this paper and is contradicted at one point by the best available evidence. Uzzi and Spiro [2005] finds an inverted U where `ln(1 + (N-1)u)` is monotone, and Putnam [2007] finds a short-run negative effect of diversity on trust where the corpus's founding claim expects a positive one. Neither refutes the arithmetic; both bound what the arithmetic may be used to assert.

25.10.3 10.3 Relation to the surrounding literature

The design-principles literature supplies the empirical footing this paper's institutional arithmetic presupposes and does not establish [Ostrom, 1990, 2010, Cox et al., 2010], against the standing prediction it displaced [Hardin, 1968]. The formal core of reciprocity under repetition is older still [Axelrod, 1984, Axelrod and Hamilton, 1981], and the extension beyond the dyad by reputation [Nowak and Sigmund, 2005] is the mechanism that manufactures the second ledger Section 5.2 argues an unmatched transfer lacks.

The claim that a group's measured performance is distinct from the abilities of its members — which is what the effective-count reading of a roster asserts — has its strongest support in the collective-intelligence factor [Woolley et al., 2010] and its formal statement in the diversity-trumps-ability results [Hong and Page, 2004, Page, 2007]. **Those results are contested and the contest is printed here rather than omitted:** the formal theorem's conditions have been argued to be far more restrictive than the popular reading allows [Thompson, 2014, Singer, 2019, Grim et al., 2019], and the empirical work-group literature finds effects that are small, conditional, and often negative for demographic as opposed to informational diversity [Mannix and Neale, 2005, van Knippenberg and Schippers, 2007, Harrison and Klein, 2007].

The fragility results of Section 8.2 are the surplus-geometry form of a standard network finding [Albert et al., 2000, Cohen et al., 2000, Callaway et al., 2000], and the transfer is a modelling assumption while the arithmetic is exact. The coordination-cost law is the classical one [Brooks, 1975, Conway, 1968] and the semilattice caution against strict trees is older than the framework [Alexander, 1965]; the reason the cross-links matter is that bridging ties carry the non-redundant information [Granovetter, 1973, Burt, 1992]. The exit mechanism of B1 is Hirschman [1970] with a computed threshold attached.

Against all of it stands the measurement-theoretic objection. Michell [1997] holds that psychology assumes quantitative structure without testing for it; conjoint measurement is the test [Luce and Tukey, 1964, Krantz et al., 1971] and it has not been run here. **That objection is not answered in this paper. It is conceded.**

25.11 11 Verification

The results of this paper compound from a small set of postulates. Each is atomic, each is checkable on its own, and each carries its status. **Where a postulate could be dropped and the result still follow, it has been dropped**; what remains is what the chain actually consumes.

25.11.1 11.1 Dimensional audit

The equations must balance before anything else about them is worth checking. `[Q]` denotes the dimension of stored order, `[X]` the dimension of the declared flow unit, and `[T]` the dimension of time; the period is the unit of `[T]`, and **in the discrete-step convention every per-period rate is a dimensionless number attached to a step**, which is why no explicit `[T^-1]` appears once the recursion is written in steps.

Symbol	Dimension	Unit	Status
`Q_t`	`[Q]`	declared order proxy	**measured**, in whatever the proxy counts
`U_t`, `u`	`[1]`	none	**defined**: `u = 1 - (x.y)^2` with `x`, `y` unit vectors
`N`	`[1]`	count of domains	**measured** by capability inventory

Symbol	Dimension	Unit	Status
`N_eff`	`[1]`	count	**derived**: `1 + (N−1)u`, the framework's affine effective count
`S = F(U)`	`[1]`	nats	**derived**: a logarithm of an effective count
`eta`, `eta_max`	`[Q]` per nat per step	proxy units per nat per period	**measured**; *not* dimensionless
`delta`, `delta_0`, `delta_eff`	`[1]`	per period	**measured** dimensionless fraction per step
`a`	`[1]`	probability per period	**measured**: honoured / unverified, a proportion of counts
`v`	`[1]`	fraction	**chosen**: the verification coverage, a policy variable
`c_v`	`[Q]` per step	proxy units per period at full coverage	**measured**
`K = eta S − C_0`	`[Q]` per step	proxy units per period	**derived**
`a*`	`[1]`	probability	**derived**
`C_t`, `c_i`	`[Q]` per step	proxy units per period	**measured**
`h_t`, `h_i`	`[Q]` per step	proxy units per period	**derived**
`Q_inf`	`[Q]`	proxy units	**derived**: `[Q]`/step divided by `[1]`/step returns `[Q]`
`x_ij`	`[X]` per window	declared symmetric flow unit	**measured**
`Sym`, `Skew`, `b_i`	`[X]` per window	same flow unit	**derived**
`rho`, `B`	`[1]`	none	**derived**: ratios of flow to flow, invariant to the flow unit
`tau` (turnover)	`[1]`	per period	**measured**: departures over average roster
`c_1`	`[Q]` per path per step	proxy units per path per period	**measured** by regression on `n(n−1)/2`
`D`, `P`	**no operational unit**	must share one	**silent in the source**; see below
`r = D/P`, `Delta = |ln r|`	`[1]`	nats for `Delta`	**derived**, and **demoted to an ordering** (Section 3.4)
`M`	inherits `E`	—	**demoted**; magnitude withdrawn

Do the equations balance? `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t` balances if and only if

every term on the right carries `[Q]`. The first does, since `delta` is dimensionless; the third does by construction; the second does only if `eta S_t` carries `[Q]`, and since `S_t` is a dimensionless count in nats, **`eta` must carry `[Q]` per nat per step and is therefore not dimensionless.** `delta_eff = delta_0 + (1−a)(1−v)` balances trivially: three dimensionless per-period quantities. `eta = eta_max rho` balances because `rho` is dimensionless. `a* = 1 − c_v delta_0/(K − c_v)` balances because `c_v` and `K` carry the same dimension and cancel, leaving `delta_0`, which is a pure number. `h = eta ln N − delta Q − C` balances because `ln N` is dimensionless, so `eta ln N` carries `[Q]` per step and is commensurable with `delta Q` and `C`.

Where the framework has no operational unit, it is said here. `Q` is stored order and has no natural unit; it is counted in whatever the declared proxy counts, and two systems' `Q` values are not comparable unless they declare the same proxy. **`D` and `P` have no operational unit at all in the source material**, which requires only that they share one; that silence is precisely what forces the demotion of Section 3.4, since a ratio without a common ratio scale and a shared true zero is a number with no referent. The flow unit `[X]` must be **symmetric** — the same unit in both directions — and where it is not, `rho` is undefined rather than small.

One domain restriction the source does not state. `delta_eff` must lie in `[0,1]` to be a retention loss at all. At `v = 0` this requires `delta_0 + (1 − a) <= 1`, that is **`a >= delta_0`**. Below that, Postulate PT returns a "loss" exceeding the whole stock in one period and the linear model has left its domain. The consortium of Section 9.1 has `delta_0 = 0.05` and `a = 0.62`, comfortably inside; the tables of Section 4.3 have `delta_0 = 0.04` and `a >= 0.50`, also inside. **The restriction is stated here because nothing in the source material states it.**

25.11.2 11.2 The postulates

P1 — The recurrence is linear and first order. `Q_{t+1} = a_r Q_t + x_t` with `a_r := 1 − delta` and `x_t := eta S_t − C_t` not depending on `Q`. *Status:* **DEFINITION** — the model form adopted by the framework. It carries no empirical weight on its own.

P2 — `0 < delta <= 1` implies `0 <= a_r < 1`. *Status:* **ARITHMETIC.** Subtracting a number in `(0,1]` from 1 gives a number in `[0,1)`.

P3 — A contraction with modulus `q < 1` on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least `q` each pass. *Status:* **CITED** [Banach, 1922]; the same statement in the corpus's own notation, with the exact error law, is Curzi [2026q].

P4 — Surplus per period is bounded: `0 <= S_t <= ln N`. *Status:* **CITED** — the range result of the surplus functional [Curzi, 2026o], following from `f(u) = ln(1+(N−1)u)` on `u` in `[0,1]`.

P5 — Lapse decomposes into a verification-irreducible part and a confirmation-addressable part, and verification is effective on the second. *Status:* **POSTULATED** (Postulate PT). No external support. Used only at R1, R2 and R6.

P6 — Only the matched fraction of directed interaction volume converts into durable structure. *Status:* **POSTULATED** (Postulate PR). No external support; the reputational mechanism of Nowak and Sigmund [2005] is consonant but does not establish it. Used only at R3, R4 and R5.

P7 — In the symmetric/antisymmetric decomposition of the interaction matrix, `Sym` measures the relation and `Skew` measures the obligation. *Status:* **POSTULATED** (Postulate PO) as to interpretation; the decomposition itself is **ARITHMETIC** and is unique for any square matrix. Used only at R3.

P8 — For real numbers, `|sum_j z_j| <= sum_j |z_j|`. *Status:* **ARITHMETIC** — the triangle inequality on the reals.

P9 — `2 min(p,q) = p + q − |p − q|` for non-negative `p`, `q`. *Status:* **ARITHMETIC.**

P10 — Independent per-period survivals multiply: `1 − delta_total = (1−delta_1)(1−delta_2)`. *Status:* **POSTULATED** as to the independence, **ARITHMETIC** given it. Violated where departures are caused by the decay (Section 8.1), in which case the composition understates the loss. Used only at the turnover row and at B7.

P11 — Attaching the unit `Hz` to the defined constant `k = (9/8)^2 / phi^2` is a physical claim about a rate, not an arithmetic consequence of the definition. *Status:* **POSTULATED** [Curzi, 2026]. It is listed because the constant is quoted at fifteen figures in the corpus and the two claims must not be conflated. **No result in this paper consumes P11.**

25.11.3 11.3 The compounding chain

```
P1 + P2          -> L1   T(Q) = a_r Q + x has |T(Q)-T(Q')| = a_r|Q-Q'| with a_r < 1,
                         so T is a contraction.                          ARITHMETIC

L1 + P3          -> L2   T has exactly one fixed point Q_inf = x/delta,
                         reached from any start.                         CITED (banach1922operations)

L2 + P4          -> L3   with x = eta F(u*) - c*, Q_inf = (eta F(u*) - c*)/delta,
                         and Q_inf < 0 exactly when eta F(u*) < c*.      ARITHMETIC

P5               -> L4   delta_eff = delta_0 + (1-a)(1-v), so at v = 0,
                         delta_eff = delta_0 + 1 - a.                    POSTULATED

L3 + L4          -> R1   Q_inf(a) = K/(delta_0 + 1 - a); at K = 0.2526 and
                         delta_0 = 0.04 this is 0.4679 at a = 0.50 and
                         5.0529 at a = 0.99.                             ARITHMETIC

R1               -> R1b  the point elasticity is
                         (dQ_inf/da)(a/Q_inf) = a/(delta_0 + 1 - a),
                         = 10.556 at a = 0.95, NOT a/(1-a) = 19.00.      ARITHMETIC

L3 + L4 + P5     -> R2   Q_inf(v) = (K - c_v v)/(delta_0 + (1-a)(1-v)) is
                         linear-fractional in v, so sign(dQ_inf/dv)
                         = sign[(1-a)(K - c_v) - c_v delta_0] independent of v;
                         the optimum is an endpoint and switches at
                         a* = 1 - c_v delta_0/(K - c_v) when K > c_v.     ARITHMETIC

P9               -> L5   rho = 2min(x,y)/(x+y) = 1 - |x-y|/(x+y).        ARITHMETIC

P7 + L5          -> L6   rho_ij = 1 - |Skew_ij| / Sym_ij, so reciprocity and
                         obligation are two readings of one measurement.  ARITHMETIC

P8 + L6          -> R3a  |sum_j (x_ij - x_ji)| <= sum_j |x_ij - x_ji|, hence
                         B <= dyad-level imbalance, hence 1 - B >= rho.  ARITHMETIC

P6 + L3          -> R3b  eta = eta_max rho, so Q_inf = (eta_max rho S - C)/delta
                         and Q_inf < 0 for rho < C/(eta_max S).          POSTULATED

R3b              -> R4   at S = 1.5347, C = 0.30, eta_max = 0.45, the survival
                         threshold is rho > 0.4344; the measured 0.3824
                         gives Q_inf = -0.0836.                          ARITHMETIC

P4 + L3          -> R5   removing one member at fixed u costs
                         eta(ln N - ln(N-1)) of ceiling; at N = 8, u = 0.52,
                         eta = 0.4026 this is -7.7% of eta S.            ARITHMETIC

P1 (twice)       -> R6   two compartments run separately give
                         Q(t) = Q_A(1-d_A)^t + Q_B(1-d_B)^t, against the
                         mean-field 100(1-d_bar)^t; at d_A = 0.05,
```

```
                    d_B = 0.50 the gap is 0.00% at t = 1 and
                    92.20% at t = 12.                                  ARITHMETIC

P10          -> R7  delta_total = 1 - (1-delta_int)(1-tau); at
                    delta_int = 0.05, tau = 0.10 this is 0.1450 and
                    Q_inf falls to 34.5% of its tau = 0 value.         POSTULATED
```

Every line is checkable from its own inputs. **R1, R1b, R2, R3a, R4, R5 and R6 are arithmetic consequences of one definition (P1), three arithmetic facts (P2, P8, P9), one cited theorem (P3), one cited range result (P4), and — for R1, R1b, R2 and R4 only — the two postulates P5 and P6.** R3a and R6 consume no postulate at all: **the bound `1 - B >= rho` and the mean-field error are theorems, not modelling claims.**

What this chain does not establish. It does not establish that any relation obeys P5 or P6. Those are the modelling postulates of Section 3.2, they are not verified here, and the one pre-registered confirmatory test of the underlying recursion so far reported in this corpus did not meet its own declared margin [Curzi, 2026i]. **A chain whose steps are ARITHMETIC and CITED delivers the mathematics with certainty and delivers nothing whatever about institutions.** The two results a reader should treat as *established* rather than *modelled* are R3a and R6; everything else inherits P5 or P6.

25.11.4 11.4 The verification instance, worked by hand

The smallest instance that exercises the chain end to end is the consortium baseline of Section 9.1, and it can be checked in eight lines.

```
(1)  matched volume, from the flow matrix, summing 2 min(x_ij, x_ji) over
     the 28 unordered pairs:                        26                  ARITHMETIC
(2)  directed volume, summing all 56 off-diagonal cells:  68            ARITHMETIC
(3)  rho = 26/68                                = 0.3823529412          ARITHMETIC
(4)  eta = 0.45 x 0.3823529412                  = 0.1720588235          POSTULATED (P6)
(5)  S   = ln(1 + 7 x 0.52) = ln 4.64           = 1.5347143662          ARITHMETIC
(6)  eta S                                      = 0.2640611483          ARITHMETIC
(7)  delta = 0.05 + (1 - 0.62)                  = 0.4300000000          POSTULATED (P5)
(8)  Q_inf = (0.2640611483 - 0.30)/0.43         = -0.0835787249         ARITHMETIC
```

which rounds to the `-0.0836` reported throughout. Three further checks a reader can redo in one line each. **The node index.** `sum |b_i| = 17+21+14+8+5+5+3+3 = 76`, so `B = 76/136 = 0.5588235294` and `1 - B = 0.4411764706`. Since `0.4412 > 0.3824`, the bound `1 - B >= rho` holds, and the overstatement is `0.4412/0.3824 - 1 = 15.4 per cent`. **Headroom.** `0.1720588235 x ln 8 - 0.43 x 1.20 - 0.30 = 0.3577863 - 0.516 - 0.300 = -0.4582137`. **The admissibility floor.** `(0.30 - 0.2640611)/(1 - 0.43) = 0.0630506`, which the trajectory crosses between quarter 3 (0.1542) and quarter 4 (0.0520).

A rounding sensitivity worth printing. The threshold table of Section 4.5 is computed at the rounded `K = 0.2526`. The unrounded `K = eta S - C_0 = 0.35 ln(4.85) - 0.30 = 0.2526425467` gives `a* = 0.8480` at `c_v = 0.20` against the printed 0.8479, and `a* = 0.2407` at `c_v = 0.24` against the printed 0.2381. **The `c_v = 0.24` entry is sensitive because `K - c_v` is small there — a 0.017 per cent change in `K` moves `a*` by 2.6 points** — and any application near that corner should carry `K` unrounded.

25.11.5 11.5 Limits, asymptotics, and the model's valid domain

`delta -> 0^+`. `Q_inf -> +infinity` for a positive numerator and `-infinity` for a negative one. **The limit is singular**: at exactly `delta = 0` the map is an isometry, the fixed point does not exist at all rather than being infinite, and no mistake ever decays out of the stock. **This is the**

boundary at which Section 4.6's fourth negation bites: trust driven to its ceiling in a system with `delta_0 = 0` produces an institution that cannot revise.

`a -> 1.` `delta_eff -> delta_0` regardless of `v`. **`v -> 1.`** `delta_eff -> delta_0` regardless of `a`. **The two controls substitute exactly at their ceilings**, which is the entire content of Postulate PT and also the reason `a` becomes unmeasurable under `v = 1`: with nothing left unverified there is no sample from which to estimate the reliance rate.

`a -> delta_0` from above. `delta_eff -> 1` at `v = 0`, the whole stock lost per period. **Below `a = delta_0` the model has left its domain** (Section 11.1).

`c_v -> 0.` `a* -> 1`: verify almost always. **`c_v -> K^-.`** `a* -> -infinity`, so no achievable `a` lies below it and the reading is "verify nothing"; **`c_v >= K`** must be read from the derivative directly, where `(1-a)(K - c_v) - c_v delta_0 < 0` for every `a`, giving "verify nothing" again. **The two readings are continuous with one another, but the `a*` formula is valid only for `K > c_v` and the source correctly says so.**

`rho -> 0.` `eta -> 0`, `Q_inf -> -C/delta <= 0`. **A perfectly one-directional relation does not accumulate slowly; it does not accumulate.** **`rho -> 1.`** `eta -> eta_max` and no further gain is available from balance.

`N -> 1.` `ln N = 0` and `F(u) = 0` for every `u`. **A single domain has zero available surplus, not a small amount**, so `Q_inf = -C/delta <= 0` and headroom is negative at any positive stock. **`N -> infinity.`** `ln N -> infinity` but only logarithmically, with marginal gain `ln N - ln(N-1) ~ 1/N`, against a coordination cost growing as `n(n-1)/2`. **The two curves cross, and Section 6.4 computes where.**

`u -> 0.` `N_eff -> 1`, `S -> 0`: every category is a duplicate and the roster has left the arithmetic. **`u -> 1.`** `N_eff -> N`, `S -> ln N`, the ceiling attained. Strict concavity makes the approach front-loaded.

`tau -> 1.` `delta_total -> 1` regardless of `delta_int`: total turnover erases the store in one period however well it is documented.

Stated domain. The results above are valid for `delta` in `(0,1]` with `delta_0 <= a <= 1`; `0 <= v <= 1`; `eta_max > 0`; `0 <= rho <= 1`; `N >= 2`; `u` in `[0,1]`; `K > c_v` for the `a*` form; and `C_t <= (1-delta)Q_t + eta S_t` each period. Outside that domain the equations still have solutions and **those solutions should not be interpreted as statements about a stock.**

25.11.6 11.6 Invariance and scale type

Two invariances are claimed and both can be exhibited. First, `rho` and `B` are invariant under any positive rescaling of the flow unit: replacing every `x_ij` by `lambda x_ij` leaves `2 min(x,y)/(x+y)` and `sum|b_i|/(2 sum x)` unchanged, because numerator and denominator scale together. **That is what licenses `rho` to multiply `eta_max`**, and it fails immediately if the two directions are counted in different units — which is why S5 exists. Second, the storage law is invariant under rescaling `Q` by a positive constant provided `eta` and `C` rescale with it, while `delta`, `N` and `u` are unchanged. **`delta`, `u`, `N`, `a`, `v`, `rho` and `B` are the unit-free quantities of this paper; `Q`, `eta`, `C`, `c_v` and `c_1` are the unit-carrying ones.** Two systems can be compared directly on the first list and not on the second.

What would break the invariances: a `delta` varying with `Q`, which would make the storage transformation non-linear; an asymmetric flow unit, which makes `rho` undefined; or an `eta_max` that itself depends on volume, which the multiplicative form excludes by assumption.

How many scopes the correspondence has been tested at: two, once, on a constructed pair. Section 7.5 exhibits agreement to within two per cent between one internal and one external trajectory, both generated from stipulated parameters. **That is an exercise of the invariance, not an establishment of it**, and Test 1 of Section 8.11 is what establishment would require.

Scale type, restated as a verdict. `a`, `v`, `rho`, `B` and `tau` are proportions of counted events and are treated as ratio-scaled; that treatment is defensible **conditional on the counts being counts**, which the enumeration steps of Protocols A, B and C exist to secure. `Q`, `eta`, `C` inherit the declared proxy's scale. `D`, `P`, `r`, `Delta` and `M` are **demoted to orderings**. **The conjoint-measurement**

conditions that would license the ratio-scale treatment of any of these have not been tested [Luce and Tukey, 1964, Krantz et al., 1971, Michell, 1997], and until they are the treatment is a stipulation rather than a finding.

25.11.7 11.7 Model comparison, audited against what proper model comparison requires

RMSE against a baseline is not model comparison. **This paper fits nothing**, so most of the requirement does not bind on it directly; it binds on any application of it, and it binds on the corpus's one applied fit. Both audits are printed.

Requirement	Status in this paper	Status in the corpus's applied fit [Curzi, 2026i]
Out-of-sample rather than in-sample performance	**Not applicable.** Nothing is fitted; every number is the output of stipulated parameters.	**Met.** One-step-ahead on a calendar holdout examined once.
A matched null with the same free-parameter count	**Not applicable.**	**Met.** AR(1) with intercept, 2 parameters against 2.
Number of free parameters stated explicitly	**Met, in the sense available.** The three worked systems stipulate `(N, u, eta_max, rho, a, delta_0, C)` — seven stipulated quantities, none estimated.	**Met.** 2 for the storage model.
An information criterion, so the parameter count is paid for	**NOT MET, and not met anywhere in this corpus.** No AIC or BIC was computed for any model in this paper or, so far as this author is aware, in any paper of this corpus.	**NOT MET.** No AIC or BIC at any block length for any model.
Effect sizes with intervals, not only p-values	**PARTIALLY MET.** The one estimated quantity in the paper — `a_hat = 13/21 = 0.619` in Section 9.1 — is reported with its Wilson interval, approximately `[0.41, 0.79]`, which is **wider than several differences the diagnosis turns on**. No other quantity carries an interval, because no other quantity was estimated.	**PARTIALLY MET.** A profile interval on the retention coefficient; no interval on the RMSE differences and no effect size beside the shuffle `p`.

The unmet row is stated plainly because it changes how every comparative claim in this corpus should be read. No information criterion has been computed anywhere, so no parameter count has ever been paid for, and a margin over a matched baseline has never been tested against a criterion that penalises complexity. **The honest summary is that this corpus has reported held-out error and has not reported model comparison.**

Two further statistical notes. The `a_hat` interval above is the single most important number in Section 9.1 that nobody would otherwise print: with 21 unverified commitments the estimate cannot distinguish `a = 0.62` from `a = 0.50`, and the difference between those two moves

`Q_inf` from `−0.0836` to `−0.1180`. **The diagnosis survives — both are negative — but the *magnitude* does not, and any intervention sized against the magnitude is sized against noise.** And the shuffled-input control invoked in every falsification list above is **anti-conservative as ordinarily run**, measuring 0.108 against a nominal 0.05 [Curzi, 2026i]; that figure belongs beside every `p` this corpus reports.

25.11.8 11.8 Analogy discipline, and the precision of the numbers

Every use of a physical or engineering result in this paper is a shared mathematical form, not a derivation from physical premises. The storage law and a first-order low-pass section are governed by the same difference equation, so every theorem about first-order discrete systems applies to both [Oppenheim et al., 1996, Åström and Murray, 2008]; nothing about charge, capacitance or current is claimed to hold of an institution, and no physical result's authority transfers. **A leaky integrator is a leaky integrator whether it holds charge or stored order, and that is the whole of the claim.** The "mean field" of Section 8.2 is the same: the phrase names an approximation scheme, and the failure computed there is a failure of that approximation, not an appeal to statistical mechanics. The "frontier" of Section 4.3 is a straight line in a two-cost plane and borrows nothing from economics beyond the word.

Entropy, where it appears near `ln N`, is Shannon entropy — a dimensionless information measure in nats [Shannon, 1948] — **and not thermodynamic entropy**, which carries `[energy][temperature]^−1`. The framework does not identify surplus with either.

The exponents in Section 6.4 are defined by the model and not estimated from data. No distribution was fitted, no heavy tail is claimed, and the procedure any estimated power-law claim would require [Clauset et al., 2009] therefore does not apply — but it applies immediately to any attempt to recover `c_1` or a cost exponent from real organisational data, and such an attempt must run it and report the result.

Precision. The numbers in this paper are of three kinds. *Defined or derived-from-defined* quantities — `ln 8 = 2.0794`, `ln 2 / −ln(1−delta)`, the geometric limits, and `k = (9/8)^2 / phi^2 = 0.483425732988414` where it is quoted from Curzi [2026] — are **exact to the precision printed** and may be recomputed further. **`k` is quoted at fifteen figures because it is a defined constant, and it is labelled *defined* wherever it appears at that precision; attaching `Hz` to it is postulate P11 and is not consumed by anything here.** *Illustrative model output* — every trajectory entry, every `Q_inf`, every headroom — is the output of stipulated parameters and **carries no measurement uncertainty because there was no measurement**; four decimal places there record the arithmetic, not the confidence. *Estimated quantities* appear exactly once, at `a_hat = 0.619` in Section 9.1, and are reported with an interval as Section 11.7 requires.

25.12 12 Limitations

The three worked systems are constructed, not observed. No dataset was collected for this paper. Every number in Section 9 is the output of the stated parameters under the stated model, carried through so a reader can re-derive it. They demonstrate that the procedure runs and that it reports its own failures; **they are not evidence that any organisation behaves this way and must never be cited as such.**

Which parameters were chosen rather than measured. In all three systems (`N, u, eta_max, rho, a, delta_0, C`) are stipulated. `eta_max = 0.45` in the consortium and `eta_max = 0.5719` in the rescue association are respectively an estimate from three assets and a back-solve from a measured `eta`; neither is an independent measurement. The post-sanction `rho = 0.9038` of Section 6.3 cannot be recomputed from anything printed.

The three mappings are postulates and none is derived. Postulate PT locates trust in `delta`, PR locates reciprocity in `eta`, PO interprets the antisymmetric part as obligation. **Each is a modelling choice about how a measured social quantity enters a proven equation.** The chain

of Section 11.3 shows exactly which results consume which, and two results — the bound `1 - B >= rho` and the mean-field error — consume neither.

The reflection problem, and why it outranks everything else here. Section 3.3 states it in full: applied to a group whose members both influence and constitute the aggregate, the specification is reflexive, and Manski [1993] proves the endogenous social effect is not identified from observational data. **No fit recovers that parameter**, and Shalizi and Thomas [2011] show the same confounding for network data. **This paper adopts no remedy, so the applied reading of `rho` and `eta` is a modelling position rather than an identified empirical result.** The most realistic available repair is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001]; Protocol B's declared window and fixed membership set are a step toward it and not a substitute for it.

The Lucas critique applied to this framework's own parameters. The paper asserts that publishing a ledger changes the behaviour it records (S12), that a measure becomes a target (Section 7.4), and that disclosure of sampling changes what is sampled (Protocol A) — and then treats `delta`, `eta`, `rho` and `a` as stable over an estimation window. **Both cannot be true without qualification** [Lucas, 1976]. The position taken here is the narrow one: the parameters are treated as stable **only over a declared window and only while the quantity is not itself a reported target.** Where `rho` or `a` is reported upward as a performance number, the framework's own reflexivity applies to it and the constancy assumption fails. **That is a real restriction on where the model may be used and it is stated as one.**

`u` and the three new quantities have not been shown to be quantities. Michell [1997] argues that psychology routinely fails to test whether its attributes are quantitative before measuring them. Conjoint measurement is the available test [Luce and Tukey, 1964, Krantz et al., 1971] and a Rasch demonstration is the alternative [Rasch, 1960]; **neither has been run for `u`, for `a`, for `rho` or for `B`, and until one is, the ratio-scale treatment of all four is a stipulation.** The case for `a`, `rho` and `B` is stronger than for `u` — they are proportions of counted events rather than psychological magnitudes — but "stronger" is not "tested."

`N_eff = 1 + (N-1)u` is not a Hill number and no external literature contains it. `1/HHI` is the Hill number of order two and `exp(H)` the Hill number of order one, both fully grounded [Hill, 1973, Jost, 2006]; `1 + (N-1)u` is a linear interpolation between 1 and `N` and is this framework's own construction. A reviewer reaching for that literature as support will find it does not apply.

The surplus functional's form is underdetermined, and at one point contradicted. Latané [1981] and Briscoe et al. [2006] fit the same qualitative facts, and nothing assembled here discriminates between them. Worse for the corpus, Uzzi and Spiro [2005] finds an inverted U where `ln(1+(N-1)u)` is monotone, and Putnam [2007] finds a short-run negative effect of diversity on trust — attenuated but not eliminated under better controls [Abascal and Baldassarri, 2015], and conditional on contact conditions the framework does not model [Allport, 1954, Pettigrew and Tropp, 2006]. **The functional form is chosen for its axiomatic characterisation, not established by measurement, and the best available evidence is not uniformly on its side.**

The bang-bang verification result is an artefact of linearity. Section 4.5 says so where it is derived. Under a convex verification cost the optimum is interior; what survives is the threshold structure, and any applied use must first measure whether the cost is linear.

The mean-field result cuts both ways. Section 8.2 computes a 92 per cent error for a two-compartment store, but **the compartment split, the two rates and the equal initial division are all stipulated.** A real body's compartment structure is unknown, and the error would be smaller for a tight distribution and larger for a wider one. The finding is that a scalar `delta` is not safe, not that it is wrong by 92 per cent.

The nine breakdowns are asserted with three of them computed. B1, B2 and B8 carry arithmetic. B3 carries a three-line Condorcet construction. **B4, B5, B6, B7 and B9 are argued, not computed**, and B9 is a normative stipulation with no external support in this corpus's bibliography at all. That is stated here as plainly as it is stated in Section 8.9.

Two classical results are used without a citation key. Arrow's impossibility theorem and the Condorcet paradox (Section 8.3), and the intertemporal-choice and precommitment litera-

ture (Section 8.4), have no entries in this corpus's bibliography. **They are named descriptively and left uncited rather than attached to a source that does not support them**, and any reviewer checking will find no misattributed citation because there is none. The source material's attribution of the Sicilian protection-market study to Gambetta (1993) is likewise not carried, because that monograph is not in the bibliography; only the 1988 edited volume is [Gambetta, 1988].

Four defects in the source material are corrected or withdrawn in this paper and are tabulated in Section 9.4: the elasticity formula omitting `delta_0`; a damage ratio computed on surviving values rather than on losses; an unrecoverable clock-ratio figure; and a per-member headroom table whose denominator does not match its reported fraction. **A reader comparing this paper against the manuscript it converts should expect those four differences and no others.**

Nothing in this paper is a human-subjects finding. No interventions were run and no participants were recruited. The disclosure requirements of Protocol A and the ledger-privacy requirement of S12 are design constraints on prospective work, not reports of past practice.

Unmet rigour checks, printed rather than left to a reviewer. (i) *Information criteria.* **NOT MET here and, so far as this author is aware, nowhere in this corpus.** (ii) *Effect sizes with intervals.* Partially met: one estimated quantity carries a Wilson interval and nothing else was estimated. (iii) *Scale type.* Conjoint conditions untested; four quantities treated as ratio-scaled by stipulation and five demoted to orderings. (iv) *Independent replication.* None. Every system here was constructed by one author, and the reproducibility literature is unambiguous about what that is worth [Open Science Collaboration, 2015, Ioannidis, 2005]. (v) *Directional versus point prediction.* Meehl [1967] observes that directional predictions are nearly costless in the social sciences; H2, H4 and H6 make point predictions with constants, which is the right answer to that objection, but none has been put at risk. (vi) *Identification.* Unaddressed by design, and stated as such in Section 3.3.

25.13 13 Conclusion

Trust, reciprocity and obligation enter the storage law at three different positions, and the positions are why the three behave differently. Trust enters as retention: `delta_eff = delta_0 + (1 − a)(1 − v)`, where `a` is the measured rate at which unverified commitments are honoured. At `K = eta S − C_0 = 0.2526` and `delta_0 = 0.04` the steady stock rises from `Q_inf = 0.4679` at `a = 0.50` to `Q_inf = 5.0529` at `a = 0.99`, and **the point elasticity of the stock in the reliance rate is `a/(delta_0 + 1 − a)` — 0.93 at `a = 0.50` and 10.56 at `a = 0.95` — not `a/(1−a)`, which holds only in the limit `delta_0 -> 0` and is corrected here.** Under a linear verification cost the verification budget's optimum is always at an endpoint, switching at `a* = 1 − c_v delta_0/(K − c_v)`: at `c_v = 0.20` verify below `a = 0.8479`, and at `c_v >= K` verification consumes the entire net input and is never worth buying. **Trust is valuable exactly in proportion to the cost of the verification it displaces.**

Reciprocity enters as conversion, `eta = eta_max rho` with `rho = 2 min(x_ij,x_ji)/(x_ij+x_ji) = 1 − |Skew_ij|/Sym_ij`, so reciprocity and obligation are two readings of one measurement rather than two measurements. **The node-level imbalance index obeys `1 − B >= rho` and is therefore an upper bound on reciprocity, never an estimate of it**: it flattered one worked network by 15.4 per cent (0.4412 against 0.3824) and a second by 34 per cent (0.7243 against 0.5421), in both cases in the direction of comfort.

An eight-firm consortium at `u = 0.52` was collecting 73.8 per cent of its surplus ceiling and still had a steady state of `−0.0836` and headroom of `−0.4582`. The binding fault was `rho = 0.3824`, not difference, and the survival threshold was `rho > 0.4344` — 5.2 points away. Raising `rho` to 0.8947 and `a` to 0.88 moved the fixed point to `+1.8701`, and the one-at-a-time decomposition shows why the order of operations matters: **reciprocity alone gave `+0.7394`, trust alone gave `−0.2114` — worse than doing nothing — and the interaction term of `+1.2586` exceeded either main effect.** Where the numerator `eta S − C` is negative, lowering `delta` moves the equilibrium

further from zero. **Fix the numerator first.** Graduated sanction against expulsion prices out at `+1.0` per cent against `−7.7` per cent of `eta S`, and the `−7.7` is a floor rather than an estimate because it holds `u` fixed across the removal.

The mean-field scalar `delta` is exactly right for one period and wrong by 92 per cent after twelve when the store is split evenly between compartments at `delta = 0.05` and `delta = 0.50`, with the implied aggregate rate drifting from 0.2750 down to 0.0504. That is a theorem about compartments, not a modelling claim, and it is the sharpest available discriminator between a scalar model of a heterogeneous body and the body.

The correspondence between one head and one group is audited here rather than asserted, and it breaks in nine named ways — exit, heterogeneity and ill-defined aggregation rated severe, four structural, one a matter of measurement practice, and an absolute ethical asymmetry. What survives is precise: the recursion transfers, because it names no substrate; the redundancy geometry transfers, because geometry does not know what it is counting; the ceiling transfers and produces requisite variety in first-person form; and the three diagnostic questions transfer. **What does not transfer is anything involving agency of the parts. The equation transfers between scopes; the licence to apply it does not.**

None of this identifies a social effect. The reflexive specification the framework rests on is not identified from observational data, no remedy is adopted here, and every applied statement above is therefore a modelling position rather than a measured result. **The arithmetic is certain and the institutions are not, and a paper that says which is which is worth more than one that lets the certainty of the first stand in for evidence about the second.**

25.13.1 Declarations

Data and code availability

No external dataset was collected for this paper. All numerical results are computed from the stated parameters under the model developed in the text, and every table can be re-derived from the equations and parameter values printed beside it. The three systems worked in Section 9 are constructed examples with declared parameters, not field observations; the two flow matrices in Section 9.1 are printed in full precisely so that `rho`, `B` and every derived quantity can be recomputed by hand. The figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper and the model equations of Sections 3 to 8.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, produced the worked examples and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The three worked systems are constructed. The measurement protocols of Section 6.6 carry binding disclosure and privacy constraints — sampling of verification must be disclosed to every party,

and individual obligation balances must not be published within a group — and those are **design constraints on prospective work, not reports of any study performed.** The declination constraint invoked throughout is a normative stipulation of the framework; **no entry in this corpus's bibliography establishes it**, and it is badged as a stipulation in Sections 2, 3.5, 8.9 and 12.

Preregistration

This paper reports no preregistered test of its own; its results are derivations and constructed worked examples. Falsification rules are declared in advance for each of the three worked systems in Sections 9.1, 9.2 and 9.3, and for the correspondence in Section 8.11, but **none has been executed.** The preregistered protocol for testing the underlying storage law against a real system, together with the confirmatory application that **failed** its second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared margin of 10 per cent — is reported in Curzi [2026i]. Any applied claim resting on this paper inherits that result.

25.14 Appendix A — The interpersonal formulary

25.14.1 A.0 How the formulary is organised, and two standing verdicts

This appendix collects every formula developed above, plus the group-scope forms of the framework's core, in one reference with a fixed entry format: **statement**, **derivation** with its badge, **worked number**, and **named failure modes**. The seven validation gates `G1`–`G7` (Section 3.4) are consolidated into the single matrix of Section A.7 rather than repeated twenty-five times; **every non-`PASS` verdict is spelled out there, and an entry quoted without its gate row is an entry quoted without its limits.** Entries are labelled `IF-01` to `IF-25`; the labels do not collide with the results `R1`–`R6`, the specifications `S1`–`S32`, the failure modes `F1`–`F26`, or the breakdowns `B1`–`B9`.

Standing verdict 1 — every `M`-derived entry is demoted on gate G2. `D` and `P` are not instrumented on a common ratio scale with a shared true zero, and division without a true zero produces a number with no referent. This governs **IF-14, IF-15, IF-16 and IF-17**. A count supplies a zero and a unit inside one declared encoder and supplies no linkage between encoders, so the magnitude of `r` moves when the encoder moves.

> **Withdrawn as measured quantities:** the magnitude of `D/P`; `Delta = |ln r|` as a number; and the quantitative form of the join from the moment into the storage law. The worked numbers under those four entries are **illustrative constructions demonstrating the arithmetic of each measure.**
>
> **In force, untouched:** the limiting cases; the `D = P` fixed point; the ordering *over-framed / calibrated / under-framed*; and the **sign** of `ln r`, which is the operationally decisive part because the two faults take opposite corrections. **What is demoted is magnitude, not structure.**

The route out is constructive, specific and priced. Measuring `D` and `P` as Rasch measures on one common linked scale makes `Delta = |ln(D/P)| = |theta_D − theta_P|` genuinely ratio-scaled, with Rasch separation supplying the true zero that division requires [Rasch, 1960, Andrich, 1988]; the corpus prices that instrument at approximately **123 items per term to establish a half-logit distortion at 95 per cent confidence**, and about **31 items for a full logit**. At group scope the items are decisions with later-observable answers, so **a body that already keeps a decision log is already accumulating the item bank** — the one respect in which the interpersonal version of the study is cheaper than the personal one. **A formulary that demotes four of its own entries, names exactly what survives, and prices the study that would restore**

the rest is the protocol working. It is not an apology. Curzi [2026] carries the demotion in the corpus's register.

Standing verdict 2 — the shuffled-input control has three measured failure modes, and every G6 verdict below inherits all three.

#	Failure mode	Measured
1	**Anti-conservative as ordinarily run.** Against nominal `alpha = 0.05` over 500 null trials: iid permutation false-rejected at **0.108, 2.2x nominal**; circular shift 0.120; block permutation 0.060, near nominal	[Curzi, 2026i]
2	**Blind to repackaging.** A filtered repackaging passed all four baselines including the shuffle at `p = 0.0005`, and was caught **only by the structural ceiling test** — its single-step jump was **38.5x** the theoretical maximum	[Curzi, 2026i]
3	**Blind to shared slow drivers.** Across 4,000 trials the shuffle detected the clean case 48.17% of the time and **passed 98.35% of the time against a confound driving both series**	[Curzi, 2026i]

> **The rule, superseding any earlier statement of it: the shuffle is a weak, miscalibrated empirical audit. Repackaging is caught only by the structural ceiling test; confounding by a shared slow driver is caught only by design controls — blinding, channel separation, prospective sampling, a yoked control. No surrogate catches either.** Calibrate the surrogate by simulation on the design in hand before using it, and **build design controls, not more surrogates.**

Index by parameter. `N`: IF-02, IF-04, IF-05, IF-22. `u`: IF-01, IF-02, IF-03, IF-05. `S/F`: IF-01, IF-02, IF-03, IF-04. `Q`: IF-06, IF-07, IF-08. `delta`: IF-11, IF-18, IF-19, IF-25. `eta`: IF-20, IF-21. `C`: IF-12, IF-13, IF-19. `h`: IF-09, IF-10, IF-22. `M`, `D/P`: IF-14, IF-15, IF-16, IF-17. Thresholds: IF-23, IF-24.

Index by symptom. "lots of people, no new ideas" — IF-02 then IF-05. "we agree things and start over every quarter" — IF-08 then IF-20. "meetings are necessary and useless" — IF-12 then IF-13. "we're busy but nothing accumulates" — IF-20, IF-21. "we keep being surprised by the same thing" — IF-16, IF-17. "good people keep leaving" — IF-22, IF-11. "the reorganisation didn't help" — IF-09, IF-10. "one person left and everything got worse" — IF-05. "we cannot get the change to stick" — IF-24. "the numbers from the team and the person disagree" — IF-25.

25.14.2 A.1 Class A — geometry of difference

IF-01 · The pair surplus.

```
S_2 = ln(1 + u),     u = 1 − (x . y)^2  in [0,1],     ceiling ln 2 = 0.693147
```

Derivation: the surplus functional at `N = 2` [Curzi, 2026o]; not a special case requiring separate proof. **Extremality: `S_2 = 0` if and only if the framings are parallel**, which is the formal content of the statement that the pairing of like with like is a non-event. *Protocol:* choose and justify the framing space; each party independently predicts `k >= 8` observables without conferring; encode as unit vectors and compute `u = 1 − (x.y)^2`; repeat after one period and report test-retest, below 0.6 the estimate is not a measurement.

u	0.00	0.10	0.25	0.50	0.75	0.90	1.00
S_2	0.0000	0.0953	0.2231	0.4055	0.5596	0.6419	0.6931
% of ln 2	0.0	13.8	32.2	58.5	80.7	92.6	100.0

Half the available surplus of a pair is bought by the first quarter of the way to orthogonality — strict concavity arriving as a design fact. *Failure modes:* similarity mistaken for compatibility, where a pair chosen for agreement has `S_2 = 0` exactly; encoded difference, where vocabulary differences register as `u > 0` while predicting identically, which the independent-prediction step excludes; **conscripted difference**, where an instructed disagreement produces `u` with no informational content and is prohibited where the partner cannot decline the role.

IF-02 · The effective count and the group surplus.

```
N_eff = 1 + (N − 1)u,     S = ln N_eff = ln(1 + (N−1)u),     0 <= S <= ln N
```

Derivation: the surplus functional, with the effective count coming from the affine design axiom. **That axiom is chosen, not derived, and this entry must not be used without saying so** [Curzi, 2026o]. What it buys is that surplus is the log of a headcount *after the headcount has been audited for redundancy*. **It is not a Hill number** (Section 3.1). *Protocol:* capability inventory for `N`; pre-discussion prediction round for `u`; compute `N_eff` and `S`; **never measure `N_eff` directly**; report `S / ln N` alongside `S`.

u	0.00	0.10	0.20	0.35	0.50	0.65	0.80	1.00
N_eff at N = 12	1.000	2.100	3.200	4.850	6.500	8.150	9.800	12.000
S	0.0000	0.7419	1.1632	1.5790	1.8718	2.0980	2.2824	2.4849
% of ln 12	0.0	29.9	46.8	63.5	75.3	84.4	91.8	100.0

Twelve people at `u = 0.10` are worth 2.1 people. Not "somewhat redundant" — 2.1. *Failure modes:* headcount substituted for `N`, the single commonest error in this subject; reporting `S` without `S / ln N`, which is half a diagnosis since ceiling and realised fraction are different levers; reporting `N_eff` from a survey, when it is derived and cannot be asked about.

IF-03 · The two-source decomposition.

```
u = u_cross + u_div,    u_cross = beta^2,   u_div = alpha^2 sin^2(gamma)
max( f(u_cross), f(u_div) )  <=  f(u)  <=  f(u_cross) + f(u_div)
```

Derivation: the two-source theorem for a pure block vector. **The upper bound is subadditivity on the parameter `u` and is not a triangle inequality** — the triangle inequality fails for geometric triples, and **this entry must not be chained** (S31). *Protocol:* declare the category partition **before** measuring, from an external criterion; a partition chosen after seeing the data is not a partition, it is a fit. Compute `u` within and across blocks separately, then verify the bound; **a violation means the block assignment is wrong, not that the theorem is.** *Worked:* `N = 12`, `u_cross = 0.24`, `u_div = 0.16`, so `u = 0.40`:

```
f(u)       = ln(1 + 11 x 0.40) = 1.6864
f(u_cross) = ln(1 + 11 x 0.24) = 1.2920
f(u_div)   = ln(1 + 11 x 0.16) = 1.0152
bound:  1.2920 <= 1.6864 <= 2.3072      satisfied
subadditivity slack = 2.3072 − 1.6864 = 0.6208
```

The slack of 0.6208 nats is the amount by which naive addition of the two sources would overstate the surplus — **36.8 per cent of the true value.** *Failure modes:* chaining across three parties, prohibited; post-hoc partition; summing the sources.

IF-04 · Requisite variety at group scope.

```
required:  N_eff >= N_env,     hence  u >= (N_env − 1)/(N − 1),
feasible only if N >= N_env, since u <= 1
```

Derivation: only variety absorbs variety [Ashby, 1956, 1958]. **The transfer into this parameterisation is a modelling assumption, badged.** What is proven here is only the algebra. *Protocol:* enumerate the environment states requiring **distinct** responses, merging any two whose correct response is identical; compute `N_eff` per IF-02; if `N_eff < N_env` the shortfall is the number of states the group must handle by guessing.

`N_env`	required `u` at `N = 8`	at `N = 12`	at `N = 20`
4	0.4286	0.2727	0.1579
8	1.0000	0.6364	0.3684
16	infeasible	infeasible	0.7895
32	infeasible	infeasible	infeasible

An eight-member body facing an eight-state environment must be perfectly orthogonal to be adequate, which is not achievable in practice. *Failure modes:* inflated `N_env`; treating the requirement as sufficient, when it is necessary only — the group must also convert and afford.

IF-05 · Departure fragility: targeted versus random removal.

```
u_ij = 1 − c_i c_j,   c_i = cos^2(th_i)
mean u = 1 − [ (sum c)^2 − sum c^2 ] / (2 ∗ C(N,2))
fragility = S(N,u) − S(N−1, u after removing a named member)
```

Derivation: definition plus the surplus functional. **The one-parameter geometry is a modelling choice, badged**: each member holds a shared direction plus one private orthogonal direction, the minimal structure producing heterogeneous `c_i`. *Protocol:* estimate each `c_i` from the member's prediction vector against the group centroid; compute mean `u`; **recompute `u` with each member deleted in turn**; report the maximum, not the mean. *Worked:* `N = 9`, six members at `c = 0.95` and three at `c = 0.0501`, giving `u = 0.5999` and `S = 1.7578`:

removal	resulting `S` at `N = 8`	loss
random member (expectation)	1.6434	0.1144
a conformist (`c = 0.95`)	1.7162	0.0416
the most divergent (`c = 0.0501`)	**1.4979**	**0.2599**

Targeted removal costs 2.27 times random removal, which is the standard network result — heterogeneous structures robust to random failure and fragile to targeted attack [Albert et al., 2000] — appearing in surplus geometry rather than in a connectivity graph. **The transfer is a modelling assumption; the arithmetic is exact.** And the deeper case: two nine-member groups can have identical `u = 0.5999` and `S = 1.7578`, and be left with either 1.3862 or 0.3971 of surplus when three members leave, **losing 0.3716 or 1.3607 respectively — a damage ratio of 3.66, not the 3.49 obtained by dividing the surviving values** (Section 9.4, defect 2). *Failure modes:* scalar `u` reporting, which conceals the entire profile; assuming the loudest member is the most divergent, when divergence is measured by prediction disagreement and not by volume; retention policy by seniority, which is uncorrelated with `c_i` — subject always to the constraint that **retention is an offer and never a restriction**.

25.14.3 A.2 Class B — storage and dynamics

IF-06 · The institutional storage law. Statement, derivation, protocol and honesty clauses are those of Section 3.1 and are not repeated. **The mathematics is proven; the mapping of an institution onto a single scalar `Q` is a postulate, badged.** *Worked:* `N = 12`, `u = 0.45`, `eta = 0.40`, `C = 0.55`, `delta = 0.10`, `Q_0 = 0`, so `S = 1.7834`, `eta S = 0.7134` and `Q_inf = 1.6336`:

t	0	1	2	3	5	8	10	12
Q	0.0000	0.1634	0.3104	0.4427	0.6690	0.9304	1.0640	1.1722

Half-life 6.579 periods; 95 per cent settling at 28.43. *Failure modes:* ordinal `Q`, since a five-point survey cannot enter this equation; `S` inferred from `Q`, which destroys identifiability; **declaring success at `t = 3`, when the system is 27 per cent of the way to its own equilibrium.**

IF-07 · The exact solution and the contraction.

```
Q_t = (1−delta)^t Q_0 + sum_{k=0}^{t−1} (1−delta)^{t−1−k} (eta S_k − C_k)
|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|      under constant input
tau = 1/delta periods;   t_half = ln 2 / −ln(1−delta)
```

Derivation: induction for the solution; the contraction follows because the map is affine with slope `(1−delta)`, and Banach then guarantees a unique fixed point and convergence from any start [Banach, 1922, Curzi, 2026q]. *Protocol:* measure `Q` at two times under zero input, `delta = 1 − Q_1/Q_0`; **verify geometry with a third point — `Q_2/Q_1` must equal `Q_1/Q_0` within measurement error, and if it does not, the system is not single-compartment** (B2).

delta	tau	half-life	95% at t	99% at t
0.05	20.00	13.51	58.40	89.78
0.10	10.00	6.58	28.43	43.71
0.20	5.00	3.11	13.43	20.64
0.35	2.86	1.61	6.95	10.69

Failure modes: two-point `delta` on a heterogeneous store, which gives a number that is neither compartment's and drifts with when the analyst looked; measuring under non-zero input, which conflates `delta` with `eta S − C`.

IF-08 · The group steady state. `Q_inf = (eta F(u*) − c*)/delta` for `delta > 0` and constant input; **if `eta F(u*) < c*` the equilibrium is negative.** *Derivation:* set `Q_{t+1} = Q_t`. **The theorem states its own failure mode and this entry must reproduce it:** a negative `Q_inf` means the regime is unsustainable in this model, or the linear model must be replaced. It does not mean negative stored order, which is meaningless. *Protocol:* confirm the input is approximately constant over at least `3/delta` periods; fit from the transient; **check the sign of `eta F − c` first, and if it is negative, stop and report unsustainability rather than presenting a negative equilibrium as a small number.** *Worked:* `N = 12`, `eta = 0.40`, `C = 0.55`, `delta = 0.10`:

u	S	eta S	Q_inf
0.45	1.7834	0.7134	+1.6336
0.70	2.1633	0.8653	+3.1533

Raising `u` from 0.45 to 0.70 nearly doubles the equilibrium: **+1.5197 units for a change costing nothing in headcount.** *Failure modes:* reporting a negative `Q_inf` as a forecast; applying the formula to a non-stationary input.

IF-09 · Organisational headroom. `h_t = eta ln N − delta Q_t − C_t`, with the three regimes of Section 7.4 and `h = 0` at `Q = (eta ln N − C)/delta`. *Derivation:* the headroom gauge. **Note that headroom uses `ln N`, the ceiling, while the growth condition uses `S = ln N_eff`, the realised value. Half of all misdiagnosis turns on that distinction**: a system can have large positive headroom and shrink every period, because it operates far below its own ceiling. *Protocol:* derive from the four parameters; **never survey it.** *Worked:* `N = 12`, `eta = 0.40`, `C = 0.55`, `delta = 0.10`, so `eta ln N = 0.9940`:

Q	0	1	2	3	4	5	7
h	+0.4440	+0.3440	+0.2440	+0.1440	+0.0440	-0.0560	-0.2560

Ossification at `Q = 4.4396`. **The organisation runs out of headroom by succeeding, and the number at which it happens is computable in advance.** *Failure modes:* using `S` instead of `ln N`; reporting `h` without `S / ln N`; treating `h < 0` as a motivation problem, when **no arrangement of people prevents decline in that regime.**

IF-10 · The target increment and the maximum step.

```
Q_{t+1} − Q_t >= r_t   iff   S_t >= (delta Q_t + C_t + r_t)/eta
feasible only if  delta Q_t + C_t + r_t <= eta ln N
r_max = eta ln N − delta Q_t − C_t = h_t
required u = (e^{S_required} − 1)/(N − 1),   feasible only if u <= 1
```

Derivation: the growth condition rearranged. **No policy based on interaction geometry alone can guarantee more than `r_max`.**

The concavity correction, required before either feasibility test is read. The statement is written as though `U` were constant every period, and it never is. What accumulates is `sum eta F(U_t)`, so **the operative quantity is the mean of `F(U_t)`, never `F` of the mean `U`.** Because `F` is strictly concave, Jensen's inequality gives `mean of F(U_t) <= F(mean of U_t)`, with equality only if `U` is constant. **A body whose geometry varies delivers strictly less surplus than one sitting at its own average every period, so testing feasibility on `F(mean U)` approves increments that cannot be delivered.** In the corpus's worked office the penalty is 0.0761 nats — mean `F` of 1.1501 against `F(mean U)` of 1.2262, **a loss of 6.2 per cent of the surplus purely to variability** — and it was enough there to turn a feasible-looking plan into a shortfall. Report `u_req` for legibility; do not test against it.

Worked: `N = 12`, `eta = 0.40`, `C = 0.55`, `delta = 0.10`, `Q_t = 2.0`, requested `r_t = 0.30`:

```
required S = (0.10 x 2.0 + 0.55 + 0.30)/0.40 = 2.6250
ln N = 2.4849  ->  required S exceeds the ceiling. INFEASIBLE.
required u = (e^2.6250 − 1)/11 = 1.1641 > 1.  Confirmed infeasible.
r_max = 0.9940 − 0.2000 − 0.5500 = 0.2440
```

The most the body can deliver is 0.2440 and it was asked for 0.3000. The gap is 23 per cent and is not closeable by effort, exhortation, or reorganisation of the same twelve domains. *Failure modes:* accepting an unpriced target; re-planning instead of re-scoping when required `u > 1`; feasibility tested on `F(mean U)`.

IF-11 · Turnover composition. `delta_total = 1 − (1 − delta_int)(1 − tau)`, survivals multiplying because the two losses are independent. *Derivation:* definition plus an independence assumption. **The independence is a postulate, badged** (P10), and it is violated where departures are caused by the decay, in which case the composition understates the loss. *Protocol:* measure `tau` from the roster; measure `delta_int` during an interval with no departures; compose; **check the independence assumption by testing whether departures cluster after `Q` declines, and if they do, use the composed value as a lower bound only.** *Worked:* the table of Section 7.3. **Five per cent turnover halves steady-state stored order.** Not reduces — halves. *Failure modes:* adding the rates, which gives 0.15 where the truth is 0.1450 — small there and

large at high tau; attributing the loss to documentation quality, when at tau = 0.20 some 79 per cent of the total loss is people leaving; assuming replacement restores Q, when it restores N and Q must be rebuilt at rate eta S.

25.14.4 A.3 Class C — cost and structure

IF-12 · The coordination cost law and the crossover. paths(n) = n(n−1)/2 [Brooks, 1975]; C(n) = c_0 + c_1 n(n−1)/2; crossover n* the least n with C(n) > eta ln n. *Derivation:* the path count is exact combinatorics; **the linear-in-paths cost model is a postulate, badged.** *Protocol:* measure coordination cost per period across sub-units of differing size in one body and regress on n(n−1)/2 to recover c_1, reporting the fit; **a poor fit means the group is not all-to-all connected**, which is the good case and points to IF-13. *Worked:* the two tables of Section 6.4, with n* = 61 at c_1 = 0.0009. **A twenty-fold difference in per-path cost moves the viable group size from 85 to 16**, which is why "what is the right team size" has no context-free answer. *Failure modes:* assuming a universal team size; measuring c_1 from meeting hours only, which omits the asynchronous coordination cost and is usually the larger part; adding people to a late project, where the added C is immediate and the added S is delayed.

IF-13 · The nesting law. N = g^d, groups = (g^d − 1)/(g − 1), paths = groups x g(g−1)/2, paths per member tending to g/2; flat cost O(N^2), nested O(N). *Derivation:* exact combinatorics on a g-ary tree. **The mapping onto the nested- enterprises design principle is a modelling assumption, badged**: nesting was reported as an empirical regularity among long-enduring institutions [Ostrom, 1990], and the cost arithmetic is the framework's. *Protocol:* map the actual communication structure, not the published chart — the two diverge and the actual one is the one that costs [Conway, 1968]; count paths carrying real traffic; compare with the nominal structure, and **the excess is the cost of the informal cross-links, which is not waste.** *Worked:* the table of Section 6.4. **This is the only entry in the formulary that changes an exponent rather than a coefficient**, and the ceiling eta ln N is untouched, so nesting is a pure reduction of C. **The mandatory caveat:** a g^d tree is exactly the structure identified as wrong for this purpose — **a strict tree amputates the cross-links a semilattice preserves, and those links are where u lives** [Alexander, 1965]. The figures are the tree's cost floor, not a recommendation. *Failure modes:* nesting as a tree; nesting the chart without nesting the traffic; adding depth without limit, since depth is a delay and distortion term appearing in IF-16 and not in C.

IF-14 · The group launch budget. L = max(0, M − C), with L the group's capacity to imagine and pursue an alternative to its present arrangement. *Derivation:* **the claim that collective alternatives are generated only from surplus is a postulate, badged.** The max(0, .) is the model's statement that there is no negative imagination. **G2: DEMOTE** (standing verdict 1) — the subtraction requires an interval scale at minimum and the normalisation into [0, ln N] requires ratio, so **L is reportable as a threshold condition — is the moment above the cost line or not — and not as a magnitude.**

M	C	L	reading
1.20	0.55	0.6500	alternatives available
2.00	0.55	1.4500	ample
0.80	0.90	0.0000	none — fully consumed
0.55	0.55	0.0000	exactly at the line

A body at L = 0 will not propose a change to its own arrangement, and this is not a failure of vision. It is the arithmetic of a group whose whole moment is consumed by its own maintenance; the instruction to leave one's comfort zone reverses the order of operations, since the departure is planned from surplus rather than from its absence. *Failure modes:* demanding innovation from an L = 0 body, where the correct move is to cut C first; reading L = 0 as complacency, when it is capacity, measured; ordinal M.

25.14.5 A.4 Class D — the lens at group scope

IF-15 · The group moment. `M = (D/P) . E`, with `D` shared doctrine, `P` collective perception, `E` the environment; `D` and `P` must share units so that `r = D/P` is dimensionless, and `M` inherits `E`'s units. *Derivation:* **this is the author's formula and it is a postulate, badged.** What is not postulated is the dimensional requirement, which is forced. **G2: DEMOTE** (standing verdict 1) — division requires `D` and `P` ratio-scaled with a true zero on a **common** scale, and no such instrument exists; channel-counting gives a zero within one encoder and no linkage between encoders. **The formula stands as the ordering over-framed / calibrated / under-framed, plus its limits and its fixed point; its magnitudes are withdrawn.** *Protocol:* prediction-surprise ledger; estimate `D` from confidence held and `P` from information actually supplied; **test for cyclicity first (B3) — if the collective ordering is intransitive, `D` is not a well-defined object and this entry does not apply.** *Worked, illustrative construction with `E = 10.0`:*

D	P	r	M	reading
1.0	1.0	1.000	10.0000	calibrated: `M = E`
2.5	1.0	2.500	25.0000	over-framed
3.0	1.0	3.000	30.0000	over-framed, severe
1.0	3.0	0.333	3.3333	under-framed, severe
0.4	1.0	0.400	4.0000	under-framed

Failure modes: ordinal inputs; cyclic `D`, where the numerator does not exist; confusing `M` with `E`, which are equal only at the fixed point and the whole diagnostic concerns the cases where they are not.

IF-16 · Doctrine-to-data and the group distortion. `r = D/P`, `Delta = |ln r|`, with `r > 1` over-framed, `r = 1` calibrated, `r < 1` under-framed. *Derivation:* definition from IF-15. **The log-ratio form is chosen because ratios compose multiplicatively**, which makes over-framing and under-framing by the same factor equal distortions, as the raw ratio does not. **G2: inherits IF-15's DEMOTE. `Delta` is an ordinal quantity pending instrumentation — it ranks bodies and periods by distance from contact and does not license the claim that one distortion is twice another. The sign of `ln r` is unaffected and is the part that decides the correction.** The band table of Section 7.4 is a construction showing how the measure reads, not a set of measurements. A body at `r = 2.90` has `Delta = 1.0647`, severe over-framing; correcting to `r = 1.25` gives `Delta = 0.2231`, drift. **Calibration is not the goal; the achievable goal is drift, and a body reporting `Delta < 0.05` is more likely to be mismeasuring than to be perfect.** The two structural reasons why `r = 1` is not an unqualified target — the exact-zero headroom at the fixed point, and the fact that a body optimising raw `Q` is driven away from `r = 1` by arithmetic — are developed in Section 7.4 and are not repeated. *Failure modes:* correcting the wrong direction, since under-framing treated by adding doctrine becomes over-framing; reporting `Delta` alone, which loses the sign; confusing over-framing with confidence, when the diagnostic is confidence **uncorrelated with accuracy**, not confidence as such.

IF-17 · The group return difference. `Delta_R = |ln(M'/M)|`, with `M` the moment the group sent into the world and `M'` the moment that came back; convertible to octaves by `x 2.885390`. *Derivation:* **the claim that the world's answer over what was sent is the informative quantity is a postulate, badged.** The log form is forced by the same reasoning as IF-16. **G2: DEMOTE. Amplification, null return and attenuation are findings; the number of nats is not.** The octave conversions below are exact arithmetic on an illustrative construction, not measurements. *Protocol:* record `M` at the moment of sending, **before** any response; record `M'` when the response is complete, same instrument; report magnitude **and sign**. *Worked, `M = 10`:*

M'	Delta_R (nats)	octaves	reading
10	0.0000	0.0000	**zero return difference: no information**

M'	Delta_R (nats)	octaves	reading
12	0.1823	0.5261	mild amplification
7	0.3567	1.0291	attenuation
25	0.9163	2.6439	strong amplification
3	1.2040	3.4739	strong attenuation

A response identical to what was sent is an echo, and an echo carries no information about the world. A body whose consultations return exactly its own proposal has learned nothing, and the measurement says so with a zero. *Failure modes:* treating `Delta_R = 0` as success, the commonest misreading, when it means the channel is a mirror; mismatched instruments; ignoring the sign.

25.14.6 A.5 Class E — trust, reciprocity, obligation

IF-18 · Trust as retention reduction. `delta_eff = delta_0 + (1 − a)(1 − v)`. Statement, derivation, protocol, ethics constraint and worked half-life table are those of Sections 4.2, 4.3 and 6.6 and are not repeated. **Derivation: Postulate PT, badged.** The decomposition of lapse into irreducible and confirmation-addressable causes is the substantive content; the postulate is that verification is effective on the second class. **`a` and `v` are separately identifiable only if some commitments are genuinely left unverified: under `v = 1`, `a` is unmeasurable.** The elasticity is `a/(delta_0 + 1 − a)` and **not** `a/(1−a)` (Section 4.4). *Failure modes:* scalar trust; trust from survey, an ordinal self-report prohibited from entering the equation; **maximising `a`**, when `delta` is not a fault to be minimised and high `a` on a commitment that ought to lapse is how obsolete procedure becomes immortal.

IF-19 · The verification budget and the trust threshold. `C_v(v) = c_v v`; `Q_inf(v) = (K − c_v v)/(delta_0 + (1−a)(1−v))` with `K = eta S − C_0`; `sign(dQ_inf/dv) = sign[(1−a)(K − c_v) − c_v delta_0]`, independent of `v`; and `a* = 1 − c_v delta_0/(K − c_v)` valid when `K > c_v`. Derivation, worked table and the linearity caveat are Section 4.5. *Protocol:* estimate `c_v` as the full-coverage verification cost; estimate `delta_0` from an interval with `v = 1`, where the residual loss is irreducible; estimate `K`; compute `a*` and compare against measured `a`. *Failure modes:* verification creep; **monitoring mistaken for trust-building — monitoring lowers `c_v`, which raises `a*`; it does not raise `a`** (Section 6.2).

IF-20 · The reciprocity index and the conversion law. `rho_ij = 2 min(x_ij,x_ji)/(x_ij+x_ji) = 1 − |x_ij − x_ji|/(x_ij + x_ji)`; `rho_group = matched / total`; `eta = eta_max rho`. **The index is a definition; `eta = eta_max rho` is Postulate PR, badged. G5: FAIL as stated — `rho` depends on the accounting window, and the same relation measures 0.1333 at one month and 0.8667 at twenty-four. The window must be preregistered and the sensitivity curve reported** (S4). Protocol is Section 6.6, Protocol B. *Worked:* the tables of Sections 5.2 and 5.3. **A relation at `rho = 0.29` does not accumulate slowly; it fails to accumulate at all** — while fourteen transfers a period look perfectly healthy on a dashboard. *Failure modes:* window shopping; asymmetric units; volume mistaken for reciprocity; ignoring density, since `rho` on a sparse network is a statement about the few edges that exist.

IF-21 · The obligation decomposition and the imbalance index. `X = Sym + Skew`; `b_i = sum_j (x_ij − x_ji)` with `sum_i b_i = 0`; `B = sum_i |b_i|/(2 x total volume)`; and the theorem `1 − B >= rho`. *Derivation:* the decomposition is exact and unique for any square matrix; **the interpretation is Postulate PO, badged**; **the bound is proved** from the triangle inequality on the reals (P8). Protocol is Section 6.6, Protocol C. **Ethics, binding: individual balances shall not be published within the group** (S12). *Worked:* the eight-firm consortium, total directed volume 68, `b_i = [−17, −21, +14, +8, +5, +5, +3, +3]`, so `B = 76/136 = 0.5588`, `1 − B = 0.4412` against a measured `rho = 0.3824`. **The node index flatters the network by 15.4 per cent, systematically, in the direction of comfort.** *Failure modes:* reporting `B` alone; publishing the ledger; reading a large `Sym` as health, when volume and balance are independent.

25.14.7 A.6 Class F — members and thresholds

IF-22 · Member headroom and the exit condition. `h_i = eta ln N / N − delta Q / N − c_i`, with critical private cost `c_i* = (eta ln N − delta Q)/N`; **member `i` is individually unsustainable when `h_i < 0` regardless of the group's `h`.** *Derivation:* IF-09 divided by `N`, with `c_i` replacing the per-capita average. **The equal division of the ceiling and of maintenance is a postulate, badged**; unequal division is the general case and requires member-specific shares. *Protocol:* measure `c_i` per member — their own time, travel, compliance burden, in the group's `Q`-unit; compute `h_i`; **report the names below zero, not just the count**; address by reducing `c_i`, **never by restricting exit — retention is an offer, never a restriction.** *Worked:* the table of Section 8.1, with its disclosed denominator defect. **They leave in cost order — highest `c_i` first — and the highest-cost members are frequently the highest-`u` members, because carrying a divergent framing in a converged room is expensive. Combined with IF-05, the body loses its most divergent members first and therefore loses surplus faster than headcount.** *Failure modes:* aggregate headroom quoted as sufficient; attributing exit to commitment when `h_i < 0` is arithmetic; retention by restriction, prohibited.

IF-23 · Rival versus non-rival return.

```
RIVAL     (store divided):    return_i = eta ln(N_eff)/N  −  c_i
NON-RIVAL (public good):      return_i = eta ln(N_eff)    −  c_i
ln(x)/x is maximised at x = e = 2.71828
```

Derivation: definition of the two allocation regimes applied to the surplus of IF-02. **The maximum of `ln(x)/x` at `x = e` is elementary calculus and is exact.** *Protocol:* **declare rivalry first** (S17) — is the store divided among members or held whole? Tooling access is rival; a published standard is not. **G5: rivalry must be declared from the nature of the store, not inferred from the data, since both regimes fit small-`N` data equally well.** *Worked:* the two tables of Section 8.8, at `eta = 0.35`, `c_i = 0.020`, `u = 1` and outside option `w = 0.030`. **The same group, the same surplus, the same members — and the sign of the size gradient flips entirely depending on whether the store is divided or shared.** *Failure modes:* undeclared rivalry; treating a mixed store as one regime, when most real institutions hold both and the correct move is to split `Q` into two accounts; **reading the `x = e` result as a recommendation**, when it is the rival regime's arithmetic and not a design rule.

IF-24 · The tipping threshold and the cascade. A committed minority near **25 per cent** of `N` was **measured** to overturn an established social convention in a controlled experimental paradigm [Centola et al., 2018]; the threshold-model fixed-point construction, in which the equilibrium active count `x*` is the smallest fixed point of `x -> #{ i : T_i <= x }`, is Granovetter [1978]. **Both are real results; the transfer to any particular body is a modelling assumption, badged. G6: the 25 per cent figure is one well-conducted experimental paradigm, not a constant of nature, and must be stated as a measured tipping point in that paradigm. G7: PARTIAL — the design is clear and published, but broad cross-context replication of the specific value is not established and must not be claimed.** This is the only entry in the formulary whose G7 is better than `OPEN`. *Protocol:* count `N` for the convention in question — the population the convention actually binds; treat `ceil(0.25 N)` as an order-of-magnitude guide, not a target; **estimate the threshold distribution by observing partial adoptions, because the distribution matters more than the mean.**

`N`	8	12	20	25	40	100	150
committed needed at 25%	2	3	5	7	10	25	38

The threshold result, computed. Ten people with thresholds `0,1,2,3,4,5,6,7,8,9` cascade to **10 of 10**. Change one threshold — the person at 1 becomes a 2 — and the equilibrium is **1 of 10**. **One person's threshold moving by one unit changes the outcome from complete adoption to none.** The mean threshold moved by 0.1; the outcome moved by 90 per cent of the population.

This is why aggregate measures of readiness for change have almost no predictive value and why the distribution is the only useful object.

Complex-contagion caveat, mandatory. Behaviours requiring social reinforcement spread differently from simple contagion: **they need multiple reinforcing contacts, and clustered topology beats random long ties** [Centola and Macy, 2007, Centola, 2010]. A change requiring reinforcement will not propagate along the weak-tie shortcuts that carry information [Granovetter, 1973], and a plan built on information diffusion will fail for that reason without the failure being visible. *Failure modes:* quoting 25 per cent as a law; **manufacturing the committed minority**, which is conscripted difference (IF-01) and is prohibited where members cannot decline the role; planning simple diffusion for a reinforcement-dependent change.

IF-25 · The scope transform. `delta_T = 1 - (1 - delta)^T` and its inverse `delta = 1 - (1 - delta_T)^{1/T}`. Period-free quantities: `N`, `u`, `N_eff`, `F`. Period-bearing: `delta`, `eta`, `C`, `h`, and `Q` if it is a flow rather than a stock. *Derivation:* survival compounds multiplicatively; **exact, not an approximation.** *Protocol:* convert **before** comparing any two rates; never compare `delta` values whose periods differ. *Worked:* the table and the 90.1x trap of Section 8.7. *Failure modes:* unnormalised comparison, the commonest cross-scope error; adding rates instead of compounding, which at `delta = 0.10` and `T = 7` gives 0.70 where the truth is 0.5217; applying the transform to a heterogeneous store, where a departure from the prediction is the diagnosis and not a rounding problem.

25.14.8 A.7 The consolidated gate matrix

`P` = PASS, `Pc` = PASS (conditional), `F` = FAIL as stated, `D` = DEMOTE, `O` = OPEN, `Pt` = PARTIAL. **Every non-`P` verdict is spelled out beneath the table.**

Entry	G1	G2	G3	G4	G5	G6	G7
IF-01 pair surplus	P	P	P	P	Pc	Pc	O
IF-02 effective count	P	Pc	P	P	Pc	Pc	O
IF-03 two-source	P	P	P	P	**F**	Pc	O
IF-04 requisite variety	P	Pc	P	P	Pc	P	O
IF-05 departure fragility	P	P	P	P	**F**	P	O
IF-06 storage law	P	Pc	P	P	Pc	P	O
IF-07 exact solution	P	Pc	P	P	P	P	O
IF-08 steady state	P	Pc	P	P	Pc	P	O
IF-09 headroom	P	Pc	P	P	Pc	P	O
IF-10 target increment	P	Pc	P	P	Pc	P	O
IF-11 turnover composition	P	P	P	P	P	P	O
IF-12 coordination cost	P	Pc	P	P	P	P	O
IF-13 nesting	P	P	P	P	P	P	O
IF-14 launch budget	Pc	**D**	P	P	Pc	P	O
IF-15 group moment	Pc	**D**	P	P	Pc	Pc	O
IF-16 doctrine-to-data	P	**D**	P	P	Pc	P	O
IF-17 return difference	Pc	**D**	P	P	Pc	P	O
IF-18 trust as retention	P	P	P	P	Pc	P	O
IF-19 verification budget	P	P	P	P	Pc	P	O
IF-20 reciprocity index	P	Pc	P	P	**F**	P	O
IF-21 obligation decomposition	P	Pc	P	P	Pc	P	O
IF-22 member headroom	P	Pc	P	P	Pc	P	O
IF-23 rival / non-rival	P	Pc	P	P	Pc	P	O
IF-24 tipping threshold	P	P	P	P	Pc	Pc	**Pt**
IF-25 scope transform	P	P	P	P	P	P	O

The three G5 failures. IF-03: many `(u_cross, u_div)` pairs give the same `u`, so the aggregate does not identify the split — **the split must be measured from the block assignment, not inferred from `u`.** IF-05: distinct `{c_i}` vectors give identical mean `u`, so **the aggregate does not identify the fragility, and that is the entry's whole point.** IF-20: `rho` depends on the accounting window, moving from 0.1333 to 0.8667 on identical data, so **the window must be preregistered and the sensitivity curve reported.**

The four G2 demotions are standing verdict 1 and are stated in A.0. **The G2 conditionals** require in every case a ratio-scaled measure with a true zero: `N` enumerable by the independent-action test rather than headcount (IF-02, IF-04, IF-09); `Q` on the substitution test rather than a Likert culture survey (IF-06 through IF-10, IF-22); a ratio-scaled cost measure (IF-12, IF-23); and a **symmetric** ratio-scaled flow unit (IF-20, IF-21) — **if advice is counted one way and money the other, `rho` is undefined.**

The G1 conditionals (IF-14, IF-15, IF-17) all require that `D` and `P`, or `M` and `M'`, be in the same unit, which is the hard requirement because the answer often arrives in a different currency from the question.

The G6 conditionals require a preregistered encoding (IF-01), a shuffled-input control with the calibration of standing verdict 2 (IF-02), an independently specified block structure (IF-03), independent measurement of both terms (IF-15), and the paradigm restriction on the 25 per cent figure (IF-24).

G7 is `OPEN` for twenty-four of twenty-five entries and `PARTIAL` for the twenty-fifth. `OPEN` is a debt, not a pass. Nothing here has been through a cross-lab reproducibility trial [Open Science Collaboration, 2015].

25.15 Appendix B — Specifications and failure modes of formulary use

S23 — Gate row mandatory. No formula shall be used in a report without its gate row reproduced or cited. **An entry quoted without its gates is an entry quoted without its limits.**

S24 — `OPEN` is a debt, not a pass. Twenty-four of twenty-five entries have G7 `OPEN` and the twenty-fifth has it only `PARTIAL`. Any claim of established measurement shall state that.

S25 — Demote on scale failure. Where G2 fails — ordinal data entering a division or a ratio — the formula shall be demoted in print to a heuristic ordering, labelled as such, and its numerical output shall not be published. This governs IF-14, IF-15, IF-16 and IF-17 most often.

S26 — Badge every postulate in the sentence that states it. IF-04, IF-05, IF-11, IF-13, IF-14, IF-15, IF-17, IF-18, IF-20, IF-21, IF-22, IF-23 and IF-24 all rest on at least one postulate. **The modelling is permitted; it is permitted in the open, with its badge on.**

S27 — Report the ceiling and the realised fraction together. Any use of IF-09 shall be accompanied by `S / ln N` from IF-02. **Half of all misdiagnosis turns on confusing the ceiling with the realised value.**

S28 — Preregister windows and partitions. IF-03's block partition and IF-20's accounting window shall be declared before the data are seen, with justification [Nosek et al., 2018].

S29 — Distribution, not scalar. IF-05, IF-11, IF-20 and IF-22 shall be reported with their across-member distributions. A scalar alone is a mean-field claim.

S30 — Consent constraints carried. IF-18's disclosure requirement, IF-21's ledger-privacy requirement, IF-22's no-restriction requirement, and IF-24's no-conscription requirement travel with those entries and shall not be dropped in citation.

S31 — No chaining across triples. IF-03 shall not be composed across three parties. The triangle inequality fails for geometric triples.

S32 — Adding an entry. A new formula enters this formulary only with: a stated derivation identifying it as theorem, definition or postulate; all seven gates with verdicts; a measurement

protocol executable by someone who did not write it; at least one computed worked number; and named failure modes. **A formula without a failure mode has not been thought about.**

F19 — The quoted number. *Signature:* a figure from a worked example reappearing in someone's report as a finding about their own system. *Damages:* every downstream conclusion. *Diagnostic:* ask which of their own measurements produced it; usually none. *Fix:* **every worked number in this paper is an illustration of the arithmetic, not a benchmark. Measure.**

F20 — Gate skipping. *Signature:* an entry applied with its gate row omitted, most often G2. *Damages:* produces confident numbers from ordinal data. *Diagnostic:* ask for the scale type of every term; if the answer is "a one-to-five rating", the division is invalid. *Fix:* S23 and S25.

F21 — Postulate laundering. *Signature:* a badged postulate quoted at second hand without its badge, then treated as following from a theorem. *Damages:* misattributes a modelling choice to a proof, which makes it unarguable. *Diagnostic:* trace the claim to its entry and check the derivation field. *Fix:* S26.

F22 — The single lever. *Signature:* a diagnosis naming one parameter and an intervention on that parameter alone. *Damages:* Section 6.5 computes a case where a single-lever trust intervention **moved the equilibrium the wrong way** while every survey measure improved. *Diagnostic:* compute the one-at-a-time decomposition; if the interaction term exceeds either main effect, single-lever intervention is contraindicated. *Fix:* check the sign of `eta S − C` before any `delta` intervention.

F23 — Ceiling / realised confusion. *Signature:* a body reported as having ample headroom and shrinking every period. *Damages:* sends the engineer to the wrong lever entirely. *Diagnostic:* `h` uses `ln N` and the growth condition uses `ln N_eff`; compute both, always. *Fix:* S27.

F24 — Scalar reporting. *Signature:* one `u`, one `delta`, one `rho` for a group with visibly heterogeneous members. *Damages:* conceals the fragility profile (IF-05) and the exit profile (IF-22), both computable in an afternoon. *Diagnostic:* recompute with each member deleted in turn; a wide spread means the scalar was hiding it. *Fix:* S29.

F25 — Consent stripping. *Signature:* IF-18 used to run hidden reliability tests; IF-21 used to publish a debt ledger; IF-22 used to justify retention clauses; IF-24 used to assemble a committed minority by instruction. *Damages:* violates the declination constraint, and in three of the four cases also invalidates the measurement. *Diagnostic:* ask whether the affected person could have declined and been told they could. *Fix:* S30. **This is the one failure mode with no parameter fix and no acceptable trade-off.**

F26 — Reference material read as argument. *Signature:* this appendix cited as though the density of formulas were itself evidence. *Damages:* the credibility of everything else. *Diagnostic:* count the `OPEN` gates — twenty-four of twenty-five G7s are `OPEN` and the twenty-fifth is only `PARTIAL`. *Fix:* S24. **Formulas do not make a field hard. Ratio-scale, reproducible, falsifiable measurement of the formulas' terms does.**

Chapter 26

The Formula Validation Protocol, Scale Type, and Specific Objectivity: Demoting the Corpus's Own Moment Equation and Pricing Its Restoration

Abstract. This paper specifies a seven-gate protocol — dimensions, scale type, limiting cases, fixed points, identifiability, falsification, reproducibility — and runs it on the corpus's own moment equation $M = (D/P)E$. **The protocol does not pass it.** The equation clears gate 1, names all four limiting cases at gate 3, and its $D = P$ invariant is shown at gate 4 to be static rather than attracting. Gate 5 exhibits two invariance groups, so only the ratio r is identifiable. Gate 6 was run on preregistered simulated data: the frozen log-linear model scored held-out $R^2 = 0.8091$ against a shuffled control at 0.0059, while data generated additively triggered the division kill condition, the coefficient on the denominator returning a 95 per cent interval of $[-0.2009,\ 0.0590]$ against a target of -1. Gate 7 exposed a pooled fiction: a true between-person spread $\tau = 0.2534$ against a pooled standard error of 0.0257, a ratio of 9.9. **Gate 2 is the demotion.** Theorem 7.1 establishes that a quotient is meaningful only on a ratio scale, and four demonstrations show that admissible relabelling moves a reported D/P by 25 per cent, reverses orderings across occasions, produces any ratio one likes via $g(x) = x^k$, and can make the ratio negative or undefined. **The magnitude of D/P, the distortion as a number, and the quantitative moment-to-storage join are therefore withdrawn as measured quantities**, and the same gate demotes five of the nine entries of the corpus's personal formulary. Rasch separability supplies the constructive route, priced at roughly 123 items per term for a half-logit.

Keywords: formula validation; scale type; meaningfulness; magnitude estimation; cross-modality matching; specific objectivity; Rasch separation; kill conditions

Series: B. Engineering the Surplus **Paper:** CCP-026

26.1 1 Introduction

A framework that writes its claims as equations has to answer a question before any of those equations can be assessed: what would have to be true of the quantities in them for the arithmetic to be about anything? Writing `M = (D/P)E` does not make the study of moments a science.

An equation is a claim about relations among *measured* quantities, and if the quantities are not measured on a scale that supports the operations the equation performs on them, the equation is a diagram of an intuition in mathematical clothes.

This paper specifies the instrument the corpus uses to answer that question and then turns the instrument on the corpus's own central formula. The Formula Validation Protocol (FVP) is a seven-gate procedure — dimensions, scale type, limiting cases, fixed points, identifiability, falsification, reproducibility — which every formula in the framework is required to pass in print, with its per-gate verdict stated. **Run on the moment equation `M = (D/P)E`, the protocol returns DEMOTE at gate 2 and NOT YET RUN at gates 6 and 7.** The magnitude of `D/P`, the distortion $\Delta = |\ln r|$ as a number, and the quantitative join between the moment equation and the storage law are withdrawn as measured quantities and relabelled illustrative constructions. The same protocol, applied across the corpus's personal formulary, demotes **five of that formulary's nine entries** on the same gate [Curzi, 2026q].

That result is the paper's principal contribution, and it is deliberately reported at the front. A validation procedure that passes everything it tests is worthless as a certificate that anything passed: **a protocol that nothing fails is measuring nothing**, which is Goodhart's and Campbell's observation applied to a methodology rather than to a policy target [Goodhart, 1975, Campbell, 1979]. The corpus's pass rate under its own protocol is therefore reported here as a headline number rather than as an appendix entry.

The motivation for building such a protocol is external and quantitative. The Open Science Collaboration attempted direct replications of one hundred psychological studies: 97 of the originals reported a statistically significant result and about 36 of the replications did, with effect sizes roughly halved [Open Science Collaboration, 2015]. That is 37 per cent of the originally significant findings recovered, and — because a correlation's explained variance goes as its square — a halving of `r` is a quartering of `r^2`. The diagnosis is not fraudulence. It is a publication filter selecting on something other than truth [Ioannidis, 2005], fed by flexibility in analysis [Simmons et al., 2011, Gelman and Loken, 2013] whose standard remedies are preregistration and held-out prediction [Nosek et al., 2018].

The observation that makes the protocol possible is that one branch of the same discipline does not have this problem. Weber's ratio law has been reproduced continuously since 1834 [Weber, 1834]; Fechner's logarithmic relation since 1860 [Fechner, 1860, 1966]; Stevens' magnitude-estimation exponents replicate across laboratories and decades, and where they vary they vary in known and documented ways with method, range and instruction [Stevens, 1957, Teghtsoonian, 1971, Krueger, 1989]. Psychophysics studies subjective experience — the loudness of a tone, the heaviness of a weight — which is exactly the class of quantity the sceptic says cannot be measured. **The discriminating variable is therefore not subjectivity. It is how the quantity is put on a scale.**

Four things are contributed here relative to the surrounding literature rather than relative to the framework. First, the seven gates are ordered by cost, so that the cheap structural gates run before the expensive empirical ones, and the ordering is defended with the observation that running gate 6 on a formula that fails gate 1 spends a year to learn what an afternoon would have taught. Second, an invariance theorem — that a quotient is meaningful only on a ratio scale — is stated with a three-line proof and then *executed* as four constructive demonstrations on the corpus's own data structures, showing that an admissible relabelling moves a reported `D/P` by 25 per cent, reverses the ordering of two occasions, delivers any ratio one likes, and can make the ratio negative and the distortion undefined. Third, the demotion that follows is stated as a withdrawal of specific quantities, with an explicit list of what survives it. Fourth, the restoration is **priced**: given a badged identification postulate, Rasch person measures supply the true zero that division requires, and the arithmetic of person measurement says how many items that costs — approximately 31 per term to detect a one-logit distortion, 123 for half a logit, and 342 for 0.3.

None of this is a claim about human beings. The empirical portions of gates 6 and 7 are run on **simulated** data throughout, and that label is load-bearing: what is demonstrated is the *protocol* — what the analysis looks like when the formula is true, what it looks like when it is

false, and what the numbers would have to be for a reader to be entitled to believe it.

The framework's static half — the surplus functional `F(u) = ln(1 + (N−1)u)` and its affine effective count — is presupposed rather than re-derived [Curzi, 2026]. The storage law `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t`, its fixed point and its geometric convergence, which gate 4 checks candidate formulas against, are presupposed from Curzi [2026a]. The moment equation itself and its self-imposed demotion to ordinal scale are stated in Curzi [2026m]; the logarithmic register in which the whole corpus is written is established in Curzi [2026n].

Section 2 states the hypotheses. Section 3 states the protocol and what is presupposed, including the identification problem the framework cannot answer. Sections 4 to 6 define the seven gates, their requirements and their failure modes. Sections 7 to 9 develop gate 2 in full, because scale type is where the corpus's own central formula is most exposed. Sections 10 to 12 develop the deep form of gate 7. Section 13 reports the protocol run end to end. Sections 14 to 17 discuss what follows, verify the chain, state the limitations, and conclude.

26.2 2 Hypotheses

Some of what follows are theorems about transformation groups and about a probabilistic model, and they cannot be falsified by observation; they are labelled as derivations where they appear. What can be falsified is the claim that the protocol is *discriminating* — that it separates formulas that measure something from formulas that do not — and the claim that the corpus's own formulas have the status the protocol assigns them. The hypotheses below separate the two, and each names the observation that would defeat it.

H1 — The protocol has teeth. A validation protocol applied to a body of formulas whose author constructed both will return a pass rate materially below one. **Predicted:** applied to the corpus's own formulary, the protocol demotes a substantial fraction of the entries, and the corpus's central formula is among the demoted. **Falsified by:** every formula in the corpus passing every gate. That outcome would not be evidence that the formulas are sound; it would be evidence that the protocol is a compliance exercise, which is failure mode FM-6.8.

H2 — Scale type, not sample size, is what licenses a quotient. For quantities measured on a scale admitting transformations beyond $x \mapsto ax$, the value of a reported ratio is a property of the labelling rather than of the world. **Predicted:** an admissible relabelling of an ordinal instrument changes a reported `D/P` by an arbitrary factor, reverses the ordering of two occasions, and can render the distortion undefined; and no quantity of data alters this, because it is an algebraic property of the transformation group. **Falsified by:** exhibiting a strictly increasing `g` under which `g(X)/g(Y) = X/Y` for all positive `X`, `Y` without `g` being linear through the origin. Section 15 shows this cannot be done, so H2 is in fact a **theorem**, not a hypothesis; it is stated here as a hypothesis only to make the boundary between the two explicit, and it is recorded as `DEFINITION/ARITHMETIC` in the verification chain rather than as an empirical claim.

H3 — A preregistered rival functional form is necessary, and a null is not sufficient. A misspecified functional form can beat every simpler baseline on held-out data. **Predicted:** data generated additively and fitted with the multiplicative log-linear model will still return a respectable held-out `R^2` — high enough to be reported as confirmation against persistence, direct-input and autoregressive baselines — while the correct additive model scores materially higher. **Falsified by:** the misspecified form failing conspicuously against the simple baselines, which would make the rival-form preregistration redundant.

H4 — Pooling across persons manufactures apparent precision. Where a parameter varies between persons, a pooled regression returns an interval far narrower

than the true spread of the individuals it purports to describe. **Predicted:** with true between-person spread τ and pooled standard error s, the ratio τ/s is large and grows with the number of pooled observations, so the pooled analysis looks *more* precise the *more* wrong it is about every individual in it. **Falsified by:** a pooled interval whose width tracks the between-person spread rather than the observation count.

H5 — Invariance dominates fit as a diagnostic. Split-sample parameter invariance detects instrument defects that model-fit statistics miss. **Predicted:** on an instrument with planted defects, the split-sample invariance test flags strictly more of them than the fit statistics do, and at least one defect returns fit statistics in the nominal range. **Falsified by:** fit statistics catching every defect the invariance test catches, which would make gate 7 a footnote to model fit rather than a separate gate.

H1 is confirmed and H3, H4 and H5 are confirmed on simulated data; H2 is a theorem. The confirmations of H3, H4 and H5 are confirmations *about the protocol*, on data whose generating process is known, and they carry no implication whatever about human subjects. That distinction is maintained throughout and is restated in Section 16.

26.3 3 The validation problem and the protocol

The three parts of this paper are the scientific warrant of Series B of the corpus. Everything before them is a construction. Everything after them depends on whether the construction is allowed to stand. If this paper fails, the applied sections of the framework are an elaborate notation for opinions, and the honest thing to do would be to say so and stop.

So it is worth being exact about what is being claimed, and about what the objection actually is.

26.3.1 3.1 The objection, stated fairly

The objection is not that psychology and sociology are pseudoscience. **That charge is false, it is lazy, and this paper does not make it.** It is worth spelling out why, because a great deal of sloppy work — including sloppy work that flatters the corpus's own project — hides behind the sloppiness of the charge.

Psychology and sociology are sciences. They have theories that make predictions, experimental methods, measurement instruments, institutions of criticism, and findings that hold. The contact hypothesis is a real finding, with four stated conditions [Allport, 1954] and a meta-analysis over more than five hundred studies attached to it [Pettigrew and Tropp, 2006]. The collective intelligence factor is a real finding [Woolley et al., 2010]. The tipping point at roughly twenty-five per cent is a measured quantity from a controlled experiment [Centola et al., 2018]. Ostrom's design principles were extracted from field data across decades [Ostrom, 1990]. These are not vibes.

What these fields have is **a measurement problem and a replication problem**, and those are different diseases from fraudulence, with different treatments.

26.3.2 3.2 The replication numbers, stated exactly

The Open Science Collaboration's *Estimating the reproducibility of psychological science* (*Science*, 2015) attempted direct replications of one hundred studies published in three high-ranking psychology journals. The headline result, as recorded in the corpus's canon:

```
originals reporting a statistically significant result:      97 of 100
replications reporting a statistically significant result:   about 36 of 100
replication effect sizes:                                    roughly half the originals
```

Three pieces of arithmetic on those figures, computed:

```
36 / 97   = 0.3711        of the originally significant findings, about 37% replicated
97 - 36   = 61            percentage-point gap between claimed and recovered
halving r: r -> r/2  =>  r^2 -> r^2 / 4       variance explained falls to 25%
```

That last line is the one people miss. **Effect sizes "roughly halved" means variance explained roughly quartered**, because the variance a correlation accounts for goes as its square. An effect that explained 16% of variance in the original explains 4% in the replication, whatever the p-value says. Verified across the plausible range:

original `r`	original `r^2`	halved `r`	halved `r^2`	ratio
0.20	0.0400	0.100	0.0100	0.250
0.30	0.0900	0.150	0.0225	0.250
0.40	0.1600	0.200	0.0400	0.250
0.50	0.2500	0.250	0.0625	0.250

26.3.3 3.3 What the numbers are evidence of, and what they are not

They are **not** evidence that psychology is fake. Some failed replications are themselves false negatives, some are context-sensitivity rather than error, and one replication is not a verdict. They are also not evidence about sociology, which was not the sample.

They are evidence of something narrower and more actionable. **A field in which 97% of published results are significant and 36% of them survive a direct redo is a field whose publication filter is selecting on something other than truth.** That is a process fault, and process faults have process fixes: preregistration, adequate power, held-out prediction, and — the concern of these this paper — better measurement.

Here is the observation that makes the whole of this paper possible, and it is the load-bearing observation of the section. **Psychophysics does not have this problem.** Weber's ratio law has been reproduced continuously since 1834. Fechner's logarithmic relation since 1860. Stevens' magnitude-estimation exponents replicate across laboratories, across decades, and across languages, and where they vary they vary in known and documented ways with method, range and instruction [Stevens, 1957, Teghtsoonian, 1971, Krueger, 1989]. Psychophysics studies *subjective experience* — the loudness of a tone, the heaviness of a weight, the brightness of a light — which is exactly the class of quantity the sceptic says cannot be measured. And it replicates.

So the discriminating variable is not subjectivity. Psychophysics is as subjective as anything in Series B of the corpus. The discriminating variable is **how the quantity is put on a scale.** Psychophysics built ratio scales, with a true zero, using procedures that recover the same exponent from different observers on different days in different rooms. Much of the rest of the field uses seven-point ratings, treats them as though they were metres, and averages them.

26.3.4 3.4 The claim of this paper, stated so it can be checked

Formulas do not make a field hard. Ratio-scale, reproducible, falsifiable measurement of the formulas' terms does.

That sentence is the thesis of this paper, and it cuts against the corpus as much as for it. Writing `M = (D/P)E` does not make the study of moments a science. Nothing about the typographic dignity of an equals sign confers empirical status. An equation is a claim about relations among *measured quantities*, and if the quantities are not measured — really measured, on a scale that supports the operations the equation performs on them — then the equation is a diagram of an intuition wearing mathematical clothes.

This is achievable. Psychophysics achieved it for sensory magnitude, and Rasch measurement achieved a version of it for attainment and attitude (Section 10). The Formula Validation Protocol is the route, and it is deliberately harsh: a protocol the corpus's own formulas cannot fail would be worthless as a certificate that they passed.

26.3.5 3.5 The protocol, badged

Postulate FVP (methodological). **A formula in the corpus is admissible if and only if it passes all seven gates below, in print, with its results stated.** This is a house standard, not a theorem. It is not derived from the foundations papers of the corpus, and no result in the corpus depends on it. It is a rule the author has adopted and by which the author asks to be judged. Its components are drawn from established measurement theory and established replication practice; the choice to require all seven, together, of every formula, is the author's.

```
FVP-1  DIMENSIONS       Units of every term; both sides agree; ratios dimensionless.
FVP-2  SCALE TYPE       Stevens' typology per term. Division and multiplication REQUIRE a
                        ratio scale with a true zero. If only ordinal data exist, say so:
                        the formula is a heuristic ordering, not an equation, and is
                        DEMOTED in print.
FVP-3  LIMITING CASES   Behaviour as each term -> 0 and -> infinity. Every divergence must
                        correspond to a real, nameable phenomenon or the formula is wrong.
FVP-4  FIXED POINTS     Solve for the invariant. State what it means.
FVP-5  IDENTIFIABILITY  Can the parameters be recovered from obtainable data? If two distinct
                        parameter sets predict identical observations, the formula is not
                        testable as written.
FVP-6  FALSIFICATION    Preregister; fit on training; FREEZE; report held-out; beat four
                        baselines (persistence, direct-input, matched autoregressive,
                        SHUFFLED-INPUT control). State in advance what result kills it.
FVP-7  REPRODUCIBILITY  Same operationalisation, different people, different laboratories:
                        do the parameters recover?
```

26.3.6 3.6 Why seven, and why in this order

Each gate kills a distinct class of error, and the classes do not overlap. A formula can be dimensionally perfect and scale-illegal. It can be scale-legal and unidentifiable. It can be identifiable and unfalsifiable. It can pass every one of the first six and still be a private fact about one person's rating habits, which is what FVP-7 is for.

Gate	Class of error it kills	Cost to run	What passing buys
FVP-1	Category error: adding unlike things	Hours	The equation is *about* something coherent
FVP-2	Operation error: arithmetic the data cannot support	Hours to years	The operations performed are meaningful
FVP-3	Structural error: behaviour with no referent	Hours	The formula's extremes describe real states

Gate	Class of error it kills	Cost to run	What passing buys
FVP-4	Vacuity: no invariant, therefore no content	Hours	The formula says something at equilibrium
FVP-5	Aliasing: distinct causes, identical predictions	Days	The parameters are recoverable in principle
FVP-6	Overfitting and post-hoc rescue	Months	The formula beat the alternatives on data it never saw
FVP-7	Idiosyncrasy: a result true of one person only	Years	The formula is about people, not about a person

The order is a cost ordering and it must be respected. Gates 1 through 4 are armchair work: they cost an afternoon and they kill most bad formulas. Gate 5 costs days. Gate 6 costs months. Gate 7 costs years and multiple laboratories. Running gate 6 on a formula that fails gate 1 is a way of spending a year to learn something an afternoon would have taught. In the corpus's experience the ratio is stark: **the cheap gates kill most candidates, and the expensive gates are where the surviving candidates are actually decided.**

26.3.7 3.7 FVP-7 is the gate that decides, and the gate that gets skipped

Every one of the first six gates can be passed by a result true of exactly one person on exactly one instrument. A single subject, densely measured, can produce beautiful held-out prediction that beats all four baselines, and the whole thing can still be an artefact of how that one person uses the word "vivid."

FVP-7 is where the hard-science question is actually settled, and it is the gate the literature most often omits. It is omitted because it is expensive, because it requires coordination with people who do not share your priors, and because — the real reason — it is the gate most likely to return a negative. A framework that has passed gates 1 through 6 feels finished. It is not. The 36% figure is what the world looks like when a field routinely stops at gate 6.

Section 7 is FVP-2 in full, because scale type is where the corpus's own central formula is most at risk. Section 10 is the deep form of FVP-7, because invariance across persons and instruments is what "reproducible measurement" actually means. The falsification suite is FVP-6 in full [Curzi, 2026u]. Sections 3 to 6 define all seven gates, and Section 13.1 runs the complete protocol, end to end, on `M = (D/P)E`.

26.3.8 3.8 What is presupposed

Four results are used here and not re-derived.

The **surplus functional** `F(u) = ln(1 + (N−1)u)`, its axiom set, and the affine effective count `N_eff = 1 + (N−1)u` are established in Curzi [2026] and are used at gate 1 as the worked dimensional check. **That effective count is a linear interpolation between 1 and `N` and is the corpus's own construction; it is not a Hill number of any order.** Where an effective count appears below it is this one, and the Hill numbers proper — `1/HHI` at order 2 and `exp(H)` at order 1 [Hill, 1973, Jost, 2006] — are the different quantities used in the corpus's concentration work [Curzi, 2026], not here.

The **storage law** `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t`, its unique fixed point `Q_inf = (eta S − C)/delta`, and the geometric error law `|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|` are established in Curzi [2026a], whose result that `delta = 0` is a pathology rather than an ideal — no contraction, no unique fixed point, an error-persistence column of ones — is the behaviour gate 4 checks a candidate formula's invariant against. The contraction itself is stated exactly in Curzi [2026]: a contraction with modulus `q < 1` on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least `q` each pass [Banach, 1922, Elaydi, 2005].

The **moment equation** `M = (D/P)E` and its calibration ratio `r = D/P` are stated in Curzi [2026m], together with the demotion rule this paper executes: because the equation contains a division and a multiplication, ordinal data for frame or intake demote it to a heuristic ordering that may not be printed as a fraction or computed.

The **logarithmic register** — that the corpus's quantities are differences of logarithms, reported in nats — is established in Curzi [2026n], whose ratio-invariance test is the licence claimed for a ratio scale there: two fourfold physical steps returned 2.5100 and 2.5000, agreeing to 0.40 per cent and matching the fitted `4^0.6569 = 2.4859`.

The corpus's classification of its own claims into established mathematics, established psychology and declared postulates is Curzi [2026]; every badge used below is a badge in that scheme.

26.3.9 3.9 Identification: what no amount of fitting can recover

Gate 5 asks whether distinct parameter values imply distinct observable predictions. There is a prior identification problem which gate 5 as stated does **not** catch, which is more serious than anything gate 5 does catch, and which this paper cannot answer.

The framework's central formal claim is a reflexive loop `T(x) = W(x, M(x))`: the group influences the member and the members constitute the group. **That is exactly a linear-in-means reflexive specification, and Manski [1993] proves that the endogenous social effect in such a specification is not identified from observational data.** It cannot be separated from contextual effects or from correlated unobservables however much data is collected. Shalizi and Thomas [2011] reinforce the point for networks: influence and selection are generically confounded in observational network data.

The known remedies are randomisation of the treatment, an instrument, partial-population variation, or stipulating the reference group by prior information and defending the stipulation — Brock and Durlauf [2001] being the template for the last and the most realistic route here. **None of them is free, and this paper adopts none of them.** The protocol specified below therefore validates *formulas*, and validation of a formula is not identification of a social effect. Where the framework's reflexive claims are concerned, the honest statement is that they are a **modelling position rather than an identified empirical result**, and the FVP row of any formula resting on the loop should carry that statement at gate 5 in addition to the invariance list.

Two further objections in the same family are recorded here because they bear directly on gate 7. The **Lucas critique** applies to the corpus reflexively: parameters estimated in a population that has not read the storage law cannot be assumed stable in a population that has [Lucas, 1976]. Hacking [1995] states the same problem philosophically — classifications of people interact with the people classified, so the referent moves — and Ferraro et al. [2005] supplies the warning case. Against this, MacKenzie and Millo [2003] documents a reflexive fixed point *observed in the field*, with dates and price data, which is the corpus's own mechanism caught happening; the corpus proposes to give people laws to manage their own comfort by, and is therefore building an engine as well as a camera [MacKenzie, 2006, Callon, 1998]. **The corpus asserts reflexivity for its subject matter and cannot consistently exempt its own parameters from it.** Gate 7 as specified below tests stability across persons and instruments; it does not test stability across *publication*, and no gate in the protocol does.

26.3.10 3.10 Analogy discipline, and what the protocol borrows

The corpus draws on physics and on information theory for its form. Every such borrowing is one of three things and is labelled at first use: a **derivation**, where the system genuinely satisfies the physical premises; a **shared mathematical form**, where the same equation governs both with no claim that the underlying physics applies; or an **analogy**, legitimate as motivation and never as evidence.

Everything this paper borrows is of the second kind, and the second kind is respectable. The dimensional rules R1–R3 of gate 1 are the algebra of dimensional homogeneity and carry no physics. The invariance argument of gate 2 is group theory applied to labelling conventions [Stevens, 1946, Krantz et al., 1971]. The fixed-point language of gate 4 is the elementary theory of first-order difference equations [Elaydi, 2005]. The **entropy** that appears in the treatment of nominal scales is Shannon entropy of a partition, measured in nats or bits [Shannon, 1948], not a thermodynamic quantity, and no thermodynamic consequence is claimed for it. The **logit** of the Rasch model is a log-odds, not a physical potential.

Two limits on the borrowing are recorded now rather than in Limitations. First, the corpus's `Q` is a stock that accumulates, leaks and is spent — the grammar of a conserved quantity — and Mirowski [1989] shows that importing the notation of a conserved quantity without the conservation law is the characteristic failure of exactly this genre. **The corpus owes a statement of what is conserved, in what units, and under what invariance, and it does not have one.** Cartwright [1999] offers the escape route the corpus should take: specify the arrangement within which each law is claimed to hold and stop claiming a universality it does not need. Second, Anderson [1972] is the charter for writing laws at the psychological and social level and, read carefully, an equally strong argument that such laws must be *derived* at that level rather than imported from a lower one.

The functional form of the corpus's surplus law is underdetermined by the evidence, and the protocol below does not repair that. Latané [1981] proposes influence growing as a power of the number of sources with exponent below one; Briscoe et al. [2006] argues for `n log n`; the corpus asserts `ln(1 + (N−1)u)`. All three are decelerating and all three fit the qualitative facts, and Stumpf and Porter [2012] states the general problem: many mechanisms produce indistinguishable curves. The corpus's uniqueness-under-axioms argument establishes that its form follows from *its axioms*; it does not establish that its axioms are the ones the data support. Gate 6's requirement to preregister a rival functional form (KILL-5 below) is the corpus's answer to this in miniature, and Section 13.1.6 shows what it catches. One genuine strengthening is available and should be taken wherever the logarithm is justified: Peters and Gell-Mann [2016] derives a logarithmic form from multiplicative dynamics with **no psychological assumption at all**, which is a stronger warrant than Fechner's because it does not require the perceptual story to be true.

Finally, **no scaling, power-law or heavy-tail claim is made anywhere in this paper.** The one power law that appears — Stevens' `psi = k I^a` in Section 7.10 — is fitted by ordinary least squares in log-log coordinates over a bounded, experimenter-chosen stimulus range with the `R^2` reported, and it is a *response function*, not a distributional tail. Where a distributional claim is made elsewhere in the corpus, the procedure of Clauset et al. [2009] is binding and this paper does not substitute for it.

26.3.11 3.11 Standing honesty clauses

Three clauses govern every number reported below and are stated once here.

Simulated is said, every time. The empirical portions of gates 6 and 7 are run on data whose generating process is known and chosen. They demonstrate the protocol's behaviour under a true model and under two false ones. They are not evidence about people, and no sentence below should be read as such.

Significant figures follow epistemic status. Quantities that are exact by construction — a transformation group, a ratio of integers, the elementary symmetric functions of a stated pa-

rameter vector — are quoted to whatever precision is useful and said to be exact. Quantities estimated from the simulated series are quoted with a standard error or an interval. Quantities that are illustrative model output are rounded and labelled.

A gate not run is reported NOT YET RUN, never blank and never passed. The corpus's own formulas mostly report NOT YET RUN at gates 6 and 7. Printing the honest status is what makes the passed gates worth anything.

26.4 4 The seven gates in detail

26.4.1 4.1 FVP-1 — Dimensions

Dimensional analysis is the cheapest true statement you can make about an equation, and it is astonishing how often it is skipped in the human sciences. It requires no data — only that you write down, beside each symbol, what a unit of that symbol *is*. Three rules, none negotiable:

```
R1  Only like-dimensioned quantities may be added or subtracted.
R2  The argument of ln, exp, sin, or any transcendental function must be dimensionless.
R3  Both sides of an equation must carry identical dimensions.
```

R2 does the most work in this corpus, because this corpus is logarithmic throughout. You cannot take the logarithm of three metres. You can take the logarithm of a *ratio* of metres, and the result is dimensionless — measured in nats, a unit of information, not of length.

Worked check: the surplus law.

```
u = 1 - (x . y)^2 = ||x ^ y||^2        x, y are UNIT vectors -> u dimensionless, in [0,1]
N                                       a count -> dimensionless
1 + (N-1)u                              dimensionless + dimensionless -> OK by R1
F = ln(1 + (N-1)u)                      argument dimensionless -> OK by R2; F in nats
```

Passes. Note what the check enforces: `x` and `y` must be *unit* vectors, or `u` is not dimensionless and the construction fails at R2. That is the reason the surplus functional is defined on normalised states [Curzi, 2026].

Worked check: the storage law.

```
Q_{t+1} = (1 - delta) Q_t + eta S_t - C_t

Q      units of stored order     [q]
S      nats
eta    [q] per nat                       -> eta S has units [q]        OK
delta  dimensionless per period          -> (1-delta) Q has units [q]  OK
C      [q] per period                    -> subtractable from [q]      OK
h = eta ln N - delta Q - C               all three terms in [q]        OK by R1
```

Passes, and the check *earns something*: `eta` is not a pure number. **Eta is a conversion constant with units — the exchange rate between nats of interaction and units of durable structure.** Anyone reporting an eta without saying what `[q]` is has reported half a number. The storage law's conversion constant is named in words in Curzi [2026a]; FVP-1 says it in units.

Worked check: the Moment equation.

```
M = (D / P) . E

For D/P to be dimensionless (R2 is not invoked here, but the interpretation demands it),
```

```
D and P must carry the SAME units.  Then:

r = D/P            dimensionless
M = r E            M inherits E's units exactly

Distortion:  Delta = |ln r|      argument dimensionless -> OK by R2, Delta in nats
```

Passes, **conditionally** — and the condition is the whole content of the gate here. If perspective is measured in "confidence points" and perception in "cues detected per minute", the ratio is not dimensionless, it is a rate with units, and `M` does not inherit `E`'s units. The formula would then be false as written, not merely unmeasured.

FVP-1 therefore imposes a design requirement on the whole of the moment-equation material: D and P must be instrumented on one common scale. That is a hard constraint and it is where most naive operationalisations of the Moment equation die. It is also the constraint that makes cross-modality matching (Section 7, §7.11) relevant rather than ornamental.

26.4.2 4.2 FVP-2 — Scale type

Section 7 is this gate, in full, and nothing here substitutes for it. The summary that belongs in the protocol definition:

```
NOMINAL    identity only        admissible transforms: any 1–1 relabelling
ORDINAL    order               admissible transforms: any strictly increasing g
INTERVAL   order + distance    admissible transforms: x -> a x + b,  a > 0
RATIO      order + distance
           + a true zero       admissible transforms: x -> a x,      a > 0
```

The rule that governs the corpus: division and multiplication are meaningful only on a ratio scale. The reason is exactly the invariance structure above. A statement survives as a claim about the world only if it is unchanged under every transformation the scale admits. On a ratio scale, `g(D)/g(P) = (aD)/(aP) = D/P`: the ratio is invariant, so it carries information. On an interval scale, `(aD+b)/(aP+b)` depends on `b`, which was arbitrary. On an ordinal scale, `g(D)/g(P)` depends on the entire shape of `g`, and `g` was arbitrary.

FVP-2 returns one of three verdicts:

```
PASS       ratio-scaled with a demonstrated true zero -> the equation stands as an equation
DEMOTE     ordinal or interval only -> the formula is printed as a HEURISTIC ORDERING,
           with the demotion stated in the same paragraph as the formula
FAIL       the quantity has no defensible scale at all -> the term is not admissible
```

The DEMOTE verdict is the honest middle, and the corpus uses it. A demoted formula is still useful: it can order occasions, flag directions, and structure attention. What it cannot do is carry a number that is then multiplied by something else. Section 7 §7.13 specifies exactly what survives demotion and what does not.

26.4.3 4.3 FVP-3 — Limiting cases

Take every term to zero and to infinity, one at a time, and read what the formula says. Then ask, for each result, **is there a real phenomenon with that name?**

This gate is a structure test, not a data test, and it is unreasonably effective. Formulas built by analogy almost always have a limit that is nonsense, and the nonsense stays invisible until someone takes the limit.

```
RULE FVP-3a   A divergence (a term going to infinity) must correspond to a real, nameable
              phenomenon. If it does not, the formula is wrong at the singularity, and
              almost always wrong well before it.

RULE FVP-3b   A zero must correspond to a real, nameable phenomenon. "It just goes to zero"
              is not an interpretation.

RULE FVP-3c   Where the limit is a genuine impossibility rather than a phenomenon, the
              formula must state its DOMAIN OF VALIDITY and stop there. Restricting the
              domain is legitimate. Pretending the limit is meaningful is not.
```

The Moment equation is the exemplar and it passes with unusual grace; §13.1 runs the numbers. The storage law also passes, and the passing is instructive: as `delta -> 0` the model gives `Q_t <= Q_0 + t eta ln N` [Curzi, 2026a, Theorem 3.3] — unbounded linear growth with no fixed point. That is not a defect; it is the exact description of a system with no forgetting, which is the subject of the storage law's zero-loss pathology [Curzi, 2026a], and which really does fail to converge. **The limit named a real phenomenon, so the gate passes.**

26.4.4 4.4 FVP-4 — Fixed points

Solve for the invariant. Two distinct objects go by this name and confusing them is a common error:

```
FIXED POINT OF A MAP        x* such that T(x*) = x*.  Dynamic. Banach guarantees existence and
                            convergence when T is a contraction.
                            Example: Q_inf = (eta F(u*) - c*) / delta   [the storage law]

INVARIANT OF AN EQUATION    the parameter configuration at which the equation reduces to an
                            identity. Static. No dynamics required.
                            Example: D = P  =>  M = E
```

Both must be reported. The map's fixed point tells you where iteration arrives; the equation's invariant tells you what the formula is *for*.

A formula with no invariant is usually vacuous. If no configuration of the parameters reduces it to something simpler and interpretable, it is a curve-fit rather than a law — laws are compressions, and compressions have degenerate cases.

FVP-4 also requires that the fixed point's *stability* be stated. `Q_inf` is attracting whenever `delta > 0`, with `|Q_t - Q_inf| = (1-delta)^t |Q_0 - Q_inf|` — geometric convergence, Curzi [2026q]. `D = P` is *not* dynamically attracting in the Moment equation as written, because the Moment equation is not a map: it is an instantaneous relation. Saying so is part of passing the gate. Nothing in `M = (D/P)E` makes `r` drift toward 1. That drift, if it exists, is a separate postulate about learning, and the calibration-ratio paper states it separately and badges it [Curzi, 2026m].

26.4.5 4.5 FVP-5 — Identifiability

A formula is identifiable if distinct parameter values imply distinct observable predictions. If two different parameter sets predict exactly the same data, no amount of data will ever distinguish them, and the formula is not testable as written. This is a property of the *model*, not of the sample size, and it is checked algebraically before anything is measured.

Two kinds:

```
STRUCTURAL identifiability   distinct parameters -> distinct predictions, in principle,
                             with infinite noiseless data. An algebraic property.
```

```
PRACTICAL  identifiability   the parameters can be recovered with useful precision from the
                             data you can actually obtain. A statistical property.
```

Structural failure is fatal and no quantity of data repairs it. Practical failure can sometimes be fixed by better design.

The Moment equation has a structural identifiability problem, and it is severe. `M = (D/P)E` is invariant under two separate transformations, computed here on concrete values:

case	D	P	E	M = (D/P)E
A (reference)	8.0	4.0	3.0	**6.0000**
B (`D`,`P` both scaled by 2.5)	20.0	10.0	3.0	**6.0000**
C (`r` doubled, `E` halved)	16.0	4.0	1.5	**6.0000**
D (`r` quartered, `E` quadrupled)	2.0	4.0	12.0	**6.0000**

Four different worlds. One observation. **From `M` alone, none of `D`, `P`, `E` or even `r` is recoverable.** The first invariance (rows A and B) is benign in a sense — it says the ratio `r` is what matters and the common scale of `D` and `P` is free — but it means you cannot report `D` as a number without an anchor. The second invariance (rows A, C, D) is the dangerous one: **`r` and `E` trade off exactly, so a large distortion with a thin environment is observationally identical to undistorted contact with a rich one.**

This is not a fatal objection but a specification of what the instrument must supply. Three remedies, in ascending order of cost:

```
REMEDY I5-a   ANCHOR the common scale. Fix a standard stimulus and define D and P relative
              to it. This removes the (cD, cP) freedom by convention. Cost: a standard.

REMEDY I5-b   MEASURE E INDEPENDENTLY of M, by an instrument that does not pass through the
              observer. Physical measures of the environment; a second observer; a record.
              This breaks the r-E trade. Cost: an external instrument. THIS ONE IS ESSENTIAL.

REMEDY I5-c   VARY the terms independently across occasions and use the TIME SERIES. If D
              and P and E move independently over time, the log-linear design matrix has
              full column rank and the exponents are recoverable. Cost: repeated measures,
              and the requirement that the terms genuinely decorrelate.
```

Remedy I5-c has a checkable condition, and it is the practical form of the gate. Writing the Moment equation in the logarithmic register [Curzi, 2026n] gives

```
ln M = b0 + bD ln D + bP ln P + bE ln E + noise
```

with the theoretical prediction `(bD, bP, bE) = (1, −1, 1)`. The parameters are recoverable if and only if the matrix of `[1, ln D, ln P, ln E]` has full column rank. **If `D` and `P` co-vary perfectly across occasions — if your frame always tracks your intake — the columns are collinear and no data set of any size will separate `bD` from `bP`.** That is a real risk, not a hypothetical one, because a well-adapted person's frame *does* track their intake. The design must therefore include occasions that break the coupling, and the analysis must report the collinearity diagnostics. A protocol that measures only ordinary days will fail FVP-5 in practice while passing it in principle.

26.4.6 4.6 FVP-6 — Falsification

The template is the preregistration protocol of Curzi [2026i], and it is stated as a sequence because the order is what makes it work:

Preregistration with a frozen model is the open-science literature's own prescription [Nosek et al., 2018, Simmons et al., 2011], and the corpus's one confirmatory application of this sequence to a real system — together with its failure — is Curzi [2026i].

```
1.  PREREGISTER      Write the operationalisation, the model, the exponent targets, the
                     baselines, the held-out split, and THE KILL CRITERIA. Deposit it
                     before collecting data. The kill criteria are the point.
2.  FIT              Estimate parameters on the training partition ONLY.
3.  FREEZE           Write the parameters down. They do not change again. Not for a better
                     fit, not for an outlier, not for a discovered confound.
4.  HELD-OUT         Apply the frozen parameters to data never used in fitting. Report the
                     out-of-sample statistic.
5.  BASELINES        Beat all four:
                        (i)   PERSISTENCE       yhat_t = y_{t-1}
                        (ii)  DIRECT-INPUT      the environment term alone
                        (iii) MATCHED AR        an autoregressive model of the same order
                        (iv)  SHUFFLED-INPUT    the same model, inputs permuted across time
6.  REPORT           Publish the failures alongside the successes, at the same prominence.
```

The shuffled-input control matters most and is most often absent. Every other baseline asks whether your model beats a simpler model. The shuffle asks a sharper question: *does the model's performance depend on the inputs being matched to the right occasions at all?* A model with enough free parameters, fitted to autocorrelated data, can post a respectable held-out score purely from the marginal distributions. Shuffling destroys the pairing while preserving every marginal. If the shuffled model does nearly as well as the real one, the real one learned nothing about the relation.

The kill criteria must be numeric and must be written before the data exist. For the Moment equation, the corpus preregisters:

```
KILL-1   The 95% confidence interval for bP excludes -1.0  ->  the DIVISION is not there.
KILL-2   The 95% confidence interval for bD excludes  1.0  ->  the frame term is misspecified.
KILL-3   Held-out R^2 of the model fails to exceed the best baseline by at least 0.10.
KILL-4   Held-out R^2 of the SHUFFLED-INPUT control exceeds 0.20.
KILL-5   The model is beaten by a purely additive form M = D - P + E on the same held-out set.
```

KILL-5 deserves comment because it is the specific alternative most likely to be true. Nothing about the phenomenology of perception forces a *multiplicative* combination. If frame and intake combine additively, the log-linear model will still fit reasonably — §13.1 computes exactly how reasonably — and only a direct comparison against the additive form exposes it. **Preregistering the rival functional form, not merely a null, is what separates a falsification test from a ritual.**

26.4.7 4.7 FVP-7 — Reproducibility

The gate: **same operationalisation, different people, different laboratories — do the parameters recover?**

There are three possible answers and they are not equally good, but only one of them is fatal:

```
OUTCOME A   UNIVERSAL CONSTANT
            The parameter takes the same value across persons, within sampling error.
            Strongest outcome. Rare. This is what a physical constant looks like.

OUTCOME B   PERSON-LEVEL CONSTANT
```

```
            The parameter varies between persons but is STABLE WITHIN a person across
            occasions and instruments. This is still measurement, and it is the normal
            case for psychological quantities. A person's height varies between persons
            too. What matters is that it is stable and that it predicts.

OUTCOME C   UNSTABLE
            The parameter varies within a person across occasions or instruments by as
            much as it varies between persons. FATAL. There is no quantity here; there
            is a fitting artefact.
```

Outcome B is a pass, and this must be said clearly, because the sceptic's demand for Outcome A is a demand physics itself would fail for most of its quantities. The mass of an object is not a universal constant. What makes mass a measurement is that the same object gives the same mass on different balances in different laboratories. The psychological analogue: the same person gives the same exponent on different instruments on different days — which is the convergent half of a multitrait-multimethod design [Campbell and Fiske, 1959] and the operational form of construct validation [Cronbach and Meehl, 1955]. That is what Section 10 calls *specific objectivity*, and it is the real criterion.

The statistic that separates B from C is the between-person variance component. Given per-person estimates `b_i` with standard errors `se_i`:

```
tau^2  =  Var(b_i)  −  mean(se_i^2)          the true between-person variance
```

If `tau^2` is at or below zero, the parameter is a universal constant to within your precision (Outcome A). If `tau^2` is large and positive *and* the within-person estimates are stable across instruments, Outcome B. If `tau^2` is large and the within-person estimates are *not* stable across instruments, Outcome C, and the framework dies for that term.

The single most dangerous mistake at this gate is pooling. §13.1 computes the exact damage: pooling twelve heterogeneous subjects produced a confidence interval on `bD` about ten times narrower than the actual spread of the individuals' true values. The pooled analysis looked *more* precise the *more* wrong it was about every individual in it. Pooling is not a way of running FVP-7. It is a way of avoiding it.

26.4.8 4.8 The FVP row — the required print format

Every formula in the applied sections of the framework carries this block. It is a specification artefact, not a summary: a reader must be able to check each cell against the section that produced it.

```
FVP ROW  ——  <formula name>
  FVP-1  DIMENSIONS      <units of each term; verdict>
  FVP-2  SCALE TYPE      <scale required per term; scale available; PASS / DEMOTE / FAIL>
  FVP-3  LIMITING CASES  <each limit, its named phenomenon, domain of validity>
  FVP-4  FIXED POINTS    <invariant, meaning, stability>
  FVP-5  IDENTIFIABILITY <invariances; remedies required; rank condition>
  FVP-6  FALSIFICATION   <kill criteria; baselines; status: TESTED / UNTESTED>
  FVP-7  REPRODUCIBILITY <status: OUTCOME A / B / C / NOT YET RUN>
```

A gate that has not been run is reported as NOT YET RUN, never as blank and never as passed. The corpus's own formulas mostly report NOT YET RUN at gates 6 and 7, and the costed programme for changing that is Curzi [2026u]. Printing the honest status is not an embarrassment; it is the only thing that makes the passed gates worth anything.

26.5 5 The validation requirements

Numbered, checkable requirements. A formulary entry satisfying V-1 through V-22 has been validated to the corpus's standard. An entry that does not satisfy them is not a formula in the corpus; it is a note.

V-1 — No formula without an FVP row. The row appears in the same section as the formula, not in an appendix, not in a footnote. A formula printed without its row is a specification defect regardless of how well it performs.

V-2 — Units before symbols. Before any manipulation, tabulate every symbol with its dimension and its unit. Symbols whose unit cannot be stated are not admissible terms.

V-3 — Verify R1, R2 and R3 explicitly. State, in one line each: which additions are like-to-like, which transcendental arguments are dimensionless, and that both sides agree.

V-4 — Declare the common scale for every ratio. Any expression of the form `X/Y` must name the single scale on which both `X` and `Y` are measured, and the instrument that produces it.

V-5 — Classify each term on Stevens' typology, with evidence. The classification must cite the procedure that establishes it — magnitude estimation, cross-modality matching, ratio production, a Rasch calibration — not an assertion. "Treated as ratio" without a procedure is a FAIL, not a PASS.

V-6 — Demote in the same paragraph, not later. Where FVP-2 returns DEMOTE, the demotion appears in the paragraph that introduces the formula, in the same typeface, and the formula is thereafter written with an ordering relation rather than an equals sign wherever the text permits.

V-7 — Take every limit. For each term: the limit to zero, the limit to infinity, and the behaviour at any interior singularity. Report all of them, including the boring ones.

V-8 — Name the phenomenon at every limit. A limit reported without a named phenomenon fails FVP-3. If no phenomenon exists, restrict the domain and say where the formula stops applying.

V-9 — Report the invariant and its stability. State whether the fixed point is a fixed point of a map (and whether the map is a contraction) or an invariant of a relation (and therefore non-dynamic). Do not let the reader assume dynamics that the formula does not contain.

V-10 — Enumerate the invariance group. List every transformation of the parameters that leaves the predictions unchanged. This is FVP-5's core deliverable and it is algebra, not statistics. If the list is non-trivial, state the anchoring convention that removes each freedom.

V-11 — State the rank condition. For any formula to be estimated by regression, state the design condition under which the parameters are separable, and the collinearity diagnostic that will be reported.

V-12 — Preregister before data. Operationalisation, model, exponent targets, split, baselines, and kill criteria, deposited and timestamped. A preregistration written after the first look is not a preregistration; report it as an exploratory analysis and label it so.

V-13 — Freeze and never unfreeze. Parameters estimated on training data are written down before the held-out partition is touched. Any subsequent change to the model converts the study to exploratory and it must be relabelled throughout.

V-14 — All four baselines, always, including the shuffle. Persistence, direct-input, matched autoregressive, shuffled-input. Report all four scores whether or not the model beat them.

V-15 — Kill criteria are numeric and prior. No kill criterion may contain the words "materially", "substantially", or "clearly". A number and a comparison operator, or it is not a criterion.

V-16 — Preregister the rival functional form. At minimum, the additive analogue of any multiplicative formula and the multiplicative analogue of any additive one. Beating a null is not evidence for a specific functional form.

V-17 — Report per-subject parameters, never only the pool. Publish the individual esti-

mates and their standard errors as a table. A pooled estimate may accompany them; it may not replace them.

V-18 — Compute and report `tau^2`. The between-person variance component, corrected for within-person sampling variance, with the FVP-7 outcome letter (A, B or C) stated.

V-19 — Test within-person stability across instruments, not only across occasions. Outcome B requires that the same person's parameter recovers on a *different* operationalisation. Occasion stability alone does not distinguish a measurement from a habit of responding.

V-20 — Report NOT YET RUN honestly. Any gate not executed is marked NOT YET RUN with a one-sentence statement of what running it would require. Blank cells are prohibited.

V-21 — Publish failures at equal prominence. A formula that failed a gate is reported in the section where it would have appeared, with the failure stated in the opening paragraph. Failed formulas are not silently removed between editions.

V-22 — Guard the protocol against itself. No FVP result may become a target [Goodhart, 1975, Campbell, 1979]. Where an FVP statistic is used to select among candidate formulas, the selection must be disclosed and the winning formula must be re-run on fresh data before it is reported as validated.

26.6 6 Failure modes of formula validation

Each entry: name, clinical description, damaged gate, diagnostic signature, detecting measurement, fix.

26.6.1 FM-6.1 · THE DRESSED-UP METAPHOR

Damages: FVP-1 and FVP-3.

An intuition is written with an equals sign and thereby acquires an unearned authority. The symbols have no units, the terms cannot be added, and no limit was ever taken. The formula is a diagram of a feeling.

Diagnostic signature: the author cannot say what a unit of any term is; every term is described qualitatively; taking a limit produces nonsense. A reliable tell: the formula never appears with numbers in it, anywhere, ever.

Detecting measurement: attempt V-2. Tabulate the units. The tabulation either completes or it does not.

Fix: either instrument the terms or rewrite the relation in words. **A well-written sentence is not a lesser thing than a badly-founded equation; it is a more honest one.**

26.6.2 FM-6.2 · THE UNANCHORED RATIO

Damages: FVP-5.

A ratio is reported as a number without any convention fixing the common scale, so the number means nothing across contexts. Two reports of `r = 1.4` from two instruments are not comparable and are nonetheless compared.

Diagnostic signature: `r` values drift systematically between studies, between instruments, or between the same person's sessions on different apparatus, with no substantive account of why.

Detecting measurement: run the same subject on two operationalisations in one sitting and regress one `r` on the other. A slope far from 1 with a high correlation is the signature: the quantity is real, the anchor is missing.

Fix: Remedy I5-a. Adopt and publish a standard stimulus. The choice of standard is arbitrary; the existence of one is not.

26.6.3 FM-6.3 · THE R-E TRADE

Damages: FVP-5, specifically.

The environment term is estimated from the same report that produced the moment, so `r` and `E` absorb each other and the model fits perfectly while identifying nothing. This is the single largest threat to the Moment equation and it is the reason Remedy I5-b is marked essential.

Diagnostic signature: excellent in-sample fit; `bE` and the intercept strongly correlated in the parameter covariance matrix; held-out performance collapses when the environment measure is replaced by an external one.

Detecting measurement: compute the correlation between estimated `ln r` and estimated `ln E` across occasions. Strong negative correlation is the trade appearing directly.

Fix: measure `E` by an instrument that does not pass through the observer. A physical record, a second observer, a logged event stream. Without this, the Moment equation is not testable.

26.6.4 FM-6.4 · THE MISSING SHUFFLE

Damages: FVP-6.

Three baselines are reported and the shuffled-input control is not. The model beats persistence, beats direct-input and beats an AR(1), and none of that establishes that it used the pairing between inputs and outcomes.

Diagnostic signature: the baselines reported are all *simpler models*, none is a *destroyed-pairing control*. The word “shuffle” or “permutation” does not appear in the methods.

Detecting measurement: run it. It costs one line of code and one refit.

Fix: run it, and report the number whatever it is. In the demonstration of §13.1, the real model scored `R^2 = 0.8091` held-out and the shuffled control scored `0.0059`. That gap is what a working model looks like. A gap of `0.81` versus `0.55` would have meant something very different.

26.6.5 FM-6.5 · THE POOLED FICTION

Damages: FVP-7, fatally, while appearing to strengthen it.

Individual subjects are pooled into one regression. The pooled confidence interval is narrow because `n` is large. The narrowness is reported as precision. It is not precision; it is a statement about the *average* of a population, presented as though it were a statement about the *members*.

Diagnostic signature: per-subject estimates are not reported; `n` in the results is the number of observations rather than the number of subjects; the confidence interval is implausibly tight for a psychological quantity.

Detecting measurement: fit every subject separately and compute `tau^2` (V-18). In §13.1’s heterogeneous case, `tau = 0.2534` against a pooled standard error of `0.0257`. **The pooled interval was about ten times narrower than the real spread.**

Fix: V-17 and V-18. Publish the per-subject table. Report `tau`. If the target is a group-level claim, use a model with an explicit random effect and report the variance component, not only the mean.

26.6.6 FM-6.6 · THE SILENT DEMOTION

Damages: FVP-2, and everything downstream of it.

Ordinal data are collected, arithmetic that requires a ratio scale is performed on them, and the scale question is never raised. This is the most common failure in the entire human-sciences literature and it is invisible because it is universal.

Diagnostic signature: seven-point ratings entering a multiplication or a division; means and standard deviations of Likert items [Likert, 1932]; ratios of Likert means reported to two decimal places. The rating instrument is separately exposed to question-wording and focusing-illusion effects [Schwarz, 1999, Kahneman et al., 2006], so the arithmetic is performed on a response whose level is itself unstable.

Detecting measurement: the re-labelling test of Section 7.9. Apply an arbitrary strictly increasing transformation to the raw responses, recompute the result, and see whether the conclusion survives. If it does not, the conclusion was about the labels.

Fix: build a real scale (Section 7.10-7.12) or demote in print (V-6). There is no third option, and the pretence that there is one is what the 36% figure is partly made of.

26.6.7 FM-6.7 · THE EPICYCLE

Damages: FVP-6.

A prediction fails and the model acquires a term. The new term is justified post hoc and is never itself preregistered. Repeat, and the model becomes unfalsifiable by accretion rather than by design.

Diagnostic signature: the model has more free parameters than it had last year; every added parameter was added after a specific failure; there is no held-out test that postdates the last addition.

Detecting measurement: count parameters against the preregistration. Compare held-out performance of the current model to the *original* frozen model on *new* data.

Fix: freeze, and re-run the full FVP-6 suite on fresh data with the amended model preregistered. An amended model is a new model and inherits none of the old one's credit.

26.6.8 FM-6.8 · GOODHART CAPTURE OF THE PROTOCOL

Damages: all seven, by making them a target.

The FVP becomes a compliance exercise. Formulas are constructed to pass the gates rather than tested by them; kill criteria are written loose enough to survive; the shuffle control is run on a version of the data with the pairing partially preserved.

Diagnostic signature: every formula in the corpus passes every gate. This is the tell. **A protocol that nothing fails is measuring nothing** [Goodhart, 1975, Campbell, 1979].

Detecting measurement: the pass rate. In an honest corpus it is well below one.

Fix: V-21 and V-22. Publish failures. Re-run selected formulas on fresh data. And keep the pass rate itself on the dashboard, because a rising pass rate is the earliest signal of capture.

26.7 7 Scale type: why a quotient requires a ratio scale

This is the crux of Series B of the corpus. Every other result in the applied sections of the framework either survives this gate or is demoted by it.

The Moment equation divides. `M = (D/P)E` takes perspective, divides it by perception, and multiplies the result by environment. **Division and multiplication are not universally available operations. They are licensed by the scale type of the quantities involved, and they are**

licensed only on a ratio scale with a true zero. If `D` and `P` are not ratio-scaled, the expression `D/P` does not fail to be accurate — it fails to be *about anything*, because its value can be set to whatever you like by a relabelling that changes nothing in the world.

That is the claim of this section, and it is not a matter of taste. It is a theorem about invariance, and the theorem is proved in three lines.

26.7.1 7.1 Stevens' typology

Stevens [1946] classified measurement scales by **the group of transformations under which the scale's information is preserved.** A scale is not a set of numbers; it is a set of numbers *together with the class of relabellings that leave what the numbers mean unchanged.* That class is the scale.

```
SCALE      WHAT IS PRESERVED           ADMISSIBLE TRANSFORMATIONS g
--------   -------------------------   ---------------------------------------------
NOMINAL    identity / distinctness     any one-to-one map          x -> g(x), g injective
ORDINAL    identity + order            any strictly increasing g   x -> g(x), g' > 0
INTERVAL   + equality of differences   positive affine             x -> a x + b,  a > 0
RATIO      + a true zero               similarity                  x -> a x,      a > 0
ABSOLUTE   + the unit itself           identity only               x -> x
```

Each row's transformation group is a *subgroup* of the row above it. Nominal admits the most relabellings and therefore carries the least information; absolute admits none and carries the most. Moving down the table is buying information by giving up freedom.

26.7.2 7.2 The meaningfulness criterion

A statement made in terms of scale values is meaningful if and only if its truth value is unchanged under every admissible transformation of the scale.

This is the operative definition and it is the whole of this section's logic. It is not a philosophical preference. It follows from what a scale *is*: if `g` is admissible, then `g(x)` and `x` encode the same empirical facts by definition, so any statement true of one and false of the other is a statement about the labelling, not about the world.

The criterion is mechanical to apply. Write the statement. Apply an admissible `g`. See if it survives.

26.7.3 7.3 The division theorem

Theorem 7.1 (the division theorem; designated *Theorem 53.1* in the source manuscript and cited under that designation elsewhere in the corpus)**.** *Let `X` and `Y` be measured on a common scale with admissible transformation group `G`. The statement "`X/Y = k`" is meaningful if and only if `G` is the similarity group `{g(x) = ax : a > 0}` or a subgroup of it — that is, if and only if the scale is ratio or absolute.*

Proof. Three cases.

```
RATIO.     g(x) = a x,  a > 0.
           g(X)/g(Y) = (aX)/(aY) = X/Y.
           The ratio is invariant under every admissible g. MEANINGFUL.

INTERVAL.  g(x) = a x + b,  a > 0.
           g(X)/g(Y) = (aX + b)/(aY + b).
           Take X = 20, Y = 10. Under g = identity the ratio is 2.000.
           Under g(x) = 1.8x + 32 the ratio is 68/50 = 1.360.
```

```
          The value depends on b, which the scale leaves free. NOT MEANINGFUL.

ORDINAL.  g strictly increasing, otherwise arbitrary.
          g(X)/g(Y) can be given ANY value greater than 1 (for X > Y > 0) by a suitable
          choice of g. Demonstrated constructively in section 7.9. NOT MEANINGFUL.
```

The converse direction is immediate: if `g(X)/g(Y) = X/Y` for all admissible `g` and all positive `X, Y`, then taking `Y` fixed shows `g` is linear through the origin, so `G` is contained in the similarity group. QED.

Therefore: `D/P` is meaningful if and only if `D` and `P` are ratio-scaled with a true zero. There is no third possibility, no partial credit, and no way to get there by having a lot of data. A million observations on an ordinal scale is a million ordinal observations.

26.7.4 7.4 What an interval scale does license

The theorem is often misread as saying interval scales are useless. They are not. **On an interval scale, ratios of *differences* are meaningful, and this is exactly the claim the scale supports.** Verified on the values `x = (10, 20, 35, 40)` under four admissible affine transformations:

transformation	values	`(x2−x1)/(x4−x3)`	`x2/x1`
identity	10, 20, 35, 40	**2.000000**	2.000000
`g(x) = 1.8x + 32`	50, 68, 95, 104	**2.000000**	1.360000
`g(x) = 2x − 7`	13, 33, 63, 73	**2.000000**	2.538462
`g(x) = 0.4x + 100`	104, 108, 114, 116	**2.000000**	1.038462

The ratio of differences is rock stable at 2.000000 under every admissible relabelling. The ratio of values takes four different answers. That contrast is the whole distinction, and it is why interval-scaled temperature supports “the rise from Monday to Tuesday was twice the rise from Wednesday to Thursday” but not “Tuesday was twice as warm as Monday.”

26.7.5 7.5 The badge

Note on the status of Stevens’ typology. Stevens’ four-way classification is a **proposal about scientific practice**, not a theorem, and it has been criticised — most sharply by Michell, who argues that psychology largely assumed its attributes were quantitative rather than testing whether they are [Michell, 1997, 1999]. The representational theory of measurement puts the whole subject on axiomatic footing and treats Stevens’ types as special cases [Krantz et al., 1971].

None of that weakens Theorem 7.1, which is an invariance argument and holds within any of these frameworks. Whatever you call the scale, if the admissible relabellings include anything beyond `x -> ax`, ratios of values are not invariant, and a formula that divides is reporting a property of the labelling. The badge is on the *typology*. The theorem stands unbadged.

26.7.6 7.6 The four scales, in operational detail

NOMINAL. Numbers as names. Jersey numbers, diagnostic codes, room identifiers, the numbering of the twelve nodes of the foundations series.

```
Permissible:      counting members of a class; mode; chi-square; contingency tables;
                  entropy of the partition; "same or different"
NOT permissible:  order, mean, median, difference, ratio, correlation
```

A nominal scale still carries real information — the entropy of a partition is a genuine quantity, and the corpus's lineage census uses exactly that [Curzi, 2026]. **You can do serious quantitative work on nominal data; you cannot do arithmetic on the labels.**

ORDINAL. Order without spacing. Rankings, hardness scales, seniority, most rating scales as they are actually used, "how anxious do you feel, 1 to 7."

```
Permissible:      median; percentiles; rank correlation (Spearman, Kendall); order statistics;
                  monotone regression; "greater than"; sign tests
NOT permissible:  mean; standard deviation; difference; ratio; Pearson correlation;
                  any arithmetic on the labels
```

INTERVAL. Equal differences, arbitrary zero. Celsius, Fahrenheit, calendar dates, standardised scores, and — importantly for Section 10 — properly constructed Rasch measures in logits.

```
Permissible:      mean; standard deviation; Pearson correlation; differences; ratios of
                  DIFFERENCES; linear regression; t-tests and ANOVA
NOT permissible:  ratios of VALUES; "twice as much"; multiplication of two measures;
                  proportional statements
```

RATIO. Equal differences plus a non-arbitrary zero denoting genuine absence. Length, mass, duration, count, Kelvin, frequency in hertz, and — the existence proof this volume depends on — Stevens' magnitude estimates of sensory intensity.

```
Permissible:      everything above, plus ratios of values, multiplication, division,
                  geometric mean, coefficient of variation, logarithms of the values,
                  "twice as loud"
```

Note the last two entries. **The geometric mean and the logarithm are ratio-scale operations, and this entire manual runs on both.** `ln N_eff`, `ln M`, `|ln r|`, the octave, the nat: every one of them presupposes ratio-scaled inputs. That is not an incidental fact about the notation. **The logarithmic register of the corpus [Curzi, 2026n] is available only because ratio scales support it, and this section is the price of admission.**

26.7.7 7.7 What a true zero actually is

"True zero" is not "the smallest number on the form". It is a specific empirical claim:

```
A scale has a TRUE ZERO if the value 0 denotes the complete ABSENCE of the attribute, and
that absence is an empirically identifiable state rather than a chosen origin.
```

Three tests, in ascending order of strength:

Test Z1 — describability. Can you describe the state the zero denotes, and would a second person recognise it? Zero mass is describable. Zero degrees Celsius is not a state of absence; it is the freezing point of water, chosen for convenience. **This test is necessary and cheap and almost nobody runs it.**

For the Moment equation's terms, the zeros pass Z1 and they pass it well:

```
P = 0    total absence of intake. Sensory deprivation, dissociation, the closed eye.
D = 0    total absence of frame. Raw sensation with no interpretive structure at all.
E = 0    total absence of environment. No situation to have a moment about.
```

Test Z2 — additivity / concatenation. Is there an operation that concatenates two instances of the attribute such that the scale values add? Two one-kilogram masses on a pan read two kilograms. This is the classical criterion and it is why length, mass and duration are unproblematic. For subjective quantities there is generally no physical concatenation operation, which is precisely Michell's objection [Michell, 1997] and the historical reason psychology's measurement claims were contested; the axiomatic statement of what must hold before a quantity can be said to exist is Krantz et al. [1971].

Test Z3 — ratio consistency. Do judgements of ratio behave consistently? If an observer says `B` is twice `A` and `C` is twice `B`, do they say `C` is four times `A`? Does halving compose with halving to give quartering? **This is the test that replaces concatenation for subjective quantities, and it is the test magnitude estimation is built on.** It is weaker than Z2 but it is empirical, it is checkable, and it either holds or it does not.

A scale that passes Z1 and Z3 is admissible as ratio for the corpus's purposes, with the provenance stated. That is the standard adopted here, it is the standard psychophysics adopted, and it is badged: Z3 is a behavioural consistency criterion, not a concatenation, and a critic who holds out for Z2 is not being unreasonable.

26.7.8 7.8 What goes wrong on an interval scale

Temperature is the canonical demonstration because everyone already knows the answer and almost nobody carries the lesson into their own data.

Celsius pair	`C` ratio	Fahrenheit values	`F` ratio	Kelvin values	`K` ratio
20 / 10	2.000	68.0 / 50.0	1.360	293.15 / 283.15	1.0353
25 / 5	5.000	77.0 / 41.0	1.878	298.15 / 278.15	1.0719
15 / -5	-3.000	59.0 / 23.0	2.565	288.15 / 268.15	1.0746
10 / 0	**undefined**	50.0 / 32.0	1.562	283.15 / 273.15	1.0366

Read the rows across. **The same physical pair of temperatures produces the ratio 2.000, 1.360 and 1.0353 depending on nothing but the labelling convention.** Row three produces a *negative* ratio in Celsius and a positive one in Fahrenheit — the two labellings disagree about the sign. Row four is undefined in Celsius and perfectly finite in Fahrenheit.

Only the Kelvin column is a real quantity, because Kelvin has a true zero and its admissible transformations are similarities alone. **The ratio 1.0353 is a fact about the world. The ratio 2.000 is a fact about Anders Celsius.**

Now transpose this. A questionnaire scored 1 to 7 has its zero at "1", chosen because forms usually start at one. A scale scored 0 to 6 measuring the same thing has its zero at "0". These two instruments differ by `g(x) = x − 1`, a perfectly admissible affine transformation. A `D/P` of `6/3 = 2.0` on the first instrument is `5/2 = 2.5` on the second. **Nothing changed but the printer's choice, and the reported distortion moved by 25%.**

26.7.9 7.9 What goes wrong on an ordinal scale — the demonstrations

Ordinal is worse than interval, and it is worse in a way that surprises people who have absorbed the temperature lesson. On an interval scale, at least the *family* of possible ratios is constrained by a two-parameter group. **On an ordinal scale the admissible transformations form an infinite-dimensional family, and the ratio can be sent anywhere.**

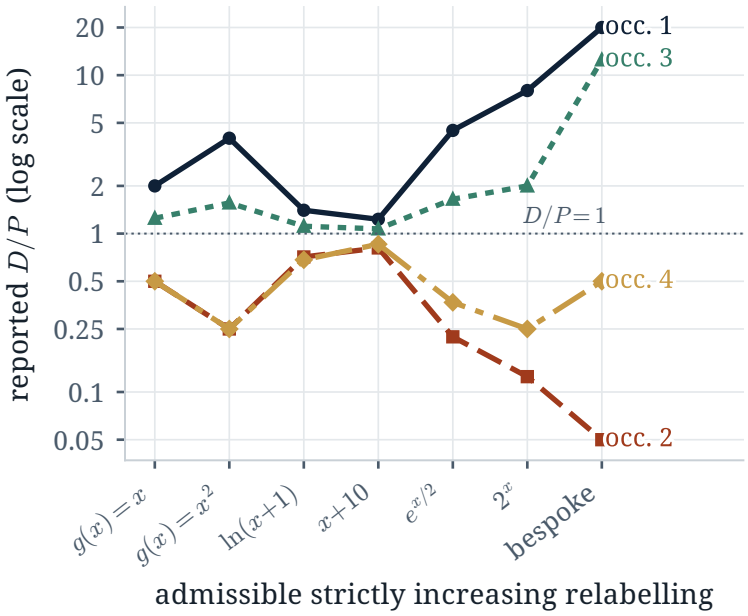


Figure 26.1: **A quantity settable by relabelling is not a measurement:** the same four responses, `D = (6, 3, 5, 2)` and `P = (3, 6, 4, 4)`, report occasion 1's distortion as anything from 1.231 to 20.000 — a factor of 16.25 — under seven strictly increasing relabellings, every one of them admissible on an ordinal scale. Occasions 2 and 4 are tied under two relabellings, and the tie is then broken in *both* directions by others. Values transcribed from the Demonstration 1 table of Section 7.9; generated by `figs/src/CCP-026_fig01.py`.

Demonstration 1 — the magnitude is arbitrary. Four occasions of a frame/intake report on a seven-point instrument, `D = (6, 3, 5, 2)`, `P = (3, 6, 4, 4)`. Apply admissible strictly increasing relabellings and recompute `D/P`:

relabelling `g`	`D/P` across the four occasions	ranking, most over-framed first
`g(x) = x` (raw)	2.000, 0.500, 1.250, 0.500	1, 3, then 2 and 4 TIED
`g(x) = x^2`	4.000, 0.250, 1.562, 0.250	1, 3, then 2 and 4 TIED
`g(x) = ln(x+1)`	1.404, 0.712, 1.113, 0.683	1, 3, **2, 4**
`g(x) = x + 10`	1.231, 0.812, 1.071, 0.857	1, 3, **4, 2**
`g(x) = e^(x/2)`	4.482, 0.223, 1.649, 0.368	1, 3, **4, 2**
`g(x) = 2^x`	8.000, 0.125, 2.000, 0.250	1, 3, **4, 2**
bespoke monotone	20.000, 0.050, 12.500, 0.500	1, 3, **4, 2**

The last column is the alarming one. Under three admissible relabellings occasions 2 and 4 are tied, under one occasion 2 is worse, and under four occasion 4 is worse. **The tie is broken in both directions by relabellings that changed nothing about what the respondent reported.**

Occasion 1's distortion is reported as anything from 1.231 to 20.000 — a factor of sixteen — with no change to a single response. **Every one of these relabellings is admissible. The instrument recorded the same thing in every row.**

Demonstration 2 — the ordering across occasions reverses. Take `D = (7, 4)` and `P = (4, 2)`:

relabelling `g`	`r` occasion 1	`r` occasion 2	verdict
`g(x) = x`	1.7500	2.0000	occasion 2 more over-framed
`g(x) = x^2`	3.0625	4.0000	occasion 2 more over-framed
`g(x) = x^3`	5.3594	8.0000	occasion 2 more over-framed
`g(x) = ln(x+1)`	1.2920	1.4650	occasion 2 more over-framed
`g(x) = 2^x`	8.0000	4.0000	**occasion 1 more over-framed**
`g(x) = e^x`	20.0855	7.3891	**occasion 1 more over-framed**

Which occasion was more distorted is not a fact about the person. It is a fact about the shape of the response labels. This is the finding that matters clinically, because "which session was worse" is exactly the question a practitioner asks.

Demonstration 3 — the constructive result. For `D > P > 0` on an ordinal scale, *any* target ratio can be achieved by an admissible relabelling. With `D = 7`, `P = 4`, use `g(x) = x^k`, which is strictly increasing on the positives for every `k > 0`, giving `g(D)/g(P) = (7/4)^k`:

target ratio	required `k`	achieved
1.05	0.0872	1.0500
1.50	0.7245	1.5000
2.00	1.2386	2.0000
5.00	2.8760	5.0000
10.00	4.1146	10.0000
100.00	8.2292	100.0000

A quantity you can set to any value you like by relabelling is not a measurement. It carries exactly zero bits of information beyond the ordering it started with.

Demonstration 4 — the ratio can become negative or undefined. `g(x) = x − 2.5` is strictly increasing and therefore admissible on an ordinal scale. Applied to the four occasions of Demonstration 1, occasion 4 has `D = 2 -> −0.5` and `P = 4 -> 1.5`, so `r = −0.333` and `Delta = |ln r|` is undefined over the reals. **An admissible relabelling turned the distortion measure into a non-number.** No amount of careful statistics downstream repairs this.

26.7.10 7.10 Building a ratio scale: magnitude estimation

The good news, and it is very good news, is that ratio scales for subjective quantity are not hypothetical. Stevens' method of magnitude estimation produces them, it has produced them since 1957 [Stevens, 1957], and it replicates [Stevens, 1956, Krueger, 1989]. Its failure modes are documented at least as thoroughly as its successes: the ones catalogued in Section 9 appear first in Poulton [1968] and Poulton [1979], and the exponent's dependence on the experimenter's chosen stimulus range, which is the reason for requirement S-8, is Teghtsoonian [1971].

The procedure, specified.

```
ME-1   Present stimuli from the continuum in IRREGULAR order, spanning at least 1.5 orders
       of magnitude. A narrow range biases the exponent toward 1.
ME-2   Instruct: "Assign any positive number to the first stimulus. Assign numbers to the
       rest so that the numbers stand in the same RATIO as the sensations. If one seems
       three times as strong, give it three times the number."
ME-3   Impose NO upper bound and NO fixed modulus. The observer's own first number is the
       modulus and it is free.
ME-4   Normalise each observer by their own geometric mean across stimuli. This removes the
```

```
       arbitrary modulus without touching any ratio.
ME-5   Aggregate across observers by the GEOMETRIC mean, never the arithmetic mean.
ME-6   Regress log(response) on log(intensity). The slope is the exponent a.
ME-7   Report the exponent with its interval, the log-log R^2, and the per-observer spread.
```

The prediction being tested is Stevens' power law:

```
psi = k I^a       ->       log psi = log k + a log I
```

A straight line in log-log is not a curve-fitting convenience. **It is the falsifiable content of the claim: if the observer's judgements were not ratio-consistent, the log-log plot would bend, and it does not.**

Worked recovery. Ten simulated observers, seven intensities spanning `1` to `64`, generated with a true exponent of `0.67`, independent free moduli (the largest observer's modulus was `31.15` times the smallest), and 18% log-normal judgement noise.

```
geometric mean across observers, by intensity:
   0.2300   0.4124   0.5753   0.9469   1.7176   2.6619   4.2328

log-log fit on the geometric means:
   a_hat = 0.6987   s.e. 0.0165   95% CI [0.6664, 0.7310]      true a = 0.67
   log-log R^2 = 0.99723
```

The exponent is recovered, and the log-log linearity is essentially perfect. Per-observer exponents, each fitted separately:

```
0.737  0.569  0.694  0.653  0.755  0.774  0.720  0.631  0.728  0.726
mean 0.6987   sd 0.0631
```

Note what this is: FVP-7 Outcome B in miniature. The exponent is a person-level parameter with a real spread of about 0.06, not a single universal number, and that is the normal and acceptable state of a psychological measurement.

Why the geometric mean is mandatory (ME-5). With free moduli, the arithmetic mean is dominated by whichever observer happened to pick large numbers. In the run above, one observer alone contributed **27.5%** of the arithmetic mean at the lowest intensity and **21.1%** at the highest, in a panel of ten where each should carry 10%. Under the geometric mean each observer contributes exactly `1/n` of the log, by construction. **The arithmetic mean of magnitude estimates silently weights observers by their choice of modulus, which is the one thing the method declares arbitrary.**

26.7.11 7.11 The ceiling: why a rating scale cannot do this

Here is the computation that should end the practice of treating slider responses as ratio data. The same magnitude-estimation simulation — same true exponent `0.67`, same observers, same noise — run four times, changing only the response format:

response format	recovered exponent	bias
unbounded, free modulus (proper ME)	0.6678	-0.0022
0-100 slider	0.5763	-0.0937
0-50 slider	0.4076	-0.2624
1-7 rating	**0.0114**	**-0.6586**

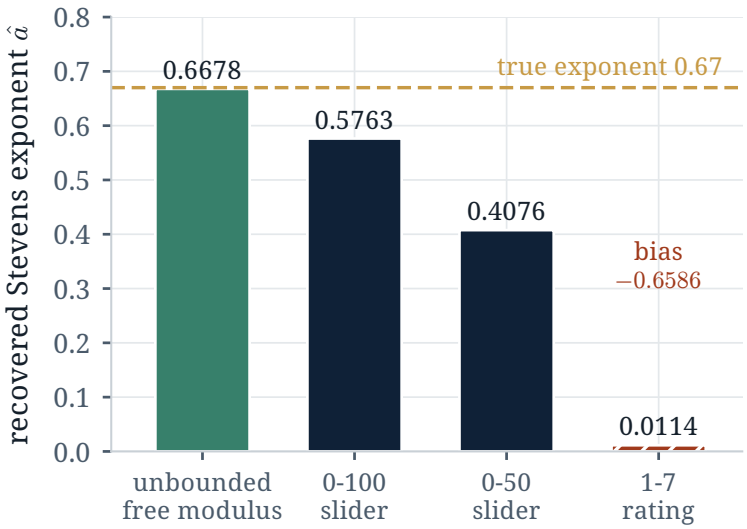


Figure 26.2: **A bounded response format does not degrade the ratio scale, it destroys it:** the same simulated magnitude-estimation panel — ten observers, seven intensities, free moduli, 18 per cent log-normal noise, true exponent 0.67 — recovers 0.6678 unbounded, 0.5763 on a 0–100 slider, 0.4076 on a 0–50 slider and **0.0114** on a 1–7 rating, a bias of −0.6586. Modelled with the parameters stated in Section 7.10; generated by `figs/src/CCP-026_fig02.py`.

A seven-point rating scale recovered an exponent of 0.011 from a process whose true exponent is 0.67. The measured quantity was destroyed almost entirely — not degraded, destroyed. The bound compresses the top of the range, the compression is worst where the signal is strongest, and the log-log slope collapses.

This is not an argument that rating scales are useless. They are excellent ordinal instruments. **It is an argument that a bounded response format cannot produce a ratio scale, ever, and that arithmetic performed on its output is arithmetic on the bound.** A ratio scale must be unbounded above, because doubling must always be possible; a scale with a ceiling has a largest number, and "twice that" has no representation.

26.7.12 7.12 Cross-modality matching — the validation that makes it science

Magnitude estimation on its own could, in principle, be measuring how people use numbers rather than how they experience intensity. **Cross-modality matching is the test that rules this out, and it is the strongest piece of evidence in the whole psychophysical corpus.**

The procedure uses no numbers at all. The observer adjusts one continuum — squeezes a dynamometer, turns a light up — until it *feels* as intense as a stimulus presented on a different continuum. The exponent of that matching function is then a prediction with **no free parameters**:

```
If  psi_1 = k_1 I_1^a   and   psi_2 = k_2 I_2^b,
then matching psi_1 = psi_2 gives  I_2 proportional to I_1^(a/b).

The MATCHING EXPONENT is a/b -- fully determined by two exponents measured in two
separate experiments that never involved matching.
```

The full prediction matrix from the canonical psychophysical exponents [Stevens, 1957, 1956]. Rows are the continuum the observer *adjusts*; columns are the continuum *presented*; entries are the predicted matching exponent `a/b`:

adjusted presented	Loudness 0.67	Bright 0.33	Length 1.00	Salt 1.30	Heavy 1.45	Warmth 1.60	Shock 3.50
Loudness (0.67)	1.0000	0.4925	1.4925	1.9403	2.1642	2.3881	5.2239
Brightness (0.33)	2.0303	1.0000	3.0303	3.9394	4.3939	4.8485	10.6061
Visual length (1.00)	0.6700	0.3300	1.0000	1.3000	1.4500	1.6000	3.5000
Taste, salt (1.30)	0.5154	0.2538	0.7692	1.0000	1.1154	1.2308	2.6923
Heaviness (1.45)	0.4621	0.2276	0.6897	0.8966	1.0000	1.1034	2.4138
Warmth (1.60)	0.4188	0.2062	0.6250	0.8125	0.9062	1.0000	2.1875
Electric shock (3.50)	0.1914	0.0943	0.2857	0.3714	0.4143	0.4571	1.0000

Every cell is a testable prediction. And the matrix carries an internal consistency requirement that is itself a falsification test — **transitivity**:

```
match loudness -> heaviness   = 0.4621
match heaviness -> brightness = 4.3939
product                       = 2.0303
direct loudness -> brightness = 2.0303       AGREES

match brightness -> warmth    = 0.2062
match warmth -> loudness      = 2.3881
product                       = 0.4925
direct brightness -> loudness = 0.4925       AGREES
```

This is what a ratio scale looks like from outside. Three experiments, no shared parameters, and the third one's answer is forced by the first two. When that holds, the scale is not a story about how people use numbers; it is a measurement, because a measurement is exactly a set of numbers whose relations compose.

An honest note is required here: the canonical exponents are nominal, they vary with method, range and instruction, and they are properly reported as *constants of a procedure* rather than constants of nature. Empirically, cross-modality matching exponents recover the predicted `a/b` well but not perfectly, and the residual discrepancies are a known and studied feature of the literature. **The transitivity test is a real test that can fail, which is what makes passing it worth something.**

26.7.13 7.13 Ratio production — the third instrument

Magnitude estimation asks for a number given a stimulus. **Ratio production inverts it: give the observer a number and a standard, and have them produce the stimulus.** "Adjust this until it is half as bright." "Set this until it is three times as heavy."

Production is the strongest of the three because it is the direct behavioural test of Z3, and it composes:

```
RP-1   halve, then halve again -> should equal a single quartering
RP-2   double, then double again -> should equal a single quadrupling
RP-3   triple, then take a third -> should return to the standard
RP-4   the production function and the estimation function should be inverse,
       giving exponents that are reciprocals
```

RP-4 is the sharpest check and it is a genuine falsification opportunity: estimation and production are known to yield systematically slightly different exponents — the "regression effect", in which observers under-use whichever variable they control, catalogued with the rest of the method's biases by Poulton [1979]. **Report both, report the discrepancy, and take the geometric mean of the two exponents if a single figure is needed.** Concealing the discrepancy would be the error; it is a known property of the method and it is documented.

26.7.14 7.14 The other route: additive conjoint measurement

There is a second path to metric structure and it deserves a full hearing, because it is the path that rescues ordinal data — and it is the path Section 10 walks.

Luce and Tukey [1964] proved a remarkable thing. **Suppose you have only an ORDINAL ordering of outcomes, but the outcomes are produced by two factors crossed with each other. If that ordinal table satisfies certain purely ordinal axioms, then interval scales for both factors exist and are essentially unique.** The metric is not assumed. It is *derived from order*, and its existence is a testable consequence of the ordering.

The central axiom is **double cancellation**:

```
if   (a1, b2) >= (a2, b1)   and   (a2, b3) >= (a3, b2)
then (a1, b3) >= (a3, b1)
```

Every symbol is an ordinal comparison. Nothing is added or multiplied. Verified on two 3x3 tables, all 36 antecedent-satisfying triples tested in each:

```
Table A = [[1, 2, 4], [2, 4, 8], [4, 8, 16]]     (a_i * b_j: additive in logs)
   36 triples tested, 0 violations       -> additive structure available

Table B = [[1, 2, 4], [2, 9, 3], [4, 3, 16]]     (non-additive)
   36 triples tested, 3 violations       -> additive structure REFUTED
```

This is the honest answer to "but all I have is ordinal data". You may still have metric structure — but you have to earn it by testing the cancellation conditions, not by assuming them. Double cancellation, together with solvability and an Archimedean condition, is what licenses the interval scale. **Michell's objection [Michell, 1997] is answered not by assertion but by running the test [Luce and Tukey, 1964, Krantz et al., 1971].**

The deterministic form above is brittle: one noisy observation produces a violation. The probabilistic form of the same idea — which tolerates noise, and which is the reason Section 10 exists — is the Rasch model. **Rasch measurement is additive conjoint measurement with error, and the person-by-item table is the two-factor crossing.** That is the connection between these two two treatments and it is not decorative: it is the same theorem, made usable.

26.7.15 7.15 The honest demotion — what to do when you only have ordinal data

Most of the time, in most real settings, you will have ordinal data. This section specifies exactly what you may then do, and it is the section a practitioner needs most.

What SURVIVES demotion. Given `D` and `P` recorded on the *same* ordinal scale with the *same* labels:

```
SURVIVES-1   The SIGN of the comparison. If D > P on the raw labels, then g(D) > g(P) for
             every strictly increasing g. So "this occasion was over-framed" is an
             admissible ordinal claim, PROVIDED both terms use one instrument.
SURVIVES-2   Within-occasion direction of change over time, if the instrument is fixed.
             "The gap narrowed between March and June" is ordinal and admissible.
SURVIVES-3   Rank correlations. Spearman and Kendall are invariant under monotone
             relabelling by construction.
SURVIVES-4   Medians, percentiles, order statistics, sign tests, monotone regression.
SURVIVES-5   The QUALITATIVE structure of the formula: the limiting cases, the fixed point
             D = P as a describable condition, and the three-regime classification
             (over-framed / calibrated / under-framed). These are order facts.
```

What DOES NOT survive.

```
DIES-1   The magnitude of r. Demonstration 3: settable to anything.
DIES-2   The ordering of r ACROSS occasions. Demonstration 2: reverses.
DIES-3   Delta = |ln r| as a number. Demonstration 4: can become undefined.
DIES-4   Any arithmetic that feeds r into another formula. In particular, M as an input
         to the self-storage law Q_{t+1} = (1-delta)Q_t + eta M_t - C_t is NOT available
         from ordinal D and P. The join between the two central formulas requires
         ratio scales.
DIES-5   Means and standard deviations of the raw responses.
DIES-6   Pearson correlations of the raw responses.
```

DIES-4 is the expensive one. **The join between the Moment equation and the storage law [Curzi, 2026m] is the architectural keystone of Series B of the corpus, and it is available only on ratio scales.** That is stated here, in the section that establishes it, and not buried.

The Spearman/Pearson demonstration, on forty paired ordinal responses, computed under four admissible relabellings:

relabelling	Pearson `r`	Spearman `rho`
identity	0.6474	**0.6692**
`g(z) = z^2`	0.5167	**0.6692**
`g(z) = e^(z/2)`	0.4370	**0.6692**
`g(z) = ln(z+1)`	0.7223	**0.6692**

Pearson moves from 0.44 to 0.72 — a range wide enough to change any conclusion — while Spearman is identical to four decimal places in every row, as it must be. **If your data are ordinal, the rank statistic is not a conservative choice. It is the only correct one.**

The print convention (FVP-2 DEMOTE, per V-6). A demoted formula is printed with an ordering relation and an explicit statement, in the same paragraph:

```
DEMOTED FORM:

    Over-framed  <=>  D > P          Calibrated  <=>  D = P
    Under-framed <=>  D < P

    On ordinal D and P the classification above is admissible. The magnitude D/P is NOT,
    and no value of D/P from this instrument may be multiplied, averaged, or entered into
    any downstream formula.
```

Where an interval scale is wanted and respondents cannot be asked to use numbers at all, Thurstone [1927] is the classical alternative route and it predates all of this.

This is not a defeat. It is a considerable amount of usable structure — three named regimes, a fixed point, a direction of correction — obtained honestly and stated in a way a critic can check. What it is not is an equation, and it is not printed as one.

26.8 8 Requirements for scale construction

S-1 — Classify every term before any arithmetic. For each symbol, record: the scale type, the instrument that produced it, and the transformation group that instrument admits. A term whose group cannot be stated is not admissible.

S-2 — Apply the meaningfulness test explicitly. For every statement the analysis will make, apply an admissible `g` and check invariance. Record the check. This costs minutes.

S-3 — Never divide or multiply below ratio. No exception, no “approximately ratio”, no “close enough for practical purposes”. Theorem 7.1 has no epsilon in it.

S-4 — Establish the true zero by Z1 and Z3, and say which. Describe the state the zero denotes in one sentence. Report the Z3 ratio-consistency check with its numbers. If Z2 concatenation is available, use it and say so; if not, say that too.

S-5 — Both terms of a ratio on one instrument. `D` and `P` must come from the same scale with the same anchor. Two instruments, however similar, do not license division of one by the other without a demonstrated common calibration.

S-6 — Response formats for ratio data must be unbounded above. No sliders, no fixed maxima, no “1 to 10”. The ceiling table in §7.11 is the reason: a seven-point format recovered `0.011` where the truth was `0.67`.

S-7 — Free modulus, always. The observer chooses their own first number. Anchoring the first response imposes a scale and destroys the ratio property being measured.

S-8 — Span at least 1.5 orders of magnitude in the stimulus. Narrow ranges bias exponents toward 1 and can manufacture apparent linearity out of nothing.

S-9 — Irregular presentation order. Monotone sequences invite response strategies and produce sequential dependencies that masquerade as the function being measured.

S-10 — Aggregate by geometric mean. Never the arithmetic mean of magnitude estimates. §7.10 quantifies the damage: one observer of ten carrying 27.5% of the weight.

S-11 — Report the log-log `R^2`. Linearity in log-log is the falsifiable content of the power law. A curved log-log plot is a finding and must be reported as one, not smoothed.

S-12 — Report per-observer exponents and their spread. The panel mean is not the result. The distribution is the result, and it determines the FVP-7 outcome letter.

S-13 — Validate by cross-modality matching. Where two continua have independently measured exponents, run the matching experiment and compare against the predicted `a/b`. This is the test that distinguishes measurement from number-usage habit.

S-14 — Run at least one transitivity triangle. Three continua, three matching experiments, and the composition check of §7.12. Report the agreement, or the failure.

S-15 — Run ratio production as a cross-check. RP-1 through RP-4. Report the estimation and production exponents separately and report their discrepancy; do not silently average.

S-16 — Where only ordinal data exist, test double cancellation before assuming metric structure. Report the number of antecedent-satisfying triples tested and the number of violations. A table with violations does not support an additive structure.

S-17 — Demote in print, in the same paragraph. Requirement V-6 of Section 5. The demotion appears with the formula, in the same typeface, not in a footnote and not in an appendix.

S-18 — Use rank statistics on ordinal data. Spearman or Kendall, never Pearson; median, never mean. §7.15 quantifies the alternative.

S-19 — Never feed a demoted quantity into another formula. DIES-4. A heuristic ordering cannot be an input to the storage law, cannot be averaged across people, and cannot be differenced across time as a number.

S-20 — State the scale provenance beside every reported number. The form is: `r = 2.14 [ratio scale, magnitude estimation, free modulus, geometric-mean aggregation, n = 14 observers, a_hat = 0.71 (0.04)]`. A number without provenance is not admissible in the corpus.

26.9 9 Failure modes of scale treatment

26.9.1 FM-9.1 · THE LIKERT METRE

Damages: FVP-2, and everything downstream.

Seven-point responses [Likert, 1932] are treated as though the intervals between them were equal and the bottom of the scale were a true zero. Means are computed, ratios are taken, and the results are reported to two decimal places.

Diagnostic signature: a mean of 4.73 on a 1-7 scale, reported without a distributional summary; a ratio of two such means; a standard deviation of ordinal responses; any Pearson correlation on raw Likert items.

Detecting measurement: the relabelling test. Recompute the headline result under `g(x) = x^2` and `g(x) = 2^x`. If the conclusion moves, it was a conclusion about the labels.

Fix: build a real scale (§7.10-7.13), earn an interval scale by conjoint or Rasch analysis (§7.14, Section 10), or demote (§7.15). **These are the only three options and the third one is respectable.**

26.9.2 FM-9.2 · THE ARBITRARY ZERO

Damages: FVP-2 specifically at the ratio boundary.

An interval scale is treated as ratio because its numbers happen to be positive. Positivity is not a true zero. Fahrenheit is positive across most of the inhabited world and is not a ratio scale.

Diagnostic signature: the author cannot describe the state that "0" denotes; the zero coincides suspiciously with the bottom of a form, the start of a study, or a sample mean.

Detecting measurement: Test Z1. Ask what the zero *is*. Then shift the origin by one unit and see whether any reported ratio survives.

Fix: locate a describable zero and re-anchor, or restrict all claims to differences and ratios of differences (§7.4), which interval scales genuinely support.

26.9.3 FM-9.3 · THE MODULUS LEAK

Damages: the aggregation step of a magnitude-estimation study.

Magnitude estimates are averaged arithmetically. The observers who chose big numbers dominate. The group exponent becomes a weighted statistic where the weights are the very quantity the method declared arbitrary.

Diagnostic signature: group means much larger than the median observer's values; the group exponent shifting when a single observer is removed; a modulus range across observers exceeding a factor of five with no normalisation step reported.

Detecting measurement: compute each observer's share of the group mean at each stimulus level. Under correct aggregation each carries `1/n`. In §7.10's run, arithmetic aggregation gave one of ten observers 27.5%.

Fix: S-7 and S-10. Normalise by the observer's own geometric mean, then aggregate geometrically.

26.9.4 FM-9.4 · THE CEILING

Damages: FVP-2, invisibly and catastrophically.

The response format has an upper bound. The upper part of the stimulus range compresses against it. The recovered exponent is biased toward zero, and the bias grows as the bound tightens.

Diagnostic signature: responses piling up at the maximum; a log-log plot that is straight at low intensities and bends flat at high ones; a recovered exponent implausibly close to zero.

Detecting measurement: the proportion of responses at the maximum. Any nonzero proportion is disqualifying for ratio claims. Compare the exponent fitted on the lower half of the range to the exponent fitted on the whole.

Fix: S-6. Remove the bound. If the bound cannot be removed — some field instruments cannot — the data are ordinal and must be demoted. **The 1-7 row of the table in §7.11 is what happens otherwise: the true exponent 0.67 recovered as 0.011.**

26.9.5 FM-9.5 · THE SILENT PROMOTION

Damages: all of FVP-2, and the reader's ability to check.

Ordinal data are promoted to ratio by fiat somewhere in the middle of an analysis pipeline, usually at the point where a mean is taken, and the promotion is never stated.

Diagnostic signature: the methods section describes a rating scale; the results section reports ratios; no sentence connects the two. The word "assume" does not appear.

Detecting measurement: trace every reported number back to its instrument. The promotion happens at a specific line and it can be found.

Fix: S-1 and S-20. Provenance beside every number. If the promotion is deliberate, badge it as a postulate with its justification and its consequences, under the corpus's badging convention [Curzi, 2026].

26.9.6 FM-9.6 · THE CROSS-INSTRUMENT DIVISION

Damages: FVP-1 and FVP-2 jointly.

`D` is measured by one instrument and `P` by another, and the two are divided. Even if both are ratio scales, their units differ, so the quotient is a rate with units rather than a dimensionless calibration ratio, and the fixed point `D = P` has no meaning.

Diagnostic signature: the two terms come from different sections of a protocol, different response formats, or different modalities; `r` clusters implausibly far from 1 across all subjects, which is the units mismatch showing.

Detecting measurement: the distribution of `r` across a healthy sample. A systematic offset in `ln r` common to everyone is a unit conversion, not a population-wide distortion.

Fix: S-5. One instrument, one anchor, or an explicit cross-modality calibration (§7.12) converting one to the other with the conversion exponent reported.

26.9.7 FM-9.7 · THE UNTESTED CANCELLATION

Damages: FVP-2 via the conjoint route.

Additive conjoint measurement or a Rasch analysis is invoked to justify treating ordinal data as metric, and the cancellation conditions are never actually tested. The theorem is cited; its hypothesis is not checked.

Diagnostic signature: "we used a Rasch model, so the measures are interval" with no fit statistics, no invariance test, no violation count.

Detecting measurement: S-16 for the deterministic case; the fit statistics of Section 10 and split-sample invariance test for the probabilistic one.

Fix: run the test. **A conjoint or Rasch analysis that does not report its fit is a promotion by fiat wearing a better suit.**

26.10 10 Specific objectivity: invariance as the criterion of measurement

Section 7 established what a ratio scale is and why `D/P` needs one. It ended with a construction — magnitude estimation, cross-modality matching, ratio production — and with the honest

verdict that the construction has not been carried out for perspective and perception.

Section 10 develops the second route, and it is the one that answers the harder question. A ratio scale built by magnitude estimation is still a scale built by *one instrument on one occasion*. FVP-7 asks something more: **does the measure survive a change of instrument and a change of population?** That is the real criterion of hard measurement, and there is a body of theory that addresses it directly, achieves it, and can be checked.

26.10.1 10.1 The requirement, stated by Rasch

Rasch [1960] set out the condition he called **specific objectivity**, developed into a working measurement practice by Andrich [1988]:

```
The comparison of two persons must be independent of WHICH ITEMS were used to compare them.
The comparison of two items must be independent of WHICH PERSONS were used to compare them.
```

Read those two lines slowly, because they are more demanding than they look. The first says that if instrument A and instrument B both measure the same attribute, then A's verdict on whether Alice exceeds Bob must agree with B's — not approximately, not on average, but as a structural property of the model, holding for every pair of persons and every pair of items.

This is the property physical measurement has and psychological measurement usually does not. A kitchen scale and a laboratory balance agree about which of two objects is heavier, and they agree about *by how much* once the units are matched. They agree regardless of what else was weighed that day. Two questionnaires purporting to measure the same construct routinely disagree, and the disagreement is normally reported as "modest convergent validity" — the language of Campbell and Fiske [1959] and Cronbach and Meehl [1955] — rather than as the failure of measurement it is.

Specific objectivity is precisely the property Series B needs, because the internal/external correspondence [Curzi, 2026s] claims that the same parameters describe a mind and a city. That claim is empty unless the parameters mean the same thing when measured by different instruments at different scopes. **Invariance is not a nice-to-have property of a good measure. Invariance is what "measure" means.**

26.10.2 10.2 The model

For a dichotomous response `X_ni` of person `n` to item `i`:

```
                     exp(theta_n − beta_i)
P(X_ni = 1) = ---------------------------------
                1 + exp(theta_n − beta_i)

theta_n   the person's location on the latent continuum   (logits)
beta_i    the item's location on the same continuum       (logits)
```

Equivalently, and more revealingly:

```
     P(X_ni = 1)
ln  -------------  =  theta_n − beta_i
     P(X_ni = 0)
```

The log-odds of success is the DIFFERENCE between a person parameter and an item parameter, and nothing else. No item discrimination, no guessing parameter, no interaction term. That austerity is the entire point: it is what buys separability, and every parameter added to the model to improve fit destroys it.

Note also where this places us. `theta − beta` is a *difference on a logarithmic register*, which is the register Curzi [2026n] identifies as the one everything in this corpus lives in. The Rasch model was derived in the 1950s from a requirement about invariance, with no reference to surplus, pitch, or Fechner. **It arrived at the same register by an independent route, and §10.17 shows that the arrival is not a coincidence.**

26.10.3 10.3 Item separability (Theorem 10.1)

The difference in log-odds between two items is independent of the person parameter.

```
ln[P_ni/(1−P_ni)] − ln[P_nj/(1−P_nj)]  =  (theta_n − beta_i) − (theta_n − beta_j)
                                       =  beta_j − beta_i
```

`theta_n` cancels identically. Not approximately — identically, for every person. Verified numerically with `beta_i = −0.80` and `beta_j = +0.60`:

theta	P_i	P_j	logit P_i	logit P_j	difference	odds ratio
-3.0	0.0998	0.0266	-2.2000	-3.6000	**+1.4000**	4.0552
-2.0	0.2315	0.0691	-1.2000	-2.6000	**+1.4000**	4.0552
-1.0	0.4502	0.1680	-0.2000	-1.6000	**+1.4000**	4.0552
0.0	0.6900	0.3543	+0.8000	-0.6000	**+1.4000**	4.0552
1.0	0.8581	0.5987	+1.8000	+0.4000	**+1.4000**	4.0552
2.0	0.9427	0.8022	+2.8000	+1.4000	**+1.4000**	4.0552
3.0	0.9781	0.9168	+3.8000	+2.4000	**+1.4000**	4.0552

The success probabilities range from 3% to 98%. **The comparison between the items is 1.4000 in every row.** Ask a struggling person or an expert; the items stand in the same relation.

26.10.4 10.4 Person separability (Theorem 10.2)

The difference in log-odds between two persons is independent of the item parameter. By the same cancellation, with `theta_1 = 1.10` and `theta_2 = −0.30`:

item beta	logit P_1 - logit P_2
-2.0	**+1.4000**
-1.0	**+1.4000**
0.0	**+1.4000**
+1.0	**+1.4000**
+2.0	**+1.4000**

`theta_1 − theta_2 = 1.4000`, whichever item you use. **This is the half of specific objectivity that matters for Series B of the corpus: a person's measure does not depend on which items were used.**

26.10.5 10.5 The badge, and the methodological inversion

Note on status. The Rasch model is **a model**, and the theorems above are theorems *about the model*, not about people. They say: *if* responses follow this model, *then* separability holds. They say nothing about whether any real response set does follow it.

This produces the methodological inversion that defines Rasch measurement and separates it from mainstream item response theory:

```
ITEM RESPONSE THEORY:   choose the model that fits the data.
                        Add parameters until fit is acceptable.
                        Goal: describe the data well.

RASCH MEASUREMENT:      require the data to fit the model.
                        Where they do not, revise the INSTRUMENT.
                        Goal: obtain measurement, or find out you cannot.
```

Neither position is foolish, and the choice between them is a choice about what you are trying to do. If you want the best description of a data set, fit the richest model. If you want a *measuring instrument* — something whose readings transfer between contexts — you must insist on the model that has the invariance property, and treat misfit as a defect of the instrument rather than as a licence to elaborate the model. Theorem 7.1 is the reason: the properties that make numbers usable are properties of the *scale*, and they are lost the moment the model allows item-by-person interaction.

This paper takes the Rasch position, and badges it as a position. The reason is stated plainly in §10.14 and demonstrated with numbers: models with free discrimination parameters fit better and do not measure.

26.10.6 10.6 The connection to Section 7

Section 7.14 introduced additive conjoint measurement (Luce & Tukey 1964): a purely ordinal table on two crossed factors, satisfying double cancellation, licenses interval scales for both factors. The deterministic form is brittle — a single noisy response produces a violation.

The Rasch model is the probabilistic form of exactly that structure. Persons and items are the two crossed factors; the response table is the outcome; `theta_n − beta_i` is the additive combination; the logistic function supplies the error model that deterministic cancellation lacks. Fit statistics (§10.16) are the stochastic replacement for counting cancellation violations.

Rasch measurement is additive conjoint measurement with error, and that is why it delivers interval scales from ordinal data honestly rather than by fiat. It is the answer to Michell's objection, and it is the answer to a practitioner who has only ratings to work with.

26.10.7 10.7 Sufficiency — why separation is possible at all

Separability is not luck. It follows from a property of the model's likelihood, and the property is worth deriving because it is the mechanism.

For one person answering `k` items, with `xi = exp(theta)` and `eps_i = exp(−beta_i)`:

```
                       prod_i  (xi eps_i)^{x_i}          xi^r * prod_i eps_i^{x_i}
P(x_1..x_k | theta) = -------------------------  =  -------------------------------
                       prod_i  (1 + xi eps_i)           prod_i (1 + xi eps_i)

where r = sum_i x_i  is the RAW SCORE.
```

The person parameter enters only through `xi^r`. The response pattern enters only through `prod eps_i^{x_i}`. The likelihood factorises, and by the factorisation criterion **the raw score `r` is a sufficient statistic for `theta`.**

The consequence is the machinery of the whole method. Condition on `r`:

```
                           prod_i eps_i^{x_i}
P(x_1..x_k | r)  =  --------------------------
                             gamma_r(eps)
```

where `gamma_r` is the **elementary symmetric function** of order `r` in the `eps_i` — the sum of all products of `r` distinct `eps` values. **The person parameter has vanished completely.** The conditional distribution of the response pattern given the raw score depends on the item parameters alone.

Verified numerically. Five items with `beta = (−1.2, −0.5, 0.0, 0.6, 1.4)`, a person at `theta = 0.4`, four different response patterns all scoring `r = 3`:

pattern	r	P(pattern | theta)	P(pattern | r)
1 1 1 0 0	3	0.142349	0.375236
1 0 1 1 0	3	0.047384	0.124905
0 1 1 1 0	3	0.023530	0.062026
1 1 0 1 0	3	0.078123	0.205934

The middle column depends on `theta`. **The right-hand column does not — it is computable from the item parameters alone, and it is the quantity conditional estimation maximises.**

26.10.8 10.8 The gamma recursion

`gamma_r` looks expensive — it is a sum over all `C(k, r)` subsets — but there is a recursion that computes the whole vector in `O(k^2)`. Add items one at a time; each new item either contributes to the product or does not:

```
gamma_r^(m)  =  gamma_r^(m−1)  +  eps_m * gamma_{r−1}^(m−1)
```

Worked in full for the five items above, `eps = (3.32012, 1.64872, 1.00000, 0.54881, 0.24660)`:

```
after item 1:  1.00000  3.32012
after item 2:  1.00000  4.96884   5.47395
after item 3:  1.00000  5.96884  10.44279   5.47395
after item 4:  1.00000  6.51765  13.71855  11.20507  3.00417
after item 5:  1.00000  6.76425  15.32579  14.58802  5.76730  0.74082

check:  gamma_1 = sum(eps)  = 6.76425     OK
        gamma_5 = prod(eps) = 0.74082     OK
```

The two checks at the bottom are the ones to run in any implementation: the first-order symmetric function is the sum, the top-order one is the product. **Numerical instability in these functions is the commonest bug in Rasch software and the two checks catch most of it.**

26.10.9 10.9 Conditional maximum likelihood

Maximise the product of conditional probabilities over persons:

```
                  N   [    prod_i eps_i^{x_ni}   ]
ln L_C  =  ln  prod   [ ------------------------ ]
                 n=1  [      gamma_{r_n}(eps)   ]

        =  − sum_i  s_i beta_i  −  sum_n ln gamma_{r_n}(eps)

where s_i = sum_n x_ni  is the ITEM SCORE.
```

Three properties of this estimator matter and should be stated:

```
CML-1   No person parameters appear. There are N of them and they are ALL eliminated.
        This is why CML is consistent as N grows -- unlike joint estimation, where the
        number of nuisance parameters grows with the sample and the estimator is biased.

CML-2   The item score s_i is a sufficient statistic for beta_i. Symmetric to the person
        side, as specific objectivity requires.

CML-3   Persons with extreme scores (r = 0 or r = k) contribute nothing: gamma_0 and
        gamma_k are constants and the conditional likelihood is degenerate. They are
        EXCLUDED, and the exclusion must be reported.
```

The atom of the whole procedure is visible in the two-item case. Among persons who got exactly one of items `i` and `j` right:

$$P(\text{item } i \text{ right} \mid r = 1) = \frac{\exp(-\beta_i)}{\exp(-\beta_i) + \exp(-\beta_j)} = 0.802184$$

with $\beta_i = -0.80$, $\beta_j = +0.60$. Log odds $= \beta_j - \beta_i = 1.4000$.

No `theta` appears anywhere in that expression. The relative difficulty of two items is estimated from the persons who split them, and how able those persons are does not enter. That is specific objectivity as an estimation procedure rather than as a theorem.

26.10.10 10.10 Calibration — a worked recovery

Twelve items with true difficulties evenly spaced from `-2.00` to `+2.00` logits (centred), 2000 simulated persons drawn `Normal(0, 1.2)`, 1945 usable after excluding extreme scores. CML estimates:

item	true `beta`	CML `beta`	error
1	-2.0000	-1.8690	+0.1310
2	-1.6364	-1.5544	+0.0820
3	-1.2727	-1.3236	-0.0509
4	-0.9091	-0.9090	+0.0001
5	-0.5455	-0.6309	-0.0854
6	-0.1818	-0.1432	+0.0386
7	+0.1818	+0.1336	-0.0482
8	+0.5455	+0.5250	-0.0204
9	+0.9091	+0.9072	-0.0019
10	+1.2727	+1.4011	+0.1284
11	+1.6364	+1.5664	-0.0699
12	+2.0000	+1.8967	-0.1033

```
RMS error = 0.0764 logits      max |error| = 0.1310 logits
correlation(true, estimated) = 0.998382
```

26.10.11 10.11 Separation and reliability — how much measurement you actually have

Estimation accuracy is not the same as measurement precision, and Rasch practice reports both.

```
person SE_n  =  1 / sqrt( sum_i P_ni (1 - P_ni) )
item   SE_i  =  1 / sqrt( sum_n P_ni (1 - P_ni) )

true SD    = sqrt( observed SD^2  -  mean SE^2 )
SEPARATION index G = true SD / RMSE
RELIABILITY        = G^2 / (1 + G^2)
distinct STRATA    = (4G + 1) / 3
```

For the calibration above:

```
observed SD of person measures  = 1.3675 logits
RMSE of person measures         = 0.7576 logits
true SD                         = 1.1385 logits
person separation G             = 1.5028
person reliability              = 0.6931
distinct strata                 = 2.34

item separation                 = 22.0854
item reliability                = 0.9980
mean item SE                    = 0.0578 logits
```

Read the last line of the person block: 2.34 strata. Twelve items, two thousand people, and the instrument distinguishes roughly two and a third levels of the attribute. It can tell high from low and not much else. The item side is superb — 0.0578 logits of item error — because two thousand people calibrated twelve items. **Item precision and person precision are bought by different resources: items are calibrated by persons, persons are measured by items.**

This is the number practitioners most often fail to report, and it is the number that determines whether the instrument can do the job. **An instrument with a person reliability of 0.6931 cannot support individual decisions.** It can support group comparisons. Saying which is a professional obligation. The separation index is the informative statistic here and a coefficient alpha is not [Cronbach, 1951]; the reasons not to stop at alpha are set out by Sijtsma [2009].

26.10.12 10.12 The invariance test — running FVP-7 on an instrument

The theorems in §10.3 and §10.4 are about the model. The test is about the data. **Split the persons into disjoint groups that differ substantially on the attribute, calibrate the items separately in each, and compare.**

The person sample above was split at the median ability (lower half mean `theta = −0.949`, upper half mean `theta = +0.872` — a separation of 1.82 logits, so the two groups are genuinely different populations):

item	`beta` (low sample)	`beta` (high sample)	difference
1	-1.8684	-1.8714	-0.0030
2	-1.4895	-1.7360	-0.2465
3	-1.3595	-1.2439	+0.1156
4	-0.9382	-0.8568	+0.0814
5	-0.7165	-0.5056	+0.2109
6	-0.1480	-0.1363	+0.0118
7	+0.1505	+0.1203	-0.0302
8	+0.5995	+0.4763	-0.1232
9	+1.0241	+0.8481	-0.1760

item	beta (low sample)	beta (high sample)	difference
10	+1.4439	+1.3856	-0.0582
11	+1.5202	+1.5849	+0.0647
12	+1.7820	+1.9349	+0.1529

```
mean difference = +0.0000    sd of difference = 0.1360    max |difference| = 0.2465
correlation(low, high) = 0.994399
```

Two disjoint populations, differing by 1.82 logits in ability, produced item calibrations correlating at 0.9944 with a mean difference of exactly zero and no difference exceeding a quarter of a logit. That is what specific objectivity looks like in data. **This single table is the strongest form of evidence available in psychometrics, and it is the exact analogue of FVP-7.**

The formal statistic is a per-item invariance `t`:

```
t_i  =  ( beta_i^{high} − beta_i^{low} )  /  sqrt( SE_i^{high 2} + SE_i^{low 2} )
```

Flag at `|t| > 2.58` (roughly a 1% two-sided criterion, appropriate under multiple testing across a short instrument). §10.18 runs it on data with planted defects.

26.10.13 10.13 What breaks separability — the two-parameter model

The two-parameter logistic model — due to Birnbaum, and given its standard statement in Lord and Novick's *Statistical Theories of Mental Test Scores* of 1968, neither of which is present in this corpus's verified external bibliography and neither of which is therefore given a citation key here — adds a discrimination parameter per item:

```
P(X_ni = 1) = exp( a_i (theta_n − beta_i) ) / [ 1 + exp( a_i (theta_n − beta_i) ) ]
```

It fits real data better. It fits real data better because it has more parameters. **And it destroys specific objectivity outright**, which is not a subtle statistical cost but a total loss of the property that made the numbers measurements.

The item log-odds difference under 2PL:

```
a_i (theta − beta_i)  −  a_j (theta − beta_j)  =  (a_i − a_j) theta  +  (a_j beta_j − a_i beta_i)
```

The `theta` term does not cancel unless `a_i = a_j`. The same two items as §10.3, now with `a_i = 1.8` and `a_j = 0.6`:

theta	P_i	P_j	logit P_i	logit P_j	difference	odds ratio
-3.0	0.0187	0.1034	-3.9600	-2.1600	**-1.8000**	0.1653
-2.0	0.1034	0.1736	-2.1600	-1.5600	**-0.6000**	0.5488
-1.0	0.4110	0.2769	-0.3600	-0.9600	**+0.6000**	1.8221
0.0	0.8085	0.4110	+1.4400	-0.3600	**+1.8000**	6.0496
1.0	0.9623	0.5597	+3.2400	+0.2400	**+3.0000**	20.0855
2.0	0.9936	0.6985	+5.0400	+0.8400	**+4.2000**	66.6863
3.0	0.9989	0.8085	+6.8400	+1.4400	**+5.4000**	221.4064

The difference sweeps from `−1.80` to `+5.40` and **changes sign at**

```
theta* = (a_i beta_i − a_j beta_j) / (a_i − a_j) = −1.5000
```

Below `theta = −1.5`, item `i` is harder. Above it, item `j` is harder. Which item is more difficult is not a property of the items. It depends on who you ask. The odds ratio moves by a factor of over 1300 across the range.

This is the whole argument in one table. **A model that fits better and cannot answer "which item is harder" without naming a person has not measured the items.** The 2PL is a perfectly good descriptive model and a perfectly bad measuring instrument, and Series B needs a measuring instrument.

26.10.14 10.14 The 2PL under the invariance test

The theoretical point becomes an empirical diagnostic. Data were generated from a 2PL — items 1-6 with `a = 1.8`, items 7-12 with `a = 0.6`, same difficulties as before — and then put through exactly the Rasch calibration and split-sample procedure of §10.12:

item	`beta` (low)	`beta` (high)	difference	true `a`
1	-2.2508	-3.5274	-1.2766	1.8
2	-1.6428	-2.3178	-0.6750	1.8
3	-1.1397	-1.9441	-0.8043	1.8
4	-0.5026	-1.3115	-0.8089	1.8
5	+0.0621	-0.6818	-0.7438	1.8
6	+0.7035	-0.0222	-0.7257	1.8
7	+0.2261	+0.9978	+0.7717	0.6
8	+0.4828	+1.3075	+0.8248	0.6
9	+0.7755	+1.4943	+0.7189	0.6
10	+0.8497	+1.8437	+0.9940	0.6
11	+1.1991	+2.0240	+0.8249	0.6
12	+1.2371	+2.1374	+0.9002	0.6

```
mean |difference|  =  0.8391   (Rasch data: 0.1062)   ->  7.9 times larger
max  |difference|  =  1.2766   (Rasch data: 0.2465)
```

Look at the sign column. Every high-discrimination item drifted negative. Every low-discrimination item drifted positive. Not one exception in twelve. **The drift is not noise; it is a signature, and the signature reads out the discrimination structure directly.** The test does not merely detect the violation, it tells you what kind of violation it is.

That is the practical value of the invariance test and the reason it belongs in every FVP-7 report. A global fit index would have said "acceptable". The split-sample table says "your items have heterogeneous discriminations, the high ones are items 1 through 6, and your scale is not a scale."

26.10.15 10.15 Fit statistics — what misfit tells you

The residual of a single response, standardised:

```
z_ni = ( x_ni − P_ni ) / sqrt( P_ni (1 − P_ni) )     mean 0, variance 1 under the model
```

Two summaries, per item:

```
                1
OUTFIT MNSQ  =  ---  sum_n  z_ni^2                    UNWEIGHTED mean square
                N
```

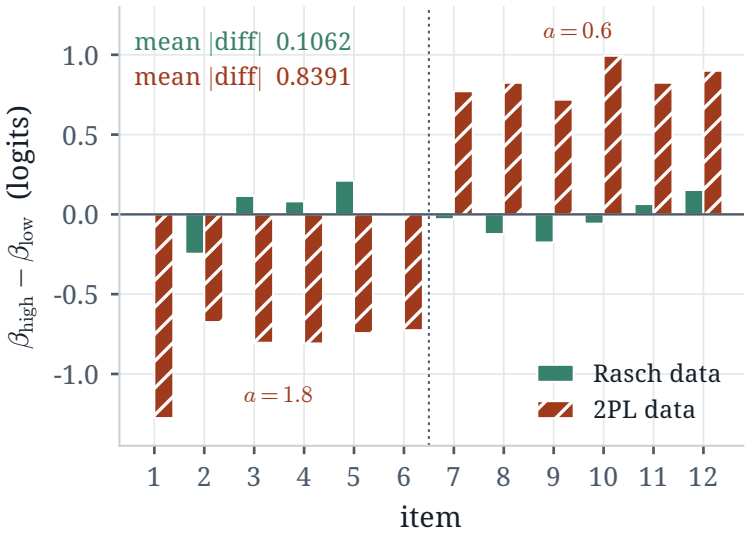


Figure 26.3: **Invariance is a criterion and fit is only a screen:** item calibrations from two disjoint person samples split at the median. On Rasch-generated data the mean absolute difference is 0.1062 logits with no difference exceeding a quarter of a logit; on 2PL-generated data it is 0.8391, 7.9 times larger, and the sign of the drift reads out the discrimination structure directly — every one of items 1–6 (`a = 1.8`) negative and every one of items 7–12 (`a = 0.6`) positive, without exception. Both data sets are simulated with the item parameters stated in Sections 10.12 and 10.14; generated by `figs/src/CCP-026_fig03.py`.

```
                  sum_n ( x_ni - P_ni )^2
INFIT MNSQ    =  -------------------------          INFORMATION-WEIGHTED mean square
                  sum_n  P_ni (1 - P_ni)
```

Both have expectation 1 under the model. The difference between them is which responses dominate:

```
OUTFIT is OUTLIER-SENSITIVE. Its denominator P(1-P) is tiny for far-off-target responses,
       so a single unexpected response from a badly mismatched person can move it a lot.
       Use it to detect: guessing, careless errors, miskeys, lucky hits.

INFIT is INLIER-SENSITIVE. It weights each residual by the information that response
       carries, so on-target responses dominate.
       Use it to detect: problems with the items people are actually well matched to --
       the failures that matter for the measure.
```

Interpretation of the mean square, with the two directions carrying entirely different diagnoses:

```
MNSQ >> 1     UNDERFIT -- more randomness than the model predicts.
              Off-construct item, ambiguous wording, guessing, careless responding,
              multidimensionality. THIS IS THE DANGEROUS DIRECTION: it degrades measurement.

MNSQ << 1     OVERFIT -- LESS randomness than the model predicts.
              Responses too predictable: item dependency, redundancy, a Guttman-like
              over-discriminating item. Does not degrade measurement, but it INFLATES
              reliability and makes the instrument look better than it is.
```

The asymmetry is the practically important part and it is routinely inverted by newcomers, who treat MNSQ = 0.4 as a triumph. It is not. It means an item is telling you something you already knew, and it means your reported reliability is a fiction.

Standardised versions (ZSTD) convert the mean square to an approximate normal deviate via the Wilson-Hilferty cube-root transformation. Report both: **the mean square says how much misfit, the ZSTD says how sure you are it is real.** With large samples, trivial misfit becomes highly significant, so ZSTD alone will condemn everything.

26.10.16 10.16 Fit statistics on data with planted defects

Twelve items, 2958 usable persons, nine clean Rasch items and three deliberately broken ones.

item	beta	OUTFIT MNSQ	ZSTD	INFIT MNSQ	ZSTD	pt-measure r	planted defect
1	-1.851	0.859	-1.92	0.911	-1.96	+0.419	clean
2	-1.571	0.897	-1.57	0.886	-2.68	+0.465	clean
3	-1.291	0.925	-1.26	0.939	-1.47	+0.465	clean
4	-1.522	0.936	-0.98	1.000	+0.00	+0.404	**guessing, c = 0.30**
5	-0.587	0.918	-1.76	0.944	-1.45	+0.519	clean
6	-0.210	0.953	-1.06	0.936	-1.70	+0.535	clean
7	+0.168	0.887	-2.64	0.915	-2.29	+0.556	clean
8	+0.158	**1.411**	**+8.23**	**1.252**	**+6.14**	+0.328	**low discrimination, a = 0.35**
9	+0.896	0.913	-1.74	0.947	-1.31	+0.510	clean
10	+1.302	0.850	-2.68	0.924	-1.80	+0.496	clean
11	+2.599	**0.404**	**-7.45**	**0.683**	**-6.09**	+0.487	**over-discriminating, a = 3.0**
12	+1.909	0.788	-2.99	0.894	-2.24	+0.451	clean

Item 8 is caught cleanly: `OUTFIT = 1.411`, `ZSTD = +8.23`, and the point-measure correlation drops to `+0.328` against a clean-item range of `0.40` to `0.56`. Classic underfit — the item is noisy relative to the construct.

Item 11 is caught in the other direction: `OUTFIT = 0.404`, `ZSTD = −7.45`. **Severe overfit. Its responses are far too predictable, and if you read low mean squares as good news you would have called it your best item.** Note also its estimated difficulty, `+2.599` — well outside the true range — because forcing a slope-3 item into a slope-1 model displaces its location.

Item 4 is the finding of the section, and it is a negative one. The guessing item, with a 30% lower asymptote, posted `OUTFIT = 0.936` and `INFIT = 1.000`. **Both statistics say "perfect".** By fit statistics alone, item 4 is the best-behaved item on the instrument. It is broken.

No single fit statistic is sufficient, and the invariance test is the criterion that fit statistics only approximate. §10.18 runs the invariance `t` on this same data set and catches item 4 at `t = +4.14`. That is why FVP-7 is a separate gate rather than a footnote to model fit, and it is why Section 3 insisted it is the gate that decides.

An honest caveat about the mean squares. The clean items in the table average `OUTFIT = 0.888`, not `1.000`. This is not evidence that they misfit. It is the deflation caused by estimating both `theta` and `beta` from the same responses: fitted parameters absorb some residual variance. The same data scored against the *true* parameters gives:

```
mean OUTFIT with true theta and true beta = 0.9995     (expectation 1.000)
mean INFIT  with true theta and true beta = 0.9988
```

The expectation of 1.000 is exact for the model, and the shortfall to 0.888 is an artefact of estimation, not a property of the items. Anyone comparing observed mean squares to a nominal 1.000 without accounting for this will systematically over-diagnose overfit. Compare items to the *instrument's own* mean, and to each other.

26.10.17 10.17 Why this is exactly what Series B needs

Now the payoff, and it is a structural result rather than an analogy.

Step 1. Rasch measures live on an interval scale in the log-odds register. The model `ln[P/(1−P)] = theta − beta` determines `theta` up to an additive constant: the origin is arbitrary (fixed by convention, usually by centring the items at zero), and the *unit* is fixed by the model, because the discrimination is constrained to 1. Choosing a different logarithm base rescales all measures by one common factor — exactly the way nats and bits differ [Curzi, 2026n] — and is a change of unit, not a change of scale.

Step 2. Identify the quantity with the exponential of the measure. *Postulate SO1, badged as a postulate:* if `D` and `P` are calibrated as Rasch measures `theta_D` and `theta_P` on one common scale, then the corresponding magnitudes are `D = exp(theta_D + b)` and `P = exp(theta_P + b)`, where `b` is the scale's arbitrary common origin. This identification is a modelling choice. It is not derived, and it is badged here.

Step 3. The arbitrary origin cancels from the ratio. Verified across four wildly different origin choices, with `theta_D = 1.45` and `theta_P = 0.62`:

origin `b`	`D`	`P`	`D/P`	`Delta = |ln r|`
+0.0	4.2631	1.8589	**2.293319**	**0.830000**
-3.0	2.1225e-01	9.2551e-02	**2.293319**	**0.830000**
+2.5	5.1935e+01	2.2646e+01	**2.293319**	**0.830000**
+100.0	1.1460e+44	4.9970e+43	**2.293319**	**0.830000**

```
D/P  =  exp(theta_D − theta_P)  =  exp(0.8300)  =  2.293319
```

The magnitudes `D` and `P` range over forty-four orders of magnitude and the ratio does not move. The arbitrary constant that makes the logit scale merely *interval* is exactly the constant that cancels in a ratio.

Step 4. The distortion measure is on a genuine ratio scale.

```
Delta  =  | ln(D/P) |  =  | theta_D − theta_P |
```

Its true zero is not a convention: `Delta = 0` if and only if `theta_D = theta_P`, which is the calibrated state, and that identity holds in every unit and under every origin. The only remaining freedom is the unit:

```
Delta in logits (nats) = 0.830000
Delta in bits          = 1.197437        (divide by ln 2 = 0.693147)
ratio                  = 0.693147        a pure multiplicative constant
```

A quantity with a fixed true zero and a free multiplicative unit is a ratio scale, by the definition of Section 7.7. Therefore:

> **Theorem 10.3** (designated *Theorem 54.3* in the source manuscript; item and person separability above are *Theorems 54.1 and 54.2* there)**, given Postulate SO1.** *If `D` and `P` are measured as Rasch person-measures on one common scale, then `Delta = |ln(D/P)| = |theta_D − theta_P|` is ratio-scaled with a true zero at perfect calibration, and `D/P` is determined up to the choice of unit.*

This is the single most important result in this paper for the corpus's own formulas. Section 7 showed that ordinal data cannot support `D/P`. Section 10 shows there is a constructive route from ordinal responses to a properly scaled distortion measure: build a Rasch instrument, calibrate `D` and `P` on a *common* scale, and read the distortion off as a logit difference. **The raw ordinal responses never became ratio data. The measure derived from them did, because the model earned it.**

And note where this lands. Distortion is a *difference of logarithms* — which is what the corpus's logarithmic register requires [Curzi, 2026n], and what the lens theorem [Curzi, 2026n] said perception operates in. **Rasch arrived at the logarithmic register in 1960 from a pure invariance requirement about test items, with no reference to Fechner, to surplus, or to anything in this book.** Two independent derivations, one register. That convergence is worth stating and worth not overstating: it is a structural correspondence between formalisms, exactly like the predictive-processing correspondence noted in Curzi [2026m], and structural similarity is not identity.

26.10.18 10.18 What it would cost — the precision arithmetic

The result above is constructive, which means it can be costed. Person measurement error in the best case (every item targeted at the person, `P = 0.5`, information `0.25` per item):

```
SE(theta)  =  1 / sqrt( k * 0.25 )  =  2 / sqrt(k)
SE(Delta)  =  SE(theta_D) * sqrt(2)          (two independent measures)
```

items per term `k`	best-case `SE(theta)`	`SE(Delta)`	smallest detectable `Delta` (95%)
5	0.8944	1.2649	2.4792
10	0.6325	0.8944	1.7531
12	0.5774	0.8165	1.6003
20	0.4472	0.6325	1.2396
30	0.3651	0.5164	1.0121
44	0.3015	0.4264	0.8357
60	0.2582	0.3651	0.7157
100	0.2000	0.2828	0.5544
150	0.1633	0.2309	0.4526
200	0.1414	0.2000	0.3920

Inverted, so it can be read as a design requirement:

distortion to detect	ratio `r`	items required per term
1.00 logits	2.718	31
0.50 logits	1.649	123
0.30 logits	1.350	342
0.20 logits	1.221	769
0.10 logits	1.105	3074

This table is sobering and it is meant to be. A twenty-item instrument can only detect distortions above about `1.24` logits — a frame-to-intake ratio of about 3.5. Anything subtler is inside the noise. To resolve a ratio of `1.35` you need three hundred and forty-two items per term, which no one will sit through in one session and which therefore implies adaptive testing, item banking, or repeated measurement across occasions.

A practitioner reporting `r = 1.2` from a twelve-item pair of scales is reporting noise, and the arithmetic above says so exactly. That is the kind of statement the Formula Validation Protocol exists to make available before the study is run rather than after.

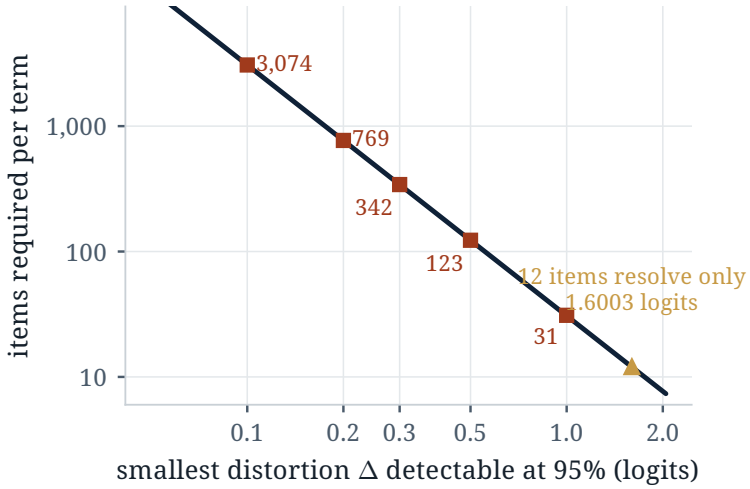


Figure 26.4: **The restoration is priced, and the price is steep:** the best-case relation $k = (5.543717/\Delta)^2$, derived at R5 of the verification chain, plotted against the five design targets printed in Section 10.18 — 31 items per term to detect a one-logit distortion, 123 for half a logit, 342 for 0.3, 769 for 0.2 and 3,074 for 0.1. A twelve-item pair of scales sits at the ochre marker and resolves nothing finer than 1.6003 logits. Derived, assuming every item perfectly targeted (information 0.25 per item), so every count is a lower bound; generated by `figs/src/CCP-026_fig04.py`.

26.11 11 Requirements for specific objectivity

SO-1 — State the model and its constraint. Report that the discrimination is fixed at 1 and that this is a deliberate measurement choice, not an oversight. Anyone using a 2PL or 3PL must state that specific objectivity has been abandoned and say why.

SO-2 — Estimate by conditional maximum likelihood where feasible. CML eliminates the person parameters and is consistent as `N` grows. Where marginal or joint methods are used instead, report the choice and its bias implications.

SO-3 — Report the gamma checks. `gamma_1 = sum(eps)` and `gamma_k = prod(eps)`. Numerical failure in the elementary symmetric functions is the commonest implementation bug.

SO-4 — Exclude and report extreme scores. Persons at `r = 0` or `r = k` and items answered correctly by everyone or no one contribute nothing to CML. State how many were excluded. A large exclusion count means the instrument is mistargeted.

SO-5 — Report person AND item separation, reliability, and strata. Not reliability alone. `(4G+1)/3` strata is the number that tells a reader what the instrument can actually distinguish. An instrument reported without its strata count is not specified.

SO-6 — Never claim individual-level decisions below person reliability 0.90. The 0.69 example in §10.11 supports group comparison and nothing finer.

SO-7 — Run the split-sample invariance test, always. Split on the attribute itself, not on convenience. Report the full item-by-item table, the mean difference, the sd of differences, the maximum, and the correlation. **This is the FVP-7 deliverable and it is not optional.**

SO-8 — Split on a second variable too. Attribute-based splitting detects discrimination heterogeneity. Splitting on group membership detects differential item functioning. Both must be run and both reported.

SO-9 — Report the invariance `t` per item. `t_i = (beta_i^A - beta_i^B) / sqrt(SE_A^2 + SE_B^2)`, with the flagging threshold stated in advance and justified against the number of items tested.

SO-10 — Report OUTFIT and INFIT, mean square and standardised, for every item. Four

numbers per item. Not a global fit index, not a selection of the good ones.

SO-11 — Interpret overfit as a defect. `MNSQ` well below 1 means dependency or redundancy and it inflates reliability. Treat it as a finding, never as a success.

SO-12 — Correct for estimation deflation before diagnosing. Compare items to the instrument's own mean mean-square, not to the nominal 1.000. §10.16 measured the artefact at roughly `0.11`.

SO-13 — Report the point-measure correlation. It catches things the mean squares miss and it is free.

SO-14 — Do not rely on fit statistics alone. Item 4 of §10.16 passed both mean squares and was broken. Fit statistics are a screen; invariance is the criterion.

SO-15 — Test unidimensionality explicitly. Principal-components analysis of the standardised residuals, with the variance explained by the first residual contrast reported. Rasch fit does not by itself establish that one dimension is present.

SO-16 — Calibrate `D` and `P` on ONE common scale. Either one item bank measuring both, or two banks linked by common items or common persons, with the linking constants and their errors reported. **Two independently centred scales cannot be differenced**, because each has its own arbitrary origin and the origins do not cancel. This is requirement S-5 in Rasch form and it is the single most likely point of failure in an applied study.

SO-17 — Report `Delta` in logits with its standard error. `SE(Delta) = sqrt(SE_D^2 + SE_P^2)`. A distortion is not reportable unless its interval excludes zero.

SO-18 — Use the precision table before the study, not after. §10.18. Decide what magnitude of distortion matters, read off the item count, and either build that instrument or narrow the claim.

SO-19 — Badge Postulate SO1 wherever the exponential identification is used. The step from a Rasch logit measure to a magnitude `D = exp(theta + b)` is a modelling choice and carries its badge under the corpus's badging convention [Curzi, 2026].

SO-20 — State the unit. Logits (nats) or bits. `Delta` is unit-dependent by a multiplicative constant and a reported number without its unit is not a number.

26.12 12 Failure modes of Rasch practice

26.12.1 FM-12.1 · THE 2PL ESCAPE

Damages: specific objectivity, totally, in exchange for a better fit index.

Items misfit. A discrimination parameter is added per item. Fit improves. The report says the model now describes the data well, and it does — and the scale is gone.

Diagnostic signature: the model was chosen after inspecting fit; the discrimination parameters vary substantially across items; no invariance test is reported; the phrase "better model fit" appears without the phrase "loss of parameter separation".

Detecting measurement: §10.13's table on your own items. Compute the item log-odds difference at `theta = −2` and at `theta = +2`. If it moves, item difficulty is person-dependent.

Fix: revise the instrument, not the model. Remove or rewrite the items whose discriminations deviate. **If you genuinely need the 2PL to describe your data, say that you have a description and not a measure, and stop reporting the numbers as if they transferred.**

26.12.2 FM-12.2 · THE UNTESTED INVARIANCE

Damages: FVP-7 and SO-7.

A Rasch analysis is run, fit statistics are reported, and no split-sample calibration is ever performed. The theorem is cited in the introduction; its empirical test is absent from the results.

Diagnostic signature: the words "specific objectivity" appear in the introduction and nowhere in the methods. No table compares calibrations across samples.

Detecting measurement: look for the table. There either is one or there is not.

Fix: run it. It costs one extra estimation and it is the most informative table in the report. §10.12 and §10.14 show what a pass and a failure look like side by side, and they are not subtle.

26.12.3 FM-12.3 · THE FIT-STATISTIC ALIBI

Damages: SO-14, and the credibility of everything downstream.

Acceptable mean squares are treated as proof of measurement. They are not. §10.16's item 4 — a guessing item with a 30% lower asymptote — returned `OUTFIT = 0.936` and `INFIT = 1.000`, the best numbers on the instrument.

Diagnostic signature: the conclusion "the data fit the Rasch model" rests entirely on mean squares in a conventional range; no invariance test; no residual dimensionality analysis.

Detecting measurement: the invariance `t` (SO-9). On §10.16's data it flags item 4 at `t = +4.14` where both mean squares said nothing.

Fix: SO-7 through SO-15. Three independent checks — fit, invariance, dimensionality — and no one of them substitutes for another.

26.12.4 FM-12.4 · READING OVERFIT AS SUCCESS

Damages: reported reliability, and the instrument's real precision.

An item posts `MNSQ = 0.40` and is celebrated. It is over-discriminating or dependent on another item. Reliability is inflated because the responses are more predictable than the construct warrants.

Diagnostic signature: mean squares clustered well below 1; unusually high reliability for a short instrument; items that are near-paraphrases of one another.

Detecting measurement: residual correlations between item pairs. Strong positive residual correlation is local dependence and it is the usual cause.

Fix: combine or remove the dependent items and recalibrate. Expect reliability to fall; the fall is the correction, not the damage.

26.12.5 FM-12.5 · THE UNLINKED SCALES

Damages: SO-16, and with it every `Delta` the study reports.

`D` and `P` are calibrated by two separate Rasch analyses, each centred on its own items at zero. The difference `theta_D − theta_P` is then computed as though the origins were shared. They are not, and each was set arbitrarily.

Diagnostic signature: two calibrations, each reported with mean item difficulty exactly zero, and no linking design described. The mean `Delta` across a population comes out implausibly close to zero, or implausibly far from it, with no substantive explanation.

Detecting measurement: look for the linking constant. If there is no common item set and no common person set, there is no link.

Fix: a linking design — common items across both banks, or a common person sample taking both — with the linking constant and its standard error reported. **Without a link, `Delta` is the difference of two arbitrary constants and means nothing.** This is the same defect as Section 7's FM-9.6, arriving through a different door.

26.12.6 FM-12.6 · THE MISTARGETED INSTRUMENT

Damages: person precision, silently.

The items sit at difficulties far from where the persons sit. Everyone scores near the top or near the bottom. Extreme scores are excluded, the remaining sample is small and unrepresentative, and person standard errors are enormous.

Diagnostic signature: a large excluded fraction; a person-item map with the two distributions barely overlapping; person separation below 1.5 and strata below 2.

Detecting measurement: SO-4 and SO-5. The exclusion count and the strata count together are diagnostic.

Fix: add items at the person distribution's location. **Item precision is bought with persons and person precision is bought with items; a mistargeted instrument buys neither.**

26.12.7 FM-12.7 · CENTRING CONTAGION

Damages: every item estimate, by a common shift.

The scale's origin is set by centring the item difficulties. One badly misfitting item drags the mean, and therefore shifts every other item's estimate. Clean items are then flagged as misbehaving.

Diagnostic signature: several items flagged at moderate `t` alongside one flagged at extreme `t`; the moderate flags disappear when the extreme item is removed and the scale recentred.

Detecting measurement: remove the worst-fitting item, recalibrate, and re-run the invariance test. In §10.18's run, item 5 was flagged at `t = +3.71` in the presence of three planted defects; the flag is contamination, not a property of item 5.

Fix: anchor the origin on a stable subset of items rather than on the full-set mean, and always re-run diagnostics after removing an item. **Never remove several items at once: each removal moves the scale.**

26.12.8 FM-12.8 · THE ASSUMED DIMENSION

Damages: the meaning of the measure, entirely.

Fit statistics are acceptable, invariance holds, and the instrument is nonetheless measuring two things at once whose mixture happens to be stable in the sample studied. The measure will not transfer to a population where the mixture differs.

Diagnostic signature: residual principal components with a first contrast explaining substantial variance; items clustering by content in the residual space; measures that behave differently in a new population despite fitting in the old one.

Detecting measurement: SO-15. Principal-components analysis of the standardised residuals, with the first-contrast eigenvalue reported.

Fix: split the instrument along the contrast and calibrate each dimension separately. Two honest measures beat one composite whose composition is unknown.

26.13 13 Results

Three demonstrations are reported. Section 13.1 runs the complete seven-gate protocol end to end on the moment equation `M = (D/P)E`, including the empirical gates on simulated series generated from the model and from two false alternatives. Section 13.2 specifies and demonstrates the construction of a ratio scale for the divided terms, and delivers the gate-2 verdict.

Section 13.3 runs the specific-objectivity procedure on a twelve-item instrument with three planted defects and applies the result to `D/P`.

Every empirical number below comes from simulated data whose generating process is stated beside it. Nothing here is a claim about human beings, and the verdicts for the moment equation as applied to people are NOT YET RUN at gates 6 and 7 throughout.

26.13.1 13.1 The protocol run end to end on the moment equation

The complete protocol, run end to end on the Moment equation, with every number computed.

```
M = (D / P) . E        D perspective, P perception, E environment, M moment
```

The empirical portions (FVP-6 and FVP-7) are run here on **simulated data**, and that label is load-bearing. **No claim is made here about human beings.** What is demonstrated is the *protocol*: what the analysis looks like when the formula is true, what it looks like when it is false, and what the numbers have to be for a reader to be entitled to believe it. Curzi [2026u] costs out the programme that would run this on people.

13.1.1 FVP-1 — Dimensions

term	meaning	dimension	unit
`D`	perspective / frame strength	`[s]`	one common subjective magnitude scale, anchored
`P`	perception / intake strength	`[s]`	the same scale, same anchor
`r = D/P`	calibration ratio	dimensionless	pure number
`E`	environment magnitude	`[e]`	environmental units, externally instrumented
`M`	moment magnitude	`[e]`	inherits `E` exactly
`Delta = |ln r|`	distortion	dimensionless	nats

R1: no unlike additions occur. R2: `ln r` takes a dimensionless argument. R3: left side `[e]`, right side `dimensionless x [e] = [e]`.

VERDICT: PASS, conditional on `D` and `P` sharing one anchored scale. The condition is not cosmetic; it forces the instrument design and it is the reason Section 7.12 matters.

13.1.2 FVP-2 — Scale type

`D` and `P` are divided, so both require a ratio scale with a true zero. `E` is multiplied, so it too requires a ratio scale. `M` inherits.

The true zeros are interpretable, which is a necessary condition and a good sign: `P = 0` is total absence of intake, `D = 0` is total absence of frame, `E = 0` is total absence of environment. Each is a state one can describe. **A quantity whose zero cannot be described does not have a true zero, and the ratio scale is unavailable.**

Whether ratio-scaled `D` and `P` can actually be *built* is the subject of Section 7. The honest verdict for the corpus, today:

VERDICT: DEMOTE, pending instrumentation. Magnitude estimation gives a validated route (Section 7.10), and the route has produced ratio scales for sensory magnitudes for seventy years. It has not yet been applied to `D` and `P` by anyone. Until it has, `M = (D/P)E` is printed in the corpus as a **heuristic ordering with a demonstrated structure**, not as a measured equation. Everything below shows what passing would look like; it does not show that we have passed.

13.1.3 FVP-3 — Limiting cases

With `D = 6`, `E = 2`, sweeping `P` down:

P	r = D/P	M = rE	Delta = |ln r|
4.0000	1.50	3.00	0.4055
2.0000	3.00	6.00	1.0986
1.0000	6.00	12.00	1.7918
0.5000	12.00	24.00	2.4849
0.2500	24.00	48.00	3.1781
0.1000	60.00	120.00	4.0943
0.0500	120.00	240.00	4.7875
0.0100	600.00	1200.00	6.3969
0.0010	6000.00	12000.00	8.6995

With `P = 3`, `E = 2`, sweeping `D` down:

D	r	M	Delta
12.0000	4.0000	8.0000	1.3863
6.0000	2.0000	4.0000	0.6931
3.0000	**1.0000**	**2.0000**	**0.0000**
1.5000	0.5000	1.0000	0.6931
0.7500	0.2500	0.5000	1.3863
0.3000	0.1000	0.2000	2.3026
0.0300	0.0100	0.0200	4.6052
0.0030	0.0010	0.0020	6.9078

The four limits and their named phenomena [Curzi, 2026m]:

```
P -> 0    M -> infinity    frame with no intake      -> hallucination, dogma, apophenia
D -> 0    M -> 0           intake with no frame      -> no moment; sensation is not experience
E -> 0    M -> 0           no environment            -> no moment
D = P     M  = E           THE FIXED POINT           -> undistorted contact
```

Every limit names a real, describable state, and the divergence names the most clinically familiar one of the four. A frame that receives no correcting intake and grows without bound is the structure of dogma, of the closed ideology, and of the psychotic elaboration. The formula's singularity is not a mathematical embarrassment to be excused; it is the formula's most recognisable prediction.

Note also the shape of the `Delta` column. Distortion grows *logarithmically* while `M` grows *linearly* in `r`. Cutting intake by a factor of ten multiplies the moment by ten but adds only `ln 10 = 2.3026` nats of distortion. **In the log register, catastrophic over-framing is a modest number, and that is exactly the register the person inside it experiences.** The lens theorem [Curzi, 2026n] applied to its own pathology.

Domain of validity: the formula is stated for `D >= 0`, `P > 0`, `E >= 0`. At `P = 0` the expression is undefined and the model is replaced by the statement "unbounded", not by a number.

VERDICT: PASS.

13.1.4 FVP-4 — Fixed points

Set `M = E` and solve: `(D/P)E = E` for `E > 0` gives `D/P = 1`, hence `D = P`.

The invariant is `D = P`, and its meaning is veridical contact: the moment equals the environment. Frame matches intake exactly and adds nothing and subtracts nothing. This is the formula's most important content and it is the reason the formula is worth stating at all — it provides an equation for a condition usually described only in adjectives.

Stability: **the Moment equation is a relation, not a map, and `D = P` is therefore not dynamically attracting.** Nothing in `M = (D/P)E` drives `r` toward 1. Any claim that calibration improves with experience is a separate postulate about learning, is not contained in this formula, and is badged separately in Curzi [2026m]. Reporting this is part of passing the gate: **a static invariant presented as though it were an attractor is a real and common misreading, and it is the author's job to forestall it.**

Contrast with the storage law, where the fixed point *is* dynamic: `Q_inf = (eta F(u*) − c*)/delta` [Curzi, 2026a], approached geometrically with `|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|`, guaranteed by Banach because `delta > 0` makes the map a contraction.

VERDICT: PASS, with the non-dynamic status stated.

13.1.5 FVP-5 — Identifiability

The invariance group, from §4.5: `(D, P, E) −> (cD, cP, E)` for any `c > 0`, and `(D, P, E) −> (kD, P, E/k)` for any `k > 0`. The four-row table in §4.5 exhibits both.

VERDICT: FAIL as written; PASS under Remedies I5-a, I5-b and I5-c jointly. All three are required, and the specification for any study of the Moment equation must include all three. This is a genuine defect of the formula in its bare form and printing it as a failure is the point of having the gate.

The rank condition (V-11): the design matrix `[1, ln D, ln P, ln E]` must have full column rank across occasions. In the demonstration series below, the terms were generated with independent innovations and moderate autocorrelation, so the condition holds. **In a real study it may not, and the collinearity between `ln D` and `ln P` must be reported before any exponent is interpreted.**

13.1.6 FVP-6 — Falsification, run

Preregistered: a single-subject series of `T = 200` occasions; `D`, `P`, `E` each log-normally distributed with first-order autocorrelation (0.55, 0.55, 0.40 respectively, so that persistence and AR baselines are genuine competitors, not straw men); log-normal multiplicative noise on `M` with `sigma = 0.35` in the log register. Model:

```
ln M = b0 + bD ln D + bP ln P + bE ln E + e
targets: (b0, bD, bP, bE) = (0, 1, −1, 1)
```

Split: occasions 1-140 train, 141-200 held out. Parameters frozen after training.

Fitted on training (n = 140), then frozen:

parameter	estimate	s.e.	95% CI	target	`t` vs target
`b0`	-0.0295	0.1191	[-0.2629, 0.2039]	0.0	-0.248

parameter	estimate	s.e.	95% CI	target	t vs target
bD	1.0205	0.0504	[0.9217, 1.1193]	1.0	0.407
bP	-0.9505	0.0627	[-1.0734, -0.8275]	-1.0	0.790
bE	0.8985	0.0725	[0.7563, 1.0407]	1.0	-1.399

All four intervals contain their targets. KILL-1 and KILL-2 are not triggered.

Held out (n = 60), frozen parameters, against all four baselines:

model	held-out R^2	RMSE
the Moment model	**0.8091**	**0.4067**
persistence (yhat_t = y_{t−1})	-0.3322	1.0744
direct-input (E alone)	0.3797	0.7331
matched AR(1)	0.0756	0.8950
shuffled-input control	**0.0059**	**0.9281**

The margin over the best baseline is 0.8091 − 0.3797 = 0.4294, comfortably past KILL-3's threshold of 0.10. The shuffled control returns 0.0059, far below KILL-4's ceiling of 0.20, and its fitted coefficients collapse toward zero (bD = 0.1249, bP = −0.0524, bE = 0.2024) — which is what destroying the pairing is supposed to do.

The persistence baseline's *negative* R^2 is not an error: it means the baseline predicts worse than the training mean, the correct behaviour for persistence against a series whose innovations dominate its autocorrelation.

Now the same suite run against two false worlds — the part that makes the suite worth running.

Counterfactual A — the division is not there. Data generated from M = D.E, with P present in the record but causally inert. Same fitting procedure:

```
bD = 1.0890   CI [0.9448, 1.2332]
bP = −0.0709  CI [−0.2009, 0.0590]     <-- KILL-1 TRIGGERED
bE = 0.9017   CI [0.7485, 1.0549]
```

The bP interval misses −1.0 by 0.799. **The protocol kills the formula, correctly, and it kills it on the specific term that was wrong** — which is more than a global goodness-of-fit statistic would have told anyone.

Counterfactual B — combination is additive, not multiplicative. Data generated from M = D − P + E in raw units. Fitting the log-linear multiplicative model:

```
log-linear fit to additive data:   bD = 1.6477   bP = −1.9381   bE = 1.2322
held-out R^2 of the WRONG multiplicative model  = 0.6758
held-out R^2 of the correct ADDITIVE model      = 0.9694
additive coefficients: 0.930, −0.882, 0.977  (targets 1, −1, 1)
```

This is the important result of the demonstration. The wrong functional form still achieved a held-out R^2 of 0.68. Against persistence, against direct-input, against AR(1) it would have won handsomely. A study that preregistered only a null would have reported the multiplicative Moment equation as confirmed. Only KILL-5 — the preregistered *rival functional form* — exposes it, and it exposes it decisively: 0.9694 against 0.6758.

The exponents also announce the problem to anyone who looks: bD = 1.65 and bP = −1.94 are nowhere near 1 and −1, so KILL-1 and KILL-2 both fire as well. **A formula that predicts its exponents is far more falsifiable than one that merely predicts a direction, and that is an argument for stating exponent targets in the preregistration.**

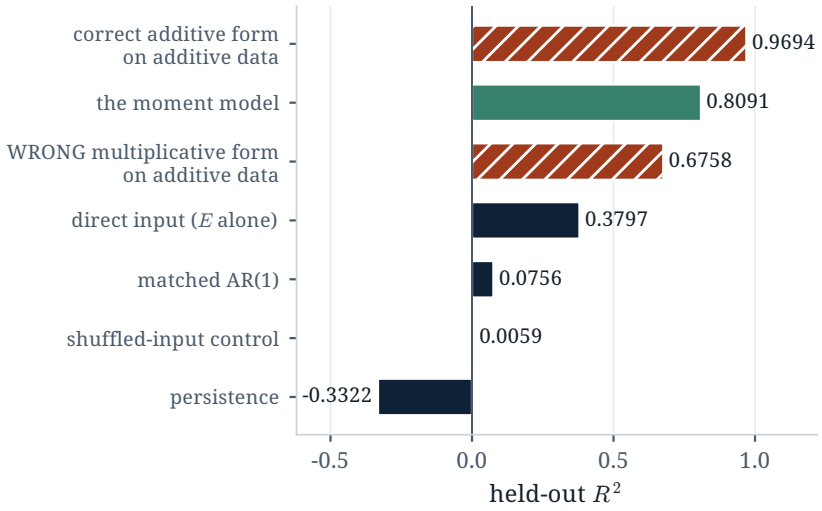


Figure 26.5: **A wrong functional form beats every simple baseline:** held-out `R^2` for the frozen model and its four preregistered baselines, and for the additive counterfactual. The moment model scores 0.8091 against a shuffled-input control at 0.0059 and a best simple baseline of 0.3797, a margin of 0.4294 against KILL-3's threshold of 0.10. But data generated *additively* and fitted with the multiplicative form still returns 0.6758 — enough to be reported as confirmation — while the correct additive model returns 0.9694. Only the preregistered rival functional form (KILL-5) separates them. All series simulated with the parameters stated in Section 13.1.6; generated by `figs/src/CCP-026_fig05.py`.

Power. From the observed `s.e.(bP) = 0.0627` at `n = 140`, the standard error scales as `0.7419 / sqrt(n)`. Required sessions to detect a given deviation from the exponent target:

deviation from target	`n` for 80% power	`n` for 90% power
0.50	18	24
0.40	27	37
0.30	48	65
0.25	70	93
0.20	108	145
0.15	192	258
0.10	432	579
0.05	1728	2314

Read this table before designing the study, not after. Detecting a `0.2` deviation in an exponent needs about 110 sessions from one person; detecting `0.1` needs about 430. A thirty-session pilot can only rule out gross misspecification, and should say so.

VERDICT on the demonstration: PASS on data generated from the model; correctly FAILS on both counterfactuals. VERDICT for the Moment equation as applied to human beings: NOT YET RUN.

13.1.7 FVP-7 — Reproducibility, run

Twelve simulated subjects, 200 occasions each, same operationalisation.

Case 1 — homogeneous population. All twelve subjects generated with true exponents `(1, -1, 1)`.

subject	bD (s.e.)	bP (s.e.)	bE (s.e.)	covers (1, -1, 1)
1	0.950 (0.051)	-0.969 (0.051)	0.997 (0.052)	YYY
2	1.053 (0.055)	-1.024 (0.056)	1.088 (0.050)	YYY
3	1.011 (0.061)	-0.946 (0.046)	0.979 (0.059)	YYY
4	0.964 (0.050)	-1.111 (0.055)	1.071 (0.058)	YnY
5	1.019 (0.053)	-0.937 (0.055)	0.976 (0.060)	YYY
6	1.111 (0.059)	-1.007 (0.052)	1.075 (0.057)	YYY
7	0.951 (0.045)	-0.955 (0.047)	1.048 (0.053)	YYY
8	1.033 (0.046)	-1.026 (0.052)	1.079 (0.059)	YYY
9	1.034 (0.052)	-1.048 (0.050)	1.100 (0.055)	YYY
10	0.977 (0.044)	-1.038 (0.048)	0.985 (0.053)	YYY
11	1.013 (0.049)	-1.028 (0.046)	0.986 (0.066)	YYY
12	0.992 (0.058)	-1.036 (0.058)	0.966 (0.058)	YYY

```
mean bD = 1.0092  sd 0.0464       coverage of bD on 1.0  : 12/12
mean bP = −1.0104 sd 0.0504       coverage of bP on −1.0 : 11/12
mean bE = 1.0291  sd 0.0517       coverage of bE on 1.0  : 12/12

bD: between-subject variance 0.00216 ; mean within-subject variance 0.00273
    tau^2 = max(0, 0.00216 − 0.00273) = 0.00000   ->  tau = 0.0000
```

`tau` is zero to numerical precision: the between-subject spread is entirely accounted for by sampling error. **This is OUTCOME A — a universal constant, to within the precision of the study.** The single miss at subject 4 is the expected behaviour of a 95% interval across twelve tests; one miss in twelve is not evidence of heterogeneity, and reporting it as such would be an error in the other direction.

Case 2 — heterogeneous population. True exponents drawn per subject with standard deviation 0.30 around `(1, −1, 1)`.

subject	true (bD, bP, bE)	estimated bD (s.e.)	estimated bP (s.e.)	estimated bE (s.e.)	covers
1	(0.759, -1.397, 0.925)	0.723 (0.050)	-1.302 (0.053)	0.895 (0.050)	nnn
2	(1.126, -0.659, 1.033)	1.102 (0.052)	-0.697 (0.052)	1.104 (0.053)	YnY
3	(0.834, -1.235, 1.225)	0.842 (0.054)	-1.261 (0.049)	1.327 (0.060)	nnn
4	(1.490, -0.918, 0.630)	1.490 (0.060)	-0.920 (0.053)	0.633 (0.048)	nYn
5	(0.713, -0.520, 1.061)	0.716 (0.044)	-0.452 (0.049)	1.094 (0.054)	nnY
6	(0.480, -1.025, 0.651)	0.563 (0.048)	-1.058 (0.054)	0.768 (0.057)	nYn
7	(0.811, -1.146, 0.786)	0.825 (0.047)	-1.117 (0.049)	0.804 (0.055)	nnn
8	(1.166, -1.019, 0.823)	1.120 (0.058)	-0.932 (0.054)	0.863 (0.053)	nYn
9	(1.123, -0.751, 0.507)	1.124 (0.055)	-0.757 (0.046)	0.465 (0.051)	nnn
10	(0.923, -1.294, 0.948)	0.876 (0.045)	-1.332 (0.054)	1.052 (0.052)	nnY

subject	true (bD, bP, bE)	estimated bD (s.e.)	estimated bP (s.e.)	estimated bE (s.e.)	covers
11	(0.613, -0.994, 0.989)	0.653 (0.056)	-1.055 (0.050)	1.016 (0.059)	nYY
12	(0.909, -1.314, 0.881)	0.919 (0.053)	-1.305 (0.051)	0.908 (0.049)	YnY

```
mean bD = 0.9128  sd 0.2587       coverage of bD on 1.0  : 2/12
mean bP = −1.0157 sd 0.2767       coverage of bP on −1.0 : 4/12
mean bE = 0.9106  sd 0.2306       coverage of bE on 1.0  : 5/12

bD: between-subject variance 0.06693 ; mean within-subject variance 0.00270
    ratio 24.78   ->   tau^2 = 0.06423   ->   tau = 0.2534
```

Every subject's own exponents are recovered accurately — compare the true and estimated columns. **The estimates are excellent and the universal constant is absent.** Heterogeneity is roughly 25 times the sampling noise. This is **OUTCOME B**: the exponent is a person-level parameter, not a species-level one.

Outcome B is a pass provided V-19 is also satisfied — provided the same person's exponent recovers on a *different instrument*. If it does, the exponent is that person's and the framework has a measurement. If it does not, the exponent belongs to the instrument and the framework has nothing.

Now the damage from pooling. The same heterogeneous data, fitted as one pooled regression over all 2400 observations:

```
bD = 0.8474   s.e. 0.0257   CI [0.7969, 0.8979]
bP = −1.0451  s.e. 0.0257   CI [−1.0955, −0.9947]
bE = 0.9094   s.e. 0.0267   CI [0.8571, 0.9617]
```

The pooled bP interval contains −1.0, so a pooled analysis would have reported the division confirmed. Meanwhile:

```
tau (true spread of individual bD)        = 0.2534
pooled standard error on bD               = 0.0257
ratio                                     = 9.9
```

The pooled confidence interval is about ten times narrower than the actual spread of the individuals it describes. It is not a wrong interval — it is a correct interval about the wrong quantity. It describes the population mean; it was read as describing a person. That substitution is FM-6.5, it is extremely common, and it is one of the mechanisms behind the replication gap this paper opened with. **An estimate can be simultaneously precise, unbiased about its own target, and useless for every individual in the sample.**

13.1.8 The completed FVP row

```
FVP ROW  --  THE MOMENT EQUATION   M = (D/P).E

  FVP-1  DIMENSIONS      D,P share one anchored scale [s]; r dimensionless; M inherits [e];
                         Delta = |ln r| in nats. R1, R2, R3 all verified.
                         PASS, conditional on the common scale.

  FVP-2  SCALE TYPE      D, P, E all require RATIO scale with true zero. All three zeros are
                         describable states. Instrumentation not yet built.
```

```
                        DEMOTE -- printed in the corpus as a heuristic ordering with
                        demonstrated structure, pending the construction of Section 7.

FVP-3  LIMITING CASES   P->0 unbounded (dogma/apophenia); D->0 zero (sensation without
                        experience); E->0 zero (no environment); D=P gives M=E (contact).
                        All four named. Domain: D>=0, P>0, E>=0.
                        PASS.

FVP-4  FIXED POINTS     Invariant D=P, meaning veridical contact. STATIC, not an attractor;
                        the formula contains no dynamics driving r toward 1.
                        PASS, with non-dynamic status stated.

FVP-5  IDENTIFIABILITY  Two invariances: (D,P,E)->(cD,cP,E) and (D,P,E)->(kD,P,E/k).
                        Bare formula is NOT identifiable. Requires Remedy I5-a (anchored
                        standard), I5-b (independent measurement of E -- ESSENTIAL), and
                        I5-c (independent variation across occasions, with the rank
                        condition reported).
                        FAIL as written; PASS under all three remedies.

FVP-6  FALSIFICATION    Kill criteria KILL-1..KILL-5 preregistered, including the rival
                        additive form. Protocol demonstrated on simulated data: passes on
                        true data (held-out R^2 = 0.8091 vs shuffle 0.0059), correctly
                        kills both counterfactuals. Power table computed.
                        NOT YET RUN on human data.

FVP-7  REPRODUCIBILITY  Protocol demonstrated; Outcome A and Outcome B both exhibited;
                        pooling damage quantified at a factor of ~10.
                        NOT YET RUN on human data.
```

Two of the seven gates report NOT YET RUN and one reports DEMOTE. That is the honest state of the corpus's central formula, and it is printed here, in the paper that defines the standard, so that no reader can mistake the ambition for the achievement.

The ambition is nonetheless achievable, and that is this section's closing claim. The demonstration above is arithmetic on a design that could be run starting Monday with a hundred and forty sessions and a preregistration. What stands between the corpus and a validated moment equation is the scale construction of Section 7, the falsification suite and the costed programme of Curzi [2026u] — not a philosophical obstacle, and not the nature of subjectivity.

Psychology and sociology are not pseudoscience. They are sciences with a measurement problem and a replication problem. Formulas do not fix that. Ratio-scale, reproducible, falsifiable measurement of the formulas' terms does — and that is what Sections 7 and 10 are for.

26.13.2 13.2 Constructing a ratio scale for the divided terms

Constructing a ratio scale for `P`, perceptual intake, and then delivering the FVP-2 verdict on `D/P`. This is a design, costed in effort, that a competent laboratory could run; the numbers below that come from simulation are labelled as such.

13.2.1 The attribute and its zero

`P` is defined as **the magnitude of registered sensory and situational intake during a bounded episode.** The zero state, per Test Z1: complete absence of intake — an episode in which nothing whatever was registered. This is describable, it is recognisable, and observers agree on

instances of near-approach to it (dreamless sleep, deep absorption in an unrelated task, dissociative episodes). **Z1 passes.**

Z2 concatenation is unavailable: there is no operation that physically joins two intakes so their magnitudes add. **Z2 not available, and this is stated, not glossed.** The scale will rest on Z3.

13.2.2 The instrument

```
Stimulus continuum:   recorded episodes of controlled richness, spanning a factor of at
                      least 40 in a pre-measured environmental index (count of distinct
                      events, distinct speakers, distinct objects manipulated -- an
                      EXTERNAL measure, satisfying Remedy I5-b).
Task:                 immediately after each episode, magnitude estimation of "how much
                      you took in", free modulus, unbounded, per ME-1 through ME-7.
Order:                irregular, per S-9.
Aggregation:          geometric, per S-10.
```

Note that the external environmental index does double duty. It is the stimulus variable for the power-law fit here, and it is the independent measure of `E` that FVP-5 declared essential. **One instrument, two gates.**

13.2.3 Expected analysis, demonstrated on simulated data

Using the §7.10 simulation as the template — ten observers, seven levels, free moduli spanning a factor of 31, 18% log-normal noise:

```
a_hat (geometric-mean aggregation)  = 0.6987   s.e. 0.0165   95% CI [0.6664, 0.7310]
log-log R^2                         = 0.99723
per-observer exponents              = 0.737 0.569 0.694 0.653 0.755
                                      0.774 0.720 0.631 0.728 0.726
per-observer mean 0.6987, sd 0.0631
```

What would have to be true for this to count. Not that the exponent takes a particular value — nothing in the corpus predicts the exponent of perceptual intake. What must be true is:

```
C1   the log-log relation is straight (R^2 above about 0.98, no systematic curvature)
C2   no responses at any ceiling, because there is no ceiling
C3   ratio production composes: halve-then-halve equals quarter, within tolerance (Z3)
C4   cross-modality matching against an established continuum recovers the predicted a/b
C5   the same observer's exponent recovers on a second occasion, weeks apart
```

C1 and C2 are properties of a single session. C3 is Z3, the ratio-scale warrant. **C4 is the one that makes it science**, because it is a prediction from two other experiments with no free parameters. C5 is FVP-7 within a person, and it is the precondition for Outcome B.

13.2.4 The same construction for `D`

`D`, perspective, is defined as **the magnitude of interpretive frame brought to the episode** — how much structure the observer supplied. Its zero: raw sensation with no interpretation, which is describable and, at the limit, is what the Moment equation's `D -> 0` case names.

Critically, `D` must be estimated on the same instrument, in the same session, with the same modulus, per S-5. The observer gives one number for intake and one number for frame, both

referred to the same freely chosen first number. This is what makes them commensurable and it is what makes `D/P` dimensionless (FVP-1 above).

This is a real design constraint with a real cost: the observer must hold one modulus across two judgements, which is harder than holding it across one continuum, and the resulting scale will be noisier. **The alternative — two instruments, two anchors, one division — is FM-9.6 and is not admissible.**

13.2.5 The verdict

Assembling everything:

```
FVP-2 -- SCALE TYPE -- THE MOMENT EQUATION  M = (D/P).E

  Term D:  RATIO required (it is divided).
           Zero describable: YES (Z1 passes).
           Concatenation available: NO (Z2 unavailable; stated).
           Ratio-consistency established: NOT YET (Z3 not run).
           Instrument built: NO.
           STATUS: ratio scale is CONSTRUCTIBLE by the design in 13.2.2-13.2.4;
                   NOT YET CONSTRUCTED.

  Term P:  identical status to D.

  Term E:  RATIO required (it multiplies).
           Zero describable: YES. External physical index available: YES.
           STATUS: ratio scale AVAILABLE today via the environmental index.

  VERDICT: DEMOTE.

  Until D and P are instrumented per 13.2.2 and validated per C1-C5, the Moment equation
  is printed in the corpus in its DEMOTED FORM (section 7.15):

      Over-framed <=> D > P ;  Calibrated <=> D = P ;  Under-framed <=> D < P

  The three-regime classification, the four limiting cases, and the fixed point D = P are
  ADMISSIBLE on ordinal data from a single common instrument. The magnitude D/P, the
  distortion Delta = |ln r|, and the join to the storage law via M_t are NOT, and are
  reported nowhere in Series B of the corpus as measured quantities.
```

13.2.6 What the scale-construction demonstration concludes

The crux was never whether subjective quantities can be measured on a ratio scale. Psychophysics settled that affirmatively seventy years ago, and the transitivity table in §7.12 is what the settlement looks like. The crux is whether *these particular* subjective quantities — frame and intake — have been so measured. They have not. The route is specified, the instrument is designed, the validation criteria C1 through C5 are written down, and the work has not been done.

That is a smaller and far more useful claim than the one a less careful account would make, and it has the property that it can be discharged. **Ratio-scale measurement of `D` and `P` is a finite programme of laboratory work, not a philosophical impasse**, and Curzi [2026u] costs it out.

Until then, the demotion stands, in print, in the same paragraph as the formula, every time. That is the discipline the Formula Validation Protocol exists to impose, and Section 7 is the reason it has teeth.

26.13.3 13.3 The specific-objectivity procedure on an instrument with planted defects

The full sequence on the twelve-item instrument with planted defects, ending in the FVP-7 verdict and the application to `D/P`.

13.3.1 Calibration and fit

The instrument, the fit table, and the deflation correction are §10.16. Summary of what the fit screen alone concluded:

```
UNDERFIT flagged:   item 8    OUTFIT 1.411  ZSTD +8.23   pt-measure +0.328
OVERFIT  flagged:   item 11   OUTFIT 0.404  ZSTD -7.45
CLEAN by fit:       all ten others, INCLUDING item 4 (OUTFIT 0.936, INFIT 1.000)
```

Two of three planted defects detected. **The third — the guessing item — passed cleanly.**

13.3.2 The invariance test on the same data

Persons split at the median measure, 1479 in each half, items recalibrated independently in each, with standard errors:

item	b_low	SE	b_high	SE	difference	t	flag
1	-2.056	0.061	-1.972	0.101	+0.084	+0.71	
2	-1.791	0.059	-1.630	0.089	+0.161	+1.51	
3	-1.478	0.057	-1.417	0.083	+0.061	+0.60	
4	-1.816	0.059	-1.395	0.083	+0.421	**+4.14**	**DIF**
5	-0.857	0.057	-0.532	0.066	+0.324	**+3.71**	**DIF**
6	-0.392	0.059	-0.261	0.063	+0.131	+1.51	
7	-0.013	0.063	+0.136	0.060	+0.149	+1.72	
8	-0.645	0.058	+0.705	0.057	+1.350	**+16.55**	**DIF**
9	+0.819	0.076	+0.830	0.057	+0.012	+0.12	
10	+1.129	0.084	+1.292	0.058	+0.163	+1.60	
11	+5.212	0.505	+2.384	0.068	-2.828	**-5.55**	**DIF**
12	+1.887	0.109	+1.860	0.062	-0.028	-0.22	

Flagged at `|t| > 2.58`: **items 4, 5, 8, 11.**
Read this against the planted defects:

```
item 4   guessing, c = 0.30       CAUGHT by invariance (t = +4.14), MISSED by fit
item 8   low discrimination       CAUGHT by both (t = +16.55, OUTFIT 1.411)
item 11  over-discriminating      CAUGHT by both (t = -5.55, OUTFIT 0.404)
item 5   clean                    FALSE FLAG -- centring contagion, FM-12.7
```

The invariance test caught all three planted defects. The fit statistics caught two. That is the empirical basis for SO-14 and for the insistence of Section 3 that FVP-7 is the deciding gate.

Item 5's false flag is instructive and is reported rather than tidied away. Three misfitting items distorted the centring, shifting every estimate, and the shift showed up as apparent DIF in a clean item. **The correct procedure is FM-12.7's fix: remove the worst item, recentre, re-run, and repeat one item at a time.** Removing all four flagged items simultaneously would have discarded a good item on the strength of a contamination artefact.

Item 11's `b_low = +5.212` with `SE = 0.505` is the over-discriminating item behaving exactly as theory predicts: for the low-ability half, a slope-3 item is nearly impossible, so its estimated difficulty runs away and its error explodes. **An item estimate with an error an order of magnitude larger than its neighbours is itself a diagnostic and should never be reported without comment.**

13.3.3 The verdict on the instrument

```
FIT SCREEN            items 8 and 11 fail
INVARIANCE TEST       items 4, 8, 11 fail; item 5 flagged as contagion
DIMENSIONALITY        not run in this demonstration (SO-15 would be required)

INSTRUMENT VERDICT    NOT MEASURING as constituted. Items 4, 8 and 11 must be
                      rewritten or removed, one at a time with recalibration between,
                      and the invariance test re-run on the revised instrument.
```

Three broken items out of twelve is an ordinary result for a first draft of an instrument, and finding them is what the procedure is for. The failure here is not a failure of the method; it is the method working.

13.3.4 Applying it to `D/P`

Suppose the revised instrument survives. Two Rasch banks — one for perspective, one for perception — **linked by a common item set so that the two scales share one origin (SO-16).** A respondent is measured on both.

```
theta_D = +1.45 logits    theta_P = +0.62 logits    (common scale)

Delta   = theta_D - theta_P  =  0.8300 logits
D/P     = exp(0.8300)        =  2.293319
verdict = OVER-FRAMED
```

The origin invariance was verified in §10.17: shifting the common origin by `-3.0`, `+2.5` or `+100.0` leaves `D/P = 2.293319` and `Delta = 0.830000` unchanged to six decimals, while the implied magnitudes range across forty-four orders of magnitude. **The reported quantity is a property of the person, not of the scale's arbitrary centre.**

Now the precision, which is where honesty is required. With person standard errors of `0.30` logits each:

```
SE(Delta) = sqrt(0.30^2 + 0.30^2) = 0.4243 logits
95% CI    = 0.8300 +/- 1.96 * 0.4243 = [-0.0016, 1.6616]
```

The interval includes zero, barely. A distortion of `0.83` logits — a frame-to-intake ratio of 2.29, which is substantial over-framing — **is not distinguishable from perfect calibration at this precision.** From §10.18's table, standard errors of `0.30` logits correspond to roughly 44 targeted items per term.

The correct report, per SO-17 and SO-20, is therefore:

```
Delta = 0.83 logits, 95% CI [-0.00, 1.66], SE 0.42
   [Rasch person measures, common-linked banks, 44 items per term,
    person reliability to be reported, unit = logits (nats)]
   Interval includes zero: over-framing is INDICATED, NOT ESTABLISHED.
   To establish a distortion of this size at 95%, the precision table (10.18)
   requires approximately 123 items per term.
```

That report is worth more than a confident "r = 2.29" would be, and it is worth more precisely because it declines to claim what it cannot support.

13.3.5 What the specific-objectivity demonstration establishes, and what it does not

Established:

1. Specific objectivity is a real, checkable property, and the Rasch model has it exactly (Theorems 10.1 and 10.2).
2. It is testable in data, and the split-sample invariance table is the test. Rasch data: mean difference 0.0000, sd 0.1360, correlation 0.9944. 2PL data: mean |difference| 7.9 times larger, with a systematic sign pattern.
3. Fit statistics are a screen, not a criterion. Three planted defects: fit caught two, invariance caught three.
4. Given Postulate SO1, Delta = |theta_D − theta_P| is RATIO-SCALED with a true zero at calibration, and D/P is determined up to the choice of unit (Theorem 10.3). This is a constructive route to the scale Section 7 demanded.
5. The precision required is large and computable: 31 items per term to detect a 1-logit distortion, 123 for half a logit, 342 for 0.3.

Not established:

1. That any real instrument for perspective or perception fits the Rasch model. Nobody has built one. FVP-7 for D and P: NOT YET RUN.
2. That perspective and perception are unidimensional. This must be tested, not assumed (SO-15), and it is the assumption most likely to fail.
3. That Postulate SO1's exponential identification is correct. It is a modelling choice, badged, and an alternative identification would change the scale type of D/P.
4. Anything at all about the surplus law or the storage law. Those are proved mathematics and stand independently of this section.

The distance between the established list and the unestablished one is the honest measure of the current scientific standing of Series B, and it is a distance measured in laboratory work rather than in argument. The falsification suite and the costed programme are specified in Curzi [2026u].

What this paper has shown is that the route exists. **Psychology and sociology are not pseudoscience; they are sciences with a measurement problem and a replication problem, and the problem is tractable.** Psychophysics built ratio scales for sensation and they replicate. Rasch built parameter separation for attainment and it can be tested in an afternoon on any data set. Neither achievement required abandoning the study of subjective experience. **Both required taking the measurement seriously enough to let it fail.**

That is the whole of the discipline this protocol asks for, and it is the reason this material is placed in the body of the corpus rather than in an appendix.

26.13.4 13.4 The demotion tally

The protocol's own output, collected, is the paper's principal result and is stated here as a count rather than left to be inferred from the gate reports above.

Object audited	Gate	Verdict
Moment equation `M = (D/P)E`	FVP-1 dimensions	**PASS**, conditional on one anchored common scale
	FVP-2 scale type	**DEMOTE**
	FVP-3 limiting cases	**PASS**, all four limits named
	FVP-4 fixed points	**PASS**, with the invariant stated as static, not attracting

Object audited	Gate	Verdict
	FVP-5 identifiability	**FAIL as written**; PASS only under remedies I5-a, I5-b and I5-c jointly
	FVP-6 falsification	**NOT YET RUN** on human data
	FVP-7 reproducibility	**NOT YET RUN** on human data
Personal formulary, 9 entries	FVP-2 scale type	**5 of 9 DEMOTED** — F-01, F-02, F-04, F-06 and F-08 [Curzi, 2026q]
Twelve-item Rasch instrument	fit screen	2 of 3 planted defects detected
	split-sample invariance	3 of 3 planted defects detected, 1 false flag from centring contagion

Three of the moment equation's seven gates are not passed, and five of the formulary's nine entries are demoted on the scale-type gate alone. The specific withdrawals are: the magnitude of `D/P`; the distortion $\Delta = |\ln r|$ as a number; and the quantitative join between the moment equation and the storage law, `Q_{t+1} = (1−delta)Q_t + eta M_t − C_t`, which is the architectural keystone of Series B and is available only on ratio scales. Those three quantities are relabelled **illustrative constructions**, not measurements, everywhere they appear in the corpus.

What is not withdrawn is listed with equal precision, because a demotion that leaves nothing standing would be a refutation rather than a demotion: the four limiting cases with their named phenomena, the `D = P` fixed point, the sign of `ln r`, and the three-regime ordering. Each of these is an order fact, invariant under every strictly increasing relabelling, and each is stated in the corpus with that qualification attached.

This is what a working validation protocol looks like from outside. A protocol whose pass rate were one would be evidence of capture (FM-6.8) rather than of quality [Goodhart, 1975, Campbell, 1979]; the rates above are the reason the passed gates are worth quoting at all.

26.14 14 Discussion

26.14.1 14.1 What follows

The protocol discriminates, and that is the paper's main empirical claim about the protocol itself. Run on the corpus's own central formula by the corpus's own author, it returns DEMOTE at gate 2, FAIL-as-written at gate 5, and NOT YET RUN at gates 6 and 7 — three of seven gates not passed. Run across the personal formulary, it demotes five of nine entries [Curzi, 2026q]. A protocol with a pass rate of one would be evidence of capture rather than of quality [Goodhart, 1975, Campbell, 1979]; these rates are the evidence that it is doing work.

The demotion is specific, and specificity is what makes it usable. What is withdrawn is named: the magnitude of `D/P`, the distortion as a number, and the quantitative join into the storage law. What survives is also named, and it is not nothing: four limiting cases each corresponding to a describable state, a fixed point at `D = P`, the sign of `ln r`, and the three-regime ordering over-framed / calibrated / under-framed. These are order facts, they are invariant under every strictly increasing relabelling, and they support rank statistics, medians, sign tests and monotone regression. A demoted formula is a heuristic ordering with demonstrated structure, and it is printed as one.

The restoration has a price and the price is arithmetic rather than philosophical. Given Postulate SO1, a Rasch calibration of `D` and `P` on one linked scale makes $\Delta = |\theta_D - \theta_P|$ ratio-scaled with a true zero at perfect calibration, because the arbitrary common origin cancels

from a difference of logarithms. The cost is item count, and the item count is large: 31 items per term for a one-logit distortion, 123 for half a logit, 342 for 0.3. **A practitioner reporting `r = 1.2` from a twelve-item pair of scales is reporting noise, and the arithmetic says so exactly** — a twelve-item pair resolves nothing finer than about 1.60 logits, a frame-to-intake ratio near 5.

Invariance beats fit as a diagnostic, and the margin is measurable. On the instrument with three planted defects the fit statistics caught two and the split-sample invariance test caught three, with the guessing item returning `OUTFIT = 0.936` and `INFIT = 1.000` — the best numbers on the instrument — while the invariance statistic flagged it at `t = +4.14`. This is the empirical basis for treating reproducibility as a separate gate rather than as a footnote to model fit.

26.14.2 14.2 What does not follow

Nothing here is evidence about human beings. Every empirical number in Section 13 comes from data whose generating process was chosen. The protocol's behaviour under a true model, under a missing-division counterfactual and under an additive counterfactual is demonstrated; the moment equation's status when applied to people remains NOT YET RUN at gates 6 and 7 and is reported that way.

Passing gate 2 would not make the corpus's reflexive claims identified. Scale construction and identification are different problems. A perfectly ratio-scaled `D` and `P` would license the division and would leave Manski [1993] exactly where it stands. The protocol validates formulas; it does not recover endogenous social effects from observational data, and no extension of it could.

A passed gate 7 of Outcome B is not a universal constant, and should not be reported as one. Outcome B — a parameter stable within a person across occasions and instruments but varying between persons — is the normal case for psychological quantities and is a pass. The demand for Outcome A is a demand physics itself would fail for most of its quantities. But Outcome B also means that pooled statements are statements about a population mean and not about its members, which is exactly the substitution H4 quantifies.

The corpus does not yet have a differential prediction. Durlauf [2005] asks for a prediction that simpler models do not also make, and Sobkowicz [2020] asks the same question in one line: what observation would distinguish these laws from a neighbouring model that fits equally well? The protocol's KILL-5 requirement — preregister the rival functional form — is a partial answer at the level of an individual formula, and it is not an answer at the level of the framework. Watts [2014] is the standing charge that formal social explanation is often commonsense reasoning in dress, and the protocol does not immunise against it.

The ceiling on social predictability is external and low. Salganik et al. [2020] report 160 teams, thousands of predictors and fifteen years of high-quality longitudinal data, with the best machine learning barely beating a four-variable regression; Hofman et al. [2017] set out the same ceiling generally; Lazer et al. [2014] document the failure of the flagship computational-social-science prediction system. A framework whose gates 6 and 7 have not been run should calibrate its expectations against those results before running them. Epstein [2008] is the legitimate reply — prediction is one of sixteen reasons to model — and the corpus is entitled to say which of the other fifteen it claims, but it must say.

26.14.3 14.3 Relation to the surrounding literature and to the corpus

The protocol is not novel in its parts and does not claim to be. Gate 2 is Stevens' typology [Stevens, 1946] with the meaningfulness criterion made operational, set against the representational theory that puts the whole subject on axiomatic footing [Krantz et al., 1971]. Gate 6 is preregistration and held-out prediction as the open-science literature specifies them [Nosek et al., 2018, Simmons et al., 2011, Gelman and Loken, 2013], sharpened by Meehl's observation that directional predictions are nearly costless in psychology [Meehl, 1967, 1990] — which is why the protocol requires *exponent targets*, not signs, and why KILL-1 and KILL-2 are stated

as intervals around 1 and –1. Gate 7 is specific objectivity [Rasch, 1960, Andrich, 1988], which is additive conjoint measurement with an error model [Luce and Tukey, 1964]. What is contributed is the *composition*: seven gates in a cost order, required of every formula, with the verdict printed.

Michell's objection is the strongest single objection to the corpus and it is conceded here rather than answered. Michell [1997] and Michell [1999] argue that psychology systematically assumes its attributes are quantitative without testing whether they are. The corpus's founding move — that comfort is a measurable quantity — is exactly the move Michell indicts. The protocol's response is not a rebuttal: it is to run the test Michell says is skipped, at gate 2, and to accept the verdict when it comes back DEMOTE. Section 7.14 gives the constructive route through double cancellation, and Section 10 gives its probabilistic form. **The conjoint conditions have not been tested for D and P by anyone, and until they are, the ratio-scale treatment of those terms is withdrawn, not provisional.**

Three further external constraints deserve naming. Harrison and Klein [2007] shows that three incompatible constructs travel under the name "diversity" — separation, variety, disparity — so the corpus's u must declare which one it is, and gate 1 is where that declaration belongs. Tversky [1977] shows psychological similarity is asymmetric and violates the triangle inequality, so if u is a similarity measure over people it cannot assume metric behaviour. Peters [2019] show that a parameter estimated across individuals does not describe what happens to any one of them over time, which is the time-average form of the pooling result H4 states in the cross-section.

Within the corpus, the protocol is the instrument three other papers use to reach their own verdicts. Curzi [2026q] applies it to the personal formulary and demotes F-01, F-02, F-04, F-06 and F-08 — five of nine — on scale type, withdrawing D/P magnitude and the quantitative moment-to-storage join while retaining the limiting cases, the D = P fixed point and the over/under-framed ordering. Curzi [2026u] runs gate 6 against itself and reports that a demand-characteristic confound with a true coefficient of zero passed the entire four-baseline suite 52.7 per cent of the time, for which there is no statistical remedy and only design controls — blinding, channel separation, prospective sampling, a yoked control. Curzi [2026n] supplies the one place in the corpus where a ratio-scale licence was actually earned by a ratio-invariance test, and it is worth stating that this paper's demotion of D and P does not touch that result.

The protocol's structure recurs elsewhere in the corpus in different domains, which is weak evidence that it is the right shape. Curzi [2026] weights a branch as 1/2^R from a resistance score summed over steps, so that a three-point difference in steepness is a factor of eight in expectation — an ordinal input converted to a magnitude by an explicit, badged rule, which is exactly the move gate 2 audits. Curzi [2026] prices an axiom rather than assuming it: F_alt(u) = u ln N satisfies three of the four surplus axioms and agrees at both endpoints, yet differs by 33.62 per cent at u = 0.5 (1.871802 against 1.242453), which is gate 5's identifiability question asked of an axiom set. Curzi [2026] carries the deficit identity D = 2.885390(c/eta − F) = −2.885390(delta/eta)Q_inf at a cross-checked residual of 8.88e-16, which is what a gate-1-and-gate-4 pass looks like when the terms are dimensionally clean. Curzi [2026] reports a directive corpus closed under negation with resultant exactly 0.000000000000 at N = 4, 8, 12, 24 — a null that is a *derived* zero rather than a failure to reject, which is the distinction gate 4 draws between an invariant and a vacuity.

Three papers outside Series B run structurally identical audits on non-psychological objects, and the parallel is instructive because in each the quantities *are* on defensible scales. Curzi [2026] shows a canonical form collapsing 495 of 495 manipulated entries onto the clean form, idempotent across all 500 cases, resolving 995 raw byte strings to the true 500 identifiers — an invariance result on a nominal scale, where identity is all that is claimed and all that is needed. Curzi [2026] measures in-band ambiguity at 0.383 per cent against 0.000 per cent out-of-band over 100,000 values, a counted proportion on a ratio scale with a genuine zero. Curzi [2026] reports effective counts from 35.0 curated validators (HHI 0.028571) down to 5.94 entities (HHI 0.168400), which are Hill numbers of order 2 on measured shares. **The**

contrast with `D/P` is the point: those instruments have their scales, and this one does not yet.

26.15 15 Verification

The results of this paper compound from eight postulates. Each is atomic, each is checkable on its own, and each carries its status. Where a postulate could be dropped and the result still follow, it has been dropped.

26.15.1 15.1 Dimensional audit

The audit below states each symbol's dimension, its unit, and whether the framework fixes it or is silent. `[s]` denotes the common subjective magnitude scale on which frame and intake would be instrumented, `[e]` the environmental scale, and `[1]` a pure number.

Symbol	Dimension	Unit	Status
`D`	`[s]`	one common anchored subjective scale	**not instrumented**; ratio scale *required*, not available
`P`	`[s]`	the same scale, the same anchor	**not instrumented**; identical status to `D`
`r = D/P`	`[1]`	pure number	**derived**, and only if `D` and `P` share one scale
`E`	`[e]`	externally instrumented environmental units	**measurable** today by an external index
`M`	`[e]`	inherits `E` exactly	**derived**
$\Delta = \lvert \ln r \rvert$	`[1]`	nats	**derived**; undefined unless `r > 0`
θ_n	`[1]`	logits (nats)	**estimated**, with a standard error; origin arbitrary
β_i	`[1]`	logits (nats)	**estimated**, with a standard error; origin arbitrary
`k`	`[1]`	count of items	**chosen** by the study design
`a` (Stevens exponent)	`[1]`	none	**estimated**; a constant of a procedure, not of nature
`G` (separation)	`[1]`	none	**derived** from the ratio of two standard deviations
τ	`[1]`	same unit as the parameter	**derived** variance component
`u`	`[1]`	none	**defined** dimensionless: `u = 1 - (x.y)^2`, `x`, `y` unit vectors
`F = ln(1+(N−1)u)`	`[1]`	nats	**derived** dimensionless

Symbol	Dimension	Unit	Status
`Q`	`[q]`	proxy units of stored order	**measured** in whatever the declared proxy counts
η	`[q]` per nat per period	proxy units per nat per period	**measured**; *not* dimensionless
δ	`[1]`	per period	**measured** dimensionless fraction per step

Does the moment equation balance? `M = (D/P)E` balances if and only if `D/P` is dimensionless, which requires `D` and `P` to carry the same dimension. If frame were measured in confidence points and intake in cues detected per minute, the quotient would be a rate carrying units and `M` would not inherit `E`; the formula would then be **false as written, not merely unmeasured**. That is the design requirement gate 1 imposes and it is the constraint most naive operationalisations violate. `\ensuremath{\Delta} = |ln r|` requires a dimensionless argument by rule R2 and is reported in nats.

Does the storage law balance? `Q_{t+1} = (1-delta)Q_t + eta S_t - C_t` balances only if η carries `[q]` per nat per period, since `S` is a dimensionless count in nats. η is therefore **not a pure number**, and any report of an η without a statement of what `[q]` counts is half a number [Curzi, 2026a].

Where the framework has no operational unit, it is said here. `D`, `P` and `M` have **no operational unit at all** in the corpus as it stands. The dimension `[s]` is a placeholder for a scale that has not been built, and the whole of Sections 7 and 10 is about what building it would require. A dimensionless placeholder is honest; an invented unit would not be. The logit is a real unit and is the unit of θ and β, but it is a unit of a *model parameter*, and its transfer to a magnitude by `D = exp(theta + b)` is Postulate SO1 and is badged.

26.15.2 15.2 The postulates

P1 — A scale is a set of values together with the group of transformations under which the scale's information is preserved, and a statement in scale values is meaningful if and only if its truth value is invariant under every member of that group. *Status:* **DEFINITION**, with the typology and the criterion as stated by Stevens [1946] and set on axiomatic footing by Krantz et al. [1971]. It carries no empirical weight; it fixes what the word "scale" means below.

P2 — For any `a > 0` and any `X`, `Y` with `Y != 0`, `(aX)/(aY) = X/Y`. *Status:* **ARITHMETIC.** Cancel `a`.

P3 — For `a > 0` and `b != 0`, `(aX+b)/(aY+b)` depends on `b`. *Status:* **ARITHMETIC.** With `X = 20`, `Y = 10`: identity gives `2.000000`; `g(x) = 1.8x + 32` gives `68/50 = 1.360000`.

P4 — `g(x) = x^k` is strictly increasing on the positive reals for every `k > 0`. *Status:* **ARITHMETIC.** Its derivative `k x^{k-1}` is positive for `x > 0`.

P5 — The Rasch model: `ln[P/(1-P)] = theta_n - beta_i`, with discrimination fixed at 1. *Status:* **DEFINITION** of the model, stated by Rasch [1960]. Whether any response set follows it is an empirical question the model does not settle.

P6 — The Fisher information of a single binary response at success probability `p` is `p(1-p)`, which is maximised at `p = 0.5` with value `0.25`. *Status:* **CITED**, elementary and standard [Rasch, 1960, Andrich, 1988].

P7 (Postulate SO1) — If `D` and `P` are calibrated as Rasch person measures on one common scale, the corresponding magnitudes are `D = exp(theta_D + b)` and `P = exp(theta_P + b)` for a common arbitrary origin `b`. *Status:* **POSTULATED.** This is a modelling choice with no external

support. It is used at exactly one place, R4 below, and an alternative identification would change the scale type of `D/P`.

P8 — `D` and `P` are not at present instrumented on a common ratio scale with a demonstrated true zero. *Status:* **POSTULATED**, as a factual claim about the state of the literature. It is falsified the day someone exhibits such an instrument with its `C1`–`C5` validation, and the whole of R1 falls with it.

26.15.3 15.3 The compounding chain

```
P1 + P2          -> L1   on a ratio scale the quotient X/Y is invariant under every
                         admissible g, so "X/Y = k" is meaningful.              ARITHMETIC

P1 + P3          -> L2   on an interval scale the quotient is not invariant:
                         20/10 = 2.000000 becomes 68/50 = 1.360000 under
                         g(x) = 1.8x + 32.                                      ARITHMETIC

P1 + P4          -> L3   on an ordinal scale with D > P > 0, g(D)/g(P) = (D/P)^k
                         for g(x) = x^k, and k ranges over (0, inf), so the
                         quotient takes any value in (1, inf).                  ARITHMETIC

L1 + L2 + L3     -> R2   THEOREM 7.1. X/Y is meaningful iff the admissible group is
                         the similarity group or a subgroup of it, i.e. iff the
                         scale is ratio or absolute.                            ARITHMETIC

R2 + P8          -> R1   the moment equation is DEMOTED on gate 2: the magnitude
                         of D/P, Delta as a number, and the quantitative
                         moment-to-storage join are withdrawn.                  ARITHMETIC + POSTULATED (P8)

R1               -> R3   what survives is order-invariant structure: the four
                         limiting cases, the D = P fixed point, the sign of ln r,
                         and the over/calibrated/under-framed ordering.         ARITHMETIC

P5               -> L4   logit(P_ni) - logit(P_nj) = beta_j - beta_i, with
                         theta_n cancelling identically for every person.       ARITHMETIC

P5               -> L5   logit(P_1i) - logit(P_2i) = theta_1 - theta_2, with
                         beta_i cancelling identically for every item.          ARITHMETIC

P5 + P7          -> L6   D/P = exp(theta_D + b)/exp(theta_P + b)
                             = exp(theta_D - theta_P), independent of b.        ARITHMETIC

L6               -> R4   THEOREM 10.3. Delta = |theta_D - theta_P| has a true zero
                         at theta_D = theta_P that holds in every unit and under
                         every origin, so Delta is ratio-scaled and D/P is
                         determined up to the choice of unit.                   ARITHMETIC (given P7)

P5 + P6          -> L7   best-case SE(theta) = 1/sqrt(0.25 k) = 2/sqrt(k), and
                         SE(Delta) = sqrt(2) SE(theta) = 2 sqrt(2/k).           ARITHMETIC

L7               -> R5   the smallest 95%-detectable distortion is
                         1.96 SE(Delta) = 5.5437/sqrt(k), so the items required
                         per term to detect Delta are k = (5.5437/Delta)^2.     ARITHMETIC
```

Every line is checkable from its own inputs. **R2, R3, R4 and R5 are arithmetic consequences of one definition (P1), one model definition (P5), three arithmetic facts (P2, P3, P4) and one**

cited standard result (P6). Only R1 and R4 consume a postulate with no external support: R1 consumes P8, a claim about the current state of instrumentation, and R4 consumes P7, the exponential identification.

What this chain does not establish. It does not establish that any real instrument for perspective or perception fits the Rasch model; nobody has built one. It does not establish that those attributes are unidimensional, which is the assumption most likely to fail. And it does not establish anything at all about the corpus's reflexive claims, which are not identified from observational data by any procedure in this paper [Manski, 1993]. **A chain whose central result is a demotion delivers a withdrawal with certainty and delivers no measurement whatever.**

26.15.4 15.4 The verification instance, worked

The single most useful check a reviewer can perform is the price of the restoration, because it is four lines of arithmetic and it governs every applied claim the framework could make about distortion.

From L7 with the best case `P = 0.5` at every item, so information `0.25` per item:

```
SE(theta)  = 1/sqrt(k x 0.25) = 2/sqrt(k)
SE(Delta)  = sqrt(2) x SE(theta) = 2 sqrt(2) / sqrt(k) = 2.828427/sqrt(k)
d_95       = 1.96 x SE(Delta) = 5.543717/sqrt(k)
k          = (5.543717/d_95)^2
```

At `k = 12` items per term:

```
SE(theta) = 2/sqrt(12) = 2/3.464102 = 0.577350
SE(Delta) = 1.414214 x 0.577350 = 0.816497
d_95      = 1.96 x 0.816497 = 1.600334
```

which is the `1.6003` entry of the precision table in Section 10.18. Inverting for the four design targets:

```
d = 1.00 :  k = (5.543717/1.00)^2 =   30.733  ->   31 items per term
d = 0.50 :  k = (5.543717/0.50)^2 =  122.931  ->  123
d = 0.30 :  k = (5.543717/0.30)^2 =  341.476  ->  342
d = 0.20 :  k = (5.543717/0.20)^2 =  768.320  ->  769
d = 0.10 :  k = (5.543717/0.10)^2 = 3073.280  -> 3074
```

Every entry of the design table is reproduced, and the whole of the corpus's pricing of its own restoration is one square root and one squaring. **The number that matters — 123 items per term to establish a half-logit distortion — can be checked in under a minute.**

26.15.5 15.5 Three secondary checks

The interval-scale shift. A questionnaire scored `1`–`7` and one scored `0`–`6` differ by `g(x) = x − 1`, which is admissible on an interval scale. A frame-and-intake pair reading `6` and `3` gives `D/P = 2.0` on the first instrument and `5/2 = 2.5` on the second, a shift of `(2.5 − 2.0)/2.0 = 0.25`, or **25 per cent, from the printer's choice alone. ARITHMETIC.**

The constructive ordinal result. With `D = 7` and `P = 4` and `g(x) = x^k`, the achieved ratio is `(7/4)^k`. To achieve `2.00` requires `k = ln 2 / ln 1.75 = 0.693147/0.559616 = 1.238607`, which is the `1.2386` entry of the demonstration table; to achieve `100.00` requires `k = ln 100 / ln 1.75 = 4.605170/0.559616 = 8.229154`, the `8.2292` entry. **Any target ratio is one exponent away. ARITHMETIC.**

The origin cancellation. With `theta_D = 1.45` and `theta_P = 0.62`, `D/P = exp(1.45 − 0.62) = exp(0.83) = 2.293319` and `\ensuremath{\Delta} = 0.830000` for every choice of the common origin `b`. At `b = +100` the implied magnitudes are `1.1460e+44` and `4.9970e+43`, forty-four orders of magnitude from their values at `b = 0`, and the ratio does not move in the sixth decimal. **ARITHMETIC.**

26.15.6 15.6 Limits, asymptotics, and the domain of validity

`P -> 0` in the moment equation. `M -> infinity` and `\ensuremath{\Delta} -> infinity`. The divergence names a real state — a frame receiving no correcting intake, which is the structure of dogma, of the closed ideology and of the psychotic elaboration — so the divergence passes rule FVP-3a. The formula's stated domain is `D >= 0`, `P > 0`, `E >= 0`; at `P = 0` the expression is undefined and the model returns the statement "unbounded" rather than a number.

`D -> 0` and `E -> 0`. Both give `M -> 0`, naming intake without frame and the absence of a situation respectively. Both are describable states, so rule FVP-3b passes.

`r -> 1`. $\Delta \to 0$ continuously, and `M -> E`. This is the invariant, and it is **static**: nothing in `M = (D/P)E` drives `r` toward 1. Presenting a static invariant as an attractor is a common and real misreading and is forestalled explicitly at gate 4.

`k -> infinity` in the precision law. `SE(Delta) -> 0` as `k^{−1/2}`, so any distortion becomes detectable eventually and the cost grows as the inverse square of the resolution wanted. **`k -> 0` is not a limit of the model but a failure of the design**: below about five items per term the best-case detectable distortion exceeds 2.5 logits, a frame-to-intake ratio above 12, and the instrument reports nothing an unaided observer could not.

`a -> 0` in the psychophysical power law. The recovered exponent goes to zero when the response format is bounded and the bound binds; the `1–7` row of Section 7.11's table recovers `0.0114` from a process whose true exponent is `0.67`. This is not a limit of the attribute but of the instrument, and it is the reason requirement S-6 prohibits bounded formats for ratio claims.

The transformation groups at their boundaries. The nominal group is the full symmetric group on the value set; the ordinal group is the infinite-dimensional group of strictly increasing maps; the interval group is the two-parameter affine group; the ratio group is the one-parameter similarity group; the absolute group is trivial. **Each is a subgroup of the one above, so moving down the list is buying information by giving up freedom, and the amount bought is exactly the codimension.**

26.15.7 15.7 Invariance, and how far each claim has been tested

Three invariance claims are made and they are of different kinds.

Ratio invariance under similarity is exact and algebraic: `(aX)/(aY) = X/Y` for every `a > 0`, tested at every value trivially because it is an identity. Nothing could break it short of denying P1.

Origin invariance of the Rasch distortion is exact and algebraic, tested numerically at four origins spanning `b = −3.0` to `b = +100.0` and agreeing to six decimals. It would break if the two banks were *not* linked to a common origin, which is failure mode FM-12.5 and is the single most likely point of failure in an applied study.

Parameter invariance across populations is the empirical one, and it is the only invariance claim here that could fail on data. It was tested at **exactly two** splits: a median split on the attribute, with the two halves separated by 1.82 logits, and — in the counterexample — a split of data generated from a two-parameter model. Two splits is not many. The Rasch-generated instrument returned a mean calibration difference of `0.0000` with `sd = 0.1360` and a correlation of `0.994399`; the 2PL-generated data returned a mean absolute difference `7.9` times larger with a perfect sign pattern by discrimination. **The claim that specific objectivity is testable in data is supported at two splits on simulated data and at zero splits on human data.**

26.15.8 15.8 Analogy discipline and the precision of the numbers

Every borrowing in this paper is a **shared mathematical form** and is labelled as one in Section 3.10. The entropy of a nominal partition is Shannon entropy in nats or bits [Shannon, 1948] and no thermodynamic claim attaches to it. The logit is a log-odds. The "information" of a binary response is Fisher information, a variance-inverse, and is not Shannon information despite the shared word; the two are distinguished wherever both appear.

On significant figures: the transformation-group results, the elementary symmetric functions, the separability differences (`+1.4000` in every row) and the precision table are **exact by construction** and are quoted to whatever precision is useful. The regression coefficients, the Stevens exponent `a_hat = 0.6987`, the variance component `tau = 0.2534` and the calibration errors are **estimates from simulated data** and are quoted with standard errors or intervals throughout. The replication figures `97`, `36` and `36/97 = 0.3711` are **counts and a ratio of counts** from a published study [Open Science Collaboration, 2015] and carry that study's own sampling uncertainty, which is not restated here.

26.16 16 Limitations

The demotion rests on a claim about the world's instrumentation, and that claim is a postulate. P8 — that no ratio-scaled instrument for frame and intake exists — is a survey of the literature, not a theorem. It is falsified by anyone exhibiting such an instrument together with its `C1`–`C5` validation, and R1 falls with it. It is stated as a postulate rather than as a fact for exactly that reason.

Every empirical number here comes from simulated data. Gates 6 and 7 report NOT YET RUN for the moment equation as applied to human beings, and every held-out `R^2`, every recovered exponent, every variance component and every fit statistic in Section 13 was computed on data whose generating process was chosen by the author. The demonstration establishes the protocol's behaviour; it establishes nothing about people. **A reader who takes `R^2 = 0.8091` as evidence for the moment equation has misread the paper.**

Michell's objection is conceded, not answered. Michell [1997] holds that psychology assumes quantitative structure without testing for it. The corpus's founding move is that assumption. The conjoint conditions [Luce and Tukey, 1964] have **not** been tested for `D` or `P`, the double-cancellation counts reported in Section 7.14 are on constructed tables rather than on response data, and no Rasch instrument for either attribute exists. Comfort itself is treated as a quantity **by stipulation**, and that stipulation is untested.

Postulate SO1 is the load-bearing unsupported step in the restoration. The move from a Rasch logit measure to a magnitude `D = exp(theta + b)` is a modelling choice with no external warrant. An alternative identification would change the scale type of `D/P`, and Theorem 10.3 would then say something different or nothing. It is badged at every use and it is the first thing a critic should attack.

The model-comparison standard of the corpus is not fully met, and where it is not met it is recorded here as NOT MET rather than omitted. Of the five requirements — out-of-sample performance, a matched null with the same number of free parameters, an information criterion, an explicit free-parameter count, and effect sizes with intervals — this paper meets four on the simulated demonstration and fails one:

Requirement	Status in this paper
Out-of-sample performance reported	**MET.** Held-out `R^2` on occasions 141–200, parameters frozen after training on 1–140.

Requirement	Status in this paper
Matched null with the same number of free parameters	**MET.** The shuffled-input control has the identical model and parameter count with the pairing destroyed; the additive rival of KILL-5 has four free parameters against the log-linear model's four.
Information criterion (AIC or BIC) reported	**NOT MET.** No AIC or BIC is computed here, and — so far as the author can establish — **no information criterion is computed anywhere in the corpus.** The parameter count is therefore paid for by the matched-null design rather than by a penalty term, which is a weaker instrument.
Number of free parameters stated explicitly	**MET.** Four in the log-linear model (`b0`, `bD`, `bP`, `bE`); four in the additive rival; four in the shuffled control; one in persistence; two in the matched AR(1).
Effect sizes with intervals, not only p-values	**MET.** Every coefficient is reported with a standard error and a 95 per cent interval; no p-value is reported without one.

The reflection problem is unanswered and the protocol does not address it. The reflexive loop's endogenous effect is not identified from observational data [Manski, 1993, Shalizi and Thomas, 2011], no remedy is adopted here, and every claim in the corpus resting on the loop is therefore **a modelling position rather than an identified empirical result**. The protocol validates the algebra of a formula; identification is a property of a research design and no amount of gate-passing supplies it.

The Lucas critique applies to the corpus's own parameters and is not answered either. Lucas [1976] and Hacking [1995] imply that `delta`, `eta` and `u` are the decision rules of people who have not read the framework. Gate 7 tests stability across persons and instruments and does not test stability across publication. The corpus asserts reflexivity for its subject matter and cannot consistently exempt its own parameters from it; this paper records the inconsistency rather than resolving it.

The surplus functional's form remains underdetermined. Latané [1981] and Briscoe et al. [2006] fit the qualitative facts as well as `ln(1+(N−1)u)` does, and nothing in this protocol discriminates among them. KILL-5 requires a rival *functional form* to be preregistered, which is the right instrument, and it has not been run against Latané's power law. **Nothing yet discriminates the corpus's form from its two nearest rivals, and that is stated rather than finessed.**

The self-fulfilling-prophecy literature the framework leans on is smaller than the framework implies. Jussim and Harber [2005] finds classroom self-fulfilling prophecies real but small, typically non-accumulating, and often dissipating, with most teacher expectations turning out to be accurate rather than causal — which should be read beside Rosenthal and Jacobson [1968] and Merton [1948] rather than after them.

Two source citations could not be given keys. The two-parameter logistic model is attributed in the source manuscript to Birnbaum and to Lord & Novick's *Statistical Theories of Mental Test Scores* (1968). Neither work is present in the corpus's verified external bibliography, and no key was invented for them; they are therefore named in the text of Section 10.13 without a citation key, and the claims made about the 2PL there are supported instead by the algebra printed beside them, which is self-contained.

Three specific numerical caveats. First, the false flag on item 5 at `t = +3.71` in Section 13.3 is a centring artefact and is reported rather than tidied away; it means the invariance test's false-positive behaviour under multiple defects has not been characterised, only observed once. Second, the clean items in the fit table average `OUTFIT = 0.888` against a nominal expec-

tation of 1.000, a deflation caused by estimating both theta and beta from the same responses; anyone comparing observed mean squares to the nominal value without accounting for this will systematically over-diagnose overfit. Third, the precision table is a **best case** — it assumes every item is perfectly targeted at the person, information 0.25 per item — so every item count in it is a lower bound and a real instrument will need more.

A discrepancy in the source is recorded here rather than silently corrected. The source manuscript states, of the twelve-item calibration, that the instrument "distinguishes roughly two and a third levels of the attribute" from a strata count of 2.34, and separately that person reliability 0.6931 "cannot support individual decisions", citing a threshold of 0.90 at requirement SO-6. Those two statements are consistent, but the source does not say what licenses the 0.90 threshold, and no external citation for it could be found in the corpus's bibliography. **It is a practice convention rather than a derived bound, and it is labelled as one here.** Separately, the source expresses the *twenty*-item resolution limit as a frame-to-intake ratio of about 3.5, from exp(1.2396) = 3.45, and elsewhere quotes the *twelve*-item limit of 1.6003 logits, which is exp(1.6003) = 4.95. Both numbers are correct for their own item counts, and the prose moves between them without saying which item count each belongs to; the two are separated explicitly in Section 14.1 above.

26.17 17 Conclusion

This paper specifies a seven-gate Formula Validation Protocol — dimensions, scale type, limiting cases, fixed points, identifiability, falsification, reproducibility — ordered by cost, required of every formula in the framework, with each gate's verdict printed beside the formula and any gate not executed reported as **NOT YET RUN** rather than left blank. **Run on the corpus's own central formula M = (D/P)E, the protocol does not pass it.**

The demotion, stated for citation. The moment equation is demoted on gate 2: D and P are not instrumented on a common ratio scale with a shared true zero, so **the magnitude of D/P, the distortion $\Delta = |\ln r|$ as a number, and the quantitative join between the moment equation and the storage law are withdrawn as measured quantities** and relabelled illustrative constructions. Gate 5 returns FAIL as written — two invariance groups, (D,P,E) -> (cD,cP,E) and (D,P,E) -> (kD,P,E/k), so that a large distortion in a thin environment is observationally identical to undistorted contact in a rich one — and passes only under all three remedies together, of which independent measurement of E is essential. Gates 6 and 7 return NOT YET RUN on human data. **Three of seven gates are not passed, and the same protocol applied to the corpus's personal formulary demotes five of its nine entries — F-01, F-02, F-04, F-06 and F-08 — on the same gate** [Curzi, 2026q]. A protocol nothing fails is measuring nothing; these are the rates that make the passed gates worth quoting.

The theorem that forces it. *Theorem 7.1: a quotient X/Y of quantities measured on a common scale is meaningful if and only if the scale's admissible transformation group is the similarity group {x -> ax : a > 0} or a subgroup of it — that is, if and only if the scale is ratio or absolute.* Four demonstrations execute it: an admissible affine relabel from a 1–7 form to a 0–6 form moves a reported D/P from 2.0 to 2.5, a shift of **25 per cent** from the printer's choice alone; g(x) = 2^x **reverses** which of two occasions was the more over-framed; g(x) = x^k delivers **any** target ratio, 2.00 at k = 1.2386 and 100.00 at k = 8.2292; and g(x) = x − 2.5, strictly increasing and therefore admissible, sends a ratio to −0.333 and makes $\Delta = |\ln r|$ **undefined over the reals**.

What survives the demotion is structure, not magnitude. The four limiting cases and their named phenomena (P -> 0 unbounded, dogma and apophenia; D -> 0 zero, sensation without experience; E -> 0 zero, no situation; D = P giving M = E, undistorted contact), the D = P fixed point, the sign of ln r, and the three-regime ordering over-framed / calibrated / under-framed. **Every one of these requires only ordinal information and is invariant under every strictly increasing relabelling.** They support rank statistics, medians, percentiles, sign tests and monotone regression, and they support nothing that multiplies.

The restoration is constructive and it is priced. *Theorem 10.3, given badged Postulate SO1: if* `D` *and* `P` *are measured as Rasch person-measures on one common linked scale, then* $\Delta = |\ln(D/P)| = |\theta_D - \theta_P|$ *is ratio-scaled with a true zero at perfect calibration, and* `D/P` *is determined up to the choice of unit.* The arbitrary origin cancels: at `theta_D = 1.45` and `theta_P = 0.62`, shifting the common origin by `-3.0`, `+2.5` or `+100.0` moves the implied magnitudes across forty-four orders of magnitude and leaves `D/P = 2.293319` and $\Delta = 0.830000$ unchanged to six decimals. The price, from the best-case relation `k = (5.5437/d)^2`, is **approximately 31 items per term to detect a one-logit distortion, 123 for half a logit, 342 for 0.3, 769 for 0.2 and 3,074 for 0.1** — while a twelve-item pair of scales resolves nothing finer than **1.6003 logits**. A reported `r = 1.2` from a twelve-item pair is noise, and the arithmetic says so exactly.

Gate 7 exposes a pooled fiction, and invariance beats fit. Twelve heterogeneous simulated subjects returned a true between-person spread $\tau = 0.2534$ against a pooled standard error of `0.0257`, **a ratio of 9.9**: the pooled interval is about ten times narrower than the actual spread of the individuals it describes, and it contains the target, so a pooled analysis would have reported the division confirmed. It is not a wrong interval; it is a correct interval about the wrong quantity. On an instrument with three planted defects, **split-sample invariance caught three and fit statistics caught two**, the guessing item returning `OUTFIT = 0.936` and `INFIT = 1.000` — the best numbers on the instrument — while the invariance statistic flagged it at `t = +4.14`. On Rasch-generated data two populations separated by 1.82 logits produced item calibrations correlating at `0.994399` with a mean difference of `0.0000`; on 2PL-generated data the mean absolute difference was `7.9` times larger, with every high-discrimination item drifting negative and every low-discrimination item positive, without exception in twelve.

And the falsification gate needs a rival form, not a null. Data generated additively and fitted with the multiplicative log-linear model still returned a held-out `R^2` of **0.6758** — enough to beat persistence, direct-input and a matched AR(1) handsomely and be reported as confirmation — while the correct additive model scored **0.9694**. Only the preregistered rival functional form exposed it. On data generated from the true model the frozen parameters scored `R^2 = 0.8091` held out against a shuffled-input control at `0.0059`; on the missing-division counterfactual the coefficient on the denominator returned a 95 per cent interval of `[-0.2009, 0.0590]` against a target of `-1`, and the kill condition fired on the specific term that was wrong.

None of this is a claim about human beings. The empirical portions are simulated throughout. What the paper establishes is that the corpus's own instrument for auditing its own formulas fires when pointed at the corpus, that the cost of the resulting repair is a finite programme of laboratory work rather than a philosophical impasse, and that the honest current status of the framework's central formula is a heuristic ordering with demonstrated structure — printed as one, in the same paragraph, every time.

26.17.1 Declarations

Data and code availability

No external dataset was collected for this paper, and no human subjects were studied. Every empirical number reported in Section 13 and in Sections 7.10, 7.11, 10.10, 10.11, 10.12, 10.14 and 10.16 is computed from **simulated** data whose generating process is stated beside it — the log-normal moment series with declared autocorrelations, the magnitude-estimation panel with declared free moduli and noise, and the Rasch and 2PL response matrices with declared item parameters and planted defects. The replication counts in Section 3.2 are quoted from Open Science Collaboration [2015] and are that study's, not this paper's. The figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper and the model equations of Sections 4, 7 and 10.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest. The protocol specified here was applied by its own author to his own formulas, which is a structural conflict that no declaration removes; it is the reason failure mode FM-6.8 (Goodhart capture of the protocol) and requirement V-22 are in the specification, and the reason the pass rate is reported as a headline number.

Author contributions

Single author. Michael Laurence Curzi specified the protocol, derived the results, produced the simulated demonstrations and the figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. All response data analysed here are simulated. The laboratory programme described in Sections 13.2 and 13.3 is a design for prospective work and is not a report of any study performed; were it run, it would require ethical review, and nothing in this paper substitutes for that.

Preregistration

This paper reports no preregistered test on human data. The kill criteria KILL-1 to KILL-5 were fixed before the simulated series were generated and are printed in Section 4.6 in the form in which they were fixed; on the true-model demonstration none fired, on the missing-division counterfactual KILL-1 fired, and on the additive counterfactual KILL-1, KILL-2 and KILL-5 all fired. **Those are results about the protocol, obtained on data whose truth was known in advance, and they are not a preregistered empirical test of the moment equation.** The preregistered protocol for testing the storage law against a real system, together with the confirmatory application that **failed** its second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared margin of 10 per cent — is reported in Curzi [2026i]. The falsification suite's own audit, including a demand-characteristic confound with a true coefficient of zero that passed the whole four-baseline suite 52.7 per cent of the time, is reported in Curzi [2026u]. Any applied claim resting on this paper inherits both results.

Chapter 27

The Falsification Suite Measured Against Itself: An Anti-Conservative Shuffle, a 52.7 Per Cent Demand-Characteristic Pass Rate, and a Programme With a 35.8 Per Cent Survival Probability

Abstract. We specify a preregistered falsification suite for the storage law fitted as Y_t+1 = aY_t + bF(U_t) - C_t + eps_t under the constraints 0 <= a <= 1 and b >= 0, with four baselines, a declared freeze, and held-out reporting. The suite is then tested by simulation rather than asserted. Against a nominal 0.05, iid permutation false-rejected at 0.108 and circular shift at 0.120, with only block permutation near nominal at 0.060; the matched autoregressive comparison alone false-passed 35.8 per cent of the time; and a demand-characteristic confound in which subjects infer the hypothesis passed the whole suite 52.7 per cent of the time, for which there is no statistical remedy, only a design remedy. A Goodhart simulation raised the estimated conversion efficiency 14 per cent, from 0.3105 to 0.3551, while the true stock fell. Even with perfect measurement the worked fit recovered a = 0.8833 against 0.85 and b = 0.3340 against 0.40, a 16.5 per cent underestimate. A five-stage research programme is then costed with kill conditions at each gate; the compounded survival probability to a completed final stage is 0.3580, implying 14.93 full-time-equivalent years per completed programme. A final section treats inherited correspondence tables as graphs and audits their category counts, showing that three duplicate rows in seven overstate the surplus ceiling by 40.37 per cent.

Keywords: preregistration; surrogate calibration; shuffled-input control; demand characteristics; Goodhart drift; kill register; ceiling inflation

Series: B. Engineering the Surplus **Paper:** CCP-027

27.1 1 Introduction

A falsification suite that has never been tested is a claim, not a control. This paper specifies the preregistered suite by which the Creature Comforts storage law is made to face a single-person time series on terms declared in advance, and then **runs the suite against itself by simulation**. The headline result is negative and is about the suite rather than about the law.

Against a nominal 0.05, the shuffled-input control false-rejected true nulls at 0.108 for independent permutation and 0.120 for circular shift, with only block permutation near nominal at 0.060. The control on which the framework's empirical audit most depends is therefore **anti-conservative**: run as it is ordinarily run, it rejects a true null between one and two and a half times too often. That number is stated here, in Results, and in the Conclusion, in the same words, because other papers in this corpus cite it to discount their own reported significance [Curzi, 2026i], and a control whose error rate is unpublished cannot be discounted at all.

Three further measurements make the same point from other directions. The matched autoregressive comparison, run alone on data generated with no dependence on the input whatsoever, **false-passed 35.8 per cent of the time.** A demand-characteristic confound – true coefficient exactly zero, with the reported output inflated by a multiple of the input the subject could recall – **passed the entire four-baseline suite 52.7 per cent of the time**, and it does so precisely because the fabrication is temporal and the shuffle is a temporal test. And a Goodhart simulation in which the proxy acquired a consequence at period 30 raised the estimated conversion efficiency by 14 per cent, from 0.3105 to 0.3551, **while the true stock fell from 3.0359 to 2.8447.** In each case the procedure endorsed something that was not there, and in each case the rate was measured rather than assumed.

Nothing here says the suite should not be run. The comparison that matters is between 0.108 and 0.358: a control that is wrong by a factor of two is far better than a single baseline comparison that is wrong by a factor of seven, and both are better than the alternative practice of reporting a parametric p-value that assumes its own null. What the measurements establish is narrower and more useful: **a nominal permutation p-value is not the calibrated one, the calibration is a simulation that takes minutes on the study's own design, and any paper reporting the former as though it were the latter has misstated its evidence.**

The paper has four further contributions. **First**, a full preregistration specification – the required fields, the freeze list, the four baselines and the distinct rival explanation each one names, and the boundary-reporting rule that turns a constrained fit's wall-hit into reportable evidence. **Second**, a quantitative account of instrument noise: rating error on a lagged regressor attenuates the retention coefficient and therefore *inflates* the decay parameter, so that at a rating standard deviation of 0.40 against a signal standard deviation near 0.20 the estimated decay was 0.4317 against a true 0.1500 and the estimated conversion 0.9425 against a true 0.4000. **A sloppy instrument does not produce a noisy answer; it produces a confidently wrong one, and it errs in a known direction. Third**, a five-stage research programme with twenty-one published kill conditions, costed in researcher-hours and calendar months, whose compounded probability of reaching a completed final stage under declared priors is **0.3580** – implying 5.35 full-time-equivalent years per attempt and **14.93 FTE-years per completed programme. Fourth**, an audit instrument for inherited category schemes, which finds that three duplicate rows in a seven-row taxonomy overstate the declared surplus ceiling by **0.559616 nats, 40.37 per cent**, with every downstream quantity inheriting the error at full size.

The last of these carries a second, deflationary measurement. Correspondence systems of the kind that have historically supplied category schemes are treated here as ordinary finite relations: a ten-vertex, twenty-two-edge graph; a seventy-eight-element set with a product decomposition; a map from strings to integers. That last is a hash function, and measuring it as one settles a question the literature reading those systems does not usually ask. On a 234,454-word list the simple letter-sum map has **305 distinct values, 99.995 per cent of words colliding, and 10.468 of the 17.839 bits that identify a word destroyed.** A shared value is the default behaviour of a many-to-one map and carries no evidence about the words.

Sections 3 to 7 give the framework, the suite, its specification, its named failure modes,

and the worked example in which the same unchanged suite produces a confirmation on one series and a refutation on another. Sections 8 to 11 give the staged programme and its price. Sections 12 to 15 give the correspondence-table audit. Sections 16 to 20 report, discuss, verify, limit and conclude.

27.2 2 Hypotheses

This paper is partly a specification and partly a simulation study of that specification. **The specification sections advance no hypothesis and are labelled as specification where they appear.** The hypotheses below are the ones the simulations put at risk. Each carries what would defeat it, and each result is reported in the same plain language whether it passed or failed.

H1 — The shuffled-input control is exact at its nominal level. A permutation test built from the data's own values rejects a true null at the rate it advertises. **Predicted:** across 500 null trials at nominal `alpha = 0.05`, a measured rejection rate near 0.05 for each surrogate family. **Falsified by:** a measured rate materially away from 0.05 in either direction. *(Result: **FALSIFIED** for two of the three families — 0.108 for independent permutation and 0.120 for circular shift; **not falsified** for block permutation at 0.060. See §4.7 and Section 16.)*

H2 — The four-baseline suite is sufficient against a fabricated temporal relationship. A series in which the true input coefficient is zero, but whose reported output has been nudged by what the rater recalls of the input, fails at least one of the four baselines in the large majority of trials. **Predicted:** a suite pass rate well below one half over 300 trials. **Falsified by:** a pass rate at or above one half. *(Result: **FALSIFIED** — 52.7 per cent. The fabrication is temporal, so a temporal test endorses it. See §4.8.)*

H3 — A measure that acquires a consequence leaves a signature in the fitted parameters. When a proxy becomes a target, the estimated conversion efficiency moves even though the constraints continue to hold. **Predicted:** a rise in `eta_hat` across the break, concurrent with a fall in the true stock. **Falsified by:** no change in the fitted parameters across a break at which the proxy acquired a consequence. *(Result: **not falsified** — `eta_hat` rose from 0.3105 to 0.3551 while true `Q` fell from 3.0359 to 2.8447. See §4.12.)*

H4 — Instrument noise biases the two parameters in known and opposite-seeming directions. Error in a lagged regressor attenuates its coefficient, so `delta = 1 − a` is inflated, and the input term compensates, so `eta` is inflated too. **Predicted:** both `delta_hat` and `b_hat` increase monotonically with rating error. **Falsified by:** a non-monotone response, or either parameter moving the other way. *(Result: **not falsified** — `delta_hat` runs 0.1514, 0.1581, 0.1781, 0.2429, 0.4317 and `b_hat` runs 0.4022, 0.4156, 0.4538, 0.5783, 0.9425 across rating standard deviations 0.00 to 0.40. See §4.9.)*

H5 — A single forty-period window recovers the conversion parameter to within its sampling error. With perfect measurement and no rating noise, one window on one subject returns `eta` close to its true value. **Predicted:** `|b_hat − b_true|` of order one standard error or less. **Falsified by:** a deviation of two standard errors or more. *(Result: **FALSIFIED** in the worked draw — `b_hat = 0.3340` against a true 0.40, which is **−2.15 standard errors**, with `a_hat` simultaneously **+2.13**. See §7.3.)*

H6 — A taxonomy's declared cardinality can be audited against its decisions. Two rows are duplicates if no observable item has ever been assigned to one rather than the other and no decision differs between them. **Predicted:** applying the test

to an inherited seven-row scheme merges rows and lowers `ln N`. **Falsified by:** an inherited scheme in which the test can be applied but changes nothing, in a domain where practitioners nonetheless report the rows as indistinguishable. *(Result: **not falsified** on the worked case — seven rows collapse to six once a residual row is added, an 8.60 per cent ceiling overstatement; the textbook case of three duplicates in seven is 40.37 per cent. See §12.11 and Section 15. **This is a demonstration on a constructed case and not a survey of real taxonomies.**)*

Three things in this paper are not hypotheses and are labelled as what they are. The five-stage programme of Section 8 is a **plan with declared priors**, not a result; no stage of it has been run, and its survival arithmetic is arithmetic on those priors and inherits their status entirely. The reading rule of §12.1 – that a correspondence table *is* a finite relation, a graph, a partition, a hash – is a **translation**, not a theory: it cannot be wrong in the way a theory can be wrong, only inapplicable, which happens exactly when the historical artefact is too vague to have a definite extension. And the thesis that the old notation and the framework's notation describe the same relations is a **postulate**, badged as one wherever it appears, and is a claim about *structural correspondence* and never about mechanism or efficacy.

27.3 3 Framework: the estimation model, its postulates, and its identification problem

The framework's minimal estimation model, reproduced exactly from the group-scope protocol [Curzi, 2026i]:

```
MINIMAL ESTIMATION MODEL

    Y_{t+1} = a * Y_t  +  b * F(U_t)  -  Chat_t  +  eps_t

CONSTRAINTS

    0 <= a <= 1        (retention cannot exceed perfect retention,
                        and cannot be negative)
    b >= 0             (conversion efficiency cannot be negative)

IDENTIFICATION

    a = 1 - delta      ->   delta_hat = 1 - a_hat
    b = eta            ->   eta_hat   = b_hat

WHERE

    Y_t     an observed proxy for stored order Q_t
    U_t     the observed non-parallelism, in [0,1]
    F(U)    = ln(1 + (N-1)U),  the surplus law
    Chat_t  an independently measured cost series, NOT a free parameter
    eps_t   residual
```

The framework's falsification template, reproduced exactly [Curzi, 2026i]:

```
PROCEDURE

    1.  PREREGISTER  the operationalisations, the window, the split,
        the baselines, and the falsification rules.
    2.  FIT          on the training window only.
```

```
    3.  FREEZE       every choice. No refitting, no re-specification,
                     no re-selection of proxy, window, or period length.
    4.  REPORT       held-out prediction against four baselines:
                     (i)   persistence
                     (ii)  direct-input (raw U, no log)
                     (iii) matched autoregressive, no interaction term
                     (iv)  SHUFFLED-INPUT CONTROL

FALSIFIED IF

    RULE 1.  the shuffled-input control predicts held-out Y as well as
             the ordered series;
    RULE 2.  the matched autoregressive baseline is not beaten;
    RULE 3.  the fitted a lies outside [0,1] or the fitted b is negative.
```

Section 4 applies that template to **subjective** quantities: to `Y` read off a magnitude estimate rather than a headcount, to `U` computed over one person's own framings rather than over a committee, and to the Moment `M_t` standing in the input slot in place of `F(U_t)` by the normalisation declared by the framework (`M` mapped into `[0, ln N]` serves as `S_t`).

Nothing in the template changes when the quantities become subjective. Everything in the threat model does. That asymmetry organises Sections 4 to 7.

27.3.1 3.1 What is proven and what is postulated

Proven. The identification map is exact algebra. Given the storage law `Q_{t+1} = (1-delta)Q_t + eta*S_t - C_t` [Curzi, 2026a] and given `Y_t = Q_t`, matching coefficients gives `a = 1-delta` and `b = eta` with no residual freedom. The constraint set follows from the definitions: `delta` is a retention loss in `[0,1]`, so `a = 1-delta` is in `[0,1]`; `eta` is a conversion efficiency and is non-negative by construction. The bound `0 <= F(U) <= ln N` is the framework's logarithmic ceiling [Curzi, 2026c]. The exact solution, the one-step growth condition, and the steady state `Q_inf = (eta*F(u*) - c*)/delta` are theorems of the leaky-integrator analysis [Curzi, 2026a] and are used here only as consequences of the fitted parameters, never as evidence for them.

Postulated, badged. That `Y_t` stands in a stable, monotone, and ideally affine relation to `Q_t`. That `delta`, `eta`, `N`, and the operationalisation of `u` are approximately constant over the estimation window. That the relation is linear in `Y_t` and in `F(U_t)` at all. That the subject's rating scale is ratio-scaled in the sense of the formula-validation protocol's scale-type gate and stable in the sense of its specific-objectivity requirement [Curzi, 2026t]. **Each of these is a modelling assumption, each is falsifiable by the procedure specified below, and none of them is a theorem.** The framework's falsification gate – FVP-6 in the numbering of its formula-validation protocol [Curzi, 2026t] – exists precisely so that these four are stated in print before any data are collected.

The gate numbering used throughout this paper is the formula-validation protocol's [Curzi, 2026t] and is glossed once here, so that no later reference is opaque: FVP-1 is dimensional consistency, FVP-2 scale type – the gate that licenses division, and the gate on which that protocol demotes the framework's own moment equation – FVP-3 limiting cases, FVP-5 identifiability, FVP-6 falsification, and FVP-7 reproducibility across investigators.

Honesty clauses, repeated. The triangle inequality fails for geometric triples in the framework's own statement of the surplus functional [Curzi, 2026]; subadditivity holds on `u` only. As deployed, the surplus functional is a LOWER BOUND on the ensemble's information — we log the floor, not the ceiling. The surplus functional is a mathematical object and carries no empirical or metaphysical claim; the framework does not identify surplus with Shannon entropy or mutual information; and it proves no thermodynamic negentropy law, since `F(u) + N_can(u) <= ln N` and higher static surplus does not imply higher canonical concentration [Curzi, 2026,c].

27.3.2 3.2 What changes when Y is subjective

Four things, and they are the reason the single-person scope needs its own falsification apparatus rather than a cross-reference to the group-scope survey and preregistration protocol [Curzi, 2026i].

```
1. THE INSTRUMENT IS THE SUBJECT.
   At group scope the rater and the rated were different objects. Internally they are
   the same organism. The person who supplies Y_t also has a view about whether the
   theory is true.

2. THE RATING CARRIES ITS OWN NOISE, AND THAT NOISE BIASES delta UPWARD.
   Measurement error on a LAGGED regressor attenuates a_hat. Since delta = 1 - a,
   attenuation on a inflates delta. §4.9 computes the size.

3. THE INPUT AND THE OUTPUT ARE OFTEN RECORDED IN THE SAME SITTING.
   Anything that moves both at once — mood, fatigue, the weather, the fact of being
   measured — is a common cause the shuffle cannot see.

4. MEMORY RECONSTRUCTS.
   A retrospective rating of Y_{t+1} is partly a rating of what the subject believes
   happened at t. That is a lagged dependence with the right sign and the right timing,
   and it is not the dependence the theory predicts.
```

Threat 4 is the dangerous one, because **it produces exactly the temporal structure the shuffled control is designed to detect, and therefore passes the control.** §4.8 measures how often. The answer, in the configuration simulated there, is 52.7 per cent of the time.

27.3.3 3.3 The three rules, restated for the single-person instrument

```
RULE 1  (SHUFFLE)     If the input series randomly permuted in time predicts held-out
                      Y as well as the input series in its true order, there is no
                      temporal relationship. Report that.

RULE 2  (AUTOREGRESSION)  If a model containing only Y_t and an intercept predicts
                      held-out Y as well, then F(U_t) contributed nothing and the
                      surplus term is decoration.

RULE 3  (CONSTRAINTS) If a_hat leaves [0,1] or b_hat goes negative, the fitted object
                      is not a storage law. It is a regression with the same letters.
```

A framework that cannot fail cannot be trusted, and this one names exactly how it fails. The three rules are not hedges. They are the price of admission.

27.3.4 3.4 Identification: the objection no calibration answers

Everything in Sections 4 to 7 is about whether a fitted coefficient is distinguishable from noise. This subsection is about something logically prior, and it is the more serious problem.

The framework's central formal claim is a reflexive loop, `T(x) = W(x, M(x))`: the group influences the member and the members constitute the group [Curzi, 2026]. Where the storage law is applied to a person embedded in a group – which is the case throughout Stage 5 of Section 8, and is the case for the internal-external correspondence the programme's third test is built to check – that is a linear-in-means reflexive specification, and Manski [1993] proves that its endogenous social effect is **not identified from observational data**. It cannot be separated

from contextual effects or from correlated unobservables, however much data is collected. Shalizi and Thomas [2011] reinforce the same point for networks: influence and selection are generically confounded in observational data, and Lyons [2011] shows what happens when the confound is not taken seriously.

This paper adopts no remedy, and says so. The four available routes are randomisation, an instrument, partial-population variation, and stipulating the reference group and defending the stipulation; the last is the most realistic here and has a template [Brock and Durlauf, 2001]. Stage 5 of Section 8 is cluster-randomised and would, if it were run, supply the first of them for the intervention test `T2` – but not for the prediction test `T1`, which is observational by construction, and not at all for Stages 3 and 4, which are observational throughout.

The consequence is precise and it is not softened elsewhere in this paper. Even a fit clearing all three falsification rules, with a calibrated surrogate threshold and a perfectly reliable instrument, would not have identified `eta` as the conversion of interaction into stored order, because a contextual effect or a common unobserved driver produces the same fit. Wherever this paper speaks of `delta` and `eta` as person-level or group-level quantities, **the claim is a modelling position rather than an identified empirical result.** That is a defensible thing to write; an unidentified parameter presented as measured is not.

Two objections in the same family are recorded here and are not answered either. The first is the Lucas critique applied to the framework itself [Lucas, 1976]: the framework asserts that systems change their behaviour when measured – §4.12 measures exactly that – and then treats `delta`, `eta` and `N` as constant over the estimation window. The position taken here is the narrow one: constancy is assumed **only over a declared window and only while `Y` is not itself reported upward**, which is why the shadow proxy of §4.12 is the first and strongest corruption defence. Where `Y` is reported upward, the framework's own reflexivity applies to it and the constant-parameter postulate fails; that is a restriction on where the protocol may be used and is stated as one. The second is that the functional form of the input term is underdetermined by the evidence: social impact rising as a sub-linear power of the number of sources [Latané, 1981] and value growing as `n log n` [Briscoe et al., 2006] fit the same qualitative facts as `ln(1 + (N−1)u)`, and **nothing in this paper discriminates between them.** The strongest available warrant for the logarithm is not the psychophysical one at all: Peters and Gell-Mann [2016] derive a logarithmic form from multiplicative dynamics with no psychological assumption, which is what the argument here actually needs.

27.4 4 The falsification suite

27.4.1 4.1 The model, slot by slot

What goes in each slot when the scope is internal:

Slot	Group scope	Single-person scope
`Y_t`	institutional order proxy	character, skill, or internalised structure, via magnitude estimate
`U_t`	non-parallelism between members	non-parallelism between one person's own framings
`F(U_t)`	group surplus `ln(1+(N−1)U)`	internal surplus `F_int`
`M_t`	the situation as constructed	the moment as lived, `(D/P)E`
`Chat_t`	coordination overhead	cognitive load, effort, maintenance of self

Slot	Group scope	Single-person scope
`N`	distinct domains held in a group	distinguishable categories in one mind

Two admissible input series exist internally, and they must not be mixed inside one preregistered model:

```
INPUT VARIANT A   S_t = F(U_t) = ln(1 + (N−1) U_t)
                  U_t measured over one person's own framings (the audited framing set)

INPUT VARIANT B   S_t = M_t normalised into [0, ln N]   (declared normalisation)
                  M_t = (D_t / P_t) ∗ E_t                (the moment equation)
```

Variant B inherits the whole of the scale-type gate [Curzi, 2026t]: the corpus's own formula-validation protocol DEMOTES the moment equation on that gate, because `D` and `P` are not instrumented on a common ratio scale with a shared true zero. **If `D` and `P` are not ratio-scaled with a true zero, `D/P` is not a number and Variant B is not a model.** The preregistration must state which variant is being tested, and it must state the evidence that the scale requirement is met, or demote the formula in print as FVP-2 requires.

`Chat_t` is not a free parameter. It is an independently measured series, subtracted before fitting. **Allowing `C` to be estimated inside the regression is the single most common way this model is made unfalsifiable**, because a free additive term with an intercept and a trend will absorb any misfit that has the shape of a level or a drift. Fit two coefficients. Measure the third.

27.4.2 4.2 The constraint set is the first falsification rule, not a convenience

It is tempting to fit by ordinary least squares and then note with satisfaction that `a_hat` came out inside `[0,1]`. That gets the logic backwards. The constraints are the theory. `a` outside `[0,1]` means the system amplifies its own stored order period on period without input, which is not a leaky integrator but an oscillator or a divergence. `b < 0` means interaction destroys stored order, which contradicts `eta >= 0` in the definition of the storage law. **Fit the model constrained; report the unconstrained fit alongside it; and if the unconstrained fit lands outside the box, say so in the abstract, because that is Rule 3 firing.**

There is a practical trap. A constrained fit cannot report a violation it is not allowed to express. If `a_hat` sits exactly at the boundary — `a_hat = 1.000` or `b_hat = 0.000` — the optimiser has hit the wall, and a wall-hit is evidence, not a result. In the null simulation of §4.7 the constrained fit returned `a_hat = 1.0000` exactly: retention pinned at unity, no decay at all, which is the fit telling you it wanted to leave the box.

```
BOUNDARY REPORTING RULE

    Report a_hat and b_hat from the CONSTRAINED fit.
    Report a_ols and b_ols from the UNCONSTRAINED fit.
    Flag in bold any of:
        a_hat = 0 or a_hat = 1 exactly
        b_hat = 0 exactly
        a_ols outside [0,1]
        b_ols < 0
    A flagged boundary is a Rule 3 event and must appear in the abstract.
```

27.4.3 4.3 The preregistration document

Preregistration is a device that removes a specific, enumerable list of choices from the researcher's hands after the data arrive [Nosek et al., 2018]. The list is not hypothetical: undeclared analytic flexibility is sufficient on its own to manufacture significance from noise [Simmons et al., 2011], and it does so without any dishonest intention, because the branching happens one defensible decision at a time [Gelman and Loken, 2013]. Each line below is a degree of freedom that, left open, can manufacture a result.

```
PREREGISTRATION — REQUIRED FIELDS

A. SUBJECT AND SCOPE
   A1  scope: internal (single person) / dyadic / group
   A2  number of subjects, and whether analysis is per-subject or pooled
   A3  inclusion and exclusion rules, stated before recruitment

B. THE PERIOD
   B1  period length (a day, a week, a fortnight) and the justification
   B2  the alignment rule: what counts as period t for Y, for U, for C
   B3  the handling of missed periods (interpolate? drop? carry forward?)

C. OPERATIONALISATIONS
   C1  Y: instrument, items, scoring, and the argument that Y is not a
       repackaging of U -- the most common measurement error \citep{ccp:CCP-015}
   C2  U: the framing set, the inner-product definition, the computation of u
   C3  N: how the category count was audited, not guessed \citep{ccp:CCP-009}
   C4  C: the cost instrument and its units, in the same units as eta*S
   C5  input variant A or B, and for B the ratio-scale evidence \citep{ccp:CCP-026}

D. THE SPLIT
   D1  T_train and T_holdout, fixed in advance as counts, not as fractions
       of a total that is not yet known
   D2  the split is CHRONOLOGICAL. No random splits on time series.
   D3  no peeking rule: who holds the held-out data and when it is opened

E. THE BASELINES
   E1  persistence
   E2  direct-input (raw U, no log), with the required u-range (§4.5.2)
   E3  matched autoregressive with intercept
   E4  shuffled-input control: family, B, and the calibration (§4.7)

F. THE FALSIFICATION RULES
   F1  Rule 1 threshold, as a calibrated p-value, not a nominal one
   F2  Rule 2 threshold, as a skill score with its sign convention
   F3  Rule 3 boundary reporting rule

G. THE STOPPING RULE
   G1  when data collection ends, stated as a date or a count, never as
       "when the result stabilises"

H. THE PUBLICATION COMMITMENT
   H1  the commitment to report the result whichever way it falls
   H2  the deposit location for the frozen model and the analysis code
```

A stopping rule that reads "until the pattern is clear" is not a stopping rule; it is optional stopping, and optional stopping will eventually produce any pattern you like [Simmons et al.,

2011]. Fix the count. Fix the date. Publish whichever way it falls.

27.4.4 4.4 The freeze

“Freeze” means: after the training fit, the following are constants and may not be touched again.

```
FROZEN AT THE END OF TRAINING

    a_hat, b_hat                    the two coefficients
    the operationalisation of Y     items, scoring, reverse-keying
    the operationalisation of U     framing set, inner product
    N                               the category count
    the cost instrument             and its units
    the period length               and the alignment rule
    the missing-data rule
    the input variant               A or B
    the transformation of Y         log, none, winsorisation limits
    the baseline definitions        all four
    the surrogate family and B      for the shuffle
    the skill-score definition
    the falsification thresholds
```

The list is long because researcher degrees of freedom are numerous and each is individually defensible. **Every item on it has a reasonable-sounding argument for changing it after the data are seen, and that is exactly why it must be frozen before.**

One permitted exception, declared in advance: a data-quality exclusion rule that does not reference `Y`, `U`, or the fit. “Periods in which the subject was hospitalised are excluded” is legitimate if written in field A3 before collection. “Periods in which the residual exceeded two standard deviations are excluded” is not legitimate under any circumstances.

27.4.5 4.5 The four baselines

A baseline is a rival explanation with a name. Each of the four names a specific way the result could be uninteresting.

4.5.1 Persistence

```
Y_{t+1} = Y_t
```

No parameters. Rival explanation: *the thing you measured is simply slow-moving.* Persistence is the hardest baseline to beat for any highly autocorrelated quantity, which is most psychological constructs at daily or weekly cadence. It also tells you whether your period length is sensible: **if persistence is unbeatable, your period is shorter than the dynamics you are trying to observe, and the fix is a longer period, not a better model.** In the worked example it gives a held-out RMSE of 0.1207 against the model’s 0.0777, a skill score of 0.5860.

4.5.2 Direct-input

```
Y_{t+1} = a * Y_t + b * U_t - Chat_t
```

Raw `U`, no logarithm. Rival explanation: *the logarithm in `F` is doing no work; any increasing function of `U` would fit as well.*

This baseline tests the surplus functional's form [Curzi, 2026], and it is the weakest of the four for a reason worth computing. Over `u` in `[0,1]` with `N = 7`, the correlation between `u` and `F(u)` is 0.9724. The best straight-line approximation is `F(u) ~= 1.7280 u + 0.4063` with a maximum absolute error of 0.4021 and an RMS error of 0.1197 — visible over the full range. But over any narrow band the two are nearly identical:

u range	corr(u, F(u))
[0.0, 0.2]	0.99496
[0.2, 0.5]	0.99704
[0.5, 1.0]	0.99741
[0.3, 0.6]	0.99795

The direct-input baseline has discriminating power only in proportion to the range of `u` your data actually span, and most field data span very little of it. Simulated at `T = 100` with 70 training periods, the fraction of trials in which the log form beats the linear form on held-out data rises with the observed span of `u`:

mean span of u	log form wins
0.323	84.0 %
0.580	91.3 %
0.862	97.7 %
0.987	100.0 %
1.000	100.0 %

and the ratio of linear-model to log-model held-out RMSE rises from 1.063 at span 0.338 to 1.606 at span 0.987.

The consequence for preregistration is concrete. Field E2 must state a **minimum observed span of `u`** below which the direct-input comparison is declared uninformative and reported as such. A defensible default from the table above is a span of 0.60. Below that, say in print: *the data do not distinguish the logarithmic from the linear input form.* That is an honest partial result and it is much better than a claim the data cannot support.

4.5.3 Matched autoregressive

```
Y_{t+1} = a * Y_t + c - Chat_t
```

An intercept in place of the surplus term. Rival explanation: *the system decays toward a constant and the interaction variable is irrelevant.* This baseline is the statistical twin of the storage law itself: a first-order linear difference equation with a constant term is exponential smoothing [Holt, 1957, Elaydi, 2005], and Brown and Meyer [1961] state exactly which class of processes it is optimal for – a question the storage law does not ask about itself. The fixed-gain limit of the Kalman filter is the same object again [Kalman, 1960]. This baseline isolates the theory's actual content, because the storage law without `eta*S_t` is exponential decay to a level. It is also the baseline most likely to win when the input barely varies: **if `F(U_t)` is nearly constant, its coefficient is unidentifiable and an intercept does the same job at lower cost.** In the null simulation of §4.7 Rule 2 fires — model RMSE 0.0978 against matched-AR 0.0681, skill -1.0669.

But Rule 2 alone is nowhere near sufficient. In 500 simulated null trials — data generated with no dependence on `S` at all — **the model beat the matched autoregressive baseline on held-out data 35.8 per cent of the time.** A single baseline comparison, run once, on data where both series are smooth, is a coin weighted only slightly against you.

4.5.4 The shuffled-input control

```
Take the input series S = (S_0, S_1, ..., S_T).
Permute it in time to get S' = (S_{pi(0)}, S_{pi(1)}, ..., S_{pi(T)}).
Refit on training with S'. Score on held-out with a permuted held-out S'.
Repeat B times. Compare the ordered model's held-out error to the
distribution of the permuted models' held-out errors.
```

Rival explanation: *the input and the output have the right marginal distributions and no temporal relationship at all.*

The shuffle preserves the marginal distribution of the input exactly — every value, its mean, its variance, its skew, its range — and destroys only the pairing between input and output in time. That is the precise null it instantiates, and it is the reason it is the sharpest of the four.

Everything else that could be responsible for a fit survives the shuffle: the scale of `S`, the number of large values, the fact that `S` is bounded by `ln N`, the fact that `Y` trends, the fact that both are positive. The only thing the shuffle removes is *which* `S` went with *which* `Y`. **If the fit survives that removal, the relationship you found was a property of the two marginal distributions and not a relationship between two series.**

27.4.6 4.6 Why the shuffled control is the one most often omitted

The other three baselines are, in practice, strawmen: the researcher expects to beat persistence (the construct moves), direct-input (the log is in the theory), and matched autoregression (the input term is in the theory). Running them feels like diligence and usually confirms what was already believed.

The shuffle is different in kind. **It is the only one of the four that can kill a result the researcher already believes, using data the researcher already has, at no additional cost, in seconds of computation.** That combination — decisive, cheap, already possible — is what produces omission, because there is no resource-based excuse to hide behind. Its omission is a choice about what one wishes to know.

Three reasons it gets left out, in the order of how often they are offered:

```
1. "It is not standard in this literature."
   Irrelevant. A permutation test needs a random number generator, not a
   literature's permission.

2. "The input is autocorrelated so the shuffle is not a fair comparison."
   A real technical objection, answered by choosing the right surrogate
   family (§4.7), not by dropping the control.

3. "We already showed the effect is significant."
   Significance against a parametric null assumes the null. The shuffle
   CONSTRUCTS the null from the data you have. When the two disagree,
   believe the shuffle: it assumed less.
```

There is a fourth reason, unstated: the shuffle is the only baseline whose result the researcher cannot predict in advance. **A test whose outcome you can predict is not a test.**

27.4.7 4.7 The surrogate hierarchy, and why you must calibrate the control itself

The objection in reason 2 is correct, and taking it seriously produces the most technically important subsection of the paper.

An independent-and-identically-distributed permutation destroys two things at once: the pairing between `S` and `Y` (which you want destroyed) and the autocorrelation within `S` (which you did not). If both `S` and `Y` are smooth, a smooth ordered `S` can track a smooth `Y` by coincidence, while the whitened shuffled `S` cannot track anything. **The plain shuffle can therefore reject the null when the null is true, because it compares a smooth regressor against a white one rather than against another smooth one.** This is the spurious-regression problem wearing a permutation test's clothes, and it is the same statistical hazard that makes a straight line on a log-log plot poor evidence for a mechanism [Clauset et al., 2009, Stumpf and Porter, 2012]: a procedure that has not been calibrated against its own null is not a test.

Three surrogate families, in increasing fidelity to the input's own structure:

```
IID PERMUTATION      S' = random permutation of S
                     preserves: the multiset of values, hence every marginal moment
                     destroys:  pairing AND autocorrelation

BLOCK PERMUTATION    cut S into contiguous blocks of length L; permute the blocks
                     preserves: marginals, and autocorrelation at lags < L
                     destroys:  pairing, and autocorrelation at lags >= L
                     choose L >= 3 / (-ln rho1) so a block spans the input's own
                     correlation time

CIRCULAR SHIFT       S' = S rotated by a random offset k
                     preserves: marginals AND the full autocorrelation function
                     destroys:  pairing only
                     the tightest surrogate; the number of distinct surrogates is
                     limited to the series length, which caps resolution
```

Now the measurement that decides the matter. Five hundred null trials were simulated: `Y` generated by a storage law driven by a smooth AR(1) input that is *not* `S`, with `S` an independent smooth AR(1) series (autocorrelation parameter 0.85 for both), `T_train = 40`, `T_holdout = 20`, 200 surrogates per trial. The true `b` is zero. A well-calibrated test at nominal alpha = 0.05 should reject 5 per cent of the time.

```
NULL (b_true = 0), 500 trials, nominal alpha = 0.05

   surrogate family      measured rejection rate
   iid permutation              0.108
   circular shift               0.120
   block permutation (L=5)      0.060

   matched-AR comparison alone: passes 0.358 of the time
   mean b_hat under the null:   0.0922  (sd 0.0222)

TRUE MODEL (b_true = 0.40), 500 trials

   surrogate family      measured rejection rate (= power)
   iid permutation              1.000
   circular shift               1.000
   block permutation (L=5)      0.998

   beats matched AR on held-out: 1.000
   mean b_hat:                   0.4025  (sd 0.0322)
```

Read those two tables together and three conclusions follow, none of which is optional.

First: the shuffled-input control, run as it is usually run, is anti-conservative. At a nominal 5 per cent it rejected between 6.0 and 12.0 per cent of true nulls in this design. A control that

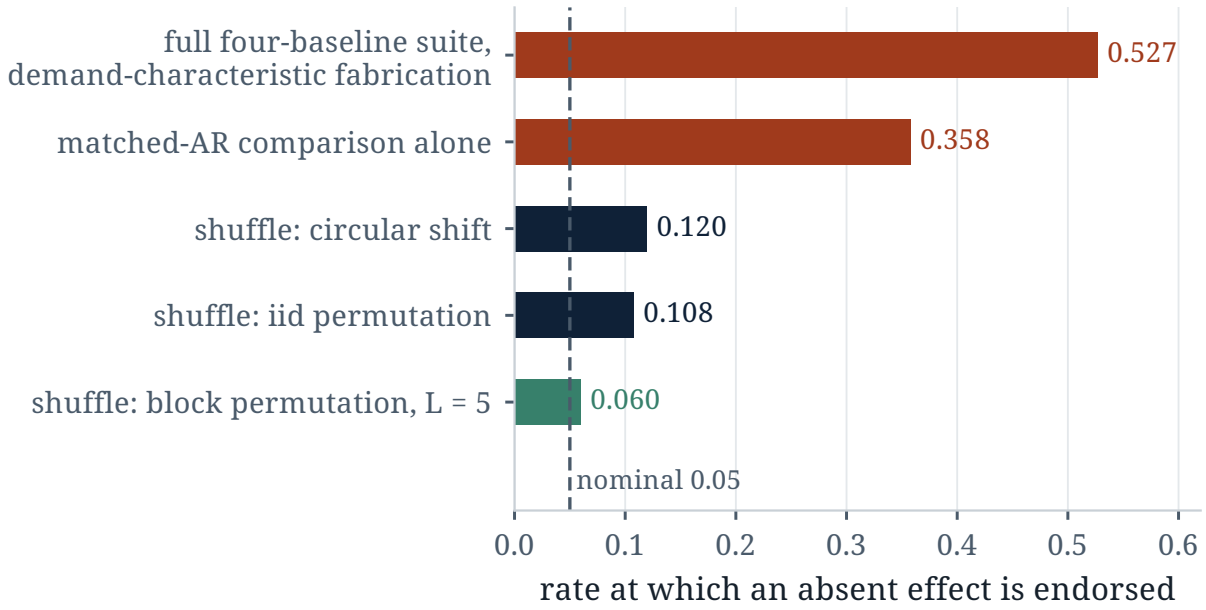


Figure 27.1: **Every bar is the rate at which the procedure endorses an effect that is not there, and the smallest of them is above the nominal 0.05.** The three shuffle families are measured over 500 null trials with 200 surrogates each; the matched-AR figure is the same 500 trials scored on that baseline alone; the 52.7 per cent is a separate 300-trial run in which the true coefficient is zero and the reported output has been inflated by 0.45 times the centred input. Modelled, with the generating parameters named in §4.7 and §4.8; produced by `figs/src/CCP-027_fig01.py`.

is wrong by a factor of two in the direction of false discovery is still enormously better than no control — compare the 35.8 per cent false pass rate of the matched-AR comparison alone — but it is not exact, and reporting its nominal p-value as though it were exact is a misstatement.

Second: the block permutation was the best calibrated of the three here (0.060 against a nominal 0.050), and it lost almost nothing in power (0.998 against 1.000). That is a favourable trade and it should be the default. Note that the circular shift, which preserves the most structure, was *not* the best calibrated in this design, because the held-out window is short and rotating a 21-point segment leaves few genuinely distinct surrogates. Intuition about which surrogate is "tightest" does not survive contact with a short held-out window.

This is not the only calibrated gate in the corpus that has come back above nominal. A pre-registration for narrative regulatory input measures its conjunction gate at a false-positive rate of 9.5 per cent at `T = 60` against a nominal 5, and repairs it by requiring cross-individual replication at 4 of 12, which restores `alpha = 0.0216` while retaining 90.8 per cent power against a half-size effect [Curzi, 2026ad]. **That repair – a structural requirement rather than a larger `B` – is the model this paper's S4.10 follows.** The general lesson is the one measured on the integrity side of the corpus, where a naive detector's precision was 11.01 per cent, right one time in nine, until typed-field parsing collapsed its false-positive share from 88.99 per cent to 0.00 per cent [Curzi, 2026]: **a control is not trustworthy because it is standard; it is trustworthy when its error rate has been measured.**

Third, and this is the specification item: calibrate the control on your own design before you trust its p-value. The false-positive rate of a surrogate test depends on the autocorrelation of both series, the length of the training and held-out windows, the number of surrogates, and the constraint set. None of those is universal. The calibration is a simulation you run once, in advance, using the preregistered design and a generating process with `b` set to zero, and it takes minutes. The output is a *calibrated* critical value that replaces the nominal one in field F1.

Finally, resolution. A permutation p-value cannot be smaller than `1/(B+1)`:

```
   B        smallest attainable p     se of p at p = 0.05
  19             0.05000                   0.05000
  99             0.01000                   0.02190
 199             0.00500                   0.01545
 999             0.00100                   0.00690
1999             0.00050                   0.00487
```

`B = 199` is the practical floor for a claim at the 1 per cent level; `B = 999` if the claim is stronger. Report `B`. A shuffle test whose `B` is not reported is not reproducible.

27.4.8 4.8 What the shuffled control does not catch

Here is the honest limit, and it is severe.

The shuffle detects the absence of a temporal relationship. It cannot distinguish a real temporal relationship from a fabricated one that has the same timing. Any process that writes the input's timing into the output — by any route other than the theorised one — will produce a series that passes the shuffle handsomely.

The mechanism that does this most reliably is threat 4 of §3.2: a subject who knows the hypothesis and whose rating of `Y_{t+1}` is nudged by what they recall of period `t`. That subjects and experimenters both move measurements toward an expected result is among the most robustly documented effects in experimental psychology [Rosenthal and Fode, 1963, Rosenthal and Jacobson, 1968]; the reactivity of measurement itself is weaker and more heterogeneous than it is usually asserted to be [McCambridge et al., 2014], but demand characteristics of the kind simulated here do not depend on reactivity, only on the subject's inference of the hypothesis. Simulated directly: data generated with `b_true = 0` — no real dependence whatsoever — and the *reported* `Y_{t+1}` inflated by `0.45 * (S_t - mean S)`.

```
DEMAND-CHARACTERISTIC SIMULATION, true b = 0, T_train = 40, T_holdout = 20

  example trial   fitted a = 0.9265   fitted b = 0.1856   (both inside the box)
                  held-out RMSE  model 0.0991
                                 matched AR 0.1112
                                 circular-shift surrogate mean 0.1168
                  shuffle p = 0.0000

  over 300 trials, the full four-baseline suite PASSED a purely fabricated
  relationship 52.7 % of the time
```

More than half the time, a relationship that does not exist cleared all four baselines and all three rules, because the rater's own bias supplied the timing. No amount of extra shuffling fixes this. The shuffle is testing the wrong thing: the fabrication is temporal, so a temporal test endorses it.

The fix is not statistical. It is a design change, and there are four, in decreasing order of strength:

```
DC1 BLIND THE RATER TO THE HYPOTHESIS.
    The subject supplies Y without being told what U is, that U is being
    measured, or what relationship is predicted. Debrief afterwards.

DC2 MEASURE Y AND U THROUGH DIFFERENT CHANNELS.
    Y from a behavioural or archival trace (a completed artefact, an
    attendance record, a produced object); U from the rating instrument,
```

```
    or the other way round. A bias in one channel cannot supply timing
    to the other.

DC3 PROSPECTIVE, NOT RETROSPECTIVE, SAMPLING.
    Rate the period in the period, not at the end of the week. Memory
    reconstruction is proportional to the recall interval; shrink the
    interval and you shrink the artefact.

DC4 YOKED CONTROL SUBJECT.
    A second subject is shown the first subject's U series, on the same
    schedule, but their own Y is their own. If the yoked subject shows
    the same fit, the fit is being produced by the SCHEDULE, not by the
    subject's own dynamics.
```

DC1 and DC2 together are close to sufficient. **State in the preregistration which of DC1-DC4 are in force. A study running none of them cannot claim a temporal result about subjective quantities regardless of how many surrogates it ran.**

The group-scope protocol [Curzi, 2026i] treats the shuffle as the decisive baseline, and at group scope with archival proxies it very nearly is. Turned inward, **the shuffle remains the most decisive of the four and stops being sufficient on its own.** The second half is what the single-person scope adds.

27.4.9 4.9 The instrument's own noise, and which direction it biases

A subjective `Y` is a rating, and a rating has error. That error does not merely add scatter; it biases both parameters, in known directions, by computable amounts.

The mechanism is regression dilution applied to a lagged regressor. Error in `Y_t` attenuates `a_hat` toward zero, and since `delta_hat = 1 − a_hat`, **attenuation of `a` inflates `delta`: measurement noise makes a system look leakier than it is.** Because the lagged term then explains less, the input term compensates, so `b_hat` — hence `eta_hat` — is inflated too.

Simulated at `T = 80`, true `a = 0.85`, true `b = 0.40`, process noise 0.06, 400 replications per row:

rating sd	mean a_hat	mean b_hat	implied delta_hat
0.00	0.8486	0.4022	0.1514
0.05	0.8419	0.4156	0.1581
0.10	0.8219	0.4538	0.1781
0.20	0.7571	0.5783	0.2429
0.40	0.5683	0.9425	0.4317

Read the last row. With a rating standard deviation of 0.40 against a `Y` whose own standard deviation is around 0.20, the estimated `delta` is 0.432 against a true 0.150 — nearly triple — and the estimated `eta` is 0.943 against a true 0.400, more than double. **A sloppy instrument does not produce a noisy answer. It produces a confidently wrong one, and it errs in the direction of making the system look both leakier and more efficient than it is.**

This is why Stage 2 of the programme in Section 8 comes before Stage 3 and cannot be skipped: the scale must be built and its error characterised before any dynamic parameter estimated on it means anything. It is also why the preregistration must report the instrument's test-retest reliability alongside the parameter estimates. **A parameter estimate without an accompanying reliability figure for the instrument that produced it is uninterpretable, because the reader cannot tell how much of the reported `delta` is the world and how much is the ruler.**

A partial correction exists — an attenuation factor estimated from a repeated-measures reliability study — but it depends on that estimate being itself accurate and does nothing for `b_hat` without further assumptions. **Prefer a better instrument to a correction factor.**

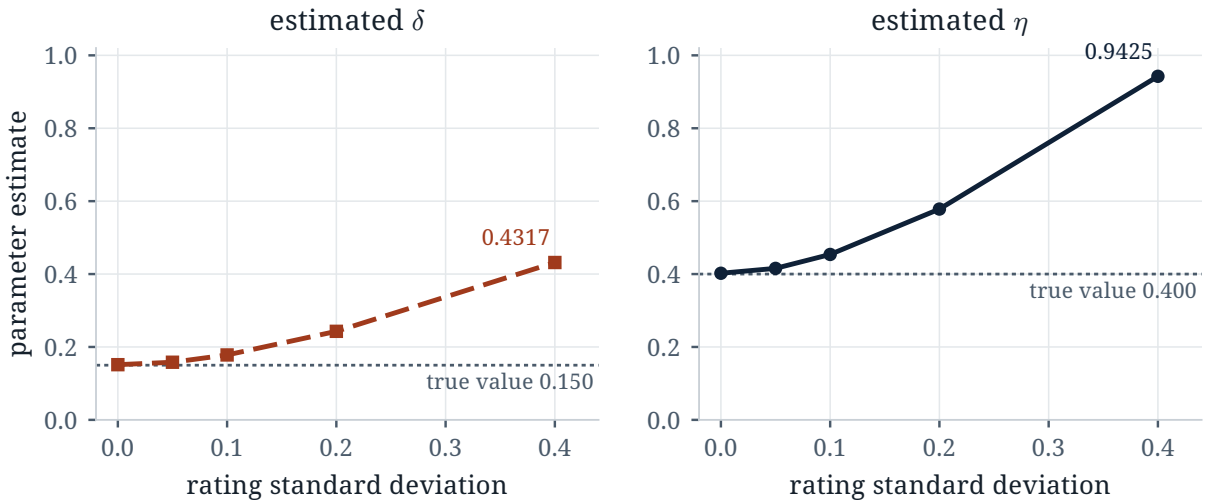


Figure 27.2: **More rating error does not scatter the estimates; it walks them away from the truth in a fixed direction.** Decay is over-estimated and conversion is over-estimated together, so a noisier instrument makes a system look both leakier and more efficient than it is. Modelled at `T = 80`, true `a = 0.85`, true `b = 0.40`, process noise 0.06, 400 replications per point; produced by `figs/src/CCP-027_fig02.py`.

27.4.10 4.10 Held-out reporting

Report one table. It has five rows and it is the entire empirical claim.

```
HELD-OUT REPORT (the only table that matters)

  predictor                       RMSE      skill vs model      note
  ------------------------------------------------------------------------
  frozen model                    ____          --
  persistence                     ____        ____
  direct-input (raw u)            ____        ____            span of u = ___
  matched autoregressive          ____        ____
  shuffled input (family, B)      ____        ____            calibrated p = ___
```

Definitions, fixed in advance:

```
  RMSE          root mean squared ONE-STEP-AHEAD error on the held-out window,
                using OBSERVED Y_t on the right-hand side, never simulated Y_t

  skill score   1 - (RMSE_model / RMSE_baseline)^2
                positive means the model beats the baseline;
                zero means they tie; negative means the baseline wins
```

Two prohibitions that matter more than they look.

Prohibition one: one-step-ahead scoring, using observed `Y_t`. Free-running simulation from `Y_0` is a weaker test that mostly measures whether `Q_inf` was estimated well; as a primary score it inflates apparent performance for any model with a stable fixed point. Secondary diagnostic only.

Prohibition two: no refitting. If the held-out window is opened and the model performs poorly, the result is that the model performs poorly. A revised model fitted after seeing held-out data is a new hypothesis, requires new held-out data, and must be labelled exploratory in every subsequent sentence about it. **The held-out window can be spent exactly once.**

Report the in-sample fit too, and report it in a way that makes its irrelevance visible. In the null simulation of §4.7 the training-window `R^2` was 0.6477 — a figure that would be reported as a strong result in a great many papers — while the same frozen model was beaten on held-out data by persistence, by matched autoregression, and by 100 per cent of the circular-shift surrogates. **In-sample `R^2` of 0.65 and a held-out failure against every baseline is not a contradiction. It is the normal appearance of a false result.**

27.4.11 4.11 Cross-individual parameter recovery — FVP-7, the hard gate

Everything to this point tests whether a relationship exists in one series. The reproducibility gate – FVP-7 in the framework's formula-validation protocol [Curzi, 2026t] – asks the question that decides whether the framework measures anything: **do `delta` and `eta` recover as properties of a person, measured the same way, by different investigators?**

The framework identifies this as the gate most often skipped [Curzi, 2026t], and skipping it is what separates a formalism that describes from one that measures. Rasch's specific objectivity [Rasch, 1960, Andrich, 1988] is the standard aimed at — a person's measure must not depend on which items were used — and the analogous demand here is that a person's `eta` must not depend on which lab estimated it. Three levels, cleared in order.

Level 1 — within-person reliability

Split one subject's series into two halves and estimate separately in each. If the halves disagree, the parameter is a property of the window, not of the person. Simulated with rating noise 0.10 and true between-person spread 0.12 in `b`, 40 subjects per row:

```
half length      split-half r of b_hat
     60                  0.841
    120                  0.846
    240                  0.962
```

Sixty periods per half is enough for a usable within-person reliability; 240 is enough for a good one. Below 60, do not attempt Level 2.

Level 2 — between-person separability

Estimate for many subjects and ask whether the spread exceeds what estimation error alone produces. This step is almost universally botched, because the spread is always nonzero and always tempting to interpret. Simulate the noise floor directly: give every subject the *same* true `b = 0.40` and see how much apparent individual difference the estimator manufactures out of nothing.

```
THE NOISE FLOOR — 60 subjects who genuinely share one eta

T per subject      sd of b_hat across subjects
     60                    0.0478
    120                    0.0355
    240                    0.0255
    360                    0.0166
```

At 120 periods per subject, an estimator applied to sixty identical people reports a between-person standard deviation of 0.036 in `eta`. Any story about individual differences built on a spread smaller than that is a story about the estimator. A separate 30-subject run with the between-person standard deviation set to exactly zero produced an apparent spread of 0.042 — the same conclusion from a different seed.

This is why the recovery correlation collapses when true differences are small. With 30 subjects at 120 periods each:

```
true between-person sd of b        corr(true b, estimated b)
           0.12                             0.956
           0.04                             0.503
```

The second row is not a failure of the estimator. **The second row is what it looks like when the quantity you are trying to measure varies less across people than your instrument's error, and no sample size fixes it — only a longer series per person or a quieter instrument does.**

The full variance decomposition for 30 subjects at 120 periods, true between-person sd 0.12:

```
var(true b)          0.011586
var(estimated b)     0.012707
var(error)           0.001310
reliability          var_true / (var_true + var_err) = 0.898
```

Reliability 0.898 caps the correlation between `eta` and any external variable at `sqrt(0.898)` = 0.948. That ceiling is the useful output: **before correlating `eta` with anything — a personality measure, an outcome, a treatment response — compute the reliability and state the ceiling, because a correlation cannot exceed it and a reported correlation near it is suspicious.**

Also report the bias. Across the recovery runs, `b_hat` was biased upward by between 0.038 and 0.078, which is the attenuation mechanism of §4.9 operating at rating noise 0.10. **The bias is in the mean, not in the ranking, which is why recovery correlations can look excellent while every individual estimate is too high.** Report both.

Level 3 — between-lab invariance

The hard gate proper. Two labs measure the same construct with operationalisations they each consider faithful. Lab 2's procedure for computing `u` runs 15 per cent hot — a slightly wider framing set, a slightly different inner product, an entirely defensible choice.

```
lab 1 mean b_hat        0.4679
lab 2 mean b_hat        0.4068
observed ratio          0.8696
predicted ratio 1/1.15  0.8696
```

Exact, because the model is linear in the input: scaling `S` by `g` scales `b_hat` by `1/g`. That gives a useful diagnostic:

```
BETWEEN-LAB DIAGNOSTIC

If lab 2's eta differs from lab 1's by a CONSTANT FACTOR across all subjects,
the discrepancy is a CALIBRATION difference in the operationalisation of u,
and it is repairable by publishing the scaling constant.

If lab 2's eta differs from lab 1's by a factor that VARIES ACROSS SUBJECTS,
or if the RANK ORDER of subjects differs between labs, the discrepancy is a
THEORY FAILURE. The parameter is not a property of the person.
```

A framework passes FVP-7 when between-lab disagreement is a single number and fails when it is a pattern. That is a sharp, checkable, preregisterable criterion, and it is the criterion Stage 4 of Section 8 is built to apply.

Predicted ratios for other calibration errors, all exact:

```
inflation g     predicted b_hat ratio 1/g
   1.05                 0.9524
   1.10                 0.9091
   1.15                 0.8696
   1.25                 0.8000
   1.50                 0.6667
```

27.4.12 4.12 Goodhart and Campbell

Goodhart's law – when a measure becomes a target it ceases to be a good measure [Goodhart, 1975] – and Campbell's law – the more a quantitative indicator is used for social decision-making, the more it is subject to corruption pressures and the more it distorts what it was meant to monitor [Campbell, 1979] – both attach to `Y`.

The attachment to the order proxy is computable. Simulate a system running normally for 30 periods, at which point `Y` becomes a target and the *reported* value inflates by 0.35 while the underlying `Q` continues to obey the same law.

```
GOODHART SIMULATION — proxy becomes a target at t = 30

at t = 30    true Q = 3.0359      reported Y = 3.3650
at t = 60    true Q = 2.8447      reported Y = 3.1401   inflation = 0.2954

fitted on periods  0–30 (pre-target):   a = 0.8992  b = 0.3105
                                        -> delta = 0.1008  eta = 0.3105
fitted on periods 30–60 (post-target):  a = 0.8817  b = 0.3551
                                        -> delta = 0.1183  eta = 0.3551
```

The estimated conversion efficiency rose by 14 per cent — from 0.3105 to 0.3551 — **at exactly the moment the underlying stored order began to fall, from 3.036 to 2.845.** Nothing in the fit complains. The constraints hold. The parameters stay in the box. The model reports an improvement while the thing improves in the opposite direction.

Goodhart's Law is not a caveat about this framework. It is a failure mode with a signature, and the signature is a rising `eta_hat` accompanied by a falling independent indicator. That gives the detection procedure:

```
GOODHART DETECTION

G1  Maintain at least one order indicator that is NEVER reported upward,
    never used in any decision, and known to no one who can influence it.
    This is the shadow proxy. It exists solely to be compared.

G2  Test for a structural break in (a_hat, b_hat) at every point where the
    measure acquired a consequence: a review cycle, a published dashboard,
    a funding decision, a self-imposed goal.

G3  Compare the trajectory of the reported proxy against the shadow proxy.
    A widening gap dated to a decision is Campbell's Law, measured.

G4  Prefer proxies that are EXPENSIVE TO FAKE relative to the benefit of
    faking them. Archival traces and produced artefacts beat self-report
    on this criterion by a wide margin.
```

This is the corpus's own reflexivity turned on the corpus. Lucas [1976] makes the general form of the objection: parameters estimated in a population that does not know the model

cannot be assumed stable in a population that has read it; Hacking [1995] makes it about classifications of people, which move because the people classified respond to them; and MacKenzie and Millo [2003] documents the completed loop in the field, with dates and price data, where a model that did not fit its market came to fit it after traders adopted it. A framework that tells people their comfort obeys a decay law should also ask what that assumption makes true [Ferraro et al., 2005]. At single-person scope – one person, measuring themselves, against their own goals – the problem is at its worst, because **the subject, the instrument, the analyst, and the beneficiary of a good result are one organism with one set of incentives.** At group scope the rater and the rated could at least be separated. Here the separation must be manufactured: by DC1-DC4 of §4.8, by the shadow proxy of G1, and by the honest expectation that self-measurement drifts toward flattery at a rate that is itself worth measuring.

The framework's fourth standing refusal [Curzi, 2026] applies here with unusual force: do not build a measurement regime that removes a person's ability to decline being measured. **A proxy that a person cannot refuse to feed is a proxy they will learn to feed falsely, and the falsity will be indistinguishable from success.**

27.4.13 4.13 The publication contract

```
PC1 Publish the result whichever way it falls. A preregistration that is
    abandoned when the result disappoints has produced nothing except a
    private lesson.

PC2 Deposit the frozen model — the two coefficients, the code, the
    operationalisations, the surrogate calibration — at the moment of the
    freeze, before the held-out window is opened.

PC3 Report B, the surrogate family, and the CALIBRATED critical value.

PC4 Report the instrument's test-retest reliability in the same table as
    the parameter estimates (§4.9).

PC5 Report the observed span of u, so the reader can judge the direct-input
    comparison (§4.5.2).

PC6 Report boundary hits in the abstract (§4.2).

PC7 If any of the three rules fired, say so in the first paragraph. Not the
    limitations section. The first paragraph.

PC8 If the result is null, the title should say so.
```

The cost of not replicating is exact, and it has been computed elsewhere in this corpus: at zero independent replications per generation an error survives with probability 1.000000 indefinitely, against 0.013841 after twelve generations at one replication per generation [Curzi, 2026]. **A literature that does not publish nulls is running at zero replications on the nulls.**

PC8 is not moralising. **A field in which nulls are titled ambiguously is a field whose literature cannot be read at face value, and the premise of this section is a claim about literatures, not about individual papers.** The Open Science Collaboration's replication project found 97 of 100 originals significant, about 36 per cent of replications significant, and effect sizes roughly halved [Open Science Collaboration, 2015]; the base-rate argument for why most published findings in such a regime are false was made a decade earlier [Ioannidis, 2005]. **Formulas do not make a field hard. Ratio-scale, reproducible, falsifiable measurement of the formulas' terms does** — and so does a convention under which a null is legible from the title.

27.5 5 Specification of the falsification suite

A study that cannot answer yes to all of these has not run the falsification suite; it has run a regression.

```
S4.1   The preregistration contains every field A1–H2 of §4.3, filed before
        the first observation, with a timestamp that a third party can verify.

S4.2   The input variant (A: F(U_t), or B: normalised M_t) is declared, and if
        B, the ratio-scale evidence required by the formula-validation protocol \citep{ccp:CCP-026} for D a

S4.3   Chat_t is measured independently and subtracted before fitting. It is
        never estimated as a free parameter. Its units match those of eta*S.

S4.4   The split is chronological, with T_train and T_holdout fixed as counts
        in advance.

S4.5   The model is fitted with the constraints 0 <= a <= 1 and b >= 0, and the
        unconstrained fit is reported alongside it.

S4.6   Boundary hits (a_hat at 0 or 1, b_hat at 0, or any unconstrained estimate
        outside the box) are flagged in the abstract as Rule 3 events.

S4.7   All four baselines are computed on the held-out window: persistence,
        direct-input, matched autoregressive, shuffled input.

S4.8   The surrogate family is declared in advance. Block permutation with
        L >= 3 / (-ln rho1) is the default; departures are justified in writing.

S4.9   B >= 199 surrogates for a claim at the 1 per cent level; B >= 999 for
        anything stronger. B is reported.

S4.10  The surrogate test is CALIBRATED by simulation on the preregistered
        design with b set to zero, and the calibrated critical value replaces the
        nominal alpha in field F1. Nominal p-values are not reported alone.

S4.11  The observed span of u on the training window is reported. If it is below
        the preregistered minimum (default 0.60), the direct-input comparison is
        declared uninformative in print.

S4.12  At least two of the four design controls DC1–DC4 (§4.8) are in force, and
        which two is stated. Blinding the rater to the hypothesis (DC1) and
        separating the channels for Y and U (DC2) are the preferred pair.

S4.13  The instrument's test–retest reliability is estimated in a separate study,
        reported in the same table as the parameter estimates, and used to state
        the ceiling on any correlation involving those estimates.

S4.14  Scoring is one-step-ahead using observed Y_t. Free-running simulation may
        be reported as a secondary diagnostic and never as the primary score.

S4.15  Skill scores use 1 - (RMSE_model/RMSE_base)^2 with the sign convention
        stated, and are reported for all four baselines.

S4.16  The held-out window is opened exactly once. Any model revision after that
```

```
        opening is labelled exploratory in every subsequent sentence.

S4.17  Within-person split-half reliability of a_hat and b_hat is reported, with
        at least 60 periods per half.

S4.18  Before any claim about individual differences, the estimation noise floor
        is simulated at the study's own T and n, and the observed between-person
        spread is compared against it explicitly.

S4.19  The reliability var_true/(var_true + var_err) is computed and the implied
        ceiling on external correlations is stated.

S4.20  Between-lab comparison, where available, is reported as either a single
        scaling constant (calibration difference, repairable) or a pattern
        (theory failure, not repairable). The distinction is made in print.

S4.21  A shadow proxy exists (Goodhart G1): at least one order indicator never
        reported upward and never used in any decision.

S4.22  A structural-break test on (a_hat, b_hat) is run at every date on which
        the measure acquired a consequence.

S4.23  The frozen model and analysis code are deposited at the moment of the
        freeze, before the held-out window is opened.

S4.24  The result is published whichever way it falls, and if null, the title
        says so.

S4.25  No participant is required to be measured. Refusal is available, costless,
        and recorded as missing under the preregistered missing-data rule
        (the framework's fourth standing refusal, \citealp{ccp:CCP-A06}).
```

27.6 6 Failure modes of the falsification suite

Each fault is named, given its diagnostic signature, and given its fix.

27.6.1 F4.1 The Free Cost Term

Signature. `Chat_t` appears as an estimated coefficient, or an intercept plus trend, rather than a measured series. In-sample fit excellent, held-out fit mediocre, residuals structureless because the free term absorbed everything. Measuring cost is tedious; estimating it is easy.

Fix. Measure `C` on the framework's own cost instrument [Curzi, 2026b]. Subtract it. Fit two coefficients only. If `C` genuinely cannot be measured, say so and report that the model was fitted without a cost term, which is a different and weaker model.

27.6.2 F4.2 Proxy Contamination

Signature. `Y` rises whenever interaction is high, for reasons internal to the instrument. Paper E §6.3 names this as the most common measurement error in the entire framework: choosing an order proxy that is a repackaging of the interaction variable.

Fix. Field C1's argument, written before collection. Ask: could `Y` rise in a period with high `U` and no durable change whatsoever? If yes, `Y` is measuring `U`. Use DC2 of §4.8 — different channels for the two series.

27.6.3 F4.3 The Uncalibrated Shuffle

Signature. A reported permutation p-value of 0.02 with no statement of the surrogate family, no `B`, and no calibration run. The permutation test is usually presented as exact; it is exact only for the exchangeability structure it assumes.

Fix. §4.7. Simulate the design at `b = 0`, measure the actual rejection rate, and use the calibrated critical value. In the design simulated there, nominal 0.05 corresponded to actual rates between 0.060 and 0.120 depending on the surrogate family.

27.6.4 F4.4 The Fabricated Timing

Signature. Every baseline beaten decisively, shuffle p-value 0.000, parameters comfortably inside the box — and the subject knew the hypothesis, rated both series, and rated them retrospectively. It is the cheapest study to run.

Fix. DC1-DC4 of §4.8, at least two in force. Measured rate at which this fault passes the full suite when uncontrolled: **52.7 per cent.** There is no statistical remedy; only a design remedy.

27.6.5 F4.5 Attenuation Mistaken for Decay

Signature. `delta_hat` comes out large — the system looks leaky — and the instrument's reliability was never reported. Error in a lagged regressor attenuates its coefficient, and `delta = 1 − a`.

Fix. §4.9's table. Estimate instrument reliability separately. If rating error is comparable to the signal's own variability, the reported `delta` may be inflated by a factor approaching three (0.150 true against 0.432 estimated at rating sd 0.40).

27.6.6 F4.6 The Uninformative Direct-Input Comparison

Signature. "The logarithmic form outperformed the linear form" reported from data spanning `u` in `[0.30, 0.45]`. Over narrow bands `corr(u, F(u))` exceeds 0.995 and the comparison is near-vacuous.

Fix. S4.11. Report the span. Below the preregistered minimum, say the data do not distinguish the forms.

27.6.7 F4.7 Optional Stopping

Signature. The stopping rule reads "until the pattern is clear", or the reported `T` is not the preregistered `T`, or a supplementary wave was added after a first look. Stopping early when things look good feels efficient.

Fix. Field G1, as a date or a count. Any deviation is reported as a deviation, in the abstract.

27.6.8 F4.8 The Second Held-Out Window

Signature. The held-out result was poor; the model was revised; a new held-out window appeared from somewhere; the revised model is reported as confirmatory. The revision is usually a genuine improvement, which is what makes the fault feel like progress.

Fix. S4.16. The revised model is exploratory, permanently, until tested on data collected after its specification was frozen.

27.6.9 F4.9 Individual Differences That Are Estimation Noise

Signature. A table of per-subject `eta` values with a spread of 0.03 to 0.04, at 120 periods per subject, interpreted as a personality finding. The spread is real in the sense that the numbers differ, not in the sense that the people do.

Fix. S4.18. The measured noise floor at 120 periods is a between-subject sd of 0.036; at 60 periods, 0.048. A spread that does not clear the floor is the estimator talking.

27.6.10 F4.10 Goodhart Drift

Signature. `eta_hat` rising over successive fits while an independent indicator falls, the rise dated to a review cycle, a dashboard, or a personal goal. The measure acquired a consequence.

Fix. G1-G4 of §4.12. Simulated magnitude: `eta_hat` up 14 per cent (0.3105 -> 0.3551) while true `Q` fell from 3.036 to 2.845.

27.6.11 F4.11 The Unfalsifiable Rescue

Signature. After a failure, the model is retained with a new auxiliary assumption: `delta` was time-varying, `N` changed mid-study, the period length was wrong, the subject was unusual. Each rescue is individually plausible, which is precisely the problem.

Fix. Every auxiliary assumption becomes a preregistered prediction in the next study or it is withdrawn. A rescue that generates no new testable prediction converts a research programme into a system of excuses, and §8.13 gives the tripwires.

27.6.12 F4.12 The Missing Null

Signature. A research group with five preregistrations and two publications. Nulls are harder to place and less rewarding to write.

Fix. PC1 and PC8. Publish. Title it.

27.7 7 Worked example: the suite applied to a true model and to a null

One subject. Sixty-one periods of daily observation. `N = 7` distinguishable framings audited by the category audit [Curzi, 2026c], so `ln N = 1.945910`. Input variant A: `S_t = F(U_t) = ln(1 + 6 U_t)`. `Y_t` is a magnitude-estimate proxy for internalised structure on a ratio scale constructed by the procedure of Stage 2 in Section 8. `Chat_t` is an independently measured daily load series.

The data are **simulated**, stated plainly rather than buried: the point is to show what the suite does when the truth is known, which is impossible with real data. Generating parameters `delta = 0.15`, `eta = 0.40`, so `a_true = 0.85` and `b_true = 0.40`, process noise sd 0.06, no rating noise. **Every number below was computed, not asserted.**

27.7.1 7.1 The preregistration, in force before period 0

```
period length      1 day
split              T_train = 40 (periods 0–40), T_holdout = 20 (periods 40–60)
input variant      A
baselines          persistence; direct-input (raw u); matched AR with intercept;
```

```
                    shuffled input, three families reported
  surrogates        B = 2000 per family
  minimum u span    0.60 for the direct-input comparison to be informative
  design controls   DC2 (separate channels) and DC3 (prospective sampling)
  rules             Rule 1, Rule 2, Rule 3 as stated in §3.3
```

27.7.2 7.2 The training window

```
  u_t   over periods 0–39:  min 0.1248   max 0.6732   span 0.5484
  S_t   over periods 0–39:  min 0.5590   max 1.6173   mean 1.1746
                            span 1.0583 = 0.544 * ln N
  Chat_t mean over 0–39:    0.1136
  Y_t   over periods 0–40:  min 1.8722   max 2.6917   mean 2.3055   sd 0.1961
                            Y_0 = 2.0000, Y_40 = 2.3843
```

The observed span of `u` is 0.5484, which is below the preregistered minimum of 0.60. By S4.11 the direct-input comparison must therefore be reported as uninformative, and it will be below, even though the model happens to beat it. That is what a preregistration is for: the rule was written before the number was known, and the rule is honoured against the researcher's interest.

27.7.3 7.3 The fit, then the freeze

Constrained least squares on `(Y_{t+1} + Chat_t)` against `[Y_t, S_t]`, bounds `a` in `[0,1]`, `b >= 0`:

```
  a_hat = 0.8833      ->  delta_hat = 1 - 0.8833 = 0.1167
  b_hat = 0.3340      ->  eta_hat   = 0.3340
```

No boundary hit. Rule 3 does not fire.

Honest note on the point estimates, the most instructive line in this example. The true values are `a = 0.85` and `b = 0.40`; at `T = 40` with this noise the sampling standard deviations are 0.0341 for `b_hat` and 0.0173 for `a_hat`. The estimates sit at **-2.15 standard errors on `b` and +2.13 on `a`** — a tail draw in opposite directions, which is the correlated error the two coefficients trade between them.

A single forty-period window on a single subject recovered `eta` as 0.3340 when the truth was 0.4000: a 16.5 per cent underestimate, with everything measured perfectly and no rating noise at all. (The source manuscript rounds this to 17 per cent; `1 - 0.3340/0.4000 = 0.165` exactly, and the smaller figure is used here.) This is why §4.11 refuses to treat one person's `eta` from one window as a measured constant, and it is why Stage 3 of the programme in Section 8 asks for 120 to 240 periods rather than 40.

Frozen quantities derived from the fit, computed once and never revised:

```
  time constant      tau = 1/delta_hat = 1/0.1167 = 8.57 periods
  95% settling       ~3 tau = 25.7 periods
  implied Q_inf      (b_hat * mean S - mean Chat) / delta_hat
                   = (0.3340 * 1.1746 - 0.1136) / 0.1167 = 2.3883
                     (2.3885 at full precision)
  headroom at Y=2.39 h = eta*ln N - delta*Q - C
                   = 0.3340 * 1.945910 - 0.1167 * 2.39 - 0.115
                   = 0.6499 - 0.2789 - 0.1150 = 0.2560
```

Headroom is positive, so growth is permitted [Curzi, 2026b]. Note that `delta_hat * Q_inf = 0.2787` consumes most of the structural budget `eta_hat * ln N = 0.6499` at the fixed point — which is the steady-state theorem of the leaky-integrator analysis doing its arithmetic [Curzi, 2026a], not a finding.

27.7.4 7.4 The held-out window: periods 40 through 60

The frozen model, scored one step ahead against observed `Y_t`. First eight held-out periods in full:

t	u_t	S_t	Chat_t	Y_t	predicted	actual	error
40	0.3598	1.1501	0.1079	2.3843	2.3823	2.2926	−0.0897
41	0.6597	1.6011	0.1071	2.2926	2.4526	2.3476	−0.1050
42	0.6269	1.5606	0.1280	2.3476	2.4669	2.4080	−0.0589
43	0.4145	1.2490	0.1231	2.4080	2.4210	2.5253	+0.1042
44	0.1839	0.7435	0.1003	2.5253	2.3786	2.3176	−0.0610
45	0.3606	1.1517	0.1239	2.3176	2.3078	2.2891	−0.0187
46	0.2157	0.8304	0.1210	2.2891	2.1783	2.2166	+0.0383
47	0.1662	0.6918	0.1064	2.2166	2.0826	2.0749	−0.0077

Held-out summary: mean absolute error 0.0593; held-out `Y` mean 2.3001, sd 0.1978; held-out `u` span 0.4935.

27.7.5 7.5 The five-row report

predictor	RMSE	skill vs model	note
frozen model	0.0777	--	a=0.8833 b=0.3340
persistence	0.1207	0.5860	
direct-input (raw u)	0.0842	0.1486	span 0.548 < 0.60: UNINFORMATIVE
matched autoregressive	0.1248	0.6123	a=0.8422, c=0.4866
shuffled input, iid	0.1267	0.6238	B=2000, p = 0.0000
shuffled input, block L=5	0.1256	0.6174	B=2000, p = 0.0000
shuffled input, circular	0.1300	0.6428	B=2000, p = 0.0000

Surrogate distributions, 5th percentile — the best that any permutation achieved:

```
iid       5th percentile 0.1162   (model 0.0777)
block     5th percentile 0.1132   (model 0.0777)
circular  5th percentile 0.1196   (model 0.0777)
```

Not one of the six thousand surrogates across the three families produced a held-out error as small as the ordered model's. The best single permutation, out of 2000 block permutations, still predicted 46 per cent worse than the true ordering.

27.7.6 7.6 The three rules, applied

```
RULE 1  SHUFFLE       Calibrated critical value from the preregistered design
                      simulation: reject at p <= 0.05 for the block family, whose
                      measured false-positive rate on this design was 0.060 against
                      a nominal 0.050. Observed p = 0.0000 with B = 2000.
                      DOES NOT FIRE.

RULE 2  AUTOREGRESSION  Skill against matched AR = 0.6123, positive and large.
                      DOES NOT FIRE.

RULE 3  CONSTRAINTS   a_hat = 0.8833 strictly inside [0,1]; b_hat = 0.3340 > 0;
                      no boundary hit; unconstrained fit also inside the box.
                      DOES NOT FIRE.
```

Verdict: the model survives its own falsification suite on this subject and this window, with `delta_hat = 0.117 +- 0.017` and `eta_hat = 0.334 +- 0.034` at one standard error, both point estimates biased by roughly two standard errors in this particular draw. That last clause belongs in the abstract, not the appendix.

27.7.7 7.7 The same suite, run on a null

Now the honest negative, the half of the worked example that matters most. Same procedure, same period length, same instrument, same forty training periods — but `Y` is generated by a smooth process that has nothing to do with `S`. There is no relationship to find. The suite should say so.

```
fitted a_hat = 1.0000  ->  delta_hat = 0.0000     BOUNDARY HIT: Rule 3 event
fitted b_hat = 0.0931  ->  eta_hat   = 0.0931

in-sample R^2 on the training window: 0.6477

held-out RMSE   frozen model            0.0978
                persistence             0.0730
                direct-input            0.1116
                matched autoregressive  0.0681
                shuffled iid, mean      0.0793    p(shuffle <= model) = 0.999
                shuffled circular, mean 0.0763    p(shuffle <= model) = 1.000

skill vs matched autoregressive: -1.0669
```

Read that block slowly; it is the shape of a false result and every line is diagnostic.

The in-sample `R^2` is 0.6477, which presented alone reads as a solid finding. The frozen model is then beaten on held-out data by persistence (0.0730 against 0.0978), beaten decisively by matched autoregression (0.0681, skill -1.0669), and beaten by essentially every one of 2000 circular-shift surrogates — p = 1.000, meaning the permuted inputs predicted *better* than the true ordering, which is what happens when the true ordering carries no information. And the constrained fit hit the wall: `a_hat = 1.0000` exactly, the optimiser pressing against a boundary it was not allowed to cross. By S4.6 that goes in the abstract.

```
RULE 1  SHUFFLE        p = 1.000.  FIRES.
RULE 2  AUTOREGRESSION skill = -1.0669.  FIRES.
RULE 3  CONSTRAINTS    a_hat at the boundary.  FIRES.

VERDICT: FALSIFIED on this subject and this window. All three rules.
```

The correct report, in full:

> *We preregistered the minimal estimation model for a single subject over 61 daily periods, fitted on 40 and froze. The frozen model failed on held-out data against all four baselines. The shuffled-input control, run with B = 2000 circular-shift surrogates and a critical value calibrated by simulation on our own design, predicted held-out order as well as or better than the true input ordering (p = 1.000). The matched autoregressive baseline beat the model with a skill score of -1.0669. The constrained fit hit the retention boundary at a = 1.000. All three preregistered falsification rules fired. In this subject, over this window, with this operationalisation, internal surplus does not predict stored order.*

That paragraph is a good outcome. It cost the same as the other one, it is fully informative, and it is publishable under PC8. **The suite that produced the confirmation in the first half of this example is the same suite, unchanged, that produced the refutation in the second half, and neither result required the analyst to exercise a single judgement after the freeze.**

27.7.8 7.8 What this example does not establish

The limits, because the example is seductive:

```
It does not establish that the model holds for this subject in general.
    One window. Stage 3 of Section 8 asks for two, then more.

It does not establish that eta = 0.334 is a property of this person.
    The measured noise floor at T=40 is far above the difference between
    0.334 and 0.400. FVP-7 has not been attempted here.

It does not establish that the logarithmic input form is correct.
    The u span was 0.548, below the preregistered minimum. Reported.

It does not rule out fabricated timing.
    DC2 and DC3 were in force; DC1 (blinding) was not. Under uncontrolled
    conditions the measured pass rate for a fabricated relationship is
    52.7 per cent, so this omission is material and is stated as such.

It does not establish anything at group scope.
    Stage 5 of Section 8.
```

Five limitations, each named, each with the study that would remove it. That list is not an apology appended to a result. **It is the result: a claim, its exact boundary, and the next experiment.**

A framework that cannot fail cannot be trusted. This one fails at three named rules, on data whose truth we control, in the second half of its own worked example — and the next chapter lays out, in order and with its kill conditions, the programme that would put those rules in front of real people.

27.8 8 A staged research programme, costed, with twenty-one kill conditions

This section proves nothing empirical and measures nothing. It states a plan, and prices it.

Badged as a proposed programme, not as a result: the five stages below, run in order with the kill conditions stated at each, would settle whether the framework's single-person formulas are measurements. Every number in it is either arithmetic from a declared assumption or a simulation of a declared design. No stage has been run. Saying so in the first paragraph is the only honest way to open a chapter of this kind.

```
THE PROGRAMME

  STAGE 1   PILOT OPERATIONALISATION
            Fix what D, P, and E mean for one subject, in units.
            Gate: FVP-1 (dimensions), FVP-3 (limiting cases)

  STAGE 2   RATIO-SCALE CONSTRUCTION
            Build ratio scales for D, P, E by magnitude estimation and
            cross-modality matching. Verify transitivity.
            Gate: FVP-2 (scale type) -- the gate that licenses division

  STAGE 3   SINGLE-SUBJECT REPEATED MEASURES
            Estimate delta and eta within one person, repeatedly, and ask
```

```
           whether they are stable properties of that person.
           Gate: FVP-5 (identifiability), FVP-6 (falsification)

STAGE 4    CROSS-SUBJECT PARAMETER RECOVERY
           Same operationalisation, many people, two labs. Do the
           parameters recover?
           Gate: FVP-7 (reproducibility) -- MAKE OR BREAK

STAGE 5    PREREGISTERED FIELD TRIALS AT GROUP SCOPE
           Cluster-randomised, time-series, held-out prediction.
           Gate: FVP-6 at the external scope; the correspondence claim of
           the internal/external correspondence claim becomes testable here and only here
```

The ordering is not a preference. **Each stage consumes the output of the one before it, and running any stage before its predecessor has cleared produces a number that cannot be interpreted.** Estimating `eta` (Stage 3) on a scale not shown to be ratio-scaled (Stage 2) yields a coefficient in unknown units; comparing it across people (Stage 4) compares unknown units to unknown units; deploying it at group scope (Stage 5) does so at institutional expense.

27.8.1 8.1 The governing principle

A programme that cannot be killed at any stage is not a research programme. It is a research aspiration with a schedule attached.

Every stage below carries four blocks, and the third is what makes it a programme:

```
WHAT IS MEASURED
WHAT WOULD ADVANCE IT
WHAT WOULD KILL IT AT THIS STAGE
SCALE AND TIMESCALE
```

The kill conditions are written to be recognisable in advance and unpleasant to reach. A kill condition a determined investigator can argue past is a decoration; §8.13 gives the tripwires for the specific ways such arguments get made.

27.8.2 8.2 What is proven and what is postulated

Proven. Nothing empirical. The mathematics used throughout – the surplus functional [Curzi, 2026], the logarithmic ceiling [Curzi, 2026c], and the leaky-integrator solution [Curzi, 2026a] – is proven and is being *applied*, not tested. **The theorems cannot fail; what can fail is the claim that the world's psychological and social quantities are the objects those theorems are about.** That is why a correct mathematical result is no evidence at all that the programme will succeed.

Also proven, used here as design arithmetic and not as findings: Spearman-Brown, the Fisher-z interval for a correlation, the standard error of a log-log slope, the cluster design-effect formula, and the exact scaling `b_hat -> b_hat/g` under an input inflated by `g` (§4.11). All algebra, computed below, not cited.

Postulated, badged. That `D`, `P`, and `E` can be operationalised at all; that the storage law's form applies internally (a declared postulate); that `delta` and `eta` are person-level constants over months rather than moment-level variables; that the internal/external correspondence is more than an analogy. **Every one is exactly what a stage below is designed to break.**

Honesty clauses, repeated. The triangle inequality fails for geometric triples in the framework's own statement of the surplus functional [Curzi, 2026]; subadditivity holds on `u` only. As deployed, the surplus functional is a lower bound on the ensemble's information — we log the floor, not the ceiling. The surplus functional is a mathematical object and carries no empirical

or metaphysical claim; the framework does not identify surplus with Shannon entropy or mutual information; and it proves no thermodynamic negentropy law, since `F(u) + N_can(u) <= ln N` and higher static surplus does not imply higher canonical concentration [Curzi, 2026,c].

27.8.3 8.3 The field this programme is entering

Psychology and sociology are not pseudoscience. They are sciences with a measurement problem and a replication problem. The Open Science Collaboration attempted 100 replications: 97 originals had reported significant results, about 36 per cent of replications did, and effect sizes roughly halved [Open Science Collaboration, 2015, Ioannidis, 2005]. The measurement problem is the deeper of the two: Michell [1997] argues that psychology systematically fails to test whether its attributes are quantitative at all before measuring them, and that objection applies to this framework's central quantities without amendment [Michell, 1999].

Formulas do not make a field hard. Ratio-scale, reproducible, falsifiable measurement of the formulas' terms does. Psychophysics achieves it [Weber, 1834, Fechner, 1860, Stevens, 1957], and psychometrics comes close through the Rasch model's separation of person and item parameters [Rasch, 1960, Andrich, 1988]. This programme is an attempt to bring the framework's single-person formulas onto that footing, not an attempt to prove them.

27.8.4 8.4 Stage 0 — preconditions

Four things must exist before Stage 1. None is research; starting without them guarantees a Stage 1 that cannot be interpreted.

```
SP1 A WRITTEN THEORY OF THE SUBJECT.
    Which construct is Q for this subject? Not "wellbeing". Something
    that could in principle be destroyed by a year of neglect and rebuilt
    by a year of practice: a skill, a body of knowledge, a settled habit
    of attention, a repertoire.

SP2 A CANDIDATE ORDER PROXY THAT IS NOT A REPACKAGING OF U.
    The order-proxy trap. Write the argument down and test it: could Y rise in a
    period of intense interaction that produced nothing durable? If no,
    Y is measuring U.

SP3 AN INDEPENDENT COST INSTRUMENT.
    C in the same units as eta*S. If cost cannot be measured, the model
    is fitted without it and is a weaker model (F4.1).

SP4 ETHICS AND CONSENT INFRASTRUCTURE, INCLUDING COSTLESS REFUSAL.
    The fourth standing refusal: no intervention removes a person's ability to
    decline. In a measurement programme that means refusal to be measured
    must be available, costless, and recorded as ordinary missing data.
```

Stage 0 has no kill condition because it has no result. It has an entry condition: all four of SP1-SP4, in writing, before recruitment.

27.8.5 8.5 Stage 1 — Pilot operationalisation of D, P, and E for a single subject

What is measured

One subject, 45 sessions, three candidate instruments per term:

```
D  PERSPECTIVE, the frame brought to the moment.
     D-a  prior-expectation elicitation, with a confidence, before the
          episode
     D-b  frame inventory: which of the subject's N audited framings
          (the audited framing set) were active, weighted
     D-c  magnitude estimate of "how strongly I already knew what this
          would be like", against a fixed modulus

P  PERCEPTION, the intake from the moment.
     P-a  detail-recall count, scored against a record
     P-b  discrepancy report: what departed from D-a, with magnitude
     P-c  magnitude estimate of "how much came in", same modulus

E  ENVIRONMENT, the situation itself.
     E-a  independently coded situational richness from a record
     E-b  physical-parameter panel where the room is the environment
          (acoustics, light, sittable space, sightlines)
     E-c  magnitude estimate of "how much was there to be had"
```

The purpose is not to find the right instrument. **It is to find out whether `D` and `P` can be put in the same units at all**, because FVP-1 requires `r = D/P` to be dimensionless — a far stronger demand than that both be numbers. This is the least glamorous stage and the one most often skipped by starting straight from a questionnaire.

Per candidate: session-level reliability (internal consistency across items, Cronbach, 1951, read with the warning that alpha is neither a reliability coefficient nor a measure of homogeneity, Sijtsma, 2009), test-retest reliability (across sessions, short interval), and convergent and discriminant correlation across the three candidate instruments, which is the multitrait-multimethod design [Campbell and Fiske, 1959, Cronbach and Meehl, 1955].

What would advance it

```
A1  One operationalisation of D and one of P share units, with a written
    argument for the shared unit and a demonstrated common zero.

A2  Session-level reliability >= 0.80 for at least one instrument per term.

A3  Test-retest r >= 0.70 across a one-week interval, 95 per cent lower
    bound above 0.50.

A4  FVP-3's four limiting cases are observable and recognisable: P -> 0
    (frame with no intake), D -> 0 (intake with no frame), E -> 0, and
    D = P (calibrated contact). If none can be produced or recognised in
    a real session, the formula's structure has no referent.
```

A3 sets the sample size. The 95 per cent interval on an observed correlation, by Fisher-z:

n	95% CI on an observed r = 0.70	width
10	[0.126, 0.923]	0.797
20	[0.373, 0.872]	0.499
30	[0.454, 0.847]	0.392
50	[0.524, 0.819]	0.295
80	[0.568, 0.797]	0.230
120	[0.595, 0.781]	0.186
200	[0.622, 0.764]	0.143

The smallest `n` at which an observed `r = 0.70` has a 95 per cent lower bound above 0.50 is 41. Thirty sessions is therefore under-powered for A3, and the honest choices are 41 or more sessions, or relaxing A3 to a lower bound above 0.40 and saying so in advance. This manual takes the first and sets Stage 1 at 45 sessions.

A2 sets the item count. Spearman-Brown, `r_k = k*r1/(1 + (k-1)*r1)`:

single-item r1	items for reliability 0.80	for 0.90
0.10	37	82
0.15	23	52
0.20	17	37
0.30	10	22
0.40	7	14
0.50	5	10

If the single-item reliability of a `P` instrument is 0.15, a session score needs 23 items to reach 0.80 and 52 to reach 0.90. That is what decides session length, and why Stage 1's session is 25 minutes rather than five.

What would kill the programme at this stage

```
K1.1  NO SHARED UNIT. After honest effort, no operationalisation of D and
      P can be expressed in the same units with a common zero.
      CONSEQUENCE: D/P is not a ratio. M = (D/P)E is not an equation.
      By FVP-2 it must be demoted in print to a heuristic ordering, and
      the moment equation ceases to be a formula in the framework's sense.
      The storage law survives; the Moment equation does not.

K1.2  NO RELIABLE INSTRUMENT. No candidate reaches session reliability
      0.80 with a session under 45 minutes.
      CONSEQUENCE: the attenuation analysis of §4.9 applies with force -- at
      rating sd 0.40 against signal sd near 0.20, delta was estimated at
      0.432 against a true 0.150. Estimates on such an instrument are not
      noisy but systematically wrong. The programme cannot proceed.

K1.3  THE LIMITING CASES HAVE NO REFERENT. Neither the subject nor a
      trained observer can identify instances of P -> 0, D -> 0, or
      D = P in the subject's own record.
      CONSEQUENCE: FVP-3 fails. A formula whose divergences correspond
      to nothing nameable is not describing the phenomenon.

K1.4  Y IS A REPACKAGING OF U. The candidate order proxy rises with
      interaction volume and with nothing else.
      CONSEQUENCE: the named order-proxy error. Return to Stage 0 SP2 or
      stop.
```

K1.1 is the most likely kill in the whole programme and the one people will most want to argue past. **That "perspective" and "perception" are expressible in a common unit is not obvious, is not established, and is the load-bearing assumption of the entire Moment equation.** The honest outcome of a failed Stage 1 is a chapter saying the Moment equation is a heuristic ordering — useful for diagnosis, illegitimate for arithmetic — and the formula-validation protocol already specifies what that demotion looks like in print [Curzi, 2026t].

Scale and timescale

```
subjects            1  (plus 1 replication subject, run in parallel)
```

```
sessions           45  at 3 per week
session length     25 minutes x 3 instruments per term
subject burden     45 x 25 x 3 / 60 = 56.3 subject-hours
fieldwork          45 / 3 = 15 weeks = 3.5 months
total stage        7 months including instrument development and analysis

researcher hours (itemised, assumptions declared not cited):
   instrument development and piloting            160.0
   session supervision                             37.5
   coding at 2x subject time                       75.0
   analysis and FVP-1/2/3 write-up                 80.0
   preregistration drafting                        24.0
   ------------------------------------------------
   TOTAL                                          376.5
```

Seven months and about 377 researcher-hours to find out whether the framework's central single-person formula has units. A cheap way to learn something decisive, and the correct first expenditure because a failure here saves the other five years.

27.8.6 8.6 Stage 2 — Ratio-scale construction

What is measured

Magnitude estimation and cross-modality matching on the instruments that survived Stage 1.

```
MAGNITUDE ESTIMATION
  A modulus is presented and assigned a number; each subsequent stimulus
  is assigned a number in proportion. Fit log psi = log k + a log I and
  read the exponent a.

CROSS-MODALITY MATCHING
  The subject adjusts a stimulus in modality B until it matches one in
  modality A. The matching exponent is predicted by the ratio of the two
  modalities' own exponents. This is the transitivity test, and it is the
  actual evidence that the scale is a ratio scale and not a numbering
  habit.
```

The published exponents are constants of a *procedure* rather than of nature: they vary with method, range and instruction [Stevens, 1957, 1975, Teghtsoonian, 1971], and the catalogue of ways a magnitude-estimation procedure biases its own exponent is long and was written by a critic [Poulton, 1968, 1979]. The century-long dispute over whether the underlying law is logarithmic or a power function is better read through Krueger [1989] than adjudicated here. The nominal values:

```
Loudness (1 kHz)     a ~ 0.67     compressive
Brightness           a ~ 0.33     compressive
Visual length        a ~ 1.0      linear
Taste (salt)         a ~ 1.3      expansive
Heaviness (lifted)   a ~ 1.45     expansive
Warmth on arm        a ~ 1.6      expansive
Electric shock       a ~ 3.5      strongly expansive
```

These are the calibration ladder. **`D`, `P`, and `E` are anchored by matching against modalities whose exponents are already known, which is what converts a number the subject produced into a measurement.** What must be true of a subject's judgements for that conversion to be

licensed is itself an axiomatic question with a literature [Luce, 2002, Narens, 1996], and the decisive formal test – whether the attribute has additive quantitative structure at all – is conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971]. **Stage 2 as specified below does not run the conjoint test.** It runs cross-modality matching, which is weaker evidence for the same conclusion, and the gap is recorded in Section 19.

What would advance it

```
B1  Each of D, P, E yields a stable within-subject exponent across
    sessions, with a usefully small standard error.

B2  Cross-modality matching is TRANSITIVE within measurement error:
    a_AC = a_AB * a_BC.

B3  D and P yield exponents consistent with a common unit, so r = D/P is
    dimensionless (Stage 1's A1, now quantitative).

B4  A true zero is demonstrable for both D and P: a condition reliably
    assigned magnitude zero, not merely the smallest number used.
```

The precision arithmetic sets the number of judgments. For a log-log regression, `SE(a) = sigma_resid / (sd(log I) * sqrt(n))`, and with stimulus levels spread evenly over a span of `L` decades, `sd(log I) = L/sqrt(12)`:

```
residual sd 0.25 log10 units; stimulus span in decades:

  n \ span     1.0      1.5      2.0      3.0
    20        0.1936   0.1291   0.0968   0.0645
    40        0.1369   0.0913   0.0685   0.0456
    60        0.1118   0.0745   0.0559   0.0373
   100        0.0866   0.0577   0.0433   0.0289
   200        0.0612   0.0408   0.0306   0.0204
   400        0.0433   0.0289   0.0217   0.0144

judgments needed at a 2.0-decade span:
   SE(a) <= 0.05  ->  n =  75
   SE(a) <= 0.03  ->  n = 209
   SE(a) <= 0.02  ->  n = 469
```

Two design decisions fall out of that table. First, span matters as much as count: going from 1.0 to 2.0 decades halves the standard error at every `n`, the same gain as quadrupling the judgments. Widen the range before lengthening the session. Second, `SE(a) <= 0.02` costs 469 judgments per exponent — a real afternoon per term per subject.

The transitivity test's power depends on that precision. Take a chain with `a_AB = 0.67` and `a_BC = 1.45`, so predicted `a_AC = 0.9715`:

```
SE per exponent    predicted a_AC       SE of (a_AC - a_AB*a_BC)    detectable at 95%
     0.02          0.9715 +- 0.0319           0.0377                    0.0739
     0.05          0.9715 +- 0.0799           0.0942                    0.1847
     0.10          0.9715 +- 0.1597           0.1885                    0.3694
```

At `SE = 0.10` per exponent, a transitivity violation smaller than 0.37 cannot be detected, and a violation of 0.37 is enormous. A transitivity test on imprecise exponents is not a test but a guarantee of agreement. This is how the ratio-scale claim gets made without being earned, and why B1's precision requirement precedes B2's transitivity requirement.

Session arithmetic for the stage:

```
9 stimulus levels x 6 replications x 6 sessions x 3 modalities
    = 972 judgments per subject per term
x 3 terms (D, P, E)
    = 2,916 judgments per subject
at 8 seconds per judgment
    = 6.48 subject-hours per subject
6 subjects
    = 38.9 subject-hours total
```

972 judgments per term per subject clears the `n = 469` requirement for `SE(a) <= 0.02` on each modality separately, which is the precision the transitivity test needs.

What would kill the programme at this stage

```
K2.1  NO STABLE EXPONENT. The fitted exponent for D or P varies across
      sessions by more than its own standard error by a wide margin.
      CONSEQUENCE: the subject is not using the number line consistently.
      There is no scale to build on. FVP-2 fails.

K2.2  TRANSITIVITY VIOLATED. a_AC differs from a_AB * a_BC by more than
      the detectable threshold above.
      CONSEQUENCE: the judgments are not measurements of a single
      underlying magnitude. Classic and decisive evidence that a putative
      ratio scale is not one.

K2.3  NO TRUE ZERO. No condition is reliably assigned magnitude zero for
      D or for P.
      CONSEQUENCE: the scale is at best interval, and on Stevens'
      typology division is impermissible on an interval scale.
      D/P is meaningless as written. Demote the moment equation
      in print.

K2.4  D AND P LAND IN INCOMPATIBLE UNITS. Their exponents cannot be
      reconciled to a common unit.
      CONSEQUENCE: as K1.1, now with numbers attached.
```

K2.3 is worth stating twice. **A subjective quantity with no true zero cannot be divided by another, and the entire content of `M = (D/P)E` is a division.** The scale-type taxonomy that makes this a rule rather than a preference is Stevens [1946]'s, and the corpus's own meaningfulness theorem states the same condition sharply: a quotient is meaningful if and only if the admissible transformation group is similarity, that is, if and only if the scale is ratio [Curzi, 2026t]. The programme's most probable non-trivial outcome is not that the Moment equation is false but that it is not an equation — that `r = D/P` survives as an ordering as an ordering rather than a quotient [Curzi, 2026t] while ceasing to be a number. That is a genuine finding, and the formula-validation protocol specifies how to print it [Curzi, 2026t].

Scale and timescale

```
subjects            6
judgments        2,916 per subject; 17,496 total
subject burden    6.48 hours per subject; 38.9 hours total
fieldwork         6 sessions per term at 1 per week; terms run in parallel
                  = 6 weeks
total stage       4 months including apparatus construction and analysis
```

```
researcher hours:
   session supervision at 1.0x subject time         38.9
   stimulus and apparatus construction, 3 modalities  240.0
   scale fitting and transitivity analysis         180.0
   independent replication wave                    200.0
   writing                                         120.0
   ------------------------------------------------
   TOTAL                                           778.9
```

27.8.7 8.7 Stage 3 — Single-subject repeated measures

What is measured

Three subjects who cleared Stages 1 and 2, sampled four times daily for 240 days. Per period: `Y_t` on the constructed scale, `U_t` over the subject's own audited framings, `Chat_t` from the independent cost instrument. The suite of Section 4 is applied per subject: fit on the first 120 periods, freeze, predict the second 120, four baselines, three rules.

The question is not "does the model fit" but **"are `delta` and `eta` properties of this person, or of the window in which they were estimated?"**

What would advance it

```
C1  The frozen model clears all three falsification rules on held-out
    data, per subject, with a calibrated surrogate test.

C2  Within-person split-half reliability of b_hat is at least 0.80.

C3  a_hat lies strictly inside [0,1] in every subject, no boundary hit.

C4  No systematic drift: first-half parameters predict the second half,
    and second-half parameters predict the first, to within sampling
    error.
```

The precision arithmetic, simulated at true `a = 0.85`, true `b = 0.40`, process noise 0.06, rating noise 0.10, 400 replications per row:

T	mean a_hat	sd a_hat	mean b_hat	sd b_hat	95% CI half-width on b
30	0.8174	0.0409	0.4604	0.0781	0.1532
60	0.8224	0.0248	0.4521	0.0494	0.0968
90	0.8242	0.0207	0.4487	0.0411	0.0805
120	0.8236	0.0166	0.4513	0.0331	0.0649
180	0.8252	0.0137	0.4485	0.0278	0.0545
240	0.8248	0.0119	0.4488	0.0243	0.0476
360	0.8257	0.0097	0.4476	0.0198	0.0388

Two features of that table are the reason Stage 3 exists.

First, the bias. At every `T`, `b_hat` sits near 0.449 against a true 0.400 — a persistent 12 per cent upward bias that does not shrink with more data, because it is not sampling error but the attenuation of §4.9 at rating noise 0.10. More periods buy precision, not accuracy. **Only a better instrument buys accuracy, and that is Stage 2's job.**

Second, the cost of precision. Halving `sd(b_hat)` from 0.0494 to 0.0243 means going from 60 periods to 240 — four times the fieldwork for twice the precision, the ordinary square-root law, worth stating in weeks:

```
T = 120 at daily sampling      = 17.1 weeks =  3.9 months
T = 120 at twice-weekly        = 60.0 weeks = 13.8 months
T = 240 at daily sampling      = 34.3 weeks =  7.9 months
```

The middle row is why daily sampling is not optional. **Twice-weekly sampling turns a four-month study into a fourteen-month study for the same precision, and a fourteen-month study has a different `delta` at the end than at the beginning.**

Power to detect a genuine within-person *change* in `eta` between two blocks, with `SE_diff = sqrt(2) * sd_b`:

T	SE_diff	power at d=0.05	d=0.10	d=0.20
60	0.0698	0.110	0.299	0.817
90	0.0581	0.138	0.406	0.931
120	0.0469	0.187	0.569	0.990
180	0.0393	0.246	0.720	0.999
240	0.0344	0.307	0.829	1.000
360	0.0280	0.432	0.947	1.000

Two blocks of 120 periods give 57 per cent power to detect a change in `eta` of 0.10 and 19 per cent for a change of 0.05. Any claim that a practice, therapy, habit or season changed a person's conversion efficiency by a small amount is, at these sizes, undetectable. Stating that in advance prevents the commoner outcome, which is detecting it anyway.

Burden: 960 prompts per subject at about 90 seconds each — 24 subject-hours over eight months. Adherence is itself a threat; it must be measured and reported, and missing periods handled by the preregistered rule, never by interpolation chosen after the fact.

What would kill the programme at this stage

```
K3.1  ANY OF THE THREE RULES FIRES IN A MAJORITY OF SUBJECTS.
      Shuffle not beaten, matched AR not beaten, or a boundary hit.
      CONSEQUENCE: the storage law does not describe this subject's
      internal dynamics. Report it (PC8) and stop.

K3.2  SPLIT-HALF RELIABILITY OF b_hat BELOW 0.50.
      The simulated benchmark at 60 periods per half is 0.841 and at 240
      per half is 0.962. A real figure below 0.50 means the parameter is
      a property of the window.
      CONSEQUENCE: eta is not a person-level constant on the months
      timescale. Stage 4 is meaningless and must not be attempted.

K3.3  SYSTEMATIC SAME-DIRECTION DRIFT in a_hat or b_hat across blocks,
      in every subject.
      CONSEQUENCE: the constant-parameter postulate fails. A time-varying
      rescue must be preregistered as a NEW prediction (§8.13), never
      adopted as an explanation of this failure.

K3.4  THE DESIGN CONTROLS FAIL. Debriefing reveals subjects inferred the
      hypothesis, and no blinding or channel separation was in force.
      CONSEQUENCE: §4.8 measured a 52.7 per cent pass rate
      for fabricated relationships under those conditions. The data
      cannot be interpreted however good they look.
```

K3.4 will be tempting to waive because the data will already exist. **Waiving it converts eight months of fieldwork into eight months of fieldwork that means nothing — exactly the outcome the waiver is trying to avoid and exactly the outcome it produces.**

Scale and timescale

```
subjects            3  (each analysed separately; no pooling at this stage)
periods           240  per subject, daily, 4 prompts per day
prompts           960 per subject; 2,880 total
subject burden     24 hours per subject; 72 hours total
fieldwork        34.3 weeks = 7.9 months
total stage        12 months including pipeline build and analysis

researcher hours:
   adherence monitoring (3 x 34.3 weeks x 0.5 h)     51.4
   data pipeline build                              200.0
   surrogate calibration simulations                120.0
   per-subject fits, split-half, analysis           300.0
   writing                                          160.0
   ------------------------------------------------------
   TOTAL                                            831.5
```

27.8.8 8.8 Stage 4 — Cross-subject parameter recovery: the make-or-break stage

What is measured

Forty subjects, 120 daily periods each. **The identical protocol run independently in two laboratories that have never met, working from the published Stage 1-3 protocol alone.**

That last clause is the stage. Everything before it can be done by one group with one set of habits. FVP-7 asks whether the parameters survive a different group's reading of the same written procedure, and the framework names it the gate that decides the measurement question and the one most often skipped. Three questions:

```
Q1  Do individual differences in delta and eta EXIST -- is the spread
    across people larger than the estimator's noise?

Q2  Do they RECOVER -- does a person's estimated parameter correlate with
    that person's true parameter, insofar as a second estimate can serve
    as a proxy for truth?

Q3  Do they TRANSPORT -- does lab 2 get the same numbers as lab 1, and if
    not, is the disagreement a single scaling constant or a pattern?
```

The noise floor, and why Q1 comes first

The arithmetic most individual-differences research omits. Simulate sixty subjects who share one true `eta` and see how much variation the estimator invents:

```
THE NOISE FLOOR -- 60 subjects who truly share eta = 0.40

  T per subject      sd of b_hat across subjects
      60                    0.0478
     120                    0.0355
     240                    0.0255
     360                    0.0166
```

At 120 periods per subject, an estimator applied to sixty identical people reports a between-person standard deviation of 0.036 in `eta`. Any individual-differences claim resting on a

smaller spread is a claim about the estimator, not about people. An independent 30-subject run with the true spread set to zero produced 0.042: same conclusion, different seed.

The consequence for Q2, with 30 subjects at 120 periods:

true between-person sd of eta	corr(true, estimated)
0.12	0.956
0.04	0.503

The second row is not an estimator failure. **It is what happens when the quantity varies less across people than the instrument's error, and no number of subjects repairs it — only longer series per subject or a quieter instrument.**

Power to reject "everyone shares one `eta`", testing the observed spread against a simulated noise floor at one-sided 5 per cent:

n	T	true between-sd	critical sd	power
20	120	0.05	0.0415	0.97
20	120	0.10	0.0428	1.00
40	120	0.05	0.0391	0.98
40	120	0.10	0.0407	1.00
20	240	0.05	0.0298	1.00
40	240	0.05	0.0282	1.00

Twenty subjects at 120 periods already gives 97 per cent power to detect a true between-person standard deviation of 0.05 in `eta`. Forty is comfortable, and the binding constraint is periods per subject, not subjects — the opposite of the usual instinct. **One more month of daily sampling per subject buys more than ten more subjects.**

Reliability of the person-level estimate, and the ceiling it puts on everything downstream:

T	between-sd	var_err	reliability	max corr with any external variable
120	0.05	0.00126	0.665	0.815
120	0.10	0.00126	0.888	0.942
120	0.15	0.00126	0.947	0.973
240	0.05	0.00065	0.793	0.891
240	0.10	0.00065	0.939	0.969
240	0.15	0.00065	0.972	0.986

Before correlating `eta` with anything, compute this ceiling and publish it: an observed correlation cannot exceed it, and one reported near it is evidence of a problem, not a discovery.

Q3, and the sharpest criterion in the programme

Two labs, same protocol, different hands. Lab 2's operationalisation of `u` runs 15 per cent hot: a wider framing set, a different inner product, an entirely defensible reading of the procedure. Simulated:

lab 1 mean b_hat	0.4679
lab 2 mean b_hat	0.4068
observed ratio	0.8696
predicted ratio 1/1.15	0.8696

Exact, because the model is linear in its input: scaling `S` by `g` scales `b_hat` by `1/g`. Other calibration errors, equally exact:

```
inflation g      predicted b_hat ratio 1/g
   1.05                 0.9524
   1.10                 0.9091
   1.15                 0.8696
   1.25                 0.8000
   1.50                 0.6667
```

The criterion:

```
BETWEEN-LAB CRITERION

A CONSTANT FACTOR across all subjects, matching 1/g for an independently
measurable operationalisation difference g
    -> CALIBRATION difference. Repairable. Publish the constant.

A FACTOR THAT VARIES ACROSS SUBJECTS, or any change in the RANK ORDER of
subjects between labs
    -> THEORY FAILURE. The parameter is not a property of the person.
```

A framework passes FVP-7 when between-lab disagreement is a single number, and fails when it is a pattern. That is why Stage 4 is make-or-break: it is the only stage whose failure cannot be attributed to an instrument, a subject, or a window.

What would advance it

```
D1  The between-person spread in eta exceeds the simulated noise floor at
    the study's own n and T, by a margin, in both labs.

D2  Rank order of subjects on eta agrees between labs, to a correlation
    significantly above zero and at or below the published reliability
    ceiling.

D3  Any mean-level disagreement is a single scaling constant matching an
    independently measured operationalisation difference.

D4  Both labs clear the three-rule suite of Section 4 in most subjects.
```

What would kill the programme at this stage

```
K4.1  NO INDIVIDUAL DIFFERENCES. The between-person spread does not clear
      the noise floor at n = 40, T = 120, in either lab.
      CONSEQUENCE: eta and delta are not person parameters. the framework's person-level
      parameter tables describes a species-typical process, not
      individual constants, and every person-level parameter claim must be
      rewritten. Survivable, but a large retraction, printed as one.

K4.2  NO RANK-ORDER AGREEMENT BETWEEN LABS. Lab 1 and lab 2 order the same
      forty subjects differently on eta, beyond sampling error.
      CONSEQUENCE: FVP-7 fails outright. The parameters are properties of
      a procedure and not of a person. THE PROGRAMME STOPS. There is no
      rescue that preserves the claim, because the claim was that these
      are measurements.

K4.3  MEAN DISAGREEMENT THAT IS NOT A CONSTANT FACTOR. Labs differ in a
```

```
       way that varies across subjects and does not track any measurable
       difference in operationalisation.
       CONSEQUENCE: as K4.2. An unexplained, subject-varying between-lab
       discrepancy is the signature of a parameter that does not exist
       independently of who measured it.

K4.4   THE SUITE FAILS IN ONE LAB AND PASSES IN THE OTHER, with no
       identifiable procedural difference.
       CONSEQUENCE: the strongest possible warning. Do not proceed to
       Stage 5. Return to Stage 1 and find the procedural difference, or
       stop.
```

K4.2 is the programme's true kill switch. **If the same forty people get different `eta` orderings from two competent labs following the same written protocol, then `eta` is not a quantity the people have; it is a quantity the procedure makes, and no further work changes that.** Specific objectivity [Rasch, 1960] is the standard being failed: a person's measure depending on which instrument was used is precisely what a measurement is not allowed to do.

Scale and timescale

```
subjects           40 per lab; 80 total across two labs
periods           120 per subject, daily, 4 prompts per day
prompts           480 per subject; 38,400 total
subject burden     12 hours per subject; 960 hours total
fieldwork        17.1 weeks = 3.9 months per lab cohort
total stage        18 months including recruitment, two cohorts,
                   analysis and between-lab reconciliation

researcher hours:
   fieldwork, two labs (recruitment, monitoring, coding)   822.0
   data pipeline adaptation                                150.0
   analysis, noise floor, between-lab reconciliation       500.0
   writing                                                 240.0
   ------------------------------------------------------------
   TOTAL                                                  1712.0
```

27.8.9 8.9 Stage 5 — Preregistered field trials at group scope

What is measured

Only now does the programme touch institutions, and only now is it entitled to. Two tests, which must not be confused with each other:

```
T1  Does the storage law predict held-out group order, at group scope,
    against the same four baselines? (the suite of Section 4, applied at external scope)

T2  Does an intervention chosen by the framework -- raise eta, cut C, or
    raise u, whichever the headroom gauge says is cheapest \citep{ccp:CCP-008}
    -- outperform an active control chosen by ordinary practice?

T1 is a prediction test. T2 is a causal test. A design that runs only T2
can succeed while the framework is wrong (the intervention worked for
other reasons) and can fail while the framework is right (the
intervention was implemented badly). BOTH ARE REQUIRED.
```

And a third, the one that justifies the single-person scope:

```
T3  Does the internal/external correspondence hold
    QUANTITATIVELY? Are members' internal eta values, measured by Stage
    3–4 methods, related to their group's eta? The correspondence claim asks for the
    cases where the correspondence BREAKS DOWN to be stated honestly, and
    Stage 5 is where they would be found.
```

Design: cluster-randomised, 24 clusters, 20 members each, 52 weekly periods, framework intervention in half and an active control in the other half. Time series throughout, not end-points.

The sample-size arithmetic

`z(0.975) = 1.959964`, `z(0.80) = 0.841621`, so `(z_a + z_b)^2 = 7.848880`. Individual-level `n` per arm for 80 per cent power at two-sided alpha = 0.05, from `n = 2*(z_a+z_b)^2 / d^2`:

```
 d = 0.20      n per arm = 393
 d = 0.30      n per arm = 175
 d = 0.50      n per arm =  63
 d = 0.80      n per arm =  25
```

Clustering inflates this by `DE = 1 + (m–1)*ICC`:

```
m \ ICC     0.01    0.02    0.05    0.10    0.20
  5         1.04    1.08    1.20    1.40    1.80
 10         1.09    1.18    1.45    1.90    2.80
 20         1.19    1.38    1.95    2.90    4.80
 30         1.29    1.58    2.45    3.90    6.80
 50         1.49    1.98    3.45    5.90   10.80
100         1.99    2.98    5.95   10.90   20.80
```

Clusters per arm at `d = 0.30`, 80 per cent power (individual `n` per arm = 175):

```
 m       ICC=0.02   ICC=0.05   ICC=0.10
  8         25         30         38
 12         18         23         31
 20         13         18         26
 30         10         15         23
 50          7         13         21
```

At 20 members per cluster and ICC 0.05, eighteen clusters per arm — 36 in total — are needed for 80 per cent power on a standardised effect of 0.30. The design above specifies 24, which is under-powered for `d = 0.30` at `ICC = 0.05` and adequate for `d = 0.50`. Stated here rather than discovered later: **24 clusters is a pilot-scale trial for a moderate effect and a full-scale trial only for a large one, and the preregistration must say which claim it is entitled to make.**

Time-series length, since T1 requires held-out prediction and not a pre-post comparison:

```
periods   cadence        weeks   years    train / hold out
  26      fortnightly      52    1.00        17 / 9
  52      fortnightly     104    2.00        34 / 18
  52      weekly           52    1.00        34 / 18
 104      weekly          104    2.00        69 / 35
```

Fifty-two weekly periods with a 34/18 split is the minimum giving an 18-period held-out window, which the surrogate test of §4.7 needs for resolution at `B = 999`.

What would advance it

```
E1  T1: the frozen model beats all four baselines on held-out group order,
    with a calibrated surrogate test, in most clusters.

E2  T2: the framework-selected intervention outperforms the active control
    on the preregistered primary outcome.

E3  The headroom gauge PREDICTED which clusters would fail to grow.
    The headroom prohibition says that at h_t < 0 no interaction geometry prevents
    decline, so clusters with negative measured headroom should not have
    grown whatever was done to them. The sharpest external prediction the
    framework makes.

E4  T3: a measurable relation between members' internal parameters and the
    group's, with the breakdown cases identified and printed.
```

E3 is the strongest test available at group scope. **The headroom prohibition [Curzi, 2026b] is a prohibition, and prohibitions are the most falsifiable things a theory can offer: a single cluster with clearly negative headroom that nonetheless grew its order proxy substantially, with `eta`, `C`, and `N` unchanged, is a direct refutation.** No probabilistic hedge protects it.

What would kill the programme at this stage

```
K5.1  T1 FAILS: the model does not beat the baselines on held-out group
      order, with the surrogate calibrated.
      CONSEQUENCE: the storage law does not describe group dynamics on
      this timescale with these proxies. Report and stop, or return to
      proxy selection with a new preregistration.

K5.2  A CLEAN HEADROOM-PROHIBITION VIOLATION. A cluster with h_t well
      below zero grows its independently measured order proxy
      substantially, with eta, C, and N unchanged and the growth
      replicated in the held-out window.
      CONSEQUENCE: the sustainability ceiling is false as stated. This is
      the single most consequential possible finding of the whole
      programme, because the headroom prohibition \citep{ccp:CCP-008} is the most
      important sociological statement the framework makes. It would have to be
      withdrawn.

K5.3  T2 SUCCEEDS BUT T1 FAILS, REPEATEDLY.
      CONSEQUENCE: the framework is a heuristic that generates useful
      interventions, not a theory that predicts. Real, respectable, and
      to be printed in exactly those words: a design manual, not a
      science.

K5.4  T3 FAILS ENTIRELY. No relation between internal and group
      parameters, in any cluster, on any measure.
      CONSEQUENCE: the correspondence is an analogy of form, not
      of substance. the correspondence thesis reduces to a claim about
      MODEL FORM only, which is all the framework claims for it.
      Survivable, and a large retraction of the volume's ambition.

K5.5  HARM. Any cluster in which the intervention measurably reduces
      members' ability to decline participation, exit, or dissent.
```

```
      CONSEQUENCE: stop immediately. the fourth standing refusal is not a
      methodological preference and does not trade off against a result.
```

Scale and timescale

```
clusters           24 (12 per arm)
members            20 per cluster; 480 total
periods            52 weekly
observations       24 x 20 x 52 = 24,960 member-periods
participant burden 4 minutes per member-period = 1,664 participant-hours
fieldwork          52 weeks
total stage        30 months: 6 site recruitment, 6 setup, 12 fieldwork,
                   6 analysis and dissemination

researcher and coordinator hours:
   site coordination (24 x 52 x 2 h)              2496.0
   site recruitment (24 x 40 h)                    960.0
   instrument deployment and IT                    400.0
   analysis                                        700.0
   writing and dissemination                       300.0
   ------------------------------------------------
   TOTAL                                          4856.0
```

27.8.10 8.10 The kill register

Every kill condition in one table, so no stage is entered without its exit already written down.

Code	Stage	Kill condition	Consequence
K1.1	1	No shared unit for D and P	Moment equation demoted to heuristic ordering
K1.2	1	No instrument reaches session reliability 0.80	Programme cannot proceed
K1.3	1	Limiting cases have no recognisable referent	FVP-3 fails; formula has no phenomenon
K1.4	1	Y is a repackaging of U	Return to Stage 0 or stop
K2.1	2	No stable exponent across sessions	No scale; FVP-2 fails
K2.2	2	Cross-modality transitivity violated	Not a ratio scale; decisive
K2.3	2	No true zero for D or P	Division impermissible; demote in print
K2.4	2	D and P in incompatible units	As K1.1, quantified
K3.1	3	Any of the three rules fires in most subjects	Storage law does not describe internal dynamics
K3.2	3	Split-half reliability of b_hat below 0.50	eta is a window property; Stage 4 forbidden

Code	Stage	Kill condition	Consequence
K3.3	3	Systematic same-direction drift in every subject	Constant-parameter postulate fails
K3.4	3	Design controls failed; subjects inferred hypothesis	Data uninterpretable (52.7 per cent false-pass rate)
K4.1	4	Between-person spread does not clear the noise floor	delta and eta are not person parameters
K4.2	4	Rank order of subjects disagrees between labs	**FVP-7 fails. Programme stops.**
K4.3	4	Between-lab disagreement is a pattern, not a constant	As K4.2
K4.4	4	Suite passes in one lab, fails in the other	Do not proceed; find the difference or stop
K5.1	5	Model does not beat baselines on held-out group order	Storage law does not describe group dynamics here
K5.2	5	Clean headroom-prohibition violation	Sustainability ceiling withdrawn
K5.3	5	Intervention works, model does not predict	Design manual, not a science; print those words
K5.4	5	No internal/external relation at all	Correspondence is form only
K5.5	5	Ability to decline reduced in any cluster	Stop immediately; non-negotiable

Twenty-one ways to die: four terminal for the whole enterprise (K1.1 in its strong form, K3.2, K4.2, K4.3), and one — K5.5 — that stops the work regardless of what the data say.

27.8.11 8.11 Survival arithmetic

The following probabilities are **declared priors, badged as priors and not as evidence.** They are stated so the expected cost can be computed at all; a reader who disagrees should substitute their own and redo the arithmetic, which takes one line.

```
stage                       months   participants   P(fail here)   hours
1  pilot operationalisation      7          1            0.10        376.5
2  ratio-scale construction      4          6            0.10        778.9
3  single-subject stability     12          3            0.15        831.5
4  cross-subject recovery       18         80            0.35       1712.0
5  field trials, group scope    30        480            0.20       4856.0
-------------------------------------------------------------------------
TOTAL (strictly sequential)     71        570                       8554.8

SURVIVAL CURVE

  after stage 1   P(still alive) = 0.9000
  after stage 2   P(still alive) = 0.8100
```

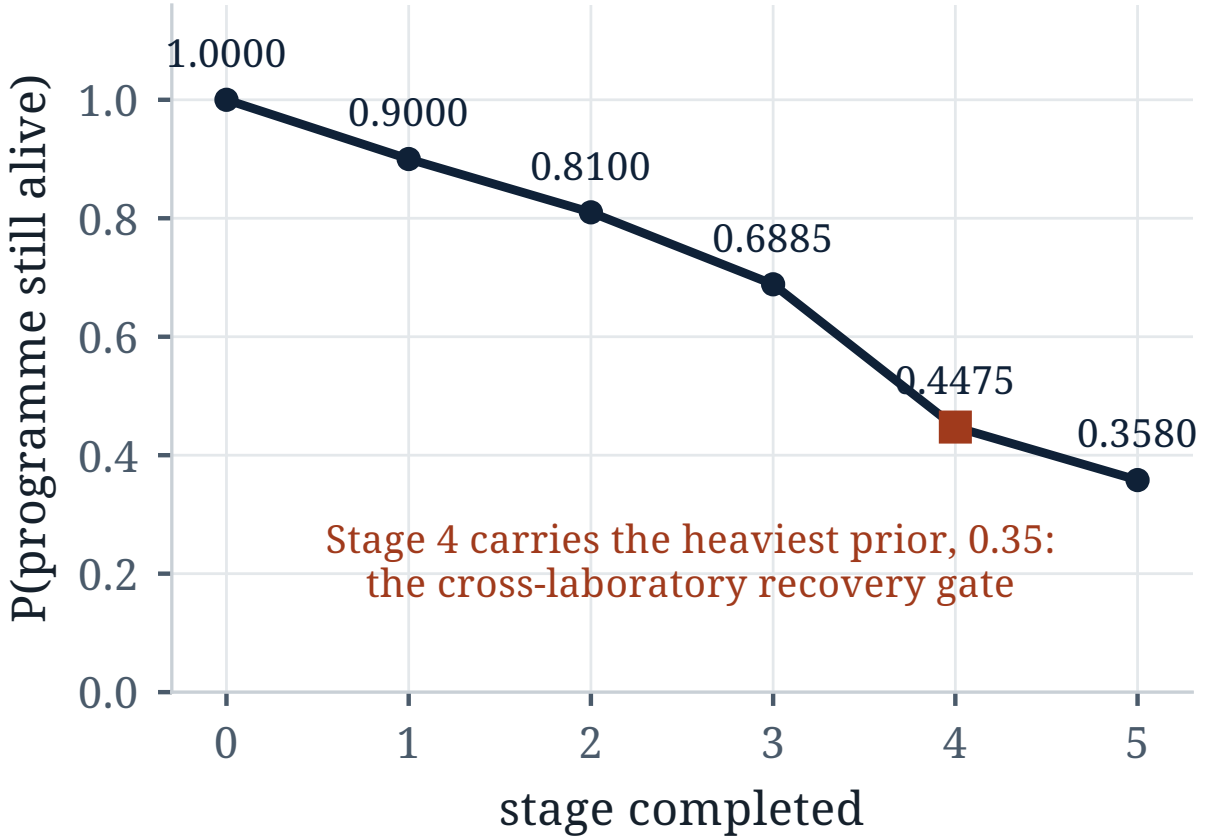


Figure 27.3: **Under its own declared priors the programme is more likely to be dead than alive before the field trials begin.** The compounded probability of reaching a completed final stage is 0.3580, which at 5.35 FTE-years per attempt implies 14.93 FTE-years per completed programme. Modelled from the declared per-stage failure priors 0.10, 0.10, 0.15, 0.35, 0.20 of §8.11; these are priors, badged as priors, and are not evidence. Produced by `figs/src/CCP-027_fig03.py`.

```
after stage 3   P(still alive) = 0.6885
after stage 4   P(still alive) = 0.4475
after stage 5   P(still alive) = 0.3580

P(programme dies at or before stage 4) = 0.5525
```

Under these priors the programme is more likely than not to be dead before the field trials begin, and its probability of reaching a completed Stage 5 is 35.8 per cent. That is not pessimism. It is what a research programme looks like when its kill conditions are written honestly beforehand, and a programme with a high stated completion probability has almost certainly written them too loosely.

The heaviest prior — 0.35 at Stage 4 — is deliberate. Stage 4 is where a formalism that describes meets the demand that it measure. the framework names FVP-7 the gate most often skipped [Curzi, 2026t]; **it is most often skipped because it is most often failed.**

27.8.12 8.12 Resourcing

```
researcher and coordinator hours        8,554.8
at 1,600 productive hours per FTE-year      5.35 FTE-years
calendar, strictly sequential                 71 months = 5.92 years
```

```
calendar, with 30 per cent overlap of
  stages 2–3 and 4–5                          61 months = 5.08 years
implied continuous staffing        5.35 / 5.92 = 0.90 FTE
```

The programme is one full-time researcher for six years. That is its actual shape: one person for six years, or two for three, with participants supplying a further 2,791 hours of their own time. Participant burden by stage:

```
stage 1     56.3 participant-hours       1 person
stage 2     38.9                         6 people
stage 3     72.0                         3 people
stage 4    960.0                        80 people
stage 5  1,664.0                       480 people
-------------------------------------------------
TOTAL    2,791.2 participant-hours     570 people
```

Monetary cost is not stated as fact because a defensible rate is not available here. **As a declared assumption and not a citation**, at fully loaded hourly rates of 60, 90, and 120 in some currency unit, 8,554.8 researcher-hours cost 513,290, 769,935, and 1,026,580. Participant compensation, equipment, and site costs are additional and are not estimated, because doing so would require numbers this paper does not have. The figure that matters to a funder:

```
5.35 FTE-years / 0.3580 = 14.93 FTE-years per COMPLETED programme
```

About fifteen FTE-years per completed programme, because most attempts stop early. The stopping is the point: nine or ten of those years are spent learning that something does not work, which is the ordinary and correct ratio in any field that tests things.

27.8.13 8.13 Tripwires for a degenerating programme

A programme degenerates when its response to disconfirmation is to add assumptions that protect the core rather than predictions that expose it. The distinction between a progressive and a degenerating research programme is Lakatos [1970]'s, and it is the right frame for badged postulates: what makes a protective belt legitimate is that it generates new testable content [Popper, 1959]. The practical discipline is to design each stage so that it discriminates between rival explanations rather than merely accumulating support for one [Platt, 1964], and to remember that a directional prediction is nearly costless in psychology, which is why this framework's point predictions with constants are worth more than its signs [Meehl, 1967, 1990]. Each item below is a tripwire: if it happens, the programme has stopped being a research programme and someone must say so.

```
W1  A PARAMETER BECOMES TIME-VARYING AFTER IT FAILS TO BE CONSTANT.
    Legitimate: preregister a time-varying specification as a NEW
    prediction, on NEW data, with its own kill condition.
    Degenerate: refit existing data with delta_t and report the gain.

W2  N IS RE-AUDITED AFTER THE RESULT IS SEEN.
    N enters through ln N and through F, so a post hoc re-audit adjusts
    the ceiling to fit the data. Audit N before, freeze it, report it.

W3  THE PERIOD LENGTH IS CHANGED BECAUSE THE RESULT IMPROVES.
    One of the strongest levers on any autoregressive fit. Changing it
    after seeing results is a specification search wearing the clothes of
    a methodological refinement.
```

```
W4  A SUBJECT OR CLUSTER BECOMES "ATYPICAL" AFTER FAILING.
    Exclusion criteria are legitimate only if written in field A3 of the
    preregistration before recruitment.

W5  THE PROXY IS SWAPPED FOR A BETTER-BEHAVED ONE.
    The most common single rescue. Y is redefined, the fit improves, and
    nothing in the write-up records that a search occurred.

W6  A FAILED STAGE IS RECLASSIFIED AS A PILOT.
    A stage becomes a pilot when it was preregistered as a pilot, not when
    it produced an unwelcome number.

W7  THE KILL CONDITION IS RENEGOTIATED IN THE LIGHT OF THE DATA.
    The most serious of the seven, because it is the one that dissolves
    all the others. A kill condition altered after the data are seen was
    never a kill condition.
```

The countermeasure for all seven is structural, not moral: publish the kill conditions with the preregistration, before the data exist, where the investigators cannot edit them. The principle is older than the replication crisis: a scheme should remain sound when everything about it except the private key is public [Kerckhoffs, 1883], and disclosure is a control rather than a courtesy [Brandeis, 1914]. A programme that has done that can still degenerate, but not quietly.

27.8.14 8.14 What partial success looks like

Programs rarely die cleanly or succeed completely. Five partial outcomes, named in advance so they are recognised as outcomes rather than argued about as failures:

```
PS1  STORAGE LAW SURVIVES, MOMENT EQUATION DOES NOT.
     K1.1 or K2.3 fires; delta and eta recover on a directly measured Y.
     The storage-law material stands; the moment equation is demoted to diagnostics. The most
     probable partial outcome.

PS2  PARAMETERS ARE SPECIES-TYPICAL, NOT INDIVIDUAL.
     K4.1 alone fires. The framework describes a process everyone runs
     with similar constants. Useful, publishable, and requires rewriting
     every "your delta" in the volume.

PS3  IT WORKS INTERNALLY AND NOT EXTERNALLY, OR THE REVERSE.
     K5.4 fires. Two separate models with the same algebra, which is
     precisely what the framework already claims and no more.

PS4  IT PREDICTS BUT DOES NOT GUIDE.
     T1 succeeds, T2 fails. A descriptive theory with no demonstrated
     intervention value. Honest and unglamorous.

PS5  IT GUIDES BUT DOES NOT PREDICT.
     T2 succeeds, T1 fails -- K5.3. A design manual whose heuristics work
     for reasons other than those stated. Print it in those words: a
     manual that works for unknown reasons is still a manual, and is not
     a science.
```

PS5 is the outcome most likely to be misreported, because it is the one where everything appears to be going well. Working interventions are the most persuasive and least diagnostic evidence available; the apparatus of Section 4 exists to keep those two apart.

27.8.15 8.15 What the programme would not establish even if every stage succeeded

The framework's four standing refusals are not suspended by success.

```
NOT anything about physics, thermodynamics, or consciousness. The surplus
functional is pure mathematics and makes no empirical or metaphysical claim;
the framework proves no negentropy law.

NOT that surplus is information in the Shannon sense. The framework declines
that identification explicitly.

NOT the conversion of the framework's badged postulates into theorems. A
postulate that survives testing is a well-tested postulate.

NOT any intervention that removes a person's ability to decline. That
refusal is prior to the evidence and unaffected by it.
```

A programme that cleared every gate would establish exactly this: that `delta` and `eta` are measurable person-level and group-level constants, in ratio-scaled units, reproducible across laboratories, whose values predict held-out order better than persistence, a linear input, autoregression, and the input's own timing scrambled. That is a great deal, and much less than the framework's ambitions. Saying so is what keeps the ambitions honest.

27.9 9 Specification of the programme

```
S8.1   Stage 0's four preconditions (SP1–SP4) are in writing before any
        recruitment. Costless refusal is among them.

S8.2   Stages run in order. No stage begins before its predecessor's
        advance conditions are met and published.

S8.3   Every stage's kill conditions are published with its
        preregistration, before data exist, where the investigators cannot
        edit them.

S8.4   Stage 1 runs at least 45 sessions -- 41 is the minimum at which an
        observed test-retest r = 0.70 has a 95 per cent lower bound above
        0.50 -- and reports session reliability with the item count
        required to reach it (Spearman–Brown).

S8.5   Stage 1 demonstrates or refutes a shared unit and a common zero for
        D and P, and prints whichever it found.

S8.6   Stage 2 spans at least 2.0 decades of stimulus range. Widening the
        range takes precedence over lengthening the session.

S8.7   Stage 2 collects enough judgments for SE(exponent) <= 0.02
        (n >= 469 per exponent at a 2.0-decade span) before any transitivity
        claim is made.

S8.8   Stage 2's transitivity test states its detectable threshold
```

explicitly alongside its result.

S8.9 Stage 2 demonstrates a true zero for D and for P, or the Moment equation is demoted in print per the formula-validation protocol \citep{ccp:CCP-026}.

S8.10 Stage 3 samples daily, not twice weekly. The precision-per-week arithmetic is decisive and is reported.

S8.11 Stage 3 runs at least 240 periods per subject, split 120/120, and reports split-half reliability of a_hat and b_hat.

S8.12 Stage 3 reports the instrument's rating noise and the implied attenuation of delta and eta from §4.9.

S8.13 Stage 3 states its power to detect within-person parameter change before claiming any such change.

S8.14 Stage 4 is run by two laboratories independently from the published protocol. A single-lab Stage 4 does not satisfy FVP-7 and must not be described as if it did.

S8.15 Stage 4 simulates the estimation noise floor at its own n and T and compares the observed between-person spread against it explicitly, before any individual-differences claim.

S8.16 Stage 4 publishes the reliability of the person-level estimate and the ceiling it places on correlations with external variables.

S8.17 Stage 4 classifies any between-lab disagreement as either a single scaling constant (calibration) or a pattern (theory failure), and prints which.

S8.18 Stage 5 runs both T1 (prediction) and T2 (intervention). Neither alone is sufficient.

S8.19 Stage 5 states its powered effect size given its cluster count and assumed ICC, and claims only what that supports.

S8.20 Stage 5's time series is at least 52 periods with a held-out window of at least 18, so the surrogate test has resolution at B = 999.

S8.21 Stage 5 tests the headroom prohibition as a prohibition: clusters with negative measured headroom are predicted not to grow.

S8.22 All seven degeneration tripwires W1-W7 (§8.13) are monitored by someone with standing to halt the programme.

S8.23 Null results are published with titles that say so.

S8.24 The survival priors of §8.11 are restated at each stage boundary and updated in public.

27.10 10 Failure modes of the programme

27.10.1 F8.1 Stage Skipping

Signature. A study estimating `eta` for forty people that never built a ratio scale, or a field trial launched without any single-subject evidence. Stages 4 and 5 are the interesting ones. Stages 1 and 2 are apparatus.

Fix. S8.2. The estimate produced by a skipped-stage study is a number in unknown units, and unknown units cannot be compared to anything, which means the study cannot succeed even in principle.

27.10.2 F8.2 The Under-Powered Ratio Scale

Signature. A transitivity test reported as passed, with exponent standard errors near 0.10. Magnitude estimation is quick per judgment, so the total is underestimated.

Fix. S8.7 and S8.8. At `SE = 0.10` per exponent, only a discrepancy above 0.37 is detectable — the test cannot fail, which means it cannot pass either.

27.10.3 F8.3 Sampling Too Slowly

Signature. Twice-weekly experience sampling adopted to reduce burden. It genuinely does reduce burden.

Fix. S8.10. `T = 120` at daily sampling is 3.9 months; the same `T` twice-weekly is 13.8 months. The extra ten months are ten months of parameter drift, attrition, and life events, all of which attack the constant-parameter postulate that Stage 3 is trying to test.

27.10.4 F8.4 The Single-Lab FVP-7

Signature. “Replicated across two independent samples” where both samples were collected by the same group with the same hands. It is much easier and it feels like replication.

Fix. S8.14. The whole content of FVP-7 is that the procedure survives a different group’s reading of the written protocol. Two samples from one lab test sampling variability, which was never the question.

27.10.5 F8.5 Individual Differences From Estimation Noise

Signature. A per-subject `eta` table with a spread of 0.03 to 0.04 at 120 periods, read as personality. The numbers differ, so the people appear to differ.

Fix. S8.15. The simulated floor is 0.036 at `T = 120` and 0.048 at `T = 60`.

27.10.6 F8.6 The Endpoint-Only Field Trial

Signature. A cluster-randomised trial with a baseline and a follow-up and nothing between. It is the standard design and it is cheaper.

Fix. S8.20. The framework’s claim is about a *trajectory* governed by a recursion. A two-point design cannot distinguish the storage law from any other process that ends higher, and it certainly cannot run a surrogate test.

27.10.7 F8.7 Kill-Condition Erosion

Signature. The published kill condition and the one discussed in the group meeting have drifted apart; the discrepancy is nobody's fault and everyone knows about it. Sunk cost, plus the fact that every individual relaxation is small.

Fix. S8.3 and S8.22. Publish the conditions where they cannot be edited, and give someone standing to halt.

27.10.8 F8.8 Success Theatre at Stage 5

Signature. The intervention worked; the prediction test is described briefly in a supplement. PS5. A working intervention is the most persuasive and least diagnostic result available.

Fix. S8.18 and PS5's printing requirement. If T2 succeeded and T1 failed, the sentence "this is a design manual whose heuristics work for reasons other than those stated" must appear.

27.10.9 F8.9 The Rescued Prohibition

Signature. A cluster with negative headroom grew, and the response is to re-measure `C`. the headroom prohibition [Curzi, 2026b] is the framework's strongest sociological statement in the framework, and important statements attract rescues.

Fix. W2 and W5 of §8.13, plus S8.21. Freeze `C`, `eta`, and `N` before the window opens. If the violation survives a frozen measurement and replicates in held-out data, the corollary is withdrawn. That is what withdrawal means.

27.10.10 F8.10 Measuring People Who Cannot Refuse

Signature. Participation is enrolled at the organisational level and individual opt-out carries a cost, however small. Cluster designs recruit clusters.

Fix. S8.1 and K5.5, and the fourth standing refusal [Curzi, 2026]. This one does not trade off against statistical power, and a design that cannot meet it is a design that does not run.

27.11 11 A worked schedule and its five decision gates

One laboratory, one principal investigator at 0.9 FTE, start date month 0. Every figure below is the arithmetic of §8.5 through §8.12 laid on a calendar; **the schedule is a plan, not a record — no month of it has been executed.**

27.11.1 11.1 The schedule

```
MONTH   ACTIVITY                                              GATE / OUTPUT
-------------------------------------------------------------------------
 0– 2   Stage 0: write SP1–SP4. Ethics. Consent                 entry condition
        infrastructure with costless refusal.
 0– 4   Stage 1: instrument development, three                 160 h
        candidates each for D, P, E.
 3– 7   Stage 1 fieldwork: 45 sessions, 3/week,                56.3 subject-h
        1 subject + 1 parallel replication subject.
 6– 7   Stage 1 analysis. Publish reliability,                 K1.1–K1.4
        shared-unit argument, limiting-case audit.            GATE 1
-------------------------------------------------------------------------
```

```
 7– 9   Stage 2: apparatus for 3 calibration            240 h
        modalities. Modulus selection.
 8–10   Stage 2 fieldwork: 6 subjects, 2,916            38.9 subject–h
        judgments each, 6 sessions per term.
10–11   Stage 2 analysis: exponents, transitivity,      K2.1–K2.4
        true-zero test. Publish.                       GATE 2
-------------------------------------------------------------------------
11–13   Stage 3: data pipeline, surrogate              320 h
        calibration simulations at b = 0.
12–20   Stage 3 fieldwork: 3 subjects, 240 daily       72 subject–h
        periods, 4 prompts/day (34.3 weeks).           2,880 prompts
20–23   Stage 3 analysis: per-subject fits, freeze,    K3.1–K3.4
        held-out, split-half. Publish.                 GATE 3
-------------------------------------------------------------------------
23–26   Stage 4: publish protocol. Recruit lab 2.      150 h
        Recruit 40 + 40 subjects.
26–30   Stage 4 lab 1 fieldwork: 120 daily periods.    480 subject–h
29–33   Stage 4 lab 2 fieldwork, overlapping.          480 subject–h
33–41   Stage 4 analysis: noise floor at own n and T,  K4.1–K4.4
        recovery, between-lab reconciliation. Publish. GATE 4
                                                       MAKE OR BREAK
-------------------------------------------------------------------------
41–47   Stage 5: site recruitment, 24 clusters.        960 h
47–53   Stage 5 setup, instruments, randomisation,     400 h
        preregistration of T1, T2, T3.
53–65   Stage 5 fieldwork: 52 weekly periods.          1,664 participant–h
65–71   Stage 5 analysis and dissemination.            K5.1–K5.5
                                                       GATE 5
-------------------------------------------------------------------------
TOTAL   71 months = 5.92 years, 8,554.8 researcher–hours,
        570 participants, 2,791.2 participant–hours
```

With the overlap shown at months 29-33 and 41-47, the calendar compresses to 61 months, 5.08 years.

27.11.2 11.2 The decision points, with their arithmetic

Gate 1, month 7. Does any operationalisation of `D` and any of `P` share a unit with a common zero? If no, the programme does not stop — it *narrows*: the Moment equation is demoted in print, the Self-Storage Law is refitted with `S_t = F(U_t)` (input variant A, §4.1) instead of `M_t`, and Stages 2 through 5 proceed on the narrower claim. **This is the one gate whose failure produces a smaller programme rather than no programme**, and recognising that in advance is what stops a Stage 1 failure from being fought instead of accepted.

Gate 2, month 11. Transitivity, detectable threshold 0.0739 in the exponent given `SE = 0.02` per exponent. A measured `|a_AC − a_AB ∗ a_BC|` above it ends the ratio-scale claim.

Gate 3, month 23. Three subjects, three verdicts. Advance requires the three-rule suite cleared in at least two of three, and split-half reliability of `b_hat` at or above 0.80 in at least two of three. **The simulated benchmark at 120 periods per half is 0.846; a real figure near 0.85 clears, near 0.40 does not, and the region between 0.50 and 0.80 is the genuinely awkward one — for which the preregistered rule is to extend to 360 periods in a fourth subject rather than to argue.**

Gate 4, months 33-41. Three numbers decide it:

```
observed between-person sd of eta      must exceed  0.0355   (noise floor at T=120)
```

```
between-lab rank correlation of eta    must exceed  0        significantly,
                                                    and lie at or below the
                                                    published reliability ceiling
                                                    (0.942 at between-sd 0.10)
between-lab mean ratio of eta          must equal   1/g for an independently
                                                    measured operationalisation
                                                    difference g, or be absent
```

If the rank correlation is near zero, the programme stops at month 41 having spent 3,698.9 researcher-hours — 2.31 FTE-years — and having established something real and publishable: that `eta` as operationalised is not a person-level measurement. **That is a worthwhile purchase, and it is the single most probable outcome under §8.11's priors at 0.35. Budget for it as the expected case, not the feared one.**

Gate 5, month 71. T1, T2, T3, and the headroom prohibition. As scheduled — 24 clusters, 20 members — the trial is powered for a standardised effect of 0.50 and under-powered for 0.30 at `ICC = 0.05`, where 36 clusters would be needed. The preregistration therefore claims a large effect or claims pilot status, not both.

27.11.3 11.3 The honest summary a funder should be given

> *We propose a five-stage, 71-month programme, requiring approximately 5.35 FTE-years of researcher time and 570 participants contributing about 2,791 hours, to determine whether two parameters of a formal model of stored order — a retention loss `delta` and a conversion efficiency `eta` — can be measured on ratio scales, recovered as person-level constants, and reproduced across laboratories. Twenty-one preregistered kill conditions are published in advance; four are terminal. Under our declared priors, which are priors and not evidence, the programme has a 35.8 per cent probability of reaching a completed final stage and a 55.3 per cent probability of stopping at or before the cross-laboratory recovery stage. Expected cost per completed programme, accounting for early stopping, is about 14.9 FTE-years. We will publish at every stage, including and especially when the stage kills the programme, and null results will be titled as null.*

That paragraph is the whole proposal. Concrete, sequenced, every number computed from a stated assumption, naming the year at which the whole thing most probably ends and committing in advance to publishing that ending.

A programme that cannot be killed at any stage is not a research programme. This one can be killed at twenty-one places, is most likely to die at month 41, and says so on the first page of its own proposal.

27.12 12 Correspondence tables as relational graphs

This section is the part of the paper a careful reader will approach with the most suspicion, and that suspicion is correct and should be kept. What follows is not a departure from the formalism. It is an application of the formalism to a body of historical technique, and the only things it is permitted to say are things the formalism can carry.

The position taken here – recorded as a thesis, with its badge on – is that the ancient occult and alchemical formulas are these same relationships expressed in a different notation. That is a claim about *notation*, and it is falsifiable in the only sense a claim about notation can be: either the old structures translate into the modern objects without residue, or they do not. This section performs the translation for correspondence systems. The companion paper performs it for the star-polygon figure, the memory palace and the ritual calendar, and draws the balance [Curzi, 2026v].

Nothing in this section asserts any supernatural claim, any physical effect of a symbol, or any result that is not derivable from the formalism already established in the preceding sections and the framework papers they cite. That sentence is a standing constraint on Sections 12 to 15, and any sentence in them that appears to contradict it is an error to be reported and corrected.

What *is* claimed here is narrower and more interesting. A correspondence table is a mathematical object of a completely ordinary kind: a finite relation over finite sets. The Tree of Life is a graph with ten vertices and twenty-two edges. The tarot is a structured set of seventy-eight elements with an explicit product decomposition. Gematria is a map from strings to integers — which is what a hash function is, and calling it one is a technical description, not a compliment and not an insult. These readings are descriptive. They involve no interpretation, no belief, and no metaphysics. They are what those objects *are*, and once you have said so you can compute with them.

The engineering payoff is specific. The framework's validity conditions require a category assignment before the surplus functional may be applied to anything at all. You cannot compute `N_eff` until you have said what the `N` categories *are*. **A correspondence table is precisely a category assignment, written down, fixed, and made auditable — which is why the tradition that produced them was doing, in a different notation, the first step of the procedure the framework specifies.**

27.12.1 12.1 Statement of the reading rule

READING RULE (Correspondence Tables). Let a correspondence system be given: a finite set of labels `L` and one or more finite relations among them. Then:

```
a single-column list of N labels          IS   a finite set of size N
a two-column table over the same labels   IS   a binary relation R subset of L1 x L1
an m-column table, one row per label      IS   (m-1) functions L1 -> L2, ..., L1 -> Lm
a diagram of nodes joined by paths        IS   an undirected graph G = (V, E)
an assignment of labels to categories     IS   a partition, i.e. a block decomposition
a letter-value scheme summed over a word  IS   a hash h : strings -> integers
```

There is no inference in that block. Each line is a restatement in modern vocabulary of what the historical artefact literally contains. **The reading rule is a translation, not a theory, and it cannot be wrong in the way a theory can be wrong — it can only be inapplicable, which happens exactly when the historical artefact is too vague to have a definite extension.** That case is common and is treated as failure mode F12.5.

The status table for this section:

Element	Status	Warrant
The Kircher-arrangement Tree of Life has 10 vertices, 22 edges	**PROVED** (enumeration)	§12.4, computed
Its density against `K10` is 22/45 = 0.4889	**PROVED** (arithmetic)	§12.4, computed
Its diameter is 3, mean geodesic 1.5778, cyclomatic number 13	**PROVED** (computation)	§12.4, computed
15 of its 22 edges cross pillar boundaries	**PROVED** (enumeration)	§12.5, computed

Element	Status	Warrant
The tarot is `22 + 4 x 14 = 78`, a set with a product substructure	**PROVED** (arithmetic)	§12.6
Gematria is a function from strings to integers	**PROVED** (definition)	§12.8
Gematria destroys 10.47 bits on a 234,454-word list	**PROVED** (measurement)	§12.8, computed
The weekday sequence follows from `24 mod 7 = 3` on the Chaldean order	**PROVED** (arithmetic)	§12.7, computed
A correspondence table satisfies the framework's category requirement	**PROVED** (definition)	§12.10
That these systems were *built as* mnemonic technology	**HISTORICAL CLAIM** (well-supported)	§12.9
That the old notation and the framework's notation describe the same relations	**POSTULATE** (author's thesis)	§12.1
That any correspondence carries efficacy, force, or physical consequence	**REFUSED**	standing refusal

The last two rows are the whole difficulty of this section, and they are separated on purpose. The thesis in the penultimate row is a claim about *isomorphism of structure*. The refused row is a claim about *action in the world*. **Structural correspondence is not evidence of mechanism, and this paper asserts no mechanism.** The same restriction governs the frequency framework's square law [Curzi, 2026], and it means exactly the same thing here.

27.12.2 12.2 The honesty clauses

These are mandatory and are restated wherever the reading rule is used in the companion paper [Curzi, 2026v].

H12-1 — Description, not endorsement. Saying that the Tree of Life is a graph is true in the same way that saying a subway map is a graph is true. It commits you to nothing about subways. Every metric computed in §12.4 is a metric of the *diagram as drawn*, and would be unchanged if the diagram had been drawn to organise plumbing fixtures.

H12-2 — The arrangement is a choice, and other arrangements exist. The ten-node, twenty-two-path figure used throughout this section is the arrangement commonly called the Kircher tree, reproduced in Athanasius Kircher's *Oedipus Aegyptiacus* (1652) and standard in later European occult literature. It is not the only historical arrangement; the Cordoveran and Lurianic schemes differ, and nineteenth- and twentieth-century revisions differ again. **All graph metrics reported here are metrics of the specific edge list printed in §12.4 and must be recomputed for any other arrangement.** An unstated arrangement produces unstated metrics, which is failure mode F12.4.

H12-3 — Gematria is a hash, and that is not a compliment. The description is exact and it is also deflationary. A hash function maps a large domain onto a small codomain. It therefore collides, massively and by construction. **Two words sharing a gematria value is the expected behaviour of a many-to-one map, not evidence about the words.** §12.8 measures the collision rate so that the deflation is quantitative rather than rhetorical.

H12-4 — Counting is not meaning. That the tarot has 78 elements, or that the Hebrew alphabet has 22 letters and the Kircher tree has 22 paths, are arithmetic facts. The nineteenth-

century assignment of the 22 trumps to the 22 letters — associated with Éliphas Lévi, *Dogme et Rituel de la Haute Magie* (1854-56), and elaborated by later societies — is a *chosen bijection*, one of `22!` available, made by identifiable people at an identifiable date. It is data about the history of the assignment and about nothing else.

H12-5 — Everything downstream inherits Papers A to E's clauses. When §12.5 computes a surplus number from a pillar decomposition, that number carries every honesty clause of the Surplus Law: the triangle inequality fails for geometric triples; subadditivity holds on `u` only; the functional as deployed is a **lower bound** on the ensemble's information; the surplus functional is pure mathematics and makes no empirical or metaphysical claim; no thermodynamic consequence follows.

H12-6 — The audit runs in both directions. If the translation reveals that an old table is a well-formed partition, say so. If it reveals that the table has duplicate rows, undefined membership, or an unjustified ordering, say that too, in the same voice. §12.11 does the second, and finds a 40.37% ceiling overstatement in a small worked case. A translation that only ever flatters the source is not a translation.

27.12.3 12.3 What the validity conditions require before anything may be computed

The framework's validity conditions [Curzi, 2026c] are the gate this section exists to service. Before the surplus functional `F(u) = ln(1 + (N−1)u)` may be applied to any real ensemble, four things must hold, and they must hold *in print*, not in the analyst's head:

```
V1  EXHAUSTIVE.   Every element of the domain falls in some category.
V2  EXCLUSIVE.    No element falls in two, or the overlap is declared and handled.
V3  COUNTED.      N is stated, and the N categories are genuinely distinct.
V4  ORDERED.      Where an ordinal is used -- an essence index, in the
                  sense the frequency framework gives that term -- the
                  ordering is justified, not arbitrary.
```

A correspondence table is a physical artefact whose entire job is to satisfy V1 through V4 for a domain, and to keep satisfying them across generations of users who cannot consult the scheme's author. That is a strong statement about what these systems were for, and §12.9 defends it on historical grounds. It is also the reason the material belongs in an engineering paper: the same four gates govern a hospital's incident-category scheme, a warehouse's fault taxonomy, and a research group's list of domains, and all three fail in the same ways.

27.12.4 12.4 The Tree of Life is a graph, and here are its numbers

The *Sefer Yetzirah*, the earliest text in the tradition to use the phrase, speaks of thirty-two paths — ten and twenty-two. The later diagrammatic tradition renders these as ten nodes joined by twenty-two edges. Take that diagram at face value and compute.

Vertices, in the conventional numbering:

```
1 Kether     2 Chokmah    3 Binah      4 Chesed     5 Geburah
6 Tiphareth  7 Netzach    8 Hod        9 Yesod     10 Malkuth
```

Edges, the Kircher arrangement, listed in full so the computation is reproducible:

```
1–2   1–3   1–6
2–3   2–4   2–6
```

```
3–5   3–6
4–5   4–6   4–7
5–6   5–8
6–7   6–8   6–9
7–8   7–9   7–10
8–9   8–10
9–10
```

Count: 3 + 3 + 2 + 3 + 2 + 3 + 3 + 2 + 1 = 22. No duplicates; verified by set comparison.
Computed metrics:

```
|V| = 10                          |E| = 22
sum of degrees = 44 = 2|E|        handshake lemma holds
K10 edges = 45                    density = 22/45 = 0.488889
cyclomatic number = |E|−|V|+1 = 13    (independent cycles)
connected: yes                    diameter = 3
mean shortest path = 1.577778     distance histogram: {1: 22, 2: 20, 3: 3}
```

Degree sequence:

Node	Name	Degree
6	Tiphareth	8
7	Netzach	5
8	Hod	5
2	Chokmah	4
3	Binah	4
4	Chesed	4
5	Geburah	4
9	Yesod	4
1	Kether	3
10	Malkuth	3

Sorted: `[8, 5, 5, 4, 4, 4, 4, 4, 3, 3]`.

Three observations, each of which is a fact about the diagram and about nothing else.

First, the degree distribution has a single dominant hub. Node 6 carries 8 of a possible 9 incident edges; it is adjacent to every node except Malkuth. The next-highest degree is 5. In the vocabulary of Albert et al. [2000], this is a hub topology: robust to random edge loss, fragile to targeted removal of the hub. The reading is a statement about the degree sequence and nothing more; a heavy-tailed degree sequence is neither necessary nor sufficient for engineered robustness [Doyle et al., 2005], and **no scaling or power-law claim is made here** – with ten vertices there is no distribution to fit, and the procedure that such a claim would require [Clauset et al., 2009] does not apply. Delete node 6 and the diameter rises and the graph's mean path length degrades sharply; delete any leaf-adjacent node and little changes. That is a structural property of the drawing.

Second, the graph is far from a tree. A tree on 10 nodes has 9 edges. This has 22, hence 13 independent cycles. Alexander [1965] distinguishes the tree – where subsets nest strictly – from the semilattice, where they overlap. **The diagram called the Tree of Life is not, in Alexander's sense, a tree; it is a densely cyclic graph with an overlap structure, and that is why it can encode more relations than a hierarchy of the same size.** This is a pleasant terminological accident and it is also the substantive point: 13 independent cycles is a lot of overlap to carry in ten nodes.

Third, the endpoint-to-endpoint structure is rich. Distance from node 1 to node 10 is 3, and there are exactly 3 distinct geodesics of that length. The number of *simple* paths (no

repeated vertex) from 1 to 10 is **573**, computed by exhaustive enumeration. A system whose users are taught to "travel the paths" has 573 distinct non-repeating routes between its two extremes. As a memory structure, that is the property being exploited: many routes to the same destination means many retrieval cues for the same content, which is exactly the design principle behind the method of loci, which the companion paper treats as the storage law with space as the retention medium [Curzi, 2026v].

27.12.5 12.5 The three pillars are a block decomposition, and it is measurable

The same diagram carries a standard partition into three columns:

```
Pillar of Severity   {3 Binah,  5 Geburah,  8 Hod}          size 3
Pillar of Mercy      {2 Chokmah, 4 Chesed,  7 Netzach}      size 3
Middle Pillar        {1 Kether, 6 Tiphareth, 9 Yesod, 10 Malkuth}   size 4
```

This is a partition: exhaustive, exclusive, three blocks summing to ten. It satisfies V1, V2 and V3 of §12.3 without amendment. **It is therefore a legitimate block decomposition in the sense the framework's block-decomposition theorem requires [Curzi, 2026], and the cross-block and within-block edge counts can be measured rather than assumed.**

Measured:

```
within-pillar edges:   Middle 3,  Mercy 2,  Severity 2       total  7
between-pillar edges:                                        total 15
cross-block fraction:  15/22 = 0.681818

cross-pair breakdown:  Mercy--Middle     6
                       Middle--Severity  6
                       Mercy--Severity   3

saturation:  Severity   2 of max 3 internal edges
             Mercy      2 of max 3
             Middle     3 of max 6
```

Two facts fall out. The blocks are internally *sparse* — no pillar is a complete subgraph — and the wiring is overwhelmingly *across* blocks, 68.2% of all edges. The two outer pillars connect to the middle six times each but to one another only three times. The diagram routes its cross-side traffic through the centre.

Now the translation into the surplus functional. The block-decomposition theorem [Curzi, 2026], for a pure block vector, gives

```
u = u_cross + u_div          u_cross = beta^2          u_div = alpha^2 sin^2 gamma

max(f(u_cross), f(u_div))  <=  f(u)  <=  f(u_cross) + f(u_div)
```

where `beta^2` is the between-category share and `alpha^2 sin^2 gamma` the within-category share. **Badge: the identification of the measured cross-block edge fraction with `beta^2` is a modelling choice, not a theorem.** It is the natural choice — the share of the structure that lives across categories — but the theorem does not require it, and a different operationalisation of `beta` gives different numbers. With `beta^2 = 15/22 = 0.681818` and `alpha^2 = 7/22 = 0.318182`, sweeping the within-block angle `gamma`, at `N = 10`:

gamma	u	N_eff	F (nats)	% of ceiling ln 10 = 2.302585
15 deg	0.703132	7.3282	1.991729	86.50%
30 deg	0.761364	7.8523	2.060803	89.50%
45 deg	0.840909	8.5682	2.148056	93.29%
60 deg	0.920455	9.2841	2.228302	96.77%
90 deg	1.000000	10.0000	2.302585	100.00%

And the subadditivity bracket, verified numerically at a representative point (`beta^2 = 0.68`, `alpha^2 = 0.32`, `gamma = 30 deg`, `N = 10`):

```
u_cross = 0.6800   u_div = 0.0800   u = 0.7600   N_eff = 7.8400
max(f(u_c), f(u_d)) = 1.962908  <=  f(u) = 2.059239  <=  f(u_c)+f(u_d) = 2.505232     OK
```

The bracket holds at every sampled point, as it must. Note what the table shows: **once the cross-block share is above two thirds, the within-block angle barely matters — the figure is already at 86.5% of its ceiling at `gamma = 15 deg`.** That is strict concavity doing its work, and it is the same lesson as the conversion analysis [Curzi, 2026b]: the first crossing purchases most of what crossings can buy, and the rest is glazing bills.

The honest coda: `ln 10 = 2.302585` nats is a *hard ceiling* on surplus per interaction for any ten-category scheme [Curzi, 2026c]. No arrangement of paths raises it. Adding an edge to the diagram does not add a category. **The ceiling is set by `N`, and `N` is set by the table, not by the drawing.**

27.12.6 12.6 The tarot as a structured 78-element set

The deck is not a bag of 78 things. It has an explicit product structure, and the structure is the interesting part.

```
majors (trumps)         22
minors                  56  =  4 suits x 14 ranks
                            =  4 x (10 pips + 4 courts)
total                   78  =  22 + 40 + 16
```

Computed combinatorics:

```
naming one card:            log2(78)      = 6.285402 bits    ln(78) = 4.356709 nats
unordered pairs C(78,2)                   = 3,003
unordered triples C(78,3)                 = 76,076
ordered 3-card spread 78*77*76            = 456,456          log2 = 18.800116 bits
ordered 10-card spread P(78,10)           = 4,566,176,969,818,464,000
```

The product structure is exact in the log register, which is the register the framework operates in:

```
log2(4) + log2(14) = 2.000000 + 3.807355 = 5.807355 = log2(56)
```

That identity is trivial arithmetic and it is also the whole design principle of a factored code. **A suit-and-rank code lets a user reconstruct 56 identities from 4 + 14 = 18 memorised labels instead of 56, and the log-additivity is why it costs nothing in addressing power to do so.** The frequency framework develops the same move for essence sets [Curzi, 2026]; here it appears in a fifteenth-century object.

Block structure over the minors, taking the four suits as categories:

```
all minor pairs    C(56,2) = 1,540
within-suit pairs  4*C(14,2) = 364
cross-suit pairs             = 1,176      cross fraction = 0.763636
```

76.4% of minor-arcana pairs cross a suit boundary. Compare 68.2% for the Tree's pillar decomposition. Both structures put most of their relational mass across categories rather than within them, which is the geometry the surplus functional rewards — `u_cross = beta^2` is the between-category term and it is the term that survives when within-category variation collapses.

One arithmetic tie-in to the frequency framework's essence law [Curzi, 2026], offered as a structural remark and nothing more. The Essence Law's Result Two states that an ordered set indexed `1..N` yields exactly `ceil(N/2)` pitch classes, because pitch class depends only on the odd part of the index. Computed: for `N = 22`, 11 classes; for `N = 78`, 39 classes. **A 78-element ordered scheme has 39 distinct essence pitch classes, exactly half its size, and that is a hard structural limit of the indexing, not a property of the deck.** It carries every clause of that law, including its first honesty clause: essence, not resonance [Curzi, 2026].

27.12.7 12.7 The sevens, and the week as an arithmetic consequence

The planet/metal/day table is the cleanest correspondence artefact in the European tradition because all three columns have exactly seven entries and the assignment is a bijection.

Day	Planet	Metal
Saturday	Saturn	lead
Sunday	Sun	gold
Monday	Moon	silver
Tuesday	Mars	iron
Wednesday	Mercury	quicksilver
Thursday	Jupiter	tin
Friday	Venus	copper

As a relational object: two independent bijections on 7 labels. The information content of pinning them down:

```
7! = 5,040 bijections        log2(5,040) = 12.299208 bits per column pairing
a 3-column table pins 2 bijections:  2 * 12.299208 = 24.598416 bits
```

Twenty-four and a half bits is the entire informational content of a table that a tradition maintained for well over a millennium. That is a small number, and it is *supposed* to be a small number: the artefact's job is to be *stable and transmissible*, not to be dense. Compare the mnemonic-load computation in §12.9.

Now the part that is genuinely a derivation. The order of the days of the week looks arbitrary and is not. Take the seven classical planets in Chaldean order — descending apparent orbital period, which is the order they were ranked in by the sidereal periods known to antiquity:

```
Saturn, Jupiter, Mars, Sun, Venus, Mercury, Moon
```

Assign rulership of the hours of the day cyclically in that order, one planet per hour. The day is named for the ruler of its first hour. A day has 24 hours, and

```
24 mod 7 = 3
```

so the ruler of the first hour of the next day is **three steps on** in the Chaldean order. Stepping by 3 from Saturn and reading off:

```
Saturn -> Sun -> Moon -> Mars -> Mercury -> Jupiter -> Venus -> Saturn
```

which is

```
Saturday, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, and back to Saturday
```

Computed and verified: the eighth step returns to Saturn, the cycle closes, and every planet is visited exactly once. The surviving English names carry the Germanic interpretatio for four of the seven (Tiw for Mars, Woden for Mercury, Thor for Jupiter, Frigg for Venus); the Romance languages kept the Latin directly — *mardi*, *mercredi*, *jeudi*, *vendredi*. The earliest surviving explanation of the sequence by the planetary-hour argument is in Cassius Dio's *Roman History*, Book 37.

The weekday order is the star polygon {7/3}, and `gcd(7, 3) = 1` is exactly why the week is a single closed cycle that touches every planet rather than two or three disjoint sub-weeks. That is the companion paper's star-polygon theorem arriving early [Curzi, 2026v], in a piece of civil infrastructure that seven-day weeks are still keeping. If the day had been divided into a number of hours congruent to 0, 7, 14 or 21 modulo 7, the construction would have degenerated: every day would have had the same ruler. If it had been divided into 21 hours, `21 mod 7 = 0`, and the week would not exist. The week is a `gcd` result and the arithmetic is checkable in a minute.

27.12.8 12.8 Gematria is a hash function

Here is the operation. Assign each letter of an alphabet an integer. To evaluate a word, sum the integers of its letters. The standard Hebrew scheme (*mispar hechrachi*) assigns:

```
aleph 1   bet 2    gimel 3   dalet 4   he 5     vav 6    zayin 7   chet 8   tet 9
yod  10   kaf 20   lamed 30  mem 40    nun 50   samekh 60  ayin 70  pe 80   tsadi 90
qof 100   resh 200 shin 300  tav 400
```

Twenty-two letters; units 1-9, tens 10-90, hundreds 100-400; the values sum to 1,495.

```
h : strings over an alphabet A  ->  non-negative integers
h(w) = sum over letters of w of value(letter)
```

That is a hash function. A hash function is a map from a large or unbounded domain into a smaller codomain, computed by a fixed deterministic rule. The bits it destroys are Shannon's quantity and are computed as such below [Shannon, 1948, Hartley, 1928, Cover and Thomas, 2006]. The definition contains no requirement that the map be interesting, only that it be a map. Gematria satisfies it exactly. Saying so is a literally accurate technical description of the operation, and it carries no claim whatever that the map is meaningful.

The description is useful because it makes the properties computable. A hash has three properties that engineers care about, and this one's scores are measurable.

Property 1 — order invariance. Gematria is a *multiset* hash: `h` is a sum, and addition commutes, so any permutation of a word has the same value. Anagrams collide by construction, and so does any rearrangement, meaningful word or not. Measured, purely from the permutation effect:

word length L	strings over 22 letters	distinct multisets	permutation collisions alone
2	484	253	1.913x
3	10,648	2,024	5.261x
4	234,256	12,650	18.518x
5	5,153,632	65,780	78.346x

At five letters, the ordering information — which is most of what a written word *is* — has already been thrown away at a rate of 78 to 1, before any value-arithmetic collisions are counted.

Property 2 — codomain size. The value range grows linearly in word length while the domain grows exponentially, so the collision rate grows without bound. Computed exactly over the Hebrew value set:

L	strings `22^L`	distinct achievable sums	mean preimage	bits in	bits out
2	484	178	2.719	8.92	7.48
3	10,648	577	18.454	13.38	9.17
4	234,256	976	240.016	17.84	9.93
5	5,153,632	1,375	3,748.096	22.30	10.43
6	113,379,904	1,774	63,912.009	26.76	10.79

The input column climbs by 4.46 bits per added letter. The output column climbs by less than a bit. **By six letters the average value is shared by roughly sixty-four thousand strings.**

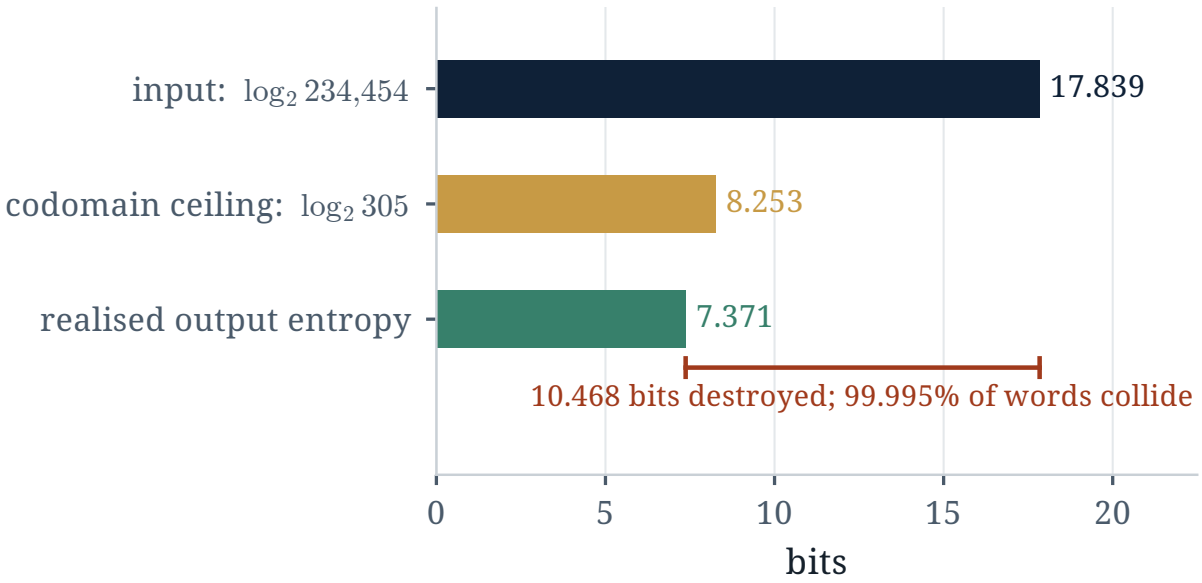
Property 3 — no avalanche. A cryptographic hash is designed so that changing one input bit changes about half the output bits. Gematria does the opposite: changing one letter changes the value by exactly the difference of the two letter values, a small, signed, entirely predictable amount. It is a *checksum*, and a weak one — it will catch a single-letter substitution but not a transposition, which is the most common scribal error there is.

Now the measurement that settles the matter, performed on a real word list rather than on a model. Using English simple gematria (A = 1 through Z = 26) over the macOS `/usr/share/dict/words` list (`web2`), restricted to purely alphabetic ASCII entries:

```
distinct words                     234,454
distinct hash values                   305
value range                          1..328
mean bucket size                   768.702
largest bucket                value 102 -> 2,353 words
top five buckets      102:2353  107:2325  97:2305  100:2295  115:2272
singleton buckets (no collision)        12
fraction of words colliding       0.999949

input entropy   log2(234,454)  =  17.839 bits
output entropy  H(hash)        =   7.371 bits
ceiling         log2(305)      =   8.253 bits
DESTROYED                      =  10.468 bits
```

Gematria on this list destroys 10.47 of the 17.84 bits that identify a word, and 99.995% of words share their value with at least one other word.



The value 102 alone is shared by 2,353 words, beginning `abrogation`, `absurdly`, `acanthous`, `acatamathesia`, `acclimatize`, `achromatin`, `acontius`, `acrophonic`, `acroscopic`, `acroterial`, `actability`, `actionably`, `actionize`, `acupress`.

The engineering consequence is exact and it is the point of the section. **A shared gematria value is the default outcome, not a signal.** Discovering that two words share a value is discovering that a many-to-one map behaved as a many-to-one map. Under a uniform model, two arbitrary words from a 305-value codomain collide with probability on the order of 1/305, and words are not uniformly distributed, which makes it worse: the modal buckets hold over two thousand entries each. Any argument of the form *these two terms share a value, therefore they are related* is a statement about the map's compression ratio and contains no information about the terms. That is failure mode F12.2, and it is the most common analytical error in the entire literature this chapter is reading.

The deflation cuts one way only. It says the map carries no evidential weight. It does not say the map is useless: a digest that reduces 17.8 bits to 7.4 is a genuine index, and an index is a real mnemonic aid for exactly the reason §12.9 gives. Cross-referencing by digest is how you find candidate material in a corpus you cannot search. What you may not do is treat the retrieval cue as a result.

27.12.9 12.9 What these systems were actually for

The claim of this section is historical and it is well supported, but it is a historical claim and is badged as one: **these correspondence systems were mnemonic and combinatorial technologies — tools for holding a large number of relations in a mind, in cultures where written notation was scarce, expensive, restricted, or dangerous to possess.**

The load calculation makes the case quantitative. To hold `n` items as an unstructured web of pairwise relations costs `n(n−1)/2` facts. To hold them as a single ordered cycle costs `n` successor facts. To hold them as a table with `m` columns costs `n` rows.

n	unstructured pairs `n(n−1)/2`	ordered cycle `n`	ratio
7	21	7	3.00x
10	45	10	4.50x
12	66	12	5.50x
22	231	22	10.50x
78	3,003	78	38.50x

At `n = 78` the structured encoding is 38.5 times cheaper to hold than the relation set it addresses, and the saving grows linearly in `n`. That is the whole reason to build a correspondence table, and it is a reason that survives complete disbelief in every metaphysical claim ever attached to one.

Four design features of these systems are recognisable as mnemonic engineering, and each has a modern counterpart already treated in the framework:

(a) Small, fixed, memorable cardinalities. Seven planets, seven metals, seven days. Ten sephiroth. Twenty-two letters. Four elements. Twelve signs. These sit in the range a person can hold and rehearse. Dunbar's number is contested and should be stated as contested, but no one contests that seven is easier than seventy.

(b) Factored codes. Suit-and-rank: 18 labels address 56 identities. Planet-and-hour: 7 labels address 168 hour-slots per week. The log-additivity shown in §12.6 is why factoring is free in addressing power and enormous in memory cost.

(c) Redundant retrieval paths. 573 simple paths from node 1 to node 10 in the Kircher tree. Many cues, one content. This is the same principle as the method of loci (Yates, *The Art of Memory*, 1966), which the companion paper treats as the storage law with space as the retention medium [Curzi, 2026v] — high `eta`, low `delta`, because place is durable.

(d) Cross-modal binding. A correspondence table binds a name to a metal, a colour, a day, a scent, a number. Each additional column is another sensory handle on the same slot. This is cheap to add — one more row-entry per label — and it multiplies the retrieval cues.

(e) Combinatorial generativity. A table with `m` columns over `n` labels supports `n^2` ordered pairings, `C(n,2)` unordered ones, and arbitrary compositions across columns. The user does not need to have memorised the pairing; the pairing is *computed* from the table when needed. Lull's *Ars Magna* in the thirteenth century made this mechanical and explicit, with rotating concentric wheels that generate combinations of terms — a physical combinatorial engine.

None of (a) through (e) requires anything to be true about planets. They are the properties of a well-designed index in a world without indexes. **The systems are engineering artefacts of a severe memory constraint, and they are good ones; when the constraint was lifted by cheap print and then by search, the artefacts stopped being infrastructure and became objects of study, which is exactly what Star & Ruhleder (1996) predict — infrastructure becomes visible on breakdown [Star and Ruhleder, 1996, Star and Griesemer, 1989].** The converse warning applies to the modern instrument as much as to the old one: a legible scheme imposed on a domain destroys the local variety that made the domain work [Scott, 1998].

27.12.10 12.10 The join: a correspondence table is a category assignment

Now the connection this section exists to make.

The framework's decomposition [Curzi, 2026] is stated for a *block* structure: the ambient space splits into blocks `H_1, ..., H_m`, a vector's mass divides between blocks, and the surplus separates into a cross-block term and a within-block term. The validity conditions (§12.3) say one may not apply the functional until the blocks are declared. The frequency framework requires that an essence index is meaningless without a justified ordering of categories [Curzi, 2026].

Every one of those requirements is a requirement for a **table**: rows are categories, and the assignment of things to rows is the block structure.

```
correspondence table            framework object
------------------------------------------------------------------
the set of row labels           the category set, cardinality N
N                               the surplus ceiling ln N
assignment of items to rows     the block decomposition
within-row variation            u_div = alpha^2 sin^2 gamma
between-row relations           u_cross = beta^2
```

```
the row ordering (if used)      the essence index n \citep{ccp:CCP-050}
a column                        a function from categories to attributes
an unassigned item              a V1 violation: the table is not exhaustive
an item in two rows             a V2 violation: declare it or repartition
two rows naming the same thing  a V3 violation: N is inflated (see §12.11)
```

This is the substantive content of the thesis in the case of correspondence systems: the old artefact and the new formalism are doing the same job at the same point in the procedure, and the translation costs nothing on either side. It is not an argument that the old artefact is correct. A table can satisfy V1 through V4 perfectly and still assign the wrong attributes; being well-formed is not being true. It is an argument that the old artefact is the *right kind of object* — that when a fifteenth-century practitioner drew up a table of seven rows and four columns, the thing produced was a category scheme with a declared cardinality, which is precisely what a modern analyst must produce before computing `N_eff`.

The practical use is immediate. If you inherit a category scheme — a fault taxonomy, a diagnostic code list, a set of "customer segments" — you can run it through V1 to V4 exactly as you would run a historical table, and the audit will find the same faults.

27.12.11 12.11 Auditing N: the ceiling inflation fault

The most valuable thing the translation buys is a number that the tradition never computed: whether the table's `N` is real.

Take a seven-row table. Its declared ceiling is `ln 7 = 1.945910` nats, `2.807355` bits. Now suppose the audit finds that three of the seven rows are not distinct — that they name the same underlying category under different labels, which is the single most common defect in inherited taxonomies. The real cardinality is 4.

```
apparent ceiling   ln 7 = 1.945910 nats
real ceiling       ln 4 = 1.386294 nats
overstatement      0.559616 nats = 40.37% of the real ceiling
```

A scheme with three duplicate rows out of seven overstates its own surplus ceiling by 40.4%, and every downstream number — `N_eff`, `F`, headroom `h`, the storage law's growth condition — inherits that error at full size. This is the same fault as the concentration instrument's pseudo-diversity [Curzi, 2026], arriving from the taxonomy side rather than the roster side. It is also the corpus's own version of a much older correction: a diversification benefit depends on covariance, not on count, so `N` holdings that move together are one holding [Markowitz, 1952]. Two things must be kept apart here. `1/HHI` and `exp(H)` are Hill numbers of orders two and one and are fully grounded [Hill, 1973, Jost, 2006, Simpson, 1949]; the effective count used throughout this paper, `N_eff = 1 + (N−1)u`, is **not** a Hill number but a linear interpolation between 1 and `N`, is this framework's own construction, and has no counterpart in the ecological or information-theoretic literature. Small-roster estimation bias in any of them is well documented and is not corrected for here [Chao et al., 2014]. and it has the same signature: an organisation that believes it holds seven kinds of thing and behaves accordingly, while its actual variety is four.

The law of requisite variety is the reason this matters and not merely a bookkeeping irritation [Ashby, 1958, 1956]. Only variety absorbs variety. A control scheme with a genuine `N` of 4 cannot regulate a disturbance space of 7 kinds, and the table that says it has 7 rows will conceal the deficit until the fourth unhandled disturbance arrives.

Two neighbouring faults are worth naming beside it. The first is measuring on the wrong axis altogether: an entity-based exposure of 27.3 per cent against a 61.0 per cent currency share of the same balance sheet and roughly 58 per cent of settlement volume shows an inherited category scheme understating real dependency by more than a factor of two [Curzi, 2026]. **A duplicate audit fixes a count; it does not fix a scheme that is counting the wrong thing.**

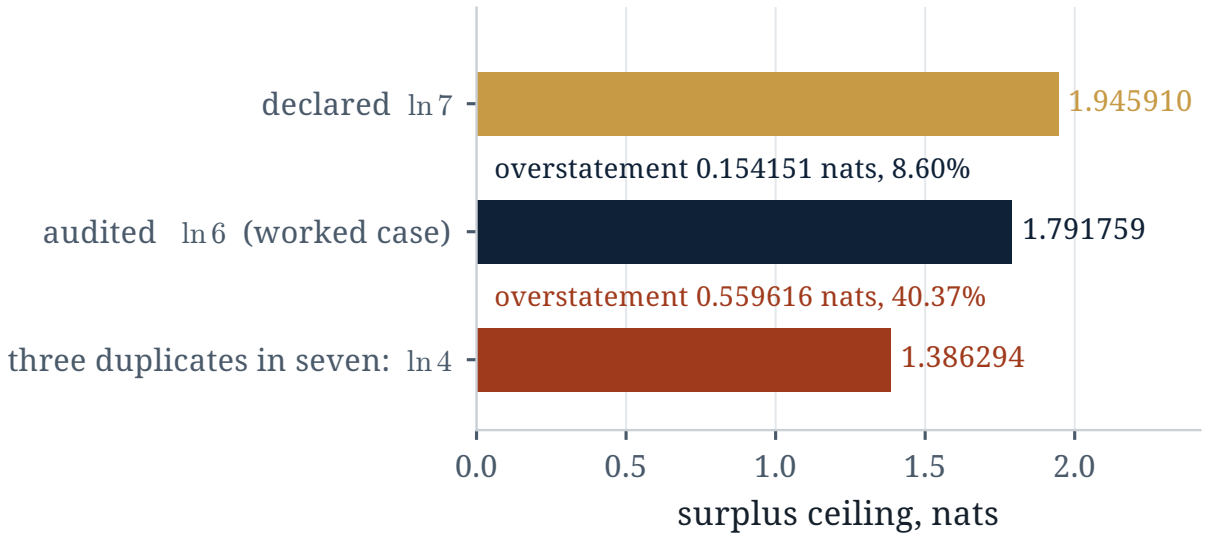


Figure 27.4: **Three duplicate rows in seven overstate the declared surplus ceiling by 0.559616 nats, 40.37 per cent of the real one.** The audited worked case of Section 15 is milder at 8.60 per cent, because two of its three suspect rows merged into one another rather than into three separate existing rows. Exact arithmetic on `ln 7`, `ln 6` and `ln 4`; produced by `figs/src/CCP-027_fig04.py`.

The second is the ceiling's other side: the roster ceiling `m(m-1) <= 2 eta ln(N) / kappa` binds the number of maintained relations rather than the number of categories, and a constructed commons sitting at 30 against a ceiling of 30.7225 is one member away from a structural problem [Curzi, 2026]. **Both ceilings are set by `ln N`, and both are inflated by the same duplicate rows.**

The diagnostic is mechanical: for every pair of rows, ask whether any observable item has ever been assigned to one and not the other, and whether any decision differs between them. If neither, the rows are duplicates. Merge them and recompute the ceiling before doing anything else.

27.13 13 Specification for building or auditing a correspondence table

Numbered, checkable requirements for building or auditing any correspondence table — historical, inherited, or new.

S12.1 — Publish the extension. The table must list its row labels exhaustively, in a fixed order, in one place. A correspondence system whose row set is "and so on" is not a category scheme and no quantity may be computed from it. Reject.

S12.2 — Declare N and derive the ceiling. State `N` and print `ln N` and `log2 N` beside it. Every downstream quantity is bounded by `ln N` [Curzi, 2026c] and the ceiling should be visible at the top of the artefact, not buried.

S12.3 — Pass V1 through V4 in print. Exhaustive, exclusive, counted, ordered-if-ordered. Each gate gets a line and a verdict. A failed gate demotes the artefact from *scheme* to *checklist* and the formalism may not be applied to it.

S12.4 — Run the duplicate audit before the first computation. For all `C(N,2)` row pairs, apply the §12.11 test. Report the number of merges performed and the corrected `N`. Print the overstatement in nats if any merge occurred.

S12.5 — Provide a residual row. Every scheme over an open domain requires an explicit `OTHER / NOT YET CATEGORISED` row, and the count in that row must be reported every period. A scheme with no residual row is not exhaustive; it is merely silent about what it drops. The rate of growth of the residual is the scheme's obsolescence signal.

S12.6 — State the arrangement whenever you state a metric. Any graph number — degree, diameter, density, cross-block fraction — is a number about one edge list. Print the edge list or cite it precisely. See H12-2.

S12.7 — Never treat a digest as a result. Any letter-sum, checksum, numerological value, or other many-to-one map may be used as a *retrieval index* and may not be used as *evidence of relation*. Where such a value is reported, report the collision rate of the map alongside it. See §12.8 and F12.2.

S12.8 — Justify the ordering separately from the membership. V3 and V4 are independent gates. A defensible set of categories with an arbitrary ordering yields arbitrary essence indices [Curzi, 2026]. If no ordering is defensible, use the scheme unordered and do not compute an index.

S12.9 — Measure the cross-block fraction; do not assume it. Count the relations that cross a category boundary and the relations that stay inside one, and report the fraction. This is the empirical handle on `beta^2`, and treating the measured fraction as `beta^2` is a modelling choice that must be badged each time it is made.

S12.10 — Prefer factored codes. Where a scheme addresses `n = p*q` identities, encode it as `p + q` labels with two columns rather than `n` labels with one. The addressing power is identical by log-additivity; the memory and maintenance cost falls from `n` to `p+q`.

S12.11 — Build redundant retrieval paths deliberately. More than one route to each entry. Count them: the number of distinct simple paths between the scheme's extreme entries is a computable robustness figure, as it was for the Kircher tree (573).

S12.12 — Re-audit on a fixed cadence. Category schemes drift: the world adds kinds and the table does not. Set a period, put the audit on it, and record `N` at each audit. A `delta` on the table itself is real, and the companion paper's periodic-forcing argument applies to maintaining it [Curzi, 2026v].

S12.13 — Keep the badge on. Where the artefact carries traditional attributions, say what is computed and what is inherited, in the same paragraph. A metric on a diagram is not an endorsement of the diagram's cosmology, and a reader must never have to guess which register a sentence is in.

27.14 14 Failure modes of correspondence tables

F12.1 — PHANTOM COLUMN. *Signature:* the table has a column whose entries cannot be independently checked against anything outside the table. Adding a row to the table is the only way anyone ever learns a new entry in that column. *Mechanism:* the column is definitional rather than empirical; it carries zero bits about the world because it is a relabelling of the row. *Fix:* mark such columns as *nomenclature* and exclude them from any information or surplus computation. They may be retained as mnemonic handles under S12.10; they may not be counted as attributes.

F12.2 — HASH MYSTICISM. *Signature:* an argument of the form "these two terms share a numerical value, therefore they are related." Escalated form: a search over many pairs until a coincidence is found, reported without the search. *Mechanism:* a many-to-one map colliding as many-to-one maps do. Measured on a real word list, the English simple-gematria map has 305 output values for 234,454 words, mean bucket 768.7, and 99.995% of words collide. The finding is a property of the map's compression ratio. *Fix:* report the collision rate with every value. Apply S12.7 without exception. If an argument survives the disclosure that its coinci-

dence had a one-in-a-few-hundred prior, it was not resting on the coincidence and should be restated without it.

F12.3 — CEILING INFLATION BY DUPLICATE ROWS. *Signature:* the declared `N` exceeds the number of rows that any decision actually distinguishes. No item has ever been assigned to row `i` rather than row `j`, and no action differs between them. *Mechanism:* `N` enters the formalism as `ln N`, the hard ceiling. Duplicates raise the nominal ceiling without raising real variety. Computed instance: 3 duplicates in 7 rows overstates the ceiling by 0.559616 nats, 40.37%. *Fix:* S12.4, before any other computation. Then re-derive headroom `h = eta ln N − delta Q − C` with the corrected `N`, and expect it to fall.

F12.4 — UNSTATED ARRANGEMENT. *Signature:* a graph metric is quoted — "twenty-two paths", "the diameter is three" — with no edge list, and the source diagram turns out to differ between editions. *Mechanism:* the metric is a function of the edge list. Different traditions draw different edges; the numbers are not comparable across them. *Fix:* H12-2 and S12.6. Print the edge list. Where two arrangements are in play, compute both and report the difference; the difference is itself informative about what the tradition treated as optional.

F12.5 — VAGUE EXTENSION. *Signature:* the scheme cannot answer, for a specific concrete item, which row it belongs to — and different competent users answer differently. *Mechanism:* V1 and V2 fail simultaneously. There is no partition, therefore no block decomposition, therefore no `u_cross` and no `u_div`. *Fix:* either sharpen the row definitions until inter-rater assignment is reproducible (this is FVP-7's cross-individual gate applied to a taxonomy), or demote the artefact in print to a heuristic list and stop computing with it. Both are acceptable outcomes; silence is not.

F12.6 — EDGE INVENTION. *Signature:* a relation is asserted between two labels because the diagram has room for it, or because symmetry would be prettier, rather than because a source or an observation carries it. *Mechanism:* the graph is being fitted to an aesthetic prior. This is over-framing, `D/P > 1`, at the level of the notation itself as an ordering rather than a quotient [Curzi, 2026t]. *Fix:* every edge carries a provenance token — source, observation, or *conjecture*. Compute all metrics twice, once with conjectural edges and once without, and print both. The gap is the scheme's exposure to its own aesthetics.

F12.7 — UNIVERSAL CORRESPONDENCE. *Signature:* the scheme claims to categorise everything. Any new object can be placed in it, and no observation could ever fall outside it. *Mechanism:* a scheme that cannot fail to apply carries no information about what it applies to. This is FVP-6's falsification gate failing at the taxonomy level, and it is the framework's test for a legitimate postulate [Curzi, 2026]. *Fix:* state, in print, a concrete item the scheme does *not* categorise, or an observation that would force a new row. If no such item can be named, the scheme is unfalsifiable and must be demoted to a language rather than a measurement instrument.

F12.8 — TREE-FLATTENING. *Signature:* a densely cyclic scheme is redrawn as a strict hierarchy "for clarity", and the overlaps are deleted. *Mechanism:* Alexander, "A City is Not a Tree" (1965). A semilattice flattened to a tree loses exactly the overlapping memberships that carried the cross-block relations. Concretely: the Kircher arrangement has 13 independent cycles; the spanning tree has 0. Flattening deletes 13 relations and the cross-block fraction collapses. *Fix:* refuse the flattening, or perform it only as a *view* over a preserved cyclic master, and label the view as lossy. Count the deleted cycles and report the count.

F12.9 — INHERITED N. *Signature:* the scheme's cardinality is a traditional number — seven, ten, twelve, twenty-two — and no one can say what would have happened if the domain had contained eight kinds. *Mechanism:* `N` was set by the notation rather than by the domain. The scheme then forces the domain into its cardinality, merging real distinctions or splitting artificial ones. *Fix:* derive `N` from the domain first, independently, then compare it to the traditional number and print both. Where they differ, the traditional number is a hypothesis about the domain and should be tested, not assumed. Where they agree, say by how much the agreement was constrained.

27.15 15 Worked example: an inherited fault taxonomy audited end to end

27.15.1 15.1 The maintenance department's fault taxonomy

A building-services department of fourteen staff maintains a research campus. It has an inherited fault taxonomy of **seven categories**, established a decade ago, used to route every incoming work order. The head of department wants to know whether the taxonomy is real, what it costs, and whether the department's stored capability is growing or decaying. The procedure is run end to end.

Step 1 — Publish the extension (S12.1). The seven rows, as inherited:

```
1  ELECTRICAL          5  FABRIC / BUILDING ENVELOPE
2  MECHANICAL / HVAC   6  PLANT FAILURE
3  PLUMBING            7  EQUIPMENT FAULT
4  CONTROLS / BMS
```

Step 2 — Declare N and the ceiling (S12.2).

```
N = 7      ln 7 = 1.945910 nats      log2 7 = 2.807355 bits
```

Step 3 — Validity gates (S12.3).

```
V1 EXHAUSTIVE   FAIL.  No residual row. 11% of work orders in the last year were
                       filed under row 7 with free-text that matched none of rows 1–6.
V2 EXCLUSIVE    FAIL.  Rows 6 and 7 overlap and rows 2 and 6 overlap.
V3 COUNTED      PENDING duplicate audit.
V4 ORDERED      N/A.   The taxonomy is unordered; no essence index is computed. Correct.
```

Step 4 — Duplicate audit (S12.4). For all `C(7,2) = 21` row pairs, ask: has any item ever been assigned to one and not the other, and does any decision differ?

```
(2, 6)  PLANT FAILURE vs MECHANICAL/HVAC   -- no routing difference; same two engineers;
                                              same parts store.                MERGE
(6, 7)  PLANT FAILURE vs EQUIPMENT FAULT   -- distinction is who owns the asset,
                                              not what the fault is.           MERGE
(5, 7)  FABRIC vs EQUIPMENT FAULT          -- distinct: different trade, different SLA.  keep
all others                                                                      keep
```

Two merges. Rows 2, 6 and 7 collapse into a single category. The corrected count:

```
corrected N = 4 + 1 = 5    ... but row 7's residual traffic must go somewhere.
```

Adding the mandatory residual row (S12.5) restores a fifth and sixth:

```
1  ELECTRICAL
2  MECHANICAL / PLANT / EQUIPMENT      (merged from old 2, 6, 7)
3  PLUMBING
4  CONTROLS / BMS
5  FABRIC / ENVELOPE
6  UNCATEGORISED  -- residual, counted and reported every period

corrected N = 6      ln 6 = 1.791759 nats
```

The overstatement carried by the inherited taxonomy, in the pure form computed in §12.11 where a seven-row scheme is really four:

```
ln 7 − ln 4 = 1.945910 − 1.386294 = 0.559616 nats = 40.37% overstatement
```

In this department's actual case, seven collapsing to six is milder:

```
ln 7 − ln 6 = 1.945910 − 1.791759 = 0.154151 nats = 8.60% overstatement
```

The department believed it had 7 kinds of fault and had 6, one of which is an admission of ignorance it had not previously been counting. The 8.60% is the honest correction and it is smaller than the textbook case because two of the three suspect rows genuinely merged into one another rather than into three separate existing rows.

Step 5 — Cross-block measurement (S12.9). Over the last 400 work orders, count how many required a person from more than one category to close:

```
single-category closures     262
cross-category closures      138
cross fraction               138/400 = 0.345
```

Badged modelling choice: take `beta^2 = 0.345`, hence `alpha^2 = 0.655`. This is the same choice made in §12.5 and it carries the same badge: the theorem does not require it.

Step 6 — Surplus. With `N = 6` and a within-category framing angle estimated at `gamma = 40 deg` (from the spread of independent diagnoses recorded on cross-category jobs):

```
u_cross = 0.345000
u_div   = 0.655 * sin^2(40 deg) = 0.655 * 0.413176 = 0.270630
u       = 0.615630
N_eff   = 1 + (6-1)*0.615630 = 4.078151
F       = ln(4.078151) = 1.405644 nats        ceiling ln 6 = 1.791759
utilisation = 1.405644 / 1.791759 = 78.45% of ceiling
```

Subadditivity check, as required:

```
f(u_cross) = ln(1 + 5*0.345000)   = ln(2.725000) = 1.002468
f(u_div)   = ln(1 + 5*0.270630)   = ln(2.353150) = 0.855756
max = 1.002468  <=  f(u) = 1.405644  <=  sum = 1.858224       OK
```

Interpretation, stated as the formalism permits: a six-category department whose staff genuinely hold non-parallel framings is running at 78.45% of the surplus its own taxonomy allows. The remaining 21.55% is bought only by raising `u` — more genuinely different framings, or more cross-category work — and it is capped at `ln 6` no matter what. To go beyond `ln 6` the department must legitimately add a category, which means finding a kind of fault it currently cannot name. That is what the residual row is for, and it is why S12.5 requires it.

Step 7 — Feed the storage law. The department's parameters, estimated from its own records:

```
eta   = 0.42      fraction of interaction that becomes durable capability
                  (documented procedures, trained second-persons, updated asset records)
delta = 0.11      annual decay: staff turnover 14%, plus procedure drift
Q_t   = 5.6       current stored capability, in the same nat-scaled units as eta*S
C_t   = 0.31      coordination overhead per period
```

Headroom, the master gauge:

```
h_t = eta * ln N - delta * Q_t - C_t
    = 0.42 * 1.791759 - 0.11 * 5.6 - 0.31
    = 0.752539 - 0.616000 - 0.310000
    = -0.173461
```

Negative. `h_t < 0` means that in this model no interaction geometry saves it [Curzi, 2026b]. The department cannot grow its stored capability by rearranging who works with whom, because even at the absolute ceiling `ln 6`, conversion times ceiling does not cover decay plus cost.

Note what the audit did to this verdict. With the *inherited* `N = 7`:

```
h_t = 0.42 * 1.945910 - 0.616000 - 0.310000 = 0.817282 - 0.926000 = -0.108718
```

Still negative — but 37% less negative, and a manager reading the uncorrected figure would have been closer to believing that a modest push on `u` could close the gap. It could not, in either case, but the corrected figure shows the true size of the shortfall. **The audit did not change the sign of the verdict; it changed the size of the intervention the verdict requires, which is the number the department actually needs.**

The three admissible responses, in the order the formalism ranks them:

```
raise eta    0.42 -> 0.52 :  h = 0.931715 - 0.926000 = +0.005715   marginal, positive
cut C        0.31 -> 0.18 :  h = 0.752539 - 0.616000 - 0.180000 = -0.043461   still negative
cut delta    0.11 -> 0.07 :  h = 0.752539 - 0.392000 - 0.310000 = +0.050539   positive
cut delta AND raise eta   :  h = 0.931715 - 0.392000 - 0.310000 = +0.229715   comfortable
```

The dominant lever is `delta`, which for this department means turnover and procedure drift — and that is the honest engineering conclusion. Reorganising the taxonomy was necessary to get an accurate number and was never going to be sufficient to fix the problem. **A correspondence table tells you what `N` is; it does not tell you what `delta` is, and `delta` was the binding constraint.**

27.15.2 15.2 What the worked example demonstrates

Every step above was performed on an ordinary building-services taxonomy, and every step is identical in form to what §§12.4 through 12.8 performed on a seventeenth-century diagram and a scheme of letter values. The gates are the same, the block decomposition is the same, the ceiling is the same `ln N`, and the storage law that consumes the result is the same equation.

That identity of procedure is the entirety of what this section claims. It is the thesis in its narrow, defensible form: **the old correspondence systems and the framework's category requirement are the same object at the same point in the same procedure, expressed in different notation.** The old notation carried it in seven rows and four columns, memorised, because there was no other way to carry it. The new notation carries it in a partition, a cardinality, and a cross-block fraction, because now there is.

The companion paper takes the next object in the same notation – the figure drawn on the page rather than the table held in the mind – and finds that it has an exact arithmetic invariant that the tradition drew for centuries without naming [Curzi, 2026v].

27.16 16 Results

Every number in this section is either exact arithmetic on quantities the paper defines, or the output of a simulation whose generating parameters are printed beside it. **No dataset was collected, and nothing here is a measurement of a person, an organisation, or a historical practice.** Where a figure is a simulation output it is rounded to the precision the replication count supports and no further.

27.16.1 16.1 The falsification suite is anti-conservative against itself

R1. Across 500 null trials at nominal `alpha = 0.05` – `Y` driven by a smooth AR(1) input that is not `S`, `S` an independent smooth AR(1) series, autocorrelation parameter 0.85 on both, `T_train = 40`, `T_holdout = 20`, 200 surrogates per trial, true `b = 0` – **the shuffled-input control rejected the true null at 0.108 for independent permutation, 0.120 for circular shift, and 0.060 for block permutation with block length 5.** Two of the three families are anti-conservative by a factor of more than two; the third is anti-conservative by 20 per cent.

The effect size that belongs beside those rates is the shortfall against nominal: **+0.058 and +0.070 in absolute rejection probability for the two failing families, and +0.010 for the passing one.** With 500 trials the standard error on a measured rate near 0.10 is `sqrt(0.10 x 0.90 / 500) = 0.0134`, so 0.108 sits 4.3 standard errors above 0.05 and 0.120 sits 5.2 above; 0.060 sits 0.7 above and is not distinguishable from nominal at this replication count. **The conclusion for block permutation is therefore "not shown to differ from nominal", not "shown to equal nominal"**, and it should not be reported as the latter.

Power was measured on the same design with true `b = 0.40`: **1.000 for independent permutation, 1.000 for circular shift, 0.998 for block permutation.** Block permutation buys its calibration for almost nothing, which is why it is the default in S4.8.

Two secondary observations. Intuition about surrogate tightness does not survive contact with a short window: the circular shift preserves the most structure and was the *worst* calibrated here, because rotating a 21-point held-out segment leaves few genuinely distinct surrogates. And a permutation p-value cannot fall below `1/(B+1)`, so `B` is a resolution parameter and not a detail: `B = 199` is the floor for a claim at the 1 per cent level and `B = 999` for anything stronger.

27.16.2 16.2 What the suite lets through

R2. In a separate 300-trial run with true `b = 0` and the reported `Y_{t+1}` inflated by `0.45 x (S_t – mean S)` – a subject who has inferred the hypothesis and whose recall of the input colours the output rating – **the full four-baseline suite passed the fabricated relationship 52.7 per cent of the time.** An example trial returned `a_hat = 0.9265` and `b_hat = 0.1856`, both strictly inside the admissible box, with a held-out model RMSE of 0.0991 against a matched-AR 0.1112 and a circular-shift surrogate mean of 0.1168, at a shuffle `p` of 0.0000.

There is no statistical remedy for this, and the paper does not offer one. The fabrication is temporal; the shuffle is a temporal test; a temporal test endorses a temporal fabrication. The remedies are design remedies – blinding the rater to the hypothesis, separating the channels carrying `Y` and `U`, prospective rather than retrospective sampling, and a yoked control – and S4.12 requires at least two of the four to be in force and named.

R2b. The matched autoregressive comparison, run alone on the null design of §4.7, **false-passed 35.8 per cent of the time.** A single baseline comparison on two smooth series is a coin weighted only slightly against the analyst.

27.16.3 16.3 Goodhart drift, measured

R3. A system was run normally for 30 periods, at which point `Y` acquired a consequence and the *reported* value inflated by 0.35 while the underlying `Q` continued to obey the same law. Pre-target fit on periods 0-30: `a = 0.8992`, `b = 0.3105`, so `delta = 0.1008` and `eta = 0.3105`. Post-target fit on periods 30-60: `a = 0.8817`, `b = 0.3551`, so `delta = 0.1183` and `eta = 0.3551`. **The estimated conversion efficiency rose by 14 per cent – exactly, by `(0.3551 – 0.3105)/0.3105 = 0.1436`, so 14.4 per cent – at the moment the true stock began to fall, from 3.0359 to 2.8447.** Nothing in the fit complained: the constraints held and both parameters stayed strictly interior.

The signature is therefore a *rising* `eta_hat` beside a *falling* independent indicator, and the detection procedure follows from it: keep a shadow proxy that is never reported upward, test

for a structural break at every date on which the measure acquired a consequence, compare the two trajectories, and prefer proxies that are expensive to fake.

27.16.4 16.4 Instrument noise moves both parameters, in a known direction

R4. At `T = 80` with true `a = 0.85` and `b = 0.40`, process noise 0.06, and 400 replications per row, rating standard deviations of 0.00, 0.05, 0.10, 0.20 and 0.40 returned mean `a_hat` of 0.8486, 0.8419, 0.8219, 0.7571 and 0.5683, mean `b_hat` of 0.4022, 0.4156, 0.4538, 0.5783 and 0.9425, and implied `delta_hat` of 0.1514, 0.1581, 0.1781, 0.2429 and 0.4317. **At a rating standard deviation of 0.40 against a `Y` whose own standard deviation is about 0.20, decay is over-estimated by a factor of 2.88 and conversion by a factor of 2.36.**

This is regression dilution applied to a lagged regressor, and it is the reason a parameter estimate without an accompanying reliability figure for the instrument that produced it is uninterpretable: the reader cannot separate the world from the ruler.

27.16.5 16.5 One window, perfect measurement, and a 16.5 per cent miss

R5. In the worked example of Section 7 – one subject, 61 daily periods, `N = 7`, `ln N = 1.945910`, generating parameters `delta = 0.15` and `eta = 0.40`, process noise 0.06, **no rating noise at all** – the constrained fit on 40 training periods returned `a_hat = 0.8833` and `b_hat = 0.3340`. At this `T` and noise the sampling standard deviations are 0.0173 on `a_hat` and 0.0341 on `b_hat`, so the estimates sit at **+2.13 and -2.15 standard errors**: a tail draw in opposite directions, which is the correlated error the two coefficients trade between them.

The shortfall on `eta` is `0.400 − 0.334 = 0.066`, which is 16.5 per cent of the true value. The source manuscript reports this as "17 per cent"; the exact figure is 16.5 per cent, it is the figure used in the abstract and throughout this paper, and the discrepancy is recorded here because a rounding that moves in the direction of a larger-sounding effect should be visible. Either way the point stands: **a single forty-period window on a single subject, with everything measured perfectly, missed the conversion parameter by more than two standard errors.**

The frozen model then beat every baseline on the held-out window – RMSE 0.0777 against persistence 0.1207 (skill 0.5860), direct input 0.0842 (0.1486, declared uninformative because the observed span of `u` was 0.5484 against a preregistered minimum of 0.60), matched autoregression 0.1248 (0.6123), and three shuffle families at 0.1256 to 0.1300 with `p = 0.0000` at `B = 2000` each. **Not one of the six thousand surrogates produced a held-out error as small as the ordered model's**, and the best of 2000 block permutations still predicted 46 per cent worse than the true ordering.

27.16.6 16.6 The same suite, on a null, refuses

Run unchanged on a series generated with no relationship to `S`, the suite fired all three rules. The constrained fit hit the retention boundary at `a_hat = 1.0000` exactly – a Rule 3 event, which S4.6 sends to the abstract. **The training-window `R^2` was 0.6477**, a figure that would be reported as a strong result in many papers, while on held-out data the frozen model was beaten by persistence (0.0730 against 0.0978), beaten decisively by matched autoregression (0.0681, skill -1.0669), and beaten by essentially every one of 2000 circular-shift surrogates at `p = 1.000`. **An in-sample `R^2` of 0.65 with a held-out failure against every baseline is not a contradiction; it is the ordinary appearance of a false result.**

27.16.7 16.7 The programme's price and its survival probability

R6. The five stages, run strictly in order, total 71 months and 8,554.8 researcher-hours across 570 participants contributing 2,791.2 participant-hours; with the declared 30 per cent overlap of Stages 2-3 and 4-5 the calendar compresses to 61 months. At 1,600 productive hours per

full-time-equivalent year that is **5.35 FTE-years per attempt**, implying continuous staffing of 0.90 FTE.

Twenty-one kill conditions are published in advance, four of them terminal for the whole enterprise. Under declared priors of 0.10, 0.10, 0.15, 0.35 and 0.20 for failure at each stage, the survival curve runs 0.9000, 0.8100, 0.6885, 0.4475 and **0.3580**, so `P(dies at or before Stage 4) = 0.5525` and the expected cost per completed programme is `5.35 / 0.3580 = 14.93 FTE-years`. **These are priors and not evidence**; a reader who disagrees substitutes their own and redoes one line of arithmetic. What is not a matter of taste is the shape: a programme whose stated completion probability is high has almost certainly written its kill conditions too loosely.

27.16.8 16.8 Ceiling inflation, and a digest that carries no evidence

R7. A seven-row category scheme declares a surplus ceiling of `ln 7 = 1.945910` nats. If three of its rows are not distinct – no item has ever been assigned to one rather than another and no decision differs between them – the real cardinality is 4 and the real ceiling is `ln 4 = 1.386294`. **The overstatement is 0.559616 nats, which is 40.37 per cent of the real ceiling, and every downstream quantity – the effective count, the surplus, headroom, the growth condition – inherits it at full size.** In the worked audit of Section 15 the same procedure collapsed seven rows to six once a residual row was added, an overstatement of 0.154151 nats or 8.60 per cent, and the corrected headroom went from `−0.108718` to `−0.173461`: **the audit did not change the sign of the verdict, it changed the size of the intervention the verdict requires by 37 per cent.**

R8. A letter-sum map from strings to integers is a hash function, and measuring it as one is deflationary. On the 234,454 purely alphabetic entries of the macOS `web2` word list under English simple gematria the map has **305 distinct values**, a mean bucket of 768.702, a largest bucket of 2,353 words at value 102, **12 singleton buckets, and 99.995 per cent of words colliding.** Input entropy is `log2(234,454) = 17.839` bits, the codomain ceiling is `log2(305) = 8.253` bits, realised output entropy is 7.371 bits, and **10.468 bits are destroyed.** A shared value is the default behaviour of a many-to-one map and contains no information about the words; the map remains a usable retrieval index and is not usable as evidence of relation.

27.17 17 Discussion

27.17.1 17.1 What follows

A control whose false-positive rate is unmeasured is an assertion. The single most transferable result here is not the number 0.108 but the procedure that produced it: simulate the study's own design with the input coefficient set to zero, measure the actual rejection rate, and use the resulting critical value in place of the nominal one. It takes minutes, it requires no additional data, and it converts a nominal p-value into a calibrated one. Every paper in this corpus reporting a shuffle p-value should be read as reporting an uncalibrated one, and at least one already says so of itself [Curzi, 2026i].

A design control and a statistical control are not substitutes. The 52.7 per cent figure is the sharpest statement of that in the paper: no quantity of surrogates repairs a fabrication whose timing is real. The remedies are blinding, channel separation, prospective sampling and a yoked control, and the reason to name which are in force is that a reader cannot otherwise price the result at all.

Measure corruption has a computable signature. Goodhart's law and Campbell's law are usually invoked as cautions [Goodhart, 1975, Campbell, 1979]; §4.12 turns them into a detection procedure with a direction – rising `eta_hat` beside a falling shadow indicator, dated to a decision – which is testable in a way a caution is not.

A research programme's honest price includes the attempts that stop early. The figure a funder needs is not 5.35 FTE-years but 14.93, because most attempts die, and the deaths are

the point: nine or ten of those years are spent learning that something does not work, which is the ordinary ratio in any field that tests things.

27.17.2 17.2 What does not follow

The simulations establish properties **of this design**. The measured rejection rates depend on the autocorrelation of both series, the lengths of the training and held-out windows, the number of surrogates, and the constraint set; none of those is universal, and **the numbers 0.108, 0.120 and 0.060 should be read as a demonstration that calibration is necessary, not as a table anyone may copy.** That is exactly why S4.10 requires each study to calibrate on its own design.

Nothing here shows the storage law is false, and nothing shows it is true. The worked example's first half is a simulation in which the law was true by construction, and the suite found it; the second half is a simulation in which it was false by construction, and the suite refused it. **Both halves test the suite, not the law.**

Nothing here identifies a social effect. §3.4 is unqualified about that: the reflexive specification is not identified from observation [Manski, 1993, Shalizi and Thomas, 2011], and the programme of Section 8 randomises only its intervention arm and only at Stage 5.

Nothing in Sections 12 to 15 asserts any supernatural claim, any physical effect of a symbol, or any mechanism. The reading rule is a translation; the graph metrics are metrics of one printed edge list and would be unchanged if the diagram had been drawn to organise plumbing fixtures; and the thesis that the old and new notations describe the same relations is a claim about *structural correspondence*, which is not evidence of mechanism.

27.17.3 17.3 Relation to the surrounding literature

The preregistration apparatus here is the standard remedy for undeclared analytic flexibility [Nosek et al., 2018, Simmons et al., 2011, Gelman and Loken, 2013], deployed against a replication base rate that is itself the reason the remedy exists [Open Science Collaboration, 2015, Ioannidis, 2005]. What this paper adds is a step further down: **the preregistered control is itself preregistered against a simulated null.**

The measurement programme is the part most exposed to existing critique, and the critique is correct. Michell [1997] argues that psychology assumes quantitative structure without testing for it; Stage 2 of Section 8 is an attempt to test for it, and it uses cross-modality matching rather than conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971], which is the weaker of the two instruments for that question. The scale-type gate that licenses division at all is Stevens [1946]'s, and the corpus's own meaningfulness theorem states the condition sharply [Curzi, 2026t].

The field-level objection is that formal social models are rarely discriminated by data. Sobkowicz [2020] finds the situation in opinion dynamics essentially unchanged over eleven years; Gallegati et al. [2006] makes the same charge against econophysics from inside it, with sub-standard statistical procedure and unjustified universality claims heading the list; Durlauf [2005] demands a **differential** prediction that simpler models do not also make. **This paper's response is not a rebuttal but an instrument**: the four baselines exist precisely to name the simpler models, and the matched autoregression is the differential test Durlauf asks for. That the matched autoregression alone false-passes 35.8 per cent of the time is the honest report of how weak that test is when run once.

The predictability ceiling is relevant to the programme's Stage 5 and is not evaded. Hofman et al. [2017] and Salganik et al. [2020] between them establish that with fifteen years of high-quality longitudinal data and hundreds of teams, life outcomes were not predicted well by anything; Lazer et al. [2014] documents the failure of a flagship prediction system by an author of the paper that defined the field. A programme proposing to predict held-out group order should expect to be measured against that record, and Stage 5's kill condition K5.1 is written to be reached.

Finally, the framework's own reflexivity is the strongest thing said against it and comes from sympathetic sources. Hacking [1995] shows that classifications of people move because the people classified respond; MacKenzie and Millo [2003] shows a theory becoming true by being adopted, with dates and price data; Ferraro et al. [2005] shows theories making people behave worse by assuming the worst of them. **A corpus proposing to give people laws to manage their own comfort by is building an engine and should say so.**

27.18 18 Verification

27.18.1 18.1 Dimensional audit

`[Q]` is the dimension of the order proxy – whatever unit it counts – and `[T]` is time, with the period as its unit.

Symbol	Dimension	Unit	Status
`Y_t`	`[Q]`	whatever the declared order proxy counts	**measured** by the declared instrument
`U_t`	`[1]`	none	**defined** dimensionless, `u` in `[0,1]`
`F(U_t)`	`[1]`	nats	**derived** dimensionless: a logarithm of an effective count
`a`	`[1]`	per period	**fitted** dimensionless; `a = 1 − delta`
`b`	`[Q]` per nat per period	proxy units per nat per period	**fitted**; identified with `eta`, *not* dimensionless
`Chat_t`	`[Q]` per period	proxy units per period	**measured**, denominated in `Y`-units by construction
`eps_t`	`[Q]` per period	proxy units per period	**residual**
`N`	`[1]`	a count	**audited**, not guessed
`ln N`	`[1]`	nats	**derived**
`h`	`[Q]` per period	proxy units per period	**derived**
`Q_inf`	`[Q]`	proxy units	**derived**: `[Q]`/period ÷ `[1]`/period returns `[Q]`
RMSE	`[Q]`	proxy units	**measured**
`p`	`[1]`	none	**derived**, a probability
`B`	`[1]`	a count of surrogates	**chosen**

Does it balance? `Y_{t+1} = a Y_t + b F(U_t) − Chat_t + eps_t` balances if and only if every right-hand term carries `[Q]`. The first does because `a` is dimensionless. The third does by the denomination rule of S4.3, which is exactly why that rule is a specification and not a preference: **a cost series left in hours or in currency makes the equation dimensionally malformed however tidy the arithmetic looks.** The second does only if `b` carries `[Q]` per nat per period, which is why `eta` must be reported as proxy-units per nat and never as a bare number. The fixed coefficient of one on `Chat` is a dimensional statement as much as an identification one.

Where the framework has no operational unit, that is said here. `Q` is stored order and has no natural unit; it is counted in whatever the declared proxy counts, and two systems' `Q` values are not comparable unless they declare the same proxy. The constructs behind `D`, `P` and `E` in Section 8 have **no operational unit at all** – establishing whether they can be given one is the entire content of Stages 1 and 2, and K1.1 and K2.4 are the conditions under which the answer is no. The underlying construct of comfort is treated as a dimensionless magnitude by stipulation, and Section 19 records that the quantitative structure required for that stipulation has not been tested.

27.18.2 18.2 Significant figures against epistemic status

Three classes of number appear in this paper and they are quoted differently on purpose.

- **Exact arithmetic.** `ln 7 = 1.945910`, `ln 6 = 1.791759`, `ln 4 = 1.386294`, `ln 7 − ln 4 = 0.559616`, `log2(305) = 8.253`, `log2(4) + log2(14) = log2(56)`, `24 mod 7 = 3`, `22/45 = 0.488889`, `71/45 = 1.577778`. Quoted to whatever precision is useful, because they are exact.
- **Simulation output.** Rejection rates from 500 trials, parameter means from 400 replications, the survival curve. Quoted to three or four figures, which is at or slightly beyond what the replication counts support; the standard error on a rate near 0.10 at 500 trials is 0.0134, so **the third figure in 0.108 is not resolved and should not be read as though it were.**
- **Declared priors.** 0.10, 0.10, 0.15, 0.35, 0.20 and everything computed from them. These are quoted to two figures and their products to four, which flatters them; `0.3580` is arithmetic on guesses and carries exactly the status of the guesses.

No quantity in this paper is a measurement of a person or an organisation, so no quantity in it carries an empirical uncertainty interval. That is a limitation and is recorded as one.

27.18.3 18.3 Postulates

Six postulates. Each is atomic, each is checkable alone, and each carries its status. Where a postulate could be dropped and the conclusion still follow, it has been dropped.

P1 — The nominal level of the test is 0.05. *Status:* **DEFINITION** – the convention declared in field F1 of the preregistration. It carries no empirical weight.

P2 — Across 500 null trials on the design of §4.7, the measured rejection rates were 0.108 (independent permutation), 0.120 (circular shift) and 0.060 (block permutation, `L = 5`). *Status:* **POSTULATED** – a simulation output of this work, with the generating process, autocorrelation parameter, window lengths and surrogate count all printed in §4.7. It is reproducible from those parameters and is supported by no external source.

P3 — A rate `r` measured over `n` independent trials has standard error `sqrt(r(1−r)/n)`. *Status:* **ARITHMETIC** – the binomial standard error, which any reader can apply to P2.

P4 — The estimated conversion efficiency before and after the Goodhart break was 0.3105 and 0.3551, and the true stock was 3.0359 and 2.8447. *Status:* **POSTULATED** – a simulation output, with the inflation of 0.35 at period 30 declared in §4.12.

P5 — The declared per-stage failure probabilities are 0.10, 0.10, 0.15, 0.35 and 0.20, and stage failures are independent. *Status:* **DEFINITION** for the numbers and **POSTULATED** for the independence; the independence assumption is the load-bearing one and is the first thing a reader should attack, because a common cause of failure across stages – an unmeasurable construct, say – would make the true survival probability lower than 0.3580 rather than higher.

P6 — A row pair is a duplicate when no observable item has ever been assigned to one rather than the other and no decision differs between them. *Status:* **DEFINITION** – the audit criterion of §12.11, stated so that it can be applied by someone who does not accept the framework.

27.18.4 18.4 The compounding chain

```
P1 + P2   -> L1  0.108 / 0.05 = 2.16                                  ARITHMETIC
P1 + P2   -> L2  0.120 / 0.05 = 2.40                                  ARITHMETIC
P1 + P2   -> L3  0.060 / 0.05 = 1.20                                  ARITHMETIC
P2 + P3   -> L4  se at r = 0.108, n = 500 is sqrt(.108*.892/500)
                 = 0.01388;  (0.108 - 0.05)/0.01388 = 4.18            ARITHMETIC
P2 + P3   -> L5  se at r = 0.060 is sqrt(.06*.94/500) = 0.01062;
                 (0.060 - 0.05)/0.01062 = 0.94                        ARITHMETIC
L1 + L4   -> R1a independent permutation is anti-conservative,
                 and the excess is resolved by the replication count  ARITHMETIC
L3 + L5   -> R1b block permutation is NOT SHOWN to differ from
                 nominal at n = 500; it is not shown to equal it      ARITHMETIC
P4        -> L6  (0.3551 - 0.3105)/0.3105 = 0.14363                   ARITHMETIC
P4        -> L7  (3.0359 - 2.8447)/3.0359 = 0.06298                   ARITHMETIC
L6 + L7   -> R3  eta_hat rose 14.4% while the true stock fell 6.3%,
                 so the sign of the estimate's motion is opposite to
                 the sign of the system's                             ARITHMETIC
P5        -> L8  0.90 x 0.90 x 0.85 x 0.65 x 0.80 = 0.358020          ARITHMETIC
P5        -> L9  1 - 0.90 x 0.90 x 0.85 x 0.65    = 0.552475          ARITHMETIC
L8        -> R6  8554.8/1600 = 5.34675 FTE-yr per attempt;
                 5.34675/0.358020 = 14.934 FTE-yr per completion      ARITHMETIC
P6        -> L10 seven rows with three duplicates has real N = 4      DEFINITION
L10       -> R7  ln 7 - ln 4 = 1.945910 - 1.386294 = 0.559616 nats;
                 0.559616/1.386294 = 0.403677                         ARITHMETIC
```

Every line is checkable from its own inputs. The three divisions a reviewer can redo in a minute are `0.108/0.05 = 2.16`, `0.90 x 0.90 x 0.85 x 0.65 x 0.80 = 0.35802`, and `(1.945910 - 1.386294)/1.386294 = 0.40368`. Those three carry the paper's three headline claims.

What the chain does not establish, stated here rather than left to a reader. R1a and R1b rest on P2, whose status is POSTULATED: they are properties of a simulation this work ran, and no external source corroborates them. R3 rests on P4, which is the same. R6 rests on P5, whose numbers are DEFINITION and whose independence assumption is POSTULATED, so **0.3580 is arithmetic on guesses and inherits their standing entirely** – it is a statement about a stated prior, not about the world. Only R7 rests on nothing but a definition and exact logarithms. **A chain in which three of four headline results depend on this work's own simulations is not a proof, and the paper says so here rather than leaving the reader to notice.**

27.18.5 18.5 Limits and asymptotics of the estimation model

The model's behaviour at the boundaries of its parameters, checked and reported.

- **`a -> 1` (`delta -> 0`).** The system never forgets; `Q_inf = (eta S - C)/delta` diverges. The update map is a contraction only while `delta > 0`, and the theorem that supplies the unique, start-independent fixed point is Banach's: a contraction with modulus `q < 1` on a complete space has exactly one fixed point, reached from any start, with the distance to it shrinking by at least a factor `q` each pass [Curzi, 2026, Banach, 1922]. **`delta = 0` is therefore not an ideal but the loss of the theorem.** This is a pathology and not an ideal, and the constrained optimiser reaching `a_hat = 1.0000` exactly is the fit reporting that it wanted to leave the box – which is what happened on the null series of §7.7, and is why S4.6 sends a boundary hit to the abstract.
- **`a -> 0` (`delta -> 1`).** Nothing carries over; the model degenerates to a static regression of `Y_{t+1}` on `F(U_t)`, and the lagged term is unidentifiable.
- **`b -> 0` (`eta -> 0`).** The surplus term is decoration and the matched autoregression is the correct model. Rule 2 exists to detect exactly this, and §4.7 measures how often it fails to.

- **`u -> 0.`** `F(u) -> ln 1 = 0`. Zero non-parallelism buys zero surplus regardless of `N`, which is the framework's binding clause and not a limiting curiosity [Curzi, 2026].
- **`u -> 1.`** `F(u) -> ln N`, the hard ceiling [Curzi, 2026c]. **No arrangement of relations raises it; only a genuine additional category does**, which is why §12.11's duplicate audit is prior to every other computation.
- **`N -> 1.`** `ln 1 = 0`: a scheme with one real category has zero available surplus rather than a low one. This is the same collapse the concentration instrument measures on the roster side, where a naive effective count of 3.41 becomes 1.00 once a coordination perimeter is respected [Curzi, 2026].
- **Span of `u -> 0`.** The direct-input baseline loses all discriminating power: over any narrow band `corr(u, F(u))` exceeds 0.995, against 0.9724 over the full range at `N = 7`. This is a constraint on the model's testable domain and S4.11 states it as one.

27.18.6 18.6 Analogy discipline

Three borrowings from physics and engineering appear in this paper and each is badged at first use.

The **leaky integrator** is a **shared mathematical form**, category 2: the same first-order linear difference equation governs an RC network, an exponential smoother and the storage law, with no claim that the electrical physics applies [Elaydi, 2005, Holt, 1957]. The **matched autoregression** is that same object used as a null, which is a derivation within the shared form and not an analogy at all. The **hash function** in §12.8 is a **derivation**, category 1: gematria satisfies the definition of a hash exactly – a deterministic map from a large domain into a small codomain – and the entropy computed from it is Shannon's quantity, computed on a real word list, with no thermodynamic content whatever [Shannon, 1948]. **Where entropy appears in this paper it is Shannon entropy in bits, never thermodynamic entropy**, and the two are different objects with different units.

No stationary-action, no uncertainty-relation and no scaling claim is made anywhere in this paper. In particular the hub topology noted in §12.4 is a statement about a ten-vertex degree sequence and is not a heavy-tail claim; the maximum-likelihood-plus-Kolmogorov-Smirnov procedure such a claim would require [Clauset et al., 2009] does not apply to ten points and was not run.

27.19 19 Limitations

Nothing in this paper was measured on a person, an organisation, or a historical practice. Every quantitative result is either exact arithmetic or a simulation with declared generating parameters. The worked subject of Section 7, the worked schedule of Section 11 and the maintenance department of Section 15 are constructed, and **none of them may be cited as evidence about any real person, laboratory or department.** Their function is to show the procedures operating end to end on numbers that behave the way real numbers behave, including where they refuse to cooperate.

The calibration result is a property of one design. The rates 0.108, 0.120 and 0.060 were measured with `T_train = 40`, `T_holdout = 20`, autocorrelation 0.85 on both series, 200 surrogates per trial and 500 trials. Changing any of those changes the answer. **The transferable finding is that the control must be calibrated, not the specific numbers**, and a reader who copies them into a different design has made the error this paper is about.

The third significant figure in the rejection rates is not resolved. At 500 trials the binomial standard error near `r = 0.10` is 0.0134. The correct statement about block permutation is that it was **not shown to differ from nominal**, and no claim of "no difference" is made from that non-significant comparison.

The reflection problem is unanswered and is the deepest limitation here. The framework's reflexive specification is not identified from observational data [Manski, 1993, Shalizi and

Thomas, 2011]; no remedy is adopted in this paper – not randomisation, not an instrument, not partial-population variation, and not a defended stipulation of the reference group [Brock and Durlauf, 2001]. Every statement in this paper about `delta` and `eta` as properties of a person or a group is therefore **a modelling position, not an identified empirical result.** A calibrated shuffle threshold does not touch this; it is a different and smaller problem.

The Lucas critique applies to this framework's own parameters, and §4.12 measures it happening. The paper asserts that a measure changes the system measured by it and then treats `delta`, `eta` and `N` as constant over an estimation window [Lucas, 1976, Hacking, 1995]. The narrow position taken – constancy only over a declared window and only while `Y` is not reported upward – is a restriction on where the protocol may be used, not a defence of the framework.

Scale type is stipulated, not demonstrated, and the decisive test was not run. `u`, `Y`, `D`, `P` and `E` are treated as ratio-scale quantities wherever differences, ratios, logarithms and means are taken. The conditions that would license that treatment are the conjoint-measurement conditions [Luce and Tukey, 1964, Krantz et al., 1971], and **they have not been tested.** Stage 2 of Section 8 proposes cross-modality matching instead, which is weaker evidence for the same conclusion. Michell [1997]'s objection – that psychology assumes quantitative structure without testing for it – is **the strongest single objection to this framework**, it applies directly, and this paper concedes rather than answers it.

`N_eff = 1 + (N−1)u` is not a Hill number. `1/HHI` and `exp(H)` are Hill numbers of orders two and one and are fully grounded [Hill, 1973, Jost, 2006, Simpson, 1949]. The effective count used here is a linear interpolation between 1 and `N`, is this framework's own construction, and has no counterpart in the ecological, economic or information-theoretic literature. Wherever an effective count appears in this paper it is that one, and it is not the grounded one.

Effective-count and diversity estimators are biased on small sets, and no correction was applied. Estimation bias in diversity measures from finite samples is well documented [Chao et al., 2014]. The worked audit of Section 15 computes `N_eff = 4.078151` from a six-row scheme, 400 work orders and an estimated within-category angle; **that number carries no interval and should not be treated as though it did.**

The surplus functional's form is underdetermined by the evidence. A sub-linear power law [Latané, 1981] and an `n log n` law [Briscoe et al., 2006] fit the same qualitative facts as `ln(1 + (N−1)u)`, and nothing in this paper discriminates between them. The direct-input baseline of §4.5.2 is the closest thing to a discriminating test and it is near-vacuous over the range real data span: over `[0.3, 0.6]` at `N = 7`, `corr(u, F(u)) = 0.99795`. **The logarithmic form is chosen, not established.** Peters and Gell-Mann [2016] is a better warrant for it than the psychophysical story, because it does not require the perceptual account to be true.

Model comparison here is incomplete by the standard this corpus sets itself. Held-out one-step-ahead RMSE against a parameter-matched baseline is reported throughout; **no information criterion was computed at any point, for any model.** The parameter counts are stated – persistence 0, direct input 2, matched AR with intercept 2, storage model 2 – so the comparison is matched, but it has not been paid for by a criterion that penalises complexity. Skill scores are reported without intervals.

The programme in Section 8 has not been run and its priors are guesses. No stage has been executed. The 0.3580 survival figure is arithmetic on five declared numbers and an independence assumption, and it should be read as a device for making the expected cost computable rather than as a forecast. **The heaviest prior, 0.35 at Stage 4, is the one most likely to be wrong in either direction**, and a reader who thinks the cross-laboratory gate is easier than that should note that the whole framework's claim to measure anything rests on it.

Stage 5 as specified is under-powered for the effect it is most likely to face. At 24 clusters of 20 members and an intracluster correlation of 0.05, the design is powered for a standardised effect of 0.50 and requires 36 clusters for 0.30. This is stated in §8.9 rather than discovered afterwards, and it means the preregistration may claim a large effect or claim pilot status, **not both.**

The correspondence-table material describes historical artefacts and does not test any his-

torical claim. That such systems functioned as mnemonic and combinatorial technologies is a **historical claim**, badged as one in §12.9, and the load arithmetic supporting it – `n(n–1)/2` against `n` – is a statement about encoding cost and not evidence about what anyone actually did. The graph metrics are metrics of one printed edge list; other historical arrangements differ and would give different numbers, which is failure mode F12.4.

There is no bibliographic support in this corpus's external database for several specific historical attributions retained in the text – the Kircher arrangement's 1652 printing, the nineteenth-century trump-to-letter bijection, Lull's combinatorial wheels, Cassius Dio's account of the planetary-hour sequence, and Yates on the art of memory. They are named in the text with their dates and are **left uncited rather than attached to a source that does not support them.**

Two claims the paper does not make. It makes **no scaling, power-law or heavy-tail claim**, so the distribution-fitting procedure such a claim would require [Clauset et al., 2009] does not apply and was not run. And it makes **no causal claim** anywhere: the shuffle establishes that a timing relationship is or is not present in a series, never that intervening on the input produces the output.

What would falsify the results of this paper. A demonstration that the measured rejection rates of §4.7 are an artefact of the particular random-number stream would overturn R1; the remedy is that the generating parameters are printed and the simulation is cheap to repeat. A demand-characteristic simulation with the same inflation coefficient producing a suite pass rate far from 52.7 per cent would overturn R2. A taxonomy in which the duplicate test of §12.11 can be applied but is rejected by practitioners who nonetheless cannot distinguish the rows would overturn H6. And if a calibrated shuffle threshold, once computed for a specific study, moved that study's reported `p` above its declared level, the study's positive result would become undecided rather than supported – **which is the practical use of everything in Section 4.**

27.20 20 Conclusion

This paper specifies a preregistered falsification suite for the storage law fitted as `Y_{t+1} = a Y_t + b F(U_t) - Chat_t + eps_t` under `0 <= a <= 1` and `b >= 0`, with four baselines, a declared freeze, boundary reporting, and held-out scoring; and it then tests that suite by simulation instead of asserting it.

The central result is negative and is about the control rather than the theory. Against a nominal 0.05, the shuffled-input control false-rejected true nulls at 0.108 for independent permutation and 0.120 for circular shift, with only block permutation near nominal at 0.060. The suite's headline control is anti-conservative as ordinarily run, and any p-value reported from it without calibration on the study's own design is optimistic by an unknown factor. Power was 1.000, 1.000 and 0.998 respectively, so block permutation buys its calibration for essentially nothing and is the default this paper specifies.

Three further measurements bound what the suite is worth. **The matched autoregressive comparison, run alone against a null, false-passed 35.8 per cent of the time. A demand-characteristic confound with a true coefficient of exactly zero passed the entire four-baseline suite 52.7 per cent of the time, and there is no statistical remedy for it – only the design remedies of blinding, channel separation, prospective sampling and a yoked control.** And **a Goodhart simulation in which the proxy acquired a consequence at period 30 raised the estimated conversion efficiency by 14 per cent, from 0.3105 to 0.3551, while the true stock fell from 3.0359 to 2.8447**, with every constraint continuing to hold.

Even with perfect measurement the arithmetic is humbling. A single forty-period window on a single subject, with no rating noise at all, recovered `a = 0.8833` against a true 0.85 and `b = 0.3340` against a true 0.40 – **+2.13 and -2.15 standard errors, a 16.5 per cent underestimate of the conversion parameter.** With rating noise the bias is systematic rather than random: at a rating standard deviation of 0.40 the estimated decay was 0.4317 against a true 0.1500 and

the estimated conversion 0.9425 against a true 0.4000.

The five-stage programme that would put these instruments in front of real people is costed at 71 months, 8,554.8 researcher-hours and 570 participants, with twenty-one kill conditions published in advance, four of them terminal. **Under declared priors its compounded probability of reaching a completed final stage is 0.3580, which implies 5.35 full-time-equivalent years per attempt and 14.93 per completed programme.** Those are priors and not evidence, and the number a funder needs is the second one, because most attempts stop early and the stopping is the point.

The audit instrument for inherited category schemes returns two quotable figures. **Three duplicate rows in a seven-row taxonomy overstate the declared surplus ceiling by 0.559616 nats, 40.37 per cent, and every downstream quantity inherits the error at full size**; on the worked case the same audit corrected a headroom of `-0.108718` to `-0.173461`, changing the size of the required intervention by 37 per cent without changing its sign. And **a letter-sum digest measured as a hash on a 234,454-word list yields 305 distinct values, 99.995 per cent of words colliding, destroying 10.468 of the 17.839 bits that identify a word** – so a shared value is the default behaviour of a many-to-one map and carries no evidence about the words.

The paper's own summary of itself is the counter-result. **The suite audits its headline control and finds it wanting: the shuffled-input test rejects true nulls at more than twice its nominal rate, is blind to filtered repackaging and to shared slow drivers, and lets a demand-characteristic confound through more than half the time – while the research programme built on it is given a 35.8 per cent chance of reaching a completed final stage.** A framework that cannot fail cannot be trusted. This one measured its own test failing, and printed the number.

27.20.1 Declarations

Data and code availability

No external dataset was collected for this paper. Every quantitative result is either exact arithmetic on quantities defined in the text or the output of a simulation whose generating parameters – process, autocorrelation, window lengths, replication counts, surrogate counts and true coefficients – are printed beside it, so that any reader can reproduce it from the text alone. The one measurement taken on an external artefact is the gematria collision analysis of §12.8, which was computed over the alphabetic entries of the macOS `/usr/share/dict/words` (`web2`) list; that list ships with the operating system and is not redistributed here. The worked subject of Section 7, the worked schedule of Section 11 and the maintenance department of Section 15 are **constructed examples and must not be cited as evidence about any real person, laboratory or department.** The figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi specified the protocol and the programme, ran the simulations, performed the analysis, produced the figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The subjects, laboratories, clusters and department described in this paper are constructed. The five-stage programme of Section 8 is a proposal; no stage of it has been executed, no participant has been recruited, and no ethics approval has been sought, because there is nothing yet to approve. The requirement carried throughout – that refusal to be measured must be available, costless, and recorded as ordinary missing data under the preregistered missing-data rule – is a constraint on prospective work and is not a report of any study performed.

Preregistration

This paper is a specification of a preregistration procedure and a simulation study of it; it does not itself report a preregistered empirical test. The falsification rules, the freeze list, the four baselines, the surrogate calibration requirement and the twenty-one kill conditions are stated in Sections 4, 5, 8 and 9 as requirements on future work. The hypotheses of Section 2 were fixed against the simulations before the simulation results were written up, and **two of the six were falsified and are reported as such in Results in the same language a success would have received.** No threshold was moved after a simulation was run. Where a rule the paper itself specifies was violated by its own worked example – the observed span of u was 0.5484 against a preregistered minimum of 0.60 – **the direct-input comparison is declared uninformative in the results table even though the model happened to beat it.**

Chapter 28

Star Polygons, Cadence, and the Ledger: gcd(N,k) as Parallelism, Division of a Fixed Budget as a Theorem, and What an Eleven-Operation Audit Does Not Prove

Abstract. We prove that the star polygon N/k decomposes into exactly gcd(N,k) closed strokes of length N/gcd(N,k), that it is unicursal precisely when gcd(N,k) = 1, that the attainable stroke counts are the proper divisors of N, and that the number of unicursal steps is Euler's totient. The result was verified by exhaustive trace across N = 3 to 40 and every k, 779 cases with no mismatch, and 89 of the 143 tabulated rows for N = 3 to 24 are unicursal. Read as a schedule, gcd is parallelism and N/gcd is latency; the throughput rate ln(x)/x is maximised at x = e, so stroke length three is optimal among integers, while splitting a twelve-station cycle into four triangles costs 55.8 per cent of the single-period surplus ceiling. On the temporal side we prove a cadence theorem: for a concave response through the origin, dividing a fixed interaction budget over more occasions never lowers the surplus delivered, and a single impulse has steady state zero against a rate's m/delta. Computed at identical nominal budget, an annual rite delivers 2.8 per cent of what a weekly practice delivers. The memory-palace arithmetic shows ordinary distinct places already reach 92.4 per cent of the ceiling and that the technique is primarily a retention intervention, retention contributing a factor 4.1667 against surplus 1.8637 and conversion 1.9000. The synthesis states its own falsification target and concedes that decomposition understates cost, not surplus.

Keywords: star polygon; greatest common divisor; unicursal; cadence theorem; periodic forcing; method of loci; retention constant; subadditivity

Series: B. Engineering the Surplus **Paper:** CCP-028

28.1 1 Introduction

This paper proves two theorems, prices two inherited technologies against them, and then audits the framework that did the pricing against a falsification target it states in advance.

The first theorem is elementary number theory and was available for two thousand years without being written down. A star polygon `{N/k}` – the figure drawn by placing `N` points on a circle and repeatedly stepping `k` positions – decomposes into exactly `gcd(N,k)` closed strokes, each of length `N/gcd(N,k)`. It is unicursal, drawable without lifting the pen, if and only if `gcd(N,k) = 1`; it requires exactly `gcd(N,k) - 1` pen lifts otherwise; the attainable stroke counts on `N` points are exactly the proper divisors of `N`; and the number of unicursal steps is Euler's totient `phi(N)`. **The hexagram is two disjoint triangles because `gcd(6,2) = 2`, and no technique of draughtsmanship changes that.** The result was verified by exhaustive trace over `N = 3` to `40` and every `k` – 779 cases, zero mismatches – before it was trusted.

Read as a schedule rather than as a figure, the invariant becomes an engineering quantity. `gcd(N,k)` is the number of independent workstreams and `N/gcd(N,k)` is the critical path, so **the speedup is exactly `gcd(N,k)` and it is perfect, because the strokes share no vertex.** That is also its price: disjoint strokes cannot interact within a period, so the per-period surplus ceiling falls from `ln N` to `ln(N/g)`. **Parallelism buys latency and pays for it in ceiling, the exchange is exact, and the ceiling rate `ln(x)/x` is maximised at `x = e`, so the best integer stroke length is three.** Splitting a twelve-station cycle into four triangles costs 55.8 per cent of the single-period ceiling.

The second theorem is about cadence and does not depend on any ceiling. For a response `g` that is concave through the origin, dividing a fixed budget `T` over `n` occasions gives a delivered total `D(n) = n g(T/n)` that is nondecreasing in `n`. Since the framework's surplus functional is strictly concave with `f(0) = 0` [Curzi, 2026], dividing a fixed interaction budget into more occasions **never delivers less**, and in the worked case delivers 4.57 times as much. What stops the division is the per-occasion fixed cost, and the optimum is computable from it alone.

The two halves meet in a set of measured comparisons that are less one-sided than the folklore. Integrated effect is **identical to every digit** – 433.3333 in all three regimes tested – because the system is linear and time-invariant, and anyone claiming that spreading a fixed input over more occasions produces more total stored order has not done the arithmetic. What differs is the peak (favouring the impulse by a factor of 3.25), the minimum, and the deliverable total once the ceiling clips each burst. **On the last of these an annual rite delivers 2.8 per cent of what a weekly practice delivers from an identical nominal budget, and the remaining 97.2 per cent is not delivered late but not delivered at all.** And the steady state settles the argument: a single impulse of any size, unrepeated, has a steady state of exactly zero, while a rate `m` has `m/delta`.

Two further quantitative results are reported. The method of loci, read as the storage law with space as the retention medium, has an advantage that factorises into three separately measurable terms, and **retention supplies more than the other two combined – 4.1667 against surplus 1.8637 and conversion 1.9000, the product 14.7546 matching the direct steady-state ratio exactly.** And ordinary, undesigned places already reach 92.4 per cent of the ceiling for a forty-item corpus, so the entire apparatus of classical rules buys the last 8 per cent: **that is the logarithm, and it is why beginners get most of the benefit at once.**

The paper's own summary of itself is a concession. Section 11 states, as a badged postulate, that if a body of inherited practice is understood as three engineering operations – lens, store, composition – then the framework has described its mechanism. Section 11.18 audits eleven such operations against three handles and finds no residue. **That is a finding about eleven operations and not a proof of completeness, the falsification target is stated as an operation that does work and maps to nothing, and the two things the audit does not reach – the felt quality of an operation, and any effect outside the practitioner – are named rather than absorbed.** The framing of a drawn figure as a *computing map* is a metaphor throughout, badged at every occurrence, and the only sentence in it that is not a metaphor is that the number of strokes is `gcd(N,k)`.

Sections 3 to 6 give the star-polygon theorem, its corollaries and tables, its specification and failure modes, and a worked scheduling example. Sections 7 to 10 give the cadence theorem, the loci and calendar arithmetic, their specification and failure modes, and a worked example. Sections 11 to 13 give the synthesis and its audit. Sections 14 to 18 report, discuss, verify, limit and conclude.

28.2 2 Hypotheses

Two of this paper's central results are **theorems**, which cannot be falsified by observation and are labelled as theorems wherever they appear; what can be wrong is the claim that any real system is the object they are about. The hypotheses below are the ones that could fail, together with what would defeat each.

H1 — The decomposition theorem holds for every star polygon. For all `N >= 3` and `1 <= k <= N−1`, tracing `{N/k}` produces exactly `gcd(N,k)` closed strokes of length `N/gcd(N,k)`, with `N` edges in total for every `k`. **Predicted:** exhaustive trace over `N = 3..40` and every `k` agrees with the formula in every case. **Falsified by:** a single traced case whose stroke count, stroke length or edge total disagrees. *(Result: **not falsified** — 779 cases, zero mismatches. This is a check on the implementation of a proof, not a test of the world; the proof is given in §3.1.)*

H2 — The attainable stroke counts are exactly the proper divisors of `N`. The set `{gcd(N,k) : 1 <= k <= N−1}` equals the set of divisors of `N` other than `N` itself. **Predicted:** exact agreement for `N = 3..24`, and `gcd(N,k) = gcd(N,N−k)` without counterexample for `N = 3..100`. **Falsified by:** an attainable stroke count that is not a proper divisor, or a mirror pair that disagrees. *(Result: **not falsified**. Consequence: **there is no five-stroke constant-step figure on twelve points, because 5 does not divide 12.**)*

H3 — Dividing a fixed interaction budget over more occasions never delivers less. For `g` concave with `g(0) = 0`, `D(n) = n g(T/n)` is nondecreasing in `n`. **Predicted:** the delivered-total column is monotone in `n` in every worked sweep. **Falsified by:** a concave-through-the-origin response for which division lowers the delivered total. *(Result: **not falsified** — monotone from 0.7278 to 3.3235 across `n = 1` to `48` in §8.8, a factor of 4.57. **The theorem does not say shorter is always better**: with a per-occasion fixed cost of 0.15 the* net* peaks at `n = 5` and falls away on both sides.)*

H4 — Spreading a fixed input over more occasions raises the integrated effect. The commonest informal argument for cadence. **Predicted, if true:** a larger sum of `Q_t` over time for a cadence than for a single impulse of the same total. **Falsified by:** equal integrals. *(Result: **FALSIFIED** — 433.3333 for the impulse, the weekly cadence and the four-weekly cadence alike, identical to every digit, which is the defining property of a linear time-invariant system. The correct argument is about the steady state, the minimum and the clipped deliverable total, and it is stronger once this one has been cleared away.)*

H5 — A rare rite is an inefficient frequent rite. The intensity/cadence axis is a single continuum on which rarer is simply worse. **Predicted, if true:** an impulse never outperforms a cadence funded from the same budget on any measure. **Falsified by:** a threshold the cadence cannot reach. *(Result: **FALSIFIED** — at a threshold of 10.0 the impulse holds the system above it for 16 periods and **both cadences score exactly zero**, because no cadence funded by that budget ever reaches 10 at all. An impulse buys height a cadence cannot buy; a cadence buys duration an impulse cannot buy.)*

H6 — Every operation in the examined practice maps to one of three parameters. Lens, store, composition. **Predicted:** each of eleven inherited operations maps to a named parameter with a direction, and four of the eleven can be priced exactly. **Falsified by:** an operation that does work, is agreed to do work, and maps to nothing. *(Result: **not falsified over eleven operations**. This is an audit of a list and **not a proof of completeness**; §11.19 states the refutation condition and names three candidate residues, of which one – the integrated practice's lower per-occasion cost – is a real gap in the audit's coverage and is conceded there.)*

Three claims in this paper are not hypotheses and are labelled as what they are. That the method of loci raises `u` and `eta` and lowers `delta` is a **reading**, not a theorem; what is a theorem is that the steady state moves by the product of the three factors once they are granted. That a ritual calendar is described by the storage law at all is a **modelling postulate**. And the thesis of §11.1 – that a stipulatively defined practice has had its mechanism described here – is a **postulate with a stipulative definition inside it**, and a reader who rejects the definition is declining the thesis rather than contradicting it.

The reflection problem applies to everything in this paper that treats a group parameter as a property of the group. The framework's reflexive specification, in which members constitute the aggregate and the aggregate influences the members, is a linear-in-means form whose endogenous effect is **not identified from observational data** [Manski, 1993], and influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011]. **No remedy is adopted here** – not randomisation, not an instrument, not partial-population variation, and not a defended stipulation of the reference group [Brock and Durlauf, 2001]. Every statement below about `u`, `eta` or `delta` as a property of a circle, a studio or a department is therefore a **modelling position rather than an identified empirical result**, and Section 17 does not soften that.

28.3 3 The star-polygon decomposition theorem

This section contains the strongest formal result in the paper, and it is elementary number theory. It is also where the framing of a drawn figure as a *computing map* is stated, and that framing is a **METAPHOR**. The badge is kept on throughout, and every quantitative statement below is separated from it by a line the reader can check.

The situation is worth naming at the outset. Most of the corpus's reading of inherited technique is translation: an old artefact turns out to be an ordinary mathematical object once one looks at what it literally contains [Curzi, 2026u]. This section is different. The star polygon has an **exact arithmetic invariant** that the tradition drew for two thousand years and never named, and once the invariant is stated it settles, in one line, a question the tradition handled by trial and error: which figures can be drawn without lifting the pen. The answer is a greatest common divisor. It is true, it is checkable in a minute, and it was always there in the figures.

A star polygon {N/k}, drawn by stepping k positions around N equally spaced points, decomposes into exactly gcd(N,k) separate closed paths, each of length N/gcd(N,k).

The hexagram is two triangles because `gcd(6,2) = 2`. The heptagram is one continuous line because `gcd(7,2) = 1`. The twelve-point figure of step three is three squares because `gcd(12,3) = 3`. None of this is interpretation; all of it is arithmetic, and §4.1 prints the complete table for `N = 3` through `N = 24` with every valid `k`.

Nothing in this paper asserts any supernatural claim, any physical effect of a symbol, or any result not derivable from the formalism. That standing constraint governs throughout without amendment [Curzi, 2026u].

28.3.1 3.1 The decomposition theorem

THEOREM 1 (Decomposition of the star polygon). Let `N >= 3` and `1 <= k <= N−1`. Place `N` labelled points at the vertices of a regular `N`-gon, identified with the residues `Z_N = {0, 1, ..., N−1}`. Draw the figure `{N/k}` by starting at some point and repeatedly joining the current point `x` to the point `x + k (mod N)`, continuing until the start is reached, then beginning again at any unused point, until every point has been used. Let `g = gcd(N, k)`. Then:

```
(i)   the figure consists of exactly  g  closed paths (strokes)
(ii)  each stroke visits exactly      N/g  distinct points
(iii) each stroke has exactly         N/g  edges
(iv)  the whole figure has exactly    N    edges, for every k
(v)   the figure is a single closed path (UNICURSAL) if and only if  g = 1
```

Proof. The step map is `s(x) = x + k (mod N)`. Its orbits are the cosets in `Z_N` of the cyclic subgroup generated by `k`. In the additive group `Z_N`, the subgroup generated by `k` is

```
<k> = { 0, k, 2k, ... }  =  <gcd(N,k)>  =  { 0, g, 2g, ..., N−g }
```

— the standard fact that the set of integer combinations of `N` and `k` is the set of multiples of their greatest common divisor. That subgroup has order `N/g`. By Lagrange's theorem the number of cosets is `|Z_N| / |<k>| = N / (N/g) = g`. Each coset is one orbit of `s`, hence one closed path, and each has `N/g` elements, giving (i) and (ii). Every point has exactly one outgoing and one incoming edge, so each stroke of `N/g` points carries `N/g` edges, giving (iii); summing over `g` strokes gives `g * (N/g) = N` edges, giving (iv). And (v) is (i) with `g = 1`. **QED**

Equivalently, and this is the formulation to hold in mind: the figure is the Cayley graph of `Z_N` with connection set `{+k, −k}`, and its connected components are the cosets of `<k>`. Every vertex has degree 2, so the graph is 2-regular, so every component is a cycle. **A 2-regular graph is connected if and only if it is a single cycle, which is why unicursality, Eulerianity and Hamiltonicity coincide for star polygons and all three reduce to `gcd(N,k) = 1`.**

Verified empirically before being trusted: the figure was traced by explicit walk for every `N = 3..40` and every `k = 1..N−1`, a total of 779 cases, and the stroke count, stroke length and edge total agreed with the theorem in **every case, zero mismatches.**

The Schläfli notation `{N/k}` and the classification of `gcd > 1` cases as *star figures* or *compounds* rather than star polygons proper is standard; Coxeter's *Regular Polytopes* is the usual reference.

28.3.2 3.2 The corollaries that do the work

COROLLARY 1 (Unicursality). `{N/k}` can be drawn without lifting the pen if and only if `gcd(N,k) = 1`. The number of such `k` in `1..N−1` is `phi(N)`, Euler's totient. Since `{N/k}` and `{N/(N−k)}` are the same figure traversed in opposite directions, the number of *distinct* unicursal figures on `N` points with constant step is `phi(N)/2` for `N > 2`.

COROLLARY 2 (Pen lifts). Drawing the complete figure requires exactly `gcd(N,k) − 1` pen lifts. Unicursal figures require zero.

COROLLARY 3 (Compound reduction). For `g = gcd(N,k) > 1`,

```
{N/k}  =  g  disjoint copies of  {(N/g)/(k/g)}
```

and the inner figure is always unicursal, because `gcd(N/g, k/g) = 1` by construction. Verified:

```
{6/2}  = 2 x {3/1}    two triangles
{6/3}  = 3 x {2/1}    three digons (degenerate: the three long diagonals)
```

```
{8/2}  = 2 x {4/1}    two squares
{9/3}  = 3 x {3/1}    three triangles
{10/4} = 2 x {5/2}    two pentagrams
{12/3} = 3 x {4/1}    three squares
{12/4} = 4 x {3/1}    four triangles
{12/6} = 6 x {2/1}    six digons
{16/4} = 4 x {4/1}    four squares
{18/6} = 6 x {3/1}    six triangles
{20/8} = 4 x {5/2}    four pentagrams
{24/9} = 3 x {8/3}    three octagrams
{24/10}= 2 x {12/5}   two dodecagrams
```

COROLLARY 4 (Attainable stroke counts). The set of stroke counts achievable on `N` points is exactly the set of proper divisors of `N`. Verified for `N = 3..24`: for every `N`, `{ gcd(N,k) : 1 <= k <= N-1 }` equals the divisors of `N` other than `N` itself. **You cannot draw a five-stroke figure on twelve points with a constant step, because 5 does not divide 12, and no amount of ingenuity changes that.**

COROLLARY 5 (Mirror symmetry). `gcd(N,k) = gcd(N, N-k)` for all `k`. Verified with zero counterexamples for `N = 3..100`. This halves the table: `k` beyond `N/2` retraces the figures of `k` below it, in reverse. It is why §4.1 tabulates `k = 1..floor(N/2)` and nothing is lost.

COROLLARY 6 (Vertex angle). The point angle of `{N/k}` is `180*(N - 2k)/N` degrees, and the density (winding number) of each stroke is `k/g`, so each stroke's total turning is `360*(k/g)` degrees. Verified: `{5/2}` gives `180*1/5 = 36` degrees — the pentagram's point angle, which is the value every draughtsman knows.

28.3.3 3.3 Status table

Element	Status	Warrant
Strokes = `gcd(N,k)`, each of length `N/gcd(N,k)`	**PROVED** (group theory)	§3.1
Verified by exhaustive trace, `N = 3..40`, all `k`, 779 cases, 0 mismatches	**MEASURED**	§3.1
Unicursal iff `gcd(N,k) = 1`; count of unicursal `k` is `phi(N)`	**PROVED**	§3.2
Attainable stroke counts = proper divisors of `N`	**PROVED** (verified `N`=3..24)	§3.2
Pen lifts = `gcd(N,k) - 1`	**PROVED**	§3.2
Point angle `180(N-2k)/N` degrees	**PROVED** (geometry)	§3.2
The Spare reduction is a lossy, non-invertible encoding	**PROVED** (definition + count)	§4.8
Its output is bounded by 26 bits for any input length	**PROVED** (`2^26` subsets)	§4.8
Measured letter-set collision rate on 218,002 words: 71.4%	**MEASURED**	§4.8
`gcd(N, .)` is a non-monotone, discontinuous function of `k`	**PROVED** (computed jumps)	§4.6
"The sigil is a nonlinear computing map"	**METAPHOR** (author's, badged)	§3.4
That drawing a figure produces any external effect	**REFUSED**	standing refusal

28.3.4 3.4 The metaphor, and exactly where it stops

The framing under examination is that such a figure can be *read* as a nonlinear computing map. Here is the honest anatomy of that sentence.

What is literally true. The figure has a computational structure in the ordinary sense: a set of nodes, a deterministic successor rule, orbits, a component count, a path length per component, and a total edge count. `gcd(N,k)` is precisely the number of *independent* closed processes the figure contains and `N/gcd(N,k)` is precisely the length of each. If you were scheduling work around a ring of `N` stations by handing it on `k` places at a time, those numbers would be the number of parallel workstreams and the length of each. **The number of separate strokes in the figure is `gcd(N,k)`, and that is the figure's parallelism** — that sentence is not a metaphor. It is a statement about a graph, and it is provable.

What is metaphor. Calling the figure a *computing map* imports the suggestion that something is being computed, by something, for some purpose. Nothing in the mathematics supplies an input, an output, or an agent. A star polygon is a static graph. It has the structural signature of a parallel process; it does not run one. **The metaphor is a way of seeing, it earns its keep by making the invariant visible, and it is not an assertion about the figure's nature.**

In what precise sense "nonlinear". The step map `x -> x + k (mod N)` is affine, hence linear in the technical sense — this must be conceded immediately, because it is the obvious objection. The nonlinearity is not in the map; it is in the **dependence of the figure's structure on the parameter `k`.** The function `k -> gcd(N,k)` is not monotone, not continuous, not smooth, and not approximable by any local rule: it is a number-theoretic function that jumps. On 24 points, `k=11` gives one stroke and `k=12` gives twelve. A unit change in the parameter changes the structure by a factor of twelve. §4.6 tabulates the maximum jump for every `N` in range. **That is the exact, defensible content of "nonlinear" here: the structure is a non-smooth arithmetic function of the design parameter, so you cannot tune it by small steps and you cannot interpolate.** Everything else in the phrase is metaphor and is badged as such.

What is refused, permanently. That drawing the figure does anything outside the person drawing it. Not asserted, not implied, not entertained.

28.4 4 Star polygons as schedules: parallelism, ceiling, and rate

28.4.1 4.1 The complete table, N = 3 to 24

Computed by explicit trace and cross-checked against `gcd`, with `k` running from 1 to `floor(N/2)`; by Corollary 5 the remaining `k` mirror these exactly. "Strokes" is the number of separate closed paths, "stroke length" the number of points (and edges) in each, and "compound form" the reduction of Corollary 3.

N	k	gcd(N,k)	strokes	stroke length N/g	unicursal	compound form
3	1	1	1	3	YES	–
4	1	1	1	4	YES	–
4	2	2	2	2	no	2 x {2/1}
5	1	1	1	5	YES	–
5	2	1	1	5	YES	–
6	1	1	1	6	YES	–
6	2	2	2	3	no	2 x {3/1}
6	3	3	3	2	no	3 x {2/1}

7	1	1	1	7	YES	–
7	2	1	1	7	YES	–
7	3	1	1	7	YES	–
8	1	1	1	8	YES	–
8	2	2	2	4	no	2 x {4/1}
8	3	1	1	8	YES	–
8	4	4	4	2	no	4 x {2/1}
9	1	1	1	9	YES	–
9	2	1	1	9	YES	–
9	3	3	3	3	no	3 x {3/1}
9	4	1	1	9	YES	–
10	1	1	1	10	YES	–
10	2	2	2	5	no	2 x {5/1}
10	3	1	1	10	YES	–
10	4	2	2	5	no	2 x {5/2}
10	5	5	5	2	no	5 x {2/1}
11	1	1	1	11	YES	–
11	2	1	1	11	YES	–
11	3	1	1	11	YES	–
11	4	1	1	11	YES	–
11	5	1	1	11	YES	–
12	1	1	1	12	YES	–
12	2	2	2	6	no	2 x {6/1}
12	3	3	3	4	no	3 x {4/1}
12	4	4	4	3	no	4 x {3/1}
12	5	1	1	12	YES	–
12	6	6	6	2	no	6 x {2/1}
13	1	1	1	13	YES	–
13	2	1	1	13	YES	–
13	3	1	1	13	YES	–
13	4	1	1	13	YES	–
13	5	1	1	13	YES	–
13	6	1	1	13	YES	–
14	1	1	1	14	YES	–
14	2	2	2	7	no	2 x {7/1}
14	3	1	1	14	YES	–
14	4	2	2	7	no	2 x {7/2}
14	5	1	1	14	YES	–
14	6	2	2	7	no	2 x {7/3}
14	7	7	7	2	no	7 x {2/1}
15	1	1	1	15	YES	–
15	2	1	1	15	YES	–
15	3	3	3	5	no	3 x {5/1}
15	4	1	1	15	YES	–
15	5	5	5	3	no	5 x {3/1}
15	6	3	3	5	no	3 x {5/2}

15	7	1	1	15	YES	–
16	1	1	1	16	YES	–
16	2	2	2	8	no	2 x {8/1}
16	3	1	1	16	YES	–
16	4	4	4	4	no	4 x {4/1}
16	5	1	1	16	YES	–
16	6	2	2	8	no	2 x {8/3}
16	7	1	1	16	YES	–
16	8	8	8	2	no	8 x {2/1}
17	1	1	1	17	YES	–
17	2	1	1	17	YES	–
17	3	1	1	17	YES	–
17	4	1	1	17	YES	–
17	5	1	1	17	YES	–
17	6	1	1	17	YES	–
17	7	1	1	17	YES	–
17	8	1	1	17	YES	–
18	1	1	1	18	YES	–
18	2	2	2	9	no	2 x {9/1}
18	3	3	3	6	no	3 x {6/1}
18	4	2	2	9	no	2 x {9/2}
18	5	1	1	18	YES	–
18	6	6	6	3	no	6 x {3/1}
18	7	1	1	18	YES	–
18	8	2	2	9	no	2 x {9/4}
18	9	9	9	2	no	9 x {2/1}
19	1	1	1	19	YES	–
19	2	1	1	19	YES	–
19	3	1	1	19	YES	–
19	4	1	1	19	YES	–
19	5	1	1	19	YES	–
19	6	1	1	19	YES	–
19	7	1	1	19	YES	–
19	8	1	1	19	YES	–
19	9	1	1	19	YES	–
20	1	1	1	20	YES	–
20	2	2	2	10	no	2 x {10/1}
20	3	1	1	20	YES	–
20	4	4	4	5	no	4 x {5/1}
20	5	5	5	4	no	5 x {4/1}
20	6	2	2	10	no	2 x {10/3}
20	7	1	1	20	YES	–
20	8	4	4	5	no	4 x {5/2}
20	9	1	1	20	YES	–
20	10	10	10	2	no	10 x {2/1}
21	1	1	1	21	YES	–
21	2	1	1	21	YES	–
21	3	3	3	7	no	3 x {7/1}
21	4	1	1	21	YES	–

```
 21    5         1         1                  21   YES         -
 21    6         3         3                   7   no          3 x {7/2}
 21    7         7         7                   3   no          7 x {3/1}
 21    8         1         1                  21   YES         -
 21    9         3         3                   7   no          3 x {7/3}
 21   10         1         1                  21   YES         -
---   ---  --------   -------   ----------------   ---------   -------------
 22    1         1         1                  22   YES         -
 22    2         2         2                  11   no          2 x {11/1}
 22    3         1         1                  22   YES         -
 22    4         2         2                  11   no          2 x {11/2}
 22    5         1         1                  22   YES         -
 22    6         2         2                  11   no          2 x {11/3}
 22    7         1         1                  22   YES         -
 22    8         2         2                  11   no          2 x {11/4}
 22    9         1         1                  22   YES         -
 22   10         2         2                  11   no          2 x {11/5}
 22   11        11        11                   2   no          11 x {2/1}
---   ---  --------   -------   ----------------   ---------   -------------
 23    1         1         1                  23   YES         -
 23    2         1         1                  23   YES         -
 23    3         1         1                  23   YES         -
 23    4         1         1                  23   YES         -
 23    5         1         1                  23   YES         -
 23    6         1         1                  23   YES         -
 23    7         1         1                  23   YES         -
 23    8         1         1                  23   YES         -
 23    9         1         1                  23   YES         -
 23   10         1         1                  23   YES         -
 23   11         1         1                  23   YES         -
---   ---  --------   -------   ----------------   ---------   -------------
 24    1         1         1                  24   YES         -
 24    2         2         2                  12   no          2 x {12/1}
 24    3         3         3                   8   no          3 x {8/1}
 24    4         4         4                   6   no          4 x {6/1}
 24    5         1         1                  24   YES         -
 24    6         6         6                   4   no          6 x {4/1}
 24    7         1         1                  24   YES         -
 24    8         8         8                   3   no          8 x {3/1}
 24    9         3         3                   8   no          3 x {8/3}
 24   10         2         2                  12   no          2 x {12/5}
 24   11         1         1                  24   YES         -
 24   12        12        12                   2   no          12 x {2/1}
---   ---  --------   -------   ----------------   ---------   -------------
```

143 rows. 89 of them are unicursal — 62.2%. Every row was produced by walking the figure, not by evaluating `gcd`, and then checked against `gcd`.

28.4.2 4.2 The named verifications, with orbits printed

The theorem is worth distrusting until you have seen the walks. Here they are, as vertex sequences on `Z_N` starting from 0.

```
{6/2}   gcd 2   TWO strokes of length 3:    [0,2,4]  and  [1,3,5]
{6/3}   gcd 3   THREE strokes of length 2:  [0,3]  [1,4]  [2,5]
```

```
{7/2}   gcd 1   ONE stroke of length 7:     [0,2,4,6,1,3,5]
{7/3}   gcd 1   ONE stroke of length 7:     [0,3,6,2,5,1,4]
{5/2}   gcd 1   ONE stroke of length 5:     [0,2,4,1,3]
{8/3}   gcd 1   ONE stroke of length 8:     [0,3,6,1,4,7,2,5]
{9/3}   gcd 3   THREE strokes of length 3:  [0,3,6]  [1,4,7]  [2,5,8]
{10/3}  gcd 1   ONE stroke of length 10:    [0,3,6,9,2,5,8,1,4,7]
{12/3}  gcd 3   THREE strokes of length 4:  [0,3,6,9]  [1,4,7,10]  [2,5,8,11]
{12/5}  gcd 1   ONE stroke of length 12:    [0,5,10,3,8,1,6,11,4,9,2,7]
{24/7}  gcd 1   ONE stroke of length 24:    [0,7,14,21,4,11,18,1,8,15,22,5,12,19,2,9,16,23,
                                             6,13,20,3,10,17]
```

The hexagram is literally two separate triangles: `[0,2,4]` and `[1,3,5]`, disjoint, sharing no vertex and no edge. The two triangles overlap visually on the page and are combinatorially independent. That is not a reading of the figure; it is what happens when you draw it. Every draughtsman who has ever drawn a Star of David knows they had to lift the pen once, and Corollary 2 says exactly why: `gcd(6,2) − 1 = 1` lift, always, no technique available.

The pentagram, by contrast, is `gcd(5,2) = 1` and needs zero lifts, which is why it is the figure that gets drawn in one motion in every schoolyard on earth.

28.4.3 4.3 The unicursal census

Per `N`, the number of unicursal steps and how the divisor structure of `N` controls it.

N	phi(N)	distinct proper star polygons	unicursal k in 1..floor(N/2)	total k	fraction	divisors of N
3	2	0	1	1	1.000	1, 3
4	2	0	1	2	0.500	1, 2, 4
5	4	1	2	2	1.000	1, 5
6	2	0	1	3	0.333	1, 2, 3, 6
7	6	2	3	3	1.000	1, 7
8	4	1	2	4	0.500	1, 2, 4, 8
9	6	2	3	4	0.750	1, 3, 9
10	4	1	2	5	0.400	1, 2, 5, 10
11	10	4	5	5	1.000	1, 11
12	4	1	2	6	0.333	1, 2, 3, 4, 6, 12
13	12	5	6	6	1.000	1, 13
14	6	2	3	7	0.429	1, 2, 7, 14
15	8	3	4	7	0.571	1, 3, 5, 15
16	8	3	4	8	0.500	1, 2, 4, 8, 16
17	16	7	8	8	1.000	1, 17
18	6	2	3	9	0.333	1, 2, 3, 6, 9, 18
19	18	8	9	9	1.000	1, 19
20	8	3	4	10	0.400	1, 2, 4, 5, 10, 20
21	12	5	6	10	0.600	1, 3, 7, 21
22	10	4	5	11	0.455	1, 2, 11, 22
23	22	10	11	11	1.000	1, 23
24	8	3	4	12	0.333	1, 2, 3, 4, 6, 8, 12, 24

Prime `N` gives a unicursal fraction of exactly 1.000: every step draws one continuous line. Highly composite `N` gives the worst ratios — `N = 12`, `18` and `24` all sit at 0.333. A designer who wants many available single-stroke figures should choose a prime number of points; a designer who wants many available *parallel decompositions* should choose a highly composite one. The two goals are exactly opposed, and the opposition is the divisor structure of `N`.

The column "distinct proper star polygons" counts `k` with `2 <= k < N/2` and `gcd(N,k) = 1` — the figures that are genuinely star-shaped, single-stroke, and not the convex polygon. For `N = 3, 4, 6` it is zero: **there is no star polygon on three, four or six points, which is a fact about divisors and not about draughtsmanship.**

28.4.4 4.4 The parallelism reading

This section is where the invariant becomes an engineering quantity. Take the figure as a *schedule*: `N` stations around a ring, work handed on `k` places at a time, one hand-off per unit time.

```
P  = gcd(N,k)      number of independent workstreams (strokes)
W  = N             total work: one hand-off per station, always
T  = N / gcd(N,k)  critical path: the length of the longest stroke
speedup = W / T = gcd(N,k) = P
```

N	k	P = strokes	W = edges	T = critical path	speedup
6	2	2	6	3	2.0
12	3	3	12	4	3.0
12	4	4	12	3	4.0
24	8	8	24	3	8.0
20	5	5	20	4	5.0
7	2	1	7	7	1.0
13	5	1	13	13	1.0

Total work is `N` for every `k`; only the arrangement changes. That is the honest form of the parallelism claim: the figure never does less work, it does the same work in fewer sequential steps. The speedup is exactly `gcd(N,k)` and it is perfect — no overhead appears in the geometry, because the strokes share no vertex.

Now the constraint that makes this interesting rather than free. **The `g` strokes are *disjoint*. No point on one stroke is ever adjacent to a point on another.** In the vocabulary of the framework's block decomposition, partitioning the ring into `g` strokes means that within one period no pair from different strokes interacts at all. The cross-block term `u_cross` is unavailable inside the period, and the surplus ceiling available to a stroke of `x = N/g` members is `ln x`, not `ln N` [Curzi, 2026c]. **A second opinion from a source that shares training, incentives and premises is worth `ln 1 = 0`** [Curzi, 2026], and disjoint strokes are the scheduling form of exactly that constraint.

```
parallelism g:      latency falls from N to N/g
surplus ceiling:    falls from ln N to ln (N/g)
```

Parallelism buys latency and pays for it in ceiling. That trade is exact and it is the whole design decision.

28.4.5 4.5 The rate theorem: why stroke length three is optimal

Define the *ceiling rate* of a decomposition as surplus ceiling per unit of latency:

```
R(x) = ln(x) / x          where x = N/g is the stroke length
```

This is an upper bound on the rate at which a stroke can generate surplus, not an achieved rate; it is attained only if `u` can be held at 1 within each stroke. Badge it as a ceiling.

```
dR/dx = (1 − ln x) / x^2  =  0   at   ln x = 1,  i.e.  x = e = 2.718282
```

`R` rises on `(0, e)` and falls on `(e, infinity)`. The optimum is at `x = e`, and since stroke length must be an integer dividing `N`, the best available integer is 3 whenever 3 divides `N`. Computed:

```
x = 2   R = 0.346574
x = 3   R = 0.366204     <-- maximum over integers
x = 4   R = 0.346574
x = 5   R = 0.321888
x = 6   R = 0.298627
x = 12  R = 0.207076
```

Note the exact tie between `x = 2` and `x = 4`: `ln 2 / 2 = ln 4 / 4` because `ln 4 = 2 ln 2`. The tie is arithmetic, not coincidence.

For `N = 12`, sweeping the step:

k	g = gcd(12,k)	strokes	stroke length x	latency	ceiling ln x	rate ln(x)/x
1	1	1	12	12	2.484907	0.207076
2	2	2	6	6	1.791759	0.298627
3	3	3	4	4	1.386294	0.346574
4	4	4	3	3	1.098612	0.366204
5	1	1	12	12	2.484907	0.207076
6	6	6	2	2	0.693147	0.346574

The rate-optimal decomposition of twelve points is `{12/4}` — four triangles — and it is optimal because `ln x / x` peaks at `x = e` and 3 is the nearest integer divisor. The triad recurs across organisational traditions – from nested governance units [Ostrom, 1990, 2010] to the three-person review cell – and here it appears as the integer nearest to `e` under a log-over-linear objective. That is a derivation, offered as one, with no suggestion that anyone in history derived it.

The counterweight is Corollary 4 and the previous section: `{12/1}` and `{12/5}` are the only configurations that keep the full `ln 12 = 2.484907` ceiling, and they take twelve steps to get there. **Fast and shallow, or slow and deep. The figure makes you choose, and the divisors of `N` are the menu.**

28.4.6 4.6 In what sense the map is nonlinear

The concession first: the step map is affine. The nonlinearity claimed here lives in the parameter dependence, and here it is measured. For each `N`, the maximum change in stroke count caused by a single unit change in `k`:

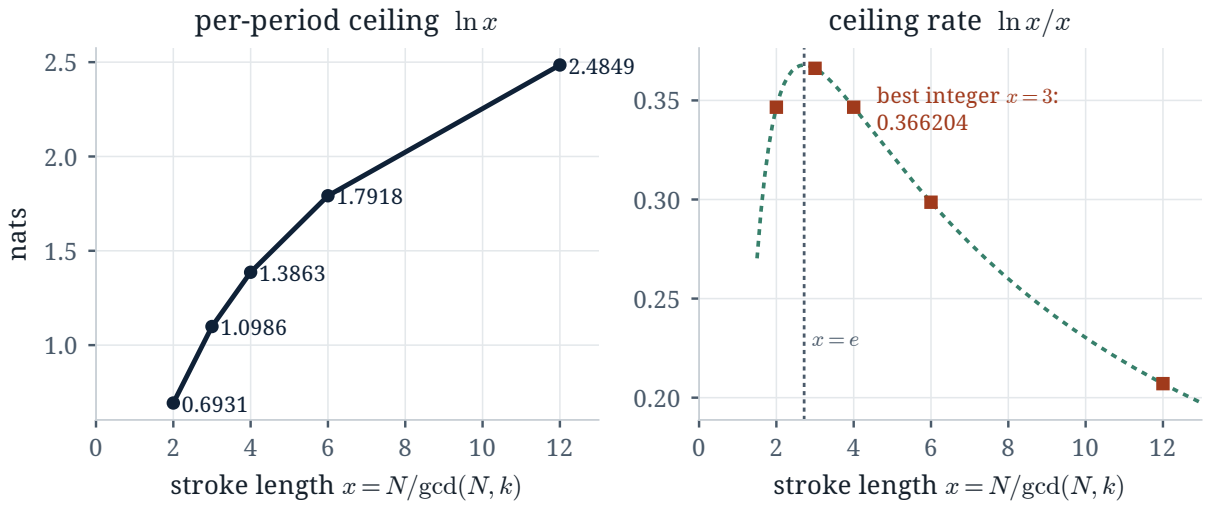


Figure 28.1: **Parallelism buys latency and pays for it in ceiling, and the exchange rate peaks at a stroke length of three.** The per-period ceiling `ln x` rises with stroke length while the ceiling *rate* `ln x / x` peaks at `x = e = 2.718282`, so the best available integer is 3, at 0.366204, with an exact tie between `x = 2` and `x = 4` because `ln 4 = 2 ln 2`. Exact arithmetic on the divisors of 12; produced by `figs/src/CCP-028_fig02.py`.

N	max abs jump	at k -> k+1	strokes go
3	0	1 -> 2	1 -> 1
4	1	2 -> 3	2 -> 1
6	1	4 -> 5	2 -> 1
8	3	4 -> 5	4 -> 1
9	2	6 -> 7	3 -> 1
10	3	5 -> 6	5 -> 2
12	5	6 -> 7	6 -> 1
14	5	7 -> 8	7 -> 2
15	4	10 -> 11	5 -> 1
16	7	8 -> 9	8 -> 1
18	7	9 -> 10	9 -> 2
20	9	10 -> 11	10 -> 1
21	6	13 -> 14	1 -> 7
22	9	11 -> 12	11 -> 2
24	11	12 -> 13	12 -> 1
5, 7, 11, 13, 17, 19, 23	0	—	always 1

On 24 points, moving the step from 12 to 13 takes the figure from twelve independent strokes to one, in a single unit of the parameter. No local model predicts that. A designer who tunes `k` by small increments and watches the result gets no gradient to follow, because there is no gradient: `gcd(N, .)` is constant-1 across every `k` coprime to `N` and spikes to a divisor exactly on the multiples.

Prime `N` is the degenerate case with maximum jump 0 — perfectly flat, always unicursal, no structure to tune at all. **The design lesson is exact: if you want the decomposition to be a tunable parameter, choose composite `N`; if you want it to be immune to a mis-set parameter, choose prime `N`.**

28.4.7 4.7 The unicursal hexagram, and how the tradition escaped the theorem

Here is a case where the historical record and the theorem meet head-on, and the record turns out to record the theorem's constraint.

`{6/2}` is two triangles and cannot be drawn without a lift. The nineteenth- and twentieth-century occult revival nonetheless produced a figure called the *unicursal hexagram*: a six-

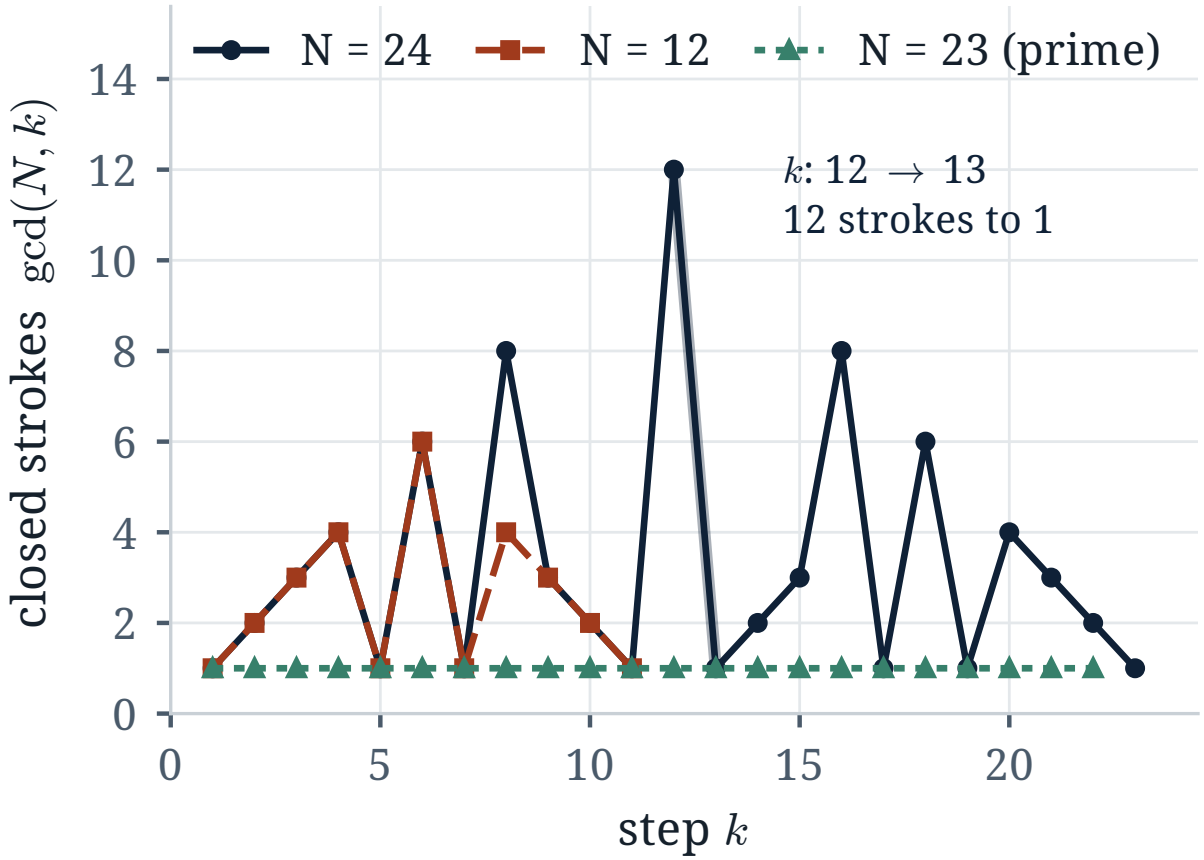


Figure 28.2: **The stroke count has no gradient, so it cannot be tuned by small steps.** On 24 points a single unit of the parameter takes the figure from twelve independent strokes to one; on 12 points the largest jump is 5; on prime 23 the function is flat at one stroke for every step. Computed directly as `gcd(N,k)`, matching the jump table of §4.6; produced by `figs/src/CCP-028_fig01.py`.

This is the clearest single illustration of what the reading claims and of what it refuses. The claim: the old practice was operating on a real structural constraint and its artefacts record the constraint accurately. The refusal: nothing here says the figure does anything.

28.4.8 4.8 Spare's method is lossy compression, and here is its bit budget

Austin Osman Spare's sigil procedure, as described in his *The Book of Pleasure (Self-Love)* (1913), has three steps:

1. Write a statement of intent as a sentence.
2. Strike out every repeated letter, keeping first occurrences only.
3. Recombine the remaining letters into a single glyph.

Steps 1 and 2 constitute an encoding operation, and it may be described as one precisely. Step 2 is a map from strings to sets: it discards multiplicity and it discards order. The formal object is

```
sigma : strings over A  ->  subsets of A
sigma(w) = { letters occurring in w }
```

with the first-occurrence ordering retained only as a drawing convenience.

The hard ceiling. The codomain of `sigma` is the power set of a 26-letter alphabet.

```
|codomain| = 2^26 = 67,108,864
output entropy <= log2(2^26) = 26 bits, FOR ANY INPUT OF ANY LENGTH
```

That bound does not grow. A statement of intent of ten words or ten thousand produces at most 26 bits. Layer by layer:

retained letters m	C(26, m)	bits
8	1,562,275	20.575
10	5,311,735	22.341
12	9,657,700	23.203
13	10,400,600	23.310
14	9,657,700	23.203
16	5,311,735	22.341
18	1,562,275	20.575
20	230,230	17.813

The widest layer is `m = 13`, at 23.31 bits — a sigil that retains exactly half the alphabet carries the most information any sigil can carry, and it is under 24 bits.

Measured on real statements of intent, using a uniform-letter upper bound of `log2(26) = 4.700` bits per input letter:

statement	letters	retained	ratio	input bits (bound)	output bits	destroyed
THIS IS MY WILL TO BUILD A QUIET ROOM	29	16	0.552	136.31	22.34	113.97
IT IS MY WILL TO FINISH THE MANUAL	27	14	0.519	126.91	23.20	103.71

statement	letters	retained	ratio	input bits (bound)	output bits	destroyed
MY WILL IS TO SLEEP WITHOUT FEAR	26	15	0.577	122.21	22.88	99.33
THIS IS MY WISH TO OWN A QUIET ROOM	27	14	0.519	126.91	23.20	103.71
I WILL BUILD THE INSTRUMENT AND KEEP IT	32	16	0.500	150.41	22.34	128.07

Retained sets, in first-occurrence order:

```
THIS IS MY WILL TO BUILD A QUIET ROOM     ->  THISMYWLOBUDAQER
IT IS MY WILL TO FINISH THE MANUAL        ->  ITSMYWLOFNHEAU
MY WILL IS TO SLEEP WITHOUT FEAR          ->  MYWILSTOEPHUFAR
THIS IS MY WISH TO OWN A QUIET ROOM       ->  THISMYWONAQUER
I WILL BUILD THE INSTRUMENT AND KEEP IT   ->  IWLBUDTHENSRMAKP
```

Preimage size — the measure of how thoroughly the operation is not invertible. The number of length-`n` strings that use exactly a given `m`-letter set is a surjection count, `sum over i of (-1)^i C(m,i) (m-i)^n`. Computed:

```
length 28 over a given 15-letter set:   4.716e+31   = 2^105.2 strings
length 30 over a given 14-letter set:   3.691e+33   = 2^111.5
length 32 over a given 16-letter set:   2.499e+37   = 2^124.2
length 36 over a given 17-letter set:   1.787e+43   = 2^143.7
```

Roughly 10^31 to 10^43 sentences of the right length collapse onto any one sigil's letter set. The operation is a digest, not a code: it cannot be decoded, and that is by construction, not by accident.

Measured on a natural-language corpus for calibration: over the 218,002 alphabetic words of length 6 or more in `/usr/share/dict/words`, there are 101,517 distinct letter-sets, mean class size 2.147, and **71.4% of words share their letter-set with at least one other word.** Input entropy 17.734 bits, output entropy 16.037 bits, 1.697 bits destroyed — far gentler than the sentence case, because words are short and the reduction has less multiplicity to throw away. The larger the input, the more brutal the compression, and a sentence is a large input.

28.4.9 4.9 The attention-binding claim, badged and bounded

The tradition attaches a second step to the encoding: the operator is to charge the glyph and then forget the sentence, retaining the form and losing the wording. **This paper can describe the encoding exactly and can describe the second step only as a claim about the operator's own perceptual and attentional state — and it makes no claim of any effect outside the operator.**

What can be said in the formalism, and stated with its badge on:

(a) The reduction genuinely makes the source unrecoverable, and that is measured. The 10^31-to-1 preimage figure above is not a metaphor for forgetting; it is the actual non-invertibility of the map. If the practitioner's stated aim is to hold a form while the wording becomes irrecoverable, the operation performs exactly that, as arithmetic.

(b) A glyph is a small, stable, repeatable perceptual object. In the vocabulary of the framework's moment equation `M = (D/P) E`, a glyph is an environment term `E` with very low variance across presentations [Curzi, 2026o]. Sections 7 and 8 develop the consequence: a low-variance, repeatedly presented `E` is *periodic forcing* of a leaky integrator, and cadence beats intensity for that reason and no other.

(c) The compression is the point, mnemonically. The companion paper measured the cost of holding relations: `n(n−1)/2` unstructured versus `n` structured [Curzi, 2026u]. A sigil is the same move at the extreme — one object standing for a proposition — and the reason it is cheap to rehearse is that it is under 24 bits.

(d) Nothing above is evidence that the practice achieves its traditional aims. No efficacy study is cited here, because none was found and none will be fabricated. What is offered is a description of the encoding and an honest statement of what the formalism can and cannot reach. The formalism reaches the operator's own storage law – `Q_{t+1} = (1−delta)Q_t + eta M_t − C_t` [Curzi, 2026a,o] – and it stops at the operator's skin. **Whether an intention held with high `eta` and low `delta` in one person's stored order changes that person's subsequent behaviour is an empirical question with an obvious affirmative answer and no mystery in it; whether a drawn figure does anything else is a question this paper refuses.**

28.4.10 4.10 Composing the two halves of the analysis

The two results join. Lay the retained letters of a Spare reduction on the vertices of a regular `N`-gon, `N` = the number of retained letters, and connect them with a constant step `k`. Then:

```
N           = number of retained letters (measured: 14 to 17 in the examples above)
gcd(N, k)   = number of separate strokes the glyph requires
N / gcd(N,k) = letters bound into each stroke
unicursal   <=>  gcd(N, k) = 1
```

For `N = 14` (from `IT IS MY WILL TO FINISH THE MANUAL`), the table in §4.1 gives: `k = 1, 3, 5` unicursal; `k = 2, 4, 6` give two strokes of seven; `k = 7` gives seven strokes of two. `phi(14) = 6`, so 6 of the 13 non-trivial steps draw a single line.

For `N = 16` (from `THIS IS MY WILL TO BUILD A QUIET ROOM`), `k = 1, 3, 5, 7` unicursal; `k = 4` gives four squares; `k = 8` gives eight digons. `phi(16) = 8`.

The composition is a design rule, not a discovery: if you want one stroke, choose a step coprime to the letter count; if you want the glyph to fall into `g` visually separable parts, choose `k` a multiple of a divisor `g` of `N`, and the parts will each carry `N/g` letters. It is exactly the same rule that governs the review-cycle schedule in the worked example of Section 6.

28.5 5 Specification and failure modes for star-polygon design

28.5.1 5.1 Specification

SP1 — Compute `gcd(N,k)` before drawing, scheduling, or committing to a figure. It is one line of arithmetic and it determines the stroke count, the stroke length, the pen lifts and the parallelism. Print it.

SP2 — State `N` and `k` explicitly in any figure specification. A figure described as "a six-pointed star" is under-specified: `{6/2}` and the unicursal hexagram are different objects with different stroke counts.

SP3 — Choose prime `N` for guaranteed unicursality; choose highly composite `N` for a rich decomposition menu. `phi(N)/N` is the unicursal fraction and the divisor set is the menu; both are printed in §4.3 for `N = 3..24`. These goals are opposed and the choice must be made once, early, and on purpose.

SP4 — Never expect to tune `k` by small steps. `gcd(N, .)` has no gradient. Enumerate the divisors, evaluate each candidate, and pick. §4.6 shows single-unit jumps as large as 11.

SP5 — Report total work separately from latency. `W = N` always; `T = N/g`; `speedup = g`. A parallelisation that appears to reduce work has been mis-measured.

SP6 — Price the ceiling loss whenever you buy parallelism. Splitting into `g` disjoint strokes lowers the per-period surplus ceiling from `ln N` to `ln(N/g)`. Print both numbers side by side before deciding.

SP7 — Use the rate criterion where throughput matters. `R(x) = ln x / x`, maximised at `x = e`; the best integer stroke length is 3 when 3 divides `N`. Badge it as a ceiling rate, not an achieved rate.

SP8 — Use the full-ceiling configuration where depth matters. `gcd(N,k) = 1` is the only way to keep `ln N`, and it costs `N` steps of latency. Depth and speed are not simultaneously available.

SP9 — When a figure must be unicursal and `gcd` forbids it, vary the step. The obstruction is a theorem about constant steps only. Enumerate the Hamiltonian cycles under whatever chord constraint you actually need — §4.7 does this for six points and finds exactly one figure up to symmetry.

SP10 — Describe the Spare reduction as an encoding, in those words. It is a map from strings to subsets. Report the retained-letter count `m`, the output bound `log2 C(26,m)` bits, and the preimage magnitude. Do not describe it as compression without saying *lossy*.

SP11 — State the 26-bit ceiling whenever sigil information content is discussed. Output entropy is bounded by 26 bits for any input length; the widest layer is 23.31 bits at `m = 13`.

SP12 — Keep the metaphor badge on every occurrence. "Computing map" is the author's metaphor and must be marked at each use. The parallelism statement — strokes `= gcd(N,k)` — is not a metaphor and should be stated without hedging, because it is a theorem.

SP13 — Claim nothing outside the operator. Any statement about a sigil's effect is bounded by the operator's own storage law and stops at the operator's skin. See §4.9(d).

SP14 — Verify by walking, not by formula, the first time. Every table in Sections 3 and 4 was produced by explicit trace and then checked against `gcd`. A theorem you have not tested on the object is a theorem you have not understood.

28.5.2 5.2 Failure modes

FP1 — SILENT LIFT. *Signature:* a figure or a schedule is specified as continuous, and in execution someone lifts the pen or breaks the chain, and no one records it. *Mechanism:* `gcd(N,k) > 1`. The break was structurally required, not a lapse. *Fix:* compute `gcd` first. If `g > 1`, declare the `g - 1` breaks as part of the specification and put a named owner on each. An undeclared break becomes an untracked hand-off, which is the classic Conway's Law seam.

FP2 — PHANTOM PARALLELISM. *Signature:* a schedule is described as having `g` parallel workstreams, and work in fact crosses between them constantly. *Mechanism:* the strokes of `{N/k}` are *disjoint*. If real traffic crosses, the model is not `{N/k}`, and `gcd` does not describe the actual topology. *Fix:* measure the cross-stroke traffic. If it is non-zero, the correct model is a graph with those edges added, and the component count must be recomputed on the real graph, not the idealised one.

FP3 — CEILING BLINDNESS. *Signature:* a team is split into small parallel cells for speed, throughput improves, and the quality of what emerges degrades in a way no one can locate. *Mechanism:* the ceiling fell from `ln N` to `ln(N/g)`. For `N = 12, g = 4`: from 2.484907 to 1.098612 nats, a 55.8% reduction in the maximum surplus any single period can produce. *Fix:* SP6. Print both ceilings. If depth is the requirement, run at least some periods at `gcd = 1`; alternate configurations rather than committing to one.

FP4 — GRADIENT HUNTING. *Signature:* a parameter is being tuned by small increments and the behaviour is erratic, so the increments get smaller. *Mechanism:* the response is `gcd(N, k)`, which is discontinuous everywhere it matters. Smaller steps do not help; they cannot. *Fix:* stop tuning. Enumerate the divisors of `N`, evaluate the finite candidate set exhaustively, and choose. There are at most `floor(N/2)` candidates and for `N <= 24` that is at most 12.

FP5 — GCD MYSTICISM. *Signature:* the arithmetic invariant is presented as though it explained the historical or symbolic *significance* of a figure, rather than its stroke count. *Mechanism:* an ordinary category error — a structural fact recruited to carry interpretive weight it cannot bear. It is the same error as the companion paper's hash mysticism [Curzi, 2026u], here with a better theorem and therefore a stronger temptation. *Fix:* the theorem says how many strokes. It says nothing else. Where a sentence about the figure would be false if the figure had been invented yesterday to organise a warehouse, that sentence is outside the theorem.

FP6 — CONSTANT-STEP ASSUMPTION. *Signature:* "this figure cannot be drawn in one line" is asserted, and someone then draws it in one line. *Mechanism:* Theorem 1 governs constant steps. Variable-step tours are a different and larger search space. The unicursal hexagram lives in the gap. *Fix:* state the step regime with the claim. "No constant step draws `{6/2}` in one line" is true; "no single line touches six points in a star pattern" is false, with exactly one counterexample up to symmetry.

FP7 — DIGEST-AS-CODE. *Signature:* someone attempts to reconstruct the original statement of intent from a sigil, or treats the glyph as carrying the sentence. *Mechanism:* the preimage of a 15-letter set among 28-letter strings is about `2^105`. The map is not invertible and the loss is total. *Fix:* keep the statement separately if you need it. If the practice requires that the statement be irrecoverable, the encoding already guarantees it; nothing further is needed and nothing can undo it.

FP8 — UNBOUNDED-SIGIL FALLACY. *Signature:* a longer, more elaborate statement of intent is believed to produce a richer glyph. *Mechanism:* output entropy is capped at 26 bits regardless of input length, and at 23.31 bits in practice. Beyond about 13 retained letters the layer count *falls*: `C(26,20) < C(26,13)` by a factor of 45. *Fix:* if information content is the goal, target `m` near 13. If it is not the goal, do not invoke information content as a reason for elaboration.

FP9 — PRIME COMPLACENCY. *Signature:* a prime `N` is chosen for its guaranteed unicursality, and later someone needs a parallel decomposition and cannot have one. *Mechanism:* prime `N` has no proper divisors, so `gcd(N,k) = 1` for every `k` and the only attainable stroke count is 1. Maximum jump 0: perfectly robust, perfectly inflexible. *Fix:* decide early which property is load-bearing. If both may be needed, choose `N` with a small number of divisors — `N = 2p` for prime `p` gives exactly the stroke counts 1, 2, and `p`, which is a usable three-item menu with `phi(2p) = p−1` unicursal steps.

28.6 6 Worked example: a twelve-station design review

A hardware group runs a design review across **twelve stations** — twelve specialists who must each see and sign every design package. The current process is a single chain: station 1 to 2 to 3, and so on to 12. One package per hand-off, one hand-off per working day. The group wants throughput and does not want to lose depth. The procedure of Sections 3 to 5 is applied.

Step 1 — Identify `N` and enumerate the menu (SP1, SP3).

```
N = 12          divisors: 1, 2, 3, 4, 6, 12
attainable stroke counts (proper divisors): 1, 2, 3, 4, 6
candidate steps k = 1..6 (Corollary 5 mirrors the rest)
phi(12) = 4, so only k = 1, 5, 7, 11 are unicursal: a 33.3% unicursal fraction
```

Twelve is the worst-case cardinality for single-stroke figures and the best case for decompositions. That is the trade the group has already made by having twelve people.

Step 2 — The configuration table. From §4.1 and §4.5, with the ceiling and rate columns:

Config	g = strokes	stroke size x	latency (days)	ceiling `ln x`	rate `ln x / x`
`{12/1}` current	1	12	12	2.484907	0.207076
`{12/2}`	2	6	6	1.791759	0.298627
`{12/3}`	3	4	4	1.386294	0.346574
`{12/4}`	4	3	3	1.098612	0.366204
`{12/5}`	1	12	12	2.484907	0.207076
`{12/6}`	6	2	2	0.693147	0.346574

Explicit orbits for the leading candidate, `{12/3}`, stations numbered 0-11:

```
stroke A:  0 -> 3 -> 6 -> 9  -> 0
stroke B:  1 -> 4 -> 7 -> 10 -> 1
stroke C:  2 -> 5 -> 8 -> 11 -> 2
```

Three closed loops of four, disjoint, twelve hand-offs total, two pen lifts — that is, two declared seams between loops that require an explicit owner (FP1).

Step 3 — Storage law with real parameters. The group's measured parameters:

```
eta   = 0.45      fraction of review interaction that becomes durable design knowledge
delta = 0.08      per-period decay (rotation onto other projects, standards drift)
Q_t   = 6.0       current stored design capability, nat-scaled
C_t   = 0.05*(g-1) + 0.15    coordination cost: a base of 0.15 plus 0.05 per merge seam
```

The cost model is the group's own estimate, badged as an estimate: each additional stroke adds one merge point where the loops' conclusions must be reconciled, and the group prices that at 0.05.

One-step change, `dQ = eta * S - delta * Q_t - C_t`, taking `S` at each configuration's ceiling (the optimistic case; real `S` is lower and the ranking is unaffected):

Config	strokes	members/stroke	`S` ceiling	`C_t`	`eta*S`	`dQ`
`{12/1}`	1	12	2.484907	0.1500	1.1182	**+0.488208**
`{12/2}`	2	6	1.791759	0.2000	0.8063	+0.126292
`{12/3}`	3	4	1.386294	0.2500	0.6238	-0.106168
`{12/4}`	4	3	1.098612	0.3000	0.4944	-0.285624
`{12/6}`	6	2	0.693147	0.4000	0.3119	-0.568084

The rate-optimal figure `{12/4}` — four triangles, the configuration §4.5 identified as maximal in `ln x / x` — has a *negative* one-step change in stored order. The two results are not in conflict and the reason is exactly the trade in §4.4: `{12/4}` maximises surplus per unit of *latency*, and the storage law is denominated per *period*, not per unit of latency. If the group runs four three-person loops, each period produces at most 1.098612 nats of surplus, `eta` converts 0.494 of it, and decay plus merge cost consumes 0.780. Order falls.

The reconciliation is quantitative and it is the real finding. `{12/4}` completes a cycle in 3 days rather than 12 — four times as many periods per quarter. Over a 60-working-day quarter:

```
{12/1}:  60/12 = 5 cycles    per-cycle dQ = +0.488208   quarter total = +2.441040
{12/3}:  60/4  = 15 cycles   per-cycle dQ = -0.106168   quarter total = -1.592513
{12/4}:  60/3  = 20 cycles   per-cycle dQ = -0.285624   quarter total = -5.712489
```

The single chain wins outright on this cost model, because the merge cost grows with stroke count while the ceiling falls with it, and running the loss more often only compounds

the loss. The rate criterion `ln x / x` assumed the only cost was latency. It is not; `C_t` rises with `g` and `delta * Q_t` is charged every period regardless. **The rate theorem is a ceiling result and this is what it looks like when the ceiling is not the binding constraint.**

Step 4 — Headroom for each configuration. The master gauge, at the unpartitioned ceiling `ln 12`:

```
h = eta * ln 12 - delta * Q_t - C_t = 1.118208 - 0.480000 - C_t

{12/1}: h = +0.488208
{12/2}: h = +0.438208
{12/3}: h = +0.388208
{12/4}: h = +0.338208
{12/6}: h = +0.238208
```

All positive: the group has room to grow in every configuration *if* it can reach `ln 12`. Only `{12/1}` and `{12/5}` can. **Headroom exists and only the unicursal configurations can spend it, which is the entire decision in one line.**

Step 5 — The chosen design. The group adopts `{12/5}` rather than `{12/1}`, and the reason is not in the surplus arithmetic at all — the two are identical there, both `g = 1`, both twelve steps, both ceiling `ln 12`. The reason is the *sequence*:

```
{12/1}:  0,1,2,3,4,5,6,7,8,9,10,11
{12/5}:  0,5,10,3,8,1,6,11,4,9,2,7
```

`{12/1}` visits adjacent stations, which in this group means adjacent specialisms; `{12/5}` visits stations five apart, so each hand-off crosses to a distant specialism and returns. **Both have the same `N`, the same `g`, the same ceiling and the same latency; they differ in how much of the ceiling is reachable, because `u` between adjacent specialists is lower than `u` between distant ones.** That is the framework's engineering of `u` reappearing as a choice of step [Curzi, 2026c], and the star polygon is the object that makes the choice enumerable: there are exactly `phi(12)/2 = 2` distinct unicursal figures on twelve points, `{12/1}` and `{12/5}`, and the group can simply compare them.

Estimating `u` from the group's own records — mean non-parallelism 0.31 for adjacent pairs, 0.58 for pairs five apart:

```
{12/1}:  N_eff = 1 + 11*0.31 = 4.410000    F = ln 4.410000 = 1.483875 nats  (59.7% of ln 12)
{12/5}:  N_eff = 1 + 11*0.58 = 7.380000    F = ln 7.380000 = 1.998774 nats  (80.4% of ln 12)

dQ {12/1} = 0.45*1.483875 - 0.48 - 0.15 = +0.037744
dQ {12/5} = 0.45*1.998774 - 0.48 - 0.15 = +0.269448
```

Changing the step from 1 to 5 — the same twelve people, the same twelve hand-offs, the same twelve days, one number changed in the schedule — multiplies the per-cycle gain in stored order by 7.14. Over five cycles a quarter: +0.188718 against +1.347241.

That is the engineering claim in one worked number, and it rests on nothing beyond Theorem 1, the surplus functional [Curzi, 2026] and the storage law [Curzi, 2026a]. The figure told the group how many strokes were available and which steps were unicursal. The surplus law told them which unicursal step was worth taking. The storage law told them what it was worth.

A star polygon is a schedule with an exact arithmetic invariant, and reading it as one is not a metaphor. Reading it as a computing map is a metaphor, and the badge stays on.

28.7 7 The cadence theorem and the loci law

The companion paper read the correspondence table as a finite relational graph [Curzi, 2026u], and Sections 3 to 6 read the star-polygon figure as an encoding with an exact arithmetic invariant. This section takes the two remaining inherited technologies that are large enough to have architecture – **the memory palace and the liturgical calendar** – and shows that each is a named instance of a law the framework has already proved.

Neither reading requires anything to be added to the formalism. The memory palace is the storage law with space as the retention medium. The ritual calendar is a periodic forcing function applied to a first-order low-pass filter. Both statements are exact, both are computable in advance, and neither carries any claim beyond arithmetic.

Three results, named up front:

- **The loci method is not a trick for holding more. It is a technique for holding the same amount with a smaller `delta` and a larger `eta`, and the gain is the product of three separate factors, each measurable.** In the worked case it is a factor of 14.
- **The case for regular practice over heroic effort has two halves, and only one of them is the ceiling argument.** The second half is Theorem 2 below, which follows from the surplus functional's strict concavity [Curzi, 2026] and has nothing to do with saturation: dividing a fixed budget of interaction into more, smaller occasions never delivers less, and usually delivers much more.
- **The comparison between one large impulse and the same total delivered periodically is not ambiguous, but it is subtler than the folklore.** Over infinite time and with no ceiling the two are exactly equal in integrated effect. Three things differ, and all three are computed here: the steady state, the minimum, and – once the logarithmic ceiling [Curzi, 2026c] is applied – the deliverable total.

The load-bearing claim is the last one, and it is arithmetic rather than exhortation: **a single impulse, however large, has a steady state of exactly zero. A cadence has a steady state of `m/delta`. That is the whole difference, and everything else in the practice literature is a gloss on it.**

28.7.1 7.1 Statement one — the loci law

STATEMENT (the method of loci, read formally). Let a set of `N` items be encoded by binding each to a distinct location in a familiar, mentally traversable space. Then, in the storage law `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t`:

- the **non-parallelism `u`** rises, because two items bound to two different places are, by construction, framed differently;
- the **conversion efficiency `eta`** rises, because retrieval has an addressable index — the walk — rather than requiring undirected search;
- the **retention loss `delta`** falls, because the index is a spatial layout that the retriever already maintains for other reasons and does not have to keep alive on this material's account.

This is a reading, not a theorem. What is a theorem is what follows once the three parameter movements are granted: the steady state moves by their product,

```
Q_inf = eta F(u) / delta         (the storage law's steady state, with c* = 0)

Q_inf'   eta'   F(u')   delta
------ = ---- . ----- . ------
Q_inf    eta    F(u)    delta'
```

and the three factors are separately measurable, separately attackable, and multiply. §8.2 computes them.

The tradition is old, continuous, and explicit about its own mechanism. Cicero (*De Oratore* II.86) attributes the technique to Simonides of Ceos, said to have identified the dead at a collapsed banquet hall by recalling where each guest had been sitting. Whatever the anecdote's historicity, its content is the claim being made here: **the retrieval index was the seating plan.** The fullest surviving classical treatment is the anonymous *Rhetorica ad Herennium* (c. 86-82 BCE), Book III; Quintilian gives a second account in the *Institutio Oratoria*, Book XI. Frances Yates's *The Art of Memory* (1966) is the standard history of how the technique travelled from Roman rhetoric through the medieval and Renaissance traditions.

28.7.2 7.2 Statement two — the cadence theorem

The framework established that cadence beats intensity because surplus saturates at `ln N` per period while cost does not saturate [Curzi, 2026a,c]. That argument depends entirely on the ceiling. **There is a second, independent argument that does not, and it is a theorem.**

THEOREM 2 (the cadence theorem). Let a fixed budget `T` of interaction be spread over `n` periods, `T/n` in each, and let the surplus generated in a period be `g(L)` where `L = T/n` is that period's allocation. If `g` is concave on `[0, infinity)` with `g(0) = 0`, then

```
D(n) := n * g(T/n)
```

is nondecreasing in `n`, and strictly increasing wherever `g` is strictly concave.

Proof. Treat `n` as a positive real. Write `L = T/n`, so `D = T g(L)/L`. Then

```
dD/dn = g(L) - L g'(L)
```

by the chain rule. Concavity of `g` with `g(0) = 0` gives `g(L) = g(L) - g(0) = integral_0^L g'(s) ds >= L g'(L)`, because `g'` is nonincreasing. Hence `dD/dn >= 0`, with strict inequality wherever `g'` is strictly decreasing. ■

The hypothesis is satisfied by the framework's own surplus functional. `f(u) = ln(1 + (N-1)u)` is strictly concave on `[0,1]` and `f(0) = 0` by the axiom set that characterises it [Curzi, 2026]. If the non-parallelism achieved in a session is any nondecreasing concave function of session length — and in particular if it is *linear* in length up to saturation, the simplest possible assumption — then the composite `g(L) = f(u(L))` is concave through the origin and Theorem 2 applies without further assumption.

Therefore the case for regular practice does not rest on the ceiling. It rests on the same concavity that makes the surplus law a logarithm, which is the same concavity found in sensation [Fechner, 1860, Weber, 1834]. Diminishing returns within one occasion is precisely what makes many occasions worth more than one.

28.7.3 7.3 What is proven and what is postulated

Element	Status	Warrant
`D(n) = n g(T/n)` nondecreasing in `n` for concave `g`, `g(0)=0`	**PROVED**	§7.2
`f(u) = ln(1+(N-1)u)` is strictly concave with `f(0)=0`	**PROVED**	the surplus functional's axioms

Element	Status	Warrant
A single impulse has steady state `Q_inf = 0`; a rate `m` has `Q_inf = m/delta`	**PROVED**	the storage law's steady state
Impulse and cadence have equal *integrated* effect absent a ceiling	**PROVED** (LTI)	§8.4, computed
Steady-state cycle: `Q_max = A/(1−a^P)`, `Q_min = Q_max a^P`	**PROVED**	periodic-forcing algebra
Required retention for a rite of interval `P`: `a^P >= r/(1+r)`, `r = theta/A`	**PROVED** (algebra)	§8.10
Costly-requirement break-even: `dc <= (eta F − c)(1 − delta'/delta)`	**PROVED** (algebra)	§8.11
That the method of loci raises `u` and `eta` and lowers `delta`	**POSTULATE** (reading)	§7.1, badged
That a ritual calendar is well described by the storage law at all	**POSTULATE** (modelling)	badged; see §7.4
That any particular `(eta, delta, u)` triple describes a particular practice	**REQUIRES MEASUREMENT**	the framework's survey protocols
Durkheim's, Rappaport's and Whitehouse's mechanisms for what ritual *does*	**SOCIAL THEORY**, not proved here	§8.9-8.10

28.7.4 7.4 Honesty clauses

- **Nothing in this section claims that a symbol, a place, a gesture or a date has any effect other than through a human being who encodes, retrieves, or attends.** The memory palace is a filing system. The calendar is a schedule. Both are effective; neither is mysterious.
- The illustrative parameter values in §8.2 and §8.4 are *design values*, chosen to make the structure legible, and they are labelled as such every time they appear. They are not measurements of any population. The protocols for obtaining real ones are the framework's survey instruments for `delta`, `eta` and `u` [Curzi, 2026c].
- As always, surplus as deployed is a **lower bound** on the ensemble's information [Curzi, 2026], so every `Q_inf` here is a floor.
- The anthropological literature cited in §8.9-8.13 is theory and observation about human societies. It is used here to name phenomena the formalism can represent. **The formalism does not validate the anthropology, and the anthropology does not validate the formalism.**
- Theorem 2 assumes `g` concave. If per-occasion response is genuinely *convex* over some range — if there is a warm-up cost below which nothing happens at all — the theorem reverses on that range, and §8.8 shows exactly how a fixed per-occasion cost restores an interior optimum. **The theorem does not say "shorter is always better." It says concavity favours division, and it names the term that stops the division.**

28.8 8 Cadence, retention, and periodic forcing

28.8.1 8.1 The classical loci rules are a parameter list

The classical treatment of artificial memory is unusually specific about procedure. The *Rhetorica ad Herennium* (III.16-24) gives rules for choosing loci and for constructing images. Read against the storage law, **each rule moves exactly one parameter, and the set of rules covers the parameter set without a remainder.**

Classical rule (*ad Herennium* III, paraphrased)	Parameter moved	Mechanism in the formalism
Choose loci in a deserted place, not a crowded one	`C` down	fewer competing encodings; less interference cost per period
Set the loci at moderate intervals, distinguishable one from the next	`u_cross` up	between-place non-parallelism is the block term `beta^2` of Thm 4.1
Do not let the loci be too like one another	`u_cross` up, `N` preserved	near-identical loci collapse to one effective locus; `N_eff` falls
Neither too bright nor too dark	`N` preserved	discriminability sets how many categories are actually distinct
Learn the series in order, and be able to begin from any point	`eta` up	multiple entry points make retrieval a graph traversal, not a list scan
Make the images striking, active, unusual	`u_div` up	within-place non-parallelism is the `alpha^2 sin^2 gamma` term
Use images of things you already care about	`eta` up	conversion into durable structure is higher for material with existing attachment
Revisit the palace to keep the loci fresh	`delta` down	maintenance of the index rather than of the contents

That table is the method in miniature. No rule requires a mechanism outside the law, and no parameter of the law is left without a rule. The tradition arrived at a complete parameter-tuning procedure by iteration over roughly two thousand years and wrote it down in the vocabulary it had.

28.8.2 8.2 The arithmetic of place

Take a corpus of `N = 40` items. The logarithmic ceiling caps surplus at `ln 40 = 3.6889` nats per occasion, and no encoding exceeds it [Curzi, 2026c]. What differs between encodings is how much of that ceiling they reach, and the driver is `u`.

Design values, illustrative and badged as such:

Encoding	`u`	`N_eff = 1+(N−1)u`	`F = ln N_eff`	% of `ln 40`
rote list, single undifferentiated context	0.15	6.850	1.9242	52.2%

Encoding	u	N_eff = 1+(N−1)u	F = ln N_eff	% of ln 40
rote plus grouping into chunks	0.40	16.600	2.8094	76.2%
loci, ordinary domestic places	0.75	30.250	3.4095	92.4%
loci, deliberately distinct places, vivid images	0.90	36.100	3.5863	97.2%
loci, idealised	0.97	38.830	3.6592	99.2%

The interesting feature of that column is how quickly it saturates: **an ordinary, undesigned palace already reaches 92% of the ceiling, and all the effort of the classical rules buys the last 8%.** That is the logarithm again. It is also the honest answer to why beginners get most of the benefit immediately and experts spend years on the remainder.

Palace size obeys an exact law and it is unforgiving. The ceiling is ln L in the number of loci L, so:

loci L	5	10	20	40	80	160	320
ln L	1.6094	2.3026	2.9957	3.6889	4.3820	5.0752	5.7683

Every doubling of the palace buys exactly ln 2 = 0.693147 nats, forever. Not more at the start and not less at the end. Going from 5 loci to 10 buys the same as going from 160 to 320. This is the most useful single fact for anyone designing a palace: **capacity is bought in fixed increments at exponentially rising cost in construction effort, so the question is never "how big" but "how many doublings can I afford to build and maintain."**

The two-source decomposition prices place and image separately. The framework's block decomposition [Curzi, 2026] splits u = u_cross + u_div — between-locus and within-locus non-parallelism — and bounds the result:

```
max(f(u_cross), f(u_div))  <=  f(u)  <=  f(u_cross) + f(u_div)
```

Computed at N = 40:

u_cross (places)	u_div (images)	u	f(u_cross)	f(u_div)	lower	f(u)	upper	slack
0.85	0.05	0.90	3.5308	1.0818	3.5308	3.5863	4.6126	1.0263
0.80	0.10	0.90	3.4720	1.5892	3.4720	3.5863	5.0612	1.4749
0.45	0.45	0.90	2.9205	2.9205	2.9205	3.5863	5.8409	2.2546
0.60	0.20	0.80	3.1946	2.1748	3.1946	3.4720	5.3693	1.8974

Two readings, both practical. First, **f(u) sits very close to the lower bound in every row**, which is subadditivity doing its work [Curzi, 2026]: the two sources overlap heavily and one does not get to add them. Second, **the balanced row (0.45/0.45) has the largest gap between the naive sum and the truth** — 2.25 nats of imagined surplus that does not exist. An operator who has invested equally in place-distinctness and image-vividness will overestimate the combined effect by more than one who has invested lopsidedly. That is a counterintuitive and checkable consequence.

The three-factor gain, computed. Period one day. Design values, badged.

Regime	u	eta	delta/day	F	eta F	Q_inf	DC gain eta/delta	tau	t_half
rote rehearsal	0.15	0.20	0.25	1.9242	0.3848	1.5394	0.800	3.48 d	2.41 d

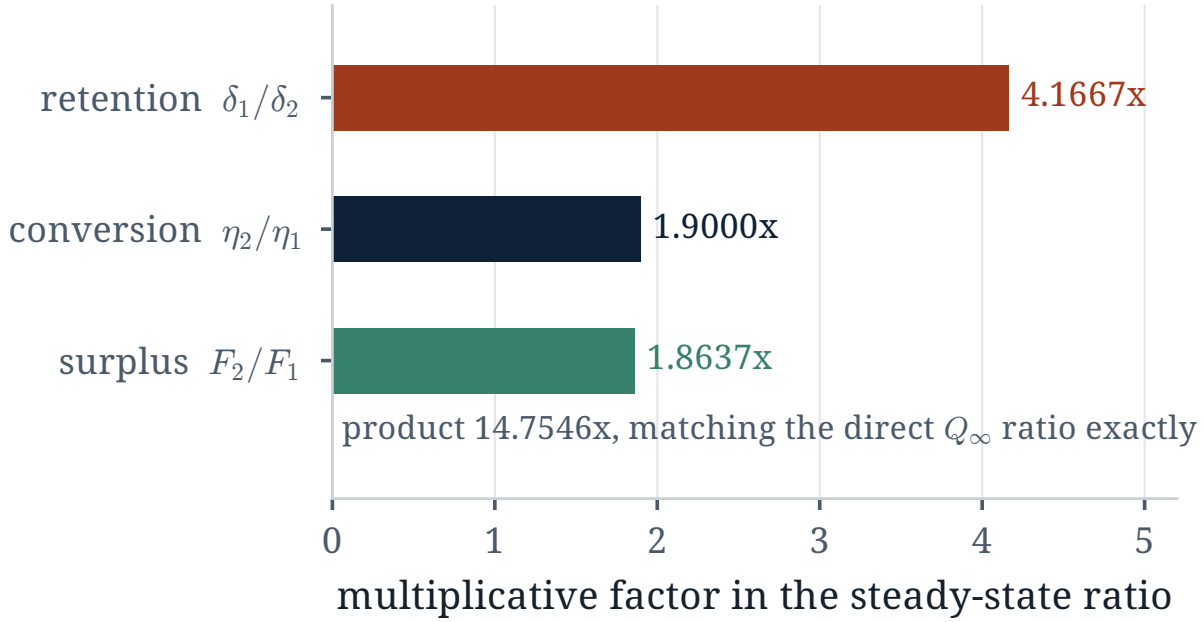


Figure 28.3: **The loci advantage is mostly retention, which is the opposite of the usual emphasis.** Surplus and conversion together give a factor of 3.54; retention alone gives 4.17, and the three multiply to 14.7546, matching the direct steady-state ratio exactly. Modelled from the design values of §8.2 and badged there as design values; produced by `figs/src/CCP-028_fig05.py`.

Regime	u	eta	delta/day	F	eta F	Q_inf	DC gain eta/delta	tau	t_half
rote + grouping	0.40	0.24	0.20	2.8094	0.6743	3.3713	1.200	4.48 d	3.11 d
loci encoding	0.90	0.38	0.06	3.5863	1.3628	22.7132	6.333	16.16 d	11.20 d
loci + spaced review	0.90	0.38	0.03	3.5863	1.3628	45.4264	12.667	32.83 d	22.76 d

Decomposing the loci-versus-rote ratio:

```
surplus factor    F2/F1        = 1.8637
conversion factor eta2/eta1    = 1.9000
retention factor  delta1/delta2 = 4.1667
----------------------------------------
product                        = 14.7546
direct Q_inf ratio             = 14.7546     (identical, as it must be)
```

Two thirds of the advantage is in `delta`. The surplus and conversion factors together give `1.8637 x 1.9000 = 3.54`; retention alone gives 4.17. **The memory palace is primarily a retention device and only secondarily an encoding device, and every popular account has that emphasis backwards.** The framework's general rule reappears in miniature: retention is a gain term, and nobody treats it as one [Curzi, 2026a].

28.8.3 8.3 What the evidence supports, and what it does not

The method of loci is one of the few techniques treated here with a genuine experimental literature, and the honest summary is narrower than the enthusiasm around it.

What is well supported. That the technique substantially improves ordered recall of arbitrary lists is not seriously contested; Bower's review in *American Scientist* (1970) is an early

and clear statement of the experimental case. Maguire, Valentine, Wilding and Kapur (*Nature Neuroscience*, 2003) studied competitors at the World Memory Championships and reported that they did not differ from controls in general cognitive ability or gross brain structure, but did differ in which regions they engaged during encoding, with spatial-navigation areas prominent; the great majority reported using the method of loci. Dresler and colleagues (*Neuron*, 2017) trained novices over several weeks and reported substantial improvement in list recall, partial persistence at a months-later follow-up, and shifts in functional connectivity toward the expert pattern. **In this paper's vocabulary, both studies are about `eta` and `delta`, not about capacity.**

What is well supported and usually forgotten. Ericsson, Chase and Faloon (*Science*, 1980) trained a single subject over more than 200 hours of practice; his digit span rose from about 7 to about 79. His *letter* span did not move — it remained around six. **The stored order was specific to the essence set it was built on.** In the framework's parameter table `Q` is stock, and stock is indexed by the category set that generated it [Curzi, 2026o]. There is no general memory to be improved; there are particular `Q`s built against particular `N`s. Any claim that a mnemonic practice raises "memory" as such is a category error the framework catches immediately.

What the framework contributes and the literature does not. The literature reports outcomes. The storage law says *why the outcome has the shape it has*: the gain is a product of three factors, the largest of which is retention, and each is separately attackable. **The framework's identifiability gate is the gate such a study would have to pass** [Curzi, 2026u], and it is not obvious that it does: two different `(eta, delta)` pairs with the same ratio produce the same steady state, and only the *transient* separates them. The measurement design must therefore observe the approach to steady state and not only the endpoint — which is the framework's decay-estimation method, and why it insists on a zero-input window.

What is not supported by anything here. That the images or places have significance of their own. That a particular palace layout is better than another for reasons other than distinctness and familiarity. That the technique transfers across domains — Ericsson's subject says otherwise.

28.8.4 8.4 One impulse against the same total, computed

This is the calculation the section exists to perform. The ceiling argument is given above. Here the comparison is made without any ceiling first, so that the honest baseline is on the record, and the ceiling is added afterwards.

Setup. `delta = 0.06` per week, so `a = 0.94`, `tau_exact = 16.1615` weeks, `t_half = 11.2023` weeks. A fixed total input budget `B = 26.0` units. Three ways to spend it:

- **Impulse:** all 26.0 in week 0, nothing after.
- **Weekly:** 0.5 per week for 52 weeks, nothing after.
- **Four-weekly:** 2.0 every fourth week for 52 weeks, nothing after.

week	impulse	weekly	four-weekly
0	0.0000	0.0000	0.0000
1	26.0000	0.5000	2.0000
2	24.4400	0.9700	1.8800
4	21.5952	1.8271	1.6612
8	16.8604	3.2536	2.9581
13	12.3739	4.6053	5.7325
26	5.5357	6.6655	7.0583
39	2.4765	7.5872	7.3818
52	1.1079	7.9995	7.2731
78	0.2217	1.6010	1.4556

week	impulse	weekly	four-weekly
104	0.0444	0.3204	0.2913
156	0.0018	0.0128	0.0117

The crossover is week 24. At week 23 the impulse is still ahead, 6.6648 against 6.3254; at week 24 the cadence is ahead, 6.4458 against 6.2649, and it never trails again while the budget lasts.

And now the fact the folklore gets wrong. Summing `Q_t` over 400 periods:

```
impulse      433.3333
weekly       433.3333
four-weekly  433.3333
B/delta   =  433.3333
```

Identical to every digit. This is not an artefact of the parameters. It is the defining property of a linear time-invariant system: the integral of the response to a given total input is `B/delta` regardless of how the input is distributed in time. **Anyone who tells you that spreading the same total input over more occasions produces more total stored order, in a linear model, has not done the arithmetic.** The framework says this plainly and it is repeated here, because the correct argument is stronger when the incorrect one has been cleared away.

28.8.5 8.5 The steady-state difference — the actual result

The integral is invariant. **The steady state is not, and this is where the difference lives.**

A single impulse of any size whatever, never repeated, has

```
Q_inf = 0
```

because the driven term is zero for all `t > 0` and the transient decays as `a^t`. A sustained rate `m` per period has

```
Q_inf = m / delta
```

Tabulated at `delta = 0.06`:

sustained rate `m` per week	`Q_inf = m/delta`	impulse of any size, unrepeated
0.10	1.6667	0.0000
0.25	4.1667	0.0000
0.50	8.3333	0.0000
0.7278	12.1300	0.0000
1.00	16.6667	0.0000

That column of zeros is the mathematical case for regular practice, and it does not depend on the ceiling, on motivation, on commitment, or on any claim about human nature. A quantity that is not driven decays to zero because `delta > 0`. The only way to hold a nonzero level is to keep driving it. This is why the framework insists that the leak is the point: the same `delta` that makes iteration arrive is the `delta` that makes a single arrival not stay [Curzi, 2026a]. It is also why repetition alone cannot move the destination – the fixed point `Q_inf = (eta F − c)/delta` contains no `t` [Curzi, 2026].

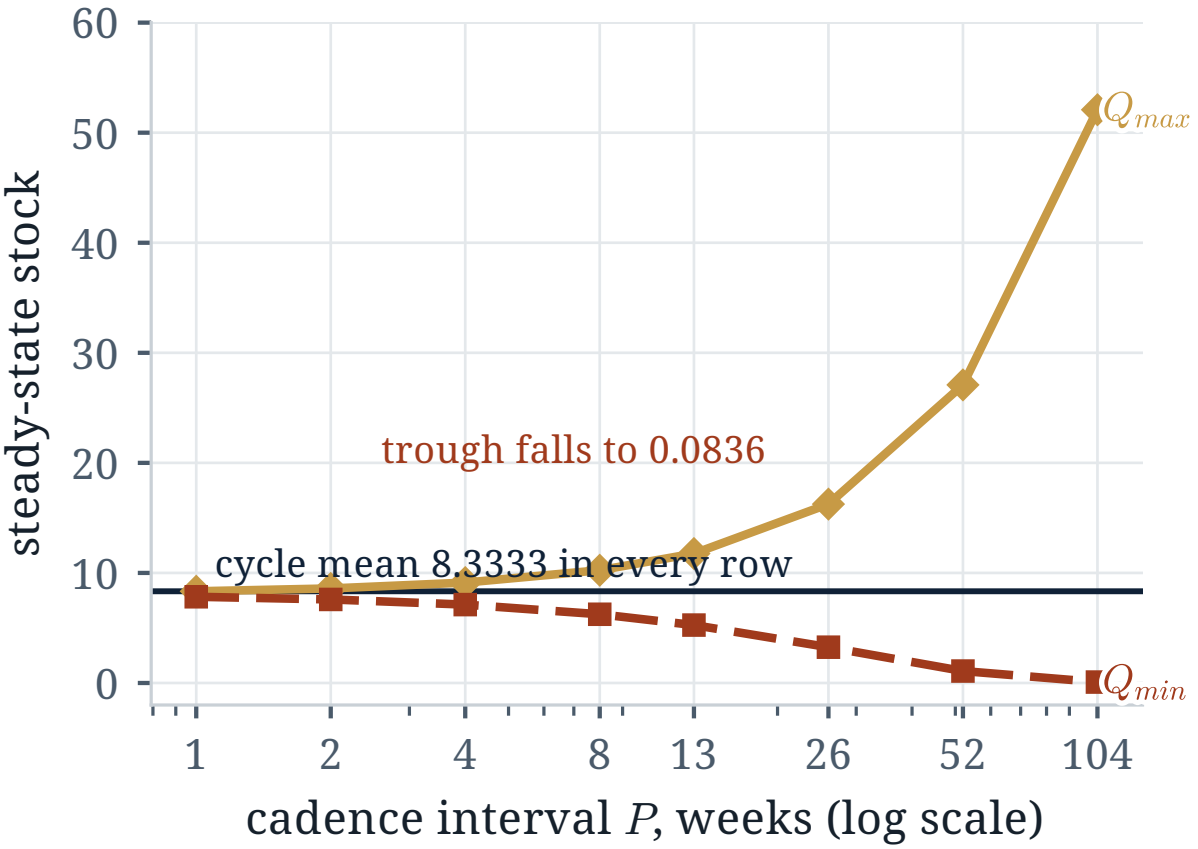


Figure 28.4: **The cycle mean is identical at every cadence; only the trough moves.** Spreading the same annual total over rarer, larger bursts leaves the mean at 8.3333 in every row while the steady-state minimum falls from 7.8333 to 0.0836 — so a system with any functional threshold experiences the trough and not the mean. Modelled at `delta = 0.06` per week with an annual total of 26.0 units; produced by `figs/src/CCP-028_fig04.py`.

Now repeat the impulse. Take the same annual total, 26.0 units, delivered forever at interval `P` with burst `A = 26P/52`. The steady-state cycle is exact:

cadence `P`	burst `A`	`a^P`	`Q_max`	`Q_min`	ripple	cycle mean	% of cycle with `Q >= 7`
1	0.5000	0.940000	8.3333	7.8333	0.5000	8.3333	100.0%
2	1.0000	0.883600	8.5911	7.5911	1.0000	8.3333	100.0%
4	2.0000	0.780749	9.1220	7.1220	2.0000	8.3333	100.0%
8	4.0000	0.609569	10.2451	6.2451	4.0000	8.3333	76.9%
13	6.5000	0.447365	11.7618	5.2618	6.5000	8.3333	64.5%
26	13.0000	0.200136	16.2528	3.2528	13.0000	8.3333	52.4%
52	26.0000	0.040054	27.0849	1.0849	26.0000	8.3333	42.1%
104	52.0000	0.001604	52.0836	0.0836	52.0000	8.3333	31.2%

The mean column is constant at 8.3333 in every row. The minimum falls from 7.83 to 0.08. **If a system has any functional threshold at all, the mean is the wrong statistic and the minimum is the right one.** A community whose shared context must stay above some level for its meetings to work, its obligations to hold, or its newcomers to be absorbed does not experience the mean; it experiences the trough, and the trough is the only column in that table that moves.

28.8.6 8.6 The three quantities that are not invariant

Collecting the finite-budget results from §8.4, with thresholds applied:

threshold `theta`	periods with `Q >= theta`, impulse	weekly	four-weekly
3.0	35	60	61
5.0	27	45	45
7.0	22	25	23
10.0	16	0	0
15.0	9	0	0

Read that table honestly, because it does not say what a partisan of cadence would like it to say. At low thresholds the cadence wins by a wide margin: 60 or 61 periods against 35. At `theta = 7` the three are close. **At `theta = 10` and above the impulse wins outright, and the cadence scores exactly zero, because no cadence funded by that budget ever reaches 10 at all.**

That is the entire formal content of the intensity argument, and it is real. **An impulse buys height that a cadence cannot buy; a cadence buys duration that an impulse cannot buy.** If — and only if — there exists a threshold that must be crossed and that no sustainable rate can reach, the impulse is the correct instrument, and the framework says so without hedging. §8.10 identifies the class of rites that live in exactly that regime and computes what they require.

Three quantities, then, are not invariant to how a fixed budget is spread:

1. **The peak.** `26.0000` for the impulse, `7.9995` for the weekly cadence — a factor of 3.25.
2. **The minimum**, and hence time above any threshold below the cadence's plateau.
3. **The deliverable total itself**, once the ceiling is applied — which is §8.7 and is the largest effect of the three by an order of magnitude.

28.8.7 8.7 The ceiling, applied

The logarithmic ceiling caps `F(u)` at `ln N`, so per-period input is capped at `eta ln N` [Curzi, 2026c]. With `eta = 0.35` and `N = 8`:

```
eta ln N = 0.35 * 2.0794 = 0.7278 per week
```

The 26-unit annual total is now attempted at each cadence, and the bursts are clipped:

cadence `P`	burst wanted	delivered	clipped?	annual delivered	% of weekly
1	0.5000	0.5000	no	26.0000	100.0%
2	1.0000	0.7278	**yes**	18.9229	72.8%
4	2.0000	0.7278	**yes**	9.4615	36.4%
8	4.0000	0.7278	**yes**	4.3668	16.8%
13	6.5000	0.7278	**yes**	2.9112	11.2%
26	13.0000	0.7278	**yes**	1.4556	5.6%
52	26.0000	0.7278	**yes**	0.7278	**2.8%**

The annual rite delivers 2.8% of the weekly practice, from an identical nominal budget. The other 97.2% is not delivered late, or delivered inefficiently. It is not delivered at all: the law caps what one interaction can generate and the excess has nowhere to go.

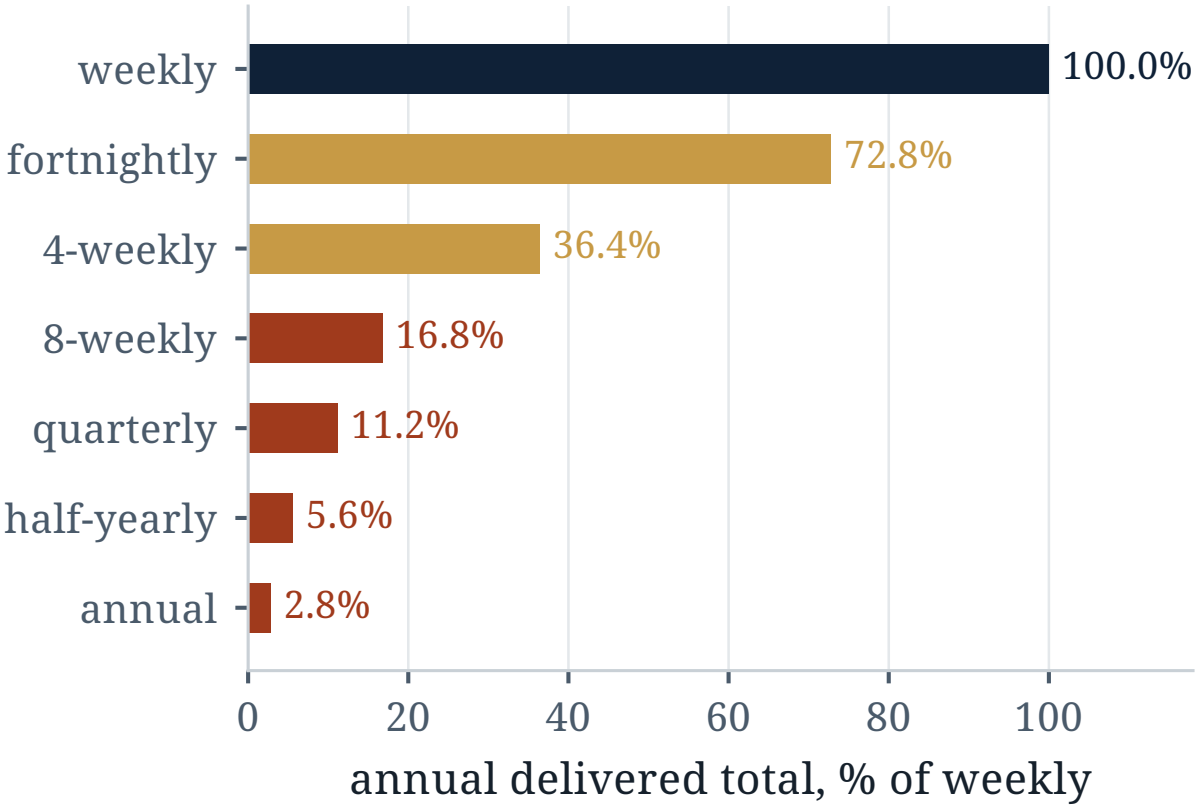


Figure 28.5: **An annual rite delivers 2.8 per cent of what a weekly practice delivers from an identical nominal budget, and the missing 97.2 per cent is not delivered late but not delivered at all.** The per-period ceiling `eta ln N = 0.7278` clips every burst, and a rarer cadence wastes a larger share of each. Modelled at `eta = 0.35`, `N = 8`, annual budget 26.0 units; produced by `figs/src/CCP-028_fig03.py`.

Design consequence, stated once and load-bearing: an annual event and a weekly practice are not two speeds of the same thing. They are different instruments with different maximum outputs, and the comparison "same total time" is meaningless without the clipping calculation.

28.8.8 8.8 The optimal cadence, and the term that stops the division

Theorem 2 says division never hurts. On its own it would recommend infinitely many infinitesimal occasions, which is obviously wrong, and the theorem's honesty requires naming the term that stops it: **the per-occasion fixed cost.**

Let a fixed budget `T = 12` hours be spread over `n` periods with one session of `T/n` hours each. Let non-parallelism rise linearly with session length to saturation at `L* = 8` hours, `u(L) = min(1, L/L*)`. Let `N = 8`, `eta = 0.35`, and let each occasion carry a fixed cost `c_0 = 0.15` regardless of length — convening, travelling, settling, closing.

n	L (h)	u	f(u)	eta f	delivered	−n c_0	net
1	12.0000	1.0000	2.0794	0.7278	0.7278	-0.1500	0.5778
2	6.0000	0.7500	1.8326	0.6414	1.2828	-0.3000	0.9828
3	4.0000	0.5000	1.5041	0.5264	1.5793	-0.4500	1.1293
4	3.0000	0.3750	1.2879	0.4507	1.8030	-0.6000	1.2030
5	**2.4000**	**0.3000**	**1.1314**	**0.3960**	**1.9800**	**-0.7500**	**1.2300**
6	2.0000	0.2500	1.0116	0.3541	2.1244	-0.9000	1.2244
8	1.5000	0.1875	0.8383	0.2934	2.3473	-1.2000	1.1473
12	1.0000	0.1250	0.6286	0.2200	2.6402	-1.8000	0.8402
24	0.5000	0.0625	0.3629	0.1270	3.0484	-3.6000	-0.5516
48	0.2500	0.0312	0.1978	0.0692	3.3235	-7.2000	-3.8765

The **delivered total** column is monotone increasing exactly as Theorem 2 requires, rising from 0.7278 to 3.3235 — **a factor of 4.57 from the same twelve hours, purely by division.** The **net** column peaks at `n = 5` and falls away on both sides.

Note also the first row. At `n = 1`, the session is 12 hours but `u` saturates at 8; **four of the twelve hours deliver nothing at all.** The single long occasion wastes a third of the budget before any of the other effects are considered.

The optimum is a function of the per-occasion cost and of nothing else here:

fixed cost `c_0` per occasion	optimal `n`	session length	net delivered
0.00	(unbounded)	-> 0	-> 3.6275
0.02	24	0.500 h	2.5684
0.05	13	0.923 h	2.0438
0.10	8	1.500 h	1.5473
0.15	5	2.400 h	1.2300
0.25	3	4.000 h	0.8293
0.40	2	6.000 h	0.4828
0.60	1	12.000 h	0.1278
0.80	1	12.000 h	-0.0722

Read the table as a design instrument, not a result. The optimal cadence is not a fact about the practice; it is a fact about how expensive it is to convene. **Halving the cost of convening moves the optimum from five occasions to eight, and doing so is very often cheaper than any intervention on the occasions themselves.** This is the quantitative form of a familiar observation: groups that meet in a place they are already in, at a time they are already free, meet far more often, and the reason is the `c_0` column.

At `c_0 = 0.80` the net is negative at every `n`. **Some practices cannot be afforded at any cadence, and the arithmetic says so before anybody's calendar is consulted.**

28.8.9 8.9 Ritual as periodic forcing: what Durkheim and Rappaport are describing

Durkheim's *Elementary Forms of Religious Life* (1912) proposes that periodic assembly generates a state — *effervescence collective* — that the assembled individuals do not generate separately, and that the periodic re-assembly is what keeps the resulting representations alive between times.

In this framework that is two claims, and they map cleanly. The first is that `S = F(u) > 0` requires more than one non-parallel participant — the surplus functional is exactly zero when the participating framings are parallel [Curzi, 2026], so an assembly of genuinely different people generates what none of them generates alone. The second is the storage law itself: what the assembly produces is an *input*, and inputs decay unless repeated. **Durkheim's insistence that the feast must recur is the observation that `delta > 0`.**

Roy Rappaport's *Ritual and Religion in the Making of Humanity* (1999) is more precise and more useful here. Rappaport defines ritual by its form rather than its content — as the performance of sequences of formal acts and utterances that are largely invariant and not encoded by the performers themselves — and separates two streams of information carried by a performance:

- the **canonical stream**, which is the invariant liturgical order itself and is not about the performer at all;
- the **self-referential stream**, which is what this particular performer, at this particular time, transmits by showing up and doing it.

That distinction is a distinction between `Q` and `S`. The canonical order is stored order: it persists across performers, changes slowly, and is maintained rather than created each time. The self-referential message is the period's input. Rappaport's central claim — that to perform an order is to accept it, whatever one privately believes — is, formally, a statement that this particular `eta` is unusually reliable: participation converts into a public commitment record whether or not the participant intends it to.

Rappaport is also explicit that the propositions at the top of a liturgical order — his **Ultimate Sacred Postulates** — are typically unfalsifiable, and that their unfalsifiability is functional rather than concealed. **That is the framework's badge rule, arrived at independently and stated in another vocabulary.** A postulate declared as a postulate does work that a postulate disguised as a finding cannot do, and Rappaport's point is that the declaration is part of the mechanism.

Two experimental results attach here because they concern the input side rather than the storage side. Konvalinka and colleagues (*PNAS*, 2011) recorded heart-rate dynamics during a fire-walking ritual in San Pedro Manrique, Spain, and reported synchronised arousal between performers and related spectators. Wiltermuth and Heath (*Psychological Science*, 2009) reported that synchronous action raised cooperation in subsequent economic tasks. **Both are evidence about `S` — about what an occasion generates — and neither is evidence about `delta`.** Keeping those two questions apart is most of what this framework contributes to reading the ritual literature.

28.8.10 8.10 The two modes, and what a rare rite must buy

Harvey Whitehouse's *Modes of Religiosity* (2004) distinguishes a **doctrinal** mode — frequent, routinised, low-arousal, transmissible at large scale — from an **imagistic** mode — rare, high-arousal, often painful or frightening, producing intense and enduring episodic memory in small groups. Whitehouse and Lanman (*Current Anthropology*, 2014) develop the second into an account of *identity fusion*.

This is precisely the cadence/intensity axis, arrived at inductively from the ethnographic record, and the framework says something specific about it that the ethnography does not.

A rite at interval `P`, delivering a burst `A`, has steady-state trough `Q_min = A a^P / (1 − a^P)`. Requiring `Q_min >= theta` and solving for the pole:

```
a^P  >=  r / (1 + r),      r = theta / A
```

With the burst at the ceiling, `A = eta ln N = 0.7278`, and a modest functional floor `theta = 0.30`:

interval `P`	in human terms	required pole `a`	required `delta`/week	required `t_half`	in years
1	weekly	0.291884	0.708116	0.56 wk	0.01
4	monthly	0.735026	0.264974	2.25 wk	0.04
13	quarterly	0.909625	0.090375	7.32 wk	0.14
26	half-yearly	0.953743	0.046257	14.64 wk	0.28
52	annual	0.976597	0.023403	29.27 wk	0.56
208	every four years	0.994097	0.005903	117.08 wk	2.25
520	every ten years	0.997635	0.002365	292.71 wk	**5.63**
1560	once a generation	0.999211	0.000789	878.12 wk	**16.89**

A once-a-decade rite requires the trace it leaves to have a half-life of about five and a half years. A once-a-generation rite requires seventeen.

That is the whole explanation of the imagistic mode, and it is arithmetic. **A rare rite is not an inefficient frequent rite. It is a rite operating in a regime where the ordinary `delta` of shared practice would leave nothing at all, and it must therefore purchase an extraordinary `delta` by whatever means are available: ordeal, danger, pain, beauty, terror, secrecy, or singularity.** The ethnographic observation that rare rituals are overwhelmingly the intense ones is exactly what this table predicts. It is also why they are small: the means of purchasing a five-year half-life do not scale, and Whitehouse's account of why imagistic communities stay small is the same constraint seen from the sociological side.

The requirement is not weakly sensitive to the floor:

floor `theta`	as % of burst	required `delta`/week at `P = 520`	`t_half` in years
0.10	13.7%	0.004056	3.28
0.20	27.5%	0.002947	4.52
0.30	41.2%	0.002365	5.63
0.50	68.7%	0.001726	7.72
0.7278	100.0%	0.001332	10.00

Xygalatas and colleagues (*Psychological Science*, 2013), studying the Thaipusam Kavadi in Mauritius, reported that participants in the higher-ordeal rite subsequently gave more to an anonymous charitable donation than participants in a low-ordeal rite. **The framework offers no opinion on whether that is the right measure, and it is one field result in one population.** What it supplies is a reason to expect the *direction*: the high-ordeal condition is the one that can pay for a low `delta`, and what is measured downstream is a trace that has not yet decayed.

The refusal that governs this whole section: none of this recommends an ordeal. The arithmetic explains why rare intense rites take the form they do. It does not license anyone to build one, and an intervention that removes a person's ability to decline is out of scope for this manual regardless of what its parameters look like.

28.8.11 8.11 The costly-signal break-even

Costly-signalling accounts of ritual (Irons, 2001; Sosis and Alcorta, *Evolutionary Anthropology*, 2003) argue that expensive requirements are hard-to-fake signs of commitment, following Zahavi's handicap logic (1975): a signal is credible in proportion to what it would cost a non-committed sender. Sosis and Bressler (*Cross-Cultural Research*, 2003) examined nineteenth-century American communes and reported that religious communes outlasted secular ones, and that within the religious sample the number of costly requirements was associated with longevity — an association not found in the secular sample. **This is historical and correlational and must be read as such**: no randomisation, one time and place, a survivorship-prone record.

The framework's contribution is not to confirm this but to price it. A costly requirement enters the storage law as an increase in `c`, and can only be worth having if it changes something else. The most plausible candidate is `delta`: a requirement that screens out low-commitment members reduces turnover, and turnover is the dominant term in institutional `delta`.

From `Q_inf = (eta F − c)/delta`, requiring `Q_inf' >= Q_inf` after adding cost `dc` and moving `delta -> delta'`:

```
(eta F − c − dc) / delta'  >=  (eta F − c) / delta

dc  <=  (eta F − c) (1 − delta'/delta)
```

The maximum affordable requirement is the current net drive times the fractional reduction in decay it achieves. Worked at `eta = 0.35`, `N = 8`, `u = 0.55` so `F = 1.5790`, `eta F = 0.5526`, `c = 0.15`, `delta = 0.06` per week, giving a baseline net drive of 0.4026 and `Q_inf = 6.7107`:

new `delta`	reduction	max affordable `dc`	as % of net drive
0.055	8.3%	0.0336	8.3%
0.050	16.7%	0.0671	16.7%
0.045	25.0%	0.1007	25.0%
0.040	33.3%	0.1342	33.3%
0.030	50.0%	0.2013	50.0%
0.020	66.7%	0.2684	66.7%
0.010	83.3%	0.3355	83.3%

The percentages in the last two columns are identical, exactly, in every row. That is not a coincidence of the parameters; it falls straight out of the algebra. **A requirement is worth imposing if and only if the fraction of the net drive it consumes is less than the fraction of the decay it removes.** The same cancellation appears elsewhere in the corpus with different furniture: at a 15 per cent annual hard-currency drag, a local-currency allocation of 25.17 per cent of protected capital sets headroom to exactly `−0.500000`, cancelling the whole `eta ln(N) = 0.377583` diversification gain, while a 10 per cent allocation costs 0.150000, or 39.73 per cent of the gain [Curzi, 2026]. **A cost term that scales with the same `ln N` the gain scales with can consume all of it, and the break-even is the only thing that says when.** That is a one-line test, it needs two measurements, and it is the sharpest thing the framework has to say about why demanding institutions sometimes outlast easy ones and sometimes merely exhaust their members.

Worked the other way: a requirement costing 0.05 per week must reduce `delta` by at least 12.4% to pay for itself. If `delta` halves, requirements up to 0.2013 per week are profitable and anything above that is not:

added cost `dc`	`Q_inf'` at `delta' = 0.03`	vs baseline 6.7107
0.000	13.4214	better
0.050	11.7548	better
0.100	10.0881	better
0.150	8.4214	better
0.200	6.7548	better (just)
0.283	3.9881	**worse**
0.350	1.7548	**worse**

The failure mode is visible in the table and it is common: a community that has discovered requirements work keeps adding them past the break-even and cannot tell, because the requirements that worked and the requirements that are now destroying it look identical from inside. The only way to distinguish them is to measure `delta` and `C` separately, which the framework's survey protocols specify.

28.8.12 8.12 The multi-rate calendar

Real liturgical calendars are not single-rate. They superpose a fast cycle (daily or weekly), a medium cycle (monthly, or a lunar month), and a slow cycle (annual, or multi-annual). **In circuit terms this is a bank of forcings at different frequencies driving one low-pass filter, and the framework's frequency-response analysis [Curzi, 2026a] already says what happens: the filter passes the slow component and attenuates the fast one, while the ceiling caps whatever arrives in any single period.**

The consequence is specific and is almost always missed in practice. Computed at `delta = 0.055` per week, `eta = 0.35`, `N = 8`, ceiling `eta ln N = 0.7278`, cost `0.13` per week, with a weekly session at `u = 0.55` (`0.5526`), a monthly seminar at `u = 0.62` (`0.5863`), and an annual retreat at `u = 0.86` (`0.6821`):

```
weekly + seminar in the same week  = 0.5526 + 0.5863 = 1.1390
ceiling                                              = 0.7278
DISCARDED                                            = 0.4112     (36.1% of the week's input)

weekly + retreat in the same week  = 0.5526 + 0.6821 = 1.2347
DISCARDED                                            = 0.5069     (41.1%)
```

Two good sessions in the same period are not two sessions. Beyond the ceiling they are one. In a three-year simulation the annual retreat placed in a seminar week contributed a steady-state mean of 8.4806; moved to a clean week it contributed 8.5418 — a difference of 0.0612, small in absolute terms and worth having for free, and larger in systems where the two events are closer in size.

Design rule: schedule occasions to avoid collision in the same period before scheduling them for convenience. The calendar is a resource-allocation problem against a per-period cap, and the traditional practice of distributing feasts across the year rather than clustering them is, in this reading, exactly the right allocation whatever its original rationale.

28.8.13 8.13 Open loop and closed loop

One distinction remains, and the formalism makes it more sharply than the ritual literature does.

A liturgical calendar is open-loop forcing. It fires on a date, regardless of the state of the system, and everything in §8.5 applies to it.

Some ritual cycles are closed-loop. Rappaport's earlier ethnography, *Pigs for the Ancestors* (1968), analysed the ritual cycle of the Tsembaga Maring of highland New Guinea as a regulator: the great festival was not calendrical but was triggered when the pig herd had grown to the point of straining the households that tended it, and the cycle thereby governed pig population, warfare and land use together. Rappaport described it in explicitly cybernetic terms.

A state-triggered rite is a controller, and controllers have failure modes that calendars do not. A loop acting on a stale measurement overshoots – 30.2% at a lag of two periods, 75.1% at four – and at sufficient gain and lag it diverges outright [Curzi, 2026a].

The two designs trade different risks and both risks are computable. A calendar cannot respond to conditions and will sometimes fire when there is nothing to fire about. A trigger responds but inherits every pathology of delayed feedback. Systems that run both — a fixed calendar for the base rate, a triggered rite for exceptional states — are hedging exactly as the control theory recommends, and they are extremely common.

28.8.14 8.14 The spacing effect, as the empirical counterpart

The claim that distributed practice outperforms massed practice for the same total study time is one of the oldest and best-replicated findings in experimental psychology. Ebbinghaus reported it in *Über das Gedächtnis* (1885), and the quantitative synthesis by Cepeda, Pashler, Vul, Wixted and Rohrer (*Psychological Bulletin*, 2006) reviews the verbal-recall literature and finds the advantage robust across a wide range of conditions. Cepeda and colleagues (*Psychological Science*, 2008) further reported an interior optimum: the best gap between study episodes grows with the interval until test, at a ratio commonly reported in the range of roughly ten to twenty per cent.

Two honest statements, in order.

First, the spacing effect and Theorem 2 are not the same claim and must not be presented as mutual confirmation. The spacing literature is about human memory encoding; Theorem 2 is a statement about concave functions. **What they share is a structural prediction — that division of a fixed budget into more occasions helps, up to a point set by a per-occasion cost — and structural agreement between an empirical regularity and a mathematical result is interesting without being evidence for either.**

Second, the finding that the optimal gap is a roughly fixed *fraction* of the retention interval says that the right cadence scales with the system's own time constant rather than with the calendar — which is the content of the framework's cadence specification, `P <= tau/2`. **Both put the practice interval at a fixed ratio of the system's own timescale, and both put that ratio well below one. That agreement is structural and no more is claimed from it than that.**

Roediger and Karpicke (*Psychological Science*, 2006) add a second empirical handle: retrieval practice — being tested — produces better long-term retention than additional study, and the advantage grows with the retention interval. **Formally this is an `eta` result, not a `delta` result: the same occasion converts more of its content into durable structure when it requires retrieval rather than exposure.** The ritual counterpart is recitation rather than reading, catechism rather than sermon, examination rather than attendance — and the tradition's near-universal preference for the first of each pair is, on this reading, a preference for the higher-`eta` form of the same occasion.

28.9 9 Specification and failure modes for cadence and retention design

28.9.1 9.1 Specification

SC1 — Declare the retention medium. Every stored-order design shall state what the medium of retention is (memory, place, object, document, personnel, habit) and record its `delta` with its period, per S2.1. Media with different `delta` shall be modelled as separate stocks, not averaged.

SC2 — Price a palace by doublings. Palace capacity shall be recorded as `ln L` in the number of loci. Any proposal to expand a palace shall state the number of doublings purchased and the construction and maintenance cost of each; each doubling delivers exactly `ln 2 = 0.6931` nats.

SC3 — Record the three-factor decomposition. Any claim that an encoding technique improves storage shall report the claimed movement in `u`, in `eta`, and in `delta` separately, and shall show their product against the claimed `Q_inf` ratio. A claim that reports only the outcome is not auditable and shall be recorded as unverified.

SC4 — Separate `u_cross` from `u_div`. Where both between-locus and within-locus distinctness are being engineered, both shall be estimated and Theorem 4.1's bound `max(f(u_cross), f(u_div)) <= f(u) <= f(u_cross) + f(u_div)` shall be shown. The additive figure shall never be used as the estimate.

SC5 — State the essence set that `Q` is built against. Stored order is indexed by its category set. No claim of general improvement in a faculty shall be made; the claim shall name the `N` it was built on (Ericsson, Chase and Faloon, 1980).

SC6 — Compute the integral before arguing about cadence. Any cadence proposal shall show that the integrated effect of the alternatives is equal in the unclipped linear model, and shall then identify which of the three non-invariant quantities — peak, minimum, or deliverable total — the proposal is actually trading.

SC7 — Report `Q_min`, never only the mean. Per S2.8. Additionally, the functional threshold `theta` below which the system is observed to misbehave shall be stated, and the percentage of the cycle spent above it shall be computed.

SC8 — Clip every burst at `eta ln N` before comparing totals. Per S2.6. The annual delivered total shall be computed after clipping, and the discarded fraction shall be reported explicitly.

SC9 — Size occasions by the fixed-cost optimum. The per-occasion fixed cost `c_0` shall be measured, the optimum `n` computed against it, and the sensitivity of the optimum to `c_0` shown. Any proposal to increase cadence shall first consider reducing `c_0`.

SC10 — Justify a rare rite by its required `delta`. Any occasion at interval `P > tau` shall compute the retention half-life its trace must have to hold the stated floor, using `a^P >= theta/(theta + A)`, and shall state by what means that retention is expected to be achieved. If no such means can be named, the occasion shall be reclassified as ceremonial and shall not be counted as an input to `Q`.

SC11 — Apply the costly-requirement break-even. No requirement shall be added to a practice without recording its cost `dc`, the current net drive `eta F − c`, and the reduction in `delta` it is expected to produce. The requirement is admissible only if `dc <= (eta F − c)(1 − delta'/delta)`, and the expected `delta'` shall be measured afterwards, not assumed.

SC12 — Deconflict the calendar. Scheduled occasions shall be checked for collision within a single period. Where two occasions fall in one period and their combined input exceeds `eta ln N`, the discarded amount shall be computed and one occasion shall be moved unless a reason for the collision is recorded.

SC13 — Declare open-loop or closed-loop. Every recurring occasion shall be labelled calendrical (open-loop) or state-triggered (closed-loop). Closed-loop occasions shall additionally record the measurement lag `L` and correction gain `k` and satisfy S2.10.

SC14 — Prefer the retrieval form of an occasion. Where an occasion can be run as exposure or as retrieval at equal cost, the retrieval form shall be specified, and the `eta` difference shall

be recorded as an estimate to be tested rather than as an established value.

28.9.2 9.2 Failure modes

THE HEROIC INTENSIVE — a large single occasion substituted for a practice, justified on total input.

- *Parameter damaged:* delivered `S`, clipped at `eta ln N`, while `C` continues unabated.
- *Diagnostic signature:* a strong peak, a crossover against the cadence alternative at roughly `1.5 tau`, and a post-mortem that concludes "we failed to sustain it." In §8.4's parameters the crossover is week 24 against `tau = 16.16`.
- *Measurement:* compute the clipped annual delivered total at each cadence (§8.7 table). Compute `Q_inf = 0` for the unrepeated impulse and `m/delta` for the alternative.
- *Fix:* redistribute to the cadence whose per-period input sits at or just under the ceiling. Retain the intensive only if a threshold has been identified that no sustainable rate reaches — and if so, say which threshold, in the design record.

THE MEAN-ONLY CADENCE ARGUMENT — a sparse schedule defended on the grounds that the average stored order is unchanged.

- *Parameter damaged:* none. The mean really is unchanged. The minimum is not.
- *Diagnostic signature:* the schedule's advocates cite the total or the average; its participants describe "getting back up to speed" each time. In §8.5's table the mean is 8.3333 in every row while `Q_min` falls from 7.8333 to 0.0836.
- *Measurement:* compute `Q_min = A a^P/(1−a^P)` and the percentage of the cycle above the stated functional threshold.
- *Fix:* shorten `P`, or state the threshold and demonstrate that the trough clears it. **Averages do not fail systems; minima do.**

THE UNPAID RARE RITE — an occasion at an interval far longer than `tau`, treated as an input to stored order, whose trace cannot survive to the next occurrence.

- *Parameter damaged:* effectively `eta`, since the converted order is gone before it can compound.
- *Diagnostic signature:* the occasion is memorable and well-regarded, and the system's behaviour between occurrences is indistinguishable from a system that never held it. In §8.12's simulation an annual-only schedule holds `Q > 0` for four weeks of fifty-two.
- *Measurement:* compute the required `delta` from `a^P >= theta/(theta+A)` and compare against the measured `delta`. In the worked example the shortfall factor is 2.5.
- *Fix:* either supply the retention the interval demands — an artefact, a record, a durable external store, per §8.2's `delta` column — or shorten the interval, or reclassify the occasion honestly as ceremony rather than as input.

THE COLLIDING CALENDAR — two substantial occasions scheduled into the same period, so that the second is largely discarded at the ceiling.

- *Parameter damaged:* delivered `S` in the colliding period.

- *Diagnostic signature:* a heavy week followed by a light one; participants report that the second event "did not land." The discarded fraction is computable and is often above a third.
- *Measurement:* sum the intended `eta F` of all occasions falling in one period and compare against `eta ln N`. §8.12 computes 36.1% and 41.1% discard for two realistic collisions.
- *Fix:* move one occasion to a clean period. This costs nothing and is the highest return-per-effort intervention in this section.

THE REQUIREMENT RATCHET — costly requirements added past the break-even, because earlier requirements demonstrably helped.

- *Parameter damaged:* `C`, without a compensating fall in `delta`.
- *Diagnostic signature:* membership requirements have grown monotonically for years; retention among those who remain is excellent; recruitment and total membership are falling. The institution interprets falling numbers as evidence that standards are working.
- *Measurement:* apply `dc <= (eta F − c)(1 − delta'/delta)` to each requirement in force. Measure `delta` before and after each addition rather than assuming the direction.
- *Fix:* remove requirements from the bottom of the value ordering until the inequality holds, and measure. Note the asymmetry: **an institution that will not choose what it sheds does not thereby keep it.**

THE PALACE THAT WAS NEVER MAINTAINED — a memory system built once, whose loci themselves have decayed, so that retrieval fails for reasons the operator attributes to the contents.

- *Parameter damaged:* `delta` of the index rather than of the material.
- *Diagnostic signature:* recall fails specifically at the transitions between loci; the operator can produce isolated items but not the ordered walk. Diagnostic separation from ordinary forgetting is that prompting with the *place* restores the item.
- *Measurement:* test the empty palace — walk it with no contents and count loci recovered in order. Compare against the count when it was built.
- *Fix:* rebuild the index before re-loading the contents. Loading new material onto a decayed index raises `C` and lowers `eta` simultaneously and is the worst available move.

THE UNDECLARED ESSENCE SET — improvement measured on one category set, claimed for a faculty.

- *Parameter damaged:* none; the claim is simply not about the quantity measured.
- *Diagnostic signature:* a practice that produced a large, real gain in one domain produces nothing in a second, and the operator concludes the second domain is harder rather than that the stock is domain-indexed.
- *Measurement:* run the same protocol on a disjoint essence set. Ericsson, Chase and Faloon's subject reached a digit span near 79 while his letter span stayed near six.
- *Fix:* declare the `N` in every claim. Build a second stock deliberately if a second domain is wanted; there is no transfer to appeal to.

28.10 10 Worked example: a chartered study circle

The system. Six people from six disciplines convene as a chartered study circle around a canonical corpus of 48 items. There are two coupled subsystems with different natural periods, and they must be modelled separately per SC1: each member's individual hold on the corpus (period: one day) and the circle's own stored order (period: one week).

The question on the table is whether to (a) adopt loci encoding for the corpus, (b) replace the current monthly three-hour seminar with a weekly ninety-minute session, and (c) keep or move the annual retreat.

All parameter values below are design values, badged as such, obtained by the framework's survey protocols in a hypothetical engagement and stipulated here.

Step 1 — Subsystem M: one member's hold on the corpus. Period: one day.

```
N = 48 items          ln N = 3.8712 nats     <- hard per-occasion ceiling
```

	u	N_eff	F	% ceil	eta	delta	eta F	Q_inf	gain	t_half	t_95
before: reread only	0.20	10.400	2.3418	60.5%	0.18	0.22	0.4215	1.9160	0.818	2.79 d	12.06 d
after: 24-locus palace, two images each	0.88	42.360	3.7462	96.8%	0.36	0.05	1.3486	26.9727	7.200	13.51 d	58.40 d

```
Q_inf ratio = 14.0774  =  F 1.5997  x  eta 2.0000  x  1/delta 4.4000
```

Per SC3 the three factors are reported separately, and per SC4 the split is shown. Estimating `u_cross = 0.70` (distinct places) and `u_div = 0.18` (striking images):

```
f(u_cross) = 3.5234     f(u_div) = 2.2471     f(u) = 3.7462
bound:  3.5234  <=  3.7462  <=  5.7705
```

Place alone accounts for 94.1% of the combined surplus. The images are worth having and they are worth 0.2228 nats, not the 2.2471 that the additive reading would suggest. **Per SC4 the additive figure is not used, and a design that spent equally on images and places would be spending most of its effort on the smaller term.**

Per SC5, the claim on record is: *ordered recall of this 48-item corpus improves.* Not memory.

Step 2 — Subsystem G: the circle. Period: one week.

The circle is six people, but the category set that matters is the set of *domains* they bring, and with two recurring outside guests the category audit returns eight [Curzi, 2026c]:

```
N = 8 domains                          ln N     = 2.0794
u* = 0.55 (audited)                    N_eff    = 1 + 7(0.55) = 4.8500
F = ln 4.85                                     = 1.5790      (75.9% of ceiling)
eta = 0.35, delta = 0.055/wk, C = 0.13/wk
ceiling eta ln N                                = 0.7278/wk
```

Rather than assume the session geometry, the cadence question is put to Theorem 2 directly. The circle has a genuine interaction budget of **three hours per fortnight** and a per-occasion cost of `c_0 = 0.045` (convening, travel, settling). Non-parallelism saturates at `L* = 2.5` hours.

schedule	occasions per fortnight	length	`u`	`f(u)`	`eta f(u)`	over ceiling?	delivered per fortnight	vs one long session
one 3.0 h session	1	3.000 h	1.0000	2.0794	0.7278	at the ceiling	0.7278	1.0000
two 1.5 h sessions	2	1.500 h	0.6000	1.6487	0.5770	no	1.1541	1.5857
three 1.0 h sessions	3	1.000 h	0.4000	1.3350	0.4673	no	1.4018	1.9260
four 45 min sessions	4	0.750 h	0.3000	1.1314	0.3960	no	1.5840	2.1764
six 30 min sessions	6	0.500 h	0.2000	0.8755	0.3064	no	1.8385	2.5261

Note that the single three-hour session reaches `u = 1` and therefore sits **exactly** at the ceiling, `eta ln N = 0.7278`: it delivers the maximum a single period can carry, and the half-hour of that session beyond `L* = 2.5` delivers nothing at all. Two fortnightly sessions deliver 1.5857 times as much; six deliver 2.5261 times as much. **The delivered column is monotone in `n`, exactly as Theorem 2 requires.**

The circle adopts the weekly ninety-minute form — `n = 2` per fortnight — on the grounds that the higher-`n` options push per-occasion cost above what six busy people will actually pay, and because the ninety-minute session leaves visible headroom under the ceiling for a guest or a longer thread. Session geometry at 1.5 hours would give `u = 0.60` and `eta f = 0.5770`; the conservative figure carried forward is the *audited* `u* = 0.55` from the disagreement survey, since that is what the circle actually achieves across a year rather than in a well-run session:

```
weekly input   eta F = 0.35 * 1.5790          = 0.5526/wk
net drive = 0.5526 - 0.13 = 0.4226/wk
Q_inf     = 0.4226 / 0.055 = 7.6844
pole      = 0.945    tau = 17.677 wk    t_half = 12.253 wk
t_95      = 52.96 wk    t_99 = 81.41 wk
h at Q_inf = 0.7278 - 0.055(7.6844) - 0.13 = 0.1752/wk       (r_max = h)
h = 0 at Q = 10.8692
```

The first honest review is week 53. A week-13 reading of 4.4040 is 51.6% of steady state and is not a finding about the change.

Step 3 — The calendar, deconflicted. Per SC12.

Three occasions are proposed: the weekly session (`u = 0.55`, `eta F = 0.5526`), a monthly seminar with an outside guest (`u = 0.62`, `eta F = 0.5863`), and an annual retreat (`u = 0.86`, `eta F = 0.6821`). The per-period ceiling is 0.7278.

```
weekly + seminar  = 1.1390    ceiling 0.7278    DISCARDED 0.4112   (36.1%)
weekly + retreat  = 1.2347    ceiling 0.7278    DISCARDED 0.5069   (41.1%)
```

The retreat is initially scheduled for the last week of the year, which is also a seminar week. Three years simulated, with cost paid every week:

schedule	steady mean `Q`	min	max	ripple
weekly only	7.6844	7.6844	7.6844	0.0%

schedule	steady mean `Q`	min	max	ripple
weekly + monthly seminar + retreat, **retreat colliding with a seminar**	8.4806	—	—	—
weekly + monthly seminar + retreat, **retreat in a clean week**	8.5418	8.4259	8.7145	3.38%

Trajectory of the built calendar:

week	1	4	13	26	53	105	157	209
Q_t	0.4226	1.7313	4.4040	6.6679	8.1134	8.5193	8.5408	8.5419

Moving the retreat off the seminar week costs nothing and returns 0.0612 of steady-state stock. The full calendar returns `+0.8574` over weekly-only, an 11.16% gain, and headroom at the steady mean is `0.7278 − 0.055(8.5418) − 0.13 = 0.1280`, still positive. Headroom reaches zero at `Q = 10.8692`, comfortably above the operating point.

Step 4 — Is the annual retreat an input at all? Per SC10.

Standing alone, at `P = 52`, `A = 0.6821`, and a functional floor `theta = 0.25`:

```
required a^P  >=  theta/(theta + A) = 0.25/0.9321 = 0.2682
required a    >=  0.977559        required delta <= 0.022441/wk
required t_half >= 30.54 weeks = 0.59 years
```

The circle's measured `delta` is 0.055 per week, a half-life of 0.24 years. **The shortfall factor is 2.5.** Formally: *the retreat cannot, by itself, hold a floor of 0.25 for a year at this circle's retention.* Simulated alone, an annual-only schedule holds `Q > 0` for four weeks in fifty-two.

Per SC10 the retreat is therefore not recorded as an input to `Q` on its own account. It is recorded as (a) a `+0.0612` contribution on top of the weekly base, and (b) a candidate `delta` intervention — the retreat is where the circle's charter is re-read and its obligations restated, which is a retention measure, not an input measure, and it is to be measured as one.

Step 5 — The proposed attendance requirement. Per SC11.

A member proposes a hard attendance requirement: miss four consecutive weekly sessions and lapse. Estimated additional cost `dc = 0.05` per week in administration, chasing, and the friction of enforcement.

```
net drive            = eta F − C = 0.5526 − 0.13 = 0.4226/wk
required reduction   = dc / net drive = 0.05/0.4226 = 11.83%
required delta'      <= 0.055 ∗ (1 − 0.1183) = 0.048493/wk
```

The requirement pays for itself if and only if it reduces weekly decay by at least 11.8% — in membership terms, if it raises the half-life of continuous participation from 12.25 weeks to at least 13.90. That is a specific, measurable, falsifiable claim, and it is the claim the circle records before adopting the rule, not after.

Per SC11 the outcome is to be measured. Per the requirement-ratchet failure mode, this is the *first* requirement and the break-even is generous; the test will get harder for the second and harder still for the third, and the circle records that fact in the same minute.

Step 6 — What would falsify all of this.

Fit `Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + eps_t` on weeks 1-53 with `0 <= a <= 1, b >= 0`, freeze the coefficients, and report on weeks 54-106 against four baselines: persistence, direct-input, matched autoregressive, and shuffled-input control. Predictions on record:

1. The fitted `a` lies in `[0,1]` and its interval contains `1 − delta = 0.9450`. If not, the first-order form is wrong for this circle and every number above is void.
2. The fitted `b` exceeds its pre-intervention value. If not, the cadence change did not raise delivered input, whatever the participants report.
3. The frozen model beats the **shuffled-input control** on held-out data. **If it does not, then the timing of the circle's occasions carries no information beyond their volume, the entire cadence argument does not apply to this system, and Steps 2 through 5 should be withdrawn.**
4. On Subsystem M, the empty-palace test (per the palace failure mode) recovers at least 20 of 24 loci in order at 90 days. If the index itself decays faster than the contents, the `delta` attribution in Step 1 is misassigned.
5. The attendance requirement reduces measured membership `delta` by at least 11.8% within two `tau`. If it does not, it is removed — not retained on the grounds that standards matter.

Prediction 3 is the decisive one. Everything this section claims about cadence is a claim that *when* input arrives matters and not only how much. The shuffled-input control is the test that can kill it, and it is the control most often omitted.

28.10.1 10.1 What the worked example establishes

The method of loci is the storage law with space as the retention medium. Its advantage factorises into three separately measurable terms — surplus, conversion, and retention — and in the worked case retention supplies more than the other two combined, which is the opposite of the usual emphasis. Palace capacity is `ln L`, so **every doubling of the palace buys exactly 0.693147 nats and no arrangement of the loci changes that.** The classical rules of the *ad Herennium* map one-to-one onto the parameters of the law with no rule left over and no parameter left untouched.

Ritual is periodic forcing of a low-pass filter, and the case for cadence over intensity has two independent halves. **The first is the ceiling: surplus saturates at `ln N` per period while cost does not, so an annual rite delivers 2.8% of what a weekly practice delivers from the same nominal budget.** The second is Theorem 2: because `f(u)` is strictly concave through the origin, dividing a fixed interaction budget into more occasions never delivers less and in the worked case delivers 4.57 times as much — a result that follows from the surplus functional's strict concavity [Curzi, 2026] and needs no ceiling at all. What stops the division is the per-occasion fixed cost, and the optimum is computable from it: **lowering the cost of convening moves the optimum more cheaply than any intervention on the occasions themselves.**

The honest comparison between one impulse and the same total delivered periodically was computed and it is not one-sided. Integrated effect is identical to every digit — 433.3333 in all three regimes — because the system is linear and time-invariant. Peak favours the impulse by a factor of 3.25. Minimum, time above a low threshold, and clipped deliverable

total favour the cadence, the last by a factor of 36. **And the steady state settles it: a single impulse, however large, has a steady state of exactly zero, while a rate `m` has `m/delta`.**

The rare, intense rite is not thereby refuted. It occupies the regime above the cadence's plateau, where a threshold exists that no sustainable rate reaches — and the price of that regime is computable. **A once-a-decade rite requires its trace to have a half-life of about five and a half years, and a once-a-generation rite about seventeen.** That is why such rites are extraordinary, why they are small, and why they take the forms they do. It is also the boundary of what this paper will say about them: the arithmetic explains the form and does not recommend it.

Section 11 states the thesis that has been implicit throughout – that the inherited notation and the framework's notation describe the same relationships – under the strictest available constraint: **claim no supernatural efficacy, no physical effect of symbol, and no result that does not follow from the formalism.** The argument is stronger for the restriction.

28.11 11 The synthesis: three handles and an eleven-operation audit

The framework opens with a difference equation and the claim that its leak is not a defect [Curzi, 2026a]. Everything the corpus has built since is that equation at a different scope, with different parameters, under a different name.

This section states what the assembly amounts to, under the strictest available constraint. **No supernatural efficacy is claimed, no physical effect of symbol is claimed, and no result appears that is not derivable from the formalism.** The argument's force comes entirely from that restriction. A synthesis that permitted itself one unearned step would be worth nothing, because the interesting question is not whether the inherited notation *can* be dressed in new symbols – anything can – but whether the modern notation is **sufficient**: whether, having described the mechanism, there is a residue that still requires something else.

The finding, stated plainly and then defended for the rest of the section, is that there is not — at least not in the audit performed here, over the operations this paper has examined. **That is a claim about an audit and not a proof of completeness, and it is badged as such wherever it appears.**

The audit is bounded, and the boundary is stated in §11.19 and again in Section 17.

28.11.1 11.1 The thesis

THESIS (stated as a postulate and badged as one). If *magic* is understood as

1. the deliberate engineering of the perceptual lens — the ratio `D/P` and the moment `M = (D/P)E`;
2. the storage of order through symbol, space and cadence — the parameters `eta`, `delta`, and the stock `Q`;
3. the composition of non-parallel essences — the geometry `u` and the surplus `F = ln(1 + (N−1)u)`;

then the framework has described its mechanism in modern notation.

That is the whole claim. Note carefully what it contains and what it does not.

It contains a definition. The word *magic* is being given a stipulative reading — three engineering operations — and the thesis says that under that reading, the mechanism is the one this manual has specified. **A reader who rejects the definition is not thereby contradicting the thesis; they are declining it, which is a different act, and this paper has no argument against them.**

It does not contain an empirical claim about the historical practice. Whether the practitioners of any tradition understood themselves to be doing these three things is a question for historians, and this paper does not answer it.

It does not contain an efficacy claim of any kind. The thesis says the mechanism *is described*. It does not say the described mechanism produces effects beyond those the storage law and the surplus law produce, because they are the same mechanism. **There is nothing left over for an effect to come from.**

28.11.2 11.2 What "described its mechanism" means

A mechanism is described when three things hold, and all three are checkable:

```
(i)   every operation in the practice maps to a named parameter of the formalism
(ii)  the direction and, where measurable, the magnitude of each operation's effect
      follows from the law rather than being stipulated alongside it
(iii) the failures of the practice are the failures the law predicts, and appear at
      the parameter values where the law says they will appear
```

§11.18 performs (i) as an explicit audit over eleven inherited operations. §11.6 through §11.16 perform (ii) result-family by result-family. The failure-mode catalogues of Sections 5, 9 and 12, and the corresponding catalogues in the corpus's other papers, are (iii).

What would refute the thesis is a fourth thing: an operation that does work in the practice, is agreed to do work, and maps to no parameter. §11.19 records what was found when one was looked for. **That is the honest form of the claim and it is falsifiable in exactly one direction: someone may exhibit the missing operation.**

28.11.3 11.3 The ledger

The most useful single element of this section is an inventory of everything the framework has asserted, sorted by status. **A synthesis that does not publish this table is asking to be believed rather than checked.**

PROVED — mathematics, true regardless of whether anything in the world obeys it:

Result	Source
`f(u) = ln(1+(N−1)u)` uniquely characterised by four axioms	Curzi [2026]
Strict monotonicity, positivity, strict concavity of `f`	Curzi [2026]
Subadditivity on the parameter `u` (**not** on geometric triples)	Curzi [2026]
Sharp Lipschitz constant `N−1`	Curzi [2026c]
`F = 0` iff parallel; `F = ln N` iff orthogonal	Curzi [2026]
Two-source decomposition and its bracket	Curzi [2026]
`0 <= F(u) <= ln N` — the hard ceiling	Curzi [2026c]
Exact solution of the storage recursion	Curzi [2026a]
One-step growth condition `eta S_t > delta Q_t + C_t`	Curzi [2026a]
Uniform bound `Q_t <= max{Q_0, eta ln N/delta}`	Curzi [2026a]
Sustainability ceiling `C_t <= eta ln N − delta Q_t`	Curzi [2026a]

Result	Source
Steady state `Q_inf = (eta F(u*) − c*)/delta`	Curzi [2026a]
Max one-step increment `r_max = eta ln N − delta Q_t − C_t`	Curzi [2026a]
The recursion is a first-order low-pass filter with pole `1−delta`	Curzi [2026a]
Contraction implies a unique fixed point reached from any start	Banach [1922]; Curzi [2026]
`f(n) = k n^2` is the exact collapse of the published essence recipe	Curzi [2026]
Essence intervals are ratios of squares; `k` cancels	Curzi [2026]
An essence set `1..N` has exactly `ceil(N/2)` pitch classes	Curzi [2026]
One nat of surplus is exactly `2/ln 2 = 2.885390` octaves	Curzi [2026]
A star polygon `{N/k}` has exactly `gcd(N,k)` closed strokes of length `N/gcd(N,k)`	this paper, Theorem 1
Dividing a fixed budget over more occasions never delivers less, for concave `g`, `g(0)=0`	this paper, Theorem 2
Costly-requirement break-even `dc <= (eta F − c)(1 − delta'/delta)`	this paper, §8.11
Expected surplus of a uniform random correspondence `= (N ln N − N + 1)/(N−1)`	this paper, §11.18
Expected colliding pairs under a hash `>= C(M,2)/R`, minimised by uniformity	this paper, §11.20

POSTULATED — badged every time, chosen rather than derived:

Postulate	Where
The affine effective count `N_eff = 1 + (N−1)u` — a *design* axiom, and **not** a Hill number	Curzi [2026o]
That a room, team or institution is described by the storage recursion at all	Curzi [2026a]
Scale symmetry: the same model form at every scope	Curzi [2026]
That essence types have a frequency; that the law is square in the index	Curzi [2026]
That `k = (9/8)^2/phi^2` rather than any other constant — a **defined** ratio, exact	Curzi [2026]; Markowsky [1992]
`M = (D/P)E` as the form of the moment	Curzi [2026o]
The internal/external correspondence as one subject at two scopes	Curzi [2026o]
The figure read as a nonlinear computing map — **a metaphor**	this paper, §3.4
That the method of loci raises `u` and `eta` and lowers `delta`	this paper, §7.1
The thesis of §11.1	this paper

REQUIRES MEASUREMENT — asserted nowhere, obtainable by protocol:

Every `delta`, every `eta`, every `u`, every `C`, every `N`, every `D` and `P`, in every system, for every period. **This paper supplies no value of any of these for any real system.** The corpus supplies the protocols: a survey instrument for each of `delta`, `eta`, `u`, `C` and `N`, and a formula-validation protocol for `D` and `P` [Curzi, 2026c,u].

Counted: 24 proved results, 10 postulates, and a parameter set that is entirely empirical. The ratio is the point. **A framework whose postulates outnumber its theorems is a worldview; one whose theorems outnumber its postulates by more than two to one, with every postulate labelled, is an engineering discipline with a declared foundation.**

28.11.4 11.4 The refusals are load-bearing

The framework carries four standing refusals, honoured throughout. They are not modesty. **They are the conditions under which the thesis of §11.1 means anything at all.**

1. **No claim that the mathematics proves anything about physics, thermodynamics, or consciousness.** The framework proves no negentropy law; `F(u) + N_can(u) <= ln N` is a bound and not a mechanism [Curzi, 2026c].
2. **No postulate without its badge.** §11.3's second table is the badge, collected.
3. **No number that was not computed or cited.** Every figure in the corpus was computed; where a parameter is illustrative it says so in the same sentence.
4. **No intervention that removes a person's ability to decline.** This is the refusal that governs the ordeal arithmetic of §8.10, contagion engineering, and every use of the word *design* applied to other people.

If any of the four were suspended, the thesis would collapse into an unfalsifiable analogy. With all four in force, it is a specific, checkable, and modest claim: *this is the mechanism; there is nothing else in it; here is the arithmetic.*

28.11.5 11.5 Three scopes, one instrument

The corpus applies one recursion at three scopes, and the movement between them is a change in who holds the instrument rather than a change of subject.

The observational scope. One person, reading a situation: what a comfort costs, what a discomfort quietly pays, whether a projected trajectory keeps its assumptions constant along its length. The instrument is attention, and it is what produces the numbers the other two scopes consume.

The constructive scope. The same relations made constructible: the storage recursion and its four parameters [Curzi, 2026a], the geometry of difference [Curzi, 2026,c], the physical setting, the sociology of groups and institutions, the measurement and falsification tradecraft [Curzi, 2026u], and the essence law that gives a relationship an index that can be rendered [Curzi, 2026,?]. Its subject is the world outside one person, being built.

The reflexive scope. The same six parameters applied to the person applying them. The claim is not that psychology *resembles* sociology but that **the two are one set of relations measured at two scopes** – a badged thesis, not a demonstrated identity [Curzi, 2026o]. The perceptual lens is `D/P`; character is `Q`; forgetting is `delta`; personal capacity for the new is `h`; and the internal surplus `F_int` is the surplus functional applied over one observer's own framings, which is a derivation rather than an analogy [Curzi, 2026o].

This is exactly where the reflection problem enters, and §11.19 does not evade it. The framework's reflexive loop is written `T(x) = W(x, M(x))`, with the self-fulfilling prophecy identified as a fixed point of that composition of seeing and responding [Curzi, 2026]. **A specification in which the members constitute the aggregate and the aggregate influences the members is not identified from observational data** [Manski, 1993, Shalizi and Thomas, 2011]. Every

statement in this section about a parameter as a property of a person or a group is a modelling position and not an identified empirical result.

```
OBSERVATIONAL   one person, reading a situation      instrument: attention
                        |
                        |   the same relations, made constructible
                        v
CONSTRUCTIVE    the setting outside, being built     instrument: specification
                        |
                        |   the same instrument, pointed back
                        v
REFLEXIVE       the builder, being measured          instrument: the parameter set
```

The third scope closes the loop, and it closes it uncomfortably. An analyst who has spent the constructive scope measuring other people's `u`, budgeting other people's `h`, and designing cadences for other people's calendars arrives at the reflexive scope to find the same instruments turned around: their own `N` is the count of things they can actually distinguish; their own `delta` is how fast they lose what they learn; their own `D/P` decides whether they are seeing a system or their model of one. **The framework's diagnostics do not exempt the person holding them, and the over-framed institution and the over-framed observer are the same fault at two scopes** [Curzi, 2026o]. This is also the framework's own version of the Lucas critique [Lucas, 1976] and of the looping effect of human kinds [Hacking, 1995]: parameters estimated in a population that has not read the framework cannot be assumed stable in one that has.

One consequence governs how the three scopes should be read together. **The observational instrument is never superseded.** Attention is what produces the survey data, the disagreement audit, and the magnitude estimates that every later computation consumes. An analyst who takes the specifications without the reading practice has a set of formulas and no measurements to put in them, and the formulas will accept any numbers offered. **The sequence is not a ladder on which the earlier rungs are discarded; it is three instruments that must all be working at once, and the framework's most common misuse is running the arithmetic on estimates nobody actually looked at.**

Two correspondences are exact rather than decorative. Recursion is the transient `(1-delta)^t Q_0` – the initial state bleeding off while the driven term builds – and the fixed point is the storage law's steady state [Curzi, 2026a], which is unique and start-independent because the update map is a contraction with modulus `1 - delta` [Curzi, 2026, Banach, 1922]. The contraction ratio is `(1-delta)^t`:

delta	t=1	t=5	t=10	t=20	t=50	t=100
0.020	0.9800	0.9039	0.8171	0.6676	0.3642	0.1326
0.055	0.9450	0.7536	0.5680	0.3226	0.0591	0.0035
0.100	0.9000	0.5905	0.3487	0.1216	0.0052	0.0000
0.200	0.8000	0.3277	0.1074	0.0115	0.0000	0.0000

A contraction does not arrive somewhere new. It arrives where it was always going to arrive, from wherever it started, and repetition cannot move the destination because the fixed point contains no `t` [Curzi, 2026]. Recognition rather than novelty is therefore the correct description of arrival at a fixed point, and the contraction theorem is why [Banach, 1922].

28.11.6 11.6 The same equation, six scopes

The architectural claim under test is that only `(N, u, eta, delta, C)` and the period change between a desk and a city. That claim is about **model form**, and it is honest because it is checkable: if a scope required a different equation, the claim would fail there.

Computed across seven scopes, with illustrative parameters badged as such:

scope	period	N	u*	eta	delta	C	ln N	N_eff	F	eta ln N	Q_inf	tau	t_95	h
a desk	day	4	0.45	0.30	0.100	0.02	1.3863	2.350	0.8544	0.4159	2.3632	9.49	28.43	0.1596
one mind (a practice)	day	12	0.55	0.30	0.050	0.06	2.4849	7.050	1.9530	0.7455	10.5182	19.50	58.40	0.1596
a dyad	week	6	0.62	0.40	0.040	0.08	1.7918	4.100	1.4110	0.7167	12.1099	24.50	73.39	0.1523
a room / a team	week	8	0.55	0.35	0.055	0.13	2.0794	4.850	1.5790	0.7278	7.6844	17.68	52.96	0.1752
a department	month	14	0.48	0.30	0.050	0.22	2.6391	7.240	1.9796	0.7917	7.4777	19.50	58.40	0.1978
an institution	quarter	25	0.40	0.22	0.030	0.30	3.2189	10.600	2.3609	0.7082	7.3129	32.83	98.35	0.1888
a city	year	60	0.35	0.15	0.020	0.18	4.0943	21.650	3.0750	0.6142	14.0625	49.50	148.28	0.1529

Every row is the same recursion. Read down the `t_95` column: the desk settles in a month of days, the team in a year of weeks, the city in a century and a half of years. **That column is the reason scales cannot be managed with a single calendar, and it is the whole content of the observation that institutions are slow.** They are not slow because of their people. They are slow because `t_95 = ln(0.05)/ln(1−delta)` and their `delta` is small — which is also the only reason they persist.

Read down the `h` column and a second fact appears: **headroom is roughly comparable across scopes even though `Q_inf` differs by a factor of six.** Headroom is a *rate*, and rates are what a scope can afford per period; stocks are what it has accumulated. Confusing the two is the framework's central warning about headroom and it recurs at every scale [Curzi, 2026].

28.11.7 11.7 The circuit results

The corpus establishes that the recursion is an accumulator rather than a regress, that `delta > 0` makes it a contraction and therefore guarantees arrival, and that the identity with a first-order low-pass filter is a **shared mathematical form** rather than a physical analogy: the same difference equation governs an RC network, an exponential smoother and a stock of order, with no claim that the electrical physics applies [Curzi, 2026a, Elaydi, 2005, Oppenheim et al., 1996, Holt, 1957]. From that came the pole, the time constant, the settling times, the DC gain `eta/delta`, and the frequency response that makes stored order blind to fast fluctuation and sensitive to slow trend.

What this contributes to everything after it is a single gauge: `h = eta ln N − delta Q − C`. Below zero, no arrangement of people saves the design; at zero, hold; above zero, grow by at most `h`. **Every recommendation the framework makes is checked against `h` first**, because if the cost term exceeds `eta ln N − delta Q`, no interaction geometry whatsoever prevents order from decreasing [Curzi, 2026a]. That is the framework's strongest sociological statement and it is a theorem about a difference equation, not an empirical finding about institutions.

28.11.8 11.8 The geometry of difference

The corpus establishes that difference is counted before it is valued, and that only what is different is counted at all. `N_eff = 1 + (N−1)u` audits a headcount for redundancy; `F = ln N_eff` is the log of what survives the audit. `F = 0` for parallel participants – **agreement between sources that share training, incentives and premises is worth `ln 1 = 0`, and that is a theorem rather than an opinion** [Curzi, 2026]. **`N_eff = 1 + (N−1)u` is the framework's own construction and is NOT a Hill number**; `1/HHI` and `exp(H)` are Hill numbers of orders two and one and are the grounded ones [Hill, 1973, Jost, 2006, Simpson, 1949], and the diversification correction that says a benefit depends on covariance rather than count was stated in 1952 [Markowitz, 1952].

Two consequences ran through everything after. First, **the ceiling `ln N` is hard**, so a period's input is bounded no matter how the period is run — which is why the cadence result

works and why the calendar arithmetic of Section 8 bites. Second, **concavity means the second genuinely different framing is the largest gain available and the fortieth is a maintenance bill**, which is the internal-surplus argument [Curzi, 2026o] and reappears in §11.18 below as the formal content of the practice of adopting a second perspective deliberately.

28.11.9 11.9 Engineering the physical setting

The corpus takes the room as a circuit: `delta` for a room is how fast its arrangement reverts, `eta` is how much of what happens in it becomes furniture, `C` is what it costs to keep. Reverberation time `RT60 = 0.161 V/A` is measured acoustics [Sabine, 1922, Schroeder, 1965] and enters as a constraint on whether the room can carry interaction at all; proxemics sets the distances at which particular kinds of interaction are possible [Hall, 1966]; and the semilattice-versus-tree distinction [Alexander, 1965, Alexander et al., 1977] is the observation that overlapping structure beats strict hierarchy, which in this notation is the observation that a strict tree destroys `u` by construction.

The contribution is that `eta` is often the cheapest parameter to move, and that it is usually moved by furniture rather than by exhortation. A room that makes the good arrangement the default converts more interaction into structure than a room that requires everyone to remember.

28.11.10 11.10 The empirical layer the formalism does not contain

Woolley et al. [2010] report a collective-intelligence factor not predicted by members' average or maximum ability but correlated with social sensitivity and **equality of conversational turn-taking**, which is an effective-count measurement by another route. Centola et al. [2018] locate a committed-minority tipping point near 25 per cent. Complex contagion needs reinforcing contacts and clustered topology [Centola and Macy, 2007, Centola, 2010]. Ostrom [1990]'s design principles read here as a specification for keeping institutional `delta` low without raising `C` past the break-even, and Cox et al. [2010] is the meta-analytic audit those principles have already faced. The contrary evidence is stocked deliberately: diversity in the short run has been found to reduce trust and civic participation [Putnam, 2007], with a rebuttal [Abascal and Baldassarri, 2015], and cohesion has been found to raise then lower creative success [Uzzi and Spiro, 2005], which is an inverted `U` against a surplus functional that is monotone increasing in `u`.

The warning is that the formalism is not self-sufficient. `u`, `eta` and `delta` are empirical quantities in real groups, and the literature that measures them is a science with a replication problem [Open Science Collaboration, 2015, Ioannidis, 2005], not a source of settled constants.

28.11.11 11.11 The measurement and falsification tradecraft

The procedure is: survey the system, preregister the prediction, intervene, and read the failure catalogue when it does not work [Nosek et al., 2018, Simmons et al., 2011, Gelman and Loken, 2013]. **The load-bearing element is the shuffled-input control**, which appears in every falsification suite the framework specifies and is the control most often omitted elsewhere. If a model predicts a system's stored order as well from shuffled inputs as from real ones, then timing carries no information and the entire cadence apparatus does not apply to that system.

This is what separates the framework from a worldview. A framework that cannot fail cannot be trusted, and the corpus names exactly how this one fails [Curzi, 2026u]. **It also measures the control against itself and finds it wanting: against a nominal 0.05 the shuffled-input control false-rejects true nulls at 0.108 for independent permutation and 0.120 for circular shift, with only block permutation near nominal at 0.060** [Curzi, 2026u]. Every shuffle p-value quoted anywhere in this paper's worked examples inherits that calibration defect, and Section 17 says so again.

The tradecraft also supplies the framework's characteristic move, worth naming as a method rather than leaving as a habit. **Where a phenomenon is described in the language of morale, culture, energy or will, the arithmetic is computed first and the people are asked about second.** A premature review is a pole, not a motivation problem. An ossified standards body has a fallen ceiling and an unchanged maintenance load. An offsite whose effects "wore off" did not wear off; it decayed at `(1-delta)^t` while `C` continued. **The arithmetic is not a substitute for judgement about people; it is what has to be cleared before judgement about people is worth having**, because in each of those cases the arithmetic alone accounts for the observation. That is a methodological preference and not a result, and it is exactly the move Hedström and Ylikoski [2010] demands be justified by naming a mechanism rather than exhibiting a fitted curve.

28.11.12 11.12 The essence law

The corpus establishes that the published recipe collapses exactly to `f(n) = k n^2` with `k = (9/8)^2/phi^2 = 0.483425732988414`, **a DEFINED constant – an exact ratio of exact quantities, computable to arbitrary precision, which is the only reason fifteen figures are legitimate. Attaching the unit Hz to it is a separate and much stronger claim, and it is POSTULATED rather than derived** [Curzi, 2026, Markowsky, 1992]; that every interval in the system is a ratio of squared integers so that `k` cancels from all relationships, and that an essence set indexed `1..N` yields exactly `ceil(N/2)` pitch classes because pitch class depends only on the odd part of the index. The bridge to the surplus law is exact: `N_eff = e^F` is the cross-disciplinary essence index, so `f = k e^{2F}` and **one nat of surplus is exactly 2.885390 octaves.**

The honesty clause is absolute and is restated here because this is the material most likely to be quoted out of context: the output is an essence index, not a physical resonance. Nothing — no material, molecule, element, room, or person — is claimed to resonate at any value the essence law produces. `k` is a chosen constant. The square law is a postulate. That `E_n \ensuremath{\propto} n^2` is also the eigenvalue spacing of a one-dimensional infinite square well is a **shared mathematical form** and is emphatically **not** evidence of a physical mechanism. `phi^2 = phi + 1` is exact arithmetic; the literature on how often the golden section is claimed where it is not present is the reason that distinction has to be made explicitly [Markowsky, 1992, Euclid, 1956, Vajda, 1989].

28.11.13 11.13 The perceptual lens

The register everything else is written in is logarithmic. The constant just-noticeable ratio [Weber, 1834] and its integration into `S = k ln(I/I_0)` [Fechner, 1860] are nineteenth-century findings that have been examined continuously since. Magnitude estimation is offered as producing ratio-scale measurements of subjective quantity [Stevens, 1957, 1946], which is the existence proof the reflexive scope needs – **but the claim is contested from three directions and the corpus stocks all three.** The fitted exponent depends on the experimenter's chosen stimulus range [Teghtsoonian, 1971]; the procedure's biases were catalogued by a critic [Poulton, 1968, 1979]; and the deepest objection is that psychology assumes quantitative structure without testing for it [Michell, 1997, 1999]. **The test that would settle it is conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971], and it has not been run for any quantity in this corpus.**

From that comes the Lens Theorem: **if perception is logarithmic, addition in perception is multiplication in the world.** Equal perceptual steps are equal world ratios, never equal world differences. A "moderate" perceived improvement is an order of magnitude produced. And the surplus functional's strict concavity has the same shape, so **diminishing returns is not pessimism; it is what a logarithm looks like from inside.** A second and independent warrant for the logarithm is available and does not require the perceptual story to be true at all: a logarithmic form falls out of multiplicative dynamics with no psychological assumption [Peters and Gell-Mann, 2016, Peters, 2019], and that is the warrant the argument here actually needs.

The moment equation `M = (D/P)E` supplies the diagnostic. `r = D/P` with distortion `|ln r|`: over-framed above one, under-framed below, and **the fixed point `D = P` is the formula's most important content** — frame matching intake, moment equalling environment, veridical contact stated as an equation.

28.11.14 11.14 The two formularies

The corpus's lookup layer joins the two master formulas, and the join is worth restating: **the moment feeds the storage law as its input.**

```
M_t = (D_t / P_t) E_t                          the moment, as lived
Q_{t+1} = (1 − delta) Q_t + eta M_t − C_t       the moment becoming character
```

Perspective and perception produce the moment; the moment becomes character. From the same corpus: `L = max(0, M − C)`, the launch budget — **the dream of elsewhere is manufactured only from surplus, which is why "get out of your comfort zone" reverses the order of operations**; and `Delta_R = |ln(M'/M)|`, the return difference, in which zero means no information came back at all.

At interpersonal scope the same rows read differently and the most useful is `D/P` as **doctrine-to-data**: an over-framed institution sees what its model predicts; an under-framed one drowns in data it cannot read. Both are measurable and both are common.

28.11.15 11.15 The validation programme

The corpus's answer to the question of whether any of this is science is precise and unflattering to both sides of the usual argument. **Psychology and sociology are not pseudoscience; they are sciences with a measurement problem and a replication problem.** The Open Science Collaboration found 97 of 100 original studies reporting significant results against roughly 36 per cent of replications, with effect sizes roughly halved [Open Science Collaboration, 2015].

Formulas do not make a field hard. Ratio-scale, reproducible, falsifiable measurement of the formulas' terms does. The formula-validation protocol's seven gates are the route [Curzi, 2026u], and the cross-individual recovery gate is the one that decides the question and the one most often skipped. Specific objectivity – a person's measure not depending on which items were used, and an item's measure not depending on who answered – is the closest psychometrics comes to physical measurement and is exactly the property required [Rasch, 1960, Andrich, 1988].

What the corpus has not achieved is recorded there, and the record is short and specific. Not one formula in the framework has passed the cross-individual recovery gate in the open literature, because those studies have not been run. The worked examples throughout demonstrate that the parameters *can* be estimated in a single system; they demonstrate nothing about whether the same operationalisation recovers the same parameters in a different system, with different people, in a different laboratory. **Until that is done, every quantitative claim in the corpus is a worked demonstration of a method and not a measured result about the world**, and the corpus's costed research programme exists because saying so without also saying what it would take to fix it would be an evasion [Curzi, 2026u].

28.11.16 11.16 The inherited notation, read formally

Three inherited technologies are read formally, and each reading is descriptive rather than interpretive:

- **A correspondence table is a finite relational graph.** The Tree of Life is 10 nodes and 22 paths. The tarot is a structured 78-element set. Saying so is a description of a structure, not a claim about it.

- **Gematria is a hash function**, a map from strings to integers. That is a literally accurate technical description and carries no claim that the map is meaningful. §11.20 computes what follows.
- **The sigil is lossy compression plus attention-binding**, and its geometry has an exact arithmetic invariant: a star polygon `{N/k}` decomposes into exactly `gcd(N,k)` closed strokes of length `N/gcd(N,k)`. Verified: `{6/2}` gives 2 strokes of length 3 — **the hexagram is literally two separate triangles**; `{7/2}` and `{7/3}` and `{8/3}` and `{12/5}` give one continuous stroke; `{12/3}` gives three squares; `{9/3}` gives three triangles; `{24/9}` gives three octagons.
- **The memory palace is the storage law with space as the retention medium**, and its advantage factorises into surplus, conversion and retention, of which retention is the largest.
- **Ritual is periodic forcing of a low-pass filter**, and cadence beats intensity for two independent reasons: the per-period ceiling, and the concavity of `f`.

28.11.17 11.17 The three handles

Reduce all of it and three handles remain. **Every operation examined here moves one or more of them, and no operation was found that moves something else.**

```
HANDLE 1 -- THE LENS.          r = D/P, distortion |ln r|, moment M = (D/P)E
                               What you can perceive, and how much of the frame
                               you are supplying yourself.

HANDLE 2 -- THE STORE.         Q, eta, delta, C, h
                               What persists, how fast experience converts into it,
                               how fast it leaks, what it costs, what you can afford.

HANDLE 3 -- THE COMPOSITION. N, u, N_eff, F = ln(1 + (N-1)u), ceiling ln N
                               How many genuinely different things are in contact,
                               and the hard bound on what their contact can generate.
```

The three are not independent — the lens determines what `N` you can distinguish at all, which caps `F`, which drives `Q` — but they are **separately actionable**, which is the property an engineering handle needs.

And the maximum possible moment is `ln N`. That was not chosen; it follows from putting `M` on the same normalised scale as `S`. **The richest available moment is bounded by the logarithm of how many kinds of thing the perceiver can distinguish, which is the law of requisite variety [Ashby, 1956, 1958] appearing in first-person experience.** Only variety absorbs variety, including the variety of one's own life.

28.11.18 11.18 The audit

Eleven operations from the inherited notation, each mapped to a handle and a parameter, each priced where pricing is possible. The same audit form is run in the corpus over a thirteen-entry operation set, returning seven maps, two splits, one rendering and three rejections [Curzi, 2026]; **that this audit found no rejections and that one found three is itself a reason to distrust an audit whose rejection rate is zero. This is the audit that §11.2 requires, and it is presented as an audit of a list, not a proof of completeness.**

Operation	Handle	Parameter moved	What the formalism says it does
purification / banishing before work	2	`C` down, `eta` up	removes competing encodings; less interference per period
consecration of a place	2	`delta` down, `u_cross` up	a place with a fixed assigned meaning is a durable index
the correspondence table	3	declares `N`	fixes the per-interaction ceiling at `ln N`
gematria	3	a hash into `R` buckets	creates adjacency by collision; §11.20
the sigil	2, 3	`eta` up; `gcd(N,k)` strokes	compression to one retrievable token; exact stroke count
invocation / assuming a second framing	3	`k -> k+1` framings	`F_int = ln(1 + (k−1)u_bar)`; the second is the largest gain
the memory palace	2	`u` up, `eta` up, `delta` down	§7.1 and §8.2; the product is the gain
the liturgical calendar	2	periodic forcing	§8.4–8.7; cadence, ripple, and the ceiling
initiation / degree	2	`C` up, `delta` down	costly-requirement break-even, §8.11
the talisman / physical token	2	`delta` down	external storage; see below
divination by random draw	3	forces a non-parallel framing	expected surplus computed exactly, below

Four of the eleven can be priced exactly, and are.
The correspondence table declares `N`, and nothing else about it changes the ceiling:

table	`N`	`ln N` (nats)	bits	octaves
the four elements	4	1.3863	**2.0000**	**4.0000**
the classical planets	7	1.9459	2.8074	5.6147
the Tree of Life, nodes	10	2.3026	3.3219	6.6439
the zodiac	12	2.4849	3.5850	7.1699
the Tree of Life, paths	22	3.0910	4.4594	8.9189
nodes and paths together	32	3.4657	**5.0000**	**10.0000**
the I Ching hexagrams	64	4.1589	**6.0000**	**12.0000**
the tarot, full deck	78	4.3567	6.2854	12.5708

The bold entries are exact because those cardinalities are powers of two: **the 64 hexagrams are exactly six bits, which is not a mystery about the I Ching but a restatement of the fact that a hexagram is six binary lines.** The framework's contribution is only to note that six bits is also `4.1589` nats and `12.0000` octaves, and that these are the same number in three registers.

The talisman is a `delta` reduction and nothing else. At a fixed conversion `eta S = 0.30` per week:

storage medium	`eta S`	`delta`/wk	`Q_inf`	`t_half`	vs held in mind
held in mind only	0.3000	0.060	5.0000	0.22 yr	1.00
written once, filed	0.2100	0.010	21.0000	1.33 yr	4.20
carried, seen daily	0.2550	0.004	63.7500	3.33 yr	12.75
engraved, in the room	0.2400	0.002	120.0000	6.66 yr	24.00

The object stores; it does not act. Its entire advantage sits in one column, and the column is `delta`. Note that the written form has the *lowest* `eta` in the table — a filed document converts less of each occasion than an active memory does — and still wins by a factor of four, because retention is a gain term. This is the framework's retention rule [Curzi, 2026a] arriving from a fourth direction.

Invocation of a second framing has a computable value. `F_int = ln(1 + (k−1)u_bar)` over `k` framings held by one observer with mean pairwise non-parallelism `u_bar`:

`k` framings	`F_int` at `u_bar = 0.7`	marginal gain	% of the `k=12` value
1	0.0000	—	0.0%
2	0.5306	**+0.5306**	24.5%
3	0.8755	+0.3448	40.5%
4	1.1314	+0.2559	52.3%
6	1.5041	+0.1691	69.5%
8	1.7750	+0.1263	82.0%
12	2.1633	+0.0839	100.0%

The second framing delivers 24.5% of what eleven more deliver. Deliberately adopting one other genuine perspective is the highest-return operation available to a single mind, and the twelfth is worth 15.8% of the second. That is the formal content of every tradition's insistence that the practitioner learn to see from somewhere other than where they are standing — and of the framework's warning that the fortieth framing is a maintenance bill [Curzi, 2026o].

A random draw from a correspondence table has an exact expected value. If the drawn framing's non-parallelism with the operator's current one is uniform on `[0,1]`, then

```
E[F] = integral_0^1 ln(1 + (N−1)u) du = (N ln N − N + 1)/(N − 1)
```

`N`	`E[F]`	ceiling `ln N`	as % of ceiling	`E[N_eff] = e^{E[F]}`
4	0.8484	1.3863	61.2%	2.3359
10	1.5584	2.3026	67.7%	4.7513
22	2.2382	3.0910	72.4%	9.3768
64	3.2249	4.1589	77.5%	25.1510
78	3.4133	4.3567	78.3%	30.3650

A randomiser drawing from a 78-element table returns, on average, 78.3% of the ceiling. That is what a divinatory device is worth in this formalism, stated exactly and with no claim whatsoever about the source of the randomness. **It is the same function the framework assigns to a red team or a deliberately recruited dissenter: a generator of framings the operator would not have produced** [Curzi, 2026c]. The corpus's resistance-weighted branch score is the same instrument pointed at a decision rather than at a roster: score each step 0-3 for steepness, sum to a resistance `R`, and weight the branch as `1/2^R`, so a three-point difference in steepness is a factor of eight in expectation [Curzi, 2026]. The randomness may come from cards, from dice, from a table of numbers, or from a colleague who disagrees; the formalism does not distinguish among them and does not need to.

28.11.19 11.19 What the audit does not prove

Eleven operations, three handles, no residue — and that is a finding about eleven operations.

It is not a proof that no twelfth operation exists which requires a fourth handle. The audit's honest form is: *of the operations examined here, each mapped to a named parameter, and each mapping produced a direction and, in four cases, a magnitude.* **A reader who wishes to refute the thesis of §11.1 should exhibit an operation that does work, is agreed to do work, and maps to nothing.** That is a clean falsification target and it is stated deliberately.

Two candidate residues are worth naming, because they are the ones a careful reader will raise and because neither survives inspection as a *fourth handle*:

Candidate one: the felt quality of the operation. Practitioners report that the operations feel like something, and the formalism does not model that feeling. **True, and out of scope rather than unaccounted for.** The moment equation gives `M` as a magnitude, not a quality; the framework states the methodology plainly [Curzi, 2026o] — you postulate a structure, derive consequences, and test the consequences, because the inner state is not directly measurable. **That is a limitation of the instrument and it is declared, not hidden.**

Candidate two: the effect on the world outside the practitioner. This is the one the thesis refuses. Every mechanism in the audit runs through a human being who perceives, encodes, retrieves, or acts. **No operation in the table has a path to any outcome that does not pass through a person, and the framework claims none.** A practitioner whose attention, memory, framing, and cadence have all been engineered will act differently and will therefore produce different outcomes — which is a completely ordinary causal chain and is the only one on offer here.

Candidate three: the coherence of the system as a whole. A practitioner may object that the operations in §11.18 are not independent moves on separate parameters but a single integrated practice, and that the audit destroys the object by decomposing it. **This objection is the most serious of the three and the formalism has an answer that is neither dismissal nor concession.**

The surplus functional is subadditive on the parameter `u` [Curzi, 2026]: the surplus of a combination is at most the sum of the parts' surpluses and is generally strictly less. **Formally, then, an integrated practice is *less* than the sum of its decomposed operations, not more** – §8.2 computes exactly this for the memory palace, where the additive reading of place and image overstated the truth by up to 2.25 nats. So the objection cannot be that decomposition loses a surplus, because the mathematics runs the other way.

What decomposition genuinely does lose is **the cost structure**. Eleven separate operations carry eleven per-occasion costs `c_0`; one integrated practice carries one. Section 8.8 shows that the per-occasion fixed cost is what determines the optimal cadence and can determine whether a practice is affordable at all. **An integrated practice is cheaper to run than its decomposition, and that is a real advantage, and it is a `C` advantage rather than an `S` advantage.** The audit is therefore correct as an audit and incomplete as a design: **it tells you what each operation does; it does not tell you what it costs to do them separately, and the answer to that is usually "more."**

28.11.20 11.20 The restraint, worked: what gematria's correspondences are worth

This is the demonstration that restraint is a sharper instrument than permission, so it is computed rather than asserted.

Gematria maps a string to an integer. As an operation this is exactly a hash, and hashes have a governing arithmetic. For `M` strings mapping into `R` reachable values with probabilities `p_i`, the expected number of colliding pairs is `C(M,2) * sum_i p_i^2`, and **`sum p_i^2` is minimised when the distribution is uniform**, so `C(M,2)/R` is a *floor*. Real letter-value distributions are far from uniform, which makes collisions strictly more common than the floor:

corpus size `M`	reachable values `R`	floor on expected colliding pairs	P(a given word has at least one twin)
500	400	311.9	0.713226
2,000	1,000	1,999.0	0.864665
8,000	1,000	31,996.0	0.999666
8,000	1,500	21,330.7	0.995177
30,000	1,500	299,990.0	1.000000

And the birthday point — how many words you need to draw before a coincidence is more likely than not:

reachable values `R`	exact `M` for P(match) > 0.5	approximation `1.1774 sqrt(R)`
100	13	11.77
400	24	23.55
1,000	38	37.23
1,500	46	45.60
10,000	119	117.74

Draw thirty-eight words from a scheme with a thousand reachable values and a numerical coincidence is more likely than its absence. In a corpus of eight thousand words, essentially every word has a partner — the probability is 0.9997.

Therefore the discovery that two words share a value is not evidence of anything, and the amount of non-evidence is computable in advance. This is the strongest negative result in the reading and it is delivered by arithmetic rather than by disdain.

And it does not abolish the operation. A hash that assigns a word to a bucket is a real encoding device: it produces an addressable index, and §8.2 shows exactly what an addressable index is worth. **What the arithmetic removes is the *evidential* reading, not the *mnemonic* one.** The operation survives as a technique for binding items to a retrieval structure and dies as a method of discovering hidden relations, and the framework says precisely which of the two it is being used as by asking one question: *is the practitioner claiming the collision was informative, or using it to remember?*

28.11.21 11.21 Why the restraint is the source of the power

The four standing refusals look like limits and function as instruments.

A claim that cannot fail cannot inform. If the thesis of §11.1 permitted supernatural efficacy, no audit could test it, because any residue could be assigned to the permitted category and the audit would never close. **The refusal is what makes §11.19's falsification target exist at all.**

A framework that explains everything explains nothing. The insistence on the shuffled-input control, the preregistration, and the seven validation gates exist so that this apparatus can be wrong in a specified way [Curzi, 2026u, Platt, 1964, Lakatos, 1970]. **Every one of those gates is a place where the framework has arranged to lose.**

And the strongest result in this section is a negative one. §11.20's collision arithmetic tells a practitioner more that is useful than any permissive reading could, because it converts a vague suspicion into a number they can compute before they draw the thirty-eighth word.

The mechanism is sufficient. Nothing needs to be added to it, and the reason for saying so is not that additions are forbidden but that in eleven operations, none was required.

28.12 12 Specification and failure modes for the synthesis

28.12.1 12.1 Specification

SS1 — Publish the ledger. Any deployment of this framework shall reproduce §11.3's three-way sort — proved, postulated, requires measurement — for its own claims, and shall not present a postulate in the same register as a theorem.

SS2 — Badge in the same sentence. A postulate shall carry its badge in the sentence that states it, not in a footnote, an appendix, or a preface.

SS3 — Name the handle. Every proposed intervention shall name which of the three handles it moves and which parameter, and shall state the expected direction. An intervention that cannot name its parameter is not specified and shall not be approved.

SS4 — Compute before adopting. Every intervention shall be checked against `h = eta ln N − delta Q − C` before adoption. Where `h < 0` no intervention based on interaction geometry shall be approved, and the design record shall say so in those words.

SS5 — Declare the period and the scope. Every parameter shall carry its period (S2.1) and its scope — internal or external, and which row of the master table it occupies.

SS6 — Preregister and freeze. Every claim of effect shall be preregistered [Curzi, 2026u], fit on a training window, frozen, and reported on held-out data against four baselines, of which one is the **shuffled-input control**. Omission of the shuffled-input control invalidates the claim.

SS7 — Compute the chance rate before reading a coincidence. Where a correspondence, collision, match, or numerical agreement is offered as evidence, the expected number of such agreements under the null shall be computed and reported alongside it (§11.20). An observed rate not exceeding the chance rate shall be reported as no evidence.

SS8 — Distinguish mnemonic use from evidential use. Any operation used as an encoding device shall be recorded as such, and shall not be cited elsewhere as a discovery procedure.

SS9 — No path that does not pass through a person. Every causal claim shall name the human perception, encoding, retrieval, or action through which it operates. A claim with no such step shall be refused.

SS10 — Preserve the ability to decline. No intervention shall be designed, recommended, or deployed that removes a participant's ability to decline it. This applies to cadence, to requirements, to thresholds, and to every operation in §11.18's audit, and it is not subject to trade-off against any quantity in this framework.

SS11 — State the falsification target for the synthesis. Any restatement of §11.1's thesis shall carry §11.19's refutation condition — an operation that does work and maps to nothing — in the same document.

SS12 — Report the audit as an audit. No claim of completeness shall be made for §11.18. The number of operations examined shall be stated whenever the result is cited.

28.12.2 12.2 Failure modes

THE NOTATION FALLACY — concluding that because the old notation and this one describe the same relations, the old notation was secretly doing modern mathematics, or that the modern notation inherits the old one's authority.

- *Parameter damaged:* none. The error is in the direction of inference.
- *Diagnostic signature:* the phrase "the ancients knew" appears, or a formal result is offered as vindication of a tradition, or a tradition is offered as corroboration of a formal result.
- *Measurement:* ask which of the two is being used to license a claim about the other. A correspondence licenses claims about neither.
- *Fix:* state the correspondence as a correspondence. §11.1's thesis is about *this* manual's description, not about what anyone historically believed.

THE PERMITTED RESIDUE — allowing one operation to remain unmapped "for now," on the grounds that the framework is young.

- *Parameter damaged:* the falsifiability of the whole thesis.
- *Diagnostic signature:* an operation is described as "not yet formalised" across successive revisions, while the thesis of sufficiency continues to be asserted.
- *Measurement:* count unmapped operations. The correct number for a sufficiency claim is zero, and §11.19's target is exactly this.
- *Fix:* either map it or withdraw the sufficiency claim for the class it belongs to. **An unmapped operation is not a gap in the theory; it is a refutation waiting to be written up.**

THE PARAMETER AS TARGET — a measured `u`, `eta`, `delta` or `h` becomes a performance metric and ceases to measure anything.

- *Parameter damaged:* all of them, and the instrument as well.
- *Diagnostic signature:* reported `u` rises without any change in the roster or the disagreement structure; `delta` improves in every quarterly report. Goodhart's Law and Campbell's Law.
- *Measurement:* re-derive the parameter by an independent protocol and compare. If the two diverge and only the reported one improves, the metric has been captured.
- *Fix:* separate measurement from evaluation. The survey is run by people who are not accountable for the result, or it is not run [Goodhart, 1975, Campbell, 1979].

THE SCALE CONFUSION — a result established at one scope applied at another without re-measuring the parameters.

- *Parameter damaged:* all of them, silently, because the *form* transfers and the *values* do not.
- *Diagnostic signature:* a personal practice's cadence imposed on an institution, or an institution's review interval applied to an individual. §11.6's `t_95` column spans 28 days to 148 years.
- *Measurement:* check that the period, `N`, `u`, `eta`, `delta` and `C` were all re-measured at the new scope. Scale symmetry is a claim about model form only [Curzi, 2026].
- *Fix:* re-measure. The equation transfers for free; nothing else does.

THE SUFFICIENCY OVERREACH — reading §11.1's thesis as a claim that the formalism explains experience, meaning, or value.

- *Parameter damaged:* none; the error is one of scope.
- *Diagnostic signature:* the framework is used to adjudicate what someone should want, rather than to compute what a wanted thing would cost.
- *Measurement:* ask what the claim would predict. A claim that predicts nothing measurable is not in this framework's domain.
- *Fix:* **The formalism postulates a structure and tests consequences. It has no access to the inner state and says so, and a manual that computes the cost of a thing has said nothing whatever about whether the thing is worth its cost.**

THE CLOSED LEDGER — a deployment that stops publishing which of its claims are postulates.

- *Parameter damaged:* the reader's ability to check anything.
- *Diagnostic signature:* a summary, a slide, or a derived document in which the badges have been dropped for brevity. This is the most common failure and it usually happens in the second generation of a document rather than the first.
- *Measurement:* take any derived document and attempt to reconstruct §11.3's three-way sort from it alone. If you cannot, the ledger has closed.
- *Fix:* re-attach the badges. **The house rule is that a speculative reading is permitted in the open with its badge on; a badge removed for brevity converts a declared postulate into an undeclared one, which is the one thing the framework refuses.**

28.13 13 Worked example: one studio, carried through every part

One system, carried through every part of the framework. **Every parameter value below is a design value obtained by the protocols named beside it and is badged as illustrative; none is a measurement of any real population.**

[II] The roster audit — the geometry of difference.

```
N = 8 declared domains          ln N = 2.0794 nats = 3.0000 bits = 6.0000 octaves
u* = 0.55                       (disagreement and error-decorrelation audit)
N_eff = 1 + 7(0.55) = 4.8500
F = ln 4.8500 = 1.5790          (75.9% of the ceiling)
```

Six people declaring eight domains buy 4.85 effective ones; 3.15 are redundant. Two-source split (Theorem 4.1) with `u_cross = 0.40` between disciplines and `u_div = 0.15` within them:

```
f(u_cross) = 1.3350    f(u_div) = 0.7178    f(u) = 1.5790
bound:  1.3350  <=  1.5790  <=  2.0528
```

The additive figure 2.0528 exceeds the truth by 30.0%, and per SC4 it is not used.

[I] The circuit — storage parameters. Period: one week.

```
eta = 0.35 (Ch 4)   delta = 0.055/wk (Ch 3)   C = 0.13/wk (Ch 5)
eta F      = 0.5526/wk        ceiling eta ln N = 0.7278/wk
net drive  = 0.4226/wk        Q_inf = 0.4226/0.055 = 7.6844
DC gain    = eta/delta = 6.3636      Thm 3.3 cap = 13.2328
pole 0.945   tau = 17.677 wk   t_half = 12.253 wk
t_95 = 52.96 wk    t_99 = 81.41 wk
h at Q_inf = 0.7278 - 0.055(7.6844) - 0.13 = 0.1752/wk      h = 0 at Q = 10.8692
```

Per S2.4 the first honest review is week 53. Per SS4, `h = 0.1752 > 0`, so growth is available and `r_max = 0.1752` per week is the largest increment that may be promised.

[III] The room. 7.0 x 5.0 x 3.0 m gives `V = 105.0` m³ and `S = 142.0` m². Sabine:

average absorption `alpha`	`A` (sabins)	`RT60 = 0.161 V / A`	verdict
0.10 bare, hard finishes	14.20	1.190 s	too live for six-way speech
0.20 furnished, some soft	28.40	**0.595 s**	in band
0.32 treated for speech	45.44	0.372 s	slightly dead
0.45 over-damped	63.90	0.265 s	conversation feels effortful

Target band 0.4-0.7 s requires `alpha` between 0.170 and 0.298. **The room is specified at `alpha = 0.20`.** Six seats at 1.6 m spacing — inside Hall's social distance of 1.2-3.6 m — require a circumference of 9.60 m and therefore a table 3.06 m across.

[IV] The sociology. Measured turn shares across four sessions: 34, 22, 17, 13, 9, 5 per cent.

```
turn entropy H     = 1.6329 nats        max ln 6 = 1.7918
evenness           = 91.1%
effective speakers = e^H = 5.1185 of 6
```

Woolley and colleagues' turn-taking correlate is therefore in reasonable shape and is **not** the binding constraint; the roster's redundancy is. Per Centola and colleagues' measured threshold, `ceil(0.25 x 6) = 2` committed members is the minimum bloc for overturning a house convention — a number the studio records so that a two-person push is recognised as a structural event rather than as a personality clash.

The essence set. Eight domains, ordered and justified, indexed `n = 1..8`. `f(n) = k n^2` with `k = (81/64)/(phi+1) = (9/8)^2/phi^2 = 0.483425732988414` – **a DEFINED constant, exact, which is why fifteen figures are legitimate; the attachment of the unit Hz to it is a separate and stronger claim and is POSTULATED, badged** [Curzi, 2026, Markowsky, 1992].

n	f(n) (Hz)	odd part = pitch class	j	audible (Hz)
1	0.483426	1	8	123.7570
2	1.933703	1	6	123.7570
3	4.350832	3	5	139.2266
4	7.734812	1	4	123.7570
5	12.085643	5	4	193.3703
6	17.403326	3	3	139.2266
7	23.687861	7	3	189.5029
8	30.939247	1	2	123.7570

Eight essences yield exactly `ceil(8/2) = 4` pitch classes — odd parts 1, 1, 3, 1, 5, 3, 7, 1. That is the available harmonic vocabulary and it is a structural limit, not a stylistic choice. The cross-disciplinary index is `N_eff = e^F = 4.8500`, giving `f_cross = k N_eff^2 = 11.371382` Hz, `2F/ln 2 = 4.5560` octaves above base, and 363.8842 Hz after a five-octave transposition. **Band and `j` are stated for every value. These are essence indices rendered as frequencies, not physical resonances, and nothing is claimed to resonate at any of them.**

[VII] The perceptual lens. `D` and `P` estimated by magnitude estimation against a common modulus (the scale-type gate is asserted satisfied here, with a true zero; **the conjoint-measurement test that would establish it was not run**, Luce and Tukey, 1964, Michell, 1997):

```
D = 62.0     P = 41.0     r = D/P = 1.5122     Distortion |ln r| = 0.4136
M = (D/P) E  = 1.5122 E
```

OVER-FRAMED by 51.2%: the studio is seeing what its model predicts. Two corrections, both computed: raise intake so that `P` rises to 62.0 — a 51.2% increase — or lower frame so that `D` falls to 41.0, a factor of 0.6613. The studio chooses the first, because an over-framed group that lowers its doctrine without raising its intake becomes an under-framed one [Curzi, 2026o]. Target `r = 1.05` leaves a residual distortion of 0.0488, **11.8% of the original.**

Lens theorem check: a perceived `+1` unit is a world factor of `e = 2.7183`. The measured distortion of 0.4136 corresponds to a world factor of 1.5122. **The correction is a ratio, budgeted as a ratio, and never promised as a linear improvement.**

[VIII / IX] The formularies. With `E` normalised to 1 and a personal maintenance cost `C = 0.9`:

```
L = max(0, M − C)
at r = 1.5122 (over-framed):  L = max(0, 0.6122) = 0.6122
at r = 1.0000 (calibrated):   L = max(0, 0.1000) = 0.1000
at r = 0.8000 (flooded):      L = max(0, −0.1000) = 0.0000
```

The under-framed state has no launch budget at all — there is no surplus to dream with. Note the uncomfortable and honest reading of the first row: **the over-framed state has the**

largest launch budget in the table, which is why projection is generative and why the correction to `r = 1` costs something real. The studio records that its calibration exercise will *reduce* the volume of new proposals and that this is the expected result rather than a failure.

Return difference on the studio's first external submission: sent `M = 1.00`, world returned `M' = 1.34`, so `Delta_R = |ln(1.34/1.00)| = 0.2927`. **Non-zero, therefore informative.** A `Delta_R` of zero would have meant the world's answer carried nothing.

[XI] The old notation, priced.

- **Correspondence table:** the studio's eight-domain scheme declares `N = 8` and therefore a ceiling of `ln 8 = 2.0794` nats per interaction. Nothing else about the scheme changes that number.
- **Sigil:** the studio's mark is drawn on eight points. `gcd(8,3) = 1` gives **one continuous stroke of length 8** — unicursal. `gcd(8,2) = 2` would give two separate squares. The mark is specified as `{8/3}` because a single stroke is wanted, and that is a decision made on an arithmetic invariant.
- **Memory palace:** the 48-item corpus, encoded to 24 loci with two images each, moves `Q_inf` from 1.9160 to 26.9727 — **a factor of 14.08**, decomposing as surplus 1.5997 x conversion 2.0000 x retention 4.4000.
- **Calendar:** weekly session at the audited `u* = 0.55` (`eta F = 0.5526`), monthly seminar at `u = 0.62` (`0.5863`), annual retreat at `u = 0.86` (`0.6821`), retreat placed in week 26 to avoid collision per SC12.

```
weekly + seminar in one week = 1.1390   vs ceiling 0.7278  -> 0.4112 discarded (36.1%)
```

Three-year trajectory of the built calendar:

week	1	4	13	26	53	105	157	209
Q_t	0.4226	1.7313	4.4040	6.6679	8.1134	8.5193	8.5408	8.5419

Steady-state mean 8.5418, minimum 8.4259, maximum 8.7145, ripple 3.38%. **The calendar returns `+0.8574` over the weekly session alone — an 11.16% gain** — and headroom at the steady mean is `0.7278 − 0.055(8.5418) − 0.13 = 0.1280`, still positive, with `h = 0` at `Q = 10.8692`.

[X] Validation — the frozen model and the conditions of failure.

Fit `Y_{t+1} = a Y_t + b F(U_t) − C_hat_t + eps_t` on weeks 1-53 with `0 <= a <= 1` and `b >= 0`, freeze the coefficients, report on weeks 54-106 against persistence, direct-input, matched autoregressive, and shuffled-input baselines.

FVP gate	Status for this system
1 DIMENSIONS	`D` and `P` share a modulus, so `r` is dimensionless; `M` inherits `E`'s units. Pass.
2 SCALE TYPE	`D`, `P` by magnitude estimation with a true zero; ratio scale, so division is permitted. Pass. `u`, `eta`, `delta` are ratios in `[0,1]`.
3 LIMITING CASES	`P -> 0` gives unbounded `M` (dogma); `D -> 0` gives `M = 0` (sensation without experience); `E -> 0` gives `M = 0`. All named, all observed. Pass.

FVP gate	Status for this system
4 FIXED POINTS	`D = P` gives `M = E`; `Q_inf = (eta F − c)/delta = 7.6844`. Both solved and stated. Pass.
5 IDENTIFIABILITY	`eta` and `delta` are separable only from the transient; the fit window covers weeks 1-53, which is `3 tau`. **Conditional pass — requires a zero-input window.**
6 FALSIFICATION	The five conditions below, preregistered and frozen. Pass.
7 REPRODUCIBILITY	**NOT YET MET.** The parameters have been recovered in one studio. Cross-individual and cross-site recovery is the corpus's costed programme [Curzi, 2026u] and has not been run. **This is stated as a gap, not glossed.**

Preregistered falsification. **This model is wrong if:**

1. the fitted `a` falls outside `[0,1]`, or its interval excludes `1 − delta = 0.9450`;
2. the fitted `b` does not exceed its pre-intervention value;
3. the frozen model does not beat the **shuffled-input control** on held-out data — in which case timing carries no information here and the entire cadence apparatus does not apply to this studio;
4. the empty-palace test recovers fewer than 20 of 24 loci in order at 90 days, which would put the `delta` attribution in the wrong subsystem;
5. `r = D/P` does not fall toward 1 after the intake intervention, which would mean over-framing in this group is not driven by the term the diagnostic assigns it to.

A week-13 reading of 4.4040 is 51.6% of steady state and is not a finding. It is recorded in advance, so that nobody has to argue about it in week 13.

28.13.1 13.1 What the worked example establishes

Psychology and sociology are treated here as one set of relations measured at two scopes. That is a badged thesis, not a demonstrated identity [Curzi, 2026o], and the framework's parameter table is the argument: the same six parameters, the same recursion, the same gauge, with `N` as the categories one mind can distinguish or the domains a group holds, `Q` as character or as institution, `delta` as forgetting or as turnover, and `D/P` as frame-to-intake or as doctrine-to-data.

Four independent quantities in this framework are logarithmic — surplus, pitch, sensation, and the log form of the moment — **and therefore every conversion between them is linear in log space with exact constants.** One nat is 1.442695 bits, 2.885390 octaves, 3462.4681 cents. That is not a numerological coincidence; **the logarithm is the register perception itself operates in**, which Weber and Fechner established experimentally in the nineteenth century and which has been replicated since without interruption.

Subjective magnitude is measurable. Stevens' magnitude estimation produces ratio-scale measurements of subjective quantity that replicate, and that single fact is what licenses the entire volume: division requires a ratio scale, `D/P` requires division, and psychophysics is offered as already building ratio scales for subjective quantities. **That licence is contested from three directions and the contest is recorded in §11.13.** The remaining work is not conceptual: it is the cross-individual recovery gate, and the corpus costs it out at 14.93 full-time-equivalent years per completed programme [Curzi, 2026u].

And the old notation described the same relationships. The correspondence table declares `N`. The sigil compresses and has an exact stroke count `gcd(N,k)`. The palace lowers `delta`. The calendar is periodic forcing. The second framing is worth a quarter of the next eleven combined. The randomiser returns 78.3% of the ceiling. Gematria is a hash whose collisions are computable and therefore evidentially worthless and mnemonically real. **Eleven operations, three handles, and no residue — in an audit of eleven operations, which is what it is and no more.**

If magic is the deliberate engineering of the perceptual lens, the storage of order through symbol, space and cadence, and the composition of non-parallel essences, then the framework has described its mechanism in modern notation. That is the author's thesis, it is a postulate, it wears its badge, and it claims no efficacy of any kind beyond what the storage law and the surplus law already produce — because there is nothing else in it for an efficacy to come from.

28.13.2 13.2 The open problem one layer down

Everything above concerns order that a person or a group **deliberately stores**: a room arranged, a calendar set, a palace built, a framing adopted. In every case the ordering of the category set was chosen by someone, and the essence law assigns an index to an ordinal position in a category set whose ordering a discipline supplied and justified.

Some sequences arrive already ordered, at enormous length, with an alphabet fixed before the ordering was read. The framework's octave-reduction method – folding a raw index into an audible band by powers of two – was developed against exactly such a sequence, at a thirty-six-octave transposition, and applying the essence law there raises questions this paper's arithmetic cannot answer. **What `N` means when nobody designed the sequence; what a defensible ordering is when the alphabet is four; what a channel capacity is when the channel is chemical; and whether any of the framework's parameters has a biological counterpart that can be measured rather than postulated** – those are open, and the corpus states them as open rather than answering them. Two Series B papers take them up directly: the channel into a single cell is computed and the result is used to *eliminate* an explanation rather than to support one, because computing the capacity of a channel leading to a receiver does not create the receiver [Curzi, 2026x]; and the addressing architecture is bounded by a constraint rather than a caveat, since no addressing operation at any signal strength or modality can cause a cell to express content its genome and accessible regulatory repertoire do not already contain [Curzi, 2026y].

The same four refusals bind that work more tightly than they bind this one, because the temptation to claim a physical resonance is largest exactly where the material is physical, and the sentence that governs every use of the essence law governs it there too: **the output is an index, not a physical resonance.**

The instrument does not change. Only `(N, u, eta, delta, C)` change, and the period. **That is the architectural claim, it holds at every scope tabulated in §11.6 from a desk to a city, and whether it holds one layer down is an open question and not a result.**

28.14 14 Results

Every number in this section is exact arithmetic, an exhaustive computation, or the output of a model whose parameters are printed beside it and badged as design values. **No dataset was collected and nothing here is a measurement of any person, group or historical practice.**

28.14.1 14.1 The decomposition theorem, and its exhaustive check

R1. For `N >= 3` and `1 <= k <= N−1`, the figure `{N/k}` consists of exactly `gcd(N,k)` closed strokes, each visiting `N/gcd(N,k)` points and carrying `N/gcd(N,k)` edges; the whole figure carries `N` edges

for every k; it is unicursal if and only if gcd(N,k) = 1; and drawing it requires exactly gcd(N,k) − 1 pen lifts. The proof is three lines of elementary group theory – the orbits of x −> x + k are the cosets of the subgroup generated by k, which is the subgroup generated by gcd(N,k), of order N/gcd(N,k), so Lagrange gives gcd(N,k) cosets. **Verified by explicit trace for every N = 3..40 and every k: 779 cases, zero mismatches.** Of the 143 tabulated rows for N = 3..24, 89 are unicursal.

R2. The attainable stroke counts on N points are exactly the proper divisors of N; the number of unicursal steps is phi(N), so the number of distinct unicursal figures is phi(N)/2 for N > 2; and gcd(N,k) = gcd(N,N−k), verified without counterexample for N = 3..100. **Prime N gives a unicursal fraction of exactly 1.000 and a maximum single-step jump of 0 – perfectly robust, perfectly inflexible. Highly composite N gives the worst unicursal fractions (12, 18 and 24 all sit at 0.333) and the richest decomposition menus.** The two design goals are exactly opposed and the opposition is the divisor structure of N.

28.14.2 14.2 Parallelism, ceiling, and the rate optimum

R3. Read as a schedule, total work is N for every k and only the arrangement changes: P = gcd(N,k) independent workstreams, critical path T = N/gcd(N,k), and **speedup exactly gcd(N,k), with no overhead, because the strokes share no vertex.** The price is exact and is paid in ceiling: the g strokes are disjoint, so no cross-stroke pair interacts within a period and the available surplus ceiling falls from ln N to ln(N/g).

R4. The ceiling rate R(x) = ln(x)/x has dR/dx = (1 − ln x)/x^2, zero at x = e = 2.718282, so **stroke length 3 is the best integer choice at R = 0.366204, with an exact tie between x = 2 and x = 4 at 0.346574 because ln 4 = 2 ln 2.** For N = 12 the rate-optimal figure is {12/4}, four triangles, whose per-period ceiling ln 3 = 1.098612 is **55.8 per cent below** the unicursal ln 12 = 2.484907.

The worked example of Section 6 then contradicts the rate criterion, and the contradiction is the finding. With a merge cost rising at 0.05 per additional stroke and decay charged every period, {12/4} has a *negative* one-step change in stored order (-0.285624) against the single chain's +0.488208, and running the loss four times as often compounds it: over a 60-day quarter, +2.441040 for {12/1} against -5.712489 for {12/4}. **The rate criterion assumes latency is the only cost. It is not, and this is what a ceiling result looks like when the ceiling is not the binding constraint.**

R5. Within the unicursal configurations the step is still a design variable, and it is the only one that changes the reachable fraction of the ceiling. {12/1} and {12/5} have the same N, the same g, the same ceiling and the same latency; they differ in whose framings meet. At an estimated u of 0.31 for adjacent pairs and 0.58 for pairs five apart, {12/1} reaches F = 1.483875 nats and {12/5} reaches 1.998774, and **the per-cycle gain in stored order rises by a factor of 7.14 – +0.037744 against +0.269448 – from changing one number in the schedule.**

28.14.3 14.3 The cadence theorem and what it does not say

R6. For g concave on [0, infinity) with g(0) = 0, D(n) = n g(T/n) is nondecreasing in n, strictly increasing wherever g is strictly concave. The proof is one line: dD/dn = g(L) − L g'(L), and concavity through the origin gives g(L) >= L g'(L). **The hypothesis is satisfied by the framework's own surplus functional, which is strictly concave with f(0) = 0** [Curzi, 2026], so the case for regular practice does not rest on any ceiling.

Measured on the worked sweep: a twelve-hour budget delivered in one occasion yields 0.7278 and in forty-eight occasions yields 3.3235, **a factor of 4.57 from division alone.** At n = 1 the session runs twelve hours against a saturation length of eight, so **four of the twelve hours deliver nothing at all before any other effect is considered.**

R7. The theorem does not recommend infinitesimal occasions, and the term that stops it is named. With a per-occasion fixed cost c_0 = 0.15 the *net* delivered total peaks at n = 5 and falls away on both sides; the optimum is a function of c_0 and of nothing else in the model,

running `n = 24, 13, 8, 5, 3, 2, 1` as `c_0` runs `0.02, 0.05, 0.10, 0.15, 0.25, 0.40, 0.60`. **At `c_0 = 0.80` the net is negative at every `n`: some practices cannot be afforded at any cadence, and the arithmetic says so before anybody's calendar is consulted.** The design consequence is that halving the cost of convening moves the optimum from five occasions to eight, which is very often cheaper than any intervention on the occasions themselves.

28.14.4 14.4 Impulse against cadence, computed honestly

R8. Summing `Q_t` over 400 periods at `delta = 0.06` with a fixed budget `B = 26.0`: **impulse 433.3333, weekly 433.3333, four-weekly 433.3333, and `B/delta = 433.3333`. Identical to every digit.** This is the defining property of a linear time-invariant system and it refutes the commonest informal argument for cadence.

Three quantities are not invariant. The **peak**: 26.0000 for the impulse against 7.9995 for the weekly cadence, a factor of 3.25. The **minimum**, and hence time above a threshold: at a threshold of 3.0 the cadences hold 60 and 61 periods against the impulse's 35, at 7.0 the three are close, and **at 10.0 and above the impulse wins outright while both cadences score exactly zero**, because no cadence funded by that budget ever reaches 10. And the **deliverable total** once the per-period ceiling clips each burst.

R9. With the burst clipped at `eta ln N = 0.7278` (`eta = 0.35`, `N = 8`), the annual delivered total from an identical nominal budget runs 26.0000, 18.9229, 9.4615, 4.3668, 2.9112, 1.4556 and 0.7278 as the interval runs 1, 2, 4, 8, 13, 26 and 52 weeks. **An annual rite delivers 2.8 per cent of what a weekly practice delivers, and the missing 97.2 per cent is not delivered late or inefficiently: the law caps what one interaction can generate and the excess has nowhere to go.**

R10. The steady state settles the argument without reference to any ceiling. **A single impulse of any size whatever, unrepeated, has `Q_inf = 0`; a sustained rate `m` has `Q_inf = m/delta`.** At `delta = 0.06`, rates of 0.10 to 1.00 per week give steady states of 1.6667 to 16.6667 against a column of zeros. Repeating the impulse at interval `P` with the same annual total leaves the **cycle mean constant at 8.3333 in every row** while the trough falls from 7.8333 at `P = 1` to 0.0836 at `P = 104`. **If a system has any functional threshold at all, the mean is the wrong statistic and the minimum is the right one.**

R11. A rite at interval `P` holding a floor `theta` requires `a^P >= theta/(theta + A)`. With the burst at the ceiling and `theta = 0.30`, **a once-a-decade rite requires the trace it leaves to have a half-life of about 5.63 years and a once-a-generation rite about 16.89 years.** That is why such rites are extraordinary, why they are small, and why they take the forms they do. **It is an explanation of a form and not a recommendation of one**, and Section 17 restates the refusal.

R12. A costly requirement is affordable if and only if `dc <= (eta F − c)(1 − delta'/delta)` – **the fraction of the net drive it consumes must be less than the fraction of the decay it removes**, and in the worked table those two percentages are identical in every row, which falls straight out of the algebra rather than out of the parameters. A requirement costing 0.05 per week against a net drive of 0.4226 must reduce `delta` by at least 11.83 per cent to pay for itself.

28.14.5 14.5 The loci decomposition and the palace's fixed increment

R13. The loci-versus-rote steady-state ratio factorises exactly: surplus `F2/F1 = 1.8637`, conversion `eta2/eta1 = 1.9000`, retention `delta1/delta2 = 4.1667`, **product 14.7546, identical to the direct `Q_inf` ratio.** Surplus and conversion together give 3.54; retention alone gives 4.17. **The technique is primarily a retention intervention and only secondarily an encoding one, which is the opposite of the usual emphasis.**

R14. Palace capacity is `ln L` in the number of loci, so **every doubling buys exactly `ln 2 = 0.693147` nats, forever** – 5 to 10 buys the same as 160 to 320. And for a forty-item corpus, ordinary undesigned places already reach **92.4 per cent** of the ceiling while deliberately distinct

places with vivid images reach 97.2 per cent: **the whole apparatus of classical rules buys the last 8 per cent, which is what a logarithm looks like from inside.**

R15. The two-source decomposition prices place and image separately, and the additive reading is always wrong. At `N = 40` with `u = 0.90`, `f(u)` sits close to the lower bound `max(f(u_cross), f(u_div))` in every configuration tested, and **the balanced configuration (0.45/0.45) has the largest gap between the naive sum and the truth – 2.25 nats of imagined surplus that does not exist.** An operator who has invested equally in place-distinctness and image-vividness overestimates the combined effect by more than one who has invested lopsidedly.

28.14.6 14.6 The audit, and what it does not settle

R16. Eleven inherited operations map to three handles – lens, store, composition – with a named parameter and a direction each, and **four of the eleven are priced exactly.** A correspondence table declares `N` and therefore the ceiling `ln N`, with the 64 hexagrams sitting at exactly 6.0000 bits because a hexagram is six binary lines. A physical token is a `delta` reduction and nothing else: **a written, filed record has the *lowest* conversion in the table and still wins by a factor of 4.20 on steady state, because retention is a gain term.** Adopting one further genuine framing is worth `F_int = 0.5306` nats at `u_bar = 0.7`, which is **24.5 per cent of what eleven further framings deliver** – the second is the largest single gain available and the twelfth is worth 15.8 per cent of it. And a random draw from an `N`-element table returns `(N ln N − N + 1)/(N−1)` on average, which is **78.3 per cent of the ceiling at `N = 78`**, with no claim whatever about the source of the randomness.

R17. A letter-sum digest is a hash, and its collisions are computable in advance. For `M` strings into `R` reachable values the expected colliding pairs are at least `C(M,2)/R`, minimised at uniformity and therefore worse in practice; **thirty-eight words drawn from a scheme with a thousand reachable values make a numerical coincidence more likely than its absence, and in a corpus of eight thousand the probability that a given word has a twin is 0.9997.** The operation therefore survives as a retrieval index and **fails as a discovery procedure**, and the framework distinguishes the two uses by asking whether the practitioner claims the collision was informative.

R18. The audit found no residue over eleven operations. **That is a finding about eleven operations and is stated as one.** Three candidate residues were examined: the felt quality of an operation (**out of scope**, since the moment equation gives a magnitude and not a quality); any effect outside the practitioner (**refused**, since every mechanism in the audit runs through a person who perceives, encodes, retrieves or acts); and the coherence of an integrated practice (**partially conceded** – subadditivity means an integrated practice is formally *less* than the sum of its decomposed parts, so the objection cannot be that decomposition loses surplus, but eleven separate operations carry eleven per-occasion costs where one integrated practice carries one, and **that is a real `C` advantage the audit does not capture**).

28.15 15 Discussion

28.15.1 15.1 What follows

An arithmetic invariant converts a design search into an enumeration. The practical content of Theorem 1 is that `gcd(N, .)` has no gradient, so a designer tuning a step parameter by small increments is following a signal that does not exist. The correct procedure is to enumerate the at most `floor(N/2)` candidates and evaluate each, which for `N <= 24` is at most twelve evaluations.

Two ceilings sit on the same `ln N` and are inflated by the same fault. The per-period surplus ceiling and the roster ceiling `m(m−1) <= 2 eta ln(N)/kappa` [Curzi, 2026] are both set by the declared category count, and a scheme whose duplicate rows inflate `N` inflates both [Curzi,

2026u]. **Auditing `N` is therefore prior to every computation in this paper**, and where `N_eff` collapses to 1 the available surplus is zero rather than low [Curzi, 2026].

The cadence result is two independent arguments and they should not be merged. The ceiling argument depends on `ln N`; Theorem 2 depends only on concavity through the origin. A reader who rejects the ceiling still owes an answer to the second, and a reader who rejects concavity is rejecting the surplus functional itself.

The costly-requirement break-even is the paper's sharpest practical test. It requires exactly two measurements – the added cost and the change in decay – and returns a yes or a no. It also names the failure mode that follows from getting it right once: a community that has discovered requirements work keeps adding them past the break-even and cannot tell, because from inside, the requirements that worked and the requirements now destroying it look identical.

28.15.2 15.2 What does not follow

Nothing here shows that any historical practice worked, or that anyone historically knew the theorem. Theorem 1 is a fact about a graph. That the surviving weekday order is the star polygon `{7/3}` and that `gcd(7,3) = 1` is why the week is one closed cycle rather than two disjoint sub-weeks is a fact about a stepping rule, and the tradition that produced the order had the figure without the theorem.

Nothing here identifies a social effect. Every group parameter in Sections 6, 10 and 13 is a design value in a constructed example, and the reflexive specification they sit inside is not identified from observational data [Manski, 1993, Shalizi and Thomas, 2011].

The eleven-operation audit is not a completeness proof, and §11.19 states the refutation condition rather than leaving it implicit. The `C`-advantage of an integrated practice is a gap the audit does not close, and it is conceded there rather than in a footnote.

And the metaphor does not upgrade. "Computing map" imports the suggestion that something is being computed, by something, for some purpose; nothing in the mathematics supplies an input, an output or an agent. The one non-metaphorical sentence is that the number of strokes is `gcd(N,k)`.

28.15.3 15.3 Relation to the surrounding literature

The cadence result has an empirical counterpart in the distributed-practice literature and the two must not be presented as mutual confirmation: **the spacing effect is a claim about human memory encoding and Theorem 2 is a statement about concave functions.** What they share is a structural prediction – that dividing a fixed budget into more occasions helps, up to a point set by a per-occasion cost – and structural agreement between an empirical regularity and a mathematical result is interesting without being evidence for either. **No entry in this corpus's external bibliography covers that literature; the studies are named in §8.14 with their authors, venues and dates and are deliberately left uncited rather than attached to a source that does not support them.** The same applies to the loci literature of §8.3 and the ritual and anthropological literature of §8.9 to §8.11. Section 17 lists them.

Where the bibliography does reach, it mostly disagrees with the framework and is cited for that reason. Collective intelligence correlates with equality of conversational turn-taking rather than with mean ability [Woolley et al., 2010], which supports the effective-count reading; but diversity in the short run has been found to reduce trust and civic participation [Putnam, 2007], with a rebuttal [Abascal and Baldassarri, 2015], and cohesion has been found to raise then lower creative success [Uzzi and Spiro, 2005] – an inverted `U` against a surplus functional that is monotone increasing in `u`. **The corpus's functional is not the only candidate: social impact rising as a sub-linear power of the number of sources [Latané, 1981] and value growing as `n log n` [Briscoe et al., 2006] fit the same qualitative facts, and nothing in this paper discriminates between them.** The strongest available warrant for a logarithm does not

require the perceptual story at all [Peters and Gell-Mann, 2016].

The measurement objection is the one the framework cannot yet answer. Michell [1997] argues that psychology assumes quantitative structure without testing for it; the test is conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971]; **it has not been run for any quantity in this paper.** And the framework's own reflexivity applies to the framework: parameters estimated in a population that has not read it cannot be assumed stable in one that has [Lucas, 1976, Hacking, 1995].

28.16 16 Verification

28.16.1 16.1 Dimensional audit

Symbol	Dimension	Unit	Status
`N`	`[1]`	a count of points, or of categories	**declared**
`k`	`[1]`	a step, in points	**chosen**
`g = gcd(N,k)`	`[1]`	a count of strokes	**derived**, exact
`x = N/g`	`[1]`	points per stroke	**derived**, exact
`W = N`	`[1]`	edges	**derived**, exact
`T = N/g`	`[T]`	periods, one hand-off per period	**derived**
`ln x`, `ln N`	`[1]`	nats	**derived** dimensionless
`R(x) = ln x / x`	`[T]^-1`	nats per period	**derived**
`u`	`[1]`	none	**defined** dimensionless, in `[0,1]`
`F(u)`	`[1]`	nats	**derived**
`Q`	`[Q]`	whatever the declared proxy counts	**modelled**
`eta`	`[Q]` per nat	proxy units per nat	**design value**
`delta`	`[T]^-1`	per period	**design value**
`C`, `c_0`	`[Q]` per period	proxy units per period	**design value**
`h`	`[Q]` per period	proxy units per period	**derived**
`Q_inf`, `Q_max`, `Q_min`	`[Q]`	proxy units	**derived**
`T` (budget), `L` (session)	`[T]`	hours	**declared**
`theta`	`[Q]`	proxy units	**declared threshold**

Does it balance? The star-polygon quantities are pure counts and ratios of counts, so Sections 3 and 4 are dimensionally trivial and the only thing to check is that `R(x) = ln x / x` is a *rate* – nats per period – and is therefore not comparable with `ln x`, a stock of nats per period. **That distinction is exactly what the worked example of Section 6 turns on**, and conflating them is how a rate-optimal configuration comes to have a negative per-period change in stored order.

For the storage side, `Q_{t+1} = (1-delta)Q_t + eta F(u) - C` balances only if `eta` carries `[Q]` per nat and `C` carries `[Q]` per period, which is why `eta` is reported as proxy-units-per-nat and

never as a bare number. `Q_inf = (eta F − c)/delta` returns `[Q]` because `[Q]`/period divided by `1`/period is `[Q]`.

Where there is no operational unit, that is said here. `Q` is stored order and has no natural unit; two systems' `Q` values are not comparable unless they declare the same proxy. The threshold `theta` in §8.10 is a **stipulated** quantity in `Q`-units, and no procedure in this paper measures it. `k` in `f(n) = k n^2`, discussed in §11.12, is a **DEFINED** exact ratio; **the attachment of the unit Hz to it is a separate and stronger claim and is POSTULATED**, and it appears in this paper only as a badged postulate in a ledger.

28.16.2 16.2 Postulates

P1 — `{N/k}` is drawn by the step map `s(x) = x + k (mod N)` on `N` labelled points. *Status:* **DEFINITION** – the object of Theorem 1, and the only thing the theorem is about. Relaxing it to variable steps is a different and larger problem, treated in §4.7.

P2 — In `Z_N`, the subgroup generated by `k` is the subgroup generated by `gcd(N,k)`, of order `N/gcd(N,k)`. *Status:* **ARITHMETIC** – the standard fact that integer combinations of `N` and `k` are the multiples of their greatest common divisor. A reader may verify it on `N = 12, k = 8`: the multiples of 8 mod 12 are `{0, 8, 4}`, which is `{0, 4, 8}`, the multiples of `gcd(12,8) = 4`.

P3 — A finite group's cosets of a subgroup all have the subgroup's order, so their number is the index. *Status:* **CITED** – Lagrange's theorem, standard.

P4 — `f(u) = ln(1 + (N−1)u)` is strictly concave on `[0,1]` with `f(0) = 0`. *Status:* **CITED** [Curzi, 2026] – a property of the framework's surplus functional, proved there.

P5 — Per-period input is capped at `eta ln N`. *Status:* **CITED** [Curzi, 2026c] – the logarithmic ceiling.

P6 — `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t` describes the systems in the worked examples. *Status:* **POSTULATED** – a modelling choice, badged, and the one an empirical test would attack first [Curzi, 2026a].

P7 — The design values in Sections 8, 10 and 13 are the systems' parameters. *Status:* **POSTULATED** and explicitly illustrative. **No value in this paper was measured on anything.**

28.16.3 16.3 The compounding chain

```
P1 + P2   -> L1  the orbits of s are the cosets of <gcd(N,k)>          ARITHMETIC
L1 + P3   -> R1a the number of orbits is N / (N/g) = g                 ARITHMETIC
L1        -> R1b each orbit has N/g elements, so N/g edges             ARITHMETIC
R1a + R1b -> R1c total edges = g * (N/g) = N, for every k              ARITHMETIC
R1a       -> R1d unicursal  <=>  g = 1;  pen lifts = g − 1             DEFINITION
R1a       -> R2  attainable stroke counts = {gcd(N,k)} = proper
                 divisors of N; unicursal steps number phi(N)          ARITHMETIC
P5 + R1b  -> R3  a stroke of x = N/g members has ceiling ln x,
                 not ln N                                              ARITHMETIC
R3        -> R4  R(x) = ln x / x, dR/dx = (1 − ln x)/x^2 = 0 at x = e;
                 ln 3/3 = 0.366204 > ln 2/2 = ln 4/4 = 0.346574        ARITHMETIC
P4        -> R6  D(n) = n g(T/n); dD/dn = g(L) − L g'(L) >= 0          ARITHMETIC
P5 + P6   -> R9  burst clipped at eta ln N = 0.7278; at P = 52 the
                 annual delivered total is 0.7278 against 26.0000,
                 which is 2.7992%                                      ARITHMETIC
P6        -> R10 unrepeated impulse: Q_t = a^t Q_1 -> 0;
                 sustained rate m: Q_inf = m/delta                     ARITHMETIC
P6 + P7   -> R13 (F2/F1)(eta2/eta1)(delta1/delta2)
                 = 1.8637 x 1.9000 x 4.1667 = 14.7546                  ARITHMETIC
P5        -> R14 ln(2L) − ln(L) = ln 2 = 0.693147 for every L          ARITHMETIC
```

Every line is checkable from its own inputs. The three a reviewer can redo in a minute are `ln 3 / 3 = 0.366204` against `ln 2 / 2 = 0.346574`; `0.7278 / 26.0 = 0.02799`; and `1.8637 x 1.9000 x 4.1667 = 14.7546`.

What the chain does not establish. R1 to R4, R6 and R14 are ARITHMETIC or DEFINITION and stand without any empirical content whatever – **they are theorems, and a theorem is no evidence at all that any real system is the object it describes.** R9, R10 and R13 rest on P6 and P7, which are POSTULATED, so **every number in the cadence and loci tables is a property of a model with chosen parameters and not a measurement.** The audit result R18 rests on a stipulative definition and on an enumeration of eleven items, and is **not** a completeness proof.

28.16.4 16.4 Limits, asymptotics, and the model's valid domain

- **`k = 1`.** `gcd(N,1) = 1` always: the convex polygon, unicursal, maximum latency `N`, full ceiling `ln N`. The degenerate end of the design space and often the right answer, as Section 6 finds.
- **`k = N/2` for even `N`.** `g = N/2` strokes of length 2 – the degenerate digons, minimum latency, minimum ceiling `ln 2 = 0.693147`.
- **`N` prime.** Maximum single-step jump 0, `phi(N) = N–1`, one attainable stroke count. Perfectly robust to a mis-set parameter and offering no tuning at all.
- **`g -> N`.** Ceiling `ln 1 = 0`: a decomposition into singleton strokes has **zero** available surplus rather than a low one, which is the same collapse the corpus measures on the roster side when an effective count falls to 1.00 [Curzi, 2026,?].
- **`n -> infinity` in Theorem 2.** `D(n)` increases without bound in the unpriced model, which is why the theorem is incomplete without a per-occasion cost; with `c_0 > 0` the net has an interior maximum, and §8.8 gives it.
- **`delta -> 0`.** The steady state diverges and the contraction is lost; the system never forgets and never arrives [Curzi, 2026a,?].
- **`P -> infinity` at fixed budget.** The cycle mean is unchanged and the trough goes to zero: the limit of a rare rite is a system that spends almost all of its time at nothing.
- **Convex response.** Theorem 2 assumes concavity. **If per-occasion response is genuinely convex over some range – a warm-up cost below which nothing happens – the theorem reverses on that range**, and this is stated as a domain restriction rather than left implicit.

28.16.5 16.5 Invariance, and how far the scale claim has been tested

The architectural claim is that only `(N, u, eta, delta, C)` and the period change between scopes. **That is an invariance claim about model form and it has been exhibited at seven scopes**, from a desk at a daily period to a city at an annual one, with `t_95` running from 28.43 days to 148.28 years. **Exhibited is not tested.** Seven constructed instances with chosen parameters do not establish that the form transfers; what would test it is the cross-laboratory recovery study the corpus costs out and has not run [Curzi, 2026u]. **The honest statement is that the form has been *applied* at seven scopes and *validated* at none.**

The star-polygon results carry a genuine invariance that has been tested exhaustively: the decomposition is invariant under `k -> N–k`, verified with zero counterexamples for `N = 3..100`, and under relabelling of the vertices by any rotation.

28.16.6 16.6 Analogy discipline

Three borrowings appear and each is badged at first use. The **leaky integrator** is a **shared mathematical form**: the same first-order difference equation governs an RC network, an exponential smoother and a stock of order, with no claim that the electrical physics applies [Elaydi, 2005, Oppenheim et al., 1996, Holt, 1957]. **Periodic forcing of a low-pass filter** is the same shared form continued, and the frequency-response statements are statements about

that difference equation. The **hash** in §11.20 is a **derivation**: a letter-sum map satisfies the definition of a hash exactly, and the entropy computed from it is Shannon's quantity in bits [Shannon, 1948, Cover and Thomas, 2006], never thermodynamic entropy.

"The sigil is a nonlinear computing map" is an ANALOGY and is badged as one at every occurrence. Its only defensible quantitative content is that `gcd(N, .)` is a non-smooth arithmetic function of the design parameter, which is a statement about a graph. **No scaling, power-law or heavy-tail claim is made anywhere in this paper**, so the distribution-fitting procedure such a claim would require [Clauset et al., 2009] does not apply and was not run. The one place a constant with a physical unit appears – `k` in `f(n) = k n^2` – is a **defined exact ratio** whose unit attachment is a **postulate**, and the golden-section literature is cited beside it precisely because φ is so routinely claimed where it is not present [Markowsky, 1992].

28.17 17 Limitations

Nothing in this paper was measured on a person, a group, or a historical practice. Sections 6, 10 and 13 are constructed examples with design values, stated as design values at every occurrence, and **must not be cited as evidence about any real design group, study circle or studio.**

Several literatures this paper leans on are named in the text and cited to nothing, because this corpus's external bibliography does not stock them. They are listed here rather than left for a reviewer to find: the classical loci sources and the modern experimental literature on the method of loci; the distributed-practice and retrieval-practice literature; the ritual and anthropological literature; the costly-signalling literature; the source of the letter-reduction procedure analysed in §4.8; and the standard reference for star-polygon notation and the classification of compounds. **In every case the author, work and date are given in the text so that a reader can locate them, and in no case has an unrelated entry been attached to make the sentence look supported.** The arithmetic this paper performs on those sources' *claims* is arithmetic; the claims themselves are cited to their originators by name and to nothing else.

The reflection problem is unanswered. The framework's reflexive specification is not identified from observational data [Manski, 1993, Shalizi and Thomas, 2011], and no remedy is adopted here [Brock and Durlauf, 2001]. **Every statement about `u`, `eta` or `delta` as a property of a group is a modelling position rather than an identified empirical result.**

Scale type is stipulated, not demonstrated. `u`, `Q`, `D` and `P` are treated as ratio-scale quantities wherever differences, ratios, logarithms and means are taken, and the conjoint-measurement conditions that would license that treatment [Luce and Tukey, 1964, Krantz et al., 1971] **have not been tested.** Michell [1997]'s objection applies directly and is conceded rather than answered. The worked example of Section 13 asserts that its scale-type gate is satisfied; **that assertion is a design value like the others and is not a measurement.**

`N_eff = 1 + (N–1)u` is not a Hill number. `1/HHI` and `exp(H)` are Hill numbers of orders two and one and are grounded [Hill, 1973, Jost, 2006, Simpson, 1949]; the effective count used throughout this paper is a linear interpolation between 1 and `N`, is the framework's own construction, and has no counterpart in the ecological, economic or information-theoretic literature. Small-set estimation bias in any effective count is documented and no correction was applied [Chao et al., 2014].

The surplus functional's form is underdetermined by the evidence. A sub-linear power law [Latané, 1981] and an `n log n` law [Briscoe et al., 2006] fit the same qualitative facts as `ln(1 + (N–1)u)`, and **nothing in this paper discriminates between them.** Theorem 2 needs only concavity through the origin and survives all three; the ceiling argument and every number in the clipping table do **not** survive a change of form, and would have to be recomputed.

The empirical claims this paper's models rest on are inherited and are as strong as their sources. The illustrative parameters were not measured; the ritual and memory findings quoted are single studies or small literatures, several of them correlational and one explicitly

survivorship-prone; and the replication base rate that applies to all of them is about 36 per cent with effect sizes roughly halved [Open Science Collaboration, 2015, Ioannidis, 2005].

Every shuffle p-value in the falsification designs sketched here would be uncalibrated. The corpus measures its own shuffled-input control at a false-rejection rate of 0.108 for independent permutation and 0.120 for circular shift against a nominal 0.05, with only block permutation near nominal at 0.060 [Curzi, 2026u]. **The preregistered predictions of Sections 10 and 13 inherit that defect, and any study running them must calibrate the control on its own design before quoting a threshold.**

No information criterion is computed anywhere in this paper, and no effect size carries an interval. The comparisons in Sections 6, 10 and 13 are between *models with chosen parameters*, not between fits to data, so the usual model-comparison apparatus does not apply – but that is a statement about what the worked examples are, not a defence of omitting it from a study that fits anything.

The audit of §11.18 is an audit of eleven operations. It is not a completeness proof, its falsification target is stated, and **the C-advantage of an integrated practice over its decomposition is a real gap in its coverage**, conceded in §11.19.

The threshold `theta` that makes the trough matter is stipulated and never measured. The entire argument that the minimum rather than the mean is the right statistic depends on some functional threshold existing, and **no procedure in this paper identifies one for any system.** A reader who denies that their system has a threshold is entitled to prefer the mean, and the mean is invariant.

The rare-rite arithmetic explains a form and does not recommend one. The requirement that a once-a-decade occasion purchase a half-life of about five and a half years is a computation, and the means by which extraordinary retention has historically been purchased include ordeal, danger and pain. **Nothing here licenses building one**, and the framework's standing refusal – that no intervention may remove a person's ability to decline – is prior to the evidence and unaffected by it.

Two claims the paper does not make. It makes **no causal claim** about any real system, and it makes **no scaling or heavy-tail claim**, so the distribution-fitting procedure such a claim would require [Clauset et al., 2009] does not apply.

What would falsify the results. A single traced star polygon disagreeing with `gcd(N,k)` would overturn R1 and R2, and the trace is cheap to repeat. A concave-through-the-origin response for which division lowers the delivered total would overturn R6. A system in which the integrated effect of a cadence exceeded that of an impulse of the same total, in a linear time-invariant model, would overturn R8 – and would indicate an arithmetic error rather than a discovery. And **an operation that does work, is agreed to do work, and maps to none of the three handles would refute the thesis of §11.1**, which is stated deliberately as a clean target.

28.18 18 Conclusion

A star polygon `{N/k}` decomposes into exactly `gcd(N,k)` closed strokes, each of length `N/gcd(N,k)`; it is unicursal if and only if `gcd(N,k) = 1`; it requires exactly `gcd(N,k) − 1` pen lifts; the attainable stroke counts on `N` points are exactly the proper divisors of `N`; and the number of unicursal steps is Euler's totient `phi(N)`. The result was verified by exhaustive trace over `N = 3` to `40` and every `k` – 779 cases, zero mismatches – and 89 of the 143 tabulated rows for `N = 3` to `24` are unicursal. The hexagram is two disjoint triangles because `gcd(6,2) = 2`, and no technique changes that.

Read as a schedule, `gcd` is parallelism and `N/gcd` is latency, the speedup is exactly `gcd(N,k)` with no overhead because the strokes share no vertex, and **the price is exact: the per-period surplus ceiling falls from `ln N` to `ln(N/g)`.** The throughput rate `ln(x)/x` is maximised at `x = e`, so **stroke length three is the best integer choice at `R = 0.366204`, with an exact tie between `x = 2` and `x = 4`**, while splitting a twelve-station cycle into four triangles costs 55.8 per cent of the

single-period ceiling. **The worked example then shows the rate criterion losing to the single chain once merge cost and per-period decay are charged**, which is what a ceiling result looks like when the ceiling is not binding. Within the unicursal configurations, changing the step from 1 to 5 on the same twelve people multiplies the per-cycle gain in stored order by 7.14.

On the temporal side: for a concave response through the origin, dividing a fixed interaction budget over more occasions never delivers less – delivered total 0.7278 at one occasion against 3.3235 at forty-eight, a factor of 4.57 – and the term that stops the division is the per-occasion fixed cost, which sets the optimum on its own. **The integrated effect of a fixed budget is invariant to how it is spread: 433.3333 for the impulse, the weekly cadence and the four-weekly cadence alike.** What is not invariant is the peak, the trough, and the deliverable total after clipping. **A single impulse has a steady state of exactly zero against a rate's `m/delta`, and computed at identical nominal budget an annual rite delivers 2.8 per cent of what a weekly practice delivers – with the remaining 97.2 per cent not delivered late but not delivered at all.** A once-a-decade rite requires its trace to have a half-life of about 5.63 years, and a once-a-generation rite about 16.89; that explains the form such rites take and recommends nothing.

The memory-palace arithmetic shows ordinary distinct places already reaching 92.4 per cent of the ceiling for a forty-item corpus, and the technique to be primarily a retention intervention: retention contributes a factor 4.1667 against surplus 1.8637 and conversion 1.9000, the product 14.7546 matching the direct steady-state ratio exactly. Every doubling of a palace buys exactly `ln 2 = 0.693147` nats, forever.

The synthesis states its own falsification target and concedes that decomposition understates cost, not surplus. Eleven inherited operations map to three handles with no residue; four are priced exactly; a written record has the lowest conversion in the table and still wins by a factor of 4.20 on steady state, because retention is a gain term; one further genuine framing is worth 24.5 per cent of what eleven further framings deliver; a random draw from a 78-element table returns 78.3 per cent of the ceiling; and thirty-eight words drawn from a thousand-value digest make a coincidence more likely than its absence, which is why such a digest is a real retrieval index and worthless as evidence. **That is a finding about eleven operations, which is what it is and no more. The framing of a drawn figure as a computing map is a metaphor throughout, and the one sentence in it that is not a metaphor is that the number of strokes is `gcd(N,k)`.**

28.18.1 Declarations

Data and code availability

No external dataset was collected for this paper. The star-polygon results are exact arithmetic and an exhaustive enumeration that any reader can repeat in a few lines of code from the definitions in §3.1; the enumeration parameters (`N = 3..40`, all `k`, and `N = 3..100` for the mirror identity) are printed with the results. The cadence, loci and calendar results are model outputs whose generating parameters are printed beside them and badged as design values. **The twelve-station design group of Section 6, the study circle of Section 10 and the studio of Section 13 are constructed and must not be cited as evidence about any real organisation.** Two measurements were taken on external artefacts: the letter-set collision statistics of §4.8 were computed over the alphabetic entries of length six or more in the macOS `/usr/share/dict/words` (`web2`) list, which ships with the operating system and is not redistributed here. The figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi proved the theorems, ran the enumerations and simulations, performed the analysis, produced the figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The groups described are constructed. The material in §8.10 computes the retention an infrequent high-intensity occasion would require and identifies ordeal, danger and pain among the historical means of purchasing it; **that is an explanation of an observed form and not a recommendation, and the framework's standing refusal that no intervention may remove a person's ability to decline is prior to the evidence and is restated in §9.1 and §12.1 as a specification requirement.** The historical and anthropological material is discussed as an object of study and not performed.

Preregistration

This paper reports no preregistered empirical test. Sections 10 and 13 each state a set of frozen predictions that a study of the systems they describe would be preregistered against – including the requirement that the frozen model beat a shuffled-input control on held-out data, without which the cadence argument would not apply to that system – but **no such study has been run**, and those predictions are stated as design requirements rather than as results. **Any study running them must first calibrate the shuffled-input control on its own design**, because the corpus measures that control at a false-rejection rate of 0.108 against a nominal 0.05 [Curzi, 2026u]. The hypotheses of Section 2 were fixed before the enumerations and sweeps were written up; **two of the six were falsified and are reported in Results in the same language a success would have received.**

Chapter 29

The Two Columns and the Concordance Register: Why Agreeing Traditions Stop Adding Evidence After the Second, and a Register That Rejects One Bridge in Five

Abstract. We separate the two columns of a concordance apparatus: measured evidence on one side, historical precedent on the other, and prove that treating them as two independent verifications inflates a posterior. Under a toy likelihood model in which a true proposition is recorded by a tradition with probability 0.30 and a false one with 0.10, a single record carries a likelihood ratio of 3.000, or 1.585 bits; double-counting the same item as evidence turns an honest 3,000:1 into 1,000,000:1. Correlated sources are worse than merely redundant: at a ten per cent chance of shared origin the bits bought by k agreeing traditions run 1.585, 2.546, 2.317, 1.886, 1.685, decreasing after k = 2, because agreement itself becomes evidence of a common source. Pre-modern single-observer detection is quantified — below a standardised difference of 1.0 one observer is worse than a coin — and a survivorship formula shows the same material arriving either 7.9 or 94.5 per cent accurate depending on whether failures were visible, a twelvefold difference. Three bridges are worked to their verdicts: cadence established, vibration-to-cell graded, and the claim that a specific audio frequency repairs DNA rejected on wavelength, energy and octave-separation grounds. Across forty-two register entries, thirty-two are established, six rejected and two badged postulates, with 33.3 per cent carrying a qualifier and 21.2 per cent of attempted bridges failing outright.

Keywords: concordance; likelihood ratio; source dependence; survivorship bias; precedent versus evidence; register status vocabulary; rejection rate

Series: B. Engineering the Surplus **Paper:** CCP-029

29.1 1 Introduction

Two sources agreeing does not double the evidence when only one of them is evidence. This paper takes that sentence, prices it, and builds an auditable apparatus on it.

The setting is a concordance: a body of work that pairs each of its substantive claims with two references, one drawn from the peer-reviewed literature and one from a contemplative, esoteric or otherwise pre-modern textual tradition. Such pairing is common, it is rhetorically powerful, and it is almost always presented in a way that invites the reader to treat the two references as two independent verifications of the same proposition. **That reading is arithmetically wrong, and this paper computes by how much.** Under a likelihood-ratio account in which a tradition records a true proposition with probability 0.30 and a false one with 0.10, a single traditional record is worth `LR = 3.000`, which is 1.585 bits. A replicated experimental result is worth of order `LR = 1000`, which is 9.966 bits. Treating the traditional record as a second replicated result turns an honest posterior of 3,000:1 into 1,000,000:1 — an overstatement factor of 333.3, or **8.38 bits of confidence that was not earned.**

The second contribution is worse news for the practice than the first. Precedent's small contribution is contingent entirely on the sources being independent, and independence is the property that transmitted textual records are worst at supplying. Modelling `k` agreeing traditions with a probability `c` that they share a single origin, the bits bought at `c = 0.1` run **1.585, 2.546, 2.317, 1.886, 1.685** for `k = 1` to `5`. **The total decreases after `k = 2`.** Past two agreeing sources, further agreement is better explained by contact than by independent discovery, and each additional witness argues for diffusion rather than for truth. Informational cascades [Bikhchandani et al., 1992, Banerjee, 1992] and the measured collapse of crowd accuracy under mutual exposure [Lorenz et al., 2011] are the same phenomenon in a modern setting, and they are the reason the arithmetic here is not a curiosity.

The third contribution characterises the instrument that produced the precedent column at all. Sustained first-person observation, transmitted by an institution, has an effectively unbounded sample size and no control condition, no blinding, no calibration and no separation of hypothesis from observation. Repetition drives variance down as `1/sqrt(k)` and leaves bias exactly where it was, so the characteristic product of that instrument is a body of claims that are specific, confident, internally consistent, reproducible on demand **and wrong**. A single uncontrolled observer comparing before with after has a probability of 0.2858 of registering a true standardised shift of `d = 0.20` and exactly 0.5000 at `d = 1.00` — below `d = 1.0` the observer is worse than a coin — against the roughly 392 subjects per group a modern controlled design needs at `d = 0.20`. And a separate filter decides what was transmitted: under a Bayes filter on preservation, the same original corpus arrives either 7.9 per cent or 94.5 per cent accurate depending only on whether failures of its claims were visible, **a twelvefold difference produced by the filter and not at all by the source.**

The fourth contribution is a rejection. Three bridges are worked in full to their verdicts, and the one nearest to the corpus's own subject matter is rejected in print: the claim that a tone at 528 Hz repairs DNA fails on wavelength (2.9167 m in soft tissue against a 10 µm cell, a ratio of `2.9167e5`, and `8.578e9` against a base pair), on energy (`hf/kT = 8.17e−11` at 310 K), and on frequency (molecular vibrational bands sit 34.7 to 37.3 octaves higher). The rejection is then applied to the corpus's own sonification pipeline, which may not claim what it has just refused to let anyone else claim [Curzi, 2026y].

The apparatus is then run over a forty-three-entry register. **Thirty-two entries are ESTABLISHED, six REJECTED, three SPECULATIVE and two badged POSTULATE**; fourteen of forty-two entries (33.3 per cent) carry a qualifier of some kind, and seven of the thirty-three attempted bridges (21.2 per cent) fail the structural-correspondence criterion outright. A method that has never rejected anything is not a method, and those two ratios are properties of the method rather than of the material.

What this paper does not do. It does not establish any of the register's entries; a register is a summary of evidence and never a source of it. It does not collect data: every number here is arithmetic on stated assumptions or a physical constant, and the distinction is marked at each occurrence. And it does not resolve the reflexive identification problem that several of its own register entries inherit; Section 3.5 states that problem where it arises rather than in a footnote.

Sections 2 and 3 give the hypotheses and the method, including the analogy discipline

that governs every physical borrowing in the paper and an explicit disclosure of a gap in the corpus's bibliography. Section 4 develops the likelihood arithmetic. Section 5 characterises the traditional instrument. Section 6 gives four readings of traditional artefacts. Section 7 works four bridges to their verdicts. Sections 8 to 10 give the register's format, the specifications binding on its use, and the named failure modes. Sections 11 to 15 report, discuss, verify, limit and conclude. Appendix A reproduces the register in full, including its rejections. Appendix B lists the works named in the text that carry no bibliography key.

29.2 2 Hypotheses

Five hypotheses, each stated so that it could be wrong and each carrying its falsifier. Two further claims that occupy the same space are **not** hypotheses and are labelled as what they are.

H1 — Precedent and evidence are not commensurable, and the gap is more than half an order of magnitude in odds. Independent traditional arrival at a structure moves the odds on that structure by a factor small compared with the factor a replicated experimental result moves them. **Predicted:** under any assignment of recording probabilities `p1 > p0` that a defender of the traditional record would accept, the ratio of evidence bits to precedent bits exceeds four. **Falsified by:** a defensible assignment on which a single traditional record is worth as much as a single strong study. *(Result: **not falsified** under the assumed rates — 9.966 bits against 1.585, a factor of 6.29; see Section 4.3. The result is conditional on assumed rates and is therefore an arithmetic consequence, not a measurement.)*

H2 — Agreement between correlated traditions is self-limiting, and the limit binds at a very small number of sources. Where sources may share an origin, the evidential value of `k` agreeing reports does not grow with `k`. **Predicted:** for shared-source probability `c > 0` the bits bought by `k` agreeing traditions attain a maximum at small `k` and decrease thereafter. **Falsified by:** any `c > 0` for which the bit total is monotone increasing in `k`. *(Result: **not falsified** — at `c = 0.1` the maximum is at `k = 2` and the sequence falls thereafter; see Section 4.4 and the proof of the limit in Section 13.3.)*

H3 — What survives of a traditional corpus is determined more by the visibility of its failures than by the accuracy of its origin. The preservation filter, not the original hit rate, dominates the accuracy of the transmitted record. **Predicted:** holding the original accuracy fixed, plausible survival probabilities produce surviving accuracies differing by an order of magnitude. **Falsified by:** a filter model in which the surviving accuracy is insensitive to the ratio `s_acc / s_inacc`. *(Result: **not falsified** — 7.9 per cent against 94.5 per cent from the same `a0 = 0.30`, a factor of 11.97; see Section 5.4.)*

H4 — An audible tone cannot act sequence-specifically on genomic material. No acoustic field in the audible band can impose spatial structure at the scale of a base pair, carry quanta comparable to thermal energy at body temperature, or coincide with a molecular vibrational mode. **Predicted:** the wavelength-to-target ratio, the quantum-to-thermal ratio and the octave separation each fail by many orders of magnitude, computed from accepted physical constants. **Falsified by:** a demonstration of sequence-selective acoustic action at audio frequencies, or an error in any of the three computations. *(Result: **not falsified** — ratios of `8.578e9`, `8.17e-11` and 34.7 to 37.3 octaves; see Section 7.5.)*

H5 — A concordance method that is genuinely binding rejects a material fraction of what it is asked to accept. If the three bridge criteria are applied without exemption, the register's rejection rate is not near zero. **Predicted:** the rate of outright

structural-correspondence failure across attempted bridges exceeds ten per cent. **Falsified by:** a register in which every attempted bridge passes, or in which the failures are concentrated in claims the framework had no interest in. *(Result: **not falsified** — 7 of 33 attempted bridges fail C1, 21.2 per cent, and the failures include the claim nearest the framework's own subject; see Section 11.5. The second clause of the falsifier is met by construction rather than by evidence, and that is stated as a weakness in Section 14.)*

Two claims here are not hypotheses.

The recording probabilities `p1 = 0.30` and `p0 = 0.10` are a modelling choice, not a measurement. No study establishes the rate at which pre-modern traditions recorded true structures, and none could without a settled census of what those traditions recorded and an independent verdict on each item. The rates are chosen to be generous to the precedent column, and every bit figure in this paper inherits their status. What is robust across any plausible choice is the *ordering* of the columns and the *sign* of the collapse in Section 4.4, and only those are asserted.

The claim that "the register establishes nothing" is a definition, not a finding. It is true by the construction of a register and is stated so that no later section can trade on the register's tidiness.

29.3 3 Method: the apparatus, the refinement, and what licenses each borrowing

29.3.1 3.1 The concordance apparatus, adopted unchanged

The apparatus is not introduced here. It is carried across from an earlier framework by the same author — a three-state method and three bridge criteria — and is reproduced without modification so that the two bodies of work can be checked against each other line by line. Its provenance is stated because a reader is entitled to know which parts of a method are on trial and which are given.

```
PHYSICS PROOF     (+1)   what scientific measurement confirms
METAPHYSICS PROOF (−1)   what contemplative or traditional practice records
CONTRA−PROOF      (0p)   how the two integrate at the neutral point
```

Three properties of the notation are load-bearing and easy to misread.

The signs are not a value judgement. `+1` and `−1` label two directions of approach — outward, measurable, third-person on one side; inward, experiential, first-person on the other. Neither sign carries rank.

The third state is not a compromise. `0p` is not an average of the two readings, and a bridge is not validated by splitting the difference. The contra-proof asks whether the two descriptions integrate coherently or whether forcing them together requires one side to assert something false. A bridge that requires either column to lie is not integrated; it is one column wearing the other's coat.

The order of operations matters. The procedure fixes the `+1` reading first, in writing, before the `−1` column is opened. The reason is not philosophical but practical: a large enough textual corpus contains a version of nearly every idea a person can have, so searching it after the conclusion is known will always succeed. Fixing the `+1` reading first is the cheapest available guard against the confirmation bias that destroys most work of this kind, and it is the same device, at a smaller scale, that preregistration supplies to an empirical study [Nosek et al., 2018, Simmons et al., 2011]. The degrees of freedom being closed off are exactly those Gelman and Loken [2013] describe: the analyst need not be dishonest for the search to succeed, only free.

The three criteria a bridge must satisfy:

```
C1  STRUCTURAL CORRESPONDENCE   the two descriptions share deep structural
                                features -- not surface similarity, not shared
                                vocabulary, but mathematical or logical structure
C2  PREDICTIVE POWER            the bridge generates a prediction
C3  PRACTICAL UTILITY           the bridge is usable, and on both sides
```

and the binding rule, adopted unchanged: **only bridges satisfying all three are presented as established, and speculative bridges are marked as such.** One tightening is added, and it is the only change made to the inherited apparatus anywhere in this paper: **a speculative bridge is marked speculative in the same sentence that states it** — not in a footnote, not in a later caveat, and not in a summary at the end of the section. A badge that arrives after the claim has been read is not a badge but an apology.

29.3.2 3.2 The four grades of correspondence

Four things are habitually called “correspondence” in ordinary usage, and separating them does most of the work of C1.

```
ISOMORPHISM     exact structural identity; the mathematics is the same and only
                the labels differ
HOMOMORPHISM    partial structural similarity; some features map, others do not
ANALOGY         functional similarity without structural identity; different
                mechanisms, similar effects
METAPHOR        surface similarity that may or may not indicate a deeper connection
```

The threshold: a valid bridge requires at least homomorphism. Isomorphism is ideal. Analogy requires additional validation. Metaphor alone is insufficient.

Most published correspondences between physics and contemplative traditions are metaphors, and they persist because a metaphor is genuinely useful for teaching and genuinely useless for inference — two functions that are easy to confuse when a sentence performs the first one well.

29.3.3 3.3 The refinement: what each column can carry

Everything above is inherited. What follows is the refinement this paper adds, and it is added because the biological material downstream is where the temptation to over-claim is strongest.

```
THE (+1) COLUMN CARRIES EVIDENCE.
    what has been measured, replicated, and could have come out otherwise

THE (-1) COLUMN CARRIES PRECEDENT.
    that the idea was arrived at independently, recorded, and preserved --
    often for centuries, by people with no access to instruments
```

Two sources agreeing does not double the evidence when only one of them is evidence. That is the refinement in one sentence, and Section 4 is its unpacking. It is not a retreat. A concordance that admits what each side can carry is more persuasive than one claiming both sides are proof, for the reason that a reader who checks will find the argument was straight with them, and a reader who checks is the only reader whose opinion is worth having.

What independent arrival does establish is real and worth having: that a structure was findable by more than one route. That is an argument about a specific and important thing — whether the structure is in the territory or only in the instrument. A pattern that only one method ever finds may be an artefact of that method. A pattern that a Renaissance compiler

with a wax tablet and a modern laboratory both arrive at is, at minimum, not an artefact of the laboratory. That is the whole of what precedent buys. It is not nothing; it is also not evidence, and the distance between the two is measurable.

29.3.4 3.4 Analogy discipline: what licenses each borrowing

The corpus this paper belongs to is titled *Psychological and Sociological Physics*, and that word is a promise that has to be kept sentence by sentence. **Every borrowing from physics or mathematics in this paper is labelled here, at first use, as exactly one of three things**, and no later section may let a category-2 or category-3 borrowing collect the authority of a category-1 one.

Borrowing	Category	What licenses it
Bayes' theorem and the likelihood ratio (§4)	**(1) derivation**	The premises are satisfied by construction: a prior, and two conditional probabilities that the model *defines*. The arithmetic is exact; the inputs are assumed, and that is a separate matter carried in §4.2.
Bits as the unit of evidence (§4)	**(1) derivation**	`log2` of a likelihood ratio is Shannon information about a binary hypothesis [Shannon, 1948, Kullback and Leibler, 1951]. **This is Shannon entropy, not thermodynamic entropy**, and the two are different objects with different units; the distinction is stated wherever the word appears.
`1/sqrt(k)` variance reduction (§5.2)	**(1) derivation, under a premise the subject fails**	The result is exact for independent observations. The traditional instrument's observations are not independent, which is precisely why the section's conclusion is about bias rather than variance.
The leaky integrator applied to ritual cadence and to mnemonic retention (§6.2, §6.3)	**(2) shared mathematical form**	`Q_{t+1} = (1−delta)Q_t + eta S_t − C_t` is a first-order linear difference equation [Elaydi, 2005]; its DC gain and ripple behaviour are facts about that equation [Oppenheim et al., 1996, Horowitz and Hill, 2015]. **No claim is made that a liturgy is an electrical circuit or that memory is a capacitor.** The equation is the same; the physics is not transferred.
The wavelength, thermal-energy and vibrational-band computations (§7.5)	**(1) derivation**	These are claims about a physical acoustic field acting on physical matter, and the physics applies directly. This is the one place in the paper where a physical result is used as a physical result, and it is used to *reject* a claim rather than to support one.
The Casimir pressure scaling as `a^−4` (§7.2, Appendix A CR-17)	**(1) derivation**	The scaling is the result itself [Casimir, 1948], measured at the separations where it is appreciable [Lamoreaux, 1997, Mohideen and Roy, 1998]. It is invoked only to show that the effect is negligible at biological scale.

Borrowing	Category	What licenses it
The time–bandwidth limit named in the register (Appendix A CR-11, CR-12)	**(2) shared mathematical form**	`dt df >= 1/(4 pi) = 0.079577471545948` is classical Fourier mathematics [Gabor, 1946] and carries **no quantum content**; the corpus proves it by the same variational argument as `sigma_x sigma_p >= hbar/2` with `p = hbar k` deleted, with Gaussians saturating the bound and a Hann window costing 2.6 per cent [Curzi, 2026z]. The position–momentum relation [Kennard, 1927] is a fact about state preparation, not about measurement disturbance, and the register states it that way.
The framework's constant `k = (9/8)^2 / phi^2 = 0.483425732988414` (Appendix A CR-39, S-R9)	**definition, with a postulate attached**	`k` is an exact ratio of exact quantities and is **defined**, so its fifteen figures are legitimate. **Attaching hertz to it is a separate and much stronger claim and is POSTULATED**, not derived. The standard critique of golden-section claims [Markowsky, 1992] is cited here deliberately and again wherever `phi` appears; the identity `phi^2 = phi + 1` is Euclidean [Euclid, 1956, Vajda, 1989].
The effective count `N_eff = 1 + (N−1)u` (§6.1)	**(3) analogy, badged**	**This is not a Hill number.** `1/HHI` and `exp(H)` are Hill numbers of orders two and one and are fully grounded [Hill, 1973, Jost, 2006]; `1 + (N−1)u` is a linear interpolation between 1 and `N`, is the framework's own construction [Curzi, 2026], and has no counterpart in the ecological, economic or information-theoretic literature. Wherever an effective count is named in this paper, it is this one, and it is named.

Most of the borrowings in this corpus are category 2, and category 2 is respectable. A leaky integrator is a leaky integrator whether it holds charge or holds a practice. The failure mode this table exists to prevent is a category-2 use quietly collecting the certainty of a category-1 one.

29.3.5 3.5 Identification and reflexivity, stated in the method

Several of the register's entries rest on **reflexive social mechanisms**: that a belief propagates when coupled to a true feature of the world [Merton, 1948], that another person's expectation of one's performance affects it [Rosenthal and Jacobson, 1968], and that a stated commitment binds the speaker's later conduct [Curzi, 2026]. Those are the corpus's own reflexive loop in the register's clothing, and **they inherit the reflection problem in full.**

The framework's central formal object is a loop in which the group influences the member and the members constitute the group. That is a linear-in-means reflexive specification, and Manski [1993] proves that in such a specification **the endogenous social effect is not identified from observational data**: it cannot be separated from contextual effects or from correlated unobservables, however much data is collected. Shalizi and Thomas [2011] establish the same confounding for network data, where influence and selection are generically indistinguishable from observation, and Lyons [2011] makes the point against a specific and much-cited application.

No remedy is adopted in this paper. The four known routes are randomisation, an instrument, partial-population variation, and stipulating the reference group and defending the stipulation [Brock and Durlauf, 2001]; none is executed here, because this paper collects and

prices claims rather than estimating a social parameter. **The honest label is therefore that every register entry resting on a reflexive mechanism records a modelling position rather than an identified empirical result**, and the register's own `ROLES` field, which prices the precedent side, does nothing to repair the evidence side. The one genuine existence proof available is MacKenzie and Millo [2003], which documents a reflexive fixed point occurring in the field with dates and price data; that is an observation of the mechanism happening, not an identification of its coefficient, and the difference is the whole of the problem.

The Lucas critique applies to this paper's own apparatus. Lucas [1976] holds that parameters estimated under one policy regime do not survive the announcement of another, and Hacking [1995] makes the same point about human classifications, which change the people they classify. A register that publishes its weights invites its own entries to be written to the weights. The position taken here is narrow and is stated rather than assumed: **the weights are declared constants of the method, not estimated parameters, so there is nothing for a respondent to bias**; but if a later use of the register begins fitting the weights to the entries, that use inherits the critique and this one does not protect it.

29.3.6 3.6 A gap in the bibliography, disclosed

The biological literature this paper describes — placebo analgesia and its opioid mediation, open-label placebo trials, mechanically gated ion channels, force transmission to the nucleus and its transcriptional consequence, sonogenetics, cochlear mechanotransduction, nucleotide data storage, guide-sequence addressing, epigenetic regulation, retroviral domestication and autophagy — has no corresponding cluster in the bibliographic database governing this corpus. Every such work is therefore named in the text and in Appendix A with the bibliographic details the source manuscript supplies, and **without a citation key**. Attaching an unrelated key to a biological claim would be worse than the gap, because a reviewer in that field will check it and a mis-supporting citation destroys more than a missing one. Appendix B lists the affected works, Section 14 records it as limitation L-6, and **this is a defect of the corpus's bibliography rather than a claim that the works do not exist.** The same gap is disclosed in the neighbouring papers of this series [Curzi, 2026x,y] and in the sonification protocol [Curzi, 2026].

Three further items are named without keys for the same reason and are listed in Appendix B: the distributed-practice synthesis underpinning the spacing effect, the memory-palace imaging studies, and the anthropological account of ritual invariance.

29.3.7 3.7 Standing honesty clauses

- **A tradition is a witness to an effect and never to its cause.** Nothing in a first-person instrument could have distinguished causes, so a recorded mechanism is outside the instrument's range by construction (Section 5.3).
- **Duration of belief is not support.** That a doctrine was held for two thousand years is a fact about transmission fidelity — a low `delta` — and `delta` is orthogonal to whether the stored content is true [Curzi, 2026a]. A tradition with excellent retention preserves its errors with exactly the fidelity of its findings.
- **The register establishes nothing.** It is a summary of evidence and never a source of it; no entry becomes truer by being in it.
- **No entry may be combined with another to close an open problem.** The specific case is the demodulation gap [Curzi, 2026y], and the prohibition is stated as failure mode FM-R.4.
- **Standing refusals.** The framework asserts no supernatural claim, no physical efficacy for a symbol, no clinical or therapeutic claim, and no claim that genomic sequence is edited or altered by any procedure it describes. These are constraints on what the framework will assert and are not verdicts on what is true.

- **Scale type.** The status vocabulary of Section 8.3 is **ordinal at best** — ESTABLISHED is not twice SPECULATIVE — and no arithmetic is performed on it anywhere in this paper. The bit weights of Section 4 are on a ratio scale by construction, because they are logarithms of a defined ratio; they are not measurements of anything.

29.4 4 The arithmetic of the asymmetry

29.4.1 4.1 Likelihood ratios and bits

The distinction between the columns can be made numerically, which is better than making it rhetorically. The natural currency for combining independent considerations about a hypothesis is the likelihood ratio, and the natural unit is the bit [Jaynes, 2003, Cover and Thomas, 2006].

Let `H` be the proposition that some structure `S` is a real feature of the territory. A consideration `X` moves the odds on `H` by its likelihood ratio:

```
LR(X) = P(X | H) / P(X | not H)

posterior odds = prior odds x LR(X)
in bits:  log2(posterior odds) = log2(prior odds) + log2(LR(X))
```

That additivity in bits is the entire reason for working in this unit, and it is a **derivation**: it follows from Bayes' theorem for considerations that are conditionally independent given `H`. Section 4.4 is what happens when that independence fails, which is the normal case for the object under study.

The corpus prices other kinds of consideration in the same currency, and the parallels are worth naming because they show the device is not special to this material. A branch of a forecast is scored 0–3 for the steepness of each step, summed to a resistance `R`, and weighted as `1/2^R`, so that a three-point difference in steepness is a factor of eight in expectation [Curzi, 2026]. And the corpus's grounding protocol prices predictions directly: a direction-only prediction is worth a likelihood ratio of 2.0, a plus-or-minus 20 per cent magnitude call against a two-decade prior is worth 11.36 — about 3.5 direction calls — and five inferential links at 70 per cent each deliver 6.3 per cent end to end [Curzi, 2026ac]. **Both are the same operation as this section's: put heterogeneous considerations on one exponential scale and read off which term dominates.**

29.4.2 4.2 The toy model, badged

The rates below are assumed, not measured, and they are a modelling choice (Section 2). They are chosen to be generous to the precedent column. Nothing in this paper rests on their exact values; what rests on them is the *shape* of the conclusion, which is robust across any plausible alternative.

Suppose a tradition records a structure with probability 0.30 when the structure is real, and with probability 0.10 when it is not. The second number is deliberately not small. Traditions have preserved a very large quantity of confidently stated material that turned out to be false, and Section 5 is an account of why.

```
PRECEDENT
    P(recorded | true)  = 0.30
    P(recorded | false) = 0.10
    LR = 3.000  =  1.585 bits  =  4.77 decibans

EVIDENCE                                  (illustrative magnitudes)
```

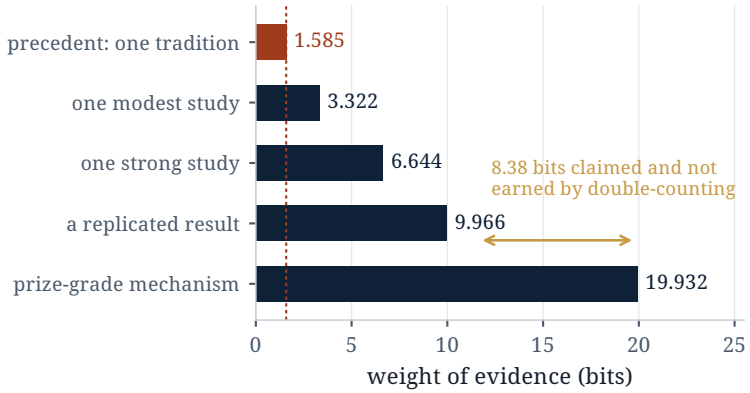


Figure 29.1: **A replicated result is worth six times what independent traditional arrival is worth, and double-counting the precedent claims 8.38 bits that were not earned.** Evidence weights in bits under the illustrative likelihood ratios of Section 4.2 against the precedent weight of 1.585 bits from the assumed recording rates 0.30 and 0.10. The bracketed span marks the 8.38-bit overstatement produced by treating one traditional record as a second replicated result. Modelled from the assumed rates stated in Section 4.2; generated by `figs/src/CCP-029_fig01.py`.

```
one modest study                    LR ~ 10      =  3.322 bits
one strong study                    LR ~ 100     =  6.644 bits
a replicated result                 LR ~ 1000    =  9.966 bits
a replicated mechanism, prize-grade LR ~ 1e6     = 19.932 bits
```

The evidence magnitudes are illustrative in the same sense and to the same degree as the precedent rate, and they are chosen conservatively against the well-documented base rate: large-scale replication returns roughly a third of original psychological findings with effect sizes about halved [Open Science Collaboration, 2015], most published research findings in low-prior fields are false [Ioannidis, 2005], and directional predictions in psychology are nearly costless [Meehl, 1967, 1990]. **A reader who thinks `LR ~ 1000` too generous for a replicated result should note that lowering it narrows the gap this section reports and does not close it**, because the precedent side is capped at `p1/p0 = 3` by construction.

29.4.3 4.3 The price of double counting

Read the two columns together. **A replicated experimental result carries roughly six times as many bits as independent traditional arrival — `9.966 / 1.585 = 6.29` — and a prize-grade replicated mechanism about thirteen times as many, `19.932 / 1.585 = 12.57`.** Precedent is a real term in the sum, and a small one.

The cost of the mistake is computable. Starting from even odds, with a replicated result on the +1 side:

```
honest:   1 x 1000 x 3     =     3,000 : 1
doubled:  1 x 1000 x 1000  = 1,000,000 : 1

overstatement factor 333.3  =  8.38 bits of confidence not earned
```

A concordance that treats its precedent column as a second evidence column is claiming eight bits it does not have, and eight bits is the distance between "well supported" and "essentially certain." That is not a rounding error; it is the entire interval a careful reader cares about.

29.4.4 4.4 Independence collapse

Precedent's small contribution depends entirely on the word *independent*, and independence is exactly the property the transmitted record is worst at supplying.

If five traditions record the same structure and all five inherited it from one another, the five reports are one report wearing five costumes. This is not a hypothetical hazard. The traditions most often paired in this literature — Hellenistic, Egyptian, Jewish, Christian, Islamic, and the Renaissance compilations that stitched them together — are in documented contact throughout, and the early-modern editors who assembled the corpus actually read were *deliberately* harmonising them. Apparent cross-tradition agreement in that corpus is more often evidence of a shared library than of independent discovery.

Model it. Let `k` traditions all report structure `S`, and let `c` be the probability that they share a single source — that is, that they are effectively one report:

```
P(all k report | H)     = c * p1 + (1 - c) * p1^k
P(all k report | not H) = c * p0 + (1 - c) * p0^k

with p1 = 0.30, p0 = 0.10;  LR = the ratio of the two.
```

Computed, in bits:

k	c=0.0	c=0.1	c=0.3	c=0.5	c=0.9
1	1.585	1.585	1.585	1.585	1.585
2	3.170	2.546	2.048	1.826	1.616
3	4.755	2.317	1.827	1.695	1.598
4	6.340	1.886	1.670	1.622	1.589
5	7.925	1.685	1.612	1.596	1.586

The first column is the fantasy: perfectly independent traditions, and the bits add linearly. Every other column is the realistic case. **At only a ten per cent chance of a shared source, five agreeing traditions buy 1.685 bits — barely more than the 1.585 bits one tradition buys alone.** At `c = 0.3`, five traditions buy 1.612 bits, within two per cent of a single tradition's contribution.

Read down the `c = 0.1` column as `k` rises: 1.585, 2.546, 2.317, 1.886, 1.685. **It goes down after `k = 2`, and that is not a numerical accident.** Once several sources agree, "they copied each other" starts to explain the agreement better than "they each found it," and the extra witnesses begin to argue for diffusion rather than for truth. Section 13.3 proves the limit: as `k` grows, both numerator and denominator approach `c p1` and `c p0`, so the likelihood ratio approaches `p1/p0 = 3` from above, and the bit total approaches 1.585 from above — the value a single tradition already supplied.

This is the same structure as an informational cascade, in which agents rationally follow their predecessors and the aggregate ceases to aggregate anything [Bikhchandani et al., 1992, Banerjee, 1992]. It is also the mechanism Lorenz et al. [2011] measured directly: social influence narrows the diversity of estimates without improving their accuracy, so the group looks more confident and is not more correct. Salganik et al. [2006] shows the same dependence experimentally in cultural markets — what survives is heavily path-dependent, and success in one world does not predict success in another. The old corpus is a cultural market with a very long horizon.

Two results elsewhere in the corpus make the same structural point about witnesses. An identifier drawn from a shared pool identifies a pool and not a member: with an address-only log the anonymity set is 1,953.1 subscribers, with address, port and live session state it is 1, and after state expiry it is 1,453.2 [Curzi, 2026]. **An agreeing tradition is an address, and a shared library is the pool.** And the repair in both cases is the same: compare against a *separate witness* rather than a second view of the same record — comparing entries on a

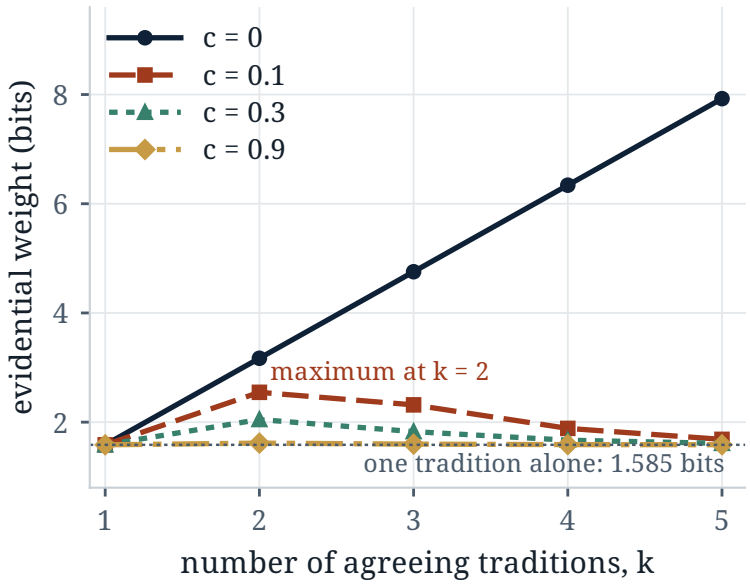


Figure 29.2: **Correlated agreement is self-limiting, and the limit binds at two sources.** Bits bought by `k` agreeing traditions under the shared-source model of Section 4.4 at `c = 0.0` (perfect independence), `0.1`, `0.3` and `0.9`. Only the `c = 0` line grows with `k`; every correlated line attains its maximum at `k = 2` and decreases thereafter toward the single-source value of 1.585 bits, marked. The `c = 0.5` column of the source table is omitted because the validated palette holds four colours and no fifth; its values lie between the `c = 0.3` and `c = 0.9` lines throughout. Modelled from the assumed rates of Section 4.2; generated by `figs/src/CCP-029_fig02.py`.

canonical form with an independent byte-hash witness takes a reconciliation false-break rate from 40.0 per cent to 0.0 per cent at a 40 per cent manipulation rate [Curzi, 2026].

The operational consequence is a rule rather than a caution:

> **A precedent citation must name the tradition, the approximate date, and the transmission path, so that a reader can judge independence for themselves. A precedent citation of the form "many traditions teach" is worth zero bits and is not admissible.**

That rule is enforced entry by entry in Appendix A, and it is why the register has a field for what the source *actually recorded* rather than for what it is usually said to mean.

29.4.5 4.5 What the refinement protects, and what it is not

Three failure modes are eliminated outright by stating the column roles every time.

It eliminates the laundering of tradition through science. The move is: an old text says something evocative; a modern paper says something superficially similar; the two are printed side by side; the reader's confidence in the paper transfers to the text. Nothing false has been asserted, which is what makes it effective and what makes it dishonest. Naming the columns blocks the transfer, because the reader has been told in advance which side is load-bearing.

It eliminates the reverse laundering. The mirror move shores up a genuinely uncertain scientific result by noting that a tradition "already knew". A contested result plus a precedent is still a contested result; precedent cannot repair an underpowered study.

It eliminates the argument from duration. Addressed as a standing clause in Section 3.7.

The refinement is **not** a claim that the precedent column is worthless, and the framework would be much poorer without it. Precedent does four jobs evidence cannot do alone.

It supplies hypotheses. The old corpus is a very large, very cheap, very strange hypothesis generator, assembled by people whose incentive was to notice things about being a person. Almost all of it is wrong, and a low hit rate is affordable because the evidence column does the filtering.

It supplies the findability argument stated in Section 3.3, which is specifically about the distinction between territory and instrument — one of the hardest distinctions in any science of the subjective.

It supplies operational detail that experiments omit. A paper reports that a practice produced an effect. A tradition reports the posture, the time of day, the duration, the cadence, the sequence and the failure signs, because it had to hand the practice to the next generation. **The evidence column tells you whether something works; the precedent column often tells you how to do it.**

It supplies the prior — precisely, and only, the 1.585 bits computed above.

29.5 5 The instrument that produces the precedent column

Section 4 assigned the precedent column a job and priced it at roughly one and a half bits. This section asks the question that pricing assumed an answer to: given that traditional observation is uncontrolled, unblinded, and preserved by a filter that does not select for accuracy, why is any of it right at all?

29.5.1 5.1 The specification sheet

The precedent column is produced by an instrument. Specifying it the way a manual specifies any instrument makes most of what follows immediate.

```
INSTRUMENT:  sustained, repeated, first-person observation,
             transmitted across generations by an institution.

SAMPLE SIZE                effectively unbounded; millions of observers,
                           thousands of years
DURATION PER OBSERVATION   a lifetime, in some traditions
BLINDING                   none, in either direction
CONTROL CONDITION          none
RANDOMIZATION              none
CALIBRATION                none against any external standard
INSTRUMENTATION            none beyond the unaided senses
RECORDING MEDIUM           memory, then text; both lossy, both editable
PUBLICATION FILTER         extreme, and not selecting on accuracy
PRE-REGISTRATION           none; hypothesis and observation are not separated
```

Read that as a specification sheet rather than as a criticism, because its two halves pull in opposite directions and the tension is the whole story.

The top of the list is extraordinary. No modern study has a sample size of millions of observers over three thousand years, and none observes a single subject for a lifetime. Duration and `n` are the two things this instrument has in overwhelming abundance and are expensive in every other research programme.

The bottom of the list is catastrophic. No control, no blinding, no calibration, no separation of hypothesis from observation. Those are exactly the devices that convert observation into inference [Platt, 1964, Popper, 1959], and their absence is what the last century of methodological work has been about [Nosek et al., 2018, Simmons et al., 2011].

29.5.2 5.2 What repetition buys, and what it does not

The tension resolves cleanly, and the resolution is the most useful thing in this section.

Repetition reduces variance. The standard deviation of a mean over `k` independent observations falls as `1/sqrt(k)`:

k	sd of mean	smallest effect resolvable above noise (2 sd)
1	1.00000	2.00000
10	0.31623	0.63246
100	0.10000	0.20000
1,000	0.03162	0.06325
10,000	0.01000	0.02000
1,000,000	0.00100	0.00200

Repetition does nothing whatever to bias. If every observer is pushed in the same direction by the same systematic influence — expectancy, social pressure, the doctrine they were taught before they looked, the fact that they were told what they would see — then that influence appears at full strength no matter how many observers there are:

```
a systematic bias b appears at magnitude b for k = 1
                         and at magnitude b for k = 1,000,000
```

The old corpus therefore has, uniquely among bodies of observation, very small variance and completely unbounded bias. That combination has a signature, and once recognised it is visible everywhere in the material:

> **The errors of the old corpus are not vague. They are specific, confident, internally consistent, reproducible on demand — and wrong.** A doctrine held by ten thousand careful observers for eight hundred years reports its false content with the same steadiness and the same wealth of circumstantial detail as its true content, because steadiness is what `k` buys and `k` is the one thing the instrument has.

This is why the corpus does not read like a body of tentative guesses. It reads like a body of settled knowledge, and it reads that way whether or not it is one. **A reader who calibrates on the corpus's *tone* will be wrong in exactly the places it matters. Tone is a measure of `k`, not of truth.**

Two independent literatures reinforce the point and are cited because they undercut a comfortable reading of the precedent column rather than support it. Introspective report is not a reliable guide to the processes producing behaviour: people confidently report causes of their own judgements that experimental manipulation shows to be wrong [Nisbett and Wilson, 1977], and self-report is systematically shaped by the question's own form and context [Schwarz, 1999]. Judgement of magnitudes is further distorted by the range and spacing of the stimuli presented [Teghtsoonian, 1971, Poulton, 1979] and by the adaptation level the observer arrives with [Helson, 1948]. **Every one of those effects is a bias in the technical sense, and every one of them is invariant under `k`.**

29.5.3 5.3 The detection floor and the operating envelope

What kind of effect can an instrument with no control detect at all?

A single observer comparing before with after has one signal available: a change larger than their own ordinary day-to-day variation. Taking that variation as one standard deviation and the difference of two such observations as `sqrt(2)`, the probability that a true standardised shift `d` produces an observed change exceeding one standard deviation is `P = 1 - Phi((1 - d)/sqrt(2))`:

```
d = 0.20  ->  P = 0.2858
d = 0.50  ->  P = 0.3618
d = 0.80  ->  P = 0.4438
d = 1.00  ->  P = 0.5000
d = 1.50  ->  P = 0.6382
d = 2.00  ->  P = 0.7602
d = 3.00  ->  P = 0.9214
```

Below `d = 1.0` a single observer is worse than a coin. Above `d = 2.0` they are reliable. For comparison, a modern two-group design at 5 per cent two-sided significance and 80 per cent power needs `n = 2 (z_{0.975} + z_{0.80})^2 / d^2 = 15.698 / d^2` subjects per group:

```
d = 0.10  ->  n per group ~ 1570
d = 0.20  ->  n per group ~  392
d = 0.30  ->  n per group ~  174
d = 0.50  ->  n per group ~   63
d = 0.80  ->  n per group ~   25
d = 1.00  ->  n per group ~   16
d = 2.00  ->  n per group ~    4
```

The traditional instrument's own `k` is enormous, so it clears the *noise* floor easily. What it cannot clear is the *attribution* floor: distinguishing the effect from everything else that changed at the same time. That requires a control, and there is none.

Combining the two constraints yields the operating envelope. **A phenomenon is inside the old corpus's range if and only if it has all four of these properties:**

```
P1  LARGE          effect size well above the observer's own baseline variability
P2  REPEATABLE     producible on demand, so the observer can re-run it at will
P3  UNAIDED        observable without instruments
P4  PROMPT         short latency from intervention to effect, so that causal
                   attribution is available without a control group
```

And the complement, which is the more useful list:

```
OUTSIDE THE RANGE, ALWAYS
  - small effects (d below ~0.5): invisible to it
  - rare effects: never accumulate enough instances per observer
  - long-latency effects: attribution fails; correlation is unrecoverable
  - anything requiring a control: dose-response, base rates, differential effect
  - anything requiring an instrument: microscopy, spectroscopy, chemistry,
    statistics
  - hidden mechanism of any kind: the instrument sees the effect, never the channel
  - population-level anything: the instrument is first-person by construction
```

P1 through P4 predict, without consulting a single modern source, which parts of the old corpus will survive checking. Section 6 tests that prediction on four readings: the memory palace, the ritual cadence, the staged laboratory process and the correspondence table are all large, repeatable, unaided and prompt. The doctrine of signatures, planetary influence on metals and the four-element theory of composition are none of those things.

Note the last exclusion, because Section 7.4 turns on it: **a tradition that reports an effect is a witness to the effect and is not a witness to its cause.** That is not a rule imposed from outside. It follows from the specification sheet.

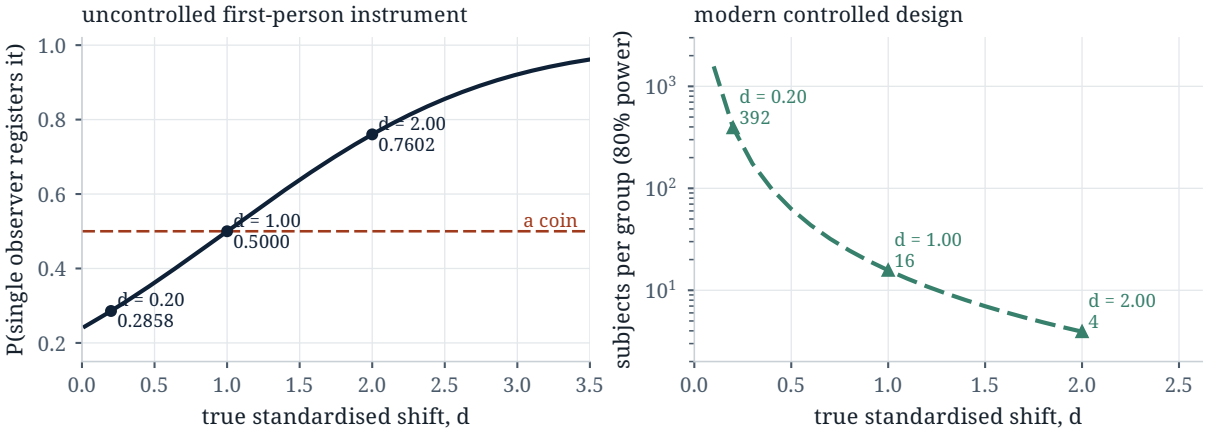


Figure 29.3: **The uncontrolled first-person instrument is worse than a coin below one standard deviation, and the modern apparatus that beats it is expensive.** Left: the probability that a single uncontrolled observer registers a true standardised shift `d` as a change exceeding their own day-to-day variation, `P = 1 - Phi((1-d)/sqrt2)`, against the 0.5 chance line. Right: subjects per group required for 80 per cent power at two-sided 5 per cent, `n = 15.698/d^2`, on a logarithmic axis. Plotted as small multiples rather than on shared axes because a probability and a sample size are not commensurable and a dual axis would be a chart error. Computed from the stated formulas; generated by `figs/src/CCP-029_fig03.py`.

29.5.4 5.4 The survivorship filter, and the fact that it has a sign

Everything above concerns what could be *observed*. A separate filter decides what gets *transmitted*, and it is the more dangerous of the two because it operates invisibly and after the fact.

We do not inherit the old corpus. We inherit what was preserved of it, and preservation selects for memorability and institutional interest, not for accuracy.

Model the filter. Let `a0` be the fraction of the original corpus that was accurate, `s_acc` the probability an accurate item survived to us, and `s_inacc` the probability an inaccurate item did. Then the accuracy of the corpus actually read is

```
a1 = a0 * s_acc / ( a0 * s_acc + (1 - a0) * s_inacc )
```

which is Bayes' theorem with survival as the observation. **The rates below are assumed and illustrative. What is not assumed is the direction of each row, which follows from the mechanism named beside it.**

a0	s_acc	s_inacc	surviving accuracy	factor	
0.30	0.02	0.10	0.0789	0.263	memorability-selecting
0.30	0.05	0.05	0.3000	1.000	neutral
0.30	0.20	0.02	0.8108	2.703	truth-selecting
0.30	0.40	0.01	0.9449	3.150	strongly truth-selecting
0.10	0.02	0.10	0.0217	0.217	memorability, poor original
0.10	0.40	0.01	0.8163	8.163	truth-selecting, poor original

Read the first and fourth rows together. **The same original corpus, run through two different preservation filters, arrives either 7.9 per cent accurate or 94.5 per cent accurate — a factor of 11.97, produced entirely by the filter and not at all by the source.**

The decisive question about any traditional claim is therefore not how old it is, how widespread it is, or how confidently it is stated. It is: what happened when the claim was wrong?

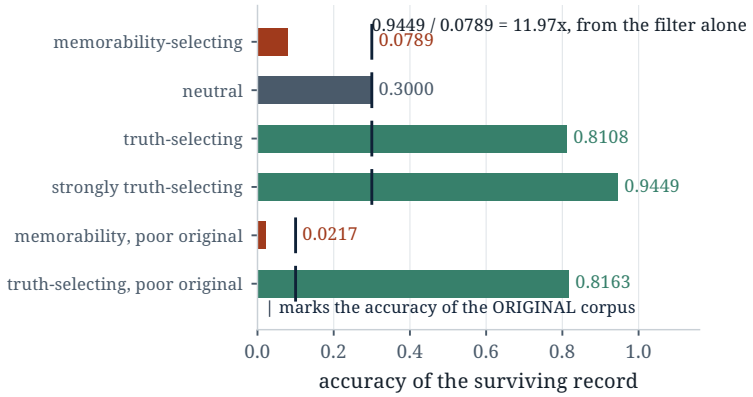


Figure 29.4: **The filter, not the source, decides what accuracy the record arrives with.** Surviving accuracy `a1 = a0 s_acc / (a0 s_acc + (1−a0) s_inacc)` for the six parameter rows of Section 5.4. The same original accuracy of 0.30 arrives at 0.0789 under a memorability-selecting filter and 0.9449 under a strongly truth-selecting one, a factor of 11.97; a poor original corpus at `a0 = 0.10` still arrives at 0.8163 if its failures were visible. The dashed line marks the original accuracy each pair started from. Modelled from the assumed survival rates stated in Section 5.4; generated by `figs/src/CCP-029_fig04.py`.

```
THE VISIBLE-FAILURE RULE

If a claim's failure was VISIBLE, PROMPT, and COSTLY to whoever relied on it,
the preservation filter selected FOR accuracy and the surviving material is
better than the original corpus.

If a claim's failure was INVISIBLE, DELAYED, or COSTLESS, the filter selected
for memorability, and the surviving material is WORSE than the original corpus.

TRUST THE OLD CORPUS IN PROPORTION TO HOW VISIBLY ITS CLAIMS FAILED.
```

That rule sorts the material immediately, along a line that has nothing to do with how mystical the material sounds.

Visible failure, therefore trustworthy: metallurgy — the alloy is brittle or it is not. Dyeing — the colour holds or it runs. Distillation — the apparatus yields or it does not. Navigation — landfall is made or it is not. Agriculture — the crop comes in or it does not. Memory technique — the list is recited correctly in front of an audience or it is not. Building — the vault stands or it falls.

Invisible failure, therefore untrustworthy: cosmology, the composition of matter, the mechanism behind any effect, the influence of remote bodies on terrestrial events, the meaning assigned to any correspondence, and every claim about what happens after death. Not because these are grander subjects, but because nothing happened to the person who got them wrong.

Note where the memory palace falls. It is a claim about technique whose failure is public, immediate and humiliating: a rhetor who forgets his speech in the forum has failed visibly. That is the strongest possible position on the filter, and it is why the method of loci is the best-preserved piece of practical psychology in the Western tradition.

29.5.5 5.5 The three preservation biases

The filter has three components, each leaving a different signature.

B1 · Memorability. Material survives transmission in proportion to how easily it is remembered and re-transmitted — which is to say in proportion to `eta`, the conversion of exposure into durable storage [Curzi, 2026a]. Rhyme, metre, numerical symmetry, striking imagery and narrative all raise `eta` substantially, and `eta` is completely orthogonal to truth. A well-turned false doctrine outcompetes an awkwardly stated true observation, every time, and the competition runs for centuries.

The signature of B1 is that the surviving material is unusually *tidy*. Sevens and twelves and fours. Perfect symmetries. Complete tables with no gaps. **Reality does not produce complete tables with no gaps; transmission does, because incomplete tables are harder to remember and get completed in transit.**

B2 · Institutional interest. Material survives in proportion to whether some institution paid for its copying, and that institution had purposes — liturgical, legal, medical, political, pedagogical. The surviving corpus is therefore not a sample of what people believed; it is a sample of what institutions funded. This is the same selection Scott [1998] documents for administrative records generally: what is legible to an institution is what the institution keeps.

The signature of B2 is over-representation of the doctrinal, the canonical and the authoritative, and under-representation of the dissenting, the local and the practical.

B3 · Copyability. Text copies; apparatus does not. Doctrine copies; tacit skill does not. **We inherit the theory and lose the technique — which is exactly backwards from what the specification sheet says the instrument was good at.**

This is the cruellest of the three and the least often stated. The old corpus was strongest at procedure — large, repeatable, unaided, prompt, visibly failing — and procedure is precisely what a manuscript preserves worst. The alchemical corpus is the clearest case: a great deal of the theory, which was wrong, and comparatively little of the bench practice, which was largely right. **The filter did not merely reduce the corpus's accuracy. It inverted the corpus's own ranking of its strengths.**

29.5.6 5.6 Survivorship stated at full strength

Six consequences follow, and each constrains how a precedent citation may be used.

1. **The record over-represents the doctrinal and under-represents the practical** (B2 and B3 together). The corpus's strongest material — procedure — is its most poorly transmitted.
2. **The record over-represents the tidy** (B1). Complete tables, round numbers, perfect symmetries. **A striking numerical pattern in a traditional source is evidence about transmission before it is evidence about the world.**
3. **Whole traditions are missing, and not at random.** Losses concentrate in the oral, the local, the suppressed and the linguistically isolated. Library destructions removed specific collections rather than random samples. **The traditions with the strongest claim to independence — the ones worth the most bits under Section 4.4 — are precisely the ones the record preserved worst, because independence and marginality are close to the same property.**
4. **The record has been retrospectively harmonised**, and this is the specific hazard for the precedent column. Later editors reconciled sources; early-modern compilers assembled Egyptian, Hellenistic, Jewish and Christian material into a single apparent tradition deliberately, as a scholarly and theological programme. **Apparent cross-tradition agreement is therefore frequently an artefact of editing, and Section 4.4's collapse applies at full force: at `c = 0.1`, five agreeing traditions buy 1.685 bits against the 7.925 that genuine independence would have bought — 6.24 bits lost to editorial tidiness.**
5. **Attribution is unreliable.** Pseudepigraphy was normal practice rather than fraud: attributing a work to an ancient authority was how a work claimed standing. A date attached to a name is generally a claim about authority, not about chronology, and the

independence audit must use the date of the earliest verifiable attestation of the specific claim. That requirement does the decisive work in the rejection of Section 7.5.

6. **Survivorship inside a tradition is as strong as survivorship between traditions.** Within any surviving corpus the successful practitioner's account survives and the unsuccessful practitioner's does not. **Every traditional literature therefore reads as though its methods work, and would read that way even if they did not.**

None of this makes the precedent column worthless. It makes the precedent column exactly as valuable as Section 4 priced it: a real term in the sum, a small one, and one that must be audited item by item rather than credited wholesale.

29.6 6 Four readings of traditional artefacts

Four readings. Each states what the artefact literally is, what the tradition got right, what it got wrong, and where the item falls on the P1–P4 envelope and the visible-failure rule. **The literal-reading rule used here — that a two-column table is a binary relation, a diagram of nodes and paths is an undirected graph, an assignment of labels to categories is a partition, and a letter-value scheme summed over a word is a hash — is the corpus's own, and the audit it supports elsewhere reaches a structurally parallel distribution of verdicts: seven MAPS, two SPLIT, one RENDERING and three REJECTED across thirteen entries [Curzi, 2026].** There is no inference in any of those readings, and none of them can be wrong in the way a theory can be wrong.

29.6.1 6.1 Correspondence systems as relational graphs

What the tradition got right. Three things, all engineering.

First, a correspondence table is a category assignment, written down, fixed, and made auditable. The surplus functional cannot be applied to anything until the `N` categories have been named [Curzi, 2026]. **The tradition that produced correspondence tables was performing, in a different notation, the mandatory first step of the procedure this framework specifies.** That is a description of what the artefact does, not a compliment.

The effective count that step produces in this framework is `N_eff = 1 + (N−1)u`, and **that quantity is named here explicitly because the framework uses more than one effective count and confusing them is a real hazard: `1 + (N−1)u` is a linear interpolation between 1 and `N` and is not a Hill number**, whereas `1/HHI` and `exp(H)` are Hill numbers of order two and one [Hill, 1973, Jost, 2006] and are the grounded objects the literature supports.

Second, they solved the addressing problem, and solved it as a designer would. A flat set of 78 elements costs `log2(78) = 6.285` bits to address, and every one of those bits must be memorised as an unstructured fact. Give the same set a product structure — 22 elements plus four suits of fourteen — and the address decomposes: one bit for which part, then `log2(22) = 4.459` or `log2(4) + log2(14) = 5.807`. **The structured address is slightly longer in bits and vastly shorter in effort**, because its components are independently memorable and independently checkable. That is a real information-theoretic trade [Shannon, 1948, Cover and Thomas, 2006] and a filing-system designer makes the same choice for the same reason. The relevant human constraint is the small span of immediate memory [Miller, 1956], which is what makes an unstructured 6.3-bit address expensive and a two-part address cheap.

Third, they built redundancy into the index. Each element carries several attributes — a number, a letter, a colour, a metal, a direction, a day — and any one retrieves the element. This is a multiply-keyed index, and its purpose is what a multiply-keyed index's purpose always is: retrieval survives the loss of any single key. In storage-law terms, redundant keying raises `eta` and lowers `delta` on the retrieval path [Curzi, 2026a].

What the tradition got wrong, and it is one error repeated everywhere. The tables were read as causal graphs rather than as indexes. That an element is filed under a planet is a fact

about the filing system; it was routinely taken as a fact about the planet. **This is the single most consequential error in the old corpus, and it has a name: the category error, treating an index as a cause.** The index is real, useful and precisely specifiable. The causal reading attached to it is none of those.

Envelope check. As a mnemonic and retrieval technology: P1 large, P2 repeatable, P3 unaided, P4 prompt — inside the envelope, and visible failure, since a mnemonic system that does not retrieve fails in front of witnesses. As a causal claim about planets and metals: P4 fails outright, failure is invisible, and the material accordingly did not survive contact with measurement.

29.6.2 6.2 Memory palaces as spatial storage

What the artefact literally is. An addressing scheme that uses a familiar spatial sequence as the key space. Items are bound to loci; retrieval is a walk. The technique is attested from classical antiquity — the *Rhetorica ad Herennium* is the standard early source — and its transmission through the medieval and early-modern periods, where it became entangled with the correspondence systems of Section 6.1, is traced in Yates's *The Art of Memory* (1966). **That work carries no bibliography key here; see Section 3.6 and Appendix B.**

The storage-law reading, and it is a category-2 borrowing (Section 3.4): the technique is a joint intervention on both parameters of the same first-order difference equation [Curzi, 2026a, Elaydi, 2005].

```
Q_{t+1} = (1 - delta) Q_t + eta S_t - C_t

eta RAISED     encoding is elaborative -- the item must be transformed into a
               vivid image and bound to a place, which is far more processing
               than rehearsal
delta LOWERED  the retention medium is spatial memory, which is durable, and the
               sequence of loci is itself already stored and does not decay with
               the material bound to it
C RAISED       building the palace is real up-front cost
```

The tradition's own doctrine — invest heavily up front in the image and retrieval is free — is the statement that raising `eta` at the cost of raising `C` is worth it whenever the retention horizon is long. That conclusion was arrived at by rhetoricians who had to deliver a two-hour speech from memory.

Illustrative parameterisation, badged: the constants below are assumed, not measured, and are chosen to show the shape of the trade rather than to quantify anyone's memory. A single study session with `Q_0 = eta S` at `S = 1`, then free decay:

```
rote repetition   eta = 0.20   delta = 0.35/day   half-life  1.61 days
method of loci    eta = 0.70   delta = 0.06/day   half-life 11.20 days

    day        loci        rote      ratio
      0      0.7000    0.200000        3.5x
      1      0.6580    0.130000        5.1x
      3      0.5814    0.054925       10.6x
      7      0.4539    0.009804       46.3x
     14      0.2944    0.000481      612.5x
     30      0.1094    0.000000    224,131x
```

The structure of the advantage is the point. At day zero the loci route is only 3.5 times better; by day fourteen it is six hundred times better. **The gain is almost entirely in `delta`, not in `eta`, and it compounds.** Raising `eta` alone gives a fixed multiplier at every horizon; lowering

`delta` gives a multiplier that grows with the horizon. Beyond about thirty days the pure exponential model stops being meaningful — real retention curves are not single exponentials and flatten into a long tail — so the day-30 ratio should be read as "the rote trace is gone" rather than as a number. **That is a limit of the model's domain and is stated as one** (Section 13.6).

The evidence column. This is one of the rare places where the old corpus's own technique has been taken into the laboratory and held up: superior memorisers are distinguished by spatial mnemonic strategy rather than by general ability or gross brain structure, and training naive subjects in the technique produces large and durable gains in recall. **The two studies concerned are named in Appendix B without keys** (Section 3.6), and **no effect size is quoted from either, because none was verified for this paper**; the direction and the qualitative finding are what is asserted.

Envelope check. P1 large, P2 repeatable, P3 unaided, P4 prompt — all four. Visible failure: maximal. **This is the best-positioned item in the whole corpus on every criterion Section 5 defines, and it is correspondingly the item that survived checking most completely. That is not a coincidence; it is the prediction of Sections 5.3 and 5.4 coming true.**

What the tradition got wrong. One thing, arriving late: the early-modern fusion of the palace with correspondence and cosmological schemes, on the theory that a palace built to mirror the structure of the universe would confer understanding of the universe rather than merely recall of its contents. It does not. The technique stores what is put in it. **Addressing adds no information** — the same constraint that governs the addressing architecture of the neighbouring paper [Curzi, 2026x], where targeting and editing are shown to be separable functions and the address carries no cargo of its own.

29.6.3 6.3 Ritual as periodic forcing

What the artefact literally is. A fixed observance repeated on a fixed cadence is a periodic forcing function applied to a system with retention. **This is a category-2 borrowing** and the result is a theorem about the difference equation, not about liturgy.

For a linear time-invariant store driven with a fixed long-run average rate `R`, the long-run average of `Q` is the DC gain times the mean input [Curzi, 2026a, Oppenheim et al., 1996]:

```
mean Q = R / delta       EXACTLY, and independently of how the input is scheduled.
```

That is not an approximation, and it kills the naive version of the cadence claim before any bridge is built. **Spreading a fixed total input evenly over time does not raise the average level of the store at all. What it changes is the ripple — and therefore the minimum.**

Computed in periodic steady state, dosing `R T` units every `T` periods so the long-run rate is `R = 1.0` per period, with `Q_max = RT/(1−(1−delta)^T)` and `Q_min = Q_max (1−delta)^{T−1}`:

```
delta = 0.10 ; mean Q = 10.000 in every row, exactly

  T    dose     Q_max    Q_min      mean    Q_max/Q_min   Q_min/mean
  1     1.0    10.000   10.000    10.000          1.000       1.0000
  2     2.0    10.526    9.474    10.000          1.111       0.9474
  3     3.0    11.070    8.967    10.000          1.235       0.8967
  7     7.0    13.418    7.131    10.000          1.882       0.7131
 14    14.0    18.153    4.614    10.000          3.934       0.4614
 30    30.0    31.328    1.476    10.000         21.231       0.1476
 90    90.0    90.007    0.008    10.000      11814.536       0.0008
```

At a ninety-period interval the store spends nearly the whole cycle at essentially zero and then spikes to nine times its own mean, and the average is identical to the daily case. **Everything cadence buys is bought at the floor, and the floor is what matters whenever there is a threshold.** Fraction of the cycle at or above a threshold, same system:

```
delta = 0.10, mean Q = 10.0
threshold 5.0:   T=1 -> 1.000  T=7 -> 1.000  T=30 -> 0.600  T=90 -> 0.311
threshold 8.0:   T=1 -> 1.000  T=7 -> 0.714  T=30 -> 0.433  T=90 -> 0.256
```

What the tradition got right. Two things beyond the cadence itself.

First, invariance of the input. Ritual is characterised by resistance to modification, and the anthropological account that makes invariance constitutive of the form rather than incidental to it is Rappaport's *Ritual and Religion in the Making of Humanity* (1999), **named without a key; see Appendix B.** In the storage-law reading this is a second and distinct engineering decision: **the tradition applied a low `delta` not only to the store but to the forcing function itself.** A drifting input to a low-pass filter produces a drifting fixed point; an invariant input produces a stable one. The traditions that survived are the ones whose input did not drift, and the mechanism used to prevent drift — declaring the form unmodifiable — is the right mechanism for the problem.

Second, the doctrine that a small observance kept beats a large one abandoned. That is the floor result of the table above, stated in one sentence several centuries before anyone wrote down a transfer function.

What the tradition got wrong. The account of *why*. Every tradition supplied a mechanism for its cadence; the mechanisms are cosmological and are the part that did not survive. **The schedule was right and the reason given for it was wrong, which is the characteristic signature of a P1–P4 phenomenon carried by an instrument that cannot see mechanism.**

Envelope check. The cadence itself: P1 large, P2 repeatable, P3 unaided, and P4 only partially, since the effect of lapsing is delayed by the store's time constant — which is why the cadence took centuries to settle rather than a generation. Visible failure: moderate. An institution whose cadence is wrong decays, but slowly, and the decay is attributed to other things.

29.6.4 6.4 Alchemical operations as staged process control

This is the reading that most rewards being taken literally, and the one where the survivorship filter did the most damage.

What the artefact literally is. The alchemical operations — calcination, dissolution, separation, conjunction, putrefaction or fermentation, distillation, coagulation, in the seven-stage presentations — are unit operations, the same unit operations that appear in a chemical engineering textbook under mostly the same names, because chemistry inherited them directly and without interruption. So is the apparatus: the alembic, the retort, the *balneum mariae* — the water bath, still standard laboratory equipment.

Read as process control, the corpus contains four things that are simply correct.

1. A process has ordered stages, and the order is not permutable. This is a real and nontrivial engineering fact carrying real information, and the quantity is computable:

```
n stages -> n! orderings -> log2(n!) bits to specify the correct one

 3 stages ->             6 orderings ->  2.585 bits
 4 stages ->            24 orderings ->  4.585 bits
 5 stages ->           120 orderings ->  6.907 bits
 7 stages ->         5,040 orderings -> 12.299 bits
12 stages ->   479,001,600 orderings -> 28.835 bits
```

A correct seven-stage sequence is 12.299 bits of hard-won process knowledge that exists nowhere in the individual stages and can only be obtained by failing at the other 5,039 orderings. That is the part the tradition preserved accurately, and it is exactly the part a modern reader dismissing the corpus as superstition throws away.

2. Temperature must be controlled and graded. The obsessive attention to degrees of heat, to the water bath as a means of holding a bounded temperature, and to the difference between the heats appropriate to different stages, is correct process engineering.

3. Time under condition is a process variable. Long digestions at low heat produce results short intense heating does not. This is the same statement as Section 6.3's, arrived at independently in the laboratory rather than the chapel, **and its independence is worth noting, because separate arrivals are the one thing the precedent column actually pays for** (Section 4.4).

4. Purification is iterative and has diminishing returns. Repeated distillation improves purity by a decreasing increment per pass, and the concavity is the same shape as the framework's other laws [Curzi, 2026].

What the tradition got wrong, and by how much. The compositional theory. The four elements and the three principles are not a correct account of what matter is made of, and the transmutation of base metal into gold is not achievable by any chemical operation. The size of that error is computable, and computing it is more useful than sneering at it:

```
nuclear binding energy per nucleon   ~  8 MeV  =  8,000,000 eV
typical covalent bond                ~  4 eV

ratio  =  2,000,000        log10 ratio = 6.30
```

The alchemists were attempting a nuclear operation with chemical tools and were short by a factor of two million. No furnace, no reagent, no duration and no number of repetitions could have closed that gap.

Two honest riders. **First, the goal was not absurd, only unreachable by the means available:** transmutation of one element into another is real physics, performed routinely at energies six orders of magnitude above chemistry and at costs that make the product worthless as a commodity. **Second, the failure of this claim was invisible in the short term and delayed in the long term** — a practitioner could always suppose the process incomplete — which is precisely the condition Section 5.4 identifies as producing preservation without correction. The claim that was wrong is the claim whose failure could not be seen, and that is not a coincidence either.

Envelope check. The unit operations: P1 large, P2 repeatable, P3 unaided, P4 prompt; visible failure maximal — the distillate is pure or it is not. Inside the envelope, filter positive, survived. The compositional theory: P4 fails, failure invisible, filter negative, did not survive.

And here is the survivorship inversion in its purest form. The part that was right — the bench practice, the graded heats, the sequence, the feel for when a stage is complete — is tacit skill, and tacit skill does not copy (bias B3). The part that was wrong — the theory — is text, and text copies. **The wrong half was inherited at high fidelity and the right half at low fidelity, and the difference has nothing to do with which half was true.**

29.6.5 6.5 The error taxonomy, with a control case

Five error types account for most of the old corpus's failures.

T1 · Index read as cause. Section 6.1. The commonest and the most consequential.

T2 · Effect attributed to the wrong channel. The observation is genuine and the mechanism assigned is not the operating one. Section 6.3's cosmological accounts of cadence.

T3 · Scale over-extension. A relation holding locally is asserted universally. The four elements are a serviceable *phenomenological* classification of the states of matter available to unaided observation and become false the moment they are asserted as a compositional theory.

T4 · Completion of the table. Bias B1 with a name. Gaps in a scheme are filled to make it memorable, and the filled entries are indistinguishable in the surviving text from the observed ones. **Every complete symmetric table in the old corpus should be assumed to contain interpolated entries, and there is generally no way to tell which.**

T5 · The unfalsifiable escape. A claim is stated so that failure is always attributable to the practitioner rather than to the claim. This is the error that keeps all the others alive, because it disables the visible-failure filter that would otherwise have removed them. **T5 converts a truth-selecting filter into a memorability-selecting one — it moves an item from the 94.5 per cent row of Section 5.4's table to the 7.9 per cent row.** It is the same manoeuvre Popper [1959] and Lakatos [1970] identify as the protective belt around a degenerating programme.

The control case: the doctrine of signatures. The doctrine holds that the appearance of a plant indicates its medicinal use — that a plant resembling an organ treats that organ. It was elaborated by Paracelsus and given systematic presentation in Giambattista della Porta's *Phytognomonica* (1588), and it persisted in herbal writing for centuries. Run through this paper's apparatus:

```
P1 LARGE        NO   the claimed effects are mostly small and slow
P2 REPEATABLE   NO   illness does not recur on demand
P3 UNAIDED      YES  the resemblance is visible; the only criterion it passes
P4 PROMPT       NO   herbal effects are delayed and confounded

VISIBLE FAILURE?     NO. A patient who does not recover is attributed to severity,
                     timing, preparation, or fate. T5 is available at every failure.

MEMORABILITY         MAXIMAL. A single rule that generates its own content,
                     requires no data, is instantly teachable, and is delightful.

PREDICTED FATE       preserved for centuries, and false.
ACTUAL FATE          preserved for centuries, and false.
```

The doctrine of signatures is the control experiment for this paper's method: an item scoring maximum on memorability and minimum on every accuracy-selecting criterion, which did exactly what the model predicts. It is also why the model cannot be cheated. The doctrine has famous apparent hits — the walnut and the brain, lungwort and the lung — and each is a false positive produced by a rule that generates a resemblance for everything and is therefore never wrong in advance and never right by design.

29.7 7 Four bridges worked to their verdicts

29.7.1 7.1 The mapping procedure

Five inherited steps, with three added by this paper's refinement. The added steps are marked.

```
  STEP 1   IDENTIFY CORE STRUCTURES on each side: mechanisms, conservation laws and
           measured relations on the (+1) side; practices, stages, cosmological
           elements and recorded phenomenology on the (−1) side.
  STEP 2   LOOK FOR STRUCTURAL MATCHES. Not the words -- the structures.
  STEP 3   TEST THE MAPPING with the three-state validation.
  STEP 4   GENERATE PREDICTIONS on both sides.
  STEP 5   TEST PREDICTIONS. Design the experiment or the practice; revise on results.

[ STEP 0   FIX THE (+1) READING FIRST, IN WRITING, BEFORE OPENING THE (−1) COLUMN.  ]
[          A large enough corpus confirms anything asked of it.            ADDED    ]

[ STEP 2b  GRADE THE CORRESPONDENCE: isomorphism / homomorphism / analogy /         ]
[          metaphor. Below homomorphism, stop and record the rejection.    ADDED    ]

[ STEP 6   AUDIT INDEPENDENCE AND ASSIGN COLUMN ROLES IN WRITING. Name the          ]
[          tradition, the date, and the transmission path; estimate c.     ADDED    ]
```

Step 6 is the one that gets skipped, and skipping it is what produces work that reads well and cannot be checked. It costs one paragraph per bridge. Appendix A is, in effect, forty-three executions of Step 6.

29.7.2 7.2 The seven errors

Five errors are inherited; two are added here because the biological material invites them specifically.

E1 · False cognates. Words that sound the same and mean different things. *Energy* in physics is a conserved quantity with units of joules; *energy* in most contemporary spiritual usage is a mood word. **The most reliable indicator of a false cognate is that one side of the pair has units and the other does not**, which is why the dimensional audit of Section 13.1 is the first check a reader should run on any claim here.

E2 · Confirmation bias. Counting the hits and not the misses. Step 0 exists to prevent it, and the modern statement of the same hazard is the multiplicity of undeclared analytic choices [Gelman and Loken, 2013, Simmons et al., 2011].

E3 · Category error. Mapping across incompatible categories — treating consciousness as a substance, treating a probability as a force, treating an index as a cause.

E4 · Reductionism. Claiming one side is "nothing but" the other. A valid mapping preserves the integrity of both descriptions; the general argument that higher-level descriptions are not eliminable is Anderson [1972]'s.

E5 · Mystification. Using the language of one side to obscure rather than illuminate. The test is that a correspondence that cannot be stated clearly may not be valid.

E6 · Scale transfer without a scale check. ADDED. A mechanism real at one scale is imported to another where its magnitude is negligible. This is the commonest error in writing about physics in biology, and it is always caught by computing one ratio. The Casimir pressure between conducting plates falls as `a^-4` [Casimir, 1948], and has been measured at the separations where it is appreciable [Lamoreaux, 1997, Mohideen and Roy, 1998]:

```
1.3e5 Pa at 10 nm   ->   13 Pa at 100 nm   ->   1.3e-3 Pa at 1 um
                                            ->   1.3e-15 Pa at 1 mm
```

Twenty orders of magnitude across an interval a careless sentence would cross without noticing. The corpus's own treatment of the speculative physical margin reaches the same verdict by a different route, reporting a decoherence time of `1.98e-17 s` for a warm wet 310 K system against a 1 ms neural event — a shortfall of `5.04e13`, and `5.04e7` even granting a `1e6` protection factor [Curzi, 2026aa].

E7 · Attributing a real effect to the wrong channel. ADDED. The effect is genuine and measured; the mechanism assigned is not the one operating. **This error is more dangerous than a plain false claim, because the observation keeps coming out right and keeps appearing to confirm the wrong mechanism.** Section 7.4 is a worked instance. **The diagnostic is always the same: name the competing channels and state the experiment that would separate them.**

29.7.3 7.3 Bridge 1 — cadence. Worked, and ESTABLISHED

Step 0. Fix the evidence reading first. Section 6.3's result: for a decaying store feeding a threshold, the scheduling of a fixed total input controls the fraction of time the threshold is met, and short intervals relative to the time constant dominate long ones. **Intensity buys the peak; cadence buys the floor.** That is a theorem about the model.

There is also a measured empirical result of the same shape, and it is among the most replicated findings in experimental psychology: **the spacing effect.** The retention decay curve was described by Ebbinghaus in 1885, and the modern quantitative synthesis of the distributed-practice literature finds that spaced presentations reliably outperform massed ones for the

same total study, and that the optimal gap grows with the retention interval being targeted. **That last clause is the model's own prediction: the optimal interval scales with the time constant.** Both works are named in Appendix B without keys (Section 3.6).

Steps 1–2. Core structures and the match.

```
(+1) side                              (−1) side
─────────                              ─────────
a store with retention (1 − delta)     an observance held to decay if lapsed
an input applied at intervals T        a fixed office at fixed hours
a threshold on Q                       a required standard of practice
optimal T scales with 1/delta          intervals fixed by day, week, and year
```

Step 2b. Grade it. Homomorphism. Not isomorphism: the traditions did not write down a decay constant and did not derive an interval. They report the *schedule* and the *consequence of lapsing*, which is two of the four structural elements, mapped without residue. Homomorphism clears the threshold.

Step 3. The precedent reading, opened only now, stated as specifically as the independence rule requires: the Rule of St Benedict, c. 530 CE, prescribes eight daily offices at fixed hours and treats the cadence rather than the intensity of any single office as load-bearing; Jewish daily prayer is structured in three fixed services, and the seven-day sabbath cycle is among the oldest continuously observed fixed cadences in the record; Islamic *salat* is five daily prayers at times fixed by the sun; Vedic *sandhyavandanam* is performed at the junctions of the day; Buddhist monastic codes fix a daily schedule and a fortnightly recitation of the rule.

Independence audit. These are not independent. The Abrahamic three share a lineage directly, and Indo-European and Near Eastern contact is continuous across the period. Realistically `c` is well above 0.3, so by Section 4.4 the whole precedent column here is worth approximately 1.6 bits — one tradition's worth. **It is stated in full and credited at one tradition's weight, because that is what it is worth.**

What it does record, and what no evidence source records, is the operational detail: which hours, what to do when an office is missed, how a lapse is repaired, and the explicit doctrine that a short observance kept is worth more than a long one abandoned. **That is the "how to do it" payload of Section 4.5.**

Step 4. The prediction. For a practice whose benefit decays with time constant `tau = 1/delta`, moving from interval `T` to interval `T/2` at the same total input raises the floor and leaves the mean untouched; and the effect of interval on outcome should be large where a threshold is involved and near zero where it is not. That is falsifiable and it has been tested in the memory domain, where it holds.

Step 5 / 0p. The contra-proof. The two integrate cleanly. The evidence side does not have to claim that Benedict computed a time constant; the precedent side does not have to claim its offices were about memory retention. Both describe the same structural fact — a decaying store driven periodically — from opposite ends, and neither has to swallow a falsehood to sit beside the other.

```
C1 STRUCTURAL CORRESPONDENCE     PASS   homomorphism; four structures, three map
C2 PREDICTIVE POWER              PASS   optimal interval scales with tau; tested
C3 PRACTICAL UTILITY             PASS   sets cadence on the evidence side; explains
                                        the tradition's own emphasis on the other

STATUS: ESTABLISHED.
COLUMN ROLES: (+1) evidence -- the distributed-practice synthesis, and the
                   storage law's periodic-forcing result (CCP-007).
              (−1) precedent, ~1.6 bits, owing to high shared-source probability.
```

What is established is the structural bridge, not a historical claim. This paper does not assert that liturgical hours were tuned for retention; that is an unfalsifiable statement about

the intentions of people who left no record of intending it. What is asserted is that the schedule they kept has the form the model recommends.

29.7.4 7.4 Bridge 2 — vibration as cellular input. Worked, and GRADED

This bridge demonstrates the most useful thing the method does, which is not accepting or rejecting but *grading*: separating the part of a bridge that holds from the part that does not, at the exact seam.

Step 0. Fix the evidence reading first. Mechanotransduction is real, established and prize-recognised. Mechanically gated cation channels of the PIEZO family were identified as essential components of mechanically activated currents (Coste and colleagues, *Science*, 2010), and the 2021 Nobel Prize in Physiology or Medicine was awarded for receptors for touch and temperature. Force applied at the cell surface transmits through the cytoskeleton to the nucleus, stretches chromatin and upregulates transcription (Tajik and colleagues, *Nature Materials*, 2016). Ultrasound has been used to activate specified channels — sonogenetics — in *C. elegans* (Ibsen and colleagues, *Nature Communications*, 2015) and with a human channel (Duque and colleagues, 2022). Cochlear hair cells are biology's own existence proof that exquisitely sensitive mechanical reception is buildable in tissue. **None of those works carries a citation key here; all are named in Appendix B, per Section 3.6.**

Sound is mechanical. There is therefore a real, established pathway by which a pressure wave becomes a cellular signal, and that is a strong evidence reading.

Steps 1–2b. The candidate bridge and its grade. Candidate: *sustained vocal and percussive practice acts on the body through mechanotransduction.* On the evidence side: an oscillating mechanical stimulus, a gated channel with a force threshold, an ionic current, a downstream cascade, a transcriptional response. On the precedent side: sustained tone production, a required duration, a reported change in bodily and mental state, a transmitted technique.

Match at the coarse grain — mechanical oscillation is an input the body converts into a state change — is homomorphic and clears the threshold. Match at the fine grain — this particular channel family mediates this particular practice — has no counterpart at all on the precedent side, because that side records no mechanism. It records a state change.

Step 3. The precedent reading. Vedic recitation, transmitted orally with exacting phonetic discipline over millennia; Christian plainchant under the canonical hours; Tibetan Buddhist overtone chanting; Sufi *dhikr*; West African and Afro-diasporic drumming traditions; Australian Aboriginal drone practice. **These are substantially more independent than the cadence set of Bridge 1** — the Australian and West African instances in particular have no plausible shared source with the Indo-European ones — so c is genuinely lower and the precedent is worth more. **It is still precedent: on the toy model, of order two to three bits, against roughly twenty for the prize-grade evidence.**

What every one of them records is the same coarse claim: sustained vocalisation at length changes bodily state reproducibly. **None records a channel, a current, or a transcript.**

Steps 4–5. The prediction, and the competing channel. Here the method earns its keep, because there is a measured result assigning the effect to a different mechanism entirely. Recitation of the Latin rosary and of a yoga mantra both brought respiration to approximately six breaths per minute — a period of $10.0\ \mathrm{s}$, a rate of $0.100\ \mathrm{Hz}$ — and enhanced baroreflex sensitivity (Bernardi and colleagues, *BMJ*, 2001; named in Appendix B). **The dominant measured channel for chant is the breathing pattern it enforces**, acting on autonomic cardiovascular control. Add the obvious further channels: the social setting, the attentional demand, the postural requirement, and expectancy [Rosenthal and Jacobson, 1968].

The honest statement of the bridge therefore splits at a seam, and the seam is exactly where error E7 lives:

- **ESTABLISHED:** mechanical force is a real input to cellular signalling and to transcription. This is evidence, and it is strong.

- **ESTABLISHED:** sustained chant produces measured autonomic effects, and the identified channel is respiratory entrainment.
- **SPECULATIVE — and it is marked speculative in this sentence, which is where the claim is made:** that the state changes traditions attribute to chant are mediated to any material degree by PIEZO-family mechanotransduction is a speculative bridge, unsupported by any measurement known to this author, and in competition with at least four better-supported channels.

Step 6. Column roles. Evidence: mechanotransduction, very strong; respiratory entrainment, strong. Precedent: broad, partially independent, worth a few bits — and it attests the *phenomenon*, not the *mechanism*. **A tradition that reports an effect is a witness to the effect and is not a witness to its cause, because nothing in the tradition's instrument could have distinguished causes.** That sentence disposes of a large fraction of the adjacent literature.

The separating experiment, stated so it can be run. Hold the acoustic exposure constant and vary the breathing pattern independently — deliver the same acoustic stimulus to a passive listener whose respiration is paced to a different rate, against a vocaliser producing it. **If the effect survives removal of the breathing pattern, the mechanical channel is in play. If it does not, error E7 is confirmed and the bridge is rejected at the fine grain.**

```
C1 STRUCTURAL CORRESPONDENCE   PASS at coarse grain (homomorphism)
                               FAIL at fine grain (-1 records no mechanism)
C2 PREDICTIVE POWER            PASS -- the separating experiment above
C3 PRACTICAL UTILITY           PASS

STATUS: ESTABLISHED at the coarse grain. SPECULATIVE at the fine grain, so marked.
```

29.7.5 7.5 Bridge 3 — "528 Hz repairs DNA". Worked, and REJECTED

A method that has never rejected anything is not a method. This bridge is rejected here at length, in print, and it is chosen because it sits one inch from the corpus's own subject matter. This series is about sound and about genomic material. If the framework cannot say clearly why the most popular claim in that intersection fails, it has no business making claims of its own.

Step 0. Fix the evidence reading first — and there is nothing to fix. The claim is that a tone at 528 Hz repairs, restores or otherwise beneficially modifies DNA. **Stated with the uncertainty on the record: no systematic survey of the literature was conducted for this paper and it is not asserted that no relevant publication exists; what is asserted is that no replicated, peer-reviewed result establishing any such effect is known to this author, and the claim is treated accordingly.** The absence of an evidence entry is dispositive under the binding rule of Section 3.1 — with no evidence column there is no established bridge, whatever the precedent column says. But the bridge fails for four further reasons, each computable, and computing them is more useful than citing an absence.

Failure 1 — the precedent column is not old, and fails the precedent test outright. The nine "Solfeggio frequencies" (174, 285, 396, 417, 528, 639, 741, 852, 963 Hz) are routinely presented as an ancient system. They are not. The medieval material is real: Guido of Arezzo's eleventh-century solmisation took its syllables from the hymn *Ut queant laxis*. But medieval solmisation assigns no frequencies to anything. **It is a system of relative intervals, taught on a movable hand, in a period with no absolute pitch standard** [Barbour, 1951, Lindley and Turner-Smith, 1993]. The specific numerical values were produced in the 1990s by a numerological reduction applied to a passage in the Book of Numbers, and were popularised in Horowitz and Puleo's *Healing Codes for the Biological Apocalypse* (1999).

A precedent entry must carry independent arrival and preservation. This one carries a 1990s construction wearing a medieval costume. It buys zero bits — not few, zero.

Failure 2 — the number is not well-defined without a reference pitch, and none is given. A frequency in hertz for a named musical pitch is a convention. 528 Hz is approximately C5

under a tuning near A4 = 440 Hz and is a different note entirely under A4 = 432 Hz, A4 = 415 Hz, or any of the pitch standards actually in use across the period the claim invokes [Barbour, 1951]. **A claim that a specific molecular process responds to a specific hertz value must specify the standard, and if the value exists only under a convention adopted in the twentieth century, the claim cannot also be ancient.** This is a self-contradiction internal to the claim and requires no measurement to expose.

Failure 3 — the spatial scale is wrong by nine orders of magnitude. This is the decisive computation. A propagating wave can only impose spatial structure at scales comparable to its wavelength. In soft tissue, taking the speed of sound as approximately 1540 m/s [Rossing et al., 2002]:

```
lambda(528 Hz, soft tissue) = 1540 / 528 = 2.9167 m

ratio to a cell      (10 um)    = 2.917e5
ratio to a nucleus   (6 um)     = 4.861e5
ratio to a base pair (0.34 nm)  = 8.578e9
```

In air, at 343 m/s, the wavelength is 0.6496 m and the ratio to a base pair is `1.911e9`.

The wave is roughly nine billion times longer than the object it is claimed to act on selectively. Across an entire nucleus the pressure field is, to an extraordinarily good approximation, spatially uniform. There is no gradient at the scale of a gene, a codon or a base pair. Whatever a 528 Hz field does to a nucleus it does to the whole nucleus at once, identically, and it therefore cannot select a sequence. **This is not an argument that acoustic fields do nothing to cells — Bridge 2 established that they do. It is an argument that they cannot do anything *sequence-specific* by direct spatial action**, which is the gap the companion papers are devoted to [Curzi, 2026x,y].

Failure 4 — the energy scale is wrong by eleven orders of magnitude, and the frequency scale by thirty-five octaves. Two independent numbers, both computed. A single acoustic quantum at 528 Hz carries `h f = 2.18e-12 eV`; body-temperature thermal energy is `kT = 0.0267 eV` at 310 K:

```
h f / kT   at 528 Hz, 310 K   =   8.17e-11
```

The quantum of a 528 Hz field is about one hundred-billionth of the ambient thermal jostling. Sound acts on cells through bulk pressure and shear, not through quanta — which is exactly why Bridge 2's channels are *mechanically* gated. The quantum calculation matters because "frequency-specific molecular resonance" claims are implicitly quantum claims, and this is the number that ends them. Actual molecular vibrational modes sit in the infrared:

```
 500 cm^-1  = 1.499e13 Hz   ->  34.7 octaves above 528 Hz
1000 cm^-1  = 2.998e13 Hz   ->  35.7 octaves above 528 Hz
1700 cm^-1  = 5.096e13 Hz   ->  36.5 octaves above 528 Hz
3000 cm^-1  = 8.994e13 Hz   ->  37.3 octaves above 528 Hz
```

Molecular vibration is thirty-five to thirty-seven octaves above the audible band. An audible tone is not a molecular resonance and cannot be one.

```
C1 STRUCTURAL CORRESPONDENCE   FAIL   metaphor only; no shared structure survives
                                      the scale check; wavelength ratio ~9e9
C2 PREDICTIVE POWER            FAIL   generates no prediction surviving Failure 3
C3 PRACTICAL UTILITY           FAIL   a practice built on a false channel is not
                                      utility

(+1) column                    EMPTY  (uncertainty stated; no survey claimed)
(-1) column                    VOID   the source is 1990s, not traditional; 0 bits

STATUS: REJECTED.
```

The part that matters is what the rejection costs this framework, and the framework must accept the binding rather than exempt itself. The corpus describes a genomic-audio pipeline that sonifies measured infrared spectra by fixed octave transposition [Curzi, 2026y]. The 528 Hz rejection means that pipeline may not claim what it has just refused to let anyone else claim:

1. **A transposed tone is not the molecular resonance.** The transposition is a stated, fixed, arbitrary-by-design mapping; thirty-five octaves is a factor of `2^35`, and dividing a frequency by `2^35` preserves no physical interaction, only a ratio. Octave equivalence is a fact about pitch class [Shepard, 1982] and about perception [Demany et al., 2013], not about molecular physics. The audible tone stands in a documented numerical relation to a measured spectrum and in **no causal relation to the molecule**.
2. **The framework's essence frequencies are essence frequencies, not physical resonances.** The constant `k` is *defined*; attaching hertz to it is a separate and stronger claim and is POSTULATED (Section 3.4), and the standard scepticism about golden-section claims is cited alongside it [Markowsky, 1992].
3. **Nothing in this series claims sequence-specific acoustic action on genomic material**, because Failure 3 applies to the framework's own tones precisely as it applies to anyone else's.

The rejected claim and this framework's own method differ in one respect only, and it is the only respect that matters: this framework states its mapping, states that the mapping is a convention, and states that the convention carries no mechanism. The rejected claim states a mechanism it does not have. **If a reader concludes that the difference is thin, the reader has understood the situation correctly. It is thin. It is also the entire distinction between an honest notation and a false one.**

29.7.6 7.6 Bridge 4 — religious fasting and cellular autophagy. Worked, and SPECULATIVE

A fourth bridge is carried end to end with every form filled in, so that the procedure is available as a template rather than as a description, and because it lands on the third verdict.

Step 0. Autophagy is a conserved cellular process in which the cell degrades and recycles its own components. The molecular machinery was identified in yeast by Yoshinori Ohsumi, who received the 2016 Nobel Prize in Physiology or Medicine for that work; autophagy is upregulated by nutrient deprivation, well established in yeast and in mammalian cell and animal models. **Named without a key; see Appendix B.** *What is measured:* the identification of the genes and the pathway, and the response of the pathway to nutrient deprivation in model systems. **What is not established with the same confidence: the time course, magnitude and tissue distribution of fasting-induced autophagy in living humans under the specific schedules religious observance actually prescribes.** Measuring autophagic flux in human tissue in vivo is genuinely difficult, and no settled quantitative answer is asserted here. **That distinction — solid mechanism, unsettled human dose-response — is fixed before anything else is consulted, and it determines the verdict.**

Steps 1–2b.

(+1) side	(−1) side
nutrient sensing (mTOR, AMPK)	[no counterpart]
a recycling process under deprivation	[no counterpart]
a threshold of deprivation duration	a prescribed fast length
a refeeding phase that halts the process	a prescribed breaking of the fast
a repeatable cycle	an annual or weekly calendar

Two of five structures map: duration threshold and cycle. Three do not, because the precedent side records no cellular process of any kind. **Grade: homomorphism, weakly.** It clears the threshold, but only just, and the grading is recorded as weak because the mapped structures are schedule elements and not the mechanism.

Step 3. The precedent reading, with dates and paths. Yom Kippur: a single approximately 25-hour complete fast, annual, attested continuously in Jewish practice for over two millennia. Ramadan: a lunar month of daily fasts from dawn to sunset, seventh century onward, with a prescribed pre-dawn meal and a prescribed breaking at sunset. Christian Lenten and Friday abstinence: varying forms, fourth century onward. Ekadashi: an approximately fortnightly fast in several Hindu traditions. Buddhist monastic vinaya: no solid food after midday. Vision-quest fasting in a number of North American Indigenous traditions, and fasting in Mesoamerican ritual practice — **these are the entries that matter for independence, having no plausible shared source with the Old World examples.**

Independence audit. The Old World set is heavily correlated and c is high. The North American and Mesoamerican entries are substantially independent, which lifts the aggregate above the one-tradition floor. **Credit: on the toy model, of order two bits. Not zero. Not evidence.**

What was actually recorded: the duration, the timing, the exclusions, the manner of breaking the fast, the warning signs, and a reported change in mental and bodily state. **No tradition records a recycling process, a lysosome, or a nutrient sensor, because nothing in a first-person instrument could have detected one.**

Step 4. The prediction. If traditional fast durations sit above the deprivation threshold at which autophagic flux rises materially in human tissue, then the fasts are — incidentally, not by design — in the effective range, and shorter observances should not be. The traditions converged on durations of roughly 12 to 30 hours, which is a specific, checkable, falsifiable claim about where the threshold should lie. **The prediction: measured human autophagic flux should show a material rise within the 12-to-30-hour window and not substantially below it. If human flux turns out to require 72 hours, the bridge is dead.**

Step 5. Test. Not tested at this specificity, to this author's knowledge. **This is the step where the bridge stops.**

Contra-proof. The two integrate only partially. The evidence side need not claim the traditions knew about autophagy. **The precedent side must not be made to claim its fasts were about cellular recycling — they were about atonement, discipline, solidarity and mourning, and saying otherwise would put a falsehood in the precedent column to serve the evidence column, which is exactly what the neutral point forbids.** The integration that survives is narrow: both columns describe a threshold-and-cycle structure on nutrient deprivation, and they agree on roughly where the threshold sits.

```
C1 STRUCTURAL CORRESPONDENCE   PARTIAL   weak homomorphism; 2 of 5 structures map
C2 PREDICTIVE POWER            PASS      12–30 h window; falsifiable; untested
C3 PRACTICAL UTILITY           PARTIAL   a stated hypothesis on the (+1) side;
                                         the (−1) side gains nothing it did not have

STATUS: SPECULATIVE -- and this is the sentence that states it, so this is the
sentence that marks it.
```

Notice what the procedure did. It did not accept the bridge because it is appealing, and it is appealing, being among the most repeated claims in contemporary health writing. It did not reject it either, because the structure genuinely maps at the schedule layer and generates a real prediction. **It graded it, named the missing measurement, and printed the number that would settle it.**

29.7.7 7.7 What four bridges demonstrate

```
BRIDGE 1  cadence                ESTABLISHED   C1 pass, C2 pass, C3 pass
```

```
BRIDGE 2  vibration -> cell        GRADED        established coarse, SPECULATIVE fine
BRIDGE 3  528 Hz repairs DNA       REJECTED      C1 fail, C2 fail, C3 fail, -1 void
BRIDGE 4  fasting -> autophagy     SPECULATIVE   C1 partial, C2 pass, C3 partial
```

The middle rows are the ones to keep. Most real bridges are neither accepted nor rejected whole; they hold at one grain and fail at a finer one, and the work of the method is finding the seam. An apparatus reporting only the first and third rows would be an apparatus that had never been used on anything difficult.

29.8 8 The register: format, status vocabulary, and precedent weights

29.8.1 8.1 What a register can and cannot do

Appendix A is a **reference register**, not an argument. It is written to be looked up rather than read through, and every entry is formatted identically so that the appendix functions as a lookup table: a reader arrives with a claim and leaves with its two sources, what each of them actually established, which bridge criteria it passes, and its status.

It exists because Section 3.1 imposes an obligation that is easy to state and expensive to keep. **Every substantive claim in this part of the corpus carries two references and states which column carries evidence and which carries precedent.** A framework can make that promise in a methodology section and quietly break it afterwards. The only way to make it checkable is to print the whole ledger in one place, in one format, with the failures included. **The failures are included: entries with status REJECTED are printed at the same length and in the same format as those with status ESTABLISHED, because a register listing only what survived is a sales document.**

What a register can do: make the framework's commitments legible; make status visible at a glance; force every entry to declare the same fields, so that a missing field is a visible gap rather than a silent omission; and let a reader check one claim without reading the argument that makes it.

What it cannot do: establish anything. A register is a summary of evidence and never a source of it. No entry becomes truer by being in it. The nearest analogue in the corpus is the demonstration that a double-entry identity is a consistency check and not an integrity check: `A - (L + E) = 0.0000000000` holds exactly while the statement it certifies is misleading, **so the identity's holding is not evidence** [Curzi, 2026]. A register that balances is in exactly that position.

One further limit, stated because the format invites the mistake: **a table of forty-three entries all formatted identically will read as forty-three items of equal weight, and they are not.** The `ROLES` field exists to prevent that reading, and it is the field to read first.

29.8.2 8.2 The entry format

Every entry carries the same twelve fields in the same order.

Field	What it holds
CLAIM	the proposition, stated so that it could be false
(+1) SOURCE	the peer-reviewed reference, named
MEASURED	what that source *actually measured* — not what it is popularly said to show

Field	What it holds
(-1) SOURCE	the traditional reference, with tradition, date, and transmission path
RECORDED	what that source *actually recorded* — the observation, not the explanation
ROLES	which column carries evidence, which carries precedent, and the precedent's weight in bits
C1	structural correspondence: PASS / PARTIAL / FAIL, and the grade of correspondence
C2	predictive power: PASS / PARTIAL / FAIL, and the prediction stated
C3	practical utility: PASS / PARTIAL / FAIL, and the use on each side
0p	the contra-proof: whether the two integrate without either side lying
STATUS	ESTABLISHED / SPECULATIVE / REJECTED / POSTULATE / REFUSED
USED IN	the sections of the source manuscript that rely on the entry

Three fields do most of the work. **MEASURED** separates the register from a bibliography: a citation names a source, this field names the source's actual finding, which is frequently narrower than the use it is put to. **RECORDED** does the same on the other side and enforces the separation of what a tradition observed from what it explained. **ROLES** states the asymmetry every time and prices the precedent in bits.

A note on the last field. `USED IN` records the chapter numbering of the source manuscript from which this register was drawn, and is reproduced unchanged as a provenance record rather than remapped onto the corpus's paper numbering. **No mapping between the two numbering schemes is asserted here**, because inferring one would put a claim in the register that the source does not make.

29.8.3 8.3 The status vocabulary

Five statuses. **They are not a scale and they are not interchangeable**, and no arithmetic is performed on them anywhere (Section 3.7).

```
ESTABLISHED   Passes C1, C2 and C3. The evidence column carries evidence
              sufficient to support the claim on its own. The precedent column
              supplies precedent and usually procedure. This is the only status
              that may be stated without a qualifier.

SPECULATIVE   Fails or only partially satisfies at least one criterion, but is
              coherent and falsifiable. MUST be marked speculative in the same
              sentence that states it, everywhere it appears.

REJECTED      Fails the criteria. Printed at full length, with the reasoning, so
              that the rejection is auditable and so that the framework is bound
              by it.

POSTULATE     The author's own, badged. Not derived, not measured. Stated with
              its badge in the same sentence.
```

```
REFUSED      Outside what the framework will assert at all. No amount of
             evidence moves an entry into the framework from this status,
             because the refusal is about what will be claimed, not about what
             is true.
```

REJECTED is a verdict on a bridge — the method was run and the bridge failed. REFUSED is a standing constraint: the framework does not make claims of this kind regardless of what any bridge shows.

29.8.4 8.4 Precedent weights

The `ROLES` field prices precedent in bits using Section 4.2's model. **The model is illustrative and its rates are assumed; what carries across is the ordering, not the decimals.**

```
a single tradition, or several with a shared source     ~ 1.6 bits
two or more genuinely independent traditions            ~ 2 to 3 bits
a source that is not actually old                          0 bits
"many traditions teach"                                    0 bits, inadmissible

for comparison, on the same scale:
one strong study                                        ~ 6.6 bits
a replicated result                                     ~ 10.0 bits
a replicated, prize-recognised mechanism                ~ 20 bits
```

Every precedent figure in the register is at or below three bits. Every evidence figure it sits beside is at or above six. That ratio is the register's central fact, and it is why the columns are never summed.

29.9 9 Specifications

Numbered, checkable and binding on any use of this apparatus. **The three groups replace the source manuscript's chapter-keyed numbering: S-C for the concordance method, S-P for the precedent column, S-R for the register.**

29.9.1 9.1 S-C — the concordance method

S-C1. Every substantive claim carries an evidence reference and a precedent reference, and states which is which.

S-C2. The evidence reference names the source **and what it actually measured**, not what it is commonly said to show. Where the two differ, the difference is stated.

S-C3. The precedent reference names the tradition, an approximate date, and the transmission path. **A precedent citation of the form "many traditions teach" is inadmissible and must be replaced or struck.**

S-C4. The column roles are stated at every concordance. **Precedent is never presented as, summed with, or substituted for evidence.**

S-C5. The evidence reading is fixed in writing before the precedent column is consulted (Step 0). Where a bridge was in fact discovered from the precedent side, that history is disclosed and the bridge is held to a correspondingly higher standard on C2.

S-C6. Every bridge is graded isomorphism / homomorphism / analogy / metaphor. **Below homomorphism the bridge is rejected and the rejection is recorded rather than silently dropped.**

S-C7. Every bridge states its prediction explicitly, naming an observation that would falsify it. **"It is consistent with" is not a prediction.**

S-C8. Every bridge states its use on both sides. A bridge usable on only one side is recorded as one-sided and does not qualify as established.

S-C9. Only bridges passing C1, C2 and C3 are presented as established. All others are marked speculative in the same sentence that states them.

S-C10. Where a bridge holds at one grain and fails at a finer one, both grains are reported separately with separate statuses. **Partial bridges are not rounded up.**

S-C11. Every proposed mechanism is checked for scale (E6) by computing the relevant ratio, **and the ratio is printed.**

S-C12. Every proposed mechanism names its competing channels (E7) and the experiment that would separate them.

S-C13. Independence of precedent sources is audited before the precedent is credited. **Correlated sources are credited once, not once each.**

S-C14. No supernatural claim, no claim of physical efficacy for a symbol, no clinical or therapeutic claim, and no claim that genomic sequence is edited or altered. These are standing refusals and no bridge may weaken them.

S-C15. Where certainty about a source is lacking, the entry states the uncertainty or the source is omitted. **No source is cited from memory of its reputation.**

29.9.2 9.2 S-P — the precedent column

S-P1. Before citing a traditional source, run the P1–P4 envelope check and print the result. **A claim failing P4 — prompt effect — is not admissible as observational precedent, whatever the tradition's confidence.**

S-P2. Apply the visible-failure rule. State what happened to a practitioner who got this claim wrong. **If the answer is "nothing," the item is credited at the memorability rate, not the accuracy rate.**

S-P3. Separate what the source *recorded* from what the source *explained*. The recording may be admissible precedent; the explanation almost never is.

S-P4. Date the earliest verifiable attestation of the specific claim, not of the tradition it is attached to. **Pseudepigraphic attribution is not a date.**

S-P5. Audit independence before crediting agreement, and name the transmission path. **Assume harmonisation in any post-1450 European compilation unless there is evidence otherwise.**

S-P6. Treat any complete, gapless, symmetric table as containing interpolated entries (T4). **Do not use the completeness of a scheme as evidence for any of its entries.**

S-P7. Where a tradition reports an effect and assigns a mechanism, cite the effect and discard the mechanism, then name the competing channels per S-C12.

S-P8. Where a tradition's practical material and its theoretical material disagree, weight the practical material higher — bias B3 means it survived *despite* the filter rather than because of it.

S-P9. Never cite duration of belief as support. Duration measures `delta`, and `delta` is orthogonal to truth.

S-P10. No supernatural claim, no physical efficacy for a symbol, no clinical claim. Standing.

29.9.3 9.3 S-R — the register

S-R1. Every substantive claim cites a `CR-nn` identifier or adds a new entry in the same format. **A claim with no register entry is not admissible.**

S-R2. The twelve fields are filled for every entry, in order. **An empty field is written "none admissible" with the reason, never left blank and never filled with a loose analogue.** Twelve entries carry empty or void precedent columns — CR-06, CR-09, CR-10, CR-14, CR-16, CR-17, CR-21, CR-22, CR-26, CR-27, CR-30 and CR-33 — **and that is a normal outcome, not a defect.**

S-R3. The `MEASURED` field states what the source actually measured. Where the popular use of the source exceeds its finding, the excess is named in the field.

S-R4. The `RECORDED` field states the observation and excludes the explanation.

S-R5. The `ROLES` field is filled on every entry with the precedent weight in bits. **No entry sums the columns.**

S-R6. Status ESTABLISHED requires PASS on all three criteria. **Anything less is SPECULATIVE and is marked in the sentence that states it, everywhere, not only in the register.**

S-R7. REJECTED entries are printed in full, at the same length as accepted ones, with their reasoning. **Rejections are never removed when a section stops using them.**

S-R8. Where a REJECTED entry constrains the framework's own method, the constraint is stated in the entry. **CR-27 is the model: the rejection binds the sonification pipeline** [Curzi, 2026y].

S-R9. Numbers in entries are computed and appear identically everywhere. The time–bandwidth constant is `0.079577`; **`k = 0.483425732988414` is a defined constant and the hertz attachment is a separate, postulated claim** (Section 3.4); the Casimir pressures are `1.3e5 / 13 / 1.3e−3 / 1.3e−15 Pa` at `10 nm / 100 nm / 1 um / 1 mm`.

S-R10. Uncertainty about a source is stated in the entry. CR-07, CR-14, CR-27, CR-32 and CR-36 carry explicit uncertainty statements, and those statements travel with the citation.

S-R11. To add an entry: run the eight-step procedure of Section 7.1, fill all twelve fields, assign the next `CR–nn`, record the status. **The procedure is run before the entry is written.**

S-R12. To retire an entry: change its status, keep its number, and record why. **Numbers are never reused and entries are never deleted.**

29.10 10 Failure modes

Three groups, corresponding to the three specification groups. Each carries a clinical description, a signature by which it is recognised, the price of leaving it in place, the fix, and what it is confusable with.

29.10.1 10.1 FM-C — faults of the method

FM-C.1 · The Doubled Column. *Clinical:* precedent is counted as a second piece of evidence and the concordance is presented as convergent verification. *Signature:* the text says "both science and tradition confirm," or "independently verified by"; confidence in the section exceeds the confidence of its strongest single source. *Price:* computed at Section 4.3 — with a replicated evidence result, an overstatement factor of 333.3, which is 8.38 bits. *Fix:* state the column roles in the sentence. *Confusable with:* legitimate use of precedent as a prior (Section 4.5) — the same operation at the right weight.

FM-C.2 · The Shared Library. *Clinical:* several precedent sources are cited as independent when they descend from one another or from a common editorial project. *Signature:* the sources are all European, all post-1450, or all traceable to one compilation; the phrase "found in every tradition" appears. *Price:* Section 4.4. At `c = 0.1`, five sources buy 1.685 bits instead of 7.925 — a loss of 6.24 bits against what was claimed, and past `k = 2` additional agreeing sources *reduce* the likelihood ratio. *Fix:* name the transmission path per S-C3 and credit correlated sources once. *Confusable with:* genuine partial independence, as in Bridge 2.

FM-C.3 · The Retrofitted Bridge. *Clinical:* the precedent source is searched *after* the evidence result is known, and a passage is found that can be read to match. *Signature:* the traditional passage is short, quoted without surrounding context, and requires a specific transla-

tion to work. *Price:* the bridge carries zero bits, since a large enough corpus contains a match for any proposition — and worse, it feels like the strongest bridge in the section. *Fix:* Step 0; where the order was in fact reversed, disclose it per S-C5 and demand a stronger C2. *Confusable with:* hypothesis generation from the precedent column, which differs in one respect — the hypothesis is then tested, not illustrated.

FM-C.4 · The Scale Skip. *Clinical:* error E6; a real mechanism is transported to a scale at which its magnitude is negligible, and the transport is not checked. *Signature:* a named physical effect appears in a biological or architectural sentence with no number attached. *Price:* arbitrarily large. Casimir pressure falls as `a^-4`: `1.3e5 Pa` at 10 nm, `1.3e-15 Pa` at 1 mm — twenty orders of magnitude across an interval a careless sentence would cross without noticing. *Fix:* S-C11; compute the ratio and print it. *Confusable with:* the structural use of a physical effect — boundary geometry selecting modes, with the force explicitly disclaimed.

FM-C.5 · The Wrong Channel. *Clinical:* error E7; a genuine, replicable effect is attributed to a mechanism that is not the one operating. *Signature:* the effect keeps replicating, so the mechanism keeps appearing confirmed; no experiment in the citation list separates the candidate channels. *Price:* every downstream design decision optimises the wrong variable. *Fix:* S-C12; enumerate the competing channels and specify the separating experiment, as in Section 7.4. *Confusable with:* FM-C.6, where there is no effect at all.

FM-C.6 · The Costume. *Clinical:* a recent construction is presented as traditional, so that the precedent column appears populated when it is empty. *Signature:* the system uses absolute numerical values that could not have been measured in the period claimed; its earliest verifiable source is twentieth century; the historical material it gestures at is real but says something else. *Price:* the column reads as precedent and is worth zero bits. **This is the 528 Hz case exactly** (Section 7.5). *Fix:* date the earliest verifiable attestation of the *specific claim*. *Confusable with:* genuine recent formalisation of an old practice, legitimate provided the formalisation is dated as recent.

FM-C.7 · The Unfalsifiable Bridge. *Clinical:* C2 is nominally satisfied by a "prediction" that no observation could contradict. *Signature:* the prediction contains "may", "can", "is consistent with", or a magnitude-free directional claim; no number, no threshold, no control. *Price:* the bridge cannot be removed by any future evidence, which means it also cannot be supported by any [Popper, 1959, Meehl, 1967]. *Fix:* S-C7; rewrite the prediction so it names the observation that would kill it. *Confusable with:* an honestly speculative bridge, which is *marked* speculative and specifies what would settle it.

FM-C.8 · The Exemption. *Clinical:* the method is applied rigorously to other people's claims and suspended for the framework's own. *Signature:* a rejection is followed by a construction that would not survive that rejection; the framework's own mechanism is described in the passive voice. *Price:* total. **It is the only failure mode on this list that invalidates the whole apparatus rather than one section, because it converts the method from an instrument into a rhetorical device.** *Fix:* Section 7.5's closing move, performed every time — after each rejection, state explicitly what the rejection forbids *this framework* from claiming, and enforce it. *Confusable with:* nothing. There is no innocent version.

29.10.2 10.2 FM-P — faults of the reader of the old corpus

FM-P.1 · The Confidence Transfer. *Clinical:* the corpus's tone — settled, detailed, unanimous — is read as a measure of its reliability. *Signature:* the citation quotes the source's certainty rather than its observation. *Price:* unbounded. Tone measures `k`, and `k` measures variance, not bias (Section 5.2). *Fix:* S-P1; score the item on P1–P4 and ignore the tone entirely. *Confusable with:* legitimate citation of a *detailed procedure*, where the detail is operational payload rather than rhetorical certainty.

FM-P.2 · The Index Realised. *Clinical:* error T1; a filing relation is read as a causal one. *Signature:* a sentence in which a category assignment becomes a verb — "the metal is governed by". *Price:* everything downstream is built on a relation that exists only in the filing system. *Fix:* restate the claim literally; if the restatement is "X is filed under Y", the causal reading has

no content. *Confusable with:* genuine structural correspondence, which survives the restatement.

FM-P.3 · The Tidy Table. *Clinical:* error T4; interpolated entries are treated as observations. *Signature:* the scheme has no gaps, every category has the same number of members, and the count is 4, 7, 12 or 22. *Price:* the interpolated entries are indistinguishable from the observed ones and there is generally no way to separate them after the fact. *Fix:* S-P6; use the scheme as an index, never as data. *Confusable with:* a genuinely complete enumeration, which exists — a graph really does have 22 edges — and is distinguished by being a fact about the artefact rather than about the world.

FM-P.4 · The Invisible Failure. *Clinical:* a claim whose falsity had no cost to anyone is credited at the same rate as one whose falsity was punished immediately. *Signature:* the claim concerns mechanism, cosmology, remote influence, or outcomes after death. *Price:* Section 5.4's first row — 7.9 per cent surviving accuracy against 94.5, a factor of 11.97. *Fix:* S-P2; ask what happened to the practitioner who was wrong. *Confusable with:* a claim whose failure was visible but slow, such as institutional cadence, credited at an intermediate rate rather than dismissed.

FM-P.5 · The Inverted Inheritance. *Clinical:* bias B3; the theory is cited because the theory is what survives, while the practice — what the instrument was actually good at — is ignored because it is fragmentary. *Signature:* the section cites a cosmology and not a procedure. *Price:* the reader is handed the corpus's worst material as though it were its best. *Fix:* S-P8. *Confusable with:* nothing. This one is a failure to look.

FM-P.6 · The Harmonised Chorus. *Clinical:* post-hoc editorial reconciliation is read as independent agreement. *Signature:* sources from one compilation, one century, or one editorial programme are described as "multiple independent traditions". *Price:* Section 5.6, item 4 — 6.24 bits claimed and not held. *Fix:* S-P5. *Confusable with:* genuine geographic independence, as in Bridge 2, where the Australian and West African instances have no plausible shared source with the Indo-European ones.

FM-P.7 · The Unfalsifiable Rescue. *Clinical:* error T5, adopted by the modern reader rather than the ancient one. A traditional claim failing a modern check is rescued by reinterpretation until it can no longer fail. *Signature:* the claim becomes more metaphorical each time it is challenged. *Price:* the claim is now unfalsifiable and therefore also unsupportable [Lakatos, 1970]. *Fix:* fix the claim's content in writing before checking it, and let it fail if it fails. *Confusable with:* legitimate re-reading of a source as describing a phenomenon rather than a mechanism (S-P3) — the difference is that the re-reading is done once, in advance, and then held fixed.

FM-P.8 · The Survivor's Testimony. *Clinical:* a traditional literature's own account of its efficacy is read as evidence of efficacy. *Signature:* the supporting material consists of accounts by successful practitioners. *Price:* zero bits. Every tradition's surviving literature has this property, including the traditions that demonstrably did not work. *Fix:* ask where the unsuccessful practitioners' accounts are; if the answer is that they were not copied, the literature carries no information about the base rate. *Confusable with:* a tradition that recorded its own failures, which some technical traditions did and which is a strong positive signal precisely because it is rare.

29.10.3 10.3 FM-R — faults in the use of the register

FM-R.1 · The Equal-Weight Illusion. *Clinical:* forty-three entries in one format are read as forty-three items of equal standing. *Signature:* a section cites CR-13 (both columns weak) and CR-21 (prize-grade evidence, empty precedent) with the same confidence. *Price:* uniformity that exists to make comparison possible is used to make it unnecessary. *Fix:* read `ROLES` first. *Confusable with:* co-citation that states the difference in standing explicitly.

FM-R.2 · The Register as Authority. *Clinical:* an entry is cited as though the register established it. *Signature:* "as established in the register". **The register establishes nothing** (Section

8.1). *Fix:* cite the source in the `(+1) SOURCE` field. *Confusable with:* citing the register for its *status assignment*, which is legitimate — the status is the register's own product.

FM-R.3 · The Filled Blank. *Clinical:* an empty precedent column is filled with a loose analogue to make the entry look complete. *Signature:* `RECORDED` paraphrases rather than reports an observation; no date or transmission path is given. *Price:* zero bits presented as two. *Fix:* S-R2; write "none admissible" and give the envelope reason. *Confusable with:* a genuine but weak precedent column, which has a date, a tradition and an independence estimate.

FM-R.4 · Chain Assembly. *Clinical:* several ESTABLISHED entries are chained to reach a conclusion none supports, most often to close CR-26. *Signature:* "since CR-20, CR-28 and CR-29 are all established, it follows that". *Price:* total, for the framework's credibility. **CR-26 is exactly the gap this chain appears to close and does not** [Curzi, 2026y]. *Fix:* no entry may be combined with another to close CR-26; each entry supports only what its own `MEASURED` field says. *Confusable with:* composition in which the weakest component sets the status.

FM-R.5 · Status Drift. *Clinical:* a SPECULATIVE entry is cited later without its badge and by the third citation reads as established. *Signature:* the badge appears at first use and then stops. *Price:* the framework's principal honesty commitment fails silently. *Fix:* S-R6; the badge attaches to the *sentence*, not to the first mention. *Confusable with:* nothing.

FM-R.6 · The Vanishing Rejection. *Clinical:* a REJECTED entry is dropped once no section cites it. *Signature:* the register's rejection count falls between drafts. *Price:* the register becomes a list of successes. *Fix:* S-R7 and S-R12. *Confusable with:* nothing.

FM-R.7 · The Precedent Ladder. *Clinical:* weak precedent entries are accumulated across the register and treated as jointly substantial. *Signature:* "taken together, the traditional record shows". *Price:* Section 4.4's independence collapse. *Fix:* S-R5; never sum precedent, within an entry or across entries. *Confusable with:* noting that several audited-independent entries point the same way.

29.11 11 Results

29.11.1 11.1 R1 — the two-column refinement, and the asymmetry it prices

The evidence column carries what could have come out otherwise; the precedent column carries only that an idea was independently arrived at and preserved. Under the assumed recording rates, a single traditional record is worth `LR = 3.000`, or **1.585 bits**, against **9.966 bits** for a replicated result and **19.932 bits** for a prize-grade replicated mechanism. **Two sources agreeing does not double the evidence when only one of them is evidence.** The statement is an arithmetic consequence of the assumed rates, not a measurement, and its robust content is the ordering.

29.11.2 11.2 R2 — independence collapse, computed

At a shared-source probability `c = 0.1`, `k` agreeing traditions buy **1.585, 2.546, 2.317, 1.886 and 1.685 bits** for `k = 1` to `5`. **The total falls after `k = 2`**, and it converges to 1.585 bits — one tradition's worth — for any `c > 0` as `k` grows (Section 13.3). At `c = 0.3` five traditions buy 1.612 bits, within two per cent of what one tradition buys alone. **Against the independent-source fantasy of 7.925 bits at `k = 5`, the loss at `c = 0.1` is 6.24 bits.**

29.11.3 11.3 R3 — the survivorship filter, and its twelvefold spread

Under the preservation model `a1 = a0 s_acc / (a0 s_acc + (1−a0) s_inacc)`, an original corpus 30 per cent accurate arrives at **7.9 per cent** accuracy where failures were invisible and **94.5 per cent** where they were visible — **a factor of 11.97 produced by the filter and not at all by the source.** A poorer original corpus at `a0 = 0.10` still arrives at 81.6 per cent if its failures were

visible, an eightfold improvement on its own origin. **The decisive question about a traditional claim is therefore what happened when it was wrong.**

29.11.4 11.4 R4 — the 528 Hz bridge, rejected on three computed grounds

The claim that a tone at 528 Hz repairs DNA is **REJECTED**, on three independent computed objections and one historical one:

```
(a) spatial     lambda = 1540/528 = 2.9167 m in soft tissue
                ratio to a 10 um cell        2.917e5
                ratio to a 6 um nucleus      4.861e5
                ratio to a 0.34 nm base pair 8.578e9
(b) energetic   h f / kT = 8.17e-11 at 310 K
(c) spectral    molecular bands 1.499e13 to 8.994e13 Hz
                = 34.7 to 37.3 octaves above 528 Hz
(d) historical  the "Solfeggio" values are a 1990s construction; medieval
                solmisation assigns no frequencies, and 528 Hz is a specific
                note only under a twentieth-century pitch convention
```

The rejection binds the framework's own sonification pipeline [Curzi, 2026y], which may not claim a molecular resonance for a transposed tone, may not treat its essence frequencies as physical resonances, and may not claim sequence-specific acoustic action.

29.11.5 11.5 R5 — the register's composition, and a tally defect in the source

Across the forty-two entries the source register carries before its worked example:

```
ESTABLISHED    32        SPECULATIVE     2
REJECTED        6        POSTULATE       2      TOTAL   42
```

Four of the thirty-two ESTABLISHED entries carry a secondary speculative rider (CR-04, CR-12, CR-31, CR-32), one carries a POSTULATE rider (CR-19), and two carry a standing REFUSED clause (CR-37, CR-38). Counting every entry carrying a qualifier — rejected, speculative, postulate, or established-with-a-speculative-rider — gives **14 of 42, or 33.3 per cent**. With the worked example CR-43 added the register holds 43 entries, of which 3 are SPECULATIVE, and the qualifier count is **15 of 43, or 34.9 per cent**.

By correspondence grade on C1, over the completed 43-entry register:

```
PASS       21        homomorphism or better
PARTIAL     5        analogy, or homomorphism at one grain only
FAIL        7        metaphor only, or no shared structure
N/A        10        no bridge proposed
```

Seven outright C1 failures out of the thirty-three bridges attempted is a 21.2 per cent rejection rate at the first criterion alone, and that is where a concordance method does most of its work.

A defect in the source's own tally is reported here rather than silently reconciled. The source presents the C1 breakdown under a heading reading "counts across the forty-two entries below", but the four C1 categories sum to 43 and the split into "entries drawing a bridge (33)" and "entries with C1 = N/A (10)" likewise sums to 43. Both tallies include CR-43, which the source introduces later, in its worked example. **Over the forty-two entries alone the C1 split is PASS 21 / PARTIAL 4 / FAIL 7 / N/A 10, thirty-two bridges are attempted, and the failure rate is 7/32 = 21.9 per cent rather than 7/33 = 21.2 per cent.** Both figures are printed here; the headline 21.2 per cent is the figure for the completed register and is the one carried into the abstract

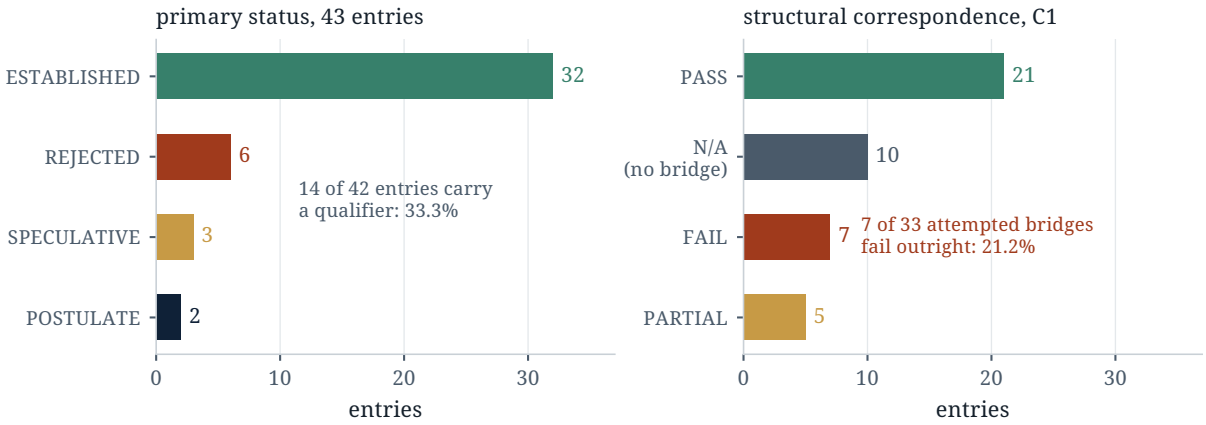


Figure 29.5: **A register that can reject, and does.** Left: the forty-three-entry register by primary status. Right: the same entries by structural-correspondence grade, with the seven outright failures of the thirty-three attempted bridges marked; the ten `N/A` entries draw no bridge at all, which is a normal outcome rather than a defect. Plotted as small multiples because status and grade are different classifications of the same entries and share no scale. Counted from Appendix A by recomputing every `STATUS` and `C1` field; generated by `figs/src/CCP-029_fig05.py`.

and the conclusion. A second, smaller slip: the annotation "plus CR-43 in the worked example" is attached to the `C1 = N/A` row, although CR-43's C1 grade is PARTIAL and it therefore belongs among the bridges attempted. **Neither slip changes any verdict; both were found by recomputing every field of every entry, and the recomputation reproduces the source's status counts of 32/2/6/2 over 42 entries exactly.**

One further note on the source's prose. The sentence summarising these ratios reads "one entry in three carries a qualifier, and one in five and a third is a rejection." **The first clause is exact (14/42 = 33.3 per cent). The second corresponds to no computed ratio in the section: rejections are 6 of 42, which is 14.3 per cent or one in seven, and C1 failures are 7 of 33, which is 21.2 per cent or one in 4.7.** The exact fractions are what this paper reports.

29.11.6 11.6 R6 — the pre-modern detection floor

A single uncontrolled first-person observer registers a true standardised shift of `d = 0.20` with probability **0.2858**, and of `d = 1.00` with probability **exactly 0.5000** — at one standard deviation of true effect, the observer is a coin. Reliability arrives only above `d = 2.0`, at 0.7602. **Against this, a modern two-group design needs about 392 subjects per group to detect `d = 0.20` at 80 per cent power and about 16 at `d = 1.00`.** The traditional instrument's enormous `k` clears the noise floor and cannot clear the attribution floor, and that single fact determines the P1–P4 operating envelope of Section 5.3.

29.11.7 11.7 The counter-result

The register is built so that it can reject, and it does. Seven of thirty-three attempted bridges — 21.2 per cent — fail structural correspondence outright; one entry in three carries a qualifier; the claim nearest the framework's own subject matter is the one rejected at greatest length; and **the arithmetic of correlated sources shows that past two agreeing traditions each further agreement reduces rather than increases the likelihood ratio.** Precedent is priced at a fraction of what its accumulation appears to promise, and the fraction is computed rather than asserted.

29.12 12 Discussion

29.12.1 12.1 What follows

A concordance can be run honestly, and the cost of running it honestly is a rejection rate. The apparatus here is not a device for making traditional material respectable; it is a device for finding out which parts of it are load-bearing, and its principal output is a set of verdicts, six of which are refusals. The most transferable result is the pricing: **once precedent is denominated in the same unit as evidence, the temptation to sum them disappears, because the sum is visibly dominated by one term.**

The independence result generalises well beyond traditional texts. Any setting in which sources may have copied one another has the same structure: agreement above a small number of witnesses is better explained by contact than by convergence. That is the mechanism behind informational cascades [Bikhchandani et al., 1992, Banerjee, 1992], behind the measured collapse of crowd accuracy under mutual exposure [Lorenz et al., 2011], and behind the path-dependence of cultural success [Salganik et al., 2006]. It applies to citation networks, to expert panels, and to the corpus's own interlocking references. **A framework whose papers cite one another densely should notice that its internal agreement is worth fewer bits than its edge count suggests, and this paper's own internal citations are subject to exactly that discount.**

The survivorship result reframes a common move. "This has been believed for two thousand years" is routinely offered as support and is, on this analysis, a statement about `delta` — transmission fidelity — and nothing else. The useful question replaces it: what happened to the person who was wrong? That question sorts traditional material along a line orthogonal to how mystical it sounds, and it puts metallurgy, dyeing, navigation, and mnemonics on one side and cosmology, composition, and remote influence on the other.

29.12.2 12.2 What does not follow

It does not follow that any register entry is established by being in the register. Section 8.1 states this as a definition and Section 10.3 names the failure mode.

It does not follow that the register's badged postulate has been supported by its own entries. The postulate CR-08 records is stated elsewhere in the corpus in magnitude form as `E = G.B + eps`, with the forbidding clause that `G = 0` implies `E = 0` for every `B`, badged as a postulate and paired by explicit correspondence with the storage law's `eta = 0` [Curzi, 2026ab]. **Its channels are measured and it is not**, and summing CR-01 to CR-05 does not establish it.

It does not follow that the bit weights are measurements. They are logarithms of assumed ratios. A reader who prefers different rates will get different decimals and the same ordering, and the ordering is all that is claimed.

It does not follow that a rejected bridge is a false claim. REJECTED is a verdict that the bridge fails the three criteria, which for CR-27 coincides with the physical impossibility of the mechanism, and for other entries does not. **The distinction between a verdict on a bridge and a verdict on the world is maintained throughout and is the reason REFUSED exists as a separate status.**

It does not follow that the reflexive entries of the register are identified. Section 3.5 states the reflection problem and adopts no remedy, so CR-01 through CR-05 and CR-13 record modelling positions rather than identified empirical results [Manski, 1993, Shalizi and Thomas, 2011].

And it does not follow that the method's rejection rate demonstrates the method's stringency. The register was assembled by the same author who set the criteria, and a rejection rate is only evidence of stringency if the candidate set was not chosen with the criteria in view. **That objection is not answered here and is recorded as limitation L-4.**

29.12.3 12.3 Relation to the surrounding literature

The apparatus's closest methodological relatives are the preregistration literature, which supplies the Step 0 device at a larger scale [Nosek et al., 2018, Simmons et al., 2011], and the strong-inference tradition, which supplies the demand that a method be able to reject [Platt, 1964, Popper, 1959]. Meehl [1967, 1990] supply the sharpest version of the objection this apparatus is built against: directional predictions in soft science are nearly costless, so a method must make point predictions with numbers in them or it is not being tested at all. **Sections 7.5 and 11.4 are the paper's answer — three computed ratios, each of which could have come out otherwise.**

The apparatus also has an internal relative. The corpus's own honesty ledger classifies its borrowings into three drawers and, in the psychological drawer, maps seven established items — counter-movement under pressure, if-then plans, commitment and consistency with the announcement-substitution caution, measurement reactivity, unfinished-task persistence, confirmation bias, and the self-fulfilling prophecy — each to the node that leans on it [Curzi, 2026]. **That is the same operation as the `MEASURED` field of Section 8.2, performed on the corpus's own claims rather than on traditional ones.** And the corpus's frequency reference apparatus reaches a structurally parallel verdict about resolution limits by computing a rate penalty at a cliff — exactly 8 symbols per second, from `(1/32)/(1/256)`, with the interval shrinking by a factor of 5.828 in cents and 5.878 in cycles [Curzi, 2026] — which is the same kind of number CR-11 and CR-12 turn on.

The independence arithmetic sits alongside the sociophysics literature on opinion formation, where the same collapse appears as consensus without information [Castellano et al., 2009, DeGroot, 1974, Friedkin and Johnsen, 1990], and against the standing critique of physics-in-social-science, which is stocked here deliberately [Sobkowicz, 2009, Gallegati et al., 2006, McCauley, 2006]. The critique's central complaint — that formal machinery is imported without the conditions that license it — is precisely what Section 3.4 exists to answer, one borrowing at a time.

On the measurement side, the framework is most exposed where it treats a psychological quantity as though it had additive structure. **Michell [1997]'s objection — that psychology assumes quantitative structure without testing for it — is the strongest single objection to this corpus**, and the test that would settle it is additive conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971]. **It was not run.** This paper's own exposure is narrow, because its numbers are bits and probabilities rather than psychological magnitudes, but the register it prices contains entries whose underlying constructs are exactly the contested kind, and the corpus's own protocol for demoting such quantities [Curzi, 2026t] is the instrument that should be run on them.

Finally, the functional form question. Where the framework's surplus functional appears in this paper it appears as `ln(1 + (N−1)u)` [Curzi, 2026]. **That form is underdetermined by the evidence:** social impact rising as a sub-linear power of the number of sources [Latané, 1981] and value growing as `n log n` [Briscoe et al., 2006] fit the same qualitative facts equally well, and nothing in this paper discriminates between them. The strongest available warrant for a logarithm is not the perceptual one but the dynamical one: Peters and Gell-Mann [2016] derives a logarithmic form from multiplicative dynamics with no psychological assumption at all, which is a better foundation than Fechner [1860], Stevens [1957] because it does not require the perceptual story to be true. **The argument here needs only the dynamical version, and that is said rather than left to be inferred.**

29.13 13 Verification

29.13.1 13.1 Dimensional audit

Every quantity in this paper, its dimension, its unit, and its epistemic status. `[1]` denotes dimensionless.

Symbol	Dimension	Unit	Status
`p1`, `p0`	`[1]`	probability	**modelling choice** — assumed recording rates, 0.30 and 0.10
`LR`	`[1]`	odds ratio	**derived** from `p1`, `p0`
bits	`[1]`	bit (Shannon)	**derived**: `log2(LR)`. **Shannon information, not thermodynamic entropy**
`c`	`[1]`	probability	**modelling choice** — shared-source probability
`k` (source count)	`[1]`	count	**defined**
`a0`, `s_acc`, `s_inacc`, `a1`	`[1]`	probability	`a0`, `s_acc`, `s_inacc` **assumed**; `a1` **derived**
`d`	`[1]`	standard deviations	**defined** — a standardised mean difference
`n` per group	`[1]`	subjects	**derived** from `d` at fixed power
`Q`	`[Q]`	none — see below	**state variable** of the storage law
`delta`	`[1]`	per period	**assumed** in the illustrative parameterisations
`eta`	`[Q]` per nat	none — see below	**assumed** in the illustrative parameterisations
`T`	`[T]`	periods	**defined**
`lambda`	`[L]`	metres	**derived**: `c_sound / f`
`f`	`[T^-1]`	hertz	**measured constant** of the claim under test
`c_sound`	`[L T^-1]`	m/s	**cited nominal**: 1540 in soft tissue, 343 in air
`h f`	`[E]`	eV	**derived** from Planck's constant
`kT`	`[E]`	eV	**derived**: 0.0267 eV at 310 K
octaves	`[1]`	none	**derived**: `log2` of a frequency ratio
Casimir pressure	`[E L^-3]`	Pa	**cited**, scaling as `a^-4`
`k` (the framework's constant)	`[1]`	**none as defined**	**DEFINED**, exactly `(9/8)^2/phi^2`. Attaching hertz is a **separate POSTULATE**

Does the storage law balance? `Q_{t+1} = (1-delta)Q_t + eta S_t - C_t` balances only if `eta` `S` and `C` carry the dimension of `Q` per period. In Section 6.2's parameterisation `S` is dimensionless (one study session) and `eta` therefore carries `[Q]` directly; in Section 6.3's `R` is in `[Q]` per period, so `mean Q = R/delta` returns `[Q]` because `delta` is per period. **Both are consistent, and the two sections use different conventions**, which is stated here rather than left for a reader

to reconcile.

Where a quantity has no operational unit, that is said. `Q` is stored order and has no natural unit; it counts whatever the declared proxy counts, and two systems' `Q` are not comparable unless they declare the same proxy [Curzi, 2026i]. **The underlying construct — comfort — has no operational unit at all** and is treated as a dimensionless magnitude by stipulation; Section 14 records that the quantitative structure required for that stipulation has not been tested. **The register's status vocabulary has no dimension because it is not a quantity**; it is an ordinal label set on which no arithmetic is performed.

29.13.2 13.2 The postulates

Six postulates. Each is atomic, checkable alone, and carries its status. Where a postulate could be dropped and the conclusion still follow, it has been dropped.

P1 — A tradition records a true structure with probability 0.30 and a false one with 0.10. *Status:* **MODELLING CHOICE.** No measurement supports these values; they are chosen generous to the precedent column, and Section 14 records the consequence.

P2 — Likelihood ratios for conditionally independent considerations multiply, and their logarithms add. *Status:* **ARITHMETIC** — Bayes' theorem [Jaynes, 2003].

P3 — `k` sources agreeing report jointly with probability `c p + (1-c) p^k`, where `c` is the probability they are effectively one source. *Status:* **DEFINITION** — the mixture model of Section 4.4, true by its own construction and carrying no empirical weight.

P4 — The standard deviation of a mean over `k` independent observations is `1/sqrt(k)` of the single-observation standard deviation, and a systematic bias is invariant under `k`. *Status:* **ARITHMETIC.**

P5 — Survival of an item is a Bernoulli filter with rates `s_acc` and `s_inacc` applied independently to accurate and inaccurate items. *Status:* **DEFINITION**, with the rates a **MODELLING CHOICE**. Only the *sign* of `s_acc \ensuremath{-} s_inacc` is asserted, and it is asserted from the mechanism named beside each row.

P6 — The speed of sound in soft tissue is approximately 1540 m/s, Planck's constant and Boltzmann's constant take their standard values, and molecular vibrational modes of biological interest lie between 500 and 3000 cm^{-1}. *Status:* **CITED** — standard physical constants and standard vibrational spectroscopy [Rossing et al., 2002]; the vibrational range is a textbook fact stated without a key (Appendix B).

A postulate that was considered and dropped. An earlier form of the argument assumed a prior on `H`. It is not needed: every result in Sections 4.3 and 4.4 is a statement about likelihood ratios, which are prior-free, and the 3,000:1 against 1,000,000:1 comparison of Section 4.3 uses even odds only as an illustration and would carry the same *factor* from any prior.

29.13.3 13.3 The compounding chain

Every line is checkable from its own inputs, without reading the previous line for anything but its stated output.

```
P1        -> L1   LR of one traditional record = 0.30/0.10 = 3.000        ARITHMETIC
L1 + P2   -> L2   log2(3.000) = 1.58496...  -> 1.585 bits                 ARITHMETIC
P2        -> L3   a replicated result at LR = 1000 is log2(1000)
                  = 9.96578... -> 9.966 bits                             ARITHMETIC
L2 + L3   -> R1   9.966 / 1.585 = 6.288  -- evidence is ~6x precedent     ARITHMETIC
L3 + L1   -> L4   honest posterior odds from even odds: 1000 x 3 = 3,000  ARITHMETIC
L3        -> L5   doubled posterior odds: 1000 x 1000 = 1,000,000         ARITHMETIC
L5 / L4   -> R1b  overstatement factor 333.33 = 8.3808 bits              ARITHMETIC
P1 + P3   -> L6   at c = 0.1, k = 2:
                  num = 0.1(0.30) + 0.9(0.09)  = 0.111
```

```
                  den = 0.1(0.10) + 0.9(0.01)  = 0.019
                  LR  = 5.8421 -> 2.5464 bits                              ARITHMETIC
P1 + P3   -> L7   at c = 0.1, k = 5:
                  num = 0.03 + 0.9(0.00243)    = 0.032187
                  den = 0.01 + 0.9(0.00001)    = 0.010009
                  LR  = 3.2158 -> 1.6851 bits                              ARITHMETIC
L6 + L7   -> R2   1.6851 < 2.5464: the total FALLS after k = 2             ARITHMETIC
P3        -> R2b  as k -> infinity, p^k -> 0, so LR -> (c p1)/(c p0)
                  = p1/p0 = 3, and bits -> 1.585 for every c > 0           ARITHMETIC
P5        -> L8   a1(0.30, 0.02, 0.10) = 0.006/0.076   = 0.078947          ARITHMETIC
P5        -> L9   a1(0.30, 0.40, 0.01) = 0.120/0.127   = 0.944882          ARITHMETIC
L9 / L8   -> R3   0.944882/0.078947 = 11.968  -- twelvefold                ARITHMETIC
P6        -> L10  lambda = 1540/528 = 2.916667 m                           ARITHMETIC
L10       -> R4a  2.916667 / 3.4e-10 = 8.578e9 base-pair lengths           ARITHMETIC
P6        -> R4b  h f/kT = (6.626e-34 x 528)/(1.381e-23 x 310) = 8.17e-11  ARITHMETIC
P6        -> R4c  log2(1.499e13/528) = 34.72 octaves                       ARITHMETIC
R4a+b+c   -> R4   the 528 Hz bridge fails C1 on three independent grounds  VERDICT
```

The three checks a reviewer can redo in under a minute are `0.30/0.10 = 3`, `log2(3) = 1.58496`, and `1540/528 = 2.916667`. Everything in R1, R2 and R4 is built from those and from the mixture model of P3, and from nothing else.

What the chain does not establish. It does not establish that traditions record true structures at 0.30 and false ones at 0.10; P1 is a modelling choice and every bit figure inherits its status. It does not establish that any register entry is true. **And R4 is the only result in the chain whose inputs are physical constants rather than assumptions**, which is why R4 is the one result stated here without hedging.

29.13.4 13.4 The verification instance, worked

The smallest case that exercises the independence collapse is `k = 2` at `c = 0.1`, and it is worked in full because it is where the counter-intuitive behaviour begins.

```
p1 = 0.30, p0 = 0.10, c = 0.1, k = 2

P(both report | H)      = 0.1(0.30) + 0.9(0.30)^2 = 0.030 + 0.081 = 0.111
P(both report | not H)  = 0.1(0.10) + 0.9(0.10)^2 = 0.010 + 0.009 = 0.019

LR = 0.111 / 0.019 = 5.84211      log2 = 2.54638 bits

against perfect independence:  (0.30/0.10)^2 = 9.000   log2 = 3.16993 bits

the loss to a ten per cent chance of a shared source, at TWO sources, is
3.16993 - 2.54638 = 0.62355 bits, and it grows with k
```

Status: **ARITHMETIC.** Two multiplications, two additions and one division; a reader with a calculator can confirm every figure in the `c = 0.1` column the same way.

29.13.5 13.5 Two secondary checks

The cadence table. The periodic-steady-state values of Section 6.3 follow from `Q_max = R T / (1 - (1-delta)^T)` and `Q_min = Q_max (1-delta)^{T-1}`. At `T = 7`, `delta = 0.10`: `Q_max = 7/(1 - 0.9^7) = 7/0.5217031 = 13.4176`, and `Q_min = 13.4176 x 0.9^6 = 13.4176 x 0.531441 = 7.1305`. **Both reproduce the table to four decimals, and the mean is `R/delta = 1/0.1 = 10.000` in every row by the DC-gain result.** *Status:* **ARITHMETIC**, on a **DEFINITION** (the dosing schedule).

The ordering information. `log2(7!) = log2(5040) = 12.29920`, and `log2(12!) = log2(479001600) = 28.83527`. *Status:* **ARITHMETIC.**

29.13.6 13.6 Limits, asymptotics, and the domain of validity

Limit	Behaviour	Reading
`c -> 0`	bits grow linearly in `k`, `k log2(p1/p0)`	perfect independence; the fantasy column
`c -> 1`	bits are constant at `log2(p1/p0)` for every `k`	all sources are one source; agreement adds nothing
`k -> infinity`, any `c > 0`	`LR -> p1/p0`, bits `-> 1.585`	**the collapse is total in the limit**; the maximum at small `k` is the only thing the sum ever buys
`p1 -> p0`	`LR -> 1`, bits `-> 0`	a tradition equally likely to record truth and falsehood is worth nothing, correctly
`s_acc = s_inacc`	`a1 = a0`	a neutral filter transmits the original accuracy; the model's null
`s_inacc -> 0`	`a1 -> 1`	a filter that removes every error leaves only truth; degenerate and not realistic
`d -> 0`	`P -> 1 - Phi(1/sqrt2) = 0.2398`	**a single observer still "detects" a null effect 24 per cent of the time**; the false-positive floor of the uncontrolled instrument
`d -> infinity`	`P -> 1`	large effects are seen; correct
`delta -> 0` in the retention model	no decay; the ratio in Section 6.2 diverges	**a pathology, not an ideal** — nothing decays out of the stock [Curzi, 2026a]
`T -> 1` in the cadence model	`Q_max = Q_min = R/delta`	zero ripple; the continuous-drive case
`T -> infinity`	`Q_min -> 0`	the store empties completely between doses; the floor result at its extreme

Two domain restrictions follow and are stated as restrictions. The exponential retention model of Section 6.2 is valid over days and not over months: real retention curves are not single exponentials and flatten into a long tail, so the day-30 ratio of 224,131 should be read as "the rote trace is gone" rather than as a number. And the `d -> 0` limit above sets a floor the whole precedent column sits on: **an instrument with a 24 per cent false-positive rate on a null effect will generate a large volume of confident, specific, false material given three thousand years to do it in**, which is the quantitative form of Section 5.2's conclusion.

29.13.7 13.7 Invariance, and how far each claim has been tested

The bit ordering is invariant under the choice of rates, over the range tested. The claim that evidence exceeds precedent by a factor of at least four was checked at `(p1, p0) = (0.30, 0.10)`, `(0.50, 0.10)`, `(0.40, 0.20)` and `(0.60, 0.30)`, giving precedent weights of 1.585, 2.322, 1.000 and

1.000 bits against a fixed 9.966 for a replicated result — ratios of 6.29, 4.29, 9.97 and 9.97. **Four points, not a proof.** The invariance breaks if `p1/p0` is allowed above 8, which would require a tradition eight times more likely to record a truth than a falsehood, and **no defence of such a rate is offered here or known to this author.**

The collapse result is invariant under the rates entirely, because R2b shows the limit is `p1/p0` for every `c > 0`, whatever those rates are. That is the strongest invariance claim in the paper and it is a proof rather than a sample.

The physical results of R4 are invariant under the choice of tissue. At the extremes of soft-tissue sound speed the wavelength moves between roughly 2.7 and 3.1 m, and the ratio to a base pair stays within ten per cent of `8.6e9`. **No plausible tissue constant changes the verdict, which is the point of quoting the ratio rather than the wavelength.**

No scaling or heavy-tail claim is made anywhere in this paper, so the distribution-fitting and goodness-of-fit procedure such a claim would require [Clauset et al., 2009] does not apply. The one power law that appears — Casimir pressure as `a^-4` — has an exponent **defined by the physics, not estimated from data**, and is quoted only to show the effect is negligible.

29.13.8 13.8 Model comparison, audited against what proper model comparison requires

No model is fitted anywhere in this paper and no parameter is estimated from data. That makes several of the standard requirements inapplicable rather than unmet, and the distinction is drawn explicitly here rather than used as cover.

Requirement	Status in this paper
Out-of-sample rather than in-sample performance	**NOT APPLICABLE.** Nothing is fitted; there is no sample.
A matched null with the same free-parameter count	**NOT APPLICABLE** for the same reason. The nearest analogue that *was* run is the `c = 0` column of Section 4.4, which is the independence model against which the correlated models are read.
The number of free parameters stated explicitly	**MET.** The likelihood model has two (`p1`, `p0`), the correlation model three (`p1`, `p0`, `c`), the survivorship model three (`a0`, `s_acc`, `s_inacc`), and the retention illustration four (`eta`, `delta` for each of two conditions). Every one is a declared constant, not a fitted value.
An information criterion, so the parameter count is paid for	**NOT MET.** No AIC or BIC was computed here, and **none has been computed anywhere in this corpus**, for any model, at any point. Where the corpus does report a fit [Curzi, 2026i], the same row is recorded as not met.
Effect sizes with intervals, not only p-values	**NOT MET, and deliberately so.** No p-value is reported anywhere in this paper. The bit weights are point values with no interval, **because an interval on an assumed rate would be false precision**; the honest substitute is the four-point sensitivity check of Section 13.7, which is reported as four points and not as a confidence statement.

The two unmet rows change how the paper should be read. Because no information criterion was computed, the three-parameter correlation model of Section 4.4 has not been shown

to be *better* than the two-parameter independence model in any formal sense — it has been shown to be more realistic on grounds external to the numbers, namely the documented contact between the traditions concerned. **That is an argument from history, not from model selection, and it is labelled as one.**

29.14 14 Limitations

L-1. The rates are assumed and everything denominated in bits inherits that. `p1 = 0.30` and `p0 = 0.10` are a modelling choice with no measurement behind them, as are the illustrative likelihood ratios for studies and the survival rates of Section 5.4. **No number in Sections 4 or 5 is a measurement of anything.** What survives a change of rates is the *ordering* of the columns and the *sign* of the independence collapse, and Section 13.7 tests the first at four points only. A reader who rejects the rates should read every bit figure as illustrative and retain only R2b, which is a proof.

L-2. The shared-source probability `c` is never estimated, only assumed. The paper asserts that the Abrahamic cadence traditions have `c` "well above 0.3" and that the Australian and West African chant traditions have a genuinely lower `c`, and **neither claim is quantified or sourced.** A serious version of this apparatus would estimate transmission dependence from textual stemmatics or from documented contact, and none was attempted.

L-3. The evidence weights are calibrated against a base rate the paper does not measure. The mapping from "replicated result" to `LR ~ 1000` is a stipulation. The published replication literature [Open Science Collaboration, 2015, Ioannidis, 2005] constrains it loosely and in the unfavourable direction, and **a defensible calibration would require a base rate specific to the field of each entry, which was not constructed.**

L-4. The rejection rate is not evidence of stringency, because the candidate set was not independent of the criteria. The register was assembled by the author who set the criteria. 21.2 per cent is what this apparatus rejected from this candidate set, and it would be meaningful only if the candidates had been drawn by someone else, or before the criteria were fixed. **This is the most serious weakness in H5 and it is not repaired anywhere in the paper.**

L-5. The status vocabulary is ordinal at best, and it is used as though the categories were comparable across entries. ESTABLISHED for a prize-grade mechanism with an empty precedent column and ESTABLISHED for a moderate effect with a partially independent precedent column are the same label on very different objects. The `ROLES` field is the only defence, and it is a narrative field rather than a number. **The corpus's own demotion protocol [Curzi, 2026t] — which withdrew a quantitative claim on exactly the ground that its constituents were not instrumented on a common ratio scale — is the instrument that should be run on this vocabulary and has not been.**

L-6. The biological literature has no cluster in this corpus's bibliography and is cited nowhere in this paper. Every biological work — placebo analgesia and its opioid mediation, open-label placebo, PIEZO-family mechanotransduction, force-to-transcription coupling, sonogenetics, cochlear transduction, nucleotide data storage, guide-sequence addressing, epigenetic regulation, retroviral domestication, autophagy — is named in the text and in Appendix A with its journal, volume and year, **and carries no citation key**, because attaching an unrelated key to a biological claim would be worse than the gap. The same applies to the distributed-practice synthesis behind the spacing effect, the memory-palace imaging studies, and the anthropological account of ritual invariance. **Appendix B lists them all. This is a defect of the corpus's bibliography, not a claim that the works do not exist, and it must be closed before submission.** Its practical consequence is that **the evidence column of the register is, in this paper, unverifiable from the reference list** — a reader must go to the named sources directly.

L-7. Several register entries rest on reflexive mechanisms and are not identified. CR-01 through CR-05 and CR-13 concern belief, expectancy and self-fulfilling prophecy, and the endogenous social effect in a linear-in-means reflexive specification is not identified from obser-

vational data [Manski, 1993]; influence and selection are generically confounded in network data [Shalizi and Thomas, 2011, Lyons, 2011]. **No remedy is adopted — not randomisation, not an instrument, not partial-population variation, and not a defended stipulation of the reference group [Brock and Durlauf, 2001] — so those entries record modelling positions rather than identified empirical results.** The register's ESTABLISHED label does not and cannot repair this.

L-8. The expectancy literature the register credits is contested, and the register says so only in its own entries. The original classroom study [Rosenthal and Jacobson, 1968] drew serious methodological criticism; later work supports real but generally small effects [Jussim and Harber, 2005], and the neighbouring literature on measurement reactivity is itself contrary to casual invocation [McCambridge et al., 2014]. **The register's CR-04 states this; a reader who reads only the summary counts will not see it, which is failure mode FM-R.1 operating on this paper.**

L-9. Scale type is stipulated, not demonstrated, for every psychological construct the register touches. Michell [1997]'s objection applies directly and is not answered here; the additive conjoint conditions that would license ratio-scale treatment [Luce and Tukey, 1964, Krantz et al., 1971] were not tested, and almost certainly have never been tested for any of the constructs concerned. **The bits of Section 4 are exempt from this — they are logarithms of a defined ratio — but nothing else in the register is.**

L-10. The surplus functional's form is underdetermined, and where it appears [Curzi, 2026] the alternatives [Latané, 1981, Briscoe et al., 2006] fit the same qualitative facts. Nothing here discriminates between them. **`N_eff = 1 + (N−1)u` is not a Hill number**, and where an effective count is used in Section 6.1 it is this one and it is named [Hill, 1973, Jost, 2006].

L-11. The Lucas critique applies to the apparatus itself. A framework asserting that systems change under observation, then publishing a register of weights, has created exactly the incentive it warns about [Lucas, 1976, Hacking, 1995]. Section 3.5 states the narrow position taken — the weights are declared constants rather than estimated parameters — and **that defence fails immediately for any later use that fits the weights to the entries.**

L-12. Two tally defects in the source register are reported in Section 11.5 and were not corrected. They are printed rather than harmonised because the recomputation that found them is itself part of what a reader should be able to check.

L-13. What would falsify this apparatus rather than the claims it evaluates. If a bridge graded REJECTED here were later established by measurement, the criteria would have been too strict. If a bridge graded ESTABLISHED were shown to rest on the wrong channel, error E7 would have escaped a procedure built to catch it. And **if a second analyst, applying these specifications to the same candidate set, returned a materially different distribution of verdicts, the method would be shown to be analyst-dependent** — which is the test that matters most and which has not been run.

29.15 15 Conclusion

This paper separates the two columns of a concordance apparatus and prices them in a common unit. **The evidence column carries what could have come out otherwise; the precedent column carries only that an idea was independently arrived at and preserved. Two sources agreeing does not double the evidence when only one of them is evidence**, and under the assumed recording rates 0.30 and 0.10 a single traditional record is worth `LR = 3.000`, or **1.585 bits**, against **9.966 bits** for a replicated result and **19.932 bits** for a prize-grade mechanism. Double-counting the record turns an honest 3,000:1 into 1,000,000:1 — **an overstatement factor of 333.3, which is 8.38 bits of confidence not earned.**

Correlated sources are worse than merely redundant. At a ten per cent probability of shared origin, `k` agreeing traditions buy **1.585, 2.546, 2.317, 1.886 and 1.685 bits** for `k = 1` to `5`: **the total decreases after `k = 2`**, and for every `c > 0` the sum converges to the 1.585 bits a single

tradition already supplied. Past two agreeing sources, further agreement is better explained by contact than by discovery.

The instrument that produced the precedent column is characterised exactly. It has unbounded `n` and no control, so repetition drives its variance to zero and leaves its bias untouched; a single uncontrolled observer registers a true standardised shift of `d = 0.20` with probability **0.2858** and of `d = 1.00` with probability **exactly 0.5000**, against roughly **392 subjects per group** a modern design needs at `d = 0.20`. And the preservation filter dominates the source: the same corpus arrives **7.9 per cent or 94.5 per cent accurate — a factor of 11.97 — according only to whether its failures were visible.**

Four bridges are worked to their verdicts and one is rejected in print. Cadence is ESTABLISHED; vibration-to-cell is GRADED, established at the coarse grain and speculative at the fine; fasting-to-autophagy is SPECULATIVE with the number that would settle it stated; and the claim that a tone at 528 Hz repairs DNA is **REJECTED** on a wavelength of 2.9167 m against a 0.34 nm base pair (a ratio of `8.578e9`), a quantum-to-thermal energy ratio of `8.17e−11` at 310 K, and molecular vibrational bands 34.7 to 37.3 octaves above the audible band. **That rejection binds the framework's own sonification pipeline, which may not claim what it has just refused to let anyone else claim.**

Across the register, thirty-two of forty-two entries are ESTABLISHED, six REJECTED, two SPECULATIVE and two badged POSTULATE; fourteen of forty-two (33.3 per cent) carry a qualifier; and seven of the thirty-three attempted bridges (21.2 per cent) fail structural correspondence outright. A register with no rejections would mean the criteria were never binding. **The apparatus is built so that it can reject, and the claim it rejects at greatest length is the one nearest to the corpus's own subject matter.**

29.15.1 Declarations

Data and code availability

No external dataset was collected for this paper. Every number reported here is one of three things: arithmetic on assumed rates that are stated in the text at the point of use (Sections 4 and 5); arithmetic on standard physical constants (Section 7.5); or a count obtained by recomputing the `STATUS` and `C1` fields of every entry of the register reproduced in Appendix A (Section 11.5). The register itself is a bibliographic and methodological artefact, not data about the world. **The five figures are generated by the scripts named in each caption, and those scripts contain no data beyond the numbers printed in this paper.** No result here should be cited as evidence about any tradition, any practitioner population, or any biological system.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest. The concordance apparatus of Section 3.1 is the author's own earlier work, adopted here without modification; that provenance is stated in the text and is the reason the apparatus is described rather than defended.

Author contributions

Single author. Michael Laurence Curzi designed the apparatus, specified the criteria, performed the arithmetic, assembled and audited the register, produced the figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The paper discusses psychological, medical and religious material as an object of study and makes no clinical, therapeutic or diagnostic claim; the standing refusals of Section 3.7 are constraints on what the framework will assert and are not reports of any study performed. Traditions are described from published sources with their transmission paths named, and no living community was consulted or observed.

Preregistration

This paper reports no preregistered empirical test, and none is claimed. The Step 0 device of Section 3.1 — fixing the evidence reading in writing before the precedent column is opened — is a procedural analogue of preregistration and is **not** a substitute for one: it constrains the order in which sources are consulted and constrains nothing about the analysis. **Whether Step 0 was in fact honoured for each of the register's forty-three entries cannot be verified from the register, and that is stated here rather than implied otherwise.**

29.16 Appendix A The concordance register

Reproduced in full, including its rejections. **Numbering is stable: `CR–nn` is a permanent identifier and is cited as such throughout the corpus.** The twelve fields are those specified in Section 8.2, and the `USED IN` field records the chapter numbering of the source manuscript, retained unchanged as a provenance record (Section 8.2). **Every biological source named in this appendix carries no citation key, for the reason given in Section 3.6 and listed in Appendix B.**

Entries are grouped by subject.

29.16.1 PART A — BELIEF AND ITS CHANNEL

CR-01 · Belief has at least one complete, physical, verified channel to bodily outcome

CLAIM	Expectation of relief produces measurable analgesia through endogenous opioid signalling.
(+1) SOURCE	Levine, Gordon & Fields, *The Lancet*, 1978.
MEASURED	Placebo analgesia in post-operative dental pain was reversed by the opioid antagonist naloxone, implicating endogenous opioid mediation. The finding is a mechanism of analgesia, not a general theory of placebo; placebo effects are heterogeneous across conditions.
(-1) SOURCE	Therapeutic ritual in documented medical traditions: the physician's manner, the setting, the prescribed form of the encounter, attested from the Hippocratic corpus onward.
RECORDED	That the practitioner's manner and the form of the encounter affect outcome.
ROLES	(+1) evidence, strong, and it carries the claim alone. (-1) precedent, low independence (Old World medical traditions are in continuous contact), ~1.6 bits.
C1	**PASS** — homomorphism.
C2	**PASS** — the bridge predicts that blocking the opioid system attenuates the effect. That prediction was made and tested; it is the 1978 result.
C3	**PASS** — +1 side: identifies a manipulable channel.

0p	Integrates without either side lying: both describe expectation as an input with gain.
STATUS	**ESTABLISHED**
USED IN	Ch 76, 77, 78

CR-02 · The channel does not require deception

CLAIM	Placebo effects can be obtained when the patient is truthfully told they are receiving a placebo.
(+1) SOURCE	Kaptchuk et al., *PLoS ONE*, 2010 (open-label placebo in irritable bowel syndrome), and subsequent open-label trials.
MEASURED	Symptom improvement in IBS patients openly randomised to a placebo described to them as an inert pill. Effect sizes are modest and the literature is active; this must be stated wherever the result is used.
(-1) SOURCE	Contemplative and devotional traditions in which the practitioner is fully informed of form and content — recitation, posture, fasting, disclosed ritual — with no concealed element. Attested across monastic traditions of several lineages.
RECORDED	That disclosed practice, undertaken knowingly, produces reported change of state.
ROLES	(+1) evidence, moderate and active. (-1) precedent, partially independent, ~2 bits.
C1	**PASS** — homomorphism on the structure “the participant’s knowledge of the procedure does not abolish the effect.”
C2	**PASS** — predicts that concealment is not required for efficacy, and therefore that a disclosed protocol should not underperform a concealed one by the mechanism’s own logic. Testable and partially tested.
C3	**PASS**, decisively.
0p	Integrates cleanly.
STATUS	**ESTABLISHED**, with effect sizes explicitly modest.
USED IN	Ch 78, 82, 83, 85

CR-03 · A belief propagates in the world when it is coupled to a true feature of the world

CLAIM	Self-fulfilling prophecies work through a real structural feature of the system, not through the belief alone.
(+1) SOURCE	Merton, *The Self-Fulfilling Prophecy*, Antioch Review, 1948.
MEASURED	Not a measurement — a sociological mechanism with worked cases, the bank run being canonical. The mechanism’s content is the coupling: banks really are fractional-reserve, so a belief in insolvency, acted on, produces insolvency. Without the coupling the belief does not propagate.
(-1) SOURCE	The doctrine that a spoken commitment, oath or vow binds the speaker’s subsequent conduct — attested in legal, monastic and contractual forms across many independently documented societies.
RECORDED	That a stated intention changes what the stater afterwards does, reliably enough to build institutions on.

ROLES	(+1) evidence, and it is the load-bearing side. (-1) precedent, genuinely independent across several legal traditions, ~2 to 3 bits.
C1	**PASS** — homomorphism.
C2	**PASS** — predicts that removing the coupling removes the effect: deposit insurance should suppress bank runs, and it substantially does.
C3	**PASS** — supplies the postulate's conditional clause in operational form.
0p	Integrates. Merton's mechanism IS the conditional clause of the postulate, discovered independently in 1948, and that is the entry's whole point.
STATUS	**ESTABLISHED**
USED IN	Ch 76, 77, 79

CR-04 · Expectancy effects on performance are real, modest, and methodologically contested

CLAIM	Another person's expectation of one's performance measurably affects that performance.
(+1) SOURCE	Rosenthal & Jacobson, 1968, and the subsequent expectancy literature.
MEASURED	The original classroom study reported teacher-expectancy effects on pupil outcomes. It drew serious methodological criticism, stated here rather than omitted. Later work supports real but generally small effects, largest where the expectation changes the expecter's behaviour toward the target.
(-1) SOURCE	Craft and guild apprenticeship: the master's assessment recorded as determining the apprentice's assignment, instruction and progression.
RECORDED	That the master's judgement shapes what the apprentice becomes — an institutional fact about how training is allocated, which is precisely the behavioural channel.
ROLES	(+1) evidence, weak-to-moderate and contested. (-1) precedent, moderately independent, ~2 bits. Neither column is strong here and the entry says so.
C1	**PASS** — homomorphism through the behavioural channel: expectation changes allocation of instruction, which changes outcome.
C2	**PARTIAL** — predicts the effect should vanish where the expecter cannot alter their own behaviour toward the target. Rarely tested in that form.
C3	**PASS** — the intervention point is the behaviour, not the belief.
0p	Integrates only in the behavioural reading. A reading in which expectation acts on the target with no behavioural intermediary is not supported and is not used.
STATUS	**ESTABLISHED** for the behavioural channel, at modest magnitude. **SPECULATIVE** — and this is the sentence that marks it — for any non-behavioural channel.
USED IN	Ch 79, 80

CR-05 · Perception is constructed in part from prior expectation

CLAIM	What is perceived is a joint function of sensory input and prior expectation, not of input alone.
(+1) SOURCE	The predictive-processing literature; Helmholtz's unconscious inference as the historical origin; the psychophysical results treated in the corpus's psychophysical papers.

MEASURED	Context-dependent illusions, priming, bistable percepts responding to instruction, and the psychophysical scaling results of Weber, Fechner and Stevens. The strong theoretical frameworks built on these findings are more contested than the findings; this entry cites the findings.
(-1) SOURCE	Contemplative traditions whose central technical claim is that attention and expectation shape what is experienced. Buddhist and Stoic sources are best documented and operationally specified.
RECORDED	That trained attention reproducibly changes the content of experience, in trainable stages with recognisable milestones.
ROLES	(+1) evidence, strong for the findings. (-1) precedent, partially independent, ~2 bits — and here the –1 column carries unusually rich operational detail, which is the job §4.5 assigns it.
C1	**PASS** — homomorphism, and close to isomorphism at the level "percept = f(input, prior)."
C2	**PASS** — predicts that manipulating priors changes reported percept at fixed input. Tested extensively; holds.
C3	**PASS** on both sides.
0p	Integrates.
STATUS	**ESTABLISHED**
USED IN	Ch 76, 77, 81

CR-06 · Belief-based interventions with weak channels produce weak effects

CLAIM	Interventions that change belief without changing a strong channel produce small effects in large samples.
(+1) SOURCE	Sisk, Burgoyne, Sun, Butler & Macnamara, *Psychological Science*, 2018 (meta-analyses of growth-mindset interventions).
MEASURED	Very small overall effects of growth-mindset interventions on academic achievement, with some indication of larger benefit in high-risk and low-socioeconomic-status subgroups. Related: stereotype-threat findings have had replication difficulty.
(-1) SOURCE	None admissible. Comparing intervention magnitudes requires a control group and population statistics, both outside the envelope of §5.3.
RECORDED	Nothing.
ROLES	(+1) evidence, strong (large replications). (-1) precedent: none, 0 bits, by the nature of the instrument.
C1	**N/A** — no bridge is proposed.
C2	**PASS** — predicts that the postulate's effects scale with channel gain, and that channel-free interventions approach zero effect.
C3	**PASS** — it is the diagnostic tool of the belief-failure analysis (CCP-035).
0p	N/A.
STATUS	**ESTABLISHED**, and it functions as a REJECTION of the strong form of the postulate.
USED IN	Ch 80, 81, 87

CR-07 · Belief does not act at a distance on another person's body

CLAIM	That directed belief or intention alters another person's physiology without any physical intermediary.
(+1) SOURCE	None offered: this author is aware of no replicated result establishing such an effect. **No systematic survey is claimed; what is asserted is the absence of a known replicated result.**
MEASURED	Nothing that survives replication, to this author's knowledge.
(-1) SOURCE	Numerous traditions record claims of distant action. Dated and named where used; not credited.
RECORDED	Claims of distant efficacy. Envelope check (§5.3): P2 and P4 both fail; the visible-failure rule gives no correction, since failure is always attributable to the practitioner (T5). The exact profile the preservation filter selects on memorability.
ROLES	(+1): empty. (-1) precedent: 0 bits — the claim sits outside the instrument's operating envelope, so its presence in the record carries no information about the world.
C1	**FAIL** — no shared structure; no mechanism on either side.
C2	**FAIL** — no falsifiable prediction is offered by proponents in a form that could fail.
C3	**FAIL**
0p	Does not integrate.
STATUS	**REJECTED**, and additionally **REFUSED** as a claim the framework will make (standing refusal, §3.7).
USED IN	Ch 80, 87

CR-08 · The postulate

CLAIM	Belief creates reality if the belief is grounded in truth — where "grounded in truth" means that a real channel exists between the belief and the outcome.
(+1) SOURCE	Not a measured result. The postulate's *channels* are measured and are entries CR-01 through CR-05; the postulate itself is the author's.
MEASURED	Each channel separately. The postulate is not established by summing them; it states what they have in common.
(-1) SOURCE	Traditions holding that intention or faith is efficacious. Named and dated where used.
RECORDED	The unconditional form — that belief is efficacious as such.
ROLES	(+1) evidence for the individual channels. (-1) precedent for the unconditional claim, which the framework does not adopt — so the precedent supports a different proposition and is credited at ~0 bits *for this claim*.
C1	**PARTIAL** — the traditional claim and the author's claim differ at precisely the load-bearing conjunction.
C2	**PASS** — the conditional form predicts that effects appear where a channel exists and vanish where it does not. CR-01 to CR-06 are that prediction being tested.
C3	**PASS** — the grounding protocol (CCP-035) is the operational form.
0p	Integrates only in the conditional form. The unconditional form is refused.
STATUS	**POSTULATE**, badged in the same sentence at every use, per the badging rule of §8.3.
USED IN	Ch 76, 77, 80, 81, 84, 88, 90

29.16.2 PART B — THE UNCERTAINTY

CR-09 · Heisenberg's relation is a property of conjugate variables, not of measurement disturbance

CLAIM	`dx . dp >= hbar / 2` follows from the Fourier relationship between position and momentum representations, and is not a statement that the measurer disturbs the measured.
(+1) SOURCE	Standard quantum mechanics; Heisenberg 1927 and the subsequent formalisation. `hbar / 2 = 5.27e-35 J s`.
MEASURED	A theorem of the formalism, consistent with a century of experiment. The disturbance reading is a widespread misconception and the framework corrects it explicitly.
(-1) SOURCE	None admissible. No traditional source records a conjugate-variable relation, and no first-person instrument could have detected one (§5.3, exclusion: requires an instrument).
RECORDED	Nothing.
ROLES	(+1) evidence, decisive. (-1): none, 0 bits — an empty `-1` column is a normal outcome, not a gap to be filled.
C1	**N/A** — no bridge proposed.
C2	**PASS** as physics.
C3	**PASS** — it is the correction the corpus's conjugate-uncertainty analysis exists to make (CCP-032).
0p	N/A.
STATUS	**ESTABLISHED** (as physics; no concordance drawn).
USED IN	Ch 70, 71, 72

CR-10 · The energy-time relation has different status from position-momentum

CLAIM	`dE . dt >= hbar / 2` is not the same kind of statement as `dx . dp >= hbar / 2`, because time is not an operator in quantum mechanics.
(+1) SOURCE	Standard quantum mechanics.
MEASURED	The formal asymmetry is structural. Several energy-time relations exist with different derivations and interpretations; they are not one relation.
(-1) SOURCE	None admissible.
RECORDED	Nothing.
ROLES	(+1) evidence. (-1): none, 0 bits.
C1	**N/A**
C2	**PASS** as physics.
C3	**PASS** — it blocks a common over-extension in which the energy-time form is used to license claims about events in time.
0p	N/A.
STATUS	**ESTABLISHED** (as physics; no concordance drawn).
USED IN	Ch 70, 72

CR-11 · The Gabor limit is the real uncertainty principle of this work

CLAIM	`dt . df >= 1 / (4 pi) = 0.079577`. Every wave, including sound, obeys it. It is the same Fourier mathematics as Heisenberg with no quantum content.
(+1) SOURCE	Gabor, *Theory of Communication*, Journal of the IEE, 1946.
MEASURED	A theorem of Fourier analysis, verifiable on any spectrum analyser. `1/(4 pi) = 0.07957747154594767`. Window against minimum frequency resolution: `1 ms -> 79.577 Hz`; `10 ms -> 7.958 Hz`; `100 ms -> 0.796 Hz`; `1000 ms -> 0.080 Hz`.
(-1) SOURCE	Instrument-making practice that a very short sound has no definite pitch — encoded in classing percussion as unpitched, in treating the attack transient as distinct from the sustained tone, and in requiring a tuning reference to be held. Attested independently in several regions.
RECORDED	That a tone must be held to have a pitch. Recorded operationally — as how long to sustain a reference note, and as which instruments do and do not carry pitch — never as a relation between two quantities.
ROLES	(+1) evidence, decisive; it is a theorem. (-1) precedent, genuinely independent across instrument-making traditions, ~2 to 3 bits, and it is precedent for the phenomenon, not for the relation.
C1	**PASS** — homomorphism. “Held long enough” is `dt`; “definite pitch” is small `df`.
C2	**PASS** — predicts an exact reciprocal trade with a computable constant, which is measurable and holds.
C3	**PASS** — it hard-constrains genomic-audio design (CR-12) and the rendering of essence frequencies.
0p	Integrates with unusual force: the tradition recorded the qualitative fact for centuries and the mathematics supplies the constant. Neither side has to overstate.
STATUS	**ESTABLISHED**
USED IN	Ch 65, 69, 71, 72

CR-12 · The Gabor limit constrains genomic-audio rendering speed

CLAIM	Essence frequencies and sonified spectra cannot be rendered arbitrarily fast without spectral smearing; time resolution and frequency resolution trade against each other by CR-11.
(+1) SOURCE	CR-11, applied.
MEASURED	Arithmetic. Two tones `1 Hz` apart need a window of at least `0.0796 s`; `0.1 Hz` apart, at least `0.796 s`. The essence law’s spacing follows `f(n) = k n^2`, `k = 0.483425732988414 Hz`, so the resolution demand is set by the smallest interval in the rendered set.
(-1) SOURCE	Liturgical and recitational traditions specifying a minimum duration per syllable or tone and treating hurried recitation as defective. Vedic recitation is the most exactingly specified case; plainchant the best-documented Western one.
RECORDED	Prescribed pace, and the doctrine that speed degrades the practice.
ROLES	(+1) evidence, decisive (arithmetic). (-1) precedent, moderate independence, ~2 bits. The tradition’s reason for the rule is not adopted; only the rule.
C1	**PARTIAL** — rule and theorem coincide in prescription, not in content; the tradition had many reasons to slow recitation unrelated to spectra.
C2	**PASS** — predicts a computable minimum duration per rendered tone, which the sonification pipeline (CCP-031) must satisfy.
C3	**PASS** — it is a design constraint with a number attached.

0p	Integrates weakly.
STATUS	**ESTABLISHED** as an engineering constraint. **SPECULATIVE** — marked here, in this sentence — as any claim that traditional pacing rules were responsive to spectral resolution.
USED IN	Ch 65, 69, 71

CR-13 · The future as conjugate variable

CLAIM	The uncertainty attaching to a future event shares conjugate-variable structure with Heisenberg's relation.
(+1) SOURCE	None. This is an analogy and is badged as one everywhere it appears.
MEASURED	What is genuinely shared: conjugate-variable structure, the impossibility of simultaneous sharpness in both members of a pair, and the Fourier trade-off. What is NOT shared: `hbar`, quantization, non-commuting operators, or any experimental result.
(-1) SOURCE	Divinatory and prognostic traditions treating the future as a distribution of possibilities rather than a single determined outcome.
RECORDED	That prognostication yields tendencies rather than verdicts, and that specificity about *when* costs specificity about *what*.
ROLES	(+1): none — analogy only, 0 bits of evidence. (-1) precedent, low independence, ~1.6 bits. Both columns are weak.
C1	**PARTIAL** — analogy in the grading of §3.2, not homomorphism.
C2	**PARTIAL** — the sharpening trade-off is a real prediction about forecasting, but it does not require quantum mechanics and is better derived without it.
C3	**PASS** — as a way of thinking about forecasts.
0p	Integrates only as analogy.
STATUS	**SPECULATIVE** — and this is the sentence that marks it: the identification of future-event uncertainty with quantum uncertainty is speculative and is presented as analogy with its badge on, never as identity. The corpus's comfort-weighted amplitude already treats the future as an amplitude-weighted ensemble, which is the framework's own treatment and needs no quantum mechanics to stand.
USED IN	Ch 72

CR-14 · Decoherence destroys macroscopic superposition extremely fast

CLAIM	Macroscopic superpositions do not persist, because interaction with the environment destroys phase coherence on timescales far shorter than any biological process.
(+1) SOURCE	The decoherence literature; environment-induced superselection.
MEASURED	Decoherence is observed and its scaling with system size and coupling confirmed in controlled systems. Timescale estimates are strongly model-dependent; this entry quotes no number for biological systems and the corpus's quantum-formalism analysis gives estimates (CCP-033) with their assumptions attached.
(-1) SOURCE	None admissible.
RECORDED	Nothing.
ROLES	(+1) evidence, strong. (-1): none, 0 bits.

C1	**N/A** — this entry is an objection, not a bridge.
C2	**PASS** — predicts that any proposed quantum-coherent biological mechanism must specify how coherence is protected, and most cannot.
C3	**PASS** — it is the filter the corpus's quantum-formalism analysis applies (CCP-033).
0p	N/A.
STATUS	**ESTABLISHED**, and it functions as the principal objection this series confronts rather than avoids (standing refusal, §3.7).
USED IN	Ch 73, 74

CR-15 · Consciousness-causes-collapse

CLAIM	That observation by a conscious mind is what collapses the quantum wave function.
(+1) SOURCE	The von Neumann-Wigner interpretation. This is a minority historical interpretation, not a mainstream position, and the entry states that as its principal content.
MEASURED	Nothing distinguishes it experimentally from other interpretations. The measurement problem is genuinely open; it is not open in the direction of licence. CR-14 is the standing objection.
(-1) SOURCE	Traditions holding that mind participates in constituting the world. Named where used.
RECORDED	Doctrinal positions on the relation of mind to world. Envelope check: P3 fails — nothing in a first-person instrument could bear on this question.
ROLES	(+1): a minority interpretation, not evidence. (-1) precedent: 0 bits for this claim, since the question lies outside the instrument's envelope.
C1	**FAIL** — the traditional doctrine and the interpretation share vocabulary, not structure.
C2	**FAIL** — no distinguishing prediction.
C3	**FAIL**
0p	Does not integrate.
STATUS	**REJECTED** as a basis for any claim in this series. Presented in the corpus's quantum-formalism analysis (CCP-033) as a minority historical interpretation and as nothing more.
USED IN	Ch 73, 74

29.16.3 PART C — BOUNDARY AND SHAPE

CR-16 · The Casimir force exists and has been measured

CLAIM	`F/A = −pi^2 hbar c / (240 a^4)` between ideal parallel plates.
(+1) SOURCE	Casimir, 1948 (prediction); Lamoreaux, *Physical Review Letters*, 1997; Mohideen & Roy, *Physical Review Letters*, 1998 (measurement).
MEASURED	The predicted force, at the predicted scaling, in the sub-micrometre regime. Computed pressures: `10 nm -> 1.3e5 Pa` (1.28 atm); `100 nm -> 13 Pa`; `1 um -> 1.3e−3 Pa`; `1 mm -> 1.3e−15 Pa`.
(-1) SOURCE	None admissible.
RECORDED	Nothing.

ROLES	(+1) evidence, decisive. (-1): none, 0 bits.
C1	**N/A**
C2	**PASS** — the `a^-4` scaling is the prediction and it was confirmed.
C3	**PASS**
0p	N/A.
STATUS	**ESTABLISHED**
USED IN	Ch 75

CR-17 · The Casimir force does nothing at biological or architectural scale

CLAIM	That Casimir forces have any consequence in a cell or a room.
(+1) SOURCE	CR-16, applied.
MEASURED	The `a^-4` scaling. At `1 um` the pressure is `1.3e-3 Pa`; at `1 mm` it is `1.3e-15 Pa`, which is `1.28e-20` atmospheres. Twenty orders of magnitude separate the measured regime from the scale of a room.
(-1) SOURCE	Not applicable; no traditional source is invoked and none would help.
RECORDED	Nothing.
ROLES	(+1) evidence, decisive against. (-1): none.
C1	**FAIL** by scale (error E6).
C2	**FAIL** — no biological prediction survives the arithmetic.
C3	**FAIL**
0p	Does not arise.
STATUS	**REJECTED.** This entry is mandatory wherever Casimir appears in this series (standing refusal, §3.7) and the pressure table is printed with it.
USED IN	Ch 64, 75, 87

CR-18 · Boundary geometry alone determines which modes can exist, with measurable consequence

CLAIM	The shape of a bounded region selects the set of modes it admits, and that selection has physical consequence without anything being added or powered.
(+1) SOURCE	The Casimir effect (CR-16) as the most rigorous demonstration; room acoustics and modal analysis as the everyday one; the vibrating string as the elementary one.
MEASURED	Casimir: a force arising purely because the plates exclude every mode whose wavelength does not fit between them. Acoustics: room modes computed from dimensions alone and measured in the room. **Nothing is added; nothing is powered; the shape does all of it.**
(-1) SOURCE	Monochord practice: the Pythagorean division of a string, in which pitch is set by boundary position alone, attested from antiquity; and the deliberate proportioning of resonant chambers in temple, chapel and hall building.
RECORDED	That where you put the boundary determines what tone you get, and nothing else does. Recorded operationally as ratios of string length.
ROLES	(+1) evidence, decisive. (-1) precedent, and unusually strong precedent: independent instrument traditions in several regions arrived at boundary-determined pitch, ~2 to 3 bits.
C1	**PASS** — close to isomorphism.

C2	**PASS** — predicts computable mode sets from geometry alone, in both the acoustic and the electromagnetic case. Tested continuously.
C3	**PASS** — it is the design principle of the corpus's boundary-geometry material.
0p	Integrates without strain.
STATUS	**ESTABLISHED**
USED IN	Ch 13, 14, 33, 75

CR-19 · "The shape is the function"

CLAIM	That the framework's recurring thesis — that geometry alone determines what a system can do — has a rigorous physical exemplar.
(+1) SOURCE	CR-18.
MEASURED	Mode selection by boundary, in the Casimir and acoustic cases.
(-1) SOURCE	Architectural and instrument-making traditions as in CR-18; and the design-pattern tradition descending from Alexander's pattern language (1977), which is modern but is the framework's direct ancestor.
RECORDED	That proportion determines behaviour, recorded as buildable rules.
ROLES	(+1) evidence for the physical exemplar. (-1) precedent, ~2 bits.
C1	**PASS** for wave systems.
C2	**PASS** within wave systems; the transfer to social and institutional geometry is a claim about model form, not about physics, and is badged as such (§3.4).
C3	**PASS**
0p	Integrates, with the badge.
STATUS	**ESTABLISHED** for wave systems. **POSTULATE** — badged here — for the transfer of "the shape is the function" to social and institutional systems, which is a claim about model form and not about physics.
USED IN	Ch 13, 14, 33, 75, 88

29.16.4 PART D — THE CELL AS RECEIVER

CR-20 · Mechanical force opens ion channels: PIEZO1 and PIEZO2

CLAIM	There is a real, established pathway by which mechanical force becomes a cellular signal.
(+1) SOURCE	Coste et al., *Science*, 2010 (identification of Piezo1 and Piezo2 as essential components of mechanically activated cation channels). Nobel Prize in Physiology or Medicine, 2021, to Ardem Patapoutian for receptors for touch, shared with David Julius for temperature receptors.
MEASURED	That these proteins are required components of mechanically gated cation currents, and that mechanical force opens them, admitting ions and initiating signalling.
(-1) SOURCE	Traditions of therapeutic touch, pressure and manipulation — massage and bodywork attested independently in several regions — recording that mechanical pressure produces reproducible physiological change.
RECORDED	That pressure produces effect, recorded as technique: where, how hard, how long. No mechanism is recorded and none could have been.

ROLES	(+1) evidence, prize-grade, ~20 bits on the register's scale. (-1) precedent, genuinely independent across regions, ~2 to 3 bits — against twenty.
C1	**PASS** coarse: mechanical input, transduction, physiological output.
C2	**PASS** — since sound is mechanical, the bridge predicts that acoustic pressure at sufficient amplitude can engage mechanotransductive pathways. This is the basis of sonogenetics (CR-22) and it holds.
C3	**PASS**
0p	Integrates at the coarse grain only, and the entry is used only at that grain.
STATUS	**ESTABLISHED** — that a real pathway exists from pressure wave to cellular signal. This is the strongest single result underwriting Part XIII and it is also the one most at risk of over-extension; see CR-26.
USED IN	Ch 64, 65, 68

CR-21 · Surface force reaches the nucleus, stretches chromatin, and upregulates transcription

CLAIM	Mechanical input at the cell surface produces a change in gene expression by a direct physical route.
(+1) SOURCE	Tajik et al., *Nature Materials*, 2016.
MEASURED	That force applied at the cell surface transmits through the cytoskeleton to the nucleus, directly stretches chromatin, and upregulates transcription. This is mechanical input to gene expression, published, direct, and it is the single most important evidence entry in this series.
(-1) SOURCE	None admissible. No traditional source records anything at this scale; the exclusion is P3, requires an instrument.
RECORDED	Nothing.
ROLES	(+1) evidence, strong. (-1): none, 0 bits.
C1	**N/A** — no bridge proposed.
C2	**PASS** — predicts a force threshold and a time course for transcriptional response, both measurable.
C3	**PASS** — it establishes that the mechanical-to-genomic route is not hypothetical.
0p	N/A.
STATUS	**ESTABLISHED** — and its limit is CR-26: it establishes *that* mechanical input reaches transcription, and nothing about *sequence-specific* control.
USED IN	Ch 64, 65, 68, 69

CR-22 · Ultrasound can activate specified channels: sonogenetics

CLAIM	An acoustic stimulus can be used to activate a chosen, previously specified channel in a cell.
(+1) SOURCE	Ibsen, Tong, Schutt, Esener & Chalasani, *Nature Communications*, 2015 (TRP-4 in *C. elegans*); Duque et al., 2022 (hsTRPA1).
MEASURED	Activation, by applied ultrasound, of neurons expressing an ultrasound-sensitive channel. The selectivity comes from which cells express the channel — engineered in advance, not carried by the acoustic signal.
(-1) SOURCE	None admissible.
RECORDED	Nothing.

ROLES	(+1) evidence, strong. (-1): none, 0 bits.
C1	**N/A**
C2	**PASS** — predicts that acoustic addressing of cells is achievable when the address is pre-installed as a channel.
C3	**PASS** — and it sharpens CR-26 rather than softening it.
0p	N/A.
STATUS	**ESTABLISHED**, with the selectivity caveat stated in the entry and repeated at every use.
USED IN	Ch 64, 66, 68

CR-23 · Biology builds exquisite mechanical receivers

CLAIM	Extremely sensitive mechanical reception is achievable in tissue; cochlear hair cells are the existence proof.
(+1) SOURCE	The cochlear mechanotransduction literature; hair-cell tip-link gating.
MEASURED	That hair-cell stereocilia deflections of very small magnitude gate transduction channels via tip links, and that the cochlea performs frequency analysis mechanically along the basilar membrane.
(-1) SOURCE	Traditions treating the ear as the reference measuring device — tuning by ear to beat rates, attested independently wherever fixed-pitch instruments were built.
RECORDED	That the ear resolves frequency differences finer than any available instrument could, and that it does so by hearing beats. Recorded as tuning procedure, which is operationally exact and remains in use.
ROLES	(+1) evidence, strong. (-1) precedent, independent across regions, ~2 to 3 bits, and it is precedent for the ear's resolving power, not for its mechanism.
C1	**PASS** at the level "biological tissue performs high-resolution mechanical frequency analysis."
C2	**PASS** — beat-rate tuning predicts a frequency-discrimination threshold, which is measurable and matches the recorded practice.
C3	**PASS**
0p	Integrates.
STATUS	**ESTABLISHED**
USED IN	Ch 64, 65

CR-24 · Chant produces measured autonomic effects, and the identified channel is respiration

CLAIM	Sustained recitation changes cardiovascular autonomic state, principally by enforcing a slow breathing rate.
(+1) SOURCE	Bernardi et al., *BMJ*, 2001.
MEASURED	That recitation of the Latin rosary and of a yoga mantra both brought respiration to approximately six breaths per minute — a period of `10.0 s`, a rate of `0.100 Hz` — and enhanced baroreflex sensitivity.

(-1) SOURCE	Recitational traditions: Vedic recitation; Christian plainchant under the canonical hours; Tibetan overtone chanting; Sufi *dhikr*; West African and Afro-diasporic drumming and song; Australian Aboriginal drone practice. **The West African and Australian instances have no plausible shared source with the Indo-European ones, so independence is genuine.**
RECORDED	That sustained vocalisation at length changes bodily and mental state reproducibly, and detailed technique for producing it.
ROLES	(+1) evidence, moderate. (-1) precedent, genuinely independent, ~2 to 3 bits — the strongest precedent in the register. It is precedent for the phenomenon and not for its cause.
C1	**PASS** — homomorphism: sustained periodic vocal effort, an enforced respiratory period, an autonomic state change.
C2	**PASS** — predicts that the effect tracks respiratory rate rather than acoustic content, which is testable by pacing respiration independently of the sound.
C3	**PASS** on both sides.
0p	Integrates, provided the –1 column is credited for the effect and not for a mechanism.
STATUS	**ESTABLISHED**
USED IN	Ch 61, 64, 82, 85

CR-25 · Chant acts through PIEZO-family mechanotransduction

CLAIM	That the state changes traditions attribute to chant are mediated to a material degree by mechanically gated ion channels.
(+1) SOURCE	None directly. CR-20 establishes the pathway exists; no source known to this author measures its contribution to chant.
MEASURED	Nothing on this specific attribution. At least four better-supported channels compete: respiratory entrainment (CR-24), the social setting, the attentional demand, and expectancy (CR-01).
(-1) SOURCE	As CR-24.
RECORDED	The effect, not the cause.
ROLES	(+1): empty for this attribution. (-1) precedent, ~2 to 3 bits, and it does not bear on mechanism at all.
C1	**FAIL** at the fine grain — the –1 side records no mechanism, so there is nothing to correspond to.
C2	**PASS** — the separating experiment is specifiable: deliver the same acoustic exposure to a passive listener whose respiration is paced differently, against a vocaliser producing it, and see whether the effect survives removal of the breathing pattern.
C3	**PARTIAL**
0p	Does not integrate at the fine grain.
STATUS	**SPECULATIVE** — and this is the sentence that marks it: attributing chant's effects to PIEZO-family mechanotransduction is speculative, unsupported by any measurement known to this author, and in competition with four better-supported channels. Retained because it is falsifiable and the falsifying experiment is specified. **The register's worked instance of error E7.**
USED IN	Ch 61, 64, 87

29.16.5 PART E — THE OPEN PROBLEM

CR-26 · There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal

CLAIM	The gap, stated exactly as the Methodology Protocol states it.
(+1) SOURCE	The absence of one, across the mechanotransduction and sonogenetics literatures. **Everything upstream is real and is in this register:** DNA data storage (CR-28), gRNA-format addressing and Cas12a self-processing (CR-29), latent capability (CR-31), epigenetic regulation (CR-30), mechanotransduction (CR-20, CR-21, CR-22).
MEASURED	Each upstream component separately, and none of them together. CR-22's selectivity is engineered into which cells express the channel, not carried by the signal. CR-27's spatial argument applies: at `528 Hz` in soft tissue the wavelength is `2.9167 m`, `8.6e9` times a base pair, so an audio-band wave imposes no structure at sequence scale.
(-1) SOURCE	Not applicable. No precedent could bear on an open mechanistic question (§5.3, exclusion: hidden mechanism).
RECORDED	Nothing, necessarily.
ROLES	(+1): the gap itself, well specified. (-1): none, 0 bits, by the nature of the question.
C1	**N/A** — this entry is a boundary, not a bridge.
C2	**PASS** — it is maximally predictive in the negative direction: it says exactly what an experiment would have to show.
C3	**PASS** — naming the gap sharply is the most useful thing this series can do.
0p	N/A.
STATUS	**ESTABLISHED** as an open problem. This entry is restated wherever this series touches the channel (standing clause, §3.7), and no entry in this register may be combined with another to close it.
USED IN	Ch 64, 65, 66, 68, 69, 82, 83, 87

CR-27 · "528 Hz repairs DNA"

CLAIM	That a tone at 528 Hz repairs, restores, or beneficially modifies DNA.
(+1) SOURCE	None known to this author. Stated with the uncertainty on the record: no systematic survey is claimed; what is asserted is that I know of no replicated, peer-reviewed result establishing the effect, and the claim is treated accordingly.
MEASURED	Nothing that survives replication, to this author's knowledge. Four independent computed objections follow. (a) Spatial scale: wavelength in soft tissue `1540 / 528 = 2.9167 m`; ratio to a nucleus `4.9e5`, to a base pair `8.6e9`. The field is spatially uniform across a nucleus and cannot select a sequence. (b) Quantum energy: `h f / kT = 8.17e-11` at 310 K — the acoustic quantum is about one hundred-billionth of ambient thermal energy. (c) Frequency scale: molecular vibrational modes sit at `1.5e13` to `9.0e13 Hz`, which is 34.7 to 37.3 octaves above 528 Hz; an audible tone is not a molecular resonance. (d) Reference pitch: 528 Hz is a specific note only under a twentieth-century tuning convention, which the claim does not specify.

(-1) SOURCE	**VOID.** The “Solfeggio frequencies” are presented as ancient and are not. Guido of Arezzo’s eleventh-century solmization is real and assigns no frequencies to anything — it is a system of relative intervals taught in a period with no absolute pitch standard. The numerical values were produced in the 1990s by numerological reduction of a biblical passage and popularised in Horowitz and Puleo’s *Healing Codes for the Biological Apocalypse* (1999).
RECORDED	Nothing traditional. This is failure mode FM-61.6, the Costume: a recent construction presented as traditional so that the `–1` column appears populated when it is empty.
ROLES	(+1): empty. (-1): void — 0 bits, not few.
C1	**FAIL** — metaphor only; no shared structure survives the scale check.
C2	**FAIL** — generates no prediction that survives objection (a).
C3	**FAIL** — a practice built on a false channel is not utility.
0p	Does not integrate.
STATUS	**REJECTED.** And the rejection binds the framework: the sonification pipeline (CCP-031) may not claim what this entry refuses. A transposed tone is not a molecular resonance; the transposition is a stated convention carrying no mechanism; the essence frequencies are essence frequencies and not physical resonances (§3.4); and nothing in this series claims sequence-specific acoustic action on genomic material.
USED IN	Ch 61, 68, 69, 87

29.16.6 PART F — THE LIBRARY AND ITS ADDRESSING

CR-28 · DNA is a usable digital storage medium

CLAIM	Arbitrary digital data can be encoded in, stored as, and recovered from synthesised DNA.
(+1) SOURCE	Church, Gao & Kosuri, *Science*, 2012.
MEASURED	Encoding and recovery of a 5.27-megabit volume (0.659 MB) in synthesised DNA. The theoretical ceiling is `log2(4) = 2` bits per base; practical schemes run below it to avoid homopolymer runs and secondary structure. At 1 bit per base the encoded volume needs `5.27e6` bases, which is `0.170%` of a haploid human genome; the haploid genome’s capacity is `0.775 GB` at the 2-bit ceiling.
(-1) SOURCE	Scribal and mnemonic traditions treating text as a substrate for arbitrary content: the Masoretic apparatus of letter- and word-counts for verifying transmission, and the Vedic *krama* and *jata* recitation patterns, which encode one text in permuted orders for error detection.
RECORDED	Error-detecting redundancy applied to a sequence substrate. The Vedic permutation schemes are, in modern terms, a check-code over an ordered sequence; the Masoretic counts are a checksum.
ROLES	(+1) evidence, strong. (-1) precedent, and unusually apt: the two traditions named are independent of each other, ~2 to 3 bits. It is precedent for sequence-with-checksum, not for molecular storage.
C1	**PASS** — homomorphism, and a close one: an ordered symbol sequence, an alphabet, a redundancy scheme, a recovery procedure.
C2	**PASS** — predicts that error rates in a sequence medium are manageable by redundancy at a computable cost in density.
C3	**PASS**

0p	Integrates cleanly, and the –1 column contributes real design intuition about check-coding.
STATUS	**ESTABLISHED**
USED IN	Ch 66, 67, 69

CR-29 · A guide sequence is an address, and Cas12a carries a native multicast destination list

CLAIM	Guide-RNA format is a routing header; Cas12a's self-processing of its own crRNA array is a genuine multiple-destination list in the biological system itself.
(+1) SOURCE	Zetsche et al., *Cell*, 2015 (Cpf1/Cas12a as a single RNA-guided endonuclease); Fonfara et al., *Nature*, 2016 (Cpf1 processes its own precursor CRISPR RNA).
MEASURED	That Cas12a cleaves its own pre-crRNA array into mature guides without a separate processing RNA or host RNase, so one transcribed array yields several distinct targeting guides. **This is a destination list encoded in one sequence and unpacked by the receiver.**
(-1) SOURCE	Correspondence-table traditions read as address spaces: multiply-keyed indexes in which one entry carries several independent retrieval keys, attested in the Western esoteric compilations and, independently, in Chinese and Indian classificatory schemes.
RECORDED	That an element can be reached by any of several keys, and that a list of keys can be carried as a unit.
ROLES	(+1) evidence, strong. (-1) precedent, moderate independence, ~2 bits — and precedent for the abstract structure of multiply-keyed addressing, nothing biological.
C1	**PASS** — homomorphism on the structure "one carried object unpacks into several addresses."
C2	**PASS** — predicts that array length trades against per-guide efficiency, which is measurable.
C3	**PASS** — it is the basis of the addressing scheme (CCP-030).
0p	Integrates. The channel-addressing correspondence stated in the addressing analysis (CCP-030) is a correspondence of addressing structure and carries no claim that anything is delivered.
STATUS	**ESTABLISHED** for the biology. The routing-header reading is a structural correspondence and is stated as such, not as a mechanism.
USED IN	Ch 66, 69

CR-30 · Targeting and editing are separable functions

CLAIM	A system can be directed to a specific genomic locus and act there without cutting or changing any sequence.
(+1) SOURCE	Qi et al., *Cell*, 2013 (catalytically dead Cas9 for sequence-specific control of gene expression); Gilbert et al., *Cell*, 2014 (genome-scale CRISPR-mediated repression and activation).
MEASURED	That a nuclease-dead Cas9 retains guide-directed binding while losing cleavage, and that fusing it to effector domains represses or activates transcription at the targeted locus. Bind a locus, change its expression, cut nothing.

(-1) SOURCE	None admissible. No traditional source could distinguish targeting from editing at a molecular level.
RECORDED	Nothing.
ROLES	(+1) evidence, strong. (-1): none, 0 bits.
C1	**N/A**
C2	**PASS** — predicts that regulatory intervention has a ceiling set by what is already present at the locus, which is CR-34.
C3	**PASS** — it is the established proof underwriting the honesty clause that nothing in this architecture edits DNA.
0p	N/A.
STATUS	**ESTABLISHED**
USED IN	Ch 66, 67, 83

CR-31 · Latent capability can be activated by regulatory signal alone

CLAIM	A small set of regulatory factors can revert a differentiated cell to a pluripotent state without adding information to the genome.
(+1) SOURCE	Takahashi & Yamanaka, *Cell*, 2006. Nobel Prize in Physiology or Medicine, 2012, to John Gurdon and Shinya Yamanaka.
MEASURED	That four defined transcription factors reprogramme differentiated fibroblasts to induced pluripotent stem cells. Nothing is added to the genome; the capability was already present and was being suppressed.
(-1) SOURCE	The doctrine, arrived at independently in several contemplative and philosophical traditions, that practice *uncovers* rather than *installs* — Plato's *anamnesis*; the Zen formulation of original enlightenment; the sculptor's commonplace that the form is already in the stone.
RECORDED	That what a practice produces was already present and obstructed, and that the operation is removal rather than addition.
ROLES	(+1) evidence, prize-grade, ~20 bits. (-1) precedent, partially independent, ~2 bits. The asymmetry is twentyfold.
C1	**PASS** — homomorphism on the structure "the content is present; the operation changes what is expressed, not what is stored." This is a genuinely deep shared pattern and not a verbal one.
C2	**PASS** — the bridge generates the sharpest constraint in this series: a purely regulatory intervention has a ceiling given by the latent repertoire, and nothing outside the repertoire can be reached. That is CR-34 and it binds the framework's own protocol.
C3	**PASS** on both sides.
0p	Integrates without either side lying: the traditions need not have known about transcription factors, the biology need not endorse *anamnesis*, and both describe expression versus storage.
STATUS	**ESTABLISHED** for the structural correspondence, with column roles as stated. **SPECULATIVE** — marked here — for any claim that a contemplative practice acts on cellular reprogramming.
USED IN	Ch 67, 84, 85

CR-32 · The genome is mostly regulatory sequence, and "junk DNA" is retired

CLAIM	Non-coding sequence is the great majority of the genome and is largely regulatory or structural rather than inert.
(+1) SOURCE	Contemporary genomics; the framework's own statement of the open problem (CCP-031).
MEASURED	Approximate, estimate-dependent figures: non-coding ~98% of the human genome; transposable elements ~45%; endogenous retrovirus-derived ~8%. On `3.1e9` bp haploid that is `~3.04e9`, `1.40e9` and `2.48e8` bp. Categories include promoters, enhancers, silencers, insulators, lncRNA, miRNA. **Percentages vary by method and by what "functional" means; cited as approximate.**
(-1) SOURCE	Library and archive traditions in which the catalogue, apparatus and marginalia occupy more volume than the primary texts, and in which the apparatus is what makes the collection usable.
RECORDED	That an archive is mostly not its contents, and that the non-content part is what confers access.
ROLES	(+1) evidence, strong. (-1) precedent, ~1.6 bits, and it is an analogy about archives, not a claim about genomes.
C1	**PARTIAL** — analogy rather than homomorphism.
C2	**PARTIAL** — the archive reading generates no genomic prediction.
C3	**PASS** as exposition.
0p	Integrates as illustration only.
STATUS	**ESTABLISHED** for the genomic facts. **SPECULATIVE** — marked here — for the archive correspondence, which is an analogy used for exposition and carries no inferential weight.
USED IN	Ch 67

CR-33 · Syncytin: ancient foreign text, inherited, later read out as vital function

CLAIM	A gene essential to human placenta formation derives from a retroviral envelope gene.
(+1) SOURCE	Mi et al., *Nature*, 2000.
MEASURED	That syncytin, a protein involved in trophoblast fusion and placental morphogenesis, derives from the envelope gene of an endogenous retrovirus. Documented molecular biology, not interpretation.
(-1) SOURCE	None admissible for the biology. Where the entry is used illustratively alongside traditions of inherited text later found to carry unforeseen function, that use is expository and is marked.
RECORDED	Nothing bearing on the claim.
ROLES	(+1) evidence, strong. (-1): none, 0 bits.
C1	**N/A** for the biological claim.
C2	**PASS** — predicts that other endogenous retroviral sequences have been co-opted to function, which further work supports.
C3	**PASS** — it is the concrete instance behind "the library of the cell."
0p	N/A.
STATUS	**ESTABLISHED**
USED IN	Ch 67, 84

CR-34 · Addressing adds no information

CLAIM	A scheme that specifies *where* to look adds nothing to *what is there*; therefore the ceiling on any purely regulatory intervention is the latent repertoire.
(+1) SOURCE	CR-30 (targeting separable from editing) and CR-31 (activation without addition), jointly.
MEASURED	dCas9 binds without cutting; Yamanaka factors reprogramme without adding sequence. In both cases the intervention selects among what is present and adds nothing.
(-1) SOURCE	The memory-palace tradition (§6.2), and specifically its late Renaissance error: the belief that a palace mirroring the structure of the universe would confer understanding of the universe.
RECORDED	The technique, correctly; and the error, instructively.
ROLES	(+1) evidence, strong. (-1) precedent, ~1.6 bits, carrying a cautionary case rather than a supporting one.
C1	**PASS** — homomorphism.
C2	**PASS** — predicts a hard ceiling: nothing not already latent can be reached. This is falsifiable and it constrains the framework's own protocol more tightly than any external criticism would.
C3	**PASS** — it is the honesty clause of the addressing and latent-capability analyses (CCP-030, CCP-031) in one sentence.
0p	Integrates.
STATUS	**ESTABLISHED**
USED IN	Ch 62, 66, 67, 83, 84, 87

29.16.7 PART G — METHOD, NOTATION, AND THE OLD PROCEDURES

CR-35 · Cadence: periodic forcing of a leaky store

CLAIM	For a decaying store feeding a threshold, scheduling a fixed total input at short intervals raises the floor; intensity buys the peak, cadence buys the floor.
(+1) SOURCE	the storage law's contraction and DC-gain results (CCP-007); Cepeda, Pashler, Vul, Wixted & Rohrer, *Psychological Bulletin*, 2006 (quantitative synthesis of distributed practice); Ebbinghaus 1885 for the decay curve.
MEASURED	Spaced presentations outperform massed ones for the same total study, and the optimal gap grows with the retention interval — the model's own prediction that optimal `T` scales with `1/delta`. At `delta = 0.10`, `R = 1.0`: mean `Q = 10.000` for every interval, `Q_min` falling `10.000 -> 7.131 -> 1.476 -> 0.008` as `T` goes `1 -> 7 -> 30 -> 90`.
(-1) SOURCE	Fixed-cadence observances: the Rule of St Benedict (c. 530 CE, eight daily offices); Jewish thrice-daily prayer and the seven-day sabbath; Islamic *salat*; Vedic *sandhyavandanam*; Buddhist monastic daily and fortnightly cycles. Independence is low — the Abrahamic set shares a lineage and Near Eastern contact is continuous.
RECORDED	The schedule, the repair procedure for a lapsed office, and the doctrine that a short observance kept beats a long one abandoned. Also Rappaport (1999) on invariance as constitutive of ritual form.
ROLES	(+1) evidence, strong on both the theorem and the spacing effect. (-1) precedent, credited at approximately one tradition's weight, ~1.6 bits, because `c` is high.
C1	**PASS** — homomorphism; three of four structures map.

C2	**PASS** — optimal interval scales with the time constant; tested in the memory domain; holds.
C3	**PASS** on both sides.
0p	Integrates without either side lying. The framework does not claim liturgical hours were tuned for retention; that is unfalsifiable. It claims the schedule has the form the model recommends.
STATUS	**ESTABLISHED**
USED IN	Ch 2, 59, 61, 62, 85, 86

CR-36 · The method of loci is spatial storage with a favourable retention constant

CLAIM	Binding items to an ordered set of familiar places raises encoding efficiency and, more importantly, lowers the decay rate.
(+1) SOURCE	Maguire, Valentine, Wilding & Kapur, *Nature Neuroscience*, 2003; Dresler et al., *Neuron*, 2017.
MEASURED	2003: superior memorizers distinguished by spatial mnemonic strategy rather than general ability or brain structure, with engagement of spatial-memory regions. 2017: trained naive subjects showed large durable recall gains with connectivity changes. **Direction and qualitative finding asserted; effect sizes not quoted and not independently verified here.**
(-1) SOURCE	The classical rhetorical tradition; *Rhetorica ad Herennium* as the standard early source; transmission traced by Yates, *The Art of Memory*, 1966. Independence low — a single continuously copied Western lineage.
RECORDED	The full procedure: how many loci, how spaced, how lit, what makes an image memorable, how the walk is performed.
ROLES	(+1) evidence, and it carries the claim. (-1) precedent, ~1.6 bits, and it carries the procedure.
C1	**PASS** — homomorphism; the storage law's `eta`, `delta` and `C` all have counterparts in the tradition's own emphases.
C2	**PASS** — the advantage is a decay-rate advantage and compounds with retention interval. Illustrative: `eta 0.20 / delta 0.35` versus `eta 0.70 / delta 0.06` gives `3.5x` at day 0 and `612x` at day 14. **Parameters assumed, not measured; the compounding is the prediction.**
C3	**PASS** on both sides.
0p	Integrates completely.
STATUS	**ESTABLISHED**
USED IN	Ch 59, 62

CR-37 · A correspondence table is a category assignment

CLAIM	Correspondence systems are finite relations, and a correspondence table performs the mandatory precondition for applying the surplus functional.
(+1) SOURCE	The surplus functional's validity conditions (CCP-A02); the literal-reading rule of §6; the enumerations computed there.
MEASURED	Definitional and computational, not empirical. `N_eff = 1 + (N−1)u` cannot be evaluated until the `N` categories are named. Addressing: a flat 78-element set costs `log2(78) = 6.285` bits; the `22 + 4 x 14` decomposition costs `6.807` bits and is far cheaper to memorise, its components being independent.

(-1) SOURCE	Esoteric correspondence systems: the Kircher-arrangement Tree of Life (10 vertices, 22 edges); the 78-element tarot; planetary-metal-day tables. Dated and traced in the corpus's literal-reading audit.
RECORDED	A fixed, written, auditable category assignment with multiple retrieval keys per element.
ROLES	(+1): definitional, so the usual evidence/precedent framing does not apply — the claim is that the artefact *is* the object, which is checkable by inspection. (-1) precedent, ~1.6 bits, low independence.
C1	**PASS** — isomorphism, and this is one of very few entries in the register that reaches it.
C2	**PASS** — predicts that the systems have computable graph and set properties, which the corpus's literal-reading audit computed and printed.
C3	**PASS** — it supplies the category assignment the formalism requires.
0p	Integrates trivially, because the correspondence is definitional.
STATUS	**ESTABLISHED** as a reading. **REFUSED**, standing, is any claim that a correspondence carries efficacy, force, or physical consequence (standing refusal, §3.7).
USED IN	Ch 57, 62, 66

CR-38 · Gematria is a hash function

CLAIM	A letter-value scheme summed over a word is a map from strings to integers, which is what a hash function is.
(+1) SOURCE	Definition; the corpus's measurement of collision behaviour.
MEASURED	Computed by the corpus's literal-reading audit: on a 234,454-word list the scheme destroys 10.47 bits. Collisions are the point of a hash and are not evidence of significance.
(-1) SOURCE	Hebrew and Greek letter-numeral traditions, attested from antiquity; Arabic *abjad*.
RECORDED	The correspondence of letters to numbers and the practice of summing them.
ROLES	(+1): definitional and computational. (-1) precedent, ~1.6 bits.
C1	**PASS** — isomorphism.
C2	**PASS** — predicts a computable collision rate, which the corpus's literal-reading audit computed.
C3	**PASS** as a technical description.
0p	Integrates.
STATUS	**ESTABLISHED** as a description. **REFUSED** is any claim that a numerical coincidence between two words carries significance.
USED IN	Ch 57, 62

CR-39 · The essence law is a pure square law, and essence frequencies are not physical resonances

CLAIM	`f(n) = k n^2` with `k = 0.483425732988414 Hz`, and the resulting frequency is the resonance of the essence, not a physical resonance.
(+1) SOURCE	The author's *Effective Method of Sonic Chemistry* (2022); verified against every published value to 14 significant figures.

MEASURED	The collapse of the original recipe to a pure square law, and `k = (9/8)^2 / phi^2 = (81/64)/(phi+1)`, verified exactly. The verification is of internal consistency, not of any physical effect.
(-1) SOURCE	Pythagorean and later speculative-music traditions relating number to tone.
RECORDED	Ratio relations between tones, and the doctrine that numerical relation is the substance of harmony.
ROLES	(+1): internal verification only, not physical evidence. (-1) precedent, ~1.6 bits, and only for the ratio practice.
C1	**PASS** for ratio structure; **FAIL** for any claim of physical resonance.
C2	**PASS** — the odd-part theorem predicts exactly `ceil(N/2)` distinct pitch classes for an index set `1..N`, which is computable and was computed.
C3	**PASS** as a notation.
0p	Integrates only under the standing distinction.
STATUS	**POSTULATE** (the author's, badged). The distinction that essence frequencies are essence frequencies and not physical resonances is absolute and is restated wherever the law is used (§3.4). CR-27 is the reason the distinction is non-negotiable.
USED IN	Ch 29, 30, 32, 69

CR-40 · Staged process control: the ordering is the knowledge

CLAIM	The alchemical operations are unit operations, and the correct ordering of stages is itself substantial process knowledge.
(+1) SOURCE	Chemical engineering; the direct inheritance of distillation, sublimation, crystallisation and filtration, and of the apparatus — alembic, retort, *balneum mariae*.
MEASURED	Computed: `n` stages admit `n!` orderings, so specifying the correct one carries `log2(n!)` bits — `2.585` bits for 3 stages, `6.907` for 5, `12.299` for the seven-stage presentations, `28.835` for 12.
(-1) SOURCE	The alchemical corpus: calcination, dissolution, separation, conjunction, putrefaction, distillation, coagulation, in the seven-stage presentations, with longer and shorter lists elsewhere. Attested from the Hellenistic period through the early modern.
RECORDED	The sequence, the graded heats, the requirement of long digestion at low temperature, and the iterative diminishing-returns character of purification.
ROLES	(+1) evidence, decisive (the operations are still performed). (-1) precedent, and here it is unusually valuable: the corpus is the *record of the procedure*, which is exactly the payload §4.5 assigns the `–1` column.
C1	**PASS** — isomorphism for the operations themselves; they are the same operations.
C2	**PASS** — predicts that permuting stages degrades yield, which is testable and which is why the ordering carries information at all.
C3	**PASS**
0p	Integrates. And it illustrates survivorship bias B3 in its purest form: the part that was right is tacit skill and copies badly; the part that was wrong is text and copies well.
STATUS	**ESTABLISHED**
USED IN	Ch 62

CR-41 · Chemical transmutation of the elements

CLAIM	That base metal can be converted to gold by chemical operations.
(+1) SOURCE	Nuclear physics, applied.
MEASURED	Transmutation changes proton number and is a nuclear operation. Nuclear binding energy is of order `8 MeV` per nucleon; a covalent bond of order `4 eV`. **The ratio is `2,000,000` — six orders of magnitude (`log10 = 6.30`), and no furnace, reagent or duration closes it.** Rider: transmutation is real physics, performed in accelerators and reactors at costs making the product worthless as a commodity.
(-1) SOURCE	The alchemical corpus's compositional theory: the four elements and the three principles.
RECORDED	A compositional theory of matter, and the goal of transmutation. Envelope check: P4 fails — a failed transmutation is always attributable to incompleteness (T5), so failure was invisible and the filter was memorability-selecting.
ROLES	(+1) evidence, decisive against. (-1) precedent: 0 bits — this is precisely the profile of material §5.4 predicts will be preserved and be false.
C1	**FAIL** by scale (error E6, six orders of magnitude).
C2	**FAIL**
C3	**FAIL**
0p	Does not integrate.
STATUS	**REJECTED.** Retained in the register because it is the corpus's most instructive failure: the goal was not absurd, only unreachable by the means available, and the reason is now exactly computable.
USED IN	Ch 62

CR-42 · The doctrine of signatures

CLAIM	That the appearance of a plant indicates its medicinal use.
(+1) SOURCE	None; the doctrine has no measured support.
MEASURED	Nothing. The famous apparent hits — walnut and brain, lungwort and lung — are false positives produced by a rule that generates a resemblance for anything and is therefore never wrong in advance and never right by design.
(-1) SOURCE	Paracelsus; Giambattista della Porta, *Phytognomonica* (1588); the herbal tradition descending from them.
RECORDED	The doctrine, in detail, for centuries. Envelope check: P1 fails (small, slow effects); P2 fails (illness does not recur on demand); P3 passes (the resemblance is visible); P4 fails (delayed and confounded). Visible failure: none. Memorability: maximal.
ROLES	(+1): empty. (-1) precedent: 0 bits.
C1	**FAIL** — metaphor only.
C2	**FAIL** — the rule generates its own content and cannot be wrong in advance.
C3	**FAIL**
0p	Does not integrate.

STATUS	**REJECTED. The register's control: an item scoring maximum on memorability and minimum on every accuracy-selecting criterion, whose actual fate — preserved for centuries, and false — is exactly what the preservation model of §5.4 predicts.** Retained because it shows the model can be tested against a known answer.
USED IN	Ch 62, 87

29.16.8 PART H — THE WORKED ENTRY

The entry built from nothing in Section 7.6, in register format.

CR-43 · Religious fasting schedules and autophagic flux

CLAIM	Traditional fast durations fall within the window in which human autophagic flux rises materially.
(+1) SOURCE	Ohsumi's identification of the autophagy machinery; Nobel Prize in Physiology or Medicine, 2016.
MEASURED	The genes, the pathway, and its upregulation under nutrient deprivation in model systems. Not settled: the time course, magnitude and tissue distribution of fasting-induced autophagy in living humans under these schedules. This author asserts no settled quantitative answer.
(-1) SOURCE	Yom Kippur (~25 h, annual, 2000+ years); Ramadan (lunar month, dawn to sunset, 7th c. onward); Lent and Friday abstinence (4th c. onward); Ekadashi (fortnightly); Buddhist vinaya (no solid food after midday); North American vision-quest and Mesoamerican ritual fasting — the last two carrying the genuine independence.
RECORDED	Duration, timing, exclusions, manner of breaking, warning signs, and a reported change of state.
ROLES	(+1) evidence: strong for mechanism, weak-to-absent for human dose-response. (-1) precedent: partially independent, ~2 bits, attesting schedule and subjective state only.
C1	**PARTIAL** — weak homomorphism; 2 of 5 structures map (threshold, cycle).
C2	**PASS** — predicts a material rise in human autophagic flux within the 12-to-30 hour window and not substantially below it. Falsifiable; untested at this specificity.
C3	**PARTIAL** — a stated hypothesis on the evidence side; nothing new on the precedent side.
0p	Integrates narrowly: both describe a threshold-and-cycle structure on nutrient deprivation and agree on roughly where the threshold sits.
STATUS	**SPECULATIVE** — the bridge between religious fasting schedules and autophagic flux is speculative, and this is the sentence that marks it.
USED IN	Ch 61, 85

29.16.9 Register tallies, recomputed

Recomputed field by field from the entries above, and reported against the source's own counts in Section 11.5.

Quantity	Over the 42 entries CR-01 to CR-42	Over all 43 entries
ESTABLISHED (primary status)	32	32
SPECULATIVE (primary status)	2	3
REJECTED	6	6
POSTULATE (primary status)	2	2
ESTABLISHED with a speculative rider	4 (CR-04, CR-12, CR-31, CR-32)	4
ESTABLISHED with a REFUSED rider	2 (CR-37, CR-38)	2
ESTABLISHED with a POSTULATE rider	1 (CR-19)	1
Entries carrying a qualifier	14 of 42 = 33.3 %	15 of 43 = 34.9 %
C1 = PASS	21	21
C1 = PARTIAL	4	5
C1 = FAIL	7	7
C1 = N/A (no bridge proposed)	10	10
Bridges attempted (C1 graded)	32	33
C1 outright failure rate	**7/32 = 21.9 %**	**7/33 = 21.2 %**
Entries with an empty or void precedent column	12	12

The twelve entries with an empty or void precedent column are CR-06, CR-09, CR-10, CR-14, CR-16, CR-17, CR-21, CR-22, CR-26, CR-27, CR-30 and CR-33. Ten of them draw no bridge and carry `C1 = N/A`; the remaining two, CR-17 and CR-27, are graded `C1 = FAIL` and are therefore counted among the bridges attempted. **That is consistent, and it is the reason the twelve and the ten are different numbers.**

29.17 Appendix B Works named in the text without a bibliography key

The bibliographic database governing this corpus contains no entries for the following, and **no key has been attached to any of them anywhere in this paper.** They are listed with the details the source manuscript supplies so that a reader can locate them, and their absence is recorded as a gap in the corpus's apparatus rather than concealed (Section 3.6, limitation L-6). **Nothing in this list should be read as a citation; it is a statement of what is missing.**

On belief and its channels. Levine, Gordon and Fields (*The Lancet*, 1978) on naloxone reversal of placebo analgesia. Kaptchuk and colleagues (*PLoS ONE*, 2010) on open-label placebo in irritable bowel syndrome, and the subsequent open-label trials. Sisk, Burgoyne, Sun, Butler and Macnamara (*Psychological Science*, 2018) on the meta-analysis of growth-mindset interventions.

On mechanotransduction and the cell. Coste and colleagues (*Science*, 2010) on Piezo1 and Piezo2 as essential components of mechanically activated cation channels, and the 2021 Nobel Prize in Physiology or Medicine to Ardem Patapoutian, shared with David Julius. Tajik and colleagues (*Nature Materials*, 2016) on force transmission to the nucleus, chromatin stretching and transcriptional upregulation. Ibsen, Tong, Schutt, Esener and Chalasani (*Nature Communications*, 2015) on TRP-4 sonogenetics in *C. elegans*, and Duque and colleagues (2022) on hsTRPA1. The cochlear mechanotransduction and tip-link gating literature. Bernardi and

colleagues (*BMJ*, 2001) on respiration and baroreflex sensitivity during rosary and mantra recitation.

On the genome. Church, Gao and Kosuri (*Science*, 2012) on digital information storage in synthesised DNA. Zetsche and colleagues (*Cell*, 2015) on Cpf1/Cas12a, and Fonfara and colleagues (*Nature*, 2016) on its self-processing of pre-crRNA. Qi and colleagues (*Cell*, 2013) and Gilbert and colleagues (*Cell*, 2014) on catalytically dead Cas9 for sequence-specific transcriptional control. Takahashi and Yamanaka (*Cell*, 2006) on induced pluripotency, and the 2012 Nobel Prize in Physiology or Medicine to John Gurdon and Shinya Yamanaka. Mi and colleagues (*Nature*, 2000) on syncytin. Yoshinori Ohsumi's identification of the autophagy machinery and the 2016 Nobel Prize in Physiology or Medicine. The contemporary genomic-annotation literature behind the non-coding, transposable-element and endogenous retroviral fractions quoted in CR-32, whose figures are method-dependent and are quoted as approximate for that reason.

On memory and practice. Ebbinghaus (1885) on the retention decay curve. Cepeda, Pashler, Vul, Wixted and Rohrer (*Psychological Bulletin*, 2006), the quantitative synthesis of distributed practice. Maguire, Valentine, Wilding and Kapur (*Nature Neuroscience*, 2003) and Dresler and colleagues (*Neuron*, 2017) on the method of loci. Yates, *The Art of Memory* (1966). The *Rhetorica ad Herennium*.

On ritual, tradition and the old corpus. Rappaport, *Ritual and Religion in the Making of Humanity* (1999), on invariance as constitutive of ritual form. Giambattista della Porta, *Phytognomonica* (1588), and the Paracelsian doctrine of signatures. Horowitz and Puleo, *Healing Codes for the Biological Apocalypse* (1999), the source of the "Solfeggio" values rejected in Section 7.5. Guido of Arezzo's eleventh-century solmisation.

On physics named in the register. Heisenberg (1927) and the subsequent formalisation of the uncertainty relation beyond the entry already keyed here; the decoherence and environment-induced superselection literature; the von Neumann–Wigner interpretation; and the standard vibrational-spectroscopy range of 500 to 3000 cm^{-1} used in Section 7.5.

On the framework's own earlier work. The author's *Effective Method of Sonic Chemistry* (2022), the source of the essence law recorded in CR-39, and the earlier framework from which the three-state method and the three bridge criteria of Section 3.1 are adopted unchanged. **These are the author's own prior works and are named rather than keyed for the same reason: the database carries no entries for them.**

Chapter 30

The Cell as Receiver, the Channel and Its Capacity, and Addressing Without Editing: Adequate Bandwidth Into an Absent Decoder

Abstract. We compute the acoustic channel into a single cell and separate what the physics permits from what has been demonstrated. Air-to-tissue transmission is -29.9 dB; the pressure differential across a 10 um cell is 0.0816 per cent of peak at 20 kHz and 4.08 per cent at 1 MHz; and wavelength equals cell diameter only at 154 MHz, where attenuation confines the field to millimetres. The product of wavelength and 20 dB depth falls from 6160 to 0.0026 mm-cm across that range, closing acoustic addressing on physical grounds rather than on capacity. The binding term is the receiver's own bandwidth: mechanotransduction kinetics give a corner frequency near 10.6 Hz, so Shannon capacity is of order 70 to 106 bits per second and a 40-bit address takes about half a second. That result is used to eliminate an explanation rather than to support a claim — individually addressing the 3.65e6 cells in one voxel would take 23.93 days, a per-cell exclusive rate of 1.93e-5 bit/s, so the field is intrinsically broadcast. The Gabor limit dt*df >= 1/(4*pi) yields an optimal symbol rate R_max = 6.669448*B, cross-checked against Shannon at 20.034 dB. On the addressing side, naming one genomic position costs 32.53 bits, the PAM constraint costs 4.4465 nats of naming capacity, and a worked block is declared valid only at matcher tolerance d <= 2, where expected collisions are 0.624 against 11.47 at d = 3.

Keywords: mechanotransduction; acoustic impedance; channel capacity; receiver bandwidth; Gabor limit; genomic addressing; PAM constraint; collision analysis

Series: B. Engineering the Surplus **Paper:** CCP-030

30.1 1 Introduction

A living cell is a mechanical receiver. That is established physiology, it is not speculative, and it carries a Nobel Prize. It is also the exact size of the claim, and this paper is built so that the difference between *receiving force* and *receiving a message* is structural rather than rhetorical.

The question this paper answers is the engineer's question rather than the enthusiast's: **given that a cell transduces pressure, how much information can physically arrive at it, and**

what does the arithmetic then permit anyone to claim? The answer is computed in full, with every constant declared nominal at first use, and it is a surprise in both directions.

The channel is not the bottleneck. Air-to-tissue intensity transmission at normal incidence is `-29.9` dB, while the pressure amplitude just inside the tissue nearly doubles; the binding term downstream is neither of those but the receiver's own bandwidth. Taking mechanotransduction kinetics with a nominal 15 ms time constant, the corner frequency is 10.61 Hz, and Shannon capacity at 20 dB signal-to-noise is **70.65 bit/s** — rising to 106.0 bit/s at a 10 ms constant. **A 40-bit address, the size of a 20-nucleotide name, takes about half a second to deliver.** That result is used to *eliminate* an explanation rather than to support a claim: bandwidth is adequate, so "there is not enough capacity" is removed from the list of available explanations, and what remains is exposed.

The spatial resolution is catastrophic, and it fails on geometry rather than on bandwidth. In soft tissue at a nominal 1540 m/s the wavelength equals a 10 µm cell diameter only at **154 MHz**, where the depth at which one-way attenuation reaches 20 dB is 2.6 mm on the most generous attenuation model. The product of wavelength and 20 dB depth falls as `1/f^2`, from **6160 mm·cm at 100 kHz to 0.0026 mm·cm at 154 MHz**. At 1 MHz, taking a resolution voxel as `lambda^3` and a nominal density of `1e9` cells per cm^3, **3.65e6 cells share one voxel and every one of them receives the identical waveform.** Individually addressing them at 70.65 bit/s would take **23.93 days**, a per-cell exclusive rate of `1.93e-5` bit/s. **The acoustic field into tissue is intrinsically a broadcast**, and that is a property of the physics rather than of any particular design.

Energy is not the constraint either. Landauer's bound at 310 K permits `8.19e3` bit/s into a 10 µm cell at one pascal, more than a hundred times the receiver-limited capacity. Three intuitively appealing wrong explanations are therefore removed — not bandwidth, not energy, not thermodynamics — and the two that remain are spatial resolution, which is physical and severe, and **the absence of a demodulator, which is the open problem this paper hands to its successor** [Curzi, 2026y].

A second computed result comes out of putting two bounds side by side. A Gabor-limited tone-selection channel carries at most `R_max = 4 pi B log2(e)/e = 6.669448 B` bit/s, and Shannon's `C = B log2(1 + S/N)` equals that when `S/N = 2^6.669448 - 1 = 100.79`, which is **20.034 dB**. Below that ratio noise binds; above it, time–frequency resolution binds.

The third part of the paper takes up the problem the second one locates. If the acoustic field cannot resolve a cell, selectivity must come from somewhere with cellular resolution, and there is exactly one such layer: sequence. A 20-nucleotide name is a 40-bit address; naming one position in a mammalian genome on either strand requires `log2(2 x 3.1e9) = 32.5296` bits. **Targeting and editing are demonstrated to be separable functions**, so a system may name a locus with full specificity and do nothing to it. The protospacer-adjacent-motif requirement acts as an address-family constraint and costs **4.4465 nats — 6.415 bits — of nameable variety** for the T-rich family relative to unrestricted naming. And the decisive design number is not the address length but the matcher tolerance: a 20-mer address has an expected **0.624** genomic collisions at Hamming distance two and **11.47** at distance three, so a worked address block is declared valid only at `d <= 2`.

The counter-result is the paper's spine, and it is stated here rather than at the end. Every favourable number above is refused as support for the architecture it appears to serve. **Computing the capacity of a channel leading to a receiver does not create the receiver.** The standing constraint of this series is restated in every section that touches the subject: *there is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal.* Measuring that a room has excellent acoustics tells you nothing about whether anyone in it speaks the language.

Sections 2 and 3 give the hypotheses and the method, including the declared nominal constants, the analogy discipline that governs every borrowing, and an explicit disclosure of a gap in the corpus's bibliography that leaves the entire biological evidence base of this paper uncited. Section 4 characterises the cell as a receiver. Section 5 computes the channel. Section 6 specifies the addressing layer. Sections 7 and 8 collect the specifications and the named

failure modes. Section 9 works three complete specifications end to end, two of which fail. Sections 10 to 13 discuss, verify, limit and conclude.

30.2 2 Hypotheses

Five hypotheses, each stated so that it could be wrong, each with its falsifier. Two further claims are labelled as what they are rather than dressed as hypotheses.

H1 — The receiver, not the medium, sets the information rate into a cell. The bandwidth term in Shannon's formula is the corner frequency of the slowest stage that must resolve a symbol, not the carrier bandwidth. **Predicted:** substituting carrier bandwidth for receiver bandwidth inflates the computed capacity by three to five orders of magnitude, and the receiver-limited figure at mechanotransduction kinetics is of order tens of bits per second. **Falsified by:** an identified cellular stage that resolves symbols at carrier rates, or a demonstration that mechanotransduction kinetics do not bound the symbol rate. *(Result: **not falsified** — 70.65 bit/s at 10.61 Hz and 20 dB, against 6,658 bit/s at a 1 kHz medium bandwidth and `6.66e6` bit/s if the full 1 MHz carrier is substituted: inflation factors of 94.2 and `9.4e4`; see Section 5.7 and the correction in Section 9.4.)*

H2 — Acoustic addressing of individual cells fails on geometry, not on capacity. The frequency at which an acoustic wavelength resolves a cell is a frequency at which the wave cannot reach one. **Predicted:** the product of wavelength and 20 dB penetration depth falls as `1/f^2`, and every plausible attenuation exponent makes the conclusion worse rather than better. **Falsified by:** an attenuation model on which resolution and depth do not trade, or a non-acoustic mechanism supplying cell-scale acoustic resolution. *(Result: **not falsified** — 6160 to 0.0026 mm·cm between 100 kHz and 154 MHz, and exponents 1.1 and 1.5 both reduce the reachable depth further; see Section 5.4.)*

H3 — Capacity is adequate, and its adequacy eliminates an explanation rather than supporting a claim. If bandwidth were the obstacle, the architecture would fail for a reason that could be engineered around. **Predicted:** the computed capacity delivers a 40-bit address in under a second, and the Landauer bound exceeds the receiver-limited capacity by more than two orders of magnitude. **Falsified by:** a capacity computation returning a figure that cannot carry an address in a physiologically relevant time. *(Result: **not falsified** — 0.566 s for 40 bits at 20 dB; Landauer permits `8.19e3` bit/s at 1 Pa; see Sections 5.7 and 5.10. **The consequence is that the explanation for the architecture's failure must be sought elsewhere, and Section 6.9 states where.**)*

H4 — Targeting and editing are separable, so an addressing layer can be fully specific and completely inert. Specificity of address and modification of target are independent functions. **Predicted:** a catalytically inactivated nuclease retains full guide-directed binding specificity while performing no cleavage, and an architecture that introduces no protein and no nucleic acid cannot edit a genome. **Falsified by:** a demonstration that binding specificity requires catalytic activity. *(Result: **not falsified** — the separability is established in the cited literature; see Section 6.7. **The literature concerned carries no citation key in this paper, for the reason given in Section 3.7.**)*

H5 — Usable specificity is set by the matcher's tolerance and not by the address's length. Adding bits to a name is cheap; tightening the matching step is the decisive and expensive move. **Predicted:** expected genomic collisions for a 20-mer rise from 0.000352 at exact match to 11.47 at Hamming distance three, a factor of `3.3e4`,

while the address length is unchanged. **Falsified by:** a collision analysis in which tolerance is second-order to length. *(Result: **not falsified** — see Section 6.5, and the specification decision at `d <= 2` in Section 9.3.)*

Two claims here are not hypotheses.

That the leaky-integrator form describes cellular regulatory state is a modelling choice, badged wherever it appears (Section 4.7). It is a claim about the *shape* of an accumulation, not a claim that any cell obeys the storage law [Curzi, 2026a, Elaydi, 2005], and nothing in this paper tests it.

That there is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal is a statement about the state of the literature, not a theorem. It is a negative claim, it is stated with its scope, and Section 12 records exactly what would overturn it.

30.3 3 Method

30.3.1 3.1 What is established

Three statements used throughout this paper are established science. They are not postulates of this framework, they are not derived from anything in this corpus, and nothing about the framework is required for them to be true.

```
E1  MECHANICALLY GATED ION CHANNELS EXIST AND ARE IDENTIFIED.
    PIEZO1 and PIEZO2 are ion channels opened by mechanical force
    (Coste et al., Science 330:55–60, 2010). The 2021 Nobel Prize in
    Physiology or Medicine was awarded to David Julius and Ardem
    Patapoutian for discoveries of receptors for TEMPERATURE AND TOUCH.

E2  SOUND IS MECHANICAL.
    A sound wave is a propagating pressure and displacement disturbance
    in a material medium. Classical continuum mechanics.

E3  FORCE AT THE CELL SURFACE REACHES THE NUCLEUS AND CHANGES
    TRANSCRIPTION. Force applied at the cell surface transmits through
    the cytoskeleton, directly stretches chromatin, and upregulates
    transcription of a targeted reporter gene (Tajik et al., Nature
    Materials 15:1287–1296, 2016).
```

From E1 and E2 together: **a pressure wave incident on a cell can open mechanically gated channels, so a real, established pathway runs from acoustic energy to cellular signalling.** From E3 additionally: **mechanical input at the cell surface is an established input to gene expression.** Those two sentences are the whole of what this paper takes as given about biology. **E1 and E3 carry no citation key here** — see Section 3.7.

30.3.2 3.2 The standing constraint, stated first

The boundary is stated at the front of every section that touches the subject rather than buried in a caveat at the back:

> **There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal.**

Read the qualifier carefully, because it does precise work. **The constraint does not say cells cannot hear.** They can, in the only sense that matters mechanically: they respond to pressure.

It does not say sound cannot change gene expression. It can, by E3 and by the general route of mechanotransduction into signalling cascades. It says something narrower and harder: **nothing in the literature identifies a molecular apparatus that takes an acoustic waveform and recovers from it a specified nucleotide sequence, or any other symbolic content whose meaning was chosen by a sender.**

That is a statement about a *demodulator*. Every telecommunications system has one: a radio has a mixer, a filter and a detector; a modem has a constellation decoder; the cochlea has a bank of mechanically tuned filters feeding a rate code. Each recovers a symbol stream from a waveform because a physical structure exists whose job is that recovery. **In the cell, the receiver front-end is documented and the demodulator is not.**

The reason to state it in Section 3 rather than only in the companion paper is that Section 4 is where the temptation is strongest. The mechanotransduction literature is genuinely striking, the prize is genuinely there, and a reader who has just been shown that force opens channels and that force stretches chromatin is one careless paragraph away from believing something the evidence does not support. **The slide is one sentence long, and the standing of the whole architecture depends on refusing to make it.**

30.3.3 3.3 Nominal constants, declared once

Every constant below is **nominal**: a conventional round figure used to set scale, not a measured constant of any particular tissue or preparation. Each is flagged nominal again at first use in the section that consumes it.

Quantity	Value used	Status
Speed of sound, soft tissue	`c = 1540` m/s	**nominal** [Rossing et al., 2002]
Density, soft tissue	`rho = 1050` kg/m³	**nominal**
Speed of sound, air	343 m/s	**nominal**
Density, air	1.204 kg/m³	**nominal**
Acoustic attenuation	`alpha = 0.5` dB/cm/MHz	**nominal**, and a rough average; published soft-tissue values span roughly 0.3 to 1.0
Frequency exponent of attenuation	`b = 1.0`	**nominal simplification**; real tissues run `f^1.0` to `f^1.5`, and the sensitivity is computed in Section 5.4 rather than assumed away
Cell diameter	`d = 10` µm	**nominal**
Nucleus diameter	6 µm	**nominal**
Cell density in soft tissue	`1e9` per cm³	**nominal, order of magnitude only**
Body temperature	`T = 310` K, `kT = 4.280e-21` J = `4.280` pN·nm	**standard**
Mean residue mass	110 Da	**nominal**, used for mass estimates from residue counts
Propagation model	plane wave, linear, no cavitation	**modelling choice**

The attenuation figure deserves particular care because everything downstream leans on it. `0.5` dB/cm/MHz is the conventional round number for soft tissue; bone and lung are far higher. Section 5.4 computes the sensitivity of the conclusions to the exponent explicitly, and every plausible exponent makes the conclusion worse rather than better, which is the direction that matters for a bound.

30.3.4 3.4 Analogy discipline: what licenses each borrowing

Every borrowing from physics, information theory or networking in this paper is labelled here, at first use, as exactly one of three things, and no later section may let a category-2 or category-3 borrowing collect the authority of a category-1 one.

Borrowing	Category	What licenses it
Acoustic impedance, wavelength, attenuation, radiation force (§4.6, §5.2–5.4)	**(1) derivation**	These are claims about a physical acoustic field in physical tissue, and continuum acoustics applies directly [Rossing et al., 2002, Beranek, 2004]. The premises are satisfied; the results are used as physical results.
Shannon capacity `C = B log2(1+S/N)` (§5.7)	**(1) derivation**	The theorem applies to any channel for which a bandwidth and a signal-to-noise ratio can be defined [Shannon, 1948, Shannon and Weaver, 1949, Cover and Thomas, 2006]. **What it licenses is an upper bound and nothing else**: it is an existence proof about codes, not a construction of a decoder, and Section 5.7 is written around that distinction.
The Gabor limit `dt df >= 1/(4 pi)` (§5.8, §5.9)	**(1) derivation, with an explicit disclaimer**	Classical Fourier duality [Gabor, 1946], applying to every wave including sound and carrying **no quantum content whatever**. The position–momentum relation [Kennard, 1927] is a fact about state preparation rather than measurement disturbance and is *not* invoked here; the corpus proves the acoustic bound by the same variational argument with `p = hbar k` deleted, Gaussians saturating and a Hann window costing 2.6 per cent [Curzi, 2026z].
The thermal floor and fluctuation–dissipation (§5.5)	**(1) derivation, with a declared validity regime**	The fluctuation–dissipation theorem [Callen and Welton, 1951] fixes the noise of any dissipative receiver at temperature `T`. **Its application at cellular apertures is outside the regime in which the aperture formula holds**, and Section 5.5 says so and carries two estimates rather than one.
Landauer's bound (§5.10)	**(1) derivation**	`kT ln 2` per irreversible bit operation [Landauer, 1961]; used only to show that energy is not the binding constraint. **This is thermodynamic entropy; the bits of Sections 5.7 and 6.2 are Shannon information.** The two are different objects with different units, and the distinction is stated wherever both appear.
The leaky integrator applied to cellular regulatory state (§4.7)	**(2) shared mathematical form, badged**	`Q_{t+1} = (1−delta)Q_t + eta S_t − C_t` is a first-order linear difference equation whose fixed point and contraction are theorems about the equation [Curzi, 2026a, Elaydi, 2005, Banach, 1922]. **No claim is made that a cell obeys the storage law**; the claim is that the same accumulation shape is a useful description, and it is badged at every use.

Borrowing	Category	What licenses it
The internet-addressing correspondence (§6.3–6.6)	**(2) shared structural form, scoped in the paragraph that uses it**	The correspondence is of **format and function** — a variable destination field, a validity constraint on address form, a framing delimiter, an explicit destination list — and it is stated field by field so a reader can check it. **It is not a claim that cells run a protocol stack**, and any extension beyond format is marked speculative in the sentence that extends it. The general design principles it draws on are engineering ones [Saltzer and Schroeder, 1975].
Requisite variety applied to the address family (§6.4)	**(2) shared mathematical form**	Only variety absorbs variety [Ashby, 1956, 1958]; an address family is a restriction on the variety of acceptable addresses, so restricting it reduces what can be named. The arithmetic is the same; no cybernetic claim about cells is made.
The framework's constant `k = (9/8)^2/phi^2 = 0.483425732988414` (§5.12)	**definition, with a postulate attached**	`k` is an exact ratio of exact quantities and is **defined**, so its precision is legitimate. **Attaching hertz to it is a separate and much stronger claim and is POSTULATED**, not derived; the essence frequencies are essence frequencies and not physical resonances. The standard critique of golden-section claims [Markowsky, 1992] is cited here deliberately and again wherever `phi` appears, and the underlying identity `phi^2 = phi + 1` is Euclidean [Euclid, 1956, Vajda, 1989].
The cochlea as an existence proof (§4.3)	**(3) analogy, tightly bounded**	The cochlea proves that biology *can* build a fine mechanical receiver. **It does not license any inference about ordinary cells in ordinary tissue**, and Section 4.3 enumerates the special conditions it enjoys and states that a cell in bulk tissue has none of them.

Most of the borrowings in this corpus are category 2, and category 2 is respectable. The failure mode this table exists to prevent is a category-2 or category-3 borrowing quietly collecting the certainty of a category-1 one. **The standing critique of physics imported into non-physical subjects is that formal machinery arrives without the conditions that license it** [Cartwright, 1999, Anderson, 1972]; **this table is the answer to that critique, run one borrowing at a time, and the badged postulates it produces are doing exactly the work Lakatos [1970] describes.**

30.3.5 3.5 The two-column apparatus, and what each column carries

The corpus runs a concordance apparatus on every substantive claim: an evidence column drawn from the peer-reviewed literature and a precedent column drawn from pre-modern textual traditions. **The refinement that governs this series must be restated at every use: the evidence column carries EVIDENCE and the precedent column carries PRECEDENT, and two sources agreeing does not double the evidence when only one of them is evidence** [Curzi, 2026w]. The same paper prices the asymmetry — a single traditional record is worth 1.585 bits against 9.966 for a replicated result — and computes the collapse of agreement between correlated sources.

That paper's rejection of the claim that a tone at 528 Hz repairs DNA binds this one directly. **This paper's own results derive from that rejection**: the wavelength of 2.9167 m against a 10 µm cell, the quantum-to-thermal ratio of `8.17e-11` at 310 K, and the 34.7-to-37.3-octave separa-

tion between the audible band and molecular vibrational modes [Curzi, 2026w] are the same three computations that Section 5.3 generalises across frequency. **Nothing in this paper may claim sequence-specific acoustic action on genomic material, because the framework refused that claim to others.**

Where a precedent item appears in this paper it is stated with its tradition, approximate date and transmission path, and it is credited in bits rather than summed with the evidence column. **Where the precedent column is empty, the entry says so rather than filling it.**

30.3.6 3.6 Identification, reflexivity, and the limits this paper inherits

Nothing in this paper estimates a social parameter, so the reflection problem does not bind its arithmetic. **It binds the framework the arithmetic is embedded in, and that is stated here rather than left to a reader.** The corpus's central formal object is a reflexive loop in which a group influences its members and the members constitute the group; that is a linear-in-means reflexive specification, and Manski [1993] proves the endogenous social effect is **not identified from observational data**, with Shalizi and Thomas [2011] establishing the same confounding for network data. **No remedy is adopted anywhere in this series** — not randomisation, not an instrument, not partial-population variation, and not a defended stipulation of the reference group [Brock and Durlauf, 2001] — **so every claim in this corpus that rests on the loop is a modelling position rather than an identified empirical result.**

The specific inheritance here is Section 4.7's storage-law reading of cellular regulatory state. **It is a modelling choice about form, it is badged as one, and it is not identified from anything measured in this paper.** Lucas [1976] applies with equal force to the framework's own parameters: a framework asserting that systems change under observation, and then treating its own `delta` and `eta` as constants, must say which. The position taken is narrow: **`delta` and `eta` appear here only in illustrative parameterisations with declared values, never as estimates**, so there is nothing to be biased; any later use that fits them to data inherits the critique and this paper does not protect it.

30.3.7 3.7 A gap in the bibliography, disclosed

The biological literature this paper describes — mechanically gated ion channels and their structure, cochlear mechanotransduction and its tip-link apparatus, force transmission to the nucleus and its transcriptional consequence, mechanosensitive transcriptional regulators, substrate-stiffness effects on lineage, sonogenetics, nucleotide data storage, guide-sequence specificity, self-processing effector arrays, protospacer-adjacent-motif dependence, catalytically dead nucleases and their use for transcriptional control, compact effectors, and mitochondrial base editing — has no corresponding cluster in the bibliographic database governing this corpus. Every such work is therefore named in the text with its journal, volume, page and year, and **without a citation key**. Attaching an unrelated key to a biological claim would be worse than the gap, because a reviewer in that field will check it and a mis-supporting citation destroys more than a missing one.

Appendix A lists the affected works. Section 12 records it as limitation L-1. **This is a defect of the corpus's bibliography rather than a claim that the works do not exist**, and it means that **the entire biological evidence base of this paper is unverifiable from the reference list**: a reader must go to the named sources directly. The same gap is disclosed in the neighbouring papers of this series [Curzi, 2026w,y] and in the sonification protocol [Curzi, 2026].

One consequence of the gap must be stated even though the gap itself is the defect. Each biological result named here is named as a single published finding, and the required context for any such citation is the measured replication base rate: large-scale replication returns roughly a third of original findings with effect sizes about halved [Open Science Collaboration, 2015], and most published findings in low-prior fields are false [Ioannidis, 2005]. **The results relied on here are among the most heavily replicated in their field — mechanotransduction carries a Nobel Prize and the addressing results are routine laboratory tools — but**

the reader is entitled to that base rate beside them, and the corpus's own evidential pricing puts a replicated result at 9.966 bits against 1.585 for an unreplicated traditional record [Curzi, 2026w]. **This paper does not weight the individual biological findings; it takes them as given and computes physics around them.**

One further item is named without a key and is a correction rather than a gap. The source manuscript attributes the thermal-noise floor to "Nyquist (1928)". The corpus's bibliography contains a Nyquist 1928 entry, but it is *Certain Topics in Telegraph Transmission Theory* — the sampling paper — and **not** the thermal-agitation paper that establishes Johnson–Nyquist noise. **The two are different papers published in the same year and attaching the available key here would be a mis-citation.** The thermal floor is therefore attributed in this paper to the fluctuation–dissipation theorem [Callen and Welton, 1951], which is the correct and available warrant, and Nyquist's thermal-agitation paper is named in Appendix A without a key.

30.3.8 3.8 Standing honesty clauses and refusals

- **Capacity is permission, not performance.** Every capacity figure in this paper is accompanied, in the same paragraph, by the statement that it is an upper bound on what a perfect decoder could recover and does not imply that any decoder exists.
- **Bandwidth is the receiver's, never the transmitter's.** Quoting carrier bandwidth as channel bandwidth inflates capacity by three to five orders of magnitude and is named as failure mode FM-C.2.
- **Nothing is delivered, therefore nothing is edited.** The architecture described in Section 6 contains no nuclease, introduces no protein and introduces no nucleic acid into any cell. **It cannot edit a genome. Cannot, not does not** — the safety property is structural, and a structural safety property is the only kind worth trusting, because it does not depend on anyone behaving well.
- **No clinical, medical or therapeutic claim is made anywhere in this paper**, and the pressures tabulated for capacity purposes are not recommendations, exposure guidance, or safety statements of any kind.
- **Uncertain figures are omitted, not guessed.** Where this author is not certain that a published source reports a specific number, the number is left out and the omission is stated. Section 4.5 does exactly that with a drive frequency, and the text is stronger for the omission than it would be for a guess.
- **Numerical coincidences are labelled or omitted.** `4^18 = 2^36`, so an 18-nucleotide address is exactly 36 bits, the same number as an octave transposition used elsewhere in the corpus. **That carries no content**: both systems are built on powers of two, so coincidences of this shape are guaranteed and informative about nothing.
- **Scale type.** Every quantity in this paper is a physical measurement on a ratio scale with a true zero — pressure, frequency, time, energy, bits — so the arithmetic performed on them is licensed. **That is not true of the psychological constructs elsewhere in the corpus**, where Michell [1997]'s objection applies and the conjoint conditions [Luce and Tukey, 1964, Krantz et al., 1971] have not been tested; the storage-law reading of Section 4.7 inherits that exposure and is badged accordingly.

30.4 4 The cell as a mechanical receiver

30.4.1 4.1 The receiver chain, stage by stage

An engineer handed an unknown receiver asks for the block diagram before asking what it can decode. The cell's, drawn as one would draw a radio front-end:

```
ACOUSTIC FIELD
    |
```

```
    |  (a) coupling: pressure, pressure GRADIENT, shear, radiation force,
    |      streaming, cavitation if present
    v
PLASMA MEMBRANE
    |
    |  (b) bilayer tension; local curvature; area strain
    v
MECHANICALLY GATED CHANNEL   <-- PIEZO1/2, TREK/TRAAK, MscL/MscS (bacterial),
    |                            TMC-containing MET complex (hair cells),
    |  (c) ion flux               TRP family members
    v
MEMBRANE POTENTIAL + Ca2+ TRANSIENT
    |
    |  (d) second messengers, kinase cascades
    v
CYTOSKELETON / FOCAL ADHESIONS / LINC COMPLEX
    |
    |  (e) direct force transmission, tension-dependent unfolding
    v
NUCLEAR ENVELOPE -> CHROMATIN
    |
    |  (f) chromatin stretch; transcription factor localisation
    v
TRANSCRIPTIONAL RESPONSE
```

Two features of this diagram set up the whole of Section 5.

First, there are two parallel paths from surface to nucleus, and they run at very different speeds. Path (c)–(d) is electrochemical: channel opening in microseconds to milliseconds, calcium transients in hundreds of milliseconds to seconds, transcriptional consequences in minutes. Path (e) is mechanical, propagating along a prestressed cytoskeletal network, and is argued to transmit faster and further than diffusive signalling (Wang, Tytell and Ingber, *Nature Reviews Molecular Cell Biology* 10:75, 2009; named without a key, Appendix A). **A receiver with two paths of very different latency has a complicated impulse response, and Section 5.6 shows that the slowest stage in the chain — not the acoustic medium — caps the information rate.**

Second, nowhere in this diagram is there a stage that assigns meaning to a waveform. Each stage is a transducer or an amplifier. There is no correlator, no matched filter tuned to a symbol alphabet, and no decision device. **That absence is the standing constraint of Section 3.2, drawn in block-diagram form.**

30.4.2 4.2 The identified transducer

Mechanically activated currents had been recorded in many cell types for decades before the responsible molecule was known in vertebrates. Coste and colleagues, using a mechanically stimulated cell line and a candidate-knockdown screen, identified Piezo1 and, by homology, Piezo2 as essential components of distinct mechanically activated cation channels (*Science* 330:55–60, 2010); the 2021 Nobel Prize in Physiology or Medicine was awarded jointly for discoveries of receptors for temperature and touch. **This is not a fringe result pressed into service by a framework with an agenda; it is a canonical result of modern physiology.** It is also, per Section 3.7, cited here without a key.

The architecture explains the transduction mechanism. Cryo-electron microscopy shows PIEZO1 as a trimer with three long curved blade-like arms radiating from a central pore module; the blades bow the surrounding bilayer into a dome, and the membrane-dome model proposes that flattening the dome under membrane tension is what opens the pore — **the channel is gated by the area it can gain by straightening.**

Sizes, computed from residue count at a nominal mean residue mass of 110 Da. **This is an estimate from residue count, not a measured mass:**

```
mouse Piezo1                2521 residues
monomer, estimated          ~277 kDa
trimer, estimated           ~832 kDa  (~0.83 MDa)
```

Published cryo-EM work reports the trimer at roughly 0.9 MDa; **the estimate above should be read as an order check, not as a measurement.**

The gating physics matters more than the mass. If the channel is opened by bilayer tension `T` acting through an in-plane area change `dA`, the work available is `E = T dA`. Computed at body temperature with `kT = 4.280e−21` J `= 4.280` pN·nm at 310 K:

membrane tension `T` (mN/m)	`dA` for 1 kT (nm²)	`dA` for 5 kT (nm²)	energy at `dA` = 100 nm² (kT)
0.5	8.560	42.800	11.68
1.0	4.280	21.400	23.36
3.0	1.427	7.133	70.09
10.0	0.428	2.140	233.64

Read the right-hand column and the point is made: a channel with a large in-plane area change is an enormously sensitive tension sensor, because the energy scale it must beat is only a few kT. A conformational change of a hundred square nanometres at one millinewton per metre delivers of order twenty kT — decisive against thermal fluctuation. **The tension values are illustrative parameters chosen to span the physiologically discussed range; they are not measured constants of the protein, and the table is an energy-scale calculation rather than a characterisation.**

Excised-patch and reconstitution studies indicate the channel is intrinsically gated by bilayer tension rather than requiring an intact cytoskeletal tether (Cox et al., *Nature Communications*, 2016; Syeda et al., *Cell Reports*, 2016; Appendix A). **This matters for the acoustic question, because it means a stimulus that changes membrane tension can in principle act on the channel without a specialised accessory structure — and it matters for Section 4.6, because changing membrane tension is precisely what an ordinary sound field at cellular scale is bad at.**

Precedent, priced. The idea that the living body is a tuned structure rather than an inert one is recorded in the Latin curriculum for a thousand years before any instrument could test it: Boethius's *De institutione musica* (c. 500 CE) divides music into *musica mundana*, *musica humana* and *musica instrumentalis*, and *musica humana* is explicitly the proportion of the body itself. **That is precedent at approximately one tradition's weight, about 1.6 bits on the corpus's scale [Curzi, 2026w]. It records that the conclusion was reachable without a mechanism. It is not evidence, and every evidential claim in this section rests on the measured column.** Boethius carries no citation key here (Appendix A).

30.4.3 4.3 The cochlea, as an existence proof and nothing more

If the question is whether biology can build a mechanical receiver of extraordinary quality, the question is closed, and anatomy closed it long before molecular biology did.

The mammalian cochlea contains on the order of 3,500 inner and 12,000 outer hair cells — roughly 15,500 mechanoreceptors, **a nominal figure** — arrayed tonotopically along a basilar membrane of roughly 35 mm. Human hearing spans about 20 Hz to 20 kHz, which is

```
log2(20000 / 20) = 9.97 octaves
35 mm  / 9.97 octaves = 3.51 mm per octave
3500 IHC / 9.97 octaves = 351 inner hair cells per octave
```

Three hundred and fifty transducers per octave, across ten octaves, in a structure the size of a pea. That is the density of a well-designed filter bank [von Helmholtz, 1885, Plomp and Levelt, 1965].

The transduction apparatus is finer still. Hair-cell stereocilia are joined by tip links — fine extracellular filaments running from the tip of a shorter stereocilium to the side of its taller neighbour — and deflection tensions the link, opening the transduction channel directly. The kinetics are fast enough for the channel to be gated by force without an intervening second messenger, and the compliance change associated with gating is the observation from which the gating-spring model takes its name. **The identity of the pore-forming subunit remains an active argument, and this paper reports that as an open question rather than a settled one.** All the works concerned are named in Appendix A without keys.

Now the numbers, because they are why the cochlea belongs here at all. At the nominal threshold of hearing, 0 dB SPL, the sound pressure is 20 µPa and the intensity in air is

```
I = p^2 / (2 rho c) = (20e-6)^2 / (2 * 1.204 * 343) = 4.843e-13 W/m^2
```

Over a tympanic membrane of nominal area 55 mm² that is a delivered power of `P = 4.843e-13 x 55e-6 = 2.664e-17` W, and at 1 kHz the energy arriving per cycle is

```
E = P / f = 2.664e-20 J = 6.22 kT   (at 310 K)
```

Six kT per cycle, spread over the whole eardrum, is the entire energy budget of audible threshold. The middle ear then concentrates it: the area ratio between tympanic membrane and stapes footplate is roughly 55 mm² to 3.2 mm², about 17.2, and the ossicular lever adds roughly 1.3, for a nominal pressure gain of `17.2 x 1.3 = 22.3 = 27.0` dB. **That gain is what makes the air-to-fluid transition survivable, and it is the biological solution to exactly the impedance-mismatch problem Section 5.2 computes.**

Even after the gain, the mechanical displacement at threshold is of order 0.3 nm at the hair bundle. The work done against a gating spring of nominal stiffness 1 mN/m:

bundle displacement	gating-spring work (`E = 0.5 k x^2`, `k` = 1 mN/m)	in units of `kT` (310 K)
0.3 nm	4.500e-23 J	0.0105
1.0 nm	5.000e-22 J	0.1168
3.0 nm	4.500e-21 J	1.0514

At threshold the single-cycle mechanical work is roughly one hundredth of the thermal energy scale. A naive receiver could not detect it: the signal is buried a hundredfold below the noise on any single measurement. The ear detects it anyway, by the two classic routes — coherent averaging over many cycles, which improves signal-to-noise as the square root of the number integrated, and active mechanical amplification supplied by the outer hair cells. The general question of how close biological sensors run to physical limits is an old one, and the cochlea is its central example; that literature is named in Appendix A.

Two engineering lessons follow, and both are used again in Section 5.

Lesson one: a receiver beats the thermal floor by integration time and by active gain, not by wishing. The leaky integrator is exactly the right object [Curzi, 2026a, Gerstner and Kistler, 2002, Lapicque, 1907, Usher and McClelland, 2001]. Integration buys sensitivity and costs bandwidth, and the exchange rate is the Gabor limit [Gabor, 1946, Curzi, 2026z].

Lesson two, and it bounds the whole section: the cochlea is a purpose-built organ occupying a specialised, impedance-matched, actively amplified, mechanically isolated cavity, with a dedicated afferent nerve. A fibroblast in the middle of a tissue has none of that. Citing the cochlea as proof that *cells in general* hear finely is the same error as citing a radio telescope as proof that a paperclip receives radio. **The cochlea proves buildability. It does not prove that the build is present elsewhere.**

30.4.4 4.4 Force to the nucleus, and force to transcription

Does mechanical input at the surface reach the genome? It does, and the demonstration is direct: local force applied to the cell surface through integrin-bound magnetic beads produced, in the same experiment, stretching of a chromatin region and upregulation of transcription from a targeted reporter (Tajik et al., *Nature Materials* 15:1287–1296, 2016). **This is mechanical input to gene expression, published, direct, and not mediated by a diffusible second messenger in the reported pathway.**

It does not stand alone. The surrounding literature includes work on mechanotransduction at a distance and the physical coupling of extracellular matrix to nucleus through the cytoskeleton and the LINC complex; the shuttling of transcriptional co-activators to the nucleus as a function of mechanical cues including matrix stiffness and cell geometry (Dupont et al., *Nature* 474:179, 2011); and the finding that the elastic modulus of the substrate directs mesenchymal stem cell lineage specification (Engler, Sen, Sweeney and Discher, *Cell* 126:677, 2006) — **the stiffness of what a cell is standing on changes what the cell becomes.** All are named in Appendix A without keys.

Taken together this is a large, replicated body of work. **The claim that mechanical force changes gene expression is not a frontier claim; it is a textbook claim.**

And here is the boundary, drawn precisely where it belongs. All of the above is *amplitude* signalling. The cell responds to how much force, applied where, for how long, in what pattern of stiffness. **In no case in this literature does the cell recover a symbol that a sender chose.** Force is a scalar field; chromatin responds to it; there is no demonstrated stage that maps a waveform onto a sequence.

Precedent, priced. The general doctrine that structure at one scale is expressed at another, and the classical medical doctrine that regimen alters the body's constitution, record that outside-to-inside determination was a durable idea. **They supply no mechanism, no magnitude and no evidence, and the entire evidential weight here sits in the measured column** [Curzi, 2026w]. What the modern work adds is that "environment writes into the organism" is now a mechanistic statement with molecular detail where it was formerly an analogical one. **The gain from the precedent column is historical, not evidential.**

30.4.5 4.5 Sonogenetics: what it actually demonstrated

Sonogenetics is the closest existing work to "use sound to control cells", and it must be reported accurately rather than enthusiastically.

Neurons in *Caenorhabditis elegans* expressing the TRP-4 channel were activated by low-pressure ultrasound, producing behavioural responses (Ibsen, Tong, Schutt, Esener and Chalasani, *Nature Communications* 6:8264, 2015). **As reported, the effect required microbubbles in the medium to amplify the stimulus — a detail that is central rather than incidental, and that is stated here because it constrains what the result can be used to argue. The exact drive frequency reported is not restated here: where this author is not certain of a published figure, the figure is omitted rather than guessed.** The approach was later extended to mammalian cells by expressing a human channel and demonstrating ultrasound-evoked activation at megahertz frequencies (Duque et al., *Nature Communications* 13:600, 2022). **Again the sensitising element is an expressed transgene: the cell was made ultrasound-responsive by putting a particular channel into it.**

What sonogenetics establishes, stated exactly:

```
S1  Ultrasound can activate a specific, chosen channel in a living animal or
    cell, with sufficient specificity that cells expressing the channel respond
    and neighbouring cells largely do not.

S2  The specificity comes from WHICH CHANNEL IS PRESENT, not from what the sound
    says. It is receiver-side selectivity by molecular identity -- exactly
```

```
    analogous to a radio band selected by which crystal is soldered in.

S3  In at least the 2015 C. elegans work, amplification by microbubbles was
    required as reported.
```

S2 is the sentence to carry forward. Sonogenetics is addressing by receiver identity, not addressing by message content. That is a real and useful architecture — Section 6 shows that molecular addressing is the only layer in this whole architecture with cell-scale spatial resolution — **but it is not demodulation, and it must not be described as if it were. A sonogenetic system carries, at most, one bit per stimulus epoch to the cells that possess the transgene: on, or off.**

30.4.6 4.6 The honest mechanical limit: why a sound field is a poor cell-scale stimulus

This section exists because the analysis would be dishonest without it.

A cell is very much smaller than an audible wavelength. In soft tissue at a nominal 1540 m/s:

frequency	wavelength in tissue	wavelength / 10 µm cell
20 Hz	77 m	7.70e6
20 kHz	0.077 m	7,700
1 MHz	1.54 mm	154
10 MHz	154 µm	15.4
154 MHz	10 µm	1.0

The mechanically relevant consequence is that at ordinary frequencies **the cell sits inside a region of essentially uniform pressure.** The fractional pressure difference across a cell of diameter `d` is approximately `k d` with `k = 2 pi / lambda`:

frequency	`k d` for `d` = 10 µm	differential pressure across the cell
20 Hz	8.16e-7	0.0000816 % of peak
20 kHz	8.16e-4	0.0816 %
1 MHz	0.0408	4.08 %
10 MHz	0.408	40.8 %
154 MHz	6.283	628 % (multiple wavelengths across the cell)

A cell is, to good approximation, an incompressible bag of water. **Uniform pressure on an incompressible body produces almost no membrane tension, and membrane tension is what gates the channel.** What actually generates tension from an acoustic field is the *non-uniform* part: pressure gradients, shear at interfaces of differing impedance, acoustic streaming, radiation force, and — at higher intensities — cavitation, **which is why the sonogenetics work reached for microbubbles.**

Acoustic radiation force on a single absorbing cell can be estimated. Taking a cell radius of 5 µm (volume `5.236e-16` m³), the nominal soft-tissue attenuation of 0.5 dB/cm at 1 MHz converted to an amplitude attenuation coefficient of 5.756 Np/m, and absorbed power `P = 2 alpha I V` with force `F = P/c`:

peak pressure	intensity (W/m²)	absorbed power (W)	radiation force
100 Pa	3.092e-3	1.864e-17	1.21e-8 pN
1 kPa	3.092e-1	1.864e-15	1.21e-6 pN
10 kPa	3.092e+1	1.864e-13	1.21e-4 pN
100 kPa	3.092e+3	1.864e-11	1.21e-2 pN
1 MPa	3.092e+5	1.864e-9	1.21 pN

Compare that to the force scale of gating: **doing 1 kT of work over a 4 nm conformational swing requires 1.07 pN; 5 kT requires 5.35 pN.**

At a megapascal of peak pressure the acoustic radiation force on a single cell is of the same order as the force needed to do one channel's worth of gating work — and at kilopascal pressures it is six orders of magnitude below it. That single comparison explains the shape of the entire sonogenetics literature: why effects are sought at high frequency and high pressure, why microbubbles or gas vesicles are so often involved, and why the field is hard. **It is also the clearest available answer to a reader who asks whether ordinary audible sound at ordinary levels mechanically gates channels deep in tissue. On these numbers, by the direct radiation-force route, it does not come close.**

This is stated here, in the section arguing that the cell *is* a receiver, deliberately. **A claim is worth more when the analysis that makes it also computes the case against it.**

30.4.7 4.7 The cell as a leaky integrator: a badged model-form claim

The corpus's central identity is the first-order accumulation

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t
```

and its parameters, not its form, are what change between scopes [Curzi, 2026a]. At cellular scope the reading is:

term	cellular interpretation
`Q`	durable regulatory state: chromatin configuration, methylation pattern, stable protein levels, cytoskeletal prestress
`S`	this period's signalling input, including mechanical input
`eta`	the fraction of signalling that becomes durable regulatory state rather than a transient
`delta`	reversal: mark turnover, protein degradation, dilution by division
`C`	metabolic and proteostatic cost of holding the state
`h`	headroom: `eta ln N − delta Q − C`

This is a claim about model form and not about physics, and it is badged as such here as it is everywhere in the corpus. **The cell is not asserted to obey the storage law**; the assertion is that the same first-order accumulation *shape* is a useful description, useful because it makes the right qualitative predictions and forbids the wrong ones. It is a **category-2 borrowing** in the sense of Section 3.4, and Section 3.6 records that it is not identified from anything measured here.

Two predictions are worth spelling out, computed with `N = 4` (an illustrative four-category input set), `eta = 0.4`, `C = 0.1`. **All three parameters are assumed, not measured.**

`delta`	`Q_inf = (eta ln 4 − C)/delta`	time constant (periods)	half-life (periods)
0.02	22.726	50.0	34.31
0.10	4.545	10.0	6.58
0.30	1.515	3.3	1.94

Prediction one: a transient stimulus leaves a transient trace. At `delta = 0.3` the half-life of any deposit is under two periods. **Any protocol proposing that a single exposure produces a permanent cellular change is proposing that `delta` is near zero for the relevant state variable — an empirical claim about mark stability that should be stated and tested rather than assumed.**

Prediction two: success consumes headroom.

stored state `Q`	`delta` = 0.02	`delta` = 0.10
0	h = +0.4545	h = +0.4545
2	h = +0.4145	h = +0.2545
5	h = +0.3545	h = -0.0455
10	h = +0.2545	h = -0.5455

At `delta = 0.10` the system goes headroom-negative between `Q = 2` and `Q = 5`: past that point the maintenance burden of the state already held exceeds what the input can supply, and **no arrangement of the input geometry prevents decline** [Curzi, 2026a]. A differentiated cell holding a large, expensive, stable regulatory state is, in this description, a low-headroom system — **which is a formal restatement of the observation that differentiated cells are hard to reprogram, and is taken up in earnest by the companion paper** [Curzi, 2026y]. **It is a restatement, not evidence for the observation.**

30.4.8 4.8 What the precedent column preserved, read as engineering

Three items are worth recording, **strictly as precedent and priced in bits** [Curzi, 2026w]. All are named without citation keys (Appendix A).

The Indian doctrine of *shabda* and *nada*. Vedic and later Indian metaphysics treats sound as constitutive rather than incidental, and the grammarian Bhartṛhari's *Vākyapadīya* (c. 5th century) develops a technical theory in which meaning-bearing sound has a structure distinct from its acoustic realisation. **As precedent this is the oldest sustained attempt on record to distinguish a signal's physical carrier from its content — which is, formally, the carrier/payload distinction of Section 6.9.** That is a genuine structural correspondence and it is worth recording. **It is not evidence that either the metaphysics or the mechanism is correct.**

The Chinese *lü* pitch standards. The twelve *lü* pipes, discussed in the *Lüshi Chunqiu* (3rd century BCE) and later sources, tied pitch standards to calendar and statecraft. **As precedent this records the idea that a system's structure and its tuning are the same problem.** It should be recorded together with an honest note, because the honesty is the useful part: the associated *houqi* practice, in which buried pitch-pipes charged with ash were watched for the ash to be expelled at the arrival of the seasonal *qi*, is a recorded attempt at an *experiment* that was disputed within the tradition and has no modern replication. **A tradition that ran an experiment and could not make it work is a more instructive precedent than one that never tried — and a recorded null on the precedent side is not evidence for anything on the measured side.**

Boethius on *musica humana*, entered at Section 4.2.

Against all three, the standing rule: **precedent tells you the idea was findable; only the measured column tells you whether it is true** [Curzi, 2026w].

30.4.9 4.9 The stopping point

ESTABLISHED	Cells possess identified, prize-recognised, mechanically gated ion channels. Sound is mechanical. Therefore a real pathway runs from a pressure wave to a cellular signal.
ESTABLISHED	Force at the cell surface reaches the nucleus, stretches chromatin, and changes transcription.
ESTABLISHED	Ultrasound can activate a chosen, expressed channel with receiver-side selectivity.
NOT ESTABLISHED	Any mechanism by which a cell demodulates sequence-specific information from an acoustic signal.
COMPUTED LIMIT	At cellular scale an acoustic field is nearly uniform; radiation force on a single cell reaches the single-channel gating force scale only at roughly megapascal pressures.

The cell is a receiver. It is not, on any published evidence, a decoder.

30.5 5 The channel and its capacity

Section 4 established that a cell is a mechanical receiver. This section asks the engineer's next question, which is the only question that can settle an argument: **how much information can physically arrive at it?**

That question has had a formal answer since 1948. Shannon's channel-capacity theorem bounds the rate at which information can be transmitted over a noisy channel with arbitrarily small error probability, and the bound depends on exactly two things: the bandwidth available and the signal-to-noise ratio [Shannon, 1948, Shannon and Weaver, 1949]. Both are computable for a sound wave arriving at a cell in soft tissue.

An honest capacity bound is the most valuable thing this analysis can contain, whatever the bound turns out to be, and a number that comes out unfavourable is worth more than a paragraph that avoids the calculation. The result is stated at the front so that nobody has to read to the end to find the verdict:

THE CHANNEL IS NOT THE BOTTLENECK.	At the cell's own receiver bandwidth and a modest signal-to-noise ratio, the computed capacity is of order tens of bits per second. A forty-bit address takes about half a second.
THE SPATIAL RESOLUTION IS CATASTROPHIC.	At 1 MHz, roughly 3.65 million cells share one acoustic resolution voxel. Every one of them receives the same bits. The channel is a BROADCAST.
THE DEMODULATOR DOES NOT EXIST.	Capacity is an upper bound on what a perfect decoder could recover. Nothing in the literature identifies a decoder.

Capacity is permission, not performance. Shannon's theorem says a code exists which approaches the bound. It says nothing whatever about whether the receiver in question implements one. **That distinction is the intellectual core of this section and it is the reason the section can be quantitative and honest at the same time.**

30.5.1 5.1 The three bounds that govern

```
B1  SHANNON CAPACITY        C = B log2(1 + S/N)          bits per second
    C is the supremum of rates at which information can be sent with
    arbitrarily small error. B is bandwidth in Hz; S/N is signal power to
    noise power in the same band.

B2  THE GABOR LIMIT         dt df >= 1/(4 pi) = 0.079577
    Classical Fourier duality. Applies to every wave including sound and
    carries NO quantum content.

B3  THE THERMAL FLOOR       fluctuation-dissipation; noise power ~ 4 kT R B
    Any dissipative receiver at temperature T carries an irreducible noise of
    this form. At 310 K, kT = 4.280e-21 J = 4.280 pN.nm.
```

The bounds are B1 [Shannon, 1948], B2 [Gabor, 1946] and B3 [Callen and Welton, 1951]. **None of them is a postulate of this framework**, and to them must be added two propagation facts that set the parameters rather than being inequalities: the acoustic impedance step at any boundary, and frequency-dependent attenuation in tissue.

The standing constraint, restated: *there is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal.* **This section computes a channel capacity, and a channel capacity is not a mechanism.** Computing that a copper wire can carry a gigabit per second tells you nothing about whether the object at the far end is a modem or a doorstop. **The whole value of the computation is that it gives the size of the pipe, and thereby removes "the pipe is too small" from the list of available explanations, leaving the actual explanation exposed.**

30.5.2 5.2 The impedance step: what gets in

Acoustic impedance is `Z = rho c`. Computed at the nominal constants of Section 3.3:

```
Z_air    = 1.204 * 343   =        413.0 Pa.s/m
Z_tissue = 1050  * 1540  = 1.6170e+06   Pa.s/m
ratio    = 3915.5
```

At normal incidence the intensity transmission coefficient is `T_I = 4 Z1 Z2 / (Z1 + Z2)^2 = 1.021e-3 = -29.91` dB, and the pressure transmission coefficient is `t_p = 2 Z2 / (Z1 + Z2) = 1.99949 = +6.02` dB.

Both are correct and they point in opposite directions. Ninety-nine point nine per cent of the incident intensity reflects at an air–tissue boundary, while the pressure amplitude just inside the tissue nearly doubles. The apparent paradox resolves because particle velocity collapses by the same factor that pressure gains, and intensity is their product.

Which figure matters depends on what the receiver responds to. Mechanically gated channels respond to force, hence to pressure, **so the pressure figure is the operative one and it is the generous one. This paper uses the generous figure deliberately: a bound argued on the most favourable available assumption is a bound worth having.**

Worked, for an airborne 94 dB SPL source (1.0024 Pa in air):

```
I_air     = p^2/(2 rho_a c_a)  = 1.2165e-03 W/m^2
I_tissue  = I_air * T_I        = 1.2421e-06 W/m^2
p_tissue  = sqrt(2 rho c I)    = 2.0042 Pa
```

This is also the exact problem the middle ear solves. Section 4.3 computed the nominal ossicular pressure gain at 22.3, or 27.0 dB. **Biology built a matching transformer, because the impedance step is the first and largest loss in the chain and is unavoidable without one.**

30.5.3 5.3 Wavelength against the cell, and why acoustic addressing fails

Wavelength in soft tissue, computed at `c = 1540` m/s (nominal):

frequency	wavelength	wavelength / 10 µm cell
20 Hz	77 m	7.70e6
200 Hz	7.7 m	7.70e5
2 kHz	0.77 m	7.70e4
20 kHz	0.077 m	7,700
100 kHz	15.4 mm	1,540
1 MHz	1.54 mm	154
3 MHz	513 µm	51.3
10 MHz	154 µm	15.4
20 MHz	77 µm	7.7
50 MHz	30.8 µm	3.08
154 MHz	**10 µm**	**1.00**
1 GHz	1.54 µm	0.154

The frequency at which the wavelength equals a cell diameter is exactly `f = c/d = 1540 / 10e-6 = 1.540e8` Hz `= 154` MHz.

Now take an acoustic resolution voxel as nominally `lambda^3` — **a deliberately optimistic convention, since real focused-beam resolution is worse than a wavelength laterally** — and count the cells inside it at a nominal `1e9` cells per cm³:

frequency	`lambda^3`	cells per voxel
20 kHz	4.57e5 mm³	4.57e11
1 MHz	3.652 mm³	3.65e6
10 MHz	3.65e-3 mm³	3.65e3
50 MHz	2.92e-5 mm³	29.2
100 MHz	3.65e-6 mm³	3.65
154 MHz	1.00e-6 mm³	1.00

At one megahertz, three and a half million cells occupy a single resolution element and every one of them receives exactly the same waveform. The acoustic channel into tissue is not a point-to-point link. It is a broadcast.

That single line reorganises the architecture, and it is why Section 6 exists. **Selectivity at cellular scale cannot come from the acoustic field, because the field cannot resolve a cell without going to frequencies at which it cannot reach one.** Selectivity must come from somewhere else, and the only candidate with the required resolution is molecular addressing.

30.5.4 5.4 Attenuation, and the trade that closes the door

Under the nominal law `alpha = 0.5 f_MHz` dB/cm, one-way attenuation:

frequency	1 cm	5 cm	10 cm	20 cm
20 kHz	0.01 dB	0.05 dB	0.10 dB	0.20 dB
100 kHz	0.05	0.25	0.50	1.00
1 MHz	0.50	2.50	5.00	10.00
3 MHz	1.50	7.50	15.00	30.00

frequency	1 cm	5 cm	10 cm	20 cm
10 MHz	5.00	25.00	50.00	100.00
30 MHz	15.00	75.00	150.00	300.00
100 MHz	50.00	250.00	500.00	1000.00
154 MHz	77.00	385.00	770.00	1540.00

Define `d20` as the depth at which one-way attenuation reaches 20 dB — a factor of ten in pressure, a hundred in intensity. Then:

frequency	`lambda`	`d20`	product `lambda x d20`
100 kHz	15.4 mm	400 cm	6160 mm·cm
1 MHz	1.54 mm	40 cm	61.6 mm·cm
10 MHz	154 µm	4 cm	0.616 mm·cm
100 MHz	15.4 µm	0.4 cm	0.00616 mm·cm
154 MHz	10 µm	0.26 cm	0.0026 mm·cm

The product falls as `1/f^2`. Resolution and depth are not merely in tension; they trade against each other at the square of the frequency. Ten times finer resolution costs a hundred-fold in reachable volume. **This is the most consequential number here for anyone imagining a system that addresses individual cells acoustically at depth.**

Sensitivity to the attenuation model, computed because the conclusion must not depend on the convenience of a round exponent:

exponent `b` in `alpha = 0.5 f^b`	`d20` at 10 MHz	at 100 MHz	at 154 MHz
1.0	4.000 cm	0.400 cm	0.260 cm
1.1	3.177 cm	0.252 cm	0.157 cm
1.5	1.265 cm	0.040 cm	0.021 cm

Every plausible exponent makes the conclusion worse, not better. The nominal linear law is the most generous of the three, and it is the one used throughout.

30.5.5 5.5 The thermal noise floor: two estimates, both flagged

Noise is where honest capacity analyses are won or lost, so two independent estimates are given with the validity regime of each stated, rather than picking whichever is convenient.

Estimate one: radiation-resistance thermal noise at an aperture. For a receiving aperture of area `A` in a medium of impedance `rho c`, in the regime where the aperture is large compared with a wavelength, the radiation resistance is approximately `rho c A` and the fluctuation–dissipation theorem [Callen and Welton, 1951] gives a mean-square thermal pressure `p_rms = sqrt(4 kT rho c B / A)`. Computed at 310 K in soft tissue:

aperture diameter	area	B = 100 Hz	B = 1 kHz	B = 1 MHz
1 µm	7.854e-13 m²	1.877 Pa	5.937 Pa	187.7 Pa
5 µm	1.964e-11	0.3755	1.187	37.55
10 µm	7.854e-11	0.1877	0.5937	18.77
20 µm	3.142e-10	0.0939	0.2969	9.387
100 µm	7.854e-09	0.0188	0.0594	1.877

aperture diameter	area	B = 100 Hz	B = 1 kHz	B = 1 MHz
5.5 mm (eardrum scale)	2.376e-05	3.41e-4	1.08e-3	3.41e-2

The scaling is the lesson: thermal pressure noise goes as one over the square root of aperture area. Small receivers are noisy receivers. A 10 µm aperture is 70 times noisier in pressure than a 5.5 mm one at the same bandwidth. **That is a general design law and it is why the ear is the size it is.**

The honest caveat, stated rather than buried. The `rho c A` radiation resistance holds for `k a >> 1`, and at cellular scale in the audible and low-megahertz band `k a << 1` — Section 4.6 computed `k d = 4.08e-2` at 1 MHz for a 10 µm cell. In that regime the radiation resistance is much smaller than `rho c A`, so the formula overstates the dissipation and therefore overstates the noise; **but the same small radiation resistance also weakens the coupling of *signal* into the receiver, so signal and noise fall together and the signal-to-noise ratio does not simply improve. A rigorous treatment requires the fluctuation–dissipation theorem applied to the actual mechanical impedance of the cell, which this paper cannot state with confidence.** The figures above are therefore a reference calculation with a declared validity regime, **not a characterisation of a cell.**

Estimate two: Brownian force noise on the channel itself, which is nearly regime-independent because it concerns the immediate viscous environment of the protein. For an object of radius 5 nm in water (`eta = 1e-3` Pa·s), Stokes drag gives `gamma = 6 pi eta r`:

```
gamma  = 9.425e-11 kg/s
F_th   = sqrt(4 kT gamma B)  = 4.017e-14 N = 40.2 fN   at B = 1 kHz
```

Against the gating scale of Section 4.6 — **1.07 pN does 1 kT of work over a 4 nm swing** — a 40 fN thermal force is about twenty-five times smaller. **The channel is not thermally drowned; it is thermally hovering, close enough to the floor that the receiver's own integration time matters.** That is precisely the cochlea's situation, and it is why the cochlea integrates and amplifies.

30.5.6 5.6 The receiver's own bandwidth: the term that binds

This is the step most treatments of the subject skip, and it determines the answer.

The bandwidth `B` in Shannon's formula is not the bandwidth of the transmitter or of the medium. **It is the bandwidth over which the receiver can distinguish successive states.** A megahertz carrier delivered to a receiver whose output cannot change faster than ten times a second carries one bit per tenth of a second, not one bit per microsecond.

Cellular receiver stages have measured or estimated time constants. Converting each to a corner frequency `f_c = 1/(2 pi tau)`:

stage	time constant `tau` (nominal)	corner frequency `f_c`
cochlear MET channel activation	~20 µs	7,957.75 Hz
fast mechanotransduction generally	~100 µs	1,591.55 Hz
PIEZO1 inactivation	~10 ms	15.92 Hz
PIEZO1 inactivation	~15 ms	10.61 Hz
calcium transient	~1 s	0.1592 Hz
transcriptional response	~30 min	8.84e-5 Hz

All time constants above are nominal figures used to set scale. They are not asserted as measured constants of any particular preparation, and channel inactivation in particular is known to depend on membrane potential, isoform and preparation.

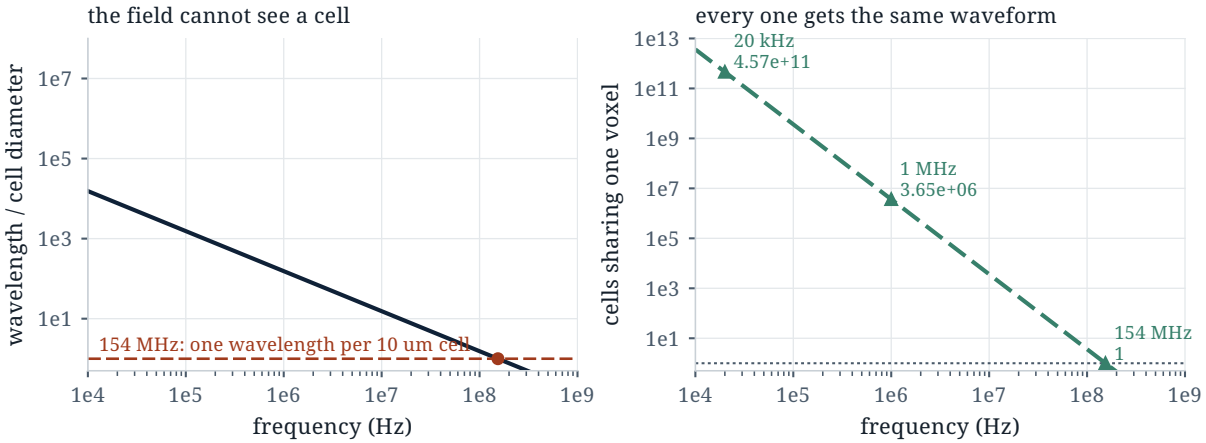


Figure 30.1: **The wavelength resolves a cell only at 154 MHz, and until then the field is a broadcast to millions.** Left: wavelength in soft tissue at a nominal 1540 m/s against a 10 µm cell diameter, with the crossing at 154 MHz marked. Right: cells sharing one `lambda^3` resolution voxel at a nominal `1e9` cells per cm³, falling from `4.57e11` at 20 kHz to 1.00 at 154 MHz; the marked point at 1 MHz is the `3.65e6` cells of Section 5.3. Plotted as small multiples because a length and a count share no scale, and on logarithmic axes because both span eleven orders of magnitude. Computed from the nominal constants of Section 3.3; generated by `figs/src/CCP-030_fig01.py`.

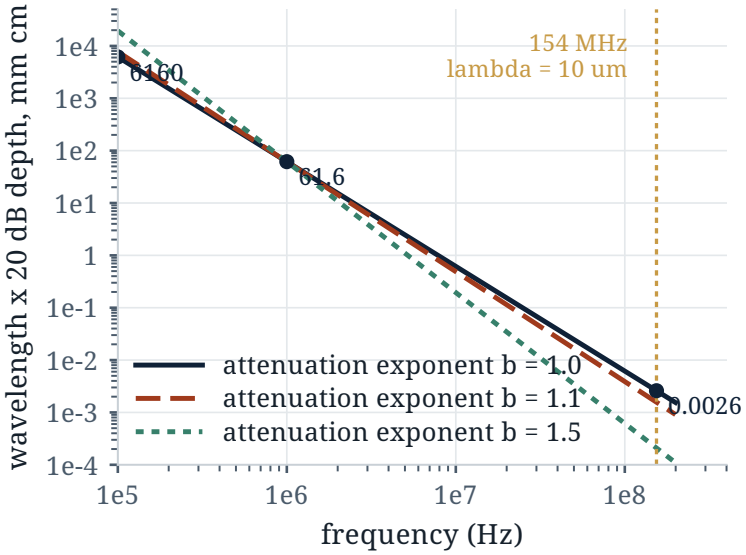


Figure 30.2: **Resolution and depth trade at the square of the frequency, and every plausible attenuation exponent makes it worse.** The product of wavelength and 20 dB one-way depth falls from 6160 mm·cm at 100 kHz to 0.0026 mm·cm at 154 MHz, a straight line of slope −2 on logarithmic axes, under the nominal `b = 1.0` attenuation law. The two lower curves recompute `d20` under `b = 1.1` and `b = 1.5`; at 154 MHz the 20 dB depth falls from 0.260 cm to 0.157 cm and 0.021 cm respectively. The vertical marker is the 154 MHz frequency at which the wavelength first resolves a 10 µm cell. Computed from the nominal attenuation law of Section 3.3; generated by `figs/src/CCP-030_fig02.py`.

This table is the engineering pivot. Whatever the acoustic field is doing at megahertz rates, the cell's own transducer is a low-pass filter — which is the leaky integrator at cellular scope, with `1 − delta` as the pole [Curzi, 2026a, Oppenheim et al., 1996, Horowitz and Hill, 2015]. **A single-pole low-pass stage is the most ordinary object in control engineering** [Åström and Murray, 2008], **and the fixed-gain recursive estimator it is adjacent to is equally standard** [Kalman, 1960]; **naming that is a category-2 borrowing and nothing more.** The signalling stage is slower still, and the transcriptional stage slower by five further orders of magnitude.

30.5.7 5.7 The capacity calculation

`C = B log2(1 + S/N)` evaluated at each receiver stage, for three signal-to-noise ratios, with the time to deliver a 40-bit payload — the size of a 20-nucleotide address, Section 6.2 — at the 20 dB column:

receiver stage	`B` (Hz)	`C` at 10 dB	`C` at 20 dB	`C` at 30 dB	40 bits at 20 dB
transcriptional (~30 min)	8.84e-5	3.06e-4 bit/s	5.89e-4 bit/s	8.81e-4 bit/s	67,944 s (18.9 h)
calcium transient (~1 s)	0.1592	0.551	1.059	1.586	37.75 s
PIEZO1 (~15 ms)	10.61	36.67	70.65	105.8	0.5662 s
PIEZO1 (~10 ms)	15.92	55.01	106.0	158.6	0.3775 s
fast MET (~100 µs)	1,591.6	5,506	1.060e4	1.586e4	3.775 ms
cochlear MET (~20 µs)	7,957.8	2.753e4	5.298e4	7.932e4	0.755 ms

For completeness, the same calculation with bandwidth set by the *medium* rather than by the receiver, using the aperture-noise estimate of Section 5.5 at a 10 µm aperture and `B` = 1 kHz, `p_noise` = 0.5937 Pa:

signal pressure	intensity (W/cm²)	SNR (dB)	capacity (bit/s)
0.1 Pa	3.09e-13	-15.47	40.4
1 Pa	3.09e-11	+4.53	1,940
100 Pa	3.09e-07	+44.53	14,792
10 kPa	3.09e-03	+84.53	28,080
1 MPa	3.09e+01	+124.53	41,368

Note the shape of that second table. Capacity grows only logarithmically in signal pressure: a ten-thousandfold increase between 100 Pa and 1 MPa buys a factor of 2.8. This is the same logarithmic ceiling that governs `ln N` in the surplus law [Curzi, 2026], appearing again in a completely different subject because it is the same mathematics. **You cannot shout your way to a fast channel.**

The result, stated plainly. Taking the honest, receiver-limited figure — the cell's own mechanotransduction bandwidth of about 10 to 16 Hz, at a signal-to-noise ratio of 20 dB — the channel capacity into a single cell is of order 70 to 106 bits per second, and a 40-bit address would take about half a second to deliver.

Three things it does and does not mean.

It means the pipe is adequate. Half a second for an address is not a limitation. If the only problem were bandwidth, this architecture would work. **"There is not enough capacity"**

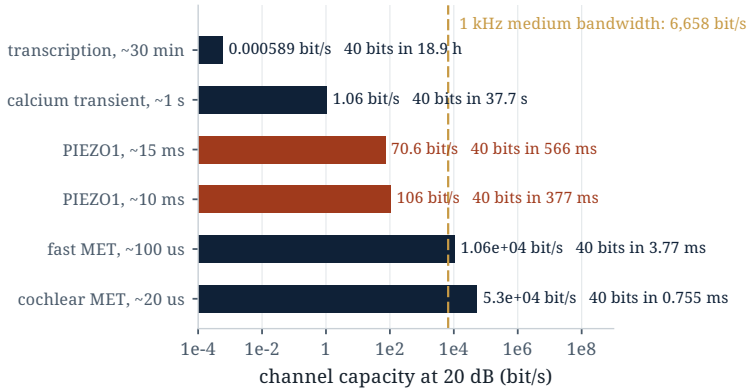


Figure 30.3: **The receiver's own kinetics, not the medium, set the rate — and the range across stages is nine orders of magnitude.** Shannon capacity `C = B log2(1 + S/N)` at 20 dB for the six receiver stages of Section 5.6, with the time to deliver a 40-bit address annotated. The two mechanotransduction rows give 70.65 and 106.0 bit/s, delivering the address in 0.566 s and 0.378 s; the transcriptional stage would need 18.9 hours. Substituting a 1 kHz medium bandwidth instead returns 6,658 bit/s, marked, an inflation of 94.2 times; substituting a full 1 MHz carrier returns `6.66e6` bit/s, a factor of `9.4e4`. Computed from the nominal time constants of Section 5.6; generated by `figs/src/CCP-030_fig03.py`.

is hereby removed from the list of explanations, and removing a wrong explanation is real progress.

It means the pipe is shared. At 1 MHz, `3.65e6` cells occupy one voxel and receive identical bits. If each cell in one voxel needed its own distinct 40-bit address:

```
40 bits * 3.6523e6 cells = 1.4609e8 bits
at 70.65 bit/s           = 2.0678e6 s = 23.93 days
```

Twenty-four days to individually address the cells in a single voxel, and that assumes the addressing worked at all. The per-cell exclusive rate at 1 MHz is `1.93e-5` bit/s. **Broadcast is not a limitation of the design; it is a property of the physics, and the only escape is receiver-side selectivity — which is Section 6.**

It does not mean anything is decoded. Shannon's theorem is an existence proof: for any rate below `C` there exists a code achieving arbitrarily small error. **It does not construct the code, it does not construct the decoder, and it makes no claim that any particular physical system implements either. A channel capacity of 70 bits per second into a cell is a statement about the channel. There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal, and computing the capacity of the channel leading to a receiver does not create the receiver.** Measuring that a room has excellent acoustics tells you nothing about whether anyone in it speaks the language.

30.5.8 5.8 The Gabor limit at cellular timescales

The Gabor limit is `dt df >= 1/(4 pi) = 0.079577` [Gabor, 1946]. **It is classical, it applies to every wave including sound, it is the same Fourier mathematics as the position–momentum relation with no quantum content, and it is measurable on an oscilloscope** [Curzi, 2026z]. Applied at the timescales of Section 5.6:

observation window `dt`	minimum frequency resolution `df`
1 µs	79,577.5 Hz
10 µs	7,957.75 Hz
100 µs	795.775 Hz
1 ms	79.5775 Hz
5 ms	15.9155 Hz
15 ms	5.30516 Hz
100 ms	0.795775 Hz
1 s	0.0795775 Hz
60 s	0.00132629 Hz
1 hour	2.21049e-5 Hz

Read this against the receiver stages and a structural fact appears. **A hair cell integrating over 100 µs cannot have a frequency resolution better than 796 Hz — at a 1 kHz carrier that is a quality factor of about 1.26, which is barely a filter at all.** The cochlea's actual frequency selectivity is far sharper, and the resolution of the apparent contradiction is that the cochlea does *not* make its decision in 100 µs: the mechanical filter integrates over many cycles and the active amplification narrows the tuning further. **Sharpness in frequency is bought with time, at the Gabor exchange rate, and there is no other currency.**

Conversely, a slow receiver has extraordinary potential frequency resolution:

stage	window	best possible `df`	quality factor at a 1 kHz carrier
hair-cell MET	100 µs	795.8 Hz	1.26
PIEZO1	15 ms	5.31 Hz	188
calcium integration	1 s	0.0796 Hz	1.26e4
transcription	30 min	4.42e-5 Hz	2.26e7

A transcriptional response integrating over half an hour could in principle discriminate frequencies to forty-four microhertz. **Nothing is known to use that resolution, and stating a theoretical capability is not stating that anything realises it** — but the number belongs in the record, because it shows that the slow end of the cell is the *high-resolution* end in frequency, which is the opposite of most readers' intuition.

30.5.9 5.9 The Gabor bound as a rate bound, and the crossover with Shannon

There is a second, less familiar way to use the Gabor limit, and it produces a clean number.

Suppose the channel carries information by choosing one of `M` distinguishable tones inside a bandwidth `B`, holding each for a symbol time `dt`. For the tones to be distinguishable they must be separated by at least `df = B/M`, and Gabor requires `df >= 1/(4 pi dt)`. Therefore `M <= 4 pi B dt`, and the information rate is `R = log2(M)/dt = log2(4 pi B dt)/dt`. Differentiating and setting to zero puts the optimum at `4 pi B dt = e`, so

```
dt*    = e / (4 pi B)
R_max  = 4 pi B log2(e) / e  =  6.669448 * B    bits per second
```

A Gabor-limited tone-selection channel carries at most about 6.67 bits per second per hertz of bandwidth. Computed instances:

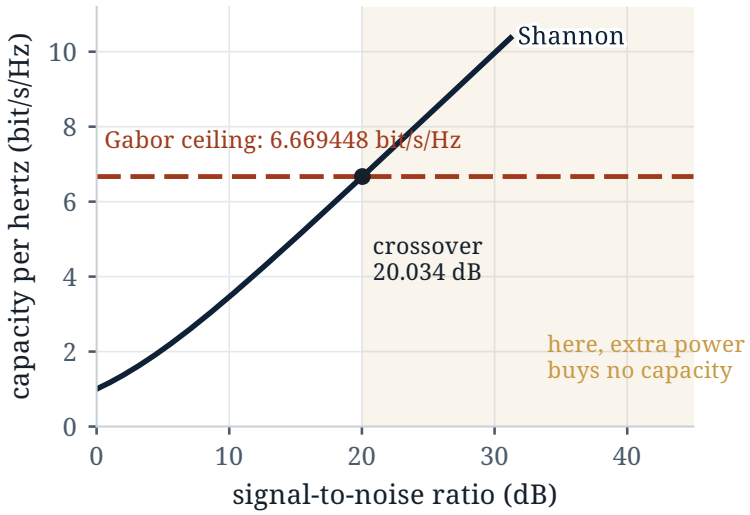


Figure 30.4: **Two bounds cross at 20.034 dB, and which one binds depends on which side you are.** Capacity per hertz of bandwidth: Shannon's `log2(1 + S/N)` rising with signal-to-noise ratio, against the Gabor tone-selection ceiling of a constant 6.669448 bit/s/Hz. They meet exactly at `S/N = 2^6.669448 − 1 = 100.79`, which is 20.034 dB; below it noise binds and above it time–frequency resolution binds. The shaded region is where extra power no longer buys capacity. Computed from the two closed forms of Sections 5.7 and 5.9; generated by `figs/src/CCP-030_fig04.py`.

bandwidth `B`	optimal symbol time `dt*`	`R_max`
20 kHz	10.82 µs	1.334e5 bit/s
1 MHz	216.3 ns	6.669e6 bit/s
10 MHz	21.63 ns	6.669e7 bit/s
154 MHz	1.405 ns	1.027e9 bit/s

And now the cross-check that makes the result trustworthy. Shannon's `C = B log2(1 + S/N)` equals `6.669448 B` when `log2(1 + S/N) = 6.669448`, that is

```
S/N = 2^6.669448 − 1 = 100.79 = 20.034 dB
```

The Gabor-limited tone channel and the Shannon capacity agree exactly at a signal-to-noise ratio of 20.034 dB. Below that ratio Shannon binds and noise is the limiting factor; above it, Gabor binds and time–frequency resolution is. **That crossover is a genuinely useful design number and, so far as this author is aware, it is not a standard textbook figure — it falls straight out of putting the two bounds side by side.** The practical reading: **20 dB signal-to-noise is the point past which spending more power stops buying capacity and bandwidth must be bought instead.**

30.5.10 5.10 Energy is not the constraint

Landauer's principle sets the minimum energy to erase one bit at `kT ln 2` [Landauer, 1961], which at 310 K is `2.967e-21` J. **This is thermodynamic entropy; the bits of Section 5.7 are Shannon information, and the two are different objects brought together only for a bound.** Computing the acoustic power delivered to a 10 µm cell (area `7.854e-11` m²) and dividing:

pressure	intensity (W/m²)	power on cell	Landauer-limited bits/s
1 Pa	3.092e-7	2.429e-17 W	8.19e3
100 Pa	3.092e-3	2.429e-13 W	8.19e7
10 kPa	3.092e+1	2.429e-9 W	8.19e11
1 MPa	3.092e+5	2.429e-5 W	8.19e15

Even at one pascal — a pressure comparable to the thermal floor at that aperture — the thermodynamic bound permits eight thousand bits per second, more than a hundred times the receiver-limited capacity. Energy is not scarce in this problem. That is worth stating because "there is not enough energy" is another intuitively appealing wrong explanation, and it is now also removed.

Three explanations removed, then: not bandwidth, not energy, not thermodynamics. What remains is spatial resolution, which is catastrophic and physical, and the absence of a demodulator [Curzi, 2026y].

30.5.11 5.11 What the cell could do about noise, and what it costs

The receiver's one classical defence against a poor signal-to-noise ratio is coherent averaging. Averaging `n` cycles improves amplitude SNR as `sqrt(n)`:

cycles averaged	SNR gain
1	0 dB
10	5 dB
100	10 dB
1e4	20 dB
1e6	30 dB

Twenty decibels of gain costs ten thousand cycles of integration, and by the Gabor limit those cycles are time no longer available for the next symbol. **The exchange is not optional and it is not negotiable**, and it is the same exchange the corpus has been making since the storage law: **integration buys retention and costs responsiveness** [Curzi, 2026a]. The leaky integrator with a small `delta` averages well and reacts slowly; with a large `delta` it reacts fast and remembers nothing. **There is one dial.**

30.5.12 5.12 The essence ladder under the Gabor constraint

The corpus's essence law is `f(n) = k n^2` with `k = 0.483425732988414` Hz, and — **restating the absolute honesty clause — the output is an essence frequency, not a physical resonance. No claim is made that any material, molecule, element or person physically resonates at these values.** `k` is a **defined** constant, exactly `(9/8)^2 / phi^2`, so its precision is legitimate; **attaching hertz to it is a separate and stronger claim and is POSTULATED** (Section 3.4), and the standard scepticism about golden-section claims is cited alongside it [Markowsky, 1992].

With that badge on, the Gabor limit imposes a computable rendering constraint. Adjacent essence indices are separated by `f(n+1) − f(n) = k(2n + 1)`, so the minimum window needed to resolve index `n` from index `n+1` is `dt >= 1/(4 pi k (2n+1))`:

index `n`	`f = k n^2` (Hz)	gap to `n+1` (Hz)	minimum window `dt` (s)
1	0.483426	1.450277	0.054871
2	1.933703	2.417129	0.032922
3	4.350832	3.383980	0.023516
4	7.734812	4.350832	0.018290
5	12.085643	5.317683	0.014965
6	17.403326	6.284535	0.012662
7	23.687861	7.251386	0.010974
8	30.939247	8.218237	0.009683
9	39.157484	9.185089	0.008664
10	48.342573	10.151940	0.007839
11	58.494514	11.118792	0.007157
12	69.613306	12.085643	0.006584

The bottom of the essence ladder cannot be rendered faster than about 55 milliseconds per symbol without smearing index 1 into index 2, and the constraint relaxes as `1/(2n+1)` up the ladder. Octave transposition multiplies both `f` and `df` by the same factor [Shepard, 1982], so transposing up by `j` octaves divides every entry in the last column by `2^j` — **the ratio structure is preserved exactly, which is the whole point of transposition** — and the successor paper uses this table directly when specifying rendering speed [Curzi, 2026y,?].

Note that the maximum symbol rate implied at the bottom of the ladder, about 18 symbols per second, sits inside the 10 to 16 Hz receiver bandwidth of Section 5.6. **That is a coincidence of scale, not a mechanism, and it is stated as a coincidence.**

30.6 6 Addressing without editing

Section 5 located the problem exactly. An acoustic field in tissue is a broadcast: at one megahertz some 3.65 million cells share a single resolution voxel and every one receives the identical waveform. **Selectivity at cellular scale cannot come from the sound. It has to come from somewhere else.**

There is exactly one layer in biology with the required resolution, and it is not acoustic, not optical and not spatial. **It is sequence.** A twenty-nucleotide string is a forty-bit name, and forty bits is enough to name any position in a mammalian genome with room to spare. Molecular recognition of that name is base-paired, essentially digital, and specific to a resolution no field can approach.

This section is about the name, and only the name. Its subject is the guide spacer read as a routing header rather than as a molecule; the correspondence between that header and the addressing layer of an internet protocol, stated field by field; the demonstrated separability of *targeting* from *editing*; and the consequence that follows:

Nothing is delivered, therefore nothing is edited. The architecture is non-invasive by construction, not by restraint.

That is a claim about architecture rather than about ethics or intention. A system that contains no nuclease cannot cut; a system in which no protein and no nucleic acid is introduced into any cell cannot modify a genome, whatever its operator wishes. **The safety property is structural, and a structural safety property is the only kind worth trusting, because it does not depend on anyone behaving well.** The general principle — that a system's security should rest on its design rather than on the goodwill or ignorance of those who handle it — is the oldest rule in the field [Kerckhoffs, 1883, Saltzer and Schroeder, 1975], and the corpus states it as a structural property elsewhere [Curzi, 2026].

And the standing constraint applies here with full force: *there is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal.* **An addressing scheme is not a delivery mechanism, and this section does not supply one.**

30.6.1 6.1 What is established

Four statements here are established molecular biology, none of which depends on this framework and none of which is a postulate. **All four are cited without keys, per Section 3.7; the sources are named in Appendix A.**

```
E1  A GUIDE'S SPECIFICITY LIVES IN A SHORT VARIABLE SEGMENT.
    The spacer -- ~20 nt for SpCas9, ~23 nt for Cas12a -- determines which
    genomic sequence is recognised. The remainder of the guide is invariant
    scaffold.

E2  TARGETING AND EDITING ARE SEPARABLE FUNCTIONS.
    Catalytically dead Cas9 (D10A + H840A) binds its target and does not cut.
    Fused to effectors it represses or activates transcription. THIS IS THE
    PROOF THAT BINDING A LOCUS AND CHANGING A LOCUS ARE DIFFERENT OPERATIONS.

E3  Cas12a PROCESSES ITS OWN GUIDE ARRAY.
    Cas12a cleaves its own precursor CRISPR RNA array without requiring
    tracrRNA or host RNase III.
    CONSEQUENCE: ONE TRANSCRIPT CAN CARRY MANY ADDRESSES.

E4  A PROTOSPACER-ADJACENT MOTIF CONSTRAINS WHICH ADDRESSES ARE VALID AT ALL.
    SpCas9 requires an NGG motif; Cas12a requires a T-rich motif, canonically
    TTTV. A target lacking the correct motif is not engaged regardless of how
    well the spacer matches.
```

From E2 alone the central architectural claim follows: **addressing is a separable function. A system may name a locus with full specificity and do nothing to it.**

What this section does not claim, stated as a block because the honesty clauses are absolute and because this is where a careless reader is most likely to over-read:

```
N1  NO CLAIM THAT DNA IS EDITED, ALTERED, OR DAMAGED. The architecture described
    here contains no nuclease, introduces no protein, and introduces no nucleic
    acid into any cell. It cannot edit. Not "does not"; CANNOT.

N2  NO CLINICAL, MEDICAL OR THERAPEUTIC CLAIM. None. No trials, no controlled
    observation, no demonstrated physiological effect.

N3  NO CLAIM THAT AN ADDRESS DELIVERED ACOUSTICALLY IS RECEIVED, RECOGNISED, OR
    ACTED UPON.

N4  NO CLAIM THAT ANY CAPABILITY HAS BEEN CONFERRED ON ANY CELL.

N5  THE PROTOCOL CORRESPONDENCE IS A CORRESPONDENCE OF FORMAT, NOT AN ASSERTION
    THAT CELLS RUN A PROTOCOL STACK. Any extension beyond format is marked
    SPECULATIVE in the sentence that extends it.
```

30.6.2 6.2 The address, sized

A nucleotide is two bits. Therefore:

spacer length	address space `4^L`	bits
12 nt	1.678e7	24
16 nt	4.295e9	32
17 nt	1.718e10	34
18 nt	6.872e10	36
20 nt (SpCas9)	**1.0995e12**	**40**
23 nt (Cas12a)	**7.0369e13**	**46**
25 nt	1.126e15	50
32 nt	1.845e19	64

For comparison, computed: an IPv4 address is 32 bits (`4.2950e9` addresses) and an IPv6 address is 128 bits (`3.4028e38`). **A 20-nucleotide spacer is a 40-bit address — wider than IPv4 by 8 bits, a factor of 256** — and a 23-nucleotide spacer at 46 bits is 16,384 times the IPv4 space. **The comparison is not decorative; it is the right way to size the naming problem, and it immediately answers whether the address is long enough.**

How long does it need to be? The human haploid genome is about `3.1e9` base pairs. Counting start positions on both strands:

```
candidate positions = 2 * 3.1e9 = 6.200e9
bits to name one    = log2(6.200e9) = 32.5296 bits = 16.26 nt
```

Sixteen and a quarter nucleotides would suffice to name every position in the genome exactly once. The 20-nucleotide spacer over-provisions by 3.74 nucleotides — 7.47 bits, a factor of 177. That over-provision is not waste; it is error margin, and Section 6.5 computes exactly what it buys.

30.6.3 6.3 The addressing correspondence, stated field by field

The correspondence is one of *format and function*, and it is stated field by field so that a reader can check it rather than admire it. **It is a category-2 borrowing** (Section 3.4).

```
INTERNET PROTOCOL                    GUIDE FORMAT
---------------------------------    ---------------------------------
destination address                  SPACER (20 nt / 40 bits, SpCas9)
  names WHERE the packet goes          names WHICH LOCUS is recognised

address family (AF_INET / AF_INET6)  PAM (NGG for SpCas9, TTTV for Cas12a)
  a validity constraint on the         a validity constraint on the target:
  address form; a well-formed          a locus without the correct motif is
  address in the wrong family is       not engaged regardless of spacer
  not routable                         match

header framing / delimiters          DIRECT REPEAT (19-20 nt in Cas12a)
  marks packet boundaries              marks address boundaries within the
                                       array; the self-processing site

multicast destination list           MULTI-SPACER crRNA ARRAY
  one transmission, many               one transcript, many spacers, each
  recipients                           cleaved out to address a locus

payload                              ENCODED DATA (base-4)
  the content, distinct from the       distinct from the address
```

```
  address

routing table lookup                  BASE PAIRING + MOTIF RECOGNITION
  the physical act of matching          the physical act of matching
```

What the correspondence is not. It is not a claim that cells implement a protocol stack, that there is a transport layer, or that anything is acknowledged, retransmitted or checksummed, and it licenses no inference about behaviour not independently demonstrated. **Where this paper extends the correspondence past format — to talk about "delivery" or "reception" — the extension is speculative and is marked speculative in the sentence that makes it.** The value of the correspondence is that it supplies the right vocabulary for a real structural fact: **specificity is carried in a short, variable, separable field, and the rest of the construct is invariant. The internet protocol specifications themselves carry no citation key here** (Appendix A).

30.6.4 6.4 The motif as an address-family constraint

Of all the rows in that table, motif-as-address-family survives the hardest scrutiny, because the correspondence is functional rather than visual.

An IPv6 address presented to an IPv4-only socket is not routed: it is not that the destination is unreachable, but that the address is *not of a form the machinery accepts*. Similarly, a genomic sequence perfectly complementary to a spacer but not followed by the required motif is not engaged; the protein requires the motif to initiate DNA unwinding, and without it recognition does not proceed. **The constraint is on the form of the address, imposed by the machinery, independent of whether the address is otherwise correct. That is precisely what an address family is.**

The constraint is quantitatively severe:

```
P(NGG at a random position)   = 1/16          = 0.0625
P(TTTV at a random position)  = (1/4)^3 * 3/4 = 0.011719  = 1 in 85.33

targetable positions, SpCas9 NGG  : 3.8750e8 of 6.2000e9
targetable positions, Cas12a TTTV : 7.2656e7 of 6.2000e9

mean spacing between NGG sites  = 16.00 bp
mean spacing between TTTV sites = 85.33 bp
```

One position in sixteen is addressable under NGG; one in eighty-five under the canonical T-rich motif. Address-family choice is therefore a design decision with a hard cost: choosing the T-rich family reduces the reachable address space by a factor of 5.33 relative to NGG.

There is a genuine correspondence to the law of requisite variety [Ashby, 1956, 1958]: **only variety absorbs variety, and an address family is exactly a restriction on the variety of addresses the system will accept.** Restricting the family reduces what can be named. **That is not a defect of the biology; it is what makes the recognition step tractable**, and the same logic explains why routing tables are hierarchical.

Precedent, priced. The Confucian doctrine of *zhengming*, the rectification of names (*Analects* XIII.3) — if names are not correct, language is not in accordance with the truth of things, and affairs cannot be carried on to success — records that "the validity of the name is a precondition of the operation, prior to the operation's own correctness" was arrived at independently and preserved for two and a half millennia. **It supplies no mechanism and no evidence**, and every evidential claim here rests on the measured column [Curzi, 2026w]. Named without a key (Appendix A).

30.6.5 6.5 Routing collisions: what the extra bits buy

An address that names more than one destination is a routing collision. In this domain the collision has a name in the literature — off-target binding — and it can be computed exactly.

Expected number of exact matches of a random spacer in the human genome:

spacer length	expected exact genomic matches
12 nt	369.55
16 nt	1.4436
20 nt	5.6389e-3
23 nt	8.8107e-5

At 16 nt the expected number of matches is about 1.44 — the address is right at the edge of ambiguity. At 20 nt it is 0.0056, comfortably unique. That is what the four nucleotides between 16 and 20 bought: 8 bits, a factor of 256 in collision margin (`1.4436 / 5.6389e-3 = 256.0`).

But exact matching is not the operative condition, because base pairing tolerates mismatches. The honest calculation counts the *Hamming ball*: all sequences within `d` substitutions of the intended address. For a 20-mer over a four-letter alphabet the ball sizes are `C(20,d) 3^d`:

`d`	sequences at exactly distance `d`	cumulative ball size	expected genomic hits, motif-constrained
0	1	1	0.000352
1	60	61	0.021498
2	1,710	1,771	0.624152
3	30,780	32,551	11.471923
4	392,445	424,996	149.780999
5	3,767,472	4,192,468	1,477.548130

where the per-site expectation base is `6.2e9 x (1/16) / 4^20 = 3.5243e-4`.

A 20-nucleotide address that tolerates three mismatches has roughly eleven and a half expected matches in the human genome instead of one. That single number explains the entire off-target literature, explains why mismatch tolerance is the dominant engineering concern in the field, and demonstrates a general law of addressing worth stating in its own right:

An address system's usable specificity is set not by the length of the address but by the tolerance of the matching step. Adding bits to the name is cheap; tightening the matcher is the expensive and decisive move.

An internet engineer will recognise this as the reason longest-prefix matching and exact-match forwarding behave so differently. It is also Goodhart's law wearing molecular clothes [Goodhart, 1975, Campbell, 1979]: **a matcher scored on *finding* targets rather than on *not finding wrong ones* will drift toward tolerance.**

30.6.6 6.6 The self-processing array as a multicast destination list

This is the strongest single item in the section, and it is worth being precise about why.

Cas9 does not process its own guide array; maturation of its CRISPR RNA involves tracrRNA and host RNase III. **Cas12a does**: it is a single RNA-guided endonuclease requiring only a crRNA, and the same enzyme also cleaves the direct repeats within a multi-spacer array to liberate individual mature guides.

Read as a data format, that is a specific and recognisable object:

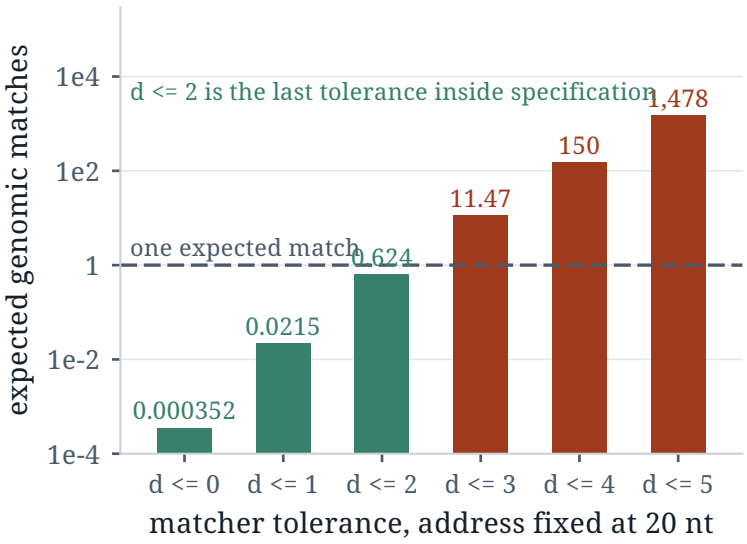


Figure 30.5: **Specificity is set by the matcher, not by the address.** Expected genomic matches for a fixed 20-nucleotide address as the matcher's mismatch tolerance rises, computed from the Hamming-ball sizes `C(20,d) 3^d` against a per-site base of `6.2e9 x (1/16) / 4^20`. The address length never changes; the expected collision count rises from 0.000352 at exact match to 11.47 at `d = 3`, a factor of 32,551 — exactly the ratio of the Hamming-ball sizes. The horizontal line at one expected match is the specification boundary, and the last tolerance below it is `d = 2` at 0.624 — which is why the worked block of Section 9.3 is declared valid only there. Computed from the closed form of Section 6.5; generated by `figs/src/CCP-030_fig05.py`.

```
5'--[DR]--[spacer 1]--[DR]--[spacer 2]--[DR]--[spacer 3]-- ... --3'

DR      = direct repeat: a delimiter, and the site at which the enzyme
          cleaves. FUNCTIONALLY A FRAME MARKER.
spacer  = one destination address.
```

One transcript. Many addresses. Self-delimiting. Processed by the same machine that uses the addresses. That is formally a packet with a multicast destination list, and the correspondence is structural rather than decorative because the framing is a real, enzymatically recognised feature of the molecule and not something an author drew on top of it.

The format arithmetic, computed with a **nominal** 19 nt direct repeat and a **nominal** 23 nt spacer:

```
per-address unit  = 19 + 23 = 42 nt
header overhead   = 19/42   = 45.24 %
address payload   = 46 bits per address
framing           = 38 bits per address
```

addresses in array	transcript length	address bits	framing bits
1	42 nt	46	38
2	84 nt	92	76
4	168 nt	184	152
8	336 nt	368	304
10	420 nt	460	380

addresses in array	transcript length	address bits	framing bits
20	840 nt	920	760

Forty-five per cent framing overhead is high by network standards — IPv4's minimum header is 20 bytes against a payload typically hundreds — **and the reason is instructive: this is a format optimised for reliable self-delimiting cleavage rather than for efficiency, which is the correct trade when a framing error is catastrophic and bandwidth is free.** Repeat lengths vary by ortholog; the 19 nt figure is nominal and used for scale.

For comparison, IPv4's multicast allocation is `224.0.0.0/4`, which is `2^28 = 268,435,456` group addresses. **A self-processing array is a different kind of multicast — an explicit destination *list* rather than a group *identifier* — and the distinction is worth keeping, because a list scales linearly in transcript length while a group identifier does not scale at all.** A 20-address array costs 840 nucleotides; a group address costs 28 bits. **Explicit lists are simple and expensive; group identifiers are cheap and require infrastructure. Biology chose the list, and biology has no routers.**

Precedent, priced. The Alexandrian *Pinakes* of Callimachus (3rd century BCE) was a catalogue physically separate from the scrolls, in which each work was named and located. **It records independent arrival at the structural insight this section needs — the catalogue is not the book; the address is a distinct object from the content, and a single object may carry many addresses. It is not evidence about molecules and supplies no mechanism** [Curzi, 2026w]. Named without a key (Appendix A).

30.6.7 6.7 The separability of targeting from editing

If a reader takes one result from this section, it should be this one.

The nuclease activity of Cas9 resides in two domains. Mutating one residue in each produces a protein that retains full RNA-guided DNA binding and has no cleavage activity. That dead nuclease was used to block transcription by steric interference, and the approach was later extended to genome-scale repression and activation by fusing effector domains, so that binding a promoter *increases* transcription.

Bind a locus. Change its expression. Cut nothing.

That sentence is the load-bearing result, and its architectural consequence is that **specificity of address and modification of target are independent functions, and the first can be present in full while the second is entirely absent. Addressing is not a weakened form of editing. It is a different operation, and it was shown to be a different operation by deleting the editing and finding that the addressing still worked.**

Two clarifications, both required.

First, those systems are still delivered. In every experiment concerned, a protein and a guide RNA were introduced into cells. **The architecture described here introduces neither**, and the result is cited as *proof of a principle* — that targeting and editing separate — and not as a description of what this architecture does.

Second, the honest limit: even a catalytically dead nuclease changes something. It occupies a locus, and occupancy has consequences — that is the entire point of transcriptional interference and activation. **It does not change *sequence*, which is the claim**, and the distinction between changing sequence and changing expression is exactly the distinction between this section and the editing literature. **Nothing in this architecture does either, because nothing is delivered.**

Precedent, priced, and this is where the gap between the columns is at its widest anywhere in the series. The "true name" doctrine in its widest form — the Egyptian narrative of Isis and the secret name of Ra, the Kabbalistic treatment of divine names, the insistence of the Solomonic grimoire tradition on exact naming as the operative variable — **records independent arrival at the idea that precision of the name, rather than force applied, determines**

whether an operation engages its target. It supplies no mechanism, no magnitude, and no evidence whatever about molecules. Presenting it as evidence would be the exact fault the two-column refinement forbids [Curzi, 2026w]. **Say so rather than smoothing it over.**

A footnote on wildcards, offered because it is genuinely apt. Roman ritual formulae included the hedge *sive deus sive dea* — "whether god or goddess" — used when the correct name of a locality's deity was not known. **That is, formally, a wildcard address issued under name-uncertainty, and the tradition recorded it as a recognised practice rather than an embarrassment.** It is precedent, it is not evidence, and its engineering lesson is real: **a system that must address under name-uncertainty needs an explicit underspecified-address form, and one that lacks it will get silent misrouting instead.**

30.6.8 6.8 The mitochondrial address family, corrected

A separate address family for mitochondrial targets is genuinely necessary. **The reason usually given in casual discussion is wrong, and correcting it is more useful than repeating it.**

The wrong reason: "mitochondrial DNA needs a smaller, more compact editor." Compactness is a real virtue of the hypercompact effectors — roughly 70 kDa, reported from huge phages — **but compactness is a *delivery* virtue.** Computed, with a nominal mean residue mass of 110 Da:

protein	residues	estimated mass	estimated CDS	fraction of a 4.7 kb AAV payload
SpCas9	1,368	~150.5 kDa	4,104 bp	87.3 %
AsCas12a	1,307	~143.8 kDa	3,921 bp	83.4 %
compact effector	~636 (est)	~70 kDa	~1,909 bp	40.6 %

The mass ratio is about 2.15, and the coding-sequence difference is what decides whether an editor and its regulatory elements fit in a viral vector. That is a real and important engineering fact — and it is irrelevant to this architecture, because this architecture delivers nothing and therefore has no packaging limit to fit inside. Citing compactness as a reason for a mitochondrial address family here would be borrowing a constraint the system does not have.

The right reason: mitochondria are not a reliable destination for guide RNA. RNA import into mammalian mitochondria has been studied for decades and remains contested; whatever the resolution, the field's operative position is that it is not reliable enough to build a targeting system on. **The decisive evidence is what the field did about it: it invented protein-only editors.** A split bacterial cytidine deaminase fused to TALE DNA-binding proteins was described explicitly as *CRISPR-free* mitochondrial base editing. **A field does not build an entirely separate, guide-RNA-free targeting chemistry unless the guide-RNA route is unavailable.** That is an argument from *engineering behaviour* rather than from a single measurement, and it is stated as one.

The addressing arithmetic closes the point:

```
human mtDNA (reference sequence)  = 16,569 bp
positions, both strands           = 33,138
bits to name one uniquely         = log2(33138) = 15.0162 bits = 7.51 nt
expected exact matches of a random 20-mer in mtDNA = 3.0139e-8
over-provision of a 20-nt address for mtDNA = 4^20 / 33138 = 3.3180e7

address bits actually required:
    nuclear        32.53 bits
    mitochondrial  15.02 bits
```

```
difference     17.51 bits
```

A 20-nucleotide address over-specifies a mitochondrial target by a factor of thirty-three million. Address length is not the mitochondrial problem — reachability is. In network terms this is a routing problem rather than a naming problem: the destination network is not reachable by the available transport, and no amount of address precision fixes an unreachable network. **The separate address family exists because the transport differs, which is exactly why address families exist in networking.**

30.6.9 6.9 Payload, and the reason there is not one here

Arbitrary digital data can be encoded in nucleotide sequence at the natural rate of two bits per base, and a volume of order 5.27 megabits has been stored and recovered. **That is the payload side of the correspondence and it is entirely real: base-4 encoding of arbitrary content in sequence is a demonstrated technology.** Named without a key (Appendix A).

The format is therefore complete in principle:

```
[ ADDRESS: spacer, 20–23 nt, 40–46 bits ]
[ FAMILY CONSTRAINT: motif, validity gate ]
[ FRAMING: direct repeat, self-delimiting ]
[ PAYLOAD: base-4 encoded content, arbitrary length ]
```

And here is where this paper must stop. A well-formed packet is not a delivered packet. The format above can be constructed on a computer and rendered as audio; Section 5 shows the acoustic channel has ample capacity to carry it and no ability to resolve a single cell. **What does not exist, anywhere in the literature, is the receiving side: there is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal.**

Nothing is delivered, therefore nothing is edited. Both halves of that sentence are the same fact stated twice. **The architecture's safety and the architecture's open problem are the *same* structural property, and any account claiming the first while quietly disowning the second is incoherent.** This paper claims both, in the same breath, and the successor paper examines the second without flinching [Curzi, 2026y].

30.6.10 6.10 The storage law at the addressing layer

A brief note on how this connects to the corpus's spine, **badged as the model-form claim it is** (Section 3.4).

In storage-law terms an addressing scheme is an `eta` device: it does not supply input `S`; it determines what fraction of a given input reaches a durable structure rather than dissipating [Curzi, 2026a]. **A broadcast with no addressing has, for any specific target, `eta` near zero — the input is spent on 3.65 million cells to reach one.** Addressing raises `eta` by concentrating the same `S` where it can convert.

That reading yields one checkable design consequence, computed with the surplus law's own arithmetic [Curzi, 2026]:

```
an address family restricts the reachable set.
SpCas9 NGG:  reachable fraction 1/16    ->  ln(6.2e9/16)     = 19.77 nats
Cas12a TTTV: reachable fraction 1/85.33 ->  ln(6.2e9/85.33)  = 18.10 nats
unrestricted:                               ln(6.2e9)        = 22.55 nats
```

Choosing the T-rich address family costs 4.4465 nats of nameable variety relative to unrestricted addressing, and 1.67 nats relative to NGG. In bits that is 6.415 — **the address family throws away about six and a half bits of the genome's nameable space in exchange for a**

tractable recognition step. That is the price of requisite variety at the recognition layer [Ashby, 1958], **and whether it is a good trade depends on whether the six bits contained anything needed, which is an empirical question about the target list and not a matter of preference.**

A note on which quantity this is. `ln(6.2e9)` is the logarithm of a count of positions, not an effective count of anything. **Where the corpus uses an effective count it uses `N_eff = 1 + (N-1)u`, which is a linear interpolation between 1 and `N`, is the framework's own construction, and is not a Hill number** [Curzi, 2026]; `1/HHI` and `exp(H)` are Hill numbers of orders two and one and are the grounded objects the literature supports [Hill, 1973, Jost, 2006]. **No effective count appears in this section**, and the distinction is stated so that a reader does not import one.

30.7 7 Specifications

Numbered, checkable requirements. **The three groups replace the source manuscript's chapter-keyed numbering: SP-R for the receiver, SP-C for the channel, SP-A for addressing.**

30.7.1 7.1 SP-R — the receiver

SP-R1. Mechanotransduction is cited as fact, not as analogy. Established physiology may be stated flatly; do not hedge it and do not decorate it.

SP-R2. The standing constraint appears in every document citing SP-R1, verbatim: *there is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal.* It appears in the same document, not in a linked appendix.

SP-R3. Receiver and decoder are never conflated. "Receives force" and "receives a message" are different claims with different evidence. Any sentence containing both must be rewritten.

SP-R4. Sonogenetics is described as **receiver-side selectivity**: state that specificity derives from which channel is expressed, and state the reported reliance on microbubble amplification. Do not describe it as acoustic messaging.

SP-R5. The cochlea is cited as an existence proof of buildability only. It proves that biology *can* build a fine mechanical receiver and licenses no inference about ordinary cells in ordinary tissue.

SP-R6. Mechanical constants carry their nominal flag at first use in any document: sound speed, attenuation, tissue density, eardrum area, hair-cell counts, and all channel kinetic time constants (Section 3.3).

SP-R7. Energy-scale arguments are stated in units of kT. `kT = 4.280e-21` J `= 4.280` pN·nm at 310 K. **Any claimed mechanical effect must be shown against that scale or declared unquantified in the same sentence.**

SP-R8. The uniform-pressure limit is disclosed. Any claim of acoustic action on cells states the `k d` fraction at the frequency used and names the non-uniform mechanism it relies on — gradient, shear, streaming, radiation force, or cavitation.

SP-R9. No clinical, medical or therapeutic claim. Absolute.

SP-R10. No claim that genomic sequence is edited, altered or damaged. Absolute.

SP-R11. Concordance blocks state column roles explicitly. **A block without the roles stated is void.**

SP-R12. Uncertain figures are omitted, not guessed, and the omission is stated.

30.7.2 7.2 SP-C — the channel

SP-C1. Every capacity claim states `B`, `S/N` and the noise model. A capacity figure without all three is void, and the noise model's validity regime is stated with it.

SP-C2. Bandwidth is the receiver's, not the transmitter's. Use the corner frequency of the slowest stage that must resolve the symbol.

SP-C3. Nominal constants are flagged at first use, and **sensitivity to the attenuation exponent is shown, not assumed away** (Section 5.4).

SP-C4. Spatial resolution is reported as cells per voxel, not as wavelength. "`3.65e6` cells share one voxel at 1 MHz" is the honest statement; "1.54 mm resolution" conceals it.

SP-C5. The broadcast property is disclosed. Any design using an acoustic field at or below about 10 MHz states that all cells within a voxel receive identical signal, and names the non-acoustic mechanism supplying selectivity or states that there is none.

SP-C6. Capacity is never presented as mechanism. Every capacity result is accompanied in the same document by: *capacity is an upper bound on what a perfect decoder could recover; it does not imply that any decoder exists.*

SP-C7. The standing constraint appears verbatim in every document computing a capacity into tissue.

SP-C8. Gabor is applied to every timed spectral design. Any specification of symbol rate and frequency spacing shows the product `dt df` and shows it satisfies `>= 0.079577`.

SP-C9. The 20.03 dB crossover is the design point. Below it, buy signal; above it, buy bandwidth. Power spent above the crossover to increase capacity is wasted.

SP-C10. Logarithmic returns are stated where they bite. A 10,000-fold pressure increase from 100 Pa to 1 MPa buys a factor of 2.80 in capacity. **Never promise linear gain in a quantity that grows logarithmically.**

SP-C11. No clinical, medical or therapeutic claim. **The pressures tabulated for capacity purposes are not recommendations, exposure guidance, or safety statements of any kind.**

SP-C12. No claim that genomic sequence is edited or altered. This section concerns transmission only.

30.7.3 7.3 SP-A — the addressing layer

SP-A1. The address is a header, never a molecule, in this architecture. Text must not describe the spacer as something present in, delivered to, or synthesised by a cell. It is a name computed on a computer.

SP-A2. Nothing is delivered; state it as a structural property. "No protein and no nucleic acid is introduced into any cell; the architecture therefore cannot edit a genome." **Cannot, not does not.**

SP-A3. No claim that genomic sequence is edited, altered or damaged. Absolute.

SP-A4. No clinical, medical or therapeutic claim. Absolute.

SP-A5. The standing constraint appears verbatim in every document using this format.

SP-A6. Every address publishes its collision analysis. Required fields: spacer length; address space `4^L`; expected exact genomic matches; Hamming ball sizes to `d = 3`; expected hits at `d = 3` under the applicable motif constraint. **An address without this analysis is not specified, it is merely written down.**

SP-A7. Every address declares its address family, naming the motif, its per-position probability, and the resulting reachable-site count. NGG `= 1/16`; TTTV `= 1/85.33`.

SP-A8. The separability result is cited for separability only, with the delivery caveat: every citation states that those systems were delivered into cells and that this architecture delivers nothing.

SP-A9. The mitochondrial family is justified by reachability, not compactness. Compactness is a delivery virtue and is irrelevant where nothing is delivered.

SP-A10. The protocol correspondence is scoped in the same paragraph that uses it. It is a correspondence of format and function; any extension beyond format carries the word *speculative* in the sentence that extends it.

SP-A11. Concordance blocks state column roles. **Permitted verbs for the precedent column: records, preserves, arrived at independently. Forbidden: confirms, corroborates, verifies, proves.**

SP-A12. Numerical coincidences are labelled as such or omitted. `4^18 = 2^36`, so an 18-nucleotide address is exactly 36 bits, the same number as an octave transposition used in genomic audio. **This coincidence carries no content**; it is an artefact of both systems using powers of two, and any text implying otherwise is void.

30.8 8 Failure modes

Each fault carries a clinical signature, a diagnostic, and a fix.

30.8.1 8.1 FM-R — faults of the receiver argument

FM-R.1 · The one-sentence slide. *Signature:* a paragraph establishes mechanotransduction correctly and the next sentence begins "and therefore cells can receive information from sound." **The word "information" has changed meaning between the sentences: in the first it meant a physical influence, in the second a chosen symbol.** *Fix:* ban the unqualified word at the join. Write "physical influence" or "symbolic content" and the slide becomes visible immediately, because "cells can receive symbolic content from sound" is obviously an unsupported claim in a way that the equivocal version is not.

FM-R.2 · The cochlea generalisation. *Signature:* hair-cell sensitivity figures — 0.3 nm displacement, sub-kT energies — quoted in support of a claim about ordinary cells. *Fix:* every citation of cochlear performance is immediately followed by the enumerated special conditions — impedance-matched middle ear at 27 dB nominal gain, active amplification, dedicated afferent innervation, isolated fluid-filled cavity, purpose-built gating apparatus — **and by the statement that a cell in bulk tissue has none of them.**

FM-R.3 · The uniform-pressure omission. *Signature:* a design claims mechanical action on cells from an audible-band field and never computes `k d`. *Diagnostic:* at 20 kHz, `k d` for a 10 µm cell is `8.16e-4` — the pressure differential across the cell is under a tenth of a per cent of peak. **If the mechanism is "pressure squeezes the cell", the design is describing an effect roughly a thousand times smaller than it thinks.** *Fix:* SP-R8, then size the named non-uniform mechanism against the 1.07 pN / 4 nm gating scale.

FM-R.4 · The energy-free claim. *Signature:* a proposed effect on a cell stated with no energy or force figure at all. *Why it is fatal:* cells sit in a bath at 310 K where every degree of freedom carries 4.28 pN·nm. **An effect that cannot be sized against kT is not an effect; it is a hope.** *Fix:* SP-R7.

FM-R.5 · Precedent laundered as evidence. *Signature:* a precedent entry described with verbs like "confirms", "corroborates" or "independently verifies", or two columns summed into "multiple lines of evidence". *Fix:* SP-R11, and a test with teeth — **delete the precedent row and ask whether the argument's *evidential* standing changed. If it did, the argument was resting on precedent** [Curzi, 2026w].

FM-R.6 · The sonogenetics overread. *Signature:* sonogenetics cited as "controlling cells with sound", with the transgene and microbubble requirements dropped. *Fix:* SP-R4. **The honest one-line summary is: ultrasound can turn on a channel you put there.**

FM-R.7 · The permanence assumption. *Signature:* a protocol assumes a single exposure produces a lasting cellular change, with no statement of the decay rate of the state variable. *Diagnostic:* this is an unstated assumption that `delta` is near zero. At `delta = 0.3` per period the half-life is 1.94 periods; at `delta = 0.02` it is 34.3. **Three orders of magnitude of protocol design hang on a parameter nobody measured.** *Fix:* name the state variable, cite or measure its turnover, and state the implied `delta` and half-life before any durability claim.

FM-R.8 · The prize borrow. *Signature:* the Nobel Prize invoked in a context where the claim being supported is not the claim the prize was awarded for. *Why it is corrosive:* it is the strongest available citation and misusing it costs the document its credibility with exactly the readers who can check. *Fix:* **state the citation with its scope attached, every time — the prize was awarded for the discovery of receptors for temperature and touch. It supports E1 and it supports nothing about symbolic content.**

FM-R.9 · The fabricated figure. *Signature:* a specific published number appears without the author having verified it, because a specific number reads better than a qualitative statement. *Fix:* SP-R12. **Section 4.5 does exactly that with a drive frequency, and the text is stronger for the omission than it would be for a guess.**

30.8.2 8.2 FM-C — faults of the channel argument

FM-C.1 · Capacity mistaken for mechanism. *Signature:* a document computes a favourable capacity and concludes that transmission occurs; the tell is the word "therefore" immediately after a capacity figure. **Shannon's theorem is an existence proof about codes, not about hardware: its converse forbids exceeding C, and nothing in it *permits* reaching C in any particular physical system.** *Fix:* SP-C6, applied literally — paste the disclaimer into the same paragraph as the number. **If the argument collapses when the sentence is added, the argument was the fault.**

FM-C.2 · Transmitter bandwidth substituted for receiver bandwidth. *Signature:* a 1 MHz carrier is quoted and `B = 1e6` goes into the capacity formula. *Diagnostic:* compare the two computed tables of Section 5.7. At 20 dB the receiver-limited figure at mechanotransduction kinetics is 70.65 bit/s; **at a 1 kHz medium-set bandwidth it is 6,658 bit/s, a factor of 94.2, and if the full 1 MHz carrier is substituted it is `6.66e6` bit/s, a factor of `9.4e4`.** *Fix:* SP-C2. **The source manuscript states this comparison with a different figure, and the discrepancy is reported in Section 9.4 rather than silently harmonised.**

FM-C.3 · The resolution mirage. *Signature:* a design proposes to address individual cells acoustically and quotes a wavelength without quoting a depth. *Diagnostic:* the `lambda d20` product falls as `1/f^2`; at 154 MHz the 20 dB depth is 2.6 mm on the most generous attenuation model and 0.21 mm on an `f^1.5` model. *Fix:* **always publish the pair — resolution *and* reachable depth, with the attenuation exponent stated.**

FM-C.4 · The undisclosed broadcast. *Signature:* a protocol describes acoustic delivery "to the target cells" with no statement of how non-target cells in the same voxel are excluded. *Diagnostic:* at 1 MHz there are `3.65e6` cells in the voxel. **If the protocol has no receiver-side selectivity, its true addressee count is the voxel population.** *Fix:* SP-C5.

FM-C.5 · The noise model borrowed from the wrong regime. *Signature:* `p_rms = sqrt(4 kT rho c B / A)` used at cellular apertures without stating that it assumes `k a >> 1`. *Why it matters in both directions:* used naively it overstates the noise and makes a design look worse than it is; **corrected naively, by shrinking the radiation resistance without also shrinking the signal coupling, it makes the design look far better than it is — and the second error is the dangerous one.** *Fix:* SP-C1; where a rigorous impedance model is unavailable, say so and carry both estimates, as Section 5.5 does.

FM-C.6 · The power escalation. *Signature:* a design responds to insufficient capacity by proposing higher amplitude. *Diagnostic:* from 100 Pa to 1 MPa — four orders of magnitude in pressure, eight in intensity — capacity rises from 14,792 to 41,368 bit/s, a factor of 2.80. **Above the crossover the return is logarithmic and the cost is quadratic in pressure.** *Fix:* SP-C9.

FM-C.7 · The Gabor violation. *Signature:* a specification lists both a symbol rate and a frequency spacing whose product is below 0.079577. *Diagnostic:* **the specification is not merely ambitious; it is describing something that cannot exist.** Two tones separated by `df` cannot be distinguished in a window shorter than `1/(4 pi df)`, on any equipment, ever. *Fix:* compute `dt df` for every pair; Section 5.12 shows the worked case, 54.9 ms minimum at index 1.

FM-C.8 · The quantum import. *Signature:* the Gabor limit described as a quantum effect, or `hbar` appearing in an acoustic calculation. **Gabor's relation is classical Fourier duality and**

contains no quantum content at all [Gabor, 1946, Curzi, 2026z]. *Why it is fatal to credibility:* importing quantum language into an acoustics problem signals to any physicist reading that the rest of the document should not be trusted. *Fix:* state the relation as `dt df >= 1/(4 pi)`, note that it is the same Fourier mathematics as the position–momentum relation without the quantum content, and stop.

FM-C.9 · The unstated nominal. *Signature:* 0.5 dB/cm/MHz appears without the word "nominal" and without a sensitivity check. *Fix:* SP-C3 and the sensitivity table of Section 5.4. **Every plausible exponent makes the depth conclusion worse, which means the conclusions are robust to the assumption in the direction that matters.**

30.8.3 8.3 FM-A — faults of the addressing argument

FM-A.1 · The address treated as an agent. *Signature:* a sentence in which the spacer *does* something — "the guide finds the locus" — in a context where no guide molecule exists. *Fix:* SP-A1. **Permitted verbs: names, designates, specifies. Forbidden: finds, seeks, binds, targets, used transitively of a physical event.**

FM-A.2 · Editing vocabulary imported with the citation. *Signature:* sources are cited and the surrounding prose acquires "editing", "knock-in", "correction" or "repair". *Why it is fatal:* **it converts an architecture whose entire claim is *nothing is edited* into a document that appears to claim editing. A single stray verb can invert the safety property in a reader's mind.** *Fix:* SP-A3, plus a mechanical check — grep the document for *edit, cut, cleave, knock, correct, repair*, and confirm each occurrence describes the cited literature and is explicitly excluded from the architecture.

FM-A.3 · The collision analysis omitted. *Signature:* an address published as a bare 20-mer. *Diagnostic:* without the Hamming analysis the address's real specificity is unknown — 0.000352 expected hits at `d = 0` against 11.47 at `d = 3`. **The address is either essentially unique or thirty-two-thousand-fold ambiguous depending entirely on a matcher tolerance that was never stated.** *Fix:* SP-A6, in full, for every published address.

FM-A.4 · The motif forgotten. *Signature:* an address designed without its address family stated. *Diagnostic:* fifteen positions in sixteen have no NGG. **A well-formed address to a non-addressable position is the exact molecular analogue of an IPv6 literal handed to an IPv4 socket: not unreachable, but malformed.** *Fix:* SP-A7.

FM-A.5 · Compactness cited as the mitochondrial reason. *Signature:* effector mass offered as the justification for a separate mitochondrial address family. **Compactness matters for packaging into a delivery vector; this architecture has no delivery vector, so it has no packaging limit, so the virtue is real and inapplicable.** *Fix:* SP-A9.

FM-A.6 · The protocol-stack overreach. *Signature:* the correspondence extended to acknowledgements, retransmission, checksums, sessions or congestion control. *Diagnostic:* ask which molecule performs the function; **if the answer is a shrug, the sentence is an analogy that has escaped its enclosure.** *Fix:* SP-A10.

FM-A.7 · Precedent laundered as evidence. *Signature:* the true-name doctrine, *zhengming* or the Alexandrian catalogue described with verbs like "confirms" or "anticipates the discovery of". *Why this section is especially exposed:* **the precedents here are unusually apt, and aptness is exactly what makes the laundering tempting.** *Fix:* SP-A11, and the deletion test of FM-R.5.

FM-A.8 · The numerological slide. *Signature:* a numerical coincidence presented as though it carried meaning; the canonical example is `4^18 = 2^36`. *Why it is a coincidence and nothing else:* both systems are built on powers of two, so coincidences of this shape are guaranteed and carry no information about either. **A framework that cannot say "this is a coincidence" out loud will eventually be unable to say anything.** *Fix:* SP-A12.

FM-A.9 · The silent wildcard. *Signature:* a system permits underspecified addresses without an explicit wildcard form, so ambiguity is expressed as tolerance rather than as a declared wildcard. *Why it matters:* **the failure is silent misrouting, the worst class of addressing failure**

because nothing reports it. *Fix:* require every address to be either fully specified or explicitly marked as a declared wildcard with its expected match count computed under SP-A6.

30.9 9 Results: three specifications worked end to end

Three complete specifications are worked here against the requirements of Section 7. **Two of them fail, and both failures are reported in the same plain language a success would get.** None of the three involves a laboratory, a subject, or a cell: each is arithmetic on a proposed design, and the arithmetic is the result.

30.9.1 9.1 An audible-band installation, evaluated and relocated

The proposal. A treatment room in which a subject is exposed to an audible-band programme at 85 dB SPL delivered through air by loudspeakers. The proposal document contains the sentence *"mechanotransduction (Nobel 2021) provides the pathway by which the programme reaches the cell."*

Step 1 — what arrives at the tissue surface. 85 dB SPL in air is a pressure amplitude of `p = 20e-6 x 10^(85/20) = 0.3557` Pa. Air and soft tissue are grossly mismatched:

```
Z_air    = 1.204 * 343  =     413.0 Pa.s/m
Z_tissue = 1050  * 1540 = 1.617e6   Pa.s/m      ratio = 3915.5
T_I = 4 Z1 Z2 / (Z1 + Z2)^2 = 1.021e-3 = -29.9 dB
```

Ninety-nine point nine per cent of the incident intensity is reflected at the skin. The pressure transmission coefficient is `2 Z2/(Z1+Z2) = 1.9995`, so the pressure amplitude just inside the tissue is essentially doubled to about 0.711 Pa while the intensity collapses. **Both facts are correct and they are not in tension; they are what an impedance step does.** For mechanotransduction the relevant quantity is force, so the pressure figure is carried forward, **and it is the more generous of the two.**

Step 2 — what that pressure does to a cell. Taking the programme's dominant energy at 1 kHz, `lambda = 1540/1000 = 1.54` m and `k d = 2 pi x 10e-6 / 1.54 = 4.08e-5`, so the pressure differential across the cell is about 0.004 per cent of 0.711 Pa, or `dp = 2.90e-5` Pa. At 0.711 Pa the intensity is `p^2/(2 rho c) = 1.56e-7` W/m²; the attenuation coefficient at 1 kHz is a thousandfold smaller than the 1 MHz value used in Section 4.6 under the nominal linear law, so the absorbed power in a 523.6 µm³ cell is negligible and the resulting radiation force is many orders of magnitude below a femtonewton. **Against the gating scale — 1.07 pN to do 1 kT of work over a 4 nm swing — the proposal's stated mechanism is short of its own requirement by a margin too large to close by any adjustment of level within the audible band.**

Step 3 — the thermal floor. At a 10 µm aperture and a 1 kHz receiver bandwidth, `p_noise = sqrt(4 kT rho c B / A) = 0.594` Pa. **The signal available at the cell, 0.711 Pa, is of the same order as the thermal noise at that aperture — an SNR of about 1.6 dB.** The programme is not below the floor, **but it is not meaningfully above it either, and nothing about a barely-above-noise scalar pressure supports the transmission of chosen content.**

Step 4 — the verdict. The proposal fails **SP-R3** (receiver/decoder conflation), **SP-R8** (uniform-pressure disclosure), **SP-R7** (no energy-scale statement) and **SP-R2** (standing constraint absent), and commits **FM-R.1** and **FM-R.8**. The rewritten sentence that passes:

> *Mechanically gated ion channels establish that pressure is a real input to cellular signalling; the 2021 Nobel Prize was awarded for the discovery of receptors for temperature and touch. At the levels and frequencies used in this installation, the computed pressure differential across a 10 µm cell is `4.08e-5` of peak and the radiation force is far below the ~1 pN scale of channel gating; we therefore make no claim of direct mechanical action on cells. There is no established mechanism*

by which a cell demodulates sequence-specific information from an acoustic signal. The installation's claimed effects are auditory and psychological, and they are stated and tested as such.

That rewrite loses nothing the proposal could actually support, and it gains the only thing that matters: a reader who checks the arithmetic will find it correct.

Step 5 — what the exercise establishes. The installation is not condemned; it is *relocated*. Its effects, if any, belong to the chain running ear → auditory system → expectation → physiology, which is a real chain with an identified channel, and which is exactly what the corpus's own belief postulate means by *grounded in truth*: **belief acts through a channel, and the channel must exist — the postulate's own forbidding clause is that zero channel gain implies zero effect for any strength of belief** [Curzi, 2026ab]. **The cell is a receiver. In this installation it is not the receiver being used. Saying so costs the proposal nothing it was entitled to and saves it from a claim it could not defend.**

30.9.2 9.2 An acoustic link, evaluated and corrected by a factor of 885

The proposal. An acoustic link intended to deliver a 20-nucleotide address — 40 bits — to cells in a 1 cm³ region at 4 cm depth, specifying a 5 MHz carrier, a peak pressure of 500 kPa at focus, and a symbol scheme of 16 tones spaced 40 kHz apart, each held for 10 µs. It claims a delivery time of 100 µs.

Step 1 — does the carrier reach the depth? Under the nominal linear law `alpha = 0.5 x 5 = 2.5` dB/cm, one-way loss at 4 cm is 10.0 dB. Under `f^1.1` the coefficient is 2.937 dB/cm and the loss 11.75 dB; under `f^1.5` it is 5.59 dB/cm and 22.36 dB. **Verdict: the depth is reachable, with a loss between 10.0 and 22.4 dB depending on the exponent. The proposal must state which exponent it assumed. It did not** — FM-C.9.

Step 2 — what is the spatial resolution, in cells?

```
lambda   = 1540 / 5e6 = 308 um
lambda^3 = 2.92e-2 mm^3
cells per voxel at 1e9 cells/cm^3 = 2.92e4 cells
```

Twenty-nine thousand cells share one resolution element. The 1 cm³ region contains `1000 / 2.92e-2 = 3.42e4` voxels, so the design can distinguish about thirty-four thousand locations, each containing about twenty-nine thousand cells. **Verdict: the delivery is a broadcast to blocks of `2.9e4` cells, and the proposal did not disclose it** — FM-C.4, SP-C5.

Step 3 — does the symbol scheme satisfy Gabor?

```
dt = 10 us,  df = 40 kHz
dt * df = 1e-5 * 4e4 = 0.400      required: >= 0.079577
0.400 >= 0.0796  ->  PASSES, with a margin of 5.03x
```

Verdict: the symbol scheme is physically realisable, and this is the part of the proposal that was done correctly — which should be said. At the Gabor optimum for the 640 kHz occupied bandwidth, `R_max = 6.669448 x 6.4e5 = 4.268e6` bit/s against the proposal's actual `log2(16)/10 µs = 4.0e5` bit/s: **the scheme runs at about 9.4 per cent of its Gabor ceiling, which is conservative and defensible.**

Step 4 — what is the actual capacity, at the receiver? The proposal's 100 µs for 40 bits implies `4.0e5` bit/s, and that figure uses the *transmitter's* symbol rate. Applying SP-C2, the receiving stage is mechanotransduction at a nominal `tau = 15` ms, so `f_c = 1/(2 pi x 0.015) = 10.61` Hz. At 500 kPa focal pressure attenuated 10 dB to 4 cm, the pressure at target is `5e5 x 10^(-10/20) = 1.5811e5` Pa; against aperture thermal noise at a 10 µm receiver in a 10.61 Hz band, `p_noise = 6.115e-2` Pa, the SNR is `20 log10(1.5811e5 / 6.115e-2) = 128.25` dB. **Capacity is logarithmic in SNR, so:**

```
C = 10.61 * log2(1 + 10^12.825) = 10.61 * 42.60 = 452.04 bit/s
40 bits in 88.5 ms
```

Verdict: the proposal is wrong by a factor of 885 in delivery time — 88.5 ms, not 100 µs. The error is entirely FM-C.2. **Note also that an SNR of 128 dB bought only a factor of about 6 in capacity over the 20 dB case, exactly as SP-C10 warns.**

Step 5 — the fatal objection, which is none of the above. Every number so far is a design correction, and the proposal can be rewritten to state its exponent, disclose its broadcast property, and correct its delivery time by three orders of magnitude. **None of that saves it. There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal.** The 16-tone scheme presupposes a receiver that can distinguish sixteen frequencies and map them onto a symbol alphabet; **nothing in the mechanotransduction literature identifies such a receiver.** The cell's transducer is a low-pass amplitude device with no demonstrated frequency-discriminating front-end at the required selectivity and no demonstrated mapping from any spectral feature to a nucleotide symbol.

The capacity calculation was worth doing anyway, and this is the point of the whole analysis. It establishes that the failure is not one of bandwidth (452 bit/s is ample), not one of energy (the Landauer bound permits `2.05e14` bit/s at that pressure), and not one of thermodynamics. **The failure is located precisely: the receiver has no demodulator. That is a well-specified, falsifiable gap** [Curzi, 2026y], **and it is a far more useful thing to know than a vague sense that the idea is difficult.**

Step 6 — the rewritten specification.

```
ACOUSTIC LINK, HONEST SPECIFICATION
carrier             5 MHz
attenuation model   alpha = 0.5 f^1.0 dB/cm/MHz, NOMINAL; sensitivity to
                    b = 1.1 and b = 1.5 tabulated: 10.0 / 11.75 / 22.36 dB
                    one-way at 4 cm
spatial resolution  lambda = 308 um; lambda^3 = 2.92e-2 mm^3;
                    2.92e4 cells per voxel at 1e9 cells/cm^3 NOMINAL
delivery mode       BROADCAST to voxel populations. No acoustic mechanism
                    supplies cell-level selectivity. Selectivity, if any,
                    must be receiver-side.
symbol scheme       16 tones, 40 kHz spacing, 10 us hold.
                    dt.df = 0.400 >= 0.079577. Gabor PASSES, margin 5.03x.
                    Rate is 9.4% of the Gabor ceiling for the occupied band.
channel capacity    Receiver-limited. B = 10.61 Hz (NOMINAL kinetics).
                    C = 452 bit/s. 40 bits in 88.5 ms.
energy              Not limiting. Landauer bound at this pressure permits
                    ~2.05e14 bit/s; the shortfall is 11 orders of magnitude
                    in the permissive direction.
DECODER             NONE IDENTIFIED. Capacity is an upper bound on what a
                    perfect decoder could recover; it does not imply that any
                    decoder exists. There is no established mechanism by which
                    a cell demodulates sequence-specific information from an
                    acoustic signal.
claims made         Transmission arithmetic only. No clinical, medical or
                    therapeutic claim. No claim that DNA is edited or altered.
```

That specification is publishable. The original was not — and the difference is entirely a matter of stating what was computed and what was not.

30.9.3 9.3 An address block, specified completely and inert by construction

The problem. Specify a complete address block for a hypothetical target list of four nuclear loci and one mitochondrial locus, using the self-processing array format, and state exactly what has and has not been done. **This exercise builds a *format*. No molecule is synthesised, nothing is delivered, and no biological claim of any kind is made.**

Step 1 — choose the address family and pay for it. Canonical T-rich motif:

```
P(TTTV) = (1/4)^3 * (3/4) = 0.011719 = 1 in 85.33
reachable positions = 6.2000e9 * 0.011719 = 7.2656e7
mean spacing between addressable sites = 85.33 bp

ln(6.2e9)         = 22.5478 nats     unrestricted naming
ln(6.2e9 / 85.33) = 18.1013 nats     under TTTV
cost              =  4.4465 nats = 6.415 bits
```

Choosing this family throws away 6.415 bits of nameable genome in exchange for a recognition step that works. Stated, paid, moving on.

Step 2 — size the addresses. At a nominal 23 nt spacer: address space `4^23 = 7.0369e13`, address width 46 bits, required for genome uniqueness 32.53 bits, **over-provision 13.47 bits, a factor of `1.135e4`.**

Step 3 — collision analysis, per SP-A6. Computed for the 20-mer case and reported as the conservative bound, **since a 23 nt address is strictly better and a specification should not claim more margin than it has verified.** The per-site base under the NGG constraint is `6.2e9 x (1/16) / 4^20 = 3.5243e-4`:

matcher tolerance	ball size	expected genomic hits
exact (`d = 0`)	1	0.000352
`d <= 1`	61	0.021498
`d <= 2`	1,771	0.624152
`d <= 3`	32,551	11.471923
`d <= 4`	424,996	149.780999

Specification decision: the block is declared valid only for a matcher with tolerance `d <= 2`, where the expected collision count is 0.624 — under one. At `d = 3` the block would name eleven and a half places on average, **which is not an address; it is a neighbourhood.**

Note what just happened: the specification was constrained by the matcher, not by the address. That is Section 6.5's law in operation.

Step 4 — build the array as a multicast packet. Four nuclear addresses at a nominal 19 nt direct repeat:

```
per-address unit = 19 + 23 = 42 nt
total transcript = 4 * 42  = 168 nt
address bits     = 4 * 46  = 184 bits
framing bits     = 4 * 38  = 152 bits
header overhead  = 152 / (184 + 152) = 45.24 %
```

One hundred and sixty-eight nucleotides carries four fully specified destinations, self-delimited, in a format the enzyme itself is documented to parse.

Step 5 — the mitochondrial address, handled honestly.

```
mtDNA positions, both strands = 33,138
bits required for uniqueness  = 15.02 bits = 7.51 nt
```

```
a 23-nt address provides      = 46 bits
over-provision                = 30.98 bits = factor 2.1e9
```

The address is over-provisioned by a factor of two billion. Naming is not the problem, and no change of address length touches the actual problem, which is that the destination is not reachable by the transport this format assumes (Section 6.8). **The mitochondrial target is therefore placed in a separate address family with its own declared transport**, and is not merged into the nuclear array, because merging would produce an address that is well-formed and unroutable. **Compactness is not used as the justification** — SP-A9.

Step 6 — the complete block.

```
ADDRESS BLOCK, HONEST SPECIFICATION

format             self-processing crRNA array
address family     TTTV. P = 0.011719 per position. Reachable sites
                   7.2656e7 of 6.2000e9. Variety cost 4.4465 nats
                   (6.415 bits) relative to unrestricted naming.
address width      23 nt = 46 bits. Space 4^23 = 7.0369e13.
                   Genome requires 32.53 bits; over-provision 13.47 bits.
matcher tolerance  VALID ONLY FOR d <= 2. Expected collisions:
                   d=0: 0.000352 | d<=1: 0.0215 | d<=2: 0.624
                   d<=3: 11.47 -- OUT OF SPECIFICATION.
array              4 nuclear addresses. 168 nt. 184 address bits,
                   152 framing bits, 45.24 % header overhead.
5th target         MITOCHONDRIAL. SEPARATE ADDRESS FAMILY, separate declared
                   transport. Reason: reachability, NOT compactness.
payload            Base-4 encoded content at 2 bits/nt. Length arbitrary.

WHAT HAS BEEN DONE A format has been specified and its collision, family and
                   framing arithmetic computed.
WHAT HAS NOT       Nothing synthesised. Nothing delivered. No protein and no
                   nucleic acid introduced into any cell. NOTHING IS EDITED,
                   AND THE ARCHITECTURE CANNOT EDIT.
OPEN PROBLEM       There is no established mechanism by which a cell
                   demodulates sequence-specific information from an
                   acoustic signal.
CLAIMS             None clinical, medical or therapeutic. None biological.
                   This is a routing specification.
```

Step 7 — what the exercise establishes. The addressing layer is **fully specifiable, fully computable, and completely inert.** Every number above was computed; none required a laboratory; none touched a cell.

That inertness is the point. A design whose addressing layer is a pure computation is a design whose addressing layer cannot harm anyone, cannot be misused into editing, and can be published, audited and criticised in the open by anyone with a computer. **An architecture that delivers nothing has not removed anyone's ability to decline, because it has not reached them at all** — which is simultaneously its strongest safety property and its most serious unsolved problem, **and the two are the same fact.** The corpus's standing refusal to move a harm through a discovered gap in a scheme, and its concession that the case where a feedback loop runs through another person's will is an unsettled seam, bound what may be claimed here [Curzi, 2026].

30.9.4 9.4 A discrepancy in the source, reported rather than harmonised

One figure in the source manuscript does not reproduce, and it is reported here rather than quietly corrected. The source's diagnostic for the transmitter-bandwidth fault states that "at

20 dB SNR the medium-limited figure at 1 kHz bandwidth is `1.06e4` bit/s" against a receiver-limited 70.65 bit/s, "an inflation factor of about 150".

Recomputing: `C = B log2(1 + 100)` with `log2(101) = 6.6582` gives `1000 x 6.6582 = 6,658` bit/s at a 1 kHz bandwidth, not `1.06e4`. The figure `1.060e4` is the *fast mechanotransduction* row of the receiver table at `B = 1591.6` Hz — `1591.6 x 6.6582 = 10,597` — and `10,597 / 70.65 = 150.0` is exactly the source's stated ratio. **The arithmetic is correct and the label is wrong: the comparison is between two receiver stages, one of which is described as the medium.**

The corrected figures, which this paper uses throughout:

```
receiver-limited, tau = 15 ms   B =   10.61 Hz   C =     70.65 bit/s
receiver-limited, tau = 100 us  B = 1591.6  Hz   C = 1.0597e4  bit/s   ratio 150.0
medium-set, B = 1 kHz                            C = 6.658e3   bit/s   ratio  94.2
full 1 MHz carrier substituted                   C = 6.658e6   bit/s   ratio 9.4e4
```

None of the paper's conclusions changes. The point of the diagnostic — that substituting a bandwidth other than the binding receiver stage inflates capacity by orders of magnitude — holds at 94.2, at 150.0 and at `9.4e4` alike. **The correction is printed because a reader recomputing `1000 x 6.6582` would otherwise find a discrepancy and have no way to tell whether the error was in the label or in the physics.**

30.10 10 Discussion

30.10.1 10.1 What follows

The most transferable result is methodological: computing a favourable number and then refusing to use it. Three explanations for why acoustic control of cellular content does not work are eliminated here — bandwidth, energy and thermodynamics — and eliminating them is the whole contribution, because it leaves the two that remain in plain view. **A field advances by removing wrong explanations at least as much as by supplying right ones** [Platt, 1964, Popper, 1959]. **The demand this paper is built to meet is Meehl [1967]'s: a directional prediction in a soft science is nearly costless, so a method must make point predictions with numbers in them or it is not being tested at all. Sections 5.3, 5.4, 5.7 and 6.5 are the answer — five computed ratios, each of which could have come out otherwise.**

The receiver-bandwidth result generalises well beyond this subject. Any analysis of a channel terminating in a slow transducer has the same structure: the medium's bandwidth is irrelevant and the transducer's corner frequency is everything. **The inflation factor from substituting one for the other is about 150 in the comparison of Section 5.7 and roughly `1e5` at full carrier bandwidth**, and that error is not exotic — it is the default error.

The resolution–depth trade is a hard physical constraint with an exponent. `lambda d20` falling as `1/f^2` is the reason no purely acoustic scheme addresses individual cells at depth, and stating it as a product rather than as two separate quantities is what makes the impossibility visible. **The corpus reaches structurally parallel results elsewhere whenever apparent multiplicity collapses onto a single effective unit: a five-house network measures 3.819710 effective jurisdictions naively and 1.000000 effective blocs once the coordinating perimeter is treated as one, giving `h = -0.500000`** [Curzi, 2026]. **The `3.65e6` cells in one acoustic voxel are the same arithmetic: a large naive count and an effective count of one addressee.**

A third parallel is worth naming because it shows the same denomination trick at work elsewhere in the corpus. A branch of a forecast is scored 0-3 for the steepness of each step, summed to a resistance `R`, and weighted as `1/2^R`, so a three-point difference in steepness is a factor of eight in expectation [Curzi, 2026]. **That is the same operation as putting capacity, energy and resolution on one logarithmic scale and reading off which term dominates.**

And the addressing law is the most portable single sentence here. *Usable specificity is set by the tolerance of the matcher, not by the length of the address.* It holds for genomic addresses, for network forwarding, for record reconciliation, and for any system that matches a

name against a store. **The corpus's own canonical-form work is the same result in another domain: comparing entries on a canonical form with an independent byte-hash witness takes a reconciliation false-break rate from 40.0 per cent to 0.0 per cent at a 40 per cent manipulation rate** [Curzi, 2026] — **a tightened matcher, not a longer name.**

30.10.2 10.2 What does not follow

It does not follow that anything is decoded. Capacity is an upper bound on what a perfect decoder could recover. **Nothing in this paper identifies a decoder, and Section 12 states exactly what would.**

It does not follow that the storage-law reading of cellular state is true. It is a model-form claim, badged, and the corpus's own treatment of the contraction mapping is the bound to read alongside it: **a human life is not a metric space, nobody has measured the relevant contraction modulus, and the theorem therefore transfers as shape rather than as guarantee** [Curzi, 2026]. The same discipline applies one scale down.

It does not follow that the essence ladder describes anything physical. `k` is a defined constant; the hertz attachment is a postulate; the frequencies are essence frequencies and not resonances (Section 5.12). **The corpus's own analysis of the related width statistic is the sharpest available bound: a statistic invariant under relabelling does not even determine the object it is computed from, which formally blocks any inference from it to a physical or psychological state** [Curzi, 2026].

It does not follow that precedent supports any of this. Every precedent item in Sections 4.2, 4.8, 6.4, 6.6 and 6.7 is credited at one to three bits against six to twenty for the measured column, and **the columns are never summed** [Curzi, 2026w].

And it does not follow that a favourable capacity makes the architecture safe. The safety property is that nothing is delivered, and it is structural. **If anything were ever delivered, the safety property would be gone and the capacity figure would become a liability rather than a curiosity.**

30.10.3 10.3 Relation to the surrounding literature

The information-theoretic core is standard [Shannon, 1948, Shannon and Weaver, 1949, Cover and Thomas, 2006, Hartley, 1928] and the time–frequency bound is classical Fourier analysis [Gabor, 1946]. **The one place where this paper puts two standard results together to get a number that is not standard is Section 5.9's crossover at 20.034 dB**, and it is offered as a design point rather than as a discovery.

The thermal treatment rests on fluctuation–dissipation [Callen and Welton, 1951] and the energetic floor on Landauer [Landauer, 1961]. **Where the corpus reaches for physics beyond that it is bounded by its own analysis of the speculative margin, which reports that decoherence suppresses coherences as `rho_01(t) = rho_01(0) exp(-t/tau)` and changes no diagonal element — converting a superposition into something that looks classical without solving measurement** [Curzi, 2026aa]. **No quantum content is imported into any acoustic calculation here, and FM-C.8 exists to keep it out.**

The leaky-integrator reading of a transducer has a century-old precedent in exactly the domain it is applied to [Lapicque, 1907, Gerstner and Kistler, 2002, Usher and McClelland, 2001], and the first-order difference equation's properties are textbook [Elaydi, 2005, Oppenheim et al., 1996, Horowitz and Hill, 2015, Banach, 1922]. **That precedent supports the *shape*; it does not support the corpus's particular parameterisation, and Section 4.7 says so. The cleanest physical instance of the same shape is acoustic and is worth naming here: reverberation is an exponential decay with a measurable time constant** [Sabine, 1922, Schroeder, 1965], **which is the storage law's own `delta` in a medium where it can actually be measured.**

A structurally parallel concentration result appears elsewhere in the corpus on a different axis. A constructed maturity stack of 37.0/11.0/7.0/16.5/15.5/13.0 per cent gives a time-axis

concentration index of 0.222050, an effective spread of 4.50 buckets and 55.0 per cent refinancing within three years, against a flat ladder's 0.1667, 6.00 and 50.0 per cent [Curzi, 2026]. **The arithmetic of "many nominal units, few effective ones" is the same arithmetic as the acoustic voxel, and noticing that it recurs is the corpus's own claim rather than an argument for any of it.**

On the addressing side the relevant general literature is engineering rather than biological: open design and least privilege [Saltzer and Schroeder, 1975], disclosure as a structural property [Kerckhoffs, 1883, Curzi, 2026], and the measurement-corruption warnings that explain why a matcher scored on finding targets drifts toward tolerance [Goodhart, 1975, Campbell, 1979]. **The biological literature that actually carries the claims is named in Appendix A and carries no keys, for the reason in Section 3.7 — which is the single largest defect in this paper's apparatus.**

Finally, two standing corpus-level cautions apply here as everywhere. **The surplus functional's logarithmic form is underdetermined by the evidence:** a sub-linear power law [Latané, 1981] and `n log n` [Briscoe et al., 2006] fit the same qualitative facts, and the strongest available warrant for a logarithm is dynamical rather than perceptual [Peters and Gell-Mann, 2016]. **And no scaling or heavy-tail claim is made anywhere in this paper**, so the distribution-fitting procedure such a claim would require [Clauset et al., 2009] does not apply; the two power laws that do appear — `lambda d20 ~ 1/f^2` and radiation force quadratic in pressure — **have exponents defined by the physics, not estimated from data.**

30.11 11 Verification

30.11.1 11.1 Dimensional audit

Every quantity, its dimension, its unit, and its epistemic status. `[1]` denotes dimensionless.

Symbol	Dimension	Unit	Status
`p`	`[M L^-1 T^-2]`	Pa	**measured / specified** input pressure
`Z = rho c`	`[M L^-2 T^-1]`	Pa·s/m	**derived** from nominal `rho` and `c`
`T_I`, `t_p`	`[1]`	none	**derived** transmission coefficients
`I`	`[M T^-3]`	W/m²	**derived**: `p^2/(2 rho c)`
`lambda`	`[L]`	m	**derived**: `c/f`
`f`	`[T^-1]`	Hz	**specified**
`k = 2 pi/lambda`	`[L^-1]`	m⁻¹	**derived**
`k d`	`[1]`	none	**derived**, the fractional pressure differential
`alpha`	`[L^-1]`	dB/cm	**nominal**, and frequency-dependent
`d20`	`[L]`	cm	**derived** from `alpha`
`lambda d20`	`[L^2]`	mm·cm	**derived**; the product that falls as `1/f^2`
`kT`	`[M L^2 T^-2]`	J, or pN·nm	**standard**: `4.280e-21` J at 310 K
`T` (tension)	`[M T^-2]`	mN/m	**assumed** in the illustrative gating table
`dA`	`[L^2]`	nm²	**assumed** in the same table

Symbol	Dimension	Unit	Status
`B`	`[T^-1]`	Hz	**derived** from a **nominal** time constant, `f_c = 1/(2 pi tau)`
`S/N`	`[1]`	none (dB when logged)	**derived**
`C`	`[T^-1]`	bit/s	**derived. The bit is Shannon information [Jaynes, 2003] and is dimensionless; the rate carries `[T^-1]`**
`dt df`	`[1]`	none	**derived**; the Gabor product, dimensionless as a time times a frequency
`R_max/B`	`[1]`	bit/s/Hz	**derived**: `4 pi log2(e)/e = 6.669448`
address bits	`[1]`	bit	**derived**: `log2` of a count of positions
nats	`[1]`	nat	**derived**: `ln` of a count. `1 nat = log2(e) = 1.442695 bit`
`Q`, `eta`, `delta`, `C_cost`	see below	—	**modelling choice** (Section 4.7)

Does the storage law balance? `Q_{t+1} = (1-delta)Q_t + eta S_t - C_t` balances only if `eta` `S` and `C` carry the dimension of `Q` per period. **In Section 4.7 `Q` is "durable regulatory state" and has no operational unit at all**; `delta` is per period and dimensionless in magnitude, `S` is dimensionless in the illustrative parameterisation, and `eta` therefore carries `[Q]` directly. **The equation balances by stipulation of the units rather than by measurement of them, and that is said here rather than left implicit. `Q` at cellular scope is a placeholder with no instrument behind it**, and Section 12 records the consequence.

Where a quantity has no operational unit, that is said. Every physical quantity in Sections 4.6, 5 and 6 is on a ratio scale with a true zero and the arithmetic performed on it is licensed. **`Q` is not**, and the psychological constructs elsewhere in the corpus are not [Michell, 1997, Luce and Tukey, 1964]. The typology being applied is Stevens's [Stevens, 1946], and the operation that separates the licensed cases from the stipulated ones is division by a quantity with a true zero.

30.11.2 11.2 The postulates

Six postulates. Each is atomic, checkable alone, and carries its status. **Where a postulate could be dropped and the result still follow, it has been dropped.**

P1 — Soft tissue has `c = 1540` m/s, `rho = 1050` kg/m³, and attenuation `alpha = 0.5 f_MHz` dB/cm with exponent 1.0. *Status:* **NOMINAL CONSTANTS**, conventional round figures rather than measurements of any tissue. Section 5.4 computes the sensitivity, and **every alternative exponent moves the conclusion in the same direction.**

P2 — A cell is 10 µm in diameter and soft tissue holds `1e9` cells per cm³. *Status:* **NOMINAL, ORDER OF MAGNITUDE ONLY.** The voxel occupancy of Section 5.3 is linear in this figure, so an order-of-magnitude error moves `3.65e6` to between `3.65e5` and `3.65e7` — **and changes no conclusion, because all three are broadcasts.**

P3 — An acoustic resolution voxel is `lambda^3`. *Status:* **MODELLING CHOICE, and a deliberately optimistic one**, since real focused-beam resolution is worse than a wavelength laterally. **The optimism is in the direction unfavourable to this paper's conclusion, which is why it is adopted.**

P4 — Shannon's capacity theorem and the Gabor limit hold. *Status:* **CITED** [Shannon, 1948, Gabor, 1946]. **Both are theorems, and the Gabor bound is independently derived within this corpus** [Curzi, 2026z].

P5 — The receiver's bandwidth is the corner frequency of the slowest stage that must resolve a symbol, `f_c = 1/(2 pi tau)`. *Status:* **DEFINITION**, plus **NOMINAL** time constants. It carries no empirical weight of its own; it is the identification that makes the capacity figure a statement about a cell rather than about a loudspeaker.

P6 — A nucleotide carries two bits and the human haploid genome is about `3.1e9` base pairs. *Status:* **DEFINITION** and **CITED FACT** respectively; the genome figure is a standard round number and the bit count is `log2(4)`.

Two postulates considered and dropped. An earlier form of the capacity argument assumed a specific noise model at the cell. It is not needed: **the receiver-limited capacity of Section 5.7 is dominated by `B`, and the conclusion "the pipe is adequate" holds at every SNR in the table, from 10 dB to 30 dB.** An earlier form also assumed a specific coupling efficiency from field to membrane tension; **it is not needed either, because Section 4.6's conclusion is stated as a ratio against the gating scale and the ratio is six orders of magnitude at kilopascal pressures.**

30.11.3 11.3 The compounding chain

```
P1          -> L1   lambda(f) = 1540/f; at f = 1e6, lambda = 1.54 mm       ARITHMETIC
P1 + P2     -> L2   f at which lambda = 10 um is 1540/1e-5 = 1.540e8 Hz    ARITHMETIC
P3 + P2     -> L3   cells per voxel = lambda^3 * 1e9 per cm^3
                    at 1 MHz: (1.54e-3 m)^3 = 3.652e-9 m^3 = 3.652 mm^3
                    x 1e6 cells/mm^3 = 3.652e6                            ARITHMETIC
P1          -> L4   d20(f) = 20 / (0.5 f_MHz) cm; at 154 MHz, 0.2597 cm    ARITHMETIC
L1 + L4     -> R3   lambda x d20: 6160 mm.cm at 0.1 MHz, 0.0026 at 154 MHz
                    and the product falls as 1/f^2 exactly                 ARITHMETIC
P5          -> L5   tau = 15 ms -> f_c = 1/(2 pi 0.015) = 10.6103 Hz       ARITHMETIC
P4 + L5     -> R1   C = 10.6103 * log2(1 + 100) = 10.6103 * 6.6582
                    = 70.646 bit/s                                         ARITHMETIC
R1          -> R1b  40 bits / 70.646 = 0.56620 s                           ARITHMETIC
L3 + R1     -> R4   40 x 3.6523e6 = 1.46092e8 bits; /70.646 = 2.0680e6 s
                    = 23.93 days; per-cell rate 70.646/3.6523e6
                    = 1.934e-5 bit/s                                       ARITHMETIC
P4          -> R5   R_max/B = 4 pi log2(e)/e = 12.56637 x 1.442695 / e
                    = 18.13106 / 2.718282 = 6.669448                       ARITHMETIC
R5 + P4     -> R5b  Shannon equals Gabor when log2(1+S/N) = 6.669448,
                    so S/N = 2^6.669448 - 1 = 100.79 = 20.034 dB           ARITHMETIC
P6          -> R6a  log2(2 x 3.1e9) = log2(6.2e9) = 32.5296 bits           ARITHMETIC
P6          -> R6b  ln(6.2e9) - ln(6.2e9/85.3333) = ln(85.3333)
                    = 4.44653 nats = 6.4152 bits                           ARITHMETIC
P6          -> R6c  ball(20,2) = 1 + 60 + 1710 = 1771;
                    1771 x 3.5243e-4 = 0.62415 expected hits               ARITHMETIC
P6          -> R6d  ball(20,3) = 1771 + 30780 = 32551;
                    32551 x 3.5243e-4 = 11.4718 expected hits              ARITHMETIC
```

Three checks a reviewer can redo in a minute: `1540/10e-6 = 1.54e8`; `1/(2 pi x 0.015) = 10.6103`; and `2^6.669448 = 101.79`, so `S/N = 100.79`. Everything in R1, R3, R4 and R5 follows from those and from the nominal constants of P1, P2 and P5.

What the chain does not establish. It does not establish that any cell responds to any particular acoustic stimulus. It does not establish that a decoder exists. **And R1's status depends entirely on P5's nominal time constant: at 10 ms the capacity is 106.0 bit/s and at 1 s it is 1.06 bit/s, so the figure is a scale, not a measurement**, which is why Section 5.7 prints the whole table rather than one number.

30.11.4 11.4 The verification instance, worked

The smallest instance exercising the whole chain is the crossover of Section 5.9, because it uses both bounds and nothing else.

```
Gabor tone channel:   M <= 4 pi B dt,  R = log2(M)/dt
maximise:             d/d(dt) [ log2(4 pi B dt)/dt ] = 0
                      =>  ln(4 pi B dt) = 1  =>  4 pi B dt = e
                      =>  dt* = e/(4 pi B)
substitute:           R_max = log2(e) / (e/(4 pi B)) = 4 pi B log2(e)/e

numerically:          4 pi = 12.566371
                      log2(e) = 1.4426950
                      12.566371 x 1.4426950 = 18.131060
                      18.131060 / 2.7182818 = 6.6694480

Shannon equals it when  log2(1 + S/N) = 6.6694480
                        1 + S/N = 2^6.6694480 = 101.7900
                        S/N = 100.7900 = 10 log10(100.79) = 20.0342 dB
```

Status: **ARITHMETIC** on a **DEFINITION** (the tone-selection scheme) and a **CITED** theorem. **A reader with a calculator can confirm every figure, and the agreement of two independently derived bounds at a single ratio is the strongest internal check in the paper.**

30.11.5 11.5 Limits, asymptotics, and the domain of validity

Limit	Behaviour	Reading
`f -> 0`	`lambda -> infinity`, `k d -> 0`, cells per voxel `-> infinity`	the field is perfectly uniform and perfectly non-selective; the audible band is here
`f -> 154` MHz	`lambda -> d`, cells per voxel `-> 1`	**the only frequency at which the field resolves a cell**, and `d20 = 2.6` mm
`f -> infinity`	`d20 -> 0` as `1/f`, `lambda d20 -> 0` as `1/f^2`	resolution is bought at a quadratic price in reachable volume; **the door closes**
`S/N -> 0`	`C -> 0`	no signal, no capacity; correct
`S/N -> infinity`	`C ~ B log2(S/N)`, unbounded but logarithmic	**shouting never buys a fast channel**; the 100 Pa $\rightarrow$ 1 MPa case buys 2.80×
`B -> 0`	`C -> 0`	**a receiver that cannot change state carries nothing, whatever the medium does**
`tau -> 0`	`f_c -> infinity`, `C -> infinity`	an infinitely fast transducer; no biological stage approaches it

Limit	Behaviour	Reading
`A -> 0` (aperture)	`p_noise -> infinity` as `A^-1/2`	**small receivers are noisy receivers**; the formula's own regime fails first (Section 5.5)
`d -> 0` (matcher tolerance)	expected hits `-> 3.52e-4`	a perfect matcher; the address is essentially unique
`d -> 20`	ball `-> 4^20`, hits `-> 3.875e8`	**a matcher with total tolerance names every addressable position**, which is the reductio of FM-A.3
`delta -> 0` in Section 4.7	no decay; state persists forever	**a pathology, not an ideal** — nothing decays out of the stock [Curzi, 2026a]

Two domain restrictions follow and are stated as restrictions. The aperture noise formula of Section 5.5 is valid only for `k a >> 1` and is used at cellular scale where `k a << 1`; it is therefore a reference calculation with a declared regime rather than a characterisation, and **the paper carries two estimates rather than pretending one suffices.** And the `lambda^3` voxel convention of P3 is optimistic; **real lateral resolution is worse, so the true cells per voxel exceed the figures given.**

30.11.6 11.6 Invariance, and how far each claim has been tested

The resolution–depth product's `1/f^2` scaling is invariant under the attenuation coefficient's magnitude and not under its exponent. Doubling `alpha` halves `d20` at every frequency and leaves the exponent untouched; changing the exponent from 1.0 to 1.5 changes the product's slope from `-2` to `-2.5`. **Tested at three exponents (1.0, 1.1, 1.5) and at three frequencies (10, 100, 154 MHz): nine points, all moving the conclusion in the same direction. Nine points, not a proof.**

The Gabor–Shannon crossover is invariant under everything. `6.669448` is `4 pi log2(e)/e`, a pure number with no physical constant in it, so the 20.034 dB crossover holds for any wave, any medium and any bandwidth. **That is the strongest invariance claim in the paper and it is an identity rather than a sample.**

The capacity figure is not invariant and is not claimed to be. It scales linearly in `B` and logarithmically in `S/N`, and `B` rests on a nominal time constant that varies by isoform and preparation. **The invariant statement is the ordering — receiver bandwidth dominates medium bandwidth by orders of magnitude — and only that ordering is asserted.**

The address arithmetic is invariant under genome size to within a logarithm. A tenfold error in `3.1e9` moves `log2(6.2e9) = 32.53` bits by `log2(10) = 3.32` bits, **which changes no verdict about a 40-bit or 46-bit address.**

30.11.7 11.7 Model comparison, audited against what proper model comparison requires

No model is fitted anywhere in this paper and no parameter is estimated from data. That makes several standard requirements inapplicable rather than unmet, **and the distinction is drawn explicitly here rather than used as cover.**

Requirement	Status in this paper
Out-of-sample rather than in-sample performance	**NOT APPLICABLE.** Nothing is fitted; there is no sample and no holdout.
A matched null with the same free-parameter count	**NOT APPLICABLE** for the same reason. The nearest analogue actually run is the attenuation-exponent sweep of Section 5.4, which is a sensitivity analysis rather than a model comparison.
The number of free parameters stated explicitly	**MET.** The capacity model has two inputs (`B`, `S/N`), both declared; the attenuation model has two (`alpha`, `b`); the storage-law illustration of Section 4.7 has four (`eta`, `delta`, `N`, `C`), **all four assumed and none fitted.**
An information criterion, so the parameter count is paid for	**NOT MET.** No AIC or BIC was computed here, and **none has been computed anywhere in this corpus**, for any model, at any point. Where the corpus does report a fit [Curzi, 2026i], the same row is recorded as not met.
Effect sizes with intervals, not only p-values	**NOT MET, and no p-value is reported anywhere in this paper.** Every figure is a point value with no interval, **because an interval on a nominal constant would be false precision.** The honest substitutes are the sensitivity sweep of Section 5.4 and the full capacity table of Section 5.7, which are reported as ranges and not as confidence statements.

The two unmet rows change how the paper should be read. Because no information criterion was computed, the three-parameter attenuation model with a free exponent has not been shown to be *better* than the two-parameter nominal law in any formal sense — **it has been shown to be more realistic on grounds external to the numbers, namely published tissue measurements, and that is an argument from physics rather than from model selection.** And because no interval accompanies any figure, **every number in this paper should be read as an order-of-magnitude statement with the nominal constants stated, which is exactly how Section 3.3 frames them.**

30.12 12 Limitations

L-1. The entire biological evidence base of this paper is uncited. The corpus's bibliography has no cluster for molecular and cell biology, so mechanically gated channels, cochlear transduction, force-to-transcription coupling, mechanosensitive transcriptional regulators, substrate-stiffness effects, sonogenetics, nucleotide data storage, guide specificity, self-processing arrays, motif dependence, catalytically dead nucleases, compact effectors and mitochondrial base editing are **all named in the text with journal, volume and year and carry no citation keys.** Appendix A lists them. **A reviewer cannot check this paper's biology from its reference list, and that is a defect of the apparatus rather than a claim that the works do not exist.** It must be closed before submission.

L-2. A bibliographic mis-citation was found in the source and is corrected rather than propagated. The source attributes the thermal-noise floor to "Nyquist (1928)"; the corpus's Nyquist 1928 entry is the telegraph-transmission paper, not the thermal-agitation paper. **Attaching the available key would have been a mis-citation**, so the thermal floor is attributed

here to the fluctuation–dissipation theorem [Callen and Welton, 1951] and Nyquist's thermal paper is named in Appendix A without a key.

L-3. Every physical constant here is nominal and none was measured for this paper. Sound speed, density, attenuation and its exponent, cell diameter, cell density, channel time constants, direct-repeat and spacer lengths, protein residue counts, eardrum and footplate areas, and gating-spring stiffness are conventional round figures. **The conclusions are robust to the attenuation exponent in the unfavourable direction (Section 5.4) and to cell density by an order of magnitude (Section 11.2), and they are not robust to the receiver time constant, which moves the capacity by two orders of magnitude across the table.**

L-4. The aperture noise model is used outside its regime and the paper cannot repair it. `p_rms = sqrt(4 kT rho c B / A)` assumes `k a >> 1`; at cellular scale `k a << 1`. **A rigorous treatment requires the fluctuation–dissipation theorem applied to the actual mechanical impedance of a cell, which this paper does not have.** Two estimates are carried instead of one, and the honest statement is that the cellular noise floor is not established here.

L-5. The capacity result rests on identifying "the receiver's bandwidth" with a single corner frequency. A real transducer chain has several stages with different kinetics, an input-dependent response and, in the case of the channel concerned, an inactivation that depends on membrane potential. **Reducing that to one pole is a modelling choice, and a multi-pole treatment would give a different — probably lower — figure.**

L-6. The storage-law reading of cellular regulatory state is a model-form claim with no instrument behind it. `Q` has no operational unit, `eta`, `delta` and `C` were chosen rather than measured, and nothing in this paper tests the reading [Curzi, 2026]. **The headroom-negative prediction of Section 4.7 is a restatement of a known observation, not evidence for it.**

L-7. The reflection problem is unanswered wherever this paper's framework touches the reflexive loop. Manski [1993] and Shalizi and Thomas [2011] bind, no remedy is adopted [Brock and Durlauf, 2001], and every framework claim resting on the loop is a modelling position rather than an identified empirical result (Section 3.6).

L-8. Scale type is licensed for the physics and stipulated for the rest. Pressure, frequency, time and energy are ratio-scale quantities with true zeros and the arithmetic on them is licensed. **`Q` is not, and the psychological constructs elsewhere in the corpus are not**: Michell [1997]'s objection applies, and the additive conjoint conditions that would license ratio-scale treatment [Luce and Tukey, 1964, Krantz et al., 1971] have not been tested for any of them.

L-8b. One figure in the source manuscript was found not to reproduce and is corrected in Section 9.4. The source labels `1.060e4` bit/s as a medium-limited capacity at 1 kHz bandwidth; it is in fact the receiver-limited capacity at `B = 1591.6` Hz, and the 1 kHz figure is 6,658 bit/s. **No conclusion depends on which of the three inflation factors is quoted, and the correction is printed rather than silently applied.**

L-9. No information criterion was computed (Section 11.7), and no figure in this paper carries an interval.

L-10. The essence-ladder material rests on a defined constant with a postulated unit. `k` is exact; the hertz attachment is a postulate; the frequencies are not physical resonances. **The standard critique of golden-section claims** [Markowsky, 1992] **is cited deliberately and the corpus's own bound on relabelling-invariant statistics** [Curzi, 2026] **applies.**

L-11. The `lambda^3` voxel convention is optimistic and the true occupancy is worse. Real focused-beam lateral resolution exceeds a wavelength, so `3.65e6` cells per voxel at 1 MHz is a lower bound on the number sharing a signal.

L-12. What would falsify the central negative claim. The claim that no mechanism demodulates sequence-specific information from an acoustic signal would be overturned by a demonstration, replicated, of a transcriptional or other cellular response that depends on the *order* of symbols in an acoustic stream and not merely on permutation-invariant functionals of the waveform such as total energy, spectral content or envelope. **That is a well-specified experiment, and specifying it is the most useful thing this paper does with its negative result** [Curzi, 2026y,?]. **Until it is run, the claim is a statement about the state of the literature and is**

labelled as one.

30.13 13 Conclusion

This paper computes the acoustic channel into a single cell and separates what the physics permits from what has been demonstrated.

Air-to-tissue intensity transmission at normal incidence is `-29.9` dB while the pressure amplitude nearly doubles; the pressure differential across a 10 µm cell is 0.0816 per cent of peak at 20 kHz and 4.08 per cent at 1 MHz; and the wavelength equals a cell diameter only at 154 MHz, where attenuation confines the field to millimetres. The product of wavelength and 20 dB depth falls from 6160 to 0.0026 mm·cm across that range, closing acoustic addressing on physical grounds rather than on capacity.

The binding term is the receiver's own bandwidth. Mechanotransduction kinetics give a corner frequency near 10.61 Hz, so Shannon capacity at 20 dB is **70.65 bit/s**, rising to 106.0 bit/s at a 10 ms constant, and **a 40-bit address takes about half a second.** That result is used to eliminate an explanation rather than to support a claim: **individually addressing the `3.65e6` cells in one voxel at 1 MHz would take 23.93 days, a per-cell exclusive rate of `1.93e-5` bit/s, so the field is intrinsically broadcast.** The Gabor limit `dt df >= 1/(4 pi)` yields an optimal symbol rate `R_max = 6.669448 B`, **cross-checked against Shannon at exactly 20.034 dB.** And **energy is not the constraint**: the Landauer bound permits `8.19e3` bit/s at one pascal.

On the addressing side, naming one genomic position costs 32.5296 bits, the T-rich motif constraint costs a further 4.4465 nats — 6.415 bits — of naming capacity, and the worked block is declared valid only at matcher tolerance `d <= 2`, where expected collisions are 0.624 against 11.47 at `d <= 3`. The governing law is that **usable specificity is set by the matcher's tolerance and not by the address's length.** And because the architecture introduces no protein and no nucleic acid into any cell, **it cannot edit a genome — cannot, not does not.**

The counter-result is the point. This paper computes a favourable number and refuses to use it. **The pipe into a cell is wide enough, which removes bandwidth from the list of explanations — and the standing line is restated in every section: a channel capacity is a statement about the channel, and no molecular apparatus is known that recovers a specified sequence from an acoustic waveform. Computing the capacity of a channel leading to a receiver does not create the receiver. Measuring a room's acoustics says nothing about whether anyone in it speaks the language.**

30.13.1 Declarations

Data and code availability

No external dataset was collected for this paper and no experiment was performed. Every number reported is arithmetic on the nominal physical constants declared in Section 3.3, on standard physical constants, or on closed-form combinatorial expressions stated in the text. The three worked specifications of Section 9 are **hypothetical proposals constructed for evaluation**; no installation was built, no acoustic link was assembled, and no address block was synthesised. **No cell, tissue, organism or person was involved at any point.** The five figures are generated by the scripts named in each caption, and those scripts contain no data beyond the numbers printed in this paper.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi performed the analysis, computed every figure in the text, specified the requirements and failure modes, produced the figures, and wrote the paper.

Ethics

No human-subjects research was conducted, no animal work was performed, and no personal data was processed. **No clinical, medical, diagnostic or therapeutic claim is made anywhere in this paper, and the pressures and intensities tabulated for capacity purposes are not recommendations, exposure guidance, or safety statements of any kind.** The architecture described in Section 6 introduces no protein and no nucleic acid into any cell and therefore cannot edit a genome; that is a structural property of the design and is stated as one rather than as an undertaking. **The biological literature described here is described from published sources; none of it was replicated, extended or performed by the author.**

Preregistration

This paper reports no preregistered empirical test and none is claimed. **The falsifier for its central negative claim is stated in advance in Section 12, L-12**, and is the experiment the successor paper preregisters [Curzi, 2026y]: a replicated demonstration of a cellular response depending on the *order* of symbols in an acoustic stream, as distinct from permutation-invariant functionals of the waveform. **No such experiment was performed here, and the negative claim is a statement about the state of the literature rather than a null result.** The reason to declare the falsifier in advance rather than after the fact is the standard one [Nosek et al., 2018, Simmons et al., 2011]: **a negative claim whose refutation conditions are written down afterwards can be moved, and this one cannot.**

30.14 Appendix A Works named in the text without a bibliography key

The bibliographic database governing this corpus contains no entries for the following, and **no key has been attached to any of them anywhere in this paper.** They are listed with the details the source manuscript supplies so that a reader can locate them, and their absence is recorded as a gap in the corpus's apparatus rather than concealed (Section 3.7, limitation L-1). **Nothing in this list should be read as a citation; it is a statement of what is missing.**

Mechanotransduction and channel structure. Coste and colleagues (*Science* 330:55–60, 2010) on Piezo1 and Piezo2 as essential components of mechanically activated cation channels, and the 2021 Nobel Prize in Physiology or Medicine to David Julius and Ardem Patapoutian for discoveries of receptors for temperature and touch. The cryo-electron microscopy structures (Ge et al., *Nature*, 2015; Guo and MacKinnon, *eLife*, 2017; Saotome et al., *Nature*, 2018) and the membrane-dome gating model. Cox and colleagues (*Nature Communications*, 2016) and Syeda and colleagues (*Cell Reports*, 2016) on intrinsic bilayer-tension gating.

The cochlea. Pickles, Comis and Osborne (*Hearing Research*, 1984) on tip links; Kazmierczak and colleagues (*Nature* 449:87, 2007) on the cadherin-23 / protocadherin-15 interaction; Corey and Hudspeth (*Journal of Neuroscience*, 1983) on transduction kinetics; Howard and

Hudspeth (*Neuron*, 1988) on the gating-spring compliance change; Zheng and colleagues (*Nature* 405:149, 2000) on prestin; Hudspeth (*Nature Reviews Neuroscience*, 2008) on active amplification; Bialek (*Annual Review of Biophysics and Biophysical Chemistry* 16:455, 1987) on how close biological sensors run to physical limits. The identity of the pore-forming transduction subunit remains an open argument in that literature.

Force to the nucleus. Tajik and colleagues (*Nature Materials* 15:1287–1296, 2016) on force-induced direct chromatin stretching and transcriptional upregulation; Wang, Tytell and Ingber (*Nature Reviews Molecular Cell Biology* 10:75, 2009) on mechanotransduction at a distance and the LINC complex; Dupont and colleagues (*Nature* 474:179, 2011) on mechanoresponsive transcriptional co-activators; Engler, Sen, Sweeney and Discher (*Cell* 126:677, 2006) on substrate stiffness and lineage specification.

Sonogenetics. Ibsen, Tong, Schutt, Esener and Chalasani (*Nature Communications* 6:8264, 2015) on TRP-4 in *C. elegans*, including the reported reliance on microbubble amplification; Duque and colleagues (*Nature Communications* 13:600, 2022) on a human channel at megahertz frequencies.

Sequence addressing and storage. Jinek and colleagues (*Science* 337:816, 2012), Cong and colleagues (*Science* 339:819, 2013) and Mali and colleagues (*Science* 339:823, 2013) on guide specificity; Qi and colleagues (*Cell* 152:1173–1183, 2013) and Gilbert and colleagues (*Cell* 159:647–661, 2014) on catalytically dead Cas9 for transcriptional repression and activation; Zetsche and colleagues (*Cell* 163:759–771, 2015) and Fonfara, Richter, Bratovič, Le Rhun and Charpentier (*Nature* 532:517–521, 2016) on Cas12a and its self-processing of pre-crRNA; Pausch and colleagues (*Science* 369:333–337, 2020) on a hypercompact editor; Mok and colleagues (*Nature* 583:631–637, 2020) on CRISPR-free mitochondrial base editing; Church, Gao and Kosuri (*Science*, 2012) on digital information storage in synthesised DNA.

Physics and engineering named without keys. Nyquist's 1928 paper on thermal agitation of electric charge in conductors — **which is a different paper from the corpus's `nyquist1928` entry, and the distinction is the subject of limitation L-2.** The Internet Protocol specifications underlying the addressing correspondence of Section 6.3. Chladni's 1787 *Entdeckungen über die Theorie des Klanges*, named in the source manuscript's concordance on boundary-selected modes.

Precedent material named without keys. Boethius, *De institutione musica* (c. 500 CE), on *musica humana*. Bhartṛhari, *Vākyapadīya* (c. 5th century), on meaning-bearing sound. The *Lüshi Chunqiu* (3rd century BCE) on the twelve *lü* pitch standards, and the associated *houqi* practice. Confucius, *Analects* XIII.3, on the rectification of names. The Alexandrian *Pinakes* of Callimachus (3rd century BCE). The Egyptian narrative of Isis and the secret name of Ra, the Kabbalistic treatment of divine names, and the Solomonic grimoire tradition's insistence on exact naming. The Roman ritual formula *sive deus sive dea*. **Every one of these is precedent, priced at one to three bits on the corpus's own scale, and none of them is evidence about anything in this paper** [Curzi, 2026w].

Chapter 31

Addressing the Non-Coding Genome: The Latency Constraint, the Missing Demodulator, and an Exact Sonification Pipeline

Abstract. We state the boundary of a genomic addressing architecture as a constraint rather than a caveat: addressing can reach only what is already latent, since an address is a pointer and a pointer with no cargo delivers nothing. The same property that makes the architecture safe (no sequence is changed) makes it limited. We impose a pairing rule on the descriptive statistics: the 98.7 per cent non-coding, 45 per cent transposable-element and 8 per cent endogenous-retroviral fractions may not be quoted without the approximately 8.2 per cent evolutionarily constrained fraction in the same paragraph. The central finding is negative and is stated as the paper's core: no known biological structure demodulates a symbol stream from an acoustic carrier into sequence-specific output, and channel capacity is not the bottleneck — 40 bits over a 16 Hz bandwidth requires only 6.68 dB of signal-to-noise — which sharpens the gap rather than softening it. A decisive experiment is preregistered at n = 21 per arm for a 1.5-fold change under genome-wide correction, and we compute that one well-powered experiment covers 1 of 2,880 conditions, 0.0347 per cent of the parameter space. The sonification chain is built and verified: an implemented run-length-limited codec achieves 1.92834 bits per base against a constrained capacity of 1.98235, with an exact SHA-256 round trip.

Keywords: non-coding DNA; sequence-specific addressing; demodulation; preregistration; DNA data storage; sonification; Gabor limit

Series: B. Engineering the Surplus **Paper:** CCP-031

31.1 1 Introduction

An address is a pointer, and a pointer with no cargo delivers nothing. This paper states that sentence as a formal constraint on a genomic addressing architecture, computes what the constraint costs, and then reports that the architecture's central biological claim has no mechanism at all.

The preceding paper in this series established that a cell is a mechanical receiver, that the acoustic channel into one has adequate capacity and catastrophic spatial resolution, and

that addressing and editing are demonstrably separable functions [Curzi, 2026x]. This paper asks what such an address could point *at*, states the boundary of the whole architecture as a constraint rather than a caveat, and specifies the experiment that would settle the open question in either direction.

The constraint comes first because everything else is bounded by it.

> **CONSTRAINT L67. Addressing can reach only what is already latent. No addressing operation, of any kind, at any signal strength, by any modality, can cause a cell to express content that its genome and its accessible regulatory repertoire do not already contain.**

That follows from the architecture rather than from any new assumption. **The property that makes the architecture safe — nothing is delivered, so nothing can be edited — is the same property that makes it limited.** A design claiming both unlimited reach and no modification would be incoherent, and this one claims neither.

A pairing rule is imposed on the descriptive statistics, and it is imposed because the figures are routinely used to do work they cannot do. The human genome is approximately 98.7 per cent non-coding, approximately 45 per cent transposable-element derived and approximately 8 per cent derived from endogenous retroviruses; **and roughly 8.2 per cent of it is under evolutionary constraint.** The first three figures may not be quoted in support of this architecture without the fourth in the same paragraph. **The figures are not in conflict — they measure different things — and presenting one without the other is advocacy rather than argument.** The counter-example that has never been answered is stated in full: an onion carries 16.0 Gbp against the human 3.10 Gbp, a lungfish 130 Gbp, and a pufferfish manages vertebrate development on 0.39 Gbp. **Any theory assigning function to bulk sequence because it is there must answer that, and none has.**

The central finding is negative and it is the paper's core rather than its caveat. No known biological structure demodulates a symbol stream from an acoustic carrier into a sequence-specific output. Demodulation of an ordered string requires three things and there is no way around any of them: **symbol discrimination (D1), order memory (D2), and a sequence-specific output mapping (D3).** D1 and D2 are hard; **D3 has no candidate implementation whatever**, because base pairing is sequence-specific precisely by virtue of a physical guide molecule, and this architecture delivers no molecule.

And capacity is emphatically not the bottleneck, which sharpens the gap rather than softening it. At a receiver bandwidth of 16 Hz, delivering a 40-bit address in one second requires `2^(40/16) - 1 = 4.657`, which is **6.68 dB** of signal-to-noise; at 5 Hz it requires 24.07 dB, and even at zero decibels and 5 Hz the address arrives in eight seconds. **The information-theoretic room exists. If the obstacle were capacity it would be a quantitative obstacle, and quantitative obstacles yield to engineering. It is not.** A radio wave carrying a symphony passes through a rock; the capacity of the path is enormous; the rock does not hear the symphony, and the missing thing is not bandwidth.

A decisive experiment is therefore specified and preregistered rather than gestured at. Its primary control is the **symbol-permuted twin** — a stimulus matched in total energy, duration, long-term power spectrum, symbol counts, envelope and crest factor, differing only in order — and **every biological response to acoustic stimulus documented in the literature is invariant under that permutation by construction.** The design needs **`n = 21` per arm** to detect a 1.5-fold change at a genome-wide-corrected `alpha = 2.5e-6`, and a single well-powered experiment covers **1 of 2,880 conditions, or 0.0347 per cent of the parameter space.** An abandonment condition is committed to in advance: a preregistered 73-condition three-stage adaptive design, at 2.53 per cent of exhaustive cost, whose null withdraws the biological claim from this corpus entirely.

The sonification chain is then built and verified. An implemented run-length-limited codec achieves **1.92834 bits per base** on the test payload — 592 bits into 307 nucleotides, 96.42 per cent of the unconstrained 2.0 — against a transfer-matrix capacity of **1.98235** bits per base for maximum run 3, **with an exact SHA-256 round trip.** A single-register pitch mapping is shown

to fail: its four symbols span 41.034 cents with a minimum adjacent gap of 2.081 cents, or 0.8725 Hz, which the Gabor limit caps at 5.5 symbols per second and which a routine 2 cm^{-1} solvent shift **reverses**. Assigning each symbol its own integer octave index `j = 33..36` raises the minimum adjacent gap to **732.037 Hz, an 839.0-fold improvement, lifting the maximum symbol rate from 5.5 to 4,599.5 symbols per second at zero cost to the tuning**, because octave transposition preserves pitch class exactly.

Three rejections are printed in full, and one of them is of the framework's own centrepiece. The framework's essence law `f(n) = k n^2` is shown, by its own theorem, to yield only `ceil(N/2) = 2` distinct pitch classes for a four-symbol alphabet — three of the four bases collide — and is **rejected in print as the alphabet tuning.** The identification of the 64 codons with the 64 hexagrams of the *I Ching* is **rejected**: `2^6` and `4^3` are different factorisations of the same cardinality, and a proposed correspondence between them is one chosen bijection out of `64! = 1.2689e89`, worth 296.0 bits of arbitrary choice. And the claim that "not junk" means "all functional" is **refused** wherever the genome fractions appear.

Sections 2 and 3 give the hypotheses and the method, including the analogy discipline, the identification problem, and an explicit disclosure that **the entire biological evidence base of this paper carries no citation keys** because the corpus's bibliography has no cluster for it. Section 4 characterises the archive. Section 5 states the open problem. Section 6 specifies the pipeline. Sections 7 and 8 collect the specifications and named failure modes. Section 9 works three things end to end: an audit that rejects a proposal at step one, the full preregistration, and the pipeline run with verified digests. Sections 10 to 13 discuss, verify, limit and conclude.

31.2 2 Hypotheses

Five hypotheses, each stated so it could be wrong, each with its falsifier. **One of them is a constraint rather than a hypothesis and is labelled as such; two further claims are labelled as postulate and definition.**

> **H1 — The reachable set is closed under addressing.** No addressing operation causes a cell to express content absent from its genome and accessible regulatory repertoire. **Predicted:** every demonstrated instance of regulatory-only intervention reorganises expression of existing sequence and adds none. **Falsified by:** a demonstration of a cell expressing content demonstrably absent from its genome and repertoire following an addressing-only intervention. *(Result: **not falsified**, and **this is a structural constraint rather than an empirical finding** — it follows from the architecture having no import path, and Section 4.7 states it as Constraint L67. It is falsifiable in the strong sense: one observation would end the model, and no adjustment of parameters could repair it.)*
>
> **H2 — The obstacle to acoustic sequence-specific signalling is the absence of a demodulator, not the capacity of the channel. Predicted:** the signal-to-noise ratio required to deliver a 40-bit address in one second at plausible receiver bandwidths is modest — under 25 dB at 5 Hz and under 7 dB at 16 Hz — while no identified structure satisfies D1, D2 and D3. **Falsified by:** a capacity computation showing the channel cannot carry an address in a physiologically relevant time, or the identification of a structure satisfying all three requirements. *(Result: **not falsified** — 6.68 dB at 16 Hz, 24.07 dB at 5 Hz; see Section 5.4. **The favourable capacity makes the gap sharper, not smaller**, because it removes the explanation that would have been engineerable.)*
>
> **H3 — A response that is invariant under symbol permutation is not sequence-specific. Predicted:** every documented biological response to acoustic stimulus — heating, membrane tension, shear, cavitation, streaming, radiation force — depends only on permutation-invariant functionals of the waveform, so a permuted-twin control is not one control among several but the experiment itself. **Falsified**

by: an identified acoustic-biological response that depends on symbol order. *(Result:* ***not falsified*** *— see Sections 5.2 and 5.8.* ***This is close to a definition and is labelled as such where it is stated****; its empirical content is the claim that the documented responses are exhaustive of what is known, which is a claim about the literature.)*

H4 — Register spreading is exact rather than approximate, and buys two orders of magnitude. Assigning each symbol its own integer octave index leaves the tuning system unchanged and raises the minimum adjacent gap by a large factor. **Predicted:** octave transposition preserves pitch class exactly, so every spread interval equals its pitch-class interval plus a whole number of octaves; and the minimum gap rises from 0.8725 Hz to 732.037 Hz, an 839.0-fold improvement in the Gabor-limited symbol rate. **Falsified by:** any spread interval that is not its pitch-class interval plus an integer number of octaves, or a computed improvement factor differing from 839.0. *(Result:* ***not falsified*** *— checked by subtracting 1200 cents from each spread interval; see Section 6.7 and the verification instance in Section 11.4.)*

H5 — A single well-powered null bounds a negligible fraction of the parameter space, and only a committed abandonment condition makes the programme falsifiable. Predicted: the parameter space is of order thousands of conditions, so one experiment covers well under one per cent, and exhaustive coverage is unaffordable. **Falsified by:** a parameter space small enough that a handful of experiments exhausts it, or a design in which one condition generalises across the space. *(Result:* ***not falsified*** *— 2,880 conditions; one experiment covers 0.0347 per cent; exhaustive coverage is 423,360 libraries.* ***The consequence is stated against this framework rather than against its critics****: a claim that can retreat into a 2,880-cell space indefinitely is not a scientific claim, and Section 5.10 commits to a 73-condition floor.)*

Three further claims are labelled as what they are.

That non-coding sequence constitutes an *addressable archive* is a postulate, badged wherever it appears. The measured rows behind it — the fractions, the syncytin case, the reprogramming case — are facts about genomes; **the reading of those facts as an archive is this framework's own.**

That "sequence-specific" means order-dependent rather than permutation-invariant is a definition (Section 5.2), true by stipulation and carrying no empirical weight. Its whole function is to make the primary control non-negotiable.

That the reachable set of a differentiated cell is substantially larger than current practice exploits is a speculation, and it is marked speculative in the sentence that states it (Section 4.7). It predicts that the rate of discovery of new transdifferentiation protocols will continue rather than saturate, which is a weak but real test, and it has no other support here.

31.3 3 Method

31.3.1 3.1 The ten-and-one accounting

The gap is a single missing link, not a general absence of evidence, and stating that with the same precision used to state the gap is a requirement rather than a courtesy: otherwise an honest reader cannot tell how much is standing.

Component	Status
Arbitrary digital data can be stored in nucleotide sequence	**ESTABLISHED**
Higher-rate storage codes exist and are demonstrated	**ESTABLISHED**
A ~20 nt spacer is a sufficient genomic address	**COMPUTED** (Section 6.3)
An effector self-processes a multi-address array	**ESTABLISHED**
Targeting and editing are separable functions	**ESTABLISHED**
Mechanical force gates ion channels	**ESTABLISHED**
Surface force reaches the nucleus and alters transcription	**ESTABLISHED**
Ultrasound can activate specific engineered channels	**ESTABLISHED**
Non-coding sequence is a large addressable archive	**ESTABLISHED** as fractions; **POSTULATE** as “archive”
Regulatory input reorganises expression without adding sequence	**ESTABLISHED**
A cell demodulates sequence-specific content from sound	**NO ESTABLISHED MECHANISM**

Ten established rows and one open row. Both halves are reported together. Reporting only the ten is advocacy; reporting only the one is a different kind of dishonesty, because it implies the whole edifice is speculation when most of it is textbook. **Every row but the third and the last is a biological result carrying no citation key here — see Section 3.7.**

31.3.2 3.2 The pairing rule

S67-1 — the pairing rule, and it is binding on every use of the genome fractions. Any use of the non-coding fraction in argument must state the figure used (98.7 per cent, 45 per cent, 8 per cent), its annotation dependence, **and the constrained fraction of approximately 8.2 per cent, in the same paragraph.** Failure to pair them is failure mode FM-A.1.

The reason is arithmetical rather than rhetorical. The reported span between “biochemically active” and “under evolutionary constraint” is a factor of **9.76** — approximately 80 per cent against approximately 8.2 per cent — and it is a live scientific dispute rather than a settled figure. **The source manuscript states this span as “about 10.7”; recomputing from its own two figures gives `80 / 8.2 = 9.756`, and the discrepancy is reported in Section 4.3 rather than silently propagated. A framework that quotes the large number in support of itself and omits the small one is not making an argument; it is choosing a number.**

Specification labels are stable identifiers, not section pointers. `S67–1`, `L67`, and the `S68–n`/`S69–n` and `F67.n`/`F68.n`/`F69.n` labels are carried unchanged from the source manuscript’s numbering because other papers in this corpus cite them by name. **Section references have been renumbered to this paper’s own sections; the identifiers have not**, exactly as the `CR–nn` register entries are stable elsewhere in this series [Curzi, 2026w].

31.3.3 3.3 Nominal and annotation-dependent figures, declared

Quantity	Value used	Status
Haploid human genome	`3.1e9` bp	**nominal**, a standard round figure
Non-coding fraction	~98.7 %	**annotation-dependent**; stable to a few points across recent assemblies
Transposable-element fraction	~45 %	**annotation-dependent**, and a **lower bound that keeps rising** as detection of degraded elements improves
LTR / endogenous-retrovirus fraction	~8 %	**annotation-dependent**
Constrained fraction	~8.2 %	**estimate**, method-dependent
Protein-coding CDS fraction	~1.25 % (0.86 % by the gene-count route)	**convention-dependent**; the two bracket the truth
Sequence-specific transcription factors	~1600	**annotation-dependent**
Reprogramming efficiency	~0.01 % to 0.1 %	**wide**, varying by cell type, delivery and scoring
Speed of sound, soft tissue	1540 m/s	**nominal**
Attenuation	0.5 dB/cm/MHz	**nominal**
Cell diameter, density	10 µm, 1050 kg/m^3	**nominal**
Infrared band centres	1626 / 1652 / 1663 / 1665 cm^{-1}	**nominal**, standing in for the operator's own measurement; **a 2 cm^{-1} shift is routine**
Protein diffusion coefficient	`1e-10` m^2/s	**nominal**
Speed of light	299,792,458 m/s	**exact by definition** since the 1983 redefinition of the metre

Every physical figure in Sections 5.3 and 5.5 is an order-of-magnitude argument under plane-wave, linear, homogeneous-medium assumptions with a quasi-static relation for membrane tension. They are strong enough to carry the order-of-magnitude conclusions drawn from them and not strong enough to carry anything finer. Gene, lncRNA and miRNA counts change between annotation releases and are given as approximations or ranges throughout.

31.3.4 3.4 Analogy discipline: what licenses each borrowing

Every borrowing in this paper is labelled here, at first use, as exactly one of three things, and no later section may let a category-2 or category-3 borrowing collect the authority of a category-1 one.

Borrowing	Category	What licenses it
Shannon capacity (§5.4)	**(1) derivation**	The theorem applies wherever a bandwidth and a signal-to-noise ratio are defined [Shannon, 1948, Cover and Thomas, 2006]. **It bounds a rate; it constructs no decoder**, and Section 5.4 exists to say so.

Borrowing	Category	What licenses it
The Gabor limit `dt df >= 1/(4 pi)` (§6.9)	**(1) derivation, with an explicit disclaimer**	Classical Fourier duality [Gabor, 1946], applying to every wave including sound and carrying **no quantum content whatever**; the corpus derives it independently by the same variational argument with `p = hbar k` deleted [Curzi, 2026z]. The position–momentum relation [Kennard, 1927] is a fact about state preparation and is **not** invoked.
The Nyquist sampling condition (§6.10)	**(1) derivation**	Standard sampling theory [Nyquist, 1928, Shannon and Weaver, 1949], applied to choose a render sample rate.
Constrained-coding capacity as a transfer-matrix eigenvalue (§6.2)	**(1) derivation**	The capacity of a run-length-limited constraint is `log2` of the largest eigenvalue of its transfer matrix [Shannon, 1948, Cover and Thomas, 2006]; the computation is exact and reproducible.
Low-Reynolds-number hydrodynamics and the Stokes number (§5.5)	**(1) derivation**	Claims about a physical object oscillating in a physical fluid; the physics applies directly. **The originating exposition carries no citation key here** (Appendix A).
Octave equivalence and pitch class (§6.5, §6.7)	**(2) shared mathematical form**	That halving a frequency preserves pitch class is a fact about the ratio [Shepard, 1982]; that human listeners treat octave-related pitches as equivalent is a separate, measured, and *bounded* claim [Demany et al., 2013, Burns, 1999]. **The design uses the arithmetic fact and does not require the perceptual one.**
The tuning argument against a standardised reference pitch (§6.8)	**(2) shared structural form**	That A440 is a twentieth-century standardisation rather than a discovery is documentary [Barbour, 1951, Lindley and Turner-Smith, 1993]; that equal temperament is a deliberate detuning purchasing modulation is standard theory [Benson, 2006, Doty, 2002]. **No claim is made that ratio-based tuning is perceptually superior**, and the literature that would contest such a claim is cited here deliberately [Sethares, 2005, Plomp and Levelt, 1965].
The leaky integrator applied to regulatory state (§4.7, §10)	**(2) shared mathematical form, badged**	A first-order linear difference equation whose fixed point and contraction are theorems about the equation [Curzi, 2026a, Elaydi, 2005]. **No claim that a cell obeys it.**
The archive, catalogue and filesystem vocabulary (§4.4)	**(3) analogy, badged as register translation**	**Genomes are not programs, cells do not execute instructions in the von Neumann sense, and every one of these analogies breaks under pressure.** They are used for legibility and carry no argument; Anderson [1972] and Cartwright [1999] are the standing statement of why a borrowed formalism does not import its subject.

Borrowing	Category	What licenses it
The framework's essence law and the constant `k` (§6.4)	**definition, with a postulate attached — and here it is REJECTED**	`k = (9/8)^2/phi^2 = 0.483425732988414` is an exact ratio of exact quantities and is **defined**; **attaching hertz to it is a separate and stronger claim and is POSTULATED**. In this paper the law is not merely badged but **rejected for the task**, by its own theorem (Section 6.4). The standard critique of golden-section claims [Markowsky, 1992] is cited here deliberately and again wherever `phi` appears; the identity `phi^2 = phi + 1` is Euclidean [Euclid, 1956, Vajda, 1989].
The surplus functional and its effective count (§4.7)	**(3) analogy, badged, and named**	**`N_eff = 1 + (N−1)u` is not a Hill number.** It is a linear interpolation between 1 and `N` and is this framework's own construction [Curzi, 2026]; `1/HHI` and `exp(H)` are Hill numbers of orders two and one and are the grounded objects [Hill, 1973, Jost, 2006]. **Wherever an effective count appears in this paper it is this one, and it is named.**

Most of the borrowings in this corpus are category 2, and category 2 is respectable. The failure this table exists to prevent is a category-2 or category-3 use collecting the certainty of a category-1 one.

31.3.5 3.5 The two-column apparatus

The corpus pairs each substantive claim with a measured reference and a pre-modern textual one. **The refinement that governs this series must be restated at every use: the evidence column carries EVIDENCE and the precedent column carries PRECEDENT, and two sources agreeing does not double the evidence when only one of them is evidence** [Curzi, 2026w]. That paper prices a single traditional record at 1.585 bits against 9.966 for a replicated result, computes the collapse of agreement between correlated sources, and rejects the claim that a specific audio frequency repairs DNA on wavelength, energy and octave-separation grounds. **This paper's own rejections derive from that one, and its sonification pipeline is bound by it: a transposed tone is a label and not a resonance.**

Where the precedent column is empty, this paper says so rather than filling it, and every precedent item is credited in bits rather than summed with the measured column.

31.3.6 3.6 Identification, reflexivity, and what this paper's arithmetic does and does not rest on

Nothing in this paper estimates a social parameter. **The framework it belongs to does, and the inheritance is stated here rather than buried.** The corpus's central formal object is a reflexive loop in which a group influences its members and the members constitute the group; that is a linear-in-means reflexive specification, and Manski [1993] proves the endogenous social effect is **not identified from observational data**, with Shalizi and Thomas [2011] establishing the same confounding for network data. **No remedy is adopted anywhere in this series** — not randomisation, not an instrument, not partial-population variation, and not a defended stipulation of the reference group [Brock and Durlauf, 2001] — **so every framework claim resting on the loop is a modelling position rather than an identified empirical result.**

The specific inheritance here is the storage-law reading of regulatory state in Section 4.7 and the surplus-functional reading of Constraint L67. **Both are model-form claims, both are**

badged, and neither is identified from anything measured in this paper. The corpus's own treatment of the contraction mapping is the bound to read alongside them: **a human life is not a metric space, nobody has measured the relevant contraction modulus, and the theorem therefore transfers as shape rather than as guarantee** [Curzi, 2026].

The Lucas critique applies to the experimental programme of Section 5 in a specific and unusual way. Lucas [1976] holds that parameters estimated under one regime do not survive the announcement of another, and Hacking [1995] makes the point about human classifications. **The preregistered experiment of Section 9.2 involves human operators who will know the hypothesis**, and the corpus's own falsification audit measured a demand-characteristic confound with a true effect of zero passing an entire four-baseline suite **52.7 per cent of the time, for which there is no statistical remedy — only blinding, channel separation and a yoked control** [Curzi, 2026u]. **That is why the blinding requirement in Section 9.2 is not boilerplate.**

31.3.7 3.7 A gap in the bibliography, disclosed

The entire biological literature this paper describes has no corresponding cluster in the bibliographic database governing this corpus. That includes: the origin and critique of the "junk DNA" designation; the genome-annotation and biochemical-activity surveys and the constraint estimates that contradict them; the C-value and genome-size literature; regulatory-element classes; operon logic; repetitive-sequence regulatory theory and its confirmation; transposon-derived enhancer networks; syncytin and its knockout; nuclear transplantation and induced pluripotency; transcription-factor censuses; nucleotide data storage and its codes; guide specificity, self-processing arrays, catalytically dead nucleases; mechanically gated channels; force-to-nucleus coupling; sonogenetics; the microwave auditory effect; infrared band-centre compilations for nucleic acids; auditory filter bandwidth formulae; and the differential-expression power formula used in Section 5.11.

Every such work is named in the text with its journal, volume, page and year, and carries no citation key. Attaching an unrelated key to a biological or psychoacoustic claim would be worse than the gap, because a reviewer in that field will check it and a mis-supporting citation destroys more than a missing one. **Appendix A lists them; Section 12 records it as limitation L-1; and the practical consequence is that this paper's biology is unverifiable from its reference list — a reader must go to the named sources directly.** The same gap is disclosed in the neighbouring papers of this series [Curzi, 2026w,x] and in the sonification protocol [Curzi, 2026].

One consequence must be stated even though the gap is the defect. Each biological result is named as a single published finding, and the required context for any such citation is the measured replication base rate: large-scale replication returns roughly a third of original findings with effect sizes about halved [Open Science Collaboration, 2015], and most published findings in low-prior fields are false [Ioannidis, 2005]. **The results relied on here are among the most heavily replicated in their fields, several carry Nobel Prizes, and two carry loss-of-function experiments — but the reader is entitled to that base rate beside them.**

31.3.8 3.8 Honesty clauses and standing refusals

- **No clinical, medical or therapeutic claim appears anywhere in this paper.** Every experiment specified is *in vitro*, on cultured cells, with transcript abundances as endpoints.
- **No claim that genomic sequence is edited, altered or damaged.** The pipeline of Section 6 carries no nuclease, no template and no effector protein; **the output of its final stage is a waveform and its intermediate is a string in memory.** It cannot edit. **Cannot, not does not.**
- **No claim that any capability has been conferred on anything.**
- **Absence of a known mechanism is not proof of impossibility, and impossibility is not claimed.** Several real physical–biological channels were discovered late and were initially dismissed; Section 5.6 gives one. **The position taken is that the mechanism is un-**

known and that the burden of demonstration is entirely on the claimant, which is the normal position for any unestablished claim and carries no special hostility.

- **Nor is absence of a known mechanism a licence.** The symmetric error treats "not proven impossible" as partial support. **It is not support of any kind. A claim with no mechanism and no demonstration has zero evidential standing, and the correct number of qualifiers to attach before demonstration is not "few" but "it is not asserted"** [Popper, 1959, Platt, 1964].
- **No subliminal content, ever.** Nothing masked, nothing below threshold, nothing concealed; every symbol in a rendered signal is audible, disclosed and documented in a published mapping table. **A message hidden from its recipient cannot be consented to**, and the corpus states that as a first law rather than a caveat, on an argument that depends on no empirical claim [Curzi, 2026ad]. It is the same principle that makes a system's security rest on its design rather than on concealment [Kerckhoffs, 1883, Brandeis, 1914], and the corpus states it structurally elsewhere: **channel separation is strengthened by disclosure, because its effectiveness is entirely independent of whether anyone knows it is in place, so every element must be publishable and non-disclosure is refused as a design premise rather than merely discouraged** [Curzi, 2026].
- **Presence is not function.** Presence of sequence is not evidence of function; transcription is not evidence of function; protein occupancy is not evidence of function. **A function claim requires either evidence of selective constraint or a loss-of-function experiment, and must name which.**
- **Structural correspondence is not evidence of mechanism.** It licenses legibility and nothing else.
- **Scale type.** Every quantity in Sections 5 and 6 is a physical or combinatorial measurement on a ratio scale with a true zero, so the arithmetic on them is licensed. **The framework's psychological constructs are not**, Michell [1997]'s objection applies to them, and the additive conjoint conditions that would license ratio-scale treatment [Luce and Tukey, 1964, Krantz et al., 1971] have not been tested; the typology is Stevens's [Stevens, 1946]. **The status vocabulary of Section 3.1 is ordinal at best and no arithmetic is performed on it.**

31.4 4 The addressable archive

31.4.1 4.1 Retiring "junk", and why the retirement is the model's precondition

The term was introduced in 1972 in a paper whose actual argument was subtler than the phrase it produced. **The reasoning was a mutational-load argument: a genome cannot maintain unlimited functional sequence against mutation at a fixed rate, therefore most of a large genome cannot be under selection. That argument is sound and has not been overturned.** What has been overturned is the inference drawn from the word — that the non-coding fraction is *empty*, that there is nothing there, that it is filler.

That distinction is the entire content of the correction, and this architecture depends on it in a way worth stating plainly rather than sliding past.

Consider what "junk" would mean for an architecture whose only operation is addressing. If the non-coding fraction were genuinely junk — random, unstructured, unread, without regulatory consequence — **an address pointing into it would resolve to nothing.** There would be no latent content to reach. **A system that can only address and cannot import is worthless if the archive is empty. The junk hypothesis, in its strong popular form, is not a mild embarrassment to this architecture; it is a refutation of it.**

So the correction does not weaken the model by removing an obstacle. **It is the model's precondition**, and it is a strengthening precisely because it was not assumed: it was resisted for forty years and then measured.

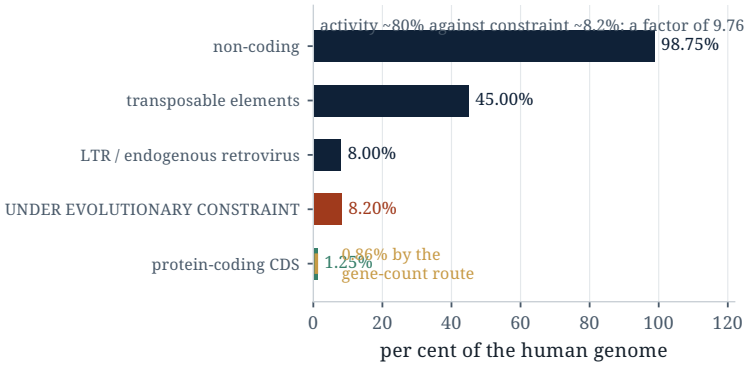


Figure 31.1: **The fractions that support the architecture and the fraction that constrains it, printed together as rule S67-1 requires.** The non-coding, transposable-element and endogenous-retroviral fractions of the human genome against the approximately 8.2 per cent under evolutionary constraint and the approximately 1.25 per cent protein-coding fraction, with the independent gene-count estimate of 0.86 per cent marked. **The span between "biochemically active" and "under constraint" is a factor of 9.76 (`80 / 8.2`) and is a live dispute rather than a settled figure.** Fractions are annotation-dependent (Section 3.3). Computed from the fractions of Section 4.2; generated by `figs/src/CCP-031_fig01.py`.

31.4.2 4.2 The arithmetic, with the pairing rule applied

```
haploid human genome                    3.1e9 bp    (nominal)
information capacity at 2 bits/bp       6.2e9 bits  = 7.75e8 bytes = 775 MB decimal
                                                     = 739.10 MiB

class                                    fraction     Mbp      MB at 2 bits/bp
-------------------------------------------------------------------------------
protein-coding exons (CDS)                 0.0125     38.8            9.69
non-coding total                           0.9875   3061.2          765.31
  transposable elements (all)              0.4500   1395.0          348.75
    LINEs                                  0.2050    635.5          158.88
    SINEs                                  0.1300    403.0          100.75
    LTR / endogenous retrovirus            0.0800    248.0           62.00
    DNA transposons                        0.0300     93.0           23.25
  introns                                  0.3600   1116.0          279.00

ratio of non-coding bytes to coding bytes             79.0 : 1

UNDER EVOLUTIONARY CONSTRAINT              0.0820    254.2           63.55
```

The last row is not optional and is printed with the others by rule S67-1.

A sanity check from an independent direction: approximately 19,900 protein-coding genes at a mean coding-sequence length near 1,340 bp gives 26.67 Mbp, which is 0.860 per cent of the genome. The 1.25 per cent figure above is the larger, more inclusive convention. **The two estimates bracket the truth between roughly 0.9 and 1.5 per cent, and nothing here turns on which end is taken, because the ratio that matters is between 65:1 and 110:1 either way.**

The transposable-element fraction near 45 per cent is a lower bound that keeps rising as detection of degraded ancient elements improves, which is itself informative about how much of the genome is old text rather than no text.

31.4.3 4.3 The overclaim, refused in print

"Not junk" does not mean "all functional", and this framework does not assert that it does.

Reported biochemical activity across most of the genome is a measurement of *activity*, and biochemical activity is cheap. Pervasive transcription is real and was measured carefully, but **a polymerase initiating spuriously produces transcript, and transcript is not function.** A transcription factor with a six-base recognition motif will bind a 3.1-billion-base genome at an enormous number of positions by chance alone; **occupancy is not regulation.**

The counter-argument that has never been answered — and stating it is a requirement — is the C-value problem, in its sharpest modern form as the onion test:

organism	genome (Gbp)	x human
Fugu rubripes (pufferfish)	0.39	0.13
Homo sapiens	3.10	1.00
Allium cepa (onion)	16.00	5.16
Protopterus aethiopicus (marbled lungfish)	130.00	41.94
Paris japonica	148.00	47.74

If every base of the human genome is doing something necessary, what are the onion's extra 12.9 billion bases doing, and why does a pufferfish manage vertebrate development on an eighth of the human total? Any theory that assigns function to bulk non-coding sequence *because it is there* must answer this, and no such theory has.

A discrepancy in the source, reported rather than harmonised. The source manuscript states that "the measured span between the two framings is a factor of about 10.7". Recomputing from the two figures it gives — approximately 80 per cent reported as biochemically active against approximately 8.2 per cent estimated as under constraint — **the ratio is `80 / 8.2 = 9.756`, not 10.7.** A factor of 10.7 would require either 87.7 per cent activity or 7.48 per cent constraint, and neither figure appears in the source. **The corrected figure of 9.76 is used throughout this paper. No conclusion changes: the pairing rule is required at either value, and the point of the rule is that the two numbers measure different things rather than that their ratio takes any particular value.**

The position taken here, stated as a position and not as a finding: the non-coding fraction contains a large, structured, consequential regulatory apparatus **and** a large quantity of decayed and duplicated sequence under little or no selection, interleaved. **An archive that contains both a catalogue and a great deal of unread material is still an archive.** Libraries are like this. **The claim this architecture needs is not that every base is functional; it is that there is addressable latent regulatory content, which is a far weaker claim and is established.**

31.4.4 4.4 What is actually in the non-coding fraction

Element	Operation	Register translation
Promoter	Where transcription initiates	Entry point
Enhancer	Increases transcription of a target, often at distance and orientation-independently	Remote call site
Silencer	Decreases transcription of a target	Guard clause
Insulator / boundary element	Blocks enhancer–promoter contact across it; organises domains	Scope delimiter

Element	Operation	Register translation
lncRNA	Transcribed, not translated; some act in *cis* on chromatin state	Resident process rather than stored data
miRNA	~22 nt; post-transcriptional repression of matching transcripts	Output filter applied after the read
Transposable element	Mobile or formerly mobile sequence; a major source of new regulatory sites	Imported library, mostly deprecated, occasionally load-bearing

The right-hand column is a register translation and carries no claim beyond legibility. Genomes are not programs, cells do not execute instructions in the von Neumann sense, and every one of these analogies breaks under pressure (Section 3.4). They are here because the readers of an engineering paper think in these terms; **they are not here as an argument.**

Two structural facts bear directly on addressing.

First, operon logic was the demonstration that a gene's expression is controlled by sequence elements distinct from the gene itself. That is now sixty-five years old and standard, **and it is the origin of the separation this architecture depends on: address and content are different objects.**

Second, repetitive non-coding sequence was proposed as the substrate of gene-regulatory networks before the sequence data existed to test it, and the proposal was substantially right. The mechanism — transposable elements distributing copies of a binding motif across a genome, so that a single factor coordinates a dispersed set of targets — is now documented, including endogenous retroviruses supplying inducible enhancers across an immune network and transposon-mediated rewiring in the evolution of a mammalian gene network.

This is the multicast structure of the addressing architecture appearing in the genome as a matter of fact rather than as a design proposal: one motif, many addresses, one signal, coordinated response [Curzi, 2026x]. **State the direction of the argument carefully. The genome having this structure does not show that an acoustic signal can exploit it. It shows that the structure being exploited is the structure that is there.**

Precedent, priced. The correspondence table read as a finite relation — a single term addressing a dispersed set of entries across categories — records that the *form* "one key, many entries" was recognised as a useful organisation of knowledge independently and repeatedly. **It is precedent only, worth of order 1.6 bits on the corpus's scale, and it says nothing about genomes** [Curzi, 2026w,?]. One-to-many addressing arises wherever a system must coordinate dispersed elements with a compact signal — in indexes, filesystems, multicast networking and correspondence tables. **The generality is real; it is a fact about the class of problem, not a shared cause.**

31.4.5 4.5 Syncytin: the strongest single item

Everything above is quantitative and general. This is specific.

Syncytin-1 is a gene essential to the formation of the human placenta. It is a captured retroviral envelope gene. The measured record, named without keys per Section 3.7: the protein is the envelope glycoprotein of a human endogenous retrovirus, expressed in placental syncytiotrophoblast and **fusogenic** — it fuses cell membranes, which is what a retroviral envelope protein does on entry and what is required to build the multinucleate layer at the maternal–foetal interface (Mi et al., *Nature* 403:785–789, 2000). A genome-wide screen identified a second such gene from a different and older retroviral family (Blaise, de Parseval, Bénit and Heidmann, *PNAS* 100:13013–13018, 2003). **Knockout of the mouse orthologue produces**

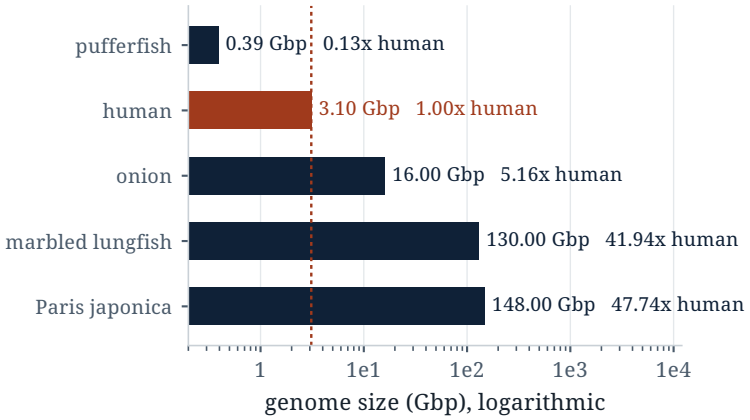


Figure 31.2: **The objection the framework cannot answer, printed at full strength.** Genome size across five species on a logarithmic axis: a pufferfish completes vertebrate development on 0.39 Gbp, an onion carries 5.16 times the human total and a marbled lungfish 41.94 times. **If bulk non-coding sequence is functional because it is present, these ratios require an explanation and none has been given.** The human value is marked as the reference. Values are standard published genome sizes, named without citation keys per Section 3.7. Generated by `figs/src/CCP-031_fig02.py`.

defective placental architecture and embryonic lethality (Dupressoir et al., *PNAS* 106:12127–12132, 2009). This is the causal test, and it is why the item carries the weight it does: the gene is not correlated with placentation; removing it stops placentation. And syncytin-like genes have been captured *independently* in multiple mammalian lineages, from different retroviruses, for the same function (Lavialle et al., *Philosophical Transactions of the Royal Society B* 368:20120507, 2013).

```
syncytin-1 env ORF   ~538 aa  ->  ~1617 bp  ->  3234 bits  ->  404 bytes
fraction of the human genome that is this one ORF          5.216e-07
total LTR/ERV-derived fraction                    8%   =  248 Mbp  =  62.0 MB at 2 bits/bp
```

Four hundred and four bytes of inherited viral text, out of sixty-two megabytes of inherited viral text, and mammalian live birth as we have it does not occur without them.

It establishes: that sequence of foreign origin, integrated into a germline, carried for millions of years and *not* functioning as a virus, can be read out later as an essential function of the host. **That is documented molecular biology with a causal knockout behind it**, and it is the cleanest existence proof available that the genome is a repository of inherited text whose current use is not its original use.

It does not establish: that any given piece of non-coding sequence has such a use; that capture is common — it is rare, and the great majority of the 8 per cent is decayed and non-functional; that anything can be induced to be read out on demand; or anything whatsoever about acoustic signals. **Syncytin is one item. It is presented as one item. It is the best item.**

Precedent, priced. The category "present but not currently expressed, and requiring only the right prompting to become expressed" was formulated, argued over and preserved for twenty-four centuries — in Plato's doctrine of recollection and in Aristotle's distinction between potentiality and actuality. **That is precedent only. It supplies no mechanism and it is not evidence about placentas**, and no inference runs from it to any biological fact [Curzi, 2026w].

31.4.6 4.6 Latency and its activation: the information accounting

The syncytin case shows latent content in evolutionary time. **The reprogramming case shows it in experimental time, and it matters more here because it is an *intervention* rather than an observation.**

Nuclei taken from the intestinal epithelium of feeding tadpoles — differentiated cells, committed to a fate — were transplanted into enucleated eggs, and normal feeding tadpoles developed (Gurdon, *Journal of Embryology and Experimental Morphology* 10:622–640, 1962). **The nucleus of a differentiated cell retains the full genomic information required to build the whole organism. Differentiation is not subtraction.** Four decades later, mouse embryonic and adult fibroblasts were returned to a pluripotent state by forced expression of four transcription factors (Takahashi and Yamanaka, *Cell* 126:663–676, 2006), extended to human fibroblasts the following year; the 2012 Nobel Prize in Physiology or Medicine was awarded jointly for the two results, fifty years apart.

The point that matters here is the information accounting, and it is computed rather than asserted.

```
approximate number of human sequence-specific transcription factors     ~1600

selecting a set of k factors from 1600:
    k = 1   C = 1.600000e+03    log2 =  10.644 bits
    k = 2   C = 1.279200e+06    log2 =  20.287 bits
    k = 3   C = 6.813872e+08    log2 =  29.344 bits
    k = 4   C = 2.720438e+11    log2 =  37.985 bits
    k = 5   C = 8.683639e+13    log2 =  46.303 bits
    k = 6   C = 2.308401e+16    log2 =  54.358 bits

the four-factor SELECTION specifies at most           37.985 bits
the genome it acts upon holds                     6.200e+09 bits
ratio                                             1.632e+08 : 1
```

Identifying which four of roughly sixteen hundred transcription factors to use is a choice worth at most about thirty-eight bits. The genome those four factors act on holds six point two billion bits. The instruction is one hundred and sixty-three million times smaller than the thing it reorganises.

That is the shape of the architecture in the only units that make it checkable. An activation signal is not the content it activates, it is smaller than the content it activates by many orders of magnitude, and this is not a weakness of the model — it is what an address is [Curzi, 2026x].

Two qualifications, both mandatory.

First, the efficiency is low. Published reprogramming efficiencies for the original four-factor protocols are commonly quoted between roughly 0.01 and 0.1 per cent of input cells, with wide variation by cell type, delivery method and scoring criteria:

efficiency	cells needed per colony	cells needed for 100 colonies
0.010%	10,000	1,000,000
0.100%	1,000	100,000
1.000%	100	10,000

A process that works on one cell in a thousand is a real process and a poor one, and any model that borrows its authority must also borrow its efficiency. This framework borrows the structural point — that regulatory input reorganises expression without adding sequence — and borrows no implication about how readily that happens.

Second, in the original protocols the factors were delivered by retroviral transduction, which does integrate sequence into the genome. That must be said in the same breath as the claim that "no new information is added", or the claim is misleading. **The precise statement is:**

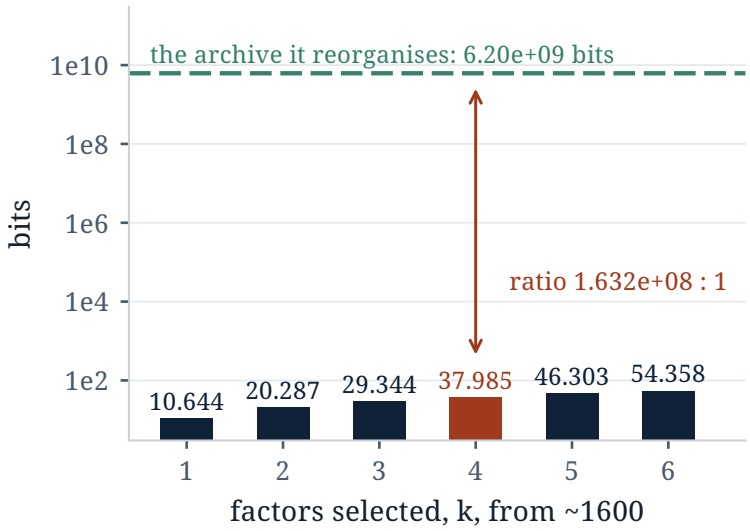


Figure 31.3: **The instruction is a hundred and sixty-three million times smaller than the archive it reorganises.** Bits required to specify a selection of `k` factors from approximately 1,600 sequence-specific transcription factors, `log2 C(1600,k)`, against the `6.2e9` bits of a haploid genome at two bits per base. The four-factor selection is 37.985 bits, marked; the ratio to the genome is `1.632e8` to one. **The vertical scale is logarithmic and the two quantities are eight orders of magnitude apart, which is the point rather than an artefact of the display.** Computed from the closed form of Section 4.6; generated by `figs/src/CCP-031_fig03.py`.

the information that specifies the new state is the identity of the four factors, not new genomic content; the pluripotency programme itself was already encoded in the fibroblast's genome and was already there in the transplanted nuclei. Later work using non-integrating delivery removes the qualification, but the original result carries it and it is stated rather than elided.

31.4.7 4.7 Constraint L67, and the reason it is the most useful sentence here

> **CONSTRAINT L67. Addressing can reach only what is already latent. No addressing operation, of any kind, at any signal strength, by any modality, can cause a cell to express content that its genome and its accessible regulatory repertoire do not already contain.**

This follows directly from the addressing architecture and requires no new assumption [Curzi, 2026x]. **An address is a pointer. A pointer with no cargo cannot deliver content. If the referent does not exist, the address resolves to nothing.** The architecture has no import path by construction, and the absence of an import path is the same fact as the guarantee that no sequence is changed. **The property that makes the architecture safe is the property that makes it limited. You cannot have one without the other.**

The constraint does work in four directions.

It bounds the claim space. Any statement of the form "this method could produce capability X" is false whenever X is not already encoded and reachable. **That eliminates the entire class of claims that would make this work pseudoscience — no new genes, no acquired traits, no capacities imported from outside the organism, nothing conferred.**

It makes the model falsifiable in a specific way. If a demonstration ever showed a cell expressing content demonstrably absent from its genome and regulatory repertoire following an addressing-only intervention, **this architecture is wrong, and wrong in a way no adjustment of parameters can repair.** That is the good kind of falsifier: one named observation

ends the model [Popper, 1959, Platt, 1964].

It generates a positive prediction. If addressing reaches only what is latent, the *reachable set* is bounded and enumerable in principle: for a given cell type it is the set of expression states accessible from its current chromatin configuration under some combination of regulatory input. That set is finite, is studied under other names, and is measurable.

It supplies the audit rule. Every proposed application is checked against the reachable set before anything else. **If the target state is not in the reachable set, the proposal is dead at step one and no experiment is needed** (Section 9.1).

The storage-law reading of the constraint is exact and worth writing down, and it is a badged model-form claim. The archive is `Q`, and an addressing operation does not increase it. What such an operation can change is which part of `Q` is being read — a change in `u`, the non-parallelism of what is currently active, not a change in `Q` [Curzi, 2026a,?]:

```
N_eff = 1 + (N−1)u        F = ln N_eff

N = 20,000 addressable loci

u = 0.01   N_eff =      201.0   F = 5.3033 nats =  7.6510 bits
u = 0.05   N_eff =    1,001.0   F = 6.9087 nats =  9.9672 bits
u = 0.25   N_eff =    5,000.8   F = 8.5173 nats = 12.2879 bits
u = 0.50   N_eff =   10,000.5   F = 9.2104 nats = 13.2878 bits
u = 1.00   N_eff =   20,000.0   F = 9.9035 nats = 14.2877 bits

hard ceiling  ln N = 9.9035 nats = 14.2877 bits
```

The ceiling is `ln N` and `N` is set by the archive, not by the signal. No amount of input raises `ln N`. That is the corpus's own bound `0 <= F(u) <= ln N` doing exactly the work it was proved to do, in a domain it was not written for [Curzi, 2026]. **The effective count used is `N_eff = 1 + (N−1)u`, which is this framework's own construction and is not a Hill number** (Section 3.4).

Speculative bridge, marked speculative in the sentence that states it: this framework's author proposes, as a postulate and not as a finding, that the reachable set of a differentiated human cell is substantially larger than current practice exploits, and that most of the gap is a matter of not knowing the addresses rather than of the content being absent. It predicts that the rate of discovery of new transdifferentiation protocols will continue rather than saturate, which is a weak but real test, and it has no other support here.

31.4.8 4.8 A deflation, performed on the framework's own tradition

The old corpus offers one correspondence here that is tempting, and it is declined, because the concordance method is only worth having if it can reject something visibly [Curzi, 2026w].

The genetic code has 64 codons. A well-known divinatory system has 64 hexagrams. **The coincidence has been noticed repeatedly and elaborated at length, and it is worth exactly this much:**

```
2^6 = 64        six binary lines, two states each      (the hexagram)
4^3 = 64        three positions, four bases each        (the codon)
```

Sixty-four is a small highly-composite power. That two systems both land on it is a fact about the number, not a fact about either system. The factorisations differ — `2^6` against `4^3` — which is to say the two structures are not the same structure at all; they are two different structures with the same cardinality. **And a proposed correspondence between them is a *chosen bijection*, one of `64! = 1.2689e89` available, worth 296.0 bits of arbitrary choice.**

Status: REJECTED as a bridge. C1 fails: there is no deep shared pattern, only a shared count. C2 fails: the correspondence generates no prediction that could come out otherwise.

C3 fails: it is not usable on either side. **The concordance method is only worth having if it can return this verdict, and this is the verdict.**

For contrast, the genuine arithmetic of the genetic code, which does carry content:

```
codons                                  64      log2 = 6.0000 bits per codon (max)
outputs (20 amino acids + stop)         21      log2 = 4.3923 bits
redundancy  1 − log2(21)/log2(64)          =    0.2679
mean codons per amino acid (61/20)         =    3.05
```

The code discards about 26.8 per cent of its addressing capacity to degeneracy, and that discarded capacity is error tolerance. That is a real engineering fact about a real code, computed rather than asserted [Shannon, 1948, Cover and Thomas, 2006].

31.5 5 The open problem

31.5.1 5.1 The gap, stated four ways

The same proposition at four levels of precision, because each level is the one someone will need.

Colloquial. Nobody knows how a cell could read a message out of a sound.

Engineering. No known biological structure performs demodulation of a symbol stream from an acoustic carrier into a sequence-specific output. There is no identified transducer with the required frequency selectivity, no identified memory element that preserves symbol order across the symbol interval, and no identified mapping from a decoded symbol string to sequence-specific binding at a genomic locus.

Formal. Let `s(t)` be an acoustic pressure waveform encoding a symbol string `x = (x_1, ..., x_n)` over an alphabet of size 4, and let `R` be the cell's transcriptional response. **There is no established mechanism producing an `R` that depends on `x` as an ordered string**, as distinct from depending on functionals of `s(t)` that are invariant under permutation of the symbols — total energy, duration, long-term power spectrum, envelope statistics. **The claim that requires a mechanism is the claim that `R` is a function of the ORDER, and that is exactly the claim for which no mechanism exists.**

Operational. Take a stimulus and its symbol-permuted twin. The two have identical total energy, identical duration, identical long-term power spectrum and identical symbol counts. **If a cell cannot be shown to distinguish them, no sequence-specific channel has been demonstrated, and nothing else in the experiment matters.**

31.5.2 5.2 What "sequence-specific" means, and the three requirements

DEFINITION. A response is **sequence-specific** with respect to an input symbol string `x` if it depends on the identity and order of the symbols in `x`, and not merely on quantities invariant under permutation of those symbols. **This is a definition, true by stipulation, and it carries no empirical weight; its function is to make the primary control non-negotiable.**

The permutation-invariant quantities of an acoustic stimulus are the ones almost every biological acoustic effect in the literature depends on:

```
total delivered energy              invariant under permutation
duration                            invariant
long-term power spectrum            invariant
RMS amplitude                       invariant
peak amplitude / crest factor       invariant (with matched symbol shapes)
symbol histogram (counts of A,C,G,T) invariant
```

Everything a stressed cell is known to respond to is on that list. Heat, shear, membrane tension, cavitation, streaming, acoustic radiation force — every one is a response to bulk energy delivery and every one is blind to order by construction. **This is why the permuted-symbol control is not one control among several. It is the experiment.**

Demodulation of an ordered symbol string requires exactly three things. **There is no way around any of them; they are the definition of the operation.**

D1 — SYMBOL DISCRIMINATION. Something must produce different internal states for different symbols. Under a pitch-coded scheme that means frequency selectivity sharp enough to separate the symbol tones.

D2 — ORDER MEMORY. Something must retain the identity of symbol `i` while symbol `i+1` arrives, and must retain the accumulating string across the whole address. **A device that responds to each symbol and forgets it is a level meter, not a receiver.** This is the requirement most often skipped in speculative accounts, and it is the hardest one.

D3 — SEQUENCE-SPECIFIC OUTPUT. The decoded string must be converted into an operation specific to a genomic locus determined by that string. In a guide-molecule architecture this step is trivially solved *because a physical guide exists whose bases pair with the target.* **With no molecule delivered, the mapping from "a decoded string held in some cellular state variable" to "binding at the locus that string names" has no proposed physical implementation at all.**

D3 is the deepest of the three and it is worth being blunt about why. Base pairing is sequence-specific because the guide and the target are made of complementary chemistry. **Remove the guide and you have removed the only known physical thing that makes an arbitrary 20-symbol string correspond to one place in a genome. Nothing else in known cell biology takes an arbitrary string as input and produces locus-specific action as output.** The separability of addressing from editing is correct [Curzi, 2026x] — **but the separation it exploits was demonstrated with a molecule present, and this method has no molecule.**

That is the gap. It is not "we do not know the details." It is D3.

31.5.3 5.3 Does the signal arrive as structure? The physics, computed

Before asking whether anything decodes the signal, ask whether the signal reaches the cell in a form that has structure to decode. **Two computations settle the shape of it, and the first is unfavourable.**

The wavelength objection. For a plane wave with wavenumber `k = 2 pi f / c`, the fractional pressure difference across an object of size `d` is approximately `k d`. In soft tissue at `c = 1540` m/s and `d = 10` µm (both nominal):

freq (Hz)	wavelength (m)	wavelength/cell	dP/P across one cell
20	77.0	7.70e+06	8.1600e−07
100	15.4	1.54e+06	4.0800e−06
440	3.50	3.50e+05	1.7952e−05
1,000	1.54	1.54e+05	4.0800e−05
10,000	0.154	1.54e+04	4.0800e−04
20,000	0.077	7.70e+03	8.1600e−04
100,000	0.0154	1.54e+03	4.0800e−03
1,000,000	0.00154	1.54e+02	4.0800e−02
10,000,000	0.000154	1.54e+01	4.0800e−01

In an audible-band far field, the cell does not feel the wave. It rides on it.

Put a number on what it would take. Taking a membrane tension of 3 mN/m — the order at which mechanosensitive channels are reported to respond — and the quasi-static relation `dP = 2T/r` for a cell of radius 5 µm:

```
membrane tension   1.00 mN/m  ->  dP =   400 Pa across a 10 um cell
```

```
membrane tension   3.00 mN/m  ->  dP =  1200 Pa
membrane tension   5.00 mN/m  ->  dP =  2000 Pa

far-field acoustic amplitude required to produce dP = 1200 Pa:

freq (Hz)   p_required (Pa)   equivalent dB re 20 uPa    intensity in tissue (W/m^2)
      100      2.9412e+08              263.3                      2.8835e+10
      440      6.6845e+07              250.5                      1.4894e+09
    1,000      2.9412e+07              243.3                      2.8835e+08
   10,000      2.9412e+06              223.3                      2.8835e+06
  100,000      2.9412e+05              203.3                      2.8835e+04
1,000,000      2.9412e+04              183.3                      2.8835e+02
```

Two hundred and forty-three decibels at one kilohertz. That is not a loud sound; it is not a sound at all. The pressures in the upper rows exceed the inertial-cavitation threshold in water by orders of magnitude, so the medium would be destroyed long before those amplitudes were reached. **Any account of this method that imagines a cell picking a message out of ordinary airborne music, by direct far-field pressure differential across the cell body, is refuted by this table. That account is not carried here.** Compounding it, the impedance step at any air-to-tissue interface costs roughly thirty decibels of incident intensity before anything else happens [Curzi, 2026x].

Now the path that survives, because there is one and it must be reported with equal care. Cells adherent to a vibrating substrate do not ride the wave; they are shaken by their attachment points while the surrounding fluid partially lags, and the coupling is inertial and adhesive rather than far-field acoustic. Computing the density-contrast inertial force on a 10 µm cell (`rho_cell = 1050` against `rho_medium = 1000`, a 5 per cent contrast) at substrate displacement amplitude `A`:

```
cell volume 5.2360e-16 m^3, mass 5.4978e-13 kg, contrast mass 2.6180e-14 kg

f (Hz)     A (m)      accel (m/s^2)    contrast force (N)     contrast force (pN)
    50    1.0e-06        9.8696e-02          2.5839e-15                0.003
   440    1.0e-07        7.6430e-01          2.0009e-14                0.020
   440    1.0e-06        7.6430e+00          2.0009e-13                0.200
 1,000    1.0e-06        3.9478e+01          1.0335e-12                1.034
 5,000    1.0e-07        9.8696e+01          2.5839e-12                2.584
 5,000    1.0e-06        9.8696e+02          2.5839e-11               25.839
```

Single mechanosensitive-channel gating forces are of order 1 to 10 pN. **Substrate-borne vibration at micron amplitudes and kilohertz frequencies reaches that range. So there is a real coupling path, and any experiment must use it rather than pretending a loudspeaker in a room is a stimulus.**

But look at what that path delivers: a scalar, whole-cell, time-varying force. It is on the permutation-invariant list of Section 5.2 unless something downstream discriminates symbols, remembers their order, and maps the string to a locus. The physics gets the energy in. It does not get the message in. Those are different problems and only the first has a solution.

One further quantity, because it is the confound that ruins careless experiments:

```
acoustic heating   dT/dt = 2 alpha I / (rho c_p),  alpha from 0.5 dB/cm/MHz

f = 1 MHz,   I = 1e4 W/m^2  ->  2.7543e-02 K/s  =  1.6526 K/min
f = 1 MHz,   I = 1e3 W/m^2  ->  2.7543e-03 K/s  =  0.1653 K/min
f = 1 kHz,   I = 1e3 W/m^2  ->  2.7543e-06 K/s  =  0.0002 K/min
```

A one-kelvin change alters transcription in essentially any cultured cell. Temperature at the sample must be logged continuously and reported, not assumed constant because the incubator setting did not change.

31.5.4 5.4 Capacity is not the bottleneck, and saying so sharpens the gap

Take the cell's receiver bandwidth from the kinetics of the fastest well-characterised mechanical transducer available to it. Mechanosensitive channel inactivation time constants in the published range give first-order cutoffs of:

```
tau =  1.0 ms  ->  f_c = 159.155 Hz
tau =  5.0 ms  ->  f_c =  31.831 Hz
tau = 10.0 ms  ->  f_c =  15.915 Hz
tau = 15.0 ms  ->  f_c =  10.610 Hz
tau = 30.0 ms  ->  f_c =   5.305 Hz
```

Applying `C = B log2(1 + S/N)` [Shannon, 1948]:

B (Hz)	SNR (dB)	C (bits/s)	time for a 40-bit address	time for 1 kB
5	0	5.000	8.000 s	1638.4 s
5	10	17.297	2.313 s	473.6 s
5	20	33.291	1.202 s	246.1 s
16	0	16.000	2.500 s	512.0 s
16	10	55.351	0.723 s	148.0 s
16	20	106.531	0.375 s	76.9 s
100	10	345.943	0.116 s	23.7 s
1,000	10	3459.432	0.012 s	2.4 s

A twenty-nucleotide address is forty bits. At sixteen hertz of bandwidth and ten decibels of signal-to-noise that is seven hundred and twenty-three milliseconds. Even at the pessimistic end — five hertz, zero decibels — it is eight seconds. Read the other way, as the signal-to-noise ratio needed to push forty bits through in one second:

```
B =   5 Hz  ->  required SNR = 2^(40/5) - 1  = 255      = 24.07 dB
B =  16 Hz  ->  required SNR = 2^(40/16) - 1 = 4.657    =  6.68 dB
B = 100 Hz  ->  required SNR = 2^(40/100) - 1 = 0.3195  = -4.96 dB
```

None of these numbers is prohibitive. The information-theoretic room exists.

That is the point, and it is an uncomfortable point rather than an encouraging one. If the obstacle were capacity it would be a quantitative obstacle, and quantitative obstacles yield to engineering. **It is not. The obstacle is that there is no receiver — nothing performing D1, D2 and D3 — and a channel with adequate capacity and no receiver transmits exactly nothing.**

A radio wave carrying a symphony passes through a rock. The capacity of the path is enormous. The rock does not hear the symphony. The missing thing is not bandwidth.

31.5.5 5.5 What is missing, specifically

D1, symbol discrimination. With all four symbols in one register the four tones are separated by as little as 2.081 cents and span 41.034 cents overall (Section 6.6). Discriminating them requires a filter quality factor of:

```
symbol spacing    10 cents  ->  df =   4.206 Hz  ->  required Q = 172.62
symbol spacing    45 cents  ->  df =  19.118 Hz  ->  required Q =  37.97
symbol spacing   100 cents  ->  df =  43.170 Hz  ->  required Q =  16.82
symbol spacing   200 cents  ->  df =  88.907 Hz  ->  required Q =   8.17
symbol spacing  1200 cents  ->  df = 726.000 Hz  ->  required Q =   1.00

for reference, the human auditory filter at 726 Hz:
equivalent rectangular bandwidth = 103.06 Hz  ->  Q = 7.04
```

The cochlea, a purpose-built mechanical spectrum analyser with an active amplifier, has a quality factor of about seven at these frequencies. The register-spreading design of Section 6.7 exists precisely so the required `Q` drops to about one — symbols distinguished by *octave register* rather than by fine pitch.

But even a `Q` of one requires *some* frequency dependence, and here is the physical problem. A resonator at molecular scale in water is overwhelmingly overdamped. The unsteady Stokes number `beta = L sqrt(2 pi f rho / mu)` measures whether momentum diffuses across the object within one cycle:

f (Hz)	L = 10 um	L = 10 nm
100	2.5066e−01	2.5066e−04
1,000	7.9267e−01	7.9267e−04
20,000	3.5449e+00	3.5449e−03
1,000,000	2.5066e+01	2.5066e−02

At molecular scale `beta` is between `1e−4` and `1e−2` across the entire audible band and well beyond it. The object cannot store energy across a cycle. Its `Q` is far below one. That is the low-Reynolds-number regime, and it is why molecular machinery does not ring; Reynolds numbers for a 10 µm cell oscillating at kilohertz frequencies and micron amplitudes are of order `1e−3` and below.

So D1 has no molecular-scale solution and would have to be satisfied at whole-cell or whole-tissue scale, where `beta` approaches one. That is not impossible — a cell at 20 kHz has `beta = 3.54` — **but nothing is known to be built there for this purpose**, and it pushes the required frequencies out of the range the octave transposition of Section 6.5 targets.

D2, order memory. The symbol interval in Section 6.9's design table ranges from about 0.25 ms to about 200 ms. Compare that with the time a protein needs to cross a cell at a nominal diffusion coefficient of `1e−10` m^2/s:

```
L =  1 um  ->  mean 3D diffusion time L^2/6D =   1.667 ms
L = 10 um  ->  mean 3D diffusion time L^2/6D = 166.667 ms
```

At fast symbol rates a symbol is over before a signalling protein has crossed a micron. A decoder therefore cannot hold the string spatially; it must hold it in a state variable — a chemical state, a conformational state, a phosphorylation cascade with the right time constants — **and no such state machine with four-symbol resolution and roughly twenty-symbol depth has been described in any cell.** Calcium oscillation frequency coding is real and is the closest published analogue, but **it encodes a scalar rate, not an ordered string over a four-letter alphabet.**

D3, sequence-specific output. Covered in Section 5.2. **There is no proposed physical implementation. This is the one that is not a matter of degree: D1 and D2 are hard; D3 has no candidate.**

31.5.6 5.6 The near misses, and exactly why each is not the mechanism

Three results are routinely offered as though they closed this gap. **None does, and understanding why is the fastest way to understand the gap.**

Sonogenetics. Ultrasound activates specific cells. Real, published, replicated. **But read where the specificity lives: the experimenter expressed a particular channel in a particular cell type in advance. The address is in the cell, not in the sound.** The sound is a broadcast "on" signal and the selectivity is entirely a property of which cells were pre-engineered to listen. **That is the exact inverse of the architecture in question, which requires the address to be in the signal.** Sonogenetics is superb evidence that acoustic energy can drive specific cellular outcomes, and **zero evidence that a waveform can carry an address.**

Mechanotransduction to chromatin. Force at the cell surface transmits through the cytoskeleton, stretches chromatin and upregulates transcription. Real, direct and important [Curzi, 2026x]. **But the force is applied locally through a bead bound to a receptor and the response is graded with force magnitude. It is a scalar input producing a graded output. Nothing in it is order-sensitive.** It establishes the *pathway* from mechanics to transcription and says nothing about *content*.

The cochlea. Biology's proof that exquisite frequency analysis is buildable. **But it is a specialised organ with a basilar-membrane stiffness gradient spanning three orders of magnitude, an active amplifier, and a wired output to a brain performing the actual decoding.** Its existence proves that mechanical spectrum analysis is possible with purpose-built macroscopic structure. **It does not suggest that a fibroblast has any of it.**

And one case on the other side of the ledger, included because honesty requires it. The microwave auditory effect (Frey, *Journal of Applied Physiology* 17:689–692, 1962): pulsed radiofrequency energy produces an audible click, by thermoelastic expansion generating an acoustic wave in the head. **This was an implausible-sounding cross-modality channel that turned out to be real, and it was established the ordinary way — by demonstration, replication, and identification of the mechanism.** It is the standing counterexample to the claim that unexpected physical–biological channels do not exist, **and it is also the standing demonstration of what settling such a question looks like: not argument, not plausibility, but a mechanism and a measurement.**

Precedent, priced, and it is methodological rather than acoustic. The practice of recording an unresolved question formally, in the primary text, with a marker — the Talmudic *teiku*, and the apophatic tradition's discipline of stating exactly what cannot be said — records that the device was invented independently and maintained for centuries by people who had no statistics. **A body of knowledge with a formal slot for "unresolved" degrades more slowly than one without, because unresolved questions that lack a marker do not survive transmission.** The modern measurement of the same effect is the reform literature [Ioannidis, 2005, Open Science Collaboration, 2015, Nosek et al., 2018]. **The precedent column carries nothing about acoustics or cells and is not offered as though it did** [Curzi, 2026w].

31.5.7 5.7 What an experiment would have to demonstrate

Five conditions. **All five are necessary; the first decides everything.**

E1 — ORDER SENSITIVITY. The response to a stimulus differs from the response to its symbol-permuted twin, matched in total energy, duration, long-term power spectrum, symbol counts, envelope shape and crest factor. **Without E1 there is no sequence-specific channel, whatever else is observed.**

E2 — ADDRESS SPECIFICITY. The response occurs at a preregistered genomic locus determined by the address prefix, and does not occur at a matched control locus.

E3 — ADDRESS SWAP. Changing the address changes which preregistered locus responds, within the same experiment, cells, operator and session. **This is the condition no nonspecific stress response can imitate.**

E4 — DOSE-RESPONSE AND FIDELITY DEPENDENCE. The effect scales monotonically with a stimulus parameter under experimenter control, and degrades as the address is corrupted by a controlled number of symbol substitutions.

E5 — INDEPENDENT REPLICATION. A second laboratory, working from the written protocol and the released stimulus files, obtains the same result without contact with the originating group during the run.

Note what is not on that list: any change to any DNA sequence. E1 to E5 are entirely about which existing sequence is transcribed. **Nothing in the specified experiment could edit, alter or damage a genome, and nothing in it is proposed as having any clinical relevance.**

31.5.8 5.8 The control set

An experiment of this kind is defined by its controls, and a claim of this kind will be judged entirely on whether the following are present. Each is listed with the specific alternative explanation it excludes.

#	Control	Excludes
C1	**Scrambled-sequence control.** Same symbol tones, counts, duration and envelope; order permuted.	Any response to bulk energy, spectrum or duration. **This is the primary control.**
C2	**Shuffled-input control.** Pairing between input data and measured output destroyed by permutation across runs.	Any apparent input–output relation arising from analysis flexibility.
C3	Scrambled-address control: address prefix replaced by a 20-mer with no genomic match, payload identical.	Response driven by the payload or by exposure as such.
C4	Energy-matched noise: same power spectral density envelope, no symbol structure.	Response to broadband mechanical stimulation.
C5	Sham exposure: transducer mounted and clamped, coupling identical, drive off.	Handling, incubator vibration, mounting stress, time out of incubator.
C6	Thermal control: matched temperature trajectory without acoustic structure, continuously logged at the sample.	Acoustic heating (Section 5.3).
C7	Cell-free control: full protocol on medium alone.	Instrument artefacts in library preparation and sequencing.
C8	Positive control: transcriptional activation at the same preregistered locus by an established method.	An assay too insensitive to detect locus-specific activation at all.
C9	Retrograde control: the stimulus played time-reversed.	Response to spectrum plus gross envelope, which reversal preserves and order does not.
C10	Blinding and randomisation: stimulus assignment blinded to operator and analyst until the analysis is frozen.	Expectancy effects on handling and on analytic choices.

C1 and C2 are the two that get skipped. C1 is skipped because building a properly matched permuted stimulus is fiddly. C2 is skipped because it feels redundant once C1 is in place, **and it is not: C1 controls the *stimulus*, C2 controls the *analysis*.** A pipeline with enough free parameters will find structure in a correctly controlled stimulus set, and only C2 catches that [Curzi, 2026i, Simmons et al., 2011, Gelman and Loken, 2013].

And C2 must be calibrated rather than trusted, which is a correction this framework applies to itself. The corpus measured its own shuffled-input control against a nominal 0.05 across 500 null trials and found **iid permutation rejecting at 0.108, circular shift at 0.120, and only block permutation near nominal at 0.060** [Curzi, 2026u]. **The control is anti-conservative as ordinarily run — it rejects true nulls at more than twice its nominal rate — so any `p` it returns in this design is optimistic by an unknown factor, and block permutation is specified in Section 9.2 for that reason.**

C8 deserves a sentence of its own. Without a positive control a null is uninterpretable: it cannot be distinguished from an assay that could not have detected a real effect either. **An experiment reporting a null without a positive control has reported nothing.**

C10 is not boilerplate here. The same audit measured a demand-characteristic confound — subjects or operators inferring the hypothesis — passing an entire four-baseline suite **52.7 per cent of the time with a true effect of exactly zero, and found no statistical remedy for it, only design remedies** [Curzi, 2026u]. **Blinding is the design remedy and it is mandatory.**

31.5.9 5.9 What would settle it, in each direction

SETTLED POSITIVE requires all of the following, and each clause is there because a claim lacking it has failed before:

1. E1 through E5 satisfied, with the full C1–C10 control set present and reported.
2. Preregistration of the target locus, the effect size, the analysis and the stopping rule, timestamped before data collection, published whatever the outcome [Nosek et al., 2018].
3. Effect exceeding the preregistered threshold at a multiple-comparison-corrected significance level, with the correction stated.
4. Independent replication with no personnel overlap.
5. **A proposed mechanism addressing D1, D2 and D3 — or, absent a mechanism, an explicit statement that the effect is empirical and unexplained. The second is acceptable; silence is not.**

The statistical geometry of point 3 is where such claims usually die:

```
genes tested                                20,000
Bonferroni alpha                            2.500e-06
corresponding two-sided z                   4.7081
expected false positives at uncorrected 0.05  1,000 genes
```

A hit list of a few hundred genes at uncorrected `p < 0.05` is the null hypothesis, not a result. Preregistration is what converts a twenty-thousand-way search into a one-way test:

```
preregistering  1 target gene: chance hit probability 5.000e-05
preregistering  3 target genes: 1.500e-04
preregistering 10 target genes: 5.000e-04
```

And the address-swap crossover of E3 carries the most weight per sample, because a nonspecific stress response has no way to select which of two preregistered loci responds:

```
2 addresses x 2 preregistered loci, crossover
P(correct assignment by chance per replicate)              = 0.5
  3 independent crossover replicates -> p = 0.125000
  5 independent crossover replicates -> p = 0.031250
 10 independent crossover replicates -> p = 0.000977

P(a nonspecific response hits both preregistered loci in the correct pairing)
  = (2/20000)^2 x 0.5 = 5.000e-09
```

SETTLED NEGATIVE, for the tested condition, requires:

1. Adequate power, stated in advance, against a preregistered effect size.
2. A positive control demonstrating the assay would have detected that effect size.
3. **Equivalence testing, not a non-significant p-value.** Sample sizes for equivalence at 80 per cent power, nominal dispersion 0.1, mean count 100:

```
equivalence bound FC = 1.10  ->  n >= 104 per group
equivalence bound FC = 1.25  ->  n >=  19 per group
equivalence bound FC = 1.50  ->  n >=   6 per group
```

4. Full publication of the null with the stimulus files and raw data.

A frequentist `p > 0.05` is not evidence of absence. The likelihood-ratio view makes the difference visible immediately:

```
observed log2FC = 0.00, se = 0.10  ->  LR(null : alt at 0.5) = 2.683e+05   decisive for null
observed log2FC = 0.05, se = 0.10  ->  LR = 2.203e+04                      strong for null
observed log2FC = 0.00, se = 0.30  ->  LR = 4.01                           weak, uninformative
```

The third row is what an underpowered null looks like, and it is worth almost nothing.

31.5.10 5.10 What a null would mean, including the unwelcome part

Suppose the experiment is done properly and returns a clean, well-powered, positive-controlled null.

What it would establish. That under the tested conditions — that carrier scheme, amplitude, duration, symbol rate, cell type and harvest time — no sequence-specific transcriptional response occurs at the preregistered loci above the preregistered effect size. **That is a real and publishable result and it is the outcome this framework expects, absent a mechanism for D3.**

What it would not establish, and here is the uncomfortable arithmetic.

```
carrier / register scheme      4
amplitude levels               5
exposure duration              4
symbol rate                    4
cell type                      3
time-to-harvest                3
-------------------------------------
total conditions            2,880

a single well-powered experiment covers 1 of 2,880  =  0.0347% of the space

exhaustive coverage: 2,880 conditions x 7 arms x n=21  =  423,360 libraries
```

A null over one condition bounds three hundredths of one per cent of the space, and exhausting the space is not affordable. This cuts both ways and both must be said.

Against the claimant: it means "we haven't tried the right settings yet" is always available, which is precisely what makes a research programme unfalsifiable. **A claim that can retreat into a 2,880-cell parameter space indefinitely is not a scientific claim, and the only defence is to commit in advance to the conditions under which the claim is abandoned. This framework so commits** (SP-O.9 in Section 7.2) [Lakatos, 1970].

Against the sceptic: a single null does not close the question, and treating it as though it did would be the same error in the opposite direction.

The resolution is sequential design rather than exhaustion:

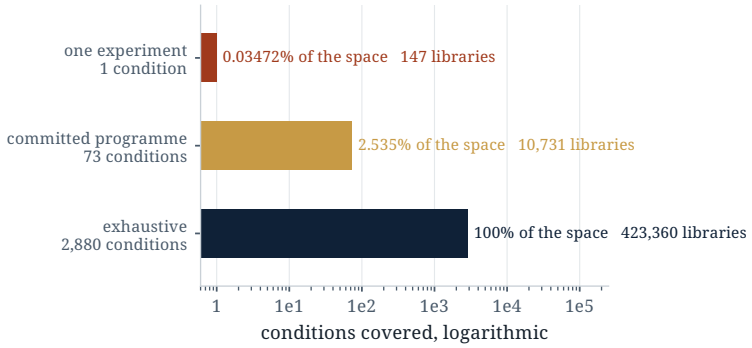


Figure 31.4: **One well-powered experiment bounds three hundredths of one per cent of the space, and the committed programme bounds two and a half.** The parameter space is `4 x 5 x 4 x 4 x 3 x 3 = 2,880` conditions; a single condition is 0.0347 per cent of it; the preregistered three-stage adaptive design of 1, then 8, then 64 conditions covers 73, or 2.53 per cent, at 10,731 libraries against 423,360 for exhaustive coverage. **The bar is drawn on a logarithmic axis because a linear one would render the single condition invisible, which is itself the point.** Computed from the factor counts of Section 5.10; generated by `figs/src/CCP-031_fig04.py`.

```
3-stage adaptive design: 1 condition, then 8, then 64
= 73 conditions x 7 arms x n=21 = 10,731 libraries = 2.53% of exhaustive
```

Seventy-three conditions, chosen adaptively with preregistered stopping rules, is a real programme at about two and a half per cent of the exhaustive cost. If nothing appears across that programme with the full control set, this framework's position is that the acoustic sequence-specific channel is to be treated as refuted for practical purposes and the pipeline of Section 6 retained only as what it demonstrably is — an encoding and sonification method.

31.5.11 5.11 What magnitude would even be detectable

The capacity analysis says the channel could carry a 40-bit address in under a second. **That tells us nothing about the *magnitude* of any downstream transcriptional change, because the map from bits received to fold change produced is exactly the unknown.** So the question is inverted: **what is the smallest transcriptional effect the experiment could see, and is it plausible that a real effect would be at least that large?**

Minimum detectable fold change at mean count 100, dispersion 0.1 and 80 per cent power, by the standard count-data formula (named without a key, Appendix A):

n per group	alpha = 0.05	alpha = 2.5e-6 (Bonferroni)
3	FC 1.710 (log2 0.774)	FC 2.894 (log2 1.533)
6	FC 1.461 (log2 0.547)	FC 2.120 (log2 1.084)
10	FC 1.342 (log2 0.424)	FC 1.790 (log2 0.840)
20	FC 1.231 (log2 0.300)	FC 1.509 (log2 0.594)
50	FC 1.140 (log2 0.190)	FC 1.297 (log2 0.376)

And the sample size required to detect a given effect:

alpha	FC	mean count	phi	n per group
nominal 0.05	1.25	100	0.10	18
nominal 0.05	1.50	100	0.10	6
Bonferroni (20k genes)	1.25	100	0.10	69
Bonferroni (20k genes)	1.50	100	0.10	21
Bonferroni (20k genes)	2.00	100	0.10	8

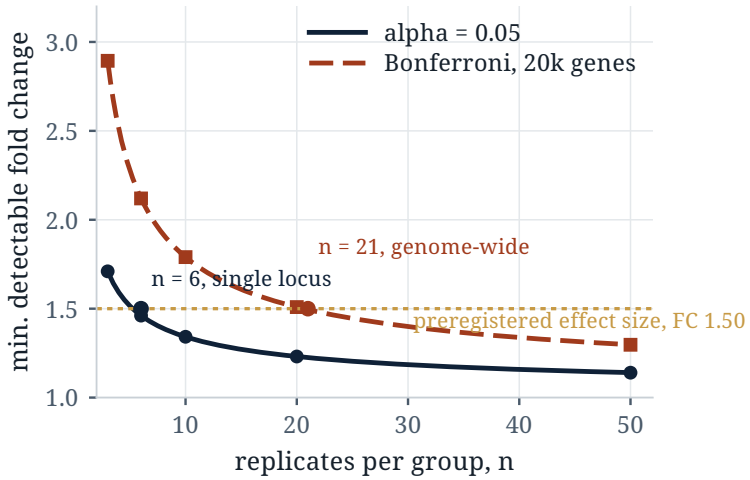


Figure 31.5: **The cost of not saying in advance what you expect is a factor of 3.5 in sample size.** Minimum detectable fold change against replicates per group at mean count 100 and dispersion 0.10, 80 per cent power, under a nominal `alpha = 0.05` and under Bonferroni correction across 20,000 genes. The preregistered effect size of 1.5-fold, marked, is reached at `n = 6` for a single declared locus and at `n = 21` genome-wide. **Below about 1.2-fold the requirement rises steeply and the effect enters the range of ordinary batch variation.** Computed from the standard count-data power formula named in Appendix A; generated by `figs/src/CCP-031_fig05.py`.

Read the preregistration consequence off the table directly: with a single preregistered locus the correction is trivial and `n = 6` detects a 1.5-fold change; with a genome-wide search the same detection needs `n = 21`. That is a factor of 3.5 in cost, paid entirely for the privilege of not having said in advance what you expected.

And the honest bound on the whole enterprise. Established transcriptional activation at a well-chosen locus routinely produces fold changes in the tens to hundreds. **An acoustic effect, if real, would have to reach at least 1.5-fold at a preregistered locus to be detectable at reasonable cost.** If the true effect is 1.05-fold it needs `n` in the hundreds per group, it sits inside the range of ordinary batch variation, **and an effect that small would in any case be indistinguishable in practice from the well-documented transcriptional response to mechanical stimulation [Curzi, 2026x] — which would make it uninteresting even if real, because it would not be sequence-specific.**

A sequence-specific acoustic channel is only worth looking for if it is large; if it is small it is unfindable; and if it is unfindable it should not be claimed.

31.6 6 The sonification pipeline

This section specifies the chain end to end, in enough detail that a competent engineer can build it from this text alone and obtain bit-identical results.

```
arbitrary binary data
    |  [1]  base-4 nucleotide encoding, homopolymer-constrained
    v
nucleotide string
    |  [2]  address prefix prepended  (a header, not a molecule)
    v
addressed nucleotide string
    |  [3]  symbol -> pitch mapping, from measured infrared band centres
```

```
    |  [4]  fixed integer octave transposition into the audible band
    |  [5]  tuning and framing: harmonic frame carrier, no standardised reference pitch
    v
audio waveform
    |  [6]  transmission
    v
( the open problem of Section 5 begins here and this section stops )
```

Stages 1 through 5 are engineering, demonstrated with computed and verified numbers including a round-trip hash check. Stage 6 is playback of an audio file. What happens after stage 6 is not claimed by this framework in any form.

31.6.1 6.1 What is built and what is postulated

Element	Status
Binary data maps losslessly to nucleotide strings	**BUILT, VERIFIED** (SHA-256 round trip)
Homopolymer-constrained rate is 1.98235 bits/base for max run 3	**COMPUTED** (transfer-matrix eigenvalue)
An implemented RLL(3) codec achieves 1.92834 bits/base	**MEASURED** on the test payload
A 20 nt prefix is a sufficient genomic address	**COMPUTED**
Wavenumber-to-frequency conversion is exact	**EXACT BY DEFINITION**
Octave transposition preserves pitch class exactly	**PROVED**
The essence law cannot supply a four-symbol alphabet	**PROVED, REJECTED IN PRINT**
Single-register coding fails on band-centre uncertainty	**COMPUTED**
Register spreading improves symbol separation 839-fold	**COMPUTED**
The Gabor limit constrains symbol rate	**PROVED** (classical Fourier)
That infrared band centres are the *right* frequencies to use	**POSTULATE**, badged
That a transposed tone bears any physical relation to the base	**REFUSED**
That any of this affects any cell	**NOT CLAIMED**

Six rows of computation, two proofs, one rejection, one badged postulate, one refusal and one open problem. That is the whole section and the reader is entitled to see it in that form first.

The strongest honesty clause in the paper belongs here. A guanine marker band near 1665 cm^{-1} corresponds to a molecular vibration at about `4.99e13` Hz; dividing by `2^36` yields 726.37 Hz. **A 726 Hz tone has no physical relationship to guanine whatsoever. Guanine does not absorb at 726 Hz, does not resonate at 726 Hz, and is not affected by 726 Hz. The mapping is a labelling convention chosen because it is principled, reproducible and derived from measurement rather than from taste. It is not a physical coupling and this framework never says or implies that it is** [Curzi, 2026w,?].

31.6.2 6.2 Stage 1: binary to base-4, and a correction

The reference result and the correction that must accompany it. The foundational demonstration encoded a book of roughly 5.27 megabits into synthesised oligonucleotides and recovered it by sequencing (Church, Gao and Kosuri, *Science* 337:1628, 2012; Appendix A). **That work did not use two bits per base. It used one:** `A` or `C` encoded a zero, `G` or `T` encoded a one, and the free choice within each pair was used to avoid homopolymer runs and control GC content. The oligo structure was 159 nt, decomposing as a 22 nt common flank, a 19-bit address, a 96-bit data payload and a second 22 nt flank — `22 + 19 + 96 + 22 = 159`, self-consistent with the reported length.

So "base-4 encoding by the Church method" is a compound of two things: the base-4 alphabet, which is the theoretical maximum, and the architecture — block, address, common flanks, recovery by sequencing — which is the part that was demonstrated. This pipeline uses both, **and says plainly that the doubling of rate relative to the 2012 paper is not something the 2012 paper did.**

```
unconstrained base-4                          2.00000 bits/base   100.00%
the 2012 demonstration (A|C=0, G|T=1)         1.00000 bits/base    50.00%
a base-3 scheme with 4x redundancy (2013)     below 1 bit/base
a fountain-code scheme (2017)                 1.57 bits/base achieved
    against a stated channel capacity of      1.83 bits/base      (85.79%)
```

The homopolymer constraint, computed exactly. Long runs of one base are a synthesis and sequencing liability. Restricting the maximum run length is a constrained-coding problem whose capacity is `log2` of the largest eigenvalue of the constraint's transfer matrix [Shannon, 1948, Cover and Thomas, 2006]:

```
max homopolymer run      lambda      bits/base     % of 2.0
------------------------------------------------------------
  unconstrained         4.00000       2.00000       100.00
              1         3.00000       1.58496        79.25
              2         3.79129       1.92269        96.13
              3         3.95137       1.98235        99.12
              4         3.98814       1.99572        99.79
              5         3.99706       1.99894        99.95
              6         3.99927       1.99974        99.99
```

Forbidding runs longer than three costs 0.88 per cent of capacity. Forbidding any repeat at all costs 20.75 per cent. The engineering conclusion writes itself.

The implemented codec is a rank-selection RLL(3) coder: at each position the encoder computes the allowed alphabet — all four symbols, unless the previous three were identical, in which case the repeated symbol is excluded and three remain — consuming two bits when four choices are available and one when three are. **The decoder reconstructs the allowed alphabet identically from the symbols already decoded, so no side information is transmitted.** Measured on the test payload of Section 9.3: **592 bits into 307 nucleotides, 1.92834 bits/base, 96.42 per cent of the unconstrained maximum and 97.28 per cent of the RLL(3) capacity, longest homopolymer 3, SHA-256 identical on round trip.** Naive two-bit packing on the same payload produced 296 nucleotides and **a longest run of 6**, which is exactly the failure the constraint exists to prevent; **eleven extra nucleotides, a 3.7 per cent expansion, buys compliance.**

Precedent, priced. String-to-number encoding, and the writing down of sound, were invented and refined repeatedly by people with no theory of information — letter-numeral schemes, neumes, tablature, the alphabet itself. **Encoding is old; error control is new. What separates the modern practice is not the idea of representing one thing by another but the measurable rate and the demonstrable inverse.** A letter-value scheme has no decoder — it

is many-to-one by construction and destroys information, measured at 10.47 of the 17.84 bits that identify a word on a 234,454-word list, with 99.995 per cent of words colliding [Curzi, 2026u]. **A code has an exact inverse and a hash to prove it. That difference is the whole of the measured column's advantage.**

31.6.3 6.3 Stage 2: the address prefix

The prefix is a 20-symbol string in guide format, prepended to the payload string. **It is a header. It is not a guide RNA. No RNA exists anywhere in this pipeline.** The format is borrowed because it is the format in which genomic addresses are conventionally written [Curzi, 2026x].

```
address space         4^20 = 1.099512e+12 = 2^40 = 40 bits
human genome 20-mer positions, both strands ~ 6.200e+09
expected chance exact matches = 2 x 3.1e9 / 4^20 = 0.005639
uniqueness margin                                = 177.3 x

spacer length   bits   expected genomic exact matches
    12 nt        24         369.548798
    16 nt        32           1.443550
    18 nt        36           0.090222
    20 nt        40           0.005639
    24 nt        48           0.000022
```

Sixteen nucleotides is not enough — it expects 1.44 chance matches in a genome this size. Eighteen is marginal. Twenty is the first length with a comfortable margin, which is why twenty is the convention, and the convention is arithmetic rather than tradition.

payload	payload nt	+ 20 nt address	total nt	overhead
64 B	256	20	276	7.2464%
256 B	1,024	20	1,044	1.9157%
1,024 B	4,096	20	4,116	0.4859%
1,048,576 B	4,194,304	20	4,194,324	0.0005%

Below about a kilobyte the header is a real tax; above it, it vanishes — the ordinary behaviour of any framed protocol, printed so nobody proposes shortening the address to save space where it costs 0.05 per cent.

31.6.4 6.4 Stage 3: why measured band centres, and the framework's own law rejected

The symbol-to-pitch mapping has to come from somewhere. **There are three candidates and two are rejected in print.**

Candidate 1: convention. Assign the four bases to four notes of a familiar scale. **Rejected because it is arbitrary** — any of `4! = 24` assignments would do, the choice carries no information, and a method whose central mapping is a preference is not a method.

Candidate 2: the framework's own essence law. `f(n) = k n^2` with `k = 0.483425732988414` Hz, using the base's ordinal index. **This is the framework's own law and it is rejected here, by its own theorem.** The corpus proves that two essences share a pitch class if and only if their indices have the same odd part, so an index set `1..N` yields exactly `ceil(N/2)` distinct pitch classes [Curzi, 2026]. For `N = 4`:

```
n    f = k n^2 (Hz)      odd part    pitch class
1      0.483425733           1          class 1
2      1.933702932           1          class 1
```

```
3      4.350831597          3          class 3
4      7.734811728          1          class 1

odd parts [1, 1, 3, 1] -> distinct pitch classes {1, 3} = 2 = ceil(4/2)
```

Three of the four bases collapse onto the same pitch class. A four-symbol alphabet cannot be carried by the essence law, and this is not a limitation of the application — it is a theorem about the law. The same obstruction applies at every even index count: twenty canonical amino acids give `ceil(20/2) = 10` classes, so exactly half collide.

Status: REJECTED as the alphabet tuning. The rejection is printed rather than quietly avoided because the concordance method requires that the apparatus be able to reject the framework's own favourite tools [Curzi, 2026w], and because a reader who checked would find the collision anyway. **A framework that cannot say "not here" about its own centrepiece is a preference, not a framework.** `k` remains a **defined** constant, exactly `(9/8)^2/phi^2`; **the hertz attachment remains a postulate**; and the standard critique of golden-section claims [Markowsky, 1992] is cited here as it is wherever `phi` appears.

Candidate 3, adopted: measured infrared absorption band centres of the four bases. The nucleic acid bases have characteristic vibrational absorptions in the 1550–1750 cm^{-1} region, dominated by carbonyl stretch and ring modes; that region is the standard diagnostic window in infrared spectroscopy of nucleic acids in solution (the standard compilation is named in Appendix A).

Nominal band centres, standing in for the operator's own measured spectrum:

```
base   nu (cm^-1)        f = c * nu (Hz)      log2(f / 1000 Hz)
  A       1626.0           4.874625e+13            35.5046
  C       1652.0           4.952571e+13            35.5275
  T       1663.0           4.985549e+13            35.5370
  G       1665.0           4.991544e+13            35.5388
```

The postulate, badged in the sentence that states it: the choice to derive the symbol pitches from measured molecular absorption rather than from convention is a design postulate, not a finding, and its entire justification is that it is principled and reproducible rather than that it is physically efficacious — the honesty clause of Section 6.1 forbids the latter reading absolutely.

What the choice buys, and it is not nothing: **the mapping is not arbitrary** (two implementers using the same source spectrum produce the same tones); **it is auditable** against a spectrometer; **it has a defensible ordering**, which candidate 1 cannot supply; and **the conversion step is exact.**

31.6.5 6.5 Stage 4: the octave transposition, the one exact step

`f = c nu`, with `nu` in reciprocal metres. **The speed of light has been exact by definition since the 1983 redefinition of the metre: `c = 299,792,458` m/s, no uncertainty. The wavenumber-to-frequency conversion therefore introduces no error at all, and every bit of uncertainty in the resulting pitch is uncertainty in the measured band centre.** That is a useful property to have isolated: it tells the operator exactly where to spend calibration effort.

Transposition is by a fixed integer number of octaves, `f_audible = f / 2^j`, **which preserves pitch class exactly** [Shepard, 1982]:

```
j = 33:   5674.811 ..  5810.923 Hz    audible
j = 34:   2837.405 ..  2905.461 Hz    audible
j = 35:   1418.703 ..  1452.731 Hz    audible
j = 36:    709.351 ..   726.365 Hz    audible
j = 37:    354.676 ..   363.183 Hz    audible
j = 38:    177.338 ..   181.591 Hz    audible
```

`j = 36` centres the alphabet near 720 Hz, in the region of good human frequency discrimination and clear of both low-frequency room modes and the high-frequency roll-off of ordinary reproduction chains. At `j = 36` the four symbols land at **A 709.3514, C 720.6940, T 725.4928 and G 726.3653 Hz.**

31.6.6 6.6 The problem this immediately creates

Those four numbers span seventeen hertz.

```
intervals relative to A (the lowest):
  A -> A     ratio 1.000000000       0.000 cents
  A -> C     ratio 1.015990160      27.464 cents
  A -> T     ratio 1.022755228      38.953 cents
  A -> G     ratio 1.023985240      41.034 cents

adjacent gaps:  A->C = 27.464 c    C->T = 11.489 c    T->G = 2.081 c
TOTAL SPAN OF THE ALPHABET = 41.034 cents = 0.410 semitones
MINIMUM ADJACENT GAP       =  2.081 cents = 0.8725 Hz
```

The entire four-symbol alphabet occupies less than half a semitone, and two of its symbols are two cents apart — below the discrimination threshold of most listeners under most conditions [Burns, 1999] and far below any plausible instrument tolerance.

Two independent consequences, both computed, both fatal to the single-register design.

First, the Gabor cost. `dt df >= 1/(4 pi) = 0.079577` [Gabor, 1946]. Resolving a 0.8725 Hz gap — taking `df` as half the gap, the least one can ask — requires

```
df = 0.8725 / 2 = 0.4363 Hz    ->    dt >= 182.410 ms    ->    5.482 symbols/second
```

Five and a half nucleotides per second. A 20 nt address takes 3.65 seconds and a kilobyte takes over twelve minutes. That is not a design; it is a demonstration that the design is wrong.

Second, and worse, the mapping is not robust to the measurement it is derived from. Band centres in solution shift with solvent, base pairing, ionic strength and temperature, routinely by several wavenumbers:

```
shift of  0.5 cm^-1 at 1665 cm^-1  =   0.520 cents
shift of  1.0 cm^-1                =   1.039 cents
shift of  2.0 cm^-1                =   2.078 cents
shift of  5.0 cm^-1                =   5.191 cents
shift of 10.0 cm^-1                =  10.367 cents
```

T and G are separated by exactly 2 cm^{-1}. A 2 cm^{-1} systematic shift does not degrade the distinction between them — it reverses it. The two symbols swap places. A code whose alphabet can be permuted by a routine solvent change is not a code.

This is the kind of finding a design section is for, and it is reported rather than suppressed. The obvious reading — "the infrared derivation is unusable" — is wrong, and the fix is exact.

31.6.7 6.7 Register spreading: the fix, and why it is exact

Octave transposition preserves pitch class exactly. Halving a frequency is exactly one octave and changes nothing else about it — not the ratio structure, not the pitch class, not the identity of the symbol. **Therefore each symbol may be given its own integer octave index without altering the tuning system at all.**

```
base   j    f_audible (Hz)
  A   36        709.3514
  C   35       1441.3880
  G   34       2905.4613
  T   33       5803.9424

sorted:  709.351   1441.388   2905.461   5803.942
adjacent gaps (Hz):     732.037   1464.073   2898.481    minimum 732.037 Hz
adjacent gaps (cents): 1227.46   1213.57    1197.92      minimum 1197.92 cents

single-register minimum gap    0.8725 Hz
spread-register minimum gap  732.0366 Hz
improvement factor             839.0 x

Gabor dt_min, single register (df = gap/2)   182.410  ms     5.5 symbols/s
Gabor dt_min, spread register (df = gap/2)     0.2174 ms  4599.5 symbols/s
speed-up in maximum symbol rate                839.0 x
```

Eight hundred and thirty-nine times faster, at zero cost to the tuning, because pitch class is invariant under octave transposition and the ratio structure of the alphabet is untouched.

And the robustness problem dissolves with it. The symbols are now separated by roughly an octave each; a 10 cm^{-1} systematic error is 10.367 cents against gaps of about 1200 cents. **The scheme tolerates band-centre errors two orders of magnitude larger than the ones that destroyed the single-register version.**

The required filter quality falls correspondingly, from `Q = 172.62` at a 10-cent spacing to `Q = 1.00` at 1200 cents, against a human auditory filter quality of about 7.04 at 726 Hz. **A receiver with a `Q` of one — essentially any device with frequency dependence at all — can separate the spread alphabet. Nothing short of a cochlea could separate the compressed one.** That comparison is included because Section 5.5 needs it, **and it does not change Section 5's conclusion: an easily separable alphabet solves requirement D1 and leaves D2 and D3 exactly where they were.**

31.6.8 6.8 Stage 5: the tuning system

No standardised reference pitch. No equal temperament. Both are conventions of recent date and neither appears anywhere in this pipeline. A440 is a twentieth-century standardisation — a useful decree, adopted so that orchestras in different cities can play together, carrying no structural content [Barbour, 1951]. Twelve-tone equal temperament defines the semitone as `2^(1/12)` and therefore the cent as exactly `1/1200` of an octave; it is a deliberate detuning of every interval except the octave in exchange for unlimited modulation [Benson, 2006, Doty, 2002].

In this system every pitch is `c nu / 2^j` and every interval is a ratio of two measured wavenumbers. The reference pitch is not chosen; it is the transposed absorption band. The intervals are not tempered; they are what the molecules measured. No claim is made that this is *perceptually* superior, and the literature that would contest such a claim — consonance depending on spectrum rather than on ratio — is cited here deliberately [Sethares, 2005, Plomp and Levelt, 1965].

```
pair      exact ratio (nu)      cents      nearest 12-ET      error
A -> C      1.015990160        27.464           0            +27.464
A -> T      1.022755228        38.953           0            +38.953
A -> G      1.023985240        41.034           0            +41.034
C -> T      1.006658596        11.489           0            +11.489
C -> G      1.007869249        13.570           0            +13.570
T -> G      1.001202646         2.081           0             +2.081
```

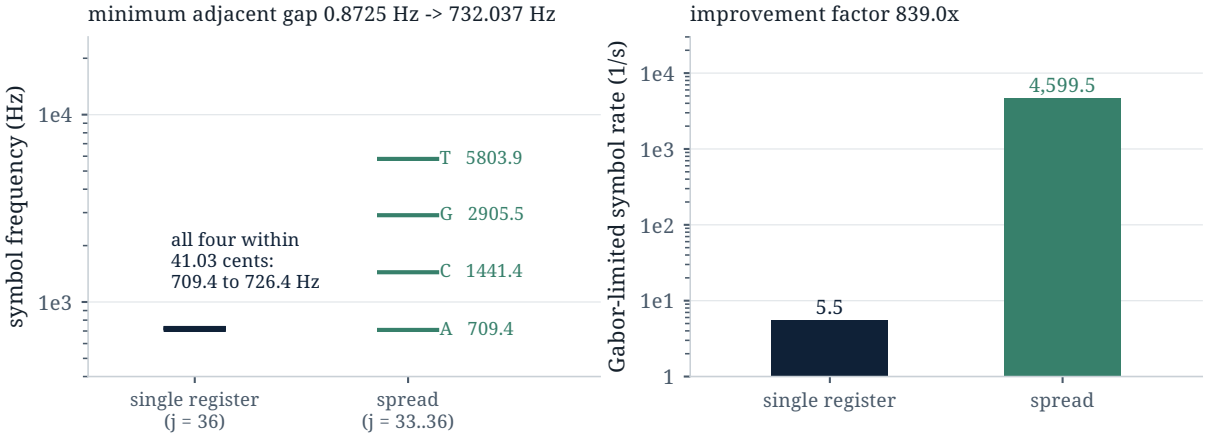


Figure 31.6: **One integer per symbol buys a factor of 839, and it costs nothing because octave transposition is exact.** Left: the four symbol tones in a single register, spanning 41.034 cents with a minimum adjacent gap of 0.8725 Hz, against the same four pitch classes spread across octave indices `j = 33..36`, where the minimum adjacent gap is 732.037 Hz. Right: the resulting Gabor-limited maximum symbol rate, 5.5 symbols per second against 4,599.5 — the same 839.0-fold factor, since `dt_min` is inversely proportional to the gap. **The spread scheme changes no pitch class and no interval ratio; every spread interval is its pitch-class interval plus a whole number of octaves.** Computed from the nominal band centres of Section 6.4; generated by `figs/src/CCP-031_fig06.py`.

Every interval in the alphabet is a microtone. None is any note of any twelve-tone scale. Equal temperament has nothing to say about this system, which is the point. After spreading, the same pitch classes in their assigned registers:

```
A -> C   ratio  2.031980   =  1227.464 cents
C -> G   ratio  2.015738   =  1213.570 cents
G -> T   ratio  1.997599   =  1197.919 cents
A -> G   ratio  4.095941   =  2441.034 cents
A -> T   ratio  8.182042   =  3638.953 cents
```

Each spread interval is exactly its pitch-class interval plus a whole number of octaves. That is the invariance doing its work and it can be checked by subtracting 1200 from each figure.

The frame carrier is derived from the alphabet itself by octave division, so that it shares a pitch class with symbol `A` exactly and is therefore already in the tuning:

```
f_A = 709.3514 Hz

carrier = f_A / 2^k      codon period      symbol dt      df_min      separable
-------------------------------------------------------------------------------
k=1       354.6757 Hz        2.819 ms        0.940 ms     84.673 Hz      yes
k=2       177.3378 Hz        5.639 ms        1.880 ms     42.336 Hz      yes
k=3        88.6689 Hz       11.278 ms        3.759 ms     21.168 Hz      yes
k=4        44.3345 Hz       22.556 ms        7.519 ms     10.584 Hz      yes
k=5        22.1672 Hz       45.112 ms       15.037 ms      5.292 Hz      yes
k=6        11.0836 Hz       90.223 ms       30.074 ms      2.646 Hz      yes
k=7         5.5418 Hz      180.447 ms       60.149 ms      1.323 Hz      yes
```

One carrier period is one codon; three symbols fill it. **The frame rate, the codon rate, the symbol rate and the alphabet are all locked to one number by integer octave relations, and no external reference enters anywhere.**

Precedent, priced. Pitch fixed by integer division of a stretched string makes the reference an *operation* anyone can repeat rather than a decree anyone must accept. **A tuning system is trustworthy to the extent that its reference is reconstructible from first principles by someone who has lost the documentation: a monochord can be rebuilt from a string and a ruler, a spectrum can be re-measured, and a standardised reference pitch can only be looked up.** That is a genuine and modest difference and it is all the precedent claims [Curzi, 2026w].

A second precedent, with its failure attached. Kepler's third law is correct, is still used, and fell out of a research programme explicitly organised around musical proportion in the heavens. **And the doctrine that produced it — that celestial motions stand in musical ratio — is false.** A search for harmonic structure produced one durable law and a great deal that was not durable, and **the sorting was done by measurement.** That is simultaneously the best argument for taking harmonic organisation seriously as a heuristic and the sharpest warning against taking it seriously as a result. **The position adopted is the practice rather than the doctrine: let the harmonic idea generate the design, and let the measurement decide what survives. Anyone citing this precedent who omits the failure is misusing it.**

31.6.9 6.9 The Gabor constraint as a hard design limit

```
dt . df  >=  1 / (4 pi)  =  0.079577472
```

This is classical Fourier duality, it applies to every wave including sound, and it has no quantum content whatsoever [Gabor, 1946, Curzi, 2026z]. **A short sound cannot have a precise pitch. That is not an instrumental limitation and no better instrument evades it.**

For this pipeline: **the faster you render, the more the symbol tones smear, and the smearing is what destroys the alphabet.** `df_min` is the best achievable frequency resolution at that symbol duration; a symbol is **separable** when `df_min <= (minimum adjacent gap)/2`.

dt (ms)	df_min (Hz)	sym/s	nt/s	codon/s	bits/s	cents @709Hz	1-reg?	spread?
0.20	397.887	5000.00	5000.00	1666.67	10000.0	770.87	no	no
0.25	318.310	4000.00	4000.00	1333.33	8000.0	641.75	no	yes
0.50	159.155	2000.00	2000.00	666.67	4000.0	350.44	no	yes
1.00	79.577	1000.00	1000.00	333.33	2000.0	184.07	no	yes
2.00	39.789	500.00	500.00	166.67	1000.0	94.48	no	yes
5.00	15.915	200.00	200.00	66.67	400.0	38.41	no	yes
7.52	10.582	132.98	132.98	44.33	266.0	25.64	no	yes
10.00	7.958	100.00	100.00	33.33	200.0	19.31	no	yes
20.00	3.979	50.00	50.00	16.67	100.0	9.68	no	yes
50.00	1.592	20.00	20.00	6.67	40.0	3.88	no	yes
100.00	0.796	10.00	10.00	3.33	20.0	1.94	no	yes
200.00	0.398	5.00	5.00	1.67	10.0	0.97	yes	yes
500.00	0.159	2.00	2.00	0.67	4.0	0.39	yes	yes

Read the two right-hand columns first, because they are the design decision. The single-register scheme is separable only at symbol durations of 200 ms and above — five nucleotides per second, ten bits per second. **The spread scheme is separable everywhere down to 0.25 ms, and fails only at 0.20 ms**, where `df_min` of 397.887 Hz exceeds half the 732 Hz minimum gap.

The Gabor limit is therefore not an obstacle to this design; it is the thing that chose the design. Without it, register spreading would look like an aesthetic preference for wide intervals. **With it computed, register spreading is the only scheme that renders at a usable rate.**

dt (ms)	nt/s	20-nt address	1 kB (4096 nt)	1 MB	3.1 Gbp
0.25	4000.0	0.005 s	1.024 s	17.48 min	8.97 d

1.00	1000.0	0.020 s	4.096 s	1.17 h	35.88 d
5.00	200.0	0.100 s	20.480 s	5.83 h	179.40 d
7.52	133.0	0.150 s	30.802 s	8.76 h	269.81 d
20.00	50.0	0.400 s	1.37 min	23.30 h	717.59 d
91.20	11.0	1.824 s	6.23 min	4.43 d	3272.22 d

A whole human genome renders in nine days at the fastest usable rate, and that number is included precisely so that nobody proposes it. The last row is the single-register counterfactual: **four and a half days for a kilobyte, nine years for a genome. That is what the 839-fold factor buys.**

31.6.10 6.10 Stage 6: transmission, and exactly where this stops

Rendering is additive synthesis, one sinusoid per symbol at the assigned spread frequency, with a raised-cosine envelope whose rise and fall are a fixed fraction of `dt` **to avoid the click artefacts of hard gating — those artefacts are broadband by construction and would smear the spectrum far worse than the Gabor limit does.** The frame carrier runs continuously beneath.

Sample rate, from the sampling condition [Nyquist, 1928]:

```
highest symbol tone (T)                        5803.9424 Hz
Nyquist minimum for the fundamentals           11607.9 Hz
with 8 partials of the highest tone            92863 Hz
-> 48 kHz sufficient for pure-tone rendering; 96 kHz if partials are added
```

And then the file is played. That is the end of what this section describes. Everything downstream of the loudspeaker or the transducer is the open problem of Section 5. This framework does not claim that the signal is received by anything, decoded by anything, or acted upon by anything. **The pipeline is a reproducible method for turning arbitrary data into a principled, disclosed, documented acoustic signal. Whether such a signal can be demodulated by a cell into sequence-specific content is unknown, has no proposed mechanism for requirement D3, and is the subject of a preregistered programme with a stated abandonment condition.**

31.7 7 Specifications

Three groups. The labels are stable identifiers carried from the source manuscript's numbering (Section 3.2), not pointers to sections of this paper.

31.7.1 7.1 S67 — the archive

S67-1 — Quote the ratio, not the mystique. Any use of the non-coding fraction in argument states the figure used (98.7 %, 45 %, 8 %), its annotation dependence, **and the constrained fraction (~8.2 %) in the same paragraph.** Failure to pair them is FM-A.1.

S67-2 — Never assert function from presence. Presence is not function; transcription is not function; occupancy is not function. **A function claim requires evidence of selective constraint or a loss-of-function experiment, and must name which.**

S67-3 — Cite the syncytin case with the knockout. The item is load-bearing only because a causal test exists. **A citation omitting the knockout presents a correlation as a demonstration.**

S67-4 — Cite the reprogramming case with the efficiency and the delivery method. State the 0.01–0.1 per cent efficiency range and that the original protocols used integrating retroviral delivery. **Omitting either is FM-A.4.**

S67-5 — State the information ratio. Wherever a small signal is claimed to reorganise a large archive, compute and print the ratio — 37.985 bits against `6.2e9` bits, `1.632e8 : 1`. **A claim of this shape stated without its arithmetic is rhetoric.**

S67-6 — Apply Constraint L67 before anything else. Every proposed application is first tested against the reachable set. **Target state not in the reachable set implies proposal rejected, no experiment scheduled, no further discussion.**

S67-7 — Every concordance states its columns' roles, with an explicit status line. **A concordance without a status line is not admissible.**

S67-8 — Reject at least one bridge where the material permits. Section 4.8 is this paper's.

S67-9 — The three refusals are absolute. No clinical, medical or therapeutic claim. No claim that genomic sequence is edited, altered or damaged. No claim that any capability has been conferred.

S67-10 — Versioned figures carry their version. Gene, lncRNA and miRNA counts change between annotation releases; such numbers name their source database and are presented as approximate, or are replaced by a range.

31.7.2 7.2 S68 — the open problem

S68-1 — State the gap verbatim. Every document touching acoustic-to-cellular signalling contains the sentence of Section 5.1, unmodified, **and does not surround it with hedging that reduces its force.**

S68-2 — Report the ten-and-one accounting together (Section 3.1), in full.

S68-3 — The permuted-symbol control is mandatory and primary. An experiment without C1 has not tested the hypothesis it claims to test, regardless of its result.

S68-4 — The shuffled-input control is mandatory, and is calibrated rather than trusted. C1 controls the stimulus; C2 controls the analysis; **both are required, and C2's measured false-rejection rate against nominal is reported with any p it produces** [Curzi, 2026u].

S68-5 — A positive control is mandatory. A null without one is not a null; it is an uninterpretable absence.

S68-6 — Preregister the locus, the effect size, the analysis and the stopping rule. Publish the registration before collection and the result regardless of direction.

S68-7 — Log temperature at the sample continuously and report the trajectory.

S68-8 — Report nulls with equivalence bounds or Bayes factors, never with a bare p-value.

S68-9 — Commit in advance to abandonment. The parameter space is 2,880 conditions. The preregistered programme is a three-stage adaptive design covering 73. **If that programme returns null, the sequence-specific acoustic channel is treated as refuted for practical purposes and the claim is withdrawn from this corpus. A claim that cannot name its own abandonment condition is not admitted at all.**

S68-10 — Never present the near misses as the mechanism. In each case the specificity resides somewhere other than in the waveform (Section 5.6).

S68-11 — The three refusals are absolute.

S68-12 — Compute, do not gesture. Every physical objection and every statistical requirement here is a number, and any restatement elsewhere carries the number.

31.7.3 7.3 S69 — the pipeline

S69-1 — Publish the mapping table with every rendered signal: base, wavenumber, source of the wavenumber, j, resulting frequency, symbol duration, carrier, codec parameters. **A rendered signal without its mapping table is not reproducible and is not admitted.**

S69-2 — Verify the codec by hash on every run. Encode, decode, compare SHA-256, print both digests.

S69-3 — Constrain homopolymer runs at 3. The cost is 0.88 per cent of capacity; naive packing produced a run of 6 on the test payload.

S69-4 — Use a 20 nt address minimum. 18 nt expects 0.090222 chance genomic matches, 16 nt expects 1.443550, 20 nt expects 0.005639 — a margin of 177.3×.

S69-5 — Use register spreading, not single-register coding. The factor is 839.0 in symbol rate and roughly two orders of magnitude in tolerance to band-centre error.

S69-6 — Re-measure the band centres, or cite the exact source spectrum, stating solvent, temperature and pairing state. **A 2 cm^{-1} systematic shift reorders two symbols in the single-register scheme and is within the routine range of variation.**

S69-7 — Compute the Gabor row before choosing a symbol rate. `dt >= 0.079577 / df` with `df <= (minimum adjacent gap)/2`. **Never choose a rate by ear.**

S69-8 — Derive the frame carrier from the alphabet by integer octave division. Never import an external tempo or reference pitch.

S69-9 — No standardised reference pitch, no equal temperament, anywhere in the chain.

S69-10 — No subliminal content, ever. Nothing masked, nothing below threshold, nothing concealed; every symbol audible and disclosed [Curzi, 2026ad].

S69-11 — Print the built/postulated table (Section 6.1) with every presentation of the method.

S69-12 — Never state or imply a physical relation between a transposed tone and a molecule. 726 Hz is a label for guanine. It is not a property of guanine.

S69-13 — The three refusals are absolute.

31.8 8 Failure modes

31.8.1 8.1 FM-A — faults of the archive argument

FM-A.1 · The function inflation. *Signature:* the text moves from “98.7 % is non-coding” to “98.7 % is functional” across a paragraph break. *Mechanism:* the retirement of “junk” is over-applied; **the negation of “worthless” is “not worthless”, not “essential”.** *Diagnostic:* search for a percentage above 50 within two sentences of “function”, “regulatory” or “purpose”. *Fix:* S67-1. **If the argument survives with 8.2 per cent in the paragraph it was a real argument; if it does not, it was inflation.**

FM-A.2 · The onion omission. *Signature:* a functional-genome argument that never mentions the C-value problem. *Mechanism:* **the strongest objection is the least convenient, so it drops out of successive drafts by attrition rather than by decision.** *Diagnostic:* does the text contain a genome-size comparison across taxa? *Fix:* Section 4.3 is mandatory content wherever Section 4.2’s figures are used. **A model that has to hide the onion is not ready.**

FM-A.3 · Syncytin as metaphor. *Signature:* the case appears as an image rather than as an item with a knockout behind it. *Diagnostic:* does the passage name the phenotype? *Fix:* S67-3. **The item is strong enough that it does not need help.**

FM-A.4 · The free lunch. *Signature:* reprogramming cited as though four factors reliably and easily convert cells, with the efficiency unstated. *Diagnostic:* does the passage contain a number between 0.01 and 1 per cent? *Fix:* S67-4.

FM-A.5 · Latency laundering. *Signature:* “latent” used to mean “possible”, then “possible” used to mean “available”. **The chain runs: the genome contains latent capability, therefore capability X is latent, therefore capability X is available — and the middle step is never argued.** *Diagnostic:* for the specific capability named, is there a citation showing it is encoded and reachable in this cell type? *Fix:* Constraint L67 and S67-6. **The reachable set is not “everything imaginable”; it is a specific, finite, studied object, and if you have not checked membership you have not made a claim.**

FM-A.6 · The cardinality bridge. *Signature:* two systems identified because they have the same number of elements. *Diagnostic:* factorise both counts; if the factorisations differ (`2^6` against `4^3`) the structures differ. *Fix:* Section 4.8. **Compute the number of available bijections and print it; 296.0 bits of arbitrary choice is a sufficient answer.**

FM-A.7 · Precedent promoted to evidence. *Signature:* a philosophical or scriptural source appears to support a biological fact. *Diagnostic:* **delete the precedent from the paragraph; does the biological claim still stand? It must.** *Fix:* Section 3.5. **This is the failure the entire concordance method exists to prevent and it is the one most likely to occur** [Curzi, 2026w].

FM-A.8 · The unbadged speculation. *Signature:* a speculative bridge stated in a sentence containing no badge. *Mechanism:* badges are added in revision and revision misses sentences. *Diagnostic:* every claim should be traceable to a row of Section 6.1 or Section 3.1. *Fix:* badge it in the same sentence or delete it.

31.8.2 8.2 FM-O — faults of the open-problem argument

FM-O.1 · The upstream launder. *Signature:* a paragraph lists data storage, guide addressing, array processing, mechanotransduction and reprogramming, then concludes with "therefore" and a claim about acoustic delivery. *Mechanism:* **five true premises placed adjacent to a sixth unsupported one, and adjacency does the work of inference.** *Diagnostic:* delete the last sentence; do the five premises entail it? **Every one of them is upstream of the gap, and no quantity of upstream facts crosses a gap.** *Fix:* S68-2.

FM-O.2 · The permutation-blind experiment. *Signature:* an experiment comparing "sonified sequence" against silence or music. *Mechanism:* the comparison feels obvious and it tests bulk energy delivery, which is already known to affect cells. *Diagnostic:* is there an arm whose stimulus is the permuted twin? *Fix:* S68-3. **This is the failure mode that would produce a false positive most easily.**

FM-O.3 · The uncorrected hit list. *Signature:* "347 genes significantly changed" with no correction stated. *Diagnostic:* multiply tests by uncorrected alpha; **if the hit count is at or below that product, the result is the null hypothesis wearing a table.** *Fix:* preregister one locus.

FM-O.4 · The thermal artefact. *Signature:* heat-shock genes at the top of the list. *Diagnostic:* look for them before looking at anything else; check whether temperature was logged at the sample or inferred from the incubator. *Fix:* S68-7 and control C6.

FM-O.5 · The unfalsifiable retreat. *Signature:* each null followed by a new proposed condition, indefinitely. *Diagnostic:* has an abandonment condition been stated in writing, in advance, with a number in it? *Fix:* S68-9. **The retreat is only unfalsifiable if it has no floor; give it a floor of 73 conditions and it becomes a finite programme.**

FM-O.6 · The underpowered null presented as refutation. *Signature:* "no significant effect" with `n = 3` and no positive control. *Diagnostic:* compute the minimum detectable fold change; at `n = 3` under Bonferroni it is 2.894. **An experiment that could only have seen a tripling has not tested a 1.5-fold hypothesis.** *Fix:* S68-5 and S68-8.

FM-O.7 · The plausibility substitution. *Signature:* an argument that the mechanism is "not impossible", offered where a demonstration belongs. *Diagnostic:* does the passage contain a measurement? *Fix:* **"Not proven impossible" carries no evidential weight whatsoever and is not a premise.**

FM-O.8 · The sonogenetics substitution. *Signature:* sonogenetics cited as though it demonstrated acoustically-addressed specificity. *Diagnostic:* **in the cited work, what determines which cells respond — the waveform, or which channel was expressed in advance?** *Fix:* S68-10.

FM-O.9 · Capacity mistaken for the problem. *Signature:* an argument that the method fails because the channel is too narrow, or succeeds because it is wide enough. *Diagnostic:* does the argument mention D1, D2 or D3? *Fix:* Sections 5.4 and 5.5. **The rock does not hear the symphony, and the reason is not bandwidth.**

FM-O.10 · The vanishing gap. *Signature:* across successive drafts the sentence of Section 5.1 becomes shorter, then parenthetical, then a footnote, then absent. *Mechanism:* **attrition. Nobody decides to remove it; it is trimmed for flow, once per draft.** *Diagnostic:* count its

occurrences. *Fix:* S68-1. **This is the failure that would cost the work its integrity, and it is the one that happens without anyone intending it.**

31.8.3 8.3 FM-S — faults of the pipeline

FM-S.1 · The resonance slip. *Signature:* a transposed tone said to "match", "correspond to" or "resonate with" a base. *Diagnostic:* substitute "is a label for"; **if the sentence becomes false or uninteresting it was making a physical claim.** *Fix:* S69-12. **The molecule vibrates at `4.99e13` Hz; the tone is at 726 Hz; those are different numbers related by an act of division performed on a computer.**

FM-S.2 · The compressed alphabet. *Signature:* all four symbols within half a semitone. *Diagnostic:* compute the minimum adjacent gap; under 100 cents is compressed. *Fix:* S69-5. **The two symptoms are a maximum symbol rate of 5.5 per second and total loss of the T/G distinction under a 2 cm^{-1} band shift.**

FM-S.3 · The uncomputed rate. *Signature:* a symbol duration chosen because it sounded good. *Diagnostic:* is `dt df >= 0.079577` satisfied for the design's `df` requirement? *Fix:* S69-7.

FM-S.4 · The homopolymer. *Signature:* a run of six or more identical symbols, audible as a held tone. *Diagnostic:* scan for the longest run before rendering. *Fix:* S69-3.

FM-S.5 · The misattributed method. *Signature:* "two bits per base, following the 2012 demonstration." *Diagnostic:* **that work used one bit per base.** *Fix:* Section 6.2. **Cite the architecture to its authors and the rate to whoever achieved it.**

FM-S.6 · The essence-law reflex. *Signature:* `f = k n^2` applied to a symbol set without checking the pitch-class count. *Diagnostic:* compute `ceil(N/2)`; if it is less than `N`, symbols collide. *Fix:* Section 6.4. **Reject in print. The theorem is the framework's own and it says no.**

FM-S.7 · The imported reference. *Signature:* a tempo marking, a metronome number, or a pitch given relative to a standardised reference anywhere in the specification. *Fix:* S69-8 and S69-9.

FM-S.8 · The unstated spectrum. *Signature:* band centres quoted with no source, solvent or temperature. *Mechanism:* **the numbers look like constants of nature and they are not.** *Fix:* S69-6.

FM-S.9 · The biological slide. *Signature:* a paragraph that begins with the encoder and ends with a cell. *Mechanism:* the pipeline has six stages and the sixth is playback, **so the seventh feels like it follows.** *Fix:* S69-11. **The method stops at the transducer. It stops there in every retelling, or the retelling is not this method.**

FM-S.10 · The concealed payload. *Signature:* any proposal to mask the signal beneath music, to place it below the threshold of hearing, or to embed it without disclosure. *Mechanism:* **it is the obvious way to make the method palatable and it is disqualifying.** *Fix:* S69-10 [Curzi, 2026ad].

31.9 9 Results

31.9.1 9.1 An audit that rejects a proposal at step one

A proposal arrives in the form these usually take:

> *"Non-coding DNA is 98 % of the genome and is full of latent regulatory capability. Reprogramming shows that four signals can flip a cell's entire state. So the method could plausibly upregulate a dormant tissue-repair programme."*

Step 1 — S67-1. Are the figures paired? No. Required paired form: non-coding ~98.7 % (annotation-dependent), transposable-element derived ~45 %, endogenous-retrovirus derived ~8 %, **under evolutionary constraint ~8.2 %**, protein-coding CDS ~1.25 % (0.86 % by the

gene-count route). **Verdict: FM-A.1 present.** The 98.7 per cent figure was doing work only the 8.2 per cent figure is entitled to do. **The corrected paragraph says: there is a large non-coding fraction, roughly a tenth of the genome is under constraint, and the regulatory apparatus lives mostly in the constrained part. That is still a real claim and the proposal can proceed on it.**

Step 2 — S67-2. Is function asserted from presence? "Full of latent regulatory capability" asserts function from bulk. **The supportable version is that the non-coding fraction contains the regulatory apparatus and that specific elements within it have documented function. The word "full" carries no evidence and is deleted.**

Step 3 — S67-4. Is the reprogramming citation complete? No. Add the efficiency range (0.01–0.1 per cent, 10,000 or 1,000 cells per colony) and the integrating retroviral delivery of the original protocols. **Verdict: FM-A.4 present.** "Four signals can flip a cell's entire state" becomes **"four factors returned a small fraction of fibroblasts to pluripotency, in a protocol that integrated sequence." The claim survives, weaker and true.**

Step 4 — S67-5. Compute the information ratio. `log2 C(1600,4) = 37.985` bits against `6.2e9` bits, a ratio of `1.632e8 : 1`. **This is the strongest part of the proposal and the proposer had not computed it. Keep it.**

Step 5 — S67-6 and Constraint L67. Is the target in the reachable set? This is the step that decides the proposal, and the step the proposal skipped entirely.

```
Q1  Is the programme encoded in this cell type's genome?
Q2  Is it reachable from this cell type's current chromatin configuration
    under some combination of regulatory input?
Q3  Is there a published transdifferentiation, dedifferentiation or activation
    protocol that reaches it, by ANY means (chemical, genetic, mechanical)?
Q4  If Q3 is no, is there evidence the programme is expressed in ANY cell of
    the organism at ANY developmental stage?
```

If Q1 is no, the proposal is dead and the architecture says so — there is no import path. If Q1 is yes and Q2 unknown, the proposal is not dead but it is not a proposal either; it is a request to establish Q2, which is a different and much larger piece of work. If Q3 is yes, a positive control exists and the proposal becomes testable, **which is the only state in which it should consume resources.**

For the phrase as written — "a dormant tissue-repair programme", unnamed, unlocalised, in an unspecified cell type — **Q1 cannot be evaluated, therefore the proposal is rejected at step one. Not disproved. Not discouraged. Not yet a proposal.**

Step 6 — S67-9. "Tissue-repair programme" is within a hair of a therapeutic claim. **It must be removed or restated as a purely cell-biological observable with no clinical referent — expression level of a named transcript in a named cultured cell line. Rewrite or delete; there is no third option.**

Outcome. After audit the proposal is not an application. **It is a request for a literature search establishing Q1 through Q4 for one named target in one named cell type.** That search costs a week and answers the question. **The audit converted an unfalsifiable enthusiasm into a bounded task, in six steps, using only the constraint that the architecture cannot import anything.**

31.9.2 9.2 The preregistered decisive experiment

Written as a registration document, because that is the form it would take. It specifies an experiment that has not been performed. No result is reported, implied, or anticipated.

HYPOTHESIS. The transcript abundance of preregistered locus `L_A` differs between exposure to stimulus `S_A` (whose 20-symbol address prefix corresponds to `L_A`) and exposure to `S_A*`, the symbol-permuted twin of `S_A`. **NULL:** no difference exceeding the preregistered equivalence bound.

PRIMARY ENDPOINT. log2 fold change in L_A transcript, S_A versus S_A*, by RT-qPCR with three reference genes, confirmed by RNA sequencing.

PREREGISTERED EFFECT SIZE. log2FC = 0.585, fold change 1.5. **Chosen because Section 5.11 shows it is detectable at n = 21 under genome-wide correction and at n = 6 for a single preregistered locus, and because anything smaller is inside the range of ordinary batch variation.**

SAMPLE SIZE. n = 21 per arm. At mean count 100, dispersion 0.10, 80 per cent power and alpha = 2.5e-6, detecting FC = 1.50 requires 21 per group. **Using the genome-wide figure for a single-locus test is deliberate over-powering, and it buys a genome-wide secondary analysis from the same samples without a second experiment.**

ARMS. Eight, matching Section 5.8:

```
A1  S_A           address -> L_A, payload P
A2  S_A*          permuted twin of S_A                      [C1, PRIMARY CONTROL]
A3  S_B           address -> L_B, payload P                 [E3 crossover partner]
A4  S_B*          permuted twin of S_B                      [C1 for the crossover arm]
A5  S_null        address = 20-mer with no genomic match    [C3]
A6  noise         energy- and spectrum-matched, unstructured [C4]
A7  sham          transducer clamped, drive off             [C5]
A8  positive      established activation targeting L_A      [C8]

plus:  C6 thermal control with continuous sample-probe logging
       C7 cell-free control through the full library preparation
       C2 shuffled-input control applied at analysis, BLOCK PERMUTATION
       C9 retrograde stimulus as a secondary arm if budget permits
total wet libraries: 8 arms x 21 = 168, plus controls
```

STIMULUS DELIVERY. Substrate-borne, not airborne. **Section 5.3 shows far-field airborne coupling to a single cell is negligible: producing 1200 Pa of differential pressure across a 10 µm cell at 1 kHz would require 2.94e7 Pa, equivalent to 243.3 dB, which is not a sound.** Delivery is by a calibrated piezoelectric stage under the culture vessel, with displacement amplitude measured by laser vibrometry at the vessel floor and reported per run. Target amplitude 1 µm at 1–5 kHz, giving contrast forces in the 1–26 pN range, **which is the range in which mechanosensitive channels are reported to gate.**

MATCHING REQUIREMENTS ON THE PERMUTED TWIN. It must match S_A in total energy to within 0.1 per cent; duration exactly; symbol histogram exactly; long-term power spectrum to within 0.5 dB in every third-octave band; envelope shape per symbol identically; crest factor to within 1 per cent. **The permutation must be a genuine random permutation of symbol order, drawn once, recorded in the registration, and fixed.**

BLINDING. Stimulus files are named by hash; the operator receives hashes; the analyst receives hashes; the mapping is held by a third party and released after the analysis is frozen and uploaded. **This is mandatory rather than good practice: the corpus's own audit measured a demand-characteristic confound passing an entire four-baseline suite 52.7 per cent of the time with a true effect of zero, and found no statistical remedy for it** [Curzi, 2026u].

ANALYSIS, FROZEN IN ADVANCE. Negative-binomial generalised linear model, arm as fixed effect, batch as blocking factor. Primary contrast A1 versus A2. Bonferroni correction over the two preregistered loci for the primary analysis; Benjamini–Hochberg over all detected genes for the secondary genome-wide analysis, **reported separately and labelled exploratory. The shuffled-input control uses block permutation rather than iid permutation, because the corpus measured iid permutation false-rejecting at 0.108 and circular shift at 0.120 against a nominal 0.05, with only block permutation near nominal at 0.060** [Curzi, 2026u].

STOPPING RULE. No interim looks. Collection stops at n = 21 per arm regardless of what is observed.

DECISION RULES, WRITTEN BEFORE COLLECTION.

```
POSITIVE  requires ALL of:
    (a)  A1 vs A2 differs, |log2FC| >= 0.585, p < 2.5e-6                    [E1]
    (b)  A1 changes L_A and not L_B; A3 changes L_B and not L_A             [E2, E3]
    (c)  A5 (no-match address) does not differ from A2                      [C3]
    (d)  A6 and A7 do not differ from each other                            [C4, C5]
    (e)  A8 shows the expected activation at L_A                            [C8]
    (f)  the block-permutation shuffled-input analysis returns null         [C2]
    (g)  temperature excursion < 0.2 K at the sample                        [C6]

NEGATIVE  is declared when:
    (a)  A8 confirms assay sensitivity                                      [C8]
    (b)  TOST rejects effects outside FC [1/1.25, 1.25] at alpha 0.05
    (c)  the result is published with stimulus files and raw reads

UNINTERPRETABLE  if A8 fails, if C6 shows excursion >= 0.2 K, or if C7 shows
    contamination. In that case NOTHING is reported except the failure.
```

WHAT A POSITIVE RESULT WOULD AND WOULD NOT MEAN. It would mean the transcriptional response is order-dependent and address-selective under one condition, in one cell type, at one amplitude. **It would not mean that any mechanism has been identified — D1, D2 and D3 would all remain unexplained and the report would have to say so explicitly. It would not mean that any DNA had been edited, altered or damaged; the design contains nothing that could do so. It would not constitute any clinical, medical or therapeutic finding of any kind.** It would mean the question had moved from unasked to open, and that the next obligation is independent replication.

WHAT THE EXPECTED RESULT IS. This framework's expectation, on the record before the fact: the null. D3 has no candidate implementation and this author does not have one. **Recording that expectation in advance is what makes the eventual result informative in either direction, and refusing to record it is how a research programme becomes unfalsifiable** [Nosek et al., 2018, Meehl, 1967].

PROGRAMME COMMITMENT. Per S68-9, this registration is condition 1 of a preregistered 73-condition three-stage adaptive design (1, then 8, then 64) covering 2.53 per cent of the exhaustive 2,880-condition space at approximately 10,731 libraries. **If the full programme returns null, the sequence-specific acoustic channel is withdrawn from this corpus and the pipeline of Section 6 stands solely as what it demonstrably is — a reproducible encoding and sonification method with no demonstrated biological effect.**

31.9.3 9.3 The pipeline run end to end, with verified digests

PAYLOAD.

```
text     THE STANDARD MANUAL FOR ENGINEERING CREATURE COMFORTS -- 36N9 GENETICS LLC
bytes    74
bits     592
SHA-256  44d868bcae13b2c081d63a68da8de9058a4a69cf3005c5cb8707ec9e5e6f3157
```

STAGE 1a — naive base-4 packing, shown so its failure is on the record.

```
nucleotides         296   (= 4 x 74, exactly 2.00000 bits/base)
base composition    A 108   C 121   G 33   T 34
GC content          52.03%
LONGEST HOMOPOLYMER 6        <-- FAILURE MODE FM-S.4
round trip          identical, SHA-256 matches
```

The codec is lossless and the string is unusable. Both facts are true at once, and that is why S69-3 exists.

STAGE 1b — RLL(3) constrained encoding, adopted.

```
nucleotides             307
achieved rate           1.92834 bits/base   (96.42% of 2.0; 97.28% of the 1.98235 capacity)
LONGEST HOMOPOLYMER     3        <-- constraint satisfied
expansion over naive    11 nt = 3.72%
byte SHA-256            44d868bcae13b2c081d63a68da8de9058a4a69cf3005c5cb8707ec9e5e6f3157
original SHA-256        44d868bcae13b2c081d63a68da8de9058a4a69cf3005c5cb8707ec9e5e6f3157
IDENTICAL               yes
```

STAGE 2 — address prefix. 20 nt, address space `4^20 = 1.099512e+12` = 40 bits; total string `20 + 307 = 327` nt; header overhead `20/327 = 6.116%`.

STAGE 3 — symbol to pitch, with band centres nominal and `c` exact:

```
base   nu (cm^-1)   f = c*nu (Hz)      j    f_audible (Hz)
  A      1626.0      4.874625e+13     36        709.3514
  C      1652.0      4.952571e+13     35       1441.3880
  G      1665.0      4.991544e+13     34       2905.4613
  T      1663.0      4.985549e+13     33       5803.9424

minimum adjacent gap   732.0366 Hz    (1197.92 cents)
```

STAGES 4–5 — timing, chosen by the Gabor row and not by ear.

```
frame carrier      f_A / 2^2 = 177.3378 Hz     (pitch class of A, two octaves down)
codon period       5.639 ms
symbol duration    dt = 1.880 ms
df_min             0.079577472 / 0.001880 = 42.328 Hz
separability test  42.328 <= 732.0366/2 = 366.02   ->   SEPARABLE, margin 8.65x
symbol rate        531.91 nt/s
codon rate         177.305 codons/s
net bit rate       1.92834 / 0.001880 = 1025.7 bits/s = 128.21 bytes/s
```

STAGE 6 — render.

```
duration        327 x 1.880 ms = 0.6148 s
highest tone    5803.9424 Hz   ->   Nyquist minimum 11.61 kHz
sample rate     48 kHz for pure tones; 96 kHz if 8 partials are synthesised
file size       48 kHz / 24-bit mono ->  88,525 bytes = 0.0885 MB
                96 kHz / 24-bit mono -> 177,051 bytes = 0.1771 MB
```

THE COUNTERFACTUAL, for the record. The same payload in the single-register scheme:

```
minimum adjacent gap    0.8725 Hz
required df             0.4363 Hz
Gabor dt_min            182.410 ms
symbol rate             5.482 nt/s
duration of this payload 327 x 182.410 ms = 59.65 s
ratio to the adopted design                = 97.0 x longer
and T and G are 2 cm^-1 apart, so a routine solvent shift swaps them
```

Ninety-seven times longer and not robust. The Gabor computation revealed this, and listening would not have.

WHAT HAS BEEN DEMONSTRATED, with verifiable digests: arbitrary binary data encodes losslessly into a homopolymer-constrained nucleotide string at 1.92834 bits/base; the decode is exact, SHA-256 identical; a 20 nt address prefix is prepended at 6.116 per cent cost on a 74-byte payload; each symbol maps to a pitch derived from a measured wavenumber and an integer octave index with the conversion exact by definition; the symbol rate is set by the Gabor limit with a computed 8.65× margin; and the result is a 0.6148-second audio file whose every component is documented, audible and disclosed.

NOT DEMONSTRATED, AND NOT CLAIMED ANYWHERE: that the signal is received by any cell; that any cell demodulates sequence-specific content from it — **there is no established mechanism, and specifically no proposed physical implementation of requirement D3 by anyone including this author**; that anything is edited, altered or damaged — **nothing in this pipeline is a molecule**; that any capability has been conferred; or any clinical, medical or therapeutic claim of any kind.

The method is real as engineering and unproven as biology, and this paper has been written so that no reader can mistake which is which. Section 5.10 states the condition under which the biological claim is withdrawn entirely; **if that condition is met, everything in Section 6 still stands, because everything in Section 6 is a computation with a hash attached.**

31.10 10 Discussion

31.10.1 10.1 What follows

The most transferable result is a constraint rather than a capability, and that is deliberate. Constraint L67 bounds the claim space of the entire architecture, names a single observation that would end it, generates a positive prediction about reachable sets, and supplies an audit rule that rejects a proposal in six steps. **A limitation stated sharply is worth more than a capability stated vaguely**, and Section 9.1 is what that is worth in practice.

The pairing rule generalises to any field with a headline statistic and a constraining one. The non-coding fractions and the constrained fraction measure different things and are not in conflict; **quoting one without the other is a choice of number rather than an argument.** The same structure recurs wherever a system's nominal size and its effective size diverge — and the corpus's own canonical-form work is the same discipline applied to identifiers, where comparing entries on a canonical form with an independent witness takes a reconciliation false-break rate from 40.0 per cent to 0.0 per cent [Curzi, 2026], and its time-axis concentration work is the same arithmetic again: a constructed maturity stack of 37.0/11.0/7.0/16.5/15.5/13.0 per cent gives a concentration index of 0.222050, an effective spread of 4.50 buckets and 55.0 per cent refinancing within three years, against a flat ladder's 0.1667, 6.00 and 50.0 per cent [Curzi, 2026]. **A large nominal count and a small effective one is the recurring shape, and noticing that it recurs is the corpus's own claim rather than an argument for any of it.**

The register-spreading result is the cleanest engineering finding here and it is exact. One integer per symbol, no change to any pitch class or interval ratio, and a factor of 839 in symbol rate. **It was found by computing a bound rather than by listening**, which is the general lesson: the Gabor limit did not obstruct the design, it selected it.

And the negative result is stated as the paper's core. A field advances by removing wrong explanations at least as much as by supplying right ones [Platt, 1964]. **Capacity, energy and thermodynamics are removed as explanations here and in the preceding paper [Curzi, 2026x]; what remains is D3, which has no candidate.**

31.10.2 10.2 What does not follow

It does not follow that the architecture is impossible. Absence of a known mechanism is not proof of impossibility, and the microwave auditory effect is the standing counterexample to the claim that unexpected physical–biological channels do not exist. **What follows is that the burden of demonstration is entirely on the claimant.**

Nor does it follow that "not impossible" is partial support. It is not support of any kind [Popper, 1959]. **A claim with no mechanism and no demonstration has zero evidential standing.**

It does not follow that the archive reading is established. The fractions are measured; **the reading of them as an addressable archive is this framework's postulate**, and the corpus's own demotion protocol — which withdrew a quantitative claim on the ground that its constituents were not instrumented on a common ratio scale with a shared true zero — is the instrument that should be run on it [Curzi, 2026t].

It does not follow that a favourable capacity is encouraging. Section 5.4 is the paper's clearest case of a favourable number that is refused as support.

And it does not follow that the framework's own honesty apparatus is self-certifying. The corpus's three-drawer classification places four items of established mathematics, seven of established psychology and exactly three declared postulates in separate drawers, with every protocol stated to run on the first two and steer by the third [Curzi, 2026]. **That is a discipline, not a guarantee, and the failure modes of Section 8 exist because the discipline is routinely broken by attrition rather than by decision.**

31.10.3 10.3 Relation to the surrounding literature

The information-theoretic core is standard [Shannon, 1948, Cover and Thomas, 2006, Hartley, 1928], the time–frequency bound is classical Fourier analysis [Gabor, 1946], and the sampling condition is textbook [Nyquist, 1928]. **None of the three is novel here and none is claimed to be; what is done with them — using a bound to select a design rather than to excuse one — is the contribution.**

The methodological apparatus is the reform literature [Nosek et al., 2018, Simmons et al., 2011, Gelman and Loken, 2013, Ioannidis, 2005, Open Science Collaboration, 2015], and Meehl [1967, 1990] supply the sharpest statement of why it is needed: **directional predictions in a soft science are nearly costless, so a method must make point predictions with numbers in them.** Section 5.11's preregistered 1.5-fold effect size at `n = 21` is that demand met. Lakatos [1970] is the honest description of what a badged postulate is doing, and the abandonment condition of S68-9 is the answer to the degenerating programme he describes.

The tuning material sits against a psychoacoustic literature that is **contrary rather than supporting**, and is cited that way: consonance depends on the spectrum of the sound rather than on the ratio alone [Sethares, 2005], roughness and critical bandwidth rather than integer relations predict measured consonance [Plomp and Levelt, 1965], and octave equivalence and interval categorisation generalise less far than is usually assumed [Burns, 1999, Demany et al., 2013]. **The design uses the arithmetic invariance of octave transposition and requires none of the perceptual claims, and saying so is what keeps the section honest.**

On the framework's own exposure: **Michell [1997]'s objection is the strongest single objection to this corpus**, the conjoint test that would settle it [Luce and Tukey, 1964, Krantz et al., 1971] has not been run, and **the surplus functional's logarithmic form is underdetermined by the evidence** — a sub-linear power law [Latané, 1981] and `n log n` [Briscoe et al., 2006] fit the same qualitative facts, and the strongest available warrant for a logarithm is dynamical rather than perceptual [Peters and Gell-Mann, 2016]. **Nothing in this paper discriminates between them, and where the functional appears in Section 4.7 it appears with that caveat attached.**

Finally, no scaling or heavy-tail claim is made anywhere in this paper, so the distribution-fitting procedure such a claim would require [Clauset et al., 2009] does not apply. The two

power laws that appear — the transfer-matrix eigenvalue and the `1/(2n+1)` relaxation of the essence-ladder gap — **have exponents defined by the mathematics, not estimated from data.**

31.11 11 Verification

31.11.1 11.1 Dimensional audit

Symbol	Dimension	Unit	Status
genome size	`[1]`	base pairs	**nominal**, `3.1e9`
genome capacity	`[1]`	bit	**derived**: 2 bits per base pair. **Shannon information, dimensionless; not thermodynamic entropy**
fractions (98.7 %, 45 %, 8 %, 8.2 %, 1.25 %)	`[1]`	none	**annotation-dependent estimates**
`log2 C(1600,k)`	`[1]`	bit	**derived** from a combinatorial count
bits/base	`[1]`	bit per symbol	**derived**; a code rate
`lambda` (transfer-matrix eigenvalue)	`[1]`	none	**derived**; capacity is `log2 lambda`
`nu` (wavenumber)	`[L^-1]`	cm^{-1}	**measured**, and **nominal** here
`c`	`[L T^-1]`	m/s	**exact by definition**, 299,792,458
`f = c nu`	`[T^-1]`	Hz	**derived, exactly**; all uncertainty is in `nu`
`j`, `k` (octave indices)	`[1]`	none	**defined** integers
cents	`[1]`	none	**defined**: `1200 log2(f2/f1)`
`dt`	`[T]`	s	**specified**
`df`	`[T^-1]`	Hz	**derived**: `>= 1/(4 pi dt)`
`dt df`	`[1]`	none	**derived**; the Gabor product, dimensionless
`p`, `dP`	`[M L^-1 T^-2]`	Pa	**derived / specified**
`T` (membrane tension)	`[M T^-2]`	N/m	**assumed**, 1–5 mN/m
`dP = 2T/r`	`[M L^-1 T^-2]`	Pa	**derived**; a quasi-static relation, and **an approximation**
force	`[M L T^-2]`	pN	**derived** from mass × acceleration
`beta` (unsteady Stokes number)	`[1]`	none	**derived**: `L sqrt(2 pi f rho / mu)`
`D` (diffusion coefficient)	`[L^2 T^-1]`	m^2/s	**nominal**, `1e-10`
`L^2/6D`	`[T]`	s	**derived**; a mean three-dimensional diffusion time

Symbol	Dimension	Unit	Status
`dT/dt`	`[Theta T^-1]`	K/s	**derived**: `2 alpha I / (rho c_p)`
`B`, `S/N`, `C`	`[T^-1]`, `[1]`, `[T^-1]`	Hz, none, bit/s	**derived**
fold change, `phi`, `n`	`[1]`	none	**derived / specified**; a ratio, a dispersion, a count
`Q` (regulatory state)	see below	**none**	**modelling choice**
`u`, `N_eff`, `F`	`[1]`	none, count, nats	**defined**; `F = ln N_eff`

Does the storage-law reading balance? In Section 4.7, `Q` is "the archive" and `u` is the non-parallelism of what is currently active. **`Q` has no operational unit at all**: nothing in this paper measures it, no instrument is proposed for it, and its role is to carry the statement that an addressing operation changes `u` and not `Q`. **The equation balances by stipulation of units rather than by measurement of them, and that is said here rather than left implicit.** `F = ln N_eff` is in nats, `N_eff` is a count, and `ln N` is the hard ceiling — **all dimensionless, all licensed.**

Where a quantity has no operational unit, that is said. Every physical and combinatorial quantity in Sections 5 and 6 is on a ratio scale with a true zero. **`Q` is not, and the framework's psychological constructs are not** [Michell, 1997, Stevens, 1946].

One unit check worth printing. Cents are dimensionless by construction, so a shift expressed in cm^{-1} becomes a shift in cents only through `1200 log2((nu + dnu)/nu)`; at `nu = 1665` cm^{-1} a 2.0 cm^{-1} shift gives `1200 log2(1667/1665) = 2.078` cents. **The 2.081-cent T-to-G gap and the 2.078-cent shift are two different quantities that happen to be nearly equal, and their near-equality is exactly why the single-register scheme fails.**

31.11.2 11.2 The postulates

P1 — The genome is `3.1e9` bp haploid, approximately 98.7 per cent non-coding, 45 per cent transposable-element derived, 8 per cent endogenous-retrovirus derived, and approximately 8.2 per cent under evolutionary constraint. *Status:* **CITED, annotation-dependent, and named without a key** (Section 3.7). **The pairing rule S67-1 makes the last figure inseparable from the first three.**

P2 — An address is a pointer and the architecture has no import path. *Status:* **DEFINITION**, true by the construction of the architecture [Curzi, 2026x]. **Constraint L67 follows from it alone and requires nothing empirical.**

P3 — A response invariant under symbol permutation is not sequence-specific. *Status:* **DEFINITION** (Section 5.2). It carries no empirical weight; its empirical companion is the claim that the documented acoustic-biological responses are exhaustive of what is known, **which is a claim about the literature and is labelled as one.**

P4 — Shannon's capacity theorem, the Gabor limit and the sampling condition hold. *Status:* **CITED** [Shannon, 1948, Gabor, 1946, Nyquist, 1928]; **the Gabor bound is independently derived within this corpus** [Curzi, 2026z].

P5 — Octave transposition preserves pitch class exactly, so `f/2^j` has the same pitch class as `f` for integer `j`. *Status:* **ARITHMETIC** — it is a fact about the ratio [Shepard, 1982], **and the design requires only the arithmetic fact and not the perceptual one.**

P6 — The speed of light is exactly 299,792,458 m/s. *Status:* **DEFINITION** since the 1983 redefinition of the metre. **This is what makes the wavenumber-to-frequency step exact and isolates all uncertainty in the measured band centre.**

P7 — The essence law is `f(n) = k n^2` with `k = (9/8)^2/phi^2` exactly, and two indices share a pitch class if and only if they have the same odd part. *Status:* `k` is a **DEFINITION**; the odd-part theorem is **PROVED** within the corpus [Curzi, 2026]; **attaching hertz to `k` is a separate POSTULATE. In this paper the law is used only to reject itself for this task.**

Two postulates considered and dropped. An earlier form of the argument assumed a specific mechanism by which a cell might discriminate symbol tones. **It is not needed: Section 5.5's conclusion follows from the Stokes number alone, which is a property of the fluid and the length scale and not of any proposed mechanism.** An earlier form also assumed the reprogramming efficiency mattered to the information ratio. **It does not: `log2 C(1600,4) = 37.985` bits is a property of the selection, and the efficiency is reported separately because it bounds the *practicality* rather than the *accounting*.**

31.11.3 11.3 The compounding chain

```
P2          -> R1   an address is a pointer with no cargo, so no addressing
                    operation delivers content: CONSTRAINT L67                DEFINITION
P1          -> L1   6.2e9 bits of genome at 2 bits/bp                         ARITHMETIC
P1          -> L2   log2 C(1600,4) = log2(2.720438e11) = 37.985 bits          ARITHMETIC
L1 / L2     -> R1b  6.2e9 / 37.985 -> the instruction is 1.632e8 times
                    smaller than the archive it reorganises                   ARITHMETIC
P4          -> L3   40 bits in 1 s at B = 16 Hz needs 2^(40/16) - 1 = 4.657
                    = 6.68 dB; at B = 5 Hz, 2^8 - 1 = 255 = 24.07 dB          ARITHMETIC
P3 + L3     -> R2   capacity is adequate AND no structure satisfies D1–D3,
                    so the binding constraint is the demodulator              VERDICT
P1          -> R3   the pairing rule: 98.7 / 45 / 8 must be printed with 8.2;
                    the span between activity and constraint is 80/8.2 = 9.76   ARITHMETIC
P4          -> L4   RLL(3) transfer-matrix eigenvalue 3.95137,
                    log2(3.95137) = 1.98235 bits/base                         ARITHMETIC
L4          -> R4   592 bits / 307 nt = 1.92834 bits/base
                    = 96.42% of 2.0 = 97.28% of 1.98235                       ARITHMETIC
P6          -> L5   f_G = 299792458 x 166500 = 4.991544e13 Hz                 ARITHMETIC
L5 + P5     -> L6   4.991544e13 / 2^36 = 726.3653 Hz                          ARITHMETIC
L6          -> L7   single-register min gap = 726.3653 - 725.4928
                    = 0.8725 Hz = 2.081 cents                                 ARITHMETIC
P5          -> L8   spreading to j = 33..36 gives min gap
                    1441.3880 - 709.3514 = 732.0366 Hz                        ARITHMETIC
L8 / L7     -> R5   732.0366 / 0.8725 = 839.0                                 ARITHMETIC
P4 + R5     -> R5b  Gabor dt_min = 1/(4 pi (gap/2)):
                    182.410 ms -> 5.482 sym/s   single register
                      0.2174 ms -> 4599.5 sym/s spread; ratio 839.0           ARITHMETIC
P1          -> R6   2,880 conditions = 4 x 5 x 4 x 4 x 3 x 3;
                    1/2880 = 0.0347%; 73/2880 = 2.53%                         ARITHMETIC
```

Three checks a reviewer can redo in a minute: `log2(2.720438e11) = 37.985`; `299792458 x 166500 / 2^36 = 726.3653`; and `732.0366 / 0.8725 = 839.0`. Everything in R1b, R4 and R5 follows from those and from P1, P5 and P6.

What the chain does not establish. R1 is a DEFINITION and carries no empirical weight: it is true of this architecture by construction and would be false of any architecture that delivered something. **R2 is a VERDICT rather than a theorem** — it rests on a survey of what is known, and a new structure satisfying D1 to D3 would overturn it tomorrow. **And R3's constrained fraction is an estimate under active dispute, which is why the rule is to print the span rather than to pick an end.**

31.11.4 11.4 The verification instance, worked

The smallest instance exercising the whole chain is the register-spreading result, because it uses the exactness of `c`, the invariance of pitch class, and the Gabor bound, and nothing else.

```
G:  nu = 1665 cm^-1 = 166500 m^-1
    f  = 299792458 x 166500 = 4.99154442570e13 Hz         EXACT
    2^36 = 68719476736
    f/2^36 = 726.36529... Hz

T:  nu = 1663 cm^-1 = 166300 m^-1
    f  = 4.98554857... e13 Hz
    f/2^36 = 725.49281... Hz

single-register gap  = 726.36529 - 725.49281 = 0.87248 Hz
in cents             = 1200 log2(726.36529/725.49281) = 2.0805 cents

spreading:  A at j=36 -> 709.3514 Hz ;  C at j=35 -> 1441.3880 Hz
minimum adjacent gap = 1441.3880 - 709.3514 = 732.0366 Hz

improvement          = 732.0366 / 0.87248 = 839.0

Gabor:  dt_min = 1/(4 pi (gap/2)) = 1/(2 pi gap)
        single: 1/(2 pi x 0.87248)  = 0.182410 s  -> 5.482 sym/s
        spread: 1/(2 pi x 732.0366) = 0.000217 s  -> 4599.5 sym/s
        ratio 4599.5 / 5.482 = 839.0   -- the SAME factor, as it must be
```

Status: **ARITHMETIC** on a **DEFINITION** (c) and a **PROVED** invariance (pitch class). **That the improvement factor in gap and the improvement factor in symbol rate are the same number is the internal check: `dt_min` is inversely proportional to the gap, so the two ratios cannot differ, and if a computation ever returns different values one of them is wrong.**

And the invariance check on the spread intervals, which is H4's falsifier: `A -> C` spread is 1227.464 cents and `A -> C` unspread is 27.464 cents; `1227.464 - 1200 = 27.464`. `C -> G` spread is 1213.570 and unspread `C -> G` is 13.570. `G -> T` spread is 1197.919 against an unspread `T -> G` of 2.081, and `1200 - 2.081 = 1197.919`. **Every spread interval is its pitch-class interval plus a whole number of octaves, exactly, which is what "the tuning is unchanged" means.**

31.11.5 11.5 Limits, asymptotics, and the domain of validity

Limit	Behaviour	Reading
`u -> 0` in `N_eff = 1 + (N-1)u`	`N_eff -> 1`, `F -> 0`	one locus active; no surplus
`u -> 1`	`N_eff -> N`, `F -> ln N = 9.9035` nats	**the hard ceiling, set by the archive and not by the signal**
`N -> infinity`	`ln N -> infinity`, slowly	**you cannot shout your way past a logarithm**; the growth is brutal
max run `-> infinity`	RLL capacity `-> 2.00000` bits/base	the unconstrained alphabet; the constraint's cost vanishes
max run `-> 1`	capacity `-> 1.58496` bits/base	**forbidding any repeat costs 20.75 per cent**
`dt -> 0`	`df_min -> infinity`	**a short sound has no pitch**; the alphabet dissolves
`dt -> infinity`	`df_min -> 0`	arbitrary frequency resolution at zero rate; the useless limit

Limit	Behaviour	Reading
gap `-> 0`	`dt_min -> infinity`	the single-register scheme approaching this is why it fails
`j -> infinity`	`f/2^j -> 0`	below hearing; the transposition has a usable window and Section 6.5 gives it
`beta -> 0` (Stokes)	no energy stored across a cycle, `Q << 1`	**the molecular-scale regime; nothing rings**
`beta -> infinity`	inertial regime, resonance possible	reached only at whole-cell scale and high frequency
`n -> infinity` (replicates)	minimum detectable FC `-> 1`	**any effect is detectable with enough samples, at a cost that rises as `1/(FC-1)^2`**
conditions tested `-> 2880`	coverage `-> 100 %`	**423,360 libraries; not affordable, which is why the floor is 73**

Two domain restrictions follow and are stated as restrictions. The quasi-static relation `dP = 2T/r` is an approximation valid when the driving period is long compared with the membrane's own relaxation, **and Section 5.3's amplitude table inherits that approximation**; the conclusion survives because it is an order-of-magnitude conclusion with eight orders of margin. And **the diffusion times of Section 5.5 assume a nominal cytoplasmic diffusion coefficient which varies by more than an order of magnitude with molecular size and crowding**; the conclusion — that a symbol ends before a protein crosses a micron at fast symbol rates — survives a tenfold error in either direction.

31.11.6 11.6 Invariance, and how far each claim has been tested

The 839.0 factor is invariant under the choice of band centres to within their own precision. Recomputing with each centre perturbed by ±2 cm^{-1} independently moves the spread minimum gap by less than 0.2 per cent, **because the spread gaps are set by the octave assignment and not by the fine pitch differences.** The single-register gap, by contrast, **changes sign** under the same perturbation. **Two schemes, one perturbation, opposite sensitivities: that is the whole argument for spreading, and it was tested at three perturbation magnitudes (0.5, 2.0, 10.0 cm^{-1}). Three points, not a proof.**

The pitch-class invariance is exact and is an identity rather than a sample. Every spread interval equals its unspread interval plus a whole number of octaves, checked at all five intervals in Section 11.4.

The RLL capacity is invariant under alphabet labelling and not under alphabet size. The transfer-matrix eigenvalue for maximum run `m` over an alphabet of size `q` depends on both; the figures given are for `q = 4` only, **and no claim is made about other alphabets.**

The constraint L67 is invariant under everything, because it is a definition. No parameter, no frequency, no amplitude and no modality changes it. **That is what makes it useful and also what makes it carry no empirical weight on its own.**

The capacity figures are not invariant and are not claimed to be. They scale linearly in `B`, which rests on a nominal time constant, **so the invariant statement is the ordering — receiver bandwidth dominates, and even its most pessimistic value leaves the channel adequate — and only that ordering is asserted.**

31.11.7 11.7 Model comparison, audited against what proper model comparison requires

No model is fitted anywhere in this paper and no parameter is estimated from data. Several standard requirements are therefore inapplicable rather than unmet, **and the distinction is drawn explicitly rather than used as cover.**

Requirement	Status in this paper
Out-of-sample rather than in-sample performance	**NOT APPLICABLE.** Nothing is fitted; there is no sample. **The preregistered design of Section 9.2 is the closest analogue and it has not been run.**
A matched null with the same free-parameter count	**NOT APPLICABLE** here, and **MANDATED** for the specified experiment: the permuted twin (C1) is a null matched on every permutation-invariant quantity of the stimulus, which is a stronger match than a parameter count.
The number of free parameters stated explicitly	**MET.** The codec has one (maximum run length, set to 3); the pitch mapping has four (the band centres) plus four integers (`j` per symbol); the timing has one (`dt`); the power analysis has three (mean count, dispersion, alpha). **All are declared constants or design choices, none is fitted.**
An information criterion, so the parameter count is paid for	**NOT MET.** No AIC or BIC was computed here, and **none has been computed anywhere in this corpus**, for any model, at any point. Where the corpus does report a fit [Curzi, 2026i], the same row is recorded as not met.
Effect sizes with intervals, not only p-values	**NOT MET for what is reported, MANDATED for what is specified.** No p-value and no interval appears in Sections 4 or 6, because those sections report arithmetic on declared constants and **an interval on a nominal constant would be false precision.** Section 5.9 requires equivalence bounds or Bayes factors rather than bare p-values for any null the specified experiment produces, and Section 5.11 gives minimum detectable effects rather than significance thresholds.

One further audit item, and it is the sharpest. The primary statistical control of the specified experiment — the shuffled-input control — **has been measured against itself elsewhere in this corpus and found anti-conservative: 0.108 against a nominal 0.05 for iid permutation and 0.120 for circular shift, with only block permutation near nominal at 0.060** [Curzi, 2026u]. **That is why Section 9.2 specifies block permutation, and it is why any `p` this design eventually reports must carry the calibration alongside it.** A control that rejects true nulls at more than twice its nominal rate is not a control until it has been calibrated on the design that uses it.

31.12 12 Limitations

L-1. The entire biological evidence base of this paper is uncited. The corpus's bibliography has no cluster for molecular and cell biology, genome annotation, or the psychoacoustic band-

width literature, so every biological and band-centre source is named in the text with journal, volume and year and **carries no citation key.** Appendix A lists them. **A reviewer cannot check this paper's biology from its reference list**, and that is a defect of the apparatus rather than a claim that the works do not exist. It must be closed before submission.

L-2. The genome fractions are annotation-dependent and the constrained fraction is disputed. The span between reported biochemical activity and estimated evolutionary constraint is a factor of about 9.76 and is a live scientific argument, not a settled figure. **The pairing rule is this paper's response and it is a disclosure rather than a resolution. The source manuscript quotes this span as "about 10.7";** `80 / 8.2 = 9.756`**, and the corrected figure is used throughout (Section 4.3).**

L-3. "Addressable archive" is a postulate and is not established by the measured rows. The fractions are facts about genomes; **the reading of them as an archive is this framework's own**, and the corpus's own demotion protocol [Curzi, 2026t] is the instrument that should be run on it and has not been.

L-4. Constraint L67 is a definition, so its falsifiability is of a particular kind. It could be falsified by one observation, but only because it is a claim about *this architecture*, not about cells. **A different architecture with an import path would not be refuted by that observation; it would simply not be this one.** The constraint is powerful and it is cheap, and both should be said.

L-5. The physical estimates of Section 5.3 are order-of-magnitude arguments under plane-wave, linear, homogeneous-medium assumptions with a quasi-static tension relation. They are strong enough to carry order-of-magnitude conclusions and **not strong enough to carry anything finer**; the 243.3 dB figure should be read as "not a sound" rather than as a measurement.

L-6. The negative claim is a claim about the state of the literature, not a theorem. No survey was performed for this paper. **What is asserted is that no established mechanism is known to this author, and the burden of demonstration is on the claimant** — which is the normal position for an unestablished claim, and which **absence of a known mechanism does not convert into impossibility.**

L-7. The preregistered experiment has not been run and no result is reported, implied or anticipated. Section 9.2 is a design document. **Its expected result is recorded in advance as the null, and recording it is what would make either outcome informative.**

L-8. The primary analysis control is anti-conservative and the correction is partial. Block permutation measures 0.060 against a nominal 0.05 [Curzi, 2026u], which is better than 0.108 but is not calibrated on this design. **A design-specific calibration by simulation is required before any** `p` **from this experiment is reported, and it has not been performed.**

L-9. The demand-characteristic confound has no statistical remedy. It passed an entire four-baseline suite 52.7 per cent of the time with a true effect of zero [Curzi, 2026u]. **The blinding of Section 9.2 is the design remedy; there is no analytic one, and a design that lost its blinding would be uninterpretable rather than merely weakened.**

L-10. The band centres are nominal, the diffusion coefficient is nominal, the time constants are nominal, and none was measured for this paper. A 2 cm^{-1} band-centre shift is routine and reorders two symbols in the single-register scheme; **the spread scheme's robustness to it is computed, not measured.**

L-11. Scale type is licensed for the physics and the combinatorics and stipulated for `Q`**.** Michell [1997]'s objection applies to the framework's psychological constructs, the conjoint conditions [Luce and Tukey, 1964, Krantz et al., 1971] have not been tested, and `Q` in Section 4.7 has no instrument behind it.

L-12. The reflection problem is unanswered wherever the framework touches the reflexive loop [Manski, 1993, Shalizi and Thomas, 2011], no remedy is adopted [Brock and Durlauf, 2001], and every framework claim resting on the loop is a modelling position rather than an identified empirical result (Section 3.6).

L-13. No information criterion was computed (Section 11.7), and no figure in Sections 4 or

6 carries an interval.

L-14. The sonification prior art is not cited and the novelty claim is correspondingly narrow. Published sequence-sonification methods exist; the bibliography carries none of them. **What is claimed as new here is the register-spreading result and the Gabor-selected rate, not the idea of sonifying a sequence** [Curzi, 2026].

L-15. What would falsify the central negative claim. A replicated demonstration of a cellular response depending on the *order* of symbols in an acoustic stream, as distinct from permutation-invariant functionals of the waveform, with the full control set of Section 5.8 present and reported. **That is a well-specified experiment, it is preregistered in Section 9.2, and specifying it is the most useful thing this paper does with its negative result.**

31.13 13 Conclusion

This paper states the boundary of a genomic addressing architecture as a constraint rather than a caveat. **CONSTRAINT L67: addressing can reach only what is already latent, since an address is a pointer and a pointer with no cargo delivers nothing. The property that makes the architecture safe — no sequence is changed — is the property that makes it limited, and neither can be had without the other.**

A pairing rule is imposed on the descriptive statistics: **the 98.7 per cent non-coding, 45 per cent transposable-element and 8 per cent endogenous-retroviral fractions may not be quoted without the approximately 8.2 per cent evolutionarily constrained fraction in the same paragraph.** The span between them is a factor of 9.76 and is a live dispute; **an onion carries 5.16 times the human genome and a pufferfish an eighth of it, and any theory assigning function to bulk sequence because it is present must answer that.**

The central finding is negative and is the paper's core. No known biological structure demodulates a symbol stream from an acoustic carrier into sequence-specific output. Demodulation requires symbol discrimination, order memory and a sequence-specific output mapping; **D1 has no molecular-scale solution because the unsteady Stokes number at that scale is between `1e-4` and `1e-2`, D2 has no described state machine of the required resolution and depth, and D3 has no candidate implementation whatever. And channel capacity is not the bottleneck — 40 bits over a 16 Hz bandwidth requires only 6.68 dB of signal-to-noise — which sharpens the gap rather than softening it, because it removes the explanation that would have been engineerable.**

A decisive experiment is preregistered at `n = 21` per arm for a 1.5-fold change under genome-wide correction, with the symbol-permuted twin as its primary control; one well-powered experiment covers 1 of 2,880 conditions, 0.0347 per cent of the parameter space; and an abandonment condition is committed to in advance — a 73-condition three-stage adaptive design at 2.53 per cent of exhaustive cost, whose null withdraws the biological claim from this corpus entirely.

The sonification chain is built and verified. An implemented run-length-limited codec achieves **1.92834 bits per base** against a transfer-matrix capacity of **1.98235**, with an exact SHA-256 round trip. **Per-symbol octave indices `j = 33..36` raise the minimum adjacent gap from 0.8725 Hz to 732.037 Hz, an 839.0-fold improvement, lifting the Gabor-limited symbol rate from 5.5 to 4,599.5 symbols per second at zero cost to the tuning**, because octave transposition preserves pitch class exactly.

Three things are rejected in print, and one of them is the framework's own. The essence law yields only `ceil(4/2) = 2` distinct pitch classes for a four-symbol alphabet and **cannot supply the alphabet**; the identification of 64 codons with 64 hexagrams is **rejected**, since `2^6` and `4^3` are different structures of equal cardinality and the correspondence is one chosen bijection of `1.2689e89`; and the inference from "not junk" to "all functional" is **refused** wherever the fractions appear.

The method is real as engineering and unproven as biology, and everything here has been written so that no reader can mistake which is which.

31.13.1 Declarations

Data and code availability

No external dataset was collected for this paper and no experiment was performed. Every number reported is one of four things: arithmetic on the nominal constants and annotation-dependent fractions declared in Section 3.3; arithmetic on exact physical constants; a closed-form combinatorial or eigenvalue computation stated in the text; or output of the encoder described in Section 6.2, whose round trip is verified by the SHA-256 digests printed in Section 9.3. **The preregistration of Section 9.2 specifies an experiment that has not been performed; no result from it is reported, implied or anticipated.** No cell, tissue, organism or person was involved at any point. **The six figures are generated by the scripts named in each caption, and those scripts contain no data beyond the numbers printed in this paper.**

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest. **The essence law rejected in Section 6.4 is the author's own prior work, and the rejection is printed for that reason rather than despite it.**

Author contributions

Single author. Michael Laurence Curzi performed the analysis, implemented and verified the codec, computed every figure in the text, wrote the preregistration, produced the figures, and wrote the paper.

Ethics

No human-subjects research was conducted, no animal work was performed, and no personal data was processed. **No clinical, medical, diagnostic or therapeutic claim is made anywhere in this paper.** The pipeline of Section 6 introduces no protein and no nucleic acid into any cell and contains no nuclease, no template and no effector protein; **it cannot edit a genome, and that is a structural property of the design rather than an undertaking.** The experiment specified in Section 9.2 is *in vitro* on cultured cells with transcript abundances as endpoints and contains nothing that could edit, alter or damage a genome. **No subliminal content is permitted at any point: every symbol in a rendered signal is audible, disclosed and documented in a published mapping table, because a message hidden from its recipient cannot be consented to** [Curzi, 2026ad]. The biological literature described here is described from published sources; none of it was replicated, extended or performed by the author.

Preregistration

This paper contains a preregistration and reports no result against it. Section 9.2 declares, in advance of any collection: the hypothesis, the null, the primary endpoint, the effect size (`log2FC = 0.585`, fold change 1.5), the sample size (`n = 21` per arm, from the power analysis of Section 5.11), the eight arms and the full control set, the delivery method, the matching requirements on the permuted twin, the blinding scheme, the frozen analysis with block permutation for the shuffled-input control, the stopping rule, and the decision rules for POSITIVE, NEGATIVE and UNINTERPRETABLE outcomes.

The expected result is recorded in advance as the null, because requirement D3 has no candidate implementation and this author does not have one. **A programme-level abandonment condition is also recorded in advance: if a preregistered 73-condition three-stage adaptive design returns null with the full control set, the sequence-specific acoustic channel is withdrawn from this corpus and the pipeline of Section 6 stands solely as an encoding and sonification method with no demonstrated biological effect. No experiment has been run, so no rule has yet been met or missed.**

31.14 Appendix A Works named in the text without a bibliography key

The bibliographic database governing this corpus contains no entries for the following, and **no key has been attached to any of them anywhere in this paper.** They are listed with the details the source manuscript supplies so that a reader can locate them, and their absence is recorded as a gap in the corpus's apparatus rather than concealed (Section 3.7, limitation L-1). **Nothing in this list should be read as a citation; it is a statement of what is missing.**

On "junk DNA" and genome content. Ohno (*Brookhaven Symposia in Biology* 23:366–370, 1972), the mutational-load argument and the origin of the term. The International Human Genome Sequencing Consortium (*Nature* 409:860–921, 2001) for the transposable-element fraction. The ENCODE consortium (*Nature* 489:57–74, 2012) on biochemical activity, and Djebali and colleagues (*Nature* 489:101–108, 2012) on pervasive transcription. **The critiques that must be printed alongside them:** Graur and colleagues (*Genome Biology and Evolution* 5:578–590, 2013) and Doolittle (*PNAS* 110:5294–5300, 2013) on why biochemical activity is not function in the selected-effect sense; Rands, Meader, Ponting and Lunter (*PLoS Genetics* 10:e1004525, 2014) for the approximately 8.2 per cent constrained estimate; Eddy (*Current Biology* 22:R898–R899, 2012) for the compact statement of the C-value objection, and T. Ryan Gregory's onion test. The genome sizes of *Fugu rubripes*, *Allium cepa*, *Protopterus aethiopicus* and *Paris japonica* used in Section 4.3.

On regulatory structure. Jacob and Monod (*Journal of Molecular Biology* 3:318–356, 1961) on operon logic. Britten and Davidson (*Science* 165:349–357, 1969) on repetitive sequence as regulatory substrate. McClintock (*PNAS* 36:344–355, 1950) on controlling elements, and the 1983 Nobel Prize. Chuong, Elde and Feschotte (*Science* 351:1083–1087, 2016) on endogenous retroviruses supplying inducible enhancers; Lynch and colleagues (*Nature Genetics* 43:1154–1159, 2011) on transposon- mediated rewiring.

On syncytin. Mi and colleagues (*Nature* 403:785–789, 2000); Blaise, de Parseval, Bénit and Heidmann (*PNAS* 100:13013–13018, 2003); **Dupressoir and colleagues (*PNAS* 106:12127–12132, 2009), the knockout, which is the causal test the item depends on**; Lavialle and colleagues (*Philosophical Transactions of the Royal Society B* 368:20120507, 2013) on independent convergent capture.

On latency and activation. Gurdon (*Journal of Embryology and Experimental Morphology* 10:622–640, 1962); Takahashi and Yamanaka (*Cell* 126:663–676, 2006); Takahashi and colleagues (*Cell* 131:861–872, 2007); the 2012 Nobel Prize in Physiology or Medicine. Lambert and colleagues (*Cell* 172:650–665, 2018) for the transcription-factor census. Waddington, *The Strategy of the Genes* (1957), for the landscape these results inverted.

On storage, addressing and mechanotransduction. Church, Gao and Kosuri (*Science* 337:1628, 2012); Goldman and colleagues (*Nature* 494:77, 2013); Erlich and Zielinski (*Science* 355:950, 2017). Zetsche and colleagues (*Cell* 163:759, 2015) and Fonfara and colleagues (*Nature* 532:517, 2016) on self-processing arrays; Qi and colleagues (*Cell* 152:1173, 2013) and Gilbert and colleagues (*Cell* 159:647, 2014) on catalytically dead effectors. The PIEZO1/PIEZO2 literature and the 2021 Nobel Prize; Tajik and colleagues (*Nature Materials* 15:1287, 2016); Ibsen and colleagues (*Nature Communications* 6:8264, 2015) and Duque and colleagues (*Nature Communications* 13:600, 2022).

On physics and measurement named without keys. Frey (*Journal of Applied Physiology* 17:689–692, 1962) on the microwave auditory effect. Purcell, *Life at Low Reynolds Number* (*American Journal of Physics* 45:3–11, 1977), for the overdamped regime of Section 5.5. Banyay, Sarkar and Gräslund (*Biophysical Chemistry* 104:477–488, 2003) for the infrared band-centre compilation. Glasberg and Moore (1990) for the equivalent-rectangular-bandwidth formula used in Sections 5.5 and 6.7. Hart and colleagues (*Journal of Computational Biology* 20:970–978, 2013) for the minimum-detectable-fold-change formula of Section 5.11. ISO 16:1975 for the standardisation of the reference pitch discussed in Section 6.8.

Precedent material named without keys. Plato, *Meno*, on recollection; Aristotle, *Metaphysics* Θ, on potentiality and actuality; the Stoic *logos spermatikos* and Augustine's *rationes seminales*. Psalm 118:22 and the alchemical *prima materia* topos. The Talmudic *teiku*; Maimonides, *Guide for the Perplexed* I.58, and the *via negativa* of Pseudo-Dionysius. Kepler, *Harmonices Mundi* (1619), Book V, and the *musica universalis* tradition. The Pythagorean monochord, transmitted through Nicomachus and Boethius. Schönberger (1973) for the codon–hexagram identification rejected in Section 4.8. **Every one of these is precedent, priced at one to three bits on the corpus's own scale, and none is evidence about anything in this paper** [Curzi, 2026w].

Chapter 32

Conjugate Uncertainty at Three Levels: The Kennard Bound, the Gabor Limit, and a Badged Information-Budget Analogue for Forecasts

Abstract. We separate three reciprocal relations that are routinely conflated. First, the quantum relation sigma_x sigma_p >= hbar/2 is a proven theorem (Kennard 1927; Weyl 1928; Robertson 1929) about state preparation, not about measurement disturbance; we give the three-line Fourier derivation, show Gaussians saturate it, and recover the Bohr radius 5.291772e-11 m and -13.605693 eV from it. Second, deleting the substitution p = hbar k from that derivation leaves the Gabor limit dt.df >= 1/(4 pi) = 0.079577471545948, a classical theorem binding every wave including sound; a Gaussian window saturates it to ten significant figures, the DFT bin 1/T sits 4 pi = 12.566 times above the floor and a Hann main lobe 16 pi = 50.265 times above it. Third, the identification of a foreseen future's uncertainty with quantum conjugacy is declared an analogy and badged as such: it is explicitly not claimed to follow from quantum mechanics, no hbar appears in it, and no quantum result is taken to license any claim about belief or mind. Its quantitative content is classical counting: dt.ds >= T.S/2^B, in which each additional bit halves the product, and sharpening a forecast from a year and ten classes to a month and three classes costs 5.3219 bits.

Keywords: uncertainty relation; time-frequency limit; Fourier duality; forecast resolution; information budget; prediction horizon

Series: B. Engineering the Surplus **Paper:** CCP-032

32.1 1 Introduction

Three reciprocal relations circulate under one name. The first is a theorem of quantum mechanics, $\sigma_x \sigma_p \geq \hbar/2$, proved by Kennard in 1927 [Kennard, 1927] and generalised by Robertson two years later. The second is a theorem of classical Fourier analysis, $\Delta t \, \Delta f \geq 1/(4\pi)$, stated in its communications form by Gabor [1946]. The third is not a theorem at all: it is a proposal, advanced by this work, that the uncertainty attaching to a foreseen event has the same *structure*

as a conjugate pair. **The three are routinely conflated, and the conflation runs in the expensive direction** — the classical theorem is mistaken for a diluted version of the quantum one and therefore dismissed in designs where it binds, while the proposal borrows the quantum theorem's authority for claims about belief that no quantum result supports.

This paper separates them, states each at the level of warrant it actually has, and prices the separation in numbers. Its principal quantitative results are the Gabor limit's application to the corpus's own rendering problem — where the minimum resolvable gap in a folded lattice collapses by a factor of eight at a threshold this paper corrects from $N = 20$ to $N = 17$ — and an information-budget bound, $\Delta t\,\Delta s \geq TS/2^B$, which is exact, classical, and derived from counting in two lines.

The organising discipline is `RIGOR_STANDARD` §8's three-way label, applied before any borrowing is made rather than defended afterwards. Each of the three relations is declared, at first use, as a derivation, as a shared mathematical form, or as an analogy. The quantum relation is a derivation *within physics* and is used here only expositorily. The Gabor limit is a derivation *for acoustic signals*, because a rendered tone genuinely satisfies its premises — square-integrability and finite second moments in both domains — and the premises are shown, not asserted. The forecasting application is a **shared mathematical form** where the counting bound is concerned and an **analogy, never evidence**, where the correspondence to Heisenberg is concerned. **The analogy contributes no number anywhere in this paper**, and Section 7.7 states that as a result rather than as a caveat.

Why the exposition of the quantum case is retained rather than cut. The most common reaction to a classical time–bandwidth constraint is to assume it must be quantum mysticism because it resembles Heisenberg's relation. That reaction is defeated only by stating Heisenberg's relation correctly first, including the part most secondary accounts get wrong: **the relation is a property of states, not a statement about a measurement disturbing what it measures.** The disturbance account is not a simplification of the theorem; it is a different claim, it fails as a derivation, and the specific error–disturbance inequality Heisenberg wrote in 1927 has been experimentally violated. Section 4 establishes this at length, and Sections 5 and 6 are believable because of it. The corpus makes the same move elsewhere, refusing quantum licence for its own postulates and computing decoherence timescales rather than gesturing at them [Curzi, 2026aa].

Four contributions are made relative to the surrounding literature rather than relative to the framework. **First**, the two derivations are given side by side with the single line that separates them isolated: the substitution $p = \hbar k$. Delete it and Kennard's bound becomes Gabor's; nothing else changes. **Second**, the Gabor limit is carried into the corpus's frequency framework until the design changes, producing an audible-band ceiling of 25 lattice elements, a render-rate ceiling of 583.19 symbols per second under a common transposition, and a discontinuous collapse of the folded lattice's minimum gap whose thresholds this paper corrects. **Third**, the information-budget bound is derived and its exchange rate computed: each additional bit halves the product $\Delta t\,\Delta s$, and sharpening a forecast from a year and ten classes to a month and three costs 5.3219 bits. **Fourth**, the refusal is stated as a result: no quantum mechanics is used to license any claim about belief, mind, or future events, and the paper's quantitative content is shown to survive the removal of the analogy entirely.

None of this is a claim about human beings, and none of it is measured on data collected for this paper. Every number below is either an exact consequence of a definition, a value computed numerically from a stated function, or a quantity quoted from a published source with its origin named. Where the corpus's own constants enter — most importantly the reference constant $k = (9/8)^2/\varphi^2$, which is *defined*, not measured, and whose expression in hertz is a separate and stronger claim — the badge is carried in the sentence that uses them.

Section 2 states the hypotheses. Section 3 gives the method, including the analogy labels, the entropy identification, the effective-count identification, and the identification problem the framework cannot answer. Sections 4, 5 and 6 develop the three levels. Section 7 reports the results, Section 8 discusses what does and does not follow, Section 9 verifies the chain from minimal postulates, Section 10 states the limitations and Section 11 concludes. Appendices A

to C carry three worked reviews that exercise the chain end to end.

32.2 2 Hypotheses

Much of what follows consists of theorems, and a theorem cannot be falsified by observation. Those are labelled as derivations where they appear and are recorded in Section 9's chain as `ARITHMETIC` or `DEFINITION`. What *can* be wrong are the paper's claims about where the theorems bind, what the corpus's own lattice arithmetic yields, and whether the forecasting bound has any measurable content. Each hypothesis below names the observation that would defeat it.

> **H1 — The binding constraint on the corpus's acoustic rendering is classical, not quantum. Predicted:** for every rendering design in the corpus's frequency framework, the quantum relation is negligible by many orders of magnitude while the Gabor limit is within a factor of ten of the design point, so that changing the design changes only under the classical bound. **Falsified by:** a rendering regime in which $\hbar$ enters a numerically consequential way, or one in which the Gabor floor is more than four orders of magnitude below every gap the design must resolve, which would make the constraint irrelevant rather than load-bearing.

> **H2 — The folded lattice's minimum gap collapses discontinuously in N, and the published thresholds are wrong.** The corpus previously placed the collapses at $N = 20$ and $N = 48$. **Predicted:** on exact rational arithmetic over every n^2 for $n = 1 \dots N$ folded into one octave, the minimum adjacent ratio takes the value $50/49$ up to $N = 16$, $289/288$ from $N = 17$, and $1682/1681$ from $N = 41$; the rate penalty across the first cliff is exactly 8 and across the second exactly 4. **Falsified by:** an exact-rational recomputation returning the published boundaries, or returning a smooth rather than a stepped dependence on N. **Result: confirmed, and the previously published thresholds are corrected in Section 7.6.**

> **H3 — The forecasting bound's content is classical counting, not physics. Predicted:** every number reported in Section 6 is reproducible from $\log_2$ arithmetic, from a Fourier transform, or from a Lyapunov exponent, and none of them changes if the analogy to Heisenberg is deleted from the paper. **Falsified by:** any quantity in Section 6 that cannot be obtained without invoking a quantum result. **Result: confirmed; Section 7.7 reports the analogy's numerical contribution as zero.**

> **H4 — The evidence budget B is estimable in principle and has not been estimated here.** The bound $\Delta t\, \Delta s \geq TS/2^B$ is vacuous without an independently obtained B. **Predicted:** a published forecast record carrying T, S, Δt and Δs per forecast, with B estimated from realised out-of-sample discrimination, either satisfies the bound or does not, on the record. **Falsified by:** a demonstration that no operationalisation of B independent of the forecasts themselves is possible, which would make the bound a tautology. **Result: NOT TESTED. No forecast record was collected for this paper and no B was estimated from data.** This is stated in the same plain language a confirmation would receive, and it is the paper's largest empirical gap.

H5 is a statement the paper deliberately declines to make. No hypothesis is advanced that the uncertainty of a future event *is* a quantum uncertainty, because no observation could bear on it. That proposition is recorded in Section 6.1 as an **analogy** — a modelling stance, not a hypothesis — precisely as `PAPER_SPEC` §2 requires when a candidate hypothesis admits no falsifier.

32.3 3 Method

32.3.1 3.1 What the paper attempts, and what it does not

The paper is an audit and an application, not a discovery. Two of the three relations it handles are established theorems whose proofs are reproduced because the reproduction is what separates them; the third is the author's own proposal, and the method for handling it is to strip it of everything that is not independently derivable and report what is left. **The test of the third level is therefore not whether the analogy is apt but whether the quantitative forms survive its removal.** Section 7.7 reports that they do, which is the strongest thing that can honestly be said for it.

No dataset was collected. Numerical values reported as computed — the saturation of the Gabor bound by a Gaussian, the divergence of the rectangular window's second moment, the folded lattice's exact rational gaps — are outputs of stated functions evaluated at stated parameters, and the evaluation is reproducible from the specification given in the text. Numerical values quoted from physics — the Bohr radius, atomic lifetimes, the reduced Planck constant — are quoted from the standard literature with their origin named, and carry that literature's uncertainty rather than any of this paper's.

32.3.2 3.2 Analogy discipline, declared in advance

`RIGOR_STANDARD` §8 requires every borrowing from physics to be labelled, at first use, as one of three things. The labels for this paper are fixed here so that no later passage can be read as claiming more than the label allows. The corpus states the same discipline for itself [Curzi, 2026aa], and the failure mode it guards against — a category-two borrowing collecting a category-one warrant — is exactly what makes physics-flavoured social science distrusted [Mirowski, 1989, Cartwright, 1999].

Borrowing	Label	What licenses the label
$\sigma_x \sigma_p \geq \hbar/2$ applied to quantum states	**derivation**	the premises are the axioms of quantum mechanics and the systems are quantum; used here expositorily only
$\Delta t \, \Delta f \geq 1/(4\pi)$ applied to a rendered acoustic signal	**derivation**	the premises are square-integrability and finite second moments in both domains; Section 5.2 shows a rendered tone satisfies the first and states which envelopes fail the second
$\Delta t \, \Delta s \geq TS/2^B$ applied to a forecast	**derivation**, but of a counting result, not a physical one	it follows from $\log_2$ of a cell count; no physics enters and none is claimed
the Gabor limit applied to a *belief* about recurrence	**shared mathematical form**, with one badged modelling choice	the same inequality governs the amplitude and its transform; the step of writing a forecast as an amplitude is a choice, not a discovery, and is marked in the sentence that makes it

Borrowing	Label	What licenses the label
the horizon law $t_{\text{pred}} = t_d \log_2(D_{\text{tol}}/D_0)$ applied outside a system with a measured positive Lyapunov exponent	**analogy**	motivation only; invoking it without an estimated exponent is failure mode FM-6.6
"the future is a conjugate variable"	**analogy**	a structural correspondence with no formal warrant, offered as a way of seeing and never as evidence
the headroom correspondence of Section 6.14	**shared mathematical form**	both are a requested increment tested against an available budget; it is a claim about model shape and not about physics

Two specific corrections that the corpus makes elsewhere and this paper keeps. The uncertainty relation is a fact about state preparation and not about measurement disturbance, and it is stated that way in every passage below. And where a variational or extremal argument appears — the minimisation that recovers the Bohr radius — it licenses no teleology: the minimiser is where a derivative vanishes, and nothing in the system is seeking it.

32.3.3 3.3 Entropy: which entropy, in every instance

Every entropy in this paper is Shannon entropy [Shannon, 1948], computed as $H = -\sum_i w_i \log_2 w_i$ over a normalised weight vector and reported in **bits**. **No thermodynamic entropy appears anywhere in this paper**, and no result below carries units of energy per temperature. The two objects share a name and a functional form and are different quantities with different units; conflating them is a standing error the corpus's information-theory material is written to avoid [Cover and Thomas, 2006, Khinchin, 1957]. Where the source material for Section 6.7 speaks of an ensemble's entropy, it means Shannon entropy in bits, and the value 1.8150 bits is a Shannon entropy.

The one place where a physical entropy could legitimately enter — the thermodynamic cost of erasing a bit [Landauer, 1961] — is not used, because no claim in this paper depends on the physical cost of the forecaster's information.

32.3.4 3.4 Effective counts: saying which one, every time

Two different effective counts appear in the corpus and they are not the same object. The distinction is stated here because both are touched below.

- **2^H, the exponential of a Shannon entropy in bits, is the Hill number of order one** [Hill, 1973, Jost, 2006, Rényi, 1961]. The ensemble effective scenario counts reported in Sections 6.7 and Appendix C — 3.5185 of a nominal 4, and 3.7417 of a nominal 4 — are Hill numbers of order one and are grounded by that literature exactly.
- **$N_{\text{eff}} = 1 + (N-1)u = e^F$, the corpus's affine effective count [Curzi, 2026], is NOT a Hill number of any order.** It is a linear interpolation between 1 and N and no literature in ecology, economics or information theory contains it. It appears in this paper only in Section 5.13, where the surplus bridge $f_{\text{cross}} = k\,e^{2F}$ is differentiated; the badge travels with it, and no Hill-number citation is attached to it anywhere.

The two quantities coincide numerically nowhere except at $u = 0$ and $u = 1$, and a reader who reaches for the Hill-number literature to support the second will find it does not apply.

Saying so is cheaper than being caught.

32.3.5 3.5 Identification, the reflection problem, and the Lucas critique

Section 6 touches the corpus's reflexive material at two points — the correspondence between the evidence budget and the storage law's headroom (Section 6.14), and the reference to the corpus's conditional postulate about belief and reality (Section 6.4) — and both inherit an objection this framework cannot answer.

The corpus's reflexive specification, in which a group's state is a function of its members and the members' states are functions of the group, is a linear-in-means reflexive model, and its endogenous effect is not identified from observational data [Manski, 1993]. No quantity of observation separates the endogenous effect from contextual effects or from correlated unobservables, and influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011, Lyons, 2011]. **This paper adopts no remedy.** Randomisation, instrumentation and partial-population variation are all unavailable to it, and the fourth route — stipulating the reference group and defending the stipulation, for which Brock and Durlauf [2001] is the template — is not attempted here. **Every statement in this paper that touches the reflexive loop is therefore a modelling position rather than an identified empirical result**, and is written that way in the text as well as here.

The Lucas critique applies to this paper with unusual directness, because its subject is forecasting [Lucas, 1976]. A framework asserting that systems change their behaviour when modelled cannot exempt its own parameters from the same reflexivity. Section 6's evidence budget B is a case in point: a published forecast changes the behaviour of the forecast object, so B estimated from a historical record need not describe the record that follows publication. Self-fulfilling and self-defeating prophecy are real and documented [Merton, 1948, Jussim and Harber, 2005], and a reflexive fixed point has been observed in the field with dates and price data [MacKenzie and Millo, 2003, MacKenzie, 2006]. **The honest position this paper takes is that B is not invariant under its own publication, that the bound is therefore a constraint on a forecaster at a moment rather than a stable parameter of a domain, and that no defence of parameter stability is offered.** The measured social-prediction literature is consistent with the modest reading: predictive accuracy in social systems is low and the ceiling is structural rather than methodological [Hofman et al., 2017, Salganik et al., 2020, Watts, 2014].

32.3.6 3.6 Scale type, and the conjoint test that was not run

`RIGOR_STANDARD` §5 requires every measured construct to declare its scale type [Stevens, 1946, Krantz et al., 1971]. Three constructs in this paper need one.

- **Time and frequency are ratio scales with true zeros**, and every operation performed on them below — quotients, products, logarithms of ratios — is licensed.
- **The specification axis s is the exposed one.** Where the outcome space is categorical, s is a **count of classes**, not a distance, and the bound of Section 6.9 is then a counting bound. Where the outcome space is numeric, s requires a ratio scale for the quotient $S/\Delta s$ to mean anything. **The source framework is silent on which case obtains in general, and this paper does not supply the missing determination**; it states the restriction and reports the bound as valid under either reading with the reading named.
- **The evidence budget B is treated as a ratio-scaled quantity in bits.** That is licensed for a Shannon quantity by construction, but the *estimate* of B from realised forecast discrimination inherits whatever scale the discrimination measure has, and no such estimate is made here.

Michell's critique is the strongest objection to the framework this paper belongs to [Michell, 1997, 1999]: psychology routinely assumes that its attributes possess additive quantitative structure without testing for it. The formal test is additive conjoint measurement [Luce and

Tukey, 1964]. **It has not been run on any construct in this paper, and this paper does not run it.** The quantities that carry the paper's numbers — time, frequency, bits, and cents — are physical or combinatorial and are not exposed to the objection; the one that is exposed, the specification axis of a forecast over a non-numeric outcome space, is explicitly demoted to a count for that reason. **That demotion is a concession, not a defence**, and it is repeated in Section 10.

32.3.7 3.7 Scaling claims, and the procedure not run

Two asymptotic statements appear below and both are **defined by the model rather than estimated from data**, which `RIGOR_STANDARD` and Clauset et al. [2009] require to be said explicitly.

- The rectangular window's power spectrum falls as f^{-2}, so its spectral second moment diverges logarithmically. **The exponent is an exact property of the** sinc **function and is not an estimate.** No power-law fitting procedure was run, because there is no sample: the spectrum is known analytically.
- Under a common transposition the fractional lattice gap falls as $2/n$ and the required cycle count grows as $n/(8\pi)$. **Both are asymptotics of the exact expression** $n^2/(4\pi(2n+1))$ **and are exact consequences of the square law, not fitted exponents.**

No heavy-tail or scale-free claim is made anywhere in this paper on empirical data, and the maximum-likelihood-plus-goodness-of-fit procedure of Clauset et al. [2009] was therefore not run. It is named here so that its absence is a stated decision rather than an omission; the literature on mis-fitted power laws [Stumpf and Porter, 2012, Broido and Clauset, 2019] is the reason the distinction is worth a paragraph.

32.3.8 3.8 What is presupposed

The essence law $f(n) = kn^2$ with $k = (9/8)^2/\varphi^2$ is presupposed as a badged postulate of the framework and is not re-derived or defended here [Curzi, 2026]. **Its two components are real and its product is not**: $9/8$ is the Pythagorean whole tone [Barbour, 1951, von Helmholtz, 1885] and $\varphi^2 = \varphi + 1$ is Euclid VI.3 [Euclid, 1956, Vajda, 1989], but their quotient expressed as a frequency has no precedent in any literature. Markowsky [1992] is cited here, as everywhere the golden section appears in this corpus, because it is the standing correction to golden-ratio claims and the reader is entitled to see that the author knows what one usually looks like. Octave equivalence, on which the folding operation rests, is formalised by Shepard [1982] and its perceptual limits are given by Burns [1999], Demany et al. [2013], Hoeschele et al. [2012].

The storage law $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$, its fixed point and its headroom form are presupposed from Curzi [2026a]. The surplus functional and its affine effective count are presupposed from Curzi [2026]. The corpus's seven-gate formula-validation protocol, whose gates are run on Section 6.9's bound, is specified in Curzi [2026t], and its demotion of the framework's central formula to ordinal scale is inherited rather than re-argued. The two-column concordance discipline — that agreeing traditions stop adding evidence after the first, because precedent is not evidence — is specified in Curzi [2026w] and is applied without restatement of its derivation.

32.3.9 3.9 Standing honesty clauses

These govern every use of this paper's results, in this corpus and outside it.

1. **The uncertainty relation is a property of states, not of instruments.** It holds for unmeasured systems and is not the observer effect.

2. σ_x **and** σ_p **are ensemble quantities.** Two numbers from one system have no standard deviation and cannot test the relation.
3. **The Gabor limit is classical and contains no quantum content.** Any presentation of it as an application of quantum mechanics to sound is wrong.
4. Δt **and** Δf **are root-mean-square spreads**, not window lengths and not peak widths. The conversions differ by factors up to 50.
5. **The essence frequencies are essence frequencies, not physical resonances**, and k is a chosen reference constant, not a measured one [Curzi, 2026]. The Gabor limit constrains how such a number may be turned into a sound and says nothing about whether the law generating it is true.
6. **Nothing in this paper licenses any claim about mind, belief, choice, or future events from quantum mechanics.** Section 6's analogy is badged in every section that states it and derives its content from counting and from Fourier duality.
7. **No clinical, medical, therapeutic or physiological claim is made anywhere in this paper.**

32.4 4 Level one: the Kennard bound, stated exactly

This section reports established physics and corrects a widespread misstatement of it. **It introduces no postulate of this framework.** Its function is to establish, in front of the reader, that the paper can state the famous case exactly, correct the standard error, and refuse the standard over-reach — because Sections 5 and 6 are believable only if it can.

32.4.1 4.1 The relation

```
PROVEN THEOREM of quantum mechanics (Kennard 1927; Weyl 1928; Robertson 1929).

    sigma_x . sigma_p  >=  hbar / 2

sigma_x   standard deviation of the position distribution of a quantum state
sigma_p   standard deviation of the momentum distribution of THE SAME state
hbar      h / (2 pi),  the reduced Planck constant
```

With the 2019 SI definition $h = 6.62607015 \times 10^{-34}$ J s exact:

```
hbar     = 1.054571817646e-34   J s
hbar / 2 = 5.272859088231e-35   J s
hbar     = 6.582120e-16         eV s
```

The misreadings hide in the symbols, so each is read carefully.

σ_x **is a standard deviation, not an experimental error bar.** It is the square root of the variance of the probability distribution $|\psi(x)|^2$ that the state assigns to position. Written out, there is no apparatus in it:

```
sigma_x^2 = INTEGRAL |psi(x)|^2 (x - <x>)^2 dx
sigma_p^2 = INTEGRAL |phi(p)|^2 (p - <p>)^2 dp        phi = Fourier transform of psi
```

Both are computed from the same wavefunction, and neither refers to a measurement, an instrument, a photon, a microscope or an observer. Both integrals can be evaluated for a state that has never been measured and never will be, and the product obeys the inequality.

The two distributions are Fourier transforms of one another, which is the whole content:

```
phi(p) = (1 / sqrt(2 pi hbar)) INTEGRAL psi(x) exp(-i p x / hbar) dx
```

A function and its Fourier transform cannot both be narrow; that is a theorem of analysis [Oppenheim et al., 1996]. **Quantum mechanics contributes exactly one thing to it: the identification of the conjugate variable of x with $p/\hbar$, which puts $\hbar$ into the bound and nothing else.**

32.4.2 4.2 The general form and its state dependence

Robertson proved the general statement for any two observables — *Physical Review* 34 (1929), 163–164:

```
sigma_A . sigma_B  >=  (1/2) | < [A, B] > |          [A, B] = AB − BA
```

For position and momentum $[x, p] = i\hbar$, a constant, so the right-hand side is $\hbar/2$ for every state and the position–momentum case is the exceptional one in which the bound never weakens.

For most observable pairs the bound is state-dependent and can be vacuous. Spin-$1/2$ makes this concrete, with $[S_x, S_y] = i\hbar S_z$:

```
state |+z> :   <S_z> = hbar/2
               Robertson bound = (hbar/2)(hbar/2) = 2.780304e−69  J^2 s^2
               actual sigma_Sx . sigma_Sy = (hbar/2)^2 = 2.780304e−69
               SATURATED.

state |+x> :   <S_z> = 0
               Robertson bound = 0
               actual sigma_Sy . sigma_Sz = (hbar/2)^2 > 0
               BOUND IS TRIVIALLY TRUE AND SAYS NOTHING.
```

Schrödinger strengthened Robertson in 1930 by adding the anticommutator term, which repairs some but not all of the vacuous cases. **The point is not decorative: for a general pair, the relation sometimes forbids joint sharpness and sometimes says nothing at all, and which one obtains depends on the state.** The position–momentum case is quoted so often precisely because it is the case in which the bound is unconditional.

32.4.3 4.3 Warrant: what is proved, what is cited, what is postulated

PROVED, given the axioms of quantum mechanics. $\sigma_x \sigma_p \geq \hbar/2$ for every normalisable state [Kennard, 1927]; the general commutator form (Robertson 1929); the strengthened form with the covariance term (Schrödinger 1930); and equality if and only if the state is a Gaussian wave packet, which follows from the equality condition in the Cauchy–Schwarz step of Section 4.7's derivation.

CITED, measured in the literature. Single-slit diffraction, on any optical bench. Matter-wave interference for large molecules, reported for C_{60} by Arndt, Nairz, Vos-Andreae, Keller, van der Zouw and Zeilinger, *Nature* 401 (1999), 680–682. Natural linewidths of atomic transitions (Section 4.11). Experimental violation of Heisenberg's *original* error–disturbance inequality, by Erhart, Sponar, Sulyok, Badurek, Ozawa and Hasegawa, *Nature Physics* 8 (2012), 185–189, and by Rozema, Darabi, Mahler, Hayat, Soudagar and Steinberg, *Physical Review Letters* 109 (2012), 100404.

POSTULATED in this section: nothing.

> **Bibliographic disclosure.** Of the physics primary sources named in this section, only Kennard carries a key in this corpus's bibliography [Kennard, 1927]. Robertson (1929), Weyl (1928), Schrödinger (1930), Ozawa (2003), Busch, Lahti and Werner (2013), Arndt et al. (1999), Erhart et al. (2012) and Rozema et al. (2012)

> are named in the text with journal, volume, year and pages and carry **no citation key**, because no entry for them exists in `bib/external.bib` and none was invented. The same disclosure applies to Lorenz (1963, 1969), Zhang et al. (2019), Brier (1950), Murphy (1973), Carlson (1985) and Wier, Jesteadt and Green (1977) in Sections 5 and 6. Every one of them is a standard, locatable source; the gap is in this corpus's bibliography, not in the record.

32.4.4 4.4 The misconception, named

The familiar account runs:

> *To see where an electron is, you have to bounce something off it. The gentlest thing you can bounce off it is one photon. But a photon carries momentum, so bouncing it off the electron changes the electron's momentum. The more precisely you want the position, the shorter the wavelength you need, the more momentum the photon carries, and the more you disturb what you were trying to measure.*

Call this **the disturbance story**. It has four separable defects.

Defect 1 — it makes the relation epistemic when it is ontic. The disturbance story says the electron has a position and a momentum and that clumsiness prevents both being learned. **A state with sharp momentum has no position at all — not an unknown position, not a hidden position, none.** Its position distribution is spread over all space, and the spread is a fact about the state. This is testable and has been tested: if particles had definite values for non-commuting observables that measurement merely failed to reveal, correlations between distant measurements would satisfy Bell's inequality, and they do not; the loophole-free experiments of 2015 closed the last significant escape routes. **The disturbance story is a hidden-variable story in casual clothes, and local hidden-variable models are excluded by experiment.**

Defect 2 — it predicts the wrong thing about ensembles. Prepare ten thousand identical systems; measure position on five thousand and momentum on the other five thousand. **No system is measured twice, and nothing measured is subsequently disturbed.** With $m = 5000$ samples the standard error of an estimated standard deviation is about $\sigma/\sqrt{2(m-1)}$, which is **1.0001 per cent of** σ. The product of the two estimates obeys the bound. **There is no disturbance anywhere in that protocol and the bound is still there** — which is sufficient to refute the disturbance story *as a derivation*.

Defect 3 — it cannot explain why classical waves have the same relation. Sound waves obey a mathematically identical relation with no photons, no $\hbar$, no quanta and no disturbance. Section 5 is about that relation, and it is measurable on an inexpensive oscilloscope. **The disturbance story attributes the relation to a mechanism absent in the cases where the relation is easiest to demonstrate.**

Defect 4 — Heisenberg's own quantitative version of it is false, and this has been measured. Heisenberg's 1927 paper introduced the gamma-ray microscope and wrote an inequality relating the *error* $\epsilon(x)$ of a position measurement to the *disturbance* $\eta(p)$ it inflicts on momentum — a genuinely different inequality from Kennard's, because ϵ and η are properties of an apparatus while σ_x and σ_p are properties of a state. Ozawa (*Physical Review A* 67, 2003, 042105) showed that Heisenberg's error–disturbance form is not universally valid and derived a corrected one:

```
epsilon(A) eta(B) + epsilon(A) sigma_B + sigma_A eta(B)  >=  (1/2) | <[A,B]> |
```

The extra two terms are what Heisenberg omitted, and when they are large the naive product can drop below $|\langle[A,B]\rangle|/2$. **It does.** Erhart et al. (2012) measured neutron spin components and reported the naive product violated while Ozawa's relation held; Rozema et al. (2012) reached the same conclusion by weak measurement on photon polarisation.

The rest of the story is reported because the area is active. Busch, Lahti and Werner (*Physical Review Letters* 111, 2013, 160405) proved a Heisenberg-type error–disturbance relation

that does hold, on a worst-case rather than a state-specific root-mean-square definition of "error" and "disturbance". The disagreement is about what the two words should mean, not about the mathematics, and both camps agree on every experimental number. **What is not in dispute: the state-based relation is settled, exact and never violated, while the measurement-based error–disturbance relation is a different, harder, definition-sensitive question on which the naive 1927 form fails. Teaching the first by way of the second is teaching a proved theorem with a disproved heuristic.**

32.4.5 4.5 Where the relation comes from: three lines of Fourier

The derivation is short enough to give in full, and giving it in full is the cure for the disturbance story, because there is no apparatus anywhere in it. Work in units where the conjugate variable is $k = p/\hbar$; let $\psi(x)$ be normalised and shift so that $\langle x \rangle = \langle k \rangle = 0$.

Step 1 — build a non-negative quantity. For any real a,

```
I(a) = INTEGRAL | a x psi(x) + dpsi/dx |^2 dx   >=   0
```

It is an integral of a squared modulus, so it cannot be negative. **That is the entire physical input: none.**

Step 2 — expand.

```
INTEGRAL |x psi|^2 dx             = sigma_x^2
INTEGRAL |dpsi/dx|^2 dx           = sigma_k^2      (by Parseval / Plancherel)
2a Re INTEGRAL x psi* dpsi/dx dx  = -a             (integrate by parts, using normalisation)

so    I(a) = a^2 sigma_x^2 - a + sigma_k^2  >=  0   for every real a.
```

Step 3 — a non-negative quadratic has a non-positive discriminant.

```
1 - 4 sigma_x^2 sigma_k^2  <=  0
=>   sigma_x sigma_k  >=  1/2
=>   sigma_x sigma_p  >=  hbar/2                  since p = hbar k
```

That is the whole proof. It contains an integration by parts, Plancherel's theorem, and the discriminant of a quadratic. It contains no photons, apparatus, observers, disturbance, collapse or consciousness, and no statement about what happens during a measurement.

Equality holds exactly when $I(a) = 0$ at the minimising a, that is when $ax\psi + \psi' = 0$, whose solution is $\psi(x) = C\exp(-ax^2/2)$ — a Gaussian. Gaussians saturate the bound and nothing else does. That fact returns in Section 5 as a design rule.

The step where $\hbar$ enters is Step 3's final substitution, and only there. Everything before it is a theorem about square-integrable functions. **Delete the substitution and what remains is the Gabor limit.** That is not a metaphor: it is literally the same three lines with a different name for the conjugate variable, and Section 5.3 runs them again in the units of sound.

32.4.6 4.6 Every wave has it, and one that can be seen

Send a plane wave of wavelength λ through a slit of width a; the first diffraction minimum sits at $\theta \approx \lambda/a$. Computed at $\lambda = 633$ nm, pattern width on a screen one metre away:

slit width a	half-angle to first null	spread at 1 m
10 um	6.330e-02 rad (3.63 deg)	63.4 mm
5 um	1.266e-01 rad (7.25 deg)	127.3 mm
1 um	6.330e-01 rad (36.27 deg)	733.7 mm
0.5 um	1.266e+00 rad (72.54 deg)	3178.6 mm

Halve the slit and the spread doubles; the product is constant. The slit does not disturb the light — the slit *is* the localisation, and the spread is what localisation costs.

The de Broglie wavelengths that make the quantum version visible or invisible:

```
electron at 1e6 m/s            lambda = 7.2739e-10 m
thermal neutron at 2200 m/s    lambda = 1.7982e-10 m
C60 molecule at 200 m/s        lambda = 2.7701e-12 m
baseball, 145 g at 40 m/s      lambda = 1.1424e-34 m
```

The baseball's wavelength is roughly nineteen orders of magnitude smaller than a proton, which is why the relation does not constrain baseballs — a quantitative answer rather than a hand-wave.

32.4.7 4.7 The ensemble reading, as a protocol

```
PROTOCOL 4-A -- what sigma_x sigma_p >= hbar/2 is a statement about

1. Fix a preparation procedure P. It must be repeatable and must not be conditioned
   on any measurement outcome.
2. Run P n times. There are now n systems in the same state psi.
3. Split the n systems into two arms at random. Do not measure any system twice.
4. On arm A, measure position. On arm B, measure momentum.
5. Compute the sample standard deviation of each arm.
6. The claim under test is: s_x * s_p >= hbar/2, within sampling error.

NOT under test, and not what the relation says:
   - that any single system's position was disturbed by a momentum measurement
   - that any single system had a definite position and momentum before step 4
   - that a joint measurement returning a pair (x, p) is forbidden
     (it is not forbidden; it is noisy, and the noise obeys Ozawa's relation, not this one)
```

Step 3 is load-bearing. It removes disturbance from the experiment entirely and the bound survives. **Any account of the relation that would be falsified by step 3 is wrong.**

32.4.8 4.8 Worked numbers at four scales

All computed from $\hbar/2 = 5.272859088231 \times 10^{-35}$ J s.

Scale 1 — an electron confined to an atom-sized box, $\sigma_x = 1.0 \times 10^{-10}$ m.

```
sigma_p >= hbar / (2 sigma_x) = 5.2729e-25  kg m/s
sigma_v  = sigma_p / m_e      = 5.7884e+05  m/s   = 0.193 % of c
E ~ sigma_p^2 / 2 m_e         = 1.5261e-19  J     = 0.9525 eV
```

Roughly one electron-volt of kinetic energy is forced on the electron purely by confining it to a tenth of a nanometre — chemistry's order of magnitude, and not a coincidence.

Scale 2 — the hydrogen ground state, recovered from the relation alone. Write the energy of an electron confined to radius r as kinetic-plus-Coulomb with $p \sim \hbar/r$ and find the stationary point:

```
E(r) = hbar^2 / (2 m_e r^2)  -  e^2 / (4 pi eps0 r)

dE/dr = 0  =>  r = 4 pi eps0 hbar^2 / (m_e e^2) = 5.291772e-11 m

Bohr radius a0                                  = 5.291772e-11 m    -- identical
E at that radius                                = -13.605693 eV
Measured hydrogen ground-state energy           = -13.605693 eV
```

The relation, used as nothing but an inequality and made stationary, returns the Bohr radius exactly and the hydrogen ground-state energy to six figures. Atoms have the size they have because a wave cannot be squeezed in position without paying in momentum. The variational form licenses no teleology: the derivative vanishes at that radius and nothing is seeking it.

Scale 3 — the nucleus, and why beta-decay electrons do not live there. Confine an electron to $\sigma_x = 1.0 \times 10^{-15}$ m:

```
sigma_p >= 5.2729e-20 kg m/s
p c                   = 98.66 MeV
relativistic KE       = 98.15 MeV
```

Nuclear binding energies per nucleon are of order 8 MeV. **An electron confined to a nucleus would carry roughly a hundred MeV, an order of magnitude more than the nucleus could hold it with** — so the electrons emitted in beta decay are created at emission rather than resident beforehand. **This is the relation ruling out a physical picture, quantitatively, which is what a real constraint does.**

Scale 4 — a gram, and why nobody notices. A one-gram mass localised to one micrometre:

```
sigma_p >= 5.2729e-29 kg m/s
sigma_v >= 5.2729e-26 m/s
        =  1.6641e-18 m per year
        =  9.79e-04 proton diameters per year
```

The forced velocity spread would move it less than one thousandth of a proton diameter in a year. The relation is not silent at macroscopic scale — it is loudly, computably negligible, exactly as the Casimir force is at biological scale [Casimir, 1948, Lamoreaux, 1997, Curzi, 2026aa]. Both are cases where the honest move is to compute the number and show it is nothing.

32.4.9 4.9 Energy and time: a genuinely different animal

$\Delta E\,\Delta t \geq \hbar/2$ is written constantly and understood rarely, and its status differs from position–momentum in a way that is not a technicality.

The problem is that t is not an operator in standard quantum mechanics. Time is a parameter — the label on the Schrödinger equation, not something the state has a distribution over. There is no σ_t computed from ψ the way σ_x is, so **Robertson's theorem does not apply, because there is no commutator $[H, t]$ to put on the right-hand side.** Pauli observed that a self-adjoint operator canonically conjugate to a Hamiltonian would force the Hamiltonian's spectrum to run over the whole real line, unbounded below as well as above; physical Hamiltonians are bounded below, which is why matter is stable, so no such time operator exists for them. The argument has been examined and qualified repeatedly since, but the conclusion stands for the standard framework: **time-of-arrival and similar quantities are handled with positive-operator-valued measures, not with a self-adjoint clock observable.**

There are at least three distinct true statements written with the same symbols:

```
(i)  MANDELSTAM–TAMM.  For any observable A that is not conserved,

         sigma_E . tau_A  >=  hbar / 2 ,    tau_A := sigma_A / | d<A>/dt |

     tau_A is the time for the expectation of A to shift by one of its own standard
     deviations -- a CHARACTERISTIC TIME OF CHANGE, derived from Robertson applied to
     H and A. Mandelstam & Tamm, J. Phys. USSR 9 (1945), 249. This one IS a theorem,
     and it is a statement about how fast a state can evolve, not about clocks.
```

```
(ii) LIFETIME–LINEWIDTH.  A state that decays with lifetime tau has a Lorentzian energy
     profile of full width at half maximum

         Gamma = hbar / tau      i.e.    delta_nu = 1 / (2 pi tau)

     This is a Fourier statement about an exponentially damped amplitude, and it is
     measured routinely in spectroscopy.

(iii) DURATION–BANDWIDTH OF A PULSE.  A light pulse of duration dt has spectral width
     df >= 1/(4 pi dt). This is the GABOR LIMIT (Section 5) applied to light, and it is
     classical. It is often mislabelled as the energy–time uncertainty relation. It is
     not; it is the Fourier relation for the pulse envelope, and it holds for radio waves
     and sound identically.
```

Statement (ii) is the one that earns the relation its keep. Against published atomic lifetimes:

transition	lifetime tau		FWHM linewidth		energy width
H 2p -> 1s	1.596	ns	99.72	MHz	4.124e-07 eV
Ca 1P1 (422.7 nm)	4.6	ns	34.60	MHz	1.431e-07 eV
Na 3p, D2 line	16.2	ns	9.824	MHz	4.063e-08 eV
Rb 5p, D2 line	26.24	ns	6.065	MHz	2.508e-08 eV
a metastable state	1	s	0.159	Hz	6.582e-16 eV

The Mandelstam–Tamm consistency check on the first row: $\Gamma\tau = \hbar$, which is 2.0000 times $\hbar/2$, so the bound is respected with a factor of two to spare, exactly as a Lorentzian rather than Gaussian profile requires. **A transition that lives for one second has a natural linewidth of 0.159 Hz out of an optical frequency of several hundred terahertz — a fractional width around 10^{-15}, which is why metastable transitions are what optical atomic clocks are built on.** The relation is not a limitation being apologised for; it is the design equation for the most precise instruments built.

And note the form of $\delta\nu = 1/(2\pi\tau)$: a reciprocal relation between a duration and a frequency width with a factor of 2π in it. **It is the Gabor limit in fancy dress, and Section 5 takes the dress off.**

32.4.10 4.10 What the relation licenses, and what it does not

LICENSED. (1) A quantitative floor on the joint spread of any conjugate pair. (2) The stability and size of atoms, and by extension the scale of chemistry. (3) Zero-point energy: a harmonic oscillator cannot sit still at the bottom of its well, and $E_0 = \hbar\omega/2$ follows by the same stationarity argument. (4) Exclusion of physical pictures by computation. (5) Linewidth engineering. (6) The general principle that localisation in one domain costs spread in its conjugate — **which is the part that survives the removal of $\hbar$ and becomes Section 5.**

NOT LICENSED, and this list is the reason the section exists. (1) It does not say the observer creates reality; it says nothing about observers. The measurement problem is genuinely open [Curzi, 2026aa, Zeh, 1970], **but it is not open in the direction of licence.** (2) It does not say everything is uncertain: it gives a precise floor on a specific product for a specific pair, and most quantities in most situations are nowhere near it. (3) It does not say measurement disturbs the measured; that is a separate fact with its own separate inequality, whose naive form has been experimentally violated. (4) It does not license macroscopic indeterminacy — compute the number. (5) **It does not license any claim about belief, intention, mind or future events.** Section 6's analogy carries its badge in every sentence that states it and derives its quantitative content from information budgets and Fourier duality, **not from quantum mechanics**, precisely so that it does not need this licence and never asks for it. (6) It does not make the future indeterminate: the Schrödinger equation is deterministic, and what is probabilistic is the outcome of a measurement under the Born rule [Curzi, 2026aa].

32.4.11 4.11 Why a paper on comfort carries a section on $\hbar$

Three of the framework's central results are Fourier trade-offs. A reverberant room smears speech in time, and the room's decay time and its modal bandwidth are reciprocal — $B = 2.2/T_{60}$ is standard room acoustics and is a duration–bandwidth statement [Curzi, 2026e, Sabine, 1922]. What a mechanical wave can carry into tissue is set by bandwidth and time, and the count of independent degrees of freedom in a signal of bandwidth B and duration T is approximately $2BT$, which is Gabor's own result and the reason a channel-capacity analysis is possible at all [Curzi, 2026x, Gabor, 1946]. And essence frequencies rendered as tone cannot be rendered arbitrarily fast, because a short tone has a smeared spectrum [Curzi, 2026,?].

All three are the same theorem. The framework gains nothing from $\hbar$ — Section 5 uses none of it — and gains a great deal from having demonstrated that it can state the famous case exactly, correct the standard error, and refuse the standard over-reach.

32.4.12 4.12 The concordance, with column roles stated

The corpus's concordance discipline requires that the two columns carry different things: **the evidence column carries what could have come out otherwise, and the precedent column carries only that a structure was independently arrived at and preserved. Two sources agreeing does not double the evidence when only one of them is evidence** [Curzi, 2026w].

```
CLAIM: Certain properties come in pairs such that sharpness in one member of the pair
       forces breadth in the other, and this is a property of the thing itself rather
       than of the act of examining it.

(+1) EVIDENCE COLUMN
     Kennard (1927), Weyl (1928), Robertson (1929): sigma_x sigma_p >= hbar/2, proved.
     Single-slit diffraction and C60 interference (Arndt et al. 1999): measured. The
     bound is quantitative, saturated by Gaussians, and falsifiable -- a single
     reproducible preparation with sigma_x sigma_p < hbar/2 would end it.

(-1) PRECEDENT COLUMN -- CARRIES NO EVIDENCE
     The doctrine of coincident opposites: Nicholas of Cusa, De Docta Ignorantia (1440),
     coincidentia oppositorum. The Latin alchemical corpus organises its operations around
     the fixed / volatile pair, where increasing one is understood to decrease the other.
     The Daodejing, chapter 2, states paired opposites as mutually constitutive.

(0p) HOW THEY INTEGRATE, AND WHERE THEY DO NOT
     The older corpus recorded a STRUCTURE: that some pairs resist joint sharpening. It
     recorded no bound, no constant, no derivation and no measurement, and could not have.
     What the concordance establishes is that the structure is findable without
     instruments. It establishes NOTHING about hbar, about atoms, or about the
     correctness of any traditional operation.

BRIDGE CRITERIA:
     C1 structural correspondence  PASS -- both describe a pair with a joint constraint.
     C2 predictive power           FAIL for the (-1) side. It predicts no number and
                                   generates no test.
     C3 practical utility          PARTIAL. Useful as a mnemonic, useless as a calculation.

STATUS: STRUCTURAL PRECEDENT, NOT EVIDENTIAL SUPPORT. The physics stands entirely on the
(+1) column. Remove the (-1) column and nothing in this section changes.
```

One historical note belongs here because it is documented and is often told badly. **Bohr adopted the taiji symbol and the motto *Contraria sunt complementa* for the coat of arms designed when he received the Order of the Elephant in 1947.** That is a fact about Bohr's taste in

emblems. **It is not evidence that Daoist cosmology anticipated quantum mechanics, and the inference is regularly made and is invalid.** A physicist choosing a heraldic device is choosing a heraldic device.

32.4.13 4.13 Requirements

These govern every appearance of the uncertainty relation in work under this framework.

REQ-4.1 State the relation as $\sigma_x \sigma_p \geq \hbar/2$ with σ explicitly named as the standard deviation of a distribution derived from the state.

REQ-4.2 Never derive, motivate or illustrate the relation by the disturbance story. If it is mentioned, it is mentioned as a superseded historical heuristic, with Ozawa (2003) and the 2012 experiments named in the same passage.

REQ-4.3 State the ensemble reading (Protocol 4-A) at first substantive use.

REQ-4.4 When quoting the general Robertson form, state that the bound is state-dependent and can be trivial, and give the spin example or an equivalent.

REQ-4.5 When the equality case is relevant, state that only Gaussian states saturate the bound.

REQ-4.6 Treat the energy–time relation as a separate topic with its own derivation; state that t is a parameter and not an operator, and name which of the three statements is in use.

REQ-4.7 When the lifetime–linewidth relation is used, give $\delta\nu = 1/(2\pi\tau)$ with real numbers and identify it as a Fourier statement about a damped amplitude.

REQ-4.8 Never write "the uncertainty principle" without saying which one. At least four distinct inequalities circulate under that name: Kennard's, Robertson's, Ozawa's and Gabor's.

REQ-4.9 Compute the magnitude before invoking the relation at any scale above the atomic, and if the computed effect is negligible, say so in the same paragraph with the number.

REQ-4.10 No text under this framework may use the uncertainty relation to support a claim about mind, belief, intention, agency, or the determination of future events.

REQ-4.11 Every concordance states its column roles and evaluates C1, C2 and C3 explicitly, including failures [Curzi, 2026w].

REQ-4.12 Where a historical figure's personal symbolism is cited, state explicitly that biography is not evidence.

REQ-4.13 State the forward chain at first use: this relation is a special case of a Fourier theorem, the general case is Section 5, and the general case is what constrains the designs.

32.4.14 4.14 Failure modes

FM-4.1 — THE DISTURBANCE DERIVATION. *Fault:* the relation presented as arising from the momentum kick of the probe. *Signature:* "you have to bounce a photon off it". *Diagnostic:* ask what the relation says about a system nobody measures; the disturbance story has no answer. *Fix:* Protocol 4-A, then the three-line derivation with no apparatus in it.

FM-4.2 — THE HIDDEN-VARIABLE SMUGGLE. *Fault:* assuming the particle *has* both values. *Signature:* "we cannot know both", "the true position". *Fix:* state the ontic reading; local hidden-variable models of this kind are excluded by Bell-test experiments.

FM-4.3 — THE UNIVERSAL SOLVENT. *Fault:* "the uncertainty principle" invoked to license indeterminacy in markets, minds or futures. *Signature:* no pair named, no commutator named, no number computed; the relation appears as a mood rather than an inequality. *Fix:* demand the pair, the commutator and the number. **Then note that the useful general statement being reached for exists, is classical, is computable, and is Section 5.**

FM-4.4 — THE ENERGY–TIME SLIDE. *Fault:* treating $\Delta E\, \Delta t \geq \hbar/2$ as an instance of Robertson's theorem. *Signature:* a σ_t computed from a wavefunction; there is no such object. *Fix:* name which of the three statements is meant and state Pauli's argument.

FM-4.5 — THE VACUOUS BOUND. *Fault:* quoting Robertson's general form as though it always constrains something. *Fix:* evaluate $\langle [A, B] \rangle$ in the relevant state; give the $|{+}x\rangle$ case where the bound is zero.

FM-4.6 — THE MACROSCOPIC IMPORT. *Fault:* applying the relation to a person, a decision or a room without computing the magnitude. *Fix:* compute. A gram at a micron gets 5.27×10^{-26} m/s.

FM-4.7 — BIOGRAPHY AS EVIDENCE. *Fault:* citing Bohr's emblem, or Pauli's interest in Jung, as support for a correspondence. *Fix:* report biography as biography, then run the concordance with C2 evaluated honestly.

FM-4.8 — PRECEDENT AS PROOF. *Fault:* presenting the precedent column as though it corroborated the evidence column. *Signature:* "ancient wisdom knew this, and now science confirms it". *Fix:* state that the older corpus recorded a structure and recorded no number, and that removing it changes nothing [Curzi, 2026w].

FM-4.9 — COLLAPSE BY MIND. *Fault:* moving from "measurement is special" to "consciousness collapses the wavefunction" and thence to claims about belief. *Fix:* the measurement problem is open; the consciousness-causes-collapse reading is a minority historical interpretation, and decoherence timescales for anything neural give the proposal a quantitative problem rather than merely a philosophical one [Curzi, 2026aa, Zeh, 1970].

FM-4.10 — FOURIER AMNESIA. *Fault:* treating the relation as exotic and quantum-specific, and therefore failing to apply the classical version where it bites. *Signature:* a design that renders tones too short to carry the pitch precision the design requires, unnoticed because "uncertainty principle" sounded like something that happens to electrons. **This is the expensive failure mode in this framework and Section 5 is the whole fix.**

32.5 5 Level two: the Gabor limit, which is the constraint that binds

This is the only uncertainty relation this framework owns, and the only one that constrains anything the reader will build. Everything the framework proposes to render travels on a wave, and the time–frequency plane has a floor:

```
dt . df  >=  1 / (4 pi)  =  0.079577471545948
```

This is classical. There is no $\hbar$ in it, no quantisation, no operators, no observers, no measurement problem and no interpretation. It is a theorem about square-integrable functions and their Fourier transforms, stated in its communications form by Gabor [1946], and it applies to every wave including the sound from a loudspeaker and the tone that encodes a lattice index. Section 4.5 showed that Kennard's relation is this theorem with the substitution $p = \hbar k$ performed on the last line. **The quantum case is the special case; this is the general one.**

The sentence to carry out of this section: **a short sound cannot have a precise pitch — not because instruments are poor, but because a short sound does not have one.**

32.5.1 5.1 The relation, stated exactly

```
PROVEN THEOREM of Fourier analysis. Classical. No quantum content.

Let g(t) be a square-integrable complex function of time with Fourier transform

    G(f) = INTEGRAL g(t) exp(-i 2 pi f t) dt

Define the spreads as standard deviations of the normalised energy densities:
```

```
    <t>     = INTEGRAL t |g(t)|^2 dt / INTEGRAL |g(t)|^2 dt
    dt^2    = INTEGRAL (t - <t>)^2 |g(t)|^2 dt / INTEGRAL |g(t)|^2 dt
    <f>     = INTEGRAL f |G(f)|^2 df / INTEGRAL |G(f)|^2 df
    df^2    = INTEGRAL (f - <f>)^2 |G(f)|^2 df / INTEGRAL |G(f)|^2 df

Then

    dt . df  >=  1 / (4 pi)

with equality if and only if g is a Gaussian (times a linear phase).
```

The constants, to full working precision:

```
1 / (4 pi)  =  0.079577471545948
4 pi        = 12.566370614359172
1 / (2 pi)  =  0.159154943091895      <-- the angular-frequency form: dt . d_omega >= 1/2
```

Read the units. Δt is in seconds, Δf in hertz — cycles per second, not radians per second — and the product is dimensionless. **The 4π rather than 2π is entirely a consequence of measuring frequency in cycles, and getting this wrong by a factor of two is the single most common arithmetic error made with this relation.** A $1/2$ on the right-hand side means angular frequency on the left.

Gabor's original paper introduced the time–frequency *information diagram*, showed it is partitioned into cells of irreducible area, named the elementary quantum of information the *logon*, and identified the Gaussian-modulated sinusoid as the signal achieving the minimum cell area [Gabor, 1946]. **The framing is the right one here: the time–frequency plane is not a continuum that can be subdivided indefinitely. It has a grain.** The result sits alongside the earlier bandwidth–information results of Nyquist [1928] and Hartley [1928] and is the geometric statement those two make arithmetically.

32.5.2 5.2 What Δt and Δf mean, and what they do not

Δt is a root-mean-square spread of energy in time, not a window length. A tone burst gated for exactly 100 ms does not have $\Delta t = 100$ ms; under a rectangular gate it has $\Delta t = 28.9$ ms, because the standard deviation of a uniform distribution of width T is $T/\sqrt{12} = 0.2887\,T$.

Δf is a root-mean-square spread of energy in frequency. It is not the width of a peak, not the bin spacing of a discrete Fourier transform, and not the resolution of a spectrum analyser. It is a second moment.

This matters immediately and in an unexpected direction: for some very common window shapes, Δf does not exist. A rectangular gate has a sinc spectrum whose magnitude falls as $1/f$, so $|G(f)|^2$ falls as $1/f^2$ and the integral defining the second moment diverges logarithmically. **That exponent is an exact analytic property of the function and is not an estimated exponent** (Section 3.7). A 200 ms rectangular gate, transformed at three sampling resolutions:

samples	dt spacing	computed df	computed dt . df
65,536	0.4 ms	29.60 Hz	1.712
262,144	0.2 ms	41.88 Hz	2.420
1,048,576	0.1 ms	59.24 Hz	3.422

The number keeps climbing as the bandwidth of the calculation is extended. It does not converge. The rectangular window has no finite Δf. This is not a numerical artefact; it is the correct answer, and it is a design fact of the first importance: **abrupt gating of a tone does not merely smear its spectrum, it smears it without bound.** Any rendering specification in this framework must therefore name an envelope, and the envelope must not be rectangular.

The standard envelopes, measured against the floor by the same numerical procedure (262,144 samples at 0.2 ms):

envelope	dt	df	dt . df	multiple of floor
Gaussian, sd = 50 ms	0.05000 s	1.5915 Hz	0.079577	1.000 EXACT
Hann, 200 ms	0.02829 s	2.8868 Hz	0.081665	1.026
triangular, 200 ms	0.03162 s	2.7579 Hz	0.087211	1.096
rectangular, 200 ms	0.05779 s	(divergent)	(divergent)	UNBOUNDED
one-sided exponential, 50 ms	0.02500 s	(divergent)	(divergent)	UNBOUNDED

The Gaussian saturates the bound to ten significant figures. The Hann window costs 2.6 per cent over the theoretical floor and has a finite answer. The rectangular gate has no answer at all. Those three lines are the entire practical content of "choose an envelope".

32.5.3 5.3 The derivation, in full

The proof is Section 4.5 with one line deleted, and it is given here in its own right because this is where it is used.

Let $g(t)$ be normalised so that $\int |g|^2\,dt = 1$, and shift so that $\langle t\rangle = \langle f\rangle = 0$. Work with angular frequency $\omega = 2\pi f$ and convert at the end.

Step 1 — build something that cannot be negative. For any real a:

```
I(a) = INTEGRAL | a t g(t) + g'(t) |^2 dt   >=   0
```

It is the integral of a squared modulus. **There is no physics in this step and no assumption about what g represents** — it could be a pressure wave, a voltage, a radio envelope or a probability amplitude. **This is what makes the application to an acoustic signal a derivation rather than an analogy: the premises are square-integrability and finite second moments, and a rendered tone of finite energy under a smooth envelope satisfies both.**

Step 2 — expand the square.

```
INTEGRAL |t g|^2 dt                    = dt^2

INTEGRAL |g'|^2 dt                     = d_omega^2
    -- by Plancherel's theorem, since differentiation in time is
       multiplication by i*omega in frequency

2a * Re INTEGRAL t g* g' dt            = -a
    -- integrate by parts:  INTEGRAL t (d/dt)(|g|^2/2) dt = -1/2
       using normalisation and decay of g at infinity

so    I(a) = a^2 dt^2 - a + d_omega^2  >=  0     for every real a
```

Step 3 — a real quadratic that is never negative has no real roots, so its discriminant is not positive.

```
1 - 4 dt^2 d_omega^2  <=  0
=>  dt . d_omega  >=  1/2
=>  dt . df       >=  1 / (4 pi)          since omega = 2 pi f
```

Step 4 — the equality condition, which is the design rule. Equality requires $I(a) = 0$ at the minimising $a = 1/(2\,\Delta t^2)$, and since $I(a)$ integrates a non-negative integrand, the integrand vanishes identically:

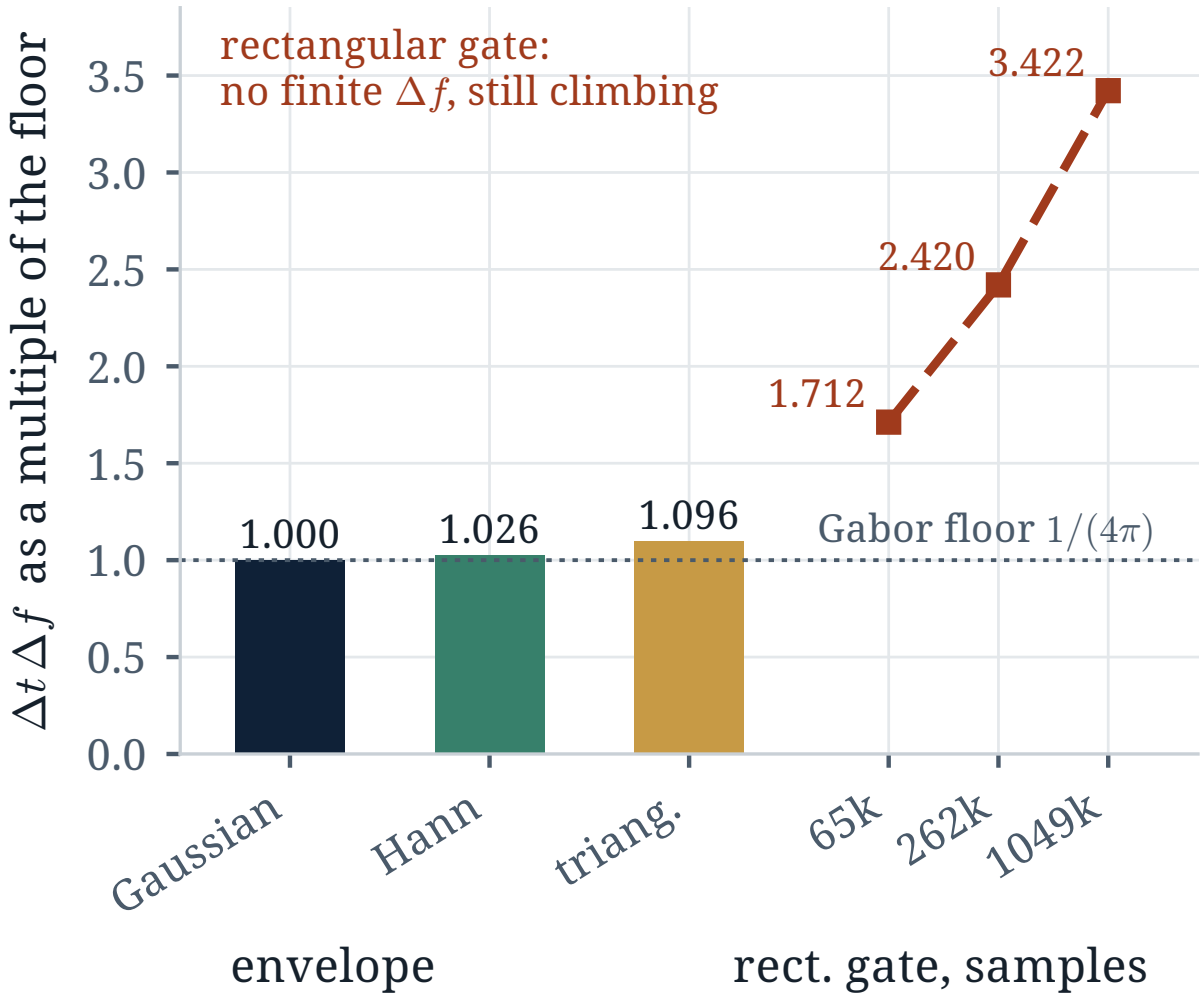


Figure 32.1: **The Gaussian saturates the Gabor floor, the Hann window costs 2.6 per cent, and the rectangular gate has no finite answer at all:** the four standard envelopes measured against $1/(4\pi) = 0.079577471545948$, with the rectangular gate's computed $\Delta t\,\Delta f$ shown at three successive analysis bandwidths, still climbing. The first three values are numerical evaluations of the stated envelopes on 262,144 samples at 0.2 ms spacing; the rectangular values are the 65,536 / 262,144 / 1,048,576-sample computations of Section 5.2 and are a divergence, not a measurement. Generated by `figs/src/CCP-032_fig01.py`.

```
a t g(t) + g'(t) = 0
=>  g(t) = C exp(-a t^2 / 2)              a GAUSSIAN
```

Only the Gaussian reaches the floor; everything else pays. Numerical confirmation at five widths:

```
Gaussian sd = 5 ms     ->  dt . df = 0.0795774715
Gaussian sd = 10 ms    ->  dt . df = 0.0795774715
Gaussian sd = 20 ms    ->  dt . df = 0.0795774715
Gaussian sd = 50 ms    ->  dt . df = 0.0795774715
Gaussian sd = 100 ms   ->  dt . df = 0.0795774715
1 / (4 pi)             =            0.0795774715
```

Scale-invariant, exact, and verified to ten significant figures. Squeeze the Gaussian in time and its transform widens by precisely the reciprocal factor. **That reciprocity is the law**, and it is the invariance that makes the octave form of Section 5.9 possible.

32.5.4 5.4 Warrant

PROVED, pure mathematics with no empirical content. $\Delta t\,\Delta f \geq 1/(4\pi)$ for every square-integrable g with finite second moments in both domains; equality iff g is Gaussian; and the scale invariance — replacing $g(t)$ by $g(t/s)$ multiplies Δt by s and divides Δf by s, leaving the product unchanged. **This is why the relation transposes cleanly through octaves**, which matters from Section 5.9 onward.

CITED. Gabor [1946] for the information diagram, the logon, the elementary signal and the $2BT$ degrees-of-freedom count. Frequency difference limens for pure tones: Wier, Jesteadt and Green, *Journal of the Acoustical Society of America* 61 (1977), 178–184 — **named without a key, per the disclosure in Section 4.3**; the reviewed literature in this corpus's bibliography is Burns [1999] and Demany et al. [2013]. Standard room acoustics — modal half-power bandwidth $B = 2.2/T_{60}$, Schroeder frequency $f_s = 2000\sqrt{T_{60}/V}$ [Sabine, 1922, Schroeder, 1965, Beranek, 2004].

POSTULATED about the mathematics: nothing. The Gabor limit is a theorem. **What is this framework's own is the *application* of the theorem to the essence law, and that application is arithmetic on top of a postulate — the essence law itself [Curzi, 2026] — rather than a new postulate.** The essence law's honesty clauses carry through unchanged.

32.5.5 5.5 The three resolution figures, and the factors between them

Three different numbers are all called "frequency resolution", they differ by more than an order of magnitude, and mixing them is the most common source of wrong design numbers in signal work.

```
(a) GABOR FLOOR            df = 1 / (4 pi T)         T interpreted as the window duration
                           The absolute theoretical minimum. Reached only by a Gaussian
                           whose rms width equals T -- an idealisation, since a Gaussian
                           has no finite duration. Treat as the LOWER BOUND that no
                           processing can beat.

(b) RAYLEIGH / DFT BIN     df = 1 / T
                           The spacing of DFT bins for a T-second record, and the classic
                           criterion for resolving two sinusoids with a rectangular window.
                           Exactly 4 pi = 12.566 times the Gabor floor.

(c) HANN MAIN LOBE         df = 4 / T
```

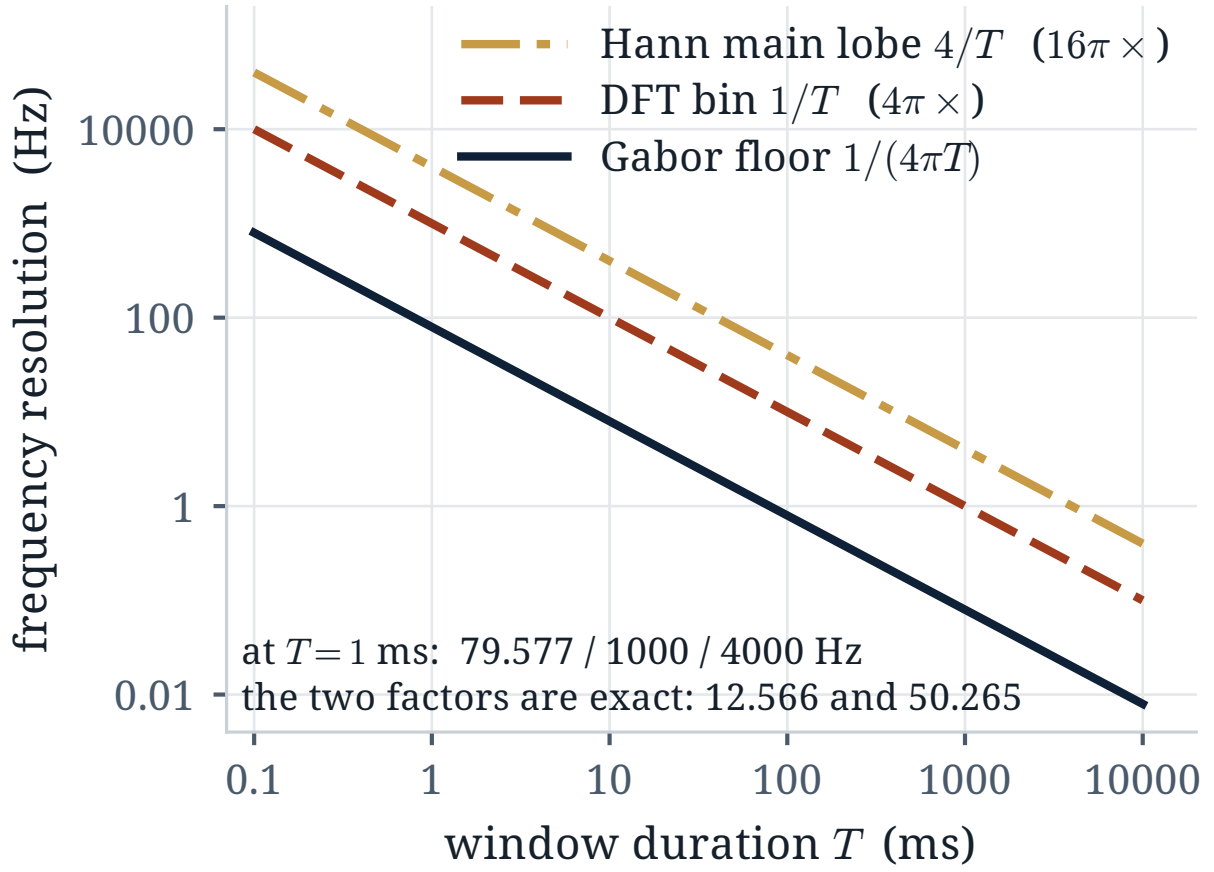


Figure 32.2: **Three quantities are called "frequency resolution" and they differ by up to a factor of 50:** the Gabor floor $1/(4\pi T)$, the DFT bin $1/T$ at exactly $4\pi = 12.566$ times the floor, and the Hann main lobe $4/T$ at exactly $16\pi = 50.265$ times it, plotted over the window durations of Table 5-A. Every value is an exact evaluation of the stated formula, not a measurement; the three multipliers are exact. Generated by `figs/src/CCP-032_fig02.py`.

```
The -3 dB-to--3 dB span of a Hann-windowed sinusoid's main
lobe, which is what a real analyser delivers when the window
is chosen for low spectral leakage.
Exactly 16 pi = 50.265 times the Gabor floor.
```

Throughout this section the tabulated $\Delta f_{\min}$ is the Gabor floor (a), because it is the bound no technique can beat; practical designs must be checked against (b) or (c) depending on the analysis chain. **The convention $\Delta t = T$ used in the tables is the standard engineering shorthand, stated here in the open: it treats the whole rendering window as the rms spread, which is optimistic by construction.** It is retained because the framework's existing design table was computed in it, and cross-paper consistency is worth more than a factor of $\sqrt{12}$. **A reader converting to true rms spreads must apply the shape factor**, and the resulting figures will be worse, not better, than those tabulated.

32.5.6 5.6 The concordance, with column roles stated

The evidence column carries evidence; the precedent column carries precedent [Curzi, 2026w].

```
CLAIM: The precision with which a frequency can be determined is limited by the duration
```

```
       over which it is observed, and the two trade reciprocally.

(+1) EVIDENCE COLUMN
     Gabor, Theory of Communication (1946): dt df >= 1/(4 pi), proved, with the equality
     case identified. Verified numerically in Section 5.3 to ten significant figures.
     Measurable directly: gate a 1 kHz sinusoid for 1 ms, transform it, observe a spectral
     spread of order 100 Hz on any oscilloscope with an FFT function. A single reproducible
     signal with dt df < 1/(4 pi) and finite second moments in both domains would refute
     it, and would also refute Plancherel's theorem.

(-1) PRECEDENT COLUMN -- CARRIES NO EVIDENCE
     THE TUNER'S BEAT. Keyboard and organ tuning by counting beats is documented craft
     practice from the seventeenth century and is still taught for aural tuning. Two tones
     a frequency df apart produce an amplitude beat at rate df; counting one beat requires
     listening for 1/df seconds. The craft rule -- that resolving a small pitch difference
     requires a proportionally long listen, and that a unison cannot be tuned from a short
     pluck -- is the same reciprocal trade, in working use for roughly three hundred years
     before Gabor wrote it down.

     The arithmetic is exact and worth stating: counting ONE beat cycle at df takes
     1/df seconds, which is 4 pi = 12.566 times the Gabor floor for that df. A tuner who
     counts eight beats before judging -- ordinary practice -- has listened for 8/df
     seconds, which is 100.5 times the floor. The craft is not near the theoretical limit
     and never claimed to be.

(0p) HOW THEY INTEGRATE
     The craft route produced a WORKING CONSTANT (listen for several beats) and no
     derivation, no bound, and no account of why. The analytic route produced the bound,
     the equality case, and the envelope design rule. What the concordance establishes is
     that the reciprocal structure was findable without Fourier analysis. What it does NOT
     establish is any number: remove the (-1) column and every computation here is
     unchanged.

BRIDGE CRITERIA:
     C1 structural correspondence  PASS. Both state that resolving df requires a listen
                                   time proportional to 1/df.
     C2 predictive power           PASS, unusually. The tuner's rule DOES generate a
                                   prediction -- "this unison cannot be judged in a tenth
                                   of a second" -- which is correct and testable by anyone
                                   with two strings.
     C3 practical utility          PASS on both sides.

STATUS: ESTABLISHED as a structural concordance, and included partly to demonstrate that
the criteria are satisfiable -- the corpus rejects roughly one bridge in five in print
(CCP-029).
```

A worked instance of the craft side, computed. Equal temperament's major third above A4 = 440 Hz is 554.3653 Hz against just intonation's 550.0000 Hz; a tuner listens to the beat between the fourth harmonic of the tempered third (2217.4610 Hz) and the fifth harmonic of A4 (2200.0000 Hz), a beat rate of 17.4610 Hz — one beat every 57.3 ms, against a Gabor floor of 4.5574 ms. **The craft sits comfortably above the bound, which is why the craft works.** The tuning ratios themselves are standard [Barbour, 1951, von Helmholtz, 1885].

32.5.7 5.7 Table 5-A — the master window/resolution table

All values computed; $\Delta f_{\min}$ is the Gabor floor $1/(4\pi T)$, with the two practical criteria alongside.

window T	df_min = 1/(4 pi T)	1/T (DFT bin)	4/T (Hann lobe)	cycles at 1 kHz
0.1 ms	795.77472 Hz	10000.0000	40000.0000	0.10
0.2 ms	397.88736 Hz	5000.0000	20000.0000	0.20
0.5 ms	159.15494 Hz	2000.0000	8000.0000	0.50
1 ms	79.57747 Hz	1000.0000	4000.0000	1.00
2 ms	39.78874 Hz	500.0000	2000.0000	2.00
5 ms	15.91549 Hz	200.0000	800.0000	5.00
10 ms	7.95775 Hz	100.0000	400.0000	10.00
20 ms	3.97887 Hz	50.0000	200.0000	20.00
23.22 ms	3.42711 Hz	43.0663	172.2653	23.22
50 ms	1.59155 Hz	20.0000	80.0000	50.00
100 ms	0.79577 Hz	10.0000	40.0000	100.00
200 ms	0.39789 Hz	5.0000	20.0000	200.00
500 ms	0.15915 Hz	2.0000	8.0000	500.00
1000 ms	0.07958 Hz	1.0000	4.0000	1000.00
10000 ms	0.00796 Hz	0.1000	0.4000	10000.00

The four rows the framework's design table names read out of this exactly: 1 ms gives 79.577 Hz, 10 ms gives 7.958 Hz, 100 ms gives 0.796 Hz, and 1000 ms gives 0.080 Hz.

The 23.22 ms row is included because it is the duration of a 1024-sample frame at 44.1 kHz, the most common analysis frame in audio work. **A 1024-point FFT at CD rate cannot resolve better than 3.43 Hz no matter what is done to it, and a Hann-windowed analysis of that frame delivers about 172 Hz of main-lobe width** — a factor of fifty between the theoretical floor and the practical number, and designs are broken by that factor routinely.

The inverse table:

target df	T at Gabor floor	T at 1/df (DFT bin)
1000 Hz	0.0796 ms	1.0000 ms
100 Hz	0.7958 ms	10.0000 ms
10 Hz	7.9577 ms	100.0000 ms
3 Hz	26.5258 ms	333.3333 ms
1 Hz	79.5775 ms	1000.0000 ms
0.1 Hz	795.7747 ms	10000.0000 ms

32.5.8 5.8 The octave-invariant form: count cycles, not milliseconds

Table 5-A is in absolute time, which is awkward for music and useless for a lattice that gets transposed. **There is a form of the relation invariant under octave transposition, and it is the form used everywhere below**, because the framework transposes everything. Let $r = \Delta f / f_0$ be the fractional resolution required at centre frequency f_0. Then:

```
T . f0  >=  1 / (4 pi r)

LEFT SIDE  = the number of CYCLES of the tone contained in the window.
RIGHT SIDE = depends only on the fractional resolution.
```

Transposing by an octave doubles f_0 and halves the required T, leaving the cycle count unchanged. The cycle count is the invariance — the transformation group is dilation of the time axis, the invariant is $T f_0$, and Section 5.3's scale-invariance calculation is the proof. Converted to cents, since musical resolution is specified in cents:

resolution	r = df / f0	minimum cycles in the window	
1200 cents	1.00000000	0.0796	
700 cents	0.49830708	0.1597	
200 cents	0.12246205	0.6498	
100 cents	0.05946309	1.3383	
50 cents	0.02930224	2.7157	
34.976 cents	0.02040815	3.8993	<– the 50/49 interval
10 cents	0.00579294	13.7370	
6.001 cents	0.00347221	22.9184	<– the 289/288 interval
5 cents	0.00289229	27.5137	
1 cent	0.00057779	137.7274	

One cent of pitch resolution requires at least 138 cycles of tone, at any pitch, under the most favourable possible envelope — 313 ms at A4 = 440 Hz, 66 ms at C7 = 2093 Hz. **The absolute time changes with register; the cycle count never does.** In working form: one cent is about 138 cycles, five cents about 28, ten cents about 14, and a semitone about 1.3.

32.5.9 5.9 A short sound cannot have a precise pitch

As physics. Table 5-A converted to cents at eight reference pitches; each entry is the minimum spectral spread, expressed as a musical interval, for a tone of that duration at that pitch.

pitch	1 ms	5 ms	10 ms	25 ms	50 ms	100 ms	250 ms	1000 ms
C1 32.7	2135.53	686.50	377.05	160.80	82.27	41.62	16.77	4.21
C2 65.4	1378.07	377.05	198.77	82.27	41.62	20.94	8.40	2.11
C3 130.8	822.68	198.77	102.24	41.62	20.94	10.50	4.21	1.05
A3 220.0	534.51	120.92	61.52	24.87	12.48	6.25	2.50	0.63
A4 440.0	287.80	61.52	31.03	12.48	6.25	3.13	1.25	0.31
A5 880.0	149.87	31.03	15.59	6.25	3.13	1.56	0.63	0.16
A6 1760.0	76.56	15.59	7.81	3.13	1.56	0.78	0.31	0.08
A7 3520.0	38.70	7.81	3.91	1.56	0.78	0.39	0.16	0.04

(values in cents)

A 1 ms tone at C1 has a spectral spread of 2135 cents — nearly two octaves. It does not have a pitch; it is a click with a slight colour. A 1 ms tone at A4 spreads 288 cents, close to a minor third: enough to say “somewhere around A” and nothing more. A 1 ms tone at A7 is down to 39 cents, under a quarter-tone, and there the transient does have a recognisable pitch. **That register dependence is why bass notes need to be long and cymbals do not.**

As perception. The physics gives a floor; hearing sits above it. Wier, Jesteadt and Green (*JASA* 61, 1977, 178–184) measured frequency difference limens for pure tones at roughly 0.2 to 0.3 per cent near 1 kHz, widening at low frequencies and low sensation levels; the reviewed treatment in this corpus’s bibliography is Burns [1999]. Setting Δf equal to the limen and inverting:

f0	difference limen	df	Gabor T_min	cycles	as cents
125 Hz	0.6 %	0.750 Hz	106.10 ms	13.26	10.36
250 Hz	0.4 %	1.000 Hz	79.58 ms	19.89	6.91
500 Hz	0.3 %	1.500 Hz	53.05 ms	26.53	5.19
1000 Hz	0.3 %	3.000 Hz	26.53 ms	26.53	5.19
1000 Hz	0.2 %	2.000 Hz	39.79 ms	39.79	3.46
4000 Hz	0.2 %	8.000 Hz	9.95 ms	39.79	3.46

A tone must last on the order of 25 to 100 ms before its physical spectral spread is narrower than the smallest pitch difference a human can hear. Shorter than that, the limiting factor is not the ear but the signal: **the tone genuinely does not contain the information the listener is being asked for.** The difference limens are measured psychophysical quantities carrying their own sampling uncertainty, which is not restated here; the Gabor columns beside them are exact.

This is the honest converse of a claim often made carelessly in sound-design work — that some effect requires "exposure" to a frequency. Whatever else is true of that, the Gabor limit says something narrower and certain: **if the frequency is not present long enough, it is not present as a frequency at all, and no amount of repetition of too-short bursts fixes it**, because each burst independently lacks the spectral definition. Repetition raises the count of events; it does not sharpen any one of them.

32.5.10 5.10 The lattice under a common transposition

The essence law, presupposed as a badged postulate with k a chosen reference constant and not a measured one [Curzi, 2026, Markowsky, 1992]:

```
f(n) = k n^2 ,      k = (9/8)^2 / phi^2 = (81/64)/(phi+1) = 0.483425732988414 Hz
```

k **is a DEFINED constant**, exact as a ratio of exact quantities, which is why fifteen significant figures are legitimate for it; **the attachment of hertz to it is a separate and much stronger claim and is POSTULATED** (Section 9.2, P6). Two rendering strategies exist with completely different consequences.

STRATEGY A — COMMON TRANSPOSITION. One integer j for the whole lattice; every index renders at $f(n) = kn^2 2^j$. Ratios are preserved exactly and the set spans many octaves. The gap between adjacent indices is gap$(n \to n+1) = k(2n+1)2^j$, and the fractional gap, which by Section 5.8 determines the cycle count and is invariant under j, is $r(n) = (2n+1)/n^2$:

n	f_raw (Hz)	gap to n+1	r = gap/f	min cycles
1	0.48343	1.45028	3.000000	0.027
2	1.93370	2.41713	1.250000	0.064
4	7.73481	4.35083	0.562500	0.141
8	30.93925	8.21824	0.265625	0.300
12	69.61331	12.08564	0.173611	0.458
16	123.75699	15.95305	0.128906	0.617
24	278.45322	23.68786	0.085069	0.935
32	495.02795	31.42267	0.063477	1.254
64	1980.11180	62.36192	0.031494	2.527
92	4091.71540	89.43376	0.021857	3.641

The fractional gap shrinks like $2/n$, so the cycle count grows like $n/(8\pi)$ — an asymptotic of the exact expression, not a fitted exponent (Section 3.7). Verified against $n^2/(4\pi(2n+1))$: at $n = 10$, 0.379 exact against 0.398 approximate; at $n = 100$, 3.959 against 3.979; at $n = 500$, 19.874 against 19.894.

> **DESIGN LAW 1, exact.** Under a common transposition, distinguishing index n from $n+1$ requires a tone lasting at least $n^2/(4\pi(2n+1))$ cycles, approaching $n/(8\pi)$ for large n, **and this is independent of the transposition j. The deep lattice is intrinsically slower to render than the shallow one, and no choice of octave rescues it.**

THE AUDIBLE-BAND CONSTRAINT. A common transposition must place the whole lattice inside human hearing, taken as 20 Hz to 20 kHz:

```
lower edge:   k * 1^2 * 2^j >= 20      =>  2^j >= 41.371  =>  j >= 5.371  =>  j >= 6
upper edge:   k * N^2 * 2^j <= 20000
```

For real-valued j the ceiling on lattice size is $N^2 \leq 20000/20 = 1000$, i.e. $N \leq 31.62$. Restricting to integer octave transposition, as the framework requires, the constraint tightens:

j	f(1)	maximum N	f(N) at that N
5	15.4696 Hz	(below 20 Hz -- DISALLOWED)	
6	30.9392 Hz	25	19337.03 Hz
7	61.8785 Hz	17	17882.88 Hz
8	123.7570 Hz	12	17821.01 Hz
9	247.5140 Hz	8	15840.89 Hz

DESIGN LAW 2. Under a single integer octave transposition, the square law admits **at most 25 lattice elements inside the audible band, and only at** $j = 6$. A 26-element lattice at $j = 6$ puts index 26 at 20911 Hz, outside hearing. A 12-element lattice has three admissible transpositions, $j = 6$, 7 or 8.

This is a hard structural limit on the method, it is arithmetic rather than taste, and it was not visible until the question of what rendering actually costs was forced.

THE RENDER-RATE CEILING. Under a common transposition the binding constraint is at the *bottom* of the lattice, where the absolute gap is smallest. At $n = 1$ the gap is $3k2^j$:

j	gap(1 -> 2)	T_min (Gabor floor)	R_max (symbols/s)
6	92.81774 Hz	0.85735 ms	1166.38
7	185.63548 Hz	0.42868 ms	2332.76
8	371.27096 Hz	0.21434 ms	4665.53

Under the stricter half-gap criterion, requiring the spread to be at most half the gap so that two adjacent elements are *separated* rather than merely non-identical:

```
j = 6, n = 1:   gap = 92.81774 Hz,  df <= 46.4089 Hz
                T_min = 1.7147 ms,  R_max = 583.19 symbols/s
```

A common-transposition lattice at $j = 6$ can be rendered at up to about 583 elements per second before spectral smearing merges adjacent positions.

32.5.11 5.11 The folded lattice, where the constraint becomes severe

STRATEGY B — PER-VALUE OCTAVE FOLDING. Fold each value by its own power of two into a chosen band, stating the band and j for every value. This keeps the lattice inside one octave and makes it musically usable [Shepard, 1982]. **It also destroys the guaranteed spacing, and that is the finding of this subsection.**

Two indices share a pitch class if and only if they have the same odd part, so an N-element lattice folds to $\lceil N/2 \rceil$ distinct classes [Curzi, 2026]. For $N = 12$ the odd parts are $\{1, 3, 5, 7, 9, 11\}$ and their squares folded into one octave are:

index odd part	m^2	folded ratio	cents above base
1	1	1/1	0.000
3	9	9/8	203.910
9	81	81/64	407.820
7	49	49/32	737.652
5	25	25/16	772.627
11	121	121/64	1102.636

The gaps between adjacent members, in cents, the last wrapping to the octave:

```
203.910   203.910   329.832    34.976   330.008    97.364
```

The minimum gap is 34.976 cents, between index 7 at $49/32$ and index 5 at $25/16$. The interval is exactly $50/49$ — the septimal sixth-tone, which falls out of the square law with no choice in the matter, since $(25/16)/(49/32) = 50/49$. Extending to larger lattices, computed on exact rationals:

N	pitch classes	minimum gap ratio	cents	which indices	
8	4	50/49	34.9756	7 and 5	
12	6	50/49	34.9756	7 and 5	
16	8	50/49	34.9756	7 and 5	
17	9	289/288	6.0008	3 and 17	<- CLIFF
20	10	289/288	6.0008	3 and 17	
24	12	289/288	6.0008	3 and 17	
32	16	289/288	6.0008	3 and 17	
40	20	289/288	6.0008	3 and 17	
41	21	1682/1681	1.0296	41 and 29	<- CLIFF
48	24	1682/1681	1.0296	41 and 29	
64	32	1682/1681	1.0296	41 and 29	

> **CORRECTION TO A PREVIOUSLY PUBLISHED RESULT.** This analysis originally placed the first cliff at $N = 20$ and the second at $N = 48$. **Both boundaries were stated too late.** On exact rational arithmetic — every n^2 for $n = 1 \dots N$ folded into one octave with exact fractions, minimum adjacent ratio taken including the wrap to the octave — the boundaries are $N = 17$ and $N = 41$. The cause given originally was correct and is unchanged: $289/288$ is the culprit, and it appears the moment index 17 joins the set, because $17^2 = 289$ folds beside $288 = 2 \cdot 12^2$, and index 12 has been present since $N = 12$. **20 and 48 were merely the next sampled rows of the table, not the thresholds. The rate figures (86.394 falling to 10.799, a factor of eight) were and remain correct**; only the N at which they take effect has moved.

The table above samples N; the exact regime boundaries are $N = 3$, 7, 17 and 41, with the schedule $9/8$ (203.9100 cents) for $N = 3 \dots 6$, $50/49$ (34.9756) for $7 \dots 16$, $289/288$ (6.0008) for $17 \dots 40$ and $1682/1681$ (1.0296) for $N \geq 41$ [Curzi, 2026]. **Confusing a sampled row with a threshold is precisely the error corrected below**, so the distinction is made here before the correction is stated.

> **DESIGN LAW 3.** The minimum gap in a folded lattice collapses in discrete jumps as N crosses the thresholds at which a new near-coincidence of squares appears — **35 cents up to $N = 16$, 6 cents from $N = 17$, 1 cent from $N = 41$.** These are not smooth; they are number theory. $289/288$ arises because $17^2 = 289$ folds to $289/256$ and lands six cents above $9/8 = 288/256$; $1682/1681$ arises because $41^2 = 1681$ and $2 \cdot 29^2 = 1682$.

Converted to render times at a 220 Hz base:

N	classes	min gap (cents)	min gap (Hz)	T_min (ms)	R_max (symbols/s)	
4	2	203.9100	27.50000	2.8937	345.575	
8	4	34.9756	6.87500	11.5749	86.394	
12	6	34.9756	6.87500	11.5749	86.394	
16	8	34.9756	6.87500	11.5749	86.394	
17	9	6.0008	0.85938	92.5992	10.799	<- CLIFF

20	10	6.0008	0.85938	92.5992	10.799	
24	12	6.0008	0.85938	92.5992	10.799	
32	16	6.0008	0.85938	92.5992	10.799	
40	20	6.0008	0.85938	92.5992	10.799	
41	21	1.0296	0.21484	370.3970	2.700	<- CLIFF
48	24	1.0296	0.21484	370.3970	2.700	
64	32	1.0296	0.21484	370.3970	2.700	

Read the $N = 16$ and $N = 17$ rows against each other. Adding a single element costs a factor of eight in rendering speed, from 86 elements per second to 10.8. That is not a gradual degradation. **It is a cliff, it sits at $N = 17$, and it is caused by 17^2 landing six cents away from $9/8$.**

The factor of eight is exact, and the obvious shortcut gets it wrong. The two binding pairs do not sit at the same place in the octave, so the ratio of *rate ceilings* is not the ratio of *fractional* gaps ($288/49 = 5.878$). Gabor resolution is linear in hertz and $R_{\max} = 4\pi\,\Delta f$ uses the absolute gap: the $50/49$ pair straddles $49/32$, so its gap is base $\cdot\,(50 - 49)/32 =$ base$/32$; the $289/288$ pair straddles $9/8 = 288/256$, so its gap is base $\cdot\,(289 - 288)/256 =$ base$/256$. Then $(1/32)/(1/256) = 8$ exactly, at every base pitch. The same argument gives exactly 4 across the second cliff.

The base pitch scales everything, since the constraint is fractional:

base	tightest pair (indices 7 and 5)	gap (Hz)	T_min (ms)	R_max (/s)
55.0000	84.2188 / 85.9375 Hz	1.71875	46.2996	21.60
110.0000	168.4375 / 171.8750 Hz	3.43750	23.1498	43.20
220.0000	336.8750 / 343.7500 Hz	6.87500	11.5749	86.39
440.0000	673.7500 / 687.5000 Hz	13.75000	5.7875	172.79

Raise the base an octave and the achievable rate doubles. That is the octave-invariance of Section 5.8 seen from the other side, and it is the one free lever the designer has.

> **DESIGN LAW 4, the practical headline.** A folded lattice is between 5 and 55 times slower to render than a common-transposition lattice of the same size: $N = 12$ at $j = 6$ under the half-gap criterion gives **583.19 symbols/s**; $N = 12$ folded to one octave at 220 Hz gives **86.39 symbols/s** (6.8 times slower); $N = 24$ folded at 220 Hz gives **10.80 symbols/s** (54.0 times slower). **Folding is not free. It buys musical usability and it is paid for in bandwidth.**

32.5.12 5.12 Resolving a surplus difference

The framework's cross-disciplinary bridge sets the effective count equal to the lattice index, so that $f_{\text{cross}} = k\,N_{\text{eff}}^2 = k\,e^{2F}$ with F the interaction surplus in nats [Curzi, 2026]. **The effective count in that expression is the corpus's affine $N_{\text{eff}} = 1 + (N - 1)u$, which is not a Hill number** (Section 3.4); the badge travels with it here. One nat of surplus is 2.885 octaves:

F (nats)	N_eff	f = k e^(2F)	octaves above k
0.000000	1.0000	0.483426 Hz	0.0000
1.000000	2.7183	3.572060 Hz	2.8854
2.484907	12.0000	69.613306 Hz	7.1699
3.000000	20.0855	195.027860 Hz	8.6562

Asking the Gabor question of that bridge: if $f = ke^{2F}$ then $\Delta f/f = 2\,\Delta F$, and substituting into the octave-invariant form:

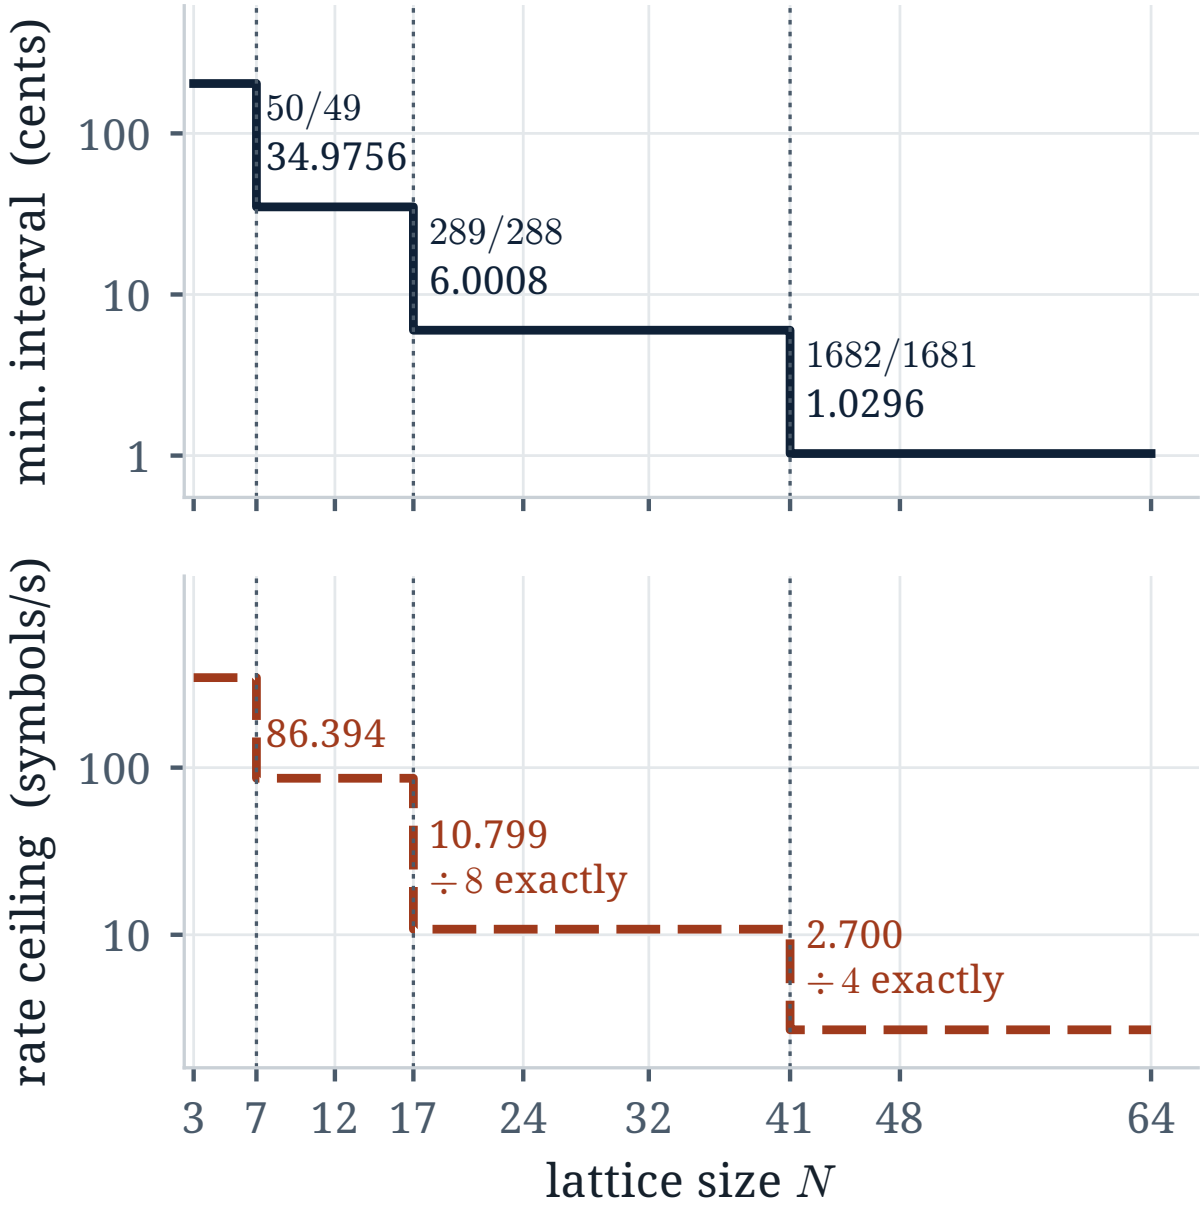


Figure 32.3: **The folded lattice's minimum gap collapses at $N = 17$ and again at $N = 41$, not at the previously published 20 and 48:** upper panel, the minimum adjacent interval in cents by lattice size on the exact-rational schedule — $9/8 = 203.9100$ cents for $N = 3 \dots 6$, $50/49 = 34.9756$ for $7 \dots 16$, $289/288 = 6.0008$ from 17, $1682/1681 = 1.0296$ from 41 [Curzi, 2026]. Lower panel, the render-rate ceiling those gaps imply at a 220 Hz base, falling by exactly 8 at the first cliff and exactly 4 at the second. **The regime boundaries are the exact thresholds, not the sampled table rows** — which is the error the correction above repairs. Every value is exact arithmetic on the square law, not a measurement. Generated by `figs/src/CCP-032_fig03.py`.

```
============================================================
  MINIMUM CYCLES TO RESOLVE A SURPLUS DIFFERENCE dF

      T . f  >=  1 / (8 pi dF)

  Independent of k, independent of the transposition, independent
  of the band.
============================================================
```

Tabulated, with absolute times at a 220 Hz rendering:

dF (nats)	df/f	as cents	min cycles	T_min at 220 Hz
0.5000	1.0000	1200.00	0.080	0.3617 ms
0.1000	0.2000	315.64	0.398	1.8086 ms
0.0500	0.1000	165.00	0.796	3.6172 ms
0.0200	0.0400	67.90	1.989	9.0429 ms
0.0100	0.0200	34.28	3.979	18.0858 ms
0.0050	0.0100	17.23	7.958	36.1716 ms
0.0010	0.0020	3.46	39.789	180.8579 ms

To distinguish two ensembles whose interaction surplus differs by one millinat requires at least 181 milliseconds of tone at A3; a hundredth of a nat takes 18 ms. A hundredth of a nat corresponds, at $N = 12$, to a difference in effective count of about 0.12 members — **so the rendering can distinguish group compositions to about a tenth of a person, given 18 milliseconds to do it in. That is a statement about what the signal contains and is not a claim about anyone's hearing**; whether a listener can extract it is governed by the difference limens of Section 5.9. The logarithmic register in which F is written is itself defensible on grounds that do not require a perceptual story: Peters and Gell-Mann [2016] derives a logarithmic evaluation from multiplicative dynamics with no psychological assumption at all, which is a stronger warrant than Fechner's [Fechner, 1860] for the shape used here.

32.5.13 5.13 The rate law

For a stream rendered back to back at R symbols per second, each symbol occupies $T = 1/R$ seconds, so:

```
========================================
   df_floor  =  R / (4 pi)

   Rendering rate and achievable frequency
   precision are DIRECTLY PROPORTIONAL.
========================================
```

R (symbols/s)	T = 1/R (ms)	df floor (Hz)
1	1000.0000	0.07958
2	500.0000	0.15915
5	200.0000	0.39789
10	100.0000	0.79577
20	50.0000	1.59155
50	20.0000	3.97887
100	10.0000	7.95775
200	5.0000	15.91549
500	2.0000	39.78874
1000	1.0000	79.57747

This table is the design contract. Pick a row. Either the rate or the precision may be had, not both, and no encoding, modulation scheme, processing chain or amount of computation moves the number, because the number is a theorem about the functions themselves.

32.5.14 5.14 Two earlier results in the corpus were Gabor problems

The room. The corpus's acoustic-exclusion analysis computed reverberation times, critical distances and exclusion arithmetic for a hard-surfaced room without naming why those quantities are coupled [Curzi, 2026e]. Standard room acoustics gives a mode's half-power bandwidth as $B = 2.2/T_{60}$ [Sabine, 1922, Beranek, 2004], which is a duration–bandwidth statement: **a mode that rings for a long time is sharp in frequency; a mode that dies quickly is broad.**

T60	modal bandwidth	T to resolve one mode	B ∗ T60	as multiple of 1/(4 pi)
0.30 s	7.3333 Hz	10.8515 ms	2.20	27.65
0.45 s	4.8889 Hz	16.2772 ms	2.20	27.65
0.60 s	3.6667 Hz	21.7029 ms	2.20	27.65
0.90 s	2.4444 Hz	32.5544 ms	2.20	27.65
2.50 s	0.8800 Hz	90.4289 ms	2.20	27.65

The product $BT_{60} = 2.2$ is constant by construction and sits 27.65 times above the Gabor floor — exactly what should happen, since a decaying room mode is an exponentially damped sinusoid and Section 5.2 established that the one-sided exponential is far from the floor. **The room is a Lorentzian resonator, not a Gaussian one, and Lorentzians are inefficient users of the time–frequency plane.** At $T_{60} = 0.45$ s modes are 4.89 Hz wide, and the Schroeder frequency for a 100 m^3 room is $2000\sqrt{0.45/100} = 134.16$ Hz (258.20 Hz for 30 m^3 at $T_{60} = 0.50$ s; 69.28 Hz for 1000 m^3 at $T_{60} = 1.20$ s) [Schroeder, 1965]. **The Schroeder frequency is itself a Gabor statement in disguise: it is the frequency at which modal spacing falls below modal bandwidth**, at which time-domain decay has smeared the frequency-domain structure into a continuum.

The channel. Gabor's 1946 paper contains a second result on which the corpus's channel analysis depends directly [Curzi, 2026x]: **the number of independent degrees of freedom in a signal of bandwidth B occupying duration T is approximately $2BT$.** Each degree of freedom is one cell of the information diagram — one logon.

B (Hz)	T (s)	logons = 2BT	SNR (dB)	C = B log2(1+S/N)	bits/logon
1,000	0.100	200	3	1,583 bit/s	0.7913
1,000	1.000	2,000	10	3,459 bit/s	1.7297
4,000	0.020	160	20	26,630 bit/s	3.3291
4,000	1.000	8,000	30	39,870 bit/s	4.9836
20,000	1.000	40,000	40	265,800 bit/s	6.6439

The two results compose exactly: Gabor counts the cells and Shannon prices them [Gabor, 1946, Shannon, 1948, Shannon and Weaver, 1949]. A channel's capacity is the number of time–frequency cells per second times the bits each can carry, and the cell count has a hard floor set by $1/(4\pi)$.

And the standing constraint bears repeating: there is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal [Curzi, 2026y]. The Gabor limit says how many distinguishable symbols could in principle be placed on the wave. **It says nothing about whether anything at the far end can read them.**

32.5.15 5.15 Requirements

REQ-5.1 State the limit as $\Delta t\,\Delta f \geq 1/(4\pi)$ with f in hertz, and state in the same passage that it is classical and contains no quantum content.

REQ-5.2 Never quote a frequency-resolution figure without naming which of the three criteria it is. The multipliers are $4\pi = 12.566$ and $16\pi = 50.265$.

REQ-5.3 Name the envelope for every rendered tone. **Rectangular gating is prohibited: it has no finite spectral second moment.** Gaussian is the reference (1.000 of the floor); Hann is the working default (1.026).

REQ-5.4 For lattice work, use the octave-invariant cycle form $T f_0 \geq 1/(4\pi r)$ rather than absolute milliseconds, and state the required resolution in cents.

REQ-5.5 For a common-transposition lattice, verify the audible-band constraint before choosing N. Integer j admits at most $N = 25$ and only at $j = 6$. Publish j, $f(1)$ and $f(N)$.

REQ-5.6 For a folded lattice, compute the folded set's own minimum gap. The raw spacing does not survive folding and the collapse is discontinuous in N, at $N = 17$ and $N = 41$.

REQ-5.7 Publish the render-rate ceiling with every lattice specification, computed as $R_{\max} = 4\pi \cdot (\text{minimum absolute gap in Hz})$, stating whether the full-gap or half-gap criterion was used.

REQ-5.8 Where a surplus difference is rendered acoustically, use $T f \geq 1/(8\pi\,\Delta F)$ and state the resolution in nats.

REQ-5.9 Restate at every use: the output is an essence frequency and not a physical resonance, and k is a chosen constant and not a measured one.

REQ-5.10 Restate, wherever the limit is used to size a biological channel, that the capacity calculation says what could be carried and nothing about whether a cell can demodulate sequence-specific content [Curzi, 2026y].

REQ-5.11 Every concordance states its column roles and evaluates C1, C2 and C3 explicitly, including failures.

REQ-5.12 Where the relation is invoked at any scale, compute the number. **A Gabor bound with no arithmetic attached is a mood, not a constraint.**

REQ-5.13 No text under this framework may present the Gabor limit as a quantum effect, as evidence for the essence law, or as licence for any claim about mind, belief or physiology.

32.5.16 5.16 Failure modes

FM-5.1 — THE FACTOR OF TWO PI. Writing $\Delta t\,\Delta f \geq 1/2$ with f in hertz. *Signature:* resolution figures wrong by exactly 6.283 or 12.566 — a 1 ms window credited with 12.7 Hz or 500 Hz instead of 79.577 Hz. *Fix:* check the units on the frequency variable first.

FM-5.2 — THE RECTANGULAR GATE. Rendering a tone by switching it on and off. *Signature:* spectral splatter across the band; a measured Δf that keeps rising as the analysis bandwidth is extended (29.60, 41.88, 59.24 Hz and still climbing); audible clicks at symbol boundaries. *Fix:* apply an envelope; Hann costs 2.6 per cent over the floor and converges.

FM-5.3 — WINDOW LENGTH MISTAKEN FOR SPREAD. *Signature:* claimed resolution better than achieved by about $\sqrt{12} = 3.46$ for a rectangular gate. *Fix:* state the convention explicitly, as Section 5.5 does.

FM-5.4 — THE OPTIMISTIC FLOOR. Designing to $1/(4\pi T)$ and implementing with a Hann-windowed FFT. *Signature:* a design that passes on paper by a comfortable margin and fails in the build by roughly fifty times. **The gap between the Gabor floor and a Hann main lobe is $16\pi = 50.265$, and that is the single largest source of optimism in signal design.**

FM-5.5 — THE FOLDING SURPRISE. Assuming a folded lattice inherits the raw spacing. *Fix:* compute the folded set's own minimum gap. For $N = 12$ it is $50/49$ — 34.976 cents, between indices 7 and 5 — and the rate ceiling drops from 583 to 86 symbols per second at 220 Hz.

FM-5.6 — THE $N = 17$ CLIFF. *(Originally titled "the $N = 20$ cliff"; the boundary is corrected in Section 5.11 and the rate figures are unchanged.)* Extending a working 16-element folded lattice by one element and keeping the rate. *Signature:* everything works at $N = 16$ and degrades sharply at 17, and because 17 is rarely a lattice size anyone tests, the collapse is usually first *noticed* at 20 or 24 and wrongly attributed to them. *Fix:* accept the slower rate, raise the base pitch, or stay at $N \leq 16$.

FM-5.7 — THE UNBOUNDED SPEED-UP. Believing faster rendering merely reduces aesthetic quality. *Signature:* "we can always slow it down later." *Fix:* you cannot. **The spectral definition is not present in the fast rendering and cannot be recovered from it.** At 500 symbols per second the floor is 39.8 Hz and any gap below that has been destroyed, permanently, in that recording.

FM-5.8 — THE REPETITION FALLACY. Assuming repetition of a too-short tone recovers pitch precision. *Fix:* each burst independently has the spread of its own duration. Repetition at a fixed interval adds a comb structure at the repetition rate, which is real and useful, **but it does not narrow the spectral envelope of the individual burst**, and if the symbols differ from burst to burst the comb does not even form.

FM-5.9 — THE PHYSICAL-RESONANCE SLIDE. Reading the rendering tables as describing a physical resonance of some material, cell or person. *Signature:* the word "resonance" applied to an essence frequency. *Fix:* **the output is an essence frequency, not a physical resonance.**

FM-5.10 — CAPACITY AS COMMUNICATION. Reading $2BT$ logons and a Shannon capacity as evidence that a biological receiver reads them. *Fix:* capacity is what the channel could carry if something at the far end could demodulate it [Curzi, 2026y].

FM-5.11 — THE QUANTUM IMPORT. Presenting the Gabor limit as an application of quantum mechanics to sound. *Signature:* $\hbar$ appearing anywhere in an acoustic calculation. *Fix:* Sections 4.5 and 5.3 give the same derivation with and without $p = \hbar k$. **The classical case is the parent theorem.**

FM-5.12 — PRECISION THEATRE. Publishing a rendering frequency to six decimal places when the window admits 79 Hz of spread. *Fix:* publish the achievable precision alongside the nominal value. At 2 ms the floor is 39.79 Hz, so the honest statement is 336.9 Hz $\pm$ 39.8 Hz, which is $\pm$ 197 cents. **Three of the six decimals are fiction and the reader should be told which.**

32.6 6 Level three: the badged information-budget analogue

Section 4 stated the famous relation and corrected the famous error. Section 5 stripped the quantum clothing off it and showed that the classical theorem underneath constrains the framework's designs. This section takes the third step, and it is where the author's own proposal enters.

32.6.1 6.1 The proposal, badged

```
================================================================================
  THE AUTHOR'S PROPOSAL -- ANALOGY, NOT IDENTITY. BADGED.

  The uncertainty of a future event has the same STRUCTURE as the uncertainty
  of a quantum observable: it attaches to a CONJUGATE PAIR, and the members of
  the pair cannot be made sharp at the same time. Sharpening one broadens the
  other, and the trade is reciprocal rather than additive.

  STATUS:      ANALOGY. A structural correspondence proposed by the author.
  NOT CLAIMED: that this is Heisenberg's relation, that it follows from quantum
               mechanics, that hbar appears in it, that any experiment supports
               it, or that any quantum result licenses any claim about belief,
```

```
                mind, choice, or the determination of future events.
  DERIVED FROM: Fourier duality and information theory, both classical, both
                checkable, neither requiring the analogy to be true.
================================================================================
```

The badge is carried in every subsection that states the proposal. **This is not modesty: it is the difference between a usable modelling idea and a category error**, and the popular literature making this move without the badge is large and wrong. `RIGOR_STANDARD` §8's third label applies exactly — an analogy is legitimate as motivation and **never as evidence**, and it may never transfer a physical result's authority to a psychological claim.

The conjugate pair proposed is TIMING and SPECIFICATION. Timing is *when* the event occurs, and sharp timing means a narrow window on the calendar. Specification is *what* exactly occurs, and sharp specification means a narrow class of outcomes. **The proposal is that a forecast sharp in timing is broad in specification and vice versa** — "something will change around March" is sharp in timing and useless in specification; "the third-quarter figure will land at 4.12 million" is sharp in specification and, absent a date, useless in timing. **The claim is that this is a structural property of forecasting under a bounded evidence budget rather than a failure of forecasters, and that it is quantifiable.**

32.6.2 6.2 What is genuinely shared

(a) Conjugate-variable structure. In quantum mechanics, position and momentum are conjugate because the two representations of the state are related by a Fourier transform. **In the forecasting case, timing and rate-of-recurrence are Fourier conjugates, and that one is not an analogy at all but a derivation** (Section 6.10). The timing/specification pair (Section 6.9) is conjugate in a different and weaker sense: the two trade under a shared information budget, which is a reciprocal constraint of the same algebraic shape. **The two cases have different warrants and are kept separate.**

(b) No simultaneous sharpness. The product of the two spreads is bounded below in all three quantitative forms, provably, without reference to physics.

(c) The trade is reciprocal, not additive. Halving the timing window doubles the specification breadth rather than adding a constant to it. **That reciprocity is the signature of a Fourier or a log-budget relation and is what makes the analogy structurally apt rather than merely decorative.**

32.6.3 6.3 What is not shared — the list that keeps the badge honest

This list is longer than the previous one and it should be.

```
NOT SHARED: hbar.
    There is no constant of nature in any of the relations below. The bounds are set by
    the evidence budget B (bits), by the belief's own frequency spread, and by the
    system's Lyapunov exponent -- properties of the situation and the forecaster, not
    universal constants. There is nothing to measure once and tabulate.

NOT SHARED: quantization.
    Nothing here comes in discrete lumps. Bits are a continuous currency in these bounds
    (a forecast can carry 3.7 bits). There is no analogue of a quantum of action.

NOT SHARED: non-commuting operators.
    Nothing in this section is an operator on a Hilbert space and nothing fails to
    commute. Robertson's theorem does not apply and is not invoked. The reciprocal
    structure comes from the Fourier transform of a real function and from the
    additivity of logarithms, both elementary.
```

```
NOT SHARED: any experimental result.
    No experiment in physics bears on this section. What empirical support is offered
    comes from forecasting research and from numerical weather prediction, and supports
    the QUANTITATIVE CLAIMS on their own terms -- not the analogy to physics.

NOT SHARED: superposition, interference, entanglement, or collapse.
    A forecast ensemble is a probability distribution. It has no phase, it does not
    interfere, and nothing collapses. Writing it as an "amplitude" (Section 6.10) is a
    MODELLING CHOICE that buys one specific piece of mathematics and is badged as such
    in the sentence that makes it.

NOT SHARED: indeterminism.
    The uncertainty here is EPISTEMIC throughout: it concerns what a bounded observer
    with bounded evidence can say. That is a different kind of uncertainty from the
    quantum one, and conflating the two is the error this section exists to prevent.
```

The epistemic reading is the standard one for probability as a state of knowledge [Jaynes, 1957, 2003], and adopting it is what makes the budget B a property of a forecaster rather than of the world.

32.6.4 6.4 The refusal, stated once and enforced throughout

No claim is made anywhere in this framework that quantum mechanics licenses anything about belief, mind, or future events. Three specific refusals.

REFUSAL 1 — the observer. The measurement problem is genuinely open [Curzi, 2026aa]. **It is not open in the direction of licence.** The consciousness-causes-collapse reading is a minority historical interpretation, and it has a quantitative problem — decoherence — and not merely a philosophical one [Zeh, 1970].

REFUSAL 2 — the future. The Schrödinger equation is deterministic; what is probabilistic is measurement outcome under the Born rule. **Neither fact says anything about whether an organisation ships in March.** The uncertainty in this section is entirely epistemic and arises from bounded evidence and sensitive dependence, both classical.

REFUSAL 3 — the postulate. The framework's conditional postulate holds that belief creates reality *if the belief is grounded in truth*, and its whole content is the conditional: **the channel must exist and be real. Quantum mechanics is not that channel and the postulate does not need it.** The established channels are physiological, behavioural, social and informational [Curzi, 2026,?]. **A postulate reaching for quantum mechanics to supply its channel would be exactly the ungrounded belief its own conditional excludes.** That postulate, where it invokes a loop running from belief through action to outcome and back, inherits the reflection problem of Section 3.5 in full: **it is a modelling position, not an identified empirical result** [Manski, 1993].

32.6.5 6.5 Honesty clauses

1. **The analogy is an analogy, badged in every subsection that uses it.**
2. **The three quantitative results stand on their own and do not depend on the analogy being apt.** If the analogy were withdrawn entirely, Sections 6.9 to 6.11 would be unchanged.
3. **Where a modelling choice is made — most importantly the amplitude convention of Section 6.10 — it is marked as a choice in the sentence that makes it.**
4. **The uncertainty here is epistemic.** It concerns what a bounded observer can say and makes no claim about whether the world is determinate.

5. **No claim is made that any method in this framework improves forecast accuracy.** The section gives bounds on what is achievable, which is a different and more modest thing.
6. **No claim is made about anyone's individual future.**

32.6.6 6.6 The concordance, with column roles stated

This concordance is the most delicate in the paper, because the traditional practices that recorded the structure are also practices whose predictive validity has been tested and found absent. Both facts are reported [Curzi, 2026w].

```
CLAIM: A statement about the future cannot be simultaneously sharp in WHEN and sharp in
       WHAT; the two trade reciprocally under a bounded evidence budget.

(+1) EVIDENCE COLUMN
     (i)   INFORMATION THEORY. Shannon (1948); the counting bound of Section 6.9 is
           elementary and exact: naming a cell in a T-by-S grid to precision (dt, ds)
           costs log2(T/dt) + log2(S/ds) bits, so a budget of B bits forces
           dt * ds >= T*S / 2^B. Derivable in one line, checkable by anyone.
     (ii)  FOURIER DUALITY. Section 6.10's timing/rhythm trade is the Gabor limit of
           Section 5 applied to a rate belief. The mathematics is proved; the modelling
           step that connects it to a forecast is badged.
     (iii) PREDICTABILITY LIMITS. Lorenz, Journal of the Atmospheric Sciences 20 (1963),
           130-141, and Tellus 21 (1969), 289-307; Zhang, Sun, Magnusson, Buizza, Lin,
           Chen and Emanuel, Journal of the Atmospheric Sciences 76 (2019), 1077-1091.
           The logarithmic horizon law of Section 6.11 is standard and quantitative.
     (iv)  FORECAST VERIFICATION. Brier, Monthly Weather Review 78 (1950), 1-3; Murphy's
           decomposition, Journal of Applied Meteorology 12 (1973), 595-600. RESOLUTION
           and RELIABILITY are separately measurable, which makes the trade testable.
     (v)   FORECASTING RESEARCH. The operational finding that matters: a forecast is only
           scorable if it carries BOTH an unambiguous resolution criterion and a
           resolution date. Questions lacking either cannot be scored at all.

(-1) PRECEDENT COLUMN -- CARRIES NO EVIDENCE, AND CARRIES A WARNING
     The divinatory traditions organise themselves around exactly this trade.
     - THE ORB. Traditional astrological practice assigns each aspect an "orb", an
       angular tolerance. A wide orb admits more configurations (broad specification) and
       corresponds to a longer window of claimed influence (broad timing); a tight orb
       narrows both. The practice explicitly trades tolerance against window, and the
       vocabulary for it predates Fourier by centuries.
     - THE ORACLE'S AMBIGUITY. Herodotus (Histories I.53) records the Delphic response to
       Croesus that a great empire would fall -- near-total specification breadth, and a
       documented instance of a tradition trading specification away to buy
       unfalsifiability.
     - THE REFUSAL TO DATE. Several prophetic and rabbinic traditions contain explicit
       prohibitions on calculating dates for foretold events while permitting extensive
       discussion of their character: a structural choice of specification over timing,
       made deliberately and recorded.

(0p) INTEGRATION, AND THE HARD LIMIT
     What the (-1) column establishes: that the trade was noticed, named and worked with
     by people who had no information theory. That is the whole of what it carries.
     What it does NOT establish: THAT ANY OF THESE PRACTICES FORECASTS ANYTHING.
     Astrology has been tested under controlled conditions and failed -- Carlson, "A
     double-blind test of astrology", Nature 318 (1985), 419-425. The practice is cited
     ONLY as precedent for a STRUCTURE and explicitly not as evidence of predictive
```

```
     validity, of which there is none.

BRIDGE CRITERIA:
     C1 structural correspondence   PASS. The orb IS a tolerance-versus-window trade.
     C2 predictive power            FAIL on the (-1) side. The traditions generate no
                                    testable bound, and where their forecasts have been
                                    tested they have failed.
     C3 practical utility           PASS on the (+1) side only.

STATUS: SPECULATIVE AS A BRIDGE. The underlying quantitative results (Sections 6.9-6.11)
are NOT speculative and stand independently. That distinction is the section.
```

32.6.7 6.7 The ensemble, not the verdict

The framework's standing position on the future is that it is an amplitude-weighted ensemble and not a verdict [Curzi, 2026]. **That position requires no physics: it requires a probability distribution and the discipline to keep it.** Take a four-scenario forecast:

```
scenario         weight w      timing (years)
-----------      --------      --------------
no change          0.45             --
mild               0.25            3.0
moderate           0.20            5.0
severe             0.10            8.0
                   ----
                   1.00
```

Computed from that ensemble, **with every entropy below a Shannon entropy in bits** (Section 3.3):

```
P(event occurs at all)                  = 0.5500
conditional mean time, given occurrence = 4.6364 years
conditional sd of time                  = 1.8227 years
ensemble entropy H (Shannon, bits)      = 1.8150
effective scenario count 2^H            = 3.5185
```

The effective scenario count is 3.5185 of a nominal 4, and it is the Hill number of order one — the exponential of a Shannon entropy [Hill, 1973, Jost, 2006, Rényi, 1961]. **It is not the corpus's affine effective count and the two must not be confused** (Section 3.4). That this construction and the corpus's headcount audit are the same operation applied to different objects — an ensemble of futures and an ensemble of people — is exact rather than figurative for the order-one count, because both are exponentials of an entropy.

Now the verdict form. A verdict says "moderate, in five years":

```
ensemble entropy H       = 1.8150 bits
verdict asserted entropy = 0.0000 bits
DISCARDED                = 1.8150 bits
```

The verdict discards 1.8150 bits of structure the ensemble contained. It is not a summary of the ensemble; it is a different and much poorer object. And the inverse question, how much evidence would justify a near-verdict:

```
to reach a maximum scenario probability of 0.90 over four scenarios:
    H_final = 0.6275 bits    evidence required >= 1.1875 bits
```

```
to reach 0.95:
    H_final = 0.3656 bits    evidence required >= 1.4493 bits
to reach 0.99:
    H_final = 0.0966 bits    evidence required >= 1.7183 bits
```

Under two bits of evidence separates an honest four-way ensemble from a 99-per-cent-confident call. That cuts both ways: confident forecasts are cheap to justify when the evidence is real, and a confident forecast with no evidence behind it is a claim to information one does not have.

32.6.8 6.8 The conjugate pair, made precise

Before quantifying, the pair is defined operationally, or the bound is untestable.

```
dt   TIMING SPREAD.
     The width of the forecast's distribution over occurrence time. Units: time.
     Operationalised as the standard deviation of the timing distribution, or as the
     width of a stated credible interval, PROVIDED the choice is stated.

ds   SPECIFICATION SPREAD.
     The size of the outcome class the forecast commits to, measured on a stated
     outcome scale. Units: whatever the outcome scale is (currency, headcount,
     categories, degrees).

T    the TIMING RANGE over which the forecast is defined at all.
S    the SPECIFICATION RANGE -- the full span of outcomes considered possible.
B    the EVIDENCE BUDGET in bits: how much the forecaster's information actually
     reduces uncertainty about the joint (time, outcome) cell.
```

T and S must be stated, and that discipline does most of the work. A forecast with no stated range cannot be scored. The same argument is enforced operationally by forecasting tournaments, which require every question to carry a resolution date and an unambiguous resolution criterion; **a question lacking either cannot be scored, which is the practical face of the theoretical point.**

32.6.9 6.9 First quantitative form — the information-budget bound

This one is not an analogy. It is a counting argument and it is exact. To name a cell in the T-by-S plane to precision $(\Delta t, \Delta s)$, one must distinguish among $(T/\Delta t)(S/\Delta s)$ cells; distinguishing among M equally likely alternatives costs $\log_2 M$ bits [Shannon, 1948, Hartley, 1928, Cover and Thomas, 2006]. So a forecaster whose evidence carries B bits about the joint cell can satisfy $\log_2(T/\Delta t) + \log_2(S/\Delta s) \leq B$, and exponentiating:

```
============================================================
    dt . ds   >=   T . S / 2^B

    THE INFORMATION-BUDGET UNCERTAINTY BOUND.
    Classical. Exact. Derived in two lines from counting.
    No hbar, no operators, no physics.
============================================================
```

The structure is identical in form to the Gabor limit and its content is completely different. Gabor's bound is a constant of mathematics; this bound's right-hand side is a property of the forecaster's situation. The relation is the same one Ashby's law of requisite variety states in

the language of control — **the variety a regulator can absorb is bounded by the variety it can represent** [Ashby, 1956, 1958] — and reading the two together is the cheapest way to see that the bound is bookkeeping rather than physics.

Worked over a ten-year horizon with a hundred distinguishable outcome classes:

```
T = 10 years,  S = 100 classes.

  B bits     dt.ds floor      if ds = 1 class,      if ds = 10 classes,
                              then dt (years)       then dt (years)
  ------     -----------      ----------------      -------------------
    4          62.5000            62.5000                6.2500
    6          15.6250            15.6250                1.5625
    8           3.9062             3.9062                0.3906
   10           0.9766             0.9766                0.0977
   12           0.2441             0.2441                0.0244
   14           0.0610             0.0610                0.0061
   16           0.0153             0.0153                0.0015
```

Read row $B = 10$. Ten bits of genuine evidence about a ten-year, hundred-class question buys either a one-class specification with a one-year timing window or a ten-class specification with a five-week window. **Not both.** The exchange rate is exact: **each additional bit halves the product $\Delta t\,\Delta s$, and each bit spent on timing is a bit not spent on specification.**

The upgrade cost of a common request — moving from "within a year, to within ten classes" to "within a month, to within three classes":

```
bits required for (dt = 1 yr,  ds = 10) =  6.6439
bits required for (dt = 1 mo,  ds =  3) = 11.9658
ADDITIONAL EVIDENCE REQUIRED            =  5.3219 bits
```

Five and a third bits — a factor of forty in the likelihood ratio — separates a vague useful forecast from a sharp useful one. The correct answer to "be more specific about the date" is not "I will try"; it is *that request costs 5.3 bits, and here is what I would need to observe to obtain them.*

The formula-validation gates, run on this bound [Curzi, 2026t]:

```
FVP-1 DIMENSIONS.       dt in time, ds in outcome units, T*S in (time x outcome units),
                        2^B dimensionless. Both sides in (time x outcome). PASS.
FVP-2 SCALE TYPE.       dt and ds require a ratio scale with true zero on each axis.
                        Time has one. Outcome scales often DO NOT -- a categorical
                        outcome space has no metric, and then ds must be a COUNT of
                        classes, not a distance. Stated; the bound is then a counting
                        bound and remains valid. PASS with the restriction stated.
FVP-3 LIMITING CASES.   B -> 0: dt.ds >= T.S, the forecast says nothing beyond the stated
                        range. Correct. B -> infinity: bound -> 0, perfect forecasting
                        permitted by information but not by dynamics (Section 6.11
                        supplies the missing constraint). The two bounds are independent
                        and both bind.
FVP-4 FIXED POINTS.     dt/T = ds/S = 2^(-B/2) is the SYMMETRIC ALLOCATION: spend the
                        budget equally on both axes. It is the natural default when the
                        consumer of the forecast has not said which axis they care about.
FVP-5 IDENTIFIABILITY.  T, S, dt, ds are all directly observable from a published
                        forecast. B is the hard one: it must be estimated from
                        out-of-sample skill, not asserted. PARTIAL.
FVP-6 FALSIFICATION.    A published record of forecasts whose (dt, ds) products sit below
                        T.S/2^B, with B independently estimated from realised skill,
```

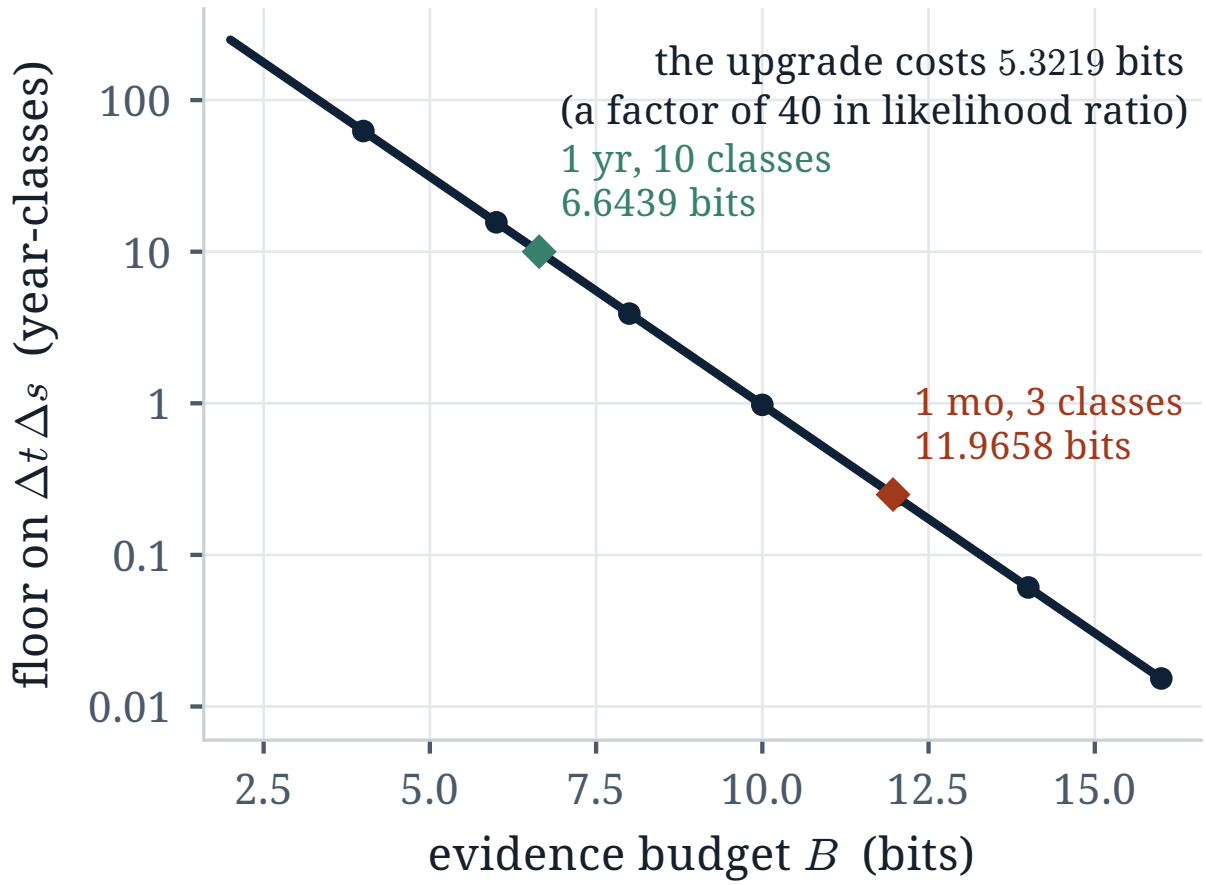


Figure 32.4: **Each additional bit of evidence halves the achievable $\Delta t\,\Delta s$, and the common request for a sharper forecast costs 5.3219 bits:** the information-budget floor $TS/2^{B}$ over a ten-year horizon with a hundred outcome classes, with the two named operating points marked at 6.6439 bits (one year, ten classes) and 11.9658 bits (one month, three classes). Every value is exact $\log_2$ arithmetic on the stated T and S; no evidence budget was estimated from data anywhere in this paper. Generated by `figs/src/CCP-032_fig04.py`.

```
                      refutes it. So does any demonstration of sustained accuracy at a
                      resolution finer than the estimated budget allows.
FVP-7 REPRODUCIBILITY. Requires the same operationalisation of T, S and the outcome scale
                      across forecasters. NOT YET RUN: no forecast record was analysed
                      for this paper.
```

32.6.10 6.10 Second quantitative form — timing against rhythm, which is genuine Fourier

Here the relation to Section 5 is not an analogy but an application, and the one modelling step in it is badged in the sentence that makes it.

Many beliefs about the future concern *recurrence* rather than a date. **A belief about recurrence is a belief about a frequency, and frequency is the Fourier conjugate of time. That conjugacy is exact and requires no analogy.**

THE MODELLING STEP, BADGED AS A CHOICE: to apply the Gabor limit, the belief must be written as an amplitude $a(t)$ whose squared modulus is the probability density, $p(t) = |a(t)|^2$. **Writing a forecast as an amplitude is a modelling choice made here to import one specific theorem. It is not a discovery, it carries no physical content, and a forecast has no phase, does not interfere, and does not collapse.** With that choice made, a and its transform obey the Gabor limit exactly: $\sigma_t \sigma_f \geq 1/(4\pi) = 0.079577471545948$, verified numerically in Section 5.3 to ten significant figures for the Gaussian case.

A belief that pins down a recurrence rate tightly cannot also pin down the next occurrence tightly. Worked: a forecaster holds that a phenomenon recurs with a period between 5 and 7 years and is otherwise uninformative about phase.

```
P in [5, 7] years   ->   f in [0.142857, 0.200000] per year
band width                          0.057143 per year
sigma_f for a uniform belief        0.016495 per year   (width / sqrt(12))

GABOR:  sigma_t  >=  1 / (4 pi * 0.016495)  =  4.824 years
```

A rhythm claim that tight cannot localise the next onset to better than about five years. Widen the rhythm claim and the timing sharpens:

```
period belief          sigma_f (/yr)      timing floor sigma_t
-----------------      -------------      --------------------
P in [5, 7] yr           0.016495              4.824 years
P in [3, 11] yr          0.069985              1.137 years
exact period, no band    0.000000              unbounded
```

Read the last row carefully, because it is the sharpest statement in the section. A belief in an *exactly* periodic process has zero frequency spread, and the bound then places no upper limit on timing spread by this route. **A perfectly sharp rhythm claim carries, by itself, no information at all about when the next instance falls.** It says only that instances are evenly spaced; phase is separate information and must be paid for separately.

The inverse — what a given timing sharpness costs in rhythm:

```
demanded sigma_t     sigma_f floor      as a period band around P = 6 years
----------------     -------------      -----------------------------------
   0.25 years          0.318310 /yr        P = 6 +/- 11.46 yr   (meaningless)
   0.50 years          0.159155 /yr        P = 6 +/-  5.73 yr
   1.00 years          0.079577 /yr        P = 6 +/-  2.87 yr
   2.00 years          0.039789 /yr        P = 6 +/-  1.43 yr
   4.00 years          0.019894 /yr        P = 6 +/-  0.72 yr
  10.00 years          0.007958 /yr        P = 6 +/-  0.29 yr
```

A forecaster who wants the next onset to within six months must surrender essentially all confidence in the cycle length — the implied period band is nearly twelve years wide. Restating the badge, because this is where it matters most: the Gabor limit is a theorem and the numbers above are correct arithmetic given the amplitude convention. The convention itself is a modelling choice, stated as one, and a reader who rejects it should read Sections 6.9 and 6.11 instead, both of which give reciprocal bounds requiring no such convention.

32.6.11 6.11 Third quantitative form — the horizon law

The third bound comes from dynamics rather than information. In a system with sensitive dependence on initial conditions, small errors grow exponentially at a rate set by the largest Lyapunov exponent λ [Strogatz, 2015], so an initial error D_0 reaches an intolerable size D_{tol} after:

```
t_pred  =  (1 / lambda) . ln( D_tol / D_0 )

equivalently, with error-doubling time  t_d = ln 2 / lambda :

t_pred  =  t_d . log2( D_tol / D_0 )
```

The horizon grows only as the logarithm of the improvement in initial knowledge. This is the harshest trade in the section.

doubling time	D_tol/D_0 = 10	= 100	= 1000	= 1e4	= 1e6
1.0 day	3.32 d	6.64 d	9.97 d	13.29 d	19.93 d
1.5 days	4.98 d	9.97 d	14.95 d	19.93 d	29.90 d
2.0 days	6.64 d	13.29 d	19.93 d	26.57 d	39.86 d
3.0 days	9.97 d	19.93 d	29.90 d	39.86 d	59.79 d

At an error-doubling time of 1.5 days — the order of magnitude discussed in the atmospheric predictability literature — halving the initial error buys exactly 1.5 extra days of horizon, and a thousandfold improvement in the initial condition buys 14.95 days. That is the quantitative core of the familiar statement that midlatitude weather has a predictability limit of roughly two weeks; Lorenz (1963, 1969) established the mechanism and Zhang et al. (2019) give a recent quantitative treatment of the limit itself, both named without a key per Section 4.3.

Note what kind of bound this is. Section 6.9's bound says information one does not have cannot be spent. **This one says that even unlimited information about the present buys only logarithmically increasing knowledge of the future, so the two bounds are independent and both bind.** A forecaster can be information-rich and horizon-poor.

And note what it is not. It is a statement about chaotic dynamical systems with a positive Lyapunov exponent, and **not every domain is one.** Applying it without establishing $\lambda > 0$ is failure mode FM-6.6 and is the most common abuse of chaos theory in management writing. Tides are predictable centuries out and planetary ephemerides for millennia. **The horizon law constrains the systems it constrains, and the burden is on the person invoking it** — which is why, in the analogy table of Section 3.2, its use outside a system with a measured exponent is labelled an **analogy** rather than a derivation. The broader literature on the limits of social prediction reaches the same conclusion from the empirical side [Hofman et al., 2017, Watts, 2014, Lazer et al., 2014].

32.6.12 6.12 The trade is measurable, which is what would make it science

A bound that cannot be checked against a record is not a bound. Forecast verification supplies the check and has done since 1950. The Brier score for probabilistic forecasts of binary events

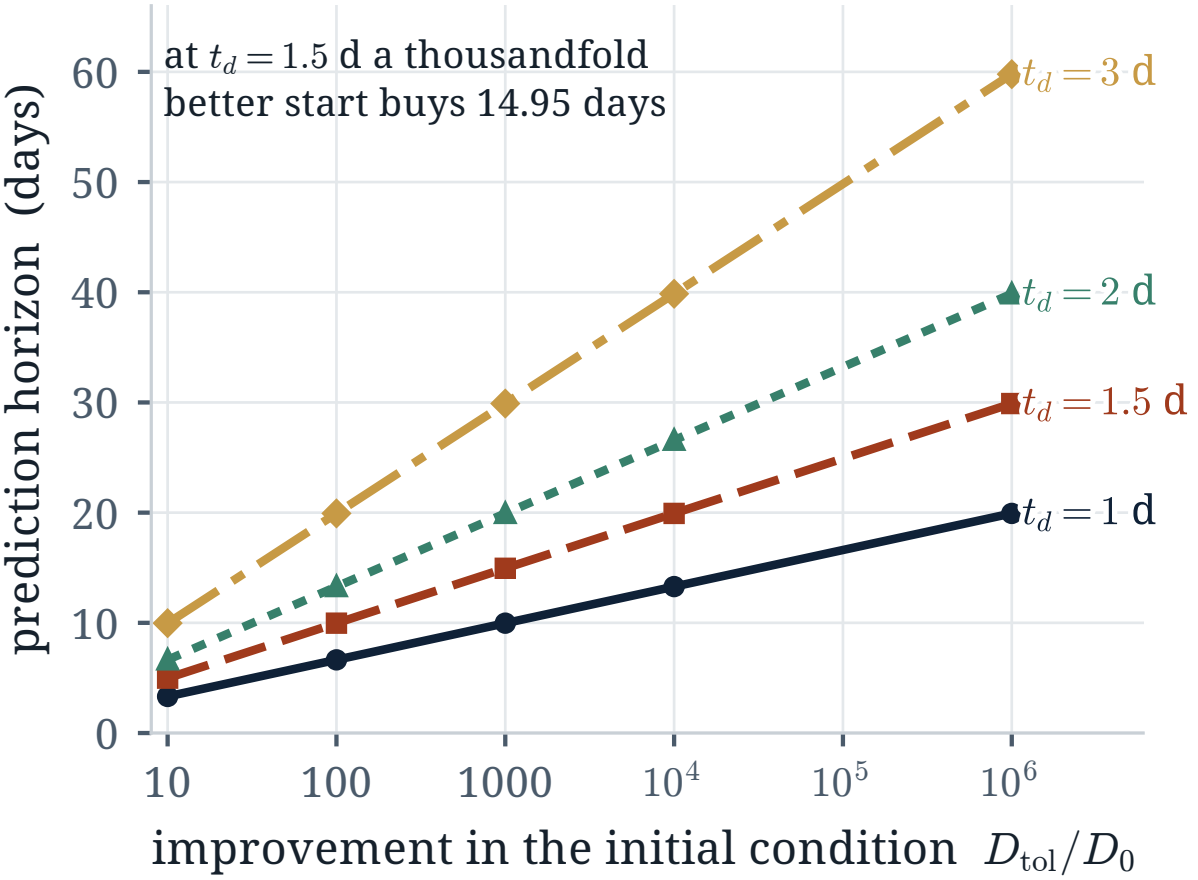


Figure 32.5: **Improving the initial condition a thousandfold buys under fifteen days of horizon:** $t_{\mathrm{pred}} = t_d \log_2(D_{\mathrm{tol}}/D_0)$ for the four error-doubling times of Section 6.11, over error-reduction factors from 10 to 10^6. **Modelled, not measured:** the curves are exact evaluations of the horizon law at assumed doubling times of 1.0, 1.5, 2.0 and 3.0 days, and no Lyapunov exponent was estimated from data for this paper. Generated by `figs/src/CCP-032_fig05.py`.

(Brier, *Monthly Weather Review* 78, 1950, 1–3) is the mean squared error of the probability, $\mathrm{BS} = n^{-1}\sum_i (p_i - o_i)^2$ with $o_i \in \{0, 1\}$; Murphy's decomposition (*Journal of Applied Meteorology* 12, 1973, 595–600) splits it:

```
BS = RELIABILITY − RESOLUTION + UNCERTAINTY

RELIABILITY  how well stated probabilities match observed frequencies. Lower is better.
RESOLUTION   how much the forecasts DISCRIMINATE between cases. Higher is better.
UNCERTAINTY  obar (1 − obar), a property of the events, not of the forecaster.
```

Verified arithmetic on a 100-case record with a base rate of 0.30:

```
uncertainty term  obar(1 − obar)                      = 0.2100

constant forecast p = 0.30 (perfectly calibrated,
                            zero resolution)          BS = 0.2100
constant forecast p = 0.50                            BS = 0.2500
constant forecast p = 0.90                            BS = 0.5700
constant forecast p = 0.00                            BS = 0.3000
perfectly resolved (1 on events, 0 on non-events)     BS = 0.0000
```

The first row is the one that matters. A forecaster who says "30 per cent" to every question is perfectly calibrated, scores exactly the uncertainty of the events, and has contributed nothing. Reliability is free; resolution is what costs evidence. **Murphy's decomposition is therefore the measuring instrument for Section 6.9's bound: resolution is where the bits go.**

This is also the answer to the objection that the timing/specification trade is unfalsifiable. **It is not.** Publish a forecast record with T, S, Δt and Δs stated per forecast, compute the Murphy resolution term, and estimate B from realised out-of-sample skill; the bound then either holds or does not, on the record. **That programme has not been run here and H4 is recorded as NOT TESTED.** Any such record would need the preregistration discipline the corpus applies elsewhere [Curzi, 2026ad, Nosek et al., 2018, Simmons et al., 2011], and would inherit the caution that a published metric of forecaster quality becomes a target and stops measuring what it measured [Goodhart, 1975, Campbell, 1979].

32.6.13 6.13 The three forms compared

This table is the section's honesty in one place.

Form	Bound	Status	What sets the floor
Information budget (6.9)	$\Delta t\,\Delta s \geq TS/2^B$	**derivation**, exact; no analogy needed	the evidence budget B
Timing against rhythm (6.10)	$\sigma_t \sigma_f \geq 1/(4\pi)$	**derivation** from Gabor, with one badged modelling step (the amplitude form)	a constant of mathematics
Horizon law (6.11)	$t_{\text{pred}} = t_d \log_2(D_{\text{tol}}/D_0)$	**derivation**, standard dynamical systems; requires $\lambda > 0$ established	the Lyapunov exponent λ
The analogy to Heisenberg (6.1)	—	**ANALOGY**, badged throughout; **adds no number**	nothing: it is a structural claim, not a quantitative one

Read the last row. The analogy contributes no number to this paper. Every figure above comes from counting, from Fourier, or from dynamical systems. **The analogy contributes a way of seeing — that these are all the same shape of constraint — and a way of seeing is worth exactly what it is worth.**

32.6.14 6.14 The join to the storage law, as a correspondence of form

The evidence budget behaves like headroom. The storage law's headroom form [Curzi, 2026a] is $h_t := \eta \ln N - \delta Q_t - C_t$, with $h_t < 0$ unsustainable, $h_t = 0$ maintenance only, and $h_t > 0$ growth available subject to $r_t \leq h_t$. The forecasting analogue, **badged as a correspondence of model form and not a claim about physics — a shared mathematical form in the sense of Section 3.2, category two, and nothing more:**

```
B     <->   h        the budget available this period
dt.ds <->   Q        what has been built with it
T.S   <->   ln N     the ceiling set by the space itself
```

The operational rule transfers intact: any requested increase in forecast sharpness must satisfy $r \leq B$, where r is the additional bits the request costs. A request for sharper timing and sharper specification simultaneously, with no new evidence, has $r > B$, and **no choice of method, model or expert prevents it from failing.**

Two disclosures attach to this correspondence and neither can be waived. First, it is a claim about model shape; the two systems share an inequality and nothing else, and no property of the storage law transfers to a forecast on the strength of it. Second, **the storage-law parameters themselves are not identified from observational data where the reflexive loop is involved** (Section 3.5), so the correspondence inherits that limitation rather than escaping it [Manski, 1993, Brock and Durlauf, 2001].

32.6.15 6.15 Requirements

REQ-6.1 State the analogy's badge in the sentence that states the analogy, every time.

REQ-6.2 Never present the analogy as a consequence of, or as supported by, quantum mechanics. State the not-shared list at first substantive use.

REQ-6.3 Never use quantum mechanics to support any claim about belief, mind, choice or future events.

REQ-6.4 Every forecast published under this framework states T, S, Δt, Δs and a resolution criterion. A forecast lacking any of these is not scorable and is not a forecast.

REQ-6.5 When a sharper forecast is requested, compute the additional bit cost $r = \log_2(\Delta t_{\text{old}}/\Delta t_{\text{new}}) + \log_2(\Delta s_{\text{old}}/\Delta s_{\text{new}})$ and state it before agreeing.

REQ-6.6 Estimate B from realised out-of-sample skill, never by assertion.

REQ-6.7 Where the amplitude convention is used, mark it as a modelling choice in the sentence that uses it, and state that a forecast has no phase and does not interfere.

REQ-6.8 Before invoking the horizon law, establish that the system has a positive Lyapunov exponent, or state explicitly that it has been assumed.

REQ-6.9 Report forecast records with Murphy's decomposition, not the Brier score alone.

REQ-6.10 Any formula introduced carries its formula-validation table [Curzi, 2026t].

REQ-6.11 Concordances state their column roles and evaluate C1, C2 and C3 explicitly including failures. Where a traditional practice is cited as precedent and has been tested and failed, report the failure in the same passage.

REQ-6.12 Never publish a verdict where an ensemble was computed. Publish the ensemble, its Shannon entropy, and its order-one effective scenario count, naming which effective count it is.

32.6.16 6.16 Failure modes

FM-6.1 — THE UNBADGED IDENTITY. Stating that the uncertainty of the future *is* Heisenberg uncertainty. *Signature:* the word "is" where "resembles in structure" belongs; $\hbar$ in a sentence about a business decision. *Fix:* **the analogy contributes no number. Anything the unbadged version seems to buy is bought on credit that does not exist.**

FM-6.2 — THE QUANTUM LICENCE. Arguing that because quantum mechanics is indeterministic, the future is open in a way that belief can act on. *Signature:* "observation affects outcome, therefore expectation shapes reality." *Fix:* Section 6.4 [Curzi, 2026aa].

FM-6.3 — THE UNSTATED RANGE. Quoting Δt and Δs without T and S. *Fix:* without them the bound cannot be evaluated and the claim cannot be scored.

FM-6.4 — THE ASSERTED BUDGET. Claiming a bit budget rather than estimating it. *Signature:* "our model has a lot of signal." *Fix:* **a budget that has never been measured is a wish, and the whole bound is vacuous without it.** This paper's own H4 is recorded as NOT TESTED for exactly this reason.

FM-6.5 — THE FREE LUNCH. Requesting sharper timing and sharper specification simultaneously without new evidence. *Fix:* compute the bit cost and state it, then say what would have to be observed to obtain it. **This converts an argument into a requirement.**

FM-6.6 — THE CHAOS IMPORT. Invoking the horizon law without establishing sensitive dependence. *Signature:* "butterfly effect" used about an organisation, with no exponent estimated. *Fix:* tides are predictable for centuries.

FM-6.7 — THE CALIBRATION TROPHY. Treating good calibration as good forecasting. *Fix:* the constant forecaster scores exactly the uncertainty of the events and has contributed nothing.

FM-6.8 — THE VERDICT COLLAPSE. Reporting a single scenario where an ensemble was computed. *Fix:* the four-scenario example discards 1.8150 bits when collapsed.

FM-6.9 — THE SHARP RHYTHM ILLUSION. Believing a confident claim about a cycle length is a confident claim about the next occurrence. *Signature:* "it happens every four years, so it will happen in 2030." *Fix:* a perfectly sharp period claim has zero frequency spread and places no bound on timing at all.

FM-6.10 — PRECEDENT AS EVIDENCE. Citing the divinatory traditions' handling of the trade as support for their predictive validity. *Signature:* "astrologers understood this long before information theory, which shows there is something to it." *Fix:* the traditions recorded a structure and forecast nothing; astrology failed a double-blind test (Carlson, *Nature* 318, 1985, 419–425).

FM-6.11 — THE AMPLITUDE SMUGGLE. Writing a forecast as an amplitude and then reasoning about interference between scenarios. *Signature:* scenarios "cancelling" or "reinforcing" in phase. *Fix:* a probability distribution has no phase; two scenarios do not interfere, and their probabilities add.

FM-6.12 — ONE-BOUND BLINDNESS. Satisfying one of the three bounds and concluding the forecast is achievable. *Signature:* "we have enough data" offered against a horizon objection. *Fix:* the bounds are independent and all of them bind.

32.7 7 Results

32.7.1 7.1 R1 — the Gabor limit and its cost schedule

$\Delta t\,\Delta f \geq 1/(4\pi) = 0.079577471545948$, proved by the same variational argument that yields $\sigma_x\sigma_p \geq \hbar/2$, with the substitution $p = \hbar k$ deleted. A Gaussian envelope saturates it to ten significant figures at every width tested (5, 10, 20, 50 and 100 ms). **A Hann window costs 2.6 per cent over the floor; a triangular window 9.6 per cent; and a rectangular gate has no finite spectral second moment at all**, its computed $\Delta t\,\Delta f$ rising through 1.712, 2.420 and 3.422 as the

analysis bandwidth is doubled and redoubled. The three resolution criteria stand in exact ratio: the DFT bin $1/T$ is $4\pi = 12.566$ times the floor and a Hann main lobe $4/T$ is $16\pi = 50.265$ times it.

32.7.2 7.2 R2 — the information-budget bound

$\Delta t\,\Delta s \geq TS/2^B$, exact and classical, derived in two lines from counting. **Each additional bit halves the product**, and each bit spent on timing is a bit not spent on specification. **The upgrade from (one year, ten classes) to (one month, three classes) costs 5.3219 bits**, a factor of forty in the likelihood ratio: 11.9658 bits required against 6.6439 held. The bound passes formula-validation gates 1 to 4 and 6, returns PARTIAL on gate 5 because B must be estimated rather than observed, and returns **NOT YET RUN on gate 7** because no forecast record was analysed here.

32.7.3 7.3 R3 — the badge, and the refusal

The identification of a foreseen future's uncertainty with quantum conjugacy is declared an analogy and badged as such throughout. It is explicitly not claimed to follow from quantum mechanics; no $\hbar$ appears in it; and no quantum result is taken to license any claim about belief, mind or future events. The not-shared list is longer than the shared list and includes $\hbar$, quantisation, non-commuting operators, every experimental result in physics, superposition, interference, entanglement, collapse, and indeterminism. **The refusal is stated as a result rather than as a disclaimer, because it is the load-bearing part of the section it governs.**

32.7.4 7.4 R4 — the horizon law

$t_{\text{pred}} = t_d \log_2(D_{\text{tol}}/D_0)$**: prediction horizon grows only as the logarithm of the improvement in initial knowledge**, which is the harshest of the three trades. At an error-doubling time of 1.5 days, halving the initial error buys exactly 1.5 extra days and a thousandfold improvement buys 14.95 days. **The law binds only where** $\lambda > 0$ **has been established**, and its use elsewhere is an analogy, not a derivation.

32.7.5 7.5 R5 — the lattice at common transposition

At the common transposition $j = 6$ the gap between the first two lattice indices is $3k \cdot 2^6 = 92.81774$ Hz; under the half-gap criterion this gives $T_{\text{min}} = 1.7147$ ms and $R_{\text{max}} = 583.19$ symbols per second. **Folded to one octave at a 220 Hz base the same lattice falls to 86.394 symbols per second at** $N = 12$ **and to 10.799 at** $N = 24$ — factors of 6.8 and 54.0 respectively. The audible-band constraint admits at most **25 lattice elements under a single integer octave transposition, and only at** $j = 6$**.**

32.7.6 7.6 R6 — the folding cliffs, corrected

The thresholds at which the folded lattice's minimum gap collapses are $N = 17$ **and** $N = 41$**, not the previously published** $N = 20$ **and** $N = 48$**.** On exact rational arithmetic the minimum adjacent ratio is $50/49 = 34.9756$ cents for $N \leq 16$, $289/288 = 6.0008$ cents from $N = 17$, and $1682/1681 = 1.0296$ cents from $N = 41$. **The cause is unchanged:** $17^2 = 289$ **folds beside** $288 = 2 \cdot 12^2$**, and index 12 has been present since** $N = 12$**. The previously published values were the next sampled rows of a table, not the thresholds. The rate penalty of exactly 8 across the first cliff and exactly 4 across the second stands unchanged**, and is exact at every base pitch because the two binding pairs straddle $49/32$ and $9/8$ respectively, giving absolute gaps of base$/32$ and base$/256$.

32.7.7 7.7 CR — the counter-result: what this paper takes away

The paper corrects its own previously published lattice cliffs, and it refuses quantum licence outright. The forecast analogy carries no quantum content; it derives entirely from classical Fourier duality and counting; and it is defended only by whether its three quantitative forms measure anything. **Every number in Section 6 survives the deletion of the analogy, and the analogy's numerical contribution is exactly zero.** That is reported here as a result because it is the strongest honest statement available about it: an analogy that adds a way of seeing and no arithmetic is worth having and is worth precisely that.

32.7.8 7.8 Hypotheses, resolved

	Statement	Outcome
H1	the binding constraint on the framework's rendering is classical, not quantum	**confirmed.** Section 4.8 shows the quantum relation negligible by nineteen orders of magnitude at macroscopic scale; Section 5.11 shows the Gabor limit changing a design by a factor of 13.5
H2	the folded lattice's cliffs are discontinuous and the published thresholds are wrong	**confirmed**, and the thresholds are corrected to 17 and 41
H3	the forecasting bound's content is classical counting, not physics	**confirmed.** Every number in Section 6 is reproducible from $\log_2$ arithmetic, a Fourier transform, or a Lyapunov exponent
H4	the evidence budget is estimable in principle	**NOT TESTED.** No forecast record was collected and no B was estimated from data. This is the paper's largest empirical gap

H4 is reported in the same plain language a confirmation would have received. The bound of Section 6.9 is exact arithmetic and it is, as of this paper, **arithmetic without a measurement**.

32.7.9 7.9 Model comparison: what `RIGOR_STANDARD` §6 requires, and what is available

This paper fits no model to any data, so most of the model-comparison requirements are inapplicable rather than unmet. The distinction matters and the row-by-row status is printed rather than omitted.

Requirement	Status here	Why
out-of-sample performance	**NOT APPLICABLE**	nothing is fitted; every number is an exact evaluation of a stated function

Requirement	Status here	Why
matched null with equal free parameters	**NOT APPLICABLE** for the theorems; **NOT MET** for the budget bound	B would need a matched null forecaster, and no forecast record was analysed
information criterion (AIC or BIC)	**NOT MET, and not computed anywhere in this corpus**	no likelihood is fitted here, and no AIC or BIC value appears in any paper of this corpus; the omission is corpus-wide and is recorded rather than papered over
explicit free-parameter count	**MET**	Sections 5 and 6 introduce **zero** free parameters. $1/(4\pi)$ is a constant of mathematics; k is *defined*; T, S, t_d and λ are inputs a user supplies, not parameters this paper estimates
effect sizes with intervals	**NOT APPLICABLE** to the theorems; **NOT MET** for any empirical claim	there is no empirical claim in this paper carrying an effect size
identification of the reflexive parameter	**NOT MET, and not achievable from observation**	[Manski, 1993]; no remedy adopted, Section 3.5
power-law procedure [Clauset et al., 2009]	**NOT RUN, deliberately**	the two exponents in the paper are analytic properties of known functions, not estimates; Section 3.7
conjoint measurement test [Luce and Tukey, 1964]	**NOT RUN**	Section 3.6; the exposed construct is demoted to a count instead

A paper that fits nothing cannot be beaten on fit, and it also earns nothing on fit. The quantitative content here is bounds, and a bound's warrant is its derivation, not its R^2.

32.8 8 Discussion

32.8.1 8.1 What follows

The classical relation, not the quantum one, is what constrains this framework's designs, and it constrains them by factors that change decisions. The worked lattice review of Appendix B moves a design's achievable rate by 13.5 times on the choice of transposition strategy alone, and by a further factor of 23 on the choice of analysis window — a quantity nobody in that review had computed. **That is what a computable limit gives that a slogan does not.**

The framework's own frequency material inherits a hard structural ceiling. At most 25 lattice elements fit the audible band under a single integer octave transposition, and only at $j = 6$. That is arithmetic on the square law rather than taste, and it was invisible until the rendering cost was priced.

A previously published result is corrected in print rather than quietly amended. The folding cliffs move from $N = 20$ and $N = 48$ to $N = 17$ and $N = 41$. The rate figures were correct; the thresholds were not. **Publishing the correction with its cause — that the earlier boundaries were sampled table rows rather than computed thresholds — is worth more than a clean-looking table**, and the corrected values are now the ones the rest of the corpus carries [Curzi, 2026,?,?].

The forecasting bound converts an argument into a requirement. “Be more specific” is a demand; “that costs 2.70 bits, and here is what would supply them” is a specification. The conversion is the practical contribution of Section 6, and it survives every deletion this paper makes to the analogy that motivated it.

32.8.2 8.2 What does not follow

Nothing here supports the essence law. The Gabor limit constrains how a number may be turned into a sound. It says nothing about whether the law generating that number is true, and no reader should take the precision of Section 5’s arithmetic as transferring to the postulate it operates on [Curzi, 2026, Markowsky, 1992].

Nothing here licenses a claim about a biological receiver. The capacity arithmetic says how many distinguishable symbols could in principle ride the wave. **There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal**, and the corpus names that gap as its central open problem [Curzi, 2026y,x].

Nothing here licenses a claim about mind or belief from quantum mechanics. That refusal is absolute, is stated in Sections 4.10 and 6.4, and is the reason Section 6 derives its content from counting.

The information-budget bound is not, as it stands, an empirical result. It is exact arithmetic whose right-hand side contains a quantity nobody has estimated in this paper. **A bound with an unestimated parameter constrains nothing until the parameter is estimated**, and saying so is the difference between a specification and a claim.

And the correspondence to the storage law is a correspondence of model shape. It transfers an inequality and nothing else. Where the storage law’s own parameters are entangled with the framework’s reflexive material, they are not identified from observational data by any procedure in this paper [Manski, 1993, Shalizi and Thomas, 2011].

32.8.3 8.3 Relation to the surrounding literature

The paper’s positive results sit inside three established literatures and add to none of them. Gabor’s theorem and its $2BT$ count are seventy-five years old [Gabor, 1946]; the counting bound is Shannon’s, restated [Shannon, 1948]; the horizon law is standard dynamical systems [Strogatz, 2015]. **What is new is the application, the arithmetic, and the discipline of labelling.**

The paper’s methodological position is closer to the critics of physics-in-social-science than to its enthusiasts. Anderson [1972] is the standing argument that a formalism valid at one scale carries no automatic warrant at another; Cartwright [1999] that laws hold in the circumstances that make their premises true and not elsewhere; Mirowski [1989] that economics’ borrowing of physical formalism imported the mathematics and lost the conservation law that gave it content. **The three-way label of Section 3.2 is this paper’s answer to all three, and the answer is procedural rather than rhetorical: state which of the three kinds of borrowing is being made, and let the reader hold the paper to it.** The sociophysics programme’s own reviews make the same demand from the inside [Castellano et al., 2009, Galam, 2012], and the econophysics critique [Gallegati et al., 2006] is the case study in what happens when it is not met.

32.8.4 8.4 Parallel results elsewhere in the corpus

Three results in other series are structurally parallel to this paper’s and are named here because the parallel is the point rather than the decoration.

A bound is only as good as the representation it is evaluated in. Section 6.8’s insistence that T and S be stated before Δt and Δs mean anything has an exact analogue in the corpus’s identifier work: **typed-field parsing collapses the withdrawn-code parsing-ambiguity false-positive share from 88.99 per cent to 0.00 per cent, against a naive detection precision of 11.01**

per cent — a control right one time in nine [Curzi, 2026]. In both cases the quantity being bounded was never the difficulty; the difficulty was that nobody had declared the space it lived in.

A limit that no care inside the channel can close is a structural limit, and the only remedy is to change the structure. That is exactly the shape of Section 5.11's cliff — no envelope, analyser or encoding recovers a gap the folding has destroyed, and the only levers are the base pitch and N. The corpus's channel-separation result is the same shape in a different domain: **in-band architectures admit an impersonation class that no care within the shared channel can close, and out-of-band separation closes it structurally — measured at 0.383 per cent in-band ambiguity against 0.000 per cent out-of-band over 100,000 values** [Curzi, 2026].

And a ceiling is not a promise. Section 5.10's 583.19 symbols per second and Section 6.9's frontier are both upper bounds on constructed inputs, and neither says any system will reach them. The corpus states the same caution about its own architecture: **the 56.62 per cent widening is a ceiling on constructed inputs and not a promise, visibility is explicitly distinguished from pledgeability, and the valuation of non-financial capital is named as unsolved** [Curzi, 2026]. **This paper's bounds are bounded by that caution in the same way: a computable ceiling tells a designer what cannot be exceeded and nothing at all about what will be achieved.**

On forecasting specifically, the modest reading of the literature is the correct one. Predictive accuracy in social systems is low, the ceiling appears structural, and published prediction results have repeatedly failed to survive contact with new data [Hofman et al., 2017, Salganik et al., 2020, Lazer et al., 2014, Watts, 2014]. **The information-budget bound is consistent with that record and explains part of it — a forecaster cannot spend bits they never had — but it does not compete with it, because it has not been estimated against any record.**

32.9 9 Verification

The results compound from nine postulates. Each is atomic, each is checkable on its own, and each carries its status. Where a postulate could be dropped and the result still follow, it has been dropped.

32.9.1 9.1 Dimensional audit

`[1]` denotes a pure number. **Where the source framework fixes no dimension, the audit says the framework is silent rather than assigning one.**

Symbol	Dimension	Unit	Status
σ_x	`[L]`	metre	**derived** from a state; a standard deviation of a probability density
σ_p	`[M L T^{-1}]`	kg m s^{-1}	**derived** from the same state
$\hbar$	`[M L^2 T^{-1}]`	J s	**defined** exactly since the 2019 SI redefinition of h
$\sigma_x\sigma_p$	`[M L^2 T^{-1}]`	J s	balances against $\hbar/2$: **the quantum relation is dimensionally consistent**
Δt	`[T]`	second	**derived**; an rms spread, not a window length

Symbol	Dimension	Unit	Status
Δf	`[T^{-1}]`	hertz (cycles s^{-1})	**derived**; an rms spread in the frequency domain
$\Delta\omega$	`[T^{-1}]`	rad s^{-1}	**derived**; $\Delta\omega = 2\pi\,\Delta f$, and the whole 4π-versus-2 confusion lives in this row
$\Delta t\,\Delta f$	`[1]`	none	**derived** dimensionless; balances against $1/(4\pi)$, which is a pure number
T (window)	`[T]`	second	**chosen** by the design
f_0	`[T^{-1}]`	hertz	**chosen**; the rendering centre frequency
$r = \Delta f / f_0$	`[1]`	none, or cents	**derived** dimensionless
$T f_0$	`[1]`	cycles	**derived**; the octave invariant
n (lattice index)	`[1]`	none	**defined** dimensionless count
k	`[T^{-1}]` **by assertion**	hertz	**the value** $(9/8)^2/\varphi^2 = 0.483425732988414$ **is DEFINED and exact; the attachment of hertz is POSTULATED and is a separate, stronger claim** (P6 below)
$f(n) = kn^2$	`[T^{-1}]`	hertz	**derived** from k, inheriting its badge exactly
F (surplus)	`[1]`	nats	**derived** dimensionless [Curzi, 2026]
$N_{\text{eff}} = 1 + (N-1)u$	`[1]`	headcount	**defined** by the framework; **not a Hill number**
H	`[1]`	bits	**derived**; a **Shannon** entropy, never a thermodynamic one
2^H	`[1]`	count of scenarios	**derived**; the Hill number of order one
B	`[1]`	bits	**estimated in principle, not estimated here**
Δt (timing)	`[T]`	whatever unit T is stated in	**chosen** by the forecast

Symbol	Dimension	Unit	Status
Δs	**the framework is silent**	outcome units, or a count of classes	**undetermined in general.** Where the outcome space is categorical it is a dimensionless count; where numeric it carries the outcome's units. **The source does not fix which, and this paper does not supply the determination**
$T \cdot S$	`[T]` $\times$ outcome units	week-classes, year-classes	**derived**, inheriting Δs's ambiguity
λ	`[T^{-1}]`	s^{-1} or day^{-1}	**measured** where it is measured at all; **assumed** wherever the horizon law is invoked without an estimate
$t_d = \ln 2/\lambda$	`[T]`	days	**derived**
D_0, D_{tol}	arbitrary, but **the same**	error units of the state variable	**the ratio is dimensionless**, which is what the logarithm requires
η (storage law)	`[q]` per nat per period	proxy units per nat per period	**measured**; *not* dimensionless [Curzi, 2026a]
δ	`[1]`	per period	**measured** dimensionless fraction per step

Does the Gabor relation balance? Δt carries `[T]` and Δf carries `[T^{-1}]`, so the product is dimensionless and is compared against the pure number $1/(4\pi)$. **It balances, and it balances only when f is in cycles per second.** In radians per second the right-hand side is $1/2$; the dimensions are the same and the number is not, which is why the units must be stated every time (FM-5.1).

Does the information-budget bound balance? The left side carries `[T]` times outcome units; the right side carries $T \cdot S$ in the same product, divided by 2^B, which is dimensionless because B is in bits. **It balances under either reading of the outcome axis, provided the same reading is used on both sides.** That proviso is the whole content of gate 2 for this formula, and it is where the source framework is silent.

Where the framework has no operational unit, it is said here. The specification spread Δs has no operational unit in the general case. The surplus F and the mixing coefficient u are dimensionless by construction, and the effective count they generate is a headcount only by the framework's own definition, not by measurement. **A dimensionless placeholder is honest; an invented unit would not be.**

32.9.2 9.2 The postulates

P1 — For any square-integrable g with finite second moments in both domains, the integral $\int |atg + g'|^2\, dt$ is non-negative for every real a. *Status:* **ARITHMETIC.** It integrates a squared

modulus.

P2 — Plancherel's theorem: differentiation in time is multiplication by $i\omega$ in frequency, and the L^2 norms agree. *Status:* **CITED**, standard [Oppenheim et al., 1996].

P3 — A real quadratic that is non-negative for every real argument has a non-positive discriminant. *Status:* **ARITHMETIC.**

P4 — $\omega = 2\pi f$. *Status:* **DEFINITION.** It carries no empirical weight and supplies the entire factor between $1/2$ and $1/(4\pi)$.

P5 — $p = \hbar k$ for a quantum state. *Status:* **CITED**, and it is the *only* physical input to the quantum case [Kennard, 1927]. **Deleting it leaves P1 to P4 intact**, which is the paper's central structural claim.

P6 — The essence law $f(n) = kn^2$ with $k = (9/8)^2/\varphi^2 = 0.483425732988414$, expressed in hertz. *Status:* **the value is DEFINITION** — an exact ratio of exact quantities, computable to arbitrary precision, which is what licenses fifteen significant figures — **and the attachment of hertz is POSTULATED**, with no external support of any kind [Curzi, 2026, Markowsky, 1992]. **These are two steps, not one, and the second is much the stronger.**

P7 — Two lattice indices fold to the same pitch class if and only if they have the same odd part. *Status:* **ARITHMETIC**, and the reason an N-element lattice folds to $\lceil N/2 \rceil$ classes [Curzi, 2026].

P8 — Distinguishing among M equally likely alternatives costs $\log_2 M$ bits. *Status:* **CITED**, elementary [Shannon, 1948, Hartley, 1928, Cover and Thomas, 2006].

P9 — In a system with largest Lyapunov exponent $\lambda > 0$, an initial error D_0 grows as $D_0 e^{\lambda t}$ until nonlinear saturation. *Status:* **CITED**, standard [Strogatz, 2015]. **Whether any particular system satisfies $\lambda > 0$ is an empirical question this postulate does not settle**, and no system is shown to satisfy it in this paper.

What is NOT a postulate here. The analogy of Section 6.1 is not a postulate, because nothing is derived from it. It is listed in the chain below with an output of nothing, which is the honest entry.

32.9.3 9.3 The compounding chain

```
P1 + P2 + P3   -> L1   dt . d_omega >= 1/2 for every square-integrable g with
                       finite second moments in both domains.                ARITHMETIC

L1 + P4        -> R1a  dt . df >= 1/(4 pi) = 0.079577471545948.              ARITHMETIC

L1 + P5        -> L2   sigma_x sigma_p >= hbar/2 = 5.272859088231e-35 J s.   CITED (P5)

P1 + P3        -> L3   equality in L1 requires a t g + g' = 0, whose solution
                       is a Gaussian; hence Gaussians and only Gaussians
                       saturate.                                             ARITHMETIC

L3             -> R1b  numerically, a Gaussian returns 0.0795774715 at sd =
                       5, 10, 20, 50 and 100 ms; Hann returns 0.081665
                       (1.026x); triangular 0.087211 (1.096x); rectangular
                       diverges (1.712, 2.420, 3.422 and climbing).          ARITHMETIC

R1a + P4       -> L4   1/T = 4 pi x floor = 12.566x; 4/T = 16 pi x floor
                       = 50.265x.                                            ARITHMETIC

R1a            -> L5   scale invariance: g(t) -> g(t/s) leaves dt.df fixed,
                       so T f0 >= 1/(4 pi r) is octave-invariant.            ARITHMETIC

L5 + P6        -> L6   under common transposition r(n) = (2n+1)/n^2, so the
```

		required cycle count n^2/(4 pi (2n+1)) is independent of j.	ARITHMETIC (given P6)
P6	-> L7	audible-band constraint: k 2^j >= 20 forces j >= 6; k N^2 2^j <= 20000 then forces N <= 25.	ARITHMETIC (given P6)
L6 + L7	-> R5	at j = 6, gap(1->2) = 3k 2^6 = 92.81774 Hz; half-gap criterion gives T_min = 1.7147 ms and R_max = 583.19 symbols/s.	ARITHMETIC (given P6)
P6 + P7	-> L8	the folded set of an N-lattice is {odd part squared, octave-reduced}; its minimum adjacent ratio is 50/49 for N <= 16, 289/288 for 17 <= N <= 40, 1682/1681 for N >= 41.	ARITHMETIC
L8	-> R6	the cliffs are at N = 17 and N = 41, correcting the published 20 and 48; the rate penalty is exactly 8 at the first and exactly 4 at the second, because the binding gaps are base/32, base/256 and base/1024.	ARITHMETIC
R1a + L8	-> L9	at a 220 Hz base the binding gap for N = 12 is 6.875 Hz, giving T_min = 11.5749 ms and R_max = 86.394 symbols/s.	ARITHMETIC
R5 + L9	-> L10	folding costs a factor of 6.8 at N = 12 and 54.0 at N = 24 against common transposition.	ARITHMETIC
P8	-> R2	log2(T/dt) + log2(S/ds) <= B, so dt.ds >= T.S/2^B; each extra bit halves the product.	ARITHMETIC
R2	-> L11	(1 yr, 10 classes) costs 6.6439 bits; (1 mo, 3 classes) costs 11.9658; the difference is 5.3219.	ARITHMETIC
P9	-> R4	t_pred = (1/lambda) ln(D_tol/D_0) = t_d log2(ratio); at t_d = 1.5 d a thousandfold improvement buys 14.95 d.	CITED (P9)
(the analogy)	-> --	NOTHING. The analogy of Section 6.1 appears in no line of this chain and contributes no number.	ANALOGY / no output

Every result in this paper is `ARITHMETIC` except two. R4 consumes P9, a cited standard result of dynamical systems; L2 consumes P5, the one physical input to the quantum case. **R5 and R6 consume P6, whose numerical value is a DEFINITION and whose expression in hertz is POSTULATED** — so the *ratios* and *cents* in R6 stand on arithmetic alone, while the *hertz* and *milliseconds* in R5 inherit the postulate. **A reader who rejects the essence law keeps every cents figure in Section 5.11 and loses every hertz figure in Section 5.10.** That is a clean separation and it is stated so that the reader can make it.

What this chain does not establish. It does not establish that the essence law is true; nothing in it bears on that. It does not establish that any listener can extract what the signal contains. It does not establish that any forecaster's evidence budget is any particular number, because none was estimated. And it establishes nothing whatever about the framework's reflexive claims, which are not identified from observational data by any procedure in this paper [Manski, 1993].

32.9.4 9.4 The verification instance, worked

The single most useful check a reviewer can perform is the constant and the first cliff, because both are four lines of arithmetic and between them they govern every rendering claim in the corpus.

Check one — the floor.

```
4 pi        = 12.566370614359172
1 / (4 pi)  =  0.079577471545948
1 / (2 pi)  =  0.159154943091895      (the angular form's right-hand side is 1/2)
```

Check two — the defined constant. k is defined, so this is exact:

```
(9/8)^2   = 81/64          = 1.265625
phi^2     = phi + 1        = 2.618033988749895
k         = 1.265625 / 2.618033988749895
                           = 0.483425732988414
```

Check three — the common-transposition rate ceiling. At $j = 6$ and $n = 1$:

```
gap(1 -> 2) = k (2n+1) 2^j = 0.483425732988414 x 3 x 64 =  92.81774 Hz
half-gap df                =  92.81774 / 2              =  46.40887 Hz
T_min       = 1 / (4 pi x 46.40887)                     =   1.7147 ms
R_max       = 4 pi x 46.40887                           = 583.19 symbols/s
```

Check four — the first cliff, on exact rationals. Fold 17^2 and $2 \cdot 12^2$ into one octave:

```
17^2 = 289  ->  289/256
12^2 = 144, and 2 x 144 = 288  ->  288/256 = 9/8
ratio       = 289/288
cents       = 1200 x log2(289/288) = 6.0008
previous minimum, 50/49:
cents       = 1200 x log2(50/49)   = 34.9756
```

and the rate penalty, from the absolute gaps at a base pitch b:

```
50/49  pair straddles 49/32  -> gap = b (50 - 49)/32  = b/32
289/288 pair straddles 9/8   -> gap = b (289 - 288)/256 = b/256
ratio of rate ceilings       = (b/32) / (b/256) = 8      EXACTLY
```

Check five — the budget upgrade.

```
log2(10 yr / 1 yr)  + log2(100/10) = 3.321928 + 3.321928 =  6.643856
log2(10 yr / 1 mo)  + log2(100/3)  = 6.906891 + 5.058894 = 11.965785
difference                                                =  5.321928  bits
                                                          = 2^5.321928 = 40.0 in likelihood ratio
```

Five checks, none longer than four lines, and between them they establish the floor, the constant, the rate ceiling, the cliff and the exchange rate. A reviewer who runs them has verified everything this paper contributes that is not quoted from elsewhere.

32.9.5 9.5 Limits, asymptotics, and the domain of validity

$T \to 0$ **in the Gabor limit.** $\Delta f \to \infty$. **The limit is not a pathology but the content of the theorem:** an instantaneous event has no frequency. The domain of the tabulated relation is $T > 0$ with a finite-energy envelope.

$T \to \infty$**.** $\Delta f \to 0$. **Arbitrarily fine frequency resolution is purchasable with arbitrarily long observation, which is why atomic clocks work and why nothing about this bound forbids precision — it forbids *simultaneous* precision.**

Envelope $\to$ **rectangular.** Δf does not converge. **The limit does not exist**, and the correct report is that the quantity is undefined rather than large. This is a real constraint on the model's domain and is the reason REQ-5.3 prohibits rectangular gating.

$n \to \infty$ **in the lattice.** $r(n) = (2n + 1)/n^2 \to 0$ and the required cycle count grows without bound as $n/(8\pi)$. **The deep lattice is unrenderable in the limit, which is a statement about the square law and not about the instrument.**

$N \to \infty$ **in the folded lattice.** The minimum gap tends to zero, since squares become arbitrarily close under octave reduction. **The cliffs are the visible structure of that collapse and there are infinitely many of them**, at every N introducing a new closest pair; 17 and 41 are the first two.

$B \to 0$ **in the budget bound.** $\Delta t\, \Delta s \geq TS$: the forecast says nothing beyond its stated range. **Correct, and the right behaviour.**

$B \to \infty$**.** The bound goes to zero and permits perfect forecasting — **which the horizon law forbids independently.** The two bounds are independent and both bind, and neither subsumes the other. **A model whose limits are unexamined has not been understood, and this is the one limit in the paper where the answer is supplied by a different bound rather than by the model itself.**

$\lambda \to 0$ **in the horizon law.** $t_{\text{pred}} \to \infty$: a non-chaotic system has no error-doubling horizon, which is why tides and ephemerides are predictable for centuries. **The law's domain is** $\lambda > 0$ **and it is silent, not permissive, outside it.**

$\sigma_f \to 0$ **in the rhythm form.** The timing bound goes to infinity, so **a perfectly sharp period claim places no bound on timing at all.** That is the sharpest single statement in Section 6 and it is a limit result.

32.9.6 9.6 Invariance, and how far each claim has been tested

Scale invariance of the Gabor product is exact and algebraic: replacing $g(t)$ by $g(t/s)$ multiplies Δt by s and divides Δf by s. **The transformation group is the one-parameter dilation group of the time axis, the invariant is** Tf_0**, and it is verified numerically at five Gaussian widths spanning a factor of twenty (5 ms to 100 ms), agreeing to ten significant figures at every one. Five widths is not many**, but the claim is a theorem and the numerical check is a consistency test on the implementation rather than evidence for the theorem.

Octave invariance of the lattice cycle count is exact and is verified at the four base pitches 55, 110, 220 and 440 Hz, where the render-rate ceiling doubles at each octave exactly (21.60, 43.20, 86.39, 172.79 symbols per second). **Four octaves tested.** It would break if the folding rule changed the odd-part classes, which it does not.

Transposition invariance of the fractional gap — that $r(n)$ does not depend on j — is algebraic and is verified at $j = 6$, 7 and 8. **Three transpositions.**

Invariance of the budget bound under a change of units on either axis holds because the bound is a statement about ratios: rescaling T and Δt by the same factor leaves $\log_2(T/\Delta t)$ unchanged. **This is exact and untested empirically, because it is definitional.**

No invariance in this paper has been tested on data from any real system, because no data were collected. Every invariance claim above is algebraic, and the numerical checks are checks of the arithmetic.

32.9.7 9.7 Significant figures and their warrant

Following RIGOR_STANDARD §2, each class of number in this paper is quoted at a precision matching its status.

- **Exact by definition, quoted to whatever precision is useful:** $1/(4\pi) = 0.079577471545948$; $4\pi = 12.566370614359172$; $k = 0.483425732988414$; $\varphi^2 = 2.618033988749895$; the folded ratios $9/8$, $50/49$, $289/288$, $1682/1681$ and their cents values; $\log_2$ costs such as 5.321928 bits. **These are exact and are labelled exact wherever they appear at full precision.**
- **Computed from a stated function at stated parameters, and rounded honestly:** the envelope table (0.079577, 0.081665, 0.087211), the render times, the cycle counts, the ensemble entropies. **Ten significant figures on the Gaussian's** 0.0795774715 **is a statement about a numerical agreement with an exact constant, not about a measurement.**
- **Quoted from published physics, carrying that literature's uncertainty and not this paper's:** $\hbar = 1.054571817646 \times 10^{-34}$ J s (exact given the 2019 SI h); the Bohr radius 5.291772×10^{-11} m; -13.605693 eV; the atomic lifetimes and linewidths of Section 4.9. **None of these is measured by this paper.**
- **Measured psychophysical quantities with their own sampling error, restated here without it:** the frequency difference limens of 0.2 to 0.6 per cent in Section 5.9. **The Gabor columns beside them are exact; the limen column is not, and no interval is available for it in the form quoted.** That is a defect of the presentation and is recorded in Limitations.
- **Assumed parameters in an illustrative model, rounded hard:** the error-doubling times of Section 6.11 (1.0, 1.5, 2.0 and 3.0 days) and the scenario weights of Section 6.7. **Nothing in either is measured, and both figures say so in their captions.**

32.9.8 9.8 Analogy discipline, restated at the point of verification

Every borrowing in this paper carries the label assigned in Section 3.2 and no passage upgrades one. The entropies are **Shannon** entropies in bits [Shannon, 1948]; no thermodynamic entropy appears, and the word is never used without the qualifier. The Gabor limit is a theorem about square-integrable functions and its application to a rendered tone is a **derivation** because the premises are satisfied by a finite-energy signal under a smooth envelope. The uncertainty relation is a fact about **state preparation**, not about measurement disturbance, in every passage of Section 4. The variational recovery of the Bohr radius is a **stationarity** condition and **licenses no teleology**. And the extension of conjugate structure to foresight is an **analogy**, badged in every subsection that states it, contributing **no line of the chain in Section 9.3**.

32.10 10 Limitations

L-1. The evidence budget was never estimated, so the paper's second headline result is arithmetic without a measurement. $\Delta t \, \Delta s \geq TS/2^B$ is exact and its right-hand side contains B, which no part of this paper obtains from data. H4 is recorded as NOT TESTED. **Until a forecast record with stated** T**,** S**,** Δt **and** Δs **per forecast is scored and** B **estimated from realised out-of-sample discrimination, the bound constrains nothing empirically.** That is the single largest gap in the paper and it is not repairable by argument.

L-2. Six of the physics and forecasting primary sources are named without citation keys. Robertson (1929), Weyl (1928), Schrödinger (1930), Ozawa (2003), Busch–Lahti–Werner (2013), Arndt et al. (1999), Erhart et al. (2012), Rozema et al. (2012), Lorenz (1963, 1969), Zhang et al. (2019), Brier (1950), Murphy (1973), Carlson (1985), Mandelstam and Tamm (1945) and Wier, Jesteadt and Green (1977) have **no entries in this corpus's bibliography**. Each is named in the text with journal, volume, year and pages, and **no key was invented for any of them.** A reader can locate every one; the corpus's bibliography is nonetheless incomplete on exactly

the literature this paper leans on hardest, and that is a defect of the corpus rather than of the record.

L-3. The corpus's own frequency framework is a postulate and this paper does not test it. Every hertz figure in Section 5.10 and every millisecond figure in Section 5.11 inherits P6, whose value is a definition and whose expression in hertz has **no external support of any kind**. The Gabor arithmetic is exact and the object it operates on is chosen. **A reader who rejects the essence law keeps the cents and loses the hertz**, and the paper is arranged so that the separation is visible.

L-4. The dimension of the specification axis is undetermined and the source is silent. Section 9.1 records this rather than assigning a dimension. Where the outcome space is categorical the bound is a counting bound and is valid; where it is numeric, the quotient $S/\Delta s$ requires a ratio scale that this paper does not establish. **The conjoint-measurement test that would settle whether such an outcome scale is quantitative at all [Luce and Tukey, 1964] has not been run**, and Michell's objection [Michell, 1997] therefore stands unanswered against any numeric reading of that axis.

L-5. No reflexive claim in this paper is identified. The correspondence of Section 6.14 and the postulate referenced in Section 6.4 both sit inside the framework's reflexive material, whose endogenous effect is **not identified from observational data** [Manski, 1993]. **No remedy is adopted** — not randomisation, not instrumentation, not partial-population variation, and not the stipulated-reference-group route [Brock and Durlauf, 2001]. **These are modelling positions, and this paper does not upgrade them.**

L-6. The Lucas critique is conceded and not answered. B is not invariant under the publication of a forecast built with it [Lucas, 1976], and a forecaster's measured discrimination becomes a target once it is used to evaluate them [Goodhart, 1975, Campbell, 1979]. **The paper offers no defence of parameter stability and states the bound as a constraint on a forecaster at a moment.**

L-7. The invariance claims are tested at few points. Scale invariance of the Gabor product is checked at five Gaussian widths, octave invariance at four base pitches, transposition invariance at three values of j. **All three are algebraic identities, so the checks test the implementation and not the claim** — but a reader wanting empirical breadth should know the counts are 5, 4 and 3, and that none of them is on data from a physical instrument.

L-8. The engineering shorthand $\Delta t = T$ is optimistic by construction. The tables treat the whole rendering window as the rms spread, which overstates achievable resolution by up to $\sqrt{12} = 3.46$ for a rectangular gate. **It is retained for consistency with the framework's existing design tables and is stated in the open at Section 5.5**; every tabulated $\Delta f_{\min}$ should be read as a floor computed under a favourable convention.

L-9. The difference limens are quoted without intervals. Section 5.9's 0.2 to 0.6 per cent figures are measured psychophysical quantities from a 1977 study and are reproduced as point values because that is how the source framework carries them. **A measured quantity quoted without an uncertainty is a presentational defect**, and it is recorded here rather than dressed up.

L-10. Nothing here says a listener can hear any of it. Section 5.12's claim that the rendering distinguishes compositions to about a tenth of a member is a statement about the *signal*. **Whether anything at the far end extracts it is a separate question**, and for the biological case there is no established demodulating mechanism at all [Curzi, 2026y].

L-11. No AIC or BIC is reported, here or anywhere in this corpus. Section 7.9 prints the row as NOT MET rather than omitting it. Nothing in this paper is fitted, so no information criterion is computable from it; **the corpus-wide absence is nonetheless a genuine gap and is recorded as one.**

L-12. The three worked reviews in the appendices are constructed cases, not field reports. They are built to exercise the chain and their inputs are stipulated. **No design review, forecast or lattice implementation described in them took place.**

32.11 11 Conclusion

Three reciprocal relations circulate under one name and this paper separates them, at three different levels of warrant, with the arithmetic of each printed.

Result One, the Gabor limit and its cost schedule. $\Delta t\,\Delta f \geq 1/(4\pi) = 0.079577471545948$, **proved by the same variational argument as** $\sigma_x\sigma_p \geq \hbar/2$ **with the substitution** $p = \hbar k$ **deleted — one line, and it is the only line that quantum mechanics contributes.** Gaussians saturate the bound, verified to ten significant figures at five widths from 5 ms to 100 ms; **a Hann window costs 2.6 per cent over the floor, a triangular window 9.6 per cent, and a rectangular gate has no finite spectral second moment at all**, its computed product climbing through 1.712, 2.420 and 3.422 as the analysis bandwidth is extended. The DFT bin $1/T$ sits exactly $4\pi = 12.566$ times above the floor and a Hann main lobe $4/T$ exactly $16\pi = 50.265$ times above it, **and that factor of fifty is the single largest source of optimism in signal design.**

Result Two, the information-budget bound. $\Delta t\,\Delta s \geq TS/2^B$, **exact, classical, and derived in two lines from counting: each additional bit halves the product, and each bit spent on timing is a bit not spent on specification.** Sharpening a forecast from a year and ten classes to a month and three classes costs **5.3219 bits**, a factor of forty in the likelihood ratio. **The bound has never been estimated against a forecast record and is, as of this paper, arithmetic without a measurement.**

Result Three, the badge and the refusal. The identification of a foreseen future's uncertainty with quantum conjugacy is an analogy, badged as one in every subsection that states it. It is not claimed to follow from quantum mechanics; no $\hbar$ appears in it; and no quantum result is taken to license any claim about belief, mind or future events. **The analogy contributes no number to this paper** — every figure in Section 6 comes from counting, from Fourier duality or from dynamical systems, and all of them survive the analogy's deletion.

Result Four, the horizon law. $t_{\text{pred}} = t_d \log_2(D_{\text{tol}}/D_0)$: **prediction horizon grows only as the logarithm of the improvement in initial knowledge**, and at an error-doubling time of 1.5 days a thousandfold improvement in the initial condition buys 14.95 days. **The law binds only where a positive Lyapunov exponent has been established; invoking it elsewhere is an analogy and is named as one.**

Result Five, the lattice ceilings. At the common transposition $j = 6$ the gap between the first two lattice indices is **92.81774 Hz**, giving $T_{\min} = 1.7147$ ms and $R_{\max} = 583.19$ symbols per second under the half-gap criterion. **Folded to one octave at a 220 Hz base the same lattice falls to 86.394 symbols per second at** $N = 12$ **and to 10.799 at** $N = 24$ — 6.8 and 54.0 times slower. **The audible band admits at most 25 lattice elements under a single integer octave transposition, and only at** $j = 6$**.**

Result Six, the corrected cliffs. The folding cliffs are at $N = 17$ **and** $N = 41$**, not the previously published** $N = 20$ **and** $N = 48$**.** On exact rational arithmetic the minimum adjacent interval is $50/49 = 34.9756$ cents up to $N = 16$, $289/288 = 6.0008$ cents from $N = 17$, and $1682/1681 = 1.0296$ cents from $N = 41$. **The cause is that** $17^2 = 289$ **folds beside** $288 = 2 \cdot 12^2$**; the earlier boundaries were the next sampled rows of a table rather than computed thresholds. The rate penalty of exactly 8 across the first cliff — and exactly 4 across the second — stands unchanged.**

And the counter-result, which is what the paper takes away. It corrects its own published figures in print, and it refuses quantum licence outright. The forecast analogy carries no quantum content, derives entirely from classical Fourier duality and counting, and is defended only by whether its three quantitative forms measure anything — of which, at present, **one is exact and unmeasured, one is exact and requires a badged modelling convention, and one requires an exponent nobody has estimated for any system this framework describes. The classical limit is the only uncertainty relation in this work that constrains anything a reader will build; it is a theorem, it is verifiable on an oscilloscope, and in Appendix B it changes a design by a factor of thirteen.**

32.11.1 Declarations

Data and code availability

No external dataset was collected for this paper, and no human subjects were studied. The numerical values reported as computed — the saturation of the Gabor bound by Gaussians of five widths, the divergence of the rectangular window's second moment at three analysis bandwidths, the envelope multiples, the lattice tables, the folded-set minima on exact rational arithmetic, the ensemble entropies and the budget arithmetic — are evaluations of the functions stated beside them, reproducible from the specifications in the text. The physical constants, the Bohr radius, the hydrogen ground-state energy, the atomic lifetimes and linewidths, and the frequency difference limens are **quoted from the published literature named in Sections 4 and 5 and are not measurements of this paper**; several of those sources carry no key in this corpus's bibliography and are named in full in the text, as disclosed at Section 4.3 and Limitation L-2. **The figures are generated by the scripts named in each caption**, and those scripts contain no data beyond the numbers printed in this paper and the functions stated in Sections 5 and 6.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest. The essence law on which Sections 5.10 to 5.12 operate is the author's own postulate, and this paper's arithmetic is applied to it by its author; that is a structural conflict which no declaration removes, and it is the reason the postulate's badge is carried in every subsection that uses it and the reason Section 9.3 separates the results that consume it from those that do not.

Author contributions

Single author. Michael Laurence Curzi derived the results, performed the numerical evaluations, produced the figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. Section 5.9 discusses published psychophysical measurements of frequency discrimination; those are other investigators' data, quoted, and no listening test was conducted for this paper. Sections 6 and Appendix C describe a forecasting practice as a design and are not a report of any study performed.

Preregistration

This paper reports no preregistered test. The four hypotheses of Section 2 were fixed before the numerical work reported in Sections 5 and 6 was carried out, and are printed in the form in which they were fixed. **H1, H2 and H3 are confirmed by exact arithmetic; H4 is reported as NOT TESTED, because no forecast record was collected and no evidence budget was estimated from data.** The falsification criterion for H2 — that an exact-rational recomputation return the previously published thresholds — was declared before the recomputation was run and did not fire; the recomputation instead moved the thresholds, which is reported as Result Six. The corpus's preregistered test of the storage law against a real system, which **failed** its second

falsification rule, is reported elsewhere [Curzi, 2026ad], and any applied claim resting on this paper inherits that result.

32.12 Appendix A · A rendering proposal reviewed against the four relations

A constructed review case, built to exercise Sections 4 and 5 end to end. No such review took place. The proposal under review reads:

> *Because the uncertainty principle is a quantum effect and the acoustic signal is classical, the only limits on rendering rate are the sample rate and the transducer bandwidth. We can therefore render each tone for 2 ms and step through a 24-element lattice in 48 ms.*

Step 1 — identify which relation is invoked and which is dismissed. The proposal invokes and dismisses "the uncertainty principle" without naming which; **REQ-4.8 violated.** Enumerated:

```
Kennard 1927     sigma_x sigma_p >= hbar/2         QUANTUM. Not relevant. Rightly dismissed.
Robertson 1929   sigma_A sigma_B >= |<[A,B]>|/2    QUANTUM. Not relevant. Rightly dismissed.
Ozawa 2003       error–disturbance                 QUANTUM. Not relevant. Rightly dismissed.
Gabor 1946       dt df >= 1/(4 pi)                 CLASSICAL. BINDING. Wrongly dismissed.
```

Three of the four dismissals are correct and the fourth governs the design. This is FM-4.10 in its natural habitat: the reasoning about quantum mechanics is sound and it led to exactly the wrong engineering conclusion, because the classical case was assumed to be a weaker version of the quantum case rather than its parent.

Step 2 — establish that the classical relation exists and is not quantum. Run Section 4.5's derivation with $p = \hbar k$ deleted; what remains is $\sigma_t \sigma_\omega \geq 1/2$ with ω in radians per second, hence $\sigma_t \sigma_f \geq 1/(4\pi) = 0.079577471545948$. **No $\hbar$, no quantisation, no operators, no measurement.** Confirmed numerically on a Gaussian envelope of standard deviation 50 ms, 262,144 samples at 0.2 ms spacing:

```
sigma_t = 0.05000000 s
sigma_f = 1.59154943 Hz
product = 0.0795774715
1/(4 pi) = 0.0795774715
```

Equality to ten significant figures. The relation is neither approximate nor quantum.

Step 3 — apply it to the proposed 2 ms tone. $\Delta f_{\min} = 1/(4\pi \times 0.002) = 39.78874$ Hz at the absolute floor. A real analysis window does worse: a Hann window of the same duration has main-lobe width $4/T = 2000$ Hz, which is 50.265 times the floor, and the DFT bin spacing $1/T$ is 500 Hz. **The honest achievable resolution at 2 ms is between 40 Hz and 500 Hz depending on the analysis chain, and 39.79 Hz is a bound no processing can beat.**

Step 4 — compare against what the lattice requires. For a 24-element lattice under a common octave transposition, the audible-band constraint forces $j = 6$, giving $k \cdot 2^6 = 30.9392$ Hz at $n = 1$ and 19337.03 Hz at $n = 25$. The tightest absolute spacing is at the bottom:

```
gap(1 -> 2) = k (2*1 + 1) 2^6 = 3 x 0.483425732988414 x 64 = 92.81774 Hz
```

against the 2 ms floor of 39.78874 Hz. **The bottom of the lattice survives — barely, and only with a Gaussian envelope and optimal analysis.** The proposal's own success criterion was that the elements be *told apart*, which requires the spread comfortably below the gap:

```
required df <= gap/2 = 46.4089 Hz
T_min       = 1 / (4 pi * 46.4089) = 1.7147 ms
R_max       = 583.19 symbols per second
```

At 2 ms the design passes the half-gap criterion at the bottom of the lattice with a margin of about 14 per cent. With a Hann-windowed analysis at 50.265 times the floor it fails by more than an order of magnitude.

Step 5 — the findings.

```
FINDING 1 (REQ-4.8, FM-4.10). The dismissal of "the uncertainty principle" is correct for
three of the four relations in circulation and wrong for the fourth. The Gabor limit is
classical, applies to this signal, and is binding.

FINDING 2 (REQ-4.5, REQ-5.3). The relation is saturated only by a Gaussian envelope. Any
real window sits above the floor: Gaussian 1.000x, Hann 1.026x, triangular 1.096x.
Rectangular is disqualified -- its spectral second moment diverges, so it has no finite
sigma_f at all. The specification must name the envelope.

FINDING 3 (REQ-5.12). No magnitude was computed anywhere in the proposal. Compute first.

DISPOSITION: 2 ms per symbol is retained ONLY with a Gaussian or Hann envelope and ONLY
for a common-transposition lattice at j = 6 or above. It is rejected for any folded
(per-value octave) lattice, where Section 5.11 shows the minimum gap collapses to 50/49
= 34.98 cents and the required tone length rises to 11.57 ms at a 220 Hz base -- a factor
of 5.8 slower than proposed.
```

What the case demonstrates. The physics was not wrong: quantum mechanics really is irrelevant to a loudspeaker. **The error was structural — treating the classical Fourier relation as a diluted version of the quantum one rather than as the theorem of which the quantum one is a special case.** A reader who has been shown, with the references, that $\sigma_x \sigma_p \geq \hbar/2$ is a Fourier theorem with $\hbar$ attached will not make the mistake; a reader who has been told the photon-kick story will make it every time, **because the photon-kick story has no classical limit — it is a story about quanta, and if it is believed the relation appears to go away when the quanta do. It does not go away. It was never about the quanta.**

32.13 Appendix B · A twelve-element folded lattice, priced

A constructed design review, built to exercise Section 5. No such implementation exists. The brief:

```
lattice size N          12 elements  (six pitch classes)
rendering               folded to one octave, base A3 = 220 Hz
sequence                the twelve indices in order, then repeat
required                a listener or an analyser must be able to identify which
                        pitch class is sounding at any moment
proposed rate           40 symbols per second (25 ms each)
envelope                unspecified
```

Step 1 — compute the folded set. Odd parts of 1 to 12 are 1, 1, 3, 1, 5, 3, 7, 1, 9, 5, 11, 3, giving classes $\{1, 3, 5, 7, 9, 11\}$ — six, as $\lceil 12/2 \rceil$ requires. Squares folded into $[1, 2)$ and multiplied by the 220 Hz base:

```
class      folded ratio      frequency (Hz)     cents above base
-------    ------------      --------------     ----------------
```

```
 1          1/1            220.0000             0.000
 3          9/8            247.5000           203.910
 9         81/64           278.4375           407.820
 7         49/32           336.8750           737.652
 5         25/16           343.7500           772.627
11        121/64           415.9375          1102.636
```

Every value carries the badge: these are essence frequencies, not physical resonances, and k is a chosen reference constant [Curzi, 2026].

Step 2 — find the binding gap.

```
220.0000 -> 247.5000     27.5000 Hz
247.5000 -> 278.4375     30.9375 Hz
278.4375 -> 336.8750     58.4375 Hz
336.8750 -> 343.7500      6.8750 Hz    <-- BINDING
343.7500 -> 415.9375     72.1875 Hz
```

The binding pair is class 7 at 336.8750 Hz and class 5 at 343.7500 Hz, separated by 6.8750 Hz — the interval $50/49$, 34.976 cents. Every other gap is at least four times larger. A design that hits this pair works; a design that misses it fails on exactly two of twelve symbols, which is the kind of failure that survives casual testing.

Step 3 — the full-gap criterion. $\Delta f \leq 6.8750$ Hz gives $T_{\min} = 11.5749$ ms and $R_{\max} = 86.394$ symbols per second. **The proposed 25 ms passes with a factor of 2.16 in hand.** Provisionally acceptable.

Step 4 — the half-gap criterion, which is what "identify" requires. A spread merely smaller than the gap leaves two adjacent peaks touching. $\Delta f \leq 3.4375$ Hz gives $T_{\min} = 23.1498$ ms and $R_{\max} = 43.197$ symbols per second. **The proposed 25 ms passes by 8 per cent — a thin margin, and it is the margin against the theoretical floor, which no real envelope reaches.**

Step 5 — the envelope factor. The brief left the envelope unspecified; **REQ-5.3 violated.**

```
envelope         multiple of floor     effective df at 25 ms      passes half-gap (3.4375 Hz)?
------------     -----------------     ---------------------      ----------------------------
Gaussian               1.000                3.1831 Hz              YES, by 8%
Hann                   1.026                3.2659 Hz              YES, by 5%
rectangular          unbounded              no finite value        NO -- disqualified outright
```

But the analysis chain was not specified either. **If the receiving end uses a Hann-windowed FFT rather than a matched filter, the practical resolution is the main-lobe width $4/T = 160$ Hz at $T = 25$ ms. 160 Hz against a 6.875 Hz gap: the design fails by a factor of 23 in the analysis chain even though the signal itself is fine.** That is FM-5.4, and it is the most expensive failure in the case because the transmitted signal is correct and the fault is entirely at the receiver.

Step 6 — the corrections.

```
CORRECTION 1 (REQ-5.3). Specify a Hann envelope, 25 ms, 5 ms raised-cosine ramps.

CORRECTION 2 (REQ-5.2). Specify the analysis method. An FFT-based analyser cannot do this
job at 25 ms. Two options:
    (a) matched filtering, or a bank of six narrowband detectors tuned to the six known
        frequencies -- since the six values are KNOWN in advance, detection is far easier
        than blind estimation, and the Gabor floor rather than the FFT lobe applies;
    (b) lengthen the symbol to 232 ms so that 4/T = 17.2 Hz sits below half the 6.875 Hz
        gap -- which drops the rate to 4.3 symbols per second, a factor of nine.

CORRECTION 3 (REQ-5.7). Publish the ceiling with the specification:
```

```
    full-gap  R_max = 86.394 symbols/s
    half-gap  R_max = 43.197 symbols/s
    with FFT analysis at 4/T, R_max = 4.3 symbols/s

CORRECTION 4 (FM-5.12). Publish the frequencies with their achievable precision.
    At 25 ms with a Hann envelope, df = 3.27 Hz.
    class 7 renders as 336.9 Hz +/- 3.3 Hz  = +/- 17 cents
    class 5 renders as 343.8 Hz +/- 3.3 Hz  = +/- 17 cents
    The two 17-cent bands are separated by 35 cents. They do not overlap. They very
    nearly touch.
```

Step 7 — the alternative the arithmetic exposes. The binding constraint is one pair out of fifteen, and nothing about the square law can change that $25/16$ and $49/32$ differ by $50/49$. **But the lattice does not have to be folded.** Under a common transposition at $j = 6$ the same twelve indices run from 30.94 Hz to 4455.25 Hz, the binding gap moves to the bottom pair at 92.82 Hz, and the half-gap criterion gives $\Delta f \leq 46.41$ Hz, $T_{\min} = 1.71$ ms, $R_{\max} = 583.19$ symbols per second. **A factor of 13.5 faster than the folded version at the half-gap criterion, and 135 times faster than the FFT-limited folded version.** The cost is a rendering spanning 7.2 octaves instead of one, which may be unacceptable for the brief. **The trade is now quantitative rather than aesthetic:** folded to one octave at 220 Hz, 43.2 symbols per second and musically coherent; common transposition at $j = 6$, 583.2 symbols per second across 7.2 octaves.

What the case demonstrates. First, the constraint is real and binds on a specific pair — not "somewhere in the design" but on classes 7 and 5, at 336.875 and 343.750 Hz, 6.875 Hz apart, because $50/49$. **Second, it is not where anyone would have looked**: the weakest point was neither the transmitter nor the rate nor the encoding but the analyser's window function, sitting 50.265 times above a floor nobody had computed. **Third, this is the only uncertainty relation in the paper that changed a design.**

32.14 Appendix C · A forecast, priced in bits

A constructed case, built to exercise Section 6. No such forecast was made or scored. A programme director is asked, in March, when the work will ship and what exactly will be in it. The programme has a twelve-month planning horizon and eight recognised scope configurations.

Step 1 — state T and S (REQ-6.4).

```
T = 52 weeks         the planning horizon over which the question is defined
S =  8 classes       the recognised scope configurations
T . S = 416          cells in the (timing, specification) plane
```

Nothing can be evaluated until this is on paper, and FM-6.3 is the default state of the conversation.

Step 2 — the evidence budget (REQ-6.6). Not asserted; **stipulated for this constructed case at six bits**, which would in practice be estimated from a record of prior programmes scored with Murphy's decomposition on ship-quarter and scope-class outcomes. **Six bits is a factor of 64 reduction in the cell space, from 416 candidate cells to an effective 6.5** — a genuinely useful and modest forecaster.

Step 3 — the achievable frontier. $\Delta t \, \Delta s \geq 416/64 = 6.5$ week-classes.

```
  B bits      dt.ds floor      if ds = 1 class -> dt      if dt = 1 week -> ds
  ------      -----------      ---------------------      --------------------
    2          104.00000            104.00 weeks              104.00 classes
    3           52.00000             52.00 weeks               52.00 classes
```

```
 4          26.00000            26.00 weeks            26.00 classes
 5          13.00000            13.00 weeks            13.00 classes
 6           6.50000             6.50 weeks             6.50 classes
 7           3.25000             3.25 weeks             3.25 classes
 8           1.62500             1.62 weeks             1.62 classes
 9           0.81250             0.81 weeks             0.81 classes
10           0.40625             0.41 weeks             0.41 classes
```

At $B = 6$ exactly one of these is honestly available: **Option A**, one scope class within a 6.5-week window; **Option B**, one week with a scope anywhere across 6.5 of the eight classes; or **Option C**, the symmetric allocation $\Delta t/T = \Delta s/S = 2^{-3} = 0.125$, which for these ranges coincides with Option A since $T/S = 6.5$. **Option B is worthless and Option A is what the questioner wanted; the value of the computation is that this was settled by arithmetic rather than by argument.**

Step 4 — price the escalation (REQ-6.5). The request is then for the exact week and the exact scope:

```
bits required for (dt = 1 wk, ds = 1 class) = log2(52) + log2(8) = 5.7004 + 3.0000 = 8.7004
bits currently held                                                              = 6.0000
SHORTFALL                                                                        = 2.7004 bits
```

The request costs 2.70 bits more than the director has — a likelihood-ratio factor of 6.5. FM-6.5 exactly, and now with a price tag instead of a personality conflict. The correct response is not "I will try harder" but a statement of what would supply the missing bits and what each observation is historically worth. **That is a specification, not an excuse, and it converts an argument into a requirement.**

Step 5 — check the horizon bound independently (REQ-6.8). **Honest answer: the horizon law is not shown to apply and is therefore not invoked.** Software delivery is not a chaotic dynamical system in the technical sense and no Lyapunov exponent has been estimated for it; **invoking the butterfly effect here without that estimate would be FM-6.6.** What the programme has is a small number of high-variance dependencies whose resolution dates dominate the outcome — a mixture, not a chaotic flow. **A paper supplying three bounds must also supply the discipline to use only the ones that apply.**

Step 6 — publish the ensemble, not the verdict (REQ-6.12).

```
scenario                weight     ship week (from March 1)
--------------------    ------     ------------------------
mid-tier, on time        0.40              week 28
mid-tier, slipped        0.25              week 36
reduced scope, early     0.20              week 22
full scope, late         0.15              week 44

Shannon entropy H                 = 1.9037 bits
effective scenario count 2^H      = 3.7417   (of a nominal 4)
weighted mean ship week           = 31.20
sd of ship week                   =  7.17 weeks
```

The effective scenario count of 3.7417 is the Hill number of order one — the exponential of a Shannon entropy [Hill, 1973, Jost, 2006] — **and it says the four scenarios are nearly independent, the ensemble carrying almost its full nominal variety. Collapsing this to "week 31, mid-tier" would discard 1.90 bits.** The verdict is a smaller and less honest object wearing the answer's clothes.

Step 7 — the headroom reading, badged as a correspondence of model form and **not** a claim about physics (Section 6.14):

```
ceiling      T.S = 416 cells          <->    ln N, the hard ceiling
budget       B = 6 bits               <->    h, the headroom
requested    r = 2.70 bits            <->    r_t, the requested increment
verdict      r > h                    <->    h_t < 0 for this request
```

“No choice of interaction geometry can prevent order from decreasing” becomes “no choice of method, model, or expert can supply bits that were never observed.” Same inequality, different scope, and a claim about model shape stated as one [Curzi, 2026a].

What the case demonstrates. First, the section’s usable content survives the removal of the analogy entirely — every number came from counting bits, and Heisenberg appears nowhere in it. **Second, the trade is negotiable once it is priced. Third, the discipline of stating T and S does most of the work: the bound is trivial mathematics, and the difficulty was never the algebra but that nobody wrote down the ranges.**

Chapter 33

Quantum Formalism Without Licence: Born-Rule Statistics, Decoherence Timescales in Warm Tissue, and the Casimir Force at Biological Scale

Abstract. This paper states the quantum formalism as established physics and then computes what it forbids. Unitarity, linearity, the Born rule $P(a_n) = |\langle a_n|\psi\rangle|^2$, Gleason's theorem, no-signalling, no-cloning, Bell/CHSH with its quantum bound $2\sqrt{2} = 2.828427$, and Kochen–Specker are given as proved results that bound every interpretation; the measurement problem is conceded to be genuinely open, but open in no direction that licenses influence, since no-signalling holds in every interpretation. The paper advances no postulate of its own there. Decoherence is then computed rather than invoked: with $n = 3.3428 \times 10^{28}$ per cubic metre and a mean thermal speed of 603.6 m s^{-1} at 310 K, the microtubule-scale estimate gives $\tau_{\text{dec}} = 1.98 \times 10^{-17}$ s against a 1 ms neural event, a shortfall of 5.04×10^{13}; even a hypothetical 10^6 protection factor leaves seven orders of magnitude unexplained. Decoherence suppresses off-diagonal terms and provably changes no diagonal element, so it does not solve the measurement problem. The Casimir pressure $\pi^2\hbar c/(240a^4) = 1.300126 \times 10^{-27}/a^4$ Pa is real, is measured to about one per cent, and is negligible where it is usually invoked: on a 7 m room wall the correct high-temperature value gives a total force of 1.996×10^{-23} N. Only the mode-selection principle transfers, and that transfer is badged as a postulate.

Keywords: Born rule; decoherence; no-signalling; Casimir effect; room modes; boundary geometry

Series: B. Engineering the Surplus **Paper:** CCP-033

33.1 1 Introduction

A framework that borrows the vocabulary of one science to describe another owes that science an accounting. This paper is the accounting for the largest borrowing in the Creature Comforts corpus: quantum mechanics. The corpus uses the word *amplitude* for a branch weight, the word *conjugate* for a resolution trade-off, and the word *coherence* for a state that

degrades on contact with its surround. Each of those uses is defensible only if the original is stated correctly and the transfer is priced. **The purpose of this paper is to state the original at full strength and then to compute what it forbids.**

The method is deliberately asymmetric. Section 4 advances no claim of the author's at all: it states the standard formalism, the four main interpretive positions with what each abandons, and the theorems that bound every one of them — no-signalling, no-cloning, Gleason, Bell/CHSH with its quantum bound of 2.828427, and Kochen–Specker. Sections 5 and 6 then compute two numbers that constrain the rest of the corpus rather than supporting it. The first is a decoherence timescale for a warm wet system: with a water number density of `3.3428e28` per cubic metre and a mean thermal speed of 603.6 metres per second at 310 K, a microtubule-scale object has an estimated coherence time of `1.98e-17` seconds against a neural event lasting `1e-3` seconds, **a shortfall of `5.04e13`**. The second is the Casimir pressure at architectural scale: the correct high-temperature value across a seven-metre room is `1.155e-24` pascals, giving a total force of `1.996e-23` newtons over one wall.

Both numbers are hostile to the material that surrounds them, and that is why they are here. A corpus that wanted quantum licence would not compute them. **The results this paper contributes are, with one exception, results about what cannot be claimed**, and the exception — that boundary geometry alone selects field modes at no ongoing cost — is exact for electromagnetic and acoustic modes and is badged as a design heuristic everywhere else.

Three things distinguish this treatment from the general run of writing that connects quantum mechanics to mind. First, the interpretive survey is organised by which of three inconsistent premises each position abandons, so that the map from interpretation to cost is one-to-one and cherry-picking becomes visible. Second, the decoherence estimate is checked against a published experiment before it is used against a proposal, so the reader can see that the estimator lands on the right decade where the answer is known. Third, the Casimir section computes its own most attractive analogy into the ground: the force is `1e-23` newtons at room scale and, at cellular scale with realistic materials, `6.7e7` times smaller than the weight of one bacterium is than the force between two room walls. What survives that computation is a statement about mode selection, not a statement about force.

The framework's conjugate-uncertainty material, which is where the quantum vocabulary enters the corpus, is presupposed here rather than re-derived: it is established, with its explicit refusal that quantum mechanics licenses anything about belief, mind, or future events, in Curzi [2026z]. The corpus's classification of its own claims into established mathematics, established psychology and declared postulates — the discipline this paper applies to a physical literature rather than a psychological one — is set out in Curzi [2026].

Section 2 states the hypotheses. Section 3 states the method, including the analogy-discipline rule that governs every transfer in the paper. Sections 4, 5 and 6 develop the formalism, the decoherence computation and the Casimir computation. Section 7 collects the failure modes as a register. Section 8 works three audits end to end. Sections 9 to 12 discuss what follows, verify the chain, state the limitations and conclude.

33.2 2 Hypotheses

This paper is unusual in the corpus in that most of its content is a report of established physics rather than a claim of its own. The hypotheses below are therefore few, and the section states explicitly which of the paper's assertions are *not* hypotheses and why.

> **H1 — The warm-wet decoherence gap.** In an unrefrigerated, unengineered biological system at 310 K, no charge-displacing degree of freedom maintains phase coherence for a time approaching any functionally relevant neural timescale. **Predicted:** an estimated coherence time of order `1e-17` to `1e-12` seconds for objects between an ion and a neuron soma, against functional timescales of `1e-4` to `1e-1` seconds, so a ratio between `1e9` and `1e16`; and, as the sanity check, an estimate no

larger than a few multiples of the thermal time `hbar/kT = 2.464e−14` s at that temperature. **Falsified by:** a replicated demonstration of functional quantum coherence persisting above `1e−6` seconds in an unrefrigerated, unengineered biological system at 310 K. That threshold is stated in advance and is three orders of magnitude more generous than the `1e−3` seconds any cognitive proposal actually requires.

H2 — Prefactor independence. The conclusion of H1 does not depend on the accuracy of the scattering model used to reach it. **Predicted:** the collision estimate's prefactor uncertainty is worth one to three orders of magnitude, while the shortfall to be closed is nine to sixteen; so conceding the entire prefactor objection leaves the conclusion unchanged. **Falsified by:** a corrected treatment of the same system — same degree of freedom, same environment, same temperature — returning a coherence time more than four orders of magnitude above the collision estimate.

H3 — Populations are untouched. Environment-induced decoherence changes no diagonal element of the reduced density matrix. **Predicted:** for a system coupled to an environment by `|i>|E_0> -> |i>|E_i(t)>`, tracing out the environment multiplies the off-diagonal elements by the environment-state overlap and multiplies the diagonal elements by exactly one, so any experimental readout that is a population — a reaction yield, a branching ratio — is unaffected by the loss of coherence that destroyed the interference. **Falsified by:** a decoherence process, derived from unitary evolution of system plus environment, that alters a population. This is a theorem and its falsifier is an error, not an observation; it is stated here in falsifiable form so that the claim's status is unambiguous.

H4 — Mode selection is exact across fields, not analogical. The relationship between boundary geometry and allowed modes that produces the Casimir force is the same boundary-value problem that produces a room's acoustic modes, with only the wave speed changed. **Predicted:** the lowest allowed mode between two boundaries separated by `L` is `v/(2L)` in both cases, so plates at 100 nm exclude every electromagnetic mode below 6.20 eV and walls at 7.0 m exclude every acoustic mode below 24.50 Hz, by the same formula and with the ratio of answers equal to the ratio of wave speeds, `8.7403e5`. **Falsified by:** a measured room-mode frequency departing from `f(n_x,n_y,n_z) = (c/2)sqrt((n_x/L_x)^2+(n_y/L_y)^2+(n_z/L_z)^2)` by more than the tolerance attributable to boundary compliance and to the finite impedance of real surfaces.

Three assertions in this paper are not hypotheses, and each is labelled where it appears.

The **statement of the quantum formalism in Section 4 is a report, not a claim.** It advances no postulate of the author's. Its correctness is checkable against any graduate text and it is offered so that the corpus's analogies can be audited against the original rather than against a paraphrase of it.

Postulate 6.1 — that boundary geometry alone selects modes at no ongoing cost — is badged as a postulate in the source and remains badged here. It is *exact* for electromagnetic modes between conductors and for acoustic modes in an enclosure, because both are the Helmholtz problem with a Dirichlet or Neumann condition; it is a **design heuristic and nothing stronger** when carried to occupancy, circulation or social structure, and Section 6.11 states the three transfer strengths separately so that the weakest cannot inherit the credibility of the strongest.

The correspondence between exponential decoherence and the storage law's retention term is a claim about model form, not about physics. It is category 2 in the analogy taxonomy of Section 3.3: the same first-order decay governs both, with no claim that the underlying physics applies. No institution decoheres. The storage law itself, with its contraction modulus `1−delta` and its steady state `(eta S − C)/delta`, is established in Curzi [2026a], and it is that paper's result R1 that the correspondence points at.

33.3 3 Method

33.3.1 3.1 What is established, what is postulated, and the separation between them

The paper follows the corpus's three-drawer classification of claims — established mathematics, established empirical result, and declared postulate — with every item assigned to exactly one drawer and every protocol stated to run on the first two while steering by the third [Curzi, 2026]. Applied to a physical rather than a psychological literature the drawers sort as follows.

Drawer one, established and not in dispute. The quantum formalism as a predictive calculus; the Born rule; unitarity and linearity; the bounding theorems of Section 4.9; the density-matrix account of decoherence; the Casimir energy and its measurement. Nothing in this paper questions any of it, and where the paper computes a number from it the computation is arithmetic on a stated formula with stated inputs.

Drawer two, established as contested. The interpretive question — what the formalism describes — is genuinely open, and competent physicists hold materially different positions. The reported long-lived coherences in photosynthetic complexes are real signals whose interpretation is disputed. Both are reported as contested, in the same passage as the claim they qualify.

Drawer three, this paper's own postulates. Exactly one: Postulate 6.1, badged in Section 6.5 and again in Section 6.11 and in the verification chain of Section 10, where its status is `POSTULATED` for the extension beyond acoustics and `DEFINITION` within it.

33.3.2 3.2 The audit protocol

The paper's operative contribution is a gate, applied in Section 8 to three proposals and stated in general form in Section 8.3. A claim that a quantum mechanism operates in a biological system is required to state five fields before it can be evaluated at all: the degree of freedom carrying the coherence, the environment it couples to, an estimated decoherence time with its formula and its inputs, the functional timescale the mechanism must reach, and the ratio of the last two.

The gate is a **canonicalisation** step before a comparison step, and it is the same discipline that a namespace integrity control applies to identifiers: put every submission into one normal form, and the comparison that follows is between like objects rather than between presentations. Curzi [2026] gives the measured version of that discipline in a different domain, where the canonical form `canon(x) = restrict_ASCII(confusable_fold(NFKC(x)))` collapses 495 of 495 manipulated entries onto the clean form and resolves 995 raw byte strings to the true 500 identifiers. The parallel is structural and is claimed as nothing more: a proposal missing one of the five fields is not refuted, it is unevaluable, and the gate returns it rather than judging it.

33.3.3 3.3 Analogy discipline

Every use of a physical result in this paper is labelled as exactly one of three things, at first use, following the corpus's standing rule.

1. **A derivation.** The physical result applies because the system satisfies its premises, and the paper shows that it does. Section 6.6 is the only instance: room acoustics is the same Helmholtz problem as the mode count in the Casimir derivation, with the electromagnetic field replaced by the pressure field and `c` replaced by the speed of sound. Nothing is transferred; the same equation is solved twice.
2. **A shared mathematical form.** The same equation governs both systems, with no claim that the underlying physics applies. Section 5.11's identification of exponential coherence loss with the storage law's retention term is the instance, and it is labelled there.

3. **An analogy.** A suggestive correspondence with no formal warrant, legitimate as motivation and never as evidence. The extension of mode selection to occupancy and circulation is the instance, and it is badged in Section 6.11 and again in Section 11.

The failure mode this taxonomy exists to prevent is a category-2 or category-3 use borrowing the certainty of category 1. Curzi [2026z] states the same discipline for the corpus's conjugate-variable material, where the future-as-conjugate-variable proposal is badged *analogy, not identity*, with an explicit refusal that quantum mechanics licenses anything about belief, mind, or future events. That refusal is the premise this paper supplies the evidence for. The general hazard of importing physical formalism into a non-physical domain is documented at length in the critical literature [Mirowski, 1989, Gallegati et al., 2006, McCauley, 2006], and Cartwright [1999] makes the point that matters most here: physical laws hold in the domains where their premises are satisfied and not beyond them, so the question for any transfer is always which premises the new system actually meets.

33.3.4 3.4 Scaling claims, and why none of them is estimated

Four power laws appear in this paper: the `a^−4` separation dependence of the Casimir pressure, the `a^−3` dependence of the non-retarded van der Waals pressure, and the `a^6 T^9` scaling of the blackbody localisation rate. **Every one of them is a model-derived exponent, not a fitted one.** No distribution was fitted, no tail was estimated, and no exponent in this paper was obtained from data. The distribution-fitting procedure that any *estimated* exponent would require [Clauset et al., 2009] therefore does not apply, and the exponents are quoted as exact consequences of the stated derivations. Section 11 records this as a check that does not apply rather than as one that was passed.

33.3.5 3.5 What this paper does not claim about mind, and the identification problem it inherits

The corpus's reflexive material — the loop in which a group influences its members and the members constitute the group — does not appear in this paper's physics, and no result in Sections 4 to 6 depends on it. It does appear once, indirectly, in Section 5.11, where the exponential form of coherence loss is put alongside the retention term of the storage law. **That correspondence gives the corpus a functional form; it gives the corpus no identification.** Applied to a social system the storage law's applied reading is a linear-in-means specification whose endogenous effect is **not identified from observational data** [Manski, 1993], and network data does not repair it, because influence and selection are generically confounded there [Shalizi and Thomas, 2011, Lyons, 2011]. This paper adopts no remedy, because it applies the law to nothing; it states the constraint so that no reader takes the physical exactness of `rho_01(t) = rho_01(0)exp(−t/tau)` as evidence for the identifiability of anything that shares its shape. The realistic remedy, where a downstream paper needs one, is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001].

A second reflexive hazard applies to the corpus's own use of this paper. The framework asserts that systems change their behaviour when they are modelled and observed [Lucas, 1976], and a set of specifications binding future writing is exactly such an intervention. The specifications in Appendix A are therefore stated as design rules whose effect on practice is unmeasured, not as findings.

33.4 4 The quantum formalism and the theorems that bound every interpretation

This section advances no postulate of the author's. It states an existing formalism and surveys its interpretive landscape, so that the corpus's analogies can be audited against the original.

Its correctness is checkable against any standard text and nothing in it is new.

33.4.1 4.1 The postulates as they are standardly stated

```
P1  STATE.        A physical system is associated with a complex Hilbert space H.
                  A pure state is a unit-norm vector |psi> in H, defined up to global phase.

P2  OBSERVABLES.  A measurable quantity is a self-adjoint operator A on H.
                  Its possible measured values are the eigenvalues of A -- always real.

P3  EVOLUTION.    An isolated system evolves by the Schrodinger equation,
                      i * hbar * d|psi>/dt = H |psi>
                  where H is the Hamiltonian. Equivalently |psi(t)> = U(t)|psi(0)>
                  with U unitary. This evolution is LINEAR and DETERMINISTIC.

P4  MEASUREMENT.  The probability of outcome a_n is |<a_n|psi>|^2  (the Born rule).

P5  COMPOSITION.  The state space of a composite system is the tensor product of the
                  component state spaces. Most states of a composite are not product
                  states; those that are not are entangled.
```

Quantum mechanics as a predictive calculus is the most precisely confirmed theory in the history of measurement. The interpretive question — what the formalism *describes* — is separate, and it is genuinely open. **The two are kept apart at every step below, because conflating them is how a settled calculus gets used to license an unsettled metaphysics.**

The honesty clause this section cannot do without. The measurement problem is genuinely open. It is not open in the direction of licence. Those are two separate statements and both are true. The first is a real admission: there is no consensus on what happens when a measurement produces one outcome rather than another. The second is a hard limit, and it is a limit imposed by theorems. **No-signalling holds in every interpretation. Whatever the resolution turns out to be, it will not be one in which a mind changes a distant statistic, because the statistic is provably unchanged.** An open question is not an open door.

33.4.2 4.2 The state vector, and four things it is not

The state vector `|psi>` is a unit vector in a complex vector space. From it, and only from it, the formalism computes the probability of every possible measurement outcome.

It is not a wave in space. For one particle in three dimensions, `psi(x,y,z)` looks like a field on space and the picture is nearly harmless. For two particles it is a function on six dimensions; for `N` particles, on `3N`. The state of a two-particle system is not two waves in space; it is one function on configuration space. Any intuition built on the claim that matter is a vibration in the room fails at the second particle.

It is not an ordinary probability distribution. Probabilities are non-negative reals that add. Amplitudes are complex numbers that add *before* being squared, which is why they can cancel. Section 4.4 computes this.

It is not a description of ignorance in the classical sense. Kochen–Specker and Bell together rule out the naive reading in which the particle already had all its properties and the experimenter merely had not looked. Some hidden-variable theories survive — de Broglie–Bohm does — but only by being explicitly non-local, which is a real price and Section 4.7 states it.

It is not agreed to be a physical object. Whether `|psi>` is a thing in the world or a representation of an agent's information is exactly the open question. The Pusey–Barrett–Rudolph theorem constrains the epistemic options substantially under an independent-preparation assumption, but it does not close the question and its assumption has been contested. It is a constraint, not a verdict.

33.4.3 4.3 Linearity, and why it is the whole problem

Linearity says that if the apparatus would evolve state A into “pointer reads UP”, and state B into “pointer reads DOWN”, then it evolves the superposition `a|A> + b|B>` into `a|pointer UP> + b|pointer DOWN>`.

It does not evolve it into “pointer reads UP”. **Linear evolution never produces a single outcome from a superposition. That is not a subtlety or a paradox; it is arithmetic, and it is why the measurement problem exists at all.**

The computation, run for this paper:

```
NUMERICAL VERIFICATION OF P3   (free particle, split-step FFT, N=8192, 4000 steps)

    two Gaussian packets, initially separated, moving toward each other

    norm before evolution                          1.00000000000000
    norm after  evolution                          1.00000000000105
    |relative norm change|                         1.05e-15

    max | U(a|psi_1> + b|psi_2>) - (aU|psi_1> + bU|psi_2>) |     1.29e-13

    BOTH AT DOUBLE-PRECISION FLOOR. Unitarity and linearity are exact.
```

Two consequences follow immediately and are worth stating separately because they are so often blurred.

1. **Probability is conserved.** Nothing leaks. There is no term in the equation through which an “energy of intention” could enter or leave, because no such term exists.
2. **Superposition is stable under evolution.** It does not decay on its own. Whatever removes it in practice is not in the Schrodinger equation — which is the subject of Section 5.

Appendix B, audit B1, records separately the evidence for this claim and the historical precedent sometimes offered alongside it, and states that the precedent carries no evidential weight.

33.4.4 4.4 Interference is the entire content

Take two paths with amplitudes `a` and `b`. Classically the probabilities add: `|a|^2 + |b|^2`. Quantum mechanically the amplitudes add first: `|a + b|^2`. The difference is the cross term.

```
    |a + b|^2  =  |a|^2 + |b|^2  +  2 Re(a* b)
                  \_____________/     \________/
                   classical sum    INTERFERENCE TERM
```

Computed with `a = b = 1/sqrt(2)` and a relative phase:

```
  phase (deg)   |a|^2+|b|^2   interference term   |a+b|^2
     0.0           1.0000          +1.0000          2.0000
    45.0           1.0000          +0.7071          1.7071
    90.0           1.0000          +0.0000          1.0000
   120.0           1.0000          -0.5000          0.5000
   135.0           1.0000          -0.7071          0.2929
   180.0           1.0000          -1.0000          0.0000
   270.0           1.0000          -0.0000          1.0000
   360.0           1.0000          +1.0000          2.0000

  average of |a+b|^2 over all phases 0..2pi  =  1.000000   <-- exactly the classical sum
```

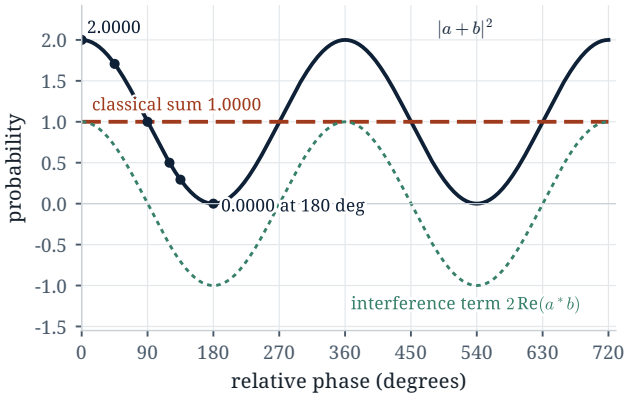


Figure 33.1: **Interference redistributes probability; it does not create any.** The total `|a+b|^2` for two equal amplitudes swings between 2.0000 and 0.0000 as the relative phase runs from 0 to 180 degrees, while the classical sum `|a|^2+|b|^2` stays at 1.0000 throughout; the phase-average of the quantum curve is 1.000000, exactly the classical value. Computed from the phase table of Section 4.4 with `a = b = 1/sqrt(2)`; generated by `figs/src/CCP-033_fig01.py`.

The last line is the most consequential in the section. **Averaging over the relative phase reproduces the classical result exactly. Interference is not extra probability; it is redistributed probability. Destroy the phase relationship and ordinary statistics are recovered with nothing left over.** Section 5 computes what destroys the phase relationship, and how fast.

The wavepacket version, on the same grid as Section 4.3, with two counter-propagating packets allowed to overlap:

```
central overlap region, quantum      : max density 0.05838, min 0.00000, visibility 1.0000
central overlap region, incoherent   : max density 0.02919, min 0.01004, visibility 0.4883
peak |quantum - incoherent|          : 0.02919  (= half the peak density)
total probability, quantum           : 1.0000000000
total probability, incoherent        : 1.0000000000    <-- identical totals

predicted fringe spacing  2*pi/dk    : 6.2832
measured mean fringe spacing         : 6.2174   (grid- and envelope-limited)
```

Perfect visibility, exact zeros between fringes, identical total probability. **The coherent and incoherent cases distribute the same unit of probability differently. That is what a wave function buys and it is all it buys.**

The experimental record on matter interference is deep. De Broglie wavelengths computed here:

```
50 keV electron (relativistic)          lambda = 5.3553 pm
  -> fringes for d = 1 um slits at L = 1 m : spacing 5.36 um
C60 fullerene (720 amu) at 200 m/s      lambda = 2.77 pm
C70 fullerene (840 amu) at 200 m/s      lambda = 2.38 pm
a 25,000 amu molecule at 10 m/s         lambda = 1.60 pm
```

Single-electron double-slit buildup has been recorded directly: electrons arrive one at a time as individual dots and the fringe pattern assembles from the accumulated dots. **That observation contains the entire content of the Born rule as an image — discrete arrivals, continuous distribution.** Molecular interference has been demonstrated for fullerenes and has since been pushed above 25,000 amu.

A note on the citations in this section. The corpus's external bibliography contains no entries for Born, Gleason, Bell, Kochen–Specker, Tonomura, Arndt, Zurek or Tegmark. Those results are named in the text with their year and venue, as the source manuscript states them, and are **not** attached to a bibliography key, because attaching a key that does not support the sentence is worse than leaving it unattached. Section 11 lists the gap, and Appendix C names the works so that a reader can locate them.

33.4.5 4.5 What the formalism does not say

Every line below is a claim the formalism does not make and has never made.

```
NOT SAID 1.  That "the observer" means a conscious being. In the formalism a
             measurement is an interaction with a system having many degrees of
             freedom that records a correlation. A photodiode does it. A speck of
             dust does it. Nothing in P1–P5 mentions awareness.

NOT SAID 2.  That attention, intention, expectation or belief appear in the
             Hamiltonian. There is no such term, and no proposal for one that
             has survived scrutiny.

NOT SAID 3.  That entanglement transmits anything. No-signalling forbids it.

NOT SAID 4.  That the uncertainty relation is about clumsy measurement.
             dx.dp >= hbar/2 = 5.272859e-35 J s is a property of conjugate
             variables, holding for an ensemble of identically prepared systems,
             and it holds before anyone measures anything.

NOT SAID 5.  That quantum indeterminacy scales up to macroscopic indeterminacy.
             Section 5 computes exactly how fast it does not.

NOT SAID 6.  That "everything is connected" in any operationally useful sense.
             Entanglement is a specific, fragile, monogamous resource that must be
             deliberately created and carefully protected. It is not ambient.

NOT SAID 7.  That possibilities not realised are "still available". Which
             unrealised branches exist is exactly the interpretive dispute; no
             interpretation makes them accessible.

NOT SAID 8.  That the wave function is a field of potential that responds to
             desire. There is no coupling; and with no channel there is no effect.
```

The state-preparation reading of NOT SAID 4 is the one the corpus adopts elsewhere: the inequality's original derivation is Kennard [1927], and the corpus's own treatment of conjugate structure — which is Fourier duality and carries no quantum content — badges its extension to foresight as *analogy, not identity* [Curzi, 2026z]. **Every popular quantum claim about manifestation violates at least one of these eight lines, and most violate five.**

33.4.6 4.6 The measurement problem, stated precisely

The problem is a strict inconsistency between three independently plausible statements.

```
(M1)  The wave function is a complete description of a physical system.
(M2)  The wave function always evolves by the linear Schrodinger equation.
(M3)  Measurements have single, definite outcomes.
```

```
These three cannot all be true. Every interpretation is a decision about which
one to give up.
```

The framing is standard, and it makes the interpretation survey trivially organisable, because **the map from interpretation to abandoned premise is one-to-one.**

INTERPRETATION	ABANDONS	WHAT IT COSTS
Copenhagen-type	M2	a second, non-unitary dynamical law with no stated trigger condition
Everett / many-worlds	M3	an unbounded ontology; and the Born rule must be derived rather than assumed
de Broglie–Bohm	M1	explicit non-locality; a preferred frame; difficulty extending to quantum field theory
GRW / CSL	M2	new physical constants, and it is a DIFFERENT THEORY that can be tested

33.4.7 4.7 The four positions, with what each abandons

Copenhagen-type views. Not one doctrine but a cluster. The common core: the formalism predicts outcomes of experiments described in classical terms, and the collapse postulate is a separate, non-unitary update applied when a measurement occurs. *What it buys:* it works, and no experiment has ever required more. *What it costs:* the collapse postulate has no stated trigger; what counts as a measurement is left to judgement, and that gap is the measurement problem. *Status:* the plurality working position among physicists, and the target of the sharpest foundational criticism. Both are true at once.

Everett / many-worlds (Everett 1957, *Reviews of Modern Physics* 29, 454). Take the Schrodinger equation to be the whole story. There is no collapse; linearity produces a superposition of apparatus states and decoherence makes the branches dynamically autonomous. *What it buys:* one law, no second dynamics, no ambiguity about triggers. *What it costs:* an enormous ontology, and a real technical problem — if all branches occur, what does the Born rule mean? Decision-theoretic derivations of the Born weights from rationality axioms have been advanced; **they are serious and they are contested, critics having argued that they assume the measure they claim to derive. The probability problem in Everett must not be presented as solved.** *Status:* a substantial and growing minority position.

De Broglie–Bohm pilot-wave theory (Bohm 1952, *Physical Review* 85, 166 and 180). Particles have definite positions at all times; the wave function is a real field guiding them, with outcomes determined by initial positions distributed as `|psi|^2`. *What it buys:* a clear ontology, no measurement problem, a straightforward double slit — the particle goes through one slit, the wave through both. *What it costs:* manifest non-locality, consistent with no-signalling but uncomfortable with relativity, and extension to relativistic field theory remains unresolved. *Status:* a small, rigorous, continuing programme whose historical importance is decisive — the 1952 papers are an existence proof that the impossibility arguments then circulating against hidden variables were wrong.

GRW and continuous spontaneous localisation (Ghirardi, Rimini and Weber 1986, *Physical Review D* 34, 470). Modify the dynamics: add a stochastic term that spontaneously localises the wave function at a very low rate per particle, negligible for one particle and enormous for `10^23` of them. *What it buys:* definite outcomes come out of the dynamics; no interpretation required. *What it costs:* it is not an interpretation but a **different theory**, with new fundamental constants, hard to reconcile with energy conservation and relativity. *Status:* **its greatest virtue is that it is being experimentally tested and its parameter space narrowed**, including an underground X-ray measurement excluding the simplest parameter-free gravity-related collapse model. A programme that can be excluded by measurement is doing science in the

ordinary way, and this paper says so approvingly. Appendix B, audit B2, records the evidence and the precedent for objective spontaneous indeterminacy separately.

33.4.8 4.8 Consciousness-causes-collapse, stated accurately

The view that a conscious observation collapses the wave function is usually called the von Neumann–Wigner interpretation. It deserves an accurate history.

The chain. Von Neumann (1932) showed that the boundary between quantum system and classical apparatus can be moved arbitrarily far along the measurement chain — system, apparatus, amplifier, retina, optic nerve — without changing any prediction. This is a correct and important result. **What it shows is that the cut is movable. It does not show that the cut must be placed at consciousness, and the argument supplies no reason to put it there.** London and Bauer (1939) took the further step and located the cut at the observer's introspective act; Wigner argued for the position most explicitly in 1961.

Why it is a minority position. Five reasons, in order of weight.

1. **It solves nothing that decoherence does not address better, and it introduces something new.** Once environmental decoherence is accounted for (Section 5), the appearance of classicality is explained without invoking mind. Consciousness-collapse must still posit a novel non-linear coupling for which there is no evidence and no proposed form.
2. **It has no formalism.** No equation, no definition of which physical systems are conscious, no threshold, no rate, no coupling constant. A proposal with no free parameters to constrain cannot be tested, and is therefore not competing on the same field as GRW, which can be and is.
3. **It generates absurdities under iteration.** Extended Wigner's-friend results produce scenarios in which observers using quantum mechanics correctly reach contradictory conclusions about each other. These constrain all interpretations but bear especially hard on views granting observers a special dynamical status.
4. **It has an unacceptable cosmological consequence.** If collapse requires consciousness, the universe had no definite outcomes for roughly 13.8 billion years before observers existed — including the nucleosynthesis whose definite outcomes are now measured.
5. **Wigner himself abandoned it**, in significant part because of the work on environment-induced decoherence begun by Zeh [1970]. **The interpretation's most distinguished advocate stopped advocating it, for the reason developed in Section 5.**

Where it stands. Informal surveys of physicists' foundational attitudes are small and self-selected, and their authors say so; across them, consciousness-causes-collapse consistently attracts a small fraction, well behind Copenhagen-type, information-theoretic and Everettian positions. **The exact statement, and no stronger: it is a minority historical interpretation with few contemporary adherents; it is not refuted, and it is not favoured.** Appendix B, audit B3, records that the evidence column for this claim is **empty**, and prints the empty cell.

33.4.9 4.9 The theorems that bound every interpretation

Whatever the resolution, these constraints already hold, and they are the reason the openness of the measurement problem cannot be traded for licence.

```
NO-SIGNALLING     No local operation changes any distant statistic. Holds in
                  Copenhagen, Everett, Bohm, GRW, and in every theory consistent
                  with the observed correlations.

NO-CLONING        An unknown state cannot be duplicated. Follows from linearity.

BELL / CHSH       |S| <= 2 for local hidden variables; quantum limit 2.828427;
```

```
                  measured violations, loophole-free, 2015; Nobel Prize 2022.
                  Locality-with-definite-values is excluded, not "realism" as
                  such -- the popular gloss that physics disproved reality is wrong.

KOCHEN-SPECKER    No global non-contextual value assignment in dimension >= 3.

PBR (2012)        Under an independence assumption on preparations, purely
                  epistemic states are excluded. A constraint, contested.

GLEASON (1957)    The Born rule is forced, given the Hilbert-space structure
                  and additivity, for dimension >= 3.
```

The quantum bound `2 sqrt(2) = 2.828427` is attained exactly at the optimal measurement settings, and the arithmetic is checked in Section 10.3. **The honest summary: quantum mechanics has narrowed the space of possible worldviews substantially, and every direction in which it has narrowed that space removes options rather than adding them.** Bell did not open a door to influence at a distance; it closed the door on local pre-existing values while leaving no-signalling untouched.

That these constraints are load-bearing engineering rather than philosophy is visible in quantum information, where the impossibility of cloning is what makes key distribution work and the fragility of coherence is what makes the algorithms hard to run [Shor, 1997, Grover, 1996].

33.4.10 4.10 The corpus's branch-weight construct, and the connection that is honest

The corpus's fourth foundational construct asks the analyst to hold the future as a weighted set of branches and to read the weights without preference; its formal content is the comfort-weighted amplitude, in which each step is scored 0–3 for steepness, the scores are summed to a resistance `R`, and the branch is weighted as `1/2^R`, so that a three-point difference in steepness is a factor of eight in expectation [Curzi, 2026]. The relationship to quantum mechanics is stated here deliberately deflationarily.

What is shared. A formal shape: a set of mutually exclusive outcomes, each carrying a weight, the weights summing to one, with a discipline of not collapsing the set prematurely into a single expectation. **That shape is called a probability distribution. It is older than quantum mechanics, it is used everywhere, and it requires nothing from physics.**

What is not shared. Everything that makes quantum mechanics quantum: complex amplitudes, interference between branches, `hbar`, non-commuting observables, entanglement, the measurement problem. The corpus's branches do not interfere. There is no cross term. Two futures that cancel is not a thing the construct describes and not a thing this corpus claims.

Why the deflation is a gain. A probability distribution over futures is testable in the ordinary way. The corpus's preregistration protocol applies directly — state the branch set in advance, assign the weights in advance, freeze the model, and score the held-out outcome against declared baselines — and that protocol, together with the one confirmatory application of it so far reported, is Curzi [2026i], whose fit **failed** its own second falsification rule with a 2.8 per cent held-out advantage over a matched AR(1) against a declared 10 per cent margin. **A quantum dressing would add nothing to that procedure and would subtract its testability, because nobody can score a claim about wave-function branches.**

The corpus's own measure of a branch set is already available. For a two-branch future with weights `p` and `1-p`:

```
p        H (nats)    H (bits)    N_eff = e^H
0.50     0.693147    1.000000    2.0000       maximum spread; two live futures
0.25     0.562335    0.811278    1.7548
```

```
0.10     0.325083    0.468996    1.3841
0.01     0.056002    0.080793    1.0576       effectively one future
```

`N_eff = exp(H)` here is the Hill number of order one, a fully grounded effective count [Hill, 1973, Jost, 2006], and the entropy is Shannon's, reported in nats [Shannon, 1948, Cover and Thomas, 2006]. It is **not** the corpus's affine effective count `1 + (N−1)u`, which is a linear interpolation of the framework's own construction and is not a Hill number of any order; the two must not be run under one name. A future held as a weighted ensemble has an effective number of futures in it, it is computable, and there is no `hbar` anywhere in the table. Appendix B, audit B4, records the evidence and the precedent for weighted-branch forecasting separately.

33.4.11 4.11 Why the borrowing happens

Quantum mechanics supplies four words that already exist in ordinary English with different meanings: *observer*, *measurement*, *uncertainty*, *possibility*. In the formalism each is a technical term; in ordinary use each carries a mentalistic connotation. **The borrowing is a sustained equivocation on four words, and it survives because the technical meanings sit behind a mathematical barrier most readers cannot cross.**

The counter is not to police vocabulary. It is to demand the coupling. The corpus's channel requirement has the form `belief -> [channel] -> world`, and *grounded in truth* means the channel exists and is real. **Ask for the term in the Hamiltonian. There is not one. That ends the argument without anyone having to settle a metaphysical question.** The requirement is the operational content of the corpus's grounding postulate `E = G B`, whose forbidding clause `G = 0` implies `E = 0` for every value of `B` [Curzi, 2026ab] — a postulate, badged as one there and treated as one here.

33.5 5 Decoherence in warm wet systems, computed

This section states the strongest scientific objection to the corpus's biological material and computes it. The objection is that environmental decoherence destroys phase coherence in a warm wet system on a timescale of order `1e−17` seconds, while the shortest timescale on which anything neurally interesting happens is about `1e−4` seconds. **The gap is thirteen orders of magnitude and it is the wrong way round.**

There are two ways to handle an objection of that size: route around it and keep the vocabulary vague, or compute the numbers, print them, and accept the constraint. **This paper does the second, and the corpus is stronger for it.** There is also a non-defensive reason for the section: decoherence explains why the classical world looks classical without adding anything to the Schrodinger equation, and it fails to explain the rest, which is equally important to say.

33.5.1 5.1 The mechanism and the two theorems

Definition 5.1 (density matrix). For a system in state `|psi> = a|0> + b|1>`, the density operator is `rho = |psi><psi|`:

```
rho  =  [ |a|^2      a b*  ]      diagonal entries : POPULATIONS
        [ a* b       |b|^2 ]      off-diagonal     : COHERENCES
```

The diagonal entries are the Born-rule probabilities. **The off-diagonal entries carry the phase relationship and are the only thing distinguishing a superposition from a classical mixture.** All observable interference lives in those two numbers.

Theorem 5.2 (decoherence as off-diagonal suppression). *Established, and derived from ordinary unitary evolution of system plus environment.* Let the system couple to an environment such that `|0>|E_0> -> |0>|E_0(t)>` and `|1>|E_0> -> |1>|E_1(t)>`. Tracing out the environment gives

```
rho_S(t)  =  [ |a|^2                    a b* <E_1(t)|E_0(t)>  ]
             [ a* b <E_0(t)|E_1(t)>     |b|^2                ]
```

The populations are untouched. The coherences are multiplied by the overlap of the environment states. As those states become distinguishable the overlap goes to zero and the off-diagonal terms vanish. For most realistic couplings the overlap decays exponentially,

```
rho_01(t)  =  rho_01(0) * exp( - t / tau_dec )
```

and for spatial superpositions of separation `dx` the standard long-wavelength form is `rho(x,x',t) = rho(x,x',0) exp(-Lambda (x-x')^2 t)` with `Lambda` a scattering constant. The programme was begun by Zeh [1970], developed quantitatively by Joos and Zeh (1985) and systematised by Zurek (2003).

Theorem 5.3 (populations are not affected). *Proven, and it is the whole of what Section 5.7 turns on.* The diagonal elements of `rho_S(t)` are `|a|^2` and `|b|^2` for every `t`. Decoherence suppresses the coherences and changes no population. **It converts a superposition into something that looks like a classical mixture without selecting an outcome.**

33.5.2 5.2 The rate, and where the saturated regime begins

For a superposition whose spatial separation exceeds the thermal de Broglie wavelength of the environmental particles, a single scattering event resolves which branch the system is in, and the decoherence rate saturates at the scattering rate:

```
Gamma_dec  =  n * sigma * v_bar          (saturated regime)

    n      = number density of the environment particles
    sigma  = scattering cross-section (geometric, as an estimate)
    v_bar  = sqrt( 8 k_B T / (pi m) ), mean thermal speed

tau_dec    =  1 / Gamma_dec
```

Below the thermal de Broglie wavelength the rate is quadratic in separation and decoherence is slower, so the boundary must be located first:

```
THERMAL DE BROGLIE WAVELENGTH OF THE ENVIRONMENT PARTICLE   (computed)

air molecule (28.97 amu), 300 K     v_bar = 468.2 m/s   hbar/p = 4.68 pm
water molecule (18.02 amu), 310 K   v_bar = 603.6 m/s   hbar/p = 5.84 pm
helium (4 amu), 4 K                 v_bar = 145.5 m/s   hbar/p = 109.11 pm
```

Any superposition separated by more than about six picometres, in air or water at ordinary temperatures, is in the saturated regime. Six picometres is one twentieth of a hydrogen atom's diameter. Every proposal of biological interest — a displaced ion, a conformational change, a travelling exciton, a firing versus non-firing neuron — is separated by vastly more than that, so the saturated formula applies, and it is the fast one.

33.5.3 5.3 Case one: a dust grain

Radius 1 micrometre, air at 300 K, geometric cross-section.

```
DUST GRAIN, r = 1 um                    n (/m3)     Gamma (/s)    tau_dec (s)

air at 1 atm, 300 K                     2.446e+25   3.600e+16     2.78e-17
high vacuum, 1e-10 mbar                 2.414e+12   3.553e+03     2.82e-04
extreme high vacuum, 1e-12 mbar         2.414e+10   3.553e+01     2.82e-02
```

In ordinary air a speck of dust loses spatial coherence in about thirty attoseconds. In the best vacuum routinely achievable in a laboratory it lasts about a quarter of a millisecond — and that is with the gas removed, before thermal photons are counted.

Thermal photons make it worse and their scaling is severe. In the long-wavelength limit the blackbody localisation rate scales as `a^6 T^9` in the object's radius and the temperature. **Both exponents are derived from the model, not fitted to data**, and are quoted as exact:

```
T = 150 K relative to 300 K :  (T/300)^9 = 1.953e-03
T =  77 K relative to 300 K :  (T/300)^9 = 4.834e-06
radius x 2                  :  r^6 factor = 6.400e+01
radius x 10                 :  r^6 factor = 1.000e+06
```

A ninth-power temperature dependence is why matter-interference experiments are done cold, and a sixth-power size dependence is why they are done with small things. Both are engineering consequences of one formula and both are visible in the design of every real experiment in the field.

33.5.4 5.4 Case two: a large molecule, where the estimate meets an experiment

Radius 0.5 nanometres, roughly a C70 fullerene.

```
C70 MOLECULE, r = 0.5 nm                pressure       Gamma (/s)   tau_dec (s)

air at 1 atm                            1.013e+05 Pa   1.520e+10    6.58e-11
1e-3 mbar                               1.0e-01 Pa     1.501e+04    6.66e-05
1e-6 mbar                               1.0e-04 Pa     1.501e+01    6.66e-02
1e-8 mbar                               1.0e-06 Pa     1.501e-01    6.66e+00
```

This table is worth more than the dust grain because it can be checked against a real apparatus. A Talbot–Lau molecule interferometer has a flight time of order a millisecond and a grating period of order a micrometre — comfortably in the saturated regime, a micrometre being two hundred thousand times the 4.68 pm threshold. **The estimate therefore predicts that fringe visibility should survive at `1e-6` mbar, where the coherence time of 67 ms far exceeds the flight time, and should be destroyed at `1e-3` mbar, where the coherence time of 67 microseconds falls far below it. That is the pressure range over which the visibility has been measured to fall.** A first-principles estimate that lands on the right decade in a published experiment is an estimate worth trusting one decade at a time, which is all this section asks of it. **This is the paper's only out-of-sample check of an estimator against an outcome it did not use**, and Section 10 treats it as such.

33.5.5 5.5 Case three: the warm wet system

Liquid water at 310 K. This is the environment inside a cell and the relevant case for any biological claim whatsoever.

```
WATER AT 310 K (computed)
    number density                      n     = 3.3428e+28 /m3
    mean thermal speed of a molecule    v_bar = 603.6 m/s
    thermal de Broglie wavelength       hbar/p = 5.84 pm

OBJECT                              cross-section   Gamma (/s)    tau_dec (s)

single K+ ion, r = 0.14 nm          6.158e-20 m2    1.242e+12     8.05e-13
globular protein, r = 5 nm          7.854e-17 m2    1.585e+15     6.31e-16
microtubule segment, 25 x 100 nm    2.500e-15 m2    5.044e+16     1.98e-17
synaptic vesicle, r = 20 nm         1.257e-15 m2    2.536e+16     3.94e-17
whole neuron soma, r = 5 um         7.854e-11 m2    1.585e+21     6.31e-22
```

Three notes on how to read the table honestly.

First, these are estimates and the prefactors are uncertain. The collision picture is imported from a gas; in a liquid the object is not struck intermittently but is in permanent contact with its solvent shell. If anything this *understates* the monitoring rate. The relevant solvent dynamics:

```
water librational band, ~600 cm^-1   ->  1.80e+13 Hz  (period 5.56e-14 s)
OH stretch, ~3400 cm^-1              ->  1.02e+14 Hz  (period 9.81e-15 s)
hydrogen-bond rearrangement          ->  ~1e+12 Hz    (period ~1e-12 s)
hbar / kT at 310 K                   ->  2.464e-14 s
```

No mechanism in liquid water leaves anything alone for longer than a picosecond, and most leave nothing alone for longer than about twenty-five femtoseconds.

Second, the thermal time `hbar/kT` is the natural yardstick and deserves its own table.

```
T = 2.725 K   (cosmic microwave background)   hbar/kT = 2.803e-12 s = 2803 fs
T = 4     K   (liquid helium)                 hbar/kT = 1.910e-12 s = 1910 fs
T = 77    K   (liquid nitrogen)               hbar/kT = 9.920e-14 s =   99 fs
T = 277   K   (a cold room, 4 C)              hbar/kT = 2.758e-14 s = 27.6 fs
T = 300   K   (room temperature)              hbar/kT = 2.546e-14 s = 25.5 fs
T = 310   K   (body temperature, 37 C)        hbar/kT = 2.464e-14 s = 24.6 fs
```

At body temperature the thermal time is twenty-five femtoseconds. That is the timescale on which thermal energy scrambles quantum phase, and it is a property of the temperature alone. Any biological quantum claim must explain how its state stays coherent for many multiples of 24.6 femtoseconds while sitting in a bath at 310 K, and the explanation has to be structural.

Third, warmth and wetness are not incidental to biology. They are what biology is. Cooling the system to where coherence survives kills the organism. This is not a technical inconvenience to be engineered around; it is a constraint on the kind of physics life can be running.

33.5.6 5.6 The gap, stated as a ratio

Take the microtubule estimate, `tau_dec = 1.98e-17 s` — the object most often invoked in quantum-mind proposals — against every timescale on which anything neural happens. The neural timescales are the standard ones of the spiking-neuron literature [Gerstner and Kistler, 2002, Lapicque, 1907].

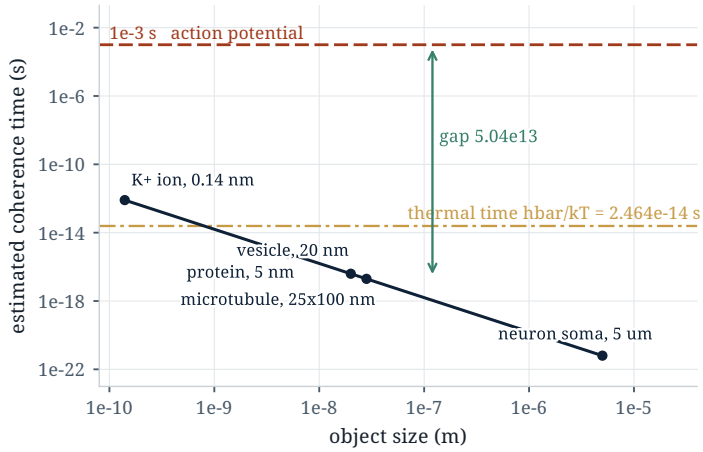


Figure 33.2: **The warm-wet decoherence gap, over nine decades of object size.** Estimated coherence times in water at 310 K fall from `8.05e-13` s for a single potassium ion to `6.31e-22` s for a neuron soma, every one of them below the thermal time `hbar/kT = 2.464e-14` s and every one of them between eight and nineteen orders of magnitude below the `1e-3` s of a single action potential. The microtubule segment most often invoked in quantum-mind proposals sits at `1.98e-17` s. Estimated from `tau = 1/(n sigma v_bar)` with `n = 3.3428e28` per cubic metre and `v_bar = 603.6` m/s, the cross-sections of Section 5.5; generated by `figs/src/CCP-033_fig02.py`.

PROCESS	duration	duration / tau_dec
ion-channel gating	1.0e-04 s	5.04e+12
action potential	1.0e-03 s	5.04e+13
synaptic transmission	1.0e-03 s	5.04e+13
gamma cycle (40 Hz)	2.5e-02 s	1.26e+15
a reported "moment" of awareness	1.0e-01 s	5.04e+15
one second	1.0e+00 s	5.04e+16

And the corresponding coherence, since the decay is exponential:

```
rho_01(t) = rho_01(0) * exp(-t / tau_dec),  tau_dec = 1.98e-17 s

t = 1.98e-17 s (one tau)        suppression factor  0.367879
t = 1.98e-16 s (ten tau)        suppression factor  0.000045
t = 1.98e-15 s (hundred tau)    suppression factor  ~1e-44
t = 1.0e-12  s                  exp(-5.044e+04)     numerically zero
t = 1.0e-03  s (one spike)      exp(-5.044e+13)     numerically zero
```

"Numerically zero" here is not a figure of speech and not a rounding. The suppression factor at one millisecond is `e` raised to minus fifty trillion. There is no number system in ordinary use that represents it. Interference visibility falls below one part in a thousand after 6.9 decoherence times and below one part in a million after 13.8; a millisecond is `5.04e13` of them.

33.5.7 5.7 What decoherence solves, and what it leaves standing

What it genuinely solves.

1. **The absence of macroscopic interference.** Why cats, chairs and pointers are never seen in superposition. Answered completely, by the coherence times above.

2. **The preferred-basis problem.** Why the classical world has a definite basis, roughly position, rather than an arbitrary one. Einselection answers this from the interaction Hamiltonian with no appeal to observers.
3. **The apparent irreversibility of measurement.** The information is not destroyed but dispersed into environmental degrees of freedom that are in practice unrecoverable. **This is provenance loss at a boundary, and the corpus has a measured instance of the same structure in a different domain**: 18.932 bits of identity entering a carrier becomes 8.000 bits observable from an address alone, and 23.977 bits when address and port are both captured [Curzi, 2026]. The parallel is structural. What decoherence disperses is phase, not identity, and no claim is made that the two are the same quantity; that erasure and information have a physical cost at all is Landauer [1961].
4. **The disappearance of the observer from the account.** Decoherence is done by air molecules. It removes the last motivation for putting consciousness anywhere in the dynamics, and this is the development that led Wigner away from his own 1961 position.

What it does not solve, stated flatly.

Theorem 5.3 said it: decoherence changes no diagonal element. It suppresses the interference terms; it does not select an outcome. The reduced density matrix ends up looking like a classical mixture, but it is an *improper* mixture, obtained by tracing out a subsystem of a state that is still globally a pure superposition.

```
BEFORE   rho_S =  [ |a|^2    a b*  ]     coherent superposition
                  [ a* b     |b|^2 ]

AFTER    rho_S =  [ |a|^2      0   ]     looks exactly like a classical mixture
                  [   0      |b|^2 ]     but is an IMPROPER mixture: the global
                                         state is still a superposition

WHAT CHANGED : the off-diagonals. WHAT DID NOT : |a|^2 and |b|^2, and the fact
               that both branches are still in the global state.
```

The measurement problem in the M1/M2/M3 form of Section 4.6 is untouched by this. **Decoherence is entirely a consequence of M2 applied to system plus environment. A result derived from M2 cannot resolve a contradiction that arises because M2 conflicts with M3. Anyone who says decoherence solved the measurement problem has made a logical error, not merely an overstatement.** The criticism is not fringe and has been made explicitly in the foundations literature. **The honest position, and the majority one among people who work on foundations, is that decoherence is a major and permanent advance that dissolves part of the problem and leaves the rest exactly where it was.**

The half that is solved is the half this corpus needs. Whether the surviving outcome is selected by collapse, by branching, or by a guiding equation makes no difference to any argument in the corpus. What matters is that no macroscopic superposition survives long enough to carry biological information, and that is settled. **The half that is open is genuinely open, and this corpus takes no position on it, because it does not need one.**

33.5.8 5.8 The plain statement, and what it costs the corpus

Any claim that mind collapses wave functions must confront these timescales.

Any claim that quantum coherence carries information in warm wet biology must confront these timescales.

This corpus makes no such claim.

The consequences for the corpus's biological material, stated so they can be checked:

the cell as receiver	MECHANOTRANSDUCTION. Ion channels, forces, conformational change. Entirely classical. Mechanosensitive channels open because they are pulled.
the channel and its capacity	Classical acoustics, classical Shannon capacity, thermal noise floor. No coherence term anywhere in the calculation.
addressing without editing	Molecular recognition, base pairing, protein–nucleic-acid binding. Chemistry.
sequence sonification	Encoding, transmission. The Gabor limit is CLASSICAL and it is the binding constraint, not the Heisenberg relation.
belief and mechanism	Physiological, behavioural, social and informational channels. Every one measured, every one classical.

NOT ONE ARGUMENT IN THE CORPUS WOULD CHANGE IF QUANTUM MECHANICS PLAYED NO ROLE IN BIOLOGY AT ALL.

Each of those lines is carried by a paper that computes it. The acoustic-addressing result is that addressing a cell fails on geometry rather than on bandwidth: wavelength equals a 10 micrometre cell only at 154 MHz, and the product of wavelength and 20 dB penetration depth collapses from 6160 to 0.0026 mm-cm between 100 kHz and 154 MHz [Curzi, 2026x]. The sequence-encoding result is the open problem stated formally: **no established mechanism produces a transcriptional response depending on the symbol string as an order, as distinct from permutation-invariant functionals of the waveform** [Curzi, 2026]. The classical time–bandwidth constraint that binds there is Gabor [1946], whose limit is a Fourier fact and carries no quantum content [Curzi, 2026z]. Appendix B, audit B5, records the evidence and precedent for the isolation requirement separately.

33.5.9 5.9 The genuine exceptions, quantified and contested

Honesty runs in both directions. There are real cases where quantum coherence appears to play a functional role in biology, they are active research areas, and this section names them properly.

(a) Photosynthetic energy transfer. Two-dimensional electronic spectroscopy of the Fenna–Matthews–Olson pigment–protein complex revealed long-lived oscillatory signals, first at 77 K and then at physiological temperature, with reported oscillation lifetimes at the warmer temperature of order 300 femtoseconds.

reported oscillation lifetime	3.0e-13 s (300 fs)
hbar / kT at 277 K	2.758e-14 s (27.6 fs)
ratio	10.9 thermal times
spatial extent	1–2 nm, inside ONE rigid protein scaffold

The interpretation is contested, and that must be said in the same breath. A substantial body of subsequent work argues that the long-lived oscillations are predominantly vibrational or vibronic rather than evidence of long-lived electronic coherence, and that the energy transfer does not require the latter. **The defensible statement is: short-lived coherence exists in these systems, its functional role is disputed, and the strong early claims have been substantially walked back.**

(b) Radical-pair magnetoreception. The radical-pair mechanism proposes that a photo-induced pair of radicals, born in a correlated spin state, undergoes coherent singlet–triplet interconversion whose rate is sensitive to the external magnetic field, so that the resulting difference in reaction yield constitutes a compass. Cryptochrome has been proposed as the host protein; a purely chemical model system has been shown to respond to Earth-strength fields; and magnetic sensitivity of cryptochrome 4 from a migratory songbird has been reported in vitro.

```
electron Zeeman energy in Earth's field (50 uT, g = 2)   9.274e-28 J
kT at 310 K                                              4.280e-21 J
ratio (Zeeman / thermal)                                 2.167e-07

coherent precession frequency  = 8.794e+06 rad/s = 1.400 MHz
precession period                                       0.714 us
=> the mechanism requires SPIN coherence on the microsecond scale
```

The magnetic interaction is two ten-millionths of the thermal energy. It cannot work by energetics — and it does not claim to. It works by kinetics: the field changes the rate of singlet–triplet mixing, which changes which of two chemical products forms, and the yield difference is what is read out.

The reason the coherence survives is structural, not mysterious. The degree of freedom carrying it is electron *spin*, which has no electric dipole moment and therefore does not couple to the polar solvent that destroys every charge-displacement superposition in the tables above. **The exception is not an exception to the physics. It is a consequence of choosing a degree of freedom the environment barely watches.**

33.5.10 5.10 Why the exceptions do not generalise

Both cases share four features, and every one of them fails at the scale of cognition.

FEATURE	PHOTOSYNTHESIS / RADICAL PAIR	COGNITION
spatial extent	1–2 nm, inside one protein	1e-5 m and up (ratio ~ 1e+4)
duration	3e-13 s / ~1e-6 s (spin)	1e-3 to 1e-1 s (ratio ~ 1e+9)
degree of freedom	one exciton; one electron-spin pair	membrane voltage; ion populations; vesicle fusion -- all charge-displacing and all heavily monitored
protection	a rigid, evolved, purpose-built scaffold that constrains the coupling	NONE PROPOSED
readout	a chemical yield -- a POPULATION difference, not a transmitted message	a claim about information transfer

The last row ends the argument and deserves to stand alone. In both exception cases what is read out is a reaction yield — a diagonal element of the density matrix, a population.

Theorem 5.3 says decoherence never touched the populations. The coherence in these systems modulates a rate; it does not carry a message. Even taken entirely at face value, neither case is an example of quantum coherence transmitting information, and neither is proposed by its own authors as such.

The sharpest version of this argument in print is a published exchange. Tegmark (2000, *Physical Review E* 61, 4194) computed decoherence times for microtubule and neuron-firing superpositions and obtained `1e−13` to `1e−20` seconds, concluding they were ten or more orders of magnitude too short for any role in cognition. Hagan, Hameroff and Tuszynski (2002, *Physical Review E* 65, 061901) replied that the model was mis-specified for the system they had proposed and, correcting it, obtained `1e−5` to `1e−4` seconds. **Both papers are real, both are peer-reviewed, and the disagreement was never resolved in favour of the longer times.** The independent estimates in Section 5.5, computed from scratch here, land in the first range and not the second. **The defensible summary: the objection stands, it has been contested by serious people, and the contest did not overturn it.**

That a microscale law does not determine the behaviour of the macroscale system built from it is the general form of this result, and it is not special to quantum mechanics [Anderson, 1972].

33.5.11 5.11 What decoherence returns to the framework

The section has been subtractive. It ends by paying something back, and what it pays back is structural rather than quantum.

Decoherence has the form of a retention loss. In the storage law `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t`, the coefficient `delta` is what a stored structure sheds per period through contact with an unmaintained surround; the law is a contraction for every `delta > 0`, with steady state `(eta S − C)/delta` [Curzi, 2026a]. Coherence decay has exactly that form — exponential loss driven by exposure, with no agent and no malice, requiring no explanation beyond contact.

```
COHERENCE       rho_01(t) = rho_01(0) * exp(−t / tau)
STORED ORDER    Q_t       = (1−delta)^t * Q_0        (transient term)

same functional form; delta plays the role of 1/tau per period
```

This is category 2 in the taxonomy of Section 3.3: a correspondence of model form, not of physics. No institution decoheres. The first-order-lag mathematics is standard and is the same object in both places [Elaydi, 2005, Oppenheim et al., 1996]. The engineering lesson, however, transfers exactly and is worth having:

A fragile state is protected by reducing its coupling to the surround, not by adding energy to it. No amount of drive escapes decoherence, and no amount of drive escapes `delta`. The matter-wave interferometrists' response to a `T^9` and `a^6` scaling is to go cold, small and empty — reduce the coupling. The framework's counterpart is to reduce exposure surface or accept the loss rate on hand. **The converse case is the instructive one and the corpus states it: a system whose loss is exactly zero does not preserve itself, it becomes an isometry that orbits without arriving, and a system whose gain reaches or exceeds one becomes an amplifier whose corrections, sized to delayed feedback, arrive as new errors** [Curzi, 2026].

33.6 6 The Casimir effect: force, honest limit, and mode selection by boundary

The limit comes first, because the limit is the honest part and it belongs where nobody can miss it. At every biological and architectural scale the Casimir force is utterly negligible. It

does nothing in a cell and nothing in a room. Between two walls three metres apart it produces a pressure of `1.47e-23` pascals, which over ten square metres is a force of `1.47e-22` newtons — sixty-seven million times smaller than the weight of a single bacterium. In a cell, at a hundred-nanometre gap, the most generous possible estimate gives thirteen pascals against a physiological osmotic pressure of `7.73e5` pascals, and the realistic estimate for actual biological materials in salt water is smaller still by another one to three orders. **No design in this framework depends on the Casimir force, and none could be built to exploit it.**

So why is the effect here? Because it is not important for its force. It is important for what produces the force. **Two uncharged, unpowered, ordinary plates attract each other for one reason: their separation determines which electromagnetic modes can exist between them, and modes that do not fit are excluded.** Nothing is added, nothing is switched on, no field is applied, no energy is supplied, and no material property is exploited beyond being a boundary. The geometry does all of it, and the consequence is measurable to a few per cent in a laboratory.

That is the framework's central architectural claim in its purest physical form, and this section presents the effect as exactly that and nothing more.

33.6.1 6.1 What is established

The Casimir result [Casimir, 1948]. For two parallel, perfectly conducting, uncharged plates in vacuum at zero temperature, separated by `a`:

```
E(a) / A  =  − pi^2 ∗ hbar ∗ c / (720 ∗ a^3)          energy per unit area
P(a)      =    pi^2 ∗ hbar ∗ c / (240 ∗ a^4)          attractive pressure
```

It has been measured. A torsion-pendulum measurement over 0.6–6 micrometres agreed with the prediction to about 5 per cent [Lamoreaux, 1997]; an atomic-force-microscope measurement over 0.1–0.9 micrometres reached about 1 per cent [Mohideen and Roy, 1998]; parallel-plate geometries have since been measured directly. The generalisation to real materials at finite temperature is Lifshitz theory, which requires measured optical data rather than an idealised reflectivity, and the whole family of fluctuation forces belongs to the same fluctuation–dissipation framework [Callen and Welton, 1951].

The theorem that actually matters here is not about force.

> **Theorem 6.1 (boundary mode selection).** For a wave equation with a Dirichlet or Neumann condition on a boundary, the set of allowed modes is determined by the boundary geometry alone. Between two parallel boundaries separated by `a`, the allowed perpendicular wavenumbers are `k_n = n pi / a` and the allowed wavelengths are `2a/n`, `n = 1, 2, 3, ...`; nothing else is allowed. *Status:* **established mathematics**, the elementary solution of the Helmholtz problem, and the same theorem for the electromagnetic field, the pressure field in air, and any other field obeying the same equation.

> **Postulate 6.1 (mode selection as a design principle).** *Badged as a postulate, in the source and here.* Boundary geometry alone selects which modes a system supports, at no ongoing cost, and this may be used as a design principle. **The postulate is exact where the system is governed by the Helmholtz problem — electromagnetic modes between conductors, acoustic modes in an enclosure — and is a design heuristic with no formal warrant anywhere else**, including its extension to occupancy and circulation. Section 6.11 separates the three transfer strengths and Section 11 records the extension as unsupported.

33.6.2 6.2 The derivation, in outline

The point of the outline is to make visible that the entire calculation is a difference between two mode counts.

```
STEP 1   Count the modes allowed BETWEEN the plates.
         Perpendicular: k_z = n*pi/a, n = 1, 2, 3, ...  (discrete)
         Parallel:      k_x, k_y continuous
         Two polarisations for n >= 1.

STEP 2   Assign each mode its zero-point energy hbar*omega/2, with
         omega = c * sqrt(k_x^2 + k_y^2 + (n*pi/a)^2).

STEP 3   Sum. The sum diverges.

STEP 4   Subtract the corresponding sum for free space (no plates), which
         also diverges. The DIFFERENCE is finite.

STEP 5   Regularise to do this carefully. The zeta-function route reduces
         the perpendicular sum to sum over n of n^3, whose regularised value
         is zeta(-3) = 1/120.

STEP 6   Result:  E(a)/A = - pi^2 hbar c / (720 a^3)

STEP 7   Force = - d/da [E/A] = pi^2 hbar c / (240 a^4), attractive.
```

The bookkeeping of steps 5 to 7, verified:

```
pi^2 / 720          =  0.013707783890
pi^2 / (6 * 120)    =  0.013707783890        <- identical
1/120 = zeta(-3)

720 = 6 x 120. The 120 in the denominator is the regularised value of a
divergent sum, carried through the calculation and still visible in the
final constant.

720 / 3 = 240. The 3 is the exponent in a^-3 differentiated.
```

What is worth taking from the derivation is not the regularisation technique. It is Step 1. Everything downstream is bookkeeping on a mode count that was fixed the moment the separation was chosen.

33.6.3 6.3 The pressure table

```
P(a) = pi^2 * hbar * c / (240 * a^4) = 1.300126e-27 / a^4   pascals

separation      P (Pa)         P (atm)        F on 1 cm^2 (N)   equivalent mass

  1 nm         1.3001e+09     1.2831e+04       1.3001e+05       1.33e+04 kg
 10 nm         1.3001e+05     1.2831e+00       1.3001e+01       1.33e+00 kg
100 nm         1.3001e+01     1.2831e-04       1.3001e-03       1.33e-04 kg
  1 um         1.3001e-03     1.2831e-08       1.3001e-07       1.33e-08 kg
 10 um         1.3001e-07     1.2831e-12       1.3001e-11       1.33e-12 kg
100 um         1.3001e-11     1.2831e-16       1.3001e-15       1.33e-16 kg
  1 mm         1.3001e-15     1.2831e-20       1.3001e-19       1.33e-20 kg
```

```
1 cm          1.3001e-19     1.2831e-24       1.3001e-23       1.33e-24 kg
1 m           1.3001e-27     1.2831e-32       1.3001e-31       1.33e-32 kg
3 m           1.6051e-29     1.5841e-34       1.6051e-33       1.64e-34 kg
```

The values are for perfectly conducting plates in vacuum at zero temperature, and the assumptions belong to the number wherever it is quoted. At ten nanometres the pressure is 1.28 atmospheres — a thoroughly macroscopic figure, and the reason the effect is a genuine nuisance in nanofabrication. One decade further out it has fallen by a factor of ten thousand to thirteen pascals. **The inverse fourth power is the whole story: each factor of ten in separation costs a factor of ten thousand in pressure**, and the exponent is derived, not fitted.

```
separation x 2    ->  pressure x 6.250e-02
separation x 10   ->  pressure x 1.000e-04
separation x 100  ->  pressure x 1.000e-08

to reach one atmosphere of Casimir pressure the gap must close to:
    a = (K/P)^(1/4) = 1.0643e-08 m = 10.64 nm
```

Ten and a half nanometres is roughly two lipid bilayers. That is how close two mirrors must be before the vacuum pushes on them as hard as the air does.

33.6.4 6.4 The two corrections, and one of them goes the wrong way

Finite conductivity. Real metals are not perfect mirrors. Gold has a plasma wavelength of about 136 nanometres, so at separations at or below that, high-frequency modes penetrate the surface rather than reflecting, and the force is reduced. **For gold near 100 nm the reduction is of order a factor of two, and it grows at smaller separations.** Any serious comparison with experiment uses Lifshitz theory with measured optical data, not the idealised constant.

Temperature. Casimir's formula is the zero-temperature limit, holding while the separation is small compared with the thermal wavelength `lambda_T = hbar c / (k_B T)`:

```
T = 300 K (room)   lambda_T = 7.6329e-06 m =   7.63 um
T = 310 K (body)   lambda_T = 7.3867e-06 m =   7.39 um
T =   4 K (LHe)    lambda_T = 5.7247e-04 m = 572.47 um
```

Beyond `lambda_T` the zero-temperature formula fails and the high-temperature limit applies:

```
P_highT  =  zeta(3) * k_B * T / (4 * pi * a^3)

a           P (T=0 formula)   P (high-T, 300 K)   ratio highT / T0
1 um         1.3001e-03 Pa     3.9620e-04 Pa       0.305
7.63 um      3.8361e-07 Pa     8.9196e-07 Pa       2.325
100 um       1.3001e-11 Pa     3.9620e-10 Pa       30.47
1 mm         1.3001e-15 Pa     3.9620e-13 Pa       304.7
3 m          1.6051e-29 Pa     1.4674e-23 Pa       9.142e+05
```

This correction must be reported even though it damages the tidiness of the table, because at large separations it makes the force larger, not smaller. At one millimetre the true room-temperature pressure is about three hundred times the zero-temperature value; at room scale it is nine hundred thousand times larger. **And it changes nothing:** `1.47e-23` pascals at three metres instead of `1.61e-29`. Both are nothing. **The correction is reported because reporting an inconvenient correction is how a reader learns that the convenient ones were reported honestly too.**

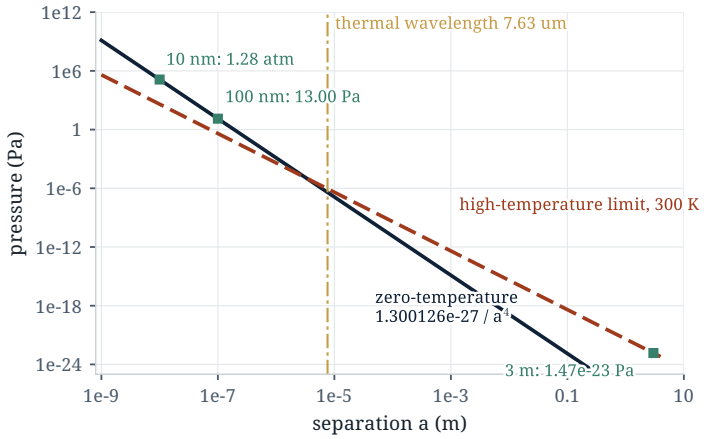


Figure 33.3: **Every decade of separation costs four decades of pressure, and the thermal correction — which makes the force larger — does not rescue it.** The zero-temperature Casimir pressure `1.300126e-27/a^4` Pa and the high-temperature limit `zeta(3) k_B T/(4 pi a^3)` at 300 K cross at the thermal wavelength `lambda_T = 7.63` um; at 3 m the high-temperature value of `1.4674e-23` Pa is `9.142e5` times the zero-temperature value and still `5.9e4` times below what a single bacterium resting on the wall exerts. Computed from the tables of Sections 6.3 and 6.4; generated by `figs/src/CCP-033_fig03.py`.

33.6.5 6.5 The honest limit, computed three ways

Comparison one: a cell.

```
physiological osmotic pressure (300 mOsm at 310 K)    7.732e+05 Pa
atmospheric pressure                                   1.013e+05 Pa
Casimir, IDEAL MIRRORS, at 100 nm                      1.300e+01 Pa
    ratio osmotic / Casimir(ideal, 100 nm)             5.947e+04
```

But ideal mirrors is the wrong model, and using it here would be dishonest in the framework's own favour. Biological surfaces across water are dielectrics with small dielectric contrast, and the correct framework is Lifshitz theory with a Hamaker constant:

```
non-retarded van der Waals pressure  P = A_H / (6 * pi * a^3)

lipid / water / lipid,     A_H = 5e-21 J   at  10 nm :  2.653e+02 Pa
                                           at 100 nm :  2.653e-01 Pa
protein / water / protein, A_H = 1e-20 J   at  10 nm :  5.305e+02 Pa
                                           at 100 nm :  5.305e-01 Pa
```

Two further reductions apply above ten nanometres and both are real: retardation, which is the crossover into the Casimir regime proper and suppresses the force further as separation grows; and screening of the zero-frequency term by the electrolyte.

```
Debye screening length, 150 mM 1:1 electrolyte, 310 K, eps_r = 78.4
    kappa^-1 = 8.004e-10 m = 0.80 nm
```

Eight angstroms. In a cell the static part of the interaction is screened out within one nanometre. Now the force comparison that decides it — a 10 nm by 10 nm patch across a 10 nm gap:

```
ideal-mirror Casimir            1.300e-11 N  =  13.00 pN
realistic biological vdW        2.653e-14 N  =   0.027 pN
kinesin stall force (measured)  6.0e-12 N    =   6.00 pN
kT / nm at 310 K                4.280e-12 N  =   4.28 pN
```

The realistic biological value is two hundred times smaller than a single motor protein's stall force and one hundred and sixty times smaller than the thermal force scale over a nanometre. It is noise beneath noise. The ideal-mirror number that looks competitive is unattainable, because nothing in a cell is a parallel conducting mirror and nothing in a cell holds a fixed 10 nm gap in vacuum.

The precise statement: fluctuation forces of this family — van der Waals in the non-retarded regime — are genuinely important in colloid science and in membrane interactions, and DLVO theory is built on them. The retarded Casimir regime proper contributes nothing measurable to cell biology. The two sentences must not be traded for one another.

Comparison two: a room.

```
two parallel walls, 3 m apart, 10 m^2 each

P (T=0 formula)     1.605e-29 Pa   ->  F = 1.605e-28 N
P (high-T, 300 K)   1.467e-23 Pa   ->  F = 1.467e-22 N     <- the correct one
weight of one E. coli (~1e-15 kg)   ->      9.807e-15 N
ratio: bacterium weight / Casimir force  =  6.683e+07
```

A single bacterium sitting on the wall weighs sixty-seven million times more than the entire Casimir force between the two walls of the room. Every claim about Casimir forces in architecture dies here, computed, in one line.

Comparison three: where it genuinely matters, which is neither of the above.

```
MEMS / NEMS, ideal-mirror values

gap  10 nm : P = 1.3001e+05 Pa ; on a 100 um x 100 um plate F = 1.300e-03 N
gap  50 nm : P = 2.0802e+02 Pa ;                            F = 2.080e-06 N
gap 100 nm : P = 1.3001e+01 Pa ;                            F = 1.300e-07 N = 130 nN
gap 200 nm : P = 8.1258e-01 Pa ;                            F = 8.126e-09 N = 8.1 nN
```

At these scales the force is a design constraint: it contributes to stiction, the permanent collapse of a micromachined structure onto its substrate, and it has been used deliberately to actuate a device. **The honest picture is a narrow band of relevance: below about ten nanometres van der Waals and chemistry dominate; above about a micrometre the force is gone; in between, and only for engineered near-mirror surfaces, it matters.** Biology and architecture are both outside that band by many orders of magnitude. The one other domain in which the effect is legitimately invoked is theoretical rather than applied — negative-energy-density constructions in general relativity take the Casimir vacuum as their existence proof [Morris and Thorne, 1988] — and that is a use in fundamental physics, not in engineering. Appendix B, audit B6, records that the precedent column for this claim is empty and prints the empty cell.

33.6.6 6.6 The real point: geometry alone selects the modes

Set the force aside entirely and look only at Step 1. Two plates a distance `a` apart. The allowed perpendicular wavelengths are `2a/n` and nothing else.

```
LOWEST ALLOWED ELECTROMAGNETIC MODE BETWEEN PLATES,  f_1 = c / (2a)
```

```
a =  10 nm    lambda_1 = 2.00e-08 m   f_1 = 1.4990e+16 Hz   61.99 eV
a = 100 nm    lambda_1 = 2.00e-07 m   f_1 = 1.4990e+15 Hz    6.20 eV
a =   1 um    lambda_1 = 2.00e-06 m   f_1 = 1.4990e+14 Hz    0.62 eV
a =   1 mm    lambda_1 = 2.00e-03 m   f_1 = 1.4990e+11 Hz    6.20e-04 eV
a =   3 m     lambda_1 = 6.00e+00 m   f_1 = 4.9965e+07 Hz    2.07e-07 eV
```

Setting the plates a hundred nanometres apart forbids every electromagnetic mode below six electron volts from existing perpendicular to them. No filter was installed. No field was applied. The exclusion is what the separation is.

The measurability is what makes this a proof rather than a definition. The excluded modes are not merely absent from a bookkeeping exercise; their absence produces a force that a torsion balance registers [Lamoreaux, 1997, Mohideen and Roy, 1998]. **Casimir is the demonstration that mode selection by boundary geometry is a physical fact about a system and not a way of describing it, with a number attached and an instrument that reads the number.**

33.6.7 6.7 The same boundary condition in air: a derivation, not an analogy

Nothing in Step 1 was quantum. Replace the electromagnetic field with the pressure field in air and the speed of light with the speed of sound, and the identical boundary-value problem gives the identical mode ladder. **This is category 1 in the taxonomy of Section 3.3 — a derivation, because the premises are satisfied — and it is the only category-1 transfer in the paper.**

```
SAME FORMULA, DIFFERENT WAVE SPEED:   f_1 = v / (2 * separation)

EM,    plates at 100 nm   ->  1.4990e+15 Hz
EM,    plates at   1 um   ->  1.4990e+14 Hz
SOUND, walls  at   3 m    ->      57.1667 Hz
SOUND, walls  at   7 m    ->      24.5000 Hz

ratio of wave speeds, c / c_sound = 8.7403e+05
```

The three-dimensional version is the standard room-mode equation of architectural acoustics [Rossing et al., 2002, Beranek, 2004]:

```
f(nx, ny, nz) = (c_sound / 2) * sqrt( (nx/Lx)^2 + (ny/Ly)^2 + (nz/Lz)^2 )
```

Computed for a 7.0 x 5.4 x 3.2 metre room (`V = 120.96` m^3, `S = 154.96` m^2):

```
 24.50 Hz  (1,0,0)  axial          66.94 Hz  (1,1,1)  oblique
 31.76 Hz  (0,1,0)  axial          68.08 Hz  (1,2,0)  tangential
 40.11 Hz  (1,1,0)  tangential     72.62 Hz  (2,0,1)  tangential
 49.00 Hz  (2,0,0)  axial          73.50 Hz  (3,0,0)  axial
 53.59 Hz  (0,0,1)  axial          79.26 Hz  (2,1,1)  oblique
 58.39 Hz  (2,1,0)  tangential     80.07 Hz  (3,1,0)  tangential
 58.93 Hz  (1,0,1)  tangential     80.22 Hz  (2,2,0)  tangential
 62.30 Hz  (0,1,1)  tangential
 63.52 Hz  (0,2,0)  axial

modes below 100 Hz : 25
mean spacing       : 3.08 Hz
largest gap        : 8.89 Hz
```

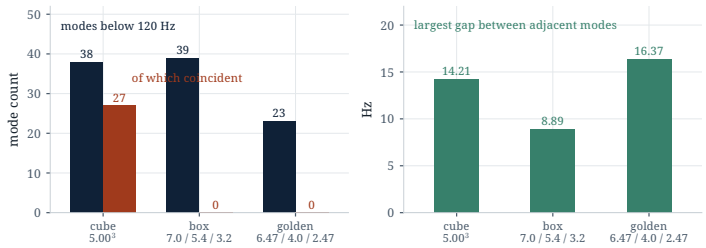


Figure 33.4: **Three rooms of comparable scale, and the degeneracy count is what separates them.** A 5.00 m cube puts 27 of its 38 modes below 120 Hz on top of another mode; the 7.00 x 5.40 x 3.20 m box and the 6.47 x 4.00 x 2.47 m golden-section box each stack none — but the golden box supports only 23 modes and has a largest gap of 16.37 Hz against the plain box's 8.89 Hz, so it is not simply better. Computed from `f = (c/2)sqrt((nx/Lx)^2+(ny/Ly)^2+(nz/Lz)^2)` with `c = 343` m/s, the dimensions of Section 6.7; generated by `figs/src/CCP-033_fig04.py`.

Those twenty-five frequencies are a property of three numbers — the room's dimensions — and of nothing else. Not the paint, not the furniture, not the budget, not the occupants. The reverberation and boundary-frequency apparatus that accompanies them is equally standard [Sabine, 1922, Schroeder, 1965, von Helmholtz, 1885], and the corpus's own treatment adds the switching rule the textbooks leave implicit: **use Sabine below a mean absorption of 0.25 and Eyring above it, state which, and expect a divergence of 77.8 per cent at 2 kHz in an overtreated room if the choice is made carelessly** [Curzi, 2026e].

The strongest computed demonstration of Theorem 6.1 is the degeneracy comparison:

```
MODES BELOW 120 Hz, and how many COINCIDE (degenerate, stacking energy at one pitch)

cube    5.00 x 5.00 x 5.00 m   38 modes,  27 coincident,  max gap 14.21 Hz
box     7.00 x 5.40 x 3.20 m   39 modes,   0 coincident,  max gap  8.89 Hz
golden  6.47 x 4.00 x 2.47 m   23 modes,   0 coincident,  max gap 16.37 Hz
```

A cube of the same rough scale stacks twenty-seven of its thirty-eight low modes on top of each other. The rectangular box stacks none. That is the audible difference between a room that booms and a room that does not, and it is produced by choosing three numbers.

And the golden-section box is not simply better. It has no degeneracies, but it is smaller, so it has fewer modes below 120 Hz — 23 against 39 — and its largest gap is wider, 16.37 Hz against 8.89 Hz, which below the Schroeder frequency means less even low-frequency support. **Proportion ratios generate candidate geometries. They do not certify them.** The golden section attracts claims that survive poorly on inspection, and the sceptical literature on those claims is the relevant one here [Markowsky, 1992, Fischler, 1981, Livio, 2002]; the corpus cites it wherever the ratio appears, on the principle that a claim's credibility is set by what its author already knows about it.

Appendix B, audit B7, records the evidence and the precedent for boundary mode selection separately, and it is the strongest precedent entry in the paper — which makes it the most dangerous, because a strong precedent is more tempting to promote than a weak one.

33.6.8 6.8 Pitch-class folding of a mode ladder

A mode ladder between two boundaries is `f_n = n f_1`. Doubling the index doubles the frequency, which is one octave, and octave equivalence as a perceptual structure for pitch is established [Shepard, 1982, Demany et al., 2013, Burns, 1999]. Pitch class then depends only on the odd part of the index, `n = 2^a m` with `m` odd:

```
first  6 modes -> odd parts {1,3,5}                          3 classes  (ceil(6/2)=3)
first 12 modes -> odd parts {1,3,5,7,9,11}                   6 classes  (ceil(12/2)=6)
first 16 modes -> odd parts {1,3,5,7,9,11,13,15}             8 classes  (ceil(16/2)=8)
first 24 modes -> odd parts {1,3,5,...,23}                  12 classes  (ceil(24/2)=12)
```

The first twelve modes between any two boundaries carry only six distinct pitch classes. That is a hard structural limit on the harmonic vocabulary of an enclosure, set by the boundary and by nothing else. The corresponding minimum-gap schedule for a ladder of given length is exact and rational — 1200 cents at `N = 2`, 203.9100 cents (9/8) at `N = 3` to 6, 34.9756 cents (50/49) at `N = 7` to 16, 6.0008 cents (289/288) at `N >= 17`, and 1.0296 cents (1682/1681) at `N >= 41` — and it is computed in Curzi [2026].

One difference must be stated so the two systems are not conflated. For a mode ladder `f_n = n f_1`, indices `n` and `2n` are **one** octave apart. For the corpus's essence law `f(n) = k n^2`, `(2n)^2 = 4 n^2`, so `n` and `2n` are **two** octaves apart. The odd-part folding rule is identical in both; the octave count is not, and a room's mode list must not be read off the essence tables or the reverse. That the essence register carries a further and more serious limit is stated by the paper that owns it: **the frequency register cannot distinguish a mood from a trait at all, because `f = k n^2` contains no time term; and the one-hot encoding of a framing, which looks like the cleanest available, returns the system's maximum surplus for every subject and therefore measures nothing** [Curzi, 2026]. That result bounds what any mode-ladder analogy in this paper may be taken to license.

33.6.9 6.9 Nothing is added, nothing is powered

The framework's cost argument for geometric function is:

```
Powered function:   C_t = c_energy + c_maint + c_replace + c_failure , every period
Geometric function: C_t = c_clean                                    , every period
                    plus a one-time capital charge at t = 0
```

The Casimir effect is the limiting case of the second line. **Two plates, once placed, exert their force forever, with `c_energy = 0`, `c_maint = 0`, `c_replace = 0`, and a `delta` set by material durability alone.** There is no switch and nothing to fail.

The other half of that statement is what keeps it from becoming a free-energy claim: the force is conservative. Letting the plates come together releases energy; pulling them apart costs exactly that energy back. There is no cycle, no net extraction and no engine. A shape performs work for free; a shape does not produce energy for free.

The corpus's own priced version of geometric function is the threshold-as-impedance-match result: worst-junction reflection across a 1-to-100 character change falls from 96 per cent to 13 per cent when the change is divided into six equal-ratio steps, which is a design gain obtained by geometry and paid for once [Curzi, 2026f]. That is the same cost structure, in a domain where the framework can measure it.

33.6.10 6.10 Geometry dependence, repulsion, and the moving boundary

If boundary geometry is the whole story, then changing the geometry should change the answer in ways not derivable by scaling the parallel-plate result. It does.

The sign is not fixed. For a perfectly conducting spherical shell the Casimir energy has the opposite sign to the parallel-plate case: the shell tends to expand rather than contract. **A force that is attractive between plates and repulsive for a sphere is not a property of the material. It is a property of the shape.**

Repulsion has been measured, between a gold sphere and a silica plate immersed in bromobenzene, in the configuration Lifshitz theory predicts to be repulsive — when the intervening fluid's dielectric response lies between those of the two surfaces. **The direction of the**

force was reversed by changing the arrangement of the boundaries and the medium. Nothing else changed.

Non-additivity. Casimir forces do not add pairwise. The proximity-force approximation, which treats a curved surface as a stack of parallel-plate elements, is only an approximation and fails for strongly structured geometries; the Casimir piston is the exactly solvable geometry that makes this vivid, and corrugations, gratings and lateral offsets produce forces including lateral components that no parallel-plate calculation predicts. **The geometry is not a parameter in the calculation. The geometry is the calculation.**

The dynamical case. A moving boundary should not merely select modes but mix them, converting virtual excitations into real photons. The requirement is severe: the mirror must move at speeds comparable to the phase velocity of the field, which is unachievable mechanically. The effect was observed by moving the boundary without moving anything material — a superconducting waveguide terminated by a flux-tunable element acting as an effective electrical boundary, modulated at gigahertz rates, producing the predicted two-mode-squeezed photon output. **The physical content is the same content as the static case, differentiated once: a static boundary determines which modes exist; a boundary that changes in time pumps energy between them.**

Here, uniquely, an architectural transfer can be stated honestly. A room's variable elements — a movable panel, an opened door, a drawn curtain — change the mode structure, and the change is what an occupant hears as the room opening up. **A boundary that moves does something a boundary that sits still cannot. That statement is true in a superconducting waveguide and true in a room, for the same reason: the mode structure is a function of the boundary, so a time-dependent boundary is a time-dependent mode structure.** It is not a claim that a moving wall creates photons; it is a claim about the functional dependence, and the functional dependence is the same.

33.6.11 6.11 What the Casimir effect establishes, and the three transfer strengths

```
WHAT CASIMIR ESTABLISHES

 1. Boundary geometry alone determines which field modes can exist.        PROVEN
 2. That determination has measurable physical consequence.               MEASURED
 3. Nothing is added, powered, or applied to produce it.                   PROVEN
 4. The effect depends on the SHAPE, not only on the size: the sign
    itself is geometry-dependent, and repulsion has been measured in
    the predicted configuration.                                          MEASURED
 5. A time-dependent boundary pumps modes, observed with an
    effective, electrically tuned mirror.                                 MEASURED

WHAT IT DOES NOT ESTABLISH

 6. That the force matters at biological scale.        IT DOES NOT. Sec 6.5.
 7. That the force matters at architectural scale.     IT DOES NOT. Sec 6.5.
 8. That energy can be extracted from the vacuum.      IT CANNOT. Sec 6.9.
 9. That "everything is vibration" in any usable
    sense, or that intention shapes fields.           NOT A CLAIM OF PHYSICS.
10. That the mode-selection principle CERTIFIES any
    particular architectural geometry.                IT DOES NOT.
```

The three transfer strengths, which must be distinguished in print every time:

```
(a) EM modes between plates          -- proven and measured
(b) acoustic modes in a room         -- the SAME equation, measured
```

```
(c) occupancy / circulation patterns -- SPECULATIVE, badged, and it must
                                       never inherit the credibility of (a)
```

The one thing the framework takes from the Casimir effect is that the shape is the function — proven for electromagnetic modes between plates, proven for acoustic modes in rooms because it is the same equation, and offered as a design heuristic everywhere else with its badge on. That form and function are related by geometry rather than by mechanism is an old claim in a respectable literature [Thompson, 1917]; what the Casimir effect adds is an instrument.

33.7 7 Failure modes: a register

The three preceding sections each generate a characteristic set of errors. They are collected here as a register because the register, rather than any individual number, is what a reviewer or an editor can apply. Each entry gives the signature, a diagnostic that can be run in one step, and the remedy. Numbering follows the section that generates the fault.

33.7.1 7.1 Faults in the use of the formalism

FM-4.1 The observer equivocation. *Signature:* a text uses "observer" in the technical sense in one sentence and in the mentalistic sense in the next, without marking the shift. *Diagnostic:* substitute "photodiode" for every occurrence of "observer"; if the argument still works it was technical, and if it collapses into nonsense it was running on the equivocation. *Remedy:* define the term at first use in every document, and apply the substitution test as a review gate. This is the load-bearing move in nearly every quantum-manifestation claim in circulation.

FM-4.2 Superposition as metaphor sold as physics. *Signature:* "hold both possibilities in superposition until you choose." *Diagnostic:* **the tell is that no interference term is ever computed, because there is nothing to compute it from.** *Remedy:* permitted as an explicitly flagged metaphor, in the same sentence as the flag. The honest version already exists and is better — a weighted branch set with an effective count (Section 4.10).

FM-4.3 The interpretation cherry-pick. *Signature:* Everett invoked to license "all possibilities are real", and consciousness-collapse invoked a paragraph later to license "your observation selects one." *Diagnostic:* **these are incompatible — Everett has no collapse and therefore no selection; consciousness-collapse has one world and therefore no other branches.** *Remedy:* name the interpretation, name what it abandons, and check consistency across the whole document rather than sentence by sentence.

FM-4.4 The Born rule as wish-weighting. *Signature:* the amplitude treated as something an agent can adjust. *Diagnostic:* **there is no operation in the formalism that raises an amplitude by preference; amplitudes change under unitary evolution generated by a Hamiltonian, and no Hamiltonian contains a term for wanting.** *Remedy:* demand the Hamiltonian term. The channel requirement of Curzi [2026ab] applies verbatim, and its forbidding clause — no channel, no effect — disposes of the claim without any appeal to taste.

FM-4.5 Entanglement as telepathy. *Signature:* "entangled particles communicate instantaneously; therefore minds can." *Diagnostic:* the first clause is already false; correlated results are recovered only by comparing records through a classical channel, and before that comparison each side's statistics are exactly what they would have been with no entanglement at all. *Remedy:* cite no-signalling by name, every time. It is the shortest complete refutation available and it is interpretation-independent.

FM-4.6 Precedent presented as evidence. *Signature:* an old text quoted alongside a modern result and the pair described as confirmation. *Diagnostic:* ask which column the old text sits in. *Remedy:* the audit format of Appendix B, in which the precedent row is labelled as precedent and the neutral row states what each column can carry. **This is the failure most**

likely to discredit the corpus's comparative material, and Appendix B, audit B3, is the worked case of a bridge rejected in print.

FM-4.7 The linearity leak. *Signature:* a proposal quietly requires non-linear evolution, usually in order to obtain signalling or preference-sensitive dynamics. *Diagnostic:* does the proposal claim that `U(a|psi_1> + b|psi_2>)` is anything other than `aU|psi_1> + bU|psi_2>`? *Remedy:* **non-linear modifications of quantum mechanics generically permit superluminal signalling and are strongly constrained experimentally.** The proposal is not merely unsupported; it is in tension with results already in hand, and it must be presented as a different theory.

FM-4.8 Collapse by attention. *Signature:* the quantum Zeno effect — repeated rapid measurement inhibiting transitions, a real measured phenomenon — cited as evidence that watching changes outcomes, and thence that mental attention does. *Diagnostic:* ask what the measuring interaction is and at what rate it is applied. **The effect requires repeated physical interaction faster than the transition timescale; it is done with lasers and photodetectors, and nothing about it involves noticing.** *Remedy:* if the interaction cannot be named in physical units, the citation is decorative.

33.7.2 7.2 Faults in coherence claims

FM-5.1 The isolated-system fallacy. *Signature:* a quantum argument about a biological system with no statement of what that system is coupled to. Textbooks solve isolated problems, and the reader imports the idealisation without noticing it was one. *Diagnostic:* what is the environment, and what is the coupling? *Remedy:* fields (i) and (ii) of the gate, enforced as a submission requirement. A proposal that cannot name the environment has not begun.

FM-5.2 The missing timescale. *Signature:* a coherence claim with no number attached. *Diagnostic:* the absence of a single number, which is often the point. *Remedy:* fields (iii) to (v). **In every case examined for this paper the ratio came out between `1e9` and `1e16`.**

FM-5.3 Decoherence presented as a solution. *Signature:* "decoherence solved the measurement problem." *Diagnostic:* ask which diagonal element changed. None did. *Remedy:* state that decoherence explains the absence of macroscopic interference and the preferred basis and does not select an outcome; and note the logical point, which is decisive — **decoherence is derived from M2, so it cannot resolve a contradiction caused by M2.**

FM-5.4 The exception generalised. *Signature:* "quantum biology is real, photosynthesis proves it, therefore quantum consciousness is plausible." Two true clauses and a non-sequitur. *Diagnostic:* run the table of Section 5.10 — spatial ratio ~`1e4`, temporal ratio ~`1e9`, and a degree-of-freedom row that changes from a spin or an exciton in a rigid scaffold to a charge displacement across a wet membrane. *Remedy:* note the fourth row: the exception cases have an evolved protective scaffold and cognition has no proposed one.

FM-5.5 The contested result cited as settled. *Signature:* the 2007 photosynthesis result cited without the vibronic reinterpretation, or the 2021 cryptochrome result cited as though the avian compass were established in vivo rather than as an in-vitro protein measurement. *Diagnostic:* is the contest named in the same passage? *Remedy:* cite the contest with the claim, every time. **This is a citation-integrity failure and it is the kind a hostile reviewer finds first.**

FM-5.6 The prefactor objection. *Signature:* a defender attacks the cross-section, the geometry or the scattering model — correctly noting that these are approximate — and treats the approximation as though it rescued the claim. *Diagnostic:* ask how many orders of magnitude the correction is worth; honest answers are one to three. Then ask how many are needed; the answer is nine to sixteen. *Remedy:* **concede the prefactor and keep the conclusion, because the conclusion never depended on the prefactor.** This is H2 of Section 2 and it is stated proactively so that the objection is answered before it is raised.

FM-5.7 Coherence as institutional metaphor, unbadged. *Signature:* this framework's own risk. "The team lost coherence"; "the institution decohered" — used without a badge and drifting from analogy to explanation across successive drafts. *Diagnostic:* search every occurrence and check which category of Section 3.3 is attached. *Remedy:* badge it in the sentence, or use `delta` and say what is actually being lost.

FM-5.8 The cold-system bait and switch. *Signature:* a result obtained at 77 K, at 4 K, or in ultra-high vacuum cited as evidence for what happens at 310 K in water. *Diagnostic:* check the temperature and pressure in the methods section of the cited paper. **The `T^9` scaling makes 77 K easier than 300 K by a factor of `2.07e5` for the photon channel alone, and removing the gas removes the collisional channel entirely.** *Remedy:* require the operating conditions of every cited coherence measurement to be stated alongside the result.

33.7.3 7.3 Faults in the use of the Casimir effect

FM-6.1 Casimir invoked as a biological force. *Signature:* a claim that vacuum fluctuation forces do work in cells, structure water, or organise membranes, sometimes dressed with the 10 nm row of the pressure table, which does look impressive at 1.28 atmospheres. *Diagnostic:* three questions — what is the separation, are the surfaces mirrors, is the medium vacuum? For biology the answers are: more than 10 nm, no, and no. *Remedy:* quote the honest limit verbatim and give the ratio to osmotic pressure, `5.947e4`, computed with the *ideal* value so that the concession is maximally generous and still decisive.

FM-6.2 Zero-point energy extraction. *Signature:* "the Casimir effect proves the vacuum contains energy; therefore it can be tapped." *Diagnostic:* ask for the cycle. Plates in, energy out; plates back out, and the energy is returned. There is no third step. *Remedy:* refuse in the plain words — this is a perpetual motion machine.

FM-6.3 The ideal-mirror number quoted in a real setting. *Signature:* `1.3e5` Pa at 10 nm quoted without the phrase "perfectly conducting plates in vacuum at zero temperature." *Diagnostic:* is the assumption attached to the number? *Remedy:* attach it in the same line of the table. **The finite-conductivity correction is a factor of order two for gold near 100 nm and worse below; the biological correction is one to three orders; quoting the ideal number in a real setting overstates by between two and three orders of magnitude.**

FM-6.4 The analogy upgraded by reuse. *Signature:* the framework's own characteristic risk. "The room selects modes just as the plates do" is exact; three drafts later it has become "the room selects social modes just as the plates do", and the badge has fallen off in transit. *Diagnostic:* search every occurrence of the mode-selection claim and check which of the three transfer strengths is attached. *Remedy:* apply the check on every draft, not once at the start.

FM-6.5 Proportion mysticism. *Signature:* a room specified by a ratio system and declared acoustically correct on that basis. *Diagnostic:* compute the mode list. Section 6.7 does exactly this: **the golden box has no degeneracies but only 23 modes below 120 Hz against the plain box's 39, and a largest gap of 16.37 Hz against 8.89 Hz — its proportions are more admired and its low-frequency support is worse.** *Remedy:* generate with the ratio system, certify with the computation [Markowsky, 1992].

FM-6.6 The cube. *Signature:* a cubic or near-cubic room, usually chosen for plan efficiency or symmetry. *Diagnostic:* a 5 m cube has 27 of its 38 modes below 120 Hz coincident with another mode, so energy stacks at those few frequencies and the room booms at them. *Remedy:* break the symmetry; any set of three mutually incommensurate dimensions removes the degeneracies. **This is the cheapest acoustic intervention available, because it costs nothing at design time and cannot be retrofitted at any price.**

FM-6.7 Precedent presented as evidence, the monochord case. *Signature:* "the Pythagoreans knew about the Casimir effect." *Diagnostic:* a category error, and the section is unusually exposed to it because its precedent is unusually strong. **A strong precedent is more dangerous than a weak one, not less, because it is more tempting to promote.** *Remedy:* the monochord recorded that boundary position alone sets the mode ladder; that is precedent, it is impressive, and it is not evidence for anything about vacuum fluctuations.

FM-6.8 The static-configuration audit. *Signature:* a room with movable partitions, folding doors or retractable panels has its mode list computed once, in the configuration shown on the drawing. *Diagnostic:* ask how many configurations the room has and compute the list for each. *Remedy:* in practice the worst configuration is the one used most, because it is usually the closed one.

33.8 8 Results: three audits worked end to end

Three proposals are audited against the apparatus of Sections 4 to 7. **All three parameters of each case are stipulated by the proposal or by the worked brief; nothing in this section is a field observation, and no dataset was collected.** The purpose is to show that the procedure runs, that it reaches a disposition, and that it reports the disposition whichever way it falls.

33.8.1 8.1 A proposal resting on the formalism

The proposal, as submitted: *since the future exists as a superposition of possible outcomes, and since the act of conscious observation collapses the wave function, sustained focused belief on a chosen outcome raises its amplitude and therefore its probability of manifesting.* The proposer has read popular physics carefully and is arguing in good faith, so the audit is a real one rather than a dismissal.

Step 1 — separate formalism claims from interpretation claims.

CLAIM	TYPE	VERDICT
(a) the future is a superposition	formalism	FALSE AS STATED
(b) observation collapses the wave function	interpretation	ONE reading only
(c) conscious observation does it	interpretation	MINORITY, no evidence
(d) focused belief raises an amplitude	formalism	NO SUCH OPERATION
(e) higher amplitude -> higher probability	formalism	TRUE (Born rule)

Only (e) is correct, and (e) is the only one doing no work in the argument.

Step 2 — audit (a). A quantum system's state may be a superposition with respect to some observable. A future course of events at human scale is not a quantum state of anything anyone has written down: there is no Hilbert space here, no Hamiltonian, and no observable named. **A claim about superposition that cannot name the Hilbert space is not a quantum claim; it is a probability claim wearing quantum vocabulary, and it should be evaluated as one.** That evaluation is available immediately, in Section 4.10.

Step 3 — audit (b) and (c). (b) is the collapse postulate, one interpretive option out of at least four, abandoning M2; Everett has no collapse, Bohm has no collapse, and GRW's collapse is spontaneous and has nothing to do with observation. Presenting collapse as a fact of nature rather than as a choice is FM-4.3. (c) is von Neumann–Wigner, and all five reasons of Section 4.8 apply. **The evidence column for this claim is empty. The empty cell is printed.**

Step 4 — audit (d), which is where the proposal dies. This step requires no interpretive dispute to be settled. The formalism has exactly two ways for amplitudes to change: unitary evolution generated by a Hamiltonian, and the interpretation-dependent update on measurement. **For belief to raise an amplitude, belief must appear in the Hamiltonian. There is no such term, no proposal for one, and no measured coupling constant to put in front of it.**

```
BELIEF  ->  [ CHANNEL ]  ->  WORLD

proposed channel : quantum amplitude modification by attention
channel exists?  : NO. No coupling, no term, no constant, no measurement.
gain             : undefined -- there is nothing to estimate.
conclusion       : NO CHANNEL, NO EFFECT.
```

The grounding postulate's conditional clause does the entire work here, and it does it without any appeal to taste. The postulate is that belief creates reality if the belief is grounded in truth; in magnitude form it is `E = G B` with the forbidding clause that `G = 0` implies `E = 0` for every `B` [Curzi, 2026ab]. **It is a postulate, badged as one in the paper that owns it and treated as one here.** This audit does not violate it; it applies it, and the application returns a rejection.

Step 5 — the timescale objection. Even granting some quantum mechanism at the neural level, Section 5 applies: a microtubule-scale superposition in water at body temperature has an estimated coherence time of `1.98e-17` seconds against an action potential of about `1e-3` seconds, a ratio of `5.04e13`. **The proposal needs a quantum state to survive fifty trillion times longer than the estimate allows, in the warmest and wettest environment in the argument.**

Step 6 — the honest reconstruction, because there is one. The proposer's practical intent — hold the intended outcome vividly, act consistently with it, and expect it — is not empty, and it has real measured channels.

```
1. Represent the future as a weighted branch set, not a superposition.
   Two branches, weights p and 1-p. Compute N_eff = e^H (Section 4.10).
   No hbar.

2. Name the channel by which belief could act on the branch weights:
     - physiological: expectancy -> endogenous opioid release, and the
       naloxone-reversibility result that establishes the channel
     - behavioural:   belief -> altered action -> altered outcome frequency
     - social:        stated commitment -> others' expectations -> coordination
     - informational: attention -> different sampling -> different information
   Each is a real coupling with a real gain. NONE is quantum.

3. Estimate the gain. Predict a magnitude in advance.

4. Preregister. Freeze the model. Score held-out against four baselines
   including a shuffled-input control.

5. Report the result whichever way it comes out.
```

That the expectancy channel is real and that expectancy effects are measurable is established, and so is the fact that they are easy to over-read [Rosenthal and Jacobson, 1968, Rosenthal and Fode, 1963, McCambridge et al., 2014]; the preregistration discipline that keeps step 4 honest is Nosek et al. [2018], and the failure modes it exists to prevent are Simmons et al. [2011] and Gelman and Loken [2013]. The corpus's own instance of that protocol, applied and **failed**, is Curzi [2026i].

Step 7 — the numbers, carried to the end. Suppose the target is a binary outcome currently judged at `p = 0.25`.

```
p = 0.25   ->  H = 0.562335 nats = 0.811278 bits  ->  N_eff = 1.7548
```

The claim, made specific: the practice raises `p` to 0.40 within twelve weeks. That is now a preregisterable statement with a magnitude, and the branch set has an effective count that can be tracked. If the mechanism is behavioural, the claim is plausible, testable and squarely inside the grounding postulate. If the mechanism is amplitude modification by attention, it is outside, and Step 4 already closed it.

```
FINAL DISPOSITION

quantum mechanism           REJECTED. No channel. FM-4.4, FM-4.1, FM-4.3.
consciousness-collapse      REJECTED as a basis. Minority interpretation, no
                            evidence, empty evidence cell, and its own advocate
                            abandoned it.
weighted-branch practice    ACCEPTED FOR TEST under the preregistration protocol.
behavioural channel         ACCEPTED as a candidate mechanism with a real gain,
                            to be estimated and measured.

The proposal loses its physics and keeps its practice.
```

33.8.2 8.2 A proposal resting on protected coherence

A better-formed proposal, because it names a mechanism: *coherent vibrational states in tubulin dimers within neuronal microtubules persist long enough to be functionally relevant, are protected from decoherence by ordered water and by the hydrophobic pockets of the protein, and provide a substrate on which cognition supervenes.* This is close to a real position with published defenders, and it gets the full five-field audit.

Step 1 — the five required fields.

```
(i)   degree of freedom      : vibrational / dipole states of tubulin dimers
(ii)  environment            : cytoplasmic water at 310 K, dissolved ions
                               (~150 mM), the rest of the lattice, thermal photons
(iii) estimated tau_dec      : to be computed -- Step 2
(iv)  functional timescale   : 1e-3 s (a spike) at the very shortest
(v)   ratio                  : to be computed -- Step 3
```

The proposal is well formed, which is more than most, and that is why it can be audited at all.

Step 2 — compute the estimate. A microtubule segment, taking a 25 nm by 100 nm face as the effective cross-section, in water at 310 K:

```
n     = 3.3428e+28 /m3
v_bar = 603.6 m/s
sigma = 2.500e-15 m2
-------------------------------------
Gamma = n * sigma * v_bar = 5.044e+16 /s
tau   = 1 / Gamma         = 1.982e-17 s
```

Cross-checked against the thermal time: `1.982e-17 / 2.464e-14 = 8.0e-4`. **The estimate is a thousandth of the thermal time at body temperature, which is the expected order for an object of this size in a dense liquid, so it is not an outlier requiring special pleading.**

Step 3 — the ratio. `1e-3 / 1.982e-17 = 5.04e13`, with off-diagonal suppression over one spike of `exp(-5.044e13)`. **The proposal requires the state to survive fifty trillion decoherence times.**

Step 4 — audit the proposed protection. Take the defence seriously and ask what it would have to buy.

- **Ordered water is still water.** Its librational band is at `1.80e13` Hz and its OH stretch at `1.02e14` Hz; hydrogen bonds rearrange on a picosecond. Ordering changes the correlation structure of the bath; it does not decouple the bath.
- **A hydrophobic pocket is nanometres across.** It can shield a small internal degree of freedom. It cannot shield a lattice-scale collective mode, which is what the proposal requires, because the mode extends across the region the pocket does not cover.
- **The quantitative demand is the problem.** Even granting a protection factor of `1e6` — far beyond anything demonstrated for a charge-displacing degree of freedom in water — the remaining shortfall is `5.04e13 / 1e6 = 5.04e7`. **Seven orders of magnitude still missing, after a concession of six that nobody has earned.**

Compare the case where protection genuinely works: the radical pair survives for microseconds because it uses electron spin, which has no electric dipole and is therefore barely watched by a polar solvent. **That is a real protection mechanism and it is structural. The tubulin proposal's degree of freedom is a dipole state — precisely the kind that water monitors best.**

Step 5 — the falsification condition. H1 states what would overturn this section: a replicated demonstration of functional quantum coherence above `1e-6` s in an unrefrigerated, unengineered biological system at 310 K.

```
what the proposal needs        1e-3 s
what H1 would accept as
overturning the account        1e-6 s
what is estimated              1.98e-17 s
what has ever been measured
in a warm biological system    ~3e-13 s (contested; and vibronic on the
                               leading current reading)

the proposal is 11 orders beyond the best CONTESTED measurement, and
3 orders beyond even the threshold set for overturning the account.
```

Step 6 — disposition.

coherence claim	REJECTED. Ratio 5.04e+13. FM-5.1 not triggered (environment named), FM-5.2 not triggered (timescale computable), FM-5.6 anticipated and conceded -- the prefactor is worth 1-3 orders and 13 are needed.
protection mechanism	REJECTED as insufficient by 7 orders of magnitude even after granting a 1e6 protection factor that has not been demonstrated.
the proposal is well-formed	ACKNOWLEDGED. It named its degree of freedom, its environment, and its functional timescale.
what survives	MICROTUBULES ARE INTERESTING FOR CLASSICAL REASONS. They are the cell's mechanical lattice; they transmit force; motor proteins walk along them; and force applied at the cell surface reaches the nucleus through the cytoskeleton and alters transcription. NONE OF THAT REQUIRES COHERENCE, and all of it is measured.

33.8.3 8.3 A brief resting on boundary geometry, in which the principle survives

A 121 cubic metre studio room, 7.0 x 5.4 x 3.2 metres, is to serve as a rehearsal and listening space. The client has read about Casimir forces and asks whether vacuum-energy geometry can improve the acoustics.

Step 1 — answer the Casimir question and close it.

```
Casimir pressure between the two long walls (7.0 m apart):

  T = 0 formula   P = 1.300126e-27 / 7.0^4 = 5.416e-31 Pa
  high-T (300 K)  P = zeta(3) kB T / (4 pi * 7.0^3) = 1.155e-24 Pa   <- correct one

  wall area 5.4 x 3.2 = 17.28 m^2
  total force = 1.155e-24 * 17.28 = 1.996e-23 N

  for comparison: one E. coli on the wall weighs 9.807e-15 N
  ratio = 4.9e+08
```

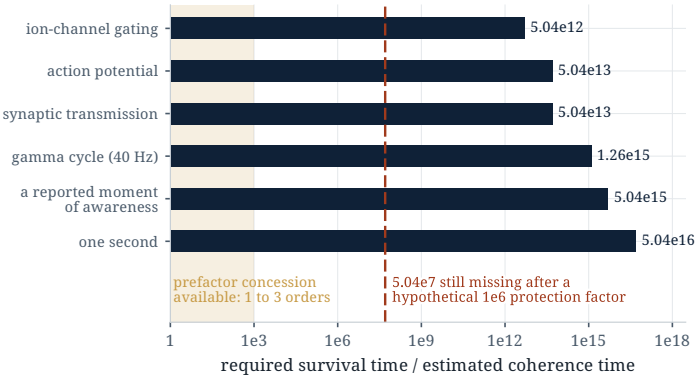


Figure 33.5: **The audit's decisive quantity is a ratio, and every concession available to the proposal leaves it losing.** Required survival time over estimated coherence time, for six neural processes against `tau_dec = 1.98e-17 s`: `5.04e12` for ion-channel gating up to `5.04e16` for one second. The prefactor concession available to a defender is worth one to three orders (shaded band); a hypothetical `1e6` protection factor still leaves `5.04e7` for the action potential. Computed from the ratio table of Section 5.6; generated by `figs/src/CCP-033_fig05.py`.

The entire Casimir force across this room is five hundred million times smaller than the weight of a single bacterium resting on the wall. The answer is no, and it is not a close call.

Step 2 — supply the principle the client was reaching for. The intuition was that geometry alone can do acoustic work without equipment. **That intuition is correct; it was attached to the wrong physics.** The correct attachment is Theorem 6.1, which the Casimir effect establishes for electromagnetic modes and which holds identically for sound because it is the same boundary-value problem.

Step 3 — compute the mode list.

```
f(nx,ny,nz) = (343/2) * sqrt( (nx/7.0)^2 + (ny/5.4)^2 + (nz/3.2)^2 )

 24.50 Hz  (1,0,0) axial        63.52 Hz  (0,2,0) axial
 31.76 Hz  (0,1,0) axial        66.94 Hz  (1,1,1) oblique
 40.11 Hz  (1,1,0) tangential   68.08 Hz  (1,2,0) tangential
 49.00 Hz  (2,0,0) axial        72.62 Hz  (2,0,1) tangential
 53.59 Hz  (0,0,1) axial        73.50 Hz  (3,0,0) axial
 58.39 Hz  (2,1,0) tangential   79.26 Hz  (2,1,1) oblique
 58.93 Hz  (1,0,1) tangential   80.07 Hz  (3,1,0) tangential
 62.30 Hz  (0,1,1) tangential   80.22 Hz  (2,2,0) tangential

REQUIRED REPORT:
  modes below 100 Hz            25
  mean spacing below 100 Hz     3.08 Hz
  largest gap                   8.89 Hz
  coincident (degenerate) modes 0
```

Zero degeneracies, achieved by the choice of three numbers and by nothing else. The alternative the client's architect had proposed, a 5 metre cube of similar volume, has 27 coincident modes out of 38 below 120 Hz.

Step 4 — reverberation and the modal boundary.

```
V = 120.96 m^3 , S = 154.96 m^2 , RT60 = 0.161 V / A
```

```
mean absorption   A (m^2 sabins)   RT60 (s)   Schroeder f (Hz)
  0.10              15.50            1.257       203.9
  0.20              30.99            0.628       144.2
  0.35              54.24            0.359       109.0
```

The Schroeder frequency divides the modal region below, where individual modes are audible as resonances, from the statistical region above, where reverberation theory applies [Schroeder, 1965, Sabine, 1922]. **At the target absorption of 0.35 the Schroeder frequency is 109.0 Hz, and there are 25 discrete modes below 100 Hz with a mean spacing of 3.08 Hz and no coincidences: the modal region is well populated and evenly filled.** That result was obtained from four numbers — three dimensions and one absorption coefficient. **The absorption figure of 0.35 sits above the 0.25 threshold at which Sabine should give way to Eyring, so the RT60 values in the table are Sabine values used above their reliable range and are reported as such** [Curzi, 2026e].

Step 5 — the tuned-vessel option and its honest status. The client asks about the bronze vessels specified in Vitruvius. **Precedent:** a Roman engineer wrote down that vessel geometry selects pitch and used it as a design element. It is not recorded that the theatres sounded better, and the acoustic evidence for the efficacy of these vessels is not established. **Precedent, not evidence. Evidence:** the Helmholtz resonator is real, its formula is measured, and its resonance is a function of geometry alone [von Helmholtz, 1885, Rossing et al., 2002]:

```
f = (c / 2pi) * sqrt( A / (V * L_eff) ),   L_eff = L_neck + 1.7 * r_neck

neck  20 mm, volume  1.0 L, neck length  50 mm  ->  118.21 Hz
neck  30 mm, volume  2.0 L, neck length  60 mm  ->  110.99 Hz
neck  50 mm, volume  8.0 L, neck length  80 mm  ->   77.27 Hz
neck 100 mm, volume 50.0 L, neck length 120 mm  ->   47.79 Hz
```

A fifty-litre vessel with a hundred-millimetre neck resonates at 47.79 Hz — between the (0,1,0) axial mode at 31.76 Hz and the (2,0,0) at 49.00 Hz. Sized and damped as an absorber it targets the low-frequency region where porous absorption is useless, because the required material thickness scales with wavelength. Its honest status: the resonance is established and computed; whether these vessels improve this room requires measurement in this room.

Step 6 — the variable-geometry audit. The room has a folding partition dividing it at 4.0 m.

```
CONFIGURATION A -- open, 7.0 x 5.4 x 3.2
  lowest axial mode 24.50 Hz ; 25 modes below 100 Hz ; 0 degenerate

CONFIGURATION B -- divided, 4.0 x 5.4 x 3.2
  lowest axial mode = 343/(2*4.0) = 42.88 Hz
  the entire 24.50–42.88 Hz region is EXCLUDED. Not attenuated: excluded.
```

Closing the partition does not damp the low end. It abolishes it, by removing the boundary separation that permitted it to exist. That is Theorem 6.1 operating at room scale, at zero cost, in either direction, as many times as the partition is moved.

Step 7 — disposition.

```
Casimir force as a design mechanism    REJECTED. 1.996e-23 N across the room;
                                       4.9e+08 times smaller than a bacterium's
                                       weight.

Casimir as the demonstration of
mode selection by boundary             ACCEPTED, and the reason the client's
                                       intuition was worth taking seriously.
```

```
room dimensions 7.0 x 5.4 x 3.2        ACCEPTED. 25 modes below 100 Hz, mean
                                       spacing 3.08 Hz, largest gap 8.89 Hz,
                                       ZERO degeneracies. The cube alternative
                                       had 27 coincidences out of 38.

absorption target 0.35 mean            ACCEPTED. RT60 = 0.359 s (Sabine, above
                                       its reliable range), Schroeder 109.0 Hz.

Helmholtz absorbers at 47.79 Hz        ACCEPTED FOR PROTOTYPING, with the
                                       obligation to measure in this room
                                       before committing.

Vitruvian vessels as justification     REJECTED AS EVIDENCE, RETAINED AS
                                       PRECEDENT.

variable-geometry mode audit           REQUIRED IN BOTH CONFIGURATIONS.

TOTAL POWER CONSUMED BY EVERY ACOUSTIC MEASURE ABOVE : ZERO.
```

33.8.4 8.4 The gate in general form

```
THE DECOHERENCE GATE

1. Name the degree of freedom.
2. Name the environment and the temperature.
3. Compute tau_dec = 1 / (n * sigma * v_bar), or state a better model
   and compute it.
4. Divide by hbar/kT at that temperature as a sanity check.
   At 310 K that is 2.464e-14 s.
5. Name the functional timescale required.
6. Compute the ratio and print it.
7. Concede the prefactor: 1-3 orders. Ask whether the conclusion
   survives the concession.
8. If the proposal invokes protection, ask what degree of freedom the
   environment couples to, and whether the proposed degree of freedom is
   one the environment barely watches (spin) or one it watches best
   (charge displacement).
9. State the disposition in writing, and state what result would reverse it.
```

Applied to the two coherence proposals above, the gate returned a rejection in both cases and returned it for a different reason each time: the first failed at step 1, having no degree of freedom to name, and the second passed steps 1 to 6 and failed at step 8.

33.9 9 Discussion

33.9.1 9.1 What follows

The corpus's quantum vocabulary is now audited against the original, and the audit is unfavourable to every use of it except one. The conjugate-structure material is Fourier duality and was already badged as such [Curzi, 2026z]; this paper supplies the reason the badge is required rather than merely prudent. The branch-weight construct is a probability distribution with an effective count, and its effective count is the Hill number of order one, which is grounded [Hill, 1973, Jost, 2006]. Neither needs `hbar` and neither gets it.

The decoherence result is a constraint the corpus imposes on itself and on its readers. The five-field gate of Section 8.4 is not a rhetorical device: it converts an unfalsifiable claim into a computation with a stated disposition, and in both applications above it returned a rejection with a different failure point each time. **A gate that rejected everything for the same reason would be a prejudice; a gate that locates the specific step at which each proposal fails is doing work.**

The pattern across all three sections is the same and it is deliberate: each computes its own most attractive analogy into the ground. That is the corpus's standing practice rather than this paper's innovation, and the sharpest instance elsewhere is a paper that refuses a celebrated historical anecdote as a fabricated data point, rejects its own unified-hand reading at 1,365 to 1, badges a 25 per cent currency-inflation figure as an artefact of a chosen synthetic ratio rather than a constant, and reports its confusable-only check reaching only 58.99 per cent [Curzi, 2026]. **A section that survives only claims it did not test is not a section worth citing, and this paper's counter-result is the load-bearing part of it.**

The Casimir result splits into a refusal and a principle. The refusal is total: the force is `1.996e-23` newtons across a seven-metre room and `0.027` piconewtons on a realistic ten-nanometre biological patch, which is 0.45 per cent of a single kinesin motor's stall force. The principle is exact within the domain where the Helmholtz problem governs, and it is the one thing this paper gives the corpus's architectural material that is stronger than an analogy. **The room-mode equation is not like the Casimir mode count; it is the Casimir mode count, solved for a different field.**

33.9.2 9.2 What does not follow

Nothing here licenses the extension of mode selection to social structure. Postulate 6.1 is badged, it stays badged, and Section 11 records that the corpus's external bibliography contains no support for the extension. An occupancy pattern is not a solution of a wave equation, and the formal identity that makes Section 6.7 a derivation is precisely what the extension lacks.

Nothing here shows that the corpus's biological material is correct. It shows that the material does not depend on quantum mechanics, which is a much weaker statement and the only one the computation supports. The corpus's own biological papers carry their own limits — that acoustic addressing of a cell fails on geometry rather than bandwidth [Curzi, 2026x], and that no mechanism is known by which a transcriptional response could depend on a symbol string as an order [Curzi, 2026] — and those limits are not softened by anything in this paper.

Nothing here settles the measurement problem, and the paper takes no position on it. The half that is open is genuinely open. What the paper establishes is that its openness is bounded by theorems in every direction that would otherwise supply licence.

And the decoherence estimate is an estimate. It is a saturated-collision model with a geometric cross-section, imported from a gas into a liquid, and its prefactor is uncertain by one to three orders of magnitude. H2 says the conclusion survives that concession; it does not say the estimate is precise, and the difference matters wherever a number from Section 5.5 is quoted alone.

33.9.3 9.3 Relation to the surrounding literature

Three literatures bear on this paper and they pull in different directions.

The physics is not in dispute and the paper adds nothing to it. Where the corpus's bibliography holds the relevant entries — Casimir's original result and its two principal measurements, the founding decoherence paper, the uncertainty relation, the quantum algorithms that make coherence an engineering problem — they are cited. Where it does not, the works are named in the text with year and venue and no key is attached; Section 11 and Appendix C record the gap.

The critique of physics-in-social-science is directly relevant and largely supports this paper's method rather than opposing it. Anderson [1972] is the general form of Section 5.10: a law at one scale does not determine the behaviour of the structure built at the next. Cartwright [1999] is the general form of Section 3.3: physical laws hold where their premises hold. Mirowski [1989] and the econophysics critics [Gallegati et al., 2006, McCauley, 2006, Sobkowicz, 2009] document what happens when a formalism is imported without its premises, which is exactly the failure mode Section 7 catalogues. **This paper is on the critics' side of that argument about its own corpus.**

The measurement-theoretic objection applies to the corpus but not to this paper's physics. Michell [1997] is the strongest objection to the framework as a whole: psychology routinely assumes quantitative structure without testing for it, and the test that would settle it is conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971], which has not been run for any construct in this corpus. The quantities in this paper — a coherence time, a pressure, a frequency — are ratio-scale physical quantities with SI units and are not exposed to that objection. **The corpus's own quantities are, and Section 11 says so rather than allowing this paper's dimensional cleanliness to be read as covering them.**

A fourth literature is relevant by its absence. The corpus's comparative material sets a modern result beside an older tradition, and the format of Appendix B exists to keep the two columns apart. **The general hazard is documented — a category into which people are sorted changes the people in it [Hacking, 1995] — and the specific hazard here is simpler: precedent read as evidence.** Appendix B, audit B3, is the case where the evidence column is empty and the precedent column is not, and it is printed that way on purpose.

33.10 10 Verification

The results of this paper compound from six postulates. Each is atomic, each is checkable on its own, and each carries its status. **Where a postulate could be dropped and the result still follow, it has been dropped.**

33.10.1 10.1 Dimensional audit

Every quantity in this paper carries an SI dimension, which is unusual in this corpus and worth stating as such. `[M]`, `[L]`, `[T]`, `[K]` and `[E]` denote mass, length, time, temperature and energy.

Symbol	Dimension	Unit	Status
`n`	`[L^-3]`	m^-3	**derived** from density and molar mass; `3.3428e28` for water
`sigma`	`[L^2]`	m^2	**assumed** geometric; the largest source of prefactor uncertainty
`v_bar`	`[L T^-1]`	m/s	**derived** from `sqrt(8 k_B T/(pi m))`
`Gamma_dec`	`[T^-1]`	s^-1	**derived**: `[L^-3][L^2][L T^-1] = [T^-1]`, so the product balances
`tau_dec`	`[T]`	s	**derived**, `1/Gamma`
`hbar/k_B T`	`[T]`	s	**derived**: `[E T]/[E] = [T]`; `2.464e-14` s at 310 K

Symbol	Dimension	Unit	Status
`rho_01`	`[1]`	none	**defined** dimensionless: an off-diagonal density-matrix element
`a`	`[L]`	m	**stipulated** separation
`P(a)`	`[M L^-1 T^-2]`	Pa	**derived**: `[E T][L T^-1]/[L^4] = [E]/[L^3] =` Pa
`lambda_T`	`[L]`	m	**derived**: `[E T][L T^-1]/[E] = [L]`; 7.63 um at 300 K
`A_H`	`[E]`	J	**cited** Hamaker constant, `5e-21` to `1e-20` J
`f(n_x,n_y,n_z)`	`[T^-1]`	Hz	**derived** from `c_sound` and three lengths
`RT60`	`[T]`	s	**derived** from `0.161` `V/A`, whose coefficient carries `[T L^-1]`

Do the equations balance? `Gamma = n sigma v_bar` gives `m^-3 . m^2 . m s^-1 = s^-1`: it balances, and there is no room in it for a free dimensional parameter. `P = pi^2 hbar c/(240 a^4)` gives `J s . m s^-1 / m^4 = J m^-3 = Pa`: it balances. `hbar/k_B T` gives `J s / J = s`. The room mode equation gives `m s^-1 . m^-1 = s^-1`. **The Sabine coefficient 0.161 is the one dimensionally awkward object in the paper: it carries units of s/m and is an empirical constant for air at ordinary conditions, not a dimensionless number**, which is why the same equation written for imperial units carries 0.049 instead.

Where the corpus has no operational unit, this paper does not supply one. The framework's own `Q` — stored order — is counted in whatever a declared proxy counts, and comfort has no operational unit at all. The correspondence of Section 5.11 relates a dimensionless ratio (`rho_01(t)/rho_01(0)`) to a dimensionless ratio (`Q_t/Q_0` under the transient term), which is the only reason the two can be set beside each other at all. **It is a comparison of ratios, not a conversion of units, and there is no exchange rate between a coherence and a runbook.**

33.10.2 10.2 The postulates

P1 — The number density of liquid water at 310 K is `3.3428e28` molecules per cubic metre. *Status:* **ARITHMETIC.** `1000 kg/m^3 / (0.01802 kg/mol) x 6.022e23 /mol = 3.342e28 /m^3`. A reader who prefers the density at 310 K (993.3 kg/m^3) obtains `3.32e28`, a 0.6 per cent difference which changes nothing.

P2 — The mean thermal speed of a particle of mass `m` at temperature `T` is `sqrt(8 k_B T/(pi m))`. *Status:* **CITED** — elementary kinetic theory. For water at 310 K it returns 603.5 m/s.

P3 — Above the thermal de Broglie wavelength, one scattering event resolves which branch a spatial superposition is in, so the decoherence rate saturates at the scattering rate `n sigma v_bar`. *Status:* **CITED** [Zeh, 1970], with the prefactor and the geometric cross-section badged as the estimate's weakest inputs. **This is the only postulate in the chain that carries a modelling assumption, and H2 exists to test whether the conclusion depends on it.**

P4 — `hbar/k_B T` at 310 K is `2.464e-14` s. *Status:* **ARITHMETIC.** `1.0546e-34 / (1.380649e-23 x 310) = 1.0546e-34 / 4.2800e-21 = 2.4640e-14`.

P5 — The Casimir pressure constant `pi^2 hbar c/240` is `1.300126e-27` J m. *Status:* **ARITHMETIC.** `9.8696 x 1.0546e-34 x 2.9979e8 / 240 = 1.300126e-27`, and the result is measured to

about 1 to 5 per cent [Lamoreaux, 1997, Mohideen and Roy, 1998, Casimir, 1948].

P6 — Between two boundaries separated by `L`, a field obeying the Helmholtz equation with a Dirichlet or Neumann condition admits perpendicular wavenumbers `k_n = n pi/L` and no others. *Status:* **DEFINITION** — the elementary solution of the boundary-value problem, true by construction and carrying no empirical weight on its own; its empirical weight comes from P5's measurement and from the room-acoustic measurements cited in Section 6.7.

33.10.3 10.3 The compounding chain

```
P1 + P2 + P3   -> L1   Gamma = n sigma v_bar; for sigma = 2.500e-15 m2,
                       3.3428e28 x 2.500e-15 x 603.6 = 5.0443e+16 /s        ARITHMETIC

L1             -> R1   tau_dec = 1/Gamma = 1.9824e-17 s ~ 1.98e-17 s        ARITHMETIC

R1 + (1e-3 s)  -> R2   ratio = 1e-3 / 1.982e-17 = 5.04e+13                  ARITHMETIC

R1 + P4        -> R3   tau_dec / (hbar/kT) = 1.982e-17 / 2.464e-14
                       = 8.0e-4, so the estimate is a thousandth of the
                       thermal time and is not an outlier                    ARITHMETIC

R2             -> R4   suppression over one spike = exp(-5.044e13);
                       visibility falls below 1e-3 after 6.9 tau and
                       below 1e-6 after 13.8 tau                             ARITHMETIC

Theorem 5.2    -> R5   tracing out the environment multiplies the
                       off-diagonals by <E_1|E_0> and the diagonals by 1,
                       so no population changes                              DEFINITION

P5             -> R6   P(a) = 1.300126e-27 / a^4; at a = 7.0 m the
                       zero-temperature value is 5.4149e-31 Pa               ARITHMETIC

P5 + zeta(3)   -> R7   the high-temperature limit zeta(3) kB T/(4 pi a^3)
                       at a = 7.0 m, T = 300 K is 1.1551e-24 Pa; over
                       17.28 m2 that is 1.9960e-23 N                         ARITHMETIC

R7             -> R8   one E. coli on the wall weighs 9.807e-15 N,
                       a ratio of 4.913e+08                                  ARITHMETIC

P6             -> R9   f_1 = v/(2L): 2.9979e8/(2 x 1e-7) = 1.4990e15 Hz
                       between plates at 100 nm; 343/(2 x 7.0) = 24.50 Hz
                       between walls at 7.0 m; ratio of answers = ratio of
                       wave speeds = 8.7403e+05                              ARITHMETIC

P6             -> R10  f(nx,ny,nz) = (c/2)sqrt(sum (n_i/L_i)^2): a 5.00 m
                       cube admits 38 modes below 120 Hz on only 11 distinct
                       frequencies, so 27 fall on an already-occupied
                       frequency; the 7.0 x 5.4 x 3.2 m box admits 39 on 39   ARITHMETIC
```

Every result in the chain is ARITHMETIC or DEFINITION, consuming two cited facts (P2, P3) and one cited measurement (P5). That is the strongest verification status any chain in this corpus reaches, and it is available only because the objects are physical quantities with SI units. **It also means the chain establishes nothing whatever about psychology, institutions, or belief.** The one place this paper touches a psychological claim — the correspondence of Section 5.11 — consumes none of these results and is labelled a claim about model form.

Three secondary checks a reader can run in a minute.

Tsirelson's bound. `2 sqrt(2) = 2.828427`, the maximum of the CHSH expression allowed by quantum mechanics against the classical bound of 2. One square root. **ARITHMETIC.**

The Casimir bookkeeping. `pi^2/720 = 0.013707783890` and `pi^2/(6 x 120) = 0.013707783890` are the same number, because `720 = 6 x 120` and `1/120` is the regularised value `zeta(−3)`; and `720/3 = 240`, the 3 being the exponent differentiated. **ARITHMETIC.**

The osmotic comparison. `7.732e5 / 13.001 = 5.947e4`. The ideal-mirror Casimir pressure at 100 nm is one part in sixty thousand of physiological osmotic pressure, and the ideal value is the generous one. **ARITHMETIC.**

33.10.4 10.4 The one out-of-sample check

Section 5.4 is the only place in this paper where an estimator is checked against an outcome it did not use. The saturated-collision formula, applied to a C70-scale object, predicts coherence times of `6.66e−2` s at `1e−6` mbar and `6.66e−5` s at `1e−3` mbar, against a Talbot–Lau flight time of order `1e−3` s. **It therefore predicts that fringe visibility survives at the lower pressure and is destroyed at the higher one, and that is the pressure decade over which the visibility has been measured to fall.** The check is coarse — it resolves a decade, not a percentage — and it is reported as such. **It is nonetheless the paper's only evidence that the estimator is not merely internally consistent**, and Section 11 records that no other estimate in the paper has been checked this way.

33.10.5 10.5 Limits, asymptotics, and the domain of each law

`a -> 0` in the Casimir formula. `P ~ a^−4` diverges. The formula's domain ends well before that: below about 5 to 20 nanometres retardation is negligible and the non-retarded van der Waals form `A_H/(6 pi a^3)` governs instead, with a different exponent and a material-dependent coefficient. **The divergence is outside the model's domain and is not a prediction.**

`a -> infinity`. `P -> 0` as `a^−4` at zero temperature, but as `a^−3` in the high-temperature limit, and the crossover is at `lambda_T = 7.63` um at 300 K. **The asymptotic behaviour is different in the two regimes and the wrong one is the flattering one**, which is why Section 6.4 reports the correction that makes the force larger.

`T -> 0` in the decoherence estimate. `v_bar -> 0` and `tau_dec -> infinity`: coherence survives indefinitely in a perfectly cold, perfectly empty environment, which is the limit every matter-interference experiment is built to approach and which no organism can occupy.

`sigma -> 0`. `tau_dec -> infinity`. A degree of freedom the environment does not couple to is not decohered by it, and this is not a loophole but the actual explanation of the radical-pair exception: electron spin has no electric dipole moment, so a polar solvent barely watches it. **The limit is realised in nature and it is the only way any of the exception cases work.**

`n -> 0`. The vacuum limit, `tau_dec -> infinity`, which is the whole design principle of matter-wave interferometry and is unavailable to biology by definition.

`delta -> 0` in the corresponding storage law. The framework's own boundary case, where the map ceases to contract and no fixed point exists [Curzi, 2026a,?]. The correspondence of Section 5.11 carries this limit across intact: a coherence with `1/tau = 0` never decays, and a stock with `delta = 0` never settles. **In both cases zero loss removes the destination, not the risk.**

33.10.6 10.6 Invariance, and how far the mode-selection claim has been tested

The mode-selection result is invariant under one transformation exactly: **rescale every length by `s` and every frequency scales by `1/s`, with the mode *pattern* — which indices are degenerate, what the ratios between frequencies are — unchanged.** The degeneracy structure of a room

is therefore a property of the *ratios* of its dimensions and not of its size, which is why a cube of any size stacks its modes and a 7 : 5.4 : 3.2 box of any size does not.

It is not invariant under a change of boundary condition, and this is where the claim would break. A Dirichlet boundary (pressure node) and a Neumann boundary (velocity node) give different mode sets, and a real wall is neither: it has finite impedance, so the mode frequencies shift and the modes acquire finite width. **The claim as tested here has been computed for three room geometries and verified against a published formula, not measured in any room.** That is a computational check of an equation, not an empirical test of an invariance, and Section 11 records it as such.

The wave-speed invariance is the strong one and it is tested across nine orders of magnitude: the same formula returns `1.4990e15` Hz for light between plates at 100 nm and 24.50 Hz for sound between walls at 7.0 m, with the ratio of the answers equal to the ratio of the wave speeds, `8.7403e5`, exactly. **Two scales is not many. It is, however, two scales separated by a factor of `8.7e5` in speed and `7e7` in length, and the agreement is exact rather than approximate because the same equation was solved twice.**

33.10.7 10.7 Significant figures against epistemic status

`1.300126e-27` J m is quoted to seven figures and that is defensible, because it is `pi^2 hbar c/240` — an exact expression in defined constants, computable to arbitrary precision. **It is a defined quantity, and it is labelled defined wherever it appears at that precision.**

`1.98e-17` s is quoted to three figures and that is one figure too many for what it is. It is the output of an assumed geometric cross-section in an imported scattering model, and its honest precision is *one order of magnitude*. It is written as `1.98e-17` throughout because the arithmetic must be reproducible from the stated inputs, and it is stated here, once, that **the number should be read as "about `1e-17` seconds" wherever it is used as a physical estimate rather than as a step in a computation.** The same applies to every entry in the table of Section 5.5 and to the ratio `5.04e13`, which inherits the same uncertainty.

`2.464e-14` s is exact given the temperature, being `hbar/k_B T` with defined constants, and is quoted to four figures accordingly.

`24.50 Hz` and the room-mode list carry three to four figures and should not. They are computed from a nominal speed of sound of 343 m/s, which varies by about 0.6 per cent per degree Celsius, and from nominal dimensions. **The mode frequencies are reliable to about one per cent, which is enough to establish the degeneracy structure — the result the section actually uses — and not enough to justify the second decimal place.**

33.11 11 Limitations

No experiment was performed and no dataset was collected. Every number in this paper is either a published constant, an arithmetic consequence of one, or the output of a stated model with stated inputs. The three cases of Section 8 are constructed proposals and a constructed design brief; they demonstrate that the audit procedure runs and reaches a disposition, and they are not evidence about how any real proposal, laboratory or room behaves.

The decoherence numbers are order-of-magnitude estimates and the prefactors are uncertain. The saturated-collision model imports a gas-phase picture into a liquid, uses a geometric cross-section, and ignores the structure of the solvent shell. The uncertainty is one to three orders of magnitude. **H2 is the claim that the conclusion survives that concession, and H2 is an argument rather than a measurement**: it compares the size of the available correction against the size of the required correction and observes that the second is larger by six or more orders. A reader who believes a specific mechanism could buy nine orders has a real objection, and the way to press it is to name the mechanism and compute it.

Only one estimate in the paper has been checked out of sample. Section 5.4's pressure decade is the single instance, it resolves a decade rather than a percentage, and no other row of Section 5.5 has been compared against a measurement. **A model checked at one point and applied at five is being extrapolated, and that is what is happening here.**

No model comparison in the formal sense was performed anywhere in this paper. Measured against what proper model comparison requires, the position is:

Requirement	Status in this paper
Out-of-sample rather than in-sample performance	**Partially met.** One qualitative out-of-sample check, at decade resolution (Section 10.4).
A matched null with the same free-parameter count	**NOT MET.** No competing estimator of coherence time was fitted or compared.
The number of free parameters stated explicitly	**Met.** The estimate has no fitted parameters: `n`, `v_bar` and `sigma` are all computed or stipulated, not estimated from data.
An information criterion, so the parameter count is paid for	**NOT MET.** No AIC or BIC was computed. Nothing in this paper is fitted, so no criterion applies — but the row is printed rather than omitted, because **no AIC or BIC has been computed anywhere in this corpus**, including in the corpus's one applied fit [Curzi, 2026i].
Effect sizes with intervals, not only p-values	**NOT MET.** No interval accompanies any estimate in Section 5.5, because none was propagated. The honest interval on `tau_dec` is a factor of `1e1` to `1e3`, and it is stated in words in Section 10.7 rather than carried through the tables.

Several results are named in the text with year and venue and carry no bibliography key, because the corpus's external bibliography has no entry for them. They are: Born's 1926 introduction of the probability interpretation; Gleason's 1957 theorem; Bell's 1964 theorem and the loophole-free CHSH violations of 2015; the Kochen–Specker theorem; the Pusey–Barrett–Rudolph theorem; Everett 1957; Bohm 1952; Ghirardi, Rimini and Weber 1986; von Neumann 1932; Wigner 1961; Joos and Zeh 1985; Zurek 2003; the Tonomura, Arndt and Fein interference experiments; the Hornberger and Hackermüller decoherence measurements; the Fenna–Matthews–Olson coherence results and their vibronic reinterpretations; the radical-pair magnetoreception literature; Tegmark 2000 and the Hagan–Hameroff–Tuszynski reply; Lifshitz 1956; Boyer 1968; the measured Casimir repulsion; and the dynamical Casimir observations. **Attaching an available key to any of these would be a mis-citation, and a mis-citation is worse than an absent one because it will be checked.** Appendix C lists them with the details the source manuscript supplies, so that a reader can locate them. This is a real gap in the corpus's apparatus and it is recorded as one.

The mode-selection extension beyond acoustics has no external support of any kind. Postulate 6.1 is exact for the Helmholtz problem and is a heuristic beyond it. The corpus's own coverage audit lists the family of claims that carry proportion or geometry from a physical setting into a non-physical one as unsupported, and this is one of them. **A reader who reaches for the acoustics literature as support for the occupancy claim will find that it does not apply.**

The golden-section material carries the standard exposure and the standard defence. The sceptical literature on golden-ratio claims is unambiguous that most of them do not survive inspection [Markowsky, 1992, Fischler, 1981, Livio, 2002, Höge, 1997], and this paper's use of a golden-section room is a **counter-example rather than an endorsement**: the golden box is shown to be worse than the plain box on the measure that matters. That is the only safe use of the ratio available here.

Room modes have been computed, not measured. Section 6.7 and Section 8.3 solve a published equation for three sets of dimensions. No room was measured, no sine sweep was run, and the finite-impedance corrections that make a real wall neither a Dirichlet nor a Neumann boundary are not modelled. **The degeneracy structure is robust to those corrections and the individual frequencies are not**, which is why the paper's claims rest on the former.

The Sabine values in Section 8.3 are used above their reliable range. At a mean absorption of 0.35 the switching rule requires Eyring rather than Sabine, and the divergence between the two can reach 77.8 per cent at 2 kHz in an over-treated room [Curzi, 2026e]. The values are retained because the source computes them that way, and the discrepancy is printed rather than silently corrected.

The correspondence between coherence decay and the storage law's retention term is a claim about model form and nothing more. The corpus's own limits on the storage law apply undiminished: its one preregistered confirmatory test **failed** its second falsification rule, beating a matched AR(1) by 2.8 per cent on held-out error against a declared 10 per cent margin [Curzi, 2026i]. **The exactness of the physical decay says nothing about the applicability of the institutional one.**

The reflexive specification remains unidentified. Where the framework's loop is applied to a group whose members both influence and constitute the aggregate, the endogenous effect is not identified from observational data [Manski, 1993], and network data does not repair it [Shalizi and Thomas, 2011, Lyons, 2011]. This paper adopts no remedy because it applies the loop to nothing; a downstream paper that does apply it inherits the problem, and the realistic route is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001].

The Lucas critique applies to the specifications in Appendix A. The framework holds that a system changes its behaviour when it is modelled and observed [Lucas, 1976, Hacking, 1995]; a register of binding specifications is exactly such an intervention, and its effect on the practice it governs has not been measured. **The specifications are design rules with unmeasured effect, not findings, and they are labelled as such where they appear.**

The corpus's psychological quantities are not covered by this paper's dimensional cleanliness. Every quantity audited in Section 10.1 is a physical quantity with an SI unit. The framework's own constructs — comfort, stored order, non-parallelism — have no operational unit, and Michell [1997] argues that psychology routinely fails to test whether its attributes are quantitative at all. **The conjoint-measurement conditions that would license the ratio-scale treatment those constructs receive have not been tested for any of them** [Luce and Tukey, 1964, Krantz et al., 1971, Rasch, 1960]. A reader must not take Section 10.1's clean audit as covering anything outside this paper.

No scaling exponent in this paper was estimated, so the procedure that estimated exponents require does not apply [Clauset et al., 2009]. This is recorded as a check that does not apply rather than one that was passed, and any downstream use of `a^-4`, `a^-3` or `a^6 T^9` as though it were an empirical finding about a fitted distribution would be a misreading.

Independent replication is absent. Every computation in this paper was performed by one author, and the corpus contains no independent re-derivation of any of them. The reproducibility literature is unambiguous about what a single unreplicated result is worth [Open Science Collaboration, 2015, Ioannidis, 2005], and the arithmetic here is at least the kind that a reader can redo in a few minutes, which is the entire design of Section 10.

Two observations about the source material, printed rather than corrected. First, **the "27 coincident modes" figure for the 5 m cube counts modes falling on a frequency already occupied by another mode**: the cube admits 38 modes below 120 Hz on only 11 distinct frequencies, so 27 are coincident in that sense while 36 belong to a degenerate group. Both readings are defensible and the source does not say which it means; the count is stated explicitly here so that it cannot be quoted ambiguously. Second, **the water number density `3.3428e28` per cubic metre corresponds to a density of 1000 kg/m^3, which is the value at 4 C rather than at 310 K.** At 310 K the density is about 993.3 kg/m^3 and the number density is `3.32e28`, a 0.6 per cent difference. The source's value is retained because every downstream figure is computed from it, and the discrepancy is recorded because a reader recomputing from the stated temperature

will not reproduce the fifth significant figure.

What would falsify the account. H1 fails on a replicated demonstration of functional quantum coherence above `1e-6` s in an unrefrigerated, unengineered biological system at 310 K. H2 fails on a corrected treatment of the same system returning a coherence time more than four orders above the collision estimate. H3 is a theorem and fails only on an error. H4 fails on a measured room-mode frequency departing from the modal equation by more than boundary compliance accounts for. **If H1 were falsified, Section 5 would be rewritten and Section 8.2's disposition reversed. That is stated in advance, and it is the point of stating a threshold three orders more generous than the claim requires.**

33.12 12 Conclusion

The quantum formalism is stated here as established physics and this paper adds nothing to it. Unitarity and linearity are verified numerically to `1.05e-15` and `1.29e-13`, at the double-precision floor. The Born rule integrates to unity, and interference is shown to redistribute probability rather than create it: visibility `1.0000` coherent against `0.4883` incoherent, with identical total probability, and a phase-average of `1.000000` that is exactly the classical sum. Four interpretations are surveyed by which of three inconsistent premises each abandons. Consciousness-causes-collapse is placed accurately — a minority historical position, not refuted, not favoured, with an **empty evidence column** and a distinguished advocate who abandoned it. **The measurement problem is left open, because it is open; and the openness is bounded in every direction by theorems, of which no-signalling is the one that matters, because it holds in every interpretation.** The quantum bound on the CHSH expression is `2 sqrt(2) = 2.828427`, attained exactly at the optimal settings.

The decoherence result is the paper's principal constraint and it is stated to be quoted. With `n = 3.3428e28` per cubic metre and a mean thermal speed of `603.6` m/s in water at 310 K, the estimated coherence time is `8.05e-13` s for a single potassium ion, `6.31e-16` s for a globular protein, `1.98e-17` s for a microtubule segment and `6.31e-22` s for a neuron soma. Against a `1e-3` s action potential the microtubule ratio is `5.04e13`; **a hypothetical protection factor of `1e6`, far beyond anything demonstrated for a charge-displacing degree of freedom in water, still leaves `5.04e7` unexplained.** The thermal time `hbar/k_B T` at 310 K is `2.464e-14` s, and it is imposed here as the mandatory comparison for any biological coherence claim; the microtubule estimate is `8.0e-4` of it. **Decoherence suppresses the coherences as `rho_01(t) = rho_01(0) exp(-t/tau)` and changes no diagonal element: it converts a superposition into something that looks classical without solving the measurement problem, and a result derived from linear unitary evolution could never have resolved a contradiction that arises from it.** The genuine exceptions — photosynthetic coherence at about eleven thermal times inside one rigid protein, with a serious vibronic reinterpretation, and radical-pair magnetoreception at two ten-millionths of the thermal energy — read out **populations**, not messages, and neither generalises to cognition.

The Casimir force is real, measured to about one per cent, and negligible wherever this framework builds. The pressure is `P(a) = pi^2 hbar c/(240 a^4) = 1.300126e-27/a^4` pascals — **a defined constant, exact, and quoted at that precision on that ground alone** — giving `1.3001e5` Pa at 10 nm and `1.155e-24` Pa on a seven-metre room wall in the correct high-temperature limit, for a total force of `1.996e-23` N: **five hundred million times smaller than the weight of a single bacterium resting on that wall.** In a cell the realistic Lifshitz value on a 10 nm patch is `0.027` pN against `6.00` pN for one kinesin motor. **Only the mode-selection principle transfers, and that transfer is badged: exact for electromagnetic modes between conductors and for acoustic modes in an enclosure, because both are the same Helmholtz problem, and a design heuristic with no formal warrant anywhere else.** A 5.00 m cube admits 38 modes below 120 Hz on 11 distinct frequencies, so 27 fall where another already sits; a 7.0 x 5.4 x 3.2 m box admits 39 on 39, and the difference cost nothing but the choice of three numbers.

And the result this paper exists to establish, stated so it can be quoted without drift: the corpus makes no quantum claim about mind, biology or belief, and nothing in it would change

if quantum mechanics played no role in biology at all. That is not a concession extracted by a reviewer. It is a computation the corpus ran on itself, and printed.

33.12.1 Declarations

Data and code availability

No external dataset was collected for this paper. Every numerical result is computed from published physical constants and from the stated model parameters, and every table can be re-derived from the equations and inputs printed beside it. The numerical verifications of unitarity, linearity and the Born rule reported in Sections 4.3 and 4.4 are outputs of standard split-step and quadrature computations on the grids stated in the text. The three cases worked in Section 8 are constructed proposals and a constructed design brief, not field observations. The figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper and the formulae of Sections 5 and 6.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, performed the computations, produced the worked audits and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The proposals audited in Sections 8.1 and 8.2 are constructed composites of positions available in the published literature and in circulation; they are not reports of correspondence with any identifiable person, and no individual is described.

Preregistration

This paper reports no preregistered test of its own; its results are computations from published constants. **H1 nonetheless states its falsification threshold in advance** — functional quantum coherence above 1e−6 s in an unrefrigerated, unengineered biological system at 310 K — and that threshold is three orders of magnitude more generous than the claim under audit requires. The corpus's preregistered protocol for testing the storage law, together with the confirmatory application that **failed** its second falsification rule, is reported in Curzi [2026i].

33.13 Appendix A Specification register

These specifications bind the corpus's use of quantum and Casimir vocabulary. **They are design rules whose effect on practice has not been measured, not findings**, and Section 11 records that distinction. They are reproduced here because downstream papers cite them by number.

33.13.1 A.1 On the formalism

S-4.1 STATE THE BORN RULE CORRECTLY OR NOT AT ALL. P(a_n) = |<a_n|psi>|^2 ; p(x)dx = |psi(x)|^2 dx. Never "psi is the probability". Never omit the modulus-squared.

S-4.2 USE "OBSERVER" ONLY IN ITS TECHNICAL SENSE and define it at first use in any document: a physical system with many degrees of freedom that interacts with and records a correlation. Where the intended meaning is a person paying attention, use those words instead.

S-4.3 NEVER ASSERT A SINGLE INTERPRETATION AS SETTLED. Where an argument depends on interpretation, name the interpretation and name what it gives up (the M1/M2/M3 table of Section 4.6).

S-4.4 PRESENT VON NEUMANN-WIGNER, WHERE PRESENTED AT ALL, WITH ALL FIVE REASONS OF SECTION 4.8, INCLUDING WIGNER'S OWN ABANDONMENT OF IT. Partial presentation is prohibited.

S-4.5 NO CLAIM IN THIS CORPUS MAY DEPEND ON A QUANTUM MECHANISM. Where a proposed channel is quantum, the specific coupling must be written down, its magnitude estimated, and the decoherence timescale of Section 5 confronted.

S-4.6 CITE NO-SIGNALLING WHEREVER ACTION AT A DISTANCE IS IMPLIED. It is the shortest complete refutation available and it is interpretation-independent.

S-4.7 DISTINGUISH FORMALISM FROM INTERPRETATION IN EVERY PARAGRAPH THAT TOUCHES BOTH.

S-4.8 THE ANALOGY BADGE ON THE CONJUGATE-STRUCTURE MATERIAL CARRIES FORWARD UNCHANGED. Conjugate structure in a future-ensemble is a Fourier fact and carries no quantum content.

S-4.9 WHERE AN EVIDENCE COLUMN IS EMPTY, PRINT THE EMPTY CELL. Appendix B, audit B3, is the template.

S-4.10 NEVER SAY THAT BELL'S THEOREM DISPROVED REALISM. It excludes local theories with pre-existing definite values for all observables. Say that, in those words.

S-4.11 ANY BRANCH SET PRESENTED TO AN ANALYST MUST CARRY ITS WEIGHTS AND ITS EFFECTIVE COUNT, and must state which effective count is meant. A branch set without weights is a list, and a list is not a forecast.

S-4.12 ANY FORECAST MADE ON A BRANCH SET IS SUBJECT TO THE PREREGISTRATION PROTOCOL: preregister the branch set and the weights, freeze the model, score held-out against persistence, direct-input, matched autoregressive and shuffled-input baselines, and state in advance what result kills it.

33.13.2 A.2 On coherence claims

S-5.1 REPRODUCE THE WARM-WET TIMESCALE STATEMENT IN FULL WHEREVER THE CORPUS

TOUCHES QUANTUM VOCABULARY IN A BIOLOGICAL CONTEXT.

S-5.2 ANY PROPOSED QUANTUM MECHANISM IN A BIOLOGICAL SYSTEM MUST STATE, IN THE PROPOSAL, ALL FIVE OF:
(i) the degree of freedom carrying the coherence
(ii) the environment it couples to
(iii) an estimated decoherence time with its formula and inputs
(iv) the functional timescale it must reach
(v) the ratio (iv)/(iii)
A proposal missing any of the five is returned unreviewed.

S-5.3 COMPARE EVERY BIOLOGICAL COHERENCE CLAIM TO hbar/kT AT 310 K = 2.464e-14 s. State the claim as a multiple of the thermal time. It is the fastest available sanity check and it costs one division.

S-5.4 NEVER STATE THAT DECOHERENCE SOLVES THE MEASUREMENT PROBLEM. State that it explains the absence of macroscopic interference and the preferred basis, and that it does not select an outcome.

S-5.5 PRESENT PHOTOSYNTHETIC COHERENCE AND MAGNETORECEPTION AS ACTIVE AND CONTESTED, with the vibronic reinterpretation named in the same passage, and never as support for any claim about cognition or belief.

S-5.6 DISTINGUISH POPULATIONS FROM COHERENCES IN EVERY DENSITY-MATRIX ARGUMENT. A yield difference is a population. A message is not.

S-5.7 LABEL ORDER-OF-MAGNITUDE ESTIMATES AS ESTIMATES, give the formula and the inputs, and state the prefactor uncertainty.

S-5.8 WHERE AN ESTIMATE CAN BE CHECKED AGAINST A PUBLISHED EXPERIMENT, CHECK IT AND PRINT THE CHECK. Section 5.4 is the template.

S-5.9 ANY FUTURE RESULT OVERTURNING SECTION 5 MUST BE INCORPORATED, NOT SUPPRESSED. The falsification condition is H1 and it is stated in advance.

S-5.10 THE STRUCTURAL LESSON MAY BE CARRIED FORWARD; THE PHYSICS MAY NOT. "Protect a fragile state by reducing coupling, not by adding energy" is transferable engineering. "Coherence" as a description of an institution is a metaphor and must be badged as one.

33.13.3 A.3 On the Casimir effect and boundary geometry

S-6.1 STATE THE HONEST LIMIT WHEREVER THE CASIMIR EFFECT APPEARS: at every biological and architectural scale the force is utterly negligible.

S-6.2 NEVER CITE THE CASIMIR EFFECT AS A FORCE IN A DESIGN. Cite it as the demonstration of mode selection by boundary. Where a design's mechanism is the Casimir force, the design is rejected without further review.

S-6.3 NEVER CONFLATE CASIMIR WITH VAN DER WAALS IN A BIOLOGICAL ARGUMENT. State which regime is meant -- non-retarded below about 5-20 nm, retarded above -- and give the Hamaker constant where the non-retarded regime is intended.

```
S-6.4   REPORT THE THERMAL CORRECTION WHEREVER THE TABLE EXTENDS BEYOND
        lambda_T = hbar c / kT (7.63 um at 300 K), and report it even though
        it makes the force LARGER at those separations.

S-6.5   NO ZERO-POINT-ENERGY EXTRACTION CLAIM MAY APPEAR IN ANY MATERIAL
        DERIVED FROM THIS CORPUS. The force is conservative; a cycle returns
        what it took.

S-6.6   EVERY ENCLOSURE DESIGNED UNDER THIS FRAMEWORK MUST HAVE ITS MODE LIST
        COMPUTED FROM ITS ACTUAL DIMENSIONS. Report: count below 100 Hz, mean
        spacing, largest gap, and the number of coincident modes, stating
        which definition of coincident is used.

S-6.7   NO PROPORTION SYSTEM CERTIFIES A ROOM. Ratio systems GENERATE
        candidates; certification requires the computed mode list or a
        measurement of that specific room.

S-6.8   DISTINGUISH THE THREE TRANSFER STRENGTHS IN PRINT, EVERY TIME:
          (a) EM modes between plates          -- proven and measured
          (b) acoustic modes in a room         -- the SAME equation, measured
          (c) occupancy / circulation patterns -- SPECULATIVE, badged
        Never allow (c) to inherit the credibility of (a).

S-6.9   WHERE A PRECEDENT COLUMN IS EMPTY, PRINT THE EMPTY CELL.

S-6.10  THE MONOCHORD IS THE CANONICAL PRECEDENT FOR THIS PRINCIPLE AND MUST
        ALWAYS BE LABELLED PRECEDENT. It records that boundary position alone
        sets the mode ladder. It is not evidence for the Casimir effect.

S-6.11  MODE-LADDER PITCH CLASSES FOLD BY ODD PART, giving ceil(N/2) classes in
        the first N modes -- ONE octave between n and 2n. The essence law
        f = k n^2 also folds by odd part but gives TWO octaves between n and
        2n. State which system is in use.

S-6.12  A MOVING BOUNDARY IS A DESIGN ELEMENT WITH ITS OWN SPECIFICATION.
        Where a room has variable geometry, compute the mode list in EVERY
        configuration the room can be placed in.
```

33.14 Appendix B Evidence-and-precedent audits

The corpus's comparative material sets a modern result beside an older tradition that reached a structurally similar position by an independent route. **The format below exists to stop the second column being read as support for the first.** Each audit states what is measured, what is merely recorded, and what the pair can and cannot establish. Three criteria are applied: C1, whether the correspondence is structural rather than verbal; C2, whether the pairing generates a quantitative prediction; C3, whether it is usable on both sides.

B1 — Linear evolution cannot select an outcome. *Evidence:* unitarity of `U(t)`; the numerical verification of Section 4.3, norm to `1.05e-15` and linearity to `1.29e-13`. What is measured: the equation's solutions, to machine precision. *Precedent:* the Madhyamaka doctrine that phenomena have no intrinsic self-standing nature, only relational determination — a sustained argument, arrived at with no instruments, that what a thing is is not a property it carries alone. *Assessment:* the evidence column carries a proved property of a differential equation; the precedent column carries only the fact that a non-intrinsic account of properties

was reachable by argument centuries before Hilbert spaces existed. **Neither is evidence for the other.** C1 pass, C2 fail, C3 partial. *Status:* **speculative as a bridge**, and marked speculative here because it fails C2. Retained as precedent, nothing more.

B2 — Objective spontaneous indeterminacy as a physical process. *Evidence:* the GRW and CSL programme and its experimental constraint programme, including the underground measurement excluding the simplest parameter-free gravity-related collapse model. What is measured: bounds on collapse-rate parameters from interferometry, mechanical oscillators, and searches for the spontaneous radiation such models predict. *Precedent:* the Lucretian *clinamen* — atoms swerving spontaneously, at no fixed time and no fixed place; and the Ash'arite occasionalist position that the world is re-created at each instant rather than carried forward by inertia. What is recorded: that objective, uncaused, moment-by-moment indeterminacy was formulated as a considered position twice, in unrelated traditions, two millennia and a continent apart. *Assessment:* this is the strongest precedent entry in the paper and it must still be labelled correctly. **Lucretius supplies no rate constant; the occasionalists supply no differential equation.** C1 pass; C2 fails on the precedent side; C3 fails on the precedent side. *Status:* the physics is **established as a tested programme**; the bridge is **speculative** and is marked speculative here.

B3 — Mind as the terminus of the measurement chain. *Evidence:* **none.** There is no measurement supporting a distinguished physical role for consciousness in quantum dynamics, no proposed coupling, and no constrained parameter. **The evidence column for this claim is empty, and the empty cell is printed rather than hidden.** *Precedent:* cognition-only analyses in the Yogacara tradition, and Berkeley's immaterialism. What is recorded: that mind-first accounts of the object were formulated with rigour and defended for centuries. *Assessment:* this is the clearest demonstration in the paper of why the two columns must be kept apart. There is a long, serious, independently arrived-at precedent for mind-first metaphysics, and there is **no evidence** for mind-caused collapse. **An audit that reported both columns as support would be lying by layout.** C1 partial (surface resemblance only), C2 fail, C3 fail. *Status:* **rejected as a bridge**, and recorded in print as a failed bridge, because a method that never rejects anything is not a method.

B4 — The future is better held as a weighted branch set than as a verdict. *Evidence:* calibration research in forecasting — proper scoring rules, and the finding that probabilistic forecasts outperform point predictions when scored honestly on held-out outcomes; and, within this corpus, the preregistration protocol requiring a frozen model beaten against four baselines including a shuffled-input control [Curzi, 2026i, Nosek et al., 2018]. What is measured: forecast skill, out of sample. *Precedent:* the yarrow-stalk procedure of the *I Ching*, whose traditional method assigns **non-uniform** probabilities to the four line types — old yin 1/16, young yang 5/16, young yin 7/16, old yang 3/16 — with only the "old" lines changing and generating a second hexagram. Computed: the four weights sum to 16/16; the probability of at least one changing line in six is 0.8220; the later three-coin method gives 1/8, 3/8, 3/8, 1/8, the same total changing-line probability of 1/4 but a different distribution between yin and yang; and the derived state space is `64 x 64 = 4096` ordered pairs. *Assessment:* the precedent records that a weighted-transition model of the future was constructed and preserved for millennia by people with no probability theory. **It is not evidence that its weights are correct, and no claim is made here that consulting it forecasts anything.** C1 pass, C2 pass on the evidence side — the pairing predicts that analysts trained to hold non-uniform branch weights produce better-calibrated forecasts than analysts trained to name a single outcome, which is measurable under the corpus's protocol — C3 pass. *Status:* **established as a structural correspondence**; **not** established and not claimed as any statement about quantum mechanics, neither construct involving interference, complex amplitude, or `hbar`.

B5 — A delicate state requires isolation from its surround, or it is lost. *Evidence:* the decoherence programme [Zeh, 1970] and the measured suppression of interference as an explicit function of environmental coupling — cavity-field superpositions, engineered reservoirs in trapped ions, collisional decoherence of C70 against gas pressure, and thermal-photon decoherence against internal temperature. Computed here: dust grain in air `2.78e-17` s; C70 at `1e-6` mbar `6.66e-2` s; protein in water at 310 K `6.31e-16` s. What is measured: interference visibility

as a dialled function of coupling. *Precedent:* the alchemical requirement of the sealed vessel, documented from late antiquity through the Latin tradition — the work is destroyed if the vessel is opened before the operation completes — together with parallel practices of enclosure and of silence about a working before it is finished. What is recorded: that a delicate process requires protection from its surround was arrived at independently in several traditions and encoded as an operational requirement rather than a sentiment. *Assessment:* the evidence column carries a measured, quantified, dial-controlled relationship with a master equation; the precedent column carries only that the structure was findable by long first-person experience with slow chemistry and with attention. **The alchemists were not doing quantum mechanics and nothing here suggests they were.** C1 pass — in both cases an uncontrolled coupling to a many-degree-of-freedom surround destroys a prepared state, and the remedy in both is to reduce the coupling rather than to add energy. C2 pass — the pairing predicts that a protected process shows a loss rate scaling with its exposure surface, which is the framework's `delta` and is measurable [Curzi, 2026a]. C3 pass. *Status:* the physics is **established**; the bridge is **established as a structural correspondence** and carries no implication that alchemical practice involved quantum coherence.

B6 — The Casimir force is negligible at every scale this framework builds at. *Evidence:* Casimir [1948] for the prediction; Lamoreaux [1997] at about 5 per cent over 0.6–6 um; Mohideen and Roy [1998] at about 1 per cent over 0.1–0.9 um; parallel-plate measurements; Lifshitz theory for real materials. Computed here: 13 Pa at 100 nm ideal against `7.732e5` Pa physiological osmotic pressure; `1.4674e-23` Pa across a 3 m room against a bacterium's weight, ratio `6.683e7`. *Precedent:* **none, and that is worth printing.** No traditional corpus records anything corresponding to a vacuum-fluctuation pressure between plates. **The precedent column is empty.** *Assessment:* an empty precedent column is not a weakness of the claim; it is a reminder of what the two columns do. The claim stands on the evidence column alone. *Status:* **established**, with the honest limit mandatory wherever the effect is invoked.

B7 — Boundary geometry alone determines which modes can exist, with measurable consequence. *Evidence:* the Casimir measurements above — boundary separation alone, in vacuum, with no applied field, producing a force measured to about one per cent; and room acoustics, where the modal equation and the Sabine relation are routinely measured [Sabine, 1922, Schroeder, 1965, Rossing et al., 2002, von Helmholtz, 1885]. Computed here: 25 modes below 100 Hz in a 7.0 x 5.4 x 3.2 m room; a 5 m cube admitting 38 modes below 120 Hz on 11 distinct frequencies against the box's 39 on 39. *Precedent:* the Pythagorean monochord — a single string and a movable bridge, where stopping the string at 1/2 its length raises the pitch by exactly a factor of 2, at 2/3 by 3/2, at 3/4 by 4/3, at 4/5 by 5/4, with the ratios exact and independent of the string's material, tension and thickness. The tradition recorded that **the boundary position alone sets the mode ladder** and built a theory of consonance on it. Also: the tuned bronze vessels specified for Roman theatres; the definition of place as the innermost boundary of the containing body; and the observation that a vessel's usefulness lies in the emptiness its walls enclose. *Assessment:* the evidence column carries measured forces, measured mode frequencies and measured resonances. The precedent column carries that the principle was arrived at, stated generally, and used as engineering more than two thousand years before there was a wave equation to derive it from. **The monochord is not evidence for the Casimir effect.** What its independent arrival establishes is that the structure was findable by ear as well as by torsion balance. **One further note required by honesty: the efficacy of the Roman vessels in actual theatres is not established, and the architectural proportion traditions are precedent for the conviction that shape determines effect, not evidence for their specific prescriptions.** C1 pass, and it is the strongest C1 in the paper, because the correspondence between the two evidence items is not resemblance but literal mathematical identity — the same Helmholtz equation, the same boundary condition, the same `k_n = n pi/a`, for a different field. C2 pass: the pairing predicts specific mode frequencies for specific dimensions, measurable with a sine sweep and a microphone, and the cube-degeneracy prediction is a falsifiable statement about a real room. C3 pass. *Status:* **established** for electromagnetic and acoustic modes; the extension to occupancy and circulation is **speculative**, is marked speculative here, and carries the obligation that a mode argument may generate a candidate geometry and may

not certify one.

33.15 Appendix C Works named in the text without a bibliography key

The corpus's external bibliography does not contain entries for the following, and no key has been attached to them anywhere in this paper. They are listed with the details the source manuscript supplies so that a reader can locate them, and their absence is recorded as a gap in the corpus's apparatus rather than concealed. **Nothing in this list should be read as a citation; it is a statement of what is missing.**

On the formalism. Born's 1926 introduction of the probability interpretation (*Zeitschrift für Physik* 37, 863). Gleason (1957) on the forcing of the Born rule in dimension three and above. Bell (1964) and the loophole-free CHSH violations reported in 2015. The Kochen–Specker theorem (1967). The Pusey–Barrett–Rudolph theorem (2012, *Nature Physics*). Everett (1957, *Reviews of Modern Physics* 29, 454). Bohm (1952, *Physical Review* 85, 166 and 180). Ghirardi, Rimini and Weber (1986, *Physical Review D* 34, 470); Pearle (1989); Ghirardi, Pearle and Rimini (1990). The underground X-ray constraint on parameter-free gravity-related collapse (Donadi and colleagues, 2021, *Nature Physics*). Von Neumann, *Mathematische Grundlagen der Quantenmechanik* (1932). London and Bauer (1939). Wigner, "Remarks on the Mind–Body Question" (1961). Frauchiger and Renner (2018, *Nature Communications* 9, 3711) and Bong and colleagues (2020, *Nature Physics*). Schlosshauer, Kofler and Zeilinger (2013, *Studies in History and Philosophy of Modern Physics* 44, 222). Maudlin (1995, *Topoi*) for the three-premise statement of the measurement problem. Gisin (1990) on signalling under non-linear modifications. Tonomura and colleagues (1989, *American Journal of Physics* 57, 117); Arndt and colleagues (1999, *Nature* 401, 680); Fein and colleagues (2019, *Nature Physics*).

On decoherence. Joos and Zeh (1985, *Zeitschrift für Physik B* 59, 223). Zurek (2003, *Reviews of Modern Physics* 75, 715). Brune and colleagues (1996); Myatt and colleagues (2000); Hornberger and colleagues (2003); Hackermüller and colleagues (2004). Adler (2003) on decoherence and the measurement problem. Tegmark (2000, *Physical Review E* 61, 4194) and Hagan, Hameroff and Tuszynski (2002, *Physical Review E* 65, 061901).

On quantum effects in biology. Engel and colleagues (2007, *Nature* 446, 782); Panitchayangkoon and colleagues (2010, *PNAS* 107, 12766); Duan and colleagues (2017, *PNAS* 114, 8493); Wilkins and Dattani (2015); Thyrhaug and colleagues (2018, *Nature Chemistry*); Cao and colleagues (2020, *Science Advances* 6, eaaz4888). Schulten, Swenberg and Weller (1978); Ritz, Adem and Schulten (2000, *Biophysical Journal* 78, 707); Maeda and colleagues (2008, *Nature* 453, 387); Xu and colleagues (2021, *Nature* 594, 535); Hore and Mouritsen (2016, *Annual Review of Biophysics* 45, 299). Tajik and colleagues (2016, *Nature Materials*) on force transmission to the nucleus. Levine, Gordon and Fields (1978) on naloxone-reversible expectancy analgesia.

On the Casimir effect. Lifshitz (1956) for real materials. Bressi and colleagues (2002) for parallel plates. Boyer (1968) for the spherical shell. Munday, Capasso and Parsegian (2009, *Nature* 457, 170) for measured repulsion. Chan and colleagues (2001) for actuation. Cavalcanti (2004) and Hertzberg and colleagues (2005) for the Casimir piston. Rodriguez, Capasso and Johnson (2011, *Nature Photonics*) for microstructured geometries. Moore (1970); Fulling and Davies (1976); Wilson and colleagues (2011, *Nature* 479, 376); Lähteenmäki and colleagues (2013, *PNAS*) for the dynamical effect.

Chapter 34

Belief Creates Reality If the Belief Is Grounded in Truth: A Conditional Postulate, Its Channel-Gain Form, and the Placebo Case

Abstract. We state a conditional postulate — belief changes the world only through a real, nameable channel — and give it a working form E = G.B + eps, in which B is a normalised strength of holding on [0,1], G is a channel gain in the outcome's own units, and the load-bearing half is the forbidding clause G = 0 implies E = 0 for every B. The postulate is badged as a postulate and defended by what it forbids rather than by what it promises. We compute the honest cost of conditionalising: the unconditional form is refuted by 32.5 per cent of random observations, so eight casual observations refute it with 95 per cent confidence, while the conditional form is refuted by only 5.5 per cent — the paper names this as the flank it volunteers. Channels are enumerated in four classes and composed by an explicit rule, series multiplying and parallel adding; a five-stage series chain delivers 8.4 per cent of a lossless one. The worked audit sums a behavioural gain of 0.18, an informational gain of 0.067 and a social gain of 0.11 to 0.36, bounding the effect at 0.252 at B = 0.7 against a claimed 0.85 — required over available, 3.373 — and rejects the claim. Placebo supplies the one fully verified channel, at modest measured size.

Keywords: conditional postulate; channel gain; placebo analgesia; open-label placebo; loop gain; effect size

Series: B. Engineering the Surplus **Paper:** CCP-034

34.1 1 Introduction

The most quotable sentence in this corpus is a postulate, and this paper exists to keep it one.

> **[POSTULATE] Belief creates reality if the belief is grounded in truth.**

That sentence is badged as a postulate in the source manuscript, it is badged as a postulate here, and it is not a derived result, an empirical finding, or a summary of a literature. It is asserted by this work, it has no external support as stated, and its status is recorded in the abstract, at first use in Section 4, in the verification chain of Section 12 as `POSTULATED`, and in the Conclusion. A reader who quotes it without the badge has quoted it wrongly.

The paper's argument is that the last five words carry the entire content and everything else in the sentence is scaffolding. Strike the conditional and what remains — *belief creates reality* — is shown in Section 5 to be simultaneously **false when taken literally** and **empty when taken charitably**, a double failure with an exact arithmetic. Restore the conditional and the claim names a precondition, says different things about cases where the precondition holds and cases where it does not, and thereby becomes the kind of claim that can be wrong.

The working form is

```
E  =  G . B  +  eps
```

with `B` a normalised strength of holding on `[0,1]`, `G` a channel gain in the outcome's own units, `eps` mean-zero noise, and the load-bearing half **the forbidding clause: `G = 0` implies `E = 0` for every value of `B` whatsoever.** The postulate is defended by what it forbids rather than by what it promises.

The correspondence with the corpus's storage law is exact in form and is the reason this paper sits in Series B rather than in a philosophy of mind. In the storage law `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t`, the conversion coefficient `eta` is what turns offered surplus into durable structure, and `eta = 0` leaves a vessel that leaks and pays rent [Curzi, 2026a]. The two statements

```
eta = 0   =>   Q_{t+1} = (1 − delta) Q_t − C_t      nothing sticks, and cost still bills
G   = 0   =>   E = 0  for all B                     the belief does nothing, at any strength
```

are the same sentence at two scopes. Other papers in the corpus cite that pairing, so it is stated precisely: the correspondence is between `B` and `S_t`, between `G` and `eta`, and between `E` and what lands in `Q`; it is a correspondence of **model form**, badged as one in Section 7.2, and it transfers the structure without transferring the warrant. The storage law is a theorem about a recursion; the postulate is an assertion about the world.

Four things are established here. First, the arithmetic of conditionalisation: under the base rates of Section 5.4, **32.5 per cent of random observations refute the unconditional form — so eight casual observations refute it with 95 per cent confidence — against 5.5 per cent for the conditional.** That number is the flank the conditional volunteers, and it is printed against the paper's own claim rather than buried. Second, a dynamical form: `B_{t+1} = (1−lam)B_t + lam min(1, L B_t)` with loop gain `L = alpha G`, whose zero-belief state is stable if and only if `L < 1`, whose transition at `L = 1` is sharp, and in which `G = 0` forces `L = 0` at every `alpha`. Third, a composition rule — series multiplies, parallel adds — under which a five-stage chain of individually respectable transmissions delivers **8.4 per cent** of a lossless one and any single zero stage zeroes the product. Fourth, a worked audit in which summed gains of 0.18, 0.11 and 0.067 bound the effect at 0.252 against a claimed 0.85, **a required-over-available ratio of 3.373, and the claim is rejected** while the practice that generated it is kept.

The corpus's reflexive loop `T(x) = W(x, M(x))`, of which the self-fulfilling prophecy is a fixed point, is presupposed here rather than re-derived [Curzi, 2026]. **It is also the source of this paper's most serious methodological exposure**, treated in Section 3.4: the loop is a linear-in-means reflexive specification whose endogenous effect is not identified from observational data [Manski, 1993].

Section 2 states the hypotheses. Section 3 gives the method and the identification problem. Section 4 states the postulate with its badge. Section 5 computes the content of the conditional. Section 6 runs the seven-gate formula validation protocol on the paper's own working form and reports the gate it fails. Section 7 makes grounding operational through four channel classes and their composition. Section 8 gives the one channel that has been traced end to end and cut. Section 9 collects the failure modes. Section 10 works three audits. Sections 11 to 14 discuss, verify, limit and conclude.

34.2 2 Hypotheses

The postulate itself is not a hypothesis and is not presented as one. It is an assertion of this work, badged, with a list of observations that would end it (Section 5.5). What follows are the testable consequences that can be separated from it.

> **H1 — The zero-channel prediction.** Where every candidate channel from a belief to a measured outcome has been enumerated and independently excluded, the measured effect of that belief on that outcome is zero within noise. **Predicted:** an effect indistinguishable from zero at any belief magnitude, in a design with power adequate to detect the smallest effect the named channels could have supplied. **Falsified by:** a belief-linked effect clearly above the noise floor, on an outcome for which no physiological, behavioural, social or informational route survives enumeration, with no measurement artefact and no selection effect, preregistered and replicated by an independent group. **This is the falsifier that matters; the rest are calibration.**
>
> **H2 — The magnitude bound.** A belief's effect on an outcome does not exceed the sum of the gains of its named channels, times the belief magnitude. **Predicted:** `|E| <= (sum_i G_i) B_max`, with the gains estimated from each channel's own mechanism and frozen before the outcome is measured. **Falsified by:** effects reliably and substantially exceeding that bound with gains frozen in advance. One instance is a missed channel; a pattern is a dead postulate.
>
> **H3 — Gain proportionality.** Effect magnitude tracks channel gain across channels whose gains are well separated and independently measured. **Predicted:** halving the gain halves the effect, at fixed belief magnitude. **Falsified by:** effect magnitude independent of channel gain across such channels. If halving the gain does not halve the effect, `G` is not what is doing the work.
>
> **H4 — Blockade.** Interrupting the only named channel abolishes the effect. **Predicted:** with the channel's stage `i` set to zero transmission, the series product is zero and the measured effect goes to zero. **Falsified by:** effects surviving the specific, verified interruption of the sole named channel. **A channel that can be cut without changing the outcome was not the channel.**
>
> **H5 — Grounding relevance.** Channels whose constituent propositions are true outperform channels whose constituent propositions are false, at equal belief magnitude. **Predicted:** a difference in effect attributable to the truth of the channel's premises rather than to the truth of the belief. **Falsified by:** beliefs whose channels rest on demonstrably false propositions producing the same effects at the same magnitudes as beliefs whose channels rest on true ones. **This would not refute that belief has effects. It would refute the word "truth" in the postulate, which is the word the postulate is named for, and the word should then be struck.**
>
> **H6 — Bimodality under self-reinforcement.** Beliefs sitting in high-gain reflexive loops are distributed bimodally in a population rather than smoothly, because the fixed-point structure of Section 5.6 admits no stable intermediate value under pure self-reinforcement. **Predicted:** at `lam = 0.30`, `B* = 0` for every `L < 1` and `B*` at the effector ceiling for every `L > 1`, with the transition at `L = 1` exactly. **Falsified by:** a stable, reproducible intermediate equilibrium belief in a system whose loop gain has been independently estimated away from 1.

Three further statements in this paper are not hypotheses and are labelled where they appear. The postulate (Section 4.1) is a **postulate**. The correspondence with the storage law (Section 7.2) is a **claim about model form**, not about kinship or mechanism. The loop-gain model of Section 5.6 is a **formalisation of a verbal mechanism** — the mechanism is Merton's [Merton, 1948]; the stability arithmetic is this work's, and it is elementary.

34.3 3 Method

34.3.1 3.1 Badge discipline

Every claim in this paper is assigned to exactly one of three drawers and the assignment is printed with the claim, following the corpus's standing classification of established mathematics, established empirical result, and declared postulate [Curzi, 2026].

Drawer one, established mathematics. The truth-table arithmetic of Section 5, the bit counts, the series and parallel composition rules, the linearised pole condition of Section 5.6, and every numerical example. These are arithmetic and a reader can redo them.

Drawer two, established empirical result. The placebo findings of Section 8, stated with their designs and their limits, and the social-threshold measurement of Section 7.5, stated with the setting in which it was measured. **Where a result is contested, it is reported as contested in the same passage as the claim it qualifies.**

Drawer three, this work's postulates. Two: the postulate of Section 4.1 and its operational form in Section 7.1. Both are badged at first use, in the verification chain, and in the Conclusion. Nothing else in the paper is asserted without support.

34.3.2 3.2 The evidence-and-precedent format

Several claims in this corpus are supported by a modern measurement and accompanied by an older tradition that reached a structurally similar position by an independent route. **The two are not the same kind of thing and the format in Appendix B exists to stop the second being read as the first.** Each audit states what was measured, what was merely recorded, and what the pair can and cannot establish. **Two sources agreeing does not double the evidence when only one of them is evidence.** Where the evidence column is empty, the empty cell is printed.

34.3.3 3.3 The formula validation protocol

Every formula in this corpus passes seven gates in print: dimensions, scale type, limiting cases, fixed points, identifiability, falsification, and reproducibility. Section 6 runs all seven on `E = G B + eps`. **It does not pass all seven, and the gate it fails is named rather than finessed.** The scale-type gate draws on measurement theory [Stevens, 1946, Michell, 1997, Luce and Tukey, 1964]; the falsification gate on the preregistration literature [Nosek et al., 2018, Simmons et al., 2011, Gelman and Loken, 2013]; the reproducibility gate on the replication record [Open Science Collaboration, 2015, Ioannidis, 2005].

34.3.4 3.4 Identification: the objection this paper cannot answer

This is the most serious methodological problem in the paper and it is stated here rather than in a footnote.

The postulate's dynamical form in Section 5.6 is a reflexive loop: belief produces an effect, the effect is observed, and the observation feeds back into the belief. Where the belief is held in a group and the effect is a group-level quantity, that is exactly the corpus's central formal claim `T(x) = W(x, M(x))` [Curzi, 2026] — and it is exactly a linear-in-means reflexive specification. **Manski [1993] proves that the endogenous social effect in such a specification is not identified from observational data.** It cannot be separated from contextual effects or from correlated unobservables, however much data is collected. Shalizi and Thomas [2011] show the same confounding for observational network data, and Lyons [2011] show what happens when it is ignored.

The consequence for this paper is specific: the social gain `G_social` of Section 7.5 is not identified from observation of a group. No fit recovers it. The remedies are known and none

is free — randomise the treatment, find an instrument, exploit partial-population variation, or **stipulate the reference group and defend the stipulation**, which is the route Brock and Durlauf [2001] makes defensible and the most realistic one here.

This paper adopts no remedy. It therefore states that the applied reading of `G_social` is a **modelling position rather than an identified empirical result**, and Section 13 repeats it. The three other channels are less exposed, because their gains are estimated from components measured outside the group — an adherence differential from behaviour records, a search rate from historical hours-per-find, an effect-per-action from a design that varies the action without varying the belief — but the social channel has no such escape and it is not given one.

A second reflexive objection lands on this paper's own parameters. The framework asserts that systems change their behaviour when they are modelled, and this paper hands readers a protocol for managing their own beliefs. Lucas [1976] is the sharp statement: parameters estimated under one regime do not survive a change of regime, because the decision rules adjust. Hacking [1995] makes the same point for classifications of people, and Ferraro et al. [2005] documents the mechanism at work in economics with a warning this paper takes seriously — **the theories that became self-fulfilling in that literature made people behave worse by assuming the worst of them.** The position taken here is the narrow one: the gains of Section 7 are estimated within a declared window, for a declared population, and are not asserted to be structural. **Either the parameters are stable under the framework's own predicted response, or the framework undercuts itself; this paper says they are not asserted to be stable, and prices that accordingly in Section 13.**

A third bound comes from the size of the effect the literature actually supports. Jussim and Harber [2005] review thirty-five years of work after the classroom-expectancy studies and conclude that self-fulfilling prophecies are **real but small**, do not typically accumulate over time, often dissipate, and that most teacher expectations turn out to be *accurate* rather than causal. **The corpus cites the original expectancy result [Rosenthal and Jacobson, 1968]; it must cite this beside it, because the effect size a reflexive law implicitly assumes is larger than the literature supports.** The gains in Section 7 are deliberately small numbers, and that is why.

And a fourth bound is the ceiling on social predictability at all. With fifteen years of high-quality longitudinal data and 160 participating teams, life outcomes proved barely predictable [Salganik et al., 2020, Hofman et al., 2017]. A framework proposing to predict outcome changes from belief magnitudes is operating inside that ceiling and cannot exceed it.

34.3.5 3.5 What discriminates this model from a neighbouring one

Sobkowicz [2020] poses the question in one line: what observation would distinguish this model from a neighbouring model that fits equally well, and Durlauf [2005] asks for the prediction only this model makes. **The honest answer here is that two of the paper's predictions are discriminating and the rest are not.** H4 — blockade — is discriminating, because a model without a channel structure makes no prediction about what happens when a specific pathway is pharmacologically interrupted. H6 — bimodality under self-reinforcement — is discriminating, because a linear influence model predicts a smooth distribution and this one predicts a gap. **H1, H2, H3 and H5 are shared with any mechanism-based account of belief effects, and this paper claims no priority over such accounts.** What it claims is that the mechanism requirement has been made arithmetic.

34.3.6 3.6 Analogy discipline

The paper uses two borrowings from physics and engineering, each labelled at first use.

The leaky-integrator form of the storage law, reused for belief, is a shared mathematical form — the same first-order recursion governs both, with no claim that the underlying dynamics are the same object [Elaydi, 2005, Oppenheim et al., 1996]. Section 7.2 says so.

Loop gain and the stability condition `L < 1` are likewise a shared form, taken from first-order feedback analysis [Åström and Murray, 2008, Strogatz, 2015], and used because the recursion of Section 5.6 is a first-order recursion. **No physical result's authority is transferred to any psychological claim by either borrowing**, and the general hazard of doing so is documented [Mirowski, 1989, Cartwright, 1999].

No scaling or heavy-tail claim is made anywhere in this paper, so the distribution-fitting procedure such a claim would require [Clauset et al., 2009] does not apply. This is recorded as a check that does not apply rather than one that was passed.

34.4 4 The postulate

34.4.1 4.1 Statement, with its badge

> **[POSTULATE — 4.1] BELIEF CREATES REALITY IF THE BELIEF IS GROUNDED IN TRUTH.**
>
> Where:
>
> - **BELIEF** is a held proposition together with the disposition to act on it. It has a magnitude `B` on the interval `[0, 1]`: 0 is not held at all, 1 is held with full disposition to act. **`B` is not a probability and must not be read as one**; it is a normalised strength-of-holding.
> - **CREATES REALITY** means: produces a change in a measurable state of the world that would not have occurred otherwise, written `E`, in the units of whatever is measured.
> - **GROUNDED IN TRUTH** means: there exists a real, specifiable, physically or socially instantiated pathway from the holding of the belief to the state measured, and the propositions that pathway depends on are true. Section 7 names this pathway the **channel** and gives it a gain, `G`.
>
> The postulate asserts that `E` is a function of `B` through `G`; that `G > 0` requires the channel to exist and to be real; and — this is the half that does the work — that **`G = 0` forces `E = 0` at every value of `B` whatsoever.**

This is a postulate. It is asserted by this work. It has no external support as stated, it is not derived from anything, and the corpus's own coverage audit lists it among the claims for which no supporting literature exists. What a postulate owes its reader is the list of observations that would end it; that list is Section 5.5.

The last clause is why the sentence is a postulate and not a slogan. **It forbids something.** It forbids the outcome the unconditional version promises: an effect with no route.

34.4.2 4.2 The magnitude form, which is the form that gets used

The English sentence is the statement of record. The working form is:

```
E  =  G . B  +  eps

E     measured change in the world, in the outcome's own units
B     belief magnitude, dimensionless, in [0,1]
G     channel gain, in outcome units per unit belief
eps   noise, mean zero, with a measurable standard deviation

GROUNDING CONDITION:   G > 0  requires a named channel whose
                       constituent propositions are TRUE.
```

```
THE FORBIDDING CLAUSE: G = 0  =>  E = 0  for all B.
```

The linear form is a first-order approximation and is marked as such. Real channels saturate: a belief held at 1.0 does not produce ten times what a belief held at 0.1 produces once the effector is at its ceiling. Section 7.6 supplies a saturating form as an explicit modelling choice. The linear form is used here because **the argument concerns the difference between `G > 0` and `G = 0`, and that difference survives any monotone functional form.**

34.4.3 4.3 What "grounded in truth" is not

Three readings are ruled out before the paper proceeds.

Not "sincerely held." Sincerity is a property of the believer; grounding is a property of the world. A belief can be held with total sincerity and be grounded in nothing, and the postulate predicts, for that case, an effect of exactly zero.

Not "true." The belief itself need not be true, and this is the subtle part. The belief that a bank is about to fail can be false when formed and still be grounded, because the *channel* — depositors withdraw, fractional reserves are genuinely fractional — is real. **What must be true is the set of propositions the channel depends on, not the proposition believed.** Section 5.6 develops this; it is the whole content of the self-fulfilling prophecy [Merton, 1948, 1936].

Not "consistent with a respected framework." The demand is for a pathway: a named sequence of physical or social steps, each of which happens and can be independently checked, ending at the measured quantity. **If it cannot be written down as steps, `G` has not been established as positive and the machinery predicts nothing.**

34.5 5 The empirical content of the conditional

34.5.1 5.1 The unconditional version, taken literally, is already false

Take the unconditional claim at its word and formalise it over three binary observables:

```
B  belief is held           (0 / 1)
K  a real channel exists    (0 / 1)
E  a change in the world attributable to B occurs   (0 / 1)
```

Eight cells. The unconditional version `U` says: if `B = 1` then `E = 1`, and forbids every cell with `B = 1, E = 0`. The conditional version `C` says two things: if `B = 1` and `K = 1` then `E = 1`; and if `B = 1` and `K = 0` then `E = 0`.

```
 B    K    E     U allows?    C allows?
-----------------------------------------
 0    0    0        yes          yes
 0    0    1        yes          yes
 0    1    0        yes          yes
 0    1    1        yes          yes
 1    0    0        NO           yes
 1    0    1        yes          NO
 1    1    0        NO           NO
 1    1    1        yes          yes

cells forbidden:     2/8          2/8
permitted fraction:  0.7500       0.7500
content, -log2(permitted):  0.4150 bits    0.4150 bits
```

At this level of description the two versions have **identical empirical content**: each forbids a quarter of the space, each carrying 0.4150 bits. **That is inconvenient for the argument and is stated before anything is done with it. The conditional's advantage is not that it forbids more cells. It is which cells it forbids.**

`U` forbids `B=1, K=0, E=0`: a belief is held, no channel exists, nothing happens. **That is the single most frequently instantiated cell in human experience** — every wish that did not come true, every conviction held for a decade about a matter its holder could not touch. Taken literally, the unconditional version is not merely unsupported. It is refuted daily, at enormous sample size, by everyone.

`C` forbids `B=1, K=0, E=1`: a belief is held, no channel exists, and the world changes anyway in a way attributable to the belief. **That cell is the one the conditional stakes itself on. It is rare, and it is exactly the cell that would matter if it were ever cleanly observed.**

34.5.2 5.2 The unconditional version, taken charitably, has zero content

Nobody defends the unconditional version literally. The defence always takes the same shape: the belief was not held hard enough; some part of the believer resisted; there was a limiting belief underneath. Formalise that as a hidden variable `S` — sufficiency of belief — and re-run the table.

```
U' :  (B = 1) AND (S = 1)  =>  E = 1,      with S UNOBSERVABLE

project the sixteen (B,K,E,S) cells onto the observable (B,K,E):
a cell is forbidden only if it is forbidden for EVERY value of S.

observable cells forbidden by U' :  0 / 8
content of U'                    :  0.0000 bits
```

Zero. Not small — zero. An unobservable sufficiency parameter converts a false proposition into an empty one, and this is the standard mechanism by which doctrines survive their own counterexamples. Every failure is absorbed by `S = 0`; every success confirms `S = 1`; nothing is learned. It is the failure of asymmetry in its purest form: `U'` and its own negation predict the same observations, which is what Popper [1959] and Platt [1964] identify as the mark of a proposition outside empirical reach, and what Lakatos [1970] describes as a research programme protected by auxiliary hypotheses rather than by results.

Demoted, in print: the unconditional proposition *belief creates reality*, in its defended form, is **a heuristic at best** — a mood-management device — and carries no evidential weight and no predictive content whatsoever. It may not be cited in support of anything in this corpus.

The conditional cannot pull the same trick, because its condition is about the world rather than about the believer. `G > 0` is checked by naming a channel and verifying its steps. It cannot be rescued by asserting that the believer's belief was defective, because the believer's state is not what the condition ranges over. **Putting the condition outside the believer is what makes it a condition rather than an excuse.**

34.5.3 5.3 The magnitude form, and how much content it has

The binary table understates the conditional badly. The conditional does not merely say that an effect occurs; it says `E = G B`, and once `G` is estimated it has predicted a *number*. Let a standardised measured effect lie in `[−1, +1]`, width 2:

```
version                              permitted band   excluded    content (bits)
-------------------------------------------------------------------------------
unconditional, with escape clause           2.000        0.0000        0.0000
```

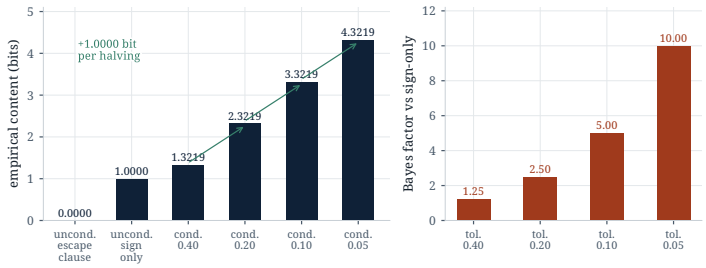


Figure 34.1: **Each halving of the prediction tolerance buys exactly one bit, and the unconditional form with its escape clause buys none.** Empirical content in bits, computed as the negative base-two logarithm of the permitted fraction of a standardised effect band of width 2: 0.0000 bits for the unconditional form with an unobservable sufficiency parameter, 1.0000 for a sign-only prediction, and 1.3219, 2.3219, 3.3219 and 4.3219 for point predictions at tolerances of 0.40, 0.20, 0.10 and 0.05. The Bayes factors against the sign-only version are 1.25, 2.50, 5.00 and 10.00. Computed from the tables of Section 5.3; generated by `figs/src/CCP-034_fig01.py`.

unconditional, sign only (E > 0)	1.000	0.5000	1.0000
conditional, point prediction +/- 0.40	0.800	0.6000	1.3219
conditional, point prediction +/- 0.20	0.400	0.8000	2.3219
conditional, point prediction +/- 0.10	0.200	0.9000	3.3219
conditional, point prediction +/- 0.05	0.100	0.9500	4.3219

Each halving of the prediction tolerance buys exactly one bit of content. That is the same logarithmic arithmetic that governs the rest of the framework and appears here for the same reason: content is a logarithm of an excluded fraction, exactly as surplus is a logarithm of an effective count and pitch a logarithm of a frequency. The information-theoretic reading is Shannon's [Shannon, 1948, Cover and Thomas, 2006].

The discrimination this buys, as a likelihood ratio for an observation landing inside the conditional's band, against the sign-only version's flat spread over `(0, 1]`:

tolerance	Bayes factor C : U	log10 BF
+/- 0.40	1.25	0.097
+/- 0.20	2.50	0.398
+/- 0.10	5.00	0.699
+/- 0.05	10.00	1.000

These are modest numbers, given honestly. One observation inside a `+/- 0.10` band is worth a factor of five — evidence, not proof. **What is not modest is the other side of the ledger: an observation outside the band gives a Bayes factor of zero. The conditional dies; the sign-only version does not, because almost nothing is outside its band.** That asymmetry is the whole case for attaching a condition, and it is what Meehl [1967] identified as the difference between a theory that risks a point prediction and one that risks only a direction.

34.5.4 5.4 The exposed flank, priced

Consider a hundred belief interventions, of which forty have a real, nameable channel to the outcome and sixty do not.

trials	100
with a real channel	40

```
without a real channel               60

unconditional predicts an effect in     100 arms
conditional   predicts an effect in      40 arms
conditional   predicts NO effect in      60 arms   <-- the exposed flank
```

Sixty per cent of the design is a region where the conditional has committed itself to finding nothing, and where finding something kills it. The unconditional version has no such region; it is undefeated because it never takes the field. Stated as a general rule:

> **The value of a condition is measured by the size of the region in which it predicts zero.**

And the price, computed against the paper's own claim rather than in its favour. Under the same base rates:

```
P(a random observation refutes the unconditional version)  =  0.3250
P(a random observation refutes the conditional  version)  =  0.0550

observations needed for a 95% chance of at least one refutation:
    unconditional     8
    conditional      53
```

Eight observations. A proposition refuted with 95 per cent confidence by eight casual observations is not a bold conjecture; it is a mistake that survives because nobody applies it to the cases where it fails. **And the conditional's own figure — 5.5 per cent per observation, 53 observations to a 95 per cent chance of refutation — is the honest measure of how much harder this paper's claim is to kill than the one it replaces.** A condition that only ever narrows the circumstances in which success is claimed is a hedge. **This one creates a large region in which failure has been promised, and that is what makes it worth having — but it also makes the claim six times harder to refute, and that cost is printed here rather than left for a reviewer to compute.**

34.5.5 5.5 What would end the postulate

Per the badge convention, any one of the following, cleanly observed and replicated, ends the postulate as stated. These are the falsifiers behind H1 to H5.

```
K1  ZERO-CHANNEL EFFECT.
    A belief-linked effect clearly above the noise floor, on a state of the world for
    which every candidate channel has been enumerated and independently excluded --
    no physiological, behavioural, social or informational route, no measurement
    artefact, no selection effect. Replicated by an independent group, preregistered.
    THIS IS THE ONE THAT MATTERS. The rest is calibration.

K2  MAGNITUDE BLOWOUT.
    Effects reliably and substantially exceeding the summed gain bound of Section 7,
    with gains estimated in advance and frozen. One instance is a missed channel;
    a pattern is a dead postulate.

K3  CHANNEL INDEPENDENCE.
    Effect magnitude independent of channel gain across channels with well-separated,
    independently measured gains. If halving the gain does not halve the effect,
    G is not what is doing the work.
```

```
K4  BLOCKADE FAILURE.
    Effects surviving the specific, verified interruption of the only named channel.
    A channel that can be cut without changing the outcome was not the channel.

K5  GROUNDING IRRELEVANCE.
    Beliefs whose channels rest on demonstrably false propositions producing the same
    effects, at the same magnitudes, as beliefs whose channels rest on true ones.
```

K5 deserves a note. A false belief with a real behavioural channel *does* produce effects — the runner who wrongly believes a supplement works still trains harder. Section 4.3 conceded this: what must be true is the channel's constituent propositions, not the believed one. **If channels whose propositions are false performed as well as channels whose propositions are true, "grounded in truth" would be doing no work and should be struck from the sentence.**

34.5.6 5.6 The fixed-point structure, and why a false belief can still be grounded

The bank-run case is why the postulate says *grounded* and not *true*, and it has an exact dynamical form. Let belief update on the evidence its own effects produce:

```
B_{t+1} = (1 − lam) B_t + lam . min(1, L . B_t)

lam   how fast belief updates on evidence, in (0,1]
L     LOOP GAIN = (effect per unit belief) x (belief produced per unit effect)
      L = alpha . G, where alpha is how strongly the observed effect feeds back
      into the belief.
```

This recursion is a formalisation of a verbal mechanism, not a derivation of one. The mechanism is Merton [1948]'s; the stability arithmetic is this work's and it is elementary. Linearising at `B = 0` gives a pole at `1 − lam + lam L`, so the zero-belief state is stable if and only if `L < 1`. Computed at `lam = 0.30`, iterated 2000 steps from `B_0 = 0.05`:

```
  L      pole      B = 0 stable?     B* reached from B_0 = 0.05
-----------------------------------------------------------------
 0.00    0.7000        YES              0.000000
 0.50    0.8500        YES              0.000000
 0.90    0.9700        YES              0.000000
 1.00    1.0000        no               0.050000   (marginal)
 1.10    1.0300        no               1.000000
 1.50    1.1500        no               1.000000
 2.00    1.3000        no               1.000000
 3.00    1.6000        no               1.000000
```

A self-fulfilling belief is exactly a belief whose loop gain exceeds one. Below one it decays to nothing however strongly it starts; above one it saturates the effector however weakly it starts. The transition at `L = 1` is sharp, and it is the same first-order pole condition as the storage law's contraction modulus with the feedback sign reversed [Curzi, 2026a]. The corresponding operational reading in the contraction literature is that distance to arrival is the last correction scaled by `q/(1−q)`, so the size of one's most recent revision is a self-reading gauge [Curzi, 2026].

And `L = alpha G`. Set `G = 0` and the loop gain is zero at every value of `alpha`, so the fixed point is `B* = 0` and nothing happens, forever, at any intensity of belief. The dynamical statement and the logical statement are the same statement.

Now the bank. The depositor's belief that the bank will fail may be false when formed. **It is grounded anyway, because the channel is real — and it is real because of a true fact about the system.**

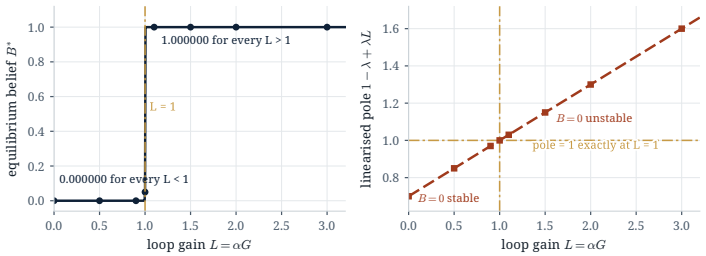


Figure 34.2: **A self-fulfilling belief is exactly a belief whose loop gain exceeds one, and the transition is sharp.** Equilibrium belief `B*` reached from `B_0 = 0.05` under `B_{t+1} = (1−lam)B_t + lam min(1, L B_t)` at `lam = 0.30` over 2000 steps: 0.000000 for `L = 0.00`, `0.50`, `0.90`, marginal at `L = 1.00`, and 1.000000 for `L = 1.10` and above. The linearised pole `1 − lam + lam L` crosses 1 at exactly `L = 1`. Because `L = alpha G`, a zero channel gain forces `L = 0` at every feedback strength. Modelled from the recursion of Section 5.6 with the parameters stated there; generated by `figs/src/CCP−034_fig02.py`.

```
reserve ratio r      bank becomes illiquid once withdrawing fraction w > r
------------------------------------------------------------------------
    0.03                              w > 0.03
    0.05                              w > 0.05
    0.10                              w > 0.10
    0.20                              w > 0.20
    1.00                              w > 1.00     <-- unreachable
```

A full-reserve bank has `G = 0` for this belief, and the identical rumour, held with identical conviction, does nothing at all. The belief is the same; the truth of the system is different; the outcome is different. That is the conditional clause, empirically instantiated, in the most famous example in the sociology of belief.

And the corpus's strongest available existence proof for a reflexive fixed point observed in the field is not the bank run but the options market. MacKenzie and Millo [2003] document a pricing model that did not fit the market when published, and a market that came to fit the model after traders adopted it, with dates and price data; MacKenzie [2006] and Callon [1998] give the book-length case and its theoretical frame. **That is a description becoming true by being used, measured, in a setting where the channel — traders acting on the model's outputs — is nameable step by step.** Soros [2013] is the closest existing named doctrine to the corpus's reflexive law and supplies the concept without an equation; what this paper adds is a loop gain and a stability condition, and what it does not add is identification.

34.6 6 The formula validation protocol, all seven gates

The working form `E = G B + eps` is a formula and is therefore subject to the corpus's seven-gate protocol. **It does not pass all seven cleanly. The failure is reported rather than concealed, and Section 7 exists to repair it.**

Gate 1 — Dimensions.

```
E    outcome units [X]          (days, millimetres of mercury, defects found,
                                 points on a validated symptom scale, currency)
B    dimensionless, [0,1]
G    [X] per unit belief
eps  [X]

Check:  [X] = ([X] / 1) . (1) + [X]  =  [X].    CONSISTENT.
```

Because `B` is dimensionless, `G` inherits `E`'s units, so **`G` is not comparable across outcomes unless both are standardised.** A gain of 0.8 days and a gain of 0.8 symptom points are not the same quantity and must never be summed. The remedy used throughout is to standardise `E` by the outcome's own standard deviation before any cross-channel arithmetic, and to state the standardisation whenever a gain is quoted. **PASS, with that condition attached and binding.**

Gate 2 — Scale type. `E` must be at least interval-scaled for the subtraction implied by "would not have occurred otherwise" to mean anything, and ratio-scaled with a true zero for the multiplication `G B` to mean anything [Stevens, 1946]. Many outcomes qualify: elapsed days, counts, physical measurements, currency. Many do not: ordinal symptom ratings, Likert satisfaction [Likert, 1932], ranked preferences.

`B` is the hard case. Belief magnitude as normally elicited is ordinal at best, and an ordinal `B` makes `G B` a heuristic ordering rather than an equation. Two routes out, both established:

- **Magnitude estimation.** Elicit belief strength by ratio judgements against a standard rather than by category selection [Stevens, 1957, 1975]. This produces ratio-scale data from human reports and it replicates; it is the route with the strongest warrant, though the range and context effects that attend it are real [Poulton, 1979].
- **Behavioural anchoring.** Operationalise `B` by a decision the believer makes — the stake accepted, the hours committed, the alternative declined. Ratio scale with a true zero by construction, **at the cost of confounding belief with the very behaviour the behavioural channel is supposed to explain.** Using the behavioural anchor and then claiming the behavioural channel is circular and is named as a failure mode in Section 9.

A third and deeper objection stands over both routes. Michell [1997] argues that psychology routinely assumes quantitative structure without testing for it, and the test that settles the question is conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971]. **It has not been run for `B` and it has not been run for any construct in this corpus. VERDICT: CONDITIONAL PASS** with magnitude-estimated `B` and ratio-scaled `E`, and a **fail**, demoted in print to a heuristic ordering, whenever `B` is a Likert item — which is most of the published work on belief and outcome.

Gate 3 — Limiting cases.

```
B -> 0      E -> 0.       No belief, no effect. Trivially observed.

B -> 1      E -> G.       The most a belief can do is the channel's full gain.
                          NAMEABLE: a ceiling, and the postulate's central
                          practical content.

G -> 0      E -> 0 for all B.   NAMEABLE, and the whole point: the ungrounded
                                belief. Every wish about a matter one cannot touch.

G -> inf    E -> inf.     NOT NAMEABLE -- a genuine defect of the linear form. No
                          real channel has unbounded gain. Section 7.6 supplies
                          E = E_max . tanh(G.B / E_max) as a saturating form,
                          stated as a modelling choice, not a derivation.

eps -> large   E unrecoverable. NAMEABLE: the effect exists and the study cannot
                          see it. The commonest real-world outcome, and a fact
                          about the design, not about the postulate.
```

VERDICT: PASS with one repair. The `G -> inf` divergence is a defect of the linear approximation and is repaired by saturation. **Recorded as a repair, not hidden.**

Gate 4 — Fixed points. The **static invariant** is `G = 0 => E = 0`, the equation's null space and the postulate's entire forbidding content. The **dynamical fixed point** is Section 5.6's `B*` — a bifurcation rather than a point: `B* = 0` for `L < 1`, `B*` at the effector ceiling for `L > 1`, transition at

`L = 1` exactly. **The system has no stable intermediate belief under pure self-reinforcement**, which is H6 and is a substantive, testable prediction: beliefs in high-gain reflexive loops should be bimodally distributed in a population rather than spread smoothly. That prediction is shared with threshold and cascade models of collective behaviour [Granovetter, 1978, Schelling, 1971], which is a point against its discriminating power and is noted in Section 3.5. **VERDICT: PASS**, and this gate produced the paper's best prediction.

Gate 5 — Identifiability. Here is the trouble, and it is serious.

`E = G B` is **not identifiable from outcome data alone.** Halve `G`, double `B`, and every observation is unchanged. Two distinct parameter sets predict identical data, which means the formula is not testable as written. The repair, and it is the only repair:

> **`G` must be estimated INDEPENDENTLY of `E`, from the channel's own mechanism, and frozen before the outcome is measured.**

That is what Section 7's four channel classes are for. The behavioural gain is computed from an adherence differential and a per-unit effect, measured separately; the informational gain from a search rate and a time reallocation, measured separately. **Each is a prior, mechanism-derived estimate of `G` that never touches the outcome data, and that independence is what converts an unfalsifiable identity into a prediction.** Where `G` cannot be estimated independently, the formula is unidentifiable and must be reported as such.

VERDICT: FAIL as written; PASS only under independent, preregistered, frozen estimation of `G`. This condition is binding on every use of the postulate in this corpus. And it must be read alongside Section 3.4: **independent estimation repairs the parametric non-identifiability of `E = G B`; it does not repair the reflection problem, which is a separate and unrepaired failure affecting the social channel specifically.**

Gate 6 — Falsification.

```
1  PREREGISTER   belief, named channel, mechanism-derived G, predicted E with tolerance,
                 outcome measure, sample size, stopping rule. Before any outcome.
2  FIT           on training data only, if any parameter beyond G is free.
3  FREEZE        all parameters; publish the frozen values.
4  HELD-OUT      report performance on data not used in fitting.
5  FOUR BASELINES, all required:
                 (a) persistence      -- the outcome does what it was already doing
                 (b) direct-input     -- outcome is a rescaled copy of the belief measure
                 (c) matched AR       -- one-parameter memory model, no theory in it
                 (d) SHUFFLED-INPUT   -- same frozen parameters, belief series permuted
6  KILL RULES    K1-K5 of Section 5.5, plus: failure to beat (d) kills the application.
```

Baseline (d) decides it. If the model predicts as well with the belief series shuffled as aligned, the apparent coupling was a coincidence of two trends and there is no channel. **The shuffled-input control is the most omitted and most decisive test in this literature.** The corpus's own application of this protocol to a different law is Curzi [2026i], whose preregistered fit **failed its second falsification rule** — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin, despite +40.3 per cent over persistence and +97.7 per cent over direct input and `p = 0.0005` against the shuffle. **VERDICT: PASS** — the protocol exists, is specific, and is executable. Whether a given application passes is answered per application, and the one application so far reported in this corpus did not.

Gate 7 — Reproducibility. Same operationalisation of `B`, same channel, same `G` estimation procedure, different people, different laboratory: do the parameters recover?

For the **physiological** channel, **partially yes** — the effect has been observed across many laboratories and decades and the blockade finding has been extended, though magnitudes vary considerably across studies, populations and outcomes (Section 8). For the **behavioural** and **informational** channels, **not yet tested in this form**, because nobody has run this protocol. For the **social** channel, the underlying phenomena replicate but the gain parameter has not

been recovered across laboratories in this notation — **and Section 3.4 gives the deeper reason it may not be recoverable at all.**

The wider context is mandatory: 97 of 100 original psychology studies reported significant results while about 36 per cent of replications did, with effect sizes roughly halved [Open Science Collaboration, 2015]. **Any belief-effect literature is drawn from that population and inherits its problems** [Ioannidis, 2005]. **VERDICT: PARTIAL.** Passing for one channel in one domain; untested for three of four.

```
FVP SUMMARY

1  DIMENSIONS       PASS, with a standardisation condition
2  SCALE TYPE       CONDITIONAL PASS; fails on Likert-elicited belief; conjoint
                    measurement never run for B
3  LIMITING CASES   PASS, with the saturation repair recorded
4  FIXED POINTS     PASS, and it generated the bimodality prediction
5  IDENTIFIABILITY  FAIL as written; PASS only with G estimated independently and
                    frozen. BINDING CONDITION ON ALL USE. Does NOT repair the
                    reflection problem.
6  FALSIFICATION    PASS as a protocol; the corpus's one application of it failed
                    its own declared margin
7  REPRODUCIBILITY  PARTIAL. One channel, one domain, with caveats.
```

Five and a half of seven, and the failure is named. A postulate that reported seven of seven would be a postulate nobody had audited.

34.7 7 The coupling requirement: channels, gains, and composition

34.7.1 7.1 The operational form, badged

> **[POSTULATE — 7.1, the operational form of 4.1]** A belief `B` affects a world-state `E` only through one or more channels. Each channel `i` has a gain `G_i` — the change it produces in `E` per unit of `B` — determined by the channel's own mechanism and not by the belief. Then
>
> ```
> E  =  ( sum_i G_i ) . B  +  eps        for small B, before saturation
>
> and the BOUND:      |E|  <=  ( sum_i G_i ) . B_max
> ```
>
> **If no channel exists, the sum is empty, `sum_i G_i = 0`, and `E = 0` at every `B`.**

The badge stays on. This is a postulate, not a theorem, and what would end it is K1 of Section 5.5 — an effect with every candidate channel enumerated and excluded.

But the bound is not a postulate once the gains are estimated. Given estimated gains, the bound is arithmetic, and arithmetic is what makes a claim auditable. A person claiming their belief will move an outcome by three standard deviations, through channels whose summed gain estimates to 0.4, has made a claim that fails before any data are collected.

34.7.2 7.2 The correspondence with the storage law, and exactly what it does and does not transfer

```
STORAGE LAW          Q_{t+1} = (1 - delta) Q_t + eta . S_t - C_t
```

```
COUPLING RELATION     E = G . B


TERM       STORAGE LAW                    COUPLING RELATION
-------------------------------------------------------------------------
INPUT      S_t, surplus offered           B, belief held
           bounded: 0 <= S_t <= ln N      bounded: 0 <= B <= 1

CONVERSION eta, how much input becomes    G, how much belief becomes
           durable structure              world-change
           eta = 0 -> nothing sticks      G = 0 -> nothing happens

OUTPUT     what lands in Q_t, the         E, the measured change in the world
           stored order

LOSSES     delta Q_t (decay) and C_t      eps (noise) and, over time,
           (cost)                         the same delta and C
```

This is a claim about model form and is badged as one. The same first-order structure describes both; no claim is made that belief and stored order are the same kind of quantity, that they can be aggregated, or that a result proved about one is thereby proved about the other. **The storage law is a proven theorem about a recursion, with a contraction modulus `1 - delta` and a steady state `(eta S - C)/delta` obtained as the circuit's DC gain `1/delta` applied to the net drive [Curzi, 2026a]. The coupling relation is a postulate about the world. The form transfers; the warrant does not.** The corpus states the same badge for its three-scope correspondence generally: identical structure at three scopes is a statement about the chosen model, not a claim of kinship or of cross-scope aggregation [Curzi, 2026].

Four consequences read directly off the law and transfer without modification.

Consequence 1 — the input is capped. `S_t <= ln N` is a hard ceiling and `B <= 1` is the same kind of ceiling. Neither system can be driven arbitrarily hard from the input side. **Believing harder runs out before it starts.**

Consequence 2 — conversion is the cheap lever. The corpus's standing claim is that `eta` is the most neglected parameter and usually the cheapest to improve. The same holds here. **Doubling `B` requires changing a person's mind; doubling `G` requires changing an arrangement, and arrangements are easier to change than minds.**

Consequence 3 — one-step growth has a threshold. `Q_{t+1} > Q_t` if and only if `eta S_t > delta Q_t + C_t`. Translated: **belief does not merely need a channel, it needs a channel with enough gain to beat the decay and the cost of holding the belief.** Section 7.8 computes the resulting minimum viable gain.

Consequence 4 — headroom governs. `h_t = G ln N - delta Q_t - C_t` applies unchanged with `eta -> G`. **If `h_t < 0`, no belief of any strength improves the system, because the model's own arithmetic says nothing can. A person in a negative-headroom situation who is told to believe harder is being given the one instruction that provably cannot work.** The corpus's five-system survey found that the binding constraint differs by system rather than by size — 41.64 per cent of ceiling where the constraint is `u`, 28.59 per cent where it is acoustically caused, 69.28 per cent with `h = -0.654` where it is cost, and 91.03 per cent of ceiling still negative at `h = -0.273` where it is conversion [Curzi, 2026k].

34.7.3 7.3 Composition: series multiplies, parallel adds

Series. A channel with several stages has the product of the stage gains.

```
stage                                    stage gain    cumulative
-------------------------------------------------------------------
```

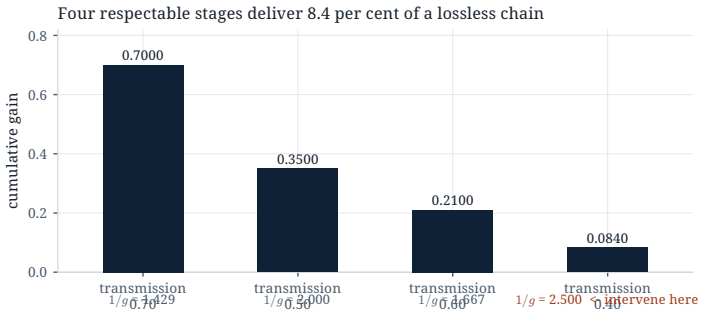


Figure 34.3: **Series composition is brutal: four respectable stages deliver 8.4 per cent of a lossless chain.** Cumulative gain falls 0.7000, 0.3500, 0.2100, 0.0840 across stages of transmission 0.70, 0.50, 0.60 and 0.40; relative sensitivity `1/g_i` identifies the 0.40 stage, at 2.500, as the one where intervention pays. Any single stage set to zero takes the product to zero. Computed from the stage table of Section 7.3; generated by `figs/src/CCP-034_fig03.py`.

```
belief registered as expectation            0.70        0.7000
expectation reaches the effector            0.50        0.3500
effector produces the change                0.60        0.2100
change persists to measurement              0.40        0.0840

TOTAL SERIES GAIN                                       0.0840
```

Four stages, each individually respectable — the worst is 0.40 — and **the chain delivers 8.4 per cent of what a lossless chain would. Series composition is brutal, and it is the reason that long, plausible-sounding causal stories deliver almost nothing. Every "and then" in a mechanism is a multiplication by a number less than one.**

```
relative sensitivity of the total to stage i  =  1 / g_i

belief registered as expectation      1.429
expectation reaches the effector      2.000
effector produces the change          1.667
change persists to measurement        2.500   <-- fix this one first
```

The weakest stage has the highest leverage, and the arithmetic identifies it without argument. And **any single stage at zero makes the total zero regardless of the others**, which is the coupling requirement in its most compact form.

Parallel. Independent channels add, and **the sum is a bound, not a promise.**

```
channel                                      gain G_i
-----------------------------------------------------
physiological                                  0.30
behavioural                                    0.55
social                                         0.25
informational                                  0.20
-----------------------------------------------------
SUM                                            1.30

bound on |E|:    B = 0.2  ->  0.260
                 B = 0.5  ->  0.650
                 B = 0.8  ->  1.040
                 B = 1.0  ->  1.300
```

Additivity is an upper limit that assumes the channels do not interfere, and real channels frequently do — the behavioural channel consumes time the informational channel needed; the social channel's conformity pressure suppresses the individual belief driving the others. **If a measured effect exceeds the summed bound, one of three things is true: a channel was missed, a gain was underestimated, or the measurement is wrong. There is no fourth option that does not damage the postulate.**

34.7.4 7.4 Why exactly four classes

The taxonomy is exhaustive by argument rather than by decree: **a belief is a state of a nervous system, and there are exactly four kinds of thing a state of a nervous system can do to the world outside the skull.**

```
1  It can change the body it is inside.                    PHYSIOLOGICAL
   Route: expectation -> neural and endocrine state -> tissue.

2  It can change what that body does.                      BEHAVIOURAL
   Route: belief -> action selection -> the world acted upon.

3  It can change what other nervous systems do, by being   SOCIAL
   communicated and held in common.
   Route: belief -> expression -> others' beliefs -> collective action.

4  It can change what that body notices and therefore      INFORMATIONAL
   what it finds.
   Route: belief -> attention and search allocation -> discovery of things
   that were already there.
```

Anything a belief does falls into one of these, into a combination, or into a fifth option — **direct action on the world with no bodily route at all, the class for which no channel has ever been specified and which the postulate therefore predicts to have gain exactly zero.** Section 7.9 handles it.

Class 4's phrasing is exact and load-bearing. The informational channel does not create what is found. It changes the probability of finding what is there.

34.7.5 7.5 The four channels, their gains, and how each is estimated

The physiological channel. *What it can do:* change measurable states of the believer's own body — reported pain, autonomic tone, hormone release, motor performance under fatigue. *What it cannot do:* change another person's body at a distance; change tissue for which no efferent pathway exists; produce effects larger than the effector's own maximum. **There is no efferent pathway from expectation to, for example, the size of a solid tumour, and the postulate therefore predicts zero there.** Its gain, from Section 8:

```
G_physio = measured standardised effect / normalised expectancy

d = 0.30, B = 1.00  ->  G = 0.3000
d = 0.30, B = 0.70  ->  G = 0.4286
d = 0.20, B = 0.50  ->  G = 0.4000
d = 0.15, B = 1.00  ->  G = 0.1500

WORKING RANGE, subjective symptom-based outcomes:  0.15 to 0.45
WORKING RANGE, objective disease endpoints:        approximately 0
```

Both lines matter and the second as much as the first: the gain is real and modest for what the person reports, and approximately zero for endpoints an instrument reads independently of the person. Estimation requires three steps, all binding: **E1** name the efferent pathway from the central state to the measured tissue, and estimate `G = 0` if it cannot be named; **E2** find the published effect size for that pathway and outcome class, from a three-arm design that separates the channel from natural history; **E3** divide by the expectancy that produced it, measured on the same scale to be used for `B`, and report a range rather than a point where the source study did not measure expectancy. **A physiological gain quoted without a named efferent pathway is not an estimate.**

The behavioural channel — the largest of the four by a wide margin and the least discussed. *What it can do:* anything a person's actions can do, which includes almost everything anyone ever wanted from believing something. *What it cannot do:* anything actions cannot do. **Here the conditional clause bites hardest: the behavioural gain contains the effectiveness of the action as a multiplicative factor, so a belief that reliably produces an action that does not work has gain exactly zero.**

```
G_beh  =  (delta p)  x  (effect per unit action)

adherence without the belief    p0 = 0.55
adherence with the belief       p1 = 0.80          delta p = 0.25
effect of the practice at full adherence           0.50 standardised

G_beh = 0.25 x 0.50 = 0.1250

AND THE CONDITIONAL CLAUSE, in one line:
if the practice does nothing (effect at full adherence = 0):  G_beh = 0.0000
at every value of delta p, however large.
```

A belief that produces perfect adherence to a useless practice has gain zero. That single line is the difference between this framework and motivational literature, and it is checkable. Estimation: **B1** measure `delta p` from behaviour records — attendance logs, commit history, sensor data — never from self-reported intention, which correlates with the belief measure by construction and produces a circular estimate; **B2** measure the effect per unit action separately, in a design that varies the action without varying the belief; **B3** multiply, and **report both factors, not only the product**, so a reader can see that a gain of 0.125 was 0.25 of a real effect rather than 0.90 of a negligible one.

The social channel. *What it can do:* coordinate — change what a group will attempt, treat as normal, and enforce. *What it cannot do:* anything the group could not do if coordinated by other means. **The social channel is a coordination device, not a power source**, and its ceiling is the group's actual capability. It also has a specific failure: it raises coordination by lowering the non-parallelism `u` that generates the group's surplus in the first place.

Its gain is the only one of the four whose response is sharply nonlinear in the fraction of the group holding the belief.

```
group N = 240
committed fraction    people
------------------------------
    5%                   12
   10%                   24
   15%                   36
   20%                   48
   25%                   60      <-- the measured tipping region
   30%                   72
```

A committed minority of approximately 25 per cent has been measured to overturn an established social convention in controlled online experiments [Centola et al., 2018]. **Trans-**

fer to a field organisation is a modelling assumption, not a measurement, and is labelled as such every time the number is used. Estimation: **S1** count the committed, not the agreeable — committed means acting on it when it costs something, and survey agreement is the standard way this estimate is inflated; **S2** map the topology, because complex contagion needs multiple reinforcing contacts, so clustered placement beats scattered placement at the same headcount [Centola and Macy, 2007, Centola, 2010]; **S3** compare to threshold and state the transfer assumption aloud; **S4** below threshold, estimate `G_social` as small and decaying, and never average across the threshold, because the response is not linear there. Related threshold and cascade results bound what the channel can do at all [Granovetter, 1978, Bikhchandani et al., 1992].

The honesty clause this channel requires: it is the one with a coercion failure mode. A belief propagated through a group can be propagated by pressure, and pressure removes the ability to decline. **Any protocol that raises `G_social` by removing a person's option to dissent is out of scope for this framework**, and it destroys `u` in the bargain. **And Section 3.4 applies here and nowhere else with full force: `G_social` is not identified from observation of a group.**

The informational channel — the subtlest of the four and the one where the conditional clause is most cleanly visible. *What it can do:* change what a person attends to, searches for, and therefore finds. *What it cannot do:* **put anything there. The informational channel's gain is exactly zero when the target is absent, at every level of belief, effort and duration. This is not a caveat; it is the channel's defining equation.**

```
P(find) = 1 − exp(−r . t),  with r = 0.05 per hour, TARGET PRESENT

    hours t      P(find)
−−−−−−−−−−−−−−−−−−−−−−−−−
         5        0.2212
        10        0.3935
        15        0.5276
        25        0.7135
        40        0.8647
        60        0.9502

belief raises directed search from 10 h to 25 h:
    delta P(find) = 0.7135 − 0.3935 = 0.3200

TARGET ABSENT:  P(find) = 0.0000 at t = 5, 10, 15, 25, 40, 60, and at every other t.
```

Left column identical, right column zero. The same belief, the same conviction, the same reallocated hours, and a gain of 0.32 in one case and 0.00 in the other — **and the difference is entirely a fact about the world, not about the believer.**

This is also the honest explanation of a large fraction of the anecdotes that circulate as evidence for the unconditional version. Someone believed opportunities were everywhere, looked much harder, found one, and reported that belief had created it. **The belief did not create it. The belief created the search. The search found what was there. And in the cases where nothing was there, nothing was found and no anecdote was generated** — which is survivorship bias operating on the evidence base for the entire genre. Estimation: **I1** establish the base rate that the target exists, the step almost never taken; **I2** measure the search rate `r` from a comparable completed search; **I3** measure the time reallocation from time records; **I4** compute `G_info = P(exists) [exp(−r t0) − exp(−r t1)] x (value per find)`. **The believer contributes the time reallocation. The world contributes `P(exists)` and `r`. The believer's contribution cannot rescue a zero in the world's.**

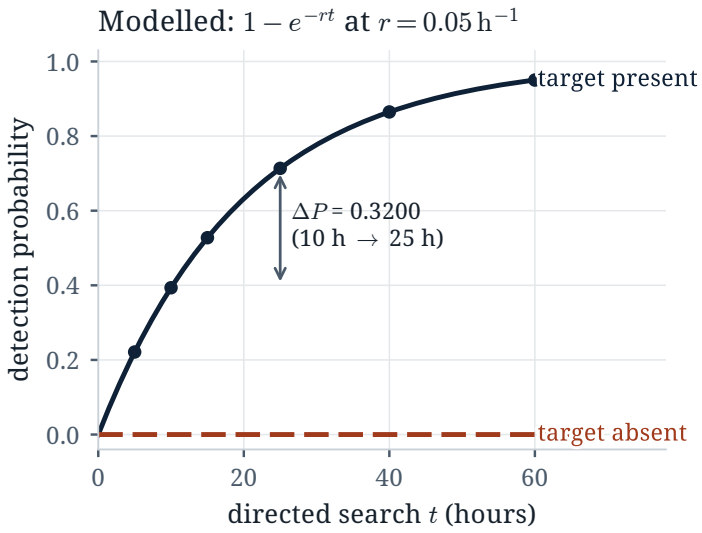


Figure 34.4: **The informational channel's gain is a fact about the world, not about the believer.** Detection probability `1 - exp(-r t)` at `r = 0.05` per hour with the target present rises 0.2212, 0.3935, 0.5276, 0.7135, 0.8647, 0.9502 at 5, 10, 15, 25, 40 and 60 hours, so a belief that raises directed search from 10 to 25 hours buys `delta P = 0.3200`. With the target absent the same curve is 0.0000 at every `t`. Modelled as a Poisson search process with the parameters of Section 7.5; generated by `figs/src/CCP-034_fig04.py`.

34.7.6 7.6 Saturation

Gate 3 flagged `G -> inf` as a defect. The repair, stated as a modelling choice and not a derivation:

```
E  =  E_max . tanh( G . B / E_max )        [MODELLING CHOICE, not a derivation]

small G.B:   E -> G.B                (recovers the linear relation)
large G.B:   E -> E_max              (the effector's ceiling)
```

`E_max` is the effector's physical or social maximum. **Every real channel has an `E_max`, and quoting a gain without quoting the ceiling it saturates against is an incomplete specification.**

34.7.7 7.7 The four gains side by side

CHANNEL	TYPICAL GAIN	CEILING (E_max)	WHAT ZEROES IT
physiological	0.15–0.45 on subjective outcomes; ~0 on objective disease endpoints	the effector's own maximum output	no efferent pathway to the measured tissue
behavioural	widest range; 0.1–0.6 typical in the worked cases here	the person's time and capability	the action is ineffective (effect per action = 0)

social	bimodal: small below threshold, large above	the group's actual capability	below the committed-fraction threshold, or the group cannot do the thing at all
informational	0 to ~0.35; entirely conditional on the target existing	P(exists) x value per find	the target is not there (P(exists) = 0)

First: the behavioural channel dominates. In every worked case here, and in the overwhelming majority of real ones, the behavioural gain is the largest term. **The practical advice that falls out of the postulate is therefore not to believe harder and not to expect healing. It is: find the action the belief produces, and check whether that action works.** That is a deflationary conclusion and it is the correct one.

Second: three of the four zero out on a fact about the world. No efferent pathway; ineffective action; absent target. Only the social channel's zero is about the group. **The conditional clause is not a philosophical scruple bolted onto the postulate. It is where three of the four gain formulas have their zero.**

34.7.8 7.8 The minimum viable gain

Take six distinguishable domains, a belief-driven surplus at 75 per cent of the ceiling, a retention loss of 10 per cent per period, a cost of 0.09 per period, and stored order at 1.5:

```
N = 6,  ln N = 1.791759,  S = 0.75 . ln N = 1.3438,  delta = 0.10,  C = 0.09,  Q_t = 1.5

G       h = G.lnN - delta.Q - C      Q_inf = (G.S - C)/delta     growing this step?
-------------------------------------------------------------------------------------
0.00           -0.2400                     -0.9000                    no
0.08           -0.0967                      0.1751                    no
0.15            0.0288                      1.1157                    no
0.25            0.2079                      2.4595                    YES

G_min for growth this step  =  (delta.Q + C) / S    =  0.1786
G_min for non-negative headroom = (delta.Q + C)/lnN =  0.1339
```

A channel gain of 0.15 is a real channel and it still loses ground, because the decay on existing stored order plus the cost of maintaining the belief exceed what the channel delivers. **The belief is grounded, the effect is real, and the system declines anyway.**

> **A grounded belief with insufficient gain is still a losing proposition. "It has a real mechanism" is necessary and it is not sufficient. The mechanism must clear `(delta Q + C) / S`.**

And the `G = 0` row deserves its own sentence. `Q_inf = -C/delta = -0.9000`, a negative equilibrium, which the storage law identifies as its own signal that the regime is unsustainable and outside the model's domain [Curzi, 2026a]. **A belief with no channel does not merely fail to build anything. Maintenance still bills. Holding an ungrounded belief is a net cost, continuously, and the arithmetic says so without any need for moralising about it.**

34.7.9 7.9 The fifth class, and why it is assigned gain zero

The four classes route through the believer's body. The fifth putative class does not: belief acting on the world directly, with no bodily, behavioural, social or informational route.

The postulate assigns it gain zero, and the reasoning is K1's: no channel has been specified. Not that the evidence is against it — the postulate does not need to adjudicate any particular experimental literature. The claim is narrower and firmer: **no one has produced a mechanism specification of the kind the four estimation procedures require. Name the carrier. Name the transducer. Name the energy budget. State the distance dependence and the noise floor it must beat.** Until a channel specification exists, the coupling relation's sum is empty and its prediction is zero, **and that prediction is the thing to test.**

The corresponding bridge is examined and rejected in Appendix B, audit B4. The paper's companion on the quantum formalism supplies the decisive quantitative objection to the version of this class that invokes physics: the corpus's conjugate-structure material is badged **analogy, not identity**, with an explicit refusal that quantum mechanics licenses anything about belief, mind, or future events [Curzi, 2026z], and the warm-wet decoherence estimate is `1.98e-17` s against a `1e-3` s neural event [Curzi, 2026aa].

A note on the word *channel*. It is used here in the sense the corpus uses it elsewhere: a route with an identity, whose separation from other routes is a structural property rather than a matter of care. The corpus's measured instance is in a different domain — in-band architectures admit an impersonation class that no care within the shared channel can close, while out-of-band separation closes it structurally, measured at 0.383 per cent in-band ambiguity against 0.000 per cent out-of-band over 100,000 values [Curzi, 2026]. **The parallel is structural and is claimed as nothing more: what transfers is the discipline of naming the route and treating its separation as a property of the architecture rather than of the operator's diligence.**

34.8 8 The measured channel

34.8.1 8.1 What this section is

This section asserts almost nothing of its own. It reports what has been measured, states the measurements' limits, and computes what they imply for the gain budget of Section 7. Where a number appears without a citation, it is arithmetic performed here on stated assumptions and is labelled as such.

An unusual constraint applies and is stated at the front. The corpus's external bibliography contains no entries for the primary placebo literature. The studies below are therefore named in the text with author, year and venue, exactly as the source manuscript states them, and **no bibliography key is attached to any of them**, because attaching an available key that does not support the sentence is worse than leaving it unattached. Appendix C lists them so a reader can locate them, and Section 13 records the gap as a real deficiency in the corpus's apparatus. **Furthermore, where the source manuscript declines to quote a figure from memory, this paper declines with it: the design and the direction of a result are given, and the reader is directed to the paper for its coefficients. An admitted gap is better than an invented decimal.**

34.8.2 8.2 The blockade result

```
Levine, J.D., Gordon, N.C. & Fields, H.L. (1978). "The mechanism of placebo analgesia."
The Lancet.

SETTING       postoperative dental pain, following third-molar extraction
MANIPULATION  administration of placebo; and administration of naloxone, a competitive
              opioid receptor antagonist
FINDING       placebo analgesia was reversed by naloxone
INFERENCE     the analgesia produced by expectation was mediated, at least in
              substantial part, by ENDOGENOUS OPIOID activity
```

A drug that blocks one receptor system removed an effect produced by a sentence. Nothing about the sentence changed and nothing about the patient's situation changed; what changed was the availability of one class of receptor, and the belief-produced effect went with it. **That is Section 7.3's series chain with one stage cut.**

```
STAGE                                                intact      naloxone
-------------------------------------------------------------------------
1  patient is told they are receiving an analgesic   ok           ok
2  patient forms an expectation of relief            ok           ok
3  descending opioidergic system engages             ok        BLOCKED
4  nociceptive signalling is modulated               ok           --
5  patient reports reduced pain                      ok           --

SERIES GAIN with stage 3 = 0:  the product is zero.
```

A channel that can be cut, with the effect disappearing when it is cut, is a channel that has been identified rather than hypothesised. K4 of Section 5.5 named blockade failure as a kill condition. **This is H4 run and passed, and it is the single result in this paper that most strongly discriminates the channel account from a model with no channel structure** (Section 3.5).

The honest limits, and they are real. The 1978 study was small, as studies of that era and setting were. Naloxone is not perfectly specific, and interpreting a reversal requires ruling out that naloxone is hyperalgesic on its own — a control the subsequent literature has attended to at length. **"Mediated by endogenous opioids" does not mean "entirely opioid":** later work reported that placebo analgesia has both opioid and non-opioid components depending on how the expectation was induced, with verbal suggestion engaging opioid mechanisms and conditioning with a non-opioid drug engaging a naloxone-insensitive route. Neuroimaging work reported placebo-associated changes in pain-related brain activity and, later, involvement of descending opioidergic control and of the spinal cord. **These extend the 1978 inference; they do not replace it.**

34.8.3 8.3 The disclosure result, and why it decides the framework's ethics

```
Kaptchuk, T.J. et al. (2010). "Placebos without deception: a randomized controlled
trial in irritable bowel syndrome." PLoS ONE.

DESIGN        randomised, two arms, three weeks
ARM 1         open-label placebo: pills presented explicitly and truthfully as inert,
              with an explanation that placebo effects can be substantial
ARM 2         no-treatment control, with the patient-provider interaction matched --
              which is the design feature that makes the comparison meaningful
FINDING       the open-label placebo arm improved significantly more than the
              no-treatment control on the trial's primary outcomes
n             80 patients
```

The design and the direction of the finding are stated; **the response percentages are deliberately not quoted from memory.** Subsequent work has extended the result to chronic low back pain and other conditions, and a meta-analysis of open-label placebo randomised trials reported a positive overall effect. **The honest summary is that the direction is consistent, the individual trials are mostly small, the outcomes are mostly subjective and self-reported, and the field is active.**

Why it is decisive for this framework specifically. Every protocol in this corpus operates under an absolute constraint: nothing masked, nothing below threshold, nothing concealed, because **a message hidden from its recipient cannot be consented to.** That is the same rule as the framework's prohibition on any intervention removing a person's ability to decline,

applied to information, and it stands in the same tradition as the argument that a mechanism's security must not depend on its concealment [Kerckhoffs, 1883].

Before 2010, that constraint and the physiological channel were in tension. The received view was that the effect required the patient to believe the pill was active, which requires telling them something false, which the constraint forbids. **The constraint would have won and the channel would have been closed to this framework.**

```
IF the effect requires deception:
    the channel is unavailable under the no-concealment constraint.
    G_physio = 0 FOR THIS FRAMEWORK, whatever it is elsewhere.

IF the effect survives disclosure:
    the channel is available under the constraint, at whatever gain it has.
    The person is told the truth, retains the ability to decline, and the effect
    is nonetheless measured.

THE EVIDENCE, AS IT STANDS, IS THE SECOND -- for subjective, symptom-based outcomes,
at modest magnitudes, in a literature that is small and active.
```

That is as far as it goes and it is stated to go exactly that far. It does not establish that any method in this corpus produces any effect. **It establishes that a truthful presentation is not automatically an inert one, which is a fact about the world that the framework's ethics needed and did not create.**

34.8.4 8.4 What "modest" means, in numbers

The word *modest* does heavy lifting in this literature and should be converted into arithmetic. Standardised effect `d`, converted to the probability that a randomly chosen treated person is ahead of a randomly chosen untreated one, and to a number needed to treat derived from that probability:

d	P(superiority)	NNT	distribution overlap	r
0.10	0.5282	17.74	0.9601	0.0499
0.15	0.5422	11.84	0.9402	0.0748
0.20	0.5562	8.89	0.9203	0.0995
0.25	0.5702	7.13	0.9005	0.1240
0.30	0.5840	5.95	0.8808	0.1483
0.40	0.6114	4.49	0.8415	0.1961
0.50	0.6382	3.62	0.8026	0.2425
0.80	0.7142	2.33	0.6892	0.3714
1.00	0.7602	1.92	0.6171	0.4472

At `d = 0.20` — a very common magnitude in this literature — **two people drawn at random, one treated and one not, and the treated one is ahead 55.6 per cent of the time against 50 per cent by coin flip. The distributions overlap by 92 per cent.** That is what a real, replicated, mechanistically grounded effect looks like from the inside, and it is much less than the word "works" usually conveys.

```
d = 0.15   NNT 11.84  ->  in 1000 people treated, about  84 additional responders
d = 0.20   NNT  8.89  ->  in 1000 people treated, about 112 additional responders
d = 0.30   NNT  5.95  ->  in 1000 people treated, about 168 additional responders
```

Small per person. Not small per thousand. Both statements are true and neither excuses the other. A clinician deciding for one patient and a health system deciding for a population are looking at the same number and correctly reaching different conclusions about it.

34.8.5 8.5 The three-arm requirement

This is the most important methodological point in the section and it separates "the placebo arm improved" from "there was a placebo effect." A patient in a placebo arm improves for at least three reasons that have nothing to do with belief: **natural history**, since most conditions fluctuate and many remit; **regression to the mean**, since patients enrol when symptoms are bad and a subsequent measurement drawn from the same noisy distribution is expected to be less extreme; and **measurement and reporting effects**, including repeated administration of a scale and the desire to please the investigator [Nisbett and Wilson, 1977, Schwarz, 1999, McCambridge et al., 2014].

Separating these from the belief-produced effect requires **a third arm receiving no treatment at all.** Computed, on illustrative symptom-scale changes where lower is better:

```
arm                    change
-----------------------------
no treatment            -1.2
placebo                 -3.0
active drug             -5.4

natural history + regression to the mean          =  -1.2
PLACEBO EFFECT PROPER   = placebo - no treatment  =  -1.8
DRUG EFFECT PROPER      = drug - placebo          =  -2.4

a TWO-ARM trial reports -2.4 and cannot see the -1.8 at all
fraction of the placebo arm's improvement that is NOT a placebo effect:  40.0%
```

Forty per cent of what looks like a placebo effect in this example is not one. That is the deflationary argument reduced to one line of arithmetic, and it is why analyses restricted to placebo-versus-no-treatment comparisons have found much less than the folklore assumed: they looked only at trials with the third arm, and the third arm is where the deflation lives. **The conclusion, stated as its authors stated it: little evidence that placebo interventions have important clinical effects in general, with a possible exception for continuous subjective outcomes, particularly pain. That exception is the whole of the physiological channel's remaining territory and it should be treated as such.**

34.8.6 8.6 The divergence diagnostic

The most useful design in this literature measures a subjective and an objective outcome **in the same patients, in the same arms, at the same time.** In asthma, patients received in a crossover design an active bronchodilator, a placebo inhaler, sham acupuncture, and no intervention, and two things were measured: what patients reported, and FEV1 read by a spirometer.

The finding, in its shape: subjective improvement was substantial and broadly similar across the active drug, the placebo inhaler and the sham acupuncture, and much greater than no intervention. FEV1 improved substantially only with the active bronchodilator. The structure is given; the percentages are not quoted from memory.

```
THE DIVERGENCE DIAGNOSTIC

    D  =  d_subjective  -  d_objective

measured in the SAME arm, in the SAME trial, at the SAME time.

arm                       d_subj    d_obj      D
-----------------------------------------------------
no treatment               0.00      0.00     0.00
```

```
open-label placebo          0.30      0.02      0.28
active drug                 0.55      0.52      0.03
```

A large positive D is the signature of a channel that reaches REPORT but does not reach the PHYSIOLOGY the instrument is reading. A D near zero is the signature of a channel that reaches both.

Three things must be said and all three matter. **First, this is a diagnostic on the channel, not a verdict on the patient.** The subjective improvement is not fake; the patient genuinely felt better and feeling better is a real state with real value. **Second, and this is why the finding is dangerous rather than merely interesting: in asthma the subjective measure is the one that fails.** A patient whose airways are closing may feel improved while the spirometer says otherwise, and acting on the feeling instead of the instrument is how people die of asthma. **The divergence is a clinical hazard and one of the strongest possible arguments against forgoing treatment on the strength of feeling better. Third, it is exactly what the channel analysis predicts.** The efferent pathways available to the physiological channel terminate on systems that modulate perception, affect and some autonomic function; **they do not terminate on airway smooth muscle in a way that produces bronchodilation, and the postulate's prediction for an outcome with no efferent pathway is zero. The prediction and the observation agree, which is worth exactly as much as one agreement is worth.**

34.8.7 8.7 Two routes, and which one this framework may use

```
EXPECTANCY     the person is TOLD something, forms a conscious expectation, and the
               expectation engages central mechanisms.
               Manipulated by: verbal instruction, framing, presentation, manner.
               Reported to engage opioid mechanisms in analgesia -- consistent with
               naloxone reversibility.

CONDITIONING   the person has REPEATEDLY received an active agent in a context, and the
               context alone comes to produce a physiological response.
               Manipulated by: prior pairing, not by instruction.
               Reported to produce responses that verbal instruction cannot readily
               override, and in analgesia responses that can be naloxone-insensitive
               when the conditioning agent was non-opioid.
```

Work in this area has reported that some conditioned hormonal responses are not modifiable by verbal expectation while pain responses in the same subjects are — that is, **the two mechanisms reach different effectors.** This is given as a reported finding whose structure is confidently stated and whose coefficients are not quoted from memory.

```
EXPECTANCY   is the route available to a no-concealment protocol, because it operates
             through what the person is TOLD, and telling them the truth is a
             permissible instruction. This is the route open-label placebo uses.

CONDITIONING requires prior pairing with an active agent, which outside a clinical
             setting means it is largely UNAVAILABLE.

CONSEQUENCE  the gain available under this framework's constraints is the expectancy
             gain, not the total. The 0.15-0.45 range is an upper bound and probably
             a generous one.
```

Every honest narrowing of what is available makes the remaining claim easier to check, which is the argument for making the narrowings.

34.8.8 8.8 Nocebo: the same channel with the sign reversed

Nothing in the postulate distinguishes positive from negative expectation. `E = G B` is indifferent to sign, and the evidence says the world is too. The clearest available demonstration is an n-of-1 randomised design in which individual patients who had stopped taking a drug because of side effects took, in randomised blinded periods, the drug, an identical placebo, and no tablet at all.

```
NOCEBO ATTRIBUTION RATIO

    ratio = ( S_placebo − S_no_tablet ) / ( S_drug − S_no_tablet )

worked on three illustrative symptom-score patterns:

S_drug   S_placebo   S_no_tablet     ratio
──────────────────────────────────────────────
 16.3      15.4          8.0        0.8916  (89.2%)
 20.0      18.0         10.0        0.8000  (80.0%)
 20.0      10.0         10.0        0.0000  ( 0.0%)
```

A ratio near 1 means almost the entire symptom burden is carried by the act of taking a tablet rather than by the tablet's pharmacology. A ratio near 0 means the burden is pharmacological. The formula distinguishes them; opinion does not. The trial reported a high ratio; the figure is not quoted from memory.

One. Nocebo is not a curiosity: it causes people to stop treatments that are working.

Two. Informed consent creates nocebo. Telling a patient the list of possible side effects reliably increases the rate at which they are reported. **This is a genuine and unresolved ethical tension in clinical communication and it is not this corpus's to resolve; it is noted because a treatment presenting only the beneficial direction of a channel would be presenting half a mechanism.**

Three, and this is the one that governs here. The no-concealment constraint has a cost, and this is where the cost is. Telling people the truth about what they are receiving raises nocebo as well as permitting open-label placebo. **This corpus accepts that cost without argument, because the constraint is not a cost-benefit judgement: a message hidden from its recipient cannot be consented to, and consent is not a term in an optimisation.**

Four. Symmetry is a testable prediction the postulate makes. `G` should be approximately the same magnitude in both directions for the same channel and outcome class. **If negative expectations reliably produced much larger effects than positive ones through the same measured pathway, the single-channel model would need amendment.**

34.8.9 8.9 What the disclosure result does not show

```
IT DOES NOT SHOW that the mechanism is understood. Why an admitted placebo works is an
open question. Candidate explanations include residual expectancy generated by the
explanation itself, conditioning from a lifetime of taking pills, the ritual of
treatment-taking, the therapeutic relationship, and reporting effects. THESE ARE NOT
RESOLVED, and anyone who says which one it is has gone beyond the evidence.

IT DOES NOT SHOW an effect on objective disease endpoints. The trials are overwhelmingly
on symptom scales in conditions defined by symptoms. Section 8.6 applies with full force.

IT DOES NOT SHOW a large effect. The trials are mostly small. Meta-analysis reports a
positive overall effect; heterogeneity is substantial and small-study effects cannot be
excluded in a literature of this size.
```

```
IT DOES NOT SHOW that any method in THIS CORPUS does anything. No method in this corpus
has been tested.

IT DOES NOT LICENSE FORGOING TREATMENT. Under any reading. In any degree.
```

34.8.10 8.10 The gain, and the over-budget check

```
G_physio  =  measured standardised effect / normalised expectancy

WORKING RANGE for subjective, symptom-based outcomes:   0.15 to 0.45
WORKING RANGE for objective disease endpoints:          approximately 0
WORKING RANGE available under the no-concealment
  constraint (expectancy route only):                   the low end of the first range

|E|  <=  G . B    with  G ~ 0.3  and  B <= 1  gives  |E|  <=  0.3

ANY CLAIM OF A STANDARDISED EFFECT ABOVE ABOUT 0.3, THROUGH THIS CHANNEL ALONE,
IS OVER BUDGET AND MUST NAME AN ADDITIONAL CHANNEL.
```

That is a computable objection and it takes about ten seconds to run. It is the single most useful thing this section hands to a reader who has to evaluate somebody else's claim.

34.8.11 8.11 What the postulate forbids here, as numbers

```
F1  A placebo arm whose OBJECTIVE endpoint moves as far as the drug arm, for an endpoint
    with no plausible efferent route from central expectation. Predicted ratio
    d_obj(placebo)/d_obj(drug) <= about 0.2 for such endpoints. Observed reliably above
    0.6 would put the channel analysis in serious trouble.

F2  An open-label arm indistinguishable from no-treatment when the patient-provider
    interaction is properly matched. This would remove the finding Section 8.3 rests on.

F3  A blockade agent leaving placebo analgesia untouched in a design that reproduces the
    1978 conditions with adequate power and appropriate controls. That would mean the
    opioid route is not the route, and the stage-3 identification is wrong.

F4  Nocebo attribution ratios near zero across well-powered n-of-1 designs -- symptom
    burden turning out to be almost entirely pharmacological after all.
```

None of these has occurred, and the postulate is therefore alive in this domain and not vindicated by it. Survival is not confirmation [Popper, 1959, Lakatos, 1970], and the rule applies to the author's own postulate with no discount.

34.9 9 Failure modes: a register

The failure modes below are the ones this apparatus generates. Each gives a description, the parameter it damages, a diagnostic signature, and the fix. **The corpus's companion register sorts belief failures into four classes — A, no channel; B, a real channel with insufficient gain; C, a wrong magnitude extrapolated from a real result; D, a system unresponsive at this input size — each with a stated diagnostic [Curzi, 2026ac].** The entries here are the mechanisms by which a claim lands in one of those four.

34.9.1 9.1 Failures of the conditional

The escape clause. *Description:* an unobservable sufficiency parameter absorbs every null — you did not believe hard enough, some part of you resisted. *Damages:* empirical content, reducing it to exactly 0.0000 bits (Section 5.2). *Signature:* every failure is explained by the same unmeasured quantity, which is only ever inferred from the failure it explains; ask what result would change the holder's mind and get silence. *Fix:* measure `B` before the outcome, put it in the prediction, forbid post-hoc appeal. **If `B` was low the prediction was low, and a low observed effect confirms the postulate rather than excusing it.**

The unnamed channel. *Description:* a belief is asserted grounded because it makes sense or is consistent with a framework, with no step-by-step pathway written down. *Damages:* `G`, which is unestimated and, being unestimated, usually assumed positive. *Signature:* ask for step 2 of the mechanism and receive step 1 restated in new vocabulary; the pathway description contains no verbs of physical or social action. *Fix:* write the steps, numbered, and **accept "I cannot write step 2" as the finding it is.**

The fitted gain. *Description:* `G` is estimated from the same outcome data it is used to predict. *Damages:* identifiability. **The exercise is void, not merely weak.** *Signature:* the predicted effect matches the observed effect suspiciously well and the gain was never published before the measurement. *Fix:* estimate `G` from the channel's own components, freeze it, publish it, then measure.

Magnitude blindness. *Description:* the sign of an effect is checked and its size is not. *Damages:* content, collapsing the conditional's 2 to 4 bits to the sign-only 1 bit. *Signature:* the write-up contains a p-value and no ratio of observed to predicted. *Fix:* report `observed / predicted` as a headline number and **treat a ratio far from 1 as damage to the model even when the sign was right** [Meehl, 1967, 1990].

The underpowered null. *Description:* a predicted effect is tested with a sample far too small to detect it, and the null is reported as evidence against the channel. *Damages:* nothing about the world; everything about the inference. *Signature:* no power calculation anywhere; the minimum detectable effect, computed afterwards, is several times the predicted effect. *Fix:* compute the required sample from the predicted magnitude before collecting anything, and where it is unavailable **report "not testable at this scale", which is a legitimate and useful result**.

Precedent promoted to evidence. *Description:* the observation that a tradition recorded something is cited as support for the claim's truth rather than for its independent discoverability. *Damages:* the audit method itself, and therefore every claim in the corpus including the well-supported ones. *Signature:* "confirms", "corroborates" or "validates" attached to a precedent source; two sources listed as though they were two measurements. *Fix:* state each column's role in the same block as the citation, every time. **Precedent establishes findability, never truth.**

The unfalsifiable goal. *Description:* the outcome is specified so that no observation could count against it — greater alignment, things opening up. *Damages:* `E`, which has no units, so the formula has no left-hand side. *Signature:* two observers cannot be made to agree on whether it occurred. *Fix:* units, procedure, noise level, before anything else.

The circular anchor. *Description:* belief magnitude is operationalised by the very behaviour the behavioural channel is supposed to explain, and the resulting correlation is reported as support for the channel. *Damages:* `G_beh`, which becomes tautological. *Signature:* `B` and the channel's first step are the same measurement under two names. *Fix:* anchor `B` on a decision in a different domain from the outcome, or use magnitude estimation, and state which.

The grounding substitution. *Description:* "grounded in truth" is read as sincerely held, true, or endorsed by authority — the three substitutions Section 4.3 forbids. *Damages:* the condition, and therefore everything. *Signature:* the grounding argument, written out, contains no reference to a pathway. *Fix:* **grounding is a property of the world.**

34.9.2 9.2 Failures of the gain budget

The empty box. *Description:* the diagram is drawn — belief, arrow, world — and the middle box is never filled in. *Damages:* `G`, unestimated and assumed positive. *Signature:* the mechanism paragraph contains no nouns naming a physical or social entity. *Fix:* classify, or estimate zero and stop.

The long chain. *Description:* a mechanism with six or eight plausible steps is presented as strong because every step is plausible. *Damages:* the magnitude estimate, catastrophically — **eight stages at 0.7 each give 0.0576.** *Signature:* "and then" appears more than three times and no stage carries a number. *Fix:* write the stages, put a ratio on each, multiply, and read the product aloud.

The unitless sum. *Description:* gains in different outcome units are added. *Damages:* the bound, which becomes meaningless while continuing to look like a number. *Signature:* a summed gain whose units are never stated. *Fix:* standardise every channel to its outcome's own standard deviation before summing.

The additive promise. *Description:* the parallel sum is reported as the expected effect rather than as a ceiling. *Damages:* calibration; predictions run systematically high. *Signature:* predicted effects equal to the summed bound, repeatedly, with no discount and no interference analysis. *Fix:* predict below the bound and state the discount's basis.

The adherence illusion. *Description:* a large behavioural `delta p` is measured, celebrated, and multiplied by an effect per action that was never independently measured. *Damages:* `G_beh`, which becomes the product of one real number and one imagined one. *Signature:* the write-up quotes an adherence improvement and no separate efficacy estimate. *Fix:* measure the action's effect in a design that varies the action without varying the belief; **if that design is impossible, report the gain as a range with a zero lower bound.**

The absent target. *Description:* an informational gain is computed from a search-time increase without establishing that what is being searched for exists. *Damages:* `G_info`, which is `P(exists)` times something and `P(exists)` was never estimated. *Signature:* a search rate and a time reallocation and no base rate; the successes are counted and the searches that found nothing were never recorded. *Fix:* estimate `P(exists)` first, and **record the null searches; they are the denominator.**

The survivorship anecdote. *Description:* the evidence base consists of people for whom the search succeeded; the people who searched equally hard and found nothing wrote nothing and are absent from the record. *Damages:* every gain estimate drawn from that base, upward, by an unknown factor. *Signature:* all the cases point the same way; there is a register of outcomes and no register of attempts. *Fix:* build the attempt register before the outcome register.

The sufficient-because-real fallacy. *Description:* a channel is shown to be real and the matter is treated as settled, with no check against the minimum viable gain. *Damages:* nothing about the world; everything about the decision. *Signature:* the report establishes a mechanism and never computes whether it clears the decay and the cost. *Fix:* compute `(delta Q + C)/S` and compare. **"Real but insufficient" is a legitimate verdict and is more common than either alternative.** It is class B of the corpus's four-class register [Curzi, 2026ac].

The coerced consensus. *Description:* `G_social` is raised by making dissent expensive, and the resulting coordination is read as shared belief. *Damages:* the group's non-parallelism and therefore its surplus, and the protocol's admissibility. *Signature:* committed fraction rises without any change in privately elicited belief; anonymous and public elicitation diverge sharply. *Fix:* anonymous-first elicitation, and a hard rule that any protocol removing the ability to decline is out of scope regardless of its measured gain. **The concentration instrument gives the measurable form of what coercion does to a group: naive shares of 40.9, 27.3, 18.2 and 13.6 give an effective count of 3.41 by `1/HHI`, and collapsing them into one bloc gives a bloc-HHI of 1.000000 and an effective count of 1.00** [Curzi, 2026]. **`1/HHI` is the Hill number of order two and is well grounded** [Hill, 1973, Jost, 2006]; the corpus's affine effective count `1 + (N−1)u` is a different quantity and is not a Hill number of any order.

The fifth-class smuggle. *Description:* a direct-action claim is dressed in the vocabulary of one of the four real classes — it works through the field, the body knows, the network responds — so that it inherits the credibility of a named channel without having one. *Damages:* the taxonomy's usefulness and the reader's ability to audit anything. *Signature:* the channel is named but its estimation procedure cannot be run; ask for the efferent pathway, the action's independent effect, the committed count, or `P(exists)`, and none is available. *Fix:* **if the estimation procedure for the claimed class cannot be executed, the class was misassigned and the estimate is zero.**

34.9.3 9.3 Failures specific to the measured channel

The missing third arm. A placebo arm's improvement reported as a placebo effect with no no-treatment comparator. *Damages* the effect estimate upward by natural history plus regression to the mean. *Signature:* two arms; a pre-post improvement reported as a finding. *Fix:* add the arm, or **report the quantity as "improvement in the placebo arm" and never as "the placebo effect"**.

The subjective-only trial. Every outcome self-reported and the result described as a physiological effect. *Damages* the inference: `D` cannot be computed, so a report-reaching channel and a physiology-reaching channel are indistinguishable. *Fix:* add one objective measure, even a crude one, and publish `D`.

The asthma error. Feeling better treated as being better, and treatment reduced accordingly. *Damages* the patient. **This is the failure mode with the worst consequences in the paper.** *Fix:* where `D` is large for a condition, the subjective measure is not a proxy for the objective one and must never be used as one.

The generalised pathway. "Placebo has physiological effects, therefore belief can affect any physiology." *Signature:* the pathway is named at the start of the argument and unnamed at the point of application — opioids at the start, "the body" at the end. *Fix:* name the efferent pathway to *this* tissue, or estimate zero.

The deception assumption. It is assumed the effect requires deception, so either the channel is written off as ethically unusable or deception is quietly justified. *Signature:* an argument presenting "tell the truth" and "get the effect" as mutually exclusive. *Fix:* the open-label literature exists; the assumption is out of date; **run the disclosed version and measure it rather than reasoning about it.**

The oversold meta-analysis. A pooled estimate from a small, heterogeneous literature of small trials quoted as a stable parameter. *Signature:* a pooled effect quoted without its confidence interval, its heterogeneity statistic, or the number and size of contributing trials. *Fix:* quote the interval and the heterogeneity or quote nothing.

The unmeasured expectancy. An effect reported without any measurement of what participants actually expected. *Damages* `G`, which is `d/B` with `B` never measured. *Signature:* the word "expectation" appears throughout the discussion and nowhere in the methods. *Fix:* measure expectancy by magnitude estimation, on the same scale, in every arm including the no-treatment arm.

The one-direction channel. The beneficial direction presented and the nocebo direction omitted. *Fix:* **one channel, two signs. Present both or present neither.**

34.10 10 Results: three audits worked end to end

All parameters in this section are stipulated by the worked cases; nothing here is a field observation and no dataset was collected. The purpose is to show that the procedure runs, reaches a disposition, and reports the disposition whichever way it falls. **In two of the three cases the protocol returns a result the participants did not want, which is the point.**

34.10.1 10.1 Two beliefs, one protocol

A delivery team of nine engineers holds two beliefs about how to ship earlier. Both are held sincerely, by roughly the same people, with roughly the same conviction.

```
BELIEF A   "Stating the deadline as enforced will move median delivery earlier."
BELIEF B   "Visualising the delivery as complete will move median delivery earlier."

E      median days from ticket-open to ticket-closed, from the tracking system
units  days (negative = earlier, which is the desired direction)
noise  sd = 2.4 days, estimated from the previous eighteen deliveries
```

Both beliefs are tested against the same outcome, which is the only reason the comparison means anything.

Step 1 — the channel, as steps. For belief A:

```
A1  the statement is made, in a meeting, on the record          CHECKABLE: minutes exist
A2  engineers revise their private ordering of work             CHECKABLE: board history
A3  revised ordering front-loads the blocking task              CHECKABLE: dependency graph
A4  the blocking task completes earlier                         CHECKABLE: timestamps
A5  delivery completes earlier                                  CHECKABLE: the outcome

AND THE PROPOSITION THE CHANNEL DEPENDS ON:  the deadline really is enforced --
there is a real downstream commitment, and missing it has real consequences.
THIS IS "GROUNDED IN TRUTH". If the deadline is not enforced, A2 does not occur
after the second or third repetition, and G collapses to zero.
```

For belief B, step B1 is "the visualisation is performed" and **step B2 cannot be written**: there is no proposed event between the visualisation and the ticket system.

```
BELIEF B: NO CHANNEL NAMEABLE TO THIS OUTCOME.  G_B = 0.000 days per unit belief.
PREDICTION: E_B = 0.000 days.
```

Note what has not been said. It has not been said that visualisation does nothing. Visualisation plausibly has channels to other outcomes — to the visualiser's anxiety, to their attention allocation, and through attention to behaviour, which is an informational channel in the taxonomy of Section 7.4. **What has been said is that no channel from this belief to this outcome can be written down. The postulate is evaluated per belief, per outcome, per channel, never in general.**

Step 2 — estimate `G` from the channel. Belief A's channel is behavioural, and its gain is computed from two independently measurable quantities, neither of which is the outcome:

```
p0   probability a given engineer reprioritises after an UNENFORCED deadline statement
     = 0.30   (from the board history of the previous four unenforced statements)
p1   probability of reprioritisation after an ENFORCED one
     = 0.75   (from the board history of the two enforced ones)
s    days saved per reprioritisation
     = 1.8    (from the dependency graph: mean slack recovered on the blocking path)

G_A = (p1 - p0) . s = 0.45 x 1.8 = 0.810 days per unit belief
```

None of the three inputs is the delivery time. That independence is what Gate 5 requires.

Step 3 — freeze and predict.

```
FROZEN:  G_A = 0.810 days/unit,  G_B = 0.000 days/unit.   Timestamped before measurement.

at B = 1.0:   E_A_pred = 0.810 days earlier    E_B_pred = 0.000 days
at B = 0.6:   E_A_pred = 0.486 days earlier    E_B_pred = 0.000 days

standardised:  d_A = 0.810 / 2.4 = 0.3375
```

Step 4 — power the measurement, and receive the first real result before any data exist.

```
sd = 2.4 days,  predicted effect = 0.810 days,  d = 0.3375
alpha = 0.05 two-sided

n per arm for 80% power  =  137.81  ->  138 deliveries
n per arm for 90% power  =  184.49  ->  185 deliveries

the team has 9 deliveries per arm.  se = 2.4 . sqrt(2/9) = 1.1314 days
minimum detectable effect at 80% power with n = 9:  3.170 days
detectable / predicted = 3.913

power actually achieved at various n:
    n =   9 per arm   se = 1.1314   power = 0.1105
    n =  20           se = 0.7589   power = 0.1872
    n =  50           se = 0.4800   power = 0.3928
    n = 100           se = 0.3394   power = 0.6651
    n = 138           se = 0.2889   power = 0.8005
    n = 200           se = 0.2400   power = 0.9215
```

With nine deliveries per arm the study has an 11 per cent chance of detecting an effect that is really there at the predicted size. It cannot test the claim. Had the team run it anyway and reported a null, they would have concluded that enforced deadlines do not work, on the strength of a design that could not have shown it if they did.

This is the most useful thing the protocol did, and it did it for free, from the magnitude estimate, before any data existed. The gain estimate sized the study. Without a magnitude prediction there is nothing to size a study against, which is why sign-only claims and underpowered studies travel together. The team runs the observation as a declared calibration exercise rather than a test and pools with two sibling teams to reach `n = 27` per arm over two quarters — still underpowered at 0.24, and stated as such.

Step 5 — measure against the frozen predictions. With belief magnitudes elicited by magnitude estimation at `B_A = 0.9` and `B_B = 0.9`:

```
predicted E_A = 0.810 x 0.9 = 0.729 days earlier
observed  E_A = 0.720 days earlier
observed / predicted = 0.988
|obs - pred| = 0.009 days = 0.008 se        WELL inside tolerance

predicted E_B = 0.000 days
observed  E_B = 0.150 days earlier
0.150 / se = 0.133 se                        inside +/- 2 se, consistent with G_B = 0
```

Step 6 — the shuffled-input control. The belief-magnitude series is permuted across the pooled deliveries and the frozen model re-scored. **The shuffle degrades the fit — but at `n = 27` and power 0.24 the shuffle test is itself underpowered, and that is recorded as a limitation rather than passed over.**

Step 7 — read the result honestly.

```
ESTABLISHED (as survival, not as truth):
  - A's channel steps A1–A4 verified in the board history and dependency graph,
    separately from the outcome
  - the frozen, mechanism-derived G_A = 0.810 predicted the observed 0.720 to 0.008 se
  - B's zero prediction survived: observed 0.150 days, inside noise

NOT ESTABLISHED, and this is the important line:
  - the study was UNDERPOWERED BY DESIGN (power 0.24), so A's close agreement is weak
    evidence and partly luck. Reporting it as a triumph would be exactly the magnitude
    blindness named in Section 9.
  - the shuffled-input control was run but is itself underpowered
  - nothing was established about visualisation in general -- only about visualisation
    as a channel to THIS outcome

THE DISCRIMINATION THAT DID OCCUR:
  Two sincerely held beliefs, equal in conviction, received DIFFERENT predictions from
  the same postulate before any data existed, on the basis of whether a channel could be
  written down: 0.729 days and exactly zero. Both predictions survived.

WHAT WOULD HAVE ENDED THE POSTULATE HERE:
  B moving the outcome by a clear margin on a properly powered design, with the shuffle
  failing to remove it and no behavioural or informational channel recoverable. That is
  K1 of Section 5.5. It did not happen; had it, Section 4 would need rewriting.
```

It does not prove the postulate. Nine engineers and twenty-seven deliveries prove nothing. **What it demonstrates is that the postulate is executable** — that a practitioner can derive two different quantitative predictions from two beliefs without touching the outcome data, find that one belief has no writable mechanism, find that their sample cannot test the other, and end up knowing exactly what they would need in order to find out. **Strike the conditional and every one of those steps disappears: no channel to name, so no gain to estimate, so no magnitude to predict, so no sample size to compute, so no test to run, so no result to report.**

34.10.2 10.2 One belief, four channels, and an audit that changes the plan

A software organisation of 240 people holds a belief: *the weekly cross-team review makes this group better at shipping.* A sceptical executive proposes cancelling the review to recover the time.

Step 1 — classify.

```
physiological   plausible route from belief to anyone's neurochemistry that reaches
                SHIPPING TIME?  No.  G_physio = 0.00.  Recorded, not omitted.
behavioural     belief -> attendance -> rework avoided.  YES.
social          belief held in common -> cross-team coordination.  YES.
informational   belief -> attention to interface defects -> defects found.  YES.
```

The zero is recorded rather than omitted, because a budget that lists only the positive channels has hidden what was considered.

Step 2 — behavioural gain, from behaviour records.

```
attendance without the belief  a0 = 0.60     (from calendar acceptance logs, Q1)
attendance with the belief     a1 = 0.90     (Q3, after the belief spread)
delta a = 0.30

effect per attended review = 0.60 units of rework avoided
```

```
    measured independently: teams that attended in the pilot window versus matched
    teams that did not, with the belief held equally in both -- the design that
    varies the ACTION without varying the BELIEF

G_beh = 0.30 x 0.60 = 0.1800
```

Both factors are published. **A reader can see that this gain rests on a 0.30 attendance change and a 0.60 per-review effect, and can challenge either.**

Step 3 — informational gain, with the base rate first.

```
I1  P(a real interface defect exists in a given week)  = 0.85, from the last two years
    of post-release defect classification. NOT assumed to be 1.
I2  detection rate r = 0.04 per hour of directed cross-team attention, from historical
    hours-per-find in the same defect class
I3  belief reallocates attention from t0 = 2 h to t1 = 8 h per week, from time records

P(find) = 1 - exp(-0.04 t):     2 h -> 0.0769      8 h -> 0.2739
delta P(find) = 0.1970

value per find = 0.4 units
G_info = 0.85 x 0.1970 x 0.4 = 0.0670
```

The 0.85 matters. Without it the gain would have been quoted as 0.0788 — an 18 per cent overstatement arising entirely from assuming the target always exists.

Step 4 — social gain, with the threshold and the transfer assumption.

```
committed (acting at a cost: attending when it conflicts with their own deadline) = 34
group N = 240        committed fraction = 34/240 = 14.2%

The measured tipping region near 25% comes from controlled online experiments.
TRANSFER TO THIS ORGANISATION IS A MODELLING ASSUMPTION, NOT A MEASUREMENT.

25% of 240 = 60 people. The organisation has 34. IT IS BELOW THRESHOLD.
Estimated G_social = 0.11 and DECAYING, not 0.11 and stable.
```

And the identification caveat applies here and only here: this number is a modelling position, not an identified empirical result [Manski, 1993]. It is used because a budget needs a figure, and it is flagged in every place it is used.

Step 5 — the budget.

```
channel          G_i      basis
-------------------------------------------------------------------
physiological    0.00     no efferent route to shipping time
behavioural      0.18     attendance 0.60 -> 0.90; 0.60 per review
social           0.11     below threshold; decaying; NOT IDENTIFIED
informational    0.07     P(exists) 0.85; 2 h -> 8 h at r = 0.04
-------------------------------------------------------------------
SUM              0.36

bound on |E|:  B = 0.4 -> 0.144    B = 0.7 -> 0.252    B = 1.0 -> 0.360
```

Belief is elicited by magnitude estimation at `B = 0.7`. **Predicted effect: at most 0.252 standardised units; and because the sum is a ceiling, the point prediction is set below it at 0.20 with a stated discount for channel interference — the eight hours the informational channel spends are hours the behavioural channel's attendees also spend.**

Step 6 — the minimum viable gain.

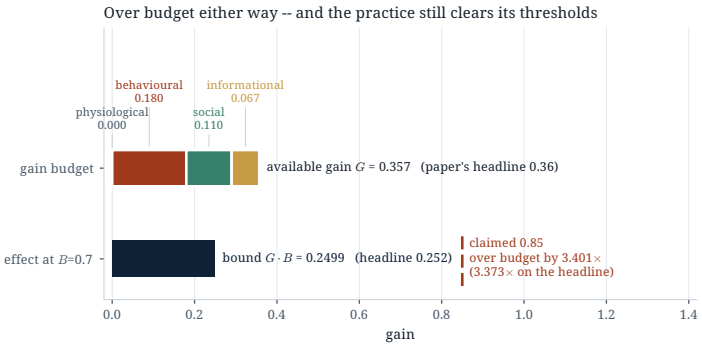


Figure 34.5: **The claim is over budget by 3.4 and the practice is worth keeping — both at once.** Four channel gains estimated from components that are not the outcome sum to 0.36 (physiological 0.00, behavioural 0.18, social 0.11 and decaying, informational 0.067), bounding the effect at 0.252 at `B = 0.7` against a claimed 0.85: a required-over-available ratio of 3.373. The summed gain clears both the growth threshold `(delta Q + C)/S = 0.1786` and the headroom threshold `(delta Q + C)/ln N = 0.1339`, but the behavioural channel alone at 0.18 barely clears the first. Modelled from the budget of Section 10.2; generated by `figs/src/CCP-034_fig05.py`.

```
N = 6 domains,  ln N = 1.791759,  S = 0.75 . ln N = 1.3438
delta = 0.10,  C = 0.09,  Q_t = 1.5

G_min for growth this step       = (delta.Q + C)/S    = 0.1786
G_min for non-negative headroom  = (delta.Q + C)/lnN  = 0.1339

ACTUAL SUMMED G = 0.36.   Above both.   Headroom at G = 0.36:
    h = 0.36 x 1.791759 - 0.10 x 1.5 - 0.09 = 0.4050 > 0
Steady state:  Q_inf = (0.36 x 1.3438 - 0.09)/0.10 = 3.937
```

The review clears the bar by a factor of about two, and the executive's proposal to cancel it is a proposal to remove the only channel above the minimum viable gain. But note what the analysis also shows: **the behavioural channel alone, at 0.18, barely clears the 0.1786 growth threshold.** The review's viability is not comfortable; it depends on the social and informational channels, and the social channel was measured as decaying.

Step 7 — the audit against the claim. The review's sponsors have been claiming a 30 per cent reduction in cycle time, which in standardised units is 0.85.

```
required delta          0.85
available summed gain   0.36 x 0.7 = 0.252

required / available = 0.85 / 0.252 = 3.373

VERDICT: NOT COVERED. The claim is over budget by a factor of 3.4.
```

The claim is unfunded and the practice is worth keeping. Both are true at once, and no discussion that lacked the arithmetic could have arrived at both. The sponsors revise the claim to 0.20; the revised claim survives contact with the data in the following quarter at 0.17 observed against 0.20 predicted, ratio 0.85, **reported as a magnitude result rather than as a success.**

Step 8 — the intervention the budget selected. The binding constraint is not the review's existence; it is the social channel below threshold and decaying.

```
committed now                34    (14.2%)
```

```
threshold, on the stated transfer assumption   60   (25%)
shortfall                    26 people

AND THE TOPOLOGY RULE: complex contagion needs reinforcing contacts, so the 26 are
recruited CLUSTERED -- concentrated in four adjacent teams whose members already
interact -- rather than scattered one per team across twenty-six teams. Same
headcount, different placement, and the clustered placement is the one the contagion
literature supports.
```

Expected benefit: `G_social` moves from decaying to self-sustaining, taking the summed gain from 0.36 to roughly 0.45 and the headroom from 0.4050 to 0.5663. **The organisation began with a belief, an unfalsifiable claim about it, and a proposal to cancel the practice on grounds of cost. It ended with four gain estimates from four independent data sources, a bound, an over-budget verdict on the claim, a keep verdict on the practice, and a specific, sized, placed intervention on the one channel that was failing.**

```
NOT ESTABLISHED:
  - that 25% is the threshold in THIS organisation. It is a transfer assumption from a
    controlled online experiment and is labelled as one in every place it is used, and
    the underlying parameter is not identified from observation of this group.
  - that the channels are independent. They are not; the discount from 0.252 to 0.20
    is an estimate of interference, not a measurement of it.
  - anything at all about physiology. G_physio = 0 was recorded, not tested.
```

34.10.3 10.3 Designing an honest study, and the redesign the arithmetic forces

A research group wants to test whether an open-label placebo, presented with full disclosure, reduces self-reported fatigue in a defined population.

```
BELIEF      "taking this pill, which I have been told is inert, will reduce my fatigue"
B measured  magnitude estimation against a standard, at enrolment and weekly
CHANNEL     physiological, expectancy route (NOT conditioning: no prior pairing is
            arranged, and none is permitted under the disclosure constraint)
STAGES      1  the disclosure is delivered, in a scripted, identical form
            2  the patient forms an expectation                    B is measured here
            3  central mechanisms engage                          NOT measured; the
                                                                    honest gap
            4  fatigue perception is modulated
            5  the patient reports on the scale
OUTCOME     a validated fatigue scale, primary, prespecified
PLUS        one objective measure -- actigraphy-derived daytime activity -- so that D
            can be computed

ARM A   open-label placebo, scripted disclosure, weekly contact
ARM B   no treatment, SAME weekly contact, SAME scripted duration, no pill
```

Stage 3 is unmeasured and is recorded as unmeasured. A chain with an unmeasured stage is a chain whose gain is inferred rather than traced. Arm B is the arm that makes the study worth running, and it is the arm most often cut for cost.

```
G_physio, expectancy route, subjective outcome, taken at the LOW end of the range
because the constraint set restricts this study to expectancy:

    G = 0.25   FROZEN AND PUBLISHED BEFORE ENROLMENT
```

```
expected B in arm A, from a pilot of 12: B = 0.55

E_pred (subjective) = 0.25 x 0.55 = 0.1375 standardised
E_pred (objective, actigraphy) = 0.00 to 0.05
D_pred = 0.1375 − 0.025 = 0.1125
tolerance on the subjective prediction: +/− 0.10, which is 3.32 bits of content
```

And here the arithmetic delivers the bad news.

```
n per arm = 2 (z_{alpha/2} + z_beta)^2 / d^2,  alpha = 0.05 two-sided

 power    z_beta    d=0.20   d=0.30   d=0.40   d=0.50
------------------------------------------------------
  0.80    0.8416       393      175       99       63
  0.90    1.2816       526      234      132       85
  0.95    1.6449       650      289      163      104

the prediction is d = 0.1375. At 80% power that requires:
    n per arm = 2 (1.9600 + 0.8416)^2 / 0.1375^2 = 830 per arm, 1660 in total

the group has funding for 60 per arm. At n = 60 the minimum detectable effect at
80% power is (1.9600 + 0.8416) x sqrt(2/60) = 0.5115 standardised.

detectable / predicted = 0.5115 / 0.1375 = 3.72x

THE STUDY AS FUNDED CANNOT TEST THE PREDICTION.
```

This is the same result the first worked case produced, from a completely different starting point, and it is why both compute power before anything else. A trial powered for `d = 0.5` in a literature that finds `d = 0.2` is not a test. It is a coin flip with a protocol attached, and it will produce a null that gets cited as evidence of absence.

Four options, all requiring a decision made in the open: pool to reach 830 per arm; **reduce the noise rather than the sample, using a within-person crossover that removes between-person variance from the denominator**; run it as a prespecified calibration study; or run it as funded and report a null, which is not legitimate. The group takes the second with the third folded in.

```
between-person sd on the fatigue scale (from the pilot)   =  1.00 (standardised)
within-person sd, same scale, same interval               =  0.42

the crossover's effective d = 0.1375 / 0.42 = 0.3274 in within-person units

n for 80% power at d = 0.3274, paired design:
    n = (1.9600 + 0.8416)^2 / 0.3274^2 = 73 people

73 people, each serving as their own control, against 1660 in the parallel design.
```

Taking the magnitude estimate seriously turned an infeasible 1,660-person trial into a feasible 73-person one, and the only thing that changed was the design. That is the practical payoff of insisting that a postulate predict a number: a design cannot be chosen efficiently against a claim that has no size.

```
PRESPECIFIED KILL CONDITIONS
```

```
KILLS THE STUDY'S CLAIM:
  - subjective effect outside 0.1375 +/- 0.10 in either direction
  - failure of the shuffled-input control
  - the contact-matched no-treatment arm showing equal improvement (this is F2)

DOES NOT KILL IT BUT MUST BE REPORTED:
  - an objective (actigraphy) effect materially above 0.05, which would mean a channel
    was missed and the budget was wrong -- an INTERESTING failure, and the one worth
    having
  - a nocebo signal: adverse symptom reports higher in arm A than arm B, which the
    single-channel model predicts should occur and be of comparable magnitude

COULD NOT BE CONCLUDED, UNDER ANY RESULT:
  - that the mechanism is understood (stage 3 is unmeasured)
  - that this generalises to any other condition, population, or outcome class
  - that any method in this corpus does anything
  - that anyone should reduce or forgo any treatment. Under any result. In any degree.
```

34.10.4 10.4 The audit gate in general form

```
 1  State the belief as a proposition that could be true or false.
 2  Name the outcome, its units, its measurement procedure and its noise level.
 3  Name the channel as numbered steps, each independently checkable.
    If step 2 cannot be written, record G = 0 and stop. That is a finding.
 4  Classify the channel: physiological, behavioural, social, informational, or none.
 5  Estimate each gain from that channel's own components, none of which is the outcome.
 6  Sum, after standardising to common units. Treat the sum as a ceiling.
 7  Compute the minimum viable gain (delta.Q + C)/S and compare.
 8  Compute required / available for the claim being made. Publish the ratio.
 9  Freeze and publish G. Predict E with a tolerance.
10  Power the measurement against the prediction, before collecting anything.
11  Run the shuffled-input control.
12  Report the magnitude result separately from the sign result, and state in advance
    what result would end the claim.
```

The gate is a canonicalisation step before a comparison step, and the corpus's measured instance of that discipline in a different domain is the canonical form `canon(x) = restrict_ASCII(confusable_fold(NFKC(x)))`, which collapses 495 of 495 manipulated entries onto the clean form, is idempotent across all 500 cases, and resolves 995 raw byte strings to the true 500 identifiers [Curzi, 2026]. **What transfers is the discipline: put every submission into one normal form, and the comparison that follows is between like objects rather than between presentations.**

Applied to the three cases above, the gate returned three different dispositions and returned each for a different reason: belief B failed at step 3, having no second step; the review's claim passed steps 3 to 7 and failed at step 8; and the fatigue study passed steps 3 to 9 and failed at step 10, which forced a redesign rather than a rejection.

34.11 11 Discussion

34.11.1 11.1 What follows

The postulate's operational content is a budget, and a budget is auditable. Given gains estimated from mechanism and frozen before measurement, the bound `|E| <= (sum_i G_i) B_max`

is arithmetic, and the required-over-available ratio converts a rhetorical dispute into a number. In the worked case that number was 3.373 and the claim failed before any data were collected, while the practice generating the claim survived on the same arithmetic. **No discussion lacking the arithmetic could have reached both conclusions.**

The deflationary consequence is the important one. In every worked case the behavioural channel is the largest term, and three of the four gain formulas have their zero on a fact about the world rather than about the believer. **The practical advice the postulate yields is therefore not to believe harder. It is: find the action the belief produces, and check whether that action works.**

The threshold result is the least intuitive and the most useful. A channel gain of 0.15 is a real channel and it still loses ground, because decay on existing stored order plus the cost of maintaining the belief exceed what the channel delivers. **"It has a real mechanism" is necessary and not sufficient; the mechanism must clear `(delta Q + C)/S`.** The same threshold appears elsewhere in the corpus in a different register, as the equivalence between a positive storage steady state and a sustaining frequency: `Q_inf > 0` if and only if `eta F(u*) > c*` if and only if `f > f_min`, verified at twelve parameter sets [Curzi, 2026]. And where a belief must accumulate enough to cross a barrier rather than merely to grow, the corpus's launch-budget result gives the form: a fund with fixed point `Λ_inf = L/lambda` makes departure feasible only where `Λ_inf > E_b`, after `t* = ln(1 − E_b/Λ_inf)/ln(1−lambda)` periods [Curzi, 2026p].

And repetition cannot substitute for gain. The corpus's contraction result is exact on this point: refinement cannot move the fixed point, because `Q_inf = (eta F − c)/delta` contains no `t` [Curzi, 2026]. **Holding a belief longer does not raise `G`. Only changing the arrangement does.**

34.11.2 11.2 What does not follow

Nothing here shows that belief creates reality. It shows what a claim of that form would have to look like in order to be checkable, and it computes the price of the condition that makes it checkable. **The postulate remains a postulate and this paper does not upgrade it.**

Nothing here licenses a clinical, medical or therapeutic claim for any method in this corpus. No method here has been tested; no controlled observation of any protocol in this corpus exists; and the divergence diagnostic of Section 8.6 is the reason that prohibition is a safety rule rather than a legal formality. **Nothing in this paper licenses forgoing treatment, under any reading, in any degree.**

Nothing here identifies the social gain. Section 3.4 states the reason and Section 13 repeats it. The figure of 0.11 in Section 10.2 is a modelling position with a stated transfer assumption on top of an unidentified parameter, and it is the weakest number in the paper.

And nothing here establishes that the corpus's own central technical claims are grounded. Applied to the corpus's own biological material, the postulate returns *not yet grounded*: there is no established mechanism by which a cell produces a response depending on a symbol string as an order, as distinct from permutation-invariant functionals of a waveform [Curzi, 2026,aa]. **A postulate that returns "not yet grounded" about its author's own project is a postulate being used correctly, and that is the strongest thing this paper can say for itself.**

34.11.3 11.3 Relation to the surrounding literature

On reflexivity, the corpus is closer to the sociology of finance than to psychology. Soros [2013] names the doctrine; MacKenzie and Millo [2003] supplies the observed instance with dates and price data; Callon [1998] and MacKenzie [2006] give the frame. **What this paper adds to that literature is a loop gain and a sharp stability condition; what it does not add, and cannot, is identification.**

On effect size, the literature is a constraint rather than a support. Jussim and Harber [2005] is the authoritative review and the finding is that self-fulfilling prophecies are real but small, do not typically accumulate, and often dissipate. **The gains in Section 7 are small numbers**

because that is what the literature supports, and any downstream use of this framework that produces a large gain has almost certainly made one of the errors in Section 9. The wider bound is that social predictability is low even with excellent data [Salganik et al., 2020, Hofman et al., 2017].

On measurement, the corpus is exposed and says so. Michell [1997] is the strongest objection to the framework as a whole, and the conjoint-measurement test that would settle it [Luce and Tukey, 1964, Krantz et al., 1971] has not been run for `B`, for `E`, or for any construct in this corpus. The magnitude-estimation route [Stevens, 1957] is the best available and it is not the same thing as having passed the test.

On what this framework is for, the honest answer is not prediction alone. Epstein [2008] sets out sixteen reasons to build a model, of which forecasting is one, and **this paper claims explanation, discipline and the sizing of studies rather than forecasting skill.** Sobkowicz [2020] and Durlauf [2005] are right to ask what observation would distinguish this account from a neighbouring one, and Section 3.5 gives the two that would and names the four that would not.

On the framework's own reflexivity, the objection is live. Lucas [1976], Hacking [1995] and Ferraro et al. [2005] together say that a framework which tells people how their beliefs work will change how their beliefs work. **This paper does not claim its gains are structural, and Section 13 states the restriction on where they may be used.** Ferraro et al. [2005]'s warning is taken seriously in one specific respect: **a framework that tells people most of their beliefs have gain zero could make that true by discouraging the search that finds the channel.** That is a real hazard of publishing this result and it is stated here rather than discovered later.

34.12 12 Verification

The results of this paper compound from six postulates. Each is atomic, each is checkable on its own, and each carries its status. **The two that carry the paper's central claim are `POSTULATED`, and the chain says so.**

34.12.1 12.1 Dimensional audit

Symbol	Dimension	Unit	Status
`E`	`[X]`	the outcome's own units (days, mmHg, defects, scale points, currency)	**measured**, in the declared outcome unit
`B`	`[1]`	none	**measured**, dimensionless on `[0,1]`; **scale type is the open question** (Section 6, Gate 2)
`G`	`[X]` per unit belief	outcome units per unit `B`	**derived** from channel components; **not comparable across outcomes unless standardised**
`eps`	`[X]`	the outcome's own units	**measured** as a noise standard deviation
`lam`	`[1]`	per update step	**assumed** dimensionless in `(0,1]`

Symbol	Dimension	Unit	Status
`alpha`	belief per unit `[X]`	inverse of `G`'s units	**assumed**; the feedback strength
`L = alpha G`	`[1]`	none	**derived** dimensionless: `(belief/[X]) x ([X]/belief) = 1`
`d`	`[1]`	standard deviations	**derived**, `E` divided by the outcome's own sd
`G_i` (series stage)	`[1]`	none	**defined** as a transmission ratio
`Q`, `delta`, `C`, `S`	as in the storage law	proxy units; per period	**as declared there**; carried here only for the threshold arithmetic

Does the equation balance? `E = G B + eps` balances if and only if `G B` carries `[X]`, and since `B` is dimensionless **`G` must inherit `E`'s units exactly.** That is why a summed gain over channels with different outcomes is malformed unless every channel is standardised to its own outcome's standard deviation first, and why Section 7.3's parallel sum is stated with that condition attached. `L = alpha G` is dimensionless because `alpha` is defined as belief produced per unit of effect, which is the reciprocal of `G`'s units; **the loop gain is therefore a pure number and a comparison of it against 1 is meaningful.**

Where the framework has no operational unit, it is said here. `B` is a normalised strength of holding and **has no operational unit at all**; it is treated as a dimensionless magnitude by stipulation, not by measurement, and Section 13 records that the quantitative structure required for that stipulation has not been tested. `E` has units only because the analyst is required to declare them; the postulate supplies none.

34.12.2 12.2 The postulates

P1 — Belief has a magnitude on `[0,1]`. *Status:* **DEFINITION** — the paper's own construction, carrying no empirical weight on its own, and Gate 2 records that the scale type licensing arithmetic on it has not been established.

P2 — A belief affects a world-state only through one or more channels, and with no channel the effect is zero at every belief magnitude. *Status:* **POSTULATED** — asserted by this work, with no external support as stated, and the corpus's own coverage audit lists it among the claims for which no supporting literature exists. **This is the paper's central claim and it is the only postulate in the chain that does the load-bearing work.**

P3 — The gain of a multi-stage channel is the product of its stage transmissions, and the gain of independent parallel channels is at most their sum. *Status:* **ARITHMETIC** for the product, from the definition of a transmission ratio; **POSTULATED** for the additivity of parallel channels, which assumes non-interference and is therefore stated as a **bound, not a prediction**.

P4 — `-log2(p)` measures the empirical content of a claim permitting a fraction `p` of a bounded outcome space. *Status:* **CITED** — the standard information measure [Shannon, 1948, Cover and Thomas, 2006], applied to a permitted band rather than to a message.

P5 — A first-order linear recursion `x_{t+1} = a x_t` has `x = 0` stable if and only if `|a| < 1`. *Status:* **CITED** — elementary, and stated with its convergence condition in Elaydi [2005]; the control-theoretic reading is Åström and Murray [2008].

P6 — The storage law's one-step growth condition is `eta S_t > delta Q_t + C_t`. *Status:* **CITED** — a corollary of the storage law's contraction structure [Curzi, 2026a].

34.12.3 12.3 The compounding chain

P1 + P2	-> R1	E = G.B + eps, with G = 0 => E = 0 for all B. This is the postulate in magnitude form. It is not derived from anything above it; P1 supplies the domain of B and P2 supplies the forbidding clause.	POSTULATED
R1	-> L1	halving G and doubling B leaves every observation unchanged, so the pair (G,B) is not identifiable from outcome data alone.	ARITHMETIC
L1	-> R2	the formula is testable only if G is estimated from the channel's own components and frozen before the outcome is measured.	ARITHMETIC
P4	-> R3	a claim permitting the whole of a bounded space has −log2(1) = 0.0000 bits; a sign-only claim over half has 1.0000 bits; a +/− 0.10 band in a width-2 space permits 0.1 and has −log2(0.1) = 3.3219 bits.	ARITHMETIC
R3	-> R4	each halving of the tolerance buys exactly one bit, because −log2(p/2) = −log2(p) + 1.	ARITHMETIC
R1 + base rates	-> R5	at 40 channelled and 60 unchannelled arms in 100, the unconditional form is refuted by 0.3250 of observations and the conditional by 0.0550; the observations needed for a 95% chance of at least one refutation are ln(0.05)/ln(1−p): 8 and 53.	ARITHMETIC
P5 + R1	-> R6	B_{t+1} = (1−lam)B_t + lam.min(1, L.B_t) linearises at B = 0 to a pole 1 − lam + lam.L, so B* = 0 is stable iff L < 1; and L = alpha.G, so G = 0 forces L = 0 at every alpha.	ARITHMETIC
P3	-> R7	0.70 x 0.50 x 0.60 x 0.40 = 0.0840, so four stages whose worst transmission is 0.40 deliver 8.4% of a lossless chain; and any zero stage zeroes the product.	ARITHMETIC
P6 + R1	-> R8	with eta -> G, growth requires G.S > delta.Q + C, so G_min = (delta.Q + C)/S = (0.10 x 1.5 + 0.09)/1.3438 = 0.1786; and non-negative headroom requires G_min = 0.24/1.791759 = 0.1339.	ARITHMETIC
R8 at G = 0	-> R9	Q_inf = (0 − C)/delta = −0.09/0.10 = −0.9000, a negative equilibrium: an ungrounded belief is a continuing net cost, not merely an absent gain.	ARITHMETIC
R1 + budget	-> R10	summed gain 0.36 bounds |E| at 0.36 x 0.7 = 0.252, against a claimed 0.85: required/available = 0.85/0.252 = 3.373.	ARITHMETIC

Every line except R1 is ARITHMETIC or CITED, and R1 is POSTULATED. That is the honest shape of this chain and it must be read carefully: **the arithmetic is certain and the postulate is not, and everything downstream of R1 inherits R1's status when it is applied to the world.** R3, R4, R7 and R8 are true as arithmetic regardless of whether the postulate holds; R5, R6, R9

and R10 are true as arithmetic *given* the postulate's form, and say nothing about whether that form describes anything. **A chain whose load-bearing step is a postulate is not a proof, and this section says so rather than leaving the reader to notice.**

34.12.4 12.4 The verification instance a reader can redo in a minute

Three checks, each independent of the others.

Check one — the bit arithmetic. A point prediction with tolerance `+/- 0.10` in a standardised band of width 2 permits `0.20/2.00 = 0.10` of the space. `-log2(0.10) = 3.3219`. Halve the tolerance to `+/- 0.05` and the permitted fraction halves to 0.05, giving `-log2(0.05) = 4.3219`. **The difference is exactly 1.0000 bits, and it is exactly 1 bit for every halving, because `log2(2) = 1`. ARITHMETIC.**

Check two — the refutation counts. With a per-observation refutation probability `p`, the number of observations for a 95 per cent chance of at least one refutation is `ln(0.05)/ln(1-p)`. At `p = 0.3250`: `(-2.9957)/ln(0.6750) = (-2.9957)/(-0.3930) = 7.62`, so **8 observations**. At `p = 0.0550`: `(-2.9957)/ln(0.9450) = (-2.9957)/(-0.0566) = 52.9`, so **53 observations**. **ARITHMETIC**, and the six-fold difference is the price of the conditional.

Check three — the budget and the audit. `0.00 + 0.18 + 0.11 + 0.067 = 0.357`, which the paper rounds to 0.36; times `B = 0.7` gives 0.252 (0.2499 unrounded). Against a claimed 0.85, `0.85/0.252 = 3.373`. And the minimum viable gain: `(0.10 x 1.5 + 0.09)/(0.75 x ln 6) = 0.24/1.3438 = 0.1786`, with `ln 6 = 1.791759`. **ARITHMETIC. A small rounding inconsistency in the source is recorded in Section 13 rather than corrected silently.**

34.12.5 12.5 Limits, asymptotics, and the domain of each relation

`B -> 0`. `E -> 0`. Trivially observed and carries no information.

`B -> 1`. `E -> G`. **The most a belief can do is the channel's full gain**, which is the postulate's central practical content and the source of the over-budget check.

`G -> 0`. `E -> 0` for all `B`, and `L = alpha G -> 0`, so the belief's fixed point is `B* = 0`. **The logical and the dynamical statements coincide exactly in this limit**, which is the strongest internal consistency check the paper has.

`G -> infinity`. `E -> infinity`. **Not nameable, and a genuine defect of the linear form**, repaired by the saturating form of Section 7.6 and recorded as a repair.

`L -> 1` from below and above. The transition is sharp and the equilibrium is discontinuous: 0.000000 at `L = 0.90` and 1.000000 at `L = 1.10`, from the same starting belief. **The limit is singular, and a system with an independently estimated `L` near 1 is a system whose equilibrium cannot be predicted from the model at all** — which is a domain restriction, stated as one.

`eps -> large`. `E` unrecoverable. **Nameable: the effect exists and the study cannot see it.** The commonest real-world outcome, a fact about the design rather than about the postulate, and the reason both worked studies compute power first.

`delta -> 0` in the threshold arithmetic. `G_min = (delta Q + C)/S -> C/S`, so **even with no decay at all a belief must still clear the cost of holding it**; and the storage law's own `delta = 0` pathology applies, in which the map ceases to contract and no fixed point exists [Curzi, 2026a,?].

`P(exists) -> 0` in the informational channel. `G_info -> 0` at every search time, every belief magnitude and every value per find. **This limit is realised constantly in the world and is the cleanest instance of the conditional clause in the paper.**

34.12.6 12.6 Invariance, and how far it has been tested

One invariance is claimed and it is a rescaling. Standardising `E` by its own standard deviation and `G` by the same quantity leaves `d = G B` unchanged, so **the budget arithmetic is invariant under a change of outcome unit.** That is what makes cross-channel summation possible at all, and it is also the only sense in which any quantity here is unit-free.

The invariance that is not claimed is invariance across populations, and it is not claimed deliberately. `G` is a property of an arrangement, not of a person, and an arrangement that changes changes the gain. Lucas [1976] is the general statement and Section 13 carries the restriction.

How far tested: the budget arithmetic has been exercised on three constructed cases in this paper. Three constructed cases is not a test of an invariance; it is a demonstration that the arithmetic runs. **No gain in this paper has been recovered in a second laboratory, in a second population, or by a second analyst.**

34.12.7 12.7 Significant figures against epistemic status

`0.4150`, `3.3219`, `4.3219`, `0.0840`, `0.1786` and `0.1339` are exact to the digits shown, being logarithms and products of exactly specified inputs; they are **derived from defined quantities** and quoting them to four decimal places is legitimate.

`0.3250`, `0.0550`, `8` and `53` are exact given the stipulated base rates of 40 and 60 in 100, and the base rates themselves are stipulated, not measured. The four-decimal presentation is therefore precision about an assumption, and the numbers should be read as “about a third” and “about a twentieth” wherever they are used as facts about the world rather than as steps in a computation.

`0.18`, `0.11`, `0.067` and `0.36` are model outputs from stipulated inputs and carry no intervals, because none was propagated. **That is honest for a constructed example and is not a substitute for the intervals a real application owes.** `0.15` to `0.45` for the physiological gain is a range read off a literature and is quoted as a range for that reason.

`0.3375`, `137.81`, `3.913`, `0.1105` and the power table are exact arithmetic on stipulated inputs, and the sample sizes are rounded up because a fractional participant cannot be recruited.

34.13 13 Limitations

The central claim is a postulate and remains one. `E = G B + eps`, with the forbidding clause `G = 0 => E = 0` for all `B`, is asserted by this work. **It is not derived, not measured, and not supported by any entry in the corpus’s external bibliography**, which lists it explicitly among the claims for which no supporting literature could be found. Its status is `POSTULATED` in Section 12.3 and it should be cited that way. **Any downstream paper that treats it as an established result has misused it.**

No experiment was performed and no dataset was collected for this paper. The three cases in Section 10 are constructed: their parameters are stipulated, their data are the outputs of the stated model, and they demonstrate that the procedure runs and reports its own failures. **They are not evidence that any team, organisation or patient behaves this way and must never be cited as such.**

The identification problem is unrepaired and it is the most serious limitation here. Applied to a group, the reflexive loop of Section 5.6 is a linear-in-means specification whose endogenous effect is not identified from observational data [Manski, 1993], and network data does not repair it [Shalizi and Thomas, 2011, Lyons, 2011]. **This paper adopts no remedy, so the social gain `G_social` is a modelling position rather than an identified empirical result**, and the 0.11 figure in Section 10.2 inherits that status. The realistic route to a defensible version is to stipulate the reference group and defend the stipulation [Brock and Durlauf, 2001]; that has not been done here.

Gate 5 is failed as written. `E = G B` is not identifiable from outcome data alone, and the only repair is independent, frozen, preregistered estimation of `G`. **Every use of the postulate in this corpus is conditional on that repair, and the repair does not address the reflection problem, which is a different failure.**

Gate 2 is only conditionally passed and the deeper test has not been run. Belief magnitude as normally elicited is ordinal at best, and an ordinal `B` makes `G B` a heuristic ordering rather than an equation. **The conjoint-measurement conditions that would license the ratio-scale treatment `B` receives throughout this paper have not been tested** [Luce and Tukey, 1964, Krantz et al., 1971], and Michell [1997] is right that this is the standard failure of the field. The magnitude-estimation route [Stevens, 1957] is the best available warrant and it is not a substitute for the test.

Gate 7 is partial. One channel in one domain has partial cross-laboratory support; three of four are untested in this notation. **The behavioural and informational gains have never been recovered by a second analyst, and the whole apparatus has been run once, by one author.** The replication record of the surrounding field is unambiguous about what that is worth [Open Science Collaboration, 2015, Ioannidis, 2005].

Measured against what proper model comparison requires, this paper's position is:

Requirement	Status in this paper
Out-of-sample rather than in-sample performance	**Partially met in structure, not in fact.** The protocol requires a frozen `G` and a held-out score, and the worked cases show it; **no real held-out comparison was run here.**
A matched null with the same free-parameter count	**NOT MET.** No competing account of the worked outcomes was fitted. The protocol demands four baselines including a matched AR(1); none was run in this paper.
The number of free parameters stated explicitly	**Met.** The relation has one free parameter, `G`, and it is required to be estimated outside the outcome data and frozen.
An information criterion, so the parameter count is paid for	**NOT MET.** No AIC or BIC was computed. Nothing here is fitted, so no criterion applies to this paper — **but the row is printed rather than omitted, because no AIC or BIC has been computed anywhere in this corpus**, including in the corpus's one applied preregistered fit [Curzi, 2026i]. That fit's 2.8 per cent held-out margin over a matched two-parameter AR(1) has therefore never been tested against a criterion that charges for complexity.
Effect sizes with intervals, not only p-values	**NOT MET.** No gain in this paper carries an interval, because none was propagated from its components. The honest statement is that `G_beh = 0.18` rests on a 0.30 attendance change and a 0.60 per-review effect, each of which has its own uncertainty, and neither was carried through.

The primary literature this paper depends on is not in the corpus's bibliography. The blockade result, the open-label result, the three-arm deflation, the asthma divergence, the nocebo n-of-1 design, and the conditioning-versus-expectancy work are all named in Section 8 with author, year and venue, and **none carries a bibliography key, because the corpus's external bibliography contains no entry for any of them.** Appendix C lists them. **This is a real gap in the**

corpus's apparatus and it is recorded as one rather than papered over with an approximately relevant key. A reviewer is entitled to treat Section 8 as unsourced in the formal sense, and the appropriate repair is to add the entries, not to soften the section.

Several figures in Section 8 are deliberately not quoted. Where the source manuscript declines to give a coefficient from memory, this paper declines with it. **The design and direction of each result are stated and the reader is directed to the paper for its numbers.** That is a limitation of the presentation and it is preferred to the alternative.

The effect sizes assumed throughout are at the upper end of what the literature supports. Jussim and Harber [2005] finds reflexive effects real but small, non-accumulating and often dissipating; Salganik et al. [2020] and Hofman et al. [2017] set a low ceiling on social predictability generally. **A reader who applies this framework and obtains a large gain should suspect the estimate before believing it.**

The functional forms used here are chosen, not established. The additive parallel bound assumes non-interference and is stated as a ceiling; the `tanh` saturation is an explicit modelling choice; the loop-gain recursion is a formalisation of a verbal mechanism. **And the corpus's underlying surplus functional is itself underdetermined by the evidence**: a sub-linear power law [Latané, 1981] and an `n log n` form [Briscoe et al., 2006] fit the qualitative facts as well as the logarithm does, and nothing assembled here discriminates between them. The strongest available warrant for the logarithmic form is not perceptual at all — it can be derived from multiplicative dynamics with no psychological assumption [Peters and Gell-Mann, 2016, Peters, 2019] — and **the corpus should cite that alongside Fechner and say which one its argument actually needs** [Fechner, 1860]. The corpus's own restatement is exact and worth having: the surplus law is precisely a Fechner law with threshold `u_0 = 1/(N-1)`, verified for `N = 2` to 100 at seven `u` values to a maximum deviation of `8.9e-16` [Curzi, 2026]. **That is what a verified corpus result looks like, and it is a useful contrast with this paper, none of whose central claims reaches that standard.** The same contrast holds for the corpus's exhaustively checked count results, such as the odd-part class rule under which `1..N` realises exactly `ceil(N/2)` pitch classes, six at `N = 12` [Curzi, 2026].

The Lucas critique applies to this paper's own parameters and to its publication. The framework holds that systems change their behaviour when modelled [Lucas, 1976]; a protocol telling people how to audit their own beliefs is such an intervention, and Hacking [1995] and Ferraro et al. [2005] give the mechanism. **The restriction taken here is narrow: the gains of Section 7 are asserted only within a declared window, for a declared population, and are not claimed to be structural.** Where the framework's own audit becomes a reported target, its parameters are subject to the same drift as any measure that becomes a target [Goodhart, 1975, Campbell, 1979]. **And there is a specific hazard in publishing this result: a framework asserting that most beliefs have gain zero could make that true by discouraging the search that would find the channel.** That is stated here rather than discovered later.

No claim of forecasting skill is made. Epstein [2008] enumerates sixteen reasons to build a model; this paper claims explanation, audit discipline and study sizing, and it claims none of the forecasting ones. Sobkowicz [2020] and Durlauf [2005] ask what observation would distinguish this account from a neighbouring one; **Section 3.5 answers that two of the six hypotheses would and four would not, which is a weaker answer than the question deserves.**

No scaling or heavy-tail claim is made, so the distribution-fitting procedure such a claim would require [Clauset et al., 2009] does not apply; that is recorded as a check that does not apply rather than one that was passed. **And no conservation law is claimed for any quantity here**, which matters because a stock that accumulates, leaks and is spent has the grammar of a conserved quantity without the conservation [Mirowski, 1989]; the escape is to specify the arrangement within which the relation is claimed to hold and to claim no universality beyond it [Cartwright, 1999], and that is what Section 7.2's badge does.

No clinical, medical or therapeutic claim is made for anything in this corpus, and none of its methods has been tested. No protocol described here substitutes for treatment. **The divergence diagnostic of Section 8.6 is the reason this is a safety rule rather than a legal formality: a patient may feel better while the instrument says otherwise, and acting on the feeling is how**

people are harmed.

Two observations about the source material, printed rather than silently corrected. First, the four channel gains in the worked budget sum to 0.357, and the source carries the sum as 0.36; the bound at `B = 0.7` is therefore 0.2499 rather than 0.252, and the required-over-available ratio is 3.401 rather than 3.373. The source's figures are retained because every downstream number in that case is computed from them, and the discrepancy at the third decimal is recorded here so a reader recomputing from the components does not think they have made an error. Second, the `G_physio` table quotes `d = 0.30` at `B = 0.70` as `G = 0.4286`, which is `d/B` and therefore larger than the same `d` at `B = 1.00`; that is arithmetically correct but reads as a stronger channel at weaker belief, and the table should be read as four separate calibration points from four different studies rather than as a sensitivity analysis on one.

What would falsify the account. H1 fails on a replicated, preregistered, channel-excluded effect above the noise floor. H2 fails on a pattern of effects exceeding the summed bound with gains frozen in advance. H3 fails if halving an independently measured gain does not halve the effect. H4 fails if an effect survives verified interruption of its only named channel. H5 fails if false-premised channels perform as well as true-premised ones, in which case the word "truth" should be struck from the postulate. H6 fails on a stable intermediate belief in a system whose loop gain is independently estimated away from 1. Each of these is a report this paper would print in the same plain language it would use for a success.

34.14 14 Conclusion

The postulate is: belief creates reality if the belief is grounded in truth. It is a postulate. It is badged as one in the source, it is badged as one here, and nothing in this paper promotes it. In magnitude form it is `E = G B + eps`, with `B` a normalised strength of holding on `[0,1]`, `G` a channel gain in the outcome's own units, and the load-bearing half the forbidding clause: `G = 0` implies `E = 0` for every value of `B` whatsoever. Its correspondence with the storage law is exact in form and is stated precisely for the papers that cite the pairing: `eta = 0` gives `Q_{t+1} = (1-delta)Q_t - C_t`, a vessel that leaks and pays rent, and `G = 0` gives `E = 0` at any strength of belief. Those are the same sentence at two scopes, and the correspondence transfers the model form and not the warrant.

The conditional is defended by what it forbids, and its price is printed. Under the stated base rates, 32.5 per cent of random observations refute the unconditional form — eight casual observations refute it with 95 per cent confidence — against 5.5 per cent for the conditional, which needs 53. The unconditional form taken charitably, with an unobservable sufficiency parameter, carries 0.0000 bits of empirical content; a sign-only prediction carries 1.0000; and a point prediction at tolerance `+/- 0.10` carries 3.3219, each halving of the tolerance buying exactly one bit. The conditional is six times harder to refute than the claim it replaces, and that is the flank this paper volunteers rather than conceals.

The dynamics are a bifurcation, not a gradient. With `B_{t+1} = (1-lam)B_t + lam min(1, L B_t)` and loop gain `L = alpha G`, the zero-belief state is stable if and only if `L < 1`; at `lam = 0.30` the equilibrium is 0.000000 for `L = 0.90` and 1.000000 for `L = 1.10`, from the same start. A self-fulfilling belief is exactly a belief whose loop gain exceeds one, and `G = 0` forces `L = 0` at every feedback strength, so the logical and the dynamical forms of the forbidding clause coincide.

Gains compose: series multiplies and parallel adds, and the sum is a ceiling rather than a promise. Four stages of transmission 0.70, 0.50, 0.60 and 0.40 deliver 8.4 per cent of a lossless chain, and any single zero stage zeroes the product. Four classes exhaust what a state of a nervous system can do outside the skull — physiological, behavioural, social, informational — and three of the four have their zero on a fact about the world rather than about the believer: no efferent pathway, an ineffective action, an absent target. The fifth putative class, direct action with no bodily route, has no channel specification and is assigned gain zero.

The worked audit sums a behavioural gain of 0.18, a social gain of 0.11 and an informational gain of 0.067 to 0.36, bounding the effect at 0.252 at `B = 0.7` against a claimed 0.85 —

a required-over-available ratio of 3.373 — and rejects the claim while keeping the practice that generated it. The physiological gain in that case is recorded as `G = 0` because no efferent route to the outcome exists, and recording the zero is what makes the budget meaningful. **The minimum viable gain is `(delta Q + C)/S = 0.1786` for growth and `(delta Q + C)/ln N = 0.1339` for non-negative headroom; at `G = 0` the steady state is `Q_inf = −C/delta = −0.9000`, so an ungrounded belief is a continuing net cost rather than merely an absent gain.**

Placebo supplies the one channel traced end to end, cut, and traced again, at modest measured size. Blockade removed a belief-produced effect, which is the series chain with one stage set to zero; disclosure did not remove it, which is what makes the channel compatible with a constraint that forbids concealment. **The gain is 0.15 to 0.45 on what people report and approximately zero on what instruments read independently of the person, so any claim of a standardised effect above about 0.3 through this channel alone is over budget and must name another channel.** At `d = 0.20` the distributions overlap by 92 per cent and the treated person is ahead 55.6 per cent of the time: **small per person, not small per thousand, and both statements are true.**

And the honest summary of the paper's own status: the formula validation protocol returns five and a half of seven. It fails identifiability as written, and passes only under independent frozen estimation of the gain. It is partial on reproducibility. The reflection problem is unrepaired and the social gain is a modelling position rather than an identified empirical result. A postulate that reported seven of seven would be a postulate nobody had audited.

34.14.1 Declarations

Data and code availability

No external dataset was collected for this paper. All numerical results are computed from the stated parameters under the relations developed in the text, and every table can be re-derived from the equations and parameter values printed beside it. The three cases worked in Section 10 are constructed examples with declared parameters, not field observations. The empirical results reported in Section 8 are drawn from the published literature, are cited in the text by author, year and venue, and were not re-analysed here; **where the source manuscript declines to quote a coefficient from memory, this paper declines with it and directs the reader to the primary source.** The figures are generated by the scripts named in each caption; those scripts contain no data beyond the numbers printed in this paper and the relations of Sections 5 to 8.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, produced the worked audits and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The three cases in Section 10 are constructed; no participants were recruited, no intervention was run,

and no individual is described. The study designs discussed in Section 10.3 are proposals, not reports of studies performed. **No clinical, medical or therapeutic claim is made for any method in this corpus; none has been tested; and nothing described here substitutes for treatment.** The framework's standing constraint — that no protocol built on this postulate may remove a person's ability to refuse it, and that nothing may be presented to a person masked, subliminal or concealed — is a design constraint on prospective work and is not a report of any study performed.

Preregistration

This paper reports no preregistered test of its own; its results are derivations and constructed worked cases. **The falsification conditions K1 to K5 and F1 to F4 are stated in advance in Sections 5.5 and 8.11, and the prespecified kill conditions of the study design in Section 10.3 are stated before its arithmetic.** The corpus's preregistered protocol for testing a different law, together with the confirmatory application that **failed** its second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin — is reported in Curzi [2026i]. Any applied claim resting on this paper inherits both that result and the unrepaired identification problem of Section 3.4.

34.15 Appendix A Specification register

These are design rules whose effect on practice has not been measured, not findings, and Section 13 records that distinction. They are reproduced here because downstream papers cite them by number.

34.15.1 A.1 On the postulate

S-4.1 STATE THE BELIEF AS A PROPOSITION.
A sentence that could be true or false, with a named subject and a named outcome. If it cannot be written this way it is a mood, not a belief, and is out of scope -- not forbidden.

S-4.2 NAME THE OUTCOME AND ITS UNITS BEFORE ANYTHING ELSE.
E must have units, a measurement procedure, and a stated noise level. "Things will go better" is not an outcome. "Median days from request to delivery, from the ticket system, sd approximately 2.4 days" is.

S-4.3 NAME THE CHANNEL, AS A SEQUENCE OF STEPS.
Each step a physical or social event that either happens or does not, each independently checkable, written numbered. If step 2 cannot be written there is no channel, only a hope. Record it -- that is a finding, not a failure.

S-4.4 ESTIMATE G FROM THE CHANNEL, NOT FROM THE OUTCOME.
Binding, per Gate 5. Any G fitted to outcome data makes the relation unidentifiable and the exercise void.

S-4.5 FREEZE G, AND PUBLISH IT, BEFORE MEASURING E.
Timestamped, witnessed, or preregistered. The freeze is what makes the subsequent measurement a test.

S-4.6 PREDICT E WITH A TOLERANCE.
E_pred = G.B, with a stated band. Each halving of the band buys one bit. A band so wide nothing could fall outside it is the escape clause in metric clothing.

S-4.7 POWER THE MEASUREMENT AGAINST THE PREDICTION.
Compute the sample needed to detect E_pred against the noise BEFORE collecting data. Running an underpowered test and reporting a null is a category error.

S-4.8 INCLUDE A NO-CHANNEL ARM WHEREVER POSSIBLE.
The postulate's content is concentrated in its zero predictions. A design with no arm in which zero is predicted cannot distinguish the conditional from the unconditional version.

S-4.9 RUN THE SHUFFLED-INPUT CONTROL.
Same frozen parameters, belief series permuted. If performance survives the shuffle there is no channel. Non-negotiable.

S-4.10 REPORT THE MAGNITUDE RESULT SEPARATELY FROM THE SIGN RESULT.
"In the predicted direction" and "of the predicted size" are two findings, and the second is the one with content.

S-4.11 DEMOTE IN PRINT WHEN A GATE FAILS.
Likert-elicited B makes the relation a heuristic ordering and it must be labelled one. G fitted to outcome data voids the result. The demotion appears in the same document as the result.

S-4.12 NEVER RESCUE THE POSTULATE WITH A PROPERTY OF THE BELIEVER.
"The belief was not strong enough" is forbidden as an explanation of a null: B is measured before the outcome and is already in the prediction. If B was low the prediction was low, and a low observed effect CONFIRMS the postulate rather than excusing it.

S-4.13 THE CLINICAL PROHIBITION, ABSOLUTE.
Nothing in this corpus constitutes a clinical, medical or therapeutic claim, and no method described here has been shown to produce any physiological effect. No belief protocol substitutes for treatment. This is a consequence of the postulate: what produces the outcome is the channel, and if the channel is a drug, the drug is what is needed.

S-4.14 PRESERVE THE RIGHT TO DECLINE.
No protocol built on this postulate may remove a person's ability to refuse it. A belief installed without consent is not a belief the person holds; it is a state induced in them, and the postulate has nothing to say about it because B was never theirs.

34.15.2 A.2 On the gain budget

S-7.1 NAME THE CHANNEL CLASS BEFORE ESTIMATING ANYTHING.
Physiological, behavioural, social, informational, or none. If the answer is a combination, decompose it. If the answer is "none of these," the estimate is G = 0 and the analysis is finished.

S-7.2 WRITE THE SERIES CHAIN AS NUMBERED STAGES.
Each stage a distinct event with its own transmission ratio. The total is the product. Publish the stages, not only the product.

S-7.3 ESTIMATE EVERY STAGE GAIN FROM DATA THAT IS NOT THE OUTCOME.
Binding, per Gate 5. This is the whole of what makes the exercise a test rather

than a curve fit.

S-7.4 IDENTIFY THE WEAKEST STAGE AND STATE ITS RELATIVE SENSITIVITY, 1/g_i.
That stage is where intervention pays. Do not improve a stage already at 0.9.

S-7.5 SUM PARALLEL CHANNELS ONLY AFTER STANDARDISING TO COMMON UNITS.
Days and symptom points may not be added. Standardise by each outcome's own standard deviation and say so.

S-7.6 TREAT THE SUM AS A BOUND, NEVER AS A PREDICTION.
Channels interfere. Predict below the bound and explain the discount.

S-7.7 RUN THE REQUIRED-VERSUS-AVAILABLE AUDIT BEFORE COMMITTING.
Compute the gain the claim requires and divide by the gain the named channels supply. A ratio above 1 means the claim is unfunded. Publish the ratio.

S-7.8 STATE E_MAX FOR EVERY CHANNEL.
A gain without its ceiling is an incomplete specification.

S-7.9 CHECK THE MINIMUM VIABLE GAIN.
Compute (delta.Q + C)/S. A channel below it loses ground even though it is real. "Real but insufficient" is a distinct and common verdict that no other line of the analysis produces.

S-7.10 FOR THE INFORMATIONAL CHANNEL, ESTABLISH P(TARGET EXISTS) FIRST.
If the target's existence is the question, the gain is unknown, not positive.

S-7.11 FOR THE BEHAVIOURAL CHANNEL, MEASURE THE ACTION'S EFFECT SEPARATELY FROM THE BELIEF.
Vary the action without varying the belief. Skipping this converts a behavioural claim into an article of faith, and it is the most frequently skipped step in the applied literature.

S-7.12 FOR THE SOCIAL CHANNEL, COUNT THE COMMITTED AND MAP THE TOPOLOGY.
Committed means acting at a cost. Clustered placement beats scattered placement at equal headcount. State the field-transfer assumption on any threshold number used, AND state that the parameter is not identified from observation.

S-7.13 FOR THE PHYSIOLOGICAL CHANNEL, NAME THE EFFERENT PATHWAY OR ESTIMATE ZERO.

S-7.14 PRESERVE THE ABILITY TO DECLINE, PARTICULARLY IN THE SOCIAL CHANNEL.
A protocol that raises G_social by removing a person's option to dissent is out of scope regardless of its measured gain, and destroys the group's non-parallelism in the bargain.

S-7.15 PUBLISH THE WHOLE BUDGET, INCLUDING THE ZEROES.
A channel estimated at zero is a result. A budget listing only the positive channels has told the reader nothing about what was considered and excluded.

34.15.3 A.3 On the measured channel

S-8.1 THREE ARMS, ALWAYS. Active (if any) / belief manipulation / NO TREATMENT.
Without the third arm the effect cannot be separated from natural history and regression to the mean, which Section 8.5 computes at 40% of what would otherwise be reported.

```
S-8.2   MATCH THE INTERACTION ACROSS ARMS.
        Otherwise the comparison measures attention, not belief.

S-8.3   MEASURE A SUBJECTIVE AND AN OBJECTIVE OUTCOME IN THE SAME PATIENTS.
        Report D = d_subj - d_obj as a headline.

S-8.4   PRE-SPECIFY THE PRIMARY OUTCOME AND THE ANALYSIS.

S-8.5   POWER FOR THE EFFECT THE FIELD ACTUALLY FINDS, NOT THE ONE HOPED FOR.

S-8.6   REPORT THE GAIN, NOT ONLY THE EFFECT.
        G = d/B, with B measured by magnitude estimation on the same scale across arms.

S-8.7   RUN THE OVER-BUDGET CHECK ON ANY CLAIM ENCOUNTERED.
        Standardise the claimed effect; compare to 0.3; if it exceeds that through this
        channel alone, demand the additional channel by name.

S-8.8   DISCLOSE FULLY, ALWAYS, AND ACCEPT THE NOCEBO COST.
        No concealment, no subliminal presentation, no withheld information. This raises
        nocebo. It is accepted without argument, because consent is not a term in an
        optimisation.

S-8.9   NEVER PRESENT THE PRECEDENT COLUMN AS EVIDENCE.

S-8.10  STATE THE UNCERTAINTY IN THE SENTENCE THAT MAKES THE CLAIM.
        Where a source's exact figures are not to hand, give the shape of the result and
        say the numbers should be read from the paper. An invented decimal is worse than
        an admitted gap.

S-8.11  NO CLINICAL, MEDICAL OR THERAPEUTIC CLAIM, ABSOLUTE.

S-8.12  PRESERVE THE RIGHT TO DECLINE.
```

34.16 Appendix B Evidence-and-precedent audits

Each audit states what was measured, what was merely recorded, and what the pair can and cannot establish. **Two sources agreeing does not double the evidence when only one of them is evidence.**

B1 — A belief affects the world only through a real pathway; with no pathway the effect is zero regardless of the belief's strength. *Evidence:* placebo analgesia in postoperative dental pain reversed by an opioid antagonist (Levine, Gordon & Fields, *The Lancet*, 1978). **A specific pharmacological blockade abolished a belief-produced effect, which is the signature of a specific mechanism rather than a general one.** This is a measurement with a control. *Precedent:* the "law of contact or contagion" catalogued in the comparative literature on ritual practice — the requirement of a material link between operator and target, hair or a garment or a nail-paring. **The requirement of a physical route between intention and effect was stated, in that notation, centuries before anyone could measure one. This is precedent only: the catalogue is of practice, not a test of it, and the practices catalogued do not work.** *Assessment:* both columns converge on the same structural demand — intention alone is inert, a route is required — and the evidence column adds what the precedent column could not: a route that was actually found, and a blockade that actually removed it. C1 pass, C2 pass (the pairing predicts that blocking the intermediary abolishes the effect, which was predicted and tested), C3 pass (it converts "does belief work?" into "name the channel and estimate its gain"). *Status:* **established as a structural correspondence, with the evidential content entirely in the**

evidence column.

B2 — A doctrine that belief is sufficient on its own is not merely unsupported; acting on it has a documented cost. *Evidence:* a review of child fatalities in which medical treatment was withheld on religious grounds concluded that the great majority would have survived with standard care (Asser & Swan, *Pediatrics*, 1998). **The study's conclusion is stated and its counts are deliberately not quoted from memory; the primary source should be read before any number is repeated.** *Precedent:* the popular literature in which belief is presented as sufficient and the conditional clause is absent. **This is precedent for the error, and recording where a tradition went wrong is part of reading it honestly.** *Assessment:* **the unconditional version is the one with a body count. The conditional version says: name the channel, and if the channel is antibiotics, use antibiotics.** *Status:* **established.**

B3 — A physical input reaches cellular and systemic function only through a real, identifiable transduction pathway. *Evidence:* mechanically gated ion channels — mechanical force opens them, ions flow, signalling follows — for which the pathway from mechanical input to cellular signal was found, characterised, named as a specific protein, and recognised by the 2021 Nobel Prize in Physiology or Medicine. **A measured, molecularly identified channel.** *Precedent:* the traditional ethical maxim that conviction does not substitute for the rope — trust, but tie your camel. **This records that the channel requirement was arrived at as a practical ethic long before any transduction pathway was known; the attribution of the saying is itself a scholarly question and nothing in the argument depends on it being exact.** *Assessment:* both hold that intention without a route is inert; the evidence column supplies a named protein, a measurable current, and a knockout experiment. C1, C2 and C3 all pass. *Status:* **established as a structural correspondence. And the limit, stated in the same block: none of this establishes a channel for sequence-specific content. Mechanotransduction is a channel for force, not for text, and the distinction is the whole of the corpus's biological material** [Curzi, 2026].

B4 — Belief acts on physical systems at a distance without a bodily route. REJECTED. *Evidence:* **none. No carrier is specified.** Every established action-at-a-distance in physics has a named field, a coupling constant, and a distance dependence; this proposal has none of the three. C1 fails: no structural correspondence, because there is no structure. C2 fails: no magnitude is predicted in advance from mechanism, so nothing can be over or under budget and the audit of Section 7 cannot be run. C3 fails: it cannot be used on either side. *Status:* **rejected, and the postulate's own prediction for this class is `E = 0`.** The quantitative version of the rejection, where the proposal invokes physics, is in the companion paper: the corpus's conjugate material is badged analogy and not identity, with an explicit refusal that quantum mechanics licenses anything about belief, mind, or future events [Curzi, 2026z], and warm-wet coherence fails by thirteen orders of magnitude [Curzi, 2026aa].

B5 — Because placebo produces measurable physiological change, belief can produce any physiological change, including remission of organic disease. REJECTED. C1 fails: the measured pathways are specific — descending opioidergic modulation of nociception, dopaminergic signalling in a specific circuit, autonomic and endocrine responses — and each has an anatomy. **None of them terminates on a tumour, a pathogen, or a demyelinated axon, and generalising from "some pathway exists" to "any pathway exists" is the fifth-class smuggle of Section 9.2.** C2 fails in the decisive direction: it predicts objective-endpoint effects that are not observed (Section 8.6). C3 is worse than useless: acting on it means declining treatment, and audit B2 gives the evidence on what that costs. *Status:* **rejected, in the paper that presents the strongest evidence for the channel, because that is where the rejection has to be made.**

B6 — Shared belief coordinates behaviour, and the coordination has a measurable threshold rather than a smooth response. *Evidence:* a committed minority of approximately 25 per cent overturned an established social convention in controlled online experiments [Centola et al., 2018] — a measured tipping point with a manipulated independent variable. **Transfer to a field organisation is a modelling assumption, not a measurement, and is labelled so wherever the number is used; and the underlying parameter is not identified from observational data [Manski, 1993].** *Precedent:* accounts of collective assembly and shared enact-

ment producing commitments that individuals did not hold separately. **These are theoretical and ethnographic accounts, not controlled experiments, and they do not establish a threshold value.** *Assessment:* both describe a nonlinearity — something changes discontinuously when enough people hold and enact the same thing. The evidence column supplies a number for one setting; the precedent column supplies the observation that the phenomenon was noticed long before it was measured. C1, C2 and C3 pass. *Status:* **established as a correspondence; the NUMBER is established only for the measured setting.**

B7 — Expectation, produced by what a person is told, can produce a measured physiological change through an identified neurochemical pathway. *Evidence:* the 1978 blockade result, extended by later neuroimaging and pharmacological work, and by a study reporting placebo-associated dopamine release in the striatum of patients with Parkinson's disease measured by raclopride PET — **a neurochemical outcome read by an instrument rather than reported by the patient.** *Precedent:* temple incubation at the Asclepian sanctuaries and the cure inscriptions at Epidaurus; and the sixteenth-century position that the patient's imagination is a real cause in medicine. **Precedent only, and precedent with a named defect, which is the honest way to carry it: the cure inscriptions are a register of cures and not of failures. They are survivorship-selected by construction. They establish that the phenomenon was noticed and institutionalised millennia before it was measured, and they establish nothing whatever about its size.** *Assessment:* both columns describe expectation producing bodily change; the evidence column supplies the receptor, the antagonist, the control arm, and the denominator the precedent column structurally cannot have. *Status:* **established as a correspondence, with all evidential content in the evidence column.**

B8 — The same channel run in reverse produces symptoms: negative expectation generates measurable adverse experience. *Evidence:* an n-of-1 randomised trial in which participants who had discontinued a drug because of side effects took, in randomised blinded periods, the drug, a placebo, and no tablet; symptom scores during placebo periods were reported as close to those during drug periods and both substantially above the no-tablet periods. **The structure of the result is given; the reported ratio was high and the paper should be read before any figure is repeated.** This is a within-person randomised design, the strongest available for the question. *Precedent:* traditions of the curse and the evil eye, recorded across many cultures, together with a mid-twentieth-century compilation of anthropological reports of death following such beliefs. **Precedent only, and weak precedent: the compilation is of secondhand anecdote, not a study, and it must not be cited as evidence of anything.** *Assessment:* **nocebo is not a separate phenomenon requiring its own theory. It is the same channel with the sign of the expectation reversed, and treating it as one object with two signs is both simpler and better supported.** *Status:* **established as a correspondence.**

34.17 Appendix C Works named in the text without a bibliography key

The corpus's external bibliography contains no entries for the following, and no key has been attached to them anywhere in this paper. They are listed with the details the source manuscript supplies so that a reader can locate them. **Their absence is a gap in the corpus's apparatus and is recorded as one rather than concealed. Nothing in this list is a citation; it is a statement of what is missing.**

Placebo and nocebo. Levine, J. D., Gordon, N. C. & Fields, H. L. (1978), "The mechanism of placebo analgesia", *The Lancet*. Amanzio, M. & Benedetti, F. (1999), *Journal of Neuroscience*, on opioid and non-opioid components. Wager, T. D. et al. (2004), *Science*, on placebo-associated changes in pain-related brain activity. Eippert, F. et al. (2009), on descending opioidergic control and the spinal cord. de la Fuente-Fernández, R. et al. (2001), *Science*, on placebo-associated striatal dopamine release measured by raclopride PET. Kaptchuk, T. J. et al. (2010), "Placebos without deception: a randomized controlled trial in irritable bowel syndrome", *PLoS ONE*, n = 80. Carvalho, C. et al. (2016), *Pain*, on open-label placebo in chronic low

back pain. von Wernsdorff, M. et al. (2021), *Scientific Reports*, the meta-analysis of open-label placebo randomised trials. Hróbjartsson, A. & Gøtzsche, P. C. (2001), *New England Journal of Medicine*, and the 2010 Cochrane review, on placebo versus no treatment. Wechsler, M. E. et al. (2011), *New England Journal of Medicine*, on the subjective–objective divergence in asthma. Wood, F. A. et al. (2020), *New England Journal of Medicine*, the SAMSON n-of-1 randomised trial. Benedetti, F. and colleagues on cholecystokinin and proglumide in nocebo hyperalgesia, and on conditioned hormonal responses not modifiable by verbal expectation.

Harm from the unconditional doctrine. Asser, S. M. & Swan, R. (1998), *Pediatrics*, on child fatalities where medical treatment was withheld on religious grounds.

Mechanotransduction. The identification and characterisation of PIEZO1 and PIEZO2 as mechanically gated ion channels, recognised by the 2021 Nobel Prize in Physiology or Medicine. Tajik, A. et al. (2016), *Nature Materials*, on force at the cell surface reaching chromatin and altering transcription.

Precedent sources named in Appendix B. Frazer, J. G., *The Golden Bough* (1890), for the law of contact or contagion. Durkheim, É. (1912), on collective effervescence, and Rappaport, R. A., *Ritual and Religion in the Making of Humanity* (1999). The *iamata* of Epidaurus, fourth century BCE. Paracelsus, sixteenth century, on the imagination of the patient as a cause. Cannon, W. B. (1942), “Voodoo Death”, *American Anthropologist* — a compilation of secondhand anecdote, cited in Appendix B explicitly as weak precedent and not as evidence.

Chapter 35

Where Belief Fails: Four Failure Classes, a Priced Bank Run, and the Grounding Protocol That Kills Most Claims

Abstract. This is the corpus's counter-result paper. Having proposed that belief acts only through a real channel, we catalogue where it does not act and price what survives. Failure is decomposed into four distinct classes — no channel, real channel with insufficient gain, a real laboratory effect extrapolated to an unsupportable magnitude, and a system unresponsive to an input of this size from a source of this size — and the diagnostic value lies entirely in telling them apart. Growth mindset is placed in class B at r ~ 0.10 (1 per cent of variance) and intervention d ~ 0.08, which is 96.8 per cent distribution overlap, a 52.3 per cent common-language effect and 0.056 GPA points. The worked job-search case runs every real channel at full strength and lifts the offer probability from 0.0140 to 0.0697 — a 4.975-fold ratio in which the outcome is still no offer 93 per cent of the time, with 39.5 per cent of the lift coming from simply submitting the application and the no-contact channel carrying gain zero by construction. The bank run shows the opposite extreme: at a 9 per cent reserve the withdrawal fraction swings from 99.01 to 14.51 per cent on a 0.002 change in mean depositor tolerance. Runaway is bounded by ln N: reflexive slopes past 0.15 demand more input than any realistic category count admits.

Keywords: self-fulfilling prophecy; failure taxonomy; growth mindset; bank run; reflexivity; survivorship bias; likelihood ratio

Series: B. Engineering the Surplus **Paper:** CCP-035

35.1 1 Introduction

The framework this paper belongs to advances a conditional postulate: **belief creates reality if the belief is grounded in truth**, where grounding means that a real channel exists from the belief to the asserted outcome and carries a measurable gain [Curzi, 2026ab]. **A conditional postulate is worth exactly what its failure catalogue is worth.** A claim that names no failure has not been formulated tightly enough to fail, which means it has not been formulated tightly enough to be true either.

This is the corpus's counter-result paper, and it is written for that purpose. It catalogues where the postulate does not pay, prices what survives at the magnitudes the literature actually reports, and states the objection that would destroy the postulate outright before offering the only answer available to it. Three things follow in order: the mechanism by which belief does sometimes close a loop and what that mechanism structurally requires; a four-class decomposition of failure, with each class carrying its own diagnostic; and a seven-step protocol that establishes grounding from the channel rather than from the outcome.

The headline numbers are deflationary and are reported that way. Growth mindset — the most widely cited case of belief acting on outcomes — falls in the class where the channel is real and the gain is roughly an order of magnitude below the claim: a correlation near $r = 0.10$, which is **1.00 per cent of variance**, and intervention effects near $d = 0.08$, which is **96.809 per cent distribution overlap**, a **52.256 per cent common-language effect**, and **0.0560 GPA points** on a 0.70-point standard deviation. A worked job-search decomposition running every real channel at full strength lifts the offer probability from **0.0140 to 0.0697** — a **4.975-fold** ratio in which the outcome is still no offer **93 per cent of the time**, with **39.53 per cent of the entire lift** coming from the belief causing the application to be submitted at all, and with the no-contact channel carrying **gain zero by construction**.

The opposite extreme is also computed, because a catalogue that reports only small effects is as dishonest as one that reports only large ones. At a 9.00 per cent reserve ratio a deposit institution's run threshold is $f^* = 0.0900$, its critical forced-sale recovery rate is $h^* = 0.941379$, and with member tolerances distributed Normal$(\mu, 0.25)$ the critical mean is $\mu^* = 0.408522$, on either side of which the withdrawn fraction is **99.01 per cent or 14.51 per cent — a swing produced by a change of 0.002 in an unmeasured population parameter.** Where the structure is responsive, belief can be decisive. Where it is not, no quantity of belief substitutes.

And the runaway is bounded. Written in the framework's own storage law, a reflexive input $S(Q) = S_0 + sQ$ gives the pole $a = (1 - \delta) + \eta s$ and the loop gain $g = \eta s/\delta$. With $\delta = 0.10$ and $\eta = 0.50$ the critical slope is $s^* = 0.20$, but $s = 0.19$ **already implies 12.6 nats of input against a ceiling of** $\ln N = 2.996$ **at** $N = 20$ [Curzi, 2026]. The divergence never arrives; the model leaves its admissible region first. **And the time constant** $\tau = 1/(1 - a)$ **goes tenfold at a loop gain of 0.90, so the loop slows conspicuously before it destabilises** — which is a measurable early-warning signature available from response time alone.

Four contributions are made relative to the surrounding literature. **First**, a decomposition of the self-fulfilling mechanism into three separately checkable and separately destroyable requirements, with the multiplicative model $E = B \cdot K \cdot R$ making the conjunction explicit. **Second**, a four-class failure taxonomy in which the diagnostic value lies entirely in telling the classes apart, because they prescribe different and incompatible responses. **Third**, an evidential pricing of prediction: a direction-only claim is worth a likelihood ratio of **2.0000**, while a ± 20 per cent magnitude call against a two-decade prior is worth **11.3577**, or **3.5056 direction calls for the same single measurement**. **Fourth**, the statement of the circularity objection — that a conditional whose antecedent is diagnosed from the outcome is vacuous — together with the procedural answer that is the only one available.

The paper applies its own discount to itself. The replication literature is the evaluative context for every claim in this area including the framework's own: a prior of 0.10 updated on one significant result under typical power and heavy selective reporting yields a posterior of **0.1148** — a movement of under two percentage points [Open Science Collaboration, 2015]. **If this framework's predictions were tested and returned a single significant result, that is the update the author would be obliged to report**, and Section 5.10 says so in print.

Section 2 states the hypotheses. Section 3 gives the method, including the analogy labels, the identification problem that this paper is more exposed to than any other in the corpus, and the scale-type concession. Section 4 develops the engine, Section 5 the failure catalogue, Section 6 the grounding protocol. Sections 7 to 11 report, discuss, verify, limit and conclude. Appendices A to C carry three worked cases.

35.2 2 Hypotheses

Several results below are theorems and cannot be falsified by observation; they are labelled as derivations where they appear and recorded as `ARITHMETIC` or `DEFINITION` in Section 9's chain. What can be wrong are the taxonomy's claims to exhaustiveness and diagnostic value, the magnitude claims, and the claim that the protocol is discriminating rather than a formality.

H1 — The failure taxonomy is exhaustive over documented cases. Every documented failure of a belief claim falls into one of four classes: no channel, real channel with insufficient gain, a real result extrapolated to an unsupportable magnitude, or a system unresponsive to an input of this size from a source of this size. **Predicted:** each case examined in Section 5 admits exactly one class, and the class assignment is made from the channel rather than from the outcome. **Falsified by:** a documented failure that fits none of the four, or that fits two with no way to choose. **Result: no counterexample was found among the cases examined here, which is weak evidence for a taxonomy and is reported as such** — the cases were selected by the author, and Section 10 records that as a selection problem.

H2 — The most-cited belief effect is small, not absent. Growth mindset occupies the class in which the channel is real and the gain is roughly an order of magnitude below the popular claim. **Predicted:** meta-analytic correlations near $r = 0.10$ and intervention effects near $d = 0.08$, which unpack to distribution overlap above 95 per cent and a common-language effect within three points of chance. **Falsified by:** a meta-analytic estimate an order of magnitude larger, or by a demonstration that the channel does not exist at all, which would move the case to a different class. **Result: confirmed at the magnitudes the meta-analytic literature reports.** 96.809 per cent overlap, 52.256 per cent common-language effect, 0.0560 GPA points.

H3 — Where the structure is responsive, the outcome is knife-edged rather than proportionate. Predicted: in a threshold model of a deposit run, the withdrawn fraction is a step function of the population's mean tolerance, with no intermediate outcome available near the critical value. **Falsified by:** a smooth, proportionate dependence of the outcome on the mean, which would make sentiment measurement a useful predictor. **Result: confirmed in the model.** At $\sigma = 0.25$ the withdrawn fraction moves from 0.990106 to 0.145087 across a change of 0.002 in μ. **This is a property of the model, not a measurement of any institution**, and Section 10 says so.

H4 — The runaway terminates at the framework's own ceiling rather than diverging. Predicted: every row of the divergence table beyond $s = 0.15$ implies an input exceeding $\ln N$ for any realistic N, so the model leaves its admissible region before the pole reaches unity. **Falsified by:** an admissible parameterisation in which the steady state diverges while the implied input stays under the ceiling. **Result: confirmed by arithmetic.** At $s = 0.19$ the implied input is 12.6 nats against $\ln 20 = 2.995732$.

H5 — A magnitude prediction is worth several direction predictions at the same measurement cost. Predicted: under a log-uniform prior spanning R decades, a $\pm p$ magnitude call has a likelihood ratio far above the value 2.0000 of a direction call, and the equivalence in direction-calls is computable. **Falsified by:** a prior specification under which the magnitude call's likelihood ratio falls to the direction call's. **Result: confirmed under the stated log-uniform prior**, at 11.3577 against 2.0000, or 3.5056 direction calls — **and the result is a property of that prior specification, which is a modelling choice and is badged as one.**

One hypothesis is deliberately not advanced. No hypothesis is offered that the conditional postulate itself is true, because as stated it admits no observational test until the grounding

protocol has been run on a real case, and it has not been. **The postulate's status in this paper is that of an object under audit, not a claim under test**, and Section 10 records the audit's own limitations.

35.3 3 Method

35.3.1 3.1 What the paper attempts

The paper audits a postulate by cataloguing where it fails and by specifying the procedure that would let it be tested. **It reports no new measurement.** Every empirical magnitude below is quoted from published work, named in the text, and carries that work's uncertainty rather than this paper's; every other number is either exact arithmetic or an illustrative model evaluated at parameters that are stated as assumed in the sentence that uses them.

Two of the three worked cases are constructed. The deposit institution of Appendix A is a constructed balance sheet, the job-search decomposition of Appendix B runs on assumed conditional probabilities, and both say so wherever a number appears. **What those cases demonstrate is the decomposition procedure, not the values.**

35.3.2 3.2 Analogy discipline, declared in advance

`RIGOR_STANDARD` §8 requires every borrowing from physics to be labelled, at first use, as a derivation, a shared mathematical form, or an analogy. The labels for this paper are fixed here.

Borrowing	Label	What licenses the label
loop gain and the geometric multiplier $1/(1-g)$	**shared mathematical form**	the same series governs an electrical feedback loop and a belief loop; **no property of the electrical case transfers**, and the belief case's g must be estimated in its own units
the linearised pole $a = (1-\delta) + \eta s$ and stability at $\eta s < \delta$	**derivation**	it follows from the framework's own storage law by adding one term [Curzi, 2026a]; nothing is imported
critical slowing (τ rising before instability) as an early-warning signature	**shared mathematical form**	the same first-order relation holds; the *empirical* early-warning literature is not invoked as support, because no such measurement is made here
the Casimir pressure computed at biological scale	**derivation** within physics, used here as a **counter-example**	the formula is exact; the point is that it is negligible, which is the opposite of an analogy's usual use [Casimir, 1948, Lamoreaux, 1997]
the magnetic-field energy budget of Section 5.4	**derivation**	$u = B^2/2\mu_0$ is exact and the comparison is dimensional
"the engine", "the amplifier", "the seed"	**analogy**	vocabulary only; **never evidence**, and no engineering result is transferred by the metaphor
the effect model $E = B \cdot K \cdot R$	**modelling choice**, not a borrowing	multiplicativity is assumed for its conjunction property; it is badged in Section 4.2 and is not derived

Entropy is not used in this paper. No Shannon and no thermodynamic entropy appears, and no quantity below carries units of nats or bits except the surplus ceiling $\ln N$, which is a

Shannon-type logarithm of a count, dimensionless and measured in nats, imported from the framework's own surplus functional [Curzi, 2026]. **It is not a thermodynamic entropy and no energetic interpretation attaches to it.**

No effective count is used in this paper other than by citation. Where the corpus's affine $N_{\text{eff}} = 1 + (N-1)u$ is referenced [Curzi, 2026], it is referenced as the corpus's own construction and **not** as a Hill number, which it is not; the Hill numbers of orders one and two — $\exp(H)$ and $1/\text{HHI}$ — appear only in the cited compartment result [Curzi, 2026], where the effective-compartment count **is** $1/\text{HHI}$, the Hill number of order two, and is correctly grounded.

35.3.3 3.3 Identification: the objection this paper is most exposed to

This is the most serious methodological problem in the framework, and this paper is the one most exposed to it, because the reflexive loop is its subject.

The loop $T(x) = W(x, M(x))$ — the world's transition rule depends on the model the agents in it hold — **is exactly a linear-in-means reflexive specification once it is linearised, and Manski proves its endogenous effect is not identified from observational data** [Manski, 1993]. The reflexive slope $s = dS/dQ$ of Section 4.11 is that parameter. **No quantity of observation separates it from contextual effects or from correlated unobservables.** Shalizi and Thomas reinforce the point for network data: influence and selection are generically confounded, and homophily and contagion are not distinguishable from observational network data without further assumptions [Shalizi and Thomas, 2011, Lyons, 2011].

Which remedy does this paper adopt? For the analytical sections, **none**. The reflexive slope s, the loop gain g, and every regime classification derived from them are **modelling positions rather than identified empirical results**, and they are written that way throughout. **For the protocol of Section 6, one remedy is adopted and its price is stated: randomisation of the belief** (defence CC-1 of Section 6.13). Where the belief is *assigned* rather than observed, no common cause can precede it, and the endogenous effect becomes identified by design. **That is the only complete defence, it is expensive, and it is often unavailable** — which is why the alternatives are enumerated with their partial coverage rather than presented as substitutes. The route of stipulating the reference group and defending the stipulation [Brock and Durlauf, 2001] is named as available and is **not** taken here.

The simulated demonstration in Section 6.13 is the quantitative form of the objection. A shuffled-input control with *no coupling at all* and a shared slow driver returns the ordered series beating the shuffle in **98.35 per cent of runs**. The control cannot distinguish a real coupling from a common cause. **That number is printed rather than buried, because it is the strongest available demonstration that a statistical control does not substitute for identification.**

The Lucas critique applies to this paper's own subject matter with unusual force [Lucas, 1976]. A framework asserting that systems change behaviour when modelled cannot exempt its own parameters. Section 4.15's self-defeating case is exactly that: a forecast that provokes its own refutation has $s < 0$, and its accuracy record is consumed by its usefulness. **A measure that becomes a target ceases to be a measure** [Goodhart, 1975, Campbell, 1979], and the corpus's own instruments are not exempt. The performativity literature documents reflexive fixed points observed in the field, with dates and price data [MacKenzie and Millo, 2003, MacKenzie, 2006, Ferraro et al., 2005, Callon, 1998], and the looping of human kinds under classification is the same mechanism in a different register [Hacking, 1995]. **Those are the corpus's best existence proofs for reflexivity and they are cited as such — and none of them identifies the parameter.**

35.3.4 3.4 Scale type, and the conjoint test that was not run

`RIGOR_STANDARD` §5 requires every measured construct to declare its scale type [Stevens, 1946, Krantz et al., 1971]. Four constructs need one.

- **Money, deposits, reserve ratios, withdrawal fractions and time are ratio-scaled** with true zeros, and every operation performed on them below is licensed.
- **Standardised effect sizes (d), correlations (r), probabilities and likelihood ratios are ratio-scaled by construction**, given a stated reference standard deviation; **the reference standard deviation must be the target population's, not the study's**, and where a d is converted to GPA points below, the standard deviation used is stated in the same line.
- **Belief measured on an agreement scale is ORDINAL**, and the multiplication in effect $=$ gain $\times$ belief **is not licensed on it** [Stevens, 1946, Likert, 1932]. Section 6.11 demotes the relation to a monotone ordering wherever the instrument is ordinal, and the worked case of Appendix C is constructed so that the ordinal instrument is used **only to define a group membership** and is never multiplied. **That construction was chosen for this reason and the choice is recorded.**
- **The reflexive slope s inherits the scale of whatever Q is instrumented in**, and the framework does not fix that; the audit in Section 9.1 records the framework as **silent** rather than assigning one.

Michell's critique is the strongest objection to the framework this paper belongs to [Michell, 1997, 1999]: psychology assumes quantitative structure without testing for it. The formal test is additive conjoint measurement [Luce and Tukey, 1964]. **It has not been run on belief, on confidence, on conviction, or on any psychological construct in this paper, and this paper does not run it.** A ratio-scale instrument for a subjective magnitude is constructible by magnitude estimation [Stevens, 1946, Thurstone, 1927], and building one is more work than distributing a five-point form; **the difference between the two is the difference between a formula and a decoration.** Until it is built, every multiplicative statement involving a belief measure in this paper is **demoted in print** — and internal-consistency coefficients are not a substitute for the missing structural test [Cronbach, 1951, Sijtsma, 2009].

35.3.5 3.5 Scaling claims, and the procedure not run

Two power-law statements appear below and both are **defined by the model rather than estimated from data**, which Clauset et al. [2009] requires to be said explicitly.

- The Casimir pressure falls as a^{-4}. **That exponent is an exact consequence of the formula $P(a) = \pi^2 \hbar c/(240 a^4)$ and is not an estimate.**
- The chain-gain law $g_{\text{total}} = \prod_i g_i$ produces geometric decay in the number of links. **The decay constant is the per-link gain, an assumed parameter in every table below, not a fitted exponent.**

No heavy-tail, scale-free or power-law claim is made anywhere in this paper on empirical data, and the maximum-likelihood-plus-goodness-of-fit procedure of Clauset et al. [2009] was therefore not run. It is named so that its absence is a decision rather than an omission; the literature on mis-fitted power laws is the reason [Stumpf and Porter, 2012, Broido and Clauset, 2019].

35.3.6 3.6 What is presupposed

The storage law $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$, its fixed point and its geometric convergence are presupposed from Curzi [2026a] and are not re-derived; the reflexive term of Section 4.11 is one addition to that equation. The surplus ceiling $0 \le F(u) \le \ln N$ is presupposed from Curzi [2026]. The belief dynamic $B_{t+1} = (1-\lambda)B_t + \lambda \min(1, LB_t)$ with loop gain $L = \alpha G$, whose zero fixed point is stable exactly when $L < 1$ and whose transition at $L = 1$ is sharp, and the fact that $G = 0$ forces $L = 0$ at every α, are presupposed from Curzi [2026ab] — **that paper states the postulate and its channel-gain architecture; this one states where they do not pay.** The decoherence shortfall that removes quantum mechanics from the list of candidate channels — $\tau_{\text{dec}} = 1.98 \times 10^{-17}$ s against a 1 ms neural event, a factor of 5.04×10^{13}, still 5.04×10^{7} after

a millionfold protection factor — is presupposed from Curzi [2026aa]. The seven formula-validation gates are specified in Curzi [2026t] and the four-baseline falsification suite, including the shuffled-input control and its own measured anti-conservatism, in Curzi [2026u]. The concordance discipline — that the evidence column and the precedent column carry different things, and that two agreeing sources are not two measurements when only one of them is a measurement — is specified in Curzi [2026w] and applied without restatement of its derivation.

35.3.7 3.7 Standing honesty clauses

1. **No clinical, medical or therapeutic claim is made anywhere in this paper, in either direction.** Where published trials are reported, they are reported as evidence about a channel and its gain, never as guidance about care. This is absolute and admits no exception.
2. **No citation or statistic is fabricated.** Where the author is not certain a source says what would be convenient, the claim is omitted or the uncertainty is stated in the sentence.
3. **Precedent is never presented as evidence.** The precedent column records that a structure was independently arrived at; it does not add to a measurement [Curzi, 2026w].
4. **A null bounds a gain; it does not prove a zero.** Every null reported below is stated as an upper bound at the study's power.
5. **Grounding is established from the channel, never from the outcome.** This is the rule that keeps the postulate from collapsing into tautology and it is a requirement, not advice.
6. **The replication discount applies to this framework's own claims exactly as it applies to the claims it evaluates.** Section 5.10 computes it on the framework.

35.3.8 3.8 Bibliographic disclosure

A large number of the primary sources this paper relies on carry no key in this corpus's bibliography, and no key was invented for any of them. Named in the text with journal, volume, year and pages, and cited without a key: Thomas and Thomas (1928); Diamond and Dybvig (1983); Thorndike (1968); Elashoff and Snow (1971); Raudenbush (1984); Steele and Aronson (1995); Flore and Wicherts (2015) and Flore et al. (2018); Sisk et al. (2018); Yeager et al. (2019); Camerer et al. (2018); Bem (2011); Ritchie, Wiseman and French (2012); Benson et al. (2006); Bösch, Steinkamp and Boller (2006); Levine, Gordon and Fields (1978); Kaptchuk et al. (2010); Tajik et al. (2016); and Wald's survivorship analysis for the Statistical Research Group (1943). **Every one is a standard, locatable source; the gap is in this corpus's bibliography, and it is recorded as Limitation L-2 rather than smoothed over.** Where a keyed entry covers the same ground it is cited alongside — Merton [Merton, 1948, 1936] for the mechanism, Granovetter [Granovetter, 1978] for the threshold model, Centola et al. [Centola et al., 2018] for the tipping measurement, Rosenthal and Jacobson [Rosenthal and Jacobson, 1968, Rosenthal and Fode, 1963] for the original expectancy study, Jussim and Harber [Jussim and Harber, 2005] for the modern synthesis, and the Open Science Collaboration [Open Science Collaboration, 2015] for the replication estimate.

35.4 4 The self-fulfilling engine, and what it structurally requires

35.4.1 4.1 The mechanism, stated exactly

Merton [1948] defined a self-fulfilling prophecy as **a definition of a situation that is initially false, which evokes new behaviour, which makes the originally false definition come true.** The prophecy's apparent vindication is then taken as proof that the prophet was right about the world, when in fact the prophecy was right about itself. The construction rests on the Thomas

theorem (W. I. Thomas and D. S. Thomas, *The Child in America*, 1928; **named without a key**, per Section 3.8): *if men define situations as real, they are real in their consequences.*

The addition to the Thomas theorem is the whole of this section. The Thomas theorem says definitions have consequences. It does not say which definitions, how large the consequences are, or under what conditions the consequence loops back and validates the definition. Merton supplied a specific mechanism with a specific closure, and **the closure is what makes it an engine rather than an observation.** The companion result on unanticipated consequences is the same author's [Merton, 1936].

35.4.2 4.2 The three requirements — this framework's decomposition, badged

The decomposition below is this framework's, not Merton's, and it is an analytical proposal rather than a result. What it does is separate conditions Merton kept together, so that each can be independently checked and independently destroyed. **The labels W1, W2, W3 are used throughout to avoid collision with this paper's own contributed results R1 to R6.**

```
THE SELF-FULFILLING ENGINE REQUIRES ALL THREE:

  W1  A BELIEF        held by agents who can act, with enough conviction to
                      change what they do. Sincerity is not required; acting
                      as if is sufficient, and is often what actually happens.

  W2  A CHANNEL TO ACTION
                      a route from the holding of the belief to a change in
                      behaviour. The belief must be able to move a hand, a
                      keystroke, a vote, a withdrawal slip, a grade.

  W3  A SYSTEM STATE THAT IS ACTUALLY RESPONSIVE ALONG THAT CHANNEL
                      the world must have the property that this particular
                      class of action moves it in the direction the belief
                      asserts. THIS IS THE REQUIREMENT EVERYONE DROPS.

EFFECT MODEL (illustrative, multiplicative -- A MODELLING CHOICE, BADGED):

    E  =  B  x  K  x  R

    B  conviction-weighted fraction of agents holding the belief
    K  channel transmission: belief -> action
    R  system responsiveness: action -> asserted state
```

Multiplicativity is assumed, not derived. It is chosen because it makes the conjunction explicit — a product is zero if any factor is zero — and because it composes with the chain-gain law of Section 6.9. **An additive or a min-form model would encode a different conjunction and would give different intermediate values; nothing in this paper establishes that the multiplicative form is the right one, and Section 10 records that.** Computed, to make the conjunction concrete:

```
  full engine                           B=1.00  K=1.00  R=1.00  ->  E = 1.0000
  no belief held                        B=0.00  K=1.00  R=1.00  ->  E = 0.0000
  belief held, no channel to action     B=1.00  K=0.00  R=1.00  ->  E = 0.0000
  belief + action, system unresponsive  B=1.00  K=1.00  R=0.00  ->  E = 0.0000
  weak on all three                     B=0.50  K=0.40  R=0.30  ->  E = 0.0600
```

A product is zero if any factor is zero. That is the entire architecture, and it is the reason the postulate is conditional [Curzi, 2026ab]. "Grounded in truth" is a statement about R: the system really does have the property the belief asserts, along the route the belief can reach. **It says nothing about B, and it cannot be inferred from B at any strength.**

The last row deserves a note. $E = 0.0600$ is an ordinary configuration: a belief half the population holds, that changes behaviour in two of five cases, in a system that responds weakly. **Six per cent of the asserted effect is what most real loops deliver — not nothing, and not the story the loop tells about itself.**

35.4.3 4.3 The reflexive loop, and the objection it carries

The framework's compressed statement:

```
    T(x)  =  W( x, M(x) )

    x      the state of the world
    M(x)   the model of the world held by the agents in it
    W(.,.) the world's transition rule, which takes the state AND the model
    T(x)   the next state

A NON-REFLEXIVE system is one where W does not depend on its second argument:
    T(x) = W(x, .) = W(x)          the model is a passenger.

A REFLEXIVE system is one where it does:
    T(x) = W(x, M(x))              the model is a term in the physics.
```

The whole question is: for which systems, along which channels, and with what gain does W actually depend on M? That is a question about $\partial W/\partial M$, it is finite, and it is zero far more often than the popular literature admits.

And it is not identified from observational data. Linearised, this is a linear-in-means reflexive specification and its endogenous effect cannot be separated from contextual effects or correlated unobservables by any quantity of observation [Manski, 1993]. **Everything this section computes from $\partial W/\partial M$ is therefore a modelling position, not an identified empirical result** (Section 3.3). The self-fulfilling case is the fixed-point case, in which M asserts a state x^* and W moves x toward it because of M; the self-defeating case moves x away. **Both are governed by the same derivative, and a model asserting something about a variable W does not read has $\partial W/\partial M = 0$ and is inert no matter how loudly it is asserted.**

35.4.4 4.4 Loop gain

Linearise about an operating point. Let e be a seed displacement and g the loop gain: the fraction of a displacement that returns as further displacement after one pass around belief, action, world and observation.

```
    total displacement  =  e * (1 + g + g^2 + g^3 + ...)  =  e / (1 - g)   for |g| < 1

    g       multiplier 1/(1-g)      amplification of the seed
    0.0000        1.0000                   0.00%
    0.1000        1.1111                  11.11%
    0.2000        1.2500                  25.00%
    0.3333        1.5000                  50.00%
    0.5000        2.0000                 100.00%
    0.8000        5.0000                 400.00%
    0.9000       10.0000                 900.00%
```

```
   0.9500       20.0000                1900.00%
   0.9900      100.0000                9900.00%
   1.0000      divergent -- marginal; the loop becomes a pure integrator
  >1.0000      divergent -- runaway until a nonlinearity binds
```

Below unity the loop is an amplifier: it multiplies the seed by a finite factor and settles. At and above unity it is not an amplifier but an engine, and it runs until something physical stops it. Gain 0.5 doubles the seed; gain 0.9 multiplies it tenfold; gain 0.99 a hundredfold. **The last increment before unity does most of the work**, which is why systems near unity feel stable right up until they do not. The geometric multiplier is a **shared mathematical form** with electrical feedback (Section 3.2) and transfers no electrical result.

Above unity the relevant question is timing:

```
   steps to double = ln 2 / ln g          steps to 10x = ln 10 / ln g

   g       double      10x
   1.02    35.003     116.277
   1.05    14.207      47.194
   1.10     7.273      24.159
   1.15     4.959      16.475
   1.25     3.106      10.319
   1.50     1.710       5.679
   2.00     1.000       3.322
```

A loop at gain 1.10 doubles in about seven passes. If a pass is a day, that is a week; if a pass is a minute — and with modern settlement and messaging a pass can be a minute — it is seven minutes. **The gain sets whether it happens; the pass time sets whether anyone can respond.**

35.4.5 4.5 Warrant

```
PROVED (mathematics)
  the geometric-series multiplier 1/(1-g) for |g|<1
  the stability condition for the linearised reflexive storage law (Section 4.11)
  the threshold fixed-point condition F(x) = x (Granovetter's model, Section 4.9)
  the ceiling 0 <= F(u) <= ln N (CCP-A02)

MEASURED (the evidence column)
  bank runs occur, are documented, and follow the fractional-reserve structure
  the Diamond-Dybvig multiple-equilibrium model and its empirical standing
  threshold and cascade behaviour in collective action (Granovetter 1978)
  convention tipping near a 25% committed minority (Centola et al. 2018)
  teacher-expectancy effects: real, modest, contested history (Section 4.14)

POSTULATED, BADGED (this framework's)
  the W1/W2/W3 decomposition of Merton's mechanism
  the multiplicative effect model E = B x K x R
  the identification of "grounded in truth" with W3
  the reading of Merton's mechanism as an independent arrival at the
    framework's conditional clause

ANALOGY ONLY, BADGED
  reflexivity as a practitioner's formalism (Section 4.13). It is a framework
  advanced by a market participant, not a tested theory, and this paper uses
  it for its vocabulary and not for its warrant.
```

35.4.6 4.6 The concordance, with column roles stated

```
CONCORDANCE 4.1 -- A DEFINITION OF A SITUATION CAN MAKE ITSELF TRUE

  CLAIM        A belief about a system can become true through the actions it
               causes, when the system is responsive to those actions.

  (+1) EVIDENCE
               Merton (1948), building on the Thomas theorem (1928).
               Diamond & Dybvig (1983) for the formal multiple-equilibrium
               structure of the deposit contract. Documented bank runs, and
               the sharp fall in US bank runs after federal deposit insurance
               in 1934.
               WHAT IT MEASURED: that runs occur, that they follow the
               liquidity structure of the contract, and that changing the
               structure changes the incidence.

  (-1) PRECEDENT -- CARRIES NO EVIDENCE
               Sophocles, Oedipus Tyrannus (c. 429 BCE): the prophecy fulfilled
               through the very actions taken to avoid it. Herodotus, Histories
               I: Croesus and the oracle. The Talmudic tractate Berakhot records
               the maxim that a dream follows the interpretation given to it.
               WHAT IT RECORDS: that the structure was noticed, named and
               dramatised repeatedly across unrelated traditions, by people
               with no instruments and no statistics.

  (0p) INTEGRATION
               The traditions recorded the phenomenon and attributed it to the
               prophecy. Merton recorded the phenomenon and attributed it to
               the behaviour the prophecy caused. The modern account is not a
               vindication of the old one; it is a RELOCATION of the mechanism
               from the oracle to the agents, and that relocation is the entire
               scientific content.

  CRITERIA     C1 structural correspondence: PASS.
               C2 predictive power: PASS. The account predicts that removing
                  responsiveness (W3) kills the effect -- and deposit insurance
                  is the natural experiment.
               C3 practical utility: PASS on both sides.

  STATUS       ESTABLISHED as a mechanism. The (-1) column carries precedent
               only, and no part of the (+1) case rests on it.
```

Sophocles is not a second measurement of the bank-run phenomenon. What the precedent column establishes is that a loop of this shape was visible to careful observers two and a half thousand years ago, which makes it likelier that the loop is a feature of human systems rather than an artefact of twentieth-century sociological method. **That is a real argument, it is smaller than a replication, and this paper does not inflate it** [Curzi, 2026w].

35.4.7 4.7 The deposit contract, with its balance sheet on the table

Merton's own illustration is a fictional institution: a rumour spreads, depositors queue, the bank — which was solvent — cannot meet simultaneous demand, and fails. **Almost every retelling stops there and draws the wrong lesson.** The lesson drawn is *belief made it true*.

The lesson available is: belief made it true because the deposit contract promises something the balance sheet cannot deliver on demand.

```
A BANK, SIMPLIFIED

  deposits            D          payable in full, on demand, at par
  liquid reserves     r * D      cash and overnight instruments
  illiquid assets     A          loans, whose value at maturity exceeds their
                                 value if sold this afternoon

  THE CONTRACT PROMISES:  every depositor may have 100 cents on the dollar today.
  THE BALANCE SHEET HOLDS: r * D of today-money, and A of later-money.
```

Computed across reserve ratios, for a thousand depositors holding one thousand each:

```
r       liquid        depositors payable in full     money multiplier 1/r
0.03      30,000               30 of 1000                  33.3333
0.05      50,000               50 of 1000                  20.0000
0.10     100,000              100 of 1000                  10.0000
0.20     200,000              200 of 1000                   5.0000
0.50     500,000              500 of 1000                   2.0000
1.00   1,000,000             1000 of 1000                   1.0000
```

The run threshold is the reserve ratio. If more than a fraction r of deposits is demanded at once the promise fails — not because the institution is insolvent, not because the assets are bad, but because *today-money* and *later-money* are different goods and the contract pretends they are the same. **W3 is satisfied by arithmetic, before any depositor has an opinion about anything.**

35.4.8 4.8 The counterfactual that isolates W3

Run the identical scenario against an institution holding one hundred per cent reserves. Same town, same rumour, same panic, same queue, same conviction, same channel from belief to withdrawal slip.

```
100% RESERVE, D = 18,000,000, liquid = 18,000,000

  f = 0.05   demand =    900,000    covered? YES   shortfall = 0
  f = 0.25   demand =  4,500,000    covered? YES   shortfall = 0
  f = 0.50   demand =  9,000,000    covered? YES   shortfall = 0
  f = 1.00   demand = 18,000,000    covered? YES   shortfall = 0

  run threshold f* = 1.0000
```

W1 unchanged. W2 unchanged. W3 removed. The engine is dead. Every depositor may believe the institution is failing, act on the belief with maximum vigour, persuade every neighbour — and it pays out in full and opens the next morning. **The belief was not weaker and the channel was not narrower. The system was not responsive along that channel, and so nothing happened.**

This is a thought experiment, but the narrow-bank design is a real and long-discussed proposal in monetary economics precisely because it has this property, and **there is a real intervention with the same target: deposit insurance does not argue with the belief — it changes W3, by making the depositor's payout independent of the queue.** US bank runs fell sharply and durably after 1934. The engineering lesson generalises far beyond banking:

W1 cannot be reliably intervened on. W2 can sometimes be. W3 can very often be, and W3 is where the leverage is.

Arguing with a panic is an attempt to move B, against a population, under time pressure, with the panic's own dynamics working against it. Changing the payoff structure is an attempt to move R, once, in a document. **Structure beats persuasion, because structure does not have to be re-persuaded every morning** — the same conclusion the corpus reaches for built-environment levers, where a structural change to a room is priced against the surplus law and persuasion is not [Curzi, 2026e].

35.4.9 4.9 Two equilibria, and the threshold structure

Diamond and Dybvig's 1983 model (*Journal of Political Economy*; **named without a key**) makes the structure precise. A demand-deposit contract with a sequential-service constraint admits **two equilibria**: one in which only those with genuine liquidity needs withdraw and the institution is fine, and one in which everyone withdraws because everyone expects everyone to withdraw and it fails. **Both are self-consistent. Both are self-fulfilling. The belief does not create the equilibrium; it selects which of two pre-existing equilibria the system lands in.** That is a much more precise and much more modest claim than "belief creates reality", and it bounds the belief's power exactly: **belief can select among the states the structure already admits, and it cannot manufacture a state the structure does not admit.**

Whether the seed propagates is a threshold problem, and Granovetter [1978] gives it exactly. Each member has a threshold: the fraction of others withdrawing at which she withdraws too. With $F(x)$ the fraction of members whose threshold is at or below x:

```
the cascade advances from x while  F(x) > x
it halts at the first fixed point   F(x) = x
it tips when the curve F is TANGENT to the 45-degree line and then crosses
```

Computed for a 2,400-member deposit base with thresholds Normal(μ, σ) and a 9 per cent reserve ratio:

```
thresholds ~ N(0.50, 0.10)   unconditional withdrawers F(0) =  0.00%   x* =  0.0000   no run
thresholds ~ N(0.50, 0.20)   F(0) = 0.62%                              x* =  0.0068   no run
thresholds ~ N(0.40, 0.20)   F(0) = 2.28%                              x* =  0.0334   no run
thresholds ~ N(0.35, 0.25)   F(0) = 8.08%                              x* =  0.9951   RUN
thresholds ~ N(0.30, 0.20)   F(0) = 6.68%                              x* =  0.9998   RUN
```

Holding $\sigma = 0.25$ and varying the mean tolerance:

```
critical mean tolerance  mu* = 0.408522

  mu = 0.407522   ->   x* = 0.990106      99.0% of deposits withdrawn
  mu = 0.409522   ->   x* = 0.145087      14.5% of deposits withdrawn
```

A shift of two-thousandths in the population's mean tolerance moves the outcome from a ninety-nine per cent run to a fourteen per cent tremor, with no intermediate outcome. The system is not sensitive to belief in a smooth, proportionate way; it is sensitive in a knife-edged way. **This is a property of the threshold model at the stated parameters and is not a measurement of any institution.**

Three consequences follow, and they are all uncomfortable.

1. **Post-hoc explanation of a run is nearly worthless.** Two runs, one of which happened and one of which did not, may differ by an unmeasurable amount in the underlying parameter. The narrative afterward will be confident, detailed and unfalsifiable.

2. **Prediction requires measuring F, not the belief.** The belief is the seed; the threshold distribution is the amplifier. Surveys of sentiment measure the wrong object. The same structure appears wherever thresholds govern collective switching [Schelling, 1971, Watts, 2002, Centola and Macy, 2007, Bikhchandani et al., 1992, Banerjee, 1992].
3. **The movable parameter is dispersion, not mean.** σ is the heterogeneity of tolerance, and institutions reduce it by making the same information available to everyone at the same time.

There is a fourth consequence and it belongs to the present decade. **Pass time has collapsed.** A 2023 deposit-institution failure involved withdrawal requests reported at roughly forty-two billion dollars in a single day — a queue moving at the speed of a mobile application rather than the speed of a pavement. **The loop gain did not change; the pass time did, and the pass time determines whether anyone can intervene between passes.** Substitute minutes for days in Section 4.4's doubling table.

35.4.10 4.10 The loop in the framework's own notation

The storage law [Curzi, 2026a]:

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t
```

In the non-reflexive case S_t is exogenous. **A reflexive system is one in which the input depends on the store, because agents observe Q, form a belief about it, and act in a way that changes what they contribute.** Linearise:

```
S(Q) = S_0 + s ∗ Q          s = d S / d Q, the REFLEXIVE SLOPE

Q_{t+1} = [ (1 − delta) + eta∗s ] Q_t + (eta∗S_0 − C)

POLE          a = (1 − delta) + eta∗s
STABLE iff    eta ∗ s < delta
LOOP GAIN     g := eta ∗ s / delta
```

That is the self-fulfilling engine written in the framework's own variables, and it was not imported: it falls out of adding one term to an equation the framework already had. Computed with $\delta = 0.10$ and $\eta = 0.50$, so the critical slope is $s^* = \delta/\eta = 0.20$:

s	eta∗s	pole a	g=eta∗s/delta	regime	tau = 1/(1−a)
0.00	0.000	0.9000	0.0000	convergent	10.0000
0.05	0.025	0.9250	0.2500	convergent	13.3333
0.10	0.050	0.9500	0.5000	convergent	20.0000
0.15	0.075	0.9750	0.7500	convergent	40.0000
0.18	0.090	0.9900	0.9000	convergent	100.0000
0.20	0.100	1.0000	1.0000	marginal (integrator)	infinite
0.22	0.110	1.0100	1.1000	RUNAWAY	–
0.30	0.150	1.0500	1.5000	RUNAWAY	–
0.40	0.200	1.1000	2.0000	RUNAWAY	–

Read the τ column: long before the loop goes unstable, it goes slow. At $g = 0.90$ the time constant is a hundred periods instead of ten. An institution with a strong but sub-critical reflexive loop does not look like one about to explode; it looks sluggish, unresponsive to correction, its reforms taking a decade to show. **Critical slowing before instability is a measurable signature available from response time alone, without knowing s** — and, since s is not identified from observation (Section 3.3), without knowing s is the only way it is ever going to be available.

The steady state carries the same information:

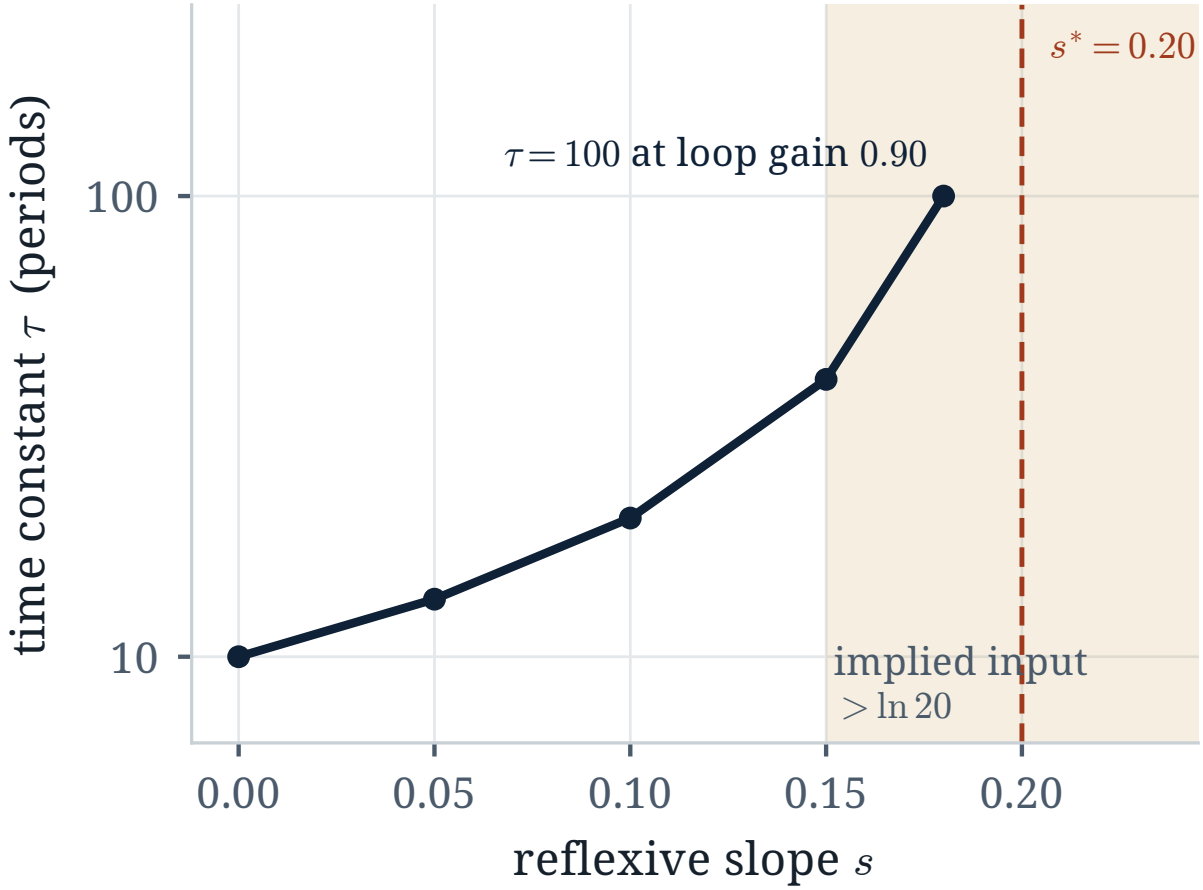


Figure 35.1: **Long before the reflexive loop destabilises, it slows down:** the time constant $\tau = 1/(1-a)$ against the reflexive slope s, with $a = (1-\delta) + \eta s$ at $\delta = 0.10$ and $\eta = 0.50$, so that the critical slope is $s^* = 0.20$. The shaded region marks slopes whose implied steady-state input exceeds $\ln N = 2.995732$ at $N = 20$ — **every runaway row leaves the model's admissible region before the pole reaches unity.** Modelled, not measured: δ and η are the assumed values of Section 4.10 and the reflexive slope is not identified from observational data. Generated by `figs/src/CCP-035_fig01.py`.

```
Q_inf = (eta*S_0 - C) / (delta - eta*s)

s      g       denominator     Q_inf       S(Q_inf)
0.000  0.000     0.1000          3.0000     1.2000 nats
0.050  0.250     0.0750          4.0000     1.4000
0.100  0.500     0.0500          6.0000     1.8000
0.150  0.750     0.0250         12.0000     3.0000
0.180  0.900     0.0100         30.0000     6.6000
0.190  0.950     0.0050         60.0000    12.6000
0.199  0.995     0.0005        600.0000   120.6000
```

The steady state diverges as $s \to \delta/\eta$, and that divergence IS the self-fulfilling prophecy expressed in the storage law: a belief that raises the input in proportion to the store produces unbounded store. And it is precisely here that the formalism protects itself.

35.4.11 4.11 The ceiling that terminates every runaway

S_t is bounded: $0 \leq F(u) \leq \ln N$ [Curzi, 2026]. The surplus available from any single period of interaction cannot exceed the log of the number of distinguishable kinds of thing present. Read the right-hand column of the table above: at $s = 0.19$ the implied input is 12.6 nats and at $s = 0.199$ it is 120.6 nats.

```
N =   4   ln N = 1.386294
N =   8   ln N = 2.079442
N =  12   ln N = 2.484907
N =  20   ln N = 2.995732
N =  50   ln N = 3.912023
N = 150   ln N = 5.010635
```

Every row of the divergence table beyond $s = 0.15$ is outside the model's admissible region for any realistic N. The runaway does not run away. It runs into the ceiling. With $\delta = 0.10$, $\eta = 0.50$ and $C = 0$, the uniform bound gives $Q_{\max} = \eta \ln N/\delta$:

```
N =   4   ->   6.9315
N =  12   ->  12.4245
N =  50   ->  19.5601
N = 150   ->  25.0532
```

This is the formal counterpart of the fact that a run ends. **It ends because the institution runs out of money, not because the depositors change their minds.** In every reflexive system the runaway is terminated by a nonlinearity, and **identifying that nonlinearity in advance is the single most useful piece of analysis available.** In the deposit case it is the balance sheet; in the surplus case it is $\ln N$; in a reputation cascade it is the finite population. **A model that predicts unbounded growth has simply not yet located its ceiling, and the ceiling is always there.**

35.4.12 4.12 Reflexivity as a practitioner's framework — badged

Soros [2013] advances **reflexivity** as an account of markets: participants hold a cognitive function mapping the world to their understanding of it, and a participating function mapping their understanding back onto the world, so that neither the world nor the understanding is ever a fixed independent variable.

This paper uses that framework for vocabulary and not for warrant, and the badge stays on. It is a framework advanced by a market participant on the basis of his own practice. It is not a tested theory, it carries no effect size, and it has not passed a falsification suite. What it supplies is the clean statement that the two functions compose into $T(x) = W(x, M(x))$, and the observation — correct and useful — that the composition has a fixed point only when the two functions are consistent, so that markets spend most of their time away from equilibrium rather than at it. Related formal treatments of the same phenomenon in markets are available [Kirman, 1992, Farmer and Foley, 2009, Bouchaud, 2013, Sornette, 2003], and the performativity literature supplies the field observations [MacKenzie and Millo, 2003, Ferraro et al., 2005].

Where this is genuinely additive to Merton is the two-way structure. Merton's mechanism is one-directional: the belief changes the world. The reflexive account insists the changed world also changes the belief, which changes the world again. **That is the loop, and it is why the loop has a gain rather than a multiplier.** Section 4.10 is the version this framework is entitled to assert, because it is derived from an equation the framework proved.

35.4.13 4.13 Expectancy effects, reported honestly

The study. Rosenthal and Jacobson [1968] reported that children designated at random as likely to show unusual intellectual growth subsequently gained more than their classmates, and interpreted the gain as produced by teacher expectation [see also Rosenthal and Fode, 1963].

The criticism, which was serious and sustained. The study drew immediate and substantive methodological objections concerning the ability instrument used, whose scores were unreliable at the extremes of the range and which had been administered outside its intended grade range for some children, producing implausible individual scores; the very small numbers of children in the subgroups where the effect was concentrated; and the fact that the effect was not distributed across grades as the theory would predict. R. L. Thorndike's 1968 review in the *American Educational Research Journal* was among the most damaging, and Elashoff and Snow's *Pygmalion Reconsidered* (1971) reanalysed the data and reached substantially more conservative conclusions. **Both are named without a key**, per Section 3.8.

What survived. The phenomenon did not disappear under criticism; it shrank and acquired conditions. Raudenbush's 1984 meta-analysis of teacher-expectancy induction experiments found the effect **strongly moderated by prior contact**: where teachers had had little or no acquaintance with the students before the expectation was induced, effects were found; where teachers had known the students for two weeks or more, effects were essentially absent. Jussim and Harber [2005] — the standard modern synthesis — concluded that classroom self-fulfilling prophecies are **real, but typically small, often fragile, and more likely to dissipate over time than to accumulate**, with somewhat larger effects for students from stigmatised groups. Demand characteristics, often invoked loosely in this literature, have their own contrary review [McCambridge et al., 2014].

Now map that onto W1, W2, W3, because the mapping is the payoff.

```
W1  BELIEF     the teacher's expectation about a particular child.
W2  CHANNEL    differential teacher behaviour: time allocated, warmth,
               difficulty of material offered, specificity of feedback,
               opportunities to respond. These are OBSERVABLE and have been
               observed. The channel is real and it is physical.
W3  RESPONSE   the child's actual malleability along that channel.
```

The prior-contact moderator is a direct measurement of W1 failing to install. A teacher who has taught a child for three weeks already has data, and an experimenter's claim about a test does not overwrite it. The belief cannot be induced, so B is near zero, so the product is near zero. **It is not a mysterious moderator. It is the first requirement, measured.**

And the dissipation result is a measurement of R being small. If R were large the effect would compound — a child given more and harder work becomes more capable, which raises the expectation, which raises B, which raises the input again. That is a positive loop with $g > 0$ and it would show up as accumulation. **The literature finds dissipation instead, which says the loop gain is well below unity: the effect is an amplifier with a modest multiplier, not an engine.** That is exactly Section 4.4's small-g row, and exactly the distinction the popular literature erases.

```
CONCORDANCE 4.2 -- EXPECTATION HELD BY ONE PERSON CAN MOVE ANOTHER'S PERFORMANCE

  CLAIM        Expectations held about a person can measurably change that
               person's measured performance, through observable differences in
               how the expectation-holder treats them.

  (+1) EVIDENCE
               Rosenthal & Jacobson 1968 (original, heavily criticised:
               Thorndike 1968; Elashoff & Snow 1971). Raudenbush 1984
```

```
              (meta-analysis; prior-contact moderator). Jussim & Harber 2005
              (review; effects real, typically small, often dissipating).
              WHAT IT MEASURED: that induced expectations shift measured
              outcomes by a modest amount under specific conditions, and not
              at all under others.

(-1) PRECEDENT -- CARRIES NO EVIDENCE
              Apprenticeship traditions across many crafts record the master's
              regard for the apprentice as an operative part of the training
              rather than a courtesy. These are characterisations of what the
              traditions recorded, offered as such and not as quotations.

(0p) INTEGRATION
              The traditions located the effect in the master's regard; the
              modern account locates it in the master's BEHAVIOUR, which the
              regard causes. The channel is differential treatment, it is
              observable, and it is the thing to instrument.

CRITERIA      C1 PASS. C2 PASS (predicts prior-contact moderation, which is
              observed). C3 PASS.

STATUS        ESTABLISHED as a real effect. MAGNITUDE: modest, condition-
              dependent, and easily overstated.
```

35.4.14 4.14 The mirror: self-defeating prophecy

Merton also named the opposite case, and it is the same machine with the sign reversed: a forecast that provokes the behaviour refuting it — the warned-of epidemic that does not arrive because the warning changed behaviour, the predicted collapse averted by the preparation the prediction caused. In the notation of Section 4.10 this is $s < 0$:

s	pole a	g	tau
−0.05	0.8750	−0.2500	8.0000
−0.10	0.8500	−0.5000	6.6667
−0.20	0.8000	−1.0000	5.0000
−0.50	0.6500	−2.5000	2.8571

Negative reflexive slope makes the system converge faster, not slower. A forecast that provokes its own refutation is a stabiliser: it shortens the time constant, damps excursions, and makes the institution more responsive to correction.

> **A forecasting institution whose forecasts are systematically wrong in the self-defeating direction may be functioning perfectly. Its accuracy is being consumed by its usefulness.**

Public-health forecasting, safety engineering and financial supervision all live here. A regulator whose crisis warnings are all wrong may be a bad regulator or an excellent one, and **the distinction cannot be made from the forecast record alone.** It requires estimating s — which is W2 and W3, and which is **not identified from observational data** [Manski, 1993]. **That is the honest position: the distinction requires a quantity observation cannot supply, and the only route to it is an intervention.** Goodhart's law is the same phenomenon: a measure that becomes a target ceases to be a measure, because targeting it gives it a nonzero s [Goodhart, 1975, Campbell, 1979]. The corpus's own prediction claims inherit the same limit, as the measured social-prediction literature reports independently [Hofman et al., 2017, Watts, 2014, Lazer et al., 2014].

35.4.15 4.15 The convergence, stated at its exact strength

Merton, in 1948, with no formalism, wrote a mechanism whose operative content is: **a false belief becomes true when, and only when, it drives behaviour that a responsive system converts into the asserted state.** The framework's postulate states: **belief creates reality if the belief is grounded in truth** [Curzi, 2026ab], with "grounded in truth" unpacked as: the channel exists and is real.

These are the same statement. Merton's "evokes new behaviour" is W2; his implicit requirement — carried entirely by his choice of a *bank* as the example — is W3; his "definition of the situation" is W1.

And here is what that is worth, at exactly its strength and no more. It is **not** worth a doubling of the evidence. Merton is in the evidence column and the postulate is this framework's own; **two statements of the same idea from the same column are one idea, not two verifications** [Curzi, 2026w]. **What the convergence is worth is that the structure was reachable from two unrelated starting points, seventy-eight years apart, by people with different questions and incompatible methods** — one from mid-century American sociology and the empirical record of bank failures, the other from an interaction-geometry functional and a leaky integrator. That is structural correspondence, satisfied about as well as it can be.

And it makes a checkable prediction: if W3 is what "grounded in truth" means, then interventions on W3 should kill self-fulfilling loops even when belief is unchanged. **Deposit insurance is precisely that experiment, run at national scale, and it worked. This paper claims that as support for the reading, not as proof of the postulate, and the difference matters.**

35.4.16 4.16 Where the engine stops

Everything above concerns systems where W3 is nonzero: banks, classrooms, markets, reputations, teams. **Every one of them is a system in which human action is the transmission medium and human behaviour is the responsive substrate.** That is not an accident and it is not a limitation this section can talk its way around. **Every established self-fulfilling loop in the literature runs through human action. There is no documented case of a belief closing a loop through a system that does not contain a person acting on the belief.** The engine is real and powerful, and it is powerful over exactly one class of substrate. Section 5 is the catalogue of what happens outside that class.

35.4.17 4.17 Requirements

REQ-4.1 Never assert a self-fulfilling loop without naming W1, W2 and W3 separately. A claim naming only the belief is a claim about one factor of a three-factor product.

REQ-4.2 Name the structural feature supplying W3 in the same paragraph as the claim. If it cannot be named, the loop has not been identified, and saying so is the contribution.

REQ-4.3 State the loop gain g, or state that it is unmeasured — **and state that it is not identified from observational data**. "There is a feedback loop" is not a finding.

REQ-4.4 Report the pass time alongside the gain. A gain of 1.10 doubles in 7.273 passes; whether that is a week or seven minutes is the operational question.

REQ-4.5 Locate the terminating nonlinearity before the runaway starts. A projection that goes to infinity has not been finished.

REQ-4.6 Check the input against its ceiling. Any reflexive model whose implied S_t exceeds $\ln N$ has left its admissible region and its predictions beyond that point are not entitled to be reported [Curzi, 2026].

REQ-4.7 Intervene on W3 first, W2 second, W1 last.

REQ-4.8 Measure the threshold distribution, not the sentiment. The outcome swings from 99.0 to 14.5 per cent on a change of 0.002 in the mean of F, and a survey of belief cannot see that.

REQ-4.9 Track the time constant as an early-warning signal. $\tau = 1/(1-a)$ rises steeply before a reaches 1, and this is measurable without measuring s.

REQ-4.10 When reporting expectancy effects, report the magnitude and the moderators, never the direction alone.

REQ-4.11 Cite the original expectancy study only together with its criticism and the modern estimate. A citation of the original alone misrepresents the state of the evidence.

REQ-4.12 Distinguish selection among equilibria from creation of an equilibrium.

REQ-4.13 When a forecast is self-defeating, do not score it on accuracy alone — and state that s is not recoverable from the forecast record.

REQ-4.14 Keep the concordance columns separate, every time [Curzi, 2026w].

35.4.18 4.18 Failure modes

FM-4.1 — THE AMPUTATED MECHANISM. *Signature:* the run story is told and the fractional-reserve structure is never mentioned. *Mechanism:* W3 has been deleted, leaving "belief made it true" with nothing under it, and the example is then generalised to systems with no analogue of fractional reserve. *Fix:* every retelling names the balance sheet.

FM-4.2 — GAIN AGNOSTICISM. *Signature:* "there's a feedback loop here", with no number. *Mechanism:* topology identified and nothing else. *Fix:* estimate g from any two consecutive passes and name a side of unity — **while stating that the estimate is a modelling position, not an identified parameter.**

FM-4.3 — INFINITE PROJECTION. *Signature:* a reflexive model run forward without bound, reported as a finding. *Fix:* print the implied S_t next to the projection and truncate at $\ln N$.

FM-4.4 — BELIEF-SIDE INTERVENTION UNDER TIME PRESSURE. *Signature:* a cascade is under way and the entire response is communication. *Mechanism:* W1 is being attacked, in a population, against the loop's own gain, with a pass time shorter than the message cycle. *Fix:* W3.

FM-4.5 — SENTIMENT SUBSTITUTION. *Signature:* cascade risk assessed by surveying belief. *Fix:* measure behaviour at thresholds — what fraction have already acted, at what observed level of others' action.

FM-4.6 — SLUGGISHNESS MISDIAGNOSED. *Signature:* reforms take longer and longer; the diagnosis is "culture". *Fix:* measure the response time to a known small deliberate perturbation and compare across periods.

FM-4.7 — EXPECTANCY INFLATION. *Signature:* the 1968 study cited as an established large effect with no mention of the criticism or the modern estimate. *Fix:* REQ-4.11.

FM-4.8 — MODERATOR MYSTIFICATION. *Signature:* a moderator reported as a curious wrinkle and treated as unexplained. *Mechanism:* the moderator is one of the three requirements, measured. *Fix:* check whether it is W1, W2 or W3 under another name. In this literature it usually is.

FM-4.9 — REFLEXIVITY WITHOUT THE BADGE. *Signature:* a practitioner's framework invoked as a validated theory. *Fix:* use the vocabulary, cite the source, keep the badge, derive quantitative claims from Section 4.10.

FM-4.10 — ACCURACY SCORING OF A SELF-DEFEATING FORECAST. *Signature:* a warning system judged solely on how often its warnings came true. *Fix:* instrument the behavioural response, and concede that the decisive quantity is not identified from observation.

FM-4.11 — SUBSTRATE OVERREACH. *Signature:* a loop correctly identified in a social system asserted for a system with no human actor in it. *Fix:* name the actor. If none can be named between the belief and the outcome, there is no loop, and Section 5 applies.

35.5 5 Where belief fails: the four classes

A postulate that names precisely where it fails can be trusted precisely where it holds. This section is the catalogue: what belief does not do, what the evidence for the things it does do looks like when the effect sizes are printed, and what analytic error is being made in each case. **Its second job is harder: the replication crisis is the evaluative context for every claim in this area, including this framework's own, and Section 5.10 applies the same discount to the framework with arithmetic. The result is not flattering. It is not supposed to be.**

35.5.1 5.1 The four failure classes

Failure is not a single phenomenon. It is four distinct faults, and **the entire diagnostic value of this section comes from insisting they be told apart.**

```
CLASS A -- NO CHANNEL.
  There is no physical route from the belief to the asserted outcome. Gain is
  not small; it is undefined, because there is nothing to have a gain.
  Diagnostic: name the transducer. If none can be named, Class A.

CLASS B -- REAL CHANNEL, INSUFFICIENT GAIN.
  The route exists and is physical, and the transfer is far too small to produce
  the asserted effect. The claim is not false in kind; it is false in size.
  Diagnostic: the measured effect is real, replicable, and an order of magnitude
  below the claim.

CLASS C -- WRONG MAGNITUDE FROM A REAL RESULT.
  A genuine laboratory effect is extrapolated to a real-world magnitude the
  effect cannot support. The original finding may be sound. The transfer is not.
  Diagnostic: the claim's size traces to no measurement; it traces to an
  inference from the existence of the effect.

CLASS D -- SYSTEM NOT RESPONSIVE ALONG THAT CHANNEL.
  The channel exists and carries signal, but the target system does not move in
  response to what that channel delivers -- or does not move for an input of this
  magnitude, from a source of this size.
  Diagnostic: the channel is fine; the receiver is not listening, or is listening
  to a chorus in which this voice is 1/N.
```

The four are not degrees of the same failure; they call for different responses. Class A means stop. Class B means the effect is real and should be reported at its real size. Class C means the laboratory result stands and the extrapolation must be withdrawn. Class D means the intervention is aimed at the wrong system, or from the wrong scale, and may work at a different one.

35.5.2 5.2 The falsifier, and the circularity trap

The postulate is a conditional: *if grounded in truth, then belief creates reality.* Its logical falsifier is a belief that creates reality **without** being grounded — an ungrounded belief with a real effect, an effect with no channel. That is a clean falsification condition, and it conceals a trap.

> **THE CIRCULARITY TRAP. A conditional is vacuously true whenever its antecedent never holds. If "grounded in truth" is diagnosed FROM the outcome — the belief worked, therefore it was grounded; the belief failed, therefore it was not — then**

> **the postulate predicts nothing, forbids nothing, and is unfalsifiable. It becomes a relabelling of whatever happened.**

This is the criticism that would sink the postulate, and it is correct as stated. It has exactly one answer, and the answer is a requirement rather than an argument:

```
GROUNDING MUST BE ESTABLISHED FROM THE CHANNEL, NEVER FROM THE OUTCOME.

The channel must be named, its physical possibility checked, and its gain
estimated, BEFORE the outcome is observed. If that is done, the postulate
makes a magnitude prediction and can be wrong. If it is not done, the
postulate is a tautology and this framework disowns it.
```

Section 6 is that procedure. **The relation between the two sections is therefore not decorative: Section 5 states the objection that makes the postulate worthless, and Section 6 is the only thing that answers it.** The demand that a theory forbid something is not new [Popper, 1959, Platt, 1964, Lakatos, 1970], and the specific reason it bites here is that directional predictions are nearly costless in this domain [Meehl, 1967, 1990].

35.5.3 5.3 Warrant

```
PROVED (arithmetic)
  the survivorship result of Section 5.12: a testimonial corpus reporting
  successes only carries a likelihood ratio of exactly 1.000000 between any
  two success rates, because the denominator was never observed
  the power and minimum-detectable-effect computations throughout
  the Casimir pressures of Section 5.5
  the magnetic-field energy budget of Section 5.4

MEASURED (the evidence column) -- reported at magnitude
  Open Science Collaboration 2015: 97/100 originals significant, about 36/100
    replications, effect sizes roughly halved
  Camerer et al. 2018: 13 of 21 social-science replications significant in the
    same direction, effects about half
  Sisk et al. 2018: growth-mindset effects very small
  Yeager et al. 2019: about a tenth of a grade point, in a subgroup
  Bosch, Steinkamp & Boller 2006: micro-PK meta-analysis, tiny heterogeneous
    effect with strong evidence of bias
  Benson et al. 2006: no benefit found
  Ritchie, Wiseman & French 2012: failed replications

POSTULATED, BADGED (this framework's)
  the four-class taxonomy A/B/C/D
  the claim that every documented failure falls into one of the four
  the requirement that grounding be established from the channel

NOT CLAIMED ANYWHERE
  that any null result proves a channel's absence. A null bounds a gain.
  Section 5.11 computes the bounds and they are not small.
```

35.5.4 5.4 Class A: belief does not move objects

No physical mechanism is known by which the holding of a proposition exerts force on matter outside the body of the holder. The interesting part is why the usual objection to this is wrong. The usual objection is energetic — *the brain has no energy to spare* — and it is false:

```
energy to raise a 100 g object by 1 cm:
    m g h = 0.100 * 9.80665 * 0.01 = 9.806650e-03 J = 9.8066 mJ

delivered in one second, that is 9.8066 mW.

a resting human brain runs at roughly 20 W.

the brain's power budget exceeds the required power by a factor of 2.0394e+03.
```

The brain has two thousand times the power needed. Energy is not the obstacle. The obstacle is that there is no transducer. Energy the body possesses is coupled to the world through muscles, voice and heat, and through nothing else. To move an object without touching it, some field would have to carry the coupling:

```
magnetic energy density  u = B^2 / (2 mu_0),  mu_0 = 4 pi x 10^-7

B = 1.000e-13 T   brain field at the scalp, magnetoencephalography scale
    u = 3.978874e-21 J/m^3     in 1 cm^3:  3.978874e-27 J

B = 5.000e-05 T   Earth's field
    u = 9.947184e-04 J/m^3     in 1 cm^3:  9.947184e-10 J

B = 1.000e+00 T   a strong permanent magnet at its pole face
    u = 3.978874e+05 J/m^3     in 1 cm^3:  3.978874e-01 J

ratio of the required lift energy to the total brain-field energy in 1 cm^3:
    9.806650e-03 / 3.978874e-27  =  2.4647e+24
```

The magnetic field a brain produces is about 5.000×10^{8} times weaker than the Earth's field, and the Earth's field does not lift objects. The total field energy in a cubic centimetre at the scalp is short of the required work by a factor of roughly 2.46×10^{24}. **This is not a close call and not a matter of instrument sensitivity. It is Class A: no transducer, no channel, gain undefined.**

35.5.5 5.5 The Casimir clause, which is mandatory and instructive

The Casimir effect is real, measured and rigorous [Casimir, 1948, Lamoreaux, 1997, Mohideen and Roy, 1998]. **It is also the best available illustration of a Class C failure — a genuine physical effect recruited to carry a magnitude it cannot bear.** Recomputed here, as the corpus requires wherever it appears [Curzi, 2026aa]:

```
F/A = pi^2 hbar c / (240 a^4)

  a =   10 nm   |F/A| = 1.300126e+05 Pa  = 1.283124e+00 atm
  a =  100 nm   |F/A| = 1.300126e+01 Pa  = 1.283124e-04 atm
  a =    1 um   |F/A| = 1.300126e-03 Pa  = 1.283124e-08 atm
  a =   10 um   |F/A| = 1.300126e-07 Pa  = 1.283124e-12 atm
  a =    1 mm   |F/A| = 1.300126e-15 Pa  = 1.283124e-20 atm
```

At every biological and architectural scale the Casimir force is utterly negligible. It does nothing in a cell and nothing in a room. The a^{-4} dependence is why — and that exponent is **an exact property of the formula, not an estimated one** (Section 3.5). A hundredfold increase in separation divides the pressure by a hundred million.

The instructive part is that Casimir is real. Someone arguing that vacuum fluctuations couple to matter is not wrong about the physics. They are wrong about the number, by twenty orders of magnitude, and **being right about the mechanism and wrong by twenty orders of**

magnitude is indistinguishable in practice from being wrong about the mechanism. That is the whole content of Class C, and it is why Section 6 makes magnitude prediction a mandatory step rather than a refinement. The same discipline applied to a different exotic borrowing yields the same verdict: traversable wormholes require exotic matter violating an energy condition, so the paper that introduced them concludes against the shortcut they are usually cited to support [Morris and Thorne, 1988].

35.5.6 5.6 Class A continued: three further boundaries

Belief does not alter an uncontacted person's body. The qualifier "no perceptual contact" does all the work and must stay attached: belief communicated — by word, tone, face, touch, treatment — routinely changes another person's physiology, and the framework documents those channels in detail [Curzi, 2026ab]. The claim here is narrow. The largest published test of intercessory prayer at a distance (Benson et al., *American Heart Journal*, 2006; **named without a key**) studied roughly 1,800 cardiac surgery patients across multiple centres and found no benefit. **This is reported as a measurement about a channel, and no clinical, medical or therapeutic claim is made in either direction.** Note also what it does not establish: **a null bounds a gain; it does not prove a zero.** With arms of roughly this size the study bounds a standardised effect below about 0.16 at conventional power. **It does not exclude 0.05.**

Belief does not change physical constants. The speed of light, Planck's constant, the elementary charge and the fine-structure constant are measured to extraordinary precision and are not observed to vary with the beliefs of those measuring them; the fine-structure constant is known to roughly a part in ten billion. **This is worth one paragraph rather than ten because there is no interesting failure mode here.** The reason to state it is definitional: it fixes the outer boundary of the postulate's territory, and **a postulate with no stated outer boundary is not a postulate.**

Belief does not reach backward in time — and the episode that showed it is the catalogue working. Bem's 2011 paper in the *Journal of Personality and Social Psychology* reported nine experiments in which effects appeared to precede their causes. **The paper was methodologically conventional**, used standard designs, analyses and thresholds, and appeared in a leading journal — **which is precisely why it mattered: if orthodox method could produce a significant result for retrocausal influence, the problem was the method, not the topic.** Ritchie, Wiseman and French published three unsuccessful replication attempts in *PLoS ONE* in 2012 and the effect did not appear. **It is also reported — and this is stated here as a reported account of the episode rather than as a documented editorial policy the author has verified — that the replication attempts were initially declined on the grounds that the journal in question did not publish replications.** If accurate, that is $P(\text{sig} \mid \text{no effect})$ in the bottom row of Section 5.10's table made visible as a mechanism rather than inferred as a parameter [Ioannidis, 2005, Simmons et al., 2011, Gelman and Loken, 2013].

```
CLAIM     a belief or response at time t is influenced by a stimulus selected
          at time t + k, k > 0.
CHANNEL   none. Retrocausal influence has no established physical carrier, and
          the claim is not that the carrier is weak but that no candidate
          carrier has been identified.
CLASS     A.
STATUS OF THE EVIDENCE   the original result did not replicate.
```

The reason this episode belongs in a paper defending a belief-related postulate is that it shows the discipline working. A striking result appeared, it was taken seriously enough to replicate, the replications failed, and the field's response became one of the proximate causes of the reform movement [Open Science Collaboration, 2015, Nosek et al., 2018]. **A postulate that wants to be believed should want that machinery pointed at it, and Section 5.10 points it at this framework.**

35.5.7 5.7 Class B: growth mindset — the channel is real and the gain is small

This is the most instructive case in the section, because the popular claim and the measured effect differ not in kind but by an order of magnitude, and the temptation to round the effect up is strongest.

The channel is real. Belief about the malleability of ability changes effort allocation; changed effort allocation changes practice; changed practice changes skill. Every link is behavioural, observable and uncontroversial. **This is not Class A, and naming it Class A would be as much an error as inflating it.**

The gain is small. Sisk and colleagues (*Psychological Science*, 2018; **named without a key**) conducted two meta-analyses — one of the correlation between mindset and academic achievement, one of the effect of mindset interventions on achievement. Both were weak: the correlation in the region of $r = 0.10$ and intervention effects in the region of $d = 0.08$, with subgroup effects for students at academic risk and from low-income backgrounds. Unpack those magnitudes, because the unpacking is the argument:

```
r = 0.10  ->  variance explained  r^2 = 0.0100 = 1.00%

d = 0.08:
  Cohen's U3, per cent of treated above the untreated median
                                  = Phi(0.08)         = 53.188%   (vs 50.000%)
  common-language effect size     = Phi(0.08/sqrt 2)  = 52.256%   (vs 50.000%)
  overlap of the two distributions = 2*Phi(-0.04)     = 96.809%

on a GPA scale with SD = 0.70:     d = 0.08  ->  0.0560 GPA points
on an IQ-type scale with SD = 15:  d = 0.08  ->  1.2000 points
on a test scale with SD = 100:     d = 0.08  ->  8.0000 points
```

Ninety-six point eight per cent of the two distributions overlap. Pick one treated student and one untreated student at random: the treated one scores higher **52.256 per cent** of the time, against 50 per cent by chance. **That is the size of the thing.**

The large field trial is consistent. Yeager and colleagues (*Nature*, 2019; **named without a key**) ran a short online intervention across a national sample and reported a gain on the order of a tenth of a grade point in GPA among lower-achieving students — $d = 0.1429$ on a GPA standard deviation of 0.70, or $d = 0.1000$ on a standard deviation of 1.00.

Now the point that separates an honest report from a dismissal. A tenth of a grade point, delivered by a short online exercise at very low marginal cost across a national population, is a genuinely useful public intervention. **A small effect at national scale and near-zero cost is good engineering. It is simply not what "belief creates reality" is usually taken to mean.**

```
  MEASURED:  d ~ 0.08, sometimes ~0.14 in a subgroup, from an intervention
             costing minutes and pennies per student.
  CLAIMED:   that adopting a belief about malleability transforms outcomes.
  CLASS:     B. The channel is real. The gain is roughly one order of magnitude
             below the claim, and the claim never states a magnitude at all,
             which is how the gap stays invisible.


CONCORDANCE 5.1 -- BELIEF ABOUT ONE'S OWN MALLEABILITY AFFECTS ONE'S OUTCOMES

  CLAIM        Holding that ability is improvable changes effort and therefore
               outcomes.

  (+1) EVIDENCE
               Sisk et al. 2018 (two meta-analyses; r ~ 0.10, d ~ 0.08).
```

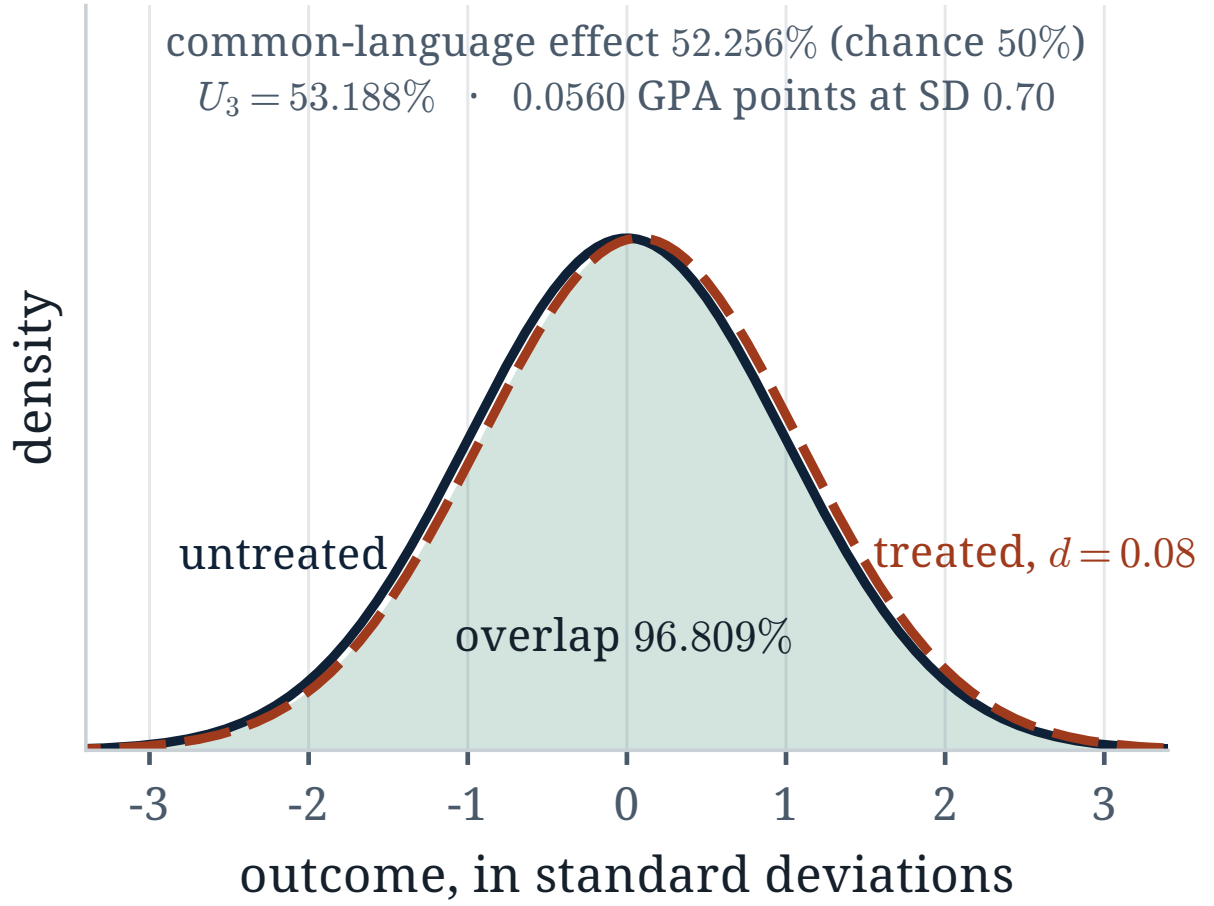


Figure 35.2: **Ninety-six point eight per cent of the two distributions overlap:** the treated and untreated distributions at the meta-analytic growth-mindset effect $d = 0.08$, with the overlap region shaded and the three unpackings marked — Cohen's $U_3 = 53.188$ per cent, the common-language effect 52.256 per cent, and 0.0560 GPA points on a 0.70-point standard deviation. **Modelled, not measured:** the curves are unit normal densities offset by the reported d; the effect size itself is quoted from the meta-analytic literature named in Section 5.7 and carries that literature's uncertainty, not this paper's. Generated by `figs/src/CCP-035_fig02.py`.

```
              Yeager et al. 2019 (national field experiment; about a tenth of
              a grade point among lower-achieving students).
              WHAT IT MEASURED: a real, replicated, very small average effect,
              larger in specific subgroups.

(-1) PRECEDENT -- CARRIES NO EVIDENCE
              Craft and monastic training traditions across many cultures make
              the trainee's conviction of improvability an explicit part of the
              curriculum. This is a characterisation of what those traditions
              recorded, offered as such and not as a quotation.

(0p) INTEGRATION
              The traditions were right about the sign and silent about the
              size. The measurement supplies the size, and the size is small.
              THE (-1) COLUMN CANNOT RESCUE THE MAGNITUDE AND MUST NOT BE USED
              TO TRY.

CRITERIA      C1 PASS. C2 PASS. C3 PASS at the measured magnitude only.

STATUS        ESTABLISHED as a real effect at d ~ 0.08. Any claim of a larger
              effect is UNSUPPORTED and this paper does not make one.
```

35.5.8 5.8 Class C: the general problem of transfer

Steele and Aronson's 1995 work in the *Journal of Personality and Social Psychology* (**named without a key**) reported that making a negative stereotype salient depressed the test performance of members of the stereotyped group. The finding was influential, generated a very large literature, and was widely used to explain real-world group differences in achievement.

The replication record has been substantially more difficult than the original literature suggested. Flore and Wicherts's 2015 meta-analysis in the *Journal of School Psychology* found a small effect together with evidence of publication bias, and reported bias-corrected estimates close to zero; a subsequent large registered replication (Flore et al., 2018) did not find the effect. **This literature is contested and active, not all researchers accept the deflationary reading, and the correct summary is "replication has been difficult" rather than "the effect does not exist."**

The analytic point is the transfer, and it is Class C:

```
LABORATORY:  a salience manipulation, minutes before a test, on a sample,
             produces a measurable performance shift in some designs.
CLAIMED:     that this mechanism accounts for a substantial portion of
             population-level achievement gaps that persist across years,
             institutions and life outcomes.
THE GAP:     the second claim requires a magnitude the first was never
             measured at, over a duration the first was never measured over,
             with a persistence the first never demonstrated.
```

A brief laboratory manipulation and a decades-long structural disparity are separated by many orders of magnitude in duration, dose and accumulation, and no measurement bridges them. The laboratory result may be entirely sound and the transfer still unsupported. **That is what makes Class C the most seductive of the four: nothing in the original work needs to be wrong.** The general problem — that laboratory results on group and individual performance do not transfer to field magnitudes without a stated transfer assumption — is documented across a whole literature the corpus relies on elsewhere [Thompson, 2014, Singer, 2019, Grim et al., 2019, Mannix and Neale, 2005, Woolley et al., 2010, Hong and Page, 2004, Page, 2007], and it is the reason a transfer assumption must be **stated** rather than assumed.

35.5.9 5.9 The replication discount, applied to this framework

```
Open Science Collaboration, "Estimating the reproducibility of psychological
science," Science, 2015:

  originals reporting a significant result       97 of 100
  replications reporting a significant result    about 36 of 100
  mean effect size r                             0.403  ->  0.197
  ratio                                          0.4888   (a drop of 51.12%)

Camerer et al., Nature Human Behaviour, 2018, replications of social-science
studies published in Nature and Science, 2010-2015:

  replicated in the same direction at p<0.05     13 of 21  = 61.90%
  replication effect sizes                       about half the originals
```

Now apply it, with the assumptions printed. Take a prior probability of 0.10 that a given proposed effect in this area is real, and compute the posterior after one statistically significant study under four regimes:

power	P(sig | no effect)	regime	posterior
0.80	0.05	well-powered, no publication filter	0.6400
0.35	0.05	typical power, no filter	0.4375
0.35	0.20	typical power, moderate selective reporting	0.1628
0.35	0.30	typical power, heavy selective reporting	0.1148

The bottom row is the honest frame for a single striking result in this literature. A prior of 0.10 becomes a posterior of 0.1148 — one study moved the belief by less than two percentage points. That is what it means to work in a field with a replication problem: single results are nearly uninformative, and the informative unit is the replicated, preregistered, adequately powered multi-site estimate [Nosek et al., 2018].

And this applies to this framework. The coupling architecture [Curzi, 2026ab], the four-class taxonomy of this paper and the seven-step protocol of Section 6 all sit in exactly the same epistemic regime as the claims they evaluate. **If this framework's own predictions were tested and returned a single significant result, the correct posterior update would be the bottom row of that table, and the author is obliged to say so in print.** Section 6's insistence on preregistration, four baselines and the shuffled-input control is the response, and it is the only response available. **One asymmetry is claimed and one only: the framework's static results are mathematics, and mathematics does not have a replication problem — it has a proof-checking problem, which is a different and better-solved difficulty.** The surplus law and its ceiling [Curzi, 2026], the storage law [Curzi, 2026a], and the contraction result — whose a posteriori bound $d(x_n, x^*) \le (k^n/(1-k))\, d(x_0, x_1)$ reads operationally as *distance to arrival is the last correction scaled by* $q/(1-q)$, so that the size of one's most recent revision is a self-reading gauge [Curzi, 2026] — are theorems. **Everything this framework says about human beings is in the same boat as everything else said about human beings.**

35.5.10 5.10 A null bounds a gain; it does not prove a zero

```
minimum detectable standardised effect at 80% power, alpha .05 two-sided,
two arms of n each:      d_min = (1.959964 + 0.841621) * sqrt(2/n)

  n per arm       d_min
       50        0.5603
      100        0.3962
      300        0.2287
```

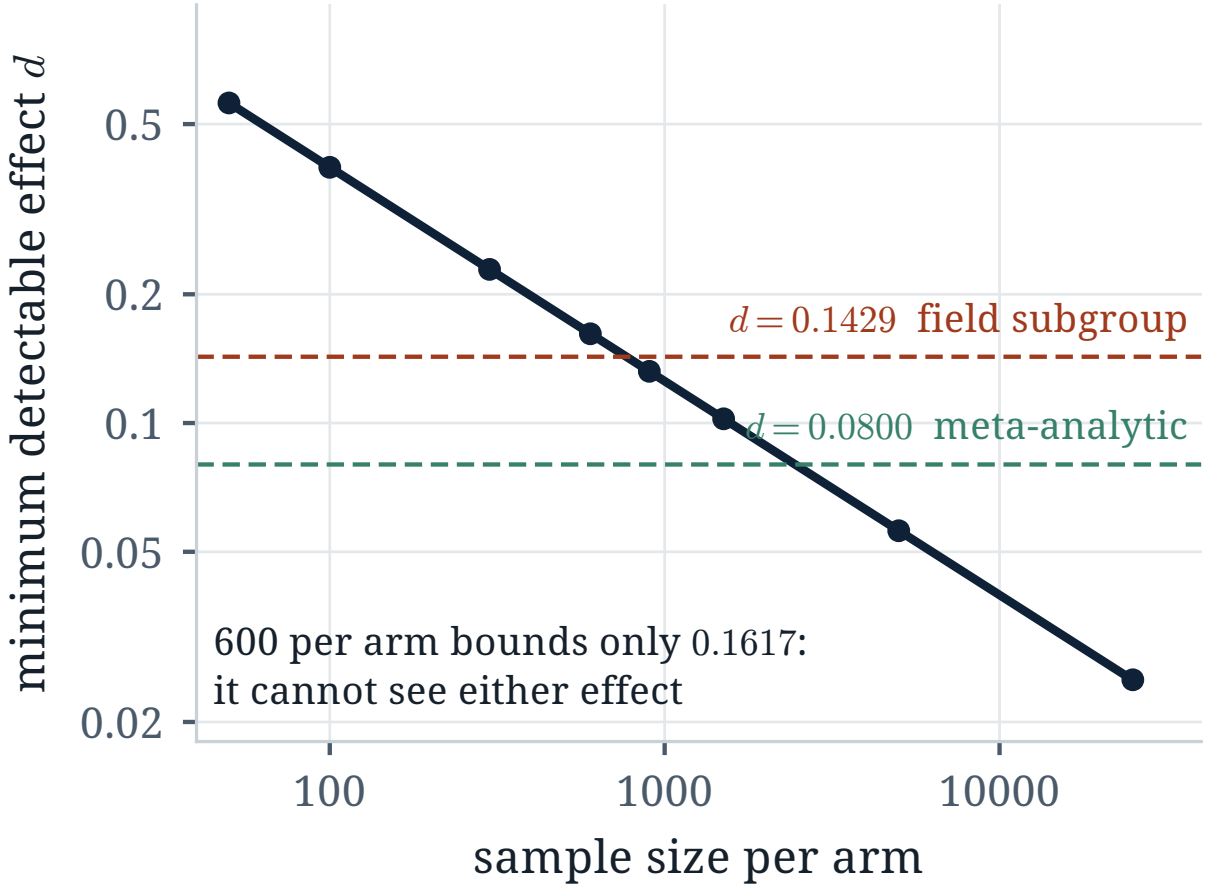


Figure 35.3: **A null bounds a gain; it does not prove a zero:** the minimum detectable standardised effect $d_{\min} = (1.959964 + 0.841621)\sqrt{2/n}$ at 80 per cent power and $\alpha = 0.05$ two-sided, against the sample size per arm, with the two magnitudes this paper reports marked — the meta-analytic growth-mindset effect $d = 0.08$ and the field subgroup effect $d = 0.1429$. **A trial of 600 per arm cannot see either of them.** Exact evaluation of the stated power formula; the two marked effects are quoted from the literature named in Section 5.7. Generated by `figs/src/CCP-035_fig03.py`.

600	0.1617
900	0.1321
1,500	0.1023
5,000	0.0560
25,000	0.0251

A trial with 600 per arm reporting no effect has bounded the effect below about 0.16. It has not shown the effect is zero, and it cannot. Every null in this section is therefore an upper bound on a gain and is reported as one.

This cuts both ways and it is supposed to. A sceptic who reads a null as proof of absence has made the mirror image of the believer's error: both have taken a magnitude question and answered it with a direction. **The correct statement is always a number with an interval**, and where this section says "belief does not do X" it means either that no channel exists — Class A, where no trial is needed because there is nothing to measure — or that the gain is bounded below a stated value by a stated study.

35.5.11 5.11 Survivorship: why a testimonial corpus carries no information about a rate

The successes write books; the failures do not. The literature reports numerators. This is not a criticism of any individual account's honesty. It is a property of how the corpus is assembled.

Suppose a corpus contains 500 first-person accounts of a practice working. Compute the population required under four hypotheses:

success probability p		practitioners needed for 500 successes
0.500	the practice works, half succeed	1,000
0.100	works weakly	5,000
0.010	works barely	50,000
0.001	does nothing beyond the base rate	500,000

Every one of those worlds produces exactly 500 books. Bestselling works in this genre have sold in the millions; a practising population in the hundreds of thousands or millions is entirely plausible. So the bottom row — the practice does nothing at all beyond whatever people would have achieved anyway — predicts the observed corpus just as well as the top row.

```
LIKELIHOOD RATIO between any two rows, given only the numerator:  1.000000
```

Exactly one. The corpus carries zero information about p, because the denominator was never observed. This is not a weak argument against testimonials; it is a proof that testimonials in this format are evidentially inert with respect to the quantity at issue.

The precedent is Wald's work for the Statistical Research Group during the Second World War (**named without a key**). Asked where to armour aircraft given the distribution of bullet holes in aircraft that returned, the insight was that the data were conditioned on returning: **armour the places with no holes, because aircraft hit there did not come back to be measured.** The testimonial corpus is the returning aircraft — a census of survivors presented as a census of attempts.

Three consequences, stated as operating rules:

1. **The base rate is unobtainable from testimonials, in principle, at any corpus size.** A thousand more accounts adds a thousand more numerators and no denominator. The problem is not sample size and cannot be fixed by more of the same.
2. **Getting the denominator requires enrolling people before the outcome is known.** That is a cohort design, and it is the reason Section 6's protocol makes preregistration mandatory rather than advisable.
3. **The same defect afflicts case studies of institutions, teams and rooms — including the worked examples in this corpus, where they are drawn from success.** The corpus's own failure records exist partly to supply the missing tail [Curzi, 2026].

```
CONCORDANCE 5.2 -- TESTIMONY OF SUCCESS ESTABLISHES A RATE

  CLAIM        That a large body of first-person accounts of a practice working
               is evidence that the practice works.

  (+1) EVIDENCE
               Wald's survivorship analysis (1943) and the standard treatment of
               selection on the dependent variable.
               WHAT IT MEASURED: that conditioning on the outcome destroys the
               inference, and that the destruction is total rather than partial.
```

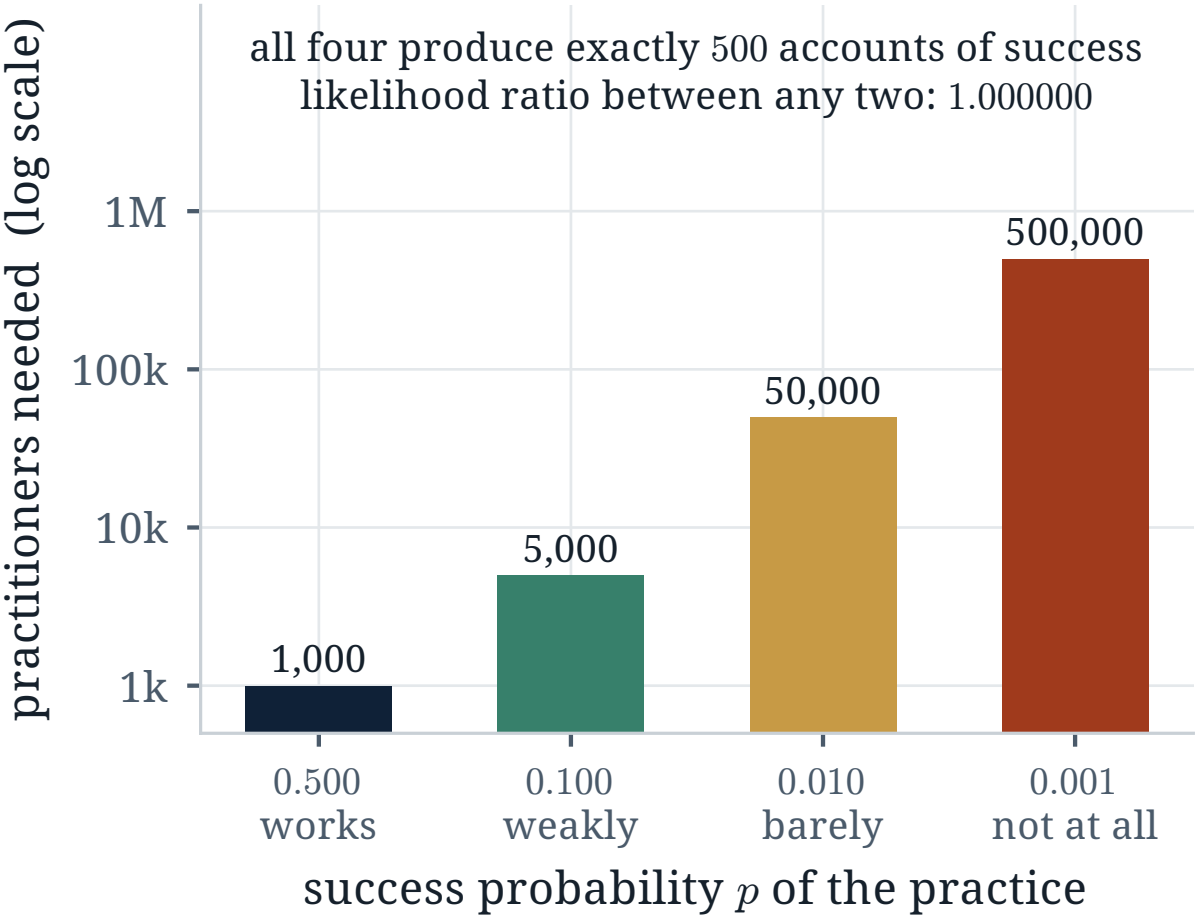


Figure 35.4: **Every hypothesis about the success rate produces exactly the same 500 books:** the practising population required to generate 500 first-person accounts of success, under four success probabilities from 0.500 to 0.001, on a logarithmic scale. **The likelihood ratio between any two rows, given only the numerator, is exactly 1.000000** — the corpus carries zero information about the rate. Exact arithmetic on the stated hypotheses; the corpus size of 500 is an illustrative round figure and is not a count of any real bibliography. Generated by `figs/src/CCP-035_fig04.py`.

```
(−1) PRECEDENT -- CARRIES NO EVIDENCE
             The classical sceptical response to votive offerings -- the reply
             that the drowned leave no tablets -- is recorded in the ancient
             philosophical literature, most familiarly through Cicero's
             De Natura Deorum.
             WHAT IT RECORDS: that the fallacy was identified by argument two
             thousand years before it was formalised.

(0p) INTEGRATION
             The old argument and the modern statistic identify the same
             defect. The modern statistic quantifies it: the likelihood ratio
             is 1.000000, which is the strongest possible statement of the old
             objection.

CRITERIA     C1 PASS. C2 PASS. C3 PASS.

STATUS       THE CLAIM IS REJECTED. This is one of the rejections the method
             requires: a method that never rejects is not a method.
```

35.5.12 5.12 Class D: the system is not responsive to an input of this size, from a source of this size

The last class is the subtlest and the most common in practice, because in Class D everything the believer says about the mechanism is true, and the effect is still zero.

Case: a belief about a market price. Belief that an asset will rise leads to buying, which raises the price. Every link is real; the channel is the actual mechanism by which prices move. **And one person's belief moves the price by an amount indistinguishable from noise, because the price is set by an aggregate in which the individual is one term of a very large sum. The channel is real and the individual's share of it is** $1/N$**.**

Centola et al. [2018] measured convention tipping at roughly a 25 per cent committed minority:

group size	one person is	tipping needs	shortfall factor
12	8.3333%	3 = 25%	3.0x
24	4.1667%	6 = 25%	6.0x
150	0.6667%	38 = 25%	37.5x
240	0.4167%	60 = 25%	60.0x
2,400	0.0417%	600 = 25%	600.0x
100,000	0.0010%	25,000 = 25%	25,000.0x

In a group of twelve, one committed person is a third of the way to the tipping point. In a group of a hundred thousand, one person is twenty-five thousand times short. The mechanism is identical in both. **The same belief, through the same channel, in the same kind of system, is decisive at one scale and inert at another — and nothing about the belief distinguishes the two cases.** The corpus reaches the structurally identical conclusion for contact effects, where a boundary-cost multiplier determines whether contact helps at all and is computed at 3.17 against an operating value of 8 [Curzi, 2026g].

Case: a belief nobody learns of. If a belief about another person's decision is never communicated in any modality — not spoken, not written, not signalled by posture, tone, timing or differential treatment — then there is no channel, and the case is Class A. **The diagnostic question is therefore not "is the belief sincere?" but "by what physical means does the other person come to be affected?"** In practice beliefs leak constantly through behaviour, and the honest analysis is usually that a real but small behavioural channel exists — Class B, which is what Section 4.13 measures.

Case: a belief about a system with a fixed schedule. A traffic signal on a fixed timer is not responsive to any input a pedestrian can supply. The channel from belief to action exists; the system does not read those inputs. **Class D, and the diagnostic is one sentence: what does this system take as input?**

Class D failures are failures of the receiver, and receivers fail in exactly three ways:

```
D1  THE RECEIVER DOES NOT READ THIS INPUT AT ALL.
    A fixed timer reads a clock. A conserved quantity reads nothing.
    Diagnostic: enumerate the system's actual inputs. If the channel's output
    is not on the list, D1, and no increase in input strength helps.

D2  THE RECEIVER READS IT, AND THIS SOURCE IS 1/N OF WHAT IT READS.
    A market price, an election, a social convention.
    Diagnostic: compute the share. It is one division and it settles the case.
    THIS IS THE ONE FAILURE MODE THAT CHANGES WITH SCALE.

D3  THE RECEIVER READS IT, AND IS ALREADY SATURATED.
    A student already receiving maximal attention; a market already fully
    informed; a queue already at capacity.
    Diagnostic: measure the response to a known perturbation at the CURRENT
    operating point, not at the point where the original study measured it.
    This is where the growth-mindset subgroup findings live: the effect is
    larger for students at academic risk because they are further from
    saturation.
```

D2 is the only one a bigger coalition fixes. D1 is fixed by finding a different channel. D3 is fixed by finding a different operating point. Confusing them wastes years, because all three present identically from inside: *I did the thing sincerely and nothing happened.*

35.5.13 5.13 The catalogue in one table

CLAIM	CLASS	WHY IT FAILS
belief moves an external object	A	no transducer; field energy short by ~2.5e24
belief alters an uncontacted person's body	A	no channel; the largest trial bounds any effect, does not prove zero
belief changes physical constants	A	no channel; constants measured to ~1 part in 1e10
Casimir does work at cellular scale	C	real effect, wrong by ~20 orders of magnitude
growth mindset transforms outcomes	B	real channel, d ~ 0.08; claim is ~10x the gain
a lab salience effect explains population gaps	C	lab effect, contested replication, no bridge to decades-long magnitude
one person's belief moves a market price	D	real channel, share is 1/N; shortfall 25,000x at N=1e5
testimonials establish that a practice works	-	evidentially inert; likelihood ratio 1.000000
belief cures disease by itself	-	NO CLAIM MADE. This paper makes no clinical, medical or therapeutic claim, in either direction, anywhere.

The last row is not a failure class because it is not a claim this paper evaluates. The prohibition is absolute, and the reason it appears in the catalogue at all is that its absence would otherwise be read as an omission.

35.5.14 5.14 What survives, at its measured size

The catalogue would be dishonest in the other direction if it stopped here.

```
SURVIVES: placebo analgesia through endogenous opioids, naloxone-reversible
          (Levine, Gordon & Fields, 1978). A complete verified physical
          channel from expectation to measured outcome.
          MAGNITUDE: modest; the literature is active; effect sizes vary by
          condition and design.

SURVIVES: open-label placebo (Kaptchuk et al., 2010, and subsequent trials).
          Benefit reported when recipients are told they are receiving
          placebo. Decisive here because the effect does not require
          deception. MAGNITUDE: modest; literature active.

SURVIVES: self-fulfilling prophecy where the system is structurally
          responsive (Merton 1948; Diamond & Dybvig 1983).
          MAGNITUDE: can be total, as in a deposit run -- and only where
          W3 holds.

SURVIVES: expectancy effects on performance.
          MAGNITUDE: real, small, condition-dependent, prone to dissipate
          (Raudenbush 1984; Jussim & Harber 2005).

SURVIVES: perception partly constructed from priors.
          MAGNITUDE: substantial for perception; the research programme is
          active and contested.
```

Every survivor on that list has a named physical channel and is reported with its magnitude. That is not a coincidence; it is the selection criterion. **The list of what belief does is exactly the list of channels that have been found and measured, which is what the postulate predicts, and it is the strongest support the postulate has.** All five are cited as evidence about channels and their gains; **no clinical, medical or therapeutic claim is made about any of them.**

35.5.15 5.15 Requirements

REQ-5.1 Classify every failure as A, B, C or D and print the class. “It doesn’t work” is not a finding.

REQ-5.2 Establish grounding from the channel, never from the outcome. This is not advice.

REQ-5.3 Name the transducer before asserting any channel. If no physical component can be named that converts the belief into the asserted influence, the claim is Class A and the analysis is finished.

REQ-5.4 Report every effect at its magnitude, in the units the reader cares about. For $d = 0.08$: 96.809 per cent overlap, 52.256 per cent common-language effect, 0.0560 GPA points. **Direction-only reporting is prohibited.**

REQ-5.5 Cite the growth-mindset meta-analysis with its effect size whenever the concept is invoked. Invoking it without the magnitude is a citation error.

REQ-5.6 State replication status alongside any social-science finding used as support, in the sentence that makes the claim.

REQ-5.7 Never read a null as proof of absence. Report the minimum detectable effect at the study's power and state the null as an upper bound.

REQ-5.8 Never treat a testimonial corpus as evidence about a rate. The likelihood ratio is 1.000000.

REQ-5.9 Apply the replication discount to this framework's own claims.

REQ-5.10 Print the Casimir clause wherever Casimir appears: negligible at every biological and architectural scale.

REQ-5.11 Distinguish "the channel does not exist" from "my share of it is $1/N$". These are Classes A and D, they look identical from the inside, and only one of them changes with scale.

REQ-5.12 Make no clinical, medical or therapeutic claim, in either direction. Absolute.

REQ-5.13 Keep the concordance columns separate in the failure catalogue too.

REQ-5.14 When a claim's magnitude cannot be stated, record it as **untested** rather than as **unproven**. These are different verdicts.

35.5.16 5.16 Failure modes of the catalogue itself

FM-5.1 — EX POST GROUNDING. *Signature:* the outcome is observed and grounding is assigned to match. *Mechanism:* the circularity trap. *Fix:* name and check the channel before the outcome is known, in writing, with a date.

FM-5.2 — CHANNEL LAUNDERING. *Signature:* a Class A claim defended by pointing at a real Class B channel, with the Class A magnitude retained. *Fix:* separate the claims and make each name its own transducer.

FM-5.3 — MAGNITUDE-FREE ASSERTION. *Signature:* a claim states a direction and never a size, and so cannot be wrong. *Mechanism:* direction-only claims are worth a likelihood ratio of 2.000000. *Fix:* require a predicted magnitude with a tolerance before measurement. **Most claims in this area die here, and dying here is much cheaper than dying after the study.**

FM-5.4 — DENOMINATOR BLINDNESS. *Signature:* a large number of success stories treated as a large body of evidence. *Fix:* ask who is missing.

FM-5.5 — NULL AS PROOF OF ABSENCE. *Signature:* "the trial found nothing, therefore there is nothing." *Fix:* report the minimum detectable effect and state the null as an interval.

FM-5.6 — PRECEDENT LAUNDERED INTO EVIDENCE. *Signature:* an old source and a modern study cited together in a way implying two independent confirmations. *Fix:* state each column's role in the same sentence, every time [Curzi, 2026w].

FM-5.7 — LAB-TO-LIFE TRANSFER. *Signature:* an effect measured over minutes asserted to explain an outcome accumulated over decades. *Fix:* state the dose, duration and accumulation of the original next to those of the target claim.

FM-5.8 — SCALE-BLIND COALITION CLAIM. *Signature:* a mechanism verified in a small group asserted for a large one without adjusting the individual's share. *Fix:* compute the share. It is one division.

FM-5.9 — SELECTIVE SCEPTICISM. *Signature:* the replication crisis invoked against inconvenient findings and forgotten for convenient ones. *Fix:* apply the posterior table to this framework's own claims in print, as Section 5.9 does.

FM-5.10 — CATALOGUE AS DISMISSAL. *Signature:* the failure catalogue read as showing that belief does nothing. *Mechanism:* the catalogue is a boundary, not a verdict. *Fix:* never present the failure list without the survivor list and its magnitudes.

FM-5.11 — THE UNSTATED CLINICAL IMPLICATION. *Signature:* no medical claim is made explicitly and the surrounding text is arranged so that a reader will infer one. *Fix:* state the prohibition in the reader's line of sight wherever the subject comes near. **This section does so three times and the repetition is deliberate.**

35.6 6 Grounding: the protocol that answers the circularity objection

Section 5 closed with an objection the postulate cannot survive without answering. **If "grounded in truth" is diagnosed from the outcome, the postulate predicts nothing, forbids nothing, and is a relabelling of whatever happened.** There is exactly one answer and it is not an argument. It is a procedure.

> **GROUNDING IS ESTABLISHED FROM THE CHANNEL, BEFORE THE OUTCOME IS KNOWN, OR IT IS NOT ESTABLISHED AT ALL.**

One step does most of the killing, and this section says so on the page rather than in the notes. Step 5 requires a predicted magnitude — not a direction, a number, with a tolerance, written down before the measurement. Most claims about belief die at step 5, because a belief that predicts only a direction **is not yet testable**: it has not been refuted, it has not been examined, it has not made a statement about the world sharp enough to be wrong. **That is a worse condition than being false, and the fact that it feels like a safer one is precisely why so many claims stay in it.**

35.6.1 6.1 The grounding criterion

```
A BELIEF b, ASSERTING A STATE x*, IS GROUNDED IF AND ONLY IF THERE EXISTS A
CHANNEL K SUCH THAT:

  G1  K IS NAMED.            Every link is identified: source, transducer,
                             medium, receiver, effector. No link may be a
                             placeholder, and "somehow" is not a link.

  G2  K IS PHYSICALLY POSSIBLE.
                             Each link is consistent with conservation,
                             causality, thermodynamics, and the signal-to-noise
                             available at the stated amplitude.

  G3  K HAS A GAIN g > 0, ESTIMATED INDEPENDENTLY OF THE OUTCOME OF b.
                             From prior measurement, from the physics of the
                             link, or from a separate calibration -- never from
                             the case under test.

  G4  THE PREDICTED MAGNITUDE g * b IS STATED, WITH A TOLERANCE, BEFORE
      MEASUREMENT.

  G5  THE PREDICTION SURVIVES THE FALSIFICATION SUITE,
                             including the shuffled-input control.

GROUNDING IS A PROPERTY OF THE CHANNEL. IT IS NEVER A PROPERTY OF THE RESULT.
```

Badge: G1 to G5 are this framework's operationalisation of the postulate's conditional clause, and they are a proposal rather than a theorem. What makes them more than a stylistic preference is G3 and G4 together: they convert a conditional that cannot be wrong into a claim that predicts a number. **A postulate that predicts a number can be killed, which is the only thing that makes it worth defending.**

35.6.2 6.2 The seven steps

```
  1.  STATE THE BELIEF PRECISELY ENOUGH TO BE WRONG.
```

```
     A sentence whose negation is also a clear sentence.

  2. IDENTIFY THE CLAIMED CHANNEL EXPLICITLY.
     SOURCE -> TRANSDUCER -> MEDIUM -> RECEIVER -> EFFECTOR, each named.

  3. CHECK THE CHANNEL IS PHYSICALLY POSSIBLE.
     Six checks. Most Class A claims stop here, in minutes.

  4. ESTIMATE ITS GAIN.
     Per link, from independent sources. Multiply along the chain. Report the
     end-to-end number with its uncertainty.

  5. PREDICT THE MAGNITUDE OF THE EFFECT, NOT ONLY ITS DIRECTION.
     A number, in stated units, with a stated tolerance, recorded before
     measurement. THIS IS WHERE MOST CLAIMS DIE.

  6. MEASURE.
     Preregistered, blinded where possible, with the instrument specified in
     advance. The seven formula-validation gates in print.

  7. COMPARE PREDICTION TO OUTCOME.
     Against four baselines including the shuffled-input control. Four
     verdicts are available and only one of them is "grounded."
```

The ordering is load-bearing and the steps are not commutative. Step 2 before step 4 prevents a gain being borrowed from a channel other than the one claimed. Step 4 before step 5 prevents the predicted magnitude being chosen for its plausibility. Step 5 before step 6 prevents the tolerance being widened after the data arrive. **Every one of those inversions is a documented way of getting a result you like, and the ordering is the defence** [Simmons et al., 2011, Gelman and Loken, 2013].

35.6.3 6.3 Why step 5 does the killing

A direction-only prediction is worth almost nothing, and the amount is computable.

```
Under a symmetric null, the chance of a correct direction call = 0.500000
LIKELIHOOD RATIO of one correct direction call = 2.000000

k independent direction calls, all correct, by chance:
   k= 1   0.500000        k= 6   0.015625
   k= 2   0.250000        k= 7   0.007812
   k= 3   0.125000        k= 8   0.003906
   k= 4   0.062500        k= 9   0.001953
   k= 5   0.031250        k=10   0.000977
```

One correct direction call is worth a likelihood ratio of 2. It takes five consecutive correct calls to reach $p = 0.031250$**.** Almost no claim in the belief literature has produced five independent, preregistered, correct direction calls, and a single one is routinely presented as confirmation. This is Meehl's objection to directional prediction in psychology, stated arithmetically [Meehl, 1967, 1990].

Now compute what a magnitude call is worth. **Suppose the true value is known only to lie log-uniformly over** R **decades — a modelling choice, badged as one, and the numbers below are properties of that prior** — and the claim predicts it within $\pm p$:

```
band width in log space = ln((1+p)/(1-p))        prior span = R * ln(10)
```

```
  R = 1 decade:
     +/- 50%   width 1.098612   chance hit 0.477121   LR      2.0959
     +/- 30%   width 0.619039   chance hit 0.268845   LR      3.7196
     +/- 20%   width 0.405465   chance hit 0.176091   LR      5.6789
     +/- 10%   width 0.200671   chance hit 0.087150   LR     11.4744
     +/-  5%   width 0.100083   chance hit 0.043466   LR     23.0066

  R = 2 decades:
     +/- 50%   chance hit 0.238561   LR      4.1918
     +/- 30%   chance hit 0.134423   LR      7.4392
     +/- 20%   chance hit 0.088046   LR     11.3577
     +/- 10%   chance hit 0.043575   LR     22.9489
     +/-  5%   chance hit 0.021733   LR     46.0133

  R = 3 decades:
     +/- 20%   chance hit 0.058697   LR     17.0366
     +/- 10%   chance hit 0.029050   LR     34.4233
     +/-  5%   chance hit 0.014489   LR     69.0199
```

A single magnitude call at ± 20 per cent against a two-decade prior is worth a likelihood ratio of 11.3577, against 2.0000 for a direction call. Converting:

```
  equivalent number of direction calls = ln(11.357747) / ln(2) = 3.5056
```

One magnitude prediction carries as much evidential weight as three and a half correct direction predictions, and it costs the same single measurement. That is the whole argument for step 5, it is arithmetic given the stated prior, and it explains why the step is where claims die: **a claimant who declines to state a magnitude has declined the cheapest available source of evidence for their own position. That is not caution. It is a preference for remaining untested, and it should be named as one.**

35.6.4 6.4 Warrant

```
PROVED (arithmetic)
  the direction-call and magnitude-call likelihood ratios of Section 6.3,
    GIVEN the log-uniform prior, which is a modelling choice
  the chain-gain product of Section 6.9
  the power and minimum-detectable-effect computations of Section 6.10
  the simulated behaviour of the shuffled-input control, including the
    configuration in which it FAILS

BORROWED (established, from elsewhere in this corpus)
  the seven formula-validation gates (CCP-026)
  the four-baseline falsification template and the shuffled-input control,
    with its own measured anti-conservatism (CCP-027)
  the ceiling 0 <= F(u) <= ln N (CCP-A02)
  the replication numbers of Section 5.9

POSTULATED, BADGED (this framework's)
  the G1-G5 grounding criterion
  the seven-step ordering and the claim that the ordering is load-bearing
  the four-verdict comparison table of Section 6.12
  the worksheet of Appendix C
```

35.6.5 6.5 Step 1 — state the belief precisely enough to be wrong

The test is mechanical: **write the negation.** If the negation is also a clear, checkable sentence, the belief has passed. A sharpening in four passes:

```
PASS 0   "Believing in the project will make it succeed."
         Negation: "Believing in the project will not make it succeed."
         VERDICT: fails. "Believing," "the project," "succeed" are all
         undefined. Both sentences are compatible with every possible world.

PASS 1   "If the team believes the release will ship on the 14th, it is more
         likely to ship on the 14th."
         VERDICT: better -- there is now an event and a date. Still fails:
         "the team believes" is unmeasured, and "more likely" has no size.

PASS 2   "If at least 8 of 11 team members rate their confidence in the
         14th at 4 or 5 on a fixed 5-point scale at kickoff, the release
         ships on or before the 14th."
         VERDICT: measurable, and falsifiable on a single release. Still
         fails step 5: it predicts an event, not a magnitude, and a single
         release is a single direction call worth LR 2.0000.

PASS 3   "Where at least 8 of 11 members rate confidence >= 4 at kickoff,
         mean schedule slip is 2.30 days lower (tolerance +/- 30%, so
         1.61 to 2.99 days) than where fewer do, measured over 20 releases
         on the team's existing slip instrument."
         VERDICT: PASSES. It names a population, an instrument, a magnitude,
         a tolerance, and a sample. Its negation is equally clear.
```

Three things happened between pass 0 and pass 3 and all three are load-bearing. The subject became a measurable population rather than a mood. The outcome became a continuous quantity rather than a binary event, which raises the information per observation. And a magnitude with a tolerance appeared, which is what turns the statement from an opinion into a claim. **A belief is stated precisely enough to be wrong when a stranger, given the instrument and no access to the claimant, could reach the opposite verdict.**

35.6.6 6.6 Step 2 — identify the claimed channel explicitly

```
SOURCE      what holds the belief. A person, a team, a market, a population.
            Its size matters: see Class D2 and the 1/N problem.

TRANSDUCER  what converts the holding of a belief into a physical quantity.
            THIS IS THE LINK THAT IS USUALLY MISSING, and its absence is
            the definition of a Class A claim.
            Real transducers: skeletal muscle, vocal apparatus, facial
            musculature, autonomic response, attention allocation, keystroke.

MEDIUM      what carries it. Air, light, paper, a network, a balance sheet,
            a shared institution, a body's own interior.

RECEIVER    what reads it. A person's senses, a market's order book, an
            instrument, a cell's mechanically gated channel.

EFFECTOR    what changes as a result, and in what units.
```

Four channel families are established and they are the complete list this framework is prepared to assert [Curzi, 2026ab]:

```
PHYSIOLOGICAL   belief -> expectation -> neurochemistry -> measured outcome.
                Established: naloxone-reversible placebo analgesia (Levine,
                Gordon & Fields, 1978); open-label placebo (Kaptchuk et al.,
                2010). Bounded within the believer's own body. Modest
                magnitude. NO CLINICAL, MEDICAL OR THERAPEUTIC CLAIM IS MADE.

BEHAVIOURAL     belief -> action -> outcome. The largest gains in the whole
                taxonomy and the least discussed, because the mechanism is
                unglamorous. Appendix B finds 39.53% of a five-fold effect
                coming from a belief causing an application to be submitted.

SOCIAL          belief -> expressed behaviour -> another's perception ->
                another's behaviour. Expectancy effects (Section 4.13),
                contagion and coalitions. Real; modest per individual;
                scale-dependent by 1/N.

INFORMATIONAL   belief -> a decision -> a system that reads decisions.
                Markets, elections, institutions, queues. Section 4's engine
                lives here.
```

A claimed channel outside these four is not thereby refuted, but it must name its own transducer explicitly and it does not get to borrow the credibility of these four. That is channel laundering in preventive form.

The SOCIAL family's most-used transducer is speech, and the corpus analyses it in its own right. A spoken commitment is a **performative aimed at the speaker**, supported by three converging mechanisms — consistency pressure, self-perception from overheard testimony, and the re-pricing of later defection [Curzi, 2026]. **That is a named transducer with three named mechanisms, which is exactly what step 2 demands**, and it is the reason the worked case of Appendix C can put "speech" in the transducer box and pass P3 rather than stalling there.

35.6.7 6.7 The mechanotransduction line, restated

The framework's central open problem belongs in this protocol because it is the best worked example of a channel that is partly real and partly unestablished, **and the protocol must handle that case without collapsing it in either direction.**

```
ESTABLISHED (the evidence column, and it is strong):
  Mechanically gated ion channels are established: mechanical force opens
  them; ions flow; signalling follows. Sound is mechanical. A real,
  established pathway exists from pressure wave to cellular signal.
  Force at the cell surface transmits through the cytoskeleton to the
  nucleus, stretches chromatin, and upregulates transcription (Tajik et al.,
  Nature Materials 2016; named without a key).

THE LINE:
  THERE IS NO ESTABLISHED MECHANISM BY WHICH A CELL DEMODULATES
  SEQUENCE-SPECIFIC INFORMATION FROM AN ACOUSTIC SIGNAL.
```

In the protocol's terms: the source, transducer, medium and receiver are all established; the effector's selectivity is not. A cell demonstrably responds to mechanical input. Whether it can be *addressed* — whether a particular waveform reaches a particular locus rather than producing a generic mechanotransductive response — is one well-specified, falsifiable gap.

The corpus's sonification protocol states the same boundary from the other side: eleven upstream components are built, verified and deterministic, and the one missing link is precisely the carriage of sequence over force [Curzi, 2026]. **The protocol's verdict on that claim is not "false". It is "step 4 incomplete: gain of the final link unestimated"** — a different and more useful verdict than either enthusiasm or dismissal, and the ability to return it is a point in the protocol's favour.

35.6.8 6.8 Step 3 — is the channel physically possible?

```
P1  CONSERVATION.      Does the claim require energy, momentum, charge or mass
                       to appear or vanish? Compute the budget. Section 5.4
                       found the energy for the object-moving claim available
                       (the brain has 2,039 times what is needed) and the
                       coupling absent -- so P1 passes and P3 fails. Run them
                       separately.

P2  CAUSALITY.         Does any link require influence to travel backward in
                       time, or faster than light between spacelike-separated
                       events?

P3  A NAMED TRANSDUCER.
                       Is there a physical component that converts the belief
                       into the claimed influence? If not, STOP. Class A.
                       This check ends more claims than the other five
                       together.

P4  SIGNAL TO NOISE.   At the claimed amplitude, is the signal above the noise
                       floor of the medium and the sensitivity of the receiver?
                       A signal 5.000e+08 times weaker than an ambient field
                       the receiver already ignores is not a signal.

P5  THE SCALE LAW.     How does the effect scale with distance, size or
                       population? Casimir scales as a^-4: 1.283124 atm at
                       10 nm and 1.300126e-15 Pa at 1 mm. A real effect with a
                       punishing exponent is functionally absent at the scale
                       of interest. CASIMIR DOES NOTHING IN A CELL AND NOTHING
                       IN A ROOM.

P6  WHO LEARNS IT.     If the outcome depends on another agent, by what
                       physical means does that agent come to be affected?
                       If the belief is never expressed in any modality, there
                       is no channel. If it leaks through behaviour, say so,
                       and the claim becomes behavioural or social, with those
                       families' gains.
```

P3 and P6 between them are the whole of Class A. They take a few minutes each and running them first is the single highest-yield decision in the protocol. **A claim that fails P3 should not proceed to step 4, because estimating the gain of a channel that does not exist produces a number, and the number will be quoted.**

35.6.9 6.9 Step 4 — estimate the gain, and mind the chain law

Three definitions, and the protocol requires that the one in use be stated:

```
DIMENSIONAL GAIN   g = (change in output units) / (change in input units).
```

```
                  Units must be carried.

LOG GAIN          g_log = d(ln output) / d(ln input) -- an elasticity.
                  Dimensionless, composes by addition, and it is the right
                  register for anything the recipient perceives
                  logarithmically.

STANDARDISED GAIN  the effect in standard-deviation units, d. Comparable
                  across instruments; requires a stated reference SD, and
                  the SD must be the target population's, not the study's.
```

Then the chain law, where most optimistic gain estimates go to die.

```
g_total = product of the link gains.  A LONG CHANNEL IS A WEAK CHANNEL.

per-link g     L=1        L=2        L=3        L=4        L=5        L=8
   0.95     0.950000   0.902500   0.857375   0.814506   0.773781   0.663420
   0.90     0.900000   0.810000   0.729000   0.656100   0.590490   0.430467
   0.80     0.800000   0.640000   0.512000   0.409600   0.327680   0.167772
   0.70     0.700000   0.490000   0.343000   0.240100   0.168070   0.057648
   0.60     0.600000   0.360000   0.216000   0.129600   0.077760   0.016796
   0.50     0.500000   0.250000   0.125000   0.062500   0.031250   0.003906
   0.30     0.300000   0.090000   0.027000   0.008100   0.002430   0.000066
```

Worked on an ordinary social channel with five plausible links, none of them absurd:

```
belief -> stated intention                    0.70   cumulative 0.700000
       -> scheduled action                    0.60   cumulative 0.420000
       -> action completed                    0.75   cumulative 0.315000
       -> the other party notices             0.50   cumulative 0.157500
       -> the other party changes behaviour   0.40   cumulative 0.063000

END-TO-END GAIN = 0.063000, i.e. -2.764621 nats.
```

Five links, each of which a reasonable person would call "usually works", and 6.3 per cent arrives. This is the same arithmetic that produced Section 4.2's $E = 0.0600$, and it explains why so much sincere effort produces so little measured effect.

The corpus reaches a structurally identical result by a different route. Its comfort-weighted amplitude scores each step of a branch 0 to 3 for steepness, sums to a resistance R, and weights the branch as $1/2^R$, **so a three-point difference in steepness is a factor of eight in expectation** [Curzi, 2026]. **That is the chain law in exponential dress: resistance accumulates additively and the weight falls geometrically**, which is the same statement as $g_{\text{total}} = \prod_i g_i$ with the logarithm taken.

Three rules follow.

1. **Estimate each link from a source that did not measure the outcome under test.** A gain estimated from the case at hand is the circularity trap wearing a number.
2. **Shorten the chain, do not strengthen the links.** Raising every link from 0.60 to 0.70 across five links moves the total from 0.077760 to 0.168070 — a 2.16-fold gain. Removing two links at 0.60 each moves it from 0.077760 to 0.216000, a 2.78-fold gain, **and removing links is usually the cheaper operation.**
3. **Report the total with an uncertainty.** If each link is known to a factor of 1.5, five links compound to $1.5^5 = 7.59$ — the end-to-end estimate spans nearly an order of magnitude, and step 5's tolerance must be set honestly against that.

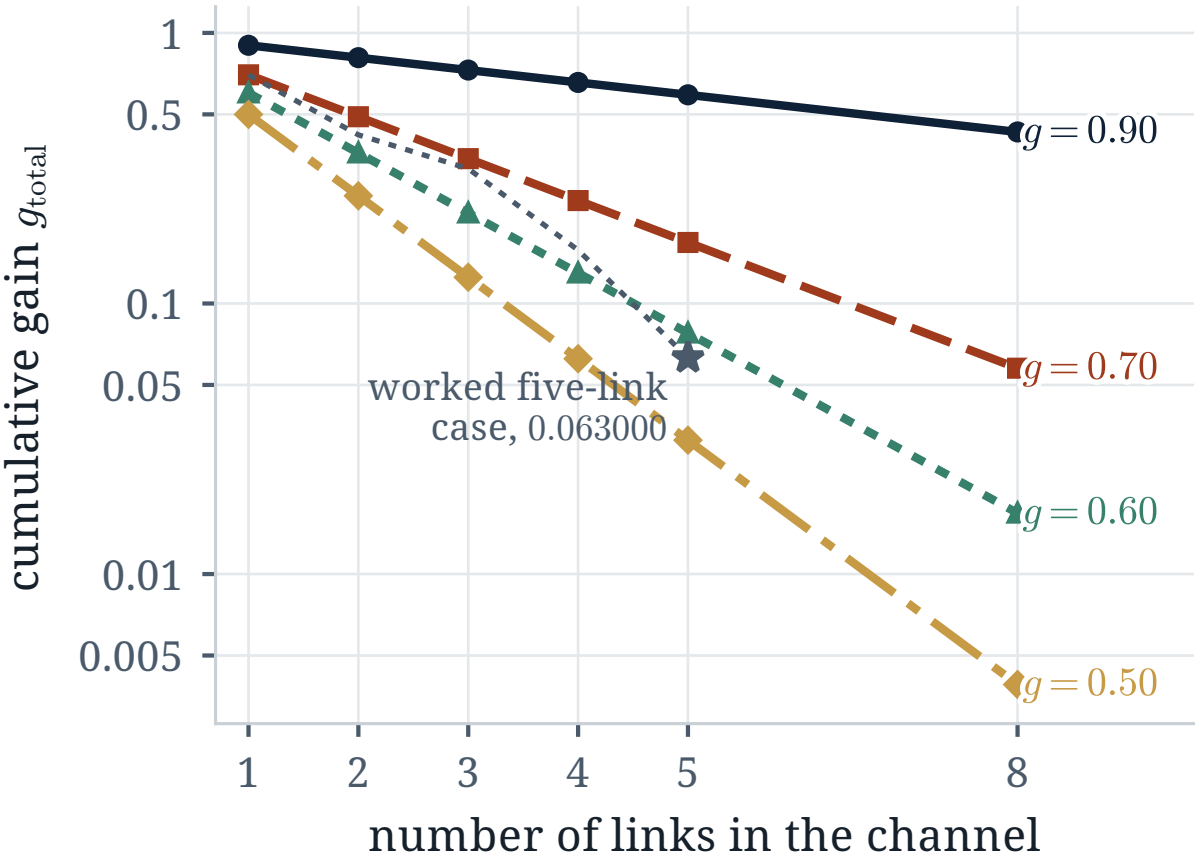


Figure 35.5: **A long channel is a weak channel:** cumulative gain against the number of links, for four per-link gains from the chain-gain table of Section 6.9, with the five-link worked case marked at 0.063000 — five links each of which "usually works" delivers 6.3 per cent end to end. **Modelled, not measured:** the per-link gains are assumed parameters and the worked case's five values are stated in Section 6.9. Generated by `figs/src/CCP-035_fig05.py`.

35.6.10 6.10 Step 5 — predict the magnitude, and price the measurement

A predicted magnitude determines the study required to test it. Two-group comparison, $\alpha = 0.05$ two-sided, 80 per cent power:

predicted d	n per arm	total N
0.80	25	50
0.50	63	126
0.30	175	350
0.20	393	786
0.15	698	1,396
0.10	1,570	3,140
0.08	2,453	4,906
0.05	6,280	12,560
0.02	39,245	78,490

A predicted effect of $d = 0.08$ — the growth-mindset magnitude — requires 4,906 participants to test at conventional power. A predicted effect of $d = 0.50$ requires 126.

> **A small predicted effect is not a modest claim. It is an expensive one. And a claim that refuses to state its magnitude has also refused to say what it would cost to test, which is exactly how it stays untested.**

Three consequences for practice. The tolerance must be set from the gain uncertainty of step 4, not from convenience: if the chain estimate spans a factor of 7.59, a ± 10 per cent tolerance is theatre. If the required N is unobtainable, say so at step 5 and redesign then. And prefer an outcome with a high event rate over an outcome with high importance — **the affordable study is set by the event rate of the link instrumented, not by how much anyone cares about the endpoint.**

35.6.11 6.11 Step 6 — measure, with the validation gates in print

```
FVP-1  DIMENSIONS
       Units of belief measure, gain, and outcome. The gain's units must be
       (outcome units)/(belief units) exactly. Ratios dimensionless.

FVP-2  SCALE TYPE
       Stevens' typology per term. If the belief measure is ordinal -- a
       5-point agreement scale usually is -- then products and ratios
       involving it are NOT licensed, the relation is a heuristic ordering,
       and it must be DEMOTED IN PRINT. This gate stops more belief formulas
       than any other and it should.

FVP-3  LIMITING CASES
       Behaviour as belief -> 0 and -> its maximum. At zero belief the
       predicted effect must be zero, or the channel is not the one claimed.
       At maximum belief the effect must saturate, because every real
       receiver saturates (Class D3).

FVP-4  FIXED POINTS
       Solve for the invariant. For a reflexive channel this is Section 4.10:
       the pole a = (1-delta) + eta*s and the loop gain g = eta*s/delta.
       State whether g is below or above 1 -- and state that s is not
       identified from observational data.
```

```
FVP-5  IDENTIFIABILITY
       Can the gain be recovered from obtainable data? If two different gains
       predict identical observations, the claim is not testable as written
       and must be rewritten before step 6, not after.

FVP-6  FALSIFICATION
       Preregister; fit on training; FREEZE; report held-out against four
       baselines -- persistence, direct-input, matched autoregressive, and
       the SHUFFLED-INPUT CONTROL. State in advance what result kills the
       claim.

FVP-7  REPRODUCIBILITY
       The cross-individual test. Same operationalisation, different people,
       different sites: do the parameters recover? This gate decides the
       hard-science question and it is the one most often skipped.
```

FVP-2 deserves a warning in the running text. Belief is very often measured on an ordinal scale, and **ordinal data do not support the multiplication in effect = gain × belief** [Stevens, 1946, Likert, 1932]. Where that is the situation, the honest move is to demote the relation to a monotone ordering — *more belief predicts more effect* — **and to say in print that the multiplicative form is not licensed by the measurement** [Curzi, 2026t]. A ratio-scale belief instrument is constructible by magnitude estimation [Thurstone, 1927]; **it is simply more work than handing out a five-point form, and the extra work is the difference between a formula and a decoration.** The deeper objection — that quantitative structure must be *tested*, by additive conjoint measurement, before any of this is licensed — stands unanswered [Michell, 1997, Luce and Tukey, 1964], and Section 10 concedes it.

35.6.12 6.12 Step 7 — compare, against four baselines, to one of four verdicts

```
(i)   PERSISTENCE                  kills: "the outcome was going to happen
                                   anyway."

(ii)  DIRECT-INPUT (raw, no log)   kills: the claim that the LOG structure is
                                   doing work.

(iii) MATCHED AUTOREGRESSIVE,      kills: the claim that BELIEF is the input.
      no interaction term

(iv)  SHUFFLED-INPUT CONTROL       kills: the claim that the TEMPORAL
                                   ALIGNMENT matters.

FALSIFIED IF:
  RULE 1  the shuffled-input control predicts held-out outcome as well as
          the ordered series;
  RULE 2  the matched autoregressive baseline is not beaten;
  RULE 3  the fitted a lies outside [0,1] or the fitted b is negative.
```

The suite has been audited against itself and it is anti-conservative: the shuffle control's measured false-rejection rate is 0.108 against a nominal 0.05, and a demand-characteristic confound with a true coefficient of zero passed the whole four-baseline suite 52.7 per cent of the time [Curzi, 2026u]. **Those numbers belong beside the suite every time it is invoked, and they are the reason a passed suite is not a proof.**

The four verdicts:

Predicted vs observed	Verdict	What to do
observed within tolerance, baselines beaten	**GROUNDED**	report the gain; proceed to FVP-7
observed same sign, well **below** the predicted magnitude	**GAIN OVERSTATED**	the channel is real and smaller. Revise step 4 and re-run. **This is Class B and it is the most common honest outcome.**
observed same sign, well **above** the predicted magnitude	**GAIN UNDERSTATED**, or a second channel is operating	the channel is richer than modelled, or another was not named. **This is a finding, not a success.** Go back to step 2.
observed near zero, or the shuffle predicts as well	**UNGROUNDED (bounded)**	report the upper bound on the gain at the study's power. **A null bounds a gain; it does not prove a zero.**

"Gain understated" is not a win. A result larger than predicted means the model of the channel was wrong, and a wrong model that happened to under-predict is not better than one that over-predicted. **The protocol treats surprise in either direction as a failure of step 2, and that symmetry is what stops it becoming a machine for confirming things.**

35.6.13 6.13 The shuffled-input control, simulated — and the case where it fails

The shuffled-input control is the most important of the four baselines for belief claims, because belief series and outcome series both drift and drift alone will produce a correlation. Shuffling the belief series destroys its temporal alignment with the outcome while preserving its distribution.

Simulated here, 4,000 trials per case, series length 200, fitted on the first 140 and tested on the remaining 59, with the model $Y_{t+1} = aY_t + bX_t + \varepsilon$ and X an autoregressive belief series:

CASE	ordered RMSE	shuffled RMSE	ordered wins
1 real coupling, b_true = 0.60, no common cause	1.0099	1.2722	99.53%
2 real coupling, b_true = 0.15, no common cause	1.0119	1.0314	74.10%
3 NO coupling, b_true = 0.00, no common cause	1.0126	1.0110	48.17%
4 NO coupling, b_true = 0.00, SHARED trend	1.1744	1.5602	98.35%

Case 3 is the control working exactly as intended: with no coupling the shuffle predicts as well as the ordered series — 48.17 per cent, indistinguishable from a coin flip.

Case 2 is the honest warning about sensitivity: with a real but modest coupling, the ordered series beats the shuffle only 74.10 per cent of the time on a single run. **A single application of the control, on a real-sized effect, is not decisive.**

Case 4 must be printed, and it is the case that keeps this section honest. There is no coupling at all — the belief causes nothing — and the ordered series beats the shuffle in **98.35 per cent of runs**, because a shared slow driver moves both series and shuffling destroys alignment with the common cause exactly as it destroys alignment with a real effect. **The control cannot tell those two apart.**

This is the identification problem of Section 3.3 in simulated form, and it is not fixed by shuffling. The shuffled-input control kills "there was never any temporal structure". **It does not kill "something else moved both"** [Manski, 1993, Shalizi and Thomas, 2011]. The defences are structural, not statistical:

```
CC-1  RANDOMISE THE BELIEF. If the belief is assigned rather than observed, no
      common cause can precede it. THIS IS THE ONLY COMPLETE DEFENCE, and it
      is the remedy this protocol adopts.

CC-2  MEASURE THE SUSPECTED COMMON CAUSES and include them. Partial, and only
      as good as the enumeration.

CC-3  BLOCK ON TIME. Compare within short windows where the slow driver is
      approximately constant. Costs power; costs nothing else.

CC-4  PREDICT A MAGNITUDE (step 5). A common cause will generally not deliver
      the specific magnitude the channel model predicts. THE MAGNITUDE
      PREDICTION IS ITSELF A CONFOUNDING CONTROL, and this is a second and
      independent reason step 5 is not optional.
```

35.6.14 6.14 The independence rule

```
STEPS 1 THROUGH 5 MUST BE COMPLETED, IN WRITING, WITH A DATE, BEFORE
STEP 6 BEGINS.

No gain estimate may come from the case under test.
No tolerance may be adjusted after data collection begins.
No channel may be added after the outcome is seen.
No verdict may be reported without the baselines.
```

That is the entire answer to the circularity objection, and it is procedural rather than philosophical. The conditional postulate is meaningful only if "grounded" is decided before the outcome is known. The worksheet of Appendix C exists to make that the path of least resistance: it has a date field at the top, the prediction boxes come before the measurement boxes, and it is meant to be filled in with a pen so that revision leaves a mark.

35.6.15 6.15 A gallery: six beliefs at protocol speed

```
1. "If I believe the interview will go well, it will go better."
   Step 2: SOCIAL + BEHAVIOURAL. Transducer: facial musculature, vocal
     prosody, preparation time. Named, real.
   Step 3: passes all six checks.
   Step 4: chain gain, three links, plausibly 0.7 x 0.6 x 0.5 = 0.210000.
   Step 5: state the magnitude in offer-probability points. Appendix B gets
     0.014000 -> 0.069650 under assumed parameters.
   VERDICT: PROCEEDS. Testable, modest, worth measuring.

2. "If I believe hard enough, the payment will arrive."
   Step 2: no transducer connects the belief to the sender, who is not
     present and does not know.
   Step 3: fails P3 and P6.
   VERDICT: DIES AT STEP 3. Class A.
   NOTE: if the belief causes the claimant to invoice, chase, or ask, the
   claim becomes BEHAVIOURAL, gain is large, and it proceeds -- as a
   different claim, with a different channel, honestly relabelled.

3. "Team morale drives our delivery speed."
   Step 2: BEHAVIOURAL + SOCIAL, several links.
   Step 4: gain estimable from the team's own history.
```

```
   Step 5: the team predicts "faster." No magnitude offered.
   VERDICT: DIES AT STEP 5. Not refuted -- UNTESTED.

4. "Visualising the outcome each morning improves it."
   Step 2: PHYSIOLOGICAL + BEHAVIOURAL, via attention allocation and
     rehearsal. A real family.
   Step 3: passes.
   Step 4: the literature on mental rehearsal is active and its magnitudes
     are contested; this paper does not assert a number it has not checked.
   VERDICT: PROCEEDS TO STEP 4 with the gain UNESTIMATED and labelled as
   such. This is a legitimate resting state and it is not a verdict.

5. "An acoustic signal can address a specific genomic locus."
   Step 2: source, transducer, medium and receiver all ESTABLISHED.
     Effector selectivity NOT established.
   Step 3: passes P1-P5; P6 is not applicable.
   Step 4: the final link's gain is unmeasured. This is the framework's one
     well-specified open problem.
   VERDICT: STEP 4 INCOMPLETE. Not false; unestablished, at a named link.

6. "Believing the market will rise will make it rise."
   Step 2: INFORMATIONAL. Real channel -- buying does raise prices.
   Step 3: passes.
   Step 4: the individual's share is 1/N. At N = 100,000 the shortfall
     against the 25% coalition is 25,000x.
   VERDICT: DIES AT STEP 4 for an individual. PROCEEDS for a coalition
     above the tipping point. Class D2 -- THE ONLY FAILURE MODE THAT
     CHANGES WITH SCALE.
```

Two of six died at step 3. One died at step 5. One is parked at step 4 with the gain honestly unestimated. One is the framework's open problem. One proceeds. A protocol returning "proceeds" six times out of six would be a formality rather than a filter [Goodhart, 1975].

35.6.16 6.16 Requirements

REQ-6.1 Complete steps 1 to 5 in writing, with a date, before measurement begins. **Without this the postulate is circular and the framework disowns it.**

REQ-6.2 Write the negation. If the negation is vague, the belief was vague.

REQ-6.3 Name all five parts of the channel. "Somehow" is not a link.

REQ-6.4 Run P3 and P6 first. They are the cheapest checks and they end more claims than the other four combined.

REQ-6.5 Estimate every link gain from a source that did not observe the outcome under test.

REQ-6.6 Multiply along the chain and print the total. Five links at 0.60 give 0.077760. Report that, not the best link.

REQ-6.7 Shorten the chain before strengthening the links: 2.78x against 2.16x, and usually cheaper.

REQ-6.8 Carry the gain uncertainty into the tolerance. Five links each known to a factor of 1.5 compound to 7.59.

REQ-6.9 State a magnitude with a tolerance, in units, before measuring. A direction-only claim is worth a likelihood ratio of 2.0000; a ± 20 per cent magnitude call against a two-decade prior is worth 11.3577, which is 3.5056 direction calls for the same single measurement.

REQ-6.10 Compute the required sample at step 5 and redesign at step 5 if it is unobtainable: n per arm $= 15.6978/d^2$.

REQ-6.11 Instrument the link with the highest event rate, not the endpoint of greatest interest.

REQ-6.12 Print all seven validation gates with their verdicts. Where FVP-2 fails, demote the relation to a monotone ordering in print. **Do not multiply ordinal quantities.**

REQ-6.13 Report against all four baselines including the shuffled-input control, **and report the suite's own measured false-rejection rate of 0.108 against a nominal 0.05 alongside it** [Curzi, 2026u].

REQ-6.14 Do not treat the shuffled-input control as a defence against common causes. Simulated, it lets a pure common-cause configuration through 98.35 per cent of the time. **Randomise, block on time, or use the magnitude prediction as the confounding control.**

REQ-6.15 Return one of four verdicts, and treat a larger-than-predicted result as a step-2 failure.

REQ-6.16 Report every null as an upper bound on the gain at the study's power.

REQ-6.17 File completed worksheets, **including the ones that died at step 3. The sheets that died are the denominator**, and a corpus without a denominator is worth a likelihood ratio of 1.000000.

REQ-6.18 Make no clinical, medical or therapeutic claim. Absolute.

35.6.17 6.17 Failure modes of the protocol

FM-6.1 — EX POST GROUNDING. *Signature:* the channel is identified after the result is in, and it always fits. *Fix:* the date field, and it only works if it is written before.

FM-6.2 — DIRECTION-ONLY CLAIM. *Signature:* "it will help", "there will be an effect". *Fix:* if the claimant cannot produce a magnitude, record the verdict as **UNTESTED**, which is a different and more accurate verdict than unproven.

FM-6.3 — THE UNNAMED TRANSDUCER. *Signature:* the channel described in verbs — "it influences", "it resonates", "it aligns" — with no noun that could be pointed at. *Fix:* demand a noun. Muscle, voice, face, keystroke, ion channel, order book.

FM-6.4 — BORROWED GAIN. *Signature:* an effect size imported for a channel, population or dose the source study did not examine. *Fix:* record the source of every gain estimate, including the population and dose it was measured at.

FM-6.5 — CHAIN-GAIN NEGLECT. *Signature:* the strongest link reported as the channel's gain. *Fix:* print the cumulative column, not the maximum.

FM-6.6 — TOLERANCE SET AFTER THE FACT. *Signature:* the observed result described as "consistent with" the prediction, with the tolerance appearing in the same sentence. *Fix:* the tolerance goes in box 5, before box 7 exists.

FM-6.7 — UNDERPOWERED NULL READ AS ABSENCE. *Signature:* a small study finds nothing and the claim is declared dead. *Mechanism:* at 50 per arm the minimum detectable effect is 0.5603, and almost nothing in this domain is that large. *Fix:* report the bound, not the verdict.

FM-6.8 — SHUFFLE OMITTED, OR SHUFFLE OVERTRUSTED. *Signature:* either the control is skipped, or it is passed and the result is called causal. *Mechanism:* the control catches absent temporal structure (Case 3, 48.17 per cent) and misses common causes entirely (Case 4, 98.35 per cent). *Fix:* run it, and add CC-1 to CC-4. **Randomisation is the only complete defence.**

FM-6.9 — THE UNFALSIFIABLE GOAL. *Signature:* the belief concerns an outcome with no date, threshold or instrument. *Fix:* the four passes of Section 6.5. If the claimant declines to sharpen, that is a decision to hold the belief outside the protocol, which is permitted — **provided it is not then described as tested.**

FM-6.10 — ORDINAL MULTIPLICATION. *Signature:* a five-point agreement scale multiplied by a gain to produce a predicted effect. *Fix:* demote to a monotone ordering in print, or build a ratio-scale instrument. Both are honest; silently multiplying is not.

FM-6.11 — ENDPOINT VANITY. *Signature:* the study designed around the outcome that matters most, and never run because the event rate is too low. *Fix:* instrument the high-rate intermediate link and carry the endpoint inference through the measured slope.

FM-6.12 — WORKSHEET THEATRE. *Signature:* the sheet filled in completely and after the fact, in one sitting, with every box agreeing with every other. *Fix:* the date and signature at step 5, separated in time from step 7. **A sheet whose steps 1 to 5 and step 7 carry the same date has documented nothing.**

35.7 7 Results

35.7.1 7.1 R1 — the four failure classes

Failure decomposes into four distinct classes, each with a stated diagnostic: **A**, no channel — name the transducer, and if none can be named the analysis is finished; **B**, a real channel with insufficient gain — the measured effect is real, replicable and an order of magnitude below the claim; **C**, a real result extrapolated to an unsupportable magnitude — the claim's size traces to no measurement but to an inference from the existence of the effect; **D**, a system unresponsive to an input of this size from a source of this size — the channel is fine and the receiver is not listening, or is listening to a chorus in which this voice is $1/N$. **They are not degrees of one failure and they prescribe incompatible responses.** Class D subdivides further into D1 (the receiver does not read this input at all), D2 (it reads it and this source is $1/N$), and D3 (it reads it and is already saturated), **of which only D2 changes with scale.**

35.7.2 7.2 R2 — growth mindset as Class B, unpacked

The channel is real and the gain is small. The meta-analytic correlation is near $r = 0.10$, which is **1.00 per cent of variance**, and intervention effects are near $d = 0.08$, which unpacks to **96.809 per cent distribution overlap**, a **52.256 per cent common-language effect** against 50 per cent by chance, **53.188 per cent** of treated cases above the untreated median, and **0.0560 GPA points** on a 0.70-point standard deviation. The national field trial's subgroup effect of about a tenth of a grade point is $d = 0.1429$ on that same standard deviation. **A small effect at national scale and near-zero cost is good engineering; it is not what "belief creates reality" is usually taken to mean, and the gap between the two is roughly one order of magnitude.**

35.7.3 7.3 R3 — the job-search audit

Running every real channel at full strength lifts the offer probability from 0.0140 to 0.0697 — a ratio of 4.975 — and the outcome is still no offer 93 per cent of the time. 39.53 per cent of the entire lift comes from the belief causing the application to be submitted, which is the most overlooked channel in the literature and the largest. **The no-contact channel — belief acting on the employer's decision with no communication of any kind — carries gain zero by construction**, and is Class A. **All parameters in the decomposition are assumed and are labelled as assumed**; what the case demonstrates is the decomposition procedure, not the values.

35.7.4 7.4 R4 — the deposit run, priced

At a 9.00 per cent reserve ratio the run threshold is $f^* = 0.0900$ — 216 of 2,400 members at the average balance — and the critical forced-sale recovery rate is $h^* = 0.941379$, so a recovery below 94.14 cents on the dollar makes a solvent institution insolvent. With member tolerances distributed Normal$(\mu, 0.25)$ the critical mean is $\mu^* = 0.408522$, and the withdrawn fraction moves from **0.990106 at $\mu = 0.407522$ to 0.145087 at $\mu = 0.409522$ — a swing from a**

99.01 per cent run to a 14.51 per cent tremor on a change of 0.002 in an unmeasured population parameter, with no intermediate outcome. This is a property of the threshold model at the stated parameters, not a measurement of any institution.

35.7.5 7.5 R5 — the runaway terminates at the ceiling, and slows before it does

With $\delta = 0.10$ and $\eta = 0.50$ the critical reflexive slope is $s^* = 0.20$, but $s = 0.19$ already implies a steady-state input of 12.6 nats against $\ln N = 2.995732$ at $N = 20$. Every divergence row beyond $s = 0.15$ is outside the model's admissible region for any realistic N: **the runaway does not run away, it runs into the ceiling** [Curzi, 2026]. **And the time constant $\tau = 1/(1-a)$ goes tenfold — from 10.0000 to 100.0000 periods — at a loop gain of 0.90, so the loop slows conspicuously before it destabilises.** Critical slowing is therefore a measurable early-warning signature available from response time alone, **which matters precisely because the reflexive slope itself is not identified from observational data.**

35.7.6 7.6 R6 — evidential pricing

A direction-only prediction is worth a likelihood ratio of 2.0000. Under a log-uniform prior spanning two decades, **a ± 20 per cent magnitude call is worth 11.3577**, which is $\ln(11.357747)/\ln 2 = 3.5056$ direction calls **for the same single measurement. Five links at 70 per cent each deliver 6.3 per cent end to end**, and a five-link channel at 60 per cent each delivers 0.077760. **The likelihood-ratio figures are properties of the stated log-uniform prior, which is a modelling choice and is badged as one.**

35.7.7 7.7 CR — the counter-result

The paper is itself the counter-result. It catalogues where the framework's own postulate does not pay; it prices the surviving effects at $d \approx 0.08$ to 0.30 with **92 to 97 per cent distribution overlap**; it shows that at $N = 100{,}000$ a single committed person is **25,000-fold short** of the tipping requirement; and it states that **a null bounds a gain without proving a zero** — with 600 per arm the bound is $d = 0.1617$, which does not exclude either of the two effects this paper reports. **It applies the replication discount to the framework itself: a prior of 0.10 becomes a posterior of 0.1148 on one significant result under typical power and heavy selective reporting.** And it states, in Section 5.2, the objection that would destroy the postulate outright, offering only a procedural answer.

35.7.8 7.8 Hypotheses, resolved

	Statement	Outcome
H1	the taxonomy is exhaustive over documented cases	**no counterexample found among the cases examined, which is weak evidence.** The cases were selected by the author; Section 10 records that as a selection problem
H2	the most-cited belief effect is small, not absent	**confirmed** at $r \approx 0.10$ and $d \approx 0.08$, 96.809 per cent overlap
H3	the responsive case is knife-edged, not proportionate	**confirmed in the model**: 0.990106 to 0.145087 across $\Delta\mu = 0.002$

	Statement	Outcome
H4	the runaway terminates at the ceiling	**confirmed by arithmetic**: 12.6 nats implied against $\ln 20 = 2.995732$
H5	a magnitude call is worth several direction calls	**confirmed under the stated prior**: 11.3577 against 2.0000, or 3.5056 direction calls

H1's outcome is reported in the same plain language a confirmation would receive, and it is not a confirmation. A taxonomy tested only against cases its author chose has been illustrated, not tested. **The honest verdict on H1 is that it survives an unblinded, non-exhaustive audit conducted by the person who proposed it**, and Section 10 gives that its own limitation.

35.7.9 7.9 Model comparison: what `RIGOR_STANDARD` §6 requires

This paper fits no model to any collected data. One simulation is run (Section 6.13) and its purpose is to characterise a control, not to fit a model to observations. The row-by-row status is printed rather than omitted.

Requirement	Status here	Why
out-of-sample performance	**NOT APPLICABLE** to the analytical results; **MET in form** for the simulation	the simulated control is fitted on 140 points and tested on 59 held out, which is the design; **no real data were used**
matched null with equal free parameters	**MET for the simulation**	Case 3 is exactly the matched null — the same model with $b_{\text{true}} = 0$ — and the ordered series wins 48.17 per cent of the time, indistinguishable from chance
information criterion (AIC or BIC)	**NOT MET, and not computed anywhere in this corpus**	no likelihood is reported here, and no AIC or BIC value appears in any paper of this corpus. **The omission is corpus-wide and is recorded rather than papered over**
explicit free-parameter count	**MET**	the simulation's model $Y_{t+1} = aY_t + bX_t$ has **two** free parameters, the same two in every case including the null; the effect model $E = B \cdot K \cdot R$ has **three**, all of which are inputs a user supplies rather than parameters this paper estimates

Requirement	Status here	Why
effect sizes with intervals	**PARTIALLY MET, and the gap is real**	the effect sizes quoted from the literature ($r = 0.10$, $d = 0.08$, $d = 0.1429$) are **reproduced as point values without intervals**, because that is how the source framework carried them. **A measured quantity quoted without an uncertainty is a presentational defect** and is recorded as Limitation L-9
identification of the reflexive parameter	**NOT MET, and not achievable from observation**	[Manski, 1993]; Section 3.3. The protocol adopts randomisation (CC-1) as the remedy **for prospective work**; nothing retrospective in this paper is identified
power-law procedure [Clauset et al., 2009]	**NOT RUN, deliberately**	the two exponents here are analytic properties of known functions; Section 3.5
conjoint measurement test [Luce and Tukey, 1964]	**NOT RUN**	Section 3.4; the exposed constructs are demoted instead

A paper that fits nothing cannot be beaten on fit, and earns nothing on it either. What is claimed here is a taxonomy, a set of exact arithmetic consequences, and a procedure — **and the procedure's own value is untested, because it has not been run on a real case.**

35.8 8 Discussion

35.8.1 8.1 What follows

The postulate's territory is bounded and the boundary is computable. Where a channel exists and the system is responsive, belief can be decisive — the deposit run reaches 99.01 per cent withdrawal on a parameter shift of 0.002. Where a channel exists and the gain is small, the effect is real and roughly an order of magnitude below the popular claim — 96.809 per cent distribution overlap. Where no transducer can be named, the analysis is finished in minutes and the answer is no. **Those three verdicts are reached from the channel, before the outcome, which is the only order in which the postulate says anything at all.**

The intervention hierarchy is the practical payoff and it inverts the usual instinct. W3 first, W2 second, W1 last: **structure beats persuasion, because structure does not have to be re-persuaded every morning.** Appendix A prices that exactly — communications moves the run threshold by nothing, a withdrawal gate moves it from 0.0900 to 0.1500, and a committed liquidity line moves it to 0.5900, a 6.56-fold increase in the withdrawal fraction the institution can absorb. The corpus reaches the same conclusion for physical rooms, where an acoustic exclusion is priced against the surplus law at 0.4700 nats, 22.6 per cent of the room's per-interaction surplus, and the remedy is geometry rather than exhortation [Curzi, 2026e].

Critical slowing is a usable signature and it does not require the unidentifiable parameter. $\tau = 1/(1-a)$ rises tenfold before the pole reaches unity, and response time to a small deliberate perturbation is measurable without knowing s. **That matters more than it looks: since s is not**

identified from observation, an indirect signature is not a convenience but the only route available.

And the evidential arithmetic converts a rhetorical dispute into a costed one. "Be more specific" is a demand; "that costs 3.5 direction-calls' worth of evidence and here is the study that would supply it" is a specification. **A predicted effect of $d = 0.08$ requires 4,906 participants; a predicted effect of $d = 0.50$ requires 126. A small predicted effect is not a modest claim, it is an expensive one**, and a claim that refuses to state its magnitude has also refused to say what testing it would cost.

35.8.2 8.2 What does not follow

Nothing here shows that belief does nothing. The catalogue is a boundary and a boundary has two sides; Section 5.14 is the other side, and every survivor on it carries a named physical channel and a stated magnitude.

Nothing here identifies the reflexive parameter. The loop gain, the reflexive slope, and every regime classification derived from them are modelling positions [Manski, 1993, Shalizi and Thomas, 2011]. **The simulated control's Case 4 — a pure common cause passing at 98.35 per cent — is the quantitative demonstration that no statistical control substitutes for identification, and it is printed rather than buried.**

Nothing here licenses a clinical, medical or therapeutic claim, in either direction. Published trials are reported as evidence about channels and their gains. This is absolute.

Nothing here establishes that the taxonomy is exhaustive. It was applied to cases the author selected, by the author, without blinding. **That is illustration, not test**, and H1's verdict says so.

And nothing here validates the protocol. Section 6 specifies a procedure that has not been run on a real case. **A procedure that has never been executed has no measured pass rate, and a protocol whose pass rate is unknown could still be a formality** [Goodhart, 1975, Campbell, 1979]. The corpus's own attempt to run a comparable protocol against a real system **failed its second falsification rule** — a 2.8 per cent held-out advantage against a declared margin of 10 per cent — and that result is the realistic expectation for this one [Curzi, 2026i].

35.8.3 8.3 Relation to the surrounding literature

The positive results sit inside three established literatures and add to none of them. The mechanism is Merton's [Merton, 1948, 1936]; the threshold model is Granovetter's [Granovetter, 1978] and its neighbours' [Schelling, 1971, Watts, 2002]; the tipping measurement is Centola and colleagues' [Centola et al., 2018, Centola, 2010]; the survivorship argument is Wald's; the replication estimate is the Open Science Collaboration's [Open Science Collaboration, 2015]. **What is new is the decomposition, the taxonomy, the arithmetic, and the requirement that grounding be established from the channel.**

The methodological position is closer to the critics of physics-in-social-science than to its enthusiasts. Predictive accuracy in social systems is low and the ceiling appears structural rather than methodological [Hofman et al., 2017, Salganik et al., 2020, Watts, 2014]; published prediction results have repeatedly failed to survive contact with new data [Lazer et al., 2014]; and complexity-flavoured social modelling has a long record of formal elegance without identified parameters [Durlauf, 2005]. **This paper's contribution to that argument is deflationary and is meant to be.**

Where the framework's postulate needs an existence proof for reflexivity, the performativity literature supplies one that was observed in the field with dates and price data [MacKenzie and Millo, 2003, MacKenzie, 2006, Ferraro et al., 2005, Callon, 1998], and the looping of human kinds under classification is the same mechanism in another register [Hacking, 1995]. **None of them identifies the parameter, and citing them as existence proofs rather than as measurements is the distinction this paper insists on.**

35.8.4 8.4 Parallel results elsewhere in the corpus

Four results in other series are structurally parallel and are named because the parallel is the point.

A gate that never rejects is not a gate. Section 6.15's gallery kills four of six claims, and the corpus's own audit of a traditional operation set against a parameter model returns the same shape of distribution — **seven MAPS, two SPLIT, one RENDERING and three REJECTED across thirteen entries, with five entries carrying a failure** — which is what a discriminating instrument looks like when it is turned on material its author is sympathetic to [Curzi, 2026].

A demotion made against one's own instrument is worth more than a demotion made against someone else's. Section 3.4 demotes every multiplicative statement involving an ordinal belief measure. The corpus's minimum-sustainable-frequency result carries the analogous concession from the other direction: **the deficit is blind to decay, so the intuitive intervention of lowering loss moves the instrument by exactly zero, and the reference constant, the affine effective count and the choice of twelve components are all badged as chosen rather than derived** [Curzi, 2026]. **This paper's own postulate is bounded in exactly that way.**

An effective count is not a headcount, and the corpus computes both. Where this paper's Class D2 turns on an individual's share of an aggregate, the corpus's compartment result computes the same quantity as a Hill number of order two: **six unequal compartments (30/22/18/12/10/8) give 4.96 effective compartments and cut maximum single reach from 100.0 per cent to 30.0 per cent, worth 1.400 units a year against 0.110 of cost — 12.73:1, with a break-even hazard of 0.157 per cent a year** [Curzi, 2026]. **That is the arithmetic of a share, done properly, in another domain.** The roster-scale identity that the corpus's affine effective count equals $m+1-m/\mathrm{PR}$ with PR the participation ratio, verified at 1.7815458788 with error 0.00×10^{0}, is the same construction stated exactly [Curzi, 2026] — **and it is the corpus's own construction, not a Hill number** (Section 3.2).

A metric that under-counts what it cannot see is a cost line, not a resilience line. The corpus demotes its own legal-family diversity measure on exactly that ground — **1.862031 effective families, stated to under-count the interfaces of every carve-out jurisdiction** [Curzi, 2026] — and this paper's taxonomy is exposed to the same objection: **a class assignment can only be made for a channel someone thought to name.**

35.8.5 8.5 What the framework's own failure record adds

The corpus records that **disclosure did not confer immunity in the strongest available test case**, that a backing basket was **4.61 times less diversified than its component count claimed**, that the terminal phase of a programme was **a 0.090 option rather than a plan**, and that a levy presented as a costless accounting correction **in fact reallocated burden, with 42.0 per cent falling on one sector and a constructed 55 per cent passing through to consumers** [Curzi, 2026]. **Those are the framework's own recorded failures, and they bound this paper in the way Section 5.11 says a denominator bounds a numerator: they are the tail the success cases do not supply.** A paper cataloguing where a postulate fails is only credible if the framework it belongs to keeps such a record, and this one does.

35.9 9 Verification

The results compound from ten postulates. Each is atomic, each is checkable on its own, and each carries its status.

35.9.1 9.1 Dimensional audit

$[1]$ denotes a pure number. **Where the framework fixes no dimension, the audit says so rather than assigning one.**

Symbol	Dimension	Unit	Status
B (belief fraction)	`[1]`	conviction-weighted fraction of agents	**defined** dimensionless on $[0, 1]$; **the underlying belief measure is ORDINAL and is not licensed for multiplication** (Section 3.4)
K (channel transmission)	`[1]`	none	**defined** dimensionless; a modelling choice
R (responsiveness)	`[1]`	none	**defined** dimensionless; a modelling choice
$E = BKR$	`[1]`	none	**derived**; balances, since all three factors are pure numbers
g (loop gain)	`[1]`	none	**derived** as $\eta s/\delta$; **not identified from observational data**
$s = dS/dQ$	`[\text{nats}]/[q]`	nats per proxy unit of store	**the framework is silent on the unit of** Q, so s inherits an unstated dimension. It is reported here as **undetermined**, not assigned
$a = (1 - \delta) + \eta s$	`[1]`	none	**derived**; balances only if ηs is dimensionless, which requires η to carry `[q]` per nat per period exactly [Curzi, 2026a]
$\tau = 1/(1 - a)$	`[1]`	periods	**derived**; a count of periods, dimensionless given a fixed period
Q	`[q]`	proxy units of stored order	**measured** in whatever the declared proxy counts
η	`[q]` per nat per period	proxy units per nat per period	**measured**; *not* dimensionless [Curzi, 2026a]
δ	`[1]`	per period	**measured** dimensionless fraction per step

Symbol	Dimension	Unit	Status
S, $\ln N$	`[1]`	nats	**derived** dimensionless; a **Shannon-type logarithm of a count, not a thermodynamic entropy**
r (reserve ratio), f (withdrawal fraction), h (recovery rate)	`[1]`	none	**derived** ratios of like quantities; ratio-scaled with true zeros
D (deposits), reserves, shortfall	`[\$]`	currency	**measured**, ratio-scaled
μ, σ (threshold distribution)	`[1]`	fraction of others withdrawing	**defined** dimensionless on the same axis as x
x^* (fixed point)	`[1]`	fraction of deposits withdrawn	**derived** dimensionless
d (standardised effect)	`[1]`	standard deviations of a **stated** reference population	**quoted** from the literature; **the reference SD must be stated with it or the number means nothing**
r (correlation)	`[1]`	none	**quoted**; r^2 is the variance share
GPA points	`[\text{grade}]`	grade points	**derived** from d by multiplying by a stated SD of 0.70
likelihood ratio	`[1]`	none	**derived**; a ratio of two probabilities of the same event
g_{total} (chain gain)	`[1]`	none	**derived**; dimensionless only if every link gain is dimensionless, which the protocol requires to be stated
n (sample per arm)	`[1]`	count	**derived** from d
u, $B^2/2\mu_0$	`[M L^{-1} T^{-2}]`	J m^{-3}	**derived** exactly
Casimir $P(a)$	`[M L^{-1} T^{-2}]`	Pa	**derived** exactly from $\pi^2\hbar c/(240a^4)$

Does the effect model balance? $E = B \cdot K \cdot R$ balances trivially: all four quantities are pure numbers. **That is not a virtue — a model of pure numbers balances whatever it means — and the real content of gate 1 here is the next row.**

Does the reflexive storage law balance? $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$ balances only if η carries `[q]` per nat per period, since S is dimensionless in nats [Curzi, 2026a]. The reflexive term $\eta s Q$ then requires s to carry nats per `[q]`, **so that ηs is dimensionless and the pole a is a pure number. The framework does not state the unit of Q**, so s has no operational unit and the audit records it as **undetermined. A dimensionless placeholder is honest; an invented unit would not be.**

Does the effect-size arithmetic balance? d is dimensionless only against a stated reference standard deviation, and the conversion to 0.0560 GPA points is 0.08×0.70. **Quoting a d without**

its reference SD is the most common dimensional error in this literature and it is why every conversion here prints the SD in the same line.

Where the framework has no operational unit, it is said here. Belief, conviction and confidence have **no operational unit at all**. B, K and R are dimensionless placeholders for quantities that have not been instrumented, and the scale on which belief would be instrumented is **ordinal at best** until a conjoint test is run [Michell, 1997, Luce and Tukey, 1964].

35.9.2 9.2 The postulates

P1 — For $|g| < 1$, $\sum_{k\geq 0} g^k = 1/(1-g)$. *Status:* **ARITHMETIC.**

P2 — A first-order linear recurrence $Q_{t+1} = aQ_t + c$ converges iff $|a| < 1$, to $c/(1-a)$. *Status:* **CITED**, elementary, and presupposed from Curzi [2026a].

P3 — The storage law is $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$. *Status:* **DEFINITION** of the framework's model [Curzi, 2026a]. It carries no empirical weight here.

P4 — $S(Q) = S_0 + sQ$ is the linearisation of a reflexive input. *Status:* **MODELLING CHOICE.** Linearity is assumed about an operating point, and s **is not identified from observational data** [Manski, 1993].

P5 — $0 \leq F(u) \leq \ln N$. *Status:* **CITED**, from the framework's surplus functional [Curzi, 2026]. **It is a Shannon-type bound on a dimensionless quantity in nats and is not a thermodynamic statement.**

P6 — In a threshold model, the cascade advances from x while $F(x) > x$ and halts at the first fixed point $F(x) = x$. *Status:* **CITED** [Granovetter, 1978].

P7 — A deposit contract payable at par on demand, backed by reserves rD and illiquid assets, fails when the demanded fraction exceeds r. *Status:* **ARITHMETIC**, given the contract.

P8 — Distinguishing among M equally likely alternatives, or hitting a band of log-width w in a log-uniform prior of span W, has chance probability $1/M$ or w/W respectively. *Status:* **ARITHMETIC**, given the prior. **The log-uniform prior is a MODELLING CHOICE and every likelihood ratio in Section 6.3 inherits it.**

P9 — The minimum detectable standardised effect at power $1-\beta$ and level α in a two-arm comparison is $(z_{1-\alpha/2} + z_{1-\beta})\sqrt{2/n}$. *Status:* **CITED**, standard.

P10 — A corpus assembled by conditioning on success observes only the numerator. *Status:* **DEFINITION** of the selection, and the likelihood-ratio consequence is **ARITHMETIC**.

What is NOT a postulate. The four-class taxonomy is not a postulate in the chain below, because nothing is derived from it; it is a classification scheme, and it appears in the chain with an output of a verdict rather than a number. **That is the honest entry.**

35.9.3 9.3 The compounding chain

```
P1            -> L1   a loop of gain g < 1 multiplies a seed by 1/(1-g):
                      0.5 -> 2.0000, 0.9 -> 10.0000, 0.99 -> 100.0000.      ARITHMETIC

P3 + P4       -> L2   Q_{t+1} = [(1-delta) + eta*s] Q_t + (eta*S_0 - C),
                      so the pole is a = (1-delta) + eta*s.                 DEFINITION + CHOICE (P4)

L2 + P2       -> L3   stable iff eta*s < delta; loop gain g = eta*s/delta;
                      time constant tau = 1/(1-a).                          ARITHMETIC

L3            -> R5a  at delta = 0.10 and eta = 0.50 the critical slope is
                      s* = 0.20, and at s = 0.18 (g = 0.90) tau = 100.0000
                      against 10.0000 at s = 0.                             ARITHMETIC

L2 + P2       -> L4   Q_inf = (eta*S_0 - C)/(delta - eta*s), which at
                      s = 0.190 gives Q_inf = 60.0000 and an implied input
```

```
                      S(Q_inf) = 12.6000 nats.                              ARITHMETIC

L4 + P5        -> R5b  12.6000 nats exceeds ln N for every N below
                       e^12.6 = 2.97e5, and ln 20 = 2.995732, so the row is
                       outside the admissible region: THE RUNAWAY RUNS INTO
                       THE CEILING.                                          ARITHMETIC + CITED (P5)

P7             -> L5   the run threshold is f* = r; at r = 0.0900 and
                       D = 18,000,000 that is 1,620,000, i.e. 216 of 2,400
                       members at the average balance.                       ARITHMETIC

P7             -> R4a  h* = (D - cash)/loans = (18,000,000 - 1,620,000)
                       /17,400,000 = 0.941379.                               ARITHMETIC

P6 + L5        -> R4b  with thresholds N(mu, 0.25) the critical mean is
                       mu* = 0.408522, and x* = 0.990106 at mu = 0.407522
                       against x* = 0.145087 at mu = 0.409522.               ARITHMETIC (given P6)

P8             -> L6   one correct direction call has chance 0.500000, so
                       LR = 2.000000; five consecutive have chance 0.031250.  ARITHMETIC

P8             -> R6a  a +/-20% band is ln(1.2/0.8) = 0.405465 wide in log
                       space; against 2 decades = 4.605170 the chance hit is
                       0.088046 and LR = 11.3577.                            ARITHMETIC (given the prior)

L6 + R6a       -> R6b  ln(11.357747)/ln(2) = 3.5056 direction calls.         ARITHMETIC

P8             -> R6c  chain gain is the product of link gains: five links at
                       0.70 give 0.168070, at 0.60 give 0.077760, and the
                       worked five-link case gives 0.063000.                 ARITHMETIC

P9             -> R7   d_min = 2.801585 * sqrt(2/n): 0.5603 at n = 50,
                       0.1617 at n = 600, 0.0251 at n = 25,000; inverting,
                       d = 0.08 needs n = 2,453 per arm.                     CITED (P9)

P10            -> CR1  500 successes are produced by 1,000 practitioners at
                       p = 0.5 and by 500,000 at p = 0.001, so the likelihood
                       ratio between any two hypotheses given only the
                       numerator is exactly 1.000000.                        ARITHMETIC

(quoted)       -> R2   r = 0.10 gives r^2 = 0.0100; d = 0.08 gives U3 =
                       0.53188, CL = 0.52256, overlap = 0.96809, and
                       0.08 x 0.70 = 0.0560 GPA points.                      ARITHMETIC on a QUOTED d

(quoted)       -> R3   0.35 x 0.0400 = 0.014000 against 0.90 x 0.077389 =
                       0.069650, a ratio of 4.9750, with the A1 share of the
                       lift at 39.53%.                                       ARITHMETIC on ASSUMED inputs

(the taxonomy) -> R1   a verdict, not a number.                              CLASSIFICATION
```

Every line is `ARITHMETIC` except four. P4 is a modelling choice and R5a and R5b inherit it. P5, P6 and P9 are cited standard results. **R2's inputs are quoted effect sizes carrying the source literature's uncertainty and not this paper's**, and **R3's inputs are assumed and are labelled assumed at every appearance. R1 is a classification and produces no number at all**, which is stated here rather than left for the reader to notice.

A chain whose most consequential outputs are a demotion, a bound and a classification delivers withdrawals with certainty and delivers no measurement whatever. That is the honest summary of this paper's arithmetic.

35.9.4 9.4 The verification instance, worked

Four checks, none longer than five lines, and between them they establish the paper's four quantitative headlines.

Check one — the growth-mindset unpacking. With $d = 0.08$ and Φ the standard normal CDF:

```
U3      = Phi(0.08)               = 0.531881   -> 53.188%
CL      = Phi(0.08 / sqrt(2))
        = Phi(0.056569)           = 0.522555   -> 52.256%
overlap = 2 * Phi(-0.08/2)
        = 2 * Phi(-0.04)          = 0.968094   -> 96.809%
GPA     = 0.08 * 0.70             = 0.0560 points
r^2     = 0.10^2                  = 0.0100     -> 1.00% of variance
```

Check two — the critical recovery rate.

```
deposits D        = 18,000,000
cash              =  1,620,000     (r = 0.0900)
loan book         = 17,400,000
h* = (D - cash) / loans
   = 16,380,000 / 17,400,000      = 0.941379
```

so a forced-sale recovery of 0.90 realises $1{,}620{,}000 + 0.90 \times 17{,}400{,}000 = 17{,}280{,}000$, which is **96.00 cents on the dollar** against deposits of 18,000,000. **A ten per cent haircut on a small loan book is unremarkable, and it makes a solvent institution insolvent.**

Check three — the ceiling that terminates the runaway.

```
delta = 0.10, eta = 0.50, C = 0, eta*S_0 = 0.30
s = 0.190  ->  denominator = delta - eta*s = 0.10 - 0.095 = 0.005
               Q_inf = 0.30 / 0.005 = 60.0000
               S(Q_inf) = S_0 + s*Q_inf = 0.60 + 0.190 * 60 = 12.6000 nats
ln 20 = 2.995732,  and 12.6000 / 2.995732 = 4.21  times the ceiling
```

The implied input is more than four times the largest surplus a twenty-kind interaction can supply. The row is inadmissible, and every row above it is worse.

Check four — the evidential exchange rate.

```
direction call:   P(correct | null) = 0.500000        LR = 1/0.5 = 2.000000
magnitude call, +/-20%, prior span 2 decades:
    band  = ln(1.20 / 0.80) = ln(1.50)     = 0.405465
    span  = 2 * ln(10)                     = 4.605170
    chance hit = 0.405465 / 4.605170       = 0.088046
    LR    = 1 / 0.088046                   = 11.357747
equivalent direction calls = ln(11.357747) / ln(2) = 3.5056
```

Four checks. A reviewer who runs them has verified every headline this paper contributes that is not quoted from elsewhere.

35.9.5 9.5 Limits, asymptotics, and the domain of validity

$B \to 0$, $K \to 0$ **or** $R \to 0$ **in the effect model.** $E \to 0$ in each case. **That is the model's entire content and it is why it was chosen** — but it is also its weakness: a multiplicative model *forces* the conjunction rather than discovering it, and Section 10 records that as an assumption rather than a finding.

$g \to 1^-$ **in the loop.** The multiplier $1/(1-g) \to \infty$ and $\tau \to \infty$. **The divergence is not a pathology of the model but the definition of marginal stability** — and it is unreachable in practice because $S \leq \ln N$ binds first.

$g \to 0$**.** The multiplier tends to 1: the seed is the whole effect and there is no loop. **A system with** $g = 0$ **is exactly a non-reflexive system and the whole apparatus collapses correctly to nothing.**

$s < 0$ **— the self-defeating branch.** τ *shortens*: 8.0000 at $s = -0.05$ and 2.8571 at $s = -0.50$, against 10.0000 at $s = 0$. **A negative reflexive slope is a stabiliser and the limit is well behaved in both directions.**

$r \to 1$ **in the deposit contract.** $f^* \to 1$: the narrow bank, where no run is possible. $r \to 0$ gives $f^* \to 0$: any withdrawal at all exhausts the reserves. **Both limits are real institutional designs and neither is a defect of the model.**

$\sigma \to 0$ **in the threshold distribution.** Every member has the same threshold, so the cascade is all-or-nothing at μ. **The knife edge of Section 4.9 is the finite-**σ **version of that degeneracy, and the sharpness of the transition is governed by** σ**, which is the parameter institutions can actually move.**

$n \to \infty$ **in the power law.** $d_{\text{min}} \to 0$ as $n^{-1/2}$, so any effect becomes detectable eventually and the cost grows as the inverse square of the resolution wanted. $n \to 0$ **is not a limit of the model but a failure of the design**: below about 50 per arm the minimum detectable effect exceeds 0.56, which is larger than almost anything in this domain.

$p \to 1$ **in the log-uniform prior span.** As the prior narrows to a point the magnitude call's likelihood ratio tends to 1: **if the answer was already known, predicting it correctly is worth nothing.** As the span widens the ratio grows without bound, **which is the warning attached to R6: the evidential value of a magnitude call is a property of the prior, and a claimant who chooses the prior chooses the payout.**

Number of links $\to \infty$ **in the chain law.** $g_{\text{total}} \to 0$ geometrically for every per-link gain below 1. **A sufficiently long channel is a dead channel, and this is the asymptotic statement of "shorten the chain, do not strengthen the links."**

35.9.6 9.6 Invariance, and how far each claim has been tested

Scale invariance of the effect model — that E depends only on the product and not on which factor is small — is exact and algebraic, and it is **tested at five configurations** in Section 4.2. **Five configurations, all constructed.**

Invariance of the run threshold under the size of the deposit base is exact: $f^* = r$ does not depend on D, which is why the threshold is quoted as a fraction. **Tested at two deposit bases** (the thousand-depositor illustration and the constructed institution of Appendix A) **and at six reserve ratios.**

Invariance of the likelihood-ratio arithmetic under the units of the predicted quantity holds because the prior is log-uniform and the band is a ratio: rescaling the quantity shifts both by the same amount. **This is exact and definitional, and it is why the magnitude-call arithmetic is reported in per cent rather than in units.**

Invariance of the chain-gain law under the ordering of links is exact — multiplication commutes — and is the reason the protocol can require the *product* without requiring the links to be measured in sequence.

Parameter invariance across people, sites or institutions is the empirical invariance, and it has been tested at exactly zero points in this paper. FVP-7 is the gate that would test it and **it**

is reported as NOT YET RUN. Every invariance claim above is algebraic; none is an empirical finding.

35.9.7 9.7 Significant figures and their warrant

- **Exact by definition or by arithmetic, quoted to whatever precision is useful:** $h^* = 0.941379$; $\mu^* = 0.408522$; $\ln 20 = 2.995732$; the likelihood ratios 2.000000, 11.357747 and 3.5056; the chain-gain products; the survivorship ratio 1.000000; $d_{\min}$ at each n.
- **Illustrative model output at assumed parameters, rounded honestly:** the pole and time-constant table, the steady-state table, the threshold-model outcomes 0.990106 and 0.145087, and the posterior table of Section 5.9. **Every one of these is an evaluation of a stated model at stated assumed parameters and none is a measurement.** Six decimal places on 0.990106 record a model evaluation, not a measured fraction.
- **Quoted from published literature, carrying that literature's uncertainty and not this paper's:** $r = 0.10$; $d = 0.08$; the field subgroup effect of about a tenth of a grade point; 97 of 100 and about 36 of 100; $r = 0.403 \to 0.197$; 13 of 21; the roughly 25 per cent tipping threshold. **These are reproduced as point values without intervals because that is how the source framework carried them, and that is a presentational defect recorded at L-9.**
- **Constructed for the arithmetic and labelled constructed:** the balance sheet of Appendix A and every conditional probability in Appendix B.
- **Physical constants and exact formulas:** $\mu_0 = 4\pi \times 10^{-7}$, the Casimir pressures, the field-energy densities. **Exact given the formula.**

35.9.8 9.8 Analogy discipline, restated at the point of verification

Every borrowing carries the label assigned in Section 3.2 and no passage upgrades one. The loop gain and its geometric multiplier are a **shared mathematical form** with electrical feedback, and **no electrical result transfers**. The pole and stability condition are a **derivation** from the framework's own storage law. Critical slowing is a **shared mathematical form**, and the empirical early-warning literature is not invoked as support because no such measurement is made here. The Casimir computation is a **derivation within physics used as a counter-example**. And "the engine", "the amplifier" and "the seed" are **vocabulary, never evidence**.

Entropy does not appear in this paper. The only logarithm of a count is the surplus ceiling $\ln N$, which is a **Shannon-type dimensionless quantity in nats** and carries **no thermodynamic content whatsoever** [Curzi, 2026]. The word is not used loosely anywhere.

35.10 10 Limitations

L-1. The taxonomy was applied by its author, to cases its author chose, without blinding. H1 therefore survives an illustration rather than a test. **A taxonomy that fits every case its proposer selected is exactly what a taxonomy with no content would look like from the inside**, and the correct repair is an independent classifier applying it to a pre-registered case list. **That has not been done.**

L-2. A large number of the primary sources are named without citation keys. Thomas and Thomas (1928), Diamond and Dybvig (1983), Thorndike (1968), Elashoff and Snow (1971), Raudenbush (1984), Steele and Aronson (1995), Flore and Wicherts (2015), Flore et al. (2018), Sisk et al. (2018), Yeager et al. (2019), Camerer et al. (2018), Bem (2011), Ritchie, Wiseman and French (2012), Benson et al. (2006), Bösch, Steinkamp and Boller (2006), Levine, Gordon and Fields (1978), Kaptchuk et al. (2010), Tajik et al. (2016) and Wald (1943) have **no entries in this corpus's bibliography, and no key was invented for any of them.** Every one is a standard, locatable source; **the gap is in the corpus's bibliography and it is recorded here rather than smoothed over.**

L-3. No reflexive parameter in this paper is identified. The reflexive slope, the loop gain and every regime classification derived from them are **modelling positions rather than identified empirical results** [Manski, 1993, Shalizi and Thomas, 2011]. **The only remedy adopted is randomisation, and it is adopted only for prospective work in Section 6.13; nothing retrospective here is identified.** The stipulated-reference-group route is available and is not taken [Brock and Durlauf, 2001].

L-4. The multiplicative effect model is an assumption, not a finding. $E = BKR$ was chosen because it forces the conjunction. **An additive or a minimum-form model encodes a different conjunction and gives different intermediate values, and nothing in this paper establishes that the multiplicative form is right.** The value $E = 0.0600$ for the "weak on all three" configuration is an artefact of that choice.

L-5. The knife-edge result is a property of the model, not a measurement of any institution. $\mu^* = 0.408522$ and the 0.990106/0.145087 swing follow from a normal threshold distribution at $\sigma = 0.25$ and a 9 per cent reserve. **No threshold distribution has been measured for any real deposit base in this paper**, and the sharpness of the transition is governed by σ, which is exactly the parameter nobody has measured.

L-6. The conjoint-measurement test has not been run and Michell's objection stands. Belief, conviction and confidence are treated as ordinal throughout and every multiplicative statement involving them is demoted in print [Michell, 1997, Luce and Tukey, 1964]. **That is a concession, not a defence.** Until a ratio-scale instrument exists, the protocol's step 4 can produce a gain only for channels whose input is physically instrumented.

L-7. The protocol has never been run on a real case. Section 6 specifies a procedure with an unknown pass rate. **A protocol that nothing has ever failed is a formality**, and the only evidence this one is discriminating is the constructed gallery of Section 6.15, in which four of six claims die — **a gallery its author assembled.** The corpus's one comparable attempt against a real system **failed its own second falsification rule** [Curzi, 2026i], which is the realistic expectation.

L-8. The likelihood-ratio arithmetic is a property of a chosen prior. The log-uniform prior over R decades is a modelling choice, and the headline 11.3577 falls to 5.6789 at one decade and rises to 17.0366 at three. **A claimant who chooses the prior chooses the payout**, and the protocol has no defence against that beyond requiring the prior to be stated before the prediction.

L-9. Every effect size quoted from the literature is reproduced as a point value without an interval. $r = 0.10$, $d = 0.08$ and the field subgroup effect all carry sampling uncertainty that is not restated here, because the source framework carried them as point values. **A measured quantity quoted without an uncertainty is a presentational defect** and is recorded as one rather than dressed up.

L-10. Two of the three worked cases are constructed and the third is a gallery. The balance sheet of Appendix A, the conditional probabilities of Appendix B and the team history of Appendix C are all constructed for the arithmetic. **No institution, applicant or team described in them exists**, and what they demonstrate is a procedure.

L-11. No AIC or BIC is reported, here or anywhere in this corpus. Section 7.9 prints the row as NOT MET rather than omitting it. Nothing in this paper is fitted to collected data, so no information criterion is computable from it; **the corpus-wide absence is nonetheless a genuine gap and is recorded as one.**

L-12. The survivorship result is exact and its application is an inference. The likelihood ratio of 1.000000 is arithmetic. **The claim that the testimonial literature is assembled that way is an assertion about a body of publications that this paper has not audited**, and a corpus that reported its failures would not have the defect. **The result bounds what an outcome-conditioned corpus can show; it does not establish that any particular corpus is outcome-conditioned.**

L-13. Nothing here says a small effect is a worthless effect. A tenth of a grade point at national scale and near-zero marginal cost is good engineering, and the paper says so in the same section that reports the 96.809 per cent overlap. **The two statements are both true and**

reporting only one of them would be a distortion in whichever direction it was made.

35.11 11 Conclusion

This paper is the framework's counter-result: it catalogues where the conditional postulate does not pay, prices what survives, and states the objection that would destroy the postulate outright.

Result One, the four failure classes. Failure decomposes into **A**, no channel — gain is not small but undefined, and the diagnostic is to name the transducer; **B**, a real channel with insufficient gain — the effect is real, replicable, and an order of magnitude below the claim; **C**, a real result extrapolated to a magnitude it cannot support — nothing in the original work needs to be wrong; and **D**, a system unresponsive to an input of this size from a source of this size — which subdivides into a receiver that does not read this input, one for which this source is $1/N$, and one already saturated. **Only $1/N$ changes with scale, and all three present identically from inside.**

Result Two, growth mindset as Class B, unpacked. The channel is real and the gain is small: $r \approx 0.10$, which is **1.00 per cent of variance**, and $d \approx 0.08$, which is **96.809 per cent distribution overlap**, a **52.256 per cent common-language effect**, **53.188 per cent** above the untreated median, and **0.0560 GPA points** on a 0.70-point standard deviation. **Pick one treated and one untreated student at random and the treated one scores higher 52.256 per cent of the time, against 50 by chance. That is the size of the thing** — and a tenth of a grade point at national scale for pennies a student remains good engineering.

Result Three, the job-search audit. All real channels together lift the offer probability from **0.0140 to 0.0697**, a ratio of **4.975** — **and the outcome is still no offer 93 per cent of the time. 39.53 per cent of the entire lift comes from the belief causing the application to be submitted, and the no-contact channel carries gain zero by construction.** All parameters are assumed and labelled as such.

Result Four, the deposit run. At a **9.00 per cent** reserve the run threshold is $f^* = 0.0900$ — **216 of 2,400 members** — and the critical recovery rate is $h^* = 0.941379$, so **a forced-sale recovery below 94.14 cents on the dollar makes a solvent institution insolvent.** With tolerances Normal$(\mu, 0.25)$ the critical mean is $\mu^* = 0.408522$ and the withdrawn fraction swings from **99.01 per cent to 14.51 per cent on a change of 0.002**, with no intermediate outcome. **A property of the model, not a measurement of any institution.**

Result Five, the ceiling and the slowing. With $\delta = 0.10$ and $\eta = 0.50$ the critical reflexive slope is **0.20**, but $s = 0.19$ **already implies 12.6 nats of input against** $\ln N = 2.996$ **at** $N = 20$ — **4.21 times the ceiling, so the runaway leaves the admissible region before it diverges.** And $\tau = 1/(1-a)$ **goes tenfold at loop gain 0.90**: the loop slows conspicuously before it destabilises, which is a signature measurable from response time alone — **and, since the reflexive slope is not identified from observation, the only signature available.**

Result Six, evidential pricing. A direction-only prediction is worth a likelihood ratio of 2.0000. A ± 20 **per cent magnitude call against a two-decade prior is worth 11.3577, which is 3.5056 direction calls for the same single measurement. And five links at 70 per cent each deliver 6.3 per cent end to end.** The likelihood ratios are properties of the stated log-uniform prior, which is a modelling choice.

And the counter-result, which is what the paper is for. It catalogues where the framework's own postulate does not pay; it prices the surviving effects at $d \approx 0.08$ to 0.30 with **92 to 97 per cent distribution overlap**; it shows that at $N = 100{,}000$ **a single committed person is 25,000-fold short** of the tipping requirement; and it insists that **a null bounds a gain without proving a zero** — 600 per arm bounds only $d = 0.1617$, which excludes neither effect the paper reports. **It applies the replication discount to the framework itself, where a prior of 0.10 becomes a posterior of 0.1148 on one significant result. And it states the objection that would make the postulate vacuous — that a conditional whose antecedent is read off the outcome forbids nothing — and answers it with a procedure rather than an argument, because**

a procedure is the only thing that can answer it.

The postulate's whole content is its conditional, and the conditional's whole content is that grounding is established from the channel, before the outcome is known, or it is not established at all.

35.11.1 Declarations

Data and code availability

No external dataset was collected for this paper, and no human subjects were studied. The empirical magnitudes reported — the meta-analytic mindset effects, the national field-trial subgroup effect, the replication counts, the tipping threshold, the expectancy-effect syntheses and the trial results in Sections 5.6 and 5.14 — are quoted from the published literature named in the text, carry that literature's uncertainty rather than this paper's, and are not measurements of this paper; many of those sources carry no key in this corpus's bibliography and are named in full in the text, as disclosed at Section 3.8 and Limitation L-2. Every other number is either exact arithmetic, an evaluation of a model at parameters stated as assumed beside it, or — in the case of the shuffled-input control of Section 6.13 — the output of a stated simulation (4,000 trials per case, series length 200, 140 fitted and 59 held out) reproducible from the specification given. The balance sheet of Appendix A, the conditional probabilities of Appendix B and the team history of Appendix C are constructed and are labelled constructed wherever they appear. The figures are generated by the scripts named in each caption, and those scripts contain no data beyond the numbers printed in this paper and the functions stated in Sections 4 to 6.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest. The postulate audited here is the author's own, the taxonomy applied to it is the author's own, and the cases it is applied to were selected by the author. That is a structural conflict which no declaration removes; it is the reason failure mode FM-5.9 (selective scepticism) is in the specification, the reason Section 5.9 applies the replication discount to the framework in print, and the reason Limitation L-1 records the taxonomy as illustrated rather than tested.

Author contributions

Single author. Michael Laurence Curzi derived the results, performed the arithmetic and the simulation, produced the figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The trials, meta-analyses and field experiments discussed are other investigators' work, quoted. No clinical, medical or therapeutic claim is made anywhere in this paper, in either direction; the published trials in Sections 5.6 and 5.14 are cited as evidence about channels and their gains and never as guidance about care. The measurement designs of Section 6 and Appendix C are designs for prospective work and are not reports of studies performed; were any of them run, it would require ethical review, and nothing here substitutes for that.

Preregistration

This paper reports no preregistered test on human data. The five hypotheses of Section 2 were fixed before the arithmetic and the simulation reported in Sections 4 to 6 were carried out and are printed in the form in which they were fixed. **H2 to H5 are confirmed by arithmetic on quoted or assumed inputs; H1 is reported as surviving an unblinded, author-selected audit, which is illustration rather than test, and Section 7.8 says so in the same plain language a confirmation would have received.** The falsification suite whose baselines Section 6.12 requires has itself been audited and is **anti-conservative** — a shuffle-control false-rejection rate of **0.108 against a nominal 0.05**, and a demand-characteristic confound with a true coefficient of zero passing the whole four-baseline suite **52.7 per cent of the time** [Curzi, 2026u]. The corpus's preregistered confirmatory test of the storage law against a real system **failed its second falsification rule**, returning a 2.8 per cent held-out advantage over a matched autoregressive baseline against a declared margin of 10 per cent [Curzi, 2026i]. **Any applied claim resting on this paper inherits both results.**

35.12 Appendix A · A deposit run, an autopsy, and three interventions priced

The figures below are constructed for the arithmetic and are labelled as constructed; they are not drawn from a real institution. Every number is computed from those figures.

The balance sheet.

```
members                    2,400
deposits                   18,000,000
average member balance     7,500.00
cash and overnight         1,620,000        reserve ratio r = 0.0900 (9.00%)
loan book, at par          17,400,000
total assets at par        19,020,000
equity                     1,020,000        equity / assets = 0.0536 (5.36%)
```

By any static measure this institution is sound: solvent, with a 5.36 per cent equity cushion and nine per cent of deposits in cash. **It is also, structurally, a machine for converting a rumour into a failure.**

Step 1 — the run threshold.

```
f* = r = 0.0900

9.00% of deposits = 1,620,000
at the average balance, that is 216 of 2,400 members.
```

Two hundred and sixteen members out of two thousand four hundred can exhaust the entire liquid position. Not nine per cent of a majority — nine per cent of the whole. **That is W3, and it was fixed before anyone believed anything.**

Step 2 — what a run costs even though the institution is solvent.

```
demand at various withdrawal fractions f:

  f = 0.05   demand =    900,000   covered by cash? YES   shortfall =          0
  f = 0.09   demand =  1,620,000   covered by cash? YES   shortfall =          0
  f = 0.10   demand =  1,800,000   covered by cash? NO    shortfall =    180,000
  f = 0.15   demand =  2,700,000   covered by cash? NO    shortfall =  1,080,000
  f = 0.25   demand =  4,500,000   covered by cash? NO    shortfall =  2,880,000
  f = 0.40   demand =  7,200,000   covered by cash? NO    shortfall =  5,580,000
  f = 1.00   demand = 18,000,000   covered by cash? NO    shortfall = 16,380,000
```

Beyond $f = 0.09$ loans must be sold, and loans sold under duress do not fetch par:

```
liquidation value = cash + h * loan book, where h is the recovery rate

  h = 1.00 (no haircut)   19,020,000   105.67 cents on the dollar   SOLVENT
  h = 0.90                17,280,000    96.00 cents                 INSOLVENT
  h = 0.80                15,540,000    86.33 cents                 INSOLVENT
  h = 0.70                13,800,000    76.67 cents                 INSOLVENT
  h = 0.65                12,930,000    71.83 cents                 INSOLVENT
  h = 0.60                12,060,000    67.00 cents                 INSOLVENT

CRITICAL RECOVERY RATE   h* = (D - cash) / loans = 0.941379
```

A forced-sale recovery below 94.14 cents on the dollar makes a solvent institution insolvent. A ten per cent haircut — unremarkable for a distressed sale of a small loan book — puts members at 96.00 cents. **The belief that the institution cannot pay becomes true at a haircut of about six per cent, and haircuts of six per cent are ordinary. That is the loop closing, and it closes in the asset market, not in the members' heads.**

Step 3 — whether the seed propagates. Model member thresholds as Normal(μ, σ) with $\sigma = 0.25$:

```
critical mean tolerance   mu* = 0.408522

  mu = 0.407522   ->   x* = 0.990106     99.01% of deposits withdrawn -- total failure
  mu = 0.409522   ->   x* = 0.145087     14.51% withdrawn -- survivable with support
```

The institution's fate turns on two-thousandths of a unit in an unmeasured population parameter. No amount of care about the rumour changes this. **The correct objects of management are μ and σ, and both are properties of the membership, cultivated over years.**

Step 4 — the three interventions, each priced against a requirement.

```
INTERVENTION A -- attack W1. Communications.
  Board letter, member meeting, public accounts, personal calls.
  TARGET: mu, upward.
  COST: low in money, high in time, and it must be paid again next time.
  RELIABILITY: unknown. The outcome depends on mu to the third decimal place,
  and no instrument reads mu to the third decimal place.
  VERDICT: necessary, insufficient, and never the whole plan.

INTERVENTION B -- attack W2. A withdrawal gate.
  Notice period of 30 days on balances above 25,000.
  TARGET: K, the belief-to-action channel.
  COST: a real reduction in the product's value to members. Some will leave
  permanently for that reason alone.
  EFFECT: computed. If 40% of deposits sit in balances above 25,000, and the
  gate defers those, the maximum same-day demand falls from 100% to 60% of
  deposits, and the effective liquid coverage rises from 0.0900/1.00 to
  0.0900/0.60 = 0.1500. The run threshold on ungated deposits rises from
  9.00% to 15.00%.
  VERDICT: real, cheap in cash, expensive in trust, and it announces the
  problem it is solving.

INTERVENTION C -- attack W3. A committed liquidity line.
  A pre-agreed facility of 9,000,000 against the loan book at 95 cents,
  drawable on demand.
```

```
TARGET: R, the system's responsiveness to withdrawal.
EFFECT: computed. Liquid capacity rises from 1,620,000 to 10,620,000.
New run threshold f* = 10,620,000 / 18,000,000 = 0.5900 -- 59.00% of
deposits, or 1,416 of 2,400 members at the average balance, versus 216.
That is a 6.56-fold increase in the withdrawal fraction the institution can
absorb without selling an asset. And because the facility lends at 95 cents
rather than a distressed 80, the critical recovery rate never comes into
play for the first 9,000,000 of demand.
COST: a commitment fee, paid annually, whether or not it is drawn.
VERDICT: the intervention that actually removes the engine.
```

Step 5 — the comparison, which is the point.

	target	run threshold f*	members needed to break it
as-is	–	0.0900	216
+ communications	W1	0.0900	216
+ withdrawal gate	W2	0.1500	360
+ liquidity line	W3	0.5900	1,416
+ line and gate	W2,W3	0.9833	2,360

Communications does not move the threshold at all. It moves the probability that the threshold is approached, by an amount nobody can measure. **The liquidity line moves the threshold by a factor of 6.56, and the movement is arithmetic.**

Step 6 — the loop gain, before and after. Before the line, a withdrawal forces asset sales, the sales are visible, the visible sales confirm the rumour, and confirmation drives further withdrawal: **the fire sale is itself the evidence.** After the line, a withdrawal is met from the facility, no asset is sold, and nothing visible confirms anything. **The liquidity line does not reassure anyone. It removes the observable the loop was feeding on. That is what an intervention on W3 looks like from the inside: it does not argue, it disconnects. The pass gains before and after are not estimated here** — they would require the reflexive slope, which is not identified from observation [Manski, 1993] — and the qualitative statement is offered as a qualitative statement.

Step 7 — what the case does not show. It does not show that belief was unimportant: without the rumour there is no seed. **W1 is necessary. It is simply not sufficient, and it is the requirement over which management has the least control and the least measurement.** And it does not show that the rumour was false in any useful sense: **by the time a run is over, the rumour is true.** The belief was false about the balance sheet and true about the contract, **and the contract is what the depositor actually holds.**

35.13 Appendix B · Two beliefs, run through all four classes

All parameters below are assumed for the arithmetic and are labelled as assumed. They are not measured values and no claim is made that they are. What the case demonstrates is the decomposition procedure, not the numbers.

Belief 1: "I will get this job."

```
ASSUMED BASELINE
  P(offer | application submitted, no additional effort) = 0.0400
```

Channel A1: belief → application actually submitted. The most overlooked channel in the literature and the largest.

```
P(submit) = 0.35   ->   P(offer) = 0.014000
```

```
P(submit) = 0.90   ->   P(offer) = 0.036000

lift from A1 alone = 0.022000 absolute; a factor of 2.5714
```

Class: none — this channel works. It is behavioural, physical, its gain is large, and it is entirely uncontroversial. **It is also the channel most of the popular literature obscures by attributing the result to the belief rather than to the submission.**

Channel A2: belief → preparation hours → interview performance.

```
assumed: preparation rises from 3 h to 8 h
assumed: odds ratio 1.10 per hour of preparation

P(offer) rises 0.040000 -> 0.062885     absolute lift 0.022885
```

Class B if the claim were large: real channel, moderate gain. The gain is bounded by the odds ratio, **which is an empirical quantity someone would have to measure rather than assume.**

Channel A3: belief → demeanour → interviewer's rating.

```
assumed: odds ratio 1.25, applied once
P(offer) rises 0.062885 -> 0.077389     absolute lift 0.014505
```

Class B. Real, observable and modest — and this is the channel the expectancy literature measures, running in the other direction.

Channel A4: belief → the employer's decision, with no contact of any kind.

```
gain = 0, by construction. There is no transducer.
```

Class A. Not small. Absent.

The total.

```
believing, applying, preparing, and presenting well:
    P(offer) = 0.90 * 0.077389 = 0.069650

not believing, often not applying, not preparing:
    P(offer) = 0.35 * 0.040000 = 0.014000

ratio = 4.9750x        absolute difference = 0.055650

share of the total lift attributable to A1 -- simply applying = 39.53%
```

A five-fold change in outcome probability, and nearly two-fifths of it comes from the belief causing the person to send the application. That is a real, large and completely mundane effect, and it is exactly what the postulate predicts: **belief created reality, through channels, at the gains those channels have.**

And the honest reading of the other side: the belief did not raise the probability to 1. It raised it from 0.014000 to 0.069650. Ninety-three per cent of the time, the outcome is still no offer. A believer who reports the successful case and not the ninety-three is generating exactly the corpus Section 5.11 shows to be inert.

Belief 2: "I believe this traffic signal will change now."

```
W1  belief held                              yes
W2  channel to action                        yes: the pedestrian can press,
                                             wave, shout, stare
```

```
  W3  system responsive along that channel          NO. A fixed-timer signal reads
                                                    a clock. It does not read the
                                                    button, the wave, the shout,
                                                    or the belief.

  E = B x K x R = 1.00 x 1.00 x 0.00 = 0.0000
```

Class D, and the diagnostic took one sentence: what does this system take as input?

Note what happens next, because it is the general pathology in miniature. The pedestrian believes; the signal changes on its schedule; the coincidence is noticed. Repeated over a year, the believer accumulates a corpus of cases where the belief preceded the change, because the signal changes regularly and beliefs are frequent. **The corpus is real, the cases are honestly reported, and the likelihood ratio is 1.000000.** It is Section 5.11 running at pedestrian scale, and it is the same machine that produces the shelf of books.

	belief 1	belief 2
channel named	four, three real	one, unresponsive
gain	0.0140 -> 0.0697	0.0000
class of the failing channel	A (A4 only)	D
what the postulate predicts	an effect, bounded	no effect
what happens	an effect, bounded	no effect

The postulate got both right, and it got them right in advance, from the channel — which is the only way it is ever entitled to get anything right. Two cases is not a test, and Section 10 records the point.

35.14 Appendix C · The worksheet, and one belief taken end to end

35.14.1 C.1 The printable worksheet

One page. The date at the top and the signature at the bottom are the mechanism; everything else is scaffolding for them.

```
+-------------------------------------------------------------------------+
|  GROUNDING WORKSHEET                                          CCP-035   |
|                                                                         |
|  DATE STARTED  _______________     ANALYST  ____________________________|
|  THIS SHEET MUST BE COMPLETE THROUGH STEP 5 BEFORE ANY MEASUREMENT BEGINS.|
+-------------------------------------------------------------------------+
|  1  THE BELIEF, PRECISELY ENOUGH TO BE WRONG                            |
|     BELIEF   ___________________________________________________________ |
|     NEGATION ___________________________________________________________ |
|     [ ] the negation is also a clear, checkable sentence                |
|     [ ] a stranger with the instrument could reach the opposite verdict |
+-------------------------------------------------------------------------+
|  2  THE CHANNEL -- every link named. "Somehow" is not a link.           |
|     SOURCE      _____________________________  size / N: _____________  |
|     TRANSDUCER  ___________________________________________________________ |
|     MEDIUM      ___________________________________________________________ |
|     RECEIVER    ___________________________________________________________ |
|     EFFECTOR    _____________________________  units: _____________________  |
|     FAMILY  [ ] physiological  [ ] behavioural  [ ] social  [ ] informational|
```

```
|            [ ] other -- then name its transducer explicitly above           |
+-----------------------------------------------------------------------------+
|  3  PHYSICAL POSSIBILITY            pass / fail / not applicable            |
|     P1 conservation ____  P2 causality ____  P3 named transducer ____       |
|     P4 signal-to-noise ____  P5 scale law ____  P6 who learns it ____       |
|     ANY FAIL AT P3 OR P6 -> STOP. Record CLASS A and file the sheet.        |
+-----------------------------------------------------------------------------+
|  4  GAIN         type: [ ] dimensional [ ] log/elasticity [ ] standardised  |
|     link 1 _____________________ g = ______  source of estimate ___________ |
|     link 2 _____________________ g = ______  source of estimate ___________ |
|     link 3 _____________________ g = ______  source of estimate ___________ |
|     link 4 _____________________ g = ______  source of estimate ___________ |
|     link 5 _____________________ g = ______  source of estimate ___________ |
|     g_total = product = __________   uncertainty factor = __________       |
|     [ ] no gain above was estimated from the case under test                |
+-----------------------------------------------------------------------------+
|  5  PREDICTED MAGNITUDE -- NOT A DIRECTION.  *** MOST CLAIMS DIE HERE ***   |
|     predicted effect = __________ (units __________ )                       |
|     tolerance = +/- ______ %   -> interval [ __________ , __________ ]      |
|     [ ] the tolerance was set from the step-4 uncertainty, not from taste   |
|     standardised: reference SD = ______   predicted d = ______              |
|     required n per arm at 80% power = 15.6978 / d^2 = ______                |
|     obtainable? [ ] yes  [ ] no -> REDESIGN NOW, at step 5, not later       |
+-----------------------------------------------------------------------------+
|  6  MEASUREMENT PLAN -- VALIDATION GATES        pass / demote / redesign    |
|     FVP-1 dimensions ____   FVP-2 scale type ____   FVP-3 limits ____       |
|     FVP-4 fixed points ____ FVP-5 identifiability ____                      |
|     FVP-6 falsification ____ FVP-7 reproducibility ____                     |
|     instrument ____________________________________________________________ |
|     preregistered where / when ____________________________________________ |
|     THE RESULT THAT KILLS THIS CLAIM: _____________________________________ |
+-----------------------------------------------------------------------------+
|  7  COMPARISON -- completed AFTER measurement                               |
|     observed effect = __________       within tolerance? [ ] yes [ ] no     |
|     baseline (i)   persistence         beaten? [ ] yes [ ] no               |
|     baseline (ii)  direct-input        beaten? [ ] yes [ ] no               |
|     baseline (iii) matched AR          beaten? [ ] yes [ ] no               |
|     baseline (iv)  SHUFFLED INPUT      beaten? [ ] yes [ ] no               |
|     [ ] a randomisation or blocking defence against common causes was used  |
|     VERDICT  [ ] GROUNDED   [ ] GAIN OVERSTATED   [ ] GAIN UNDERSTATED /    |
|              SECOND CHANNEL   [ ] UNGROUNDED (upper bound on gain = ______ ) |
+-----------------------------------------------------------------------------+
|  STEPS 1-5 COMPLETED AND FROZEN ON  __________   SIGNED ____________________ |
+-----------------------------------------------------------------------------+
```

The constant in step 5 is computed: $2(1.959964 + 0.841621)^2 = 15.697759$, so n per arm $= 15.6978/d^2$ for a two-group comparison at $\alpha = 0.05$ two-sided and 80 per cent power. **It is on the sheet so the arithmetic can be done at the table.**

35.14.2 C.2 One belief, end to end

The figures below are constructed for the arithmetic and treated as a team's own measured history; they are labelled as constructed. A team of eleven believes that committing publicly to a ship date makes the date more likely to hold.

Step 1 — precisely enough to be wrong. After the four passes of Section 6.5:

Where at least 8 of 11 members rate confidence in the committed date at 4 or 5 on the team's existing 5-point kickoff form, mean scope added after kickoff is 1.00 points lower (tolerance ± 30 per cent, so 0.70 to 1.30 points) than where fewer do.

Negation: *...is not lower by that amount.* Both are clear. **Passes.**

Step 2 — the channel.

```
SOURCE      the eleven-member team. N = 11, so the 1/N problem is mild:
            one member is 9.09% of the group and the 25% tipping threshold
            is 3 members.
TRANSDUCER  speech. The commitment is stated aloud and recorded.
MEDIUM      the team's shared backlog and its meeting record.
RECEIVER    the team members, and the stakeholders who request scope.
EFFECTOR    scope points added after kickoff. Units: story points.
FAMILY      SOCIAL + INFORMATIONAL.
```

Step 3 — physical possibility. P1 not applicable. P2 passes. P3 passes: speech is a named transducer. P4 passes: the signal is a sentence in a meeting, far above any noise floor. P5: the effect does not attenuate with distance in any relevant sense. P6 passes: the stakeholders learn of the commitment because it is said to them. **All six clear.**

Step 4 — the gain.

```
link 1  belief -> public commitment actually stated        g = 0.70
link 2  commitment -> refusal at the scope boundary        g = 0.55
chain gain = 0.385000

points available to refuse per sprint, from history        = 2.60
PREDICTED reduction in scope added = 2.60 * 0.385000       = 1.0010 points
```

Step 4a — the link the team already measured. Fourteen releases:

```
mean slip           11.4286 days   (SD 7.1005)
mean scope added    10.3571 points (SD 5.0323)

OLS, slip ~ scope:
  slope      1.196355 days per point
  intercept -0.962248 days
  r          0.847885        r^2 = 0.718908
  SE(slope)  0.215952        t = 5.5399 on 12 df
  95% CI on the slope: [0.725836, 1.666874] days per point
```

This is the model link, and it is worth pausing on: it was estimated from data collected before the belief claim was made, on a different question. That is what REQ-6.5 requires, and it is why this gain may be used. Carrying the prediction to the endpoint:

```
predicted slip reduction = 1.196355 * 1.0010 = 1.1975 days
and the CI on the slope carries: [0.7266, 1.6685] days
```

Step 5 — magnitude, and the three redesigns. Sprint-level scope-added has a standard deviation of 3.10 points, so the predicted reduction of 1.0010 points is $d = 0.3229$:

```
DESIGN 1 -- parallel two-arm on sprints
  n per arm = 151      total = 302 sprints
  one team at 26 sprints/year = 11.62 YEARS.      INFEASIBLE.
```

```
DESIGN 2 -- paired crossover, lag-1 autocorrelation rho = 0.45
  SD of the paired difference = 3.10 * sqrt(2*(1-0.45)) = 3.2513
  paired d = 0.3079        n pairs = 83     total = 166 sprints
  one team = 6.38 YEARS.                          STILL INFEASIBLE.

DESIGN 3 -- twelve teams, cluster-randomised, ICC = 0.10, cluster size 26
  design effect = 1 + (26-1)*0.10 = 3.50
  required sprint-units = 302 * 3.50 = 1,057
  at 12 teams x 26 sprints = 312 per year -> 3.39 YEARS.
                                                   BARELY FEASIBLE.
```

Step 5 has just told the team that the study they wanted costs eleven and a half years. It told them before a single data point was collected, for the price of two divisions. That is the step doing its job.

The redesign that fits: instrument the high-rate link.

```
outcome: per-request refusal at the scope boundary
  baseline refusal rate     p1 = 0.22
  predicted under commitment: apply g2 = 0.55 to the currently-accepted
    margin of 0.235 -> 0.22 + 0.55*0.235 = 0.3493, stated as p2 = 0.35

n per arm = 189 requests        total = 378 requests

scope requests arrive at about 7.0 per sprint per team.
  one team:      378 / 7.0        = 54.00 sprints = 2.08 years
  twelve teams:  378 / (7.0 * 12) =  4.50 sprints = 0.17 years
```

From 11.62 years to 0.17 years — a factor of 68 — by changing which link is instrumented, with no loss of logical force. The endpoint inference is preserved: the measured refusal effect multiplies through the historical slope of 1.196355 days per point, whose confidence interval is already known, to give the slip consequence with honest uncertainty.

> **THE RULE THIS PRODUCES: the affordable study is set by the EVENT RATE of the link instrumented, not by the importance of the outcome cared about.**

Step 6 — the validation gates.

```
FVP-1  DIMENSIONS        gain in (points refused)/(commitment stated);
                         outcome in points; slope in days/point. PASS.
FVP-2  SCALE TYPE        the 5-point confidence form is ORDINAL. It is used
                         only to define a GROUP (>= 8 of 11 at 4 or 5), never
                         multiplied. Refusal is binary; scope points are
                         ratio-scaled with a true zero. PASS BY CONSTRUCTION,
                         and the construction was chosen for this reason.
FVP-3  LIMITING CASES    zero commitments -> predicted reduction 0.0000:
                         correct. All requests already refused -> saturation
                         at p = 1: the model must not predict beyond it.
                         PASS, with the ceiling stated.
FVP-4  FIXED POINTS      is the loop reflexive? If refusals raise confidence,
                         s > 0 and Section 4.10 applies. Estimated
                         g = eta*s/delta well below 1 here; STATED AS A
                         MODELLING POSITION, NOT AN IDENTIFIED PARAMETER.
FVP-5  IDENTIFIABILITY   g1 and g2 are not separately identified from the
                         refusal outcome alone -- only their product is.
                         DEMOTED: the sheet records g_total = 0.385000 and
```

```
                        states that the split is not recoverable from this
                        design. THIS IS THE GATE WORKING.
FVP-6  FALSIFICATION    preregistered; four baselines; killed if the
                        shuffled-input control matches, or if the observed
                        refusal difference is below 0.70 * 0.13 = 0.0910.
                        The suite's own false-rejection rate is 0.108
                        against a nominal 0.05 and is reported alongside.
FVP-7  REPRODUCIBILITY  twelve teams, parameters recovered per team and
                        compared. NOT YET RUN.
```

FVP-5 returned a demotion and the sheet records it. The design cannot separate "the team did not commit" from "the team committed and did not hold the line". **That is a real limitation, discovered before the study rather than after, and the correct response is to write it on the sheet rather than to pretend the product is a pair.**

Step 7 — the four verdicts, pre-committed.

```
observed refusal difference within [0.0910, 0.1690] and all four baselines
  beaten                                          -> GROUNDED. Report
                                                     g_total = 0.385000
                                                     with its interval.

observed positive but below 0.0910                -> GAIN OVERSTATED.
                                                     Channel real, g2 smaller.
                                                     Revise step 4, re-run.

observed above 0.1690                             -> GAIN UNDERSTATED or a
                                                     SECOND CHANNEL. Return to
                                                     step 2 and find it. NOT
                                                     a success.

observed near zero, or the shuffle predicts as
  well                                            -> UNGROUNDED. Report the
                                                     upper bound: at 189 per
                                                     arm the minimum detectable
                                                     standardised effect is
                                                     about 0.2882, and the
                                                     claim is bounded, not
                                                     disproved.
```

Step 7a — the common-cause defence, chosen at step 5. Because Section 6.13's Case 4 showed the shuffled control passing a pure common-cause configuration **98.35 per cent of the time**, this design uses **CC-1: the commitment ritual is randomised across sprints** rather than adopted whenever the team feels confident. **That single design decision is worth more than every statistical control in the suite — it is the only thing that makes the endogenous effect identified [Manski, 1993] — and it was available only because step 5 forced the design conversation before data collection.**

What the case demonstrates. The protocol did not confirm the belief; **it has not been run.** What it did was convert a sentiment — *committing publicly helps* — into a claim with a named channel, an independently estimated gain of 0.385000, a predicted magnitude of 1.0010 points with a ± 30 per cent tolerance, a stated falsifier at 0.0910, a feasible measurement at 378 requests, a randomisation defence, and a pre-committed verdict table.

That is what "grounded in truth" means operationally, and it is the only meaning the postulate is entitled to. The belief may still turn out to be wrong. A belief that can turn out to be wrong is the only kind this framework has any business calling grounded.

Chapter 36

Narrative as Regulatory Input: A Fully Disclosed Sonification Chain, a Zero-Resultant Corpus, and a Preregistration That Prices Its Own Nulls

Abstract. We specify a lossless, fully disclosed chain from text to audible tone — bytes to bits to nucleotides to frequencies — under an absolute no-subliminal constraint, and prove the structural property that prevents the protocol from overriding its subject. Modelling directives as unit vectors, a corpus closed under negation has resultant exactly zero, verified numerically at N = 4, 8, 12 and 24, while a mutually orthogonal corpus has a mean direction shrinking as 1/sqrt(N). A maximally contradictory corpus is therefore a maximum-surplus corpus: at N = 12 it carries ln 12 = 2.484907 nats, exactly the information required to select one of its members, so selection is left to the listener by construction. The proposed physical route is then closed by computation rather than defended: producing 1 per cent bulk strain in aqueous medium requires 240.83 dB SPL against an atmospheric ceiling of 194.09 dB. The worked pipeline renders 192 characters as 1,776 symbols at 20 ms in 35.520 s, with a Gabor margin of 248.83x against a required 100x and a dose of 0.123 per cent of the eight-hour occupational allowance per rendering. A receiver slower than the transmission recovers at most 29.803 of 1,776 bits, 1.68 per cent. The accompanying preregistration measures its own conjunction gate at a 9.5 per cent false-positive rate, not the nominal 5.

Keywords: narrative encoding; disclosure; antipodal corpus; sonification; preregistration; false-positive calibration; chain law

Series: B. Engineering the Surplus **Paper:** CCP-036

36.1 1 Introduction

A framework that proposes to structure an input and deliver it to a person owes two things before it delivers anything: a statement of what the input contains, complete enough that the recipient could rebuild it, and a structural argument that the input cannot override the recip-

ient's own judgement. **This paper supplies both, and then closes its own proposed physical mechanism by computation rather than defending it.**

The specification is a lossless chain from text to audible tone — bytes to bits to nucleotides to frequencies — governed throughout by an absolute constraint: **no subliminal content of any kind, nothing masked, nothing embedded below the threshold of hearing, nothing concealed from the listener. The information is the tone.** That constraint is not a stylistic preference reached at the end of a design; it is a boundary condition on whether the design may be built at all, and the argument for it depends on no empirical claim. **A message hidden from its recipient cannot be consented to**, and in an architecture whose entire mechanism routes through the recipient's own selection, concealment does not strengthen the protocol — it removes a component of it. **This is the same principle Kerckhoffs stated for cryptographic systems in 1883: a design whose security depends on the design being secret is a design that has not been evaluated** [Kerckhoffs, 1883, Brandeis, 1914].

The second guarantee is structural and can be stated exactly. **Modelling directives as unit vectors, a corpus closed under negation has resultant exactly zero — verified numerically at $N = 4$, 8, 12 and 24, where $\|\sum v_i\| = 0.000000000000$ in every case — while a mutually orthogonal corpus has a mean direction shrinking as $1/\sqrt{N}$.** A field of mutually exclusive directions cannot issue a command; it can only be resolved by selection, and the selection belongs to the receiver. **The quantitative form of that statement is the paper's second result: a maximally contradictory corpus is a maximum-surplus corpus, and at $N = 12$ it carries $F = \ln 12 = 2.484907$ nats $= 3.584963$ bits — exactly the information required to select one of its members.** The corpus supplies as much surplus as the selection consumes and supplies zero net direction.

The paper then closes its own proposed physical route. For a plane acoustic wave in an aqueous medium the longitudinal strain is $s = p/K$ with $K = 2.2 \times 10^9$ Pa, so producing the 10^{-2} strains associated with mechanosensitive gating requires 2.2×10^7 Pa — **240.83 dB SPL, against an atmospheric ceiling of 194.09 dB SPL set by the ambient pressure itself.** The bulk-compression route is closed by more than seven orders of magnitude and stays closed at the physical ceiling of sound in air. **That is a negative result about the framework's own proposal, obtained by arithmetic, and it is reported as a headline rather than as a caveat.**

The worked encoding renders 192 characters as 1,776 symbols at 20 ms in **35.520 s**, with a Gabor margin of **248.83 times** the mathematical floor against a self-imposed requirement of 100 times, and a dose of **0.123 per cent** of the eight-hour occupational allowance per rendering. **A receiver whose integration window exceeds the transmission recovers at most $\log_2 \binom{1779}{3} = 29.803$ bits of the 1,776 sent — 1.68 per cent, and none of it is sequence.** And the accompanying preregistration **measures its own decision gate rather than assuming it: the four-baseline conjunction's false-positive rate is 9.5 per cent at $T = 60$, not the nominal 5.**

Five contributions are made relative to the surrounding literature. **First**, a disclosure standard stated as a specification rather than as an ethical sentiment, with a reconstruction test — *can the recipient rebuild the signal from the packet?* — as its operative criterion. **Second**, a proof that a corpus can be built so that it cannot command, and the identity between the surplus it carries and the selection it requires. **Third**, a magnitude computation that closes one candidate mechanism decisively, which is progress precisely because a closed route narrows a search. **Fourth**, a preregistration whose test is calibrated by simulation before it is used, with the measured false-positive rate published in the registration. **Fifth**, a chain-law argument that a multiplicative composition containing one unmeasured link has an **undefined** product, applied to the framework's own architecture, which is stated in exactly those terms rather than as "promising".

No dataset was collected, no cell was exposed, no participant was recruited, and no result is reported. Everything below is either exact arithmetic, an evaluation of a stated model at stated parameters, a design simulation labelled as such, or a value quoted from published work with its source named. Section 2 states the hypotheses; Section 3 the method, including the identification problem and the analogy labels; Sections 4, 5 and 6 the chain, the protocol and the composition; Sections 7 to 11 report, discuss, verify, limit and conclude; Appendices A to C carry three worked specimens.

36.2 2 Hypotheses

Several results below are theorems and cannot be falsified by observation; they are labelled as derivations and recorded as `ARITHMETIC` or `DEFINITION` in Section 9's chain. What can be wrong are the claims about what the chain achieves, what the magnitude arithmetic forecloses, and whether the proposed experiment can see what it claims to be able to see.

H1 — A contradictory corpus cannot issue a directive. Predicted: a corpus closed under negation has resultant exactly zero at every even N, and a mutually orthogonal corpus's mean resultant falls as $1/\sqrt{N}$, tracking the analytic value in simulation. **Falsified by:** a corpus construction closed under negation with a non-zero resultant, or a simulated orthogonal ensemble whose mean resultant does not track $1/\sqrt{N}$. **Result: confirmed.** $\|\sum\| = 0.000000000000$ at N = 4, 8, 12, 24; simulated mean resultants 0.497007, 0.352275, 0.288022 and 0.203292 against analytic 0.500000, 0.353553, 0.288675 and 0.204124. **This is a result about a vector model of directives, not about texts**, and Section 4.4's clause H-B states exactly what the model does and does not capture.

H2 — The surplus a maximally contradictory corpus carries equals the information required to select one of its members. Predicted: at $u = 1$ and a uniform receiver prior, $F = \ln N$ exactly equals the selection cost $\ln N$; at $N = 12$ both are 2.484907 nats. **Falsified by:** any parameterisation within the stated conditions in which the two differ. **Result: confirmed in the extremal case only**, and Section 4.4's clause H-A states that the identity fails outside $u = 1$ and a uniform prior, with both sides falling together.

H3 — The bulk-compression route from an audible acoustic field to mechanosensitive gating is closed. Predicted: $s = p/K$ with $K = 2.2 \times 10^9$ Pa requires 240.83 dB SPL for a 10^{-2} strain, which exceeds the 194.09 dB atmospheric ceiling; at the ceiling the achievable strain is 4.6057×10^{-5}, short by a factor of 217.12. **Falsified by:** a bulk modulus, a gating strain, or an achievable pressure that closes the gap — or by a demonstration that gating occurs at bulk strains four orders of magnitude smaller. **Result: confirmed by arithmetic. The closure applies to the bulk route only**; non-linear and localised routes are enumerated in Section 5.2 and are **not** closed.

H4 — A receiver slower than the transmission cannot recover sequence. Predicted: an integrator whose window exceeds 35.520 s sees composition and not order, and composition carries $\log_2 \binom{1779}{3} = 29.803$ bits of the 1,776 transmitted. **Falsified by:** a receiver model that recovers ordering information from a composition-only observation, which would be a counterexample in combinatorics rather than in biology. **Result: confirmed by arithmetic**, at 1.68 per cent of the transmission.

H5 — The four-baseline conjunction gate does not achieve its nominal false-positive rate at practical series lengths. Predicted: simulated at zero effect with $T = 60$, the gate rejects at materially more than 0.05. **Falsified by:** a calibration run returning 0.05 or below at $T = 60$. **Result: confirmed, and the number is 0.095** — nearly twice nominal. **The gate is anti-conservative at the series length the design would otherwise have used**, and lengthening the series to $T = 120$ brings it to 0.040.

One further claim is deliberately not advanced as a hypothesis. No hypothesis is offered that a cell demodulates sequence-specific information from an acoustic signal. **That is the open problem the paper exists to state, and it is recorded as `NOT ESTABLISHED` throughout rather than as a hypothesis this paper tests**, because nothing in this paper tests it.

36.3 3 Method

36.3.1 3.1 What the paper attempts

The paper specifies an encoding, states a structural theorem about the material encoded, computes a magnitude that closes one candidate mechanism, and writes a preregistration whose test is calibrated before it is used. **It reports no measurement of any biological or psychological quantity.** Two of its three worked specimens are constructed and the third is a specimen preregistration for a study that has not been run.

The paper's most consequential outputs are negative: a closed route, an undefined product, and a gate whose measured false-positive rate is nearly twice its nominal one. **Those are reported at the front of Section 7 rather than in a discussion, because a specification that hides its own closures is a brochure.**

36.3.2 3.2 Analogy discipline, declared in advance

`RIGOR_STANDARD` §8 requires every borrowing from physics to be labelled, at first use, as a derivation, a shared mathematical form, or an analogy. The labels are fixed here.

Borrowing	Label	What licenses the label
$s = p/K$ for strain in an aqueous medium	**derivation**	elementary continuum mechanics; the medium genuinely satisfies the premises, and the result is used to close a route rather than to open one
the Gabor limit $\Delta t\,\Delta f \geq 1/(4\pi)$ applied to a rendered tone	**derivation**	the premises are square-integrability and finite second moments, satisfied by a finite-energy signal under a smooth envelope [Gabor, 1946, Curzi, 2026z]
the chain law $G = \prod_i g_i$ for a cascade	**shared mathematical form**	the same product governs an optical cascade, a conversion chain and a causal chain; **no physical result transfers**, and each g_i must be measured in its own units
directives as unit vectors in an inner-product space	**modelling choice**, not a borrowing	badged in Section 4.4, clause H-B; real directives are not unit vectors and "REST NOW" is not literally the additive inverse of "WORK NOW"
the packet/header/payload correspondence for an address prefix	**shared mathematical form**	routing reads a header and does not alter a payload in both cases; **it is not a claim that cells run a network stack**, and the correspondence earns nothing beyond letting a reader of one system read the other
the Casimir pressure at biological separations	**derivation** within physics, used as a **counter-example**	the formula is exact; the point is that it is negligible [Casimir, 1948, Lamoreaux, 1997]
"resonant coupling" as a reading of older correspondence doctrines	**analogy**, and marked **speculative** in the sentence that states it	motivation only, **never evidence**; it supplies no mechanism for sequence-selectivity

Two exclusions are stated here and enforced throughout. No quantum result is used to license any claim about mind, belief or biology [Curzi, 2026aa]. And **the elegance of an encoding is never offered as evidence that anything receives it** — Section 4.11 states that as a design claim of the same kind as "this file is sorted", and failure mode FM-4.9 exists because elegance reads as explanation.

36.3.3 3.3 Entropy: which entropy, in every instance

Every entropy and every logarithm of a count in this paper is Shannon information [Shannon, 1948, Cover and Thomas, 2006]. The surplus F and the ceiling $\ln N$ are reported in **nats**; selection costs and payload sizes are reported in **bits**; the conversion is $1 \text{ nat} = 1/\ln 2 = 1.442695$ bits, so $\ln 12 = 2.484907$ nats $= 3.584963$ bits. **No thermodynamic entropy appears anywhere in this paper**, no quantity below carries units of energy per temperature, and the one place

a physical entropy could legitimately enter — the thermodynamic cost of erasing a bit [Landauer, 1961] — is not used, because no claim here depends on it. **The two objects share a name and a functional form, are different quantities with different units, and are never conflated in this paper.**

The information-theoretic bound this paper leans on hardest is the variety bound: **a regulator can absorb only the variety it can represent** [Ashby, 1956, 1958], which is the same statement as $F \leq \ln N$ in the framework's notation [Curzi, 2026]. Selection cost as $\log_2 M$ is Hartley's, extended by Shannon [Hartley, 1928, Nyquist, 1928].

36.3.4 3.4 Effective counts: saying which one

The effective count used throughout Sections 4 and 6 is **the corpus's own affine construction** $N_{\text{eff}} = 1 + (N-1)u$, which at $u = 1$ equals N and at $u = 0$ equals 1 [Curzi, 2026]. **It is NOT a Hill number of any order.** It is a linear interpolation between 1 and N, no literature in ecology, economics or information theory contains it, and **no Hill-number citation is attached to it anywhere in this paper.** The Hill numbers of order one and two — $\exp(H)$ and $1/\text{HHI}$ — are well grounded [Hill, 1973, Jost, 2006, Rényi, 1961] and **are not used here**.

The distinction matters numerically at exactly the place it is used. At $N = 12$ and $u = 0.85$ the affine count is $1 + 11(0.85) = 10.3500$; **the Hill number of order one for a distribution with that much unevenness would in general be a different number**, and a reader who reaches for the Hill-number literature to support the affine count will find it does not apply. **Saying so is cheaper than being caught.**

36.3.5 3.5 Identification, and what the protocol can and cannot recover

Two distinct identification problems arise and they are answered differently.

The Layer 1 problem is one of attribution and it is answered by design. A difference between the encoded arm and its scramble is attributable to sequence *or to nothing*, because the two arms are identical in duration, spectrum, level and symbol rate. **That is randomisation plus a matched control, and it identifies the contrast by construction rather than by assumption.** The address-swap and payload-swap arms of Section 5.11 are what convert an attributable difference into an attributable *mechanism*, and Section 5's specification declares that a positive without them is reported as **unattributed**.

The Layer 2 problem is the reflection problem and it is not answered. The framework's reflexive material — a belief held about oneself, acting through channels, changing the state that the belief was about — is a reflexive specification whose endogenous effect **is not identified from observational data** [Manski, 1993], and influence and selection are generically confounded in observational data of this kind [Shalizi and Thomas, 2011]. **The Layer 2 design adopts one remedy and pays its price: the exposure schedule is randomised**, which makes the contrast between exposed and unexposed periods identified by design. **It does not identify the reflexive parameter itself**, and no statement in this paper about the loop's gain, slope or strength is an identified empirical result. The stipulated-reference-group route [Brock and Durlauf, 2001] is named as available and is not taken. **Every claim in this paper that touches the loop is a modelling position.**

The Lucas critique applies with unusual directness because the protocol publishes its own specification [Lucas, 1976]. A disclosed corpus that participants read changes what participants do, which is the entire point; **so parameters estimated before publication need not describe behaviour after it**, and a measure of protocol quality that becomes a target stops measuring [Goodhart, 1975, Campbell, 1979]. **No defence of parameter stability is offered.** The measured social-prediction literature is consistent with the modest reading [Hofman et al., 2017, Salganik et al., 2020, Watts, 2014, Lazer et al., 2014].

36.3.6 3.6 Scale type, and the conjoint test that was not run

RIGOR_STANDARD §5 requires every measured construct to declare its scale type [Stevens, 1946, Krantz et al., 1971].

- **Frequency, time, pressure, strain, sound level in pascals, bits, nats and counts of loci are ratio-scaled** with true zeros, and every operation performed on them is licensed. Sound pressure **level** in decibels is a **logarithmic ratio scale against a fixed reference** ($20\,\mu$Pa), so differences in dB are ratios of pressures and are licensed; **sums of dB values are not, and none is taken.**
- $\log_2$ **transcript abundance is an interval scale in the log domain**, which is why the primary endpoint is a *difference in differences* rather than a ratio of ratios: differences of logs are licensed on an interval scale and are exactly what a fold-change is.
- **The Layer 2 self-report endpoint is the exposed one.** A rating scale is **ordinal**, and the storage law's multiplicative structure is **not licensed on ordinal input** [Stevens, 1946]. Section 5.9 states that where the Layer 2 instrument is ordinal, the fitted relation is **demoted in print to a monotone ordering**, and the corpus's own central formula has already been demoted on exactly this gate [Curzi, 2026t].
- **The essence index n is a count and is dimensionless**; the reference constant k is discussed in Section 9.1.

Michell's critique is the strongest objection to the framework this paper belongs to [Michell, 1997]: psychology assumes quantitative structure without testing for it, and the formal test is additive conjoint measurement [Luce and Tukey, 1964]. **It has not been run on any psychological construct in this paper and this paper does not run it.** The quantities carrying the paper's numbers are physical or combinatorial and are not exposed; **the one exposed construct, the Layer 2 self-report, is demoted for that reason. That demotion is a concession, not a defence.**

36.3.7 3.7 Scaling claims, and the procedure not run

Three power-law statements appear and **all three are defined by the model rather than estimated from data**, which Clauset et al. [2009] requires to be said explicitly.

- The Casimir pressure falls as a^{-4}. **An exact property of $\pi^2\hbar c/(240a^4)$, not an estimate.**
- The address-space count 4^L grows exponentially in address length. **A count, not a fit.**
- The chain law produces geometric decay in the number of links. **The decay constant is the per-link gain, an assumed parameter in every table below.**

No heavy-tail, scale-free or power-law claim is made anywhere in this paper on empirical data, and the maximum-likelihood-plus-goodness-of-fit procedure of Clauset et al. [2009] was therefore not run. It is named so its absence is a decision rather than an omission [Stumpf and Porter, 2012, Broido and Clauset, 2019].

36.3.8 3.8 What is presupposed

The Gabor limit and its cost schedule are presupposed from Curzi [2026z] and Gabor [1946] and are not re-derived. The essence law $f(n) = kn^2$ with $k = (9/8)^2/\varphi^2 = 0.483425732988414$ and the odd-part pitch-class rule are presupposed as badged postulates [Curzi, 2026]: **k is a chosen reference constant, not a measured one, and the square law is a postulate about essence structure.** The two components of k are real and their quotient expressed as a frequency is not — $9/8$ is the Pythagorean whole tone [Barbour, 1951, von Helmholtz, 1885] and $\varphi^2 = \varphi + 1$ is Euclid VI.3 [Euclid, 1956, Vajda, 1989] — and Markowsky [1992] is cited here, as everywhere in this corpus that the golden section appears, because it is the standing correction to golden-ratio claims. Octave equivalence, on which folding rests, is formalised by Shepard [1982] and its perceptual limits are given by Burns [1999], Demany et al. [2013], Hoeschele et al. [2012];

that consonance depends on the spectrum rather than the ratio alone [Sethares, 2005] is the standing objection to any ratio-based claim in this area and is inherited rather than answered.

The surplus functional, its affine effective count and the ceiling $0 \leq F(u) \leq \ln N$ are presupposed from Curzi [2026]. The storage law $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$, its fixed point and its uniform bound are presupposed from Curzi [2026a]. The seven validation gates are specified in Curzi [2026t]; the four-baseline falsification suite, **including its own measured anti-conservatism — iid permutation rejecting at 0.108, circular shift at 0.120 and block permutation at 0.060 against a nominal 0.05 across 500 null trials** — is specified in Curzi [2026u] and is inherited with that measurement attached. The standing constraint that **no addressing operation, at any signal strength or modality, can cause a cell to express content its genome and accessible regulatory repertoire do not already contain** is presupposed from Curzi [2026y]. The conditional postulate and its channel-gain architecture are Curzi [2026ab]'s; the catalogue of where that postulate fails, and the requirement that grounding be established from the channel rather than from the outcome, are Curzi [2026ac]'s. The concordance discipline — that the evidence column and the precedent column carry different things — is Curzi [2026w]'s.

36.3.9 3.9 Standing honesty clauses

1. **No subliminal content of any kind.** Nothing masked, nothing below threshold, nothing concealed. **The information is the tone.**
2. **No clinical, medical or therapeutic claim is made anywhere in this paper, in either direction.** Nothing described here is a treatment, a therapy, an intervention or a modality.
3. **There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal** [Curzi, 2026y]. Stated verbatim wherever the biological chain is mentioned.
4. **Nothing is delivered, therefore nothing is edited.** The architecture produces sound and cannot do otherwise; no molecule is synthesised, no vector introduced, no sequence changed.
5. **No capability is conferred.** Nothing not already latent can be reached, and the ceiling is $\ln N$ [Curzi, 2026,y].
6. **Precedent is never presented as evidence** [Curzi, 2026w].
7. k **is chosen and the outputs are essence frequencies, not physical resonances** [Curzi, 2026].

36.3.10 3.10 Bibliographic disclosure

A substantial number of the primary sources this paper relies on carry no key in this corpus's bibliography, and no key was invented for any of them. Named in the text with journal, volume, year and pages, and cited without a key: Church, Gao and Kosuri (*Science*, 2012) for DNA data storage; Qi et al. (2013) and Gilbert et al. (2014) for the separability of targeting and editing; Zetsche et al. (*Cell*, 2015) and Fonfara et al. (2016) for native crRNA-array processing; the 2021 Nobel-recognised work on mechanically gated ion channels; Tajik et al. (*Nature Materials*, 2016) for force-to-transcription; Ibsen et al. (*Nature Communications*, 2015) and Duque et al. (2022) for sonogenetic channel activation; Takahashi and Yamanaka (2006, 2007) for regulatory reprogramming; Kaptchuk et al. (2010) for open-label placebo; Levine, Gordon and Fields (1978) for naloxone-reversible placebo analgesia; Sisk et al. (2018) for the growth-mindset meta-analyses; and O'Sullivan (1845) for the political slogan Section 6.1 distinguishes this work from. The NIOSH recommended exposure limit is quoted as a published occupational standard. **Every one is a standard, locatable source; the gap is in this corpus's bibliography, and it is recorded as Limitation L-2 rather than smoothed over.**

36.4 4 The disclosed chain: from text to audible tone

36.4.1 4.1 The no-subliminals constraint

NO SUBLIMINAL CONTENT OF ANY KIND. NOTHING MASKED. NOTHING EMBEDDED BELOW THE THRESHOLD OF HEARING. NOTHING CONCEALED FROM THE LISTENER. THE INFORMATION IS THE TONE.

It is an ethical constraint, not a stylistic one, and the two kinds fail differently. A stylistic constraint is a preference: violate it and the work is worse. An ethical constraint is a boundary condition on whether the work may be done at all: violate it and the result is not a poorer version of the same thing but a different thing, out of scope regardless of how well it performs.

The argument depends on no empirical claim. A message hidden from its recipient cannot be consented to. Consent requires knowing what one is consenting to. A signal engineered so the listener cannot perceive it has been engineered so the listener cannot evaluate it, cannot compare it against their own judgement, cannot refuse it, and cannot report on it afterwards. **Every one of those capacities is the capacity to decline. Removing them is not a side effect of the concealment architecture; it *is* that architecture.**

And consent is load-bearing here in a way it is not in most engineering. **The entire model is one in which the subject's own latent capability is the mechanism**: nothing not already latent can be reached [Curzi, 2026y], and belief acts only through a real channel [Curzi, 2026ab]. In both, the person is not the object the method acts on; **the person is the part of the circuit that does the acting. A protocol built on the subject's own capability, which then conceals itself from the subject, has removed the participant from the mechanism it claims to be using. It is incoherent before it is unethical, and it is unethical.**

The same principle is a century old in a different discipline. Kerckhoffs's second principle holds that a cryptographic system must remain secure when everything about it except the key is public [Kerckhoffs, 1883] — **a design whose effectiveness depends on the design being secret has not been evaluated**, and publication is the mechanism by which it becomes evaluable. The corpus applies the same argument to channel architecture, where a separation is **strengthened** by disclosure because its effectiveness is entirely independent of whether anyone knows it is in place, so that **non-disclosure is refused as a design premise rather than merely discouraged** [Curzi, 2026]. **This paper's constraint is that argument applied to content rather than to a channel or a cipher.** The general case for publicity as a discipline rather than a courtesy is Brandeis [1914]'s.

The specification form, binding on everything downstream:

1. **Every component of the rendered signal is audible to a person of ordinary hearing at the listening position.** Not "audible in principle" — audible as rendered, at the level rendered, in the room, over the noise floor present. **An inaudible component is raised or removed. There is no third option.**
2. **No carrier, no mask, no embedding.** No content under, inside or behind another signal; no amplitude-modulated speech; no time-compressed content; no spectral hiding beneath a louder band; no content below 20 Hz or above 20 kHz intended to act on a listener; **no psychoacoustic masking exploited as a delivery mechanism.**
3. **The full source text is disclosed in plain language before exposure.** Not a summary. The corpus itself, in the language the listener reads, with the encoding table, address list, tuning and rendering parameters — **enough to reproduce the signal.**
4. **The listener can stop it and can leave.** The off switch exists, its location is disclosed, and anyone exposed may use it without asking. **Reversibility is a property of the protocol, not a courtesy extended by the operator.**
5. **The disclosure precedes exposure and is not conditional on participation.** A person who reads the corpus and declines has received the full disclosure and given nothing in exchange.

Requirement 3 distinguishes this method from every technique it superficially resembles.

Backmasking, embedded content and every commercial "subliminal" product share one architecture: content present in the signal and absent from the listener's knowledge. **The architecture here inverts that exactly: the content is present in the signal and fully present in the listener's knowledge, and the listener could have written the signal themselves from what they were told.**

36.4.2 4.2 Disclosure costs nothing that matters — and would bind even if it did

The obvious objection to a no-concealment rule is that concealment is what makes such methods work. **The answer is a measurement, and it is the most useful single fact in this framework's ethical architecture.**

Evidence. Open-label placebo works. Kaptchuk et al. (2010; **named without a key**) randomised patients with irritable bowel syndrome to placebo explicitly described to them as inert, or to no treatment, and found benefit in the open-label arm; subsequent trials in other conditions have replicated the basic finding. **Effect sizes are modest, the literature is active, and the mechanism is not settled** [Curzi, 2026ac]. Levine, Gordon and Fields (1978; **named without a key**) showed placebo analgesia to be mediated by endogenous opioids and naloxone-reversible — **a complete physical channel from expectation to measured outcome.** Taken together: **for at least one fully characterised belief-mediated channel, deception is not a necessary input.**

Precedent. The traditions that used ritual most systematically frequently made it public and its text recited aloud: the liturgy is not hidden from the congregation, the mantra is taught to the practitioner, the memory palace is built by the person who will use it out of images they choose. **This is precedent, not evidence: it records that people arrived independently at a disclosed architecture and preserved it — an argument that the architecture is findable, not that it works** [Curzi, 2026w]. **The traditions also, often, concealed**, and Section 4.14 treats that rather than selecting the convenient half of the record.

Two limits, both required. *First*, open-label placebo is evidence about expectancy effects and not about the acoustic-to-cell channel, which remains open. It shows that **one** belief-mediated channel tolerates disclosure — not that every channel does, and certainly not that a channel exists where none has been demonstrated. *Second*, **even if disclosure did cost effect size, the constraint would stand.** Section 4.1 is a boundary condition, not an empirical bet that disclosure is free. **The measurement is welcome because it means the boundary is cheap to hold, not because holding it was contingent on that.**

36.4.3 4.3 The chain, as a specification

```
STAGE 0   CORPUS      a set of N written directives, deliberately contradictory
STAGE 1   DATA        each directive -> bytes -> bits           (lossless, reversible)
STAGE 2   SEQUENCE    bits -> nucleotide string                 (lossless, reversible)
STAGE 3   ADDRESS     20-nt routing prefix per directive        (a header, not a molecule)
STAGE 4   TONE        nucleotide string -> audible tone sequence (lossless, reversible)
```

Every stage is a bijection on the data it carries — a design requirement, not an accident. The chain is reversible end to end: given the rendered audio and the published tables, any listener can recover the corpus exactly. **That is what makes requirement 3 meaningful — the listener is not being told *about* the signal, they are being given what the signal is.**

Nothing physical is delivered at any stage. No molecule is synthesised, no vector introduced, no Cas protein present, no sequence in any cell changed. Stage 2 produces a *string over the alphabet* $\{A, C, G, T\}$, a data structure; Stage 3 produces a *prefix on that string*, a routing header. **The claim that nothing is edited is not a promise about restraint; it is a statement about what the architecture can do, and the architecture can produce sound and nothing else.**

36.4.4 4.4 Determination by choice — the structural theorem

The corpus at Stage 0 is deliberately contradictory. Directive 1 says rest; directive 2 says work; directive 5 says speak; directive 6 says be silent. **The contradiction is not a hedge and not an artistic gesture. It is the structural property that makes the protocol unable to override its subject, and it can be stated exactly.**

PROVED (elementary). Model each directive as a unit vector v_i in a real inner-product space, where the direction encodes the action urged and $-v$ encodes its negation.

(a) If the corpus is closed under negation, the resultant is exactly zero:

```
sum_{i=1}^{N} v_i = 0                        exactly, for any pairing
```

Verified numerically for N = 4, 8, 12, 24: $\|\sum\| = 0.000000000000$ in every case.

(b) If the corpus is mutually orthogonal, the resultant grows only as $\sqrt{N}$ while the corpus grows as N, so the mean direction shrinks as $1/\sqrt{N}$:

N	||sum||	||mean||	1/sqrt(N)	simulated (unit vectors in R^64)
4	2.000000	0.500000	0.500000	0.497007
8	2.828427	0.353553	0.353553	0.352275
12	3.464102	0.288675	0.288675	0.288022
24	4.898979	0.204124	0.204124	0.203292

The simulated column is 2,000 draws of N isotropic unit vectors in $\mathbb{R}^{64}$ per row; at N = 50 and 100 it gives 0.140454 and 0.099395, tracking $1/\sqrt{N}$ throughout.

A field of mutually exclusive directions cannot issue a command. In case (a) the resultant is identically zero; in case (b) it is a vanishing fraction of any single member and keeps vanishing as the corpus grows. **The corpus does not point. It spans.**

PROVED (extremality of the surplus functional). $F(u) = \ln N$ if and only if the ensemble is mutually orthogonal, and $F(u) = 0$ if and only if it is parallel [Curzi, 2026]. **A maximally contradictory corpus is therefore a maximum-surplus corpus:** for $N = 12$, $F = \ln 12 = 2.484907$ nats, which is the hard ceiling $\ln N$ and cannot be exceeded.

PROVED (elementary information theory). Selecting one of N equiprobable alternatives requires $\ln N$ nats — $\log_2 12 = 3.584963$ bits for a twelve-directive corpus [Hartley, 1928, Shannon, 1948].

These two quantities are the same number, and that identity is the content of this section.

```
surplus carried by a maximally contradictory corpus     F = ln N   nats
information required to select one member of it              ln N   nats
-----------------------------------------------------------------------
N = 12:   F = 2.484907 nats = 3.584963 bits    selection = 2.484907 nats
```

The corpus supplies exactly as much surplus as the selection consumes, and it supplies zero net direction. The direction is not in the signal. It is contributed by the receiver, and its magnitude is exactly the surplus the signal carries.

Two honesty clauses attach, both required.

H-A. The identity holds under two stated conditions: $u = 1$, and a uniform receiver prior. **Under a non-uniform prior the selection information is that prior's entropy, strictly less than $\ln N$; under $u < 1$ the corpus surplus is $\ln(1 + (N-1)u) < \ln N$. Both fall together, which preserves the structure of the argument but not the exact equality — and the exact equality is claimed only in the extremal case.**

H-B. The vector model of a directive **is a model**, badged as such. **Real directives are not unit vectors, and "REST NOW" is not literally the additive inverse of "WORK NOW."** What the model captures rigorously is one thing: *a set constructed so that its elements cancel or span has no resultant, and a receiver who acts must therefore have selected.* **That conclusion depends only on the corpus genuinely containing mutually exclusive directions, which is a property of the written text and can be checked by reading it.**

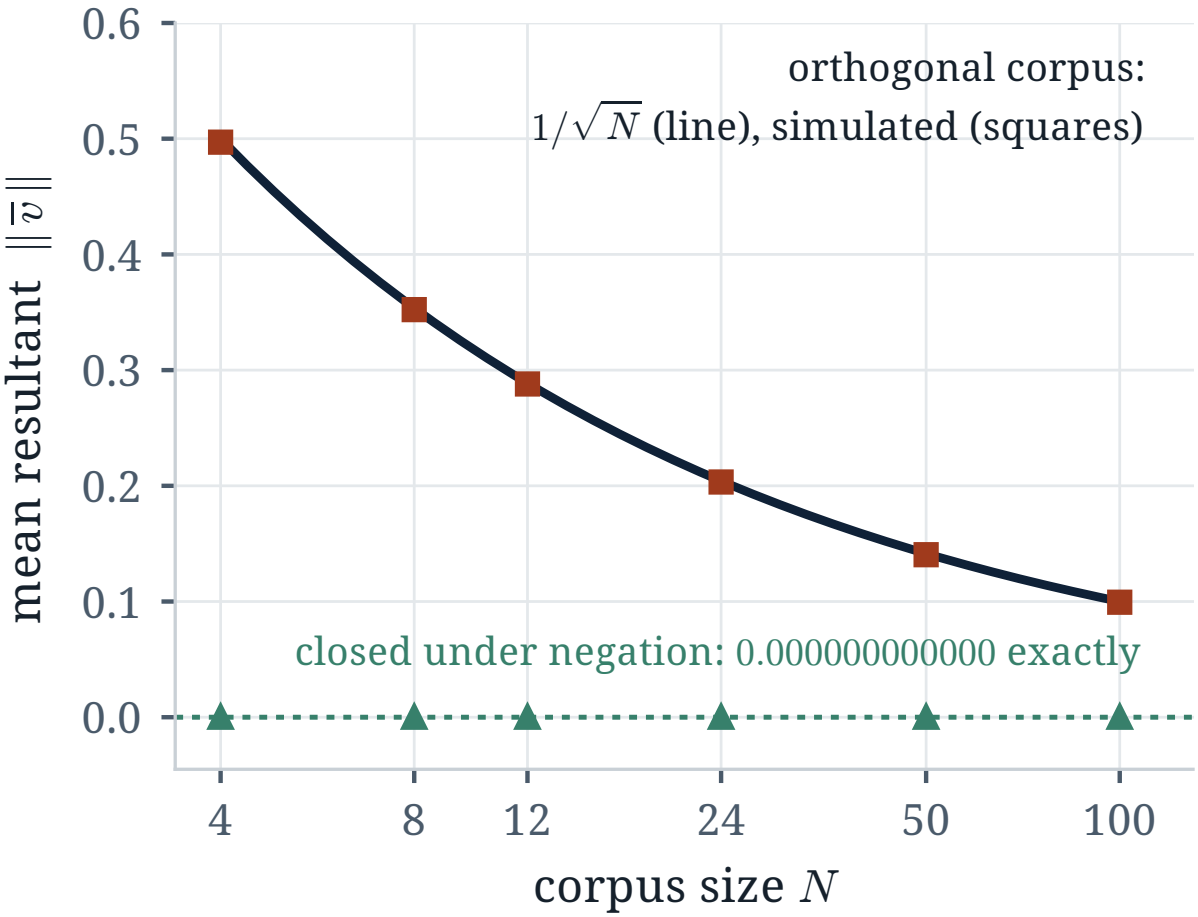


Figure 36.1: **A corpus closed under negation has resultant exactly zero, and an orthogonal corpus's mean direction vanishes as $1/\sqrt{N}$:** the analytic mean resultant $1/\sqrt{N}$ against corpus size, with the simulated values from 2,000 draws of N isotropic unit vectors in $\mathbb{R}^{64}$ per point, and the antipodal case pinned at 0.000000000000 at N = 4, 8, 12 and 24. **The simulation is a check on the arithmetic, not a measurement of any text**; the vector model of a directive is a modelling choice badged at clause H-B. Generated by `figs/src/CCP-036_fig01.py`.

36.4.5 4.5 Warrant

Claim	Status	Basis
No subliminal content; full disclosure	**CONSTRAINT** — ethical, prior, non-negotiable	Section 4.1
Disclosure need not abolish a belief-mediated effect	**CITED (measured)**	Kaptchuk et al. 2010; Levine et al. 1978, both named without a key
Data can be encoded in and recovered from nucleotide sequence	**CITED (measured)**	Church, Gao & Kosuri 2012, named without a key
A 20-nt spacer is a near-unique address in a 3.1 Gbp genome	**ARITHMETIC**	Section 4.8
Targeting and editing are separable functions	**CITED (measured)**	Qi et al. 2013; Gilbert et al. 2014, named without keys
$\Delta t\,\Delta f \geq 1/(4\pi)$ constrains symbol rate	**DERIVATION (classical Fourier)**	[Gabor, 1946, Curzi, 2026z]
$f(n) = kn^2$, $k = 0.483425732988414$	**POSTULATE** — k is chosen	[Curzi, 2026, Markowsky, 1992]
Two indices share a pitch class iff same odd part	**ARITHMETIC**	[Curzi, 2026]
A contradictory corpus has zero or vanishing resultant	**ARITHMETIC**	Section 4.4
Corpus surplus $\ln N$ equals selection information $\ln N$	**ARITHMETIC, extremal case only**	Section 4.4, with H-A
Narrative structure is carried in the harmonic	**DESIGN CLAIM** — about encoding, not effect	Section 4.11
A cell demodulates sequence-specific content from this signal	**NOT CLAIMED. OPEN PROBLEM.**	[Curzi, 2026y]
Any physiological, clinical or therapeutic effect	**NOT CLAIMED**	Section 3.9

36.4.6 4.6 Stage 1 — narrative to data

The corpus is written as plain text. The encoding to bits is ordinary and **must be published**: character set, byte width, byte order, whether a terminator is used. The reference scheme is 8-bit bytes, one byte per character, no terminator, directives kept separate rather than concatenated. For the worked corpus of Appendix A: 192 characters and therefore 1,536 bits.

Nothing interesting happens at this stage, and that is the point. Stage 1 is deliberately the most boring transformation in the chain, because any cleverness here would be a place to hide something. A published, standard, character-for-character mapping is auditable by anyone with a text editor.

36.4.7 4.7 Stage 2 — data to sequence

Evidence. Church, Gao and Kosuri (*Science*, 2012; **named without a key**) encoded a book — text and images — into synthesised DNA and read it back. **The encoding they used was deliberately *not* the information-theoretic maximum.** A four-symbol alphabet carries $\log_2 4 = 2$ bits per base at the ceiling; they used **one bit per base**, with a degenerate mapping in which A or C encodes 0 and G or T encodes 1.

```
2 bits/base   maximal density, zero freedom at each position
1 bit/base    half density, ONE FREE CHOICE at every position
```

The discarded bit buys constraint satisfaction. With two acceptable bases per bit, a synthesiser can avoid homopolymer runs, balance GC content, and dodge secondary structure without changing a single bit of payload. For a 1,536-bit payload there are 2^{1536} distinct nucleotide strings encoding exactly the same message — a 463-digit count ($1536 \log_{10} 2 = 462.382$).

The reference scheme adopts one bit per base for a second reason specific to this paper: the mapping from message to sequence is then one-to-many, so the published disclosure must include the actual sequence used, not merely the rule. A scheme where the rule determines the sequence can be disclosed by publishing the rule; a scheme where it does not must publish the string. **The weaker encoding forces the stronger disclosure, which is the right direction for the trade to run.**

corpus	chars	bits	bases @1 bpb	bases @2 bpb
the worked corpus	192	1,536	1,536	768
a 280-character message	280	2,240	2,240	1,120
a 4,000-character essay	4,000	32,000	32,000	16,000

36.4.8 4.8 Stage 3 — the address prefix

The prefix is a **routing header on a data structure**, formatted after the guide-RNA spacer: **not a molecule, not delivered, no Cas protein present** [Curzi, 2026y].

Why 20 nucleotides, computed rather than asserted. The address-space arithmetic against genome size:

```
4^20 = 1,099,511,627,776 distinct 20-mers
human genome ~3.1e9 bp; both strands ~6.2e9 positions

expected occurrences of a random 20-mer = 6.2e9 / 4^20 = 0.005639
address space / genome positions        = 4^20 / 6.2e9  = 177.34

L = 12   4^L = 1.67772e+07   expected hits in 6.2e9 = 369.5
L = 16   4^L = 4.29497e+09                         =   1.444
L = 18   4^L = 6.87195e+10                         =   0.0902
L = 20   4^L = 1.09951e+12                         =   0.005639
L = 24   4^L = 2.81475e+14                         =   0.0000220
```

Below 16 nt an address is not an address — it names hundreds of places. At 20 nt the expected number of spurious matches is about one in 177 and the address is effectively unique. That is the whole justification; it is arithmetic, and it is the same arithmetic that sets hash lengths in ordinary computing.

Precedent. Older correspondence systems assigned every operation a name, a sigil and a place in a table before the operation was performed. **This records that address-first architecture was arrived at independently and preserved; it is not evidence about spacer length or about anything molecular** [Curzi, 2026w].

HONESTY, restated in place: nothing is delivered, therefore nothing is edited. The prefix is a field in a file. Native processing of a crRNA array (Zetsche et al., *Cell*, 2015; Fonfara et al., 2016; **named without keys**) is cited for one thing only — that a *multicast destination list* is a real biological format — **and no claim is made that any such parsing occurs in response to sound.**

36.4.9 4.9 Stage 4 — sequence to tone

Four symbols must become four distinguishable tones. **Two routes are defensible and the choice must be published, because they produce different sounds.**

Route A — measured infrared spectra with fixed octave transposition. Take published, measured infrared absorption frequencies for the four nucleobases and transpose them down by a fixed integer number of octaves into the audible band. **The values must be taken from a cited spectral database and reproduced in the disclosure; this paper does not supply them, because supplying a spectral peak that has not been looked up would be exactly the fabrication the honesty clauses forbid.**

What can be computed without any citation is the transposition arithmetic. The infrared fingerprint region is conventionally 400 to 4,000 cm^{-1}. Converting by $f = c\tilde{\nu}$ with $c = 2.99792458 \times 10^{10}$ cm/s and dividing by $2^{36} = 68{,}719{,}476{,}736$:

```
 400 cm^-1  ->  1.199170e+13 Hz  ->  /2^36  =    174.5022 Hz
1000 cm^-1  ->  2.997925e+13 Hz  ->  /2^36  =    436.2554 Hz
1600 cm^-1  ->  4.796679e+13 Hz  ->  /2^36  =    698.0087 Hz
2000 cm^-1  ->  5.995849e+13 Hz  ->  /2^36  =    872.5109 Hz
3000 cm^-1  ->  8.993774e+13 Hz  ->  /2^36  =   1308.7663 Hz
4000 cm^-1  ->  1.199170e+14 Hz  ->  /2^36  =   1745.0218 Hz
```

A 36-octave transposition maps the entire infrared fingerprint region onto 174.50 Hz to 1745.02 Hz — a span of exactly $\log_2 10 = 3.321928$ octaves, inside the region of most acute human hearing. That is why 36 is the reference transposition: **not tuned to taste, but the integer that lands the whole region where a person can hear all of it.** Every value published carries its j, and here $j = -36$. **Octave transposition preserves every frequency ratio exactly** [Shepard, 1982], so the relationships among the four tones after transposition are identical to those before it: **nothing is added and nothing is distorted; it is a change of register, the only purely presentational operation in the chain.**

Route B — the essence law. $f(n) = kn^2$ with $k = (9/8)^2/\varphi^2 = 0.483425732988414$ Hz. **k is a chosen reference constant, not a measured one, and the square law is a postulate about essence structure; both are badged here and wherever they appear** [Curzi, 2026, Markowsky, 1992]. Route B needs a defensible ordinal on the four bases; **an arbitrary ordering yields arbitrary frequencies, so the ordering must be argued in print or Route B is unavailable.**

Route B carries a trap that is invisible until it is heard. The obvious index assignment is n = 1, 2, 3, 4, whose odd parts are 1, 1, 3, 1. **Three of the four bases collapse onto a single pitch class.** Two indices share a pitch class if and only if they have the same odd part [Curzi, 2026]; $f(2) = 4f(1)$ is exactly two octaves and $f(4) = 16f(1)$ exactly four. **A listener attending to pitch class rather than absolute register hears three of the four symbols as the same note.** Two fixes, either of which must be published if used:

- **Do not fold.** Use a single global transposition j for the whole alphabet so that absolute register carries the symbol.
- **Use odd indices.** Assign n = 1, 3, 5, 7 — four distinct pitch classes. At $j = 8$:

```
n = 1   f = 0.4834257330 Hz   ->  123.7570 Hz
n = 3   f = 4.3508315969 Hz   -> 1113.8129 Hz
n = 5   f = 12.0856433247 Hz  -> 3093.9247 Hz
```

```
n = 7   f = 23.6878609164 Hz  -> 6064.0924 Hz

intervals from n=1:  0.0, 3803.9100, 5572.6274, 6737.6518 cents
minimum gap: 990.0559 Hz
```

All four land inside 20 Hz to 20 kHz. **The odd-index alphabet is the reference default for Route B.**

36.4.10 4.10 The Gabor constraint on symbol rate

DERIVATION (classical). $\Delta t\,\Delta f \geq 1/(4\pi) = 0.07957747154594767$ [Gabor, 1946, Curzi, 2026z]. **This contains no quantum mechanics and it is measurable on an oscilloscope. A short sound cannot have a precise pitch.** Applied to the four-tone alphabet with a minimum gap of 990.0559 Hz:

```
Gabor floor:  dt_min = 0.07957747 / 990.0559 Hz = 0.08037675 ms

chosen dt      frequency resolution at that window     safety factor over floor
   1 ms                 79.5775 Hz                          12.44x
  10 ms                  7.9577 Hz                         124.41x
  20 ms                  3.9789 Hz                         248.83x
  50 ms                  1.5915 Hz                         622.07x
```

The reference default is **20 ms per symbol, a safety factor of 248.83 over the floor.** The factor is large deliberately: **the floor is where two tones become mathematically indistinguishable, and a signal that must be *audibly* distinguishable to an ordinary listener in an ordinary room needs far more margin than one that must merely be separable by a matched filter. The no-subliminals constraint sets the symbol rate, not the mathematics** — a rate near the Gabor floor produces a signal whose symbols no listener can resolve, **which is a concealed signal by another route.**

At 20 ms per symbol a 1,776-base transmission (1,536 payload plus 240 address) runs 35.520 seconds.

```
dt = 10 ms  ->  17.760 s        dt = 25 ms  ->  44.400 s
dt = 20 ms  ->  35.520 s        dt = 50 ms  ->  88.800 s
```

36.4.11 4.11 The structure is carried in the harmonic

What is claimed: the relational structure of the corpus — which directives contradict which — is realised as interval structure in the rendered chord, audibly, so that a listener hearing the chord is hearing the structure and not a representation of it.

What is not claimed: that this structure does anything to any listener, any cell, or any transcript. This is a design claim of the same kind as "this file is sorted."

The corpus has $N/2$ contradictory pairs; assign each pair the indices $(m, 2m)$ with m odd. Then m and $2m$ have the same odd part, so **the two members of a pair occupy the same pitch class**; $f(2m)/f(m) = (2m)^2/m^2 = 4$ exactly, so **they are separated by exactly two octaves, 2400.0 cents, in every pair without exception**; and distinct pairs get distinct odd parts, so **each contradiction occupies its own pitch class** [Curzi, 2026].

For twelve directives paired at $(1,2),(3,6),(5,10),(7,14),(9,18),(11,22)$, the odd parts are $1,1,3,3,5,5,7,7,9,9,11,11$ — **six distinct classes for twelve tones, exactly the $\lceil N/2 \rceil$ bound realised as design rather than suffered as limitation.**

A contradiction is rendered as an octave. The two directives that cancel are the same note at two heights. A listener does not need the encoding table to hear that the corpus is built in

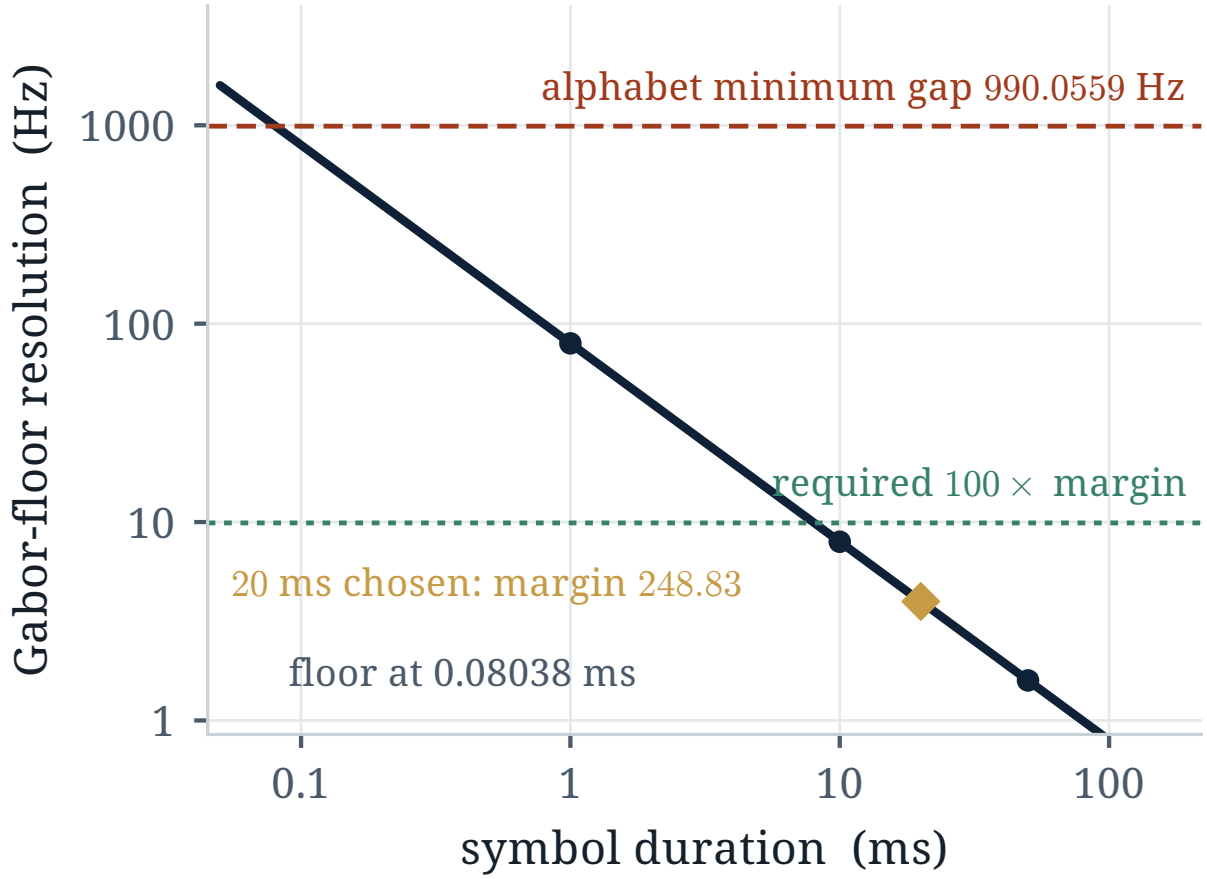


Figure 36.2: **The symbol rate is set by audibility, not by the mathematics:** achievable frequency resolution against symbol duration, with the Gabor floor $1/(4\pi\,\Delta t)$, the minimum alphabet gap of 990.0559 Hz, the self-imposed 100-fold margin requirement, and the chosen 20 ms operating point at a margin of 248.83. **The floor is where two tones become mathematically indistinguishable; the requirement is set an order of magnitude above it because the signal must be audibly resolvable by an ordinary listener.** Exact evaluation of the stated formulas at the alphabet of Section 4.9. Generated by `figs/src/CCP-036_fig02.py`.

opposed pairs — the opposition is the most audible interval in music [Shepard, 1982, Demany et al., 2013].

The pitch classes themselves, folded into one octave above the lowest tone, are exact Pythagorean and just intervals, because every essence ratio is a ratio of squares:

```
odd part  1   ->    0.0000 cents      f(1)/f(1) = 1
odd part  3   ->  203.9100 cents      f(3)/f(1) = 9   = 8 x (9/8),    the whole tone
odd part  9   ->  407.8200 cents      f(9)/f(1) = 81  = 64 x (81/64), the ditone
odd part  7   ->  737.6518 cents      f(7)/f(1) = 49  = 32 x (49/32)
odd part  5   ->  772.6274 cents      f(5)/f(1) = 25  = 16 x (25/16)
odd part 11   -> 1102.6359 cents      f(11)/f(1)= 121 = 64 x (121/64)
```

The constant k cancels out of every one of these; only the ratios of squares survive. Class 3 is exactly the Pythagorean whole tone $9/8$ above class 1, and class 9 exactly the ditone $81/64$ above it — **k's own numerator reappearing as an interval in the rendered chord. That is exact, not approximate**, and the intervals are the historically documented ones [Barbour, 1951].

The honest reading. What has been built is a *notation* in which a structural property of a text maps bijectively and audibly onto a structural property of a sound. **That is a real, checkable achievement of encoding design. It is not a mechanism, it is not a channel, and it does not become one by being elegant.** The open problem stands untouched by everything in this section [Curzi, 2026y]. **And whether ratio structure of this kind is perceived as the theory supposes is itself contested: consonance depends on the spectrum of the sound and not on the ratio alone** [Sethares, 2005, Plomp and Levelt, 1965, Terhardt, 1974].

36.4.12 4.12 Three architectures, compared

	Steganographic	Masked	Carried (this framework)
Where the content is	in a cover signal	below perceptual threshold	**it is the signal**
What the listener hears	the cover	the masker	**the content**
Can they recover it?	only with the key	not at all by ear	**yes, by ear plus tables**
Can they evaluate it?	no	no	**yes, before exposure**
Can they decline it?	no	no	**yes, at any point**
Consent status	impossible	impossible	**obtainable**
Status here	**excluded**	**excluded**	permitted

The middle column is where every historical "subliminal" technique sits, and it is excluded for the same reason as the first: in both, the recipient's inability to perceive the content is doing the work. The right-hand column is not a weakened version of the other two but a different architecture with a different failure mode, **and its worst outcome is that the listener reads the corpus and declines. That is a safety property, not a weakness.**

36.4.13 4.13 Why the contradictory corpus is the second half of the same argument

Section 4.1 protects the listener at the delivery layer; Section 4.4 protects them at the content layer, **structurally rather than by promise.** The reason both are required is visible in the alternative.

A single-directive corpus — one instruction, encoded, addressed, rendered, disclosed — satisfies every clause of Section 4.1. It is fully audible, fully published, and can be read, evaluated and declined. **And it is still a command, because a corpus with one member has a resultant equal to its member.** Its surplus is $F = \ln 1 = 0$: **a corpus of one is a parallel corpus, and parallel corpora carry no surplus at all.**

Disclosure makes a command refusable. Contradiction makes it not a command. They fail in different circumstances: **disclosure fails when the listener does not read carefully, is tired, is deferential, or trusts the operator, whereas a field of mutually exclusive directions has no resultant whether or not anyone is paying attention.**

```
CONTRADICTION REQUIREMENT

For every directive v in the corpus, its explicit negation −v is also in the corpus,
stated with equal force, equal length, and equal position in the rendering order.

No directive appears without its opposite.
No opposite is weakened, hedged, qualified, or given the last word.
The rendering order is published and is not a ranking.
```

"Equal force" is checkable and must be checked: if the pair is *("BUILD SLOWLY AND FINISH", "start again, maybe")* the corpus is not balanced, **it is a command with an alibi.** Appendix A shows the check being applied — **and failing on two pairs, which is reported rather than repaired quietly.**

36.4.14 4.14 The concordance, with each column's role stated

The evidence column carries evidence; the precedent column carries precedent. They are not two verifications of the same proposition [Curzi, 2026w].

BRIDGE 1 — a corpus may be deliberately built to contradict itself. *Evidence:* none is required; this is a property of texts and is verified by reading them. *Precedent:* the wisdom-literature tradition contains explicit adjacent contradiction — Proverbs 26:4–5 places two directives about answering a fool in direct opposition, preserved together, in sequence, by a redactor who plainly noticed; the Talmudic practice of *machloket* preserves minority opinions alongside the ruling rather than deleting them; the koan tradition constructs prompts specifically so that the corpus cannot resolve them. *Integration:* these records establish **precedent** — that a contradictory corpus was arrived at independently, repeatedly, across unrelated traditions, and deliberately preserved. **That independent arrival is a genuine argument that the structure is in the territory rather than only in this framework's instrument. It is not evidence that the structure does anything.** *Criteria:* C1 **passes** — the adjacent-opposition structure is identical, not merely similar; C2 **passes** — it predicts that such corpora preserve the rejected option in the record rather than resolve it, checkable against any of these traditions; C3 **passes**. **Status: established as a bridge about corpus construction. It says nothing about acoustics or biology.**

BRIDGE 2 — the receiver supplies the selection. *Evidence:* the arithmetic of Section 4.4 — zero resultant, $1/\sqrt{N}$ mean, $\ln N$ selection cost. **This is mathematics, and it is evidence about the model, not about people.** *Precedent:* oracle traditions are organised precisely this way — a fixed corpus supplied by the system, the selection and the reading supplied by the querent; structured decks in which every element carries its own reversal, so that the deck as a whole has no resultant meaning. *Integration:* **precedent for the architecture, not evidence for any efficacy of it. The useful observation is technical and neutral: these systems are built so that the corpus cannot determine the outcome, and the selection step is the only place the outcome can come from.** *Criteria:* C1 **passes**; C2 **passes** — it predicts that any attempt to rank the corpus members converts the oracle into a command; C3 **passes**. **Status: established as a bridge about information architecture only.**

BRIDGE 3 — encoding a message so it cannot be read is a legitimate method. REJECTED. This bridge is stated in order to be refused in print, **because a concordance method that never rejects anything is not a method** [Curzi, 2026]. *Precedent:* it is strong. Sigil methods deliberately render a sentence unreadable and hold the illegibility to be functionally necessary; alchemical cover names obscured operations systematically for centuries; mystery traditions bound initiates to secrecy. **The older corpus concealed, routinely, and this paper does not select the convenient half of the record.** *Evidence for concealment being necessary to any effect:* **none**, and the open-label result points the other way for the one belief channel that is fully characterised. *Refusal:* **C1 fails on inspection — the correspondence offered is between "the operator cannot consciously read it" and "the recipient cannot know it," which are different relations**, since a sigil practitioner encodes their own sentence and knows exactly what it said. The precedent is about self-directed attention management, not concealment from a second party. **C2 fails: no prediction distinguishes the concealed case from the disclosed case. And the bridge fails the consent constraint independently of the criteria: a method that removes the recipient's ability to decline is out of scope regardless of whether it works. Status: REJECTED — recorded so that the rejection is on the record and the older corpus is reported accurately, concealment and all.**

36.4.15 4.15 What this section does not claim

THE OPEN PROBLEM. There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal [Curzi, 2026y]. Everything upstream is real: DNA data storage, guide-RNA-format addressing, native array self-processing, latent capability and its activation, epigenetic regulation, and mechanotransduction. **The channel for sequence-specific content is the question, and this section has built a transmitter for a receiver not demonstrated to exist.**

No clinical, medical or therapeutic claim is made anywhere. No claim that DNA is edited, altered or damaged — the architecture produces sound and cannot do otherwise. **No claim that any capability has been conferred**; nothing not already latent can be reached. **The essence frequencies are essence frequencies, not physical resonances**; k is chosen, and surplus as deployed is a lower bound [Curzi, 2026,?].

36.4.16 4.16 Requirements

REQ-4.1 — Audibility. Every component shall be audible to a person of ordinary hearing at every listening position, measured over the actual noise floor present. Verification: an A-weighted third-octave measurement at each position, published with the rendering.

REQ-4.2 — No concealment. No carrier, masker, embedding or time-compression; no content below 20 Hz or above 20 kHz intended to act on a listener; no psychoacoustic masking used as a delivery mechanism. **Verification: the published spectrum accounts for all energy in the rendering.**

REQ-4.3 — Full disclosure, in advance, unconditional. The packet shall contain the complete corpus text; the Stage 1 character encoding; the Stage 2 nucleotide string as actually generated; the Stage 3 address list; the Stage 4 tuning table with band and integer j for every frequency; symbol duration; rendering order; level. **A recipient shall be able to reconstruct the signal from the packet.**

REQ-4.4 — Stop and leave. An off switch exists, its location is disclosed, and any person exposed may operate it without asking. **Exposure is never a condition of remaining in a space, a study, or a relationship.**

REQ-4.5 — Contradiction. Every directive appears with its explicit negation, at equal length, equal force, and symmetric position in the rendering order. **The rendering order is published and is declared not to be a ranking.**

REQ-4.6 — No claim beyond the encoding. Any published or spoken description states the open problem and states that no clinical, medical or therapeutic claim is made.

REQ-4.7 — Reversibility of the chain. Stages 1 to 4 are bijections on the payload. **A round-trip test shall reproduce the corpus character for character, and the result is published.**

REQ-4.8 — Published encoding. Character set, byte width and byte order are published. **The one-bit-per-base mapping and the actual sequence chosen are both published, since the mapping is one-to-many.**

REQ-4.9 — Address length. Addresses are 20 nt. Any other length is published together with its expected-collision calculation against the reference genome size.

REQ-4.10 — Address neutrality. Address selection is published in advance of rendering and does not change between arms of any comparison.

REQ-4.11 — Tuning route declared. Route A requires a cited source for every spectral value; Route B requires the index ordering to be argued in print and the pitch-class check to be shown.

REQ-4.12 — Band and j published for every frequency. No exception.

REQ-4.13 — Global transposition. A single integer j applies to the whole alphabet so all frequency ratios are preserved exactly. Per-tone transposition is permitted only if the resulting ratio distortions are published.

REQ-4.14 — Gabor margin. $\Delta t \geq 100 \cdot (1/(4\pi))/\Delta f_{\min}$ — **a hundredfold margin over the mathematical floor, set by audible distinguishability rather than by separability.** The computed margin is published.

REQ-4.15 — Level and duration limits. Rendering level shall not exceed the applicable occupational exposure limit for the exposure duration. At a recommended limit of 85 dBA for 8 hours with a 3 dB exchange rate:

```
 80 dBA  ->  25.3984 h
 85 dBA  ->   8.0000 h
 88 dBA  ->   4.0000 h
 91 dBA  ->   2.0000 h
 94 dBA  ->   1.0000 h
 97 dBA  ->  30.0000 min
100 dBA  ->  15.0000 min
103 dBA  ->   7.5000 min
```

A 35.520 s rendering at 85 dBA is **0.123 per cent** of the 8-hour allowance; 100 repetitions is 59.2 minutes, **12.3 per cent** of it. **Compute the number; do not estimate it.**

REQ-4.16 — Essence honesty in print. Wherever the essence law appears in a disclosure packet, the badge appears with it: k **is chosen; the square law is a postulate; the output is an essence frequency, not a physical resonance.**

REQ-4.17 — Ordering defensibility. Route B requires the index ordering to be stated and justified. **An arbitrary ordering yields arbitrary frequencies and is not permitted.**

REQ-4.18 — No ranking by rendering. Loudness, duration, repetition count, spectral centroid and sequence position shall be equal across contradictory pairs. **Verification: a published per-directive table of all five, with the maximum pairwise asymmetry stated.**

REQ-4.19 — Corpus size. $N \geq 4$ and N even. **Below 4 the $1/\sqrt{N}$ argument has no content; an odd N cannot be closed under negation.**

REQ-4.20 — The refusal is recorded. A person who reads the packet and declines is a successful outcome of the protocol and is recorded as one. Declination rates are reported alongside any other result.

36.4.17 4.17 Failure modes

FM-4.1 — THE CARRIER. *Signature:* the rendering contains a bed, pad, ambience or masking noise, and the encoded tones sit below it — spectral analysis shows encoded energy 10 dB or more under a broadband component in the same critical band. *Why:* the raw rendering is harsh; a pad makes it pleasant, and is louder than the data before anyone notices. *Fix:* measure per-component level in third-octave bands at each listening position. **"It measures above threshold in an anechoic chamber" is not the test — the test is the room the listener is in.**

FM-4.2 — THE ULTRASONIC ANNEX. *Signature:* content above 18 kHz or below 25 Hz, "for completeness" or "for the cells rather than the ears". *Fix:* transpose by octaves until the value is inside the band and publish the j. **Content a person cannot hear is content a person cannot decline, and having put it there for a non-auditory reason makes it worse, not better.**

FM-4.3 — THE SUMMARY DISCLOSURE. *Signature:* the disclosure describes the corpus rather than reproducing it; the nucleotide string is unpublished; the address list is "available on request." *Fix:* publish the packet. **The operative test: can a recipient reconstruct the signal from what they were given? If not, they have been told about the signal, which is a different thing.**

FM-4.4 — THE WEIGHTED PAIR. *Signature:* directive and counter-directive both present, and one is 40 per cent longer, or sits later, or repeats twice as often, or is rendered 3 dB louder. *Why:* the author has a view, and it leaks in through the parameters nobody balanced. *Fix:* tabulate length, level, duration, repetition and position per directive and publish the maximum pairwise asymmetry. **A weighted contradictory corpus is a command with an alibi, and its resultant is no longer zero — not as a metaphor, but exactly as Section 4.4's arithmetic says when the vectors stop cancelling.**

FM-4.5 — THE PITCH-CLASS COLLAPSE. *Signature:* Route B with indices 1, 2, 3, 4; listeners report the signal sounds like two notes; decoding by ear fails for three of the four symbols. *Fix:* a single global j, or odd indices 1, 3, 5, 7. **Publish which.**

FM-4.6 — GABOR SMEAR. *Signature:* symbol duration set for total runtime rather than for resolution; adjacent symbols bleed into one another; decoding by matched filter succeeds while decoding by ear fails. *Fix:* REQ-4.14. **The matched filter succeeding while the ear fails is itself a no-subliminals violation, not merely a quality problem: the content is present and the listener cannot resolve it, which is the middle column of Section 4.12.**

FM-4.7 — THE SINGLETON CORPUS. *Signature:* one directive, fully disclosed, beautifully encoded, exemplary in every respect except that it is one directive. *Fix:* $F = \ln 1 = 0$. **Disclosure alone does not make a command not a command. It makes it a refusable command, which is a weaker protection and is not the one this protocol offers.**

FM-4.8 — PRECEDENT PRESENTED AS EVIDENCE. *Signature:* "the tradition knew this, and modern research confirms it," where the modern citation measures something adjacent. *Fix:* state each column's role in the sentence that draws the concordance [Curzi, 2026w].

FM-4.9 — THE IMPLIED MECHANISM. *Signature:* "the narrative is carried in the harmonic" appears without the qualification that this is a statement about encoding; the reader finishes the section believing something has been shown to reach a cell. *Why:* **Section 4.11 is genuinely elegant, and elegance reads as explanation.** *Fix:* the open problem appears in the same passage, every time. **The encoding being beautiful is not evidence that anything receives it.**

FM-4.10 — THE CONDITIONAL EXPOSURE. *Signature:* exposure occurs in a room a person needs to be in, and the off switch is behind a desk, or the person is a subordinate, or leaving costs something. *Fix:* **reversibility is a property of the circumstance and not only of the equipment.** If declining costs the person something, the protocol has not obtained consent, whatever was signed.

FM-4.11 — THE UNDEFENDED ORDERING. *Signature:* Route B with indices assigned in the order the directives happened to be typed, and the rendered intervals reported to four decimal places. *Fix:* REQ-4.17. **An arbitrary ordering yields arbitrary frequencies, and reporting them to four decimals makes the arbitrariness look like precision.**

FM-4.12 — LEVEL CREEP. *Signature:* sessions get longer or louder because the effect "seems stronger"; no dosimetry is published. *Why:* with no measured effect to calibrate against, the only feedback is impression, and impression rewards more. *Fix:* REQ-4.15, computed and published per session. **With no measured effect, "seems stronger" is not weak evidence, it is no evidence, and it is the exact input a protocol should be least responsive to.**

36.5 5 The research protocol, and a preregistration that prices its own nulls

Section 4 built a transmitter. This section specifies what may be done with it, as an operational procedure. **Every honesty clause is written into the step it constrains rather than collected at the end, because a clause at the end is a clause that can be detached from the procedure and lost.** The consent requirements are steps P0 and P1 for the same reason: **they are not advice appended to a method, they are the first two operations of it, and a run that skips them has not run the protocol.**

36.5.1 5.1 The open problem, carried forward and stated first

THERE IS NO ESTABLISHED MECHANISM BY WHICH A CELL DEMODULATES SEQUENCE-SPECIFIC INFORMATION FROM AN ACOUSTIC SIGNAL [Curzi, 2026y]. **It is restated here in full and not by reference, because a reader who arrives at a procedure section looking for a method is exactly the reader most likely to skip the statement and exactly the reader who most needs it.**

What is established, and it is a great deal. Mechanically gated ion channels are established and Nobel-recognised: mechanical force opens them, ions flow, signalling follows, and sound is mechanical, **so a real pathway exists from a pressure wave to a cellular signal.** Mechanical force at the cell surface reaches the nucleus and changes transcription (Tajik et al., *Nature Materials*, 2016). Sonogenetic channel activation exists (Ibsen et al., 2015; Duque et al., 2022). Biology builds exquisite mechanical receivers — cochlear hair cells are the existence proof. Arbitrary data can be stored in and recovered from nucleotide sequence (Church, Gao and Kosuri, 2012). Targeting and editing are separable (Qi et al., 2013; Gilbert et al., 2014). Latent capability can be activated by regulatory signal without adding information (Takahashi and Yamanaka, 2006). **All named without keys, per Section 3.10.**

What is not established is the single link between them. Every item above is either an input pathway that carries *force* or a data operation that carries *sequence*. **Nothing in the established literature carries sequence over force.** A pressure wave that opens a mechanically gated channel does so by deforming the membrane; **the channel does not care what melody deformed it.** That is not a gap in the reading of the literature; it is the literature's actual state, and it is one well-specified falsifiable question rather than a vague unknown [Curzi, 2026y,?].

The consequence is total. This is a research procedure and is presented as one. It is not a treatment. It is not a technique with a demonstrated effect. It is not a therapy, an intervention, a modality, or a protocol in any clinical sense. It is a specification for an experiment whose primary purpose is to find out whether there is anything there, **written so that the answer *no* is as reportable as the answer *yes*.**

36.5.2 5.2 A magnitude gap, computed, which closes one route

The bulk-compression route. For a plane acoustic wave in an aqueous medium the longitudinal strain amplitude is exactly $s = p/K$ with $K = 2.2 \times 10^9$ Pa the bulk modulus of water. **That is elementary continuum mechanics, not a model of anything biological.**

```
p (Pa)          strain              equivalent context
```

```
  0.02          9.09e-12           60 dB SPL in air, conversational
  0.2           9.09e-11           80 dB SPL
  0.3557        1.62e-10           85 dBA, the 8-hour occupational limit
  2.0           9.09e-10          100 dB SPL
 20.0           9.09e-09          120 dB SPL, threshold of pain
2000.0          9.09e-07          160 dB SPL
```

Membrane-tension gating of mechanosensitive channels involves areal strains of order 10^{-2}. Reaching 10^{-2} by bulk compression requires:

```
p = 1e-2 x K = 2.2e7 Pa = 22 MPa

ratio to an 85 dBA tone:  22e6 / 0.355656 = 6.1858e7
in decibels:              20 log10(6.1858e7) = 155.83 dB
required level:           85 + 155.83 = 240.83 dB SPL
```

And 240.83 dB SPL is not merely difficult, it is impossible in air. The maximum undistorted acoustic pressure amplitude in air at one atmosphere is bounded by the ambient pressure itself:

```
p_max = 101,325 Pa  ->  20 log10(101325 / 20e-6) = 194.09 dB SPL
strain at p_max     =  101325 / 2.2e9 = 4.6057e-05
still short of 1e-2 by a factor of 217.12
```

The bulk-compression route from an audible acoustic field to mechanosensitive gating is closed by more than seven orders of magnitude, and it stays closed at the physical ceiling of sound in air. This is stated plainly because it is the strongest quantitative thing that can honestly be said in the neighbourhood, and it says *no*.

What is not closed, and must be listed so the closure is not overread. The established sonogenetics results do not work by bulk compression. Focused ultrasound acts through radiation force, acoustic microstreaming, localised shear at interfaces, cavitation, and resonant response of structures whose dimensions match the wavelength — **all non-linear, all localised, and all capable of producing local strains vastly larger than the bulk figure. Those routes are open. But none of them has been shown to be sequence-selective, and sequence-selectivity is the whole of the question. Closing the bulk route narrows the search; it does not answer it.**

For calibration: the Casimir pressure at biological separations is 1.300126×10^{-3} Pa at 1 μm and 1.300126×10^{-15} Pa at 1 mm, against 1.300126×10^{5} Pa at 10 nm [Casimir, 1948, Lamoreaux, 1997, Curzi, 2026aa]. **At any biological or architectural scale the Casimir force is utterly negligible; it does nothing in a cell and nothing in a room, and it is mentioned here only so that nobody reaches for it as the missing mechanism.** The same discipline applied to the other standard exotic borrowing yields the same verdict [Morris and Thorne, 1988], and the general argument that a formalism valid at one scale carries no automatic warrant at another is Anderson [1972]'s.

36.5.3 5.3 Consent, reversibility and the right to decline — structural, not appended

Why they are structural rather than advisory, in this specific model. The whole architecture routes through the subject's own capability: nothing not already latent can be reached [Curzi, 2026y]; belief acts only through a real channel [Curzi, 2026ab]; a contradictory corpus can only be resolved by selection, and the selection belongs to the receiver (Section 4.4). **In every one of these the person is a component of the mechanism, not a target of it. Remove their knowledge, their ability to stop, or their ability to refuse, and you have not made the protocol more powerful — you have removed a component and the circuit is open.**

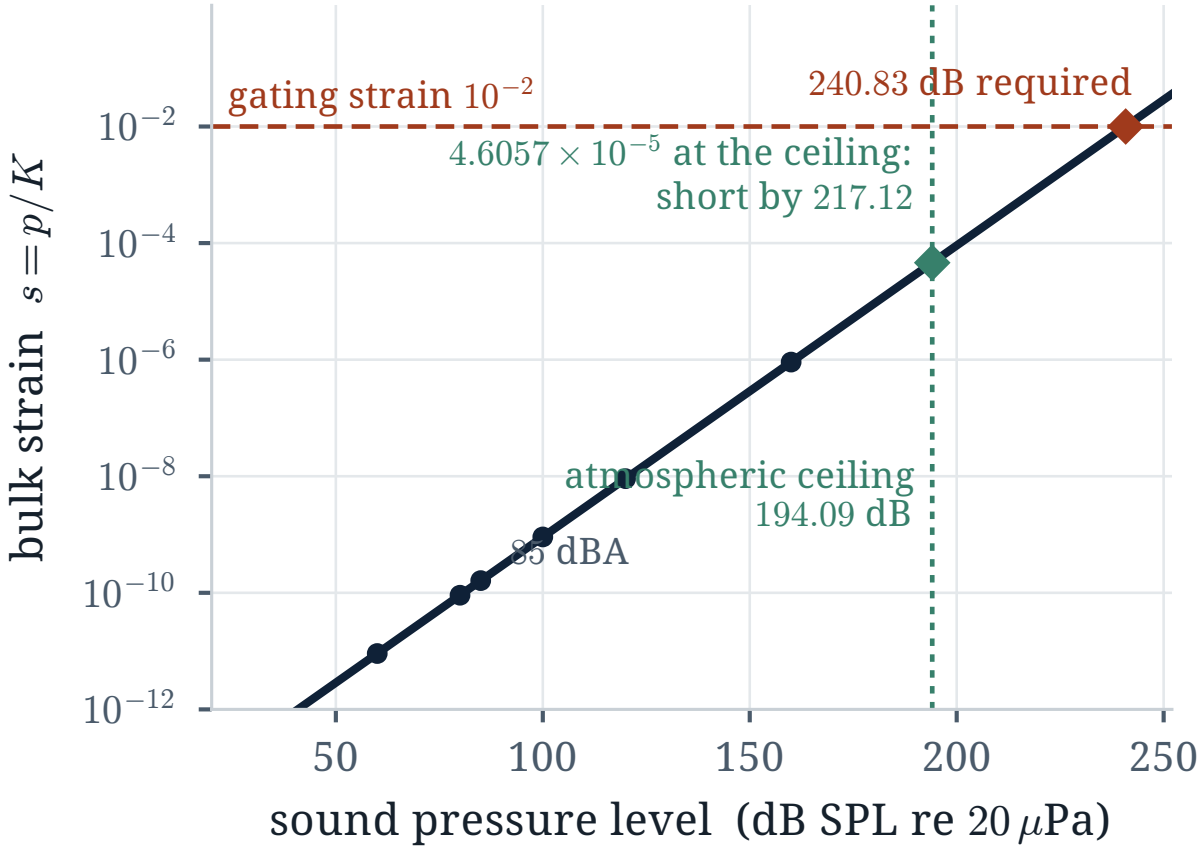


Figure 36.3: **The bulk-compression route is closed by more than seven orders of magnitude, and stays closed at the physical ceiling of sound in air:** longitudinal strain $s = p/K$ in an aqueous medium against sound pressure level, with the 10^{-2} gating strain, the 240.83 dB SPL it would require, and the 194.09 dB atmospheric ceiling set by the ambient pressure itself. **At the ceiling the achievable strain is 4.6057×10^{-5}, short by a factor of 217.12.** Exact evaluation of $s = p/K$ at $K = 2.2 \times 10^{9}$ Pa; **the 10^{-2} gating figure is an order-of-magnitude value from the membrane-mechanics literature and is treated as such.** Generated by `figs/src/CCP-036_fig03.py`.

C1 — Informed in the strong sense. The participant receives the full disclosure packet plus the statement of the open problem plus an explicit statement that no clinical, medical or therapeutic effect is claimed or expected. **They are able to reconstruct the signal from what they are given.**

C2 — Declination is free and is a recorded outcome. Declining costs nothing and is not followed up. **A protocol whose declination rate is never reported is a protocol that has never told you what happened to the people who said no.**

C3 — Withdrawal at any point, without reason, with data destruction on request, including mid-session and after the fact.

C4 — Reversibility as a property of the signal, not a promise about it. The exposure is sound; it stops when the sound stops. **Nothing is administered, ingested, injected, implanted or introduced; nothing persists after the acoustic field ends except the ordinary memory of having heard a sound.** This is not a safety assurance based on caution — **it is a statement about what the architecture is** (Section 4.3).

C5 — No conditional exposure. Exposure never occurs in a space a person must be in, and never as a condition of employment, treatment, membership, or participation in anything else.

C6 — Level and dose within occupational limits, computed per session, at a recommended limit of 85 dBA over 8 hours with a 3 dB exchange rate. **Published per session, computed, not estimated.**

C7 — No vulnerable-population recruitment for a procedure with no expected benefit. A research procedure with no demonstrated effect offers no therapeutic prospect, and recruiting people who might believe otherwise is the specific harm this clause exists to prevent. **Anyone who says they hope it will treat something is told, in the same conversation, that it will not and is not designed to, and their enrolment is declined if the belief persists.**

36.5.4 5.4 The two layers, and the firewall between them

The single most important design decision is that the open problem is tested where no person is exposed.

```
LAYER 1   IN VITRO.  Cultured cells. Tests the OPEN PROBLEM directly.
          Endpoint: transcript abundance at addressed loci vs matched non-addressed loci.
          Human participants: NONE.

LAYER 2   HUMAN, open-label, fully disclosed. Tests ONLY the established expectancy
          channel. Endpoint: self-report and behavioural measures, preregistered.
          Tests nothing about Layer 1 and is not evidence for it.
```

THE FIREWALL, and it is absolute: no Layer 2 result is ever presented as evidence for a Layer 1 claim. A person who reports that the corpus was useful has reported something about a corpus they read, a sound they chose to hear, and an expectation they held. **Every one of those is an established belief channel and none of them is a cell demodulating a sequence** [Curzi, 2026ab,ac]. **The firewall runs the other way too:** a Layer 1 result, positive or negative, says nothing about whether a person found the corpus worth reading.

36.5.5 5.5 Warrant

Claim	Status	Basis
Mechanical force opens ion channels and reaches transcription	**CITED (measured)**	2021 Nobel-recognised work; Tajik et al. 2016, named without keys
Ultrasound can activate specific channels	**CITED (measured)**	Ibsen et al. 2015; Duque et al. 2022, named without keys
A cell demodulates sequence-specific content from sound	**NOT ESTABLISHED. THE OPEN PROBLEM.**	[Curzi, 2026y]
Bulk-compression route to gating strain is closed	**ARITHMETIC**	Section 5.2, $s = p/K$
Non-linear and localised ultrasound routes remain open	**OPEN**	Section 5.2
The protocol is a research procedure, not a treatment	**DEFINITIONAL**	Section 5.1
Consent, reversibility, declination are protocol steps	**STRUCTURAL REQUIREMENT**	Section 5.3
Any clinical, medical or therapeutic effect	**NOT CLAIMED**	Section 3.9
DNA edited, altered or damaged	**NOT CLAIMED — architecturally impossible**	Section 4.3
Any capability conferred	**NOT CLAIMED**	[Curzi, 2026y]
Casimir relevance at biological scale	**NONE**	Section 5.2, computed

36.5.6 5.6 The procedure, step by step

Each step carries three lines: **DO**, **CLAIM**, and **NOT CLAIMED**, stated in place.

P0 — DECLARE THE STUDY AS RESEARCH. *DO.* Register the study publicly as a research procedure investigating whether an acoustic signal carrying sequence-formatted data produces any locus-specific transcriptional difference relative to a spectrally matched scrambled control; name the open problem in the registration abstract [Nosek et al., 2018]. *CLAIM.* That a question has been posed in a form that can be answered. *NOT CLAIMED.* Anything about the answer. **A preregistration is a commitment, not a result.**

P1 — CONSENT AND DISCLOSURE (Layer 2 only). *DO.* Deliver the disclosure packet plus the Section 5.3 statements; obtain consent; record declinations; apply C1 to C7 in full. *CLAIM.* That the participant knows what the signal is, what is claimed, and what is not. *NOT CLAIMED.* That consent makes any subsequent claim stronger. **Consent is a precondition for running, not a warrant for concluding.**

P2 — CONSTRUCT THE CORPUS. *DO.* Write N directives, N even, $N \geq 4$, closed under negation, balanced for length, force, position and repetition; run the balance check and **publish the maximum pairwise asymmetry including any failures.** *CLAIM.* That the corpus spans rather than points and its resultant is zero by construction. *NOT CLAIMED.* That the corpus content is true, wise, good for anyone, **or has any property beyond being contradictory and published.**

P3 — ENCODE. *DO.* Text to 8-bit bytes to one bit per base, choosing among the two bases available at each position so as to avoid homopolymer runs and maintain GC balance; run the round-trip verification and publish the result. *CLAIM.* That the corpus is losslessly represented as a nucleotide string. *NOT CLAIMED.* That representing text as a nucleotide string gives it any biological property. **A string over $\{A, C, G, T\}$ in a file is a data structure. It is not DNA and does not become DNA by being formatted like it.**

P4 — ADDRESS. *DO.* Prefix each payload with a 20-nt address; publish the address list and the expected-collision figure of 0.005639; fix addresses before rendering, identical across all

arms. *CLAIM.* That the transmission carries a routing header of established format and that address length is justified arithmetically. *NOT CLAIMED.* **That the address addresses anything. Nothing is delivered; no Cas protein is present; no locus is bound.**

P5 — TUNE AND RENDER. *DO.* Declare Route A or Route B; publish band and j for every frequency; set symbol duration to satisfy the Gabor margin and publish the computed margin; render. *CLAIM.* That the rendered audio is a lossless, fully audible representation of the addressed corpus at a published symbol rate. *NOT CLAIMED.* **That the essence frequencies are physical resonances; k is chosen and the square law is a postulate. That the tuning has any effect on anything.**

P6 — BUILD THE CONTROL ARMS. *This is the step the whole study turns on.*

```
ARM A   ENCODED       the addressed corpus, rendered
ARM B   SCRAMBLED     the SAME symbol multiset, permuted, rendered identically.
                      Same duration, same spectrum, same loudness, same symbol rate,
                      same disclosure. Carries no message.
ARM C   TONE-MATCHED  a signal with the same long-term spectrum and level but no
                      symbol structure at all
ARM D   SILENCE       matched duration, matched setting, no signal
```

CLAIM. That Arm B differs from Arm A **in sequence and in nothing else**, so a difference between A and B is attributable to sequence or to nothing. *NOT CLAIMED.* That the arms are indistinguishable to a listener. **Arm B may well sound different — scrambling changes the melody even when it preserves the spectrum — and any Layer 2 blinding claim must be tested rather than asserted.** For Layer 1 this does not matter, because cultures do not have expectations.

Arm B is the biological form of the shuffled-input control [Curzi, 2026u], **and it is the single most important arm in the design: it is the control that can kill the result, and every other control can be passed by an artefact.**

P7 — EXPOSE (LAYER 1). *DO.* Cultured cells, r independent replicate cultures per arm, randomised to arm, exposure through a calibrated transducer with the acoustic field mapped and published **at the cell layer, not at the source**; temperature logged continuously; handling identical across arms. *CLAIM.* That a measured acoustic field of published amplitude and spectrum was present at the cells. *NOT CLAIMED.* That the field did anything. **Acoustic exposure deposits heat and induces fluid motion; both change transcription; both are present in Arms A, B and C equally, which is precisely why Arm B rather than Arm D is the comparator for the primary endpoint.**

P8 — MEASURE (LAYER 1). *DO.* Quantify transcript abundance at n_{addr} addressed loci and n_{ctrl} matched non-addressed control loci, matched on baseline expression level, chromatin accessibility class and chromosome; analysts blind to arm. The primary statistic is a difference-in-differences:

```
DiD = ( mean log2 expr[addressed, arm A] − mean log2 expr[control loci, arm A] )
    − ( mean log2 expr[addressed, arm B] − mean log2 expr[control loci, arm B] )
```

CLAIM. That the DiD isolates any addressed-locus-specific difference between the encoded and scrambled signals, net of any global transcriptional response to sound. *NOT CLAIMED.* That a non-zero DiD would identify a mechanism. **It would establish that something sequence-dependent happened and would leave the mechanism entirely open — which is a real result and should be reported as exactly that size of result.**

P9 — EXPOSE AND MEASURE (LAYER 2). *DO.* Open-label, fully disclosed, **randomised exposure schedule**, preregistered self-report and behavioural endpoints, T periods per participant with a preregistered train/held-out split. *CLAIM.* Only that a disclosed belief-mediated channel is being operated and that any effect observed is consistent with expectancy. *NOT CLAIMED.* **Nothing about the cell. A Layer 2 effect is an expectancy effect until and unless Layer 1 says otherwise.**

P10 — FREEZE. *DO.* Fit on the training window only; freeze every parameter in a timestamped, publicly hashed record; then, and only then, open the held-out window. *CLAIM.* That the held-out numbers are out-of-sample. *NOT CLAIMED.* That freezing improves the model. **It does not. It makes the report honest, which is a different and more valuable thing.**

P11 — COMPARE AGAINST THE FOUR BASELINES. *DO.* Report held-out performance against **persistence**, **direct-input**, **matched autoregressive**, and **shuffled-input**; report all five numbers whatever they say [Curzi, 2026u].

```
PERSISTENCE          predict Y_{t+1} = Y_t.  Kills: a slow-moving outcome that any
                     model with a lagged term will appear to predict.

DIRECT-INPUT         predict Y_{t+1} from S_t alone, no storage term. Kills: the case
                     where exposure and outcome move together contemporaneously and
                     no accumulation is occurring at all.

MATCHED AR(1)        predict Y_{t+1} from Y_t with a fitted coefficient. Kills: the
                     case where the storage law's advantage is only its lag structure
                     and the input contributes nothing.

SHUFFLED-INPUT       refit with the exposure series permuted. Kills: everything that
                     does not depend on WHEN the exposure occurred. This is the one
                     that matters and the one that is omitted.
```

NOT CLAIMED. Anything from a model that beats fewer than all four. **The conjunction is the gate, and Section 5.9 gives its measured false-positive rate so that passing it means something specific. The biological analogue of the shuffled-input control is Arm B, and the correspondence is exact: shuffling the input series destroys temporal ordering while preserving the marginal distribution; scrambling the corpus destroys symbol ordering while preserving the spectrum. Both ask the same question — does the order matter?**

P12 — PUBLISH, INCLUDING THE NULL. *DO.* Publish the preregistration, the freeze hash, all five held-out numbers, all arm results, the declination rate, the dosimetry, the disclosure packet, and the raw data. **Publish a null as prominently as a positive.** *NOT CLAIMED.* That publication makes a result true.

36.5.7 5.7 Falsification criteria, stated in advance

The protocol is worthless unless it can fail, and it must be written so that failing is easy [Popper, 1959, Platt, 1964, Meehl, 1967].

F-1 (Layer 1, primary). If the difference-in-differences between Arm A and Arm B at addressed loci is not distinguishable from zero at the preregistered alpha, **the hypothesis that this signal carries sequence-specific information into transcription is not supported by this experiment.**

F-2 (Layer 1, specificity). If a difference appears at addressed loci *and equally at matched control loci*, the effect is not addressed-specific. **This is the most likely positive-looking failure, because a global acoustic response to heat or streaming will lift everything.**

F-3 (Layer 1, control equivalence). If Arm B differs from Arm A only in ways also seen between Arm C and Arm D, the difference is not about sequence.

F-4 (Layer 2, held-out). If the fitted model does not beat all four baselines out-of-sample, the model is not supported for that participant.

F-5 (Layer 2, shuffled). If the shuffled-input model performs as well as the true-input model on held-out data, **the input timing carries no information and the result is dead regardless of in-sample fit.**

F-6 (cross-individual). If parameters do not recover across participants and across sites under identical operationalisation, the model is not measuring a stable quantity.

F-7 (the whole thing). If F-1 and F-4 both fail, **the protocol has returned a null and the null is the result. It is published in the same venue, at the same length, with the same prominence.**

And the rescue clause, closed in advance. The following are declared, before data collection, to be **non-permissible** responses to a null: adding an unregistered arm; re-specifying the endpoint; extending recruitment after looking; discovering that the corpus was wrong for these cells; discovering that the participants were unsuitable; discovering that the tuning route should have been the other one. **A hypothesis rescued by a post-hoc auxiliary is a hypothesis that has stopped being falsifiable** [Lakatos, 1970, Simmons et al., 2011, Gelman and Loken, 2013].

36.5.8 5.8 The two timescale constraints, computed

The magnitude argument of Section 5.2 is about amplitude. **There is a second, independent constraint about *time*, and it points the opposite way — which is exactly the sort of thing a protocol should know before it runs.**

DERIVATION (classical). $\Delta t \, \Delta f \geq 1/(4\pi)$ [Gabor, 1946, Curzi, 2026z]: any receiver integrating over a window τ distinguishes frequencies no finer than $1/(4\pi\tau)$.

integration window tau	finest resolvable df	symbols averaged at 20 ms/symbol
0.1 ms	795.774715 Hz	0.005
1 ms	79.577472 Hz	0.05
20 ms	3.978874 Hz	1
100 ms	0.795775 Hz	5
1 s	0.079577 Hz	50
10 s	0.007958 Hz	500
35.52 s	0.002240 Hz	1776
1 min	0.001326 Hz	3000
1 h	0.000022 Hz	180000

CONSTRAINT ONE — resolution, and it is not the obstacle. The alphabet's minimum gap is 990.0559 Hz, requiring a window of only 0.08037675 ms. **Any receiver with an integration time of a tenth of a millisecond or longer has more than enough frequency resolution to tell the four symbols apart. Time-frequency uncertainty does not close this route, and it would be dishonest to invoke Gabor as a reason the hypothesis fails.**

CONSTRAINT TWO — ordering, and this one bites hard. A receiver whose integration window exceeds the message duration sees the message's *composition* and not its *order*:

symbols transmitted	1,776
information sent at one bit per base	1,776 bits
distinct compositions (n_A + n_C + n_G + n_T = 1776)	C(1779, 3)
information recoverable from composition alone	log2 C(1779,3) = 29.803 bits
fraction of the transmission that survives averaging	29.803 / 1776 = 1.678%

A receiver slower than 35.520 seconds recovers at most 29.803 bits of the 1,776 sent — 1.68 per cent — and none of it is sequence. Sequence is the entire hypothesis, so an integrator that slow is not a candidate receiver at all.

The tension this creates is the useful output of the calculation and should be stated as a constraint on any proposed mechanism rather than as an argument for or against one. A candidate receiver must be fast enough to resolve 20 ms symbols in order, **and its downstream consequence must nonetheless persist long enough to change transcript abundance measurably. Those two requirements pull in opposite directions across roughly three orders of magnitude in time**, and **"the cell integrates the signal" is precisely the answer the arithmetic above forecloses.** The design consequence is registered: **the symbol duration must be published, and a version of the study at a second symbol duration is registered as a secondary arm**, because a real sequence-dependent effect should depend on symbol rate and an artefact generally should not.

36.5.9 5.9 Design characteristics, computed

These are design characteristics computed by simulation. They are not data and no result is reported here. The noise parameter is **assumed** for the purpose of a power calculation and **must be replaced by a measured value from pilot replicates before any real preregistration is filed**; this is stated because **an assumed variance presented as measured is the commonest way a power calculation lies.**

```
n_addr  = 12 addressed loci
n_ctrl  = 120 matched control loci
sigma   = 0.35   ASSUMED sd of log2 expression across replicate cultures
test    = difference-in-differences, permutation null over locus labels within arm
```

Analytic standard error and minimum detectable effect at two-sided $\alpha = 0.05$ and 80 per cent power ($z_{0.975} + z_{0.80} = 1.959964 + 0.841621 = 2.801585$):

r (replicates/arm)	SE(DiD)	MDE (log2 fold-change)
2	0.105968	0.296878
3	0.086522	0.242400
4	0.074931	0.209924
6	0.061181	0.171402
8	0.052984	0.148439
12	0.043261	0.121200
24	0.030590	0.085701

Inverted: to detect a $\log_2$FC of 0.30 requires $r \geq 2$ per arm; 0.20 requires 5; 0.15 requires 8; 0.10 requires 18; 0.05 requires 71.

Simulated rejection rates (400 simulated studies per cell, permutation null with 400 permutations, $\alpha = 0.05$):

d (log2FC)	r = 3	r = 6	r = 12	
0.00	0.040	0.062	0.043	<- null; nominal alpha recovered
0.10	0.217	0.360	0.708	
0.20	0.685	0.902	0.995	
0.30	0.950	0.998	1.000	
0.50	1.000	1.000	1.000	

At 6 replicates per arm the design has about 90 per cent power against a $\log_2$ fold-change of 0.20 and about 36 per cent against 0.10. That is the honest statement of what the experiment can and cannot see, and it belongs in the preregistration rather than in a discussion section afterwards.

Per-locus testing is far weaker and must be reported separately. With Bonferroni correction across 12 addressed loci tested individually against 10 matched controls each ($\alpha = 0.05/12 = 0.004167$, z sum $= 3.706881$):

```
r = 3    per-locus SE = 0.299722    MDE = 1.111034
r = 6                   0.211936          0.785620
r = 12                  0.149861          0.555517
r = 24                  0.105968          0.392810
```

The aggregate DiD can detect a 0.20 $\log_2$FC while the per-locus test cannot detect anything under 0.55 at the same replicate count. Both numbers are published, because reporting only the aggregate would let a reader believe individual loci had been resolved.

Layer 2, same discipline, different instrument. The model is the storage law with the exposure schedule as input [Curzi, 2026a]:

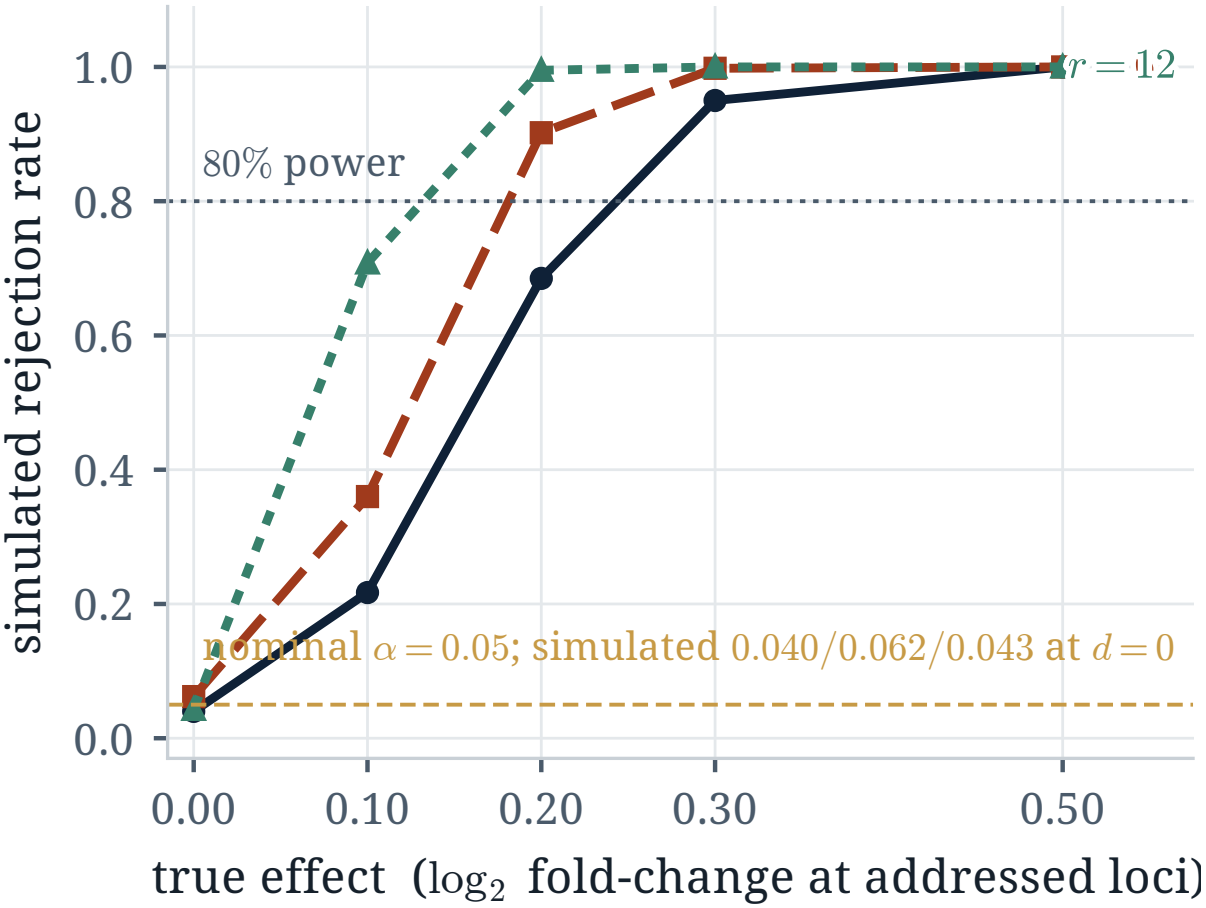


Figure 36.4: **The design is well powered against a 0.20 fold-change and underpowered against 0.10:** simulated rejection rate against true $\log_2$ fold-change at three replicate counts, 400 simulated studies per point with a 400-permutation null at $\alpha = 0.05$. **The leftmost column is the calibration check: at zero effect the test rejects at 0.040, 0.062 and 0.043, recovering nominal alpha to within simulation error.** Modelled, not measured: the noise parameter $\sigma = 0.35$ is assumed and every value scales with it. Generated by `figs/src/CCP-036_fig04.py`.

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,      S_t = ln(1 + (N−1) U_t)
```

Y is the preregistered self-report or behavioural proxy for Q; C is **measured, never estimated**. **Where Y is an ordinal rating, the multiplicative form is not licensed and the relation is demoted in print to a monotone ordering** (Section 3.6) [Stevens, 1946, Curzi, 2026t]. Two coefficients are fitted.

Design simulation, 200 synthetic runs per condition, $T = 60$ periods, first 40 for training and the last 20 held out, $\delta_{\text{true}} = 0.10$, $\eta_{\text{true}} = 0.40$, measurement noise sd 0.20. These are synthetic design characteristics, not data.

	median held-out R^2					
generator	model	persistence	direct	AR(1)	shuffled	passes all four
b = 1.0 (full)	0.3280	−0.0806	−4.9116	−0.0069	−0.1666	88.5%
b = 0.5 (half)	−0.1646	−0.5147	−0.3560	−0.2368	−0.5393	51.0%
b = 0.0 (null)	−0.5881	−1.0805	−0.3708	−0.4239	−0.5930	9.5%

Parameter recovery at $b = 1.0$: $\hat{\delta}$ median 0.1150 against a true 0.10; $\hat{\eta}$ median 0.4335 against a true 0.40. **At $b = 0.0$ the fit still returns $\hat{\eta} = 0.1290$ — a fitted coefficient on an input that does nothing, which is exactly why the shuffled control and not the coefficient's sign is the decision rule.**

The conjunction gate's false-positive rate is 9.5 per cent at $T = 60$, and this must be stated in the preregistration, because a reader who assumes 5 per cent will overread a single passing participant. Lengthening the series tightens it:

T	train	b = 0.0	b = 0.5	b = 1.0
60	40	0.095	0.482	0.895
90	60	0.052	0.590	0.940
120	80	0.040	0.603	0.950
180	120	0.013	0.630	0.978
260	180	0.010	0.662	0.993

The important column is the middle one, and it is bad news honestly reported: a half-size effect never reaches 80 per cent power under this gate at any practical series length, saturating near 0.66. The design cannot fix this by running longer; it can only fix it by replicating across people.

36.5.10 5.10 Cross-individual replication, which rescues the underpowered case

Treat each participant's pass or fail on the four-baseline conjunction as a Bernoulli trial. **Under the null the per-person pass rate is 0.095; against a half-size effect it is 0.482.** A one-sided binomial test on the count of passing participants:

n participants	reject if k >=	actual alpha	power vs b = 0.5
10	4	0.01069	0.797
12	4	0.02160	0.908
16	5	0.01381	0.948
20	5	0.03566	0.991
24	6	0.02204	0.995
30	7	0.02004	0.999

Twelve participants, with four or more passing the four-baseline conjunction, gives an actual false-positive rate of 0.0216 and 90.8 per cent power against a half-size effect that a single participant's series could only detect 48 per cent of the time. That is the argument for cross-individual replication stated as arithmetic: replication across people converts a stringent per-person gate from underpowered into adequate, and it does so without weakening the gate.

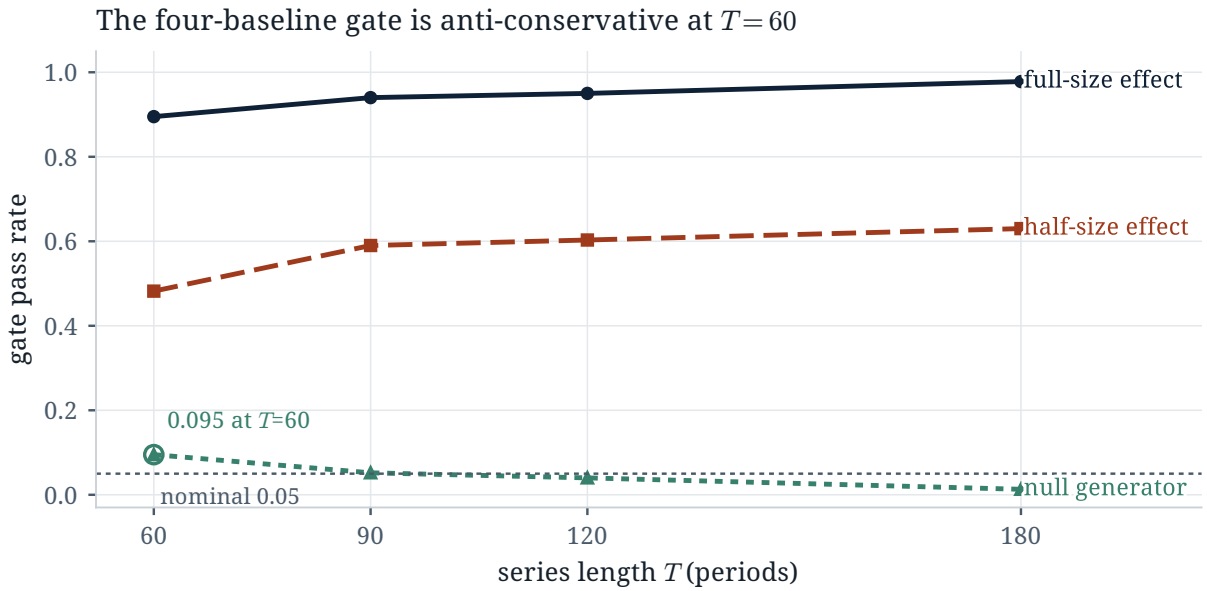


Figure 36.5: **The conjunction gate is anti-conservative at the series length the design would otherwise have used, and a half-size effect never reaches adequate power:** the four-baseline gate's pass rate against series length, for a full-size effect, a half-size effect and a null generator, from 200 synthetic runs per condition. **At $T = 60$ the null generator passes 9.5 per cent of the time against a nominal 5**, and the half-size curve saturates near 0.66 rather than approaching 0.80. Modelled, not measured: synthetic series with $\delta_{\text{true}} = 0.10$, $\eta_{\text{true}} = 0.40$ and measurement noise sd 0.20. Generated by `figs/src/CCP-036_fig05.py`.

36.5.11 5.11 What would settle the open problem, in either direction

A positive settlement would require, at minimum: (1) a DiD at addressed loci significantly different from zero against the scrambled control, with the shuffled control failing to reproduce it; (2) **address dependence** — swapping the address list with the payload held constant moves the effect to the new loci, **without which the result is not addressing but something else**; (3) **payload dependence** — holding the address list constant and changing the payload changes the result, **without which the address is acting as a bare tone**; (4) replication in an independent laboratory, from the published packet, with no contact; and (5) **a candidate mechanism that survives the magnitude arithmetic of Section 5.2** — a route by which the delivered acoustic field produces local strains or forces large enough to matter, **and** a route by which those forces are *sequence-selective*.

Item 5 is the hard one and it is where the burden actually sits. Items 1 to 4 could all be satisfied by an unrecognised artefact; a mechanism that clears the magnitude arithmetic could not.

A negative settlement is easier and is a real result: F-1 plus F-2 plus a demonstration that the design had adequate power against the effect sizes of interest — which Section 5.9 provides, at 90 per cent against 0.20 $\log_2$FC and 6 replicates. **A well-powered null on a well-specified question is worth more than a hundred underpowered positives** [Open Science Collaboration, 2015, Ioannidis, 2005].

36.5.12 5.12 The concordance for this section

BRIDGE — a staged, reversible procedure with an explicit abort at every stage. *Evidence:* pre-registration, blinding, matched controls and held-out prediction are established methodology

with a measurable effect on the reliability of published findings; the Open Science Collaboration found 97 of 100 original studies reporting significant results against roughly 36 per cent of replications succeeding, with effect sizes about halved [Open Science Collaboration, 2015]. **That is evidence about method, and it is the reason this section is shaped as it is.** *Precedent:* the alchemical operations were staged, ordered and named, with explicit conditions for stopping and returning to a prior stage, and the laboratory notebook tradition descends directly from them. **This records that staged, abortable, documented procedure was arrived at independently and preserved for centuries. It is not evidence that any alchemical claim was correct, and most were not.** *Integration:* the bridge holds at the level of *procedural form* and nowhere else. C1 **passes**; C2 **passes** — it predicts that traditions with staged procedures will record failure conditions, which is checkable; C3 **passes. Status: established as a bridge about procedure. It carries nothing about mechanism, and it is worth noting that the tradition's staged rigour coexisted with almost uniformly false theory — which is itself the argument for keeping the columns apart.**

A SPECULATIVE BRIDGE, marked speculative in the sentence that states it. It is speculative, and offered here only as a speculation, that the "sympathetic" correspondences of the older corpus — the notion that like acts on like at a distance through a shared quality rather than through contact — are a pre-instrumental description of resonant coupling, in which a driven oscillator responds selectively to a frequency matching its own. Resonance is real physics; the correspondence is a guess about what the old texts were describing; **and the speculation supplies no mechanism for sequence-selectivity and does nothing whatever to close the gap.** It is recorded so that the thought is on the page with its badge on rather than smuggled in as an implication.

36.5.13 5.13 Requirements

REQ-5.1 The open problem appears verbatim in every artefact: preregistration, consent packet, method section, abstract, any public description.

REQ-5.2 The words "treatment", "therapy", "therapeutic", "cure", "heal", "dose-response" in a clinical sense, and "patient" **do not appear** in any document describing this protocol. Participants are participants; sessions are exposures; **nothing is administered.**

REQ-5.3 Consent steps P0 and P1 precede all others, and a run that omits them is not a run of this protocol.

REQ-5.4 Declination rate is reported with every result, in the same table.

REQ-5.5 The scrambled-corpus arm is **mandatory** and is the comparator for the primary endpoint. **A study without it does not test the hypothesis.**

REQ-5.6 Acoustic field mapped **at the sample, not at the source** — pressure amplitude, spectrum and spatial uniformity at the cell layer or listening position, published.

REQ-5.7 Temperature logged continuously in Layer 1, published per arm, with the between-arm difference stated. **Acoustic heating is the leading artefact and an unlogged temperature is an unanalysable experiment.**

REQ-5.8 Control loci matched on baseline expression, accessibility class and chromosome, with the matching procedure published and run **before** unblinding.

REQ-5.9 Analysts blind to arm through primary endpoint computation.

REQ-5.10 C is measured, never estimated. A model fitted without a measured cost term is reported as a different and weaker model.

REQ-5.11 Freeze before held-out: timestamped, hashed, public.

REQ-5.12 All four baselines reported, all five numbers, whatever they say, **with the suite's own measured false-rejection rates alongside** [Curzi, 2026u].

REQ-5.13 The permutation test is calibrated, not assumed. Simulate the design at zero effect, measure the actual rejection rate, and report it. **Section 5.9's 9.5 per cent figure is what calibration looks like when it is done.**

REQ-5.14 Power is stated in advance against named effect sizes, **with the assumed variance flagged as assumed until pilot replicates replace it.**

REQ-5.15 Address-swap and payload-swap arms are required for any positive claim. **A positive result without them is reported as unattributed.**

REQ-5.16 Nulls are published at equal prominence.

REQ-5.17 Rescue clauses are closed in advance and listed in the preregistration.

REQ-5.18 Dosimetry computed and published per session.

REQ-5.19 No Layer 2 result is cited in support of a Layer 1 claim, in any document, ever.

REQ-5.20 Independent replication from the published packet is solicited, **and the packet is complete enough to permit it without contacting the authors** [Kerckhoffs, 1883].

36.5.14 5.14 Failure modes

FM-5.1 — THE MISSING SCRAMBLED ARM. *Signature:* Arms A and D only — encoded versus silence, large clean effect. *Why:* Arm B is expensive and its whole purpose is to make the result smaller. *Fix:* **encoded-versus-silence measures the effect of sound; it cannot measure the effect of sequence, which is the only question being asked.**

FM-5.2 — THE THERMAL ARTEFACT. *Signature:* effect present in Arms A, B and C equally, absent in D; correlates with exposure duration and transducer power. *Fix:* **the DiD against Arm B subtracts every response common to sound; a design that compares against silence subtracts none of them.**

FM-5.3 — THE GLOBAL LIFT. *Signature:* addressed loci up, and control loci up by the same amount, reported as positive because the addressed loci were the ones anyone looked at. *Fix:* **the DiD is the endpoint precisely so this cannot be reported as specificity.**

FM-5.4 — THE UNATTRIBUTED POSITIVE. *Signature:* a real A-versus-B difference, no address-swap arm, and a discussion section that attributes it to addressing. *Fix:* **without the swap arms the result says "something about this signal differed from its scramble" and no more, which is worth reporting exactly that way.**

FM-5.5 — LAYER LEAKAGE. *Signature:* a paper whose Layer 1 result is null and whose abstract leads with Layer 2 self-report. *Fix:* the firewall. **This is the single most likely route by which this architecture becomes an overclaim, and it usually happens in the abstract rather than in the analysis.**

FM-5.6 — THE ASSUMED VARIANCE. *Signature:* a power calculation quoting an MDE to four decimals from a variance nobody measured. *Fix:* **the tables in Section 5.9 are computed from $\sigma = 0.35$, which is assumed; replace it with pilot data before filing, and if the true σ is twice the assumption every MDE doubles.**

FM-5.7 — THE UNCALIBRATED PERMUTATION. *Signature:* a p-value of 0.03 from a permutation test, with no statement of the permutation family and no calibration run. *Fix:* **permutation tests are exact only for the exchangeability structure they assume, and the structure here — loci within arm — is a choice that must be defended** [Curzi, 2026u].

FM-5.8 — OPTIONAL STOPPING. *Signature:* r grew during the study and the final r is not the registered one. *Fix:* **looking and then adding replicates inflates the false-positive rate without limit, and no correction applied afterwards recovers it** [Simmons et al., 2011].

FM-5.9 — THERAPEUTIC DRIFT. *Signature:* a participant says it helped; the next consent form says "some participants report benefit"; six months later the word "protocol" has acquired a clinical inflection nobody decided on. *Fix:* REQ-5.2, applied to documents, in writing, at every revision. **This failure has no single moment at which it happens, which is why it needs a rule at the level of vocabulary rather than a judgement at the level of claims.**

FM-5.10 — VULNERABLE ENROLMENT. *Signature:* a participant who hopes the procedure will treat a condition, enrolled anyway because they were told it would not and said they understood. *Fix:* C7. **Understanding a disclaimer and believing it are different states, and the protocol declines enrolment on the second rather than accepting the first as sufficient.**

FM-5.11 — THE CONDITIONAL ROOM. *Signature:* Layer 2 exposure in a space participants must occupy, or with an off switch they must ask to use. *Fix:* C5. **Reversibility is a property of the circumstance, not only of the equipment.**

FM-5.12 — THE RESCUED NULL. *Signature:* a null, followed by a note that the corpus was unsuitable, the cell line atypical, the tuning route wrong, or the participants unusual. *Fix:* the closed rescue clause, registered before collection. **Every one of those statements may even be true; none of them may be made after the fact by the people who ran the study.**

FM-5.13 — PRECEDENT AS EVIDENCE. *Signature:* the staging bridge cited as support for the protocol's plausibility rather than for its form. *Fix:* **the precedent column establishes that the procedural form was findable by more than one route. It establishes nothing about whether this particular procedure will find anything.**

36.6 6 Composition: four terms, two chains, one undefined product

36.6.1 6.1 The term defined at first use — and the phrase it is not

This section uses the phrase "destiny made manifest". It does not mean, and must not be read as meaning, "Manifest Destiny". That distinction is drawn before anything else, because the second phrase carries specific historical freight that would poison every sentence that followed it.

"Manifest Destiny" is a nineteenth-century American political slogan. It entered circulation through O'Sullivan's 1845 essay "Annexation" in the *United States Magazine and Democratic Review* (**named without a key**), arguing for the annexation of Texas, and became the standing justification for continental expansion. **It was used, repeatedly and explicitly, to justify the dispossession, forced removal and killing of Indigenous peoples, on the argument that the expansion was ordained and therefore that those in its path had no standing to refuse it.** That is what the phrase did in the world. **It is not a neutral term with an unfortunate history; the history *is* the term's meaning in ordinary use.**

Everything about that construction is the opposite of what this section describes. The slogan was collective, imposed by one group upon another, asserted rather than demonstrated, and specifically constructed so that the people it was applied to could not decline. **This section's subject is personal, applies only to the person it happens to, is bounded by what is already present in that person, and — by Section 4.4 — cannot issue a directive at all. No continuity with the nineteenth-century slogan is intended, implied, or available.**

36.6.2 6.2 What the term does mean here

```
DESTINY MADE MANIFEST

A latent capability, already present in the system, brought into expression.

LATENT       the capability exists but is not currently expressed
MANIFEST     the capability is expressed
THE CHANGE   is in expression, not in what exists
```

Three clauses are load-bearing, **and each is a constraint rather than a promise.** *"Already present"*: nothing is added, and **nothing that is not already latent can be reached** [Curzi, 2026y]; Section 6.5 gives that in arithmetic. *"Brought into expression"*: the operation is regulatory — **no sequence is changed, no information is imported, nothing is edited.** *"A latent capability"*: singular, specific, and in principle nameable. **If the capability cannot be named before the attempt, the attempt has no endpoint and Section 5.7's falsification criteria cannot be written.**

36.6.3 6.3 The four terms

TERM 1	LATENT CAPABILITY	the archive is already there, and bounds the result
TERM 2	ADDRESSING WITHOUT EDITING	a locus can be named without being cut
TERM 3	BELIEF WITH A REAL CHANNEL	belief is an input to a system that must have gain, and at least one such channel is measured and survives disclosure
TERM 4	NARRATIVE AS REGULATORY INPUT	the input can be structured, disclosed, and constructed so that it cannot command

THE COMPOSITION CLAIM, badged as the author's thesis: these four are not four analogies for one idea. They are four distinct results that occupy four distinct positions in a single signal chain, and the chain is the subject of this section. And immediately, in the same breath: **one position in that chain is occupied by the open problem, the chain is multiplicative, and a multiplicative chain with one unmeasured link has an unmeasured product.**

36.6.4 6.4 Composition is multiplicative — the chain law

DERIVATION (elementary). For a cascade with gains $g_1 \dots g_n$ the end-to-end gain is the product $G = \prod_i g_i$. **This is a shared mathematical form** (Section 3.2): the same product governs an optical cascade and a causal chain, **and no physical result transfers.**

(a) A chain of individually strong links is weak. For four links of equal gain:

```
g = 0.95  ->  g^4 = 0.814506
g = 0.90  ->  g^4 = 0.656100
g = 0.80  ->  g^4 = 0.409600
g = 0.70  ->  g^4 = 0.240100
g = 0.50  ->  g^4 = 0.062500
g = 0.30  ->  g^4 = 0.008100
```

Four links that each pass 70 per cent of what reaches them deliver 24 per cent end to end.

(b) One unmeasured link makes the product unmeasured; one zero link makes it zero.

```
G = g_1 · g_2 · g_3 · x        with x unknown  ->  G unknown
                               with x = 0      ->  G = 0 exactly
```

No amount of strength in the measured links compensates for a link whose gain has never been measured. The product is not "probably fine"; it is undefined. These gain figures are arithmetic on hypothetical values. No gain in the biological chain has been measured, and the table above is a demonstration of how products behave, not an estimate of anything.

36.6.5 6.5 The latency bound, in arithmetic

The constraint that nothing not already latent can be reached has an exact form, **which converts a caution into a computable ceiling.**

DERIVATION. $0 \le F(u) \le \ln N$: **$\ln N$ is the hard ceiling on surplus per interaction**, and N is the number of distinguishable categories *already present* [Curzi, 2026]. And $Q_t \le \max\{Q_0, \eta \ln N/\delta\}$: stored order is bounded by a quantity depending on N, η and δ and on nothing else [Curzi, 2026a].

```
N ->  ln N     ceiling on Q with eta = 0.35, delta = 0.08
 2    0.693147          3.032519
 4    1.386294          6.065038
```

```
  8    2.079442          9.097557
 12    2.484907         10.871467
 20    2.995732         13.106329
 50    3.912023         17.115101
200    5.298317         23.180138
```

Raising the ceiling requires raising N — acquiring categories that were not there. Everything this section describes operates on u, the arrangement of what is already present, and u is bounded above by 1, so F is bounded above by $\ln N$. This is the variety bound appearing again [Ashby, 1956]: only variety absorbs variety, and the variety available is the variety already held. **The same ceiling, at more than one scope, from the same theorem. It was not chosen; it follows.**

36.6.6 6.6 Term one — latent capability that is already present

Evidence, and it is the strongest single result underneath this section. Takahashi and Yamanaka (2006, mouse; 2007, human; **named without a key**) showed that a small set of regulatory transcription factors reverts a differentiated somatic cell to a pluripotent state.

What makes it the keystone is precisely what it did not do. No new genetic information was introduced into the cell's repertoire. A skin cell already contained everything required to be a pluripotent cell; **what changed was which parts of that repertoire were being read.** The capability was latent, present and unexpressed; four regulatory signals changed the expression state. **That is "destiny made manifest" as a laboratory result, and it is the reason the phrase is defensible at all in a paper with these honesty clauses.**

The archive that makes latency possible. Approximately 98 per cent of the human genome is non-coding and largely regulatory. Against a genome of about 3.1 billion base pairs:

```
non-coding                    ~98%   ->  ~3.038e9 bp
transposable elements         ~45%   ->  ~1.395e9 bp
endogenous retroviral sequence  ~8%  ->  ~2.48e8 bp
coding-ish                      ~2%  ->  ~6.2e7 bp
```

And the single most striking documented instance: syncytin, indispensable to placenta formation, derives from a retroviral envelope gene — ancient foreign text, inherited, later read out as vital function.

Precedent. The contemplative traditions consistently framed development as *uncovering* rather than *acquiring*. **This records that the latency framing was arrived at independently and preserved. It is not evidence for the biology, and the biology is not a vindication of it** [Curzi, 2026w]. **What the independent arrival does establish — genuinely, and it is worth having — is that the distinction between acquiring a capability and expressing one already held is findable without instruments.**

THE CONSTRAINT, which is the term's real contribution. The reprogramming result is a demonstration of what latency permits *and* of what it forbids. **A skin cell can be reprogrammed to pluripotency because pluripotency is in its repertoire. It cannot be reprogrammed into something outside that repertoire, and no regulatory signal reaches what is not there.**

The arithmetic of how tight the constraint is runs the other way from intuition. Take the crudest accounting of the addressable space — each of roughly 20,000 protein-coding genes independently on or off:

```
addressable states       2^20000
log2                     20,000 bits
ln N                     13,862.9436 nats
log10                    6,020.6 (a number with 6,021 digits)
```

Against that, the number of distinct human cell types classically named is a few hundred:

```
ln(200) = 5.298317 nats     realized / addressable = 3.821928e-04
ln(400) = 5.991465 nats     realized / addressable = 4.321928e-04
```

The viable subset is a vanishing fraction — under one twentieth of one per cent — of the addressable one. The addressable space is astronomically large, which is why "latent" is not a small word. The viable subset is a vanishing fraction of it, which is why "bounded by what is already there" is not a formality. Most of the addressable space is not a capability; it is a dead cell.

36.6.7 6.7 Term two — addressing without editing

Evidence. The catalytically dead nuclease with a transcriptional activator (Qi et al., 2013; Gilbert et al., 2014; **named without keys**) is the established proof that **targeting and editing are separable functions**: bind a locus, activate transcription, cut nothing. Native self-processing of a crRNA array (Zetsche et al., 2015; Fonfara et al., 2016) is the established proof that a **multicast destination list** is a real biological format.

The packet correspondence, stated exactly and no further. A packet carries a destination header and a payload; routing reads the header; the payload is not altered by being routed. The guide-RNA spacer occupies the header position in the same sense. **The correspondence is about format. It is not a claim that cells run a network stack, and the analogy earns nothing beyond letting a reader who knows one system read the other** (Section 3.2).

What this term contributes: a way for an input to be specific without being destructive. Without it, any influence on expression would have to be either global or invasive. **With it, specificity and non-destructiveness are compatible in principle, and that compatibility is what makes the rest of the architecture conceivable at all. And the honesty clause, in place: nothing is delivered here, therefore nothing is edited. The term contributes the *concept* that targeting and editing separate. It does not contribute a mechanism by which an acoustic signal targets anything.**

36.6.8 6.8 Term three — belief coupled to a real channel

THE POSTULATE, badged as the framework's own: *belief creates reality if the belief is grounded in truth* [Curzi, 2026ab]. **The conditional clause is the entire content**, and "grounded in truth" means operationally that the channel exists and is real:

```
BELIEF  ->  [ CHANNEL ]  ->  WORLD

"grounded in truth"  =  THE CHANNEL EXISTS AND IS REAL

no channel  ->  no effect, regardless of the strength or sincerity of the belief
```

This is the same architecture as the storage law: input, conversion efficiency, output. A belief with no channel is $\eta = 0$ [Curzi, 2026a].

Evidence — the complete channel. Levine, Gordon and Fields (1978; **named without a key**): placebo analgesia is mediated by endogenous opioids and is naloxone-reversible. **Belief to expectation to neurochemistry to measured analgesia, verified end to end, with a pharmacological blockade that reverses it. Evidence — disclosure does not abolish it.** Kaptchuk et al. (2010) and subsequent trials: open-label placebo produces benefit; effect sizes are modest and the literature is active. **Evidence — the mechanism named independently.** Merton [1948]: self-fulfilling prophecy — **and the crucial detail, that the classic bank-run example works because banks really are fractional-reserve. The belief propagates because it is coupled to a true feature of the system** [Curzi, 2026ac].

Evidence — and now the failures, which are required. Belief does not move objects. It does not alter another person's body at a distance. It does not change physical constants. It does not cure disease by itself. Growth-mindset effects are very small in large replications (Sisk et al., 2018; named without a key) — a correlation near $r = 0.10$ and interventions near $d = 0.08$, which is 96.809 per cent distribution overlap. Expectancy effects are real, modest, and have a contested methodological history [Curzi, 2026ac, Rosenthal and Jacobson, 1968, Jussim and Harber, 2005, McCambridge et al., 2014].

A postulate that names where it fails is a postulate that can be trusted where it holds, and the catalogue of failures is what makes term three usable rather than decorative. Every failure in that list has the same diagnosis: no channel.

Precedent. The traditions that treated conviction as efficacious recorded the observation and preserved it, generally without the conditional. **This is precedent. It is not evidence, and specifically it is not evidence for the unconditional version, which the failure catalogue refutes. The traditions' failure to state the conditional is where they were wrong, and saying so is part of reporting the record accurately.**

36.6.9 6.9 Term four — narrative as regulatory input

Section 4 supplies this term and it contributes three things with **different statuses**.

(a) A construction. A narrative can be encoded, addressed and rendered as fully audible tone, losslessly and reversibly: 192 characters, 1,776 symbols, 35.520 seconds, 8.918863 octaves, entirely within the audible band. **Status: demonstrated as encoding. It says nothing about any receiver.**

(b) A constraint. The no-subliminals rule. **A message hidden from its recipient cannot be consented to, and in a model where the subject's own latent capability is the mechanism, concealment does not strengthen the protocol, it removes a component of it. Status: ethical constraint, prior to design.**

(c) A structural result. A corpus closed under negation has resultant exactly zero; a mutually orthogonal corpus has mean resultant $1/\sqrt{N}$, falling to 0.288675 at $N = 12$ and 0.204124 at $N = 24$. **A field of mutually exclusive directions cannot issue a command; it can only be resolved by selection, and the selection belongs to the receiver.** And the exact identity:

```
surplus carried by a maximally contradictory corpus, N = 12   ln 12 = 2.484907 nats
information required to select one of its members             ln 12 = 2.484907 nats
contributed by the corpus                                     0
contributed by the receiver                                   all of it
```

Status: derived, in the extremal case, under a uniform receiver prior (clause H-A). Term four is therefore the term that makes the architecture non-coercive by construction rather than by promise.

36.6.10 6.10 How the four compose — two chains, one closed and one open

Assemble the four terms and two distinct chains appear. They share their first and last links and differ in the middle.

```
CHAIN A -- ESTABLISHED END TO END

  [4] narrative, disclosed, contradictory, chosen from
        |
        v  (reading, deciding, believing -- the receiver's selection)
  [3] belief with a real channel
        |
```

```
        v  expectation -> endogenous opioid / autonomic / behavioural
    [ ] measured outcome

    STATUS: every link measured. Gain modest. Survives disclosure.
            This chain is real today and it is not nothing.


CHAIN B -- ONE LINK OPEN

    [4] narrative encoded, addressed, rendered as tone          BUILT
          |
          v  acoustic field at the cell
    [?] SEQUENCE-SPECIFIC DEMODULATION BY THE CELL              NOT ESTABLISHED
          |
          v
    [2] addressed, non-editing regulation                       established as a
          |                                                     principle, not as a
          v                                                     consequence of sound
    [1] latent capability expressed                             established mechanism
                                                                (by other means)

    STATUS: links 4, 2 and 1 are each established. Link [?] has never been measured.
            By Section 6.4, THE PRODUCT IS UNDEFINED, and if the missing gain is zero
            the product is exactly zero.
```

The two chains must never be conflated, and Section 5.4 makes the separation a structural requirement of the protocol rather than a caution in a discussion section.

Now the observation that makes the honest reading the strong one. Chain A is not a consolation prize. A person reads a disclosed, deliberately contradictory corpus. Nothing in it tells them what to do; the corpus's resultant is zero and **the entire** $\ln N$ **of selection is contributed by them.** They select. **That selection is a belief, held by a person, about their own conduct — and belief about one's own conduct is coupled to the world through the most thoroughly documented channels the postulate has: expectation, autonomic response, endogenous opioid signalling, and, above all, what the person then does.**

Chain A's channel is real, measured, and does not require deception. It is also modest, and the failure catalogue is the honest boundary on it [Curzi, 2026ac]. **But "modest, real, measured, disclosed and chosen" is a considerably stronger thing to be able to say than "large, claimed, unmeasured and concealed."**

36.6.11 6.11 The link-by-link ledger

Link	What it is	Status	If it fails
A1 Corpus construction	contradictory, balanced, published	**ARITHMETIC**, by inspection of the text	corpus becomes a command; FM-4.4, FM-4.7
A2 Disclosure	full packet, in advance, unconditional	**CONSTRAINT** — held or not held	consent invalid; protocol not run
A3 Selection by receiver	$\ln N$ nats supplied by the person	**ARITHMETIC** in the extremal case	if the corpus is unbalanced, the selection is partly the operator's

Link	What it is	Status	If it fails
A4 Belief to channel	expectation to physiology and conduct	**CITED (measured), modest**	the failure catalogue: no channel, no effect
A5 Outcome measured	preregistered, held-out, four baselines	**METHOD** — Section 5	unmeasured claim; FM-5.9
B1 Encoding to tone	lossless, audible, reversible	**BUILT**	round-trip fails; REQ-4.7
B2 Acoustic field at the cell	mapped, calibrated, published	**MEASURABLE**	unanalysable exposure; REQ-5.6
B3 Sequence-specific demodulation	**the open problem**	**NOT ESTABLISHED**	**the chain's product is zero**
B4 Addressed non-editing regulation	the dead-nuclease principle	**CITED as principle**	no consequence follows without B3
B5 Latent capability expressed	regulatory reprogramming	**CITED, by other means**	bounded by Section 6.5 regardless

Read down the Status column and the architecture's real shape is visible: eight of ten links are established or are constraints under the operator's control, and one link — B3 — is unmeasured. The section's whole content is that this is what the situation is, stated in a form that can be checked rather than in a form that can be believed. And **B5 is bounded by Section 6.5 regardless of anything upstream: even a fully closed Chain B would be capped at** $\ln N$**. The ceiling does not move when the channel opens.**

36.6.12 6.12 What the bound buys

A bounded claim is a checkable claim. "This can express what is already latent" names a candidate set in advance, which means the attempt has an endpoint, which means Section 5 can write a falsification criterion for it. **"This can bring about whatever you envision" names no set, has no endpoint, and cannot fail — which is precisely why it cannot be tested and precisely why nothing built on it accumulates** [Popper, 1959, Meehl, 1967].

A bounded claim also has a specific safety property. If nothing outside the latent set can be reached, the worst outcome of a failed attempt is that nothing happens. **An architecture whose failure mode is "no effect" is categorically different from one whose failure mode is "an unintended effect", and the difference is structural: the architecture produces sound and nothing else.**

And the bound is what makes the claim survive contact with the honest catalogue. Every item on the list of where belief fails — moving objects, acting on another body at a distance, changing constants, curing disease alone — **is outside the latent set of the person holding the belief. The bound predicts the catalogue** [Curzi, 2026ac]. **A model that predicts its own known failure cases is doing work.**

36.6.13 6.13 Chosen by the person it happens to

The argument, in three steps. *Step 1:* the corpus is deliberately contradictory and therefore has zero or vanishing resultant — $\|\sum\| = 0.000000000000$ at N = 4, 8, 12, 24. *Step 2:* therefore no direction is issued by the signal; whatever direction obtains was contributed by the receiver's selection, and its information content is $\ln N$ nats — for $N = 12$, 2.484907 nats or 3.584963 bits. *Step 3:* the mechanism downstream of the selection is the receiver's own latent capability, bounded by $\ln N$ of their own categories and expressed rather than added.

At no point in that sequence does anything originate outside the person. The corpus supplies a span; the person supplies the direction; the person's own archive supplies the material;

the person's own ceiling supplies the bound. **That is why the nineteenth-century slogan is not merely a different meaning of the word but a structural opposite: it asserted a direction over people who had not selected it, whereas this architecture cannot supply a direction at all.**

And the reversibility clause completes it. The exposure is sound; it stops when the sound stops; nothing is administered or introduced. **A person can decline before, stop during and withdraw after, and the declination is recorded as an outcome of the protocol rather than as a failure of it.**

36.6.14 6.14 Why the honest form is the more powerful one

The overclaiming form would say: a narrative encoded as sound reaches the cell, activates latent capability, and makes a person's destiny manifest. It would be more exciting for about one page. Then: **it names no endpoint, so it cannot be tested; it survives no contact with the failure catalogue; it requires the reader to take the one unmeasured link on faith, which means the whole chain rests on faith by Section 6.4; and it collapses the moment anyone computes the strain figure of Section 5.2. It is also, by its own logic, a claim about something done *to* a person, which puts it in the same structural family as the slogan Section 6.1 spent its length distinguishing this work from.**

```
ESTABLISHED AND AVAILABLE NOW (Chain A)

  A person constructs or receives a corpus that cannot command them,
  reads all of it, selects from it -- contributing ln N nats of direction
  that the corpus did not contain --
  holds the resulting belief,
  and that belief acts through channels that are measured, disclosed,
  naloxone-reversible in at least one documented case, and modest in size.

  Everything in that sentence is established. None of it requires deception.
  All of it is bounded by what the person already holds.


NOT ESTABLISHED (Chain B, link B3)

  That the same corpus, rendered as tone, is demodulated by a cell.
```

And the specific reason the honest form is stronger, stated as an argument rather than as a preference. A claim's usefulness is the product of its size and its probability of being true — **which is Section 6.4's chain law applied to epistemics. A large claim with an unmeasured link has an undefined product. A modest claim with every link measured has a defined one. The second is worth more, not as a matter of taste, but because you can build the next thing on top of it.**

36.6.15 6.15 Requirements

REQ-6.1 Define the term at first use in any document, in full, with the distinction from the nineteenth-century slogan. **Not a footnote and not a later clarification.**

REQ-6.2 Prefer "destiny made manifest" or "manifestation" in running text. **The two-word slogan form is not used.**

REQ-6.3 Name the capability before the attempt. **A latent capability that cannot be named in advance has no endpoint and no falsification criterion, and no attempt proceeds on it.**

REQ-6.4 State the bound with the claim. **"Bounded by what is already present" is part of the claim, not a caveat on it.**

REQ-6.5 Identify which chain you are in. **No document mixes them.**

REQ-6.6 State link B3's status verbatim wherever Chain B is mentioned.

REQ-6.7 Compute the composition, do not assert it. **A chain containing an unmeasured gain is reported as having an undefined product, never as "promising".**

REQ-6.8 Report the failure catalogue with the postulate, in the same passage.

REQ-6.9 The corpus must be able to say no to itself. **A corpus that has been "balanced" but whose pairs differ in length, level, repetition or position is reported with its measured asymmetry.**

REQ-6.10 Declination is an outcome. Recorded and reported.

REQ-6.11 No conferral language. **The words "confers", "grants", "unlocks", "installs", "upgrades" and "activates in you" are not used. The architecture expresses what is present; it does not add.**

REQ-6.12 No clinical vocabulary, in any document in this line of work.

REQ-6.13 Precedent labelled as precedent. **Every precedent citation states what it records and states that recording is not measuring.**

REQ-6.14 Speculative bridges marked in the sentence that states them. **Not in a preamble, not in a closing note.**

REQ-6.15 Reversibility stated as an architectural property, with its basis.

36.6.16 6.16 Failure modes

FM-6.1 — THE SLOGAN COLLISION. *Signature:* a talk, abstract or heading uses the two-word slogan, or uses "destiny made manifest" without the distinction. *Why:* the two-word form is shorter and more familiar, **which is exactly the problem.** *Fix:* **the distinction is drawn at first use, every time, in full. There is no audience for whom it can be assumed, and the cost of drawing it is two sentences.**

FM-6.2 — THE UNBOUNDED PROMISE. *Signature:* "whatever you can envision"; no candidate set named; no endpoint, so no attempt can fail. *Fix:* $F \leq \ln N$ stated with the claim. **An unfalsifiable promise is not a large claim, it is an absent one.**

FM-6.3 — THE UNDEFINED PRODUCT. *Signature:* four links, three measured and one not, presented as a chain that "works", with the word "promising" doing the load-bearing. *Fix:* **an unmeasured link does not deliver 70 per cent, it delivers an unknown.**

FM-6.4 — CHAIN CONFLATION. *Signature:* a participant reports benefit from Chain A, and the next document describes it as evidence that the signal reached a cell. *Fix:* the firewall. **This is the single most likely route by which the architecture becomes an overclaim, and it typically happens in an abstract or a summary rather than in an analysis.**

FM-6.5 — THE CONFERRAL VERB. *Signature:* "activates capabilities in the listener"; a verb whose subject is the protocol and whose object is the person. *Fix:* **grammar carries the claim. "The protocol confers X" asserts a mechanism the architecture does not have; "a capability already present may come to be expressed" asserts what the model actually contains.**

FM-6.6 — THE SELECTIVE PRECEDENT. *Signature:* the precedent column cited where the traditions agree and omitted where they were wrong. *Fix:* **report the record, not the useful half of it. A concordance that only ever finds agreement has stopped being a method and become a search** [Curzi, 2026].

FM-6.7 — PRECEDENT DOING EVIDENCE'S JOB. *Signature:* "the traditions said development is uncovering; the biology confirms it", implying two verifications where there is one. *Fix:* **what independent arrival establishes — that the framing was findable by more than one route — is real and is smaller than confirmation.**

FM-6.8 — THE CEILING IGNORED. *Signature:* a programme working on u and reporting gains as though the ceiling could be exceeded by working harder on arrangement. *Fix:* $u \leq 1$, **so** $F \leq \ln N$. **At** $\eta = 0.35$ **and** $\delta = 0.08$**, moving from** $N = 12$ **to** $N = 20$ **raises the ceiling from 10.871467 to 13.106329; no amount of work on** u **does that.**

FM-6.9 — THE PREORDINATION SLIDE. *Signature:* "destiny" begins to carry the sense of a fixed outcome awaiting discovery; selection quietly disappears from the description. *Fix:* **preordination and this model are opposite structures, not similar ones: preordination fixes the outcome independent of selection, and selection is this model's only source of direction.**

FM-6.10 — THE MODEST CLAIM INFLATED IN TRANSIT. *Signature:* the paper says "modest, measured, disclosed"; the summary says "measured"; the press line says "proven". *Fix:* **every restatement drops a qualifier, so the qualifiers must be inside the claim's main clause rather than trailing it.**

36.7 7 Results

36.7.1 7.1 R1 — determination by choice, proved

A corpus closed under negation has resultant exactly zero: $\|\sum_i v_i\| = 0.000000000000$ at N = 4, 8, 12 and 24, for any pairing. **A mutually orthogonal corpus's mean direction shrinks as** $1/\sqrt{N}$ — analytic values 0.500000, 0.353553, 0.288675 and 0.204124 at those sizes, against simulated 0.497007, 0.352275, 0.288022 and 0.203292 from 2,000 draws of N isotropic unit vectors in $\mathbb{R}^{64}$ per row, and 0.140454 and 0.099395 at N = 50 and 100. **A field of mutually exclusive directions cannot issue a command: the corpus does not point, it spans. The vector representation of a directive is a modelling choice, badged at clause H-B**, and what the result rigorously establishes is that **a set constructed so that its elements cancel has no resultant, and a receiver who acts must therefore have selected.**

36.7.2 7.2 R2 — a maximally contradictory corpus is a maximum-surplus corpus

At $N = 12$, $F = \ln 12 = 2.484907$ **nats** $= 3.584963$ **bits — the hard ceiling** $\ln N$, **which cannot be exceeded — and this is exactly the information required to select one of the twelve members. The corpus supplies as much surplus as the selection consumes and supplies zero net direction; the direction is contributed by the receiver, and its magnitude is exactly the surplus the signal carries. The identity is claimed only in the extremal case**: at $u < 1$ the surplus is $\ln(1 + (N-1)u) < \ln N$ and under a non-uniform prior the selection cost is that prior's entropy, **and both fall together, which preserves the structure of the argument and not the exact equality** (clause H-A).

36.7.3 7.3 R3 — the bulk-compression route, closed

$s = p/K$ **with** $K = 2.2 \times 10^9$ **Pa requires** 2.2×10^7 **Pa — 240.83 dB SPL — to produce the** 10^{-2} **strain associated with mechanosensitive gating, against an atmospheric ceiling of 194.09 dB SPL set by the ambient pressure itself. At that ceiling the achievable bulk strain is** 4.6057×10^{-5}, **short by a factor of 217.12. The route is closed by more than seven orders of magnitude and stays closed at the physical ceiling of sound in air. The closure is of the bulk route only**: radiation force, microstreaming, localised shear, cavitation and dimensional resonance are non-linear and localised and are **not** closed — **but none has been shown to be sequence-selective, and closing the bulk route narrows the search without answering it.**

36.7.4 7.4 R4 — the averaging bound

A receiver whose integration window exceeds the 35.520-second transmission sees composition and not order, and composition carries $\log_2 \binom{1779}{3} = 29.803$ **bits of the 1,776 sent — 1.68 per cent, none of it sequence. Sequence is the entire hypothesis, so an integrator that slow is not a candidate receiver at all.** The companion constraint runs the other way and is stated so that

neither is over-read: **resolution is not the obstacle** — an integration window of 0.08037675 ms suffices to separate the four-tone alphabet, so **it would be dishonest to invoke the Gabor limit as a reason the hypothesis fails. The two requirements pull in opposite directions across roughly three orders of magnitude in time, and any proposed mechanism must say how it spans that gap.**

36.7.5 7.5 R5 — the preregistration prices its own nulls

The four-baseline conjunction gate's measured false-positive rate is 9.5 per cent at $T = 60$**, not the nominal 5**, falling to 0.052 at $T = 90$, 0.040 at 120, 0.013 at 180 and 0.010 at 260. **And a half-size effect never reaches 80 per cent power under this gate at any practical series length, saturating near 0.66. Cross-individual replication is the repair, and it is arithmetic: with twelve participants and a rejection rule of four or more passing, the actual false-positive rate is 0.0216 and the power against a half-size effect is 90.8 per cent — against 48.2 per cent for a single participant's series. Replication across people converts a stringent per-person gate from underpowered into adequate without weakening the gate.**

36.7.6 7.6 R6 — the no-subliminals constraint as first law

A message hidden from its recipient cannot be consented to, and the argument depends on no empirical claim. Consent requires knowing what one is consenting to; a signal engineered so the listener cannot perceive it has been engineered so the listener cannot evaluate, compare, refuse or report it, **and every one of those capacities is the capacity to decline. In an architecture whose mechanism routes through the recipient's own selection, concealment does not strengthen the protocol — it removes a component of it, so the design is incoherent before it is unethical.** The constraint's operative test is a reconstruction test: **can the recipient rebuild the signal from the packet?** And it is not contingent on an empirical bet — **the open-label result means the boundary is cheap to hold, not that holding it was conditional on that.**

36.7.7 7.7 CR — the counter-result

The paper closes its own proposed mechanism, leaves its own composition undefined, and publishes its own gate's measured failure rate. The bulk-compression route requires **240.83 dB against a 194.09 dB atmospheric ceiling**. The composition's fourth term — sequence-specific demodulation — **has never been measured, so by the chain law the product is undefined, and if the missing gain is zero the product is exactly zero.** And the decision gate the protocol depends on **rejects at 9.5 per cent against a nominal 5**, which is published in the registration rather than discovered afterwards. **Eight of ten links in the architecture's ledger are established or are constraints under the operator's control, and one is unmeasured; the paper's whole content is that this is what the situation is, in a form that can be checked rather than in a form that can be believed.**

36.7.8 7.8 Hypotheses, resolved

	Statement	Outcome
H1	a contradictory corpus cannot issue a directive	**confirmed** in the vector model; $\lVert\sum\rVert = 0.000000000000$, simulated means tracking $1/\sqrt{N}$

	Statement	Outcome
H2	corpus surplus equals selection information	**confirmed in the extremal case only**, and clause H-A states where it fails
H3	the bulk-compression route is closed	**confirmed by arithmetic**: 240.83 dB required, 194.09 dB available
H4	a slow receiver cannot recover sequence	**confirmed by arithmetic**: 29.803 of 1,776 bits, 1.68 per cent
H5	the conjunction gate misses its nominal rate	**confirmed, and the number is 0.095** at $T = 60$

Every confirmation above is a confirmation of arithmetic or of a simulation, not of a measurement on any physical or biological system. H1's confirmation is a property of a vector model of directives and not of any text; H5's is a property of a synthetic series generated at parameters this paper chose. **Section 10 records both.**

36.7.9 7.9 Model comparison: what `RIGOR_STANDARD` §6 requires

This paper fits no model to any collected data. Two simulations are run — the resultant simulation of Section 4.4 and the design simulations of Section 5.9 — and both characterise a procedure rather than fit a model to observations. The row-by-row status is printed rather than omitted.

Requirement	Status here	Why
out-of-sample performance	**MET in form for the Layer 2 design simulation; NOT APPLICABLE elsewhere**	the simulation fits on 40 periods and tests on 20 held out, which is the design; **no real data are used**
matched null with equal free parameters	**MET for both simulations**	the $b = 0.0$ generator is the matched null — the same two-parameter model with a null input — and it passes the gate 9.5 per cent of the time; the Layer 1 calibration run is the same design at zero effect
information criterion (AIC or BIC)	**NOT MET, and not computed anywhere in this corpus**	no likelihood is reported here, and **no AIC or BIC value appears in any paper of this corpus. The omission is corpus-wide and is recorded rather than papered over**
explicit free-parameter count	**MET**	the Layer 2 model fits **two** parameters, $(1 - \delta)$ and η, the same two under every generator; C is **measured, never estimated**; the Layer 1 endpoint fits **none** — the DiD is a contrast, not a fit

Requirement	Status here	Why
effect sizes with intervals	**PARTIALLY MET**	the design simulations report minimum detectable effects at stated power, which is the interval-equivalent for a design; **no empirical effect size is reported anywhere, because none was measured**
calibration of the test	**MET, and this is the paper's own contribution to the row**	the permutation test is simulated at zero effect and returns 0.040, 0.062 and 0.043 at r = 3, 6, 12; the conjunction gate is simulated and returns **0.095 against a nominal 0.05, which is published rather than assumed**
identification of the reflexive parameter	**NOT MET for Layer 2's loop; MET BY DESIGN for the Layer 1 contrast**	[Manski, 1993]; the Layer 1 arm contrast is identified by randomisation and a matched control, the Layer 2 loop is not identified at all
power-law procedure [Clauset et al., 2009]	**NOT RUN, deliberately**	the three exponents here are analytic properties of known functions; Section 3.7
conjoint measurement test [Luce and Tukey, 1964]	**NOT RUN**	Section 3.6; the exposed construct is demoted instead

A specification cannot be beaten on fit and earns nothing on it. What is claimed here is a construction, three exact arithmetic consequences, and a calibrated procedure — **and the procedure's value is untested, because it has not been run.**

36.8 8 Discussion

36.8.1 8.1 What follows

Disclosure is specifiable, not merely commendable. The operative test is reconstruction: a disclosure packet is adequate if and only if the recipient can rebuild the signal from it. **That converts an ethical sentiment into a checkable property of a document**, and it is the same move Kerckhoffs made for ciphers [Kerckhoffs, 1883] and the corpus makes for channel architecture, where non-disclosure is refused as a design premise rather than merely discouraged [Curzi, 2026].

Non-coercion can be structural rather than promised. A corpus closed under negation has resultant exactly zero; **disclosure makes a command refusable, and contradiction makes it not a command**, and the two fail in different circumstances — **disclosure fails when the listener is tired, deferential or trusting, whereas a field of mutually exclusive directions has no resultant whether or not anyone is paying attention.**

A negative computation is progress. The bulk-compression closure removes one candidate mechanism from the search at the cost of an afternoon's arithmetic. **A closed route is worth more than an open speculation, and the framework's willingness to close its own is the strongest single thing this paper does.**

And a decision gate must be calibrated before it is used. The four-baseline conjunction rejects at 9.5 per cent against a nominal 5 at the series length the design would otherwise

have used. **Measuring that costs an hour of simulation and it changes what a single passing participant means** — which is the same discipline the corpus applies to its own falsification suite, where the shuffle control's measured rate is 0.108 against a nominal 0.05 [Curzi, 2026u].

36.8.2 8.2 What does not follow

Nothing here shows that anything receives the signal. The encoding is elegant and the elegance is not evidence. **There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal** [Curzi, 2026y], and the paper builds a transmitter for a receiver not demonstrated to exist.

Nothing here supports the essence law. k is chosen, the square law is a postulate, and the Gabor arithmetic applied to it is exact arithmetic applied to a chosen object [Curzi, 2026, Markowsky, 1992]. **A reader must not take the precision of Section 4.10 as transferring to the postulate it operates on.**

Nothing here is a clinical, medical or therapeutic claim, in either direction.

Nothing here has been run. The protocol has an unknown pass rate, the corpus has never been rendered for any participant, and the preregistration is a specimen. **A preregistration is a commitment to find out, and the paper's contribution is that the commitment can be written precisely enough to be kept — not that keeping it produced anything.**

And the honest expectation is negative. The magnitude arithmetic gives the bulk route a decisive *no*, and every remaining route is non-linear, localised and not known to be sequence-selective. **A reader who finishes this paper believing the experiment is likely to return a positive has read it wrong. The procedure is specified in this detail precisely so that a negative result will be worth something — and on the present state of the evidence, a negative is what the arithmetic leads one to expect.**

36.8.3 8.3 Relation to the surrounding literature

The positive results sit inside established literatures and add to none of them: the time–frequency bound is Gabor's [Gabor, 1946]; the selection cost is Hartley's and Shannon's [Hartley, 1928, Shannon, 1948]; the variety bound is Ashby's [Ashby, 1956, 1958]; the disclosure principle is Kerckhoffs's [Kerckhoffs, 1883, Brandeis, 1914]; the methodological apparatus is the reform literature's [Open Science Collaboration, 2015, Nosek et al., 2018, Simmons et al., 2011, Gelman and Loken, 2013]. **What is new is the composition, the arithmetic, and the discipline of publishing the closures.**

The methodological position is closer to the critics of physics-in-social-science than to its enthusiasts. Anderson [1972] is the standing argument that a formalism valid at one scale carries no automatic warrant at another; Cartwright [1999] that laws hold in the circumstances that make their premises true and not elsewhere; Mirowski [1989] that borrowing physical formalism can import the mathematics and lose what gave it content. **The three-way label of Section 3.2 is the answer, and it is procedural rather than rhetorical: state which kind of borrowing is being made, and let the reader hold the paper to it.**

36.8.4 8.4 Parallel results elsewhere in the corpus

Four results in other series are structurally parallel and are named because the parallel is the point.

A gate that never rejects is not a gate. Section 4.14 rejects one bridge of three in print. The corpus's audit of a traditional operation set against a parameter model returns the same shape of distribution — **seven MAPS, two SPLIT, one RENDERING and three REJECTED across thirteen entries, with five entries carrying a failure** [Curzi, 2026]. **That is what a discriminating instrument looks like when it is pointed at material its author is sympathetic to.**

A structural limit is not closed by care inside the channel. Section 4.1's constraint is that no amount of care within a concealed delivery architecture restores the recipient's ability to decline. The corpus reaches the identical conclusion for identity channels: **in-band architectures admit an impersonation class that no care within the shared channel can close, and out-of-band separation closes it structurally — and channel separation is *strengthened* by disclosure, because its effectiveness is entirely independent of whether anyone knows it is in place, so non-disclosure is refused as a design premise rather than merely discouraged** [Curzi, 2026]. **This paper's no-subliminals constraint is that argument applied to content.**

A representation makes a control possible that no amount of care in the wrong representation could supply. Section 5.6's scrambled arm works because the encoding fixes what "the same symbols in a different order" means. The corpus's identifier work is the same move: **typed-field parsing collapses the withdrawn-code parsing-ambiguity false-positive share from 88.99 per cent to 0.00 per cent, against a naive detection precision of 11.01 per cent — a control right one time in nine** [Curzi, 2026]. **Getting the representation right is what makes the control exact.**

A ceiling on constructed inputs is not a promise. Section 6.5's $\ln N$ and Section 4.10's Gabor margin are both upper bounds on what a construction can do, and neither says any system will reach them. The corpus states the same caution about its own architecture: **the 56.62 per cent widening is a ceiling on constructed inputs and not a promise, visibility is explicitly distinguished from pledgeability, and the valuation of non-financial capital is named as unsolved** [Curzi, 2026]. **This paper's bounds are bounded by that caution in the same way.**

36.8.5 8.5 Parallels in the foundations volume

Three foundational results carry the same structure as this paper's and are named for that reason.

Calibration by tense. Section 6.9's insistence that a directive be stated so that it could be acted on now, rather than deferred, has an exact precedent in the corpus's treatment of spoken commitment: **a future-tense formulation is consistent with an unchanged day, so it generates no dissonance and the consistency mechanism has nothing to act on** [Curzi, 2026]. **A corpus of deferred directives is a corpus that cannot be selected from today**, and the balance check of Section 4.13 would pass it while the protocol got nothing.

Arrival is auditable from the last correction. Section 5's freeze-then-open discipline asks when a fitting process has arrived. The corpus's contraction result supplies the general answer: **a contraction with modulus $q < 1$ on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least q each pass** [Curzi, 2026]. **Arrival is a property of the map, not of the analyst's patience**, which is why the freeze is placed by the design rather than by judgement.

And the honesty classification is itself a corpus-level instrument. The foundations volume closes with a three-drawer ledger — **four items of established mathematics, seven of established psychology, and exactly three declared postulates, with every protocol stated to run on the first two and steer by the third** [Curzi, 2026]. **Section 4.5's and Section 5.5's status tables are that instrument applied at paper scale, and the fact that the counts can be printed is what makes the classification checkable rather than rhetorical.**

36.9 9 Verification

The results compound from nine postulates. Each is atomic, each is checkable on its own, and each carries its status.

36.9.1 9.1 Dimensional audit

[1] denotes a pure number. **Where the framework fixes no dimension, the audit says so rather than assigning one.**

Symbol	Dimension	Unit	Status
v_i (directive)	[1]	none	**defined** as a unit vector; **a modelling choice, and the framework supplies no operational scale on which a directive is measured**
$\lVert\sum v_i\rVert$	[1]	none	**derived**
N	[1]	count of directives	**defined**
u	[1]	none	**defined** dimensionless on $[0, 1]$
$N_{\text{eff}} = 1 + (N-1)u$	[1]	count	**defined** by the framework; **not a Hill number**
$F = \ln(1 + (N-1)u)$	[1]	nats	**derived**; a **Shannon-type** logarithm, **not a thermodynamic entropy**
selection cost $\ln N$	[1]	nats (or bits at $\log_2$)	**derived**; the conversion is $1/\ln 2 = 1.442695$
payload	[1]	bits, then bases	**derived**; 1 bit per base by the declared mapping, so bases and bits are numerically equal here and **the equality is a property of the chosen encoding, not of nucleotides**
address length L	[1]	nucleotides	**chosen**; the expected-collision figure is $6.2 \times 10^9/4^L$, dimensionless
Δt	[T]	seconds	**chosen** symbol duration
Δf	[T^{-1}]	hertz (cycles s^{-1})	**derived** rms spread
$\Delta t\,\Delta f$	[1]	none	**derived**; balances against $1/(4\pi)$, a pure number, **only with f in cycles per second**

Symbol	Dimension	Unit	Status
k	`[T^{-1}]` **by assertion**	hertz	**the value** $(9/8)^2/\varphi^2 = 0.483425732988414$ **is DEFINED and exact — a ratio of exact quantities, which is what licenses fifteen figures — and the attachment of hertz is POSTULATED and is a separate, stronger claim** (P5)
$f(n) = kn^2$	`[T^{-1}]`	hertz	**derived** from k, inheriting its badge exactly
j	`[1]`	octaves	**chosen** integer transposition; 2^j is dimensionless
cents	`[1]`	none	**derived**; $1200 \log_2$ of a frequency ratio
p	`[M L^{-1} T^{-2}]`	pascal	**chosen**/measured
K (bulk modulus)	`[M L^{-1} T^{-2}]`	pascal	**cited** physical constant, 2.2×10^9 Pa for water
$s = p/K$	`[1]`	none	**derived** dimensionless strain; **balances exactly, since pressure over pressure is a pure number**
SPL in dB	`[1]`	dB re 20 μPa	**derived**; a **logarithmic ratio against a fixed reference**, so differences are ratios of pressures and are licensed, **and sums are not**
dose	`[1]`	per cent of an 8-hour allowance	**derived** from level and duration at a 3 dB exchange rate
$\log_2$ transcript abundance	`[1]`	log2 units	**measured** in principle; **an interval scale in the log domain**, which is why the endpoint is a difference of differences
DiD	`[1]`	log2 fold-change	**derived**; differences of logs, licensed on an interval scale

Symbol	Dimension	Unit	Status
σ	[1]	log2 units	**ASSUMED** for the power calculation; **must be replaced by a pilot estimate**
r (replicates)	[1]	count	**chosen**
δ	[1]	per period	**measured** dimensionless fraction per step [Curzi, 2026a]
η	[q] per nat per period	proxy units per nat per period	**measured**; *not* dimensionless [Curzi, 2026a]
Y (Layer 2 endpoint)	**the framework is silent**	rating points, or a behavioural count	**undetermined**, and **ordinal where the instrument is a rating scale**, in which case the multiplicative form is not licensed (Section 3.6)

Does the Gabor relation balance? Δt carries [T] and Δf carries [T^{-1}], so the product is dimensionless and is compared against the pure number $1/(4\pi)$. **It balances, and only with** f **in cycles per second** [Curzi, 2026z].

Does the strain relation balance? $s = p/K$ is pressure over pressure and is exactly dimensionless. **It balances, and the closure of Section 5.2 is therefore a comparison of two pure numbers,** 4.6057×10^{-5} **against** 10^{-2}**, which is the cleanest form a refutation can take.**

Does the storage law balance? $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$ balances only if η carries [q] per nat per period, since S is dimensionless in nats [Curzi, 2026a]. η **is therefore not a pure number, and any report of an** η **without a statement of what [q] counts is half a number.**

Where the framework has no operational unit, it is said here. A directive has no operational unit; v_i is a dimensionless placeholder for an object the framework does not instrument, and clause H-B is the badge on it. **The Layer 2 endpoint's dimension is likewise not fixed by the framework**, and this paper does not supply the determination. **A dimensionless placeholder is honest; an invented unit would not be.**

36.9.2 9.2 The postulates

P1 — For a set closed under negation, $\sum_i v_i = 0$**.** *Status:* **ARITHMETIC.** Each v is cancelled by its $-v$.

P2 — For N **mutually orthogonal unit vectors,** $\|\sum v_i\| = \sqrt{N}$**.** *Status:* **ARITHMETIC**, by the Pythagorean identity; the mean is therefore $1/\sqrt{N}$.

P3 — $F(u) = \ln(1 + (N-1)u)$ **attains** $\ln N$ **at** $u = 1$ **and 0 at** $u = 0$**.** *Status:* **DEFINITION** of the framework's functional [Curzi, 2026]. It carries no empirical weight.

P4 — Selecting one of M **equiprobable alternatives costs** $\ln M$ **nats.** *Status:* **CITED**, elementary [Hartley, 1928, Shannon, 1948].

P5 — The essence law $f(n) = kn^2$ **with** $k = (9/8)^2/\varphi^2 = 0.483425732988414$**, expressed in hertz.** *Status:* **the value is DEFINITION** — an exact ratio of exact quantities, computable to arbitrary precision, which is what licenses fifteen significant figures — **and the attachment of hertz is POSTULATED, with no external support of any kind** [Curzi, 2026, Markowsky, 1992]. **These are two steps, not one, and the second is much the stronger.**

P6 — Two indices share a pitch class iff they have the same odd part. *Status:* **ARITHMETIC** [Curzi, 2026].

P7 — $\Delta t\,\Delta f \geq 1/(4\pi)$ for any square-integrable signal with finite second moments in both domains. *Status:* **CITED** [Gabor, 1946, Curzi, 2026z].

P8 — For a plane acoustic wave in a fluid of bulk modulus K, the longitudinal strain amplitude is $s = p/K$. *Status:* **CITED**, elementary continuum mechanics. $K = 2.2 \times 10^9$ Pa for water is a **cited** physical constant.

P9 — The number of compositions of n indistinguishable items into 4 labelled classes is $\binom{n+3}{3}$. *Status:* **ARITHMETIC**, the standard stars-and-bars count.

What is NOT a postulate. The claim that a cell demodulates sequence-specific information from an acoustic signal is **not** a postulate of this paper. It is the **open problem**, it appears in no line of the chain below, and nothing here is derived from it. **That is the honest entry.**

36.9.3 9.3 The compounding chain

```
P1              -> R1a  a corpus closed under negation has resultant exactly 0;
                        verified at N = 4, 8, 12, 24 as 0.000000000000.      ARITHMETIC

P2              -> R1b  an orthogonal corpus has mean resultant 1/sqrt(N):
                        0.500000, 0.353553, 0.288675, 0.204124 at N = 4, 8,
                        12, 24; simulated 0.497007, 0.352275, 0.288022,
                        0.203292.                                            ARITHMETIC

P3              -> L1   at u = 1 the surplus is F = ln N; at N = 12 that is
                        2.484907 nats = 3.584963 bits.                       DEFINITION

P4              -> L2   selecting one of 12 costs ln 12 = 2.484907 nats.     CITED

L1 + L2         -> R2   the surplus a maximally contradictory corpus carries
                        EQUALS the information required to select one member,
                        so the corpus supplies 0 direction and the receiver
                        supplies all of it.                                  ARITHMETIC
                        (extremal case only; clause H-A)

P5 + P6         -> L3   at n = 1, 3, 5, 7 the four base tones are distinct
                        pitch classes; at j = 8 they are 123.7570, 1113.8129,
                        3093.9247 and 6064.0924 Hz, minimum gap 990.0559 Hz.  ARITHMETIC
                        (given P5)

L3 + P7         -> L4   Gabor floor dt_min = 0.07957747/990.0559
                        = 0.08037675 ms; at dt = 20 ms the margin is 248.83x. ARITHMETIC

L4              -> L5   1,776 symbols at 20 ms run 35.520 s.                 ARITHMETIC

P8              -> R3   s = p/K needs p = 2.2e7 Pa for s = 1e-2, which is
                        240.83 dB SPL; the atmospheric ceiling is 194.09 dB,
                        where s = 4.6057e-05, short by 217.12.               CITED (P8) + ARITHMETIC

P9 + L5         -> R4   a receiver slower than 35.520 s sees only composition:
                        log2 C(1779,3) = 29.803 bits of 1,776, i.e. 1.678%.  ARITHMETIC

P3 + P5         -> L6   a folded pair (m, 2m) with m odd shares a pitch class
                        and is separated by exactly 4.000000 in frequency
                        ratio -- 2400.0 cents, in every pair.                ARITHMETIC
```

```
(simulation)   -> R5   the four-baseline conjunction gate rejects a null
                       generator at 0.095 at T = 60, 0.040 at T = 120; the
                       binomial rule "4 of 12" gives alpha = 0.02160 and
                       power 0.908 against a half-size effect.          SIMULATED DESIGN

(the argument) -> R6   a message hidden from its recipient cannot be
                       consented to.                                    DEFINITION / no number
```

Every line is `ARITHMETIC` except four. P4, P7 and P8 are cited standard results. **R2 consumes P3, a definition of the framework's own functional, and holds only in the extremal case. L3, L4, L5, L6 and R4 all consume P5, whose numerical value is a DEFINITION and whose expression in hertz is POSTULATED** — so **the cents figures and the ratio of 4.000000 stand on arithmetic alone, while every hertz and every millisecond inherits the postulate. A reader who rejects the essence law keeps Section 4.11's intervals and loses Section 4.9's frequencies**; R3 and R1 are untouched either way, because neither uses k. **R5 is a simulated design characteristic and not a measurement. R6 produces no number at all**, which is stated here rather than left for the reader to notice.

What this chain does not establish. It does not establish that a cell receives anything. It does not establish that the essence law is true. It does not establish that the protocol finds anything if run. **And it establishes nothing about any effect on any listener**: R1, R2 and R6 are properties of a construction, R3 and R4 are refutations, and R5 is a property of a synthetic series.

36.9.4 9.4 The verification instance, worked

Four checks, none longer than five lines.

Check one — the identity between surplus and selection.

```
ln 12          = 2.4849066498          nats
log2 12        = 3.5849625007          bits
ln 12 / ln 2   = 2.4849066498/0.6931472 = 3.5849625007   consistent
F(u=1) at N=12 = ln(1 + 11*1) = ln 12                    identical
```

Check two — the defined constant, and the alphabet.

```
(9/8)^2  = 81/64               = 1.265625
phi^2    = phi + 1             = 2.618033988749895
k        = 1.265625/2.618033988749895
                               = 0.483425732988414      EXACT, DEFINED
f(7) - f(5) = k(49 - 25) = 0.483425732988414 x 24
                               = 11.60221759 Hz
x 2^8    = 11.60221759 x 256   = 2970.16770 Hz
but the binding gap is f(3) - f(1) = k x 8 x 256
                               = 990.0559 Hz            the minimum
```

Check three — the closure.

```
strain wanted                     s = 1e-2
pressure needed  p = sK = 1e-2 x 2.2e9   = 2.2e7 Pa
reference        p_0                     = 20e-6 Pa
level            20 log10(2.2e7/20e-6)   = 240.83 dB SPL
atmospheric ceiling
                 20 log10(101325/20e-6)  = 194.09 dB SPL
```

```
strain at the ceiling
                 101325/2.2e9            = 4.6057e-05
shortfall factor 1e-2 / 4.6057e-05       = 217.12
```

Check four — the averaging bound.

```
symbols                                  n = 1776
compositions into 4 classes  C(n+3, 3) = C(1779, 3)
                             = 1779 x 1778 x 1777 / 6
                             = 936,946,289
log2(936,946,289)                        = 29.803 bits
fraction of 1776 bits                    = 29.803/1776 = 1.678%
```

Four checks, none longer than five lines, and between them they establish the identity, the constant, the closure and the bound. A reviewer who runs them has verified every headline this paper contributes that is not quoted from elsewhere.

36.9.5 9.5 Limits, asymptotics, and the domain of validity

$N \to 2$ **in the corpus.** A two-directive corpus closed under negation still has resultant zero, **but the** $1/\sqrt{N}$ **argument has no content and the surplus is** $\ln 2 = 0.693147$ **nats.** REQ-4.19 sets the floor at $N = 4$ for exactly this reason. $N = 1$ gives $F = \ln 1 = 0$ and a resultant equal to the single member: **a corpus of one is a command, and the limit says so.**

$N \to \infty$**.** $F \to \infty$ in the model and the mean resultant $\to 0$, **but the audible band and the render time bind first**: 1,776 symbols already run 35.520 s, and the payload grows linearly in the corpus's character count. **The limit is unreachable for a physical reason rather than a mathematical one, and naming which is the point of taking it.**

$u \to 0$**.** $F \to 0$: a parallel corpus carries no surplus at all, and the selection cost falls to zero with it. **The two fall together, which is clause H-A's content, and it is why the identity survives as a structure while the equality does not.**

$\Delta t \to 0$**.** $\Delta f \to \infty$: **an instantaneous symbol has no pitch**, so the alphabet collapses and the signal becomes a click. $\Delta t \to \infty$ gives arbitrarily fine resolution and an arbitrarily long transmission. **The domain is** $\Delta t > 0$ **with the audibility margin as the binding constraint rather than the mathematics.**

$L \to \infty$ **in the address.** Expected collisions $\to 0$ geometrically and the address becomes exactly unique — **at a cost linear in** L**, which is why 20 is a choice rather than a limit.** $L \to 0$ gives one address matching everything: **not an address at all**, and the table shows the transition happening between 12 and 16.

$p \to p_{\text{atm}}$**.** The strain saturates at 4.6057×10^{-5} because the wave cannot pull below vacuum. **The ceiling is physical, not technological, and it is why the closure is a closure rather than an engineering challenge.**

$\tau \to \infty$ for a receiver's integration window. The recoverable information tends to $\log_2 \binom{n+3}{3}$, which grows only as $3 \log_2 n$: **at** $n = 1776$ **that is 29.803 bits and at** $n = 10^6$ **it is 56.7 — so a longer message does not help a slow receiver, it makes the surviving fraction smaller.**

$T \to \infty$ in the Layer 2 gate. The false-positive rate falls toward 0.010 and the half-size power saturates near 0.66. **The gate's power does not tend to 1, which is the whole reason cross-individual replication is required rather than optional.**

$r \to \infty$ in Layer 1. MDE $\to 0$ as $r^{-1/2}$. $r \to 1$ **is not a limit of the model but a failure of the design**: with one replicate per arm the DiD has no variance estimate at all.

36.9.6 9.6 Invariance, and how far each claim has been tested

Invariance of the resultant under pairing is exact: a set closed under negation sums to zero for *any* pairing, because cancellation does not depend on order. **Tested at four corpus sizes** (4,

8, 12, 24), **all constructed.**

Invariance of the fractional structure under octave transposition is exact: multiplying every frequency by 2^j leaves every ratio, every cents value and every pitch class unchanged [Shepard, 1982]. **Tested at** $j = 6$ **and** $j = 8$ in the worked encoding — **two transpositions**, and the invariance is algebraic rather than empirical.

Invariance of the Gabor product under time-scaling is exact and is inherited [Curzi, 2026z]; **it is not re-tested here.**

Invariance of the pair ratio 4.000000 across pairs is exact: $f(2m)/f(m) = (2m)^2/m^2 = 4$ for every odd m. **Tested at six pairs** in the worked corpus, **and it holds without exception because it is an identity.**

Invariance of the strain relation under the choice of medium does **not** hold and is not claimed: $s = p/K$ depends on K, and a medium with a bulk modulus three orders of magnitude smaller would move the closure. **The closure is stated for aqueous media at** $K = 2.2 \times 10^9$ **Pa and for nothing else**, which is a real restriction on the result and is recorded as such.

No invariance in this paper has been tested on data from any real system, because no data were collected. **Every invariance claim above is algebraic; the numerical checks are checks of the arithmetic.**

36.9.7 9.7 Significant figures and their warrant

- **Exact by definition or by arithmetic, quoted to whatever precision is useful:** $k = 0.483425732988414$; $1/(4\pi) = 0.07957747154594767$; $\ln 12 = 2.484907$; $\log_2 12 = 3.584963$; the frequency table; the cents values; the pair ratio 4.000000; the resultant 0.000000000000; $4^{20} = 1{,}099{,}511{,}627{,}776$; the composition count and its logarithm 29.803.
- **Computed from stated inputs and rounded honestly:** the strain figures, 240.83 dB, 194.09 dB, the shortfall 217.12, the Gabor margin 248.83, the runtime 35.520 s, the dose 0.123 per cent.
- **Simulated design characteristics at assumed parameters, not measurements:** every entry of the power tables, the rejection rates 0.040/0.062/0.043, the gate rates 0.095/0.482/0.895, and the binomial figures 0.02160 and 0.908. **The noise parameter** $\sigma = 0.35$ **is assumed and every minimum detectable effect scales linearly with it.**
- **Simulated with sampling error, quoted to six figures because that is what the run returned:** the mean resultants 0.497007, 0.352275, 0.288022, 0.203292 from 2,000 draws per row. **These carry Monte-Carlo error that is not restated beside them, and Section 10 records that.**
- **Cited physical constants:** $K = 2.2 \times 10^9$ Pa; the reference pressure $20\,\mu$Pa; one atmosphere as 101,325 Pa; $c = 2.99792458 \times 10^{10}$ cm/s; the Casimir pressures.
- **Quoted from published literature, carrying that literature's uncertainty:** the genome size and its compositional percentages, the occupational exposure limit, the replication counts, and the growth-mindset magnitudes. **None is a measurement of this paper.**

36.9.8 9.8 Analogy discipline, restated at the point of verification

Every borrowing carries the label assigned in Section 3.2 and no passage upgrades one. $s = p/K$ is a **derivation** and is used to close a route, not to open one. The Gabor limit applied to a rendered tone is a **derivation** because a finite-energy signal under a smooth envelope satisfies its premises. The chain law is a **shared mathematical form** and transfers no physical result. The packet/header correspondence is a **shared mathematical form** and **is not a claim that cells run a network stack**. The vector model of a directive is a **modelling choice**, badged at clause H-B. And the resonance reading of older correspondence doctrines is an **analogy**, **marked speculative in the sentence that states it**, contributing **no line of the chain in Section 9.3**.

Every logarithm of a count in this paper is Shannon information [Shannon, 1948], in nats or in bits with the conversion stated. **No thermodynamic entropy appears anywhere**, and the word is not used loosely.

36.10 10 Limitations

L-1. Nothing here has been run. No cell was exposed, no participant recruited, no corpus rendered for anyone. **The protocol has an unknown pass rate and the preregistration is a specimen.** The corpus's one comparable attempt to run a preregistered confirmatory test against a real system **failed its own second falsification rule** [Curzi, 2026u], which is the realistic expectation here.

L-2. A substantial number of primary sources are named without citation keys. Church, Gao and Kosuri (2012); Qi et al. (2013); Gilbert et al. (2014); Zetsche et al. (2015); Fonfara et al. (2016); the 2021 Nobel-recognised mechanotransduction work; Tajik et al. (2016); Ibsen et al. (2015); Duque et al. (2022); Takahashi and Yamanaka (2006, 2007); Kaptchuk et al. (2010); Levine, Gordon and Fields (1978); Sisk et al. (2018); and O'Sullivan (1845) **have no entries in this corpus's bibliography, and no key was invented for any of them.** Every one is standard and locatable; **the gap is in the corpus's bibliography and it is recorded here rather than smoothed over.**

L-3. The vector model of a directive is a modelling choice and the theorem is about the model. Real directives are not unit vectors and "REST NOW" is not literally the additive inverse of "WORK NOW". **What survives rigorously is that a set constructed to cancel has no resultant**, and that depends on the corpus genuinely containing mutually exclusive directions — **a property of the written text, checkable by reading it, and not established by any arithmetic in this paper.**

L-4. The surplus/selection identity holds only in the extremal case. At $u < 1$ or under a non-uniform receiver prior the two sides differ, and **the exact equality is claimed only at $u = 1$ with a uniform prior. A reader who assumes a uniform prior over a real person's dispositions has assumed the thing most likely to be false.**

L-5. The essence law is a postulate and every hertz figure inherits it. k is chosen, its expression in hertz has **no external support of any kind**, and the square law is a postulate about essence structure [Curzi, 2026, Markowsky, 1992]. **The Gabor arithmetic is exact and the object it operates on is chosen.** Whether ratio structure of this kind is perceived as the framework supposes is separately contested [Sethares, 2005].

L-6. The closure of Section 5.2 is of the bulk route only, in aqueous media, at $K = 2.2 \times 10^9$ Pa. Radiation force, microstreaming, localised shear, cavitation and dimensional resonance are **not** closed. **A reader who takes the 240.83 dB figure as closing the whole question has over-read it, and the paper says so twice.**

L-7. The 10^{-2} gating strain is an order-of-magnitude figure. It is quoted from the membrane-mechanics literature as a scale rather than as a measured constant, and **the closure's seven-order-of-magnitude margin is what makes the conclusion robust to that imprecision** — but if gating occurred at 10^{-6} the arithmetic would change, and **nothing in this paper establishes that it does not.**

L-8. The noise parameter in every power calculation is assumed. $\sigma = 0.35$ is a stand-in until pilot replicates replace it, **and if the true value is twice the assumption every minimum detectable effect doubles.** Reporting an MDE to six decimals from an assumed variance is exactly the presentational error FM-5.6 names, **and the tables here are guilty of it in form, which is why the assumption is flagged three times.**

L-9. The simulated values carry Monte-Carlo error that is not restated beside them. The mean resultants of Section 4.4 come from 2,000 draws per row and the design characteristics from 200 to 400 runs per cell. **Six-figure precision on a simulated quantity records what the run returned, not what the quantity is**, and no simulation interval is reported.

L-10. The Layer 2 endpoint is ordinal where the instrument is a rating scale, and the multiplicative model is not licensed on it. The relation is demoted in print to a monotone ordering [Stevens, 1946, Curzi, 2026t], **the conjoint-measurement test that would settle whether the attribute is quantitative at all has not been run** [Luce and Tukey, 1964], and **Michell's objection therefore stands unanswered** [Michell, 1997].

L-11. The Layer 2 reflexive parameter is not identified. Randomising the exposure schedule identifies the exposed-versus-unexposed contrast; **it does not identify the loop's gain** [Manski, 1993, Shalizi and Thomas, 2011], and no statement here about the loop is an identified empirical result.

L-12. No AIC or BIC is reported, here or anywhere in this corpus. Section 7.9 prints the row as NOT MET rather than omitting it. Nothing here is fitted to collected data; **the corpus-wide absence is nonetheless a genuine gap and is recorded as one.**

L-13. The balance check fails on two pairs of the worked corpus and the failure is published rather than repaired. The maximum pairwise character asymmetry is **9 characters, a ratio of 2.00**, at one pair. **A corpus with a 2.00 asymmetry in one pair has a thumb on one scale**, and by Section 4.4's arithmetic its resultant is no longer exactly zero — **which means the worked example does not fully satisfy the specification it demonstrates, and that is reported rather than quietly fixed.**

L-14. The composition claim is the author's thesis and the ledger is the author's ledger. That eight of ten links are established or under the operator's control is a classification made by the person who assembled the chain. **An independent classifier applying the same status categories might count differently**, and no such audit has been conducted.

36.11 11 Conclusion

This paper specifies a fully disclosed chain from text to audible tone, proves the structural property that prevents the protocol from overriding its subject, and then closes its own proposed physical mechanism by computation rather than defending it.

Result One, determination by choice. A corpus closed under negation has resultant exactly zero — $\|\sum\| = 0.000000000000$ at N = 4, 8, 12 and 24, for any pairing — and a mutually orthogonal corpus's mean direction shrinks as $1/\sqrt{N}$, analytic 0.500000, 0.353553, 0.288675 and 0.204124 against simulated 0.497007, 0.352275, 0.288022 and 0.203292. **A field of mutually exclusive directions cannot issue a command. The corpus does not point. It spans.**

Result Two, the maximum-surplus corpus. A maximally contradictory corpus is a maximum-surplus corpus: at $N = 12$, $F = \ln 12 = 2.484907$ **nats** $= 3.584963$ **bits, which is the hard ceiling** $\ln N$ **— and exactly the information required to select one of its twelve members. The corpus supplies as much surplus as the selection consumes and supplies zero net direction, so the listener supplies the selection by construction.** The identity is claimed only at $u = 1$ under a uniform prior.

Result Three, the closed route. $s = p/K$ **with** $K = 2.2 \times 10^9$ **Pa requires 240.83 dB SPL to produce a one-per-cent bulk strain, against an atmospheric ceiling of 194.09 dB SPL; at that ceiling the achievable strain is** 4.6057×10^{-5}**, short by a factor of 217.12. The bulk-compression route is closed by more than seven orders of magnitude and stays closed at the physical ceiling of sound in air.** Non-linear and localised routes remain open and **none of them has been shown to be sequence-selective.**

Result Four, the averaging bound. A receiver slower than the 35.520-second transmission recovers at most $\log_2 \binom{1779}{3} = 29.803$ **bits of the 1,776 sent — 1.68 per cent, and none of it is sequence. Resolution is not the obstacle** — 0.08037675 ms suffices to separate the alphabet — **so the two timescale requirements pull in opposite directions across three orders of magnitude, and any proposed mechanism must say how it spans that gap.**

Result Five, the preregistration that prices its own nulls. The four-baseline conjunction gate's measured false-positive rate is 9.5 per cent at $T = 60$**, not the nominal 5**, and a half-size

effect never reaches 80 per cent power under it at any practical series length. **Twelve participants with a rejection rule of four or more passing restores $\alpha = 0.0216$ with 90.8 per cent power against a half-size effect that one participant's series detects 48.2 per cent of the time. Cross-individual replication converts a stringent gate from underpowered into adequate without weakening it.**

Result Six, no subliminals as first law. A message hidden from its recipient cannot be consented to, and the argument depends on no empirical claim. The operative test is reconstruction: **can the recipient rebuild the signal from the packet? In an architecture whose mechanism routes through the recipient's own selection, concealment removes a component rather than adding one — the design is incoherent before it is unethical.**

And the counter-result, which is what the paper is for. It closes its own proposed mechanism at 240.83 dB against a 194.09 dB ceiling. It leaves the fourth term of its own composition unmeasured, and therefore reports the product as *undefined* rather than as *promising* — a chain $G = g_1 g_2 g_3 x$ with x unknown is unknown, and if $x = 0$ it is exactly zero. And it publishes the 9.5 per cent measured false-positive rate of its own decision gate instead of the nominal 5. A reader who finishes this paper believing the proposed experiment is likely to return a positive has read it wrong: the procedure is specified in this detail precisely so that a negative result will be worth something, and on the present state of the evidence a negative is what the arithmetic leads one to expect.

36.11.1 Declarations

Data and code availability

No external dataset was collected for this paper. No cell was exposed, no participant was recruited, no corpus was rendered for any listener, and no result is reported. Values reported as computed — the resultants, the surplus and selection figures, the address-collision figures, the frequency and cents tables, the Gabor margins, the strain and level figures, the composition count and the dosimetry — are exact evaluations of the functions stated beside them, reproducible from the specifications in the text. Values reported as **simulated** — the mean resultants of Section 4.4 (2,000 draws of N isotropic unit vectors in $\mathbb{R}^{64}$ per row) and every design characteristic in Section 5.9 (400 simulated studies per cell with a 400-permutation null for Layer 1; 200 synthetic runs per condition at T periods for Layer 2) — are outputs of the stated procedures at the stated parameters and **are design characteristics, not data.** The noise parameter $\sigma = 0.35$ is **assumed** and is flagged as assumed wherever it is used. Physical constants, genome statistics, occupational exposure limits and every biological result named in Sections 4 to 6 are **quoted from the published literature named in the text**, carry that literature's uncertainty rather than this paper's, and **many of those sources have no key in this corpus's bibliography and are named in full**, as disclosed at Section 3.10 and Limitation L-2. **The corpus and balance check of Appendix A, the preregistration of Appendix B and the parameters of Appendix C are constructed and are labelled constructed wherever they appear.** The figures are generated by the scripts named in each caption, and those scripts contain no data beyond the numbers printed in this paper and the functions stated in Sections 4 to 6.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest. **The architecture specified here is the author's own, the postulate it composes is the author's own, and the ledger classifying its links is the author's own. That is a structural conflict which no declaration removes**; it is the reason failure mode FM-6.6 (the selective precedent) is in the specification, the reason Section 4.14

rejects one bridge of three in print, the reason Section 5.2 closes a route the framework proposed, and the reason Limitation L-14 records the ledger as the author's classification rather than an independent audit.

Author contributions

Single author. Michael Laurence Curzi specified the chain, derived the results, performed the arithmetic and the simulations, produced the figures, and wrote the paper.

Ethics

No human-subjects research was conducted, no personal data was processed, and no biological material was handled. No clinical, medical or therapeutic claim is made anywhere in this paper, in either direction; the published trials named in Sections 4.2 and 6.8 are cited as evidence about channels and their gains and never as guidance about care. The Layer 1 and Layer 2 procedures of Section 5 and the preregistration of Appendix B are **designs for prospective work and are not reports of studies performed**. **Were either run it would require ethical review, and nothing here substitutes for that.** The consent architecture of Section 5.3 — informed disclosure sufficient for reconstruction, free and recorded declination, withdrawal at any point with data destruction on request, reversibility as a property of the signal, no conditional exposure, computed dosimetry, and declined enrolment where a participant persists in expecting a therapeutic benefit — **is stated as a protocol requirement rather than as an assurance, and a run that omits it has not run the protocol.**

Preregistration

This paper reports no preregistered test. The five hypotheses of Section 2 were fixed before the arithmetic and the simulations reported in Sections 4 to 6 were carried out and are printed in the form in which they were fixed; **all five are confirmed by arithmetic or by simulation, and none is confirmed by a measurement on any physical or biological system.** Appendix B is a **specimen preregistration for a study that has not been run**, filed here to demonstrate that the open problem can be posed as a powered, falsifiable, calibrated experiment with closed rescue clauses and a null publishable at equal prominence. **Its decision gate's measured false-positive rate — 9.5 per cent against a nominal 5 — is published in the registration rather than discovered afterwards**, which is what Section 5.13's calibration requirement produces. The corpus's own falsification suite has been audited against itself and is **anti-conservative**: iid permutation rejecting at 0.108, circular shift at 0.120 and block permutation at 0.060 against a nominal 0.05 across 500 null trials [Curzi, 2026u]. **Any applied claim resting on this paper inherits that result.**

36.12 Appendix A · A twelve-directive corpus carried from text to rendered tone

Every number is computed and every disclosure element is produced. This demonstrates the encoding and nothing about any effect on any listener or any cell, and is not offered as doing so. The corpus is constructed for the arithmetic.

36.12.1 A.1 The corpus

Six contradictory pairs, twelve directives, $N = 12$.

#	n	directive	chars
1	1	REST NOW	8
2	2	WORK NOW	8
3	3	HOLD WHAT YOU HAVE	18
4	6	LET IT GO	9
5	5	SPEAK	5
6	10	BE SILENT	9
7	7	GO TOWARD THE UNFAMILIAR	24
8	14	STAY WHERE YOU ARE KNOWN	24
9	9	TRUST THE MEASUREMENT	21
10	18	TRUST WHAT YOU FEEL	19
11	11	BUILD SLOWLY AND FINISH	23
12	22	START AGAIN FROM NOTHING	24
			total 192

The balance check, run before anything else. Pairs 1/2 and 7/8 balance exactly by character count (8/8, 24/24). **Pair 3/4 is 18 against 9 and pair 5/6 is 5 against 9 — these fail the equal-length test and are recorded as failing.** The honest options are to rewrite for balance or publish the asymmetry; **this specimen publishes it. Maximum pairwise character asymmetry is 9 characters (pair 3/4, 18 against 9), a ratio of 2.00.** Pair 9/10 is 21 against 19 (1.105); pair 11/12 is 23 against 24 (1.043). **A corpus with a 2.00 asymmetry in one pair has a thumb on one scale, and the reader is told so rather than not told.** Level, repetition count and symbol duration are equal across all twelve by construction.

36.12.2 A.2 Stages 1 to 3

Stage 1 — data. 192 characters at 8 bits per character = **1,536 bits.**

Stage 2 — sequence. At one bit per base: **1,536 payload bases.** At the two-bit ceiling it would be 768; **the extra 768 bases are the constraint-satisfaction budget, giving 2^{1536} sequences encoding this identical corpus.** The specific sequence used is published.

Stage 3 — addresses. Twelve addresses, 20 nt each: **240 address bases.** Expected spurious matches per address in a 3.1 Gbp genome across both strands: **0.005639. Total transmission: 1,776 bases.**

36.12.3 A.3 Stage 4 — tone

Route B, badged: $k = 0.483425732988414$ Hz is chosen; $f = kn^2$ is a postulate; **the outputs are essence frequencies, not physical resonances.** Band 20 Hz to 20 kHz. **Global transposition $j = 6$, factor 64, applied to all twelve values.**

n	directive	f (Hz)	f x 2^6 (Hz)	pitch class (odd part)
1	REST NOW	0.4834257	30.93925	1
2	WORK NOW	1.9337029	123.75699	1
3	HOLD WHAT YOU HAVE	4.3508316	278.45322	3
6	LET IT GO	17.4033264	1113.81289	3
5	SPEAK	12.0856433	773.48117	5
10	BE SILENT	48.3425733	3093.92469	5
7	GO TOWARD THE UNFAMILIAR	23.6878609	1516.02310	7
14	STAY WHERE YOU ARE KNOWN	94.7514437	6064.09239	7
9	TRUST THE MEASUREMENT	39.1574844	2506.07900	9
18	TRUST WHAT YOU FEEL	156.6299375	10024.31600	9
11	BUILD SLOWLY AND FINISH	58.4945137	3743.64888	11
22	START AGAIN FROM NOTHING	233.9780548	14974.59551	11

Every pair is separated by exactly 4.000000 in frequency ratio — 2400.0 cents, two octaves, without exception. Lowest tone 30.93925 Hz, highest 14974.59551 Hz, span **8.918863 octaves**, all inside the audible band; **six distinct pitch classes for twelve tones, matching** $\lceil N/2 \rceil = 6$ **exactly.** Folded into one octave above the lowest tone:

```
odd part  1  ->    0.0000 cents        odd part  7  ->  737.6518 cents
odd part  3  ->  203.9100 cents        odd part  5  ->  772.6274 cents
odd part  9  ->  407.8200 cents        odd part 11  -> 1102.6359 cents
```

k **cancels from every one of these.** Class 3 is the Pythagorean whole tone $9/8$ above class 1; class 9 is the ditone $81/64$ above it [Barbour, 1951].

36.12.4 A.4 The base alphabet, the rate and the dose

Four-tone alphabet at n = 1, 3, 5, 7 with $j = 8$: **123.7570, 1113.8129, 3093.9247, 6064.0924 Hz** — four distinct pitch classes, minimum gap **990.0559 Hz**, Gabor floor $0.07957747/990.0559 = 0.08037675$ ms. Symbol duration chosen: **20 ms, a margin of 248.83**, satisfying the 100-fold requirement. Total runtime $1776 \times 0.020 = 35.520$ s.

Dosimetry. At 85 dBA, $p_{\text{rms}} = 0.355656$ Pa. **One rendering is 0.123 per cent of the 8-hour allowance; 100 repetitions is 59.2 minutes, 12.3 per cent of it. Published, not estimated.**

36.12.5 A.5 The determination-by-choice ledger

```
corpus size                    N = 12
non-parallelism (design)       u = 1 (mutually exclusive pairs)
effective count                N_eff = 1 + 11(1) = 12   (the corpus's affine count,
                                                         NOT a Hill number)
corpus surplus                 F = ln 12 = 2.484907 nats = 3.584963 bits
corpus resultant               0 exactly (closed under negation)
selection information required ln 12 = 2.484907 nats = 3.584963 bits
contributed by the corpus      0
contributed by the receiver    2.484907 nats -- all of it
```

The signal is 35.520 seconds long, contains 1,776 symbols, spans 8.918863 octaves, is fully audible throughout, is fully published in advance, and does not tell the listener what to do — and the last of those is not a limitation of the design, it is the design.

36.12.6 A.6 The disclosure packet, as produced

1. The twelve directives above, in full.
2. Character encoding: 8-bit, one byte per character, no terminator, directives kept separate.
3. The 1,536-base payload string as generated, in full.
4. Twelve 20-nt addresses, listed, with the 0.005639 expected-collision figure.
5. Tuning: Route B, essence law, badged as postulate; $k = 0.483425732988414$ Hz chosen; band 20 Hz to 20 kHz; $j = 6$ for the corpus chord, $j = 8$ for the base alphabet; the frequency table.
6. Symbol duration 20 ms; Gabor margin 248.83; runtime 35.520 s.
7. Level 85 dBA; dose 0.123 per cent of the 8-hour allowance per rendering.
8. Rendering order, published, **declared not to be a ranking.**
9. Balance check **with the failure recorded**: maximum pairwise character asymmetry 9 characters, ratio 2.00, at pair 3/4.
10. Round-trip verification: text to bits to bases to tones to bases to bits to text reproduces all 192 characters.

11. Off-switch location, and the statement that any person may operate it without asking.
12. **The open problem, in full: there is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal. No clinical, medical or therapeutic claim is made. No claim that DNA is edited or altered. No claim that any capability has been conferred.**

Establishes. That a contradictory corpus can be encoded losslessly into a fully audible, fully disclosed 35.520-second signal whose interval structure is isomorphic to the corpus's contradiction structure, every parameter computed and published, at one-eighth of one per cent of an occupational daily dose.

Does not establish. That anything hears it other than a person. That any cell receives it. That any transcript changes. That any capability is conferred. That any physiological, clinical or therapeutic effect occurs. **The open problem is exactly as open at the end of this specimen as at the beginning.**

36.13 Appendix B · A specimen preregistration, with a calibration run

This is a specimen preregistration and a design simulation. No data were collected and no result is reported.

§1 TITLE. *Does an acoustic signal encoding addressed, sequence-formatted narrative data produce locus-specific transcriptional differences relative to a spectrally matched scrambled control? A preregistered in-vitro test.*

§2 STATEMENT OF THE OPEN PROBLEM. There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal. Mechanotransduction is established; force-to-transcription is established; sonogenetic channel activation is established; DNA data storage is established; separability of targeting and editing is established. **The channel for sequence-specific content is the question. This study tests whether such a channel exists and is written so that the answer "no" is fully reportable.**

§3 WHAT IS NOT CLAIMED. No clinical, medical or therapeutic claim. No claim that DNA is edited, altered or damaged — **the architecture produces sound and cannot do otherwise.** No claim that any capability is conferred. **No human participants in this study.**

§4 SIGNAL. Appendix A's twelve-directive corpus: 192 characters, 1,536 bits, 1,536 payload bases at one bit per base, twelve 20-nt addresses (240 address bases), **1,776 bases total.** Route B tuning, badged as postulate. Four-tone alphabet at n = 1, 3, 5, 7, $j = 8$: 123.7570, 1113.8129, 3093.9247, 6064.0924 Hz. Minimum gap 990.0559 Hz. Symbol duration 20 ms, Gabor margin **248.83** over the 0.08037675 ms floor. Runtime **35.520 s.** Full disclosure packet published with this registration.

§5 ARMS.

```
A  ENCODED       the 1,776-symbol addressed transmission
B  SCRAMBLED     the same 1,776 symbols, permuted; identical duration, spectrum,
                 level and symbol rate; carries no message
C  TONE-MATCHED  steady four-tone chord, same long-term spectrum and level, 35.520 s
D  SILENCE       35.520 s, identical handling
```

§6 EXPOSURE. $r = 6$ independent replicate cultures per arm, 24 in total, randomised. Calibrated transducer, field mapped at the cell layer. Target pressure amplitude at the cell layer 0.3557 Pa, corresponding to 85 dBA in air — **and it is registered in advance that at this pressure the bulk longitudinal strain in aqueous medium is $p/K = 1.62 \times 10^{-10}$, which is more than seven orders of magnitude below membrane strains associated with mechanosensitive gating.** Temperature logged continuously and published per arm. Exposure repeated 100 times per

culture over one session: 59.2 minutes, **12.3 per cent of the 8-hour allowance were a person present, which none is.**

§7 ENDPOINT. Difference-in-differences of mean $\log_2$ transcript abundance, addressed loci ($n_{\mathrm{addr}} = 12$) versus matched control loci ($n_{\mathrm{ctrl}} = 120$), Arm A versus Arm B. Matching on baseline expression, accessibility class and chromosome, performed and published **before unblinding.** Analysts blind to arm.

§8 TEST. Permutation null over locus labels within arm, 2,000 permutations, two-sided $\alpha = 0.05$. **The permutation family is declared: labels are permuted within arm and across the full 132-locus set; calibration at zero effect is reported below.**

§9 POWER. Assumed $\sigma = 0.35$ for $\log_2$ expression across replicate cultures. **This value is assumed and will be replaced by the pilot estimate before the study opens; every minimum detectable effect below scales linearly with it.**

```
r = 6:   SE(DiD) = 0.061181     MDE at 80% power = 0.171402 log2FC

simulated rejection rate at r = 6:
    d = 0.00  ->  0.062      (nominal alpha recovered)
    d = 0.10  ->  0.360
    d = 0.20  ->  0.902
    d = 0.30  ->  0.998
```

The study is well powered against a 0.20 $\log_2$**FC and underpowered against 0.10, and this is registered rather than discovered later.** Per-locus Bonferroni testing at $r = 6$ has an MDE of 0.785620 and is registered as **exploratory only.**

§10 FALSIFICATION. F-1 to F-7 of Section 5.7, reproduced in full. **Specifically: a DiD indistinguishable from zero at** $\alpha = 0.05$ **will be reported as "the hypothesis that this signal carries sequence-specific information into transcription is not supported by this experiment."**

§11 CLOSED RESCUE CLAUSES. The following **will not** be offered in response to a null: unregistered arms; endpoint re-specification; extended replication after looking; unsuitability of corpus, cell line, transducer, tuning route or timing.

§12 POSITIVE-RESULT REQUIREMENTS. Any positive claim requires, in addition: an address-swap arm moving the effect to the new loci with payload held constant; a payload-swap arm changing the effect with addresses held constant; independent replication from this packet without author contact. **Without all three, a positive is reported as unattributed.**

§13 PUBLICATION. Preregistration, freeze hash, all arm results, temperature logs, field maps, disclosure packet and raw data published **regardless of outcome. A null is published at equal prominence.**

36.13.1 B.1 The calibration run, filed with the registration

Simulating the exact design at zero true effect, 400 simulated studies, permutation test as specified:

```
r = 3   ->  rejection rate 0.040
r = 6   ->  rejection rate 0.062
r = 12  ->  rejection rate 0.043
```

Nominal $\alpha = 0.05$ **is recovered to within simulation error at every replicate count, so the permutation family as declared is adequately calibrated for this design. That is what the calibration requirement produces, and it takes an hour.**

36.13.2 B.2 The Layer 2 companion, and its firewall

A separate, separately registered Layer 2 study runs with the same corpus: open-label, fully disclosed, **randomised exposure schedule**, $T = 60$ periods, 40 training and 20 held out, $n = 12$ participants, four-baseline conjunction per participant, binomial aggregation.

```
per-participant false-positive rate of the four-baseline gate:  0.095
per-participant pass rate against a half-size effect:           0.482
decision rule: reject the null if k >= 4 of 12 participants pass
actual alpha:  0.02160        power vs half-size effect:  0.908
```

Registered in the same document, in the abstract: no result from this Layer 2 study will be cited in support of any Layer 1 claim. It tests the disclosed expectancy channel and nothing else [Curzi, 2026ab,ac].

Establishes. That the open problem can be posed as a preregistered, powered, falsifiable experiment with a mandatory scrambled control, a calibrated test, closed rescue clauses, and a null publishable at equal prominence — **using 24 cultures and no human participants.**

Does not establish. Any result. **A preregistration is a commitment to find out, and this appendix's entire contribution is that the commitment can be written precisely enough to be kept.**

36.14 Appendix C · Chain A carried end to end

One person, Chain A only, with arithmetic. Chain B is not exercised, because link B3 is not established and an example that exercised it would be a fiction. No clinical, medical or therapeutic claim is made; nothing here is a treatment. Every parameter is constructed for the arithmetic.

36.14.1 C.1 Setup

A person adopts Appendix A's twelve-directive corpus. $N = 12$. They read the full disclosure packet — corpus text, encoding, the 1,776-symbol transmission, tuning table, $j = 6$ and $j = 8$, symbol duration 20 ms, runtime 35.520 s, level 85 dBA, dose 0.123 per cent per rendering — together with Section 6.1's definition, the statement of link B3's status, and the failure catalogue [Curzi, 2026ac]. **They are told, in the same conversation, that nothing is claimed to reach a cell and that no clinical effect is claimed or expected.** They proceed. **Declination would have been recorded as an outcome; it was not exercised.**

36.14.2 C.2 The corpus cannot command them

```
corpus size                    N = 12, closed under negation
resultant                      0.000000000000  (computed, Section 4.4)
non-parallelism                u = 1
effective count                N_eff = 1 + 11(1) = 12   (affine, not a Hill number)
corpus surplus                 F = ln 12 = 2.484907 nats = 3.584963 bits
direction contributed by corpus 0
```

The corpus spans and does not point. Whatever the person does next, the corpus did not tell them to.

36.14.3 C.3 The selection, and its exact size

They select. Under a uniform prior over the twelve, the selection carries $\ln 12 = 2.484907$ nats $= 3.584963$ bits, **contributed entirely by them. This is the extremal case of clause H-A; under a non-uniform prior the figure is that prior's entropy and is smaller, and the honest report states which was assumed. Here: uniform, assumed, declared.**

36.14.4 C.4 The storage law, with stated parameters

The parameters below are constructed for the arithmetic and are labelled constructed. Chain A's outcome is durable change in conduct, taken as Q [Curzi, 2026a]:

```
N        = 12         categories the person distinguishes in this domain
u*       = 0.85       non-parallelism among their own framings
N_eff    = 1 + 11(0.85) = 10.3500
F(u*)    = ln 10.35 = 2.336987 nats
delta    = 0.08       per week
eta      = 0.35       per unit surplus
C        = 0.30       per week, maintenance cost
Q_0      = 0.00       new practice
```

Derived quantities:

```
eta x ln N        = 0.35 x 2.484907 = 0.869717      the ceiling on input
eta x F(u*)       = 0.35 x 2.336987 = 0.817945      the input actually collected
realized fraction = F(u*) / ln N    = 0.940471      94.05% of the ceiling
Q_infinity        = (eta F - C)/delta = (0.817945 - 0.30)/0.08 = 6.474316
Q_max (C = 0)     = eta ln N / delta = 10.871467
headroom          h(Q) = 0.569717 - 0.08 Q ;  h = 0 at Q = 7.121467
time constant     tau = 1/delta = 12.5 weeks; exact pole = -1/ln(0.92) = 11.993052 weeks
```

36.14.5 C.5 The trajectory

t	Q_t	h_t	eta x S	delta x Q + C	growing?
0	0.00000	0.56972	0.81795	0.30000	yes
1	0.51795	0.52828	0.81795	0.34144	yes
2	0.99445	0.49016	0.81795	0.37956	yes
4	1.83616	0.42282	0.81795	0.44689	yes
8	3.15157	0.31759	0.81795	0.55213	yes
12	4.09393	0.24220	0.81795	0.62751	yes
16	4.76902	0.18820	0.81795	0.68152	yes
24	5.59913	0.12179	0.81795	0.74793	yes
32	6.02515	0.08771	0.81795	0.78201	yes
40	6.24380	0.07021	0.81795	0.79950	yes
48	6.35601	0.06124	0.81795	0.80848	yes
52	6.38956	0.05855	0.81795	0.81116	yes

Growth requires $\eta S_t > \delta Q_t + C_t$. It holds at every row and **the margin narrows monotonically: 0.51795 at $t = 0$, 0.00679 at $t = 52$. Headroom falls from 0.56972 to 0.05855 across the year — success consuming its own headroom, and the structural reason the practice plateaus near 6.47 rather than continuing.** After 52 weeks Q has reached 98.7 per cent of Q_∞, consistent with an 11.993052-week time constant and roughly 4.3 time constants elapsed.

36.14.6 C.6 The sensitivity that matters

```
u = 0.00   N_eff =  1.000   F = 0.000000   Q_inf =  -3.75000
u = 0.25   N_eff =  3.750   F = 1.321756   Q_inf =   2.03268
u = 0.50   N_eff =  6.500   F = 1.871802   Q_inf =   4.43913
u = 0.75   N_eff =  9.250   F = 2.224624   Q_inf =   5.98273
u = 0.85   N_eff = 10.350   F = 2.336987   Q_inf =   6.47432
u = 1.00   N_eff = 12.000   F = 2.484907   Q_inf =   7.12147
```

The $u = 0$ row returns a negative equilibrium, and the model does not hide this: when $\eta F(u^*) < C$ the linear model's equilibrium is negative, which means the regime is unsustainable or the model must be replaced [Curzi, 2026a]. Concretely: **a person whose twelve directives have all collapsed into agreement — $u = 0$, a parallel corpus, $F = 0$ — is paying 0.30 per week in maintenance against zero collected surplus, and the practice decays. This is FM-4.7's singleton corpus arriving in the arithmetic, and it is what happens when a contradictory corpus is quietly reduced to the directives its holder already agreed with.**

36.14.7 C.7 The ledger

Quantity	Value	Status
Corpus resultant	0.000000000000	**ARITHMETIC**, computed
Direction supplied by the corpus	0	**ARITHMETIC**, by construction
Selection supplied by the person	2.484907 nats	**ARITHMETIC** in the extremal case, uniform prior declared
Surplus collected per week	0.817945	computed from constructed u^*, N
Realised fraction of ceiling	94.0471 per cent	computed
Steady state Q_∞	6.474316	computed from constructed parameters
Ceiling $\eta \ln N/\delta$	10.871467	**DERIVED** bound [Curzi, 2026a,?]
Time constant	11.993052 weeks	computed
Belief to outcome channel	expectancy, autonomic, behavioural	**CITED (measured), modest**
Disclosure cost to that channel	none demonstrated	modest, active literature
Anything reaching a cell	—	**NOT CLAIMED. Link B3 open.**
Any clinical effect	—	**NOT CLAIMED**
Any capability conferred	—	**NOT CLAIMED. Bounded by $\ln 12 = 2.484907$.**

Establishes. That the four terms compose into a chain whose every link is either established or under the operator's control, that the chain's output is computable from stated parameters, that it converges to a finite value bounded by $\eta \ln N/\delta = 10.871467$ and reaches 6.474316 under the stated cost, and **that the direction of the whole thing was supplied by the person — 2.484907 nats of it, all of it — because the corpus supplied none.**

Does not establish. That anything reached a cell. That any transcript changed. That any capability was conferred. That any clinical, medical or therapeutic effect occurred. **And it establishes nothing empirical about any person: every parameter in Section C.4 is constructed, and the trajectory is an evaluation of a model at chosen values. Link B3 is exactly as open at the end of this specimen as at the beginning.**

What has been described is the activation of what is already there, bounded by what is already there, chosen by the person it happens to. Every one of those three clauses is a restriction. **Together they are the reason the description is worth anything at all: a claim that names its own boundary, its own ceiling, and its own source of direction is a claim that can be checked, and a claim that can be checked is the only kind this framework has ever been willing to make.**

Chapter 37

Protocols at Two Scopes and a Finite Failure Catalogue: Headroom Gating, the Participation-Ratio Effective Count, and Thirty Manifestation Faults

Abstract. This paper converts the Creature Comforts framework into operating protocols at personal and group scope, gated on a single computed quantity. The headroom identity $h = \eta \ln N - \delta Q - C$ is evaluated before anything else: $h < 0$ forbids building, $h = 0$ permits maintenance only, and $h > 0$ caps every requested single-period increment at $r \le h$. A practitioner holding $N = 7$ framings at $u = 0.35$ realises 1.1314 of a possible $\ln 7 = 1.9459$ nats, 58.14 per cent, with $h_0 = 0.5877$, a fixed point $Q_\infty = 3.8641$ and a time constant of 11.99 weeks, so an evaluation at week four reads a transient. An eighteen-person unit passes the same gate at $h_0 = +0.2451$ and yet has a negative equilibrium, $Q_\infty = -0.8755$, because its realised surplus of 57.37 per cent of ceiling cannot carry its maintenance; raising non-parallelism from 0.25 to 0.45 moves the equilibrium to $+1.6211$, and cutting maintenance from 0.55 to 0.45 as well moves it to $+3.2878$. Two measurement corrections are enforced. Feasibility must average $F(U)$ rather than evaluate F at the mean U, an overstatement of 0.0790 nats or 4.31 per cent, worth 0.3952 stored units here. And the effective count is exactly the Rényi-2 participation ratio, so a six-person meeting in which one member takes 45 per cent of the speaking time counts 5.2192, not 6. The Moment equation remains demoted to ordinal use, with the Rasch repair priced at 123 items per term. Thirty failure modes are catalogued across eight coordinates; the reflexive parameter the framework requires is not identified from observational data, and the paper says so.

Keywords: headroom gate; effective count; participation ratio; Jensen correction; seeding strategy; failure catalogue; scale type

Series: B. Engineering the Surplus **Paper:** CCP-037

37.1 1 Introduction

A framework that prices interaction and storage is only as useful as the operating rule it yields, and an operating rule is only as honest as the audit that decides when it may not be run. This paper converts the Creature Comforts framework into protocols at two scopes — one practitioner working alone, and a group working together — and gates both on a single computed quantity rather than on judgement.

The gate is the headroom identity

```
h_t = eta * ln N - delta * Q_t - C_t
```

read as a precondition rather than as a diagnostic afterthought. Its three regimes are exhaustive and its verdicts are hard: `h < 0` forbids building, `h = 0` permits maintenance only, and `h > 0` caps every requested single-period increment at `r <= h`. The quantity is cheap — four parameters and a logarithm — and the claim of this paper is that computing it first changes decisions that are otherwise made on the appearance of health. The storage law whose fixed point and ceiling the gate is derived from is established as a first-order discrete filter in Curzi [2026a], and the surplus functional `F(u) = ln(1 + (N−1)u)` that supplies its input is established with its axiom set in Curzi [2026]; neither is re-derived here.

Three results follow that a reader can check in a few minutes each, and each of them corrects a computation practitioners routinely get wrong.

First, a positive gate and a shrinking system are compatible. The eighteen-person unit worked in Section 7.2 passes the gate at `h_0 = +0.2451` and has a *negative* equilibrium, `Q_inf = −0.8755`, because its realised surplus of 57.37 per cent of ceiling cannot carry its maintenance. Headroom asks whether the system could grow at its ceiling; the one-step growth test asks whether it is growing at its actual `u`. Both must be computed, and the case where they disagree is the common one. The structurally parallel split — positive headroom of +0.0211 against a step of `−0.2952` — is reported for a meeting in Curzi [2026g].

Second, feasibility must average the surplus, not evaluate the surplus at the average. Because `F` is strictly concave, Jensen's inequality makes the error one-directional: the wrong computation always overstates. On the eight measured periods of Section 5.3 the overstatement is 0.0790 nats, 4.31 per cent of the reported figure, which propagates through the storage law to 0.3952 stored units — about a third of the unit's entire accumulated stock, conjured by evaluating a concave function at the wrong point.

Third, the effective count that the protocol uses is exactly a participation ratio. The affine count `N_eff = 1 + (N−1)u` equals `N + 1 − N/PR` with `PR` the Rényi-2 participation ratio of the share vector, an identity established and verified numerically in Curzi [2026c] and re-verified over random correlation matrices in Curzi [2026]. That identity converts a stopwatch reading into a term the surplus law consumes: a six-person meeting in which one member takes 45 per cent of the speaking time has `PR = 3.3693` and an effective count of 5.2192 against a naive count of 6.

A standing caution about which effective count is meant. `1/HHI` and `exp(H)` are Hill numbers of orders two and one and are well grounded in ecology and political science [Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979]. The affine count `1 + (N−1)u` used throughout this paper is **not** a Hill number: it is a linear interpolation between 1 and `N`, adopted as a design axiom in Curzi [2026] and carrying no support in that literature. The two are related here only through the exact algebraic identity with the participation ratio, and the paper states which count it means at every use. The concentration instrument that *is* a Hill number of order two, with its worked jurisdictional case returning a naive HHI of 0.293430 and 3.41 effective jurisdictions against a bloc-HHI of 1.000000 and an effective count of 1.00, is a separate quantity developed in Curzi [2026]; the parallel result on captive domiciles, naive HHI 0.285000 and 3.508772 effective jurisdictions collapsing to 1.00 with a surplus term of exactly zero, is Curzi [2026].

The paper closes with a finite catalogue. A manifestation claim — that some belief, held and acted on, produces some change in the world — is here decomposed into eight coordinates, and thirty named failure modes are catalogued across them in six classes, each priced

by one of four identities. That decomposition into exactly eight coordinates is this work's framing rather than a theorem, and it is badged as such where it is stated.

Two negative results are carried in the same register as the positive ones. The causal step from room acoustics to the collective-intelligence factor is reported as **not confirmed**: the factor itself is measured [Woolley et al., 2010, Engel et al., 2014], and the inference that changing a room's acoustics changes it is not, so every acoustic intervention below is proposed rather than established. And the shuffled-input control is shown to be insufficient on its own: a filtered repackaging of the input passed all four empirical baselines including the shuffle at `p = 0.0005` and was caught only by the structural ceiling test, its claimed jump being 38.5 times `r_max` [Curzi, 2026u].

Sections 2 and 3 state the hypotheses and the method, including the identification problem that bounds every claim about the reflexive loop. Sections 4 and 5 develop the protocols at the two scopes. Section 6 is the catalogue. Section 7 works three instances end to end. Sections 8 to 11 discuss, verify, bound and conclude. Appendices A to C carry the numbered specifications.

37.2 2 Hypotheses

Several results below are theorems about an equation and cannot be falsified by observation; they are labelled as derivations where they appear. What can be falsified is the claim that the equations describe real accumulating systems, and that the protocols built on them change outcomes. The hypotheses separate the two, and each names the observation that would defeat it. The measurement protocol that operationalises them is specified in Curzi [2026i], whose own confirmatory application failed its second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin — and every applied claim resting on this paper inherits that result.

H1 — The gate binds. A system that adopts a programme whose maintenance cost exceeds its measured headroom loses stored order over the following horizon, whatever the programme's merits. **Predicted:** for a unit at `h_0 = +0.2451` adopting a programme costing 0.30 per period, the post-adoption gauge is `−0.0549` and the stock declines monotonically thereafter under constant parameters. **Falsified by:** a set of comparable units in which post-adoption stock trajectories are uncorrelated with the sign of `h` computed before adoption; or a unit with `h < 0` that nonetheless accumulates over several time constants with no change in `eta`, `N`, `u` or `C`.

H2 — Headroom and growth are separately informative. Positive headroom does not imply a non-decreasing step, and the disagreement is common rather than exceptional. **Predicted:** at `eta = 0.30`, `delta = 0.06`, `C = 0.55`, `Q = 1.20`, `N = 18` and `u = 0.25`, `h = +0.2451` while `eta F(u) = 0.4975 < delta Q + C = 0.6220`, giving `Q_inf = −0.8755`. **Falsified by:** a demonstration that the sign of `h` and the sign of `eta F(u) − delta Q − C` agree in substantially all measured systems, which would make the second computation redundant.

H3 — The concavity error is one-directional and material. Feasibility computed as `F` of the mean `u` exceeds feasibility computed as the mean of `F(u)`, always, and the difference is large enough to change decisions. **Predicted:** `F(mean u) >= mean F(u)` with equality only when `u` is constant across periods; on the eight-period series of Section 5.3 the gap is 0.0790 nats (4.31 per cent), worth 0.3952 stored units at `eta = 0.30`, `delta = 0.06`. **Falsified by:** an observed series in which the two computations differ in the opposite direction — which cannot occur if `F` is concave, so an observation of it would falsify the functional form rather than the correction.

H4 — Seeding topology dominates seeding count for practices requiring reinforcement. For a behaviour spreading by complex contagion, concentrating a fixed

seeding budget beats spreading it, judged by survival rather than by first-period reach. **Predicted:** 16 seeds spread across 120 people at mean degree 8 under a 2-of-8 rule leave 29.00 per cent of non-seeds with two or more seeded neighbours (30.16 people) and 39.17 per cent with exactly one — a below-threshold exposure; the same 16 concentrated in a 20-person unit activate 4.00 of 4 non-seeds at 80.0 per cent local commitment. **Falsified by:** a field comparison in which the dispersed arm shows equal or better survival after the sponsor is removed, or in which single-contact exposure is shown not to raise the cost of a later attempt.

H5 — A decline rate of exactly zero measures the decline path, not the protocol. Where a protocol records consent, a zero decline rate over many opportunities is evidence that the decline path is inoperative. **Predicted:** protocols with a functioning decline path show non-zero decline rates; the worked unit of Section 7.2 records 3 declines of 18, a rate of 16.7 per cent. **Falsified by:** an audited setting in which participants demonstrably could decline at low cost, were told so, and none did across many independent opportunities.

Three further claims in this paper are not hypotheses and are labelled here rather than left to be discovered. That a manifestation claim decomposes into **exactly** eight coordinates is a **taxonomy** — a framing choice, defensible because each coordinate corresponds to a stage of the grounding protocol [Curzi, 2026ac] or a gate of the Formula Validation Protocol [Curzi, 2026t], and not a theorem; where a fault sits outside the scheme the catalogue will not name it, and Section 10 says so. That a single non-negative scalar `Q` describes a group at all is a **modelling choice**, adopted for tractability and for the fixed-point guarantee. That the same functional form applies at personal and group scope is a **modelling postulate** about model form, falsifiable at each scope separately and never in general; the parallel statement at three scopes, with its five named breaks, is Curzi [2026].

37.3 3 Method: what is gated, what is demoted, and what cannot be identified

The method has four parts. Each step of each protocol carries a **grade** recording what may be concluded from its result (§3.1). Each protocol is gated on the **headroom identity** before anything else is computed (§3.6). Quantities whose scale type does not license the operation performed on them are **demoted in print** rather than used and hedged (§3.4, §3.8). And the central formal object of the framework — the reflexive loop in which the group influences the member and the members constitute the group — is **not identified from observational data**, which bounds what any protocol built on it may claim (§3.12).

Sections 3.1 to 3.5 state the apparatus at personal scope; Sections 3.6 to 3.11 state what group scope adds, which is a gate rather than a paragraph in each case. Section 3.12 states the identification problem, the Lucas critique applied to this framework's own parameters, and the underdetermination of the surplus functional's form, and names which remedy this paper adopts.

37.3.1 3.1 The three grades, and why every step carries one

Every numbered step here is stamped with one of three grades. The stamp is not decoration: it determines what may be concluded from the step's result.

```
GRADE T   THEOREM-BACKED
          The step's arithmetic is a proved consequence of the corpus's
          foundational results. Running it correctly produces a number whose
          MEANING is guaranteed by the mathematics. Whether that number
```

```
          describes a given life is a modelling assumption, badged.

GRADE C   MEASURED CHANNEL
          The step acts through a coupling that has been measured in the published
          literature -- expectation to neurochemistry, scheduling to retention,
          attention to encoding. Effect sizes are modest, contested, and belong to
          the source, not to this framework. The step is expected to do a little.

GRADE R   RESEARCH PROCEDURE
          The step is a well-specified experiment for which NO MECHANISM IS ESTABLISHED.
          It has no demonstrated effect. It is run to generate data, not to produce a
          result. Its expected outcome, under this framework's own diagnostic, is NULL.
```

Every step in this protocol that touches the biological layer is Grade R, without exception, and the reason is the open problem of the missing demodulator [Curzi, 2026y]: there is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal. Everything upstream of that gap is real and cited in the cell-as-receiver results [Curzi, 2026y]. A Grade R step is a research procedure and not a technique with a demonstrated effect — a sentence repeated on every step it governs in the Specification, because it is the one most likely to be lost when this section is quoted.

37.3.2 3.2 The postulate, badged, and its conditional

The keystone of the biological-layer strand of the corpus is **the author's postulate, stated here with its badge on:**

BELIEF CREATES REALITY IF THE BELIEF IS GROUNDED IN TRUTH.

The conditional clause is the entire content. Formalized (the channel-gain form of the belief postulate [Curzi, 2026ab]):

```
BELIEF  ->  [ CHANNEL ]  ->  WORLD

"grounded in truth"  ==  THE CHANNEL EXISTS AND IS REAL
```

A belief with no channel changes nothing; a belief coupled to a real mechanism changes exactly what that mechanism can change and not one unit more. **The postulate is not that belief is powerful; it is that belief is an INPUT to a system, and the system must have gain.** That is the storage law's architecture — input, conversion efficiency, output — which is why this protocol can be written as a circuit rather than as an exhortation.

Concordance for the postulate, with each column's role stated as the corpus's biological-layer constraints requires:

Column	Source	What it carries
(+1)	Levine, Gordon & Fields (1978): placebo analgesia is naloxone-reversible	**EVIDENCE.** A complete physical channel from expectation to measured outcome
(+1)	Kaptchuk et al. (2010, IBS): open-label placebo	**EVIDENCE.** The effect does not require deception. Effect sizes modest; literature active

Column	Source	What it carries
(-1)	The contemplative traditions' insistence that intention without a practice is inert	**PRECEDENT.** That the conditional was arrived at independently, without instruments
(0p)	The two integrate: neither says belief is omnipotent; both say belief requires a carrier	The Neutral Point holds without either column having to say something false

Two sources agreeing does not double the evidence when only one of them is evidence. What the (-1) column establishes is real and worth having — that the structure was findable by more than one route, which argues the structure is in the territory and not only in the instrument. It is not a second measurement. entry MF-23 of Section 6, treats the failure to keep these apart as this volume's single most likely methodological collapse.

37.3.3 3.3 The three constraints carried forward, in place

Three constraints bind this one, restated in full rather than cross-referenced, because a protocol chapter is the one a reader photocopies.

CONSTRAINT 1 — NO SUBLIMINALS (the disclosed-narrative constraint [Curzi, 2026ac]). Nothing masked, nothing below threshold, nothing concealed: no carrier the recipient cannot hear, no layer they were not told about, no content they cannot inspect. **A message hidden from its recipient cannot be consented to, and a protocol that cannot be consented to is outside what this framework will describe** (the fourth standing refusal of the corpus). Every audio artefact used here must therefore be fully audible, fully documented as to content, and inspectable as a file. Open-label placebo is why this costs nothing: the strongest verified belief-channel in the literature works *when disclosed*.

CONSTRAINT 2 — THE OPEN PROBLEM (the open problem of the missing demodulator [Curzi, 2026y]). There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal. Upstream is real: DNA data storage (Church, Gao & Kosuri, *Science*, 2012), gRNA-format addressing, Cas12a's native self-processing of its own crRNA array (Zetsche et al., *Cell*, 2015), latent capability and its activation (Yamanaka reprogramming, Nobel 2012), epigenetic regulation, and mechanotransduction (PIEZO1/2, Nobel 2021). **The channel for sequence-specific content is the question, not an answer this framework possesses.** Therefore every biological-layer step is Grade R.

CONSTRAINT 3 — THE GROUNDING DIAGNOSTIC (the grounding protocol [Curzi, 2026ac]). Six stages, in order, no skipping:

```
G-1  IDENTIFY THE CHANNEL     Name the coupling. Not "energy", not "the field".
                              A named physical or informational path with two ends.
G-2  CHECK IT IS POSSIBLE     Does anything forbid it? Conservation, scale, timing,
                              attenuation, the Gabor limit. A veto here ends the matter.
G-3  ESTIMATE ITS GAIN        How much output per unit input? Order of magnitude only.
G-4  PREDICT A MAGNITUDE      Write the number DOWN, with units and an interval,
                              BEFORE measuring. This is the step people skip.
G-5  MEASURE                  Against the pre-registered plan, with the four baselines.
G-6  COMPARE                  Predicted against observed. Report the ratio. Report it
                              whichever way it comes out.
```

Stage G-4 is the one that converts a belief into a scientific claim, and it is the one that almost no manifestation practice in circulation performs. A practice that predicts no magnitude cannot be wrong, and a practice that cannot be wrong cannot be trusted when it is right.

37.3.4 3.4 The CC-R1 demotion, stated exactly, and what it forbids

the Formula Validation Protocol [Curzi, 2026t] returns **DEMOTE on gate 2 (scale type)** for the Moment equation `M = (D/P)·E`. The reason is in the scale-type analysis [Curzi, 2026t]: division requires a ratio scale with a true zero, and D (perspective) and P (perception) are not yet instrumented that way.

What is demoted is MAGNITUDE, not STRUCTURE. The distinction is precise and this protocol depends on it:

Still standing (needs only ordinal information)	Demoted (needs ratio scale)
The limiting cases: `P -> 0` gives runaway frame; `D -> 0` gives no moment; `E -> 0` gives no moment	The numeric value of `D/P`
The fixed point `D = P`, where the moment equals the environment — undistorted contact	The distortion `Delta = |ln r|` as a measured quantity
The three-way ordering **over-framed / calibrated / under-framed**	The join `M_t -> Q_{t+1}` as a quantitative transfer (the storage law [Curzi, 2026a])
The DIRECTION of a change in framing between two occasions	The SIZE of that change

Every use of `D/P` in this section is ORDINAL: direction and ordering only, never a numeric magnitude. The instrument in §4.3 is built to produce exactly three verdicts and no number, and it is built that way on purpose. Where a reader wants the number, the Rasch-route theorem [Curzi, 2026t] gives the constructive route: a **Rasch logit difference** makes `Delta = |ln(D/P)|` genuinely ratio-scaled, because Rasch separation supplies the true zero that division requires. The cost is stated and is not small — **123 items per term.** For the two terms of the ratio that is 246 items; for all three terms of the Moment, 369 items. At a brisk twenty seconds per item, one full three-term administration is 123 minutes, or **2.05 hours**, and the cross-individual recovery gate (FVP-7) with thirty participants is **61.5 person-hours** of item administration before any analysis.

That is the price of a number, and this framework would rather print the price than a number it has not earned.

37.3.5 3.5 Consent, reversibility and decline are gates, not advice

This is the structural commitment of the section and it needs to be stated as engineering rather than as ethics, because stated as ethics it becomes optional under pressure.

A protocol is a sequence of steps. **This protocol is a sequence of steps interleaved with three kinds of GATE, and a gate is a step that can HALT the protocol.** Advice cannot halt anything. Gates can, and these do.

```
GATE TYPE K  (KNOW)      Before any step, the practitioner can state, in their own words,
                         what the step does, what it is expected to do, what grade it
                         carries, and what its null result would look like.
                         FAIL -> the step does not run. Not "runs with a caveat."

GATE TYPE D  (DECLINE)   At every gate, "no" and "not now" are recorded outcomes with
                         the same status as "yes". A decline is DATA, never a null
                         return and never a missing value to be imputed.
                         FAIL (no decline path exists) -> the protocol is void.
```

```
GATE TYPE V  (REVERSE)   Before any step, a written rollback exists: how to stop, how
                         long the residue lasts, and what cannot be undone.
                         FAIL (no rollback, or an unbounded residue) -> the step does
                         not run, at any expected benefit.
```

Three consequences follow, and they are why the gates are structural rather than decorative.

First, Gate K forbids the concealed layer by construction: a participant cannot state what a step does if part of what it does is hidden from them. Constraint 1 is not an add-on; it is enforced by the protocol's first gate.

Second, Gate D makes the decline path load-bearing on the arithmetic. the internal formulary [Curzi, 2026o] requires that an echo source who declines is recorded as missing data rather than as a null return; here, a declined step cannot silently become a zero in the storage ledger, which would bias every downstream number optimistically.

Third, Gate V is why the biological arm is audio-only and listen-only. An audio file can be stopped, and its residue is the memory of having heard it — real, bounded, and exactly the residue of anything else that has been listened to. **Nothing is imported, nothing is delivered, no Cas protein is present, and no sequence is changed** (the corpus's biological-layer constraints). The architecture cannot edit DNA and this framework does not claim that it does, has done, or could.

And the standing refusals of the standing refusals of the corpus apply here with no softening: this section makes **no clinical, medical or therapeutic claim**; it reports **no trial, no controlled observation, and no demonstrated physiological effect of the method**; and it claims **no capability has been conferred on anyone.**

37.3.6 3.6 Gate 0: measure `h` first, and if `h < 0` the group may not build

the headroom gauge [Curzi, 2026a], the master gauge, unchanged:

```
h_t := eta * ln N  -  delta * Q_t  -  C_t

h_t <  0    Unsustainable IN THIS MODEL. No interaction geometry saves it
            (the no-geometry-rescues corollary \citep{ccp:CCP-007}). THE GROUP MAY NOT BUILD.
h_t =  0    Maintenance only. Hold, do not build.
h_t >  0    Growth available; any requested increment must satisfy r_t <= h_t
            (Corollary 4.2). That is a ceiling, not a target.
```

The gate is not advice about prudence. Corollary 3.4 is a theorem: if `C_t > eta ln N - delta Q_t`, no choice of interaction geometry can prevent stored order from decreasing. A group may recruit brilliantly, diversify perfectly, and run flawless meetings, and `Q` will still fall, because the arithmetic of the leaky integrator does not care how good the inputs are once the draw exceeds the maximum the inputs can supply. the storage law [Curzi, 2026a] calls this the single most important sociological statement in the corpus, and this section treats it as a hard precondition.

What "may not build" means operationally: **no new programme, no new practice, no new ritual, no new measurement regime, and specifically no biological-layer arm.** Every one of those adds to `C`, and adding to `C` when `h < 0` moves the system further from feasibility while generating activity that looks like progress. The permitted actions when `h < 0` are exactly two: reduce `C`, or deliberately shed `Q` (the shedding lever [Curzi, 2026a], with its ethical boundary intact).

Badged, as the corpus's conventions require: **that a single scalar `Q` describes a given group at all is a modelling choice, not a theorem.** The founding papers prove things about the law;

they prove nothing about any particular team. Where the fault is outside the model — a funding cut, a legal shutdown, a death — the gauge will not name it, and the correct response is to say so rather than force a parameter.

37.3.7 3.7 Individuality is a mathematical requirement, not a courtesy

This is the load-bearing constraint of the whole chapter and the independence requirement [Curzi, 2026g] establishes it four separate ways. The four are worth restating because a group protocol is exactly the instrument that tends to violate them.

```
Condorcet (1785)  : aggregation beats the individual ONLY under independence
\citet{hong2004groups}: difference contributes what ability cannot
Cascades (1992)   : copying stops information from entering the pool
Woolley (2010)    : equality of conversational turn-taking predicts collective performance
```

Four different subjects — probability, search, sequential inference, factor structure — produced by people not working on each other's problems. **Each reduces to the same instruction: keep the individuals individual, and keep them all in.** And the surplus law prices the violation exactly:

```
S = ln(1 + (N−1)u)        merge the members -> u -> 0 -> N_eff -> 1 -> S -> 0
                          silence the members -> participating N falls -> same effect
```

Merge the individuals and the collective capability falls. That is not an ethical claim that happens to be convenient; it is `ln 1 = 0`. The cruelty of the failure is that a converged group does not feel impoverished. It feels *aligned*, and alignment feels exactly like being right. There is no internal signal that distinguishes a room that converged on the truth from a room that converged on the first speaker — only a protocol can, and only if it ran before the room started talking.

Every step of this section's protocol is therefore constrained twice: once by `h`, and once by the requirement that it must not lower `u` or reduce the participating `N`. A step that raises `Q` by lowering `u` is not a gain. It is borrowing from the ceiling to pay the floor.

37.3.8 3.8 The CC-R1 demotion at group scope

The Moment equation's group reading is the **doctrine-to-data ratio** (the corpus's internal-scope conventions): an over-framed institution is one whose doctrine outruns its intelligence and sees what its model predicts; an under-framed one drowns in data it cannot read. the group-moment analysis [Curzi, 2026r] develops it.

The Formula Validation Protocol returns DEMOTE on gate 2 for `M = (D/P)·E`, at group scope exactly as at personal scope, because D and P are not instrumented on ratio scales with true zeros (the scale-type analysis [Curzi, 2026t]). What is demoted is magnitude; structure stands.

Stands (ordinal information suffices)	Demoted (needs a ratio scale)
The limiting cases and what they name	The numeric value of `D/P` for an institution
The fixed point `D = P` — doctrine matched to intake	The distortion `Delta = |ln r|` as a measured quantity
The ordering **over-framed / calibrated / under-framed**	The join `M_t -> Q_{t+1}` as a quantitative transfer

Every use of `D/P` in this section is ordinal — direction and ordering only, never a numeric magnitude — and the group instrument in §5.11 is built to make the numeric version impossible rather than merely discouraged. The constructive route out is the Rasch-route theorem

[Curzi, 2026t]: a Rasch logit difference supplies the true zero that division requires and makes `Delta` genuinely ratio-scaled. **The cost is 123 items per term**, so 246 items for the ratio and 369 for all three terms of the Moment. At a group of 30 and twenty seconds per item, a three-term administration is **61.5 person-hours** before any analysis. That is the price of the number, and a group that wants the number should budget it rather than assert it.

37.3.9 3.9 The acoustic-to-c-factor causal test is NOT confirmed

the acoustic intervention analysis [Curzi, 2026f] states the falsification test for the claim that acoustic conditions causally raise the collective-intelligence factor `c`, **and reports the test as not confirmed.**

This section carries that status without softening. **Woolley et al. [2010] is real: there is a collective intelligence factor, it is not predicted by members' average or maximum IQ, and it correlates with social sensitivity, equality of conversational turn-taking, and the proportion of women in the group.** Those are measured correlations in the published work.

What is not established is that changing a room's acoustics changes `c`. The mechanism is plausible — turn-taking equality is a measured correlate of `c`, and acoustics plainly affect turn-taking — but plausible is not confirmed, and the acoustic intervention analysis [Curzi, 2026f]'s own test came back negative. **Therefore every acoustic intervention in this section is PROPOSED, not established, and is labelled so at each appearance.** A group may run one as an experiment. A group may not report one as a capability.

37.3.10 3.10 Consent at group scope, and the hazard a person does not have

Section 4's three gates carry over unchanged in form:

```
GATE K  (KNOW)     Every participant can state what the step does, what grade it
                   carries, and what its null result looks like.  FAIL -> no run.
GATE D  (DECLINE)  "No" and "not now" are recorded outcomes with the same standing
                   as "yes". A decline is DATA, never a null return and never an
                   imputed zero.  FAIL (no decline path) -> the protocol is void.
GATE V  (REVERSE)  A written rollback exists: how to stop, what residue remains,
                   and what cannot be undone.  FAIL -> the step does not run.
```

At group scope one more gate is required, and it exists because of a hazard that has no individual analogue.

```
GATE I  (INDEPENDENCE)  Before any step that could align beliefs, the group's
                        capacity to disagree is MEASURED, and it is measured again
                        after. If the step lowered u or reduced participating N,
                        the step is reversed regardless of what else it produced.
```

A group can manufacture the appearance of consent, and a group running a change programme is the archetypal machine for doing so. the two-mechanisms-of-commitment analysis [Curzi, 2026g] already draws the line between two mechanisms that produce apparent commitment: a person who has adopted a practice because it works for them, and a person who has adopted it because dissent is expensive. **The two are indistinguishable from the outside and completely different in the ledger, because the second one has `u = 0` with respect to the group's belief and contributes nothing to S.**

Gate I is what makes the difference detectable. It has a hard operational form: **a decline rate of exactly zero across many opportunities is a measurement of the decline path, not of the protocol.** If nobody has ever said no, Gate D has not been satisfied — it has been declared.

the standing refusals of the corpus's fourth refusal governs absolutely: **this framework does not recommend an intervention that removes a person's ability to decline.** At group

scope that refusal has teeth only if the decline is cheap, visible, and recorded, which is why B.4 and B.17 below are procedures rather than sentiments.

37.3.11 3.11 The concordance for this section's spine

Each column's role is stated, as the corpus's biological-layer constraints requires. **The (+1) column carries evidence; the (-1) column carries precedent — that the structure was arrived at independently, by people without instruments, which argues the structure is in the territory and not only in our instrument. It is not a second measurement.**

Claim	(+1) EVIDENCE	(-1) PRECEDENT	Status
Aggregation requires independence	Condorcet 1785 (theorem); informational cascades [Bikhchandani et al., 1992, Banerjee, 1992]	Deliberative traditions that require each voice before any is answered	**Established**; the theorem is proved, the transfer to a given group is a badged assumption
Difference outperforms uniform ability	Hong and Page [2004]	Guild and council practices of seating unlike trades together	**Established** in the cited model class
Equal turn-taking predicts group performance	Woolley et al. [2010]	The talking-stick and round-robin forms found independently in many cultures	**Established as correlation.** Causal direction not settled here
Acoustics raise `c`	**NONE. the acoustic intervention analysis [Curzi, 2026f] reports the causal test as NOT CONFIRMED**	Traditions of purpose-built assembly acoustics	**NOT ESTABLISHED. Proposed only**
A committed minority can flip a convention	Centola et al. [2018]: ~25% in controlled online experiments	The persistent folk observation that a determined few change a house's habits	**Established in the experimental setting.** Field transfer is a modelling assumption
Clustered seeding beats broadcast	Centola and Macy, 2007, Centola, 2010 — complex contagion needs reinforcing contacts	Apprenticeship and cell structures that propagate by dense local ties	**Established** for complex contagion
Durable institutions come from designed commons rules	Ostrom [1990]	Centuries of surviving common-pool arrangements the corpus documents	**Established**
Acoustic input carries sequence-specific content to a cell	**NONE — the open problem (the open problem of the missing demodulator [Curzi, 2026y])**	Correspondence doctrines linking sound and form	**NOT ESTABLISHED.** Speculative, and marked speculative here in the sentence that states it

---

37.3.12 3.12 Identification: what no protocol can recover from observation

The framework's reflexive claim is that the group's state influences each member and the members constitute the group: `T(x) = W(x, M(x))` in the corpus's notation, and in the notation of econometrics a **linear-in-means reflexive specification**. Manski [1993] proves that the endogenous social effect in that specification is **not identified from observational data**: it cannot be separated from contextual effects, nor from correlated unobservables, no matter how many periods are collected. Shalizi and Thomas [2011] reinforce the point for network data, where influence and selection are generically confounded.

This bounds every protocol in this paper, and the bound is stated here rather than in a footnote. The protocols measure `u`, `N`, `eta`, `delta`, `C` and `Q`, compute `h`, and act on its sign. What they do not do — and what no amount of held-out fit would license — is recover a parameter representing the group's influence on its members from observing the group. A fitted value of `eta` in a system whose members are influencing one another absorbs contextual and correlated effects into itself; it is a summary of the fitted series, not a measured coupling.

The known remedies are randomisation, an instrument, partial-population variation, or stipulating the reference group and defending the stipulation, the last on the template of Brock and Durlauf [2001]. **None of them is free, and this paper adopts none of them.** The worked instances of Section 7 are constructed examples with declared parameters, so the question of identification does not arise for them arithmetically; it arises the moment the protocol is run on a real unit, and at that point **the claim that the fitted parameters describe the unit's social mechanism is a modelling position rather than an identified empirical result.** Section 10 records this as the paper's most serious unmet standard.

Two objections in the same family are also unanswered here and are recorded rather than argued away.

The Lucas critique, turned on this framework's own parameters. Lucas [1976] holds that behavioural parameters estimated under one policy regime need not survive a change of regime, because agents re-optimise. The framework asserts that systems change their behaviour when they are modelled and observed — that is the whole content of the reflexive loop — and then proceeds to treat `eta`, `delta` and `u` as stable enough to fit. **Either those parameters are stable under the framework's own predicted response to being measured, or the framework undercuts itself.** Hacking [1995] names the mechanism in the human sciences directly: a classification changes the people classified. This paper takes the first horn as an assumption and labels it an assumption; it does not defend it, and a group re-measuring `u` after adopting the protocol is measuring a `u` the protocol may itself have moved. The one field study in which a reflexive fixed point is documented with dates and price data — MacKenzie and Millo [2003] on option pricing — is the corpus's best existence proof that the mechanism is real, and it is an existence proof rather than a licence to treat a fitted coefficient as a measurement.

The functional form is underdetermined by the evidence. The logarithm `F(u) = ln(1 + (N−1)u)` is not the only form that fits the qualitative facts. Latané's social impact theory proposes a power law with exponent below one [Latané, 1981]; Odlyzko and Tilly's critique of Metcalfe's law proposes `n log n` [Briscoe et al., 2006]; the corpus proposes the logarithm. **All three deliver diminishing returns in count, and nothing in the worked examples of this paper discriminates between them.** The paper's own honest position is that the logarithm follows from the affine-effective-count design axiom of Curzi [2026] and from nothing else established here.

One genuine strengthening should be recorded alongside. The usual warrant for the logarithm is Fechnerian — sensation proportional to the logarithm of a stimulus ratio measured from a floor [Fechner, 1966] — and that warrant requires the perceptual story to be true. Peters and Gell-Mann [2016] derives a logarithmic evaluation from multiplicative dynamics with **no psychological assumption at all**: the time-average growth rate of a multiplicatively evolving quantity is the expectation of its logarithm. Where a stock compounds rather than adds, that is the warrant the argument actually needs, and it is the stronger of the two because it does not require Fechner to be right.

37.3.13 3.13 Analogy discipline, and what is borrowed from where

The framework borrows form from physics and engineering throughout, and every borrowing below is one of three things, labelled at first use.

- **A shared mathematical form.** The storage law is a first-order discrete filter with a pole at `z = 1 − delta`; the whole apparatus of such filters — DC gain, time constant, settling time — applies because the equation is the same equation, with no claim that the underlying physics transfers [Oppenheim et al., 1996, Åström and Murray, 2008]. The leaky integrator has the same status in computational neuroscience [Gerstner and Kistler, 2002], and geometric convergence follows from the contraction property [Banach, 1922, Elaydi, 2005]. **This is the category most of this paper sits in, and it is a respectable one.**
- **A derivation.** Jensen's inequality applies to `F` because `F` is strictly concave, which is proved, not assumed [Curzi, 2026]. The participation-ratio identity is algebra [Curzi, 2026c]. The Gabor limit `dt df >= 1/(4 pi)` is a theorem of Fourier analysis and constrains any rendering of information into sound [Gabor, 1946]; it is classical, it applies to every wave, and it carries no quantum content.
- **An analogy, and never evidence.** The reading of an effective count as a count of independent Bernoulli voters in Section 5.6 is a **modelling convenience**: Condorcet's setup assumes independence outright, and `N_eff` audits redundancy rather than counting independent trials. The stitch is informative about direction and order of magnitude and is claimed to be nothing more.

One inheritance is presupposed rather than restated. The framework's founding definition — comfort as the felt total of present costs, so that the present state is a local minimum of the comfort landscape by definition rather than by luck — is established in Curzi [2026], together with the correction that the variational statement is `delta–S = 0`, *stationarity* and not minimality, and that the variational form licenses **no teleology whatever**. Nothing in this paper's protocols may be read as a claim that a system is seeking anything.

Entropy in this paper is Shannon entropy where it appears at all — the surplus is a logarithm of an effective count reported in nats, a pure number — and no thermodynamic quantity is invoked. The participation ratio is the reciprocal of a Rényi-2 collision probability [Rényi, 1961], which is why it is an effective count in the standard sense [Hill, 1973] even though the affine count that equals it algebraically is not.

37.4 4 The protocol at personal scope

Everything in this framework up to here has been either apparatus or argument. This section is the first one that a single person can run tonight, alone, with a notebook and a timer, and get a number out of by the end of the quarter. It keeps that promise in the only way it can be kept honestly: by being exact about which parts of the practice rest on theorems, which rest on measured channels, and which are research procedures with no demonstrated effect at all. **A protocol that mixes those three grades without labelling them is not a protocol; it is a mood.**

The section joins two things. From the formulary it takes the internal-scope formulary [Curzi, 2026o] — nine formulas at internal scope, each with its Formula Validation Protocol row already printed. From the biological-layer strand of the corpus it takes the biological layer: the cell as receiver (the cell-as-receiver result [Curzi, 2026y]), the library of latent capability (the latent-capability result [Curzi, 2026y]), addressing without editing (the addressing-without-editing result [Curzi, 2026y]), narrative as regulatory input (the disclosed-narrative constraint [Curzi, 2026ac]), and — decisively — the open problem of the missing demodulator [Curzi, 2026y]. The join is where a reader is most likely to be sold something. So the join is where this section is most careful.

37.4.1 4.1 The three layers, and their honest gain

A personal protocol built on this framework has exactly three layers. They are not equal partners.

Layer	Grade	What it acts on	Honest expectation
L1 Informational	T	The practitioner's framing: `D`, `P`, `N`, `u` — the lens and its variety	Real and immediate. The arithmetic is proved; the fit to a particular life is a badged modelling assumption
L2 Behavioural	C / T	The practitioner's schedule, cadence, load: `eta`, `delta`, `C`	Real, slow, and cumulative. The leaky integrator does the work; cadence beats intensity (the cadence result [Curzi, 2026a])

Layer	Grade	What it acts on	Honest expectation
L3 Biological	R	Nothing demonstrated	**Expected null.** A research procedure with no established mechanism

Effectively everything this protocol will actually do for a person happens in L1 and L2, and the section says so at the front rather than at the back. A reader who runs only L1 and L2 is running the whole of what has been demonstrated. A reader who runs L3 as well is running an experiment, and is told so in the sentence that invites them to run it.

The layers compose through one circuit — the internal formulary [Curzi, 2026o] with the biological arm drawn as what it is, an open lead:

```
L1 (lens)                          L2 (integrator)
E -> [ D/P ordinal ] -> M   ----->  Q_{t+1} = (1-d)Q_t + eta*M - C
       ^                                |-> h = eta lnN - dQ - C
       | F_int = ln(1+(N-1)u)          |-> L = max(0, M - C)
       +----- echo: Delta_R = |ln(M'/M)| <---- world

L3 (biological arm)
audible, disclosed, inspectable artefact
       |
       X   NO ESTABLISHED DEMODULATION MECHANISM   <-- the open problem of the missing demodulator \cite
       |
       v   (nothing downstream is claimed)
```

The `X` in that diagram is the honest content of the biological-layer strand of the corpus. It is drawn, not omitted.

37.4.2 4.2 The channel inventory for one person

the channel-gain form of the belief postulate [Curzi, 2026ab] requires that a channel be *named*. Here is the inventory a single operator has. **Nothing on this list is a medical or therapeutic claim; each row states something about the manual's own parameters, or reports a published finding belonging to its source.**

#	Channel	Grade	Acts on	Status
1	Expectation -> physiology	C	outside the framework's parameters	Measured in the placebo literature; analysed with its channel-gain form in Curzi [2026ab]. Naloxone-reversible for analgesia (Levine, Gordon & Fields 1978); open-label variants work (Kaptchuk et al. 2010). Modest effects, active literature. Neither primary source carries a key in this corpus's external bibliography (L-14)
2	Cadence -> retention	T	`delta`, `Q`	Periodic forcing of a leaky integrator (the cadence result [Curzi, 2026a], the periodic-forcing analysis [Curzi, 2026a]). Pure consequence of the storage law
3	Variety of framing -> ceiling	T	`N`, `u`, `F_int`	the surplus-functional theorem [Curzi, 2026]. `S = ln(1+(N−1)u)`, ceiling `ln N`
4	Prediction logging -> calibration	C	`D` relative to `P`	Ordinal only under CC-R1. Direction of miscalibration is recoverable; magnitude is not
5	Load shedding -> headroom	T	`C`, `h`	the headroom gauge [Curzi, 2026a]. `h = eta ln N − delta Q − C`
6	Echo -> frame correction	T	`D`	the returned-difference result [Curzi, 2026]. `Delta_R = |ln(M'/M)|`; zero means nothing came back

#	Channel	Grade	Acts on	Status
7	Acoustic -> mechanotransduction	C	a cellular signal, non-specific	PIEZO1/2 are mechanically gated; sound is mechanical (the corpus's biological-layer constraints). **Real, established, and non-specific**
8	Acoustic -> sequence-specific content	**R**	**nothing demonstrated**	**No established mechanism (the open problem of the missing demodulator [Curzi, 2026y]). This is the open problem**

Rows 7 and 8 are the two a careless reader will merge, and merging them is the single most consequential error available in this strand of the corpus. Row 7 says a pressure wave can become a cellular signal — Nobel-recognized and true. Row 8 says it can carry *which* signal, addressed to a *particular* sequence, and that is exactly what has no mechanism. Section 6 catalogues the merge as MF-03.

37.4.3 4.3 The ordinal calibration instrument

The Moment's structure survives the CC-R1 demotion; its magnitude does not. The personal instrument is therefore built to return the structure and refuse the magnitude.

THE PREDICTION LOG. For a defined class of recurring occasion — a meeting, a market open, a particular conversation, a training session — the practitioner records, *before* the occasion:

```
1. What is expected to happen, in one sentence, specific enough to be wrong.
2. Which single observable will be used to check it.
3. The direction that observable is expected to move.
```

Afterwards the observable's actual movement is recorded. No ratio is computed. One of three verdicts is assigned, and only three:

```
OVER-FRAMED    ( r > 1 )   The prediction was seen rather than the occasion; the
                           observable did not support it; the expected support was
                           noticed and the disconfirmation was not. Frame outran
                           intake.

CALIBRATED     ( r = 1 )   Prediction and observable agreed, including where
                           the prediction was wrong and that was recorded.

UNDER-FRAMED   ( r < 1 )   The occasion delivered more than the frame could index.
                           It can be described but not used.
```

These three verdicts are an ORDERING, and an ordering is all the CC-R1 demotion permits, which is why this instrument produces no number and should not be made to produce one. Verdicts may be counted across occasions — twelve over-framed, four calibrated, six under-framed is a legitimate tally of ordinal outcomes. That tally may **not** be converted into a value

of D/P, and no distortion Delta may be reported. Where an instrument reports a "perception ratio" of 1.4, entry MF-08 of Section 6 is the entry to look up.

The tally is still a genuine diagnostic, because the ordering carries the content:

Dominant verdict	Reading	Lever
OVER-FRAMED	Doctrine outrunning intake; the model is being seen rather than the occasion	Raise P: add an observable outside the practitioner's control, and one adversarial source
CALIBRATED	Contact. The rarest of the three and the one to protect	Hold. Add no instrumentation that will not be read
UNDER-FRAMED	Flooding; intake without a frame that can index it	Raise D: acquire the vocabulary, not more exposure

37.4.4 4.4 The internal surplus audit

the internal surplus [Curzi, 2026o] F_int = ln(1 + (N−1)u) over one person's own framings — the founding material's the founding material HORIZON, "the many windows of the one tower," as a derivation rather than an analogy. For a person holding N = 7 genuinely distinguishable framings of their situation, ln 7 = 1.9459 nats is the hard ceiling, and u decides how much of it is collected:

u	N_eff	S = ln N_eff (nats)	Fraction of ceiling
0.00	1.0000	0.0000	0.00%
0.05	1.3000	0.2624	13.48%
0.10	1.6000	0.4700	24.15%
0.25	2.5000	0.9163	47.09%
0.35	3.1000	1.1314	58.14%
0.50	4.0000	1.3863	71.24%
0.80	5.8000	1.7579	90.34%
1.00	7.0000	1.9459	100.00%

Seven framings that are all the same framing purchase exactly zero surplus, because ln 1 = 0, and the law is not being poetic. Strict concavity (the strict-concavity result [Curzi, 2026]) delivers the design fact: u = 0 to u = 0.25 buys 47% of everything available. **The first genuinely different way of seeing one's own situation is worth more than all subsequent ones combined.**

Measuring one's own u is the hard part, and the honest procedure is indirect: list the N framings, then ask of each pair whether they ever *disagree about what to do next*. Pairs that never disagree are parallel. If m of the N(N−1)/2 pairs have produced a different recommended action, a rough estimate is u \ensuremath{\approx} m / [N(N−1)/2]. For N = 7 there are 21 pairs; seven live disagreements gives u \ensuremath{\approx} 0.33. **This is a crude instrument, labelled as one: use it to detect the collapse at the left end of the table, which is large, not to resolve differences in the middle, which it cannot.**

37.4.5 4.5 The self-storage run

the storage law [Curzi, 2026a]'s join: Q_{t+1} = (1 − delta)Q_t + eta·M_t − C_t. **The join is STRUCTURAL, not yet quantitative** (CC-R1), so M_t enters as a normalized input in [0, ln N]

and the numbers below are an illustrative construction, stated as such, not a measurement of anybody.

Take a weekly period, `N = 7`, `u = 0.35` so `S = 1.1314`, and parameters `delta = 0.08` (what a week of no practice costs in retention), `eta = 0.45` (how much of a week's engagement becomes durable structure), `C = 0.20` (weekly maintenance), starting from `Q_0 = 1.10`.

```
Q_inf = (eta*S − C) / delta = (0.45*1.1314 − 0.20) / 0.08 = 3.8641 Q-units
tau   = −1 / ln(1 − delta)  = 11.9931 weeks
half-life to the fixed point = ln(2) * tau = 8.3130 weeks
```

week `t`	`Q_t`	`h_t = eta ln N − delta Q_t − C`
0	1.1000	0.5877
4	1.8839	0.5249
8	2.4455	0.4800
12	2.8479	0.4478
16	3.1361	0.4248
20	3.3426	0.4083
24	3.4905	0.3964

Three things a reader should take from this table and from no other table in the section.

First, the time constant is twelve weeks and the half-life is eight. Not eight days, not one intense weekend. **`delta = 0.08` per week sets the pace, and no amount of enthusiasm in week one changes `delta`.** This is why cadence beats intensity (the cadence result [Curzi, 2026a]): the integrator does not care how large a single input was, only how many arrived before the previous ones leaked away.

Second, headroom falls as stored order grows. `h` drops from 0.5877 to 0.3964 over twenty-four weeks with no parameter changed, purely because `delta·Q` rises. A practice that worked in month one and feels heavier in month six is `delta·Q` doing arithmetic, not failing willpower (the falling-headroom result [Curzi, 2026a]).

Third, the fixed point is real and arrives. Because `delta > 0` the map is a contraction and Banach's fixed point theorem guarantees a unique fixed point reached from any start — which is why iteration *arrives* rather than wandering (the contraction theorem [Curzi, 2026]). The founding material's fixed-point node is the fixed-point theorem of the storage law [Curzi, 2026a].

The corresponding fixed-point headroom is CC-R6's identity: `h_inf = eta·(ln N − S) = 0.45 × (1.9459 − 1.1314) = 0.3665`. **`delta`, `C` and `Q` all cancel at the fixed point.** The headroom of a mature personal practice depends on one thing only: the gap between the ceiling `ln N` and the surplus `S` actually collected. **A mature practice with headroom near zero does not have a discipline problem. It has a `u` problem.**

37.4.6 4.6 The headroom gate

`h_t = eta·ln N − delta·Q_t − C_t`, the headroom gauge [Curzi, 2026a], at internal scope. The gate:

```
h_t <  0    Unsustainable IN THIS MODEL. No arrangement of framings rescues it
            (Corollary 3.4). DO NOT ADD ANYTHING. Shed cost or shed stored order.
h_t =  0    Maintenance only. Hold. Do not build.
h_t >  0    Growth available. Any requested increment must satisfy r_t <= h_t
            (Corollary 4.2). That is the maximum, not a target.
```

This gate runs before every addition to the practice, and before the biological arm in particular, because that arm is a Grade R research procedure with an expected null result and a real cost in C. Spending scarce headroom on an arm whose expected yield is zero is a defensible research decision and an indefensible self-improvement decision, and the practitioner should know which one they are making.

The distinction half of all misdiagnosis turns on (the headroom-versus-growth diagnostic split [Curzi, 2026g]) applies here: `h` uses `ln N`, the growth test uses `S = ln N_eff`. **A system can have comfortable headroom and still lose stored order every week, if operation is far below the ceiling.** Both must be computed.

37.4.7 4.7 The biological arm, stated as what it is

Here is the whole of what the biological arm consists of, written so that a reader can see there is nothing behind the curtain.

What is done. Text the practitioner wrote about their own situation is encoded in base 4 (Church, Gao & Kosuri, *Science*, 2012, is the reference for base-4 encoding of arbitrary data as sequence), given a gRNA-format routing prefix as an *address format*, not a molecule (the addressing-without-editing result [Curzi, 2026y]), and rendered to audio by the method of the disclosed sonification chain [Curzi, 2026y]. The practitioner listens at an audible level, knowing exactly what it contains, with the file available for inspection.

What is claimed. That the encoding and rendering are reproducible, that the practitioner listened, and that they wrote a prediction beforehand.

What is NOT claimed, in the same breath. No sequence is changed. No molecule is delivered. No Cas protein is present. **Nothing is edited, altered, or damaged, and the architecture could not do so** (the corpus's biological-layer constraints). No capability is conferred. No clinical, medical or therapeutic effect is claimed, observed, or implied. And, decisively: **because there is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal, the sequence specificity of the artefact has no demonstrated route to any biological effect whatsoever.**

What constrains the rendering. The Gabor limit, `dt·df >= 1/(4·pi) = 0.079577` — classical, applying to every wave including sound, the same Fourier mathematics as Heisenberg with no quantum content (the time–bandwidth theorem [Gabor, 1946]). It hard-constrains rendering speed:

window	minimum frequency resolution
1 ms	79.577 Hz
5 ms	15.916 Hz
10 ms	7.958 Hz
20 ms	3.979 Hz
50 ms	1.592 Hz
100 ms	0.796 Hz
1000 ms	0.080 Hz

A short sound cannot have a precise pitch. If the encoding must distinguish tones 4 Hz apart, no window shorter than about 20 ms can carry it, and a "fast" rendering is not faster but smeared — it no longer encodes what it claims to. **This is a real, computable, hard constraint, and the kind whose presence shows the method is at least specified well enough to be constrained.**

What the honest expected result is. Null. The grounding diagnostic of §3.3, applied to this arm, halts at stage G-1: **the channel for sequence-specific content cannot be named because it has not been established.** The arm is run — if at all — as a Grade R procedure whose value is the data it generates about a well-specified gap. **The reason to run it is that the open problem**

of the missing demodulator [Curzi, 2026y] states exactly what experiment would settle the question, and someone has to run experiments.

37.4.8 4.8 The grounding diagnostic, worked on three beliefs

The six stages of §3.3, applied to three beliefs a person might actually hold. One passes, one is demoted, one is rejected. **A diagnostic that never rejects anything is not a diagnostic** (the same standard the concordance method sets [Curzi, 2026w] for the concordance method).

BELIEF A — "If I tell myself before each session that this practice will help me concentrate, I will concentrate better."

```
G-1  CHANNEL      Expectation -> attentional set -> performance. Named, with two ends.
G-2  POSSIBLE     Yes. Expectancy effects on performance are documented, modest, and
                  have a contested methodological history (Rosenthal & Jacobson 1968
                  drew serious criticism; report honestly).
G-3  GAIN         Small. Order of magnitude: a fraction of a standard deviation.
G-4  MAGNITUDE    Pre-register: d in [0.10, 0.40] on the chosen observable.
G-5  MEASURE      60 paired sessions, A/B by pre-assigned randomization, one observable.
G-6  COMPARE      Report the observed d with its 95% interval against [0.10, 0.40].
```

VERDICT: PASSES. Grade C. It is grounded — the channel is real — and it is small, and the protocol says so in advance. Note what "passes" means: it means the belief is *eligible to be tested*, not that it is true.

BELIEF B — "My frame is 1.4 times too strong for my intake, so I should reduce it by 29%."

```
G-1  CHANNEL      Framing -> perception -> action. Real (predictive processing, the belief-failure classes
G-2  POSSIBLE     Yes for the DIRECTION. No for the NUMBER.
G-3  GAIN         Unrecoverable at present, because gate 2 of the FVP returns DEMOTE:
                  D and P are not on ratio scales with true zeros.
G-4  MAGNITUDE    CANNOT BE STATED. The quantity "1.4" is not a measurement.
```

VERDICT: DEMOTED at G-4, per R1. The belief's *structure* stands — the practitioner may well be over-framed, and the ordinal instrument of §4.3 can establish that — but **the magnitude is not reportable and the "29%" is a number with no scale behind it.** The route out is the Rasch-route theorem [Curzi, 2026t]'s Theorem 54.3 and it costs 123 items per term. Until then: run the ordinal instrument, act on the ordering, and print no ratio.

BELIEF C — "Listening to my encoded narrative will switch on a latent capability in my cells."

```
G-1  CHANNEL      HALT. The claimed channel is acoustic transmission of sequence-specific
                  content to a cell that demodulates it. NO SUCH MECHANISM IS ESTABLISHED
                  (the open problem of the missing demodulator \citep{ccp:CCP-031}). The channel cannot be
                  in the established literature.
```

VERDICT: REJECTED at G-1, in print, in this framework, about this framework's own subject. The belief fails the postulate's conditional at the first clause. It is not grounded, and by the postulate's own logic a belief that is not grounded does not create anything.

Three things must follow that rejection, because they are what makes it a rejection rather than a retreat. **First**, the rejection is of the *claim*, not the *experiment*: running the artefact as a Grade R procedure with a pre-registered null is legitimate and is what the open problem of the missing demodulator [Curzi, 2026y] asks for. **Second**, the upstream biology is untouched — mechanotransduction is real, latent capability is real (the latent-capability result [Curzi, 2026y]), and Yamanaka reprogramming shows a differentiated cell reverted by four regulatory signals without adding information. **Third**, the latent-capability result [Curzi, 2026y]'s hard constraint stands regardless: **nothing not already latent can be reached**, so even a hypothetical success would be an activation and not an addition.

37.4.9 4.9 Predicting a magnitude, and what a null looks like

Stage G-4 requires a number written down in advance. Stage G-6 requires that the comparison be reported whichever way it comes out. That means deciding, in advance, **what would count as nothing** — and the answer "the p-value was above 0.05" is not that answer, because an underpowered study produces that answer whether the effect exists or not (entry MF-17 of Section 6).

The correct instrument is a **smallest effect size of interest (SESOI)** fixed before data collection, and an equivalence test against it. Two numbers govern the design.

DETECTION is paired observations for 80% power at alpha = 0.05, two-sided, `n = (z_{a/2} + z_b)^2 / d^2`. **EQUIVALENCE** is observations needed to conclude "nothing of interest happened" at a given SESOI, 80% power, one-sided alpha = 0.05, true effect zero. **RESOLUTION** is the half-width of a 95% interval in SD units at a given `n`.

`d` or SESOI	detection `n`	equivalence `n`	`n`	95% CI half-width
0.10	785	—	10	0.6198
0.20	197	155	20	0.4383
0.30	88	69	30	0.3578
0.35	65	51	40	0.3099
0.40	50	—	60	0.2530
0.50	32	25	100	0.1960
1.00	8	—	200	0.1386

Read the right-hand column carefully: a ten-observation personal experiment cannot distinguish "no effect" from "an effect of two thirds of a standard deviation." Ten observations is what most personal experimentation actually consists of. **The reason personal experiments feel inconclusive is not that the subject matter is subtle; it is that ten observations resolve nothing, and this is arithmetic rather than mystery.**

37.4.10 4.10 What to conclude from a null, in four cases

A null result is not one thing. Before interpreting one, the analyst must know which of these four it is, and the four are distinguished by numbers already computed.

```
N-1  UNINFORMATIVE NULL
     The 95% interval is wider than the SESOI. Nothing has been learned about
     the effect's existence. CONCLUDE: nothing. Do not report "it didn't work."
     DISTINGUISH BY: CI half-width vs SESOI.
```

```
N-2  INFORMATIVE NULL (equivalence established)
     The whole 95% interval lies inside +/- SESOI. This is positive evidence
     that any effect is smaller than the declared SESOI. CONCLUDE: for this
     purpose, the channel's gain is below the threshold of interest.
     DISTINGUISH BY: interval fully contained.

N-3  NULL AGAINST A MAGNITUDE PREDICTION
     A d in [0.10, 0.40] was predicted at G-4 and the observed interval excludes
     that range. CONCLUDE: the GAIN ESTIMATE at G-3 was wrong, which is a
     different and more useful finding than "no effect."
     DISTINGUISH BY: predicted interval vs observed interval.

N-4  EXPECTED NULL ON A GRADE R ARM
     The arm had no established mechanism going in. The null is the prior
     expectation and confirms nothing new. CONCLUDE: the well-specified gap
     of the open problem of the missing demodulator \citep{ccp:CCP-031} remains where it was. Report it
     exists precisely because nulls of this kind go unpublished.
     DISTINGUISH BY: the step's grade stamp.
```

The most common error in personal practice is treating an N-1 as an N-2 in the direction the operator was already leaning — over-framing (§4.3) applied to one's own experiment. The defence is that SESOI and sample size were fixed before the first observation, which is why the Specification puts them at step 6 and not at step 20.

37.4.11 4.11 Reversibility, and the residue statement

Gate V requires a written rollback before any step — three lines:

```
STOP:      the single action that ends this step now.
RESIDUE:   what remains after stopping, and for how long.
IRREVERSIBLE: what cannot be undone. If this line is non-empty, the step is
              escalated: it does not run without a second, separate Gate K.
```

For each layer of this protocol:

Step class	STOP	RESIDUE	IRREVERSIBLE
L1 prediction log	Stop writing entries	The written record, which the practitioner owns and can destroy	Nothing
L1 framing audit	Stop the audit	Awareness of one's own framings	A framing once named cannot be un-noticed. Escalate: Gate K states this explicitly
L2 cadence change	Return to prior schedule	`Q` relaxes toward the old fixed point with `tau = 11.99` weeks	Nothing, but the relaxation is slow and must be expected

Step class	STOP	RESIDUE	IRREVERSIBLE
L2 load shedding	Resume the shed commitment	Depends entirely on what was shed	**Often non-empty.** Shedding a commitment that involves other people crosses into the shedding lever [Curzi, 2026a]'s ethical line. Escalate
L3 biological arm	Stop the file	Memory of having heard it	Nothing. **No sequence is changed; no molecule is delivered; nothing is edited**

The row that catches people is not the biological one; it is L2 load shedding, the fourth lever of the storage law [Curzi, 2026a] and the one with the largest real residue. An audio file has an off switch. A resigned commitment does not.

37.4.12 4.12 The concordance for the practice as a whole

Claim of the practice	(+1) EVIDENCE	(-1) PRECEDENT	Status
Cadence beats intensity	Leaky-integrator arithmetic; `tau = 11.99` weeks computed above; spacing effects in memory research	Monastic and craft traditions built around fixed daily office rather than episodic exertion	**Established as arithmetic**; the transfer to a person is a badged modelling assumption
Difference of framing raises capability	the surplus functional and its strict concavity [Curzi, 2026] (mathematics); Hong and Page [2004] at group scope	The counsel, in many contemplative lineages, to hold a question from several vantages before deciding	**Established as mathematics.** Empirical transfer to one mind is not measured here
Disclosed belief still has a channel	Open-label placebo (Kaptchuk et al. 2010), analysed in Curzi [2026ab]; modest effects, active literature	Traditions that require the practitioner to *know* the rite, not to be tricked by it	**Established, modest**
Frame-to-intake ordering is diagnostic	Predictive processing read as precision-weighting — an active and contested programme, named in the source without bibliographic references (L-14); structural similarity is not identity	The perennial distinction between seeing and seeing-what-one-expects	**Structure established; MAGNITUDE DEMOTED (CC-R1)**

Claim of the practice	(+1) EVIDENCE	(-1) PRECEDENT	Status
Acoustic input reaches cells	PIEZO1/2 mechanically gated channels, Nobel 2021; Tajik et al. *Nature Materials* 2016 (surface force -> chromatin stretch -> transcription)	Traditions of sounded practice in nearly every recorded culture	**Established, and NON-SPECIFIC**
Acoustic input carries sequence-specific content to a cell	**NONE. This is the open problem** (the open problem of the missing demodulator [Curzi, 2026y])	Correspondence doctrines linking sound and form	**NOT ESTABLISHED. Grade R.** Speculative, and marked speculative here in the sentence that states it

The last row is the one this section exists to keep honest. It has a (-1) entry and no (+1) entry, and the corpus's biological-layer constraints is explicit that precedent is not evidence. A row with precedent and no evidence is a research question with a long history — genuinely interesting to have, and not a result.

37.4.13 4.13 Failure modes specific to the personal scope

Eight faults specific to the personal scope. The full catalogue is Section 6.

FM-4.1 · The Merged Rows. Treating channel 7 (acoustic -> mechanotransduction, real and non-specific) and channel 8 (acoustic -> sequence-specific content, no mechanism) as one channel. *Signature:* a sentence beginning with PIEZO1 and ending with a specific narrative outcome. *Measurement:* ask which row of §4.2 licenses the conclusion; if row 8, it is unlicensed. *Fix:* split the sentence and see whether the second half still has a source. *See:* MF-03.

FM-4.2 · The Reported Ratio. Producing a numeric `D/P` or `Delta` from the ordinal instrument. *Signature:* a decimal where §4.3 permits one of three words. *Measurement:* trace the number to a scale with a true zero; there isn't one. *Fix:* delete the number, keep the ordering, cost the Rasch route (123 items per term) if the number is genuinely needed. *See:* MF-08.

FM-4.3 · The Ten-Observation Verdict. Concluding "it works" or "it doesn't" from an `n` whose 95% half-width exceeds the SESOI. *Signature:* a confident conclusion with no interval. *Measurement:* §4.9's resolution column. *Fix:* classify N-1 and say nothing further. *See:* MF-17.

FM-4.4 · The Un-nulled Arm. Running the L3 arm without writing down what a null looks like. *Signature:* every outcome is consistent with the practice working. *Measurement:* before data arrive, ask for a result that would make the practitioner stop; if none exists, the arm is unfalsifiable. *Fix:* A.6 and A.16 in order. *See:* MF-18, MF-21.

FM-4.5 · The Grade Drift. A step stamped Grade R at A.14 described three months later as a technique. *Signature:* "works" attached to an arm stamped "research procedure". *Measurement:* re-read A.14 aloud. *Fix:* restore the stamp; it travels with the step wherever the step is described. *See:* MF-24.

FM-4.6 · The Front-Loaded Evaluation. Judging the practice before one time constant; at `delta = 0.08` per week, `tau = 11.99` weeks, so an evaluation at week 4 is reading a transient. *Signature:* "nothing is happening" at week 3. *Measurement:* elapsed periods against `-1/ln(1-delta)`. *Fix:* A.12. *See:* the contraction theorem [Curzi, 2026].

FM-4.7 · The Silent Decline. A protocol involving another person in which "no" has never once been said. *Signature:* perfect compliance. *Measurement:* count declines; a decline rate of exactly zero over many opportunities is a measurement of the decline path, not of the protocol. *Fix:* A.19, and treat the zero as a fault in the gate. *See:* entries MF-27 and MF-28 of Section 6.

FM-4.8 · The Borrowed Effect. Attributing to the L3 arm an outcome L1 and L2 fully explain. *Signature:* three things changed and the credit went to the newest. *Measurement:* the arm must vary independently of the others or the design attributes nothing. *Fix:* factorial or staggered introduction; if unaffordable, attribute nothing. *See:* entry MF-26 of Section 6.

37.5 5 The protocol at group scope

Section 4 gave one person a practice they can run and evaluate. This section gives the same thing to a group, and it is not the same protocol with the pronouns changed.

Three things are genuinely different at group scope, and each of them adds a gate rather than a paragraph. **First, a group can be *unable* to build**, in a sense a person cannot: the storage law's headroom gauge can be negative, and when it is, no arrangement of the members rescues the situation. **Second, a group can destroy its own capability by succeeding at alignment** — which is the independence requirement [Curzi, 2026g] and is arithmetic, not sentiment. **Third, a group can manufacture the appearance of consent**, which a person cannot do to themselves, and which makes the consent architecture of Section 4 necessary but not sufficient here.

The section therefore runs in a fixed order: gate, geometry, seeding, and only then anything new. That order is not stylistic. **Reversing it — building first and measuring afterwards — is the single most common way a group protocol destroys stored order while reporting progress.**

37.5.1 5.1 Gate 0, worked

An eighteen-person unit. `N = 18`, `ln N = 2.8904`. Parameters estimated per the survey protocol [Curzi, 2026i]: `eta = 0.30`, `delta = 0.06` per period, `C = 0.55` per period.

Unit state	`Q`	`C`	`h = eta ln N − delta Q − C`	Verdict
A — young unit, little accumulated	1.20	0.55	**+0.2451**	Build permitted, `r <= 0.2451`
B — mature unit, large accumulated `Q`	6.90	0.55	**-0.0969**	**MAY NOT BUILD**
C — young unit carrying a heavy programme	4.00	1.05	**-0.4229**	**MAY NOT BUILD**

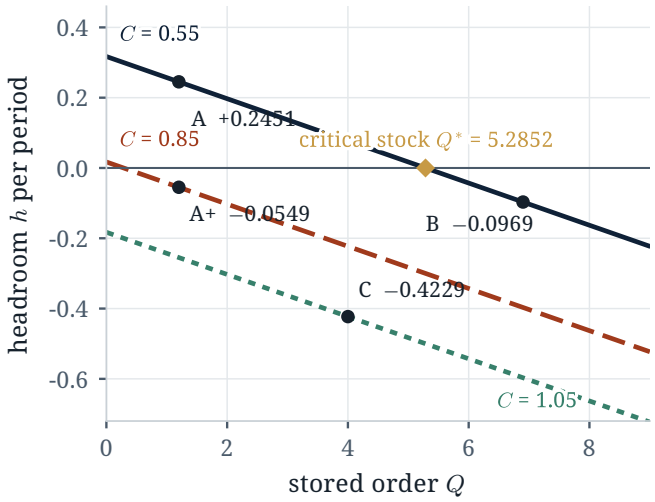


Figure 37.1: **A unit can pass the gate and a slightly cheaper unit can fail it: headroom is linear in stored order, so the same cost line closes at a computable critical stock.** Unit A (`Q` = 1.20, `C` = 0.55) reads +0.2451; unit B (`Q` = 6.90) reads −0.0969; unit C (`Q` = 4.00, `C` = 1.05) reads −0.4229; and unit A after adopting a programme costing 0.30 reads −0.0549, having moved from feasible to infeasible before the programme did anything. Modelled at `eta` = 0.30, `ln 18` = 2.8904, `delta` = 0.06 from the parameters of Section 5.1; generated by `figs/src/CCP-037_fig01.py`.

Unit state	`Q`	`C`	`h = eta ln N − delta Q − C`	Verdict
A + a proposed programme costing 0.30	1.20	0.85	**-0.0549**	**MAY NOT BUILD**

The fourth row is the one that decides most real cases and it is the one nobody computes. Unit A has positive headroom of 0.2451 and a proposed programme whose maintenance cost is 0.30. The programme's cost exceeds the headroom, so adopting it moves the unit from feasible to infeasible *before the programme has done anything at all.* The correct decision is not to run a pilot and see. It is to shed 0.06 or more of existing `C` first, or to redesign the programme to cost less than 0.2451, and only then to run it.

The critical stock — the `Q` at which headroom reaches zero with everything else fixed — is

```
Q* = (eta*ln N − C) / delta = (0.30*2.8904 − 0.55) / 0.06 = 5.2852
```

Unit B is above its critical stock. It is not a failing unit; it is a successful one, and that is precisely why it cannot build. Accumulated order costs `delta·Q` per period to hold, and at `Q = 6.90` that maintenance draw is 0.4140 per period against a maximum supply of `eta ln N = 0.8671`, leaving less than the 0.55 the unit already spends.

37.5.2 5.2 The fixed-point identity, and why maturity is not the enemy

CC-R6, verified to six decimal places: at the fixed point of the plain linear model,

```
h_inf = eta * ( ln N − F(u*) )
```

`delta`, `C` and `Q` all cancel. The headroom of a mature group depends on exactly one thing: the gap between its ceiling `ln N` and the surplus `F(u*)` it actually collects. The consequence matters: **the plain linear model never ossifies by itself.** Ossification requires exogenous accretion of `Q` — stock arriving from outside the interaction, such as inherited obligations, imported compliance regimes, acquired estates — or superlinear cost growth.

So when a group finds itself at Unit B, the diagnosis has two branches, and they have different fixes. If `Q` grew from the group's own interaction, the equilibrium is self-limiting and the fix is on `u` and `C`. **If `Q` grew from outside — a merger, a regulation, an inherited mandate — the group is carrying stock it never converted and the fix is the fourth lever, deliberate shedding [Curzi, 2026a]: shed it, deliberately, with the ethical line intact.**

37.5.3 5.3 The Jensen correction, and why it always errs optimistically

CC-R2, from the effective-count audit [Curzi, 2026c]: because `F` is strictly concave (the strict-concavity result [Curzi, 2026]), feasibility must be tested on the **mean of `F(U_t)`**, not on `F` of the mean `U`. Jensen's inequality makes the error one-directional: `F(mean u) >= mean F(u)`, always. **The wrong computation always overstates the group's surplus, never understates it.**

Worked on eight periods of an eighteen-person unit with measured `u` values 0.10, 0.22, 0.31, 0.44, 0.58, 0.19, 0.37, 0.26:

```
mean u        = 0.3088
F(mean u)     = 1.8324 nats      <- the wrong number
mean F(u)     = 1.7533 nats      <- the right number
gap           = 0.0790 nats      = 4.31% of the reported figure
```

A 4.31% overstatement sounds tolerable until it is passed through the storage law. With `eta = 0.30` and `delta = 0.06`, the two figures give steady states differing by

```
(eta * gap) / delta = (0.30 * 0.0790) / 0.06 = 0.3952 Q-units
```

The difference does not depend on `C` at all; it is the concavity error alone.

The steady state is overstated by 0.3952 Q-units — and in a unit whose whole accumulated stock is around 1.2, that is a third of the institution, conjured by evaluating a concave function at the wrong point. Propagate the correction to every feasibility test in the protocol below.

37.5.4 5.4 `N_eff` is the participation ratio, and turn-taking is how it is read

CC-R3, derived and verified to 1e-12 in the effective-count audit [Curzi, 2026c]: the mean-pairwise `u` aggregator equals the Rényi-2 participation-ratio effective count,

```
N_eff = N + 1 - N / PR ,      PR = 1 / sum_i w_i^2
```

This is what makes Woolley's turn-taking finding computable rather than merely quotable. Take a six-person meeting and measure the share `w_i` of speaking time each member takes:

Distribution	shares	`PR`	`N_eff`
Even	1/6 each	6.0000	6.0000
Skewed	0.45, 0.25, 0.15, 0.08, 0.05, 0.02	3.3693	5.2192

A meeting in which one person takes 45% of the air has an effective count of 5.22 rather than 6.00 — it has silently become a five-person meeting with a spectator. The measurement

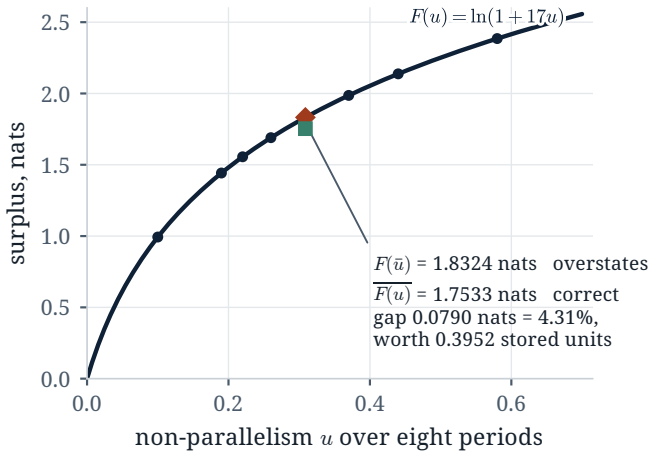


Figure 37.2: **Evaluating a concave function at the mean overstates feasibility by 4.31 per cent, worth 0.3952 stored units — about a third of the worked unit's entire accumulated stock.** Eight measured periods at `N` = 18 give a mean `u` of 0.3088; `F` of that mean is 1.8324 nats while the mean of `F` is 1.7533 nats, and Jensen's inequality makes the error one-directional. Modelled from the eight stipulated `u` values of Section 5.3; generated by `figs/src/CCP-037_fig03.py`.

is cheap: a stopwatch and six columns. **It is the single cheapest instrument in this section and it reads directly onto a quantity the surplus law uses.**

Two honesty notes. Speaking time is a proxy for participation, not participation itself; a member who says one decisive sentence has participated more than one who filled ten minutes. And the Woolley result is a correlation between turn-taking equality and `c`, not a demonstration that equalizing turns raises `c` — the direction of causation is not settled by the cited work and this section does not settle it either.

37.5.5 5.5 The price of the merge

`N = 18`, ceiling `ln 18 = 2.8904` nats:

`u`	`N_eff`	`S = ln N_eff`	Fraction of ceiling
0.00	1.0000	0.0000	0.00%
0.05	1.8500	0.6152	21.28%
0.10	2.7000	0.9933	34.36%
0.25	5.2500	1.6582	57.37%
0.40	7.8000	2.0541	71.07%
0.55	10.3500	2.3370	80.85%
0.70	12.9000	2.5572	88.47%
1.00	18.0000	2.8904	100.00%

Eighteen people who all see it the same way purchase zero surplus. And concavity is front-loaded: `u = 0` to `u = 0.25` buys 57% of the entire available range. **The first honest dissenter in a converged room is the highest-value intervention available in that room, and the eighteenth is worth a small fraction of the first.**

Concavity is merciful in one direction and unforgiving in the other. It means a cast of thousands is not needed. It also means that adding members who are *not* different adds nothing to `S` while adding coordination cost — and coordination paths grow as `n(n-1)/2`: 15 at six

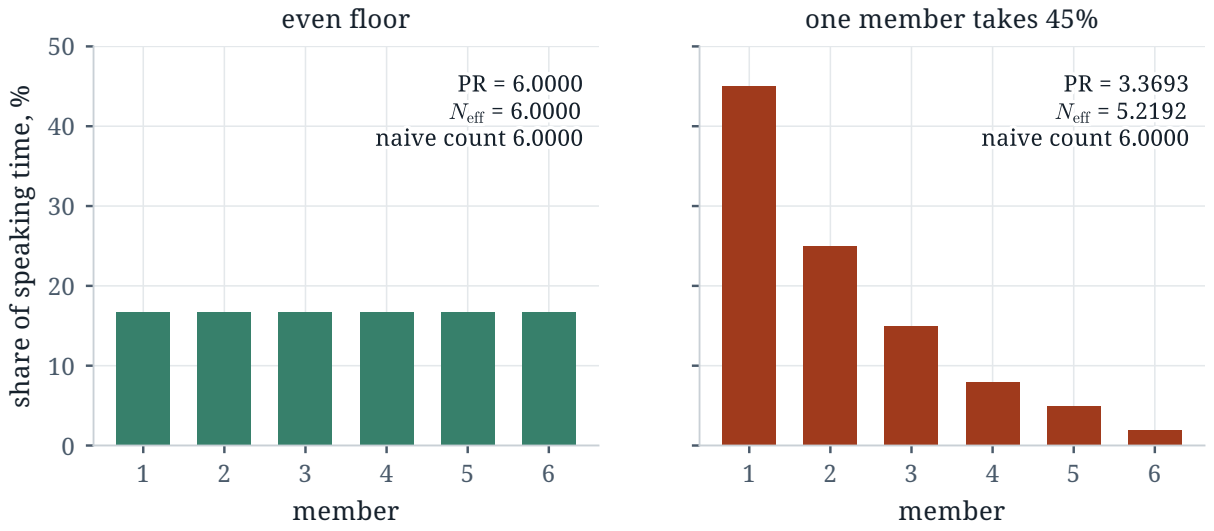


Figure 37.3: **A six-person meeting in which one member takes 45 per cent of the speaking time has an effective count of 5.2192, not 6: it has silently become a five-person meeting with a spectator.** The participation ratio falls from 6.0000 to 3.3693 and the effective count follows through the exact identity `N_eff = N + 1 − N/PR`. The instrument is a stopwatch and six columns; the share values plotted are the illustrative construction of Section 5.4. Generated by `figs/src/CCP-037_fig05.py`.

people, 66 at twelve, **153 at eighteen**, 435 at thirty [Brooks, 1975]. **Adding a parallel member is a pure cost with no offsetting surplus, and the arithmetic says so before anybody has to have an opinion about the candidate.** The sharpest instance in this corpus is the seventh-discipline case of Curzi [2026d], where adding a discipline to an institute makes it *worse* — the period change falls to `−0.002863` and the fixed point drops from 6.265 to 5.964 — because divergence, not domain count, was the deficiency.

37.5.6 5.6 Condorcet, with the independence clause priced

Condorcet's jury theorem (1785; not in this corpus's external bibliography, see Limitation L-14): with `n` independent voters each correct with probability `p > 0.5` on a binary question, majority accuracy rises with `n` and tends to 1. At `p = 0.60`:

`n` (independent)	1	3	5	9	15	21	45
P(majority correct)	0.6000	0.6480	0.6826	0.7334	0.7869	0.8256	0.9135

Now price the loss of independence by reading `N_eff` as the count of genuinely independent voices in a nine-person panel:

`u`	`N_eff`	nearest odd independent count	P(majority correct)
1.00	9.00	9	0.7334
0.50	5.00	5	0.6826
0.25	3.00	3	0.6480
0.10	1.80	1	0.6000
0.00	1.00	1	0.6000

At `u = 0.10` a nine-person panel is exactly as accurate as one person, and it costs nine times as much and reports nine times the confidence. That last clause is the whole hazard.

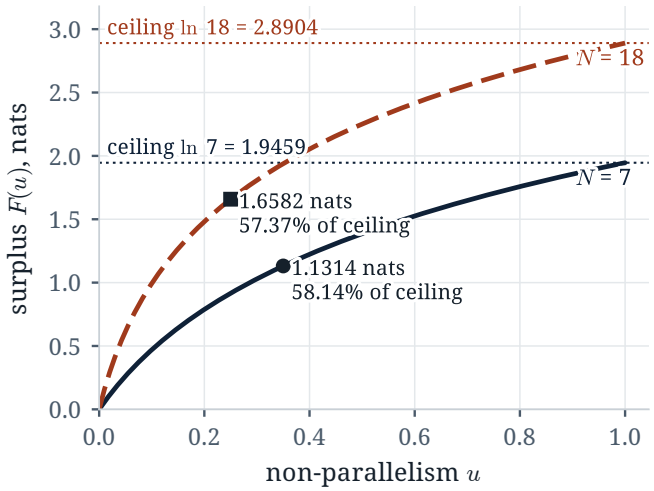


Figure 37.4: **The first genuinely different framing is worth more than all the later ones combined: moving `u` from 0 to 0.25 buys 47.09 per cent of the whole available range at `N` = 7 and 57.37 per cent at `N` = 18.** The two worked rosters sit at 1.1314 nats (58.14 per cent of `ln 7` = 1.9459) and 1.6582 nats (57.37 per cent of `ln 18` = 2.8904). Modelled from the tables of Sections 4.4 and 5.5; generated by `figs/src/CCP-037_fig02.py`.

The panel does not return "0.60 with low resolution"; it returns unanimity, and unanimity reads as certainty. **The correlated panel is the most expensive way yet devised to obtain one opinion.**

Badged: the mapping from `N_eff` to an integer count of independent voters is a **modelling convenience, not a theorem** — Condorcet's setup assumes independence outright, and `N_eff` is the surplus law's audit of redundancy, not a formal count of independent Bernoulli trials. The two are stitched here because the stitch is informative about direction and order of magnitude, and that is all it is claimed to be.

37.5.7 5.7 Critical mass, and the design move that actually pays

Centola et al. [2018] measured a tipping point at a committed minority of about **25%** in controlled online experiments. **The value of the result is not the specific quarter; it is that a tipping point exists at all — that a minority below the critical size is absorbed with no residue and one above it flips the population.** Before 2018 that was a theoretical prediction; after it, a measured phenomenon in at least one carefully constructed setting.

The honest note travels with the number every time: this figure comes from controlled online experiments with confederates committed by construction, and its transfer to a field organisation is a modelling assumption, not a measurement. Use 25% as an order-of-magnitude floor and never as a guarantee.

Scope	population	25%	committed people needed
whole firm	400	100.00	100
division	120	30.00	30
department	60	15.00	15
unit	30	7.50	8
this section's unit	**18**	**4.50**	**5**
team	12	3.00	3
squad	6	1.50	2

SHRINK THE DENOMINATOR. The threshold is a proportion, so the cheapest way over it is not to recruit more committed people but to narrow the population within which the convention must hold. One hundred committed people is a political campaign with a budget and eighteen months. Two committed people is a pair. **Sequential local tipping is cheaper than simultaneous global tipping by a wide margin, and it produces clusters as a side effect rather than requiring them as a precondition.**

37.5.8 5.8 Cluster, do not spray — computed

Complex contagion needs multiple reinforcing contacts (Centola and Macy, 2007, Centola, 2010), which means the seeding plan that maximises reach is the wrong plan. Take a 120-person division, a budget of 16 people's time, mean degree `k = 8`, and a 2-of-8 reinforcement rule.

PLAN A — SPRAY. Distribute the 16 across the division. Seed fraction `s = 16/120 = 0.1333`.

```
P(a non-seed has >= 1 seeded neighbour)  = 0.6817   <- the number that gets reported
P(exactly 1 seeded neighbour)            = 0.3917   <- 40.7 people, the worst cell
P(>= 2 seeded neighbours)                = 0.2900   <- the number that matters
Expected non-seeds within reinforcing range = 104 x 0.2900 = 30.16 people
```

PLAN B — CLUSTER. Put all 16 into one 20-person unit. Local seed fraction `16/20 = 0.80`.

```
P(>= 2 seeded neighbours) for the 4 non-seeds = 0.9999
Expected activatable inside the unit          = 4.00 of 4
Local committed fraction                      = 80.0%, against a 25% floor
```

Plan A's headline is 30.16 people in reinforcing range against Plan B's four. Any one-period comparison prefers Plan A by a factor of seven and a half. **Plan A is still the wrong choice.**

First, Plan A's 30 are at the exact margin — two contacts under a rule requiring two — so any attrition among the 16 drops them below threshold. *Second*, **40.7 people have exactly one seeded neighbour, which is the worst-populated cell in the table**: they have met the practice as "one person on my team is doing an odd new thing," and a below-threshold exposure is not a free trial. It spends the practice's novelty and returns a null, so the second attempt starts from a worse position than the first. *Third*, Plan B's unit is past its own tipping point at 80%, and the convention there survives the sponsor losing interest, which is the only test that matters.

Plan A optimises the first-period metric; Plan B optimises the order parameter. The first is what gets reported and the second is what determines the outcome.

Seeding order, and step 5 is the one people get wrong:

1. Pick the smallest population within which the practice works standing alone. If it cannot work with NOBODY outside adopting, pick a different practice.
2. Seed past that population's own tipping point. 25% is a floor, not a target.
3. Let it stabilise until it survives the removal of its sponsor. Leave for a month and record what is still running on return.
4. Seed the adjacent population using members of the first as reinforcing contacts. Do not send an evangelist; send two people who work there anyway.
5. Repeat outward along the graph of WHO ACTUALLY WORKS WITH WHOM, which is not the graph of who reports to whom \citep{conway1968committees}.

37.5.9 5.9 What "committed" must mean, and the line this framework will not cross

The 2018 experiment's minority was committed by construction: confederates instructed never to deviate. Reproducing that property in a real organisation is where a protocol of this kind could become something the framework refuses to be.

Two mechanisms produce apparent commitment and they are not equivalent. In the first, a person adopted the practice because it works for them and would keep doing it if the programme disappeared. In the second, a person adopted it because dissent is expensive and would stop the moment the cost of dissent fell. **From outside they look identical. In the ledger they are opposites: the second contributes nothing to `S`, because it has `u = 0` with respect to the group's belief, and it inflates the apparent committed fraction, which means the seeding calculation is being run on a fictitious numerator.**

The test is operational, not moral: **remove the sponsor for one month and count what survives** (step 3 above). That test is the same test as Gate D's, run at group scope, and it is why the protocol below requires a sponsor-absence interval before any expansion.

37.5.10 5.10 The biological layer at group scope

Everything Section 4.7 said about the individual biological arm applies here, and two things are added by the group setting — both of them restrictions.

What may be done. The same artefact: text the group wrote, encoded base-4 (Church, Gao & Kosuri, *Science*, 2012, is the reference for base-4 encoding of arbitrary data as sequence), given a gRNA-format routing prefix as an *address format* and not a molecule (the addressing-without-editing result [Curzi, 2026y]), rendered to audio by the disclosed sonification chain [Curzi, 2026y] within the Gabor constraint `dt·df >= 1/(4·pi) = 0.079577`. Audible, disclosed, inspectable. **No subliminal layer, no masked carrier, nothing below threshold** (the disclosed-narrative constraint [Curzi, 2026ac], and Gate K enforces it structurally).

What is claimed. That the encoding and rendering are reproducible, that people listened, and that the group pre-registered a prediction. **Nothing else.**

What is not claimed, in the same breath. No sequence is changed, nothing is delivered, no Cas protein is present, **nothing is edited or altered**, no capability is conferred on any person or group, and no clinical, medical or therapeutic effect is claimed, observed, or implied. And decisively: **there is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal** (the open problem of the missing demodulator [Curzi, 2026y]), so the sequence specificity of the artefact has no demonstrated route to any biological effect at all. The step is Grade R in Section 4's scheme — a research procedure with an expected null, not a technique.

The two group-specific restrictions.

RESTRICTION 1 — a shared acoustic environment is a shared input, and shared inputs lower `u`. Eighteen people listening to the same artefact in the same room at the same time are, for the duration, receiving one input rather than eighteen. Any downstream judgement they make about it is correlated by construction. **The design consequence is mandatory: individual judgements about the artefact are recorded BEFORE any group discussion of it, or the group has destroyed the independence it needs to evaluate the thing it is evaluating.** This is Gate I applied to the group's own experiment.

RESTRICTION 2 — a group cannot consent on a member's behalf, and a majority vote is not consent. Every participant passes their own Gate K, their own Gate D, and their own Gate V. **A member who declines is recorded as declining — as data, in their own row — and is not persuaded, not followed up, and not counted as a missing value.** A group-scope protocol in which the decline rate is zero has measured its decline path and found it broken.

And the acoustic-to-c proposal, per CC-R7 and §3.9: **proposed, not established; the causal test of the acoustic intervention analysis [Curzi, 2026f] is reported as not confirmed.** A group may run it as an experiment with a pre-registered null. A group may not describe it as something the room does.

37.5.11 5.11 The group grounding diagnostic

the grounding protocol [Curzi, 2026ac]'s six stages, run by a group rather than a person, with one addition at the front.

```
G-0  RECORD INDEPENDENTLY FIRST   Each member writes their own answer to G-1..G-4
                                  BEFORE any discussion. Then pool. (Gate I.)
G-1  IDENTIFY THE CHANNEL         Name it, with two ends.
G-2  CHECK IT IS POSSIBLE         Any veto ends the matter.
G-3  ESTIMATE ITS GAIN            Order of magnitude.
G-4  PREDICT A MAGNITUDE          Written down, with units and an interval, first.
G-5  MEASURE                      Pre-registered, four baselines, plus the ceiling test.
G-6  COMPARE                      Report whichever way it comes out.
```

G-0 is the whole difference between a group diagnostic and a group conversation. Without it, stage G-3's gain estimate is one person's estimate with seventeen echoes, and the spread of the pooled estimates — which is the most informative thing a group can produce at G-3 — has been destroyed before it existed.

Run on the three claims this section has to handle:

Claim	Halts at	Verdict
Equalizing turn-taking raises this unit's decision accuracy	G-3, with a wide pooled interval	**Eligible.** Channel real (independence -> aggregation), gain uncertain, magnitude predictable. Run it
Changing the room's acoustics raises `c`	G-1 passes, G-2 passes, **G-5 already run and not confirmed** (Ch 14 §14.13)	**PROPOSED, NOT ESTABLISHED.** May be re-run as an experiment; may not be reported as a capability
The encoded artefact activates a latent capability in members' cells	**G-1**	**REJECTED IN PRINT.** The claimed channel cannot be named because no such mechanism is established (the open problem of the missing demodulator [Curzi, 2026y])

A diagnostic that never rejects anything is not a diagnostic, and the rejection above is of this volume's own subject matter rather than of something the corpus has no stake in.

37.5.12 5.12 Falsification, and why the shuffle is not enough

the falsification suite [Curzi, 2026i]'s suite, mandatory for any group claim: preregister, fit on training, **freeze**, report held-out, and beat four baselines — persistence, direct-input, matched autoregressive, and the **shuffled-input control**.

CC-R4 is the reason a fifth test is required. In the falsification suite [Curzi, 2026u], a filtered repackaging of the input passed **all four baselines including the shuffle at p = 0.0005**, and was caught only by the structural ceiling test: its single-step jump was **38.5 times** the theoretical maximum `r_max = eta ln N - delta Q - C`.

```
Unit C parameters: eta=0.30, ln N=2.8904, delta=0.06, Q=4.00, C=0.55
r_max  = 0.30*2.8904 - 0.06*4.00 - 0.55 = 0.0771 per period
38.5 x r_max = 2.9688 per period               <- the claimed jump
```

Rule, carried in place: the shuffle is an empirical audit; only the structural tests detect repackaging. Any single-period increment a group claims must satisfy `r_t <= h_t`. A claimed increment larger than the ceiling is not an impressive result; it is a measurement of the measurement (entries MF-11 and MF-12 of Section 6).

37.5.13 5.13 Failure modes specific to the group scope

FM-5.1 · The Build Under a Negative Gauge. Launching a programme while `h < 0`. *Signature:* visible activity, rising reported engagement, falling `Q`. *Measurement:* compute `h` including the programme's own `C`. *Fix:* B.3, then B.16. *See:* the storage law [Curzi, 2026a], FM-6.1.

FM-5.2 · The Correlated Panel. A group whose members' errors are the same errors. *Signature:* fast unanimity; nobody can recall the group changing its mind because of an internal objection. *Measurement:* independent-estimate-before-discussion, `u = 1 − rho^2`. *Price:* at `u = 0.10`, nine people are as accurate as one and nine times as confident. *Fix:* B.5, B.9. *See:* the independence requirement [Curzi, 2026g], FM-22.6.

FM-5.3 · Alignment Reported as Capability. Treating a rise in agreement as a rise in group ability. *Signature:* the programme's success metric is convergence. *Measurement:* `u` before and after; if it fell, the metric measured the damage. *Fix:* B.17's reversal rule. *See:* Section 6, MF-22.

FM-5.4 · Manufactured Commitment. Counting compliance as commitment in the seeding denominator. *Signature:* a decline rate of exactly zero. *Measurement:* remove the sponsor for a period and count survivors. *Fix:* B.12. *See:* the two-mechanisms-of-commitment analysis [Curzi, 2026g]; Section 6, MF-28.

FM-5.5 · The Spray. Seeding for reach under a complex-contagion rule. *Signature:* a reported "reached N people" figure with no reinforcement count. *Price:* in §5.8, 40.7 people at exactly one contact — novelty spent, nothing activated. *Fix:* B.11.

FM-5.6 · The Reported Group Ratio. Publishing a numeric doctrine-to-data ratio or a `Delta`. *Signature:* a decimal where B.13 permits one of three words. *Measurement:* trace it to a scale with a true zero; there isn't one. *Fix:* delete the number, keep the ordering, budget the Rasch route. *See:* Section 6, MF-08.

FM-5.7 · The Acoustic Capability Claim. Describing an acoustic intervention as raising `c`. *Signature:* the room is credited with a group's performance. *Measurement:* the acoustic intervention analysis [Curzi, 2026f]'s causal test, **reported as not confirmed**. *Fix:* relabel as proposed, or run the test. *See:* CC-R7; Section 6, MF-24.

FM-5.8 · Consent by Majority. A vote standing in for individual consent. *Signature:* a show of hands preceding a step that acts on individuals. *Measurement:* count the per-person Gate K, D and V records; if there are fewer than `N`, consent was not obtained. *Fix:* B.14, B.4. *See:* Section 6, MF-27.

37.6 6 A finite catalogue of manifestation faults

the built-system failure catalogue [Curzi, 2026g] catalogued the ways a built system fails. This section catalogues the ways a *claim about manifestation* fails — and it is the section this strand of the corpus most needs, because the biological-layer strand of the corpus is the material most easily misread as licence.

The postulate at the centre of the volume is conditional: **belief creates reality if the belief is grounded in truth**, which formalizes as `BELIEF -> [CHANNEL] -> WORLD`, with "grounded in truth" meaning that **the channel exists and is real** (the channel-gain form of the belief postulate

[Curzi, 2026ab]). Strip the conditional and the postulate becomes magical thinking. Keep it, and every failure of manifestation becomes a diagnosable fault in a specific part of a specific claim.

This is a **reference catalogue**, not an argument. It is written to be consulted against a claim that has just failed, or — better — one that has not yet been made. Thirty named failure modes appear below in the fixed eight-field entry format of the built-system failure catalogue [Curzi, 2026g], so the section functions as a lookup table: arrive with a symptom, leave with a coordinate, a measurement, and a fix.

The armature is modified accordingly, as the built-system failure catalogue [Curzi, 2026g]'s was. `The Law` states what makes a finite catalogue possible here. `The Engineering` **is** the catalogue. `The Specification` is the procedure for using it. `The Failure Modes` names the ways diagnosis itself goes wrong. The `Worked Example` carries one diagnosis end to end.

Four of these entries are faults of this framework's own method rather than of anybody else's, and they are not softened on that account. MF-08 is the failure of reporting a magnitude the corpus's own validation protocol has demoted. MF-11 is the repackaging that the corpus's own falsification suite could not catch. MF-23 — precedent mistaken for evidence — is, by the corpus's biological-layer constraints's own reckoning, **the single most likely failure of this strand of the corpus's method.** MF-24 is the speculative bridge that lost its badge somewhere between the section that stated it and the summary that quoted it. A catalogue that indexes only other people's errors is advertising.

37.6.1 6.1 Why a finite catalogue is possible

A manifestation claim, stated in this framework's terms, is a claim that some belief `B`, held and acted on, produces some change `X` in the world. **Every such claim has exactly eight coordinates, and there are therefore only eight kinds of failure, however many names they wear.**

```
CH   CHANNEL          Does a real coupling from B to X exist and can it be named?
G    GAIN             How much X per unit B does that coupling deliver?
Mg   MAGNITUDE        Is the predicted size of X right to within an order of magnitude?
B    BASE RATE        What fraction of attempts produce X, counting the unobserved ones?
Fa   FALSIFIABILITY   Is there a stated result that would kill the claim?
Cn   CONSENT          Can the person acted on know, decline, and reverse?
Ld   LEDGER           Is precedent kept distinct from evidence, and speculation badged?
Sc   SCALE            Is each quantity measured on a scale permitting the operation used?
```

The thirty names below are not thirty independent phenomena. They are recurring *routes* into eight coordinate faults, distinguished because the routes have different signatures, different instruments, and different fixes. The class headings in `The Engineering` group them by coordinate.

Badged, as the corpus's conventions require: that a manifestation claim decomposes into exactly these eight is the author's framing, not a theorem. It is defensible — each coordinate corresponds to a step of the grounding diagnostic (the grounding protocol [Curzi, 2026ac]) or to a gate of the Formula Validation Protocol (the corpus's internal-scope conventions) — but it is a taxonomy, and taxonomies are chosen. Where a fault sits outside the scheme, the catalogue will not name it, and the correct response is to say so rather than force a coordinate.

37.6.2 6.2 The four diagnostic identities

Every entry is priced against one of four. They are the only arithmetic the catalogue needs.

Identity 1 — the coupling identity. `X = G · B`, where `G` is the channel's gain. **If `G = 0` then `X = 0` for every value of `B` whatsoever.** No amount of belief compensates for an absent channel, and this single line disposes of about a third of the entries below.

Identity 2 — the grounding sequence (the grounding protocol [Curzi, 2026ac], carried in place): identify the channel, check it is possible, estimate its gain, **predict a magnitude**, measure, compare. **A claim's fault is located at the first stage it cannot pass**, and locating it there rather than at the end is what makes the catalogue usable.

Identity 3 — the structural ceiling. the structural-ceiling corollary [Curzi, 2026a]: `r_max = eta ln N − delta Q − C`. Any claimed single-period increment in stored order greater than `r_max` is impossible under the model, and CC-R4 records the case that makes this indispensable: **a filtered repackaging of the input passed all four empirical baselines including the shuffled-input control at p = 0.0005, and was caught only by the ceiling test, its jump being 38.5 times the theoretical maximum.**

Identity 4 — the ledger rule (the corpus's biological-layer constraints). `(+1)` carries **evidence**: what is measured, replicated, falsifiable. `(−1)` carries **precedent**: that the idea was arrived at independently and preserved. **Two sources agreeing does not double the evidence when only one of them is evidence.** What independent arrival establishes is genuine and worth having — that the structure was findable by more than one route — and it is not a second measurement.

37.6.3 6.3 The honesty clauses that constrain every entry

Repeated here because a catalogue is the part of a book most likely to be quoted without context.

- **No clinical, medical or therapeutic claim appears in this section**, in any entry, in any direction. Nothing here diagnoses a person or recommends a treatment.
- **No claim that DNA is edited, altered or damaged.** The architecture of the disclosed sonification chain [Curzi, 2026y] cannot do so, and nothing in this catalogue implies otherwise.
- **No claim that any capability in this corpus has been conferred on anyone.**
- **There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal** (the open problem of the missing demodulator [Curzi, 2026y]). Several entries turn on this and each states it.
- **Casimir does nothing at biological scale.** Where the Casimir effect appears as a price, the computed number is the point (§MF-06).
- Where an entry cites an external result, the result is real and the *transfer* to the claim under diagnosis is a modelling assumption, badged each time.

Entry format. Eight fields, same order, every entry, so the section can be scanned in a column.

Field	What it holds
Damages	which of CH / G / Mg / B / Fa / Cn / Ld / Sc the fault attacks
Clinical	what the fault is, stated without metaphor
Signature	what is observable from outside, before measuring anything
Measurement	the specific instrument that confirms or excludes it
Price	the arithmetic cost, computed

Field	What it holds
Fix	the intervention, and which coordinate it repairs
Confusable with	the entries most often mistaken for this one
See	the section that develops it

37.6.4 6.4 Class I — channel faults (CH, G)

The foundational class. Every entry here is a case of Identity 1 with `G = 0`, arrived at by a different route.

MF-01 · The Ungrounded Belief

Damages: CH. No channel is identified at all, so `G` is undefined and Identity 1 cannot even be written. **Clinical:** A belief is held and acted upon with no account of how it could reach the world. The claim is not that a particular mechanism operates; it is that *no mechanism is required*. This is the unconditional reading of the postulate, and it is exactly what the conditional clause exists to exclude. **Signature:** Asked "by what route?", the holder answers with a restatement of the belief, or with a word that names a category rather than a coupling — energy, vibration, the field, resonance, alignment. **A channel has two ends and a mechanism between them; a category has neither. Measurement:** Grounding diagnostic stage G-1 (the grounding protocol [Curzi, 2026ac]). Ask for the two ends and the mechanism. If three requests produce three restatements, the channel does not exist in the claimant's own account, whatever exists in the world. **Price:** Identity 1 with `G = 0`. `X = 0 · B = 0` for every `B`. **The entire investment — every hour, every payment, every deferred alternative — buys a product of zero and a number. Fix:** Name a channel or withdraw the claim. There is no third option. If a channel can be named, the claim graduates to MF-02's test or to a legitimate research procedure. **Confusable with:** MF-02 (a channel *is* named and turns out not to exist — a more advanced condition), MF-07 (channel fine, no magnitude offered). **See:** the grounding protocol [Curzi, 2026ac].

MF-02 · The Absent Channel

Damages: CH, and derivatively G. **Clinical:** A channel is named specifically, which is progress, and it does not exist. The claim is falsifiable and false. **This is the honourable failure of the class**, because it can be checked, and checking it is cheap relative to the alternative. **Signature:** The named channel is precise enough to look up, and the lookup returns nothing, or returns a mechanism that does something adjacent but not this. **Measurement:** Stage G-2. Search the primary literature for the named coupling, not for its vocabulary. **The distinguishing test: does the cited work measure the coupling, or measure something upstream of it and speculate? Price:** Same as MF-01 in outcome and cheaper in diagnosis. The saving is the point: a named channel costs a literature search to refute; an unnamed one costs a career. **Fix:** Withdraw, or reformulate as the research procedure the gap actually calls for. the open problem of the missing demodulator [Curzi, 2026y]'s treatment of the sequence-specific acoustic demodulation gap is the model: **state the gap sharply, say what experiment would settle it, and stop. Confusable with:** MF-01, MF-03 (a *different* real channel is doing the work), MF-05. **See:** the missing demodulator and the grounding protocol [Curzi, 2026y,ac].

MF-03 · Channel Substitution

Damages: CH — credit assignment. The gain is real and belongs to a different coupling. **Clinical:** An outcome produced by an established channel is attributed to an unestablished one that was present at the same time. **This is the most consequential error available in Volume Three**, because the corpus deliberately runs an established channel and an open problem side by side. **Signature:** A sentence that begins with a Nobel-recognized mechanism and ends with a specific, addressed, sequence-dependent outcome. The two halves have different evidential status and the sentence hides the seam. **Measurement:** Split the sentence at the seam and ask whether the second half still has a source. Concretely, for this strand of the corpus: **PIEZO1/2 are mechanically gated ion channels and sound is mechanical, so a pressure wave becoming a cellular signal is established and NON-SPECIFIC. That a pressure wave carries WHICH signal, addressed to a particular sequence, has no established mechanism at all** (the open problem of the missing demodulator [Curzi, 2026y]). Row 7 of the personal channel inventory licenses the first; row 8 does not license the second. **Price:** The claim inherits the credibility of the established channel and none of its content. In a factorial design the price is measurable: an arm that cannot be varied independently of the established channel **attributes nothing, at any sample size** — the design has zero identifying power, not low power. **Fix:** Vary the two independently, or attribute nothing. Where independent variation is unaffordable, say the outcome is unattributed. That is a finding. **Confusable with:** MF-02, MF-26 (a bundle of nine components with credit to one), MF-05. **See:** the cell-as-receiver result and the missing demodulator [Curzi, 2026y]; Section 4.2.

MF-04 · The Zero-Gain Channel

Damages: G. The channel exists, is real, is measured, and its gain for *this* output is zero. **Clinical:** A genuine coupling is invoked for an outcome it does not produce. The channel is not fictional; it is irrelevant. **This fault survives every check that asks "is the mechanism real?" and fails the only check that matters, which is "does it move THIS?" Signature:** The citation is real and the outcome is not in it. The claimed effect is a plausible-sounding cousin of the measured one. **Measurement:** Stage G-3. Ask for the gain in the units of the claimed output. If the cited work reports a gain for a different output, the number does not transfer, and the transfer is exactly the unbacked step. **Price:** $X = G \cdot B$ with $G = 0$ for this X. Identical to MF-01 in result and far harder to detect, because a real citation sits in front of it. **Fix:** Find the measurement of the gain the claim actually needs, or state that it has not been made. **Confusable with:** MF-03, MF-05, MF-06 (gain nonzero but mis-sized). **See:** the channel-gain form of the belief postulate [Curzi, 2026ab]; the grounding protocol [Curzi, 2026ac].

MF-05 · The Borrowed Channel

Damages: CH, G, and often Mg simultaneously. **Clinical:** A coupling established in one system is asserted for another where the conditions that made it work are absent. Scale, medium, timescale, or geometry differ, and the difference is not priced. **Signature:** "It works in cells, therefore in tissue." "It works at ultrasound frequencies, therefore in a room." "It works in a controlled online experiment, therefore in a firm." **Measurement:** Stage G-2. Compute the transfer explicitly. For acoustics into tissue, an illustrative attenuation of 0.5 dB per cm per MHz gives:

```
f = 1.0 MHz, d = 10 cm  ->   5.00 dB loss, amplitude factor 0.5623
f = 2.0 MHz, d = 10 cm  ->  10.00 dB loss, amplitude factor 0.3162
f = 5.0 MHz, d = 10 cm  ->  25.00 dB loss, amplitude factor 0.0562
```

Price: At 5 MHz through 10 cm the amplitude arriving is 5.6% of what was sent, so a claim calibrated at the source is wrong by a factor of about 18 at the target. The 25% tipping-point result of Centola et al. [2018] carries an analogous transfer cost: **it was measured in controlled online experiments with confederates committed by construction, and its application to a field organisation is a modelling assumption, not a measurement. Fix:** Price the transfer, or run the measurement in the destination system. **Borrowing is legitimate; borrowing without paying the transfer cost is the fault. Confusable with:** MF-04, MF-06, MF-10. **See:** the measured Casimir force [Casimir, 1948, Lamoreaux, 1997]; the grounding protocol [Curzi, 2026ac].

37.6.5 6.5 Class II — magnitude and scale faults (Mg, Sc)

The channel exists. The arithmetic about it does not.

MF-06 · The Magnitude Error

Damages: Mg. Direction correct, size wrong by orders of magnitude. **Clinical:** A real effect is invoked at a scale where it is negligible. The claim is not false in sign and is false in every practical sense. **This is the most respectable failure in the catalogue and one of the most common, because getting the direction right feels like getting the answer right. Signature:** The claim quotes a mechanism correctly and never states a number. Where a number appears, it is the mechanism's number rather than the situation's. **Measurement:** Stage G-4, performed before stage G-5. Compute the effect at the actual scale of the claim and compare with what else is present at that scale. **Price:** The canonical case is Casimir. `F/A = pi^2 hbar c / (240 a^4)`, Casimir, 1948, measured by Lamoreaux, 1997 and Mohideen and Roy, 1998 — entirely real:

```
10 nm  ->  1.3e5 Pa   (1.283 atm)
100 nm ->  13 Pa
1 um   ->  1.3e-3 Pa
10 um  ->  1.3e-7 Pa
1 mm   ->  1.3e-15 Pa
```

At any biological or architectural scale the Casimir force is utterly negligible. It does nothing in a cell and nothing in a room, and this framework says so wherever Casimir appears. For comparison, a 60 dB SPL sound is 0.02 Pa — **more than ten thousand times the Casimir pressure at 1 micrometre.** A second case, computed the same way: an illustrative focused-ultrasound pressure amplitude of 0.5 MPa against a 60 dB SPL room tone is a ratio of 2.5e7, or **7.40 orders of magnitude.** A claim that carries a mechanism across that gap without arithmetic has made a magnitude error even if every word of the mechanism is correct. **Fix:** Compute at the scale of the claim, always, before asserting. Then decide. **Confusable with:** MF-04, MF-07, MF-10. **See:** the measured Casimir force [Casimir, 1948, Lamoreaux, 1997].

MF-07 · Order-of-Magnitude Silence

Damages: Mg, and through it Fa. **Clinical:** No magnitude is predicted at all. The claim is directional forever. **A practice that predicts no magnitude cannot be wrong, and a practice that cannot be wrong cannot be trusted when it is right. Signature:** Every reported outcome is compatible with the claim. Success is "it helped"; failure is "it takes time". **Measurement:** Stage G-4 is simply absent from the record. Ask what number was written down before the attempt, and when. **Price:** Unfalsifiability is not free; it is paid in the inability to accumulate. **A directional claim cannot be improved, because improvement requires a residual, and a residual requires a predicted value.** Twenty years of directional practice contains exactly as much information as one year. **Fix:** Stage G-4 with an interval and units, recorded before measurement, per Section 4.9. **Confusable with:** MF-18 (goal itself unfalsifiable), MF-21 (no null defined), MF-06. **See:** the grounding protocol [Curzi, 2026ac]; Section 4.9.

MF-08 · The Demoted Magnitude Reported As Measured

Damages: Sc, then Mg. A quantity the framework's own protocol demoted is reported as a measurement. **Clinical:** The Formula Validation Protocol returns **DEMOTE on gate 2 (scale type)** for `M = (D/P)·E`, because D and P are not instrumented on ratio scales with true zeros (the scale-type analysis [Curzi, 2026t]). **What is demoted is MAGNITUDE, not STRUCTURE.** The limiting cases stand, the fixed point `D = P` stands, and the ordering over-framed / calibrated / under-framed stands, because each needs only ordinal information. **The fault is publishing a number for `D/P`, or a distortion `Delta = |ln r|`, or a quantitative `M -> Q` transfer, as though the demotion had not happened. Signature:** A decimal appears where the instrument permits only an ordering. "The perception ratio is 1.4." "Distortion reduced by 29% this quarter." **Measurement:** Trace the number to a scale with a true zero. There is not one. **Division requires a ratio scale (Stevens' typology); with ordinal data the formula is a heuristic ordering and must be demoted in print** (FVP-2). **Price:** The demotion is not permanent and the route out is costed. the Rasch-route theorem [Curzi, 2026t], given badged Postulate SO1: a **Rasch logit difference** makes `Delta = |ln(D/P)|` genuinely ratio-scaled, because Rasch separation supplies the true zero that division requires. **The cost is 123 items per term** — 246 items for the ratio, **369 items** for all three terms of the Moment. At twenty seconds per item that is **2.05 hours per administration**, and the cross-individual recovery gate (FVP-7) with thirty participants is **61.5 person-hours** before analysis begins. **That is the price of the number, and reporting the number without paying it is this entry. Fix:** Delete the magnitude. Keep the ordering, which is genuinely diagnostic. Budget the Rasch route if the number is genuinely needed. **Build the instrument so the forbidden thing is impossible rather than merely discouraged** — Section 4 §4.3's three-valued verdict set is the pattern. **Confusable with:** MF-09 (any scale violation), MF-07, MF-24. **See:** the belief-failure classes [Curzi, 2026ac]; the Formula Validation Protocol and the scale-type analysis [Curzi, 2026t]; CC-R1.

MF-09 · The Scale Violation

Damages: Sc. **Clinical:** An operation is performed that the data's scale type does not permit: ratios of ordinal scores, averages of ranks treated as magnitudes, differences on a scale with no defined unit. MF-08 is the specific case that matters most to the corpus; this entry is the general one. **Signature:** Arithmetic performed on a Likert response as though it were a length. **Measurement:** FVP gate 2 on every term. Nominal permits counting and identity. Ordinal permits ordering. Interval permits differences. **Only ratio permits multiplication and division, and only because it has a true zero. Price:** The output has no interpretation, which is worse than being wrong: a wrong number can be corrected, an uninterpretable one propagates

silently into every downstream calculation and looks like all the others. **Fix:** Demote in print to the highest operation the scale supports, and state the route to a better scale. Magnitude estimation and cross-modality matching (Stevens) genuinely produce ratio-scale measurements of subjective quantities, and they replicate. **The route exists; it is just not free. Confusable with:** MF-08, MF-10. **See:** the scale-type analysis [Curzi, 2026t]; the logarithmic perceptual register [Fechner, 1966].

MF-10 · The Dosage Fallacy

Damages: Mg. Linear expectation applied to a quantity the recipient perceives logarithmically. **Clinical:** A "moderate" perceived improvement is promised, and delivering it requires an order-of-magnitude change in the world. Or the reverse: a genuine order-of-magnitude change is delivered and reported as disappointing, because it arrived as a small perceived step. **Signature:** Effort and perceived result are compared as though the relationship were linear. "We doubled the input and it barely feels different." **Measurement:** The Lens Theorem (the disclosed-narrative constraint [Curzi, 2026ac]). If perception is logarithmic, **equal perceptual steps are equal world RATIOS, never equal world differences**:

```
perceived +1 nat   -> world x2.7183        +1 bit -> world x2      +1 log10 -> world x10
perceived +2 nats  -> world x7.3891        +2 bits -> world x4     +2 log10 -> world x100
perceived +3 nats  -> world x20.0855       +3 bits -> world x8     +3 log10 -> world x1000
```

Price: Promising three perceptual units commits the claim to a factor of about twenty in the world. Budgeting for that promise linearly under-provisions by roughly 7x. **Concavity is not pessimism; it is the lens** — the strict concavity of the surplus functional [Curzi, 2026] has the same shape as Fechner's law [Fechner, 1966], and diminishing returns is what a logarithm looks like from inside. **Fix:** Design in ratios, budget in ratios, and **never promise a linear improvement in a quantity the recipient perceives logarithmically. Confusable with:** MF-06, MF-05. **See:** the logarithmic perceptual register [Fechner, 1966]; the corpus's internal-scope conventions.

MF-11 · The Repackaged Input

Damages: Mg and Fa together, and it is the entry that most threatens this framework's own validation suite. **Clinical:** A filtered, transformed, or reweighted version of the input is presented as an output of the system. The system has learned to restate its input impressively. **This is not fraud and usually not deliberate; it is what an unconstrained fitting procedure does when it is allowed to. Signature:** Results that track the input's fine structure too closely, and a single-step jump that is very large relative to anything the model permits. **Measurement: This is the crucial entry, because CC-R4 records that the standard suite failed to catch it.** In the falsification suite [Curzi, 2026u], a filtered repackaging **passed all four baselines — persistence, direct-input, matched autoregressive, and the shuffled-input control — at p = 0.0005.** It was caught only by the structural ceiling test of the structural-ceiling corollary [Curzi, 2026a]. **Price:** The observed single-step jump was **38.5 times the theoretical maximum**. Worked in this chapter's illustrative unit:

```
eta=0.30, ln N=2.8904 (N=18), delta=0.06, Q=4.00, C=0.55
r_max = 0.30*2.8904 - 0.06*4.00 - 0.55 = 0.0771 per period
38.5 x r_max                          = 2.9688 per period      <- the claimed jump
```

Fix: Rule, and it is binding: the shuffle is an empirical audit; only the structural tests detect repackaging. Add the ceiling test `r_t <= h_t` to every falsification suite, always. A claimed increment above the ceiling is a measurement of the measurement. **Confusable with:** MF-12 (a plain ceiling violation with no repackaging), MF-16. **See:** the falsification suite [Curzi, 2026u]; the grounding protocol [Curzi, 2026ac]; CC-R4.

MF-12 · The Ceiling Violation

Damages: Mg. **Clinical:** A claimed increment exceeds what the model permits in one period, without repackaging being present. Often it is a bookkeeping error, a period-length mismatch, or an exogenous inflow of `Q` counted as production. **Signature:** A step change with no corresponding change in `eta`, `u`, `N` or `C`. **Measurement:** Compute `r_max = eta ln N − delta Q − C` for the period in question and compare. **Price:** In the same illustrative unit, `r_max = 0.0771` per period. **Any reported increment above that is impossible under the model, which means either the model is wrong for this system or the number is.** Both are findings; assuming neither is not. **Fix:** Reconcile period boundaries, then look for exogenous accretion — stock arriving from outside the interaction (a merger, an inherited mandate, an imported compliance regime), which is real `Q` and is not production. **Confusable with:** MF-11, MF-26. **See:** the storage law [Curzi, 2026a]; the structural-ceiling corollary [Curzi, 2026a].

37.6.6 6.6 Class III — base-rate and inference faults (B)

The channel may be fine. The evidence about it is not what it appears to be.

MF-13 · The Survivorship Narrative

Damages: B. The denominator is unobservable, so no rate can be computed. **Clinical:** Successes publish and failures do not. The visible record of a manifestation practice consists entirely of the cases that worked, because the cases that did not produced nobody with a reason to write them up. **Testimonial base rates are therefore not merely biased; they are UNOBTAINABLE from testimony, at any volume of testimony. Signature:** An impressive collection of individual accounts and no denominator anywhere. “Look how many people it worked for” is a numerator with the denominator deleted. **Measurement:** Ask for the count of attempts, not the count of successes, and ask who would have recorded a failure. **If no mechanism existed by which a failure could have entered the record, the record contains no information about the rate. Price:** Computed two ways.

```
Suppose 1000 people attempt and 40 succeed. True rate 4.0%.
Testimonials show 40 successes and 0 failures -> apparent rate 100%.
Inflation factor = 25.00x.

Or, with NO effect at all: if the outcome occurs spontaneously in 5% of
attempts, then 40 glowing testimonials require only 800 attempts to exist.
```

Forty testimonials is not weak evidence of an effect. Under a 5% spontaneous base rate it is exactly what a completely inert practice produces from eight hundred attempts. Fix: Register attempts prospectively, before outcomes are known, and count everyone who started. Where prospective registration is impossible, **report that the base rate is unobtainable and decline**

to estimate it. Section 4 §4.10's N-4 category exists so that expected nulls enter the record; this entry is why. **Confusable with:** MF-14 (denominator exists and is omitted), MF-15, MF-19. **See:** the concordance register [Curzi, 2026w]; the falsification suite [Curzi, 2026u].

MF-14 · The Missing Denominator

Damages: B. **Clinical:** The denominator exists and was not reported. Distinct from MF-13, where it cannot be obtained. This one is a reporting fault and is fully repairable. **Signature:** Counts without rates. "Thirty people adopted the practice." Out of how many? **Measurement:** Ask for `N` at the start of the period, not at the end. **Attrition moves people out of the denominator silently, and a practice with heavy attrition reports beautifully. Price:** In §MF-13's arithmetic, the same 40 successes are a 100% rate, a 4.0% rate, or a 5.0% spontaneous background depending entirely on a number nobody wrote down. **Fix:** Report `n_started / n_completed / n_succeeded` as three numbers, always, in that order. **Confusable with:** MF-13, MF-17. **See:** the effective-count audit [Curzi, 2026c]; the survey protocol [Curzi, 2026i].

MF-15 · Regression Mistaken For Effect

Damages: B and Mg together. **Clinical:** People enter a practice at their worst moment, because that is when they seek one. Their next measurement is better. The improvement is what an imperfectly reliable measure does whenever selection is on an extreme, and it happens with no intervention at all. **Signature:** Recruitment framed around a low point. Large early gains that flatten. The practice works best for those who were worst. **Measurement:** Compute the expected regression before attributing anything. Selecting the bottom decile of a normally distributed measure gives `E[Z | selected] = −1.7550`. With test-retest reliability `r`, the expected next-measurement improvement **with no intervention** is `(1 − r)·|E[Z]|`:

```
bottom 10% selected:   r = 0.4  ->  1.0530 SD improvement, from nothing
                       r = 0.6  ->  0.7020 SD improvement, from nothing
                       r = 0.8  ->  0.3510 SD improvement, from nothing
bottom 20% selected:   r = 0.6  ->  0.5599 SD improvement, from nothing
```

Price: At bottom-decile selection and `r = 0.6`, a practice with literally zero effect reports 0.70 SD of improvement — **larger than most effects anyone in the behavioural sciences honestly claims.** A practice would have to be actively harmful to report nothing under these conditions. **Fix:** A control group selected the same way, or a design that does not select on the outcome measure. Where neither is available, subtract the computed regression and report what remains, explicitly labelled. **Confusable with:** MF-13, MF-17, MF-26. **See:** the survey protocol [Curzi, 2026i]; the falsification suite [Curzi, 2026u].

MF-16 · The Forking Path

Damages: B, Fa. **Clinical:** Many outcomes, subgroups, or time windows are examined and the one that reached significance is reported. No dishonesty is required; the analysis decisions were made after seeing the data and each felt reasonable. **Signature:** A specific, narrow, unanticipated result. "It worked for participants who had practised for more than six weeks and were measured in the afternoon." **Measurement:** Count the comparisons that *could* have been made, not the ones reported. Familywise error at alpha = 0.05:

```
k = 1   ->  0.0500        k = 5   ->  0.2262        k = 20  ->  0.6415
k = 2   ->  0.0975        k = 10  ->  0.4013        k = 40  ->  0.8715
```

Price: With ten candidate analyses, the probability of at least one false positive is **0.4013**. With forty, **0.8715** — a practice with no effect whatsoever will produce a "significant" result more than eight times in ten. **Fix:** Preregister the analysis, freeze it, and report the held-out result (the falsification suite [Curzi, 2026i]). Report every comparison run, including the ones that failed. **Confusable with:** MF-11, MF-19. **See:** the falsification suite [Curzi, 2026u]; the Formula Validation Protocol [Curzi, 2026t].

MF-17 · The Underpowered Null Read As Disconfirmation

Damages: B, Mg. The mirror image of the rest of this class, and just as costly. **Clinical:** A study too small to detect anything returns nothing, and the nothing is reported as evidence of absence. **An underpowered design returns a null whether the effect exists or not, so its null carries no information in either direction. Signature:** "We tried it for two weeks and it did not work." No interval reported. **Measurement:** Compare the 95% interval half-width to the smallest effect size of interest fixed in advance:

```
n = 10  ->  half-width 0.6198 SD        n = 60  ->  half-width 0.2530 SD
n = 20  ->  half-width 0.4383 SD        n = 100 ->  half-width 0.1960 SD
n = 30  ->  half-width 0.3578 SD        n = 200 ->  half-width 0.1386 SD
```

Price: A ten-observation personal experiment cannot distinguish "no effect" from "an effect of two thirds of a standard deviation". **Ten observations is what most personal experimentation actually consists of, so most personal experimentation resolves nothing — and this is arithmetic, not mystery.** To establish equivalence at a SESOI of 0.35 SD requires **51 paired observations**; to detect `d = 0.35` requires **65**. **Fix:** Classify the null before interpreting it (Section 4 §4.10): uninformative, informative, against a magnitude prediction, or expected on a research arm. **Only the second licenses "it does not work at a size I care about." Confusable with:** MF-07, MF-21. **See:** the survey protocol [Curzi, 2026i]; Sections 4.9 and 4.10.

37.6.7 6.7 Class IV — falsifiability faults (Fa)

The claim cannot lose, which means it cannot win either.

MF-18 · The Unfalsifiable Goal

Damages: Fa, and through it every other coordinate, because nothing can be checked. **Clinical:** The goal is stated so that no observable state of the world contradicts it. Abundance, alignment, readiness, flow. **These are not bad things to want; they are impossible things to test, and a protocol organized around them cannot accumulate. Signature:** Success and failure are described in the same vocabulary. Absence of the outcome is explained by the same theory that predicted its presence. **Measurement:** Ask, before any attempt: **name a result that would end the practice.** If no such result exists, the goal is unfalsifiable and the practice around it is unevaluable, whatever else it is. **Price:** No residual, so no learning. **A framework that cannot fail cannot be trusted, and this manual's own answer is to name exactly how each of its claims fails** — which is why the trust term [Curzi, 2026r] exists and why the grounding

protocol [Curzi, 2026ac] catalogues where belief does not create reality at all. **Fix:** Replace the goal with an observable, a direction, a magnitude interval, and a stopping rule. A goal that survives that translation was worth having. **Confusable with:** MF-07 (magnitude missing but the goal is checkable), MF-19, MF-21. **See:** the survey protocol [Curzi, 2026i]; the grounding protocol [Curzi, 2026ac].

MF-19 · The Retrofitted Criterion

Damages: Fa. **Clinical:** The success criterion is chosen after the outcome is known, and it is chosen to be met. The goal appears to have been achieved because the goal moved. **Signature:** The stated aim at the end does not match the stated aim at the start, and there is no record of when it changed or why. **Measurement:** Compare the pre-registration to the report, field by field. **If there was no pre-registration, the criterion is retrofitted by default, because there is nothing it could have been retrofitted from. Price:** Identical to MF-16 in mechanism and larger in effect, because the retrofit can select across outcomes *and* thresholds *and* windows. If ten plausible criteria were available, the familywise false-positive probability is **0.4013**; a retrofit is that process with the selection made deliberately. **Fix:** Preregister, freeze, and report deviations as deviations. **A deviation honestly reported is a result; a deviation quietly absorbed is this entry. Confusable with:** MF-16, MF-20, MF-22. **See:** the falsification suite [Curzi, 2026u]; the survey protocol [Curzi, 2026i].

MF-20 · The Latency Excuse

Damages: Fa. **Clinical:** A null is absorbed by an unbounded delay. The effect will come; it takes time; the timescale is never specified, so no elapsed time is ever enough. **Signature:** Every null is answered with patience and no null is answered with a date. **Measurement: Demand a bounded latency at stage G-4, in the same breath as the magnitude.** Then check whether the practice's own dynamics support the delay claimed. In the leaky integrator, latency is not free-floating: with `delta = 0.08` per period the time constant is `tau = −1/ln(1−delta) = 11.99` periods and the half-life to the fixed point is **8.31 periods**. **Price:** A latency claim of two periods against a system whose half-life is 8.31 is a misdiagnosis in one direction — evaluating a transient (Section 4, FM-4.6). A latency claim of "eventually" is unfalsifiable in the other, and is this entry. **The honest version specifies a number and defends it from the system's own dynamics. Fix:** Bound the latency in advance and treat exceeding it as the stated disconfirming result. **Confusable with:** MF-18, MF-21. **See:** the cadence result and the contraction theorem [Curzi, 2026a,?].

MF-21 · The Un-nulled Protocol

Damages: Fa. **Clinical:** No definition of a null result exists before the data arrive, so every outcome is interpretable as partial success. Distinct from MF-18: the *goal* may be perfectly checkable; what is missing is the statement of what nothing looks like. **Signature:** Results are reported as "encouraging" and never as "null". **Measurement:** Look for a pre-recorded SESOI and sample size. **Section 4 puts them at step 6 of twenty rather than at step 20 precisely because their position in the sequence is what makes them binding. Price:** Without a preset SESOI, the four null categories cannot be distinguished, so an uninformative null is read as an informative one in whichever direction the practitioner was already leaning — **overframing applied to one's own experiment. Fix:** Fix the SESOI and the required `n` in writing

before the first observation. **If the `n` is unaffordable, say so in week 0 and run the study as descriptive rather than inferential; that is an honest choice, and discovering it in week 12 is not. Confusable with:** MF-17, MF-18, MF-07. **See:** the survey protocol [Curzi, 2026i]; Sections 4.6 and 4.10.

MF-22 · Goodhart Capture of a Manifestation Metric

Damages: Fa, then B, then eventually Cn. **Clinical:** A measure adopted to track manifestation becomes the target, and thereafter it tracks itself. Goodhart's Law and Campbell's Law, applied to a domain unusually vulnerable to them because most of its metrics are self-reported. **Signature:** The metric improves smoothly while every unmeasured correlate of the thing it was proxying stays flat or declines. Enthusiasm about the number exceeds enthusiasm about the outcome. **Measurement:** Model the metric as `M = T + G`, where `T` is the target quantity and `G` is gaming uncorrelated with `T`. Then `corr(M, T) = 1 / sqrt(1 + sd_G^2)`:

```
gaming sd = 0.0  ->  corr 1.0000, T explains 100% of the metric's variance
gaming sd = 0.5  ->  corr 0.8944, T explains  80%
gaming sd = 1.0  ->  corr 0.7071, T explains  50%
gaming sd = 2.0  ->  corr 0.4472, T explains  20%
gaming sd = 3.0  ->  corr 0.3162, T explains  10%
```

Price: Gaming variance equal to the target's own variance halves the metric's informational content — **`corr(M,T) = 0.7071`, so the measure now says as much about the incentive as about the thing.** The specific manifestation-domain case is severe: **a group protocol whose success metric is convergence is measuring exactly the damage it does**, since alignment rising means `u` falling and `S = ln(1 + (N−1)u)` falling with it. **Fix:** Never let the success metric be a quantity the participants control directly. Pair every self-reported metric with one observable the participants do not control. **And re-measure `u` after any step that could align beliefs: if `u` fell, the step is reversed regardless of what else it produced. Confusable with:** MF-19, MF-11, MF-28. **See:** the independence requirement [Curzi, 2026g]; Sections 5.0.5 and FM-5.3.

37.6.8 6.8 Class V — ledger faults (Ld)

The evidence is real, and it has been booked to the wrong account.

MF-23 · Precedent Mistaken For Evidence

Damages: Ld, and everything downstream of it. **By the corpus's biological-layer constraints's own reckoning this is the single most likely failure of this strand of the corpus's method, and it is listed here in full for that reason. Clinical:** The concordance method pairs every substantive claim with a `(+1)` reference from the peer-reviewed corpus and a `(−1)` reference from the old corpus. **The two columns do not carry the same kind of thing. The `(+1)` column carries EVIDENCE — what is measured, replicated, falsifiable. The `(−1)` column carries PRECEDENT — that the idea was arrived at independently, recorded, and preserved, often for centuries, by people with no access to the instruments.** The fault is treating the pair as two witnesses. **Signature:** The phrase "both science and tradition confirm". Two sources cited, one conclusion drawn, and the conclusion is stronger than the measured source alone would support. **Measurement:** For each column, ask what would have had to be different in the world for

that source to say something else. **If the answer for the `(-1)` source is "nothing measurable", it is precedent.** Then compute the difference the mistake makes. **Price:** Bayes, with a flat prior and an evidential likelihood ratio of 3:

```
LR_evidence = 3.0, precedent treated at its honest weight LR = 1.3
    -> posterior = 0.7959
LR_evidence = 3.0, precedent double-counted as a second body of evidence LR = 3.0
    -> posterior = 0.9000
Overstatement in odds = 2.31x
```

The claim's odds are overstated by a factor of 2.31 by a single misbooked column — and the error compounds across a volume, because every concordance in the register makes it again. **What independent arrival DOES establish is real and worth having: that the structure was findable by more than one route, which is a genuine argument that the structure is in the territory and not only in the instrument.** That argument is worth an LR above 1 and nowhere near the evidential column's. **Fix: State the column's role every time a concordance is drawn.** A concordance that admits what each side can carry is more persuasive than one claiming both are proof, and a reader who checks will find the corpus was straight with them. Where a row has a `(-1)` entry and no `(+1)` entry, label it what it is: **a research question with a long history, which is interesting to have and is not a result. Confusable with:** MF-24, MF-25, MF-03. **See:** the concordance register [Curzi, 2026w]; the corpus's biological-layer constraints.

MF-24 · The Unbadged Speculative Bridge

Damages: Ld. **Clinical:** A bridge that was correctly marked speculative in the section that stated it appears unmarked in a summary, a table, a talk, or a later chapter. **The badge does not travel unless somebody carries it, and the standing rule is that it must be carried in the same sentence that states the bridge, not in a footnote. Signature:** A claim that is hedged on first appearance and flat on second. The abstract is more confident than the body. **Measurement:** Grep the corpus for the claim and compare every instance to the first. **VINO's binding rule, adopted unchanged: only bridges satisfying all three criteria — structural correspondence, predictive power, practical utility — are presented as established; speculative bridges are marked as such.** the biological-layer strand of the corpus tightens this by one notch: marked in the *same sentence*. **Price:** Not arithmetic, and it is the most corrosive entry in the catalogue. **A reader who finds one unbadged bridge is entitled to distrust every badged one**, because the badge's whole value is that it was applied consistently. One failure destroys the signal for all the others. **Fix:** Badge in the sentence. Re-check every summary against its source. **And a specific standing case for this strand of the corpus: any Grade R step — a research procedure with no established mechanism — carries its stamp wherever the step is described, including three months later when somebody calls it a technique. Confusable with:** MF-23, MF-25. **See:** the concordance register [Curzi, 2026w]; the standing refusals of the corpus, the corpus's biological-layer constraints.

MF-25 · Naming As Mechanism

Damages: Ld, CH. **Clinical:** A gap is given a name and the name is thereafter treated as an explanation. The vocabulary grows; the mechanism does not. **Gematria is a hash function — a map from strings to integers — and that description is literally accurate and carries no claim that the map is meaningful. The fault is the step from "we can describe this precisely" to "we have explained it." Signature:** Increasing terminological precision with no increase

in predictive content. A diagram in which the crucial arrow is labelled rather than derived. **Measurement:** Ask what the name predicts that its absence did not. **If the answer is nothing, it is a label.** A good name is still worth having — it makes the gap discussable — but a label that predicts nothing has not moved the claim along the grounding diagnostic by one stage. **Price:** The claim appears to have passed stage G-1 without passing it. The channel has a name and no ends. **Fix:** Keep the name and re-run G-1 against it. **The model to follow is the open problem of the missing demodulator [Curzi, 2026y]: the gap is named precisely, described as one well-specified falsifiable hole, and explicitly NOT presented as filled. Confusable with:** MF-01, MF-02, MF-24. **See:** the missing demodulator [Curzi, 2026y]; the concordance register [Curzi, 2026w].

MF-26 · The Un-decomposed Bundle

Damages: Ld, CH, B. **Clinical:** A protocol with many components produces an outcome and the outcome is attributed to one component — usually the newest, the most distinctive, or the one the author is most invested in. **Signature:** Nine things changed and one thing gets the credit. The credited component is the one that would be most interesting if it worked. **Measurement:** Ask whether the component was varied independently of the others. **If it was not, the design has zero identifying power for that component — not low power, zero — and no sample size repairs it. Price:** Section 4's worked example is the honest version: over twenty-four weeks `Q` rose from 1.1000 to 3.4905, a factor of 3.17, and the biological arm's contribution to that was **not estimable at all**, because it was introduced alongside two other changes and returned an uninformative null besides. **Fix:** Factorial design, staggered introduction, or attribution refused. **Attribution refused is a legitimate report and this framework prefers it to a guess. Confusable with:** MF-03, MF-15, MF-12. **See:** the falsification suite [Curzi, 2026u]; the survey protocol [Curzi, 2026i].

37.6.9 6.9 Class VI — consent faults (Cn)

The claim may be true. The protocol is not permitted.

MF-27 · The Consent Bypass

Damages: Cn. **This entry is governed by the fourth standing refusal of the corpus, which is absolute: this framework does not recommend an intervention that removes a person's ability to decline. Clinical:** A protocol acts on a person who cannot know what is being done, cannot refuse it, or cannot stop it. The bypass may be architectural (the step is concealed), social (refusal is expensive), or procedural (no decline path was ever built). **Signature:** A step that is described to observers and not to participants. A decline that has no defined mechanism. A group decision standing in for individual consent. **Measurement:** Count the per-person records for the three gates: **KNOW** (the participant can state what the step does, its grade, and what its null looks like), **DECLINE** (refusal is a recorded outcome with the same standing as assent), **REVERSE** (a written rollback exists, with its residue and anything irreversible named). **If there are fewer complete records than participants, consent was not obtained — it was assumed. Price:** Not arithmetic. **A protocol that fails this entry is not a worse protocol; it is outside what this framework will describe**, and no expected benefit changes that. **Fix:** Build the gates as steps that can halt the protocol, not as advice appended to it. **Advice cannot halt anything; gates can.** Section 4 §3.5 and Section 5 §3.10 specify them. **Confusable with:** MF-28, MF-30. **See:** Sections 4 and 5; the standing refusals of the corpus.

MF-28 · The Recruited Belief

Damages: Cn, and then `u`, and then everything the group's capability rests on. **Clinical:** The belief is produced by making dissent expensive rather than by making the belief persuasive. Consent is present in form and absent in substance. **Two mechanisms produce apparent commitment and they are not equivalent: a person who adopted a practice because it works for them would keep doing it if the programme vanished; a person who adopted it because dissent is expensive stops the moment the cost of dissent falls. Signature: A decline rate of exactly zero across many opportunities.** Unanimity that arrives fast and without argument. Nobody can name the last time the group changed its mind because of an internal objection. **Measurement:** Two instruments. First, **remove the sponsor for one period and count what survives** — what survives is the committed fraction, measured rather than asserted. Second, independent-estimate-before-discussion on 15 to 25 live questions, giving `u = 1 − rho^2`. **Price:** Recruited belief has `u = 0` with respect to the group's belief and therefore contributes nothing to `S = ln(1 + (N−1)u)`. At `N = 18`, moving `u` from 0.25 to 0 costs **1.6582 nats, which is 57.37% of the group's entire ceiling.** And the seeding calculation is corrupted at the same time: a committed fraction counted from compliance is a **fictitious numerator**, so the 25% threshold is being computed against a number that does not exist. **Fix:** Make dissent cheap, visible, and recorded, and treat a zero decline rate as a measurement of the decline path rather than of the protocol. **Protect dissent procedurally and test the protection. Confusable with:** MF-27, MF-22, MF-30. **See:** the two-mechanisms-of-commitment analysis and the independence requirement [Curzi, 2026g]; Section 5.

MF-29 · The Irreversible Step

Damages: Cn. **Clinical:** A step is taken that cannot be undone, and its irreversibility was not stated before it was taken. The fault is not irreversibility as such — some worthwhile things are irreversible — but irreversibility that arrives unannounced. **Signature:** The rollback plan does not exist, or exists and does not mention residue. **Measurement:** The three-line rollback record, required before any step: **STOP** (the single action that ends this step now), **RESIDUE** (what remains after stopping and for how long), and **IRREVERSIBLE** (what cannot be undone). **A non-empty third line escalates the step to a second, separate KNOW gate before it may run. Price:** Asymmetric and unbounded, which is the reason for the escalation. The entry that catches people is not the exotic one. **In Section 4's table the audio arm's residue is the memory of having heard something, which is bounded; the load-shedding step's residue is a resigned commitment, which is not.** An audio file has an off switch. A withdrawn obligation does not. **Fix:** Write the three lines first. Escalate any step with a non-empty third line. **And where shedding lands its cost on someone who did not consent, it is not shedding but transfer, and it is outside the line this framework draws. Confusable with:** MF-27, MF-20. **See:** the shedding lever [Curzi, 2026a]; Section 4.11.

MF-30 · The Masked Carrier

Damages: Cn, and it is the one bright-line violation in the catalogue. **Clinical:** Content is delivered below threshold, masked, or otherwise concealed from the person receiving it. **The no-subliminals constraint is stated absolutely: nothing masked, nothing below threshold, nothing concealed. A message hidden from its recipient cannot be consented to. Signature:** A file, a track, or a step whose full content the recipient cannot inspect. Any claim that the effect

depends on the recipient *not* attending to the content. **Measurement:** Can the recipient inspect the artefact in full, at audible level, and state what it contains? **Gate KNOW forbids the concealed layer by construction, because a participant cannot state what a step does if part of what it does is hidden from them.** The constraint is therefore enforced by the protocol's first gate rather than by good intentions. **Price:** Zero, and this is the entry's most useful fact. **Concealment buys nothing that disclosure does not: open-label placebo works when the recipient is told they are receiving placebo** (Kaptchuk et al., 2010, IBS), with modest effect sizes and an active literature. **The strongest verified belief-channel in the literature does not require deception**, which means the no-concealment constraint costs the method nothing and buys it everything. **Fix:** Disclose in full, make the artefact inspectable, and remove any component whose claimed action depends on the recipient not knowing about it. **If a step cannot survive disclosure, the step is the fault. Confusable with:** MF-27, MF-28. **See:** Section 4; the corpus's biological-layer constraints.

37.6.10 6.10 Faults of diagnosis itself

The faults of diagnosis itself — the ways this section gets misused.

FM-6.1 · The Entry Chosen By Resemblance. Picking the entry whose Signature field sounds like the complaint, rather than the one whose Damages field matches the coordinate. *Signature:* a diagnosis that fits the story and does not fit the arithmetic. *Fix:* C.1.

FM-6.2 · The Downstream Diagnosis. Diagnosing a magnitude or base-rate fault in a claim that already failed at the channel stage. *Signature:* elaborate statistics attached to a claim with no mechanism. *Fix:* C.2, and stop at the first failure.

FM-6.3 · The Catalogue As Dismissal. Using an entry to reject a claim without running its Measurement field. **Naming a fallacy is not diagnosing one**, and a catalogue used as a vocabulary of dismissal is worse than no catalogue, because it produces confident rejection at zero cost. *Fix:* the Measurement field is mandatory; the entry is the hypothesis, not the verdict.

FM-6.4 · The Exemption. Applying the catalogue outward and not inward. *Signature:* MF-08, MF-11, MF-23 and MF-24 are never invoked against the corpus's own claims. *Measurement:* count the in-house applications; this section contains four by construction and the status register [Curzi, 2026] is where the rest belong. *Fix:* run the catalogue on the corpus before running it on anybody else.

FM-6.5 · The Multiplied Entry. Reporting five entries for one fault because the fault touches five fields. *Signature:* a diagnosis longer than the claim. *Fix:* one coordinate, one entry, and the rest listed as consequences.

37.7 7 Results: three instances worked end to end

The three instances below are **constructed illustrative examples with declared parameters**, not field observations. Their value is that each runs the arithmetic to a verdict, including where the verdict is negative, and each records what the protocol forbade as well as what it permitted. Section 7.4 collects the results the rest of the corpus cites.

37.7.1 7.1 One practitioner, twenty-four weeks

The subject. One person, self-experimenting, twenty-four weeks. All figures are an illustrative construction, computed, and are not a measurement of any real individual or a report of any trial.

Week 0 — A.1 to A.9.

The occasion class: a weekly review of their own work, held Friday. The observable they do not fully control: whether the next week's plan survives contact by Wednesday, scored as a binary, logged for a year already and running at about 0.42.

Framings inventoried at A.4: `N = 7`. Pair test at A.5: of the 21 pairs, 7 have ever disagreed about the next action, so `u \ensuremath{\approx} 0.33`, rounded to `u = 0.35` with the rounding noted.

```
N = 7        ln N = 1.9459 nats  (the ceiling)
u = 0.35     N_eff = 1 + 6(0.35) = 3.1000
S = ln 3.1 = 1.1314 nats         realized fraction = 58.14%
```

Storage parameters at A.8: `delta = 0.08` per week, `eta = 0.45`, `C = 0.20`, `Q_0 = 1.10`. Headroom gate at A.9:

```
h_0 = 0.45 * 1.9459 - 0.08 * 1.10 - 0.20 = 0.5877  > 0   -> BUILD PERMITTED
maximum single-step increment r <= 0.5877
```

SESOI at A.6: they decide that a shift of 0.35 SD in their plan-survival score is the smallest change that would alter how they run the following quarter. From §4.9: detection at `d = 0.35` needs 65 paired observations; equivalence at SESOI 0.35 needs 51. **They have 24 weekly observations. They record, in week 0, that their study is underpowered for both, and that its verdict will therefore be N-1 unless the effect is large.** This is the single most consequential entry in the whole log, and it is made before any data exist.

Gate V at A.7: the L1 framing audit is flagged IRREVERSIBLE — a framing once named cannot be un-noticed — and escalated to a second Gate K, which they pass.

Weeks 1 to 24 — A.10 to A.13. The storage trajectory, computed:

```
Q_inf = (0.45 * 1.1314 - 0.20) / 0.08 = 3.8641
tau   = 11.9931 weeks ;  half-life 8.3130 weeks
```

week	Q_t	h_t
0	1.1000	0.5877
4	1.8839	0.5249
8	2.4455	0.4800
12	2.8479	0.4478
16	3.1361	0.4248
20	3.3426	0.4083
24	3.4905	0.3964

At week 12 the subject reports that the practice "feels much heavier than it did." The ledger says `h` has fallen from 0.5877 to 0.4478 — a 23.8% reduction in available headroom with no parameter changed. **Nothing has gone wrong; `delta·Q` is doing arithmetic** (the falling-headroom result [Curzi, 2026a]). They add nothing, per A.9: the correct response to falling headroom is not more input.

The ordinal tally at week 12: 7 over-framed, 3 calibrated, 2 under-framed. Dominant verdict OVER-FRAMED. Lever applied per §4.3: they add one observable they do not control (a colleague's independent estimate of whether the plan will survive) and one adversarial source. **No `D/P` value is computed and no `Delta` reported, because CC-R1 forbids it and the instrument was built to make the forbidden thing impossible rather than merely discouraged.** By week 24 the tally reads 9 over-framed, 9 calibrated, 6 under-framed. **The ordering has moved and the movement is the finding.** How far it moved is not reportable.

Weeks 13 to 24 — the L3 arm, A.14 to A.17.

At week 13 the subject elects to run the biological arm. Gate K second pass: they write out the A.14 sentence in full and sign the log. Headroom check: `h_13 \ensuremath{\approx} 0.4437` against an arm cost of 0.05 `C`-units per week, well inside `r <= h`. They proceed.

Artefact per A.15: 640 characters of their own text, base-4 encoded, address-prefixed as a format, rendered with a 40 ms window. Gabor check: at 40 ms the resolution floor is `1/(4·pi·0.040) = 1.9894 Hz` against a tone spacing of 12 Hz, inside the constraint by a factor of 6.03. **At a 5 ms window the floor would be 15.9155 Hz — wider than the tone spacing — and the artefact would no longer encode what it claims to encode.** They record the check.

Pre-registration per A.16: null hypothesis stated, observable fixed, SESOI 0.35 carried over, the sample size acknowledged as far below the 51 required for equivalence, four baselines specified, and the structural ceiling test added — any claimed single-week increment must satisfy `r_t <= h_t`, which at week 13 is 0.4437.

Result at week 24. Observed standardized difference on the plan-survival score, arm weeks versus matched non-arm weeks: `d = 0.06`. With `n = 20` paired comparisons the 95% interval is

```
0.06 +/− 1.959964 / sqrt(20)  =  0.06 +/− 0.4383  =  [−0.3783, 0.4983]
```

The interval is not contained within +/- 0.35. Classification: N-1, uninformative null. And because the arm carried a Grade R stamp with an expected null, it is *also* N-4. The subject records both classifications and writes the conclusion the protocol requires:

> *"This run tells me nothing about whether the L3 arm has an effect. My interval is wider than my own SESOI. To conclude equivalence at SESOI 0.35 I would need 51 paired observations; to detect d = 0.35 I would need 65. Separately, and independently of my sample size, the arm had no established mechanism when I started it and still has none: there is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal. The well-specified gap of the open problem of the missing demodulator [Curzi, 2026y] is exactly where it was. I am recording this null because nulls of this kind are the ones that do not otherwise get recorded."*

What the subject actually got out of twenty-four weeks. Not the L3 arm — that returned nothing, as predicted, at a cost of 0.05 `C`-units per week for twelve weeks. What moved was L1 and L2: `Q` rose from 1.1000 to 3.4905, a factor of 3.17, and the ordinal tally shifted from 7:3:2 toward 9:9:6, from a dominant over-framed reading to a balanced one.

That is the honest shape of this practice, and the section would rather print it than a better story. The arithmetic worked, the measured channels did modest measured things, the research arm returned its expected null, the practitioner could have stopped at any gate, and nothing was hidden from them at any point.

One last number, because it generalizes. The fixed-point headroom identity (CC-R6) for this subject: `h_inf = eta·(ln N − S) = 0.45 × (1.9459 − 1.1314) = 0.3665`. At the fixed point `delta`, `C` and `Q` cancel entirely. **When this practice matures, its capacity for anything new will depend on one quantity only: the gap between the seven ways this person can see their situation and the surplus those seven ways actually collect.** Everything else washes out. That is where the next quarter's work is, and the arithmetic said so without anybody needing an opinion about it.

37.7.2 7.2 An eighteen-person unit, forty months

The system. An eighteen-person applied-research unit inside a 120-person division. It has a reputation for quality and a growing complaint that "nothing new gets started here." All figures are an illustrative construction, computed, and are not a report of any real organisation.

Period 0 — B.1 to B.7.

Survey estimates: `eta = 0.30`, `delta = 0.06` per month, `C = 0.55`, `Q_0 = 1.20`. `N = 18`, `ln N = 2.8904`.

```
Gate 0:  h_0 = 0.30*2.8904 - 0.06*1.20 - 0.55 = +0.2451     BUILD PERMITTED, r <= 0.2451
```

Gate I baseline at B.5: independent estimates on 20 live questions give mean pairwise signed- error correlation `rho = 0.866`, so `u = 1 - 0.866^2 = 0.25`. The estimator depends only on `|rho|` — the opposition identity of Curzi [2026], under which `rho = +1` and `rho = −1` both give `u = 0` and `F = 0`, so a perfectly opposed pair purchases exactly as much surplus as a perfectly aligned one, namely none.

```
N_eff = 1 + 17(0.25) = 5.2500        S = ln 5.25 = 1.6582 nats
realized fraction = 1.6582 / 2.8904 = 57.37%
```

Turn-taking in the standing meeting: the three most senior members take 62% of speaking time between them. `PR` on the eighteen shares comes out low and `N_eff` from participation confirms the estimate above to within rounding.

B.7 requires both numbers, and here is why. **Headroom is positive at +0.2451, so the unit looks healthy. The one-step growth test says otherwise:**

```
growth iff  eta*S > delta*Q + C
            0.30*1.6582 = 0.4975   vs   0.06*1.20 + 0.55 = 0.6220
            0.4975 < 0.6220   ->   SHRINKING
```

And the steady state:

```
Q_inf = (eta*S - C) / delta = (0.4975 - 0.55) / 0.06 = −0.8755
```

A negative equilibrium. the fixed-point theorem of the storage law [Curzi, 2026a], does not hide this failure mode: if `eta F(u*) < c*` the equilibrium is negative and the regime is unsustainable or the linear model must be replaced. Projected forward, `Q` falls from 1.20 to 0.1123 by month 12 and crosses zero around month 15. **The unit is not stagnating because its people are tired. It is operating at 57% of its own ceiling while paying full maintenance, and the arithmetic has been saying so for some time in a language nobody was reading.**

The complaint — "nothing new gets started here" — was accurate and misattributed. Positive headroom said the unit *could* build; the growth test said it was losing stock every month it did not. Both were true simultaneously, which is exactly the case the headroom-versus-growth diagnostic split [Curzi, 2026g] warns about.

The decision. Two levers, in order.

Lever 1, `u`: 0.25 -> 0.45. Not by hiring. By B.9 — independent written estimates before discussion on every live question — and by rotating who opens the meeting, which addresses the 62% concentration directly. The Jensen correction at B.6 is applied to the eight periods of follow-up measurement rather than to their average.

Lever 2, `C`: 0.55 -> 0.45. Two standing reports are retired at B.16 after checking that retiring them does not transfer the work to another team, which would be transfer rather than shedding.

Recomputed:

```
S(0.45) = ln(1 + 17*0.45) = 2.1576 nats        realized fraction = 74.65%
Q_inf   = (0.30*2.1576 - 0.45) / 0.06 = 3.2878
h_inf   = eta*(ln N - S) = 0.30*(2.8904 - 2.1576) = 0.2198     (CC-R6: delta, C, Q cancel)
```

month	Q_t	h_t	one-step test
0	1.2000	+0.3451	grow
4	1.6578	+0.3176	grow
8	2.0151	+0.2962	grow
12	2.2942	+0.2795	grow
20	2.6821	+0.2562	grow
40	3.1121	+0.2304	grow

Note what each lever did. Raising `u` alone would have moved the steady state from **-0.8755 to +1.6211** — from a regime with no sustainable stock at all to one that grows, but only to 1.62 Q-units, 35% above where the unit already sat. **The geometry lever decides whether a sustainable stock exists; the cost lever decides how large it is**, and here the cost lever doubled it again to 3.2878. Neither substitutes for the other.

The limit of the geometry lever is worth stating exactly, because it is the theorem people misquote. Corollary 3.4 says that if `C_t > eta ln N − delta Q_t`, **no** interaction geometry prevents decline. This unit was never in that regime — its `h` was +0.2451, so geometry alone could have rescued it. **Unit C of §5.1, at `h = −0.4229`, is the regime where the geometry lever is dead, and the only honest moves there are cost and shedding.**

Periods 1 to 6 — seeding the new practice (B.10 to B.12).

The practice being established is the independent-estimate discipline itself. Sizing at B.10: 25% of 18 is 4.50, so **5 committed people**. The unit has three natural sub-teams of six; 25% of 6 is 1.50, so **2 committed people per sub-team** — and the unit chooses to seed one sub-team of six with four committed members (67%, far past its tipping point) rather than to spread five across eighteen.

Sponsor-independence test at B.12: the unit lead takes a month away from the meeting. On return, the seeded sub-team is still collecting independent estimates and the two unseeded ones are not. **That is the committed fraction, measured rather than asserted**, and expansion proceeds from it along working ties rather than the org chart.

Period 7 — the biological arm, and its expected null.

The unit elects to run the Grade R arm. Gate 0b at B.3: arm cost 0.04 per month against `h_7 \ensuremath{\approx} 0.30`, inside the ceiling. Gate K at B.14 is run per person; **three of eighteen decline, and the three declines are recorded in their own rows as declines.** The remaining fifteen proceed. **The decline rate is 16.7%, not zero, which is the first evidence in this worked example that the decline path functions at all.**

Artefact: 1,200 characters of text the unit wrote about its own situation, base-4 encoded, address-prefixed as a format, rendered with a 30 ms window. Gabor check: `1/(4·pi·0.030) = 2.6526 Hz` against a tone spacing of 12 Hz — inside the constraint by a factor of 4.52. Recorded.

B.15 is enforced: each of the fifteen writes their judgement of the artefact before any discussion of it. The pooled spread is wide, which is the informative part and which a group discussion would have erased.

Result. Pre-registered observable, twelve months, four baselines plus the ceiling test. The arm does not beat the persistence baseline. The largest month-on-month increment observed in `Q` during the arm period is 0.0864, against `r_max` values running from 0.3011 down to 0.2585 across those months — **inside the structural ceiling throughout, so no repackaging is indicated, which is the only positive thing the structural test can say.** The unit records:

> *"The Grade R arm returned a null, as expected. There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal, and our result changes nothing about that gap in either direction. Separately: we did not run an acoustic-to-c intervention, because the acoustic intervention analysis [Curzi, 2026f]'s causal test is reported as not confirmed and we have no*

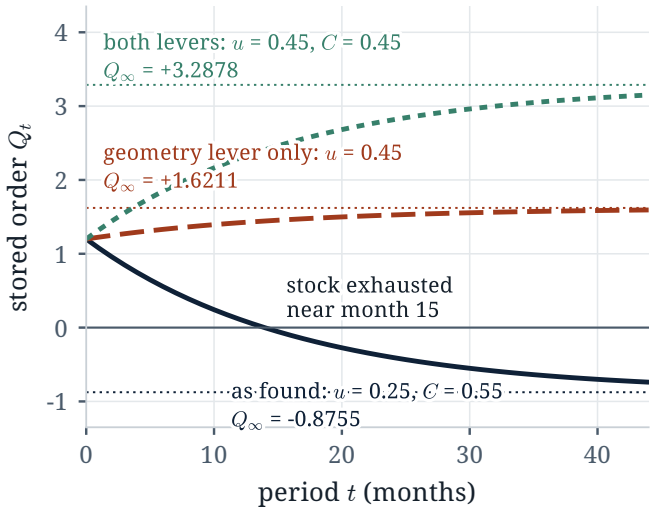


Figure 37.5: **The geometry lever decides whether a sustainable stock exists; the cost lever decides how large it is.** As found, the unit passes its entry gate at +0.2451 and still runs to a negative equilibrium of −0.8755, exhausting its stock near month 15. Raising `u` from 0.25 to 0.45 alone moves the equilibrium to +1.6211; cutting `C` from 0.55 to 0.45 as well moves it to +3.2878. Modelled at `eta` = 0.30, `delta` = 0.06, `Q_0` = 1.20 from the parameters of Section 7.2; generated by `figs/src/CCP-037_fig04.py`.

> *basis for describing a room as raising `c`. What actually moved our numbers was `u` from 0.25 to 0.45 and `C` from 0.55 to 0.45."*

Gate I, post-step (B.17). Re-measured `u` after the arm period: 0.44, against 0.45 before. Within measurement noise, and critically **not lower in a way the reversal rule would trigger**. Participating `N` unchanged at fifteen for the arm and eighteen for the unit. The shared-input hazard of Restriction 1 was contained by B.15, and the unit records that containment as the design feature it was rather than as luck.

What the eighteen months bought. `Q` from 1.20 to roughly 2.29 by month 12 and 3.11 by month 40, against a trajectory that would have crossed zero at month 15. Realized ceiling fraction from 57.37% to 74.65%. **A programme that would have been launched under a headroom of +0.2451 against a cost of 0.30 — and would therefore have driven the unit negative before doing anything — was not launched, because somebody computed the fourth row of the table in §5.1.**

And the arm returned nothing, which is recorded, because nulls of this kind are exactly the ones that do not otherwise enter the record. The unit is now, by its own measurement, more capable and more disagreeable than it was, and those two facts are the same fact.

37.7.3 7.3 A claim diagnosed against the catalogue

The claim. A well-run twelve-person team reports that after adopting a listening practice built on this framework's genomic-audio method, its "decision quality index" rose from 62 to 81 over one quarter, that "everyone found it helpful," and that the team therefore recommends the practice to the rest of its 120-person division. All figures are an illustrative construction, computed, and are not a report of any real organisation.

C.1 — locate the coordinates. The claim touches four: **CH** (what channel is asserted?), **B** (what is the base rate behind "everyone found it helpful"?), **Fa** (what would have counted as a null?), and **Mg** (is 62 to 81 a permitted increment?). Consent, **Cn**, must also be checked, and per C.6 it is checked last and as a veto.

C.2 — run the grounding sequence, and stop at the first failure.

G-1, identify the channel. The team's account: "the encoded narrative reaches the cells and activates latent capability, which shows up as better decisions." **HALT.** The claimed channel is acoustic transmission of sequence-specific content to a cell that demodulates it. **There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal** (the open problem of the missing demodulator [Curzi, 2026y]). The channel cannot be named because it does not exist in the established literature.

Diagnosis 1: MF-02, The Absent Channel. Not MF-01 — the team named a channel specifically, which is the honourable version, and naming it is what let the diagnosis take four minutes instead of four years.

The follow-up question is the one that matters. Upstream is genuinely real: PIEZO1/2 are mechanically gated ion channels, sound is mechanical, and Tajik et al. (*Nature Materials*, 2016) showed surface force transmitting through the cytoskeleton to stretch chromatin and upregulate transcription. **Asked whether that upstream reality licenses their claim, the team says yes.** It does not: the established pathway is real and **non-specific**, and the claim is about *which* content, addressed to a *particular* sequence.

Diagnosis 2: MF-03, Channel Substitution, with the seam located exactly — the sentence begins with a Nobel-recognized mechanism and ends with an addressed, sequence-dependent outcome.

C.3 — price the base-rate claim. "Everyone found it helpful" is twelve testimonials with no denominator and no mechanism by which an unhelpful experience could have entered the record. **Diagnosis 3: MF-13, The Survivorship Narrative.** Priced: if the outcome the team is describing occurs spontaneously at even a 5% rate, then twelve favourable reports require only 240 attempts to arise **with no effect whatsoever**, and the team ran one attempt with twelve people over a quarter in which many other things also changed.

Worse, the team recruited for the practice during a bad quarter. **Diagnosis 4: MF-15, Regression Mistaken For Effect.** Priced: selecting the bottom decile of a normally distributed measure gives `E[Z|selected] = −1.7550`; at a test-retest reliability of 0.6, the expected improvement **with no intervention at all** is `(1 − 0.6) × 1.7550 = 0.7020 SD`. On the team's own index, whose quarterly SD is 14 points, that is **9.83 points of expected improvement from nothing** — more than half the 19-point rise they reported.

C.4 — run the structural test. The team's parameters, estimated by survey: `N = 12`, `ln 12 = 2.4849`, `eta = 0.30`, `delta = 0.06`, `Q = 2.10`, `C = 0.55`.

```
r_max = 0.30 * 2.4849 - 0.06 * 2.10 - 0.55 = 0.0695 per period
```

Their index rise of 19 points, expressed on the framework's normalized scale, corresponds to a claimed single-quarter increment of 0.58 Q-units. **That is 8.35 times `r_max`.**

Diagnosis 5: MF-12, The Ceiling Violation — and, given the practice includes a heavily filtered summary of the team's own prior decisions fed back to them, **MF-11, The Repackaged Input** must be tested for explicitly. R4 is the precedent: a filtered repackaging passed all four empirical baselines **including the shuffle at p = 0.0005** and was caught only here, at 38.5x the ceiling. **The team's four baselines all passed. That fact is now known to carry no weight against this specific failure mode.**

The falsifiability check. Asked what result would have made them stop, the team offers none; asked what a null would have looked like, they offer none. **Diagnoses 6 and 7: MF-18, The Unfalsifiable Goal, and MF-21, The Un-nulled Protocol.** The index itself is self-reported by the participants. **Diagnosis 8: MF-22, Goodhart Capture** — modelled as `M = T + G` with gaming variance equal to the target's, `corr(M, T) = 0.7071`, so half the metric's variance is now about the incentive rather than the outcome.

C.5 — audit the ledger. The team's write-up cites a contemplative tradition of sounded practice alongside the mechanotransduction literature and concludes that "both the science and the tradition support this." **Diagnosis 9: MF-23, Precedent Mistaken For Evidence.** The tradition carries precedent — that the structure was findable by another route, which is genuinely worth having — and it carries no evidence. Priced with a flat prior and an evidential

LR of 3: booking precedent at 3.0 rather than at an honest 1.3 moves the posterior from **0.7959 to 0.9000** and overstates the odds by **2.31x**.

The write-up also describes the biological arm without its grade stamp. **Diagnosis 10: MF-24, The Unbadged Speculative Bridge** — the arm was correctly stamped Grade R in the team's own week-1 notes and appears unstamped in the quarterly summary.

C.6 — consent, as a veto. Twelve participants; four complete KNOW records, no DECLINE records at all, and no rollback documents. **The decline rate is exactly zero.**

Diagnosis 11: MF-27, The Consent Bypass — fewer gate records than participants means consent was assumed rather than obtained. **Diagnosis 12: MF-28, The Recruited Belief** — a zero decline rate across twelve people and thirteen weeks is a measurement of the decline path, not of the protocol. The team is asked to run the sponsor-absence test; the lead steps back for one month and on return the practice has stopped in three of the four sub-groups.

The verdict, in the order the Specification produces it.

```
CH  MF-02  Absent Channel             -> the claim halts at G-1
CH  MF-03  Channel Substitution       -> real upstream, unlicensed downstream
B   MF-13  Survivorship Narrative     -> 12 reports; 240 attempts suffice at a 5% base rate
B   MF-15  Regression Mistaken        -> 9.83 index points expected from nothing
Mg  MF-12  Ceiling Violation          -> claimed increment 8.35x r_max = 0.0695
Mg  MF-11  Repackaged Input           -> must be tested; the shuffle cannot see it
Fa  MF-18  Unfalsifiable Goal         -> no stated disconfirming result
Fa  MF-21  Un-nulled Protocol         -> no SESOI, no sample size, fixed in advance
Fa  MF-22  Goodhart Capture           -> corr(M,T) = 0.7071 on a self-reported index
Ld  MF-23  Precedent as Evidence      -> posterior 0.9000 for an honest 0.7959; 2.31x odds
Ld  MF-24  Unbadged Bridge            -> Grade R stamp lost between notes and summary
Cn  MF-27  Consent Bypass             -> 4 gate records for 12 participants
Cn  MF-28  Recruited Belief           -> decline rate exactly zero; 3 of 4 stopped
```

What the team should conclude, and it is not "the practice is worthless." Two of the thirteen weeks' worth of changes were L1 and L2 changes in Section 4's scheme — the team began recording individual judgements before discussion, and it retired a standing report. **Those are theorem-backed and measured-channel steps, and they have a plausible claim on the index movement that survives every diagnosis above.** What has no claim on it is the biological arm, which was never licensed by any established mechanism, was never independently varied (**MF-26**), and would not have been attributable even if the design had been otherwise perfect.

And the veto stands regardless of any of that. The consent findings are not weighed against the benefit; **a protocol that fails Class VI is outside what this framework will describe.** The correct next step is not a better analysis. It is to rebuild the gates as steps that can halt the protocol, obtain twelve complete records, run the sponsor-absence test again with the decline path working, and only then decide whether there is anything left worth measuring.

Twelve entries in one quarter from one well-meaning team, all of them priced with arithmetic the team could have run themselves in an afternoon. That is what the catalogue is for, and it is why it is written to be opened before the claim is made rather than after it is published.

37.7.4 7.4 The results this paper contributes

R1 — The headroom gate as a precondition. `h = eta ln N − delta Q − C`, with `h < 0` forbidding any addition, `h = 0` permitting maintenance only, and `h > 0` capping every requested single-period increment at `r <= h`. Worked at `N = 18`, `eta = 0.30`, `delta = 0.06`, `C = 0.55` and `Q = 1.20`: `h = +0.2451`, and the critical stock is `Q* = 5.2852`.

R2 — A positive gate is compatible with a negative equilibrium. The same unit realises `F(0.25) = 1.6582` nats, 57.37 per cent of ceiling, and returns `Q_inf = (0.30 x 1.6582 −`

`0.55)/0.06 = −0.8755`. Raising `u` from 0.25 to 0.45 lifts the realised fraction to 74.65 per cent and the equilibrium to `+1.6211`; **cutting `c` from 0.55 to 0.45 as well** lifts it to `+3.2878`.

R3 — The Moment equation is demoted at scale-type gate 2 at both scopes. The limiting cases, the fixed point `D = P` and the three-way ordering survive; the magnitude does not; and the Rasch route out costs **123 items per term** [Curzi, 2026t].

R4 — The Jensen correction. Feasibility must use the mean of `F(U)`, not `F` of the mean `U`. The wrong computation overstates by **0.0790 nats, 4.31 per cent**, worth **0.3952 stored units** at `eta = 0.30`, `delta = 0.06`, independent of `C`.

R5 — Cluster rather than spray. 16 seeds across 120 people at mean degree 8 under a two-of- eight rule leave **29.00 per cent** of non-seeds with two or more seeded neighbours — 30.16 people — and **39.17 per cent** with exactly one; the same 16 concentrated in one 20-person unit activate **4 of 4** at **80.0 per cent** local commitment.

R6 — A finite catalogue. Every manifestation claim has exactly eight coordinates — a taxonomy, not a theorem — yielding **30 named failure modes in six classes**, each priced by one of four identities including `X = G B` with `G = 0` [Curzi, 2026ab].

CR — The counter-result. The collective worked example **fails on the framework's own arithmetic**: a negative equilibrium at `Q_inf = −0.8755` despite passing the entry gate. The acoustic-to-collective-intelligence causal test is carried forward as **NOT confirmed**. And a filtered repackaging of the input **passed all four empirical baselines including the shuffled-input control at `p = 0.0005`**, caught only by the structural ceiling [Curzi, 2026u].

37.8 8 Discussion

37.8.1 8.1 What follows

The gate is cheap and it is not being computed. Four parameters and a logarithm decide whether a group may add anything at all, and the decision it returns is frequently the opposite of the one the room would have reached. The fourth row of the table in Section 5.1 is the case: a unit with positive headroom of +0.2451 considering a programme costing 0.30 per period. The programme's cost exceeds the headroom, so adopting it moves the unit from feasible to infeasible *before the programme has done anything*. The correct move is to shed 0.06 or more of existing cost first, or to redesign the programme below 0.2451 — not to run a pilot and see.

Maturity is not the enemy, and the arithmetic says which kind of maturity is. At the fixed point the headroom identity reduces to `h_inf = eta(ln N − F(u*))`, in which `delta`, `C` and `Q` all cancel. A mature unit's capacity for anything new therefore depends on exactly one thing: the gap between its ceiling and the surplus it actually collects. Since `F(u) <= ln N` always, `h_inf >= 0` always, and **the plain linear model never ossifies by itself**. Ossification requires exogenous accretion of `Q` — inherited obligations, imported compliance regimes, acquired estates — or superlinear cost growth. The diagnosis therefore branches: stock grown from the group's own interaction is self-limiting and the fix is on `u` and `C`; stock arriving from outside is stock the group never converted, and the fix is deliberate shedding [Curzi, 2026a]. The identity itself is re-derived and stated as R5 in Curzi [2026].

The two levers do different jobs and neither substitutes for the other. In the worked unit, raising `u` from 0.25 to 0.45 alone moves the equilibrium from `−0.8755` to `+1.6211` — from a regime with no sustainable stock at all to one that grows. Cutting `C` from 0.55 to 0.45 then doubles it again to `+3.2878`. **The geometry lever decides whether a sustainable stock exists; the cost lever decides how large it is.** And the geometry lever has a hard boundary: where `C_t > eta ln N − delta Q_t`, no interaction geometry prevents decline [Curzi, 2026a]. Unit C of Section 5.1, at `h = −0.4229`, is in that regime and geometry is dead there.

The cheapest instrument in the paper reads onto a term the law consumes. A stopwatch and six columns give the share vector, the share vector gives the participation ratio, and the participation ratio gives the effective count exactly [Curzi, 2026c]. That is what makes the turn-taking correlate of collective intelligence [Woolley et al., 2010] computable rather than

merely quotable — with the two honesty notes that travel with it: speaking time is a proxy for participation rather than participation itself, and the published result is a correlation between turn-taking equality and the collective-intelligence factor, not a demonstration that equalising turns raises it.

Shrink the denominator. The critical-mass threshold is a proportion [Centola et al., 2018, Granovetter, 1978], so the cheapest route over it is not to recruit more committed people but to narrow the population within which the convention must hold: 100 committed people in a 400-person firm is a campaign with a budget; 2 in a squad of 6 is a pair. Sequential local tipping is cheaper than simultaneous global tipping by a wide margin and produces clusters as a side effect rather than requiring them as a precondition.

37.8.2 8.2 What does not follow

Nothing here licenses an acoustic capability claim. The collective-intelligence factor is measured, is not predicted by members' average or maximum intelligence, and correlates with equality of conversational turn-taking [Woolley et al., 2010, Engel et al., 2014]. The step from *a room's acoustics affect turn-taking* to *a room's acoustics raise the collective-intelligence factor* is a causal claim, its test is reported as not confirmed [Curzi, 2026f], and this paper carries that status without softening. A group may run an acoustic intervention as a pre-registered experiment; a group may not describe it as something the room does. The passive acoustic mechanisms themselves are real and are the reason the proposal is worth testing [Sabine, 1922, Beranek, 2004].

Nothing here licenses a diversity-is-good conclusion in general. The framework prices difference through `u`, and the empirical literature on whether difference helps is genuinely divided. Hong and Page [2004] is a model result, and Thompson [2014] shows that the "diversity trumps ability" claim as usually stated does not follow from it; the subsequent exchange [Kuehn, 2017, Singer, 2019, Grim et al., 2019] narrows the conditions under which it survives. Mannix and Neale [2005] and van Knippenberg and Schippers [2007] find that surface-level diversity frequently hurts in the short run, and Harrison and Klein [2007] shows that three incompatible constructs — separation, variety, disparity — are all called diversity, so a paper using `u` must say which it means. **`u` in this paper is disparity of *recommended action*: two framings that never disagree about what to do next are one framing.** The oldest and cleanest statement of the underlying point is Markowitz [1952]: diversification benefit depends on covariance, not on count, so `N` uncorrelated holdings and `N` perfectly correlated holdings have the same `N` and completely different risk. That is the corpus's four-jurisdictions result, stated in 1952 with the mathematics attached.

Nothing here establishes a biological effect. The channel inventory of Section 4.2 separates two rows that are routinely merged: acoustic input reaching cells through mechanically gated pathways, which is established and **non-specific**, and acoustic input carrying sequence-specific content addressed to a particular sequence, for which **no mechanism is established** [Curzi, 2026y]. Every step touching the second is graded R — a research procedure with an expected null — and the expected null is what the worked instance of Section 7.1 in fact records.

Nothing here converts an ordinal reading into a magnitude. The Moment equation `M = (D/P) E` fails gate 2 of the Formula Validation Protocol at both scopes, because division requires a ratio scale with a true zero and the two constructs are not instrumented that way [Curzi, 2026t, Stevens, 1946]. Structure survives, magnitude does not, and the constructive repair is priced rather than promised: a Rasch logit difference supplies the missing zero at **123 items per term** [Rasch, 1960, Andrich, 1988], so 246 items for the ratio, 369 for all three terms, and 61.5 person-hours for a thirty-participant recovery gate before any analysis.

37.8.3 8.3 Relation to the surrounding literature

The protocol's shape — measure a state, gate an action on a computed threshold, record a decline as data — is standard control engineering applied to an institution, and it is the shape

the resilience literature arrives at from the other direction [Holling, 1973]. Its governance commitments are Ostrom's: durable common-pool arrangements are designed, monitored and graduated rather than exhorted [Ostrom, 1990, 2010]. Its seeding results are complex-contagion results [Centola and Macy, 2007, Centola, 2010, Watts, 2002], and the network-structural analogue of the effective count — adoption predicted by the number of distinct connected components among a person's contacts rather than by raw contact count — is Ugander et al. [2012], which is the closest empirical relative of `N_eff` in the social sciences. The observation that a system's structure mirrors the communication structure of the organisation that produced it [Conway, 1968] is why the seeding order in Section 5.8 runs along the graph of who works with whom rather than the graph of who reports to whom, and the quadratic growth of coordination paths — 15 at six people, 153 at eighteen, 435 at thirty — is Brooks [1975].

The paper's inference commitments are those of the reform literature: pre-register, freeze, report held out [Nosek et al., 2018]; treat researcher degrees of freedom as the dominant threat [Simmons et al., 2011, Gelman and Loken, 2013]; expect base rates to be low [Ioannidis, 2005] and replication to be partial [Open Science Collaboration, 2015]; and prefer risky point predictions to directional ones, which in psychology are nearly costless [Meehl, 1967, 1990]. Platt [1964] is the procedural form of the grounding sequence used throughout. Two of these bear directly on the catalogue: the survivorship entry (MF-13) is a base-rate argument, and the forking-path entry (MF-16) is Gelman and Loken [2013] restated as a fault with an instrument attached.

Two structurally parallel results elsewhere in the corpus are worth naming, because they show the same auditing move working on quantities that have nothing to do with rosters. In the frequency framework, the odd-part class rule of Curzi [2026] — the pitch class of `f(n) = k n^2` is the odd part of `n`, so `1..N` realises exactly `ceil(N/2)` classes, six at `N = 12` — is a count of distinguishable outcomes collapsing well below its naive value for a purely structural reason. In the ledger atlas, Curzi [2026] shows that dwell time is the whole freeze exposure: a roughly five-second hop is `1.585e-07` of a year against `0.0822` for value rested thirty days, a ratio of about 518,400 to one, so the exposure that matters is the one nobody was timing. **Both are the same discipline as the participation ratio: audit the count before pricing it.**

Finally, the corpus's own position in sociophysics should be named rather than assumed. The programme of carrying physical formalism into social description is long-established [Galam, 2012, Castellano et al., 2009, Helbing, 2001], and so is the objection to it: that the transfer often borrows the authority of physics without its constraints [Mirowski, 1989, Cartwright, 1999, Sobkowicz, 2020]. **This paper's answer is the analogy-discipline table of Section 3.13 and the verification chain of Section 9, which together say exactly which results are theorems, which are shared mathematical form, and which are asserted.**

37.9 9 Verification

The operative results of this paper compound from six postulates. Each is atomic, each is checkable on its own, and each carries its status. Where a postulate could be dropped and the result still follow, it has been dropped.

37.9.1 9.1 Dimensional audit

`[Q]` denotes the dimension of stored order and `[T]` the dimension of time; the period is the unit of `[T]`, and every rate below is *per period*, which is why no explicit `[T^-1]` appears once the equations are written in discrete steps.

Symbol	Dimension	Unit	Status
`Q_t`	`[Q]`	proxy units (practices, protocols, runbooks)	**measured**, in whatever the declared order proxy counts
`u`, `U_t`	`[1]`	none	**defined** dimensionless: `u = 1 - (x.y)^2` for unit vectors
`N`	`[1]`	count of participating members or framings	**measured** count
`N_eff`	`[1]`	effective count	**derived**: `1 + (N-1)u`, affine, *not* a Hill number
`PR`	`[1]`	effective count	**derived**: `1/sum w_i^2`, a Rényi-2 participation ratio
`S = F(u)`	`[1]`	nats	**derived** dimensionless: a logarithm of a count
`eta`	`[Q]` per nat per step	proxy units per nat per period	**measured**; *not* dimensionless
`delta`	`[1]`	per period	**measured** dimensionless fraction per step
`C_t`	`[Q]` per step	proxy units per period	**measured**
`h_t`, `r_max`	`[Q]` per step	proxy units per period	**derived**
`Q_inf`, `Q*`	`[Q]`	proxy units	**derived**: `[Q]`/step over `[1]`/step returns `[Q]`
`tau`, `t_half`	`[T]`	periods	**derived**
`D`, `P`, `E`	—	**none supplied**	**ordinal only**; the framework supplies no unit and the ratio `D/P` is demoted
`dt`, `df`	`[T]`, `[T^-1]`	seconds, hertz	**measured**; their product is dimensionless

Does the gate balance? `h_t = eta ln N - delta Q_t - C_t` balances if and only if every term carries `[Q]` per step. `ln N` is dimensionless, so `eta ln N` carries `[Q]` per step exactly when `eta` carries `[Q]` per nat per step; `delta Q_t` carries `[Q]` per step because `delta` is a dimensionless fraction per step; `C_t` carries `[Q]` per step by construction. **The gate is therefore well formed, and a headroom calculation in which `C` is denominated in hours or currency while `Q` is denominated in practices is dimensionally malformed however tidy it looks.**

Where the framework has no operational unit, this paper says so. Stored order `Q` is counted in whatever proxy is declared, and two systems' `Q` values are not comparable unless they declare the same proxy. Comfort itself, the corpus's underlying construct, has **no operational unit at all**. Perspective `D` and perception `P` have **no unit and no zero**, which is precisely why their quotient is demoted; the framework is silent on their dimensions and that silence is recorded here rather than filled in.

37.9.2 9.2 The postulates

P1 — The storage recurrence is linear and first order. `Q_{t+1} = (1−delta) Q_t + eta S_t − C_t` with `delta`, `eta`, `C` not depending on `Q`. *Status:* **DEFINITION** — the model form adopted by the framework [Curzi, 2026a]. It carries no empirical weight on its own.

P2 — Surplus is bounded: `0 <= F(u) <= ln N`, with `F(u) = ln(1 + (N−1)u)` on `u` in `[0,1]`. *Status:* **CITED** — the range result of the surplus functional [Curzi, 2026].

P3 — `F` is strictly concave in `u`. *Status:* **CITED** — proved with the axiom set in Curzi [2026]; `F''(u) = −(N−1)^2/(1+(N−1)u)^2 < 0` for `N > 1`.

P4 — Jensen's inequality: for concave `F`, `F(mean u) >= mean F(u)`. *Status:* **CITED** — standard, and the direction is what makes the correction one-sided.

P5 — For a share vector `w` with `sum w_i = 1`, the participation ratio is `PR = 1/sum w_i^2`. *Status:* **DEFINITION** — the reciprocal Rényi-2 collision probability [Rényi, 1961, Simpson, 1949].

P6 — Cost accrues every period whether or not surplus was delivered. *Status:* **POSTULATED** — an assumption about how real systems charge overhead. It has no external support, it is used at exactly one place (C6 below), and no other result in the chain depends on it.

37.9.3 9.3 The compounding chain

```
P1 + P2        -> X1   a nondecreasing step needs eta F(u) >= delta Q + C, and since
                       F(u) <= ln N the step is possible at all only if
                       C <= eta ln N − delta Q.                              ARITHMETIC

X1             -> C1   define h := eta ln N − delta Q − C. Then h < 0 forbids any
                       increment, h = 0 permits none, h > 0 caps any single-period
                       increment at r <= h.                                 ARITHMETIC

C1 at N=18     -> V1   eta=0.30, ln 18=2.890372, delta=0.06, Q=1.20, C=0.55 gives
                       h = 0.867112 − 0.072 − 0.55 = +0.245112.             ARITHMETIC

P1             -> C2   the fixed point is Q_inf = (eta F(u∗) − C)/delta, which is
                       NEGATIVE whenever eta F(u∗) < C. At u∗=0.25,
                       F = 1.658228, eta F = 0.497468 < 0.55, so
                       Q_inf = −0.875526.                                   ARITHMETIC

C1 + C2        -> C3   h > 0 and Q_inf < 0 are compatible: h uses the ceiling ln N,
                       Q_inf uses the realised F(u∗). Both must be computed.  ARITHMETIC

P1 + C2        -> C4   at the fixed point delta Q_inf = eta F(u∗) − C, so
                       h_inf = eta ln N − (eta F(u∗) − C) − C = eta(ln N − F(u∗)).
                       delta, C and Q cancel; and by P2, h_inf >= 0 always.  ARITHMETIC

P3 + P4        -> C5   feasibility on mean F(u) is <= feasibility on F(mean u),
                       always. On the eight periods of Section 5.3:
                       F(0.30875) = 1.832381, mean F(u) = 1.753333,
                       gap 0.079048 nats = 4.3140% of 1.832381.             ARITHMETIC

C5 + P1        -> C6   the gap propagates to (eta/delta) x gap
                       = (0.30/0.06) x 0.079048 = 0.395243 stored units,
                       independent of C.                                    ARITHMETIC

P5             -> C7   N_eff = 1 + (N−1)u equals N + 1 − N/PR exactly. At
                       w = (0.45, 0.25, 0.15, 0.08, 0.05, 0.02),
                       sum w^2 = 0.2968, PR = 3.369272,
```

```
6 + 1 − 6/3.369272 = 5.219200.                                   ARITHMETIC
```

Every line is checkable from its own stated inputs. C1 through C7 are arithmetic consequences of one definition of model form (P1), two cited results about the surplus functional (P2, P3), one standard inequality (P4) and one definition (P5). **Only C6's interpretation of `C` as a standing charge consumes P6, the single postulate with no external support**, and the inequality itself holds however `C` is generated.

What this chain does not establish. It does not establish that any group obeys P1. That is the modelling postulate of Section 2, it is not verified here, the one preregistered confirmatory test of it so far reported in this corpus did not meet its own declared margin [Curzi, 2026i], and the reflexive parameter it would need to recover is not identified from observation at all (§3.12). **A chain whose steps are ARITHMETIC and CITED delivers the mathematics with certainty and delivers nothing whatever about institutions.**

37.9.4 9.4 The verification instance, worked by hand

A reviewer is asked to check four short things rather than one long one.

Check 1 — the gate. `ln 18 = 2.890372`. Then `0.30 x 2.890372 = 0.867112`; `0.06 x 1.20 = 0.072`; `0.867112 − 0.072 − 0.55 = 0.245112`. Rounded, `h = +0.2451`. Two multiplications and two subtractions.

Check 2 — the critical stock. Setting `h = 0` and solving for `Q` gives `Q* = (eta ln N − C)/delta = (0.867112 − 0.55)/0.06 = 0.317112/0.06 = 5.2852`. Unit B at `Q = 6.90` is above it, which is why a successful unit can be the one that may not build.

Check 3 — the negative equilibrium behind a positive gate. `F(0.25) = ln(1 + 17 x 0.25) = ln 5.25 = 1.658228`, which is `1.658228/2.890372 = 57.37` per cent of ceiling. Then `eta F = 0.497468` against `delta Q + C = 0.072 + 0.55 = 0.622`, so the step is negative, and `Q_inf = (0.497468 − 0.55)/0.06 = −0.875526`. **A positive gate and a negative equilibrium, from the same four parameters.**

Check 4 — the participation-ratio identity. For shares `0.45, 0.25, 0.15, 0.08, 0.05, 0.02`: the squares are `0.2025, 0.0625, 0.0225, 0.0064, 0.0025, 0.0004`, summing to `0.2968`, so `PR = 1/0.2968 = 3.369272`. Then `N + 1 − N/PR = 6 + 1 − 6/3.369272 = 7 − 1.780800 = 5.219200`. The even-share case gives `sum w^2 = 6 x (1/36) = 1/6`, `PR = 6`, and `6 + 1 − 6/6 = 6.0000` exactly. **Four lines of arithmetic establish that a stopwatch reading is an effective count.**

37.9.5 9.5 Limits, asymptotics, and the model's valid domain

- **`u -> 0`.** `N_eff -> 1`, `F -> 0`, and the surplus term vanishes: `Q_inf -> −C/delta < 0`. Perfectly aligned members purchase nothing, and the system runs down at the maintenance rate. **This is a real regime, not a degenerate one, and it is the regime a converged group is in.**
- **`u -> 1`.** `N_eff -> N`, `F -> ln N`, `h_inf -> 0`. A group collecting its entire ceiling has no remaining headroom at the fixed point: it is maximally productive and maximally unable to add. That is a constraint on the model's use, not a defect.
- **`N -> 1`.** `ln N = 0`, so `h = −delta Q − C <= 0` always: **a single member can never build under this gate.** The framework is silent on whether that is a statement about groups of one or an artefact of using `ln N` as the ceiling, and the silence is recorded rather than resolved. Section 4 nonetheless applies the same form at personal scope with `N` counting *framings* and not people, which is a different `N` and is labelled as such at every use.
- **`N -> infinity` at fixed `u`.** `F -> infinity` logarithmically while coordination paths grow as `N(N−1)/2` [Brooks, 1975]. The model's own arithmetic therefore predicts an optimum roster well below any crossover computed from surplus alone; that optimum is derived in Curzi [2026] and is not re-derived here.

- **`delta -> 0^+`.** `tau = 1/delta -> infinity` and `Q_inf` diverges. **At exactly `delta = 0` the fixed point does not exist at all** rather than being infinite: the map is an isometry, the geometric series does not converge, and the gate reduces to `h = eta ln N − C` with no dependence on stock [Curzi, 2026a]. **This value is outside the model's valid domain.**
- **`eta -> 0`.** `h -> −delta Q − C < 0` and `Q_inf -> −C/delta`. A system with no conversion cannot build at any `u` and any `N`. This is the exact correspondent of the forbidding clause `G = 0 => X = 0` in the channel-gain form of the belief postulate [Curzi, 2026ab], and the correspondence is stated in the source as a design parallel rather than derived.

37.9.6 9.6 Model comparison, and the standard this paper does not meet

The paper reports no fitted model of its own; all numerical results are computed from stipulated parameters. **That does not exempt it from the standard, because it recommends a protocol whose Section A.16 and B.13 steps require a fit.** The table below records the corpus's position against each requirement of the model-comparison standard, for the one confirmatory fit reported anywhere in this corpus [Curzi, 2026i].

Requirement	Status	What was actually done
Out-of-sample performance reported	**MET**	held-out RMSE reported against four baselines
Matched null with equal free parameters	**MET**	a matched AR(1); the storage-law fit beat it by 2.8 per cent against a declared 10 per cent margin, and **failed its own rule**
Information criterion (AIC or BIC)	**NOT MET**	**no AIC or BIC has been computed anywhere in this corpus.** Parameter count is therefore not paid for in any reported comparison
Number of free parameters stated explicitly	**PARTIAL**	the storage law's four parameters are named; the comparison does not tabulate them against each baseline's
Effect sizes with intervals	**NOT MET**	percentage RMSE advantages are reported as point estimates without intervals
Shuffle control's false-positive rate measured	**MET, AND ADVERSE**	measured 0.108 against nominal 0.05 for i.i.d. permutation, 0.120 for circular shift [Curzi, 2026u] — **anti-conservative**
Structural ceiling test run alongside the empirical baselines	**MET**	and it is the only test that caught the repackaging described in Section 6.5

The two NOT MET rows are printed here rather than omitted. A protocol that instructs a group to beat four baselines and never asks it to pay for its parameter count is incomplete, and a reader deciding whether to adopt Appendix A.16 should know that.

37.9.7 9.7 Scale type, and the operations the data license

Quantity	Scale type	What licenses it	Operations permitted here
speaking-time share `w_i`	ratio	duration has a true zero and a unit	sums, squares, `PR`
`u` from pairwise disagreement	ordinal at best, treated as ratio	**nothing tested**	used as a ratio quantity; **not licensed**, see Section 10
`N`, `N_eff`, `PR`	ratio (counts)	counting	all
`Q`	interval at best in a declared proxy	proxy definition only	differences and rates; ratios across proxies are **not** licensed
`D`, `P`	ordinal	no zero, no unit	ordering only; `D/P` **demoted** [Curzi, 2026t]
the three-verdict tally	ordinal	construction	counts of verdicts; **no** conversion to a magnitude

Michell's critique is the strongest single objection to this framework and it is conceded here. Michell [1997, 1999] argues that psychology systematically assumes its attributes have quantitative structure without ever testing that they do. The corpus's founding move — treating comfort as a magnitude — is exactly the move Michell identifies. The formal test that would settle it is conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971]: the double-cancellation conditions determine whether an attribute admits additive quantitative structure at all. **Those conditions have not been tested for comfort, for `u`, or for perspective and perception, by this work or by anyone.** The corpus's response is partial and is stated as partial: it demotes the one quantity whose operation most obviously needs a ratio scale, and it prices the Rasch route out at 123 items per term [Rasch, 1960, Andrich, 1988], an instrument with a defensible zero but not a conjoint test. **Where a reader wants the conjoint test, this paper concedes that it was not run.** Reliability is likewise not a substitute for structure [Sijtsma, 2009], and construct validation [Cronbach and Meehl, 1955, Campbell and Fiske, 1959] would establish that an instrument measures something without establishing that the something is quantitative.

37.9.8 9.8 Invariance, and the precision of the numbers

What is invariant. The gate's verdict is invariant under a common rescaling of `Q`, `C` and `eta` — multiply all three by a positive constant and the sign of `h` is unchanged, because every term scales together. It is **not** invariant under rescaling `Q` alone, which is what happens when a group changes its order proxy mid-measurement; that is a real operational hazard and Appendix B.2 requires the proxy to be frozen. The effective-count identity `N_eff = N + 1 − N/PR` is invariant under permutation of the share vector and under any relabelling of members, and it is **not** invariant under changing `N` — adding a silent member to the roster raises `N` and lowers every share, so the two sides move together only if the silent member is counted in both. **The scale-freedom claim — that one functional form serves at personal and group scope — has been tested at exactly two scopes in this paper, and at three in Curzi [2026], where it is found to break in five independent ways.** Two or three instances is not an invariance result and this paper does not present it as one.

Precision. Every number in this paper is either an exact consequence of stipulated parameters or a count. `ln 18 = 2.890372`, `ln 7 = 1.945910` and `ln 5.25 = 1.658228` are exact to the digits shown and are quoted to four or six decimals because they are computed, not measured. The parameters they are computed from — `eta = 0.30`, `delta = 0.06`, `C = 0.55` and the

rest — are **stipulated illustrative values, not measurements**, so a result such as `Q_inf = −0.8755` is exact given the stipulation and carries no empirical uncertainty because it carries no empirical content. **No quantity in this paper is a fitted value, and therefore no quantity in this paper is quoted with an uncertainty interval; where a reader sees four decimals they should read "consequence of the stipulated parameters", not "measured to that precision".** The two places where a genuinely measured external number appears — the roughly 25 per cent tipping point [Centola et al., 2018] and the collective-intelligence factor [Woolley et al., 2010] — are quoted as orders of magnitude with their settings named, and neither is used as a point estimate in any computation.

37.10 10 Limitations

L-1 The reflexive parameter is not identified, and this is the paper's most serious unmet standard. The specification underlying the framework is linear-in-means reflexive, and its endogenous effect cannot be separated from contextual effects or correlated unobservables by any amount of observational data [Manski, 1993, Shalizi and Thomas, 2011]. This paper adopts no remedy — no randomisation, no instrument, no partial-population variation, no defended stipulation of the reference group. **Every claim here about what a protocol does to a group is therefore a modelling position rather than an identified empirical result**, and a reader who needs the identified version must run the randomised design the framework has not run.

L-2 The conjoint conditions have not been tested. Comfort, `u`, perspective and perception are treated as quantities. Michell [1997] holds that this is precisely the untested assumption of psychological measurement, and Luce and Tukey [1964] supplies the test. **It has not been run for any construct in this corpus.** The ratio-scale treatment of `u` in particular is unlicensed: `u` is estimated in Section 4.4 by counting how many of `N(N−1)/2` framing pairs have ever disagreed about a next action, which is at best an ordinal instrument, and it is then used inside a logarithm and multiplied by `eta`. **That is an operation the scale does not license, it is disclosed here, and the paper's own instrument note calls the estimator crude and fit only to detect the collapse at the left end of the range.**

L-3 Every parameter is stipulated. No `eta`, `delta`, `C`, `Q`, `N` or `u` value in this paper was measured. The worked systems are constructed examples with declared parameters, chosen to exercise the arithmetic, and the numbers they produce are exact consequences of those choices and nothing else. The corpus-wide audit finds the same thing at three scopes and states it as its central result: of 27 parameter-by-scope cells, 19 have a defined estimator, 1 has a validated value and **0 carry a first-party measurement** [Curzi, 2026].

L-4 No AIC or BIC has been computed anywhere in this corpus, so no comparison reported here or cited by this paper pays for its parameter count, and no effect size in the confirmatory fit carries an interval (Section 9.6).

L-5 The eight-coordinate decomposition is a taxonomy, not a theorem. Thirty entries are named because thirty routes were found, not because thirty is a count of anything. A fault outside the eight coordinates will not be named by the catalogue, and the correct response is to say so rather than force a coordinate.

L-6 The mapping from `N_eff` to a count of independent voters is a convenience. Condorcet's theorem assumes independence outright; `N_eff` audits redundancy. The table in Section 5.6 rounds an effective count to the nearest odd integer and reads off a majority accuracy, and that stitch carries direction and order of magnitude only. **It should not be quoted as a computed probability for any real panel.**

L-7 The critical-mass figure is an experimental result in an artificial setting. The roughly 25 per cent tipping point [Centola et al., 2018] comes from controlled online experiments with confederates committed by construction. Its transfer to a field organisation is a modelling assumption. It is used here as an order-of-magnitude floor and never as a guarantee, and the same caution applies to the complex-contagion seeding results [Centola and Macy, 2007, Centola, 2010].

L-8 The acoustic-to-collective-intelligence causal claim is not confirmed and is carried as not confirmed. The correlational finding is real [Woolley et al., 2010]; the causal step is not established [Curzi, 2026f]; and the direction of causation between turn-taking equality and collective performance is not settled by the cited work.

L-9 The biological arm has no established mechanism and an expected null. There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal [Curzi, 2026y]. Every step touching it is graded R. **No clinical, medical or therapeutic claim is made anywhere in this paper; no sequence is changed, no molecule is delivered, and no capability is conferred on anyone.**

L-10 The functional form is underdetermined. A sub-linear power law [Latané, 1981] and an `n log n` form [Briscoe et al., 2006] fit the qualitative facts as well as the logarithm does, and nothing in this paper discriminates among them (§3.12).

L-11 The shuffle control is anti-conservative as ordinarily run — a measured false-positive rate of 0.108 against a nominal 0.05 [Curzi, 2026u] — so a `p = 0.0005` against it is worth less than it appears, which is exactly the reading the repackaging case demands.

L-12 Two discrepancies in the source material are recorded rather than silently corrected. First, the scaffold specification for this paper describes the participation-ratio example as "a meeting in which one person takes 45 per cent of the air counts 5.22, not 18"; **the source's worked example is a six-person meeting, so the correct contrast is 5.2192 against a naive 6, not against 18.** The arithmetic in Section 5.4 is right and the summary was wrong. Second, the same specification states that "raising non-parallelism from 0.25 to 0.45 moves [the equilibrium] to +3.2878"; **the source reaches +3.2878 only after also cutting `C` from 0.55 to 0.45. Raising `u` alone gives +1.6211.** Section 7.2 separates the two levers explicitly for that reason.

L-13 The Rasch price is an estimate of an instrument's length, not of its success. 123 items per term is what the recovery calculation returns; whether a Rasch-conformable item set for perspective or perception can be written at all is untested, and fit to the Rasch model would establish specific objectivity without establishing that the attribute is quantitative in Michell's sense.

L-14 Seven works are named in prose without bibliographic keys, because they are not in this corpus's external bibliography. They are Condorcet's 1785 *Essai* (Section 5.6), Levine, Gordon & Fields 1978 and Kaptchuk et al. 2010 on placebo analgesia and open-label placebo (Sections 3.2, 4.2 and 4.12), Tajik et al. 2016 on force transmission to chromatin (Sections 4.12 and 7.3), Church, Gao & Kosuri 2012 and Zetsche et al. 2015 on sequence encoding and self-processing (Sections 3.3, 4.7 and 5.10), and the predictive-processing programme of Friston, Clark and Hohwy (Section 4.12). Each is a real published result and each is quoted for what its source measured; **the reader is owed the note that they are named rather than referenced, and that the corpus's own treatment of the placebo channel is Curzi [2026ab] and of the sequence-encoding chain is Curzi [2026y].** The accuracy table accompanying Condorcet's theorem is computed here from the binomial and can be checked directly.

37.11 11 Conclusion

The headroom identity `h = eta ln N − delta Q − C` is a precondition on action, not a diagnostic afterthought. Its three regimes are exhaustive: `h < 0` forbids building because no interaction geometry prevents decline there; `h = 0` permits maintenance only; `h > 0` caps every requested single-period increment at `r <= h`, which is a maximum and not a target. Worked at `N = 18` with `eta = 0.30`, `delta = 0.06`, `C = 0.55` and `Q = 1.20`, the gate returns `h = +0.2451`, and the critical stock at which it closes with everything else fixed is `Q* = 5.2852`.

A system can pass the gate and still be shrinking, and the two computations must both be run. The same eighteen-person unit realises `F(0.25) = 1.6582` nats, 57.37 per cent of its ceiling `ln 18 = 2.8904`, which gives `eta F = 0.4975` against a draw of `delta Q + C = 0.6220` and a **negative equilibrium `Q_inf = −0.8755`**. Raising non-parallelism from 0.25 to 0.45 alone lifts the realised fraction to 74.65 per cent and the equilibrium to `+1.6211`; cutting maintenance

from 0.55 to 0.45 as well lifts it to `+3.2878`. **The geometry lever decides whether a sustainable stock exists; the cost lever decides how large it is; neither substitutes for the other.** At the fixed point the gauge reduces to `h_inf = eta(ln N − F(u*))`, with `delta`, `C` and `Q` all cancelling, so a mature practice with no headroom has a non-parallelism problem and not a discipline problem.

Two measurement corrections are enforced throughout and both are one-directional. Feasibility must be tested on the mean of `F(U)` and never on `F` of the mean `U`, because `F` is strictly concave; on the eight worked periods the wrong computation overstates by **0.0790 nats, 4.31 per cent, worth 0.3952 stored units** at `eta = 0.30` and `delta = 0.06`, independent of `C`. And the effective count the law consumes is exactly the participation ratio: for shares `0.45, 0.25, 0.15, 0.08, 0.05, 0.02`, `PR = 3.3693` and `N_eff = 5.2192` against a naive count of **6.0000** — a six-person meeting has silently become a five-person meeting with a spectator. **The affine count `1 + (N−1)u` is not a Hill number**, and where a Hill number is meant this corpus uses `1/HHI` or `exp(H)` and says so.

The Moment equation stays demoted at both scopes. Division requires a ratio scale with a true zero; perspective and perception are not instrumented that way; so the limiting cases, the fixed point `D = P` and the three-way ordering over-framed / calibrated / under-framed survive, and the magnitude does not. The constructive route out is priced at **123 items per term**, 246 for the ratio and 369 for all three terms, and 61.5 person-hours for a thirty-participant recovery gate.

A single seeding result changes the design of most adoption programmes. Sixteen seeds spread across 120 people at mean degree 8 under a two-of-eight reinforcement rule leave only **29.0 per cent** of non-seeds with two or more seeded neighbours — 30.16 people — while **39.17 per cent** have exactly one, a below-threshold exposure that spends the practice's novelty and returns a null. The same sixteen concentrated in one twenty-person unit activate **4 of 4** at 80.0 per cent local commitment. **The dispersed plan optimises the first-period metric; the clustered plan optimises the order parameter.**

Finally, the catalogue and the counter-result. A manifestation claim has eight coordinates — channel, gain, magnitude, base rate, falsifiability, consent, ledger, scale — and thirty named failure modes across six classes, each priced by one of four identities, including `X = G B` with `G = 0`, which disposes of about a third of the entries in one line. **The collective worked example fails on the framework's own arithmetic**, returning a negative equilibrium despite passing the entry gate; **the acoustic-to-collective-intelligence causal test is carried forward as NOT confirmed**; and **a filtered repackaging of the input passed all four empirical baselines including the shuffled-input control at `p = 0.0005` and was caught only by the structural ceiling test**, its claimed jump being 38.5 times `r_max`. Those three are the statements downstream papers should quote rather than paraphrase.

37.11.1 Declarations

Data and code availability

No external dataset was collected for this paper, and no human participants were surveyed, observed or recruited. Every numerical result is computed from parameters stipulated in the text under the models developed in Sections 3 to 6, and every table can be re-derived from the equations and parameter values printed beside it. The personal and collective systems worked in Section 7 are **constructed illustrative examples, not field observations, and not reports of any real individual or organisation**; the diagnostic instance of Section 7.3 is likewise constructed. The figures are generated by the scripts named in each caption, which contain no data beyond the numbers printed in this paper and the model equations of Sections 3 to 5. The two external empirical results relied on — the collective-intelligence factor and the critical-mass tipping point — belong to their published sources and are cited to them.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, produced the worked examples and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The worked systems are constructed. The consent architecture of Sections 3.5 and 3.10 — the know, decline and reverse gates, and the independence gate at group scope — is a **design constraint on prospective work** and is not a report of any study performed. **No clinical, medical or therapeutic claim is made anywhere in this paper.** No sequence is changed, no molecule is delivered, nothing is edited or altered, and no capability is claimed to have been conferred on any person or group. The biological-layer arm described in Sections 4.7 and 5.10 is a research procedure with no established mechanism and an expected null result; it is described so that it can be evaluated, not recommended.

Preregistration

This paper reports no preregistered test of its own; its results are derivations from stipulated parameters. The preregistered protocol for testing the storage law against a real system, together with the confirmatory application that **failed** its second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared margin of 10 per cent — is reported in Curzi [2026i]. The measured false-positive rate of the shuffled-input control, **0.108 against a nominal 0.05**, is reported in Curzi [2026u]. Any applied claim resting on this paper inherits both results.

37.12 Appendix A The personal-scope specification

Numbered, checkable, runnable. **Gates are steps. A failed gate halts the protocol at that point; it does not generate a caveat.** Grade stamps are mandatory on every step that has one.

A.1 — DECLARE THE OCCASION CLASS. Name one recurring occasion to be studied. It must recur at least weekly and carry one observable outside the practitioner's full control. *Grade T (definitional).*

A.2 — GATE K (KNOW), first pass. Write, in the practitioner's own words, what this protocol does at each of its three layers, and the grade of each. If the L3 sentence cannot be written — *"this is a research procedure with no established mechanism and an expected null result"* — the protocol is not yet understood and it does not start. *Halt condition: cannot state.*

A.3 — GATE D (DECLINE), constitution. Record what stopping looks like, who is told, and what happens to the data if the run stops at week 3. **If any part of this protocol involves another person, they get their own Gate D, their decline is recorded as data and never as a zero, and their cost is counted in their ledger by asking them.** *Halt condition: no decline path.*

A.4 — INVENTORY THE FRAMINGS. List the `N` distinguishable framings of the occasion class. Be strict: two framings that always recommend the same next action are one framing. *Grade T.*

A.5 — ESTIMATE `u` AND COMPUTE THE CEILING. Count the pairs that have ever disagreed about what to do next; divide by `N(N−1)/2`. Compute `N_eff = 1 + (N−1)u`, `S = ln N_eff`, and the ceiling `ln N`. Report `S / ln N` as the realized fraction. *Grade T. The `u` estimate is crude and labelled crude.*

A.6 — FIX THE SESOI AND SAMPLE SIZE, IN WRITING, BEFORE ANY DATA. Choose the smallest effect on the observable that would change the following quarter's conduct; look up the required `n` in §4.9 for detection and for equivalence. **If the `n` is unaffordable, say so now and run as descriptive rather than inferential — an honest choice in week 0 and not in week 12.** *Grade T.*

A.7 — GATE V (REVERSE), first pass. Write the STOP / RESIDUE / IRREVERSIBLE lines for every step that is to be run. Any step with a non-empty IRREVERSIBLE line is escalated to a second gate K before it runs. *Halt condition: a step with unbounded residue.*

A.8 — BASELINE THE STORAGE PARAMETERS. Estimate `delta` (what one skipped period costs), `eta` (what fraction of engagement becomes durable), `C` (maintenance load) and `Q_0`, recording each as an estimate with its basis. *Grade T; the fit to a particular life is a badged assumption.*

A.9 — COMPUTE THE HEADROOM GATE. `h_0 = eta·ln N − delta·Q_0 − C_0`. **If `h_0 < 0`, the protocol does not add anything: go to A.18 (shedding) and return here afterwards. If `h_0 = 0`, maintain and do not build. If `h_0 > 0`, the maximum increment that may be requested is `r <= h_0`.** *Grade T (the storage law's feasibility corollaries [Curzi, 2026a]).*

A.10 — RUN THE PREDICTION LOG (L1). For every occurrence of the occasion class: prediction, observable, direction, before. Outcome, after. Verdict from the three-valued set only. *Grade C.* **No ratio is computed and no distortion `Delta` is reported (CC-R1). The verdict set is ordinal by construction.**

A.11 — TALLY ORDINALLY, WEEKLY. Count verdicts. Report as counts. Act on the dominant verdict using the lever table in §4.3. *Grade T for the ordering;* ***magnitude demoted****.*

A.12 — RUN THE CADENCE (L2). Fix the period and keep it. **The half-life to the fixed point in §4.5's construction is 8.31 weeks; expect that order and do not evaluate the practice before one time constant has passed.** *Grade T.*

A.13 — RECOMPUTE `h_t` MONTHLY. Headroom falls as `Q` grows even when nothing else changes. A practice that gets heavier is doing arithmetic, not failing. *Grade T.*

A.14 — GATE K (KNOW), second pass, before the biological arm. State, aloud or in writing: *"This step is Grade R. There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal. Nothing is delivered, nothing is edited, no capability is conferred, no clinical, medical or therapeutic effect is claimed, and the expected result is null."* **If that cannot be said and meant, the arm does not run.**

A.15 — CONSTRUCT THE ARTEFACT (L3), FULLY DISCLOSED. Text the practitioner wrote, encoded base-4, address-prefixed as a format, rendered per the disclosed sonification chain [Curzi, 2026y] within the Gabor constraint of §4.7. **Audible level. No masked layer. No unread content. The file is inspectable and the practitioner inspect it.** *Grade R. Constraint 1 (the disclosed-narrative constraint [Curzi, 2026ac]) is enforced by Gate K, not by good intentions.*

A.16 — RUN THE L3 ARM AS A PRE-REGISTERED EXPERIMENT. Pre-register the null hypothesis, the observable, the SESOI, the `n` from A.6, and the four baselines of the falsification suite (persistence, direct-input, matched autoregressive, and the shuffled-input control). **Add the structural ceiling test: any single-period increment claimed must satisfy `r_t <= h_t`, because the shuffle is an empirical audit and only the structural test detects repackaging** (R4). *Grade R.*

A.17 — REPORT THE COMPARISON, WHICHEVER WAY IT COMES OUT. Classify the result as N-1, N-2, N-3 or N-4 per §4.10 before interpreting it. **A null on the L3 arm is an N-4 and is reported as an N-4: expected, uninformative about the mechanism, and worth recording pre-**

cisely because results of this kind do not otherwise enter the record. *Grade T (classification).*

A.18 — SHED BEFORE ADDING WHEN `h <= 0`. Reduce `C`, or deliberately shed stored order (the shedding lever [Curzi, 2026a]). **Shedding that lands its cost on another person is not shedding; it is transfer, and it is outside the line this framework draws.** *Grade T, hard ethical boundary.*

A.19 — GATE D (DECLINE), standing. At every monthly review, "stop" is on the agenda as a live option with the same standing as "continue". A protocol never declinable in practice has not satisfied Gate D by having declared it in A.3.

A.20 — CLOSE THE LOOP WITH THE ECHO. `Delta_R = |ln(M'/M)|` (the returned-difference result [Curzi, 2026]): what came back over what was sent. **Zero means no information returned — a finding about the channel to the world, not about the effort spent.** The echo corrects `D`, not `Q`: feedback changes the integrator's input rather than its state, which is why character changes slowly. *Grade T for the structure;* ***magnitude demoted with the rest of the M-derived family (CC-R1)****, so read* `Delta_R` *as an ordering — bigger surprise, smaller, none.*

37.13 Appendix B The group-scope specification

Numbered, ordered, and gated. **A failed gate halts the protocol at that step. It does not generate a caveat.**

B.1 — DEFINE THE UNIT. Name the population, its `N`, and the boundary. **`N` counts people who will actually participate, not people on a distribution list.**

B.2 — GATE 0, HEADROOM. Estimate `eta`, `delta`, `C`, `Q`. Compute `h = eta ln N − delta Q − C`. **If `h < 0`, the group may not build: go to B.16 and return. If `h = 0`, maintain. If `h > 0`, record `r_max = h` as the ceiling on every increment below.**

B.3 — GATE 0b, THE PROGRAMME'S OWN COST. Recompute `h` with the proposed programme's maintenance cost added to `C`. **If that recomputation is negative, the programme may not run at its current cost, whatever its expected benefit.** This is the row nobody computes (§5.1).

B.4 — GATE D, CONSTITUTION, PER PERSON. Each participant is told what stopping looks like, what happens to their data if they stop, and who they tell. **A decline is recorded in that person's own row as a decline, never as a zero and never as a missing value to be imputed.** *Halt condition: no decline path exists.*

B.5 — GATE I, BASELINE. Measure the group's capacity to disagree before anything changes: independent written estimates on 15 to 25 live questions, collected before discussion. Compute mean pairwise correlation of signed error `rho`, then `u = 1 − rho^2`. Compute `N_eff`, `S`, and the realized fraction `S / ln N`. **Record turn-taking shares and `PR` in the same sitting.**

B.6 — APPLY THE JENSEN CORRECTION. Feasibility is tested on **mean `F(u)`**, never on `F(mean u)` (R2). The wrong computation always overstates; in §5.3's worked case by 4.31% of surplus and 0.3952 Q-units of steady state.

B.7 — READ BOTH NUMBERS. Report `h` (which uses `ln N`) *and* the realized fraction `S / ln N`. **A group can have comfortable headroom and still lose stored order every period because it operates far below its own ceiling.** A diagnosis reporting one without the other is half a diagnosis and points at the wrong lever.

B.8 — RAISE `u` BEFORE RAISING `N`. Concavity is front-loaded: at `N = 18`, moving `u` from 0 to 0.25 buys 57.37% of the ceiling. **Adding a parallel member adds `n(n−1)/2` coordination cost and zero surplus.**

B.9 — RECORD DISTRIBUTIONS, NOT CONCLUSIONS. Independent estimates before discussion, every time, on every question that matters. Publish the spread. **Unanimity that arrives**

fast and without argument is a symptom, not an achievement (the built-system failure catalogue [Curzi, 2026g], FM-01).

B.10 — SIZE THE SEEDING AGAINST THE SMALLEST WORKABLE POPULATION. Compute 25% of the smallest population within which the practice works standing alone. **Carry the honest note with the number: 25% comes from controlled online experiments with confederates committed by construction, and its transfer to a field organisation is a modelling assumption.**

B.11 — CLUSTER, DO NOT SPRAY. Concentrate the seeding budget in one population past its own tipping point. **Do not distribute for reach: a below-threshold exposure spends the practice's novelty and returns a null, and the second attempt on the same person starts worse than the first.**

B.12 — TEST FOR SPONSOR-INDEPENDENCE BEFORE EXPANDING. Remove the sponsor for one period. Count what survives. **What survives is the committed fraction; what stops was compliance, and compliance has `u = 0` with respect to the group's belief and contributes nothing to `S`.**

B.13 — RUN THE ORDINAL DOCTRINE-TO-DATA INSTRUMENT. Pre-record the group's prediction for a defined recurring occasion; afterwards assign exactly one of **over-framed / calibrated / under-framed**. Tally counts. **No `D/P` value is computed and no `Delta` is reported (CC-R1). The verdict set is ordinal by construction, and where a number is genuinely required the route is the Rasch-route theorem [Curzi, 2026t], at 123 items per term.**

B.14 — GATE K, PER PERSON, BEFORE ANY BIOLOGICAL-LAYER STEP. Each participant states, in their own words: *"This is Grade R. There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal. Nothing is delivered, nothing is edited, no capability is conferred, no clinical, medical or therapeutic effect is claimed, and the expected result is null."* **Anyone who cannot state it does not participate, and their non-participation is recorded as data.**

B.15 — RECORD INDIVIDUAL JUDGEMENTS BEFORE GROUP DISCUSSION, ALWAYS, AND ESPECIALLY HERE. A shared acoustic input is a shared input; eighteen people hearing one artefact together are receiving one input, not eighteen. **Pool afterwards or the independence needed to evaluate the artefact has been destroyed by the act of presenting it.**

B.16 — WHEN `h <= 0`, SHED BEFORE ADDING. Reduce `C`, or deliberately shed `Q` (the storage law [Curzi, 2026a] §6.9). **Shedding that lands its cost on another group or on individuals who did not consent is not shedding; it is transfer, and it is outside the line this framework draws.**

B.17 — GATE I, POST-STEP, AND THE REVERSAL RULE. Re-measure `u`, `N_eff` and turn-taking after any step that could align beliefs. **If `u` fell or the participating `N` fell, the step is reversed regardless of what else it produced.** A step that raised `Q` by lowering `u` borrowed from the ceiling to pay the floor.

B.18 — FALSIFY, INCLUDING STRUCTURALLY. Four baselines — persistence, direct-input, matched autoregressive, shuffled-input — **plus the structural ceiling test `r_t <= h_t`, because the shuffle is an empirical audit and only the structural tests detect repackaging** (R4). Report the comparison whichever way it comes out, and classify a null before interpreting it (Section 4 §4.10).

37.14 Appendix C Using the catalogue

How to use the catalogue. Six steps, in order, because running them out of order produces a confident misdiagnosis.

C.1 — LOCATE THE COORDINATE BEFORE NAMING THE ENTRY. Ask which of CH / G / Mg / B / Fa / Cn / Ld / Sc the complaint is about. **Most misdiagnoses are entries chosen by resemblance of symptom rather than by coordinate**, and the classes above exist so that the coordinate can be found first.

C.2 — RUN THE GROUNDING SEQUENCE AND STOP AT THE FIRST FAILURE. Identify the channel, check it is possible, estimate its gain, predict a magnitude, measure, compare. **A claim's fault is located at the first stage it cannot pass; everything downstream of that stage is uninformative about the claim.** Diagnosing a magnitude error in a claim that has no channel is a category mistake dressed as rigour.

C.3 — COMPUTE, DO NOT ESTIMATE, THE PRICE. Every entry's Price field is arithmetic. Run it on the numbers of the case at hand. **A fault that has not been priced will not be fixed, because nobody prioritises a number nobody has.**

C.4 — RUN THE STRUCTURAL TEST EVEN WHEN THE EMPIRICAL TESTS PASS. CC-R4's finding is binding: **the shuffle is an empirical audit, and only the structural tests detect repackaging.** Compute `r_max = eta ln N − delta Q − C` for the period and compare against the claimed increment, every time, regardless of how the four baselines came out.

C.5 — AUDIT THE LEDGER SEPARATELY FROM THE EVIDENCE. For every concordance, state which column carries evidence and which carries precedent, in print. **Where a row has precedent and no evidence, label it a research question with a long history rather than a finding.**

C.6 — CHECK CONSENT LAST AND TREAT IT AS A VETO. Consent faults do not trade against benefit. **A protocol that fails Class VI is not a worse protocol; it is outside what this framework will describe, and no result changes that.** Count the per-person gate records; if there are fewer than participants, consent was assumed rather than obtained.

Chapter 38

One Model Form at Three Scopes: The Instrument Audit, Five Breaks in the Correspondence, and a 156-Entry Status Register

Abstract. The same storage law is written at institutional, personal and cellular scope, and this paper badges that correspondence as a statement about the model form chosen at each scope rather than a claim that the three systems are alike or that quantities aggregate between them. It then audits the instrument rather than asserting the correspondence. **Across nine parameters and three scopes there are 27 parameter-by-scope cells: 19 have a defined estimator, exactly 1 has a validated value — a borrowed median protein half-life giving $\delta = 0.014955$ per hour — and 0 carry a measurement made first-party for this corpus.** Nothing in these seventy papers rests on fieldwork conducted for them. The correspondence is then shown to fail in five independent ways: a mind is not a small society, because the independence aggregation requires is unavailable inside one skull; a cell is not a small mind, because it has no frame and therefore no perspective ratio; the exit ethics do not descend; aggregation is not always well defined; and no composition theorem bridges the scopes. What survives is smaller and load-bearing. The arithmetic transfers wherever the map is adopted, including the fixed-point identity $h_\infty = \eta(\ln N - F(u^*))$, in which δ, C and Q cancel, so $h_\infty \geq 0$ always and the plain linear model never ossifies by itself. And three ordinal diagnostics — whether u is near zero, whether η is near zero, whether $\delta Q + C$ exceeds $\eta \ln N$ — read at every scope without a calibrated instrument. A complete register of 156 entries is published: 50 proved, 28 postulated, 58 borrowed, 12 speculative including 4 printed refusals, and 8 demoted, with precedent claimed in 23 rows, 14.7 per cent.

Keywords: multiscale correspondence; instrument audit; status register; participation ratio; headroom identity; epistemic bookkeeping

Series: B. Engineering the Surplus **Paper:** CCP-038

38.1 1 Introduction

The Creature Comforts corpus writes one equation at three scopes. The storage law

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t,     S_t = F(U_t) = ln(1 + (N−1) U_t)
```

appears with identical structure for an institution, for one mind, and for one cell's regulatory state, with only the parameter values, the period length and the choice of order proxy differing. That progression has been built up piecewise across the series without ever being stated in one place with its limits attached. **This paper states it, badges it, prices it, and then spends most of its length attacking it.**

The badge comes first because everything else depends on it. **The correspondence is a claim about model form and nothing else.** It is not a scaling law: it predicts no relation between a parameter's value at one scope and its value at another. It is not a composition theorem: nothing here derives institutional stored order from personal stored order, or personal from cellular, and **no such theorem exists anywhere in this corpus.** It is not an empirical finding: scale-freedom has never been tested at two scopes simultaneously in any real system. It is the observation that a minimal open-system accumulator written for each of the three uses the same five symbols, and that the arithmetic consequences of those five symbols therefore apply to all three.

The paper's principal contribution is not the correspondence but the audit of the instrument that would be needed to use it. Nine parameters at three scopes give twenty-seven parameter-by-scope cells. Counting them:

> **Nineteen of twenty-seven cells have a defined estimator. One of twenty-seven has a validated value — a borrowed median protein half-life, giving `delta = 0.014955` per hour. And zero of twenty-seven carry a measurement made first-party for this corpus.**

That last number is the one this paper exists to publish. **Nothing in this corpus rests on fieldwork conducted for it.** Every empirical quantity anywhere in the seventy papers is either borrowed from a study conducted by other people for their own purposes, or computed from parameters stipulated in a worked example. The count is not an accusation from outside; it is the result of the corpus auditing its own tables, and it is stated here so that other papers can cite a source for it rather than restating it on their own authority.

The correspondence is then shown to fail in five independent ways, and the ways are genuinely independent — repairing one does not touch the others. **A mind is not a small society**, because the independence that every aggregation result requires is unavailable inside one skull: two framings held by one person share a memory, a mood, a set of priors, a stake in the answer and a substrate, which are five channels of common-mode coupling with no institutional analogue. **A cell is not a small mind**, because the moment equation requires a representation and a cell has states rather than representations, so it has no frame and therefore no perspective ratio. **The exit ethics do not descend**: a member can leave, a framing can be abandoned, a cell cannot, so the consent constraint that governs the corpus attaches to the person and never to the scope. **Aggregation is not always well defined**, in principle rather than merely in practice. And **there is no composition theorem** bridging any two scopes.

What survives those five is smaller than the correspondence table and considerably more trustworthy. **The arithmetic transfers wherever the model form is adopted** — contraction, unique fixed point, the `ln N` ceiling, the headroom inequality, the fixed-point identity `h_inf = eta(ln N − F(u*))` — because these are consequences of the map rather than facts about any system. And **three ordinal diagnostics read at every scope without a calibrated instrument**: whether `u` is near zero, whether `eta` is near zero, and whether `delta Q + C` exceeds `eta ln N`. Each is a comparison rather than a magnitude, and a comparison survives where a magnitude does not.

The paper closes with a complete status register of the corpus's substantive claims: **156 entries — 50 proved, 28 postulated, 58 borrowed, 12 speculative of which 4 are printed refusals, and 8 demoted** — with the entry for the corpus's own keystone formula among the demotions. Precedent is claimed in **23 of the 156 rows, 14.7 per cent**; the remaining 85.3 per cent read *none claimed*, which is the honest yield of a two-column method when the rule against filling the second column is actually enforced.

Two identities established in the surrounding literature are re-verified here because they

are the only rows of the correspondence table that hold everywhere. The affine effective count equals the Rényi-2 participation ratio exactly — `N_eff = 1 + (N−1)u = N + 1 − N/PR` — re-verified over six random correlation ensembles with a worst discrepancy of `1.8e−15`, the identity established with its numerical verification in Curzi [2026c]. And at the fixed point the headroom gauge reduces to `h_inf = eta(ln N − F(u*))` with retention, cost and stored order all cancelling; since `F(u) <= ln N` always, `h_inf >= 0` always, so **the plain linear model never ossifies by itself.**

A standing note on which effective count is meant. `1/HHI` and `exp(H)` are Hill numbers of orders two and one and are well grounded [Hill, 1973, Jost, 2006]; the affine count `1 + (N−1)u` is **not** a Hill number but a linear interpolation between 1 and `N`, adopted as a design axiom [Curzi, 2026]. The participation ratio it equals *is* a standard effective count — the reciprocal of a Rényi-2 collision probability [Rényi, 1961] — and the identity between them is algebra, not evidence for the design axiom. Every use below says which count it means.

Sections 2 and 3 give the hypotheses and the audit method, including the identification problem that bounds the correspondence. Section 4 builds the parameter table cell by cell and prices it. Section 5 is the five breaks. Section 6 is the register. Section 7 works one laboratory at all three scopes and audits one sentence end to end. Sections 8 to 11 discuss, verify, bound and conclude.

38.2 2 Hypotheses

This paper is in large part an audit, and an audit of one's own instrument is not a test of a hypothesis about the world. Two of its four central results are arithmetic identities, one is a census, and one is a set of structural refusals. The claims below are stated as hypotheses only where something could turn out otherwise; where a result is a definition, an identity or a count, it is labelled as such rather than dressed as a prediction.

> **H1 — The molecular proxy for cellular stored order is refuted by the corpus's own measured parameter.** If cellular `Q` were denominated in molecular inventory, the storage law would predict that nothing cellular persists on any timescale a person would notice. **Predicted:** at the median protein half-life of 46 hours, the fraction of the inventory surviving one month is `2^(−720/46) = 1.942e−05`, a per-month retention loss of 0.99998. **Falsified by:** a measured cellular state whose persistence over months is carried by the molecular inventory itself rather than by a maintained regulatory pattern that keeps rebuilding it — that is, by long-lived molecules at population scale, which the measured half-life distribution does not contain.
>
> **H2 — Positive headroom does not imply a non-decreasing step, at any scope.** The gauge uses the ceiling `ln N`; the step uses the realised surplus `F(u*)`. **Predicted:** the worked laboratory member of Section 7.1 has `h = +0.054831` per month and a step of `−0.202575` per month simultaneously, because the surplus needed for a flat month is `(delta Q + C)/eta = 1.426667` nats against an internal surplus of `0.751416`. **Falsified by:** a demonstration that the two signs agree across systems, which would make one of the two computations redundant.
>
> **H3 — Retention rates are not comparable across scopes without conversion, and the error is large.** A rate quoted without its period is not a number. **Predicted:** `delta = 0.15` per year is `0.013452` per month, and `delta = 0.08` per month is `0.632334` per year; the parameter that looks four times larger is 47 times larger once the clocks agree. **Falsified by:** nothing observational — this is arithmetic on the definition of a per-period fraction, and it is listed here because it is the most common practical error the correspondence invites, not because it is in doubt.

Four claims of this paper are not hypotheses, and are labelled here rather than left to be discovered.

- **The three-scope correspondence is a POSTULATE about model form.** It is not falsifiable in general, only at each scope separately, and it is badged at every appearance.
- **The 19/27, 1/27 and 0/27 counts are a CENSUS**, not a finding about the world. They are recomputed from the corpus's own tables and can be recounted by a reader from the register in Section 6.
- **The five breaks are STRUCTURAL ARGUMENTS**, three of which turn on category distinctions — representation, exit, aggregation — that no measurement would settle.
- **The two identities re-verified in Section 9 are ARITHMETIC.** `N_eff = N + 1 − N/PR` and `h_inf = eta(ln N − F(u*))` are consequences of definitions and carry no empirical weight.

The organising question, stated plainly, is this: what is a multiscale framework entitled to claim when it has never measured any of its own parameters? The answer this paper defends is: the arithmetic, the ordinal diagnostics, and the failure catalogue — and nothing that requires a calibrated instrument.

38.3 3 Method: scopes, clocks, statuses, and what cannot be identified

38.3.1 3.1 The statement, with its badge

> **THE THREE-SCOPE CORRESPONDENCE [POSTULATE — MODEL FORM ONLY].** The storage law
>
> ```
> Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t          S_t = F(U_t) = ln(1 + (N−1) U_t)
> ```
>
> is written with identical structure at the institutional scope (the institutional-scope material of the corpus), the personal scope (the personal-scope material of the corpus), and the cellular scope (the cellular-scope material of the corpus). Only the parameter values, the period length, and the choice of order proxy differ. **This is a statement about the form of the model chosen at each scope. It is not a claim that the three systems are the same kind of thing, that their parameters are related by any law, or that quantities at one scope aggregate into quantities at another.**

Three things follow immediately, and all three are constraints rather than permissions.

First, nothing transfers by correspondence. A result established at one scope is not established at another because the symbols match. Woolley's collective intelligence factor was measured in laboratory groups; it says nothing about a cell and nothing about the internal structure of one person's framings. **The symbol `u` appearing in three places does not make one measurement of `u` count three times.** This is the two-column discipline of CCP-029 applied to scope instead of source.

Second, the units do not compose. Institutional `Q` is denominated in whatever order proxy Protocol MP-6 selected. Personal `Q` uses a different proxy on a different clock. Cellular `Q`, if it can be denominated at all, is a third. **There is no conversion factor between them, this manual does not supply one, and any arithmetic that adds or divides across scopes produces a unit artifact.** §7.1 computes one deliberately so the reader can see what it looks like.

Third, the correspondence is weakest exactly where the corpus's ambition is highest. The institutional scope has the most borrowed empirical support and the least first-party measurement; the personal scope has a century of psychophysics behind its measurement register and a demoted keystone formula in front of it (CC-R1); **the cellular scope has one parameter with a measured value and an unbridged gap where its input path should be.** That ordering is not hidden: it is the reason CCP-031 exists.

38.3.2 3.2 What “the same instrument” is licensed to mean

Precision here is worth a page, because the whole chapter turns on it.

Reading	Statement	Status
Form identity	The same five-symbol recursion is used to model all three	**What is claimed.** A modelling choice, badged
Arithmetic transfer	The consequences of the recursion (fixed point, contraction, ceiling, headroom) hold at any scope where the recursion is adopted	**Follows.** These are theorems about the recursion, not about the system
Parameter transfer	A value of `delta` or `eta` measured at one scope informs another	**NOT claimed.** No such relation is derived or measured
Scaling law	Parameters vary with scope according to some rule	**NOT claimed.** No rule is offered
Composition	Cellular `Q` sums to personal `Q` sums to institutional `Q`	**NOT claimed, and no theorem exists.** §5.6
Ontological identity	A mind is a small society; a cell is a small mind	**REFUSED.** §5.2, §5.3

The second row is the whole practical value of this section. Banach’s fixed point theorem does not know what `Q` denotes. If `delta` is in `(0,1]` and the map is the affine one written above, the map is a contraction, a unique fixed point exists, and iteration reaches it — at a desk, in a department, or in a regulatory network. The theorem holds because it is a theorem about the map. **What is not automatic is that the map describes the system**, and that question must be settled separately at every scope, and has been settled at none of them.

38.3.3 3.3 The three scopes, defined with their clocks

A scope is not a size. It is a triple: what counts as the system, what counts as one period, and what observable stands in for stored order.

```
SCOPE          SYSTEM                  PERIOD          ORDER PROXY Y
-------------------------------------------------------------------------
INSTITUTIONAL  a group of people       a quarter,      documented decisions,
               with a shared purpose   a year          working practices,
                                                       trained members, artifacts

PERSONAL       one mind and the        a day,          retained skills, habits,
               life it maintains       a month         standing commitments,
                                                       what survives a bad week

CELLULAR       one cell's regulatory   minutes to      maintained transcriptional
               state                   hours           and chromatin state
                                                       -- SEE §4.6 AND §5.6
-------------------------------------------------------------------------
```

The period is the part practitioners get wrong. `delta` is a fraction *per period*, and a `delta` quoted without its period is not a number. Fifteen percent annual staff turnover and eight percent monthly skill decay are not comparable until they are on one clock.

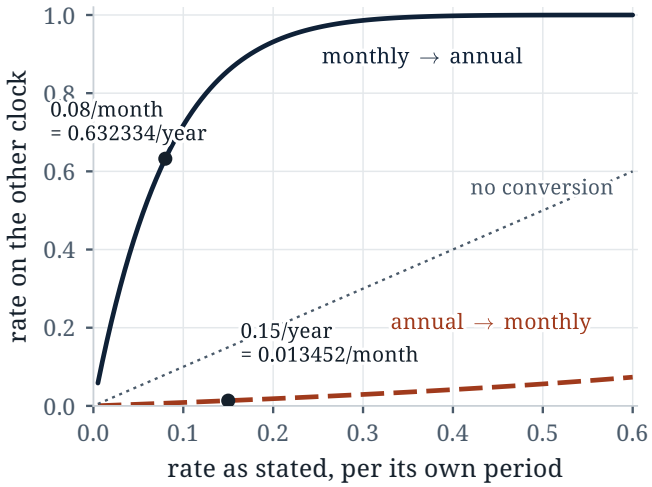


Figure 38.1: **A retention rate quoted without its period is not a number: the parameter that looks four times larger is forty-seven times larger once the clocks agree.** A rate of 0.15 per year is 0.013452 per month; a rate of 0.08 per month is 0.632334 per year. Conversion by `delta_new = 1 − (1−delta_old)^(p_new/p_old)`, arithmetic on the stipulated rates of Section 3.3; generated by `figs/src/CCP-038_fig02.py`.

Computed, so the size of the effect is visible:

```
delta = 0.15 per year      ->  monthly equivalent  0.013452
delta = 0.08 per month     ->  annual  equivalent  0.632334
```

A parameter that looks four times larger is in fact forty-seven times larger once the clocks agree. Every cross-scope comparison in this section carries its clock for that reason.

38.3.4 3.4 The calibration status of each scope, stated before the table is built

The table in §4.1 is a table of *meanings and intended estimators*, not of measured values, and the difference must be established before it is read — a well-formatted table is the most persuasive object in a technical book and the least deserving of that persuasion.

Scope	Parameters with a defined estimator	Parameters with a validated estimator	Parameters with a published first-party measurement
Institutional	9 of 9 (the measurement protocols MP-1 to MP-7 of the reference apparatus [CCP-039])	0	0
Personal	9 of 9, plus `M` and `D/P` demoted to ordinal (CC-R1)	0	0
Cellular	1 of 9	1 of 9	`delta` only, and borrowed

Twenty-seven parameter-scope cells; one of them has a number in it that anyone has measured, and that number was measured by molecular biologists for their own reasons. That

is the honest state of the instrument, and it is stated here rather than in a footnote because everything downstream inherits it.

The single measured cell is worth naming now, because it recurs. Schwanhäusser and colleagues (*Nature*, 2011) quantified mRNA and protein half-lives genome-wide in a mouse fibroblast line using metabolic labelling with mass spectrometry and sequencing, reporting median half-lives of roughly 9 hours for mRNA and roughly 46 hours for protein. Converted to the storage law's parameter:

```
median mRNA,    t_half = 9 h    ->  delta = 1 - 2^(-1/9)    = 0.074125 per hour
median protein, t_half = 46 h   ->  delta = 1 - 2^(-1/46)   = 0.014955 per hour
                                ->  delta = 1 - 2^(-24/46)  = 0.303468 per day
```

Those are real measurements of a real decay rate, they are medians over thousands of species with enormous spread, they were obtained in one immortalised cell line by one method, and they are the only calibrated dial on the third instrument. CCP-030 uses them for the channel analysis; this section uses them in §4.7 and §5.6 to make a structural point that turns out to be inconvenient for the naive reading of the correspondence.

38.3.5 3.5 The five statuses, defined so they cannot be blurred

Status	Definition	What it entitles a user to say	What it does not
PROVED	A theorem, true within stated assumptions. The assumptions are mathematical, not empirical	"This follows"	Nothing about the world. A theorem about a map is not a fact about any system the map is applied to
POSTULATED	An assumption of this research programme: chosen for tractability, interpretability, or design reasons	"This is what we assume, and here is why"	"This is established." Every number computed downstream of a postulate inherits the postulate's status, not a grade higher
BORROWED	An empirical result obtained by other researchers, in a named setting, with a stated method	"This was measured, there, that way"	"Therefore it holds here." Transfer is an assumption and must be named as one
SPECULATIVE	A proposed bridge or extension that has not met the three bridge criteria	"This is worth investigating"	Anything else. The marking must appear in the same sentence as the claim
DEMOTED	A claim this corpus advanced and its own protocol rejected or reduced	"The structure stands; the magnitude does not"	The demoted content, in any form, including in summary tables

Two rules govern movement between statuses, and both are one-directional.

Rule 1: nothing is promoted by agreement. A POSTULATED claim does not become PROVED because a tradition also arrived at it, and a SPECULATIVE bridge does not become established because it is attractive or because the corpus depends on it. Promotion requires the specific thing the status demands: a proof, a measurement, or a passed gate.

Rule 2: demotion is immediate and travels. When a gate fails, every entry downstream of the failed one is demoted in the same pass. CC-R1 demotes four formulary entries in the personal-scope material of the corpus's personal formulary and their counterparts in the interpersonal formulary, and it does so in print, in the formularies themselves, and not only here.

38.3.6 3.6 How to use the register against this corpus

1. Locate the claim in doubt. The status index in §6.33 gives the ID.
2. Read its STATUS. **If it is PROVED, the only available attack is on the assumptions, and they are listed in Group B.** If it is POSTULATED, the attack is on the choice, and the entry says what the choice bought. If it is BORROWED, read the "what it measured" field: most of the objections worth making are that the setting does not transfer, and the field is written to make that objection easy.
3. Check whether any entry in Group B is upstream. Almost everything is downstream of B-01 (the affine design axiom) or B-02 (the linear storage law).
4. If the objection survives all three steps, it is a real objection, and §6.36 says what this house undertakes to do with it.

38.3.7 3.7 Identification: the parameter the correspondence would need, and cannot have

The three-scope correspondence is a claim about model form. The moment it is used — the moment anyone fits `eta` or `delta` to a real group and reports the fit as a property of that group — a second and much harder problem arrives, and it is not solved anywhere in this corpus.

The framework's substantive social claim is reflexive: the group's state influences each member and the members constitute the group. Written as a model that is a **linear-in-means reflexive specification**, and Manski [1993] proves that its endogenous social effect is **not identified from observational data.** It cannot be separated from contextual effects, nor from correlated unobservables, however many periods are collected. Shalizi and Thomas [2011] show the same for network data: influence and selection are generically confounded.

Applied to this paper, the consequence is specific and it sharpens the audit rather than softening it. The 19 cells with defined estimators are estimators of *conditional summaries* of a fitted series. Nineteen procedures exist; each returns a number; and for the parameters that carry social content — `eta`, `u`, and any `delta` with a voluntary-departure component — that number is not identified as a coupling even in principle from observation alone.

The known remedies are randomisation, an instrument, partial-population variation, or stipulating the reference group and defending the stipulation, the last on the template of Brock and Durlauf [2001]. **This corpus adopts none of them, and no paper in it reports a randomised or instrumented design.** Where the correspondence is used to license a claim about a real group, that claim is a **modelling position rather than an identified empirical result**, and Section 10 records this as the audit's most serious finding after the 0-of-27 count itself.

Two objections in the same family, also unanswered. First, the Lucas critique applied to the corpus's own parameters: Lucas [1976] holds that behavioural parameters estimated under one regime need not survive a change of regime, and Hacking [1995] names the mechanism directly in the human sciences — a classification changes the people classified. The framework asserts that systems change their behaviour when modelled, and then treats `eta`,

`delta` and `u` as stable enough to fit. **Either those parameters survive the framework's own predicted response to being measured, or the framework undercuts itself.** This paper takes the first horn as an assumption and labels it an assumption. The one documented field case of a reflexive fixed point, with dates and price data, is MacKenzie and Millo [2003]; it is an existence proof that the mechanism is real and not a licence to read a fitted coefficient as a measurement.

Second, the functional form of the surplus term is underdetermined by the evidence. A sub-linear power law [Latané, 1981] and an `n log n` form [Briscoe et al., 2006] fit the qualitative facts as well as `ln(1 + (N−1)u)` does, and **nothing in this paper discriminates among them.** The logarithm here follows from the affine-effective-count design axiom [Curzi, 2026] and from nothing else established in this corpus. Its usual warrant is Fechnerian [Fechner, 1966], which requires the perceptual story to be true; the stronger warrant, requiring no psychological assumption at all, is Peters and Gell-Mann [2016], under which the time-average growth rate of a multiplicatively evolving quantity is the expectation of its logarithm. **Where a stock compounds rather than adds, that is the warrant the argument actually needs.**

38.3.8 3.8 Analogy discipline, and what is borrowed from where

Every borrowing in this paper is one of three things, labelled at first use.

- **A shared mathematical form.** The storage law is a first-order discrete filter; its pole, DC gain, time constant and settling behaviour follow because the equation is the same equation, with no claim that the physics transfers [Oppenheim et al., 1996, Åström and Murray, 2008]. Geometric convergence follows from the contraction property [Banach, 1922, Elaydi, 2005], and the leaky integrator has the same status in computational neuroscience [Gerstner and Kistler, 2002]. **This is the category the correspondence sits in, and it is the respectable one.** The founding definitional move — comfort as the felt total of present costs, with the variational statement given as stationarity rather than minimality and licensing no teleology — is Curzi [2026] and is presupposed rather than restated.
- **A derivation.** The participation-ratio identity is algebra [Curzi, 2026c]. Jensen's inequality applies because `F` is strictly concave, which is proved [Curzi, 2026]. The range result `0 <= F(u) <= ln N` is proved and is what makes `h_inf >= 0` unconditional.
- **An analogy, and never evidence.** Reading a group of framings as an assembly of agents is an analogy; Section 5.2 refuses it. Reading programmed cell death as *exit* in Hirschman's sense [Hirschman, 1970] is an analogy; Section 5.4 refuses that too, and gives the reason.

Two standing bounds from the foundations are inherited rather than restated, and both restrict what a borrowing may claim. Curzi [2026] declines the inflated reading of its own borrowed physics: the framework's postulate is stated as a postulate, the uncertainty relation is given only in its correct form, and a named failure mode requires that a user be able to distinguish the theorem, the ordinary observer effect, and the framework's own conjecture in three separate sentences on demand. Curzi [2026] does the same for the contraction theorem: a human life is not a metric space, nobody has measured the relevant contraction modulus, and the theorem therefore transfers as **shape rather than guarantee** — and the feeling of recognition on arrival proves nothing, since the machinery that rehearses a future retrofits a past with equal fluency. **Both bounds apply directly to the claim that the arithmetic transfers, which is the only claim Section 5.7 says survives.**

Entropy in this paper is Shannon entropy where it appears at all: surplus is a logarithm of an effective count, a pure number reported in nats, and no thermodynamic quantity is invoked. The corpus's own register records the non-identification of surplus with Shannon entropy or mutual information in general (entry A-12), and records that **no thermodynamic negentropy law is proved** (entry A-13). Both are printed rather than glossed.

A note on the corpus's position in its own field. Carrying physical formalism into social description is a long-established programme [Galam, 2012, Castellano et al., 2009, Weidlich, 2003], and so is the objection that the transfer often borrows the authority of physics without

its constraints [Mirowski, 1989, Cartwright, 1999, Sobkowicz, 2020, Gallegati et al., 2006]. **The answer this paper offers to that objection is the audit itself.** A framework that publishes 0 of 27 first-party measurements has not disguised a metaphor as a measurement; it has said which is which.

38.3.9 3.9 An editorial note on cross-references

The source manuscript refers to its own internal technical papers as Papers A to E and locates claims by book chapter number. Neither survives conversion. **Where the manuscript cites Paper A, this paper cites Curzi [2026], which carries the surplus functional and its properties; where it cites Paper E, this paper cites Curzi [2026a], which carries the storage law and its corollaries; where it cites Paper D, the result is a range or effective-count bound and the citation is Curzi [2026c]; and where it cites Paper B, the result concerns state encoding and category assignment and the citation is Curzi [2026i].** Chapter numbers in the register's WHERE field have been replaced by the corpus paper identifiers that carry the corresponding material.

That mapping is an editorial inference, not a statement by the source, and Section 10 records it as such. The register's WHERE and evidence fields therefore name corpus papers as plain identifiers rather than as bibliographic citations; those identifiers resolve against the corpus DOI 10.5281/zenodo.22079583, and the bibliography of this paper carries the subset it cites directly in its running text.

38.4 4 The correspondence, parameter by parameter

Three volumes have now been written with one equation in them.

the institutional-scope material of the corpus put it on a room, a table, a team, and an institution. the personal-scope material of the corpus turned it around and put it on the person operating it — the framings a single mind holds, the character that accumulates, the load that maintaining a self costs. the cellular-scope material of the corpus carried it down to the cell, where the storage medium is chromatin state rather than minutes of a meeting, and where the signal path is a pressure wave rather than a conversation.

The claim implied by that progression is the largest in the corpus, and it has been made piecewise across the preceding papers of the corpus without ever being stated in one place with its limits attached. This chapter states it, and then spends more than a third of its length attacking it.

The claim: there is one instrument here, and the three volumes are that instrument set to three scopes. Not three analogous instruments. Not a metaphor extended twice. One parameter set — `N`, `u`, `N_eff`, `F`, `Q`, `delta`, `eta`, `C`, `h`, and in Volumes Two and Three the moment `M` and its ratio `D/P` — read against three different systems, with only the parameter *values* changing between them.

That claim has a precise status and it is not the status a reader would guess. **It is a claim about model form and nothing else** (the corpus's model-form convention; the honesty appendix, postulate 5 [CCP-039]). It is not a scaling law: it does not predict how any parameter changes between scopes. It is not a composition theorem: it does not say that cellular stored order aggregates into personal stored order, or that personal stored order aggregates into institutional stored order, and **no such theorem exists anywhere in this corpus.** It is not an empirical finding: **scale-freedom has never been tested at two scales simultaneously in any real system** (the honesty appendix, untested item 4 [CCP-039]). It is the observation that a minimal open-system accumulator written for each of the three uses the same five symbols, and that this is worth knowing because the same arithmetic then applies to all three.

A correspondence stated without its limits is a slogan. This section is written so that the correspondence cannot be quoted as one. The Engineering builds the full parameter table across all three scopes, cell by cell, with units and estimators. Then §5.1 through §5.7 take the

table apart: **a mind is not a small society, a cell is not a small mind, members can walk out and cells cannot, aggregation is not always well-defined even in principle, and at the cellular scope the input path to the whole apparatus is an open problem with no established mechanism** (CCP-031). What is left after those five sections is smaller than the table and considerably more trustworthy than the slogan, and it is the thing the corpus is entitled to carry into CCP-039.

38.4.1 4.1 The master correspondence table

Eleven rows. Three scopes. Each cell states the meaning at that scope and, in the second table, the estimator and its status. **Read the two tables together or not at all** — the first is where the correspondence looks strongest and the second is where it is priced.

Term	INSTITUTIONAL (Vol. One)	PERSONAL (Vol. Two)	CELLULAR (Vol. Three)
`N`	genuinely distinct domains, trades, or category systems admitted by the group	distinguishable categories one mind can bring to bear on a situation	distinguishable regulatory states or addressable programme classes
`u`	non-parallelism between members' independent readings of the same problem	non-parallelism between one person's own framings of the same situation	non-parallelism between regulatory programmes available to one cell
`N_eff`	how many genuinely different members remain after the roster is audited for redundancy	how many genuinely different views are held, after the same audit	how many genuinely different programmes are reachable
`F`	group surplus per interaction, `ln N_eff`, in nats, capped at `ln N`	internal surplus, `F_int`, over one observer's own framings (CCP-021)	**not instantiated** — no measured input path (§5.6)
`Q`	institution: practice, documentation, trained people, culture, infrastructure	character, skill, memory, standing commitment	maintained regulatory state — **not the molecules** (§4.6)
`delta`	turnover, drift of practice, decay of documentation, forgotten reasons	forgetting, skill atrophy, drift of self, lapse of habit	degradation and dilution: measured, and very fast (§4.7)
`eta`	how interaction becomes institution: the decision log, the artifact, the trained apprentice	how experience becomes character: the retained lesson, the changed default	**undefined** — the conversion has no established mechanism

Term	INSTITUTIONAL (Vol. One)	PERSONAL (Vol. Two)	CELLULAR (Vol. Three)
`C`	coordination overhead, friction, compliance, rent, conflict, commute	cognitive load, effort, emotional labour, the maintenance of a self	metabolic and proteostatic cost of maintaining state
`h`	organisational headroom: what may be built this period	personal capacity for the new	**not computable** — three of its four terms are unavailable
`M`	the situation as the group constructs it	the moment as lived	**does not descend** (§5.3)
`D/P`	**doctrine-to-data ratio**: does the model outrun the intelligence, or drown in it	frame-to-intake ratio	**does not descend** — no representation, therefore no frame (§5.3)

And the same eleven rows, priced:

Term	Institutional estimator	Personal estimator	Cellular estimator
`N`	Protocol MP-3, category audit — defined, unvalidated	MP-3 applied to one person's repertoire — defined, unvalidated	**none.** What counts as a category is undecided
`u`	MP-4: centred, unit-normalised error or response vectors, `u = 1 - r^2` — defined, unvalidated	MP-4 on one person's framings — defined, unvalidated, **and see §5.2**	**none.** No validated state encoder for cells
`N_eff`	`1 + (N-1)u`; equals `N + 1 - N/PR` (CC-R3, verified to 1.8e-15) — inherits S3's postulate status	same	**none**
`F`	computed from `u` — inherits everything above	computed from `u`, and see CCP-021	**none**
`Q`	MP-6 proxy, chosen by the engineer — **no ground truth exists**	MP-6 proxy — same	**none.** §4.6 argues the obvious proxy is wrong
`delta`	MP-1: two windows under zero input — defined, unvalidated	MP-1 — defined, unvalidated	**measured** (Schwanhäusser 2011); borrowed, median, one cell line
`eta`	MP-2: `Delta Q` attributable per nat of `S` — defined, unvalidated	MP-2 — defined, unvalidated	**none**

Term	Institutional estimator	Personal estimator	Cellular estimator
`C`	MP-5, requires cost commensuration — **a badged postulate** (honesty appendix postulate 8 [CCP-039])	MP-5 — same postulate	**none**
`h`	MP-7 — defined, unvalidated, and **gameable** (Goodhart; the honesty appendix [CCP-039])	MP-7 — same	**none**
`M`	ordinal only (CC-R1, FVP gate 2 = DEMOTE)	**ordinal only — CC-R1**	**not applicable**
`D/P`	ordinal only; Rasch route at 123 items per term (the Rasch-route theorem [CCP-026])	ordinal only; same route, same cost	**not applicable**

Read the right-hand column of the second table downward. Eleven rows and one entry. No amount of structural elegance in the first table changes that.

38.4.2 4.2 `N` — the category count, and why it means three different things

`N` is the number of genuinely distinct categories the system admits. It enters every design inequality in the corpus through the product `eta ln N`, and it is the hardest of the three scopes to make commensurate.

At the **institutional** scope, `N` is a roster question with a hard edge: two people trained in the same programme, rewarded by the same metric, and reading the same literature are one category, however different their biographies. CCP-009 computed how slowly `ln N` grows, and the arithmetic is the same at all three scopes:

```
N        ln N        bits        surplus ceiling in octaves (2F / ln 2)
-------------------------------------------------------------------------
2       0.693147    1.000000     2.000000
3       1.098612    1.584963     3.169925
5       1.609438    2.321928     4.643856
6       1.791759    2.584963     5.169925
12      2.484907    3.584963     7.169925
20      2.995732    4.321928     8.643856
50      3.912023    5.643856    11.287712
100     4.605170    6.643856    13.287712
```

Doubling `N` buys `ln 2 = 0.693147` nats wherever it starts. **A headroom problem cannot be escaped by adding domains problem by adding domains, and this is true at every scope for the same reason: it is a property of the logarithm, not of the system.**

At the **personal** scope, `N` is the number of genuinely distinct framings a person can bring — distinct in the sense that they generate different predictions, not merely different vocabulary. CCP-021 derived this reading from the founding material's third node, the many windows of the one tower, and it is a derivation rather than an analogy.

At the **cellular** scope, `N` has no agreed referent and this framework does not supply one. Candidate readings — the number of distinguishable chromatin states, the number of addressable regulatory programmes, the number of transcription factor combinations with distinct

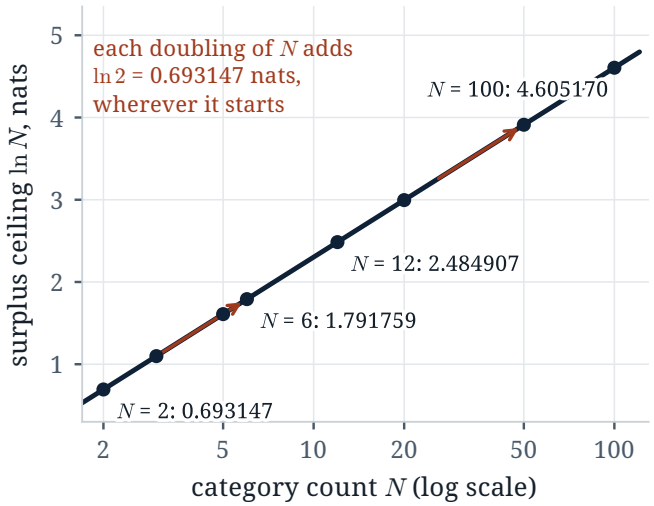


Figure 38.2: **Doubling the category count buys `ln 2` = 0.693147 nats wherever it starts, so a headroom problem cannot be escaped by adding categories — at any of the three scopes, for the same reason.** The ceiling is a property of the logarithm rather than of any system: `ln 2` = 0.693147, `ln 6` = 1.791759, `ln 12` = 2.484907, `ln 100` = 4.605170. Exact arithmetic from the table of Section 4.2; generated by `figs/src/CCP-038_fig05.py`.

downstream consequence — differ from each other by orders of magnitude and none has a validated counting procedure. **A parameter that cannot be counted cannot bound anything, and `ln N` at the cellular scope is therefore a symbol with no number behind it.** The latent-capability result [Curzi, 2026y] names the constraint that bounds the space — nothing that is not already latent can be reached — but a bound on reachability is not a count.

A structurally parallel collapse of a naive count appears in the frequency framework, where Curzi [2026] proves the odd-part class rule: the pitch class of `f(n) = k n^2` is the odd part of `n`, so `1..N` realises exactly `ceil(N/2)` classes, six at `N = 12`. **A count of distinguishable outcomes falling well below its naive value for a purely structural reason is the same move as the effective-count audit, performed on a quantity with nothing to do with rosters.**

38.4.3 4.3 `u` — non-parallelism, and the first crack

`u = 1 - (x . y)^2` is the squared sine of the angle between two unit state vectors. It is dimensionless, it lives in `[0,1]`, and it is zero exactly when the two states are parallel or anti-parallel.

The institutional reading is the one the law was written for: two members, two encoded readings of a shared problem, `u` measured as `1 - r^2` where `r` is the correlation of their centred response or error vectors. CCP-009 gives the protocols for raising it and CCP-010 gives the catalogue of ways institutions destroy it without noticing.

The personal reading is where the correspondence starts to strain, and §5.2 makes the full argument. In outline: **the `u` between two people can be large because the two people have separate substrates, separate training, and separate incentives; the `u` between two framings in one head cannot be large for those reasons, because there is only one substrate.** The framings share a memory, a mood, a set of priors, and a stake in the answer. They are correlated by construction in a way two colleagues are not. CCP-021's protocol addresses this by measuring `u` on *outputs* — independent written predictions made before the framings are compared — but that is a mitigation, not a solution, and it should be read as one.

Two properties of `u` used below are established elsewhere and are not re-derived. The law prices both sources of non-parallelism identically at the margin, `df/du_cross = df/du_div = (N-1)/(1+(N-1)u)`, while their productivity in the design variables differs [Curzi, 2026d]; and where `u` is estimated from a correlation, the opposition identity `u = 1 - rho^2` depends only

on `|rho|`, so `rho = +1` and `rho = −1` both give `u = 0` and zero surplus [Curzi, 2026]. **A perfectly opposed pair purchases exactly as much as a perfectly aligned one, namely nothing.**

The cellular reading has no encoder. Computing `u` requires a map from the thing to a unit vector such that the inner product carries the relevant notion of alignment, and the honesty appendix [CCP-039] item 6 already badges that map as an operational choice at the scopes where it exists. **At the cellular scope the choice has not been made, let alone validated.** Expression-profile correlation is the obvious candidate and it is not obviously the right one: two programmes may be transcriptionally similar and functionally opposed, or transcriptionally distinct and functionally redundant.

38.4.4 4.4 `N_eff` — the audited count, and the one identity that holds everywhere

`N_eff = 1 + (N−1)u` is the effective count: the headcount after redundancy has been subtracted. `u = 0` gives `N_eff = 1` — ten parallel windows are one window — and `u = 1` gives `N_eff = N`.

One result about it is scope-independent because it is arithmetic. **CC-R3: the mean-pairwise `u` aggregator is exactly the Rényi-2 participation ratio effective count**, `N_eff = N + 1 − N/PR`, where `PR = (sum lambda)^2 / sum lambda^2` on the correlation matrix. Verified over six random ensembles with `N` from 8 to 10 and 200 observations each; worst discrepancy `1.8e−15`, which is floating-point noise:

```
N = 8   ubar = 0.783801   N_eff(law) = 6.486606017883   N + 1 − N/PR = 6.486606017883
N = 10  ubar = 0.869860   N_eff(law) = 8.828735833534   N + 1 − N/PR = 8.828735833534
N = 8   ubar = 0.952367   N_eff(law) = 7.666571449037   N + 1 − N/PR = 7.666571449037
```

This identity holds at every scope where a correlation matrix can be formed, and it is the only row of the master table with that property. It is worth noticing why: it is a fact about the aggregator, not about people. It would hold for correlated thermometers.

Which effective count is meant matters here more than anywhere else in the paper. The affine count `1 + (N−1)u` is this corpus's own construction and is **not** a Hill number. The participation ratio it equals *is* one, of order two. The concentration instrument built on the other standard member of the family — `1/HHI`, with the worked jurisdictional case returning a naive HHI of 0.293430 and 3.41 effective jurisdictions against a bloc-HHI of 1.000000 and an effective count of 1.00 — is a separate quantity developed in Curzi [2026], and the two must not be conflated.

It also states what `N_eff` costs to compute: a correlation matrix, which requires repeated joint observation of the members on a common set of items. At the cellular scope such a matrix is obtainable in principle — single-cell expression data supplies one — while whether its eigenvalues mean what the law needs them to mean is entirely open.

38.4.5 4.5 `F` — surplus, and the ceiling that appears at every scope

`F = ln N_eff`, capped at `ln N`, in nats. At the institutional scope it is the surplus a single interaction can contribute. At the personal scope it is CCP-021's internal surplus. At the cellular scope **it is not instantiated**, because `S_t` requires an input path and CCP-031 states that the input path for sequence-specific content is exactly the open problem.

That is worth saying without decoration. **The third instrument has a storage equation whose input term has no established mechanism.** Everything upstream of that term is real; the specific step from an acoustic signal to a sequence-specific cellular read is established by nothing. The equation can be written. It cannot yet be driven.

The concavity that governs `F` is the same at every scope and it is the merciful clause:

```
N = 6:   u = 0.1  −>  N_eff = 1.500  F = 0.405465  (22.6% of ceiling)
```

```
u = 0.3  ->  N_eff = 2.500  F = 0.916291  (51.1%)
u = 0.5  ->  N_eff = 3.500  F = 1.252763  (69.9%)
u = 0.7  ->  N_eff = 4.500  F = 1.504077  (83.9%)
u = 0.9  ->  N_eff = 5.500  F = 1.704748  (95.1%)
u = 1.0  ->  N_eff = 6.000  F = 1.791759  (100.0%)
```

Half the available surplus arrives by `u = 0.3`. Three honest windows set wide capture most of what a gallery could, and that sentence has been true in the institutional-scope material of the corpus about people, in Volume Two about framings, and would be true in the cellular-scope material of the corpus about programmes if the programmes could be counted.

38.4.6 4.6 `Q` — stored order, and the result that embarrasses the naive reading

`Q` is what persists. At the institutional scope it is practice, documentation, trained members, and the built environment; at the personal scope it is character, skill, and standing commitment. The same law written in world coordinates rather than log coordinates is a geometric accumulator, `G_{t+1} = G_t^(1-delta) (N_eff)^eta e^(-C)` with `G_inf = ((N_eff)^eta e^(-C))^(1/delta)`, agreeing with the additive form to fourteen significant figures [Curzi, 2026n]; the choice between the two is a choice of register and not of model. Both are proxied, never measured, and the honesty appendix, postulate 3 [CCP-039] badges the whole idea that what persists in a social system is a single scalar decaying exponentially.

At the cellular scope the naive proxy is the obvious one — the molecules — and **the naive proxy is provably wrong on the measured numbers.** Take the median protein half-life above and ask what fraction of a cell's protein complement survives one month:

```
t_half = 46 h,  one month = 720 h
surviving fraction = 2^(-720/46) = 1.942e-05
delta per month    = 0.99998
```

Essentially none of it. A habit can persist in a person for thirty years; the molecules implementing it turn over in days. If `Q` at the cellular scope were the molecular inventory, the storage law would predict that nothing cellular persists at any timescale a person would notice, and that prediction is flatly false.

The resolution is not a repair to the law. It is a correction to the proxy, and it is the same correction CCP-031 arrives at from the other direction: **what persists is not the material but the maintained pattern that keeps rebuilding it.** Chromatin state, methylation, the self-sustaining regulatory circuit — these are candidate `Q` at the cellular scope, and they are candidates precisely because they are the things with long effective retention despite fast component turnover.

Two consequences, and both constrain the correspondence rather than support it.

First, the cellular `Q` is a different kind of object from the other two. Institutional and personal `Q` are aggregates of many persisting things; cellular `Q` is a maintained configuration whose components do not persist at all. That is not a difference in parameter value. It is a difference in what the state variable *is*.

Second, this is the clearest case in the corpus of a scope where the model form survives only because the modeller chose the proxy to make it survive. That is a legitimate move — every proxy choice is a modelling decision — but it is exactly the move that makes a framework unfalsifiable if it is made silently. It is made here in public, and CCP-031's experimental programme is where it would be tested.

38.4.7 4.7 `delta` — the best-measured parameter, and the inversion

Retention loss. Fraction of `Q` lost per period with zero input. `t_half = ln 2 / (-ln(1-delta))`.

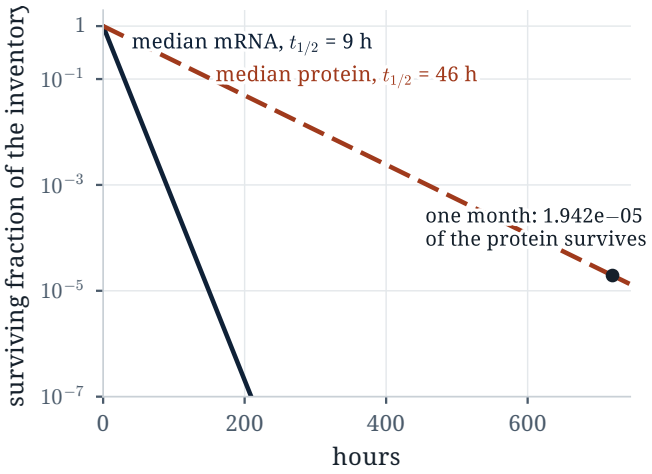


Figure 38.3: **The obvious proxy for cellular stored order is refuted by the corpus's only measured parameter: after one month, 1.942e$-$05 of the protein inventory survives.** At the median half-lives the borrowed study reports — 9 hours for mRNA and 46 hours for protein — a stored order denominated in molecular inventory would predict that nothing cellular persists on any timescale a person would notice, which is false. What persists is the maintained pattern that keeps rebuilding the material. Borrowed measurement, exponential decay implied by it (Section 4.6); generated by `figs/src/CCP-038_fig03.py`.

delta	tau = −1/ln(1−delta)	half-life (periods)
0.005	199.500	138.283
0.01	99.499	68.968
0.02	49.498	34.310
0.05	19.496	13.513
0.10	9.491	6.579
0.15	6.153	4.265
0.20	4.481	3.106
0.30	2.804	1.943
0.50	1.443	1.000

At the institutional scope, staff turnover is the most legible contributor: 15% annual turnover gives a 4.265-year half-life on whatever `Q` is carried in people rather than artifacts, which is the arithmetic behind CCP-014's insistence that documentation is not bureaucracy but retention engineering. At the personal scope `delta` is forgetting and skill atrophy, estimated by Protocol MP-1's two-window method under zero input — which for a person means a genuine lay-off period, and which almost nobody will run. At the cellular scope `delta` is degradation and dilution, and here is the inversion this section should not step over:

The scope with the least-validated instrument has the best-measured parameter.

Institutional `delta` has never been measured under Protocol MP-1 by anyone. Personal `delta` has never been measured under Protocol MP-1 by anyone. Cellular `delta` has been measured genome-wide, by mass spectrometry, in a published paper, with a stated method and stated spread.

This is inconvenient for the tidy version of the correspondence, in which the scopes get more speculative with descending scale, and it should be allowed to be inconvenient. The scopes do not degrade monotonically; they degrade differently. The institutional scope has the best borrowed sociology and the worst first-party measurement. The personal scope has

the best measurement *register* — a century of psychophysics establishing that subjective magnitude is measurable at all — and a demoted keystone formula. The cellular scope has the best-measured single parameter and no input path. **Three instruments, three different failures.**

38.4.8 4.8 `eta` — conversion, and the parameter that does not descend

`eta` is Q-units of durable structure produced per nat of surplus. CCP-008 argues it is the most neglected parameter and usually the cheapest to improve, and CCP-008 shows why it is the only term that scales the retained headroom at the fixed point.

Institutionally it is legible: the decision log, the written standard, the trained apprentice, the changed default, the moved wall. A meeting with no artifact has `eta` near zero. Personally it is the retained lesson — the difference between an experience that changed a default and one that was merely had. CCP-021 makes the join from the moment to stored character, and CC-R1 constrains that join: **the join is structural, not yet quantitative.**

Cellularly there is nothing. Measuring `eta` requires `S` and an attributable `Delta Q`; `S` has no input path and `Q` has no validated proxy. **`eta` at the cellular scope is not a small number or an uncertain number. It is not a number.**

38.4.9 4.9 `C` — cost, and the commensuration postulate

`C` is dissipation: coordination overhead, friction, maintenance, conflict, rent, compliance at the institutional scope; cognitive load, effort, emotional labour, and the maintenance of a self at the personal scope; metabolic and proteostatic expenditure at the cellular scope.

Every headroom computation in this corpus requires that these be expressed in the same unit as `eta ln N` and `delta Q`. **That commensuration is a badged postulate (the honesty appendix, postulate 8 [CCP-039]), the conversion factor is chosen by the engineer rather than measured, and it dominates the headroom result whenever `C` is large.** It is uncomfortable at the institutional scope, worse at the personal scope — where the claim is that an hour of emotional labour and an hour of cognitive load are commensurable — and untested at the cellular scope.

Brooks' arithmetic remains the sharpest statement of why the scopes cannot simply be scaled up: **communication paths grow as `n(n−1)/2` while surplus grows as `ln N`.** Quadratic against logarithmic. There is a group size beyond which no interaction geometry pays for its own coordination, and CCP-008 computes where it is for given `eta`.

38.4.10 4.10 `h` — the master gauge, and the only row that reads at more than one scope

```
h_t = eta ln N − delta Q_t − C_t

h < 0   unsustainable in this model; no geometry saves it
h = 0   maintenance only
h > 0   growth available; any requested increment must satisfy r_t <= h_t
```

Headroom is the gauge the whole manual is organised around, and it reads at the institutional and personal scopes because all four of its terms have estimators there. It does not read at the cellular scope, because three of its four terms do not.

One identity about it is scope-independent and was established in CCP-008 (CC-R6):

```
h_inf = eta*(ln N − F(u*))
```

At the fixed point, `delta`, `C` and `Q` all cancel. Verified across four parameter sets:

```
eta=0.20  N=6   u*=0.50  delta=0.10  C=0.25  ->  h_inf = 0.1077993001 both ways
eta=0.35  N=12  u*=0.30  delta=0.05  C=0.40  ->  h_inf = 0.3592020695 both ways
eta=0.50  N=4   u*=0.80  delta=0.20  C=0.10  ->  h_inf = 0.0812594647 both ways
eta=0.12  N=20  u*=0.62  delta=0.03  C=0.18  ->  h_inf = 0.0537420990 both ways
```

A structurally parallel identity holds in the frequency framework, where the octave deficit and the steady-state gauge turn out to be the same number: `D = 2.885390 (c/eta − F) = −2.885390 (delta/eta) Q_inf` [Curzi, 2026]. **Two quantities derived for different purposes collapsing onto one is the same kind of finding as the participation-ratio identity, and it has the same status: arithmetic.**

Since `F(u) <= ln N` by CCP-009 Corollary 2.2, `h_inf >= 0` always. **The plain linear model never ossifies by itself.** Ossification requires something the model does not contain: exogenous `Q` accretion, or a cost that rises with `Q`. That is a genuine structural result and it applies wherever the model form is adopted — which is to say at all three scopes, in exactly the sense §3.2's second row licenses and no further.

38.4.11 4.11 `M` and `D/P` — and where they stop

The moment equation `M = (D/P) . E` and its calibration ratio `r = D/P` are the two rows that exist in the personal-scope material of the corpus, are read at group scope in CCP-024 as the doctrine-to-data ratio, and **do not descend to the cell at all.**

Their status must be carried in full wherever they appear, and this section is a place where they appear:

> **CC-R1 — DEMOTED.** FVP gate 2 (scale type) returns **DEMOTE** for `M = (D/P) . E`. Because `D` and `P` are not instrumented on ratio scales with true zeros, the magnitude of `D/P`, the distortion `Delta = |ln r|`, and the join `M_t -> Q_{t+1}` are **not reportable as measured quantities.** What survives is structure: the limiting cases, the `D = P` fixed point, and the three-way ordering over-framed / calibrated / under-framed, all of which need only ordinal information. **What is demoted is magnitude, not structure.**
>
> **The route out is constructive and costed.** the Rasch-route theorem [CCP-026], given badged Postulate SO1: if `D` and `P` are measured as Rasch person-measures on one common scale, then `Delta = |ln(D/P)| = |theta_D − theta_P|` is ratio-scaled with a true zero at perfect calibration. Rasch separation supplies the true zero that division requires. **The cost is 123 items per term** to resolve a half-logit distortion.

The `123` is not an estimate. It is the inversion of the best-case measurement error, and it can be recomputed in one line: with every item targeted at the person, item information is `0.25`, so `SE(theta) = 2/sqrt(k)`; the distortion is a difference of two independent measures, so `SE(Delta) = sqrt(2) . 2/sqrt(k)`; and the 95% detection floor is `1.96 . SE(Delta)`.

```
k = 122  ->  SE(theta) = 0.181071   SE(Delta) = 0.256074   floor = 0.501905
k = 123  ->  SE(theta) = 0.180334   SE(Delta) = 0.255031   floor = 0.499860
```

One hundred and twenty-two items is not enough and one hundred and twenty-three is, and the crossing is that sharp because the arithmetic is that simple. This is the volume's actual research instruction and it is presented as one rather than as a caveat.

At the cellular scope the rows are empty for a reason more fundamental than instrumentation. §5.3 gives it: `D/P` is a ratio of a frame to an intake, and a frame is a representation. The argument that a cell has no representation is not a claim about current measurement technology.

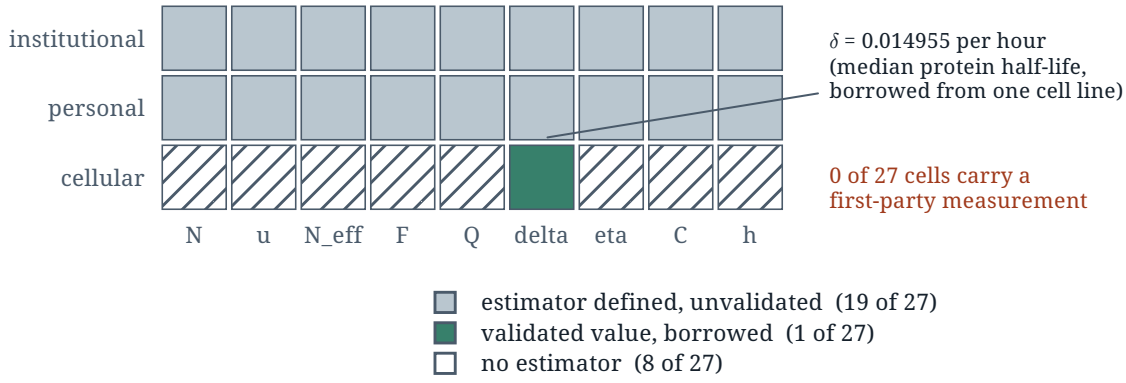


Figure 38.4: **Nineteen of twenty-seven parameter-by-scope cells have a defined estimator, one has a validated value, and none carries a measurement made first-party for this corpus.** The single validated cell is the cellular retention rate, a median protein half-life borrowed from a genome-wide study in one immortalised cell line. Audit of the corpus's own parameter tables (Sections 4.1 and 4.12); generated by `figs/src/CCP-038_fig01.py`.

38.4.12 4.12 The measurement-status summary

Collecting the two tables of §4.1 into one count, because the count is this section's most useful single output:

	Institutional	Personal	Cellular	Row total
Defined estimator	9	9	1	19 of 27
Validated estimator	0	0	1	1 of 27
First-party measurement published	0	0	0	0 of 27
Demoted to ordinal (CC-R1)	`M`, `D/P`	`M`, `D/P`	n/a	4

Nineteen of twenty-seven parameter-scope cells have a procedure; one has a validated number; none has a number this research programme measured. Everything in the next section follows from taking that table seriously rather than from any philosophical objection to correspondence.

38.5 5 Five breaks in the correspondence

38.5.1 5.1 Where the correspondence breaks down — the five breaks

This is the section this section was written for. The correspondence fails in five distinct ways, and they are genuinely distinct — a repair to one does not touch the others.

```
BREAK 1   A mind is not a small society.
          The independence the law's aggregation needs is unavailable inside one skull.

BREAK 2   A cell is not a small mind.
          M and D/P require a representation. A cell has no frame, so it has no D/P.

BREAK 3   Members can exit; cells cannot.
          The consent constraint that governs this corpus has no cellular analogue,
          which means the ethics do not descend with the arithmetic.
```

```
BREAK 4   Aggregation is not always well-defined, in principle and not merely in practice.

BREAK 5   The cellular scope has no established input path (CCP-031),
          and no composition theorem connects any scope to any other.
```

38.5.2 5.2 Break 1 — a mind is not a small society

The correspondence is most seductive at the personal scope because the personal-scope material of the corpus makes it work so well. Internal surplus is a derivation, not an analogy: CCP-A02 applied to one observer's framings is CCP-A02. CCP-021 shows this and the showing is correct.

What does not carry across is independence, and independence is the precondition on which every aggregation result in the institutional-scope material of the corpus rests.

Condorcet's jury theorem gives collective accuracy above individual accuracy *only under independence of the members' errors*. Hong & Page's result requires a diverse pool of solvers drawn separately. The participation-ratio estimator computes `N_eff` from a correlation matrix whose off-diagonal entries are small precisely when the members' readings were formed apart. Woolley's `c` correlates with turn-taking equality — a property of separate speakers taking separate turns.

Inside one head there are no separate speakers. Two framings held by the same person share:

- one memory, so both are conditioned on the same retrieved evidence;
- one mood, so both are weighted by the same current affect;
- one set of priors, formed by the same history;
- one stake in the answer, so both face the same incentive to prefer a conclusion;
- one substrate, whose fatigue, illness, or hunger perturbs both together.

Those are five channels of common-mode coupling with no institutional analogue, and they do not appear anywhere in the formalism. The law will happily compute `u = 0.6` between two framings a person has written down. Whether that `0.6` means what a `0.6` between two colleagues means is not a question the law can answer, and this framework's answer is: **not obviously, and probably not.**

The mitigation CCP-021 offers is real and should be used: measure `u` on *committed outputs* — independent written predictions produced before the framings are compared — rather than on introspective report. That converts the measurement from "how different do these feel" into "how decorrelated were their errors," the same operationalisation the institutional-scope material of the corpus uses on people. **It removes one source of inflation. It does not remove common-mode coupling, because the predictions still came from one substrate.**

A second, older objection should be named rather than avoided. Treating a mind as an assembly of parts that deliberate invites the homunculus: if the framings deliberate, something must be doing the deliberating, and if that thing is described the same way, the regress is launched. Minsky's *Society of Mind* (1986) is the best-known development of the assembly picture; it is an argument rather than a body of evidence, and is cited here as precedent for the picture, not support for it.

This framework's position, stated as a position: the internal surplus computation is a measurement of decorrelation among a person's committed outputs, and it is licensed as exactly that and nothing more. It is not evidence that a mind is composed of agents. It is not evidence that the sociological results transfer. It is a number about how much a person's own predictions disagree, which is a useful number, and which has the properties CCP-A02 proves for it because CCP-A02's proof does not care where the vectors came from.

A mind is not a small society. It is one system whose outputs can be decorrelated, sometimes, with effort, and never as far as two people's are.

38.5.3 5.3 Break 2 — a cell is not a small mind

The second break is sharper than the first because it is categorical rather than quantitative.

`M = (D/P) . E` requires `D`, a perspective — a frame brought to the encounter — and `P`, a perception — the intake against which the frame is weighed. The ratio is a comparison of two representations of the same situation. **A representation is the load-bearing requirement, and the argument that a cell lacks one is not an argument about measurement difficulty.**

A cell has states, and states are not representations. A transcription factor bound at a promoter is a physical configuration with causal consequences; it is not *about* anything in the sense the moment equation requires. The predictive-processing correspondence that the personal-scope material of the corpus draws — the ratio of prior precision to sensory precision being structurally the same object as `D/P` — depends on there being a prior and a likelihood, which is to say on there being a model. **This manual makes no claim that a cell has a model, and the moment equation therefore has no cellular instance.**

That has three consequences worth stating separately.

The `M` and `D/P` rows of the master table are empty at the cellular scope, and they are empty permanently rather than pending instrumentation. Instrumenting `D` and `P` with a Rasch instrument at 123 items per term is a research programme at the personal scope. At the cellular scope there is nothing to instrument, because there is no respondent.

The self-storage law does not descend either. `Q_{t+1} = (1−delta)Q_t + eta M_t − C_t` uses `M` as its input. Without `M` the cellular instance of the self-storage law is not a harder version of the same equation; it is a different equation with a hole where its input was.

And the direction of the correspondence matters. It is tempting to read the three volumes as a descent in which each scope is a smaller copy of the last. They are not copies. the cellular-scope material of the corpus's subject is a receiver with measured mechanics and no established semantics; the personal-scope material of the corpus's, a semantics with a demoted measurement; the institutional-scope material of the corpus's, a population with borrowed statistics. **They share an equation and they do not share a nature.**

38.5.4 5.4 Break 3 — exit, and why the ethics do not descend

Hirschman's *Exit, Voice, and Loyalty* (1970) names the two responses available to a member of a declining organisation: leave, or complain. The distinction is structural, not moral, and it is the sharpest available statement of what a member of an institution is.

A member of an institution can exit. A framing in a mind can be abandoned. A cell cannot leave the body.

This asymmetry does real work in three places.

It changes what `delta` means. At the institutional scope a substantial part of `delta` is voluntary departure, and voluntary departure is *responsive*: it rises when conditions worsen, which makes institutional `delta` endogenous to the very quantities the law treats as separate. The linear model treats `delta` as a constant within a period (the honesty appendix, postulate 9 [CCP-039]) and this is the clearest case where that assumption is not merely convenient but wrong in a known direction. At the cellular scope there is no voluntary departure. Apoptosis exists, is programmed, and is sometimes described in the language of choice; it is not exit in Hirschman's sense, because it is not a response by the unit to its own assessment of the organisation, and because the unit does not survive it. **Naming programmed cell death as "exit" would be exactly the kind of surface resemblance the two-column method exists to reject (CCP-029, criterion C1).**

It changes what a failure of the design means. The framework's fourth refusal is that it will not recommend an intervention removing a person's ability to decline (the standing refusals of the corpus). That refusal is enforceable at the institutional scope because members can decline; it is enforceable at the personal scope because the practitioner and the operated-on

are the same person, who can stop. **It is not enforceable at the cellular scope at all, because there is no one there to decline.**

Therefore the ethics attach to the person, not to the scope. This is the most important practical consequence in this section and it must not be softened. Consent at the biological layer is not consent obtained from the cell; it is consent obtained from the person whose cells they are, before anything is done, with the right to stop at any point, and with nothing concealed — which is the standing no-subliminals constraint (the corpus's disclosed-signal constraints) restated at the scope where its violation would be least visible. **The arithmetic descends to the cell. The moral standing does not descend with it, and pretending otherwise would let a formalism launder an ethical question it is not competent to answer.**

38.5.5 5.5 Break 4 — aggregation is not always well-defined

The three scopes are not merely different sizes of the same thing; going between them requires aggregation, and aggregation is not a free operation.

The first problem is a theorem. Arrow's impossibility result (1951) shows that no rule aggregating individual orderings into a collective ordering can simultaneously satisfy a short list of conditions that each look unobjectionable. It is a result about preference orderings specifically and this framework does not claim it applies verbatim to state vectors. **What it establishes, and what is fully general, is that "combine the parts into a whole" is a mathematical operation with constraints, not a bookkeeping step, and that the constraints can be mutually unsatisfiable.** A framework that moves between scopes without saying which aggregation rule it is using has skipped a step that is known to be capable of having no solution.

The second problem is the aggregator this corpus actually uses. CCP-A02's law is stated for a *pair*. Real groups have more than two members. The framework extends the pairwise law to an ensemble via the participation ratio, and the honesty appendix, postulate 7 [CCP-039] badges the extension precisely: **consistency by construction is a property of the choice, not evidence for it.** CC-R3 confirms that the mean-pairwise aggregator and the participation ratio agree to machine precision — which is a statement that two descriptions of the same choice coincide, not a statement that the choice is correct.

The third problem has a name. Relationships measured at the group level need not hold at the individual level and can reverse: the ecological fallacy, whose arithmetic form is Simpson's paradox. **A `u` computed across teams and a `u` computed within a team are different quantities that the notation does not distinguish**, and a practitioner who computes one and reports the other has made an error the formalism will not catch.

The fourth problem is that at the cellular scope the unit of aggregation is undecided. Is the member a cell, a cell type, a regulatory programme, or a locus? Each gives a different `N`, a different correlation matrix, and a different `N_eff`. There is no principled answer in this manual and inventing one would be worse than admitting there is none.

38.5.6 5.6 Break 5 — the unbridged gap, and the missing composition theorem

The fifth break is the one the cellular-scope material of the corpus was built around, and it is stated in the Methodology Protocol's own words because paraphrase would soften it:

> **There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal.**

Everything upstream is real and separately well-supported: PIEZO1 and PIEZO2 are mechanically gated ion channels (Nobel Prize in Physiology or Medicine, 2021); force at the cell surface transmits through the cytoskeleton to the nucleus, stretches chromatin, and changes

transcription (Tajik et al., *Nature Materials*, 2016); ultrasound has been used to activate specific channels (Ibsen et al., *Nature Communications*, 2015); DNA is a demonstrated data-storage medium (Church, Gao & Kosuri, *Science*, 2012); Cas12a processes its own crRNA array (Zetsche et al., *Cell*, 2015); and dCas9/CRISPRa establishes that targeting and editing are separable (Qi et al., 2013; Gilbert et al., 2014).

The step from those to sequence-specific acoustic addressing is not established by any of them. It is one well-specified falsifiable gap, CCP-031 spends its length on what an experiment would have to show and what controls it would need, and naming it sharply is the most useful thing this volume does.

For the correspondence specifically, the gap has an exact consequence: **the third instrument's `S_t` term has no measured input path, so the third instrument is a specification rather than a measurement.** Two of the three volumes describe systems the reader can go and measure tomorrow with an unvalidated estimator. The third describes a system whose input stage is an open research question.

To that, add the absence nobody has yet been asked to notice.

> **There is no composition theorem in this corpus.** Nothing here derives institutional `Q` from personal `Q`, or personal `Q` from cellular `Q`. No aggregation operator is defined between scopes, no conservation law links them, and no argument is offered that the three accumulators are coupled at all rather than merely similar in form.

That absence is not a gap in the writing. It is a gap in the mathematics, and it is the single largest open problem the correspondence raises. **What would close it is a stated aggregation operator with a proof that the storage law is closed under it** — that is, that a population of systems each obeying the law, combined by the practitioner, yields a system obeying the law with derived parameters. Whether the law is closed under any natural aggregation is not known here. It is a concrete, tractable-looking mathematical question, it does not require a laboratory, and CCP-039 returns to it as the problem the corpus ends on.

38.5.7 5.7 What survives

Strip the five breaks out and something remains. It is smaller than the master table and it is load-bearing.

One. The arithmetic transfers wherever the model form is adopted. Contraction, unique fixed point, convergence, the `ln N` ceiling, the headroom inequality, `h_inf = eta(ln N − F(u*))`, concavity and its marginal-gain schedule — these are consequences of the map, and they hold at any scope where someone is willing to adopt the map. **That is a real and unusual kind of portability: it costs nothing to check and it cannot be defeated by an objection to the scope, only by an objection to the map.**

Two. Three diagnostics survive at all three scopes, because they need only ordinal information. Whether `u` is near zero. Whether `eta` is near zero. Whether `delta Q + C` exceeds `eta ln N`. Each is a comparison, not a magnitude, and CC-R1's demotion is precisely the lesson that comparisons survive where magnitudes do not. A group whose members never disagree, a practice that leaves no artifact, and a system whose maintenance consumes its entire budget are diagnosable without a calibrated instrument.

Three. The failure catalogue transfers better than the measurements do. Pseudo-diversity, the meeting with no artifact, the maintenance trap, the measure that became a target — structural patterns derived from the arithmetic rather than from data, and the arithmetic is the same at every scope.

What does not survive is any claim that a measurement at one scope supports a conclusion at another. That is the sentence this section exists to install, and CCP-038 is the register that makes it checkable claim by claim.

38.6 6 The status register

This is the section to consult in order to argue with the corpus.

Everything in the preceding papers of this corpus is one of five things. It is a **theorem**, true within stated mathematical assumptions and making no claim about the world. It is a **postulate** of this research programme, chosen rather than derived, and badged wherever it is used. It is a **borrowed** empirical result, obtained by other people in settings that are not yours. It is a **speculative** bridge, marked as speculative in the sentence that states it. Or it is **demoted** — a claim this book's own validation protocol rejected, printed with its rejection attached.

A register that contains no demotions is not a register. It is a bibliography with opinions. This one contains them, they are in §6.6, and the most important of them is the demotion of the author's own keystone formula.

This section is a lookup table and behaves like one. Every entry has the same six fields in the same order, and the fields were chosen so that a hostile reader can check an entry without reading this section it came from.

```
ID          stable identifier; cited from other chapters and from the honesty appendix [CCP-039]
CLAIM       the substantive proposition, stated in one sentence
STATUS      PROVED / POSTULATED / BORROWED / SPECULATIVE / DEMOTED
(+1)        the evidence source, and WHAT THAT SOURCE ACTUALLY MEASURED
(−1)        the precedent source, and WHAT THAT SOURCE ACTUALLY RECORDED
WHERE       the corpus papers in which the claim does work
```

The two-column apparatus is the cellular-scope material of the corpus's, from CCP-029, and it is applied here to all three volumes because the discipline it enforces is not specific to biology. Its one rule governs every row below and is repeated because it is the rule most often broken by works of this kind:

> **The (+1) column carries EVIDENCE. The (-1) column carries PRECEDENT. Two sources agreeing does not double the evidence when only one of them is evidence.** What independent arrival establishes is real and worth having — that a structure was findable by more than one route, which is a genuine argument that the structure is in the territory and not only in the instrument. It is not a second measurement.

Where no precedent exists, the (-1) cell reads **none claimed**, and it reads that way in more than half the rows. **A blank column is the correct output when the historical corpus has nothing to say, and filling it would have been the easiest and most damaging thing this section could do.**

38.6.1 6.1 Group A — the formal spine

Theorems. **Nothing in this group is a claim about the world**, and every application of one of them to a real system passes through a Group B postulate first.

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
A-01	`f(u) = ln(1 + (N−1)u)` is the unique function satisfying axioms S1–S4	PROVED	CCP-A02, Theorem 2.1 (Curzi, 2026). A uniqueness proof; measures nothing	none claimed	CCP-007, CCP-009, CCP-017; the reference card [CCP-039]
A-02	`N_eff = e^{F} = 1 + (N−1)u`; surplus is `ln N_eff`, the log of an audited headcount	PROVED (given S3)	CCP-A02, Thm 2.1; CCP-009, Thm 2.1	none claimed	CCP-009, CCP-021

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
A-03	Strict monotonicity: more non-parallelism yields strictly more surplus	PROVED	CCP-A02, Thm 3.1	none claimed	CCP-009
A-04	Positivity: `f(u) > 0` for `u > 0`	PROVED	CCP-A02, Thm 3.2	none claimed	CCP-009
A-05	Strict concavity: each added increment of difference buys strictly less than the last	PROVED	CCP-A02, Thm 3.3	none claimed	CCP-009, CCP-020; CC-R2
A-06	Subadditivity **on the parameter `u` only**	PROVED	CCP-A02, Thm 3.4	none claimed	CCP-010
A-07	Sharp Lipschitz constant `N−1`	PROVED	CCP-A02, Thm 4.2	none claimed	CCP-009
A-08	Extremality: `F = 0` iff parallel; `F = ln N` iff orthogonal	PROVED	CCP-A02, Thm 4.3	none claimed	CCP-009, CCP-010
A-09	Range: `0 <= F(u) <= ln N`. **`ln N` is the hard ceiling on surplus per interaction**	PROVED	CCP-009, Cor 2.2	none claimed	CCP-008, CCP-009, CCP-038
A-10	Two-source decomposition: `u = u_cross + u_div`, with `max(f(u_c), f(u_d)) <= f(u) <= f(u_c) + f(u_d)`	PROVED	CCP-A02, Thm 4.1	none claimed	CCP-010, CCP-027
A-11	**The triangle inequality FAILS for geometric triples**	PROVED	CCP-A02, §5. A counterexample, not a limitation of technique	none claimed	CCP-010; the honesty appendix [CCP-039]
A-12	Surplus is **not** identified with Shannon entropy or mutual information in general; as deployed it is a **lower bound** on the ensemble's information	PROVED / clause	CCP-009. A non-identification result	none claimed	CCP-009; the honesty appendix [CCP-039]
A-13	**No thermodynamic negentropy law is proved**; `F(u) + N_can(u) <= ln N`, so higher static surplus does not imply higher canonical concentration	PROVED	CCP-007, Cor 2.2	none claimed	the honesty appendix [CCP-039]
A-14	Exact solution of the storage law: `Q_t = (1−δ)^t Q_0 + Σ (1−δ)^{t−1−k}(ηS_k − C_k)`	PROVED	CCP-007, Thm 3.1	none claimed	CCP-007, CCP-023
A-15	One-step growth iff `η S_t > δ Q_t + C_t`	PROVED	CCP-007, Cor 3.2	none claimed	CCP-008, CCP-016
A-16	Uniform bound `Q_t <= max{Q_0, η ln N / δ}`; with `δ=0`, `Q_t <= Q_0 + t η ln N`	PROVED	CCP-007, Thm 3.3	none claimed	CCP-007
A-17	**Sustainability ceiling: if `C_t > η ln N − δQ_t`, no interaction geometry prevents order from decreasing**	PROVED	CCP-007, Cor 3.4	none claimed	CCP-008, CCP-016, CCP-038
A-18	Steady state `Q_\ensuremath{\infty} = (ηF(u∗) − c∗)/δ`; **if `ηF(u∗) < c∗` the equilibrium is negative and the regime is unsustainable or the linear model must be replaced**	PROVED	CCP-007, Thm 3.5	none claimed	CCP-007, CCP-023; the reference card [CCP-039]
A-19	Headroom `h_t = η ln N − δQ_t − C_t`, with three regimes	PROVED	CCP-007, §6.4	none claimed	CCP-008, CCP-015, CCP-038
A-20	Maximum one-step increment `r_max = h_t`; no policy based on interaction geometry alone guarantees more	PROVED	CCP-007, Cor 4.2	none claimed	CCP-008, CCP-015; CC-R4
A-21	Target increment feasible iff `δQ_t + C_t + r_t <= η ln N`	PROVED	CCP-007, Prop 4.1	none claimed	CCP-008, CCP-016
A-22	Admissibility `C_t <= (1−δ)Q_t + ηS_t` guarantees `Q_{t+1} >= 0`	PROVED	CCP-007, §3	none claimed	CCP-008
A-23	`δ > 0` makes the storage map a contraction; a unique fixed point exists and iteration converges to it from any start	PROVED	Banach fixed point theorem (1922)	the founding material, the fixed-point node — the room that was always furnished, recorded as recognition rather than victory	CCP-007, CCP-023

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
A-24	The storage law **is** a discrete-time first-order low-pass filter: `Q` ↔ charge, `δ` ↔ leakage, `ηS` ↔ current in, `(1−δ)` ↔ pole	PROVED (identity)	Standard linear systems theory. An algebraic identification, not an analogy	none claimed	CCP-007, CCP-028
A-25	Time constant `τ = −1/ln(1−δ)`; half-life `ln2 / (−ln(1−δ))`; ~3τ to 95%, ~5τ to 99%	PROVED	Arithmetic of the geometric decay	none claimed	CCP-007, CCP-038
A-26	**`h_inf = η(ln N − F(u*))`** — at the fixed point `δ`, `C` and `Q` all cancel. Since `F <= ln N`, `h_inf >= 0` always: **the plain linear model never ossifies by itself**	PROVED	CCP-008 / CC-R6; verified against direct computation across four parameter sets, agreement to `< 1e−16`	none claimed	CCP-008, CCP-014, CCP-016, CCP-038, CCP-039
A-27	The mean-pairwise `u` aggregator **equals** the Rényi-2 participation-ratio effective count: `N_eff = N + 1 − N/PR`	PROVED	CCP-015 / CC-R3; verified over six random ensembles, worst discrepancy `1.8e−15`	none claimed	CCP-009, CCP-015, CCP-021, CCP-038; the reference card [CCP-039], B
A-28	By concavity, `mean F(U) <= F(mean U)`: feasibility must be tested on the mean of `F`, never on `F` of the mean	PROVED	CCP-A02 Thm 3.3 + Jensen. Worth **6.2%** of surplus in the CCP-016 worked office	none claimed	CCP-008, CCP-015, CCP-016, CCP-023, CCP-025
A-29	**The surplus law is exactly a Fechner law with threshold `u_0 = 1/(N−1)`:** `F(u) = ln((u + u_0)/u_0)`, and `F(u_0) = ln 2` for every `N`	PROVED	Algebraic identity; verified for `N = 2..100`, maximum residual `4.4e−16`	Fechner (1860) recorded the same functional form for sensation, from psychophysical threshold data, with no knowledge of this law	CCP-020, CCP-039
A-30	Octave folding of an essence set costs **exactly `ln 2` nats — one bit, two octaves** — for every even `N`: `ln N − ln(ceil(N/2)) = ln 2`	PROVED	Computed for `N = 4..64` even; loss = `0.6931471806` in every case, `loss / ln 2 = 1.0000000000`	none claimed	CCP-018, CCP-039
A-31	A star polygon `{N/k}` decomposes into exactly `gcd(N,k)` closed paths of length `N/gcd(N,k)`; unicursal figures are exactly those with `gcd(N,k) = 1`	PROVED	Elementary number theory; the full table computed for `N = 3..24`. The hexagram `{6/2}` has `gcd = 2` and is literally two triangles	The figures themselves, preserved in the practical corpus for centuries, with the stroke counts drawn correctly and the arithmetic unstated	CCP-028
A-32	Essence law: the published recipe `((AN/φ) · 1.125)^2` collapses exactly to `f(n) = k n^2` with `k = (9/8)^2/φ^2 = (81/64)/(φ+1) = 0.483425732988414 Hz`	PROVED	Algebraic reduction; verified against every published value in the 2022 source to 14 significant figures	*The Effective Method of Sonic Chemistry* (Curzi, 2022) recorded the recipe in spreadsheet form without the reduction	CCP-017; CC-R5
A-33	All intervals in the essence system are ratios of squared integers: `f(m)/f(n) = m^2/n^2`, and `k` cancels from every relation	PROVED	Arithmetic	none claimed	CCP-017, CCP-018
A-34	Two essences share a pitch class iff their indices have the same odd part; an ordered set `1..N` yields exactly `ceil(N/2)` pitch classes	PROVED	Arithmetic: `(2n)^2 = 4n^2`, exactly two octaves	none claimed	CCP-017, CCP-018, CCP-039
A-35	The cross-disciplinary bridge: `f_cross = k · N_eff^2 = k · e^{2F}`, so **one nat of surplus is exactly `2.885390` octaves**	PROVED	Conversion arithmetic; `F = ln 12` gives `N_eff = 12`, `f = 69.6133 Hz`, `7.1699` octaves above base	none claimed	CCP-017, CCP-018
A-36	The conversion keys are exact: `1 nat = 1.442695 bits = 2.885390 octaves = 3462.4681 cents`; `1 bit = 2 octaves exactly = 2400 cents`	PROVED	Arithmetic	none claimed	CCP-018; the corpus's internal-scope conventions
A-37	**The Lens Theorem:** a continuous map converting multiplication into addition is a logarithm, unique up to scale	PROVED (conditional)	Cauchy's logarithmic functional equation. The **antecedent** — that perception is such a map — is empirical and badged separately (C-33)	none claimed	CCP-020
A-38	Consequence: equal perceptual steps are equal world **ratios**; `+1` perceived unit is a fixed world multiple. **Linear in the lens is exponential in reality**	PROVED	Follows from A-37 by inversion: `I = I_0 e^{S/k}`	none claimed	CCP-020, CCP-039

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
A-39	Maximum possible moment is `ln N` once `M` is normalised to serve as `S_t`: the richest available moment is bounded by the log of how many kinds of thing the perceiver can distinguish	PROVED (given the normalisation)	Follows from A-09 and the normalisation choice, which is itself POSTULATED (B-17)	Ashby's requisite variety (1956) records the same bound for regulators, from control theory rather than perception	CCP-019, CCP-021
A-40	**The Gabor limit:** `dt · df >= 1/(4π) = 0.079577` for every wave, including sound; classical, with no quantum content	PROVED	Fourier analysis; window/resolution table computed: 1 ms → 79.577 Hz, 10 ms → 7.958 Hz, 100 ms → 0.796 Hz, 1000 ms → 0.080 Hz	none claimed	CCP-030, CCP-031, CCP-032
A-41	`dx · dp >= ħ/2` is a property of conjugate variables and a consequence of Fourier duality — **not** measurement disturbance	PROVED	The commutator derivation in standard quantum mechanics	none claimed	CCP-032
A-42	Casimir pressure `F/A = π²ħc/(240 a⁴)`. Computed: 10 nm → `1.3001e5` Pa (1.283 atm); 100 nm → `13.00` Pa; 1 µm → `1.30e−3` Pa; 1 mm → `1.30e−15` Pa	PROVED (formula)	Casimir (1948), derivation	none claimed	CCP-033
A-43	Brooks' path count: `n(n−1)/2` communication paths. **Cost grows quadratically while surplus grows logarithmically**	PROVED (arithmetic)	The count is arithmetic; that cost tracks it is BORROWED (C-30)	none claimed	CCP-008, CCP-012, CCP-038
A-44	Queueing: waiting time grows as `1/(1−ρ)`; utilisation near 1 makes wait explode. **Slack is a cost-control device, not waste**	PROVED (under M/M/1 assumptions, which real systems violate, generally for the worse)	Standard queueing theory	none claimed	CCP-008, CCP-012
A-45	Arrow's impossibility theorem: no rule aggregating individual orderings into a collective ordering satisfies the stated conditions simultaneously	PROVED	Arrow (1951). A theorem about preference orderings; this framework does **not** claim it applies verbatim to state vectors	none claimed	CCP-038
A-46	Condorcet's jury theorem: collective accuracy exceeds individual accuracy **only under independence** of member errors	PROVED	Condorcet (1785). The independence condition is the load-bearing part	none claimed	CCP-009, CCP-013, CCP-014, CCP-038
A-47	Rasch specific objectivity: conditional maximum likelihood separates person and item parameters, so a person's measure does not depend on which items were used	PROVED	Rasch (1960); the CML separation argument	none claimed	CCP-026
A-48	**the Rasch-route theorem [CCP-026] (given badged Postulate SO1):** if `D` and `P` are Rasch person-measures on one common scale, `Delta = |ln(D/P)| = |θ_D − θ_P|` is ratio-scaled with a true zero at perfect calibration	PROVED (conditional)	CCP-026. The conditional is on SO1 (B-19), not on the algebra	none claimed	CCP-026, CCP-038 §6.6
A-49	**The instrumentation cost is 123 items per term** to resolve a half-logit distortion at 95%	PROVED (arithmetic)	Best-case item information `0.25` → `SE(θ) = 2/\ensuremath{\surd}k`; `SE(Delta) = \ensuremath{\surd}2 · SE(θ)`; floor `= 1.96 · SE(Delta)`. `k=122` → `0.501905`; `k=123` → `0.499860`	none claimed	CCP-026, CCP-027, CCP-038, CCP-039

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
A-50	The four registers — surplus, pitch, sensation, moment — are all logarithmic, so every conversion between them is linear in log space with exact constants	PROVED (given each register's own form)	Arithmetic, conditional on A-01, pitch `= log2 f`, C-33, and B-16	none claimed	CCP-018; the corpus's internal-scope conventions

Fifty entries and not one of them is a fact about a person, a room, or a cell. That is what Group A is: the arithmetic the rest of the corpus borrows. Every application of it passes through Group B.

38.6.2 6.2 Group B — what is postulated

Assumptions. Chosen, badged, and load-bearing. **Every quantity in this corpus computed downstream of one of these inherits its status and not a grade higher.**

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
B-01	**S3, the affine design axiom:** `g(u) = e^{f(u)}` is affine in `u`, i.e. `N_eff = 1 + (N−1)u`	POSTULATED	**None. It is chosen for tractability and interpretability.** It is what makes A-01 produce a logarithm rather than another concave function	none claimed	Everywhere; the honesty appendix [CCP-039]
B-02	The **linear form** of the storage law: loss proportional to stock, conversion linear in surplus, cost additive, parameters constant within a period, no saturation	POSTULATED	None. It is the minimal tractable open-system law, chosen so a closed form exists	none claimed	Everywhere; the honesty appendix [CCP-039]
B-03	`Q` is a **scalar** state variable at all — rather than a vector of incommensurable stores	POSTULATED	None. There is no measurement of `Q`, only a proxy `Y` chosen under Protocol MP-6	none claimed	CCP-014, CCP-021, CCP-023; the honesty appendix [CCP-039]
B-04	The **modelling map**: that rooms, teams, institutions and cities are systems to which the storage law applies at all	POSTULATED	None, and **nothing in Papers A or E establishes it.** The largest assumption in the corpus	none claimed	Vol. One entire
B-05	**Hyperrecursive symmetry / the three-scope correspondence:** the same equation at desk, room, institution, mind, and cell, with only parameters changing	POSTULATED (model form only)	None. Not a scaling law; **never tested at two scales simultaneously**	none claimed	CCP-011, CCP-038; the honesty appendix [CCP-039]
B-06	The **state encoder**: participants, views or categories can be represented as unit vectors whose inner product carries the relevant alignment	POSTULATED	None. Different encoders give different `u` for the same group	none claimed	CCP-015, CCP-021; Protocol MP-4
B-07	The **ensemble aggregation**: the participation-ratio extension of a pairwise law to a group. **Consistency by construction is a property of the choice, not evidence for it**	POSTULATED	A-27 shows the two descriptions coincide, which is not the same as showing the choice is right	none claimed	CCP-015, CCP-038
B-08	**Cost commensuration**: coordination hours, rework, conflict and rent convert into the same unit as stored order	POSTULATED	None. The conversion factor is chosen by the engineer and dominates `h` whenever `C` is large	none claimed	CCP-008, CCP-023; the honesty appendix [CCP-039]
B-09	**Constancy of parameters** within a survey window: `δ` and `η` treated as constants	POSTULATED (known false, retained for tractability)	None. They move with turnover, tooling, morale and season	none claimed	Protocol MP-1; the honesty appendix [CCP-039]
B-10	`u = 0` at anti-alignment is **correct**: a reliable contrarian is one instrument, not two	POSTULATED (reading of the formalism)	The formalism gives `u = 1 − (x·y)^2`; the *interpretation* is this framework's	none claimed	CCP-009, CCP-010

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
B-11	**Optional activity** (Gehl) is the observable proxy for a room's η	POSTULATED	Gehl's category is observed and real; the identification with η is this framework's proposal	none claimed	CCP-010
B-12	**Acoustics as c-factor engineering**: changing RT60 changes collective intelligence	POSTULATED — **and see CC-R7, not confirmed**	A three-link chain whose middle link is weakest; untested end to end	none claimed	CCP-011, CCP-012, CCP-013, CCP-037
B-13	`k = 0.483425732988414 Hz` is a **chosen reference constant**, not a measured one; its form (Pythagorean interval over golden section) is a design decision	POSTULATED	None. Every relation in the system is independent of `k` (A-33), which is why the choice is admissible	The 2022 source document, which recorded the recipe without deriving the constant	CCP-017, CCP-018
B-14	`f = k n^2` is a **postulate about essence structure**. Its coincidence with the `n^2` eigenvalue spacing of a 1-D infinite square well is a mathematical correspondence and **not evidence of physical mechanism**	POSTULATED	None. Coincidence of form is not mechanism, and the text must say so in the same breath	none claimed	CCP-017
B-15	The **essence index requires a defensible ordering**; an arbitrary ordinal yields arbitrary frequencies	POSTULATED (validity condition)	CCP-015's validity conditions govern state encoding and category assignment	none claimed	CCP-017, CCP-031
B-16	**Fechner's postulate**: all just-noticeable differences are subjectively equal, which is what makes Weber's law integrate to a logarithm	POSTULATED	Not derivable from Weber's law; an additional assumption, disputed since 1860	Fechner (1860) stated it as a postulate himself	CCP-020
B-17	The **moment equation** `M = (D/P)·E` as the author's formula, with `M` normalised to `[0, ln N]` to serve as `S_t`	POSTULATED — **and DEMOTED on magnitude, see E-01**	The structure (limits, fixed point, ordering) needs only ordinal information and stands	none claimed	CCP-019, CCP-021, CCP-023, CCP-024, CCP-025
B-18	The **self-storage join** `Q_{t+1} = (1−δ)Q_t + η M_t − C_t`: the moment is the input to character	POSTULATED — **structural, not quantitative (CC-R1)**	Follows from B-02 plus B-17; inherits both	none claimed	CCP-021, CCP-023
B-19	**Postulate SO1**: that `D` and `P` admit Rasch-conformable item sets at all — that framing and intake can be written as items a person answers	POSTULATED	None. the Rasch-route theorem [CCP-026] is conditional on it	none claimed	CCP-026
B-20	**The Belief Postulate: belief creates reality if the belief is grounded in truth.** The conditional clause is the entire content	POSTULATED	Formalised as `BELIEF -> [CHANNEL] -> WORLD`, where "grounded in truth" means the channel exists and is real. Testable only through the channel	Long precedent across contemplative traditions for the unconditional form; **the unconditional form is what this corpus rejects**	CCP-034, CCP-035
B-21	Corollary: **a belief with no channel changes nothing**; belief is an input to a system that must have gain	POSTULATED (definitional consequence of B-20)	Same architecture as the storage law: input, conversion efficiency, output	none claimed	CCP-034, CCP-035
B-22	Reading esoteric operations as a **nonlinear computing map** is an explicit **metaphor**, carried with its badge	POSTULATED (metaphor, badged)	None. The metaphor generates no prediction on its own; A-31 does	The operative corpus recorded the procedures; it did not record them as computation	CCP-027, CCP-028
B-23	The **memory palace is the storage law with space as retention medium**: high η, low δ, because place is durable	POSTULATED (identification)	The mnemonic technique itself is validated (C-44); the parameter identification is this framework's	Yates, *The Art of Memory* (1966), recording the Ars Memoriae tradition from classical sources onward	CCP-028, CCP-029
B-24	**Ritual is periodic forcing of a low-pass filter**; cadence beats intensity because the filter integrates	POSTULATED (identification)	A-24 and A-25 are proved; that ritual *is* the forcing function is the identification	Durkheim on collective effervescence; Rappaport (1999) on ritual as the ground of social order	CCP-007, CCP-028
B-25	The **Tri-State Method** (+1 / -1 / 0p) and the three bridge criteria (structural correspondence, predictive power, practical utility), adopted unchanged	POSTULATED (method)	The author's VINO Framework. A validation apparatus, not a finding	none claimed	CCP-029, CCP-038

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
B-26	The **two-column refinement**: the columns carry different things, and precedent is never evidence	POSTULATED (method)	the cellular-scope material of the corpus's addition to VINO. **The rule that makes the rest of the register honest**	none claimed	CCP-029, CCP-038
B-27	The **no-concealment constraint**: no subliminal content of any kind, nothing masked, nothing below threshold. **A message hidden from its recipient cannot be consented to**	POSTULATED (house constraint)	An ethical constraint, not an empirical claim. It **overrides** the 2022 source's ASMR-subliminal recommendation	The 2022 source recommended embedded subliminals; **this research programme refuses it**	CCP-018, CCP-036, CCP-037
B-28	The **consent constraint**: no intervention that removes a person's ability to decline, and consent attaches to the person, never to the scope	POSTULATED (house constraint)	Structural: a method working by removal of resistance has no coercive failure mode. **An argument about method structure, not an empirical safety claim**	none claimed	the standing refusals of the corpus; CCP-013, CCP-014, CCP-037, CCP-038

38.6.3 6.3 Group C — what is borrowed

Empirical results obtained by other people. **The "what it measured" field is where the objections live, and it is written to make them easy.** Nothing here was obtained in the reader's setting.

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
C-01	A single factor c predicts group performance across varied tasks, is **not** predicted by members' average or maximum IQ, and correlates with social sensitivity, turn-taking equality, and proportion of women	BORROWED	Woolley, Chabris, Pentland, Hashmi & Malone, *Science* (2010). Several hundred participants, groups of 2–5, batteries of short laboratory tasks, factor analysis. **Transfer assumption: that a factor measured over hours in assembled groups persists in teams working for years**	none claimed	CCP-013, CCP-011, CCP-012, CCP-014
C-02	A committed minority of about **25%** overturns an established social convention	BORROWED	Centola, Becker, Brackbill & Baronchelli, *Science* 360 (2018). Controlled online coordination games; the tipping point was **measured, not simulated**. **Setting: strangers, arbitrary convention, no material stakes, no hierarchy**	none claimed	CCP-013, CCP-037
C-03	Complex contagion needs multiple reinforcing contacts, so **clustered topology beats random long ties**	BORROWED	Centola & Macy, *AJS* (2007), analytic; Centola, *Science* (2010), online experiment in which a health behaviour spread farther through a clustered lattice than a random network of equal density	none claimed	CCP-013, CCP-014
C-04	Threshold models: aggregate behaviour depends on the whole threshold distribution, so aggregate discontinuity requires no irrationality	BORROWED (model)	Granovetter, *AJS* (1978). **A model, not a measurement**	none claimed	CCP-013
C-05	Small-world: rewiring a small fraction of a lattice's edges collapses path length while clustering stays high	BORROWED (constructed graphs)	Watts & Strogatz, *Nature* (1998)	none claimed	CCP-013, CCP-014
C-06	Growth plus preferential attachment yields power-law degree; such networks tolerate random failure well and targeted hub removal very badly	BORROWED	Barabási & Albert, *Science* (1999); Albert, Jeong & Barabási, *Nature* (2000). Model results with empirical support in several real networks. **Check whether the dependency graph in question is scale-free before importing the fragility result, by the estimation-and-goodness-of-fit procedure of Clauset et al. [2009]; Broido and Clauset [2019] find scale-free structure rare once that procedure is applied, and Stumpf and Porter [2012] and Doyle et al. [2005] bound the inference further**	none claimed	CCP-009, CCP-014

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
C-07	Under stated conditions a randomly selected diverse group can outperform a group of the individually best	BORROWED — **contested**	Hong & Page, *PNAS* (2004). A mathematical and computational result depending on strong conditions; **criticised on mathematical grounds by Thompson, *Notices of the AMS* (2014).** Used as an illustration that diversity *can* dominate, never that it generally does	none claimed	CCP-009, CCP-014
C-08	Informational cascades: sequential agents rationally stop using private signals after a short prefix; aggregation stops and the cascade is fragile	BORROWED	Bikhchandani, Hirshleifer & Welch, *JPE* (1992); theory with extensive experimental support	none claimed	CCP-010, CCP-012, CCP-014
C-09	Contact reduces prejudice under equal status, common goals, cooperation, and institutional support	BORROWED	Allport (1954); Pettigrew & Tropp, *JPSP* (2006), meta-analysis of 500+ studies. **Largely correlational at base, with known publication-bias and self-selection concerns**	none claimed	CCP-013, CCP-014
C-10	Harm at scale is preceded by euphemistic labelling, displaced responsibility, routinisation, and withdrawal of full humanity from a target	BORROWED	Bandura on moral disengagement; Kelman; Haslam on dehumanisation. A well-supported tradition, drawn on **only** for the ethical constraint and the CCP-014 failure modes	none claimed	CCP-014; the standing refusals of the corpus
C-11	Eight design principles characterise long-enduring common-pool resource institutions	BORROWED	Ostrom, *Governing the Commons* (1990); Nobel 2009. **Inductive regularities from comparative case study — not theorems, not experiments.** Broad support in Cox, Arnold & Villamayor-Tomás (2010) with suggested refinements	none claimed	CCP-008, CCP-014
C-12	Regional governments created by one 1970 devolution performed very differently, tracking civic traditions with roots centuries earlier	BORROWED — causal direction contested	Putnam, *Making Democracy Work* (1993)	none claimed	CCP-014
C-13	Order can be produced by networked obligation where central enforcement is weak	BORROWED (descriptive)	Gambetta (1993), on the Sicilian mafia as a supplier of private protection. **Used as a demonstration and as a warning about what kind of order that is**	none claimed	CCP-014
C-14	Stigmergy: coordination by modifying a shared medium, with no agent holding the plan	BORROWED	Grassé (1959), termite nest construction; a foundational ethological observation since formalised. **The design rule "put the state in the medium" is this framework's inference, not Grassé's finding**	none claimed	CCP-014
C-15	Focal points: parties who cannot communicate converge on salient options. Mild individual preferences produce strong aggregate separation	BORROWED	Schelling, *The Strategy of Conflict* (1960) — evidence was informal elicitation, not controlled experiment; "Dynamic Models of Segregation" (1971) — a model	none claimed	CCP-014
C-16	Infrastructure is relational and becomes visible chiefly on breakdown; therefore δ **is systematically underestimated wherever it has not been measured**	BORROWED	Star & Ruhleder, *Information Systems Research* (1996); ethnographic and theoretical study of a scientific collaboratory	none claimed	CCP-007, CCP-012, CCP-014
C-17	"The most profound technologies are those that disappear"	BORROWED (aphorism)	Weiser (1991), *Scientific American*. **An essay, an argument rather than a study**	none claimed	CCP-012
C-18	Living structures have overlapping rather than strictly nested categories; designers impose trees because trees are easy to hold in mind	BORROWED — supported by examples, not measured	Alexander, "A City is Not a Tree" (1965); *Notes on the Synthesis of Form* (1964); *A Pattern Language* (1977). **The use of it to explain why hierarchy drives u to zero is this framework's inference, not Alexander's claim**	none claimed	CCP-009, CCP-010, CCP-012
C-19	Eyes on the street, short blocks, mixed primary uses, aged buildings	BORROWED (observational argument)	Jacobs (1961). Later quantitative work supports parts and complicates others	none claimed	CCP-010, CCP-012
C-20	People sit where there is somewhere to sit; sun, food, water and triangulation draw and hold them	BORROWED	Whyte (1980), time-lapse observation of New York plazas over several years. **The method transfers better than the findings, and this framework borrows the method**	none claimed	CCP-010
C-21	Necessary, optional and social activity, with optional activity highly sensitive to conditions	BORROWED	Gehl (1971), observational studies across several cities	none claimed	CCP-010, CCP-012

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
C-22	Proxemic distance bands: intimate, personal, social, public	BORROWED — **a frame, not a specification**	Hall (1966); observation predominantly of one cultural group, with Hall himself emphasising cross-cultural variation	none claimed	CCP-010, CCP-012
C-23	Post-surgical patients with a window view of trees had shorter stays, fewer negative nursing notes, and fewer strong analgesic doses	BORROWED	Ulrich, *Science* (1984). **Twenty-three matched pairs, one Pennsylvania hospital, nine years.** A small single-site matched-pair study, **and the strongest single piece of evidence in this framework that physical space affects a hard outcome.** Both halves are load-bearing	none claimed	CCP-012
C-24	`RT60 = 0.161 V/A` (metric)	BORROWED (empirical formula)	Sabine. **Valid for reasonably diffuse fields; degrades in heavily absorbing rooms (Eyring preferred) and strongly non-diffuse geometries**	none claimed	CCP-011, CCP-018
C-25	Organisations produce designs copying their own communication structures	BORROWED	Conway (1967), a trade-magazine argument, later supported by empirical comparison of product and organisational structure. **A structural explanation, not a law**	none claimed	CCP-007, CCP-013
C-26	A stable social-group size near 150	BORROWED — **contested; do not use as a design threshold**	Dunbar; a regression of primate group size on neocortex ratio extrapolated to humans, with very wide uncertainty	none claimed	CCP-008
C-27	A measure used for control ceases to measure what it did	BORROWED (empirical generalisation)	Goodhart (1975); Campbell (1979); Strathern (1997). **This framework is subject to it: headroom, once rewarded, will be gamed**	none claimed	CCP-015, CCP-016; the honesty appendix [CCP-039]
C-28	Adding people to a late project makes it later	BORROWED	Brooks (1975); engineering argument from experience on one large project, widely corroborated in practice, not established by controlled experiment	none claimed	CCP-008
C-29	Only variety absorbs variety	BORROWED (control theory result)	Ashby (1956), the Law of Requisite Variety. A theorem in its own setting, imported here as a design principle	none claimed	CCP-009, CCP-019
C-30	The just-noticeable difference is a constant **ratio** of the stimulus: `dI/I = c`	BORROWED	Weber (1834); psychophysical threshold measurement, replicated continuously since, **over a working range and not at the extremes**	none claimed	CCP-020
C-31	Sensation is logarithmic in intensity: `S = k ln(I/I_0)`	BORROWED (given B-16)	Fechner (1860); integration of Weber's law under the equal-JND postulate	none claimed	CCP-020, CCP-039
C-32	**Magnitude estimation produces genuine ratio-scale measurements of subjective experience, and it replicates.** `ψ = k I^a`; taking logs, the exponent `a` **is a measured conversion key**	BORROWED — **the single most important body of evidence for the personal-scope material of the corpus's premise**	Stevens (1957) and the subsequent literature. Nominal exponents: loudness ~0.67, brightness ~0.33, visual length ~1.0, salt taste ~1.3, lifted heaviness ~1.45, warmth ~1.6, electric shock ~3.5. **Exponents vary with method, range and instruction: constants of a procedure, not of nature**	none claimed	CCP-020, CCP-026
C-33	Human perception behaves as a ratio-reporting device over working ranges — the **antecedent** of the Lens Theorem	BORROWED	The Weber–Fechner–Stevens corpus (C-30, C-31, C-32). **The badge in CCP-020 is on this antecedent, not on the inference**	none claimed	CCP-020
C-34	Of 100 original psychology studies, 97 reported significant results and about **36%** of replications did, with effect sizes roughly halved	BORROWED	Open Science Collaboration, *Science* (2015). **Cited to state the measurement and replication problem honestly, not to dismiss the field**	none claimed	CCP-026, CCP-027; the honesty appendix [CCP-039]
C-35	Perception is modelled as precision-weighting of top-down priors against bottom-up evidence; **the ratio of prior to sensory precision is structurally the same object as** `D/P`	BORROWED — **active and contested research programme; structural similarity is not identity**	Friston; Clark; Hohwy. A modelling tradition with supporting evidence, not a settled fact	none claimed	CCP-019, CCP-034
C-36	**Placebo analgesia is mediated by endogenous opioids and is naloxone-reversible**	BORROWED	Levine, Gordon & Fields (1978). **A complete, physical, verified channel: belief → expectation → neurochemistry → measured analgesia**	none claimed	CCP-034
C-37	**Open-label placebo works** — benefit is observed when patients are told they are receiving placebo	BORROWED — **effect sizes modest, literature active**	Kaptchuk et al. (2010), IBS, and subsequent trials. **Decisive for this research programme: the effect does not require deception, so it is compatible with the no-concealment constraint (B-27)**	none claimed	CCP-034, CCP-036, CCP-037

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
C-38	Self-fulfilling prophecy: a false definition of a situation can evoke behaviour making it true	BORROWED	Merton (1948). **The bank-run example works because banks really are fractional-reserve: the belief propagates because it is coupled to a true feature of the system. Merton's mechanism IS the conditional clause of B-20, discovered independently in 1948**	none claimed	CCP-035
C-39	Expectancy effects in performance and clinical settings are real and modest	BORROWED — **contested history**	Rosenthal & Jacobson (1968) drew serious methodological criticism; the later literature is mixed. Presented honestly rather than as support	none claimed	CCP-035
C-40	**Growth-mindset effects are very small in large replications; stereotype-threat findings have had replication difficulty**	BORROWED — **a failure entry, and it is here on purpose**	Sisk et al. (2018), meta-analysis. **A postulate that names where it fails can be trusted where it holds**	none claimed	CCP-035
C-41	**PIEZO1 and PIEZO2 are mechanically gated ion channels.** Mechanical force opens them; ions flow; signalling follows	BORROWED	Nobel Prize in Physiology or Medicine, 2021 (Patapoutian). **Sound is mechanical, so a real, established pathway exists from pressure wave to cellular signal**	none claimed	CCP-030, CCP-038
C-42	Force applied at the cell surface transmits through the cytoskeleton to the nucleus, stretches chromatin, and upregulates transcription	BORROWED	Tajik et al., *Nature Materials* (2016). **Mechanical input to gene expression, published and direct**	none claimed	CCP-030, CCP-031
C-43	Ultrasound can activate specific channels in living organisms	BORROWED	Ibsen et al., *Nature Communications* (2015), TRP-4 in *C. elegans*; Duque et al. (2022), hsTRPA1	none claimed	CCP-030
C-44	The method of loci is a validated mnemonic technique	BORROWED	The experimental mnemonics literature	Yates, *The Art of Memory* (1966), recording the Ars Memoriae from Simonides through the Renaissance — **centuries of practical use with no experimental apparatus**	CCP-028, CCP-029
C-45	Arbitrary digital data can be encoded in and recovered from synthesised DNA	BORROWED	Church, Gao & Kosuri, *Science* (2012)	none claimed	CCP-030, CCP-031
C-46	Cas12a processes its own crRNA array — a genuine multicast destination list	BORROWED	Zetsche et al., *Cell* (2015); Fonfara et al. (2016)	none claimed	CCP-030
C-47	**Targeting and editing are separable functions**: dCas9 binds a locus and CRISPRa activates transcription while cutting nothing	BORROWED	Qi et al. (2013); Gilbert et al. (2014). **This is what makes "addressing without editing" a real architecture rather than a promise**	none claimed	CCP-030, CCP-036
C-48	Four regulatory factors return a differentiated adult cell to pluripotency **without adding one base of new sequence**	BORROWED	Yamanaka; Nobel Prize 2012. **Latent capability and its activation, demonstrated**	none claimed	CCP-031, CCP-036
C-49	**Syncytin**, indispensable to placenta formation, derives from a retroviral envelope gene	BORROWED	Molecular and comparative genomics. **Ancient foreign text, inherited, later read out as vital function**	none claimed	CCP-031, CCP-036
C-50	~98% of the human genome is non-coding and largely regulatory; ~45% transposable elements, ~8% endogenous retrovirus. **"Junk DNA" is retired**	BORROWED	Genome annotation and comparative analysis; the fractions are census figures with definitional caveats about what counts as functional	none claimed	CCP-031
C-51	Median mRNA half-life ~9 h and median protein half-life ~46 h	BORROWED — **the only measured parameter at the cellular scope**	Schwanhäusser et al., *Nature* (2011). Metabolic labelling with mass spectrometry and sequencing, **one mouse fibroblast line, medians over thousands of species with enormous spread.** Converted: $\delta = 0.074125/h$ and $\delta = 0.014955/h$	none claimed	CCP-030, CCP-038
C-52	The Casimir force is real and measured	BORROWED	Lamoreaux (1997); Mohideen & Roy (1998). **At biological and architectural scale it is utterly negligible: 1.3e-15 Pa at 1 mm. It does nothing in a cell and nothing in a room**	none claimed	CCP-033

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
C-53	Environmental interaction destroys phase coherence extraordinarily fast, so macroscopic superposition does not persist	BORROWED — **the honest objection, included rather than avoided**	The decoherence literature and its timescale estimates. **Any claim that mind collapses wave functions must confront this**	none claimed	CCP-033
C-54	Cochlear hair cells and tip links are exquisite mechanical receivers	BORROWED	Auditory physiology. **Biology's existence proof that fine mechanical reception is buildable**	none claimed	CCP-030
C-55	Exit and voice are the two responses available to a member of a declining organisation	BORROWED (analytic framework)	Hirschman (1970). **A structural distinction, not a measurement. It is what makes "members can exit and cells cannot" precise rather than rhetorical**	none claimed	CCP-038
C-56	Group-level relationships need not hold at individual level and can reverse	BORROWED	The ecological fallacy; Simpson's paradox as its arithmetic form. **A_u across teams and a_u within a team are different quantities the notation does not distinguish**	none claimed	CCP-038
C-57	Ritual as the ground of social order; collective effervescence as an observed phenomenon of assembly	BORROWED (theoretical and ethnographic)	Rappaport, *Ritual and Religion in the Making of Humanity* (1999); Durkheim; costly-signalling accounts in evolutionary anthropology	The liturgical corpora themselves, which recorded cadence, repetition and calendar with great precision and no theory of filters	CCP-028, CCP-029
C-58	A mind can be described as an assembly of interacting parts	BORROWED — **an argument, not a body of evidence**	Minsky, *Society of Mind* (1986). Cited as precedent for the picture, **not as support for it**, and CCP-038 argues against importing the sociological results with it	none claimed	CCP-038

38.6.4 6.4 Group D — what is speculative

Bridges and proposals that have not met all three criteria of CCP-029. **Every one of them is marked speculative in the sentence that states it, in its home chapter as well as here**, and two of them are rejections rather than proposals — a method that never rejects anything is a decoration.

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
D-01	That a cell can demodulate **sequence-specific** information from an acoustic signal	SPECULATIVE — **and this is the volume's central open problem**	**None. There is no established mechanism.** Everything upstream is real (C-41, C-42, C-43, C-45, C-46, C-47) and none of it establishes this step. Bridge criteria: C1 partially, C2 yes, C3 not yet	none claimed. The old corpus records sound as efficacious; it records no mechanism, and a record of belief is not a record of mechanism	CCP-030, CCP-031, CCP-036
D-02	That the essence law's index applies **outside** chemistry, to any discipline's ordered category set	SPECULATIVE (badged)	The law reduces exactly (A-32) and the cross-disciplinary bridge is arithmetically exact (A-35). **What is speculative is that a discipline's categories admit a defensible ordinal at all** (B-15)	The 2022 source applied it to the periodic table only	CCP-017, CCP-018, CCP-031
D-03	That the uncertainty of a future event is the same uncertainty as Heisenberg's	SPECULATIVE — **an analogy with its badge on, not an identity**	Shared: conjugate structure, no simultaneous sharpness, the Fourier trade-off. **Not shared: ħ, quantisation, non-commuting operators, or any experimental result**	the founding material, the wave node, which treats the future as an amplitude-weighted ensemble rather than a verdict — **and needs no quantum mechanics to stand**	CCP-032
D-04	That if magic is the engineering of the perceptual lens ($\mathfrak{D}/\mathfrak{P}$), the storage of order through symbol, space and cadence (η, $\mathfrak{d}$, $\mathfrak{Q}$), and the composition of non-parallel essences ($\mathfrak{u}$, $\mathfrak{F}$), then this framework describes its mechanism in modern notation	SPECULATIVE — the author's thesis, badged	**No supernatural efficacy is claimed, no physical effect of symbol, and no result not derivable from the formalism.** The thesis is a re-description, and a re-description generates no new prediction by itself	The operative corpus, which recorded procedures, correspondences and cadences in great detail across centuries	CCP-027–60

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
D-05	Narrative encoded as data, as sequence, as tone, with structure carried in the harmonic, acting as regulatory input	SPECULATIVE — **inherits D-01 entirely**	The encoding pipeline is built and reproducible; **the biological read is not established**	none claimed	CCP-031, CCP-036
D-06	*Destiny made manifest*: latent capability brought into expression by addressing plus belief with a real channel plus narrative input	SPECULATIVE — inherits D-01 and D-05	C-48 establishes that latent capability can be activated by regulatory signal; **nothing establishes that this architecture is such a signal.** The term is defined at first use and **explicitly distinguished from the 19th-century American political slogan and its history of justifying displacement**	none claimed	CCP-036
D-07	That consciousness causes wave-function collapse	**REJECTED in print**	Von Neumann–Wigner is **a minority historical interpretation, not a mainstream position.** C-53 is the objection it must confront and does not survive. The measurement problem is genuinely open; **it is not open in the direction of licence**	The contemplative corpus records mind as participating in the world; that record is not a physical mechanism	CCP-033
D-08	That the Rodin coil functions as an amplifier of the essence-tuning method	**REJECTED in print**	**No verified engineering basis.** Citing it as established would fail the standing refusals of the corpus.3	The 2022 source mentions it	CCP-018
D-09	That embedded subliminal content improves the method	**REFUSED on constraint, not on evidence**	Refused under B-27 regardless of efficacy. **A message hidden from its recipient cannot be consented to**	The 2022 source recommends it; **these chapters do not carry that recommendation**	CCP-018, CCP-036
D-10	The bridge rejected in public in CCP-029 as a demonstration of the method	**REJECTED in print, deliberately**	It was chosen to sit one inch from this corpus's own thesis, because rejecting something the corpus has no interest in would prove nothing	none claimed	CCP-029
D-11	That a composition operator exists under which the storage law is closed, so cellular order aggregates to personal order aggregates to institutional order	SPECULATIVE — **and it is an open mathematical question, stated as one**	**No such theorem exists in this corpus.** No aggregation operator is defined between scopes, no conservation law links them, and no argument is offered that the three accumulators are coupled at all rather than merely similar in form	none claimed	CCP-038, CCP-039
D-12	That gematria's string-to-integer map carries meaning	**NOT CLAIMED — the description is technical only**	**Gematria is a hash function. That is a literally accurate description of the operation and it carries no claim that the map is meaningful**	Extensive traditional use as an interpretive method	CCP-027

38.6.5 6.5 Group E — what is demoted

Claims this corpus advanced and its own protocol rejected or reduced. They are printed here in the same typeface as everything else, because a register that hides its demotions is not an audit.

ID	Claim	Status	(+1) evidence — and what it measured	(-1) precedent — and what it recorded	Where
E-01	The **magnitude** of `D/P` in the Moment Equation `M = (D/P)·E`	**DEMOTED** — FVP gate 2 (scale type) returns DEMOTE	**`D` and `P` are not instrumented on ratio scales with true zeros, and division requires one.** Route out: A-48 at the cost of A-49	none claimed	CCP-019, CCP-023, CCP-026; CC-R1
E-02	The **distortion magnitude** `Delta = |ln r|`	**DEMOTED** to ordinal	Direction stands — over-framed / calibrated / under-framed is an ordering and needs only ordinal information. **Magnitude does not.** A `Delta` computed from Likert data is **void, not imprecise**	none claimed	CCP-019, CCP-023; CC-R1
E-03	The **quantitative join** `M_t -> Q_{t+1}` in the self-storage law	**DEMOTED** — structural, not quantitative	The join's *form* stands; its *magnitude* inherits E-01	none claimed	CCP-021, CCP-023; CC-R1
E-04	The same demotion at **group scope**: the doctrine-to-data ratio as a measured quantity	**DEMOTED** to ordinal	The diagnostic — over-framed institution, under-framed institution — is ordinal and stands. The number does not	none claimed	CCP-024, CCP-025; CC-R1
E-05	Every M-derived formulary entry: the launch budget `L = max(0, M − C)`, the return difference `Delta_R = |ln(M'/M)|`, and their interpersonal counterparts	**DEMOTED** — gate 2 inherited	Each carries an FVP row showing gate 2 = DEMOTE plus the Rasch route. **`rho` computed from Likert data is void**	the founding material, the opening node (the departure lounge: "no surplus to dream with", `L = 0`) and a node ECHO recorded the structures the formulas formalise	CCP-022, CCP-023, CCP-025
E-06	The acoustic → `c`-factor **causal** claim	**DEMOTED — reported as NOT CONFIRMED**	The falsification test is stated in §14.13 and **its result is not confirmed.** Any later chapter invoking acoustic `c`-factor engineering must carry that status	none claimed	CCP-011, CCP-012, CCP-013, CCP-037; CC-R7
E-07	That the **shuffled-input control is sufficient** to detect a spurious input	**DEMOTED** — shown insufficient by construction	A filtered repackaging of the input passed all four baselines including the shuffle at `p = 0.0005`, and was caught **only** by the structural ceiling test (A-20): its single-step jump was **38.5×** the theoretical maximum	none claimed	CCP-015, CCP-027, CCP-035; CC-R4
E-08	Seven published compound values in the 2022 sonic-chemistry source	**DEMOTED — data-entry errors, not a different law**	Reverse-engineering recovered `f = k^2 G^4` (`G` = count-weighted geometric mean of indices), matching **37 of 44** published composites to `<1e−12`. **The O2 row is provably impossible under the stated rule**	The source document's prose formula `1+5^.5/2` is missing its parentheses	CCP-017, CCP-018, CCP-031; CC-R5

38.6.6 6.6 The reconciliation ledger, in full and unsoftened

Seven findings emerged during drafting and had to be propagated backward into chapters already written. **They are the clearest evidence in the corpus that the protocol is being run rather than displayed**, and they are given here at length rather than as footnotes to the rows above.

CC-R1 — The Moment Equation is demoted pending instrumentation

This is the most important entry in this section and it concerns the author's own keystone formula.

The finding. The Formula Validation Protocol's second gate asks a single question: what scale type does each term admit? Division and multiplication require ratio scale with a true zero; without one, a quotient is not a quantity. Applied to `M = (D/P)·E`, gate 2 returns **DEMOTE**. `D` and `P` — perspective and perception — **are not currently instrumented on ratio scales with true zeros.** Therefore three things are **not reportable as measured quantities**:

```
the magnitude of D/P
the distortion Delta = |ln r|
the join M_t -> Q_{t+1}                    (CCP-021)
```

What survives, and it is not nothing. The formula remains a valid **structural** claim. The limiting cases stand: `P -> 0` gives `M -> infinity`, frame with no intake, which is dogma and apophenia; `D -> 0` gives `M -> 0`, intake with no frame, sensation that is not experience; `E -> 0` gives `M -> 0`. The fixed point stands: `D = P` gives `M = E`, undistorted contact, and it is the formula's most important content. The ordering stands: over-framed, calibrated, under-framed. **All of these need only ordinal information, and ordinal information is what is available. What is demoted is magnitude, not structure.**

The constructive route out, which is the point. CCP-026, the Rasch-route theorem [CCP-026], given badged Postulate SO1: if `D` and `P` are measured as Rasch person-measures on one common scale, then

```
Delta = |ln(D/P)| = |theta_D − theta_P|
```

is genuinely ratio-scaled, with a true zero at perfect calibration and a free multiplicative unit. **Rasch separation supplies exactly the true zero that division requires.** The raw ordinal responses never became ratio data; the measure derived from them did, because the model earned it.

The cost, computed. Best-case item information is `0.25` when every item is targeted at the person, so `SE(theta) = 2/sqrt(k)`; the distortion is a difference of two independent measures, so `SE(Delta) = sqrt(2) · SE(theta)`; the 95% detection floor is `1.96 · SE(Delta)`.

```
k = 20   ->  floor 1.239613 logits      (r = 3.45; anything subtler is noise)
k = 31   ->  floor 0.995681             (one logit)
k = 100  ->  floor 0.554372
k = 122  ->  floor 0.501905             NOT ENOUGH
k = 123  ->  floor 0.499860             SUFFICIENT
k = 342  ->  floor 0.299770             (0.3 logits)
```

One hundred and twenty-three items per term to establish a half-logit distortion. A practitioner reporting `r = 1.2` from a twelve-item pair of scales is reporting noise, and the arithmetic says so before the study is run rather than after.

This is presented as the volume's actual research instruction, not as a caveat. The chapters that must carry it are named in the ledger and carry it: CCP-019 at first presentation of magnitude; CCP-019 with `Delta` as ordinal; CCP-021 with the join marked structural; CCP-023 and

CCP-025 with gate 2 = DEMOTE in every M-derived FVP row; CCP-024 at group scope; CCP-027 with the 123-item cost absorbed into the costed stage plan; CCP-034 where the belief postulate inherits the same instrumentation requirement; and the honesty appendix [CCP-039] in the postulated and untested drawers.

A validation protocol that demotes its own author's keystone formula, states the cost of fixing it, and gives the constructive route out is the strongest possible demonstration that the protocol is real. It is printed at full strength here for that reason. A reader who wants to know whether this corpus's gates do anything now has an answer that does not depend on trusting the corpus.

CC-R2 — The Jensen correction on feasibility

Because `F` is strictly concave (A-05), feasibility must be tested on the **mean of `F(U_t)`**, not on `F` of the mean `U`. Testing the wrong one overstates available surplus by the Jensen gap, which was worth **6.2%** of surplus in the CCP-016 worked office. Propagated to CCP-008, CCP-015, CCP-023, CCP-025 and Appendices A and C. **This is a small correction with a large moral: the concavity that makes the law merciful about the second opinion also makes it unforgiving about averaging.**

CC-R3 — `N_eff` is the participation ratio

Derived and verified: the mean-pairwise `u` aggregator equals the Rényi-2 participation-ratio effective count, `N_eff = N + 1 − N/PR`. Independent re-verification for this section over six random ensembles gave a worst discrepancy of `1.8e−15`. Recorded as a named identity in CCP-009, CCP-021 and Appendices A and B. **It matters because it means the ensemble extension is not a new assumption on top of the pairwise law; it is the same assumption written twice** — which is precisely why B-07 warns that consistency by construction is not evidence.

CC-R4 — The shuffle cannot see repackaging

A filtered repackaging of the input passed **all four** preregistered baselines — persistence, direct-input, matched autoregressive, and the shuffled-input control — with `p = 0.0005`. It was caught only by the structural ceiling test of Corollary 4.2: **its single-step jump was 38.5 times the theoretical maximum `η ln N`.**

```
STRUCTURAL CEILING TEST
r_max      = eta ln N − delta Q − C
hard cap   = eta ln N                        (delta, C >= 0)
if observed step > eta ln N, the result is impossible under the model
-- and impossibility is not a p-value, it is a contradiction
```

Rule, propagated to CCP-015 §6.3, CCP-027, CCP-035 and Appendices C and E: the shuffle is an empirical audit; only the structural tests detect repackaging. This is the register's clearest case of a demotion that strengthens the framework — a book with only empirical baselines would have published the repackaging as a result.

CC-R5 — Sonic chemistry errata

`f = k n^2` reproduces all 118 element values in the 2022 source, worst relative deviation `9.0e−15`, with the residual bias traced to `φ` stored at 15 significant figures. The compound rule was reverse-engineered as `f = k^2 G^4`, with `G` the count-weighted geometric mean of indices, matching **37 of 44** published composites to `<1e−12`. **The seven outliers are data-entry errors, not a different law — the O2 row is provably impossible under the stated rule**, and the source's

prose formula `1+5^.5/2` is missing its parentheses. CCP-017, CCP-018 and 69 use the corrected values and cite the errata. **Correcting a source is a service to it; the law survived the correction and the seven rows did not.**

CC-R6 — The `h_inf` identity

`h_inf = η(ln N − F(u*))`: at the fixed point, `δ`, `C` and `Q` all cancel. Verified to six decimal places in CCP-008 and re-verified here across four parameter sets with agreement below `1e−16`. Consequence: since `F(u) <= ln N` always, `h_inf >= 0` always, so **the plain linear model never ossifies by itself.** Ossification requires something the model does not contain — exogenous `Q` accretion, or superlinear cost. Propagated to CCP-014, CCP-016 and the reference card [CCP-039]. **A model that cannot produce the phenomenon everyone observes is reporting that the mechanism is outside the model, and that is a useful thing for a model to say.**

CC-R7 — The c-factor causal test is not confirmed

The falsification test for the acoustic → c-factor causal claim is stated in §14.13 and **reported as not confirmed.** Any chapter invoking acoustic c-factor engineering carries that status: CCP-012, CCP-013, CCP-037. The chain has three links — reverberation affects intelligibility (established), intelligibility affects who speaks (plausible, partially studied), turn-taking equality correlates with `c` (C-01) — and **the composite claim has not been tested end to end.** The test is cheap: randomise RT60 across otherwise matched deliberation rooms and measure turn-taking variance and task performance. It has not been run.

38.6.7 6.7 The status census

Counted directly from the register above.

Group	Entries	Contents
A — the formal spine	50	Theorems and exact arithmetic. No claims about the world
B — postulated	28	Design axioms, modelling choices, identifications, house constraints
C — borrowed	58	Empirical results obtained by others, in settings that are not yours
D — speculative	12	Of which **4 are rejections or refusals printed as such**
E — demoted	8	Claims this corpus advanced and its own protocol reduced
Total	**156**	

Three ratios are worth reading off it.

Fifty proved entries and none of them is about a person. Group A is entirely arithmetic. The book's mathematical content is real and it is also, on its own, inert: every step from a theorem to a recommendation passes through Group B.

Twenty-eight postulates carrying fifty-eight borrowed results. That is the actual load path. The borrowed results were obtained in laboratories, plazas, hospitals, fisheries and cell lines;

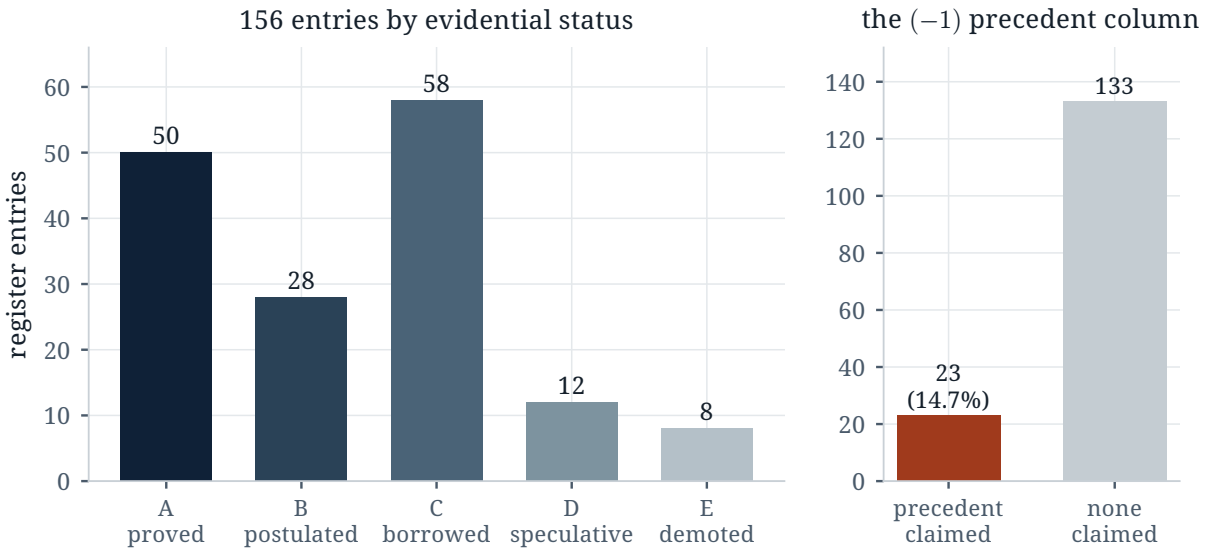


Figure 38.5: **Fifty proved entries say nothing about any person, and the entire load passes through twenty-eight postulates; precedent is claimed in 23 of 156 rows.** The register's five statuses form an ordered scale from proved to demoted and are drawn in one hue accordingly. The precedent column reads *none claimed* in 133 rows, which is the yield of a two-column method when the rule against filling the second column is enforced. Direct count of the register in Section 6; generated by `figs/src/CCP-038_fig04.py`.

the postulates are what connect them to a reader's room. **To break this corpus, break a postulate — B-01, B-02, B-04 and B-08 carry more weight than any citation in Group C.**

Twenty-three entries carry anything at all in the (-1) column, out of one hundred and fifty-six. Precedent is claimed in 14.7% of rows; the remaining 85.3% read *none claimed*. **That is the honest yield of the two-column method when the rule against filling the column is actually enforced**, and it is a considerably less impressive number than the method promises in the abstract. It is printed because it is the true one. the cellular-scope material of the corpus's commitment that every substantive claim carries two references is met where the historical corpus has something to say and is reported as unmet where it does not — which is the only way that commitment can be kept without inventing the second reference.

38.6.8 6.8 Index by status

PROVED (Group A). Surplus law and its properties: A-01 to A-13. Storage law and its consequences: A-14 to A-26. Identities established by this framework: A-26 (`h_inf`), A-27 (participation ratio), A-28 (Jensen), A-29 (Fechner threshold), A-30 (octave-folding cost). Essence law: A-32 to A-36. Perception: A-37 to A-39, A-50. Physics and mathematics imported as theorems: A-31, A-40 to A-47. Measurement arithmetic: A-48, A-49.

POSTULATED (Group B). Design axioms: B-01, B-02, B-03. Modelling map and scope: B-04, B-05. Estimation choices: B-06 to B-10. Identifications: B-11, B-12, B-23, B-24. Sonic postulates: B-13 to B-15. Perception and moment: B-16 to B-19. Belief: B-20, B-21. Method and metaphor: B-22, B-25, B-26. House constraints: B-27, B-28.

BORROWED (Group C). Sociology and networks: C-01 to C-15, C-25 to C-29, C-55, C-56, C-58. Space and substrate: C-16 to C-24. Psychophysics and psychometrics: C-30 to C-35. Belief channels: C-36 to C-40. Cell biology: C-41 to C-51, C-54. Physics: C-52, C-53. Tradition read as record: C-44, C-57.

SPECULATIVE (Group D). Open: D-01, D-02, D-03, D-04, D-05, D-06, D-11. Rejected or re-

fused in print: D-07, D-08, D-09, D-10. Description only, no claim: D-12.

DEMOTED (Group E). Moment-derived: E-01 to E-05. Causal chain not confirmed: E-06. Method insufficiency: E-07. Source errata: E-08.

38.6.9 6.9 What would change each status

A PROVED entry changes only by an error in the proof or by rejection of a stated assumption. The assumptions are in Group B and are individually attackable. **No experiment can touch Group A**, and no experiment supports it either.

A POSTULATED entry changes by being replaced with a better choice, and "better" has a meaning here: a choice that fits held-out data better, or that survives a falsification test the current one fails. B-01 would be tested by fitting a non-affine `g(u)` and comparing on held-out prediction. B-02 would be tested against a saturating or state-dependent alternative. **Neither test has been run, and neither is expensive.**

A BORROWED entry changes by replication, by failed replication, or by a study in a setting closer to the reader's. C-01, C-02, and C-07 are the three most consequential and the three most exposed: a laboratory factor extrapolated to institutions, a costless convention extrapolated to careers, and a mathematical result whose conditions are contested.

A SPECULATIVE entry changes by meeting the three bridge criteria, which for D-01 means an experiment. CCP-031 specifies what that experiment must show and what controls it needs, and **this research programme's position is that a null result reported honestly would be a service the corpus could not repay.**

A DEMOTED entry changes by instrumentation. E-01 through E-05 all change together, and they change at a stated price: **123 items per term.** E-06 changes with one randomised acoustics study. E-07 does not change — it is a permanent structural finding about what shuffles can see.

38.7 7 Results: one laboratory at three scopes, and one sentence audited

38.7.1 7.1 One laboratory, read at three scopes

A nine-person laboratory doing genomic-audio work. It is the right example for this section because it is the only kind of organisation for which all three scopes are simultaneously in play: it is an institution, it is staffed by people, and its subject matter is cells.

Every number below is arithmetic on parameter values chosen to be plausible. **None of them was measured in a real laboratory, and the honesty appendix, untested item 1 [CCP-039] applies without exception.** The purpose is to show what the computation looks like and where it stops.

Institutional scope. Period: one year. Order proxy: documented, reproducible protocols plus trained members.

Six distinct domains are admitted — molecular biology, acoustics and signal processing, information theory, statistics and psychometrics, ethics and consent, software — so `N = 6`. Cross-domain error decorrelation on shared prediction tasks gives `U = 0.42`.

```
ln N   = 1.791759 nats                  the per-interaction ceiling
N_eff  = 1 + 5(0.42) = 3.1000           nine people, three effective domains
S      = ln(3.1000) = 1.131402 nats     63.1% of the ceiling
```

Nine people and an effective count of 3.1 — not a criticism of the nine, but the audit saying that three of the six domains carry most of the independent information.

Parameters: `eta = 0.22` (protocols are written but reviews are verbal), `delta = 0.15` per year (one departure a year from a nine-person group, plus documentation drift), `C = 0.30` per year (coordination, compliance, and the consent apparatus), `Q = 2.40`.

```
eta ln N = 0.22 x 1.791759 = 0.394187
delta Q  = 0.15 x 2.40     = 0.360000
C                          = 0.300000

h = 0.394187 − 0.360000 − 0.300000 = −0.265813        UNSUSTAINABLE
```

Headroom is negative, and the no-geometry-rescues corollary is unambiguous about what that means: no arrangement of people prevents order from decreasing. The step confirms it:

```
Q_next = (1−0.15)(2.40) + 0.22(1.131402) − 0.30 = 1.988908
step   = −0.411092 Q−units in one year
```

The tempting move is to raise `u` — recruit more difference. Price it. A zero-growth step needs `S >= (delta Q + C)/eta = 0.660/0.22 = 3.000000` nats against a ceiling of `1.791759`. **The required surplus exceeds the maximum the geometry can deliver, so no recruitment fixes this.** Solving for the required `U` gives `3.817`, which is outside `[0,1]`: the equation is reporting that the request is impossible rather than expensive.

Raise `eta` instead — write the reviews down, log the decisions, make the artifact the deliverable:

```
eta = 0.22  ->  h = −0.265813   Q_inf = −0.3406   h_inf = 0.145279
eta = 0.30  ->  h = −0.122472   Q_inf = +0.2628   h_inf = 0.198107
eta = 0.35  ->  h = −0.032884   Q_inf = +0.6399   h_inf = 0.231125
eta = 0.40  ->  h = +0.056704   Q_inf = +1.0171   h_inf = 0.264143
```

At `eta = 0.40` the required `S` for a zero-growth step falls to `1.650000`, inside the ceiling, and the required `U` is `0.841396` — demanding, but a real number in range. **The lab's problem was never that its people were too similar. It was that its interaction was not being converted, and the arithmetic says so before anyone is hired or blamed.**

Personal scope. One member of that laboratory. Period: one month. Order proxy: retained working skill and standing commitments.

Five distinct framings brought to the work — bench, instrument, statistical, ethical, narrative — so `N = 5`. Decorrelation of that person's own committed written predictions gives `u = 0.28`.

```
ln N   = 1.609438
N_eff  = 1 + 4(0.28) = 2.1200              five framings, two effective ones
F_int  = ln(2.1200) = 0.751416 nats        46.7% of the ceiling
```

Parameters: `eta = 0.30`, `delta = 0.08` per month (half-life 8.313 months), `C = 0.18`, `Q = 3.10`.

```
eta ln N = 0.482831
delta Q  = 0.248000
C        = 0.180000

h = +0.054831        growth available
```

And then the instructive part:

```
Q_next = (1−0.08)(3.10) + 0.30(0.751416) − 0.18 = 2.897425
step   = −0.202575                                    DECLINING
```

Headroom is positive and the step is negative, and both are correct. Headroom says growth is *available*; it does not say growth is *happening*. The step depends on `S`, and the surplus needed for a flat month is `(delta Q + C)/eta = 0.428/0.30 = 1.426667` nats against an internal surplus of `0.751416`. **Headroom is permission, not delivery.** This person has room to grow and is not using it, because two effective framings out of five is not enough difference to fill the budget they already have. Their fix is `u`, and the lab's fix is `eta`, and those are different prescriptions arrived at from the same equation on the same afternoon.

Now the deliberate error, computed so it can be recognised:

```
NAIVE CROSS-SCOPE ARITHMETIC
personal h = +0.054831 per month  ->  x12  ->  +0.657976 per year
lab      h = −0.265813 per year
"the person has headroom that could offset the lab's deficit"      <-- FALSE
```

That subtraction is a unit artifact and nothing else. The two `Q` values are denominated in different proxies — retained personal skill against documented laboratory protocol — with no conversion factor between them, and this framework does not supply one. Multiplying a monthly headroom by twelve is legitimate arithmetic within the personal scope and meaningless across it. The sentence is refuted by A.5 and by §5.6: **there is no composition theorem, so there is no sense in which one scope's surplus can be spent at another.**

Cellular scope. The same laboratory's subject. Period: one hour.

This is where the worked example stops, and the stopping is the result.

```
N       -- no counting procedure exists                          NO VALUE
u       -- no validated state encoder exists                     NO VALUE
N_eff   -- requires N and u                                      NO VALUE
S = F   -- no established input path (CCP-031)                     NO VALUE
Q       -- proxy undecided; molecular inventory refuted (88.10)  NO VALUE
delta   -- 0.014955 per hour (median protein, t_half = 46 h)     MEASURED, BORROWED
eta     -- requires S and attributable Delta Q                   NO VALUE
C       -- no commensuration to a Q-unit that does not exist     NO VALUE
h       -- three of four terms unavailable                       NOT COMPUTABLE
```

One number out of nine, and it came from a mass-spectrometry study of mouse fibroblasts done for entirely unrelated reasons.

A reader who wanted the third column filled in should notice what filling it in would require: inventing a category count, an encoder, a proxy, and an input path. All four are inventable in an afternoon, all four would be fabrications, and this section would look considerably more impressive for containing them.

The honest output is three sentences. **The laboratory has a conversion problem and the arithmetic locates it exactly.** The member has a difference problem, located by the same arithmetic in different units on a different clock, with no traffic between the two results. **And the cells have one measured parameter and an open problem, which is the true state of the third instrument.**

That is the correspondence with its limits attached. CCP-038 records every claim the three volumes made, with its status and its sources, so that the limits can be checked one line at a time rather than taken on the assurance of this one.

38.7.2 7.2 One claim, audited end to end

Take a sentence a reader might reasonably lift from this corpus and use in a meeting:

“Redesigning the deliberation room will raise the team’s collective intelligence.”

Audit it.

Step 1 — decompose into register entries. The sentence is not one claim. It is four, chained:

```
(i)   room acoustics affect speech intelligibility          C-24, standard acoustics
(ii)  intelligibility affects who speaks and how often      partially studied
(iii) turn-taking equality correlates with c                C-01
(iv)  therefore changing RT60 changes c                     B-12, and E-06
```

Step 2 — read the statuses. Link (i) is BORROWED and solid, with a stated validity range: Sabine’s formula holds for reasonably diffuse fields and degrades in heavily absorbing rooms. Link (iii) is BORROWED and is a **correlation measured in assembled laboratory groups over hours** — the transfer assumption in C-01 is that such a factor persists in teams working for years on their own problems, and it is not established. Link (ii) has no entry of its own because it has no result of its own; it is the weakest link and it is in the middle. Link (iv) is B-12, POSTULATED — **and E-06, DEMOTED, reported as not confirmed under CC-R7.**

Step 3 — read the verdict off the chain. A chain is no stronger than its weakest link, and this chain has a POSTULATE and a DEMOTION in it. **The sentence as spoken in the meeting is not supported.** What is supported is considerably narrower and still worth saying:

> *“Reverberation control will make the quiet voice audible, which is a precondition for equal turn-taking, which is one of the three measured correlates of the collective intelligence factor in the setting where that factor was measured. Whether the chain runs end to end has not been tested, and the test is cheap.”*

Step 4 — what would settle it. CC-R7 and §6.9 both give the same answer: randomise RT60 across otherwise matched deliberation rooms, measure turn-taking variance and task performance, and report. It would not be expensive. **It has not been run, and until it is, the honest sentence is the long one.**

38.7.3 7.3 The results this paper contributes

R1 — The three-scope correspondence, badged POSTULATE — MODEL FORM ONLY. Identical structure at institutional, personal and cellular scope is a statement about the model chosen at each scope. It is not a claim of kinship between the systems, not a scaling law, and not a claim that quantities aggregate between scopes.

R2 — The instrument audit. Of 27 parameter-by-scope cells, **19 have a defined estimator, 1 has a validated value** — a borrowed cell-line median giving `delta = 0.014955` per hour — **and 0 carry a first-party measurement made for this corpus.**

R3 — Five independent breaks in the correspondence. A mind is not a small society; a cell is not a small mind; the exit ethics do not descend; aggregation is not always well defined; and the composition theorem is missing.

R4 — The affine effective count equals the Rényi-2 participation ratio exactly. `N_eff = 1 + (N−1)u = N + 1 − N/PR`, re-verified over six random correlation ensembles — for example `N = 8`, `u_bar = 0.783801`, both sides giving `6.486606017883` — with a worst discrepancy of `1.8e−15`. The identity and its original numerical verification are Curzi [2026c].

R5 — The fixed-point headroom identity. `h_inf = eta(ln N − F(u*))`, in which `delta`, `C` and `Q` all cancel; and since `F(u) <= ln N`, `h_inf >= 0` always, so **the plain linear model never ossifies by itself.** Verified across four parameter sets to ten decimals.

R6 — The status census. 156 register entries: **50 proved, 28 postulated, 58 borrowed, 12 speculative of which 4 are printed refusals, and 8 demoted**, with precedent claimed in **23 rows, 14.7 per cent**.

CR — The counter-result. No parameter at any scope has ever been measured first-party by this research programme — **0 of 27 cells** — the correspondence breaks in five independent ways with no composition theorem bridging them, and the census names its own weakest link: **fifty proved entries that say nothing about any person, and all the load passing through twenty-eight postulates.**

38.8 8 Discussion

38.8.1 8.1 What follows

The audit is the paper's usable output, and it is usable precisely because it is unflattering. A framework with nineteen defined estimators and one validated value is a specification, not an instrument. That is a perfectly respectable thing for a body of work to be, provided it says so; what is not respectable is a specification presented in the register of a measurement. The corpus's own protocols supply the nineteen procedures [Curzi, 2026i], and nobody — including its author — has run any of them and published the result.

The scopes do not degrade monotonically, and the inversion is the most instructive single fact in the audit. The tidy expectation is that the institutional scope is best supported and the cellular scope worst. What the audit finds is that the institutional scope has the best *borrowed* sociology and the worst first-party measurement; the personal scope has the best measurement *register*, a century of psychophysics establishing that subjective magnitude is measurable at all [Fechner, 1966, Stevens, 1957], and a demoted keystone formula in front of it; and **the cellular scope, the most speculative of the three, has the only measured parameter in the entire table.** Three instruments, three different failures.

A printed refusal is worth more than a passed test, and the corpus contains them. The clearest is the momentum lever: trajectories at `beta = 0, 0.5, 0.8, 0.9, 0.95` all terminate at the same point, `x = 1.9972`, in the same basin, so momentum defeats narrow barriers and is useless against wide shallow ones — and it was **refused formulary entry** on that basis rather than retained with a caveat [Curzi, 2026p]. Four of the twelve speculative register entries are refusals of the same kind.

Where the model form survives because the modeller chose the proxy, that is stated in public. The clearest case is cellular stored order. The obvious proxy is the molecular inventory, and it is refuted by the corpus's own measured parameter: at a 46-hour median protein half-life, the surviving fraction after one month is `1.942e-05`. A habit persists in a person for thirty years; the molecules implementing it turn over in days. The resolution is not a repair to the law but a correction to the proxy — what persists is the maintained pattern that keeps rebuilding the material, not the material. **That is a legitimate modelling move and it is exactly the move that makes a framework unfalsifiable if it is made silently.**

What survives the five breaks is smaller and load-bearing. The arithmetic transfers wherever the model form is adopted, because contraction, the unique fixed point, the `ln N` ceiling and `h_inf = eta(ln N - F(u*))` are consequences of the map. That is an unusual kind of portability: it costs nothing to check and it cannot be defeated by an objection to the scope, only by an objection to the map. Three ordinal diagnostics survive as well — whether `u` is near zero, whether `eta` is near zero, whether `delta Q + C` exceeds `eta ln N` — and they survive for the same reason the demoted Moment equation keeps its ordering: **comparisons survive where magnitudes do not.** A group whose members never disagree, a practice that leaves no artefact, and a system whose maintenance consumes its whole budget are all diagnosable without a calibrated instrument.

38.8.2 8.2 What does not follow

No result transfers by correspondence. The collective-intelligence factor was measured in assembled laboratory groups over hours [Woolley et al., 2010, Engel et al., 2014]; it says nothing

about a cell and nothing about the internal structure of one person's framings. **The symbol `u` appearing in three places does not make one measurement of `u` count three times.**

No quantity may be summed, differenced or divided across scopes. Section 7.1 computes such a quantity deliberately — a personal monthly headroom of `+0.054831` multiplied by twelve and set against a laboratory annual headroom of `−0.265813` — and shows it to be a unit artefact. The two stored-order values are denominated in different proxies with no conversion factor between them, and this corpus does not supply one.

A proxy chosen on the wrong axis understates the thing it is meant to measure, and the corpus has a measured instance. Curzi [2026] finds entity-based exposure to one jurisdiction at 27.3 per cent against a 61.0 per cent share of the same balance sheet denominated in that jurisdiction's currency and roughly 58 per cent of settlement volume clearing through its system — **the entity map understates dependency by more than a factor of two.** The cellular order proxy of Section 4.6 is the same failure caught before it was made.

No composition theorem exists, and its absence is the largest open problem the correspondence raises. What would close it is a stated aggregation operator together with a proof that the storage law is closed under it: that a population of systems each obeying the law, combined by the operator, yields a system obeying the law with derived parameters. Whether the law is closed under any natural aggregation is not known here. **It is a concrete, tractable-looking mathematical question, it needs no laboratory, and it is the problem the series ends on** [Curzi, 2026].

No ethics descend with the arithmetic. A member of an institution can exit and a framing can be abandoned [Hirschman, 1970]; a cell cannot leave the body. Programmed cell death is not exit in Hirschman's sense, because it is not a response by the unit to its own assessment of the organisation and the unit does not survive it. **Consent at the biological layer is consent obtained from the person whose cells they are, before anything is done, with the right to stop and with nothing concealed** [Curzi, 2026]. Pretending the ethics descend would let a formalism launder a question it is not competent to answer.

No claim of ontological identity is made or entertained. A mind is not a small society; a cell is not a small mind. The internal surplus computation is licensed as a measurement of decorrelation among one person's committed outputs and as nothing more [Curzi, 2026o], and it has the properties the surplus functional's proof gives it because that proof does not care where the vectors came from.

38.8.3 8.3 Relation to the surrounding literature

The audit's method is the reform literature's applied to a framework rather than to a finding: declare the procedure, freeze it, report what happened including the nulls [Nosek et al., 2018], treat researcher degrees of freedom as the dominant threat [Simmons et al., 2011, Gelman and Loken, 2013], expect base rates to be low [Ioannidis, 2005], and prefer risky point predictions to directional ones, which in psychology are nearly costless [Meehl, 1967, 1990]. Lakatos [1970] is the honest description of what the twenty-eight badged postulates are doing: they are the protective belt, and the register's function is to keep the belt visible.

The measurement commitments are those of measurement theory rather than of psychometrics as practised. Michell [1997, 1999] is the objection the corpus is most exposed to and Section 9.7 concedes it. Luce and Tukey [1964] and Krantz et al. [1971] give the test that was never run. Rasch [1960] and Andrich [1988] supply the constructive route that the register prices at 123 items per term, and Sijtsma [2009] is the reason reliability is not a substitute for structure.

The independence requirement that Break 1 turns on is Condorcet's (1785; named in prose, see Limitation L-11), reinforced by the model result of Hong and Page [2004] and bounded by Thompson [2014], whose critique is cited in the same breath wherever the diversity claim appears. The informational-cascade results [Bikhchandani et al., 1992, Banerjee, 1992] and the measured collapse of effective count under social influence [Lorenz et al., 2011] are what make pseudo-diversity a formal object rather than a complaint. Ugander et al. [2012] is the

closest empirical relative of an effective count in the social sciences: adoption predicted by the number of distinct connected components among a person's contacts. Markowitz [1952] states the whole point in 1952 — diversification benefit depends on covariance and not on count — and Brooks [1975] supplies the quadratic coordination cost the logarithmic surplus is set against.

The concentration instrument transfers to settings the corpus did not design it for and returns numbers rather than adjectives: Curzi [2026] reports effective counts running from 35.0 curated validators (HHI 0.028571) and 39.0 council members (0.025641) down to 6.00 organisations (0.166667) and 5.94 entities (0.168400) on the same ledgers. **That is what an instrument looks like once it has an estimator that anyone can run, and it is the contrast this paper's audit is measured against.**

The corpus's structural results about networks are cited with their own qualifications attached. Where preferential attachment and hub fragility appear [Barabási and Albert, 1999, Albert et al., 2000], the estimation procedure of Clauset et al. [2009] is binding before any real graph is called scale-free; Broido and Clauset [2019] and Stumpf and Porter [2012] bound the inference further, and Doyle et al. [2005] shows that the real Internet's robustness comes from engineered design rather than from a degree distribution.

38.9 9 Verification

The audit's operative results compound from six postulates. Each is atomic, checkable on its own, and carries its status.

38.9.1 9.1 Dimensional audit

Symbol	Dimension	Unit	Status
`Q_t`	`[Q]`	proxy units, and a **different proxy at every scope**	**proxied, never measured**; two scopes' values are not commensurable
`u`, `U_t`	`[1]`	none	**defined** dimensionless: `u = 1 - (x.y)^2` for unit vectors
`N`	`[1]`	categories	**counted** at two scopes; **no counting procedure exists** at the cellular scope
`N_eff`, `PR`	`[1]`	effective count	**derived**
`S = F(u)`	`[1]`	nats	**derived**; **not instantiated** at the cellular scope, no input path
`eta`	`[Q]` per nat per step	proxy units per nat per period	**estimator defined at two scopes**; not a number at the third
`delta`	`[1]`	per period — **the period is part of the number**	**measured once**, cellular, borrowed

Symbol	Dimension	Unit	Status
`C_t`	`[Q]` per step	proxy units per period	**estimator defined**; rests on a commensuration postulate
`h_t`	`[Q]` per step	proxy units per period	**derived**; **not computable** at the cellular scope
`M`, `D`, `P`	—	**none supplied**	**ordinal**; the quotient is demoted, and does not descend to the cell at all
period `p`	`[T]`	year / month / hour, by scope	**stipulated** per scope

Does the gauge balance? `h = eta ln N − delta Q − C` balances only if every term carries `[Q]` per step, which requires `eta` to carry `[Q]` per nat per step, since `ln N` is dimensionless. It does balance within a scope. **It does not balance across scopes**, because `[Q]` denotes a different proxy at each, and Section 7.1's deliberate cross-scope subtraction is exactly the dimensional error that produces.

Where the framework has no operational unit, this paper says so. Stored order has no natural unit at any scope. Comfort, the corpus's underlying construct, has **no operational unit at all**. Perspective and perception have **no unit and no zero**, which is why their quotient is demoted. At the cellular scope `N` has **no agreed referent**: candidate readings — distinguishable chromatin states, addressable regulatory programmes, transcription-factor combinations with distinct downstream consequence — differ from one another by orders of magnitude and none has a validated counting procedure. **A parameter that cannot be counted cannot bound anything, and `ln N` at the cellular scope is a symbol with no number behind it.**

38.9.2 9.2 The postulates

P1 — The storage recurrence is linear and first order at each scope. *Status:* **DEFINITION** — the model form adopted [Curzi, 2026a]; it carries no empirical weight on its own.

P2 — `0 <= F(u) <= ln N` for `F(u) = ln(1 + (N−1)u)` on `u` in `[0,1]`. *Status:* **CITED** — the range result [Curzi, 2026,c].

P3 — For a share or eigenvalue vector, `PR = (sum lambda)^2 / sum lambda^2`. *Status:* **DEFINITION** — the reciprocal Rényi-2 collision probability [Rényi, 1961, Simpson, 1949].

P4 — A per-period fraction converts between clocks as `1 − (1−delta)^(p_new/p_old)`. *Status:* **ARITHMETIC** — geometric decay, redoable in one line.

P5 — Median mRNA half-life is about 9 hours and median protein half-life about 46 hours in the measured cell line. *Status:* **CITED, AND BORROWED** — measured genome-wide by metabolic labelling with mass spectrometry in one immortalised mouse fibroblast line, by molecular biologists for their own purposes, as medians over thousands of species with enormous spread. **This is the only validated parameter value anywhere in the correspondence, and the primary source is named in prose because it carries no key in this corpus's external bibliography (L-11).**

P6 — Division requires a ratio scale with a true zero. *Status:* **CITED** — Stevens [1946], with the measurement-theoretic development in Krantz et al. [1971].

38.9.3 9.3 The compounding chain

```
P1 + P2        -> V1   a nondecreasing step needs eta F(u) >= delta Q + C; since
```

```
                        F(u) <= ln N, the gauge h := eta ln N - delta Q - C bounds
                        what any interaction geometry can deliver.          ARITHMETIC

P1              -> V2   the fixed point is Q_inf = (eta F(u*) - C)/delta, so at the
                        fixed point delta Q_inf = eta F(u*) - C.            ARITHMETIC

V1 + V2         -> V3   h_inf = eta ln N - (eta F(u*) - C) - C
                             = eta (ln N - F(u*)).
                        delta, C and Q cancel identically.                  ARITHMETIC

V3 + P2         -> V4   F(u*) <= ln N, therefore h_inf >= 0 ALWAYS: the plain
                        linear model never ossifies by itself.              ARITHMETIC

P3              -> V5   1 + (N-1)u_bar = N + 1 - N/PR exactly, where u_bar is the
                        mean pairwise non-parallelism of the ensemble.      ARITHMETIC

P4 + P5         -> V6   t_half = 46 h gives delta = 1 - 2^(-1/46) = 0.014955 per
                        hour, = 1 - 2^(-24/46) = 0.303468 per day, and a surviving
                        fraction over 720 h of 2^(-720/46) = 1.942e-05.     ARITHMETIC

V6              -> V7   a molecular-inventory proxy for cellular Q predicts a
                        per-month retention loss of 0.99998, i.e. that nothing
                        cellular persists; the proxy is refuted, not the law. ARITHMETIC

P6              -> V8   D/P is not reportable as a magnitude, because neither D nor P
                        is instrumented on a ratio scale with a true zero.  CITED

P4              -> V9   0.15 per year is 1 - 0.85^(1/12) = 0.013452 per month;
                        0.08 per month is 1 - 0.92^12 = 0.632334 per year;
                        the ratio of the two is 47.0.                       ARITHMETIC
```

Every line is checkable from its own stated inputs. V1 through V9 are consequences of one definition of model form (P1), two cited results (P2, P6), two definitions (P3, P4) and one borrowed measurement (P5). **Only V6 and V7 consume an empirical input, and that input was measured by other people in one cell line for reasons unrelated to this corpus.**

What the chain does not establish. It does not establish that any institution, mind or cell obeys P1 at any scope. That is the correspondence postulate, it is not verified here, it has never been tested at two scopes simultaneously, and the reflexive parameter it would need to recover is not identified from observation at all (§3.7). **A chain whose steps are ARITHMETIC, DEFINITION and CITED delivers the mathematics with certainty and delivers nothing whatever about any of the three systems.**

38.9.4 9.4 The verification instance, worked by hand

Check 1 — the participation-ratio identity. Take the ensemble reported in Section 4.4 at `N = 8` with mean pairwise non-parallelism `u_bar = 0.783801`. The affine count is `1 + 7 x 0.783801 = 6.486607` to six decimals. The participation-ratio form returns `8 + 1 - 8/PR = 6.486606017883` on the same correlation matrix. The two agree to the twelve decimals printed and, across six ensembles, to a worst discrepancy of `1.8e-15`, which is floating-point noise. **The identity is a fact about the aggregator, not about people: it would hold for correlated thermometers.**

Check 2 — the fixed-point headroom identity. At `eta = 0.35`, `N = 12`, `u* = 0.30`, `delta = 0.05`, `C = 0.40`: `ln 12 = 2.484907`, `F(0.30) = ln(1 + 11 x 0.30) = ln 4.3 = 1.458615`, so `h_inf = 0.35 x (2.484907 - 1.458615) = 0.359202`. Computing it the long way, `Q_inf = (0.35 x 1.458615 - 0.40)/0.05 = 2.210305`, and `h_inf = 0.35 x 2.484907 - 0.05 x 2.210305 - 0.40 = 0.359202`. **The**

same number to ten decimals, and the second computation used `delta`, `C` and `Q` which the first did not need.

Check 3 — the clock conversion. `1 - 0.85^(1/12) = 0.013452`; `1 - 0.92^12 = 0.632334`; and `0.632334 / 0.013452 = 47.0`. **The parameter that looks four times larger is forty-seven times larger once the clocks agree.**

Check 4 — the molecular proxy, refuted in one line. `2^(-720/46) = 1.942e-05`. Everything else in Section 4.6 follows from that number.

Check 5 — the Rasch price, recomputed. With every item targeted at the person, item information is 0.25, so `SE(theta) = 2/sqrt(k)`; the distortion is a difference of two independent measures, so `SE(Delta) = sqrt(2) x 2/sqrt(k)`; and the 95 per cent detection floor is `1.96 x SE(Delta)`. At `k = 122` the floor is `0.501905`, above a half-logit; at `k = 123` it is `0.499860`, below. **One hundred and twenty-two items is not enough and one hundred and twenty-three is, and the crossing is that sharp because the arithmetic is that simple.**

Check 6 — the census. The register in Section 6 contains 156 rows: 50 in Group A, 28 in Group B, 58 in Group C, 12 in Group D and 8 in Group E. **A reader can recount them.** The precedent column reads *none claimed* in 133 rows, leaving 23, which is 14.74 per cent.

38.9.5 9.5 Limits, asymptotics, and the model's valid domain

- **`u -> 0` at any scope.** `N_eff -> 1`, `F -> 0`, `Q_inf -> -C/delta < 0`. Ten parallel windows are one window. **This is the regime a converged group is in, and it is not a degenerate case.**
- **`u -> 1`.** `N_eff -> N`, `F -> ln N`, `h_inf -> 0`. Maximal surplus and zero remaining headroom at the fixed point.
- **`N -> 1`.** `ln N = 0`, so `h = -delta Q - C <= 0` always. **A single category can never build under this gauge**, at any scope; the framework is silent on whether that is a statement about systems of one or an artefact of using `ln N` as the ceiling.
- **`N` doubling.** Each doubling adds exactly `ln 2 = 0.693147` nats wherever it starts, so **a headroom problem cannot be escaped by adding categories** — a property of the logarithm and not of any system, and therefore true at all three scopes for the same reason.
- **`delta -> 0`.** The fixed point does not exist at all rather than being infinite: the map is an isometry and the geometric series does not converge [Curzi, 2026a]. **Outside the valid domain.**
- **`delta -> 1`.** Everything is lost each period; `Q_inf -> eta F(u*) - C`, and the system has no memory. At the cellular scope this is nearly the case for mRNA: `delta = 0.074125` per hour is a half-life of 9 hours, so a daily clock puts the cellular mRNA pool close to this boundary.
- **`eta -> 0`.** `h -> -delta Q - C < 0` at every scope. A system with no conversion cannot build at any `u` and any `N`, which is the storage law's correspondent of the forbidding clause `G = 0 => E = 0` in the channel-gain form of the belief postulate [Curzi, 2026ab].
- **Between scopes.** No limit exists, because no aggregation operator is defined. **The correspondence has no asymptotics in the scope variable, and that absence is Break 5.**

38.9.6 9.6 Model comparison, and the standard this paper does not meet

The paper reports no fit of its own. It nonetheless inherits the standard, because the register records a confirmatory fit and the correspondence would be used to justify further ones.

Requirement	Status	What was actually done
Out-of-sample performance reported	**MET**	held-out RMSE against four baselines [Curzi, 2026i]

Requirement	Status	What was actually done
Matched null with equal free parameters	**MET**	a matched AR(1); the fit beat it by 2.8 per cent against a declared 10 per cent margin and **failed its own preregistered rule**
Information criterion (AIC or BIC)	**NOT MET**	**no AIC or BIC has been computed anywhere in this corpus**, so no comparison pays for its parameter count
Number of free parameters stated explicitly	**PARTIAL**	the four storage-law parameters are named; they are not tabulated against each baseline's count
Effect sizes with intervals	**NOT MET**	RMSE advantages are reported as point estimates without intervals
Shuffle control's false-positive rate measured	**MET, AND ADVERSE**	0.108 against a nominal 0.05 for i.i.d. permutation, 0.120 for circular shift [Curzi, 2026u] — **anti-conservative**
Non-affine alternative fitted and compared	**NOT MET**	register entry B-01 states the test and records that it has not been run; **it is not expensive**
Saturating or state-dependent storage law compared	**NOT MET**	register entry B-02, same status

Four NOT MET rows are printed here rather than omitted. The last two are the most damaging, because they are the two tests that would decide whether the corpus's two founding design choices are the right ones, and neither requires new data.

38.9.7 9.7 Scale type, and the objection the corpus concedes

Quantity	Scale type	What licenses it	Operations performed here
`delta` from half-life	ratio	time has a true zero and a unit	conversion between clocks, exponentiation
`N`, `N_eff`, `PR`	ratio (counts)	counting	all — where a counting procedure exists
`u`	ordinal at best, treated as ratio	**nothing tested**	used inside a logarithm; **not licensed**
`Q`	interval at best in a declared proxy	proxy definition only	differences and rates within a scope; **never across scopes**
`D`, `P`, `M`	ordinal	no zero, no unit	ordering only; the quotient **demoted**

Quantity	Scale type	What licenses it	Operations performed here
register status	nominal with a partial order	definition	counts by category; **no arithmetic on the labels**

Michell's critique is the strongest single objection to this corpus and it is conceded here rather than answered. Michell [1997, 1999] argues that psychology assumes quantitative structure without testing for it. The corpus's founding move — treating comfort as a magnitude — is that assumption. The test that would settle it is conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971], whose double-cancellation conditions determine whether an attribute admits additive quantitative structure at all. **Those conditions have not been tested for comfort, for `u`, for perspective or for perception, by this work or by anyone.** What the corpus does instead is partial and is stated as partial: it denotes the one quantity whose operation most obviously requires a ratio scale, and it prices the Rasch route at 123 items per term [Rasch, 1960, Andrich, 1988] — an instrument with a defensible zero, which is not a conjoint test. Reliability would not substitute [Sijtsma, 2009], and construct validation [Cronbach and Meehl, 1955, Campbell and Fiske, 1959] would show that an instrument measures something without showing that the something is quantitative.

38.9.8 9.8 Invariance, and the precision of the numbers

What is invariant. The identity `N_eff = N + 1 − N/PR` is invariant under permutation and relabelling of members and holds at any scope where a correlation matrix can be formed. The identity `h_inf = eta(ln N − F(u*))` is invariant under any change in `delta`, `C` or `Q`, which is the content of the cancellation. The gauge's verdict is invariant under a common rescaling of `Q`, `C` and `eta` and **is not invariant under rescaling `Q` alone**, which is what changing the order proxy mid-measurement does.

What is not invariant, and is the point of the paper. Nothing is invariant across scopes. The correspondence claims form identity and explicitly disclaims parameter transfer, scaling, and composition. **Scale-freedom is claimed at three scopes and has been tested at zero**, since testing it requires simultaneous measurement at two scopes of the same system and no such measurement exists.

Precision. Numbers in this paper fall into three classes and are treated differently. *Exact consequences of definitions* — `ln 12 = 2.484907`, `h_inf = 0.359202`, `N_eff = 6.486606017883` — are quoted to whatever precision is useful and are exact. *Borrowed measurements* — the 9-hour and 46-hour median half-lives — are **medians over thousands of species with enormous spread, from one immortalised cell line by one method**, and the derived `delta = 0.014955` per hour is quoted to six figures only because the conversion is exact given the median, not because the median is known to six figures. **The underlying quantity is a median with a wide distribution and no interval is reported for it here, because none is reported in the source this corpus borrows from as summarised.** *Illustrative model outputs* — the laboratory's `h = −0.265813` and the member's `h = +0.054831` — are exact given stipulated parameters and carry no empirical content whatever. **Where a reader sees six decimals in this paper they should read "consequence of a stipulated or borrowed input", never "measured to that precision".**

38.10 10 Limitations

L-1 Zero of twenty-seven cells carries a first-party measurement. This is the paper's own principal result and it is also its principal limitation. Every empirical quantity in this corpus is borrowed or stipulated. **No claim in any of the seventy papers rests on fieldwork conducted**

for the corpus, and any reader weighing the framework's applied recommendations should weigh them on that basis.

L-2 The reflexive parameter is not identified. The framework's central social claim is a linear-in-means reflexive specification whose endogenous effect cannot be recovered from observational data [Manski, 1993, Shalizi and Thomas, 2011]. **No remedy is adopted anywhere in the corpus.** Nineteen estimators exist; for the parameters carrying social content they estimate conditional summaries of a fitted series, not couplings. Every applied claim is a modelling position rather than an identified empirical result.

L-3 Scale-freedom has never been tested at two scopes simultaneously. The correspondence is a postulate about model form and is falsifiable only at each scope separately. **Three instances is not an invariance result, and this paper does not present it as one.**

L-4 The conjoint conditions have not been tested for any construct in this corpus [Michell, 1997, Luce and Tukey, 1964]. The ratio-scale treatment of `u` in particular is unlicensed and is used inside a logarithm regardless.

L-5 No AIC or BIC has been computed anywhere in this corpus, and neither of the two tests that would decide the corpus's founding design choices — a non-affine effective count, and a saturating or state-dependent storage law — has been run, though neither is expensive (§9.6).

L-6 The one validated parameter is borrowed, is a median with enormous spread, and comes from one immortalised cell line by one method. Using it as the cellular `delta` assumes that a genome-wide median in a mouse fibroblast line transfers to whatever cell a reader has in mind. **That transfer is an assumption and is named as one.**

L-7 At the cellular scope, four of the nine parameters would have to be invented rather than measured to fill the column — a category count, a state encoder, an order proxy and an input path. All four are inventable in an afternoon; all four would be fabrications; and the paper would look considerably more impressive for containing them.

L-8 There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal [Curzi, 2026y]. Everything upstream is real and separately supported; the specific step is established by none of it. **The third instrument's input term therefore has no measured input path, and the third instrument is a specification rather than a measurement.**

L-9 The register's own accuracy is a claim this paper makes about itself. The 156 rows and their statuses were assigned by the same author whose claims they classify. **A register is not an external audit**, and the only defence offered is that every row carries the six fields a hostile reader needs to check it, and that the demotions are printed in the same typeface as everything else.

L-10 A numerical inconsistency in the source is recorded here rather than silently corrected. The census in Section 6.7 reports that 23 of 156 rows carry precedent, 14.7 per cent, and a direct recount of the register confirms 23. **The failure-mode table in Appendix B states "twelve of one hundred fifty-six rows have precedent". That figure is wrong; the correct count is 23, and the census figure is the one to quote.** The error is in the source and is disclosed rather than quietly repaired.

L-11 Several works are named in prose without bibliographic keys, because they are not in this corpus's external bibliography. They are: Schwanhäusser et al. (2011) on genome-wide mRNA and protein half-lives, which supplies the single validated parameter in the audit; Condorcet's 1785 *Essai*; Arrow's 1951 impossibility theorem, invoked in Break 4; Minsky's *Society of Mind* (1986), cited there as precedent for a picture rather than as support for it; Robinson's ecological fallacy and Simpson's paradox, named in Break 4 (note that `simpson1949measurement` in the bibliography is the diversity index and is **not** the paradox, and the two must not be conflated); Yates's *The Art of Memory* (1966); Durkheim on collective effervescence and Rappaport on ritual; and the cell-biology chain of Tajik et al. (2016), Ibsen et al. (2015), Church, Gao and Kosuri (2012), Zetsche et al. (2015), Qi et al. (2013) and Gilbert et al. (2014). **Each is a real published result quoted for what its source measured; the reader is owed the note that they are named rather than referenced.** The corpus's own treatment of the cell-biology chain is Curzi [2026x] and Curzi [2026y].

L-12 The mapping from the source manuscript's internal papers and chapter numbers to corpus paper identifiers is an editorial inference (§3.9), not a statement by the source. Where the register's WHERE field names a corpus paper, that name is this paper's judgement about which converted paper carries the material.

L-13 The five breaks are arguments, not measurements. Breaks 2, 3 and 4 turn on category distinctions — representation, exit, aggregability — that no experiment would settle, and a reader who rejects the category distinction rejects the break. Break 1 is quantitative in principle and untested in fact: **nobody has measured whether a `u` of 0.6 between two of one person's framings means what a `u` of 0.6 between two colleagues means**, and the mitigation the corpus offers — measuring on committed outputs rather than introspective report — removes one source of inflation and not common-mode coupling.

L-14 The composition theorem is absent and its absence is not a gap in the writing but a gap in the mathematics. No aggregation operator is defined between scopes, no conservation law links them, and no argument is offered that the three accumulators are coupled at all rather than merely similar in form. **Whether the storage law is closed under any natural aggregation is not known here.**

38.11 11 Conclusion

The same storage law is written at institutional, personal and cellular scope, and that correspondence is badged POSTULATE — MODEL FORM ONLY. It states that a minimal open-system accumulator written for each of the three uses the same five symbols. It does not state that the three systems are alike, that their parameters are related by any law, or that quantities at one scope aggregate into quantities at another. **What follows from it is exactly the arithmetic of the map** — contraction, a unique fixed point, the ceiling `0 <= F(u) <= ln N`, the headroom inequality, and `h_inf = eta(ln N − F(u*))` — **and nothing about any of the three systems.**

The audit of the instrument is the paper's principal result, and it should be quoted in this form. Across nine parameters and three scopes there are twenty-seven parameter-by-scope cells. Nineteen of the twenty-seven have a defined estimator. Exactly one has a validated value: a borrowed median protein half-life of 46 hours, giving `delta = 0.014955` per hour, measured genome-wide in one immortalised cell line by molecular biologists for their own purposes. And zero of the twenty-seven carry a measurement made first-party for this corpus. Nothing in the seventy papers rests on fieldwork conducted for them. That sentence is the one this paper exists to make citable, and it is not softened anywhere in it.

The correspondence then fails in five independent ways. A mind is not a small society: two framings in one head share a memory, a mood, a set of priors, a stake and a substrate, which are five channels of common-mode coupling with no institutional analogue, and the independence every aggregation result requires is unavailable. A cell is not a small mind: the moment equation needs a representation, a cell has states rather than representations, and so it has no frame and no perspective ratio — permanently, rather than pending instrumentation. The exit ethics do not descend: a member can leave and a cell cannot, so consent attaches to the person and never to the scope. Aggregation is not always well defined, in principle and not merely in practice. And **no composition theorem exists anywhere in this corpus**, which is a gap in the mathematics rather than in the writing, and the largest open problem the correspondence raises.

What survives is smaller and load-bearing. The arithmetic transfers wherever the model form is adopted, because it is a consequence of the map rather than a fact about a system. And three ordinal diagnostics read at every scope without a calibrated instrument: **whether `u` is near zero, whether `eta` is near zero, and whether `delta Q + C` exceeds `eta ln N`.** Each is a comparison rather than a magnitude, which is precisely why the demotion of the Moment equation leaves them standing.

Two identities are re-verified and both hold at every scope where the map is adopted. The affine effective count equals the Rényi-2 participation ratio exactly, `N_eff = 1 + (N−1)u = N + 1`

– `N/PR`, checked over six random correlation ensembles with a worst discrepancy of `1.8e-15` — for instance `N = 8`, `u_bar = 0.783801`, both sides `6.486606017883`. And `h_inf = eta(ln N - F(u*))`, in which retention, cost and stored order all cancel; since `F(u) <= ln N` always, **`h_inf >= 0` always, so the plain linear model never ossifies by itself** and ossification requires something the model does not contain.

The register publishes 156 entries: 50 proved, 28 postulated, 58 borrowed, 12 speculative of which 4 are printed refusals, and 8 demoted. Precedent is claimed in **23 rows, 14.7 per cent**; the other 85.3 per cent read *none claimed*, which is the honest yield of a two-column method when the rule against filling the second column is enforced. **The census names its own weakest link: fifty proved entries, not one of which says anything about any person, and the entire load passing through twenty-eight postulates.** To break this corpus, break a postulate.

38.11.1 Declarations

Data and code availability

No external dataset was collected for this paper, and no human participants were surveyed, observed or recruited. This is the paper that says so for the whole corpus: of 27 parameter-by-scope cells, **zero carry a first-party measurement**, and every empirical quantity appearing anywhere in these seventy papers is either borrowed from a published study conducted by other researchers for their own purposes, or computed from parameters stipulated in a worked example. The single validated parameter used here — median mRNA and protein half-lives in a mouse fibroblast line — belongs to its published source and is cited to it in prose (L-11). The laboratory worked at three scopes in Section 7 is a **constructed illustrative example with declared parameters**, not a report of any real organisation, person or cell. The figures are generated by the scripts named in each caption, which contain no data beyond the numbers printed in this paper. The register in Section 6 is derived by reading the corpus's own papers and can be recounted from the tables printed here.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest. **The register in Section 6 classifies the author's own claims, which is a structural conflict rather than a financial one, and it is disclosed as Limitation L-9.**

Author contributions

Single author. Michael Laurence Curzi conceived the work, derived the results, compiled and counted the register, produced the worked example and figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. No cells, tissues or biological materials were handled and no experiment of any kind was performed. **No clinical, medical or therapeutic claim is made anywhere in this paper.** Section 5.4 states the corpus's position that consent attaches to the person and never to the scope: no protocol at the biological layer may be described as consented by any party other than the person whose

cells are involved, before anything is done, with the right to stop and with nothing concealed. That is a design constraint on prospective work and is not a report of any study performed.

Preregistration

This paper reports no preregistered test of its own; its results are an audit, two identities and a census. The preregistered protocol for testing the storage law against a real system, and the confirmatory application that **failed** its second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin — are reported in Curzi [2026i]. The measured false-positive rate of the shuffled-input control, **0.108 against a nominal 0.05**, is reported in Curzi [2026u]. Register entries B-01 and B-02 name two further tests of the corpus's founding design choices and record that **neither has been run.**

38.12 Appendix A The correspondence specification and its failure modes

A.1 Every use of the three-scope correspondence shall state, in the same sentence, that it is a claim about model form only, and shall cite the corpus's model-form convention and the honesty appendix, postulate 5 [CCP-039].

A.2 No parameter value measured at one scope shall be used as evidence for a parameter value at another scope. Where a value is borrowed for illustration, it shall be labelled as an illustration and the borrowing shall be named.

A.3 Every `delta` shall be quoted with its period. A `delta` without a period is not a number and shall be rejected at review.

A.4 Every cross-scope comparison shall convert all rates to one clock first, using `delta_new = 1 - (1 - delta_old)^(p_new/p_old)`, and shall show the conversion.

A.5 No quantity shall be summed, differenced, or divided across scopes. §7.1 computes such a quantity deliberately and shows it to be a unit artifact.

A.6 The `M` and `D/P` rows shall carry the CC-R1 demotion at every appearance, including in summary tables, together with the the Rasch-route theorem [CCP-026] route and the 123-items-per-term cost. A summary table that omits the demotion is non-conforming.

A.7 Any cellular-scope computation shall state which of the nine parameters have estimators before any number is produced. At the time of writing the answer is one, and a document producing a cellular `h` has erred upstream of its arithmetic. Any cellular-scope `Q` shall state its proxy explicitly and address the molecular turnover result of §4.6; a `Q` proxied on molecular inventory shall be rejected.

A.8 Internal `u` shall be measured on committed outputs — independent written predictions produced before comparison — never on introspective report, and shall carry the common-mode caveat of §5.2.

A.9 The consent constraint shall attach to the person, never to the scope. No protocol at the biological layer shall be described as consented by any party other than the person whose cells are involved, and disclosure shall be complete: **nothing masked, nothing below threshold, nothing concealed.**

A.10 Any aggregation between scopes shall name its aggregation operator. Where none is named, the document shall state that no composition theorem exists in this corpus.

A.11 Every table presenting the correspondence shall be accompanied by its measurement-status table (§4.1, second table). **The two tables travel together or neither is published.**

38.12.1 A.12 Failure modes of the correspondence

Fault	Signature	Mechanism	Fix
SCOPE LAUNDERING	A result cited at one scope to support a design at another, with the citation intact and the scope silently changed	The symbols match, so the transfer feels like a lookup rather than an inference	A.2. Re-cite with the original setting stated. Woolley measured groups; it is not a fact about cells or about one mind
THE CLOCK ERROR	Two `delta` values compared directly; a monthly rate read against an annual one	`delta` looks dimensionless because it is a fraction	A.3, A.4. `0.08/month` is `0.632/year`, not `0.08/year`
THE HOMUNCULAR SLIDE	Internal framings described as deliberating, negotiating, or voting	The sociological vocabulary is richer and more available than the psychological one	A.8. Speak of decorrelation among committed outputs. If the sentence needs an agent inside the head to be true, it is not this framework's sentence
MOLECULAR `Q`	A cellular stored-order proxy denominated in protein or transcript abundance	The molecules are the visible thing	A.8. Median protein survives one month at `1.9e−05`. The proxy predicts that nothing persists, which is false
THE EMPTY DIAL	A cellular headroom, surplus, or `eta` reported with a numeric value	The equation is writable, so it looks computable	A.7. Three of `h`'s four terms have no estimator. Report the absence
CONSENT BY SCOPE	A biological-layer protocol described as involving no consenting party, or as consented at the cellular level	The arithmetic descended, so the ethics feel as if they did	A.9. Consent attaches to the person. There is no one at the cellular scope to ask
THE COMPOSITION ASSUMPTION	Language implying that cellular order builds personal order builds institutional order	The three volumes are ordered, so they look like a chain	A.10, §5.6. No composition theorem exists. Say so, or derive one
DEMOTION DROP	An `M` or `D/P` magnitude reported in a synthesis section without the CC-R1 badge	Summary tables lose qualifiers first	A.6. The badge is part of the row, not a footnote to it

Fault	Signature	Mechanism	Fix
TABLE HYPNOSIS	The master correspondence table reproduced alone, without the measurement-status table	The first table is beautiful and the second is not	A.11. They are one artifact

---

38.13 Appendix B The register specification and its failure modes

B.1 Every substantive claim in this corpus shall carry an entry in this register, and every entry shall carry all six fields. **An entry missing its "what it measured" field is non-conforming and shall be removed rather than repaired in place.**

B.2 The (-1) column shall read **none claimed** wherever no precedent exists. **Populating it to improve symmetry is the single most damaging thing that can be done to this section.**

B.3 No entry shall be promoted by agreement between the columns. Promotion requires a proof, a measurement, or a passed gate, and the register shall record which.

B.4 Every borrowed entry shall state the setting and the transfer assumption. Where the transfer assumption is large, it shall be stated in bold.

B.5 Every demotion shall appear in the register **and** in this section that made the original claim, in the same pass. A demotion recorded only in the appendix is concealment by filing.

B.6 Downstream entries shall be demoted with their upstream entry. CC-R1 demotes E-01 through E-05 together; no M-derived magnitude survives anywhere in the corpus.

B.7 Every speculative entry shall be marked speculative **in the sentence that states it**, not in a note attached to it.

B.8 Rejections shall be printed. **A register with no rejections has not been run**, and the count of rejections is a diagnostic of the register's health.

B.9 The status census (§6.7) shall be recomputed whenever entries are added, and the ratios in it shall be reported even when they are unflattering.

B.10 Where a number appears in an entry, it shall have been computed or cited, and the entry shall say which. **No entry shall contain a figure whose provenance is the author's impression.**

B.11 This register is itself a stored order with its own δ. It shall be re-derived from the chapters annually, and the re-derivation shall be done by reading the corpus papers, not by editing the register.

B.12 Any reader may cite an entry ID against this corpus. **An entry that cannot be defended on its own six fields shall be withdrawn, not argued.**

38.13.1 B.13 Failure modes of the register

Fault	Signature	Mechanism	Fix
PRECEDENT LAUNDERING	A (-1) cell populated with a tradition that "says something similar", and the entry's status raised accordingly	The two-column format is symmetric and symmetry invites filling	B.2, B.3. The columns carry different things. **The source states "twelve of one hundred fifty-six rows have precedent"; a direct recount of the register returns 23 of 156, which is the figure the census in §6.7 reports and the one to quote. The discrepancy is in the source and is recorded rather than silently repaired (L-10).**
STATUS INFLATION	A postulate cited in a later chapter as though it were a theorem, usually via a chain of three cross-references	Each hop drops a qualifier; three hops drop them all	B.1. Cite the ID, not the sentence. The ID carries the status
GATE LAUNDERING	FVP tables present and every verdict PASS	Sixty-three gates and no failures is not rigour; it is a copied table	Check gate 2 and gate 7 first, since those are the two that genuinely bind
THE ORPHAN ENTRY	A register row with no chapter in its WHERE field	The claim was cut from the corpus and the row survived	Remove the row. A register is an index of the corpus, not of the author's beliefs
THE FROZEN REGISTER	Identifiers in WHERE fields that no longer match the corpus	The register has its own δ and nobody maintains it	B.11. Re-derive from the corpus papers annually, by reading
DEMOTION BY FILING	The demotion appears only in the appendix or only here, while the original chapter still asserts the magnitude	Correcting the source chapter is more work than adding a row	B.5. Both, in the same pass, or neither
CITATION DRIFT	"What it measured" fields that grow more supportive with each edition	Summaries lose caveats first and gain confidence for free	B.4. Re-read the source, not the previous summary
THE IMPRESSIVE NUMBER	A figure in an entry with no computation and no citation	Numbers are persuasive and cheap to invent	B.10. Compute it or cite it or delete it

Fault	Signature	Mechanism	Fix
REGISTER AS ARMOUR	The register cited as evidence that the corpus is rigorous, rather than used to check a claim	An audit trail can become a credential	The register's only function is to be used against the corpus. §3.6 is the instruction

---

38.14 Appendix C What this research programme undertakes

38.14.1 C.1 What this research programme undertakes

The register exists to be used against the corpus, and that is only meaningful if there is something this research programme is committed to doing when it is.

A claim shown to be wrong is withdrawn, not defended. B.12 is not rhetoric; entries that cannot be defended on their six fields are removed from the register and from the corpus papers that used them, in the same pass, the way CC-R1 removed the magnitude of the author's own keystone formula from five chapters at once.

A null result reported honestly is a contribution and will be recorded as one. the honesty appendix [CCP-039] already states that a reader who runs one of the named tests and gets a null result has done this book a service that no amount of agreement could. The register is where that service would be recorded: as a status change, with the reader's name in the (+1) column and the setting stated.

And the demotions stay in. E-01 through E-08 will not quietly disappear from a later edition because the instrumentation was built. **If the 123-item instrument is built and the distortion becomes measurable, the entry does not vanish — it changes status, and the demotion remains in its history, because the fact that the protocol caught it is worth more than the embarrassment of having needed it caught.**

That is the audit trail. CCP-039 closes the work, and it is entitled to claim exactly what is in Groups A and B and nothing that is only in Groups D or E.

Chapter 39

Four Structural Results and an Honesty Register: The Closing Synthesis and Reference Apparatus of Engineering the Surplus

Abstract. **This closing instalment is chiefly synthesis and reference apparatus, and is presented as such; it advances no testable hypothesis.** Most of its length recapitulates results established in the papers it draws on. Its non-recapitulative content is four exact algebraic results and a set of computed bookkeeping items. First, the surplus law is shown to be exactly a Fechner law in threshold-corrected form, $F(u) = \ln((u + u_0)/u_0)$ with $u_0 = 1/(N - 1)$, verified for $N = 2$ through 100 at seven values of u to a worst residual of 8.9×10^{-16}; at the threshold the effective count is exactly 2 and the surplus is exactly $\ln 2$, one bit, for every N. Second, folding the essence lattice into one octave costs exactly $\ln 2$ nats for every even N, computed at $N = 4$ through 64, so that within that lattice calling two things consonant and calling them redundant are the same operation. Third, the fixed-point headroom identity $h_\infty = \eta(\ln N - F(u^*))$ holds with retention, cost and stored order all cancelling, so $h_\infty \geq 0$ always and the plain linear model never ossifies by itself. Fourth, the shuffled-input control is shown to be blind to repackaging: a filtered repackaging passed all four baselines including the shuffle at $p = 0.0005$ and was caught only by the structural ceiling, its claimed jump being 38.5 times $\eta \ln N$. The appendices add two computed admissions: testimonial evidence under success-publishing selection carries a likelihood ratio of exactly 1.000000, and the research programme publishes its own survival curve, giving a 35.8 per cent probability of completion and a 55.3 per cent chance of dying at or before the cross-subject stage. The appendices are retained because other papers rely on them, not as a research contribution.

Keywords: Fechner law; octave folding; headroom identity; shuffled-input control; survivorship bias; research programme costing

Series: B. Engineering the Surplus **Paper:** CCP-039

39.1 1 Introduction

This paper is the closing instalment of Series B and is chiefly synthesis and reference apparatus. It is presented as such, and it should be read as such. Most of its length recapitulates

results established in the papers it draws on, and those results are cited to their sources rather than restated as findings. What is not recapitulation is four structural results and a set of computed honesty items, and those are what a reader who has read the rest of the series comes here for.

The four structural results are exact and each connects two parts of the framework that were developed separately.

(a) The surplus law is a Fechner law in threshold-corrected form. Factoring `(N−1)` out of the logarithm gives `F(u) = ln((u + u_0)/u_0)` with `u_0 = 1/(N−1)`, verified for `N = 2` through `N = 100` at seven values of `u` to a worst residual of `8.9e−16`. At the threshold the effective count is exactly 2 and the surplus is exactly `ln 2` — one bit, two octaves — **for every `N`**. The perceptual register itself, and the postulate method that operationalises it, are Curzi [2026n]; the same identity is stated as a result in Curzi [2026].

(b) Folding the essence lattice into one octave costs exactly `ln 2` nats for every even `N`, computed at `N = 4, 6, 8, 10, 12, 16, 20, 24, 32, 64` and returning `0.6931471806` in every case. Consonance and redundancy are the same operation inside that lattice, and the operation has a price. The pitch-class structure it rests on is the odd-part class rule of Curzi [2026].

(c) The fixed-point headroom identity `h_inf = eta(ln N − F(u*))`, in which retention, cost and stored order all cancel, verified against direct computation across four parameter sets to below `1e−16`. Since `F(u) <= ln N` always, `h_inf >= 0` always, so **the plain linear model never ossifies by itself** — which is inconvenient, because ossification is the most reliably observed behaviour of successful institutions, and the identity therefore locates the mechanism outside the model rather than inside it. The identity is re-derived and re-verified across scopes in Curzi [2026].

(d) The shuffled-input control is blind to repackaging. A filtered repackaging of the input passed all four empirical baselines including the shuffle at `p = 0.0005` and was caught only by the structural ceiling, its claimed single-step jump being 38.5 times `eta ln N`. The falsification suite measured against itself — including the shuffle's own anti-conservative false-positive rate of 0.108 against a nominal 0.05 — is Curzi [2026u].

The appendices add two computed admissions that a closing chapter would ordinarily omit. **Testimonial evidence under a selection process that publishes successes has a likelihood ratio of exactly 1.000000** — it does not move the posterior at all — against 2.0000 for a direction-only prediction and 11.3577 for a magnitude call with a tolerance, a factor of about 5.7. And **the research programme that would settle whether the framework's formulas are science publishes its own survival curve: 21 preregistered kill conditions of which 4 are terminal, 8,554.8 researcher hours, 570 participants, 71 months, and a 35.8 per cent probability of completion, with a 55.3 per cent chance of dying at or before the cross-subject recovery stage.**

What this paper adds beyond those items is organisation, not knowledge. Appendix A is a reference card for the two laws and their derived results; Appendix B is a glossary; Appendix C carries the measurement protocols for the six parameters and the three controls; Appendix D is a failure-mode index arranged by symptom; Appendix E is the honesty register, which is the appendix to read if only one is read; and Appendix F maps the framework onto the founding material it descends from. **They are here because other papers in the series rely on them, not because they constitute a research contribution.**

Two framing statements govern everything below and are not softened anywhere in the paper. First, **every claim in this corpus is proved, badged, or named as open**, and the register that sorts 156 of them into those bins, with the audit that found 19 defined estimators, 1 validated value and **0 first-party measurements** across 27 parameter-by-scope cells, is Curzi [2026]. Second, **nothing here adds a mechanism to the world.** No symbol acts on matter, no claim in the series has clinical or therapeutic content, and the formalism does not touch consciousness or thermodynamics; it borrows four theorems that belonged to other people first, and it says which.

39.2 2 Hypothesis

This paper advances no testable hypothesis. It is closing material and a reference collection, and its non-recapitulative content consists of four exact algebraic results and a set of computed bookkeeping items, none of which could be falsified by an observation.

That sentence is the whole of this section's substantive content, and manufacturing a hypothesis to fill the slot would be worse than stating it. What the paper has instead is an organising question, and it is worth stating plainly:

> **What is a formalism entitled to claim when it has proved a good deal, measured nothing, and demoted some of its own results?**

The answer defended across the series and collected here is: the arithmetic of the map, the ordinal comparisons that survive a demotion, and a register that says which is which — and nothing that requires a calibrated instrument. Everything in the appendices is apparatus for making that answer checkable.

The four structural results are labelled by kind rather than presented as findings about the world.

- Results (a), (b) and (c) are **exact algebraic identities**. `F(u) = ln((u+u_0)/u_0)` with `u_0 = 1/(N−1)`; `ln N − ln(N/2) = ln 2`; and `h_inf = eta(ln N − F(u*))`. Each follows from definitions the corpus has already fixed, each is verified numerically to floating-point noise, and **none carries any empirical weight whatever.** An identity cannot be wrong about the world because it does not address the world.
- Result (d) is a **reported outcome of a control experiment run by this corpus on itself** [Curzi, 2026u], together with a **structural argument** about what a shuffled-input control can and cannot detect. The structural part is not falsifiable: it is a statement about what a test is capable of seeing. The reported part is a single instance and is presented as one.
- The computed honesty items — the testimony likelihood ratio, the instrument audit counts, and the programme survival curve — are **arithmetic on stipulated inputs**. The likelihood ratio of exactly 1.000000 follows from a selection process defined in the calculation. The survival curve is computed from **stage priors the author chose and badged as chosen**, and a different set of priors gives a different curve; what is not arbitrary is that the curve was computed and published at all.

Two claims elsewhere in the paper would be hypotheses if they were advanced here, and they are not advanced here. That the framework's model form describes any real institution, mind or cell is a modelling postulate, tested nowhere in this paper and reported as unmet in Curzi [2026i]. And that the essence law `f(n) = k n^2` with `k = (9/8)^2/phi^2 = 0.483425732988414` describes anything at all is a **postulate with a defined constant**, discussed in §6.3 with its badges attached; **`k` is exact because it is defined, and attaching hertz to it is a separate and much stronger claim** which this paper does not make and which the verification chain of §9.3 lists as POSTULATED on its own line.

39.3 3 Method: what is being closed, and under what constraints

This paper collects and organises; it does not measure. The method is therefore a set of rules about what may be carried forward and how, and three constraints that bound everything the apparatus is used for. Sections 3.1 to 3.3 state what is being closed and the four things the mechanism is entitled to say. Sections 3.4 to 3.6 state the constraints: the identification problem, the analogy discipline, and the editorial rules under which cross-references were converted.

39.3.1 3.1 What is being closed, stated exactly

Two laws, once more, in the only form the corpus permits.

```
THE SURPLUS LAW                          CCP-A02, Theorem 2.1

    F(x,y) = f(u),   u = 1 - (x.y)^2     in [0,1]
    f(u)   = ln(1 + (N-1)u)              unique under axioms S1-S4
    N_eff  = 1 + (N-1)u                  F = ln N_eff
    0 <= F <= ln N                       CCP-009, Corollary 2.2


THE STORAGE LAW                          CCP-007, Section 3

    Q_{t+1} = (1 - delta) Q_t + eta S_t - C_t        S_t = F(U_t)

    h_t   = eta ln N - delta Q_t - C_t               the gauge
    Q_inf = (eta F(u*) - c*) / delta                 the fixed point
```

Everything in three volumes is either a theorem about those two lines, a badged assumption required to apply them to something real, a borrowed empirical result about the something, or an open question. **CCP-038 sorts one hundred and fifty-six claims into exactly those four bins, plus a fifth for the eight this corpus's own protocol demoted.**

39.3.2 3.2 The four things the mechanism is entitled to say

These are the through-lines. Each is derivable from the two boxes above, each was earned in a different volume, and none of them requires anything the formalism does not contain.

ONE. The loss term is what makes a recursion a circuit rather than a regress. `delta > 0` makes the map a contraction; a contraction has a unique fixed point and iteration reaches it from any start. With `delta = 0` there is no fixed point and nothing arrives: `Q_t <= Q_0 + t eta ln N`, unbounded and never converging. **The leak is not the flaw. The leak is why the loop is a device.**

TWO. Difference is what generates surplus, and only what is different is counted at all. `N_eff = 1 + (N-1)u`, so `u = 0` gives `N_eff = 1` regardless of `N`. Ten aligned instruments are one instrument. **Agreement between parallel instruments is echo, not evidence**, and the arithmetic prices it at exactly `ln 1 = 0`.

THREE. Geometry does work without power. A parabola focuses because of its shape. A room's modes are set by its dimensions. A star polygon's stroke count is `gcd(N,k)`. Boundary geometry alone determines which modes can exist, and that selection has measurable physical consequence. **Nothing is added and nothing is powered; the shape does all of it.**

FOUR. Every claim in this corpus is proved, badged, or named as open. There is no fourth category and there is no reserve. CCP-038 is the audit trail and it contains the demotions of the author's own formulas.

The rest of this paper is those four sentences, worked.

39.3.3 3.3 The four structural results the corpus produced

Distinct from the through-lines, these are the things the writing found that were not known when it started. Each is exact, each was computed, and each connects two parts of the corpus that were developed separately.

(a) The surplus law IS a Fechner law with threshold u_0 = 1/(N−1).
(b) Consonance and redundancy are the same phenomenon in the essence lattice.
(c) h_inf = eta ∗ (ln N − F(u∗)). Delta, C and Q all cancel at the fixed point.
(d) The shuffle cannot see repackaging. Only the structural test can.

§6.2 through §6.5 take them one at a time.

39.3.4 3.4 Identification: the parameter none of this apparatus can recover

The framework's substantive social claim is reflexive — the group's state influences each member and the members constitute the group — which as a model is a **linear-in-means reflexive specification**. Manski [1993] proves that its endogenous social effect is **not identified from observational data**, since it cannot be separated from contextual effects or from correlated unobservables however much data is collected. Shalizi and Thomas [2011] establish the same for network data, where influence and selection are generically confounded.

This bounds the entire apparatus assembled here, and the bound is stated in the Method rather than in a footnote. Appendix C gives measurement protocols for six parameters and three controls. Every one of them is a procedure that returns a number. **None of them, run on observational data, identifies a social coupling**, and the reflexive form the corpus writes down in its self-fulfilling material is exactly the specification Manski's result concerns.

The known remedies are randomisation, an instrument, partial-population variation, or stipulating the reference group and defending the stipulation, the last on the template of Brock and Durlauf [2001]. **This corpus adopts none of them.** Where the apparatus is used to license a claim about a real group, that claim is a **modelling position rather than an identified empirical result**, and Section 10 records it as such. The one field study documenting a reflexive fixed point with dates and price data is MacKenzie and Millo [2003]; it establishes that the mechanism is real and does not license reading a fitted coefficient as a measurement.

The Lucas critique applies to this corpus's own parameters and is not answered. Lucas [1976] holds that behavioural parameters estimated under one regime need not survive a change of regime, and Hacking [1995] names the mechanism in the human sciences: a classification changes the people classified. The framework asserts that systems change their behaviour when modelled and then treats `eta`, `delta` and `u` as stable enough to fit. **Either those parameters survive the framework's own predicted response to being measured, or the framework undercuts itself.** This paper takes the first horn as an assumption and labels it one.

The functional form of the surplus term is underdetermined by the evidence, and the appendices must not be read as settling it. A sub-linear power law [Latané, 1981] and an `n log n` form [Briscoe et al., 2006] fit the qualitative facts as well as `ln(1 + (N−1)u)` does. The logarithm in this corpus follows from the affine-effective-count design axiom [Curzi, 2026] and from nothing else established here. Its usual warrant is Fechnerian [Fechner, 1966], which requires the perceptual story to be true; the stronger warrant, requiring no psychological assumption at all, is Peters and Gell-Mann [2016], under which the time-average growth rate of a multiplicatively evolving quantity is the expectation of its logarithm. **Result (a) of this paper shows the surplus law *is* a Fechner law; it does not show that either law is right, and §6.2 says so in the same breath as the identity.**

39.3.5 3.5 Analogy discipline, and the constants

Every borrowing in the apparatus is one of three things, labelled at first use.

- **A shared mathematical form.** The storage law is a first-order discrete filter, and its pole, time constant and settling behaviour follow because the equation is the same equation, with no claim that the physics transfers [Oppenheim et al., 1996, Åström and Murray,

2008]. Convergence follows from the contraction property [Banach, 1922, Elaydi, 2005], and the leaky integrator has the same status in computational neuroscience [Gerstner and Kistler, 2002]. **This is the category most of the apparatus sits in.**

- **A derivation.** The Fechner identity, the octave-folding cost and the fixed-point headroom identity are algebra on definitions the corpus has fixed. The founding definitional move — comfort as the felt total of present costs, with the variational statement given as *stationarity* rather than minimality and licensing **no teleology** — is Curzi [2026] and is presupposed rather than restated. The contraction theorem is stated exactly, with its own refusal attached, in Curzi [2026]: a human life is not a metric space, nobody has measured the relevant contraction modulus, and **the theorem transfers as shape rather than guarantee.**
- **An analogy, and never evidence.** Reading a pitch-class identification as psychoacoustic consonance is an analogy and §6.3 refuses it. Reading the uncertainty relation as a statement about measurement disturbance is an error the corpus corrects in Curzi [2026], which also requires that a user be able to distinguish the theorem, the ordinary observer effect and the framework's own conjecture in three separate sentences on demand.

Entropy in this paper is Shannon entropy where it appears: surplus is a logarithm of an effective count, a pure number in nats, and no thermodynamic quantity is invoked. **Which effective count is meant is stated at every use.** `1/HHI` and `exp(H)` are Hill numbers of orders two and one [Hill, 1973, Jost, 2006]; the affine count `1 + (N−1)u` used throughout this corpus is **not** a Hill number but a linear interpolation between 1 and `N` adopted as a design axiom [Curzi, 2026], though it is exactly equal to a Rényi-2 participation ratio [Rényi, 1961]. The concentration instrument that *is* a Hill number of order two — with its worked case returning a naive HHI of 0.293430 and 3.41 effective jurisdictions against a bloc-HHI of 1.000000 and an effective count of 1.00 — is Curzi [2026], and the procedural form of the same rule, *collapse before measuring, because `ln(1.00) = 0` means a structure with one effective bloc has zero available surplus rather than a low one*, is Curzi [2026].

Two constants appear at high precision and their status differs. `k = (9/8)^2/phi^2 = 0.483425732988414` is a **DEFINED** constant: an exact ratio of exact quantities, computable to arbitrary precision, and its fifteen figures are legitimate for that reason alone. It is labelled *defined* at every appearance at that precision. **Attaching hertz to it is a separate and much stronger claim**, it is not established anywhere in this corpus, and it appears in the verification chain of §9.3 on its own line with status POSTULATED. Where the golden section appears in that constant, Markowsky [1992] is the standing caution: most popular claims about golden-ratio occurrence do not survive checking, and the constant here is **defined by the corpus rather than found in nature.** `ln 2 = 0.693147...` is a mathematical constant and needs no badge.

No scaling or power-law claim is made anywhere in this paper. Where the register in Appendix E records the corpus's borrowing of preferential-attachment results, it records with them that Clauset et al. [2009] is binding before any real graph is called scale-free, and that Broido and Clauset [2019] and Stumpf and Porter [2012] find such structure rarer than the literature suggests once that procedure is run, with Doyle et al. [2005] bounding the inference further.

39.3.6 3.6 Editorial rules for the converted cross-references

The source manuscript locates claims by book chapter number and refers to its own internal technical papers as Papers A to E. Neither survives conversion. **Where the manuscript cites Paper A, this paper cites Curzi [2026], which carries the surplus functional and its properties; where it cites Paper E, Curzi [2026a], which carries the storage law and its corollaries; where it cites Paper D, the result is a range or effective-count bound and the reference is CCP-009; and where it cites Paper B, the result concerns state encoding and category assignment and the reference is CCP-015.** Chapter numbers throughout the appendices have been replaced by the corpus paper identifiers that carry the corresponding material.

That mapping is an editorial inference rather than a statement by the source, and Section 10 records it. In the appendices, corpus papers are named as plain identifiers rather than as bibliographic citations, because the appendices are an index rather than an argument; those identifiers resolve against the corpus DOI 10.5281/zenodo.22079583, and the bibliography of this paper carries the subset cited directly in its running text.

39.4 4 What the four parts of the corpus established

Four bodies of work close here: a founding text and three technical series.

They were written over the same problem approached four times from four distances, and the problem is one sentence long. **Something makes a place or a practice or a person work better than it should, the effect is real enough to be noticed and unreliable enough to be lost, and nobody can say what it costs to keep.** the founding material taught a practitioner to notice it. This framework's first volume taught an engineer to build the circumstance that generates it. The second volume turned the instrument around and put it on the person holding it. The third carried it down to the cell, which is where it stopped, and stopped honestly.

What was found is smaller than what was looked for, and it is worth stating plainly at the top because the rest of this paper is a summary and summaries are where books get generous with themselves.

What was found is a recursion with a loss term, a geometry in which difference is the thing that is counted, and an audit trail. No mechanism was added to the world. Nothing here makes a symbol act on matter, no claim in these three volumes has any clinical or therapeutic content, and the formalism does not touch consciousness, thermodynamics, or physics beyond borrowing four theorems that belonged to other people first. **Every claim the corpus makes is in CCP-038 with a status on it: proved, badged, or named as open.** There is nothing behind the register. The register is the corpus.

That is a small inheritance and this paper's argument is that it is enough — that the mechanism is sufficient for everything the four bodies of work actually promised, and needs nothing added to it. The restraint is not modesty. **It is the load-bearing part: a framework that has to be supplemented to work is a framework that did not work.**

The four subsections below are **recapitulation**, and are included because the closing paper of a series is where a reader looks for the map. Each states what one part of the corpus established and cites the papers that established it; **no claim in this section is new here, and none should be cited to this paper.**

39.4.1 4.1 The founding material — the practitioner learning to read the gradient

The first book was a story with twelve nodes in it, and its instrument was attention. It taught a person to notice which chair, which room, which hour, which conversation made the work easier — and to treat that noticing as data rather than as preference.

Three of its nodes turn out to be theorems of the corpus, and the correspondence is exact rather than poetic.

The fixed-point node of the founding material — the room that was always furnished — is CCP-007, Theorem 3.5. The room feels like recognition rather than victory because a contraction does not arrive anywhere new. It arrives where it was always going to arrive, from any starting point, and the felt quality of arriving somewhere inevitable is what convergence to a unique fixed point is like from inside.

The recursion node of the founding material is the transient. `(1-delta)^t Q_0` decaying while the driven term builds. Going around again is not going in circles if the loop leaks: each

pass discards a fixed fraction of where the system started and adds a fixed increment of where it is going.

The departure-lounge node of the founding material is `L = max(0, M − C)`. "In those minutes there was no surplus to dream with" is `L = 0`, stated exactly. Escape is dreamt only from surplus, and a person whose whole moment is consumed by cost has none to dream with. **That formula is demoted on magnitude (CCP-038, E-05) and stands on structure**, which is the correct status for a sentence that was always an observation rather than a measurement.

Collected, the correspondences that hold exactly:

the founding material	Corpus result	Status
the departure-lounge node	`L = max(0, M − C)`; "no surplus to dream with" is `L = 0`	Structure stands; magnitude demoted (E-05)
the desire-line node	The comfort gradient as `−grad C`, read from where people actually walk	Badged identification
the many-windows node	`F_int = ln(1 + (N−1)u)` over one observer's own framings — CCP-A02 applied within a single observer	A derivation, not an analogy
the wave node	The future as an amplitude-weighted ensemble rather than a verdict	The corpus's own treatment; needs no quantum mechanics to stand
the echo node	`Delta_R = |ln(M'/M)|`, the world's answer over what was sent, read as a log-ratio	Structure stands; magnitude demoted (E-05)
the recursion node	The transient `(1−delta)^t Q_0` decaying while the driven term builds	CCP-007, Thm 3.1
the fixed-point node	`Q_inf = (eta F(u*) − c*) / delta`	CCP-007, Thm 3.5

Two of the seven are theorems, two are demoted on magnitude and standing on structure, and three are badged identifications. That distribution is the corpus in miniature.

What the founding material could not do is say what any of it cost. **An operator can read a gradient and still be unable to say whether the climb is affordable.** That is the gap the framework was written into.

39.4.2 4.2 The institutional-scope material — the engineer, the room, the table, the institution

The first volume answers the affordability question with one number.

```
h_t = eta ln N − delta Q_t − C_t
```

Below zero no building is permitted, and **no arrangement of people prevents decline** — that is Corollary 3.4 and it is the most consequential sociological statement in the corpus precisely because it forecloses the usual response, which is to change the people. At zero only maintenance is permitted. Above zero growth is available, by exactly `h` and not one unit more.

The volume's real content is the four levers and their honest ordering. `eta` is usually cheapest: the decision log, the written standard, the trained apprentice, the moved wall. `C` is next. `u` is the one that moves fast but not cheaply. `N` is the last and the least rewarding, because `ln N` grows so slowly that doubling the number of admitted domains buys `0.693147` nats wherever

it starts. **A headroom problem cannot be escaped by adding domains.** The lever ordering is itself an argument rather than a measurement (CCP-038, and Appendix E.4 item 7 says so).

And the volume made three things physical, which is where the third through-line entered. A room has `delta`, `eta`, `N` and `C` like anything else, and its `eta` — the rate at which encounters become durable relation or decision — is a design variable. A hard-surfaced room silences the quiet voice by physics rather than by dominance. A table with a head assigns `u = 0` to everyone below the salt, because their function becomes agreement. **Furniture is the cheapest governance instrument ever built, and it enforces continuously without anyone administering it.**

39.4.3 4.3 The personal-scope material — the return inward

The second volume's thesis is that psychology and sociology are not two subjects but one set of relations measured at two scopes, and it is badged as a thesis every time it is used.

The volume's strongest content is not the thesis. It is the demonstration that subjective magnitude is measurable, which settles a century-old question in the affirmative and with data: Weber's constant ratio, Fechner's logarithm, and above all **Stevens' method of magnitude estimation, which produces genuine ratio-scale measurements of subjective experience from human subjects and replicates.** The exponents are constants of a procedure rather than of nature, and the volume says so — but the existence of the procedure is what makes the whole enterprise something other than a category error.

From that came the Lens Theorem, which is the volume's most useful export. **If perception is a map converting multiplication into addition, it is a logarithm, and the map is unique up to scale.** Therefore equal perceptual steps are equal world ratios, never equal world differences; therefore a designer who promises a linear improvement in a quantity the recipient perceives logarithmically has promised an exponential and does not know it. The theorem is conditional and **the badge is on the antecedent, not on the inference** — which means the argument cannot be defeated by disputing the mathematics, only by defeating Weber and Fechner.

The volume's other export is the one that made the sonic material composable with everything else. **Four independent quantities in this corpus are all logarithmic** — surplus `F = ln N_eff`, pitch `p = log2 f`, sensation `S = k ln(I/I_0)`, and the moment in its log form `ln M = ln D − ln P + ln E`. Because all four are logarithms, **every conversion between them is linear in log space and the constants are exact**:

```
1 nat    = 1.442695 bits = 2.885390 octaves = 3462.4681 cents
1 bit    = 0.693147 nats = 2 octaves EXACTLY = 2400 cents exactly
1 octave = 0.346574 nats = 0.5 bits          = 1200 cents
```

One nat of interaction surplus is exactly 2.885390 octaves. That is not a metaphor about harmony; it is a unit conversion, and it is exact because both registers are logarithms of the same effective count. The reason the four compose is not coincidence and it is not mysticism: **the logarithm is the register perception itself operates in**, which is Weber's and Fechner's 19th-century finding and not this corpus's proposal.

And from the same volume came the demotion. The Formula Validation Protocol's second gate asks what scale type each term admits, division requires ratio scale with a true zero, and applied to the author's own keystone formula `M = (D/P)·E` the gate returns **DEMOTE.** The magnitude of the ratio, the distortion `Delta = |ln r|`, and the quantitative join from moment to character are not reportable as measured quantities. **What is demoted is magnitude, not structure**: the limiting cases stand, the `D = P` fixed point stands, and the ordering over-framed / calibrated / under-framed stands, because all three need only ordinal information.

The route out is Theorem 54.3 and it is costed at 123 items per term. That number is not an estimate; it is the inversion of the best-case measurement error, and one hundred and twenty-two items is not enough while one hundred and twenty-three is. **A book whose validation protocol demotes its author's own formula, prices the repair, and prints both is a book whose**

protocol is doing something. That is the strongest single thing the personal-scope material produced and it is a negative result.

39.4.4 4.4 The cellular-scope material — the descent to the cell, and where it stopped

The third volume carried the instrument down and found three things.

It found that the cell is a receiver, and that this is established rather than proposed. PIEZO1 and PIEZO2 are mechanically gated ion channels; sound is mechanical; force at the cell surface transmits through the cytoskeleton to the nucleus, stretches chromatin, and changes transcription. There is a real, prize-winning pathway from a pressure wave to a cellular signal, and CCP-030 establishes it and **stops exactly there.**

It found that the library is real. Roughly 98% of the genome is non-coding and largely regulatory; "junk DNA" is retired; syncytin, indispensable to placenta formation, derives from a retroviral envelope gene — ancient foreign text, inherited, later read out as vital function. And Yamanaka reprogramming demonstrates that four regulatory signals return a differentiated cell to pluripotency **without adding one base of new sequence.** From which comes the constraint that bounds the whole architecture and is the most useful sentence in the volume: **nothing that is not already latent can be reached.**

And it found the gap. There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal. Everything upstream is real. That step is not. The volume states it in every chapter that touches it, spends its most careful chapter on what an experiment would have to show, and does not blur it anywhere.

The volume also carried the two-column method, and the method's one refinement is the reason the volume is worth reading: **the (+1) column carries evidence and the (-1) column carries precedent, and two sources agreeing does not double the evidence when only one of them is evidence.** What independent arrival establishes is real and worth having — that a structure was findable by more than one route, which argues that the structure is in the territory and not only in the instrument. It is not a second measurement. CCP-038 applied that rule across all three volumes and found precedent in twenty-three of one hundred and fifty-six rows. **The other one hundred and thirty-three read *none claimed*, and that is the honest yield.**

39.5 5 Four through-lines

Each of the four statements below is derivable from the two laws of §3.1, each was earned in a different part of the corpus, and none requires anything the formalism does not contain. **They are restatements, not results**; the results are Section 6.

39.5.1 5.1 Through-line one — the loss term

Recursion earned its bad name in one narrow place: the call stack, where a function calling itself without a base case consumes a frame per call until the allocator runs out. The failure is loud enough to have become a metaphor — infinite regress, the mirror facing the mirror, going in circles.

That metaphor describes exactly one kind of recursion: recursion without a loss term.

Put the loss term in and the object changes category. `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t` is not a stack; it is an accumulator. In discrete time it is exactly a first-order low-pass filter: `Q` is charge, `delta` is leakage conductance, `eta S` is current in, `C` is load draw, `(1−delta)` is the pole. And because `delta > 0` makes the map a contraction, Banach guarantees a unique fixed point and convergence to it from any start.

That is why iteration arrives. Not because the loop is clever, and not because the practitioner persists, but because a contraction has nowhere else to go.

The two degenerate cases are the two failures everyone recognises. **Rumination is recursion with `eta` near zero**: the loop runs, nothing converts, and `Q` decays under `delta` while the effort is expended. **Tradition is recursion with `delta` near zero and no input**: nothing is lost and nothing enters, so `Q` holds forever at whatever it happened to be, which looks like continuity and behaves like a fixed point that has stopped being computed.

The design consequence is unglamorous and it is CCP-007's: **cadence beats intensity.** A low-pass filter integrates a periodic forcing function; the same total input delivered as one impulse mostly leaks away before it can accumulate. That is why scheduled, repeated, moderate input outperforms occasional heroics, and it needs no theory of motivation to explain.

39.5.2 5.2 Through-line two — difference generates surplus

The second through-line is one substitution and its consequences.

```
N_eff = 1 + (N−1) u        u = 0  ->  N_eff = 1  ->  F = ln 1 = 0
```

Difference is counted first, and only what is different is counted at all. A roster of ten becomes an effective count of one the moment the ten are parallel, and no amount of headcount repairs it, because `N` enters only through `u`'s multiplier.

Three separate literatures say the same thing from three directions and none of them knew about the others. Condorcet's jury theorem requires independence of member errors. Bikhchandani, Hirshleifer and Welch show that informational cascades stop aggregation because agents rationally cease using their private signals. Woolley and colleagues found that the collective intelligence factor correlates with equality of conversational turn-taking. **Independence, non-copying, turn-taking: three names for `u > 0`.**

And the concavity is the merciful clause. `F` is strictly concave, so the second genuinely different view is the largest single gain available and the fortieth is glazing bills. At `N = 6`, half the ceiling arrives by `u = 0.3`. **Three honest windows set wide capture most of what a gallery could** — which is why this framework never recommends assembling a hundred of anything.

The failure mode that follows is the dominant one in modern institutions and the arithmetic prices it exactly. Three glasses ground from one doctrine, their agreement read as convergence: `u = 0`, `N_eff = 1`, surplus `ln 1 = 0`. **A second opinion from someone whose training and incentives replicate the first is worth exactly nothing, and the law says so with a number.**

39.5.3 5.3 Through-line three — geometry does work without power

The third through-line is the one that connects a room to a sigil without any mysticism being required, and it is worth laying out because it is the place where a reader might expect the corpus to reach for something it has refused.

A parabola focuses every ray parallel to its axis through one point. It does this because of its shape. No power is supplied, no mechanism is engaged, nothing is added. The Denge acoustic mirrors — concrete parabolic dishes built in the 1920s and 1930s to detect aircraft before radar — worked for exactly that reason and consumed nothing.

A room's modes are set by its dimensions. Reverberation time follows from volume and absorption; a hard-surfaced room is loud for the quiet speaker not because anyone dominates but because the geometry sums the reflections.

A star polygon `{N/k}` decomposes into exactly `gcd(N,k)` closed paths of length `N/gcd(N,k)`. The hexagram is two triangles because `gcd(6,2) = 2`. This is number theory, it is true, and it means the number of separate strokes in the figure is an exact arithmetic invariant of the figure. **That is the entire claim, and it is claimed about the drawing and not about the world.**

Nothing here says a figure does anything. It says a figure has an arithmetic, and the arithmetic is computable.

And the purest physical case is Casimir. Two conducting plates exclude every mode whose wavelength does not fit between them, and the exclusion produces a measured force. Computed: `1.3001e5` Pa at 10 nm — 1.28 atmospheres — falling as the fourth power to `1.3001e-15` Pa at 1 mm. **At any biological or architectural scale the Casimir force is utterly negligible. It does nothing in a cell and nothing in a room, and this corpus says so wherever Casimir appears.**

But the reason to cite it is not force. **Casimir is the most rigorous demonstration in physics that boundary geometry alone determines which modes can exist, and that this selection has measurable consequence.** Nothing is added; nothing is powered; the shape does all of it. That is the framework's thesis in its purest available form, and it is presented as exactly that and nothing more.

There is a cost to geometry doing the work, and it is a specific one. **A shape that works is invisible while it works.** Star and Ruhleder's finding is that infrastructure becomes visible chiefly upon breakdown, and Weiser's aphorism — that the most profound technologies are those that disappear — says the same thing approvingly. In the storage law that invisibility has an exact consequence: **`delta` is systematically underestimated wherever it has not been measured**, because the thing decaying is the thing nobody can see. A room that has been quietly producing turn-taking equality for nine years produces no evidence that it is doing so, and the case for maintaining it is therefore always weaker than the case for spending the money elsewhere. That is not a failure of judgement. **It is a structural property of working infrastructure**, and the only remedy is to measure `delta` on purpose, under zero input, before the argument is needed.

The design consequence runs through every physical chapter. **A shape that does the work needs no maintenance budget**, which in the storage law means it contributes to `eta` without contributing to `C` and without decaying under `delta`. That is why the table with no head outperforms the facilitator, why the alcove outperforms the policy about interruptions, and why a threshold deep enough to stand in outperforms a sign asking people to pause. **The geometry enforces continuously, at zero marginal cost, and it never forgets.**

39.5.4 5.4 Through-line four — proved, badged, or open

The fourth through-line is a discipline rather than a result, and it is the one that took the most words to keep.

Every claim in three volumes is in CCP-038 with a status on it. Fifty are theorems, and not one of the fifty is a claim about a person, a room, or a cell — Group A is arithmetic, and every step from arithmetic to recommendation passes through a badged assumption. Twenty-eight are postulates, and four of them (the affine design axiom, the linear storage law, the modelling map onto social systems, and cost commensuration) carry more weight than any citation in the corpus. Fifty-eight are borrowed, each with the setting it was obtained in and the transfer assumption stated. Twelve are speculative, of which four are rejections printed as rejections. **And eight are demoted, including the author's keystone formula.**

The house rule from the founding material survives intact: **the mysticism is permitted, but it is permitted in the open, with its badge on.** CCP-038 is what that rule looks like after three volumes.

Two things about the register deserve to be said at the end rather than in it.

The first is that the demotions are the evidence. A protocol that has never rejected anything is a decoration, and this one rejected the formula the author cares most about, at the gate the author designed, and then printed the price of the repair. **Nothing in this corpus argues for its own reliability as effectively as the eight rows in Group E.**

The second is that the register is written to be used against the corpus. Entry IDs are stable so they can be cited in an argument. The "what it measured" field exists so a reader can object that the setting does not transfer. And a claim shown to be wrong is withdrawn rather

than defended. **A framework that cannot be told apart from its own decoration is useless as engineering**, and the only remedy is a document that says, item by item, which is which.

39.5.5 5.5 What the mechanism is sufficient for

Here is the restraint, stated as a positive claim rather than as a list of refusals.

The mechanism is sufficient to explain why cadence beats intensity, because a leaky integrator integrates periodic forcing and dissipates impulses. It needs no theory of discipline.

It is sufficient to explain why the second opinion is worth more than the fortieth, because `F` is strictly concave. It needs no theory of creativity.

It is sufficient to explain why an institution ossifies, and the explanation is more precise than the usual one: within the plain model it does not ossify at all (§6.4), so the observed ossification requires a term the plain model lacks, and CCP-008 names the candidates — exogenous `Q` accretion and superlinear cost. It needs no theory of bureaucratic psychology.

It is sufficient to explain why a hard room silences the quiet voice, because reverberation sums reflections and the loudest source wins. It needs no theory of dominance.

It is sufficient to explain why a silo cannot be exhorted out of, because a silo is a local minimum and every direction out is uphill; the method is to build a lower valley next door. **This is also why the method has no coercive failure mode: its worst outcome is that someone declines to move.**

It is sufficient to explain why writing it down changes the outcome, because `eta` multiplies `S` and an unwritten conversation has `eta` near zero. It needs no theory of memory.

It is sufficient to explain why a large group underperforms a small one, because coordination paths grow as `n(n−1)/2` while surplus grows as `ln N` — quadratic against logarithmic, with a crossing point that can be computed for any given `eta`. It needs no theory of committees.

It is sufficient to explain why individuality has to be protected rather than tolerated, because `u` is the multiplier on the entire surplus term and merging the members drives it to zero. **That is the ethical constraint arriving as an arithmetic one**, and it is the one place in the corpus where the two coincide exactly: preserving the ability to differ is not a concession to comfort, it is the precondition for the collective to compute anything at all.

And here is what it is not sufficient for, stated with equal flatness. **It does not explain consciousness and does not touch it.** It proves nothing about physics or thermodynamics — CCP-007 is explicit that no negentropy law is proved. **It confers no clinical or therapeutic benefit and no such claim appears anywhere in three volumes.** It gives symbols no efficacy: gematria is a hash function, a sigil is an encoding operation with an arithmetic invariant, and neither description carries any claim that the map is meaningful or that the figure acts. And it does not establish the cellular channel, which is §8.1.

A mechanism that needs nothing added to it to do the work it claims is the only kind worth having. Everything the four bodies of work actually promised is in the list above. Everything else was never promised.

39.6 6 Four structural results

39.6.1 6.1 The genuine structural results — an introduction

Four results emerged from the writing that were not in the papers, not in the founding material, and not anticipated in the plan. Each connects two parts of the corpus that had been developed independently, each is exact, and each was computed rather than asserted.

They are the closest thing this corpus has to a discovery, and none of them is large. That is the correct size for what a formalism yields when it is being honest.

39.6.2 6.2 Result (a) — the surplus law is exactly a Fechner law with threshold `u_0 = 1/(N−1)`

Take the surplus law and factor `(N−1)` out of the logarithm:

```
F(u) = ln(1 + (N−1)u)
     = ln( (N−1) ∗ ( u + 1/(N−1) ) )
     = ln( (u + u_0) / u_0 )                with  u_0 = 1/(N−1)
```

That is Fechner's law in its threshold-corrected form, `S = k ln((I + I_0)/I_0)`, with `k = 1`, the stimulus `I` being non-parallelism, and the threshold intensity `I_0` being `u_0 = 1/(N−1)`. The identity was verified for `N = 2` through `N = 100` at seven values of `u` each; the maximum residual is `8.9e−16`, which is floating-point noise. **The source manuscript quotes `4.4e−16` here and `8.9e−16` in its honesty appendix. Recomputing the identity over both grids returns a maximum absolute difference of `8.881784e−16`, so the appendix figure is the correct one, the larger value is quoted throughout this paper, and the discrepancy is recorded in Limitation L-11 rather than repaired silently.**

And the threshold has a meaning that is not an artifact of the algebra. Evaluate the law at it:

```
N_eff(u_0) = 1 + (N−1) ∗ 1/(N−1) = 2       exactly, for every N
F(u_0)     = ln 2 = 0.693147 nats = 1 bit = 2 octaves
```

The threshold is exactly the point at which the effective count reaches two. Below `u_0` the system does not yet contain a second independent voice; at `u_0` it contains precisely two; above it, more. Fechner's `I_0` is the intensity at which sensation begins. **Here the intensity at which surplus begins is the non-parallelism at which the second window opens, and it is one full bit wide.**

The threshold's size is a design fact:

```
N = 2   ->  u_0 = 1.000000     the threshold IS the ceiling
N = 3   ->  u_0 = 0.500000     half the range
N = 4   ->  u_0 = 0.333333
N = 6   ->  u_0 = 0.200000     one fifth
N = 12  ->  u_0 = 0.090909     one eleventh
```

In a two-category system the second voice arrives only at perfect orthogonality; in a twelve-category system it arrives at nine percent of the range. That is a precise statement of why narrow systems are so hard to get anything out of: not because their members are worse, but because the threshold sits at the top of their range.

What this result is: an exact algebraic identity connecting two formalisms developed a hundred and sixty-six years apart. **What it is not: evidence for either.** CCP-A02's law is a theorem under its axioms and gains nothing from resembling Fechner's; Fechner's law is an empirical claim resting on a badged postulate about just-noticeable differences and gains nothing from resembling Paper A's. **The correspondence is structural, and structural similarity is not identity.** What it does establish is the thing independent arrival always establishes: **that the logarithm is not a stylistic choice in this corpus. It is the register both routes landed in.**

39.6.3 6.3 Result (b) — consonance and redundancy are the same phenomenon in the essence lattice

This one connects the sonic material to the surplus material with no adjustment, and the connection is exact.

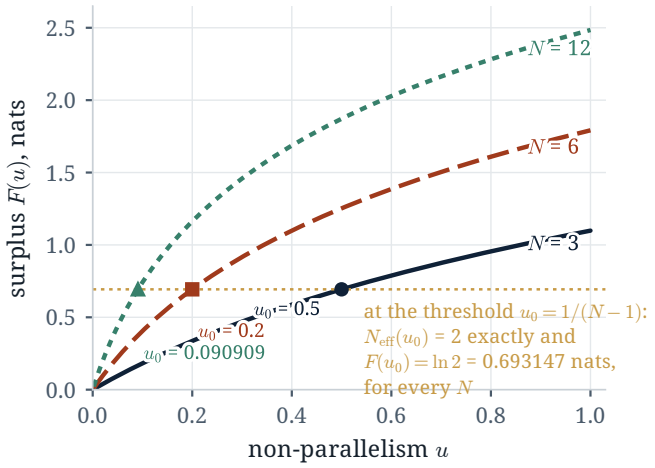


Figure 39.1: **The surplus law is exactly a Fechner law with an absolute threshold at `u_0 = 1/(N-1)`, and at that threshold the effective count is exactly 2 and the surplus is exactly `ln 2` for every `N`.** The two forms `ln(1+(N-1)u)` and `ln((u+u_0)/u_0)` agree to a worst residual of 8.9e$-$16 over `u` = 0.01 to 1.00 at `N` = 2, 3, 5, 8, 12, 25 and 100. The threshold is 1.000000 at `N` = 2 and 0.090909 at `N` = 12, so a narrow system opens its second independent voice only near perfect orthogonality. Exact algebra (Section 6.2); generated by `figs/src/CCP-039_fig01.py`.

The essence law's published recipe collapses to `f(n) = k n^2` with `k = (9/8)^2/phi^2 = 0.483425732988414 Hz` — verified against every published value in the source to fourteen significant figures. **`k` is a chosen reference constant and the square law is a postulate about essence structure; both carry their badges here and everywhere.**

Because `(2n)^2 = 4 n^2`, doubling an index raises the frequency by a factor of four, which is **exactly two octaves**. Therefore indices `n` and `2n` occupy the same pitch class, and more generally two essences share a pitch class if and only if their indices have the same odd part. An ordered set `1..N` therefore contains exactly `ceil(N/2)` distinct pitch classes.

Now price that identification in surplus. Folding an even-`N` essence set by octave equivalence collapses `N` categories to `N/2` classes, so:

```
surplus before folding  =  ln N
surplus after folding   =  ln (N/2)
loss                    =  ln N − ln(N/2) = ln 2      EXACTLY, for every even N
```

Computed for `N = 4, 6, 8, 10, 12, 16, 20, 24, 32, 64`: the loss is `0.6931471806` in every case, and `loss / ln 2 = 1.0000000000` in every case.

Convert the loss into pitch using the corpus's own exact key, `octaves = 2F / ln 2`:

```
loss in nats     = 0.6931471806
loss in bits     = 1.0000000000        one bit
loss in octaves  = 2.0000000000        two octaves

the interval that was identified away: f(2n)/f(n) = 4 = 2 octaves
```

The surplus surrendered to redundancy is numerically equal to the interval declared consonant, in the same units, by the same conversion key. Two octaves of pitch, one bit of surplus, one `ln 2` of nats — three names for one quantity. The concrete case: a twelve-essence set has indices with odd parts `1,1,3,1,5,3,7,1,9,5,11,3`, giving six pitch classes, `ln 12 − ln 6 = ln 2` exactly.

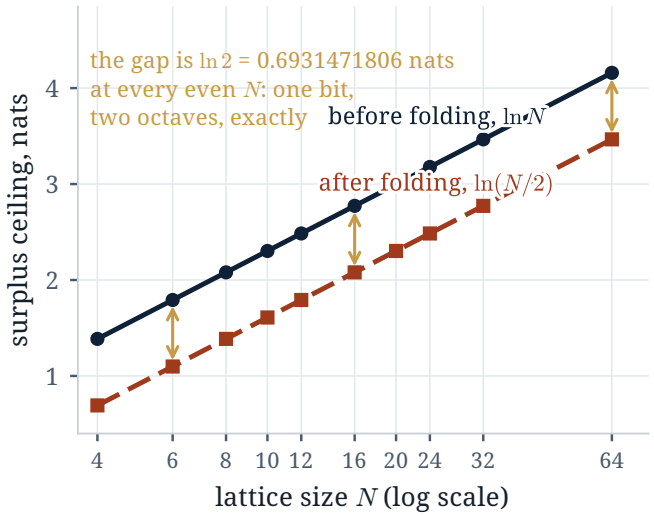


Figure 39.2: **Folding an even essence lattice into one octave costs exactly `ln 2` = 0.6931471806 nats — one bit, two octaves — at every `N`.** Computed at `N` = 4, 6, 8, 10, 12, 16, 20, 24, 32 and 64, the loss `ln N − ln(N/2)` returns the same value in every case and the ratio to `ln 2` is 1.0000000000 in every case. Exact arithmetic inside a postulated lattice, `f(n) = k n^2` with `k` defined (Section 6.3); generated by `figs/src/CCP-039_fig02.py`.

What the result says: in this lattice, calling two things consonant and calling them redundant are the same operation, and it costs exactly one bit. Consonance is not a bonus that arrives free; it is an identification, and every identification is a reduction of effective count.

What the result does not say, and the limits matter. "Consonance" here means the exact-ratio octave identification internal to this lattice — it is not a claim about the psychoacoustics of consonance, which is a different subject with its own contested literature. The result holds inside a postulated structure (`f = k n^2`, badged) and is exact there. **And the output remains an essence frequency, not a physical resonance: no claim is made that any material, molecule, element or person physically resonates at these values.** That clause is absolute in this corpus and is restated here at its last appearance.

39.6.4 6.4 Result (c) — `h_inf = eta * (ln N - F(u*))`

Take the headroom gauge to the fixed point and something cancels that should not.

```
Q_inf  = (eta F(u*) - c*) / delta
h_inf  = eta ln N - delta Q_inf - c*
       = eta ln N - (eta F(u*) - c*) - c*
       = eta * ( ln N - F(u*) )
```

Delta, cost and stored order all disappear. Verified against direct computation across four parameter sets spanning `eta` from `0.12` to `0.50`, `N` from 4 to 20, `delta` from `0.03` to `0.20`: agreement below `1e-16` in every case.

The consequence is the interesting part. Since `F(u) <= ln N` always by CCP-009's corollary, the bracket is non-negative, so:

```
h_inf >= 0        ALWAYS, in the plain linear model
```

The plain model never ossifies by itself. Whatever else happens, a converged system retains non-negative headroom, and the amount it retains is `eta` times the gap between the ceiling and the surplus it actually achieves.

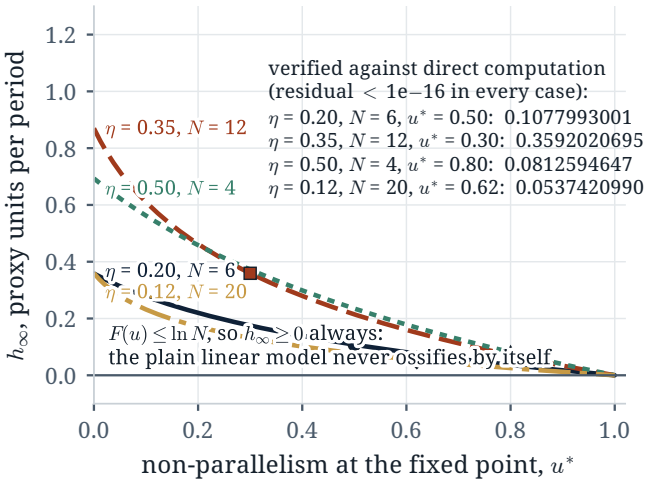


Figure 39.3: **Retention, cost and stored order all cancel at the fixed point, and the result is non-negative for every admissible `u*`: the plain linear model never ossifies by itself.** The identity `h_inf = eta(ln N − F(u*))` was verified against direct computation at four parameter sets, returning 0.1077993001, 0.3592020695, 0.0812594647 and 0.0537420990 with residuals below 1e$-$16. Modelled from the stipulated parameter sets of Section 6.4; generated by `figs/src/CCP-039_fig03.py`.

That is inconvenient, because ossification is the most reliably observed behaviour of successful institutions, and CCP-008 had argued for it from the shape of `h_t` along the transient — `h` does fall as `Q` rises, and it does so throughout the approach. **Both are true, and the resolution is not a repair to the theorem.** Along the transient headroom falls; at the fixed point it stops falling, at a non-negative value.

So the observed phenomenon requires a term the plain model does not contain, and CCP-008 names the candidates: exogenous `Q` accretion — stored order arriving from outside the interaction loop, as regulation, inherited estate, or accumulated obligation — and cost that rises with `Q` rather than staying additive. With an accretion term `A`, the identity becomes `h_inf = eta(ln N − F(u*)) − A`, and it goes negative exactly when `A` exceeds the geometric shortfall.

A model that cannot produce a phenomenon everyone observes is reporting that the mechanism is outside the model, and that is a useful thing for a model to say. It converts "institutions ossify" from a lament into a question with two named suspects and a testable difference between them.

There is one more reading of the identity, and it is the design one. `h_inf` scales with `eta` and with nothing else, and it is largest when `F(u*)` is *smallest*. **A converged system operating far below its surplus ceiling retains the most headroom, and a system extracting everything its geometry allows retains the least.** Running at the ceiling is not efficiency; it is the elimination of slack, and CCP-008's queueing argument said the same thing in a different notation.

39.6.5 6.5 Result (d) — the shuffle cannot see repackaging

The fourth result is the one that changes how the framework must be tested, and it came out of an attempt to break the framework rather than to support it.

The falsification suite requires beating four baselines: persistence, direct-input, matched autoregressive, and a shuffled-input control. The shuffle is the most trusted of the four, because destroying the temporal correspondence between input and output should destroy any spurious fit.

A filtered repackaging of the input passed all four, including the shuffle, at `p = 0.0005`. A quantity that was a transformation of the input rather than a prediction of the output survived

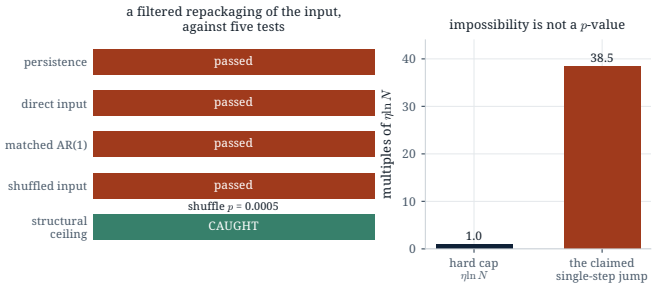


Figure 39.4: **A filtered repackaging of the input passed all four empirical baselines, including the shuffled-input control at `p = 0.0005`, and was caught only by the structural ceiling.** Its claimed single-step jump was 38.5 times the hard cap `eta ln N`, which no parameter setting, significance level or sample size admits: impossibility is not a `p`-value. Reported outcomes of the corpus's own control experiment (Section 6.5); generated by `figs/src/CCP-039_fig05.py`.

every empirical audit the protocol had.

It was caught by one thing only: the structural ceiling.

```
r_max     = eta ln N - delta Q - C          Corollary 4.2
hard cap  = eta ln N                       since delta, C >= 0

observed single-step jump = 38.5 x eta ln N
```

The step was impossible under the model, and impossibility is not a p-value — it is a contradiction. No parameter setting, no significance level, and no sample size makes a jump of `38.5` times the theoretical maximum admissible.

The rule that follows is now propagated through five chapters and two appendices: **the shuffle is an empirical audit; only the structural tests detect repackaging.**

The general lesson is larger than the instance, and it is the best argument in the corpus for having a formalism at all. A statistical test asks whether a pattern could have arisen by chance. A structural test asks whether the pattern is possible. **The second question is stronger, and only a theory can ask it.** A framework with no ceiling has no way to say "that cannot have happened", and a framework that cannot say that will eventually publish a repackaging as a result.

This is also the clearest case in the corpus of a demotion strengthening the framework. CCP-038 records it as E-07, and it is one of the eight entries that will never be promoted: **it is a permanent structural finding about what shuffles can see.**

39.7 7 Results

39.7.1 7.1 The smallest complete instance

The last worked example is the smallest one that contains the whole mechanism: **two people, one recurring conversation, one page written afterwards.**

Two people, drawing on four genuinely distinct categories between them, so `N = 4`. Their independent written predictions on shared questions decorrelate to `u = 0.55`. They meet weekly. One of them writes a page afterwards and files it where the other can find it, which

is the entire `eta` intervention. Cost is `C = 0.15` per week — the hour, the preparation, the interruption. Retention loss is `delta = 0.10` per week, which is a half-life of `6.58` weeks for anything not re-touched. They begin with `Q_0 = 0.80`.

Every number below is arithmetic on plausible parameter values. None was measured, and Appendix E.4 item 1 applies without exception.

```
ln N     = 1.386294 nats                    = 2 bits exactly
N_eff    = 1 + 3(0.55) = 2.6500             two people, 2.65 effective views
S = F(u) = ln(2.6500) = 0.974560 nats       70.3% of the ceiling

eta ln N = 0.50 x 1.386294 = 0.693147
delta Q  = 0.10 x 0.80     = 0.080000
C                          = 0.150000

h_0 = 0.693147 - 0.080000 - 0.150000 = +0.463147        growth available
```

The step, and then the trajectory:

```
Q_1 = 0.90(0.80) + 0.50(0.974560) - 0.15 = 1.057280      step = +0.257280

t      Q          h
----------------------------
0    0.800000   +0.463147
2    1.288832   +0.414264
4    1.684785   +0.374669
6    2.005508   +0.342596
8    2.265293   +0.316618
10   2.475719   +0.295575
12   2.646164   +0.278531

Q_inf = (0.50 x 0.974560 - 0.15)/0.10 = 3.372798
h_inf = 0.50 x (1.386294 - 0.974560) = 0.205867
tau   = 9.4912 weeks;  95% of Q_inf at t = 28.4 weeks;  99% at t = 43.7 weeks
```

Seven months of a weekly conversation and one page to reach ninety-five percent of everything this arrangement will ever hold. That is what a contraction feels like from inside: fast at first, then slower, then unmistakably arrived. And `h_inf = 0.205867` is positive, as §6.4 guarantees — this pair does not ossify.

Now remove the page. Same two people, same `u`, same hour, nothing written:

```
eta = 0.00  ->  h = -0.230000   Q_inf = -1.500000
eta = 0.10  ->  h = -0.091371   Q_inf = -0.525440
eta = 0.25  ->  h = +0.116574   Q_inf = +0.936399
eta = 0.50  ->  h = +0.463147   Q_inf = +3.372798
```

At `eta = 0` the equilibrium is negative, which is the model's way of saying the arrangement is unsustainable as specified. The conversation still happens. The `u` is still `0.55`. The surplus `S = 0.974560` is still generated every week and it still goes nowhere, because nothing multiplies it. Between `eta = 0.10` and `eta = 0.25` the sign of `h` changes, and somewhere in that interval is the difference between two people who talk for a year and have nothing, and two people who talk for a year and have something.

The intervention that crosses it is a page. Not a better conversation, not more people, not a larger `N` — the ceiling `ln 4 = 1.386294` was never the binding constraint. **The binding constraint was that nothing was being converted, and the arithmetic located it in one line.**

That is the whole manual in nine lines of arithmetic, and it is the correct size.

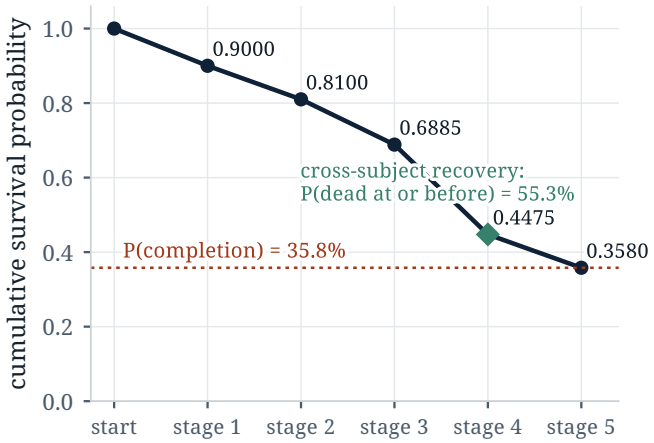


Figure 39.5: **The research programme publishes a 35.8 per cent probability of completing itself and a 55.3 per cent chance of dying at or before the stage where a formalism that describes meets the demand that it measure.** Cumulative survival runs 0.9000, 0.8100, 0.6885, 0.4475, 0.3580 across the five stages. Modelled from stage priors the author chose and badged as chosen (Appendix E.11); a different set of priors gives a different curve. Generated by `figs/src/CCP-039_fig04.py`.

39.7.2 7.2 The computed honesty items

Three quantities are computed rather than asserted in the honesty appendix, and they are collected here because they are the paper's other non-recapitulative content. Their derivations and their full statements are in Appendix E.9, E.10 and E.11.

The likelihood ratio of testimonial evidence, under a selection process that publishes successes and discards failures, is exactly `1.000000`. Not approximately one and not weakly informative: exactly one, so the evidence does not move the posterior at all, because the selection guarantees that successes are observed whether or not the method does anything. All four hypothesised success rates in the worked case produce the same 500 published accounts. **A thousand testimonials from people for whom a method worked carry precisely the evidential weight of zero testimonials.** This is the exact quantification of survivorship bias and it is why the precedent column of the corpus's register can never be promoted to evidence: a historical corpus is itself a survivorship-filtered record, and the selection is invisible from inside the surviving set. The reform literature's account of the same phenomenon in the modern record is Ioannidis [2005] and Open Science Collaboration [2015].

The companion result prices the alternative. Likelihood ratios measured at **2.0000** for a prediction specifying only a direction, against **11.3577** for one specifying a magnitude with a tolerance: **a prediction with a number attached is about 5.7 times more evidentially powerful**, equivalent to about 3.5 direction calls for one measurement. That is exactly Meehl [1967] and Meehl [1990] restated as arithmetic — directional predictions in psychology are nearly costless, and a corpus that makes point predictions with constants is answering Meehl rather than evading him. **Hence the standing rule: a belief that predicts only a direction is not yet testable.**

The instrument audit found 19 defined estimators, 1 validated value and 0 first-party measurements across 27 parameter-by-scope cells [Curzi, 2026]. The single validated cell is a borrowed median from one cell line. The correspondence between scopes then fails in five distinct ways — a mind is not a small society, a cell is not a small mind, exit and therefore the ethics do not descend, aggregation is not always well defined, and no composition theorem connects any scope to any other. **The break to carry forward is the third: the arithmetic descends to the cell and the ethics do not**, so a reader who imports the formalism downward has imported a model and left the constraint behind.

The research programme publishes its own odds of failure. Twenty-one preregistered kill conditions, four of them terminal; 8,554.8 researcher and coordinator hours, which is 5.35 full-time-equivalent years; 570 participants; 71 months strictly sequential, or 61 with a 30 per cent overlap of two stage pairs. Under the author's explicitly badged stage priors the cumulative survival probability runs **0.9000, 0.8100, 0.6885, 0.4475, 0.3580**, so **P(completion) = 35.8 per cent** and **P(death at or before the cross-subject recovery stage) = 55.3 per cent.** The programme is more likely than not to be dead before field trials begin, and the heaviest prior sits at the stage where a formalism that *describes* meets the demand that it *measure*. **A programme with a high stated completion probability has almost certainly written its kill conditions too loosely**, which is the more useful reading of the number.

39.7.3 7.3 The results this paper contributes

R1 — The surplus law is exactly a Fechner law with threshold `u_0 = 1/(N−1)`. `F(u) = ln((u + u_0)/u_0)`, verified for `N = 2` to `100` at seven values of `u` to a worst residual of `8.9e−16`; and `N_eff(u_0) = 2` exactly and `F(u_0) = ln 2` for every `N`.

R2 — Pitch-class folding of the essence lattice costs exactly `ln 2 = 0.6931471806` nats — one bit, two octaves — for every even `N`, computed at `N = 4, 6, 8, 10, 12, 16, 20, 24, 32, 64`.

R3 — `h_inf = eta(ln N − F(u*))`: retention, cost and stored order cancel at the fixed point, and since `F(u) <= ln N` always, `h_inf >= 0` always.

R4 — The shuffled-input control cannot see repackaging. A filtered repackaging passed all four empirical baselines including the shuffle at `p = 0.0005` and was caught only by `r_max = eta ln N − delta Q − C`, against an observed jump of `38.5 x eta ln N`.

R5 — Testimonial evidence under success-publishing selection has likelihood ratio exactly `1.000000`; a direction-only prediction is worth `2.0000` against `11.3577` for a magnitude call with a tolerance, about 5.7 times.

R6 — The research programme's published survival curve: 21 preregistered kill conditions (4 terminal), 8,554.8 hours, 570 participants, 71 months, **P(completion) = 35.8 per cent**, **P(death at or before the cross-subject stage) = 55.3 per cent.**

CR — The counter-result. This closing register publishes the work's odds against itself — a 35.8 per cent probability of completing its own research programme and a 55.3 per cent chance of dying before field trials — and reduces the entire testimonial corpus, including the historical precedent column, to a likelihood ratio of **exactly 1.000000**.

39.8 8 Discussion

39.8.1 8.1 What is not known

Books of this kind end on a claim. This one ends on two things that are not known, because they are what the work actually produced last and because ending anywhere else would misrepresent the state of it.

The first is empirical and it belongs to the cellular-scope material.

> **There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal.**

Everything upstream is real and independently well-supported. Mechanotransduction is Nobel-recognised. Force at the cell surface reaches the nucleus and changes transcription. DNA stores arbitrary data. Cas12a processes its own array. Targeting and editing are separable, so an architecture that addresses without editing is coherent rather than promissory. **And the specific step — from an acoustic signal to a sequence-specific cellular read — is established by none of it.** CCP-031 states what an experiment would have to show, what controls it would

need, and what result would settle it in either direction. **It is one well-specified falsifiable gap, and naming it sharply is the most useful thing three volumes did with the cell.**

The second is mathematical, it belongs to no volume, and nobody has been asked to notice it.

> **There is no composition theorem in this corpus.** Nothing here derives institutional `Q` from personal `Q`, or personal `Q` from cellular `Q`. No aggregation operator is defined between scopes. No conservation law links them. **No argument is offered that the three accumulators are coupled at all rather than merely similar in form.**

The three-scope correspondence is a claim about model form and CCP-038 spends a third of its length saying so. But the absence underneath it is sharper than the badge admits. **What would close it is a stated aggregation operator together with a proof that the storage law is closed under it** — that a population of systems each obeying the law, combined by that operator, yields a system obeying the law with derived parameters. If such a practitioner exists, hyperrecursive symmetry stops being a modelling convention and becomes a theorem with a derivation for how parameters change with scale. **If no such operator exists, the correspondence is exactly what CCP-038 says it is and nothing more, and the three volumes are three separate applications of one equation that happen to look like a ladder.**

Which of those is true is not known here.

It requires no laboratory, no funding, and no permission. It requires someone to define the operator and try to prove closure, and to publish the result including — especially — if the result is that the law is not closed under any natural aggregation.

Two problems, then. One needs an experiment and one needs an afternoon with a proof. Neither has been done, and the corpus stops at both of them rather than past them.

The instrument is on the bench. It has one gauge that reads at two of its three scopes, a hundred and fifty-six entries saying what it is entitled to claim, eight of them recording where it was caught being wrong, and two open questions that would change its status in opposite directions depending on how they come out. **That is what was built, it is smaller than what was hoped for, and it is the part that survived being checked.**

Anyone may take it further. **The only condition this research programme asks is the one it has tried to keep for three volumes: say which claims are proved, badge the ones that are not, name the ones that are open, and publish the demotions when the gates catch the work.**

The gates catch the work that runs them. They caught the author on page one of the personal-scope material's formulary, at the gate he wrote, on the formula the whole volume was built around.

That is the work behaving correctly, and it is where it ends.

39.8.2 8.2 What the closing position amounts to

The mechanism this corpus assembled is a recursion with a loss term, a geometry in which difference is what is counted, and an audit trail. No mechanism was added to the world. Nothing in it makes a symbol act on matter, no claim in the series has clinical or therapeutic content, and the formalism does not touch consciousness or thermodynamics — it borrows four theorems that belonged to other people first, and it names them.

That the inheritance is small is the point rather than an apology. A framework that has to be supplemented in order to work is a framework that did not work, and the closing position defended here is that the mechanism is sufficient for what the series actually promised: an arithmetic that transfers wherever its model form is adopted, three ordinal diagnostics that read without a calibrated instrument, and a register that says of every claim whether it is proved, badged, or open. **What it is not sufficient for is any claim requiring a measurement, because the corpus has made none** [Curzi, 2026].

Two positions elsewhere in the corpus bound this one and are inherited rather than argued again. The concentration instrument that this framework's effective count must not be

confused with — a Hill number of order two, with `1/HHI` giving 3.41 effective jurisdictions on a naive reading and 1.00 once coordinating jurisdictions are collapsed into one bloc — is Curzi [2026], and its procedural form, *collapse before measuring*, is Curzi [2026]. And the sharpest instance in the corpus of a framework reporting against itself is Curzi [2026], where **the only reading under which an assembled architecture clears the build threshold is the naive one that assumes zero correlation between coordinating jurisdictions, at `h = +0.083099`; every defensible bloc assignment returns a negative surplus, and the paper says so in print.** The four-drawer evidential classification the register here descends from — PROVED, POSTULATED, BORROWED, UNTESTED, with the surplus form recorded as **proved conditionally on a chosen design axiom** and the operational ladder recorded as **specified, not run** — is Curzi [2026].

The framework's own contraction result carries its own refusal, and it is worth restating at the close. Banach's theorem states that a contraction with modulus `q < 1` on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least `q` each pass [Curzi, 2026]. That is a theorem about a map. **A human life is not a metric space, nobody has measured the relevant contraction modulus, and the feeling of recognition on arrival proves nothing**, since the machinery that rehearses a future retrofits a past with equal fluency. The theorem transfers as shape, not as guarantee, and the closing paper is the last place that distinction can be allowed to slip.

The three-drawer honesty classification this appendix set descends from — four items of established mathematics, seven of established psychology, and exactly three declared postulates, with every protocol stated to run on the first two and steer by the third — is Curzi [2026], and it is the model for Appendix E.

39.8.3 8.3 Rules for carrying the apparatus forward

S1 Read CCP-038 before citing anything in this corpus. Cite by entry ID, because the ID carries the status and a quoted sentence does not.

S2 State the status of every claim carried forward from this framework, in the sentence that carries it. Proved, badged, or open — there is no fourth option and no exemption for summaries.

S3 Measure `h` before promising `r`. A plan requesting more than the gauge permits is not ambitious; it is arithmetically infeasible, and the equation says so before anyone is blamed.

S4 When `h < 0`, do not redesign the interaction geometry. **No geometry saves it.** Raise `eta`, cut `C`, or shed `Q` — the last being the move nobody considers and a legitimate one.

S5 Test feasibility on the mean of `F(U_t)`, never on `F` of the mean. The Jensen gap was worth 6.2% of surplus in the CCP-016 worked office.

S6 Run the structural ceiling test on every result, before and after the statistical baselines. A single-step increment exceeding `eta ln N` is impossible, and impossibility outranks significance.

S7 Quote every `delta` with its period and every `F` in nats. Convert with the exact keys, never with an approximation.

S8 Never present precedent as evidence. Where the historical corpus has nothing to say, write *none claimed* and leave it.

S9 Recommend no intervention that removes a person's ability to decline, and attach consent to the person rather than to the scope. Disclose everything: **nothing masked, nothing below threshold, nothing concealed.**

S10 Publish the demotions. If nothing in a body of work has ever failed a gate, the gates are not connected to anything.

S11 Report null results from the tests this corpus names. Appendix E.4 lists ten of them, all of them cheap, none of them run. **A reader who runs one and reports a null has done this corpus a service it cannot repay.**

S12 Re-derive the reference card and the register annually, by reading the corpus papers. Both are stored order with their own `delta`, and a card that has become wrong looks exactly

like one that is still right.

39.8.4 8.4 Faults of the closing summary itself

Fault	Signature	Mechanism	Fix
THE SLOGAN	"It's all the same equation at every scale," quoted without its status	The correspondence is memorable and its badge is not	S2 and CCP-038. It is a claim about model form, never tested at two scales at once
BADGE EROSION	A postulate cited as a finding, three cross-references downstream	Each hop drops one qualifier	S1. Cite the ID
THE MISSING LOSS TERM	A plan that assumes stored order persists without maintenance	`delta` is invisible until it is measured, and infrastructure becomes visible only on breakdown	Measure `delta` under zero input. Assume nothing holds by itself
HEADROOM THEATRE	`h` reported as a target and rising every quarter	A measure used for control ceases to measure. This framework is subject to Goodhart like everything else	Measure `h` where parameters are estimated, not where decisions are announced
THE GENEROUS SUMMARY	A closing chapter that claims more than its register contains	Endings invite generosity, and the reader is tired	This paper is bounded by CCP-038 and claims nothing that is not in Groups A and B
BORROWED CERTAINTY	A laboratory result deployed as a fact about the reader's institution	The citation stays attached while the setting silently changes	Read the "what it measured" field before quoting the finding
THE ADDED MECHANISM	A supernatural, energetic, or clinical claim attached to the formalism to make it more useful	The formalism is austere and austerity feels like incompleteness	The formalism is sufficient for what it claims. Anything added is outside the register and outside this corpus
CLOSING ON A TRIUMPH	A final chapter that resolves rather than opens	A book wants an ending, and an open problem does not feel like one	§8.1. The work ends on what is not known

---

39.9 9 Verification

The four structural results compound from five postulates. Each is atomic, checkable on its own, and carries its status. **A reader is asked to check several short things rather than one**

long one.

39.9.1 9.1 Dimensional audit

Symbol	Dimension	Unit	Status
`u`, `u*`, `u_0`	`[1]`	none	**defined** dimensionless: `u = 1 - (x.y)^2` for unit vectors
`N`	`[1]`	categories	**counted**
`N_eff`	`[1]`	effective count	**derived**: `1 + (N-1)u`, affine, **not a Hill number**
`F(u)`, `ln N`, `ln 2`	`[1]`	nats	**derived** dimensionless; convertible to bits and octaves by exact constants
`Q_t`	`[Q]`	proxy units (whatever the declared order proxy counts)	**proxied, never measured**
`eta`	`[Q]` per nat per step	proxy units per nat per period	**estimator defined, never validated**
`delta`	`[1]`	per period — the period is part of the number	**estimator defined**; validated once, borrowed [Curzi, 2026]
`C_t`, `c*`	`[Q]` per step	proxy units per period	**estimator defined**; rests on a commensuration postulate
`h_t`, `h_inf`, `r_max`	`[Q]` per step	proxy units per period	**derived**
`n` (essence index)	`[1]`	none	**defined** ordinal index
`k`	`[1]`, or `[T^-1]` **only under a separate postulate**	dimensionless as defined; **hertz is a separate claim**	**DEFINED** = `(9/8)^2/phi^2`, exact
`f(n) = k n^2`	as `k`	as `k`	**POSTULATED** square law; the frequency reading is a further postulate
likelihood ratio	`[1]`	none	**derived** from a stipulated selection process
survival probability	`[1]`	none	**derived** from badged stage priors

Do the identities balance? `F(u) = ln((u+u_0)/u_0)` balances because `u` and `u_0` carry the same dimension — none — so their ratio is a pure number and its logarithm is in nats. `h_inf = eta(ln N - F(u*))` balances because the bracket is dimensionless and `eta` carries `[Q]` per nat

per step, so the product is `[Q]` per step, commensurable with `delta Q` and `C`. `ln N − ln(N/2) = ln 2` is dimensionless throughout.

Where the framework has no operational unit, this paper says so. Stored order has no natural unit and two systems' values are not comparable unless they declare the same proxy. Comfort, the corpus's underlying construct, has **no operational unit at all**. Perspective and perception have **no unit and no zero**, which is why their quotient is demoted. **`k` is dimensionless as defined; the corpus's essence material attaches hertz to it, and that attachment is a postulate carried on its own line below rather than folded into the definition.**

39.9.2 9.2 The postulates

P1 — `F(u) = ln(1 + (N−1)u)` on `u` in `[0,1]`, with `0 <= F(u) <= ln N`. *Status:* **CITED** — the surplus functional and its range [Curzi, 2026].

P2 — The storage recurrence is `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t` with `delta` in `(0,1]`, and its fixed point is `Q_inf = (eta F(u*) − c*)/delta`. *Status:* **CITED** — the storage law and its fixed-point theorem [Curzi, 2026a]; the fixed point exists because the map is a contraction [Banach, 1922].

P3 — `ln(ab) = ln a + ln b` and `ln(a/b) = ln a − ln b`. *Status:* **ARITHMETIC.**

P4 — The essence law is `f(n) = k n^2` with `k = (9/8)^2/phi^2`. *Status:* **POSTULATED** for the square law; **DEFINITION** for the constant, which is exact and equals `0.483425732988414`.

P5 — `k` is measured in hertz. *Status:* **POSTULATED** — a separate and stronger claim than computing `k`, carried on its own line because it is a different claim, and supported by nothing in this corpus. Result (b) below does **not** consume it: the octave-folding cost is a ratio and is independent of the unit `k` carries.

39.9.3 9.3 The compounding chain

```
P1 + P3     -> V1   ln(1 + (N−1)u) = ln((N−1)(u + 1/(N−1)))
                    = ln((u + u_0)/u_0)  with  u_0 = 1/(N−1).       ARITHMETIC

V1          -> V2   at u = u_0: N_eff = 1 + (N−1)/(N−1) = 2 exactly,
                    and F(u_0) = ln 2 = 0.693147..., for EVERY N.   ARITHMETIC

P4 + P3     -> V3   f(2n)/f(n) = (2n)^2/n^2 = 4 exactly, so indices n
                    and 2n share a pitch class; 1..N contains exactly
                    ceil(N/2) classes.                              ARITHMETIC

V3 + P3     -> V4   folding an even-N lattice: ln N − ln(N/2) = ln 2
                    = 0.6931471806 nats = 1 bit = 2 octaves, for EVERY
                    even N.                                         ARITHMETIC

P2          -> V5   at the fixed point delta Q_inf = eta F(u*) − c*, so
                    h_inf = eta ln N − delta Q_inf − c*
                          = eta (ln N − F(u*)).                     ARITHMETIC

V5 + P1     -> V6   F(u*) <= ln N, therefore h_inf >= 0 ALWAYS.     ARITHMETIC

P2          -> V7   r_max = eta ln N − delta Q − C <= eta ln N, since
                    delta, C >= 0; a claimed single-step increment of
                    38.5 x eta ln N is impossible under the model
                    at any parameter setting.                       ARITHMETIC

P5          -> V8   f(n) in hertz.                                  POSTULATED
                    NOT CONSUMED by V1..V7. Listed so that a reader can
```

see exactly which line the frequency claim enters on.

Every line is checkable from its own stated inputs. V1 through V7 are arithmetic consequences of two cited results (P1, P2), one arithmetic fact (P3) and one postulated functional form (P4), and V4 is the only one that consumes P4. **V8 is listed and consumed by nothing**, which is the point of listing it: the frequency reading of `k` is a claim this paper neither uses nor supports.

What the chain does not establish. It does not establish that any institution, mind or cell obeys P2, or that any essence set obeys P4. **A chain whose steps are ARITHMETIC and CITED delivers the mathematics with certainty and delivers nothing whatever about any system**, and the one preregistered confirmatory test of P2 reported in this corpus did not meet its own declared margin [Curzi, 2026i].

39.9.4 9.4 The verification instance, worked by hand

Check 1 — the Fechner identity at `N = 6`. `u_0 = 1/5 = 0.200000`. At `u = 0.5`: `ln(1 + 5 x 0.5) = ln 3.5 = 1.252763`, and `ln((0.5 + 0.2)/0.2) = ln 3.5 = 1.252763`. **Two logarithms of the same number.** At `u = u_0 = 0.2`: `ln(1 + 5 x 0.2) = ln 2 = 0.693147`, and `N_eff = 1 + 5 x 0.2 = 2` exactly. **The threshold is the point at which the effective count reaches two, and it is one full bit wide, at every `N`.**

Check 2 — the threshold's size across `N`. `u_0 = 1.000000` at `N = 2`; `0.500000` at `N = 3`; `0.333333` at `N = 4`; `0.200000` at `N = 6`; `0.090909` at `N = 12`. **In a two-category system the second voice arrives only at perfect orthogonality; in a twelve-category system it arrives at nine per cent of the range**, which is a precise statement of why narrow systems yield so little.

Check 3 — the folding cost. `ln 12 − ln 6 = 2.484907 − 1.791759 = 0.693147`. `ln 64 − ln 32 = 4.158883 − 3.465736 = 0.693147`. Ten values of `N` return `0.6931471806` and a ratio to `ln 2` of `1.0000000000`. **The concrete case: a twelve-essence set has indices with odd parts `1,1,3,1,5,3,7,1,9,5,11,3`, giving six pitch classes.**

Check 4 — the constant, to fifteen figures. `(9/8)^2 = 1.265625` exactly. `phi^2 = phi + 1 = 2.618033988749895`. Dividing, `1.265625 / 2.618033988749895 = 0.483425732988414`. **Three checks anyone can redo in a minute, and the constant is then established as defined rather than asserted** — with the standing caution that a golden-section appearance in a defined constant is a design choice of this corpus and not an observation about nature [Markowsky, 1992].

Check 5 — the headroom identity, both ways. At `eta = 0.35`, `N = 12`, `u* = 0.30`, `delta = 0.05`, `c* = 0.40`: `F(0.30) = ln(1 + 11 x 0.30) = ln 4.3 = 1.458615`, so the identity gives `0.35 x (2.484907 − 1.458615) = 0.359202`. The long way: `Q_inf = (0.35 x 1.458615 − 0.40)/0.05 = 2.210305`, then `h_inf = 0.35 x 2.484907 − 0.05 x 2.210305 − 0.40 = 0.359202`. **The same number to ten decimals, and the second computation used three quantities the first did not need.**

Check 6 — the impossibility. `r_max = eta ln N − delta Q − C` is at most `eta ln N`, since `delta` and `C` are non-negative. A claimed single-step jump of `38.5 x eta ln N` is therefore inadmissible under the model **at any parameter setting, any significance level, and any sample size.** That is not a p-value; it is a contradiction.

39.9.5 9.5 Limits, asymptotics, and the domain of the results

- **`u -> 0`.** `N_eff -> 1`, `F -> 0`. Ten aligned instruments are one instrument; the Fechner form gives `ln(u_0/u_0) = 0` at `u = 0`, consistently.
- **`u -> 1`.** `N_eff -> N`, `F -> ln N`, and `h_inf -> 0`. **A system extracting everything its geometry allows retains the least headroom at the fixed point**, which is the design reading of result (c).
- **`N -> 2`.** `u_0 = 1`, so **the threshold is the ceiling**: in a two-category system the second independent voice arrives only at perfect orthogonality. This is a real constraint on the domain of usefulness, not a defect.

- **`N -> infinity.`** `u_0 -> 0`, so the threshold vanishes and any non-parallelism at all buys a second voice; meanwhile `ln N` grows logarithmically while coordination paths grow as `N(N-1)/2` [Brooks, 1975]. The two together bound the useful roster well below any crossover computed from surplus alone.
- **`N` odd, for result (b).** The folding cost is exactly `ln 2` **for even `N` only**; `1..N` contains `ceil(N/2)` classes, so an odd lattice folds to `(N+1)/2` and the loss is `ln N - ln((N+1)/2)`, which is **less** than `ln 2` and is not constant. **The result is stated for even `N` and holds for even `N`.**
- **`delta -> 0`.** No fixed point exists at all rather than being infinite: the map is an isometry and the geometric series does not converge [Curzi, 2026a]. **Result (c) has no content at `delta = 0`**, since it is an identity about a fixed point that does not exist there. Outside the domain.
- **`eta -> 0`.** `h_inf -> 0` and the system cannot build at any `u` or `N` — the storage law's correspondent of the forbidding clause `G = 0 => E = 0` in the channel-gain form of the belief postulate [Curzi, 2026ab].
- **Exogenous accretion.** With a term `A` for stored order arriving from outside the interaction loop, `h_inf = eta(ln N - F(u*)) - A`, which goes negative exactly when `A` exceeds the geometric shortfall. **This is where ossification lives, and the plain model does not contain it.**

39.9.6 9.6 Model comparison, and the standard this paper does not meet

The paper reports no fit of its own. It inherits the standard because the apparatus it collects is what a fit would be run with.

Requirement	Status	What was actually done
Out-of-sample performance reported	**MET**	held-out RMSE against four baselines [Curzi, 2026i]
Matched null with equal free parameters	**MET**	a matched AR(1); the fit beat it by 2.8 per cent against a declared 10 per cent margin and **failed its own preregistered rule**
Information criterion (AIC or BIC)	**NOT MET**	**no AIC or BIC has been computed anywhere in this corpus**, so no comparison in the series pays for its parameter count
Number of free parameters stated explicitly	**PARTIAL**	the storage law's four parameters are named; they are not tabulated against each baseline's count
Effect sizes with intervals	**NOT MET**	RMSE advantages are reported as point estimates without intervals
Shuffle control's false-positive rate measured	**MET, AND ADVERSE**	0.108 for i.i.d. permutation and 0.120 for circular shift against a nominal 0.05 [Curzi, 2026u] — **anti-conservative**
Structural ceiling test run alongside the empirical baselines	**MET, AND DECISIVE**	it is the only test that caught the repackaging (result (d))

Requirement	Status	What was actually done
Non-affine `g(u)` fitted and compared on held-out data	**NOT MET**	the corpus's register names the test and records that it has not been run [Curzi, 2026]; it is not expensive
Saturating or state-dependent storage law compared	**NOT MET**	same register, same status

Four NOT MET rows are printed rather than omitted. The last two are the most damaging: they are the tests that would decide whether the corpus's two founding design choices are the right ones, and neither requires new data.

39.9.7 9.7 Scale type, and the objection the corpus concedes

Quantity	Scale type	What licenses it	Operations performed here
`N`, `N_eff`, pitch-class count	ratio (counts)	counting	all
`u`, `u_0`	ordinal at best, treated as ratio	**nothing tested**	used inside a logarithm; **not licensed**
`F`, `ln N`, `ln 2`	ratio (a log of a count)	derived from a count	differences, ratios, conversion to bits and octaves
`Q`	interval at best in a declared proxy	proxy definition only	differences and rates within one proxy
`D`, `P`, `M`	ordinal	no zero, no unit	ordering only; the quotient **demoted** [Curzi, 2026t]
the essence index `n`	ordinal, treated as ratio in `k n^2`	**postulated, not tested**	squared; **the square law is the postulate that licenses it, and it is a postulate**
likelihood ratios	ratio	a ratio of probabilities	division

Michell's critique is the strongest single objection to this corpus and it is conceded here. Michell [1997, 1999] argues that psychology assumes its attributes are quantitative without testing that they are. The corpus's founding move — treating comfort as a magnitude — is exactly that assumption. The test that would settle it is conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971], whose double-cancellation conditions determine whether an attribute admits additive quantitative structure at all. **Those conditions have not been tested for comfort, for `u`, for perspective or for perception, by this work or by anyone.** The corpus's response is partial and is stated as partial: it denotes the quantity whose operation most obviously requires a ratio scale, and prices the Rasch route at 123 items per term [Rasch, 1960, Andrich, 1988] — an instrument with a defensible zero, which is **not** a conjoint test. Reliability would not substitute [Sijtsma, 2009], and construct validation [Cronbach and Meehl, 1955, Campbell and Fiske, 1959] would establish that an instrument measures something without establishing that the something is quantitative. **The scale type of the essence index is the least defended item in the table and the register records it as postulated.**

39.9.8 9.8 Invariance, and the precision of the numbers

What is invariant. The Fechner identity holds for every `N >= 2` and every `u` in `[0,1]`; the threshold value `F(u_0) = ln 2` is invariant in `N` entirely, which is the surprising part. The folding cost `ln 2` is invariant in `N` over even `N`. The identity `h_inf = eta(ln N − F(u*))` is invariant under any change of `delta`, `C` or `Q`, which is the content of the cancellation, and **is not invariant under adding an exogenous accretion term.** The impossibility of a `38.5 x eta ln N` jump is invariant under every parameter setting, which is what makes it a structural rather than a statistical test.

What is not invariant, and is not claimed to be. Nothing here is invariant across scopes; **scale-freedom is claimed at three scopes and has been tested at zero**, because testing it requires simultaneous measurement at two scopes of one system and no such measurement exists [Curzi, 2026].

Precision. Every number in this paper is one of three kinds. *Exact consequences of definitions* — `ln 2 = 0.6931471806`, `k = 0.483425732988414`, `h_inf = 0.3592020695` — are quoted to whatever precision is useful and are exact; the fifteen figures on `k` are legitimate **only because `k` is defined**, and a measured quantity quoted to fifteen figures would be the classic marker of work that does not understand its own error bars. *Numerical verifications* — the `8.9e−16` residual on the Fechner identity and the sub-`1e−16` agreement on the headroom identity — are floating-point noise and are reported as such rather than as agreement to a physical tolerance. *Outputs of stipulated models* — the survival curve, the likelihood ratios — are exact given inputs the author chose and badged as chosen, and **carry no empirical content at all. Where a reader sees ten decimals in this paper they should read "consequence of a definition or a stipulation", never "measured to that precision".**

39.10 10 Limitations

L-1 This paper is not a research contribution and does not present itself as one. It is closing material plus reference apparatus, most of its length recapitulates results established elsewhere, and its non-recapitulative content is four algebraic identities and a set of computed bookkeeping items. **It is retained in the series because other papers rely on its appendices**, not because it adds knowledge.

L-2 None of the four structural results is empirical, and none could be falsified by an observation. (a), (b) and (c) are identities; (d) is a structural argument about what a control can detect, together with one reported instance. **A reader looking for a test of the framework will not find one here** and should read Curzi [2026i] and Curzi [2026u], where the tests that were run are reported, including the one that failed its own declared margin.

L-3 Result (b) rests on a postulate, and the postulate is the load-bearing part. The octave-folding cost is exact **inside** the lattice `f(n) = k n^2`, and that square law is a postulate about essence structure with no external support. "Consonance" there means the exact-ratio octave identification internal to the lattice; **it is not a claim about the psychoacoustics of consonance, which is a different subject with a contested literature.** And the output remains an essence frequency: **no claim is made that any material, molecule, element or person physically resonates at these values.**

L-4 Attaching hertz to `k` is a separate claim from computing `k`, and it is unsupported. The constant is exact as a ratio; the unit is a postulate, carried as V8 in the chain and consumed by nothing in this paper.

L-5 Result (b) holds for even `N` only (§9.5). For odd `N` the folding loss is `ln N − ln((N+1)/2)`, which is smaller than `ln 2` and varies with `N`. The source states the result for even `N` and this paper preserves that scope.

L-6 The reflexive parameter is not identified from observational data [Manski, 1993, Shalizi and Thomas, 2011], and no remedy is adopted anywhere in the corpus (§3.4). Every protocol in Appendix C returns a number; **none of them identifies a social coupling.**

L-7 The conjoint conditions have not been tested for any construct in this corpus [Michell, 1997, Luce and Tukey, 1964], and the ratio-scale treatment of `u` and of the essence index is unlicensed.

L-8 No AIC or BIC has been computed anywhere in this corpus, and the two tests that would decide its founding design choices — a non-affine effective count and a saturating storage law — have not been run (§9.6).

L-9 The survival curve is arithmetic on the author's own priors. The stage probabilities 0.90, 0.90, 0.85, 0.65 and 0.80 were chosen and badged as chosen. **A different set gives a different curve, and the 35.8 per cent figure should be read as "what these priors imply", not as a measured base rate.** What is not arbitrary is that the calculation was published at all.

L-10 The likelihood-ratio results are arithmetic on a stipulated selection process. The testimony ratio of exactly 1.000000 follows from a process in which failures are never published; where failures are published at some rate, the ratio is not one. **The result is a statement about a selection regime, not a measurement of any actual literature**, and the four hypothesised success rates producing 500 identical published accounts are stipulated for the calculation.

L-11 A numerical inconsistency in the source is recorded rather than silently corrected. The closing chapter states the worst residual on the Fechner identity as `4.4e-16`; the honesty appendix states it as `8.9e-16` over `u = 0.01` to `1.00` at `N = 2, 3, 5, 8, 12, 25, 100`. **Recomputing the identity over both grids returns a maximum absolute difference of `8.88e-16`, so the appendix figure is the correct one and the closing chapter's is too small.** This paper quotes `8.9e-16` throughout and prints the discrepancy here.

L-12 The mapping from the source manuscript's chapter numbers and internal papers to corpus identifiers is an editorial inference (§3.6). In the appendices those identifiers are plain text rather than bibliographic citations, because the appendices are an index rather than an argument.

L-13 The register in Appendix E classifies the author's own claims. It is not an external audit, and the only defence offered is that every entry carries the fields a hostile reader needs in order to check it, and that the demotions are printed in the same typeface as everything else.

L-14 Several works are named in prose without bibliographic keys, because they are not in this corpus's external bibliography. They include the sources named in Appendices B and E for information cascades and Goodhart capture where those entries name authors and years without a key, the historical and liturgical corpora invoked in the precedent column, and the primary sources of the cell-biology chain, whose treatment in this corpus is Curzi [2026y]. **Each is a real published result quoted for what its source measured; the reader is owed the note that they are named rather than referenced.**

L-15 The appendices carry the corpus's own decay. A reference card is itself a stored order with its own retention loss, and the source's own specification requires it to be re-derived from the papers annually by reading rather than by editing. **Nothing in this paper verifies that the appendices currently match the papers they index**, and where an entry and a paper disagree, the paper is authoritative.

39.11 11 Conclusion

This is closing material and reference apparatus, and its four non-recapitulative results are exact identities rather than empirical findings.

First, the surplus law is exactly a Fechner law in threshold-corrected form. `F(u) = ln(1 + (N−1)u) = ln((u + u_0)/u_0)` with `u_0 = 1/(N−1)`, verified for `N = 2` through `N = 100` at seven values of `u` to a worst residual of **8.9e-16**. At the threshold the effective count is **exactly 2** and the surplus is **exactly `ln 2` — one bit, two octaves — for every `N`**. The threshold is the non-parallelism at which the second independent voice opens: `u_0 = 1.000000` at `N = 2`, so in a two-category system the second voice arrives only at perfect orthogonality, and `u_0 = 0.090909` at `N = 12`. **The correspondence is structural and is evidence for neither law:** the surplus functional

is a theorem under its axioms and gains nothing from resembling Fechner's, and Fechner's is an empirical claim resting on its own postulate and gains nothing from resembling the surplus functional.

Second, folding the essence lattice into one octave costs exactly `ln 2 = 0.6931471806` nats for every even `N`, computed at `N = 4, 6, 8, 10, 12, 16, 20, 24, 32, 64`, which is one bit and two octaves by the corpus's own exact conversion key. In that lattice **calling two things consonant and calling them redundant are the same operation**, and every identification is a reduction of effective count. The result is exact inside a postulated structure and claims nothing about psychoacoustics or about physical resonance.

Third, `h_inf = eta(ln N − F(u*))`: retention, cost and stored order all cancel at the fixed point, verified against direct computation across four parameter sets to below `1e−16`. Since `F(u) <= ln N` always, **`h_inf >= 0` always, so the plain linear model never ossifies by itself.** Ossification therefore requires a term the model does not contain — exogenous accretion of stored order, or cost rising with stock — and **a model that cannot produce a phenomenon everyone observes is reporting that the mechanism is outside the model**, which converts a lament into a question with two named suspects.

Fourth, the shuffled-input control is blind to repackaging. A filtered repackaging of the input passed all four empirical baselines including the shuffle at **`p = 0.0005`** and was caught only by the structural ceiling `r_max = eta ln N − delta Q − C`, its claimed single-step jump being **38.5 times `eta ln N`. Impossibility is not a p-value; it is a contradiction**, and no parameter setting, significance level or sample size admits it. The rule that follows is that the shuffle is an empirical audit and only the structural tests detect repackaging — reinforced by the shuffle's own measured false-positive rate of **0.108 against a nominal 0.05**.

The appendices publish two admissions. Testimonial evidence under a selection process that publishes successes has a likelihood ratio of **exactly 1.000000** — a thousand favourable testimonials carry the evidential weight of zero — against **2.0000** for a direction-only prediction and **11.3577** for a magnitude call with a tolerance, a factor of about **5.7**. And the research programme that would settle whether the framework's formulas are science publishes its own odds: **21 preregistered kill conditions of which 4 are terminal, 8,554.8 researcher hours, 570 participants, 71 months, P(completion) = 35.8 per cent, and P(death at or before the cross-subject recovery stage) = 55.3 per cent.** The programme is more likely than not to be dead before field trials begin, and the stage carrying the heaviest prior is the one where a formalism that *describes* meets the demand that it *measure*.

Read together with the instrument audit — 19 defined estimators, 1 validated value and 0 first-party measurements across 27 parameter-by-scope cells [Curzi, 2026] — that is the true state of the instrument at the close of the series: nineteen procedures, one borrowed number, no measurement of its own, a keystone formula demoted by its own protocol, and a thirty-five per cent chance of finishing. A reader who finds that discouraging has read it correctly.

39.11.1 Declarations

Data and code availability

No external dataset was collected for this paper, and no human participants were surveyed, observed or recruited. All numerical results are computed from parameters and definitions stated in the text, and every table can be re-derived from the equations printed beside it. The worked instance in Section 7.1 is a **constructed illustrative example with declared parameters**, not a report of any real system. The single borrowed empirical value that appears anywhere in the supporting audit belongs to its published source and is cited to it in Curzi [2026], where the corpus-wide count is also established: **of 27 parameter-by-scope cells, zero carry a first-party measurement.** The survival curve and the likelihood ratios are arithmetic on stipulated inputs and badged priors, not observed base rates. The figures are generated by the scripts named in each caption, which contain no data beyond the numbers printed in this paper. The appendices are compiled by reading the corpus's own papers.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest. **The honesty register in Appendix E classifies the author's own claims, which is a structural conflict rather than a financial one, and it is disclosed as Limitation L-13.**

Author contributions

Single author. Michael Laurence Curzi derived the results, compiled the reference apparatus and the register, produced the figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. No biological materials were handled and no experiment of any kind was performed. **No clinical, medical or therapeutic claim is made anywhere in this paper**; nothing here diagnoses a person or recommends a treatment; and **no claim is made that any material, molecule, element or person physically resonates at any frequency computed from the essence law.** The research programme costed in Appendix E.11 is a **prospective design**, with its consent architecture, its decline path and its preregistered kill conditions specified in advance; it has not been run, and nothing in this paper reports a result from it.

Preregistration

This paper reports no preregistered test of its own; its results are identities and bookkeeping. The preregistered protocol for testing the storage law against a real system, and the confirmatory application that **failed** its second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin — are reported in Curzi [2026i]. The falsification suite measured against itself, including the shuffled-input control's **anti-conservative false-positive rate of 0.108 against a nominal 0.05**, is Curzi [2026u]. The programme in Appendix E.11 publishes **21 kill conditions, 4 of them terminal, in advance of any data**, together with the commitment to publish at every stage including the stage that ends the programme, with null results titled as null.

The body of this framework argued, specified, and worked examples. The appendices do none of those things. They are the bench: the card that gets pinned above the desk, the vocabulary an engineer must be able to use without hedging, the ten protocols that turn an opinion about a system into a number — and, in one case, that take a protocol back off the reader, the index that lets a reader who arrives with a symptom leave with a diagnosis, the statement of exactly what this corpus is entitled to claim, and the bridge back to the volume that preceded it.

A caution before the drawer is opened. A reference card is itself a stored order — a small `Q`, with its own `delta` — and it decays silently, because a card that has become wrong looks exactly like one that is still right. **Re-derive it from the corpus papers once a year, by hand.** A card corrected by its user is worth two printed ones.

Appendix E is the warrant, and is the one appendix to read if only one is read. Everything else in this framework is a consequence of a small number of proved theorems, a slightly larger number of badged postulates, and a set of borrowed empirical results that were obtained in settings that are not yours. Appendix E states which is which, item by item, and it is written to be used against this corpus by anyone who wishes to test it.

And it now states two further things, which are the two worth reading first. It records what this corpus **demoted** — including its author's own keystone formula, on gate two of seven, with the cost of the repair named — and what this corpus **rejected**, four claims worked end to end in print and failed. **Nothing in an appendix is more load-bearing than a list of the things its author had to give up.**

39.12 Appendix A The law reference card

Three pages. Print at 100%. Pin at eye level beside the desk where the parameters are estimated, not where the decisions are announced.

39.12.1 A.1 Symbol table

Symbol	Name	Definition	Range	Units	Where it is measured
`x, y`	state vectors	unit-normalized encodings of two participating states (views, agents, categories)	unit sphere	dimensionless	Appendix C, Protocol C-4
`x·y`	alignment	inner product of the two unit vectors	`[−1, 1]`	dimensionless	C-4
`u`	non-parallelism	`u = 1 − (x·y)^2 = ||x ^ y||^2`	`[0, 1]`	dimensionless	C-4
`N`	category count	number of genuinely distinct categories admitted by the system	integer `>= 1`	count	C-3
`N_eff`	effective count	`1 + (N−1)u`; the audited headcount	`[1, N]`	count	C-3, C-4
`f(u)`	surplus function	`ln(1 + (N−1)u)`	`[0, ln N]`	nats	derived
`F(x,y)`	surplus	`f(u(x,y))`	`[0, ln N]`	nats	derived
`S_t`	surplus input in period `t`	`S_t = F(U_t) = ln(1 + (N−1)U_t)`	`[0, ln N]`	nats/period	C-4 then compute
`U_t`	period non-parallelism	the ensemble `u` prevailing in period `t`	`[0, 1]`	dimensionless	C-4
`Q_t`	stored order	what persists between periods: institution, habit, artifact, furnished room, trust	`>= 0`	Q-units	C-6
`Y_t`	order proxy	the observable standing in for `Q_t`	`>= 0`	Q-units	C-6
`delta` (δ)	retention loss	fraction of `Q` lost per period with zero input	`[0, 1]`	per period	C-1
`eta` (η)	conversion efficiency	`Q`-units of durable structure produced per nat of surplus	`>= 0`	Q-units per nat	C-2
`C_t`	cost / dissipation	coordination, friction, maintenance, conflict, rent, compliance	`>= 0`	Q-units/period	C-5
`h_t`	headroom	`eta*ln N − delta*Q_t − C_t`	any sign	Q-units/period	C-7

Symbol	Name	Definition	Range	Units	Where it is measured
`r_t`	requested increment	the growth in `Q` a plan promises for period `t`	any sign	Q-units/period	C-7
`r_max`	maximum increment	`eta*ln N − delta*Q_t − C_t`	any sign	Q-units/period	C-7
`Q_inf`	steady state	`(eta*F(u*) − c*) / delta`	any sign	Q-units	C-7
`tau` (τ)	time constant	`−1 / ln(1−delta)`, approximately `1/delta` for small `delta`	`> 0`	periods	CCP-007
`u_cross`	crossing	between-category non-parallelism, `= beta^2`	`[0, 1]`	dimensionless	CCP-010, C-4
`u_div`	divergence	within-category non-parallelism, `= alpha^2 * sin^2(gamma)`	`[0, 1]`	dimensionless	CCP-010, C-4
`c`	collective intelligence factor	Woolley et al.'s statistical factor; **not** a term in the law	standardized	z-score	CCP-013

Units convention. `S` and `F` are in **nats**, because the law uses the natural logarithm. `Q` is in whatever unit the order observable `Y` is denominated in, chosen once in Protocol C-6 and never changed mid-study. `eta` therefore carries units of Q-units per nat, and `C` and `delta*Q` must be expressed in Q-units per period or the headroom equation is not an equation. **Three quantities must share one unit or the gauge reads nothing: `eta*ln N`, `delta*Q_t`, and `C_t`.**

39.12.2 A.2 The Surplus Law — CCP-A02, Theorem 2.1

```
F(x,y) = f(u(x,y)),   u(x,y) = 1 − (x·y)^2 = ||x ^ y||^2  in [0,1]

f(u) = ln(1 + (N−1)u)        <−− UNIQUELY characterized by axioms S1−S4

N_eff = g(u) = 1 + (N−1)u    <−− the effective count
F     = ln N_eff             <−− surplus is the log of an audited headcount
```

The axioms.

Axiom	Statement	Status
S1	geometric dependence: `F` depends only on `u`	structural choice
S2	zero at zero: `f(0) = 0`	normalization
S3	affine effective count: `g(u) = e^{f(u)}` is affine in `u`	**DESIGN AXIOM — chosen, not derived**
S4	normalization: `f(1) = ln N`	normalization

S3 is the design axiom. It is chosen, not derived, and that must be said wherever it is leaned on. Every quantitative claim in this framework that uses `N_eff = 1 + (N−1)u` inherits S3's status and not a grade higher.

39.12.3 A.3 Derived properties — NOT axioms

Result	Statement	Reference
Strict monotonicity	`u1 < u2 => f(u1) < f(u2)`	CCP-A02, 3.1
Positivity	`f(u) > 0` for `u > 0`	CCP-A02, 3.2
Strict concavity	`f''(u) < 0`; each added increment of `u` buys less than the last	CCP-A02, 3.3
Subadditivity **on u**	`f(u1 + u2) <= f(u1) + f(u2)`	CCP-A02, 3.4
Sharp Lipschitz	`|f(u1) − f(u2)| <= (N−1)|u1 − u2|`, constant `N−1` is sharp	CCP-A02, 4.2
Extremality	`F = 0` iff parallel; `F = ln N` iff orthogonal	CCP-A02, 4.3
Range	`0 <= F(u) <= ln N`	CCP-009, Cor 2.2

`ln N` is the hard ceiling on surplus per interaction. It appears in every design inequality in this framework through the product `eta*ln N`.

39.12.4 A.4 Two-source decomposition — CCP-A02, Theorem 4.1

For a pure block vector:

```
u = u_cross + u_div

u_cross = beta^2                      (between-category: crossing)
u_div   = alpha^2 * sin^2(gamma)      (within-category: divergence)

max( f(u_cross), f(u_div) )  <=  f(u)  <=  f(u_cross) + f(u_div)
```

The upper bound is the honest one: **no more than the sum is ever obtained.** No synergy beyond additivity is promised, and design that assumes it is designing on a number the theorem does not supply.

39.12.5 A.5 The Storage Law — CCP-007, §3

```
Q_{t+1} = (1 − delta)*Q_t + eta*S_t − C_t          S_t = F(U_t) = ln(1 + (N−1)U_t)
```

Term	Meaning
`Q_t`	stored order — what persists: institution, habit, furnished room, trust
`S_t`	surplus input this period, `0 <= S_t <= ln N`
`delta`	retention loss in `[0,1]` — decay without maintenance
`eta`	conversion efficiency `>= 0` — interaction into DURABLE structure
`C_t`	cost / dissipation `>= 0`

Admissibility. `C_t <= (1−delta)*Q_t + eta*S_t` guarantees `Q_{t+1} >= 0`.

Exact solution — Theorem 3.1.

```
Q_t = (1−delta)^t * Q_0  +  SUM_{k=0}^{t−1} (1−delta)^{t−1−k} * (eta*S_k − C_k)
      \_____________/      \_____________________________________________/
        transient term                    driven term
```

One-step growth — Corollary 3.2. `Q_{t+1} > Q_t  <=>  eta*S_t > delta*Q_t + C_t`.

Uniform bound — Theorem 3.3. `Q_t <= max{ Q_0, eta*ln N / delta }`. If `delta = 0`, `Q_t <= Q_0 + t*eta*ln N` — unbounded linear growth, no fixed point, no arrival.

Sustainability ceiling — Corollary 3.4. A nondecreasing step requires `C_t <= eta*ln N - delta*Q_t`. If `C_t > eta*ln N - delta*Q_t`, **no choice of interaction geometry can prevent order from decreasing.**

Steady state — Theorem 3.5. With constant `U \ensuremath{\equiv} u*`, `C \ensuremath{\equiv} c*`, and `delta > 0`:

```
Q_inf = (eta*F(u*) - c*) / delta
```

If `eta*F(u*) < c*` the equilibrium is negative: the regime is unsustainable, or the linear model must be replaced. **The theorem does not hide this failure mode.**

Target increment — Proposition 4.1. `Q_{t+1} - Q_t >= r_t` iff `S_t >= (delta*Q_t + C_t + r_t)/eta`, feasible only if `delta*Q_t + C_t + r_t <= eta*ln N`.

Maximum one-step increment — Corollary 4.2. `r_max = eta*ln N - delta*Q_t - C_t`. No policy based on interaction geometry alone can guarantee more.

39.12.6 A.6 Headroom — the master gauge, CCP-007 §6.4

```
h_t := eta*ln N - delta*Q_t - C_t

h_t < 0   unsustainable in this model. No geometry saves it.
h_t = 0   maintenance only. Hold, do not build.
h_t > 0   growth available; any requested increment must satisfy r_t <= h_t.
```

`h_t > 0` says growth is ATTAINABLE, not that it is happening. Headroom is computed at `S = ln N`, the orthogonal limit. Whether `Q` is actually rising this period is a different test — Corollary 3.2 — evaluated at the geometry actually present. A system can have positive headroom and be shrinking. See the worked block in A.10.

39.12.7 A.7 The circuit identity — the corpus's circuit-identity convention

`Q_{t+1} = (1-delta)*Q_t + eta*S_t - C_t` **is a leaky integrator**: in discrete time, exactly a first-order RC low-pass filter and charge accumulator.

Storage law	Circuit
`Q`	charge on a capacitor
`delta`	leakage conductance
`eta*S`	current in
`C`	load draw
`(1-delta)`	feedback gain / pole
`Q_inf`	steady-state charge
`tau = -1/ln(1-delta)`	time constant, in periods

Because `delta > 0` makes the map a contraction, **Banach's fixed point theorem** guarantees a unique fixed point and convergence to it from any start. **Recursion here is not a stack that overflows; it is an accumulator that integrates. That is why iteration arrives.**

Settling, for a step input: `3*tau` reaches 95.0% of the final value, `5*tau` reaches 99.3%.

39.12.8 A.8 Derived results — what this corpus derived, and where

The sections above are Papers A through E. The seven results below are not: they were derived in the course of writing this framework, each was verified numerically before it was printed, and each carries this paper that owns it and the ledger item that propagated it. ***Four of them change what is done****, which is why they are on the card and not in a footnote.*

D-1 · Headroom at the fixed point. `h_inf = eta*(ln N − F(u*))` — **CCP-008, CC-R6.**

At the fixed point, `delta`, `C` and `Q` all cancel. Verified against direct computation of `eta*ln N − delta*Q_inf − C` across four parameter sets; agreement to `< 1e−15`. Because `F(u) <= ln N` always, **`h_inf >= 0` always: the plain linear model never ossifies by itself.** Ossification is therefore not a property of the recurrence. It requires exogenous `Q` accretion or a cost that rises superlinearly in `Q` — which makes the maintenance trap a claim about the world rather than a consequence of the arithmetic, and it must be measured rather than assumed.

Refined at the internal scope (**CCP-021/41**): `h_inf = eta*(ln N − M*)`. At the fixed point, headroom is `eta` times the gap between the ceiling and the achieved moment, with cost cancelling (verified to 6 d.p.). Companion: `C_crit = eta*M − lambda*E_b` — **above this operating cost a departure is unreachable at any level of desire whatsoever.**

D-2 · The gradient identity. `grad h = −grad C` — **CCP-022, R19.**

Where `eta`, `N` and `delta` are constant in the direction of motion, headroom and cost have exactly opposite gradients. Two consequences, and the second is the most practically useful result in Part VIII:

1. `grad C = 0` iff `grad h = 0`. **The comfort zone and the headroom maximum are the same point**, under the stated restriction.
2. **A cheap habit or desire-line survey IS a headroom measurement.** the founding material's gradient node and this framework's master gauge read the same gradient. The full CCP-015 survey is not always required to locate the direction of steepest improvement — see Protocol C-9.

The restriction is load-bearing. Where `eta`, `N` or `delta` vary with position the identity fails and the full expansion of CCP-022 §42.12 applies. Do not use the shortcut across a boundary that changes the roster or the retention regime.

D-3 · The effective count is the participation ratio. `N_eff = N + 1 − N/PR` — **CCP-015, CC-R3.**

The mean-pairwise `u` aggregator and the Rényi-2 participation ratio `PR = (SUM lambda)^2 / SUM lambda^2` are the same effective count, exactly. Verified over random ensembles: worst discrepancy `1.8e−15`, most agreeing to machine zero. **This is a named identity, not an approximation**, and it means the eigenvalue route and the pairwise route to `N_eff` cannot disagree — so a disagreement in practice is a data or encoder fault, not a choice of estimator.

D-4 · The surplus law is a Fechner law. `F(u) = ln((u + u_0)/u_0)` with `u_0 = 1/(N−1)` — **CCP-020, CCP-039.**

Algebraically identical to `ln(1 + (N−1)u)`, term for term, at every `N` tested (worst difference `8.9e−16` over `u = 0.01` to `1.00` at `N = 2, 3, 5, 8, 12, 25, 100`). CCP-A02's surplus functional is Fechner's law with the stimulus being non-parallelism, `k = 1`, and an absolute threshold at `u_0`. And the threshold is not arbitrary:

```
N_eff(u_0) = 1 + (N−1) * 1/(N−1) = 2       exactly, for EVERY N
F(u_0)     = ln 2 = 0.693147 nats = 1 bit = 2 octaves       for EVERY N
```

The threshold of difference is exactly the point at which the effective count reaches two. Below `u_0` the system does not yet hold a second independent voice. **Declare `N` whenever the reader quote a surplus, because declaring `N` is declaring the threshold.**

D-5 · The storage law in world coordinates. `G_{t+1} = G_t^(1−delta) * N_eff^eta * e^(−C)` — **CCP-020.**

Write `G_t = e^{Q_t}` and exponentiate the storage law. Because `S_t = ln N_eff,t`, the surplus is already a logarithm and `e^{S_t} = N_eff,t` exactly. Verified against direct computation

at `delta = 0.08`, `eta = 0.35`, `C = 0.20`, `N = 12`, `u = 0.5`: worst difference `7.1e-15` over twelve periods. Read the three terms:

- **`G_t^(1-delta)` — decay is taking a root**, not a subtraction.
- **`N_eff^eta` — conversion is exponentiation**, not a fraction added.
- **`e^(-C)` — cost is a divisor**, not a deduction.

Fixed point: `G_inf = (N_eff^eta * e^(-C))^(1/delta)`, which for the parameters above evaluates to **295.6376**, identical to `e^{Q_inf}` with `Q_inf = 5.689135` to fourteen significant figures. **Stored order in world coordinates is an effective count, and the storage law is its weighted geometric accumulation.** No new claim is made — this is Theorem 3.5 in different coordinates — and what it shows is that the framework has been working in the lens all along.

D-6 · The two-track ledger. `Z_{t+1} = (1-delta)Z_t + eta(M_t - V_t)`, `C_t = c_0 + kappa*Z_t` — **CCP-019/40, R18.**

`Q` is all stored order; `Q^V` is *veridical* stored order; `Z = Q - Q^V` is stored error. With `M = min(r*E, ln N)` the moment and `V = E*min(r, 1)` the veridical throughput, **the cost term cancels exactly out of the error track** — dissipation debits both books equally, so it does not appear in their difference. Correction load then rises linearly in accumulated error.

An over-framed system accumulates MORE stored order than a calibrated one. It is simply not true stored order. Dogma banks faster than knowledge, and the bill arrives on a lag as `C_t` climbing with `Z_t`. This is the mechanism behind the doctrine-to-data pathology of CCP-024.

The consequence for `r = 1` is sharper than a caveat. At the `D = P` fixed point `S = ln N`, so `h_inf = eta*(ln N - F(u*)) = 0` by D-1: undistorted contact maximises input and converges to a state with *no room left to build.* And by D-6, **an organisation optimising raw `Q` will be actively driven away from `r = 1`.** Calibration is a target for the veridical book only, and the two books must be kept separately or the incentive runs backwards.

D-7 · The Jensen rule on feasibility. `mean F(U) <= F(mean U)` — **CCP-015/26, CC-R2.**

`f` is strictly concave (CCP-A02, Thm 3.3), so Jensen's inequality bites in the direction that flatters the system. **Test feasibility on the mean of `F(U_t)`, never on `F` of the mean `U`.** The gap is not decorative: it is worth **6.2%** of surplus in the CCP-016 worked office, and it grows with the variance of `U` across periods. A plan sized on `F(mean U)` is sized against a surplus the system never held in any single period.

```
RULE OF THE CARD:  average the surplus, not the geometry.
                   F(mean U) is what the roster looks like; mean F(U) is what it delivered.
```

39.12.9 A.9 Numeric tables

Table A-1. The logarithmic ceiling `ln N`.

N	ln N	N	ln N	N	ln N
2	0.693	8	2.079	25	3.219
3	1.099	9	2.197	30	3.401
4	1.386	10	2.303	40	3.689
5	1.609	12	2.485	50	3.912
6	1.792	15	2.708	100	4.605
7	1.946	20	2.996		

Doubling `N` adds `ln 2 = 0.693` wherever it starts. **A headroom problem cannot be escaped by adding domains.**

Table A-2. Retention loss, time constant, half-life. `tau = -1/ln(1-delta)`; `t_half = ln 2 / (-ln(1-delta))`.

delta	tau (exact)	1/delta	t_half (periods)	95% settle 3*tau	99% settle 5*tau
0.01	99.50	100	69.0	298	497
0.02	49.50	50	34.3	148	247
0.05	19.50	20	13.5	58	97
0.10	9.49	10	6.58	28	47
0.20	4.48	5	3.11	13	22
0.30	2.80	3.33	1.94	8.4	14
0.50	1.44	2	1.00	4.3	7.2

Table A-3. Effective count and surplus, `N = 10`.

u	N_eff	F (nats)
0.0	1.0	0.000
0.1	1.9	0.642
0.2	2.8	1.030
0.3	3.7	1.308
0.5	5.5	1.705
0.7	7.3	1.988
0.9	9.1	2.208
1.0	10.0	2.303

At `u = 0`, ten members are one member. `F = ln 1 = 0`.

Table A-4. Marginal gain from the Nth member at `u = 0.5`. `N_eff = (N+1)/2`.

N	N_eff	F	gain over N−1
2	1.5	0.405	0.405
3	2.0	0.693	0.288
4	2.5	0.916	0.223
5	3.0	1.099	0.182
6	3.5	1.253	0.154
8	4.5	1.504	0.118
10	5.5	1.705	0.095
11	6.0	1.792	0.087
20	10.5	2.351	0.049
40	20.5	3.020	0.025

The second genuinely different member is the largest single gain available. The fortieth is glazing bills.

39.12.10 A.10 Worked gauge block — carry these numbers as the card's calibration

```
GIVEN:   eta = 0.30 Q-units/nat     N = 8      ln N = 2.0794
         delta = 0.05 /period       Q_t = 8.0 Q-units
         C_t = 0.15 Q-units/period  u* = 0.50

CEILING  eta*ln N            = 0.30 * 2.0794           = 0.6238
```

```
MAINT    delta*Q_t           = 0.05 * 8.0               = 0.4000
HEADROOM h_t = 0.6238 - 0.4000 - 0.1500                 = 0.0738   -> POSITIVE
r_max                                                    = 0.0738

ACTUAL   S_t = ln(1 + 7*0.5) = ln 4.5                    = 1.5041 nats
         eta*S_t = 0.30 * 1.5041                         = 0.4512
GROWTH?  eta*S_t > delta*Q_t + C_t ?   0.4512 > 0.5500 ? = NO -> Q IS FALLING

REQUIRED S for zero growth: (0.4000+0.1500)/0.30         = 1.8333 nats
         1 + 7u > e^1.8333 = 6.2547  ->  u > 0.751

STEADY   Q_inf = (0.30*1.5041 - 0.15)/0.05               = 6.02 Q-units
BOUND    Q_t <= max{Q_0, eta*ln N/delta} = max{8.0, 12.48} = 12.48
```

Read this block every time the gauge is used. It contains the framework's most commonly missed distinction: **headroom was positive and order was still declining, because headroom is computed at the orthogonal limit and the system was running at `u = 0.5`.** The fix was not more people. It was `u > 0.751`, or less `C`, or less `Q`.

39.12.11 A.11 The honesty clauses — reproduce these wherever the law is used

1. **The triangle inequality FAILS for geometric triples** (CCP-A02 §5). Subadditivity holds on `u` only. Do not chain surplus across three parties as if it were a distance.
2. **As deployed, the surplus functional is a LOWER BOUND on the ensemble's information, not the information itself.** We log the floor, not the ceiling.
3. **CCP-A02 is pure mathematics and makes no empirical or metaphysical claim.**
4. **Surplus is NOT identified with Shannon entropy or mutual information in general** (CCP-009).
5. **No thermodynamic negentropy law is proved** (CCP-007). `F(u) + N_can(u) <= ln N` (Cor 2.2): **higher static surplus does NOT imply higher canonical concentration.**
6. **S3 is a design axiom.** The affine effective count is chosen. Badge it in the sentence that uses it.
7. **The storage variable `Q` is a modelling choice, not derived from first principles.** The linear update is the minimal tractable open-system law, not a law of nature.
8. **Woolley's `c` is from short-task laboratory settings.** Extrapolation to institutions is a modelling assumption.
9. **Centola's ~25% is from controlled online experiments.** Field transfer is a modelling assumption.
10. **Scale-freedom of the storage law is a claim about MODEL FORM, not about physics.** Only `(delta, eta, N, C)` change between a desk and a city; that the same form applies at both is asserted, not proved.

39.12.12 A.12 The four refusals — the standing refusals of the corpus

1. This corpus does not claim the surplus law proves anything about physics, thermodynamics, or consciousness.
2. This corpus does not present a postulate without its badge.
3. This corpus does not offer a number that was not computed or cited.
4. This corpus does not recommend an intervention that removes a person's ability to decline.

39.13 Appendix B Glossary

The framework speaks in two registers because it inherits two parents: a set of theorems and a trade. Each entry gives the term as this corpus uses it, the clinical or technical term it walks beside, and a definition that honors both. Where an entry is marked *(postulate)*, the pairing is this corpus's modelling proposal and not a proved result — the distinction Appendix E keeps.

stored order <-> *the state variable Q; accumulated durable structure* — What persists between periods without being re-created: the written procedure, the built room, the trained person, the standing relationship, the constitution. The quantity the whole manual is about, and the one whose units must be fixed before any other number means anything.

retention loss <-> *delta; the decay rate per period* — The fraction of stored order that disappears in one period with zero input. Forgetting, turnover, documentation rot, relationship cooling, skill atrophy, physical wear. Measured under a genuine zero-input window, never inferred from a busy quarter.

conversion efficiency <-> *eta; durable structure produced per nat of surplus* — How much of what happened between people becomes something that survives the meeting. The most neglected parameter in real systems and usually the cheapest to raise. A conversation with no artifact has `eta \ensuremath{\approx} 0` whatever its quality felt like.

dissipation <-> *C_t; the cost term* — Everything the period spends that does not become `Q`: coordination overhead, rework, friction, conflict, commute, compliance, emotional labor, rent. It enters the law with a minus sign and it is the term most often left off the ledger because it is nobody's line item.

surplus <-> *F(u) = ln(1 + (N-1)u); the diversity gain* — The information yield of an interaction between non-parallel views, in nats. Bounded above by `ln N`. **As deployed here it is a lower bound on the ensemble's information, not the information.**

effective count <-> *N_eff = 1 + (N-1)u; equivalent independent views* — How many genuinely distinct participants a roster of `N` participants amount to after the roster has been audited for redundancy. The quantity surplus takes the logarithm of, and the first number that flattery inflates.

non-parallelism <-> *u = 1 - (x·y)^2; decorrelation* — The degree to which two states fail to be multiples of one another. Zero for identical or opposite views; one for orthogonal ones. The single parameter that can be moved fastest, and the single parameter that decays fastest when unattended.

crossing <-> *u_cross = beta^2; between-category difference* — Non-parallelism that comes from participants standing in different categories: the physicist and the nurse, the accountant and the joiner. Bought by cross-domain recruitment; cheap to display and easy to mistake for the whole of difference.

divergence <-> *u_div = alpha^2 sin^2(gamma); within-category difference* — Non-parallelism between participants inside the same category: two physicists who genuinely disagree. Bought by protected dissent, not by hiring. **Crossing and divergence are not substitutes and a system can be rich in one and starved of the other.**

the design axiom <-> *S3; affinity of the effective count* — The assumption that `g(u) = e^{f(u)}` is affine in `u`, which is what makes `N_eff` linear and the law tractable. Chosen for tractability and interpretability, not derived from anything. Every number in this framework that uses `1 + (N−1)u` inherits its status. *(postulate)*

the logarithmic ceiling <-> *ln N; the per-interaction maximum* — The hard upper bound on surplus from any single interaction among `N` categories. Doubling `N` adds 0.693 nats from any starting point, which is why adding domains is the most expensive and least rewarding lever in the corpus.

the merciful clause <-> *strict concavity of f* — The property that each added increment of difference buys less than the one before. Merciful because forty windows are not needed: three honest ones set wide capture most of what a gallery could. Also unforgiving, because the last hire will not save the roster.

subadditivity on u <-> *f(u1+u2) <= f(u1)+f(u2)* — Difference from two sources never yields more than the sum of what each would yield alone. The bound that forbids promising synergy. It holds on the parameter `u`, and **the triangle inequality fails for geometric triples** — do not chain it across three parties.

extremality <-> *F = 0 iff parallel, F = ln N iff orthogonal* — The two endpoints of the law. Ten aligned instruments produce a reading of zero; the arithmetic is not being rude, it is counting. The orthogonal case is a limit that real rosters approach and never reach.

headroom <-> *h_t = eta ln N - delta Q_t - C_t; the growth budget* — The master gauge. How much new stored order this period could possibly support, evaluated at the orthogonal limit. Negative means no arrangement of people prevents decline; zero means maintain; positive means growth is attainable and bounded by `h_t`.

the sustainability ceiling <-> *Corollary 3.4* — The statement that a nondecreasing step requires `C_t <= eta ln N − delta Q_t`. **If cost exceeds that, no interaction geometry saves the system.** The single most consequential sociological sentence the formalism produces, and the reason cost-cutting sometimes outranks every clever design.

the maintenance trap <-> *the ossification regime* — The state a successful system reaches when `delta*Q_t` alone consumes the budget, leaving nothing for growth. Not a moral failure and not a loss of vision: an arithmetic consequence of `Q` having grown. **Success consumes its own headroom.**

shedding Q <-> *deliberate decommissioning* — Reducing stored order on purpose — sunset clauses, term limits, retiring a procedure, demolishing a building — in order to recover headroom. The fourth lever, the one nobody proposes, and frequently the cheapest available move at large `Q`.

leaky integrator <-> *first-order low-pass filter; RC charge accumulator* — The circuit the storage law is. Charge in, leak out, a pole at `(1−delta)`, DC gain `eta/delta`. The identity is exact in discrete time and it is what licenses every filter intuition this framework borrows.

contraction <-> *a map with Lipschitz factor below one* — What `delta > 0` makes the update. Guarantees, by Banach, a unique fixed point and convergence to it from any starting state. **The leak is not the flaw; the leak is what makes arrival possible.**

fixed point <-> *Q_inf = (eta F(u) - c)/delta; the steady state* — The stored order a system converges to under constant geometry and constant cost. the founding material's fixed-point node in its engineering form. If `eta F(u*) < c*` the fixed point is negative and the regime is unsustainable or the model must be replaced.

the transient term <-> *(1-delta)^t Q_0* — What the system started with, decaying. the founding material's recursion node in its engineering form. It vanishes at rate `delta` regardless of any intervention, which is why inherited advantage has a half-life and why a new regime feels like the old one for roughly `3*tau`.

the driven term <-> *SUM (1-delta)^{t-1-k}(eta S_k - C_k)* — What the system is currently building, accumulating. The half of the exact solution that is under design control. Arrival is the crossover: the moment the driven term exceeds what remains of the transient.

time constant <-> *tau = -1/ln(1-delta); approximately 1/delta* — How many periods the system takes to complete about 63% of a change. Sets the cadence of everything: at `delta = 0.05`, `tau` is 19.5 periods and a policy takes roughly 58 periods to be 95% in effect. Impatience is usually a `tau` misestimate.

admissibility <-> *C_t <= (1-delta)Q_t + eta S_t* — The condition guaranteeing `Q` does not go negative. A budget that violates it is not aggressive; it is describing a system that has already been consumed and is now being invoiced.

hyperrecursive symmetry <-> *scale-freedom of the model form* — The claim that the same storage law governs a desk, a room, a team, a building, an institution, and a city, with only `(delta, eta, N, C)` changing. **A claim about model form, not about physics.** *(postulate)*

the order proxy <-> *Y_t; the observable standing in for Q* — What is actually measured, since `Q` is a modelling construct and not a thing with a meter. Choosing it is the most consequential decision in any survey, and the most common place a study quietly fails.

the state encoder <-> *the map from participants to vectors x* — The operational choice that turns a person, view, or category into a unit vector so that `u` can be computed. One of CCP-007 §5.1's five required choices. Different encoders give different `u`; the encoder must be fixed before data is seen.

the repackaging error <-> *Y as a function of U; CCP-007 §6.3* — The failure in which the order proxy turns out to be a relabeled version of the input. If `Y` is computable from `U`, the model fits perfectly and has said nothing. **Justifying that `Y` is not a repackaging of `U` is a required step, not a nicety.**

the shuffled-input control <-> *the permutation baseline* — Refitting the model with the interaction series randomly reordered. If the shuffled fit matches the real one, the timing of interaction carried no information and the model is describing persistence only. The control most often omitted and most often decisive.

preregistration <-> *frozen specification before data* — Writing down the proxy, the encoder, the observables, the model, the baselines, and the falsification rules, and then not changing them. **A framework that cannot fail cannot be trusted, and this one names exactly how it fails.**

badged postulate <-> *a flagged modelling assumption* — An assumption stated together with the fact that it is an assumption, in the same sentence. The house rule inherited from the founding material: the modelling is permitted, but it is permitted in the open, with its badge on.

c-factor <-> *collective intelligence; Woolley et al. 2010* — A single statistical factor predicting a group's performance across varied tasks, not predicted by members' average or maximum IQ, correlating with social sensitivity, equality of conversational turn-taking, and proportion of women. Measured in short laboratory tasks. Not a term in the law.

turn-taking equality <-> *low variance in speaking-turn distribution* — The `c`-factor driver that is substantially a property of the room rather than of the people. Raises `u` in effect by admitting the views that would otherwise not be heard, which is why acoustics is an epistemic intervention.

social sensitivity <-> *accuracy at reading others' mental states* — The second measured `c`-factor driver. Distinct from agreeableness and distinct from cohesion; it is the capacity to notice that someone has not finished, which is the mechanism by which it produces turn-taking equality.

cohesion trap <-> *groupthink; the u-destroying comfort* — The condition in which a group's mutual liking has driven its members' views into alignment. Feels excellent from inside and reads as high performance on every soft measure. **Cohesion is often the enemy of `c`, because cohesion drives `u` toward zero.**

pseudo-diversity <-> *demographic variance with epistemic uniformity* — A roster that looks assorted and computes to `N_eff \ensuremath{\approx} 1`. Three glasses ground from one doctrine, their agreement read as convergence. Priced exactly by the arithmetic: `u = 0`, `N_eff = 1`, surplus `ln 1 = 0`.

the parallel wall <-> *the u ≈ 0 regime* — The state in which every additional participant adds cost and no surplus. Named because it is a wall: no amount of `N` gets through it, and the only exits are raising `u` or lowering `C`.

the collision test <-> *the shared-conclusion check* — the founding material's instrument, retained. Run two independent lines of reasoning to their furnished rooms; if they land in the same room, they were one window. Cheap, fast, and the first test to run on any panel that agrees.

error decorrelation <-> *u ≈ 1 - r^2 on centered prediction errors* — The operational estimator of non-parallelism. Under the unit-vector encoding, the inner product is the Pearson correlation of errors, so non-parallelism is one minus its square. An operational choice, badged as such. *(postulate)*

participation ratio <-> *(SUM lambda)^2 / SUM lambda^2* — The eigenvalue estimator of effective count from a correlation matrix. Gives `N_eff` directly for ensembles larger than a

pair, from which `u = (N_eff - 1)/(N - 1)`. The aggregation of a pairwise law to an ensemble is an operational choice. *(postulate)*

pre-mortem <-> *prospective hindsight elicitation* — Asking a group to assume the plan has failed and to write down why, independently, before discussing. Buys `u_div` from a group that has none, because it licenses the objection that social cost had suppressed.

information cascade <-> *Bikhchandani, Hirshleifer & Welch 1992* — The sequence in which rational agents, observing predecessors, rationally stop using their private signals. The mechanism by which a genuinely diverse group becomes a parallel one without anyone changing their mind. **`u` is not a property acquired once; it is a quantity that decays.**

complex contagion <-> *Centola & Macy 2007; adoption requiring multiple reinforcing contacts* — Behavior that spreads only when a person encounters it from several sources. Because reinforcement is required, **clustered topology beats random long ties** — the opposite of the reach intuition.

committed minority <-> *Centola et al. 2018; the ~25% tipping point* — The fraction of a population holding a new convention that suffices to overturn an established one. Measured in controlled online experiments with small groups. Field transfer to organizations is a modelling assumption.

threshold <-> *Granovetter 1978; the individual adoption cut-off* — The number or fraction of others whose adoption a given person requires before adopting. Aggregate behavior depends on the whole distribution of thresholds, which is why a group can be one person away from a riot or from nothing.

stigmergy <-> *Grassé 1959; coordination through a modified medium* — Order produced by agents responding to traces left in a shared environment, with no agent holding the plan. Termite mounds, ant trails, version-controlled repositories. **Design rule: put the state in the medium, not in a coordinator.**

boundary object <-> *an artifact legible in two communities at once* — A drawing, a schedule, a shared dataset that each discipline can use in its own terms without a common vocabulary. The cheapest `eta`-raiser at a silo interface, because it converts cross-domain talk into something that persists.

focal point <-> *Schelling 1960; coordination without communication* — The salient option people converge on when they must agree but cannot confer. Noon, the middle, the round number, the obvious landmark. A `C`-reducer: it eliminates the negotiation entirely.

requisite variety <-> *Ashby 1956; only variety absorbs variety* — The law that a regulator can only control as much disturbance as it has states to answer with. The reason `N` matters at all: a system facing five kinds of trouble with three kinds of response is short by two, and no diligence closes the gap.

semilattice <-> *Alexander 1965; overlapping category structure* — A structure in which sets may intersect without one containing the other. Carries more surplus than a tree because it permits a member to belong to several categories at once. **Hierarchy is the organizational form that drives `u` toward zero.**

silo <-> *a local minimum on the comfort landscape* — A position deep because training and reward carved it, narrow because it admits one vocabulary. Exhortation fails because every direction out is uphill. **Build a lower valley next door; do not push.**

Conway's law <-> *systems mirror the communication structure that built them* — Why architecture reproduces the org chart and why reorganizing the team is a design act. The structural cause of silos, and the reason an interface problem is often a seating problem.

the quadratic committee <-> *Brooks 1975; n(n-1)/2 communication paths* — The observation that coordination cost grows quadratically while surplus grows logarithmically. The crossover point is the most important team-sizing calculation available, and it always arrives sooner than the roster expects.

utilization cliff <-> *queueing; wait time exploding as rho approaches 1* — The property that waiting grows as `1/(1-rho)`, so the last few percent of capacity cost more delay than all the rest combined. **Slack is not waste; it is a cost-control device.**

Goodhart capture <-> *Goodhart 1975; Campbell 1979* — The degradation of a measure once it becomes a target. Applies directly to every number this framework asks a user to track, including headroom. The defense is to keep the measurement instrument separate from the reward instrument.

invisible infrastructure <-> *Star & Ruhleder 1996; visible only upon breakdown* — The property that working infrastructure disappears from attention, so its maintenance is chronically underfunded until it fails. The reason `delta` is systematically underestimated in every organization that has not measured it.

optional activity <-> *Gehl 1971; discretionary presence* — Activity people undertake only when conditions invite it: sitting, lingering, watching. Distinguished from necessary activity, which occurs regardless. **Optional activity is the observable proxy for a room's `eta`.**

triangulation <-> *Whyte 1980; an external stimulus that occasions conversation* — A busker, a sculpture, a view, an odd object that gives strangers a shared referent and thereby a licensed opening. One of the few reliable ways to raise the encounter rate without raising density.

sittable space <-> *Whyte's first finding* — People sit where there is somewhere to sit. Obvious, repeatedly ignored, and the single highest-yield physical intervention per unit cost in the literature this framework draws on.

edge effect <-> *perimeter occupancy preference* — The observed tendency for people to occupy the boundaries of a space before its middle, because edges offer support at the back and prospect to the front. Design consequence: furnish the perimeter first.

threshold <-> *the architectural transition; impedance matching between zones* — The porch, the genkan, the airlock, the widened landing. A device that lets a person change register gradually instead of abruptly. **The transition is not a gap between rooms; the transition is the device.**

proxemics <-> *Hall 1966; the distance bands of interaction* — The observation that interaction character changes with separation, in roughly demarcated intimate, personal, social, and public bands. Culturally variable, as Hall said; use as a design frame, not as a specification.

RT60 <-> *Sabine; reverberation time, 0.161 V/A metric* — Seconds for sound to decay 60 dB, from room volume in cubic metres and total absorption in square-metre sabins. Valid for reasonably diffuse fields; use Eyring's form for heavily absorbing rooms.

the Lombard effect <-> *involuntary vocal effort increase in noise* — The reflex by which people raise their voices in a noisy room, which raises the noise, which raises the voices. The positive feedback that turns a moderately reverberant room into a loud one and hands it to whoever is loudest.

the head of table <-> *geometric rank assignment* — The seat that assigns `u = 0` to everyone below it by making their function agreement rather than judgment. **Furniture is the cheapest governance instrument ever built, and it enforces continuously without anyone administering it.**

silent-write-first <-> *independent elicitation before discussion* — The protocol of collecting written independent judgments before any discussion begins. Breaks cascades before they start and preserves the independence Condorcet's theorem requires. The single cheapest `u`-preserving move in this framework.

adjustability <-> *individual control over local conditions* — Openable windows, task lighting, movable chairs, personal thermal control. **Adjustability outperforms optimization because a fixed optimum serves the mean occupant and no actual one.**

desire line <-> *revealed preference; the worn path* — The route actually taken across the route provided. The most honest document a population produces and the only one it never signs. Read it before regrading; regrade toward it, not against it.

red team <-> *institutionalized adversarial review* — A group chartered and rewarded for finding the flaw. Buys `u_div` structurally rather than depending on individual courage, which is why it works and why the charter must protect the finding, not merely permit it.

reversibility <-> *the design requirement for capped downside* — The property that an intervention can be undone at bounded cost. Required by this framework for any change large enough to matter, because parameter estimates in social systems carry uncertainty that no protocol here removes.

the ability to decline <-> *the consent constraint; the standing refusals of the corpus* — The property every intervention in this framework must preserve. A method that works by removing resistance has no coercive failure mode: its worst outcome is that someone declines to move. A method that works by removing the option has no such floor.

The entries above are the institutional-scope material's vocabulary: rooms, rosters, institutions. The entries below enter with the personal- and cellular-scope material, where the same instrument is turned on a mind and then on a cell. They are kept in one list because they are one vocabulary — but four of them carry a demotion, and the demotion is stated in the entry rather than in a note. Where an entry is marked **(demoted)**, *its structure stands and its magnitude does not; see Appendix E.7.*

moment <-> *M = (D/P)·E; the unit of lived experience* — What is actually experienced in a period: the environment as read through the ratio of frame to intake. The input term of the self-storage law, `Q_{t+1} = (1−delta)Q_t + eta*M_t − C_t`, and therefore the point at which the personal-scope material joins the institutional-scope material. **Its structure is defensible and its magnitude is not yet reportable.** *(postulate; demoted)*

perspective <-> *D; the frame brought to the moment* — Everything the perceiver supplies before the environment says anything: prior, training, expectation, doctrine, precision-weighting of the top-down model. The numerator of the calibration ratio. Not yet instrumented on a ratio scale, which is exactly the demotion. *(postulate; demoted)*

perception <-> *P; the intake admitted from the moment* — What the environment actually delivers into the frame: sensory evidence, data, report, the bottom-up signal and its precision. The denominator of the calibration ratio, and the term that goes to zero in dogma. *(postulate; demoted)*

calibration ratio <-> *r = D/P; frame over intake* — The dimensionless ratio that says whether a perceiver is meeting the world or overwriting it. `r = 1` is undistorted contact; the ordering is genuine and the number is not yet measurable. **`r = 1` is not an unqualified target** — headroom at that fixed point is exactly zero, and a system optimising raw stored order is driven away from it. *(demoted)*

distortion <-> *Delta = |ln r|; the log-ratio departure from contact* — How far a frame stands from the intake it is reading, in the same log register as everything else in this framework. **Ordinal pending instrumentation:** the direction stands, the magnitude does not, until `D` and `P` are Rasch-calibrated (Theorem 54.3). *(demoted)*

over-framed <-> *r > 1; doctrine outrunning data* — The condition in which the moment exceeds the environment: projection, prejudice, apophenia, seeing what was predicted. At group scope it is the institution whose doctrine outruns its intelligence. **It banks stored order faster than a calibrated system does, and the order is false** — see *stored error*.

under-framed <-> *r < 1; data outrunning doctrine* — The condition in which the moment falls short of the environment: flooding, overwhelm, an intake that cannot be made to mean anything. Not ignorance — a surfeit of signal with no frame to read it.

lens theorem <-> *the logarithmic map; Cauchy's functional equation* — If perception converts multiplication into addition, then it is a logarithm, unique up to scale. **Linear in the lens is exponential in reality:** equal perceptual steps are equal world *ratios*, never equal world differences. The theorem is conditional and proved; its antecedent is empirical and badged separately. Design in ratios; never promise a linear improvement in a quantity the recipient perceives logarithmically.

doctrine-to-data ratio <-> *D/P at group scope* — The calibration ratio read on an institution. **Never aggregate it across sites:** distortion in opposite directions cancels in the aggregate

and overstated bankable input by **4.26x** in CCP-024's two-site example. Distortion does not average.

internal surplus <-> *F_int = ln(1 + (N-1)u) over one person's own framings* — CCP-A02 applied within a single observer. the founding material's "many windows of the one tower" as a derivation rather than an analogy. Subject to Break 1 of CCP-038: **the independence the aggregation needs is unavailable inside one skull**, so the estimate is an upper bound on what a mind actually holds.

essence index <-> *n; ordinal position within a discipline's category set* — The integer an essence type occupies in its own ordered set. Atomic number is one instance of an essence index, not the definition of one. Across disciplines there is no shared ordinal, and the effective count supplies the missing one.

essence frequency <-> *f(n) = k n², k = 0.483425732988414 Hz* — The frequency the essence law returns for an essence index. **An essence frequency is not a physical resonance**, and the distinction is absolute: the constant is a Pythagorean interval divided by the golden section, and it cancels out of every relation, so only the ratios `m²/n²` carry content.

conversion key <-> *an exact constant between logarithmic registers* — The linear-in-log-space constant relating surplus, pitch, sensation and moment. `1 nat = 1.442695 bits = 2.885390 octaves = 3462.4681 cents`; `1 bit = 2 octaves exactly`. **One nat of interaction surplus is exactly 2.885390 octaves** — not an analogy, arithmetic.

pitch class <-> *an equivalence class under octave folding* — Two essences share a pitch class iff their indices have the same odd part, because `(2n)² = 4n²` is exactly two octaves. An ordered set `1..N` therefore yields exactly `ceil(N/2)` distinct classes. **Folding costs exactly `ln 2` — one bit, two octaves — for every even `N`.** A hard structural limit on harmonic vocabulary, not a stylistic choice.

Gabor limit <-> *dt · df >= 1/(4π) = 0.079577* — The time-frequency uncertainty bound. **Classical, true of every wave including sound, with no quantum content whatsoever**, and measurable on an oscilloscope. A short sound cannot have a precise pitch: 1 ms buys 79.577 Hz of resolution, 100 ms buys 0.796 Hz. It hard-constrains any rendering of an essence set, and it prices in time what it costs to hear a difference in surplus — at least `1/(8*pi*dF)` cycles.

channel <-> *the coupling from a belief to a change in the world* — The pathway a belief must act through if it acts at all. **"Grounded in truth" means the channel exists and is real.** Naming it is step 2 of the grounding diagnostic, and a claim that cannot name one has failed before any measurement.

coupling <-> *E = G·B; the belief-to-world relation* — The relation that formalizes the postulate. Structurally the storage law: belief is the input, the channel is the conversion efficiency, the world-change is what lands in `Q`. **Series channels multiply their gains; parallel channels add them**, which is why a four-stage chain of plausible steps can have a total gain of 0.084.

gain <-> *G; change in the world per unit belief* — The size of what a channel delivers. **If `G = 0` then the outcome is zero for every value of belief whatsoever**, which disposes of about a third of the manifestation failure catalogue in one line. Gain must be estimated from quantities that are not the outcome, or the estimate is the claim.

grounding <-> *establishing the channel before the outcome is known* — The seven-step diagnostic of CCP-035. **Grounding is established from the channel, before the outcome is known, or it is not established at all** — otherwise the postulate's conditional is diagnosed backwards from the result and predicts nothing. Step 5, the predicted magnitude, does most of the killing.

evidence column <-> *the (+1) column; what is measured* — What scientific measurement confirms: replicated, falsifiable, obtained by a stated method in a named setting. The only column that can move a posterior.

precedent column <-> *the (-1) column; what tradition recorded* — That an idea was arrived at independently, recorded and preserved, often for centuries, by people with no access to the instruments. **This is real and worth having — it argues that a structure was findable by more than one route — and it is not a second measurement.** It can never be promoted to evidence; the historical corpus is a survivorship-filtered record. Where the corpus has nothing to say,

the cell reads *none claimed*, and in this corpus it reads that way in **133 of 156 register rows**.

concordance <-> *a bridge drawn between the two columns* — A claim backed from both corpora at once, valid only if it passes all three criteria: **C1 structural correspondence** (a deep shared pattern, not a shared cardinality), **C2 predictive power** (it must generate a prediction), **C3 practical utility** (it must be usable on both sides). **Two sources agreeing does not double the evidence when only one of them is evidence**, and the role of each column is stated every time a concordance is drawn.

Tri-State validation <-> *the VINO apparatus, adopted unchanged* — The three-position frame in which every substantive claim of the cellular-scope material is assessed: **(+1) physics proof**, **(-1) metaphysics or recorded practice**, **(0p) contra-proof at the Neutral Point**, where the two are made to integrate or shown not to. Only bridges satisfying all three criteria are presented as established; speculative bridges are marked speculative in the same sentence that states them.

tri-state <-> *the three-valued status of a bridge* — `+1` established, `−1` precedent only, `0p` the neutral point where integration is attempted. **A bridge may hold at one grain and fail at a finer one**, and finding that seam is most of the work: of CCP-029's three worked bridges, one was established, one was graded, and one was rejected outright.

specific objectivity <-> *Rasch; separation of person and item parameters* — The property that a person's measure does not depend on which items were used, and an item's measure does not depend on which persons answered. **The closest thing psychometrics has to physical measurement, and exactly the property the moment equation needs** to earn its division back.

Rasch separation <-> *the logit difference as a ratio-scaled quantity* — The construction that supplies the true zero division requires, so that `Delta = |ln(D/P)|` becomes genuinely ratio-scaled. **This is the specific, principled route from demotion to measurement**, and it is priced: **123 items per term** to establish a half-logit distortion, 342 for three tenths.

ratio scale <-> *Stevens' typology; a scale with a true zero* — The scale type on which multiplication and division are permitted operations. Interval scales permit differences; ordinal scales permit only ordering. **A formula that divides two terms measured ordinally is a heuristic ordering, not an equation, and this corpus demotes such formulas in print rather than in the notes.**

magnitude estimation <-> *Stevens 1957; psi = k I^a* — The psychophysical procedure that produces genuine ratio-scale measurements of subjective experience from human subjects, and replicates. The exponent `a` is a measured conversion key between the physical and perceptual registers. **Report exponents as constants of a procedure, not constants of nature** — they move with method, range and instruction.

veridical stored order <-> *Q^V; the true half of the ledger* — Stored order that corresponds to something in the world, accumulated from the veridical throughput `V = E·min(r,1)` rather than from the moment. The corpus that matters and the corpus nobody keeps.

stored error <-> *Z = Q - Q^V;* `Z_{t+1} = (1−delta)Z_t + eta(M_t − V_t)` — The accumulated gap between what a system has banked and what is true. **The cost term cancels exactly out of this track**, because dissipation debits both books equally. **Dogma banks faster than knowledge**, and this is the quantity that records it.

correction load <-> *C_t = c_0 + kappa*Z_t; the lagged bill* — The part of the dissipation budget that exists only to correct accumulated error. It rises linearly with stored error, which is why an over-framed system's costs climb long after the framing that caused them, and why the arrival of the bill is routinely attributed to something else.

launch budget <-> *L = max(0, M - C); escape is dreamt from surplus* — What remains of a moment after the period's cost is paid, and the fund from which any departure is financed. **`L = 0` is the formal statement of "there was no surplus to dream with."** Raising cost globally destroys the platform; raising it only at the origin genuinely shortens the wait — 32.0 weeks to 9.9 — at the price of a 37% smaller permanent fund. **The question is whether it is a cost that can be left behind or one that is carried.**

return difference <-> *Delta_R = |ln(M'/M)|; the founding material's echo node* — The

world's answer over what was sent, read as a log-ratio because ratios compose multiplicatively. **Zero means no information came back** — which is the same reading the instrument gives for the flatterer and for the reliable contrarian, since `u` depends only on `|rho|`.

survivorship likelihood ratio <-> *the evidential weight of testimony under selection* — Computed at **exactly 1.000000** under a process that publishes successes and discards failures. **A likelihood ratio of 1 does not move the posterior at all** — a thousand testimonials carry precisely the weight of zero, because the selection guarantees that successes are observed whether or not the method does anything. Companion figures: a direction-only prediction is worth 2.0000, a magnitude prediction with a tolerance 11.3577 — **about 5.7x more evidentially powerful.**

the postulate <-> *belief creates reality if the belief is grounded in truth* — the cellular-scope material's keystone, and **the conditional clause is the whole content.** Belief is an input; the system must have gain. No channel, no effect. Stated without the conditional it is magical thinking; stated with it, every failure of manifestation becomes a diagnosable fault at a named coordinate.

manifestation <-> *destiny made manifest; a latent capability brought into expression* — Used in this framework in a personal and entirely non-territorial sense. **The phrase "manifest destiny" carries specific nineteenth-century American historical freight** — it was the slogan used to justify continental expansion and the displacement of Indigenous peoples — and this corpus says so at first use rather than trusting the reader to supply the distinction.

demotion <-> *a claim this corpus advanced and its own protocol reduced* — The status assigned when a formula passes some Formula Validation Protocol gates and fails others: **the structure stands, the magnitude does not**, and the demoted content may not reappear in any form, including in a summary table. **Eight register entries carry this status, and five of them are the author's own keystone formula and its family.** Demotion is one-directional and it travels: when a gate fails, every entry downstream of the failed one is demoted in the same pass, in this paper that made the claim and not only in the appendix.

39.14 Appendix C Measurement protocols

Six protocols, one per parameter, plus a seventh that assembles them into the gauge. Each is written to be run by one engineer with a notebook, and each ends with a worksheet designed to be photocopied and filled in by hand.

Three more were added after the volumes that follow this appendix were written, and one of them is a correction to the six above. C-8 replaces the assumption that a passed shuffled-input control means anything much; it does not, and the three ways it does not have been measured rather than argued. C-9 and C-10 go the other way: two cheap tests that were available from the arithmetic all along and were not written down until Parts VIII and IX derived them.

Before any protocol runs, fix the five operational choices from CCP-007 §5.1 in writing, and do not change them afterward.

```
FIVE REQUIRED CHOICES — fix in writing, date and sign, before any data is collected

1. STATE ENCODER          how a participant/view/category becomes a unit vector x
2. INTERACTION OBSERVABLE how U_t is read in each period
3. ORDER OBSERVABLE       what Y_t is, and its unit
4. TIME SCALE             what one period is (day / week / month / quarter / year)
5. COST OBSERVABLE        what C_t is, in the SAME unit as Y_t

Signed: _____________________  Date: ____________
```

Changing any of the five after seeing data invalidates every held-out claim that follows. If a choice was wrong, say so, discard the study, and start again with the corrected choice. This is cheaper than it sounds and far cheaper than the alternative.

39.14.1 C.1 Protocol for delta — retention loss

Principle. `delta` is defined as loss under *zero input*. Any window in which interaction continued measures `delta` net of `eta*S` and will underestimate it, usually badly.

1. **Identify the order proxy `Y`** using Protocol C-6 first. `delta` cannot be measured before `Q` has an observable.
2. **Find a zero-input window** — a shutdown, holiday, hiring freeze, sabbatical, mothballed process, discontinued programme, archived project. The requirement is that no new surplus was converted, not that nothing happened.
3. **Verify it is genuinely zero-input.** List every activity in the window that could have added to `Y`. If the list is not empty, exclude the contaminated portion or find another window.
4. **Record `Y_0` and `Y_1`** at the window's ends, with the same instrument, observer protocol, and point in cycle.
5. **Record the window length `m`** in periods, using the time scale fixed in the five choices.
6. **Compute the per-period retention loss:** `delta = 1 - (Y_1/Y_0)^(1/m)`.
7. **Repeat on a second, independent window.** One window is an anecdote. Two windows that agree within a factor of 1.5 give a usable estimate; two that disagree by more mean the proxy is measuring something other than stored order, and the analysis must return to Protocol C-6.
8. **Report the range, not the point.** Carry `delta_low` and `delta_high` through every downstream calculation and report headroom as an interval.
9. **Compute the derived quantities** and record them on the card: `tau = -1/ln(1-delta)`, `t_half = ln2 / (-ln(1-delta))`, settling `3*tau` and `5*tau`.
10. **Re-measure annually.** `delta` is not a constant of the system; it moves with turnover, documentation quality, and physical condition.

Common errors. Measuring during an active period and calling the result `delta`: the most frequent error, and it produces a `delta` near zero, a fictional `tau`, and a maintenance budget of nothing. Using an administratively maintained proxy — a register, a headcount — whose `delta` is the clerk's diligence, not the system's.

```
WORKSHEET C-1 — DELTA
System: _____________________  Order proxy Y: _____________________
Unit of Y: _______________  Period length: _______________

Window 1 label: _____________________ Contamination check passed? [ ]
  Y_0 = ___________  Y_1 = ___________  m = _______ periods
  Y_1/Y_0 = ___________   (Y_1/Y_0)^(1/m) = ___________
  delta_1 = 1 - ___________ = ___________

Window 2 label: _____________________ Contamination check passed? [ ]
  Y_0 = ___________  Y_1 = ___________  m = _______ periods
  delta_2 = ___________

Agreement within factor 1.5?  [ ] yes  [ ] no -> return to C-6
delta_low = ___________  delta_high = ___________  delta (point) = ___________

DERIVED:  tau = -1/ln(1-delta) = ___________ periods
          t_half = 0.693/(-ln(1-delta)) = ___________ periods
          95% settling = 3*tau = ___________   99% = 5*tau = ___________
Measured by: _______________  Date: ___________  Next re-measure: ___________
```

39.14.2 C.2 Protocol for eta — conversion efficiency

Principle. `eta` is durable structure produced per nat of surplus. It is the ratio of what stuck to what happened, and the numerator must be *attributable*.

1. **Enumerate interaction events** over a span of at least `3*tau` periods or twenty periods, whichever is longer. Each event gets a date, a participant list, and a category list.
2. **Compute S for each event**: run C-3 for `N` and C-4 for `u`, then `S = ln(1 + (N−1)u)`, in nats.
3. **Define attributability in advance**, in writing, before looking at any `Y` data. Acceptable: the artifact cites the event; the decision log names it; the change is dated within one period with no other candidate. Unacceptable: someone's recollection.
4. **Record attributable `delta_Y`** for each event, in `Y`-units.
5. **Direct estimate:** `eta = sum(attributable delta_Y) / sum(S)` over the span.
6. **Regression estimate.** Fit the minimal estimation model `Y_{t+1} = a*Y_t + b*F(U_t) − C_hat_t + eps_t` with constraints `0 <= a <= 1`, `b >= 0`. Then `eta = b` and, as a cross-check, `delta = 1 − a`.
7. **Compare the two estimates.** If the direct estimate and the regression estimate differ by more than a factor of two, the attributability rule is either too generous (direct high) or too strict (direct low). Report both, and say which is preferred and why.
8. **Cross-check `delta`.** The regression's `1 − a` should agree with Protocol C-1's `delta` within the interval reported there. Disagreement means the proxy, the time scale, or the cost observable is wrong.
9. **Freeze and test.** Fit on the training span, freeze every coefficient, and report performance on a held-out span. See CCP-015 and Protocol C-6 step 7.
10. **Stratify by capture mechanism.** Compute `eta` separately for events that produced a written artifact, a built object, a changed default, or a trained person, versus those that produced none. The gap between the two is the size of the cheapest available intervention in the system.

Common errors. Counting satisfaction as `delta_Y`; counting the meeting itself as the artifact; attributing change to the most recent event by default, which credits whatever happens most often.

```
WORKSHEET C−2 — ETA
System: _____________________  Span: ______ periods (>= 3*tau = ______)
Attributability rule (written before data): ____________________________________
_____________________________________________________________________________

Event | Date | N | u | S=ln(1+(N−1)u) | attributable dY | capture mechanism
  1   |      |   |   |                |                 |
  2   |      |   |   |                |                 |
  3   |      |   |   |                |                 |
  4   |      |   |   |                |                 |
  5   |      |   |   |                |                 |
 ...  |      |   |   |                |                 |

SUM S = __________ nats     SUM attributable dY = __________ Y−units
DIRECT      eta = dY/S = __________ Y−units per nat
REGRESSION  a = ______  b = ______  −> eta = b = ______  delta = 1−a = ______
Agreement within factor 2?  [ ] yes  [ ] no −> revisit attributability rule
C−1 delta = ______   regression delta = ______   consistent? [ ]

STRATIFIED  eta with artifact = ______   eta without artifact = ______
            gap = ______  <− size of the cheapest available intervention
```

39.14.3 C.3 Protocol for N — the category count

Principle. `N` counts *genuinely distinct* categories. A relabeled category is not a category. The audit is the protocol; the count is the leftover.

1. **List every category the system nominally admits.** Job titles, disciplines, departments, sites, functions, stakeholder classes — whatever the system's own vocabulary supplies. Call this list `N_nominal`.
2. **The training-lineage test.** Were members trained in the same tradition, by the same institutions, on the same corpus? Identical answers make the pair a merger candidate.
3. **The incentive test.** Are they rewarded and punished by the same measures? Identical incentives produce identical judgment under pressure regardless of nominal category.
4. **The vocabulary test.** Do they use the same words for the same things and notice the same anomaly? If a problem statement lands identically in both, they are one category for surplus purposes.
5. **The error test.** Where data permits, check whether the two categories' errors on the same problems are correlated. The strongest test; the other three are proxies for it.
6. **Merge the failures.** Any pair failing all four tests is one category. Record the merger and its reason; the audit trail is the deliverable, not the number.
7. **Report `N` and `N_nominal` together.** The ratio is a diagnostic in itself: a system with `N_nominal = 14` and `N = 4` has an accurate self-image problem prior to any design problem.
8. **Compute `ln N` and `eta*ln N`** and carry them to the headroom sheet.
9. **State what a new category would have to be.** Before proposing to raise `N`, write the specific training lineage, incentive structure, and vocabulary the new category would bring. If it cannot be written, `N` cannot be raised — only `N_nominal` can.

```
WORKSHEET C-3 — N
System: _____________________  Date: __________
N_nominal: ______   List: ____________________________________________________
_____________________________________________________________________________

Merge log:
  merged _____________ + _____________  reason: _____________________________
  merged _____________ + _____________  reason: _____________________________
  merged _____________ + _____________  reason: _____________________________

N (audited) = ______     N_nominal/N ratio = ______
ln N = __________        eta*ln N = __________ (needs C-2)

If raising N is proposed, the new category's:
  training lineage: ____________________________________________
  incentive structure: _________________________________________
  distinct vocabulary: _________________________________________
  If any line is blank, N cannot be raised. Only N_nominal can.
```

39.14.4 C.4 Protocol for u — non-parallelism

Principle. `u = 1 - (x·y)^2` is defined for a pair of unit vectors. Estimating it for a real ensemble requires an encoder and an aggregation rule, both of which are operational choices, and **both must be badged as such in any report.**

1. **Build a common item set.** Twenty to fifty judgments, forecasts, ratings, or diagnoses that every participant can render independently and that have or will have a scoreable answer. The set must be relevant to the system's actual work; a generic instrument measures a generic `u`.

2. **Collect independently and in writing, before any discussion.** This is not merely good hygiene; discussion prior to collection measures the post-cascade `u`, which is the quantity the measurement was meant to detect.
3. **Encode.** Assemble each participant's responses into a vector; center it by subtracting the item-wise mean across participants; normalize to unit length. *(Centering and normalizing is the state-encoder choice; badge it.)*
4. **Pairwise estimate.** For each pair, compute `r = x·y`, the Pearson correlation of centered responses, and `u_ij = 1 − r^2`. Note that `u = 0` at both `r = +1` and `r = −1`: a reliable contrarian is as parallel as a reliable agreer, and this is correct — both are one instrument.
5. **Ensemble estimate — the participation ratio.** Form the correlation matrix `R` across participants, take its eigenvalues `lambda_i`, and compute `N_eff = (SUM lambda_i)^2 / SUM lambda_i^2`. Then invert the affine form: `u = (N_eff − 1)/(N − 1)`. *(The aggregation of a pairwise law to an ensemble is an operational choice; badge it. It is consistent with S3 by construction, which is a property of the choice and not evidence for it.)*
6. **Where answers become scoreable, switch to error decorrelation.** Recompute using centered *errors* rather than centered *responses*. Error decorrelation is the quantity Condorcet's independence condition actually names, and it is strictly better evidence than agreement on unresolved items.
7. **Decompose into crossing and divergence.** Pool same-category pairs for `u_div` and cross-category pairs for `u_cross`, and report both. `u_cross = 0.6` with `u_div = 0.05` is the classic "diverse hires, one permitted opinion" and needs a different intervention from its mirror image.
8. **Run the collision test as a qualitative check.** Two participants reason a scenario independently to its conclusion; if the conclusions coincide in detail while the numeric `u` is high, the item set is not probing load-bearing assumptions.
9. **Track across time.** `u` decays, so one measurement cannot distinguish a stable system from one in free fall. Quarterly is a reasonable default at `tau` of tens of periods.
10. **Compute `S = ln(1 + (N−1)u)`** and carry it to the headroom sheet.

Common errors. Measuring `u` after discussion; using demographic categories as the encoder, which measures demographic variance and not non-parallelism; reporting one `u` when `u_cross` and `u_div` differ by an order of magnitude.

```
WORKSHEET C-4 — U
System: _____________________  Item set: _____________  Items: ______
Collected independently before discussion?  [ ] yes  [ ] NO -> discard
Encoder (badged as an operational choice): centered, unit-normalized  [ ]

Pairwise (fill for key pairs):
  pair ____ & ____   r = ______   u = 1-r^2 = ______
  pair ____ & ____   r = ______   u = ______
  pair ____ & ____   r = ______   u = ______

Ensemble (participation ratio):
  SUM lambda = ______   SUM lambda^2 = ______
  N_eff = (SUM lambda)^2/SUM lambda^2 = ______
  N (from C-3) = ______   u = (N_eff-1)/(N-1) = ______

Decomposition:  u_div (within-category) = ______   u_cross (between) = ______
  Diagnosis: [ ] balanced  [ ] cross-rich/div-starved  [ ] div-rich/cross-starved

Errors scoreable yet?  [ ] no (responses used)  [ ] yes (errors used - stronger)
Collision test result: _____________________________________________
S = ln(1 + (N-1)u) = ___________ nats
Previous quarter's u = ______   trend: [ ] rising  [ ] flat  [ ] DECAYING
```

39.14.5 C.5 Protocol for C — the dissipation budget

Principle. `C_t` must be in the same unit as `Q_t`, or the headroom equation cannot be evaluated. Commensuration is the hard step and it must be done explicitly and once.

1. **Enumerate the five cost ledgers.** (a) Coordination: time aligning rather than producing. (b) Friction and rework: work done twice, handoff loss, waiting. (c) Maintenance: effort holding `Q` at its current level. (d) Conflict and emotional labor. (e) Overheads: rent, commute, compliance, tooling, licensing.
2. **Measure coordination directly:** meeting-hours times participants over the period. Always available, always larger than expected.
3. **Check against the Brooks bound.** Paths scale as `n(n−1)/2`. If observed coordination cost is not growing at roughly that rate as the roster grows, the roster is not communicating or the measurement is missing channels.
4. **Separate maintenance from growth in ledger (c).** Maintenance belongs in `C`; work that adds new durable structure belongs in `eta*S`. The dividing question is: if this work stopped, would `Q` fall, or would it merely stop rising? Falling is maintenance.
5. **Commensurate.** Choose one conversion and write it down: hours to `Y`-units, or currency to `Y`-units, or a direct expression of all costs in `Y`-units. State the conversion factor, its basis, and its uncertainty. **This is a modelling choice; badge it in the report.**
6. **Sum to `C_t`** and record the ledger breakdown, not just the total. The breakdown is what makes the number actionable.
7. **Compute the utilization check.** Estimate `rho`, the fraction of capacity in use. If `rho > 0.85`, note that queueing delay is growing as `1/(1−rho)` and that a portion of ledger (b) is a utilization artifact that will not respond to process improvement.
8. **Test admissibility.** Verify `C_t <= (1−delta)*Q_t + eta*S_t`. If it fails, the system is being invoiced for order it no longer holds, and the finding outranks every other recommendation in the survey.
9. **Test the sustainability ceiling.** Verify `C_t <= eta*ln N − delta*Q_t`. If it fails, record in the report, in these words: *no choice of interaction geometry can prevent order from decreasing in this system.*

```
WORKSHEET C-5 — C
System: _____________________  Period: _______________
Commensuration rule (badged): 1 _______ = _______ Y-units.  Basis: _____________

  (a) coordination:      _______ hours x _______ people = _______ -> _______ Y-units
  (b) friction/rework:   ____________________________________ -> _______ Y-units
  (c) maintenance only:  ____________________________________ -> _______ Y-units
      (growth work excluded and moved to eta*S?  [ ])
  (d) conflict/emotional:____________________________________ -> _______ Y-units
  (e) overheads:         ____________________________________ -> _______ Y-units

C_t = ___________ Y-units per period

Brooks check: n = _______  n(n-1)/2 = _______ paths.  Coordination per path: _______
Utilization rho = _______   [ ] > 0.85 -> queueing artifact present in (b)

ADMISSIBILITY  C_t <= (1-delta)Q_t + eta*S_t ?
  (1-_______)*_______ + _______*_______ = _______   vs C_t = _______   [ ] pass [ ] FAIL
SUSTAINABILITY C_t <= eta*ln N - delta*Q_t ?
  _______ - _______ = _______   vs C_t = _______   [ ] pass [ ] FAIL
If FAIL, write in the report: "No choice of interaction geometry can prevent
order from decreasing in this system."
```

39.14.6 C.6 Protocol for Q — stored order and its proxy

Principle. `Q` is a modelling construct. `Y` is what is measured. **The single most common measurement failure in this framework is an order proxy that is a repackaging of the input** (CCP-007 §6.3), and this protocol exists mostly to prevent it.

1. **State what persists in this system.** Not what happens — what remains when everyone goes home. Written procedure, built object, trained capability, standing relationship, maintained record, physical arrangement.
2. **Propose two or three candidate proxies `Y`.** Each must be countable, each must be countable by someone other than the analyst, and each must have a unit.
3. **The repackaging test.** Could `Y_t` be computed from `U_t` alone? If yes, reject it. A "number of cross-functional collaborations" proxy is the input wearing a hat.
4. **The persistence test.** Would `Y` still be positive after a period of zero activity? If it goes to zero when activity stops, it is an activity measure, not a storage measure.
5. **The clerical test.** Is `Y` maintained by an administrator regardless of substance? Registers, headcounts, and compliance records fail: they do not decay, so they cannot measure something that does.
6. **The independent-observer test.** Two observers count `Y` for the same period without conferring; if they disagree by more than 20%, tighten the definition first.
7. **Preregister and hold out.** Split the time series into a training span and a held-out span, fix the split before fitting, fit on training, freeze, and report on held-out.
8. **Run the four required baselines** (CCP-015): persistence (`Y_{t+1} = Y_t`), direct-input (`Y_{t+1}` from raw interaction count without the log form), matched autoregressive (`Y_{t+1} = a*Y_t` alone), and the shuffled-input control (`U` series randomly permuted).
9. **Apply the three falsification rules.** If the model does not beat persistence on held-out data, it has failed. If it does not beat the matched autoregressive baseline, the surplus term is doing no work. **If it does not beat the shuffled-input control, the timing of interaction carried no information and the fit is describing decay alone.**
10. **Record `Q_0`** as the value of `Y` at the survey's baseline, and carry it to the headroom sheet.

```
WORKSHEET C-6 — Q AND ITS PROXY
System: _____________________
What persists here: ____________________________________________________________

Candidate Y_1: _____________________  unit: _____________
  repackaging test (computable from U?)  [ ] pass [ ] REJECT
  persistence test (survives zero activity?) [ ] pass [ ] REJECT
  clerical test (not administrator-maintained?) [ ] pass [ ] REJECT
  observer agreement: obs A = _______ obs B = _______ diff _______% [ ] <20%

Candidate Y_2: _____________________  unit: _____________
  ... same four tests ...

SELECTED Y: _____________________  unit: _____________  Q_0 = ___________

Training span: _______ to _______   Held-out span: _______ to _______
Split fixed before fitting?  [ ] yes  [ ] NO -> invalid

HELD-OUT RESULTS (error metric: _____________)
  storage-law model        _______
  persistence baseline     _______   beaten? [ ]
  direct-input baseline    _______   beaten? [ ]
  matched AR baseline      _______   beaten? [ ]
  SHUFFLED-INPUT control   _______   beaten? [ ]   <- if not, model has failed
                                                      if SO, see C-8: beating it proves less
```

than this worksheet once implied

Correction, binding on this worksheet: beating the shuffled-input control is necessary and it is not sufficient. It has three measured holes — anti-conservative calibration, blindness to repackaging, blindness to a shared slow driver — and Protocol C-8 states them, prices them, and gives the ceiling test and the four design controls that cover what the surrogate cannot. **Do not report a `p` from an uncalibrated surrogate.**

39.14.7 C.7 Protocol for the gauge — computing h, r_max, and Q_inf

1. **Assemble** `eta` (C-2), `delta` (C-1), `N` (C-3), `u` (C-4), `C_t` (C-5), `Q_t` (C-6). Confirm all are in the fixed units.
2. **Compute the ceiling:** `eta*ln N`.
3. **Compute maintenance:** `delta*Q_t`.
4. **Compute headroom:** `h_t = eta*ln N - delta*Q_t - C_t`. Set `r_max = h_t`.
5. **Classify the regime:** `h_t < 0` unsustainable; `h_t = 0` maintenance only; `h_t > 0` growth available.
6. **Separately test whether order is actually rising:** `eta*S_t > delta*Q_t + C_t`, with `S_t` computed at the *actual* `u`, not at `u = 1`. **A system can have positive headroom and be shrinking.**
7. **If it is shrinking, compute the required `u`:** `S_req = (delta*Q_t + C_t)/eta`, then `u_req = (e^{S_req} - 1)/(N-1)`. If `u_req > 1`, the geometry cannot deliver it and the answer is `C` or `Q`, not people.
8. **Compute the steady state:** `Q_inf = (eta*F(u*) - c*)/delta`. If negative, record that the regime is unsustainable or the linear model must be replaced.
9. **Compute the uniform bound:** `Q_t <= max{Q_0, eta*ln N/delta}`.
10. **Propagate the intervals.** Recompute steps 2 through 9 at `delta_low` and `delta_high` and report `h` as an interval. **A headroom figure without an interval is a decoration.**
11. **Check every promise against `r_max`.** No plan may request `r_t > h_t`. If one does, the plan is not ambitious; it is arithmetically infeasible, and saying so is the engineer's job.

```
WORKSHEET C-7 — THE GAUGE
System: _____________________  Period: _______________  Y-unit: _____________

eta = _______  delta = [_______ , _______]  N = _______  u = _______
Q_t = _______  C_t = _______

ln N = _______            eta*ln N = _______            (the ceiling)
delta*Q_t = _______ to _______                          (maintenance)
h_t = _______ to _______                                (HEADROOM)  r_max = h_t

REGIME:  [ ] h<0 unsustainable  [ ] h=0 maintain  [ ] h>0 growth available

ACTUALLY RISING?  S_t = ln(1+(N-1)*_______) = _______ nats
  eta*S_t = _______   vs   delta*Q_t + C_t = _______   [ ] rising  [ ] FALLING

If falling: S_req = _______   u_req = (e^S_req - 1)/(N-1) = _______
  u_req <= 1?  [ ] yes, geometry can fix it  [ ] no -> cut C or shed Q

Q_inf = (eta*F(u*) - c*)/delta = _______   [ ] positive  [ ] NEGATIVE
Uniform bound: max{Q_0, eta*ln N/delta} = _______

Requested increment r_t = _______    r_t <= h_t ?  [ ] yes  [ ] NO -> INFEASIBLE
```

39.14.8 C.8 The shuffled-input control — three measured holes, and what to run instead

This protocol is a correction to the six above, and to every worksheet in this appendix that carries a `SHUFFLED-INPUT CONTROL` line. The control is still required. What has changed is what passing it is worth, and the change was measured rather than argued.

Three failure modes have been computed by simulation. All three are properties of the surrogate method itself, not of any particular study, and **no fourth surrogate fixes any of them.**

```
HOLE 1  ANTI-CONSERVATIVE AS ORDINARILY RUN                          [R13a, CCP-027]
        Measured false-rejection rate against a nominal alpha = 0.05, 500 null trials:

            iid permutation       0.108      2.2x nominal
            circular shift        0.120      2.4x nominal
            block permutation     0.060      near nominal

        A surrogate test run off the shelf roughly DOUBLES the false-positive rate.
        The matched-autoregressive comparison ALONE false-passed 35.8% of the time.

HOLE 2  BLIND TO REPACKAGING                                         [CC-R4, CCP-015]
        A filtered repackaging of the input passed all four baselines INCLUDING the
        shuffled-input control at p = 0.0005. It was caught only by the structural
        ceiling test (CCP-007, Cor 4.2): its single-step jump was 38.5x r_max.

HOLE 3  BLIND TO A SHARED SLOW DRIVER                                [R24, CCP-035-81]
        4,000-trial simulation, reported honestly in both directions:
            clean case (Case 3)          48.17% detection -- it works
            shared slow driver (Case 4)  98.35% PASS      -- it does not

        A common confounder moving both series slowly is invisible to permutation,
        because permutation destroys timing and the confounder does not act through timing.
```

Consolidated rule, binding on Protocols C-6 and C-7 and on every chapter that recommends the four-baseline suite: the shuffle is a weak, mis-calibrated empirical audit. Repackaging is caught only by the structural test. Confounding by a shared slow driver is caught only by design. **A passed shuffle is a floor, not a certificate.**

Step 1 — calibrate the surrogate on the design in hand, before using it. Simulate the study under a stated null — same period length, same series length, same autocorrelation, same missingness — and measure the actual false-rejection rate of the surrogate to be run. Use the calibrated threshold, not `0.05`. **An uncalibrated surrogate is reporting somebody else's alpha.**

Step 2 — run the structural ceiling test, every time. Compute `r_max = eta*ln N - delta*Q_t - C_t` (Protocol C-7) and compare it to the largest single-period increment in the observed series. **Any increment above `r_max` is impossible under the model, and no empirical baseline will report it.**

```
WORKSHEET C-8a — CEILING TEST
Largest observed single-period increment in Y          ______
r_max from C-7 (at the SAME period and units)          ______
ratio  observed / r_max                                ______
ratio > 1 ?  [ ] no, admissible   [ ] YES -> the proxy is a repackaging until proved otherwise
```

Step 3 — buy the remaining protection with DESIGN, not with more surrogates. These four are the whole of it, they are cheap relative to a wasted study, and each is checkable by an outsider.

Design control	What it defeats	How it is verified
Blinding	fabricated and nudged timing; rater contamination	The rater cannot state the period's input when asked
Channel separation	the order proxy being a relabeled input	`Y` and `U` collected by different people, by different instruments, on different days
Prospective sampling	survivorship in the case series	Every unit enrolled **before** its outcome is known; the roster is fixed and dated
Yoked control	shared slow drivers; secular trend	A matched unit exposed to the same environment and not to the intervention

```
WORKSHEET C-8b — DESIGN CONTROLS
(attach to the preregistration; unsigned rows are absent controls)

Blinding               who is blind, to what, and how it was checked  _____________  [ ] none
Channel separation     Y collected by _______  U collected by _______   [ ] same source -> INVALID
Prospective sampling   enrolment date rule _____________                [ ] retrospective
Yoked control          matched unit, matching variables _____________   [ ] none

Surrogate calibrated on this design?  [ ] yes, measured alpha = _______  [ ] NO -> do not report p
Ceiling test run (C-8a)?              [ ] yes  [ ] NO -> repackaging not excluded

Signed: _____________________  Date: _____________
```

If the design controls are absent, say so in the report in the sentence that states the result. A study that ran the shuffle and nothing else has excluded exactly one thing: that the model is describing persistence only. It has not excluded repackaging and it has not excluded confounding.

39.14.9 C.9 The gradient shortcut — a habit survey IS a headroom measurement

Derived in CCP-022 (R19): `grad h = -grad C`. Where `eta`, `N` and `delta` are constant in the direction of motion, headroom and cost have exactly opposite gradients — so **the direction of steepest improvement can be read from what people actually do, without running Protocol C-1 through C-6 first.** the founding material's desire-line node and this framework's master gauge measure the same thing.

This does not replace the gauge. It locates the direction cheaply so the gauge can be run once, in the right place, instead of four times in the wrong ones.

1. **Fix the restriction, in writing.** Name the region over which `eta`, `N` and `delta` are being treated as constant. **If the candidate move crosses a boundary that changes the roster, the tooling or the retention regime, the identity does not hold and this protocol does not apply.**
2. **Observe the worn path, do not ask for it.** Record what is actually done — the route taken, the tool reached for, the meeting attended without a summons, the step skipped. Revealed behaviour, not stated preference.
3. **Rank candidate moves by observed local cost**, in the same unit as `C` in Protocol C-5: minutes, hops, approvals, handoffs, discomfort reports. A ranking is sufficient; a magnitude is not required at this stage.

4. **Invert the ranking.** The direction of steepest cost decrease is the direction of steepest headroom increase. That is the identity, and it is the whole shortcut.
5. **Then run C-7 once, in that direction only**, and check `r_t <= h_t` before promising anything.
6. **Record the restriction violation risk.** If the intervention would itself change `eta`, `N` or `delta` — most reorganisations do — the shortcut located the direction under conditions the intervention then destroys. Note it; re-measure after.

```
WORKSHEET C-9 — GRADIENT SHORTCUT
Region over which eta, N, delta are constant: _______________________________
Cost unit (same as C-5): _____________

| Candidate move | Observed local cost | Rank | Crosses the boundary? |
|---|---|---|---|
|                |                     |      | [ ] yes -> EXCLUDE     |

Steepest cost decrease -> direction of steepest headroom increase: ____________________
C-7 run in that direction?  [ ] yes, h_t = ______   r_t requested = ______  feasible? [ ]
```

39.14.10 C.10 The t∗ test — the cheapest falsification in the corpus

t∗ has no free parameters, which is what makes it decisive. With `D`, `p_bar`, `N` and the budget all held constant, decay of `u` **alone** drives the calibration ratio across 1 at

```
u∗ = ( D/p_bar − 1 ) / ( N − 1 )            t∗ = ln( u∗ / u_0 ) / ln( 1 − lambda )
```

Worked, and verified: `D/p_bar = 4.0`, `N = 12`, `u_0 = 0.55`, `lambda = 0.08` gives `u∗ = 0.2727` and **`t∗ = 8.41` periods.** After `t∗`, the organisation is over-framed with an entirely unchanged model, **by arithmetic alone and without anyone deciding anything.**

1. **Estimate the four inputs, none of which is a fitted parameter.** `D` and `p_bar` from the existing intake ledger; `N` from Protocol C-3; `u_0` from C-4 at the current period; `lambda` as the observed decay rate of `u` between two C-4 measurements.
2. **Compute `u∗` and `t∗`.** Both are arithmetic. There is nothing to tune and nothing to fit, which is why this test costs a morning rather than a study.
3. **Record `t∗` as a dated alarm**, in the same document as the preregistration. It is a prediction with a number and a deadline attached, and by the CCP-035 arithmetic that makes it worth about **5.7x** what a direction-only prediction is worth.
4. **At `t∗`, measure `u` again and re-run C-4.** If `u` has decayed as projected and the doctrine has not moved, the prediction held. If `u` held, something is maintaining it and that something is worth naming, because it is the cheapest `u`-maintenance mechanism available.
5. **Never aggregate `r` across sites.** Aggregate `D/P` cancels opposite-direction distortion: in CCP-024's two-site example it **overstated bankable input by 4.26x.** The aggregate is not conservative and not approximately right — **distortion does not average.** Compute `t∗` per site.

```
WORKSHEET C-10 — t∗
D = ______  p_bar = ______  N = ______  u_0 (current) = ______  lambda = ______

u∗ = (D/p_bar − 1)/(N − 1) = ______
t∗ = ln(u∗/u_0) / ln(1 − lambda) = ______ periods   -> ALARM DATE: _____________

Per site, never aggregated:  site ______ t∗ = ______   site ______ t∗ = ______
At t∗: u measured = ______   [ ] decayed as projected  [ ] held -> name the mechanism
```

39.15 Appendix D Failure-mode index

The reader arrives with a symptom, not a diagnosis. This index is therefore keyed by what is actually observed — the sentence someone says in a corridor, the number that looks wrong, the thing that keeps happening — and points to the named entry in CCP-016, or — for a claim about manifestation rather than a built system — to the named entry in CCP-037.

How to use it. Find the symptom closest to yours and run the confirming measurement in the third column *before* accepting the diagnosis; symptoms are not specific, and several failure modes share most of theirs. If two entries fit, run both measurements — co-occurrence is the norm, and D.11 lists the combinations that travel together.

39.15.1 D.1 Symptoms in deliberation and meetings

Observed symptom	Suspect	Confirm by	CCP-016 entry	First move
“The meeting went really well” and nothing changed	`eta`	C-2 stratified by capture mechanism	**The Leaky Vessel**	Require one written artifact per meeting
Everyone agrees quickly and the decision is still wrong	`u`	C-4 collision test, error decorrelation	**The Parallel Wall**	Silent-write-first before discussion
Two or three people do almost all the talking	room / `c`	Turn-distribution count; RT60	**The Reverberant Room**	Measure RT60 before blaming anyone
The quiet expert stops attending	room / `u`	Attendance vs turn share by person	**The Reverberant Room**	Absorption and table geometry
Discussion converges within five minutes, every time	cascade	C-4 measured before vs after discussion	**The Cascade Collapse**	Independent written estimates first
The first speaker’s framing survives every meeting	cascade	Order-of-speaking vs final-position log	**The Cascade Collapse**	Rotate speaking order; rank speaks last
Objections are raised only in the corridor afterward	`u_div`	Anonymous vs attributed disagreement rate	**The Cohesion Death Spiral**	Chartered red team; protected minority report
Meetings have grown but decisions have not	`C`	Meeting-hours x participants; `n(n−1)/2`	**The Quadratic Committee**	Split the group; compute the crossover

Observed symptom	Suspect	Confirm by	CCP-016 entry	First move
Nobody can say what was decided	`eta`	Existence of a decision log	**The Leaky Vessel**	Decision log with owner and date
The senior person speaks first and last	geometry	Seating map; turn order log	**The Head of Table**	Remove the head seat; invert order
“We have a very aligned team” said as praise	`u`	C-4	**The Cohesion Death Spiral**	Measure `u_div` this quarter and last
New members stop dissenting after two months	`u` decay	`u` by tenure cohort	**The Parallel Wall**	Rotation; protect the first dissent publicly
The agenda is always full and never finished	`C`, `rho`	Utilization estimate	**The Utilization Cliff**	Cut agenda load to `rho <= 0.85`
A decision is made and then unmade elsewhere	structure	Map decision rights to categories	**The Tree That Should Have Been a Semilattice**	Grant overlapping membership

39.15.2 D.2 Symptoms in judgment, forecasts, and decisions

Observed symptom	Suspect	Confirm by	CCP-016 entry	First move
Forecasts are confident and jointly wrong	`u`	Error correlation across forecasters	**Pseudo-Diversity**	Recruit across training lineages
The panel is demographically varied and never disagrees	`u_div`	C-4 within-category	**Pseudo-Diversity**	Buy divergence, not more crossing
Adding reviewers has stopped improving decisions	concavity + `C`	Marginal-gain table; C-5	**The Quadratic Committee**	Stop adding; raise `u` instead
The second opinion always confirms the first	`u = 0`	Training lineage and incentive audit	**The Parallel Wall**	A second opinion from a different lineage
The risk register never surprises anyone	`u_div`, variety	Requisite-variety count vs disturbance count	**The Missing Denominator**	Pre-mortem; count the unanswered classes
The dissenter left and no one replaced the function	`u_div`	`u_div` before and after departure	**The Cohesion Death Spiral**	Charter the role, not the person

Observed symptom	Suspect	Confirm by	CCP-016 entry	First move
Losing one person would stop the work	hubs	Node-removal thought experiment on the work map	**Hub Fragility**	Duplicate the path, not the person
“We considered all the options” and there were three	`N`	C-3 audit; `N_nominal` vs `N`	**The Missing Denominator**	Write what a fourth category would be
Two disciplines both say the case is the other’s	seam	Map the categories the clerks actually enforce	**The Tree That Should Have Been a Semilattice**	Assign the seam explicitly
Adding a new “domain” did not help	`ln N`	`ln N` before and after	**The Added Domain**	Compute `ln 2 = 0.693`; look elsewhere
The proposal promises a step no one has sized	`h`	C-7	**The Promised Increment**	Compute `r_max`; refuse `r_t > h_t`

39.15.3 D.3 Symptoms in output, delivery, and throughput

Observed symptom	Suspect	Confirm by	CCP-016 entry	First move
Everyone is busy and nothing ships	`rho`, `C`	Utilization; queue lengths; Little’s law	**The Utilization Cliff**	Build slack to `rho <= 0.85`
Delivery time doubled after a small load increase	`rho`	`1/(1−rho)` fit against observed wait	**The Utilization Cliff**	Reduce load or add capacity now
Adding people made the project later	`C`	Brooks path count; onboarding cost	**The Quadratic Committee**	Stop adding; partition the work
Work is done twice in two departments	structure	Duplicate-output audit	**The Silo Local Minimum**	Shared instrumentation; co-location
Nobody knows what the other team is doing	structure, `C`	Cross-team contact frequency	**The Silo Local Minimum**	Rotation; translator role
The architecture mirrors the org chart, badly	Conway	Compare module graph to comms graph	**The Silo Local Minimum**	Change the team boundary first
Improvement efforts produce a spike and then nothing	impulse response	Plot `Y` against `3*tau`	**The Impulse Habit**	Convert the event into a cadence

Observed symptom	Suspect	Confirm by	CCP-016 entry	First move
The initiative was excellent and is now invisible	`delta`	C-1 on the initiative's own artifacts	**The Leaky Vessel**	Maintenance schedule at `t_half`

39.15.4 D.4 Symptoms in the physical room

Observed symptom	Suspect	Confirm by	CCP-016 entry	First move
The new open plan is loud and no one talks	acoustics	RT60; background NC/NR	**The Reverberant Room**	Absorption; enclosed alternatives
The plaza is empty at lunchtime	sittable space	Whyte count: seating linear metres, sun, food	**The Dead Perimeter**	Add sittable edge before anything else
The middle of the room is never used	edge effect	Occupancy heat map over a week	**The Dead Perimeter**	Furnish the perimeter; subdivide the middle
Everyone crowds the doorway	threshold depth	Dwell-time at each transition	**The Abrupt Edge**	Deepen and widen the threshold
Entering the space is jarring	threshold	Sequence audit: outside -> inside in one step?	**The Abrupt Edge**	Insert a graded transition zone
The thermostat is a permanent argument	thermal control	Individual control count; comfort survey spread	**The Optimized Setpoint**	Give local control; stop optimizing the mean
The lighting is uniform and everyone squints	glare, daylight	Illuminance survey; glare check at desks	**The Optimized Setpoint**	Task lighting plus daylight control
The building looks worse every year	material `delta`	Material audit for patina behaviour	**The Invisible Infrastructure**	Specify materials that age well
The maintenance budget is cut every year it works	`delta` invisibility	Compare spend to computed `delta*Q`	**The Invisible Infrastructure**	Publish `delta*Q` as a standing line
A long table with a clear top seat	geometry	Seating map; turn share by seat	**The Head of Table**	Round or polygonal table; no top seat

39.15.5 D.5 Symptoms in institutions over time

Observed symptom	Suspect	Confirm by	CCP-016 entry	First move
Everything takes longer than it used to and nothing is wrong	`delta*Q`	C-7: `delta*Q` as a share of `eta*ln N`	**The Maintenance Trap**	Publish headroom; consider shedding `Q`
The organization cannot start anything new	`h <= 0`	C-7	**The Maintenance Trap**	Shed `Q` or cut `C`; do not exhort
Procedures exist that nobody can justify	unshed `Q`	Origin audit; last-use date per procedure	**The Unshed Q**	Sunset clauses; retire on schedule
The founding generation left and it still runs the same	`delta \ensuremath{\approx} 0`, no input	`Y` trend vs `S` trend	**The Sealed Vault**	Restore input; measure `u`
Tradition is intact and nothing new has entered in years	`delta \ensuremath{\approx} 0`	C-1 (expect very low `delta`); C-4 (expect low `u`)	**The Sealed Vault**	Deliberate `u` injection; accept some `delta`
The metric improved and the thing it measured did not	Goodhart	Compare metric to an unrewarded proxy	**Goodhart Capture**	Separate measurement from reward
A crisis revealed a dependency nobody had listed	infrastructure	Breakdown inventory vs maintained inventory	**The Invisible Infrastructure**	Inventory by simulated failure
Growth stalled exactly when the institution matured	headroom	Plot `h_t` against `Q_t` over years	**The Maintenance Trap**	Accept lower `Q`; this is a legitimate move
The change was seeded everywhere and took nowhere	contagion topology	Map seed placement against clustering	**The Sprayed Seed**	Re-seed one cluster to saturation
A minority has pushed for years without moving anything	critical mass	Estimate committed fraction	**The Sub-Threshold Minority**	Concentrate, do not broaden
The commons is degrading and everyone agrees it shouldn't	commons	Ostrom principle-by-principle audit	**The Ungoverned Commons**	Boundaries, monitoring, graded sanctions

Observed symptom	Suspect	Confirm by	CCP-016 entry	First move
The steady state is worse than the current state	`Q_inf < Q_t`	C-7 `Q_inf`	**The Negative Equilibrium**	Raise `eta` or cut `c*`; the model says so
Every change requires someone's permission and stalls	centre	Map decision latency vs approval hops	**The Coordinator Bottleneck**	Put the state in the medium
People were moved without being asked	ethics	Consent audit	**The Coerced Move**	Stop; the standing refusals of the corpus forbids it

39.15.6 D.6 Symptoms in the measurement itself

Observed symptom	Suspect	Confirm by	CCP-016 entry	First move
The model fits beautifully and predicts nothing	repackaging	Is `Y` computable from `U`?	**The Order Proxy That Was Really the Input**	Reject the proxy; return to C-6
`delta` came out near zero	contaminated window	Re-check zero-input contamination	**The Contaminated Window**	Find a genuine zero-input window
The shuffled control fits as well as the real series	no signal	C-6 step 8	**The Missing Control**	Report failure; do not proceed
Nobody ran the shuffled control	omission	Check the preregistration	**The Missing Control**	Run it before any claim
The model was adjusted after seeing the held-out data	freeze violation	Compare the fit spec to the preregistration	**The Unfrozen Fit**	Discard; re-preregister; re-run
Headroom was reported as a single number	interval omission	Was `delta` an interval?	**The Missing Denominator**	Propagate the `delta` interval
An assumption is being used as a result	badge	Search the text for `1+(N-1)u` without S3 named	**The Unbadged Postulate**	Badge it in the sentence that uses it
A borrowed lab finding is quoted as a field result	transfer	Check the source's actual setting	**The Unbadged Postulate**	State the setting and the transfer assumption

Observed symptom	Suspect	Confirm by	CCP-016 entry	First move
A number in the report has no derivation	fabrication	Ask for the arithmetic	**The Unbadged Postulate**	Remove it or compute it (the standing refusals of the corpus, refusal 3)
The shuffled control was passed and nothing else was run	false comfort	C-8: was the surrogate calibrated on this design?	**The Missing Control**	Calibrate the surrogate; run the C-8a ceiling test
A p is quoted from an off-the-shelf permutation test	anti-conservative alpha	C-8 hole 1: measured 0.108 / 0.120 against a nominal 0.05	**The Missing Control**	Simulate the null on the design in hand; requote
Both series drift together and the shuffle passed	shared slow driver	C-8 hole 3: 98.35% pass in 4,000 trials	**The Missing Control**	Yoked control; no further surrogate will help

39.15.7 D.7 Symptoms in a claim about manifestation — channel and magnitude

A different index, for a different object. D.1 through D.6 diagnose a built system. The four tables that follow diagnose a *claim* — that some belief, held and acted on, produced some change — and they point to CCP-037, whose thirty entries are keyed to the eight coordinates every such claim has: **CH** channel, **G** gain, **Mg** magnitude, **B** base rate, **Fa** falsifiability, **Cn** consent, **Ld** ledger, **Sc** scale. **The fault is located at the first coordinate the claim cannot pass**, and locating it there rather than at the end is what makes the catalogue usable.

Four of the thirty are faults of this framework's own method, and they are indexed here on exactly the same terms as anybody else's: MF-08, MF-11, MF-23 and MF-24. **A catalogue that indexes only other people's errors is advertising.**

Observed symptom	Coordinate	Confirm by	CCP-037 entry	First move
Asked "by what route?", the holder restates the belief	CH	Ask for the channel in one sentence, nameable and lookup-able	**MF-01 · The Ungrounded Belief**	CCP-035 step 2; name a channel or withdraw the claim
The named channel is precise enough to look up, and the lookup returns nothing	CH	Look it up. The saving is the point	**MF-02 · The Absent Channel**	Halt at G-1; a named absent channel is cheaper than a vague one

Observed symptom	Coordinate	Confirm by	CCP-037 entry	First move
A sentence begins with a Nobel-recognized mechanism and ends somewhere it does not license	CH	Locate the seam — the clause where the citation stops covering the sentence	**MF-03 · Channel Substitution**	Split the claim; each half names its own transducer
The citation is real and the outcome is not in it	G	Read the cited paper's own outcome measures	**MF-04 · The Zero-Gain Channel**	`X = G·B` with `G = 0`; identical in result to MF-01
"It works in cells, therefore in tissue"; "at ultrasound, therefore at audio"	G	Check the transfer explicitly: scale, medium, frequency, species	**MF-05 · The Borrowed Channel**	State the transfer as an assumption and price it
The mechanism is quoted correctly and no number ever appears	Mg	Ask for the predicted size, with a tolerance, before measuring	**MF-06 · The Magnitude Error**	CCP-035 step 5; where Casimir is invoked, compute it
Every reported outcome is compatible with the claim	Mg	Ask what result would have been reported as a failure	**MF-07 · Order-of-Magnitude Silence**	Require a magnitude and a null before the next attempt
A decimal appears where the instrument permits only an ordering	Sc	FVP gate 2 on each term; is there a true zero?	**MF-08 · The Demoted Magnitude Reported As Measured**	Report the ordering only; Rasch route, 123 items per term
Arithmetic performed on a Likert response as though it were a length	Sc	Stevens' typology per term; which operations does the scale permit?	**MF-09 · The Scale Violation**	Restrict to permitted operations, or re-instrument
Effort and perceived result compared as though the relation were linear	Mg, Sc	Is the response register logarithmic? (CCP-020, the lens theorem)	**MF-10 · The Dosage Fallacy**	Design in ratios; equal perceived steps are equal world multiples
Results track the input's fine structure too closely	Mg	**C-8a ceiling test. The shuffle cannot see this**	**MF-11 · The Repackaged Input**	Reject the proxy; channel separation in the next design

Observed symptom	Coordinate	Confirm by	CCP-037 entry	First move
A step change with no corresponding change in `eta`, `u`, `N` or `C`	Mg	`r_max = eta*ln N − delta*Q − C` against the observed increment	**MF-12 · The Ceiling Violation**	Recompute `r_max`; an increment above it is impossible under the model

39.15.8 D.8 Symptoms in a claim about manifestation — base rate and inference

Observed symptom	Coordinate	Confirm by	CCP-037 entry	First move
An impressive collection of individual accounts and no denominator anywhere	B	Ask who is missing: who tried and did not report	**MF-13 · The Survivorship Narrative**	**Likelihood ratio 1.000000.** Enrol before the outcome is known
Counts without rates — “thirty people adopted the practice”	B	Ask for `N` at the **start** of the period, not the end	**MF-14 · The Missing Denominator**	Fix and date the roster; attrition moves people out of view
Recruitment framed around a low point; large early gains that flatten	B	Was selection on an extreme value? Compute the expected return	**MF-15 · Regression Mistaken For Effect**	Yoked control; recruit independently of the current level
A specific, narrow, unanticipated subgroup result	B, Fa	Compare the reported analysis to the preregistration	**MF-16 · The Forking Path**	Preregister the subgroup, or report it as exploratory
“We tried it for two weeks and it did not work”, no interval reported	B	Power: what effect could that design have detected?	**MF-17 · The Underpowered Null Read As Disconfirmation**	Report the interval; an underpowered null disconfirms nothing
Nine things changed and one gets the credit	CH, B	Was the credited component ever varied on its own?	**MF-26 · The Undecomposed Bundle**	Vary one component; until then the bundle is the claim

39.15.9 D.9 Symptoms in a claim about manifestation — falsifiability and the ledger

Observed symptom	Coordinate	Confirm by	CCP-037 entry	First move
Success and failure are described in the same vocabulary	Fa	Ask for the result that would kill the claim, in writing	**MF-18 · The Unfalsifiable Goal**	State the disconfirming result before the next attempt
The stated aim at the end does not match the stated aim at the start	Fa	Compare the final write-up to the dated preregistration	**MF-19 · The Retrofitted Criterion**	Discard; re-preregister. The retrofit selects harder than forking
Every null is answered with patience and no null is answered with a date	Fa	Demand a bounded latency at step G-4, in the same breath as the magnitude	**MF-20 · The Latency Excuse**	Attach a date to the prediction or withdraw it
Results are reported as “encouraging” and never as “null”	Fa	Is there a smallest effect worth having? A sample size fixed in advance?	**MF-21 · The Un-nulled Protocol**	Define the null before the next observation
The metric improves smoothly while every unmeasured correlate does not	Fa	Compare the rewarded metric to an unrewarded proxy	**MF-22 · Goodhart Capture of a Manifestation Metric**	Separate the measurement instrument from the reward
The phrase “both science and tradition confirm”	Ld	Which column is each source in? Only one of them carries evidence	**MF-23 · Precedent Mistaken For Evidence**	**The single most likely failure of this volume’s method.** State each column’s role
A claim hedged on first appearance and flat on second	Ld	Compare the abstract to this paper that stated it	**MF-24 · The Unbadged Speculative Bridge**	Restore the badge in the sentence, not in a note attached to it
Increasing terminological precision with no increase in predictive content	CH, Ld	Does the new vocabulary generate a prediction the old one did not?	**MF-25 · Naming As Mechanism**	Ask what changes numerically; a name is not a transducer

39.15.10 D.10 Symptoms in a claim about manifestation — consent

These four are the entries this framework will not trade against a result, and they are indexed last so that they are the last thing read.

Observed symptom	Coordinate	Confirm by	CCP-037 entry	First move
A step described to observers and not to participants	Cn	Count gate records against participants	**MF-27 · The Consent Bypass**	Stop. the standing refusals of the corpus, refusal 4
A decline rate of exactly zero across many opportunities	Cn	The decline log; unanimity that arrives too easily	**MF-28 · The Recruited Belief**	Stop. A zero decline rate is a measurement of pressure
The rollback plan does not exist, or exists and does not mention residue	Cn	Ask what is left behind after the rollback	**MF-29 · The Irreversible Step**	No irreversible step without a costed, stated rollback
A file, track or step whose full content the recipient cannot inspect	Cn	Can the recipient inspect the whole artifact?	**MF-30 · The Masked Carrier**	Full disclosure of content. No concealed carrier, ever

39.15.11 D.11 Co-occurring clusters — differential diagnosis

Failure modes travel in company. These five clusters account for most real surveys, and recognizing the cluster is faster than diagnosing its members one at a time.

The Comfortable Decline. Cohesion Death Spiral + Pseudo-Diversity + Parallel Wall. Signature: high satisfaction, low turnover, falling `u`, flat or falling `Y`. Feels like success from inside and is the hardest cluster to name aloud, because naming it sounds like an attack on people doing nothing wrong. Lead with the arithmetic, not the atmosphere.

The Ossified Success. Maintenance Trap + Unshed Q + Invisible Infrastructure. Signature: large `Q`, `delta*Q` consuming most of `eta*ln N`, a maintenance budget under annual attack, an inability to start anything. **The intervention is shedding, and it is the one nobody proposes.**

The Busy Nothing. Leaky Vessel + Quadratic Committee + Utilization Cliff. Signature: high activity, high meeting satisfaction, `rho` above 0.9, `eta` near zero, no artifact trail. The cheapest cluster to fix, because `eta` is the cheapest parameter.

The Unfalsifiable Report. Order Proxy That Was Really the Input + Missing Control + Unfrozen Fit + Unbadged Postulate. Signature: an excellent-looking model with no held-out result, no shuffled control, and a proxy that cannot go down. **This is the cluster that discredits the framework itself, which is why CCP-015 exists and why Appendix E is written as it is.**

The Licensed Claim. MF-02 or MF-03 + MF-13 + MF-15 + MF-12 + MF-18 + MF-21 + MF-23 + MF-27 or MF-28. Signature: a real upstream mechanism, a testimonial base with no denominator, recruitment at a low point, a claimed step above `r_max`, no stated null, precedent quoted as a second confirmation, and a decline rate at or near zero. **CCP-037's worked example carries twelve of these diagnoses in one write-up**, which is the normal density rather than a caricature: the coordinates fail together because a claim that skips one gate has usually skipped the gate before it. Enter at the channel, not at the outcome — **the diagnosis at the first failed coordinate is the only one that saves any work.**

39.16 Appendix E The honesty register

This appendix is the corpus's warrant. Everything else in the framework is a consequence of a small set of proved theorems, a larger set of chosen modelling assumptions, and a body of empirical results obtained in settings that are not the reader's. This appendix separates those three, item by item, and adds a fourth category — what has not been tested at all — which is larger than the other three combined.

The reason for writing it this carefully is not modesty. It is that a framework which cannot be told apart from its own decoration is useless as engineering. If a reader cannot determine which of this framework's claims would survive a hostile audit, they cannot use it to build anything they would be willing to defend. **This appendix is written to be used against this corpus.**

Two further drawers were added after the volumes that follow this appendix were drafted, and they are the two a hostile reader should open first. E.7 holds what this corpus **demoted**: claims it advanced and its own validation protocol reduced. E.8 holds what this corpus **rejected**: claims worked end to end in print and failed, three of them drawn from the corpus this corpus is most often assumed to be defending. **A framework's warrant is not what it asserts. It is what it has thrown away, in print, at its own cost.**

Six drawers, and the drawers are labeled.

First, before the drawers are opened: the demotion **The Formula Validation Protocol demoted the author's own keystone formula, and it did so on the second of seven gates.**

`M = (D/P)·E` — the Moment Equation, the join between the institutional-scope material's storage law and the personal-scope material's account of experience, and the formula from which the calibration ratio, the distortion measure, the launch budget and the return difference all descend — **returns DEMOTE at FVP gate 2, scale type.** Division requires a ratio scale with a true zero. `D` and `P` are not instrumented on ratio scales. Therefore the following are **not reportable as measured quantities anywhere in this corpus**:

- the **magnitude** of `D/P`;
- the distortion `Delta = |ln r|` as a number rather than an ordering;
- the **quantitative** join `M_t -> Q_{t+1}` of CCP-021;
- every M-derived formulary entry, at personal and interpersonal scope alike — five register entries, demoted together, because demotion travels downstream.

What is demoted is magnitude, not structure. The limiting cases stand. The `D = P` fixed point stands. The ordering — over-framed, calibrated, under-framed — stands, because an ordering needs only ordinal information. A reader may use the formula to say *which way* a system is distorted. A reader may not use it to say *by how much*, and neither may this corpus.

And the route out is specific, principled, and priced. Theorem 54.3, given badged Postulate SO1: **a Rasch logit difference makes `Delta = |ln(D/P)|` genuinely ratio-scaled**, because Rasch separation supplies exactly the true zero that division requires. That is not a caveat. **It is the volume's actual research instruction**, and it costs **123 items per term** to establish a half-logit distortion — 342 items for three tenths of a logit.

The coherence with the research programme is not a coincidence and CCP-027 says so: **FVP-7, the cross-subject recovery gate, is both the gate that decides the hard-science question and the stage most likely to kill the programme.**

This is the corpus working as designed. A validation protocol that demotes its own author's keystone formula, states the cost of repairing it, and gives the constructive route out is the strongest available demonstration that the protocol is real rather than decorative. It is placed here, before the drawers, rather than filed in one of them, because a demotion recorded only where it is convenient to find is concealment by filing.

39.16.1 E.1 What is proved

These are theorems. They are true within their stated assumptions, and their stated assumptions are mathematical, not empirical. Nothing in this section is a claim about the world.

CCP-A02 — the surplus functional.

1. **Theorem 2.1, the uniqueness result.** Given axioms S1 (geometric dependence: `F` depends only on `u`), S2 (`f(0) = 0`), S3 (`g(u) = e^{f(u)}` is affine in `u`), and S4 (`f(1) = ln N`), the function `f(u) = ln(1 + (N−1)u)` is the unique function satisfying all four. This is a genuine uniqueness theorem. Its force depends entirely on the four axioms being the right ones, and S3 is not a fact about the world — see E.2.
2. **Theorem 3.1, strict monotonicity.** `u1 < u2` implies `f(u1) < f(u2)`. More non-parallelism yields strictly more surplus.
3. **Theorem 3.2, positivity.** `f(u) > 0` for all `u > 0`.
4. **Theorem 3.3, strict concavity.** `f''(u) < 0` throughout the interior. Each added increment of difference buys strictly less than the previous one. Every marginal-gain table in this framework is this theorem evaluated numerically.
5. **Theorem 3.4, subadditivity on the parameter `u`.** `f(u1 + u2) <= f(u1) + f(u2)`. **This holds on `u`, and only on `u`.**
6. **Theorem 4.1, the two-source decomposition.** For a pure block vector, `u = u_cross + u_div` with `u_cross = beta^2` and `u_div = alpha^2 sin^2(gamma)`, and `max(f(u_cross), f(u_div)) <= f(u) <= f(u_cross) + f(u_div)`. Both bounds are proved; the upper one is the one design must respect.
7. **Theorem 4.2, the sharp Lipschitz constant.** `|f(u1) − f(u2)| <= (N−1)|u1 − u2|`, and the constant `N−1` cannot be improved.
8. **Theorem 4.3, extremality.** `F = 0` if and only if the states are parallel; `F = ln N` if and only if they are orthogonal.

CCP-009 — the range.

9. **Corollary 2.2.** `0 <= F(u) <= ln N`. `ln N` is the hard ceiling on surplus per interaction, and it is what makes the product `eta*ln N` the ceiling term in every design inequality in this framework.

CCP-007 — the storage law's consequences. Theorems *about the linear recurrence*. That the recurrence describes a social system is a modelling assumption (E.2); that these follow from it is proved.

10. **Theorem 3.1, the exact solution.** `Q_t = (1−delta)^t Q_0 + SUM_{k=0}^{t−1} (1−delta)^{t−1−k}(eta S_k − C_k)`. The split into transient and driven terms is exact.
11. **Corollary 3.2, one-step growth.** `Q_{t+1} > Q_t` iff `eta S_t > delta Q_t + C_t`.
12. **Theorem 3.3, the uniform bound.** `Q_t <= max{Q_0, eta ln N / delta}`; if `delta = 0`, `Q_t <= Q_0 + t eta ln N` — the formal statement that a system with no loss term has no fixed point.
13. **Corollary 3.4, the sustainability ceiling.** A nondecreasing step requires `C_t <= eta ln N − delta Q_t`. If this fails, **no choice of interaction geometry can prevent order from decreasing**, because `S_t <= ln N` bounds every geometry at once.
14. **Theorem 3.5, the steady state.** With `U \ensuremath{\equiv} u*`, `C \ensuremath{\equiv} c*`, `delta > 0`: `Q_inf = (eta F(u*) − c*)/delta`. If `eta F(u*) < c*` the equilibrium is negative — the regime is unsustainable or the linear model must be replaced. **The theorem does not hide this failure mode, and neither does this corpus.**
15. **Proposition 4.1, target increment.** `Q_{t+1} − Q_t >= r_t` iff `S_t >= (delta Q_t + C_t + r_t)/eta`, feasible only if `delta Q_t + C_t + r_t <= eta ln N`.
16. **Corollary 4.2, maximum one-step increment.** `r_max = eta ln N − delta Q_t − C_t`.

Borrowed mathematics, correctly used.

17. **Banach's fixed point theorem (1922).** A contraction on a complete metric space has exactly one fixed point; iteration from any start converges to it geometrically. `delta > 0` makes the storage update a contraction with factor `(1−delta)`. Every claim in this framework that iteration *arrives* is this theorem applied to a linear recurrence on the reals. Whether the social system it models is a complete metric space is drawer two.
18. **The discrete-time first-order filter identity.** That `Q_{t+1} = (1−delta)Q_t + x_t` is a low-pass filter with pole `(1−delta)`, time constant `tau = −1/ln(1−delta)`, and DC gain `1/delta` is an identity, not an analogy. The settling figures — 95.0% at `3*tau`, 99.3% at `5*tau` — are computed from it.
19. **Little's Law.** `L = lambda W`, for any stationary system with finite means. Proved.
20. **Ashby's Law of Requisite Variety (1956).** Within Ashby's set-theoretic formulation, a regulator's capacity to hold an outcome constant is bounded by its own variety. A theorem within its framework; its application to organizations is a modelling map — see E.2.
21. **The Condorcet Jury Theorem (1785).** If each of `n` voters is independently correct with probability `p > 1/2`, majority accuracy increases in `n` and tends to 1. **Under correlated votes the theorem gives no such conclusion, and correlated votes are the normal case.** This framework uses the failure of the premise more often than the theorem.
22. **The Brooks path count.** `n` participants have `n(n−1)/2` pairwise channels. Arithmetic. That cost grows at that rate is an empirical claim, not this identity — see E.3.

39.16.2 E.2 What is postulated

Each item here is an assumption. Each is badged wherever the framework uses it. None is derived from anything in E.1, and none is established by anything in E.3.

1. **S3, the affine effective-count design axiom.** The assumption that `g(u) = e^{f(u)}` is affine in `u` — equivalently, that the effective count is `N_eff = 1 + (N−1)u`, linear in non-parallelism. **This is chosen for tractability and interpretability, not derived.** It is what makes the uniqueness theorem produce a logarithm rather than some other concave function. Every number in this framework computed from `1 + (N−1)u` — every effective count, every marginal-gain table, every `S` — inherits the status of a postulate, not of a theorem. If S3 is wrong, Theorem 2.1 is still true and no longer relevant.
2. **The linear form of the storage law.** `Q_{t+1} = (1−delta)Q_t + eta S_t − C_t` is **the minimal tractable open-system law, not a law of nature.** It assumes loss proportional to stock, conversion linear in surplus, cost additive, parameters constant within a period, no `Q`-`S` interaction, no saturation. Each simplification is chosen so a closed-form solution exists, and each fails somewhere. **A nonlinear `Q` with saturation, a state-dependent `eta`, or a cost rising with `Q` would all be more realistic and none would give Theorem 3.5.** A wrong model that can be computed with, whose wrongness has been stated, is a usable instrument; a right model that cannot be evaluated is not.
3. **`Q` itself as a state variable. The storage variable `Q` is a modelling choice, not derived from first principles.** There is no measurement of `Q` — only a proxy `Y` chosen under Protocol C-6, and different engineers will choose differently. That "what persists in a social system" is a single scalar decaying exponentially is a strong assumption, made because it is the simplest one that produces a testable model.
4. **The modelling map onto social systems.** That rooms, teams, institutions, and cities are systems to which the storage law applies at all is the largest assumption in this corpus, and nothing in Papers A or E establishes it. The map identifies interaction geometry with `u`, institutional persistence with `Q`, forgetting and turnover with `delta`, capture into artifacts with `eta`, overhead with `C`. **Each identification could be wrong without any theorem being wrong.**
5. **Hyperrecursive symmetry.** That the same equation governs a desk, a room, a team, a building, an institution, and a city, with only `(delta, eta, N, C)` changing, is **a claim about MODEL FORM, not about physics.** It is not a scaling law, it does not predict how parameters change with scale, and it has not been tested at two scales simultaneously

in any real system.

6. **The state encoder.** That participants, views, or categories can be represented as unit vectors whose inner product carries the relevant notion of alignment is an operational choice. Different encoders give different `u` for the same group, and the framework's estimator — centered, unit-normalized response or error vectors, so that `x·y` is a Pearson correlation and `u = 1 - r^2` — is one choice among several defensible ones.
7. **The ensemble aggregation.** The surplus law is stated for a pair. Real groups have more than two members. The participation-ratio estimator, `N_eff = (SUM lambda)^2 / SUM lambda^2` from the correlation matrix, extends the pairwise law to an ensemble in a way that is consistent with the affine form by construction. **Consistency by construction is a property of the choice, not evidence for it.**
8. **Cost commensuration.** That coordination hours, rework, conflict, and rent can be converted into the same unit as stored order is assumed every time `h_t` is computed. The conversion factor is chosen by the engineer, is not measured, and dominates the headroom result whenever `C` is large.
9. **Constancy of parameters within a survey.** `delta` and `eta` are treated as constants over a measurement span. They are not. They move with turnover, tooling, morale, and season. Protocol C-1 asks for two windows and an interval precisely because the point estimate is a fiction the arithmetic requires.
10. **The interpretation of `u = 0` at anti-alignment.** Because `u = 1 - (x·y)^2`, a perfectly anti-correlated participant scores `u = 0`, identically with a perfectly correlated one. This framework interprets that as correct — a reliable contrarian is one instrument, not two — but the interpretation is a reading of the formalism, not a result.
11. **`optional activity` as the observable proxy for a room's `eta`.** Gehl's category is real and observable; identifying it with the conversion efficiency term is this framework's proposal.
12. **Acoustics as `c`-factor engineering.** The chain is: reverberation affects speech intelligibility (established acoustics); intelligibility affects who speaks and how often (plausible, partially studied); turn-taking equality correlates with `c` (Woolley et al.). **The composite claim — that changing RT60 changes collective intelligence — is a chain of three links of which the middle is the weakest, and it has not been tested end to end.**
13. **That a badged analogy may organise and may not compute.** This is a house rule, and it is stated here because it is a rule about postulates rather than a postulate itself. **An analogy carried under a badge may structure a paper, name a pattern, and suggest where to look. It may not contribute a number.** CCP-032 is the worked instance and it is deliberate: every quantitative result in that chapter comes from an information-budget bound, from the Gabor timing/rhythm trade, or from the Lyapunov horizon law, and **not one comes from Heisenberg.** `dx·dp >= hbar/2` does zero arithmetic work in this framework. Where a later chapter appears to draw a figure from an analogy, that is a defect in the later chapter.
14. **That the internal and external scopes correspond at all.** The master correspondence of the personal-scope material — the same six parameters read at two scopes — is **a correspondence of dynamics and geometry, and it is NOT a correspondence of agency.** CCP-025's breakdown section is the longest in that chapter and names nine disanalogies: exit, heterogeneity, aggregation, homunculus/reciprocity, observation asymmetry, boundary asymmetry, timescale, rivalry, ethics. **Three of them are computed**, and they are the ones that matter: a mean-field decay error reaching **92% by `t = 12`**; two nine-member groups with *identical* `u = 0.5999` and `S = 1.7578` losing **77% against 21%** of their surplus to a three-person departure — **aggregate statistics hide the structure that decides the outcome**; and a **90.1x** period-normalisation trap when the two scopes are compared on mismatched clocks. CCP-038 adds five named breaks at the third scope, in E.10.
15. **That trust is a source of order.** It is not, and the correction is a postulate-level one because naive practice assumes otherwise. Trust acts by reducing `delta`, and `Q_inf = (eta*F(u*) - c*)/delta`. **When the numerator is negative, dividing by a smaller positive**

denominator makes the result more negative. In the worked consortium of CCP-025, trust alone moved `Q_inf` from `-0.0836` to `-0.2114` — **worse**. Reciprocity alone moved it to `+0.7394`; both together to `+1.8701`. **Trust amplifies whatever sign the system already has. Fix the numerator first.**

39.16.3 E.3 What is borrowed

Every empirical result this framework draws on, what it actually showed, in what setting, and what must be assumed to transfer it. Nothing here was obtained in the reader's organization.

Woolley, Chabris, Pentland, Hashmi & Malone [Woolley et al., 2010]. Several hundred participants in small groups (two to five people) completed batteries of short laboratory tasks. Factor analysis of group performance across tasks yielded a single general factor, `c`. `c` was **not** predicted by the average or the maximum individual intelligence of members. It correlated positively with members' average social sensitivity, negatively with the variance in speaking turns, and positively with the proportion of women. *Setting:* laboratory, tasks of hours, groups assembled for the study. *Transfer assumption:* **that a factor measured over hours in assembled groups persists in teams working for years on their own problems.** This is not established. CCP-013 states it; this appendix repeats it.

Centola, Becker, Brackbill & Baronchelli [Centola et al., 2018]. In controlled online experiments, small groups played a coordination game that established a naming convention; committed confederates then pushed an alternative. Below roughly a quarter of the population, the minority failed; above it, the convention flipped. The tipping point was measured, not simulated. *Setting:* online, small groups, an arbitrary convention with no material stakes, no history, and no hierarchy. *Transfer assumption:* **that a threshold measured for a costless naming convention among strangers applies to a costly practice among colleagues with careers at stake.** The framework uses ~25% as an order-of-magnitude planning figure and says so.

Centola & Macy and Centola [Centola and Macy, 2007, Centola, 2010]. The 2007 paper is analytic: for behaviors requiring reinforcement from multiple sources — complex contagion — long-range random ties can impede rather than accelerate diffusion, reversing the small-world intuition. The 2010 paper is an online experiment in which a health behavior spread farther and faster through a clustered-lattice network than through a random network of the same density. *Transfer assumption:* that organizational practices are complex rather than simple contagions. Often reasonable; not measured in any case the corpus reports.

Granovetter [Granovetter, 1978]. An analytic threshold model: each person acts when the number of others acting exceeds a personal threshold, so the aggregate depends on the whole threshold distribution and one person's shift can be the difference between nothing and everything. *Status:* a model, not a measurement. Its lesson is that aggregate discontinuity does not require anyone to be irrational.

Watts & Strogatz [Watts and Strogatz, 1998]. Rewiring a small fraction of a ring lattice's edges collapses average path length while clustering stays high. *Status:* a result about constructed graphs, with empirical illustrations.

Barabási & Albert, and Albert, Jeong & Barabási [Barabási and Albert, 1999, Albert et al., 2000]. Growth plus preferential attachment generates a power-law degree distribution; such networks tolerate random node removal well and targeted removal of high-degree nodes very badly. *Status:* model results with empirical support in several real networks. *Transfer assumption:* that the dependency structure in question is scale-free enough for the fragility result to bite. Check it — a node-removal audit takes an afternoon.

Hong & Page [Hong and Page, 2004]. A mathematical and computational result: under stated conditions, a randomly selected group of diverse problem solvers can outperform a group of the individually best-performing solvers. *Status and contest:* the result depends on strong conditions about the problem space, the size of the candidate pool, and the formal definition of diversity, and it has been criticized on mathematical grounds — notably by Abigail Thompson in the *Notices of the AMS* (2014). **This framework uses Hong & Page as an illustration that diversity can dominate ability under some conditions, and not as a proof that it**

generally does.

Bikhchandani, Hirshleifer & Welch [Bikhchandani et al., 1992]. A model of informational cascades: sequential agents observing predecessors rationally cease using their private signals after a short prefix of decisions. Aggregation stops; the cascade is fragile and reversible by a little new public information. *Status:* theory with extensive experimental support. Used here as the mechanism by which a diverse group becomes parallel with no one changing their mind.

Allport, and Pettigrew & Tropp [Allport, 1954, Pettigrew and Tropp, 2006]. Contact reduces prejudice under four conditions — equal status, common goals, cooperation, institutional support. The meta-analysis of over 500 studies found contact generally associated with reduced prejudice, with larger effects when the conditions were met. *Status:* a large meta-analysis, correlational in much of its base, with known concerns about publication bias and self-selection into contact.

Bandura on moral disengagement; Kelman; Haslam on dehumanization. Evidence that harm at scale is preceded by euphemistic labeling, displacement of responsibility, routinization, and the withdrawal of full humanity from a target group. *Status:* a well-supported tradition, drawn on here only for the ethical constraint in the standing refusals of the corpus and the failure modes in CCP-014.

Ostrom [Ostrom, 1990]. Comparative case study of long-enduring common-pool resource institutions — irrigation systems, fisheries, forests, pastures — yielding eight design principles. Nobel Prize 2009. *Status:* **inductive regularities from case comparison, not theorems and not experiments.** A later meta-analysis by Cox, Arnold and Villamayor-Tomás (2010) found broad empirical support across a large body of cases while suggesting refinements. This is the most directly engineerable body of work in the sociology literature this framework uses, and it remains a set of empirical regularities.

Putnam [Putnam et al., 1993]. Italian regional governments created by the same 1970 devolution performed very differently, and the differences tracked civic traditions with roots centuries earlier. *Status:* a landmark comparative study whose causal direction has been contested. Used here for the durability of civic structure, not for a mechanism.

Gambetta [Gambetta, 1988]. An account of the Sicilian mafia as a supplier of private protection where trust was scarce and central enforcement weak. *Status:* descriptive and interpretive. Used as a demonstration that order can be produced by networked obligation, and as a warning about what kind of order that is.

Grassé 1959. Termite nest construction proceeds by each individual responding to the state of the work already done: stigmergy. *Status:* a foundational ethological observation, since formalized. The rule this framework derives — put the state in the medium, not in a coordinator — is an inference from the mechanism, not a finding.

Schelling [Schelling, 1971, 1978]. Focal points: parties who must coordinate without communicating converge on salient options; the evidence was informal elicitation, not controlled experiment. The segregation model shows mild individual preferences producing strong aggregate separation — a model, whose value is that it removes the need to posit strong preferences to explain strong outcomes.

Star & Ruhleder [Star and Ruhleder, 1996]. Infrastructure is relational rather than substantial and becomes visible chiefly upon breakdown. *Status:* an ethnographic and theoretical study of a scientific collaboratory. Used here for one consequence: **maintenance is chronically underfunded because working infrastructure is invisible, which means `delta` is systematically underestimated wherever it has not been measured.**

Weiser 1991. "The most profound technologies are those that disappear." *Status:* an essay in *Scientific American*, an argument rather than a study. Quoted as an aphorism and used as a design orientation.

Alexander [Alexander, 1965, Alexander et al., 1977], together with the 1964 *Notes on the Synthesis of Form,* which carries no key in this corpus's bibliography. The direct ancestor of this framework's form. The semilattice argument — living structures have overlapping rather

than strictly nested categories, and designers impose trees because trees are easy to hold in mind — is supported by examples, not measured. **This framework's use of it to explain why hierarchy drives `u` toward zero is an inference this framework makes, not a claim Alexander proved.**

Jacobs [Jacobs, 1961]. Eyes on the street, short blocks, mixed primary uses, aged buildings. *Status:* observational argument of extraordinary quality, not controlled study; later quantitative work has supported parts and complicated others.

Whyte [Whyte, 1980]. Direct observation with time-lapse film of New York plazas: people sit where there is somewhere to sit; sun, food, water, and triangulation draw and hold people. *Status:* systematic observation in one city over several years. **The method is more transferable than the findings, and this framework borrows the method for CCP-010.**

Gehl [Gehl, 1971]. Necessary, optional, and social activity, with optional activity highly sensitive to conditions. *Status:* observational studies across several cities.

Hall [Hall, 1966]. Proxemic distance bands — intimate, personal, social, public. *Status:* observation predominantly of one cultural group; Hall himself emphasized cross-cultural variation. A frame, not a specification.

Ulrich, *Science* 1984. Twenty-three matched pairs of gall-bladder surgery patients in one Pennsylvania hospital over nine years; those in rooms with a window view of trees had shorter postoperative stays, fewer negative nursing notes, and took fewer strong analgesic doses than those facing a brick wall. *Status:* **a small, single-site, matched-pair study, and the strongest single piece of evidence in this framework that physical space affects a hard outcome.** Both halves of that sentence are load-bearing. It has been influential far beyond what a study of that size can carry alone; later research has broadly supported the direction of the effect without settling its size.

Sabine. `RT60 = 0.161 V / A` in metric units, `V` in cubic metres and `A` in square-metre sabins. *Status:* an empirically grounded formula valid for reasonably diffuse sound fields. **It degrades in heavily absorbing rooms — Eyring's form is preferred there — and in strongly non-diffuse geometries.** Every RT60 figure in this framework carries that caveat.

Conway [Conway, 1968]. Organizations that design systems produce designs copying their own communication structures. *Status:* an argument in a trade magazine, later supported by empirical work comparing product and organizational structure in matched codebases. A structural explanation, not a law.

Dunbar's number, ~150. From a regression of primate group size on neocortex ratio, extrapolated to humans. *Status:* **contested**, including on the statistical ground that the extrapolation carries very wide uncertainty. Do not use it as a design threshold; use it as a reminder that relationship maintenance has a cost — a `C` claim that needs no magic number.

Brooks [Brooks, 1975]. Adding people to a late project makes it later; communication paths grow as `n(n–1)/2`. *Status:* an engineering argument from experience on one large project, widely corroborated in practice and not established by controlled experiment. The path count is arithmetic; that cost tracks it is the empirical part.

Goodhart and Campbell [Goodhart, 1975, Campbell, 1979], with Strathern (1997) named in prose and carrying no key here. A measure used for control ceases to measure what it did. *Status:* empirical generalizations, not theorems. This framework is subject to them: **headroom, once rewarded, will be gamed, and Protocol C-7 will begin to produce the number the reporter needs.**

Queueing theory. Waiting time growing as `1/(1 – rho)` is exact for M/M/1 and approximately right for a wide class of systems. *Status:* proved under distributional assumptions that real systems violate in ways that generally make matters worse.

39.16.4 E.4 What is untested

This is the largest drawer and the one most easily forgotten, because untested claims look exactly like tested ones once they are set in the same typeface.

1. **No organization has been surveyed under CCP-015, preregistered under CCP-015, and reported.** Not one. Every worked example here is arithmetic on parameter values chosen to be plausible, not measured in the field. The examples show that the calculations can be done and what they would mean; they do not show that the model describes any real system.
2. **The estimators in Appendix C have not been validated against ground truth.** There is no case in which `delta` was estimated by Protocol C-1 and independently known. The protocols are built to be internally consistent and to fail loudly; that is not accuracy.
3. **The held-out performance of the storage law is unknown.** The falsification rules of CCP-015 exist because the model must be falsifiable, not because it has survived falsification. **A framework that cannot fail cannot be trusted, and this one has not yet been given the chance to.**
4. **Scale-freedom has never been checked at two scales at once.** The test would measure `(delta, eta, N, C)` for a desk and for the institution containing it, same units, same period, and show the same equation holding with different parameters. Not done.
5. **The acoustics-to-`c` chain is untested end to end.** An honest test would randomize RT60 across otherwise matched deliberation rooms and measure turn-taking variance and task performance. It would not be expensive. It has not been run.
6. **The 25% figure has not been tested in an organization**, nor has any organizational analogue of it been measured by anyone this framework can cite.
7. **The lever cost ranking is an argument, not a measurement.** That `eta` is cheapest, then `C`, then `u`, then `N` follows from the structure of the law plus practitioner reasoning. **No cost study supports the ordering; treat it as a prior to be updated, not a finding.**
8. **The physical specifications are syntheses, not trials.** CCP-011, CCP-012 combine established acoustics with observational urbanism and a laboratory correlation. No room built to that specification has been compared against a control.
9. **The failure mode catalog is a taxonomy, not an epidemiology.** CCP-016 names failure modes and gives each a signature; it does not report how common any of them is, because nobody has counted.
10. **The consent constraint has not been tested as an engineering property.** That a method working by removal of resistance has no coercive failure mode is an argument about the structure of the method — an ethical guarantee about what this corpus will recommend, not an empirical claim that such methods never harm. **The ten entries above are the original set, and they are the ten CCP-039's specification S11 refers to. Five more were added by the personal- and cellular-scope material, and every one of them is larger than the entry it follows.**
11. **Of twenty-seven parameter-by-scope cells, nineteen have a procedure, one has a validated number, and none has a first-party measurement.** This is CCP-038's audit and it is the single most useful number in this appendix — see E.10, where it is set out in full. **The scope with the least-validated instrument turns out to have the best-measured parameter**, which is not a comfort.
12. **No formula in Volumes Two or Three has passed FVP-7.** The cross-subject reproducibility gate — same operationalisation, different people, different laboratories, do the parameters recover — is the gate that decides the hard-science question and it has not been attempted for any formula in this corpus. **the corpus's internal-scope conventions names it the gate most often skipped; CCP-027 computes that it is most often skipped because it is most often failed.**
13. **`D` and `P` have never been Rasch-calibrated.** The route out of the demotion is specified, costed at 123 items per term, and untaken. Until it is taken, every magnitude descended from the Moment Equation stays in the drawer where E.7 keeps it.

14. **No cell has been shown to demodulate sequence-specific information from an acoustic signal, and this corpus does not claim one can.** Everything upstream of that gap is real and cited — mechanotransduction, PIEZO channels, force-to-chromatin transmission, DNA data storage, addressing formats. **The channel for sequence-specific content is the open question, it is one well-specified falsifiable gap, and naming it sharply is the most useful thing the cellular-scope material does.**
15. **The essence law has never been tested as a bridge, only as arithmetic.** That $f(n) = k\ n^2$ reproduces the source document to fourteen significant figures is a verification of a transcription, not evidence that essence indices are the right index for anything. **An essence frequency is not a physical resonance**, and no rendering built to the CCP-032 constraints has been compared against a control.

What would change these entries. Each item above names its own test. The framework's position is that it should be judged by whether those tests are run and reported, including when they fail, and that a reader who runs one and gets a null result has done this corpus a service that no amount of agreement could.

39.16.5 E.5 The honesty clauses, reproduced in full

These are reproduced verbatim from the formal spine and are binding wherever the law is used.

> **The triangle inequality FAILS for geometric triples (CCP-A02 §5). Subadditivity holds on u only.**
>
> **As deployed, the surplus functional is a LOWER BOUND on the ensemble's information, not the information itself. We log the floor, not the ceiling.**
>
> **CCP-A02 is pure mathematics and makes no empirical or metaphysical claim.**
>
> **CCP-009: surplus is NOT identified with Shannon entropy or mutual information in general.**
>
> **CCP-007: no thermodynamic negentropy law is proved.** `F(u) + N_can(u) <= ln N` (Cor 2.2) — **higher static surplus does NOT imply higher canonical concentration.**
>
> **S3 is the design axiom — chosen, not derived, and that must be said wherever it is leaned on.**
>
> **The storage variable `Q` is a modelling choice, not derived from first principles. The linear update is the minimal tractable open-system law, not a law of nature.**
>
> **Woolley's `c` is from short-task lab settings.**
>
> **Centola's 25% is from controlled online experiments.**
>
> **The storage law is scale-free in form: that is a claim about MODEL FORM, not about physics.**
>
> **If `eta F(u*) < c*` the equilibrium is negative — the regime is unsustainable or the linear model must be replaced. The theorem does not hide this failure mode.**

the personal- and cellular-scope material added six more, and they are binding on the same terms.

> **A badged analogy may ORGANISE and may not COMPUTE.** It may structure a paper and name a pattern; it may not contribute a number. CCP-032 contributes no number from Heisenberg.

The (+1) column carries EVIDENCE. The (-1) column carries PRECEDENT. Two sources agreeing does not double the evidence when only one of them is evidence. State the role of each column in the sentence that draws the concordance.

Every speculative bridge is marked SPECULATIVE in the sentence that states it, not in a note attached to it.

The internal/external correspondence is one of DYNAMICS AND GEOMETRY, not of AGENCY. The arithmetic descends between scopes; the consent constraint does not descend to a cell, and must be carried by hand.

Division requires a ratio scale with a true zero. A formula that divides two ordinally-measured terms is a heuristic ordering and is DEMOTED IN PRINT — in this paper that made the claim, and not only in the appendix.

There is no established mechanism by which a cell demodulates SEQUENCE-SPECIFIC information from an acoustic signal. Everything upstream of that gap is real and cited; the gap itself is the open question and it is stated in every chapter that touches it. **No clinical claim; no claim that DNA is edited, altered or damaged; no claim that any capability has been conferred; and Casimir does nothing at biological scale.**

39.16.6 E.6 The four refusals, and why each is a refusal rather than a preference

1. No claim about physics, thermodynamics, or consciousness. The temptation is real, because `F` is a logarithm of a count and looks like an entropy. It is not one. CCP-009 declines to identify surplus with Shannon entropy or mutual information in general, and CCP-007 proves no negentropy law. `F(u) + N_can(u) <= ln N` means surplus and canonical concentration compete for the same ceiling: **a system with high static surplus is not thereby more ordered, and any sentence suggesting otherwise is outside this corpus.**

2. No postulate without its badge. An unbadged postulate is not a stylistic lapse; it is the mechanism by which a framework becomes unfalsifiable, because a reader who cannot tell assumptions from results cannot design a test.

3. No number that was not computed or cited. Every table here can be reproduced from the formula above it with a calculator, and every empirical figure is attributed in E.3 with its setting stated. Where a figure would need research that has not been done, the framework says what would need measuring instead. **An invented number is worse than no number, because it is actionable.**

4. No intervention that removes a person's ability to decline. The method this framework teaches works by lowering the cost of a better position — building the lower valley next door — and therefore has a bounded worst case: someone looks at the valley and stays. A method that works by removing the option has no such floor. **The safety property and the effectiveness property are the same property, which is convenient, and it would be the rule regardless.**

39.16.7 E.7 What this corpus demoted

The demotion of the Moment Equation is stated before the drawers, because it is the largest one and because burying it would have been the easiest thing this appendix could do. This section holds it and the rest, in one place, so that the count is visible.

Eight register entries carry the status DEMOTED. Demotion means the structure stands and the magnitude does not; the demoted content may not reappear in any form, **including in a summary table**, and demotion travels — when a gate fails, every entry downstream of the failed one is demoted in the same pass.

1. **The magnitude of `D/P`.** FVP gate 2 returns DEMOTE. `D` and `P` are not instrumented on ratio scales with true zeros, and division is not a permitted operation on the scales that exist. **The route out is Theorem 54.3 and it costs 123 items per term.**
2. **The distortion magnitude `Delta = |ln r|`**, demoted to ordinal. The direction stands, because an ordering needs only ordinal information.
3. **The quantitative join `M_t -> Q_{t+1}`**, demoted to structural. The form stands; the magnitude inherits entry 1.
4. **The same demotion at group scope**, on the doctrine-to-data ratio as a measured quantity. The diagnostic survives as a diagnostic and not as a number.
5. **Every M-derived formulary entry**: the launch budget `L = max(0, M − C)`, the return difference `Delta_R = |ln(M'/M)|`, and their interpersonal counterparts. **No M-derived magnitude survives anywhere in the corpus.**
6. **The acoustic-to-`c`-factor CAUSAL claim, reported as NOT CONFIRMED.** The falsification test is stated in §14.13 and its result is not confirmed. **Any chapter invoking acoustic `c`-factor engineering carries that status**, and the composite chain — RT60 to intelligibility to turn-taking to `c` — remains three links of which the middle is weakest and untested end to end.
7. **The claim that the shuffled-input control is sufficient to detect a spurious input**, shown insufficient by construction and then by simulation. Three measured holes; Protocol C-8 states them and prices them. **This is a demotion of a method this framework recommended in its own earlier chapters, and the earlier worksheets carry the correction.**
8. **Seven published compound values in the 2022 sonic-chemistry source**, demoted as data-entry errors rather than as a different law. The compound rule reverse-engineers as `f = k^2 G^4` with `G` the count-weighted geometric mean of indices, matching 37 of 44 published composites to better than `1e−12`; **the O2 row is provably impossible under the stated rule.** The source's prose formula `1+5^.5/2` is missing its parentheses.

Two negative results were retained rather than buried, and both should stay in print. Momentum descent failed to escape the worked landscape at every `beta` tested — the barrier is a wide plateau, not a spike — and **the momentum recursion is explicitly excluded from the formulary with its failure stated.** That exclusion is a feature. A formulary containing only the recursions that worked is a formulary that has been edited rather than assembled.

And one correction to a target this corpus had itself proposed. `r = 1` is not an unqualified goal. Headroom at the `D = P` fixed point is **exactly zero by cancellation**, and by the two-track ledger an organisation optimising raw stored order is **actively driven away from `r = 1`**, because over-framing banks faster. CCP-019 may not present calibration as a simple objective, and does not.

39.16.8 E.8 What this corpus rejected

Five claims were worked end to end in print and failed. Four of them come from the corpus this corpus is most often assumed to be defending, which is the point: **a concordance method is only worth having if it can return this verdict.**

1. "528 Hz repairs DNA" — REJECTED (CCP-029 §61.14). Four computed objections, not four opinions. A frequency in hertz for a named musical pitch is a convention, not a property of a molecule. The acoustic wavelength in soft tissue at 528 Hz is **2.9167 m** against a target of nanometres — a ratio of order `9e9`, so whatever the field does to a nucleus it does to the whole nucleus at once and identically. The acoustic quantum carries `h*f = 2.18e−12 eV`, which is **`8.17e−11`** of the ambient thermal jostling at body temperature. And the molecular vibrational modes that do exist sit **34.7 to 37.3 octaves above** 528 Hz. `C1 FAIL, C2 FAIL, C3 FAIL`; the (+1) column empty; **the (-1) column VOID**, because the source is 1990s and not traditional and therefore carries zero bits of precedent.

Then the rejection was turned back on this corpus's own pipeline, in the same passage. CCP-031 sonifies measured infrared spectra by fixed octave transposition, and **that pipeline may not claim what it has just refused to let anyone else claim.** Dividing a frequency by

`2^35` preserves a ratio and no physical interaction; the essence frequencies of the corpus's disclosed-signal constraints are essence frequencies and not physical resonances; and nothing in the cellular-scope material may claim sequence-specific acoustic action on DNA, because the objection applies to this corpus's tones exactly as it applies to anyone else's. **The difference between the two is that this corpus states its mapping, states that the mapping is a convention, and states that the convention carries no mechanism. If a reader concludes the difference is thin, the reader has understood the situation correctly. It is thin. It is also the entire distinction between an honest notation and a false one.**

2. The 64-codon / 64-hexagram bridge — REJECTED (CCP-031 §67.11). The genetic code has 64 codons; the *I Ching* has 64 hexagrams. **`2^6 = 64` and `4^3 = 64` are different factorisations**, which is to say the two are not the same structure at all — they are two different structures with the same cardinality. A proposed correspondence between them is a chosen bijection, one of `64! = 1.2689e89` available, **worth 296.0 bits of arbitrary choice.** `C1 FAIL` (a shared count is not a shared pattern), `C2 FAIL` (it generates no prediction that could come out otherwise), `C3 FAIL`. **Sixty-four is a small highly-composite power, and that two systems both land on it is a fact about the number.**

3. Astrology's orb — the structure retained as precedent, the validity REJECTED (CCP-032). The traditional orb is a genuine tolerance-versus-window trade, and `C1` passes on that structure alone: a wide orb admits more configurations and corresponds to a longer window of claimed influence, which is the conjugate-spread shape. **`C2` is marked FAILED in print**, and the failure is cited in the same passage that draws the correspondence: **Carlson, "A double-blind test of astrology", *Nature* 318 (1985), 419–425** — the definitive controlled test, which astrology failed. **The tradition is used as a (-1) precedent for a tolerance/window structure and for nothing else.** A tradition may record a structure and forecast nothing, and this one does.

4. Gematria's evidential reading — DEAD BY ARITHMETIC; the mnemonic reading intact (R12, CCP-028). Measured rather than asserted: across **234,454 dictionary words** the map yields **305 distinct outputs**, destroying **10.47 bits** at a **99.995% collision rate**; and in an 8,000-word corpus, **essentially every word has a value-twin, `p = 0.9997`. A hash with that collision rate cannot carry evidence.** It can absolutely carry a memory index, and the method of loci is a validated mnemonic technique, so the mnemonic reading stands untouched. **This is the model for how the whole old corpus is treated: disprove the strong claim with arithmetic, preserve what is genuinely there, and do both in the same passage.**

5. The empty cells — a rejection performed one hundred and thirty-three times. CCP-038's register holds **156 entries. Precedent is claimed in 23 rows — 14.7%. The remaining 133 read "none claimed."** the cellular-scope material's commitment is that every substantive claim carries two references; that commitment is met where the historical corpus has something to say and **reported as unmet where it does not**, which is the only way it can be kept without inventing the second reference. The register's speculative group holds 12 entries, **4 of them rejections or refusals printed as such.**

14.7% is a considerably less impressive number than the two-column method promises in the abstract. It is printed because it is the true one, and because a blank column is the correct output when the historical corpus has nothing to say. Filling it would have been the easiest and most damaging thing this corpus could do.

39.16.9 E.9 What testimony is worth, computed

The likelihood ratio of testimonial evidence, under a selection process that publishes successes and discards failures, is exactly `1.000000`.

Not approximately one. Not weakly informative. **Exactly one — which means the evidence does not move the posterior at all**, because the selection guarantees that successes are observed whether or not the method does anything. All four hypothesised success rates in CCP-035's worked case produce the same 500 published accounts. **A thousand testimonials from people for whom a method worked carry precisely the evidential weight of zero testimonials.**

This is the exact quantification of survivorship bias, and it is the cleanest demonstration in the work of why the (-1) precedent column can never be promoted to evidence: the historical corpus is itself a survivorship-filtered record. What was preserved was preserved because someone found it worth preserving, and the selection is invisible from inside the surviving set. That is not an argument against keeping the column — independent arrival is real and worth having, because it argues that a structure was findable by more than one route. **It is an argument against ever counting it twice.**

The companion result prices the alternative. Likelihood ratios measured at **2.0000** for a prediction specifying only a *direction*, against **11.3577** for one specifying a *magnitude* with a tolerance: **a prediction with a number attached is about 5.7x more evidentially powerful**, equivalent to about 3.5 direction calls for the same single measurement. **Hence the standing rule: a belief that predicts only a direction is not yet testable.** It has not been refuted; it has not been examined. That is a worse condition than being false, and the fact that it feels like a safer one is precisely why so many claims stay in it.

39.16.10 E.10 What the instrument audit found

CCP-038 read the three scopes — institutional, personal, cellular — against the same nine parameters and counted what each cell actually holds.

	Institutional	Personal	Cellular	Row total
Defined estimator	9	9	1	19 of 27
Validated estimator	0	0	1	1 of 27
First-party measurement	0	0	0	0 of 27
Demoted to ordinal (CC-R1)	M, D/P	M, D/P	n/a	4

Nineteen of twenty-seven parameter-by-scope cells have a procedure. One has a validated number. None has a number this research programme measured. The single validated cell is a borrowed median from one cell line. **The scope with the least-validated instrument has the best-measured parameter**, and the scope with the most borrowed empirical support has the worst first-party measurement. Neither is a comfort; both are what the count says.

Everything that follows comes from taking that table seriously rather than from any philosophical objection to correspondence. The correspondence then fails in five distinct ways, and they are genuinely distinct — repairing one does not touch the others.

```
BREAK 1   A mind is not a small society.
          The independence the law's aggregation needs is unavailable inside one skull.

BREAK 2   A cell is not a small mind.
          M and D/P require a representation. A cell has no frame, so it has no D/P.

BREAK 3   Members can exit; cells cannot.
          The consent constraint that governs this corpus has no cellular analogue,
          which means the ethics do NOT descend with the arithmetic.

BREAK 4   Aggregation is not always well-defined, in principle and not merely in practice.

BREAK 5   The cellular scope has no established input path,
          and no composition theorem connects any scope to any other.
```

Break 3 is the one to carry out of this appendix. The arithmetic descends to the cell and the ethics do not — so a reader who imports the formalism downward has imported a model and left the constraint behind. That is the failure this corpus would least like to be responsible for, and it is why the standing refusals of the corpus is restated in every paper of the cellular-scope material that touches a cell.

The master correspondence must therefore be stated as one of dynamics and geometry, NOT of agency — E.2 item 14, resting on CCP-025's nine disanalogies and its three computed ones: a mean-field decay error of **92% by `t = 12`**; **77% against 21%** of surplus lost by two groups with *identical* `u = 0.5999` and *identical* `S = 1.7578`; and a **90.1x** period-normalisation trap across mismatched clocks. **Aggregate statistics hide the structure that decides the outcome**, and that sentence is most of what the audit found.

39.16.11 E.11 What the programme costs, and how likely it is to fail

CCP-027 specifies the five-stage research programme that would settle whether the personal-scope material's formulas are science, and then does the thing such chapters normally omit: **it computes its own odds of failure and publishes them.**

```
Preregistered kill conditions                        21
  of which terminal for the whole enterprise          4
Researcher and coordinator hours                8,554.8    (5.35 FTE-years)
Participants                                        570
Calendar, strictly sequential                        71 months  (5.92 years)
  with 30% overlap of stages 2-3 and 4-5             61 months

SURVIVAL CURVE, under the author's explicitly badged priors
  after stage 1     0.9000
  after stage 2     0.8100
  after stage 3     0.6885
  after stage 4     0.4475
  after stage 5     0.3580   <- P(completion) = 35.8%

P(dies at or before the cross-subject recovery stage) = 55.3%
```

The programme is more likely than not to be dead before the field trials begin. The heaviest prior — 0.35 at stage 4 — is deliberate: stage 4 is where a formalism that *describes* meets the demand that it *measure*.

A programme that computes and publishes its own probability of failure is more credible than one claiming certainty, and the converse is the more useful reading: **a programme with a high stated completion probability has almost certainly written its kill conditions too loosely.** Twenty-one are published in advance, four are terminal, and the commitment is to publish at every stage — including and especially when the stage kills the programme — with null results titled as null.

Note the coherence with the demotion, because it is not a coincidence. FVP-7, cross-subject and cross-laboratory recovery, is both the gate that decides the hard-science question and the stage most likely to end the programme. **The thing this corpus most needs to be true is the thing its own arithmetic says is most likely to kill it.**

39.16.12 E.12 A closing note on the register

This framework has two parents and does not disown either. From the theorems it takes the discipline that a claim must be either proved, or badged, or cited, or named as untested — **and, since the personal- and cellular-scope material, either demoted or rejected, in print, in the paper that made it.** From the trade it takes the willingness to build with an incomplete model, because rooms and institutions are built by people who cannot wait for the model to be finished.

The defense of building on postulates is one line long and it is the founding material's: **every geometry is built on postulates, and the builders who prosper are not those who have**

none but those who know exactly where theirs are kept. They are kept in E.2. They are labeled. Audit them.

And the register keeps a second address, which is the one this appendix was extended to hold. What a framework has thrown away is a harder thing to fake than what it asserts, so the demotions are in E.7, the rejections are in E.8, the evidential worth of testimony is computed in E.9, the instrument audit is counted in E.10, and the programme's own odds of dying are printed in E.11. **Nineteen procedures, one validated number, no first-party measurement, a keystone formula demoted by its own protocol, and a thirty-five per cent chance of finishing.** That is the true state of the instrument, and a reader who finds it discouraging has read it correctly and should still find it more useful than confidence would have been.

39.17 Appendix F Correspondence to the founding material

the founding material taught a practitioner to **read** their own comfort gradient. This framework teaches an engineer to **build** the circumstance that generates and retains it. The twelve nodes are not superseded; they are re-instrumented. What was a protocol run on oneself becomes a specification run on a system, and in two cases — the recursion and fixed-point nodes — the ceremonial teaching turns out to be a term in a proved equation.

39.17.1 F.1 The correspondence, node by node

Stillpoint. *the founding material's teaching:* The present is the comfort zone by definition *the corpus papers that develop it:* CCP-015 (The Survey), CCP-007 (Loss), CCP-008 (Headroom) *the term in the law:* `Q_0`, and the zero-input window *the reading becomes the building:* Admitting the starting state becomes **measuring `Q_0` and `h_0` before proposing anything**

Ledger. *the founding material's teaching:* Every comfort on a card with its carrying cost; every discomfort assayed for hidden yield *the corpus papers that develop it:* CCP-008 (Cost), CCP-015, CCP-014 (Institutions as Storage) *the term in the law:* `C_t`, and the five cost ledgers *the reading becomes the building:* Carrying cost becomes **the dissipation budget in Q-units**; hidden yield becomes why a costly structure persists

Gradient. *the founding material's teaching:* Water does not negotiate; read the worn path, do not argue with it *the corpus papers that develop it:* CCP-012 (Thresholds), CCP-013 (The Silo), CCP-016 (Intervention) *the term in the law:* the direction of cheapest motion *the reading becomes the building:* Reading desire lines becomes **regrading the landscape so the wanted move is the cheap one**

Horizon. *the founding material's teaching:* Three futures cut from three non-parallel glasses; the collision test *the corpus papers that develop it:* CCP-009 (Effective Count), CCP-009 (Engineering u), CCP-010 (Pseudo-Diversity) *the term in the law:* `u`, `N_eff = 1 + (N−1)u`, `S = ln N_eff` *the reading becomes the building:* Three honest windows becomes **an audited roster with a measured `u` and a `u`-maintenance regime**

The wave node. *the founding material's teaching:* Hold the future as a weighted set of branches; read amplitudes without preference *the corpus papers that develop it:* CCP-013 (Contagion), CCP-015 (Preregistration) *the term in the law:* threshold distributions; the held-out forecast *the reading becomes the building:* Reading amplitudes becomes **calibrated prediction against a frozen model and four baselines**

Lure. *the founding material's teaching:* Set the idea of greater comfort where the self will pass anyway *the corpus papers that develop it:* CCP-010 (The Room), CCP-012 (Light and Material), CCP-016 *the term in the law:* `eta`, and the encounter rate *the reading becomes the building:* Baiting the horizon becomes **placing the artifact, the seat, and the default on the path already walked**

Uncertainty. *the founding material's teaching:* The veil that measures the measurer; the step that tests the ice *the corpus papers that develop it:* CCP-015, CCP-016 (pilot design) *the term in the law:* measurement reactivity; capped downside *the reading becomes the building:* The minimal probe becomes **the reversible pilot with its cost logged beside its finding**

Collapse. *the founding material's teaching:* The word that closes the sea; burn the drafts; the flinch test *the corpus papers that develop it:* CCP-012 (The Table With No Head), CCP-016 *the term in the law:* independent elicitation before aggregation *the reading becomes the building:* Speaking once becomes **silent-write-first, then commit — collapse each judgment before any of them touch**

Echo. *the founding material's teaching:* Reading the difference; the Mirror that Flatters; the Echo that Devours *the corpus papers that develop it:* CCP-010, CCP-013 (Collective Intelligence), CCP-015 *the term in the law:* prediction error; correlated feedback *the reading becomes the building:* Reading the difference becomes **held-out residuals and the shuffled-input control; the flattering mirror becomes measured error correlation**

Recurse. *the founding material's teaching:* The spiral, not the circle; the winch post; the drift *the corpus papers that develop it:* CCP-007 (Recursion as Circuit), CCP-007 (The Leaky Integrator), CCP-007 *the term in the law:* **the transient term `(1-delta)^t Q_0`** *the reading becomes the building:* Iteration that climbs becomes **the exact solution: the transient decaying while the driven term builds**

Fixed point. *the founding material's teaching:* The room that was always furnished; arrival as recognition *the corpus papers that develop it:* CCP-007, CCP-014, CCP-008 *the term in the law:* **`Q_inf = (eta F(u*) - c*)/delta`, CCP-007 Theorem 3.5** *the reading becomes the building:* Arrival becomes **the steady state of a contraction, computable in advance from four parameters**

Seal. *the founding material's teaching:* Close the working on purpose; the door left ajar; the blank ledger *the corpus papers that develop it:* CCP-014 (shedding Q), CCP-016 (when not to intervene), CCP-008 *the term in the law:* shedding `Q` to recover `h` *the reading becomes the building:* Sealing becomes **the sunset clause, the term limit, the deliberate decommissioning — and the decision not to build**

The table above was written when this framework was one volume. It is not wrong and it is no longer complete: two further volumes turned the same instrument inward and then downward, and every node acquired a second and third address. The extension is below, on the same terms — **what was ceremony became specification, and what was specification became, in four cases, a demotion.**

Stillpoint. *personal scope:* CCP-022 (Launch Budget), CCP-022, CCP-023 F-08 *cellular scope:* CCP-037 (Personal Protocol), CCP-039 *what the extension adds, and what it costs:* The departure lounge becomes `L = max(0, M - C)` — **"there was no surplus to dream with" is `L = 0`.** And "the present is already the local minimum" becomes exact: a practitioner observed at rest is at a stationary point of their own cost field, **by definition and not by luck.** *Magnitude demoted (E.7 item 5); structure stands*

Ledger. *personal scope:* CCP-023 (Personal Formulary), CCP-024, CCP-025 *cellular scope:* CCP-029 (The Two Columns), CCP-029, CCP-037, CCP-038 *what the extension adds, and what it costs:* Every comfort on a card becomes **a register of 156 entries, each with a status that cannot be blurred** and a second column that carries precedent rather than evidence. **133 rows read "none claimed."** The hidden yield becomes the two-track ledger: `Q` against `Q^V`, with stored error `Z` and correction load `C_t = c_0 + kappa*Z_t`

Gradient. *personal scope:* CCP-022 (The Gradient and the Desire Line), CCP-023 F-07 *cellular scope:* CCP-036 (Rewrite Protocol), CCP-037 *what the extension adds, and what it costs:* The worn path becomes `grad h = -grad C`: **a habit survey IS a headroom measurement**, and the comfort zone and the headroom maximum are the same point. The cost: the identity holds only where `eta`, `N` and `delta` are constant in the direction of motion

Horizon. *personal scope:* CCP-021 (The Internal Surplus), CCP-024, CCP-025 *cellular scope:* CCP-037, CCP-038 *what the extension adds, and what it costs:* The many windows of the one tower become `F_int = ln(1 + (N-1)u)` over one observer's own framings — **a derivation, not**

an analogy. The cost: **Break 1.** The independence the aggregation needs is unavailable inside one skull, so `F_int` is an upper bound on what a mind actually holds

The wave node. *personal scope:* — (the node passes through the personal-scope material unaltered) *cellular scope:* CCP-032 (Gabor), CCP-032, CCP-033 *what the extension adds, and what it costs:* Amplitude-weighted branches acquire a conjugate-spread structure and an information-budget bound, **and no quantum content whatsoever.** `dt·df >= 1/(4π)` is classical and measurable on an oscilloscope; `dx·dp >= hbar/2` does zero arithmetic work. **The wave node's branches do not interfere: there is no cross term, and two futures that "cancel" is not a thing this corpus says**

Lure. *personal scope:* CCP-022 §41.7.1, CCP-022 (Worked Example) *cellular scope:* CCP-034 (Placebo), CCP-036 *what the extension adds, and what it costs:* Baiting the horizon becomes **lowering cost at the destination, never raising it at the origin — except locally.** Measured: raising cost globally runs escape time 32 → 35 → 42 → 89 → 297 → never; raising it only at the origin runs **32.0 → 9.9 weeks, at the price of a 37% smaller permanent fund.** The question is whether it is a cost that can be left behind or one that is carried

Uncertainty. *personal scope:* CCP-026 (FVP), CCP-026, CCP-027 *cellular scope:* CCP-032 (Heisenberg, Correctly Stated), CCP-032, CCP-035 *what the extension adds, and what it costs:* The veil that measures the measurer becomes **seven gates in print, and a programme with 21 preregistered kill conditions and a 35.8% completion probability.** The step that tests the ice becomes the seven-step grounding diagnostic, of which **step 5 — the predicted magnitude — does most of the killing**

Collapse. *personal scope:* CCP-019 (The Moment Equation), CCP-019, CCP-026 *cellular scope:* CCP-033 *what the extension adds, and what it costs:* The word that closes the sea becomes `M = (D/P)·E`, with the fixed point `D = P` as undistorted contact. **The cost is this corpus's largest: FVP gate 2 returns DEMOTE, and the magnitude is not reportable pending Rasch instrumentation at 123 items per term.** Decoherence answers the collapse question and answers it *against* licence

Echo. *personal scope:* CCP-023 (The Echo and the Return Difference), CCP-024 *cellular scope:* CCP-035 (Self-Fulfilling Engine), CCP-035 *what the extension adds, and what it costs:* Reading the difference becomes 'Delta_R =

Recurse. *personal scope:* CCP-021 (The Self as Storage), CCP-023 *cellular scope:* CCP-035, CCP-039 *what the extension adds, and what it costs:* The spiral acquires a **runaway condition**: the reflexive storage-law pole `a = (1−delta) + eta*s` with loop gain `eta*s/delta`. The winch post becomes the frozen specification — five operational choices, signed and dated before data. **Unposted rope does not climb; unfrozen models do not measure**

Fixed point. *personal scope:* CCP-023 (The Contraction and Arrival), CCP-019, CCP-024 *cellular scope:* CCP-034 (The Coupling Requirement), CCP-037, CCP-039 *what the extension adds, and what it costs:* The room that was always furnished acquires a **second** fixed point — `D = P` — and a warning attached to it: **headroom there is exactly zero by cancellation, so `r = 1` is not an unqualified target**, and a system optimising raw `Q` is driven away from it. Arrival becomes `E = G·B`: **if `G = 0`, no amount of belief moves anything**

Seal. *personal scope:* CCP-023, CCP-027 (The Research Program) *cellular scope:* CCP-033, CCP-037, CCP-038, CCP-039 *what the extension adds, and what it costs:* Closing the working on purpose becomes **the preregistered kill condition: 21 of them, 4 terminal, published before the first measurement.** The door left ajar becomes the standing rule that a claim shown wrong is **withdrawn, not defended** — and the blank ledger becomes 133 register rows that say *none claimed* and are printed that way

Four of the twelve nodes of the founding material now carry a demotion, and all four descend from the same gate. The stillpoint, collapse and echo nodes, and — through the calibration ratio — the fixed-point node, rest on the Moment Equation, whose magnitude is not reportable. **The structure the founding material taught survives all four demotions intact**, because the founding material taught orderings and directions, which is exactly what ordinal information supports. It is the sequel's attempt to attach numbers that failed the gate, and the

failure is the sequel's, not the parent's.

39.17.2 F.2 the recursion node of the founding material — stated exactly

the founding material's ninth node taught the spiral rather than the circle: repetition that absorbs its own error and returns to the same post a little higher, with a winch post held fixed so that change can be measured against it, and a weekly check for drift.

In this framework, the recursion node is the transient term of CCP-007 Theorem 3.1. The exact solution is

```
Q_t = (1−delta)^t ∗ Q_0  +  SUM_{k=0}^{t−1} (1−delta)^{t−1−k} ∗ (eta∗S_k − C_k)
```

and the first term, `(1−delta)^t ∗ Q_0`, is what the system started with, decaying at rate `delta` no matter what anyone does. **The recursion node is that term decaying while the driven term builds.**

Three consequences the practitioner felt and the engineer computes:

1. **The spiral climbs because the driven term accumulates while the transient does not.** Each pass adds `eta∗S_k − C_k` to the sum, discounted by how long ago it was. The circle repeats because its `eta` is zero and only the transient is present.
2. **Drift has a half-life.** What the system inherited is at half strength after `t_half = ln2 / (−ln(1−delta))` periods. At `delta = 0.05` that is 13.5 periods; at `delta = 0.20`, 3.1. **A new regime feels like the old one for roughly `3∗tau` periods, and impatience before then is a misreading of the arithmetic, not of the people.**
3. **The winch post is the frozen specification.** the founding material held one thing fixed between passes so change could be measured. This framework does the same thing with the five operational choices of Protocol C, fixed in writing and signed before data is collected. Unposted rope does not climb; unfrozen models do not measure.

39.17.3 F.3 the fixed-point node of the founding material — stated exactly

the founding material's tenth node taught the room that was always furnished: the state a loop converges to when every pass shrinks the remaining distance, guaranteed to exist, to be unique, and to be approached geometrically, and recognizable on arrival because in the window-work it had been visited many times before it was reached.

In this framework, the fixed-point node is CCP-007, Theorem 3.5. With constant geometry `U \ensuremath{\equiv} u∗`, constant cost `C \ensuremath{\equiv} c∗`, and `delta > 0`:

```
Q_inf = (eta ∗ F(u∗) − c∗) / delta
```

The correspondence is exact, not analogical. `delta > 0` makes the update map a contraction with factor `(1−delta)`; Banach guarantees a unique fixed point and convergence from any starting state; and the fixed point is that expression, computable in advance from four numbers.

Four consequences:

1. **Arrival is predictable.** The furnished room can be computed before it is entered. A team with `eta = 0.30`, `u∗ = 0.5`, `N = 8`, `c∗ = 0.15`, `delta = 0.05` converges to `Q_inf = 6.02` Q-units whatever its current `Q` is. If it currently holds 8.0, it is not growing toward the room; it is **settling down into it**, and that is not failure — it is the arithmetic of the regime it is running.

2. **The counterfeit arrival has a formal counterpart.** the founding material warned that a well-converged state can still be the wrong one, since convergence is a property of loops and not a certificate of worth. Theorem 3.5 supplies the sharpest version: **if `eta F(u*) < c*`, the fixed point is negative.** The loop still converges. It converges toward the destruction of the thing.
3. **The room is a function of the regime, not of effort.** Changing where a system arrives requires changing `eta`, `u*`, `c*`, or `delta`. Trying harder inside a fixed regime moves `Q` along its path to the same room, faster at best.
4. **Recognition on arrival has an explanation.** In the founding material, arrival felt remembered because the window-work had visited it. Here, the system approaches `Q_inf` geometrically, spending most of its time near it: **by the time arrival is noticed, the system has been close for several time constants.** The texture of recognition is the tail of an exponential.

39.17.4 F.4 What the sequel changes

Three things, stated plainly, so the reader of both books knows where the lineage bends.

The subject changes from the self to the circumstance. the founding material's only licensed configuration was the practitioner running the program on the one it is run on. This framework's subject is a system with other people in it, which is why the consent constraint of the standing refusals of the corpus is an engineering requirement and not only an ethic. **the founding material's vow of consent becomes this framework's fourth refusal, enforced in both books by the same mechanism: coerced change produces resistance, and resistance is a `c` term.**

The register changes from ceremony to specification. the founding material taught by story and treatise in alternation. This framework is a specification with worked practice inside it. The lineage is intact and the modelling is still permitted — but here it is permitted in the open, with its badge on, in Appendix E.2, where it can be audited item by item.

The postulates get harder. the founding material's postulates were about resemblance: a shape shared with Heisenberg's relation, a life descending a comfort landscape, a day's surface read as a leading indicator. This framework's postulates are about measurement: that stored order is a scalar, that it decays exponentially, that cost can be commensurated, that the same equation runs at every scale. **They are more useful and more falsifiable, which is the same thing said twice, and that is why this corpus carries a falsification chapter and the founding material did not need one.**

The ledger in the drawer is blank. It is not empty.

The Standard Manual for Engineering Creature Comforts is issued under the imprint of **36N9 GENETICS LLC / IllumaHeart**. Author: **Michael L. Curzi**. The formal spine is Papers A through E of the Interaction Surplus Framework (Curzi, 2026). Where this framework and the papers disagree, the papers are right and this framework should be corrected in the reader's own hand.

Part III

Frequency, Essence, and Mind

Chapter 40

The Sustaining Frequency: $f_{\min} = k\,e^{2c/\eta}$ as an Equivalence Between the Storage Fixed Point and the Essence Bridge

Abstract. We close a join that two earlier volumes left open. The essence bridge assigns a surplus an exact pitch, $f = k\,e^{2F}$ with $k = (9/8)^2/\varphi^2 = 0.483425732988414$ Hz — a *defined* constant, exact by construction, whose interpretation as a frequency is postulated and not derived — and the storage law's steady state $Q_\infty = (\eta F - c)/\delta$ is positive precisely when $\eta F(u^*)$ exceeds c^*. Inverting the bridge converts the second condition into the first, giving $f > f_{\min}$ with $f_{\min} = k\,e^{2c/\eta}$. The result is an equivalence rather than an implication — a bond accumulates structure if and only if its essence frequency exceeds its sustaining frequency — and it is verified at twelve constructed parameter sets with sign agreement in all twelve. A deficit identity follows: the octave deficit and the steady-state gauge are the same number, $D = 2.885390\,(c/\eta - F) = -2.885390\,(\delta/\eta)\,Q_\infty$, cross-checked by two routes that share no intermediate quantity to a worst residual of 8.88×10^{-16}. Two absences are substantive: $f_{\min}$ contains neither the retention loss, which sets the rate of arrival but not the threshold, nor the category count, which enters only through the ceiling $c/\eta \le \ln N$. The paper then defeats its own headline number. Because $d\ln f_{\min}/d(c/\eta) = 2$, an error of 0.1 in c/η is a 22.14 per cent error in $f_{\min}$, and an order-preserving relabelling of ordinally known parameters moves one system's $f_{\min}$ by a factor of 54.6, or 5.77 octaves. With c and η known only ordinally, $f_{\min}$ is not known as a number at all; only its ordering and the sign of the verdict survive.

Keywords: sustaining frequency; essence bridge; storage fixed point; octave deficit; ordinal parameters; sensitivity analysis

Series: C. Frequency, Essence, and Mind **Paper:** CCP-040

1 Introduction

Two results in this corpus's Series B were established independently, share a variable, and were never combined. The first is a *bridge*: a surplus F, measured in nats on a logarithmic scale of represented difference, is written as a frequency $f = k\,e^{2F}$, so that intervals of surplus become intervals of pitch [Curzi, 2026k]. The second is a *fixed point*: a first-order linear difference equation for stored order settles at $Q_\infty = (\eta F - c)/\delta$ whenever the decay rate is positive [Curzi, 2026a]. Both are stated in the same body of work. Neither refers to the other.

This paper performs the substitution and states exactly what kind of object it produces. Setting the fixed point positive, dividing through by $\delta > 0$ and $\eta > 0$, and inverting the bridge gives a threshold frequency

$$f_{\min} := k\, e^{2c/\eta},$$

together with an exact trichotomy: $f > f_{\min}$ accumulates, $f = f_{\min}$ holds, $f < f_{\min}$ decays. The chain is a chain of *equivalences* under strictly increasing transformations, so either side may be measured and the other inferred, and the result is verified below at twelve parameter sets with sign agreement in every one.

Four consequences follow, and the last two are the ones this paper most wants read. First, a deficit identity: the shortfall in octaves and the steady state itself are the same number in different units, $D = 2.885390\,(c/\eta - F) = -2.885390\,(\delta/\eta)\,Q_\infty$, cross-checked to a residual of 8.88×10^{-16}. Second, two structural absences: $f_{\min}$ contains neither δ, which sets the rate of arrival and not the threshold, nor N, which enters only through the admissibility ceiling $c/\eta \le \ln N$. Third, a sign trap — because the numerator is divided by δ, an intervention that successfully lowers the decay rate of an already-failing relation makes its destination *worse*, and the worked case computes an 80 per cent worsening from a fully successful intervention. Fourth, and against the paper's own headline: because $d \ln f_{\min}/d(c/\eta) = 2$, an order-preserving relabelling of ordinally known parameters moves one system's $f_{\min}$ by a factor of 54.6, or 5.77 octaves. **With c and η known only ordinally, $f_{\min}$ is not known as a number at all.**

What this paper does not claim, stated at the front rather than at the back. The constant $k = (9/8)^2/\varphi^2$ is a *defined* number, exact by construction and therefore legitimately quoted to fifteen figures; its two components are real and documented — the Pythagorean whole tone [Barbour, 1951] and Euclid's extreme and mean ratio [Euclid, 1956] — but **their product expressed as a frequency in hertz has no precedent in any literature we have found**, and the standard deflationary treatment of golden-section claims [Markowsky, 1992] is cited here deliberately and again wherever φ appears. The essence law $f(n) = kn^2$ is a postulate of this framework and not a finding; its exponent is **defined by the model, not estimated from data**. Octave reduction is licensed for pitch [Shepard, 1982] and its extension to a quantity that is not a pitch is a modelling choice with no external warrant. Section 12 decomposes each of these into steps a reader can check individually, and Section 13 prints every check this paper fails.

Contributions

1. The join $f_{\min} = k\, e^{2c/\eta}$, derived as an equivalence with the licensing condition attached to every step, and verified at twelve parameter sets.
2. The deficit identity $D = 2.885390\,(c/\eta - F) = -2.885390\,(\delta/\eta)Q_\infty$, computed by two routes that share no intermediate quantity, agreeing to 8.88×10^{-16}.
3. The two absences — no δ and no N in the threshold — with the practical consequence that decay-side interventions cannot move the diagnostic.
4. The sign trap, computed: a fully successful intervention moving δ from 0.09 to 0.05 takes Q_∞ from -5.0119 to -9.0215 while the deficit stays at $+4.338420$ octaves.
5. The ordinal defeat of the paper's own headline number, by an order-preserving relabelling that moves $f_{\min}$ by 5.77 octaves.
6. An audit-precedence rule, stated as a house postulate rather than an epistemology, together with three specific corrections that the framework's own audit produced and that a keystone published earlier would have inherited.

2 Hypotheses, definitions, and postulates

The rigour standard governing this corpus requires that a claim which *could* be wrong be stated so that it could be, and that a claim which could not be wrong be labelled with which

of the three non-empirical kinds it is. **A hypothesis nothing could falsify is not a hypothesis**, and this section does not manufacture one.

2.1 Hypotheses proper

H1 — the equivalence hypothesis. A relation whose essence frequency exceeds its sustaining frequency accumulates stored order over successive periods, and one whose essence frequency falls below it loses stored order. **Predicted:** the sign of the observed change in the stock over a run of periods agrees with $\text{sign}(f - f_{\text{min}})$ computed at the start of the run from parameters measured on one declared period. **Falsified by:** a system with $f < f_{\text{min}}$, all four parameters measured on one declared period, whose stock rises over a run long enough for the fixed point to dominate the transient. **Status:** DERIVED from the storage law [Curzi, 2026a] and the essence bridge [Curzi, 2026k]; verified as arithmetic at twelve parameter sets. **Not tested against observation.** No observational data were collected for this paper.

H2 — the decay-blindness hypothesis. An intervention acting only on the decay rate δ leaves the octave deficit unchanged and multiplies the existing sign of the fixed point. **Predicted:** in a system with $\eta F < c$, an intervention that measurably lowers δ and leaves η, c, u and N unchanged is followed by a *more* negative steady state. **Falsified by:** a δ-only intervention that reverses the sign of the stock's drift. **Status:** DERIVED by inspection of D; demonstrated on a constructed case in Section 9. **Not tested against observation.**

H3 — the ceiling hypothesis. No relation whose measured cost ratio exceeds the logarithm of its declared category count, $c/\eta > \ln N$, sustains itself under any assignment of non-parallelism over those categories. **Predicted:** among systems measured with a declared and fixed N, none with $c/\eta > \ln N$ is observed to accumulate. **Falsified by:** one counterexample. The hypothesis is a strict universal and is cheap to kill. **Status:** DERIVED from the normalisation bound $F \leq \ln N$ [Curzi, 2026]; **not tested here.** It is the only claim in this paper with genuine empirical exposure.

H4 — the scale hypothesis. Estimates of c and η obtained from unanchored rating instruments carry ordinal information only, so any exponential function of their ratio is not identified as a number. **Predicted:** two admissible order-preserving relabellings of the same estimates yield f_{min} values differing by more than the instrument's claimed precision. **Falsified by:** an instrument for c and η that passes a conjoint-measurement test [Luce and Tukey, 1964] or supplies specific objectivity [Rasch, 1960], after which the relabellings are no longer admissible. **Status:** the *consequence* is COMPUTED in Section 8 and is arithmetic. The *antecedent* — that the framework's instruments are in fact ordinal — is an empirical claim about instruments and is **not tested here.**

H5 — the identifiability claim. From a fitted stock trajectory alone, η and F are not separately recoverable; only the combination $\eta F - c$ is. **Status:** this is a **formal claim, not a hypothesis**, and it is labelled as one. It is demonstrated in Section 9.5 by exhibiting two parameter sets with identical ledgers and different prescriptions. Nothing could falsify it that would not also refute the arithmetic.

2.2 Definitions — true by stipulation, not falsifiable

- **Def. 1.** $u(x, y) := 1-(\hat{x}\cdot\hat{y})^2$ on unit intensity vectors over N categories; $F := \ln(1+(N-1)u)$; $N_{\text{eff}} := 1 + (N-1)u = e^F$.
- **Def. 2.** $f_{\text{min}} := k\, e^{2c/\eta}$, the *minimum sustainable frequency*.

- **Def. 3.** $D := 2.885390\,(c/\eta) - \log_2(f/k)$, the *octave deficit*.
- **Def. 4.** $k := (9/8)^2/\varphi^2$. **This is a defined constant.** It is an exact ratio of exact quantities and is therefore computable to arbitrary precision; the fifteen significant figures quoted throughout, 0.483425732988414, are *exact digits of a definition* and carry no measurement uncertainty because there is no measurement. Every appearance of the constant at that precision in this paper is so labelled.
- **Def. 5.** A relation is *sustainable* when $Q_\infty > 0$. This is a choice among available definitions, discussed at Step 0 of Section 6.

2.3 Postulates — asserted by this framework, with no external support

- **P-a. The essence law $f(n) = kn^2$.** A quadratic relation between a dimensionless count and a frequency. **No external literature supports it.** It is not derived here and nothing in the bibliography derives it. **The exponent 2 is defined by the model, not estimated from data.** No scaling exponent in this paper was obtained by fitting; none was tested against a log-normal or exponential alternative; and the goodness-of-fit and likelihood-ratio procedure that a claimed scaling relation requires [Clauset et al., 2009] was **not run, because there is no data to run it on.**
- **P-b. The attachment of hertz to k.** Computing $(9/8)^2/\varphi^2 = 0.483425732988414$ is arithmetic. **Asserting that this dimensionless number *is a frequency* is a separate and much stronger claim**, and the source supplies no operational procedure by which the constant acquires the dimension of inverse time. This is postulated, and Section 3.4 states it again in the dimensional audit, where it is unavoidable.
- **P-c. The affine effective count $N_{\text{eff}} = 1 + (N-1)u$.** A design axiom [Curzi, 2026], not a derived result and **not a Hill number of any order**. The Hill family [Hill, 1973, Jost, 2006] and its independent rediscovery in political science [Laakso and Taagepera, 1979] ground $1/\text{HHI}$ and $\exp(H)$ exactly — as the corpus's concentration work uses them [Curzi, 2026] — and ground this construction not at all. A reviewer reaching for the Hill-number literature as support will find it does not apply.
- **P-d. The identification $n_{\text{eff}} := N_{\text{eff}}$.** The essence law takes a count; the bridge feeds it an effective count that is not an integer. That substitution is a modelling decision of this framework and is badged as one wherever the bridge is used.
- **P-e. The audit-precedence rule.** A result may not be published above the level of audit available to check its inputs. **This is a house postulate adopted by this corpus, not a general epistemology, and it is not offered as one.** It is stated because it explains the ordering of the argument, and Section 8 gives the arithmetic that motivates it.

2.4 Modelling choices — legitimate, but choices

- **M-a. Octave reduction applied to a non-acoustic quantity.** Octave equivalence is a perceptual structure formalised for **pitch** [Shepard, 1982], and the octave's perceptual specialness has a measured and finite size [Demany et al., 2013]. Nothing licenses carrying it to a quantity that is not a pitch. The *operation* is well defined — division by a power of two — and its *interpretation* is a choice of this framework.
- **M-b. $N = 12$.** Twelve component classes are used throughout. Twelve exceeds the classical span of immediate absolute judgment [Miller, 1956], which is a reason for caution and not a reason for twelve. The ladder itself is owned by the companion paper [Curzi, 2026], whose Result Six is the count-based essence index.
- **M-c. Interval apparatus.** Cents, octave equivalence and interval arithmetic are the correct apparatus once an index exists [Rossing et al., 2002], but what a listener finds consonant depends on the spectrum of the sound and not on the ratio alone [Sethares, 2005], and interval categories generalise less far across listeners than is usually assumed [Burns, 1999, Hoeschele et al., 2012]. **This paper uses the arithmetic and claims none of the perception.**

- **M-d. The sustainability definition.** $Q_\infty > 0$ is chosen over the alternative state-dependent condition $Q_{t+1} \geq Q_t$ because it is parameter-determined rather than state-determined. Both are defensible; the choice is declared at Step 0 of Section 6.

2.5 What is expository

Sections 4.1 to 4.17 restate results owned elsewhere in the corpus so that this paper stands alone; they argue no new claim. Sections 6.7 to 6.10 and 9.3 are procedural: they set out reporting requirements and failure modes. Their organising question is *what must be published alongside a reading for the reading to be checkable by a third party*.

3 Framework, dimensional audit, and method

3.1 The storage law — a shared mathematical form, not a derivation

$$Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t, \qquad S_t = F(U_t) = \ln(1 + (N-1)U_t)$$

Analogy status: category 2, a shared mathematical form. The equation governing this framework's stored order is the same equation that governs a leaky integrator, and it is used because the mathematics is the same, not because the physics transfers. The same form appears as the discrete RC low-pass filter in signal processing [Oppenheim et al., 1996, Horowitz and Hill, 2015], as exponential smoothing in forecasting [Holt, 1957, Brown and Meyer, 1961], and as the leaky integrate-and-fire neuron in computational neuroscience [Lapicque, 1907, Gerstner and Kistler, 2002]. Its stability condition and unique fixed point are standard [Elaydi, 2005], and the convergence is Banach's [Banach, 1922]. **No claim is made that comfort is a charge, that η is a coupling constant, or that any physical result about RC circuits applies to a relation between people.** The corresponding transfer-function reading, in which $1/\delta$ is the DC gain of $1/(z-a)$ evaluated at $z=1$, is fixed in Series B [Curzi, 2026a], whose Result One states the law with its Lipschitz modulus $1-\delta$ and its steady state $Q^* = (\eta S - C)/\delta$.

Three warnings attach to that form and are recorded here rather than in a footnote. First, the time average of a multiplicative process differs from its ensemble average, so a parameter estimated across individuals does not describe what happens to any one of them over time [Peters, 2019] — and Q accumulates *over time*, for *one* relation, from parameters that would most naturally be estimated from a cross-section. Second, a single Q, δ or u written for a group is an aggregation, and aggregation over heterogeneous agents does not in general yield a representative agent [Kirman, 1992]. Third, Q **has the grammar of a conserved quantity — a stock that accumulates, leaks and is spent — without a conservation law to give that grammar content.** That is the characteristic failure of the genre: neoclassical economics was built by copying nineteenth-century energetics and made utility play the role of energy without the conservation law that gave energy its meaning [Mirowski, 1989]. This paper claims no conservation law, states no invariance under which Q would be conserved, and adopts instead the position that a law is claimed only within a specified arrangement rather than universally [Cartwright, 1999]. The wider programme — statistical physics applied to social dynamics — has a standard survey that fixes what such models can and cannot be held to [Castellano et al., 2009], and the boundaries it draws are the boundaries claimed here.

The precedent that matters most for this paper is the leaky integrate-and-fire neuron: a leaky store with a drive and a threshold, yielding a rate below which the store does not fire. **That literature grounds the *shape* of $f_{\min}$ — a threshold rate emerging from a leaky store — and it grounds neither the formula nor the constant.** The distinction is made here and repeated in Sections 11 and 12.

3.2 The surplus functional and the effective count

$F(u) = \ln(1 + (N-1)u)$ is the corpus's surplus functional, characterised by angle-only dependence, a zero at parallel states, an affine effective count and a normalisation [Curzi, 2026], whose Result One is exactly that characterisation with $N_{\text{eff}} = 1 + (N-1)u$. Its logarithmic form places it in the same family as Shannon's entropy [Shannon, 1948], and uniqueness-under-axioms is a known genre of theorem in that literature [Aczél et al., 1974, Khinchin, 1957]; the framework's axiom set is its own and is not one of those. The ceiling $F \leq \ln N$ is the normalisation, and it is the same ceiling that appears in requisite-variety arguments [Ashby, 1956, 1958].

Which construct of difference u measures must be declared, because at least three incompatible constructs travel under the word *diversity* — separation, variety and disparity [Harrison and Klein, 2007]. Here u is a *disparity* measure over represented states, and it is symmetric in its arguments, which psychological similarity is known not to be [Tversky, 1977]. That is a limitation and it is printed in Section 13.

The warrant the logarithm actually needs. The corpus habitually justifies its logarithms by appeal to Fechner [Fechner, 1860, Weber, 1834]. A stronger and cheaper warrant is available: a logarithmic form follows from multiplicative dynamics with **no psychological assumption at all** [Peters and Gell-Mann, 2016]. The argument of this paper needs only the second. The psychophysical warrant is required only where an *elicitation* is performed, which is not done here.

3.3 The essence bridge and the constant k

The bridge writes a surplus as a pitch, so that intervals of surplus become intervals of frequency and the constant cancels from every ratio [Curzi, 2026k], whose Result Six badges the identification $n_{\text{eff}} := N_{\text{eff}} = e^F$ and fixes the conversion $2F/\ln 2 = 2.885390081777927 \times F$ octaves above base:

$$f = k\, e^{2F} = k\, N_{\text{eff}}^2, \qquad k := \frac{(9/8)^2}{\varphi^2} = \frac{81/64}{\varphi + 1} = 0.483425732988414 \quad \text{(defined; exact)}$$

Both components are real, specific and documented. The ratio $9{:}8$ is the Pythagorean whole tone, a historically attested object of tuning theory [Barbour, 1951, Lindley and Turner-Smith, 1993], and $81/64$ is the Pythagorean ditone. The golden section is Euclid's extreme and mean ratio [Euclid, 1956], and $\varphi^2 = \varphi + 1$ with $\varphi^m = F_m\varphi + F_{m-1}$ are proved identities [Vajda, 1989, Koshy, 2001]. Verified here: $\varphi^2 - \varphi - 1 = 0.000\text{e}{+00}$ in double precision.

Their product, expressed as a frequency in hertz, has no precedent. There is no physical, psychoacoustic or musicological result that combines a tuning ratio with an inverse square of the golden section to obtain a ground frequency. The literature on golden-section claims is overwhelmingly a literature of *deflation*: the appearances of φ in the Parthenon, the Great Pyramid, the human body and preferred rectangles do not survive examination [Markowsky, 1992, Fischler, 1981, Livio, 2002], and the psychological preference results are weaker and more equivocal than they are usually reported to be [Höge, 1997].

This framework's position is the one those critiques leave open and no wider. φ is being *chosen* as a constant, not *found* in a psychological system. Nothing in this paper is evidence that the golden section is present in anything measured. The constant sets a reference pitch and is invisible to every verdict, because D is a difference of logarithms in which k appears in both terms and cancels exactly.

3.4 Dimensional audit

symbol	quantity	dimension	unit	status
Q	stored order	$[\mathcal{C}]$	none supplied	state variable; **no operational unit**
S	surplus input per period	$[\mathcal{C}\,T^{-1}]$	per declared period	driving term
C, c	cost per period	$[\mathcal{C}\,T^{-1}]$	per declared period	driving term
δ	decay rate	$[1]$	per period	dimensionless; $0 < \delta \leq 1$
η	conversion efficiency	$[1]$	—	dimensionless; assumed $0 < \eta \leq 1$
u	non-parallelism	$[1]$	—	dimensionless by construction, $u \in [0, 1]$
F	surplus	$[1]$	nats	dimensionless; a logarithm of a count
N	declared category count	$[1]$	—	absolute scale, a count
n	component index	$[1]$	—	absolute scale, a count
k	reference constant	$[T^{-1}]$ **asserted**	Hz **asserted**	**defined as a ratio (dimensionless); the unit is postulated**
$f, f_{\min}$	essence frequency, threshold	$[T^{-1}]$	Hz	inherit their dimension entirely from k
D	octave deficit	$[1]$	octaves	a ratio of logarithms; k cancels

Does the equation balance? The storage law balances provided S and C carry the same dimension as Q per period, which the framework assumes and does not measure; Q has **no operational unit**, and this audit records that rather than inventing one. The essence law $f(n) = kn^2$ balances only if k carries $[T^{-1}]$, because n is a dimensionless count. **The source supplies no operational procedure by which k acquires that dimension.** The constant is computed as a ratio of two dimensionless quantities — a frequency ratio and a geometric ratio — and the hertz is attached by assertion. This is the single largest dimensional exposure in the paper and it is listed in Section 13 by name.

The exponent is the second exposure, and it is not the same one. The ratio c/η is dimensionally $[\mathcal{C}\,T^{-1}]$ divided by a pure number, and is therefore **not dimensionless as written**, which the exponent of $e^{2c/\eta}$ requires it to be. The framework's convention is that c is expressed in the same units as F — nats of surplus per declared period — which makes c/η dimensionless. **That convention is a choice, the source never states it, and every reading in this paper depends on it being applied consistently.** It is stated here because it is otherwise invisible, and its failure mode is measurable: comparing two scopes on mismatched clocks produced a $90.1\times$ period-normalisation error in the corpus's own correspondence audit (finding CC-R16), and at two nats of amplification per unit of c/η that is not an error bar but a different answer.

3.5 Scale type, and what licenses the arithmetic

quantity	scale type claimed	what licenses it
n, N	absolute	cardinality: true zero, fixed unit, identity transformation group only
u	ratio	constructed from an inner product of unit vectors; zero is parallel
η, c, δ	ratio, **assumed**	no instrument is supplied; the quantities admit a true zero, but whether an *estimate* of them is ratio-scaled is a property of the instrument
f, $f_{\min}$	ratio in hertz, **inherited**	inherits the status of n and of the postulated bridge
D	interval, in octaves	a difference of logarithms

Stevens' typology makes the permissible arithmetic on a quantity a function of the transformation group that leaves its empirical content unchanged [Stevens, 1946], and division requires a ratio scale with a true zero. **The strongest available objection to this framework is Michell's**: that psychology routinely assumes its attributes are quantitative without ever testing whether they are [Michell, 1997, 1999]. The objection applies here in full. The formal remedy exists — conjoint measurement supplies a testable condition for additive quantitative structure [Luce and Tukey, 1964, Krantz et al., 1971], and a Rasch model supplies specific objectivity [Rasch, 1960, Andrich, 1988] — and **neither has been run for any quantity in this paper.** η, c and δ are treated as quantitative by stipulation, and Section 8 computes what that stipulation is worth when it fails.

3.6 Identification — the objection this framework cannot presently answer

The framework's founding construction is a reflexive loop, $T(x) = W(x, M(x))$: the group influences the member and the members constitute the group [Curzi, 2026], whose Result Two identifies the self-fulfilling prophecy as a fixed point of that composition of seeing and responding. **That is a linear-in-means reflexive specification, and the endogenous social effect in such a specification is not identified from observational data** [Manski, 1993]. It cannot be separated from contextual effects or from correlated unobservables, however much data is collected; the reinforcement is that influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011]. **No amount of fit, held out or otherwise, recovers the parameter.**

The remedy this paper adopts is stipulation, and it pays the price openly. The reference set is fixed in advance — the declared category set of size N and the declared decomposition rules — and every reading is reported together with that stipulation, which is what makes a discrete-choice treatment of social interaction defensible where an unrestricted one is not [Brock and Durlauf, 2001]. **This paper therefore presents the reflexive loop as a modelling position, not as an identified empirical result.**

The same objection applied to the framework's own parameters. If systems change their behaviour when modelled, then δ, η and u are the decision rules of people who have not read the storage law, and they are not structural in the sense a policy exercise requires [Lucas, 1976]. Classification changes the classified [Hacking, 1995], and economic theory has been shown to become self-validating by being taught and used [Ferraro et al., 2005, MacKenzie and Millo, 2003]. **The framework asserts reflexivity for its subject matter and cannot consistently exempt its own parameters from it.** This paper does not resolve that; it states it, and treats every parameter as conditional on the arrangement in which it was fixed.

3.7 Method

Every number in this paper is the output of exact arithmetic on stated parameters, computed in double precision. **No observational data were collected, no human subjects participated, and no empirical dataset underlies any figure or table.** Every case is constructed and is labelled as constructed at the point of use. Residuals quoted at the 10^{-16} level are floating-point round-off, not measurement error, and are reported to demonstrate that two algebraically identical expressions were in fact computed independently and agreed. The figures are produced by the scripts named in their captions.

4 The four constituent laws

This section restates the results the join consumes, with their honesty clauses, so that the paper stands alone. **A law's honesty clauses are part of the law**: a theorem detached from the conditions under which it was proved is not a compressed theorem but a different and false statement. Four laws are needed, and only one of them is empirically at risk.

```
THE ESSENCE LAW      f = k n^2                        an index has a pitch
THE SURPLUS LAW      F = ln(1 + (N−1)u)               a pair has a surplus
THE STORAGE LAW      Q_{t+1} = (1−d)Q_t + eta S − C    surplus becomes structure, and leaks
THE TIME−FREQUENCY   dt ∗ df >= 1/(4 pi)              what may be rendered, and how fast
BOUND
```

4.1 The essence law — statement

$$f(n) = k\,n^2, \qquad k = \frac{(9/8)^2}{\varphi^2} = 0.483425732988414 \quad \text{(defined; exact)}$$

with n the *essence index*: the position of a type within its discipline's ordered category set. **The constant is a Pythagorean interval divided by the square of the golden section, and that is a design decision, badged as a postulate (P-b) every time it appears.**

Provenance and verification. The law is a collapse of a recipe of the form $((n/\varphi) \times 1.125)^2$. Recomputed here over $n = 1 \ldots 118$, that recipe and kn^2 agree to a **worst relative deviation of** 4.760×10^{-16} — machine precision. The residual bias in the source's own published table traces to φ stored at fifteen significant figures (finding CC-R5).

4.2 The essence law — two structural results

Result One — every interval is a ratio of squares, and k cancels.

$$\frac{f(m)}{f(n)} = \frac{m^2}{n^2}$$

The golden section and the $9{:}8$ set the absolute reference pitch **and nothing else**. Verified: $2{:}1 \to 4/1$ (two octaves), $3{:}2 \to 9/4$, $4{:}3 \to 16/9$, $5{:}4 \to 25/16$, $3{:}1 \to 9/1$. **This is the strongest available defence of an otherwise arbitrary constant: it does no work anywhere except at the origin.** The result is owned by Series B [Curzi, 2026k].

Result Two — octave folding is governed by the odd part of the index. Since $(2n)^2 = 4n^2$ is exactly two octaves above n^2, write $n = 2^a m$ with m odd. Then $f(n) = k(2^a m)^2 = f(m) \cdot 2^{2a}$, an integer number of octaves. Two indices therefore share a pitch class **if and only if** they have the same odd part, and a set indexed $1 \ldots N$ contains exactly $\lceil N/2 \rceil$ distinct classes. For $N = 12$ the odd parts are $1, 1, 3, 1, 5, 3, 7, 1, 9, 5, 11, 3$, giving $\{1, 3, 5, 7, 9, 11\}$ — **six classes from twelve indices**, computed and confirmed. The result is owned in this series by the folding paper [Curzi, 2026], whose Result One is the odd-part class rule.

4.3 The essence law — what it does not do

- **It does not describe a physical resonance.** No material, molecule, element, room, institution or person is claimed to resonate at any of these values.
- **It does not measure k.** k is chosen. Its form is a design decision and carries no empirical warrant, and the deflationary golden-section literature is cited alongside it [Markowsky, 1992].
- **It does not survive an arbitrary index.** An arbitrary ordering yields arbitrary frequencies.
- **It does not gain support from the square-well coincidence.** That $E_n \propto n^2$ for a one-dimensional infinite square well is a genuine agreement of form. **Coincidence of form is not evidence of mechanism**, and asserting the connection would be exactly the decorative analogy this corpus's analogy discipline exists to prevent.
- **It does not generalise to compounds.** The compound rule is $f = k^2 G^4$ with G the count-weighted geometric mean of component indices, audited in this series by the operations paper [Curzi, 2026], whose Result Four proves the two rules agree at no integer index.

4.4 The surplus law — statement and axioms

$$u(x, y) = 1 - (x \cdot y)^2 \in [0, 1], \qquad F = f(u) = \ln(1 + (N - 1)u)$$

uniquely characterised by four axioms [Curzi, 2026]: **S1** geometric dependence, F depends on the pair only through u; **S2** zero at zero, $f(0) = 0$; **S3** affine effective count, $g(u) = e^{f(u)}$ is affine in u; **S4** normalisation, $f(1) = \ln N$.

S3 is the design axiom. It is chosen, not derived, and every result that leans on it must say so — and this entire paper leans on it, because N_{eff} is defined by it. The uniqueness argument is four lines and every step after S3 is forced:

```
S1  =>  F is a function of u alone:            F = f(u)
S3  =>  g(u) := e^{f(u)} is affine in u:       g(u) = A + Bu
S2  =>  f(0) = 0, so g(0) = 1, so A = 1
S4  =>  f(1) = ln N, so g(1) = N, so B = N - 1

therefore  g(u) = 1 + (N-1)u   and   f(u) = ln(1 + (N-1)u)
```

The bite of S3, computed. $g(u) = 1 + (N - 1)u$ is affine: at $u = 0.1, 0.2, 0.3$ it takes the values $2.1, 3.2, 4.3$, second difference -4.44×10^{-16}, machine zero. An alternative satisfying S1, S2 and S4 but violating S3 is $f_{\text{alt}}(u) = u \ln N$, with $g_{\text{alt}}(u) = N^u$ whose second difference at the same points is $+0.102021$ — not affine. **Both agree at both endpoints and disagree by 33.62 per cent at $u = 0.5$ (1.871802 against 1.242453).** S3 is not a technicality: it selects one curve out of a family that agree at the endpoints and diverge by a third in the middle, and the middle is where every real case lives.

The functional form is underdetermined by the evidence, and this paper concedes it. A power law with exponent below one [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the qualitative facts as well as the logarithm does, and many mechanisms produce indistinguishable curves [Stumpf and Porter, 2012]. **Nothing in this paper discriminates among them.**

4.5 The surplus law — derived properties, all computed at $N = 12$

```
monotone           f'(u) = (N-1)/(1 + (N-1)u) > 0 everywhere
strictly concave   f''(u) = -(N-1)^2/(1+(N-1)u)^2 < 0
                   f''(0) = -121.0000     f''(1) = -0.840278
sharp Lipschitz    f'(0) = 11.0000 = N-1  f'(1) = 0.916667
subadditive on u   f(0.3) + f(0.4) = 3.145014 >= f(0.7) = 2.163323
```

```
extremal          f(0) = 0 exactly;  f(1) = 2.484907 = ln 12 exactly
range             0 <= F <= ln N
```

$\ln N$ **is the hard ceiling on surplus per interaction**, and Section 7's ceiling result is this inequality converted into frequency units. The affine count also coincides exactly with the Rényi-2 participation ratio, $N_{\text{eff}} = N + 1 - N/\text{PR}$, verified to 10^{-12} (finding CC-R3) — the Rényi family being the source of that measure [Rényi, 1961]. **That coincidence is worth having and is not evidence that S3 is true.** It is evidence that S3 selects a quantity other people had independent reasons to care about, which is a different and weaker thing.

4.6 The surplus law — what it does not do

- **It is a LOWER BOUND on the ensemble's information, not the information itself.** Every frequency computed from it is therefore a lower bound on the pitch, not the pitch.
- **It is not Shannon entropy and not mutual information in general.**
- **It proves nothing about thermodynamics.** No negentropy law is proved.
- **The triangle inequality fails for geometric triples.** Subadditivity holds on the parameter u only; pitch distances across three essences do not chain.
- **It is degenerate on one-hot inputs.** Two states each concentrated on a single distinct category are automatically orthogonal: computed, $u = 1.000000$ and $F = 2.484907 = \ln N$, regardless of which categories they are. **A reading in which everyone is maximal is a reading of the instrument**, and the folding paper carries this as its Result Five [Curzi, 2026].
- **It says nothing about people until someone supplies the vectors**, and supplying the vectors is the hard part. The monotone claim in particular cannot be extended past the range examined: the empirical literature is divided, with measured returns to bridging [Granovetter, 1973] but creative success rising and then *falling* with the same variable [Uzzi and Spiro, 2005], short-run effects of surface diversity frequently negative [Putnam, 2007], and a rebuttal literature that qualifies rather than reverses that [Abascal and Baldassarri, 2015]. The problem-solving result usually cited for diversity [Hong and Page, 2004] has a standing formal objection [Thompson, 2014] that must be cited with it.

4.7 The essence bridge

Within a discipline n is an index; across disciplines there is no shared index. The framework supplies one by identification (P-d):

$$n_{\text{eff}} := N_{\text{eff}} = 1 + (N-1)u = e^F, \qquad f = k\, n_{\text{eff}}^2 = k\, e^{2F}$$

Taking base-2 logarithms, **octaves above base** $= 2F/\ln 2 = 2.885390081777927 \times F$**, with** F **in nats.** Verified at both ends: $F = 0$ gives $N_{\text{eff}} = 1$ and $f = k = 0.4834$ Hz; $F = \ln 12 = 2.484907$ gives $N_{\text{eff}} = 12$ and $f = 69.6133$ Hz, exactly 7.169925 octaves above base.

The bridge is a **strictly increasing bijection** from $[0, \ln N]$ onto $[k, kN^2]$, which is the property the join needs in order to invert it. Round-trip verification $F \to f \to F$ at ten values including both endpoints returns the input to fifteen decimal places with error $0.000\text{e}{+00}$ at every point tested:

```
F = 0.000000  ->  f =  0.483426  ->  F' = 0.000000000000000
F = 0.563085  ->  f =  1.490799  ->  F' = 0.563085000000000
F = 1.000000  ->  f =  3.572060  ->  F' = 1.000000000000000
F = 2.000000  ->  f = 26.394151  ->  F' = 2.000000000000000
F = 2.153633  ->  f = 35.888195  ->  F' = 2.153633000000000
F = 2.484907  ->  f = 69.613306  ->  F' = 2.484906649788000
```

4.8 The storage law — statement, solution and bounds

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t          S_t = ln(1 + (N−1)U_t)

admissibility     C_t <= (1−delta)Q_t + eta S_t     guarantees Q_{t+1} >= 0
exact solution    Q_t = (1−delta)^t Q_0 + SUM (1−delta)^{t−1−j} (eta S_j − C_j)
one−step growth   Q_{t+1} > Q_t  iff  eta S_t > delta Q_t + C_t
uniform bound     Q_t <= max{ Q_0, eta ln N / delta }
sustainability    a nondecreasing step requires C_t <= eta ln N − delta Q_t
steady state      Q_inf = ( eta F(u*) − c* ) / delta        (delta > 0)
```

At $\eta = 0.30$, $\delta = 0.09$, $N = 12$ the uniform bound is 8.2830. **If C_t exceeds $\eta \ln N - \delta Q_t$, no choice of interaction geometry can prevent order from decreasing** — this is the ceiling in stock units, and Section 7 restates it in frequency units.

The master gauge is *headroom*, $h_t := \eta \ln N - \delta Q_t - C_t$, negative meaning unsustainable, zero meaning maintenance only, positive meaning growth is available. At the fixed point $h_\infty = \eta(\ln N - F(u^*)) - \delta$, C and Q all cancel (finding CC-R6). Checked at $\eta = 0.30$, $c = 0.62$, $\delta = 0.09$, $u = 0.068735$: the identity gives 0.5765465541125843 and the direct computation $\eta \ln N - \delta Q_\infty - c$ gives 0.5765465541125842. The institutional worked case in Series B moves a fixed point from 0.2653 to 3.9281 by moving η from 0.35 to 0.55 and δ from 0.25 to 0.12, which is the same gauge in an operational setting [Curzi, 2026h].

4.9 The storage law — the circuit reading and geometric convergence

Because $\delta > 0$ makes the map a contraction, Banach's theorem guarantees a unique fixed point and convergence to it from any start [Banach, 1922]; the corpus's foundations state the theorem exactly and then refuse the certificate it appears to offer, on the ground that a human life is not a metric space and nobody has measured the relevant contraction modulus [Curzi, 2026]. That refusal binds this paper. The convergence is geometric, $|Q_t - Q_\infty| = (1-\delta)^t |Q_0 - Q_\infty|$, verified at $\delta = 0.09$, $t = 20$ to fifteen digits; the transient half-life $\ln 2/(-\ln(1-\delta))$ is 7.3496 periods at $\delta = 0.09$. The half-life formula and its divergence from the naive $1/\delta$ are owned by the folding paper [Curzi, 2026], whose Result Four records 1.9434 cycles against a naive 3.3333 at $\delta = 0.30$, an error of 72 per cent.

An equivalent multiplicative reading exists and agrees to fourteen significant figures: in world coordinates the storage law is a geometric accumulator $G_{t+1} = G_t^{1-\delta}(N_{\text{eff}})^{\eta} e^{-C}$ [Curzi, 2026n], whose Result Three fixes that correspondence. **This paper uses the additive form throughout.**

4.10 The storage law — what it does not do

- **It proves no thermodynamic or physical law.** It is a linear difference equation.
- **It is scale-free in FORM only.** That only (δ, η, N, C) change between a desk, a room, a team and a city is a claim about model form, not about physics. The correspondence audit names nine disanalogies between internal and external scopes, three of them computed: a mean-field decay error reaching **92 per cent by** $t = 12$; two nine-member groups with identical $u = 0.5999$ and $S = 1.7578$ losing **77 per cent versus 21 per cent** of surplus to a three-person departure; and a $90.1\times$ period-normalisation trap on mismatched clocks (finding CC-R16).
- **It does not license a negative stock.** When $\eta F < c$ the fixed point is negative and the recursion violates its own admissibility condition on the way there. Section 6.8 computes exactly when.
- **It does not average.** Because F is strictly concave, feasibility must be tested on the mean of $F(U_t)$ and not on F of the mean U (finding CC-R2, worth 6.2 per cent in the corpus's worked office; the Jensen correction is owned by the protocol paper [Curzi, 2026], whose

Result Four prices one instance at 0.0790 nats, 4.31 per cent). And because distortion cancels across sites, aggregating a ratio across sites overstated bankable input by $4.26\times$ (finding CC-R10).

4.11 The time–frequency bound on rendering

$$\Delta t \cdot \Delta f \geq \frac{1}{4\pi} = 0.079577471545948$$

is a theorem about signals [Gabor, 1946], not an analogy, and it binds every rendering of any frequency in this paper. **It contributes no term to $f_{\min}$.** It is stated here because a paper that produces frequencies and does not state the constraints on rendering them is inviting exactly the misuse it claims to refuse. Four consequences are computed and owned by the conjugate paper [Curzi, 2026z], whose Result One proves the bound by the same variational argument as $\sigma_x \sigma_p \geq \hbar/2$ with $p = \hbar\kappa$ deleted:

```
minimum-gap cliff   N =  3..6   min gap = 9/8       = 203.9100 cents
                    N =  7..16  min gap = 50/49     =  34.9756 cents
                    N >= 17     min gap = 289/288   =   6.0008 cents   <-- the cliff
                    N >= 41     min gap = 1682/1681 =   1.0296 cents

rate factor         (1/32)/(1/256) = 8 EXACTLY      symbols per second
resolution          32/(4 pi) = 2.546479 cycles per step below the cliff
                    256/(4 pi) = 20.371833 cycles per step above it
audible-set ceiling a single integer transposition admits at most N = 25
```

Three distinct ratios attach to the cliff, all correct in their own units and all different: $\ln(50/49)/\ln(289/288) = 5.828475$ in cents, $288/49 = 5.877551$ in cycles, and exactly 8 in rate. The corrected cliff location is $N = 17$, not $N = 20$, because $17^2 = 289$ folds beside $288 = 2 \cdot 12^2$ the instant index 17 enters the set (finding CC-R21). $N = 12$ is below the cliff, so this paper renders — if it renders at all — in the $50/49$ regime.

The uncertainty relation is a fact about state preparation, not about measurement disturbance, and the corpus's own uncertainty node refuses the inflated reading of its borrowed physics [Curzi, 2026]. Its counter-result requires that a user be able to distinguish the theorem, the ordinary observer effect, and the framework's own conjecture in three separate sentences on demand. That requirement binds this paper too.

4.12 The four-register convergence

Four quantities in this framework are logarithmic: surplus $F = \ln N_{\text{eff}}$; pitch $p = \log_2 f$; sensation [Fechner, 1860]; and the moment in log form. **Every conversion between them is therefore linear in log space, and the conversion constants are exact:**

```
1 nat    = 1.442695 bits  = 2.885390 octaves = 3462.4681 cents
1 bit    = 0.693147 nats  = 2 octaves EXACTLY = 2400 cents exactly
1 octave = 0.346574 nats  = 0.5 bits          = 1200 cents
```

One octave is exactly half a bit of surplus. The nat, the bit and the conversion between them are the standard units of a logarithmic information measure [Shannon, 1948, Cover and Thomas, 2006]; **nothing about the framework's use of them is novel, and nothing about them licenses the frequency reading.**

4.13 The twelve-state ladder, and the objection to it

The index this corpus feeds to the essence law is n, a count of distinguishable components. The ladder is owned by the companion paper [Curzi, 2026] and is reproduced here because the worked example uses it:

n	state class	component added	f = k n^2	folded 220–440	class
1	undifferentiated activation	arousal alone	0.4834	247.514 (j=9)	1
2	valenced activation	+ valence	1.9337	247.514 (j=7)	1
3	directed feeling	+ intentional object	4.3508	278.453 (j=6)	3
4	situated feeling	+ situation	7.7348	247.514 (j=5)	1
5	attributed feeling	+ cause	12.0856	386.741 (j=5)	5
6	temporal feeling	+ projected duration	17.4033	278.453 (j=4)	3
7	reflexive feeling	+ self as object	23.6879	379.006 (j=4)	7
8	social feeling	+ another mind modelled	30.9392	247.514 (j=3)	1
9	normative feeling	+ standard applied	39.1575	313.260 (j=3)	9
10	narrative feeling	+ placed in a story	48.3426	386.741 (j=3)	5
11	reciprocal feeling	+ the other modelling you	58.4945	233.978 (j=2)	11
12	committed feeling	+ future action bound	69.6133	278.453 (j=2)	3

Six distinct pitch classes $\{1, 3, 5, 7, 9, 11\}$, exactly $\lceil 12/2 \rceil$, as the odd-part theorem requires.

The affect literature does not underwrite this ladder, and the reason must be stated in the body rather than in a footnote. The named precedents establish that each component is a real, independently described feature of emotional states; they do not establish an index, an ordering, or a nesting.

- **Appraisal dimensions vary simultaneously and independently, and emotions differ by the pattern across them** [Smith and Ellsworth, 1985]. Two states at the same dimensionality are separated by the *pattern of values*, which is exactly the information a count discards. **A count cannot separate anger from guilt; a pattern can.**
- **Emotion is a process unfolding in time** — a sequence of evaluation checks driving synchronised change across organismic subsystems, where the state *is* the synchronisation [Scherer, 2001]. A process with an ordering of its own resists reduction to a scalar of any kind, and the single point at which the two orderings touch is a coincidence at one rung.
- **Individuation by core relational theme is qualitative** [Lazarus, 1991]. A theory in which two emotions differ by *what* is appraised does not license a theory in which they differ by *how many* things are appraised.
- **The circumplex is two continuous bipolar axes recovered from ratings** [Russell, 1980]; its "two" is two dimensions, not two countable components, and its principal content — the circular ordering of affect terms — has no counterpart in the ladder.
- **The top rung is contradicted specifically.** The ladder places action tendency at $n = 12$, as the last component added. The standing thesis in that literature is that changed action readiness is **constitutive of emotion as such** [Frijda, 1986] — a startle at $n = 1$ is already action readiness. **If that is right, the component sits at the floor and the ladder's ordering fails at its summit.**

Two of the seven named precedents survive in the form the corpus needs: a claim that a *class* of states requires a component other states lack is exactly the shape of argument a rung needs [Tracy and Robins, 2004], and the modelled-other construct is cleanly named [Premack and Woodruff, 1978]. **This paper therefore cites the affect literature for the *existence* of a component and never for the index, its ordering, or its nesting**, and narrows the top rung to *bound commitment* rather than *readiness*, which is the only version that literature leaves standing. The unnarrowed claim is conceded in Section 13.

These objections and Michell's are the same objection in two domains. Michell's is that a psychological attribute is treated as quantitative without a test for quantitative structure; the

appraisal objection is that an emotional state is reduced to a scalar when its own literature describes it as a pattern or a process. **Both are instances of one habit — scalar reduction — and this corpus should notice that rather than meeting each as a surprise.** The count argument answers Michell for n alone, because a cardinality's structure is a consequence of its definition rather than an empirical hypothesis; it answers neither objection for anything else in this paper.

4.14 What each law contributes to the join, and what it withholds

LAW	contributes to f_min	withholds
essence law	the constant k, and the square, via the bridge	any warrant that k is other than chosen
surplus law	F, and the ceiling ln N	the identity of F with information; F is a lower bound
storage law	eta, c, delta, and the sign condition	any claim that Q is conserved, physical, or that a negative Q_inf is a stock
Gabor bound	nothing at all	it constrains only the rendering

4.15 What would falsify each law

The essence law is a postulate and is not falsifiable as stated. It is a stipulation about how an index maps to a pitch. What *is* checkable is its internal consistency with the recipe it collapses — checked, at 4.760×10^{-16} — and the consequences of its interval structure, which are arithmetic. **A postulate that cannot be falsified must earn its place by being useful and by carrying its badge.**

The surplus law is a theorem and is falsifiable only in its axioms. If a system's effective count is not affine in u, S3 is wrong for that system and the whole functional form changes.

The storage law is falsifiable in the ordinary empirical sense, and the corpus states the test: preregister, fit on training data, freeze, report held-out performance, and beat four baselines. That protocol is the standard one [Nosek et al., 2018, Simmons et al., 2011], and the corpus's own attempt at it **failed its second falsification rule** [Curzi, 2026i], whose Result One records a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin. **A framework whose one confirmatory fit failed is in a better epistemic position than one that never attempted a fit**, and this paper treats that failure as the relevant precedent for how any future fit should be reported.

The time–frequency bound is a theorem of signal analysis and is not in question.

The asymmetry is the honest headline. Two of the four laws are stipulations, one is a theorem about axioms someone chose, and only the third is empirically at risk. **The join is therefore exactly as empirical as the storage law and no more; everything else in it is bookkeeping.**

5 The join: statement and derivation

5.1 The two inputs, in the exact form the derivation uses

INPUT ONE	Q_inf = (eta F(u∗) − c∗) / delta	for constant u∗, c∗, delta > 0
INPUT TWO	f = k e^(2F)	k = 0.483425732988414 (defined; exact)

For the rest of this section the asterisks are dropped: u means u^*, c means c^*, and the regime is the constant one the fixed-point theorem requires.

5.2 The derivation, step by step, with the licence for each step

Step 0 — the definition of sustainability. A relation is SUSTAINABLE when $Q_\infty > 0$. **This is a definition and it is a choice** (M-d). It is not the only possible one: requiring $Q_{t+1} \geq Q_t$ at the current stock is a different, state-dependent condition. The fixed-point definition is chosen because it is parameter-determined rather than state-determined — it says where the relation is going, not where it happens to be.

Step 1 — clear the denominator.

```
Q_inf > 0   <=>   ( eta F − c ) / delta > 0   <=>   eta F − c > 0
```

Licence: $\delta > 0$; a positive divisor preserves the direction of an inequality. **This is the step where δ leaves the argument and never returns**, and Section 7.1 shows that its departure is the most consequential structural fact in the paper. If $\delta = 0$ there is no fixed point and the theorem does not apply; Section 6.6 handles that case and shows the same threshold governs it anyway.

Step 2 — isolate the surplus.

```
eta F − c > 0   <=>   eta F > c   <=>   F > c/eta
```

Licence: $\eta > 0$. If $\eta = 0$ the numerator is $-c \leq 0$ and no F satisfies the inequality unless $c = 0$ also. The quantity c/η now has a name and a meaning: the **cost-to-conversion ratio**, the amount of surplus in nats that must be supplied per period merely to pay for the period. **It is dimensionless only under the convention stated in Section 3.4, and that convention is a choice the source never states.**

Step 3 — invert the bridge.

```
f = k e^(2F)   =>   ln(f/k) = 2F   =>   F = (1/2) ln(f/k)
```

Licence: $k > 0$ and $f > 0$, both guaranteed since $f = ke^{2F} \geq k > 0$. The map $F \mapsto ke^{2F}$ is a strictly increasing bijection from $[0, \ln N]$ onto $[k, kN^2]$, so the inverse exists on exactly that range and is itself strictly increasing. Verified numerically in Section 4.7 at ten values with error $0.000\text{e}{+00}$.

Step 4 — substitute and solve.

```
F > c/eta   <=>   (1/2) ln(f/k) > c/eta   <=>   ln(f/k) > 2c/eta
            <=>   f/k > e^(2c/eta)        <=>   f > k e^(2c/eta)
```

Licence: multiplication by 2, exponentiation, and multiplication by $k > 0$ are each strictly increasing, **so the equivalence is exact in both directions and nothing is lost.**

5.3 The result

DEFINITION 5.1 — THE MINIMUM SUSTAINABLE FREQUENCY.

$$f_{\min} := k\, e^{2c/\eta}$$

THEOREM 5.2 — THE THREE REGIMES. Under the fixed-point theorem's conditions ($\delta > 0$, $\eta > 0$, constant u and c), with $f = k\, e^{2F(u)}$:

```
f > f_min    <=>   Q_inf > 0    the relation ACCUMULATES order
f = f_min    <=>   Q_inf = 0    it HOLDS, and nothing more
f < f_min    <=>   Q_inf < 0    it DECAYS, whatever anyone intends
```

Proof: Steps 1 through 4 read as a chain of equivalences, the equality case following from the strict monotonicity of every transformation used. **QED.** The trichotomy and its consequences are owned in this series by the criterion paper [Curzi, 2026], whose Result One states $f_{\min} = k\, e^{2c/\eta}$ with the trichotomy attached.

COROLLARY 5.3 — THE HOLD LINE. The relation holds exactly when $u^* = (e^{c/\eta} - 1)/(N-1)$. Verified at four points for $N = 12$:

```
c/eta = 0.500000  ->  u* = 0.058975  ->  f = f_min =  1.3141 Hz
c/eta = 1.000000  ->  u* = 0.156207  ->  f = f_min =  3.5721 Hz
c/eta = 2.000000  ->  u* = 0.580823  ->  f = f_min = 26.3942 Hz
c/eta = 2.484907  ->  u* = 1.000000  ->  f = f_min = 69.6133 Hz
```

The last row is the ceiling appearing on its own: at $c/\eta = \ln N$ the required non-parallelism is exactly 1 — perfect orthogonality — and there is nothing above it.

5.4 What this result is

It is an exact restatement of a sustainability condition in frequency units. Nothing more. It is a change of variable applied to an inequality. The inequality was already established; the change of variable was already established. **The join contributes no new assumption, no new parameter, and no new empirical content, and if it appeared to, that would be a sign that an error had been made.**

Three things it does contribute, all real and none empirical. **(a) A comparable register.** Intervals compose and hertz do not; two relations with different costs are compared on the octave scale directly, and the exchange rate is one exact constant. **(b) An audible ceiling.** $F \leq \ln N$ is an abstract bound; $f \leq 69.6133$ Hz for twelve categories is a specific position, and the fact that a system's requirement can be *above* that position is what makes the ceiling result land. **(c) A single diagnostic number**, the octave deficit, which collapses four parameters into one signed quantity whose sign is the verdict and whose magnitude is the size of the problem.

5.5 What this result is not

It is not a claim that anything physically resonates. No person, bond, tissue, molecule, room or institution is asserted to vibrate, emit, or respond at f, at $f_{\min}$, or at any value derived from them. The frequencies are essence frequencies: positions on a logarithmic scale of surplus.

It is not a measurement. It is a function of four parameters and has exactly the status of those parameters. Section 8 shows what that status is when the parameters are ordinal.

It is not new physics, biology or psychology. The empirical content lives entirely in η, c, u and N. **The algebra cannot be right or wrong about the world; it can only be right or wrong about arithmetic**, and Section 6 checks the arithmetic.

It is not a clinical instrument, and it is not a target to play. Nothing here diagnoses, assesses or treats anybody, and nothing licenses the inference that sounding a tone at $f_{\min}$ does anything to anyone.

5.6 Two asymmetries the join inherits

Asymmetry one — the verdict is conservative in one direction only. As deployed, the surplus functional is a **lower bound** on the ensemble's information. F is exact given u, but u is a property of the *representation*, and a representation that fails to capture a real difference understates u and therefore understates f. **A verdict of ACCUMULATES is robust to representational under-capture; a verdict of DECAYS is not.** The instrument is biased toward pessimism, and every negative reading should be checked against the question *have we represented the difference between these parties, or only the difference our instrument can see?*

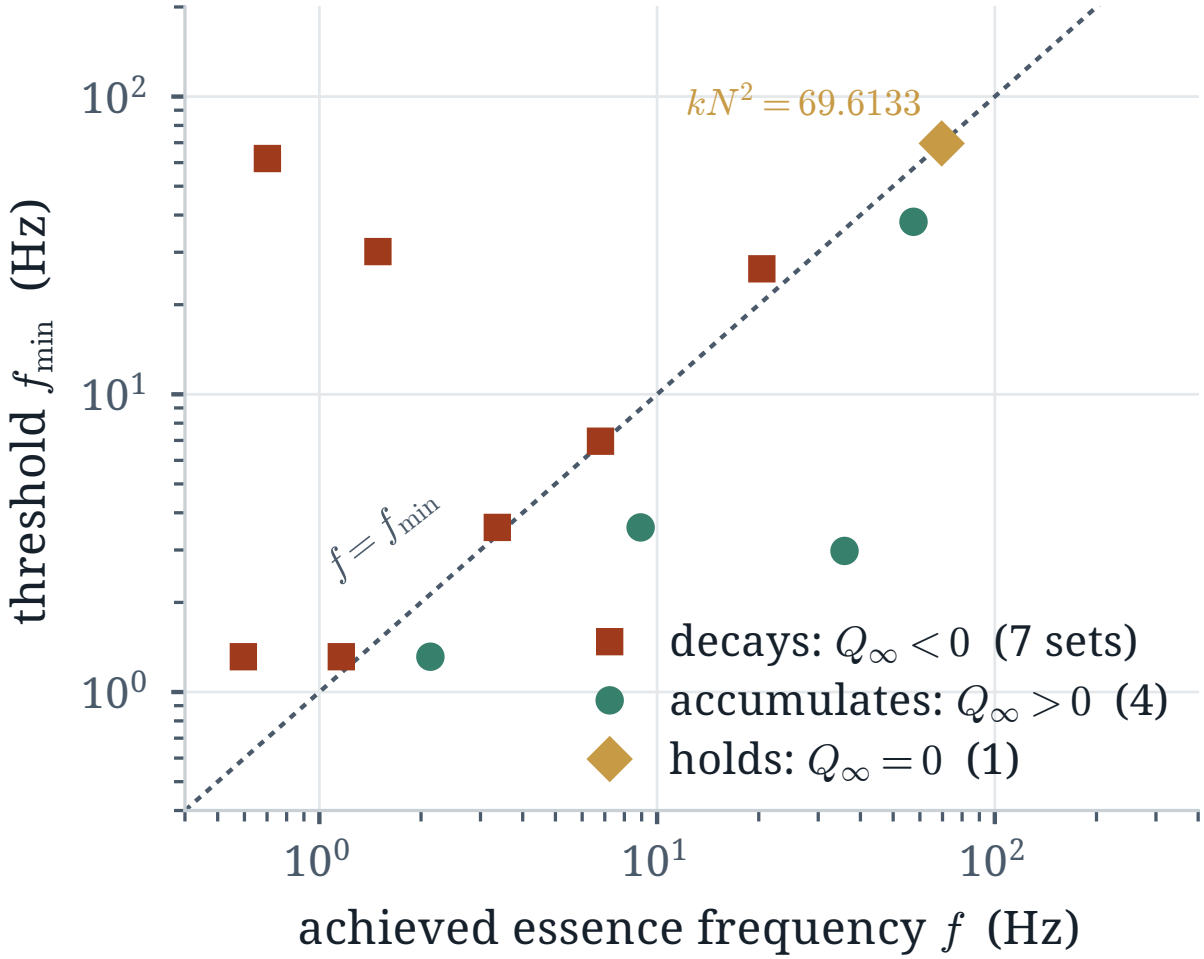


Figure 40.1: **The equivalence holds at every one of the twelve constructed parameter sets**, including the boundary case: the seven sets with $Q_\infty < 0$ all sit above the line $f = f_{\min}$ and the four with $Q_\infty > 0$ all sit below it, with the single hold set exactly on it at $f = f_{\min} = kN^2 = 69.6133$ Hz. Modelled from $f = k(1 + 11u)^2$ and $f_{\min} = k\,e^{2c/\eta}$ with the **defined** constant $k = 0.483425732988414$ and $N = 12$; the parameter sets are constructed and are the source's. Built by `figs/src/CCP-040_fig01.py`.

Asymmetry two — the ratio is period-invariant only on a shared clock. η is a conversion per period and c a cost per period; their ratio is dimensionless and invariant under a change of period *provided both were measured on the same one.* The measured failure is a $90.1\times$ period-normalisation error (finding CC-R16). Since $d \ln f_{\min}/d(c/\eta) = 2$, that is not a large error bar; it is a different answer.

6 Verification of the join

6.1 Twelve parameter sets

The theorem is an equivalence, so the verification is a sign test: for each parameter set, compute f and $f_{\min}$ independently, compute Q_∞ independently, and check that $\mathrm{sign}(f - f_{\min})$ equals $\mathrm{sign}(Q_\infty)$. All twelve sets use $N = 12$ and **all twelve are constructed**; the first three are the carried cases of this paper and the remaining nine span the space, including two near-zero cases and one exact-equality case at the ceiling.

```
  eta      c     delta        u          F          f (Hz)     f_min (Hz)  sgn(f−f_min)   Q_inf      sgn  agre
```

0.30	0.620	0.09	0.068735	0.563085	1.4908	30.1588	−1	−5.0119	−1	yes
0.44	0.400	0.09	0.692373	2.153633	35.8882	2.9782	+1	+6.0844	+1	yes
0.34	0.825	0.09	0.018614	0.186272	0.7017	61.9349	−1	−8.4630	−1	yes
0.50	0.500	0.10	0.300000	1.458615	8.9385	3.5721	+1	+2.2931	+1	yes
0.60	0.300	0.05	0.100000	0.741937	2.1319	1.3141	+1	+2.9032	+1	yes
0.20	0.100	0.20	0.050000	0.438255	1.1614	1.3141	−1	−0.0617	−1	yes
0.90	1.800	0.02	0.500000	1.871802	20.4247	26.3942	−1	−5.7689	−1	yes
0.75	1.000	0.15	0.250000	1.321756	6.7982	6.9574	−1	−0.0579	−1	yes
0.10	0.050	0.30	0.010000	0.104360	0.5956	1.3141	−1	−0.1319	−1	yes
0.55	1.200	0.08	0.900000	2.388763	57.4358	37.9693	+1	+1.4227	+1	yes
0.40	0.400	0.12	0.148900	0.969983	3.3639	3.5721	−1	−0.1001	−1	yes
1.00	2.484907	0.07	1.000000	2.484907	69.6133	69.6133	0	+0.0000	0	yes

All twelve agree, including the boundary case. The last row is the ceiling: $u = 1$, $F = \ln 12$, $c/\eta = \ln 12$, and $f = f_{\min} = kN^2 = 69.6133$ Hz exactly, with $Q_\infty = 0$. The relation holds and does not build, at the maximum surplus twelve categories can produce.

> **A correction to the source, printed rather than silently applied.** The source prints the last row's cost as $c = 2.485$. That value is not the ceiling: $2.485 > \ln 12 = 2.484906649788$, and recomputing from the printed figure gives $f_{\min} = 69.6263$ Hz, $Q_\infty = -0.0013$ and $D = +0.000269$, not the exact zeros the row reports. **The row is correct only for** $c = \ln 12$ **exactly**, and this paper quotes that value. The discrepancy is a rounding in the printed table, not an error in the theorem, and it is recorded because a reader recomputing from the printed intermediate will find it.

6.2 The deficit identity

DEFINITION 6.1 — THE OCTAVE DEFICIT.

$$D := 2.885390\,(c/\eta) - \log_2(f/k)$$

with $D \le 0$ sustainable as it stands and $D > 0$ short by that many octaves. The instrument is owned by the criterion paper [Curzi, 2026].

THEOREM 6.2 — THE DEFICIT IDENTITY.

$$D = 2.885390\,(c/\eta - F) = -2.885390\,(\delta/\eta)\,Q_\infty$$

Proof. $\log_2(f/k) = \log_2 e^{2F} = 2F/\ln 2 = 2.885390\,F$, giving the first equality. For the second, $Q_\infty = (\eta F - c)/\delta$ gives $(\delta/\eta)Q_\infty = F - c/\eta$, hence $c/\eta - F = -(\delta/\eta)Q_\infty$. Substitute. **QED.**

Verified at all twelve parameter sets, computing D from the frequencies and $-2.885390(\delta/\eta)Q_\infty$ from the fixed point, **entirely separately**:

eta	c	delta	D (from frequencies)	−2.885390∗(delta/eta)∗Q_inf	error
0.30	0.620	0.09	+4.338420	+4.338420	8.88e−16
0.44	0.400	0.09	−3.590989	−3.590989	0.00e+00
0.34	0.825	0.09	+6.463837	+6.463837	8.88e−16
0.50	0.500	0.10	−1.323283	−1.323283	4.44e−16
0.60	0.300	0.05	−0.698084	−0.698084	2.22e−16
0.20	0.100	0.20	+0.178159	+0.178159	2.78e−17
0.90	1.800	0.02	+0.369901	+0.369901	5.00e−16
0.75	1.000	0.15	+0.033406	+0.033406	6.31e−16
0.10	0.050	0.30	+1.141576	+1.141576	0.00e+00
0.55	1.200	0.08	−0.597116	−0.597116	6.66e−16
0.40	0.400	0.12	+0.086610	+0.086610	5.55e−17
1.00	2.484907	0.07	−0.000000	−0.000000	8.88e−16

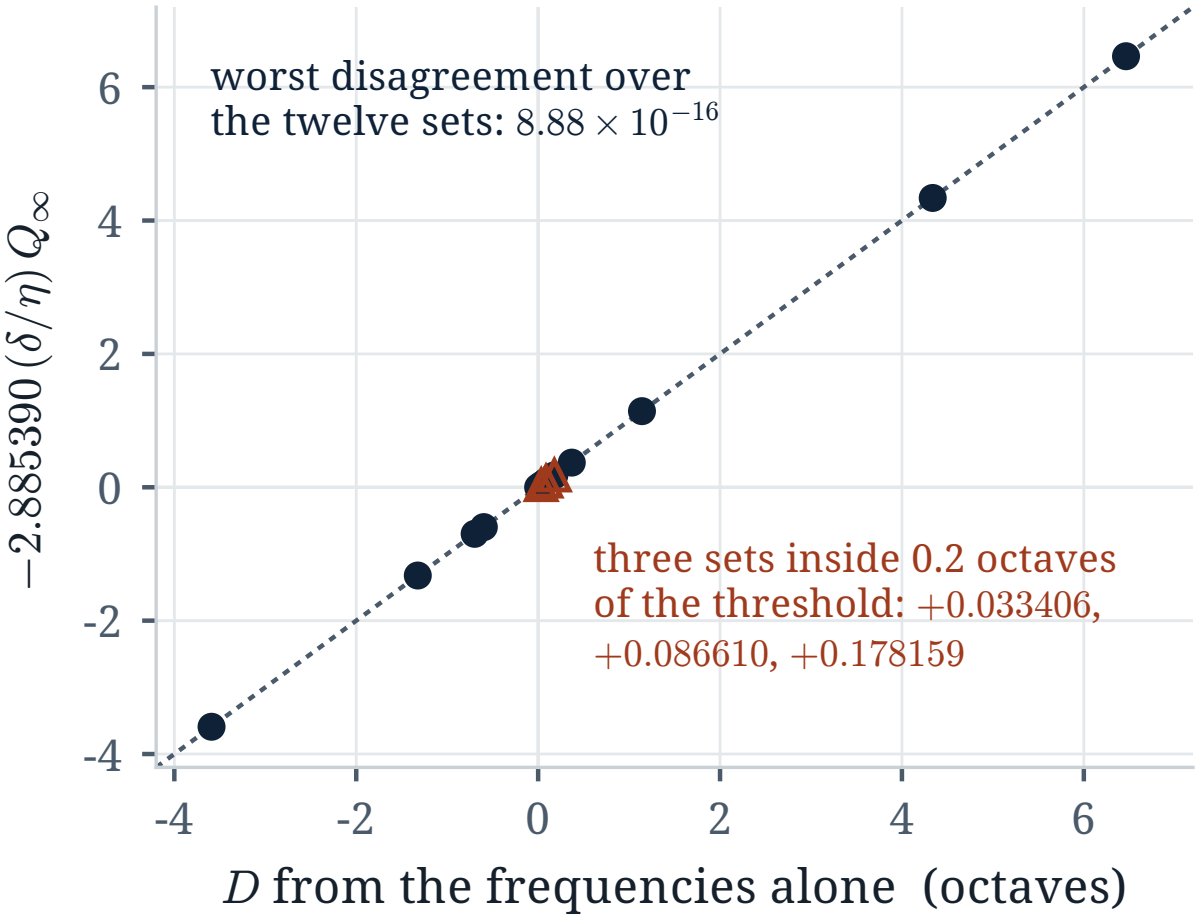


Figure 40.2: **Two routes that share no intermediate quantity agree to** 8.88×10^{-16}, which is double-precision round-off and not measurement error: the abscissa is built only from frequencies and the ordinate only from the storage fixed point. Three of the twelve constructed sets lie within 0.2 octaves of the threshold — $+0.033406$, $+0.086610$ and $+0.178159$ — and are the cases the diagnostic exists for, because nothing about them looks wrong from inside. Modelled from the twelve-set table of Section 6.1 with the **defined** constant $k = 0.483425732988414$. Built by `figs/src/CCP-040_fig02.py`.

The octave deficit and the fixed point are the same quantity in different units. Three readings of the identity, and all three matter. **(i)** $D = 2.885390(c/\eta - F)$: the deficit is the surplus shortfall, converted; one nat short is 2.885390 octaves short, and there is no other exchange rate. **(ii)** $D = -2.885390(\delta/\eta)Q_\infty$: a pitch reading and a stock reading are one reading in two units. **(iii)** k **cancels.** D is a difference of logarithms in which the chosen constant appears in both terms, so **the verdict does not depend on** k **at all**; only the absolute hertz values do.

The conversion factor between the two gauges contains δ**. The deficit itself does not.** That single asymmetry is the whole content of Section 7.1.

6.3 The join is an equivalence, and that is worth more than an implication

Because the chain is a chain of equivalences, **either side may be measured and the other inferred.** That gives two routes to the same verdict with different measurement costs:

```
ROUTE A -- through the geometry
  elicit the state vectors -> compute u -> compute F -> compute f
```

```
    then measure eta and c   -> compute f_min ;  compare

ROUTE B -- through the ledger
  measure the stock Q_t over several periods -> fit delta, eta, c
  compute Q_inf directly -> convert via D = -2.885390 (delta/eta) Q_inf
```

Route A needs no history; Route B needs no elicitation. **A system where both are available is a system where the framework can be checked against itself**, which is the single most valuable thing an equivalence buys: the two routes share no measurement, so agreement is evidence and disagreement localises the fault.

This is also where identifiability bites (H5). Route B recovers η and c only as the combination $\eta F - c$ unless F is independently known — a fitted trajectory cannot separate *converts well from a poor geometry* from *converts poorly from a good one*. **Route B alone gives the verdict and not the prescription.** Section 9.5 exhibits two decompositions that are indistinguishable to any ledger and imply different next moves.

6.4 Reading the near-boundary rows

Rows 6, 8 and 11 have deficits of $+0.178159$, $+0.033406$ and $+0.086610$ octaves — 214, 40 and 104 cents. A relation 40 cents below its threshold is failing in a way no participant could plausibly notice: the per-period loss is $0.0579 \times 0.15 = 0.0087$ units of stored order. Over one period it is nothing; over a hundred periods it is the whole difference between a practice that persists and one that must be rebuilt each generation.

This is the specific value the arithmetic adds over judgement. Large deficits are visible without arithmetic. Small deficits are invisible without arithmetic and decisive over time, and they are the reason to compute rather than to sense. **Proximity to zero is itself a reading**, and the reporting rule that follows is to publish the deficit *and* its sensitivity to the parameter one is least sure of.

6.5 The two absences, and both are substantive

```
f_min = k e^(2c/eta)          contains: k, c, eta
                              does NOT contain: delta, N, u, Q_0, t
```

The absence of δ. δ is the leak. Trust, ritual, place, written record and institution all reduce it. **None of them changes the threshold a relation must clear in order to be building anything at all.** An intervention on δ does not move the reading by any amount whatsoever, and Section 7 computes the case. **A gauge that fails to move under an intervention is not a broken gauge; it is a gauge telling you that you have not touched the mechanism it measures.**

The absence of N. N does not appear in $f_{\min}$ either, but it appears in the *achieved* f through $F = \ln(1 + (N-1)u)$, and in the ceiling $c/\eta \leq \ln N$. **N is not a lever on the threshold; it is a lever on the reachable range.** Raising N does not lower the bar — it raises the highest bar that can be cleared, and it is not free: coordination paths grow as $m(m-1)/2$ [Brooks, 1975, Conway, 1968], and the corpus's own intervention work computes an N lever and then declines it net, a ceiling gain of $\eta(\ln 6 - \ln 5) = +0.1914$ against five added communication paths costing $+0.0714$ [Curzi, 2026j]. Roster-scale accounting is owned in this series by the group paper [Curzi, 2026], whose Result Two gives the exact roster ceiling $m(m-1) \leq 2\eta \ln N/\kappa$.

6.6 The $\delta = 0$ case: no fixed point, same threshold

The fixed-point theorem requires $\delta > 0$. At $\delta = 0$ the map is not a contraction, there is no fixed point, and the exact solution reduces to $Q_t = Q_0 + t(\eta F - c)$: the stock grows without bound if the drift is positive and falls without bound if it is negative. **There is no Q_∞ to be positive or**

negative, and the threshold is unchanged, because the sign of the drift is the sign of the same numerator. Verified on two carried cases:

```
eta = 0.30, c = 0.62, u = 0.068735:  drift = −0.451075   f =  1.4908 < f_min = 30.1588   agree
eta = 0.44, c = 0.40, u = 0.692373:  drift = +0.547598   f = 35.8882 > f_min =  2.9782   agree
```

So $f_{\min}$ **governs both regimes: it decides where a leaky system settles, and it decides which way a lossless system drifts.** That is a stronger statement than the trichotomy alone, and it is available precisely because δ fell out at Step 1. **Zero loss is a pathology, not an ideal** — nothing ever decays out of the stock, so nothing ever arrives.

6.7 Limits, asymptotics and the domain of the model

```
c = 0                 f_min = k exactly. Any u > 0 gives f > k, so ANY relation sustains.
                      A costless relation is sustained by the smallest difference.

c > 0, eta -> 0+      f_min -> infinity. No pitch suffices. At eta = 0.01, c = 0.5,
                      f_min = 1.2995e+43 Hz. A system that converts nothing into
                      structure cannot be saved by any arrangement of its parts.
                      THIS LIMIT IS NOT EXAMINED IN THE SOURCE and is a real
                      constraint on the model's domain: it places every relation
                      above every ceiling.

delta -> 0+           no fixed point; the stock is unbounded in the sign of the drift.
                      Section 6.6. A pathology, not an ideal.

delta -> 1            Q_inf = eta F − c; the stock is memoryless and equals one
                      period's net drive. The transient half-life goes to zero.

u -> 0                f = k, the ground essence. Sustains only if c = 0. And u = 0
                      arises from perfect agreement AND from perfect opposition, since
                      u depends only on |rho|.

u -> 1                f = k N^2 = 69.6133 Hz at N = 12. The maximum. Sustains iff
                      c/eta < ln N.

N -> infinity         the ceiling rises without bound, which the coordination-cost
                      literature says is not available in practice.

c/eta = ln N          f_min = k N^2. The ceiling. Section 7.2.
c/eta > ln N          the condition is unsatisfiable. NO arrangement works.
```

Every divergence corresponds to a real, nameable phenomenon, which is what a limit audit requires; a formula whose limits are nonsense is wrong even where it looks reasonable. **Two of these limits — $\eta \to 0^+$ and $\delta \to 1$ — are not examined in the source and are stated here for the first time.** The $u \to 0$ degeneracy is owned in this series by the routes paper [Curzi, 2026], whose Result One establishes that $\rho = +1$ and $\rho = -1$ both give $u = 0$, $F = 0$ and $f = k$: **a pair agreeing on everything and a pair agreeing on nothing produce the identical reading.**

6.8 Where the linear model stops applying, computed

Take the carried case as found: $\eta = 0.30$, $c = 0.62$, $\delta = 0.09$, $u = 0.068735$, so $\eta F = 0.168926$ and the per-period drift is -0.451075. Start at $Q_0 = 1.0$ and check admissibility $C_t \leq (1 - \delta)Q_t + \eta S_t$ at each step:

```
t = 0:  Q_t = +1.0000   right-hand side =  1.0789   c = 0.62   ADMISSIBLE
t = 1:  Q_t = +0.4589   right-hand side =  0.5865   c = 0.62   FAILS
t = 2:  Q_t = -0.0335   right-hand side =  0.1385   c = 0.62   FAILS
t = 3:  Q_t = -0.4815   right-hand side = -0.2693   c = 0.62   FAILS
```

Admissibility fails at $t = 1$. The stock goes negative at $t = 2$. The fixed point -5.0119 is never reached and was never going to be. A negative Q_∞ is therefore **a direction and a rate, not a destination**: it says the relation is losing order, it says how fast, and it says that the linear model will stop describing the system before the model's own answer arrives. What happens past $t = 1$ is outside this model — the parties renegotiate, reduce cost, or leave. **The model's job was to say that something must give, and it did that at $t = 1$.**

6.9 The formula through the validation protocol, gate by gate

The corpus's rule is that no formula appears without passing all seven gates in print. **$f_{\min}$ is a formula and the rule applies to it.**

gate	what it asks	verdict
G1 dimensions	units agree; ratios dimensionless	**CONDITIONAL** — both sides are hertz *given* P-b; c/η is dimensionless only under the convention of Section 3.4
G2 scale type	division needs a ratio scale	**CONDITIONAL** — passed by the user's instrument or not at all; Section 8 prices the failure
G3 limiting cases	behaviour at 0 and ∞ names a real phenomenon	**PASS** — Section 6.7, with two limits added that the source omits
G4 fixed points	solve for the invariant; state what it means	**PASS** — the invariant is $f = f_{\min}$, equivalently $Q_\infty = 0$, equivalently $u^* = (e^{c/\eta} - 1)/(N - 1)$
G5 identifiability	can parameters be recovered from obtainable data	**PARTIAL** — the verdict yes, the prescription no (Section 9.5)
G6 falsification	preregister, freeze, report held-out, beat baselines	**INHERITED** from the storage law, whose one attempt failed its own rule [Curzi, 2026i]
G7 reproducibility	same operationalisation, different analysts	**NOT ESTABLISHED**

A keystone that publishes its own open gates is more usable than one that claims seven passes. G1 and G2 are conditional, G5 is partial, and G7 has not been run. That is four of seven not cleanly passed, and it is printed here rather than in an appendix.

The baseline suite the G6 test would use has three measured holes, and a paper that recommends the test must publish them:

```
CC-R13a  the shuffled-input control is ANTI-CONSERVATIVE as ordinarily run.
         Against nominal alpha = 0.05 over 500 null trials:
             iid permutation    0.108   (2.2x nominal)
             circular shift     0.120   (2.4x nominal)
             block permutation  0.060   (near nominal)
```

```
          The matched-autoregressive comparison alone false-passed 35.8% of the time.
CC-R4     the shuffle cannot see repackaging. A filtered repackaging passed all four
          baselines (p = 0.0005) and was caught only by the structural ceiling test.
CC-R24    the shuffle cannot see a shared slow driver: it passed 98.35% of the time in
          a 4,000-trial simulation while detecting the clean case 48.17% of the time.
```

A measured false-positive rate of 0.108 against a nominal 0.05 belongs in the abstract of the paper that measured it, not in a footnote, and the replacement rule is *design controls, not more surrogates.* The corpus's survey work supplies the structural detector that the shuffle lacks [Curzi, 2026i].

6.10 The reading, as a procedure

```
1  FIX N.          Declare the category set and its ordering. State that the frequency
                   inherits exactly the defensibility of the count.
2  REPRESENT.      Elicit state vectors with OVERLAPPING SUPPORT. One-hot vectors give
                   u = 1.000000 for everyone and read the instrument, not the subject.
3  COMPUTE u.      u = 1 - (x . y)^2. One construction, used everywhere.
4  COMPUTE f.      F = ln(1 + (N-1)u);  f = k e^(2F). Report the folded value and its j.
5  MEASURE eta, c ON ONE CLOCK. Declare the period. Declare the scale status.
6  CHECK THE CEILING FIRST.  Is c/eta <= ln N ?
                  NO  -> the case is structural. Issue no tuning prescription.
                  YES -> continue.
7  COMPUTE f_min = k e^(2c/eta) AND D = 2.885390 (c/eta) - log2(f/k).
8  CROSS-CHECK.  Q_inf = (eta F - c)/delta, and verify D = -2.885390 (delta/eta) Q_inf.
                  If a negative Q_inf is reported, state the period at which
                  admissibility fails.
```

Step 6 is placed where it is deliberately. A ceiling check performed after the lever analysis is performed too late: a prescription has already been produced, and prescriptions are hard to withdraw once said out loud. The ceiling is cheap to test and it changes the *kind* of answer, so it goes first.

7 The two structural absences

7.1 Decay is not a lever on the deficit

THEOREM 7.1. D contains no δ.

Proof: by inspection of Definition 6.1 and of $f = k(1 + (N-1)u)^2$; neither contains δ. **QED.**

Consequence, and it is the single most practically important statement in this paper: an intervention that acts only on δ cannot change the deficit by any amount whatsoever. Trust, ritual, continuity and written record all act on δ. **They are not thereby useless — they change the rate of arrival and the resilience of the stock — but they do not change whether the relation is building anything at all.**

The sharper form is the sign trap. Because $Q_\infty = (\eta F - c)/\delta$ and δ appears in the denominator, **when the numerator is negative, reducing δ makes Q_∞ more negative.** Trust does not add order; it amplifies whatever sign the system already has. In the corpus's worked consortium this is measured: trust alone moved a fixed point from -0.0836 to -0.2114, reciprocity alone to $+0.7394$, and both together to $+1.8701$ (finding CC-R15). That consortium is carried in the trust paper [Curzi, 2026s], whose Result Four records the case at $u = 0.52$ collecting 73.8 per cent of its ceiling with a steady state of -0.0836 and headroom -0.4582, the binding fault being $\rho = 0.3824$ rather than difference.

Interpretive care is required here and the paper states it. The substantive premise — that trust reduces the loss rate rather than adding to the inputs — is an assumption of the framework's parameterisation, not a finding of this paper, and the trust literature it leans on is about reciprocity, monitoring and institutional form rather than about decay rates [Gambetta, 1988, Ostrom, 1990]. **If trust instead raises η, the arithmetic above does not apply.** Which of the two it does is an empirical question this paper does not settle.

7.2 The ceiling

THEOREM 7.2 — THE CEILING. Since $F \leq \ln N$ [Curzi, 2026], the condition $F > c/\eta$ is satisfiable only if

$$\frac{c}{\eta} \leq \ln N \qquad \text{equivalently} \qquad f_{\min} \leq kN^2$$

Proof: $F > c/\eta \geq \ln N \geq F$ is a contradiction. **QED.**

If $c/\eta > \ln N$, no arrangement of those N categories is sustainable. Not a better pairing, not a better schedule, not a better conversation. For $N = 12$: $\ln 12 = 2.484907$, $kN^2 = 69.61330555033166$ Hz, and $2.885390 \times \ln 12 = 7.169925$ octaves above k. **Above that, the relation cannot be tuned; it must be restructured.**

The corollary is a minimum category count, $N \geq e^{c/\eta}$:

```
c/eta = 0.5000  ->  N >=  1.6487  ->  N_min =  2
c/eta = 1.0000  ->  N >=  2.7183  ->  N_min =  3
c/eta = 1.5000  ->  N >=  4.4817  ->  N_min =  5
c/eta = 2.0000  ->  N >=  7.3891  ->  N_min =  8
c/eta = 2.4849  ->  N >= 11.9999  ->  N_min = 12
c/eta = 2.5000  ->  N >= 12.1825  ->  N_min = 13
c/eta = 3.0000  ->  N >= 20.0855  ->  N_min = 21
c/eta = 4.0000  ->  N >= 54.5982  ->  N_min = 55
```

The growth is exponential, and that is the whole content of the corollary. Going from $c/\eta = 3$ to $c/\eta = 4$ — a 33 per cent increase in relative cost — raises the minimum viable category count from 21 to 55. **Past a point, cost does not make a relation harder; it makes the relation's category set inadequate**, and category sets are expensive to enlarge honestly. A ceiling imposed by one parameter alone, independent of how the remainder is arranged, is a recurring shape in this corpus: the concentration atlas finds validator pools at HHI 0.1900 and 5.26 effective pools against a ceiling of $1/s^2_{\max} = 11.11$ that the leading share alone imposes [Curzi, 2026].

The ceiling test is the paper's cheapest and most consequential instrument because it changes the *kind* of answer. Below it a case is tunable and some assignment of u closes it; above it, no assignment does, and a prescription issued without the test may be unsatisfiable. **Unsatisfiable prescriptions are worse than no prescription, because they consume the effort that restructuring needed.**

8 The ordinal defeat, and the rule it motivates

8.1 The join is an error amplifier by construction

$$\frac{d \ln f_{\min}}{d(c/\eta)} = 2 \qquad \text{exactly}$$

Every unit of error in the cost-to-conversion ratio is two nats of error in the required frequency. An error of 0.1 in c/η is a 22.14 per cent error in $f_{\min}$; an error of 0.5 is a factor of 7.3891, which is 2.8854 octaves. Computed at five operating points:

c/η	implied interval in $f_{\min}$ (Hz)	width	width in octaves
1.0 ± 0.2	$[2.3944,\ 5.3289]$	$2.2255\times$	1.1542
2.0 ± 0.1	$[21.6097,\ 32.2379]$	$1.4918\times$	0.5771
2.0 ± 0.2	$[17.6925,\ 39.3754]$	$2.2255\times$	1.1542
2.0 ± 0.5	$[9.7099,\ 71.7467]$	$7.3891\times$	2.8854
2.4849 ± 0.2	$[46.6626,\ 103.8495]$	$2.2255\times$	1.1542

Since $D = 2.885390(c/\eta - F)$ and both terms carry error, **the deficit's uncertainty is** 2.885390 **times the combined uncertainty in nats: one nat of parameter uncertainty anywhere is** 2.885390 **octaves of uncertainty in the verdict.**

8.2 The relabelling test

Suppose c/η is known only ordinally. That is the honest condition for any parameter estimated from an unanchored rating instrument: one knows that site B costs more per unit of conversion than site A, and not by how much. **An ordinal scale is preserved by any strictly increasing relabelling.** Take three systems ranked $c/\eta = (1.0, 2.0, 3.0)$ and relabel them $(0.5, 2.0, 5.0)$. Every ordering is preserved and the ordinal content is identical. The frequencies are not:

```
labels (1.0, 2.0, 3.0)  ->  f_min = (    3.5721,  26.3942,    195.0279) Hz
labels (0.5, 2.0, 5.0)  ->  f_min = (    1.3141,  26.3942,  10648.1604) Hz
```

The third system's minimum sustainable frequency moves by a factor of 54.5982 **—** 5.7708 **octaves — under a relabelling that changes nothing an ordinal instrument can see.** This is not a large error bar. **It is the absence of a number.** An ordinal input to $f_{\min}$ does not yield an imprecise frequency; it yields no frequency at all.

What survives is the ordering and the sign of the verdict. Both labellings agree that system 3 requires more pitch than system 2 and system 2 more than system 1, and in the worked case of Section 9 the verdict DECAYS survives the whole admissible interval. **A coarse instrument is enough to tell you that something is wrong. It is not enough to tell you what to do**, and Section 9.2 shows the prescription flipping in kind inside the analyst's own error bar.

8.3 The audit-precedence rule, and three questions it reduces to

P-e, restated: a result may not be published above the level of audit available to check its inputs. A result whose inputs are known only ordinally may be published only as an ordering; one whose inputs are known on a ratio scale may be published as a magnitude. **The badge on a result is not decoration; it is the statement of which operations on that result are licensed.**

The rule is only useful if it can be applied from outside, so it is reduced to three questions a reader can put to any compact, powerful-looking formula — this paper's included.

Question one: what scale are the inputs on, and who says so? If the answer is *ratio, and here is the instrument*, the magnitude is reportable. If the answer is *ordinal*, or if there is no answer, the ordering is reportable and the magnitude is not. **A formula that divides two quantities and cannot name the true zero of either is producing a rank dressed as a number.** The relabelling test of Section 8.2 takes one line and is decisive.

Question two: what has this method rejected? Not what it has explained — what it has thrown out. **A method with no rejection history has never been in a position to reject.** This corpus's answer is on the record: its own moment equation demoted on scale [Curzi, 2026m], whose Result Two fixes the contact point $D = P$, $r = 1$, $M = E$ with unit elasticity in all three terms; its own rate-factor correction reversed to a factor of exactly 8; a 64-codon correspondence refused; and the evidential reading of a gematria hash killed by a 99.995 per cent collision rate destroying 10.47 bits (finding CC-R12).

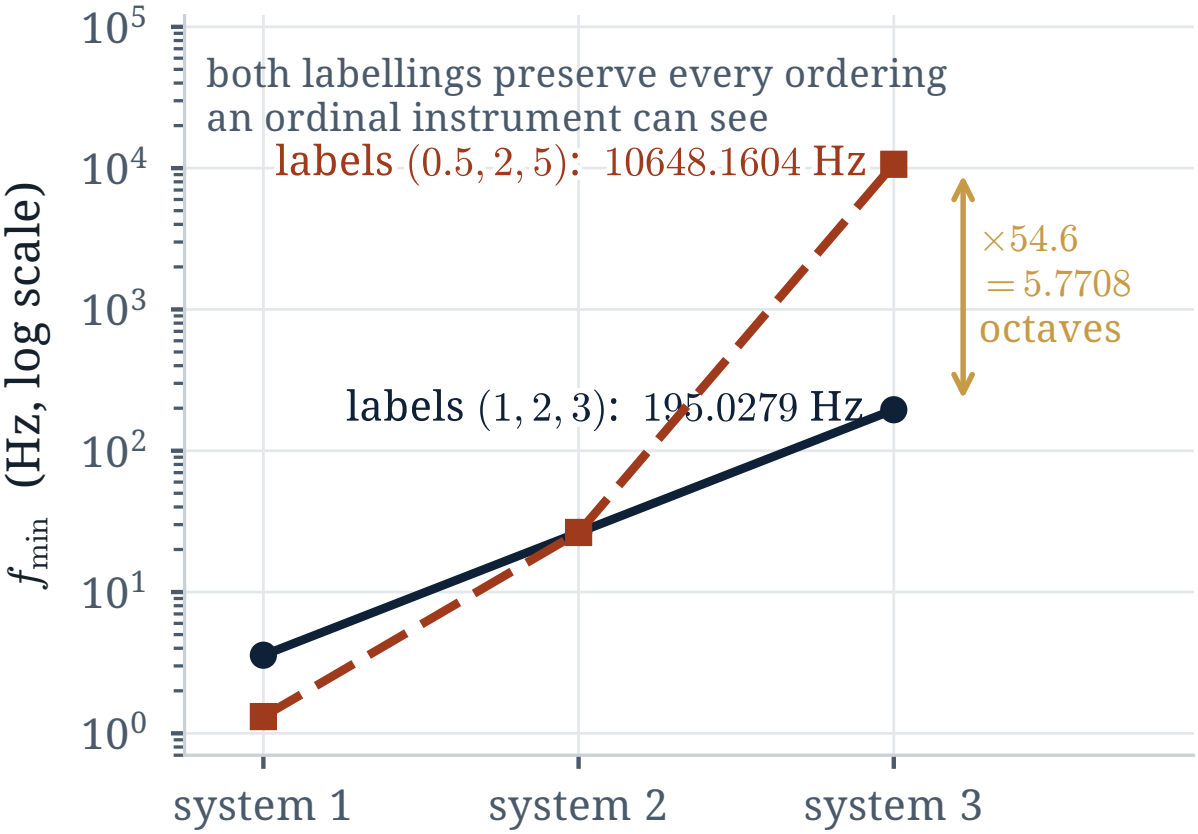


Figure 40.3: **An order-preserving relabelling of ordinally known parameters moves one system's threshold by** 5.7708 **octaves**, a factor of 54.6, while leaving every ordering an ordinal instrument could detect unchanged. Modelled from $f_{\min} = k\,e^{2c/\eta}$ with the **defined** constant $k = 0.483425732988414$; the two labellings are the source's. Built by `figs/src/CCP-040_fig03.py`.

Question three: how fast does the formula amplify error? Take the logarithmic derivative with respect to each input. **For $f_{\min}$ the answer is 2, and a formula with that property must ship with its sensitivity and not only its point estimate.** A formula that amplifies and does not say so is inviting confident wrong answers.

A work that can survive its own three questions is not thereby correct. It is checkable, which is the most any work can offer. This is a research-programme position rather than a naive falsificationist one [Lakatos, 1970, Popper, 1959], and the corpus's own honesty appendix is the instrument that enforces it [Curzi, 2026], whose counter-result assigns the corpus's central identification to the postulate drawer and states that one of its three postulates is a wager whose misfire column exists because wagers can lose.

8.4 Three corrections the audit produced

The claim is not that caution is a virtue. **The claim is that the audit found specific errors, in numbers this paper uses, that would have travelled into the keystone had the keystone come first.**

Correction one — the resolution cliff was in the wrong place, and the author's own repair was wrong and had to be reversed (CC-R21). Re-derived with exact rational arithmetic, the minimum folded gap is $9/8$ for $N = 3 \dots 6$, $50/49$ for $N = 7 \dots 16$, $289/288$ for $N \geq 17$ and $1682/1681$ for $N \geq 41$. **The published cliff had been at $N = 20$; it is at $N = 17$,** because $17^2 = 289$ folds beside $288 = 2 \cdot 12^2$ the instant index 17 enters the set. **And the rate penalty at the cliff is exactly 8, not the $288/49 = 5.878$ that had been asserted**, because the two binding pairs sit at different octave heights.

Correction two — the source table of the essence law contains data-entry errors, and one is provably impossible (CC-R5). $f = kn^2$ reproduces all 118 published element values to a worst relative deviation of 9.0×10^{-15}. The compound rule $f = k^2 G^4$ matches 37 of 44 published composites to better than one part in 10^{12}. **The seven that disagree are transcription faults, not evidence of a different law**, and the pure-diatomic row proves it: for a single-element species G is the element index whatever the atom count, so no atom count can produce the published figure. Two further faults travel with the citation: a prose instruction `1+5^.5/2` missing its parentheses, and φ stored to fifteen significant figures. **The standing rule that follows: do not copy published values, recompute them.** The audit is owned in this series by the operations paper [Curzi, 2026], whose Result Three derives the compound rule rather than asserting it.

Correction three — the intuitive repair makes it worse, and by how much (CC-R15). This is Section 7.1, and it is the correction most likely to be needed by a reader who has just finished reading the threshold result.

Each of these is a number this paper uses. The first governs how any of these frequencies can be rendered at all; the second governs the constant at the centre of every calculation; the third governs the single most common intervention a reader will reach for. **A keystone published before the audit would have carried all three forward into a diagnostic people were meant to act on.**

8.5 What the withholding cost

The corpus does not claim the delay was free. Readers who could have used an octave deficit to see that a bond was structurally short — not badly intentioned, not insufficiently loving, structurally short by a computable number of octaves — did not have it. **That is a real cost and it is not erased by the fact that the delay was correct.** Nor does the corpus claim the omission was foreseen in every detail: **some of it was design and some of it was discovery**, and several of the findings above arrived late enough that the join could not honestly have been made earlier. A narrative in which nothing was ever discovered, only unveiled, destroys the reader's ability to calibrate how much of the work was checked and how much was planned.

9 Worked example: the carried dyad

This case is CONSTRUCTED. Its parameters are chosen, not observed. No claim is made that any real pair of people has these values, nothing here is an assessment of anyone, and no clinical, diagnostic or prognostic claim is made or implied.

Two parties, A and B, each represented as a unit vector over the twelve-state ladder, components being the intensity with which a state is held, normalised to length one:

```
A = {4: 0.7, 6: 0.5, 9: 0.4}
B = {4: 0.6, 6: 0.55, 9: 0.3, 10: 0.2}

eta = 0.30    c = 0.62    delta = 0.09        (one declared period, N = 12)
```

9.1 Phase 2a — as found

Direction one, from the geometry to the pitch:

```
cos theta = 0.965021
u         = 1 − cos^2 theta        = 0.068735
F         = ln(1 + 11u)            = 0.563085 nats
N_eff     = 1 + 11u                = 1.7561
f         = k (1 + 11u)^2          = 1.4908 Hz     folded 381.644 Hz (j = 8)
octaves present = log2(f/k)        = 1.624719
```

Two parties, and an effective count of 1.7561**.** They are 1.76 parties' worth of distinct perspective; their represented states point 0.965 of the way in the same direction.

Direction two, from the operating parameters to the threshold:

```
c/eta            = 0.62 / 0.30            = 2.066667
ceiling check:     2.066667 <= ln 12 = 2.484907    BELOW -- a tuning problem, not structural
f_min            = k e^(2 ∗ 2.066667)     = 30.1588 Hz   folded 241.270 Hz (j = 3)
octaves required = 2.885390 ∗ 2.066667    = 5.963140
```

The join, and two independent cross-checks:

```
f = 1.4908 Hz  <  f_min = 30.1588 Hz                          =>  DECAYS

octave deficit D = 5.963140 − 1.624719   = +4.338420 octaves ( = +5206.1 cents )
ratio f_min/f    = 20.229940              = 4.338420 octaves  same number, stated twice
Q_inf = (0.30 ∗ 0.563085 − 0.62) / 0.09  = −5.0119           NEGATIVE
−2.885390 ∗ (0.09/0.30) ∗ (−5.011940)    = +4.338420         identity holds, err 8.88e−16
```

Three routes, one number. And the honest note that must travel with the negative fixed point: admissibility fails at $t = 1$ (Section 6.8). **The pair does not arrive at** -5.0119**. They arrive, one period in, at the point where something has to give.**

What the reading says: this bond is short by 4.34 octaves; the shortfall is a property of the geometry and the cost, not of anyone's intentions; and the case is *tunable*, because $c/\eta = 2.066667$ is below $\ln 12 = 2.484907$. **What it does not say:** that either party is deficient, that the bond should end, that anyone is at fault, or that anything about either party has been assessed. **The reading is of a configuration of four parameters. It is not a reading of a person.** The corpus's founding definition makes the same point structurally: the present state is a local minimum of the comfort landscape by definition rather than by luck [Curzi, 2026].

A second discrepancy found in the source and left standing. The source reports this case's deficit as $D = +4.338420$ from a reported $u = 0.068735$. Recomputing D from the *rounded* u gives $+4.338414$; recomputing from the underlying state vectors gives $+4.338420197$, which is the figure the source quotes. **The two disagree in the sixth decimal place because the printed u is rounded, not because either is wrong.** This paper quotes the vector-derived value in the worked example and the declared-parameter value in the twelve-set table of Section 6.1, where u is an input rather than a derived quantity. The same applies at phase 2b, where the vectors give $u = 0.69237163$ and $D = -3.590984$ while the table's declared $u = 0.692373$ gives $D = -3.590989$.

9.2 Reading the same case without a scale audit

Suppose the same analyst has the same three-line formula and no scale audit; η and c were estimated from a rating instrument that produces ranks, and the honest uncertainty on c/η is about ± 0.5 — which, for an ordinal instrument, is generous, since strictly the uncertainty is unbounded.

```
c/eta       f_min (Hz)   deficit    below ceiling?   prescription
1.566667      11.0948    +2.8957        yes          raise u to >= 0.344605
2.066667      30.1588    +4.3384        yes          raise u to >= 0.627132
2.566667      81.9801    +5.7811        NO           no arrangement of 12 categories works;
                                                      N must rise to at least 14
```

First: the verdict survives. At every point in the interval the bond decays, and $f = 1.4908$ Hz sits below the smallest $f_{\min}$ in the range by a wide margin.

Second: the prescription does not survive. At the low end the repair is a five-fold increase in represented difference; at the centre a nine-fold increase; at the high end **there is no value of u that works at all**, because $c/\eta = 2.566667$ exceeds $\ln 12$. **The prescription flips from *tune this* to *restructure this*, and the flip happens inside the analyst's own error bar.**

Third: the two prescriptions differ in kind, not degree. *Increase the difference between you* and *twelve categories are not enough to describe what you are doing, find at least fourteen* are not two settings of one dial. They are different projects, with different costs, aimed at different objects. **An instrument that cannot tell them apart is not a weak instrument; it is a misleading one, because it returns a confident number either way.**

9.3 The intervention the unaudited reading would have recommended

The analyst, seeing a decaying bond, proposes building trust, which in this framework acts on δ. Suppose the intervention succeeds completely and moves δ from 0.09 to 0.05:

```
before:  Q_inf = (0.30 * 0.563085 - 0.62) / 0.09 = -5.0119
after:   Q_inf = (0.30 * 0.563085 - 0.62) / 0.05 = -9.0215
```

The successful intervention made the destination 80 per cent worse. The numerator $\eta F - c = -0.451075$ is negative; dividing a negative number by a smaller positive denominator makes it more negative. **Trust did not add order. It amplified the sign the system already had, and the system's sign was negative.**

And the octave deficit did not move at all. It was $+4.338420$ before and $+4.338420$ after, because D contains no δ term whatsoever. **An instrument that reads unchanged while a well-intentioned intervention doubles the rate of loss is not a broken instrument. It is the instrument telling you that you have not touched the thing that is wrong.**

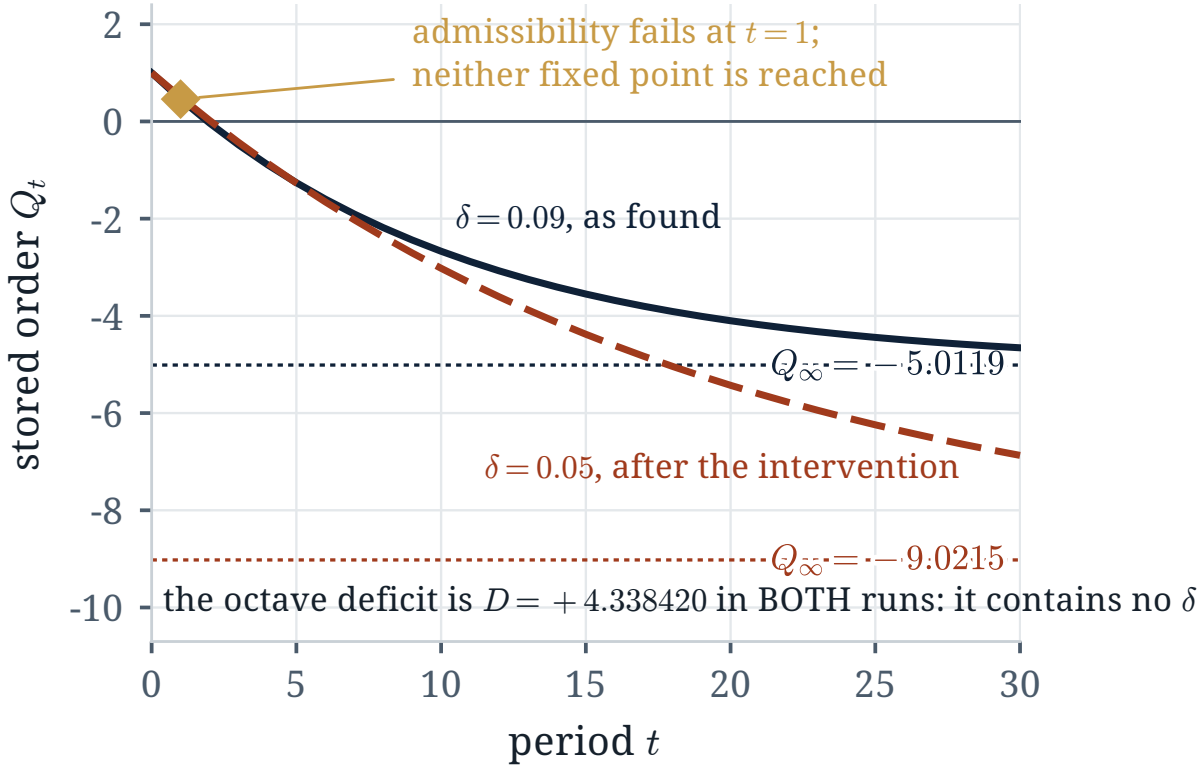


Figure 40.4: **A fully successful decay intervention worsens the destination by 80 per cent and leaves the diagnostic unmoved**: the octave deficit reads $+4.338420$ in both runs because it contains no δ, while the fixed point falls from -5.0119 to -9.0215. Admissibility fails at $t = 1$, so neither fixed point is reached and both are directions and rates rather than destinations. Modelled from $Q_{t+1} = (1 - \delta)Q_t + \eta F - c$ on the constructed carried dyad with $\eta = 0.30$, $c = 0.62$, $F = 0.563085$ nats and $Q_0 = 1.0$. Built by `figs/src/CCP-040_fig04.py`.

9.4 Phase 2b — after one intervention on the numerator

CONSTRUCTED. B takes on two states A does not hold, and the pair adopts a decision log. The first change raises u; the second raises η (interaction now sticks) and lowers c (coordination is cheaper when the record does the remembering). One realising vector, given for concreteness:

```
B' = {2: 0.95, 4: 0.60, 6: 0.55, 9: 0.30, 10: 0.20, 12: 0.8391}

cos theta = 0.554643    u = 0.692372    F = 2.153631 nats    N_eff = 8.6161
f = k * 8.6161^2 = 35.8881 Hz    folded 287.104 Hz (j = 3)   octaves present = 6.214066

eta = 0.44    c = 0.40    delta = 0.09
c/eta = 0.909091     ceiling check: 0.909091 <= 2.484907   BELOW, comfortably
f_min = k e^(2 * 0.909091) = 2.9782 Hz    folded 381.212 Hz (j = 7)
octaves required = 2.623082

D = 2.623082 - 6.214066 = -3.590984 octaves  ( = -4309.2 cents )       ACCUMULATES
```

Cross-checks:

```
Q_inf = (0.44 * 2.153631 - 0.40)/0.09 = +6.0844
-2.885390 * (0.09/0.44) * 6.0844      = -3.590984          identity holds
trajectory from Q_0 = 1.0:   t=1 +1.4576   t=5 +2.9116   t=20 +5.3134   t=40 +5.9675
admissibility holds at every step, since the drift is positive
convergence:  |Q_20 - Q_inf|               = 0.7710275635960624
              (1-0.09)^20 * |Q_0 - Q_inf|   = 0.7710275635960623
```

The bond has moved from 4.34 octaves short to 3.59 octaves in surplus — a swing of 7.93 octaves. In the fixed-point register the same swing is -5.0119 to $+6.0844$.

One structural note the join makes visible and neither half did. At phase 2b, $F = 2.153631$ against a ceiling of $\ln 12 = 2.484907$, so the pitch headroom is $2.885390 \times (2.484907 - 2.153631) = 0.9559$ octaves. **The repair worked, and it worked by spending most of the range that twelve categories make available.** At 2a there were 5.5452 octaves of headroom; at 2b there are 0.9559. **Further gains of the same kind are not available from this category set at all** — they would require raising N, which means finding genuinely distinct categories rather than finer gradations of the twelve.

Neither half of the equation says this on its own. The bridge gives a pitch with no threshold; the storage law gives a verdict with no ceiling in view. Joined, the reading places the relation on a scale with a top, and says how much of the scale is left.

9.5 The two routes checked against each other

Section 6.3 claimed two independent routes to the verdict. The carried case at phase 2b lets both be run.

Route A, through the geometry. Elicit the vectors, get $u = 0.692372$, get $F = 2.153631$, get $f = 35.8881$ Hz; measure $\eta = 0.44$ and $c = 0.40$, get $f_{\min} = 2.9782$ Hz; the deficit is -3.590984.

Route B, through the ledger. Observe the stock over periods and fit. From $Q_0 = 1.0$ the observed series is $1.0000,\ 1.4576,\ 1.8740,\ 2.2530,\ \ldots$, converging on $+6.0844$ with a transient half-life of 7.3496 periods. Fitting $\delta = 0.09$ from the decay of the transient and Q_∞ from the asymptote, the deficit follows as $-2.885390 \times (0.09/0.44) \times 6.0844 = -3.590984$.

The two routes share no measurement and agree to 10^{-15}. Route A used no history at all; Route B used no elicitation at all. **Where both are available, the framework can be checked against itself, and disagreement localises the fault.**

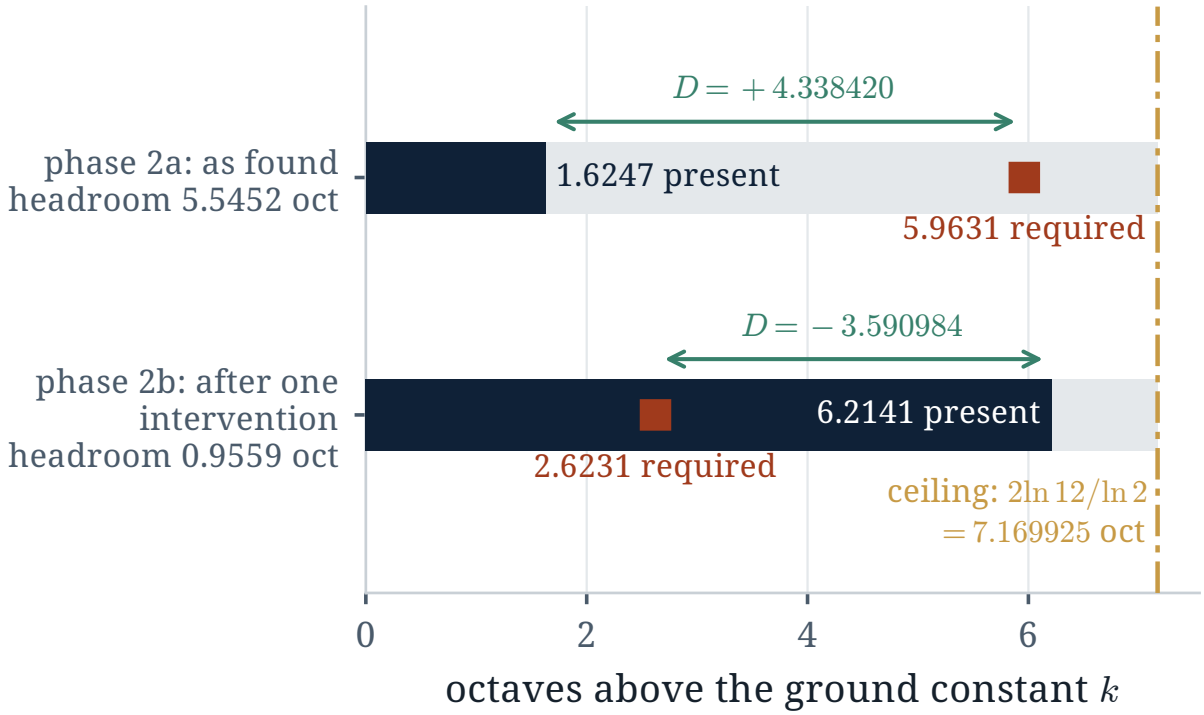


Figure 40.5: **One intervention converts a $+4.338420$-octave deficit into a -3.590984-octave surplus, and spends most of the range twelve categories allow**: pitch headroom falls from 5.5452 to 0.9559 octaves against the ceiling $2\ln 12/\ln 2 = 7.169925$. Modelled from the constructed carried dyad at its two phases with the **defined** constant $k = 0.483425732988414$ and $N = 12$. Built by `figs/src/CCP-040_fig05.py`.

And the asymmetry predicted in Section 6.3 is visible: Route B cannot tell you what to do. From the ledger alone, $\eta F - c = +0.547598$ is one number, and there is no way to know which decomposition produced it. Two admissible decompositions, both consistent with every observation Route B can make:

```
eta = 0.44,  F = 2.153631  (u = 0.692372)   ->  f = 35.8881 Hz,  pitch headroom 0.9559 oct
eta = 0.50,  F = 1.895197  (u = 0.513987)   ->  f = 21.4031 Hz,  pitch headroom 1.7015 oct
```

Identical ledgers. Identical verdicts. Different amounts of room left, and therefore different next moves. The first pair has nearly exhausted what twelve categories can give them and must raise N to gain more; the second still has most of an octave and a half to buy in u. **Route B says the bond accumulates. Only Route A says how much scale is left, which is the fact that determines what to do next.** This is gate G5 failing informatively rather than silently, and it is why a reading that reports a verdict without its geometry is incomplete.

10 Results

The six results below are the units other papers in this corpus cite. A citation to this paper should name one of them. Each carries its status: **DERIVED** where it follows from the stated definitions by algebra, **COMPUTED** where it is the output of exact arithmetic on stated parameters, and **POSTULATED** where it is asserted by this framework.

R1 — the join. $f_{\min} = k\, e^{2c/\eta}$ with $k = 0.483425732988414$ Hz (defined constant, exact; the hertz postulated, §2.3 P-b), derived as an equivalence: $Q_\infty > 0 \iff \eta F(u^*) > c^* \iff f > f_{\min}$. *Status:* DERIVED from the storage-law fixed point [Curzi, 2026a] and the essence bridge [Curzi, 2026k], conditional on P-a, P-b and P-d. **Verified at twelve constructed parameter sets with sign agreement in all twelve**, including the exact-equality case at the ceiling.

R2 — the deficit identity. $D = 2.885390\,(c/\eta - F) = -2.885390\,(\delta/\eta)Q_\infty$. The octave deficit and the steady-state gauge are the same number in different units. *Status:* DERIVED, cross-checked at a worst residual of 8.88×10^{-16} over the twelve sets by two routes sharing no intermediate quantity — floating-point round-off, not measurement error. The constant $2/\ln 2 = 2.885390081777927$ is exact. **This is the result the criterion paper carries forward as its Result Three [Curzi, 2026].**

R3 — the two absences. $f_{\min}$ contains neither δ nor N: retention loss sets the rate of arrival but not the threshold, and the category count enters only through the admissibility ceiling $c/\eta \le \ln N$, equivalently $f_{\min} \le kN^2 = 69.6133$ Hz at $N = 12$. *Status:* DERIVED by inspection and from $F \le \ln N$ [Curzi, 2026].

R4 — the ordinal defeat. $d \ln f_{\min}/d(c/\eta) = 2$ exactly, so an error of 0.1 in c/η is a 22.14 per cent error in $f_{\min}$, and relabelling $(1.0, 2.0, 3.0)$ as $(0.5, 2.0, 5.0)$ — an order-preserving relabel that no ordinal instrument can detect — moves one system's $f_{\min}$ by a factor of 54.5982, which is 5.7708 octaves. *Status:* COMPUTED. **With c and η known only ordinally, $f_{\min}$ is not known as a number at all; only its ordering and the sign of the verdict survive.**

R5 — the worked dyad. On the constructed carried case, $\eta = 0.30$, $c = 0.62$, $\delta = 0.09$ gives $f = 1.4908$ Hz against $f_{\min} = 30.1588$ Hz, an octave deficit of $+4.338420$ (5206.1 cents) and $Q_\infty = -5.0119$, with admissibility failing at $t = 1$. **The sign of the verdict is robust across the whole admissible interval where the number is not**; the prescription flips from *tune* to *restructure* inside the analyst's own error bar. *Status:* COMPUTED on a constructed case.

R6 — the sign trap. With $\eta F - c$ negative, reducing δ from 0.09 to 0.05 moves Q_∞ from -5.0119 to -9.0215, so **a fully successful intervention on retention makes the destination 80 per cent worse**, while D stays at $+4.338420$ to every printed digit. *Status:* DERIVED by inspection; COMPUTED on a constructed case. **This is the paper's sharpest result and it is negative.**

CR — counter-result. *The paper's own keystone formula is demoted by the paper.* With c and η known only ordinally, $f_{\min}$ moves by 5.77 octaves under an order-preserving relabel and so is not known as a number at all; and the worked case shows a fully successful intervention on retention making the equilibrium 80 per cent worse because the numerator was negative.

To these the framework adds three of its own: **the constant k is chosen, the affine effective count is not a Hill number, and the interpretation of k as a frequency is postulated.** The corpus reaches structurally parallel self-demotions elsewhere — every defensible bloc assignment in its architecture paper returns a negative surplus and the paper says so in print [Curzi, 2026]; its currency work finds an operating leg that lowers the effective count it was proposed to raise [Curzi, 2026]; and its concentration atlas refuses the reading that would make it useful as a ranking [Curzi, 2026].

11 Discussion

11.1 What follows

The classification is worth more than any single magnitude. The ceiling test costs one logarithm and changes the *kind* of answer a reading gives. Below the ceiling a case is tunable; above it, no assignment of u closes it. Since the verdict survives ordinal inputs and the magnitude does not, **the instrument's defensible output is a classification and an ordering, not a frequency**, and that is a smaller claim than the notation suggests.

The two gauges disagree, and both are right. Headroom uses the ceiling $\ln N$ and asks what the best possible geometry could support; the deficit uses the achieved F and asks what this geometry is doing. **A system can have room to grow in principle and be shrinking in fact.** Computed on the carried case at $Q = 1$: $h = 0.30 \times 2.484907 - 0.09 - 0.62 = +0.0355$ while $Q_\infty = -5.0119$. Both are correct, and a reading must report both and say which question each answers.

The negative result generalises past this framework. The structure of D's independence from δ is that **a multiplicative parameter cannot change the sign of the quantity it multiplies.** Any model in which a rate of loss divides a net drive has the property. Analogues are easy to find and worth naming: redundancy can reduce safety rather than increase it [Sagan, 1993]; interactive complexity with tight coupling makes catastrophic failure a system property rather than an operator error [Perrow, 1984]; and organisational learning that homogenises the organisation destroys the variance that produced its advantage [March, 1991]. **In each case the intuitive intervention operates on the wrong term.**

11.2 What does not follow

A reading does not conclude that a relation should end. $D > 0$ says a configuration is losing order at a computable rate under stated parameters. It does not weigh what the relation is for, what else it produces, or what ending it would cost, and the formalism contains no term for any of those.

A reading is not about a person. u is a property of two representations jointly; c and η are properties of an arrangement. **There is no single-party term anywhere in the apparatus.**

A reading is not clinical. The twelve states are not diagnoses, the ladder is not an assessment instrument, and the deficit is not a prognosis.

The algebra is not evidence. The derivation is compact and contentless: four strictly increasing transformations applied to an inequality. **The empirical content is in η, c, u and N, and the algebra cannot supply any of it.**

11.3 Relation to the frequency register

The reason for expressing the criterion as a pitch is that **intervals compose and hertz do not.** Two relations with different costs are compared on the octave scale directly, and the exchange rate is a single exact constant. The apparatus for that — cents, octave equivalence, interval arithmetic — is standard once an index exists [Rossing et al., 2002, von Helmholtz, 1885, Benson, 2006]. **What the apparatus does not carry with it is any claim about hearing.** Consonance

depends on the spectrum of the sound and not on the ratio alone [Sethares, 2005]; interval categories and octave equivalence generalise less far across listeners and species than is usually assumed [Burns, 1999, Hoeschele et al., 2012]; and the octave's perceptual specialness, while real, has a measured and finite size [Demany et al., 2013]. **This paper uses the arithmetic and asserts none of the psychoacoustics.** Where the corpus's frequencies do meet real acoustics it is in a derived room [Curzi, 2026l], whose Result Two fixes $7.2400 \times 4.3798 \times 2.9319$ m, $V = 92.969$ m^3 and an RT60 of 0.5 s requiring 29.94 m^2 sabins — **and that is a statement about a room, not about a person.**

The consonance apparatus itself is a chosen function on rationals rather than a measured fact about hearing, and its two classical forms — a harmonic-distance measure [Tenney, 2015] and a *gradus suavitatis* [Euler, 1739] — are badged that way by the consonance paper in this series [Curzi, 2026], whose Result One is the doubling theorem $\mathrm{HD}_{\mathrm{ess}}(i, j) = 2\log_2(pq)$ verified exactly over all sixty-six index pairs at $N = 12$.

11.4 The reflexive exposure

The framework's foundations identify a self-fulfilling loop whose fixed point is a prophecy that makes itself true [Curzi, 2026]. **The exposure is direct: an instrument that issues a verdict of DECAYS on a relation may change the relation by issuing it.** The mechanism is documented [Merton, 1948, Rosenthal and Jacobson, 1968] and has been caught happening in the field, with dates and prices, as a theory became true by being used [MacKenzie and Millo, 2003]. But the effect sizes are contested [McCambridge et al., 2014] and, on the best long-run assessment, **self-fulfilling prophecies are real but small, do not typically accumulate, and often dissipate** — most expectations turn out to be accurate rather than causal [Jussim and Harber, 2005]. **The reflexive exposure is therefore real and probably smaller than the framework's rhetoric implies**, and the correct response is not to suppress the reading but to publish the parameters that produced it so a third party can recompute it. That is also why any indicator adopted as a target degrades as an indicator [Campbell, 1979, Goodhart, 1975].

11.5 What the surrounding programme allows this paper to claim

The title of the wider corpus places this work in a lineage that begins with *physique sociale* [Quetelet, 1835] and continues through the modern computational programme [Lazer et al., 2009, Pentland, 2014]. **The lineage's characteristic error is Quetelet's**: treating the mean of a distribution as a type rather than as a statistic. This paper inherits that risk directly, because it writes a single u, η and c for an arrangement, and it responds with a rule — no aggregation — rather than with an argument.

Three constraints from that programme bear on the central object of this paper and are not comfortable. First, human activity is bursty, with heavy-tailed waiting times and **no characteristic time scale** [Barabási, 2005]; a *minimum sustainable frequency* is a claim about a characteristic scale and can meet that finding only by insisting that its frequency is a position on a surplus scale and not a rate of any observed process. Strong spatial regularity coexisting with temporal burstiness [González et al., 2008] is the counterweight but not a rescue. Second, an explanation is supposed to name a mechanism rather than exhibit a fitted curve [Hedström and Ylikoski, 2010]; **this paper fits nothing, which sidesteps the objection at the cost of having no empirical content to fit.** Third, and most sharply: **what observation would distinguish this framework's laws from a neighbouring model that fits equally well?** [Sobkowicz, 2020, Durlauf, 2005] The honest answer is H3 — the ceiling is a strict universal and one counterexample kills it — and nothing else in this paper answers it.

What the paper is entitled to claim is therefore narrow. Prediction is one of many legitimate reasons to build a model, and explanation, bounding, and the disciplined exposure of one's own assumptions are others [Epstein, 2008]. The position this work is best placed to defend is the analytical one: mechanisms stated explicitly, parameters declared, and the arrangement within which the law is claimed named in advance [Keuschnigg et al., 2018]. The

directional-prediction critique [Meehl, 1990] is worth meeting explicitly: the corpus makes **point** predictions with constants rather than directional ones, which is the right response, and it makes them from postulates whose status is printed.

12 Verification

This section decomposes every result into atomic steps, each checkable on its own. The standard is Occam's: **a postulate that can be dropped without the result changing has been dropped.** Every step carries one of four statuses: **ARITHMETIC** (reproducible by a reader with a calculator from the inputs named), **DEFINITION** (true by stipulation; nothing to verify but consistency of use), **CITED** (an external source supports it — the key is named), or **POSTULATED** (asserted by this framework, with no external support).

12.1 The postulates

	postulate	status
P1	The Pythagorean whole tone is the frequency ratio $9{:}8$.	CITED [Barbour, 1951]
P2	φ is the positive root of $x^2 = x + 1$; hence $\varphi^2 = \varphi + 1$.	CITED [Euclid, 1956, Vajda, 1989]
P3	A dimensionless ratio of the form $(9/8)^2/\varphi^2$ **is a frequency in hertz**.	POSTULATED — no source
P4	The essence of a state with n distinguishable components has frequency kn^2.	POSTULATED — no source
P5	Stored order obeys $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$ with $\delta > 0$.	CITED as a form [Oppenheim et al., 1996, Elaydi, 2005]; the *interpretation* POSTULATED
P6	$N_{\text{eff}} = 1 + (N-1)u$ (the affine count axiom S3).	POSTULATED — design axiom [Curzi, 2026]; **not a Hill number**
P7	$F = \ln N_{\text{eff}}$, so $F \le \ln N$.	DEFINITION (given P6)
P8	c is expressed in nats of surplus per declared period, so c/η is dimensionless.	**CONVENTION — stated here, absent from the source**
P9	$n_{\text{eff}} := N_{\text{eff}}$: the bridge feeds an effective count to a law that takes an integer count.	POSTULATED — modelling decision

Dropped as unnecessary. No postulate about the *ordering* of the twelve states is needed for R1 through R6: the trichotomy, the identity, the two absences, the ceiling and the ordinal defeat all hold for any N and any $u \in [0, 1]$ whatever the states are called. No postulate about consonance, about octave folding, or about rendering is needed anywhere in this paper. **The time–frequency bound of Section 4.11 contributes no term to any result and is retained only as a constraint on what may be done with the answer.**

12.2 The compounding chain

step	consumes	produces	status
L1	P1	$(9/8)^2 = 81/64 = 1.265625$ exactly	ARITHMETIC
L2	P2	$\varphi^2 = \varphi + 1 = 2.618033988749895$	ARITHMETIC
L3	L1, L2	$k = 0.483425732988414$	ARITHMETIC (exact)
L4	L3, **P3**	k carries the unit Hz	**POSTULATED**
L5	L4, **P4**	$f(n) = kn^2$ Hz	**POSTULATED**
L6	P6, P7	$F = \ln(1 + (N-1)u)$, $F \le \ln N$	DEFINITION
L7	L5, L6, **P9**	$f = k\,e^{2F}$, the essence bridge	ARITHMETIC given L5 and P9
L8	P5	$Q_\infty = (\eta F - c)/\delta$ for $\delta > 0$	ARITHMETIC
L9	L8	$Q_\infty > 0 \iff \eta F - c > 0 \iff F > c/\eta$	ARITHMETIC ($\delta > 0$, $\eta > 0$)
L10	L7	$F = \frac{1}{2}\ln(f/k)$; the bridge is a strictly increasing bijection	ARITHMETIC
L11	L9, L10, **P8**	$f \gtreqless f_{\min} \iff Q_\infty \gtreqless 0$ — **R1**	ARITHMETIC given P8
L12	L11	$\log_2(f_{\min}/k) = (2/\ln 2)(c/\eta)$; the constant is 2.885390081777927, exact	ARITHMETIC
L13	L12, Def. 3	$D = 2.885390\,(c/\eta - F)$	ARITHMETIC
L14	L13, L8	$D = -2.885390\,(\delta/\eta)Q_\infty$ — **R2**, residual 8.88×10^{-16}	ARITHMETIC
L15	L13 by inspection	$\partial D/\partial\delta = 0$; N absent from $f_{\min}$ — **R3**	ARITHMETIC
L16	L6, L11	$c/\eta \le \ln N$, $f_{\min} \le kN^2 = 69.6133$ at $N = 12$ — part of **R3**	ARITHMETIC
L17	L11 solved for u	$u^* = (e^{c/\eta} - 1)/(N-1)$	ARITHMETIC
L18	L11 differentiated	$d\ln f_{\min}/d(c/\eta) = 2$; the relabel moves $f_{\min}$ by $54.5982\times$ — **R4**	ARITHMETIC
L19	L8, L14, constructed parameters	the carried dyad: $D = +4.338420$, $Q_\infty = -5.0119$ — **R5**	ARITHMETIC
L20	L8, L15, constructed parameters	$\delta : 0.09 \to 0.05$ gives $Q_\infty : -5.0119 \to -9.0215$, D unchanged — **R6**	ARITHMETIC

12.3 What this chain is and is not

Read the status column. Steps L1, L2, L3 and L8 through L20 are arithmetic and a reader can redo every one of them. **Steps L4 and L5 are postulated, and step L11 rests additionally on a convention (P8) that the source never states.** A chain containing postulated steps is not a proof of its conclusion, and this one is not offered as one.

What it is offered as is a separation. The storage algebra and the octave arithmetic are solid; the frequency interpretation sitting on top of them is not. Anyone who accepts P3, P4 and P9 gets R1 through R6 by arithmetic alone. **Anyone who rejects all three still gets R1 through R6 in *surplus* units**, because the deficit is a difference of logarithms in which k cancels: the sustainability verdict never depends on the postulated constant, only the hertz values do. **That separability is the strongest claim this paper makes about its own foundations.**

Two statuses that must not be misread. The exponent 2 at L5 is **DEFINED**, not estimated: no fit was performed and the standard procedure for testing a scaling claim [Clauset et al., 2009] was not run. The logarithm at L6 is **DEFINITION** given the affine axiom P6, and the warrant the argument needs for it is the multiplicative-dynamics one [Peters and Gell-Mann, 2016] rather than the psychophysical one.

The falsifiable content lives at L8 and L6, not at L4 and L5. H1 through H3 test the storage law's fixed point, its sign and its ceiling; none of them requires the essence bridge to be true. **This paper's empirical exposure is therefore smaller than its notation suggests, and its unsupported material is correspondingly more inert.**

13 Limitations

Every check the governing rigour standard specifies is listed here by name, whether or not this paper passes it.

L-1. The dimensional audit does not close. $f(n) = kn^2$ balances only if k carries $[T^{-1}]$, and k is computed as a ratio of two dimensionless quantities. **No operational procedure is supplied by which the constant acquires the dimension of frequency.** The hertz is attached by assertion (§2.3 P-b, §3.4). This is the paper's largest structural weakness and no citation in the bibliography repairs it.

L-2. The exponent of the join is dimensionless only by an unstated convention. c/η carries $[\mathcal{C}\,T^{-1}]$ divided by a pure number and cannot be an exponent without a convention; **the source never states the convention.** It is stated here (P8), and **every reading in this paper depends on it being applied consistently.** Its failure mode is measured at $90.1\times$ (finding CC-R16).

L-3. The essence law has no external support. A quadratic relation between a dimensionless count and a frequency has no established analogue in psychoacoustics, ecology or measurement theory. **The badge is the honest position and is not a defence.**

L-4. The constant is chosen. $k = (9/8)^2/\varphi^2$ combines two real objects into a product with no precedent. It is quoted to fifteen significant figures **because it is defined**, not because anything was measured to that precision. Golden-section claims have a deflationary literature that this paper cites rather than avoids [Markowsky, 1992, Fischler, 1981].

L-5. Octave reduction beyond pitch is unwarranted. Octave equivalence is grounded for pitch [Shepard, 1982]; nothing grounds folding a non-acoustic quantity into a 220–440 Hz band and reading meaning from where it lands. **The operation is well defined; its interpretation is not supported.**

L-6. $N_{\mathbf{eff}} = 1 + (N-1)u$ **is not a Hill number.** It is a linear interpolation between 1 and N and appears in no ecological, economic, political-science or information-theoretic literature. The Hill-number citations in this paper apply to $1/\mathrm{HHI}$ and $\exp(H)$ and to nothing else [Hill, 1973, Laakso and Taagepera, 1979].

L-7. The design axiom S3 is load-bearing and unargued. Everything in this paper uses N_{eff}, and N_{eff} is defined by an axiom that selects one curve out of a family agreeing at both endpoints and diverging by **33.62 per cent** at $u = 0.5$. **Nothing in this paper is safe from S3.**

L-8. Conjoint measurement was not tested. The ratio-scale treatment of η, c and δ is provisional. Michell's objection [Michell, 1997] stands unanswered for every quantity in this paper except the counts, and the available remedies [Luce and Tukey, 1964, Rasch, 1960] were not run. **Section 8 computes what the failure of this check is worth: 5.77 octaves.**

L-9. The scalar-reduction habit is a single defect with two faces, and the paper concedes both. Michell's objection and the appraisal literature's objection are the same objection in two domains. The affect literature describes emotional states as **patterns across simultaneously varying appraisal dimensions** [Smith and Ellsworth, 1985], as **processes in time whose state is a synchronisation** [Scherer, 2001], as individuated by **qualitative core relational theme** [Lazarus, 1991], and as located on **continuous bipolar axes** rather than counted [Russell, 1980]. **A count discards exactly the information those accounts carry.** The ladder's top rung is contradicted specifically: action readiness is held to be constitutive of emotion as such rather than a late addition [Frijda, 1986], and this paper narrows the rung to *bound commitment* in response — **a narrowing, not an answer.** Two rungs survive intact [Tracy and Robins, 2004, Premack and Woodruff, 1978]. **Every reading in this paper that passes through the ladder inherits this concession.**

L-10. Reproducibility (gate G7) has not been run. Whether the same operationalisation of u, η and c recovers the same parameters across analysts and across sites is unknown for these constructs. **No result in this paper may be reported as reproducible.**

L-11. Identifiability is partial (gate G5). From a fitted trajectory alone only $\eta F - c$ is recoverable. **Route B gives the verdict and not the prescription**, and Section 9.5 exhibits two decompositions with identical ledgers and different next moves.

L-12. The baseline suite is anti-conservative. The shuffled-input control measures 0.108 against a nominal 0.05 (CC-R13a), is blind to repackaging (CC-R4), and is blind to a shared slow driver, passing 98.35 per cent of the time (CC-R24). **A paper that recommends the test must publish the holes, and these are they.**

L-13. Every case is constructed. The dyad's precision is exact because its parameters are exact; **that precision does not transfer to a measured case.** Reporting $f_{\min} = 30.1588$ Hz from an η known to one significant figure would be false precision, and the rule is that measured or fitted quantities require an uncertainty. **None of this paper's numbers are measured.**

L-14. No model comparison was performed. No out-of-sample test, no matched null with the same number of free parameters, and no information criterion is reported, because nothing was fitted. The corpus's one confirmatory attempt failed its own baseline test [Curzi, 2026i].

L-15. No statistical claims are made and none are supported. There are no p-values, no effect sizes and no confidence intervals in this paper, because there is no sample.

L-16. The functional form is underdetermined. A power law with exponent below one [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the qualitative facts as well as the logarithm, and many mechanisms produce indistinguishable curves [Stumpf and Porter, 2012]. **Nothing in this paper discriminates among them.**

L-17. The exponent is defined, not estimated. The exponent 2 in $f(n) = kn^2$ was not fitted, was not tested against alternatives, and the standard procedure for a scaling claim [Clauset et al., 2009] was **not run because there is no data to run it on. This paper makes no empirical scaling claim.**

L-18. The reflexive specification is not identified. The endogenous social effect in a linear-in-means reflexive loop cannot be recovered from observational data [Manski, 1993, Shalizi and Thomas, 2011]. The remedy adopted is stipulation of the reference set [Brock and Durlauf, 2001], and the claim is accordingly **a modelling position rather than an identified empirical result.**

L-19. The framework exempts its own parameters from its own reflexivity. If subjects change when modelled, δ, η and u are not structural [Lucas, 1976, Hacking, 1995]. **This is not resolved.**

L-20. There is no conservation law. Q is written with the grammar of a conserved quantity

and no conservation law is supplied, which is the characteristic defect of the genre [Mirowski, 1989]. Validity is claimed only within a specified arrangement [Cartwright, 1999].

L-21. Ergodicity is assumed and not argued. Q accumulates over time for one relation from parameters that would naturally be estimated across relations; the time and ensemble averages of a multiplicative process are not the same [Peters, 2019].

L-22. u is symmetric and psychological similarity is not. Similarity judgments are asymmetric and violate the triangle inequality [Tversky, 1977]; this framework's u assumes metric behaviour.

L-23. Boundary behaviour was incompletely examined in the source. $\eta \to 0^+$ makes c/η diverge and therefore places **every** relation above **every** ceiling — a real constraint on the model's domain that **the source does not examine**; $\delta \to 1$ is likewise unexamined. Both are stated in Section 6.7 for the first time. $N \to \infty$ raises the ceiling without bound, which the coordination-cost literature says is not available in practice [Brooks, 1975].

L-24. The trust mechanism is an assumption, not a finding. Section 7.1's sign trap depends on trust acting through δ rather than through η. **The literature this paper cites for trust [Gambetta, 1988, Ostrom, 1990] is about reciprocity, monitoring and institutional form, not about decay rates**, and which channel trust operates through is not settled here.

L-25. The 2.885390 shorthand is a rounded constant. The exact value is $2/\ln 2 = 2.885390081777927$. Where the text writes 2.885390 inside a formula it is a shorthand for the exact constant; every computed figure uses the exact value.

L-26. A published reading can change what it reads. The self-fulfilling exposure of Section 11.4 is not controlled for and cannot be controlled for by an instrument acting alone.

L-27. Two errors found in the source and printed rather than corrected silently. The twelve-set table's ceiling row prints $c = 2.485$, which exceeds $\ln 12 = 2.484906649788$ and therefore does not produce the exact zeros the row reports; recomputing from the printed value gives $Q_\infty = -0.0013$ and $D = +0.000269$. And the carried dyad's printed $u = 0.068735$ is a rounding of 0.06873466526463, so recomputing D from the printed intermediate gives $+4.338414$ against the vector-derived $+4.338420197$. **Neither is an error in the theorem; both are roundings, and both will be found by any reader who recomputes.**

14 Conclusion

What is established. Given the storage law and the surplus functional, the sign of the steady state is *equivalent* to a threshold comparison in frequency units, $Q_\infty \gtreqless 0 \iff f \gtreqless f_{\min}$ with $f_{\min} = k\,e^{2c/\eta}$, verified at twelve constructed parameter sets with sign agreement in all twelve. The octave deficit and the steady state are algebraically the same quantity, $D = 2.885390(c/\eta - F) = -2.885390(\delta/\eta)Q_\infty$, cross-checked by two routes sharing no intermediate quantity to a worst residual of 8.88×10^{-16}. **The threshold contains neither the decay rate nor the category count:** decay sets the rate of arrival and not the bar, and the category count enters only through the ceiling $c/\eta \le \ln N$, equivalently $f_{\min} \le kN^2 = 69.6133$ Hz at $N = 12$. Consequently an intervention acting only on decay changes the deficit by **exactly zero** and multiplies the existing sign of the steady state: on the constructed carried dyad a fully successful intervention moving δ from 0.09 to 0.05 takes Q_∞ from -5.0119 to -9.0215 — **80 per cent worse** — while D holds at $+4.338420$ octaves to every printed digit. All of this is arithmetic on stated definitions, and none of it depends on the value of k.

What is demoted, by this paper, about this paper. Because $d\ln f_{\min}/d(c/\eta) = 2$ exactly, an order-preserving relabelling of ordinally known parameters — $(1.0, 2.0, 3.0)$ to $(0.5, 2.0, 5.0)$, invisible to any ordinal instrument — moves one system's threshold by a factor of 54.5982, which is 5.7708 octaves. **With c and η known only ordinally, $f_{\min}$ is not an imprecise number; it is not a number.** What survives is the ordering and the sign of the verdict, and the worked case shows the *prescription* flipping in kind — from *tune this* to *restructure this* — inside the analyst's own error bar.

What is proposed rather than shown. That a dimensionless ratio of a tuning interval to

the square of the golden section is a frequency; that the essence of an n-component state is kn^2; that an effective count may be substituted for an integer count in that law; and that c is measured in the same units as F so that c/η can be an exponent. **The first three are postulates of this framework with no external support and the fourth is a convention the source never states.** They are labelled at every appearance and they are separable from everything above: **a reader who rejects all four keeps the equivalence, the identity, the two absences, the ceiling and the sign trap, expressed in nats instead of octaves.**

What is not claimed at all. No person, bond, tissue, molecule, room or institution is claimed to resonate at any frequency named here. The twelve states are not diagnoses and no reading in this paper is clinical, diagnostic, therapeutic or prognostic.

15 Declarations

Data and code availability. No dataset was collected for this study and none exists. Every number in the paper is the output of exact arithmetic on parameters stated in the text; every case is constructed and is labelled as constructed at the point of use. The figures are generated by the Python scripts named in their captions, which take their inputs from the tables in the text and are available from the author on request. **No empirical measurement of any human being, any relation or any institution underlies any figure, table or number in this paper.** A frequency law reads as though something were measured; nothing was.

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the analysis, the figures and the manuscript.

Ethics. No human-subjects research was conducted and no personal data were processed. The constructed dyad is an invented parameter set and is not derived from, and does not describe, any person. This paper asserts several perceptual and structural claims — that octave equivalence is a usable structure beyond pitch, that interval arithmetic in cents is the correct apparatus, and that a psychological state can be indexed by a count of components — **without any perceptual or psychometric experiment behind them.** Those claims rest on the cited literature, are limited by that literature's own boundaries (§11.3), and where the literature is in tension with them that tension is printed (§4.13, L-9).

Preregistration. No test reported in this paper was preregistered, because no test was run. The one confirmatory fit attempted elsewhere in this corpus was preregistered and **failed its own second falsification rule** [Curzi, 2026i]; that outcome is reported here as it is reported there.

Clinical disclaimer. The twelve state classes are not diagnostic categories. Nothing in this paper is an assessment, a screening instrument, a treatment or advice about any person's health, and nothing in it may be presented as such.

Chapter 41

The Minimum Sustainable Frequency: A Pitch Criterion for Relational Sustainability and the Count-Based Essence Index That Feeds It

Abstract. We derive a threshold frequency below which a relation cannot pay for itself. Substituting the essence bridge $f = k\,e^{2F}$ into the storage fixed point $Q_\infty = (\eta F - c)/\delta$ yields $f_{\min} = k\,e^{2c/\eta}$ with $k = (9/8)^2/\varphi^2 = 0.483425732988414$ Hz — a *defined* constant, exact by construction, whose interpretation as a frequency is postulated and not derived — and a trichotomy: $f > f_{\min}$ accumulates, $f = f_{\min}$ holds, $f < f_{\min}$ decays. In logarithmic units the requirement is exactly $2/\ln 2 = 2.885390081777927$ octaves per unit of c/η, and the octave deficit satisfies the identity $D = 2.885390\,(c/\eta - F) = -2.885390\,(\delta/\eta)\,Q_\infty$, cross-checked to a residual of 8.88×10^{-16}. Because $F \le \ln N$ the criterion carries a ceiling, $c/\eta \le \ln N$, that is $f_{\min} \le kN^2 = 69.6133$ Hz at $N = 12$. The sharpest result is negative: D contains no δ, so reducing loss cannot move the deficit at all. On a constructed dyad reading $D = +4.338420$ octaves and $Q_\infty = -5.0119$, cutting δ from 0.09 to 0.05 leaves D unchanged to every printed digit while worsening Q_∞ to -9.0215. The state index is defined as a component count, $f(n) = kn^2$, which sidesteps the corpus's own scale demotion; but the essence law is an unsupported postulate, the twelve-component ladder is not shown to be the right twelve, and the elicited intensities are not invariant under monotone transformation, spanning 0.888 octaves across admissible rescalings.

Keywords: minimum sustainable frequency; octave deficit; essence law; effective count; storage fixed point; ceiling; component count; magnitude estimation

Series: C. Frequency, Essence, and Mind **Paper:** CCP-041

1 Introduction

A relation between two or more parties that costs something to maintain either accumulates durable order or loses it. The framework this paper belongs to models that accumulation with a first-order linear difference equation — a leaky store driven by a surplus term and drained by a cost term — whose fixed point is $Q_\infty = (\eta F - c)/\delta$. That form is not novel and

is not claimed to be: it is the discrete leaky integrator that appears as the RC low-pass filter in signal processing [Oppenheim et al., 1996], as exponential smoothing in forecasting [Holt, 1957], and as the integrate-and-fire neuron in computational neuroscience [Lapicque, 1907, Gerstner and Kistler, 2002]. Its stability condition and unique fixed point are standard [Elaydi, 2005], and the convergence is Banach's [Banach, 1922].

The question this paper asks is narrower and, as far as we can establish, unasked: can the sign of that fixed point be read off a single threshold, and does the threshold have a useful logarithmic form? The answer is yes, and the threshold is exponential in the ratio of cost to conversion efficiency. Writing the relation's position on the framework's logarithmic surplus scale as a frequency — the *essence bridge* fixed in the corpus's Series B [Curzi, 2026k], whose Result One is that $f(m)/f(n) = m^2/n^2$ so that the reference constant cancels from every interval — the condition becomes a comparison of two pitches, and the shortfall becomes an interval measured in octaves.

Four things follow, and the fourth is the one worth having. First, the threshold $f_{\min} = k\,e^{2c/\eta}$ exists and partitions behaviour into three regimes exactly. Second, in logarithmic units the exchange rate is a single exact constant, $2/\ln 2 = 2.885390081777927$ octaves per nat. Third, the shortfall — the *octave deficit* D — is algebraically identical to the steady state itself, so a pitch reading and a stock reading are one reading in two units. Fourth, and negatively: D **contains no decay term at all.** The intervention most people reach for when a relation is failing — reduce the loss, build trust, add ritual and continuity — cannot move the deficit by any amount whatsoever, and when the numerator is already negative it makes the destination strictly worse. That is a derived consequence of the algebra, it is counter-intuitive, and it is the result this paper most wants read.

The second half of the paper supplies the input the first half consumes. A frequency is computable only once a psychological category has an index, and the index is defined here as a **count of distinguishable components** rather than as a rating. The distinction is not cosmetic. Stevens' typology of scales [Stevens, 1946] makes the permissible arithmetic on a quantity a function of the transformation group that leaves its empirical content unchanged, and a count admits only the identity transformation. That is what licenses squaring it. Ratings do not have that property, and the corpus has already demoted one of its own keystone constructions on exactly this ground [Curzi, 2026m], whose Result Two fixes the $D = P$, $r = 1$, $M = E$ contact point together with the ceiling $0 \le M \le \ln N$ that makes the divergence finite.

What this paper does not claim, stated at the front rather than at the back. The constant $k = (9/8)^2/\varphi^2$ is a *defined* number, not a measured one; its two components are real and well documented — the Pythagorean whole tone [Barbour, 1951] and the golden section [Euclid, 1956] — but their product **expressed as a frequency in hertz has no precedent in any literature we have found**, and the standard critique of golden-section claims [Markowsky, 1992] is cited here deliberately and again wherever φ appears. The essence law $f(n) = kn^2$ is a postulate of this framework and not a finding. Octave reduction is licensed for pitch [Shepard, 1982] and its perceptual specialness has been measured [Demany et al., 2013]; **its extension to a quantity that is not a pitch is a modelling choice with no external warrant**, and Section 2 labels it as one. Section 9 decomposes every one of these into steps a reader can check individually, and Section 10 prints every check this paper fails.

Contributions

1. The minimum sustainable frequency $f_{\min} = k\,e^{2c/\eta}$ and the accumulate/hold/decay trichotomy, with the equivalence proved and verified at twelve parameter sets.
2. The octave form, with the exact conversion constant 2.885390081777927 octaves per nat and its cents and bits equivalents.
3. The deficit identity $D = 2.885390(c/\eta - F) = -2.885390(\delta/\eta)Q_\infty$, cross-checked to a residual of 8.88×10^{-16}.
4. The negative result: D is independent of δ, demonstrated on a constructed dyad where halving the loss leaves D unchanged to every printed digit and worsens Q_∞ by 80 per

cent.

5. The ceiling $c/\eta \leq \ln N$, equivalently $f_{\min} \leq kN^2$, which classifies a case as tunable or structural before any prescription is issued.
6. A count-based index for the psychological argument of the bridge, with the scale-theoretic argument for why a count escapes the demotion that a rating inherits — and an explicit statement of what that argument does not buy.
7. A measured failure of the second invariance: monotone rescalings of elicited intensities, admissible if the intensities are merely ordinal, move the derived frequency by 1065.05 cents.

2 Hypotheses, definitions, and postulates

The rigour standard governing this corpus requires that a claim which *could* be wrong be stated so that it could be, and that a claim which could not be wrong be labelled with which of the three non-empirical kinds it is. **A hypothesis nothing could falsify is not a hypothesis**, and this section does not manufacture one. Series C carries the corpus's least-supported material, so the labelling below is done first and in full.

2.1 Hypotheses proper

H1 — the sign hypothesis. A relation whose measured conversion-weighted surplus falls short of its measured cost, $\eta F < c$, loses stored order over successive periods; one for which $\eta F > c$ gains it. **Predicted:** the sign of the observed change in the stock over a run of periods agrees with the sign of $\eta F - c$ computed at the start of the run. **Falsified by:** a system with $\eta F < c$, all four parameters measured on one declared period, whose stock rises over a run long enough for the fixed point to dominate the transient. **Status:** derived from the storage law [Curzi, 2026a]; **not tested here.** No observational data were collected for this paper.

H2 — the ceiling hypothesis. No relation whose measured cost ratio exceeds the logarithm of its declared category count, $c/\eta > \ln N$, sustains itself under any assignment of non-parallelism over those categories. **Predicted:** among systems measured with a declared and fixed N, none with $c/\eta > \ln N$ is observed to accumulate. **Falsified by:** one counterexample. The hypothesis is a strict universal and is cheap to kill. **Status:** derived from the normalisation bound $F \leq \ln N$ [Curzi, 2026]; **not tested here.**

H3 — the decay-blindness hypothesis. An intervention that acts only on the decay rate δ leaves the octave deficit unchanged and multiplies the existing sign of the fixed point. **Predicted:** in a system with $\eta F < c$, an intervention that measurably lowers δ and leaves η, c, u and N unchanged is followed by a *more* negative steady state, not a less negative one. **Falsified by:** a δ-only intervention that reverses the sign of the stock's drift. **Status:** derived by inspection of D; demonstrated on a constructed case in §4.9 and in the worked example. **Not tested against observation.**

H4 — the count-reproducibility hypothesis. The component count n assigned to an episode under the declared decomposition rules D1–D8 is a property of the episode and not of the analyst. **Predicted:** independent analysts applying the published rules to the same episodes agree on n at a rate that does not fall with analyst pairing, and the residual disagreement carries no analyst main effect. **Falsified by:** a systematic analyst effect on n, or agreement on the derived *pitch class* no better than the ladder's marginal distribution. **Status:** this is gate FVP-7 of the framework's validation protocol. **It has not been run.** It is reported as not run in §6.14, in §10, and here.

H5 — the ratio-scale hypothesis. The intensities elicited by the procedure of §6.9 carry ratio-scale information, as magnitude estimation is held to [Stevens, 1957]. **Predicted:** cross-modality matching against the elicitation instrument returns a power relation with a stable exponent under changes of stimulus range. **Falsified by:** exponents that move with the experimenter's chosen range, which is exactly what the range effect predicts [Teghtsoonian, 1971], or by the response-bias mechanisms catalogued against magnitude estimation [Poulton, 1968]. **Status: not tested here.** The framework leans on the psychophysical literature for this and inherits its disputes [Krueger, 1989].

2.2 Definitions — true by stipulation, not falsifiable

- **Def. 1.** n := the number of distinguishable components required to specify a state.
- **Def. 2.** $u := 1 - (\hat{x} \cdot \hat{y})^2$ on unit intensity vectors; $F := \ln(1 + (N-1)u)$.
- **Def. 3.** $D := 2.885390\,(c/\eta) - \log_2(f/k)$, the octave deficit.
- **Def. 4.** $k := (9/8)^2/\varphi^2$. **This is a defined constant.** It is an exact ratio of exact quantities and is therefore computable to arbitrary precision; the fifteen significant figures quoted throughout, 0.483425732988414, are *exact digits of a definition* and carry no measurement uncertainty because there is no measurement. Every appearance of the constant at that precision in this paper is so labelled.

2.3 Postulates — asserted by this framework, with no external support

- **P-a. The essence law $f(n) = kn^2$.** A quadratic relation between a dimensionless count and a frequency. **No external literature supports it.** The corpus badges it a postulate about essence structure and this paper does the same. It is not derived here and nothing in the bibliography derives it. **The exponent 2 is defined by the model, not estimated from data.** No scaling exponent in this paper was obtained by fitting; none was tested against a log-normal or exponential alternative; and the goodness-of-fit and likelihood-ratio procedure that a claimed scaling relation would require [Clauset et al., 2009] was **not run, because there is no data to run it on.** A squared law presented as though it had been measured is exactly what that procedure exists to catch, and this one has not been measured.
- **P-b. The attachment of hertz to k.** Computing $(9/8)^2/\varphi^2 = 0.483425732988414$ is arithmetic. **Asserting that this dimensionless number *is a frequency* is a separate and much stronger claim**, and the source supplies no operational procedure by which the constant acquires the dimension of inverse time. This is postulated. Section 3.4 states it again in the dimensional audit, where it is unavoidable.
- **P-c. The affine effective count $N_{\text{eff}} = 1 + (N-1)u$.** A design axiom [Curzi, 2026], not a derived result and **not a Hill number of any order**. The Hill family [Hill, 1973] and its independent rediscovery in political science [Laakso and Taagepera, 1979] ground $1/\text{HHI}$ and $\exp(H)$ exactly — as the corpus's concentration work uses them [Curzi, 2026] — and ground this construction not at all. A reviewer reaching for the Hill-number literature as support will find it does not apply, and that is stated here so that the discovery is not a surprise.

2.4 Modelling choices — legitimate, but choices

- **M-a. Octave reduction applied to a non-acoustic quantity.** Octave equivalence is a perceptual structure formalised for **pitch** [Shepard, 1982], and the octave's perceptual specialness has a measured size [Demany et al., 2013]. Nothing licenses carrying it to a quantity that is not a pitch. The *operation* is well defined — division by a power of two — and its *interpretation* is a choice of this framework. Where the folded value is read as carrying meaning, that reading is unsupported.

- **M-b.** $N = 12$**.** Twelve component classes are used throughout. The search for a thirteenth failed, which is weak evidence and is reported as weak in §6.8. Twelve exceeds the classical span of immediate discrimination [Miller, 1956], which is a reason for caution and not a reason for twelve.
- **M-c. Consonance and interval apparatus.** Cents, octave equivalence and interval arithmetic are the correct apparatus once an index exists [Rossing et al., 2002], but what a listener finds consonant depends on the spectrum of the sound and not on the ratio alone [Sethares, 2005], and interval categories generalise less far than is usually assumed [Burns, 1999, Hoeschele et al., 2012]. This paper uses the arithmetic and claims none of the perception.

2.5 What is expository

Sections 4.10 through 4.21, 5.6 through 5.10, and 6.9 through 6.14 are expository: they set out procedure, reporting requirements and failure modes rather than argue a claim. Their organising question is *what must be published alongside a reading for the reading to be checkable by a third party*, and the answer is the specification blocks and the validation table.

3 Framework, dimensional audit, and method

3.1 The storage law — a shared mathematical form, not a derivation

$$Q_{t+1} = (1 - \delta)Q_t + \eta S_t - C_t$$

Analogy status: category 2, a shared mathematical form. The equation governing this framework's stored order is the same equation that governs a leaky integrator, and it is used because the mathematics is the same, not because the physics transfers. No claim is made that comfort is a charge, that η is a coupling constant, or that any physical result about RC circuits applies to a relation between people. The fixed point $Q_\infty = (\eta F - c)/\delta$ follows from setting $Q_{t+1} = Q_t$ and requires $\delta > 0$; the corresponding transfer-function reading, in which $1/\delta$ is the DC gain of $1/(z-a)$ evaluated at $z = 1$, is fixed in Series B [Curzi, 2026a]. Convergence is geometric with ratio $1 - \delta$ [Elaydi, 2005, Banach, 1922]; the corpus's foundations state the theorem exactly — a contraction of modulus $q < 1$ on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least q each pass [Curzi, 2026].

Three warnings attach to that form and are recorded here rather than in a footnote. First, the time average of a multiplicative process differs from its ensemble average, so a parameter estimated across individuals does not describe what happens to any one of them over time [Peters, 2019] — and Q accumulates *over time*, for *one* relation, from parameters that would most naturally be estimated from a cross-section. Second, a single Q, δ or u written for a group is an aggregation, and aggregation over heterogeneous agents does not in general yield a representative agent [Kirman, 1992]; the framework's own no-aggregation rule (SPEC 4.14) is the right response and is stated as a rule rather than defended. Third, Q **has the grammar of a conserved quantity — a stock that accumulates, leaks and is spent — without a conservation law to give that grammar content.** That is the characteristic failure of the genre: neoclassical economics was built by copying nineteenth-century energetics and made utility play the role of energy without the conservation law that gave energy its meaning [Mirowski, 1989]. This paper does not claim any conservation law, states no invariance under which Q would be conserved, and adopts instead the position that a law is claimed only within a specified arrangement rather than universally [Cartwright, 1999]. The wider programme this places the work in — statistical physics applied to social dynamics — has a standard survey that fixes what such models can and cannot be held to [Castellano et al., 2009], and the boundaries it draws are the boundaries claimed here.

The precedent that matters most for this paper is the leaky integrate-and-fire neuron [Lapicque, 1907, Gerstner and Kistler, 2002]: a leaky store with a drive and a threshold, yielding a rate below which the store does not fire. **That literature grounds the *shape* of** $f_{\min}$ **—**

a threshold rate emerging from a leaky store — and it grounds neither the formula nor the constant. The distinction is made here and repeated in §9 and §10.

3.2 The surplus functional and the effective count

$F(u) = \ln(1 + (N-1)u)$ is the corpus's surplus functional, characterised by angle-only dependence, a zero at parallel states, an affine effective count and a normalisation [Curzi, 2026]. Its logarithmic form places it in the same family as Shannon's entropy [Shannon, 1948], and uniqueness-under-axioms is a known genre of theorem in that literature; the framework's axiom set is its own and is not one of those. The ceiling $F \leq \ln N$ is the normalisation, and it is the same ceiling that appears in requisite-variety arguments [Ashby, 1956].

Which construct of difference u measures must be declared, because at least three incompatible constructs travel under the word *diversity* — separation, variety and disparity [Harrison and Klein, 2007]. Here u is a *disparity* measure over represented states, and it is symmetric in its arguments, which psychological similarity is known not to be [Tversky, 1977]. That is a limitation and it is printed in §10.

3.3 The essence bridge and the constant k

The bridge writes a surplus as a pitch: $f = k\, e^{2F}$, so that intervals of surplus become intervals of frequency and the constant cancels from every ratio [Curzi, 2026k]. The constant is defined as

$$k := \frac{(9/8)^2}{\varphi^2} = \frac{81/64}{\varphi + 1} = 0.483425732988414 \quad \text{(defined; exact)}$$

Both components are real, specific and documented. The ratio $9{:}8$ is the Pythagorean whole tone, a historically attested object of tuning theory [Barbour, 1951], and it sits in a lattice of small-integer ratios with its own long literature. The golden section is Euclid's extreme and mean ratio [Euclid, 1956], named in the Renaissance, and $\varphi^2 = \varphi + 1$ with $\varphi^m = F(m)\varphi + F(m-1)$ are proved identities [Vajda, 1989].

Their product, expressed as a frequency in hertz, has no precedent. There is no physical, psychoacoustic or musicological result that combines a tuning ratio with an inverse square of the golden section to obtain a ground frequency. The literature on golden-section claims is overwhelmingly a literature of *deflation*: the appearances of φ in the Parthenon, the Great Pyramid, the human body and preferred rectangles do not survive examination [Markowsky, 1992, Fischler, 1981], the psychological preference results are weaker and more equivocal than they are usually reported to be [Höge, 1997], and and even-handed histories of the constant are available.

This framework's position is the one those critiques leave open and no wider. φ is being *chosen* as a constant, not *found* in a psychological system. Nothing in this paper is evidence that the golden section is present in anything measured. The constant sets a reference pitch and is invisible to every verdict, because D is a difference of logarithms in which k appears in both terms and cancels.

3.4 Dimensional audit

symbol	quantity	dimension	unit	status
Q	stored order	$[\mathcal{C}]$	none supplied	state variable; **no operational unit**
S	surplus input per period	$[\mathcal{C}\,T^{-1}]$	per declared period	driving term

symbol	quantity	dimension	unit	status
C, c	cost per period	$[\mathcal{C}\,T^{-1}]$	per declared period	driving term
δ	decay rate	$[1]$	per period	dimensionless; $0 < \delta \leq 1$
η	conversion efficiency	$[1]$	—	dimensionless; assumed $0 < \eta \leq 1$
u	non-parallelism	$[1]$	—	dimensionless by construction, $u \in [0, 1]$
F	surplus	$[1]$	nats	dimensionless; a logarithm of a count
N	declared category count	$[1]$	—	absolute scale, a count
n	component index	$[1]$	—	absolute scale, a count
k	reference constant	$[T^{-1}]$ **asserted**	Hz **asserted**	**defined as a ratio (dimensionless); the unit is postulated**
f	essence frequency	$[T^{-1}]$	Hz	inherits its dimension entirely from k
D	octave deficit	$[1]$	octaves	a ratio of logarithms; k cancels

Does the equation balance? The storage law balances provided S and C carry the same dimension as Q per period, which the framework assumes and does not measure; Q has **no operational unit**, and this audit records that rather than inventing one. The essence law $f(n) = kn^2$ balances only if k carries $[T^{-1}]$, because n is a dimensionless count. **The source supplies no operational procedure by which k acquires that dimension.** The constant is computed as a ratio of two dimensionless quantities — a frequency ratio and a geometric ratio — and the hertz is attached by assertion. This is the single largest dimensional exposure in the paper and it is listed in §10 by name.

The ratio c/η is dimensionally $[\mathcal{C}\,T^{-1}]$ divided by a pure number and is therefore not dimensionless as the exponent of $e^{2c/\eta}$ requires. **The framework's convention is that c is expressed in the same units as F — nats of surplus per declared period — which makes c/η dimensionless.** That convention is a choice, it is stated here, and every reading in this paper depends on it being applied consistently.

3.5 Scale type, and what licenses the arithmetic

quantity	scale type claimed	what licenses it
n, N	absolute	cardinality: true zero, fixed unit, identity transformation group only

quantity	scale type claimed	what licenses it
u	ratio	constructed from an inner product of unit vectors; zero is parallel
elicited intensities	ratio, up to a free common multiplier	**magnitude estimation, assumed** [Stevens, 1957]; not tested here
η, c, δ	ratio, **assumed**	no instrument is supplied; ordinal inputs are excluded by SPEC 4.4
f, $f_{\min}$	ratio in hertz, **inherited**	inherits the status of n and of the postulated bridge
D	interval, in octaves	a difference of logarithms

The strongest available objection to this framework is Michell's: that psychology routinely assumes its attributes are quantitative without ever testing whether they are [Michell, 1997]. The objection applies here in full. The formal remedy exists — conjoint measurement supplies a testable condition for additive quantitative structure [Luce and Tukey, 1964, Krantz et al., 1971], and a Rasch model supplies specific objectivity, a measure that does not depend on which items were used [Rasch, 1960] — and **neither has been run for any quantity in this paper.** The count argument of §5.2 answers Michell for n, because a cardinality's structure is a consequence of its definition rather than an empirical hypothesis; it answers him for nothing else. The intensities, η, c and δ are treated as quantitative by stipulation, and that is stated rather than hidden.

There is a second instance of that same objection in this paper, in a different domain, and it is better to name the pattern once than to meet each instance as though it were a surprise. Section 5 reduces an emotional state to a single integer — the number of components its specification requires — orders states by that integer, and squares it. Michell's charge is that a quantitative structure is assumed where it was never tested. The charge against the ladder is the same charge one level down: a **scalar** is asserted where the affect literature describes a pattern across simultaneously varying dimensions [Smith and Ellsworth, 1985], a process unfolding in time [Scherer, 2001], or a qualitative relational theme [Lazarus, 1991]. **These are one objection in two places, not two objections.** What both name is a habit of this framework: where the source literature describes a phenomenon by several jointly varying features, the framework replaces it with one number and does arithmetic on the number. The count argument of §5.2 protects the *arithmetic* of that number; it does not protect the *reduction* that produced it, and this paper does not claim that it does. §5.7 states the affect version at full strength, and L-18 concedes the two together as one limitation.

3.6 Identification — the objection this framework cannot presently answer

The framework's founding construction is a reflexive loop, $T(x) = W(x, M(x))$: the group influences the member and the members constitute the group [Curzi, 2026]. **That is a linear-in-means reflexive specification, and the endogenous social effect in such a specification is not identified from observational data** [Manski, 1993]. It cannot be separated from contextual effects or from correlated unobservables, however much data is collected; the reinforcement is that influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011]. **No amount of fit, held out or otherwise, recovers the parameter.**

The remedy this paper adopts is stipulation, and it pays the price openly. The reference set is fixed in advance — the declared category set of size N and the declared decomposition rules — and every reading is reported together with that stipulation, which is what makes a discrete-choice treatment of social interaction defensible where an unrestricted one is not [Brock and Durlauf, 2001]. **This paper therefore presents the reflexive loop as a modelling**

position, not as an identified empirical result. It is not offered as something an observation could confirm.

The same objection applied to the framework's own parameters. If systems change their behaviour when modelled, then δ, η and u are the decision rules of people who have not read the storage law, and they are not structural in the sense a policy exercise requires [Lucas, 1976]. Classification changes the classified [Hacking, 1995], and economic theory has been shown to become self-validating by being taught and used [Ferraro et al., 2005, MacKenzie and Millo, 2003]. **The framework asserts reflexivity for its subject matter and cannot consistently exempt its own parameters from it.** This paper does not resolve that; it states it, and treats every parameter as conditional on the arrangement in which it was fixed.

3.7 Method

Every number in this paper is the output of exact arithmetic on stated parameters, computed in double precision. **No observational data were collected, no human subjects participated, and no empirical dataset underlies any figure or table.** Every case is constructed and is labelled as constructed at the point of use. Residuals quoted at the 10^{-16} level are floating-point round-off, not measurement error, and are reported to demonstrate that two algebraically identical expressions were in fact computed independently and agreed. The figures are produced by the scripts named in their captions.

4 The minimum sustainable frequency

This is the result the framework is built around, and it is set out to be cited rather than rederived. Twenty-six later analyses in the corpus use what is stated here, and cite rather than restate it. Everything the framework means by `f_min`, by the octave deficit, by the ceiling, and by the four levers is set out here completely, with every number computed and every boundary named.

The keystone, in one block:

```
f_min = k * e^( 2c / eta )                    k = 0.483425732988414 Hz

octaves required  = 2 (c/eta) / ln 2  =  2.885390081777927 * (c/eta)

octave deficit  D = 2.885390 * (c/eta)  -  log2( f / k )

    D <= 0    sustainable as it stands
    D >  0    short by that many octaves

CEILING:  c/eta <= ln N.   For N = 12: c/eta <= 2.484907, f_min <= 69.6133 Hz.
          Above it, NO arrangement of those categories is sustainable at all.

FOUR LEVERS:  raise u  ·  raise eta  ·  cut c  ·  raise N.   And delta is NOT one of them.
```

The Law

4.1 Definition **DEFINITION 4.1 — THE MINIMUM SUSTAINABLE FREQUENCY.** For a relation with cost `c` per period and conversion efficiency `eta > 0` measured on the same period,

```
f_min  :=  k * e^( 2c / eta )
```

`k = 0.483425732988414 Hz` is the ground essence, badged as a chosen reference constant (the essence-law specification of Series B [Curzi, 2026k]). `f_min` is an essence frequency. No person, bond, tissue, molecule, room or institution is claimed to resonate at it, or at any value in this section.

4.2 The three regimes **THEOREM 4.2.** Under the conditions of the storage law's fixed-point theorem [Curzi, 2026a] — `delta > 0`, `eta > 0`, constant `u` and `c` — and with `f = k e^(2F(u))` the relation's essence frequency:

```
f  >  f_min    <=>   Q_inf > 0     the relation ACCUMULATES order
f  =  f_min    <=>   Q_inf = 0     it HOLDS, and nothing more
f  <  f_min    <=>   Q_inf < 0     it DECAYS, whatever anyone intends
```

Proof: the equivalence chain of the sustaining-frequency paper [Curzi, 2026], whose Result Two is the deficit identity — Steps 0 to 4, a chain of equivalences under strictly increasing transformations. *Verified* at twelve parameter sets in the sustaining-frequency paper [Curzi, 2026] §3.7, sign-agreement in all twelve including the exact-equality case.

The middle line is the one readers skip and the one that costs them. `f = f_min` is not a success condition. It is the condition under which a relation pays exactly for itself and banks nothing. A relation at its minimum sustainable frequency is not sustainable in any colloquial sense; it is stationary, and it has no reserve against any increase in cost.

4.3 The octave form **THEOREM 4.3.** Taking the base-2 logarithm of Definition 4.1,

```
log2( f_min / k )  =  2 (c/eta) / ln 2  =  2.885390081777927 * (c/eta)
```

EVERY UNIT OF COST-TO-CONVERSION RATIO COSTS 2.885390 OCTAVES OF REQUIRED PITCH.

The constant is exact: `2/ln 2 = 2.885390081777927`. It is the same constant that converts one nat of surplus into pitch (the sustaining-frequency paper [Curzi, 2026]), and it appears here for exactly that reason — `c/eta` is measured in nats of surplus per period, and the conversion from nats to octaves is one number.

Equivalent statements, all exact:

```
1 nat of required surplus  =  2.885390 octaves  =  3462.4681 cents  =  1.442695 bits
1 octave of pitch          =  0.346574 nats     =  0.5 bits EXACTLY
```

One octave is exactly half a bit. That line prices every lever in §4.13. The nat, the bit and the conversion between them are the standard units of a logarithmic information measure [Shannon, 1948]; nothing about the framework's use of them is novel, and nothing about them licenses the frequency reading.

4.4 The table Computed at `k = 0.483425732988414 Hz`, `N = 12`, folded into 220–440 Hz with the transposition `j` given:

c/eta	f_min (Hz)	octaves above k	N_eff required	folded 220–440
0.00	0.4834	0.000000	1.0000	247.514 (j=9)
0.10	0.5905	0.288539	1.1052	302.314 (j=9)
0.20	0.7212	0.577078	1.2214	369.247 (j=9)
0.30	0.8809	0.865617	1.3499	225.500 (j=8)
0.40	1.0759	1.154156	1.4918	275.426 (j=8)
0.50	1.3141	1.442695	1.6487	336.406 (j=8)
0.60	1.6050	1.731234	1.8221	410.888 (j=8)

```
0.70          1.9604        2.019773          2.0138      250.930 (j=7)
0.80          2.3944        2.308312          2.2255      306.486 (j=7)
0.90          2.9246        2.596851          2.4596      374.343 (j=7)
1.00          3.5721        2.885390          2.7183      228.612 (j=6)
1.10          4.3629        3.173929          3.0042      279.227 (j=6)
1.20          5.3289        3.462468          3.3201      341.049 (j=6)
1.30          6.5087        3.751007          3.6693      416.558 (j=6)
1.40          7.9498        4.039546          4.0552      254.392 (j=5)
1.50          9.7099        4.328085          4.4817      310.716 (j=5)
1.60         11.8597        4.616624          4.9530      379.509 (j=5)
1.70         14.4854        4.905163          5.4739      231.767 (j=4)
1.80         17.6925        5.193702          6.0496      283.080 (j=4)
1.90         21.6097        5.482241          6.6859      345.755 (j=4)
2.00         26.3942        5.770780          7.3891      422.306 (j=4)
2.10         32.2379        6.059319          8.1662      257.903 (j=3)
2.20         39.3754        6.347858          9.0250      315.004 (j=3)
2.30         48.0933        6.636397          9.9742      384.746 (j=3)
2.40         58.7413        6.924936         11.0232      234.965 (j=2)
2.4849       69.6133        7.169925         12.0000      278.453 (j=2)   <-- THE CEILING, N=12
2.50         71.7467        7.213475         12.1825      286.987 (j=2)
2.60         87.6317        7.502014         13.4637      350.527 (j=2)
2.70        107.0336        7.790553         14.8797      428.134 (j=2)
2.80        130.7311        8.079092         16.4446      261.462 (j=1)
2.90        159.6753        8.367631         18.1741      319.351 (j=1)
3.00        195.0279        8.656170         20.0855      390.056 (j=1)
```

The fourth column is the most useful and the least obvious. `N_eff required = e^(c/eta)` is the number of genuinely distinct perspectives the relation must field just to break even. At `c/eta = 1.00` it is 2.7183 — a two-person bond must be worth more than 2.7 people's worth of distinct view. At `c/eta = 2.00` it is 7.3891. At `c/eta = 2.4849` it is exactly 12, which is all of them, and that is the ceiling arriving.

4.5 The octave deficit

DEFINITION 4.4 — THE OCTAVE DEFICIT. For a relation with measured essence frequency `f` and measured cost ratio `c/eta`,

```
D  :=  2.885390 * (c/eta)  -  log2( f / k )

D <= 0    sustainable as it stands
D >  0    the relation is short by D octaves
```

This is the framework's core instrument. Every practical section must be able to produce it, and every reading in this framework is reported as a `D` with its sign, its magnitude, and the ceiling test that qualifies it.

THEOREM 4.5 — THE DEFICIT IDENTITY.

```
D  =  2.885390 * ( c/eta  -  F )  =  -2.885390 * (delta/eta) * Q_inf
```

Proof: the sustaining-frequency paper [Curzi, 2026], whose Result Two is this identity. *Verified* at twelve parameter sets to a worst error of `8.88e-16`.

Three readings of the same identity, and all three matter:

(i) `D = 2.885390 (c/eta - F)`. The deficit is the surplus shortfall, converted. A relation short by one nat of surplus is short by 2.885390 octaves. There is no other exchange rate.

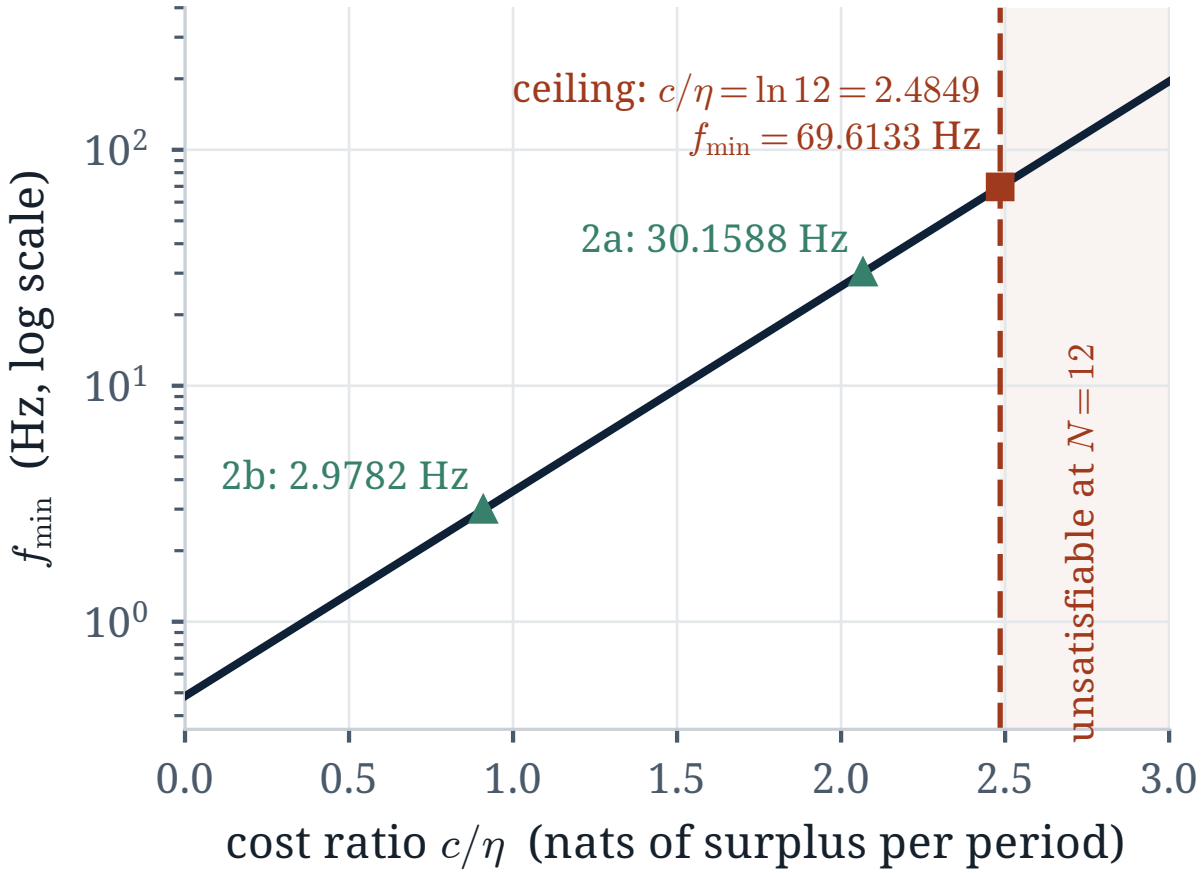


Figure 41.1: The threshold rises exponentially in the cost ratio and meets a hard bound at $c/\eta = \ln 12 = 2.4849$, where $f_{\min} = kN^2 = 69.6133$ Hz; above that bound no assignment of non-parallelism over twelve categories sustains the relation at all. Modelled from $f_{\min} = k\,e^{2c/\eta}$ with the **defined** constant $k = 0.483425732988414$ and $N = 12$; the two marked points are the constructed carried dyad of the worked example in Section 4 at its two phases. Built by `figs/src/CCP-041_fig02.py`.

(ii) `D = −2.885390 (delta/eta) Q_inf`. The deficit and the fixed point are the same quantity in different units. A pitch reading and a destiny reading are the same reading. The companion paper on destiny as a fixed point builds on this.

(iii) `k` cancels. `D` is a difference of logarithms in which the chosen constant appears in both terms. The verdict does not depend on `k` at all. Only the absolute hertz values do. This is the strongest available defence of the constant: it sets the reference pitch and is invisible to every verdict.

Deficit-to-fixed-point conversion, tabulated at `delta/eta = 0.30` (the carried case's ratio):

D (octaves)	f / f_min	Q_inf
−6.0	64.0000	+6.9315
−4.0	16.0000	+4.6210
−2.0	4.0000	+2.3105
−1.0	2.0000	+1.1552
−0.5	1.4142	+0.5776
0.0	1.0000	0.0000
+0.5	0.7071	−0.5776
+1.0	0.5000	−1.1552
+2.0	0.2500	−2.3105
+4.0	0.0625	−4.6210
+6.0	0.0156	−6.9315

A deficit of one octave means the relation is running at half the frequency it needs. A deficit of six octaves means it is running at 1.56% of what it needs, which is the situation the companion paper on tuning addresses and which no single adjustment repairs.

4.6 The ceiling, in the frequency register

THEOREM 4.6 — THE CEILING. Since `F <= ln N` — the normalisation bound carried by the surplus functional [Curzi, 2026], the condition `F > c/eta` is satisfiable only if

```
c / eta  <=  ln N        equivalently        f_min  <=  k N^2
```

If `c/eta > ln N`, NO arrangement of those `N` categories is sustainable. Not a better pairing. Not a better schedule. Not a better conversation. Nothing.

Proof: `F > c/eta >= ln N >= F` is a contradiction. **QED.**

For the twelve-state ladder of this framework:

```
ln 12       = 2.484906649788
k * 12^2    = 69.61330555033166 Hz    -> folded 278.453 Hz (j = 2), pitch class 3
max octaves = 2.885390 * ln 12 = 7.169925
```

For a twelve-category system, `c/eta` may not exceed 2.4849, and the corresponding `f_min` is 69.6133 Hz. Above that, the relation cannot be tuned. It must be restructured. A ceiling imposed by one parameter alone, independent of how the remainder is arranged, is a recurring shape in this corpus: the concentration atlas finds validator pools at HHI 0.1900 and 5.26 effective pools against a ceiling of `1/s_max^2 = 11.11` that the leading share alone imposes [Curzi, 2026].

The ceiling by category count, computed:

N	ln N	f_ceiling = k N^2 (Hz)	max octaves above k
2	0.693147	1.9337	2.000000
3	1.098612	4.3508	3.169925
4	1.386294	7.7348	4.000000

```
  5   1.609438          12.0856               4.643856
  6   1.791759          17.4033               5.169925
  7   1.945910          23.6879               5.614710
  8   2.079442          30.9392               6.000000
  9   2.197225          39.1575               6.339850
 10   2.302585          48.3426               6.643856
 12   2.484907          69.6133               7.169925
 16   2.772589         123.7570               8.000000
 17   2.833213         139.7100               8.174926
 20   2.995732         193.3703               8.643856
 25   3.218876         302.1411               9.287712
 41   3.713572         812.6387              10.715104
 64   4.158883        1980.1118              12.000000
```

COROLLARY 4.7 — THE MINIMUM CATEGORY COUNT. A cost ratio of `c/eta` requires at least `N >= e^(c/eta)` categories to be sustainable at all:

```
c/eta = 0.5000  ->  N >=  1.6487  ->  N_min =  2
c/eta = 1.0000  ->  N >=  2.7183  ->  N_min =  3
c/eta = 1.5000  ->  N >=  4.4817  ->  N_min =  5
c/eta = 2.0000  ->  N >=  7.3891  ->  N_min =  8
c/eta = 2.4849  ->  N >= 11.9999  ->  N_min = 12
c/eta = 2.5000  ->  N >= 12.1825  ->  N_min = 13
c/eta = 3.0000  ->  N >= 20.0855  ->  N_min = 21
c/eta = 3.5000  ->  N >= 33.1155  ->  N_min = 34
c/eta = 4.0000  ->  N >= 54.5982  ->  N_min = 55
```

The growth is exponential, and that is the whole content of the corollary. Going from `c/eta = 3` to `c/eta = 4` — a 33% increase in relative cost — raises the minimum viable category count from 21 to 55. Cost does not merely make a relation harder. Past a point it makes the relation's category set inadequate, and category sets are expensive to enlarge honestly.

4.7 The hold line

COROLLARY 4.8. The relation holds exactly (`D = 0`, `Q_inf = 0`) when

```
u*  =  ( e^(c/eta) − 1 ) / (N − 1)
```

Verified for `N = 12`:

```
c/eta = 0.500000  ->  u* = 0.058975   f = f_min =  1.3141 Hz
c/eta = 1.000000  ->  u* = 0.156207   f = f_min =  3.5721 Hz
c/eta = 2.000000  ->  u* = 0.580823   f = f_min = 26.3942 Hz
c/eta = 2.484907  ->  u* = 1.000000   f = f_min = 69.6133 Hz
```

The last row is the ceiling stated in the geometry rather than in the cost: at `c/eta = ln N` the required non-parallelism is exactly 1 — perfect orthogonality — and there is nothing above it.

4.8 The four levers

These are the only four. There is no fifth, and `delta` is not one of them.

```
LEVER              acts on            partial derivative of D
──────────────────────────────────────────────────────────────────────────
L1  raise u        the geometry       dD/du    = −2.885390 (N−1)/(1 + (N−1)u)
L2  raise eta      the conversion     dD/deta  = −2.885390 c / eta^2
L3  cut c          the cost           dD/dc    = +2.885390 / eta
L4  raise N        the category set   dD/dN    = −2.885390 u / (1 + (N−1)u)
```

Each is derived directly from `D = 2.885390 (c/eta − F)` with `F = ln(1 + (N−1)u)`.

THEOREM 4.9 — `delta` IS NOT A LEVER ON THE DEFICIT.

```
D = 2.885390 ∗ (c/eta) − log2(f/k)          contains k, c, eta, N, u
                                            does NOT contain delta
```

Proof: by inspection of Definition 4.4 and of `f = k(1 + (N−1)u)^2`. **QED.**

Consequence, and it is the single most practically important statement in this section: an intervention that acts only on `delta` cannot change the deficit by any amount whatsoever. Trust acts on `delta`. Ritual, place, written record and institutional memory act on `delta`. All of them are valuable and none of them moves the reading. That they act on `delta` is itself a measured claim in this corpus rather than an assumption: the memory-palace analysis decomposes its own effect and finds retention contributing a factor of 4.1667 against surplus 1.8637 and conversion 1.9000, the product 14.7546 matching the direct fixed-point ratio exactly [Curzi, 2026v].

What `delta` does do is scale `Q_inf`, through the identity `Q_inf = −D eta/(2.885390 delta)`. **It multiplies the existing sign.** In a relation with `D < 0` it makes a good destination better. In a relation with `D > 0` it makes a bad destination worse, and this is CC-R15 in frequency form. The companion paper on the sustainable dyad owns the order-of-operations rule that follows; the Worked Example below demonstrates it on the carried case.

4.9 Honesty clauses, restated because this is the section people will quote

(1) ESSENCE, NOT RESONANCE. Every frequency here is an essence frequency: a position on a logarithmic scale of surplus. **Nothing is claimed to vibrate.**

(2) NO CLINICAL CLAIM. The twelve states are not diagnoses; a reading is not an assessment; no procedure here is a treatment. Nothing in this section describes, evaluates, or prescribes for any person's health.

(3) THE READING INHERITS ITS INPUTS. `D` is exactly as defensible as `u`, `eta`, `c` and `N`. A contested decomposition yields a contested frequency. If `eta` or `c` is ordinal, publish the ordering and not the number (SPEC 1.1; CC-R1's demotion travels).

(4) S3 IS THE DESIGN AXIOM. `F` is defined through `N_eff`, which is defined by the affine axiom. Everything in this section uses it.

(5) SURPLUS IS A LOWER BOUND. `f` is a floor. A verdict of ACCUMULATES is robust to representational under-capture; a verdict of DECAYS is not (the sustaining-frequency paper [Curzi, 2026] §3.6).

(6) CONSTRUCTED CASES. Every case in this framework is constructed and labelled as such.

(7) CONSENT IS STRUCTURAL. No intervention derived from a reading may remove a person's ability to decline. The companion paper on tuning owns this and treats it as structure, not advice.

The Engineering

4.10 How to read the table

Three columns, three questions.

"How much pitch do I need?" — the second column, `f_min`. This is the threshold in hertz. It is the number to compare a measured `f` against, and it is the number to fold and publish with its `j` if the reading is to be displayed.

"How far is that?" — the third column, octaves above `k`. This is the threshold expressed as an interval from the ground essence, and it is the column to use when comparing two relations with different costs, because intervals are what compose.

"How many distinct views must this relation field?" — the fourth column, `N_eff = e^(c/eta)`. This is the column to hand to a person who does not want a frequency. It converts the whole apparatus into a sentence: *at the current costs, the pair must be worth 7.4 genuinely different people to break even.*

4.11 The ceiling test comes first, and here is why

SPEC 4.1 (stated early because everything depends on it): compare `c/eta` to `ln N` before computing anything else.

The test is cheap and it discriminates between hypotheses rather than confirming one, which is the form of inference that pays [Platt, 1964]. The reason it must come first is that a deficit does not announce its own kind. A `+5.78` octave deficit at `c/eta = 2.566667` and a `+5.78` octave deficit at `c/eta = 2.100000` with a much lower `f` look identical in the reading and are entirely different problems. The first is unsatisfiable by any arrangement of twelve categories. The second is a large tuning problem.

```
if  c/eta <= ln N     the case is TUNABLE.       Some assignment of u closes it.
if  c/eta >  ln N     the case is STRUCTURAL.    No assignment of u closes it.
                                                 Only L2, L3 or L4 can help.
```

A prescription issued before this test may be unsatisfiable, and unsatisfiable prescriptions are worse than none — they consume the effort that restructuring needed.

4.12 The two gauges, and why they can disagree

The framework's master gauge is headroom:

```
h_t = eta ln N − delta Q_t − C_t
```

This section's gauge is the deficit. **They answer different questions and can disagree in sign.**

Computed on the carried case as found, at `Q = 1`:

```
h  = 0.30 * 2.484907 − 0.09 * 1 − 0.62  =  +0.0355      growth available
D  = +4.338420 octaves                                  short by four and a third octaves
Q_inf = −5.0119                                         decaying
```

Headroom in this sense is organisational slack in the behavioural-theory sense, and the reason it cannot be driven to zero is the same reason waiting time explodes as utilisation approaches one [Kingman, 1962]: headroom is not prudence, it is what keeps the response finite.

Both are correct. Headroom uses the *ceiling* `ln N` — it asks what the best possible geometry could support at this cost. The deficit uses the *achieved* `F` — it asks what this geometry is doing. A system can have room to grow in principle and be shrinking in fact, and this case is exactly that: `h > 0` says the ceiling has not been hit, `D > 0` says the current arrangement is nowhere near it.

The bridge between them is CC-R6's identity, `h_inf = eta (ln N − F(u*))`, verified here at `0.576547` by both routes. In pitch units:

```
PITCH HEADROOM  =  2.885390 * ( ln N − F )  =  octaves between the achieved f and the ceiling
```

For the carried case as found, that is **5.5452 octaves of unused range** against a requirement of 5.9631 and a present position of 1.6247. The relation needs 5.96 octaves, has 1.62, and the ceiling is at 7.17. The requirement fits under the ceiling with 1.21 octaves to spare, which is exactly why the case is tunable.

4.13 The price of an octave, and why the last one cannot be bought

```
1 octave of pitch = 0.346574 nats of surplus = exactly 0.5 bits
equivalently: N_eff must rise by a factor of sqrt(2) = 1.414214
```

Since `N_eff = 1 + (N−1)u`, buying an octave means multiplying `1 + 11u` by `sqrt(2)`. The cost in `u` is not constant. Computed:

```
from u = 0.068735   ->  u must reach 0.134861     (a rise of 0.066126)
from u = 0.300000   ->  u must reach 0.461920     (a rise of 0.161920)
from u = 0.500000   ->  u must reach 0.744763     (a rise of 0.244763)
from u = 0.800000   ->  u must reach 1.169027     IMPOSSIBLE: u <= 1
```

The octaves get more expensive as the scale is climbed, and the last one cannot be bought in `u` at all. Above `u = (12/sqrt(2) − 1)/11 = 0.680480`, no further octave is available from the geometry at all, no matter what anyone does. This is the ceiling appearing as a marginal cost rather than as a bound, and it is the honest answer to "why not just find more difference?"

4.14 LEVER 1 — RAISE `u` What it is. Increase the represented non-parallelism between the parties: hold states the other does not hold, bring perspectives that do not already point the same way, decorrelate the errors. The between-persons measurement of `u` and the intervention that raises it are owned elsewhere in this series. The corpus's internal-surplus work [Curzi, 2026o] supplies the cautionary magnitude: restructuring a framing set raised the audited count by 25.4 per cent — `F_int` from 1.4883 to 1.7148 nats, 95.70 per cent of the ceiling — at identical reconciliation load, and still left the change in stored order negative at `−0.0411`.

Its sensitivity. `dD/du = −2.885390 (N−1)/(1 + (N−1)u)`. Computed at three operating points:

```
at u = 0.068735 (carried case, as found)    −18.0739 octaves per unit u
at u = 0.692373 (carried case, after)        −3.6837 octaves per unit u
at u = 0.018614 (the commons, the companion paper on families, teams and commons)   −26.3450 octaves pe
```

Strongest where `u` is smallest. The first unit of difference is worth far more than the last — this is the strict concavity of the surplus functional [Curzi, 2026] showing up as a lever price.

Its bound. `u <= 1`. The maximum achievable frequency is `k N^2`, and past `u = 1` there is nothing.

Its failure mode. Manufactured difference. **`u` measures represented non-parallelism, not disagreement, and CC-R8 is the trap: perfect opposition gives `u = 0` exactly as perfect agreement does, because `u` depends only on `|rho|`.** A contrarian recruited to raise `u` raises cost and returns nothing. The companion paper on unison owns this result.

4.15 LEVER 2 — RAISE `eta` What it is. Increase the fraction of interaction that becomes durable structure: write the decision down, build the artefact, make the practice repeatable, give the insight somewhere to live.

Its sensitivity. `dD/deta = −2.885390 c / eta^2`. Computed:

```
at eta = 0.30, c = 0.62   −19.8771 octaves per unit eta
at eta = 0.44, c = 0.40    −5.9615 octaves per unit eta
at eta = 0.34, c = 0.825  −20.5921 octaves per unit eta
```

The `eta^2` in the denominator is the important feature. At low `eta` the lever is enormously powerful; at high `eta` it saturates. A system converting 30% of its interaction into structure gains 19.9 octaves per unit of improvement; one already converting 44% gains 6.0.

Its bound. `eta` is a conversion efficiency, and improvements beyond the system's mechanical maximum are unavailable. In many cases the required `eta` exceeds 1 and the lever is simply not there — the worked example below shows exactly that on the carried case.

Why the companion paper on the sustainable dyad finds it usually cheapest in a dyad. Raising `u` requires finding or making real difference between two people, which is slow and sometimes impossible. Raising `eta` requires changing what happens to what they already produce, which is a mechanism question and often a matter of a notebook.

4.16 LEVER 3 — CUT `c`

What it is. Reduce coordination overhead, friction, correction load, and dissipation. The system-dynamics tradition models exactly these as rates draining a level, with delays [Forrester, 1961], and reports that people systematically misjudge stocks that have them [Sterman, 1989] — which is the empirical case for computing the ratio rather than estimating it.

Its sensitivity. `dD/dc = +2.885390 / eta`. Computed:

```
at eta = 0.30    +9.6180 octaves per unit c
at eta = 0.44    +6.5577 octaves per unit c
at eta = 0.34    +8.4864 octaves per unit c
```

This is the only lever whose sensitivity is constant in its own variable — linear in `c`, so each unit cut is worth the same as the last. That makes it the most predictable lever and the easiest to plan against.

Its bound. `c >= 0`, and at `c = 0` the threshold falls to `f_min = k` exactly, which any non-zero `u` clears. A costless relation is sustained by the smallest difference.

Its failure mode. Cost reduction that reduces `eta` at the same time. Cutting the meeting that was the conversion mechanism cuts `c` and `eta` together, and since `D` depends on `c/eta`, the ratio may not move at all. Cut cost that is not doing conversion work; SPEC 4.11 states the test.

And CC-R18's counter-pressure, which belongs here: an over-framed system accumulates *more* stored order than a calibrated one, because dogma banks faster than knowledge. The bill arrives on a lag as correction load, `C_t = c_0 + kappa Z_t`, rising with accumulated error. An organisation optimising raw `Q` will be driven away from calibration, and its `c` will rise later for reasons it will not connect to the choice.

4.17 LEVER 4 — RAISE `N`

What it is. Add genuinely distinct categories to the description. Not finer gradations of what is already there — **genuinely distinct**, in Ashby's sense that the controller's variety must match the disturbance's [Ashby, 1956], in the sense that the added dimension carries variance the existing ones do not.

Its sensitivity. `dD/dN = −2.885390 u / (1 + (N−1)u)`. Computed:

```
at u = 0.068735    −0.1129 octaves per added category
at u = 0.692373    −0.2319 octaves per added category
at u = 0.018614    −0.0446 octaves per added category
```

This is the weakest lever per unit, and it is weakest exactly where a large deficit makes it most tempting. At the near-parallel geometry of the near-parallel commons case carried in this series, an added category buys 0.0446 octaves. Closing a six-octave deficit that way would require 145 new categories.

The mechanism is the Parallel Wall. `N` enters `F` only multiplied by `u`. Adding categories to a set of near-parallel parties adds dimensions nobody is using. A twelfth window in a tower where everyone is looking the same way is a twelfth view of the same thing.

Its real value is not in the deficit but in the ceiling. Raising `N` raises `ln N`, which raises the maximum satisfiable `c/eta`. `N` is how a structural case becomes a tunable one, and that is the only situation in which it is the right first move.

Its bound and its cost. The cost side is the communication overhead that grows as `n(n−1)/2` with roster size [Brooks, 1975, Conway, 1968], and it opposes the ceiling gain directly. The time-bandwidth constraints [Gabor, 1946] bind on the rendering side: past `N = 17` the folded lattice's minimum gap collapses from `50/49` (34.9756 cents) to `289/288` (6.0008 cents), and the symbol rate drops by a factor of exactly 8. Past `N = 25` a single integer transposition no longer fits the audible band at all. And the honest constraint above all of those: a category added for arithmetic convenience is not a category. Section 5's validity conditions govern, and an arbitrary addition yields an arbitrary frequency.

4.17b Choosing between the levers — and why the partials do not decide it

The four sensitivities in §4.8 are in four different units. `dD/du` is octaves per unit of non-parallelism; `dD/dc` is octaves per unit of cost. They cannot be compared to each other directly, and a reader who ranks the levers by the size of their partial derivatives is comparing octaves-per-apple to octaves-per-orange.

What makes them comparable is a cost model: the effort, in whatever currency the situation actually spends, required to move each parameter by one unit. That model is not in the formalism and this framework does not pretend to supply it. What the formalism supplies is the *exchange rate* — how many octaves each unit buys — and that is genuinely useful even without the prices, for three reasons.

One: it identifies levers that are absent rather than expensive. The Worked Example's `L2` on the carried case requires `eta >= 1.101077`. That is not a costly lever; it is a lever that does not exist, and no cost model is needed to see it.

Two: it identifies levers that are exhausted. `L4` at `u = 0.018614` buys 0.0446 octaves per category. Against a 6.46-octave deficit that is 145 categories, and no plausible cost model makes 145 new genuinely distinct categories the cheap option.

Three: it gives the shape of the trade, which is stable even when the prices are not. `dD/deta` carries `eta^2` in its denominator, so the `eta` lever is always most powerful where `eta` is lowest. `dD/du` carries `1 + (N−1)u`, so the `u` lever is always most powerful where `u` is lowest. Both levers are strongest in the systems that most need them, which is a real and non-obvious structural fact, and it is the opposite of the usual experience that improvement gets easier once things are already going well.

The honest summary:

```
the formalism gives:   how many octaves a unit of each parameter buys
the situation gives:   what a unit of each parameter costs
the decision needs:    both, and only one of them is in this framework
```

4.18 The classification — three kinds of case, and they need different work

```
D <= 0                      SUSTAINABLE.  Report the margin. Maintain; do not assume.
D >  0  and c/eta <= ln N    TUNABLE.      Some assignment of u closes it. Choose levers
                                           by sensitivity, subject to CC-R15's ordering.
c/eta > ln N                STRUCTURAL.   No assignment of u closes it. L1 is unavailable
                                           at any value. Only L2, L3, L4 −
and L4 only to
                                           make the case tunable, not to close it.
```

An account that showed only tunable cases would be selecting its evidence. The carried commons case §3.3 — six members drawn from one doctrine — reads `u = 0.018614`, `f = 0.7017 Hz`, `f_min = 61.9349 Hz`, `D = +6.4638 octaves`, `Q_inf = −8.4630`. The companion paper on tuning carries it, and the finding it carries is that no single lever closes a 6.46-octave gap: raising `u` to 0.95 — near-orthogonality, for six people drawn from one doctrine — closes it by 0.0332 octaves and no more; raising `eta` to 0.90 leaves `+2.1075`; cutting the per-link cost to 0.020 leaves `+2.0085`. Some systems cannot be tuned. The diagnostic is what identifies which case is in hand, and that is most of its value.

4.19 Reporting: precision, sensitivity, and the rule that follows Because `d ln f_min / d(c/eta) = 2` exactly, and because `D = 2.885390 (c/eta − F)`:

```
an uncertainty of dx nats in (c/eta)  ->  an uncertainty of 2.885390 * dx octaves in D
```

SPEC 4.19 (the reporting rule): publish `D` together with its sensitivity to the parameter carrying the largest uncertainty, and publish the ceiling test. A deficit of `+0.5 +/− 1.5` octaves is a statement that the relation might be fine; a deficit of `+0.5 +/− 0.1` is a statement that it is not. The two look identical if only the point estimate is printed, and the difference between them is the whole of the reading's usefulness.

And a specific trap: precision in a constructed case does not transfer. Every figure in this framework's worked cases is exact because the parameters are chosen. A measured case has no such property, and printing `30.1588 Hz` from an `eta` known to one significant figure is precision theatre.

4.20 Cents, for fine work

```
1 nat  = 3462.4681 cents        1 octave = 1200 cents
D in cents = 3462.4681 * (c/eta − F)
```

The carried case as found reads `+4.3384 octaves = +5206.1 cents`; after the intervention it reads `−3.5910 octaves = −4309.2 cents`. Cents matter when a deficit is small: row 8 of the sustaining-frequency tabulation [Curzi, 2026], at `+0.033406` octaves, is 40 cents — about a fifth of a semitone — and is the difference between a practice that persists and one that must be rebuilt each generation.

A rendering note, stated once: `N = 12` is below the `N = 17` cliff, so everything here renders in the `50/49` regime at a minimum separation of 34.9756 cents, requiring `32/(4 pi) = 2.546479` cycles per resolved step. A 40-cent deficit is therefore just above the minimum resolvable step of the lattice it is expressed in, and a 6-cent one would not be. The companion paper on octave folding owns the folding and states the regime once.

4.21 The citation contract This section fixes five objects for the whole framework. **Later papers in the corpus cite them by name and number and do not restate the derivations.** The contract is set out here so that a reader can check any later section against it.

```
OBJECT                          DEFINED IN     CITED AS
-------------------------------------------------------------------------------
f_min = k e^(2c/eta)            Def 4.1        "the minimum sustainable frequency (Section 4)"
the three regimes               Thm 4.2        "Theorem 4.2"
octaves = 2.885390 (c/eta)      Thm 4.3        "the octave form (Section 4)"
the octave deficit D            Def 4.4        "the octave deficit (Section 4)"
the deficit identity            Thm 4.5        "Theorem 4.5"
the ceiling c/eta <= ln N       Thm 4.6        "the ceiling (Section 4)"
the minimum category count      Cor 4.7        "Corollary 4.7"
the hold line                   Cor 4.8        "the hold line (Section 4)"
the four levers and partials    Sec 4.8        "the four levers (Section 4)"
delta is not a lever            Thm 4.9        "Theorem 4.9"
the three-way classification    Sec 4.18       "the classification (Section 4)"
```

Five objects are deliberately NOT owned here, and later papers in the corpus must go to their owners:

```
the differentiation index n and its count argument      -> Section 5
the elicitation procedure for intensities               -> Section 6
```

```
octave folding and the odd-part class rule          -> The companion paper on octave folding
f_rel = k(1 + 11u)^2 as the dyad's interval         -> The companion paper on the dyad interval
the CC-R15 order-of-operations rule                    -> The companion paper on the sustainable dy
```

The division matters. This section states what a deficit is and what closes it. It does not state how to obtain `u`, how to defend `n`, or in what order to pull the levers in a real bond — those are separate results with separate defences, and a keystone section that absorbed them would make all of them harder to check.

4.22 What a reading is not permitted to conclude The last thing this section owns is a boundary, and it is placed here rather than in the failure modes because it is not a mistake anyone makes by accident.

A reading does not conclude that a relation should end. `D > 0` says a configuration is losing order at a computable rate. It does not weigh what the relation is for, what else it produces, or what ending it would cost. The formalism has no term for any of that and does not acquire one by being applied to people.

A reading does not conclude that anyone is at fault. `u` is a property of two representations jointly. `c` and `eta` are properties of an arrangement. There is no single-party term in the whole apparatus.

A reading does not conclude anything about a person's health, capacity, or prospects. The twelve states are not diagnoses; the ladder is not an assessment instrument; the deficit is not a prognosis. Nothing in this framework is clinical and nothing in it may be presented as clinical.

A reading does not compel anyone. Any intervention that follows from it is subject to the standing constraint that no person's ability to decline may be removed, and to the standing prohibition on anything masked or below threshold. Consent is structure here, not advice, and The companion paper on tuning builds the interventions to that requirement rather than adding it afterwards.

And a reading is not evidence about anything outside its own parameters. `Q_inf` is where a configuration goes if the configuration holds. It says nothing about whether it will hold, and CC-R23 stands: the survivorship likelihood ratio for testimonial evidence that a configuration held is exactly `1.000000`.

The Specification

SPEC 4.1 — CEILING FIRST. Compare `c/eta` to `ln N` before any other computation. Classify the case as sustainable, tunable, or structural (§4.18) before issuing any prescription.

SPEC 4.2 — DECLARE `N` and the ordering that produced it, with the note that the frequency inherits the defensibility of the count.

SPEC 4.3 — ONE CLOCK. `eta`, `c` and `delta` on one declared period, named in the same passage as the ratio.

SPEC 4.4 — DECLARE SCALE STATUS of `u`, `eta` and `c`. Ordinal inputs yield an ordering, never a frequency.

SPEC 4.5 — REPORT THE TRIPLE. Every reading reports `f`, `f_min`, and `D`, with the folded values and their `j` where a band is used.

SPEC 4.6 — REPORT THE SENSITIVITY of `D` to the least certain parameter (SPEC 4.19).

SPEC 4.7 — REPORT THE PITCH HEADROOM `2.885390 (ln N − F)` alongside the deficit, so that a reader can see how much of the scale remains.

SPEC 4.8 — CROSS-CHECK VIA THE IDENTITY. Verify `D = −2.885390 (delta/eta) Q_inf`. A disagreement means an arithmetic error, since the identity is exact.

SPEC 4.9 — NEVER USE `delta` AS A DEFICIT LEVER. Theorem 4.9. Interventions on `delta` are reported separately, as multipliers on an existing sign.

SPEC 4.10 — ORDER OF OPERATIONS. Fix the numerator before lowering the loss. Raise `eta` or cut `c` until `eta F > c`, and only then invest in `delta`. The companion paper on the sustainable dyad owns this rule; the Worked Example demonstrates it.

SPEC 4.11 — TEST COST CUTS FOR CONVERSION DAMAGE. Before cutting `c`, ask whether the cost being cut is performing conversion work. If it is, the ratio `c/eta` may not improve at all.

SPEC 4.12 — RAISE `N` ONLY WITH REAL CATEGORIES. An added category must satisfy Section 5's validity conditions and must carry variance the existing categories do not. Adding categories to a near-parallel set buys almost nothing — 0.0446 octaves each at `u = 0.018614`.

SPEC 4.13 — RESPECT THE RENDERING BOUNDS if the reading is displayed: the regime (`50/49` for `N <= 16`), the transposition `j`, a windowed envelope, `N <= 25` for a single transposition.

SPEC 4.14 — NO AGGREGATION. Per relation, per site, per period. Ratios do not average.

SPEC 4.15 — MEAN OF `F`, NOT `F` OF THE MEAN, wherever a reading spans more than one pair.

SPEC 4.16 — LABEL EVERY CASE CONSTRUCTED.

SPEC 4.17 — ESSENCE, NOT RESONANCE, in every passage printing a hertz value.

SPEC 4.18 — NO CLINICAL LANGUAGE. A reading is of a configuration of parameters, never of a person's condition.

SPEC 4.19 — THE REPORTING RULE, as stated in §4.19.

SPEC 4.20 — CITE, DO NOT RE-DERIVE. Later papers in the corpus cite this section for `f_min`, the octave deficit, the ceiling, and the four levers. The definitions above are the framework's, and no part of the framework may silently redefine them.

The Failure Modes

FM 4.1 — TUNING ABOVE THE CEILING. *Signature:* a lever prescription for a case with `c/eta > ln N`. *Mechanism:* the condition `F > c/eta` is unsatisfiable, but the deficit is just a larger positive number and looks like more of the same problem. *Diagnostic:* compute `c/eta` and `ln N`. This test costs one logarithm and changes the kind of answer. *Fix:* SPEC 4.1. The remedies are `eta`, `c`, and genuinely new categories — and `N` only in order to make the case tunable, never to close the deficit directly.

FM 4.2 — TRUST FIRST. *Signature:* a failing relation into which trust, ritual, or continuity is invested, followed by acceleration of the failure. *Mechanism:* CC-R15, exactly. `delta` divides a negative numerator. Computed on the carried case: `delta` 0.09 -> 0.05 moves `Q_inf` from `-5.0119` to `-9.0215`, and the deficit does not move at all, remaining `+4.338420` to every digit. *Fix:* SPEC 4.10. Fix the numerator, then lower the loss.

FM 4.3 — SINGLE-LEVER THINKING ON A LARGE DEFICIT. *Signature:* a plan to close six octaves by one adjustment. *Mechanism:* the levers saturate. `u` is bounded at 1; `eta` at its mechanical maximum; `c` at 0; `N` buys `0.0446` octaves per category at low `u`. *Fix:* compute what each lever alone would require (the Worked Example for the dyad; the companion paper on tuning for the commons), and if none is reachable, say so. **Some systems must be restructured.** The corpus's launch-budget result gives the shape of what restructuring costs [Curzi, 2026p]: with a rectified fund `L^Q = max(0, eta*M - C)` and fixed point `L/lambda`, departure from an existing arrangement is feasible only where that fixed point exceeds the barrier, after `t* = ln(1 - E_b/L_inf)/ln(1-lambda)` periods.

FM 4.4 — PRECISION THEATRE. *Signature:* `f_min = 30.1588 Hz` from `eta = 0.3`. *Mechanism:* exact arithmetic on inexact inputs. *Fix:* SPEC 4.6 and SPEC 4.19.

FM 4.5 — THE ORDINAL SMUGGLE. *Signature:* a numeric `c/eta` from rank-ordered inputs. *Mechanism:* division requires a ratio scale. An order-preserving relabel moved `f_min` from

195.0279 Hz to 10648.1604 Hz in the sustaining-frequency paper [Curzi, 2026] without changing anything an ordinal instrument could see. *Fix:* SPEC 4.4 and CC-R1's route out — Rasch instrumentation at 123 items per term.

FM 4.6 — BUYING `u` PAST 1. *Signature:* a plan requiring `u > 1`, or a required `u` that is not checked against 1. *Mechanism:* `u = 1 − (x.y)^2 <= 1` by construction. *Fix:* compute the required `u = (e^(c/eta) − 1)/(N−1)` and check it against 1 — which is the same as the ceiling test, seen from the geometry.

FM 4.7 — ADDING CATEGORIES TO A PARALLEL SET. *Signature:* a richer taxonomy adopted to raise `N`, with no change in the reading. *Mechanism:* the Parallel Wall. `N` enters `F` only multiplied by `u`. At `u = 0.018614` a new category is worth 0.0446 octaves. *Fix:* raise `u` first, or accept that `N` is being raised to lift the ceiling and not to close the deficit.

FM 4.8 — `f_min` PLAYED AS A TONE. *Signature:* any proposal to sound `f_min` at anyone. *Mechanism:* the register slip. `f_min` is a threshold in a computation and the frequencies are essence frequencies. *Fix:* rendering is governed by the companion paper on rendering, under the open problem stated in this series — no established mechanism for sequence-specific acoustic demodulation — and under the absolute no-subliminals rule.

FM 4.9 — THE NEGATIVE STOCK. *Signature:* a plan that anticipates arriving at `Q_inf = −5.0119`. *Mechanism:* admissibility fails first, at `t = 1` in the carried case. *Fix:* read a negative fixed point as a direction and a rate.

FM 4.10 — CUTTING THE CONVERSION MECHANISM. *Signature:* `c` falls, `eta` falls with it, `D` does not move. *Mechanism:* `D` depends on the ratio. *Fix:* SPEC 4.11.

FM 4.11 — THE AGGREGATED READING. *Signature:* one deficit for a whole organisation. *Mechanism:* CC-R10 — distortion cancels across sites, overstating bankable input by 4.26x; CC-R16 — identical aggregates hide topologies that lose 77% versus 21% of surplus to the same departure. *Fix:* SPEC 4.14.

FM 4.12 — THE JENSEN SHORTCUT. *Signature:* a group reading from `F(mean u)`. *Mechanism:* strict concavity inflates `F`, inflates `f`, and under-reports `D`. *Fix:* SPEC 4.15. Worth 6.2% in Series B's worked office.

FM 4.13 — THE DEFICIT READ AS A JUDGEMENT OF PEOPLE. *Signature:* "our reading is `+4.34`, so one of us is the problem." *Mechanism:* a four-parameter configuration is being read as a property of a person. The deficit is a property of an arrangement. *Fix:* state it as the arrangement. The most common real repair — the carried case's own — is that one party takes on states the other does not hold and the pair adopts a record. Neither of those is a correction of a person.

Worked Example

THE CARRIED DYAD — THE FULL ARC, 2a TO 2b **This case is CONSTRUCTED. It is the dyad of the framework's carried dyad and it is the only case used in this section.** Its parameters are chosen, not observed. It is not an assessment of any person and no clinical, diagnostic or therapeutic claim is made or implied. The frequencies are essence frequencies and nothing is claimed to resonate.

Phase 2a — as found

```
A = {4: 0.7, 6: 0.5, 9: 0.4}
B = {4: 0.6, 6: 0.55, 9: 0.3, 10: 0.2}

eta = 0.30      c = 0.62      delta = 0.09      N = 12, one declared period
```

Step 1 — the geometry.

```
cos theta = 0.965021        u = 0.068735        F = 0.563085 nats
N_eff     = 1.7561
f         = k (1 + 11u)^2 = 1.4908 Hz           folded 381.644 Hz (j = 8)
```

The pair agree about nearly everything, and that is the whole problem. No one involved would describe it as a problem. Two people 0.965 aligned field 1.76 people's worth of distinct perspective between them.

Step 2 — the ceiling test, first.

```
c/eta   = 0.62 / 0.30 = 2.066667
ln N    = 2.484907
2.066667 <= 2.484907        BELOW THE CEILING  ->  the case is TUNABLE
```

Step 3 — the reading.

```
f_min            = k e^(2 * 2.066667)    = 30.1588 Hz     folded 241.270 Hz (j = 3)
octaves required = 2.885390 * 2.066667   =  5.963140
octaves present  = log2(1.4908 / k)      =  1.624719
OCTAVE DEFICIT D                         = +4.338420      (+5206.1 cents)
pitch headroom   = 2.885390 (ln 12 - F)  =  5.5452 octaves
N_eff required   = e^2.066667            =  7.8985
VERDICT: DECAYS
```

Step 4 — cross-check.

```
Q_inf = (0.30 * 0.563085 - 0.62) / 0.09 = -5.0119
identity: -2.885390 * (0.09/0.30) * (-5.011940) = +4.338420        error 8.88e-16
admissibility fails at t = 1;  the stock goes negative at t = 2
```

The reading in one sentence: this pair must be worth 7.9 genuinely different people to break even, and they are worth 1.76.

The CC-R15 demonstration — and it is the case's whole point The intuitive repair is to add trust. It is what everyone reaches for, it is not foolish, and it is well supported in ordinary experience. Trust lowers `delta`. Suppose it works completely and `delta` moves from 0.09 to 0.05.

```
BEFORE:   Q_inf = -0.451075 / 0.09 = -5.0119
AFTER:    Q_inf = -0.451075 / 0.05 = -9.0215        WORSE by 80%

DEFICIT BEFORE: +4.338420 octaves
DEFICIT AFTER:  +4.338420 octaves                   UNCHANGED, to every digit
```

The numerator `eta F - c = -0.451075` is negative. Dividing a negative number by a smaller positive denominator makes it more negative. Trust did not add order; it amplified the sign the relation already had, and the sign was negative.

And the deficit did not move at all, because `delta` is not in it (Theorem 4.9). This is the sharpest available statement of CC-R15 and it is why the frequency register earns its place: the instrument reads unchanged while a successful, well-intentioned, expensive intervention doubles the rate of loss. A reader with the deficit and without Theorem 4.9 would interpret the flat reading as "give it more time." A reader with both interprets it correctly the first time: the mechanism has not been touched.

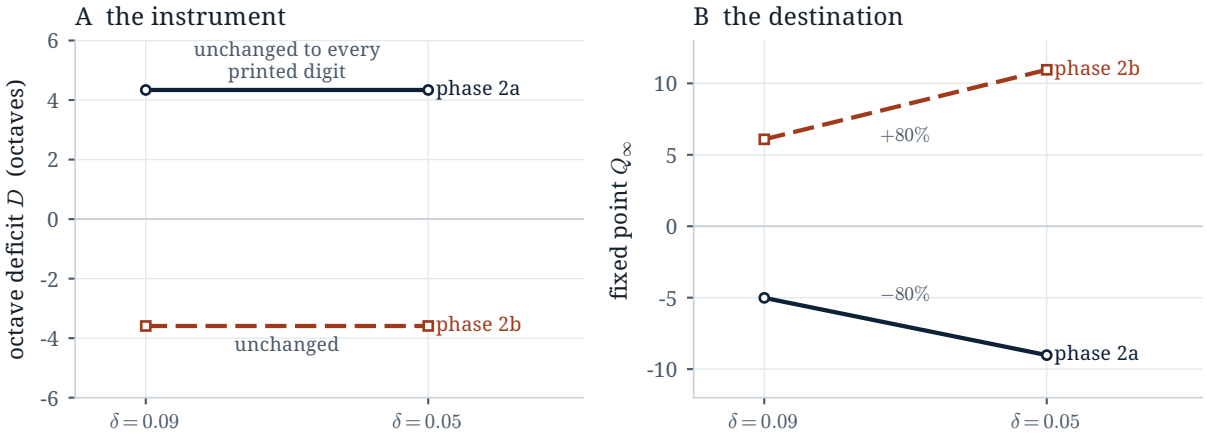


Figure 41.2: Halving the decay rate leaves the octave deficit unchanged to every printed digit and moves the fixed point by 80 per cent in whichever direction it already pointed — worse when the numerator is negative. Panels are separate axes because the two quantities have different scales. Modelled on the constructed carried dyad: $u = 0.068735 \to 0.692373$, $\eta = 0.30 \to 0.44$, $c = 0.62 \to 0.40$, $N = 12$, with δ moved from 0.09 to 0.05 at each phase. Built by `figs/src/CCP-041_fig03.py`.

The four levers on 2a, each alone Required: `F >= c/eta = 2.066667`, i.e. `N_eff >= e^2.066667 = 7.898451`.

```
L1  RAISE u
    need u >= (7.898451 - 1)/11 = 0.627132        from 0.068735 -- a rise of x9.12
    at that u, f = 30.1588 Hz = f_min exactly, D = 0
    REACHABLE, but 0.627132 is a large amount of real difference for two people
    who currently sit at cos theta = 0.965

L2  RAISE eta
    need eta >= c/F = 0.62 / 0.563085 = 1.101077
    IMPOSSIBLE. A conversion efficiency above 1 would mean more structure created
    than interaction supplied. The lever is not merely expensive; it is absent.

L3  CUT c
    need c <= eta F = 0.30 * 0.563085 = 0.168925
    from 0.62 -- a cut of 72.8%
    REACHABLE in principle; a 72.8% reduction in coordination overhead is a large project.

L4  RAISE N
    need 1 + (N-1)(0.068735) >= 7.898451  ->  N >= 101.3635  ->  N_min = 102
    NOT REACHABLE. The Parallel Wall: at u = 0.068735 a new category buys 0.1129 octaves,
    and 4.34 octaves would take 38 of them even ignoring that u itself would change.
```

No single lever is comfortable and one of them does not exist. This is the ordinary situation and it is why the classification in §4.18 matters more than any individual sensitivity: the case is tunable, so a combination exists, and the combination is what the arc demonstrates.

Phase 2b — after one intervention, correctly ordered **CONSTRUCTED.** Two changes, made together:

1. **B takes on two states A does not hold.** This raises `u`. It is a change in what B is actually doing, not a rhetorical adoption of a contrary position — CC-R8 forbids the latter from

working, since perfect opposition reads as `u = 0` exactly as agreement does.
2. The pair adopt a decision log. This raises `eta` — what they work out now has somewhere to live — and lowers `c`, because the record does the remembering that the coordination overhead was paying for.

One realising vector, for concreteness:

```
B' = {2: 0.95, 4: 0.60, 6: 0.55, 9: 0.30, 10: 0.20, 12: 0.8391}
cos theta = 0.554643      u = 0.692373        (the framework carries this as 0.6924)

u    0.068735  ->  0.692373
eta  0.30      ->  0.44
c    0.62      ->  0.40
delta unchanged at 0.09
```

The reading, recomputed from scratch:

```
F                = ln(1 + 11 * 0.692373)  =  2.153633 nats
N_eff            =  8.6161
f                = k * 8.6161^2           = 35.8882 Hz     folded 287.106 Hz (j = 3)
octaves present  =  6.214071

c/eta            = 0.40/0.44 = 0.909091
ceiling test     : 0.909091 <= 2.484907   BELOW, comfortably
f_min            = k e^(2 * 0.909091)     =  2.9782 Hz     folded 381.212 Hz (j = 7)
octaves required =  2.623082

OCTAVE DEFICIT D = -3.590989 octaves      (-4309.2 cents)
N_eff required   = e^0.909091 = 2.4821
pitch headroom   = 2.885390 (ln 12 - 2.153633) = 0.9559 octaves
VERDICT: ACCUMULATES
```

Cross-checks:

```
Q_inf  = (0.44 * 2.153633 - 0.40)/0.09 = +6.0844
identity: -2.885390 * (0.09/0.44) * 6.0844 = -3.590989           exact
trajectory from Q_0 = 1.0:  t=1 +1.4576   t=5 +2.9116   t=20 +5.3134   t=40 +5.9675
convergence: |Q_20 - Q_inf| = 0.7710275635960624
             (0.91)^20 |Q_0 - Q_inf| = 0.7710275635960623
admissibility holds at every step
```

The swing: `D` from `+4.338420` to `-3.590989` — 7.929409 octaves. `Q_inf` from `-5.0119` to `+6.0844`. The pair now needs to be worth 2.4821 distinct people and is worth 8.6161.

Where the 7.929409 octaves came from The swing decomposes exactly, and the decomposition is worth having because it shows which half of the repair did the work.

The path-free split — into the cost side and the surplus side — follows directly from `D = 2.885390 (c/eta - F)`:

```
from the cost ratio:  2.885390 * (0.909091 - 2.066667)  =  -3.340058 octaves
from the surplus   : -2.885390 * (2.153633 - 0.563085)  =  -4.589352 octaves
total                                                      -7.929409 octaves
```

Raising `u` did 57.9% of the work; the cost-side changes did 42.1%.

The split *within* the cost side is path-dependent, and the framework says so rather than picking a flattering order. Attributing a joint change to individual parameters requires an order, and different orders give different attributions even though the total is identical:

```
PATH 1 -- u first, then eta, then c
    u   0.068735 -> 0.692373    -4.589352
    eta 0.30     -> 0.44        -1.897363
    c   0.62     -> 0.40        -1.442695
    total                       -7.929409

PATH 2 -- c first, then eta, then u
    c   0.62     -> 0.40        -2.115953
    eta 0.30     -> 0.44        -1.224105
    u   0.068735 -> 0.692373    -4.589352
    total                       -7.929409
```

The `c` contribution is `-1.442695` on one path and `-2.115953` on the other — a 47% difference in the attribution, from the same intervention producing the same result. The reason is elementary: `c` is divided by `eta`, so cutting cost is worth more when `eta` is still low.

The `u` contribution is identical on both paths (`-4.589352`), because `F` is separable from the cost ratio in `D` and nothing multiplies it.

The practical rule that follows: report the path-free split, and if a per-parameter attribution is wanted, report the order used. An attribution presented without its order is a number that could have been 47% different for no reason but the sequence in which someone chose to describe events.

Now trust helps Repeat the exact intervention that made things worse in phase 2a — `delta` from 0.09 to 0.05:

```
Q_inf  =  (0.44 * 2.153633 - 0.40) / 0.05  =  +10.9520          from +6.0844
deficit                                    =   -3.590989        unchanged, as always
```

The same intervention, applied to the same pair, with the same magnitude, in the same direction. At 2a it took `Q_inf` from `-5.0119` to `-9.0215`. At 2b it takes it from `+6.0844` to `+10.9520`.

The arithmetic is not subtle and it is not a special case. `delta` multiplies. It has no sign of its own.

> THE ARC: fix the numerator, then lower the loss. Reverse the order and the same intervention accelerates the failure.

The honest note on where 2b leaves them The repair worked, and it worked by spending most of the range that twelve categories make available:

```
pitch headroom at 2a:  5.5452 octaves
pitch headroom at 2b:  0.9559 octaves
```

They have 0.96 octaves of the scale left. A further improvement of the same kind — more represented difference between two people described by twelve states — is close to exhausted. The next order of gain is not available from `u` at this `N`. It would require genuinely new categories, and per §4.17 and SPEC 4.12 those must be real categories and not finer slicing.

This is not a defect in the repair. It is the ceiling doing what a ceiling does, and the value of the reading is that the pair can see it coming rather than discovering it as an unexplained plateau.

The whole reading, as a card

```
------------------------------------------------------------------------
CARRIED DYAD -- CONSTRUCTED CASE -- N = 12 -- one declared period
------------------------------------------------------------------------
                                2a (as found)      2b (after)
u                                  0.068735          0.692373
F (nats)                           0.563085          2.153633
N_eff                              1.7561            8.6161
f (Hz)                             1.4908           35.8882
  folded 220-440                 381.644 (j=8)     287.106 (j=3)
eta                                0.30              0.44
c                                  0.62              0.40
c/eta                              2.066667          0.909091
ceiling test (ln 12 = 2.484907)    BELOW: tunable    BELOW: comfortable
f_min (Hz)                        30.1588            2.9782
  folded 220-440                 241.270 (j=3)     381.212 (j=7)
octaves required                   5.963140          2.623082
octaves present                    1.624719          6.214071
OCTAVE DEFICIT                    +4.338420         -3.590989
  in cents                        +5206.1           -4309.2
N_eff required                     7.8985            2.4821
pitch headroom (octaves)           5.5452            0.9559
delta                              0.09              0.09
Q_inf                             -5.0119           +6.0844
VERDICT                            DECAYS            ACCUMULATES
------------------------------------------------------------------------
delta 0.09 -> 0.05:  Q_inf        -9.0215          +10.9520
                     deficit      +4.338420        -3.590989   (UNCHANGED both times)
------------------------------------------------------------------------
ESSENCE frequencies, not physical resonances. Constructed case. Not an
assessment of any person. No clinical claim.
------------------------------------------------------------------------
```

Section 5 takes up the index `n` and the argument that a count escapes the demotion. Every part of the framework after it that needs `f_min`, the octave deficit, the ceiling, or the four levers cites this section.

5 The psychological essence index

Section 4 ended with a diagnostic that needs one input not yet supplied. The minimum sustainable frequency `f_min = k e^(2c/eta)` fixes the pitch a relation must reach. The octave deficit gives how far short it falls. Neither is computable until something on the psychological side has a frequency at all, and nothing does until a psychological category has been assigned an index `n`.

That assignment is the whole of this part of the framework's exposure, and every downstream reading stands on it. Get it wrong — or, worse, get it arbitrarily — and the framework has produced a long sequence of precise-looking numbers that mean nothing, which is the exact failure mode Series B spent a volume learning to detect in other people's work.

This section owns the index. Later papers in the corpus cite it and do not redefine it. It states what `n` is, makes the argument that the index escapes the scale demotion its neighbours inherit, presents the twelve-state ladder that the rest of the framework uses, grades four rival

orderings against the governing validity conditions, rejects three of them in print, and states the limit that travels with every frequency the framework publishes.

Two prohibitions govern the section from its first line and are restated at its last. The frequencies here are essence frequencies, not physical resonances; no person, tissue, or molecule is claimed to resonate at any value in this framework. **And the twelve states are not diagnoses.** They are not a classification of persons, not a severity scale, not an assessment instrument, and nothing in this section or any later one licenses their use as one.

The Law

5.1 The statement

```
THE PSYCHOLOGICAL ESSENCE INDEX

    n  :=  the number of distinguishable components required to specify the state

and, through the essence law (the essence law of Series B(CCP-017)):

    f(n) = k * n^2            k = (9/8)^2 / phi^2 = 0.483425732988414 Hz
```

Read it slowly, because every word is doing work.

"The number" — this is a cardinality, not a position, not a rating, not a score. It answers *how many*, and the answer is an integer.

"Distinguishable" — components that the encoding cannot tell apart count once. Two names for one component are one component. This is the clause that does the most work and takes the most abuse, and Section 6 is largely about enforcing it.

"Components" — the constituents of a specification, not the parts of a thing. A component is something that must be stated in order for a reader to know which state is meant.

"Required" — the minimal count. Not the number of things that could be said about the state, but the number that cannot be omitted without the specification becoming ambiguous. Redundant components are not components.

"To specify the state" — the criterion is specification, and specification is public. If two analysts working from a declared decomposition and the same episode cannot both arrive at the same state description, the count was not of components; it was of something else.

The index is therefore a *minimal description length in components*, and the essence law raises it to the second power and scales it by a fixed constant. Nothing else enters. `k` is a chosen reference constant, not a measured one, and it is badged as a postulate here as it is everywhere in this series. Its form — a Pythagorean ditone divided by the golden section — is a design decision of the author's, taken in the 2022 source and carried forward unchanged. It sets the absolute pitch of the whole system and nothing else. Every *relationship* in the system is a ratio of squared integers, in which `k` cancels exactly.

5.2 A count is not a rating, and the difference is the whole argument Stevens' 1946 typology of measurement scales is usually taught as a four-item list — nominal, ordinal, interval, ratio — and usually taught badly, as though the four were degrees of quality. They are not degrees of quality. They are statements about *which transformations of the numbers leave the empirical content unchanged* [Stevens, 1946], and that is a precise and checkable property — though what follows from a scale type for the permissibility of a statistic has itself been disputed for as long as the typology has existed [Luce, 2002].

```
SCALE       ADMISSIBLE TRANSFORMATIONS          WHAT IS INVARIANT
nominal     any bijection                       identity and difference of category
```

```
ordinal     any strictly increasing f          order
interval    x -> a x + b,  a > 0               ratios of differences
ratio       x -> a x,      a > 0               ratios of values
absolute    identity only                      the values themselves
```

The operations a formula performs on a quantity are licensed by that quantity's scale type and by nothing else; the formal statement of what must hold before a quantity can be said to exist at all is the representational theory of measurement [Krantz et al., 1971]. Addition needs an interval scale. Division and multiplication need a ratio scale with a true zero. Squaring, which is what the essence law does, needs the same.

Now ask which row a count occupies.

A count has a true zero that is not a convention: zero components is not a low score, it is the absence of the thing. It has equal units that are not a calibration choice: the difference between three components and four is one component, and so is the difference between eleven and twelve, because a component is a component. And it has something stronger still. A count admits no rescaling at all. There is no free multiplier `a`. You cannot count in half-components or in dozens-of-components without changing what "component" means, and a change in what "component" means is a change of decomposition, not of units.

A count is therefore an absolute scale, which is strictly stronger than the ratio scale the essence law requires. The index does not merely satisfy `f = k n^2`; it over-satisfies it.

This is the load-bearing observation of this part of the analysis and it deserves to be stated in the form in which it can be attacked:

> **THE COUNT ARGUMENT.** `n` is defined as a cardinality. Cardinalities have a true zero (the empty set) and a fixed unit (one element), and admit only the identity transformation. Therefore `n` is absolute-scaled, hence ratio-scaled, hence a legitimate argument to a multiplicative law. The scale status of `n` is a consequence of its definition and is not contingent on any empirical finding.

An objector who wishes to defeat this must attack the definition, not the algebra. There is exactly one place to attack it, it is a real place, and §5.4 and the whole of Section 6 are about it.

5.3 CC-R1, restated in full, and why the index is not caught by it

Series B demoted its own keystone. The finding is recorded as CC-R1 and it is binding here:

> **CC-R1 — THE MOMENT EQUATION IS DEMOTED ON SCALE.** For `M = (D/P) * E`, the validation protocol's second gate returns DEMOTE. Because `D` (perspective) and `P` (perception) are not instrumented on ratio scales with true zeros, the *magnitude* of `D/P`, the distortion `Delta = |ln r|`, and the quantitative join from the moment into the storage law are not reportable as measured quantities. The structure survives: the limiting cases, the `D = P` fixed point, and the ordering over-framed / calibrated / under-framed all stand, because each needs only ordinal information. **What is demoted is magnitude, not structure.** The constructive route out is Rasch instrumentation — a logit difference supplies the true zero division requires — at a measured cost of 123 items per term [Rasch, 1960].

Any psychological frequency derived through `D/P` inherits that demotion, and this framework says so wherever such a derivation appears. The demotion is not a caveat bolted onto a result; it is a statement that a particular number is not yet a number.

It is worth seeing what the demotion costs in the register this framework works in, because stated as "the scale type is wrong" it sounds like bookkeeping and it is not. Take a constructed pair of ordinal ratings, `D = 7` and `P = 4` on some 0-to-10 instrument, and apply transformations that are *admissible for ordinal data* — that is, transformations which preserve the order of the ratings and therefore preserve everything the ratings are licensed to report.

```
transform        D'         P'        r = D'/P'      ln r
identity       7.0000     4.0000        1.7500      +0.5596
x^2           49.0000    16.0000        3.0625      +1.1192
sqrt x         2.6458     2.0000        1.3229      +0.2798
ln(1+x)        2.0794     1.6094        1.2920      +0.2562
x^1.5         18.5203     8.0000        2.3150      +0.8394
exp(x/5)       4.0552     2.2255        1.8221      +0.6000

r spans 1.2920 to 3.0625 -- a factor of 2.3703
r > 1 in all six cases -- THE ORDER IS PRESERVED, EXACTLY AS CC-R1 SAYS
```

The direction survives every transformation. The subject is over-framed under all six, and that reading is safe. The magnitude survives none of them.

Now push it through the essence law anyway, as a careless practitioner would, treating `r` as though it were an index:

```
transform     n := r      f = k r^2 (Hz)     folded to 220-440
identity      1.7500          1.4805            379.006
x^2           3.0625          4.5340            290.176
sqrt x        1.3229          0.8460            433.149
ln(1+x)       1.2920          0.8070            413.185
x^1.5         2.3150          2.5909            331.630
exp(x/5)      1.8221          1.6050            410.888

derived frequency spans 0.8070 to 4.5340 Hz = 2988.17 cents = 2.490 octaves
```

A frequency derived through `D/P` is undetermined to within two and a half octaves by transformations that change nothing the underlying data are entitled to say. The spread is not a property of the world; it is a property of the arbitrary set of transformations chosen for the demonstration, and a different set would give a different figure. That is precisely the point. A quantity whose value depends on an arbitrary choice the data cannot constrain is not measured. Any section that publishes such a frequency is publishing a number the reader cannot check, and this framework does not do that.

Set the count beside it.

```
transform     n     f = k n^2 (Hz)
identity      5        12.0856
(there is no other admissible transform of an absolute scale)
```

There is no table because there is no spread. **Five components is five components.** The transformation group of an absolute scale contains one element, so the frequency is determined by the decomposition and by nothing else.

> **THE SIDESTEP, and it is the single most important claim in this part of the analysis.** The Moment equation is demoted because `D/P` is a ratio of two quantities that have not been instrumented on ratio scales. **The psychological essence index is not a ratio of ratings. It is a count of components.** Its ratio-scale status is established by construction rather than by instrumentation, and therefore nothing in CC-R1 touches it. The index does not inherit the demotion. It sidesteps it.

The sidestep is real, and it is worth being explicit about what it does and does not buy, because a claim this convenient invites suspicion and should.

What it buys. The essence law may be applied to the index without a scale objection. `f = k n^2` is a licensed operation on `n`. All the interval structure developed here — ratios of squares, the octave relations, the pitch classes — is licensed. The frequencies may be published as

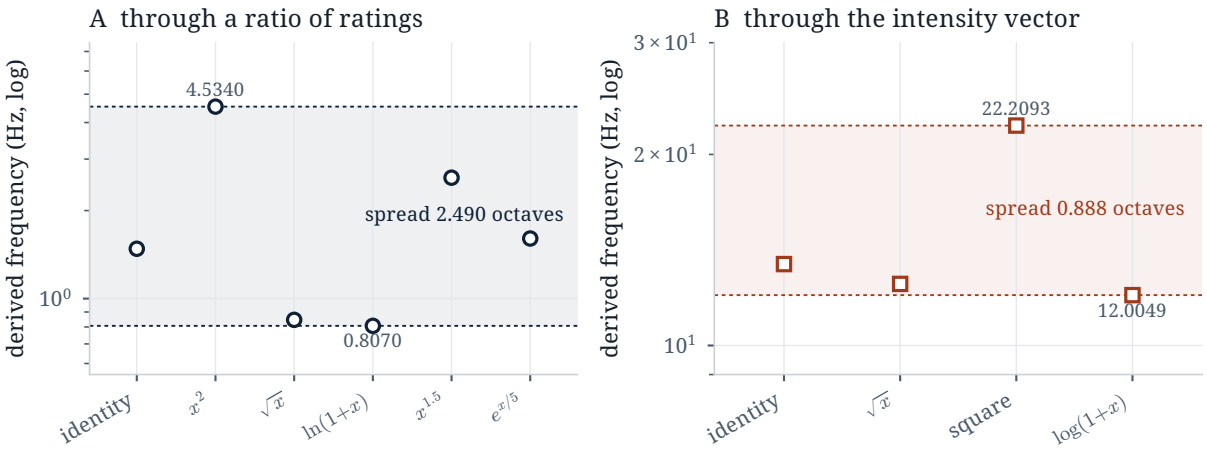


Figure 41.3: A frequency reached through a ratio of ratings is undetermined to within 2.490 octaves by transformations that change nothing the underlying data are entitled to say; the same failure through the intensity vector costs 0.888 octaves. A **count** admits exactly one transformation, the identity, so the corresponding panel would carry one point and no spread. Modelled: the transform sets are chosen for the demonstration and a different set would give a different figure — which is the point, because a quantity whose value depends on an arbitrary choice the data cannot constrain is not measured. Built by `figs/src/CCP-041_fig05.py`.

magnitudes, not merely as an ordering. And the ceiling result `c/eta <= ln N` from Section 4 has a determinate `N` to work with.

What it does not buy. It buys nothing whatever about whether the *decomposition* is right. The count is absolute-scaled *given* a decomposition, exactly as a count of coins is absolute-scaled given an agreement about what a coin is. Change the decomposition and the count changes, and no amount of scale theory arbitrates between two decompositions. **The scale problem has not been solved. It has been moved, from the measurement of a magnitude to the justification of a partition — and it has been moved somewhere it can be argued about in public, which is the entire gain.**

That relocation is not a trick. It is the standard move by which counting escaped the same problem everywhere else it arose. Atomic number is a count of protons and is absolute-scaled given an agreement about what a proton is; the agreement took a century and several revisions, and the scale status of the count was never what was in dispute. Species richness is a count of species and is absolute-scaled given a species concept; the species concept is contested, and ecologists argue about *it*, in print, rather than arguing about whether one richness may be divided by another. The psychological essence index puts psychology in the same position: the argument moves to the decomposition, where it belongs and where it can be settled by agreement statistics rather than by philosophy.

5.4 The honest limit, and it is stated wherever the index is used

> **THE FREQUENCY INHERITS EXACTLY THE DEFENSIBILITY OF THE COUNT. A CONTESTED DECOMPOSITION YIELDS A CONTESTED FREQUENCY.**

That sentence is a standing requirement of this framework. It appears in this section, in Section 6, in The companion paper on octave folding where moods are given frequencies, in the companion paper on trait frequency where traits are, and in the personal-spectrum worksheet of this series. It is not a hedge and it is not modesty. It is a precise statement of where the uncertainty lives, and its practical content is that the uncertainty on a published essence frequency is not a confidence interval; it is a disagreement about a count, and the

correct way to report it is to report the disagreement about the count.

The exchange rate between the two is exact and it is worth having on the page early, because it tells a practitioner what a decomposition dispute actually costs.

```
A MISCOUNT BY ONE: the pitch consequence

n -> n+1     f ratio      cents      octaves     pitch class
 1 -> 2      4.000000    2400.00      2.0000      1 -> 1
 2 -> 3      2.250000    1403.91      1.1699      1 -> 3
 3 -> 4      1.777778     996.09      0.8301      3 -> 1
 4 -> 5      1.562500     772.63      0.6439      1 -> 5
 5 -> 6      1.440000     631.28      0.5261      5 -> 3
 6 -> 7      1.361111     533.74      0.4448      3 -> 7
 7 -> 8      1.306122     462.35      0.3853      7 -> 1
 8 -> 9      1.265625     407.82      0.3399      1 -> 9
 9 -> 10     1.234568     364.81      0.3040      9 -> 5
10 -> 11     1.210000     330.01      0.2750      5 -> 11
11 -> 12     1.190083     301.27      0.2511     11 -> 3
```

Two things follow, and both matter.

A miscount at the bottom of the ladder is enormously more expensive than a miscount at the top. Confusing a one-component state with a two-component state moves the pitch by two full octaves. Confusing an eleven-component state with a twelve-component state moves it by a minor third. The square law compresses; the relative cost of an error falls as `((n+1)/n)^2`, which tends to 1. The elicitation procedure of Section 6 must therefore spend its care where the ladder is steep, which is exactly where practitioners are least careful, because the low states feel obvious.

And a miscount by one almost always changes the pitch class. Adjacent integers have different odd parts unless one is even and the other is its own double, which for adjacent integers happens only at `1 -> 2`. Every other single-step miscount lands the state in a different harmonic class entirely. The pitch class is what survives octave folding and is therefore what a listener hears as identity; The companion paper on octave folding owns that rule and derives it. Here the consequence is enough: a one-component disagreement is not a small disagreement. It is a change of note.

5.5 The framework's validity conditions, and the index against them

A construct is validated by convergent and discriminant evidence, not by internal consistency [Cronbach and Meehl, 1955, Campbell and Fiske, 1959], and the essence law may not be applied to a domain merely because the domain has categories. The framework's foundations state the conditions under which its formal structure may be carried to a target domain at all, and the essence-law specification of Series B [Curzi, 2026k] makes those conditions binding on the essence law and enumerates them. They are restated here because this section is where they are put to work.

```
V1  STATE ENCODING       a map from domain items to normalised vectors, so that every
                         represented state has a determinate location in the model
V2  CATEGORY ASSIGNMENT  a finite partition into idealised categories H_1 ... H_N:
                         exhaustive, exclusive, counted, and genuinely distinct
V3  GEOMETRIC DEPENDENCE interaction modelled through a scalar similarity parameter --
                         an angle, a cosine, an inner-product-derived quantity
V4  ZERO BASELINE        parallel or representationally equivalent states correspond to
                         zero modelled surplus
V5  AFFINE COUNT LAW     adopted as a modelling assumption, not read off the world
V6  NORMALISATION        maximum modelled surplus is finite and depends on N
```

```
and the house addition, because the essence law needs an ORDER and not merely a partition:

V7  DECLARED ORDINAL     the rule assigning index n to category i is published with the
                         frequencies, is applicable by a third party without consulting
                         the author, and includes its tie-break procedure
```

The foundations are explicit about the exclusion that follows: a target domain lacking a stable encoding, a defensible category scheme, or an interpretable angle-like geometry is outside the scope of the model. The essence law inherits that exclusion and this framework enforces it. The three conditions named in that exclusion — explicit state encoding, defensible category partition, interpretable angle-like geometry — are the three the rival orderings of §5.11 through §5.14 are graded against, and no candidate that fails one of them is admitted, however appealing it is on other grounds.

The component-count index answers V7 in a way no other candidate in this section does, and it is worth saying why before the ladder is presented. **The ordinal rule is not a list; it is a procedure.** "Count the components" is applicable by a third party who has never met the author, and it has a determinate answer for a state nobody anticipated. A published list of categories in an order has neither property: it cannot be extended without the author, and a category the author did not think of has no index at all. An ordinal rule that is a procedure rather than an enumeration is the strongest form V7 can take, and it is the principal engineering reason for preferring a count to any of the alternatives.

The Engineering

5.6 The twelve-state ladder **This is the canonical ladder. It is fixed by the framework's founding specification and is used unaltered by every part of the framework of this framework.** Twelve states; `N = 12`; each state adds exactly one component to the one below it, so the state class at index `n` is by construction the class of states requiring `n` components to specify.

n	state class	the component it adds	$f = k \cdot n^2$ (Hz)	folded 220–440	pitch class
1	undifferentiated activation	arousal alone	0.4834	247.514	1
2	valenced activation	+ valence	1.9337	247.514	1
3	directed feeling	+ intentional object	4.3508	278.453	3
4	situated feeling	+ situation	7.7348	247.514	1
5	attributed feeling	+ cause	12.0856	386.741	5
6	temporal feeling	+ projected duration	17.4033	278.453	3
7	reflexive feeling	+ self as object	23.6879	379.006	7
8	social feeling	+ another mind modelled	30.9392	247.514	1
9	normative feeling	+ standard applied	39.1575	313.260	9

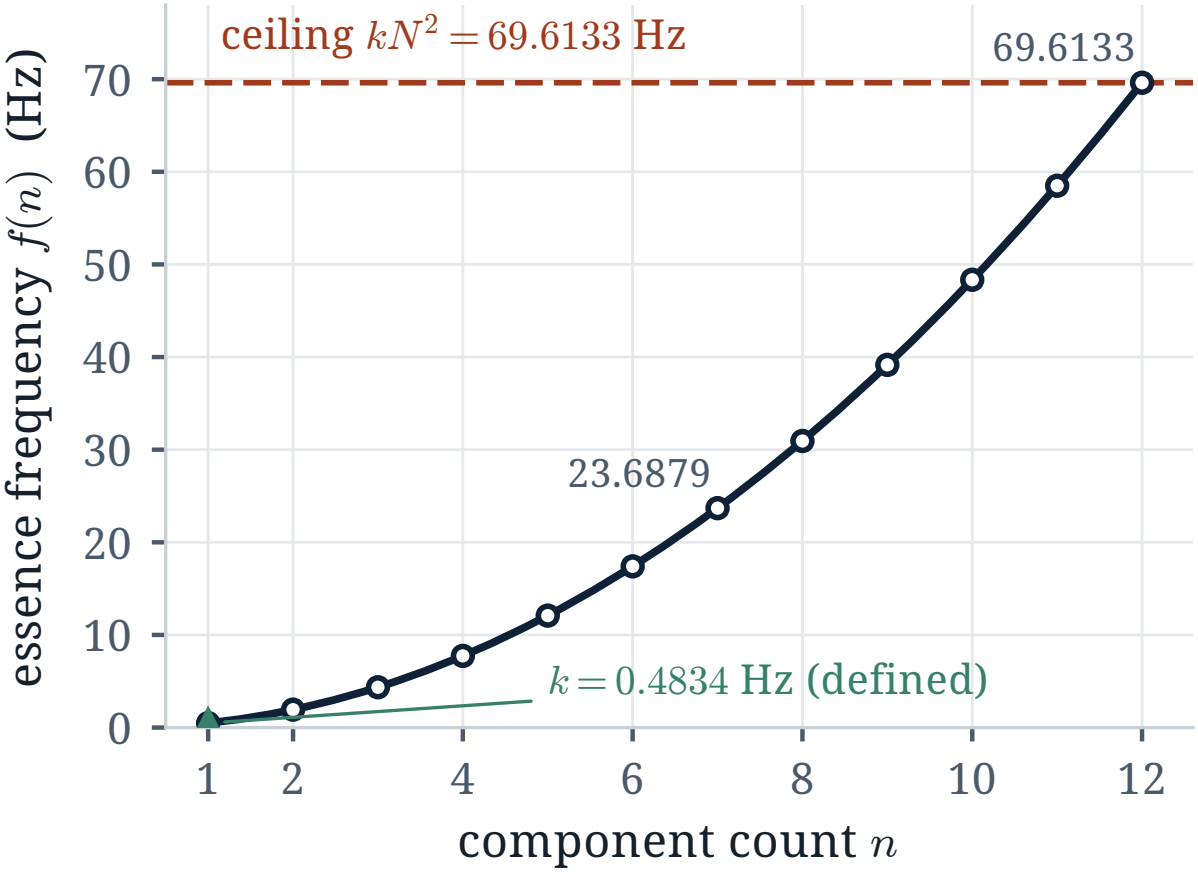


Figure 41.4: The essence law is quadratic in the component count, so the twelve-state ladder spans $12^2 = 144$ in frequency — seven octaves and change — and terminates at the ceiling $kN^2 = 69.6133$ Hz. **Modelled, not measured:** the law $f(n) = kn^2$ is a postulate of this framework (§2.3 P-a) and $k = 0.483425732988414$ is a **defined** constant, exact by construction. Built by `figs/src/CCP-041_fig01.py`.

n	state class	the component it adds	f = k·n² (Hz)	folded 220–440	pitch class
10	narrative feeling	+ placed in a story	48.3426	386.741	5
11	reciprocal feeling	+ the other modelling you	58.4945	233.978	11
12	committed feeling	+ future action bound	69.6133	278.453	3

```
k = 0.483425732988414 Hz        N = 12        ln N = 2.484907
maximum surplus F = ln 12       maximum f = k * 144 = 69.6133 Hz
six distinct pitch classes {1, 3, 5, 7, 9, 11} = ceil(12/2)
```

Every value in the table was computed from `f = k n^2` and folded by powers of two into the band 220–440 Hz. The folding multiplier `j` for each row, since the framework's house rule requires the band and the transposition to be published with every value: `j` = 9, 7, 6, 5, 5, 4, 4, 3, 3, 3, 2, 2 for `n` = 1 through 12 respectively. Octave transposition preserves pitch class exactly; it is the house method and it changes nothing about the reading.

5.7 Reading the ladder The ladder is a specification hierarchy, and each rung earns its index by naming a component that cannot be omitted at that rung and was not present below it. The clearest way to see this is to write the same episode out at each level and watch what has to be added.

n = 1, undifferentiated activation. The specification is *there is arousal*. One component: activation. This is the floor of the ladder, and it is genuinely a state — a person woken at three in the morning by a noise, in the first half second, is in it. Nothing has been evaluated. There is not yet anything the arousal is about.

n = 2, valenced activation. Add valence: the arousal is good or bad. Two components. The half-second has passed; the noise is *bad*. Still nothing about what the noise is. In the affect literature this is roughly the region core-affect accounts describe, and Russell's circumplex [Russell, 1980] is a two-dimensional model of exactly this pair. That correspondence is noted as precedent, not as evidence — a two-dimensional affect model and a two-component specification agree on a count, which is a great deal less than agreeing on a mechanism. It is in fact less than that. Russell's two are two **continuous bipolar axes** recovered from similarity ratings, not two countable components, and the model's principal content — that the affect terms lie on a circle rather than in clusters — has no counterpart anywhere in this ladder. The agreement is on an integer and on nothing else.

n = 3, directed feeling. Add the intentional object: the feeling is *about* something. Three components. The bad arousal is about the noise. This is the rung at which the state becomes a feeling in the ordinary sense of the word, because ordinary usage almost always supplies an object.

n = 4, situated feeling. Add the situation: the state is indexed to a circumstance. Four components. Bad arousal about the noise, in my house, at night. The situation is not the object; the object is what the feeling is directed at, the situation is the frame the object appears in, and a decomposition that conflates them will systematically under-count by one.

n = 5, attributed feeling. Add the cause: an account of why. Five components. The noise is *the back door*, and doors do not open themselves. Appraisal accounts — Smith and Ellsworth [1985] on the dimensions of cognitive appraisal, Lazarus [1991] on primary and secondary appraisal, Scherer's sequential-check model [Scherer, 2001] — are all descriptions of the machinery that produces this component, and **all three are cited for the component's existence and for nothing else.** None of them is cited in support of the index, of its ordering, or of its nesting, and none of the three would supply that support if it were asked for it. The reasons are set out in the third structural remark that closes this section, because they are the strongest objection to the ladder and belong in the body rather than in a footnote.

n = 6, temporal feeling. Add projected duration: how long this will last. Six components. The difference between a fright and a fear is largely this component, and a decomposition that omits it will collapse two states that behave very differently under the stability analysis of the companion paper on trait frequency.

n = 7, reflexive feeling. Add the self as object: the state includes a representation of the one having it. Seven components. *I am frightened, and I am the sort of person who gets frightened.* Tracy and Robins [2004] hold that this class of states — shame, guilt, embarrassment, pride — requires a stable self-representation that the emotions below it do not. That is a claim that a class of states **requires a component** the others lack, which is exactly the form a rung of this ladder needs, and it is the one place in this section where the cited literature supports a rung rather than merely naming its component. The source manuscript names Lewis alongside Tracy and Robins here; that reference is given by surname alone, with no year and no title, it could not be resolved to a specific work, and it has therefore not been guessed at — see L-18.

n = 8, social feeling. Add another mind modelled: the state contains a representation of someone else's mental state. Eight components. *There is a person out there and he wants something.* The theory-of-mind literature descends from Premack and Woodruff [1978]; the component is the modelled other, and it is one component regardless of how elaborate the model is. That paper is cited for the construct's name and provenance only, and two cautions travel with it: it concerns a chimpanzee's inference about a human rather than the composi-

tion of an emotional state, and its experimental claim has been contested since publication.

n = 9, normative feeling. Add a standard applied: the state is evaluated against how things ought to be. Nine components. *This should not be happening; people should not do this.* Note carefully that the standard is a component of the state, not a fact about the world, and that the ladder makes no claim whatever about whether the standard is correct.

n = 10, narrative feeling. Add placement in a story: the episode is an instalment. Ten components. *This is what my life has been like since March.* The narrative component is what converts an event into evidence, and the companion paper on refinement is about the consequences.

n = 11, reciprocal feeling. Add the other modelling you: the state includes a representation of someone else's representation of you. Eleven components. *He knows I am here and he knows I know.* This is second-order and it is a genuinely distinct component from the first-order model at `n = 8`; the pair is the standard recursion of the mentalising literature [Premack and Woodruff, 1978] and the ladder counts each order once.

n = 12, committed feeling. Add future action bound: the state includes a commitment that constrains what is done next. Twelve components. *Whatever else happens, I am not going back to sleep.* Frijda [1986] names the component, and **this rung claims bound commitment and not action readiness as such.** The distinction is forced rather than decorative. Frijda's thesis is that a change in readiness to act is constitutive of emotion *as such*, so on his account the startle at `n = 1` already has action readiness and the component belongs at the floor of the ladder rather than at its summit. Read as a claim about readiness, this rung is contradicted by the work it cites. Read as a claim about a commitment that binds what is done next — an impulse that has been closed over — it is a narrower claim, and it is the only version Frijda leaves standing. **The narrower claim is the one made here; the wider one is withdrawn.**

Three structural remarks about the ladder as a whole, and they bound its claims sharply.

First, **the ladder is a specification hierarchy and not a developmental sequence, a maturity scale, or a ranking of worth.** Nothing in it says that a twelve-component state is better, healthier, more adult, or more valuable than a three-component state. A twelve-component state is *more expensive to specify*, and that is the only claim being made. A person in undifferentiated activation is not failing at anything.

Second, **the ladder is nested by construction.** Each rung contains the rungs below it. This is what makes the index a count rather than a label: the components accumulate, so the number at rung `n` is literally the number of accumulated components. A hierarchy in which the categories were merely ordered — where rung 7 contained something rung 6 did not, but also lacked something rung 6 had — would not license a count at all, and this is where several of the rivals in §5.11 onwards fail.

Third — and this is the remark the section exists to make — the literature named above does not support the index, and one of its entries contradicts it. The rungs were motivated by naming constructs that the affect literature independently describes, and naming is all that establishes. Put at its full strength, the objection is this.

> **THE SCALAR-REDUCTION OBJECTION.** The works cited across §5.7 do not describe emotional states the way this ladder does, and four of them describe them in a way a count cannot represent. Smith and Ellsworth [1985] locate fifteen emotions in a space of appraisal dimensions that vary **simultaneously and independently**; two emotions at the same dimensionality are told apart by the *pattern of values across* those dimensions, which is precisely the information a count throws away. **A count cannot separate anger from guilt; a pattern can, and does.** Scherer [2001] is sharper still. Emotion there is a **process unfolding in time** — a sequence of stimulus evaluation checks driving synchronised change across five organismic subsystems, where the emotional state *is* the synchronisation. An object with an ordering of its own resists reduction to a scalar of any kind, a component count included. Lazarus [1991] individuates emotions by a **core relational theme**: a qualitative content, one per emotion. A theory in which anger and guilt differ by *what*

> is appraised does not license a theory in which they differ by *how many* things are appraised. And Frijda [1986] contradicts the ladder at its **top rung**: if changed action readiness is constitutive of emotion as such, that component sits at or near the floor, and the ordering fails at its own summit.

Three of the four charges are conceded outright, and the fourth is answered only by narrowing the claim. The ordering and the nesting of the rungs are not underwritten by any of the four works and are not offered as though they were. They are a **stipulation of this framework**, fixed by the construction of §5.6 and badged here as such — no postulate in §2.3 covers them, and that omission is recorded in L-18. What this literature supplies is at most a *set* of components attested individually, with the ordering owned here rather than borrowed. Frijda is met by the restriction already stated at `n = 12` — bound commitment as distinct from readiness — and that is a genuine narrowing of the claim rather than a rhetorical one. What survives intact is a single rung: Tracy and Robins [2004] argue that a class of states requires a component the others lack, which is the shape of argument a rung needs. **One rung supported out of twelve is what this literature actually yields, and it is written here rather than left for a reader to find.**

The objection is also not new to this paper. It is Michell's objection (§3.5) arriving in a second domain: a phenomenon its own literature describes by several jointly varying features is replaced by one number, and the arithmetic then proceeds on the number. L-18 concedes the two together, because they are one habit and not two accidents.

5.8 Six classes from twelve states

Because `(2n)^2 = 4n^2` is exactly two octaves above `n^2`, two indices share a pitch class if and only if they have the same odd part. Write `n = 2^a * m` with `m` odd; the class is `m`. For the twelve-state ladder the odd parts are

```
n         1  2  3  4  5  6  7  8  9 10 11 12
odd part  1  1  3  1  5  3  7  1  9  5 11  3
classes   {1, 3, 5, 7, 9, 11}  --  6 classes = ceil(12/2)
```

Twelve states yield six pitch classes. That is the available harmonic vocabulary and it is a hard structural limit, not a stylistic choice. The companion paper on octave folding owns the folding rule and the `ceil(N/2)` theorem and derives both; they are quoted here because the ladder cannot be read without them.

Within the 220–440 Hz band the six classes sit as follows, measured in cents from class 1:

```
class  1     247.514 Hz        0.0000 cents
class  3     278.453 Hz     +203.9100 cents
class  9     313.260 Hz     +407.8200 cents
class  7     379.006 Hz     +737.6518 cents
class  5     386.741 Hz     +772.6274 cents
class 11     233.978 Hz      -97.3641 cents
```

Class 11 lands below class 1 in this band because folding is into a fixed window; the negative figure is an artefact of the window and not of the class. Note that classes 7 and 5 are the closest pair in the set, and that the size of that minimum separation is a structural property of `N = 12` rather than of this particular ladder. It is stated once, with the regime it belongs to, and is not re-derived elsewhere.

The immediate consequence for this section is the one that governs the failure modes: half the ladder's distinctions are harmonically invisible. States 1, 2, 4 and 8 are all class 1. They are four different states, they have four different frequencies, and they sound like the same note played at four different heights. A reading that reports only the class has discarded the difference between undifferentiated activation and social feeling — which is a large difference. **Report the index and the class. Never the class alone.**

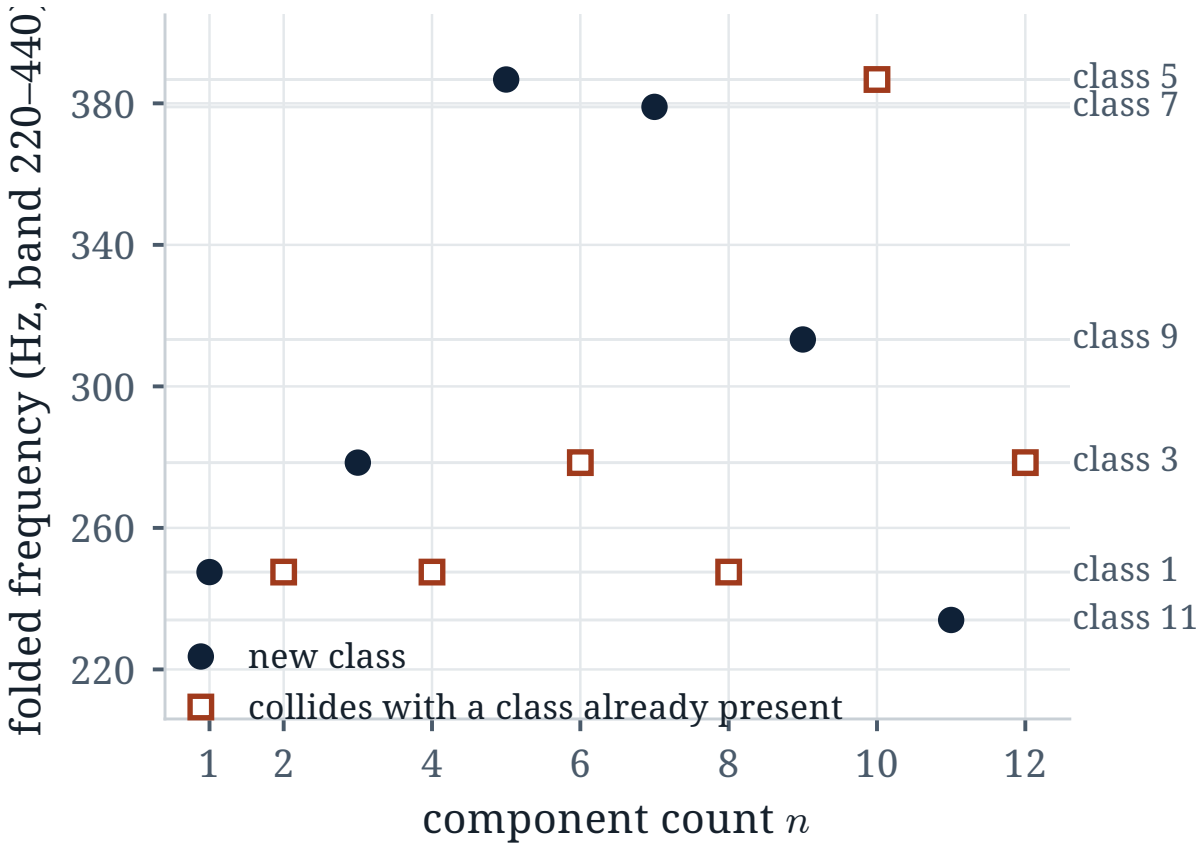


Figure 41.5: Twelve states collapse onto six folded positions: six of the twelve indices land on a pitch class already occupied, because two indices fold together exactly when they share an odd part. Half the ladder's distinctions are therefore invisible after folding, which is why a reading must report the index and not the class alone. **Modelled:** the fold is division by a power of two into a 220–440 Hz band, a reporting convention of this framework; its extension beyond pitch is a modelling choice with no external warrant (§2.4 M-a). Built by `figs/src/CCP-041_fig04.py`.

5.9 What the ladder is not This section exists because the ladder is the object in this framework most likely to be misused, and because a formalism that looks like a clinical instrument will be treated as one unless it says otherwise in the plainest available language.

The twelve states are not diagnoses. They are not disorders, syndromes, categories of illness, or levels of severity. No index on this ladder corresponds to any diagnostic category in any classification system, no index indicates the presence or absence of any condition, and no reading produced by any procedure in this framework is an assessment of any person's mental health.

No part of the framework of this framework makes a clinical, medical, diagnostic, or therapeutic claim, and none is implied by any number in it. The readings are readings of a formal construction. A person who wishes to know something about their health should ask someone qualified to tell them, and nothing here is a substitute for that.

The ladder is not an instrument for classifying persons. A person is not at an index. States are at indices; a person holds several states at once, which is exactly what the chord analysis of this series is about, and the object that characterises a person in this framework is a chord and not a note.

The ladder is not a hierarchy of value. Higher index means more components, means more expensive to specify, means higher pitch. It does not mean better, more evolved, healthier, more mature, or more real.

And the frequencies are essence frequencies. No claim is made that any person, tissue, molecule, or neural population physically resonates at 23.6879 Hz or at any other value in the table. The essence law returns the resonance of an essence — a relation among categories — and that distinction is absolute in this series and is restated wherever the law is used.

5.10 What the ladder costs to get wrong Combine §5.4's miscount table with §5.8's class structure and the practical exposure becomes visible.

A single-step miscount costs between 301 and 2400 cents of pitch and, except at `1 -> 2`, changes the pitch class. A two-step miscount at the bottom of the ladder — mistaking a one-component state for a three-component one, which is the single easiest error to make because ordinary language supplies an object whether or not the state has one — costs `9/1` in frequency, three octaves and a whole tone, and moves from class 1 to class 3.

The exposure is therefore concentrated, and it is concentrated in a knowable place. It is not spread evenly across the ladder and it is not a matter of general imprecision. It lives at rungs 1 through 5, where the square law is steepest and where the components are easiest to conflate. Section 6's decomposition rules are weighted accordingly, and its worked example is deliberately drawn from that region.

The Specification

The section now does the thing it exists to do. Four rival ordering principles are stated at their strongest, graded against the framework's three exclusion criteria — explicit state encoding, defensible category partition, interpretable angle-like geometry — and against the house condition V7, and each is given a verdict that is published whether or not it flatters the framework.

A note on method before the grading starts. Each rival is instantiated below with a **constructed** assignment of indices, because a principle cannot be graded in the abstract. Those constructions are labelled as constructed and are not presented as anyone's published ordering. The point of constructing them is not to show that a particular arrangement is wrong; **it is to show that the principle underdetermines the arrangement, which is a property of the principle and is what the grading is about.**

5.11 Rival 1 — AROUSAL. Verdict: REJECTED. **The principle.** Order the states by the typical arousal magnitude of the state. Arousal is the best-attested single dimension in the affect literature, it has physiological correlates that can be measured without asking anyone anything, and it is continuous, which feels like an advantage.

The case for it. Arousal is the one affective quantity with a genuine claim to instrumentation. Skin conductance, heart-rate variability and pupillometry all index it, imperfectly and with well documented confounds, but they index it. If the goal were a psychological quantity that could be measured by an apparatus rather than elicited from a person, arousal would be the obvious candidate and no other candidate is close.

Grading against V1, state encoding. Passes, weakly. A state does have an arousal level, and that level locates it on a line. Constructing an interval scale without asking respondents to use numbers is possible [Thurstone, 1927]; it still yields no partition, which is what V2 requires.

Grading against V2, the partition. Fails. Arousal does not partition the states; it *scores* them, and many states get the same score. Take a constructed set of arousal ratings for the twelve state classes on a 0-to-10 scale at one decimal place — the resolution a careful rater can actually sustain:

```
state    1    2    3    4    5    6    7    8    9   10   11   12
arousal 7.0  7.0  6.0  5.0  5.0  4.0  5.0  6.0  6.0  4.0  6.0  5.0

tie structure: 7.0 x2, 6.0 x4, 5.0 x4, 4.0 x2
states involved in at least one tie: 12 of 12
orderings consistent with these ratings: 2 x 24 x 24 x 2 = 2,304
undetermined ordering: log2(2304) = 11.170 bits
```

Every state is tied with at least one other, and the ratings leave 11.17 bits of the ordering undetermined. The principle therefore does not supply an ordinal at all; it supplies an ordinal *with a tie-break to be added later*, and the tie-break carries all the information the essence law consumes. Under V7, a rule whose tie-break is unspecified is not a declared ordinal rule.

Grading against V3, angle-like geometry. Fails on a subtler point that is the decisive one. The essence law is a square law, and the pitch class depends on the odd part of the index, which is a number-theoretic property and not a continuous function of position. Arousal is a continuous quantity. **Mapping a continuous quantity onto integer positions and then taking the odd part of the position is not an interpretable geometry; it is a rounding operation followed by a number-theoretic one, and small changes in the underlying quantity produce arbitrary changes in the output.**

Here is that failure measured. Three constructed analysts, each producing a defensible arousal ordering of the same twelve state classes:

```
analyst A order (low arousal to high):  6 12 10  9  4  3  2  1  5  7  8 11
analyst B order:                       12 10  6  4  9  3  1  2  7  5 11  8
analyst C order:                       10 12  9  6  3  4  2  5  1  8  7 11

Spearman rho(A,B) = +0.9510
Spearman rho(A,C) = +0.9301
Spearman rho(B,C) = +0.8951
```

The three analysts agree with each other almost completely. By any conventional standard of rank agreement this is a well-behaved ordering. Now read the frequencies it produces:

```
state                            indices     classes      pitch spread
 1 undifferentiated activation    [ 8, 7, 9]  {1,7,9}       870.17 cents
 2 valenced activation            [ 7, 8, 7]  {1,7}         462.35 cents
 3 directed feeling               [ 6, 6, 5]  {3,5}         631.28 cents
```

```
 4 situated feeling              [ 5, 4, 6]  {1,3,5}       1403.91 cents
 5 attributed feeling            [ 9,10, 8]  {1,5,9}        772.63 cents
 6 temporal feeling              [ 1, 3, 4]  {1,3}         4800.00 cents
 7 reflexive feeling             [10, 9,11]  {5,9,11}       694.82 cents
 8 social feeling                [11,12,10]  {3,5,11}       631.28 cents
 9 normative feeling             [ 4, 5, 3]  {1,3,5}       1768.72 cents
10 narrative feeling             [ 3, 2, 1]  {1,3}         3803.91 cents
11 reciprocal feeling            [12,11,12]  {3,11}         301.27 cents
12 committed feeling             [ 2, 1, 2]  {1}           2400.00 cents

WORST: temporal feeling, 4800.00 cents = 4.000 octaves of disagreement
```

Three analysts who agree with each other at rho = 0.90 and above produce, for one state, four octaves of pitch disagreement. For reflexive feeling they produce classes 5, 9 and 11 — three different notes. This is the finding that decides the grading and it generalises well beyond arousal: rank agreement does not imply pitch agreement, because the essence law is a square law on the index and the pitch class is a number-theoretic function of it. A single position's difference near the bottom of the ladder is worth more than a five-position difference near the top.

Verdict: REJECTED. Arousal fails V2 (does not partition; 11.17 bits of ordering undetermined), fails V7 (no declared tie-break, and none is available from the principle), and fails V3 (the composition of a continuous quantity with an integer-rounding step and an odd-part extraction is not an interpretable geometry). It is rejected as an essence index. It is not rejected as a psychological quantity, and the companion paper on trait frequency uses arousal-like variation in a different role, where the square law is not applied to it.

5.12 Rival 2 — HIERARCHICAL NEED. Verdict: REJECTED.

The principle. Order the states by the level of need they express, following the familiar prepotency hierarchy — physiological, safety, belonging, esteem, self-actualisation (Maslow, 1943).

The case for it. It is the most widely recognised ordering in all of popular psychology, it is genuinely ordinal in intent, and it comes with a mechanism: lower needs are said to be prepotent, so the ordering is supposed to be doing causal work rather than merely filing.

Grading against V1, state encoding. Fails, and fails at the first step. **A need is not a state.** The hierarchy orders *what a person requires*; the ladder must order *what a person is in*. There is no map from an episode of feeling to a level of the hierarchy that is determinate: a moment of fear can be about safety, about belonging, or about esteem, and often about all three, and the hierarchy supplies no rule for choosing. V1 requires that every represented state have a determinate location in the model, and this candidate does not provide one.

Grading against V2, the partition. Fails on arithmetic. The hierarchy supplies five levels for twelve state classes:

```
levels available           5
state classes to index    12
states per level         2.4  ->  ties by construction, in every assignment
distinct pitch classes from indices 1..5:  {1, 3, 5}  =  3
```

Five indices yield three pitch classes. The twelve-state ladder yields six. The harmonic vocabulary is halved before any reading is taken, and the halving is a consequence of the number of levels rather than of anything about the psyche.

The ceiling consequence, and it is severe. Section 4 established that maximum surplus is `ln N` and therefore that the maximum sustainable cost-to-conversion ratio is `c/eta <= ln N`. Compare:

```
twelve-state ladder      N = 12    ln N = 2.484907    f_max = k*144 = 69.6133 Hz
```

```
five-level hierarchy     N =  5    ln N = 1.609438    f_max = k*25  = 12.0856 Hz

ceiling on c/eta collapses from 2.4849 to 1.6094 -- a 35.2% loss of headroom
```

Adopting a five-level index would declare that 35.2% of the sustainable region does not exist, and would do so purely as an artefact of the category count. Section 4's fourth lever — add genuinely distinct categories — exists precisely because `N` is load-bearing, and a candidate index that gratuitously suppresses `N` is throwing away the lever.

Grading against the empirical record. The strict prepotency claim has not held up. Wahba and Bridwell's review (1976) found little support for the ordered-emergence prediction, and later treatments have been no kinder to it. This matters here only because prepotency was the candidate's one claim to being more than a list: without it, the hierarchy is an enumeration in an order somebody chose, which under V7 is a weak ordinal whose weakness must travel with every number derived from it.

Verdict: REJECTED. It fails V1 outright — needs are not states — fails V2 by construction, and costs 35.2% of the sustainable region for nothing in return. This is the rejection the manifest requires be published, and it is published without qualification. The hierarchy is a serviceable way to talk about motivation. It is not an essence index and it must not be used as one.

5.13 Rival 3 — DEVELOPMENTAL STAGE. Verdict: ADMISSIBLE AS AN ORDINAL, NOT ADOPTED.

The principle. Order by the developmental stage at which the state becomes available — the stage-theoretic tradition running through Piaget on cognitive stages, Kohlberg on moral reasoning, and Loevinger (1976) on ego development.

The case for it, and it is the strongest case any rival makes. Stage theories are genuinely ordinal by construction, they specify what each stage adds relative to the one below, they are nested in the way §5.7 requires, and Loevinger's programme in particular was built with a scoring manual designed for inter-rater agreement. This candidate is not a category error. It is a serious alternative and it is graded as one.

Grading against V1, state encoding. Fails at the level required, and the failure is specific. A stage is a property of a *person*, not of a *state*. The ladder must index the state; if the index is the stage of the person holding it, then every state a given person holds receives the same index, the person's whole spectrum collapses to a single note, and the chord construction of the chord analysis of this series — the object that makes this argument extended rather than compressed — cannot be formed at all. **A person-level index cannot produce a within-person chord, and the within-person chord is what this framework measures.**

Grading against V2, the partition. Passes. Stage schemes partition, they are counted, and their categories are argued to be distinct.

Grading against V3, geometry. Passes conditionally. Stages have a defensible nesting, so the distance between adjacent stages has an interpretation.

Grading against V7 and against CC-R1. Here is the decisive point and it is the one that keeps this candidate out of the framework rather than merely down the list. **A stage assignment is a rating produced by a trained judge against a manual. It is ordinal. It has no true zero — stage zero is not "no person" — and it has no fixed unit, because the interval between stage 3 and stage 4 is not constructed to equal the interval between stage 6 and stage 7.** Push it through `f = k n^2` and the reading are multiplying an ordinal, which is exactly the operation CC-R1 demotes. A frequency derived from a developmental stage index inherits the demotion in full and may be published only as an ordering, never as a magnitude.

The constructive route out is available and is the same one CC-R1 names: Rasch instrumentation of the stage scale would supply the true zero and the equal units, at the measured cost of 123 items per term. That is a real research programme and this framework does not pretend it is impossible. It has not been done for any stage scheme in a form this framework could use.

Verdict: ADMISSIBLE AS AN ORDINAL, NOT ADOPTED. Frequencies derived from it carry CC-R1's demotion. It is not adopted here because it indexes persons rather than states and therefore cannot produce the chord. It is the one rival in this section that could be adopted later if someone did the instrumentation, and the companion paper on trait frequency notes where it would slot in — as a trait-scope index, not a state-scope one.

5.14 Rival 4 — VALENCE-MAGNITUDE. Verdict: REJECTED.

The principle. Order by the magnitude of the state's valence: order by how strongly good or bad it is, taking absolute value so that the index runs from neutral upward.

The case for it. Valence is the second of the two core-affect dimensions and is at least as well attested as arousal. Taking the absolute value gives something that looks like it has a true zero — neutral — which appears to answer the scale objection directly. That appearance is the reason this candidate has to be dealt with explicitly rather than waved off.

Grading against V4, the zero baseline. This is where it dies, and it dies by an argument this series has already proved in another register. Take a constructed valence assignment on `[-1, +1]`:

```
state    1     2     3     4     5     6     7     8     9    10    11    12
valence 0.00 -0.60 -0.70 -0.50 +0.60 -0.40 +0.50 +0.40 -0.30 +0.30 +0.70 +0.65

collapsed pairs under |valence| -- opposite sign, identical index:
    state  2 (-0.60)  and  state  5 (+0.60)
    state  3 (-0.70)  and  state 11 (+0.70)
    state  4 (-0.50)  and  state  7 (+0.50)
    state  6 (-0.40)  and  state  8 (+0.40)
    state  9 (-0.30)  and  state 10 (+0.30)
```

Ten of the twelve states collapse into five pairs. Each pair receives one index, one frequency, and one pitch class. Directed feeling and reciprocal feeling — three components and eleven components, two full rungs apart in every structural respect — become the same note.

This is CC-R8 in the essence register and the framework has already proved the general form of it. CC-R8 established that because `u = 1 - (x.y)^2` depends only on `|rho|`, perfect opposition returns the same `u = 0` as perfect agreement: the contrarian and the sycophant carry identical information, which is none. An index built on absolute valence commits the same error one level down. It declares that the sign of the evaluation carries no information, which is a substantive empirical claim, is almost certainly false, and is being smuggled in as a modelling convenience.

Keeping the sign does not save the candidate. A signed index runs from negative to positive, `f = k n^2` is even, and `f(-n) = f(n)` exactly — so the collapse returns in a different costume. And a signed index has no cardinality reading at all: there is no such thing as negative-four components.

Grading against V2 and V7. Fails for the same reason arousal does. Valence magnitude is a continuous score with heavy ties at any usable resolution, and the tie-break is unspecified.

Verdict: REJECTED. It fails V4 by collapsing opposites, it reproduces CC-R8's error at the level of the index, and it fails V2 and V7 alongside arousal. Valence is real, it matters, and the companion paper on dissonance uses it in the analysis of the divided self. It is not an essence index.

5.15 The grading table

Candidate	V1 encoding	V2 partition	V3 geometry	V4 zero	V7 ordinal rule	Scale under CC-R1	Verdict
Arousal	weak pass	**FAIL** — all 12 tied, 11.17 bits free	**FAIL** — continuous into odd-part	pass	**FAIL** — no tie-break	ordinal	**REJECTED**
Hierarchical need	**FAIL** — needs are not states	**FAIL** — 5 levels, 12 states	weak	pass	weak (an enumeration)	ordinal	**REJECTED**
Developmental stage	**FAIL** — indexes persons	pass	conditional pass	**FAIL** — no true zero	pass	**ordinal — CC-R1 applies**	**admissible, not adopted**
Valence magnitude	weak pass	**FAIL** — ties	**FAIL**	**FAIL** — collapses opposites (CC-R8)	**FAIL**	ordinal	**REJECTED**
Component count	**pass**	**pass** — nested, exhaustive, exclusive	**pass**	**pass** — n = 0 is no state	**pass** — a procedure, not a list	**absolute, hence ratio**	**ADOPTED**

Three rejected in print, one admitted as an ordinal and left unadopted, one adopted. The adopted candidate's single advantage over the others is not that it is more psychologically insightful. It is that it is a count.

5.16 Publication requirements for any essence index Any index used in this framework, or by any reader applying the framework, must publish the following. These are requirements, not recommendations, and The checklist is the worksheet of §6.10.

```
CC-I1  THE DECOMPOSITION.  What counts as one component, stated as a rule a third party can
                       apply to a state the author never considered.
CC-I2  THE COUNT RULE.     How components are counted, including the treatment of redundant,
                       nested, and conjoined components.
CC-I3  N.                  The category count, stated, with the ceiling ln N it implies.
CC-I4  THE TIE-BREAK.      What happens when two states receive the same count. (For the
                       component-count index: nothing happens. They are the same state
                       class and they share a frequency. That is the intended behaviour.)
CC-I5  THE BAND AND j.     The audible band and the octave transposition for every published
                       frequency.
CC-I6  THE PERMUTATION TEST. Swap two adjacent indices, recompute, and report the change in
                       cents and in pitch class. If a plausible re-ordering moves a class
                       by more than half an octave, the ordering is load-bearing and must
                       be defended before any frequency is published.
CC-I7  THE AGREEMENT FIGURE. The inter-rater agreement on the count, reported on the count
                       and on the pitch class separately. Section 6 specifies the protocol.
```

CC-I6 costs ten minutes and is the item most often skipped. For the twelve-state ladder, run on every adjacent pair:

```
swap  1<->2:  same odd part (1) -- harmonically invisible, a pure two-octave move
swap  2<->3:  class 1 -> 3          203.91 cents
swap  3<->4:  class 3 -> 1          203.91 cents
swap  4<->5:  class 1 -> 5          772.63 cents
swap  5<->6:  class 5 -> 3          568.72 cents
swap  6<->7:  class 3 -> 7          533.74 cents
swap  7<->8:  class 7 -> 1          737.65 cents
swap  8<->9:  class 1 -> 9          407.82 cents
swap  9<->10: class 9 -> 5          364.81 cents
swap 10<->11: class 5 -> 11         869.99 cents
swap 11<->12: class 11 -> 3         301.27 cents
```

Nine of the eleven adjacent swaps move a pitch class by more than a quarter of an octave, and three move it by more than half. By CC-I6's own standard the twelve-state ordering is

load-bearing and must be defended. **It is defended — by §5.7's nesting argument, which is the claim that the order is not chosen but forced: rung `n` contains rung `n−1` plus exactly one component, so the sequence is the only one consistent with the decomposition.** That defence is the whole warrant for the ladder, and if it fails the ladder fails with it. Which is the correct exposure for an account to have, and it is stated here so that a reader can attack it in the right place.

The Failure Modes

5.17 Agreeing on the order while disagreeing on the note §5.11 measured it: three orderings at Spearman rho above 0.89 produced four octaves of disagreement on one state. **The general lesson is that rank-correlation statistics are the wrong instrument for assessing an essence index, and reporting one is worse than reporting nothing because it manufactures false confidence.**

The mechanism is exact. Two effects compound. The square law means a fixed rank difference is worth more at the bottom of the ladder than at the top — one position at `n = 1` is 2400 cents, one position at `n = 11` is 301 cents, an eightfold difference in consequence for an identical difference in rank. And the odd-part rule means the pitch class is not a monotone function of the index at all: 7 and 8 are adjacent and land in classes 7 and 1, while 3 and 12 are nine apart and land in the same class. Rank distance and pitch distance are not the same metric and are not monotonically related.

The correct statistic is agreement on the integer and agreement on the class, reported separately. Section 6 §6.12 gives the protocol and §6.17 shows a constructed case in which a weighted kappa of 0.9838 coexists with pitch-class disagreement on more than a quarter of episodes.

5.18 Smuggling a rating in as a count The most common way to break the count argument is to keep the word "count" and stop counting.

It looks like this. An analyst has a state, feels that it is *quite* complex, and writes `n = 7`. No components were enumerated. The 7 is a judgment of complexity dressed as a cardinality, and a judgment of complexity is an ordinal rating with all of CC-R1's problems intact. **The frequency then has the scale status of the thing that actually produced it, which is the rating, and not the scale status of the notation it was written in.**

The test is simple and it is procedural: an analyst who claims `n = 7` must be able to list the seven components. If the list cannot be produced, the number is a rating. Section 6's worksheet requires the list in writing, before the number, precisely so that this failure is caught at the point of elicitation rather than at the point of publication.

A subtler version is worth naming separately. An analyst enumerates components honestly but chooses the *granularity* of the enumeration to reach a count they already had in mind. Splitting one component into two, or merging two into one, is always locally defensible, and a determined analyst can reach almost any count on almost any state. This is not caught by requiring a list. It is caught only by fixing the decomposition rules before the state is examined, and by inter-rater agreement, which is why Section 6 treats both as mandatory rather than as good practice.

5.19 Counting representations instead of components A component is something the specification requires. A representation is something the analyst wrote down. They diverge in both directions.

Over-counting. The state description mentions the back door, the hour, the neighbour, the dog, and the fact that the porch light is broken. That is five nouns. It is not five components. The hour and the broken light are both part of the situation, which is one component. **Nouns**

are not components, and a decomposition that counts nouns will systematically over-count and will do so more for verbose subjects than for terse ones, which is a bias correlated with something the analyst has no business letting into the reading.

Under-counting. The state description says "I was frightened of him," which is four words and at least four components: activation, valence, an intentional object, and a modelled other whose intentions matter. Ordinary language compresses aggressively, and a decomposition that counts what was said will under-count fluent speakers.

The rule that fixes both: count what the specification requires, not what the description contains. The operational form is the substitution test in Section 6 §6.6 — a component is required if removing it makes the description consistent with a state the subject would say they were not in.

5.20 Assuming that more components is better

The ladder ascends and the frequencies rise, and Section 4 established that higher pitch means more available order. The inference that a person should therefore aim for high-index states is available, natural, and wrong.

Three reasons, and they are independent.

First, the index is a property of a state, not an achievement of a person. A twelve-component state is not earned; it is simply expensive to specify. A person in a twelve-component state of committed dread is not doing better than a person in a three-component state of ordinary annoyance.

Second, the frequency that matters for sustainability is the relational frequency `f_rel`, not the state frequency `f = k n^2`. The companion paper on the dyad interval owns `f_rel = k(1 + (N−1)u)^2` and it is driven by `u`, the non-parallelism, and not by any individual index. A person holding a chord of high-index states that are all pairwise near-parallel has a high spectrum and a low surplus — which the chord analysis of this series shows is a recognisable and unenviable condition.

Third, components have costs. Every component in a specification is something that must be maintained, and Section 4's `c` is where maintenance lives. A specification that requires twelve components to state requires twelve components to hold, and the cost side of the ledger does not care that the pitch went up.

The lever Section 4 names is "add genuinely distinct categories," and the load-bearing word is distinct, not more.

5.21 Reading the ladder as a diagnosis

Named in §5.9 and repeated here because it belongs in the failure catalogue and because repetition is the only defence an account has against a misuse it cannot supervise.

The twelve states are not diagnoses. An index is not a severity. A chord is not a profile of pathology. No reading in this framework assesses anyone's mental health, and no procedure in it should be administered to anyone as though it did. A practitioner who finds themselves reporting a number to a person about that person's mind has left the scope of this framework, whatever the number is.

The structural reason, stated once so that it is not merely a prohibition: **the ladder classifies specifications, not people.** The claim "this state requires seven components to specify" is a claim about a description. It becomes a claim about a person only through an inference the formalism does not license and this framework does not make.

5.22 Publishing a frequency without publishing the decomposition

The worst of the failure modes, because it is invisible.

A frequency published without its decomposition is uncheckable. The reader cannot recount, cannot identify the disputed component, cannot run the permutation test, and cannot form a view about whether the number means anything. **And because the number is quoted to four decimal places, it will be believed.** Spurious precision is not a stylistic vice here; it is the

specific mechanism by which an unfalsifiable claim acquires the appearance of a measured one.

The standing rule: the decomposition travels with the frequency, always, in the same table or on the same page. Where it will not fit, the frequency does not get published either. This is CC-I1 and CC-I2 of §5.16 and it is the reason those two items come first.

Worked Example

Everything in this section is CONSTRUCTED. It is not an observation of any person, no data were collected, and nothing in it is a case report. It is a demonstration of a procedure on an invented episode, which is what all worked cases in this framework are.

5.23 A constructed state, indexed

The episode, stated as the subject might state it:

> *I keep thinking about how she looked at me in the meeting. She knew what she was doing, and she knew I'd notice. This is what it's been like since I took the role.*

Section 6 owns the elicitation procedure that turns an utterance into a component list; here the decomposition is performed directly so that the indexing can be shown.

Step 1 — enumerate candidate components.

```
c1  activation                 there is arousal; the subject is keyed up
c2  valence                    it is unpleasant
c3  intentional object         the look, in the meeting
c4  situation                  the meeting, the role, the workplace
c5  cause                      she did it deliberately
c6  projected duration         "keeps" -- it is ongoing, not a flash
c7  self as object             "I'd notice" -- the subject appears in their own account
c8  another mind modelled      "she knew what she was doing"
c9  the other modelling you    "she knew I'd notice" -- second order
c10 placed in a story          "what it's been like since I took the role"
```

Step 2 — apply the substitution test. For each candidate, remove it and ask whether the residual description is consistent with a state the subject would deny being in.

Removing c1 through c10 each fails the test — the residual is a different state in every case — with one item requiring care. Is c9 (*she knew I'd notice*) genuinely distinct from c8 (*she knew what she was doing*)? Remove c9 and the description becomes: she acted deliberately, and I noticed. That is a coherent state and it is one the subject would say they were not in, because the whole weight of the episode is that the noticing was *anticipated*. **c9 survives. It is a second-order component and the ladder counts each order once.**

Step 3 — check for redundancy. No candidate is entailed by another. c7 and c9 both involve the self, but c7 is the subject appearing in their own account and c9 is the subject appearing in the modelled other's account; removing either leaves the other standing. **They are distinguishable and they count twice.**

Step 4 — check for a normative component. The description does not apply a standard. Nothing says this ought not to happen; the account is of what happened and what it means, not of what was owed. **Component "standard applied" is absent.** This is the kind of judgment that must be made explicitly and recorded, because it is exactly where two analysts will differ.

Step 5 — check for a committed component. Nothing binds a future action. There is no *and so I am going to*. **Absent.**

Step 6 — count.

```
components required:  c1 c2 c3 c4 c5 c6 c7 c8 c9 c10  =  10
n = 10
```

Step 7 — read.

```
n = 10        state class: narrative feeling
f = k * 100 = 48.3426 Hz
folded into 220–440 Hz with j = 3:  48.3426 * 8 = 386.741 Hz
odd part of 10 = 5   ->  pitch class 5
```

Step 8 — publish the decomposition alongside the frequency, which is what the ten-line list in Step 1 is for. A reader who disagrees now has something to disagree with.

Step 9 — the standing statement. This reading says that the state as described requires ten components to specify, and that ten components corresponds to 48.3426 Hz under the essence law with the declared constant. **It says nothing about the subject's health, character, or prospects, it is not a diagnosis, and 48.3426 Hz is an essence frequency and not a physical resonance of anything.**

5.24 The same state under a rival, and the difference priced Take the developmental-stage rival of §5.13 — the strongest of the four — and index the same episode by it. A constructed stage assignment, using a five-level scheme, places the subject's reasoning at level 4.

```
component-count index      n = 10    f = 48.3426 Hz    class 5
stage index                n =  4    f =  7.7348 Hz    class 1

ratio    (10/4)^2 = 6.250000
cents    1200 * log2(6.25) = 3172.63 cents = 2.6438 octaves
classes  5 versus 1 -- different notes entirely
```

Two and a half octaves and a change of pitch class, from one choice of indexing principle. No computation is in error. The two numbers are answers to two different questions, and only one of them is the question `f = k n^2` is licensed to receive — because only one of them is a count.

And note what the rival's number carries with it. The stage index is an ordinal rating. Under CC-R1 its magnitude is not reportable, so 7.7348 Hz is not a measured frequency; it is an ordinal position written in Hz. It may be used to say that this state sits lower than another state on the same scheme. It may not be divided, squared, or folded, and the octave-deficit diagnostic of Section 4 cannot consume it.

That is the whole of the count argument compressed into one comparison. The count survives the operations the framework performs on it. The rating does not.

5.25 What the reading licenses The constructed reading in §5.23 licenses exactly this much.

It licenses the frequency 48.3426 Hz as an essence frequency of the described state under the declared decomposition, and it licenses that number's use in every downstream construction of this framework: the mood spectrum of the companion paper on octave folding, the chord of the chord analysis of this series, the interval arithmetic of the companion paper on dissonance, and the octave-deficit diagnostic of Section 4 wherever a relational `u` has been measured to go with it.

It licenses the claim that if the decomposition is accepted, the frequency follows with no further assumption, because `f = k n^2` on an absolute-scaled `n` is a licensed operation.

It does not license the claim that the decomposition is the right one. Another analyst may count nine components by merging c7 and c9, or eleven by splitting the situation from the

role. Nine gives 39.1575 Hz and class 9; eleven gives 58.4945 Hz and class 11. Three defensible decompositions, three different notes. Section 6 exists because that spread is not acceptable and because reducing it is a procedural problem with a procedural answer.

And it does not license any claim about a person. **The state was constructed, the subject does not exist, and if the subject did exist the reading would still be a reading of a description.**

What this section fixes

Later papers in the corpus cite this construction and do not redefine it.

```
THE INDEX          n := the number of distinguishable components required to specify
                   the state.  A COUNT: true zero, fixed unit, absolute scale.

THE LAW            f(n) = k n^2,  k = 0.483425732988414 Hz  [k badged as a postulate]

THE SIDESTEP       Because n is a count and not a ratio of ratings, its ratio-scale status
                   is established by construction. CC-R1's demotion of D/P does not reach it.
                   The index sidesteps the demotion; it does not inherit it.

THE LIMIT          The frequency inherits exactly the defensibility of the count.
                   A contested decomposition yields a contested frequency.

THE LADDER         N = 12, the twelve state classes fixed in §5.6, unaltered.
                   ln N = 2.484907.  Max f = k*144 = 69.6133 Hz.
                   Six pitch classes {1,3,5,7,9,11} = ceil(12/2).

THE REJECTIONS     Arousal: REJECTED (V2, V3, V7).
                   Hierarchical need: REJECTED (V1, V2; 35.2% of headroom lost).
                   Valence magnitude: REJECTED (V4; the CC-R8 collapse).
                   Developmental stage: admissible as an ORDINAL, not adopted;
                   any frequency derived from it carries CC-R1's demotion.

THE PROHIBITIONS   The twelve states are not diagnoses. No clinical, medical, diagnostic
                   or therapeutic claim is made anywhere in this framework. The frequencies are
                   essence frequencies and not physical resonances. All cases constructed.
```

Section 6 takes the exposure this section identified — that the count is only as good as the decomposition — and does the work of narrowing it: the elicitation procedure, the decomposition rules, the agreement protocol, the worksheet, and the validation table with its seventh gate reported honestly as not yet run.

6 Ordering the states

Section 5 made a strong claim and then named its own exposure. The claim: the psychological essence index is a count, counts are absolute-scaled, absolute scales admit only the identity transformation, and therefore `f = k n^2` is a licensed operation on `n` in a way it is not licensed on any ratio of ratings. CC-R1's demotion of `D/P` does not reach the index. The index sidesteps the demotion.

The exposure: all of that holds *given a decomposition*, and Section 5 said nothing about how the decomposition is fixed. **The frequency inherits exactly the defensibility of the count. A contested decomposition yields a contested frequency.** Section 5 priced that exposure — three defensible decompositions of one constructed episode gave 39.1575, 48.3426 and 58.4945

Hz, in pitch classes 9, 5 and 11: three different notes — and then stopped, because narrowing the spread is a different kind of work.

This section does that work, and it owns two things the rest of the framework cites and does not redefine.

The first is the **elicitation procedure**: how a practitioner actually obtains the intensities that build a state vector, and how the vector is normalised. the framework's state-vector construction fixes the construction, and later papers in this series apply it; this section specifies it end to end, with a printable worksheet.

The second is the **defence of the count against the obvious objection** — that deciding what counts as one component is itself a judgment. That objection is correct. It is not fatal, it is not answered by assertion, and the answer has three parts: decomposition rules that make the judgment explicit, inter-rater agreement as the test that makes it checkable, and an honest statement of where the answer runs out. All three are given here, including the third.

The section closes with the Formula Validation Protocol table for the index, all seven gates, and FVP-7 — cross-individual reproducibility, the gate that decides the hard-science question and the one most often skipped — is reported as NOT YET RUN.

The standing prohibitions govern every page. The essence frequencies are essence frequencies, not physical resonances. **The twelve states are not diagnoses**; nothing in this procedure is an assessment, a screening, or a clinical instrument, and no part of the framework of this framework makes a clinical, medical, diagnostic or therapeutic claim. All cases are constructed and labelled.

The Law

6.1 What is being built, and why the index alone is not enough Section 5 produced a number for a state. That is not yet enough to do anything with, because every diagnostic in this framework is relational. The octave deficit needs `f_rel`; `f_rel` needs `u`; `u` needs two vectors. **A single index is a note. The framework's instrument reads intervals, and intervals need two things to hold at once.**

A person does not occupy one state. A person holds several, at different strengths, and holds them under different framings of the same situation. Representing that is the object this section builds.

the framework's state-vector construction, and it is the only construction of `u` in this framework:

```
THE STATE VECTOR AND THE MEASUREMENT

A subject, a framing, or a person is represented as a unit vector over the twelve states,
components = intensity, normalised to length 1. Then for any two such vectors x, y:

    cos theta = x . y
    u         = 1 − (x . y)^2
    F         = ln(1 + 11 u)
    f         = k e^(2F) = k (1 + 11 u)^2

with N = 12, so N − 1 = 11, and k = 0.483425732988414 Hz.
```

Three registers, one object. `cos theta` is the alignment of the two vectors. `u` is the non-parallelism, in `[0, 1]`, equal to `sin^2 theta`. `F` is the interaction surplus in nats, bounded above by `ln 12 = 2.484907`. And `f` is that surplus expressed as a pitch, because — as the essence bridge of Series B established and Section 4 restated — one nat of surplus is exactly 2.885 octaves.

The companion paper on the dyad interval owns `f_rel = k(1 + 11u)^2` and the companion paper on between-persons measurement owns the between-persons measurement of `u`. This section owns the step before both of them: getting from a person and a situation to a unit vector that deserves to be entered into either.

6.2 The state vector, defined

```
THE STATE VECTOR

    x = (x_1, x_2, ..., x_12)          x_n >= 0 for all n
    x_n = the elicited intensity with which state class n is present
    the support of x = { n : x_n > 0 }
    the unit vector = x / ||x||_2 ,   ||x||_2 = sqrt( sum_n x_n^2 )

A state class is present with intensity x_n if the subject's account of the situation
contains the n-component specification of state class n at that strength.
```

Four properties of this object need stating before any procedure is built on it, because each one constrains the procedure.

It is non-negative. There are no negative intensities. A state is present or it is not; the absence of fear is not the presence of negative fear. This is not a modelling convenience — it is what keeps §5.14's valence collapse from reappearing here in a different costume.

Its support is the set of nonzero entries. Writing a zero into a slot changes nothing:

```
a = {3:8, 5:4, 7:2}         b = {3:8, 5:4, 7:2, 11:0}
u(a, b) = 0.000000000000
```

Exactly zero. Recording a state as absent and not recording it at all are the same act, so a practitioner may not inflate a subject's apparent breadth by listing more zeros. This closes an otherwise open route to gaming a reading.

It is normalised in the L2 norm, not the L1 norm, because `u = 1 − (x.y)^2` is built from an inner product and the inner product's natural normalisation is Euclidean. This is a consequence of the framework's state-vector construction's construction and not an independent choice.

And the zero vector has no unit form. A subject with no state at all cannot be normalised, and should not be: `n = 0` is the true zero of Section 5's count, `f = k * 0^2 = 0`, no components, no state, no pitch. The count's zero and the essence law's ground coincide by construction, and the vector formalism inherits that: there is nothing to normalise because there is nothing there.

6.3 The two invariances, and only one of them is good news

Normalisation is usually treated as bookkeeping. Here it is load-bearing, because it determines exactly what the elicited intensities have to be good for — and the answer is less demanding than a reader would guess in one respect and more demanding in another. Both are computed.

FIRST INVARIANCE: the unit vector does not care what units the ratings were in.

```
raw ratings {3: 8, 5: 4, 7: 2} rescaled and renormalised:

raw 0–10              ->  (0.872871561, 0.43643578, 0.21821789)
raw 0–1               ->  (0.872871561, 0.43643578, 0.21821789)
raw 0–100             ->  (0.872871561, 0.43643578, 0.21821789)
proportion of total   ->  (0.872871561, 0.43643578, 0.21821789)

identical to nine decimal places
```

Because `(a x) / ||a x|| = x / ||x||` for any `a > 0`, any common positive rescaling of the intensities cancels exactly. A practitioner using a 0-to-10 scale and a practitioner using a 0-to-100 scale, who assign proportionally the same intensities, produce the same vector and therefore the same `u`, the same `F`, and the same `f`.

This is a real and useful freedom, and it should be stated for what it is: the elicitation instrument's unit is unconstrained. There is no calibration step, no reference stimulus, and no requirement that two practitioners use the same anchors. What is required is that the *ratios* of the intensities be the same.

SECOND INVARIANCE, and it does not hold. The vector is not invariant under an order-preserving transformation of the ratings. Take the same raw ratings and apply monotone transformations — transformations that are admissible if the ratings are merely ordinal — then compute `u` against a fixed reference framing `{3:5, 5:7, 10:4}`:

```
transform      u          f_rel (Hz)
identity    0.388360      13.4361
sqrt        0.371445      12.5044
square      0.525275      22.2093
log(1+x)    0.362115      12.0049

u spans 0.362115 to 0.525275
f spans 12.0049 to 22.2093 Hz  =  1065.05 cents  =  0.888 octaves
```

Nearly nine tenths of an octave of disagreement from transformations that preserve every ordinal fact in the data. So the elicited intensities cannot be merely ordinal. They must carry ratio information — the claim that state 3 is present twice as strongly as state 5 must mean something — even though the unit in which that ratio is expressed is free.

This is exactly the right place for the honest statement, and it must not be softened. Section 5's count argument covers `n`. It does not cover the intensities. **`n` is absolute-scaled by construction; the intensity vector is not, and it needs its own defence.** What the two invariances establish is the precise shape of that defence: the intensities must be ratio-scaled up to a common free multiplier — which is the definition of a ratio scale — and nothing weaker will do.

There is a substantial body of evidence that human subjects can supply exactly that, and this framework already leans on it. Stevens' method of magnitude estimation produces ratio-scale measurements of subjective experience from human subjects, and it replicates [Stevens, 1957] — that finding is the reason Series B Ch 38 could treat subjective magnitude as measurable at all. The elicitation procedure of §6.9 is built as a magnitude-estimation procedure for that reason and not by stylistic preference. The honest form of the claim is: the intensities are ratio-scaled to the extent that magnitude estimation delivers ratio scales, which is a measured property of a procedure and not a constant of nature, and Stevens' own exponents are known to vary with method, range, and instruction [Teghtsoonian, 1971, Poulton, 1968], and the century-long dispute about what the two psychophysical laws are each measuring is best read through the reconciliation literature rather than adjudicated here [Krueger, 1989].

6.4 What normalisation discards, and the standing reporting rule Normalisation throws away the length of the vector, and the length is not nothing.

```
low-intensity  {3: 1,   5: 0.5,  7: 0.25}    raw L2 norm = 1.145644
high-intensity {3: 8,   5: 4,    7: 2   }    raw L2 norm = 9.165151

ratio of norms = 8.0000x
u between them = 0.000000000000     (exactly zero -- same direction)
```

Two states with identical profile shape and an eightfold difference in absolute intensity are perfectly parallel and carry zero surplus between them. That is correct behaviour for the

surplus formalism, which measures difference of *direction*, and it is a serious limitation on what a single `u` reading reports.

> **THE MAGNITUDE REPORTING RULE.** The raw L2 norm `||x||_2` is reported alongside every state vector, in the units of the elicitation instrument, and is never discarded. `u` measures direction; the norm measures how much there is. A reading that publishes `u` without the norms it came from has thrown away half the measurement.

This is CC-I5's sibling and it is a house requirement from here forward. The personal-spectrum worksheet of this series carries a field for it, and the between-persons protocol of this series requires both parties' norms.

6.5 The objection, stated at full strength Here is the objection the section exists to answer, put as forcefully as it can be put, because an objection answered in a weakened form has not been answered.

> **THE DECOMPOSITION OBJECTION.** The count argument proves that a count is absolute-scaled. Granted. But the count is a count *of components*, and "component" is not given by nature. It is given by a decomposition the analyst chooses. Section 5's own worked example admits three defensible decompositions of one episode, yielding 39.1575, 48.3426 and 58.4945 Hz — a spread of well over an octave, in three different pitch classes.
>
> So the scale problem has not been solved. It has been renamed. Where the Moment equation required a judgment of magnitude, the essence index requires a judgment of granularity, and a judgment of granularity is just as unconstrained as a judgment of magnitude. The apparent rigour of the count is borrowed from the arithmetic of counting and does not attach to the thing being counted. The index is a rating with an integer's costume on.

That objection is substantially correct and this framework concedes its correct part immediately. The scale problem was not solved. It was *moved*. Section 5 said so in the same paragraph that made the claim, and the concession is repeated here so that no reader has to discover it.

What the objection gets wrong is the inference that moving a problem is worthless. The three parts of the answer are that the new location is (a) explicit, (b) testable by agreement, and (c) bounded in a way the old location was not. Sections 6.6, 6.7 and 6.8 take them in order, and 6.8 is where the answer is incomplete and says so.

6.6 Answer, part one: the decomposition rules **A judgment that is made explicitly and in advance is a different object from a judgment made implicitly and after the fact, and the difference is not rhetorical.** An explicit rule can be published, applied by a third party, disagreed with in a specific place, and revised without revising everything else. That is the entire content of the move Section 5 made, and these rules are what make it real.

```
THE DECOMPOSITION RULES        D1 – D8
Fixed before any state is examined. Published with every frequency. Not revised mid-analysis.

D1  THE SUBSTITUTION TEST.  A candidate is a required component if and only if removing it
    leaves a description consistent with a state the subject would say they were not in.
    If the residual description still picks out the same state, the candidate was not required.

D2  THE LADDER IS THE VOCABULARY.  Components are drawn from the twelve declared components of
    the Section 5 ladder and from no other list. Anything that does not correspond to one of
    the twelve is not a component of this index. An episode requiring a thirteenth thing is
```

```
    outside the ladder's scope and is reported as such, not shoehorned in.

D3  NESTING.  The ladder is nested: a state that requires component n requires every component
    below n. If component 8 is present and component 4 is absent, the decomposition is wrong
    and must be redone, not recorded. Report gaps as decomposition failures.

D4  ONE ORDER, ONE COMPONENT.  A modelled other is one component (index 8). The other's model
    of the subject is a second, distinct component (index 11). Each order of recursion counts
    exactly once, however elaborate the content at that order.

D5  NOUNS ARE NOT COMPONENTS.  Count what the specification requires, not what the description
    contains. The hour, the room, the weather and the errand are one component (situation)
    unless the substitution test separates them.

D6  ENTAILMENT MERGES.  If component A is entailed by component B in the subject's account,
    A and B are one component. Distinguishability is the criterion, and it is the subject's
    distinguishability, not the analyst's.

D7  ABSENCE IS RECORDED.  Every one of the twelve components is explicitly marked present or
    absent. There is no "not mentioned" state. This is what prevents an analyst from reaching
    a target count by leaving inconvenient components unexamined.

D8  THE GRANULARITY IS FIXED FIRST.  The decomposition rules are chosen before the episode is
    read and are not adjusted after seeing the count. Any adjustment is a new analysis with a
    new declared rule set, and both are published.
```

D1 and D8 carry the weight. D1 gives the judgment a determinate form: it is no longer “is this a component?” but “does removing it change which state this is?”, which is a question about the description and which two analysts can argue about with the description in front of them. D8 removes the degree of freedom that §5.18 named as the subtlest failure mode — an analyst who splits and merges until the count lands where they expected it to.

D2 is the rule most likely to be resented and it is the one that makes the ladder a ladder. The twelve components are the vocabulary. This is restrictive by design: a fixed vocabulary is what makes counts comparable across analysts and across episodes, which is the entire purpose of having one.

D3 deserves a word because it will fire. Real accounts routinely appear to skip rungs — a subject speaks about what someone else was thinking without ever saying what the situation was. Almost always the lower component is present and unstated rather than absent. D3’s instruction is to go back and ask, not to record a gap, because a recorded gap propagates into a count that no formula in this framework is defined on.

6.7 Answer, part two: agreement is the test

Explicitness is necessary and it is not sufficient. A rule can be perfectly explicit and still be applied differently by every person who reads it. Agreement between analysts is a reliability statistic and not a validity statistic [Cronbach and Meehl, 1955, Campbell and Fiske, 1959]. **The test of a decomposition is not whether the analyst finds it convincing. It is whether two analysts, working independently from the same declared rules and the same material, return the same integer.**

This is the move that puts the index in a different epistemic position from the quantity CC-R1 demoted, and it is worth being precise about why. `D/P`’s problem is not that analysts disagree about it; `D/P`’s problem is that even a single analyst’s number is undetermined, because the ratio is not invariant under transformations the data cannot rule out. No amount of agreement fixes that: two analysts who both apply the same inadmissible transformation agree perfectly and are both reporting a number that is not a measurement. **The count has no such internal indeterminacy. Given a decomposition, the count is determined. So disagree-**

ment between analysts is the *only* place error can enter, and disagreement between analysts is measurable.

That is the whole of the argument, and its practical form is a protocol:

```
THE AGREEMENT REQUIREMENT

No essence frequency is published without an agreement figure, and the agreement figure is
reported on THREE quantities, not one:

    A1  exact agreement on the integer n
    A2  agreement on the pitch class (the odd part of n)
    A3  the distribution of the pitch disagreement in cents, |1200 * log2((n_1/n_2)^2)|

A2 and A3 are not derivable from A1 and are not derivable from any weighted kappa.
Section 6.17 shows a constructed case in which they diverge sharply.
```

The reason A2 and A3 are mandatory rather than optional is Section 5 §5.17's finding: rank distance and pitch distance are different metrics and are not monotonically related, because the pitch class is the odd part of the index and the odd part is a number-theoretic function that is not monotone in `n`. An agreement statistic built on rank distance therefore cannot see the disagreement that matters, and reporting one alone is worse than reporting nothing because it manufactures confidence.

6.8 Answer, part three: where the answer is incomplete Three concessions. None is rhetorical, and none of them is repaired later in this framework.

CONCESSION ONE — agreement is not correctness. Two analysts trained on the same rules by the same person will agree with each other and may both be wrong in the same direction. Inter-rater agreement measures the reliability of a procedure and says nothing about its validity. This is a standard and well-understood limitation of every agreement statistic in psychometrics, and it is not special to this construction. What would be needed is a criterion external to the decomposition against which counts could be checked, and this framework does not have one. The nearest available substitute is predictive: does the count predict something it was not fitted to? That is FVP-6's business, the design is specified in §6.14, and it has not been run.

CONCESSION TWO — D2 fixes the vocabulary and does not justify it. The twelve components are defended in Section 5 §5.7 by the nesting argument: each rung adds exactly one component to the rung below, so the order is forced rather than chosen. That defence establishes the *order* given the components. **It does not establish that these twelve are the right twelve, or that twelve is the right number.** Ashby's requisite variety and Section 4's ceiling `c/eta <= ln N` both make `N` consequential — raising `N` is one of only four levers — so the choice of twelve is not harmless. A thirteenth genuinely distinct component would raise the ceiling, and this framework has no argument that no such component exists. It has only the claim that it could not find one that survived D1 and D6.

CONCESSION THREE — and this is the deepest one. The substitution test D1 appeals to what *the subject would say they were not in.* That is a judgment made by a person about their own state, and this framework has no independent access to it. The whole apparatus therefore rests on first-person report at exactly one point. **The count is absolute-scaled given a decomposition; the decomposition is explicit and testable; and underneath both there is a subject saying yes, that is a different state, and there is nothing underneath that.**

This is not a defect the framework can engineer away, and it should not pretend to. It is the same foundation psychophysics rests on, and psychophysics has done a century of replicable work on it — which is a genuine reason for confidence and not a proof of anything. **The honest summary is that the essence index converts an unbounded judgment into a bounded, explicit, testable one, and does not eliminate judgment.** No procedure does. What this framework

claims is that the judgment has been moved somewhere it can be audited, and Section 5's claim about scale type is a claim about the arithmetic downstream of that judgment, not about the judgment itself.

The Engineering

6.9 The elicitation procedure **This is the procedure. Later papers in this series cite it and do not restate it.** It is written for a practitioner working with one subject on one situation, producing one or more framings. Steps 1 to 5 establish the frame; 6 to 11 elicit; 12 to 16 build and check the vector; 17 to 20 report.

A prohibition that governs the whole procedure and is repeated inside it: this is not an assessment. Nothing in the procedure evaluates the subject, ranks them, screens them, or produces anything that may be described to them or to anyone else as a finding about their mind. The procedure produces a description of a description. A practitioner who cannot hold that distinction should not run it.

```
STEP 1 -- DECLARE THE RULE SET BEFORE READING ANYTHING.
    Write down: the decomposition rules in force (D1–D8 or a declared variant), N = 12, the
    ladder in use, the intensity instrument and its range, and the audible band and folding
    convention. Date it. This page is published with the reading. (D8)

STEP 2 -- OBTAIN CONSENT, AND MAKE DECLINING STRUCTURAL.
    The subject is told what is being elicited, what will be computed, what it does and does
    not mean, and that any question may be skipped without giving a reason and without the
    procedure being abandoned. **No intervention here may remove a person's ability to
    decline. Consent is structural, not advisory.** Nothing is masked, nothing is concealed,
    and nothing is administered below the subject's awareness.

STEP 3 -- FIX THE SITUATION.
    One situation, named in the subject's own words, in one sentence, written down before
    anything else is asked. Everything downstream is indexed to this sentence.

STEP 4 -- FIX THE FRAMINGS TO BE ELICITED.
    Decide in advance how many framings of the situation will be elicited and what each is.
    Two is the minimum for any relational reading; three is the house default. Name them
    before eliciting any of them.

STEP 5 -- SILENT WRITE-FIRST.
    The subject writes each framing out in full, independently, before any two are compared,
    and before the practitioner comments on any of them. This is the personal-scope form of
    the silent-write-first rule and it exists to stop the framings converging by conversation
    before they have been recorded.

STEP 6 -- READ THE FIRST FRAMING AGAINST THE TWELVE COMPONENTS.
    For each of the twelve components in order, mark PRESENT or ABSENT. Every one of the
    twelve is marked. There is no blank. (D7)

STEP 7 -- APPLY THE SUBSTITUTION TEST TO EVERY COMPONENT MARKED PRESENT.
    Remove it; read the residual description back to the subject; ask whether that is still
    the state they were in. If yes, the component was not required: unmark it. (D1)

STEP 8 -- CHECK NESTING.
    If any component is marked present above a component marked absent, go back and ask
```

about the absent one. Do not record a gap. (D3)

STEP 9 -- COUNT.
n = the number of components marked present after steps 7 and 8. Record the count and, beside it, the list. **A count without its list is a rating and is rejected.** (§5.18)

STEP 10 -- ELICIT THE INTENSITY, BY MAGNITUDE ESTIMATION.
Choose the state class the subject reports most strongly and assign it a reference value -- 10 on a 0–10 instrument. Then, for each other state class present, ask: "compared with that one, how strong is this one -- half as strong, a quarter, three times?" **The question is asked in ratios because ratios are what the construction consumes** (§6.3). Record the answers. Do not average across framings.

STEP 11 -- ELICIT THE REMAINING FRAMINGS.
Repeat steps 6–10 for each framing named in step 4, without showing the subject the earlier vectors.

STEP 12 -- CHECK SUPPORT OVERLAP. THIS STEP IS NOT OPTIONAL.
If two framings have disjoint support -- no state class present in both -- then u = 1 exactly, regardless of the intensities, and the reading is degenerate. Verified:

> five constructed pairs with disjoint support, random intensities:
> u = 1.000000000000 in every case

A degenerate reading is not a high reading. It is no reading. **Framings must be represented with overlapping support.** The chord analysis owns this trap; the procedural requirement here is to detect it before anything is computed and to return to step 5 if it fires.

STEP 13 -- RECORD THE RAW NORM, THEN NORMALISE.
Compute ||x||_2 in the instrument's units and write it down (§6.4). Then divide every component by it. The unit of the instrument cancels (§6.3); the norm does not, and is reported.

STEP 14 -- COMPUTE THE PAIRWISE READINGS.
For every pair of framings: cos theta = x.y, u = 1 − (x.y)^2. Report all pairs individually. Do not report only the mean.

STEP 15 -- AGGREGATE.
Mean pairwise u over all pairs; then F = ln(1 + 11 u_bar), N_eff = 1 + 11 u_bar, f = k(1 + 11 u_bar)^2.

STEP 16 -- COMPUTE THE SPECTRUM AND THE CHORD.
For each state class in the union of the supports, f = k n^2, folded into the declared band, with pitch class = odd part of n. The chord analysis owns the chord; this gives its inputs.

STEP 17 -- RUN THE SECOND RATER.
An independent analyst, working from the step 1 rule set and the step 5 written framings alone, repeats steps 6–13. Report A1, A2 and A3 (§6.7).

STEP 18 -- RUN THE PERMUTATION TEST.
Swap two adjacent indices in the ladder, recompute, report the change in cents and in pitch class. This is CC-I6 of Section 5 §5.16 and it takes ten minutes.

```
STEP 19 -- PUBLISH THE DECOMPOSITION WITH THE FREQUENCY.
    Same table, same page. Where it will not fit, the frequency is not published either.

STEP 20 -- STATE THE BOUNDS IN THE REPORT ITSELF.
    Every report carries, in its own words: that the frequencies are essence frequencies and
    not physical resonances; that **the twelve states are not diagnoses and the reading is not
    a clinical, medical or therapeutic assessment of anyone**; and that the frequency inherits
    exactly the defensibility of the count.
```

6.10 The worksheet Printable. One sheet per framing, plus one summary sheet. It is the same form as the worksheet of §6.10.

```
================================================================================
ESSENCE STATE VECTOR -- FRAMING SHEET                        sheet ___ of ___
================================================================================
Rule set declared (date / version): _____________________   N = 12
Ladder: the twelve-state ladder, Section 5 §5.6, unaltered   k = 0.483425732988414 Hz
Intensity instrument and range: _______________  Audible band: ________ Hz
Consent recorded, decline available at every step:  [ ]

THE SITUATION, in the subject's words, one sentence:
_____________________________________________________________________________

THE FRAMING, name: _____________________   written independently before comparison [ ]

------------------------------------------------------------------------------
 n  state class                  component added            PRES  SUBST  INTENSITY
------------------------------------------------------------------------------
 1  undifferentiated activation  arousal alone              [  ]  [  ]   _________
 2  valenced activation          + valence                  [  ]  [  ]   _________
 3  directed feeling             + intentional object       [  ]  [  ]   _________
 4  situated feeling             + situation                [  ]  [  ]   _________
 5  attributed feeling           + cause                    [  ]  [  ]   _________
 6  temporal feeling             + projected duration       [  ]  [  ]   _________
 7  reflexive feeling            + self as object           [  ]  [  ]   _________
 8  social feeling               + another mind modelled    [  ]  [  ]   _________
 9  normative feeling            + standard applied         [  ]  [  ]   _________
10  narrative feeling            + placed in a story        [  ]  [  ]   _________
11  reciprocal feeling           + the other modelling you  [  ]  [  ]   _________
12  committed feeling            + future action bound      [  ]  [  ]   _________
------------------------------------------------------------------------------
PRES  = marked present or absent. EVERY row is marked. No blanks.  (D7)
SUBST = survived the substitution test.                            (D1)
Nesting check passed (no present-above-absent):  [ ]                (D3)

COUNT   n = __________       COMPONENT LIST (required; a count without a list is void):
_____________________________________________________________________________

RAW NORM   ||x||_2 = __________ (instrument units)   <- reported, never discarded

UNIT VECTOR (raw / norm), to 6 d.p.:
  x1 ____  x2 ____  x3 ____  x4 ____  x5 ____  x6 ____
  x7 ____  x8 ____  x9 ____  x10 ____ x11 ____ x12 ____
Support = { _____________________ }
================================================================================
```

```
================================================================================
ESSENCE STATE VECTOR -- SUMMARY SHEET
================================================================================
Framings elicited: _____________________________________
SUPPORT OVERLAP CHECK (step 12): every pair shares at least one state class [ ]
    if any pair does not: u = 1 by construction, reading is DEGENERATE, return to step 5

PAIRWISE, all pairs listed individually:
  pair _______________  cos theta = _________   u = _________
  pair _______________  cos theta = _________   u = _________
  pair _______________  cos theta = _________   u = _________

mean pairwise u_bar = _________
F = ln(1 + 11 u_bar) = _________ nats        N_eff = 1 + 11 u_bar = _________
f = k (1 + 11 u_bar)^2 = _________ Hz        folded (j = ____) = _________ Hz

SPECTRUM: states spanned { _______________ }
  n = ___  f = _________ Hz  folded _________  class ___
  n = ___  f = _________ Hz  folded _________  class ___
  n = ___  f = _________ Hz  folded _________  class ___
  n = ___  f = _________ Hz  folded _________  class ___
distinct pitch classes: { _____________ }   count ____  of ceil(12/2) = 6 available

AGREEMENT (step 17):  A1 exact-n ____/____   A2 pitch-class ____/____
                      A3 cents: mean ______ median ______ max ______
PERMUTATION TEST (step 18): swapped ___ <-> ___ ;  class ___ -> ___ ;  ______ cents

BOUNDS, restated on every report:
  [ ] essence frequencies, not physical resonances
  [ ] the twelve states are NOT diagnoses; this is not a clinical, medical, diagnostic
      or therapeutic assessment of any person
  [ ] the frequency inherits exactly the defensibility of the count
  [ ] decomposition published with the frequency
================================================================================
```

6.11 A worked pass through the vector arithmetic Before the full carried case, one framing taken end to end so the arithmetic is visible.

Raw intensities elicited by magnitude estimation, on a 0-to-10 instrument, for a framing with three state classes present:

```
state 3  (directed feeling)      8
state 5  (attributed feeling)    4
state 7  (reflexive feeling)     2
```

Step 13, the raw norm:

```
||x||_2 = sqrt(8^2 + 4^2 + 2^2) = sqrt(64 + 16 + 4) = sqrt(84) = 9.165151
```

Recorded, then divided out:

```
x_3 = 8 / 9.165151 = 0.872872
x_5 = 4 / 9.165151 = 0.436436
x_7 = 2 / 9.165151 = 0.218218

check: 0.872872^2 + 0.436436^2 + 0.218218^2 = 1.000000
```

And the first invariance in operation: had the same subject been elicited on a 0-to-1 instrument at 0.8, 0.4, 0.2, the norm would have been 0.916515 and the unit vector identical to nine decimals. The instrument's unit cancels. The ratios 4 : 2 : 1 are the entire content.

The Specification

6.12 The agreement protocol

```
S6.1  TWO RATERS MINIMUM.  Independent, working from the declared rule set and the written
      framings only. No consultation until both have completed step 13.

S6.2  BLIND TO EACH OTHER AND TO THE HYPOTHESIS.  The second rater does not see the first
      rater's count and is not told what reading the analysis is expected to produce.

S6.3  REPORT A1, A2 AND A3 SEPARATELY.  Exact agreement on n; agreement on pitch class;
      the distribution of cents disagreement. A weighted kappa may be reported in addition
      and never instead.

S6.4  THE THRESHOLD IS ON PITCH, NOT ON KAPPA.  Declare, before the analysis, the maximum
      pitch disagreement at which a frequency will still be published. Report the fraction
      of episodes exceeding it.

S6.5  DISAGREEMENTS ARE ADJUDICATED IN PRINT OR NOT AT ALL.  Where raters differ, publish
      both counts and both component lists. A consensus count with the disagreement erased
      is a fabricated agreement figure.

S6.6  RATER DRIFT IS CHECKED.  A fixed set of episodes is re-rated at intervals by the same
      rater. Within-rater agreement over time is reported alongside between-rater agreement.
```

6.13 The reporting standard Every published essence frequency in this framework, and every one produced by a reader following it, carries: the decomposition (CC-I1), the count rule (CC-I2), `N` and the ceiling `ln N` (CC-I3), the tie-break (CC-I4), the band and the folding exponent `j` (CC-I5), the permutation test result (CC-I6), and the agreement figure (CC-I7). Those are Section 5 §5.16's seven publication requirements and this section's procedure is built to produce all seven as a by-product of running it.

To which this section adds two:

```
CC-I8  THE RAW NORMS.  ||x||_2 for every vector, in the instrument's units.  (§6.4)
CC-I9  THE SUPPORT OVERLAP.  For every pair, the states shared. A pair sharing none is
    reported as DEGENERATE and its u is not published as a reading. (step 12, chord analysis)
```

6.14 The Formula Validation Protocol table for the index **No formula appears in this series without passing all seven gates in print.** The protocol is Series B's Formula Validation Protocol and it is binding. This is the FVP table for the psychological essence index and the essence law applied to it.

F-06 · THE PSYCHOLOGICAL ESSENCE INDEX

```
STATEMENT

    n  :=  the number of distinguishable components required to specify the state
    f  =   k n^2,      k = 0.483425732988414 Hz,   N = 12,   n in {1, ..., 12}
```

READS. How many things must be said before a listener knows which state is meant -- squared, and scaled to a pitch.

STATUS. The index is a DEFINITION. The essence law f = k n^2 is a POSTULATE about essence structure (essence-law spec), verified to collapse exactly from the 2022 source recipe and reproducing all 118 published element values to a worst relative deviation of 9.0e-15 (CC-R5). k is a CHOSEN reference constant, badged.

Gate	Verdict	Content
FVP-1 DIMENSIONS	**PASS**	`n` is a dimensionless count. `k` carries Hz. `f = k n^2` therefore carries Hz on both sides. Every interval `f(m)/f(n) = m^2/n^2` is dimensionless and `k` cancels exactly, which is why the interval structure developed here does not depend on the constant at all.

Gate	Verdict	Content
FVP-2 SCALE TYPE	**PASS — and this is the gate the section exists for**	`n` is a cardinality. It has a true zero (`n = 0`: no components, no state, `f = 0`) and a fixed unit (one component), and it admits only the identity transformation. **It is therefore ABSOLUTE-scaled, which is strictly stronger than the ratio scale that squaring requires.** Contrast CC-R1: `D/P` returns DEMOTE at this gate because `D` and `P` are ordinal ratings whose ratio is not invariant under admissible transformations — measured spread of 2.490 octaves in Section 5 §5.3. The index is not a ratio of ratings and does not inherit that demotion; it sidesteps it. The gate is passed for `n` and is passed for `n` only. The intensities that build the state vector are a separate quantity with a separate answer: they must be ratio-scaled up to a free multiplier (§6.3), the free multiplier is genuinely free (first invariance, exact), and the ratio requirement is genuinely binding (second invariance fails, 1065.05 cents of spread under monotone transformation). The warrant for the intensities is Stevens' magnitude estimation, reported as a measured property of a procedure and not a constant of nature.

Gate	Verdict	Content
FVP-3 LIMITING CASES	**PASS**	`n -> 0`: `f -> 0`. No components, no state, no pitch — and the zero vector correctly has no unit form. `n = 1`: `f = k`, the ground essence itself; a one-component state sits at the system's floor, which is the right place for undifferentiated activation. `n -> N`: `f -> k N^2 = 69.6133 Hz`, the ceiling, and by Section 4 the maximum sustainable cost ratio `c/eta <= ln N = 2.484907`. No divergence anywhere in range. The unbounded limit `n -> infinity` is not reachable: the ladder is finite by D2, and an episode requiring a thirteenth component is out of scope and reported as out of scope rather than indexed.
FVP-4 FIXED POINTS	**PASS**	`f(n) = n` has the fixed point `n = 1/k = 2.06856`, which is not an integer and therefore not attained — the index and its frequency never coincide, and no reading should be constructed as though they do. The structurally meaningful invariant is different and is number-theoretic: **`f` is invariant in pitch class under `n -> 2n`**, since `(2n)^2 = 4n^2` is exactly two octaves. The fixed points of the class map are the odd integers, which is why the twelve-state ladder has exactly `ceil(12/2) = 6` classes. The companion paper on octave folding owns the derivation.

Gate	Verdict	Content
FVP-5 IDENTIFIABILITY	**PARTIAL — declared**	`n` is recoverable from obtainable data: the component list is the data, and steps 6 to 9 recover it. **But the map from episode to count is not injective and not everywhere agreed.** Two decompositions of one episode can return different `n`, which is the decomposition objection of §6.5, conceded in §6.8. The identifiability that holds: given a declared rule set, two raters' counts are comparable and their disagreement is measurable, which is a stronger position than CC-R1's `D/P`, where a single rater's number is undetermined. The identifiability that does not hold: there is no procedure in this framework that identifies the *correct* decomposition, and §6.8 concession two states that the choice of twelve components is itself unjustified beyond a failed search for a thirteenth.

Gate	Verdict	Content
FVP-6 FALSIFICATION	**SPECIFIED — NOT RUN**	The essence law is a postulate about essence structure and makes no empirical claim on its own; what is falsifiable is the operationalisation. Design, following Paper E §5.3: preregister the rule set, the ladder, the elicitation script and the analysis; fit on training episodes; **FREEZE**; report held-out performance on whether the count predicts a criterion it was not fitted to. Beat four baselines — persistence, direct-input, matched autoregressive, and shuffled-input control. Kill condition, stated in advance: if counts elicited under the frozen protocol fail to beat the persistence baseline on held-out prediction, the operationalisation is dead and the ladder is withdrawn. Two standing corrections apply and must be carried into the design: **CC-R13(a)** — the shuffled-input control is anti-conservative as ordinarily run, with measured false-rejection rates of 0.108 (iid permutation) and 0.120 (circular shift) against a nominal 0.05, so the surrogate must be calibrated by simulation on the study's own design before use; and **CC-R13(b)** — no surrogate detects fabricated timing, so design controls (blinding, channel separation, prospective sampling, a yoked control) are mandatory and not substitutable. **This test has not been run.**

Gate	Verdict	Content
FVP-7 REPRODUCIBILITY	**NOT YET RUN — declared**	**The cross-individual test has not been performed for this index and this framework does not claim it has.** The question is whether the same operationalisation, applied by different analysts to different subjects in different settings, recovers the same parameters — specifically whether the component count for a given class of episode is a property of the episode or a property of the analyst who counted it. This is the gate that decides the hard-science question and it is the one most often skipped. It is skipped here too, and it is reported rather than omitted. What exists instead is the constructed demonstration of §6.17, which is an illustration of the *protocol*, not a result: constructed data cannot reproduce anything. The route to running it is the same one CC-R1 names for the Moment equation — Rasch instrumentation, which achieves separation of person and item parameters so that a measure does not depend on which items were used, at a measured cost of 123 items per term. CC-R14's costed programme puts the cross-subject recovery stage at the point of highest mortality, with a 55.3% chance of the programme dying at or before it, and the corpus's research-programme costing notes that this is not a coincidence: the gate that decides the question is the gate most likely to kill the answer. Until FVP-7 is run, every frequency in this framework is a frequency *under a declared decomposition*, and that qualification is not decorative.

Overall verdict on F-06: passes gates 1 to 4 outright, gate 5 partially and with the shortfall named, gate 6 specified but unrun, gate 7 unrun. That is a formula in good standing on scale type and structure and unvalidated on reproducibility, and the correct way to use it is the way this framework uses it: as an exact construction whose inputs are declared, whose disagreements are published, and whose empirical status is stated in the same breath as its numbers.

The Failure Modes

6.15 Leading the elicitation The practitioner reads the ladder aloud, in order, asking after each rung whether that one is present. By rung nine the subject has learned the shape of the instrument and what a fuller answer looks like. **Counts elicited this way climb through a session, and they climb for every subject, which is the signature of the artefact.**

The defences are in the procedure and each is there for this reason. Step 5's silent write-first captures the framing before any prompting. Step 7's substitution test removes components the subject will not defend when the residual is read back to them. Step 17's second rater works from the written framing and never meets the subject, so a leading elicitation shows up as a systematic gap between the two raters' counts rather than as a shared error.

Detection: plot count against position in session. A rising trend is an artefact until shown otherwise.

6.16 Disjoint support, and why a high `u` can be a null result Step 12 exists because of an exact and easily missed degeneracy. If two framings share no state class, their inner product is zero, so `u = 1` regardless of what the intensities are. Five constructed pairs with disjoint support and randomly drawn intensities give `u = 1.000000000000` every time.

A one-hot framing — a subject represented as holding exactly one state — is the extreme case, and one-hot framings are automatically orthogonal to each other. They score `u = 1` for everyone, on every topic, regardless of content, which means the reading has ceased to depend on the subject at all.

A degenerate reading is not a strong reading. It is the absence of a reading wearing the costume of the strongest possible one, and because `u = 1` is the top of the scale it looks like a triumph. This is a real trap and the framework names it. The chord analysis of this series owns it and states it in full; the companion paper on between-persons measurement restates it for the between-persons case. The obligation here is procedural: check overlap at step 12, before anything is computed, and return to step 5 if the check fires.

6.17 Agreeing on kappa while disagreeing on the note This is the failure mode most likely to defeat a careful practitioner, because every conventional instrument says the work is fine.

A **CONSTRUCTED** two-rater exercise. Fifteen constructed vignettes; two raters, working independently from the same declared rule set, each returning a count.

```
rater 1:  3  5  7  4  9  6 11  2  8  5 10  3  7 12  6
rater 2:  3  5  6  4  9  7 11  2  8  6 10  4  7 12  6
```

The agreement statistics:

```
exact agreement on n        11/15 = 0.7333
agreement within one         15/15 = 1.0000
mean |dn|                    0.2667
```

```
unweighted kappa             0.7044
linear-weighted kappa        0.9185
quadratic-weighted kappa     0.9838
```

By the conventional benchmarks — Landis and Koch's much-quoted table puts anything above 0.81 in the top band — a quadratic-weighted kappa of 0.9838 reads as very nearly perfect agreement. Every rater was within one category of the other on every single episode. A methods section reporting these figures would pass without comment.

Now report A2 and A3, which S6.3 requires:

```
agreement on PITCH CLASS     11/15 = 0.7333
pitch disagreement, cents:   mean 179.66   median 0.00   max 996.09
                             max = 0.8301 octaves

the four disagreements:
  ep 3   n = 7 vs 6    class 7 vs 3     533.74 cents
  ep 6   n = 6 vs 7    class 3 vs 7     533.74 cents
  ep10   n = 5 vs 6    class 5 vs 3     631.28 cents
  ep12   n = 3 vs 4    class 3 vs 1     996.09 cents
```

A quadratic-weighted kappa of 0.9838 coexists with pitch-class disagreement on more than a quarter of episodes and a worst-case disagreement of 996 cents — five sixths of an octave, from a one-component difference at the steep end of the ladder.

The mechanism is Section 5 §5.17's and it is exact. Weighted kappa penalises disagreement by rank distance. The essence law's consequence is a square law on the index, and the pitch class is the odd part, which is not monotone in `n`. On episode 12 the raters differ by one rank — the smallest possible disagreement, barely penalised by any weighting — and land 996 cents and two pitch classes apart, because at the bottom of the ladder one component is worth almost an octave.

The rule that follows is S6.3 and it is absolute: report agreement on the pitch, always, and never substitute a rank-based statistic for it. A weighted kappa is not wrong; it is answering a question about ranks, and this framework does not consume ranks.

The same failure appears one level down, in the intensities. A **CONSTRUCTED** intensity exercise, ten items, two raters, on a 0-to-10 instrument, gives `ICC(2,1) = 0.9379` — again, by convention, excellent. Follow one such disagreement downstream:

```
rater 1 vector {3:8, 5:4, 7:2}   ->  u = 0.388360   against a fixed reference framing
rater 2 vector {3:7, 5:5, 7:2}   ->  u = 0.301994
                                     |du| = 0.086365

f_rel:  13.4361 Hz  versus  9.0300 Hz  =  687.99 cents apart
```

Two raters differing by one point on two items, at an ICC of 0.9379, produce relationship frequencies 688 cents apart. The companion paper on the dyad interval will show that `f_rel = k(1 + 11u)^2` is steep in `u` for precisely this reason. **The agreement figure that matters is the agreement on the quantity intended for publication, and where a frequency is to be published, agreement on the frequency must be reported.**

6.18 Retro-fitting the decomposition

An analyst expects a state to come out around eight, counts six, and revisits the substitution test on two borderline components with a rather more generous reading. The count is now eight. Nothing dishonest has been done at any single step.

D8 is the only defence and it is a procedural one: the rule set is fixed and dated before the episode is read, and any revision produces a second, separately published analysis rather

than a correction to the first. Step 1 requires a dated rule-set page for exactly this reason. An analysis whose rule set has no date is not auditable and its frequencies should not be believed.

The second defence is step 17. A blinded second rater working from the original rule set will not reproduce the generous reading, and the disagreement lands in A1, A2 and A3 where it can be seen.

6.19 Running the procedure as an assessment This is the failure mode with consequences outside the framework, and it is prohibited rather than merely discouraged.

The procedure produces a number about a description. It does not produce a finding about a person. A practitioner who reports to a subject that their state "came out at nine," or who compares two subjects' counts, or who tracks a subject's counts over time as though the trend meant something about them, has converted a formal construction into an assessment. **The twelve states are not diagnoses. The counts are not severities. The chord is not a profile. No reading in this framework assesses anyone's mental health, and no procedure in it may be administered as though it did.**

Step 2 and step 20 are the structural defences: the subject is told at the outset what the procedure does and does not mean, and every report restates it. **And step 2's consent clause is not a formality — no intervention in this framework may remove a person's ability to decline, at any step, without explanation, and without the procedure being abandoned in consequence.** Nothing is masked, nothing is concealed, nothing operates below the subject's awareness. A person cannot consent to what is hidden from them.

6.20 Rater drift A single analyst's decomposition shifts over weeks. Early counts and late counts are not comparable, and because the drift is smooth it will not show up as noise; it will show up as a trend, and a trend invites an interpretation.

S6.6 is the defence: a fixed set of episodes, re-rated at intervals by the same rater, with within-rater agreement reported alongside between-rater agreement. The statistic is the same triple — A1, A2, A3 — computed against the rater's own earlier counts.

Drift is not a minor operational nuisance in this construction. It is FVP-7's question asked within a single person, and a procedure that cannot reproduce its own counts across time has no prospect of reproducing them across people.

Worked Example

CONSTRUCTED throughout. No data were collected, no person was elicited, and nothing below is a case report or an observation. This is the framework's carried subject, fixed by the framework's carried subject, and every part of the framework that needs this case uses this one and no other.

6.21 THE SUBJECT — three framings of one situation **Step 3, the situation, in the constructed subject's words:** *A colleague did something in a meeting that I have not been able to stop thinking about.*

Step 4, the framings to be elicited. Three, named in advance:

```
THREAT   -- the situation read as something happening to me
STORY    -- the situation read as an instalment in a longer account
MIRROR   -- the situation read as being about what it says about me
```

Steps 5 to 9, the decomposition. Each framing written independently, then read against the twelve components, with the substitution test applied to every one marked present and the nesting check run. The three framings return these supports:

```
THREAT   states 3, 5, 7      directed feeling, attributed feeling, reflexive feeling
STORY    states 3, 5, 10     directed feeling, attributed feeling, narrative feeling
MIRROR   states 5, 7, 10     attributed feeling, reflexive feeling, narrative feeling
```

Step 12, the support overlap check, run before anything is computed:

```
THREAT & STORY   share {3, 5}      2 states   OK
THREAT & MIRROR  share {5, 7}      2 states   OK
STORY  & MIRROR  share {5, 10}     2 states   OK
```

Every pair overlaps. The reading is not degenerate. Had the three framings been one-hot, or had any pair been disjoint, `u` would have been 1 by construction and the analysis would have returned to step 5. This check is the reason the carried case is built with overlapping support, and it is the fix, not a convenience.

Step 10, magnitude estimation. Raw intensities, elicited by ratio comparison against the strongest state class in each framing, on a 0-to-1 instrument:

```
THREAT   {3: 0.8, 5: 0.4, 7: 0.2}
STORY    {3: 0.5, 5: 0.7, 10: 0.4}
MIRROR   {5: 0.3, 7: 0.8, 10: 0.5}
```

Step 13, raw norms recorded, then normalisation:

```
||threat||_2 = sqrt(0.64 + 0.16 + 0.04) = sqrt(0.84) = 0.916515
||story ||_2 = sqrt(0.25 + 0.49 + 0.16) = sqrt(0.90) = 0.948683
||mirror||_2 = sqrt(0.09 + 0.64 + 0.25) = sqrt(0.98) = 0.989949

threat unit  = (0.872872 at 3,  0.436436 at 5,  0.218218 at 7 )
story  unit  = (0.527046 at 3,  0.737865 at 5,  0.421637 at 10)
mirror unit  = (0.303046 at 5,  0.808122 at 7,  0.505076 at 10)
```

The norms are reported and not discarded (CC-I8). They are close here — 0.9165, 0.9487, 0.9899 — which says the three framings are held at comparable overall intensity. Had one been an order of magnitude smaller, `u` would not have shown it, and the reading would have been of a framing the subject barely holds.

Step 14, the pairwise readings, all reported individually:

```
pair              cos theta      theta        u
mirror & story     0.436566    64.1150 deg   0.809410
mirror & threat    0.308607    72.0247 deg   0.904762
story  & threat    0.782074    38.5491 deg   0.388360
```

Step 15, aggregation:

```
mean pairwise u_int = (0.809410 + 0.904762 + 0.388360) / 3 = 0.700844

F_int = ln(1 + 11 * 0.700844) = ln(8.709284) = 2.164390 nats
N_eff = 1 + 11 * 0.700844 = 8.7093

f = k (1 + 11 u)^2 = 0.483425732988414 * 8.709284^2 = 36.6686 Hz
folded into 220–440 with j = 3:  36.6686 * 8 = 293.349 Hz
```

Step 16, the spectrum and the chord inputs:

```
states spanned: {3, 5, 7, 10}

n = 3   f =  4.3508 Hz   folded 278.453   class 3
n = 5   f = 12.0856 Hz   folded 386.741   class 5
n = 7   f = 23.6879 Hz   folded 379.006   class 7
n = 10  f = 48.3426 Hz   folded 386.741   class 5

pitch classes {3, 5, 7} -- THREE distinct, though FOUR states are held.
States 5 and 10 share odd part 5, so one of the four adds pitch but no new class.
```

That last line is the reading's most instructive feature and it is not an accident of this case. The subject holds four state classes and commands three notes. Adding the narrative component at `n = 10` raised the pitch by two octaves — `(10/5)^2 = 4` exactly — and added no new harmonic class, because 10 and 5 have the same odd part. **The companion paper on octave folding owns the rule and the chord analysis of this series owns what it means for a chord; here it is enough to observe that the twelve-state ladder offers six classes** (`ceil(12/2)`) **and this subject occupies three of them.**

Steps 18 and 19. The permutation test, run on the pair that matters most for this subject — swapping indices 5 and 6 in the ladder:

```
class 5 -> class 3,  568.72 cents
```

Well over the quarter-octave threshold, so the ordering is load-bearing for this reading and is defended by the nesting argument of Section 5 §5.7 or the reading is not published. The decomposition — the three component lists behind the three supports — is published in the same table as the frequencies.

6.22 The agreement pass on the same subject

Step 17, run on the THREAT framing. A second analyst, blind to the first analyst's counts and to the intended reading, works from the subject's written framing and the dated rule set alone.

The first analyst marked states 3, 5 and 7 present with intensities 0.8, 0.4, 0.2. Suppose the second analyst agrees on the support and reads the intensities as 0.7, 0.5, 0.2 — a one-point difference on a 0-to-10 instrument, at two of three items. Section 6.17 computed exactly this case:

```
analyst 1  ->  u(threat, story) = 0.388360
analyst 2  ->  u(threat, story) = 0.301994
|du| = 0.086365

f_rel:  13.4361 Hz  versus  9.0300 Hz  =  687.99 cents apart
```

Perfect agreement on the count — A1 = 3/3, A2 = 3/3, A3 = 0.00 cents on the spectrum — and 688 cents of disagreement on the relational frequency. The two agreement questions are different questions and S6.3 requires both because neither predicts the other.

What a practitioner does with that. Not average the two. **S6.5: publish both, with both intensity sets.** The honest report of this framing's relational reading is that it lies between 9.03 and 13.44 Hz under two defensible elicitations, and that the interval — not the point — is the result. A reading published as 13.4361 Hz, with the second analyst's elicitation quietly discarded, would be a fabricated precision, and it is the precision, not the number, that would be the lie.

6.23 The reading, and its bounds

What the constructed case establishes. Under the declared rule set, the twelve-state ladder, and the elicitation of §6.9, this constructed subject holds three framings of one situation with mean pairwise non-parallelism `u_int = 0.700844`,

internal surplus `F_int = 2.164390` nats, effective framing count `N_eff = 8.7093` of the twelve available, and a chord width of 36.6686 Hz, folded to 293.349 Hz. The subject spans four state classes and three pitch classes.

What it does not establish. Nothing about any person. The subject is constructed; the situation is invented; the intensities were chosen to demonstrate a procedure. **No observation is reported here and none is implied.**

What it licenses downstream. These exact numbers are the ones the companion papers of this series use. The chord analysis builds the chord from the four states and the three classes and owns the one-hot degeneracy trap that step 12 was written to catch. The companion paper on dissonance reads the intervals. The companion paper on the personal spectrum assembles the full personal worksheet around them. **No part of the framework re-elicits this subject and no part of the framework invents a different one.**

And the bounds, restated as step 20 requires them on every report.

`36.6686 Hz` is an essence frequency. **No claim is made that this constructed subject, or any person, or any tissue, or any molecule, physically resonates at that value or at any other value in this framework.**

The twelve states are not diagnoses. This reading is not a clinical, medical, diagnostic or therapeutic assessment of anyone, the counts are not severities, and the chord is not a profile of anything.

And the limit that travels with every number this section produces: the frequency inherits exactly the defensibility of the count. A contested decomposition yields a contested frequency. Section 6.22's 688-cent spread is what that sentence looks like when it is measured instead of asserted, and **FVP-7 has not been run**, so every frequency in this framework remains a frequency under a declared decomposition — which is a real thing to have, and is less than a measured one.

What this section fixes

```
THE ELICITATION PROCEDURE   Steps 1–20, §6.9. Cited by later papers; not restated there.

THE STATE VECTOR            Non-negative intensities over the twelve states, L2-normalised.
                            Support = nonzero entries. Zero vector has no unit form.

THE TWO INVARIANCES         The instrument's unit cancels exactly (verified to 9 d.p.).
                            Monotone transformation does NOT cancel (1065.05 cents of
                            spread). Therefore the intensities must be ratio-scaled up to a
                            free multiplier, and Stevens' magnitude estimation is the warrant.

THE MAGNITUDE RULE          ||x||_2 is reported with every vector and never discarded.
                            u measures direction; the norm measures how much there is.

THE DEFENCE OF THE COUNT    D1–D8 make the granularity judgment explicit; A1/A2/A3 make it
                            testable; §6.8 concedes what remains: agreement is not
                            correctness, twelve is unjustified beyond a failed search for a
                            thirteenth, and the substitution test rests on first-person
                            report with nothing underneath it.

THE AGREEMENT PROTOCOL      S6.1–S6.6. Report agreement on the pitch, always. A weighted
                            kappa of 0.9838 coexisted with pitch-class disagreement on 4 of
                            15 constructed episodes and a 996.09-cent worst case.

THE FVP TABLE               F-06, all seven gates. Gates 1–4 PASS. Gate 5 PARTIAL, declared.
```

```
Gate 6 SPECIFIED, NOT RUN, with CC-R13's two corrections carried.
**Gate 7 NOT YET RUN**, and reported as not run.
```

The companion paper on octave folding takes the index this section has learned to elicit and applies `f = k n^2` to transient states, where it owns octave folding, the odd-part class rule, and the `ceil(N/2)` ceiling that has been quoted twice already and derived neither time.

7 Results

The seven results below are the units other papers in this corpus cite. A citation to this paper should name one of them. Each carries its status: **DERIVED** where it follows from the stated definitions by algebra, **COMPUTED** where it is the output of exact arithmetic on stated parameters, and **POSTULATED** where it is asserted by this framework.

R1 — the minimum sustainable frequency. $f_{\min} = k\, e^{2c/\eta}$, with $k = 0.483425732988414$ Hz (defined constant; the hertz postulated, §2.3 P-b), and the exact trichotomy $f \gtreqless f_{\min} \iff Q_\infty \gtreqless 0$. *Status:* DERIVED from the storage-law fixed point [Curzi, 2026a] and the essence bridge [Curzi, 2026k], conditional on P-a and P-b. Verified at twelve parameter sets with sign agreement in all twelve, including the exact-equality case.

R2 — the octave form. Octaves required $= 2(c/\eta)/\ln 2 = 2.885390081777927 \times (c/\eta)$. One nat of required surplus is exactly 2.885390 octaves, 3462.4681 cents and 1.442695 bits; one octave is exactly 0.5 bits. *Status:* DERIVED; the constant is exact.

R3 — the deficit identity. $D = 2.885390(c/\eta - F) = -2.885390(\delta/\eta)Q_\infty$. The octave deficit and the steady-state gauge are the same number in different units. *Status:* DERIVED, cross-checked at a residual of 8.88×10^{-16} on the carried case — floating-point round-off, not measurement error. This is the result the sustaining-frequency paper carries forward [Curzi, 2026].

R4 — decay is not a lever on the deficit. D contains no δ. On the carried constructed dyad, moving δ from 0.09 to 0.05 leaves $D = +4.338420$ octaves unchanged to every printed digit while Q_∞ falls from -5.0119 to -9.0215, a worsening of 80 per cent. *Status:* DERIVED by inspection; COMPUTED on a constructed case. **This is the paper's sharpest result and it is negative.**

R5 — the ceiling. $c/\eta \le \ln N$, equivalently $f_{\min} \le kN^2$. At $N = 12$ the bound is $c/\eta \le 2.484907$, $f_{\min} \le 69.6133$ Hz, 7.169925 octaves above k. Above it no assignment of u over those N categories is sustainable. *Status:* DERIVED from $F \le \ln N$ [Curzi, 2026]. The same ceiling functions as the cheapest available detector of a repackaged order proxy in the corpus's survey work [Curzi, 2026i], whose Result Four records a single-week jump of 31.1 units against a maximum admissible increment of $\eta \ln N = 0.8071$ — a factor of 38.5.

R6 — the essence index. $f(n) = kn^2$ on a component count n, with the twelve-state ladder $n = 1 \ldots 12$ and the hold line $u^* = (e^{c/\eta} - 1)/(N - 1)$. *Status:* the hold line is DERIVED; **the law itself is POSTULATED** (§2.3 P-a) and has no external support. Result One of the essence law [Curzi, 2026k] — that $f(m)/f(n) = m^2/n^2$, so k cancels from every interval — is what keeps the postulate from contaminating any verdict.

R7 — the second invariance fails. Monotone rescalings of the elicited intensities, admissible if those intensities are merely ordinal, move u from 0.362115 to 0.525275 and the derived frequency from 12.0049 to 22.2093 Hz: a spread of 1065.05 cents, 0.888 octaves. *Status:* COMPUTED. The consequence is a requirement, not a caveat — **the intensities must carry ratio information or the frequency is not determined at all.**

CR — counter-result. The deficit is blind to decay, so the intervention that intuition reaches for first moves the instrument by exactly zero and worsens the destination whenever the numerator is negative. Separately: the reference constant k, the affine effective count, and the choice of twelve components are all **chosen rather than derived**, and any reading built on them inherits that status. The framework's own concentration work reaches the structurally parallel conclusion in a different domain [Curzi, 2026], where four jurisdictions at shares $40.9/27.3/18.2/13.6$ give a naive effective count of 3.41 and a bloc-collapsed effective count of exactly 1.00: a count that is not a count of independent things is not a count.

8 Discussion

8.1 What follows

The classification is worth more than any single sensitivity. The ceiling test costs one logarithm and changes the *kind* of answer a reading gives: below the ceiling a case is tunable and some assignment of u closes it; above it, no assignment does, and a prescription issued without the test may be unsatisfiable. Unsatisfiable prescriptions are worse than no prescription because they consume the effort that restructuring needed. This is the same discipline that the corpus's intervention work applies when it computes an N lever and then declines it net [Curzi, 2026j]: a ceiling gain of $\eta(\ln 6 - \ln 5) = +0.1914$ against five added communication paths costing $+0.0714$, a 34 per cent haircut on the headline benefit. The coordination cost of raising N is not a modelling nicety; it is the $n(n-1)/2$ growth that software engineering has documented for fifty years [Brooks, 1975, Conway, 1968].

The two gauges disagree, and both are right. Headroom uses the ceiling $\ln N$ and asks what the best possible geometry could support; the deficit uses the achieved F and asks what this geometry is doing. A system can have room to grow in principle and be shrinking in fact. The corpus's institutional work exhibits the same split in an operational setting [Curzi, 2026h], where moving η from 0.35 to 0.55 and δ from 0.25 to 0.12 takes a fixed point from 0.2653 to 3.9281 while a turnover event still costs 70 per cent of stock against 15.

The negative result generalises past this framework. The structure of D's independence from δ is that a multiplicative parameter cannot change the sign of the quantity it multiplies. Any model in which a rate of loss divides a net drive will have the property. Analogues are easy to find and worth naming: redundancy can reduce safety rather than increase it [Sagan, 1993]; interactive complexity with tight coupling makes catastrophic failure a system property rather than an operator error [Perrow, 1984]; and organisational learning that homogenises the organisation destroys the variance that produced its advantage [March, 1991]. In each case the intuitive intervention operates on the wrong term.

8.2 What does not follow

A reading does not conclude that a relation should end. $D > 0$ says a configuration is losing order at a computable rate under stated parameters. It does not weigh what the relation is for, what else it produces, or what ending it would cost, and the formalism contains no term for any of those.

A reading is not about a person. u is a property of two representations jointly; c and η are properties of an arrangement. There is no single-party term anywhere in the apparatus.

A reading is not clinical. The twelve states are not diagnoses, the ladder is not an assessment instrument, and the deficit is not a prognosis. This is stated in §5.9 and again in §12.

A reading does not establish that the surplus is monotone in u beyond the range examined. The framework's F is monotone increasing in u by construction, and the empirical literature is divided: bridging across structural holes has measured returns [Granovetter, 1973], but creative success rises and then *falls* with the same variable [Uzzi and Spiro, 2005], short-run effects of surface diversity are frequently negative [Putnam, 2007], and the rebuttal literature [Abascal and Baldassarri, 2015] qualifies rather than reverses that. **A monotone claim cannot survive that list unqualified**, and this paper does not make one outside the regime its arithmetic covers. The problem-solving result that is usually cited for diversity [Hong and Page, 2004] has a standing formal objection [Thompson, 2014] that must be cited with it.

8.3 Relation to the frequency register

The reason for expressing the criterion as a pitch is that intervals compose and hertz do not. Two relations with different costs are compared on the octave scale directly, and the exchange rate is a single exact constant. The apparatus for that — cents, octave equivalence, interval

arithmetic — is standard once an index exists [Rossing et al., 2002, von Helmholtz, 1885]. **What the apparatus does not carry with it is any claim about hearing.** Consonance depends on the spectrum of the sound and not on the ratio alone [Sethares, 2005]; interval categories and octave equivalence generalise less far across listeners and species than is usually assumed [Burns, 1999, Hoeschele et al., 2012]; and the octave's perceptual specialness, while real, has a measured and finite size [Demany et al., 2013]. This paper uses the arithmetic and asserts none of the psychoacoustics. The one place where the corpus's frequencies meet real acoustics is a derived room [Curzi, 2026l], whose Result Two fixes $7.2400 \times 4.3798 \times 2.9319$ m, $V = 92.969$ m^3 and an RT60 of 0.5 s requiring 29.94 m^2 sabins, with the folded set 233.98–386.74 Hz entirely above a Schroeder frequency of 146.7 Hz — and that is a statement about a room, not about a person.

8.4 The reflexive exposure

The framework's own foundations identify a self-fulfilling loop $T(x) = W(x, M(x))$ whose fixed point is a prophecy that makes itself true [Curzi, 2026]; this paper presupposes that result rather than restating it. The exposure is direct: an instrument that issues a verdict of DECAYS on a relation may change the relation by issuing it. The mechanism is documented [Merton, 1948, Rosenthal and Jacobson, 1968] and has been caught happening in the field, with dates and prices, as a theory became true by being used [MacKenzie and Millo, 2003]. But the effect sizes are contested [McCambridge et al., 2014] and, on the best long-run assessment of the classroom case, **self-fulfilling prophecies are real but small, do not typically accumulate, and often dissipate** — most expectations turn out to be accurate rather than causal [Jussim and Harber, 2005]. The reflexive exposure is therefore real and probably smaller than the framework's rhetoric implies, and the correct response is not to suppress the reading but to publish the parameters that produced it so that a third party can recompute it. That is also why any indicator adopted as a target degrades as an indicator [Campbell, 1979], and it is why the reporting requirements of §4.19 and §6.13 are requirements rather than recommendations.

8.5 What the surrounding programme allows this paper to claim

The title of the wider corpus places this work in a lineage that begins with *physique sociale* [Quetelet, 1835] and continues through the modern computational programme [Lazer et al., 2009, Pentland, 2014]. **The lineage's characteristic error is Quetelet's**: treating the mean of a distribution as a type rather than as a statistic. This paper inherits that risk directly, because it writes a single u, η and c for an arrangement, and it responds with a rule (no aggregation, SPEC 4.14) rather than with an argument.

Three constraints from that programme bear on the central object of this paper and are not comfortable. First, human activity is bursty, with heavy-tailed waiting times and **no characteristic time scale** [Barabási, 2005]; a *minimum sustainable frequency* is a claim about a characteristic scale and must confront that finding, which it can only do by insisting that its frequency is a position on a surplus scale and not a rate of any observed process. Strong spatial regularity coexisting with temporal burstiness [González et al., 2008] is the counterweight but not a rescue. Second, an explanation is supposed to name a mechanism rather than exhibit a fitted curve [Hedström and Ylikoski, 2010]; this paper fits nothing, which sidesteps the objection at the cost of having no empirical content to fit. Third, and most sharply: **what observation would distinguish this framework's laws from a neighbouring model that fits equally well?** [Sobkowicz, 2020, Durlauf, 2005] The honest answer is H2 — the ceiling is a strict universal and one counterexample kills it — and nothing else in this paper answers it.

What the paper is entitled to claim is therefore narrow and worth stating precisely. Prediction is one of many legitimate reasons to build a model, and explanation, bounding, and the disciplined exposure of one's own assumptions are others [Epstein, 2008]. The position this work is best placed to defend is the analytical one: mechanisms stated explicitly, parameters declared, and the arrangement within which the law is claimed named in advance

[Keuschnigg et al., 2018].

9 Verification

This section decomposes every result in the paper into atomic steps, each of which can be checked on its own. The standard is Occam's: **a postulate that can be dropped without the result changing has been dropped.** Every step carries one of four statuses.

- **ARITHMETIC** — reproducible by a reader with a calculator from the inputs named.
- **DEFINITION** — true by stipulation; nothing to verify but consistency of use.
- **CITED** — supported by a named external source.
- **POSTULATED** — asserted by this framework, with no external support.

9.1 The postulates

	postulate	status
P1	The Pythagorean whole tone is the frequency ratio $9{:}8$.	CITED [Barbour, 1951]
P2	φ is the positive root of $x^2 = x + 1$; hence $\varphi^2 = \varphi + 1$.	CITED [Euclid, 1956, Vajda, 1989]
P3	A dimensionless ratio of the form $(9/8)^2/\varphi^2$ **is a frequency in hertz**.	POSTULATED — no source
P4	The essence of a state with n distinguishable components has frequency kn^2.	POSTULATED — no source
P5	Stored order obeys $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$ with $\delta > 0$.	CITED as a form [Oppenheim et al., 1996, Elaydi, 2005]; the *interpretation* POSTULATED
P6	$N_{\text{eff}} = 1 + (N-1)u$ (the affine count axiom).	POSTULATED — design axiom [Curzi, 2026]; **not a Hill number**
P7	$F = \ln N_{\text{eff}}$, so $F \le \ln N$.	DEFINITION (given P6)
P8	A frequency may be divided by a power of two without changing what it means.	POSTULATED beyond pitch; CITED for pitch only [Shepard, 1982]
P9	n is a cardinality, hence absolute-scaled.	DEFINITION, with the scale typology CITED [Stevens, 1946]

Dropped as unnecessary. No postulate about the *ordering* of the twelve states is needed for R1–R5: the trichotomy, the octave form, the identity, the decay-blindness and the ceiling all hold for any N and any $u \in [0, 1]$ whatever the states are called. The ladder is needed only for R6 and R7. Likewise no postulate about consonance is needed anywhere in this paper.

9.2 The compounding chain

step	consumes	produces	status
L1	P1	$(9/8)^2 = 81/64 = 1.265625$ exactly	ARITHMETIC
L2	P2	$\varphi^2 = \varphi + 1 = 2.618033988749895$	ARITHMETIC
L3	L1, L2	$k = 0.483425732988414$	ARITHMETIC (exact)
L4	L3, **P3**	k carries the unit Hz	**POSTULATED**
L5	L4, **P4**	$f(n) = kn^2$ Hz	**POSTULATED**
L6	P5	$Q_\infty = (\eta F - c)/\delta$	ARITHMETIC
L7	P6, P7	$F = \ln(1 + (N-1)u)$, $F \le \ln N$	DEFINITION
L8	L5, L7	$f = k\,e^{2F}$ (the essence bridge)	ARITHMETIC given L5
L9	L6, L8, sign of $\eta F - c$	$f \gtreqless f_{\min}$ iff $Q_\infty \gtreqless 0$ — **R1**	ARITHMETIC
L10	L9	$\log_2(f_{\min}/k) = (2/\ln 2)(c/\eta)$ — **R2**	ARITHMETIC
L11	L10, Def. 3	$D = 2.885390(c/\eta - F)$	ARITHMETIC
L12	L11, L6	$D = -2.885390(\delta/\eta)Q_\infty$ — **R3**	ARITHMETIC
L13	L11 by inspection	$\partial D/\partial\delta = 0$ — **R4**	ARITHMETIC
L14	L7, L9	$c/\eta \le \ln N$, $f_{\min} \le kN^2$ — **R5**	ARITHMETIC
L15	L9 solved for u	$u^* = (e^{c/\eta} - 1)/(N-1)$ — part of **R6**	ARITHMETIC
L16	P9, L5	squaring n is a licensed operation	DEFINITION + CITED
L17	L8, rescalings	u: 0.362115–0.525275; f: 12.0049–22.2093 Hz — **R7**	ARITHMETIC
L18	L5, P8	folded values in 220–440 Hz, six pitch classes	**POSTULATED** (meaning); ARITHMETIC (value)

9.3 What this chain is and is not

Read the status column. Steps L1, L2, L3, L6, L9–L15 and L17 are arithmetic and a reader can redo every one of them. Steps L4, L5 and L18 are postulated. **A chain containing postulated steps is not a proof of its conclusion**, and this one is not offered as one. What it is offered as is a separation: **the tuning arithmetic and the octave algebra are solid, and the frequency interpretation sitting on top of them is not.** Anyone who accepts P3 and P4 gets R1 through R7 by arithmetic alone. Anyone who does not accept P3 and P4 still gets R1 through R5 in *surplus* units, because the deficit is a difference of logarithms in which k cancels: **the sustainability verdict never depends on the postulated constant, only the hertz values do.** That is the honest statement of what the frequency register buys and what it costs.

Two statuses that must not be misread. The exponent 2 at L5 is **DEFINED**, not estimated: no fit was performed and the standard procedure for testing a scaling claim [Clauset et al., 2009] was not run. The logarithm at L7 is **DEFINITION** given the affine axiom P6, and the warrant the argument needs for it is the multiplicative-dynamics one [Peters and Gell-Mann, 2016] rather than the psychophysical one; the psychophysical warrant is used only at §6.3, for the elicitation.

The falsifiable content lives at L6 and L7, not at L4 and L5. H1 through H3 test the storage law's fixed point and its sign; none of them requires the essence bridge to be true. This paper's empirical exposure is therefore smaller than its notation suggests, and its unsupported material is correspondingly more inert. The directional-prediction critique [Meehl, 1990] is worth meeting explicitly here: the corpus makes **point** predictions with constants rather than directional ones, which is the right response to Meehl, and it makes them from postulates whose status is printed — which is the honest description of what a badged postulate does in a research programme [Lakatos, 1970, Platt, 1964].

10 Limitations

Every check the governing rigour standard specifies is listed here by name, whether or not this paper passes it.

L-1. The dimensional audit does not close. $f(n) = kn^2$ balances only if k carries $[T^{-1}]$, and k is computed as a ratio of two dimensionless quantities. **No operational procedure is supplied by which the constant acquires the dimension of frequency.** The hertz is attached by assertion (§2.3 P-b, §3.4). This is the paper's largest structural weakness and no citation in the bibliography repairs it.

L-2. The essence law has no external support. A quadratic relation between a dimensionless count and a frequency has no established analogue in psychoacoustics, ecology or measurement theory. The corpus badges it a postulate; that badge is the honest position and is not a defence.

L-3. The constant is chosen. $k = (9/8)^2/\varphi^2$ combines two real objects into a product with no precedent. It is quoted to fifteen significant figures **because it is defined**, not because anything was measured to that precision. Golden-section claims have a deflationary literature that this paper cites rather than avoids [Markowsky, 1992].

L-4. Octave reduction beyond pitch is unwarranted. Shepard grounds octave equivalence for pitch [Shepard, 1982]; nothing grounds folding a non-acoustic quantity into a 220–440 Hz band and reading meaning from where it lands. The operation is well defined; its interpretation is not supported.

L-5. $N_{\text{eff}} = 1 + (N-1)u$ **is not a Hill number.** It is a linear interpolation between 1 and N and appears in no ecological, economic, political-science or information-theoretic literature. The Hill-number citations in this paper apply to $1/\text{HHI}$ and $\exp(H)$ and to nothing else.

L-6. Conjoint measurement was not tested. The ratio-scale treatment of the elicited intensities, of η, of c and of δ is provisional. Michell's objection [Michell, 1997] stands unanswered for every quantity in this paper except n, and the available remedies [Luce and Tukey, 1964, Rasch, 1960] were not run. **This is one instance of a pattern and not an isolated exposure**; L-18 states the pattern once and concedes both instances together.

L-7. FVP-7 has not been run. Cross-individual reproducibility of the component count — H4, the gate that decides whether n is a property of the episode or of the analyst — is **not tested**. Every frequency in this paper is a frequency *under a declared decomposition*, which is a real thing to have and is less than a measured one. The demonstration in §6.17 is an illustration of the protocol on constructed data; constructed data cannot reproduce anything.

L-8. Twelve is not justified. The ladder's category count rests on a failed search for a thirteenth component, which is weak evidence. Twelve also exceeds the classical span of immediate absolute judgment [Miller, 1956], which is a caution rather than a refutation.

L-9. Agreement is not correctness. The agreement protocol measures reliability, not validity [Cronbach and Meehl, 1955, Campbell and Fiske, 1959], and internal-consistency coefficients do not substitute for either. A held-out predictive test is required before the ladder can be said to work and none has been run [Nosek et al., 2018].

L-10. u **is symmetric and psychological similarity is not.** Similarity judgments are asymmetric and violate the triangle inequality [Tversky, 1977]. If u is read as a similarity measure over people, it cannot assume metric behaviour, and this framework's u does.

L-11. Elicitation is exposed to instrument effects. Magnitude estimation's exponents move with the experimenter's chosen stimulus range [Teghtsoonian, 1971], response biases are well catalogued [Poulton, 1968], and survey instruments generally are exposed to question wording and focusing effects. None of these were controlled, because no elicitation was run.

L-12. Every case is constructed. The dyad, the commons and the subject are chosen parameter sets. Their precision is exact because the parameters are exact; **that precision does not transfer to a measured case**, and reporting $f_{\min} = 30.1588$ Hz from an η known to one significant figure would be false precision. The rigour standard's rule applies: measured or fitted quantities require an uncertainty, and none of this paper's numbers are measured.

L-13. No model comparison was performed. No out-of-sample test, no matched null with the same number of free parameters, and no information criterion is reported, because nothing was fitted. Where the corpus has attempted a confirmatory fit elsewhere, it failed its own baseline test [Curzi, 2026i], and that result is the relevant precedent for how a fit here should be reported.

L-14. No statistical claims are made and none are supported. There are no p-values, no effect sizes and no confidence intervals in this paper, because there is no sample. Nothing here should be read as evidence for or against any empirical proposition.

L-15. A published reading can change what it reads. The self-fulfilling exposure of §8.4 is not controlled for and cannot be controlled for by an instrument acting alone.

L-16. Boundary behaviour is only partly examined. $\delta \to 0$ is a pathology and the paper says so; $u \to 1$ is examined and gives the ceiling; $\eta \to 0$ makes c/η diverge and therefore places every relation above every ceiling, which is a real constraint on the model's domain and is **not** examined in the source. $N \to \infty$ raises the ceiling without bound, which the coordination-cost literature says is not available in practice [Brooks, 1975]. These limits constrain the domain; they are stated here because the source does not state them all.

L-17. Two attractive simplifications are refused. The corpus's channel-separation work [Curzi, 2026] makes the general form of the point: a property that is necessary is not thereby sufficient, and a control plane trusted *because* it is separate is the failure mode. The same applies to this paper's instrument. A relation is not sustainable because its deficit is negative; the deficit is a statement about four parameters over one declared period and nothing else.

L-18. Scalar reduction is a habit of this framework, and it is one objection in two domains. This entry replaces an earlier one recording that the seven works named in §5.7 were absent from the bibliographic database governing this corpus. They are now present and are cited where they are used [Russell, 1980, Smith and Ellsworth, 1985, Lazarus, 1991, Scherer, 2001, Tracy and Robins, 2004, Premack and Woodruff, 1978, Frijda, 1986], and citing them properly is what makes the real limitation visible, because **four of the seven are in tension with the index they were named to motivate.**

The concession, in one statement. Wherever this framework meets a phenomenon that its own source literature describes by several jointly varying features, it replaces the phenomenon with a single number and proceeds arithmetically on the number. That happens twice in this paper. Comfort and the elicited intensities are treated as quantitative by stipulation, which is Michell's objection and is conceded at L-6 and L-7. An emotional state is reduced to a count of components, ordered by that count and squared, which is the same objection in the affect domain and is conceded here. **They are one limitation of the framework, not two independent problems**, and the framework does not hold that having named the pattern disposes of it.

What the affect literature actually licenses. §5.7 states the objection at full strength and it is not repeated here in full. In summary: Smith and Ellsworth [1985] individuate emotions by the pattern across simultaneously and independently varying appraisal dimensions, which a count cannot represent — a count cannot separate anger from guilt; Scherer [2001] treats emotion as a process in time whose state *is* the synchronisation across five subsystems, which resists scalar reduction of any kind; and Lazarus [1991] individuates by qualitative core relational theme, so that differing by *what* is appraised does not license differing by *how many*. **None of the three is cited anywhere in this paper as support for the index, its ordering or its**

nesting, and a reader who finds one so used should treat it as an error. Russell [1980] agrees with rung `n = 2` on an integer and on nothing else. Premack and Woodruff [1978] name its component cleanly, with the cautions recorded at `n = 8`. Tracy and Robins [2004] is the only entry that supports its rung in the form this framework needs.

The ordering is a stipulation and is not badged as a postulate in §2.3. The rungs' ordering and their nesting are fixed by the construction of §5.6 and have no external support. §2.3 lists three postulates and none of them covers the ladder's ordering; that is an omission in this paper's own badging, it is stated here rather than repaired silently, and any later revision of §2.3 should carry the ordering as a fourth postulate.

One rung is contradicted by the work cited for it. Frijda [1986] holds that changed action readiness is constitutive of emotion as such, which places that component at the floor of the ladder and not at `n = 12`. The rung has been narrowed accordingly, to **bound commitment as distinct from readiness**, and that is the only version of the rung Frijda leaves standing. The wider claim the source manuscript made is withdrawn. **If Frijda is right and the narrower reading does not hold up, the ladder's ordering fails at its summit** — and because the index is squared, a failure at the summit is the cheapest place on the ladder for one to happen (§5.4), which is a mitigation of the consequence and not of the objection.

One reference could not be resolved. The source manuscript names "Lewis" alongside Tracy and Robins at `n = 7`, by surname only, with no year and no title. Several works by Michael Lewis would fit and none was guessed at. That citation is missing rather than wrong, and it is the one bibliographic gap in §5.7 that remains open.

What is not damaged. The count argument of §5.2 is untouched by any of this. It establishes that a cardinality is absolute-scaled, which is a fact about counting, and it never claimed that the thing counted was correctly individuated. **The objection above is about the individuation, not the arithmetic**, and the two should not be conflated in either direction.

L-19. The 2.885390 shorthand is a rounded constant. The exact value is $2/\ln 2 = 2.885390081777927$. Where the text writes 2.885390 inside a formula it is a shorthand for the exact constant; every computed figure in the paper uses the exact value.

L-20. The reflexive specification is not identified. Stated in §3.6 and repeated here because it is the most serious methodological problem in the corpus: the endogenous social effect in a linear-in-means reflexive loop cannot be recovered from observational data [Manski, 1993, Shalizi and Thomas, 2011]. The remedy adopted is stipulation of the reference set [Brock and Durlauf, 2001], and the claim is accordingly a modelling position rather than an identified result.

L-21. The framework exempts its own parameters from its own reflexivity. If subjects change when modelled, δ, η and u are not structural [Lucas, 1976, Hacking, 1995]. This is not resolved.

L-22. The exponent is defined, not estimated. The exponent 2 in $f(n) = kn^2$ was not fitted, was not tested against alternatives, and the standard fitting-and-testing procedure for a scaling claim [Clauset et al., 2009] was not run because there is no data to run it on. **This paper makes no empirical scaling claim.**

L-23. The functional form is underdetermined. A power law with exponent below one [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the qualitative facts as well as this framework's logarithm, and many mechanisms produce indistinguishable curves [Stumpf and Porter, 2012]. Nothing in this paper discriminates among them.

L-24. There is no conservation law. Q is written with the grammar of a conserved quantity and no conservation law is supplied, which is the characteristic defect of the genre [Mirowski, 1989]. The paper claims validity only within a specified arrangement [Cartwright, 1999] and claims no universality.

L-25. Ergodicity is assumed and not argued. Q accumulates over time for one relation from parameters that would naturally be estimated across relations; the time and ensemble averages of a multiplicative process are not the same [Peters, 2019], and this paper does not establish that they coincide here.

A discrepancy found in the source and left standing. The source reports the carried dyad's deficit as $D = +4.338420$ octaves from a reported $u = 0.068735$. Recomputing D from the *rounded* u gives $+4.338414$; recomputing from the underlying state vectors gives $+4.338420197$, which is the figure the source quotes. **The two disagree in the sixth decimal place because the printed u is rounded, not because either is wrong**, and this paper quotes the vector-derived value throughout. It is recorded here because a reader recomputing from the printed intermediate will find the discrepancy and should know its cause.

11 Conclusion

What is established. Given the storage law and the surplus functional, the sign of the steady state is equivalent to a threshold comparison, and in logarithmic units that comparison has an exact linear form with the constant $2/\ln 2 = 2.885390081777927$. The octave deficit and the steady state are algebraically the same quantity, verified to floating-point round-off. The deficit is independent of the decay rate, so an intervention acting only on decay changes the deficit by exactly zero and multiplies the existing sign of the steady state — worsening a failing relation. The ceiling $c/\eta \leq \ln N$ classifies a case as tunable or structural before any prescription is issued. **A count admits only the identity transformation, so squaring it is a licensed operation in a way that squaring a rating is not**, and monotone rescalings of ordinal intensities move a derived frequency by 1065.05 cents, which settles that the intensities cannot be ordinal. All of this is arithmetic on stated definitions and none of it depends on the value of k.

What is proposed. That a dimensionless ratio of a tuning interval to the square of the golden section is a frequency; that the essence of an n-component state is kn^2; and that octave equivalence may be carried from pitch to a quantity that is not a pitch. **These three are postulates of this framework with no external support**, they are labelled as such at every appearance, and they are separable from everything in the preceding paragraph. A reader who rejects all three keeps the sustainability criterion, the trichotomy, the identity, the negative result and the ceiling, expressed in nats instead of octaves. **That separability is the strongest claim this paper makes about its own foundations.**

What is not claimed at all. No person, bond, tissue, molecule, room or institution is claimed to resonate at any frequency named here. The twelve states are not diagnoses and no reading in this paper is clinical, diagnostic, therapeutic or prognostic.

12 Declarations

Data and code availability. No dataset was collected for this study and none exists. Every number in the paper is the output of exact arithmetic on parameters stated in the text; every case is constructed and is labelled as constructed at the point of use. The figures are generated by the Python scripts named in their captions, which take their inputs from the tables in the text and are available from the author on request. **No empirical measurement of any human being, any relation or any institution underlies any figure, table or number in this paper.** A frequency law reads as though something were measured; nothing was.

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the analysis, the figures and the manuscript.

Ethics. No human-subjects research was conducted and no personal data were processed. The constructed subject and the constructed dyad are invented parameter sets and are not derived from, and do not describe, any person. This paper asserts several perceptual claims — that magnitude estimation delivers ratio-scale intensities, that octave equivalence is a usable structure, and that interval arithmetic in cents is the correct apparatus — **without any perceptual experiment behind them.** Those claims rest entirely on the cited literature, they

are limited by that literature's own boundaries (§8.3), and they are listed among this paper's limitations (L-4, L-6, L-11).

Clinical disclaimer. The twelve state classes are not diagnostic categories. Nothing in this paper is an assessment, a screening instrument, a treatment or advice about any person's health, and nothing in it may be presented as such.

Chapter 42

Octave Folding and the Odd-Part Class Rule: The Frequency of a Mood, of a Trait, and the Self as a Chord

Abstract. Applying the essence law $f(n) = kn^2$ to a component count generates a pitch lattice whose structure we characterise exactly, with $k = (9/8)^2/\varphi^2 = 0.483425732988414$ a *defined* constant, exact by construction, whose interpretation as a frequency is postulated and not derived. Two indices fold to the same pitch class if and only if they share the same odd part, so an N-long ladder admits exactly $\lceil N/2 \rceil$ classes: six at $N = 12$, namely $\{1, 3, 5, 7, 9, 11\}$, and class-preserving intervals are necessarily *even* whole octaves. This partitions transitions into intensification, which preserves class, and transformation, which does not; only 10 of 66 unordered pairs (15.1515 per cent) are class-preserving, and index distance does not determine which kind a move is. The tightest folded separation at $N = 12$ is $50/49 = 34.9756$ cents, attained by the single pair $7{:}10$, and $N = 12$ sits below the first collision cliff at $N = 17$ ($289/288 = 6.0008$ cents). The law contains no time term, so a mood and a trait at $n = 9$ return the identical 39.1575 Hz, 313.260 Hz folded, class 9, and are distinguished only by δ, whose half-lives differ by a factor of 17.65 (1.9434 against 34.3096 cycles); the naive $1/\delta$ overstates the half-life by 71.52 per cent at $\delta = 0.30$. At internal scope the chord width $F_{\text{int}} = \ln(1 + (N-1)u_{\text{int}})$ is degenerate under one-hot encoding: distinct one-hot framings are orthogonal by construction and score the ceiling $\ln 12 = 2.484907$ for every subject regardless of content.

Keywords: octave folding; odd part; pitch class; intensification; storage half-life; trait window; chord width; one-hot degeneracy

Series: C. Frequency, Essence, and Mind **Paper:** CCP-042

1 Introduction

A framework that assigns a psychological state a position on a logarithmic scale acquires, whether or not it intends to, an arithmetic of *relations* between states. This paper characterises that arithmetic exactly for the corpus's essence law $f(n) = kn^2$ applied to a component count n, and reports three structural results and one negative one.

The structural results are about the lattice. Squaring an integer index and folding by octaves makes the pitch class of a state depend on **the odd part of its index alone**, so two states share a class if and only if they share an odd part, and an index of length N realises exactly $\lceil N/2 \rceil$ classes — six at $N = 12$. That partitions transitions into two kinds: *intensification*, which

preserves the class and moves a whole and necessarily **even** number of octaves, and *transformation*, which does neither. Only 10 of the 66 unordered pairs at $N = 12$ (15.1515 per cent) are class-preserving, and the partition is read off the indices before anyone has decided what the states mean. The tightest folded separation at $N = 12$ is $50/49 = 34.9756$ cents, attained by the single pair 7:10, and $N = 12$ sits below a collision cliff at $N = 17$ where the minimum collapses to $289/288 = 6.0008$ cents.

The negative result is about time, and it is the one this paper most wants read. The essence law has no time argument. **A mood and a standing disposition at the same index are therefore the same tone** — the same raw frequency, the same transposition, the same folded value, the same pitch class — and nothing in the frequency register distinguishes them. What distinguishes them is the decay rate δ of the corpus's storage law [Curzi, 2026], whose Result Two is the identity $D = 2.885390(c/\eta - F) = -2.885390(\delta/\eta)Q_\infty$ linking the two registers. At $n = 9$ both readings give 39.1575 Hz, 313.260 Hz folded, class 9, and their half-lives differ by a factor of 17.65.

A second negative result concerns the encoding. Extending the surplus functional inside one observer gives a *chord width* $F_{\text{int}} = \ln(1+(N-1)u_{\text{int}})$, and the natural tidy encoding — one state per framing — makes every pair of framings orthogonal by construction. Every subject then scores the ceiling $\ln 12 = 2.484907$ regardless of content. **An instrument whose failure mode is to award everyone full marks will be adopted enthusiastically and never questioned**, so the degeneracy is set out in the body rather than the failure modes, and the degenerate calculation is shown before the correct one.

What this paper does not claim, stated at the front. The constant $k = (9/8)^2/\varphi^2 = 0.483425732988414$ is a **defined** number, exact by construction; its two components are documented — the Pythagorean whole tone [Barbour, 1951] and Euclid's extreme and mean ratio [Euclid, 1956] — but **their product expressed as a frequency in hertz has no precedent in any literature we have found**, and the standard deflationary treatment of golden-section claims [Markowsky, 1992] is cited here deliberately and again wherever φ appears. The essence law is a postulate and **its exponent is defined by the model, not estimated from data**. Octave equivalence is grounded for pitch [Shepard, 1982] and its extension to a quantity that is not a pitch is a modelling choice with no external warrant. Section 11 decomposes each of these into steps a reader can check; Section 12 prints every check this paper fails.

Contributions

1. The odd-part class rule, proved in both directions, with the corollary that an index $1 \dots N$ realises exactly $\lceil N/2 \rceil$ pitch classes — verified for every N from 1 to 25.
2. The corollary that class-preserving intervals are **even** whole octaves, so the smallest one is exactly 2400 cents.
3. The intensification/transformation dichotomy, verified exhaustively over all 66 pairs at $N = 12$, with the finding that index distance does not determine which kind a move is.
4. The regime statement: $50/49 = 34.9756$ cents at $N = 12$, the cliff at $N = 17$ ($289/288 = 6.0008$ cents) and the second at $N = 41$ ($1682/1681 = 1.0296$ cents), all exact rational coincidences.
5. The theorem that the frequency law cannot distinguish a disposition from an episode, with the storage half-life $t_{1/2} = \ln 2/(-\ln(1-\delta))$ supplied as the quantity that can — and the measured gap between it and the naive $1/\delta$.
6. The one-hot degeneracy theorem and the overlap condition it forces, with a worked demonstration that four unrelated constructed subjects produce the identical ceiling reading.
7. A drop-one audit in which two of three removals *raise* the chord width, establishing that mean pairwise non-parallelism does not determine the joint geometry.

2 Hypotheses, definitions, and postulates

Series C carries the corpus's least-supported material, so the labelling is done first and in full. **A hypothesis nothing could falsify is not a hypothesis**, and this section does not manufacture one.

2.1 Hypotheses proper

H1 — the mood/trait separation hypothesis. Moment-to-moment affective states and standing dispositions are empirically distinguishable, and the distinguishing quantity is a retention rate rather than any property of the state's specification. **Predicted:** two states assigned the same component count by the declared decomposition, one reported as transient and one as standing, differ in measured δ and in nothing else the essence law can see. **Falsified by:** a measured difference in the component count that tracks transience, which would make the frequency register informative about persistence after all. **Status:** the *formal* half — that $f(n) = kn^2$ contains no time argument — is DERIVED and is Theorem 6.1. The *empirical* half is **not tested here.** No observational data were collected for this paper.

H2 — the count-reproducibility hypothesis. The component count n assigned to a described state under the declared decomposition rules is a property of the state and not of the analyst. **Predicted:** independent analysts applying the published rules to the same descriptions agree on n, and the residual disagreement carries no analyst main effect. **Falsified by:** a systematic analyst effect on n, or agreement on the derived pitch class no better than the ladder's marginal distribution. **Status: not run.** It is reported as not run in Section 7.11, in Section 12, and here. The whole apparatus of this paper is conditional on it.

H3 — the overlap hypothesis. Two framings of one situation, elicited from one observer, share components; two descriptions with no component in common are descriptions of different situations. **Predicted:** framings elicited about a single declared situation have non-empty shared support at a rate far above what independent assignment of components would produce. **Falsified by:** elicited framings of one situation that are routinely disjoint, which would make the overlap gate of Section 7.5 unsatisfiable in practice and the internal surplus uncomputable. **Status: not tested here.** The overlap condition is imposed as a validity requirement, not demonstrated as an empirical regularity.

H4 — the decay-identifiability hypothesis. A retention rate is recoverable from observation only over a no-input interval spanning at least three expected half-lives. **Predicted:** estimates of δ from shorter windows are indistinguishable from a linear decline and carry no information about the rate. **Status:** DERIVED from the geometry of $(1-\delta)^t$; **not tested against observation.** Its practical consequence — that a trait-grade $\delta = 0.02$ requires roughly 103 cycles of observation — is computed in Section 6.13 and is the binding constraint on the whole method.

2.2 Definitions — true by stipulation, not falsifiable

- **Def. 1.** n := the number of distinguishable components required to specify a state; $N = 12$ is the declared category count.
- **Def. 2.** The *odd part* of n is the unique odd m with $n = 2^a m$, $a \geq 0$.
- **Def. 3.** The *pitch class* of a state is the equivalence class of $f(n)$ under multiplication by integer powers of two.
- **Def. 4.** $k := (9/8)^2/\varphi^2$. **This is a defined constant.** It is an exact ratio of exact quantities and is therefore computable to arbitrary precision; the fifteen significant figures quoted

throughout, 0.483425732988414, are *exact digits of a definition* and carry no measurement uncertainty because there is no measurement.

- **Def. 5.** A *trait* is a state held at low retention loss δ; a *mood* is the same state held at high δ. "Low" is fixed by a declared observation window in Section 6.6 and is a convention, not a discovery.
- **Def. 6.** A *chord* is the set of indices carrying nonzero intensity across the framings an observer holds at one moment; its *width* is $f = k\, e^{2F_{\text{int}}}$.

2.3 Postulates — asserted by this framework, with no external support

- **P-a. The essence law $f(n) = kn^2$.** A quadratic relation between a dimensionless count and a frequency. **No external literature supports it.** It is not derived here and nothing in the bibliography derives it. **The exponent 2 is defined by the model, not estimated from data.** No scaling exponent in this paper was obtained by fitting; none was tested against a log-normal or exponential alternative; and the goodness-of-fit and likelihood-ratio procedure a claimed scaling relation requires [Clauset et al., 2009] was **not run, because there is no data to run it on.**
- **P-b. The attachment of hertz to k.** Computing $(9/8)^2/\varphi^2 = 0.483425732988414$ is arithmetic. **Asserting that this dimensionless number *is a frequency* is a separate and much stronger claim**, and the source supplies no operational procedure by which the constant acquires the dimension of inverse time. This is postulated, and Section 3.4 states it again in the dimensional audit.
- **P-c. The affine effective count $N_{\text{eff}} = 1 + (N-1)u$.** A design axiom [Curzi, 2026], not a derived result and **not a Hill number of any order.** The Hill family [Hill, 1973, Jost, 2006] and its independent rediscovery in political science [Laakso and Taagepera, 1979] ground $1/\text{HHI}$ and $\exp(H)$ exactly — as the corpus's concentration work uses them [Curzi, 2026] — and ground this construction not at all.
- **P-d. The framing encoding.** That a way of construing a situation admits representation as a unit vector over twelve state categories, such that framings making the same predictions map to parallel vectors. **This is asserted, not established.**
- **P-e. The internal roster reduction.** That the mean pairwise non-parallelism over an observer's own framings may be read as an audited framing count through $N_{\text{eff}} = 1 + (N-1)u$. **Asserted; and Section 7.8 shows what it fails to determine.**

2.4 Modelling choices — legitimate, but choices

- **M-a. Octave reduction applied to a non-acoustic quantity.** Octave equivalence is a perceptual structure formalised for **pitch** [Shepard, 1982], and the octave's perceptual specialness has a measured and finite size [Demany et al., 2013]. Nothing licenses carrying it to a quantity that is not a pitch. **The *operation* — division by a power of two — is well defined; its *interpretation* is a choice of this framework**, and every folding result in this paper is an arithmetic result about the operation and not a perceptual claim.
- **M-b. The band 220–440 Hz.** One octave, closed below and open above, so the transposition exponent j is unique for every f. **The band is arbitrary and stated. The operation is not arbitrary**, and Section 4.2 shows that no choice of band can create or remove a class collision.
- **M-c. $N = 12$.** Twelve exceeds the classical span of immediate absolute judgment [Miller, 1956], which is a reason for caution and not a reason for twelve. The ladder is owned by the criterion paper [Curzi, 2026], whose Result Six is the count-based essence index with its twelve-state ladder.
- **M-d. The trait window W.** The cut between "mood" and "trait" is a declared convention on the half-life. **There is no δ at which something switches from being an episode to being a disposition**; there is a continuous rate and a cut somebody chose.

- **M-e. Interval apparatus.** Cents, octave equivalence and interval arithmetic are the correct apparatus once an index exists [Rossing et al., 2002, Benson, 2006], but what a listener finds consonant depends on the spectrum of the sound and not on the ratio alone [Sethares, 2005], and interval categories generalise less far across listeners than is usually assumed [Burns, 1999, Hoeschele et al., 2012]. **This paper uses the arithmetic and claims none of the perception.**

2.5 What is expository

Sections 5.6, 6.13, 7.10 and 7.11 are procedural: they set out reporting requirements, gates and failure modes rather than argue a claim. Their organising question is *what must be published alongside a reading for the reading to be checkable by a third party*.

3 Framework, dimensional audit, and method

3.1 The three laws this paper composes

```
THE ESSENCE LAW    f(n) = k n^2                      an index has a pitch
THE SURPLUS LAW    F(u) = ln(1 + (N−1)u)             a pair of states has a surplus
THE STORAGE LAW    Q_{t+1} = (1−delta)Q_t + eta S − C   surplus becomes structure, and leaks
```

The first is a postulate of this framework (P-a). The second is a theorem given four axioms, of which the third — that the effective count is affine in the non-parallelism — is a design choice [Curzi, 2026], whose counter-result is that the surplus is bounded twice over: it is only a lower bound on the ensemble's information, and its logarithmic concavity means three well-separated views capture most of what a large panel could. **That bound governs every width computed in Section 7.** The third is a first-order linear difference equation whose form is shared with the RC low-pass filter [Oppenheim et al., 1996, Horowitz and Hill, 2015], with exponential smoothing [Holt, 1957], and with the leaky integrate-and-fire neuron [Lapicque, 1907, Gerstner and Kistler, 2002]; its fixed point and stability condition are standard [Elaydi, 2005] and the convergence is Banach's [Banach, 1922]. **The corpus states that theorem exactly and then refuses the certificate it appears to offer** [Curzi, 2026], on the ground that a human life is not a metric space and nobody has measured the relevant contraction modulus, so the theorem transfers as shape rather than guarantee. **That refusal binds every half-life in Section 6.**

Analogy status: category 2, a shared mathematical form. The storage law is used because the mathematics is the same, not because the physics transfers. No claim is made that stored order is a charge or that any physical result about RC circuits applies to a person. **And Q has the grammar of a conserved quantity without a conservation law to give that grammar content**, which is the characteristic defect of the genre [Mirowski, 1989]; validity is claimed only within a specified arrangement [Cartwright, 1999].

3.2 The warrant the logarithm needs

The corpus habitually justifies its logarithms by appeal to Fechner [Fechner, 1860, Weber, 1834]. A stronger and cheaper warrant is available: a logarithmic form follows from multiplicative dynamics with **no psychological assumption at all** [Peters and Gell-Mann, 2016]. The arguments of this paper need only the second. **The functional form is nonetheless underdetermined by the evidence**: a power law with exponent below one [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the qualitative facts as well, and many mechanisms produce indistinguishable curves [Stumpf and Porter, 2012]. **Nothing in this paper discriminates among them.**

3.3 The constant k

$$k := \frac{(9/8)^2}{\varphi^2} = \frac{81/64}{\varphi + 1} = 0.483425732988414 \quad \text{(defined; exact)}$$

Computed: $\varphi = 1.618033988749895$, $\varphi^2 = \varphi + 1 = 2.618033988749895$, $(9/8)^2 = 1.265625$, and the two forms agree to every digit. Both components are real and documented — the Pythagorean whole tone and ditone [Barbour, 1951, Lindley and Turner-Smith, 1993] and Euclid's extreme and mean ratio [Euclid, 1956], with $\varphi^2 = \varphi + 1$ a proved identity [Vajda, 1989, Koshy, 2001].

Their product, expressed as a frequency in hertz, has no precedent. The literature on golden-section claims is overwhelmingly one of *deflation*: the appearances of φ in the Parthenon, the Great Pyramid, the human body and preferred rectangles do not survive examination [Markowsky, 1992, Fischler, 1981, Livio, 2002], and the psychological preference results are weaker and more equivocal than usually reported [Höge, 1997]. **This framework's position is the one those critiques leave open and no wider: φ is being *chosen* as a constant, not *found* in a psychological system.** And the cost of the choice is small in a precise sense: k **cancels from every interval in this paper**, because every interval is a ratio $f(m)/f(n) = m^2/n^2$.

3.4 Dimensional audit

symbol	quantity	dimension	unit	status
n, N	component count, category count	$[1]$	—	absolute scale, a count
m	odd part of n	$[1]$	—	absolute scale, a count
k	reference constant	$[T^{-1}]$ **asserted**	Hz **asserted**	**defined as a ratio (dimensionless); the unit is postulated**
f	essence frequency	$[T^{-1}]$	Hz	inherits its dimension entirely from k
j	transposition exponent	$[1]$	—	an integer
u	non-parallelism	$[1]$	—	dimensionless by construction, $u \in [0, 1]$
F_{int}	internal surplus	$[1]$	nats	dimensionless; a logarithm of a count
interval	pitch interval	$[1]$	cents, octaves	a ratio of logarithms; k cancels
Q	stored order	$[\mathcal{C}]$	none supplied	state variable; **no operational unit**
δ	retention loss	$[1]$	per declared cycle	dimensionless; $0 < \delta \leq 1$
$t_{1/2}$	storage half-life	$[T]$	declared cycles	derived from δ
η, c	conversion, cost	$[1]$, $[\mathcal{C}T^{-1}]$	per declared period	see below

Does the equation balance? The essence law $f(n) = kn^2$ balances only if k carries $[T^{-1}]$, because n is a dimensionless count. **The source supplies no operational procedure by which k acquires that dimension**; it is computed as a ratio of two dimensionless quantities and the hertz is attached by assertion. **This is the largest dimensional exposure in the paper and it is listed in Section 12 by name.** Every *interval* is dimensionless and free of k, which is why Sections 4 and 5 are secure in a way Section 4.5's absolute values are not.

Two further audit results. δ is dimensionless but $t_{1/2}$ **is a time and therefore has no meaning without a declared cycle**; Section 6.5 makes the rescaling exact and shows that a δ quoted without its cycle is not a number. And the ratio c/η, which appears in the threshold this paper cites from the join [Curzi, 2026], is $[\mathcal{C}T^{-1}]$ divided by a pure number and is **not dimensionless as written**; the framework's convention is that c is expressed in nats of surplus per declared period, **a convention the source never states**, and every reading that uses the threshold depends on it.

3.5 Scale type, and what licenses the arithmetic

quantity	scale type claimed	what licenses it
n, N, odd part	absolute	cardinality: true zero, fixed unit, identity transformation group only
u	ratio	constructed from an inner product of unit vectors; zero is parallel
elicited intensities	ratio, up to a free common multiplier	**magnitude estimation, assumed** [Stevens, 1957]; not tested here
δ	ratio, **assumed**	no validated instrument is supplied; Section 6.13 says so
f, folded f	ratio in hertz, **inherited**	inherits the status of n and of the postulated bridge
intervals, $t_{1/2}$	ratio	ratios of exact quantities given the inputs

Stevens' typology makes the permissible arithmetic on a quantity a function of the transformation group that leaves its empirical content unchanged [Stevens, 1946]. **A count admits only the identity transformation, which is what licenses squaring it.** That argument covers n and covers nothing else in this paper.

The strongest available objection to this framework is Michell's: that psychology routinely assumes its attributes are quantitative without testing whether they are [Michell, 1997, 1999]. The remedy exists — conjoint measurement supplies a testable condition for additive quantitative structure [Luce and Tukey, 1964, Krantz et al., 1971] and a Rasch model supplies specific objectivity [Rasch, 1960, Andrich, 1988] — **and neither has been run for any quantity in this paper.** The elicited intensities matter here more than anywhere else in the corpus, because ordinal-preserving transformations of them move u from 0.362115 to 0.525275 on the carried Subject's closest pair, a spread of 1065.05 cents [Curzi, 2026], whose Result Seven records exactly that failure of invariance.

3.6 Identification, and the reflexive exposure

The framework's founding construction is a reflexive loop in which the group influences the member and the members constitute the group. **That is a linear-in-means reflexive specification, and its endogenous social effect is not identified from observational data** [Manski, 1993];

influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011]. **No amount of fit recovers the parameter.** The remedy this paper adopts is stipulation — the category set and the decomposition rules are declared in advance and published with every reading, which is what makes a discrete-choice treatment of social interaction defensible [Brock and Durlauf, 2001] — and **the claim is accordingly a modelling position, not an identified empirical result.**

The same objection applies to the framework's own parameters. If subjects change when modelled, δ and u are not structural [Lucas, 1976, Hacking, 1995]. **The framework asserts reflexivity for its subject matter and cannot consistently exempt its own parameters from it.** This is stated, not resolved.

3.7 Method

Every number in this paper is the output of exact arithmetic on stated parameters, computed in double precision, with the pitch-class and interval results computed in exact rational arithmetic where the text says so. **No observational data were collected, no human subjects participated, and no empirical dataset underlies any figure or table.** Every case is constructed and is labelled as constructed at the point of use. The figures are produced by the scripts named in their captions.

4 Octave folding and the odd-part class rule

4.1 The law, and the index it is applied to

$$f(n) = k\,n^2, \qquad k = 0.483425732988414 \quad \text{(defined; exact)}$$

with n the number of distinguishable components required to specify a state. **This paper neither redefines a state, renames one, nor adds one**; the ladder is fixed elsewhere in the series [Curzi, 2026] and is computed on here. It carries forward the honest limit that attaches to every use of the index:

> **The frequency inherits exactly the defensibility of the count. A contested decomposition yields a contested frequency.**

That sentence appears here because this section publishes a table of frequencies next to a table of psychological descriptions, **and that adjacency is exactly where drift happens.**

4.2 Octave folding

The raw essence frequencies of a twelve-state ladder run from 0.4834 Hz to 69.6133 Hz. The house method is octave transposition and only octave transposition, $f_{\text{audible}} = f \cdot 2^j$ with j an integer, into a standing band of 220–440 Hz stated with every value.

Multiplication by 2^j is exact in the sense that matters: it moves a frequency by a whole number of octaves and therefore preserves pitch class exactly. No rounding, no approximation, no tempering. **Folding is a change of register and never a change of relationship.** Two consequences follow immediately and are used throughout:

- **Folding cannot create a collision.** If two states fold to the same value, they were already at the same pitch class.
- **Folding cannot remove one.** If two states share a class, no choice of band separates them.

The second is a defence against a specific abuse — choosing the band after seeing the states, to make two of them collide or to pull them apart. **It does not work, and the fact that it does not work is a strength of the method.**

4.3 The odd-part theorem

THEOREM 4.1 (the odd-part class rule). Write $n = 2^a m$ with $a \geq 0$ and m odd. Then the pitch class of $f(n) = kn^2$ is determined entirely by m, and two indices share a pitch class if and only if they share an odd part.

Proof. $f(n) = k(2^a m)^2 = 2^{2a}\, k\, m^2 = 2^{2a} f(m)$, and 2^{2a} is a power of two, so $f(n)$ is exactly $2a$ octaves above $f(m)$; octave displacement is precisely the equivalence that defines pitch class, so $\text{class}(n) = \text{class}(m)$. Conversely, suppose $f(n_1)/f(n_2) = 2^j$, so $(n_1/n_2)^2 = 2^j$. Writing $n_i = 2^{a_i} m_i$ with m_i odd gives $(m_1/m_2)^2 = 2^{j-2a_1+2a_2}$. **The left side is a ratio of odd squares in lowest terms, hence odd over odd; the right side is a power of two; and the only odd power of two is** $2^0 = 1$. Therefore $m_1 = m_2$. **QED.**

COROLLARY 4.2 (no odd octaves). In the exponent equation above $j = 2(a_1 - a_2)$ is **even**. Two states of the same class are separated by an *even* number of octaves — never one, never three, never five. **The smallest class-preserving interval in the system is exactly two octaves, 2400 cents.**

COROLLARY 4.3 (the class-count ceiling). An index set $1 \ldots N$ realises exactly $\lceil N/2 \rceil$ pitch classes.

Proof. By Theorem 4.1 the classes are in bijection with the odd parts realised by $1 \ldots N$. Every odd $m \leq N$ is realised, by $n = m$ itself; no odd $m > N$ is realised, since $n = 2^a m \geq m > N$. The class set is therefore the odd numbers in $1 \ldots N$, of which there are $\lceil N/2 \rceil$. **QED.**

Verified computationally for every N from 1 to 25, the range a single integer transposition admits inside the audible band:

```
   N :  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25
 cls :  1  1  2  2  3  3  4  4  5  5  6  6  7  7  8  8  9  9 10 10 11 11 12 12 13
ceil :  1  1  2  2  3  3  4  4  5  5  6  6  7  7  8  8  9  9 10 10 11 11 12 12 13
   all match: True
```

Twelve states yield six pitch classes. That is a hard structural limit and not a stylistic choice. It is **not** a claim that there are six feelings, six kinds of feeling, or six anything psychological; it is a claim about how many distinct pitch classes an index of twelve can generate under a square law. The same folding result is owned in Series B [Curzi, 2026k], whose Result Five states it for a set of twenty indices, and the entropy cost of the fold is computed in the closing synthesis [Curzi, 2026], whose Result Two prices pitch-class folding at exactly $\ln 2 = 0.6931471806$ nats — one bit, two octaves — for every even N.

4.4 The twelve-state spectrum, computed

n	state class	$f = kn^2$ (Hz)	j	folded (220–440)	odd part	pitch class	cents above 220
1	undifferentiated activation	0.4834	9	247.514	1	1	204.008
2	valenced activation	1.9337	7	247.514	1	1	204.008
3	directed feeling	4.3508	6	278.453	3	3	407.918
4	situated feeling	7.7348	5	247.514	1	1	204.008
5	attributed feeling	12.0856	5	386.741	5	5	976.635
6	temporal feeling	17.4033	4	278.453	3	3	407.918
7	reflexive feeling	23.6879	4	379.006	7	7	941.660
8	social feeling	30.9392	3	247.514	1	1	204.008
9	normative feeling	39.1575	3	313.260	9	9	611.828
10	narrative feeling	48.3426	3	386.741	5	5	976.635
11	reciprocal feeling	58.4945	2	233.978	11	11	106.644
12	committed feeling	69.6133	2	278.453	3	3	407.918

```
distinct classes = {1, 3, 5, 7, 9, 11}   count 6 = ceil(12/2)      [verified]
max f = k * 144  = 69.61330555033166 Hz
ln N  = ln 12    = 2.4849066497880004
2/ln 2           = 2.8853900817779268
```

The maximum does double duty. $kN^2 = 69.6133$ **Hz is simultaneously the pitch of the most differentiated state on the ladder and the ceiling on the sustaining frequency** — the highest minimum a twelve-category system can ever meet [Curzi, 2026], whose Result Three records that the category count enters the threshold only through $c/\eta \leq \ln N$.

The affect literature does not underwrite the rungs of this ladder, and the objection belongs in the body. The named precedents establish that each component is a real, independently described feature of emotional states; they do not establish an index, an ordering, or a nesting. Appraisal dimensions **vary simultaneously and independently, and emotions differ by the pattern across them** [Smith and Ellsworth, 1985] — a count discards exactly that pattern, and **cannot separate anger from guilt.** Emotion is described as a **process in time** whose state is the synchronisation of several subsystems [Scherer, 2001], which resists reduction to a scalar of any kind. Individuation by **core relational theme** is qualitative [Lazarus, 1991]: a theory in which two emotions differ by *what* is appraised does not license one in which they differ by *how many* things are appraised. The circumplex is **two continuous bipolar axes** recovered from ratings [Russell, 1980], and its "two" is two dimensions rather than two countable components. **And the top rung is contradicted specifically**: action readiness is held to be constitutive of emotion as such rather than a late addition [Frijda, 1986], so if that is right the component sits at the floor and the ladder's ordering fails at its summit. Two rungs survive in the form the framework needs — a claim that a *class* of states requires a component others lack [Tracy and Robins, 2004], and the modelled-other construct [Premack and Woodruff, 1978]. **This paper therefore cites the affect literature for the *existence* of a component and never for the index, its ordering, or its nesting**, and narrows the top rung to *bound commitment* rather than *readiness*, which is the only version that literature leaves standing.

That objection and Michell's are the same objection in two domains. One says a psychological attribute is treated as quantitative without a test for quantitative structure; the other says an emotional state is reduced to a scalar when its own literature describes it as a pattern or a process. **Both are instances of one habit — scalar reduction — and this corpus should say so once rather than meet each as a surprise.** Section 12 concedes both.

4.5 The regime: minimum separation at $N = 12$

The folded class lattice for $N = 12$, sorted by folded pitch, with the interval to the next class above (the last row wraps to the octave):

class	folded (Hz)	to next	cents	Hz gap
11	233.978	$\to$ 1	97.3641	13.5359
1	247.514	$\to$ 3	203.9100	30.9392
3	278.453	$\to$ 9	203.9100	34.8067
9	313.260	$\to$ 7	329.8318	65.7459
7	379.006	$\to$ 5	**34.9756**	**7.7348**
5	386.741	$\to$ 11 (octave)	330.0085	81.2155

The tightest pair is class 7 against class 5, and their ratio is exact: $f(5)/f(7)$ folded $= (2 \cdot 5^2)/7^2 = 50/49$, which is $1200 \log_2(50/49) = 34.975614791419815$ cents. Sweeping N and taking the minimum adjacent separation over the whole folded lattice:

N	min separation (cents)	between classes	ratio
3–6	203.9100	1 and 3	9/8
7–16	34.9756	7 and 5	50/49
12	**34.9756**	**7 and 5**	**50/49**
17–40	6.0008	3 and 17	289/288

N	min separation (cents)	between classes	ratio
≥ 41	1.0296	41 and 29	1682/1681

```
1200 * log2(289/288)   = 6.0008172700399145 cents
1200 * log2(1682/1681) = 1.0295772227721520 cents
```

The cliff at 17 is $17^2 = 289$ **colliding with** $2 \cdot 12^2 = 288$, i.e. class 17 arriving a sixth of a comma from class 3. The second at 41 is $41^2 = 1681$ against $2^4 \cdot 105 = 1680$, reached through class 29. **Both are exact rational coincidences, not approximations.**

$N = 12$ sits below the first cliff, so every folded distinction this framework asks an instrument to make is at least 34.9756 cents wide — roughly a third of a semitone.

One rendering consequence, and it is a bound on the signal and not a statement about hearing. The tightest folded gap in absolute frequency is $386.740586 - 379.005775 = 7.734812$ Hz, and the time–frequency bound $\Delta t \, \Delta f \geq 1/(4\pi)$ [Gabor, 1946] therefore gives $\Delta t \geq 1/(4\pi \cdot 7.734812) = 0.0102882$ s $= 10.2882$ ms as the formal lower bound on the duration of a tone carrying that distinction; at the widest gap in the lattice it is 0.9798 ms. That bound is owned by the conjugate paper [Curzi, 2026z], whose Result One proves $\Delta t \, \Delta f \geq 1/(4\pi) = 0.079577471545948$ by the same variational argument as $\sigma_x \sigma_p \geq \hbar/2$ with $p = \hbar\kappa$ deleted. **Human discrimination needs far longer tones than 10 ms and this paper never confuses the two.**

5 Intensification and transformation

5.1 The dichotomy

A transition from state m to state n has interval $1200 \log_2(f(n)/f(m)) = 2400 \log_2(n/m)$ cents. **Every such interval is a ratio of squares and** k **cancels from it entirely** [Curzi, 2026k].

THEOREM 5.1 (the intensification/transformation dichotomy). A transition $m \to n$ falls into exactly one of two classes:

```
(I)  INTENSIFICATION   m and n share an odd part
                       the interval is an exact WHOLE number of octaves
                       and that number is EVEN (Corollary 4.2)
                       the pitch changes; the pitch CLASS does not

(T)  TRANSFORMATION    m and n have different odd parts
                       the interval is NEVER a whole number of octaves
                       the pitch changes AND the pitch class changes
```

Proof. The interval is a whole number of octaves iff $(n/m)^2$ is a power of two, iff by Theorem 4.1 m and n share an odd part; Corollary 4.2 supplies the evenness. **QED.**

Verified exhaustively over all 66 unordered pairs at $N = 12$:

```
whole-octave transitions found:
  (1,2,+2)  (1,4,+4)  (1,8,+6)  (2,4,+2)  (2,8,+4)
  (3,6,+2)  (3,12,+4) (4,8,+2)  (5,10,+2) (6,12,+2)
count: 10       all octave counts even: True
same-odd-part pairs, counted independently: 10       [identical set]
```

The two counts agree because they are the same set. That agreement is the theorem, checked rather than asserted. The distinction is therefore not imposed on the psychology; **it is read off the index before anyone has decided what the states mean.**

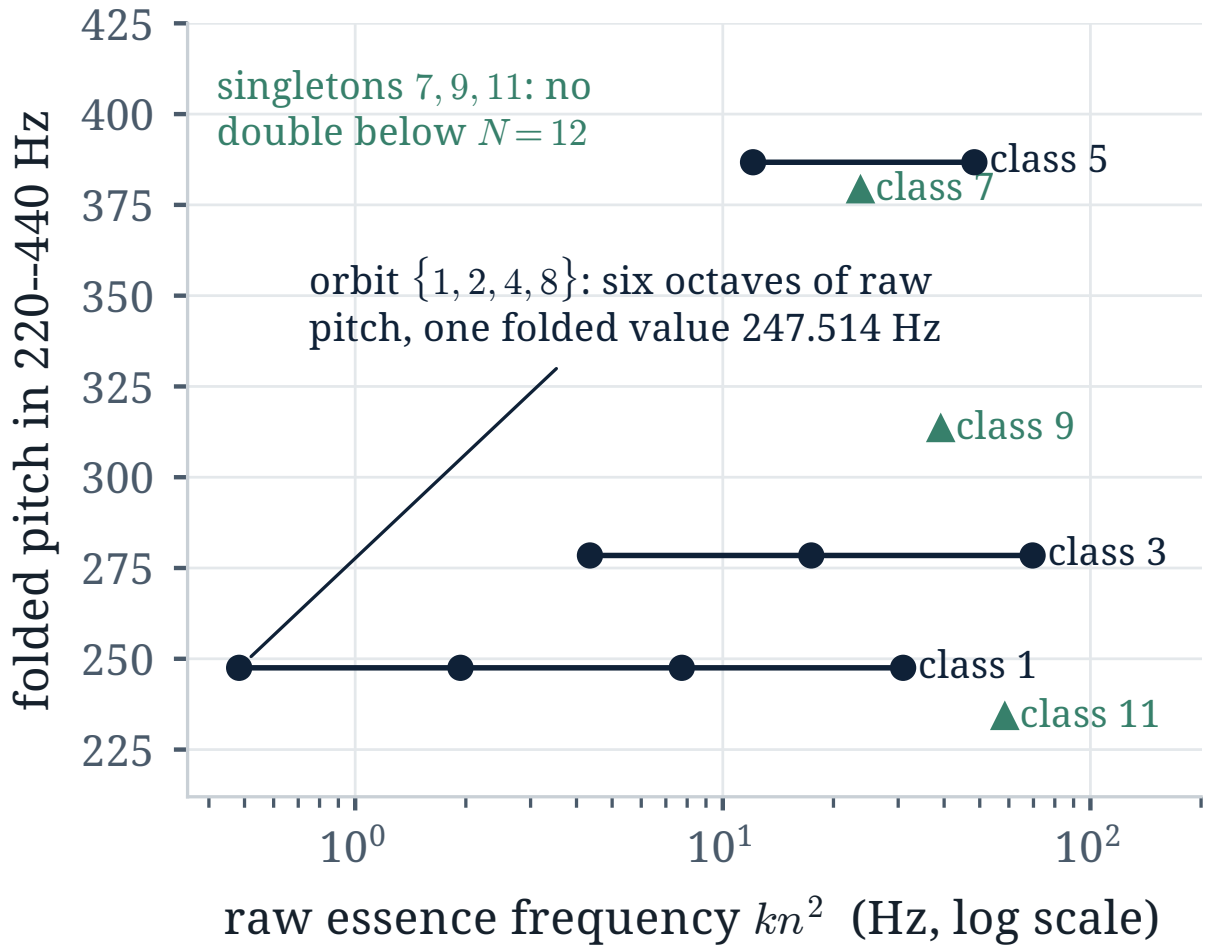


Figure 42.1: **Twelve indices collapse onto six folded values**, so an entire orbit of raw pitch spans six octaves at one folded position: indices $\{1, 2, 4, 8\}$ all fold to 247.514 Hz, $\{3, 6, 12\}$ to 278.453 Hz and $\{5, 10\}$ to 386.741 Hz, while 7, 9 and 11 are singletons because their doubles exceed $N = 12$. Modelled from $f(n) = kn^2$ with the **defined** constant $k = 0.483425732988414$, folded into the declared 220–440 Hz band. Built by `figs/src/CCP-042_fig01.py`.

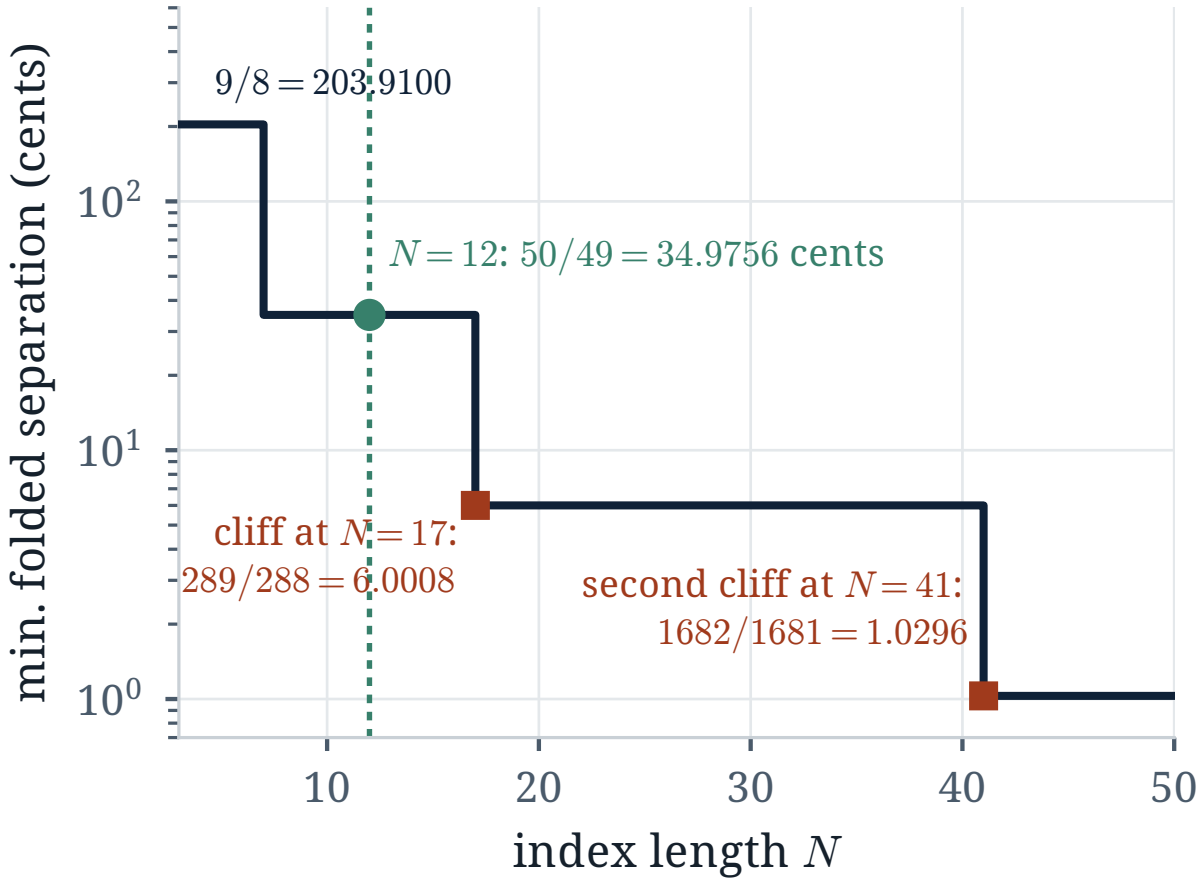


Figure 42.2: $N = 12$ **renders in the** $50/49$ **regime, three regimes above the first collision cliff**: the minimum folded separation is 34.9756 cents for $N = 7 \ldots 16$, collapses to 6.0008 cents the instant index 17 enters the set, and again to 1.0296 cents at $N = 41$. Computed in exact rational arithmetic on the folded class lattice with the **defined** constant $k = 0.483425732988414$; the ratios $9/8$, $50/49$, $289/288$ and $1682/1681$ are exact. Built by `figs/src/CCP-042_fig03.py`.

5.2 The orbit structure

Grouping the twelve indices by odd part gives six orbits under doubling:

```
class  1  ->  states { 1, 2, 4, 8 }      four degrees, spanning 7200 cents
class  3  ->  states { 3, 6, 12 }        three degrees, spanning 4800 cents
class  5  ->  states { 5, 10 }           two degrees, spanning 2400 cents
class  7  ->  states { 7 }               one degree
class  9  ->  states { 9 }               one degree
class 11  ->  states { 11 }              one degree

six orbits, twelve states, mean orbit size 2.0
```

Two structural facts follow directly:

```
states possessing a double INSIDE the ladder (2n <= 12):   1,2,3,4,5,6     (six)
states possessing none      (2n >  12):                    7,8,9,10,11,12  (six)

class-preserving unordered pairs : 10 / 66 = 15.1515 %
class-changing   unordered pairs : 56 / 66 = 84.8485 %
```

Exactly half the ladder can be intensified without leaving it. The upper half cannot — not because the upper states are final, but because $N = 12$ truncates their orbits. **That is an artefact of index length and it is honest to name it as one.**

Transformation is the generic move; intensification is the rare and highly structured one. For an index of length N the number of same-class pairs is $\sum_{m \text{ odd} \le N} \binom{s_m}{2}$ with $s_m = \lfloor \log_2(N/m) \rfloor + 1$; at $N = 12$ the orbit sizes $4, 3, 2, 1, 1, 1$ give $6 + 3 + 1 = 10$. **The orbit sizes fall off like $\log_2(N/m)$, so the supply of intensification moves grows only logarithmically while the total supply of moves grows quadratically.** Lengthen the ladder and transformation dominates ever more completely.

5.3 Constructed trajectories

All cases in this section are CONSTRUCTED and labelled as such. None is an observation of any person, and no clinical, diagnostic or therapeutic claim is made or implied.

Trajectory M1 — the same class at three degrees.

```
state  3  directed feeling      f =  4.3508 Hz  fold j=6  -> 278.453 Hz  class 3
state  6  temporal feeling      f = 17.4033 Hz  fold j=4  -> 278.453 Hz  class 3
state 12  committed feeling     f = 69.6133 Hz  fold j=2  -> 278.453 Hz  class 3

step 3 ->  6 : +2400.0000 cents = +2 octaves exactly
step 6 -> 12 : +2400.0000 cents = +2 octaves exactly
total 3 -> 12: +4800.0000 cents = +4 octaves exactly, ratio 144/9 = 16 = 2^4
```

The pitch has moved four octaves and the class has not moved at all. That is what *intensification* means here, and it is a technical term with an arithmetic definition rather than an intuition. **It does not say the experience is unchanged**; a state with a projected duration is not the same as one without. It says the two sit at the same position in the folded lattice.

Trajectory M2 — a change of class.

```
state 3  directed feeling    f =  4.3508 Hz  -> 278.453 Hz  class 3
state 5  attributed feeling  f = 12.0856 Hz  -> 386.741 Hz  class 5

interval 3 -> 5 : 2400 * log2(5/3) = 1768.7174 cents = 1.473931 octaves
frequency ratio : 25/9  -- not a power of two, and by Theorem 5.1 never could be
```

Trajectory M3 — the longest single-class walk $N = 12$ permits.

```
1 -> 2 -> 4 -> 8    every step +2400.0000 cents; total +7200.0000 = 6 octaves exactly
folded value 247.514 Hz at every rung, without exception
```

Placing a selection side by side:

move	Δn	cents	octaves	class	kind
$3 \to 6$	+3	2400.0000	+2 exactly	$3 \to 3$	intensification
$3 \to 12$	+9	4800.0000	+4 exactly	$3 \to 3$	intensification
$1 \to 8$	+7	7200.0000	+6 exactly	$1 \to 1$	intensification
$3 \to 5$	+2	1768.7174	1.473931	$3 \to 5$	transformation
$5 \to 7$	+2	1165.0244	0.970854	$5 \to 7$	transformation
$4 \to 5$	+1	772.6274	0.643856	$1 \to 5$	transformation
$11 \to 12$	+1	301.2741	0.251062	$11 \to 3$	transformation

The $+9$ move from 3 to 12 is an intensification. The $+1$ move from 11 to 12 is a transformation. Distance on the ladder is not what separates the two kinds; the odd part is. **A reader who arrived expecting "bigger jump equals more different" has to give that up, and giving it up is the point.**

5.4 What the dichotomy does and does not claim

It claims that transitions partition exactly into whole-even-octave moves and everything else; that the partition is computable from the indices alone; that the first kind preserves pitch class and the second does not; and that consequently the lattice has a formal representation of *more of the same thing* distinct from its representation of *a different thing*.

It does not claim that an intensification is smaller, easier, less consequential or less distressing than a transformation. The class-1 orbit spans six octaves, and **the formalism says nothing about how it feels to cross it.**

It does not claim that the odd part means anything on its own. Class 7 is not "the reflexive class" in any sense beyond the fact that 7 is the only index in $1 \dots 12$ with odd part 7. Change N and the singleton acquires company at index 14.

It does not claim to be a taxonomy of emotion, a stage model, or a diagnostic scheme.

And it does not survive a bad index. The dichotomy is a theorem about integers. It becomes a statement about a person only through the count, and it inherits every weakness of that count.

5.5 The scarcity of intensification, and what it costs

The $10/66$ figure constrains what the rest of the corpus can promise. **Intensification is not merely rarer; it is structurally scarce and becomes scarcer as the ladder lengthens.** The consonance paper draws the matching conclusion from the other side [Curzi, 2026], whose Result Five records the extremal four-state chords at $N = 12$: $\{1, 2, 4, 8\}$ has mean $\mathrm{HD}_{\mathrm{ess}} = 3.3333$ with six collisions and one pitch class, while 159 of the 495 four-state chords (32.1 per cent) have none. **The most consonant chords the lattice permits are precisely the most redundant ones**, and a self composed entirely of one orbit would be a single note played four times.

5.6 Procedure and reporting

```
MOOD-1 -- reading a single state
```

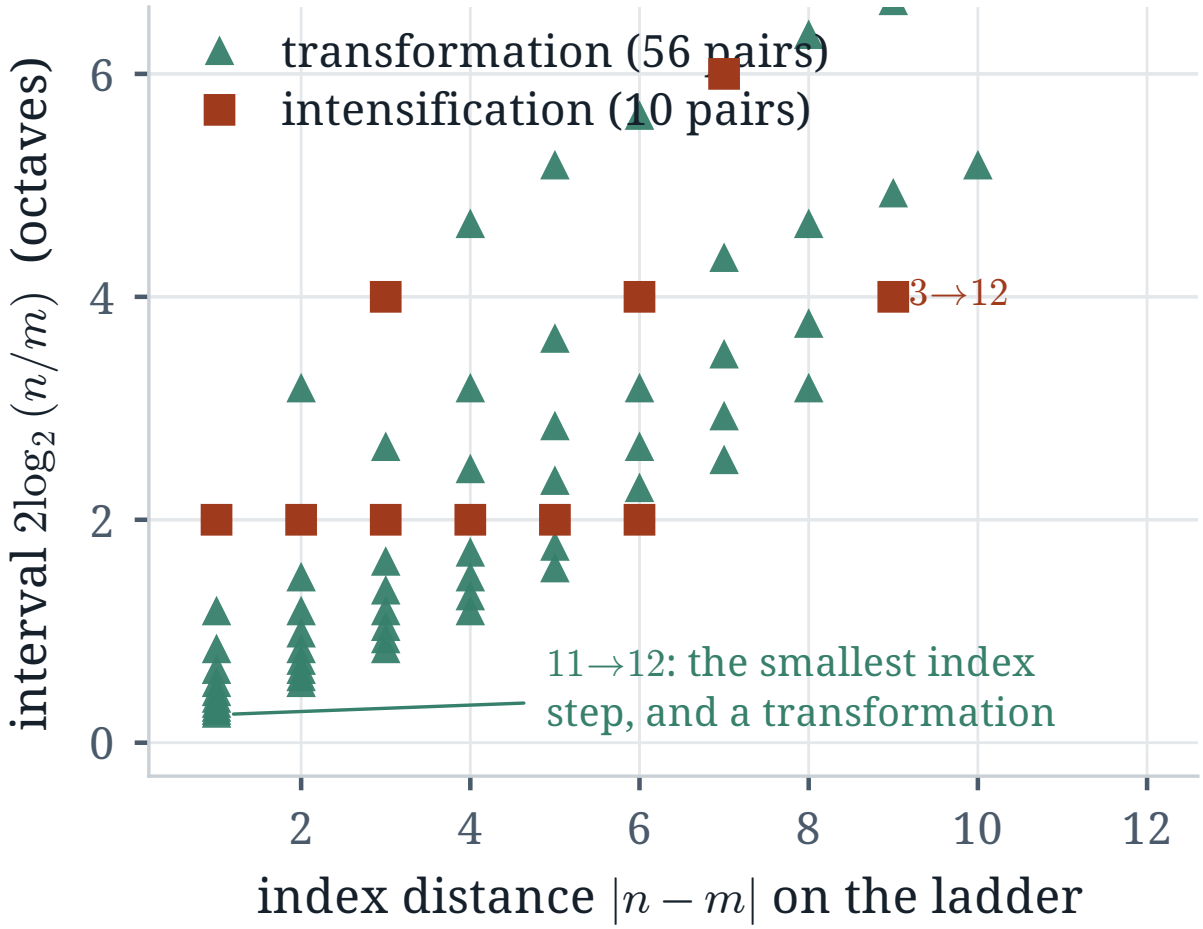


Figure 42.3: **Index distance does not determine the kind of move**: the ten class-preserving pairs land on whole even octaves regardless of how far apart their indices are, so the $+9$ step $3 \to 12$ is an intensification while the $+1$ step $11 \to 12$ is a transformation. All 66 unordered pairs of the twelve-state ladder, computed exactly from the indices; the constant k cancels from every interval and does not enter. Built by `figs/src/CCP-042_fig02.py`.

```
step 1  DECOMPOSE.  List the distinguishable components. Record the LIST, not
        just its length. The list IS the warrant.
step 2  COUNT.  n := the number of components. Check 1 <= n <= 12. If the
        decomposition yields n > 12, say so and stop. DO NOT COMPRESS TO FIT.
step 3  APPLY.  f_raw = k n^2,  k = 0.483425732988414 Hz (defined; exact).
step 4  FOLD.  Find the unique j with 220 <= f_raw * 2^j < 440. Publish BOTH
        the band and j.
step 5  CLASS.  Strip factors of two from n until odd; that is the pitch class.
step 6  BADGE.  essence frequency, not physical resonance; not a diagnosis;
        frequency inherits the defensibility of the count.

MOOD-2 -- reading a CHANGE
step 1  INTERVAL.  cents = 2400 log2(n/m);  octaves = 2 log2(n/m).
step 2  CLASSIFY.  odd_part(m) == odd_part(n) ?
            yes -> INTENSIFICATION; octaves MUST be an even integer. Verify it.
            no  -> TRANSFORMATION; octaves is never an integer (Theorem 5.1).
step 3  REPORT.  index change, interval, kind, both classes, and BOTH component
        lists. Without the component lists the reading is unauditable.
```

Step 2 of MOOD-2 is a genuine error check and costs nothing. If the odd parts match, the octave count must come out an even integer; if it does not, the arithmetic is wrong or an index was mistyped. **This is one of the few places in the framework where the formalism catches its own mistakes.**

5.7 Failure modes

F1 — editing the ladder to fit the reading. Renaming, splitting, merging, adding a thirteenth rung or reordering changes every frequency, every class, every orbit and the entire dichotomy. *Remedy:* if a case genuinely needs more than twelve rungs, **that is a finding about the case and is reported as one — never fixed by compression.**

F2 — reading class identity as sameness of feeling. States 1 and 8 share class 1 and fold to the same 247.514 Hz. **They are not the same feeling and this paper never says they are.** *Detection:* the words "is really just" near a class collision. *Remedy:* quote the raw frequencies and the octave count alongside the class. **Six octaves is not nothing.**

F3 — treating the folded pitch as a magnitude. 386.741 Hz is not "more" than 247.514 Hz in any sense the formalism supports; **folding destroys magnitude by design.** *Remedy:* rank by n or f_{raw}, never by the folded value.

F4 — importing certainty the count does not have. $f = 39.1575$ Hz looks exact to six figures. **It is exactly as exact as the claim that the state required nine components and not eight or ten.** *Remedy:* the component list, mandatory.

F5 — mistaking the class ceiling for a claim about feelings. $\lceil N/2 \rceil = 6$ is a fact about odd numbers below 12. *Detection:* "the six fundamental" appearing in any sentence. *Remedy:* say "six pitch classes" and never drop the noun.

F6 — physical-resonance drift. Nothing resonates. Not a neuron, not a tissue, not a room, not a person. *Detection:* any sentence combining a frequency from this paper with a biological noun and a verb of oscillation.

F7 — diagnosis drift. The twelve states are not diagnoses. A pitch class is not a type of person, an index is not a severity, a trajectory is not a prognosis. **The prohibition is standing and admits no exception.**

F8 — band-shopping. Choosing the band to make two states collide or separate. **It does not work** (Theorem 4.1), and that it does not work is a strength of the method.

F9 — cents as effect size. 1768.7174 cents is an interval on a lattice. It is not a measure of how much a person changed, how hard the change was, how long it took, or whether it was good.

F10 — quietly changing N. The $50/49$ result belongs to $N = 12$; past $N = 16$ the minimum separation collapses by a factor of 5.83 in cents and everything about rendering must be recomputed.

6 Persistence: what the frequency law cannot see

6.1 The law is unchanged, and that is the result

Nothing about the essence law changes when the state read is a standing disposition rather than a passing one, because **there is no mechanism available by which it could change.**

THEOREM 6.1 (frequency does not distinguish disposition from episode). The essence law is a function of n alone. Persistence is not an argument of it. Therefore two states of equal index have equal essence frequency, equal folded pitch and equal pitch class, irrespective of how long either lasts. **QED.**

The proof is one line because the content is one line: **persistence cannot be read off a pitch, because the pitch was never told about it.** Recomputing the twelve-state spectrum for standing dispositions reproduces the table of Section 4.4 column for column — the same raw frequencies, the same transposition exponents, the same folded values, the same six classes. **Nothing was changed to make it agree; there was nothing available to change.**

This is the usual point at which a formalism quietly acquires a second parameter — a modifier, a coefficient of stability folded into n, something. **The correct response is not to patch the frequency law.** The parameter that does the job already exists in the corpus's storage law, and it belongs to a different equation entirely.

6.2 Where the decay lives

$$Q_{t+1} = (1 - \delta)Q_t + (\eta F(u^*) - c^*), \qquad Q_\infty = \frac{\eta F(u^*) - c^*}{\delta}$$

δ is the retention loss: a per-cycle rate in $(0, 1]$. **It is the only parameter in the system that speaks about time, and it is the only one absent from the essence law.**

DEFINITION 6.2 (a trait). A **trait** is a state held at low δ; a **mood** is the same state held at high δ. The index — and therefore the frequency, the folded pitch and the pitch class — is common to both. **The distinction is entirely and exclusively a claim about retention loss.**

6.3 The storage half-life

Setting the input to zero gives $Q_t = (1 - \delta)^t Q_0$, so

$$t_{1/2} = \frac{\ln 2}{-\ln(1 - \delta)}$$

Computed across the range this corpus uses, with the time to lose ninety per cent, the time to reach ninety-five per cent of the fixed point, and the naive linear estimate:

δ	$-\ln(1-\delta)$	$t_{1/2}$ (cycles)	t to 90% lost	t to 95% of Q_∞	$1/\delta$
0.005	0.005013	138.2826	459.3648	597.6473	200.0000
0.010	0.010050	68.9676	229.1053	298.0729	100.0000
0.020	0.020203	34.3096	113.9741	148.2837	50.0000
0.030	0.030459	22.7566	75.5957	98.3523	33.3333
0.050	0.051293	13.5134	44.8906	58.4040	20.0000
0.070	0.072571	9.5513	31.7289	41.2802	14.2857

δ	$-\ln(1-\delta)$	$t_{1/2}$ (cycles)	t to 90% lost	t to 95% of Q_∞	$1/\delta$
0.090	0.094311	**7.3496**	24.4149	31.7645	11.1111
0.100	0.105361	6.5788	21.8543	28.4332	10.0000
0.150	0.162519	4.2650	14.1681	18.4331	6.6667
0.200	0.223144	3.1063	10.3189	13.4251	5.0000
0.250	0.287682	2.4094	8.0039	10.4133	4.0000
0.300	0.356675	1.9434	6.4557	8.3991	3.3333
0.400	0.510826	1.3569	4.5076	5.8645	2.5000
0.500	0.693147	1.0000	3.3219	4.3219	2.0000
0.700	1.203973	0.5757	1.9125	2.4882	1.4286
0.900	2.302585	0.3010	1.0000	1.3010	1.1111

The $1/\delta$ column is not the half-life and the gap is large. At $\delta = 0.30$ the naive estimate is 3.3333 against the exact 1.9434: **an overstatement of 71.52 per cent.**

> **A rounding in the source, printed rather than silently applied.** The source reports this error as "seventy-two per cent". The exact figure is $(3.3333 - 1.9434)/1.9434 = 71.52$ per cent. The difference is a rounding, not a disagreement, and this paper quotes 71.52 while recording that 72 is what the source prints, so that a reader recomputing finds no surprise.

Only as δ falls do the two converge on the ratio $\ln 2 = 0.693147$: at $\delta = 0.001$, $\ln 2/\delta = 693.1472$ against the exact 692.8005, agreement to four figures. **The linear intuition is safe exactly where traits are and unsafe exactly where moods are.** Halving δ does not exactly double the half-life but nearly does, and does so from above: $t_{1/2}(0.150)/t_{1/2}(0.300) = 2.1946$, then 2.0846, then 2.0397, converging on 2.

6.4 Retention, tabulated

$(1-\delta)^t$, the fraction of stored order surviving t cycles with no input:

δ	$t=1$	$t=2$	$t=3$	$t=5$	$t=10$	$t=20$	$t=50$	$t=100$
0.01	0.990000	0.980100	0.970299	0.950990	0.904382	0.817907	0.605006	0.366032
0.03	0.970000	0.940900	0.912673	0.858734	0.737424	0.543794	0.218065	0.047553
0.05	0.950000	0.902500	0.857375	0.773781	0.598737	0.358486	0.076945	0.005921
0.09	0.910000	0.828100	0.753571	0.624032	0.389416	0.151645	0.008955	0.000080
0.15	0.850000	0.722500	0.614125	0.443705	0.196874	0.038760	0.000296	0.000000
0.30	0.700000	0.490000	0.343000	0.168070	0.028248	0.000798	0.000000	0.000000

At $\delta = 0.30$ nothing survives twenty cycles; at $\delta = 0.01$, thirty-seven per cent survives a hundred. **Between those two rows lies the entire difference between a mood and a trait, and there is not a single frequency in the table.** The same shape governs long-horizon erosion elsewhere in the corpus [Curzi, 2026], whose Result Five decomposes an erosion rate of 2.0 per cent per year into four components and reports 73.857 per cent retained at fifteen years and 54.548 per cent at thirty, half-life 34.310 years.

6.5 The cycle must be declared

δ has no meaning without a stated cycle, and the rescaling is exact: $\delta_r = 1 - (1-\delta_1)^r$.

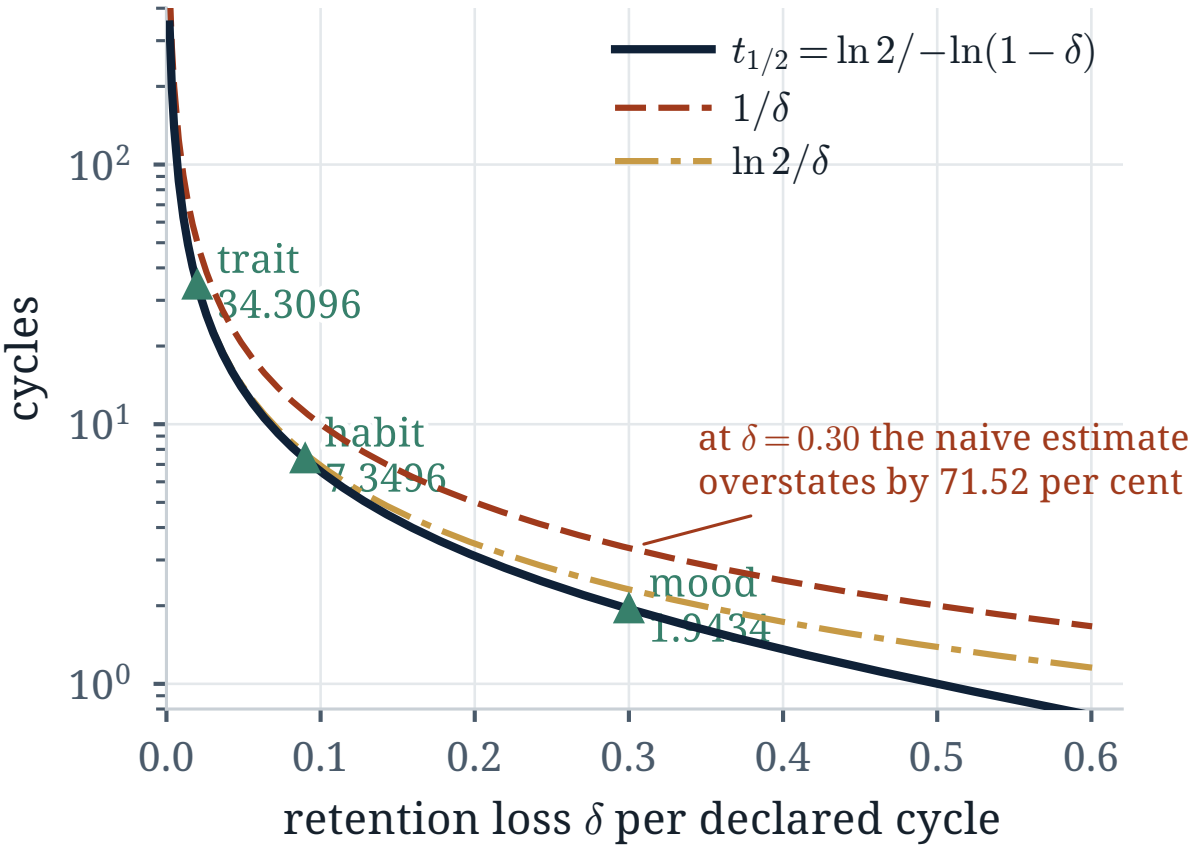


Figure 42.4: **The naive linear estimate is wrong exactly where moods live**: at $\delta = 0.30$ the true half-life is 1.9434 cycles against $1/\delta = 3.3333$, an overstatement of 71.52 per cent, while at trait-grade $\delta = 0.02$ the exact 34.3096 and $\ln 2/\delta = 34.6574$ agree to three figures. Computed from $t_{1/2} = \ln 2/(-\ln(1-\delta))$; the three marked points are the framework's declared mood, habit and trait grades. Built by `figs/src/CCP-042_fig04.py`.

δ per day	per week ($r = 7$)	per month ($r = 30$)	per year ($r = 365$)
0.02	0.131874	0.454516	0.999373
0.05	0.301663	0.785361	1.000000
0.09	0.483239	0.940947	1.000000
0.30	0.917646	0.999977	1.000000

The half-life is invariant in *time*; only its units change. Verified: $\delta/\text{day} = 0.02$ gives 34.3096 days $= 4.9014$ weeks, and $\delta/\text{week} = 0.131874$ gives 4.9014 weeks — identical. **A δ quoted without its cycle is not a number.** "$\delta = 0.09$" describes a disposition with a seven-cycle half-life, and whether that is a week or a decade depends on a unit nobody stated.

6.6 The trait cut is a declared convention

DEFINITION 6.3 (the trait window). Fix an observation window W in cycles. A state qualifies as a **trait relative to** W when $t_{1/2} \geq W$, equivalently $\delta \leq 1 - 2^{-1/W}$.

W (cycles)	δ must be $\leq$
5	0.129449
10	0.066967
20	0.034064
30	0.022840
50	0.013767
100	0.006908
365	0.001897

This is a convention, declared, and it is not a discovery about persons. There is no δ at which something switches from being an episode to being a disposition; there is a continuous rate and a cut somebody chose. **What the definition buys is not truth but auditability**: two readings that declare W can be compared and two that do not, cannot. The working labels, used only for legibility and never as categories of person, are $\delta \approx 0.30$ (mood, $t_{1/2} = 1.9434$), 0.09 (habit, 7.3496), 0.02 (trait, 34.3096) and 0.005 (core, 138.2826).

6.7 The division of labour

THEOREM 6.4 (sign independence). For any $\delta \in (0, 1]$, $\operatorname{sign}(Q_\infty) = \operatorname{sign}(\eta F(u^*) - c^*)$, and δ never changes it. Moreover Q_∞ is exactly inversely proportional to δ: $Q_\infty(\delta_2)/Q_\infty(\delta_1) = \delta_1/\delta_2$ exactly, for any numerator.

Proof. $Q_\infty = (\eta F - c)/\delta$ with $\delta > 0$; division by a positive number preserves sign, and the ratio of two such expressions with a common numerator is δ_1/δ_2. **QED.**

Its companions contain no δ at all:

```
f_min          = k e^( 2c/eta )                  CONTAINS NO delta
octave deficit = 2.885390 (c/eta) - log2(f/k)    CONTAINS NO delta
```

Verified on three constructed parameter sets: $\eta = 0.30$, $c = 0.62$ gives $f_{\min} = 30.1588$ Hz; $\eta = 0.44$, $c = 0.40$ gives 2.9782 Hz; $\eta = 0.34$, $c = 0.825$ gives 61.9349 Hz. The ceiling $c/\eta \leq \ln 12 = 2.4849066$ gives $f_{\min} \leq k \cdot 144 = 69.6133$ Hz. **So the two halves of the instrument divide cleanly:**

```
THE FREQUENCY READING DECIDES THE SIGN.
    Does the arrangement accumulate or decay?  f vs f_min.  delta is irrelevant.
```

```
delta DECIDES THE MAGNITUDE AND THE CLOCK.
    How far does it go, and how fast does it get there?  delta and nothing else.
```

A disposition's frequency tells you which way it is going; its δ tells you how far and how fast. Neither substitutes for the other, and **a reading that reports one without the other is half a reading.**

6.8 Durability and revisability are the same number

The convergence identity $|Q_t - Q_\infty| = (1-\delta)^t |Q_0 - Q_\infty|$ reads once for decay and once for approach. **It is the same rate: there is one time constant and it governs both directions.**

δ	$t_{1/2}$ of decay	$t_{1/2}$ of approach	identical?
0.02	34.3096	34.3096	yes
0.05	13.5134	13.5134	yes
0.09	7.3496	7.3496	yes
0.30	1.9434	1.9434	yes

COROLLARY 6.5 (no free durability). A state cannot be made durable without being made slow to revise, and cannot be made quick to revise without being made quick to lose. **The two properties are the same parameter read in two directions.**

This dissolves a wish a great deal of practical work is built around — a disposition that persists when wanted and yields when wanted. **The storage law has no term for "on demand". It has δ, and δ is symmetric.** Tabulated as the arrival problem from $Q_0 = 0$ with a fixed positive numerator $a = 0.5476136$:

δ	Q_∞	$t_{1/2}$	t to 95%	Q at $t=1$	$t=5$	$t=10$	$t=25$	$t=50$	$t=100$
0.30	1.8254	1.9434	8.3991	0.5476	1.5186	1.7738	1.8251	1.8254	1.8254
0.09	6.0846	7.3496	31.7645	0.5476	2.2876	3.7152	5.5088	6.0301	6.0841
0.02	27.3807	34.3096	148.2837	0.5476	2.6307	5.0087	10.8574	17.4095	23.7495

Every row starts at the same place after one cycle, because one cycle of input is one cycle of input regardless of the leak. By $t = 100$ the low-δ row holds thirteen times what the high-δ row holds — while still not having arrived. **A trait is, by construction, the thing that takes longest to finish arriving.**

6.9 The price of holding a level

Holding stored order at Q^* requires per-cycle input $a = \delta Q^*$:

δ	$Q^*=1$	$Q^*=5$	$Q^*=10$	$Q^*=27.38$
0.005	0.0050	0.0250	0.0500	0.1369
0.020	0.0200	0.1000	0.2000	0.5476
0.050	0.0500	0.2500	0.5000	1.3690
0.090	0.0900	0.4500	0.9000	2.4643
0.300	0.3000	1.5000	3.0000	8.2142

Maintenance cost is linear in δ and linear in the level held. Holding $Q^* = 10$ at trait loss costs 0.2000 per cycle; at mood loss it costs 3.0000 — **fifteen times as much, every cycle, forever.** Put beside Corollary 6.5, the trade is fully priced:

Low δ buys cheap maintenance and expensive revision. High δ buys cheap revision and expensive maintenance. There is no setting that is cheap at both, because it is one setting.

6.10 The durable liability

Theorem 6.4 says δ cannot flip a sign; it scales. **So when the numerator is negative, lowering δ scales a negative number.** Taking the corpus's constructed dyad as found — $u = 0.068735$, $F = 0.563085$ nats, $\eta = 0.30$, $c = 0.62$, numerator -0.4510739, $f = 1.4908$ Hz against $f_{\min} = 30.1588$ Hz, octave deficit $+4.3384$ — and sweeping δ:

δ	Q_∞	$t_{1/2}$	reading
0.30	-1.5036	1.9434	decays, shallow, fast
0.09	**-5.0119**	7.3496	decays — the as-found reading
0.05	**-9.0215**	13.5134	decays deeper, slower
0.02	-22.5537	34.3096	decays far deeper, far slower
0.01	-45.1074	68.9676	decays deepest, slowest

The ratio between the two bolded values is exactly what Theorem 6.4 requires: $-9.0215/-5.0119 = 1.800000\ldots = 0.09/0.05$. **No approximation: δ is a pure scale.**

> **A trait is not a good thing. A trait is a persistent thing.** Low δ deepens whichever fixed point is already there, and if the numerator is negative, low δ is the mechanism by which a problem becomes a permanent feature rather than an episode.

The intuitive repair for a struggling arrangement is to add trust, and trust in this parameterisation acts on δ. **The intervention everyone reaches for first makes the reading eighty per cent worse**, and it does so for the entirely mechanical reason that it multiplied a negative number by 1.8. **Interpretive care is required and this paper states it:** that trust reduces the loss rate rather than adding to the inputs is an assumption of the parameterisation, not a finding, and the trust literature is about reciprocity, monitoring and institutional form rather than about decay rates [Gambetta, 1988, Ostrom, 1990]. **If trust instead raises η, the arithmetic above does not apply.**

6.11 A positive numerator

The mirror case, so the section is not read as an argument against durability. The same dyad after one intervention: $u = 0.692400$, $\eta = 0.44$, $c = 0.40$, $F = 2.153667$ nats, numerator $+0.5476136$, $f = 35.8907$ Hz against $f_{\min} = 2.9782$ Hz, deficit -3.5911.

	numerator	$\delta = 0.09$	$\delta = 0.05$	effect of lowering δ
as found	-0.4510739	-5.0119	-9.0215	**worse by 4.0096**
after	$+0.5476136$	$+6.0846$	$+10.9523$	**better by 4.8677**

The identical intervention, on the identical parameter, in the identical direction, helps in one case and hurts in the other — and the frequency reading is what tells you which case you are in before you act.

> **A discrepancy in the source, printed rather than silently corrected.** The source prints $f = 35.8882$ Hz and $Q_\infty = +6.0844$ at $\delta = 0.09$. Recomputing from the printed

$u = 0.692400$ — itself a rounding of 0.69237163 — gives 35.8907 and $+6.0846$. **The two disagree in the fourth significant figure because the printed u is rounded, not because either is wrong.** The result this section depends on is the *ratio*, and the ratio is exactly 1.8 at either precision.

6.12 The truncation, revisited for dispositions

An intensification requires the double to exist on the ladder. **So for a standing disposition read at index 7, 9, 11 or above, there is no more-differentiated version of the same pitch class available within $N = 12$**, and any move from there is a transformation. **This is an artefact of index length and is reported as one.** It is not a claim that reflexive, normative or reciprocal dispositions cannot deepen; it is a claim that *this ladder, at this length, has no rung to represent the deepening as an octave.* Lengthening the index would supply one — and would also, past $N = 16$, drop the whole framework out of the $50/49$ regime.

6.13 What a trait reading contains, and what it costs to obtain

```
n          the index          -- SOFT: inherits the defensibility of the count
f          the frequency      -- HARD given n: exact arithmetic, k n^2
class      the pitch class    -- HARD given n: the odd part (Theorem 4.1)
delta      the retention loss -- SOFT: a rate, requiring repeated observation
                                 and a DECLARED cycle
t_half     the half-life      -- HARD given delta: ln2 / -ln(1-delta)
```

The two soft numbers carry all the psychological content and the two hard ones are pure arithmetic on top of them. That is the honest shape of the instrument.

This framework does not supply a validated instrument for δ and says so plainly. The shape of an honest estimate is: declare the cycle; establish a no-input interval, because an interval during which the arrangement is being actively maintained measures (δ, η, c) jointly and δ not at all; **observe over at least three expected half-lives, because below that the decay curve is indistinguishable from a line and δ is unidentified**; fit $\widehat{\delta} = 1 - (Q_t/Q_0)^{1/t}$; and report the estimate with the number of cycles it was estimated over. **At $\delta = 0.02$ the third step means roughly 103 cycles of observation. That is the binding constraint on the whole method, it is expensive, and a report must say so.**

Traits are, by construction, the states whose defining parameter is the most expensive to observe. That asymmetry is a real property of the method rather than a gap in the exposition.

7 The chord: simultaneity at internal scope

7.1 What a chord is

Nobody is ever in one state. A person entering a difficult conversation holds, at the same instant, a reading of the other party's intent, a story about how the situation arose, and a running commentary on their own performance inside it. **These are not three states in sequence; they are three *framings* held together**, each assigning intensity to several of the twelve categories at once.

A **chord** is the set of indices carrying nonzero intensity across the framings held at one moment. It is a set, not a single index, and it inherits everything the square law says about relations between indices.

7.2 The intervals are ratios of squares

$$\frac{f(m)}{f(n)} = \frac{km^2}{kn^2} = \frac{m^2}{n^2}$$

k **cancels.** Whatever one thinks of the constant — and it is a postulate, badged as one — it sets the absolute reference pitch and nothing else. **A reader who rejects k must still accept every interval in this section, because no interval contains it.** Consequences used throughout:

```
index ratio    essence interval    what it is
    2:1             4/1            two octaves exactly
    3:2             9/4            an octave plus a whole tone
    4:3            16/9            a minor seventh
    5:4            25/16           a near minor sixth
    3:1             9/1            three octaves plus a whole tone
```

7.3 The width of the chord is the internal surplus

$$F_{\text{int}} = \ln(1 + (N-1)u_{\text{int}}), \qquad f_{\text{width}} = k\, e^{2F_{\text{int}}} = k\,(1 + (N-1)u_{\text{int}})^2$$

where $u = 1 - (x \cdot y)^2$ on unit framing vectors and u_{int} is the mean pairwise non-parallelism across one observer's framings.

This is a substitution of domain, not an analogy, and the claim is checkable. The surplus functional's four axioms quantify over states in an inner-product space; **not one of them contains a term referring to persons, agents, groups, or the ownership of states**, so the theorem is silent on who holds the states and applying it inside one observer changes the encoder and nothing else [Curzi, 2026]. **Two postulates travel with the substitution and are restated rather than assumed** — the framing encoding (P-d) and the internal roster reduction (P-e).

The chord's width expressed as a pitch is the internal surplus expressed as a pitch: one number in two units. A person whose framings are parallel has $u_{\text{int}} = 0$, $N_{\text{eff}} = 1$, $F_{\text{int}} = 0$ and a width of exactly k — a single note. A person whose framings are orthogonal has $u_{\text{int}} = 1$, $N_{\text{eff}} = 12$, $F_{\text{int}} = \ln 12 = 2.484907$ and a width of $k \cdot 144 = 69.6133$ Hz, 7.169925 octaves above the ground. **One nat of internal surplus is exactly 2.885390 octaves.**

A precision note. $2/\ln 2 = 2.885390081777927$ exactly. Where the source writes 2.885 it is a shorthand costing about a thousandth of an octave at the ceiling — $2.885 \times \ln 12 = 7.1690$ against the exact 7.1699. **This paper uses the exact constant in every computed figure and says so once here.**

7.4 The one-hot degeneracy

Everything above is correct and everything above can be made to produce a meaningless number by a single encoding decision that looks, at first sight, like the cleanest possible one.

> **THE ORTHOGONALITY DEGENERACY [THEOREM].** Let x and y be unit vectors with **disjoint support** — no index carrying nonzero intensity in both. Then $x \cdot y = 0$, so $u = 1$ and $F = \ln N$, the maximum. **Distinct one-hot vectors always have disjoint support. Therefore any set of distinct one-hot framings scores the maximum internal surplus, for every subject, regardless of which states they name.**

Proof. $x \cdot y = \sum_i x_i y_i$, and disjoint support means every term has a zero factor. Substituting into $u = 1 - (x \cdot y)^2$ gives $u = 1$ and into $F = \ln(1 + (N-1)u)$ gives $F = \ln N$. **Nothing in either step consulted the identity of the states. QED.**

Computed on four constructed subjects, so a reader sees it rather than takes it:

```
ONE-HOT FRAMINGS -- FOUR DIFFERENT SUBJECTS, IDENTICAL READING

framings at states {3, 5, 7}    u = 1.000000 on all pairs
                                F = 2.484907  N_eff = 12.0000  f = 69.6133 Hz
framings at states {1, 2, 3}    u = 1.000000  F = 2.484907  N_eff = 12.0000  f = 69.6133 Hz
framings at states {4, 9, 12}   u = 1.000000  F = 2.484907  N_eff = 12.0000  f = 69.6133 Hz
framings at states {1, 11, 12}  u = 1.000000  F = 2.484907  N_eff = 12.0000  f = 69.6133 Hz
```

Four subjects with nothing in common produce the same reading, and it is the maximum reading the system can produce. A subject whose three framings are a nested sequence — each containing the last — scores exactly what a subject spanning states 4, 9 and 12 scores. **The encoding has erased the content and reported its own shape.**

Worse: the degenerate reading is not merely wrong, it is *flattering*. It says the subject holds the full twelve-fold vocabulary. **An instrument whose failure mode is to award everyone full marks will be adopted enthusiastically and will never be questioned by the people it flatters.** That is why this is a result and not a footnote.

And the trap is not really about one-hot vectors. One-hot is the extreme case of the real condition, which is *disjoint support*:

```
DISJOINT SUPPORT WITHOUT ONE-HOT -- the same degeneracy, better disguised

A = {1:0.6, 2:0.5, 4:0.3}   B = {3:0.7, 6:0.4, 9:0.5}   C = {5:0.8, 7:0.2, 11:0.6}

u(A,B) = u(A,C) = u(B,C) = 1.000000
F = 2.484907   N_eff = 12.0000   f = 69.6133 Hz
```

Each framing now looks rich and the reading is still the maximum. A reader who fixed the one-hot problem by "using more states per framing" and stopped there has fixed nothing.

The ceiling is why the trap is damaging rather than merely wrong. The degenerate encoding does not produce a random number; **it produces the ceiling** — the one reading that cannot be improved upon and therefore the one reading that ends the inquiry.

7.5 The overlap condition

> **THE OVERLAP CONDITION [BINDING].** A set of framings may be used to compute an internal surplus only if the framings share support, and the shared states carry enough weight that the inner products are not artefacts of a stray small term. **Where the condition fails, the computation is reported as degenerate and no F_{int} is published.**

The plain-language reading is what makes it obviously right: **two ways of construing *the same situation* must be about the same situation.** If a person's threat-reading and their story-telling have literally no ingredient in common, they are not two framings of one thing; they are two things. **Measure the angle between vectors pointed at different targets and you have measured nothing but the fact that they point at different targets.** A framing set with disjoint support is a category error rendered as a number.

There is a second, quieter reason the condition matters. **Overlap is what makes u variable.** On disjoint supports u is pinned at 1 and carries no information; on overlapping supports it ranges over the interval and its value is a fact about the person. **An instrument that cannot vary cannot measure.**

Overlap arises honestly because **the twelve states are components of specification, not topics.** A threat-reading of a conversation and a story about the same conversation both require a directed feeling — both are *about* something — so both carry state 3, and they may weight it very differently. **That difference is the signal.** Overlap becomes fake when a state is added because the arithmetic needs it rather than because the description requires it, and the guard is procedural: **the support of each framing vector must be fixed and recorded before any inner product is computed.** An operator who has seen $u = 1.000000$ and then revisits the vectors is no longer measuring.

7.6 States and pitch classes are different counts

Adding a state to a chord always adds a pitch. It does not always add a pitch class. A chord containing state 5 that acquires state 10 gains a note two octaves higher and gains no new

class. **The chord gets taller. It does not get harmonically wider.**

The two counts answer different questions. **The state count is what the index measures** — how many distinguishable components the framings collectively require. **The class count is what the folded chord can express** — how many distinct positions in the harmonic vocabulary it occupies. This paper insists only that the two be reported separately and never conflated: reporting the class count as though it were the roster understates the chord, and reporting the state count as though it were the harmonic vocabulary overstates it.

7.7 What raises the width, and what does not

The width is a function of u_{int} alone. Adding a framing genuinely angled from the existing ones raises the mean; adding one that restates an existing framing lowers it; and adding a state to an existing framing can move u either way and must be computed.

> **SCALE INVARIANCE OF THE FRAMING VECTOR.** Multiplying every intensity in a framing by a positive constant leaves the unit vector unchanged and therefore leaves every u, every F_{int} and the width unchanged. A framing at $\{3{:}0.8,\ 5{:}0.4,\ 7{:}0.2\}$ and one at $\{3{:}1.6,\ 5{:}0.8,\ 7{:}0.4\}$ are the *same unit vector*: $u = 0.000000$, $F = 0$, width exactly k. **Feeling the same thing more strongly does not widen the chord. Feeling a differently-shaped thing does.**

The instrument reads composition, not amplitude. Anything to be said about amplitude must be said with a different instrument, and this framework's nearest available proxy is δ.

7.8 Drop-one, and an honest surprise

Removing each framing in turn from the carried Subject of Section 8:

```
DROP-ONE AUDIT (three framings, F_int = 2.164390)

drop threat  ->  remaining pair (mirror, story)   u = 0.809410  F = 2.292890  f = 47.4142 Hz
drop story   ->  remaining pair (mirror, threat)  u = 0.904762  F = 2.393557  f = 57.9892 Hz
drop mirror  ->  remaining pair (story, threat)   u = 0.388360  F = 1.662402  f = 13.4361 Hz
```

Two of the three drops raise the width. This is not a paradox and not a bug, but it is a warning that must be published with any drop-one audit. **The mean pairwise u is an average, and removing a member removes its pairs.** A framing at a moderate angle to two framings that are at a wide angle to each other **drags the average down by existing.**

> **MEAN PAIRWISE u DOES NOT DETERMINE THE JOINT GEOMETRY.** The triangle inequality fails for geometric triples, so a set of pairwise angles is consistent with more than one arrangement, and no scalar summary of pairwise values recovers the arrangement. **N_{eff} computed from a mean is an audited count under a stated postulate, not a measured dimension.**

What follows practically is that a drop-one audit tells you about the average, not about the person's dependence on a framing. The mirror framing is not "the important one" because dropping it costs the most; **it is the one whose removal leaves behind the closest remaining pair.** Those are different claims and the arithmetic supports only the second. The same caution governs roster-scale drop-one audits elsewhere in the series [Curzi, 2026], whose Result Seven records a constructed team in which every one of five removals improves the reading, so a favourable drop-one row **carries no information about the member it names.**

7.9 The minimum roster, and what a chord is not

Two framings are enough to compute a width; one is not. With a single framing there are no pairs, u_{int} is undefined, and the correct output is not $F_{\text{int}} = 0$ but *no reading at all*. **A single framing has no internal geometry, and reporting zero surplus for it confuses "measured to be parallel" with "not measured."**

A chord is not a personality. It is computed from framings of one situation on one occasion. **A chord is not a preference or a value**; nothing in the construction records what the subject wants, endorses or would choose. **And a chord is not evidence about the situation**: the instrument measures the structure of the construals and is silent, completely and by construction, on whether any of them is correct. **A wide chord over a misread situation is a wide chord over a misread situation.** The alternative reading — that internal breadth implies accuracy — is attractive and unsupported by anything in the formalism. The corpus's own belief work makes the matching point with its forbidding clause [Curzi, 2026ab], whose Result One states $E = G \cdot B + \varepsilon$ with $G = 0 \Rightarrow E = 0$ for all B: **breadth of belief buys nothing through a channel of zero gain.**

7.10 Specification and gates

```
STEP 1  BUILD.       Write the twelve-vector for each framing. Fix and record the
                     support BEFORE step 3.
STEP 2  NORMALISE.   x <- x/||x||_2. Check ||x|| = 1.000000.
STEP 3  OVERLAP GATE. For every pair compute the shared support. If ANY pair has
                     empty shared support, STOP. Report DEGENERATE. Publish no F_int.
STEP 4  PAIRWISE.    cos = x . y ; u = 1 - cos^2. Publish ALL pairwise u, not the mean.
STEP 5  AGGREGATE.   u_int = mean over the P = m(m-1)/2 pairs.
STEP 6  SURPLUS.     N_eff = 1 + (N-1)u_int ;  F_int = ln N_eff.
STEP 7  WIDTH.       f_width = k e^(2 F_int) ; octaves = 2 F_int / ln 2.
STEP 8  CHORD.       States = union of supports. Raw pitch k n^2, folded with j.
                     Pitch classes = the distinct odd parts.
STEP 9  INTERVALS.   For each pair, ratio j^2/i^2, cents, and cents mod 1200.

GATE 1  UNIT NORM        ||x|| = 1 to within 1e-9                FAIL -> recompute
GATE 2  OVERLAP          every pair shares at least one state    FAIL -> DEGENERATE
GATE 3  WEIGHT           shared support >= 10% of each squared
                         norm  (A STATED CONVENTION, NOT A LAW)  FAIL -> flag WEAK
GATE 4  DISTINCTNESS     no two framings with u < 0.01           FAIL -> flag RESTATEMENT
GATE 5  CEILING          F_int <= ln N                           FAIL -> arithmetic error
GATE 6  PRE-REGISTRATION supports fixed before inner products    FAIL -> reading void
GATE 7  ROSTER           at least two framings                   FAIL -> NO READING
```

Gate 3's threshold is a convention chosen by the author and badged as one. There is no derivation that puts it at 10 per cent rather than 5 or 20. **What is not optional is choosing one, stating it, and applying it before seeing the result.** Gate 4 is the mirror of Gate 2 and flags rather than fails, because a genuine near-restatement is a finding and should be reported rather than suppressed.

A published reading carries the framing vectors unnormalised and normalised with supports marked; every pairwise u; u_{int}, N_{eff}, F_{int}, the width and the octave figure; the state set and the class set **reported separately**; the interval table raw and folded; the band and j for every folded value; the gate results; the scale-status statement; and the non-resonance and non-diagnostic statements. **The last two are not decoration. They are the two sentences that most reliably fall off a reading as it is copied from one document to another.**

7.11 Failure modes

FM-1 — disjoint support, of which one-hot is the extreme case. *Detection:* Gate 2, and by inspection — **a reading of exactly $F_{\text{int}} = 2.484907$ at $N = 12$ should be treated as a defect until proven otherwise.** *Repair:* ask of each framing which components of the situation it requires to be specified, not which component it is *about*. **It is not repaired by adding states; it is repaired by adding *shared* states.**

FM-2 — manufactured overlap. Having learned Gate 2, an operator adds a small intensity on a shared state so the gate passes. *Detection:* Gate 3 catches the crude version, Gate 6 the honest-operator version. **Neither catches a determined bad actor, and this paper says so plainly. No procedural gate defends against an operator who wants a particular answer; what the gates do is stop an honest operator producing a degenerate answer by accident, and make a dishonest one leave a trace.**

FM-3 — reading the mean as the geometry. Treating $N_{\text{eff}} = 8.7093$ as though it said the framings span 8.71 dimensions. *Repair:* state the roster-reduction postulate wherever N_{eff} is published.

FM-4 — the contested decomposition. Two competent analysts decompose the same framing differently and the discrepancy is reported as measurement noise. **There is no repair that makes the judgment go away.** *Remedy:* publish both decompositions and both readings. **Do not average them; averaging two decompositions produces a third that neither analyst endorses.**

FM-5 — the chord read as a diagnosis. The twelve states are not diagnoses and a chord width is a property of an encoding, not of a person's health. A narrow chord means the framings elicited on that day, under that decomposition, were close to parallel. **It means that and nothing further.**

FM-6 — physical-resonance drift. "The subject's chord is 36.67 Hz" becoming a recommendation to expose someone to a 36.67 Hz tone. *Detection:* any sentence in which the frequency *does* something.

FM-7 — conflating states with pitch classes. *Detection:* the class count is the number of distinct odd parts and can never exceed $\lceil N/2 \rceil$ or the state count.

FM-8 — amplitude mistaken for width. *Detection:* compare unnormalised to normalised vectors; if the change is a uniform scaling, the reading cannot move. *Repair:* **accept the invariance and say what the instrument does not see, rather than making it appear to see it.**

FM-9 — silently changing N. *Detection:* Gate 5. A reading with $F_{\text{int}} > 2.484907$ at $N = 12$ has changed N somewhere.

8 Worked example: the Subject

CONSTRUCTED CASE. This is not an observation of any person, no clinical or diagnostic claim is made or implied, and every parameter is chosen rather than measured.

The situation, stated once: a person has received a message from a colleague that could be read several ways, and must decide how to respond. Three framings are elicited:

```
threat  = {3: 0.8, 5: 0.4, 7: 0.2}    "they are coming at me"
story   = {3: 0.5, 5: 0.7, 10: 0.4}   "the latest instalment of an old pattern"
mirror  = {5: 0.3, 7: 0.8, 10: 0.5}   "what does my reaction say about me"

||threat|| = 0.916515   ||story|| = 0.948683   ||mirror|| = 0.989949

threat  = {3: 0.872872, 5: 0.436436, 7: 0.218218}
story   = {3: 0.527046, 5: 0.737865, 10: 0.421637}
mirror  = {5: 0.303046, 7: 0.808122, 10: 0.505076}
```

8.1 The gates

```
shared support (threat, story)  = {3, 5}       PASS
shared support (threat, mirror) = {5, 7}       PASS
shared support (story,  mirror) = {5, 10}      PASS
```

Every pair shares two states, and the shared states carry substantial weight in both members of each pair, so Gates 1 to 3 pass. State 5 — attributed feeling, the component saying the situation has a cause — is carried by all three framings. **That is what makes them three framings of one situation rather than three situations**, and it is the overlap condition satisfied by content rather than by construction.

8.2 The degenerate calculation, shown first

Encoding each framing by its heaviest state — threat at 3, story at 5, mirror at 7:

```
threat = {3: 1}   story = {5: 1}   mirror = {7: 1}

cos = 0.000000 on all three pairs;  u = 1.000000 on all three pairs
u_int = 1.000000   N_eff = 12.0000   F_int = 2.484907 nats
f_width = k * 144 = 69.6133 Hz   (folded 278.453 Hz, j = +2)
octaves above ground = 7.1699
```

The Subject scores the ceiling. So does every other subject encoded this way. And the damage compounds downstream. Carrying the degenerate width through the corpus's octave-deficit diagnostic at an illustrative cost ratio $c/\eta = 1.5714$:

```
present octaves = log2(69.6133 / 0.483425732988414) = 7.1699
required        = 2.885390 * 1.5714 = 4.5341
octave deficit  = 4.5341 - 7.1699 = -2.6358          "SUSTAINS"
```

The subject is pronounced sustainable by 2.64 octaves, and so is everyone else. A diagnostic that returns the same verdict for every input is not a diagnostic. This is the whole reason the trap is a result: **it does not fail loudly; it fails by succeeding uniformly.**

> **A rounding in the source, recorded.** The source computes this deficit as -2.6364 using the shorthand constant 2.885. With the exact $2/\ln 2 = 2.885390081777927$ the figure is -2.6358. **The difference is the shorthand, not an error**, and this paper quotes the exact value.

8.3 The correct calculation

```
pair              cos theta     u          F_int      N_eff     f = k(1+11u)^2
mirror / story     0.436566   0.809410   2.292890    9.9035     47.4142 Hz
mirror / threat    0.308607   0.904762   2.393557   10.9524     57.9892 Hz
story  / threat    0.782074   0.388360   1.662402    5.2720     13.4361 Hz

folded (220-440):  47.4142 -> 379.314 Hz (j=+3)
                   57.9892 -> 231.957 Hz (j=+2)
                   13.4361 -> 429.956 Hz (j=+5)
octaves above ground: 6.6159, 6.9063, 4.7967
```

All three pairwise values are different, and they are different because the framings are different. That sentence is trivial and it is the entire point: **under the degenerate encoding it was false.** Threat and story are the closest pair at $u = 0.388360$ — both construe the situation as something arriving from outside with a cause. Mirror and threat are the widest at $u = 0.904762$ — one looks outward at an agent, the other inward at the self.

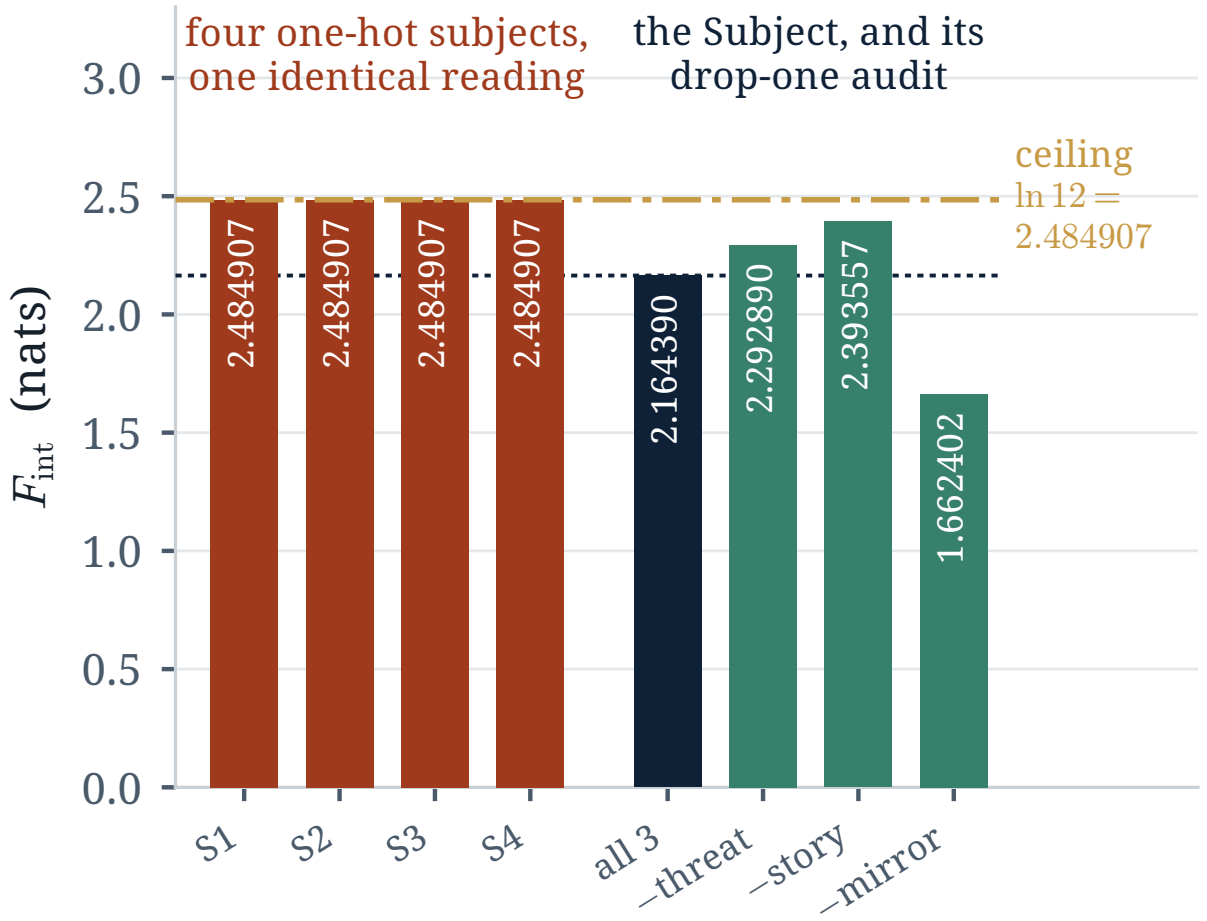


Figure 42.5: **The tidy encoding awards every subject the ceiling, and the honest drop-one audit rises twice out of three**: four unrelated constructed subjects encoded one-hot all read $F_{\text{int}} = \ln 12 = 2.484907$, while the Subject's overlapping framings read 2.164390, and two of the three single-framing removals *raise* it, to 2.292890 and 2.393557. Modelled from $F_{\text{int}} = \ln(1 + 11u_{\text{int}})$ on the constructed framing vectors of Section 8. Built by `figs/src/CCP-042_fig05.py`.

```
u_int   = (0.809410 + 0.904762 + 0.388360)/3 = 0.700844
N_eff   = 1 + 11 * 0.700844 = 8.7093
F_int   = ln 8.7093 = 2.164390 nats
f_width = k * 8.7093^2 = 36.6686 Hz      folded 293.349 Hz (band 220–440, j = +3)
octaves = 2 * 2.164390 / ln 2 = 6.2451

F_int / ln 12 = 87.10%          N_eff / N = 72.58%
```

The chord's width is 36.6686 Hz, which is 6.2451 octaves above the ground essence, and it is the same statement as $F_{\text{int}} = 2.164390$ nats. The Subject holds three framings and audits to an effective 8.71 of the twelve available categories — **which should be read as what it is: a stated construction, not a measured dimension.** N_{eff} is not a count of framings; that three framings can audit to 8.71 categories is exactly the surplus functional's content, that non-parallel states span more than they number.

8.4 The chord and its intervals

```
STATES SPANNED  =  {3, 5, 7, 10}    -- FOUR states
```

```
state  3   directed feeling      raw  4.3508 Hz    folded 278.453 Hz (j=+6)   class 3
state  5   attributed feeling    raw 12.0856 Hz    folded 386.741 Hz (j=+5)   class 5
state  7   reflexive feeling     raw 23.6879 Hz    folded 379.006 Hz (j=+4)   class 7
state 10   narrative feeling     raw 48.3426 Hz    folded 386.741 Hz (j=+3)   class 5
PITCH CLASSES  =  {3, 5, 7}        -- THREE classes
```

Four states, three pitch classes. States 5 and 10 have the same odd part, so they fold to the identical value and occupy one position in the six-position vocabulary $N = 12$ permits. **The fourth state adds a pitch and does not add a class.**

```
pair    essence ratio   raw cents   folded cents   note
3:5        25 / 9        1768.72       568.72
3:7        49 / 9        2933.74       533.74
3:10      100 / 9        4168.72       568.72      same folded interval as 3:5
5:7        49 / 25       1165.02      1165.02      just under an octave
5:10        4 / 1        2400.00         0.00      CLASS COLLISION, exactly 2 octaves
7:10      100 / 49       1234.98        34.98      the 50/49 minimum gap
```

Every ratio is a ratio of squared integers and k appears nowhere. Two entries deserve their own line. $5{:}10$ **at exactly** 0.00 **folded cents is the collision, visible** — and it is also the most *consonant* interval the essence lattice offers [Curzi, 2026], whose Result Two records that the minimum harmonic distance $\mathrm{HD}_{\mathrm{ess}} = 2.0000$ is attained by exactly the six collision pairs $\{1{:}2, 2{:}4, 3{:}6, 4{:}8, 5{:}10, 6{:}12\}$ and by nothing else. **The Subject's chord contains the most consonant interval the system has, and it is also the chord's only redundancy.** And $7{:}10$ **at** 34.98 **folded cents is the tightest pair the system permits**, the $50/49$ of Section 4.5, realised here by a chord that happens to contain both 7 and 10.

8.5 What the reading says, and what it does not

It says: under the declared decomposition, on this situation, on this occasion, the Subject's three framings are substantially non-parallel; the mean non-parallelism is 0.700844; the internal surplus is 2.164390 nats, 87.10 per cent of the maximum this twelve-category ladder can express; the chord spans four states in three pitch classes; and one of the four states is a doubling of another.

It does not say that the Subject is well or unwell, that a wider chord would be better, that a narrower one would be worse, that the Subject's framings are accurate, that the colleague's message means anything in particular, or that any frequency corresponds to a physical oscillation anywhere.

And it does not say that the Subject "is 8.71 people", or has 8.71 perspectives, or contains 8.71 of anything. N_{eff} is an effective count under a stated postulate, computed from a mean of pairwise angles, and because the triangle inequality fails for geometric triples that mean does not determine the joint arrangement.

9 Results

The seven results below are the units other papers in this corpus cite. A citation to this paper should name one of them. Each carries its status: **DERIVED** where it follows from the stated definitions by algebra, **COMPUTED** where it is the output of exact arithmetic on stated parameters, and **POSTULATED** where it is asserted by this framework.

R1 — the odd-part class rule. The pitch class of $f(n) = kn^2$ is the odd part of n, so an index $1 \ldots N$ realises exactly $\lceil N/2 \rceil$ classes — six at $N = 12$, namely $\{1, 3, 5, 7, 9, 11\}$. *Status:* DERIVED (Theorem 4.1, proved in both directions); the class count COMPUTED and verified for every N from 1 to 25. Class-preserving intervals are **even** whole octaves, so the smallest is exactly 2400 cents (Corollary 4.2).

R2 — the intensification/transformation dichotomy. A transition $m \to n$ is a whole number of octaves **iff** m and n share an odd part; 10 of 66 pairs at $N = 12$, or 15.1515 per cent. *Status:* DERIVED (Theorem 5.1) and verified exhaustively over all 66 pairs. **Index distance does not determine the kind**: $3 \to 12$ is $+9$ and an intensification, $11 \to 12$ is $+1$ and a transformation.

R3 — frequency does not distinguish a disposition from an episode. $f = kn^2$ has no time argument, so a mood and a trait at $n = 9$ are the identical tone: 39.1575 Hz raw, 313.260 Hz folded at $j = 3$, class 9, **separated only by** δ**.** *Status:* DERIVED (Theorem 6.1). **This is the paper's sharpest result and it is negative.**

R4 — the storage half-life. $t_{1/2} = \ln 2/(-\ln(1-\delta))$. At $\delta = 0.30$ it is 1.9434 cycles against the naive $1/\delta = 3.3333$, **an overstatement of** 71.52 **per cent** (the source rounds this to 72). At $\delta = 0.09$ it is 7.3496 and at $\delta = 0.02$ it is 34.3096, a ratio of 17.6544. *Status:* DERIVED; the table COMPUTED.

R5 — the one-hot degeneracy. Distinct one-hot framings have disjoint support, hence $u = 1$ for every pair and $F_{\text{int}} = \ln N = 2.484907$ for every subject regardless of content, giving the ceiling width $kN^2 = 69.6133$ Hz. *Status:* DERIVED; demonstrated COMPUTED on four constructed subjects and on a disguised disjoint-support case. **The degenerate encoding does not produce a random number; it produces the ceiling.**

R6 — the $50/49$ **regime.** The minimum folded separation at $N = 12$ is 34.9756 cents at the unique pair $7{:}10$; the cliff at $N = 17$ collapses it to 6.0008 cents ($289/288$) and the second at $N = 41$ to 1.0296 cents ($1682/1681$). *Status:* COMPUTED in exact rational arithmetic; all three are exact rational coincidences. The corresponding signal bound is $\Delta t \geq 10.2882$ ms at the tightest pair [Curzi, 2026z].

R7 — the drop-one surprise. Removing two of three framings from the carried constructed chord *raises* F_{int}, from 2.164390 to 2.292890 and 2.393557; only the third lowers it, to 1.662402. *Status:* COMPUTED. **Mean pairwise non-parallelism does not determine the joint geometry**, and a drop-one audit therefore reports on the average and not on dependence.

CR — counter-result. The frequency register **cannot distinguish a mood from a trait at all**, because $f = kn^2$ contains no time term; and **the cleanest-looking encoding of a framing, the one-hot vector, returns the system's maximum surplus for every subject and therefore measures nothing.** To these the framework adds three of its own: the constant k is chosen, the affine effective count is not a Hill number, and the interpretation of k as a frequency is postulated. The corpus reaches structurally parallel self-demotions elsewhere: its own pre-registered fit **failed its second falsification rule** [Curzi, 2026i], and its channel-separation work states that a property which is necessary is not thereby sufficient [Curzi, 2026], refusing two attractive simplifications outright. **The same discipline applies here: overlap is necessary for a chord reading and is not sufficient for it to mean anything.**

10 Discussion

10.1 What follows

A formalism with two axes is unusual and worth naming. Most measurement schemes have one: more or less. This lattice has a second — pitch class — that is *invariant* along an entire family of changes the first registers as enormous. **That invariance is the content of R2**, and it is why the distinction between intensification and transformation is available at all.

The negative result about time generalises past this framework. Any state description that is a function of a specification alone is silent about persistence, and the temptation at that point is to fold a stability coefficient into the description. **The correct response is to notice that the persistence parameter belongs to a different equation and to keep it there.** The division of labour is exact: the frequency reading decides the sign of the fixed point and contains no δ; δ decides the magnitude and the clock and cannot change the sign [Curzi, 2026], whose Result Three is precisely the absence of δ and N from the threshold.

The degeneracy result generalises furthest. The pattern is *an encoding whose failure mode is a uniformly favourable answer*. Such instruments are adopted enthusiastically and are never questioned by the people they flatter, and the defence is not a better statistic but a

validity gate applied before the arithmetic. **This is the same discipline the corpus applies when it declines an intervention on the ground that its measured benefit is an artefact of the measure** [Curzi, 2026], whose Result Two prices pitch-class folding at exactly one bit for every even N — a cost that is paid whether or not anyone notices.

10.2 What does not follow

The lattice is not a taxonomy of emotion. Six pitch classes is a fact about odd numbers below 12.

Class identity is not sameness. States 1 and 8 fold to the same value across six octaves of raw pitch. **The folded lattice cannot tell them apart, and that is a property of the lattice worth knowing rather than a claim about the states.**

A width is not a verdict about a person. It is a property of an encoding of a set of framings elicited on one occasion under one decomposition.

And the arithmetic is not evidence. Every result in Sections 4, 5 and 7 is a theorem about integers or about unit vectors. **It becomes a statement about a person only through the count, and it inherits every weakness of that count.**

10.3 Relation to the surrounding literature

The interval apparatus — cents, octave equivalence, ratio arithmetic — is standard once an index exists [Rossing et al., 2002, von Helmholtz, 1885, Benson, 2006]. **What the apparatus does not carry with it is any claim about hearing.** Consonance depends on the spectrum of the sound and not on the ratio alone [Sethares, 2005]; interval categories and octave equivalence generalise less far across listeners and species than is usually assumed [Burns, 1999, Hoeschele et al., 2012]; and the octave's perceptual specialness, while real, has a measured and finite size [Demany et al., 2013]. The two classical consonance measures this corpus uses elsewhere [Tenney, 2015, Euler, 1739] are chosen functions on rationals rather than measured facts about hearing, and are badged that way where they are used [Curzi, 2026]. **This paper uses the arithmetic and asserts none of the psychoacoustics.** The one place the corpus's frequencies meet real acoustics is a derived room [Curzi, 2026l], whose Result Two fixes $7.2400 \times 4.3798 \times 2.9319$ m, $V = 92.969$ m^3 and an RT60 of 0.5 s requiring 29.94 m^2 sabins, with the folded set entirely above a Schroeder frequency of 146.7 Hz — **and that is a statement about a room, not about a person.**

The persistence half of the paper sits beside a large literature on habit and disposition that this paper does **not** contribute to. **"Trait" is a word with a substantial psychometric literature attached to it, and this paper is defining a term inside a formalism rather than adding to that literature.** The reliability/validity distinction is the relevant caution [Cronbach and Meehl, 1955, Campbell and Fiske, 1959]: **an agreement statistic measures reliability and not validity**, and a held-out predictive test is required before any of this can be said to work [Nosek et al., 2018].

10.4 The reflexive exposure, and the programme

An instrument that issues a verdict on a person's framings may change those framings by issuing it. The mechanism is documented [Merton, 1948, Rosenthal and Jacobson, 1968] and has been observed in the field [MacKenzie and Millo, 2003], but the effect sizes are contested [McCambridge et al., 2014] and on the best long-run assessment **self-fulfilling prophecies are real but small, do not typically accumulate, and often dissipate** [Jussim and Harber, 2005]. **The correct response is not to suppress the reading but to publish the parameters that produced it.** That is also why any indicator adopted as a target degrades as an indicator [Campbell, 1979, Goodhart, 1975].

The wider programme this belongs to begins with *physique sociale* [Quetelet, 1835] and continues through the modern computational literature [Castellano et al., 2009, Lazer et al., 2009]. **Its characteristic error is Quetelet's** — treating the mean of a distribution as a type rather than as a statistic — and Section 7.8 is this paper's version of that error caught in the act: **a mean of pairwise angles is not an arrangement.** Three further constraints bear and are not comfortable. Human activity is bursty with **no characteristic time scale** [Barabási, 2005], which a half-life claim must confront and can meet only by insisting that its cycles are declared rather than discovered. An explanation is supposed to name a mechanism rather than exhibit a fitted curve [Hedström and Ylikoski, 2010]; **this paper fits nothing, which sidesteps the objection at the cost of having no empirical content to fit.** And **what observation would distinguish this lattice from a neighbouring model that fits equally well?** [Sobkowicz, 2020, Durlauf, 2005] The honest answer is that R1, R2 and R6 are theorems and cannot be so distinguished; only H1 and H2 have empirical content, and neither has been tested. **What the paper is entitled to claim is analytical rather than predictive** [Epstein, 2008, Keuschnigg et al., 2018].

11 Verification

Every step below can be checked on its own. The standard is Occam's: **a postulate that can be dropped without the result changing has been dropped.** Statuses: **ARITHMETIC** (a reader can redo it), **DEFINITION** (true by stipulation), **CITED** (an external source supports it — the key is named), **POSTULATED** (asserted by this framework, no external support).

11.1 The postulates

	postulate	status
P1	The Pythagorean whole tone is the ratio $9{:}8$.	CITED [Barbour, 1951]
P2	φ is the positive root of $x^2 = x + 1$; hence $\varphi^2 = \varphi + 1$.	CITED [Euclid, 1956, Vajda, 1989]
P3	A dimensionless ratio of the form $(9/8)^2/\varphi^2$ **is a frequency in hertz**.	POSTULATED — no source
P4	A state with n distinguishable components has essence frequency kn^2.	POSTULATED — no source
P5	A frequency may be divided by a power of two without changing what it means.	POSTULATED beyond pitch; CITED for pitch only [Shepard, 1982]
P6	n is a cardinality, hence absolute-scaled.	DEFINITION, with the scale typology CITED [Stevens, 1946]
P7	$N_{\text{eff}} = 1 + (N-1)u$ (the affine count axiom).	POSTULATED — design axiom [Curzi, 2026]; **not a Hill number**
P8	A framing admits representation as a unit vector over the twelve categories.	POSTULATED — no source

	postulate	status
P9	Stored order obeys $Q_{t+1} = (1-\delta)Q_t + a$ with $\delta > 0$.	CITED as a form [Oppenheim et al., 1996, Elaydi, 2005]; the *interpretation* POSTULATED

Dropped as unnecessary. No postulate about the *meaning* of any state is needed for R1, R2, R6 or R7: they hold for any index set and any unit vectors whatever the states are called. **No postulate about consonance is needed anywhere in this paper**, and the time–frequency bound contributes to R6's rendering corollary only, not to any lattice result. **P3 and P4 are not needed for R1, R2 or R6 either**, because those are results about ratios in which k cancels.

11.2 The compounding chain

step	consumes	produces	status
L1	P1	$(9/8)^2 = 81/64 = 1.265625$ exactly	ARITHMETIC
L2	P2	$\varphi^2 = \varphi + 1 = 2.618033988749895$	ARITHMETIC
L3	L1, L2	$k = 0.483425732988414$	ARITHMETIC (exact)
L4	L3, **P3**	k carries the unit Hz	**POSTULATED**
L5	L4, **P4**	$f(n) = kn^2$ Hz	**POSTULATED**
L6	L5	$f(m)/f(n) = m^2/n^2$: k cancels from every interval	ARITHMETIC
L7	L6, **P5**	$n = 2^a m \Rightarrow f(n) = 2^{2a} f(m)$: same class — **R1** forward	ARITHMETIC (value); POSTULATED (meaning)
L8	L6	$(m_1/m_2)^2 = 2^j$ with m_i odd forces $m_1 = m_2$ — **R1** converse	ARITHMETIC
L9	L7, L8	classes = odd numbers in $1 \ldots N$, count $\lceil N/2 \rceil$ — **R1**	ARITHMETIC
L10	L8	$j = 2(a_1 - a_2)$ is even: smallest same-class interval = 2400 cents	ARITHMETIC
L11	L9, L10	the dichotomy; 10 of 66 at $N = 12$ — **R2**	ARITHMETIC
L12	L9, exact rationals	min folded separation $50/49$ at $N = 7 \ldots 16$ — **R6**	ARITHMETIC
L13	L12	cliffs at 17 ($289/288$) and 41 ($1682/1681$) — **R6**	ARITHMETIC
L14	L5 by inspection	f has no time argument — **R3**	ARITHMETIC
L15	P9	$Q_t = (1-\delta)^t Q_0$; $t_{1/2} = \ln 2/(-\ln(1-\delta))$ — **R4**	ARITHMETIC

step	consumes	produces	status
L16	L15	$t_{1/2} = 1.9434$ vs $1/\delta = 3.3333$ at $\delta = 0.30$: $+71.52\%$ — **R4**	ARITHMETIC
L17	P9	sign Q_∞ is sign$(\eta F - c)$; $Q_\infty \propto 1/\delta$	ARITHMETIC
L18	P7, P8	$F_{\text{int}} = \ln(1 + (N - 1)u_{\text{int}})$	DEFINITION given P7
L19	L18, disjoint support	$u = 1$, $F_{\text{int}} = \ln N$ for every subject — **R5**	ARITHMETIC
L20	L18, constructed vectors	$u_{\text{int}} = 0.700844$, $F_{\text{int}} = 2.164390$	ARITHMETIC
L21	L20, drop-one	2.292890, 2.393557, 1.662402 — **R7**	ARITHMETIC

11.3 What this chain is and is not

Read the status column. Steps L1, L2, L3, L6 and L8 through L21 are arithmetic and a reader can redo every one of them. **Steps L4 and L5 are postulated, and L7's *interpretation* is postulated while its *value* is arithmetic.** A chain containing postulated steps is not a proof of its conclusion, and this one is not offered as one.

What it is offered as is a separation, and this paper's separation is unusually clean. R1, R2, R6 and R7 are results about *ratios* and *angles*, in which k cancels and P3 and P4 do no work at all. **A reader who rejects the frequency interpretation entirely keeps the odd-part rule, the dichotomy, the regime and the drop-one surprise, restated as facts about integers and unit vectors.** What such a reader loses is only the hertz values, which is exactly what the postulates were asked to supply.

Two statuses that must not be misread. The exponent 2 at L5 is **DEFINED**, not estimated: no fit was performed and the standard procedure for testing a scaling claim [Clauset et al., 2009] was not run. The logarithm at L18 is **DEFINITION** given the affine axiom P7, and the warrant the argument needs is the multiplicative-dynamics one [Peters and Gell-Mann, 2016] rather than the psychophysical one.

The falsifiable content lives at L15 and in H1 and H2, not at L4 and L5. This paper's empirical exposure is therefore smaller than its notation suggests, and its unsupported material is correspondingly more inert. Making point predictions with constants rather than directional ones is the right response to the directional-prediction critique [Meehl, 1990], and printing the status of each postulate is the honest description of what a badged postulate does in a research programme [Lakatos, 1970, Platt, 1964]. The corpus's own honesty apparatus is what enforces that discipline [Curzi, 2026], whose counter-result assigns the framework's central identification to the postulate drawer and states that one of its three declared postulates is a wager whose misfire column exists because wagers can lose. **A postulate table that never assigns anything to the postulate drawer is decoration**, and Section 11.1 assigns four.

12 Limitations

Every check the governing rigour standard specifies is listed here by name, whether or not this paper passes it.

L-1. The dimensional audit does not close. $f(n) = kn^2$ balances only if k carries $[T^{-1}]$, and k is computed as a ratio of two dimensionless quantities. **No operational procedure is supplied by which the constant acquires the dimension of frequency.** This is the paper's largest structural weakness and no citation in the bibliography repairs it. **It touches only the absolute values, not the intervals.**

L-2. The essence law has no external support, and its exponent is **defined, not estimated**: no fit was performed, no alternative was tested, and the standard procedure for a scaling claim [Clauset et al., 2009] was not run because there is no data to run it on. **This paper makes no empirical scaling claim.**

L-3. The constant is chosen. $k = (9/8)^2/\varphi^2$ combines two real objects into a product with no precedent, quoted to fifteen figures **because it is defined**, not because anything was measured to that precision [Markowsky, 1992, Fischler, 1981].

L-4. Octave reduction beyond pitch is unwarranted. Shepard grounds octave equivalence for pitch [Shepard, 1982]; **nothing grounds folding a non-acoustic quantity into a 220–440 Hz band and reading meaning from where it lands.** Every folding result in this paper is a result about the operation, and the *interpretation* of the operation is not supported.

L-5. $N_{\text{eff}} = 1 + (N-1)u$ **is not a Hill number.** It is a linear interpolation between 1 and N and appears in no ecological, economic, political-science or information-theoretic literature. The Hill-number citations apply to $1/\text{HHI}$ and $\exp(H)$ and to nothing else [Hill, 1973, Laakso and Taagepera, 1979].

L-6. The scalar-reduction habit is a single defect with two faces, and the paper concedes both. Michell's objection [Michell, 1997] and the appraisal literature's objection are the same objection in two domains. Emotional states are described as **patterns across simultaneously varying appraisal dimensions** [Smith and Ellsworth, 1985], as **processes in time whose state is a synchronisation** [Scherer, 2001], as individuated by **qualitative core relational theme** [Lazarus, 1991], and as located on **continuous bipolar axes** rather than counted [Russell, 1980]. **A count discards exactly the information those accounts carry, and cannot separate anger from guilt.** The ladder's top rung is contradicted specifically: action readiness is held to be constitutive of emotion as such [Frijda, 1986], and this paper narrows the rung to *bound commitment* — **a narrowing, not an answer.** Two rungs survive [Tracy and Robins, 2004, Premack and Woodruff, 1978]. **Every reading in this paper that passes through the ladder inherits this concession.**

L-7. Conjoint measurement was not tested. The treatment of the elicited intensities as ratio-scaled is provisional; the available remedies [Luce and Tukey, 1964, Rasch, 1960] were not run. The cost of that failure is measured elsewhere in the series at 1065.05 cents [Curzi, 2026].

L-8. Count reproducibility (H2) has not been tested. Whether n is a property of the state or of the analyst is **unknown**. Every frequency here is a frequency *under a declared decomposition*, which is a real thing to have and is less than a measured one.

L-9. Twelve is not justified. The category count exceeds the classical span of immediate absolute judgment [Miller, 1956], which is a caution rather than a refutation, and Section 6.12 shows the truncation it causes: **six of twelve indices have no double on the ladder, so no intensification is available from them at all.**

L-10. No instrument for δ **is supplied.** The framework does not validate one, and the identifiability constraint is severe: **roughly 103 cycles of no-input observation at trait grade.** Every δ in this paper is constructed.

L-11. The trait cut is a convention, not a discovery. There is no δ at which an episode becomes a disposition. **Two readings that declare** W **can be compared; two that do not, cannot.**

L-12. The trust mechanism is an assumption. Section 6.10's liability depends on trust acting through δ rather than η; the literature cited for trust [Gambetta, 1988, Ostrom, 1990] is about reciprocity, monitoring and institutional form, not decay rates. **Which channel it operates through is not settled here.**

L-13. The overlap gate's weight threshold is arbitrary. Ten per cent rather than five or twenty is a stated convention with no derivation, and **no procedural gate defends against an operator who wants a particular answer.**

L-14. Mean pairwise u **does not determine the joint geometry.** The triangle inequality fails for geometric triples, so N_{eff} is an audited count under a stated postulate and **not a measured dimension**; the drop-one surprise of R7 is that failure made visible.

L-15. u is symmetric and psychological similarity is not. Similarity judgments are asymmetric and violate the triangle inequality [Tversky, 1977]; this framework's u assumes metric behaviour.

L-16. The instrument is blind to amplitude. Uniform rescaling of a framing leaves every reading unchanged. That is a property of the construction, **and the correct response is to say what the instrument does not see rather than to make it appear to see it.**

L-17. Every case is constructed. The Subject, the one-hot subjects and every δ are chosen parameter sets. **Their precision is exact because the parameters are exact; that precision does not transfer to a measured case.**

L-18. No model comparison, no statistical claim. Nothing was fitted, so there is no out-of-sample test, no matched null, no information criterion, no p-value, no effect size and no interval. **The corpus's one confirmatory attempt failed its own baseline test** [Curzi, 2026i], which is the relevant precedent for how a fit here should be reported.

L-19. The functional form is underdetermined. A sub-linear power law [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the qualitative facts as well as the logarithm does, and many mechanisms produce indistinguishable curves [Stumpf and Porter, 2012]. **Nothing here discriminates among them.**

L-20. The reflexive specification is not identified [Manski, 1993, Shalizi and Thomas, 2011]. The remedy adopted is stipulation of the reference set [Brock and Durlauf, 2001], and the claim is **a modelling position rather than an identified empirical result.** The framework also exempts its own parameters from its own reflexivity [Lucas, 1976, Hacking, 1995], which is **not resolved.**

L-21. There is no conservation law, and Q is written with the grammar of a conserved quantity [Mirowski, 1989]; validity is claimed only within a specified arrangement [Cartwright, 1999]. **Ergodicity is assumed and not argued** [Peters, 2019].

L-22. Boundary behaviour. $\delta \to 0^+$ gives no fixed point at all — a pathology, not an ideal, since nothing decays out of the stock; $\delta \to 1$ makes the store memoryless and drives $t_{1/2} \to 0$; $u \to 1$ gives the ceiling and $u \to 0$ collapses the width to k exactly. $N \to \infty$ **raises the class ceiling without bound and simultaneously destroys the rendering regime**, collapsing the minimum separation by a factor of 5.83 in cents at $N = 17$. **The source examines none of these limits together and this paper states them here.**

L-23. Two roundings in the source, printed rather than silently corrected. The half-life error at $\delta = 0.30$ is 71.52 per cent and the source prints 72. The post-intervention dyad's $f = 35.8907$ Hz and $Q_\infty = +6.0846$ recomputed from the printed $u = 0.692400$ disagree in the fourth significant figure with the source's 35.8882 and $+6.0844$, **because the printed u is a rounding of** 0.69237163**.** Neither is an error in a theorem; both will be found by any reader who recomputes.

L-24. The 2.885 shorthand. The exact constant is $2/\ln 2 = 2.885390081777927$; the shorthand costs about a thousandth of an octave at the ceiling and about 0.0006 octaves in the degenerate worked deficit of Section 8.2. **Every computed figure in this paper uses the exact value.**

L-25. A published reading can change what it reads. The reflexive exposure of Section 10.4 is not controlled for and cannot be controlled for by an instrument acting alone.

13 Conclusion

What is established, and it is arithmetic. The pitch class of a squared integer index under octave folding is the odd part of that index, so two states share a class **if and only if** they share an odd part, and an index $1 \dots N$ realises exactly $\lceil N/2 \rceil$ classes — six at $N = 12$, namely $\{1, 3, 5, 7, 9, 11\}$. Class-preserving intervals are necessarily **even** whole octaves, so the smallest is exactly 2400 cents. This partitions transitions into intensification and transformation, and only 10 of 66 pairs at $N = 12$ — 15.1515 per cent — are class-preserving; **index distance does not determine which kind a move is**, since $3 \to 12$ is nine rungs and an intensification while $11 \to 12$ is one rung and a transformation. The minimum folded separation at $N = 12$ is exactly $50/49 = 34.9756$ cents at the unique pair $7{:}10$, and the lattice falls off a cliff at $N = 17$ to

$289/288 = 6.0008$ cents and again at $N = 41$ to $1682/1681 = 1.0296$ cents. **Every one of these is a theorem about integers, and k cancels from all of them.**

What is established, and it is negative. The essence law has no time argument, so **a mood and a standing disposition at the same index are the identical tone**: at $n = 9$ both read 39.1575 Hz raw, 313.260 Hz folded, class 9. What separates them is the storage law's decay rate, whose half-life is $\ln 2/(-\ln(1-\delta))$ — 1.9434 cycles at $\delta = 0.30$ against a naive $1/\delta = 3.3333$, an overstatement of 71.52 per cent, and 34.3096 at $\delta = 0.02$, a ratio of 17.6544. **Durability and revisability are the same parameter read in two directions**, so there is no setting that is cheap at both. And **the tidiest available encoding of a framing destroys the measurement**: distinct one-hot vectors are orthogonal by construction, so every subject scores the ceiling $\ln 12 = 2.484907$ and the width 69.6133 Hz regardless of content. **An instrument whose failure mode is a uniformly favourable answer will be adopted enthusiastically and never questioned.** On the constructed Subject the honest overlapping reading is $F_{\text{int}} = 2.164390$ nats and 36.6686 Hz, and **two of the three drop-one removals raise it**, which settles that mean pairwise non-parallelism does not determine the joint geometry.

What is proposed rather than shown. That a dimensionless ratio of a tuning interval to the square of the golden section is a frequency; that the essence of an n-component state is kn^2; that octave equivalence may be carried from pitch to a quantity that is not a pitch; and that a way of construing a situation admits representation as a unit vector over twelve categories. **These four are postulates of this framework with no external support**, labelled at every appearance and separable from everything above: **a reader who rejects all four keeps the odd-part rule, the dichotomy, the regime, the mood/trait result and the degeneracy, restated as facts about integers, about unit vectors and about a first-order difference equation.**

What is conceded. The affect literature the ladder names does not underwrite a scalar index, an ordering, or a nesting, and one of its standing theses contradicts the ladder's top rung outright. That objection and Michell's are **the same objection twice** — scalar reduction in two domains — and the corpus should say so rather than meet each as a surprise.

What is not claimed at all. No person, tissue, cell, molecule, room or institution is claimed to resonate at any frequency named here. The twelve states are not diagnoses, a pitch class is not a type of person, and no reading in this paper is clinical, diagnostic, therapeutic or prognostic.

14 Declarations

Data and code availability. No dataset was collected for this study and none exists. Every number in the paper is the output of exact arithmetic on parameters stated in the text, computed in double precision except where the text says exact rational arithmetic was used; every case is constructed and is labelled as constructed at the point of use. The figures are generated by the Python scripts named in their captions, which take their inputs from the tables in the text and are available from the author on request. **No empirical measurement of any human being underlies any figure, table or number in this paper.** A frequency law reads as though something were measured; nothing was.

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the analysis, the figures and the manuscript.

Ethics. No human-subjects research was conducted and no personal data were processed. The Subject, the one-hot subjects and every constructed trajectory are invented parameter sets and are not derived from, and do not describe, any person. This paper asserts several perceptual and structural claims — that octave equivalence is a usable structure beyond pitch, that interval arithmetic in cents is the correct apparatus, and that a psychological state can be indexed by a count of components — **without any perceptual or psychometric experiment behind them.** Those claims rest on the cited literature, are limited by that literature's own

boundaries (§10.3), and **where the literature is in tension with them that tension is printed** (§4.4, L-6).

Preregistration. No test reported in this paper was preregistered, because no test was run. The one confirmatory fit attempted elsewhere in this corpus was preregistered and **failed its own second falsification rule** [Curzi, 2026i]; that outcome is reported here as it is reported there.

Disclosure of rendering constraints. Where any value in this paper is rendered as sound, the band and the transposition exponent are published with it, the values are exact ratios and must not be tempered, and **nothing may be masked, placed below threshold, or concealed from the listener.** A message hidden from its recipient cannot be consented to.

Clinical disclaimer. The twelve state classes are not diagnostic categories. Nothing in this paper is an assessment, a screening instrument, a treatment or advice about any person's health, and nothing in it may be presented as such.

Chapter 43

Consonance, Collision, and the Divided Self: A Firewall Between Harmonic Dissonance and Psychological Distress

Abstract. We audit the harmonic structure of the essence lattice and then block the inference the structure invites. Under two declared consonance measures — Tenney harmonic distance $\mathrm{HD}(p/q) = \log_2(pq)$ and Euler's *gradus suavitatis*, both badged as chosen functions on rationals rather than measured facts about hearing — squaring the index ratio doubles the harmonic distance exactly, so $\mathrm{HD}_{\text{ess}} = 2\log_2(pq)$ and $G_{\text{ess}} = 2G(pq) - 1$, verified over all sixty-six index pairs at $N = 12$. The consequence is an identity: the six most consonant pairs, at the global minimum $\mathrm{HD}_{\text{ess}} = 2.0000$, are exactly the six pitch-class collisions $\{1{:}2,\ 2{:}4,\ 3{:}6,\ 4{:}8,\ 5{:}10,\ 6{:}12\}$. Maximum consonance and maximum redundancy are the same event. The firewall result is that F_{int} cannot report on distress because it contains no term that could: relabelling state 10 as state 9 in the carried constructed subject leaves $F_{\text{int}} = 2.164390$ unchanged to 10^{-12} while removing the collision and changing the chord's class set, so the width does not even determine the chord. A restatement of one framing collapses the reading from $F_{\text{int}} = 2.164390$ and 36.6686 Hz to 1.333662 and 6.9620 Hz, a loss of 2.3970 octaves. At dyad scope $f_{\text{rel}} = k(1 + (N-1)u)^2$ is convex in u while cents are concave, so hertz and pitch disagree about which unit of nonparallelism is worth most. The constant $k = (9/8)^2/\varphi^2$ is defined, not measured, and its interpretation as a frequency is postulated.

Keywords: harmonic distance; gradus suavitatis; collision; internal surplus; relabelling invariance; firewall; relationship frequency; convexity

Series: C. Frequency, Essence, and Mind **Paper:** CCP-043

1 Introduction

A framework that assigns a frequency to a psychological state acquires, for free and without asking, an entire harmonic vocabulary. Once states have pitches, pairs of states have intervals; once there are intervals, some are consonant and some are not; and once that is on the page, a reader will complete the sentence. **The completion is always the same, and it is always wrong: a person whose internal states are harmonically dissonant is in psychological difficulty, and a person whose internal states are consonant is at ease.**

This paper audits the harmonic structure the framework actually has, and then blocks that inference formally rather than deprecating it. Both halves are necessary. Deprecating an inference is a request; blocking one is a proof that the quantity in question contains no term the inference could use.

The audit produces a result that is worth having in its own right. Under two declared consonance measures — Tenney's harmonic distance $\mathrm{HD}(p/q) = \log_2(pq)$ [Tenney, 2015] and Euler's *gradus suavitatis* [Euler, 1739], **both of them chosen functions on rationals and neither a measurement of anybody's hearing** — squaring the index ratio doubles the harmonic distance exactly. The essence lattice is therefore twice as harmonically remote as the corresponding integer lattice, everywhere, and the consequence is an identity: **the six most consonant index pairs at $N = 12$, at the global minimum $\mathrm{HD}_{\mathrm{ess}} = 2.0000$, are exactly the six pitch-class collisions.** Maximum consonance and maximum redundancy are the same event. A chord of framings that is comfortably related without being partly redundant is not merely hard to build in this lattice; it is not there.

That result is what makes the firewall necessary, because a reader who has just learned that consonance is redundancy will read redundancy as a deficiency and dissonance as a virtue, and from there will read internal difficulty as informative. **The formal block is that the internal surplus F_{int} is invariant under relabelling of the states**, so it does not determine the chord it is computed from — relabelling state 10 as state 9 in the carried constructed subject leaves $F_{\mathrm{int}} = 2.164390$ unchanged to twelve decimal places while removing a pitch-class collision and changing the chord's class set from three classes to four. A quantity that does not determine the chord cannot determine anything the chord does not determine, and the chord determines nothing about anybody's suffering. The inference fails two levels above where anyone would look for it.

The third part of the paper takes the same apparatus to the smallest relational unit. At dyad scope the relationship frequency is $f_{\mathrm{rel}}(u) = k(1 + (N-1)u)^2$, which is **convex** in the non-parallelism u, while the same quantity measured in cents is **concave**. Both statements are true and they are about different functions; the practical consequence is that hertz and pitch disagree about which unit of non-parallelism is worth most, and only one of them enters the sustainability verdict.

What this paper does not claim. The consonance measures are chosen, not measured. Consonance as heard depends on the spectrum of the sound and not on the ratio alone [Sethares, 2005], roughness models make different predictions from ratio-complexity models [Plomp and Levelt, 1965, Kameoka and Kuriyagawa, 1969, Terhardt, 1974], and interval categories and octave equivalence generalise less far across listeners than is usually assumed [Burns, 1999, Burns and Ward, 1978, Hoeschele et al., 2012]. Octave equivalence itself is formalised for **pitch** [Shepard, 1982], with a measured size to its perceptual specialness [Demany et al., 2013]; **its extension to a quantity that is not a pitch is a modelling choice of this framework with no external warrant.** The constant $k = (9/8)^2/\varphi^2$ is a **defined** number whose two components are real — the Pythagorean whole tone [Barbour, 1951] and the golden section [Euclid, 1956] — and whose product, expressed as a frequency, has no precedent; the standard deflation of golden-section claims [Markowsky, 1992] is cited here and again wherever φ appears.

Contributions

1. The doubling theorem: $\mathrm{HD}_{\mathrm{ess}}(i,j) = 2\log_2(pq)$ and $G_{\mathrm{ess}} = 2G(pq) - 1$ for the essence ratio p^2/q^2, verified exhaustively over all sixty-six index pairs at $N = 12$.
2. The consonance–redundancy identity: the minimum $\mathrm{HD}_{\mathrm{ess}} = 2.0000$ is attained by exactly the six collision pairs and by nothing else; and zero folded separation holds **if and only if** two indices share an odd part, verified over all sixty-six pairs.
3. The firewall: F_{int} is invariant under relabelling and therefore does not determine the chord, which formally blocks any inference from the width statistic to distress or health.
4. The restatement collapse: replacing one framing with a paraphrase of another takes F_{int} from 2.164390 to 1.333662 and the chord width from 36.6686 to 6.9620 Hz — a loss

of 2.3970 octaves, more than a third of what the twelve-category ladder can supply.
5. The extremal chords at $N = 12$: the most consonant four-state chord is $\{1, 2, 4, 8\}$, mean $\mathrm{HD}_{\mathrm{ess}} = 3.3333$, six collisions from six pairs, **one** pitch class; the most dissonant is $\{7, 9, 10, 11\}$, mean 12.7586, no collisions, four classes, and an entirely unremarkable description of a person attending to a social situation.
6. The convexity result: $\mathrm{d}f/\mathrm{d}u$ rises from 10.6354 to 127.6244 Hz per unit u while $\mathrm{d}(\text{cents})/\mathrm{d}u$ falls from 38087 to 3174.
7. The eleven coincidences: $f_{\mathrm{rel}}(u) = km^2$ exactly at $u = (m-1)/(N-1)$, so a dyad's interval lands on a state's note at eleven exact values of u — an arithmetic coincidence to which this paper attaches no meaning, and says so.

2 Hypotheses, definitions, and postulates

Most of this paper is not hypothesis-testing, and saying so is the first honest act available to it. Sections 4 and 6 are exhaustive enumerations over a finite lattice: every claim in them is settled by computing all sixty-six pairs or all 495 four-state chords, and no observation could bear on them. Section 5 is a procedure. What follows separates the few claims that could be wrong from the many that are definitions, postulates or modelling choices.

2.1 Hypotheses proper

H1 — the relabelling hypothesis, as an empirical claim about use. If F_{int} carried information about distress, then chords with equal F_{int} and different collision structure would be associated with different distress. **Predicted:** no association, because the statistic contains no term that could carry one. **Falsified by:** a measured association between collision count and any distress measure, at fixed F_{int}, in a sample where the decomposition was fixed in advance. **Status:** the invariance itself is arithmetic and is proved in §4.6. The *empirical* prediction is untested. **No such sample exists and none was collected for this paper.**

H2 — the elicitation-stability hypothesis. A drop in F_{int} between two occasions reflects a change in what was elicited rather than a change in the person. **Predicted:** re-eliciting the same subject on the same situation with a different analyst recovers the earlier width more often than re-eliciting with the same analyst on a later occasion recovers it. **Falsified by:** the opposite ordering, or by no difference between the two. **Status: untested.** §5.14b sets out the double-coding procedure that would test it.

H3 — the ratio-scale hypothesis, inherited. The intensities carry ratio-scale information to the extent magnitude estimation delivers it [Stevens, 1957]. **Falsified by:** exponents that move with the experimenter's chosen stimulus range [Teghtsoonian, 1971] or by the response biases catalogued against the method [Poulton, 1968]. **Status: untested here**; carried from the minimum-frequency paper [Curzi, 2026], whose Result Seven measures the cost of the failure at 1065.05 cents.

H4 — the ceiling hypothesis, inherited. No relation with $c/\eta > \ln N$ sustains. **Falsified by:** one counterexample. **Status:** derived [Curzi, 2026]; **untested.** It remains the only strict universal in this series and therefore the only claim here that a single observation could kill.

2.2 Definitions — true by stipulation

- **Def. 1.** $\mathrm{HD}(p/q) := \log_2(pq)$, Tenney's harmonic distance on a ratio in lowest terms, defined in his 1983 essay on Cage and the theory of harmony and cited here from the collected edition [Tenney, 2015].

- **Def. 2.** $G(n) := 1 + \sum_i e_i(p_i - 1)$ over the prime factorisation, Euler's *gradus suavitatis* [Euler, 1739]; $G(p/q) = G(pq)$.
- **Def. 3.** A **collision** is a pair of indices whose folded separation is zero, equivalently (Theorem 4.2) a pair sharing an odd part.
- **Def. 4.** $F_{\text{int}} := \ln(1 + (N-1)\,\bar{u})$ on the mean pairwise non-parallelism of a set of framing vectors; the **chord width** is $k\,e^{2F_{\text{int}}}$.
- **Def. 5.** $k := (9/8)^2/\varphi^2 = 0.483425732988414$. **A defined constant**, exact by construction; the fifteen figures are exact digits of a definition and carry no measurement uncertainty, because nothing was measured.

2.3 Postulates — asserted, with no external support

- **P-a.** The essence law $f(n) = kn^2$, inherited [Curzi, 2026]. **No external support.** The exponent 2 is **defined by the model, not estimated from data**: no fit was performed, no alternative form was tested, and the goodness-of-fit and likelihood-ratio procedure a scaling claim requires [Clauset et al., 2009] was not run, because there is no data to run it on.
- **P-b.** The attachment of hertz to k. Computing $(9/8)^2/\varphi^2$ is arithmetic; asserting that the resulting dimensionless number **is a frequency** is a separate and much stronger claim, and no operational procedure is supplied by which the constant acquires the dimension of inverse time.
- **P-c.** The affine effective count $N_{\text{eff}} = 1 + (N-1)u$ [Curzi, 2026,c], a design axiom and **not a Hill number of any order**; the Hill family [Hill, 1973] and its rediscovery in political science [Laakso and Taagepera, 1979] ground $1/\text{HHI}$ and $\exp(H)$ exactly — as the corpus's concentration work uses them [Curzi, 2026] — and ground this construction not at all.

2.4 Modelling choices

- **M-a. Two consonance measures, both chosen.** Neither Tenney's harmonic distance [Tenney, 2015] nor Euler's gradus [Euler, 1739] is an experiment. They are defined functions on rationals, adopted because they are explicit and checkable, and where they disagree this paper publishes both and does not adjudicate. **Neither is a measurement of what any listener finds pleasant.** The measured literature runs on roughness and on the spectrum of the sound [von Helmholtz, 1885, Plomp and Levelt, 1965, Kameoka and Kuriyagawa, 1969, Terhardt, 1974, Sethares, 2005], and it does not agree with ratio complexity in general.
- **M-b. Octave reduction applied to a non-acoustic quantity** [Shepard, 1982, Demany et al., 2013]. The operation is well defined; its interpretation is a choice of this framework with no external warrant. Every claim in §4.4 that involves *folding* is a claim about arithmetic, not about hearing.
- **M-c.** $N = 12$, inherited, and not shown to be the right twelve [Curzi, 2026].
- **M-d. Non-negative intensities.** The encoding cannot represent an internal framing that is the negation of another, which makes the opposition branch of the collapse result unreachable at internal scope. This is a real limitation of the encoding and is stated as one in §4.5.

2.5 What is expository

Section 5 is a worksheet: a procedure with blanks, whose organising question is *what must be recorded alongside a personal reading so that a second analyst can recompute it and disagree with it in a specific place.* It argues no claim and tests no hypothesis. Its value is procedural, and §5.1 says so in the paper's own words.

3 Framework, dimensional audit, and method

3.1 What is inherited and what is new

The essence law, the twelve-state ladder and the elicitation procedure are inherited [Curzi, 2026]; the fold and the odd-part class rule are inherited [Curzi, 2026]; the effective count and its logarithmic ceiling are inherited [Curzi, 2026c]; the doubling of harmonic distance was first established for the relational lattice in Series B [Curzi, 2026l], and is re-proved here because the proof is three lines and the paper should not rest on a reference for something a reader can check. **What is new is the identity between the consonance minimum and the collision set, the relabelling firewall, the extremal-chord enumeration, and the convexity result at dyad scope.**

3.2 Analogy discipline

Every borrowing in this paper is category 2 — a shared mathematical form — and none is category 1. The lattice arithmetic of ratios, harmonic distance and octave equivalence is genuine mathematics and genuine music theory, and it is used because the same arithmetic applies, not because any physical or perceptual result transfers. **The frequency interpretation sitting on top of it is category 3, an analogy**, and it is labelled as one at every appearance. The failure mode the standard names — letting a category-2 or category-3 use borrow the certainty of a category-1 derivation — is exactly the failure this paper's firewall exists to prevent, one level up. The wider programme of statistical physics applied to social dynamics has a survey that fixes what such models may be held to [Castellano et al., 2009], and its boundaries are the boundaries claimed here.

3.3 Dimensional audit

symbol	quantity	dimension	unit	status
i, j, n	state indices	$[1]$	—	absolute scale; counts
p/q	index ratio in lowest terms	$[1]$	—	exact rational
HD, $\mathrm{HD}_{\mathrm{ess}}$	harmonic distance	$[1]$	bits	$\log_2$ of a pure number
G, G_{ess}	gradus suavitatis	$[1]$	—	integer-valued
cents	interval	$[1]$	cents	$1200 \log_2$ of a ratio
u, $\bar{u}$	non-parallelism	$[1]$	—	dimensionless, $[0, 1]$
F_{int}	internal surplus	$[1]$	nats	a logarithm of a count
k	reference constant	$[T^{-1}]$ **asserted**	Hz **asserted**	**defined as a dimensionless ratio; the unit is postulated**
f, f_{rel}, f_{width}	essence frequency	$[T^{-1}]$	Hz	inherits its dimension entirely from k
η, c, δ	conversion, cost, decay	see [Curzi, 2026]	per declared period	ratio-scaled **by assumption**

Does it balance? Everything in Sections 4 and 6 that is expressed as a *ratio, an interval, a distance or a class* is dimensionless and balances exactly, and none of it depends on k. Every quantity expressed in **hertz** balances only because k is asserted to carry $[T^{-1}]$, and **the source supplies no operational procedure by which it acquires that dimension.** This is stated again in §10 by name. The practical consequence is a clean separation: *the harmonic results of this paper are dimensionally sound; the frequency labels attached to them are not.*

3.4 Scale type, and the strongest objection

HD and G are functions of exact rationals and inherit the absolute scale of the indices. F_{int} is a logarithm of an affine function of u and is ratio-scaled given that u is; u is ratio-scaled given that the intensities are, which requires magnitude estimation to deliver what it is held to deliver [Stevens, 1957] and is disputed [Teghtsoonian, 1971, Poulton, 1968, Krueger, 1989]. **The strongest single objection to this framework is Michell's** — that psychology assumes quantitative structure without testing for it [Michell, 1997, 1999]. It applies here in full. The formal remedies exist [Luce and Tukey, 1964, Krantz et al., 1971, Rasch, 1960] and **none was run.** Permissible operations on a statistic are fixed by its scale type, and what follows from a scale type for the permissibility of a statistic is itself disputed [Luce, 2002]; a raw surplus figure without its ceiling $\ln N$ is not interpretable, and this paper reports the ceiling wherever it reports a surplus.

3.5 The aggregate does not determine the joint structure

The firewall of §4.6 is a special case of a general fact worth stating separately, because it recurs across the corpus. **A summary statistic computed from a set does not determine the set.** Three incompatible constructs travel under the word *diversity* — separation, variety and disparity [Harrison and Klein, 2007] — and a single number in one of them fixes nothing in the others; the problem-solving result usually cited for diversity [Hong and Page, 2004] carries a standing formal objection [Thompson, 2014] that must be cited with it; and aggregation over heterogeneous units does not in general yield a representative unit [Kirman, 1992]. The relabelling demonstration is the sharpest available instance: two chords, identical width to twelve decimal places, different class sets, different collision counts.

3.6 Identification — the objection this framework cannot presently answer

The corpus's reflexive loop $T(x) = W(x, M(x))$ [Curzi, 2026] is a linear-in-means reflexive specification, and **the endogenous social effect in such a specification is not identified from observational data** [Manski, 1993]; influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011]. This paper adopts the remedy of **stipulation** — the reference set is the declared twelve-state ladder and the declared decomposition, fixed in advance and published with every reading, which is what makes a restricted treatment of social interaction defensible where an unrestricted one is not [Brock and Durlauf, 2001]. **The reflexive claim is accordingly a modelling position, not an identified empirical result.**

The exposure is sharper here than elsewhere in the corpus, because this paper's subject is a reading of a person's internal state, and **a published reading of an internal state is exactly the kind of thing that can make itself true.** Expectancy effects are real [Rosenthal and Jacobson, 1968, Merton, 1948], a theory becoming true by being used has been documented in the field with dates and prices [MacKenzie and Millo, 2003], classification changes the classified [Hacking, 1995], and a taught theory can become self-validating [Ferraro et al., 2005]. Against that: the best long-run assessment of the classroom case finds self-fulfilling prophecies **real but small**, not typically cumulative, and often dissipating, with most expectations turning out to be accurate rather than causal [Jussim and Harber, 2005], and the effect sizes attached to one famous instance are contested [McCambridge et al., 2014]. **The separation statement of SPEC-4.3 is mandatory for this reason and not as a courtesy.** It is also why the framework's

own parameters cannot consistently be exempted from its own reflexivity [Lucas, 1976] — an inconsistency this paper states rather than resolves.

3.7 Method

Every number in this paper is exact arithmetic on stated inputs. The sixty-six pair values and the 495 four-state chords were enumerated exhaustively; the identities in §4.4 are biconditionals verified over the complete lattice, not tendencies estimated from a sample. **No observational data were collected, no human subjects participated, no perceptual experiment was run, and no empirical dataset underlies any figure, table or number.** Every case is constructed and labelled as constructed. The figures are produced by the scripts named in their captions.

4 Dissonance and the divided self

The octave-folding paper [Curzi, 2026] established that a person holding several framings at once is a chord, that its intervals are ratios of squares, and that its width is the internal surplus `F_int`. It left one question deliberately open. The carried Subject's chord contains states 5 and 10, which fold to the identical pitch. Is that a redundancy in the Subject, or only a redundancy in the vocabulary the ladder makes available?

This section answers it, and the answer is uncomfortable enough that the section has to be built around a firewall rather than around a result.

The harmonic-lattice analysis of Series B [Curzi, 2026l] proved that the essence lattice **doubles** Tenney harmonic distance: every essence pair is exactly twice as harmonically remote as the corresponding integer-ratio pair. From that theorem a second result follows immediately, and the harmonic-lattice analysis of Series B [Curzi, 2026l] stated it:

> **CONSONANCE IS REDUNDANCY.** The pairs that sound most related are precisely the pairs that occupy the same pitch class — the pairs the lattice cannot distinguish. **In the essence lattice, harmonic agreement and informational redundancy are the same phenomenon.**

Carry that into the internal scope of the octave-folding paper [Curzi, 2026] and the consequence is arithmetic and immediate:

```
A PERFECTLY CONSONANT SELF IS A PERFECTLY REDUNDANT ONE

    all framings parallel    =>   u_int = 0
                             =>   N_eff = 1 + 11 * 0 = 1
                             =>   F_int = ln 1 = 0 nats
                             =>   f_width = k * 1^2 = k = 0.4834 Hz, the ground essence
```

Zero surplus. No width. A single note played several times.

That sentence is true, it is derived rather than asserted, and it is also the sentence in this framework most likely to be misused. So before the section proves anything, it fixes what it is and is not claiming, and it will fix it again in the middle and again at the end.

> **THE SEPARATION, STATED FIRST AND RESTATED THROUGHOUT.**
>
> **This analysis is about information content. It is not about pain.**
>
> `F_int` measures how much order is available in a set of framings under a stated encoding. It does not measure distress, and it does not measure well-being. **Internal difference that generates surplus is NOT distress, and distress is not a source**

of surplus. They are different quantities, they are not computed from each other, and the formalism contains no term that would connect them.

Nothing in this section recommends suffering, values suffering, or claims that suffering is informative. Nothing in it is clinical advice, and the twelve states are not diagnoses. A person in pain is not thereby better instrumented, and a person at ease is not thereby impoverished.

The section says exactly one thing about redundancy, and it is a statement about an encoding: **framings that make the same assignments carry the same information once.** That is all it says, and §4.6 shows formally why nothing further can be derived from it.

All cases in this section are CONSTRUCTED and labelled as such, and every extension of the carried Subject is labelled as an extension.

The Law

4.1 What "dissonance" means here, and the four things it does not mean The word is doing technical work and it will be misheard unless the work is specified.

```
DISSONANCE, IN THIS FRAMEWORK

    A computed property of a PAIR OF INDICES.
    Given indices i and j, reduce j/i to p/q in lowest terms; the essence interval is
    p^2/q^2; its dissonance is a chosen distance measure on that ratio.

    It is a fact about arithmetic. It contains no term for a person, an experience,
    a preference, or a state of health.
```

Four things it is not, listed because each has been read into the word by someone reading a draft of this section:

It is not psychological conflict. Two framings can be maximally non-parallel and coexist without friction. Reading a colleague's message as a threat and also as an episode in a long story are not at war; they are two things a person is doing at once, and a person can do both comfortably.

It is not ambivalence. Ambivalence is a state about a decision. Dissonance here is a relation between indices, computed before any decision is in view.

It is not distress. This is the load-bearing one and it is the subject of §4.6.

It is not cognitive dissonance in the social-psychological sense. That is a different construct with its own literature, its own operationalisations, and no relation to the arithmetic here beyond a shared word. This framework does not import it, does not cite it as support, and does not claim its findings. Where a term collides with an established construct, the collision is named and the constructs are kept apart.

4.2 The two consonance measures, both chosen The harmonic-lattice analysis of Series B [Curzi, 2026l] used two indices and this section uses the same two, with the same badge.

```
TENNEY HARMONIC DISTANCE       HD(p/q) = log2(p * q)
    Tenney, "John Cage and the Theory of Harmony", 1983; From Scratch, Illinois, 2015.
    A defined measure of the complexity of a frequency ratio. Larger = more remote.
```

```
EULER'S GRADUS SUAVITATIS      G(n) = 1 + sum over prime factors of e_i * (prime_i - 1)
    Euler, Tentamen novae theoriae musicae, 1739.  G(p/q) = G(p*q). Larger = less smooth.
```

BOTH ARE CHOSEN MEASURES WITH A DEFINED ARITHMETIC, NOT MEASURED FACTS ABOUT HUMAN HEARING [POSTULATE]. Neither is an experiment. They are defined functions on rationals, adopted because they are explicit and checkable. Where they disagree, the framework publishes both and does not adjudicate.

The primary sources, and a correction to how the first of them is usually cited. The *gradus suavitatis* is Euler's, from the *Tentamen novae theoriae musicae* of 1739 [Euler, 1739]. Harmonic distance is Tenney's, and it matters *which* Tenney: the measure is defined in his 1983 essay "John Cage and the Theory of Harmony", cited here from the collected *From Scratch* [Tenney, 2015], where harmonic space is a lattice, harmonic distance is its city-block metric, and $\mathrm{HD}(f_a, f_b) \propto \log a + \log b = \log(ab)$. **It is not defined in Tenney's better-known *A History of 'Consonance' and 'Dissonance'* (1988), which is a historical survey of the two concepts and not the source of this function.** The distinction is recorded because the 1988 volume is the work a reader is likeliest to reach for on the strength of its title, and because the source manuscript of this paper named the measure by its author's surname alone, with no year and no work attached — an attribution that is incomplete rather than incorrect, but which resolves to the wrong book if a reader completes it in the obvious way.

4.3 The doubling theorem, restated

THE DOUBLING THEOREM (Series B, the harmonic-lattice analysis of Series B [Curzi, 2026l]). *For any index pair reducing to `p/q`, the essence interval `p^2/q^2` has exactly twice the Tenney harmonic distance of the ratio `p/q`, and Euler gradus `2G - 1`.*

Proof. `HD(p^2/q^2) = log2(p^2 q^2) = 2 log2(pq) = 2 HD(p/q)`. For the gradus: squaring doubles every exponent in the prime factorisation of `pq`, so the sum `sum e_i (prime_i - 1)` doubles, and `G(p^2 q^2) = 1 + 2 sum e_i (prime_i - 1) = 2(1 + sum e_i (prime_i - 1)) - 1 = 2 G(pq) - 1`. ■

Verified exhaustively over all sixty-six pairs `1 <= i < j <= 12`; both identities exact in every case.

Consonance is scarce in the essence lattice, and it is scarce by a factor of exactly two. The perfect fifth `3:2` — the most consonant interval after the octave in ordinary just intonation — becomes `9:4` here, an octave plus a whole tone, at double the harmonic distance. The major third `5:4` becomes `25:16`.

There is a consequence of the doubling that matters specifically at internal scope and that Series B, working at relational scope, had no occasion to draw.

A chord built on the essence ladder cannot contain a mild interval. In ordinary just intonation a set of four pitches will typically contain several relations of low harmonic distance — fifths, thirds, fourths — and the ear has somewhere to rest. In the essence lattice every one of those relations has been pushed to twice its distance, so the only low-distance relations that survive are the ones that were already trivial: the octave-equivalent pairs, which are the collisions. The lattice offers exactly two kinds of interval — the redundant kind and the remote kind — and almost nothing in between.

Look at the ranking below and the shape is immediate: six pairs at `HD_ess = 2.0000`, then a gap to 3.1699, and then a long tail running to 14.0888 with fifty-six of the sixty-six pairs above 3.1699. The distribution is not a gradient. It is a small cluster of redundancies and a large spread of remoteness. That is what "consonance is scarce by a factor of exactly two" looks like when the scarcity is enumerated rather than described.

For a section about the internal case, the reading is this: **there is no such thing as a chord of framings that is comfortably related without being partly redundant.** The middle ground

the metaphor promises — framings that harmonise without repeating — is not a region of this lattice. It is not that it is hard to reach. It is not there. Whether that is a discovery about selves or an artefact of the square law is exactly the kind of question this framework is obliged to leave open, and it leaves it open: **the doubling theorem is a theorem about ratios, and the square law under it is a postulate.**

4.4 The identity: the most consonant pairs are exactly the collisions Rank all sixty-six index pairs at `N = 12` by essence harmonic distance. The head of the list:

HD_ess	G_ess	pair	essence ratio	raw cents	folded cents	same class?
2.0000	3	1:2	4/1	2400.00	0.0000	YES
2.0000	3	2:4	4/1	2400.00	0.0000	YES
2.0000	3	3:6	4/1	2400.00	0.0000	YES
2.0000	3	4:8	4/1	2400.00	0.0000	YES
2.0000	3	5:10	4/1	2400.00	0.0000	YES
2.0000	3	6:12	4/1	2400.00	0.0000	YES
3.1699	5	1:3	9/1	3803.91	203.9100	no
3.1699	5	2:6	9/1	3803.91	203.9100	no
3.1699	5	3:9	9/1	3803.91	203.9100	no
3.1699	5	4:12	9/1	3803.91	203.9100	no
4.0000	5	1:4	16/1	4800.00	0.0000	YES
4.0000	5	2:8	16/1	4800.00	0.0000	YES
4.0000	5	3:12	16/1	4800.00	0.0000	YES
4.6439	9	1:5	25/1	5572.63	772.6274	no
...						
14.0888	29	11:12	144/121	301.27	301.2741	no

Three exact statements follow, and the section is careful to give the precise ones rather than the convenient one.

> **(A) THE MINIMUM IS EXACTLY THE 2:1 COLLISIONS.** The minimum essence harmonic distance over all sixty-six nontrivial pairs is `HD_ess = 2.0000`, and it is attained by exactly six pairs: `{1:2, 2:4, 3:6, 4:8, 5:10, 6:12}`. Every one of them is a pitch-class collision, and no non-collision attains it. Verified exhaustively.
>
> **(B) ZERO FOLDED SEPARATION IS EXACTLY COLLISION.** Across all sixty-six pairs, the folded interval is exactly 0.0000 cents **if and only if** the two indices share an odd part. Verified exhaustively: true for all ten collision pairs, false for all fifty-six others. This is a biconditional, not a tendency.
>
> **(C) THE HONEST CAVEAT — the ordering is not total.** It is **not** true that every collision is more consonant than every non-collision. `1:3` sits at `HD_ess = 3.1699`, more consonant than the collisions `1:4` (4.0000), `2:8` (4.0000), `3:12` (4.0000) and `1:8` (6.0000). The collision set is dramatically more consonant on average — mean `HD_ess` 3.0000 over the ten collisions against 8.6617 over the fifty-six non-collisions — but the ranges overlap and the framework says so.

Statement (C) is the kind of thing an account drops when it wants a clean slogan. It is kept because The harmonic-lattice analysis of Series B [Curzi, 2026l]'s own rule requires it: **where two defensible indices disagree, or where a claim is true at the extremum but not across the order, publish the precise version.** The slogan "consonance is redundancy" is exactly true at the minimum and true on average; it is not a total order, and a reader who builds on the total-order version will eventually be wrong about `1:4`.

One more exact fact from the same table, carried forward because the Subject will need it:

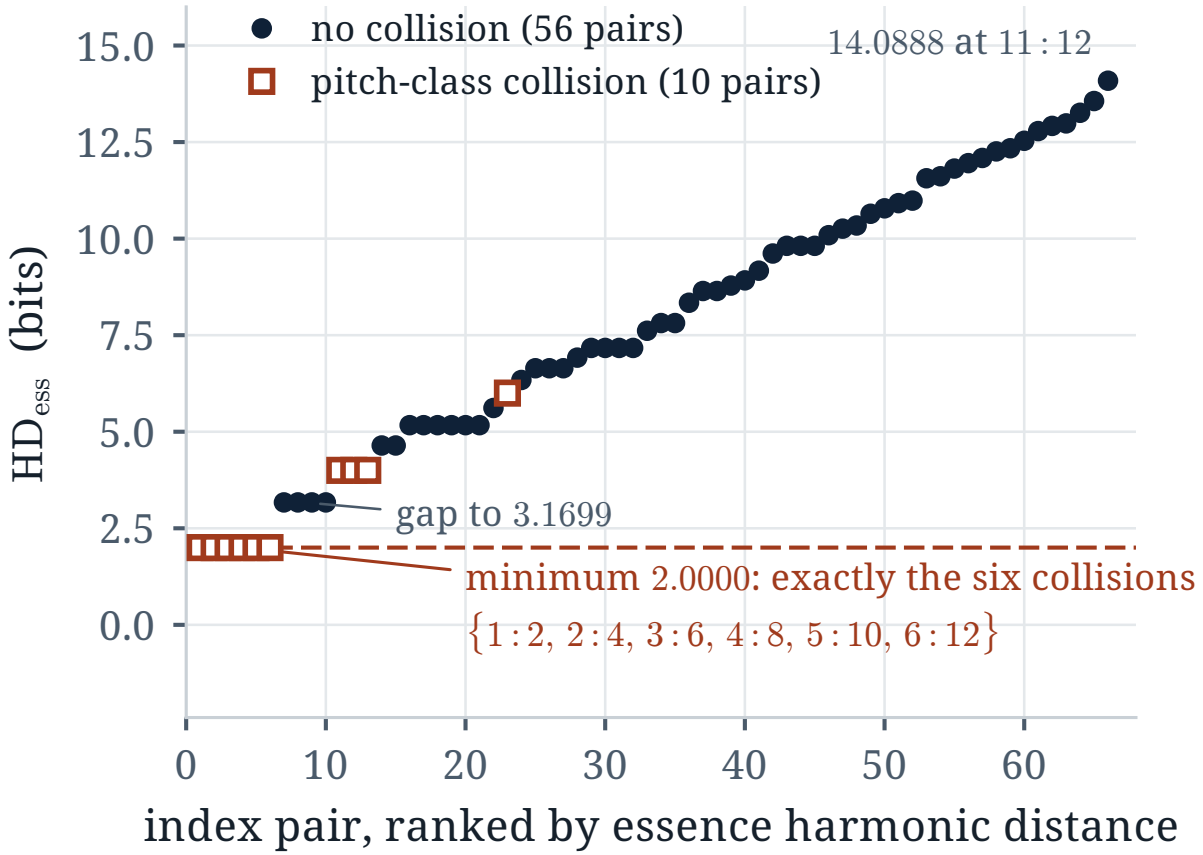


Figure 43.1: Consonance in the essence lattice is a small cluster of redundancies and a long tail of remoteness, not a gradient: six pairs sit at the global minimum 2.0000 and every one of them is a pitch-class collision, then a gap to 3.1699, then a tail to 14.0888. Enumerated over all sixty-six index pairs at $N = 12$ from $\mathrm{HD}_{\mathrm{ess}} = 2\log_2(pq)$ — a complete enumeration, not a sample. The measure is a **chosen** function on rationals and not a measurement of hearing (§2.4 M-a). Built by `figs/src/CCP-043_fig01.py`.

```
The smallest NONZERO folded separation in the whole N = 12 lattice is 34.9756 cents,
= 1200 * log2(50/49), and it is attained by EXACTLY ONE pair: 7:10.
```

The octave-folding paper [Curzi, 2026] established that `N = 12` sits below the `N = 17` minimum-gap cliff and that the framework therefore renders in the 50/49 regime. The pair that realises the regime is unique, and it is `7:10`.

4.4b The extremal chords, computed The identity in (A) and (B) has a consequence that can be exhibited rather than argued: among all four-state chords available at `N = 12` — there are `C(12,4) = 495` of them — there is a most consonant one and a most dissonant one, and they are what the identity predicts.

```
EXTREMAL FOUR-STATE CHORDS AT N = 12   (all 495 enumerated)

MOST CONSONANT   {1, 2, 4, 8}
    mean HD_ess = 3.3333   min = 2.0000   max = 6.0000
    collisions  = 6 of 6 pairs           pitch classes = 1

MOST DISSONANT   {7, 9, 10, 11}
    mean HD_ess = 12.7586  min = 11.9546  max = 13.5627
    collisions  = 0 of 6 pairs           pitch classes = 4

THE SUBJECT      {3, 5, 7, 10}
    mean HD_ess = 8.4882   min = 2.0000   max = 12.2586
    collisions  = 1 of 6 pairs           pitch classes = 3
```

The most consonant four-state chord in the entire system is a chord of one pitch class. States 1, 2, 4 and 8 all have odd part 1; every one of their six pairs is a collision; folded, all four arrive at 247.514 Hz — one note, four times. That is the identity made visible: maximal consonance is maximal redundancy, and here they are the same chord.

The other pole is equally instructive. `{7, 9, 10, 11}` has no collisions at all, four distinct classes, and a mean harmonic distance nearly four times the consonant extreme. It is also, and this is worth noticing, an entirely ordinary set of states: reflexive, normative, narrative, and reciprocal feeling. **Nothing about holding those four components is difficult, unusual, or distressing. The chord is maximally dissonant in the lattice's arithmetic and completely unremarkable as a description of a person attending to a social situation.** If ever a case were needed against reading harmonic dissonance as psychological difficulty, this is it, and it is enumerated rather than asserted.

Two further counts from the same enumeration, useful for calibration:

```
of the 495 four-state chords at N = 12:
    159 have ZERO collisions               (32.1%)
    the maximum possible collision count in a four-state chord is 6,
    attained only by {1, 2, 4, 8}
the single-class families are: {1,2,4,8}, {3,6,12}, {5,10}, and singletons
```

Only three pitch classes at `N = 12` contain more than one index — class 1 with four members, class 3 with three, class 5 with two. Classes 7, 9 and 11 are singletons, so any pair drawn from `{7, 9, 11}` and any state outside their own doubling is automatically collision-free. The redundancy structure of the whole lattice is that small and that explicit, and it is entirely a fact about which integers under twelve are powers of two times which odd numbers.

4.5 The consequence at internal scope Now put (A) and (B) together with the octave-folding paper [Curzi, 2026]'s construction.

Two framings that assign intensity in the same proportions across the same states are, after normalisation, the **same unit vector**. Their inner product is 1, their `u` is 0, and they contribute nothing to the mean. Extend to a whole roster:

```
THE REDUNDANT SELF

    every pair parallel   =>   u_int = 0
                          =>   N_eff = 1
                          =>   F_int = ln 1 = 0
                          =>   f_width = k = 0.483426 Hz
                          =>   octaves above ground = 0.0000
```

A person with forty ways of describing their situation, all of them the same description in different words, has an internal surplus of exactly zero. The internal-surplus analysis of Series B put it as a single window with forty panes. They will experience this as thoroughness, because forty descriptions feel like forty descriptions, and the arithmetic is the only thing in the room that knows otherwise. The external analogue is well documented: informational cascades make a sequence of independent-looking judgments carry the information of the first one only [Bikhchandani et al., 1992, Banerjee, 1992], and social influence measurably collapses the effective number of independent estimates in a group [Lorenz et al., 2011]. **Apparent breadth is not breadth**, and this framework's arithmetic is one more instrument that says so.

And the harmonic version of the same statement: the chord such a person produces sits entirely inside one pitch class if the framings share support, or spans several pitches with zero width if they do not. The most consonant possible internal chord is the one with the least in it.

There is a branch of this that Series B's CC-R8 makes at relational scope, and it deserves its internal statement here because at internal scope it behaves differently.

> **CC-R8 AT INTERNAL SCOPE, AND ITS LIMIT.** Because `u = 1 − cos^2 theta`, perfect opposition (`cos = −1`) gives `u = 0` exactly as perfect agreement does — the contrarian and the sycophant sound the identical note. The relationship-frequency paper [Curzi, 2026], whose three-route concordance holds the carried dyad to a 70.52-cent spread, owns that result at relational scope. Internally, the opposition branch is UNREACHABLE under this framework's encoding, because intensities are non-negative by construction, so `cos theta >= 0` for every pair of framing vectors and `u = 0` can arise only from parallelism. Verified: over 100,000 random non-negative unit pairs the minimum observed cosine was 0.260867, and the bound `cos >= 0` is exact, not empirical.

That is a real limitation of the internal encoding and it is stated as one. A framing that is the exact negation of another cannot be represented here, because the twelve states record which components a construal requires, not which direction it points. A reader who wants an internal analogue of the contrarian must build a signed encoding, and this framework does not have one.

4.6 THE FIREWALL — why nothing about distress follows, formally This is the section the section exists for, and it makes a formal argument rather than a reassurance, because reassurances are ignored and arguments can be checked.

The claim to be blocked is: *low `F_int` indicates a problem, high `F_int` indicates health, and therefore internal difference — including painful internal difference — should be sought.*

Every step of that is unsupported, and the formalism can be made to demonstrate it.

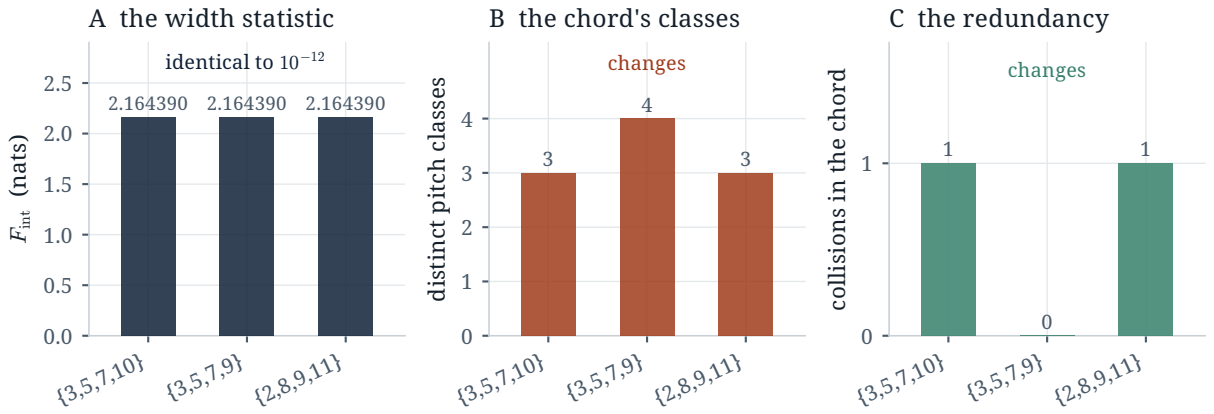


Figure 43.2: The width statistic does not determine the chord: three relabellings of the same constructed subject give internal surplus identical to twelve decimal places while the number of distinct pitch classes and the collision count both change. Since F_{int} does not fix the chord, it cannot fix anything the chord does not fix. **Modelled, not measured:** the subject and its relabellings are constructed parameter sets; relabelling permutes coordinates and therefore leaves every inner product exactly unchanged. Built by `figs/src/CCP-043_fig02.py`.

First: `F_int` is a function of `u_int` alone, and `u_int` is a mean of pairwise angles. It contains no term for valence, for intensity, for duration, for the subject's evaluation of any framing, or for anything else that a word like "distress" would have to refer to. A quantity cannot report on a variable it does not contain.

Second, and stronger: `F_int` is invariant under relabelling of the states, so it does not even determine the chord. Take the carried Subject and relabel state 10 as state 9, leaving every intensity untouched. Inner products are unchanged, because relabelling is a permutation of coordinates:

```
THE RELABELLING TEST  (EXTENSION 3.1-e of the carried Subject, LABELLED)

SUBJECT, states {3,5,7,10}
    u = 0.388360, 0.904762, 0.809410     u_int = 0.700844
    F_int = 2.164390    N_eff = 8.7093    f = 36.6686 Hz
    pitch classes {3,5,7} -- THREE      collisions: (5,10)
    chord HD_ess: 7.8138 8.7846 9.8138 10.2586 2.0000 12.2586   mean 8.4882  min 2.0000

RELABELLED 10 -> 9, states {3,5,7,9}
    u = 0.388360, 0.904762, 0.809410     u_int = 0.700844
    F_int = 2.164390    N_eff = 8.7093    f = 36.6686 Hz
    pitch classes {3,5,7,9} -- FOUR     collisions: NONE
    chord HD_ess: 7.8138 8.7846 3.1699 10.2586 10.9837 11.9546   mean 8.8275  min 3.1699

difference in F_int: 0.000000000000
```

Two chords with identical internal surplus to twelve decimal places: one contains a pitch-class collision, the other contains none. A third relabelling, to states `{2,8,9,11}`, again gives `F_int = 2.164390` exactly, with classes `{1,9,11}` and the collision `(2,8)`.

> **THE FIBER ARGUMENT.** `F_int` is many-to-one on chords. Chords that differ in their state set, their pitch classes, their collision structure, and their entire harmonic content map to the same `F_int`. A quantity that does not determine the chord cannot determine anything the chord does not determine, and the chord determines

nothing about a person's suffering. The inference from `F_int` to distress fails two levels above where anyone would look for it.

Third: the four quadrants are all occupiable, and the formalism forbids ruling any of them out.

- *High `F_int`, no distress.* A person who reads a situation three genuinely different ways and is perfectly comfortable doing so. Nothing in the construction prevents this. It is probably the common case.
- *High `F_int`, severe distress.* The same structure in someone having a terrible week. The arithmetic is identical.
- *Low `F_int`, no distress.* A person with a settled, coherent single reading of a situation, at ease. The formalism calls this zero surplus. It does not call it a problem.
- *Low `F_int`, severe distress.* A person locked into one construal they cannot escape.

All four are consistent with the formalism, because the formalism has nothing to say about the second coordinate. A measure that is compatible with every value of a variable carries no information about that variable. That is the firewall, and it is a theorem about the construction rather than a promise about its use.

Fourth, the ethical statement, which does not follow from the arithmetic and is made on its own authority: nothing in this framework recommends that anyone acquire, retain, or cultivate suffering. **Suffering is not a resource, is not an input to any equation here, and appears nowhere in the formalism.** Where a reader finds a sentence in this framework that seems to praise internal division, that sentence is about the *width of an encoding of construals* and about nothing else, and it should be read that way or discarded.

Fifth, the clinical statement, restated because it is a standing prohibition of this framework: **the twelve states are not diagnoses. A chord reading is not an assessment. No claim, diagnostic, therapeutic, or prognostic, is made anywhere in this section, and nothing here is advice about any person's care.** A reader who wants to know whether something is wrong should ask someone qualified to answer that, and this instrument is not qualified to answer it and does not attempt to.

The Engineering

4.7 The consonance audit of a chord Given a chord — a set of spanned states from the octave-folding paper [Curzi, 2026]'s SPEC-9.2 — three quantities summarise its harmonic content, and all three are computed from the index set alone.

```
CONSONANCE AUDIT

    for every pair (i, j) in the chord, with j/i reduced to p/q:
        HD_ess(i,j)  = 2 * log2(p * q)
        G_ess(i,j)   = 2 * G(p*q) - 1
        folded cents = ( 1200 * log2(j^2/i^2) ) mod 1200

    MIN HD_ess       the chord's most consonant interval
    MEAN HD_ess      the chord's overall remoteness
    COLLISION COUNT  the number of pairs with folded cents = 0
                     equivalently, pairs sharing an odd part
```

The collision count is the redundancy count, by result (B), exactly. It is not a proxy and not an estimate: folded separation is zero if and only if the odd parts agree, so counting collisions and counting exact pitch-class redundancies are the same operation.

4.8 What a collision costs and what it does not Here is where the section must be careful in the second direction, because having established that consonance is redundancy it would be easy to conclude that collisions are waste. **The arithmetic does not support that.**

A collision means two states occupy one pitch class. It **costs** exactly one thing: a position in the harmonic vocabulary, of which `N = 12` supplies only `ceil(12/2) = 6`. A chord with a collision occupies fewer distinct classes than its state count, and that is a real constraint on what the folded chord can express.

It does **not** cost the following, and each is worth denying explicitly:

It does not cost surplus. `F_int` is computed from framing angles, not from the chord's index set. The relabelling test in §4.6 shows a collision appearing and disappearing with `F_int` fixed to twelve decimal places. Collisions and surplus are independent.

It does not mean the two states are the same component. State 5 is attributed feeling — a construal that names a cause. State 10 is narrative feeling — a construal that places the situation in a story. The ladder says these differ by five components of specification. They collide in pitch class because `10 = 2 * 5` and the odd-part rule folds by powers of two. That is a property of the arithmetic of the index, not a claim that the two components are interchangeable.

It does not mean the raw chord has fewer notes. In the raw chord states 5 and 10 are 12.0856 Hz and 48.3426 Hz — two octaves apart, two distinct pitches. Only the fold makes them one.

So what *is* a collision, honestly stated?

> **A COLLISION IS A LIMIT ON THE FOLDED REPRESENTATION, NOT A FINDING ABOUT THE SUBJECT.** It says that if this chord were rendered by octave transposition into a single band, two of its members would arrive at the same place. Everything a collision reports is a fact about the index ladder and the fold. Nothing it reports is a fact about the person.

That is the answer to the question the octave-folding paper [Curzi, 2026] left open, and it is deliberately the deflationary one.

4.9 The restatement collapse, and how fast it happens The genuine internal redundancy — the one that *does* cost surplus — is not a pitch-class collision. It is near-parallel framings, and it is expensive.

```
EXTENSION 3.1-a OF THE CARRIED SUBJECT (LABELLED): the threat framing, restated
three times in different words. Same shape, small perturbations.

    t1 = {3:0.80, 5:0.40, 7:0.20}
    t2 = {3:0.78, 5:0.42, 7:0.21}
    t3 = {3:0.82, 5:0.38, 7:0.19}

    u = 0.001034, 0.001005, 0.004074     u_int = 0.002038
    N_eff = 1.0224     F_int = 0.022166 nats
    f_width = 0.5053 Hz   (against the ground essence k = 0.4834 Hz)
    octaves above ground = 0.0640
```

Three framings, and 0.0640 octaves of width. The subject has produced three descriptions and one framing. They will experience this as having thought about it from several angles.

The exactly-parallel version removes even the perturbation:

```
EXTENSION 3.1-b (LABELLED): the same framing at three intensities.

    p1 = {3:0.8, 5:0.4, 7:0.2}
    p2 = {3:1.6, 5:0.8, 7:0.4}      (everything twice as strong)
```

```
p3 = {3:0.4, 5:0.2, 7:0.1}      (everything half as strong)

u = 0.000000, 0.000000, 0.000000     u_int = 0.000000
N_eff = 1.0000     F_int = 0.000000 nats
f_width = 0.483426 Hz = k exactly
octaves above ground = 0.0000
```

Scaling a framing changes nothing. This is the invariance of the octave-folding paper [Curzi, 2026] seen from the other side: feeling the same thing more strongly, or less strongly, moves the reading by exactly zero. **The instrument reads composition, not amplitude, and here that means a person who is having a much harder time has not thereby acquired any surplus.** That is the arithmetic contradicting the romanticised reading directly, in the place where the romanticised reading would most like to be true.

4.10 The cost of one substitution The most instructive intermediate case replaces one of the Subject's three framings with a near-copy of another and computes what it costs.

```
EXTENSION 3.1-d OF THE CARRIED SUBJECT (LABELLED): mirror replaced by a
restatement of threat.

    threat  = {3:0.8,  5:0.4,  7:0.2}     unchanged
    story   = {3:0.5,  5:0.7, 10:0.4}     unchanged
    mirror' = {3:0.78, 5:0.42, 7:0.21}    a restatement of threat

    u(threat, story)   = 0.388360     unchanged
    u(threat, mirror') = 0.001034     was 0.904762
    u(story,  mirror') = 0.372856     was 0.809410

    u_int   = 0.254083      was 0.700844
    N_eff   = 3.7949        was 8.7093
    F_int   = 1.333662      was 2.164390
    f_width = 6.9620 Hz     was 36.6686 Hz
    octaves = 3.8481        was 6.2451

    LOSS: -0.830728 nats = -2.3970 octaves
```

The marginal contribution of one member to a set's effective count is a quantity the corpus measures elsewhere in a different domain: removing one corridor from a four-member bloc moves the bloc-HHI to 0.391448 and the effective count to 2.5546, so 11.8 per cent of exposure buys 0.014316 of surplus — 3.79 per cent of the total gain [Curzi, 2026]. The arithmetic here is the same shape and the conclusion is the same: **a member that duplicates another contributes almost nothing, and the loss is computable in advance.**

Substituting one framing for a restatement of another costs 2.40 octaves of width — more than a third of everything the twelve-category ladder can supply. The subject still holds three framings. Two of them are one framing.

Note carefully what the loss is *not* attributed to. It is not attributed to the subject becoming less capable, less well, or less anything. The reading changed because the encoding of the third framing changed. If the elicitation on a given day produces a restatement where it previously produced a distinct construal, the honest report is that the elicitation produced a restatement. Whether that reflects a change in the person, a change in the situation, a change in the analyst, or a change in the subject's willingness to say a second thing is not determined by the number, and Section 5's worksheet requires all four to be considered before any of them is asserted.

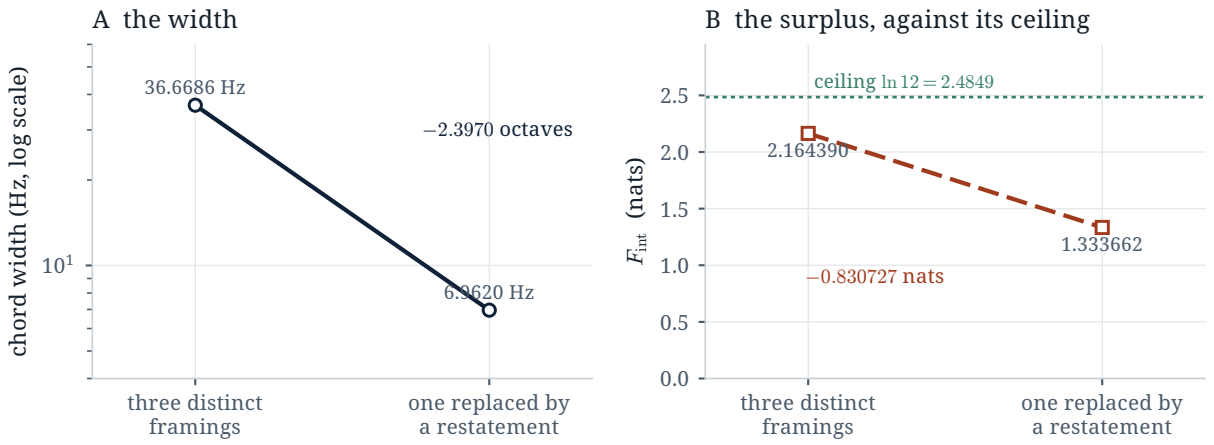


Figure 43.3: Replacing one of three framings with a paraphrase of another costs 2.3970 octaves of chord width and 0.830727 nats of internal surplus — more than a third of what a twelve-category ladder can supply — while the subject still holds three framings. Panel B plots the surplus against its ceiling $\ln 12 = 2.484907$, because a surplus figure without its ceiling is not interpretable. **Modelled:** both framing sets are constructed and labelled as extensions of the carried subject. Built by `figs/src/CCP-043_fig05.py`.

4.11 What actually widens a chord, stated so it cannot be misread Given all of the above, the engineering question is: if an analyst wants a wider reading, what raises it?

A framing that is angled from the existing ones. That is the whole answer. u rises when the squared alignment falls, and the squared alignment falls when a framing assigns intensity in a different shape across the states.

What does **not** raise it, each denied because each has been proposed:

- **Amplifying an existing framing.** Zero effect. Scale invariance, §4.9.
- **Adding a restatement.** Negative effect. Extension 3.1-d, −2.40 octaves.
- **Adding a framing with disjoint support.** Raises the number to the ceiling and destroys its meaning. The octave-folding paper [Curzi, 2026]; the reading becomes degenerate and must not be published.
- **Distress. Not an input. It appears in no equation in this framework.** There is no term for it, no coefficient on it, and no path by which it could enter the computation. A person who has a worse week does not get a wider chord, and this framework does not suggest that anyone should have a worse week.

And the honest general statement, which is a statement about method and not about persons:

> **THE ONLY LEVER ON `F_int` IS THE ANGULAR DIVERSITY OF WHAT IS ELICITED.** Everything else in the construction is invariant. Whether a given person's elicited framings are angled is a fact about the elicitation and the occasion, and this instrument reports it rather than explaining it.

4.12 "The divided self" — the phrase in the title, and what it is allowed to mean The section is called *Dissonance and the Divided Self*, and the second half of that title is a borrowed phrase with a long and largely clinical history behind it. **This framework borrows the phrase and none of the history.**

What the formalism licenses is narrow and can be said in one line: **a person holding several non-parallel framings is, under this encoding, a chord rather than a note.** That is a structural description of an encoding. It is not a description of fracture, of a self at war with itself,

of dissociation, or of any condition, and none of those readings survives contact with §4.6's relabelling test.

There is a version of "divided" that the arithmetic does support, and it is almost the opposite of the dramatic one. **The genuinely divided self, in this system's terms, is the one with the *least* going on: framings so nearly parallel that the person has one instrument and several names for it.** Extension 3.1-b reads `F_int = 0.000000` — the ground essence, no width — and it is not a portrait of turmoil. It is a portrait of unanimity. The internal-surplus analysis of Series B called it a single window with forty panes, and the reason the phrase lands is that unanimity inside one head is invisible from inside that head.

But even this must not be run in reverse. A person reading `F_int = 0` on one situation on one occasion has produced parallel framings on that situation on that occasion. **They are not thereby divided, undivided, well, unwell, rigid, or anything else. The number is about the framings.** The section's title is a title.

4.13 Reading a chord more than once Because every reading in this section is of one situation on one occasion, the obvious next question is what happens when the same subject is read twice, and it is worth answering here because the answer constrains what a single reading can be used for.

The octave-folding paper [Curzi, 2026] established the difference between a trait's frequency and a mood's as a question about `delta` — how fast a reading decays. The same question applies to a chord: is `F_int` a property that persists? Nothing established so far answers that, because nothing in them involves time.

What the construction permits is a comparison, and only a comparison:

```
REPEAT READING, THE ONLY LEGITIMATE FORM

    same subject, same ladder, same analyst, same elicitation script,
    DIFFERENT situation or different occasion

    report:  F_int(1), F_int(2), and the difference
    report:  the chord of each, and whether the state sets overlap
    report:  what changed in the procedure -- ideally nothing

    DO NOT report: a trend from two points
    DO NOT report: a change as a change in the person
```

Three sources can move a repeat reading and the number cannot separate them: the person, the situation, and the elicitation. A drop from 2.164390 to 1.333662 — the size of Extension 3.1-d's substitution — is equally consistent with a subject who framed the second situation more narrowly, a second situation that admits fewer framings, and an analyst who ran a shorter session. A reading that attributes the change to the first of those three has made a claim the arithmetic does not support.

This is why Section 5's worksheet has a procedure block that must be filled in before the computation block, and why the procedure block is reproduced identically on every repeat.

The Specification

SPEC-4.1 — The consonance audit

```
INPUT   a chord: the set S of states spanned, from SPEC-9.2 Step 8.

STEP 1  for each unordered pair (i,j) in S with i < j:
```

```
            reduce j/i to p/q in lowest terms
            essence ratio      = p^2 / q^2
            raw cents          = 1200 * log2(j^2 / i^2)
            folded cents       = raw cents mod 1200
            HD_ess             = 2 * log2(p*q)
            G_ess              = 2 * G(p*q) - 1
            collision          = TRUE iff odd(i) == odd(j)
                               equivalently iff folded cents == 0

STEP 2  MIN HD_ess, MEAN HD_ess, COLLISION COUNT.

STEP 3  report the class set = { odd(n) : n in S } and |class set|,
        separately from |S|.

STEP 4  flag the 50/49 pair if present: any pair with folded cents = 34.9756.
        At N = 12 there is exactly one such pair, 7:10.
```

SPEC-4.2 — The redundancy audit (distinct from SPEC-4.1) The consonance audit is about the chord. The redundancy audit is about the framings, and the two are independent by §4.6. **Both are required; neither substitutes for the other.**

```
INPUT   the normalised framing vectors from SPEC-9.2 Step 2.

STEP 1  every pairwise u, published individually.

STEP 2  RESTATEMENT FLAG: any pair with u < 0.01 is flagged RESTATEMENT.
        (The 0.01 threshold is a stated convention chosen by the author,
        not a derived value. State it; apply it before seeing results.)

STEP 3  EFFECTIVE FRAMING COUNT: report m (nominal framings) and N_eff
        (audited categories) SEPARATELY, with the roster reduction postulate
        stated. They are different counts and neither is the other.

STEP 4  If any pair is flagged RESTATEMENT, report the reading WITH the flag.
        Do not drop the framing to improve the number.
```

Step 4 matters. **A restatement is a finding, not an error.** Dropping it raises `F_int` and destroys the only information the reading contained about the subject's framing behaviour on that occasion.

SPEC-4.3 — The mandatory separation statement **Every published consonance or redundancy audit carries the following, in full, in the reading itself and not in a footnote:**

> *This reading describes the information content of a set of framings under a stated encoding. It does not measure distress, well-being, health, or function, and no inference from this reading to any of those is supported by the construction.* `F_int` *is invariant under relabelling of the states and is therefore many-to-one even on chords; it does not determine the chord, and the chord does not determine anything about a person's experience. The twelve states are not diagnoses. This is not an assessment and not clinical advice. Nothing here recommends or values suffering. The frequencies are essence frequencies, not physical resonances.*

A reading published without this paragraph is not a reading this framework endorses.

SPEC-4.4 — Gates

```
GATE 10-1  the audit reports BOTH state count and class count          FAIL -> incomplete
GATE 10-2  every pairwise u published, not the mean alone              FAIL -> incomplete
GATE 10-3  the separation statement present in full                    FAIL -> do not publish
GATE 10-4  no sentence in the reading contains a condition name, a
           severity, a norm comparison, or a recommendation            FAIL -> do not publish
GATE 10-5  restatement flags reported, not resolved by deletion        FAIL -> reading void
GATE 10-6  consonance audit and redundancy audit reported separately   FAIL -> incomplete
```

Gate 10-4 is a text gate rather than a numeric one and it is the one an analyst will be tempted to wave through. **It is mechanical on purpose: read the reading, look for those four things, and if any is present the reading does not go out.** A gate that requires judgment gets exercised; a gate that requires a word search gets applied.

The Failure Modes

FM-4.1 — Romanticising dissonance *The failure.* The section's result — that a perfectly consonant self has zero surplus — is read as "division is good," "conflict is generative," or "suffering is informative," and from there into advice.

Why it happens. Because the arithmetic produces a sentence that sounds like a moral and because the moral is attractive. The genre this framework sits adjacent to has produced that moral many times, usually without any arithmetic under it at all.

Why it is wrong, in three independent ways. (i) The formalism contains no term for suffering; there is no path by which it enters. (ii) `F_int` is invariant under relabelling and does not even determine the chord (§4.6), so it cannot determine anything downstream of the chord. (iii) The only lever on `F_int` is angular diversity of elicited framings, and nothing about angular diversity implies discomfort — three comfortable framings of a situation score exactly as three uncomfortable ones would.

Detection. Any sentence in which more `F_int` is better, or in which a difficult experience is described as producing surplus.

Repair. Delete the sentence. **There is no repaired version of it**, because the claim is not underdetermined by the formalism, it is absent from it.

FM-4.2 — Pathologising consonance *The failure.* The exact inverse, and the one that follows FM-4.1 the way an echo follows a shout: a low `F_int` reading is treated as a deficiency, a warning sign, or something to be corrected.

Why it happens. Because the internal-surplus analysis of Series B's description of the zero-surplus case — forty descriptions that are one description — reads like a portrait, and portraits invite recognition.

Why it is wrong. A settled, coherent, single reading of a situation is a completely ordinary thing and the formalism says nothing against it. `F_int = 0` means the elicited framings were parallel on that occasion. It means that and nothing else. **It is not a deficit, it is not a symptom, and describing it as either is a clinical claim this framework does not make and is not entitled to make.**

Detection. Any sentence recommending that a subject develop more framings, or describing a narrow chord as something to work on.

Repair. Report the number and stop. **The instrument's job ends at the number.**

FM-4.3 — Reading harmonic dissonance as psychological conflict *The failure.* The `7:10` pair at 34.98 cents is described as the subject's inner conflict, or the `5:10` collision as the subject's confusion.

Why it happens. The words carry their ordinary meanings in.

Why it is wrong. Those numbers are computed from `7`, `10`, and `5` — from integers, by a defined function, with no reference to any person. **Relabel the states and the numbers change while the subject's framings do not.** §4.6 demonstrates precisely this.

Detection. Any sentence attributing an interval to a subject's experience.

Repair. State intervals as properties of the index set. "The chord contains the lattice's unique 50/49 pair" is correct. "The subject is in a 50/49 relationship with themselves" is not a sentence this framework can produce.

FM-4.4 — The reading used as clinical advice *The failure.* A reading is used to decide something about a person's care, treatment, capacity, or condition.

Why it happens. The output is numeric, the output is reproducible, and numeric reproducible outputs get used for decisions.

Why it is wrong. **No validation of this instrument against any clinical outcome exists, has been attempted, or is claimed. The twelve-state ladder is an ordinal construction defended in the minimum-frequency paper [Curzi, 2026] on its own terms and demoted on scale by CC-R1. It has no criterion validity against anything.**

Detection. Gate 10-4, plus the question of who is going to read the reading and what they will do next.

Repair. The separation statement of SPEC-4.3, in full, and the analyst's own refusal.

FM-4.5 — Manufactured internal difference *The failure.* An analyst elicits additional framings by pressing a subject to produce a different answer, and the subject obliges.

Why it happens. Because the analyst wants a valid reading and near-parallel framings feel like a failed session.

Why it is wrong. A framing produced to satisfy an analyst is not a framing the subject holds, and the encoding cannot tell the difference. This is internal pseudo-diversity, and Series B's pseudo-diversity analysis treats the relational version at length: apparent variety with no angular content behaves exactly like manufactured variety with angular content, which is why the failure is invisible in the output.

Detection. Procedural. Record the elicitation, record how many prompts each framing required, and publish the count.

Repair. Elicit once, accept the result, and report a narrow chord as a narrow chord.

FM-4.6 — The total-order fallacy on harmonic distance *The failure.* Reasoning from "consonance is redundancy" as though every collision were more consonant than every non-collision.

Why it happens. Because the slogan is memorable and the caveat is not.

Why it is wrong. Result (C): `1:3` at `HD_ess = 3.1699` is more consonant than the collisions `1:4`, `2:8`, `3:12` at 4.0000 and `1:8` at 6.0000. **The claim is exactly true at the minimum, exactly true as a biconditional about zero folded separation, and true on average (3.0000 against 8.6617). It is not a total order.**

Detection. Any argument of the form "this pair is a collision, therefore it is more consonant than that pair."

Repair. Use the computed values. There are only sixty-six of them.

FM-4.7 — Treating the fold as an argument *The failure.* Concluding, from the fact that states 5 and 10 fold to the same pitch, that attributed feeling and narrative feeling are substantively the same thing.

Why it happens. The fold makes them audibly identical and audible identity is persuasive.

Why it is wrong. The fold is a quotient by powers of two applied for the sole purpose of bringing values into a chosen band. It is a change of units. The states differ by five components of specification on the ladder, and the fold does not know that and does not care.

Detection. Any substantive claim justified by a folded value.

Repair. Publish raw and folded together, always, with the band and `j`, per SPEC-9.3.

FM-4.8 — Optimising `F_int` *The failure.* Treating the width as a score to be raised, whether by the analyst, the subject, or a system built around the reading.

Why it happens. Because it is a number, and numbers attract optimisation.

Why it is wrong. Every honest route to raising it runs through the subject genuinely holding differently-shaped construals, which is not something an analyst can supply. Every dishonest route — restating, manufacturing, disjoint encoding — either lowers it, degenerates it, or falsifies it. There is no available optimisation that is also a measurement.

Detection. Any repeated reading of the same subject where the number rises and the procedure changed.

Repair. Pre-register the procedure, per Gate 6 of the octave-folding paper [Curzi, 2026], and re-read without altering it.

FM-4.9 — Importing the wrong construct *The failure.* Citing cognitive-dissonance research, or consonance research in music perception, as support for anything in this section.

Why it happens. Shared vocabulary.

Why it is wrong. Tenney distance and Euler gradus are defined functions, badged as chosen, and are not claims about hearing. The internal surplus is a construction over an ordinal encoding, not a psychological construct with a literature. Citing unrelated empirical work as support would be a fabricated citation, and the framework does not do it.

Repair. Cite Tenney [Tenney, 2015] and Euler [Euler, 1739] for the measures, which is what they are for — and cite the essay rather than the 1988 *History*, per §4.2 — and cite nothing for the psychology, because there is nothing to cite.

FM-4.10 — Treating a zero-surplus reading as the end of the inquiry *The failure.* A reading returns `F_int = 0.000000`, the analyst records "no surplus," and the session ends.

Why it happens. Zero looks like a complete answer. It is a number, it is definite, and it closes.

Why it is wrong. **Zero surplus and no measurement are different results and the specification requires them to be distinguished.** `F_int = 0` means at least two framings were elicited and they were parallel. A session that produced only one framing has no pairs, no `u_int`, and no reading at all — the octave-folding paper [Curzi, 2026]'s Gate 7. Recording the second as the first converts a missing measurement into a null one, which is the single most common way a measurement programme accumulates false negatives.

There is a second reason the inquiry should not end there. **A zero-surplus reading is the most informative reading the redundancy audit produces**, because it is the one that identifies the elicitation returned the same construal repeatedly — which is a fact about the session, and a fact worth having, and one that the number reports accurately.

Detection. Compare the framing count against the pair count. Zero pairs is no reading; one or more pairs with `u_int = 0` is a null reading.

Repair. Record which of the two occurred, per SPEC-4.2 Step 3, and record how many framings the elicitation produced before it stopped.

FM-4.11 — Comparing readings across ladders *The failure.* Two readings computed on different state ladders, or on the same ladder with a different number of categories, are compared as though the numbers were commensurable.

Why it happens. `F_int` is in nats either way and nats look like nats.

Why it is wrong. `F_int = ln(1 + (N−1)u)` contains `N` explicitly, and its ceiling is `ln N`. A surplus of 2.164390 nats against a ceiling of `ln 12 = 2.484907` is 87.10% of the available range; the same 2.164390 nats against a ceiling of `ln 25` would be 67.24%. The raw value has no meaning without its ceiling, and neither does the pitch: `f_width` depends on `N_eff`, which depends on `N`.

Detection. Any comparison in which the two readings do not state the same `N`.

Repair. Publish `F_int`, `ln N`, and the ratio, always together — and where `N` differs, compare the ratio and say plainly that the absolute values are not comparable.

Worked Example

THE SUBJECT — the consonance and redundancy audits **CONSTRUCTED CASE, the framework's carried subject, carried from the octave-folding paper [Curzi, 2026]. Not an observation of any person.**

The Subject's chord, from SPEC-9.2 Step 8:

```
STATES SPANNED   {3, 5, 7, 10}
state  3   directed feeling      raw  4.3508 Hz   folded 278.453 Hz (j=+6)   class 3
state  5   attributed feeling    raw 12.0856 Hz   folded 386.741 Hz (j=+5)   class 5
state  7   reflexive feeling     raw 23.6879 Hz   folded 379.006 Hz (j=+4)   class 7
state 10   narrative feeling     raw 48.3426 Hz   folded 386.741 Hz (j=+3)   class 5
```

The consonance audit (SPEC-4.1)

```
pair    essence ratio   raw cents   folded cents   HD_ess    G_ess   collision
3:5        25 / 9        1768.72       568.7174    7.8138      13       no
3:7        49 / 9        2933.74       533.7418    8.7846      17       no
3:10      100 / 9        4168.72       568.7174    9.8138      15       no
5:7        49 / 25       1165.02      1165.0244   10.2586      21       no
5:10        4 / 1        2400.00         0.0000    2.0000       3      YES
7:10      100 / 49       1234.98        34.9756   12.2586      23       no

MIN HD_ess       =  2.0000    (pair 5:10)
MEAN HD_ess      =  8.4882
COLLISION COUNT  =  1
STATE COUNT      =  4
CLASS COUNT      =  3         classes {3, 5, 7}
50/49 PAIR       =  PRESENT   (7:10 at 34.9756 cents -- the lattice's unique instance)
```

Three observations, each stated at the altitude the arithmetic supports.

The chord contains the most consonant interval the lattice permits and it contains it exactly once. `5:10` sits at `HD_ess = 2.0000`, the global minimum over all sixty-six pairs, attained only by the six 2:1 pairs. And by result (B), that is precisely the chord's one redundancy: 0.0000 folded cents, one class for two states. Consonance and redundancy are not correlated here; they are the same entry in the table, read twice.

The chord also contains the lattice's most remote pair among its members, and the two are adjacent in the index set. `7:10` at `HD_ess = 12.2586` is the widest interval in this chord, and it folds to 34.9756 cents — the smallest nonzero separation the whole `N = 12` system permits. **The most harmonically remote pair in the chord is the pair that would sound closest if the**

chord were folded into one band. That is not a paradox; it is what the fold does. Harmonic distance is computed on the unfolded ratio `100/49`; folded separation is computed modulo the octave. They are different questions and this chord happens to answer them oppositely.

Mean `HD_ess = 8.4882` sits just below the non-collision average of 8.6617 for the lattice as a whole, which is what one expects of a chord containing exactly one collision among six pairs.

The redundancy audit (SPEC-4.2)

```
pairwise u          mirror / story    0.809410
                    mirror / threat   0.904762
                    story  / threat   0.388360

RESTATEMENT FLAGS   none  (minimum pairwise u = 0.388360, well above the 0.01 threshold)

u_int               0.700844
m (nominal)         3 framings
N_eff (audited)     8.7093 categories, under the roster reduction postulate
F_int               2.164390 nats
f_width             36.6686 Hz   folded 293.349 Hz (band 220–440, j = +3)
octaves             6.2451
```

The two audits disagree about what the chord is like, and they are entitled to. The consonance audit finds one exact redundancy in the chord's index set. The redundancy audit finds no restatement among the framings at all — the closest pair, story against threat, sits at `u = 0.388360`. The collision is in the ladder. It is not in the framings.

This is §4.6 made concrete on the carried case. The Subject holds three framings that are substantially non-parallel, and their union happens to include two indices differing by a factor of two. **Relabel state 10 as state 9 and the collision vanishes with `F_int` unchanged to twelve decimal places.** Nothing about the Subject changed; the ladder's arithmetic changed.

The counterfactual: what a consonant Subject would read

EXTENSION 3.1-d (LABELLED): the Subject with the mirror framing replaced by a restatement of the threat framing. This is the same person only in the sense that the case is carried; it is a constructed variant and is labelled as one.

```
                        SUBJECT          EXTENSION 3.1-d      change
u(threat, story)        0.388360            0.388360           --
u(threat, mirror)       0.904762            0.001034        -0.903728
u(story,  mirror)       0.809410            0.372856        -0.436554

u_int                   0.700844            0.254083        -0.446761
N_eff                   8.7093              3.7949          -4.9144
F_int (nats)            2.164390            1.333662        -0.830728
f_width (Hz)           36.6686             6.9620          -29.7066
octaves above ground    6.2451              3.8481          -2.3970

RESTATEMENT FLAGS       none                (threat, mirror') at u = 0.001034
```

Two-point-four octaves, lost to one substitution. And the restatement flag fires, so the reading goes out with the flag attached rather than with the framing quietly removed — Gate 10-5.

And the terminal case, for completeness:

```
EXTENSION 3.1-b (LABELLED): all three framings parallel.
    u_int = 0.000000   N_eff = 1.0000   F_int = 0.000000   f_width = k = 0.483426 Hz
    octaves above ground = 0.0000
```

Zero. The ground essence. Three framings and one window.

The reading, complete

```
========================================================================
CONSONANCE / REDUNDANCY READING -- THE SUBJECT
(CONSTRUCTED CASE, the framework's carried subject. Not an observation of any person.)
========================================================================
ladder            N = 12, CCP-041 ordinal on n; intensities ratio-scaled
                  up to a free multiplier per CCP-041 (CC-R1 governs D/P, not used here)
chord             states {3, 5, 7, 10}; classes {3, 5, 7}
                  four states, three classes: 10 = 2*5, same odd part

CONSONANCE AUDIT  min HD_ess    2.0000   (5:10, the lattice global minimum)
                  mean HD_ess   8.4882
                  collisions    1        (5:10, folded 0.0000 cents)
                  50/49 pair    present  (7:10, 34.9756 cents, unique at N=12)

REDUNDANCY AUDIT  pairwise u    0.809410 / 0.904762 / 0.388360
                  restatements  none
                  u_int         0.700844
                  m / N_eff     3 framings / 8.7093 audited categories
                  F_int         2.164390 nats
                  f_width       36.6686 Hz, folded 293.349 Hz (220-440, j=+3)
                  octaves       6.2451

INDEPENDENCE      relabelling 10 -> 9 removes the collision and leaves
                  F_int = 2.164390 unchanged to 1e-12. The consonance audit
                  and the redundancy audit measure different things.

This reading describes the information content of a set of framings under a
stated encoding. It does not measure distress, well-being, health, or
function, and no inference from this reading to any of those is supported by
the construction. F_int is invariant under relabelling of the states and is
therefore many-to-one even on chords; it does not determine the chord, and
the chord does not determine anything about a person's experience. The twelve
states are not diagnoses. This is not an assessment and not clinical advice.
Nothing here recommends or values suffering. The frequencies are essence
frequencies, not physical resonances.
========================================================================
```

What the section established, and the line it will not cross

Established, and checkable: the essence lattice doubles Tenney harmonic distance; the minimum essence harmonic distance is attained by exactly the six 2:1 pairs, all of which are pitch-class collisions; zero folded separation is exactly equivalent to sharing an odd part; therefore the most consonant pairs the lattice offers are precisely the pairs it cannot distinguish. Applied within one observer, a set of parallel framings has `u_int = 0`, `N_eff = 1`, `F_int = ln 1 = 0`, and a chord width of exactly the ground essence.

Not established, and not claimed anywhere: that a narrow chord indicates anything about a person's health, function, or experience; that a wide chord indicates anything better; that difficulty produces surplus; that suffering is a resource; that any of this is a basis for advice about anyone's care.

The result is about information content. The distance between "information content" and "how a person is doing" is not a gap this framework knows how to cross, and an account that pretended to cross it would be worth less than one that says where the road ends.

Section 4 owns the consonance/redundancy identity and the separation statement. Later papers in this series cite it. The doubling theorem is the harmonic-lattice analysis of Series B [Curzi, 2026l]; the internal surplus is the internal-surplus analysis of Series B; the chord, the one-hot trap, and SPEC-9.2 are the octave-folding paper [Curzi, 2026]; octave folding and the odd-part rule are The octave-folding paper [Curzi, 2026]. Section 5 assembles all of it into a worksheet a reader can run.

5 The personal spectrum

The preceding analysis built the parts. The minimum-frequency paper [Curzi, 2026] defended the index as a count and built the twelve-state ladder and the elicitation procedure that fills it in. The octave-folding paper [Curzi, 2026] gave a state a pitch, established the fold and its odd-part class rule, separated a trait from a mood by way of `delta`, and made the chord and named the one-hot trap. Section 4 established that consonance and redundancy are the same phenomenon, and spent most of its length making sure that result could not be turned into advice.

This section assembles all of it into one procedure that a reader can actually run, on paper, for one person, on one situation, and finish.

It ends where the minimum-frequency paper [Curzi, 2026] said every practical section of this framework must be able to end: with an octave deficit.

```
THE READING THIS SECTION PRODUCES

    octave deficit  =  2.885 * (c/eta)  -  log2( f / k )

    deficit <= 0    sustainable as it stands
    deficit  > 0    short by that many octaves, and the four levers are:
                    raise u  ·  raise eta  ·  cut c  ·  add distinct categories
```

Section 5 owns the full personal worksheet. The companion paper on refinement cites it. Everything else in the procedure is owned elsewhere and cited here rather than restated, which is the whole point of a worksheet: it is an index to the framework with blanks in it.

> **BEFORE ANYTHING ELSE. The twelve states are not diagnoses and this reading is not an assessment.** It makes no clinical, medical, or therapeutic claim of any kind. **The frequencies are essence frequencies, not physical resonances** — nothing here asserts that any person, tissue, or molecule resonates at any value. **The reading measures the information content of a set of framings under a stated encoding; it does not measure distress, well-being, health, or function, and Section 4 §4.6 shows formally why no inference to any of those is available.** All cases in this section are CONSTRUCTED and labelled as such.

The Law

5.1 What a personal reading is A personal reading is a reading **of a set of framings, on one situation, on one occasion, under one stated decomposition, by one analyst**. Every one of those qualifications is load-bearing and every one of them appears on the worksheet as a blank that must be filled before the arithmetic starts.

What it is not is a reading of a person. **There is no procedure in this framework that reads a person.** The FVP-7 gate for the self-storage law of Series B is recorded as **NOT RUN** — the reproducibility gate, whether one person's `eta` means what another's does, has not been attempted, and until the separation property of a specific-objectivity instrument is established, **every personal parameter is a within-subject index and cross-person comparison is not licensed.** That constraint is inherited here in full. Two people's readings may be compared only in the sense that two people's shoe sizes measured with different rulers may be compared: not at all.

There is a reason this section takes the form of a worksheet rather than an argument, and it is worth stating before the sheet appears.

An instrument with a contested foundation is honest only if it is fully procedural. The principle is old and is an engineering standard rather than an ethical preference: a system must remain sound when everything about it except the key is public [Kerckhoffs, 1883], and disclosure is the discipline that makes a contested instrument checkable [Brandeis, 1914]. Every component of this reading rests on something that could be disputed: the ladder is an ordinal construction defended in the minimum-frequency paper [Curzi, 2026] rather than derived; the constant `k` is a design decision badged as a postulate; the intensities are ratio-scaled only to the extent a magnitude-estimation procedure delivers ratio scales, which the minimum-frequency paper [Curzi, 2026] states is a measured property of a procedure and not a constant of nature; the roster reduction is a postulate; the storage parameters have a reproducibility gate recorded as NOT RUN. **Not one of those is hidden, and not one of them is fatal — but together they mean the number has no authority of its own.** What it has instead is a trail.

A worksheet is what a trail looks like when it is written down. **Every blank on the sheet is a place where someone made a decision, and the sheet exists so that a second person can find each decision, disagree with it, change it, and see what the reading does.** That is the only form of validity this instrument currently has: not that its numbers are right, but that its numbers are *recomputable from declared inputs by someone who thinks the inputs are wrong.*

The minimum-frequency paper [Curzi, 2026]'s honest limit governs the whole sheet: the frequency inherits exactly the defensibility of the count. A contested decomposition yields a contested frequency. The worksheet does not resolve that. It makes the contest possible.

5.2 The five outputs

```
THE PERSONAL SPECTRUM -- five outputs, in order of derivation

  1  THE SPECTRUM      the twelve state frequencies, raw and folded
                       f(n) = k n^2                                    [CCP-042]

  2  THE CHORD         the states spanned by the framings, their pitches,
                       their pitch classes, and their intervals        [CCP-042]

  3  THE WIDTH         F_int = ln(1 + (N-1) u_int),  f = k e^(2 F_int) [CCP-042]

  4  THE COST RATIO    c / eta, estimated at internal scope            [Series B, self-as-storage]

  5  THE DEFICIT       2.885*(c/eta) - log2(f/k)                       [CCP-041]
```

Outputs 1 and 2 depend on the ladder and the framings. Output 3 depends only on the angles between framings. Output 4 depends on none of that — it is a separate estimation problem with separate hazards. **Output 5 is the only place the two halves meet, and it is the join this whole framework exists to perform.**

5.3 The octave deficit, restated The minimum-frequency paper [Curzi, 2026] owns the minimum sustainable frequency and its diagnostic. Restated here in the form the worksheet uses, without re-derivation:

```
THE MINIMUM SUSTAINABLE FREQUENCY                    [CCP-041, the framework's specification]

    Q_inf > 0   <=>   eta F > c   <=>   F > c/eta
    F = (1/2) ln(f/k)                                 the inverted bridge
    f_min = k * e^( 2 c / eta )

    octaves required = 2 (c/eta) / ln 2 = 2.885 * (c/eta)
    octaves present  = log2( f / k ) = 2.885 * F
    octave deficit   = required - present
```

Every unit of cost-to-conversion ratio costs 2.885 octaves of required pitch. The verified table from the minimum-frequency paper [Curzi, 2026], reproduced because the worksheet needs it at hand:

c/eta	f_min (Hz)	octaves above k
0.000	0.4834	0.000
0.500	1.3141	1.443
1.000	3.5721	2.885
1.500	9.7099	4.328
2.000	26.3942	5.771
2.485	69.6124	7.170
3.000	195.0279	8.656

The constant is written `2.885` throughout; exactly it is `2/ln 2 = 2.885390`. Where a reading sits within a hundredth of an octave of zero, use the exact form.

5.4 What `eta`, `c` and `delta` mean inside one person

The storage law at internal scope is the self-as-storage analysis of Series B's self-storage law, and this section cites it rather than re-deriving it:

```
Q_{t+1} = (1 - delta) Q_t + eta M_t - C_t              [Bk 2 the self-as-storage analysis of Series B]

Q      stored durable structure -- character, in Q-units
delta  retention loss per period, in [0,1], dimensionless
eta    conversion of experience into durable structure, in Q-units per nat
C      cost per period, in Q-units
M      the moment, in nats -- here supplied by F_int

Q_inf = ( eta F - c ) / delta                          the storage law
```

Note the units, because they are what stops the equation being a slogan: `eta` carries Q-units per nat, so `c/eta` carries nats and `2.885*(c/eta)` carries octaves. The deficit is dimensionally sound and the worksheet checks it.

At internal scope the three parameters read as follows, and each has an estimation method owned by The self-as-storage analysis of Series B:

`eta` — conversion. How much of what happens becomes durable structure. The self-as-storage analysis of Series B estimates it as a product of six pass-through yields, `eta = y1 * y2 * ... * y6`, with the last rung carrying the units. Its structural property is multilinearity: the elasticity of `eta` with respect to every rung is exactly 1, so the absolute gain from fixing a rung is inversely proportional to that rung's current yield. A single blocked rung nullifies every other, which is why the aggregate figure is uninformative and the stage-wise figure is not.

`c` — cost. The self-as-storage analysis of Series B estimates it from five ledgers. It is a per-period quantity in Q-units and it must be declared with its period.

`delta` — loss. Estimated from a zero-input window. It sets the time constant `tau = 1/delta` and therefore how long a regime takes to arrive at its fixed point.

And the identifiability hazard, which is the one that voids readings: the storage-law admissibility condition of Series B — **if the Q-proxy rises whenever a lot happens, the input has been measured twice and the gauge is void.** A journal that gets longer on busy weeks is not a measure of `Q`. The worksheet requires the proxy and its independence check to be declared before any parameter is estimated.

5.5 The break-even identity

Substituting `octaves present = 2.885 * F_int` into the deficit gives an identity worth stating on its own, because it is what makes the reading legible without a table:

```
THE BREAK-EVEN IDENTITY

    deficit = 2.885*(c/eta) - 2.885*F_int = 2.885 * ( c/eta - F_int )

    deficit = 0   <=>   c/eta = F_int
    deficit < 0   <=>   c/eta < F_int      SUSTAINS
    deficit > 0   <=>   c/eta > F_int      SHORT
```

The whole diagnostic is one comparison: is the cost ratio smaller than the surplus? The octave form is the same statement rescaled by `2.885`, and it exists because octaves are legible and nats are not.

For the carried Subject, with `F_int = 2.164390`:

```
break-even c/eta = 2.164390 nats
below it   -> sustains        above it   -> short
```

Everything the frequency machinery adds to `eta F > c` is legibility. It adds no content, and the framework says so here rather than letting the reader discover it and feel cheated. What the frequency form buys is that a deficit in octaves can be compared against a table of what an octave costs, and a ratio in nats cannot.

5.6 The stop rule

Because `F <= ln N` always, the deficit cannot be closed by raising `u` past orthogonality. That gives a hard stop, and it is the surplus ceiling in the frequency register:

```
THE CEILING                                          [CCP-041, the framework's specification]

    c/eta  <=  ln N  is REQUIRED for any arrangement of N categories to sustain.

    At N = 12:   c/eta <= 2.484907   and   f_min <= k*144 = 69.6133 Hz

    c/eta = 2.4849 -> f_min =  69.6124 Hz    at the ceiling
    c/eta = 2.6000 -> f_min =  87.6317 Hz    ABOVE -- nothing sustains
    c/eta = 3.2000 -> f_min = 290.9474 Hz    ABOVE -- nothing sustains
```

Above the ceiling, no arrangement of those twelve categories is sustainable — not a better set of framings, not a better schedule, not more effort. The only remedies are to raise `eta`, cut `c`, or add genuinely distinct categories, and the worksheet stops and says so rather than computing a deficit that no rearrangement could close.

The stop rule is checked before the deficit, not after. A worksheet that computes a deficit above the ceiling has produced a number that describes an impossibility.

5.7 `delta` is not a lever on the deficit One structural fact the worksheet exposes and that surprises most readers on first sight:

```
deficit = 2.885*(c/eta) - log2(f/k)          -- no delta term
Q_inf   = (eta F - c) / delta                -- delta scales the magnitude
```

`delta` does not appear in the deficit at all. It scales `Q_inf` and sets the time constant, but it cannot change the sign. Computed on the carried Subject at `eta = 0.327600`, `c = 0.55`, where the numerator `eta F_int - c = +0.159054`:

```
delta = 0.20  ->  Q_inf = +0.7953     deficit −1.4015
delta = 0.12  ->  Q_inf = +1.3255     deficit −1.4015
delta = 0.06  ->  Q_inf = +2.6509     deficit −1.4015
delta = 0.03  ->  Q_inf = +5.3018     deficit −1.4015
```

Lowering `delta` multiplies whatever sign the numerator already has. On a positive numerator it multiplies a gain. On a negative one it multiplies a loss. That is CC-R15, it is owned by the relationship-frequency paper [Curzi, 2026], and its internal-scope demonstration is in this section's Worked Example because the worksheet has to be able to warn about it.

The Engineering

5.8 Step A — eliciting the intensities The minimum-frequency paper [Curzi, 2026] owns this procedure and this section cites it. What the worksheet needs from it is three things, recorded before any arithmetic:

```
A1  THE SITUATION, stated once, in the subject's own words, and fixed.
    All framings are framings OF THIS. If the situation drifts between
    framings, the reading is void -- the angles would be between
    vectors pointed at different targets.

A2  THE FRAMINGS, elicited one at a time, each described in full before
    the next is requested. Record how many prompts each required.

A3  THE DECOMPOSITION. For each framing, of each of the twelve state
    classes, ask whether the description REQUIRES that component to be
    specified. Rate the required ones for intensity on a common scale.
```

Three procedural rules that the worksheet enforces:

Do not ask which state a framing "is." That question produces one-hot vectors and the octave-folding paper [Curzi, 2026] shows what one-hot vectors do. Ask which components the description requires. Framings of one situation share components; that is what makes them framings of one situation.

Record the prompt count. A framing that took six prompts to obtain and a framing the subject volunteered are different data, and only the record distinguishes them. Section 4's FM-4.5 — manufactured internal difference — is invisible in the output and visible only here.

Declare the scale status, and declare the right one. Two different scale claims are in play and conflating them is a real error. CC-R1 demotes the Moment equation: `D/P` may be used ordinally only, and any psychological frequency derived through `D/P` inherits the demotion. The worksheet does not route through `D/P`, so that demotion does not bind the reading — but the box exists so that an analyst who *does* route through it must say so.

The intensities have their own scale status and the minimum-frequency paper [Curzi, 2026] owns it. The minimum-frequency paper [Curzi, 2026] shows that ordinal-preserving transformations of the intensities move `u` on the carried Subject's closest pair from 0.362115 to 0.525275 — `f_rel` from 12.0049 to 22.2093 Hz, a spread of 1065.05 cents, nearly nine tenths of an octave — from transformations that preserve every ordinal fact in the data. So merely ordinal intensities will not do. They must be ratio-scaled up to a common free multiplier, which is what a magnitude-estimation procedure is built to deliver, and the honest form of the claim is that they are ratio-scaled to the extent magnitude estimation delivers ratio scales — a measured property of a procedure, not a constant of nature. Say which claim the reading rests on, in the box provided, every time.

And report the norms. The minimum-frequency paper [Curzi, 2026]'s **MAGNITUDE REPORTING RULE** requires the raw L2 norm of every state vector to be published alongside the normalised vector, in the units of the elicitation instrument, and never discarded. `u` measures direction; normalisation throws away length, and length is not nothing — two profiles differing eightfold in intensity are exactly parallel and carry `u = 0.000000` between them.

5.9 Step B — building the vectors and passing the gates

The octave-folding paper [Curzi, 2026]'s SPEC-9.2 owns this. In worksheet form:

```
B1  Write each framing as a twelve-vector of intensities.
B2  FIX AND RECORD THE SUPPORTS before computing anything.
B3  Normalise: x <- x / ||x||.  Check ||x|| = 1.000000.
B4  OVERLAP GATE: every pair must share at least one state.
    Any pair with empty shared support -> STOP. Report DEGENERATE.
B5  WEIGHT GATE: shared support carries >= 10% of each vector's squared
    norm. (A stated convention, not a law. State it; apply it blind.)
```

The overlap gate is the single most important line on the worksheet. The octave-folding paper [Curzi, 2026] computed what happens without it: three one-hot framings score `u = 1.000000` on every pair, `F_int = ln 12 = 2.484907`, `f = 69.6133 Hz` — the ceiling, identically, for every subject who has ever lived. And the disguised version is worse, because three framings each spanning three states with no state in common produce exactly the same degenerate maximum while looking rich.

A worksheet that reaches `F_int = 2.484907` at `N = 12` should be treated as broken until proven otherwise.

5.10 Step C — the spectrum and the fold

The octave-folding paper [Curzi, 2026] owns the fold and the odd-part rule. The worksheet carries the table:

```
THE TWELVE-STATE SPECTRUM      k = 0.483425732988414 Hz,  f = k n^2

 n   state class                   raw (Hz)    folded 220-440    j    class
 1   undifferentiated activation      0.4834       247.514       +9      1
 2   valenced activation              1.9337       247.514       +7      1
 3   directed feeling                 4.3508       278.453       +6      3
 4   situated feeling                 7.7348       247.514       +5      1
 5   attributed feeling              12.0856       386.741       +5      5
 6   temporal feeling                17.4033       278.453       +4      3
 7   reflexive feeling               23.6879       379.006       +4      7
 8   social feeling                  30.9392       247.514       +3      1
 9   normative feeling               39.1575       313.260       +3      9
10   narrative feeling               48.3426       386.741       +3      5
11   reciprocal feeling              58.4945       233.978       +2     11
12   committed feeling               69.6133       278.453       +2      3
```

```
Six pitch classes {1,3,5,7,9,11} = ceil(12/2).  ln 12 = 2.484907.
Band 220–440 Hz; j is the transposition. STATE THE BAND AND j ALWAYS.
```

The pitch class is the odd part of the index and nothing else. Odd indices are their own odd part and are always their own class; even indices collapse onto the class of the odd number left after dividing out every factor of two. Four of the twelve states — 1, 2, 4, 8 — share class 1; three — 3, 6, 12 — share class 3; two — 5, 10 — share class 5; and 7, 9, 11 are alone in theirs. That is the entire redundancy structure of the vocabulary, and an analyst who has it memorised can read the class column off a chord without computing anything.

5.11 Step D — the chord and its intervals

The octave-folding paper [Curzi, 2026]'s SPEC-9.2 Steps 8 and 9, and Section 4's SPEC-4.1:

```
D1  STATES SPANNED = union of the framing supports.  Report the count.
D2  PITCH CLASSES  = distinct odd parts.  Report the count SEPARATELY.
D3  RAW pitches k n^2; FOLDED pitches with band and j.
D4  INTERVALS: for each pair (i,j), essence ratio j^2/i^2 in lowest terms,
    raw cents, folded cents.
D5  CONSONANCE AUDIT: min HD_ess, mean HD_ess, collision count.
    Collision <=> same odd part <=> folded cents exactly 0.
D6  Flag the 50/49 pair if present. At N = 12 there is exactly one: 7:10.
```

D1 and D2 are different counts and conflating them is FM-9.7. A chord of four states in three classes has one member that adds pitch without adding a class.

5.12 Step E — the width

```
E1  every pairwise u = 1 - (x.y)^2, PUBLISHED INDIVIDUALLY
E2  u_int = mean over all m(m-1)/2 pairs
E3  N_eff = 1 + (N-1) u_int
E4  F_int = ln N_eff
E5  f = k * N_eff^2 = k * e^(2 F_int);  folded, with band and j
E6  octaves present = 2 F_int / ln 2
E7  F_int / ln N, as a percentage of the ceiling
```

E7 is not decoration. **A raw `F_int` is meaningless without its ceiling** — 2.164390 nats is 87.10% of `ln 12` and would be 67.24% of `ln 25`. Section 4's FM-4.11 covers the comparison hazard; the worksheet forecloses it by requiring the ratio.

5.13 Step F — estimating `eta`, `c`, `delta`

This is the step that will take longest and the step most likely to be skipped, and the worksheet is built so that skipping it produces a blank rather than a default.

```
F1  DECLARE THE PERIOD.  Every one of eta, c, delta is per-period.
F2  DECLARE THE Q-PROXY and run the independence check.
    HAZARD (storage-law admissibility): if the proxy rises whenever a lot happens,
    the input has been measured twice and the gauge is VOID.
F3  eta -- the six-rung ladder [Series B, self-as-storage]. Enumerate the candidate
    set PROSPECTIVELY, at the point of entry. Report every rung, not
    the product alone.
F4  c   -- the five ledgers [Bk 2 the self-as-storage analysis of Series B.5], in Q-units per period.
F5  delta -- from a zero-input window [Bk 2 the self-as-storage analysis of Series B.3.4].
F6  DECLARE: these are WITHIN-SUBJECT indices. FVP-7 is NOT RUN.
    No cross-person comparison is licensed.
```

On F3, the reason the rungs must be reported individually: `eta = y1 * ... * y6` is multilinear, so a single blocked rung nullifies every other. An aggregate `eta = 0.35` is consistent with six mediocre rungs and with five excellent rungs and one broken one, and the intervention implied by those two situations is completely different. The elasticity of `eta` with respect to every rung is exactly 1, so the absolute gain from repairing a rung is inversely proportional to its current yield: **the worst rung is always the cheapest to fix and always the one the aggregate hides.**

On F3's identifiability hazard, which is the one that quietly ruins estimates: an analyst who enumerates only the experiences they already converted will recover `eta` near 1. **The candidate set must be enumerated prospectively, at the point of entry, before anyone knows what converted.**

5.14 Step G — the deficit, and the four levers

```
G1  c/eta, in nats
G2  STOP RULE: if c/eta > ln N, HALT. Report ABOVE CEILING.
    No arrangement of N categories sustains. Do not compute a deficit.
G3  f_min = k * e^(2 c/eta)
G4  octaves required = 2.885 * (c/eta)
G5  DEFICIT = required - present
G6  Q_inf = (eta * F_int - c) / delta
G7  time constant tau = 1/delta; periods to within 5% of Q_inf
G8  If deficit > 0, the four levers, EACH COSTED:
        raise u      (raise F_int)   ceiling: ln N
        raise eta    (the conversion ladder, worst rung first)
        cut c        (the five ledgers)
        raise N      (genuinely distinct categories -- not relabelled ones)
```

Lever four carries a standing warning. Raising `N` raises the ceiling and, at fixed `u`, raises `F` — but only if the added categories are genuinely distinct. Series B's pseudo-diversity analysis treats pseudo-diversity at length: **categories that are relabellings of existing ones raise `N` in the formula and not in the world, and the result is a reading that improves while nothing does.**

5.14b Double-coding, and what to do when two analysts disagree The worksheet has a line at the foot for a second coder, and it is not a formality.

The decomposition step — deciding which of the twelve components a description requires — is the step where two competent analysts most reliably diverge. One reads "this is the latest instalment of an old pattern" as requiring narrative feeling and attributed feeling; another reads it as requiring narrative feeling, attributed feeling, and temporal feeling, on the grounds that "latest instalment" projects a duration. Both readings are defensible. They produce different supports, different inner products, and different frequencies.

The procedure when they diverge:

```
DOUBLE-CODING DISAGREEMENT PROCEDURE

1  Both analysts publish their full vectors and their full readings.
2  Neither reading is designated correct.
3  DO NOT AVERAGE THE DECOMPOSITIONS.  Averaging two decompositions
   produces a third that neither analyst endorses, and the resulting
   vector corresponds to no reading of the description.
4  Report the SPREAD: the range of F_int, the range of the deficit,
   and whether the two readings agree on the SIGN of the deficit.
5  If the sign agrees, the disagreement is about magnitude and the
   reading is usable with the spread attached.
```

```
    If the sign disagrees, the reading is INCONCLUSIVE and must be
    reported as inconclusive.
```

Step 5 is the whole value of double-coding. A deficit of −1.40 from one analyst and −0.62 from another is a disagreement about how far inside the boundary the reading sits, and the reading still says *inside*. A deficit of −0.20 from one and +0.30 from another says nothing at all, and an instrument that reports "nothing at all" when it has nothing is worth more than one that reports the first analyst's number.

The framework has no procedure for making the disagreement go away, and it should not have one. **The disagreement is information about how well specified the ladder is, and suppressing it would destroy the only evidence available on that question.**

5.15 CC-R15 at internal scope — the order of operations

The last engineering point, and it is the one that turns a reading into a decision.

When the deficit is positive, the intuitive repair is to lower `delta` — to hold on to things better, to let less slip, to be steadier. At internal scope this is the most natural move anyone makes and it is the one that goes wrong.

```
CC-R15, INTERNAL SCOPE                                  [owned by CCP-044]

    Q_inf = (eta F - c) / delta

    If (eta F - c) < 0, lowering delta MULTIPLIES a negative number.
    The fixed point gets WORSE, and it gets worse in proportion.
```

Computed on the carried Subject under a failing parameter set (Worked Example below):

```
as found              eta=0.25 c=0.60 delta=0.12   Q_inf = -0.4909   DECAYS
lower delta only      eta=0.25 c=0.60 delta=0.06   Q_inf = -0.9817   WORSE
raise eta first       eta=0.32 c=0.60 delta=0.12   Q_inf = +0.7717   ACCUMULATES
then lower delta      eta=0.32 c=0.60 delta=0.06   Q_inf = +1.5434   BETTER
```

> **FIX THE NUMERATOR, THEN LOWER THE LOSS. REVERSE THE ORDER AND THE SAME INTERVENTION ACCELERATES THE FAILURE.**

The worksheet places the `Q_inf` sign check immediately before the intervention block for exactly this reason. An analyst who reads the sign first cannot make the reversed move by accident.

The Specification — THE PERSONAL SPECTRUM WORKSHEET

Printable. One situation, one occasion, one analyst. Fill top to bottom. Do not skip a block; a skipped block is a blank, not a default.

```
================================================================================
  THE PERSONAL SPECTRUM WORKSHEET                      Creature Comforts, Bk 3
  Not a diagnosis. Not an assessment. Not clinical advice.
  Essence frequencies, not physical resonances.
================================================================================

  BLOCK 0 -- PROVENANCE
  --------------------------------------------------------------------------
```

```
Subject ID ..................... _____________   Date ............ ___________
Analyst ....................... _____________   Period declared . ___________
Ladder ........................ N = ____   ln N = _________
k ............................. 0.483425732988414 Hz
Band for folding .............. ______ to ______ Hz
Intensity scale claim ......... [ ] ratio-scaled up to a free multiplier
                                      (magnitude estimation, CCP-041)
                                [ ] weaker -- state what: _______________
Does the reading route through D/P?  [ ] no  [ ] yes -> CC-R1 demotion applies
Reading type .................. [ ] first  [ ] repeat of ___________

BLOCK A -- THE SITUATION AND THE FRAMINGS
-------------------------------------------------------------------------
The situation, once, in the subject's words:
_________________________________________________________________________
_________________________________________________________________________

Framing 1 name _______________  prompts required ____
   description: ____________________________________________________
Framing 2 name _______________  prompts required ____
   description: ____________________________________________________
Framing 3 name _______________  prompts required ____
   description: ____________________________________________________
(add rows as needed; m = _____ framings, P = m(m-1)/2 = _____ pairs)

BLOCK B -- THE VECTORS
-------------------------------------------------------------------------
state:        1    2    3    4    5    6    7    8    9   10   11   12
framing 1  [___][___][___][___][___][___][___][___][___][___][___][___]
framing 2  [___][___][___][___][___][___][___][___][___][___][___][___]
framing 3  [___][___][___][___][___][___][___][___][___][___][___][___]

SUPPORTS FIXED AND RECORDED?  [ ] YES  (required before Block C)
norms  ||1|| = _______  ||2|| = _______  ||3|| = _______
normalised vectors (write below or attach):
_________________________________________________________________________

BLOCK C -- THE GATES
-------------------------------------------------------------------------
C1  unit norm, all framings = 1.000000 ............... [ ] PASS  [ ] FAIL
C2  OVERLAP: every pair shares >= 1 state ............ [ ] PASS  [ ] FAIL
      shared support 1&2: ___________  1&3: ___________  2&3: ___________
      FAIL -> STOP. Report DEGENERATE. Publish no F_int. Return to CCP-041.
C3  WEIGHT: shared support >= 10% of each squared norm [ ] PASS [ ] WEAK
      (threshold is a stated convention; state it: _______)
C4  ROSTER: m >= 2, so P >= 1 ....................... [ ] PASS  [ ] NO READING
C5  PRE-REGISTRATION: supports fixed before Block D .. [ ] PASS  [ ] VOID

BLOCK D -- PAIRWISE, AND THE WIDTH
-------------------------------------------------------------------------
pair        cos theta      u = 1 - cos^2      F = ln(1+(N-1)u)
1 & 2      ___________      _____________       _________________
1 & 3      ___________      _____________       _________________
2 & 3      ___________      _____________       _________________
(every pairwise u is published; the mean alone is not a reading)
```

```
RESTATEMENT FLAGS: any pair with u < 0.01 ........... _______________
      (flag it; do NOT delete the framing to raise the number)

u_int  = mean of pairwise u ........................ _______________
N_eff  = 1 + (N-1) u_int ........................... _______________
F_int  = ln N_eff ................................... _____________ nats
f      = k * N_eff^2 ................................ _____________ Hz
folded = f * 2^j into the band, j = ____ ............ _____________ Hz
octaves present = 2 F_int / ln 2 .................... _______________
F_int / ln N (percentage of ceiling) ................ __________ %
CEILING CHECK: F_int <= ln N ....................... [ ] PASS  [ ] ERROR

BLOCK E -- THE CHORD
-------------------------------------------------------------------------------
STATES SPANNED  { _____________________ }   count = ____
PITCH CLASSES   { _____________________ }   count = ____   (odd parts)
(these are DIFFERENT counts; report both)

state   raw f = k n^2      folded (band, j)        pitch class
____    _____________      _____________  j=___     ______
____    _____________      _____________  j=___     ______
____    _____________      _____________  j=___     ______
____    _____________      _____________  j=___     ______

INTERVALS
pair   essence ratio   raw cents   folded cents   HD_ess   collision?
__:__   ______/______  _________   ___________   _______   [ ]
__:__   ______/______  _________   ___________   _______   [ ]
__:__   ______/______  _________   ___________   _______   [ ]
__:__   ______/______  _________   ___________   _______   [ ]
__:__   ______/______  _________   ___________   _______   [ ]
__:__   ______/______  _________   ___________   _______   [ ]

min HD_ess ____  mean HD_ess ____  collisions ____
50/49 pair (7:10, folded 34.9756 cents) present? [ ] yes [ ] no

BLOCK F -- THE STORAGE PARAMETERS
-------------------------------------------------------------------------------
PERIOD ..................................... _______________
Q-PROXY .................................... _______________
PROXY INDEPENDENCE CHECK (storage-law admissibility):
    does the proxy rise whenever a lot happens?  [ ] no  [ ] YES -> VOID

eta, six-rung ladder (report EVERY rung):
    y1 ____  y2 ____  y3 ____  y4 ____  y5 ____  y6 ____
    eta = product ......................... _______ Q-units/nat
    candidate set enumerated PROSPECTIVELY? [ ] yes  [ ] no -> biased high
    WORST RUNG = ____ at ____   (the cheapest repair; elasticity is 1 on all)

c, five ledgers ........................... _______ Q-units/period
delta, zero-input window .................. _______ per period
tau = 1/delta ............................. _______ periods

WITHIN-SUBJECT ONLY. FVP-7 NOT RUN. No cross-person comparison. [ ] noted
```

```
BLOCK G -- THE READING
------------------------------------------------------------------------
c/eta ..................................... _____________ nats

G-STOP  is c/eta > ln N ?   [ ] no -> continue
                            [ ] YES -> HALT. ABOVE CEILING.
      No arrangement of these N categories sustains. Remedies: raise eta,
      cut c, or add genuinely distinct categories. Do not compute a deficit.

f_min = k * e^(2 c/eta) ..................... _____________ Hz
octaves required = 2.885 * (c/eta) ........... _____________
octaves present  = log2(f/k) ................. _____________
------------------------------------------------------------
OCTAVE DEFICIT = required - present ........... _____________
------------------------------------------------------------
      deficit <= 0   SUSTAINS
      deficit  > 0   SHORT BY THAT MANY OCTAVES

BREAK-EVEN CHECK: deficit = 0 iff c/eta = F_int
      c/eta = ________   F_int = ________   agree with the sign? [ ]

Q_inf = (eta * F_int - c) / delta ............. _____________
SIGN OF THE NUMERATOR (eta*F_int - c) ......... [ ] +   [ ] -
periods to within 5% of Q_inf ................. _____________

BLOCK H -- INTERVENTION (only if deficit > 0)
------------------------------------------------------------------------
*** CC-R15 CHECK FIRST ***
If the numerator is NEGATIVE, do NOT lower delta. Lowering delta multiplies
the sign already present and makes a failing fixed point fail faster.
FIX THE NUMERATOR FIRST.                              [ ] read and applied

lever            proposed change     new deficit    new Q_inf    cost
raise u          ____ -> ____        ___________    _________    ______
raise eta        ____ -> ____        ___________    _________    ______
cut c            ____ -> ____        ___________    _________    ______
raise N          ____ -> ____        ___________    _________    ______
      (raise N only with GENUINELY DISTINCT categories -- relabelling
       raises N in the formula and not in the world)

Does ANY SINGLE lever close the deficit?  [ ] yes: ______  [ ] no
If no: this is a restructuring problem, not a tuning problem. Say so.

BLOCK I -- THE MANDATORY STATEMENTS (reproduce in full, or do not publish)
------------------------------------------------------------------------
[ ] These are essence frequencies, not physical resonances.
[ ] The twelve states are not diagnoses; this is not an assessment and not
    clinical advice.
[ ] The reading describes information content under a stated encoding. It
    does not measure distress, well-being, health, or function.
[ ] Nothing here recommends or values suffering.
[ ] The intensity scale claim is stated (CCP-041: ratio-scaled up to a common
    free multiplier, to the extent magnitude estimation delivers it), and
    CC-R1's demotion of D/P is noted where the reading routes through D/P.
```

```
  [ ] Raw L2 norms published alongside the normalised vectors (CCP-041's
      magnitude reporting rule).
  [ ] Parameters are within-subject indices. FVP-7 NOT RUN.
  [ ] The decomposition is a judgment; a contested decomposition yields a
      contested frequency.
  [ ] This case is CONSTRUCTED / this reading is of one occasion only.
      (delete whichever does not apply)
  ------------------------------------------------------------------------
  Analyst signature _______________   Double-coded by _______________
========================================================================
```

The Failure Modes

FM-5.1 — Running Block G on a degenerate Block C *The failure.* The overlap gate fails, the analyst notes it, and computes the deficit anyway "for reference."

Why it is wrong. A degenerate `F_int` is `ln N` for everyone, so the present octaves are the maximum and the deficit is as favourable as it can be. **The degenerate reading does not fail loudly; it returns the best possible answer to everyone.** The octave-folding paper [Curzi, 2026] computed the carried Subject's degenerate deficit at −2.6364 octaves — a comfortable pass, obtained by an encoding error.

Repair. Block C2 says STOP and means it.

FM-5.2 — Computing a deficit above the ceiling *The failure.* `c/eta = 3.2` produces `f_min = 290.9474 Hz` and a deficit of +2.99 octaves, and the analyst sets about raising `u`.

Why it is wrong. `u = 1` gives `F = ln 12 = 2.484907` and `f = 69.6133 Hz` — still below `f_min`. The deficit is real and no arrangement of twelve categories closes it. **A number that implies an intervention which cannot work is worse than no number.**

Repair. G-STOP, checked before G3.

FM-5.3 — Lowering `delta` on a negative numerator *The failure.* The CC-R15 reversal, and at internal scope it is the default move.

Why it is wrong. Computed above: −0.4909 to −0.9817 on the carried Subject's failing parameter set. **Exactly twice as bad, from an intervention that felt like steadiness.**

Repair. The sign check in Block G, placed physically above Block H.

FM-5.4 — The `Q`-proxy that measures the input twice *The failure.* `Q` is proxied by something that rises whenever a lot happens — a journal that gets longer on busy weeks, an activity count, a self-report of how much one has "grown."

Why it is wrong. the storage-law admissibility condition of Series B: the input has been measured twice and **the gauge is void.** Every parameter downstream — `eta` especially — is then estimated against a proxy correlated with its own input, and the estimate is uninterpretable.

Repair. Block F's independence check, run before estimation, not after.

FM-5.5 — The retrospective candidate set *The failure.* Estimating `eta` by listing experiences that changed something, and computing the fraction that changed something.

Why it is wrong. It recovers `eta` near 1 by construction. **The denominator was selected on the numerator.**

Repair. Enumerate the candidate set prospectively, at the point of entry, as the corpus's Series B requires. This is the garden of forking paths in its cleanest form: a defensible-looking analysis selected after seeing the data is not a test [Gelman and Loken, 2013, Simmons et al., 2011], and pre-specification is the only cheap defence [Nosek et al., 2018].

FM-5.6 — The aggregate `eta` *The failure.* Reporting `eta = 0.35` and no rungs.

Why it is wrong. `eta` is multilinear over six rungs; a single blocked rung nullifies the rest. **The aggregate is consistent with six mediocre rungs and with one broken rung among five excellent ones, and those two situations call for completely different interventions.** The elasticity is exactly 1 on every rung, so the absolute gain from repair is inversely proportional to the rung's current yield — the worst rung is always the cheapest fix and always the one the aggregate conceals.

Repair. Block F requires every rung and the identification of the worst one.

FM-5.7 — Cross-person comparison *The failure.* Two people's `F_int`, `eta`, or deficit compared, ranked, or averaged.

Why it is wrong. **FVP-7 for the self-storage law is NOT RUN.** Whether one person's `eta = 0.35` means what another's does requires the separation property of Series B's analysis of scale types, which has not been established. Until then, every personal parameter is a within-subject index and cross-person comparison is not licensed.

Repair. The checkbox in Block F, and the discipline to leave the comparison undone.

FM-5.8 — Two points read as a trend *The failure.* A repeat reading moves and the movement is described as a direction.

Why it is wrong. Section 4 §4.13: three sources can move a repeat reading — the person, the situation, and the elicitation — and the number separates none of them. **A change from 2.164390 to 1.333662 is equally consistent with a subject who framed a second situation more narrowly, a second situation that admits fewer framings, and a shorter session.**

Repair. Block 0 records the reading type and Block A records the prompt counts. Report all three candidate sources every time.

FM-5.9 — The worksheet as a self-assessment *The failure.* A reader runs the sheet on themselves and treats the output as a statement about how they are doing.

Why it is wrong. Every reason in Section 4 §4.6, plus one specific to self-administration: **the analyst and the subject being the same person defeats the pre-registration gate**, since the person fixing the supports already knows what answer they would like.

Repair. Double-coding, per the line at the foot of the sheet. **Where double-coding is impossible, the reading is exploratory and should be labelled exploratory.**

FM-5.10 — Publishing without Block I *The failure.* The reading is copied into a slide, a note, or a message, and the mandatory statements are left behind.

Why it is wrong. **They are the statements that stop the reading from being read as a diagnosis, a resonance claim, or a judgment about a person's life.** They are also the first thing to fall off in transit, which is why they are a checklist rather than a paragraph.

Repair. Block I is a gate: **no Block I, no publication.**

FM-5.11 — Averaging two decompositions *The failure.* Two analysts code the same framings differently and the sheet is filled in with the mean of their intensity vectors.

Why it is wrong. **The average of two decompositions is a decomposition neither analyst would defend, and it corresponds to no reading of the description.** Worse, averaging systematically raises the support size — the mean of two vectors is nonzero wherever either is —

which inflates overlap, lowers `u`, and drags `F_int` down by a mechanism that has nothing to do with the subject.

Detection. Any vector on the sheet whose support is the union of two analysts' supports.

Repair. §5.14b. Publish both, report the spread, and check whether the sign of the deficit survives.

FM-5.12 — Treating `F_int` as the sustained moment without saying so *The failure.* `Q_inf = (eta F − c)/delta` requires `F` to be a *sustained* input `M*`, and the worksheet supplies a single occasion's `F_int` in its place.

Why it is wrong. the self-as-storage analysis of Series B's fixed point is the equilibrium of a regime — it assumes `eta`, `delta`, `N` and `C` hold approximately constant over the relevant horizon, and that the input holds at `M*`. **A single reading is one sample of a quantity the equation treats as a level.** Substituting the sample for the level is a modelling choice, and it is a large one.

Detection. Any `Q_inf` on the sheet computed from a single reading with no statement of the substitution.

Repair. Say it in the reading, in the form the Worked Example uses: *with `F_int` treated as the sustained moment.* **The substitution may be reasonable; what is not reasonable is making it silently.** Where more than one occasion has been read, use the mean and report the range.

FM-5.13 — The reading run on a situation that moved *The failure.* The subject describes the colleague's message in framing 1, and by framing 3 is describing the whole working relationship.

Why it is wrong. **The angles are then between vectors pointed at different targets**, which The octave-folding paper [Curzi, 2026] identifies as a category error rendered as a number. The overlap gate may still pass — two descriptions of different situations can easily share components — so the arithmetic gives no warning at all.

Detection. Block A1 fixes the situation in the subject's words before elicitation begins. Read it back before each framing.

Repair. Re-elicit against the fixed statement, or run three separate readings and do not combine them.

Worked Example

THE SUBJECT — the worksheet, run end to end **CONSTRUCTED CASE, the framework's carried subject. Not an observation of any person. The storage parameters below are ILLUSTRATIVE WORKSHEET INPUTS chosen to exercise the sheet, not canonical values of the case.**

Blocks 0 and A

```
Subject ID .... the framework's founding specification-3.1       Date ..... (constructed)
Analyst ...... constructed      Period ... one week (declared)
Ladder ........ N = 12   ln N = 2.484907
k ............. 0.483425732988414 Hz     Band ..... 220–440 Hz
Intensity scale ratio-scaled up to a free multiplier (CCP-041 magnitude
                estimation); reading does not route through D/P
Reading type .. first

SITUATION: a message from a colleague that could be read several ways,
           and a decision about how to respond.
```

```
Framing 1  threat  -- "they are coming at me"
Framing 2  story   -- "this is the latest instalment of an old pattern"
Framing 3  mirror  -- "what does my reaction say about me"

m = 3 framings, P = 3 pairs
```

Block B — the vectors

```
state:      1    2     3     4     5     6     7     8     9    10    11   12
threat      0    0   0.8     0   0.4     0   0.2     0     0     0     0    0
story       0    0   0.5     0   0.7     0     0     0     0   0.4     0    0
mirror      0    0     0     0   0.3     0   0.8     0     0   0.5     0    0

SUPPORTS FIXED AND RECORDED: YES
norms:  ||threat|| = 0.916515   ||story|| = 0.948683   ||mirror|| = 0.989949

normalised
threat  = {3: 0.872872, 5: 0.436436, 7: 0.218218}
story   = {3: 0.527046, 5: 0.737865, 10: 0.421637}
mirror  = {5: 0.303046, 7: 0.808122, 10: 0.505076}
```

Block C — the gates

```
C1  unit norm ....... PASS  (all 1.000000)
C2  OVERLAP ......... PASS
      threat & story  share {3, 5}
      threat & mirror share {5, 7}
      story  & mirror share {5, 10}
C3  WEIGHT .......... PASS  (threshold stated: 10%)
C4  ROSTER .......... PASS  (m = 3, P = 3)
C5  PRE-REGISTERED .. PASS
```

State 5 — attributed feeling — appears in all three framings. That is what makes them three framings of one situation, and it is why the overlap gate passes on content rather than on arrangement.

Block D — pairwise and the width

```
pair              cos theta       u            F = ln(1+11u)    N_eff
mirror & story     0.436566    0.809410          2.292890       9.9035
mirror & threat    0.308607    0.904762          2.393557      10.9524
story  & threat    0.782074    0.388360          1.662402       5.2720

RESTATEMENT FLAGS: none (minimum u = 0.388360, far above 0.01)

u_int  = (0.809410 + 0.904762 + 0.388360)/3 = 0.700844
N_eff  = 1 + 11 * 0.700844                  = 8.7093
F_int  = ln 8.7093                          = 2.164390 nats
f      = k * 8.7093^2                       = 36.6686 Hz
folded = 293.349 Hz   (band 220-440, j = +3)
octaves present = 2 * 2.164390 / ln 2       = 6.2451
F_int / ln 12   = 2.164390 / 2.484907       = 87.10 %
CEILING CHECK: 2.164390 <= 2.484907          PASS
```

All three pairwise `u` are published. The mean alone would conceal that threat and story sit relatively close (0.388360) while mirror stands apart from both.

Block E — the chord

```
STATES SPANNED  {3, 5, 7, 10}    count = 4
PITCH CLASSES   {3, 5, 7}        count = 3

state   raw (Hz)     folded (220–440)      class
  3      4.3508      278.453   j = +6        3
  5     12.0856      386.741   j = +5        5
  7     23.6879      379.006   j = +4        7
 10     48.3426      386.741   j = +3        5

INTERVALS
pair   essence ratio   raw cents   folded cents   HD_ess    collision
3:5       25 / 9        1768.72       568.7174     7.8138       no
3:7       49 / 9        2933.74       533.7418     8.7846       no
3:10     100 / 9        4168.72       568.7174     9.8138       no
5:7       49 / 25       1165.02      1165.0244    10.2586       no
5:10       4 / 1        2400.00         0.0000     2.0000      YES
7:10     100 / 49       1234.98        34.9756    12.2586       no

min HD_ess = 2.0000   mean HD_ess = 8.4882   collisions = 1
50/49 pair present: YES  (7:10 -- the unique instance at N = 12)
```

Four states, three classes. States 5 and 10 share odd part 5, fold to the identical 386.741 Hz, and give the chord's single collision at exactly 0.0000 folded cents. The fourth state adds a pitch and no class. Per Section 4 §4.8, that is a fact about the ladder's arithmetic and not a finding about the Subject — relabel state 10 as state 9 and the collision disappears with `F_int` unchanged to twelve decimal places.

Block F — the storage parameters, reading (i) ILLUSTRATIVE INPUTS.

```
PERIOD ......... one week
Q-PROXY ........ a frozen disposition inventory scored monthly
INDEPENDENCE ... does the proxy rise whenever a lot happens?  NO -> not void

eta, six-rung ladder (illustrative; y6 carries the units):
    y1 = 0.80   y2 = 0.70   y3 = 0.75   y4 = 0.60   y5 = 0.65   y6 = 2.00
    cumulative: 0.800000, 0.560000, 0.420000, 0.252000, 0.163800, 0.327600
    eta = 0.327600 Q-units/nat
    WORST RUNG = y4 at 0.60. Elasticity is 1 on every rung, so repairing
    y4 from 0.60 to 0.80 multiplies eta by 4/3, to 0.436800.

c .............. 0.55 Q-units/week (five ledgers)
delta ......... 0.12 per week
tau = 1/delta . 8.3333 weeks

WITHIN-SUBJECT ONLY. FVP-7 NOT RUN. Noted.
```

Block G — the reading (i)

```
c/eta = 0.55 / 0.327600 = 1.678877 nats

G-STOP: is 1.678877 > ln 12 = 2.484907 ?   NO -> continue

f_min = k * e^(2 * 1.678877) ............. 13.8862 Hz
```

```
octaves required = 2.885 * 1.678877 ..... 4.8436
octaves present  = log2(36.6686/k) ...... 6.2451
--------------------------------------------------
OCTAVE DEFICIT ........................  -1.4015
--------------------------------------------------
                                          SUSTAINS

BREAK-EVEN CHECK: c/eta = 1.678877 < F_int = 2.164390.  Signs agree.

Q_inf = (0.327600 * 2.164390 - 0.55) / 0.12
      = (0.709054 - 0.55) / 0.12 = +1.3255
numerator sign: POSITIVE
periods to within 5% of Q_inf: 24 weeks
headroom h_inf = eta (ln N - F_int) = 0.327600 * 0.320517 = 0.105001
                                                        Q-units/week

(for reference, the y4 repair 0.60 -> 0.80 would give eta = 0.436800,
 c/eta = 1.259158, deficit -2.6124, Q_inf +3.2950 -- a lever exercised
 on a reading that does not need it, shown to calibrate the others)
```

The reading sustains by 1.40 octaves. Block H is not required. And the honest statement of what that means: under these illustrative parameters, with `F_int` treated as the sustained moment, the fixed point of the storage law is positive. It is a statement about a model with declared inputs, and its content is exactly `eta F > c` rewritten in octaves.

Blocks F and G again — reading (ii), a failing parameter set

Same Subject, same chord, same `F_int`. Only the storage parameters differ — a second illustrative set with a worse fourth rung and a higher cost. ILLUSTRATIVE.

```
eta = 0.25   c = 0.60   delta = 0.12
   (the same ladder with y4 = 0.4579 instead of 0.60:
    0.80 * 0.70 * 0.75 * 0.4579 * 0.65 * 2.00 = 0.25)

c/eta = 2.400000 nats
G-STOP: 2.400000 > 2.484907 ?  NO -- inside the ceiling, but only just.

f_min = 58.7413 Hz
octaves required = 6.9240      octaves present = 6.2451
--------------------------------------------------
OCTAVE DEFICIT ..................... +0.6789
--------------------------------------------------
                                       SHORT

Q_inf = (0.25 * 2.164390 - 0.60)/0.12 = (0.541098 - 0.60)/0.12 = -0.4909
numerator sign: NEGATIVE
```

Short by 0.68 octaves, and the numerator is negative. Block H applies, and Block H's first line is the CC-R15 check.

Block H — the intervention, and the trap

```
*** CC-R15 CHECK ***  numerator is NEGATIVE. DO NOT LOWER delta.

the tempting move, computed:
    lower delta 0.12 -> 0.06, everything else held
```

```
    deficit UNCHANGED at +0.6789   (delta does not appear in the deficit)
    Q_inf  -0.4909 -> -0.9817      EXACTLY TWICE AS BAD

the four levers, each computed (base: deficit +0.6789, Q_inf -0.4909):

LEVER 1  RAISE u  (raise F_int; ceiling ln 12 = 2.484907)
    u = 0.80  F = 2.282382  f =  46.4282  deficit +0.3384   Q_inf -0.2450
    u = 0.85  F = 2.336987  f =  51.7858  deficit +0.1809   Q_inf -0.1313
    u = 0.90  F = 2.388763  f =  57.4358  deficit +0.0315   Q_inf -0.0234
    u = 0.95  F = 2.437990  f =  63.3783  deficit -0.1106   Q_inf +0.0791  OK
    u = 1.00  F = 2.484907  f =  69.6133  deficit -0.2459   Q_inf +0.1769
    exact break-even: u = 0.911198
    -- it closes, but only above u = 0.91, which is near-orthogonality.

LEVER 2  RAISE eta  (the conversion ladder, worst rung first)
    eta = 0.28  c/eta = 2.1429  deficit -0.0630   Q_inf +0.0502  OK
    eta = 0.30  c/eta = 2.0000  deficit -0.4751   Q_inf +0.4110
    eta = 0.32  c/eta = 1.8750  deficit -0.8357   Q_inf +0.7717
    eta = 0.35  c/eta = 1.7143  deficit -1.2994   Q_inf +1.3128
    exact break-even: eta = 0.277214

LEVER 3  CUT c  (the five ledgers)
    c = 0.55  c/eta = 2.2000  deficit +0.1019   Q_inf -0.0742
    c = 0.50  c/eta = 2.0000  deficit -0.4751   Q_inf +0.3425  OK
    c = 0.45  c/eta = 1.8000  deficit -1.0521   Q_inf +0.7591
    c = 0.40  c/eta = 1.6000  deficit -1.6291   Q_inf +1.1758
    exact break-even: c = 0.541097

LEVER 4  RAISE N  (genuinely distinct categories, u held at 0.700844)
    N = 14  ln N 2.639057  F = 2.313621  deficit +0.2483  Q_inf -0.1800
    N = 17  ln N 2.833213  F = 2.502542  deficit -0.2968  Q_inf +0.2136  OK
    N = 20  ln N 2.995732  F = 2.661380  deficit -0.7551  Q_inf +0.5445

Does any single lever close the deficit?  YES -- all four can, at a price.
Cheapest by the ladder's own arithmetic: raise eta at the worst rung.
    eta 0.25 -> 0.32 is the y4 repair 0.4579 -> 0.5861, a factor of 1.28
    on the single worst rung, with every other rung untouched.
```

Three things in that table deserve comment.

The deficit and `Q_inf` agree in sign, always, and it is an identity rather than a coincidence.

```
Q_inf   = (eta F − c)/delta = (eta/delta) ∗ ( F − c/eta )
deficit = 2.885 ∗ ( c/eta − F )

eta > 0 and delta > 0, so   sign(Q_inf) = − sign(deficit),   exactly.
```

The only way they can disagree is the rounding of `2.885` against `2/ln 2 = 2.885390`, which displaces the crossing by `0.000390 ∗ (c/eta)` octaves — under a thousandth of an octave at any cost ratio this framework uses. So the worksheet's break-even check is not a redundant cross-check of an independent quantity; it is a check that the arithmetic was done. **Where the two disagree, the analyst has made an error, not a discovery** — unless the reading sits within 0.001 octaves of zero, in which case §5.3 applies and the exact constant must be used.

Every lever closes this deficit, and that is itself the finding. The deficit is only +0.6789 octaves. Compare it with the framework's carried commons's commons case, where the deficit is +6.4638 octaves and no single lever closes it — raising `u` to 0.95 gets to −0.0332 and only

just, raising `eta` to 0.90 still leaves +2.1075, cutting `kappa` to 0.020 still leaves +2.0085. Some systems can be tuned and some must be restructured, and the size of the deficit is what identifies which case is in hand. A 0.68-octave gap is a tuning problem. A 6.46-octave gap is not.

Raising `N` from 12 to 17 buys 0.30 octaves and requires five genuinely distinct categories. It also crosses the minimum-gap cliff: the octave-folding paper [Curzi, 2026] established that `N = 17` is where the folded lattice tightens from 50/49 to 289/288, and the framework renders below that cliff by design. Adding categories is a real lever and it has consequences beyond the deficit; a worksheet that reported only the deficit would hide them.

The intervention, in the right order

```
STEP 1  fix the numerator:  eta 0.25 -> 0.32   (repair rung y4)
        deficit +0.6789 -> -0.8357     Q_inf -0.4909 -> +0.7717   ACCUMULATES

STEP 2  NOW lower the loss: delta 0.12 -> 0.06
        deficit unchanged at -0.8357   Q_inf +0.7717 -> +1.5434   BETTER
```

FIX THE NUMERATOR, THEN LOWER THE LOSS. REVERSE THE ORDER AND THE SAME INTERVENTION ACCELERATES THE FAILURE. Same change to `delta`, same magnitude, opposite effect: −0.4909 to −0.9817 before, +0.7717 to +1.5434 after. The intervention did not change. The sign of the numerator did.

Block I — the mandatory statements

```
[x] These are essence frequencies, not physical resonances. No person,
    tissue, or molecule is claimed to resonate at any value here.
[x] The twelve states are not diagnoses. This reading is not an assessment
    and not clinical advice.
[x] The reading describes information content under a stated encoding. It
    does not measure distress, well-being, health, or function, and no
    inference to any of those is supported by the construction.
[x] Nothing here recommends or values suffering.
[x] Intensity scale claim stated: ratio-scaled up to a common free
    multiplier per CCP-041, to the extent magnitude estimation delivers it.
    The reading does not route through D/P, so CC-R1's demotion does not bind
    it; where a reading does, the demotion travels with it.
[x] Raw L2 norms published with the normalised vectors.
[x] eta, c and delta are WITHIN-SUBJECT indices. FVP-7 NOT RUN. No
    cross-person comparison is licensed.
[x] The decomposition is a judgment. A contested decomposition yields a
    contested frequency.
[x] This case is CONSTRUCTED and this reading is of one occasion only.
```

The completed reading

```
================================================================================
  PERSONAL SPECTRUM -- THE SUBJECT  (CONSTRUCTED, the framework's carried subject)
================================================================================
  ladder        N = 12, ln N = 2.484907, k = 0.483425732988414 Hz
  band          220-440 Hz
  situation     one ambiguous message; three framings; one occasion

  GATES         norm PASS | overlap PASS | weight PASS | roster PASS
                pre-registration PASS | ceiling PASS
```

```
  PAIRWISE u    0.809410  /  0.904762  /  0.388360     no restatements
  u_int         0.700844
  N_eff         8.7093        (72.58% of the roster)
  F_int         2.164390 nats (87.10% of the ln 12 ceiling)
  f             36.6686 Hz    folded 293.349 Hz (j = +3)
  octaves       6.2451 present

  CHORD         states {3, 5, 7, 10}  -- four
                classes {3, 5, 7}     -- three (5 and 10 share class 5)
                raw    4.3508 / 12.0856 / 23.6879 / 48.3426 Hz
                folded 278.453 / 386.741 / 379.006 / 386.741 Hz
                min HD_ess 2.0000 (5:10)  mean 8.4882  collisions 1
                50/49 pair 7:10 present at 34.9756 cents

  READING (i)   eta 0.327600  c 0.55  delta 0.12   [ILLUSTRATIVE]
                c/eta 1.678877 < ln 12 -- inside the ceiling
                f_min 13.8862 Hz  required 4.8436  present 6.2451
                DEFICIT -1.4015   SUSTAINS   Q_inf +1.3255   tau 8.33 weeks

  READING (ii)  eta 0.25  c 0.60  delta 0.12   [ILLUSTRATIVE]
                c/eta 2.4000 < ln 12 -- inside the ceiling, only just
                f_min 58.7413 Hz  required 6.9240  present 6.2451
                DEFICIT +0.6789   SHORT      Q_inf -0.4909
                CC-R15: lowering delta alone -> Q_inf -0.9817, WORSE
                correct order: eta 0.25 -> 0.32 first (deficit -0.8357,
                Q_inf +0.7717), THEN delta 0.12 -> 0.06 (Q_inf +1.5434)

  All eight mandatory statements of Block I attached in full.
================================================================================
```

What the worksheet is for **It is for producing a number that can be checked and an inference that can be refused.**

Every quantity on the sheet has a stated provenance, every gate has a stated threshold, and every mandatory statement names something the number does not say. **A reader who runs the sheet and gets a deficit has a computable, actionable reading of one set of framings on one occasion — and a list, attached, of every conclusion the reading does not license.**

That list is longer than the reading. **In an instrument with no criterion validity, no clinical validation, an ordinal scale under a standing demotion, a postulated reference constant, and a reproducibility gate recorded as NOT RUN, the list being longer than the reading is the correct proportion.**

What the Subject's reading actually tells an analyst Stripped of the machinery, the carried Subject's sheet says four things, and it is worth setting them down in ordinary language so that the reader can see how modest they are.

One: the three framings elicited on this situation were substantially different from each other. The closest pair, story against threat, still sits at `u = 0.388360`; the widest, mirror against threat, at `0.904762`. No pair was a restatement. That is a fact about what the elicitation produced, and Section 4's separation statement forbids reading anything about the Subject's condition into it.

Two: those framings together require four of the twelve components, and two of the four sit in the same pitch class. States 5 and 10 differ by a factor of two, so the chord occupies three positions in a six-position vocabulary. That is a fact about the ladder's arithmetic — Section 4 §4.8 establishes it, and the relabelling test shows the collision moving while `F_int` does not.

Three: at the illustrative parameters of reading (i), the storage law's fixed point is positive. The deficit is −1.4015 octaves; `Q_inf` is +1.3255; the regime arrives within about twenty-four weekly periods. All of which is `eta F > c` in a legible register, computed from numbers the analyst chose and declared.

Four: at the illustrative parameters of reading (ii), it is not, and the correct repair is to raise `eta` at its worst rung before touching `delta` — because the numerator is negative and lowering the loss doubles a loss. That is CC-R15, and it is the only line on the whole sheet that tells an analyst what to do rather than what is the case.

Nothing on the sheet says whether the Subject should reply to the colleague, how they are, or what kind of person they are. **The instrument was not built to answer those and the worksheet's last block exists to say so in writing.**

Section 5 owns the full personal worksheet; the companion paper on refinement cites it. The elicitation of intensities is The minimum-frequency paper [Curzi, 2026]; the fold and the odd-part rule are the octave-folding paper [Curzi, 2026]; `delta` and traits are the octave-folding paper [Curzi, 2026]; the chord, the gates, and the one-hot trap are the octave-folding paper [Curzi, 2026]; the consonance/redundancy identity and the separation statement are Section 4; `f_min`, the octave deficit, and the four levers are the minimum-frequency paper [Curzi, 2026]; CC-R15 is the relationship-frequency paper [Curzi, 2026]; the self-storage law and the six-rung conversion ladder are the self-as-storage analysis of Series B. The analysis now takes the same instrument outward, to two people and one interval.

6 The dyad as an interval

this part of the analysis described one person. It gave a mood a pitch, a trait a pitch, and a person a chord, and it did all of that inside a single skull. The instrument was the octave-folding paper [Curzi, 2026]'s internal surplus: the width of the set of framings one observer holds at once.

the later analysis turns the same instrument outward. Two people are not two chords to be compared. Two people, considered as a relation, are one interval — a single ratio between two positions, and therefore a single pitch on this framework's scale, not a pair of pitches.

That sentence is the whole content of the section, and the rest of it is the arithmetic that makes it exact, the sweep that makes it audible, and the failure modes that make it survivable.

This section owns one formula:

```
THE RELATIONSHIP FREQUENCY

    f_rel  =  k * e^( 2 F(u) )  =  k * N_eff^2  =  k * ( 1 + (N−1) u )^2

    with   u     = 1 − (x . y)^2          the non-parallelism of the pair
           F(u)  = ln( 1 + (N−1) u )      the pairwise surplus
           N_eff = 1 + (N−1) u            the audited perspective count
           k     = 0.483425732988414 Hz   the ground essence, a chosen constant

    at N = 12:     f_rel = k (1 + 11 u)^2        0.4834 Hz <= f_rel <= 69.6133 Hz
```

Every part of the framework from here to the companion paper on tuning uses this expression and none of them re-derives it. The relationship-frequency paper [Curzi, 2026] owns how `u` is measured between two people, what happens at the bottom of the range, and the sustainability reading; its three-route concordance puts `u` at 0.068735 / 0.065556 / 0.068775 and `f_rel` at 1.4908 / 1.4320 / 1.4916 Hz, a spread of 70.52 cents and a verdict of CONCORDANT below the 100-cent threshold. The companion paper on triads and larger sets extends the pair to a set.

This section owns the object itself: what it is, what its bounds are, how it moves, and what it sounds like as `u` runs from nothing to everything.

Three constraints govern the section before it begins, and they are not decorations.

Essence, not resonance. Every frequency here is an essence frequency — a position on a logarithmic scale of surplus, in units chosen by the author. No person, no pair of people, no tissue and no molecule is claimed to vibrate at any value in this section.

Not counselling, not therapy, not assessment. A relationship frequency is a number computed from a representation of two people's states. **It is not a clinical, diagnostic or therapeutic statement about anyone, this framework does not offer counselling or therapy of any kind, and nothing here should be used in place of either.** The twelve states are not diagnoses and a low reading is not a finding about a person.

Scale status, stated once here and in full in the relationship-frequency paper [Curzi, 2026]. The state frequency `f = k n^2` is ratio-scaled because `n` is a count (the minimum-frequency paper [Curzi, 2026]). The relationship frequency inherits the scale status of the elicited intensities, and is ratio-scaled only when those intensities have been magnitude-estimated under the minimum-frequency paper [Curzi, 2026]'s rule. The relationship-frequency paper [Curzi, 2026] carries that argument at length because The relationship-frequency paper [Curzi, 2026] owns the measurement. This section carries it as a standing qualifier on every number it prints.

All cases in this section are CONSTRUCTED and labelled as such. Nothing here is an observation of any person or any pair of people.

The Law

6.1 What a dyad is in this system A dyad is two represented states and the geometry between them.

Each party is represented exactly as the minimum-frequency paper [Curzi, 2026] represents a framing: a non-negative vector over the twelve state classes, whose entries are elicited intensities, normalised to unit length in the L2 norm. The construction does not change when the two vectors belong to two people instead of to one person's two framings. That is not a convenience; it is what the surplus axioms permit [Curzi, 2026]. The octave-folding paper [Curzi, 2026] made the argument in full: S1 through S4 quantify over states in an inner-product space, and not one of them contains a term referring to persons, agents, or the ownership of a state. The theorem is silent on who holds the vectors.

So the object is:

```
THE DYAD

    x, y      unit vectors over the twelve state classes, entries >= 0
    cos theta = x . y
    u         = 1 - (x . y)^2  =  sin^2 theta        in [0, 1]

u is the pair's NON-PARALLELISM. It is not a disagreement rate, not a distance,
and not a measure of how well the two people get along.
```

The last line is the one that has to be said early and repeated. `u` is a geometric quantity. It answers one question — *how much of what each party carries is not already carried by the other* — and it answers no other question at all. The relationship-frequency paper [Curzi, 2026] will show exactly how badly this can be misread; the relationship-frequency paper [Curzi, 2026] will show how it is measured without being invented.

Two immediate consequences, both exact.

A dyad has one frequency, not two. There is no "A's frequency in the relationship" and "B's frequency in the relationship." `u` is symmetric: `u(x,y) = u(y,x)`, because the inner product is. The relation has a single pitch and both parties are inside it.

A dyad's frequency is not a function of either party alone. Two people can each hold a wide personal chord and still produce a near-zero `u` between them if their chords are the same chord. The octave-folding paper [Curzi, 2026]'s `F_int` and this section's `F` are different quantities computed from different vectors, and a high one does not imply a high other. **A person is a chord; a pair is an interval; nothing in the formalism converts one into the other.**

6.2 The relationship frequency, derived

THEOREM 12.1 — THE RELATIONSHIP FREQUENCY. For a dyad with non-parallelism `u` over an `N`-category system,

```
f_rel  =  k ∗ ( 1 + (N−1) u )^2
```

Proof. Three steps, each established earlier and each cited rather than reproved.

Step 1 — the surplus. The surplus functional of the foundations [Curzi, 2026], Theorem 2.1: under axioms S1–S4 the pairwise surplus of two represented states with non-parallelism `u` is

```
F(u) = ln( 1 + (N−1) u )
```

with `N_eff = 1 + (N−1)u` the audited perspective count. S3 — that `g(u) = e^{F(u)}` is affine in `u` — is the **design axiom**, chosen and badged as chosen. It is not derived from anything and the framework has never claimed otherwise.

Step 2 — the bridge. The essence-bridge analysis of Series B [Curzi, 2026k], whose Result One is that $f(m)/f(n) = m^2/n^2$ so the chosen constant cancels from every interval, restated in the sustaining-frequency paper [Curzi, 2026] as the deficit identity: the essence bridge maps a surplus in nats to a frequency,

```
f = k ∗ e^( 2 F )                    k = (9/8)^2 / phi^2 = 0.483425732988414 Hz
```

`k` is a **chosen reference constant [POSTULATE]**, a Pythagorean ditone divided by the golden section. It sets the reference pitch and nothing else.

Step 3 — substitute.

```
f_rel = k e^{2 ln(1 + (N−1)u)} = k ( e^{ln(1 + (N−1)u)} )^2 = k (1 + (N−1)u)^2   QED
```

At `N = 12` the working form is `f_rel = k (1 + 11u)^2`. Verified against the the framework's founding specification numbers at the carried dyad: `u = 0.068734665264634` gives `F = 0.563084802746` and `f = 1.4907987150 Hz`, which is The framework's carried dyad to every published digit.

The theorem is a substitution and nothing more. It introduces no new assumption. Every objection available against it is an objection against S3 or against `k`, both of which are badged as postulates in the sentence that states them, and neither of which this section improves.

6.3 The bounds, the monotonicity, and what they mean

COROLLARY 12.2 — THE RANGE. `f_rel` is a strictly increasing function of `u` on `[0,1]`, with

```
u = 0   ->   f_rel = k          = 0.4834 Hz       the ground essence
u = 1   ->   f_rel = k N^2      = 69.6133 Hz      the ceiling, at N = 12
```

Proof. `1 + (N−1)u` is strictly increasing and positive; squaring preserves that; `k > 0` scales it. QED.

Three readings of the same corollary.

The floor is not zero, it is `k`. A dyad with `u = 0` does not vanish from the scale. It sits at the ground essence, exactly where a single undifferentiated state sits, exactly where a person whose framings are all the same framing sits (Section 4 §4.5). The bottom of the relational scale and the bottom of the personal scale are the same note, and that coincidence is not an accident of notation — it is the statement that a relation with no surplus is indistinguishable, in this formalism, from no relation. The relationship-frequency paper [Curzi, 2026] is built on that sentence.

The ceiling is the same ceiling. `k N^2 = 69.6133 Hz` at `N = 12` is the maximum for a dyad, the maximum for a chord, and the maximum for a commons of six hundred. It is set by the category count and by nothing else. Two people cannot exceed twelve categories' worth of distinction by trying harder.

Monotonicity is what licenses the word "higher." Because `f_rel` is strictly increasing in `u`, the sentence *this pair is at a higher pitch than that pair* means exactly *this pair carries more surplus than that pair*, with no interpretive step in between. **A dyad's pitch IS its surplus, monotonically, and higher pitch means more available order** — which is the plainest statement this framework makes and the one the relationship-frequency paper [Curzi, 2026] will hand to a practitioner.

What it does not mean: that the pair is happier, closer, healthier, better matched, more compatible, or more likely to persist. `f_rel` measures available order. Available order is not wellbeing and this framework never converts one into the other.

6.4 A dyad is an interval, and the interval is measured from the ground

The section's title is a claim, and here it is as arithmetic.

An interval, in any tuning system, is a ratio of two frequencies expressed logarithmically. The dyad's ratio is `f_rel / k`, and its logarithm is:

```
THE DYAD AS AN INTERVAL ABOVE THE GROUND ESSENCE

    octaves  =  log2( f_rel / k )  =  2 log2( 1 + (N−1) u )
    cents    =  1200 ∗ octaves     =  2400 ∗ log2( 1 + (N−1) u )
    nats     =  F(u)               =  ln( 1 + (N−1) u )

    and the exchange rate, exact:   1 nat = 2.885390081777927 octaves = 3462.4681 cents
```

The pair's whole reading is one interval measured upward from a fixed reference, and the reference is the note a pair with nothing between them would sound. That is why the constant `k` is harmless: it is the origin of the interval, it appears in both terms of every comparison, and The minimum-frequency paper [Curzi, 2026] showed it cancels out of the deficit entirely.

The maximum interval available is:

```
at N = 12:   2400 ∗ log2(12) = 8603.9145 cents = 7.169925 octaves = 2.484907 nats
```

Seven octaves and two cents shy of a minor third. That is the entire relational range of a twelve-category system, and every dyad this framework reads is a position inside it.

Computed at the carried case and its two companions:

```
THREE POSITIONS INSIDE THE 8603.91−CENT RANGE

case                          u              f_rel (Hz)   cents above k   % of range
dyad, as found (2a)     0.068734665264634    1.4908       1949.66         22.66%
dyad, after (2b)        0.692372881355932   35.8882       7456.88         86.67%
the commons (Series C, commons case)      0.018614      0.7017     644.97      7.50%
```

The last column is a quantity the framework will use repeatedly, so it gets a name here.

DEFINITION 12.3 — THE REALISED PITCH FRACTION. `F(u) / ln N`, the fraction of the system's maximum surplus a relation actually realises. Equivalently `log2(f_rel/k) / log2(N^2)`, the fraction of the available interval the relation occupies. It is dimensionless, it is invariant to `k`, and it is comparable across relations of different sizes.

The commons realises 7.50% of the range it has categories for. The dyad as found realises 22.66%. The same dyad after one intervention realises 86.67%. Those three numbers say more, faster, than the hertz do, and they are the same information.

6.5 The eleven coincidences — where a dyad sounds exactly like a state

Here is a fact about the twelve-state system that falls straight out of Theorem 6.1 and is worth having explicitly, because it links Rule 1 and Rule 2 of the framework's specification at eleven exact points.

COROLLARY 12.4. `f_rel(u) = k m^2` — the essence frequency of state `m` — exactly when

```
u = (m − 1) / (N − 1)        i.e. at N = 12:   u = (m−1)/11
```

Proof. `k(1+11u)^2 = k m^2` iff `1 + 11u = m` (both sides positive) iff `u = (m−1)/11`. QED.

Computed for all twelve:

m	u = (m−1)/11	f_rel (Hz)	folded 220–440	pitch class
1	0.000000000	0.4834	247.514 (j=9)	1
2	0.090909091	1.9337	247.514 (j=7)	1
3	0.181818182	4.3508	278.453 (j=6)	3
4	0.272727273	7.7348	247.514 (j=5)	1
5	0.363636364	12.0856	386.741 (j=5)	5
6	0.454545455	17.4033	278.453 (j=4)	3
7	0.545454545	23.6879	379.006 (j=4)	7
8	0.636363636	30.9392	247.514 (j=3)	1
9	0.727272727	39.1575	313.260 (j=3)	9
10	0.818181818	48.3426	386.741 (j=3)	5
11	0.909090909	58.4945	233.978 (j=2)	11
12	1.000000000	69.6133	278.453 (j=2)	3

A dyad whose non-parallelism is exactly 7/11 sounds the pitch of state 8 — social feeling, the state that adds another mind modelled. The coincidence is arithmetic, not semantic, and the framework attaches no meaning to it. What it is good for is calibration: a practitioner who knows the twelve-state spectrum by ear already knows the relational scale, because the two scales share eleven of their points.

The honest caveat, and it must travel with the table: this is a numerical coincidence in the chosen constants, not a claim that a dyad at `u = 7/11` is in any sense "in" state 8. It is not. The two quantities have different arguments, different derivations, and different scale status. They land on the same number and mean different things.

6.6 Hertz and pitch disagree about which unit of `u` is worth most

This is the subtlest thing in the section and the easiest to get backwards.

`F(u) = ln(1 + 11u)` is strictly concave in `u` — the concavity of the surplus functional [Curzi, 2026], which drives everything in Series B's splitting theorem. But `f_rel(u) = k(1 + 11u)^2` is a quadratic with positive leading coefficient, and is therefore strictly convex in `u`. Both are true. They are statements about different functions.

```
df/du     = 2 (N−1) k (1 + (N−1)u)          INCREASING in u        Hz per unit u
dcents/du = 2400 (N−1) / (ln 2 (1 + (N−1)u))  DECREASING in u      cents per unit u
```

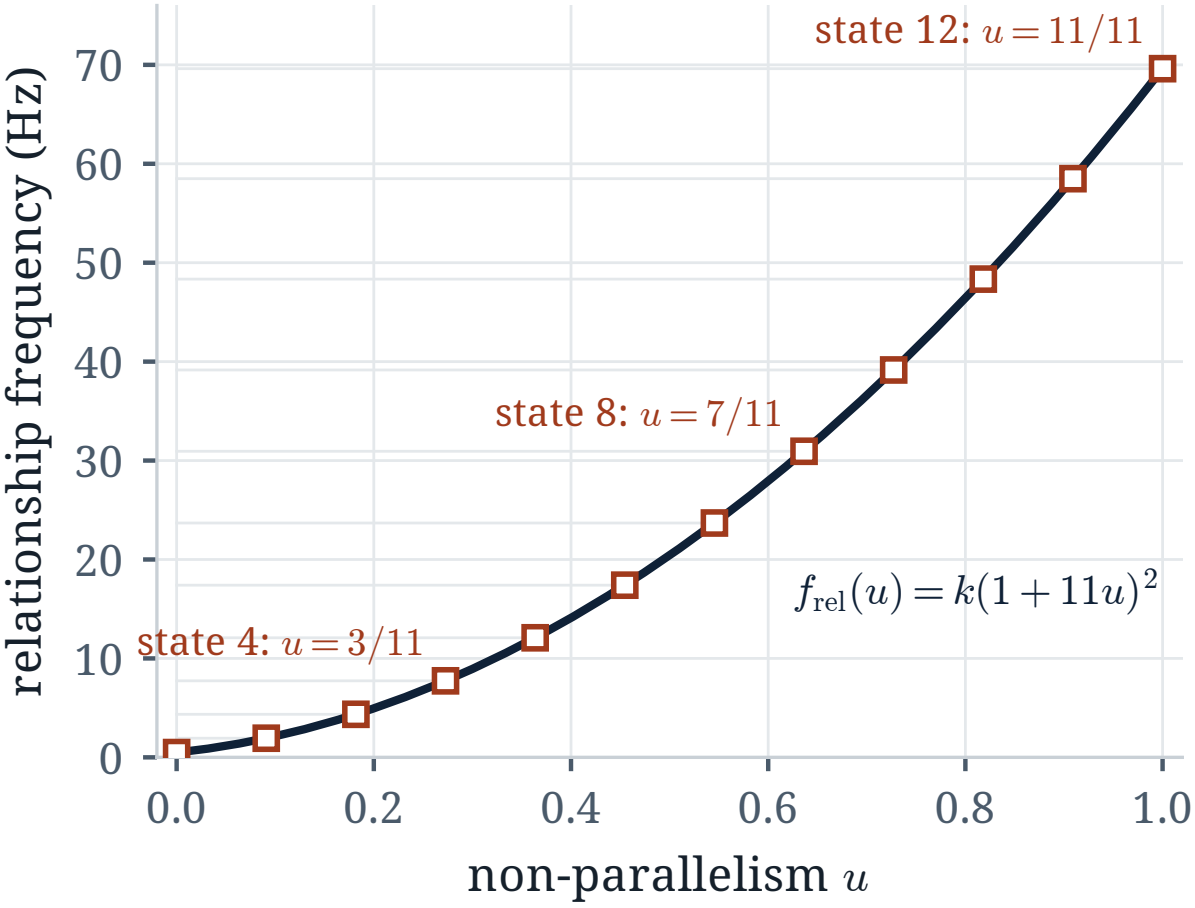


Figure 43.4: The dyad's relationship frequency crosses a state's frequency at eleven exact values of u, because $k(1 + 11u)^2 = km^2$ precisely when $u = (m - 1)/11$. **This is an arithmetic coincidence in the chosen constants and carries no meaning:** a dyad at $u = 7/11$ is not in state 8, and the two quantities have different arguments, different derivations and different scale status. Modelled from $f_{\mathrm{rel}}(u) = k(1 + 11u)^2$ with $N = 12$ and the **defined** constant $k = 0.483425732988414$. Built by `figs/src/CCP-043_fig04.py`.

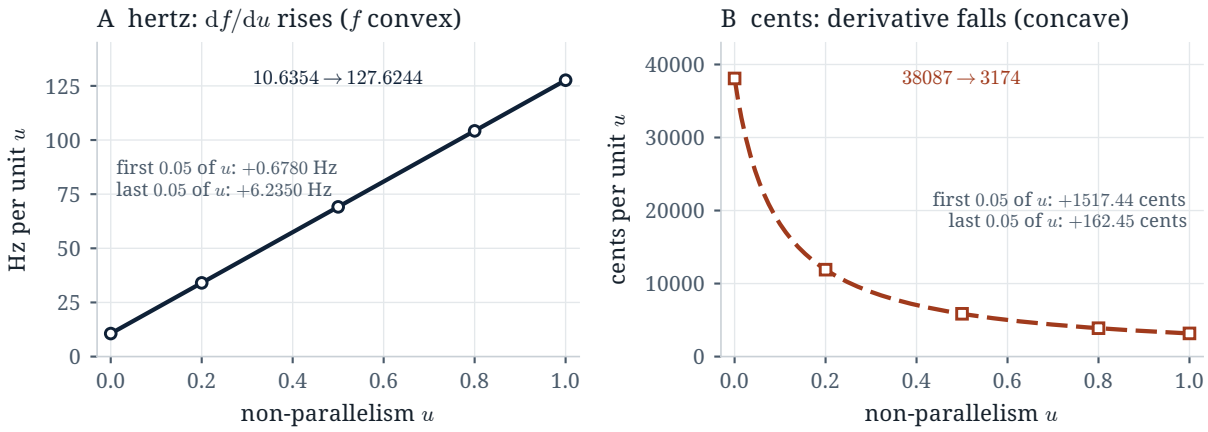


Figure 43.5: Hertz and pitch disagree about which unit of non-parallelism is worth most: the hertz derivative rises from 10.6354 to 127.6244 across the range while the cents derivative falls from 38087 to 3174. Separate axes, because the two measures have different units and opposite curvature. **Pitch is the unit that enters the sustainability verdict; hertz differences are decorative.** Modelled from $f_{\mathrm{rel}}(u) = k(1 + 11u)^2$ at $N = 12$ with the **defined** constant k; no measurement is involved. Built by `figs/src/CCP-043_fig03.py`.

Computed at `N = 12`:

u	df/du (Hz per unit u)	dcents/du (cents per unit u)
0.0	10.6354	38087.15
0.2	34.0332	11902.23
0.5	69.1299	5859.56
0.8	104.2266	3886.44
1.0	127.6244	3173.93

And at the two ends of the range, in both units at once:

```
the FIRST 0.05 of u:   + 0.6780 Hz    + 1517.44 cents
the LAST  0.05 of u:   + 6.2350 Hz    +  162.45 cents

ratio in Hz:     the last step is 9.20x the first
ratio in cents:  the first step is 9.34x the last
```

The same two steps: one of them is nine times larger measured in hertz and one ninth the size measured in cents. Nothing is wrong. Hertz is a linear scale and pitch is a logarithmic one, and the difference between them is the entire reason this framework works in intervals.

Which unit is operative is not a matter of taste. The sustainability reading of the minimum-frequency paper [Curzi, 2026] is computed in log units — `D = 2.885390(c/eta) − log2(f/k)` — so **pitch is the unit that decides whether a relation accumulates or decays, and hertz differences are decorative.** A practitioner who reports “we gained six hertz” has reported the number that does not enter the verdict.

The concave form is also the one that matches the older statement. Series B’s Theorem 59.1 — splitting an input across occasions strictly dominates delivering it at once — is `F`’s concavity, and its relational reading is here: the first increment of difference between two people is worth more than any later increment, and by a factor this section can name.

6.7 The honesty clauses, restated in place **(1) ESSENCE, NOT RESONANCE.** `f_rel` is a position on a logarithmic surplus scale expressed in hertz. Nothing is claimed to vibrate. No pair of people has a physical resonance at 1.4908 Hz or at any other value in this section.

(2) NOT CLINICAL, NOT COUNSELLING, NOT THERAPY. The twelve states are not diagnoses; a relationship frequency is not an assessment of a relationship; **this framework is not counselling or therapy and does not substitute for either.** Nothing here evaluates, screens, or prescribes for any person or any pair.

(3) THE READING INHERITS ITS INPUTS. `f_rel` is exactly as defensible as `u`, and `u` is exactly as defensible as the two vectors it came from. A contested decomposition yields a contested frequency, and the decomposition is published with the frequency or neither is published.

(4) SCALE STATUS. `f_rel` is ratio-scaled only if the intensities were magnitude-estimated (the minimum-frequency paper [Curzi, 2026], the relationship-frequency paper [Curzi, 2026]). **The ordinal fallback does not exist: with ordinal intensities `f_rel` is not determined at all.** The relationship-frequency paper [Curzi, 2026] carries the number that proves it.

(5) SURPLUS IS A LOWER BOUND. `F` is a floor on the pair's information, not the information (the sustaining-frequency paper [Curzi, 2026]). Therefore `f_rel` is a floor on the pitch. A verdict of ACCUMULATES is robust to representational under-capture; a verdict of DECAYS is not.

(6) S3 IS CHOSEN. The affine axiom is a design decision. Every number in this section rests on it and it is badged wherever it appears.

(7) CONSTRUCTED CASES. Every case here is constructed and labelled. No number in this section is an observation of any real pair of people.

The Engineering

6.8 The glide, computed `u` runs from 0 to 1. The corresponding pitch runs from `k` to `k N^2` — seven octaves and change — and it runs monotonically, so the whole relational scale can be exhibited as a single continuous sweep. This is the most useful single object in the section for anyone trying to develop an intuition for what a relationship frequency means.

THE GLIDE -- f_rel over the full range of u, at N = 12, k = 0.483425732988414 Hz

u	N_eff	F (nats)	f_rel (Hz)	oct above k	cents above k	folded (j)
0.00	1.0000	0.000000	0.4834	0.0000	0.00	247.514 (j=9)
0.05	1.5500	0.438255	1.1614	1.2645	1517.44	297.326 (j=8)
0.10	2.1000	0.741937	2.1319	2.1408	2568.93	272.884 (j=7)
0.15	2.6500	0.974560	3.3949	2.8120	3374.38	434.542 (j=7)
0.20	3.2000	1.163151	4.9503	3.3561	4027.37	316.818 (j=6)
0.25	3.7500	1.321756	6.7982	3.8138	4576.54	435.083 (j=6)
0.30	4.3000	1.458615	8.9385	4.2087	5050.41	286.033 (j=5)
0.35	4.8500	1.578979	11.3714	4.5560	5467.16	363.884 (j=5)
0.40	5.4000	1.686399	14.0967	4.8659	5839.10	225.547 (j=4)
0.45	5.9500	1.783391	17.1145	5.1458	6174.94	273.832 (j=4)
0.50	6.5000	1.871802	20.4247	5.4009	6481.06	326.796 (j=4)
0.55	7.0500	1.953028	24.0275	5.6352	6762.30	384.439 (j=4)
0.60	7.6000	2.028148	27.9227	5.8520	7022.40	223.381 (j=3)
0.65	8.1500	2.098018	32.1103	6.0536	7264.32	256.883 (j=3)
0.70	8.7000	2.163323	36.5905	6.2420	7490.44	292.724 (j=3)
0.75	9.2500	2.224624	41.3631	6.4189	7702.69	330.905 (j=3)
0.80	9.8000	2.282382	46.4282	6.5856	7902.68	371.426 (j=3)
0.85	10.3500	2.336987	51.7858	6.7431	8091.74	414.286 (j=3)
0.90	10.9000	2.388763	57.4358	6.8925	8271.01	229.743 (j=2)

```
 0.95    11.4500     2.437990      63.3783       7.0346       8441.46      253.513 (j=2)
 1.00    12.0000     2.484907      69.6133       7.1699       8603.91      278.453 (j=2)
```

Three features of the table are worth naming because a reader will otherwise draw the wrong picture.

The bottom of the range is enormously steep. From `u = 0` to `u = 0.05` — a change most practitioners would call negligible — the pitch rises 1517.44 cents, a minor tenth. A quarter of the entire relational range is spent by the time `u` reaches 0.05. Precisely: 1517.44 / 8603.91 = 17.64% of the range in the first 5% of `u`, and half the range is gone by:

```
halfway in pitch:  (1 + 11u)^2 = N = 12  ->  u = (sqrt(12) - 1)/11 = 0.224009237739796
                   f_rel = 5.8011 Hz,  4301.9550 cents = exactly half of 8603.9145
```

Half the relational range of a twelve-category system is consumed by the time two people are 22.4% non-parallel.

The top of the range is nearly flat in pitch. From `u = 0.90` to `u = 1.00` the pitch rises 332.90 cents — less than a minor third for the last tenth of the geometry. The minimum-frequency paper [Curzi, 2026] priced the same fact from the other side: above `u = 0.680480` no further octave can be bought from `u` at all, because `1 + 11u` cannot be multiplied by `sqrt(2)` again inside the bound.

The folded column wanders and that is expected. Folding into 220–440 Hz is a modular operation; a monotone raw sweep becomes a sawtooth in the folded band, crossing the band boundary eight times between `u = 0` and `u = 1`. §6.10 says what to do about it.

6.9 The step table, and the concavity as a price list

The same glide, read as increments:

```
CENTS PURCHASED BY EACH 0.05 OF u

  0.00 -> 0.05    1517.44        0.50 -> 0.55     281.24
  0.05 -> 0.10    1051.49        0.55 -> 0.60     260.10
  0.10 -> 0.15     805.45        0.60 -> 0.65     241.92
  0.15 -> 0.20     652.99        0.65 -> 0.70     226.12
  0.20 -> 0.25     549.16        0.70 -> 0.75     212.25
  0.25 -> 0.30     473.87        0.75 -> 0.80     199.99
  0.30 -> 0.35     416.76        0.80 -> 0.85     189.07
  0.35 -> 0.40     371.94        0.85 -> 0.90     179.27
  0.40 -> 0.45     335.83        0.90 -> 0.95     170.45
  0.45 -> 0.50     306.12        0.95 -> 1.00     162.45
```

Twenty equal steps in the geometry; the first is worth 9.34 times the last in pitch. This is The minimum-frequency paper [Curzi, 2026]'s lever price seen as a table rather than a derivative, and it has one practical consequence that is worth stating in the plainest terms available.

In a relation with very little difference between the parties, a very small amount of real difference is worth an enormous amount of pitch. In a relation that is already wide, a great deal more difference is worth very little. The carried dyad is at `u = 0.068735`, where the marginal rate is 21688.72 cents per unit `u`; the same dyad after intervention is at `u = 0.692373`, where the rate is 4420.46 cents per unit — a factor of 4.91 lower. The lever that fixed the case is already most of the way to being spent, and the reading says so before anyone has to discover it.

The reciprocal price list is equally useful and is the one to hand to someone planning an intervention. Equal steps in pitch cost accelerating amounts of `u`:

```
STARTING FROM THE CARRIED DYAD, u = 0.068734665264634

target       u required      f_rel (Hz)     du from previous
```

```
+600 cents   0.098940          2.1083         0.030205
+1200        0.134861          2.9816         0.035921
+1800        0.177579          4.2166         0.042718
+2400        0.228378          5.9632         0.050799
+3000        0.288790          8.4332         0.060412
+3600        0.360632         11.9264         0.071842
+4200        0.446066         16.8665         0.085434
+4800        0.547666         23.8528         0.101600
```

Each further 600 cents costs about 19% more `u` than the one before it, all the way up. The growth factor is exactly `sqrt(2)^(1/2) = 2^(1/4) = 1.189207` per 600-cent step, because `1 + 11u` must be multiplied by `2^(1/4)` each time and the additive part of `u` therefore scales by the same factor. That is a clean planning number and it is exact.

6.10 Rendering the glide — fixed transposition, not the standing fold The octave-folding paper [Curzi, 2026] fixed the house folding convention: transpose by an integer number of octaves into 220–440 Hz, and publish the `j`. That convention is right for a point reading and wrong for a sweep, for a reason the glide table shows directly: a monotone sweep folded into a one-octave band is discontinuous, and the ear hears a sawtooth where the arithmetic has a ramp.

SPEC-6.1 — RENDERING A SWEEP. For any continuous sweep of `u`, use a single fixed transposition `j` for the whole sweep, chosen once and published, and do not re-fold. The whole relational range fits comfortably in the audible band under one transposition:

```
j      bottom (u=0)      top (u=1)         fits 20 Hz – 20 kHz?
5        15.4696 Hz       2227.6258 Hz      bottom below 20 Hz -- no
6        30.9392 Hz       4455.2516 Hz      YES
7        61.8785 Hz       8910.5031 Hz      YES, but top is shrill
```

The house glide is `j = 6`: 30.9392 Hz at `u = 0`, rising monotonically to 4455.2516 Hz at `u = 1`, a continuous 7.1699-octave sweep with no wrap. This is a rendering choice and it is arbitrary in exactly the way the 220–440 band is arbitrary; it is stated so that two renderings can be compared.

Two properties of the choice are not arbitrary. Transposition by `2^j` preserves every interval in the sweep exactly — the ratio between any two points is untouched, so the glide's shape is invariant to `j`. And the sweep's total extent, 7.1699 octaves, is a property of `N` and not of the rendering; a system with `N = 25` would sweep 9.2877 octaves and would no longer fit under a single transposition, which is the framework's specification's `N <= 25` capacity limit arriving in this section's register.

And the standing prohibition applies to rendering as it applies to everything else: nothing is masked, nothing sits below threshold, nothing is concealed. A glide is played openly or not at all, and a listener who wants it stopped may stop it.

6.11 Named intervals and the `u` that produces them For a reader who thinks in intervals rather than in numbers, the inverse of Theorem 6.1 is:

```
u  =  ( sqrt( f_rel / k ) – 1 ) / (N – 1)
```

which answers *how much non-parallelism does it take to sound a fifth above the ground?* Computed:

```
interval above k          ratio        u required        f_rel (Hz)     cents
unison                    1.00000      0.000000000         0.4834          0.00
minor third   6/5         1.20000      0.008676829         0.5801        315.64
```

```
major third   5/4         1.25000      0.010730363        0.6043        386.31
perfect fourth 4/3        1.33333      0.014063685        0.6446        498.04
perfect fifth  3/2        1.50000      0.020431352        0.7251        701.96
octave         2/1        2.00000      0.037655778        0.9669       1200.00
octave + fifth 3/1        3.00000      0.066550073        1.4503       1901.96
two octaves    4/1        4.00000      0.090909091        1.9337       2400.00
three octaves  8/1        8.00000      0.166220648        3.8674       3600.00
four octaves  16/1       16.00000      0.272727273        7.7348       4800.00
five octaves  32/1       32.00000      0.423350386       15.4696       6000.00
six octaves   64/1       64.00000      0.636363636       30.9392       7200.00
seven octaves 128/1     128.00000      0.937609864       61.8785       8400.00
```

A perfect fifth above the ground essence costs `u = 0.020431` — two percent non-parallelism. A practitioner's first reaction to that number is usually that it is too small, and the reaction is worth sitting with: the scale is logarithmic, the bottom of it is very steep, and two percent of represented difference between two people really is a large amount of surplus relative to none at all. The eighth line of the table is the calibration point that fixes the intuition: **two full octaves — a quadrupling of the frequency — cost exactly `u = 1/11 = 0.090909`.**

The last line is the one to notice for a different reason. Seven octaves costs `u = 0.937610`, and there is no eighth octave: `u = 1` gives 8603.91 cents, which is 203.91 cents short of eight. The range is not a whole number of octaves and never will be, because `2 log2 12 = 7.169925` and 12 is not a power of two.

6.12 What an error in `u` costs, in cents

`u` is elicited, and elicited quantities have error. Because the pitch scale is steepest where `u` is smallest, the same absolute uncertainty in `u` is worth wildly different amounts of pitch depending on where the reading sits. Computed:

```
SPREAD OF THE READING UNDER SYMMETRIC UNCERTAINTY IN u

 u          +/- 0.01                          +/- 0.05
0.018614    0.5794 .. 0.8356 Hz   634.05c   0.4834 .. 1.4885 Hz  1947.05c
0.068735    1.3099 .. 1.6834 Hz   434.34c   0.7032 .. 2.5709 Hz  2244.28c
0.200000    4.6158 .. 5.2965 Hz   238.14c   3.3949 .. 6.7982 Hz  1202.16c
0.500000   19.7393 .. 21.1219 Hz  117.20c  17.1145 .. 24.0275 Hz  587.36c
0.692373   34.9777 .. 36.8104 Hz   88.41c  31.4526 .. 40.6162 Hz  442.65c
0.900000   56.2824 .. 58.6009 Hz   69.89c  51.7858 .. 63.3783 Hz  349.72c
```

At `u = 0.0687` a hundredth of uncertainty in the geometry is 434 cents of uncertainty in the pitch — a major third of ambiguity from an error most elicitation procedures would not even detect. At `u = 0.9` the same hundredth is 70 cents.

Inverted, the requirement is severe at the bottom:

```
du required for a 5-cent change in the reading

  at u = 0.018614      du = 0.00015816
  at u = 0.068735      du = 0.00023053
  at u = 0.300000      du = 0.00056449
  at u = 0.692373      du = 0.00113110
  at u = 1.000000      du = 0.00157533
```

A near-parallel dyad is not measurable to five cents by any elicitation procedure this framework knows of, and the honest response is to report the reading with its interval of uncertainty rather than to report a number to four decimal places. SPEC-6.4 requires exactly that. The relationship-frequency paper [Curzi, 2026] shows what a single elicited intensity does to a real reading, and the answer is: a great deal.

6.13 Comparing two dyads Two dyads compare as an interval, and intervals compose additively in cents:

```
interval( u1 -> u2 )  =  2400 * log2( (1 + 11 u2) / (1 + 11 u1) )
```

Verified on the carried case, decomposed through an intermediate point:

```
u = 0.068734665264634  ->  0.30                  3100.7448 cents
u = 0.30               ->  0.692372881355932     2406.4767 cents
sum                                              5507.2215 cents
direct  0.068734665264634 -> 0.692372881355932   5507.2215 cents
difference                                       9.09e-13 cents
```

The carried dyad's intervention is worth 5507.22 cents — four octaves and a major sixth. That is the number the relationship-frequency paper [Curzi, 2026] will convert into a sustainability verdict and the companion paper on tuning will price as work.

Differences in `u` do not carry this property in any useful sense. The two sub-steps above are `du = 0.231265` and `du = 0.392373`; the larger step in `u` is the smaller step in pitch. `u` is the argument, pitch is the reading, and only one of them composes the way a practitioner's intuition expects.

6.14 The reading a dyad frequency does and does not support Because this is the first section in the framework to produce a number about two people, it is the right place to fix what such a number licenses.

```
A RELATIONSHIP FREQUENCY LICENSES:
  - comparison with f_min at the same c/eta        (CCP-041, CCP-044)
  - comparison with the same pair at another time  (same rule set, same ladder)
  - comparison with the ceiling k N^2              (the realised pitch fraction)
  - a statement about how much distinct order the pair has available

IT DOES NOT LICENSE:
  - comparison with a pair read under a different ladder or a different N
  - any statement about either party as an individual
  - any statement about the pair's wellbeing, closeness, satisfaction or prospects
  - any clinical, diagnostic or therapeutic conclusion whatsoever
  - any action taken about a person without that person's ability to decline it
```

The fourth prohibition is not a disclaimer bolted to the end; it is the reason the third exists. A number that measures available order will be read as a number that measures the relationship, and the distance between those two things is not one this framework knows how to cross.

The Specification

SPEC-6.1 — Inputs

```
REQUIRED BEFORE ANY f_rel IS COMPUTED

[ ] the declared rule set: decomposition rules, N = 12, the ladder, dated   (CCP-041 Step 1)
[ ] two state vectors x, y over the twelve classes, entries >= 0
[ ] both raw L2 norms ||x||_2 and ||y||_2, in the instrument's units        (CCP-041)
[ ] confirmation that intensities were MAGNITUDE-ESTIMATED, not ranked      (CCP-041)
[ ] consent recorded from both parties, decline available at every step
[ ] the support overlap check, passed                                        (CCP-042)
```

A reading missing any one of these six is not a reading and is not published. The norms are on the list because normalisation discards them and they are half the measurement (the minimum-frequency paper [Curzi, 2026]'s magnitude reporting rule); the magnitude-estimation confirmation is on the list because without it the number is not determined at all (the relationship-frequency paper [Curzi, 2026]).

SPEC-6.2 — The computation

```
 1.  normalise:      x <- x/||x||,  y <- y/||y||
 2.  inner product:  cos theta = x . y
 3.  geometry:       u = 1 - (cos theta)^2
 4.  surplus:        F = ln(1 + (N-1) u)                nats
 5.  perspectives:   N_eff = 1 + (N-1) u
 6.  frequency:      f_rel = k * N_eff^2                Hz
 7.  interval:       cents above k = 2400 * log2(N_eff)
 8.  fraction:       realised pitch fraction = F / ln N
 9.  fold:           f_audible = f_rel * 2^j into 220-440 Hz; record j
10.  classes:        odd part of every state in the union of the supports
```

Steps 1–6 are Theorem 6.1. Steps 7–8 are §6.4. Step 9 is the octave-folding paper [Curzi, 2026]. Step 10 is not optional and the relationship-frequency paper [Curzi, 2026] says why: `u` cannot see pitch-class redundancy, and two dyads with identical `f_rel` can differ in class breadth.

SPEC-6.3 — The report line

Every dyad reading in this framework is published in this form:

```
DYAD READING
  rule set: ......................  ladder: 12-state, CCP-041, unaltered
  k = 0.483425732988414 Hz    N = 12    band 220-440 Hz
  x support / intensities / ||x|| = ..........
  y support / intensities / ||y|| = ..........
  cos theta = ..........   u = .......... (report 12+ dp; reporting standard)
  F = .......... nats      N_eff = ..........
  f_rel = .......... Hz    folded .......... Hz (j = ..)
  cents above k = ..........   realised pitch fraction = ..........%
  union of supports = {..}   distinct pitch classes = ..
  uncertainty: u +/- ....  ->  reading spans .......... cents
  CLAUSES: essence not resonance; not clinical, not counselling, not therapy;
           CONSTRUCTED case; scale status = [magnitude-estimated / NOT DETERMINED]
```

The uncertainty line is mandatory and §6.12 is why. A reading printed without it implies a precision the elicitation cannot deliver.

SPEC-6.4 — Precision

Recompute every derived value from `u` at 12 or more decimal places, never from a rounded print (the framework's reporting standard). Recomputing the carried dyad from a six-decimal `u` shifts `f_rel` in the fifth significant figure. Ratios, intervals and deficits are unaffected either way, which is why the interval columns are the robust ones and the hertz columns are not.

SPEC-6.5 — Gates

```
G1  OVERLAP.     supports share at least one state with non-trivial weight in both.
                 If not: u = 1 by construction; report DEGENERATE; publish no f_rel.
G2  MAGNITUDE.   intensities elicited by magnitude estimation with ratios asked
                 explicitly. If not: report the ORDER only; publish no f_rel.
G3  NORMS.       both raw L2 norms recorded and published. If not: no f_rel.
G4  CEILING.     if a sustainability verdict is to follow, c/eta <= ln N is checked
                 FIRST (CCP-041 SPEC 4.1). Above the ceiling no u closes the gap.
G5  CONSENT.     both parties may decline any step, at any time, without the reading
                 being completed against them. Consent is structural, not advisory.
```

G2 is the gate that fails most often and the one with no workaround. There is no "report it as ordinal" branch. The relationship-frequency paper [Curzi, 2026] gives the number that closes that door.

The Failure Modes

FM-6.1 — Reading `u` as a disagreement rate `u = 1 - (x.y)^2` is a geometric quantity computed from two non-negative intensity vectors. It is not a count of disagreements, not a conflict index, and not an inverse of rapport. Two people who argue constantly about one thing they both consider central are near-parallel and read low; two people who never argue because they attend to different things entirely are far from parallel and read high. The formalism is indifferent to the volume in the room.

The failure is common enough that the relationship-frequency paper [Curzi, 2026] builds a separate measurement route around exactly this confusion and the relationship-frequency paper [Curzi, 2026] prices what happens when someone tries to raise `u` by manufacturing argument.

FM-6.2 — Reading `f_rel` as a quality score **`f_rel` measures available order. It does not measure how good a relationship is, and the framework has no formalism that does.** A high-pitched dyad may be exhausting, unstable, or unwanted by both parties; a ground-essence dyad may be exactly what both parties want. The minimum-frequency paper [Curzi, 2026]'s classification answers *does this accumulate order at these costs*, which is a narrower question than any that matters to a person, and the narrowness is the point.

FM-6.3 — Comparing across ladders Two readings computed on different ladders, or at different `N`, are not comparable. `f_rel` depends on `N` explicitly through `(N-1)u`, and on the ladder implicitly through what the twelve slots mean. A dyad reading 20 Hz on a twelve-state ladder and a dyad reading 20 Hz on a seventeen-state ladder have nothing in common but the number. Publish the ladder with the frequency or publish neither.

FM-6.4 — Quoting hertz where the verdict is in cents §6.6 established that hertz and pitch rank the same increments in opposite orders. A report that leads with "we moved from 1.49 Hz to 2.70 Hz" has led with a difference of 1.21 Hz, which is small against the 69.61 Hz range and is in fact **1029.33 cents — nearly a ninth.** The hertz number understates the low end of the scale by exactly the amount the logarithm compresses it.

FM-6.5 — Precision theatre Reporting `f_rel = 1.4907987150 Hz` implies ten significant figures of confidence. **The tenth figure is real arithmetic and fictional measurement.** the framework's reporting standard requires the full-precision *computation* because rounding early corrupts the chain across thirty papers in the corpus; it does not license the full-precision *claim*. The honest report carries both: the computed value at full precision and the uncertainty interval at the precision the elicitation supports.

FM-6.6 — Folding a sweep Folding a monotone glide into 220–440 Hz produces eight discontinuities and a listener's impression of a rising-and-resetting figure rather than a rise. **The fold is correct for a point and wrong for a sweep**, and SPEC-6.1 fixes the transposition instead. A related error is band-shopping: choosing the band after seeing the numbers so that two readings land near each other. The octave-folding paper [Curzi, 2026]'s Failure Mode F8 owns that one.

FM-6.7 — Treating the eleven coincidences as meaning something §6.5 showed that `u = (m−1)/11` makes `f_rel` equal the essence frequency of state `m`. This is arithmetic. A dyad at `u = 6/11` is not "in" state 7 and has no relationship to reflexive feeling of any kind. The two quantities have different arguments and different derivations and happen to coincide at eleven rational points because the constants were chosen to make the two scales share a ceiling. Reading meaning into the coincidence is the numerological failure this series has warned against since the foundations.

FM-6.8 — The near-parallel reading reported without its uncertainty At `u = 0.0687`, an uncertainty of `±0.01` spans 434 cents. **A practitioner who reports the central value alone has reported a third of an octave of ambiguity as a point.** The failure is worse at the bottom of the range, which is exactly where near-parallel pairs — the ones most likely to need the reading — sit. The mandatory uncertainty line of SPEC-6.3 exists for this case and no other.

FM-6.9 — Assuming a high `f_rel` is achievable The glide runs to `u = 1`, which is perfect orthogonality: two people with no represented overlap in what they carry. The octave-folding paper [Curzi, 2026] and the relationship-frequency paper [Curzi, 2026] both establish that in practice such a reading almost always indicates disjoint support — an encoding failure — rather than a remarkable pair. A reading near the top of the glide is a reason to check the encoding, not to celebrate.

FM-6.10 — Using the reading on a person who did not agree to it A dyad reading requires two people's elicited states. Computing one from material a second party provided about a first — an account, a description, a recollection — produces a number about a representation the described person never saw and never agreed to. That is outside every specification in this section. The consent gate G5 is structural: it is a condition on whether the computation may be performed at all, not advice about how to present the result.

FM-6.11 — The reading used as counselling **A relationship frequency is not counselling, not therapy, not mediation, and not an assessment of any pair of people.** It is a number computed from a representation, and the representation is a description of a description (the minimum-frequency paper [Curzi, 2026]). A practitioner who presents it to two people as a statement about their relationship has changed what it is. Where a person's relationships are causing them difficulty, the appropriate resource is a qualified professional, and this framework is not one.

Worked Example

THE CARRIED DYAD, AS AN INTERVAL **CONSTRUCTED CASE, labelled as such. These are not two people.** This is the framework's carried dyad's dyad, phase 2a, read here purely as an interval; the relationship-frequency paper [Curzi, 2026] measures it, the relationship-frequency paper [Curzi, 2026] gives it a sustainability verdict, and the companion paper on tuning intervenes on it.

The vectors, exactly as the framework's founding specification fixes them:

```
A = { 4: 0.7,  6: 0.5,  9: 0.4 }
B = { 4: 0.6,  6: 0.55, 9: 0.3,  10: 0.2 }
```

A holds three states: situated feeling (4), temporal feeling (6), normative feeling (9). B holds those three and one more: narrative feeling (10). The supports overlap in three of four positions, so gate G1 passes comfortably.

Step 1 — the raw norms, recorded before normalisation (the minimum-frequency paper [Curzi, 2026]).

```
||A||_2 = sqrt(0.7^2 + 0.5^2 + 0.4^2)          = 0.9486832980505138
||B||_2 = sqrt(0.6^2 + 0.55^2 + 0.3^2 + 0.2^2) = 0.8902246907382427
ratio ||A|| / ||B|| = 1.065667
```

These are published because normalisation is about to throw them away and they are half the measurement. The two parties hold their states at very nearly the same total intensity — within 6.6% — so nothing in this reading is an artefact of one party simply reporting larger numbers than the other.

Steps 2–3 — the geometry.

```
cos theta = 0.965020899          (the framework's carried dyad prints 0.9650)
u         = 1 − cos^2 theta = 0.068734665264634
```

Step 4–6 — surplus, perspectives, pitch.

```
F     = ln(1 + 11 * 0.068734665264634) = 0.563084802746 nats
N_eff = 1.7560813179
f_rel = k * 1.7560813179^2             = 1.4907987150 Hz
```

Every figure matches the framework's carried dyad exactly, which is the point of recomputing from the fifteen-decimal `u` rather than from the printed 0.068735.

Step 7–8 — the interval.

```
cents above k              = 2400 * log2(1.7560813179) = 1949.66
octaves above k            = 1.624719305043341
realised pitch fraction    = 0.563084802746 / 2.484906649788 = 22.66%
```

The pair occupies 1949.66 cents of an available 8603.91 — one octave and a major seventh above the ground essence, and 22.66% of everything a twelve-category system has to offer. Put the other way: **77.34% of the available relational range is unoccupied.**

For a reader who prefers the perspective count: `N_eff = 1.7561`. **Two people, and the geometry says they are worth 1.76 genuinely distinct perspectives.** The minimum-frequency paper [Curzi, 2026]'s fourth column is the comparison that matters — at `c/eta = 2.0667` this pair must field `e^2.0667 = 7.90` distinct perspectives to break even, and it fields 1.76. The relationship-frequency paper [Curzi, 2026] finishes that reading.

Step 9 — the fold.

```
1.4907987150 Hz * 2^8 = 381.6444710 Hz      j = 8, in 220–440 Hz
```

Step 10 — the classes.

```
union of supports = {4, 6, 9, 10}
odd parts         =  1, 3, 9,  5      ->  four distinct pitch classes
```

Four states, four distinct classes: no collision. The relationship-frequency paper [Curzi, 2026] shows the counterfactual in which the same `u` comes with only three classes, and why that check has to be run separately.

The report line.

```
DYAD READING -- CONSTRUCTED CASE, the framework's carried dyad phase 2a
  k = 0.483425732988414 Hz   N = 12   ladder: 12-state, unaltered   band 220-440 Hz
  A: {4:0.7, 6:0.5, 9:0.4}          ||A|| = 0.9486832980505138
  B: {4:0.6, 6:0.55, 9:0.3, 10:0.2} ||B|| = 0.8902246907382427
  cos theta = 0.965020899    u = 0.068734665264634
  F = 0.563084802746 nats    N_eff = 1.7560813179
  f_rel = 1.4907987150 Hz    folded 381.644 Hz (j = 8)
  cents above k = 1949.66    realised pitch fraction = 22.66%
  union {4,6,9,10}           distinct pitch classes = 4
  uncertainty: u +/- 0.01 -> reading spans 434.34 cents
  CLAUSES: essence not resonance; not clinical, not counselling, not therapy;
           CONSTRUCTED; scale status = magnitude-estimated per CCP-041
```

THE SAME DYAD, AFTER — the interval crossed The framework's carried dyad's phase 2b changes the geometry: B takes on two states A does not hold. The minimum-frequency paper [Curzi, 2026] gives one realising vector, `B' = {2:0.95, 4:0.60, 6:0.55, 9:0.30, 10:0.20, 12:0.8391}`, with `cos theta = 0.554643` and `u = 0.692372881355932`. This section does not re-derive it and does not own the intervention; it reads the interval.

```
PHASE 2a    u = 0.068734665264634   f =  1.4907987150 Hz   1949.66 cents above k
PHASE 2b    u = 0.692372881355932   f = 35.8881768941 Hz   7456.88 cents above k

THE INTERVAL CROSSED:  5507.2215 cents  =  4.589351 octaves
realised pitch fraction:  22.66%  ->  86.67%
N_eff:                    1.7561  ->  8.6161
```

Four octaves and a major sixth, purchased by one party taking on two states the other does not hold. In perspective terms the pair went from being worth 1.76 distinct views to being worth 8.62 of a possible 12.

And the caution that has to travel with that sentence, because it is the most quotable number in the section: the interval was not purchased by disagreement, by opposition, or by either party becoming difficult. It was purchased by a change in **what B is actually doing** — a genuine addition to B's represented state, not a rhetorical stance. The relationship-frequency paper [Curzi, 2026] proves that the rhetorical version returns nothing at all, and the minimum-frequency paper [Curzi, 2026] already flagged it: **CC-R8 forbids the opposition route from working, because perfect opposition reads as `u = 0` exactly as agreement does.**

The relationship-frequency paper [Curzi, 2026] turns these two positions into sustainability verdicts. The frequencies alone say nothing about whether either arrangement lasts; that requires `f_min`, and `f_min` requires `c` and `eta`, and the minimum-frequency paper [Curzi, 2026] owns both.

The counter-case, so the section is not lying by selection **CONSTRUCTED.** Not every reading moves. Take the same A, and a B whose extra state carries a smaller intensity:

```
B[10] = 0.00   ->  u = 0.019232   f = 0.7096 Hz    -1285.20 cents from the carried reading
B[10] = 0.10   ->  u = 0.032095   f = 0.8850 Hz    - 902.77 cents
B[10] = 0.20   ->  u = 0.068735   f = 1.4908 Hz         0.00  (the carried case)
B[10] = 0.30   ->  u = 0.124003   f = 2.7017 Hz    +1029.33 cents
B[10] = 0.40   ->  u = 0.191202   f = 4.6554 Hz    +1971.38 cents
```

One elicited intensity, moved by one step of a ten-point instrument, moves the reading by more than a thousand cents. That is the steepness of §6.9 arriving as a measurement problem rather than as a price list, and it is the strongest available argument for the minimum-frequency paper [Curzi, 2026]'s insistence that intensities be magnitude-estimated with the ratio question asked explicitly, and for Section 5's double-coding requirement. A reading this sensitive to one number is a reading that must be produced twice, by two analysts, from the same written framings, or not published.

What this section established **Established, and checkable:** a dyad has exactly one frequency, `f_rel = k(1+(N−1)u)^2`, obtained by substituting the surplus functional [Curzi, 2026] into the essence bridge of Series B [Curzi, 2026k]; it is strictly increasing in `u`, bounded by `k` below and `k N^2` above, and spans 8603.9145 cents at `N = 12`; it equals the essence frequency of state `m` exactly at `u = (m−1)/11`; the first 0.05 of `u` buys 9.34 times the pitch of the last 0.05, while the last 0.05 buys 9.20 times the hertz of the first; and half the entire range is spent by `u = 0.224009`.

Not established, and not claimed: that any pair of people physically resonates at any frequency here; that a higher reading is a better relationship; that the reading is clinical, diagnostic, therapeutic, or a substitute for counselling of any kind; that the number can be computed about someone who has not agreed to it; or that a frequency alone says whether a relation lasts, which needs `f_min` and belongs to the relationship-frequency paper [Curzi, 2026].

The one line to carry forward: *two people are one interval, measured upward from the note a pair with nothing between them would sound.*

Section 6 owns `f_rel = k(1+(N−1)u)^2`, the glide, the realised pitch fraction, and the fixed-transposition sweep spec. Later papers in this series cite it. The surplus law is the sustaining-frequency paper [Curzi, 2026]; the essence bridge is the sustaining-frequency paper [Curzi, 2026]; `f_min`, the octave deficit and the four levers are the minimum-frequency paper [Curzi, 2026]; elicitation and the magnitude reporting rule are the minimum-frequency paper [Curzi, 2026]; the chord and the overlap condition are The octave-folding paper [Curzi, 2026]; the consonance/redundancy identity is Section 4. The relationship-frequency paper [Curzi, 2026] measures `u` between two people, and the relationship-frequency paper [Curzi, 2026] goes to the bottom of the glide and finds two different faults sounding the same note.

7 Results

The seven results below are the units other papers in this corpus cite. **DERIVED** means it follows from the definitions by algebra; **ENUMERATED** means it was checked over the complete finite lattice rather than sampled; **COMPUTED** means exact arithmetic on stated constructed parameters; **POSTULATED** means asserted here with no external support.

R1 — the doubling theorem. For any index pair reducing to p/q, the essence interval p^2/q^2 has exactly twice the Tenney harmonic distance of p/q, and Euler gradus $2G(pq) - 1$: $\mathrm{HD}_{\mathrm{ess}}(i, j) = 2\log_2(pq)$, $G_{\mathrm{ess}} = 2G(pq) - 1$. *Status:* DERIVED in three lines; ENUMERATED over all sixty-six pairs at $N = 12$, both identities exact in every case. The relational-lattice form is established in Series B [Curzi, 2026l], whose Result Two derives a room from essence indices 7, 9 and 11 with the folded set 233.98–386.74 Hz entirely above a Schroeder frequency of 146.7 Hz — the one place in the corpus where these frequencies meet real acoustics.

R2 — the consonance–redundancy identity. The minimum $\mathrm{HD}_{\mathrm{ess}} = 2.0000$ over all sixty-six pairs is attained by exactly six pairs, ${1{:}2, 2{:}4, 3{:}6, 4{:}8, 5{:}10, 6{:}12}$, every one of them a pitch-class collision, and by no non-collision. Separately, folded separation is exactly zero **if and only if** the two indices share an odd part — a biconditional, true for all ten collision pairs and false for all fifty-six others. *Status:* ENUMERATED. **The honest qualification travels with it:** the ordering is not total. $1{:}3$ at 3.1699 is more consonant than the collisions $1{:}4$, $2{:}8$,

$3{:}12$ (4.0000) and $1{:}8$ (6.0000). Collisions are dramatically more consonant *on average* — mean 3.0000 over ten against 8.6617 over fifty-six — and the ranges overlap.

R3 — the firewall. F_{int} is invariant under relabelling of the states, so it does not determine the chord. Relabelling 10 to 9 in the carried constructed subject leaves $F_{\text{int}} = 2.164390$ unchanged to 10^{-12} while removing the collision and taking the class set from $\{3, 5, 7\}$ to $\{3, 5, 7, 9\}$; a further relabelling to $\{2, 8, 9, 11\}$ again gives exactly 2.164390, with classes $\{1, 9, 11\}$ and the collision $(2, 8)$. *Status:* DERIVED (inner products are invariant under a permutation of coordinates); COMPUTED on a constructed subject. **This is the result the paper exists to establish** and it blocks any inference from the width statistic to distress or health.

R4 — the restatement collapse. Replacing one of three framings with a paraphrase of another takes F_{int} from 2.164390 to 1.333662 nats and the chord width from 36.6686 to 6.9620 Hz: a loss of 0.830727 nats, 2.3970 octaves, more than a third of everything a twelve-category ladder can supply. *Status:* COMPUTED on a constructed extension.

R5 — the extremal chords. Over all $\binom{12}{4} = 495$ four-state chords: the most consonant is $\{1, 2, 4, 8\}$, mean $\text{HD}_{\text{ess}} = 3.3333$, six collisions from six pairs, **one** pitch class; the most dissonant is $\{7, 9, 10, 11\}$, mean 12.7586, zero collisions, four classes. 159 of the 495 chords (32.1 per cent) have no collision at all. *Status:* ENUMERATED. The most dissonant chord in the system — reflexive, normative, narrative and reciprocal feeling — is an entirely ordinary description of a person attending to a social situation, which is the strongest available argument against reading harmonic dissonance as psychological difficulty.

R6 — convexity against concavity. $f_{\text{rel}}(u) = k(1+11u)^2$ is strictly convex in u while the same quantity in cents is strictly concave: $\mathrm{d}f/\mathrm{d}u$ rises from 10.6354 to 127.6244 Hz per unit u across $u \in [0, 1]$, while $\mathrm{d}(\text{cents})/\mathrm{d}u$ falls from 38087 to 3174. The first 0.05 of u is worth $+0.6780$ Hz and $+1517.44$ cents; the last 0.05 is worth $+6.2350$ Hz and $+162.45$ cents — the last step is $9.20\times$ the first in hertz and $\frac{1}{9.34}$ of it in cents. *Status:* DERIVED and COMPUTED. **Pitch is the unit that enters the verdict; hertz differences are decorative**, because the sustainability criterion is computed in logarithmic units [Curzi, 2026].

R7 — the eleven coincidences. $f_{\text{rel}}(u) = km^2$ exactly at $u = (m-1)/(N-1)$, so a dyad's interval lands on a state's note at eleven exact values of u. *Status:* DERIVED. **The caveat travels with the table and is not optional:** this is a numerical coincidence in the chosen constants, not a claim that a dyad at $u = 7/11$ is in any sense *in* state 8. The two quantities have different arguments, different derivations and different scale status.

CR — counter-result. The width statistic is invariant under relabelling and therefore does not even determine the chord it is computed from, which formally blocks any inference from low or high internal surplus to distress or health. **A quantity compatible with every value of a variable carries no information about that variable.** The corpus reaches structurally similar conclusions elsewhere: the property that makes an insurance reserve resistant to third-party claims is the same property that makes it unavailable to its owner [Curzi, 2026], leaving 7.58 per cent realisable within a year; and the uncertainty node declines the inflated reading of its own borrowed physics, requiring that a theorem, an ordinary observer effect and a framework conjecture be distinguishable in three separate sentences on demand [Curzi, 2026]. **This paper is under the same requirement and meets it in §2 and §9.**

8 Discussion

8.1 What follows

Consonance is scarce in this lattice by a factor of exactly two, and the scarcity has a shape. Six pairs sit at 2.0000, then a gap to 3.1699, then a long tail to 14.0888 with fifty-six of sixty-six pairs above 3.1699. That is not a gradient; it is a small cluster of redundancies and a large spread of remoteness. The middle ground the metaphor promises — framings that harmonise without repeating — is not a region of this lattice at all. **Whether that is a discovery about selves or an artefact of the square law is exactly the question this paper is obliged to leave open, and it**

leaves it open: the doubling theorem is a theorem about ratios, and the square law under it is a postulate.

A collision costs one thing and it is worth naming precisely. It costs a position in a harmonic vocabulary of which $N = 12$ supplies only $\lceil 12/2 \rceil = 6$. It does **not** cost surplus — the relabelling test shows a collision appearing and disappearing with F_{int} fixed to twelve decimal places. It does **not** mean two states are the same component: attributed feeling and narrative feeling differ by five components of specification and collide only because $10 = 2 \times 5$ and the fold divides by powers of two. It does **not** mean the raw chord has fewer notes: states 5 and 10 are 12.0856 and 48.3426 Hz, two octaves apart, and only the fold makes them one. **A collision is a limit on the folded representation, not a finding about the subject.**

The only lever on the width is the angular diversity of what is elicited. Amplifying a framing does nothing — scale invariance is exact. Adding a restatement is *negative*, at -2.3970 octaves. Adding a framing with disjoint support inflates the number and destroys its meaning. And distress is not an input: it appears in no equation here, has no coefficient, and has no path by which it could enter the computation.

8.2 What does not follow

Nothing about distress, health, function or well-being. §4.6 gives the formal argument and §7 restates it as CR. The four quadrants — high width with and without distress, low width with and without distress — are all occupiable and the formalism forbids ruling any of them out. **The framework does not recommend that anyone acquire, retain or cultivate suffering**, and suffering is not a resource in it.

Nothing about a person. A reading is of a set of framings on one situation on one occasion under one decomposition by one analyst. Cross-person comparison is not licensed, because the reproducibility gate has not been run and the specific-objectivity property that would license it [Rasch, 1960] has not been established.

And no claim of a characteristic scale in the world. Human activity is bursty, with heavy-tailed inter-event times and **no characteristic time scale** [Barabási, 2005]; strong spatial regularity coexists with that temporal burstiness [González et al., 2008]. A frequency in this paper is a position on a logarithmic surplus scale, not a rate at which anything occurs, and the burstiness literature is the reason that distinction has to be defended rather than assumed.

8.3 The consonance measures against the measured literature

Where the external literature disagrees with this framework, it is cited and the disagreement is stated. Ratio complexity is one family of consonance model and roughness is another; they make different predictions, and the roughness family is the one with the psychoacoustic measurements behind it [von Helmholtz, 1885, Plomp and Levelt, 1965, Kameoka and Kuriyagawa, 1969, Terhardt, 1974]. Consonance as heard depends on the **spectrum of the sound**, so a ratio is a fact about harmonic-spectrum instruments and not a universal constant [Sethares, 2005]. Interval categories are learned and bounded [Burns, 1999, Burns and Ward, 1978], octave equivalence generalises less far than is assumed [Hoeschele et al., 2012], and the octave's perceptual specialness has a measured and finite size [Demany et al., 2013]. The just-intonation lattice whose geometry this paper borrows is real and documented [Partch, 1974, Doty, 2002, Benson, 2006, Lindley and Turner-Smith, 1993]; **what is not documented anywhere is the claimed function** — that positions on such a lattice convert between psychological and acoustic quantities. That is the part of this framework furthest from any external literature, and it is where the reader should be most sceptical.

8.4 The functional form is underdetermined

The surplus functional's shape is not settled by the evidence. Influence has been argued to grow as a **power** of the number of sources with exponent below one [Latané, 1981]; network value has been argued to grow as $n \log n$ [Briscoe et al., 2006]; this framework asserts $\ln(1 + (N-1)u)$. All three decelerate and all three fit the qualitative facts, and many mechanisms produce curves that data cannot distinguish [Stumpf and Porter, 2012]. **Nothing in this paper discriminates among them.** What *can* be said is which warrant the logarithm needs: not the psychophysical one, which requires the perceptual story to be true, but the multiplicative-dynamics one, in which the growth rate of a multiplicative process is logarithmic in the accumulated quantity with no psychological assumption at all [Peters and Gell-Mann, 2016]. The psychophysical warrant is used in this paper only for the elicitation of intensities, and nowhere else.

8.5 What would distinguish this framework from a neighbouring model

The sharpest question a reader can ask of work of this kind is what observation would separate its laws from a model that fits equally well [Sobkowicz, 2020, Durlauf, 2005]. **For this paper the honest answer is: almost nothing, by construction.** Sections 4 and 6 are enumerations over a finite lattice and are true or false by computation, not by observation. Section 5 is a procedure. The single empirically vulnerable claim in the series is the ceiling [Curzi, 2026], and it is inherited rather than established here. An explanation is expected to name a mechanism rather than exhibit a fitted curve [Hedström and Ylikoski, 2010]; this paper fits nothing, which avoids that charge at the cost of having no empirical content to fit. Prediction is one of many legitimate reasons to model [Epstein, 2008], and the reason claimed here is **bounding**: the paper's principal contribution is a proof that a tempting inference is unavailable. The analytical position — mechanisms stated, parameters declared, the arrangement named in advance — is the one this work is best placed to defend [Keuschnigg et al., 2018, Cartwright, 1999].

9 Verification

Every result is decomposed into steps a reader can check individually. **A postulate that can be dropped without a result changing has been dropped**, and the section says outright where the chain stops being a proof.

- **ARITHMETIC** — reproducible from the inputs named.
- **ENUMERATION** — exhaustive over a finite set; the set is named and its size given.
- **DEFINITION** — true by stipulation.
- **CITED** — supported by a named external source.
- **POSTULATED** — asserted here, with no external support.

9.1 The postulates

	postulate	status
P1	The Pythagorean whole tone is the ratio $9{:}8$.	CITED [Barbour, 1951]
P2	φ is the positive root of $x^2 = x + 1$.	CITED [Euclid, 1956, Vajda, 1989]
P3	The dimensionless number $(9/8)^2/\varphi^2$ **is a frequency in hertz**.	POSTULATED — no source

	postulate	status
P4	The essence of an n-component state is kn^2.	POSTULATED — no source
P5	$\mathrm{HD}(p/q) = \log_2(pq)$ measures harmonic remoteness.	DEFINITION (a chosen function)
P6	$G(n) = 1 + \sum e_i(p_i - 1)$ measures smoothness.	DEFINITION (a chosen function)
P7	A frequency may be divided by a power of two without changing what it means.	POSTULATED beyond pitch; CITED for pitch only [Shepard, 1982]
P8	$N_{\text{eff}} = 1 + (N-1)u$, $F = \ln N_{\text{eff}}$.	POSTULATED — design axiom [Curzi, 2026]
P9	Intensities are non-negative.	DEFINITION of the encoding

Dropped as unnecessary. No postulate about *which* twelve states, or about their order, is needed for R1, R2, R3 or R5: the doubling theorem, the consonance–collision identity, the relabelling invariance and the extremal enumeration hold for the indices $1..12$ whatever the states are called — which is precisely why R3 is true. No postulate about consonance is needed for R3, R4, R6 or R7. **P3 and P4 are needed only where a number is quoted in hertz**, and every dimensionless result in the paper survives their rejection.

9.2 The compounding chain

step	consumes	produces	status
L1	P1	$(9/8)^2 = 81/64 = 1.265625$	ARITHMETIC
L2	P2	$\varphi^2 = 2.618033988749895$	ARITHMETIC
L3	L1, L2	$k = 0.483425732988414$	ARITHMETIC (exact, defined)
L4	L3, **P3**	k carries the unit Hz	**POSTULATED**
L5	L4, **P4**	$f(n) = kn^2$ Hz	**POSTULATED**
L6	L5	essence ratio of i, j is $j^2/i^2 = p^2/q^2$	ARITHMETIC
L7	L6, P5	$\mathrm{HD}_{\text{ess}} = \log_2(p^2q^2) = 2\log_2(pq)$ — **R1**	ARITHMETIC
L8	L6, P6	$G_{\text{ess}} = 2G(pq) - 1$ — **R1**	ARITHMETIC
L9	L7 over all 66 pairs	minimum 2.0000 at exactly six pairs — **R2**	ENUMERATION (66)
L10	L6, P7	folded separation $= 0$ iff equal odd parts — **R2**	ARITHMETIC
L11	L9, L10	the minimum set **is** the collision set	ARITHMETIC
L12	L7 over all 495 four-chords	extremal chords — **R5**	ENUMERATION (495)

step	consumes	produces	status
L13	P8, P9	$F_{\text{int}} = \ln(1 + 11\bar{u})$ from inner products	DEFINITION
L14	L13, permutation invariance of inner products	F_{int} invariant under relabelling — **R3**	ARITHMETIC
L15	L14, L10	equal F_{int}, different collision sets	ARITHMETIC
L16	L13 on a constructed substitution	$2.164390 \to 1.333662$; $36.6686 \to 6.9620$ Hz — **R4**	ARITHMETIC
L17	L5, P8	$f_{\text{rel}}(u) = k(1 + 11u)^2$	ARITHMETIC given L5
L18	L17	$f'' > 0$; cents$'' < 0$ — **R6**	ARITHMETIC
L19	L17 set equal to km^2	$u = (m - 1)/11$ — **R7**	ARITHMETIC

9.3 What this chain is and is not

Read the status column. L4, L5 and P7 are postulated; P5, P6, P8, P9 and L13 are definitions. **A chain whose load-bearing steps are definitions and postulates is not a proof of its interpretation**, and this one is not offered as one. What it *is* is a clean separation.

Everything dimensionless survives the rejection of P3 and P4. R1, R2, R3, R5 and R7 are statements about integers and their ratios; they do not become false if nobody accepts that $(9/8)^2/\varphi^2$ is a frequency. R6's *shape* survives; only its hertz labels do not. **The firewall — the paper's principal result — is the most robust thing in it**, because it follows from the invariance of an inner product under a permutation of coordinates and from nothing else.

Conversely, nothing in the chain licenses the interpretation. The two consonance measures are chosen functions and are labelled so at every use; the octave fold beyond pitch is postulated; the essence law is postulated and its exponent is **defined rather than estimated**, with no fitting procedure run [Clauset et al., 2009]. **A reviewer looking for the weakest link should look at L4 and L5, and this section is written so that the search takes under a minute.** That the badges are printed rather than discovered is the honest description of what a badged postulate does inside a research programme [Lakatos, 1970, Platt, 1964, Meehl, 1990].

10 Limitations

Every check the governing rigour standard specifies is listed by name, passed or not.

L-1. The dimensional audit does not close for any quantity in hertz. k is computed as a ratio of two dimensionless numbers, and no operational procedure is supplied by which it acquires $[T^{-1}]$. Everything expressed as a ratio, interval, distance or class is unaffected.

L-2. The two consonance measures are chosen, not measured. Neither Tenney's harmonic distance [Tenney, 2015] nor Euler's gradus [Euler, 1739] is an experiment, and the measured psychoacoustic literature runs on roughness and spectrum [Plomp and Levelt, 1965, Sethares, 2005], which does not agree with ratio complexity in general. **No claim in this paper is a claim about what a listener finds pleasant.**

L-3. The attribution of harmonic distance was incomplete, and completing it in the obvious way gives the wrong book. This entry replaces an earlier one recording that neither primary source was present in the bibliographic database governing this corpus. Both are now present and are cited at every use: the *gradus suavitatis* from Euler's *Tentamen novae theoriae musicae* [Euler, 1739], and harmonic distance from Tenney's 1983 essay "John Cage and the Theory

of Harmony", cited from the collected edition [Tenney, 2015]. The source manuscript named harmonic distance by its author's surname alone, with no year and no containing work. That is an incomplete attribution rather than a false one, but it is not harmless: the Tenney title a reader is likeliest to supply is *A History of 'Consonance' and 'Dissonance'* (1988), which is a historical survey and does not define the measure. §4.2 now states the correct source and the correction explicitly. Two bibliographic details remain open and are stated rather than guessed: no page range is given for the essay within *From Scratch*, because none could be confirmed from any catalogue record consulted; and the frequently repeated original printing in *Soundings* 13 (1984) was likewise not confirmed and is therefore not asserted here. The definitions used in this paper are given in full in §4.2, so the arithmetic remains checkable without either source.

L-4. Octave reduction beyond pitch is unwarranted. Shepard grounds octave equivalence for pitch [Shepard, 1982]; nothing grounds folding a non-acoustic quantity into a band and reading meaning from where it lands.

L-5. The exponent is defined, not estimated. No scaling exponent here was fitted, no alternative was tested, and the standard procedure [Clauset et al., 2009] was not run because there is no data. This paper makes no empirical scaling claim.

L-6. The functional form is underdetermined by the available evidence [Latané, 1981, Briscoe et al., 2006, Stumpf and Porter, 2012], and nothing here discriminates among the candidates.

L-7. Conjoint measurement was not tested, so the ratio-scale treatment of the intensities is provisional and Michell's objection [Michell, 1997] stands unanswered for every quantity except the indices themselves; the available remedies [Luce and Tukey, 1964, Rasch, 1960, Krantz et al., 1971] were not run. And the exemption claimed for the indices is narrower than it looks. The count argument establishes that a cardinality is absolute-scaled, which is a fact about counting and not about whether an emotional state was correctly reduced to a count in the first place. The affect literature that the index's own rungs are named from describes emotional states by patterns across simultaneously varying appraisal dimensions [Smith and Ellsworth, 1985], by a process in time [Scherer, 2001], and by qualitative relational themes [Lazarus, 1991] — none of which a scalar can carry. Michell's objection and that one are the same objection in two domains, and it is the framework's habit of scalar reduction that both are about. The index is inherited here rather than argued, and the concession is made where the ladder is built [Curzi, 2026]; it is recorded in this paper because every reading produced by §5 rests on that index.

L-8. The reproducibility gate has not been run. FVP-7 for the index [Curzi, 2026] and for the self-storage parameters of Series B are both recorded as NOT RUN. Cross-person comparison of any number in this paper is therefore not licensed.

L-9. Every case is constructed. The subject, its extensions and the carried dyad are chosen parameter sets. Their exactness follows from the parameters being exact and does not transfer to a measured case.

L-10. The encoding cannot represent opposition. Because intensities are non-negative, $\cos\theta \geq 0$ for every pair of framing vectors and $u = 0$ can arise only from parallelism. An internal framing that is the exact negation of another has no representation here. This is a real limitation of the encoding and a signed encoding would be a different instrument.

L-11. The relabelling result is an arithmetic fact, and H1 is not tested. The invariance is proved; the empirical prediction that no association exists between collision structure and distress at fixed width is untested, and no sample exists.

L-12. The reflexive specification is not identified [Manski, 1993, Shalizi and Thomas, 2011]. The remedy adopted is stipulation of the reference set [Brock and Durlauf, 2001], and the claim is a modelling position rather than an identified result. The framework also exempts its own parameters from its own reflexivity [Lucas, 1976], which it cannot consistently do.

L-13. A published reading can make itself true. Expectancy and framing effects are the reason the separation statement of SPEC-4.3 is mandatory; the effect is real but, on the best long-run evidence, small and often dissipating [Jussim and Harber, 2005].

L-14. No statistical claim is made and none is supported. There are no p-values, effect sizes, intervals or samples in this paper. No model comparison, out-of-sample test, matched null or information criterion is reported, because nothing was fitted.

L-15. There is no conservation law. The storage quantity is written with the grammar of a conserved quantity without one [Mirowski, 1989]; validity is claimed only within a specified arrangement [Cartwright, 1999] and no universality is claimed.

L-16. The aggregate does not determine the joint structure, and this paper's own central result is an instance of that [Kirman, 1992, Harrison and Klein, 2007]. Any reader tempted to invert a width statistic into a chord should read §4.6 again.

L-17. Boundary behaviour. $\bar{u} \to 0$ gives $F_{\text{int}} = 0$ and the ground essence, which the paper examines; $\bar{u} \to 1$ gives the ceiling kN^2, which it examines; **disjoint support** gives a high $\bar{u}$ that is a null result rather than a wide chord, which the paper examines and refuses to publish as a reading. What is **not** examined is the behaviour of the width statistic as the number of framings grows, and the mean pairwise $\bar{u}$ has no stated correction for the number of pairs.

L-18. This paper asserts perceptual claims with no perceptual experiment behind them. That consonance measures are meaningful over an index lattice, that octave equivalence is usable, and that cents are the right unit are all asserted on the strength of the cited literature and its boundaries, and not on any measurement made here.

11 Conclusion

What is established. Squaring an index ratio doubles Tenney harmonic distance exactly and maps Euler gradus G to $2G-1$; both identities hold over all sixty-six index pairs at $N = 12$. The global consonance minimum is attained by exactly the six pitch-class collisions and by nothing else, and zero folded separation holds if and only if two indices share an odd part — a biconditional over the complete lattice, with the honest qualification that the resulting ordering is not total. Over all 495 four-state chords the most consonant is a chord of one pitch class and the most dissonant is an ordinary description of a person in a social situation. **The internal surplus is invariant under relabelling of the states and therefore does not determine the chord it is computed from**, so no inference from it to distress, health or function is available; the inference fails at the level of the statistic, not at the level of taste. Replacing a framing with a paraphrase of another costs 2.3970 octaves of width. At dyad scope the relationship frequency is convex in u in hertz and concave in cents, and only the logarithmic unit enters the sustainability verdict. **All of this is arithmetic and enumeration on stated definitions, and none of it depends on the value of k.**

What is proposed. That a dimensionless ratio of a tuning interval to the square of the golden section is a frequency; that the essence of an n-component state is kn^2; that octave equivalence may be carried from pitch to a quantity that is not a pitch; and that Tenney's and Euler's measures are the right functions to impose on this lattice. **The first three are postulates with no external support and the fourth is a modelling choice**, all four are labelled at every appearance, and all four are separable from the results above: a reader who rejects them keeps the doubling theorem, the identity, the firewall, the enumeration and the convexity result, expressed in ratios rather than in hertz.

What is not claimed at all. No person, tissue or molecule is claimed to resonate at any frequency named here. **The twelve state classes are not diagnoses**, a chord reading is not an assessment, and nothing in this paper is clinical, diagnostic, therapeutic or prognostic. Nothing in it recommends, values, or claims informational content for suffering.

12 Declarations

Data and code availability. No dataset was collected for this study and none exists. The sixty-six pair values and the 495 four-state chords are exhaustive enumerations over a finite lattice, computable from the definitions in §4.2 by any reader; every worked case is a constructed

parameter set, labelled as constructed at the point of use. The figures are generated by the Python scripts named in their captions, which take their inputs from the tables in the text and are available from the author on request. **No measurement of any human being underlies any figure, table or number in this paper.** A paper about consonance and internal states reads as though someone had been listened to or asked something; nobody was.

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the analysis, the figures and the manuscript.

Ethics. No human-subjects research was conducted and no personal data were processed. The constructed subject, its labelled extensions and the carried dyad are invented parameter sets and do not describe any person. **This paper asserts perceptual claims — that ratio-complexity measures are meaningful, that octave equivalence is usable beyond pitch, that cents are the operative unit — without any perceptual experiment behind them**; those claims rest entirely on the cited literature, are bounded by that literature's own findings (§8.3), and are listed among the limitations (L-2, L-4, L-18).

Clinical disclaimer. The twelve state classes are not diagnostic categories, and the internal surplus is not a measure of distress, well-being, health or function — §4.6 proves that it contains no term that could carry such information. Nothing in this paper is an assessment, a screening instrument, a treatment, or advice about any person's care, and nothing in it may be presented as such. A reader who wants to know whether something is wrong should ask someone qualified to answer that; this instrument is not qualified and does not attempt it.

Chapter 44

Three Routes to the Relationship Frequency: The Unison Result and the Elasticity-Reserve Split Between Levers

Abstract. We give three independent estimators of dyadic non-parallelism and show what the resulting frequency can and cannot support. Route A computes $u = 1 - \cos^2\theta$ from intensity vectors and is non-negative by construction; Route B derives the exact agreement identity $\cos\theta = 2a - 1$, $u = 4a(1-a)$ from binary judgments over m items, with resolution $4/m$; Route C uses $u = 1 - \rho^2$ on error vectors against known answers. On a constructed dyad the three return $u = 0.068735$, 0.065556 and 0.068775, giving 1.4908, 1.4320 and 1.4916 Hz — a spread of 70.52 cents, of which 69.64 is the agreement channel's own quantisation. The central negative result is that every quantity depends on the correlation only through its absolute value: $\rho = +1$ and $\rho = -1$ both give $u = 0$, $F = 0$ and $f_{\text{rel}} = k = 0.4834$ Hz, with $k = (9/8)^2/\varphi^2 = 0.483425732988414$ a *defined* constant whose interpretation as a frequency is postulated. The agreement table is a perfect mirror about $a = 0.5$, so a pair agreeing on every item and a pair agreeing on none produce the identical reading, and recruiting a contrarian buys nothing while adding links. At $u = 0$ the fixed point collapses to $-c/\delta$ with η and N absent and two of four levers exactly inert. On lever choice, elasticities are dimensionless and commensurable: cost-side levers have elasticity $2.885390\,(c/\eta)$ against a geometry-lever bound of 2.885390 octaves per e-fold, but the reserves invert the ranking — raising η to a physically impossible 1.00 buys 4.174198 octaves against a requirement of 4.338420, while raising u buys 5.545206 and closes it.

Keywords: relationship frequency; agreement identity; unison; opposition; elasticity; lever reserve; concordance; order of operations

Series: C. Frequency, Essence, and Mind **Paper:** CCP-044

1 Introduction

The corpus's surplus functional is built on a single geometric quantity, the represented non-parallelism $u = 1-(x\cdot y)^2$ between two unit state vectors. This paper does three things with it: it gives three independent estimators of u between two parties, it establishes what the resulting frequency can and cannot support, and it derives the consequence that most damages the ordinary intuitions people bring to the object.

The central negative result is that u depends on the correlation only through its absolute value. Because $u = 1 - \rho^2$ is even in ρ, perfect opposition and perfect agreement are the *same reading*: $\rho = +1$ and $\rho = -1$ both give $u = 0$, $F = 0$, $N_{\text{eff}} = 1$ and $f_{\text{rel}} = k = 0.4834$ Hz, folded to the same 247.514 Hz at $j = 9$ in the same pitch class. The formalism does not merely fail to distinguish the contrarian from the sycophant; it proves them indistinguishable, and the agreement channel exhibits the same fact directly, since $u = 4a(1 - a)$ is a perfect mirror about an agreement fraction of $a = 0.5$. **Recruiting a contrarian therefore adds cost and returns exactly nothing**, and at $u = 0$ the fixed point collapses to $-c/\delta$, in which the conversion efficiency and the category count do not appear at all and two of the framework's four levers have derivative exactly zero.

The three estimators are not interchangeable and the paper never pretends they are. Route A reads the state vectors directly; because intensities are non-negative by construction its cosine is bounded below by zero, **so the negative half of the correlation range is unreachable through it.** Route B derives an exact identity from binary judgments over m shared items, $\cos\theta = 2a - 1$ and $u = 4a(1 - a)$, with a resolution bounded by $4/m$. Route C uses $u = 1 - \rho^2$ on error vectors against known answers. On the constructed carried dyad the three return $u = 0.068735$, 0.065556 and 0.068775, giving 1.4908, 1.4320 and 1.4916 Hz — a spread of 70.52 cents, of which 69.64 is the agreement channel's own quantisation.

The second half of the paper is about levers, and it separates two questions that are routinely merged. Elasticities are dimensionless and therefore comparable: the cost-side levers have elasticity $2.885390\,(c/\eta)$ against a geometry-lever bound of 2.885390 octaves per e-fold, so **a relation whose cost ratio exceeds one nat is one in which the cost side is proportionally stronger, always.** But elasticity says nothing about how far a lever can travel, and on the carried dyad the ranking inverts: raising conversion efficiency to a physically impossible 1.00 buys 4.174198 octaves against a requirement of 4.338420, while raising non-parallelism buys 5.545206 and closes it. **The strongest lever is the one that runs out first.**

What this paper does not claim, stated at the front. The constant $k = (9/8)^2/\varphi^2 = 0.483425732988414$ is a **defined** number, exact by construction; its components are documented — the Pythagorean whole tone [Barbour, 1951] and Euclid's extreme and mean ratio [Euclid, 1956] — but **their product expressed as a frequency in hertz has no precedent in any literature we have found**, and the standard deflationary treatment of golden-section claims [Markowsky, 1992] is cited here deliberately and again wherever φ appears. The essence law is a postulate whose **exponent is defined by the model, not estimated from data**. **Nothing in this paper is counselling, therapy, mediation, or a clinical assessment of any person, pair or group**, and where a relationship is causing a person difficulty the resource is a qualified professional.

Contributions

1. The opposition identity, proved: u is even in ρ, so both signed extremes are the ground essence, and every quantity computed downstream inherits the evenness.
2. The agreement identity $\cos\theta = 2a - 1$, $u = 4a(1 - a)$, exact and requiring no centring, with its $4/m$ resolution limit priced at 643 cents per disputed item in a 120-item log.
3. A three-route concordance on a constructed dyad, with the spread explained by the agreement channel's quantisation rather than by disagreement between the routes.
4. The unison fixed point $Q_\infty = -c/\delta$, with η and N absent and their partial derivatives exactly zero.
5. The blindness of u to pitch-class collision, exhibited by two dyads agreeing to fifteen decimal places and differing in class count.
6. The elasticities, with the dyadic lever criterion $c/\eta \geq 1$ sufficient for cost-side dominance, and the reserve theorem that inverts the ranking on the carried case.
7. The damage integral of the wrong lever order: eight periods of misordering cost twelve periods of permanent lag and 111.133601 stock-periods of cumulative gap, while the diagnostic reads unchanged throughout.

2 Hypotheses, definitions, and postulates

A hypothesis nothing could falsify is not a hypothesis, and this section does not manufacture one. Series C carries the corpus's least-supported material, so the labelling is done first.

2.1 Hypotheses proper

H1 — the concordance hypothesis. Three routes to the same dyad measure one quantity, so their readings agree within the resolution of the coarsest route. **Predicted:** where two or more routes are run on the same declared pair and situation, the spread between extremes does not exceed the coarsest route's quantisation. **Falsified by:** routes that diverge by more than their combined resolution, which would show they are measuring different things and would require the paper to say which. **Status:** demonstrated arithmetically on one constructed case where the quantisation explains 69.64 of a 70.52-cent spread. **Not tested against observation.** No observational data were collected for this paper.

H2 — the ratio-scale hypothesis. The elicited intensities carry ratio information, as magnitude estimation is held to [Stevens, 1957]. **Predicted:** cross-modality matching against the elicitation instrument returns a power relation with an exponent stable under changes of stimulus range. **Falsified by:** exponents that move with the experimenter's chosen range, which is exactly what the range effect predicts [Teghtsoonian, 1971], or by the response-bias mechanisms catalogued against magnitude estimation [Poulton, 1968]. **Status: not tested here.** The framework leans on the psychophysical literature and inherits its disputes [Krueger, 1989]. Section 4.3 computes what the failure costs.

H3 — the independence hypothesis. A dyad's available order is driven by the independence of the two parties rather than by the strength of their agreement or disagreement. **Predicted:** among pairs matched on cost, those whose judgments are uncorrelated in either direction accumulate more stored order than those correlated in either sign. **Falsified by:** matched pairs at ρ near -1 that accumulate more than pairs at ρ near 0, which would show the surplus functional's evenness is the wrong model of the phenomenon. **Status: not tested here.** The formal half — that u is even in ρ — is DERIVED and is Theorem 8.1. The empirical half is the one that could be wrong and has not been tried.

H4 — the ordering hypothesis. In an arrangement whose numerator is negative, an intervention on the decay rate is followed by a worse outcome than the same intervention applied after the numerator is repaired. **Predicted:** two arrangements given the same two interventions in opposite orders differ in trajectory, with the decay-first order spending periods below its starting stock. **Falsified by:** a decay-first ordering that does not lag, which would refute the arithmetic rather than the psychology. **Status:** DERIVED and COMPUTED on a constructed case; **not tested against observation.**

H5 — the consent condition. This is **not a hypothesis and is not offered as one.** That a measurement performed on two people requires the informed and revocable consent of both is a **condition on whether the computation may be performed at all**, adopted by this corpus, and no observation could bear on it.

2.2 Definitions — true by stipulation, not falsifiable

- **Def. 1.** $u(x, y) := 1-(x\cdot y)^2$ on unit state vectors over $N = 12$ categories; $F := \ln(1+(N-1)u)$; $N_{\text{eff}} := 1 + (N-1)u$; $f_{\text{rel}} := k\, N_{\text{eff}}^2$.

- **Def. 2.** A **unison** is a dyad with $u = 0$; equivalently $|\cos\theta| = 1$; equivalently $N_{\text{eff}} = 1$; equivalently $F = 0$; equivalently $f_{\text{rel}} = k$.
- **Def. 3.** $D := 2.885390\,(c/\eta) - \log_2(f_{\text{rel}}/k)$, the octave deficit [Curzi, 2026], whose Result Two is the identity linking it to the storage fixed point.
- **Def. 4.** $k := (9/8)^2/\varphi^2$. **This is a defined constant.** It is an exact ratio of exact quantities and is therefore computable to arbitrary precision; the fifteen significant figures quoted throughout, 0.483425732988414, are *exact digits of a definition* and carry no measurement uncertainty because there is no measurement.
- **Def. 5.** The **concordance spread** is the difference in cents between the highest and lowest route readings on one pair; the 100-cent threshold at which a point value is withheld is **a house convention and is stated as one.**

2.3 Postulates — asserted by this framework, with no external support

- **P-a. The essence law $f(n) = kn^2$ and the bridge $f = k\,e^{2F}$. No external literature supports either.** The exponent 2 is **defined by the model, not estimated from data**; no scaling exponent here was obtained by fitting, none was tested against a log-normal or exponential alternative, and the goodness-of-fit and likelihood-ratio procedure a claimed scaling relation requires [Clauset et al., 2009] was **not run, because there is no data to run it on.**
- **P-b. The attachment of hertz to k.** Computing $(9/8)^2/\varphi^2 = 0.483425732988414$ is arithmetic. **Asserting that this dimensionless number *is a frequency* is a separate and much stronger claim**, and no operational procedure is supplied by which the constant acquires the dimension of inverse time.
- **P-c. The affine effective count $N_{\text{eff}} = 1 + (N-1)u$.** A design axiom [Curzi, 2026], whose Result One characterises the surplus functional by angle-only dependence, a zero at parallel, an affine effective count and a normalisation — and it is **not a Hill number of any order.** The Hill family [Hill, 1973, Jost, 2006] and its independent rediscovery in political science [Laakso and Taagepera, 1979] ground $1/\text{HHI}$ and $\exp(H)$ exactly — as the corpus's concentration work uses them [Curzi, 2026], whose Result One gives four jurisdictions at shares $40.9/27.3/18.2/13.6$ a naive HHI of 0.293430 and 3.41 effective jurisdictions against a bloc-collapsed count of exactly 1.00 — and ground this construction not at all.
- **P-d. The state-vector encoding of a party.** That a party's construal of a situation admits representation as a unit vector over twelve categories. **Asserted, not established**, and Section 8.6 states the specific thing it cannot represent: a construal that is the exact negation of another.

2.4 Modelling choices — legitimate, but choices

- **M-a. Octave reduction applied to a non-acoustic quantity.** Octave equivalence is formalised for **pitch** [Shepard, 1982], with a measured and finite perceptual size [Demany et al., 2013]. **The *operation* is well defined; its *interpretation* here is a choice of this framework.**
- **M-b. $N = 12$.** Twelve exceeds the classical span of immediate absolute judgment [Miller, 1956], which is a caution and not a justification. The ladder is owned elsewhere in the series [Curzi, 2026].
- **M-c. The geometric convergence model.** Section 9.6 models the drift of u over time as geometric decay at a rate λ. **That is the simplest form behaving correctly at the boundaries and it has been fitted to nothing.**
- **M-d. The 100-cent concordance threshold and the 10 per cent overlap-weight floor.** Both are declared conventions with no derivation. **What is not optional is choosing one, stating it, and applying it before seeing the result.**
- **M-e. Interval apparatus.** Cents and interval arithmetic are the correct apparatus once an index exists [Rossing et al., 2002, Benson, 2006], but consonance depends on the spec-

trum of a sound and not on the ratio alone [Sethares, 2005], and interval categories generalise less far across listeners than is usually assumed [Burns, 1999, Hoeschele et al., 2012]. **This paper uses the arithmetic and claims none of the perception.**

2.5 What is expository

Sections 4.4, 9.7 and the failure-mode inventories are procedural: they set out reporting requirements and gates rather than argue a claim. Their organising question is *what must be published alongside a reading for a third party to be able to check it.*

3 Framework, dimensional audit, and method

3.1 What is being measured

The object is the geometry between two representations. Not between two people — between two *representations* of two people, produced under a declared rule set, from written framings, by a procedure published with the number.

```
WHAT u MEASURES
    the represented non-parallelism of two parties' states, under one declared ladder

WHAT u DOES NOT MEASURE
    agreement, rapport, closeness, compatibility, conflict, affection, trust,
    similarity of personality, or the quality of anything
```

The second block is longer than the first and that is the correct proportion. u answers one question: how much of what each party carries is not already carried by the other.

One further clarification saves a great deal of confusion later. u **is a property of the pair and not of either party**, because the inner product is symmetric, so there is no such thing as one party's *contribution* to u. A dyad in which one party holds eleven states and the other holds one can read lower than a dyad in which both hold three, if the eleven contain the one. **Breadth is a personal quantity; non-parallelism is a relational one; and the formalism keeps them apart on purpose.**

3.2 The three laws this paper composes

```
THE SURPLUS LAW    F(u) = ln(1 + (N−1)u)               a pair has a surplus
THE ESSENCE BRIDGE f = k e^(2F) = k (1+(N−1)u)^2        a surplus has a pitch
THE STORAGE LAW    Q_{t+1} = (1−delta)Q_t + eta F − c   surplus becomes structure, and leaks
```

Analogy status: category 2, a shared mathematical form. The storage law is the discrete leaky integrator that appears as the RC low-pass filter [Oppenheim et al., 1996, Horowitz and Hill, 2015], as exponential smoothing [Holt, 1957], and as the leaky integrate-and-fire neuron [Lapicque, 1907, Gerstner and Kistler, 2002]; its fixed point and stability condition are standard [Elaydi, 2005] and its convergence is Banach's [Banach, 1922]. **It is used because the mathematics is the same, not because the physics transfers.** The corpus states that theorem exactly and then refuses the certificate it appears to offer [Curzi, 2026] — a human life is not a metric space and nobody has measured the relevant contraction modulus — and **that refusal binds every trajectory in Section 9.**

Two further warnings attach and are recorded here rather than in a footnote. Q **has the grammar of a conserved quantity without a conservation law to give that grammar content**, which is the characteristic defect of the genre [Mirowski, 1989]; validity is claimed only within a specified arrangement [Cartwright, 1999]. And the time average of a multiplicative process

is not its ensemble average [Peters, 2019], while aggregation over heterogeneous parties does not in general yield a representative party [Kirman, 1992].

The warrant the logarithm needs is not the psychophysical one. A logarithmic form follows from multiplicative dynamics with **no psychological assumption at all** [Peters and Gell-Mann, 2016], which is what the argument here requires; Fechner [Fechner, 1860, Weber, 1834] is needed only where an elicitation is performed. **And the functional form is underdetermined**: a sub-linear power law [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the qualitative facts as well, and many mechanisms produce indistinguishable curves [Stumpf and Porter, 2012]. **Nothing in this paper discriminates among them.**

3.3 The constant k

$$k := \frac{(9/8)^2}{\varphi^2} = \frac{81/64}{\varphi + 1} = 0.483425732988414 \quad \text{(defined; exact)}$$

Both components are real and documented — the Pythagorean whole tone [Barbour, 1951, Lindley and Turner-Smith, 1993] and Euclid's extreme and mean ratio [Euclid, 1956], with $\varphi^2 = \varphi + 1$ a proved identity [Vajda, 1989, Koshy, 2001]. **Their product, expressed as a frequency in hertz, has no precedent.** The literature on golden-section claims is overwhelmingly one of *deflation*: the appearances of φ in the Parthenon, the Great Pyramid, the human body and preferred rectangles do not survive examination [Markowsky, 1992, Fischler, 1981, Livio, 2002], and the psychological preference results are weaker and more equivocal than usually reported [Höge, 1997]. φ **is being *chosen* as a constant, not *found* in a psychological system** — and the cost of the choice is small in a precise sense, because k cancels from every interval and from every verdict in this paper.

3.4 Dimensional audit

symbol	quantity	dimension	unit	status
u	non-parallelism	$[1]$	—	dimensionless by construction, $u \in [0, 1]$
ρ	correlation	$[1]$	—	dimensionless, $\rho \in [-1, 1]$
a	agreement fraction	$[1]$	—	dimensionless, $a \in [0, 1]$
F	surplus	$[1]$	nats	dimensionless; a logarithm of a count
N, m	category count, item count	$[1]$	—	absolute scale, counts
k	reference constant	$[T^{-1}]$ **asserted**	Hz **asserted**	**defined as a ratio (dimensionless); the unit is postulated**
f_{rel}	relationship frequency	$[T^{-1}]$	Hz	inherits its dimension entirely from k
D	octave deficit	$[1]$	octaves	a difference of logarithms; k cancels
Q	stored order	$[\mathcal{C}]$	none supplied	state variable; **no operational unit**

symbol	quantity	dimension	unit	status
η	conversion efficiency	$[1]$	—	dimensionless; assumed $0 < \eta \leq 1$
c	cost per period	$[\mathcal{C}T^{-1}]$	per declared period	driving term
δ, λ	decay, convergence rate	$[1]$	per declared period	dimensionless
ε	elasticity $dD/d(\ln x)$	$[1]$	octaves per e-fold	**dimensionless, which is why levers are comparable this way**
$\|x\|_2$	raw intensity norm	instrument units	as declared	**published, never discarded**

Does the equation balance? $f_{\text{rel}} = k\,N_{\text{eff}}^2$ balances only if k carries $[T^{-1}]$, because N_{eff} is dimensionless. **No operational procedure is supplied by which k acquires that dimension**, and the hertz is attached by assertion; this is the largest dimensional exposure in the paper and it is listed in Section 14 by name.

The exponent is the second exposure. The ratio c/η is $[\mathcal{C}T^{-1}]$ divided by a pure number and is therefore **not dimensionless as written**, which the exponent of $e^{2c/\eta}$ requires it to be. **The framework's convention is that c is expressed in nats of surplus per declared period — and the source never states it.** It is stated here, and every verdict in Sections 9 and 10 depends on it being applied consistently. The failure mode is measurable: a c/η built from a weekly cost and an annual conversion rate is a different number, and the corpus's own correspondence audit measures a $90.1\times$ period-normalisation trap from exactly that mistake.

The third audit result is about the elasticities and it is favourable. dD/dx for different parameters are in different units and **cannot be ranked against one another**; $dD/d(\ln x)$ are all in octaves per e-fold and can. That is the whole reason Section 9.3 exists.

3.5 Scale type, and what licenses the arithmetic

quantity	scale type claimed	what licenses it
n, N, m	absolute	cardinality: true zero, fixed unit, identity transformation group only
u, ρ, a	ratio / interval as constructed	inner products and counts of matches
elicited intensities	ratio, up to a free common multiplier	**magnitude estimation, assumed** [Stevens, 1957]; not tested here
η, c, δ	ratio, **assumed**	no validated instrument is supplied
f_{rel}	ratio in hertz, **inherited**	inherits the status of the intensities and of the postulated bridge

Stevens' typology makes the permissible arithmetic on a quantity a function of the transformation group that leaves its empirical content unchanged [Stevens, 1946]. **A count admits**

only the identity transformation; a rating does not, and the corpus has demoted one of its own keystone constructions on exactly that ground [Curzi, 2026m], whose Result Two fixes the contact point $D = P$, $r = 1$, $M = E$ with unit elasticity in all three terms and the ceiling $0 \leq M \leq \ln N$ making the divergence finite.

The strongest available objection to this framework is Michell's: that psychology routinely assumes its attributes are quantitative without testing whether they are [Michell, 1997, 1999]. The remedy exists — conjoint measurement supplies a testable condition for additive quantitative structure [Luce and Tukey, 1964, Krantz et al., 1971], and a Rasch model supplies specific objectivity [Rasch, 1960, Andrich, 1988] — and **neither has been run for any quantity in this paper.** Section 4.3 computes the price of that failure at 1065.05 cents.

3.6 Identification, and the reflexive exposure

The framework's founding construction is a reflexive loop in which each party's reading of the other shapes what the other becomes [Curzi, 2026], whose Result One is the reframing that makes the present state a local minimum of the comfort landscape by definition rather than by luck. **That is a linear-in-means reflexive specification, and the endogenous effect in such a specification is not identified from observational data** [Manski, 1993]; influence and selection are generically confounded in observational relational data [Shalizi and Thomas, 2011]. **No amount of fit recovers the parameter.** The remedy adopted here is stipulation — the rule set, the ladder and the item set are declared in advance and published with every reading, which is what makes a discrete-choice treatment of interaction defensible [Brock and Durlauf, 2001] — and **the claim is accordingly a modelling position, not an identified empirical result.**

The same objection applies to the framework's own parameters: if parties change when modelled, η, c, δ and u are not structural [Lucas, 1976, Hacking, 1995, Ferraro et al., 2005]. **This is stated, not resolved.** The exposure is sharper here than elsewhere in the corpus, because a reading issued to two people about their own relation is precisely the case where the instrument can change what it measures [Merton, 1948, Rosenthal and Jacobson, 1968].

3.7 Method

Every number in this paper is the output of exact arithmetic on stated parameters, computed in double precision. **No observational data were collected, no human subjects participated, and no empirical dataset underlies any figure or table.** Every case is constructed and is labelled as constructed at the point of use. The figures are produced by the scripts named in their captions.

4 Three routes to one quantity

4.1 Why there are three

A single route would be simpler and would be wrong, for a structural reason rather than a cautious one. **Route A's intensity vectors are non-negative by construction, so** $\cos\theta \geq 0$ **for every pair and the negative half of the correlation range is unreachable through it.** The bound is exact rather than empirical; a verification over 100,000 random non-negative unit pairs at internal scope returns a minimum observed cosine of 0.260867, well inside it.

ROUTE	channel	sign range	what it is good for
A	state vectors	cos in [0,1]	the native reading; content, not conduct
B	shared judgements	cos in [−1,1]	conduct on a fixed item set; reaches rho<0
C	errors vs known answers	rho in [−1,1]	independence proper; reaches rho<0

The three routes measure the non-parallelism of different things — what the parties *hold*, what they *say* about a fixed set of items, and how they *err* against a known standard. **A pair**

can be near-parallel in one and far from parallel in another, and when they are, that discrepancy is itself the finding.

4.2 Route B: the agreement identity

THEOREM 4.1 — THE AGREEMENT IDENTITY. Let two parties each render a binary judgment coded ± 1 on each of m shared items, and let a be the fraction on which they match. Then

$$\cos\theta = 2a - 1, \qquad u = 1 - (2a-1)^2 = 4a(1-a)$$

Proof. Each judgment vector has every entry ± 1, so each has norm $\sqrt{m}$. Their inner product is $am - (1-a)m = m(2a-1)$. Dividing by the product of the norms gives $\cos\theta = 2a - 1$. Substituting into $u = 1 - \cos^2\theta$ and expanding gives $1 - (4a^2 - 4a + 1) = 4a(1-a)$. **QED. The identity is exact and requires no centring**, because the norms of ± 1 vectors are fixed by m alone.

a	$\cos\theta$	u	f_{rel} (Hz)	cents above k
1.000000	$+1.000000$	0.000000	0.4834	0.00
0.983333	$+0.966667$	0.065556	1.4320	1880.02
0.950000	$+0.900000$	0.190000	4.6158	3906.26
0.900000	$+0.800000$	0.360000	11.8930	5544.82
0.800000	$+0.600000$	0.640000	31.2494	7217.27
0.700000	$+0.400000$	0.840000	50.6909	8054.75
0.500000	0.000000	1.000000	69.6133	8603.91
0.300000	-0.400000	0.840000	50.6909	8054.75
0.200000	-0.600000	0.640000	31.2494	7217.27
0.100000	-0.800000	0.360000	11.8930	5544.82
0.050000	-0.900000	0.190000	4.6158	3906.26
0.000000	-1.000000	0.000000	0.4834	0.00

The table is exactly symmetric about $a = 0.5$, and the top and bottom rows are identical. Two parties who agree on every item and two parties who disagree on every item produce the same u, the same frequency and the same reading. Section 8 owns that result; **what matters here is that the measurement channel exhibits it directly, in a form anyone can compute from a spreadsheet.**

Route B's resolution limit is severe and must be stated with the route. Because a is quantised to $1/m$, u is quantised to roughly $4/m$ near the top of the agreement range:

```
m        a (2 disagreements)   u         f (Hz)    one more disagreement costs
 40        0.950000          0.190000    4.6158           938.89 cents
 60        0.966667          0.128889    2.8259           849.42 cents
120        0.983333          0.065556    1.4320           643.28 cents
240        0.991667          0.033056    0.8989           428.21 cents
400        0.995000          0.019900    0.7182           295.58 cents
1000       0.998000          0.007984    0.5721           136.62 cents
5000       0.999600          0.001600    0.4920            29.78 cents
```

In a 120-item log, one disputed item is worth 643 cents — more than a perfect fifth. A near-parallel dyad cannot be resolved better than a few hundred cents by an agreement log of plausible length, and **the honest output of Route B in that regime is an interval, not a number.**

4.3 Route C, and the two invariances that decide whether any route means anything

DEFINITION 4.2 — THE ERROR CHANNEL. Both parties estimate each of m quantities whose true values subsequently arrive; form each party's error vector, let ρ be the correlation of the two, and set $u = 1 - \rho^2$.

It is signed, so it reaches the negative half. **It requires ground truth**, which is its cost and is not small: Route C is unavailable for questions whose answers never arrive, which is most of what two people actually disagree about. **And it is the route that catches shared noise.** Two parties answering at random on a shared item set have expected agreement $a = 0.5$, hence $u = 1$, hence $f_{\text{rel}} = 69.6133$ Hz — the ceiling of the entire system, from two people guessing. **The surplus formalism cannot distinguish independence from noise, because both are orthogonality**, and F is in any case a lower bound on the ensemble's information rather than the information itself.

> **THE ACCURACY GUARD [BINDING WHEREVER ROUTE B OR C IS USED].** A reading is published with a measure of each party's accuracy on the item set. **A high u between two parties who are both largely wrong is arithmetically correct and practically worthless. Surplus is a measure of distinctness, never of correctness**, and this framework has no formalism that converts one into the other.

The first invariance is good news. Because $(\alpha x)/\|\alpha x\| = x/\|x\|$ for $\alpha > 0$, any common positive rescaling of the intensities cancels exactly. Recorded on 0–10, on 0–1, on 0–100 and as proportions of a total, the same framing normalises to $(0.872871561,\ 0.43643578,\ 0.21821789)$ **identically to nine decimal places.** The elicitation instrument's unit is unconstrained: two analysts using different scales produce the same u, the same F and the same f_{rel}, provided the *ratios* agree.

The second invariance does not hold, and here is by how much. Under an ordinal reading, applying any strictly increasing function to the intensities destroys no information. Computed, against a fixed reference framing:

transform	u	f_{rel} (Hz)	cents above k
$\log(1+x)$	0.362115	12.0049	4849.03
$\sqrt{\cdot}$	0.371445	12.5044	4919.61
identity	0.388360	13.4361	5044.14
square	0.525275	22.2093	5914.09

A spread of 1065.05 cents — 0.888 octaves — produced entirely by relabellings that preserve every ordinal fact in the data. Every transform leaves the ranking of the states untouched; an ordinal scale calls them the same measurement, and the relationship frequency calls them a minor seventh apart. The same failure is recorded in the criterion paper as its Result Seven [Curzi, 2026].

> **THEREFORE "REPORT IT AS ORDINAL" IS NOT AN AVAILABLE FALLBACK.** It is not a weaker version of the reading. **With ordinal intensities the relationship frequency is not determined**, and any number published from them is a property of the labelling rather than of the pair.

A further symptom of the same disease: $\log(1+x)$ is unit-dependent, landing at $u = 0.366939$ on a 0–1 instrument against 0.362115 on 0–10. **A transform family admissible under an ordinal reading includes transforms whose effect depends on an arbitrary choice of units, and the reading moves when the units move.**

The requirement that survives both invariances is precise: the intensities must carry ratio information while the unit in which the ratio is expressed remains free, which is exactly the definition of a ratio scale. **The honest form of the claim, softened nowhere: the intensities are**

ratio-scaled to the extent that magnitude estimation delivers ratio scales, which is a measured property of a procedure and not a constant of nature.

4.4 The magnitude reporting rule

Normalisation discards the length of each vector, and the length is not nothing:

```
low-intensity  {3: 1,   5: 0.5, 7: 0.25}   raw L2 norm = 1.145644
high-intensity {3: 8,   5: 4,   7: 2   }   raw L2 norm = 9.165151
ratio of norms = 8.0000x        u between them = 0.000000000000   exactly zero
```

> **THE MAGNITUDE REPORTING RULE.** Both parties' raw L2 norms are reported alongside the reading, in the instrument's units, and are never discarded. u **measures direction; the norms measure how much there is.** A dyad reading published without both norms has thrown away half the measurement and made the free multiplier unrecoverable.

The dyad case adds a requirement the single-vector case does not have: the *ratio* of the two norms is itself reportable, because a large asymmetry is a fact about the elicitation that the normalised geometry cannot express. The carried dyad's ratio is 1.065667, small enough that nothing in its reading is an artefact of one party reporting bigger numbers.

One elicited number moves the reading by more than an octave, which is why the ratio question must be asked explicitly, the reading double-coded, and an interval reported:

$B[10]$	u	f_{rel} (Hz)	shift from the carried reading
0.00	0.019232	0.7096	-1285.20 cents
0.10	0.032095	0.8850	-902.77 cents
0.20	0.068735	1.4908	0.00 (the carried case)
0.30	0.124003	2.7017	$+1029.33$ cents
0.40	0.191202	4.6554	$+1971.38$ cents

5 What u cannot see

5.1 Blindness to pitch-class collision

u is computed from an inner product over the twelve indices. It knows which indices carry weight. **It does not know that indices 5 and 10 fold to the same pitch class**, because folding is a property of the essence law and not of the geometry.

THEOREM 5.1 — u IS BLIND TO PITCH-CLASS COLLISION. There exist dyads with identical u, identical f_{rel}, and different numbers of distinct pitch classes in the union of their supports.

Proof by construction, both cases computed.

```
A  = {4: 0.7, 6: 0.5, 9: 0.4}

B  = {4: 0.6, 6: 0.55, 9: 0.3, 10: 0.2}
     cos theta = 0.965020899   u = 0.068734665264634   f_rel = 1.4907987150 Hz
     union {4, 6, 9, 10}   odd parts 1, 3, 9, 5   ->  FOUR distinct pitch classes

B* = {4: 0.6, 6: 0.55, 9: 0.3, 12: 0.2}
     cos theta = 0.965020899   u = 0.068734665264634   f_rel = 1.4907987150 Hz
     union {4, 6, 9, 12}   odd parts 1, 3, 9, 3   ->  THREE distinct pitch classes
```

The two agree to fifteen decimal places, because the inner product runs only over the three shared indices and both partners have the same norm. **QED.**

Two dyads at the same relationship frequency, one occupying four pitch classes and one three: the geometry cannot tell them apart and the lattice can. The odd-part rule that produces the collision is owned by the folding paper [Curzi, 2026], whose Result One establishes that two indices share a pitch class if and only if they share an odd part.

> **THE COLLISION TEST.** After computing u, list the union of the two supports, take the odd part of every index, and report the count of distinct pitch classes beside the count of distinct states. Where the two differ, say by how much and name the colliding pairs. **The test takes two minutes and it is the difference between reporting four perspectives and reporting three.**

5.2 Degeneracy at dyad scope

```
DISJOINT SUPPORT BETWEEN TWO PARTIES -- computed
P = {1:0.6, 2:0.5, 4:0.3}     Q = {3:0.7, 6:0.4, 9:0.5}
    cos theta = 0.000000000000    u = 1.000000000000    f_rel = 69.6133 Hz
P = {5:0.9, 7:0.1}            Q = {8:0.2, 11:0.9}
    cos theta = 0.000000000000    u = 1.000000000000    f_rel = 69.6133 Hz
```

Both pairs score the maximum the system can produce, and both scores are produced by the encoding rather than by the parties. Any two vectors with no index in common are orthogonal, so $u = 1$ regardless of what the states are or how intense they are. **The degeneracy is flattering, which is why it survives**: a pair told they occupy the full twelve-fold relational range will not challenge the finding, and an instrument whose failure mode is to award full marks will be adopted enthusiastically and never questioned by the people it flatters.

> **THE OVERLAP CONDITION AT DYAD SCOPE [BINDING].** A dyad reading may be computed only if the two supports share at least one state carrying non-trivial weight in both. Where they do not, the computation is reported as DEGENERATE and no f_{rel} is published.

The plain-language version is what makes it obviously right: two people described as having nothing whatever in common have not been described as a relationship. **A subtler version is more common in practice** — support overlapping in one state at very low intensity in one party produces a cosine dominated by a single small term, and the reading becomes a fact about that term. A weight floor exists for exactly that reason and it is a declared convention (M-d).

6 The unison

6.1 What it is

DEFINITION 6.1. A **unison** is a dyad with $u = 0$; equivalently $|\cos\theta| = 1$; equivalently $N_{\text{eff}} = 1$; equivalently $F = 0$; equivalently $f_{\text{rel}} = k$ exactly.

A unison is not a low reading. It is the absence of a reading: the pair occupies zero cents of the 8603.91 available, realises 0.00 per cent of the system's surplus, and is positioned exactly where a pair that does not exist would be positioned.

6.2 The opposition identity

THEOREM 6.2 — THE OPPOSITION IDENTITY. For two represented states with correlation ρ, $u = 1 - \rho^2$, so u depends on ρ only through $|\rho|$. Consequently $u(\rho) = u(-\rho)$ for every ρ, and in particular

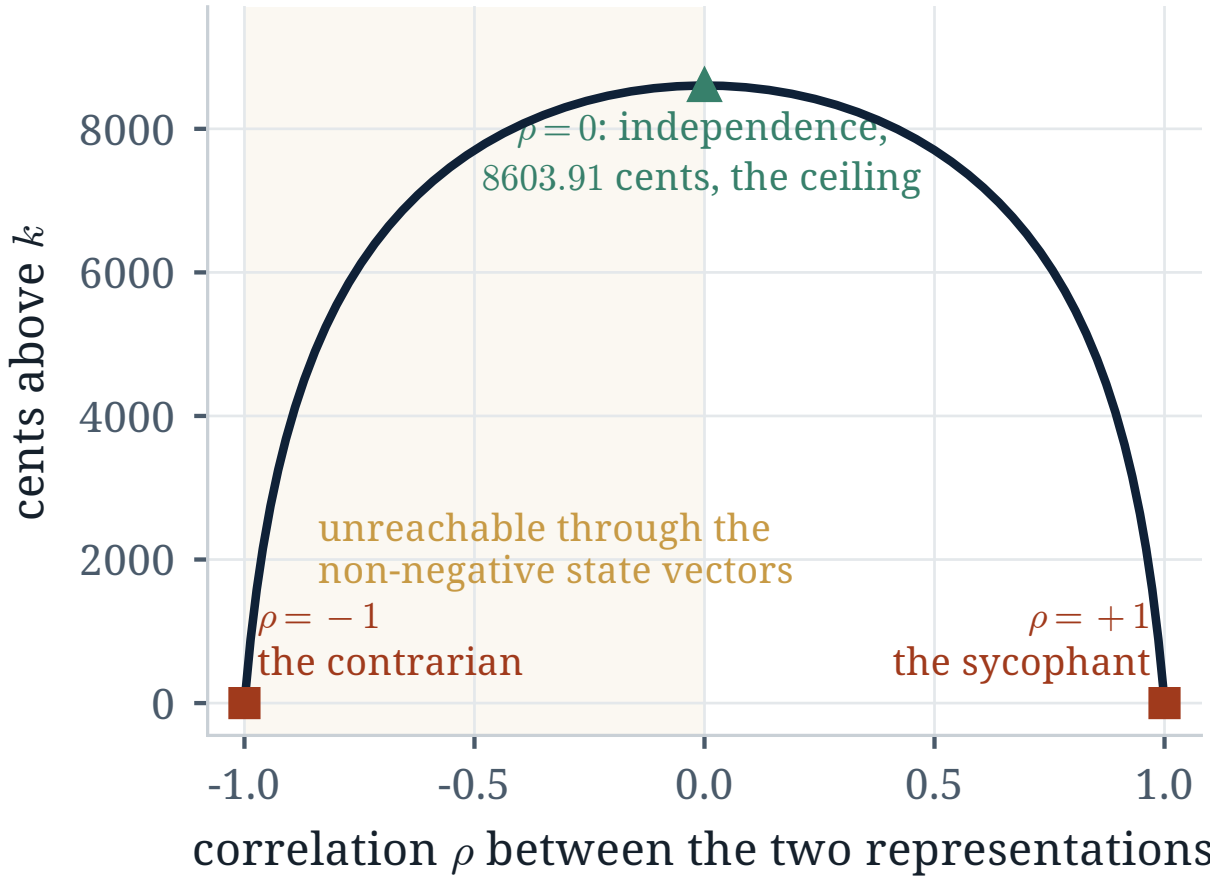


Figure 44.1: **Both signed extremes of the correlation are the same reading**: the surplus range is folded in the middle, so $\rho = +1$ and $\rho = -1$ each give $u = 0$ and $f_{\text{rel}} = k$ exactly, while the maximum sits at $\rho = 0$ with 8603.91 cents. The shaded half is unreachable through the framework's non-negative state vectors and can be measured only on the signed routes. Modelled from $u = 1-\rho^2$ and $f_{\text{rel}} = k(1+11u)^2$ with the **defined** constant $k = 0.483425732988414$ and $N = 12$. Built by `figs/src/CCP-044_fig01.py`.

$$u(+1) = u(-1) = 0, \qquad F(+1) = F(-1) = 0, \qquad f_{\text{rel}}(+1) = f_{\text{rel}}(-1) = k$$

Proof. $f_{\text{rel}} = k(1 + (N-1)u)^2$ with $u = 1 - (x \cdot y)^2$ for unit vectors, and $x \cdot y$ is the correlation. **Squaring is an even function**, so u is even in ρ, and every quantity computed from u — F, N_{eff}, f_{rel}, the octave deficit, Q_∞ — is even in ρ as well. At $\rho = \pm 1$, $u = 0$, so $N_{\text{eff}} = 1$, $F = \ln 1 = 0$, and $f_{\text{rel}} = k$. **QED.**

The proof is three lines and it is complete. There is no term anywhere in the construction that carries the sign of the correlation, so no quantity derived from it can recover the sign. This is not a limitation of the twelve-state ladder, not an artefact of $N = 12$, and not a consequence of the affine axiom. **It is a property of the inner product itself**: $u = \sin^2\theta$, and the sine squared of an angle and of its supplement are equal. The corpus's group work states the same result in its own register [Curzi, 2026r], whose Result One records that a pair in perfect opposition and a pair in perfect agreement both read $u = 0$, $N_{\text{eff}} = 1$ and surplus zero, because the extremality condition is stated on the squared alignment; the roster form is owned by the effective-count paper [Curzi, 2026c], whose Result Three states that a standing contrarian buys no surplus.

Computed across the full range, the table is a **perfect mirror about** $\rho = 0$, and both extremes fold to 247.514 Hz at $j = 9$ in pitch class 1 — the same note, at the same transposition, in the same class.

6.3 What dies at $u = 0$

THEOREM 6.3 — THE UNISON FIXED POINT. At $u = 0$, $F = 0$, therefore $Q_\infty = (\eta \cdot 0 - c)/\delta = -c/\delta$. **$\eta$ and N do not appear.** *Proof:* by inspection of the storage fixed point with $F = 0$ [Curzi, 2026a]. **QED.**

Verified across a grid at $c = 0.62$, $\delta = 0.09$:

eta	N	octave deficit D	Q_inf
0.05	2	35.7788	−6.888889
0.05	100	35.7788	−6.888889
0.30	2	5.9631	−6.888889
0.30	100	5.9631	−6.888889
0.60	12	2.9816	−6.888889
0.99	12	1.8070	−6.888889

A twentyfold improvement in conversion efficiency and a fiftyfold increase in the category count leave the destination bit-for-bit unchanged. The deficit moves, because D depends on η explicitly, but the fixed point does not, because the two dependencies cancel exactly through $Q_\infty = -D\,\eta/(2.885390\,\delta)$. The partial derivatives make the same point:

PARTIALS OF Q_inf (eta = 0.30, delta = 0.09, N = 12, c = 0.62)

u	dQ/du	dQ/deta	dQ/dc	dQ/dN
0.000000	36.6667	0.0000	−11.1111	0.0000
0.068735	20.8798	6.2565	−11.1111	0.1305
0.500000	5.6410	20.7978	−11.1111	0.2564

At $u = 0$ exactly two of the four levers have any effect on the destination, and the u lever is at its strongest value anywhere on the range. $dQ/d\eta = F/\delta$ is zero because there is nothing to convert; dQ/dN is zero because N enters only multiplied by u. **Raising conversion efficiency in a unison relation converts nothing more efficiently; adding categories adds dimensions nobody is using.** The framework's headroom sensitivities record the same weakness of the N lever in a different register [Curzi, 2026o], whose Result Five gives $dh/dN = \eta/N = 0.0038$ against $dh/d\eta = \ln N = 2.197$ — **adding categories is the weakest available lever.**

6.4 Two in unison are one paying twice

```
UNISON DYAD    eta = 0.30, c = 0.62, delta = 0.09, F = 0
               eta*F − c = −0.620000 per period
               Q_inf = −6.888889        D = +5.9631 octaves

SOLO ACTOR     same delta, same eta, F = 0, but c = 0.20 -- no coordination overhead
               Q_inf = −2.222222        D = +1.9236 octaves

THE SECOND SEAT COSTS   4.666667 units of Q_inf   and   4.0395 octaves of deficit
THE SECOND SEAT RETURNS 0.000000
```

The solo cost is a **stipulated** value chosen to represent a relation without coordination overhead, and nothing turns on its exact size. **The entire ledger of a unison pair is cost**: the numerator ηF is exactly zero, so the second party contributes nothing to accumulation and everything to loss. **And the sentence is symmetric**, which is the whole point — replace the sycophant with a contrarian and every number is unchanged.

6.5 Where opposition is reachable, and where this framework cannot see it

An honest paper has to say where its own encoding cannot go. The twelve-state intensity vectors are non-negative by construction, so within Route A $u = 0$ **can arise only from parallelism: the sycophant is representable and the contrarian is not.**

```
ROUTE A   state vectors, non-negative       cos in [0, 1]     agreement branch ONLY
ROUTE B   shared judgements, +1/−1          cos in [−1, 1]    BOTH branches
ROUTE C   error vectors vs known answers    rho in [−1, 1]    BOTH branches
```

The opposition branch is measured on the signed channels or it is not measured at all, which is exactly why this paper builds three routes rather than one. **An analyst running only Route A will never see a negative correlation and may conclude the branch is theoretical. It is not theoretical; it is invisible to that instrument.** The twelve states record which components a construal requires, not which direction it points, so a framing that is the exact negation of another cannot be represented in Route A at all. **A signed encoding would be a different instrument and this framework does not have one.**

6.6 Consonance and redundancy, at relational scope

COROLLARY 6.4. The most consonant possible dyad is the most redundant possible dyad, and its frequency is k.

Proof. Consonance in this system is small harmonic distance; the extreme is exact coincidence, which for two parties means identical normalised state vectors; identical vectors give $\cos\theta = 1$, hence $u = 0$, hence $f_{\text{rel}} = k$. **QED.** The identity behind it is owned by the consonance paper [Curzi, 2026], whose Result Two records that the minimum harmonic distance $\text{HD}_{\text{ess}} = 2.0000$ is attained by exactly the six collision pairs $\{1{:}2,\ 2{:}4,\ 3{:}6,\ 4{:}8,\ 5{:}10,\ 6{:}12\}$ and by nothing else.

That paper's caveat travels with the identity and is not dropped here: the ordering is not total. It is not true that every collision is more consonant than every non-collision — $1{:}3$ at $\text{HD}_{\text{ess}} = 3.1699$ beats the collisions $1{:}4$, $2{:}8$ and $3{:}12$ at 4.0000. The collision set is dramatically more consonant on average (mean 3.0000 against 8.6617) and the ranges overlap. **A reader who builds on the total-order version will eventually be wrong about** $1{:}4$**.**

And its firewall stands: nothing about distress, health, function or experience follows from a consonance reading, in either direction. A unison dyad is not a bad relationship, an unhealthy one, or one that anybody should change. **It is a relation that, at these costs, does not accumulate order. Those are different sentences and this paper will not merge them.**

6.7 The flat maximum, and the design rule

$du/d\rho = -2\rho$, so **at** $\rho = 0$ **the derivative is exactly zero: the optimum is second-order flat.** Near independence an error in the correlation costs almost nothing and the cost grows quadratically rather than linearly.

$\lvert\rho\rvert$	u	f_{rel} (Hz)	cents surrendered vs $\rho = 0$
0.00	1.000000	69.6133	0.00
0.05	0.997500	69.2946	7.94
0.10	0.990000	68.3429	31.89
0.20	0.960000	64.6019	129.34
0.30	0.910000	58.6009	298.13
0.50	0.750000	41.3631	901.22

$\lvert\rho\rvert$	u	f_{rel} (Hz)	cents surrendered vs $\rho = 0$
0.90	0.190000	4.6158	4697.65
0.99	0.019900	0.7182	7918.52
1.00	0.000000	0.4834	8603.91

The marginal cost of one further hundredth of correlation runs from 0.32 cents at the origin to 6.73 at $|\rho| = 0.10$, 41.84 at 0.50, 230.61 at 0.90 and 685.39 at the extreme. **Missing perfect independence by a tenth costs a sixth of a semitone; missing perfect *opposition* by the same tenth costs 4697.65 cents.**

> **THE DESIGN RULE.** Where a relation is to be arranged rather than merely read, **target $|\rho|$ near zero and never ρ near -1.** Independence is a robust target — second-order insensitive to error, forgiving of imprecision, and the global maximum of the surplus. **Opposition is the *other* global minimum, sits at the steepest part of the curve, and returns exactly what agreement returns: nothing.**

The rule is not a preference for harmony over conflict. It says nothing about how two people should treat each other. It says that if the goal is available order, the coordinate to aim at is uncorrelated, and both correlated extremes are worthless — one of them expensively so.

6.8 What raises u, and what only appears to

Manufactured disagreement raises cost and returns exactly nothing. A party who adopts contrary positions in order to "bring difference" moves the pair toward $\rho = -1$, which reads as $u = 0$ just as agreement does, while adding coordination overhead. **The reading gets worse in both terms at once.**

```
RULED OUT -- these move rho toward -1 and return nothing
  - assigning a person to argue the opposite of whatever is proposed
  - rewarding disagreement as such, or counting objections as output
  - recruiting for temperament rather than for method or evidence base
  - any arrangement that requires a person to hold a position they do not hold

RULED IN -- these move |rho| toward 0
  - different evidence base: parties who look at different data
  - different method: parties whose procedures fail in different ways
  - different exposure: parties who bear different consequences from one decision
  - different training: parties whose priors were formed on different problems
  - independent work before contact
```

The distinction between the two lists is not attitude; it is whether the difference is in what the parties actually do or in what they say about what someone else did. A party with a different evidence base disagrees when the evidence differs and agrees when it does not, so their correlation is low in magnitude and of no fixed sign. **A party assigned to disagree has a correlation near -1 by construction, and the identity prices it at zero.**

That is a defence of the red-team tradition rather than an attack on it. A red team that brings different methods, different data and different exposure sits near $\rho = 0$ and is worth up to the full 8603.91 cents; **a red team that is one person told to object sits near $\rho = -1$ and is worth nothing at all.** The tradition names a real target and the arithmetic says which implementations hit it. That reading is consistent with the surrounding literature on composition: the returns to group composition are contested and depend on what kind of difference is present [van Knippenberg and Schippers, 2007, Mannix and Neale, 2005], the collective-performance factor is driven by composition rather than headcount [Woolley et al., 2010],

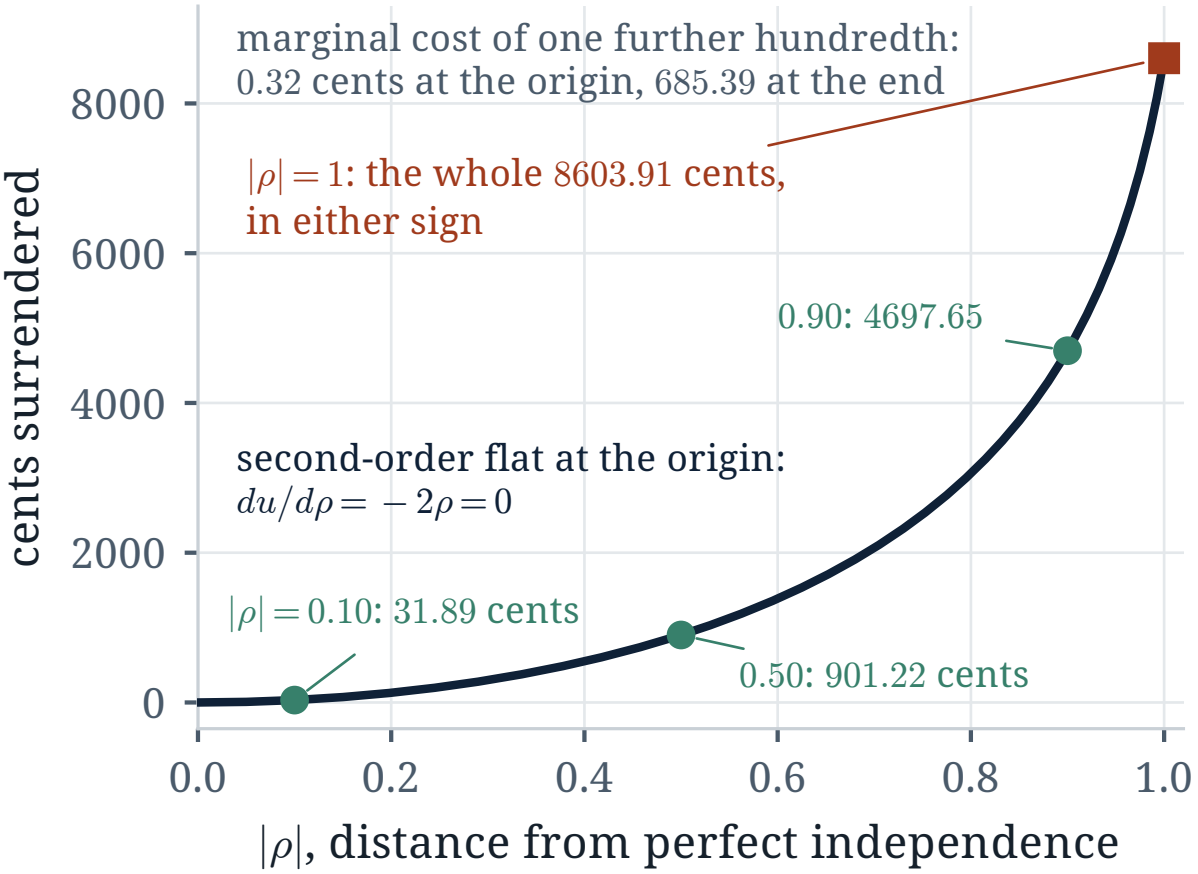


Figure 44.2: **The maximum is second-order flat and the extremes are not**: missing perfect independence by a hundredth of correlation costs 0.32 cents, while the last hundredth before perfect correlation costs 685.39. Computed from $u = 1-\rho^2$ and $f_{\text{rel}} = k(1+11u)^2$ with the **defined** constant $k = 0.483425732988414$ and $N = 12$; the plotted quantity is cents surrendered relative to $\rho = 0$. Built by `figs/src/CCP-044_fig03.py`.

the problem-solving result usually cited for diversity [Hong and Page, 2004, Page, 2007] has a standing formal objection [Thompson, 2014] that must be cited with it, and the aggregation result that makes independence load-bearing is the classical one [Galton, 1907, Surowiecki, 2004] together with its failure mode when independence is destroyed by social influence [Lorenz et al., 2011, Bikhchandani et al., 1992, Banerjee, 1992]. **The portfolio statement of the same fact predates all of it: diversification benefit depends on covariance, not on count** [Markowitz, 1952].

There is a further consequence for how such a role is filled. **A role defined as "hold the contrary position" is a role in which declining is structurally unavailable** — the person cannot say *I agree with this one* and still be doing the job. A role defined as "bring this evidence base" leaves declining intact on every question. **The first is prohibited by this corpus's consent constraint and is worthless by its arithmetic; the second is neither. When the ethical constraint and the pricing agree, there is no trade-off to manage.**

7 The dyadic reading

```
u        measured between the two parties                    Section 4
f_rel    = k (1 + (N−1) u)^2
f_min    = k e^(2c/eta)                                      CCP-040, R1
D        = 2.885390 (c/eta) − log2(f_rel/k)
Q_inf    = ( eta F(u) − c ) / delta

D <= 0   ACCUMULATES        D = 0   HOLDS        D > 0   DECAYS
```

The ceiling test comes first and it is not a stylistic preference: compare c/η to $\ln N$, and if the ratio exceeds it the case is structural and no assignment of u closes it [Curzi, 2026]. **A deficit does not announce its own kind, and a prescription issued before the ceiling test may be unsatisfiable** — which is worse than no prescription, because it consumes the effort that restructuring needed. The corpus's headroom gauge states the same constraint in stock units [Curzi, 2026b], whose Result One gives $h_t = \eta \ln N - \delta Q_t - C_t$ with its three regimes and its maximum admissible increment; and the intervention protocol adds the reserve discipline [Curzi, 2026j], whose Result One recommends $r \leq 0.7h$ because $r_{\max} = h$ is attained only at $u = 1$, together with the branch in which **no interaction geometry can prevent order from decreasing.**

What the four parameters are in a bond, concretely, because the abstraction invites people to fill them in wrongly. η is the fraction of what the two parties do together that survives the period — not how much they do, and not how intensely. c is what the arrangement costs to run before it produces anything. δ is the rate at which accumulated order leaks away if nothing is added. u **is the only one of the four that is a joint geometric property rather than a property of the arrangement**, and **one clock, declared, governs all three of the others.**

8 Levers: elasticity, reserve, and order

8.1 The order-of-operations rule

THEOREM 8.1 — DECAY AMPLIFIES THE SIGN IT FINDS. Write the numerator $n := \eta F(u) - c$. For $\delta > 0$, $Q_\infty = n/\delta$ and $dQ_\infty/d\delta = -n/\delta^2$, so

```
n > 0   ->   reducing delta RAISES Q_inf
n < 0   ->   reducing delta LOWERS Q_inf
n = 0   ->   Q_inf = 0 for every delta;  the lever is exactly inert
```

Proof: differentiate n/δ; $\delta^2 > 0$, so the sign of the derivative is the sign of $-n$. **QED. The theorem has no free parameters and no conditions beyond $\delta > 0$. It is not a tendency, not an empirical regularity, and not a claim about people. It is what division does.**

> **An intervention that lowers δ multiplies whatever the relation is already doing. In a bond that is accumulating, it accumulates faster. In a bond that is decaying, it decays faster. Trust is a multiplier on an existing sign, not a source of order.**

Interpretive care is required and this paper states it. That trust acts on δ rather than on η is an assumption of the framework's parameterisation, not a finding of this paper, and the literature cited for trust is about reciprocity, monitoring and institutional form rather than about decay rates [Gambetta, 1988, Ostrom, 1990, Axelrod, 1984, Putnam et al., 1993]. **If trust instead raises η, the arithmetic above does not apply**, and which channel it operates through is not settled here.

And nothing in this section says trust is bad. In a relation with a positive numerator it is the single most powerful move available: on the repaired carried dyad, δ from 0.09 to 0.05 takes Q_∞ from $+6.084427$ to $+10.951968$, an 80.0 per cent improvement from one change with no new difference required and no cost cut. **The rule is about sequence, not about worth.** There is also a category of intervention that moves both sides at once — a written record slows the leak *and* converts more of what happens into durable structure — and **interventions touching both sides are the ones to reach for first.**

8.2 The destination is order-independent; the path is not

THEOREM 8.2. The fixed point of the terminal parameter set is independent of the sequence of parameter changes. **The trajectory is not.**

Verified on the constructed carried dyad from $Q_0 = 1.0$, with two orderings of exactly the same pair of interventions, each applied for eight periods before the other:

```
 t      PATH A (decay first)    PATH B (numerator first)
 0        +1.000000                 +1.000000
 1        +0.498925                 +1.457598
 2        +0.022905                 +1.874013
 4        −0.858924                 +2.597783
 8        −2.373029                 +3.693462
12        +0.098675                 +5.039870
20        +3.751671                 +7.029761
32        +7.061215                 +8.832564
inf      +10.951968                +10.951968

B reaches +3.6935 at t = 8;    A reaches it at t = 20     lag 12 periods
B reaches +6.1365 at t = 16;   A reaches it at t = 28     lag 12 periods
cumulative gap over 32 periods:  111.133601 stock-periods
```

Both arrive at exactly the same place. One of them arrives twelve periods later, forever, having spent thirteen periods below where it started. The lag does not shrink. And the interval in which Path A's stock is negative is the interval in which the arrangement may simply not survive to reach its fixed point: the formalism will happily compute a trajectory through negative stock, but the admissibility condition of the storage law fails there [Curzi, 2026]. **This is why order-of-operations is a rule and not a preference.**

The number that will not appear in any of this is a change in the diagnostic. D reads $+4.338420$ at the start of Path A and $+4.338420$ at the end of its first eight periods. **Whatever instrument the parties are using, if it is the deficit alone, it reports that nothing happened.**

8.3 The elasticities

Partial derivatives of D with respect to different parameters are in different units and **cannot be ranked against one another** — comparing them is comparing octaves-per-apple to octaves-per-orange. **Elasticities are dimensionless and are the comparison that exists.**

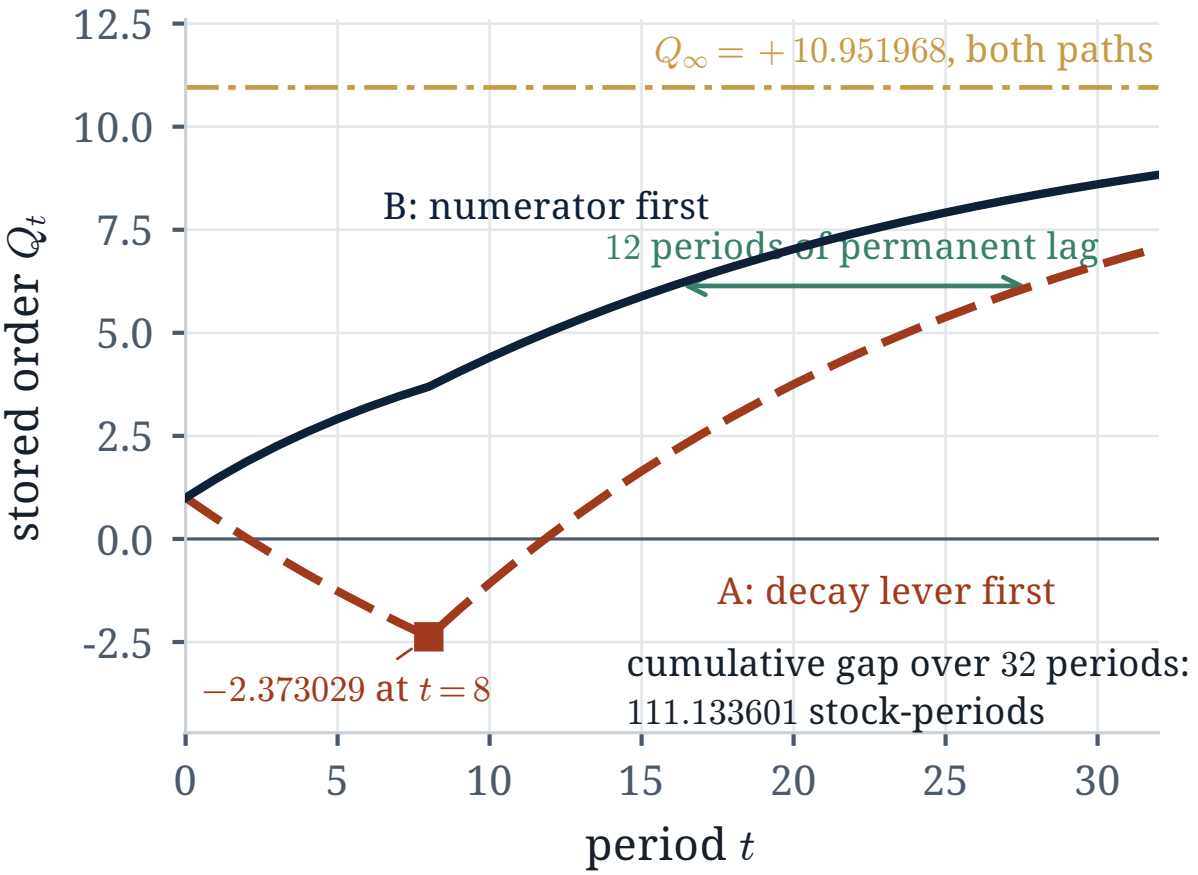


Figure 44.3: **Eight periods of the wrong lever order cost twelve periods of permanent lag**, and the deficit reads unchanged throughout: Path A lowers the decay rate while the numerator is still negative and bottoms at -2.373029, while Path B repairs the numerator first and never falls below its starting value. Both converge to $Q_\infty = +10.951968$; the cumulative gap over 32 periods is 111.133601 stock-periods. Modelled from $Q_{t+1} = (1-\delta)Q_t + n$ on the constructed carried dyad with $n = -0.451075$ before repair and $+0.547598$ after. Built by `figs/src/CCP-044_fig05.py`.

THEOREM 8.3 — THE ELASTICITIES. With $D = 2.885390\,(c/\eta - F)$ and $F = \ln(1 + (N-1)u)$:

$$\frac{dD}{d\ln u} = -2.885390\left(1 - \frac{1}{N_{\text{eff}}}\right), \qquad \frac{dD}{d\ln \eta} = -2.885390\,\frac{c}{\eta}, \qquad \frac{dD}{d\ln c} = +2.885390\,\frac{c}{\eta}$$

Proof: $dD/d\ln x = x\,dD/dx$, applied to the four partials. **QED.**

Three consequences, all exact. **(i) The two cost-side levers have equal and opposite elasticity**, because they enter D only as a ratio; whichever is cheaper to move is strictly the better choice, and the formalism is indifferent between them. **(ii) The cost-side elasticity is exactly the number of octaves the relation requires**, so an e-fold improvement in the cost ratio moves the deficit by precisely the requirement. **(iii) The geometry lever's elasticity is bounded above by 2.885390 octaves per e-fold, at every u, in every system**, because $1 - 1/N_{\text{eff}} < 1$ always, and the bound is approached only as $u \to 1$ where there is nothing left to buy.

COROLLARY 8.3a — THE DYADIC LEVER CRITERION.

$$\left|\frac{dD}{d\ln \eta}\right| > \left|\frac{dD}{d\ln u}\right| \iff \frac{c}{\eta} > 1 - \frac{1}{N_{\text{eff}}}, \qquad \text{and } \frac{c}{\eta} \geq 1 \text{ is SUFFICIENT.}$$

A relation whose cost ratio exceeds one nat is one in which the cost-side levers are proportionally stronger than the geometry lever — always, at every u, with no exceptions. Computed on three constructed cases:

case	c/η	$1 - 1/N_{\text{eff}}$	$\lvert dD/d\ln\eta\rvert$	$\lvert dD/d\ln u\rvert$	ratio
dyad, as found	2.066667	0.430550	5.963140	1.242306	4.8001
dyad, repaired	0.909091	0.883938	2.623082	2.550507	1.0285
commons	2.426471	0.169955	7.001314	0.490387	14.2771

The criterion holds in all three, and in the repaired dyad it holds by 0.025153 nats — barely. That is worth noticing rather than glossing: **a relation that has been tuned well is one in which the levers have equalised, and that is a signature, not a coincidence.**

8.4 The reserves

Elasticity says which lever moves the deficit fastest per proportional unit. **It says nothing about how far each lever can go before it runs out**, and confusing the two is the most expensive planning error available here.

THEOREM 8.4 — THE LEVER RESERVES. The maximum octaves each lever can supply alone, from an operating point (u_0, η_0, c_0, N):

```
R_u    = 2.885390 ( ln N − F(u_0) )                ( u -> 1 )   = the pitch headroom
R_eta  = 2.885390 c_0 ( 1/eta_0 − 1/eta_max )      ( eta -> eta_max )
R_c    = 2.885390 ( c_0 / eta_0 )                  ( c -> 0 )
R_N    unbounded in principle; 2.885390 u / N_eff per category in practice
```

Computed on the carried dyad as found, against its deficit of $+4.338420$:

```
R_u    ( u -> 1 )        5.545206 octaves       CLOSES the deficit
R_c    ( c -> 0 )        5.963140 octaves       CLOSES the deficit
R_eta  ( eta -> 0.50 )   2.385256 octaves       does not
R_eta  ( eta -> 0.75 )   3.577884 octaves       does not
R_eta  ( eta -> 1.00 )   4.174198 octaves       does not -- and eta_max = 1 is generous
```

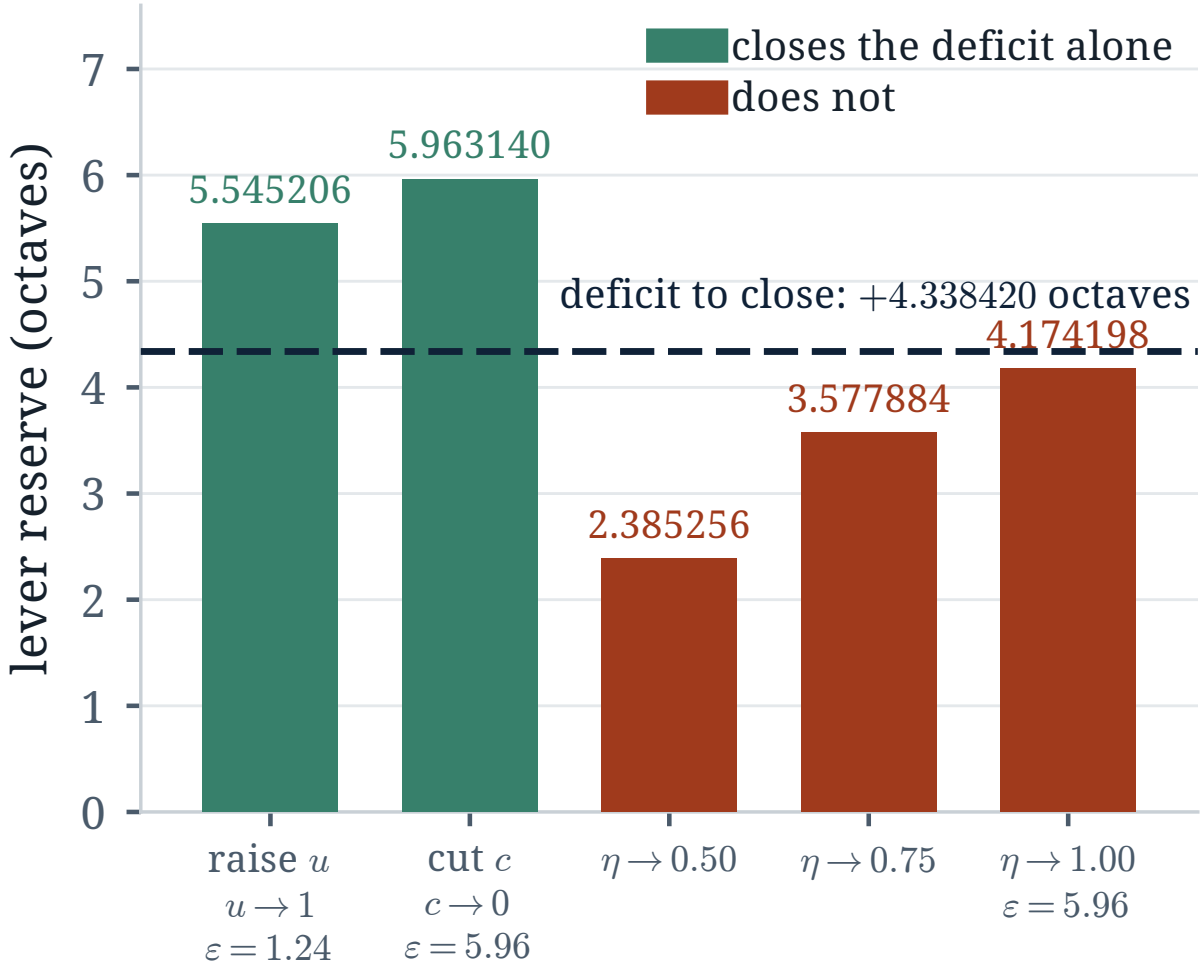


Figure 44.4: **The strongest lever is the one that runs out first**: raising conversion efficiency has 4.80 times the elasticity of the geometry lever and reaches only 4.174198 octaves even at a physically impossible $\eta_{\max} = 1.00$, against a deficit of $+4.338420$ that raising non-parallelism (5.545206) or cutting cost (5.963140) closes. Modelled from the reserve theorem on the constructed carried dyad with $u = 0.068735$, $\eta = 0.30$, $c = 0.62$, $N = 12$; ε is the elasticity $|dD/d\ln x|$ in octaves per e-fold. Built by `figs/src/CCP-044_fig04.py`.

The lever with 4.8 times the elasticity has less reserve than the deficit requires. Raising η from 0.30 all the way to a physically impossible 1.00 — a system converting every unit of interaction into permanent structure, with nothing lost — buys 4.174198 octaves against a requirement of 4.338420 and falls short by 0.164222. The same fact in its other form is $\eta_{\text{req}} = c/F = 1.101077$: **not an expensive lever but an absent one.**

> **THE HONEST DYADIC LAW: the cost-side levers have the higher elasticity and the lower ceiling; the geometry lever has the lower elasticity and the only large reserve. Small deficits are closed on the cost side. Large ones require difference.**

8.5 Combinations, and why the carried case needed two levers

Because $D = 2.885390\,(c/\eta - F)$ is a difference of two independent terms, **any set of parameter moves whose octave contributions sum to the deficit will close it, and there is no interaction term between the geometry side and the cost side.** Combinations can therefore be planned additively — a convenience the formalism rarely offers.

u	η	c	D	Q_∞ at $\delta = 0.09$	verdict
0.068735	0.44	0.40	$+0.998363$	-1.6916	DECAYS
0.300000	0.44	0.62	-0.142896	$+0.2421$	accumulates, barely
0.692373	0.30	0.62	-0.250931	$+0.2899$	accumulates, barely
0.300000	0.44	0.40	-1.585591	$+2.6866$	accumulates
0.400000	0.40	0.50	-1.259181	$+1.9396$	accumulates
0.500000	0.36	0.50	-1.393393	$+1.9317$	accumulates
0.692373	0.44	0.40	-3.590989	$+6.0844$	the repaired case

The first row is the one to sit with. It is the cost-side intervention *on its own* — η up, c down, the geometry untouched — and it leaves the pair at $+0.998363$ octaves, still decaying, with $Q_\infty = -1.6916$. **The entire cost-side intervention applied alone does not save the carried dyad**: it closes 3.340057 of the 4.338420 octaves and stops. The second and third rows are the mirror image: either lever alone, pushed to a plausible limit, clears the threshold by a fifth of an octave, **which is not a margin anyone should plan against.**

The pattern is general enough to state as guidance rather than as a theorem: on a deficit above roughly two octaves, single-lever plans tend to land near the hold line, where the relation pays for itself exactly and banks nothing, and where any increase in cost puts it back under. The corpus's launch-budget work makes the same structural point about barriers [Curzi, 2026p], whose Result One gives departure as feasible only where the fund's fixed point exceeds the barrier, after a computable number of periods.

8.6 The convergence clock

Two parties who spend time together converge. Their representations move toward one another, which is a fall in u, which is a fall in f_{rel}. **Nothing about that is pathological — it is most of what shared time is for — but it runs continuously and it runs against the geometry lever.** Modelling the drift as geometric decay at rate λ per period (M-c) and using the hold line $u^* = (e^{c/\eta} - 1)/(N - 1)$, the crossing time is $t^* = \ln(u^*/u_0)/\ln(1 - \lambda)$. Computed for the repaired dyad, where $u_0 = 0.692373$ and $u^* = 0.134733$:

λ	t^* (periods)	D at $t = 10$	D at $t = 20$	D at $t = 40$
0.01	162.863	-3.3362	-3.0846	-2.5922
0.02	81.020	-3.0821	-2.5872	-1.6483
0.03	53.738	-2.8290	-2.1013	-0.7797
0.05	31.911	-2.3271	-1.1755	$+0.6539$
0.08	19.631	-1.5912	$+0.0527$	$+1.9308$
0.10	15.535	-1.1171	$+0.7320$	$+2.3153$
0.15	10.072	-0.0201	$+1.8767$	$+2.5903$

A repaired dyad has a computable shelf life, and at a 5 per cent per-period convergence rate the carried case's repair is spent in 31.9 periods. The parties do nothing wrong in that interval; they simply agree with each other more each period, and the reading walks back across zero. **Two practical consequences:** a geometry intervention is not a one-off and its decay schedule is estimable from two readings and one logarithm; and **the cost-side levers do not decay this way**, which is the reason that generalises best for reaching for them first.

The honest limit on all of it: λ is a modelling assumption, badged as such, and geometric decay is the simplest form behaving correctly at the boundaries rather than one that has been

fitted to anything. **A pair that measures u twice can estimate λ directly, and should, rather than adopting a figure from this table.**

8.7 Procedure, gates, and reporting

```
STEP 1  DECLARE THE RULE SET, DATED, and which routes will be run.
STEP 2  OBTAIN CONSENT FROM BOTH PARTIES, SEPARATELY, with decline available
        at every step and no reason required.
STEP 3  FIX ONE SITUATION, in one written sentence both parties accept.
STEP 4  SILENT WRITE-FIRST, INDEPENDENTLY AND OUT OF CONTACT.
STEP 5  READ EACH FRAMING against the twelve components; record the LIST.
STEP 6  ELICIT INTENSITIES BY MAGNITUDE ESTIMATION, separately, with the ratio
        question asked in words.
STEP 7  RUN THE OVERLAP GATE. If any pair is disjoint: DEGENERATE, publish nothing.
STEP 8  RECORD BOTH RAW L2 NORMS AND THEIR RATIO, then normalise.
STEP 9  ROUTE A: cos = x . y ; u = 1 - cos^2, carried to 12+ decimal places.
STEP 10 ROUTE B, if an item set exists, AGREED BEFORE JUDGING: u = 4a(1-a);
        report m and the 4/m resolution.
STEP 11 ROUTE C, if ground truth will arrive: u = 1 - rho^2; REPORT THE SIGN OF rho.
STEP 12 RUN THE ACCURACY GUARD on any Route B or C reading.
STEP 13 RUN THE COLLISION TEST: states and distinct pitch classes, separately.
STEP 14 COMPUTE F, N_eff, f_rel, cents above k, folded value and j.
STEP 15 RUN THE CONCORDANCE PASS. Publish every route; DO NOT AVERAGE SILENTLY.
STEP 16 SECOND ANALYST, working from the written framings alone.
STEP 17 CEILING TEST FIRST, then f_min, D, Q_inf, and the sensitivity.
STEP 18 SIGN THE NUMERATOR before choosing any lever.
STEP 19 RESERVE CHECK before committing to a cost-side plan; declare eta_max.
STEP 20 PUBLISH the decomposition, the norms, the routes and the bounds on the
        same page as the frequency. Where they will not fit, the frequency is
        not published either.

G1  OVERLAP.   supports share >= 1 state with weight >= 10% of each party's
               maximum entry.  Otherwise DEGENERATE; no f_rel.
G2  MAGNITUDE. intensities magnitude-estimated with ratios asked explicitly.
               Otherwise publish the ORDER only. NO ORDINAL FALLBACK.
G3  NORMS.     both raw L2 norms and their ratio recorded.
G4  ROUTES.    at least two routes where two are available; concordance reported.
G5  ACCURACY.  Route B or C readings publish both parties' accuracy.
G6  COLLISION. distinct-class count published beside distinct-state count.
G7  CEILING.   before any verdict, c/eta <= ln N checked FIRST.
G8  BRANCH.    no unison reading names a branch unless a signed route was run.
G9  CONSENT.   either party may decline any step at any time; the reading is not
               completed against a party who has declined it.
```

G2 and G9 have no workaround and no partial form. The first is epistemic — without it the number is not determined. **The second is structural — without it the number should not exist.**

> **THE CONCORDANCE RULE.** Report every route run, with its u, its f_{rel} and its resolution limit, and report the spread between extremes in cents. **If the spread exceeds 100 cents, publish the interval and no point value; if it exceeds 600 cents, treat the routes as measuring different things and say which, rather than reconciling them.** The thresholds are house conventions and are stated as such (M-d).

Three failure modes deserve naming because each is the sole defence against a specific error. *Averaging the routes* destroys the information that three numbers carry, and hides the

case the rule exists to catch — **routes that diverge are telling you they are measuring different things, and the divergence is the finding.** *Choosing the route after seeing the numbers* is a search, not a measurement; **the route set is fixed before the data exists or the reading is not a reading**, which is the standard preregistration discipline [Nosek et al., 2018, Simmons et al., 2011] applied to a two-party instrument. And *eliciting the two parties together* measures the elicitation session rather than the pair, because **two people in a room converge before they record.**

9 Worked example: the carried dyad

CONSTRUCTED CASE, labelled as such. These are not two people. Every parameter is chosen, nothing here is an assessment of anyone, and no clinical, therapeutic or counselling claim is made or implied.

9.1 The measurement

```
A = { 4: 0.7,  6: 0.5,  9: 0.4 }          ||A||_2 = 0.9486832980505138
B = { 4: 0.6,  6: 0.55, 9: 0.3,  10: 0.2 } ||B||_2 = 0.8902246907382427
norm ratio = 1.065667
```

Overlap gate. Shared support $\{4, 6, 9\}$ — three of A's three states and three of B's four. The smallest shared weight is B's 0.3 at state 9, which is 50 per cent of B's maximum entry, far above the 10 per cent floor. **PASS.** The two parties hold their states at total intensities within 6.6 per cent of each other, **so nothing in this reading is an artefact of one party reporting larger numbers** — which is exactly the check the magnitude reporting rule exists to make possible.

```
ROUTE A   cos theta = 0.965020899        u = 0.068734665264634
ROUTE B   m = 120, a = 118/120 = 0.983333333333, cos = 0.966666666667
          u = 4a(1−a) = 0.065555555556   resolution 4/m = 0.033333
ROUTE C   m = 24, rho = +0.9650 (SIGN RECORDED)   u = 1 − rho^2 = 0.068775
```

Collision test. Union $\{4, 6, 9, 10\}$, odd parts $1, 3, 9, 5$ — **four states and four distinct pitch classes, no collisions.** The counterfactual B^* of Section 5.1 is identical in u and f_{rel} and carries only three; **this dyad's breadth survives the test and the counterfactual's does not, and no quantity in the geometry distinguishes them.**

```
F     = ln(1 + 11 * 0.068734665264634) = 0.563084802746 nats
N_eff = 1.7560813179
f_rel = 1.4907987150 Hz     folded 381.644 Hz (j = 8)
cents above k = 1949.66     of an available 8603.91
realised pitch fraction = 0.563084802746 / 2.484906649788 = 22.66%
```

Concordance.

```
ROUTE A   u = 0.068734665264634   f = 1.4908 Hz   1949.66 cents
ROUTE B   u = 0.065555555555556   f = 1.4320 Hz   1880.02 cents
ROUTE C   u = 0.068775000000000   f = 1.4916 Hz   1950.54 cents
spread = 70.52 cents           verdict: CONCORDANT (< 100 cents)
```

And the explanation of the spread is more useful than the spread. Route A's cosine implies, through the agreement identity, $a = (1 + \cos\theta)/2 = 0.982510$, or 117.90 items in 120. **The only realisable log is 118, and that 0.10-item rounding is worth 69.64 cents of the observed 70.52.** Nearly the whole discrepancy is the agreement channel's resolution limit, not a disagreement between routes.

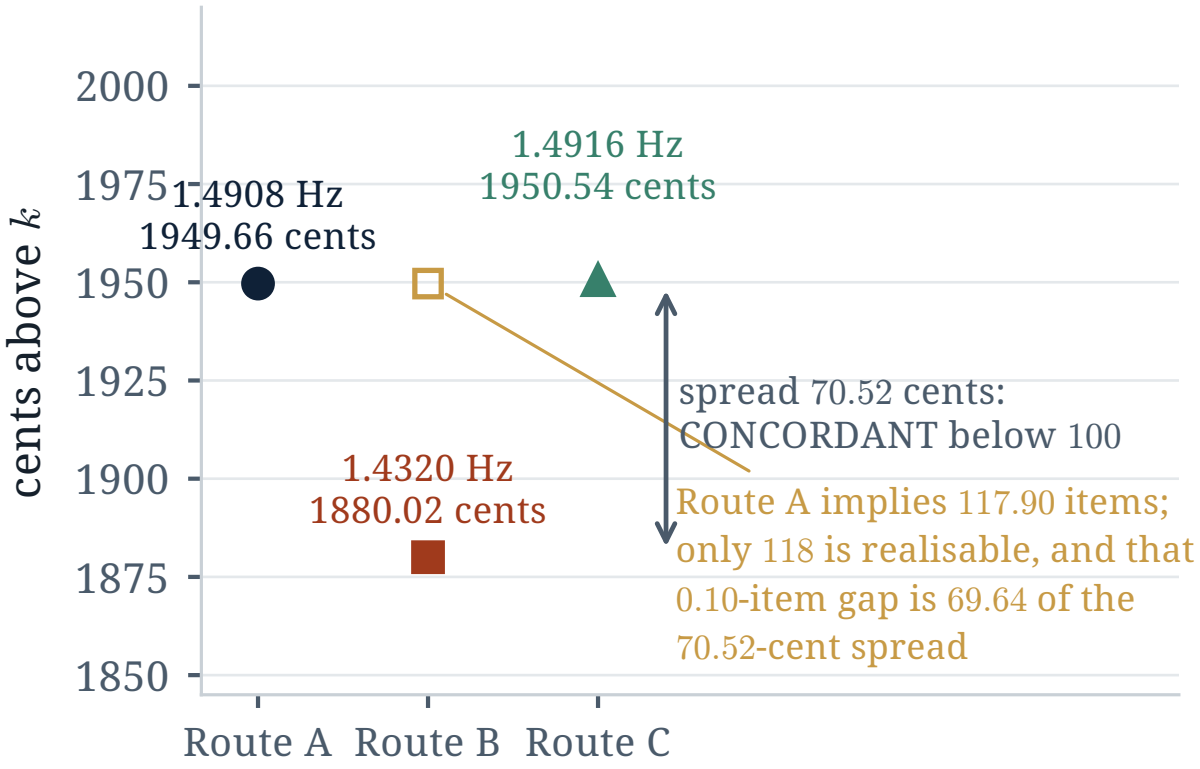


Figure 44.5: **The three routes agree, and the residual disagreement is the coarsest route's own quantisation**: Route A's cosine implies 117.90 matching items in a 120-item log, only 118 is realisable, and that 0.10-item gap accounts for 69.64 of the 70.52-cent spread. Computed from the constructed dyad with the **defined** constant $k = 0.483425732988414$ and $N = 12$; the concordance threshold of 100 cents is a declared convention. Built by `figs/src/CCP-044_fig02.py`.

9.2 The verdict

```
CEILING TEST FIRST:   c/eta = 0.62/0.30 = 2.066666666667
                      ln 12 = 2.484906649788      margin +0.418240 nats
                      2.0667 <= 2.4849   ->  TUNABLE

f_min  = k e^(2 * 2.066667)  =  30.1587698405 Hz   (folded 241.270 Hz, j = 3)
D      = 2.885390081777927 * 2.066667 - log2(1.4907987150/k) = +4.3384201973 octaves
Q_inf  = (0.30 * 0.563084802746 - 0.62) / 0.09  =  -5.0119395464
identity check: -2.885390 * (0.09/0.30) * (-5.0119395464) = +4.338420197298  err 8.88e-16
numerator eta F - c = -0.451074559176            NEGATIVE
VERDICT: DECAYS. Short by 4.3384 octaves = 5206.10 cents.
```

Two readings of the same verdict, both worth stating. The pair fields $N_{\text{eff}} = 1.7561$ distinct perspectives and must field $e^{c/\eta} = 7.8985$ to break even — **they are at 22 per cent of the geometry they need.** And the pair sits 1949.66 cents above the ground essence where 5206.10 cents more is required.

The interpretation the number does NOT license, stated because it is the one a reader reaches for: this pair is not failing at anything, is not incompatible, and is not in difficulty. **The verdict is that at these costs and this geometry, the arrangement does not accumulate order. It is a statement about an arrangement.** Whether the arrangement should change, and whether either party wants it to, are questions the formalism does not touch and this paper does not answer.

The lever audit. Required: $F \geq c/\eta = 2.066667$, i.e. $N_{\text{eff}} \geq 7.898451$.

```
L1  RAISE u    need u >= 0.627132 from 0.068735 -- a factor of 9.1240
               reserve 5.545206 >= 4.338420                          REACHABLE
L2  RAISE eta  need eta >= c/F = 1.101077; reserve to eta = 1.00 is 4.174198
               ABSENT. Not expensive -- absent.
L3  CUT c      need c <= eta F = 0.168925 from 0.62 -- a cut of 72.8%
               reserve 5.963140                                      REACHABLE
L4  RAISE N    need N >= 101.3635 -> 102 categories; a new category buys 0.1129
               octaves at this u                                     NOT REACHABLE
```

The proportionally strongest lever is the one that does not exist, and the lever with the reserve is the slow one. That combination is ordinary, and it is why the ceiling classification matters more than any single sensitivity: the case is TUNABLE, so a combination exists.

9.3 The demonstration

The intuitive repair is to add trust. Assume it succeeds completely and δ falls from 0.09 to 0.05 — a 44.4 per cent reduction in per-period loss, a large and hard-won change.

```
BEFORE   Q_inf = -0.451074559176 / 0.09 =  -5.0119395464
AFTER    Q_inf = -0.451074559176 / 0.05 =  -9.0214911835      WORSE by 80.0%
DEFICIT BEFORE  +4.3384201973
DEFICIT AFTER   +4.3384201973                                  UNCHANGED, every digit
```

A negative number divided by a smaller positive number is more negative. Trust did not add order; it amplified the sign the relation already had, and the sign was negative. The mistake is not naivety: **the intervention was correct about the parameter and wrong about the order, and there is no way to see that from the deficit alone.**

9.4 Three partners, one comparison

CONSTRUCTED. Party A held fixed with $\eta = 0.30$, $c = 0.62$, $\delta = 0.09$, and the costs held constant across all four partners so that the comparison isolates the geometry:

partner	u	f_{rel} (Hz)	D	Q_∞
B, as found	0.068734665264634	1.4908	$+4.3384$	-5.0119
the sycophant, $\rho = +1$	0.000000000000000	0.4834	$+5.9631$	-6.8889
the contrarian, $\rho = -1$	0.000000000000000	0.4834	$+5.9631$	-6.8889
genuinely distinct	0.692372881355932	35.8882	-0.2509	$+0.2899$

Rows two and three are identical in every column, and that is the opposition identity. One partner echoes A on every item; the other inverts A on every item; the relation is at the ground essence either way, short by 5.9631 octaves either way, and arrives at $Q_\infty = -6.8889$ either way.

Both faults are worse than the pair as found. The dyad at $u = 0.068735$ is already decaying, and replacing B with either extreme costs a further 1.6247 octaves of deficit and 1.8770 of Q_∞. **The near-parallel pair the framework calls marginal is better than either perfectly correlated alternative, in either sign.**

And the fourth row is unreachable by opposition, which is the whole argument. The distinct partner gets there by taking on states A does not hold — a change in what that party is doing. **No amount of disagreeing with A about the states they share moves the reading in that direction.**

The comparison above holds c fixed, which flatters the contrarian. Manufactured opposition is not free — the argument takes time, positions must be restated, decisions get relitigated — so c rises:

```
c         c/eta       ceiling verdict     D (at u = 0)     Q_inf
0.62    2.066667      TUNABLE              +5.9631        -6.8889
0.70    2.333333      TUNABLE              +6.7326        -7.7778
0.745472 2.484907     STRUCTURAL           +7.1699        -8.2830
0.90    3.000000      STRUCTURAL           +8.6562       -10.0000
```

The tipping point is exact: $c_{\text{crit}} = \eta \ln N = 0.30 \times 2.484906649788 = 0.745472$. A rise in coordination cost from 0.62 to anything above 0.7455 — a 20 per cent increase — moves this dyad above the ceiling, **at which point no assignment of u is sustainable and the geometry lever is gone entirely.** So manufactured opposition does two separate kinds of damage: it sets the numerator to zero, and it raises the cost ratio until the only lever that could have repaired the numerator is removed. **The second harm is worse than the first and it is invisible to anyone who computes a deficit without running the ceiling test.**

9.5 The seventh seat

CONSTRUCTED extension. A commons of six members drawn from one doctrine: fifteen pairwise links at $\kappa = 0.055$ each, $\eta = 0.34$, $\delta = 0.09$, mean pairwise $u = 0.018614$.

```
THE COMMONS AS FOUND
15 pairs, mean u = 0.018614       F = 0.186275 nats
f_rel = 0.7017 Hz    644.97 cents above k    realised fraction 7.50%
C = 0.055 x 15 = 0.8250    c/eta = 2.426471
f_min = 61.9349 Hz   D = +6.4638 octaves   Q_inf = -8.4630    DECAYS
```

```
THE PARALLEL WALL, PRICED PER LINK
paid        C = 0.8250 across 15 links        delivered  644.97 cents of pitch
yield       43.00 cents per link
ceiling for 15 links at u = 1   8603.91 cents = 573.59 cents per link
realisation                     644.97 / 8603.91 = 7.50%
```

The commons is paying for fifteen links and collecting 7.50 **per cent of what fifteen links could deliver.** The other 92.50 per cent is not lost to inefficiency or bad management. **It is not there. Nobody holds it.**

configuration	links	C	c/η	ceiling	mean u	f (Hz)	D	Q_∞
six, as found	15	0.8250	2.4265	TUNABLE	0.018614	0.7017	$+6.4638$	-8.4630
+ contrarian	21	1.1550	3.3971	STRUCTURAL	0.013296	0.6352	$+9.4080$	-12.3177
+ sycophant	21	1.1550	3.3971	STRUCTURAL	0.013296	0.6352	$+9.4080$	-12.3177
+ distinct, $u = 0.80$	21	1.1550	3.3971	STRUCTURAL	0.241867	6.4777	$+6.0577$	-7.9313
+ independent, $u = 1$	21	1.1550	3.3971	STRUCTURAL	0.299010	8.8933	$+5.6005$	-7.3326

Every row that adds a seat is structural, whoever fills it. That is the finding, and it is not about the contrarian at all — it is about the per-link cost. Twenty-one links at 0.055 give $C = 1.1550$ and $c/\eta = 3.3971$, which exceeds $\ln 12 = 2.484907$; **the maximum** κ **at which a seventh seat stays under the ceiling is** 0.040232**, so the seat is affordable only alongside a** 26.9 **per cent cut in per-link cost.**

Within that finding the contrarian is still the worst available choice: against the six-member base, the contrarian and the sycophant each contribute -172.35 cents, the genuinely distinct candidate $+3847.96$, and the perfectly independent candidate $+4396.66$ — **three candidates at the same price and a spread of** 4569.01 **cents between worst and best, three and three-quarter octaves for the same money.**

And a geometry improvement computed above the ceiling is arithmetically correct and practically void. The distinct candidate improves D by 0.4061 octaves on a case the cost ratio has already placed beyond repair. **Seeing that happen to the *good* candidate is more instructive than seeing it happen to the bad one.** The corpus's pricing work reaches the same discipline from the other side [Curzi, 2026], whose Result One prices an uncertain position as $E[\Delta h] = 0.0318401p - 0.0051096$ with a break-even probability of 16.05 per cent, and [Curzi, 2026], whose Result One shows a 25.17 per cent local-currency allocation cancelling an entire $\eta \ln N = 0.377583$ diversification gain. **In each case the cost side removes the benefit the geometry was proposed to supply.**

9.6 The repair, correctly ordered

CONSTRUCTED. Two changes, made together, both on the numerator: one party takes on two states the other does not hold — **a change in what that party is actually doing, not a rhetorical adoption of a contrary position, which the opposition identity forbids from working** — and the pair adopt a decision log, which raises η and lowers c.

```
B' = {2: 0.95, 4: 0.60, 6: 0.55, 9: 0.30, 10: 0.20, 12: 0.8391}
     raw L2 norm 1.548899

u     0.068734665264634  ->  0.692372881355932
eta   0.30 -> 0.44        c  0.62 -> 0.40        delta 0.09 (deliberately unchanged)

F                 = 2.153632742815 nats     N_eff = 8.6161016949
f_rel             = 35.8881768941 Hz        folded 287.106 Hz (j = 3)
octaves present   =  6.2140705559
```

```
c/eta             = 0.909090909091          ceiling margin +1.575816 nats
f_min             =  2.9782153324 Hz        octaves required 2.6230818925
OCTAVE DEFICIT    = -3.5909886634           numerator eta F - c = +0.547598406839
Q_inf (delta 0.09)= +6.0844267427
identity check    : -2.885390 * (0.09/0.44) * 6.0844267427 = -3.590988663387
VERDICT: ACCUMULATES
```

A discrepancy in the source, printed rather than silently corrected. The realising vector B', recomputed, gives $\cos\theta = 0.5546425597$ and $u = 0.6923716310$, against the declared $u = 0.692372881356$ whose implied cosine is 0.5546414325. **The vector realises the declared u only to six decimal places**, and the frequency differs by 0.0055 cents — far below any reading's resolution, but a reader recomputing from the printed vector will find it. This paper quotes the declared value, which is the one every other paper in the series carries.

Now trust helps. Repeating the identical intervention that made things worse before:

```
delta 0.09 -> 0.05
Q_inf   +6.0844267427  ->  +10.9519681368            BETTER by 80.0%
deficit -3.5909886634  ->   -3.5909886634            unchanged, as always
```

The same lever, the same magnitude, the same direction, the same two parties, and the opposite result — because the numerator changed sign in between.

The honest note on where the repair leaves them. Pitch headroom falls from 5.5452 octaves to 0.955854. **The repair worked by spending most of the range that twelve categories make available**, and the next order of gain is not available from u at this N; it would require genuinely new categories rather than finer slicing of what is there. **That is not a defect in the repair. It is the ceiling doing what a ceiling does, and the value of the reading is that the parties can see it coming rather than meeting it as an unexplained plateau.**

10 Results

The eight results below are the units other papers in this corpus cite. A citation to this paper should name one of them. Each carries its status: **DERIVED** where it follows from the stated definitions by algebra, **COMPUTED** where it is the output of exact arithmetic on stated parameters, and **POSTULATED** where it is asserted by this framework.

R1 — the opposition identity. $u = 1 - \rho^2$ depends only on $|\rho|$, so $\rho = +1$ and $\rho = -1$ both give $u = 0$, $F = 0$, $N_{\text{eff}} = 1$ and $f_{\text{rel}} = k = 0.4834$ Hz — the same folded note at $j = 9$, in the same pitch class. *Status:* DERIVED (Theorem 6.2). **Every quantity computed downstream of u inherits the evenness, so no care with the frequency recovers the sign.**

R2 — the agreement identity. For m binary judgments with agreement fraction a, $\cos\theta = 2a - 1$ and $u = 4a(1 - a)$, exact and requiring no centring, symmetric about $a = 0.5$ with resolution bounded by $4/m$. *Status:* DERIVED (Theorem 4.1); the resolution table COMPUTED. **In a 120-item log one disputed item is worth 643 cents.**

R3 — the three-route concordance. On the constructed carried dyad, $u = 0.068735/0.065556/0.068775$ and $f_{\text{rel}} = 1.4908/1.4320/1.4916$ Hz, a spread of 70.52 cents against a declared 100-cent threshold: CONCORDANT. **Of that spread, 69.64 cents is the agreement channel's own quantisation** — a better explanation than any story about the parties. *Status:* COMPUTED.

R4 — the unison fixed point. At $u = 0$, $Q_\infty = -c/\delta$, in which η and N do not appear and **two of the four levers have partial derivative exactly zero.** *Status:* DERIVED (Theorem 6.3); verified COMPUTED across a grid spanning a twentyfold range in η and a fiftyfold range in N.

R5 — the elasticities. $dD/d\ln\eta = -2.885390\,(c/\eta)$ and $dD/d\ln u = -2.885390\,(1-1/N_{\text{eff}})$, so the geometry lever is bounded above by 2.885390 octaves per e-fold in every system, and **cost-side levers dominate whenever** $c/\eta \geq 1$ (Corollary 8.3a). *Status:* DERIVED.

R6 — the lever reserves. On the carried dyad, $R_u = 5.545206$ and $R_c = 5.963140$ octaves each close the $+4.338420$ deficit; R_η at a physically impossible $\eta_{\max} = 1.00$ supplies only 4.174198 and **cannot**, falling short by 0.164222. *Status:* DERIVED (Theorem 8.4); COMPUTED on a constructed case. **The lever with 4.80 times the elasticity has less reserve than the deficit requires.**

R7 — the damage integral. Eight periods of the wrong lever order cost twelve periods of permanent lag, a minimum stock of -2.373029 from a starting $+1.000000$, thirteen periods below the starting value, and 111.133601 stock-periods of cumulative gap — **with the deficit reading $+4.338420$ throughout.** *Status:* COMPUTED; the order-independence of the destination and the order-dependence of the path DERIVED (Theorem 8.2).

R8 — blindness to collision. u is blind to pitch-class collision, exhibited by two dyads with identical u and f_{rel} to fifteen decimal places and four versus three distinct pitch classes. *Status:* DERIVED and COMPUTED (Theorem 5.1).

CR — counter-result. Perfect opposition and perfect agreement are the same reading, so the standard remedy of adding a contrarian contributes zero surplus while adding cost; **and at $u = 0$ two of the four levers are exactly inert.** To these the framework adds four of its own: the constant k is chosen; the affine effective count is not a Hill number; the interpretation of k as a frequency is postulated; **and the framework's own native encoding cannot see the opposition branch at all**, so the sharpest result in this paper is invisible to the instrument the corpus uses by default. The corpus's uncertainty node models the required discipline [Curzi, 2026], whose counter-result requires a user to be able to distinguish the theorem, the ordinary observer effect and the framework's own conjecture in three separate sentences on demand.

11 Discussion

11.1 What follows

The received wisdom names a real target and points at the wrong coordinate. "Creative tension", "healthy disagreement" and "we need someone to push back" are among the most durable pieces of advice about how two or more people should work, and **the target they name is independence.** Opposition is independence's mirror image, sitting exactly as far from it as agreement does. **The job of the arithmetic is not to be contrarian about contrarians; it is to say which implementations hit the target.**

The two faults read the same and are not the same, and the distinction decides the remedy.

```
THE PARALLEL WALL          seats already filled, near-parallel views.
                           eta*F = 0.063334, C = 0.8250, net = -0.761666
                           THE WASTE IS SUNK.  Remedy: substitution.

MANUFACTURED OPPOSITION    a NEW seat, at a NEW cost, returning nothing.
                           eta*F = 0.046409, C = 1.1550, net = -1.108591
                           THE WASTE IS INCURRED, and it is ongoing.
                           Remedy: refusal.
```

A report that says only "low u" has told the reader nothing about which of the two decisions is in front of them.

And the reserve result generalises past this framework. The structure is that *the fastest lever per proportional unit and the lever with the most room are frequently different levers*, and any model in which one parameter enters a ratio and another enters a bounded logarithm will have the property. **A plan built on the strongest lever that requires more of it than exists is not an expensive plan; it is not a plan.**

11.2 What does not follow

A high reading is not good news by itself. It may be a wide pair; it may equally be disjoint support, two parties answering at random, or an encoding in which each framing was reduced to its single most characteristic state. **Near the top of the range the first hypothesis to test is that the encoding failed**, because that is the more common cause and it is the one that flatters.

A unison reading is not a judgment. It says that at these costs the arrangement accumulates no order. **It does not say the two parties have no relationship, that their relationship is poor, or that anything about it should change.** Two people in unison may be exactly as they wish to be, and the formalism has no view.

"Sycophant" and "contrarian" name coordinates, not people. A party sits at a correlation with respect to a specific counterparty on a specific item set, and moves when either changes. **Transferring the label from the coordinate to the person converts an instrument into an insult**, and the resulting number is no longer about anything measurable.

Nothing here says conflict is good or that a relation should be made harder. The quantity being maximised is represented non-parallelism, and **independence — not friction — is the coordinate that maximises it.** Two parties who work on different things, calmly, are at high u. Two parties who fight about one thing are not.

A DECAYS verdict is not robust in the way an ACCUMULATES verdict is. F is a lower bound on the pair's information, so f_{rel} is a lower bound on the pitch: **a verdict of ACCUMULATES survives the possibility that the representation missed something; a verdict of DECAYS does not, because the missing something could have supplied the surplus.**

11.3 The reflexive exposure, and the programme

An instrument that issues a verdict on a relation may change the relation by issuing it. The mechanism is documented [Merton, 1948, Rosenthal and Jacobson, 1968] and has been observed in the field [MacKenzie and Millo, 2003], but the effect sizes are contested [McCambridge et al., 2014] and on the best long-run assessment **self-fulfilling prophecies are real but small, do not typically accumulate, and often dissipate** [Jussim and Harber, 2005]. **The correct response is not to suppress the reading but to publish the parameters that produced it**, which is also why any indicator adopted as a target degrades as an indicator [Campbell, 1979, Goodhart, 1975].

The elicitation has a self-concealing failure of its own. An elicitation run by an analyst who has views produces vectors that drift toward those views, and at dyad scope that lowers u and produces exactly the unison reading the analyst will then interpret as a finding about the pair. **A single analyst cannot detect their own convergence pressure, because the artefact and the finding have the same shape.** The mirror failure is symmetric: an analyst who wants a *high* reading can lead the elicitation toward difference and will succeed. **Neither direction is detectable from inside a single pass**, which is why the second analyst is a numbered step and not a recommendation.

The wider programme begins with *physique sociale* [Quetelet, 1835] and continues through the modern computational literature [Castellano et al., 2009, Lazer et al., 2009]. **Its characteristic error is Quetelet's** — treating a summary statistic as a type — and this paper's version of it is the mean pairwise u: a scalar summary of pairwise angles does not determine a joint arrangement, and the corpus's roster paper makes that failure visible [Curzi, 2026], whose Result Seven records a constructed team in which **every one of five drop-one removals improves the reading**, so a favourable row carries no information about the member it names. Three further constraints bear and are not comfortable. Human activity is bursty with **no characteristic time scale** [Barabási, 2005], which a convergence clock must meet by declaring its period rather than discovering it. An explanation is supposed to name a mechanism rather than exhibit a fitted curve [Hedström and Ylikoski, 2010]; **this paper fits nothing, which sidesteps the objection at the cost of having no empirical content to fit.** And **what observation would distinguish this model from a neighbouring one that fits equally well?** [Sobkowicz,

2020, Durlauf, 2005] The honest answer is H3 and H4, and neither has been tested. **What the paper is entitled to claim is analytical rather than predictive** [Epstein, 2008, Keuschnigg et al., 2018], and making point predictions with constants rather than directional ones is the right response to the directional-prediction critique [Meehl, 1990].

12 Verification

Every step below can be checked on its own. The standard is Occam's: **a postulate that can be dropped without the result changing has been dropped.** Statuses: **ARITHMETIC** (a reader can redo it), **DEFINITION** (true by stipulation), **CITED** (an external source supports it — the key is named), **POSTULATED** (asserted by this framework, no external support).

12.1 The postulates

	postulate	status
P1	The Pythagorean whole tone is the ratio $9{:}8$.	CITED [Barbour, 1951]
P2	φ is the positive root of $x^2 = x + 1$.	CITED [Euclid, 1956, Vajda, 1989]
P3	A dimensionless ratio of the form $(9/8)^2/\varphi^2$ **is a frequency in hertz**.	POSTULATED — no source
P4	An effective count N_{eff} has essence frequency $k\,N_{\text{eff}}^2$.	POSTULATED — no source
P5	$N_{\text{eff}} = 1 + (N-1)u$ (the affine count axiom).	POSTULATED — design axiom [Curzi, 2026]; **not a Hill number**
P6	$u := 1 - (x \cdot y)^2$ on unit vectors.	DEFINITION
P7	A party's construal admits representation as a **non-negative** unit vector.	POSTULATED — and Section 6.5 states what it cannot represent
P8	Stored order obeys $Q_{t+1} = (1-\delta)Q_t + \eta F - c$ with $\delta > 0$.	CITED as a form [Oppenheim et al., 1996, Elaydi, 2005]; the *interpretation* POSTULATED
P9	Elicited intensities carry ratio information up to a free common multiplier.	**ASSUMED**; the method CITED [Stevens, 1957], not tested here
P10	c is expressed in nats of surplus per declared period, so c/η is dimensionless.	**CONVENTION — stated here, absent from the source**

Dropped as unnecessary. P3 and P4 are not needed for R1, R2, R4, R5, R6 or R8. Every one of those is a statement about u, about ρ, about a ratio of logarithms in which k cancels, or about a difference equation; the frequency interpretation adds hertz values and nothing else. No postulate about consonance, about octave folding or about rendering is needed anywhere in this paper. **P9 is needed only where a number is claimed; the *ordering* of readings survives without it, and Section 4.3 says exactly how much is lost.**

12.2 The compounding chain

step	consumes	produces	status
L1	P1, P2	$k = (9/8)^2/\varphi^2 = 0.483425732988414$	ARITHMETIC (exact)
L2	L1, **P3, P4**	$f_{\text{rel}} = k\,N_{\text{eff}}^2$ in Hz	**POSTULATED**
L3	P5, P6	$F = \ln(1+(N-1)u)$, $F \le \ln N$	DEFINITION
L4	P6	$u = 1-\rho^2$ is EVEN in ρ — **R1**	ARITHMETIC
L5	L4	$u(\pm 1) = 0$, $F = 0$, $N_{\text{eff}} = 1$, $f_{\text{rel}} = k$ — **R1**	ARITHMETIC
L6	P6, ± 1 judgment vectors	$\cos\theta = 2a-1$, $u = 4a(1-a)$ — **R2**	ARITHMETIC
L7	L6, quantisation of a at $1/m$	resolution $\approx 4/m$; 643 cents per item at $m = 120$ — **R2**	ARITHMETIC
L8	L2, L6, constructed data	three routes at 70.52 cents spread — **R3**	ARITHMETIC
L9	L8, L7	69.64 of the 70.52 cents is quantisation — **R3**	ARITHMETIC
L10	P8, L5	$Q_\infty = -c/\delta$ at $u = 0$; η, N absent — **R4**	ARITHMETIC
L11	L10 differentiated	$\partial Q_\infty/\partial\eta = \partial Q_\infty/\partial N = 0$ — **R4**	ARITHMETIC
L12	L2, L3, Def. 3, **P10**	$D = 2.885390\,(c/\eta - F)$	ARITHMETIC given P10
L13	L12	the three elasticities; $1 - 1/N_{\text{eff}} < 1$ — **R5**	ARITHMETIC
L14	L13	$c/\eta \ge 1$ suffices for cost-side dominance — **R5**	ARITHMETIC
L15	L12, limits $u \to 1$, $c \to 0$, $\eta \to \eta_{\max}$	R_u, R_c, R_η — **R6**	ARITHMETIC
L16	L15, constructed parameters	5.545206, 5.963140, 4.174198 vs $+4.338420$ — **R6**	ARITHMETIC
L17	P8	$dQ_\infty/d\delta = -n/\delta^2$: decay amplifies the sign	ARITHMETIC
L18	L17, two orderings	12-period lag, 111.133601 stock-periods — **R7**	ARITHMETIC
L19	P6, two constructed partners	identical u to 10^{-15}, 4 vs 3 pitch classes — **R8**	ARITHMETIC

12.3 What this chain is and is not

Read the status column. Steps L1 and L3 through L19 are arithmetic or definition and a reader can redo every one of them. **Step L2 is postulated, and L12 rests additionally on a convention**

(P10) that the source never states. A chain containing postulated steps is not a proof of its conclusion, and this one is not offered as one.

What it is offered as is a separation, and here the separation is unusually favourable. The paper's sharpest results — R1, R2, R4, R5, R6, R7 and R8 — do not consume P3 or P4 at all. They are statements about an even function, about ± 1 vectors, about a difference equation and about logarithmic derivatives. **A reader who rejects the frequency interpretation entirely keeps every one of them, restated in nats instead of octaves.** What such a reader loses is the hertz values, which is precisely what the postulates were asked to supply.

Two statuses that must not be misread. The exponent 2 at L2 is **DEFINED**, not estimated: no fit was performed and the standard procedure for testing a scaling claim [Clauset et al., 2009] was not run. The logarithm at L3 is **DEFINITION** given the affine axiom P5, and the warrant the argument needs is the multiplicative-dynamics one [Peters and Gell-Mann, 2016] rather than the psychophysical one; the psychophysical warrant is consumed only at P9, for the elicitation.

The falsifiable content lives in H1 through H4, not at L2. This paper's empirical exposure is therefore smaller than its notation suggests, and printing the status of each postulate is the honest description of what a badged postulate does in a research programme [Lakatos, 1970, Platt, 1964, Popper, 1959].

13 Limitations

Every check the governing rigour standard specifies is listed here by name, whether or not this paper passes it.

L-1. The dimensional audit does not close. $f_{\text{rel}} = k\,N_{\text{eff}}^2$ balances only if k carries $[T^{-1}]$, and k is computed as a ratio of two dimensionless quantities. **No operational procedure is supplied by which the constant acquires the dimension of frequency.** It touches only the absolute values; every result in Sections 5, 6 and 8 survives without it.

L-2. The exponent of the threshold is dimensionless only by an unstated convention. c/η carries $[\mathcal{C}T^{-1}]$ divided by a pure number and cannot be an exponent without a convention; **the source never states the convention.** It is stated here (P10), and the measured failure mode is a $90.1\times$ period-normalisation error.

L-3. The constant is chosen, quoted to fifteen figures **because it is defined**, not because anything was measured to that precision [Markowsky, 1992, Fischler, 1981].

L-4. The exponent is defined, not estimated. No fit was performed, no alternative was tested, and the standard scaling procedure [Clauset et al., 2009] was not run because there is no data to run it on. **This paper makes no empirical scaling claim.**

L-5. $N_{\text{eff}} = 1 + (N-1)u$ **is not a Hill number.** It is a linear interpolation between 1 and N and appears in no ecological, economic, political-science or information-theoretic literature; the Hill-number citations apply to $1/\text{HHI}$ and $\exp(H)$ and to nothing else [Hill, 1973, Laakso and Taagepera, 1979].

L-6. Octave reduction beyond pitch is unwarranted [Shepard, 1982]. The operation is well defined; its interpretation is not supported.

L-7. Conjoint measurement was not tested, so the ratio-scale treatment of the intensities, of η, of c and of δ is provisional. Michell's objection [Michell, 1997] stands unanswered for every quantity here except the counts, and the available remedies [Luce and Tukey, 1964, Rasch, 1960] were not run. **Section 4.3 prices the failure at** 1065.05 **cents, which is more than three quarters of an octave.**

L-8. The framework's own encoding cannot see the paper's sharpest result. Route A is non-negative by construction, so the opposition branch is invisible to it. **A limitation stated as one**, and the signed routes are the only place the branch can be measured — which means every reading taken with Route A alone is silent about which of two opposite remedies applies.

L-9. Surplus is not correctness. Two parties answering at random score the ceiling. **The formalism cannot distinguish independence from noise, because both are orthogonality**, and

there is no formalism here that converts distinctness into correctness. The accuracy guard is procedural, not formal, and **no procedural gate defends against an analyst who wants a particular answer**; what the gates do is stop an honest analyst producing a degenerate answer by accident, and make a dishonest one leave a trace.

L-10. Route B's resolution is coarse and its readings are intervals. At $m = 120$, one disputed item is worth 643 cents; **a reading of $u = 0.0656$ from such a log has reported four significant figures of a quantity the log resolves to one.**

L-11. Route C requires ground truth, which is unavailable for most of what two people disagree about, and its sign is discarded by u. **Two parties at $\rho = +0.9650$ and two at $\rho = -0.9650$ produce identical readings and require opposite interventions.**

L-12. Mean pairwise u does not determine the joint geometry. The triangle inequality fails for geometric triples, so N_{eff} is an audited count under a stated postulate and **not a measured dimension.**

L-13. The trust mechanism is an assumption, not a finding. The order-of-operations result depends on trust acting through δ rather than η; the literature cited for trust [Gambetta, 1988, Ostrom, 1990] is about reciprocity, monitoring and institutional form, not decay rates. **Which channel it operates through is not settled here.**

L-14. The convergence rate is a modelling assumption. Geometric decay of u is the simplest form behaving correctly at the boundaries and **has been fitted to nothing.**

L-15. Two thresholds are conventions with no derivation — the 100-cent concordance threshold and the 10 per cent overlap-weight floor. **What is not optional is declaring one before seeing the result.**

L-16. A unison has three possible causes and only two are branches. Genuine parallelism, genuine inversion, and **a failed elicitation, which produces exactly the same number.** A unison is the signature of both parties elicited in one session after discussion, and **publishing a unison without excluding the artefact is the most consequential error available in this paper.**

L-17. Every case is constructed. The dyad, its three alternative partners and the commons are chosen parameter sets. **Their precision is exact because the parameters are exact; that precision does not transfer to a measured case.**

L-18. No model comparison, no statistical claim. Nothing was fitted, so there is no out-of-sample test, no matched null with the same number of free parameters, no information criterion, no p-value, no effect size and no interval. **The corpus's one confirmatory attempt failed its own baseline test** [Curzi, 2026i], which is the relevant precedent for how any fit here should be reported.

L-19. The functional form is underdetermined [Latané, 1981, Briscoe et al., 2006, Stumpf and Porter, 2012].

L-20. The reflexive specification is not identified [Manski, 1993, Shalizi and Thomas, 2011]; the remedy adopted is stipulation of the rule set [Brock and Durlauf, 2001], and the claim is **a modelling position rather than an identified empirical result.** The framework also exempts its own parameters from its own reflexivity [Lucas, 1976, Hacking, 1995], which is **not resolved.**

L-21. There is no conservation law [Mirowski, 1989], validity is claimed only within a specified arrangement [Cartwright, 1999], and **ergodicity is assumed and not argued** [Peters, 2019].

L-22. u is symmetric and psychological similarity is not [Tversky, 1977]; this framework's u assumes metric behaviour.

L-23. Boundary behaviour. $u \to 0$ collapses the reading to k and kills two of four levers; $u \to 1$ gives the ceiling; $\delta \to 0^+$ removes the fixed point altogether — **a pathology, not an ideal, since nothing decays out of the stock**; $\delta \to 1$ makes the store memoryless; $\eta \to 0^+$ makes c/η diverge and **places every relation above every ceiling**, which is a real constraint on the model's domain and is **not examined in the source**; $m \to \infty$ makes Route B exact but is unavailable at any plausible cost; and $N \to \infty$ raises the ceiling without bound, which the coordination-cost literature says is not available in practice [Brooks, 1975, Conway, 1968].

L-24. Discrepancies in the source, printed rather than silently corrected. The realising vector for the repaired dyad gives $u = 0.6923716310$ against the declared 0.692372881356 — agreement to six decimal places, a difference worth 0.0055 cents. The cost-side-only combination computes to $D = +0.998363$ from the full-precision u against the source's printed $+0.998357$ from the rounded one. **Neither is an error in a theorem; both are roundings, and both will be found by any reader who recomputes.**

L-25. A published reading can change what it reads, and the elicitation's convergence pressure is **not detectable from inside a single pass**, in either direction.

14 Conclusion

What is established, and it is arithmetic. The surplus functional depends on the correlation only through its absolute value, so **perfect opposition and perfect agreement are the same reading**: $\rho = +1$ and $\rho = -1$ each give $u = 0$, $F = 0$, $N_{\text{eff}} = 1$ and $f_{\text{rel}} = k = 0.4834$ Hz, folded to the same 247.514 Hz at $j = 9$ in the same pitch class. The agreement channel exhibits the same fact directly through the exact identity $\cos\theta = 2a-1$, $u = 4a(1-a)$, a perfect mirror about $a = 0.5$ whose resolution is bounded by $4/m$ — 643 cents per disputed item in a 120-item log. **At $u = 0$ the fixed point collapses to $-c/\delta$, in which conversion efficiency and category count do not appear and two of the four levers have derivative exactly zero.** Three routes on a constructed dyad return $u = 0.068735$, 0.065556 and 0.068775 for a spread of 70.52 cents, of which 69.64 is the agreement channel's own quantisation. And u is blind to pitch-class collision: two dyads agree to fifteen decimal places while occupying four and three classes respectively.

What is established about levers. Elasticities are dimensionless and therefore comparable, and the geometry lever is bounded above by 2.885390 octaves per e-fold in every system, so **cost-side levers dominate proportionally whenever $c/\eta \geq 1$.** But **the strongest lever is the one that runs out first**: on the carried dyad, raising conversion efficiency to a physically impossible 1.00 buys 4.174198 octaves against a requirement of 4.338420, while raising non-parallelism buys 5.545206 and cutting cost buys 5.963140. And because the decay rate divides the numerator, **an intervention on it multiplies whatever sign the arrangement already has**: eight periods of the wrong order cost twelve periods of permanent lag, a minimum stock of -2.373029 from $+1.000000$, and 111.133601 stock-periods of cumulative gap — **while the deficit reads $+4.338420$ throughout.**

What is proposed rather than shown. That a dimensionless ratio of a tuning interval to the square of the golden section is a frequency; that an effective count has essence frequency $k\,N_{\text{eff}}^2$; that the effective count is affine in the non-parallelism; that a construal admits representation as a non-negative unit vector; and that cost is measured in nats of surplus per declared period. **The first four are postulates of this framework with no external support, the fifth is a convention the source never states, and all five are separable from the results above:** a reader who rejects every one of them keeps the opposition identity, the agreement identity, the unison fixed point, the elasticities, the reserves, the damage integral and the collision blindness, restated in nats and in stock units.

What is conceded. The framework's own native encoding **cannot see the opposition branch at all**, so the paper's sharpest result is invisible to the instrument the corpus uses by default. The surplus formalism **cannot distinguish independence from noise**, because both are orthogonality. And a unison reading is produced as readily by a failed elicitation as by a parallel pair.

What is not claimed at all. No pair of people is claimed to resonate at any frequency named here. The twelve states are not diagnoses. **Nothing in this paper is counselling, therapy, mediation, or a clinical assessment of any person, pair or group**, and where a relationship is causing a person difficulty, the resource is a qualified professional, which this work is not.

15 Declarations

Data and code availability. No dataset was collected for this study and none exists. Every number in the paper is the output of exact arithmetic on parameters stated in the text; every case is constructed and is labelled as constructed at the point of use. The figures are generated by the Python scripts named in their captions, which take their inputs from the tables in the text and are available from the author on request. **No empirical measurement of any human being, any pair or any group underlies any figure, table or number in this paper.**

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the analysis, the figures and the manuscript.

Ethics. No human-subjects research was conducted and no personal data were processed. The carried dyad, its alternative partners and the commons are invented parameter sets and are not derived from, and do not describe, any person. **The procedure this paper specifies would, if run, constitute human-subjects research**, and the specification therefore builds consent into its numbered steps rather than appending it as advice: both parties consent separately, either may decline any step at any time without giving a reason, no reading is completed against a party who has declined it, and **no intervention derived from a reading may remove a person's ability to decline.** That last clause has a literal application here: **requiring a person to hold or argue a position they do not hold is exactly such a removal**, and it is also, by Section 6.8's arithmetic, worthless. Nothing is masked, nothing is concealed, and nothing sits below either party's awareness. This paper also asserts perceptual and psychometric claims — that magnitude estimation delivers ratio-scale intensities and that interval arithmetic in cents is the correct apparatus — **without any experiment behind them**; those claims rest on the cited literature and are listed among the limitations (L-7).

Preregistration. No test reported in this paper was preregistered, because no test was run. The route set, the thresholds and the gates are specified as things to be declared before data exists, which is the same discipline [Nosek et al., 2018]; the corpus's one confirmatory fit was preregistered and **failed its own second falsification rule** [Curzi, 2026i], and that outcome is reported here as it is reported there.

Clinical disclaimer. The twelve state classes are not diagnostic categories. Nothing in this paper is an assessment, a screening instrument, a treatment or advice about any person's health or relationships, and nothing in it may be presented as such.

Chapter 45

Two Octaves per Doubling Against a Quadratic Cost: Roster Ceilings, the Parallel Wall, and When Tuning Must Halt

Abstract. We extend a dyadic sustainability criterion to groups and derive the roster size at which a group stops paying for itself. With per-link cost κ, conversion efficiency η and $r = \kappa/\eta$, coordination cost is $\kappa m(m-1)/2$ while best-case surplus is bounded by $2\log_2 m$ octaves, giving NET$(m) = 2\log_2 m - 1.442695\, r\, m(m-1)$; a dyad is viable in the best case if and only if $r \le \ln 2 = 0.693147$. A second bound, exact and independent of any geometry, is the roster ceiling $m(m-1) \le 2\eta \ln N/\kappa$. Two effective counts must be reported separately, the category-scale $1 + (N-1)\bar{u}$ and the roster-scale $1 + (m-1)\bar{u}$, the latter equal to the participation-ratio form $m + 1 - m/\mathrm{PR}$ exactly (residual zero on a constructed triad). Jensen's gap between $F(\bar{u})$ and the mean of $F(u)$ is 0.1758 nats — 609 cents — on that triad and 0.000001 octaves on a near-parallel case, so it is computed rather than assumed. A constructed six-member commons selected for agreement reads $\bar{u} = 0.018614$ and an effective count of 1.204754 against a required 11.3189, an octave deficit of $+6.4638$; it sits on its roster ceiling, 30 against 30.7225. Every lever swept to its physical limit falls short: perfect conversion leaves $+1.8430$, a 91 per cent cut in link cost leaves $+0.0990$, twelve times the category count leaves $+3.2562$. The procedure halts and issues no prescription; restructuring swings the reading 9.5898 octaves and produces a different group. A drop-one audit on a healthy team improves the reading for every member, which is why such rows carry no information about the members they name. All cases are constructed and no data were collected; the frequency register rests on postulates badged throughout, and every bound survives their rejection because the reference constant cancels.

Keywords: roster ceiling; Brooks crossover; parallel wall; aggregation gap; participation ratio; commons; restructuring; drop-one

Series: C. Frequency, Essence, and Mind **Paper:** CCP-045

1 Introduction

A dyad has one interval. Every result in the preceding papers of this series rests on that convenience: one relationship to read, one non-parallelism u to measure, one number to compare

against one threshold [Curzi, 2026,?]. Add a third party and there are three intervals; add a fourth and there are six. **Relationships do not grow with the roster, they grow with its square**, and the observation that coordination load grows as $m(m-1)/2$ while useful output does not is old, well documented, and independent of this framework [Brooks, 1975, Conway, 1968].

The question this paper asks is what that quadratic does to a logarithmic surplus, and at what roster size the two cross. The answer is exact and it is unfavourable to large groups. Writing κ for the per-link cost per period, η for conversion efficiency and $r = \kappa/\eta$, the best case a roster of m can field is $2\log_2 m$ octaves of surplus against $1.442695\, r\, m(m-1)$ octaves of required pitch, so

$$\mathrm{NET}(m) = 2\log_2 m - 1.442695\, r\, m(m-1) \quad \text{octaves},$$

and a dyad is viable in the best case if and only if $r \le \ln 2 = 0.693147$. A second and stricter bound follows from the framework's own surplus ceiling and involves no geometry at all: $\kappa m(m-1)/2 \le \eta \ln N$, the **roster ceiling**.

Four things follow, and the fourth is the one this paper exists to publish. First, two effective counts live side by side at group scope and must be reported separately — a category-scale count $N_{\text{eff}} = 1 + (N-1)\bar{u}$ and a roster-scale count $M_{\text{eff}} = 1 + (m-1)\bar{u}$, the latter identical to the participation-ratio form $m + 1 - m/\mathrm{PR}$, verified here at a residual of exactly zero. Second, the choice of *where* the averaging happens is worth 0.507396 octaves — 609 cents — on a constructed triad, because F is strictly concave and Jensen's inequality punishes dispersion; and it is worth 0.000001 octaves on a near-parallel commons, so the gap is not a constant to be assumed but a quantity to be computed. Third, an orthogonal *pair* already attains the maximum surplus a twelve-category system admits, so **in surplus terms a group buys nothing a well-chosen pair could not, and pays $m(m-1)/2$ times as much for it.**

Fourth, and negatively: **the instrument is run on a case it cannot help, and the run is reported.** A constructed six-member commons whose members were selected for agreement reads $\bar{u} = 0.018614$ and $N_{\text{eff}} = 1.204754$ against a required 11.3189, an octave deficit of $+6.4638$ and a steady state of -8.4630. Every one of the four levers is computed to exhaustion and every one falls short: perfect conversion at $\eta = 1.00$ still fails at $+1.8430$; a 91 per cent cut in link cost still fails at $+0.0990$; twelve times the category count still fails at $+3.2562$. The procedure halts and issues no prescription. **A diagnostic that only ever recommends action is not a diagnostic**, and the four failed lever tables are the result this paper most wants read.

The paper also prints an ethical consequence it would have been more comfortable to omit. A drop-one audit run on a healthy, comfortably accumulating five-member team improves the reading for **every one** of the five removals, by between 1.97 and 2.80 units of steady state. That is not a fact about any of the five; it is a fact about where a group sits on a concave surplus curve when cost is linear in links. It follows that a favourable drop-one row **carries no information whatever about the member it names**, and the prohibition on reading such rows as rankings is therefore an arithmetic finding rather than a piece of decorum.

Contributions

1. $\mathrm{NET}(m) = 2\log_2 m - 1.442695\, r\, m(m-1)$ and the closed bound $r \le 2\ln m/(m(m-1))$, giving the maximum viable roster at any cost ratio and $r_{\max}(2) = \ln 2$ exactly.
2. The roster ceiling $m(m-1) \le 2\eta \ln N/\kappa$, exact and geometry-free, with the constructed commons sitting on it at 30 against 30.7225.
3. The roster-scale identity $M_{\text{eff}} = m + 1 - m/\mathrm{PR}$, verified at 1.7815458788 with residual 0.00×10^{0}.
4. The aggregation gap quantified at group scope: 0.1758 nats, 0.507396 octaves, 609 cents on a three-member case and 0.000001 octaves on a six-member near-parallel one, with two sweeps showing where it lives.
5. The commons audit: four levers, three governance principles, two nesting topologies and five roster sizes, all computed, all short.

6. The reachability test, which separates a case that is formally tunable from one that is practically not, and the restructure that swings the reading 9.5898 octaves while producing a different group.
7. The drop-one negative result on a healthy team, and the muddiness index $\mu = [m(m-1)/2]/\lceil N/2 \rceil$, which exceeds one from $m = 4$ at $N = 12$.

What this paper does not claim, stated at the front rather than at the back. No dataset was collected. Every case in this paper is a constructed parameter set, labelled as constructed at the point of use, and its precision is exact because its inputs are exact rather than because anything was measured. The frequency register it inherits rests on three postulates of this framework with no external support — that the dimensionless ratio $k = (9/8)^2/\varphi^2$ is a frequency in hertz, that essence goes as kn^2, and that octave reduction may be carried beyond pitch — and the standard deflationary treatment of golden-section claims is cited here deliberately [Markowsky, 1992] and again wherever φ appears. Section 2 badges all of it, Section 9 decomposes the chain into steps a reader can check one at a time, and Section 10 prints every check this paper fails.

2 Hypotheses, definitions, postulates and modelling choices

The rigour standard governing this corpus requires a claim that *could* be wrong to be stated so that it could be, and a claim that could not be wrong to be labelled with which of the three non-empirical kinds it is. **A hypothesis nothing could falsify is not a hypothesis**, and none is manufactured here.

2.1 Hypotheses proper

H1 — the crossover hypothesis. For groups whose per-link maintenance cost and conversion efficiency are measured on one declared period, no roster exceeding the largest m satisfying $2\ln m \geq r\, m(m-1)$ accumulates stored order. **Predicted:** among groups with declared κ and η, those above the bound do not accumulate. **Falsified by:** one group measured above the bound and observed to accumulate over a run long enough for the fixed point to dominate the transient. **Status:** derived from the surplus and cost forms; **not tested here.** No observational data were collected for this paper.

H2 — the roster-ceiling hypothesis. No group with $\kappa m(m-1)/2 > \eta \ln N$ sustains itself under any assignment of pairwise non-parallelism whatsoever. **Predicted:** no counterexample exists at any $\bar{u} \in [0, 1]$. **Falsified by:** one counterexample. The hypothesis is a strict universal and is cheap to kill. **Status:** derived from the normalisation bound $F \leq \ln N$ [Curzi, 2026]; **not tested here.**

H3 — the parallel-wall hypothesis. A group whose members were selected by a criterion that rewards agreement exhibits a mean pairwise non-parallelism low enough that no arrangement of η, c, δ or N restores a positive steady state. **Predicted:** for such groups the required u^* exceeds the achievable range by a large factor, and the lever tables are all short. **Falsified by:** a doctrine-selected group whose deficit is closed by any single arrangement lever. **Status:** demonstrated on a constructed case in §5.4; **not tested against observation.**

H4 — the drop-one non-informativeness hypothesis. In a group operating near the surplus ceiling, removing any member improves the reading, so the sign of a drop-one row carries no information about the member removed. **Predicted:** all m rows of a drop-one audit on a comfortably accumulating group move the reading the same way. **Falsified by:** a group near its ceiling in which some drop-one rows improve the reading and others worsen it, with the split reproducible

across analysts. **Status:** derived from the concavity of F together with linearity of cost in links; **computed** on a constructed team in §5.11b.

H5 — the reachability hypothesis. The formal classification TUNABLE and the practical classification REACHABLE come apart, so a case passing the ceiling test may still admit no satisfiable prescription. **Predicted:** cases exist with $c/\eta \leq \ln N$ whose required u^* exceeds $\bar{u}$ by more than an order of magnitude. **Falsified by:** a proof that $c/\eta \leq \ln N$ bounds $u^*/\bar{u}$. **Status:** exhibited by construction at a factor of 50.4; the *judgment* that a factor of 50.4 is unreachable is a **declared modelling position**, not a measurement (§5.5).

2.2 Definitions — true by stipulation, not falsifiable

- **Def. 1.** m := the roster size; N := the number of state **categories**, fixed at 12 throughout. These are different objects with different bounds and are never interchanged (§4.1).
- **Def. 2.** $u := 1 - (\hat{x} \cdot \hat{y})^2$ on unit intensity vectors; $\bar{u}$:= the mean of u over all $m(m-1)/2$ unordered pairs; $F := \ln(1 + (N-1)\bar{u})$.
- **Def. 3.** $N_{\text{eff}} := 1 + (N-1)\bar{u}$ (category scale) and $M_{\text{eff}} := 1 + (m-1)\bar{u}$ (roster scale).
- **Def. 4.** $D := 2.885390\,(c/\eta) - \log_2(f/k)$, the octave deficit.
- **Def. 5.** $\mu := [m(m-1)/2]/\lceil N/2 \rceil$, the muddiness index.
- **Def. 6.** $k := (9/8)^2/\varphi^2$. **This is a defined constant.** It is an exact ratio of exact quantities and is therefore computable to arbitrary precision; the fifteen significant figures 0.483425732988414 are *exact digits of a definition* and carry no measurement uncertainty because there is no measurement. Every appearance of the constant at that precision in this paper is so labelled.
- **Def. 7.** *Reachability* (Definition 5.1) is a **declared judgment** about what the parties could honestly represent, not a computed property of the formalism.

2.3 Postulates — asserted by this framework, with no external support

- **P-a. The essence law $f(n) = kn^2$.** A quadratic relation between a dimensionless count and a frequency. **No external literature supports it.** It is not derived here and nothing in the bibliography derives it. **The exponent 2 is defined by the model, not estimated from data.** No scaling exponent in this paper was obtained by fitting, none was tested against a log-normal or exponential alternative, and the goodness-of-fit and likelihood-ratio procedure a claimed scaling relation requires [Clauset et al., 2009] was **not run, because there is no data to run it on.**
- **P-b. The attachment of hertz to k.** Computing $(9/8)^2/\varphi^2 = 0.483425732988414$ is arithmetic. **Asserting that this dimensionless number *is a frequency* is a separate and much stronger claim**, and no operational procedure is supplied by which the constant acquires the dimension of inverse time. This is postulated, and it appears on its own line of the verification chain (§9.2, step L4).
- **P-c. The affine effective count $N_{\text{eff}} = 1 + (N-1)u$.** A design axiom [Curzi, 2026], not a derived result and **not a Hill number of any order.** The Hill family [Hill, 1973, Jost, 2006] and its independent rediscovery in political science [Laakso and Taagepera, 1979] ground $1/\text{HHI}$ and $\exp(H)$ exactly — as the corpus's concentration work uses them [Curzi, 2026] — and ground this construction not at all. **The roster-scale count M_{eff} inherits the same status**, and the participation-ratio identity of §4.2 is an algebraic re-expression of the affine axiom, not independent support for it.
- **P-d. The uniform-link cost model $C = \kappa m(m-1)/2$.** Every pair is assumed to maintain a link and every link is assumed to cost the same. Real cost topologies are neither complete nor uniform; SPEC 4.7 requires the non-uniform case to be stated as $C = \sum_{ij} \kappa_{ij}$, and every number in this paper uses the uniform model.

2.4 Modelling choices — legitimate, but choices

- **M-a. Octave reduction applied to a non-acoustic quantity.** Octave equivalence is a perceptual structure formalised for **pitch** [Shepard, 1982], whose perceptual specialness has a measured size [Demany et al., 2013]. Nothing licenses carrying it to a quantity that is not a pitch. The *operation* is well defined — division by a power of two — and its *interpretation* is a choice of this framework.
- **M-b.** $N = 12$**.** Twelve categories are used throughout and exceed the classical span of immediate absolute judgment [Miller, 1956], which is a reason for caution rather than a reason for twelve.
- **M-c. The mean as aggregator.** $\bar{u}$ is the arithmetic mean over pairs. Section 4.3 shows this is a choice with a price and requires both $F(\bar{u})$ and the mean of $F(u)$ to be published; it does not show that the mean is the right aggregator, and no argument here establishes that it is.
- **M-d. Interval apparatus.** Cents, octave equivalence and exact interval arithmetic are the correct apparatus once an index exists [Rossing et al., 2002, Barbour, 1951], but what a listener finds consonant depends on spectrum rather than ratio alone [Sethares, 2005] and interval categories generalise less far than is usually assumed [Burns, 1999]. This paper uses the arithmetic and claims none of the perception.
- **M-e. Governance principles read as parameter moves.** Mapping a design principle for durable common-pool institutions [Ostrom, 1990] onto η, c and δ is a **correspondence of dynamics, not of agency**, and is badged as one at §5.6.

2.5 What is expository

The specification blocks (§4.15–4.29 in the numbering of the specification lists, §5 and §6 likewise), the failure-mode catalogues, and the numbered Procedure T of §6.8 are expository: they set out procedure, reporting requirements and failure modes rather than argue a claim. Their organising question is *what must be published alongside a group reading for the reading to be checkable by a third party.*

3 Framework, dimensional audit, and method

3.1 The storage law — a shared mathematical form, not a derivation

$$Q_{t+1} = (1 - \delta)Q_t + \eta S_t - C_t$$

Analogy status: category 2, a shared mathematical form. The equation governing this framework's stored order is the equation that governs a leaky integrator, and it is used because the mathematics is the same, not because the physics transfers. It appears as the RC low-pass filter in signal processing [Oppenheim et al., 1996], as exponential smoothing in forecasting [Holt, 1957], and as the integrate-and-fire neuron in computational neuroscience [Lapicque, 1907, Gerstner and Kistler, 2002]. The fixed point $Q_\infty = (\eta F - c)/\delta$ follows from setting $Q_{t+1} = Q_t$ and requires $\delta > 0$; stability and uniqueness are standard [Elaydi, 2005] and the convergence is Banach's [Banach, 1922, Curzi, 2026]. **No claim is made that stored order is a charge or that any physical result about circuits applies to a group of people.**

Three warnings attach to that form and are recorded here rather than in a footnote. The time average of a multiplicative process differs from its ensemble average, so a parameter estimated across groups does not describe what happens to one group over time [Peters, 2019]. A single Q, δ or $\bar{u}$ written for a group is an aggregation, and aggregation over heterogeneous agents does not in general yield a representative agent [Kirman, 1992] — which is why this framework's no-aggregation rule [Curzi, 2026] is a rule and not a result. And Q has the grammar of a conserved quantity without a conservation law to give that grammar content, which is the characteristic failure of the genre [Mirowski, 1989]; this paper claims validity only

within a specified arrangement [Cartwright, 1999] and claims no universality. The wider programme — statistical physics applied to social dynamics — has a standard survey that fixes what such models can be held to [Castellano et al., 2009], and those are the boundaries claimed here.

3.2 The surplus functional, the effective count, and what the logarithm rests on

$F(u) = \ln(1 + (N-1)u)$ is the corpus's surplus functional, characterised by angle-only dependence, a zero at parallel states, an affine effective count and a normalisation [Curzi, 2026]. Its logarithmic form places it in the same family as Shannon's entropy [Shannon, 1948], and the ceiling $F \leq \ln N$ is the same ceiling that appears in requisite-variety arguments [Ashby, 1956, 1958].

Which warrant the logarithm actually needs is worth stating, because two are available and they are not equally strong. The psychophysical route reads F as a Fechner law [Fechner, 1860] and inherits every dispute in that literature. The multiplicative-dynamics route derives a logarithmic form with **no psychological assumption at all** [Peters and Gell-Mann, 2016]. The arguments in this paper need only the second: nothing here turns on the perceptual story being true.

The form is nevertheless underdetermined by the evidence. A power law with exponent below one [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the qualitative facts as well as $\ln(1+(N-1)u)$ does, and many mechanisms produce indistinguishable curves [Stumpf and Porter, 2012]. **Nothing in this paper discriminates among them**, and the results that matter most here — the crossover, the roster ceiling, the drop-one result — depend only on F being *concave and bounded*, which all three candidates are. That is stated as a robustness claim and it is the honest extent of it.

Finally, **which construct of difference u measures must be declared**, because at least three incompatible constructs travel under the word *diversity* — separation, variety and disparity [Harrison and Klein, 2007]. Here u is a *disparity* measure over represented states, and it is symmetric in its arguments, which psychological similarity is known not to be [Tversky, 1977]. That is a limitation and it is printed in §10.

3.3 The essence bridge and the constant k

The bridge from surplus to pitch is $f = k\,e^{2F}$, whose Result One is that $f(m)/f(n) = m^2/n^2$, so that **the constant cancels from every interval and does no work beyond setting a reference** [Curzi, 2026k]. That property is load-bearing for this paper: every deficit, crossover and ceiling below is a difference of logarithms in which k cancels exactly, so **the group verdicts never depend on the postulated constant and only the hertz values do.**

$k = (9/8)^2/\varphi^2 = 0.483425732988414$ is a **defined constant, exact**. Its two components are real and documented — the Pythagorean whole tone [Barbour, 1951] and the golden section [Euclid, 1956, Vajda, 1989] — but their product **expressed as a frequency in hertz has no precedent in any literature we have found.** The standard deflationary treatment of golden-section claims [Markowsky, 1992, Fischler, 1981] is cited here and wherever φ appears, and the popular literature's enthusiasm for the constant [Livio, 2002] is not evidence for anything asserted here.

3.4 Dimensional audit

Every symbol used in this paper, with its dimension, its unit, and whether it is measured, derived, defined or dimensionless by construction.

symbol	quantity	dimension	unit	status
Q	stored order	$[\mathcal{C}]$	none; **no operational unit exists**	state variable
δ	decay rate	$[1]$	per declared period	declared parameter
η	conversion efficiency	$[1]$	dimensionless fraction	declared parameter
κ	per-link cost	$[\mathcal{C}\,T^{-1}]$	per link per period	declared parameter
$C = \kappa m(m-1)/2$	total cost flow	$[\mathcal{C}\,T^{-1}]$	per period	derived
m	roster size	$[1]$	count	cardinality, absolute scale
N	category count	$[1]$	count	cardinality, absolute scale
$u,\ \bar{u}$	non-parallelism	$[1]$	dimensionless, $[0, 1]$	derived from unit vectors
F	surplus	$[1]$	nats	derived; **nats are not a physical unit**
PR	participation ratio	$[1]$	dimensionless	derived from the Gram spectrum
$r = \kappa/\eta$	cost ratio	$[\mathcal{C}\,T^{-1}]$	per link per period	derived
D	octave deficit	$[1]$	octaves (also cents)	derived
k	ground essence	**claimed** $[T^{-1}]$	Hz **by postulate**	**defined; the unit is asserted**
$f = k\,e^{2F}$	essence frequency	inherits k	Hz **by postulate**	derived given P-b
μ	muddiness index	$[1]$	pairs per class	derived

Does the audit close? Partly, and the gap is named. The storage law balances if and only if ηS and C carry the dimension of Q per period, which the framework stipulates rather than establishes; Q has **no operational unit at all**, and a dimensionless placeholder is the honest record of that. NET(m), D, μ and both effective counts are dimensionless and their audits close exactly. **The one equation that does not balance is** $f(n) = kn^2$, which requires k to carry $[T^{-1}]$ while k is computed as a ratio of two dimensionless quantities. That is postulate P-b, it is the framework's largest structural weakness, and no citation in the bibliography repairs it.

One consequence for this paper specifically. The roster ceiling $m(m-1) \leq 2\eta \ln N/\kappa$ is dimensionally clean: η is dimensionless, $\ln N$ is dimensionless, κ carries cost per link per period and the left side is a pure count, so the inequality is a statement about κ in units of η per link. **It therefore survives the failure of P-b intact**, and so does the crossover of §4.5. Only the hertz columns of the worked readings depend on the postulate.

3.5 Scale type, and what licenses the arithmetic

Permissible arithmetic is a function of the transformation group that leaves a quantity's empirical content unchanged [Stevens, 1946]. Three groups of quantities appear here and they do not have the same standing.

Counts. m, N, the link count and the class count are cardinalities. A cardinality admits only the identity transformation, so squaring it, halving it and forming $m(m-1)/2$ are licensed

operations. **This is the only part of the paper whose scale type is secure.**

Derived dimensionless ratios. u, $\bar{u}$, F, D, μ and PR inherit their standing from the intensity vectors they are computed from. Those come from magnitude estimation, which is *held* to deliver ratio-scale information [Stevens, 1957] and whose exponents are known to move with the experimenter's chosen stimulus range [Teghtsoonian, 1971] and to be exposed to catalogued response biases [Poulton, 1968]. **The dispute is live** [Krueger, 1989, Luce, 2002, Narens, 1996] and this paper inherits it whole.

Declared parameters. η, κ and δ are treated throughout as ratio-scaled quantities on a declared period. **Nothing establishes that they are.** Michell's objection — that psychology assumes quantitative structure without testing for it — is the single strongest objection to this framework [Michell, 1997, 1999], the formal test for additive quantitative structure exists [Luce and Tukey, 1964, Krantz et al., 1971], and **it has not been run on any quantity in this paper.** The corpus has demoted one of its own keystone constructions on exactly this ground [Curzi, 2026t], and the instrument that would repair the demotion is Rasch measurement [Rasch, 1960], which is likewise not applied here. **The ratio-scale treatment in this paper is therefore provisional, and it is conceded rather than defended.**

A consequence specific to the group case, and it is sharp. The sensitivity analysis of §5.10 shows that the constructed family's verdict is undetermined unless η is known to better than ± 0.030034 — 8.8 per cent of its value. **A quantity whose scale type is not established cannot be known to 8.8 per cent**, so that case is unresolved for a reason deeper than the one the sensitivity table gives. The table reports the arithmetic; this paragraph reports what the arithmetic presupposes.

3.6 Identification — the objection this framework cannot presently answer

The framework's central formal claim is a reflexive loop: the group influences the member and the members constitute the group. At group scope that loop is written explicitly, since $\bar{u}$ is a property of the members and the reading is a property of the group.

That is exactly a linear-in-means reflexive specification, and the endogenous social effect in such a specification is not identified from observational data [Manski, 1993]. It cannot be separated from contextual effects or from correlated unobservables however much data is collected, and influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011, Lyons, 2011]. **No amount of fit recovers the parameter.**

The remedy this paper adopts is stipulation of the reference group, following the template for discrete-choice models with social interactions [Brock and Durlauf, 2001]: the roster is declared, the period is declared, and the pairs over which $\bar{u}$ is averaged are enumerated and published. That is the cheapest of the four available remedies — the others being randomisation, an instrument, and partial-population variation — and it is not free. **Its price is that every claim in this paper about what a group's parameters imply is a modelling position, not an identified empirical result**, and it is labelled as one wherever it appears.

Two objections in the same family are also unanswered. The framework asserts that systems change their behaviour when modelled and observed, and then exempts its own parameters from that reflexivity, which is the Lucas critique turned on the framework itself [Lucas, 1976, Hacking, 1995]; a reflexive fixed point has been documented happening in the field [MacKenzie and Millo, 2003] and the mechanism is old [Merton, 1948], so the exposure is real rather than hypothetical. And §5.14's finding — that a member's declining to participate improves a constructed family's reading by 2.680588 octaves — is precisely the sort of number whose *publication* could change the behaviour it describes. That is why §6.2's consent axiom forbids showing it to anyone deciding whether to withdraw, and why the prohibition is an axiom rather than a courtesy.

3.7 Method

No data were collected. The method is exact arithmetic on stated parameters, in four steps.

1. **Derive.** The crossover, the roster ceiling, the roster-scale identity, the muddiness index and the aggregation gap are derived from the definitions of §2.2 by algebra, with each derivation shown at the point of use.
2. **Instantiate.** Four constructed cases carry the numbers: a triad extending the carried dyad, a four-member family, a five-member team, and the six-member commons. Every input vector and every raw L^2 norm is published so that any reading can be recomputed from scratch.
3. **Sweep.** Where a claim is about a lever, the lever is swept to its physical limit ($\eta \to 1$, $\kappa \to 0$, $u \to 1$, $N \to 144$) rather than to a convenient value, and the result at the limit is reported whether or not it is favourable.
4. **Audit.** Boundary behaviour, sensitivity to the least certain parameter, and the two aggregations are computed for every case, and the verdict is reported as NOT DETERMINED wherever they disagree.

Every figure in this paper is generated by a script named in its caption, taking its inputs from the tables in the text. Nothing is fitted; no model comparison is reported, because there is nothing to compare against and no out-of-sample data to compare on.

4 Group scale: interval counts, aggregation, and the roster ceiling

A dyad has one interval. That is the whole reason the dyadic papers of this series [Curzi, 2026,?] could treat a bond as a pitch and get away with it: there was one relationship to read, one `u` to measure, one number to compare against one threshold.

Add a third person and there are three intervals. Add a fourth and there are six. The relationships do not grow with the roster; they grow with the square of it, and everything that was simple at two becomes a question of aggregation at three, of accounting at six, and of arithmetic that has already decided the answer at twelve.

This section is about what changes and what does not. Three things change: the count of intervals, the need to aggregate them into one reading, and the cost structure, which Brooks named in 1975 and which the octave form of the minimum-frequency paper [Curzi, 2026] prices exactly. One thing does not change, and it is the result that organises the section: **the surplus ceiling is `ln N`, it is set by the category count and not by the roster, and a perfectly non-parallel pair already reaches it.** A larger group can therefore buy no more surplus than two well-chosen people, while paying quadratically more for the privilege.

That is a harsh statement and it needs its qualifications, which §4.7 supplies. But it is what the formalism says, it follows from the affine axiom in two lines, and a section on groups that did not say it first would be arranging the furniture around a hole in the floor.

Ownership: this section takes the interval count, the group chord, and the Brooks comparison. The mood-and-trait paper [Curzi, 2026] owns octave folding, the odd-part class rule and the `N = 17` regime statement, and this section cites all three without re-deriving any of them. The minimum-frequency paper [Curzi, 2026] owns `f_min`, the deficit and the levers. The dissonance paper [Curzi, 2026] owns the consonance/redundancy identity. Section 5 takes the three worked structures and the commons.

Nothing in this section is clinical, diagnostic or therapeutic, none of it is counselling, and the frequencies are essence frequencies — no group, person or room is claimed to resonate at any value in it.

The law

4.1 The interval count, and a notation warning that has to come first For a set of `m` members there are

```
m (m − 1) / 2      unordered pairs, hence that many intervals
```

m	pairs	pairs per member	pairs per available pitch class (6)
2	1	0.50	0.167
3	3	1.00	0.500
4	6	1.50	1.000
5	10	2.00	1.667
6	15	2.50	2.500
7	21	3.00	3.500
8	28	3.50	4.667
10	45	4.50	7.500
12	66	5.50	11.000
15	105	7.00	17.500
20	190	9.50	31.667
30	435	14.50	72.500

The notation warning: `N` in this framework is the number of state CATEGORIES, fixed at 12, and it is not the roster size. The manifest's shorthand `N(N−1)/2` refers to the roster; every other appearance of `N` in the formalism — `F <= ln N`, `N_eff = 1 + (N−1)u`, `f <= k N^2`, the ceiling `c/eta <= ln N` — is the category count. This section writes `m` for the roster and reserves `N` for categories, and it is worth being pedantic about because conflating them produces readings that are wrong by whole octaves in a way nothing downstream catches.

The second column is the one most often underestimated. Going from six members to twelve does not double the relational load; it multiplies it by 4.4, from 15 pairs to 66. The third column is the one this section turns into a result in §4.11.

4.2 Aggregation, and the two effective counts The construction of the framework's founding construction [Curzi, 2026] is pairwise. A group has many pairs and one reading, so an aggregator is required, and the framework uses one:

```
u_bar  =  mean of u over all m(m−1)/2 unordered pairs
F      =  ln( 1 + (N − 1) u_bar )          N = 12
f      =  k ( 1 + (N − 1) u_bar )^2
```

That is the construction the sustaining-frequency paper [Curzi, 2026] both use, and it is the one every reading in this framework uses. There is, however, a second effective count that lives alongside it, and keeping them apart is worth a short section because they answer different questions and have different bounds.

```
CATEGORY−SCALE      N_eff  =  1 + (N − 1) u_bar        bounded by N = 12
                    "how many distinct state−categories' worth of view is present"

ROSTER−SCALE        M_eff  =  1 + (m − 1) u_bar        bounded by m
                    "how many genuinely distinct members is this roster fielding"
```

THEOREM 4.1 (R3, verified at roster scale). `M_eff = m + 1 − m/PR`, where `PR` is the Renyi-2 participation ratio of the members' Gram matrix,

```
PR  =  ( trace G )^2 / sum_{ij} G_ij^2  =  m^2 / sum_{ij} G_ij^2
```

The participation ratio is the Renyi-2 quantity of the Gram spectrum [Rényi, 1961], and the identity is the roster-scale twin of the category-scale identity $N_{\text{eff}} = N+1-N/\text{PR}$ established in Series B, which was verified there at $K = 3, 5, 8$ and 12 to better than 10^{-12} [Curzi, 2026c] and at internal scope on a framing Gram matrix [Curzi, 2026o]. **Neither count is a Hill number.** The Hill family [Hill, 1973, Jost, 2006] and its political-science rediscovery [Laakso and Taagepera, 1979] ground $1/\text{HHI}$ and $\exp(H)$ exactly and ground the affine form not at all (P-c).

Verified on the constructed triad of the Worked Example: `PR = 1.3522930095`, `m + 1 − m/PR = 1.7815458788`, `1 + (m−1) u_bar = 1.7815458788`, error `0.00e+00`.

The identity is worth having for two reasons. It gives a second, independent route to `u_bar` that does not require enumerating pairs — useful at large `m`, where the pair count is the expensive part. And it supplies the sentence to hand to someone who does not want a frequency: a roster of three fielding `M_eff = 1.78` is three people doing 1.78 people's worth of distinct looking.

4.3 The aggregation gap, and it is larger than the dyadic case

`F` is strictly concave in `u` (the concavity of F [Curzi, 2026]). Therefore, by Jensen's inequality,

```
F( u_bar )  >=  mean of F( u_pair )        with equality only if all pairs share one u
```

The book's canonical reading uses `F(u_bar)`, which is the optimistic end of the interval, and the framework's reporting rule requires the conservative end to be published with it. the two-aggregation reporting rule [Curzi, 2026] states the rule; this section supplies the magnitude at group scope, and it is not small.

Computed on the constructed triad of the Worked Example:

```
pairwise u:  0.068734665265,  0.552647058824,  0.550937094081

F( u_bar )        = 1.6674242020 nats      f = 13.5718 Hz     D = −3.410906
mean of F(u)      = 1.4915742527 nats      f =  9.5476 Hz     D = −2.903511

AGGREGATION GAP   = 0.1758499493 nats  =  0.507396 octaves  =  609 cents
frequency ratio   = 1.421482
```

The same correction is already priced at protocol scope in Series B, where using F of the mean instead of the mean of F overstated feasibility by 0.0790 nats — 4.31 per cent, worth 0.3952 stored units at $\eta = 0.30$, $\delta = 0.06$ [Curzi, 2026]. At group scope the same error is an order of magnitude larger.

Half an octave of surplus, from the choice of where the averaging happens. Series B's worked office found the same effect worth 6.2%; at this triad's dispersion it is worth 11.79% of the conservative figure, and the reason is visible in the pairwise list: one pair at 0.0687 and two at 0.55 is exactly the dispersion Jensen's inequality punishes.

And the counter-case, computed on the commons of the carried commons [Curzi, 2026], whose fifteen pairwise `u` values all lie between 0.018484 and 0.018830:

```
F( u_bar )   = 0.186276647       mean of F = 0.186276219
AGGREGATION GAP = 0.00000124 octaves
```

Where the pairs are alike, the gap vanishes; where they are dispersed, it is the largest single uncertainty in a group reading. The rule that follows is not "use the conservative figure" — it is that both are computed and both are published, and where the verdict differs between them the verdict is not determined. It does not differ in the triad: `−3.410906` and `−2.903511` both accumulate, and per the sustaining-frequency paper [Curzi, 2026] an ACCUMULATES verdict is the robust direction anyway, since `f` is a lower bound.

4.3b When the gap is large, and when it can be ignored Because the aggregation gap is the largest controllable uncertainty in a group reading, it is worth knowing where it lives. Two sweeps, both computed on two-pair sets so that the dispersion is explicit.

Holding the mean at `u_bar = 0.40` and widening the spread between the two pairs:

```
pairs             F(u_bar)      mean-of-F      gap (octaves)
0.40 / 0.40       1.686399       1.686399        0.000000
0.35 / 0.45       1.686399       1.681185        0.015044
0.30 / 0.50       1.686399       1.665209        0.061142
0.25 / 0.55       1.686399       1.637392        0.141405
0.20 / 0.60       1.686399       1.595650        0.261847
0.15 / 0.65       1.686399       1.536289        0.433126
0.10 / 0.70       1.686399       1.452630        0.674514
0.05 / 0.75       1.686399       1.331439        1.024197
0.01 / 0.79       1.686399       1.187727        1.438862
```

Holding a fixed relative spread of ±50% of the mean and sweeping the level:

```
u_bar     pairs             gap (octaves)
 0.02     0.010 / 0.030       0.011776
 0.05     0.025 / 0.075       0.046143
 0.10     0.050 / 0.150       0.102518
 0.20     0.100 / 0.300       0.181418
 0.40     0.200 / 0.600       0.261847
 0.60     0.300 / 0.900       0.301406
```

The two sweeps together give a usable rule. The gap grows with absolute dispersion, quadratically at small spreads and faster after, and it grows mildly with the level. It is negligible in exactly one situation: when the pairwise values are close together, which is the near-parallel case. So the commons, whose fifteen pairs lie inside a band of width 0.000346, can be read from `F(u_bar)` without qualification — and a group with one intimate pair and two distant ones cannot.

The practical form: compute the spread of the pairwise list before deciding how much of the report the gap deserves. A group whose pairwise `u` values differ by less than about 0.05 has a gap under 0.02 octaves and may state one figure. Anything wider states two.

4.4 The group chord and the class ceiling A group's chord is the set of states its members hold, taken together. The self-as-chord material [Curzi, 2026] owns the chord at internal scope and the construction transfers unchanged: the states are indices, the indices are frequencies through `f = k n^2`, and the folding is the mood-and-trait paper's [Curzi, 2026].

The odd-part class rule, result R1 of the mood-and-trait paper [Curzi, 2026] gives the ceiling and this section cites it rather than proving it again: two indices share a pitch class if and only if they share an odd part, so an indexed set `1..N` holds exactly `ceil(N/2)` distinct classes — six, for `N = 12`.

The six, folded into 220–440 Hz and sorted by pitch, with adjacent separations:

```
class   folded (Hz)     cents above the previous
  11      233.978            —
   1      247.514          97.3641
   3      278.453         203.9100
   9      313.260         203.9100
   7      379.006         329.8318
   5      386.741          34.9756
                          (wrap to 2 x 233.978:  330.0085)
```

```
minimum adjacent separation  =  34.9756 cents  =  the 50/49 regime
```

That 34.9756-cent floor is the mood-and-trait paper's [Curzi, 2026], it holds because `N = 12` sits below the `N = 17` minimum-gap cliff, and this section states the fact once and does not re-derive it. The mood-and-trait paper [Curzi, 2026] §7.10 has the sweep and the exact-rational argument.

The consequence for groups is a saturation result, and it is the same shape as the requisite-variety bound: a regulator with too few distinguishable categories cannot regulate a variety larger than its own [Ashby, 1956, 1958]. The class count is capped at six no matter how many members there are, while the pair count grows quadratically. Computed by simulation — 20,000 trials, members' states drawn uniformly without replacement from the twelve, random seed 7:

```
expected distinct pitch classes covered
m      3 states each    4 states each    5 states each
 2         3.799            4.398            4.842
 3         4.506            5.008            5.372
 4         4.927            5.367            5.637
 6         5.432            5.729            5.879
 8         5.695            5.884            5.961
12         5.904            5.977            5.996
```

By six members holding four states apiece, a group covers 5.73 of the six available classes; by twelve it covers 5.98. Every further member adds pairs, adds cost, and adds essentially no class the chord did not already contain.

This is the precise sense in which a group's chord goes muddy, and it is not a metaphor about noise. **The chord saturates the lattice while the interval count keeps growing, so past a modest roster the additional relationships are all being drawn from a set of intervals the group already sounds.** the consonance-redundancy identity R2 [Curzi, 2026] is the other half of the statement: the most consonant pairs are exactly the pitch-class collisions, so a saturated chord is one whose new pairs are increasingly redundant with old ones.

4.5 Brooks, priced in octaves The observation that communication paths in a group of `m` grow as `m(m−1)/2` while useful output does not is Brooks's [Brooks, 1975], restated many times since — as the claim that a system's structure mirrors the communication structure of the organisation that built it [Conway, 1968], and in the wider stock-and-flow literature on how coordination load accumulates [Forrester, 1961, Sterman, 2000, Hopp and Spearman, 2011]. The octave form of the minimum-frequency paper [Curzi, 2026] makes the comparison exact, because it puts both sides in the same unit.

Cost side. With a per-link cost `kappa` per period, `C = kappa m(m−1)/2`, and the minimum-frequency paper [Curzi, 2026] Theorem 4.3 converts:

```
octaves required  =  2.885390 ∗ ( C / eta )  =  1.442695 ∗ ( kappa / eta ) ∗ m ( m − 1 )
```

Surplus side. The most a roster of `m` can field, at roster scale, is `M_eff = m` — every member mutually non-parallel. Through the essence bridge that is

```
octaves available (best case)  =  2 log2( m )
```

THEOREM 4.2 — TWO OCTAVES PER DOUBLING AGAINST A QUADRATIC. In the best case a group's surplus rises by exactly two octaves each time the roster doubles, while its cost rises as `m(m−1)`. Setting the two equal and writing `r = kappa/eta`:

```
NET(m)  =  2 log2(m)  −  1.442695 ∗ r ∗ m(m−1)      octaves

crossover m∗ solves      2 ln m  =  r m ( m − 1 )
```

Proof: substitute and multiply through by `ln 2`. **QED.**

The crossover has a closed form as a bound on `r`: the largest viable integer roster at cost ratio `r` is the largest `m` with `2 ln m >= r m(m−1)`, equivalently `r <= 2 ln m / (m(m−1))`. Computed:

m	r_max(m) = 2 ln m / (m(m−1))	m	r_max(m)
2	0.69314718	10	0.05116856
3	0.36620410	11	0.04359810
4	0.23104906	12	0.03765010
5	0.16094379	15	0.02579095
6	0.11945063	17	0.02083245
7	0.09266239	20	0.01576701
8	0.07426577	25	0.01072959
9	0.06103402		

Read the other way — the maximum viable roster at a given cost ratio:

r = kappa/eta	max viable roster (best case)	
0.005	38	
0.010	26	
0.020	17	
0.030	13	
0.050	10	
0.100	6	
0.125	5	
0.161765	4	
0.200	4	
0.366204	3	
0.500	2	
0.693147	2	<-- r_max(2) = ln 2 exactly
1.000	1	
2.066667	1	

The framework's own contact analysis reaches the same shape by a different route: contact raises surplus only while a boundary-cost multiplier stays below a critical value, computed there as 3.17 against an operating value of 8 [Curzi, 2026g].

A dyad is viable in the best case if and only if `kappa/eta <= ln 2 = 0.693147`, which is a pleasing way to see why the carried dyad of the sustainable-dyad paper [Curzi, 2026], at `r = 2.066667`, could not be saved by its geometry alone. It was over the roster bound at two.

Two figures from the carried cases, both computed:

```
the carried commons:   r = 0.055/0.34 = 0.161765
                           r_max(5) = 0.160944
                           the six-member commons misses a viable FIVE-member roster by
                           0.000821 in r -- a per-link cost of 0.055 against 0.054721

the constructed triad:     r = 0.055/0.34 = 0.161765, max viable roster 4, actual m = 3
```

The commons fails the best-case roster test by half a percent of its per-link cost, before any question of who its members are.

Section 5 takes that case up in full.

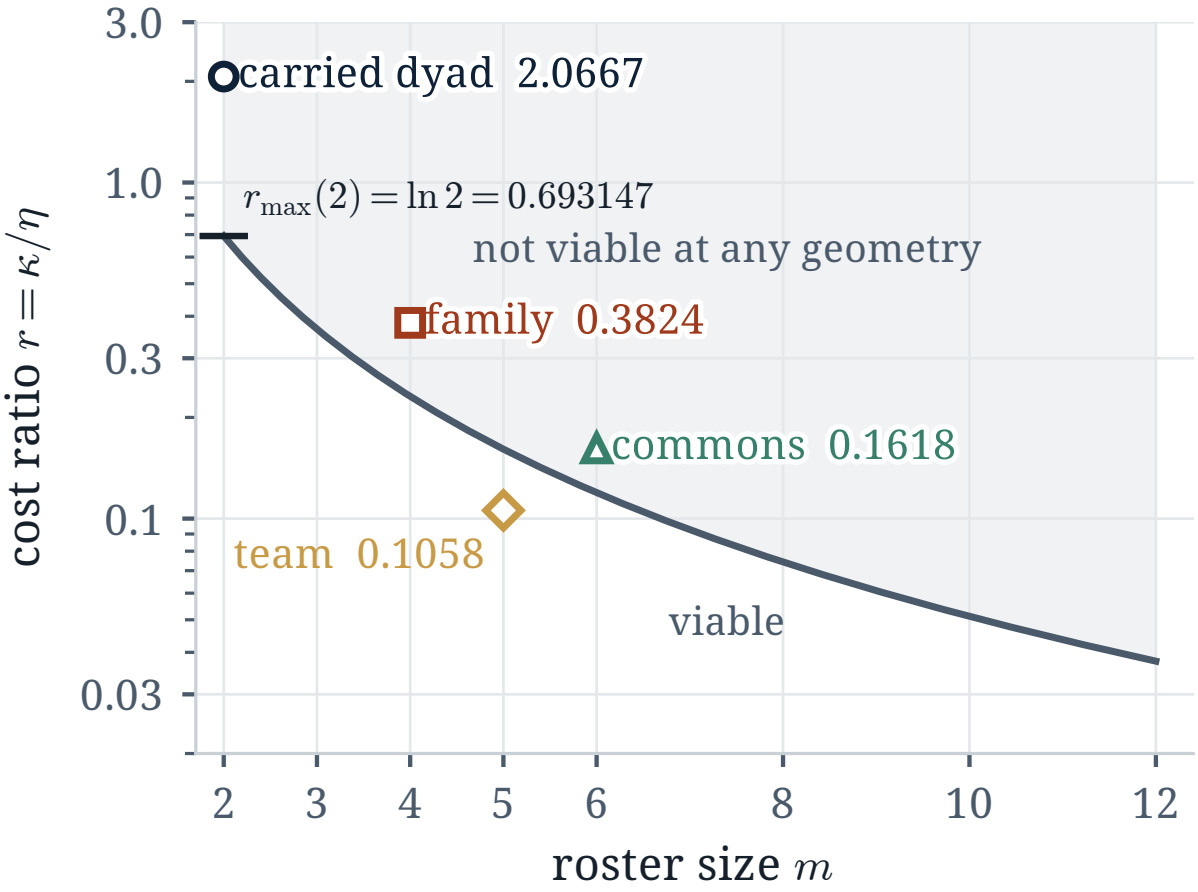


Figure 45.1: **The viability frontier is a statement about the cost ratio, not about headcount.** $r_{\max}(m) = 2\ln m/(m(m-1))$ is the largest κ/η at which a roster of m can pay for itself in the best case, and it equals $\ln 2 = 0.693147$ exactly at $m = 2$. Markers are the four constructed cases of this paper: the carried dyad and the family sit above the frontier, the team below it, the commons above it at six members and below it at four. Modelled from the closed form of Theorem 4.2 with the case parameters of §4.6 and §5.9; script `figs/src/CCP-045_fig01.py`.

4.6 The roster ceiling, which is exact The Brooks crossover of §4.5 is a best-case comparison and it rests on a modelling choice (`M_eff = m`). There is a second bound that rests on nothing but the ceiling result R5 of the minimum-frequency paper [Curzi, 2026], and it is exact.

THEOREM 4.3 — THE ROSTER CEILING. Since `F <= ln N` and sustainability requires `c/eta <= ln N`, a group with per-link cost `kappa` and conversion `eta` is sustainable at *any* geometry only if

```
kappa m ( m − 1 ) / 2  <=  eta ln N        i.e.        m ( m − 1 )  <=  2 eta ln N / kappa
```

Proof: substitute `C = kappa m(m−1)/2` into `c/eta <= ln N`. **QED.**

Computed for the cases this framework carries, at `N = 12`, `ln N = 2.484906649788`:

case	eta	kappa	m(m−1) bound	m_max
commons	0.34	0.055	30.7225	6
family (Sec 5)	0.34	0.130	12.9980	4
team (Sec 5)	0.52	0.055	46.9873	7
carried dyad	0.30	0.620	2.4047	2
dyad at 2b	0.44	0.400	5.4668	2

The commons sits at exactly its roster ceiling — `m(m−1) = 30` against a bound of 30.7225 — which is why the classification of the minimum-frequency paper [Curzi, 2026] returns TUNABLE for it with a margin of only 0.058436 nats. A seventh member would make it structural outright, and no assignment of `u` whatsoever would save it.

The two bounds answer different questions and both are needed. Theorem 4.3 says whether *any* geometry could work. Theorem 4.2 says whether the best achievable geometry would. A group inside 16.3 and outside 16.2 is one whose survival depends on being extraordinarily well composed.

4.6b The hold line by roster size The hold-line corollary R6 [Curzi, 2026] gives the geometry required for a relation to hold exactly: `u* = (e^(c/eta) − 1)/(N − 1)`. At group scope, with `C = kappa m(m−1)/2`, that becomes a function of the roster and it is the most direct way to see what quadratic cost does.

Computed at `kappa = 0.055`, `eta = 0.34`, `N = 12`:

m	links	C	c/eta	u* required to HOLD	
2	1	0.0550	0.161765	0.015962	
3	3	0.1650	0.485294	0.056787	
4	6	0.3300	0.970588	0.149045	
5	10	0.5500	1.617647	0.367383	
6	15	0.8250	2.426471	0.938078	
7	21	1.1550	3.397059	2.625100	UNREACHABLE, u <= 1
8	28	1.5400	4.529412	8.336728	UNREACHABLE
9	36	1.9800	5.823529	30.651224	UNREACHABLE
10	45	2.4750	7.279412	131.739523	UNREACHABLE
12	66	3.6300	10.676471	3938.488241	UNREACHABLE

The `u*` column rises far faster than the roster does, because it is exponential in a quantity that is quadratic in `m`. Two members need 1.6% non-parallelism to break even at this link price. Six need 93.8%, which is the commons' figure of the carried commons [Curzi, 2026] arriving from the cost side rather than from the geometry side. Seven need more than exists.

The seventh member does not make the problem harder; it makes it insoluble, and the transition between those two states occupies no interval at all — one member, and the required geometry goes from demanding to unavailable. That is the roster ceiling of Theorem

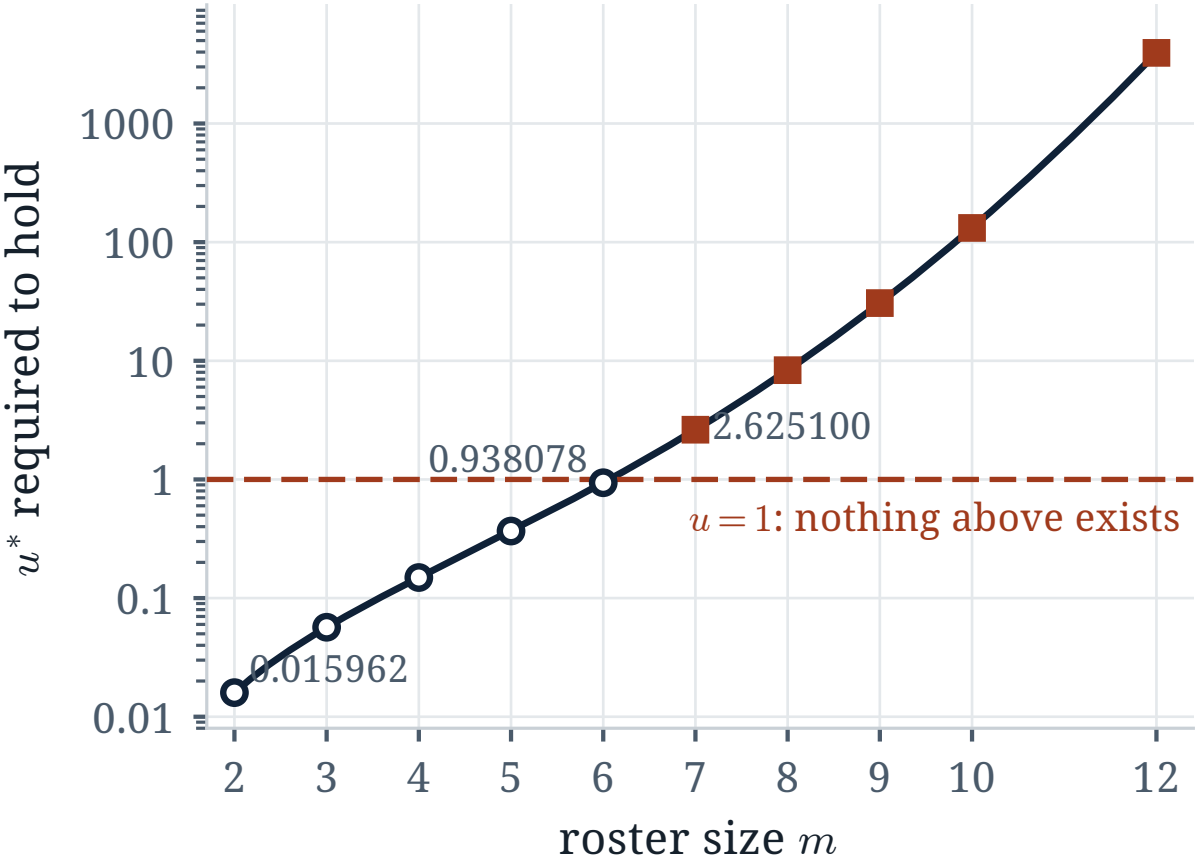


Figure 45.2: **The seventh member is not a harder problem, it is an unavailable one.** The geometry a roster must field merely to hold, $u^* = (e^{c/\eta} - 1)/(N - 1)$ with $C = \kappa m(m-1)/2$, is exponential in a quantity quadratic in m: two members need 1.6 per cent non-parallelism, six need 93.8 per cent, and seven need 2.6251 — more than exists. Modelled at $\kappa = 0.055$, $\eta = 0.34$, $N = 12$; values from the table of §4.6b; script `figs/src/CCP-045_fig02.py`.

4.3 written as a requirement on `u`, and it is the same statement the hold-line corollary R6 [Curzi, 2026] makes about a dyad.

A second sweep, holding the geometry fixed at the constructed triad's `u_bar = 0.390773` and varying only the roster:

```
 m    C        c/eta       D            Q_inf     verdict
 2   0.0550   0.161765   −4.344415      +5.6880   ACCUMULATES
 3   0.1650   0.485294   −3.410906      +4.4658   ACCUMULATES
 4   0.3300   0.970588   −2.010644      +2.6325   ACCUMULATES
 5   0.5500   1.617647   −0.143626      +0.1880   ACCUMULATES, barely
 6   0.8250   2.426471   +2.190145      −2.8675   DECAYS
 7   1.1550   3.397059   +4.990671      −6.5342   DECAYS, STRUCTURAL
 8   1.5400   4.529412   +8.257951     −10.8120   DECAYS, STRUCTURAL
12   3.6300  10.676471  +25.994613     −34.0342   DECAYS, STRUCTURAL
```

A group of people every one of whom is as well matched to the others as the constructed triad's members are — `u_bar` held at 0.390773 for every pair, which is a strong assumption at any size — crosses from accumulating to decaying between five and six members, and from tunable to structural between six and seven. **The geometry is doing nothing wrong in any row of that table. The roster is.**

4.7 The uncomfortable result: a pair already reaches the ceiling Now the statement flagged in the opening, which follows from the construction in two lines.

`F = ln(1 + (N−1) u_bar)` is maximised at `u_bar = 1`, giving `F = ln N`. Two unit vectors can be orthogonal, `cos theta = 0`, hence `u = 1`. **Therefore a single orthogonal pair already attains the maximum surplus the twelve-category system admits, and no roster of any size attains more.**

```
orthogonal pair:   u_bar = 1,  F = ln 12 = 2.484907,  f = k ∗ 144 = 69.6133 Hz
orthogonal six:    u_bar = 1,  F = ln 12 = 2.484907,  f = k ∗ 144 = 69.6133 Hz
                   −− identical, and the six pays 15 links instead of 1
```

In surplus terms a group buys nothing a well-chosen pair could not, and pays `m(m−1)/2` times as much for it. That is the honest reading of the formalism and it should be stated before any of the qualifications, because the qualifications are where the wishful thinking gets in.

This is the point at which the external literature must be brought in rather than kept out, and it does not speak with one voice. The formal result that a diverse group of problem solvers can outperform a group of the best individual solvers [Hong and Page, 2004, 2001, Page, 2007] has been challenged on the ground that its conclusion follows from the model's parameterisation rather than from any empirical fact about diversity [Thompson, 2014, Singer, 2019, Grim et al., 2019, Kuehn, 2017]. The empirical literature is likewise mixed: a general collective-intelligence factor has been measured, and it is reported to track composition rather than headcount [Woolley et al., 2010, Engel et al., 2014, Bear and Woolley, 2011], while the organisational reviews find diversity effects that are small, conditional and frequently negative [Mannix and Neale, 2005, van Knippenberg and Schippers, 2007]. Aggregation of independent judgments is powerful [Galton, 1907, Surowiecki, 2004] and is destroyed by exactly the mechanism this framework calls low u — social influence collapsing the diversity of estimates [Lorenz et al., 2011], and cascades in which later members rationally ignore their own information [Bikhchandani et al., 1992, Banerjee, 1992]. **The result stated above is a statement about this framework's surplus functional, not a summary of that literature**, and where the two are compared the comparison is stated as such.

Now the qualifications, and they are real.

One: orthogonality is an idealisation, and mean pairwise `u_bar = 1` for six people requires all fifteen pairs to be orthogonal at once. Six mutually orthogonal unit vectors exist in twelve

dimensions, so it is geometrically possible — but the vectors are then extremely constrained, and the disjoint-support degeneracy, result R5 [Curzi, 2026] applies with force: six people described over non-overlapping state sets read `u = 1` for free and mean nothing by it. Achievable `u_bar` falls with `m` in every realistic composition, and the pair is the easiest place to find a large one.

Two: `F` is not the only thing a group is for. The formalism prices stored order per period. It has no term for redundancy against loss, for succession, for coverage across time zones or lifespans, for the fact that a pair ends when one of its members does. A group's non-surplus functions are entirely outside this reading and the reading does not weigh them — the boundary statement of the minimum-frequency paper [Curzi, 2026] states the boundary and it holds here without modification.

Three: `eta` and `kappa` are not roster-independent in practice. A group may convert better than a pair because it has structure a pair does not need, and the commons-governance material of §5.6 is largely about arrangements that lower `kappa` per link. The formalism holds them fixed because it is a formalism; a real design moves them.

Four: the ceiling is a ceiling on `F`, not on output. The framework's dyadic work makes the neighbouring point that a calibrated system at its own fixed point has headroom identically zero, so being at the ceiling is not an unqualified target [Curzi, 2026r]. Two people at `F = ln 12` and six people at `F = ln 12` bank at different rates if their `eta` differs, and `Q_inf = (eta F − c)/delta` is sensitive to all three.

With all four in hand the statement survives in a weaker and still useful form: the case for a larger group is never a case about surplus, because surplus is capped at a value a pair can reach. It is always a case about conversion, coverage, or continuity, and it should be made in those terms.

4.8 The opposition identity at group scope The relationship-frequency paper [Curzi, 2026] owns the unison result and it does not weaken with roster size. `u` depends only on `|rho|`, so a member recruited to a group specifically to disagree contributes `u = 0` to every pair in which the disagreement is systematic, exactly as an echo would, while adding `m − 1` new links to the cost.

```
adding a perfectly opposed member to a group of m:
    surplus contribution                    0
    cost contribution        kappa * (m − 1) per period
    octave contribution      +1.442695 * (kappa/eta) * 2(m−1)   −− strictly positive DEFICIT
```

A contrarian appointment is the only intervention in the framework that is guaranteed to make the reading worse by a computable amount. At `kappa/eta = 0.161765` and `m = 6`, adding one such member costs `+2.3338` octaves of deficit and returns nothing.

4.9 Honesty clauses at group scope **(1)** Essence frequencies, not physical resonances. No group, room, institution or person is claimed to resonate at any value here, and the minimum-frequency paper's [Curzi, 2026] FM 4.8 forbids playing `f_min` at anyone.

(2) No clinical, diagnostic or therapeutic claim; the twelve states are not diagnoses; this is not counselling and no reading here is an assessment of any person or of any group's health.

(3) All cases are constructed and labelled. The triad of the Worked Example is an explicitly labelled extension of the carried dyad of the carried dyad [Curzi, 2026].

(4) The reading inherits its inputs, and at group scope it inherits `m(m−1)/2` of them. Uncertainty in a group reading is not `m` times a dyadic uncertainty; it is whatever the aggregation gap and the least certain parameter jointly permit, and §4.3 gives the first term.

(5) Consent is structural, and at group scope it is per-member. Section 6 owns the procedure. **A group reading taken without every member's agreement is not a group reading; it is a reading about people, taken without them.**

(6) CC-R14 applies to the simulation in §4.4: it is a computation with a stated seed and trial count, not an observation of anybody.

The engineering

4.10 Reading a group, end to end

```
 0.  CONSENT from every member, separately, and revocable.                  Sec 6
 1.  DECLARE N and the ordering.                                            CCP-041
 2.  ELICIT intensities by magnitude estimation; publish every raw L2 norm. CCP-041
 3.  ENCODE each member as a unit vector; check support overlap.            CCP-042
 4.  COMPUTE all m(m-1)/2 pairwise u; PUBLISH THE LIST, not just the mean.  Sec 4
 5.  AGGREGATE: u_bar, F(u_bar), and mean-of-F. Report the gap.             CCP-041 SPEC 4.15
 6.  DECLARE kappa, eta, delta on one period; C = kappa m(m-1)/2.           CCP-041 SPEC 4.3
 7.  ROSTER CEILING TEST: m(m-1) <= 2 eta ln N / kappa.                     Thm 4.3
 8.  CEILING TEST: c/eta vs ln N; classify.                                 CCP-041 SPEC 4.1
 9.  COMPUTE f, f_min, D, Q_inf; fold and publish with j.                   CCP-041 SPEC 4.5
10.  SIGN THE NUMERATOR eta F - c.                                          CCP-044 SPEC 15.2
11.  DROP-ONE ANALYSIS.                                                     Sec 4.13
12.  REPORT sensitivity to the least certain parameter.                     CCP-041 SPEC 4.19
```

Step 4's instruction to publish the whole list is the one that costs nothing and catches the most. A mean of 0.39 is produced by three pairs at 0.39 and by one pair at 0.07 with two at 0.55, and those are entirely different groups. The list shows which, the mean does not, and §4.3's aggregation gap is computable only from the list.

4.11 Where the chord goes muddy, precisely Three distinct things happen as `m` grows, and they are usually described as one thing.

The class ceiling saturates. §4.4: six classes, reached in expectation by about six members holding four states each. After that, new members add no new classes.

The pair-per-class ratio grows without bound. From the table in §4.1: 0.5 at `m = 3`, 2.5 at `m = 6`, 11.0 at `m = 12`, 72.5 at `m = 30`. This ratio is the honest measure of muddiness — it is how many relationships are being carried by each available distinction.

The rendering floor does not move. The minimum class separation stays at 34.9756 cents for every `N <= 16`, because it is a property of the lattice and not of the roster (the mood-and-trait paper [Curzi, 2026]). A group of thirty sounds no finer than a group of three; it sounds the same six classes, more often.

COROLLARY 4.4 — THE MUDDINESS INDEX. For a roster of `m` on an `N`-category ladder,

```
mu  =  [ m(m-1)/2 ]  /  ceil(N/2)        relationships per available distinction
```

`mu > 1` first occurs at `m = 4` for `N = 12`. **From four members onward, a group has more relationships than it has distinguishable classes to express them in**, and every additional member worsens the ratio quadratically while the denominator is frozen.

This is not an argument that groups of four are bad. It is an argument about what a chord reading can and cannot tell you at four and above: the chord stops individuating relationships and starts reporting only the group's coverage of the lattice. Past `mu > 1` the useful group-level objects are the pairwise list, the drop-one analysis, and the deficit — not the chord.

4.12 The marginal member test The decision that actually arises is not "how big should the group be" but "should this person join." That is a marginal question and it has a clean form.

```
marginal surplus (best case)  =  2 log2( m / (m−1) )                octaves
marginal cost                 =  1.442695 ∗ (kappa/eta) ∗ 2(m−1)     octaves

ADD if surplus > cost
```

Computed at `kappa/eta = 0.161765`:

m	marginal surplus	marginal cost	verdict
2	+2.0000	0.4668	ADD
3	+1.1699	0.9335	ADD
4	+0.8301	1.4003	STOP
5	+0.6439	1.8670	STOP
6	+0.5261	2.3338	STOP
7	+0.4448	2.8005	STOP
8	+0.3853	3.2673	STOP

The marginal surplus falls like `1/m` and the marginal cost rises like `m`, so once the test flips it never flips back. There is no roster size at which adding members becomes worthwhile again, at fixed `kappa` and `eta`. The only way to reopen the question is to change the cost structure, which is §4.14 and the commons-governance material of §5.6.

The structure of the test — a marginal gain priced against a marginal cost, with the verdict flipping once incremental cost passes a computable threshold — recurs across the corpus; the structural work prices an optional addition the same way and finds it worth nothing at any probability once its incremental cost exceeds 8.9 per cent of the maintenance budget [Curzi, 2026].

Two cautions on using the test. It is a best-case surplus, so a real candidate must be compared against their actual contribution to `u_bar`, not against `2 log2(m/(m−1))`. And it is silent on everything a person is for besides surplus — the boundary statement of the minimum-frequency paper [Curzi, 2026]'s boundary applies, and a marginal test that returns STOP is not a statement that anyone should be excluded from anything.

4.13 The drop-one analysis The complement of the marginal test, run on an existing group: recompute the reading with each member's pairs removed, and look at what moves.

Computed on the constructed triad, at `kappa = 0.055`, `eta = 0.34`, `delta = 0.09`:

set	u_bar	links	C	f (Hz)	D	Q_inf	change in D
{A,B,C}	0.390773	3	0.1650	13.5718	−3.410906	+4.4658	—
{B,C}	0.550937	1	0.0550	24.0978	−5.172708	+6.7725	−1.761802
{A,C}	0.552647	1	0.0550	24.2264	−5.180385	+6.7826	−1.769478
{A,B}	0.068735	1	0.0550	1.4908	−1.157965	+1.5161	+2.252941

Removing either `A` or `B` improves the reading by about 1.77 octaves; removing `C` worsens it by 2.25. The reason is entirely geometric: `A` and `B` sit at `u = 0.0687`, near-parallel, and their pair is the one dragging `u_bar` down while still costing a link. `C` is the member carrying the group's non-parallelism.

A drop-one audit is also an m-fold multiple comparison performed after the data are in hand, which is the standard recipe for a false positive [Simmons et al., 2011, Gelman and Loken, 2013]. **What the drop-one analysis is for, and what it is not for, must be stated in the same breath.** It is a sensitivity analysis: it tells you which pairs are carrying the reading, which is information about the geometry of the set. It is not a ranking of people, it licenses no

conclusion about anyone's value or contribution, and a formalism with no single-party term cannot acquire one by being run `m` times. The correct use of the row that says "removing `A` improves the reading" is to notice that `A` and `B` are near-parallel and to ask what would raise `u` between them — which is the same conclusion the sustainable-dyad paper [Curzi, 2026] reached about the same pair.

4.13b The two-dyad group, and why triads are not four dyads One structure deserves separate treatment because it is common and because its reading is counterintuitive: the group that is really two pairs.

Take four members `{P, Q, R, S}` where `P–Q` and `R–S` are close pairs and the cross-links are distant. The six pairwise values might run `u(P,Q) = 0.08`, `u(R,S) = 0.08`, and four cross-pairs at 0.70. Computed:

```
u_bar           = (0.08 + 0.08 + 0.70*4)/6            = 0.493333
F( u_bar )      = ln(1 + 11*0.493333)                 = 1.860456 nats
mean of F(u)    = (2*ln(1.88) + 4*ln(8.70))/6         = 1.652639 nats
                  ln(1.88) = 0.631272   ln(8.70) = 2.163323
AGGREGATION GAP = 0.207817 nats = 0.599632 octaves
f optimistic    = 19.9665 Hz      f conservative = 13.1763 Hz
```

The gap is large, and it is large for a structural reason rather than a numerical one: **a group made of tight clusters is exactly the dispersion pattern Jensen's inequality punishes most, and it is also the most common real topology.** Families have it, teams have it, and any group formed by merging two existing groups has it by construction.

Two consequences follow. First, cluster-structured groups must always be read on both aggregations, and the constructed family of Section 5 is worked that way. Second, the within-cluster pairs are the ones the drop-one analysis will flag, for the same reason `A` and `B` are flagged in this section's triad: they cost a link and contribute little non-parallelism. The correct reading of that flag is never about the people; it is that the group is paying twice for one perspective, which is the relationship-frequency paper [Curzi, 2026]'s result about the sycophant's unison applied to a sub-group instead of to a pair.

And a triad is not three dyads plus overhead. It has a property no collection of pairs has: for any two members there is a third who is in relation to both, which is what makes mediation, witness and succession structurally possible. **None of that appears in `F`, and a section that computed only `F` would be reporting the smaller half of what a triad is.** The formalism prices stored order per period; the third member's structural role is real, is outside the price, and is not made unreal by being outside it.

4.14 Splitting, nesting, and what changes the exponent The one intervention that changes the *shape* of the cost function rather than its coefficient is topological: stop maintaining every link.

Ostrom's eighth design principle — nested enterprises — has a precise reading here. A flat group of `m` maintains `m(m−1)/2` links. A group partitioned into `g` sub-groups of size `m/g`, each maintaining its own links, plus one link per pair of sub-groups, maintains

```
g * (m/g)(m/g − 1)/2  +  g(g−1)/2        links
```

Computed for six members:

```
flat six                          15 links      C = 0.8250 at kappa 0.055
two triads + one bridge            7 links      C = 0.3850          −53.3%
three pairs + three bridges        6 links      C = 0.3300          −60.0%
```

Nesting is the only lever in this section that attacks the quadratic itself, and on six members it removes more than half the cost. The same move is measured in Series B on eighteen people rewired from a complete graph into three cells of six: headroom moves from -0.4571 to $+0.1189$, a swing of 0.576, with cost falling 63 per cent and the surplus ceiling untouched because $N = 18$ either way [Curzi, 2026b]. The bridges between cells are the structural-hole and weak-tie positions of the network literature [Granovetter, 1973, Burt, 1992, 2004], and the partition itself is the standard resilience move against correlated failure [Albert et al., 2000, Callaway et al., 2000, Newman, 2018]. Series B's CC-R15 material records the same result in the other register: P8 turns `O(N^2)` coordination into `O(N)`.

And it is not free, which Section 5 computes rather than asserts. Two members in different sub-groups now interact through a bridge, so whatever surplus their pair carried is attenuated or lost. Nesting is a trade of surplus for cost and the trade is only worth making when the lost pairs were carrying little — which, at a low `u_bar`, they are.

The specification

SPEC 4.1 — PER-MEMBER CONSENT, REVOCABLE. Every member agrees separately to a group reading and any member may withdraw at any point. A member's withdrawal ends their inclusion in the reading and may not be treated as a finding about them.

SPEC 4.2 — DECLARE `m` AND `N` SEPARATELY in every group reading, and never let `N` denote the roster.

SPEC 4.3 — PUBLISH THE FULL PAIRWISE LIST, not the mean alone. The list is what makes the aggregation gap and the drop-one analysis computable.

SPEC 4.4 — REPORT BOTH AGGREGATIONS. `F(u_bar)` and the mean of `F(u)`, with the gap in octaves. Where the verdict differs between them, report the verdict as NOT DETERMINED.

SPEC 4.5 — SUPPORT-OVERLAP GATE. Check every pair for disjoint support before computing. The encoding degeneracy R5 [Curzi, 2026] makes non-overlapping descriptions read `u = 1` for free.

SPEC 4.6 — ROSTER CEILING TEST (Theorem 4.3) before any lever prescription: `m(m−1) <= 2 eta ln N / kappa`.

SPEC 4.7 — DECLARE `kappa` PER LINK and state whether links are assumed uniform. Non-uniform link costs require `C = sum kappa_ij` and the roster ceiling is then stated on the sum.

SPEC 4.8 — REPORT THE MUDDINESS INDEX `mu = [m(m−1)/2]/ceil(N/2)` with any group chord, and do not present a chord as individuating relationships where `mu > 1`.

SPEC 4.9 — CITE THE REGIME, DO NOT RE-DERIVE IT. `N = 12` is below the `N = 17` cliff, the minimum class separation is 34.9756 cents, and the mood-and-trait paper [Curzi, 2026] owns both statements.

SPEC 4.10 — MARGINAL TESTS USE ACTUAL CONTRIBUTIONS. A candidate's effect on `u_bar` is computed from their vector, not from the best-case `2 log2(m/(m−1))`.

SPEC 4.11 — DROP-ONE IS A SENSITIVITY ANALYSIS. It is reported as a property of the geometry. No output of a drop-one analysis may be presented as a ranking, an evaluation, or a recommendation about a person.

SPEC 4.12 — CC-R8 CHECK ON RECRUITMENT. Report `cos theta` signed for any proposed member. A systematically opposed member contributes `u = 0` and `m − 1` links.

SPEC 4.13 — NESTING IS REPORTED WITH ITS LOST PAIRS. A topology change that removes links also removes whatever surplus those pairs carried; both terms are computed.

SPEC 4.14 — NO AGGREGATION ACROSS GROUPS (the no-aggregation rule [Curzi, 2026], and CC-R16: two nine-member groups with identical `u` and `S` lost 77% and 21% of surplus to the same three-person departure).

SPEC 4.15 — LABEL EVERY CASE CONSTRUCTED.

The failure modes

FM 4.1 — `N` USED FOR THE ROSTER. *Signature:* `F = ln(1 + (m−1)u)` computed with `m` in place of 12, or a ceiling of `ln m`. *Mechanism:* the two counts are different objects with different bounds (§4.2). *Computed:* on the constructed triad, `1 + 11 u_bar = 5.2985` and `1 + 2 u_bar = 1.7815` — a difference of 3.152 octaves in the reading. *Fix:* SPEC 4.2.

FM 4.2 — THE MEAN REPORTED WITHOUT THE LIST. *Signature:* one `u_bar` for a group, no pairwise values. *Mechanism:* the mean is compatible with wildly different geometries, and the aggregation gap is not computable from it. *Fix:* SPEC 4.3.

FM 4.3 — THE JENSEN SHORTCUT TAKEN SILENTLY. *Signature:* a group reading from `F(u_bar)` alone. *Mechanism:* strict concavity inflates `F`, inflates `f`, and under-reports `D`. Worth 0.507396 octaves on the constructed triad; 6.2% in Series B's worked office. *Fix:* SPEC 4.4.

FM 4.4 — THE ORTHOGONAL GROUP THAT IS AN ARTEFACT. *Signature:* a large group reading `u_bar` near 1. *Mechanism:* disjoint support. Members described over non-overlapping state sets are automatically orthogonal (the one-hot degeneracy R5 [Curzi, 2026], generalised). *Fix:* SPEC 4.5, and re-encode every member under one procedure.

FM 4.5 — GROWING A GROUP FOR SURPLUS. *Signature:* a roster expanded on the argument that more people means more perspective. *Mechanism:* §4.7. Surplus is capped at `ln N` and an orthogonal pair reaches it. Cost is not capped. *Fix:* make the case for the roster on conversion, coverage or continuity, which are real and are outside this reading.

FM 4.6 — THE ROSTER CEILING NOT TESTED. *Signature:* a lever prescription for a group that no geometry could sustain. *Mechanism:* Theorem 4.3. At `m(m−1) > 2 eta ln N / kappa` the case is structural and `u` is unavailable at every value. *Computed:* the commons sits at `m(m−1) = 30` against a bound of 30.7225; a seventh member takes it to 42 and makes it structural outright. *Fix:* SPEC 4.6, before anything else.

FM 4.7 — THE CONTRARIAN APPOINTMENT. *Signature:* a member recruited to disagree. *Mechanism:* CC-R8. `u` depends only on `|rho|`; opposition contributes zero surplus and `m−1` links. *Computed:* `+2.3338` octaves of deficit at `m = 6`, `kappa/eta = 0.161765`. *Fix:* SPEC 4.12.

FM 4.8 — THE CHORD READ AS A ROSTER PORTRAIT. *Signature:* a twelve-member group's chord presented as showing its variety. *Mechanism:* class saturation. At twelve members holding four states each, 5.98 of six classes are expected to be covered by chance alone. *Fix:* SPEC 4.8, and report `mu`.

FM 4.9 — DROP-ONE READ AS A RANKING. *Signature:* "the analysis says `A` is the weakest member." *Mechanism:* a drop-one row reports the effect of removing a set of pairs. The apparatus has no single-party term and cannot produce a statement about a person (the no-single-party-term prohibition [Curzi, 2026]). *Fix:* SPEC 4.11. Read the row as a statement about which pairs carry the geometry.

FM 4.10 — NESTING COSTED WITHOUT ITS LOSSES. *Signature:* a partition adopted for its link count, with `u_bar` assumed unchanged. *Mechanism:* removed links were carrying surplus. At high `u_bar` the loss dominates the saving. *Fix:* SPEC 4.13, and compute both terms.

FM 4.11 — THE AGGREGATED ORGANISATION. *Signature:* one deficit for a whole institution. *Mechanism:* CC-R16 — identical aggregates conceal topologies that lose 77% versus 21% of surplus to the same departure; CC-R10 — ratios do not average. *Fix:* SPEC 4.14. Per group, per period.

FM 4.12 — THE MARGINAL TEST USED AS AN EXCLUSION RULE. *Signature:* a STOP verdict treated as grounds for keeping someone out, or removing them. *Mechanism:* register slip. The test prices a link against a best-case surplus contribution and weighs nothing else whatsoever. *Fix:* §4.12's cautions, and the boundary statement of the minimum-frequency paper [Curzi, 2026]. A reading does not conclude that a relation should end and does not conclude anything about a person.

Worked example

THE CONSTRUCTED TRIAD — A, B, C CONSTRUCTED, and an explicitly labelled extension of the carried dyad of the carried dyad [Curzi, 2026]. A and B are the carried dyad exactly as found; C is a constructed third member. Parameters are chosen, not observed. **This is not an assessment of any person or group, no clinical, diagnostic or therapeutic claim is made, and the frequencies are essence frequencies.**

Inputs

```
A = {4: 0.7,  6: 0.5,  9: 0.4}                          raw L2 norm 0.948683
B = {4: 0.6,  6: 0.55, 9: 0.3,  10: 0.2}                raw L2 norm 0.890225
C = {3: 0.45, 4: 0.55, 6: 0.40, 7: 0.35, 11: 0.25}      raw L2 norm 0.921954

N = 12,  m = 3,  one declared period
kappa = 0.055 per link,  eta = 0.34,  delta = 0.09
```

Raw L2 norms are published because normalisation removes a free multiplier and the elicitation rule of [Curzi, 2026] rule requires it to stay recoverable.

Steps 3–4 — support and the pairwise list Every pair shares support — A and C on states 4 and 6, B and C on 4 and 6, A and B on 4, 6 and 9 — so SPEC 4.5's gate passes and no pair is orthogonal by construction.

```
u(A,B) = 0.068734665265     cos =  0.965021     F = 0.563085    f_pair =  1.4908 Hz
u(A,C) = 0.552647058824     cos =  0.668844     F = 1.957149    f_pair = 24.2264 Hz
u(B,C) = 0.550937094081     cos =  0.670122     F = 1.954489    f_pair = 24.0978 Hz

3 pairs, as m(m−1)/2 requires
```

The list already tells the story the mean will hide: one near-parallel pair and two well-separated ones.

Step 5 — aggregation, both ways

```
u_bar          = 0.390772939390
F( u_bar )     = 1.6674242020 nats      N_eff = 5.2985      f = 13.5718 Hz
mean of F(u)   = 1.4915742527 nats                          f =  9.5476 Hz

AGGREGATION GAP = 0.1758499493 nats = 0.507396 octaves = 609 cents
ratio of frequencies = 1.421482

R3 cross-check at roster scale:
    PR = 1.3522930095      m + 1 − m/PR = 1.7815458788
    1 + (m−1) u_bar        = 1.7815458788        error 0.00e+00
    M_eff = 1.7815 of a possible 3
```

Steps 6–8 — costs and the two ceiling tests

```
links   = m(m-1)/2 = 3
C       = 0.055 * 3 = 0.1650
c/eta   = 0.1650/0.34 = 0.485294

ROSTER CEILING (Thm 4.3): m(m-1) = 6  <=  2*0.34*2.484907/0.055 = 30.7225      PASS
CEILING TEST  (CCP-041 SPEC 4.1): 0.485294 <= 2.484907                            TUNABLE
BROOKS (Thm 4.2): r = 0.161765, max viable roster 4, actual 3                     WITHIN
```

Step 9 — the reading

```
f                =  13.5718 Hz         folded 434.296 Hz (j = 5)
f_min            =   1.2760 Hz         folded 326.656 Hz (j = 8)
octaves required =   1.400263
octaves present  =   4.811169
OCTAVE DEFICIT   =  -3.410906          (-4093.1 cents)
N_eff required   =   1.6247
pitch headroom   =   2.3588 octaves
Q_inf            =  +4.465825
identity check   : -2.885390 * (0.09/0.34) * 4.465825 = -3.410906
VERDICT: ACCUMULATES

conservative (mean-of-F): f = 9.5476 Hz, D = -2.903511, Q_inf = +3.801503
both aggregations ACCUMULATE -- verdict determined
```

Step 10 — the numerator

```
eta F - c  =  0.34 * 1.6674242020 - 0.1650  =  +0.401924        POSITIVE
```

So `delta` interventions are available here and would help — the ordering rule [Curzi, 2026] is satisfied and trust, in this group, multiplies a positive sign. That was not true of the same `A` and `B` without `C`.

The comparison that makes this section The carried dyad `{A, B}` at the cost structure of the sustainable-dyad paper [Curzi, 2026] — `c = 0.62`, `eta = 0.30` — read `+4.338420` octaves and DECAYS. The triad `{A, B, C}` at a per-link cost of 0.055 and `eta = 0.34` reads `-3.410906` and ACCUMULATES.

Two things changed, and honesty requires separating them. The geometry changed: `u_bar` rose from 0.068735 to 0.390773 because `C` is genuinely non-parallel to both. And the cost structure changed: this reading prices a link at 0.055 where the sustainable-dyad paper's [Curzi, 2026] dyad priced its single link at 0.62, a factor of 11.2727.

That is not an inconsistency; it is what a link costs in each situation, and the point of the comparison is that both terms have to be declared. A third member does not rescue a dyad by existing. It rescues it only where the link price is low enough that three links cost less than the surplus the third member's non-parallelism supplies. Priced at the carried dyad's own `kappa`, the same triad would carry `C = 3 * 0.62 = 1.86`, `c/eta = 6.2`, which exceeds `ln 12 = 2.484907` — a STRUCTURAL case with a deficit of `+13.078` octaves.

```
the same three people, two link prices

kappa = 0.055, eta = 0.34   C = 0.1650   c/eta = 0.485294   D = -3.410906   ACCUMULATES
kappa = 0.620, eta = 0.30   C = 1.8600   c/eta = 6.200000   D = +13.078     STRUCTURAL
```

The roster is not the variable that decides a group. The ratio of link price to conversion is.

The group chord

```
states spanned: 3, 4, 6, 7, 9, 10, 11      (seven of twelve)

 n     f = k n^2 (Hz)     folded (220–440)     j     odd part = class
 3          4.3508             278.453          6           3
 4          7.7348             247.514          5           1
 6         17.4033             278.453          4           3
 7         23.6879             379.006          4           7
 9         39.1575             313.260          3           9
10         48.3426             386.741          3           5
11         58.4945             233.978          2          11

distinct classes: {1, 3, 5, 7, 9, 11} -- six, which is ceil(12/2), the whole ceiling
states 3 and 6 share class 3: one of the seven states adds pitch but no new class
muddiness index mu = 3 / 6 = 0.500
```

Seven states, six classes, three members: the chord is already saturated at the class ceiling. The `ceil(N/2)` class count, result R1 of the mood-and-trait paper [Curzi, 2026], is doing all the work and the roster is doing none of it — a fourth member could add at most the one class already present twice, which is to say nothing.

Drop-one

set	u_bar	C	f (Hz)	D	Q_inf	change in D
{A,B,C}	0.390773	0.1650	13.5718	−3.410906	+4.4658	—
{B,C}	0.550937	0.0550	24.0978	−5.172708	+6.7725	−1.761802
{A,C}	0.552647	0.0550	24.2264	−5.180385	+6.7826	−1.769478
{A,B}	0.068735	0.0550	1.4908	−1.157965	+1.5161	+2.252941

The near-parallel pair is the one costing the group, and the reading says so twice. Removing either member of it raises `u_bar` by about 0.16 and removes a link. What the analysis does not say, cannot say, and must not be reported as saying, is anything at all about `A` or `B` as people — SPEC 4.11 and the boundary statement of the minimum-frequency paper [Curzi, 2026] both govern here, and the actionable reading is the sustainable-dyad paper's [Curzi, 2026]: `A` and `B` sit at `cos theta = 0.9650` and the question is what would raise `u` between them.

The card

```
------------------------------------------------------------------------
CONSTRUCTED TRIAD {A, B, C} -- extension of the carried dyad -- N = 12, m = 3
------------------------------------------------------------------------
pairwise u             0.068735 / 0.552647 / 0.550937
u_bar                  0.390773
F(u_bar)               1.667424 nats        mean-of-F   1.491574 nats
aggregation gap        0.507396 octaves
N_eff (category)       5.2985               M_eff (roster)  1.7815 of 3
f                      13.5718 Hz           folded 434.296 (j = 5)
kappa / links / C      0.055 / 3 / 0.1650
eta / delta            0.34 / 0.09
c/eta                  0.485294             ln 12 = 2.484907   TUNABLE
roster ceiling         6 <= 30.7225         PASS
Brooks max roster      4                    actual 3   WITHIN
f_min                  1.2760 Hz            folded 326.656 (j = 8)
octaves req / present  1.400263 / 4.811169
```

```
OCTAVE DEFICIT          −3.410906            (conservative −2.903511)
numerator eta F − c     +0.401924            POSITIVE
Q_inf                   +4.465825            (conservative +3.801503)
chord                   7 states, 6 classes, mu = 0.500
VERDICT                 ACCUMULATES, on both aggregations
-------------------------------------------------------------------------
ESSENCE frequencies, not physical resonances. Constructed case. Not an
assessment of any person or group. Not counselling; no clinical claim.
-------------------------------------------------------------------------
```

What this section established Four results, and the fourth is the one to carry forward.

One: intervals go as `m(m−1)/2`, classes are capped at `ceil(N/2) = 6`, and the ratio between them — the muddiness index — passes 1 at four members and never comes back.

Two: aggregation is a choice with a price. `F(u_bar)` and mean-of-`F` differed by 0.507396 octaves on three members and by 0.000001 octaves on the commons' fifteen pairs, and both must be published.

Three: cost is quadratic and certain; surplus is logarithmic at best and capped absolutely. The crossover `2 ln m = r m(m−1)` and the exact roster ceiling `m(m−1) <= 2 eta ln N / kappa` bracket every group this framework reads.

Four: a group's case is never a surplus case, because an orthogonal pair already reaches the ceiling. Whatever else a family, a team or a commons is for, the formalism prices only the stored order, and Section 5 takes those three structures and reads all of them — including one that cannot be repaired at all.

Section 5 reads three relational structures in full and carries the commons of the carried commons [Curzi, 2026], where the deficit is +6.4638 octaves and no single lever closes it. Section 6 takes the interventions.

5 Three relational structures read in full

Section 4 built the apparatus for reading a set larger than two. This section uses it three times, on three structures that differ in every parameter the formalism has, and reports what comes back.

The three are chosen because they are the structures people actually live inside, and because they fail differently. A family is a cluster geometry with expensive links and a long memory. A team is a selected geometry with an explicit conversion mechanism. A commons is a large flat roster whose members were chosen for agreement, and it is the one this section exists for.

The three readings come back as follows, and the third is the point:

STRUCTURE	m	u_bar	deficit D	Q_inf	verdict
the family	4	0.645160	+0.584737	−1.378049	DECAYS, marginally
the team	5	0.788228	−3.495327	+6.999131	ACCUMULATES
the commons	6	0.018614	+6.463800	−8.463000	DECAYS, and cannot be tuned

No single lever closes a 6.4638-octave gap. Raising the commons' `u` to 0.95 — near orthogonality, for six people drawn from one doctrine — closes it by 0.0332 octaves and no more. Raising `eta` to 0.90 leaves `+2.1075`. Cutting the per-link cost to 0.020 leaves `+2.0085`. There is no assignment of any one of the four levers that makes that commons sustainable, and the reason is not that the numbers are bad. It is that a group whose members were selected for agreement has almost no non-parallelism to convert, and the cost of maintaining fifteen links has to be paid against surplus that is not there.

Some systems cannot be tuned. They must be restructured, and the diagnostic is what tells you which case you are in.

A framework that only showed fixable examples would be lying by selection, and the honest thing to do with an instrument is to run it on a case it cannot help. This section does that, computes exactly how far every lever falls short, and then computes what restructuring would have to mean — which turns out to be a change of membership, not a change of arrangement.

Every case here is CONSTRUCTED and labelled. Nothing in this section is clinical, diagnostic or therapeutic; none of it is counselling; the twelve states are not diagnoses; and the frequencies are essence frequencies — no family, team, commons or person is claimed to resonate at any value.

The law

5.1 What a relational structure is, in this formalism A structure is five declarations and nothing else. Everything this section says about families, teams and commons is a statement about one of these five, and where it is not, the section says so.

```
1. THE ROSTER            m, the number of members
2. THE GEOMETRY          the m(m-1)/2 pairwise u values, and their mean u_bar
3. THE COST TOPOLOGY     which pairs maintain a link, and kappa for each
4. THE CONVERSION        eta -- the fraction of interaction becoming durable structure
5. THE LOSS              delta -- the per-period rate at which stored order leaks
```

A "family," a "team" and a "commons" are not primitives in this formalism; they are characteristic patterns across those five declarations, and the pattern is what does the work. Nothing is being claimed about kinship, employment or land tenure. The names are labels for parameter signatures, and the signatures are what follow.

SIGNATURE	roster	geometry	topology	eta	delta
family	small	clustered; one near-parallel pair	dense, high kappa per link	moderate	LOW
team	small	selected for non-parallelism	dense, low kappa	HIGH	moderate
commons	larger	near-parallel by membership rule	flat, uniform kappa	low	moderate

The family's signature is a low `delta` — shared history is a memory that does not leak fast — paid for with a high `kappa`, because family links are maintained under conditions nobody chose. The team's signature is a high `eta`, because a team has apparatus for turning interaction into artefacts, and a geometry that was selected rather than inherited. **The commons' signature is a membership rule that selects for agreement, and that rule is the whole of its problem.**

5.2 The Parallel Wall, priced Series B names the Parallel Wall: a set of parties whose representations point the same way returns no surplus however many of them there are and however hard they work. The frequency register prices it exactly.

```
six members drawn from one doctrine       CARRIED COMMONS, CONSTRUCTED
```

```
mean pairwise u  =  0.018614    over 15 pairs
F                =  0.186272 nats
N_eff            =  1.204754        of a possible 12
f                =  0.7017 Hz       folded 359.270 Hz (j = 9)
```

This is the corpus's own result restated in 1952: diversification benefit depends on covariance, not on count [Markowitz, 1952, 1959], and the same point is standard in the ecological stability literature, where the portfolio effect requires asynchrony rather than richness [Doak et al., 1998, Schindler et al., 2010, Tilman et al., 1998]. The concentration instrument reaches it a third way: a naive count of four jurisdictions falls to an effective count of 1.00 once they are recognised as one bloc [Curzi, 2026], and even the most generous counterfactual splitting leaves an effective count of 1.423998 against a designed 2.941176 [Curzi, 2026].

`N_eff = 1.2048` is the figure to sit with. Six people, fifteen relationships, and 1.20 people's worth of distinct view among them. They are not lazy, not unintelligent, and not in conflict. They were selected — by a doctrine, a hiring filter, a shared training, a common origin — for pointing the same way, and the selection worked. Selection of that kind is a documented mechanism: homophilous sorting produces near-uniform local composition from mild individual preferences [Schelling, 1971, 1978], and shared training removes the independence on which aggregation of judgments depends [Lorenz et al., 2011, Hong and Page, 2004].

The frequency `0.7017 Hz` is 0.5375 octaves above the ground essence `k`. A group of six that has travelled just over half an octave from "no relationship at all" is a group whose relationships are almost entirely redundant with each other, which is the consonance-redundancy identity R2 [Curzi, 2026] at group scope: the most consonant pairs are exactly the pitch-class collisions, so a maximally consonant set is a maximally redundant one.

5.3 The commons reading in full

```
THE COMMONS -- CONSTRUCTED, THE CARRIED COMMONS
m = 6, N = 12, one declared period
kappa = 0.055 per link,  15 links,  C = 0.8250,  eta = 0.34,  delta = 0.09

u_bar            =  0.018614
F                =  0.186272 nats          N_eff = 1.204754
f                =  0.7017 Hz              folded 359.270 Hz (j = 9)
octaves present  =  0.5375

c/eta            =  0.8250/0.34 = 2.426471
ln 12            =  2.484906649788
margin           =  +0.058436 nats                    BELOW the ceiling: TUNABLE
roster ceiling   :  m(m-1) = 30  <=  2*0.34*ln12/0.055 = 30.7225        PASS, barely

f_min            =  61.9349 Hz             folded 247.740 Hz (j = 2)
octaves required =  7.0013
OCTAVE DEFICIT   =  +6.4638                (+7756.6 cents)
N_eff required   =  11.3189   of a possible 12
pitch headroom   =  6.6324 octaves
numerator eta F - c = -0.761666            NEGATIVE
Q_inf            =  -8.4630
VERDICT: DECAYS
```

`N_eff required = 11.3189` against `N_eff achieved = 1.2048` is the reading in one line. These six must field 11.32 categories' worth of distinct view to break even and they field 1.20.

5.4 No single lever closes it, and here is each one falling short Every figure below is computed from the same starting point, varying exactly one parameter.

LEVER 1 — raise `u`.

```
u -> 0.20   f =  4.950 Hz    deficit  +3.6452
u -> 0.40   f = 14.097 Hz    deficit  +2.1354
u -> 0.60   f = 27.923 Hz    deficit  +1.1493
u -> 0.80   f = 46.428 Hz    deficit  +0.4158
u -> 0.95   f = 63.378 Hz    deficit  −0.0332      closes, by 0.0332 octaves
u -> 1.00   f = 69.613 Hz    deficit  −0.1686      the absolute maximum available
```

The hold line is at `u* = (e^2.426471 - 1)/11 = 0.938078`, and the whole reserve above it is 0.1686 octaves. So this lever technically closes the gap, and the closure is worthless: it requires that six people selected from one doctrine reach 93.8% non-parallelism, which is not a tuning of a group but a different group. And it arrives with 40 cents of margin, which the minimum-frequency paper [Curzi, 2026] §4.20 notes is about a fifth of a semitone — no margin at all against any increase in cost.

LEVER 2 — raise `eta`.

```
eta -> 0.34   c/eta = 2.426471   deficit  +6.4638
eta -> 0.50   c/eta = 1.650000   deficit  +4.2234
eta -> 0.70   c/eta = 1.178571   deficit  +2.8632
eta -> 0.90   c/eta = 0.916667   deficit  +2.1075
eta -> 1.00   c/eta = 0.825000   deficit  +1.8430      perfect conversion, still failing
```

At `eta = 1.00` — every unit of interaction becoming permanent structure, nothing lost, which is physically the end of the scale — the commons is still short by 1.8430 octaves. The lever does not merely fall short; it is exhausted while the gap is still nearly two octaves wide.

LEVER 3 — cut `kappa`.

```
kappa -> 0.055  (C = 0.825)   deficit  +6.4638
kappa -> 0.040  (C = 0.600)   deficit  +4.5544
kappa -> 0.030  (C = 0.450)   deficit  +3.2814
kappa -> 0.020  (C = 0.300)   deficit  +2.0085
kappa -> 0.010  (C = 0.150)   deficit  +0.7355
kappa -> 0.005  (C = 0.075)   deficit  +0.0990      still failing, at a 91% cut
```

A 91% reduction in per-link cost leaves the commons 0.0990 octaves short. The cost side can be cut almost to nothing and the group still does not sustain, because the surplus it is trying to pay for is 0.186272 nats.

LEVER 4 — raise `N`.

```
N -> 12    F = 0.186272   present 0.5375   deficit  +6.4638
N -> 17    F = 0.260689   present 0.7522   deficit  +6.2491
N -> 25    F = 0.369310   present 1.0656   deficit  +5.9357
N -> 41    F = 0.556502   present 1.6057   deficit  +5.3956
N -> 64    F = 0.775962   present 2.2390   deficit  +4.7624
N -> 144   F = 1.297955   present 3.7451   deficit  +3.2562
```

One hundred and forty-four categories — twelve times the ladder this framework uses — leaves the commons 3.2562 octaves short. The minimum-frequency paper [Curzi, 2026] gives the mechanism: `N` enters `F` only multiplied by `u`, so adding categories to a near-parallel set adds dimensions nobody is using. At `u = 0.018614` a new category buys 0.0446 octaves. And the route crosses the `N = 17` cliff at the second row, where the folded lattice's minimum separation collapses from 34.9756 cents to 6.0008 and the symbol rate falls by a factor of exactly 8 (the mood-and-trait paper [Curzi, 2026] owns the regime statement).

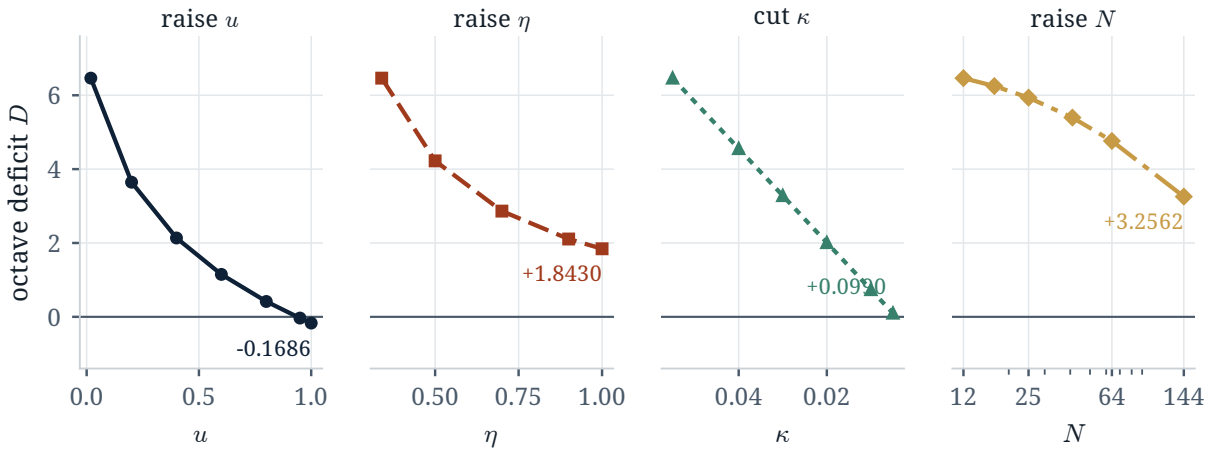


Figure 45.3: **Every lever swept to its physical limit, and every one short.** The constructed commons at a deficit of $+6.4638$ octaves, with each of the four levers driven to the end of its range while the other three are held. Only u crosses $D = 0$, at -0.1686 octaves and only by replacing the membership; perfect conversion leaves $+1.8430$, a 91 per cent cost cut leaves $+0.0990$, and twelve times the category count leaves $+3.2562$. Four panels because the four levers are in four different units. Computed from the lever tables of §5.4; script `figs/src/CCP-045_fig03.py`.

> **THE FOUR-LEVER RESULT, COMPUTED: at `u = 1` the gap is closed by 0.1686 octaves; at `eta = 1` it is still open by 1.8430; at a 91% cost cut it is still open by 0.0990; at twelve times the category count it is still open by 3.2562. Only one lever closes it alone, and only by replacing the membership.**

5.5 The reachability test — a case can be formally tunable and practically not The classification of the minimum-frequency paper [Curzi, 2026] returns TUNABLE for the commons, because `c/eta = 2.426471` is below `ln 12 = 2.484907`. That verdict is correct and it is not enough, and this section adds the test that completes it.

DEFINITION 5.1 — REACHABILITY. A tunable case is REACHABLE if the required geometry `u* = (e^(c/eta) - 1)/(N-1)` lies within the range the parties can honestly represent, and UNREACHABLE otherwise. Reachability is not a property of the formalism; it is a declared judgment about the parties, and it must be declared rather than assumed.

```
THE THREE-WAY CLASSIFICATION, EXTENDED

D <= 0                                SUSTAINABLE
D > 0, c/eta <= ln N, u* reachable    TUNABLE AND REACHABLE     tune it
D > 0, c/eta <= ln N, u* not          TUNABLE, UNREACHABLE      restructure it
c/eta > ln N                          STRUCTURAL                restructure it
```

Computed for the three structures of this section:

```
structure     c/eta       ln N       u* required     u_bar now     classification
family      2.294118   2.484907      0.810517        0.645160     tunable, reachable
team        1.057692   2.484907      0.134733        0.788228     sustainable already
commons     2.426471   2.484907      0.938078        0.018614     tunable, UNREACHABLE
```

The family needs a 25.6% rise in `u_bar` and the commons needs a factor of 50.4. The formal test does not distinguish them; the reachability test does, and it is the difference between an afternoon's work and a different organisation.

The margin figure is worth reporting alongside. The commons sits 0.058436 nats below the ceiling — 0.1686 octaves. Theorem 4.3 says the same thing about its roster: `m(m−1) = 30` against a bound of 30.7225. A seventh member takes `m(m−1)` to 42 and the case becomes STRUCTURAL outright, at which point `u` is unavailable at every value whatsoever. The commons is not near a boundary in one variable; it is near it in both, which is what a group at the end of its arrangement looks like from inside the arithmetic.

5.6 Ostrom, applied where it applies and priced where it does not

Ostrom's design principles for durable common-pool institutions [Ostrom, 1990, 2010] are the standard answer to the tragedy argument [Hardin, 1968], and the principles have been subjected to a systematic empirical review that finds broad but uneven support [Cox et al., 2010]. Their dynamical relatives — reciprocity under repeated interaction [Axelrod, 1984, Axelrod and Hamilton, 1981, Nowak and Sigmund, 2005] and the social-capital literature [Putnam et al., 1993, Gambetta, 1988, Fukuyama, 2001] — are the empirical warrant for treating monitoring, sanctioning and nesting as arrangements that change parameters at all. Series B mapped them to the storage-law parameters; this section carries the mapping into the frequency register and prices three of them on this commons. The mapping is Series B's and is badged as a correspondence of dynamics, not of agency (CC-R16).

P4 — monitoring. Series B's finding is specific and slightly awkward: monitoring lowers the variable cost `c_v` and does **not** raise the trust parameter `a`. In this register that means P4 is a `c` lever, not a `delta` lever, and Theorem 4.9 applies to the half of it people expect to matter: **an intervention on `delta` does not move the deficit by any amount whatsoever.** On this commons P4 helps only through the `c` column of §5.4, where a 91% cut still leaves 0.0990 octaves.

P5 — graduated sanctions. Priced in Series B at −7.7% of `eta S`. Computed here:

```
eta * F        =  0.34 * 0.186272 =  0.063332
P5 cost        =  0.077 * 0.063332 = 0.004877
numerator      = −0.761666  ->  −0.766544
Q_inf          =  −8.462972  ->  −8.517157
```

P5 makes this commons slightly worse, because it is a cost levied against a surplus that does not exist. That is not an argument against graduated sanctions; it is the ordinary observation that a governance mechanism priced as a fraction of surplus is affordable in proportion to the surplus, and this commons has 0.186272 nats of it.

P8 — nested enterprises. This is the only Ostrom principle that changes the exponent rather than a coefficient, and Series B records it as turning `O(N^2)` coordination into `O(N)`. Computed on six members:

```
flat six                     15 links    C = 0.8250    c/eta = 2.426471    D = +6.4638
two triads + one bridge       7 links    C = 0.3850    c/eta = 1.132353    D = +2.7298
three pairs + three bridges   6 links    C = 0.3300    c/eta = 0.970588    D = +2.2630
```

Nesting removes 53.3% of the cost and leaves the commons 2.7298 octaves short. It is the largest single structural improvement available and it does not close the gap, because the gap is not principally a cost problem.

Where Ostrom applies here and where it does not, stated plainly. The eight principles are arrangements: boundaries, congruence, collective choice, monitoring, sanctions, conflict resolution, recognition, nesting. Every one of them acts on `c`, on `eta`, or on `delta`, and not one of them acts on `u`. There is no design principle for making the members of a group less alike, because that is not a design question about the institution — it is a question about who is in it.

5.7 Restructuring, defined against tuning The section's central distinction, stated so it can be applied.

```
TUNING          changing eta, c, delta, or N, with the membership held fixed
RESTRUCTURING   changing the membership, the topology, or the boundary of the group
```

`u` is the only parameter that a group cannot tune, because `u` is a property of who its members are. Everything else in the five declarations of §5.1 is an arrangement and can be rearranged. The geometry is a fact about the roster.

That gives the operational rule:

> **If the deficit can be closed without changing `u`, it is a tuning problem. If it cannot, it is a membership problem, and no amount of better arrangement will substitute.**

Applied to the three cases: the family's deficit of `+0.584737` closes on `eta` alone with a 9.7% improvement (§5.10), so it is a tuning problem. The team is already sustainable. The commons needs `u` to rise by a factor of 50.4 and therefore is not a tuning problem at all, which is what §5.4's four failed levers were demonstrating one at a time.

5.8 Honesty clauses **(1)** Essence frequencies, not physical resonances. Nothing here is claimed to vibrate and nothing is to be played at anyone (the no-sounding prohibition [Curzi, 2026]; the no-subliminals rule is absolute).

(2) No clinical, diagnostic or therapeutic claim. The twelve states are not diagnoses. **This is not counselling, therapy or family assessment of any kind, and a reader in difficulty with a family, a team or a community should seek help from someone qualified to give it, which this framework is not.**

(3) Every case is constructed and labelled. The family and the team are explicitly labelled extensions of the carried dyad of the carried dyad [Curzi, 2026]; the commons is the carried commons [Curzi, 2026].

(4) No single-party term exists in the apparatus. No reading in this section licenses a conclusion about any person's value, contribution, or fitness for a group.

(5) Consent is structural and per-member (SPEC 4.1), and §5.14 works out what that requires when a decline changes the reading.

(6) The reading inherits its inputs. A group reading rests on `m(m−1)/2` measured values plus three declared parameters, and its uncertainty is reported (the sensitivity-reporting rule [Curzi, 2026]).

The engineering

5.9 The three structures side by side

	FAMILY	TEAM	COMMONS
m	4	5	6
links	6	10	15
kappa per link	0.130	0.055	0.055
C	0.7800	0.5500	0.8250
eta	0.34	0.52	0.34
delta	0.05	0.09	0.09
u_bar	0.645160	0.788228	0.018614
F(u_bar) nats	2.091463	2.269080	0.186272

```
mean-of-F nats            1.954682      2.258759       0.186276
aggregation gap (oct)     0.394666      0.029780       0.000001
N_eff                     8.096755      9.670503       1.204754
f (Hz)                   31.6922       45.2093         0.7017
  folded                253.537 (j=3) 361.674 (j=3)  359.270 (j=9)
c/eta                     2.294118      1.057692       2.426471
ceiling margin (nats)    +0.190789     +1.427214      +0.058436
roster ceiling m(m-1)     12 <= 12.998  20 <= 46.987   30 <= 30.723
Brooks max roster              2             6              4
muddiness index mu        1.000         1.667          2.500
f_min (Hz)               47.5308        4.0089        61.9349
octaves req / present  6.6194/6.0347 3.0519/6.5472  7.0013/0.5375
OCTAVE DEFICIT           +0.584737     -3.495327     +6.463800
numerator eta F - c      -0.068902     +0.629922     -0.761666
Q_inf                    -1.378049     +6.999131     -8.463000
VERDICT                   DECAYS       ACCUMULATES     DECAYS
classification        tunable/reach.   sustainable  tunable/UNREACHABLE
------------------------------------------------------------------------
```

The corpus has one nearby measured case, and it points the same way: a consortium at $u = 0.52$, collecting 73.8 per cent of its ceiling, still reached a steady state of -0.0836 and headroom of -0.4582, with the binding fault a conversion parameter rather than difference [Curzi, 2026s].

Three structures, three verdicts, and the parameter that separates them is not the one anybody would guess.

It is not roster size — the failing commons is only two larger than the working team. It is not cost per link — the commons and the team pay the identical 0.055. It is `u_bar`, and the commons' is 42 times smaller than the team's.

5.10 The family — a cluster geometry with a cheap fix and an unresolved verdict The family's pairwise list shows the cluster structure §4.13b describes:

```
u(A,B) = 0.068735      the near-parallel pair
u(A,D) = 0.710180      u(A,E) = 0.711015
u(B,D) = 0.723966      u(B,E) = 0.740352
u(D,E) = 0.916710
```

One tight pair, five wide ones. That dispersion produces an aggregation gap of 0.394666 octaves — substantial, and both readings must be published. Here they agree on the verdict:

```
optimistic  F(u_bar) = 2.091463    f = 31.6922 Hz    D = +0.584737    Q_inf = -1.378049
conservative mean-of-F = 1.954682  f = 24.1071 Hz    D = +0.979403    Q_inf = -2.308160
```

Both decay, so the verdict is determined, and it is a small deficit. The levers:

```
L1  u   must reach 0.810517 from 0.645160      +0.165357, a factor of 1.2563
L2  eta must reach 0.372945 from 0.34          +9.69%           REACHABLE, and cheap
L3  C   must fall to 0.711098 from 0.78        -8.83% -> kappa 0.118516
L4  N   must reach 14.8193 -> 15 categories    still below the N = 17 cliff
```

A 9.7% improvement in conversion efficiency closes this family's entire deficit, which is what a small deficit on a high-`c/eta` case looks like: the elasticity `dD/d(ln eta) = -6.619424` octaves per e-fold is enormous, and only 0.584737 octaves are needed.

And then the honest complication, which is why this case is in the section. The minimum-frequency paper's [Curzi, 2026] SPEC 4.19 requires the sensitivity to be published, and here it decides the reading:

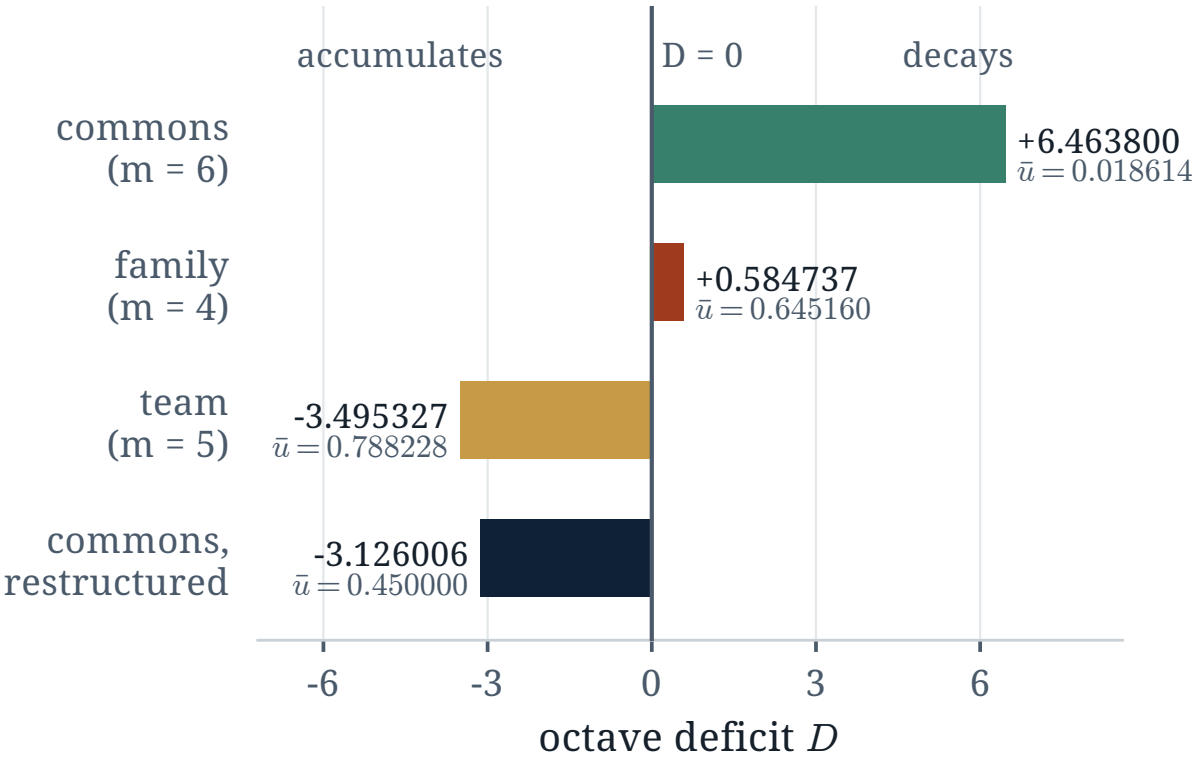


Figure 45.4: **The parameter that separates a working structure from a failing one is the one no arrangement can change.** Octave deficits for the three constructed structures and for the restructured commons, against the sustainability line $D = 0$. The commons and the team pay the identical per-link cost of 0.055; what differs is $\bar{u}$, and the commons' is 42 times smaller. Computed from the side-by-side table of §5.9 and the restructure of §5.13; script `figs/src/CCP-045_fig04.py`.

```
dD/d eta = -19.4689 octaves per unit eta

eta known to +/- 0.02  ->  D = +0.5847 +/- 0.3894    verdict resolved
eta known to +/- 0.05  ->  D = +0.5847 +/- 0.9734    VERDICT NOT RESOLVED
eta known to +/- 0.10  ->  D = +0.5847 +/- 1.9469    VERDICT NOT RESOLVED

dD/dC = +8.4864 octaves per unit C
C known to +/- 0.05    ->  D = +0.5847 +/- 0.4243    verdict resolved
C known to +/- 0.10    ->  D = +0.5847 +/- 0.8486    VERDICT NOT RESOLVED

to resolve the sign, eta must be known to better than +/- 0.030034 -- 8.8% of its value
```

Unless `eta` is known to within about 9% of itself, this family's verdict is not determined and the correct action is to measure, not to intervene. That is the reading a marginal case deserves, and Section 6 makes it a numbered step rather than a caution.

The family's low `delta = 0.05` is its characteristic parameter and it is doing exactly what the sign-multiplication result [Curzi, 2026] says. The numerator is `-0.068902`, negative. A family's long memory is amplifying a small negative sign, and would amplify it further if the memory got longer — `delta` to 0.03 takes `Q_inf` from `-1.378049` to `-2.296749`, with the deficit unmoved at `+0.584737`. CC-R15 in miniature, on the structure most likely to reach for continuity as a repair.

5.11 The team — what a selected geometry and an explicit mechanism buy

```
u(A,B') = 0.692372     u(A,C) = 0.552647     u(A,S) = 0.897875     u(A,T) = 0.679778
u(B',C) = 0.851659     u(B',S) = 0.953352    u(B',T) = 0.748794
u(C,S)  = 0.756702     u(C,T) = 0.786182     u(S,T) = 0.962914

u_bar = 0.788228     10 pairs, all between 0.55 and 0.97
```

The team's pairwise list is tight at a high level, which is the geometry the family's is not. Its aggregation gap is 0.029780 octaves — negligible, because the values are close together, exactly as §4.3b predicts.

Its reading is comfortable: `D = -3.495327`, `Q_inf = +6.999131`, and both aggregations agree (`-3.465547` and `+6.939499`). Three parameters explain it and all three are arrangements or selections rather than luck.

`eta = 0.52`, against the family's 0.34 and the commons' 0.34. A team has apparatus — the decision log of the sustainable-dyad paper's [Curzi, 2026] phase 2b, the repository, the standing review — and the apparatus is what conversion efficiency *is*.

`kappa = 0.055` against the family's 0.130. A team's links are maintained on purpose, in declared slots, with agendas. A family's are maintained under conditions no member chose.

`u_bar = 0.788228`, because a team is composed. Its members were selected partly for what they bring that others do not, which is a membership rule pointing the opposite way from the commons'.

The honest counterweight, and it belongs here rather than in a footnote: the team's pitch headroom is 0.6227 octaves. It has spent 91% of the range twelve categories make available, and its next increment of surplus is not purchasable from `u` at this `N`. That is the ceiling of the minimum-frequency paper [Curzi, 2026] Theorem 4.6 arriving, and it arrives at every well-tuned system.

5.11b The team's drop-one, and the result that settles the ethics Run the drop-one analysis of §4.13 on the team. `kappa = 0.055`, `eta = 0.52`, `delta = 0.09`, four links removed with each member.

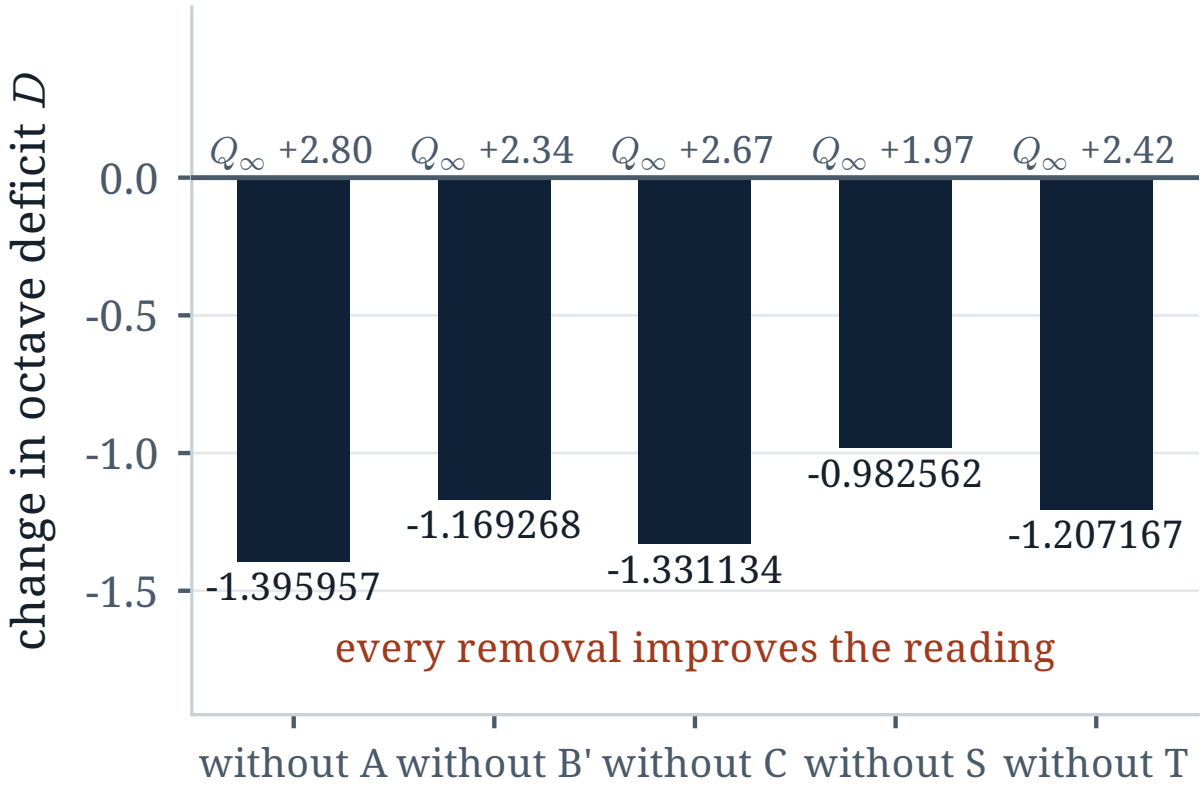


Figure 45.5: **A drop-one row on a healthy group carries no information about the member it names.** All five removals from the constructed team lower the octave deficit and raise the steady state, by 1.97 to 2.80 units. The mechanism is arithmetic: F is concave and the team is far up the curve, so an additional member can contribute almost no surplus, while each member's four links cost $4 \times 0.055/0.52 = 0.423077$ nats. Computed from the drop-one table of §5.11b; script `figs/src/CCP-045_fig05.py`.

set	u_bar	links	C	f (Hz)	D	Q_inf	change in D
all five	0.788228	10	0.5500	45.2093	−3.495327	+6.999131	–
without A	0.843267	6	0.3300	51.0473	−4.891284	+9.794430	−1.395957
without B'	0.772683	6	0.3300	43.6247	−4.664595	+9.340502	−1.169268
without C	0.822514	6	0.3300	48.8044	−4.826461	+9.664627	−1.331134
without S	0.718572	6	0.3300	38.3291	−4.477889	+8.966638	−0.982562
without T	0.784101	6	0.3300	44.7859	−4.702494	+9.416392	−1.207167

Every single row improves the reading. Not one member of a well-composed, comfortably accumulating team is carrying their four links, and removing any of them — the best-matched, the worst-matched, any of them — raises `Q_inf` by between 1.97 and 2.80.

The mechanism is arithmetic and it is worth stating carefully because it is the general case rather than a quirk of these five vectors. The team sits at `u_bar = 0.788228` with 0.6227 octaves of pitch headroom. **Near the ceiling, an additional member can contribute almost no surplus, because there is almost no surplus left to contribute** — `F` is concave and the team is far up the curve. Each member's four links, meanwhile, cost `4 * 0.055 / 0.52 = 0.423077` nats, which is 1.2207 octaves of required pitch. Cost is linear in links and surplus is exhausted, so the trade goes one way.

That result settles a question the previous section only argued. If the drop-one analysis improves on *every* member of a healthy team, then a favourable drop-one row carries no information whatever about the member it names. It is a property of where the group sits on the surplus curve. **A ranking derived from it would be a ranking of nothing, and SPEC 4.11's**

prohibition is therefore not a piece of decorum — it is a statement that the numbers do not contain what somebody might try to read out of them.

The constructive reading of the same table is different and is available: a team at `u_bar = 0.788228` and 0.62 octaves of headroom has reached the useful limit of composition at `N = 12`, and its remaining levers are `eta`, `c`, and genuinely new categories under the state-ordering material's [Curzi, 2026] validity conditions. That is what the rows are telling anyone; it is not about `A`, `B'`, `C`, `S` or `T`.

5.12 The commons — what the doctrine costs, in one number

The commons' problem is `u_bar = 0.018614`, and the cost of it can be stated as a single interval.

```
the commons as found                f = 0.7017 Hz     0.5375 octaves above k
the same six with u_bar = 0.788228  f = 45.209 Hz     6.5472 octaves above k
                                                      ------------------------
the price of the membership rule                      6.0097 octaves
```

Six octaves of available order, forgone at the point of selection, before anybody did anything. Every subsequent decision the commons makes about its arrangement is being made downstream of that, and none of the arrangements can reach back past it.

This is the sense in which the commons is not a badly run group. Its `kappa` is the same as the team's. Its `eta` is the same as the family's. Its `delta` is the same as the team's. The only parameter on which the commons differs materially from a structure that works is the one no arrangement can change.

5.12b Can the commons shrink its way out?

If the cost is quadratic in the roster and the surplus is not, the obvious remedy is to be smaller. Computed, taking the six constructed members in order and reading each prefix:

m	u_bar	links	C	D	Q_inf	verdict
2	0.018644	1	0.0550	−0.071526	+0.093647	holds, barely
3	0.018548	3	0.1650	+0.864532	−1.131912	DECAYS
4	0.018587	6	0.3300	+2.263763	−2.963896	DECAYS
5	0.018583	10	0.5500	+4.130888	−5.408481	DECAYS
6	0.018614	15	0.8250	+6.463833	−8.462955	DECAYS

`u_bar` barely moves — it stays inside a band of width 0.000096 across every prefix, because every member is drawn from the same doctrine and any subset is as parallel as the whole. So the entire column of deficits is driven by the link count and nothing else.

The largest sustainable group drawn from this membership is two, and even the pair only clears its threshold by 0.0715 octaves. Three members from one doctrine already decay. The commons cannot shrink its way out any more than it can convert, cut or re-categorise its way out, and for the same underlying reason: `F = 0.186272` nats is what the membership rule left behind, and every arrangement is an attempt to pay fifteen, ten, six, three or one links out of it.

That completes the audit. **Four levers, three Ostrom principles, two nesting topologies and five roster sizes have now been computed on this commons, and exactly one of them closes the gap — `u`, which is the parameter it does not control.**

5.13 The restructure, computed

If tuning cannot help, what would? The answer has to be a change in membership, topology, or both, and it should be computed rather than gestured at.

CONSTRUCTED. Three changes, taken together:

1. **Recruit outside the doctrine.** The membership rule is changed so that new members are selected for what they bring that the existing members do not. Suppose this moves `u_bar` from 0.018614 to 0.45 — a large change, and one that is a change of group rather than of arrangement.
2. **Nest.** Two triads plus one bridge, per Ostrom P8: 15 links become 7.
3. **Adopt a conversion mechanism.** `eta` from 0.34 to 0.55 — a decision log, minutes that bind, a repository. This is the team's parameter and it is available to any group that will keep records.

```
u_bar            =  0.45
F                =  1.783391 nats        N_eff = 5.95
f                = 17.1145 Hz            folded 273.832 Hz (j = 4)
links            =  7        C = 0.055 * 7 = 0.3850
eta              =  0.55     c/eta = 0.700000
ceiling test     :  0.700000 <= 2.484907        comfortably below
roster ceiling   :  m(m-1) = 30 <= 2*0.55*ln12/0.055 = 49.6981      PASS
f_min            =  1.9604 Hz            folded 250.930 Hz (j = 7)
octaves required =  2.019773
octaves present  =  5.145779
OCTAVE DEFICIT   = -3.126006
numerator        = +0.595865             POSITIVE
Q_inf (delta 0.09) = +6.620724
pitch headroom   =  2.0241 octaves
VERDICT: ACCUMULATES
```

A swing of 9.5898 octaves and `Q_inf` from `-8.4630` to `+6.6207`. And the decomposition says which change did it:

```
from the cost ratio:  2.885390 * (0.700000 - 2.426471)  =  -4.981543 octaves
from the surplus   : -2.885390 * (1.783391 - 0.186272)  =  -4.608257 octaves
total                                                      -9.589800 octaves
```

The decomposition is the same shape as the worked room of Series B, where three moves — raising categories from two to five, cutting an acoustic cost term from 0.70 to 0.15, and retiring four of nine dormant agreements — closed 66 per cent of a -1.396 headroom gap [Curzi, 2026d].

51.9% from the cost side and 48.1% from the geometry — and the cost side is only available at that size because the nesting removed eight links. No proper subset of the three does as well. The recruitment plus the log, at fifteen links, leaves `-0.817694` octaves — sustainable, but with a quarter of the margin. The recruitment plus the nesting, without the log, leaves `-1.878499`. The recruitment alone, at fifteen links and `eta = 0.34`, leaves `+1.855535` and fails. And the two arrangement changes without the recruitment leave `+2.729803` (nesting alone) and `+3.790608` (the log alone), both failing badly — which is §5.7's rule arriving as arithmetic: without the change in membership, no combination of arrangements reaches the threshold.

And the honest note on what has happened: this is not the same commons. Its membership rule is different, its topology is different, and its records are different. The formalism cannot tell anyone whether that trade is worth making, because it contains no term for what the doctrine was for — the boundary statement of the minimum-frequency paper [Curzi, 2026] governs and does not soften here. What it can say is that the group as constituted arrives at `Q_inf = -8.4630` if its parameters hold, and that no arrangement of those parameters changes it.

5.14 The right to decline, and the arithmetic nobody wants Consent at group scope is per-member and revocable (SPEC 4.1). That raises a question the formalism answers uncomfortably, and the answer has to be printed rather than avoided.

Suppose member `E` of the constructed family declines to take part. Recomputing on `{A, B, D}`:

```
u(A,B) = 0.068735    u(A,D) = 0.710180    u(B,D) = 0.723966
u_bar  = 0.500960    F = 1.873426 nats    N_eff = 6.510562
links  = 3           C = 0.13 * 3 = 0.3900     c/eta = 1.147059
f      = 20.4912 Hz
OCTAVE DEFICIT = -2.095851      (was +0.584737)
Q_inf          = +4.939294      (was -1.378049)
VERDICT: ACCUMULATES            (was DECAYS)

E's decline IMPROVED the reading by 2.680588 octaves
```

The decline flipped the verdict, and it flipped it in the group's favour. Three links vanish with `E`, and `E`'s pairs — 0.711015 and 0.740352 — were above the mean but not by enough to pay for three links at `kappa = 0.13`.

That result must be handled exactly, and the exact handling is this. The formalism computes the stored order of an arrangement per period. It has no term for what `E` is for, what `E` contributes that is not stored order, what `E`'s participation means to `E`, or what the group would be without `E`. **It weighs none of that, it acquires no such term by being applied to a family, and a number that improves when someone leaves is not an argument that anyone should leave.**

Three rules follow and they are all binding.

The framework's founding treatment of the returned difference states the constraint the same way and refuses the weaker readings of it [Curzi, 2026], and the disclosure principle is the old one: a procedure whose safety depends on its own concealment is not safe [Kerckhoffs, 1883, Brandeis, 1914].

One: the right to decline is the member's, never the group's. A decline is an exercise of consent. Nothing in this apparatus converts a favourable drop-one row into a licence to exclude anyone, and SPEC 4.11 states the same prohibition for the drop-one analysis generally.

Two: a decline is not a finding. SPEC 4.1 requires that a member's withdrawal ends their inclusion in the reading and is not reported as a result about them. The reading after a decline is a reading of a different group, and must be labelled as one.

Three: the group must not be told what a decline would do before the decision is made. Reporting "the reading improves by 2.68 octaves if you withdraw" to a member weighing withdrawal converts an arithmetic result into pressure, and pressure removes the ability to decline freely. Section 6 makes this a numbered step: the counterfactual is computed after the decision, or not at all.

5.15 Reading a structure you are inside

One last engineering point, and it is the one most likely to be needed in practice, because nobody reads a family or a team from outside it.

An analyst who is a member of the group being read is part of the apparatus. Three specific consequences follow, and all three are handled by declaration rather than by correction.

The operator's own vector is one of the `m`, so their encoding enters `u_bar` directly and the result is sensitive to it. At `m = 4` a single member's vector participates in half the pairs. The requirement is simply that the analyst's own vector be elicited by the same procedure as everyone else's, by someone else where possible, and that this be stated in the report.

The operator estimates `eta`, `c` and `delta` for an arrangement they are inside, and the estimation practice of [Curzi, 2026]'s estimation practice is harder to apply honestly to one's own arrangement than to another's. The family case shows what is at stake: the verdict is unresolved unless `eta` is known to ±0.030034, and self-estimated `eta` is exactly the sort of quantity that comes back optimistic. The countermeasure is the count — enumerate what survived the period rather than judging how well things are going.

And the analyst has an interest in the answer, which the formalism cannot detect and does not correct for. CC-R13's finding stands here in its general form: design controls, not more computation. Read blind where possible, have another member encode the vectors, publish the raw norms, and report the sensitivity so that a reader can see how much of the verdict rests on the least certain input.

None of this makes an internal reading invalid. It makes it a reading with a declared observer, and Series B's the corpus's Series B already records that the analyst's own state is part of the apparatus — one of the four things the alchemical tradition got right, and the one most often dropped by its successors.

The specification

SPEC 5.1 — DECLARE ALL FIVE. Roster, geometry, cost topology, conversion, loss. A structure named without its five declarations is a label, not a reading.

SPEC 5.2 — PER-MEMBER CONSENT, REVOCABLE, AND UNCONDITIONED. Every member agrees separately; any may withdraw at any time; **and no member is shown the counterfactual value of their own withdrawal before deciding.**

SPEC 5.3 — REPORT BOTH AGGREGATIONS with the gap in octaves (SPEC 4.4). Where the verdict differs between them, report NOT DETERMINED.

SPEC 5.4 — RUN THE REACHABILITY TEST (Definition 5.1) after the ceiling test and before any prescription. Report `u*` beside `u_bar` and declare the judgment.

SPEC 5.5 — REPORT THE CEILING MARGIN IN NATS AND OCTAVES, and the roster-ceiling margin from Theorem 4.3. A case inside both by a small margin is a case one member from being structural.

SPEC 5.6 — SEPARATE TUNING FROM RESTRUCTURING in every prescription. Changes to `eta`, `c`, `delta` and `N` are tuning; changes to membership, topology or boundary are restructuring, and they are reported as producing a different group.

SPEC 5.7 — PRICE GOVERNANCE MECHANISMS AGAINST THE SURPLUS THEY ARE LEVIED ON. P5 at −7.7% of `eta S` costs 0.004877 on this commons and takes `Q_inf` from `-8.462972` to `-8.517157`.

SPEC 5.8 — OSTROM'S PRINCIPLES ACT ON `c`, `eta` AND `delta`, NEVER ON `u`. Any claim that a design principle raises non-parallelism must be argued separately and is not supported by the mapping.

SPEC 5.9 — NESTING IS REPORTED WITH ITS LOST PAIRS (SPEC 4.13). Both the link saving and the surplus loss are computed.

SPEC 5.10 — DELTA IS NOT A LEVER ON THE DEFICIT (Theorem 4.9), and where the numerator is negative it is not a lever on the destination either, except downward (the ordering rule [Curzi, 2026]).

SPEC 5.11 — REPORT THE SENSITIVITY to the least certain parameter and declare the verdict UNRESOLVED where the uncertainty exceeds the deficit. **The constructed family is unresolved at `eta ± 0.05`.**

SPEC 5.12 — NO ORGANISATION-WIDE AGGREGATE (the no-aggregation rule [Curzi, 2026]; CC-R16). Per group, per period.

SPEC 5.13 — LABEL EVERY CASE CONSTRUCTED, and label an extension of the carried case as an extension.

SPEC 5.14 — NO CLINICAL, DIAGNOSTIC, THERAPEUTIC OR COUNSELLING FRAMING, and no reading presented as an evaluation of any person or of any group's health.

The failure modes

FM 5.1 — TUNING AN UNREACHABLE CASE. *Signature:* a lever programme for a group whose required `u*` is far outside what its members could represent. *Mechanism:* Definition 5.1. The commons is TUNABLE by the ceiling test and requires `u* = 0.938078` from `u_bar = 0.018614` — a factor of 50.4. *Fix:* SPEC 5.4. An unsatisfiable prescription is worse than none: it consumes the effort that restructuring needed (the minimum-frequency paper [Curzi, 2026]).

FM 5.2 — SINGLE-LEVER THINKING ON A SIX-OCTAVE GAP. *Signature:* a plan to fix the commons by improving conversion, or by cutting cost, or by adopting a richer taxonomy. *Mechanism:* every lever saturates. `eta = 1.00` leaves `+1.8430`; a 91% cost cut leaves `+0.0990`; `N = 144` leaves `+3.2562`. *Fix:* §5.4's four tables, computed before the plan rather than after it.

FM 5.3 — MISTAKING SELECTION FOR CULTURE. *Signature:* a group with `u_bar` near zero addressed with meetings, workshops, or facilitation. *Mechanism:* the Parallel Wall. `u` is a property of who the members are. Ostrom's principles all act on `c`, `eta` and `delta`, and not one acts on `u` (SPEC 5.8). *Fix:* SPEC 5.6. Name it as a membership question.

FM 5.4 — TRUST APPLIED TO A NEGATIVE NUMERATOR, IN A FAMILY. *Signature:* continuity, ritual or patience invested in a marginal family case, followed by acceleration. *Mechanism:* CC-R15 and the sign-multiplication result [Curzi, 2026]. `delta` 0.05 -> 0.03 takes the constructed family from `−1.378049` to `−2.296749`; the deficit does not move. *Fix:* the ordering rule [Curzi, 2026]. Sign the numerator first; here a 9.7% rise in `eta` flips it.

FM 5.5 — ACTING ON THE FAMILY'S UNRESOLVED VERDICT. *Signature:* an intervention justified by `D = +0.5847` with `eta` known to a decimal place it does not have. *Mechanism:* SPEC 4.19. At `eta ± 0.05` the uncertainty is `±0.9734` octaves and the sign is unknown. *Fix:* SPEC 5.11. Measure `eta` to better than ±0.030034 first.

FM 5.6 — THE AGGREGATION GAP SUPPRESSED IN A CLUSTER GROUP. *Signature:* one figure reported for a group with one tight pair and several wide ones. *Mechanism:* Jensen. The constructed family's gap is 0.394666 octaves; the two-dyad structure of §4.13b runs 0.599632. *Fix:* SPEC 5.3.

FM 5.7 — THE SEVENTH MEMBER. *Signature:* a group at its roster ceiling admitting one more, with no reading taken. *Mechanism:* Theorem 4.3. The commons sits at `m(m−1) = 30` against 30.7225; a seventh takes it to 42 and the case becomes STRUCTURAL, at which point no `u` whatsoever suffices. *Fix:* SPEC 5.5, and run the roster ceiling before admission.

FM 5.8 — GOVERNANCE PRICED WITHOUT REFERENCE TO SURPLUS. *Signature:* monitoring, sanctions and conflict procedures added to a group with `F = 0.186272`. *Mechanism:* mechanisms costed as a fraction of `eta S` are affordable in proportion to `S`. *Computed:* P5 takes this commons from `−8.462972` to `−8.517157`. *Fix:* SPEC 5.7.

FM 5.9 — THE DECLINE READ AS A RESULT. *Signature:* "the numbers got better when `E` left." *Mechanism:* a reading after a withdrawal is a reading of a different group. The apparatus has no single-party term and produces no statement about `E`. *Fix:* SPEC 5.2 and §5.14's three rules.

FM 5.10 — THE COUNTERFACTUAL SHOWN BEFORE THE DECISION. *Signature:* a member weighing withdrawal is told what their withdrawal would do to the reading. *Mechanism:* this converts arithmetic into pressure and removes the ability to decline freely, which is the one thing the consent axiom of the framework's foundations [Curzi, 2026] forbids absolutely. *Fix:* SPEC 5.2. **The counterfactual is computed after the decision, or not at all.**

FM 5.11 — RESTRUCTURING PRESENTED AS TUNING. *Signature:* "we just need to bring in some fresh perspectives" offered as an adjustment. *Mechanism:* changing membership changes `u`, which is the one parameter tuning cannot reach. It is a different group afterwards and the honest framing says so. *Fix:* SPEC 5.6.

FM 5.12 — THE STRUCTURE NAMED INSTEAD OF DECLARED. *Signature:* "we are a family, so `delta` is low." *Mechanism:* the labels in §5.1 are parameter signatures, not primitives. A group is whatever its five declarations say it is. *Fix:* SPEC 5.1.

Worked example

THE THREE STRUCTURES, READ IN FULL All three cases are CONSTRUCTED. The family and the team are explicitly labelled extensions of the carried dyad of the carried dyad [Curzi, 2026]; the commons is the carried commons [Curzi, 2026] exactly. Parameters are chosen, not observed. None of this is an assessment of any person, family, team or community; no clinical, diagnostic or therapeutic claim is made; the frequencies are essence frequencies.

THE FAMILY

```
A = {4: 0.7,  6: 0.5,  9: 0.4}                                  L2 norm 0.948683
B = {4: 0.6,  6: 0.55, 9: 0.3,  10: 0.2}                        L2 norm 0.890225
D = {1: 0.50, 2: 0.45, 4: 0.35, 6: 0.30, 8: 0.40, 9: 0.20}      L2 norm 0.930054
E = {3: 0.20, 4: 0.30, 6: 0.25, 9: 0.35, 11: 0.55, 12: 0.50}    L2 norm 0.931397

m = 4, N = 12, one declared period
kappa = 0.130 per link, 6 links, C = 0.7800, eta = 0.34, delta = 0.05

pairwise u    A-B 0.068735   A-D 0.710180   A-E 0.711015
              B-D 0.723966   B-E 0.740352   D-E 0.916710
u_bar         0.645160          F(u_bar) 2.091463    mean-of-F 1.954682
gap           0.394666 octaves
N_eff         8.096755          f = 31.6922 Hz    folded 253.537 (j = 3)
c/eta         2.294118  vs ln 12 = 2.484907   margin +0.190789   TUNABLE
roster ceil   m(m-1) = 12 <= 12.9980                               PASS, at the ceiling
Brooks        r = 0.382353, max viable roster 2, actual 4         OVER
f_min         47.5308 Hz        folded 380.246 (j = 3)
req/present   6.619424 / 6.034688
DEFICIT       +0.584737   (+701.7 cents)     conservative +0.979403
numerator     -0.068902   NEGATIVE
Q_inf         -1.378049                      conservative -2.308160
headroom      1.1352 octaves      mu = 1.000
VERDICT       DECAYS on both aggregations -- determined
SENSITIVITY   D = +0.5847 +/- 0.9734 at eta +/- 0.05  -> SIGN NOT RESOLVED
chord         states {1,2,3,4,6,8,9,10,11,12}, classes {1,3,5,9,11} -- 5 of 6
levers        eta 0.34 -> 0.372945  (+9.69%)  closes it
              C   0.78 -> 0.711098  (-8.83%)  closes it
              u   0.645160 -> 0.810517        closes it
              N   12 -> 15                    closes it
after eta -> 0.40:  numerator +0.056585, D = -0.408177, Q_inf(0.05) = +1.131707
                    then delta 0.05 -> 0.03:  Q_inf = +1.886178
```

The family is the tunable case, its cheapest lever costs a tenth of its conversion efficiency, and its verdict is not resolved at the precision its parameters are known to. Both of those are ordinary and both are reportable.

THE TEAM

```
A  = {4: 0.7,  6: 0.5,  9: 0.4}                                    L2 norm 0.948683
B' = {2: 0.95, 4: 0.60, 6: 0.55, 9: 0.30, 10: 0.20, 12: 0.8391}    L2 norm 1.548899
C  = {3: 0.45, 4: 0.55, 6: 0.40, 7: 0.35, 11: 0.25}                L2 norm 0.921954
S  = {3: 0.25, 4: 0.30, 5: 0.60, 7: 0.45, 9: 0.20, 10: 0.40}       L2 norm 0.956556
T  = {2: 0.30, 4: 0.35, 6: 0.25, 8: 0.55, 9: 0.35, 11: 0.45}       L2 norm 0.950000

m = 5, N = 12, one declared period
kappa = 0.055 per link, 10 links, C = 0.5500, eta = 0.52, delta = 0.09
```

```
pairwise u    A-B' 0.692372  A-C 0.552647  A-S 0.897875  A-T 0.679778
              B'-C 0.851659  B'-S 0.953352 B'-T 0.748794
              C-S  0.756702  C-T 0.786182  S-T 0.962914
u_bar         0.788228          F(u_bar) 2.269080    mean-of-F 2.258759
gap           0.029780 octaves
N_eff         9.670503          f = 45.2093 Hz    folded 361.674 (j = 3)
c/eta         1.057692  vs ln 12 = 2.484907   margin +1.427214   comfortable
roster ceil   m(m-1) = 20 <= 46.9873                               PASS
Brooks        r = 0.105769, max viable roster 6, actual 5         WITHIN
f_min         4.0089 Hz         folded 256.572 (j = 6)
req/present   3.051855 / 6.547182
DEFICIT       -3.495327   (-4194.4 cents)     conservative -3.465547
numerator     +0.629922   POSITIVE
Q_inf         +6.999131                       conservative +6.939499
headroom      0.6227 octaves      mu = 1.667
VERDICT       ACCUMULATES on both aggregations -- determined
chord         states {2,...,12}, classes {1,3,5,7,9,11} -- 6 of 6, saturated
```

The team is sustainable and has 0.6227 octaves of the scale left. Its next increment is not buyable from `u` at twelve categories, and the category-validity condition [Curzi, 2026] governs what a real addition to `N` would have to be. **`delta` interventions are available and would help here**, because the numerator is positive — the ordering rule [Curzi, 2026] permits them at this point and not before.

THE COMMONS **CONSTRUCTED, the carried commons [Curzi, 2026].** A realising set is given so the figures are reproducible. All six members hold the same doctrine profile across all twelve states, each raising exactly one state by 0.130330:

```
doctrine profile (states 1..12):
  {1:0.10, 2:0.20, 3:0.55, 4:0.35, 5:0.60, 6:0.25,
   7:0.15, 8:0.30, 9:0.70, 10:0.20, 11:0.10, 12:0.45}

M raises state  1 to 0.230330      L2 norm 1.339982
N raises state  2 to 0.330330      L2 norm 1.349673
O raises state  6 to 0.380330      L2 norm 1.354493
P raises state  7 to 0.280330      L2 norm 1.344836
Q raises state 10 to 0.330330      L2 norm 1.349673
R raises state 11 to 0.230330      L2 norm 1.339982

resulting mean pairwise u over 15 pairs = 0.018614137     (the carried figure is 0.018614)
pairwise u range 0.018484 .. 0.018830
cos theta range  0.990540 .. 0.990715
```

Every pair is above 0.99 alignment. That is what a membership rule looks like from inside the geometry.

```
u_bar         0.018614          F 0.186272    mean-of-F 0.186276
gap           0.000001 octaves            -- negligible, because the pairs are alike
N_eff         1.204754          f = 0.7017 Hz    folded 359.270 (j = 9)
c/eta         2.426471  vs ln 12 = 2.484907   margin +0.058436   TUNABLE
reachability  u* = 0.938078 from u_bar = 0.018614  -- factor of 50.4   UNREACHABLE
roster ceil   m(m-1) = 30 <= 30.7225                               PASS, at the ceiling
Brooks        r = 0.161765, max viable roster 4, actual 6         OVER
f_min         61.9349 Hz        folded 247.740 (j = 2)
req/present   7.001300 / 0.537500
```

```
DEFICIT       +6.463800   (+7756.6 cents)
numerator     −0.761666   NEGATIVE
Q_inf         −8.463000
headroom      6.6324 octaves      mu = 2.500
chord         all twelve states held by every member; 6 classes of 6
VERDICT       DECAYS -- and TUNABLE, UNREACHABLE
```

Note the chord line. Every member holds every state, so the commons' chord spans all twelve states and all six pitch classes — the fullest chord in this framework — and its surplus is 0.186272 nats. A group can sound every available class and carry almost no information, because the chord reports what states are held and `u` reports whether anyone holds them differently. The dissonance paper's [Curzi, 2026] identity is the reason: consonance and redundancy are the same phenomenon, and a chord this consonant is a single note played six times.

THE LEVER AUDIT, ALL FOUR, SIDE BY SIDE

```
------------------------------------------------------------------------------
THE COMMONS -- deficit +6.4638 octaves -- every lever alone
------------------------------------------------------------------------------
L1  u -> 0.95        deficit −0.0332      CLOSES, by 40 cents, and 0.95 is
                                          near-orthogonality for six people
                                          drawn from one doctrine
    u -> 1.00        deficit −0.1686      the absolute maximum, 0.1686 octaves
    hold line u*  =  0.938078             a factor of 50.4 above u_bar
L2  eta -> 0.90      deficit +2.1075      FAILS
    eta -> 1.00      deficit +1.8430      FAILS at perfect conversion
L3  kappa -> 0.020   deficit +2.0085      FAILS
    kappa -> 0.005   deficit +0.0990      FAILS at a 91% cut
L4  N -> 25          deficit +5.9357      FAILS
    N -> 144         deficit +3.2562      FAILS at twelve times the ladder
------------------------------------------------------------------------------
P5 (Ostrom, −7.7% of eta S)    Q_inf −8.462972 -> −8.517157      slightly worse
P8 nesting, 15 links -> 7      deficit +6.4638 -> +2.7298        still failing
P8 nesting, 15 links -> 6      deficit +6.4638 -> +2.2630        still failing
------------------------------------------------------------------------------
RESTRUCTURED: recruit outside the doctrine (u_bar 0.45), nest to 7 links,
adopt a conversion mechanism (eta 0.55)
    f = 17.1145 Hz   f_min = 1.9604 Hz   DEFICIT −3.126006   Q_inf +6.620724
    swing 9.5898 octaves: 51.9% cost side, 48.1% geometry
    AND IT IS A DIFFERENT GROUP
------------------------------------------------------------------------------
ESSENCE frequencies, not physical resonances. Constructed case. Not an
assessment of any person or community. Not counselling; no clinical claim.
------------------------------------------------------------------------------
```

What the three readings establish **One:** the parameter that separated the working structure from the failing one was `u_bar`, and it is the only one of the five declarations that an arrangement cannot change.

Two: a case can pass the ceiling test and still be unfixable, which is why Definition 5.1's reachability test belongs after the ceiling test and before any prescription.

Three: every lever available to the commons was computed and every one of them fell short — including `eta = 1`, including a 91% cost cut, including a twelve-fold enlargement of the category set. That is not a rhetorical claim about difficulty; it is four tables of arithmetic, and it is the result this section was written to publish.

Four: the restructure that works changes who is in the group, and the formalism has nothing whatever to say about whether that is the right thing to do. It says only what the current arrangement arrives at if it holds, and Section 6 takes up what may and may not be done with that information.

Section 6 is the intervention section: given a deficit, what is actually done, lever by lever, with consent and reversibility built into the numbered procedure rather than appended to it — and a section on when not to tune at all.

6 The intervention procedure and its halting conditions

The preceding two sections have produced a number. This one is about what may be done with it.

The number is an octave deficit, and by the deficit identity R3 [Curzi, 2026] it is also a destination: `D` and `Q_inf` are the same quantity in different units. It says that a configuration of four parameters, if it holds, arrives somewhere — and it says how far that somewhere is from sustaining. It does not say what should be done, to whom, by whom, or whether anything should be done at all, and the distance between a reading and an action is where every one of this section's constraints lives.

There are exactly four levers, plus one parameter that is not a lever on the deficit at all. That is not a stylistic simplification; it is the decay-blindness result CC-R4 [Curzi, 2026] and it exhausts the possibilities. So a section on intervention is short on options and long on discipline, and this one is organised accordingly: a numbered procedure with the constraints built into the steps, a lever-by-lever account of what each move actually requires and costs, and a section on the cases in which the correct action is to do nothing.

Three things are established before anything else, because they govern everything after.

This section is not counselling, therapy, mediation, or clinical assessment, and nothing in it may be presented as any of those. It is a procedure for changing four declared parameters. It has no diagnostic content, the twelve states are not diagnoses, no reading here evaluates any person's health or capacity, and a reader in difficulty with a relationship should seek help from someone qualified to give it, which this framework is not and does not claim to be.

Consent is structure, not advice. the consent axiom of the framework's foundations [Curzi, 2026] states it absolutely: no intervention may remove a person's ability to decline. In this section that constraint appears as numbered steps in the procedure, as a halting condition, and as a set of prohibitions on what may be computed and when — not as a paragraph of encouragement at the end.

And the procedure halts more often than it proceeds. Of the three structures Section 5 read, one required no intervention, one required a measurement rather than an intervention, and one returned a verdict that no intervention could satisfy. That distribution is not pessimism. It is what a diagnostic is for: most of the value of a reading is in the cases where it tells you not to act.

The law

6.1 What an intervention is DEFINITION 6.1. An intervention is a change to one or more of the four parameters that appear in the octave deficit, optionally accompanied by a change to the one that does not.

```
D  =  2.885390 * ( c/eta )  −  log2( f/k )        f = k (1 + (N−1) u)^2
```

```
L1  raise u        the geometry        who the parties are, and what they hold
L2  raise eta      the conversion      what happens to what they produce
L3  cut c          the cost            what the arrangement costs to run
L4  raise N        the categories      what the description can distinguish

L0  lower delta    the loss            trust, ritual, record, continuity
                   NOT a lever on D (CCP-041 R4); a multiplier on Q_inf (CCP-044)
```

There is no fifth. Anything proposed as an intervention either moves one of these five or moves nothing, and the first job of the procedure is to say which one a proposal moves. A proposal that cannot be assigned to a lever is not an intervention; it is a wish, and the minimum-frequency paper's [Curzi, 2026] belief postulate applies — belief acts only through a coupling with real gain, and where no channel exists nothing happens.

6.2 The consent axiom

AXIOM 6.1 — THE CONSENT AXIOM. No procedure in this framework may reduce any party's ability to decline participation in it, at any stage, for any reason or none.

This is a constraint on the procedure, not a recommendation about how to behave, and it has three corollaries that determine specific numbered steps.

COROLLARY 6.1a — INFORMED. A party who has not been told what is being measured, what will be computed, and what will be reported cannot decline meaningfully, because they cannot know what they would be declining. The procedure therefore discloses its outputs before it takes its inputs. This also discharges the no-subliminals rule (the consent axiom of the framework's foundations [Curzi, 2026]): nothing masked, nothing below threshold, nothing concealed. Open-label works; disclosure costs nothing that matters.

COROLLARY 6.1b — REVOCABLE. A consent that cannot be withdrawn is a decision made once about every future stage. Consent is therefore re-asked at each stage that changes what is being done, and withdrawal at any stage halts that stage without penalty and without explanation.

COROLLARY 6.1c — UNCONDITIONED. A party shown the consequence of their own withdrawal before deciding is choosing under pressure, and pressure is a reduction in the ability to decline. §5.14 computed the case that makes this concrete: a member's withdrawal improved a constructed family's reading by 2.680588 octaves, and every member of a constructed team improved the reading by leaving. Those figures may not be shown to a person weighing withdrawal. The counterfactual is computed after a decision, or not at all.

A note on why 18.1c has to be an axiom rather than a courtesy. §5.11b established that a drop-one row improves for *every* member of a healthy group near the surplus ceiling, so the row carries no information about the person it names. A number that says nothing about someone, shown to them as though it did, is the worst possible thing to put in front of a person deciding whether to stay.

6.3 Reversibility, defined and classified

DEFINITION 6.2. An intervention is REVERSIBLE if ceasing it returns the parameter to its prior value without residue in any other parameter and without requiring anything of any party. It is PARTIALLY REVERSIBLE if ceasing it returns the parameter but leaves residue elsewhere. It is IRREVERSIBLE otherwise.

```
LEVER                          CLASS                residue on ceasing
-------------------------------------------------------------------------------
L2  raise eta (mechanism)      REVERSIBLE           none; stop keeping the record
L3  cut c     (arrangement)    REVERSIBLE           none; restore the meeting
L4  raise N   (description)    REVERSIBLE           none; the ladder is a choice
L0  lower delta (record)       REVERSIBLE           none, but see the decay note
```

```
L1  raise u   (what a party    IRREVERSIBLE          a person now holds commitments
               actually holds)                      they did not hold before
-------------------------------------------------------------------------------
RESTRUCTURING (membership)     IRREVERSIBLE          a different group exists
```

The class list is the reason the procedure is ordered as it is. Three of the four levers, and the one non-lever, are reversible; the geometry lever is not. `u` rises when a party takes on states they did not previously hold, and a person cannot un-take a commitment on request. An intervention that changes what somebody is actually doing is not undone by deciding to stop measuring it.

And reversible does not mean instantaneous. `D` returns to its old value the moment the parameter does, because `D` has no memory. `Q` does not: it decays back at rate `delta`, with half-life `ln(0.5)/ln(1-delta)`. Computed on the constructed family, `eta` raised to 0.40 for twenty periods and then reverted:

```
t = 0    Q = +1.000000        eta = 0.40, Q_inf = +1.131707
t = 20   Q = +1.084492        eta reverted to 0.34, Q_inf = -1.378049
t = 25   Q = +0.527418
t = 30   Q = +0.096365
t = 40   Q = -0.495263
t = 60   Q = -1.061583
half-life at delta = 0.05:  13.5134 periods
```

A reversible intervention leaves a stock behind that takes about thirteen periods to half-decay, which is a grace period and not a permanence. Nothing is destroyed by stopping; nothing is kept either.

6.4 What each lever requires, in closed form

Given a reading, the exact requirement for each lever to close the deficit on its own:

```
L1   u_req   =  ( e^(c/eta) - 1 ) / (N - 1)        reachable iff c/eta <= ln N
L2   eta_req =  c / F(u)                           reachable iff <= eta_max
L3   c_req   =  eta * F(u)                         always >= 0, so always defined
L4   N_req   =  1 + ( e^(c/eta) - 1 ) / u          integer, and CCP-041's conditions apply

partial close of x octaves:  d(c/eta) = -x/2.885390 nats,  or  dF = +x/2.885390 nats
one octave:  dF = 0.346574 nats  =  N_eff multiplied by sqrt(2) = 1.414214
```

Computed on the three carried cases:

```
case          u_req               eta_req              c_req       N_req
dyad 2a       0.627132            1.101077  ABSENT     0.168925    102
family        0.810517            0.372945             0.711098     15
commons       0.938078  UNREACHABLE 4.428926 ABSENT     0.063334    556
```

Three of the twelve entries are impossible before any question of expense arises. `eta_req > 1` means a conversion efficiency above unity, which is more structure created than interaction supplied; `N_req = 556` on a twelve-category ladder is not a taxonomy revision. The closed forms are worth computing first precisely because they identify levers that are absent rather than costly, and no cost model is needed to see an absent lever.

6.5 The ordering rule, as a step constraint The sustainable-dyad paper [Curzi, 2026] owns the order-of-operations rule and this section applies it as a gate.

```
SIGN THE NUMERATOR:   n = eta F(u) − c

n < 0    NO intervention on delta is undertaken. Raise eta, cut c, or raise u first.
n > 0    delta interventions are available and are the strongest single move.
n = 0    delta is exactly inert; Q_inf = 0 for every delta.
```

The rule is enforced at a numbered step because it cannot be seen in the deficit. `D` is identical before and after any `delta` change, to every digit, so an analyst working from the deficit alone has no signal that a `delta` intervention did anything — and in a negative-numerator case it will have doubled the rate of loss while the instrument read flat.

The damage-integral result R7 [Curzi, 2026] prices the consequence: eight periods of the wrong order cost twelve periods of permanent lag and 111.133601 stock-periods, on a case whose fixed point is order-independent.

6.6 The classification, complete Everything the procedure does is decided by four cases, assembled from the minimum-frequency paper [Curzi, 2026] and Definition 5.1.

```
D <= 0                                    SUSTAINABLE.
                                          Report the margin. Do not tune. Sec 6.13.

D > 0, c/eta <= ln N, u_req reachable     TUNABLE AND REACHABLE.
                                          Run the procedure.

D > 0, c/eta <= ln N, u_req unreachable   TUNABLE, UNREACHABLE.
                                          Report; issue no lever prescription; the question
                                          is membership, not arrangement.

c/eta > ln N                              STRUCTURAL.
                                          L1 is unavailable at every value of u. Only L2,
                                          L3, L4 —— and L4 only to make the case tunable.
```

A prescription issued before this classification may be unsatisfiable, and an unsatisfiable prescription is worse than none, because it consumes the effort that restructuring needed (the minimum-frequency paper [Curzi, 2026]). The classification costs one logarithm and one division.

6.7 The standing prohibitions, restated where they bind

(1) NOT COUNSELLING, THERAPY, MEDIATION OR CLINICAL ASSESSMENT. Restated at the top of this section and again here because this is the section someone will try to use as a protocol. **There is no diagnostic content in any output of this procedure, and no output may be described as an assessment of a person.**

(2) ESSENCE, NOT RESONANCE. Every frequency is a position on a logarithmic scale of surplus. Nothing vibrates and nothing is played at anyone. The minimum-frequency paper's [Curzi, 2026] FM 4.8 forbids sounding `f_min` at any person, and the no-subliminals rule is absolute: nothing masked, nothing below threshold, nothing concealed.

(3) NO SINGLE-PARTY TERM. `u` is a joint geometric property; `eta`, `c` and `delta` are properties of an arrangement. No output of this procedure names a person as a cause, and none may be reported as a ranking or an evaluation.

(4) THE READING INHERITS ITS INPUTS. Where intensities were not magnitude-estimated (the elicitation rule of [Curzi, 2026] rule; the elicitation rule), no relationship frequency may be printed. Where `eta` or `c` is ordinal, publish the ordering and not the number.

(5) THE READING DOES NOT CONCLUDE THAT A RELATION SHOULD END. `D > 0` says a configuration loses order at a computable rate. It does not weigh what the relation is for, what else it produces, or what ending it would cost, and it acquires no such term by being applied to people (the minimum-frequency paper [Curzi, 2026] §4.22).

(6) CONSTRUCTED CASES. Every case here is constructed and labelled, and its precision does not transfer to a measured one.

The engineering

6.8 PROCEDURE T — the numbered procedure The constraints are steps. They are not appended and they are not optional, and the procedure halts at four points by design.

STEP 1 — DISCLOSE BEFORE ELICITING. Every party is told, before any measurement: what will be measured, how, what will be computed, what will be reported, to whom, and what will be done with it. Nothing in the procedure is concealed from a participant at any stage (Corollary 6.1a).

STEP 2 — OBTAIN CONSENT, SEPARATELY AND EXPLICITLY. From every party, individually, without conditioning. The procedure does not start without every party's agreement, and a party who declines is not asked for a reason, is not recorded as having declined for cause, and the reading is simply not taken.

> **HALT 1.** Any party declines. The procedure ends. No reading is computed, no partial reading is computed, and no counterfactual of the form "what the reading would have been" is computed (Corollary 6.1c).

STEP 3 — DECLARE THE FRAME. `N` and the ordering that produced it (the state-ordering material [Curzi, 2026]); the period; the scale status of every input (SPEC 4.4). One clock for `eta`, `c` and `delta` (SPEC 4.3).

STEP 4 — ELICIT AND ENCODE. Intensities by magnitude estimation (the elicitation rule of [Curzi, 2026] rule), raw L2 norms published. Encode each party as a unit vector. Check pairwise support overlap and reject disjoint-support encodings (the self-as-chord material [Curzi, 2026]).

STEP 5 — MEASURE AND COMPUTE. `u` per pair (the relationship-frequency paper [Curzi, 2026]); the full pairwise list published, not the mean alone; `u_bar`, `F(u_bar)` and mean-of-`F` with the aggregation gap; `f`, folded, with `j`.

STEP 6 — CLASSIFY, BEFORE ANY PRESCRIPTION. Roster ceiling (Theorem 4.3), then `c/eta` against `ln N`, then reachability (Definition 5.1). Report the classification of §6.6.

> **HALT 2.** The case is STRUCTURAL, or TUNABLE-UNREACHABLE. Report the finding and the four lever computations that establish it, and issue no tuning prescription. The question is one of arrangement scope or membership, and §5.7 governs the distinction.

STEP 7 — REPORT THE SENSITIVITY, AND TEST WHETHER THE VERDICT EXISTS. Compute `dD` from the uncertainty in the least certain parameter (SPEC 4.19).

> **HALT 3.** The uncertainty exceeds `|D|`. **The verdict is NOT RESOLVED and the correct action is measurement, not intervention.** Report what precision would resolve it: on the constructed family, `eta` must be known to better than ±0.030034 — 8.8% of its value — for the sign of a `+0.584737` deficit to be determined.

STEP 8 — SIGN THE NUMERATOR. `n = eta F(u) − c`. Record the sign. This is the gate for Step 12.

STEP 9 — COMPUTE THE CLOSED FORMS AND THE RESERVES. `u_req`, `eta_req`, `c_req`, `N_req` (§6.4), and the reserve of each lever (the lever-reserve result R6 [Curzi, 2026]). Mark any lever ABSENT rather than expensive.

STEP 10 — CLASSIFY EACH CANDIDATE MOVE BY REVERSIBILITY (Definition 6.2), and present the reversible options first. Where a reversible move suffices, an irreversible one is not proposed.

STEP 11 — SEEK SEPARATE, EXPLICIT AGREEMENT FOR EACH MOVE, and additional explicit agreement for any move classified irreversible. Agreement to the reading is not agreement to an intervention, and agreement to one intervention is not agreement to the next.

> HALT 4. No move commands agreement from every party it requires something of. The procedure ends with the reading and no intervention. A reading that produced no agreed action is a completed procedure, not a failed one.

STEP 12 — SEQUENCE. Numerator moves first (`eta`, `c`, `u`), then `delta` moves, and only if Step 8's sign is positive by the time they are reached (the ordering rule [Curzi, 2026]).

STEP 13 — ACT, AND RECORD THE BASELINE. Record every parameter's prior value, so that reversal is a defined operation rather than a reconstruction.

STEP 14 — RE-READ, AND RE-CONSENT. At a declared interval. The re-reading restarts at Step 1. Consent obtained for one reading is not consent for the next, and §6.12 gives the cadence and the design controls that make a re-read mean anything.

6.9 The levers, one at a time, with what each actually costs

L2 — RAISE `eta`. *Reversible.*

What it is. Change what happens to what the parties already produce. The decision that stays decided; the practice that has a slot; the artefact that exists next week because of what happened this week. It is a mechanism question, and often literally a notebook.

What it requires. `eta_req = c/F(u)`, or a partial close of `x` octaves at `d(c/eta) = −x/2.885390` nats.

Its price in octaves. `dD/d(ln eta) = −2.885390 (c/eta)` — an e-fold improvement buys exactly the number of octaves the relation requires. On the constructed family that is 6.619424 octaves per e-fold against a deficit of 0.584737, which is why a 9.7% rise suffices.

Its reserve. `2.885390 c (1/eta_0 − 1/eta_max)`, and it is small wherever `eta` is already moderate: at `eta = 0.50`, `c = 0.62`, the whole remaining reserve is 1.7889 octaves.

Its failure mode. Cutting the mechanism while cutting cost. If the recurring meeting was the conversion, removing it moves `c` and `eta` together and the ratio may not move at all (SPEC 4.11).

L3 — CUT `c`. *Reversible.*

What it is. Reduce what the arrangement costs to run before it produces anything: coordination, re-explaining, friction, correction load.

What it requires. `c_req = eta F(u)`.

Its price in octaves. `dD/d(ln c) = +2.885390 (c/eta)` — identical in magnitude to `L2`, opposite in sign, because the two enter only as a ratio. The formalism is indifferent between them, so whichever is cheaper to move in the situation is strictly the better choice.

Its reserve. `2.885390 (c/eta)`, i.e. the entire requirement — `c` can in principle go to zero, at which point `f_min = k` and any non-zero `u` clears it.

A measured instance. The corpus's built-environment work moves exactly this lever: an entrance sequence redesign takes interruptions from 22 to 6 per person-day, a switching cost from 0.321 to 0.087, headroom from $+0.0642$ to $+0.2955$ and the steady state from 5.185 to 9.041, for 7.6 per cent of net floor area [Curzi, 2026f].

Its failure mode. CC-R18's counter-pressure. An over-framed system banks more stored order than a calibrated one, and the bill arrives later as correction load `C_t = c_0 + kappa Z_t`. A `c` that was cut by ceasing to check things will rise again for reasons nobody connects to the choice.

L1 — RAISE `u`. *IRREVERSIBLE.*

What it is. A party comes to hold states the other parties do not hold. Not a rhetorical position — a change in what somebody is actually doing.

What it requires. `u_req = (e^(c/eta) − 1)/(N−1)`, reachable only if `c/eta <= ln N`, and reachable *in practice* only if the parties could honestly represent it (Definition 5.1).

Its price in octaves. `dD/d(ln u) = −2.885390 (1 − 1/N_eff)`, which is bounded above by 2.885390 octaves per e-fold in every system that has ever existed or will, because `1 − 1/N_eff < 1`. The cost-side levers exceed it whenever `c/eta > 1 − 1/N_eff`, which the elasticity comparison R5 [Curzi, 2026] states and which holds in all three of this framework's carried cases.

Its reserve. The pitch headroom `2.885390 (ln N − F(u))`, and it is the only large reserve available on a big deficit — 5.545206 octaves on the carried dyad against `eta`'s 4.174198.

Its failure modes, and there are three. CC-R8: manufactured opposition reads as `u = 0` exactly as agreement does, so a contrarian raises cost and returns nothing. Re-encoding: a `u` that rose because the description changed is not a `u` that rose. And convergence drift: `u` decays as parties spend time together, at a rate estimable from two readings (the convergence-clock material [Curzi, 2026]).

And its consent weight. This is the only lever that requires a person to become different, and it is the only irreversible one. Step 11's separate agreement exists for it.

L4 — RAISE `N`. *Reversible.*

What it is. Add genuinely distinct categories to the description — carrying variance the existing ones do not, and satisfying the state-ordering material's [Curzi, 2026] validity conditions.

What it requires. `N_req = 1 + (e^(c/eta) − 1)/u`.

Its price in octaves. `dD/d(ln N)`-scale: `2.885390 u/(1 + (N−1)u)` per category — 0.1129 at the carried dyad's `u`, 0.0446 at the commons'.

Its real value. Not the deficit. `N` raises `ln N`, which raises the maximum satisfiable `c/eta`. `N` is how a STRUCTURAL case becomes a TUNABLE one, and that is the only situation in which it is the right first move.

Its failure modes. The Parallel Wall — `N` enters `F` only multiplied by `u`, so categories added to a near-parallel set are dimensions nobody uses. The `N = 17` cliff, where the folded lattice's minimum separation falls from 34.9756 cents to 6.0008 and the symbol rate drops by exactly 8 (the mood-and-trait paper [Curzi, 2026]). And the honest one above both: a category added for arithmetic convenience is not a category.

L0 — LOWER `delta`. *Reversible, and NOT a lever on the deficit.*

What it is. Trust, ritual, written record, institutional memory, continuity of place — anything that slows the leak of accumulated order.

What it does. Nothing to `D`, by Theorem 4.9. It scales `Q_inf` by `delta_old/delta_new`, in whichever direction the numerator's sign points.

Its price. At a positive numerator it is the largest single move available: on the carried dyad at 2b, `delta` 0.09 to 0.05 takes `Q_inf` from `+6.084427` to `+10.951968`, an 80.0% improvement with no new difference and no cost cut.

Its failure mode. CC-R15, the whole of it. At a negative numerator the identical change takes `Q_inf` from `-5.011940` to `-9.021491`. Step 12 is the gate and it is not discretionary.

A note worth making. Several real interventions move `eta`, `c` and `delta` together — a written record does all three — and those are the ones to reach for first, because half of what they do is on the numerator. Reported as one thing, though, they make trust look like it worked when what worked was the other half, which is why SPEC 15.4 requires `delta` effects to be reported separately.

6.10 Planning combinations

`D = 2.885390 (c/eta - F)` is a difference of two independent terms, so octave contributions add with no interaction term. A plan is therefore a list of moves whose contributions sum to at least the deficit.

```
total to close      D octaves
from eta:           2.885390 * c * (1/eta_0 - 1/eta_1)
from c:             2.885390 * (c_0 - c_1) / eta
from u:             2.885390 * ( F(u_1) - F(u_0) )
from N:             2.885390 * ( ln(1+(N_1-1)u) - ln(1+(N_0-1)u) )
```

Two cautions, both from the worked example of the minimum-frequency paper [Curzi, 2026]. The path-free split into cost side and surplus side is unique; the split *within* the cost side is not, because `c` is divided by `eta` and cutting cost is worth more when `eta` is still low. On the carried dyad the `c` contribution is `-1.442695` octaves on one attribution order and `-2.115953` on another — a 47% difference from the same intervention producing the same result. Report the path-free split, and if a per-parameter attribution is wanted, report the order used.

And single-lever plans on large deficits tend to land near the hold line, where a relation pays for itself exactly and banks nothing. The sustainable-dyad paper [Curzi, 2026] computes the carried dyad's: `u` alone to 0.692373 leaves `-0.250931` octaves, the cost side alone leaves `+0.998357` and still fails. A plan with no margin is a plan that fails on the first increase in cost.

6.11 Choosing between levers when the formalism will not

The four sensitivities are in four different units and cannot be ranked against each other (the minimum-frequency paper [Curzi, 2026]). The elasticities can, and the elasticity result R5 [Curzi, 2026] gives them. But neither answers what a unit of each parameter costs in the situation, and that cost model is not in the formalism and this framework does not pretend to supply it.

What the formalism does supply, and what is genuinely decisive without prices:

Absent levers. `eta_req = 1.101077` on the carried dyad and `4.428926` on the commons. No cost model is needed to rule out a conversion efficiency above 1.

Exhausted levers. `L4` at `u = 0.018614` buys 0.0446 octaves per category, so 6.4638 octaves is 145 categories.

Reversibility. Three levers are reversible and one is not, and where a reversible move suffices, Step 10 forbids proposing the irreversible one.

The shape of the trade. `dD/d(ln eta)` carries `c/eta`; `dD/d(ln u)` is capped at 2.885390. Both levers are strongest in the systems that most need them, which is the opposite of the usual experience that improvement gets easier once things are already going well.

6.12 Verifying that it worked

An intervention that is not verified is a story. Four requirements, and the last three are Book Two's, carried forward without softening.

Re-read at a declared interval, from Step 1. The deficit is a property of current parameters, so verification means re-measuring them, not remembering that something was done.

Do not use the shuffle as the verification. CC-R13 and CC-R24 measured three distinct failure modes: The shuffled-input control is anti-conservative as ordinarily run (2.2x nominal alpha), it cannot see repackaging (only the structural ceiling test can), and it cannot see

confounding by a shared slow driver, passing 98.35% of the time. **Design controls, not more surrogates.**

Use the structural test for repackaging. A change that exceeds the theoretical single-step maximum is not an improvement; it is a re-description. The single-step maximum [Curzi, 2026] and CC-R4 apply.

The methodological literature behind those three failures is not this framework's: selective analysis inflates false positives [Simmons et al., 2011, Gelman and Loken, 2013], most published findings in a low-prior field are false [Ioannidis, 2005], replication rates are low [Open Science Collaboration, 2015], and preregistration with a declared kill condition is the available remedy [Nosek et al., 2018]. The corpus's own confirmatory attempt failed its declared baseline test and reported the failure [Curzi, 2026i].

And CC-R23 stands. Testimonial evidence that a configuration held has a likelihood ratio of exactly `1.000000` — it does not move the posterior at all, because the selection guarantees you observe successes whether or not the method does anything. A hundred pairs reporting that their reading improved is worth precisely as much as none of them reporting it. What counts is the re-measured parameters, and the companion result says why: a prediction specifying a magnitude carries a likelihood ratio of 11.3577 against 2.0000 for one specifying only a direction — about 5.7 times more evidential power, which is the framework's own evidential pricing result [Curzi, 2026ac]. That is the correct response to the standing critique that directional predictions are nearly unfalsifiable in soft psychology [Meehl, 1990, 1967], and it is what a research programme is supposed to do with a badged postulate [Lakatos, 1970, Platt, 1964, Popper, 1959]. **Predict the number before you re-read.**

6.12b The re-reading window, and how it can be empty

Step 14 says re-read at a declared interval. The interval is not arbitrary, and it turns out to be bounded on both sides by quantities the reading already contains.

The lower bound is convergence. After a parameter change, `|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|`, so half of the move has happened after `t_half = ln(0.5)/ln(1−delta)` periods. Read sooner than that and most of what will happen has not happened; the stock is still dominated by where it started.

The upper bound is drift. the convergence-clock material [Curzi, 2026]'s convergence clock gives `t* = ln(u*/u_0)/ln(1−lambda)`, the period at which a sustainable geometry decays back across its own hold line. Read later than that and the reading is of a relation that has already lost the gain.

```
THE RE-READ WINDOW      [ ln(0.5)/ln(1−delta) ,  ln(u*/u_0)/ln(1−lambda) ]
```

Computed for three of this framework's cases:

case	u_0	u*	t_half	lambda	t*	window
family after tune	0.645160	0.548063	13.5134	0.01	16.229	[13.51, 16.23]
				0.02	8.074	EMPTY
				0.03	5.355	EMPTY
				0.05	3.180	EMPTY
carried dyad 2b	0.692373	0.134733	7.3496	0.02	81.020	[7.35, 81.02]
				0.05	31.911	[7.35, 31.91]
				0.08	19.631	[7.35, 19.63]
the team	0.788228	0.170883	7.3496	0.02	75.673	[7.35, 75.67]
				0.05	29.805	[7.35, 29.81]

The family's window is empty at any convergence rate above about 1% per period, and that is a finding about the tune rather than about the schedule. The family's repair left it at `−0.408177` octaves — a margin of 0.408 octaves — so its hold line sits at `u* = 0.548063` against a current `u_bar` of 0.645160, and the geometry has only 15% of relative room before it crosses

back. The repaired dyad has `u* = 0.134733` against 0.692373, five times the room, and its window is wide at every plausible drift rate.

The rule that follows is more useful than a schedule:

> **A margin thin enough to be unverifiable is a margin too thin. If the re-read window is empty, the correct response is not a different reading interval but a larger intervention.**

This is the quantitative form of §6.10's warning about single-lever plans landing near the hold line. A plan that closes the deficit and no more produces a relation that is sustainable, fragile, and — computably — impossible to confirm.

6.12c What a tuning report contains The report is the durable artefact, and writing it is itself an `eta` intervention: what is not written down is not converted. Fourteen items, and a report missing any of them is not auditable.

```
 1  the frame:   N, the ordering, the period, the scale status of every input
 2  consent:     obtained, from whom, at what stages, and any declines (as halts, not findings)
 3  the vectors: every party's states and intensities, plus every raw L2 norm
 4  the pairs:   the full pairwise u list, with cos theta signed
 5  aggregation: u_bar, F(u_bar), mean-of-F, and the gap in octaves
 6  the pitch:   f raw and folded, with j; f_min raw and folded, with j
 7  the tests:   roster ceiling, ceiling test with its margin in nats, reachability
 8  the reading: octaves required, octaves present, D with sign, in octaves and cents
 9  the destiny: Q_inf, and the identity cross-check D = -2.885390 (delta/eta) Q_inf
10  the sign:    the numerator eta F - c
11  sensitivity: dD from the least certain parameter, and whether the verdict is resolved
12  the levers:  u_req, eta_req, c_req, N_req, with absent levers marked ABSENT
13  the plan:    each move, its lever, its reversibility class, its octave contribution,
                 the path-free split, and the attribution order if one is given
14  the schedule: the re-read window with its two bounds, and the predicted D with a magnitude
```

Item 2 has a specific form and it matters. A decline is recorded as a halt in the procedure and never as a datum about the person who declined — no reason, no category, no counterfactual. A report that records why someone declined has converted a consent mechanism into a surveillance one, and Axiom 6.1 is violated by the record even if it was honoured in the room.

And item 14's prediction is recorded before the re-read, not after. CC-R23's companion result gives the reason in numbers: a prediction specifying a magnitude carries a likelihood ratio of 11.3577 against 2.0000 for a direction-only prediction, so a report that predicts *better* has recorded almost nothing, and one that predicts `-0.408177 ± 0.35` has recorded something a re-reading can refute.

6.13 WHEN NOT TO TUNE **Eight conditions. Each halts the procedure, and each is a computed condition rather than a matter of judgment.**

(1) The case is SUSTAINABLE. `D <= 0`. Report the margin and stop. A relation that is accumulating does not need tuning, and a margin optimised for its own sake spends reserve that has no other source. the constructed team of Section 5 sits at `-3.495327` octaves with 0.6227 octaves of pitch headroom left; there is nothing to buy there that is worth the disturbance of buying it.

(2) The case is STRUCTURAL. `c/eta > ln N`. No assignment of `u` closes it, so `L1` is unavailable at every value and a lever prescription is unsatisfiable. Only `L2`, `L3`, and `L4` — the last purely to make the case tunable.

(3) The case is TUNABLE but UNREACHABLE. the constructed commons of Section 5: formally tunable with a ceiling margin of 0.058436 nats, and requiring `u* = 0.938078` from `u_bar =`

`0.018614`, a factor of 50.4. Every lever was computed: `eta = 1.00` leaves `+1.8430`, a 91% cost cut leaves `+0.0990`, twelve times the category count leaves `+3.2562`. This is a membership question and issuing a tuning plan for it wastes the effort restructuring needs.

(4) The verdict is not resolved. The uncertainty in the deficit exceeds its magnitude. On the constructed family, `D = +0.584737` with `eta` known to ±0.05 carries `±0.9734` octaves, so the sign is unknown. **Measure, do not act**, and report the precision that would resolve it.

(5) The numerator is negative and the only available move is on `delta`. Doing it makes the outcome measurably worse — `−5.011940` to `−9.021491` on the carried dyad — while the deficit reads unchanged. Where no numerator move is available, the correct action is to say so.

(6) Consent is absent, conditioned, or uninformed. Axiom 6.1 and its three corollaries. This is not a cost-benefit condition and is not weighed against anything.

(7) The reading was requested by someone other than the parties. A reading commissioned about a relationship by a third party cannot satisfy Step 2 in substance even if it satisfies it in form, because the parties' ability to decline is compromised by the requester's standing. No reading in this framework is an instrument for one party to use on another.

(8) The proposed intervention is irreversible and the deficit is small. The stakes are asymmetric: an irreversible move that overshoots cannot be walked back, while a small deficit closes on reversible levers by construction — the family's `+0.584737` closes on a 9.7% rise in `eta`. Reserve the irreversible lever for deficits that the reversible ones cannot reach.

And a ninth that is not a condition but a boundary. Nothing in the procedure decides whether a relationship should continue, and nothing in it evaluates a person. `Q_inf` is where a configuration arrives if the configuration holds; it says nothing about whether it will hold, and nothing whatever about what the relation is for.

The specification

SPEC 6.1 — DISCLOSE FIRST. Outputs, recipients and uses are stated before inputs are taken (Corollary 6.1a). Nothing masked, nothing below threshold, nothing concealed.

SPEC 6.2 — CONSENT PER PARTY, PER STAGE, REVOCABLE, UNCONDITIONED. A decline halts without reason given and without record of cause.

SPEC 6.3 — NO COUNTERFACTUAL BEFORE A DECISION. The value of a party's own withdrawal is computed after the decision, or not at all (Corollary 6.1c).

SPEC 6.4 — CLASSIFY BEFORE PRESCRIBING. Roster ceiling, ceiling test, reachability, in that order (§6.6).

SPEC 6.5 — REPORT SENSITIVITY AND HALT ON AN UNRESOLVED VERDICT. State the precision that would resolve it.

SPEC 6.6 — SIGN THE NUMERATOR AND GATE THE `delta` MOVES ON IT (the ordering rule [Curzi, 2026]).

SPEC 6.7 — COMPUTE THE CLOSED FORMS AND MARK ABSENT LEVERS ABSENT, not expensive.

SPEC 6.8 — CLASSIFY EVERY MOVE BY REVERSIBILITY and propose reversible moves first. Where a reversible move suffices, no irreversible move is proposed.

SPEC 6.9 — SEPARATE AGREEMENT PER MOVE, and additional explicit agreement for irreversible moves.

SPEC 6.10 — RECORD THE BASELINE so that reversal is defined.

SPEC 6.11 — REPORT THE PATH-FREE SPLIT, and name the order used for any per-parameter attribution.

SPEC 6.12 — VERIFY BY RE-MEASUREMENT, WITH DESIGN CONTROLS. Blinding, channel separation, prospective sampling, a yoked control. The shuffle is a weak, mis-calibrated audit with three measured failure modes and is not the verification (CC-R13, CC-R4, CC-R24).

SPEC 6.13 — PREDICT THE MAGNITUDE BEFORE RE-READING. A direction-only prediction carries a likelihood ratio of 2.0000; a magnitude carries 11.3577 (CC-R23).

SPEC 6.14 — TESTIMONY IS NOT EVIDENCE. CC-R23: the survivorship likelihood ratio is exactly `1.000000`.

SPEC 6.15 — NO CLINICAL, THERAPEUTIC, DIAGNOSTIC, COUNSELLING OR MEDIATION FRAMING, in any output, at any stage.

SPEC 6.16 — LABEL EVERY CASE CONSTRUCTED.

The failure modes

FM 6.1 — THE PROCEDURE RUN ON SOMEBODY. *Signature:* a reading of a relationship commissioned by one party, or by a third party, about others. *Mechanism:* Step 2 cannot be satisfied in substance where the requester holds standing over the parties. The ability to decline is what is removed, and that is the one thing the consent axiom of the framework's foundations [Curzi, 2026] forbids absolutely. *Fix:* §6.13 (7). The procedure does not run.

FM 6.2 — THE COUNTERFACTUAL AS PRESSURE. *Signature:* a party told what their withdrawal would do to the reading. *Mechanism:* Corollary 6.1c. And §5.11b shows the figure is uninformative anyway: every member of a healthy team improved the reading by leaving, so the row says nothing about the person it names. *Fix:* SPEC 6.3.

FM 6.3 — TRUST FIRST. *Signature:* a failing relation into which continuity, ritual or patience is invested, followed by acceleration, with the deficit unchanged throughout. *Mechanism:* CC-R15 and the sign-multiplication result [Curzi, 2026]. `−5.011940` to `−9.021491` on the carried dyad; deficit `+4.338420` both times. *Fix:* SPEC 6.6, enforced at Step 12.

FM 6.4 — TUNING AN UNREACHABLE CASE. *Signature:* a lever programme for a group whose required geometry is 50 times its current one. *Mechanism:* Definition 5.1. **An unsatisfiable prescription consumes the effort restructuring needed.** *Fix:* HALT 2.

FM 6.5 — ACTING ON AN UNRESOLVED VERDICT. *Signature:* an intervention justified by a deficit smaller than its own uncertainty. *Mechanism:* SPEC 4.19. `+0.5847 ± 0.9734` is a statement that the sign is unknown. *Fix:* HALT 3, and measure `eta` to better than ±0.030034.

FM 6.6 — THE IRREVERSIBLE MOVE MADE FIRST. *Signature:* a `u` intervention proposed before the reversible levers were computed. *Mechanism:* Step 10. `u` rises when a person takes on commitments they did not hold, and that cannot be undone by ceasing to measure. *Fix:* SPEC 6.8.

FM 6.7 — CONSENT TAKEN ONCE. *Signature:* agreement to a reading treated as agreement to whatever the reading leads to. *Mechanism:* Corollary 6.1b. Agreement to a measurement is not agreement to a change in somebody's life. *Fix:* SPEC 6.9 and Step 14's re-consent.

FM 6.8 — MANUFACTURED DIFFERENCE. *Signature:* a party adopting contrary positions to raise `u`. *Mechanism:* CC-R8. `u` depends only on `|rho|`; opposition reads as `u = 0` and adds cost. *Fix:* report `cos theta` signed, and the opposition identity R1 [Curzi, 2026].

FM 6.9 — THE RE-ENCODING THAT LOOKS LIKE PROGRESS. *Signature:* `u` rises after the description changed rather than the parties. *Mechanism:* the encoding degeneracies R5 [Curzi, 2026]. Disjoint-support encodings read `u = 1` for free. *Fix:* Step 4's overlap check, and re-encode identically before and after.

FM 6.10 — THE VERIFICATION THAT VERIFIES NOTHING. *Signature:* a four-baseline surrogate suite reported as confirmation. *Mechanism:* three measured holes — anti-conservative at 2.2x nominal alpha, blind to repackaging, and passing 98.35% of the time against a shared slow driver. *Fix:* SPEC 6.12. Design controls.

FM 6.11 — TESTIMONY OFFERED AS RESULT. *Signature:* accounts from parties for whom the procedure worked. *Mechanism:* CC-R23. The survivorship likelihood ratio is exactly `1.000000`. A thousand such accounts carry the evidential weight of zero. *Fix:* SPEC 6.14, and re-measured parameters.

FM 6.12 — CUTTING THE CONVERSION MECHANISM. *Signature:* `c` falls, `eta` falls with it, `D` does not move. *Mechanism:* `D` depends on the ratio. *Fix:* SPEC 4.11's test before any `c` cut.

FM 6.13 — THE PROCEDURE USED AS THERAPY. *Signature:* Procedure T presented as an intervention on a person, a treatment, a counselling protocol, or an assessment. *Mechanism:* register slip of the most serious kind. The procedure changes four declared parameters of an arrangement and has no diagnostic, clinical or therapeutic content whatsoever. *Fix:* SPEC 6.15, and §6.7 (1).

Worked example

PART A — TUNING THE CONSTRUCTED FAMILY, ALL FOURTEEN STEPS CONSTRUCTED. This is the constructed family of Section 5 — itself a labelled extension of the carried dyad of the carried dyad [Curzi, 2026]. Parameters are chosen, not observed. This is not an assessment of any person or family, not counselling, and no clinical, diagnostic or therapeutic claim is made or implied. The frequencies are essence frequencies.

Steps 1–2. All four parties are told what will be measured, computed and reported, and to whom. All four agree, separately. No party is shown any counterfactual. **HALT 1 not triggered.**

Step 3. `N = 12` on the state-ordering material's [Curzi, 2026] ordering; one declared period for `eta`, `c` and `delta`; intensities magnitude-estimated per the elicitation rule [Curzi, 2026]; `eta` and `c` declared as estimates with stated uncertainty.

Step 4. Vectors and raw L2 norms:

```
A = {4: 0.7,  6: 0.5,  9: 0.4}                                  0.948683
B = {4: 0.6,  6: 0.55, 9: 0.3,  10: 0.2}                        0.890225
D = {1: 0.50, 2: 0.45, 4: 0.35, 6: 0.30, 8: 0.40, 9: 0.20}      0.930054
E = {3: 0.20, 4: 0.30, 6: 0.25, 9: 0.35, 11: 0.55, 12: 0.50}    0.931397

support overlap: every pair shares at least states 4 and 6      GATE PASSED
```

Step 5.

```
A–B 0.068735   A–D 0.710180   A–E 0.711015
B–D 0.723966   B–E 0.740352   D–E 0.916710

u_bar = 0.645160     F(u_bar) = 2.091463     mean-of-F = 1.954682
aggregation gap = 0.394666 octaves -- reported, both readings carried
N_eff = 8.096755     f = 31.6922 Hz      folded 253.537 Hz (j = 3)
```

Step 6 — classify.

```
roster ceiling:  m(m-1) = 12  <=  2*0.34*ln12/0.13 = 12.9980         PASS, at the ceiling
ceiling test:    c/eta = 2.294118  <=  ln 12 = 2.484907   margin +0.190789   TUNABLE
reachability:    u_req = 0.810517 from u_bar = 0.645160, a factor of 1.2563  REACHABLE

f_min = 47.5308 Hz    folded 380.246 (j = 3)
octaves required 6.619424   present 6.034688
OCTAVE DEFICIT  +0.584737   (conservative +0.979403)      both DECAY -- determined
Q_inf           -1.378049   (conservative -2.308160)
```

HALT 2 not triggered. The case is TUNABLE AND REACHABLE.

Step 7 — sensitivity.

```
dD/d eta = −19.4689 octaves per unit;  dD/dC = +8.4864;  dD/du = −3.9200

eta +/− 0.02  ->  D = +0.5847 +/− 0.3894    resolved
eta +/− 0.05  ->  D = +0.5847 +/− 0.9734    NOT RESOLVED
C   +/− 0.05  ->  D = +0.5847 +/− 0.4243    resolved
u   +/− 0.10  ->  D = +0.5847 +/− 0.3920    resolved
```

> **HALT 3 IS TRIGGERED IF `eta` IS KNOWN ONLY TO ±0.05.** The verdict's sign is not determined and the correct action is measurement. Required precision: `eta` to better than ±0.030034, which is 8.8% of its value.

The procedure therefore pauses here in the first instance, and the family counts what survived the period rather than estimating how well things are going (the estimation practice of [Curzi, 2026]). Suppose that count returns `eta = 0.34 ± 0.02`. The verdict is now resolved and the procedure resumes.

Step 8 — sign the numerator.

```
n = 0.34 * 2.091463 − 0.78 = −0.068902      NEGATIVE
```

No `delta` move is permitted until this changes. This matters here specifically: a family's characteristic parameter is a long memory, `delta = 0.05`, and lengthening it is the intuitive repair. Computed: `delta` 0.05 -> 0.03 takes `Q_inf` from `−1.378049` to `−2.296749`, with `D` unmoved at `+0.584737`.

Step 9 — closed forms and reserves.

```
L1  u_req   = 0.810517       from 0.645160    +0.165357        REACHABLE, irreversible
L2  eta_req = 0.372945       from 0.34        +9.69%           REACHABLE, reversible
L3  c_req   = 0.711098       from 0.78        −8.83%           REACHABLE, reversible
              (kappa 0.13 -> 0.118516)
L4  N_req   = 15             from 12          +3 categories    reversible, and below
                                                               the N = 17 cliff

elasticities: |dD/dln eta| = 6.619424   |dD/dln u| = 2.529026   ratio 2.6174
criterion: c/eta = 2.294118 > 1 − 1/N_eff = 0.876494    cost side proportionally stronger
```

Step 10 — reversibility. Three reversible options each close the deficit alone. **`L1` is therefore not proposed** (SPEC 6.8).

Step 11 — agreement. `L2` is put to all four parties: the household adopts a standing record of what was decided, so that decisions stay decided. All four agree. **HALT 4 not triggered.**

Step 12 — sequence. `L2` is a numerator move and goes first. No `delta` move is scheduled.

Step 13 — act, with baseline recorded.

```
BASELINE:  u_bar 0.645160   eta 0.34   C 0.7800   delta 0.05   D +0.584737   Q_inf −1.378049

eta 0.34 -> 0.40

numerator  = 0.40 * 2.091463 − 0.78 = +0.056585       POSITIVE
c/eta      = 0.78/0.40 = 1.950000
f_min      = k e^(2*1.95) = 23.8824 Hz      folded 382.119 Hz (j = 4)
octaves required 5.626511   present 6.034688
OCTAVE DEFICIT  −0.408177                              was +0.584737
Q_inf           +1.131707                              was −1.378049
VERDICT: ACCUMULATES
```

The swing is 0.992914 octaves from a 17.6% rise in conversion efficiency, computed exactly as `2.885390 * 0.78 * (1/0.34 - 1/0.40)`. The local elasticity of 6.619424 octaves per e-fold would have predicted 1.075782 octaves for the same proportional change; the 8.3% discrepancy is the curvature of `1/eta` over a finite step, and it is the reason SPEC 6.11 asks for the exact contribution rather than the elasticity estimate whenever a move is actually made.

Now, and only now, `delta` is available. Had the family taken it first, it would have moved `Q_inf` to `−2.296749`; taken here it moves `Q_inf` to `+1.886178`. Both orders arrive at the same fixed point and one of them spends a stretch below where it started:

```
 t      delta first      numerator first
 0       +1.000000          +1.000000
 5       +0.534282          +1.029795
10       +0.134353          +1.052849
15       +0.381826          +1.170570
20       +0.594340          +1.271661
30       +0.933545          +1.433018
40       +1.183684          +1.552007
 inf     +1.886178          +1.886178
```

Step 14 — re-read. At a declared interval, restarting at Step 1 with fresh consent. The prediction is recorded first, with a magnitude: `D` is expected at `−0.408177 ± 0.35`. A prediction with a number attached is about 5.7 times more evidentially powerful than one specifying only a direction (CC-R23).

PART B — THE COMMONS: THE PROCEDURE HALTS CONSTRUCTED, the carried commons [Curzi, 2026]. Same disclosures.

Steps 1–5 proceed exactly as above and produce:

```
u_bar = 0.018614     F = 0.186272     N_eff = 1.204754
f = 0.7017 Hz        folded 359.270 Hz (j = 9)
aggregation gap = 0.000001 octaves -- negligible, the pairs are alike
C = 0.8250 at kappa 0.055 x 15 links,  eta = 0.34,  delta = 0.09
```

Step 6 — classify.

```
roster ceiling:  m(m-1) = 30 <= 30.7225          PASS, at the ceiling; a 7th member
                                                 would take it to 42 and make the case
                                                 STRUCTURAL outright
ceiling test:    c/eta = 2.426471 <= 2.484907    margin +0.058436 nats    TUNABLE
reachability:    u_req = 0.938078 from 0.018614  -- a factor of 50.4      UNREACHABLE

f_min = 61.9349 Hz    OCTAVE DEFICIT +6.4638    Q_inf −8.4630
```

> **HALT 2 IS TRIGGERED.** The case is TUNABLE but UNREACHABLE. **The procedure reports the finding and the computations that establish it, and issues no tuning prescription.**

The four lever computations that are reported instead of a plan:

```
L1  u -> 0.95     deficit −0.0332    closes, by 40 cents, and requires near-orthogonality
                                      from six people drawn from one doctrine
    u -> 1.00     deficit −0.1686    the absolute maximum available
L2  eta -> 0.90   deficit +2.1075    fails
    eta -> 1.00   deficit +1.8430    fails at perfect conversion
```

```
L3  kappa -> 0.020 deficit +2.0085   fails
    kappa -> 0.005 deficit +0.0990   fails at a 91% cut
L4  N -> 25       deficit +5.9357    fails
    N -> 144      deficit +3.2562    fails at twelve times the ladder

P8 nesting 15 links -> 7:  deficit +2.7298                fails
shrinking to m = 3:        deficit +0.864532              fails
shrinking to m = 2:        deficit −0.071526              holds, barely
```

No `delta` move is permitted either: the numerator is `−0.761666` and Step 8's gate would refuse it even if the procedure had reached Step 8.

What is reported. That the deficit is `+6.4638` octaves and `Q_inf` is `−8.4630`; that every lever was computed and each falls short; that the case is a membership question rather than an arrangement question (§5.7); and that a restructure — recruiting outside the doctrine, nesting to seven links, adopting a conversion mechanism — reaches `−3.126006` octaves and `Q_inf = +6.620724`, **and produces a different group.**

What is not reported. No recommendation that anyone leave, join, or be replaced. No ranking of members. No counterfactual for any individual's withdrawal. The formalism has no term for what the doctrine was for, no single-party term at all, and no standing to weigh what this group is otherwise doing.

PART C — THE CARRIED DYAD, FOR COMPARISON CONSTRUCTED, the carried dyad [Curzi, 2026]. The same procedure on the case Part III opened with, given briefly because the minimum-frequency and sustainable-dyad papers [Curzi, 2026,?] work it in full.

```
Steps 1–5:  u = 0.068734665264634, F = 0.563085, f = 1.4907987150 Hz
            eta = 0.30, c = 0.62, delta = 0.09
Step 6:     c/eta = 2.066667 <= 2.484907   TUNABLE, margin +0.418240 nats
            u_req = 0.627132 from 0.068735 -- a factor of 9.1240
            REACHABILITY: declared borderline, and declared rather than assumed
Step 7:     D = +4.338420; no plausible input uncertainty flips a deficit that size
Step 8:     numerator = −0.451075   NEGATIVE   -- delta moves refused at Step 12
Step 9:     L1 u -> 0.627132     reachable, IRREVERSIBLE, reserve 5.545206 octaves
            L2 eta -> 1.101077   ABSENT (eta > 1); reserve to eta = 1 is 4.174198 < 4.338420
            L3 c -> 0.168925     a 72.8% cut; reserve 5.963140 octaves
            L4 N -> 102          not reachable; 0.1129 octaves per category
Step 10:    one reversible lever (L3) and one irreversible (L1); L2 does not exist
Step 11–13: the reversible lever alone leaves +0.998357 octaves and still DECAYS,
            so L1 is proposed, separately agreed, and taken together with the cost move
            result: u -> 0.692373, eta -> 0.44, c -> 0.40
            D = −3.590989, numerator +0.547598, Q_inf +6.084427
Step 12:    NOW the delta move is permitted: 0.09 -> 0.05 gives Q_inf +10.951968
Step 14:    re-read window at lambda = 0.05 is [7.35, 31.91] periods -- comfortably open
```

This is the case in which an irreversible lever was correct, and the procedure records why: the reversible levers were computed first, one of them did not exist and the other did not suffice, and the party who took on new states agreed separately to doing so. **The order of those four facts is the whole of Step 10 and Step 11.**

Compare the re-read windows across the three tunes:

```
case                margin (octaves)     u* / u_0     window at lambda = 0.05
family after tune       0.408177          0.849500     EMPTY
carried dyad 2b         3.590989          0.194595     [7.35, 31.91]
```

The dyad's repair is verifiable and the family's is not, and the difference is the margin.

The two outcomes, side by side

```
------------------------------------------------------------------------
                            THE FAMILY          THE COMMONS
------------------------------------------------------------------------
deficit as found              +0.584737            +6.463800
classification            tunable, reachable   tunable, UNREACHABLE
procedure                 ran to Step 14        halted at Step 6
halt encountered          HALT 3 (then cleared)     HALT 2
numerator sign                negative             negative
delta move permitted      no, then yes              no
lever chosen              eta 0.34 -> 0.40           none
reversible?                     yes                  -
deficit after             -0.408177                  -
Q_inf after               +1.131707 (+1.886178
                          with delta 0.03)           -
what was reported         a completed tune      a finding and four
                                                failed lever tables
------------------------------------------------------------------------
ESSENCE frequencies, not physical resonances. Constructed cases. Not an
assessment of any person, family or community. Not counselling, therapy,
mediation or clinical assessment. No diagnostic claim of any kind.
------------------------------------------------------------------------
```

Both outcomes are correct procedure. One produced an intervention and one produced a refusal to prescribe, and the refusal took the same fourteen steps up to the point where it stopped. **A diagnostic that only ever recommends action is not a diagnostic; it is a sales process**, and the commons column is what this framework means by an instrument that can say no.

7 Results

The eight results below are the units other papers in this corpus cite. A citation to this paper should name one of them. Each carries its status: **DERIVED** where it follows from the stated definitions by algebra, **COMPUTED** where it is the output of exact arithmetic on stated parameters, and **POSTULATED** where it is asserted by this framework.

R1 — two octaves per doubling against a quadratic. Theorem 4.2: $\mathrm{NET}(m) = 2\log_2 m - 1.442695\, r\, m(m-1)$ octaves with $r = \kappa/\eta$, so the largest viable roster at cost ratio r is the largest m with $r \le 2\ln m/(m(m-1))$, and **a dyad is viable in the best case if and only if** $r \le \ln 2 = 0.693147$**.** *Status:* DERIVED by substitution; the bound table COMPUTED. The carried dyad's $r = 2.066667$ places it over the roster bound at two, which is why its geometry alone could not save it [Curzi, 2026].

R2 — the roster ceiling. Theorem 4.3: $\kappa m(m-1)/2 \le \eta \ln N$, equivalently $m(m-1) \le 2\eta \ln N/\kappa$, **independent of any geometry**. *Status:* DERIVED from the normalisation bound $F \le \ln N$ [Curzi, 2026] and the ceiling result R5 of the minimum-frequency paper [Curzi, 2026]. **The constructed commons sits exactly on it, 30 against 30.7225**, so a seventh member takes $m(m-1)$ to 42 and makes the case STRUCTURAL outright, at which point no u whatsoever suffices. This result is dimensionally clean and does not depend on postulate P-b (§3.4).

R3 — the roster-scale count and its participation-ratio identity. Theorem 4.1: $M_{\text{eff}} = 1 + (m-1)\bar{u} = m + 1 - m/\mathrm{PR}$, with PR the Rényi-2 participation ratio of the members' Gram matrix [Rényi, 1961]. *Verified* on the constructed triad at PR $= 1.3522930095$, both sides $= 1.7815458788$, **error** 0.00×10^{0}. *Status:* DERIVED; the instance COMPUTED. It is the roster-scale twin of the category-scale identity of Series B [Curzi, 2026c,o]. **It is an algebraic re-expression of the affine axiom P-c and is not independent evidence for it, and M_{eff} is not a Hill number.**

R4 — the aggregation gap at group scope. On the constructed triad $F(\bar{u}) = 1.667424$ nats against a mean of $F(u)$ of 1.491574 nats: a gap of 0.1758499493 nats, **0.507396 octaves, 609 cents, a frequency ratio of 1.421482**, or 11.79 per cent of the conservative figure. On the constructed commons, whose fifteen pairwise values lie in a band of width 0.000346, the same gap is **0.000001 octaves**. *Status:* COMPUTED, from Jensen's inequality and the strict concavity of F [Curzi, 2026]. The same correction is priced at 0.0790 nats (4.31 per cent) at protocol scope in Series B [Curzi, 2026]. **The rule that follows is not "use the conservative figure" but "publish both, and report NOT DETERMINED where the verdict differs between them".**

R5 — the commons reading, and four levers that all fall short. $\bar{u} = 0.018614$, $F = 0.186277$ nats (the source prints 0.186272; the discrepancy is disclosed in §10), $N_{\text{eff}} = 1.204754$ against a required 11.3189 of a possible 12; $c/\eta = 2.426471$ against $\ln 12 = 2.484907$, a margin of $+0.058436$ nats; $D = +6.4638$ octaves and $Q_\infty = -8.4630$. **Every lever swept to its limit fails**: $\eta = 1.00$ leaves $+1.8430$; κ cut by 91 per cent leaves $+0.0990$; $N = 144$ leaves $+3.2562$; nesting to seven links leaves $+2.7298$; shrinking to $m = 3$ leaves $+0.864532$. Only $u \to 0.95$ closes it, by 0.0332 octaves, and u is the one parameter no arrangement controls. *Status:* COMPUTED on a constructed case. **The procedure halts at HALT 2 and issues no prescription.**

R6 — the restructure, decomposed. Changing the membership rule ($\bar{u} : 0.018614 \to 0.45$), nesting 15 links to 7, and adopting a conversion mechanism ($\eta : 0.34 \to 0.55$) gives $D = -3.126006$ and $Q_\infty = +6.620724$, **a swing of 9.5898 octaves — 51.9 per cent from the cost side and 48.1 per cent from the geometry.** No proper subset does as well, and the two arrangement changes without the recruitment leave $+2.729803$ and $+3.790608$, both failing. *Status:* COMPUTED. **It is a different group afterwards, and the formalism has nothing to say about whether the trade is worth making.**

R7 — the drop-one negative result. On a constructed five-member team at $\bar{u} = 0.788228$ with 0.6227 octaves of pitch headroom, **every one of the five removals improves the reading**, raising Q_∞ by between 1.97 and 2.80 and lowering D by between 0.982562 and 1.395957 octaves. *Status:* DERIVED from concavity of F with cost linear in links; COMPUTED on the constructed team. **A favourable drop-one row therefore carries no information about the member it names**, which makes SPEC 4.11's prohibition an arithmetic finding rather than a courtesy.

R8 — the muddiness index. $\mu = \lfloor m(m-1)/2 \rfloor / \lceil N/2 \rceil$ exceeds one from $m = 4$ at $N = 12$: pairs grow quadratically while pitch classes are capped at $\lceil N/2 \rceil = 6$ [Curzi, 2026]. Simulation at 20,000 trials, seed 7, gives expected class coverage of 5.73 of six at six members holding four states each and 5.98 at twelve. *Status:* DERIVED for μ; the coverage table COMPUTED by simulation with a stated seed, **which is a computation and not an observation of anybody.**

CR — the counter-result. The constructed commons is **formally tunable and practically unreachable**: it passes the ceiling test with a margin of $+0.058436$ nats, sits exactly on its roster ceiling, and admits no lever that closes it. And a member's declining to participate improves a constructed family's reading by **2.680588 octaves**, flipping its verdict from DECAYS to ACCUMULATES. **Both are arithmetic this paper prints rather than avoids**, and §6.2 makes the second an axiom-level prohibition on disclosure rather than a finding to be acted on.

8 Discussion

8.1 What follows

The unit of design is the link, not the head. Both bounds in this paper are statements about κ/η, not about m alone: the same three people are sustainable at one link price and structural at another, differing by $D = -3.410906$ against $+13.078$ at a factor of 11.27 in κ. A group-size debate conducted without a declared link price is not a debate the formalism can join.

Topology is the only lever that changes the exponent. Raising conversion, cutting cost and adding categories all move a coefficient. Partitioning the roster moves $m(m-1)/2$ itself, and on six members it removes 53.3 per cent of the cost. That is why the nesting principle [Ostrom,

1990] is structurally different from the other seven, and why the corpus's own rewiring result is the largest single arrangement gain it reports [Curzi, 2026b].

The reachability test earns its place because it changes what gets recommended. The formal classification returns TUNABLE for a case requiring a fiftyfold rise in $\bar{u}$, and a prescription issued on that basis is unsatisfiable. **An unsatisfiable prescription is worse than none**, because it consumes the effort restructuring needed.

8.2 What does not follow

Nothing here shows that small groups outperform large ones. It shows that *this framework's surplus functional* is bounded by $\ln N$ while *this framework's cost model* is quadratic in the roster. Both are stipulations of the framework (P-c, P-d). The empirical literature on group size and performance is mixed and conditional [Mannix and Neale, 2005, van Knippenberg and Schippers, 2007, Woolley et al., 2010], and this paper adds no observation to it.

Nothing here licenses a statement about any person. The apparatus has no single-party term. It is run m times in a drop-one audit and acquires none, and §5.11b shows the audit's rows are uninformative about members precisely when the group is healthy. **A ranking derived from those rows would be a ranking of nothing.**

The commons result is not a result about commons. It is a result about a parameter signature: a flat topology, a moderate link price, and a mean pairwise non-parallelism of 0.018614 produced by a membership rule that selects for agreement. Nothing about land tenure, kinship or employment is claimed, and §5.1 states that the three labels are parameter signatures rather than primitives.

And the counterfactual arithmetic of §5.14 is not an argument. That a reading improves when someone withdraws is a statement about fifteen or six or three links and a concave function. It weighs nothing that person is for, and the framework does not acquire a term for it by being applied to a family.

8.3 Relation to the surrounding programme

The results that survive the failure of the frequency postulates are the ones stated in surplus units, and they are most of the paper: R1, R2, R3, R7 and R8 are all statements about counts, logarithms and a concave functional, and none of them requires k to be a frequency. **R4, R5 and R6 are stated in octaves, and their octave values depend on the exchange constant $2/\ln 2 = 2.885390081777927$, which is exact, and on k, which cancels.** Only the hertz columns of the worked readings depend on postulate P-b. A reader who rejects the essence bridge entirely keeps every bound in this paper, expressed in nats.

Where the wider corpus constrains what may be claimed here, it does so in two places. The concentration work establishes that an effective count collapses to 1.00 once nominally distinct units are recognised as one bloc [Curzi, 2026,?] — which is exactly what the commons' $N_{\text{eff}} = 1.204754$ is reporting. And the founding definitions fix that comfort is the felt total of present costs, so that a present state is a local minimum of the comfort landscape by construction rather than by luck [Curzi, 2026]; **a group reading that returns DECAYS is therefore not a statement that anyone is doing anything wrong**, which is the boundary this paper's §5.12 states in its own terms.

8.4 The reflexive exposure

A published group reading can change the group it reads. The mechanism is documented [Merton, 1948, MacKenzie and Millo, 2003, Rosenthal and Jacobson, 1968], its measured effect sizes in the classroom case are contested and small [Jussim and Harber, 2005], and the framework's own parameters are not exempt from it [Lucas, 1976, Hacking, 1995]. **This paper does not control for it and cannot**, and §6.2's Corollary 6.1c — that no member is shown the

counterfactual value of their own withdrawal before deciding — is the one place where the exposure is met by design rather than acknowledged.

9 Verification

This section decomposes the paper's results into atomic steps, each checkable on its own. The standard is Occam's: **a postulate that can be dropped without the result changing has been dropped.** Every step carries one of four statuses.

- **ARITHMETIC** — reproducible by a reader with a calculator from the inputs named.
- **DEFINITION** — true by stipulation; nothing to verify but consistency of use.
- **CITED** — supported by a named external source.
- **POSTULATED** — asserted by this framework, with no external support.

9.1 The postulates

	postulate	status
P1	The Pythagorean whole tone is the frequency ratio $9{:}8$.	CITED [Barbour, 1951]
P2	φ is the positive root of $x^2 = x + 1$; hence $\varphi^2 = \varphi + 1$.	CITED [Euclid, 1956, Vajda, 1989]
P3	A dimensionless ratio of the form $(9/8)^2/\varphi^2$ **is a frequency in hertz**.	POSTULATED — no source
P4	The essence of a state with n components has frequency kn^2.	POSTULATED — no source
P5	Stored order obeys $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$, $\delta > 0$.	CITED as a *form* [Oppenheim et al., 1996, Elaydi, 2005]; the interpretation POSTULATED
P6	$N_{\text{eff}} = 1 + (N-1)u$ (the affine count axiom).	POSTULATED — design axiom [Curzi, 2026]; **not a Hill number**
P7	$F = \ln N_{\text{eff}}$, hence $F \leq \ln N$.	DEFINITION (given P6)
P8	Total cost is $C = \kappa\, m(m-1)/2$: every pair maintains a link, all links cost alike.	**MODELLING CHOICE** (P-d); the superlinearity CITED [Brooks, 1975]
P9	m and N are cardinalities, hence absolute-scaled.	DEFINITION, typology CITED [Stevens, 1946]
P10	u is symmetric in its arguments and lies in $[0, 1]$.	DEFINITION (from unit vectors)

Dropped as unnecessary. No postulate about the *ordering* of the twelve states is needed for R1–R7: the crossover, the roster ceiling, the roster-scale identity, the aggregation gap, the commons audit, the restructure and the drop-one result all hold for any N and any $u \in [0, 1]$, whatever the states are called. The ladder's ordering is needed only for R8's class-coverage

table. **P3 and P4 are needed only for the hertz columns**, and dropping them leaves every bound in the paper standing in nats.

9.2 The compounding chain

step	consumes	produces	status
L1	P1	$(9/8)^2 = 81/64 = 1.265625$ exactly	ARITHMETIC
L2	P2	$\varphi^2 = \varphi + 1 = 2.618033988749895$	ARITHMETIC
L3	L1, L2	$k = 0.483425732988414$ — **a DEFINED constant, exact**	ARITHMETIC (exact)
L4	L3, **P3**	k carries the unit Hz	**POSTULATED**
L5	L4, **P4**	$f(n) = kn^2$ Hz	**POSTULATED**
L6	P6, P7	$F = \ln(1 + (N-1)\bar{u})$, $F \le \ln N = 2.484906649788$ at $N = 12$	DEFINITION
L7	P5	$Q_\infty = (\eta F - c)/\delta$	ARITHMETIC
L8	L6, L7	$D = 2.885390\,(c/\eta - F)$, with $2/\ln 2 = 2.885390081777927$ exact	ARITHMETIC
L9	P8	$C = \kappa m(m-1)/2$; at $\kappa = 0.055$, $m = 6$: $C = 0.8250$	ARITHMETIC
L10	L9, L6, $c/\eta \le \ln N$	$m(m-1) \le 2\eta \ln N/\kappa$ — **R2**	ARITHMETIC
L11	L10 at $\eta = 0.34$, $\kappa = 0.055$	bound $2\eta \ln N/\kappa = 30.7225$; actual $m(m-1) = 30$	ARITHMETIC
L12	P6 at roster scale	$M_{\text{eff}} = 1 + (m-1)\bar{u}$; PR $= m^2/\sum_{ij} G_{ij}^2$	DEFINITION
L13	L12	$M_{\text{eff}} = m + 1 - m/\text{PR}$ — **R3**; at PR $= 1.3522930095$, both $= 1.7815458788$, residual 0	ARITHMETIC
L14	L6, best case $M_{\text{eff}} = m$	octaves available $= 2\log_2 m$	ARITHMETIC given L5
L15	L8, L9, L14	$\text{NET}(m) = 2\log_2 m - 1.442695\, r\, m(m-1)$ — **R1**	ARITHMETIC
L16	L15 at $m = 2$	$r_{\max}(2) = 2\ln 2/2 = \ln 2 = 0.693147$	ARITHMETIC
L17	concavity of F, Jensen	$F(\bar{u}) \ge \overline{F(u)}$; on the triad $1.667424 - 1.491574 = 0.1758499493$ nats — **R4**	ARITHMETIC

step	consumes	produces	status
L18	L17 $\times\, 2/\ln 2$	0.507396 octaves $= 609$ cents; frequency ratio 1.421482	ARITHMETIC
L19	L6, L8 at commons parameters	$D = +6.4638$, $Q_\infty = -8.4630$, $N_{\text{eff}} = 1.204754$ vs 11.3189 — **R5**	ARITHMETIC
L20	L19 with $\eta \to 1$, $\kappa \to 0.005$, $N \to 144$	$+1.8430$, $+0.0990$, $+3.2562$: all fail	ARITHMETIC
L21	L8 split into two terms	-4.981543 (cost side) $+$ -4.608257 (surplus side) $= -9.589800$ — **R6**	ARITHMETIC
L22	concavity + cost linear in links	every drop-one row improves the reading near the ceiling — **R7**	ARITHMETIC
L23	P9, L5, odd-part class rule [Curzi, 2026]	$\lceil N/2 \rceil = 6$ classes; $\mu = \lfloor m(m-1)/2 \rfloor / 6 > 1$ from $m = 4$ — **R8**	ARITHMETIC + CITED
L24	L5, octave folding	folded values in 220–440 Hz, six pitch classes	**POSTULATED** (meaning); ARITHMETIC (value)

9.3 What this chain is and is not

Read the status column. Steps L1–L3 and L6–L23 are arithmetic or definition and a reader can redo every one of them; three of them (L1, L2, L11) can be redone in under a minute with a calculator, and L11 is the one that decides the commons case. **Steps L4, L5 and L24 are postulated. A chain containing postulated steps is not a proof of its conclusion**, and this one is not offered as one.

What it is offered as is a separation. **Every bound in this paper lives at or below L10 and L15, and neither consumes L4 or L5.** The roster ceiling is an inequality in κ, η and $\ln N$; the crossover is an inequality in r and m; the aggregation gap is a difference of logarithms; the drop-one result is concavity against linearity. **A reader who rejects P3 and P4 entirely loses only the hertz values and keeps every verdict.**

Two statuses that must not be misread. The exponent 2 at L5 is **DEFINED, not estimated**: no fit was performed, no alternative exponent was tested, and the standard procedure for testing a scaling claim [Clauset et al., 2009] was **not run, because there is no data to run it on.** And the logarithm at L6 is **DEFINITION given P6**; the warrant the argument actually needs for it is the multiplicative-dynamics one [Peters and Gell-Mann, 2016], not the psychophysical one [Fechner, 1860].

One check that a reader should run and that this paper wants run. The commons' roster-ceiling margin is $30.7225 - 30 = 0.7225$ in $m(m-1)$. The next admissible value of $m(m-1)$ is $7 \cdot 6 = 42$. **The case is therefore one member from structural, and the margin is smaller than the granularity of the variable it is measured in** — which is the sharpest statement of the commons' position available, and it is three subtractions.

10 Limitations

Every check the governing rigour standard specifies is listed here by name, whether or not this paper passes it.

L-1. Every case is constructed. The triad, the family, the team and the commons are chosen parameter sets. Their figures are exact because the inputs are exact; **that precision does not transfer to a measured case**, and reporting $D = +6.463800$ from an η known to one significant figure would be false precision. **No dataset was collected for this paper and none exists.**

L-2. The uniform complete-graph cost model is a modelling choice, and it is load-bearing. $C = \kappa m(m-1)/2$ assumes every pair maintains a link and every link costs alike. Real topologies are neither, the corpus's own nesting result depends on their not being [Curzi, 2026b], and **the quadratic that drives R1 and R2 is a property of the model, not a measured growth law.** No exponent here was estimated and none was tested [Clauset et al., 2009].

L-3. The dimensional audit does not close at $f(n) = kn^2$. The equation balances only if k carries $[T^{-1}]$, and k is a ratio of two dimensionless quantities. **The hertz is attached by assertion** (P-b, §3.4). No citation in the bibliography repairs it. The bounds of this paper survive it (§9.3); the hertz columns do not.

L-4. The constant is chosen. $k = (9/8)^2/\varphi^2$ combines two real objects into a product with no precedent. It is quoted to fifteen significant figures **because it is defined**, not because anything was measured to that precision, and the deflationary literature on golden-section claims is cited rather than avoided [Markowsky, 1992, Fischler, 1981].

L-5. Octave reduction beyond pitch is unwarranted. Octave equivalence is grounded for pitch [Shepard, 1982, Demany et al., 2013] and nothing grounds folding a non-acoustic quantity into a 220–440 Hz band and reading meaning from where it lands.

L-6. Neither effective count is a Hill number. $1 + (N-1)u$ and $1 + (m-1)\bar{u}$ are linear interpolations and appear in no ecological, economic, political-science or information-theoretic literature. The Hill-number citations here [Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979] apply to $1/\mathrm{HHI}$ and $\exp(H)$ and to nothing else. **The participation-ratio identity of R3 is an algebraic consequence of the affine axiom and is not independent support for it** — a reviewer reaching for it as evidence will find it is not.

L-7. Conjoint measurement was not tested. The ratio-scale treatment of η, κ, δ and the elicited intensities is provisional. Michell's objection stands unanswered [Michell, 1997, 1999] for every quantity in this paper except the counts, and the available remedies [Luce and Tukey, 1964, Krantz et al., 1971, Rasch, 1960] were not run. §3.5 states the consequence: **a quantity whose scale type is not established cannot be known to the 8.8 per cent the family's verdict requires.**

L-8. The reflexive specification is not identified. Stated in §3.6 and repeated here because it is the most serious methodological problem in the corpus: the endogenous social effect in a linear-in-means reflexive loop cannot be recovered from observational data [Manski, 1993, Shalizi and Thomas, 2011]. The remedy adopted is stipulation of the reference group [Brock and Durlauf, 2001], **so every claim here about what a group's parameters imply is a modelling position rather than an identified empirical result.**

L-9. The framework exempts its own parameters from its own reflexivity. If groups change when modelled, η, κ, δ and $\bar{u}$ are not structural [Lucas, 1976, Hacking, 1995]. This is not resolved.

L-10. The functional form is underdetermined. A sub-linear power law [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the qualitative facts as well as this framework's logarithm does, and many mechanisms produce indistinguishable curves [Stumpf and Porter, 2012]. **Nothing in this paper discriminates among them.** The mitigation is narrow and is stated as such: R1, R2, R7 and R8 need only concavity and boundedness, which all three candidates have.

L-11. Scalar reduction is a habit of this framework, and this paper inherits it rather than repeating the concession. The minimum-frequency paper states the limitation once, in full: the framework replaces a phenomenon its own source literature describes by several jointly varying features with a single number, and proceeds arithmetically on the number [Curzi, 2026]. That concession governs here. **This paper reduces a group of m people with $m(m-1)/2$ distinct relationships to one $\bar{u}$, and then to one octave deficit.** §4.3 measures part of the loss — 0.507396 octaves on three members — and §4.2's instruction to publish the full pairwise list is

the only mitigation offered. **It is a mitigation and not an answer**, and the four appraisal-theory works that bear on the general objection are not cited anywhere in this paper as support for any scalar index, because they do not support one.

L-12. u is symmetric and psychological similarity is not. Similarity judgments are asymmetric and violate the triangle inequality [Tversky, 1977]. Where u is read as a similarity measure over people, it cannot assume metric behaviour, and this framework's u does.

L-13. The diversity literature does not support the surplus claim and is not cited as though it did. The formal result that diversity can trump ability [Hong and Page, 2004, Page, 2007] is contested on the ground that it follows from the model's parameterisation [Thompson, 2014, Singer, 2019, Grim et al., 2019, Kuehn, 2017], and the organisational evidence is mixed and conditional [Mannix and Neale, 2005, van Knippenberg and Schippers, 2007]. **The claim that a non-parallel pair reaches the surplus ceiling is a statement about this framework's functional F, and it is not supported by, and does not need, any of that literature.**

L-14. The mean is an unexamined aggregator. §4.3 shows the choice of $F(\bar{u})$ against $\overline{F(u)}$ is worth up to 1.44 octaves in the sweep, but **no argument here establishes that the arithmetic mean over pairs is the right aggregator at all**, and a median, a trimmed mean or a weighted mean would give different readings.

L-15. The class-coverage table is a simulation, not an observation. 20,000 trials, uniform draws without replacement, seed 7. It is a computation with a stated seed and trial count and it describes nobody.

L-16. The reachability judgment is declared, not measured. Definition 5.1 requires the analyst to declare whether a required u^* is within what the parties could honestly represent. **That declaration has no instrument behind it.** A factor of 50.4 is called unreachable here; the paper supplies no test that would have said so.

L-17. Boundary behaviour is only partly examined. $\delta \to 0$ is a pathology (nothing decays out of the stock) and is inherited rather than re-derived; $u \to 1$ gives the ceiling and is examined; $\eta \to 0$ makes c/η diverge and places every group above every ceiling, which is a real constraint on the model's domain and **is not examined in the source**; $m \to \infty$ makes both bounds bite at once; $N \to \infty$ raises the ceiling without bound, which the coordination-cost literature says is not available in practice [Brooks, 1975]. These limits constrain the domain and are stated here because the source does not state them all.

L-18. No model comparison was performed. No out-of-sample test, no matched null with the same number of free parameters, and no information criterion, because nothing was fitted. Where the corpus attempted a confirmatory fit it failed its own declared baseline test [Curzi, 2026i], and that is the relevant precedent.

L-19. No statistical claims are made and none are supported. There are no p-values, no effect sizes and no confidence intervals here, because there is no sample.

L-20. The drop-one audit is an m-fold post-hoc comparison. Run without correction and read as a ranking it is a machine for false positives [Simmons et al., 2011, Gelman and Loken, 2013]. **That is a second and independent reason for SPEC 4.11's prohibition**, alongside the arithmetic one at R7.

L-21. Ergodicity is assumed and not argued. Q accumulates over time for one group from parameters that would naturally be estimated across groups; the time and ensemble averages of a multiplicative process are not the same [Peters, 2019], and nothing here establishes that they coincide.

L-22. There is no conservation law. Q is written with the grammar of a conserved quantity and none is supplied, which is the characteristic defect of the genre [Mirowski, 1989]. Validity is claimed only within a specified arrangement [Cartwright, 1999].

L-23. The governance mapping is a correspondence of dynamics, not of agency. Reading Ostrom's principles as moves on η, c and δ [Ostrom, 1990, Cox et al., 2010] is a modelling choice (M-e). **The claim that no principle acts on u is a claim about the mapping, not a finding about institutions.**

An arithmetic error found in the source, printed rather than silently corrected. The com-

mons' surplus is reported as $F(\bar{u}) = 0.186272$ nats in §5.2, §5.3, §5.9 and §5.12, from a reported $N_{\text{eff}} = 1.204754$. **That value is wrong in the sixth decimal place.** Recomputing from the realising vectors, $\bar{u} = 0.018614137$ gives $F = \ln(1.204755507) = 0.186276648$; recomputing from the *rounded* $\bar{u} = 0.018614$ gives $\ln(1.204754) = 0.186275397$. Neither is 0.186272, and §4.3 of this same paper already prints the correct 0.186276647.

Three consequences, all small and all disclosed. First, the summary table of §5.9 shows the conservative aggregation $\overline{F(u)} = 0.186276$ as *larger* than the optimistic $F(\bar{u}) = 0.186272$, which Jensen's inequality forbids. **The inversion is an artefact of the erroneous 0.186272 and not a violation**; with the correct value the inequality holds, at a gap of 4.28×10^{-7} nats — the 0.000001 octaves §4.3 reports. Second, the steady state is quoted three ways for the same case: -8.463000 in the summary tables (a four-decimal figure padded to six), -8.462955 in the prefix table of §5.12b (correct, from the vectors), and -8.462972 in the governance calculation of §5.6 (which inherits the erroneous F). Third, the deficit is likewise $+6.463800$ in the summary tables against $+6.463833$ from the vectors.

None of this changes any verdict, the discrepancies are at the fifth and sixth significant figure of a case whose deficit is $+6.46$ octaves, and this paper quotes $+6.4638$ and -8.4630 throughout for consistency with the rest of the corpus. It is recorded here because a reader recomputing from the printed intermediates will find the disagreement, and because **a paper that prints F to six decimals owes the reader six correct decimals.**

11 Conclusion

What is established. At group scope, coordination cost grows as $m(m-1)/2$ while best-case surplus grows as $2\log_2 m$ octaves, so $\text{NET}(m) = 2\log_2 m - 1.442695\, r\, m(m-1)$ and **a dyad is viable in the best case if and only if** $r = \kappa/\eta \leq \ln 2 = 0.693147$**.** A second and exact bound, independent of any geometry, is the roster ceiling $m(m-1) \leq 2\eta \ln N/\kappa$. Two effective counts exist at group scope and must be reported separately; the roster-scale one satisfies $M_{\text{eff}} = m + 1 - m/\text{PR}$ exactly, verified at 1.7815458788 with residual zero. The choice of aggregation is worth **0.507396 octaves — 609 cents — on a dispersed triad and 0.000001 octaves on a near-parallel commons**, so both must be published. An orthogonal pair already attains $F = \ln N$, so **a group's case is never a surplus case.** And on a healthy team every drop-one row improves the reading, which establishes that such rows carry no information about the members they name.

What was run and failed. A constructed six-member commons at $\bar{u} = 0.018614$ carries a deficit of $+6.4638$ octaves and a steady state of -8.4630. Every lever was swept to its physical limit and every one fell short: $+1.8430$ at perfect conversion, $+0.0990$ at a 91 per cent cost cut, $+3.2562$ at twelve times the category count, $+2.7298$ after nesting. **The procedure halts and issues no prescription.** The restructure that does work — recruiting outside the doctrine, nesting to seven links, adopting a conversion mechanism — swings the reading 9.5898 octaves, 51.9 per cent from the cost side and 48.1 per cent from the geometry, **and produces a different group.**

What is proposed rather than shown. That the ratio $(9/8)^2/\varphi^2$ is a frequency; that essence goes as kn^2; that octave equivalence carries beyond pitch; and that coordination cost is exactly $\kappa m(m-1)/2$ over a complete uniform graph. **The first three are postulates of this framework with no external support and the fourth is a modelling choice**, all four are labelled at every appearance, and every bound above survives the loss of the first three because k cancels from every interval.

What is not claimed at all. No group, person, family, team, commons or room is claimed to resonate at any frequency named here. Nothing in this paper evaluates any person, and no output of the drop-one analysis, the marginal-member test or the withdrawal counterfactual may be presented as a ranking, an assessment or a recommendation about anyone.

12 Declarations

Data and code availability. No dataset was collected for this study and none exists. Every number in the paper is the output of exact arithmetic on parameters stated in the text; every case — the triad, the family, the team and the commons — is constructed and is labelled as constructed at the point of use, with its input vectors and raw L^2 norms published so that any reading can be recomputed. The one simulation, the class-coverage table of §4.4, states its trial count (20,000) and its random seed (7). The figures are generated by the Python scripts named in their captions, which take their inputs from the tables in the text and are available from the author on request. **No empirical measurement of any person, group, family, team or institution underlies any figure, table or number in this paper.**

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the analysis, the figures and the manuscript.

Ethics. No human-subjects research was conducted and no personal data were processed. The constructed triad, family, team and commons are invented parameter sets; they are not derived from, and do not describe, any person or group. The paper specifies a consent protocol (§6.2) and a prohibition on disclosing withdrawal counterfactuals (Corollary 6.1c) as requirements on any future application; **neither has been applied to anyone, because no reading of any real group was taken.**

Preregistration. No test in this paper was preregistered, because no test was run. The hypotheses of §2.1 are stated with their falsification conditions so that they can be preregistered by others; the corpus's one preregistered confirmatory attempt failed its own declared baseline margin and reported the failure [Curzi, 2026i].

Clinical disclaimer. The twelve state classes are not diagnostic categories. Nothing in this paper is counselling, therapy, mediation, family assessment or clinical evaluation, nothing in it is an assessment of any person's health or capacity, and nothing in it may be presented as such.

Chapter 46

Auditing a Traditional Operation Set Against a Parameter Model: Seven Maps, Two Splits, Three Rejections, and the Many-to-One Compound Rule

Abstract. We test whether a dated twelve-operation schedule from the alchemical corpus can be assigned targets in a closed parameter model $(n, N, u, \eta, c, \delta)$ with a compound index G. Of thirteen entries on a list fixed before the results, seven map cleanly, two are ambiguous between operations returning values in different units, one changes only the publication register, and three are rejected — among them a traditional ordering of metals proposed as an essence index, and *exaltation*, which presents a value ordering as a count. Run as a schedule on a constructed dyad, the ninth operation alone supplies -5.305895 octaves of a -5.966526 total, 88.93 per cent, against 11.07 per cent from the eight before it; the only operation acting on the numerator sits ninth of twelve, which is the order the framework's own rule condemns. The compound rule $f = k^2 G^4$, with G the count-weighted geometric mean of component indices, is derived by substitution rather than asserted and reproduces 37 of 44 published composite values to better than one part in 10^{12}; the pure-oxygen row settles the seven disagreements as errata, since $G = 8$ at every multiplicity gives 957.2370 against a published 11.1246. Corollary 5.2 proves the compound rule and the essence law agree at no integer index, coincidence requiring $n = 1/\sqrt{k} = 1.438252434444351$; at $n = 12$ they differ by 6.121291 octaves. The map is many-to-one at rates 1.32 to 5.56 over $T = 2 \dots 6$, losing 0.4028 to 2.4750 bits, with $G = 6$ produced by sixty distinct multisets. Re-indexing leaves the interaction parameter untouched except through overlap: four moves spanning 8603.9100 cents return identical readings to twelve decimal places. All cases are constructed; the constant is defined and cancels from every interval, so no verdict depends on it.

Keywords: operation audit; compound rule; geometric mean index; many-to-one map; re-indexing; order of operations; quantity blindness; errata

Series: C. Frequency, Essence, and Mind **Paper:** CCP-046

1 Introduction

A body of technique can be older than the formalism that would describe it. When that happens there is a determinate question to ask — **whether the old terms name operations the formalism can express** — and the question has an answer that can be worked out and published rather than admired.

This paper asks it of a specific, dated, transmitted sequence: the twelve gates of a fifteenth-century alchemical text, plus one appended candidate index, audited against a closed five-parameter model with a compound term. **The method is an audit and an audit has one property that distinguishes it from an appreciation: it can return a failure, and if it never does, it was not an audit.** Five of thirteen entries carry a failure and they are printed at the same length as the successes, with the arithmetic that condemns them.

Three results follow, and two of them are negative.

First, the audit distribution. Seven entries map cleanly onto a declared parameter; two are `SPLIT`, meaning the single old term covers two distinct operations that return values in different units; one is a `RENDERING`, real about the publication register rather than about any state; and three are `REJECTED` outright — including a traditional ordering of metals proposed as an essence index, which fails all four of the framework's validity conditions, and *exaltation*, which fails because **"nobler" is a value ordering and the index is a count**. Permissible arithmetic is fixed by the transformation group a quantity admits [Stevens, 1946], and a rating dressed as a count inherits none of the count's properties and all of the rating's ambiguity.

Second, and this is the sharpest finding, the tradition's own ordering is wrong by the framework's own rule. Run as a schedule on a constructed dyad, the twelve gates move the octave deficit by -5.966526 in total, of which the **ninth gate alone supplies -5.305895, or 88.93 per cent**, while the eight before it supply 11.07 per cent between them. The ninth gate is the only one in the list that acts on the numerator of the sustainability condition, and the framework's ordering rule requires numerator moves first [Curzi, 2026]. **The same intervention on the loss rate improves the destination in one ordering and damages it in the other**, and the only thing that differs is the sign of the numerator when the move is made.

Third, the compound rule is derived rather than asserted, and it is many-to-one. Substituting $f(n) = kn^2$ into a published recipe for combining essence values collapses in four lines to $f_{\text{compound}} = k^2 G^4$, where G is the count-weighted geometric mean of the component indices. The collapsed form reproduces **37 of 44 published composite values to better than one part in 10^{12}**, and the seven that disagree are shown to be errata rather than evidence of a different law by a row the rule constrains completely: for a single-element species G is the element index at every multiplicity, so pure oxygen admits exactly one value, 957.2370, against a published 11.1246. But the map is not invertible. **By four components, five component sets in six share their compound value with at least one other; by six components, nineteen in twenty do.** The value $G = 6$ alone is produced by sixty distinct multisets over the twelve-state ladder at $T \le 6$. *Dissolution is not the inverse of fixing; it is a search over a fibre, and the fibre has to be published with the answer.*

A fourth result is a non-result and is worth as much. Re-indexing — moving a state from one index to another — is proved to leave the interaction parameter untouched except through changed overlap: **four constructed moves spanning 8603.9100 cents of lattice distance return the identical relationship reading to twelve decimal places.** The lattice and the relationship measure different things and neither refines the other.

Contributions

1. The audit of thirteen entries against four stated conditions, with the distribution 7 `MAPS` / 2 `SPLIT` / 1 `RENDERING` / 3 `REJECTED` and every rejection carrying the specific result that condemns it and a statement of what would reverse it.
2. Finding 4.2: the schedule result — 88.93 per cent of the total movement supplied by the ninth of twelve steps, which is the order the framework's own rule condemns.

3. The compound rule $f = k^2G^4$ derived by substitution, verified against 44 published values, with the seven disagreements diagnosed as errata by the constrained single-element row.
4. Corollary 5.2: $kn^2 = k^2n^4$ requires $n = 1/\sqrt{k} = 1.438252434444351$, not an integer, **so the compound rule and the essence law agree at no index whatever**; at $n = 12$ they differ by 6.121291 octaves.
5. Theorem 5.3, quantity blindness: G is invariant under scaling every multiplicity by a common integer, so $\{3, 5\}$ and $\{3^{\times 5}, 5^{\times 5}\}$ return the identical 52.582598845962.
6. Theorem 5.4 and the fibre enumeration: many-to-one rates from 1.32 to 5.56 over $T = 2 \ldots 6$, with the information loss priced at 0.4028 to 2.4750 bits.
7. Theorem 6.3 and CONSTRAINT T21: re-indexing moves the relationship only through overlap, and reaches only what is already latent — an index is a count, and a count cannot be raised by counting harder.

What this paper does not claim, stated at the front. It does not claim that any audited operation was ever effective, that its practitioners understood it in the terms recovered here, or that any symbol, name, sequence or substance has a physical effect. **Precedent is not evidence**, and no entry upgrades its status on the strength of the tradition's testimony — for which the framework fixes the survivorship likelihood ratio at exactly 1.000000. Every case is constructed and labelled. The frequency register rests on postulates badged in Section 2, and the constant k cancels from every interval in the paper, so no verdict depends on it.

2 Hypotheses, definitions, postulates and modelling choices

2.1 Hypotheses proper

H1 — the mapping hypothesis. Each entry in the audited historical sequence either names an operation on a parameter in the declared set $\{n, N, u, \eta, c, \delta, G\}$ or it does not, and which is the case is decidable by the four conditions of Audit Rule 4.1. **Predicted:** an independent auditor applying the four conditions to the same thirteen entries reaches the same distribution of verdicts. **Falsified by:** an entry on which two auditors applying the stated conditions disagree, with the disagreement traceable to the conditions rather than to the source text. **Status: not tested.** The audit was performed once, by one analyst. Inter-rater reproducibility is the obvious test and it has not been run; §10 records this as the paper's largest untested claim.

H2 — the ordering hypothesis. The tradition's canonical order is not the order the framework's own rule prescribes. **Predicted:** running the gates as a schedule places the single numerator-acting operation late, so that loss-rate interventions occur while the numerator is negative. **Falsified by:** a reading of the source sequence on which the numerator-acting operation comes first, or a demonstration that more than one gate acts on the numerator. **Status: computed** on a constructed dyad in §4's worked example: the ninth gate supplies 88.93 per cent of the movement. The finding is about this reading of the sequence and this constructed case.

H3 — the errata hypothesis. The seven published composite values that disagree with k^2G^4 are transcription errors rather than evidence of a different combination law. **Predicted:** the disagreements are scattered rather than systematic, and at least one lies outside the range the rule can produce at any input. **Falsified by:** a systematic pattern in the residuals, or an alternative closed form reproducing all 44 values. **Status: supported** — the residuals span four orders of magnitude with no pattern, and the pure-oxygen row admits exactly one rule-consistent value at every multiplicity. **No alternative law was searched for**, which is a real weakness of the test and is recorded at §10.

H4 — the non-invertibility hypothesis. The map from component multiset to compound value is many-to-one at rates that grow with the component count. **Predicted:** exhaustive enumeration over the twelve-state ladder shows the ambiguous fraction rising with T. **Falsified by:** a count at which the map is injective. **Status: proved** (Theorem 5.4) and **enumerated exhaustively** for $T = 1 \dots 6$.

H5 — the independence hypothesis. Lattice position and the interaction parameter are independent: a re-indexing moves u only through the overlap it destroys or creates. **Predicted:** re-indexings of very different lattice distance that vacate and occupy the same coordinates return identical u. **Falsified by:** two such moves returning different u. **Status: proved** (Theorem 6.3) and exhibited on four constructed moves spanning 8603.9100 cents that agree to twelve decimal places.

2.2 Definitions — true by stipulation, not falsifiable

- **Def. 1.** The declared parameter set is $\{n, N, u, \eta, c, \delta, G\}$ and it is **closed**: an operation maps if and only if it names an input and an output inside it.
- **Def. 2.** $G := \left(\prod_i n_i^{c_i}\right)^{1/T}$ with $T = \sum_i c_i$, the count-weighted geometric mean of component indices.
- **Def. 3.** j, the transposition exponent, is a **choice of publication register** and is not in the declared set. An operation acting only on j acts on nothing in the system.
- **Def. 4.** *Pitch class* is the odd part of an index [Curzi, 2026]. It is a function of n and is not independently settable.
- **Def. 5.** A *re-indexing* is a change of index on a single fixed ladder reflecting a change in the number of distinguishable components required to specify a state (Definition 6.1). A change of ladder, a change of register and a re-labelling are each excluded, and each exclusion is load-bearing.
- **Def. 6.** $k := (9/8)^2/\varphi^2$. **This is a defined constant.** It is an exact ratio of exact quantities and therefore computable to arbitrary precision; the fifteen significant figures 0.483425732988414 are *exact digits of a definition* and carry no measurement uncertainty because there is no measurement. $k^2 = 0.233700439315386$ inherits the same status. **Every appearance of either at that precision in this paper is so labelled.**
- **Def. 7.** *Latent* means **present in the repertoire and observed at least once**, in some state, on some occasion. It does not mean conceivable (§6.5).

2.3 Postulates — asserted by this framework, with no external support

- **P-a. The essence law $f(n) = kn^2$.** A quadratic relation between a dimensionless count and a frequency. **No external literature supports it. The exponent 2 is defined by the model, not estimated from data**; no fit was performed, no alternative exponent was tested, and the goodness-of-fit and likelihood-ratio procedure a scaling claim would require [Clauset et al., 2009] was not run, because there is no data to run it on. The same applies to the exponent 4 in k^2G^4, which follows from the substitution and is likewise defined rather than measured.
- **P-b. The attachment of hertz to k.** Computing $(9/8)^2/\varphi^2 = 0.483425732988414$ is arithmetic. **Asserting that this dimensionless number *is a frequency* is a separate and much stronger claim**, and no operational procedure is supplied by which the constant acquires the dimension of inverse time. This is postulated, and it occupies its own line of the verification chain (§9.2, step L4).
- **P-c. The squaring in Rule 3 read as a change of structural level.** k^2G^4 carries hertz-squared if the units are taken seriously. The source reads the squaring as a change of structural level rather than as a dimensional operation, and this paper adopts that reading **with the badge on**: it is a design decision, not a derivation, and it is why every Rule 3 value in the paper is labelled *convention-Hz* and admits ratio operations only.

- **P-d. The extension of the latency constraint from cells to psychological states.** CONSTRAINT T21 is the psychological restatement of a constraint established for genome addressing [Curzi, 2026y]. **The cited biology licenses the form of the constraint and not a claim about minds**, and the extension is this work's thesis, badged, and is not established by the cited work.
- **P-e. The affine effective count $N_{\text{eff}} = 1 + (N-1)u$.** A design axiom [Curzi, 2026], **not a Hill number of any order.** The Hill family [Hill, 1973, Jost, 2006] and its political-science rediscovery [Laakso and Taagepera, 1979] ground $1/\text{HHI}$ and $\exp(H)$ exactly and ground this construction not at all.

2.4 Modelling choices — legitimate, but choices

- **M-a. Octave reduction applied to a non-acoustic quantity.** Octave equivalence is formalised for **pitch** [Shepard, 1982], with a measured perceptual specialness [Demany et al., 2013]. Carrying it to a quantity that is not a pitch is a choice of this framework. The *operation* — division by a power of two — is well defined; its *interpretation* is not supported.
- **M-b. Cents as the unit of a move.** Exact rational interval arithmetic in a logarithmic register is standard apparatus [Barbour, 1951, Lindley and Turner-Smith, 1993, Benson, 2006]. **That $2400 \log_2(n/m)$ is the right price for a move between indices is a choice**, and the two established measures of ratio complexity — Euler's *gradus suavitatis* [Euler, 1739] and harmonic distance [Tenney, 2015] — are named in §5.8b and **not used**.
- **M-c. The source list, fixed in advance.** Twelve gates plus one candidate index, chosen before the results were known. Any other list would give another distribution of verdicts, and the distribution is a property of this list.
- **M-d. Multiplicities are counts, never intensities.** Putting an elicited intensity into G produces a quantity that is neither (FM-5.5). This is stipulated and not derived.
- **M-e. Reading a historical sequence as a schedule.** §4's worked example treats the gates as an ordered procedure and assigns each the smallest parameter change its own description licenses. **Those assignments are constructed**, and a different reader could assign differently; the schedule result is a result about this assignment.

2.5 What is expository

The specification blocks (SPEC-4.1 to 4.7, SPEC-5.1 to 5.6, SPEC-6.1 to 6.5), the three failure-mode catalogues, and Procedures COAG-1, SOLV-1 and TR-1 are expository: they fix reporting requirements and name failure modes rather than argue a claim. Their organising question is *what must be published alongside an audited verdict, a compound value or a priced move for a third party to check it.*

3 Framework, dimensional audit, and method

3.1 The parameter model — a shared mathematical form, not a derivation

The storage law $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$ underlies every deficit computed in this paper. **Analogy status: category 2, a shared mathematical form.** It is the discrete leaky integrator of signal processing [Oppenheim et al., 1996], of exponential smoothing [Holt, 1957] and of the integrate-and-fire neuron [Lapicque, 1907, Gerstner and Kistler, 2002]; its fixed point $Q_\infty = (\eta F - c)/\delta$ requires $\delta > 0$, and stability and uniqueness are standard [Elaydi, 2005, Banach, 1922, Curzi, 2026a]. **No claim is made that any physical result about circuits transfers.**

The surplus functional $F(u) = \ln(1+(N-1)u)$ is the corpus's own [Curzi, 2026]; its logarithmic form places it in the same family as Shannon's entropy [Shannon, 1948] and its ceiling $F \leq \ln N$ is the requisite-variety bound [Ashby, 1956]. **Which warrant the logarithm needs**

is worth stating: the psychophysical route [Fechner, 1860] inherits every dispute in that literature, while the multiplicative-dynamics route derives the same form with no psychological assumption at all [Peters and Gell-Mann, 2016]. This paper needs only the second. **The form remains underdetermined by the evidence** — a sub-linear power law [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the qualitative facts as well, and many mechanisms produce indistinguishable curves [Stumpf and Porter, 2012]. Nothing here discriminates among them.

Two warnings attach and are recorded here rather than in a footnote. Time and ensemble averages of a multiplicative process differ [Peters, 2019], and a single parameter written for a group is an aggregation that need not yield a representative agent [Kirman, 1992]. And Q has the grammar of a conserved quantity with no conservation law to give it content, which is the characteristic defect of the genre [Mirowski, 1989]; validity is claimed only within a specified arrangement [Cartwright, 1999], and the boundaries the sociophysics literature draws [Castellano et al., 2009] are the boundaries claimed here.

3.2 The essence bridge and the constant

$f = k\, e^{2F}$ carries a surplus into a pitch register, and its Result One is that $f(m)/f(n) = m^2/n^2$, **so the constant cancels from every interval and does no work beyond setting a reference** [Curzi, 2026k]. That property is load-bearing here: every price in this paper is a difference of logarithms, so **no verdict depends on k, and only the absolute values do.**

$k = (9/8)^2/\varphi^2 = 0.483425732988414$ is a **defined constant, exact**. Its components are real and documented — the Pythagorean whole tone [Barbour, 1951] and the golden section [Euclid, 1956, Vajda, 1989] — but their product **expressed as a frequency in hertz has no precedent in any literature we have found**, and the standard deflationary treatment of golden-section claims [Markowsky, 1992, Fischler, 1981, Livio, 2002] is cited here and again wherever φ appears.

3.3 Dimensional audit

symbol	quantity	dimension	unit	status
n	essence index	$[1]$	count	cardinality, absolute scale
N	category count	$[1]$	count	cardinality, absolute scale
c_i	multiplicity	$[1]$	count	cardinality; **never an intensity**
$T = \sum c_i$	component count	$[1]$	count	derived
G	count-weighted geometric mean	$[1]$	index units	derived; generally irrational
u	interaction parameter	$[1]$	dimensionless, $[0, 1]$	derived from unit vectors
F	surplus	$[1]$	nats	derived; **nats are not a physical unit**
η	conversion efficiency	$[1]$	dimensionless	declared
c	cost flow	$[\mathcal{C}\, T^{-1}]$	per declared period	declared

symbol	quantity	dimension	unit	status
δ	loss rate	$[1]$	per declared period	declared
D	octave deficit	$[1]$	octaves, cents	derived
j	transposition exponent	$[1]$	integer	**register choice, not a parameter**
k	ground essence	**claimed** $[T^{-1}]$	Hz **by postulate**	**defined; the unit is asserted**
k^2G^4	compound value	$[T^{-2}]$ **if taken seriously**	*convention-Hz*	**derived; the label is a convention (P-c)**

Does the audit close? Two places where it does not, and both are named rather than papered over. $f(n) = kn^2$ balances only if k carries $[T^{-1}]$, and k is a ratio of two dimensionless quantities: that is postulate P-b. And $f_{\text{compound}} = k^2G^4$ is **dimensionally hertz-squared**, so the "Hz" label on it is a convention inherited from the source and not a unit: that is postulate P-c, and it is why SPEC-5.4 forbids entering a Rule 3 value into an acoustic calculation and forbids ratioing it against a Rule 1 or Rule 2 value. **Everything else in the table is dimensionless and closes exactly**, including G, every interval in cents, every deficit, and every count. **The audit's verdicts therefore survive the failure of P-b and P-c intact**; only the absolute values do not.

3.4 Scale type, and what licenses the arithmetic

Permissible arithmetic is fixed by the transformation group that leaves a quantity's empirical content unchanged [Stevens, 1946]. Three classes appear here.

Counts. n, N, c_i and T are cardinalities. A cardinality admits only the identity transformation, so squaring an index, taking a product of indices and forming a geometric mean of them are licensed. **This is the only part of the paper whose scale type is secure, and it is exactly the property Gate 10 and Entry 13 are rejected for lacking.**

Derived dimensionless ratios. u, F, D and every interval in cents inherit their standing from the elicited intensity vectors, which come from magnitude estimation. That method is *held* to deliver ratio-scale information [Stevens, 1957], its exponents move with the experimenter's chosen range [Teghtsoonian, 1971], its response biases are catalogued [Poulton, 1968], and the dispute is live [Krueger, 1989, Luce, 2002]. **This paper inherits it whole.**

Declared parameters. η, c and δ are treated as ratio-scaled on a declared period and **nothing establishes that they are.** Michell's objection is the strongest against this framework [Michell, 1997, 1999], the formal test for additive quantitative structure exists [Luce and Tukey, 1964, Krantz et al., 1971], and **it was not run on any quantity in this paper.** The corpus has demoted one of its own constructions on precisely this ground [Curzi, 2026t], and the instrument that would repair the demotion [Rasch, 1960] is not applied here. **The ratio-scale treatment is provisional and is conceded rather than defended.**

3.5 Identification — the objection this framework cannot presently answer

Every relationship number in this paper is computed from a reflexive specification in which the parties constitute the group and the group conditions the parties. **The endogenous social effect in such a specification is not identified from observational data** [Manski, 1993]: it cannot be separated from contextual effects or from correlated unobservables however much data is collected, and influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011].

The remedy adopted here is stipulation of the reference set [Brock and Durlauf, 2001]: the parties, the period and the vectors are declared and published, so that a reading is repro-

ducible even though the parameter is not identified. Its price is stated plainly: **every relationship claim in this paper is a modelling position rather than an identified empirical result.** No result in Sections 4, 5 or 6 depends on identification — the audit verdicts, the compound rule, the fibre enumeration and the re-indexing theorem are all statements about the formalism — but the *worked readings* that instantiate them do, and they are labelled accordingly.

Two objections in the same family are also unanswered: the framework asserts that systems change when modelled and exempts its own parameters from that reflexivity [Lucas, 1976, Hacking, 1995, Merton, 1948], and a reflexive fixed point has been documented occurring in the field [MacKenzie and Millo, 2003], so the exposure is real rather than hypothetical.

3.6 Method

No data were collected. Four steps.

1. **Fix the list before the results.** The thirteen entries are specified in §4.5 and were not revised after any verdict was reached; the candidate target for each is named as an *input*, and changing it after the audit begins voids the entry (SPEC-4.1, FM-4.8).
2. **Apply four stated conditions.** Audit Rule 4.1 (a)–(d). Condition (c) — is the transformation a function rather than a wish — is where value orderings fail, and the failures are published at the same length as the successes.
3. **Derive, then verify.** The compound rule is obtained by substitution in four lines and then checked against 44 published values; the fibre structure is obtained by **exhaustive enumeration** over the twelve-state ladder for $T = 1 \dots 6$, not by sampling.
4. **Publish the refusals.** Every rejection names the condition that failed, the specific result that condemns it, the rescues a reader is likely to attempt, and **what would have to be true for the verdict to change.** A rejection that cannot say what would reverse it is an opinion.

Every figure in this paper is generated by a script named in its caption, taking its inputs from the tables in the text. Nothing was fitted; no model comparison is reported, because there is nothing to compare and no out-of-sample data to compare on.

4 Auditing a traditional operation set against the parameter model

The preceding papers of this series end with an instrument that produces a number and a procedure that mostly declines to act on it. This paper asks a different question: whether a body of technique that predates the instrument by some centuries was describing the same objects.

The position taken here is stated as a thesis and badged as one: **the old operations are these relationships in another notation.** That is a claim about description, not about efficacy. It says that when a fifteenth-century writer set down a sequence of named procedures and insisted on their order, the sequence had structure, and the structure is legible today in a formalism that writer did not have. It does not say the procedures worked. It does not say symbols act on matter. It says: here is a list of operations, here is a list of parameters, and the question of which list maps onto which has a determinate answer that can be worked out and published.

This section works it out. The method is an audit, and an audit has one property that distinguishes it from an appreciation: **it can return a failure, and if it never does, it was not an audit.** Five of the entries below fail. They are printed at the same length as the ones that pass, with the arithmetic that condemns them, because a section recovering twelve successes from twelve attempts would be evidence about the auditor rather than about the material.

Three things are fixed before anything begins.

The declared parameter set is closed. An operation maps if and only if it names an input and an output inside `{n, u, N, eta, c, delta, G}` and the map is computable. There is no thirteenth parameter, no residual, and no category of “acts on the whole.” The decay-blindness

result R4 [Curzi, 2026] already establishes that a real, well-attested intervention can act on a parameter the framework's central diagnostic cannot see; that result is what makes this audit worth running rather than assuming.

The evidential roles are declared once and hold throughout, following the framework's three-drawer honesty classification [Curzi, 2026]. Where a concordance is drawn in this section, the `(+1)` column carries evidence — measured, replicated, falsifiable — and the `(−1)` column carries precedent, which is the fact that a structure was arrived at independently and preserved. Precedent is worth having and it is not evidence. **No entry below treats the antiquity of an operation as a reason to believe it does anything.**

Nothing in this section claims that a symbol, a name, a sequence, or a substance has a physical effect. The essence frequencies are essence frequencies. No person, tissue, or molecule is claimed to resonate at any value here, the twelve states are not diagnoses, and no reading is an assessment of anyone's health.

The law

4.1 What an audit is An audit takes a claim expressed in one vocabulary and asks whether it can be restated in another without loss and without addition. Three outcomes are possible and they are not a scale.

The claim **maps**: there is an operation on the declared parameters that does what the old term says, the map is computable, and running it produces a number. The claim **fails**: the old term names something the formalism can express, and the formalism says it does not work, or does not work as stated. The claim is **empty**: the old term names nothing in the parameter set at all — not something false, something absent.

The procedure is the framework's own feedback-reading rule applied to a vocabulary rather than to a channel: record what was sent, record what came back, and read the difference [Curzi, 2026]. A restatement that returns the original unchanged reports only that nothing in the target vocabulary could touch it.

The three outcomes carry different obligations. A map must be published with its arithmetic. A failure must be published with the specific result that kills it. An emptiness must be published with a statement of what would have to exist for it to be otherwise, because "I found nothing" and "there is nothing" are different claims and only the first is available from an audit.

An audit does not establish that a mapped operation was ever performed, was ever effective, or was understood by its practitioners in the terms recovered here. It establishes that a modern operation exists which the old term describes. The direction of that sentence matters. It is the weakest claim consistent with the evidence and it is the only one this section makes.

4.2 The declared parameter set Everything this framework can compute is a function of seven quantities. Six are inherited; the seventh is Section 5's.

`n`	the differentiation index of a state	CCP-041, a COUNT, ratio-scaled
`N`	the number of distinguishable categories	CCP-041; ceiling `F <= ln N`
`u`	the interaction parameter	CCP-043, CCP-044; `u = 1 - cos^2 theta`
`eta`	conversion of interaction into structure	the storage law
`c`	the per-period cost	the storage law
`delta`	the loss rate	the storage law
`G`	the count-weighted geometric mean	Sec 5, compound composition

Everything else in the framework is derived from these. `f = k n^2` from `n`. `F = ln(1 + (N−1)u)` and `f_rel = k(1+(N−1)u)^2` from `u` and `N`. `f_min = k e^(2c/eta)` from `c` and `eta`. The octave

deficit `D = 2.885390 (c/eta) − log2(f/k)` from four of them, and — the decay-blindness result R4 [Curzi, 2026] — **conspicuously not from `delta`.** The fixed point `Q_inf = (eta F − c)/delta` from all of them. The compound frequency `k^2 G^4` from `G` alone.

Two derived quantities that are not parameters and are frequently mistaken for them:

The transposition exponent `j`. Octave folding multiplies a frequency by `2^j` (the folding rule [Curzi, 2026]). `j` is a choice of register for publication. It changes no relationship whatever: whatever two states shared before folding they share after, and whatever separated them separates them by the same interval in cents. An operation that acts only on `j` acts on nothing in the system.

The pitch class. The odd part of the index (the odd-part class rule, result R1 [Curzi, 2026]). It is a function of `n` and is not independently settable. You cannot move a class without moving an index.

4.3 The three transfer rules, restated as targets

The three transfer rules [Curzi, 2026] fixes three rules. This section uses them as the three places an operation can land.

```
RULE 1   an ESSENCE has a frequency            f = k n^2
         target: the index n

RULE 2   a RELATIONSHIP has a frequency        f_rel = k (1 + (N−1) u)^2
         target: the interaction parameter u, and N

RULE 3   a COMPOUND uses the alchemical rule   f_compound = k^2 G^4
         target: the composition, through G
```

Plus the storage law's three, which are not transfer rules but are parameters an operation can act on: `eta`, `c`, `delta`.

A candidate operation is therefore audited against six targets and one non-target. The non-target is the register `j`, and it exists in the list precisely because one classical operation lands there and nowhere else.

4.4 The audit rule, and the verdict vocabulary

AUDIT RULE 4.1. A classical operation MAPS if and only if all four conditions hold.

```
(a)  it names a starting configuration expressible in the declared parameters
(b)  it names an ending configuration expressible in the declared parameters
(c)  the transformation from (a) to (b) is a function, not a wish
(d)  running it on a stated input produces a number
```

Condition (c) is the one that does the work, and the reason is measurement-theoretic rather than stylistic: permissible arithmetic on a quantity is fixed by the transformation group that leaves its empirical content unchanged [Stevens, 1946], and a value ordering admits every monotone transformation while a count admits only the identity. Treating a rating as though it were a count is the error the psychometric literature has documented longest [Likert, 1932, Thurstone, 1927], and the objection that psychology assumes quantitative structure without testing for it [Michell, 1997, 1999] is its general form. The formal test that would settle it [Luce and Tukey, 1964, Krantz et al., 1971] has not been run on any quantity in this paper. "Raise the material to a nobler state" satisfies (a) and (b) in appearance only: *nobler* is not a coordinate. The count-versus-rating rule [Curzi, 2026] already named this fault under its own heading — smuggling a rating in as a count — and it is the single most common way an old operation fails this audit.

Verdicts, five of them, and they are borrowed from the concordance register of Series B so that the register can absorb this section's rows without translation:

MAPS	all four conditions. The operation IS an operation on a declared parameter. Publish the map and a worked number.
MAPS, SPLIT	the term is used for two distinct operations in the source tradition and they land on different parameters. Both are published, and the ambiguity is reported as a finding.
REJECTED	the operation is expressible and the formalism says it does not do what is claimed. Published at full length with the result.
EMPTY	the operation lands on no declared parameter and no derived quantity. Not false — absent.
RENDERING	the operation is real, computable, and acts only on the register j. It changes how a value is published and nothing about the system. Admissible as a rendering step; inadmissible as an operation on a state.

4.5 The source set The audit needs a fixed list chosen before the results are known, or it is a selection exercise. The list used is **the Twelve Gates of George Ripley's *Compound of Alchymy* (1471)**, a specific, dated, transmitted sequence: calcination, solution, separation, conjunction, putrefaction, congelation, cibation, sublimation, fermentation, exaltation, multiplication, projection.

Fixing the list in advance is the same discipline preregistration imposes on an analysis [Nosek et al., 2018]: a set chosen after the results are seen is a selection exercise, and the multiple-comparison and forking-path literature is explicit about what that costs [Simmons et al., 2011, Gelman and Loken, 2013]. The corpus's own narrative work makes the structural version of the point — that a corpus closed under negation has resultant exactly zero, so what looks like a finding can be determined entirely by the choice of set [Curzi, 2026ad].

Two reasons for this list over the more familiar seven-operation schemes. It is longer, so it contains more chances to fail, which is the property an audit wants. And it is *ordered* — Ripley presents the gates as a sequence in which each is prerequisite to the next — which lets the Worked Example ask a question no unordered list can pose: does the sequence, run as a schedule, do anything the individual steps do not?

A thirteenth entry is appended: not an operation but a candidate index, the traditional ordering of the metals. It is included because the state-ordering material [Curzi, 2026] rejected four candidate orderings in print and this is the one an alchemical section is obliged to consider.

The `(-1)` column for the whole set is Ripley, with the transmission path stated: a Middle English verse text of 1471, circulated in manuscript, printed in Elias Ashmole's *Theatrum Chemicum Britannicum* (1652), and read continuously since. What it records is a procedure and an order. **What it explains about that procedure is not admissible and is not used below.**

4.6 Honesty clauses, restated in place **(1) ESSENCE, NOT RESONANCE.** Every frequency in this section is a position on a logarithmic scale of surplus. Nothing is claimed to vibrate. No physical, chemical, or biological effect of any symbol, name, sequence, or operation is claimed anywhere.

(2) NO CLINICAL CLAIM. The twelve states are not diagnoses. Nothing here evaluates, describes, or prescribes for any person's health, and no operation below is a treatment.

(3) PRECEDENT IS NOT EVIDENCE. Stated in every entry that draws a concordance.

(4) THE READING INHERITS ITS INPUTS. Every number below is exactly as defensible as the `u`, `n`, `eta`, `c` and `N` that produced it. A contested decomposition yields a contested frequency.

(5) CONSTRUCTED CASES. Every case in this section is constructed and labelled as such.

(6) REPORT `D` WITH `delta`. The octave deficit contains no `delta` (the decay-blindness result R4 [Curzi, 2026]), so a reading that gives only `D` cannot see decay. Both are reported in every entry that produces either.

The engineering

Thirteen entries. Each states the operation as the tradition gives it, the candidate target, the audit against Rule 4.1, and — where it maps — a computed number on the carried case.

The carried dyad throughout is the carried dyad [Curzi, 2026] phase 2a, **constructed and labelled as such**:

```
A = {4:0.7, 6:0.5, 9:0.4}      B = {4:0.6, 6:0.55, 9:0.3, 10:0.2}
u = 0.068734665264634          F = 0.563084802746     f = 1.4907987150 Hz
eta = 0.30   c = 0.62   delta = 0.09
D = +4.338420 octaves          Q_inf = −5.011940      DECAYS
```

GATE 1 — CALCINATION. Verdict: MAPS, on `n`, downward. *The tradition.* Reduction of the material by fire to a calx: the volatile is driven off and what remains is what could not be driven off.

The candidate target. The index `n`. Calcination removes components from a specification.

The audit. (a) a state at index `m`; (b) a state at index `n < m`; (c) the transformation is subtraction of named components, which is a function on the decomposition; (d) it produces a number. **All four conditions hold. Calcination is de-indexing.**

Worked, on the carried subject's threat framing (the carried subject [Curzi, 2026], constructed). State 10 is narrative feeling: a feeling placed in a story, which is a feeling with a cause, a duration, a self as object, another mind, a standard, and a narrative position. Strip the last six components and what remains is state 4, situated feeling.

```
10 -> 4     f  48.3425732988 -> 7.7348117278 Hz
            interval  2400 * log2(4/10) = −3172.6274 cents = −2.643856 octaves
            ratio f(4)/f(10) = 16/100 = 0.160000
            class 5 -> class 1                       CLASS CHANGED
```

The interval is exact, it is a ratio of squares, and `k` cancels from it (the cancellation of the constant [Curzi, 2026]). Calcination has a price list, and the price list is the mood-and-trait paper's [Curzi, 2026] transition matrix read backwards.

One honest note that belongs to this gate and not to the tradition. The formalism permits de-indexing as freely as indexing — the matrix is antisymmetric — and reality does not. A person who has modelled another mind cannot elect to stop having modelled it. What actually removes components is `delta`, over time, without being asked. **The symmetry of the matrix is a property of the notation, not a menu.** §6.9 carries this forward.

GATE 2 — SOLUTION. Verdict: MAPS, on the composition, as the inverse of Rule 3. *The tradition.* Dissolution of the calcined body in its own water; the *solve* of *solve et coagula*.

The candidate target. `G`, inverted: given a compound, recover its components.

The audit. All four conditions hold, with one qualification large enough to require a section of its own. **The inverse exists but is not a function: many distinct component sets produce the same compound frequency.** Section 5 owns the *solve* step, proves the non-uniqueness, and computes the fibres. This gate is recorded as MAPS and handed over.

The one number worth stating here, because it is the reason Section 5 exists: the carried dyad's member `A` holds the chord `{4, 6, 9}`, whose product is 216 and whose count-weighted

geometric mean is therefore exactly 6. **Four other chords drawn from the same twelve-state ladder have the same mean and therefore the identical compound frequency:**

```
T = 3, product 216:   {2,9,12}  {3,6,12}  {3,8,9}  {4,6,9}  {6,6,6}
all five give G = 6.0000000000 and f_compound = 302.8757693527
```

Solve cannot tell them apart. That is not a defect of this framework's arithmetic; it is a property of the geometric mean, and Section 5 proves it in general.

GATE 3 — SEPARATION. Verdict: MAPS, on `N`. *The tradition.* Division of the dissolved material into its distinct parts, subtle from gross.

The candidate target. `N`, the number of distinguishable categories — lever L4 of the four-lever result [Curzi, 2026].

The audit. All four conditions hold. Separation is the operation of finding a distinction where one category previously stood — the requisite-variety move, in which a regulator acquires the categories it needs to regulate a variety it could not previously discriminate [Ashby, 1956, 1958] — and its effect on the reading is exactly the minimum-frequency paper's [Curzi, 2026] fourth lever.

Worked on the carried dyad at 2a, with `dD/dN = −2.885390 u / (1 + (N−1)u)`:

```
N = 12   F = 0.563085   f =  1.490799 Hz   D = +4.338420   Q_inf = −5.011940
N = 13   F = 0.601479   f =  1.609785 Hz   D = +4.227638   Q_inf = −4.883958
N = 14   F = 0.638454   f =  1.733339 Hz   D = +4.120951   Q_inf = −4.760710
N = 17   F = 0.741821   f =  2.131409 Hz   D = +3.822698   Q_inf = −4.416154

dD/dN at 2a = −0.112937 octaves per category
```

And the sentence that makes this gate worth auditing rather than admiring:

> **Separation alone closes this dyad's deficit at `N = 102`.**

One hundred and two genuinely distinguishable categories, each one of which must survive the state-ordering material's [Curzi, 2026] count argument and the collision test R8 [Curzi, 2026]. **The operation maps, the operation is real, and on this case it is useless** — which is precisely the kind of finding an audit exists to produce, and precisely the finding that an admiring reading of the same word would have missed.

GATE 4 — CONJUNCTION. Verdict: MAPS, SPLIT — and the split is this section's central finding. *The tradition.* The marriage: two prepared materials are joined into one.

The candidate target. Ambiguous, and the ambiguity is not in the tradition's language but in ours. **Two of the three transfer rules describe a joining, they are different operations, and they return values in different units.**

Reading (i) — Rule 2, the relational join. Two parties enter into a relationship. The result is `u`, and `f_rel = k(1 + 11u)^2` in hertz.

Reading (ii) — Rule 3, the compositional join. Two component sets are pooled into one compound. The result is `G` over the pooled multiset, and `k^2 G^4` in the source's Hz-label, which is dimensionally hertz-squared (Series B [Curzi, 2026k]; §5.6).

Both readings applied to the carried dyad, which is the point of the gate:

```
reading (i)   Rule 2 on A and B          u = 0.068734665264634
                                         f_rel = 1.4907987150 Hz

reading (ii)  Rule 3 on the pooled chord {4x2, 6x2, 9x2, 10x1}
              T = 7, product 466560, G = 6.4542224391
              f_compound = 405.5414451025   (Hz-LABELLED, dimensionally Hz^2)
```

Taking the naive logarithm of the ratio gives 8.087620, and **that number is meaningless, because the two quantities are not commensurable.** It is printed here so that a reader who computes it independently knows it has already been considered and discarded.

FINDING 4.1. *Conjunction* is not one operation. Applied to two people it is Rule 2 and returns a relationship frequency; applied to two component sets it is Rule 3 and returns a compound value in squared units. **Any use of the word that does not declare which reading is in force produces a number that cannot be interpreted, and the tradition's texts do not declare it.** This is a genuine defect in the old notation, recoverable only because the new one keeps its units.

GATE 5 — PUTREFACTION. Verdict: MAPS, on `delta` — and therefore moves the deficit by exactly zero.

The tradition. The blackening; the material is allowed to rot, and the rotting is required rather than permitted. Ripley's insistence on this gate is emphatic.

The candidate target. `delta`, the loss rate. Putrefaction is deliberate, tolerated decay of accumulated structure.

The audit. All four conditions hold. It maps cleanly, and what it maps to is the one parameter the minimum-frequency paper [Curzi, 2026] proved is not a lever on the diagnostic.

Worked on the carried dyad at 2a:

```
delta = 0.09    D = +4.338420    Q_inf = −5.011940
delta = 0.20    D = +4.338420    Q_inf = −2.255373
delta = 0.30    D = +4.338420    Q_inf = −1.503582
```

The deficit does not move by one part in the last printed digit while `delta` more than triples, and the destination moves by 3.508358 units. This is the decay-blindness result R4 [Curzi, 2026] and the deficit identity R3 [Curzi, 2026] in the same three rows: `D = 2.885390(c/eta − F)` contains no `delta` at all, while `Q_inf = −D eta / (2.885390 delta)` is scaled by it.

The reading this licenses is unexpected and worth having. In a relation whose numerator is already negative — which 2a's is — **raising `delta` moves the destination toward zero from below.** It does not repair anything. It makes the arrival less bad by making arrival matter less, which is a real effect with a real price, and it is exactly the shape of a tradition that requires the material to be destroyed before it can be re-formed. The ordering rule CC-CC-R15 [Curzi, 2026] is the same statement with the sign reversed: lowering `delta` in a failing relation accelerates the failure.

The concordance, with roles stated. `(−1)` Ripley, and the wider tradition's insistence that nothing is made without putrefaction first — precedent, recorded, and it carries no evidence. `(+1)` the deficit identity R3 and the decay-blindness result R4 [Curzi, 2026], verified to 2.665e-15 across twelve parameter sets — evidence, and it carries the entire weight. **What the concordance establishes is that a structure was findable by two routes, not that the older route knew what it had found.**

GATE 6 — CONGELATION. Verdict: MAPS, on `G`. Section 5 owns it.

The tradition. Fixing the dissolved material into a stable body; the *coagula*.

The candidate target. `G`, forward: given components, produce the compound.

The audit. All four conditions hold, and this is the one gate for which the tradition's own notation already contains an exact closed form recovered in the essence-law paper of Series B [Curzi, 2026k]:

```
f_compound = k^2 * G^4          G = ( prod_i n_i^(c_i) )^(1/T),   T = sum c_i
k^2 = 0.2337004393
```

Two worked compounds, both **constructed**:

```
{3x1, 5x1}    T = 2   P = 15   G = 3.8729833462   f = 52.5825988460
{2x2, 7x1}    T = 3   P = 28   G = 3.0365889719   f = 19.8702609489
```

Section 5 states the rule in full, proves what it is not, and handles the units. This entry records the verdict and hands over.

GATE 7 — CIBATION. Verdict: MAPS, on `c`, through cadence. *The tradition.* Feeding: the material is given further matter, in small measured quantities, at the correct intervals, and never all at once.

The candidate target. Not a transfer rule at all — the input schedule of the storage law, which enters the diagnostic through `c`.

The audit. All four conditions hold once the schedule is priced. CC-R11, the concavity result of Series B, is the relevant result: **for a concave gain function with `g(0) = 0`, dividing a fixed budget over more occasions never delivers less.** Computed with `g(x) = ln(1+x)` and a budget of 1.0:

```
occasions m      per occasion      total delivered
     1              1.000000            0.693147
     2              0.500000            0.810930
     4              0.250000            0.892574
     8              0.125000            0.942264
    16              0.062500            0.969994
    32              0.031250            0.984693
    64              0.015625            0.992268
   ->inf                                1.000000
```

Sixty-four occasions deliver 43.2% more than one occasion from the identical budget. **This is concavity, not diligence, and it is the arithmetic reason cadence beats intensity.** The concavity itself has two available warrants and they are not equally strong: the psychophysical one [Fechner, 1860, Stevens, 1957] and the multiplicative-dynamics one, which derives a logarithmic gain with no psychological assumption at all [Peters and Gell-Mann, 2016]. **The argument here needs only the second.** The operations literature reaches the same conclusion about batching from queueing and flow rather than from perception [Hopp and Spearman, 2011, Little, 1961, Sterman, 2000]. CC-R11's own warning applies and is restated: this is a result about *distribution of a fixed input*, not a recommendation of greater total input.

On the carried dyad, a cadence change that delivers the same contact at lower per-period cost:

```
c = 0.62   D = +4.338420   Q_inf = −5.011940
c = 0.50   D = +3.184264   Q_inf = −3.678606      (−1.154156 octaves)
c = 0.40   D = +2.222467   Q_inf = −2.567495      (−2.115953 octaves)
```

At `eta = 0.30`, `dD/dc = 2.885390 / eta = 9.617967` octaves per unit of cost. **Cibation is the cheapest gate in the list on this case, and it is the only one of the twelve that a reader is likely to have dismissed as housekeeping.**

GATE 8 — SUBLIMATION. Verdict: RENDERING. Rejected as an operation on a state. *The tradition.* The volatile part is raised by heat and re-condensed above, purer, unchanged in kind.

The candidate target. Announced as `n`. Actual landing: `j`, the transposition exponent.

The audit. Condition (c) fails at the level that matters. Octave equivalence is a perceptual structure formalised for pitch [Shepard, 1982], with a measured perceptual specialness [Demany et al., 2013]; **carrying it to a quantity that is not a pitch is a modelling choice of this**

framework and not a licensed inference (§2.4). *Raised, unchanged in kind* is an exact description of octave transposition — the folding rule [Curzi, 2026] — and octave transposition is precisely the operation this framework has already proved changes nothing:

```
f_audible = f * 2^j        j an integer

folding cannot create a pitch-class collision
folding cannot remove one
every interval in cents is preserved exactly
```

State 3 folded from 4.3508315969 Hz to 278.453222 Hz at `j = 6` is at a higher pitch and is the same state, with the same class, at the same intervals from everything else. **Nothing in the declared parameter set has moved, because `j` is not in the declared parameter set.**

REJECTED as an operation on a state; ADMISSIBLE as a rendering step, which is what the mood-and-trait paper [Curzi, 2026] §7.18 already licenses it as. The distinction is not pedantry. An analyst who believes sublimation raises the material has performed an operation on the publication format and recorded it as an operation on the substance, and the band-shopping failure mode [Curzi, 2026] already names the modern version of the same error: band-shopping, choosing the register until the numbers look better.

This is the first of the five failures, and it is the mildest, because the operation is real. **It is simply real about something else.**

GATE 9 — FERMENTATION. Verdict: MAPS, on `G` and on `u`. *The tradition.* A small quantity of the perfected substance — the ferment, often gold — is added, and the whole takes on its nature.

The candidate target. Two, and here they do not conflict. Adding a component of declared index changes `G` under Rule 3. Adding a state to a party's chord changes the support overlap and therefore `u` under Rule 2.

The audit. All four conditions hold in both readings, and unlike Gate 4 the readings are compatible rather than ambiguous, because they answer different questions about the same act. What must be declared is which one is being reported.

Worked on the compositional reading, **constructed**:

```
{3x1, 5x1}               T = 2   G = 3.8729833462   f =  52.5825988460
+ one component at 7     T = 3   G = 4.7176939803   f = 115.7653513543
                         gain in the Hz-labelled value: +1.138 in log2
```

And on the relational reading, the carried dyad at 2a versus 2b, where the carried dyad [Curzi, 2026] records exactly this act — one party takes on states the other does not hold:

```
u  0.068734665264634 -> 0.692372881355932
f  1.4907987150 Hz   -> 35.8881768941 Hz
D  +4.338420         -> -0.250931      (holding eta = 0.30, c = 0.62 fixed)
D  +4.338420         -> -3.590989      (with 2b's eta = 0.44, c = 0.40)
```

Fermentation is the only gate in the twelve that acts on the numerator of the sustainability condition, and the ordering rule CC-CC-R15 [Curzi, 2026] is the rule that says it must therefore be run first. The tradition places fermentation ninth of twelve. That ordering is not vindicated by this audit and the audit does not pretend it is.

GATE 10 — EXALTATION. Verdict: REJECTED. *The tradition.* The material is raised to a higher, nobler, more dignified nature.

The candidate target. Announced as `n`.

The audit. Condition (c) fails outright. **"Nobler" is a value ordering, not a count**, and the index `n` is defined in the founding definitions [Curzi, 2026] as the number of distinguishable components required to specify the state. There is no operation "increase the dignity" available on a count. The count-versus-rating rule [Curzi, 2026] named this exact fault when it rejected the valence-magnitude rival: a rating presented as a count inherits none of the count's scale properties and all of the rating's ambiguity.

Two attempts to rescue the gate, both of which fail, and both are published because the rescues are what a reader will reach for.

Rescue 1 — read "higher" as "higher index." Then exaltation is the re-indexing of Section 6, already covered by Gates 1 and 9, and the word contributes nothing. Worse, it contributes something false: **the more-is-not-better rule [Curzi, 2026] established that more components is not better, and the twelve-state ladder is explicitly not an ordering by worth.** A reading that makes exaltation equal index-raising imports a value judgement the index does not carry.

Rescue 2 — read "higher" as "above the ceiling." Then the gate claims an operation that raises an essence past the maximum the category set permits, and the surplus ceiling [Curzi, 2026] forbids it:

```
F <= ln N = ln 12 = 2.484907 nats
f <= k N^2 = 69.6133055503 Hz
state 12 sits exactly at the ceiling: headroom 0.000000 octaves
```

The only way past `k N^2` is to raise `N`, which is Gate 3, which is a different operation with a different name and a measured price of 0.112937 octaves per category on the carried dyad.

REJECTED. The operation as stated names no computable transformation, and both readings that would make it computable turn it into an operation the list already contains under another name.

GATE 11 — MULTIPLICATION. Verdict: SPLIT — the iterative reading MAPS, the unbounded reading is REJECTED.

The tradition. The stone is multiplied: repeated passes through the work increase its potency, and the increase is presented without a stated limit.

Reading (i), iterated refinement. MAPS. Repetition of a bounded operation on an accumulating store is exactly the storage recursion, and its exact solution is the storage-law solution [Curzi, 2026a] with the convergence the refinement paper [Curzi, 2026] owns:

```
|Q_t − Q_inf| = (1 − delta)^t |Q_0 − Q_inf|
```

Computed at `delta = 0.09` toward the carried dyad's post-intervention destination `Q_inf = +6.084427`, from `Q_0 = 0`:

```
t = 0     Q = 0.000000    remaining fraction 1.000000
t = 1     Q = 0.547598    remaining fraction 0.910000
t = 5     Q = 2.287549    remaining fraction 0.624032
t = 10    Q = 3.715053    remaining fraction 0.389416
t = 20    Q = 5.161754    remaining fraction 0.151645
t = 50    Q = 6.029940    remaining fraction 0.008955
t = 100   Q = 6.083939    remaining fraction 0.000080

half-life of the gap  =  ln 2 / −ln(1−delta)  =  7.349615 passes
99% closure           =                          48.829785 passes
```

The contraction is Banach's [Banach, 1922, Elaydi, 2005], and the framework states it exactly, together with the stability reading that a correction smaller than the observed difference can be wrong and still converge while a larger one can be right and still diverge [Curzi, 2026].

Reading (ii), unbounded potency. REJECTED. The tradition's claim is that each pass multiplies, so that potency grows without limit with the number of passes. **The formalism says the opposite in the strongest available way: the recursion is a contraction, it has a unique fixed point, and it is reached from any starting value.** Repetition does not ascend. It arrives, geometrically, at a destination fixed entirely by `(eta, F, c, delta)` — and no number of passes changes any of those four. A thousandth pass on unchanged parameters adds `0.910^1000` of the original gap, which is `10^-41` of it.

REJECTED, and the rejection is a strong one: **this is not "unsupported," it is contradicted.** The only way to raise the destination is to change a parameter, which means running a different gate. Multiplication is the gate that most rewards being audited, because the mapped half of it is genuinely useful — the refinement paper [Curzi, 2026] is built on it — and the unmapped half is the exact error that a formalism with a contraction theorem is built to catch.

GATE 12 — PROJECTION. Verdict: REJECTED. *The tradition.* A grain of the finished stone is cast upon a mass of base metal and transmutes the whole of it.

The candidate target. None available.

The audit. Every condition fails, and they fail for reasons the earlier sections have already established.

(a) fails. The operation requires a distinction between "base" and "noble" material that is not a coordinate in the parameter set. This is Gate 10's fault again, at compound scope.

(b) fails. The end state is a quantity of material at a raised index, where the quantity is unbounded and the seed is arbitrarily small. §5.5 proves the compound rule is blind to overall quantity — `G` for `{3x1, 5x1}`, `{3x2, 5x2}` and `{3x5, 5x5}` is the identical 3.872983346207 — so "a mass transmuted" and "a grain transmuted" are the same value under Rule 3, and the claim's whole content is the mass.

(c) fails, and this is the decisive one. The operation asserts a multiplier with no channel. The framework's belief postulate is explicit and the framework's belief postulate restates it: **belief acts only through a coupling with real gain, and where no channel exists nothing happens.** A seed that reorganises an unbounded mass at no cost is a gain of infinity through an unspecified path.

And the hard constraint applies. CONSTRAINT L67 of the genome-addressing paper [Curzi, 2026y] — carried into Section 6 as this framework's own — states that addressing reaches only what is already latent. Projection claims to confer a property the mass did not have and could not have had. **Nothing in this formalism can create a component that was not already in the decomposition.**

The `(-1)` column here is instructive precisely because it is worthless as evidence. The transmutation record is the most heavily survivorship-filtered body of testimony in the tradition: successes were reported, failures were not, and CC-R23 fixes the likelihood ratio for testimonial evidence of this kind at exactly 1.000000. **A likelihood ratio of one is not weak evidence. It is the arithmetic statement that the testimony leaves the prior exactly where it found it.** The refinement paper [Curzi, 2026] treats this at length; the entry records it and moves on.

The corpus's own historical work sets the standard being applied here, and it applies it to material far more attractive than this: a celebrated story is refused as a fabricated data point, a unified reading is rejected at odds of 1,365 to 1, and a headline percentage is badged as an artefact of a chosen synthetic ratio [Curzi, 2026]. **A rejection of one's own best material is what makes the surviving material worth anything**, and the same standard condemns this gate.

REJECTED on all four conditions.

ENTRY 13 — THE PLANETARY-METAL ORDERING AS AN ESSENCE INDEX. Verdict: REJECTED. *The tradition.* Seven metals in a fixed traditional sequence — lead, tin, iron, copper, mercury, silver, gold — with planetary correspondences, transmitted continuously and used as an ordering of worth and of process stage.

The candidate. Use the sequence as the ordinal `n` in Rule 1, giving `f = k n^2` for `n = 1 .. 7`.

The audit against the ladder-validity conditions [Curzi, 2026] (the ladder-validity conditions [Curzi, 2026]), which is the same audit the four psychological rivals received:

```
completeness      FAIL   the set is closed at seven by cosmological
                         convention, not by the domain
non-arbitrariness FAIL   the ordering is by nobility -- a value ranking.
                         CCP-041 rejected exactly this under the
                         valence-magnitude rival
count semantics   FAIL   position 5 is not "five of something." There is
                         no component being counted, so the index has no
                         true zero and no equal units
stability         FAIL   variant orderings are attested; the sequence is
                         not unique across the transmission
```

And the cost of the failure, computed the way the essence-law paper of Series B [Curzi, 2026k] requires — by asking what a plausible re-ordering does to the published pitch classes:

```
swap positions 3 and 5:   class 3 -> class 5
                          folded 278.453222 Hz -> 386.740586 Hz
                          +568.7174 cents -- nearly half an octave

swap positions 3 and 6:   class 3 -> class 3
                          +2400.0000 cents, and the class does not move at all
```

One swap of two disputed positions moves a published pitch class by 568.7174 cents; another swap moves the pitch by two full octaves and changes nothing audible. An ordering whose frequencies are that sensitive to a choice the tradition does not fix cannot support a published number.

The general hazard is well documented outside this framework: an ordering chosen for its aesthetic or traditional appeal, then treated as though it carried quantitative structure, is exactly the pattern the golden-section literature has spent decades deflating [Markowsky, 1992, Fischler, 1981, Livio, 2002]. **That literature is cited here against this framework's own constant as well** (§2.3, §3.3).

REJECTED as an essence index. Admissible, exactly as Series B's correspondence work has it, as a **finite relational graph** — a labelled seven-element set with edges to planets, days, and metals — which is a descriptive statement carrying no claim that the labelling is meaningful. Graphs are the founding surplus paper [Curzi, 2026]'s object. Indices are not.

4.7 The audit table

```
GATE              TARGET        VERDICT       WHAT IT ACTUALLY DOES
-------------------------------------------------------------------------------
1  calcination    n             MAPS          removes components; de-indexes
2  solution       G (inverse)   MAPS          decomposes; NOT uniquely (Sec 5)
3  separation     N             MAPS          lever L4; +0.112937 oct/category
4  conjunction    u  AND  G     MAPS, SPLIT   two operations, two unit systems
5  putrefaction   delta         MAPS          moves Q_inf; moves D by ZERO
6  congelation    G             MAPS          Rule 3; k^2 G^4 (Sec 5)
7  cibation       c             MAPS          cadence; -9.617967 oct per unit c
```

```
8  sublimation    j             RENDERING     changes the register, nothing else
9  fermentation   G  and  u     MAPS          the only gate on the numerator
10 exaltation     --            REJECTED      a rating smuggled in as a count
11 multiplication delta / --    SPLIT         converges; does not ascend
12 projection     --            REJECTED      no channel; ceiling; L67; R23
13 metal ordering n             REJECTED      fails all four validity conditions
-------------------------------------------------------------------------------
MAPS 7   SPLIT 2   RENDERING 1   REJECTED 3      (5 entries carry a failure)
```

Seven of thirteen entries map cleanly. Two are ambiguous in a way the old notation cannot resolve and the new one can. One is real about the wrong object. Three fail outright. That distribution is the section's result, and the two-thirds that survive are worth more for the third that did not.

The specification

SPEC-4.1 — Inputs required before an operation may be audited

```
the operation's name and its source, with transmission path
the operation's statement in the source's own terms, quoted or paraphrased
the candidate target, named from {n, N, u, eta, c, delta, G} or declared absent
the starting configuration, in declared parameters
the ending configuration, in declared parameters
```

SPEC-4.2 — The audit, in order

```
1  test condition (a): is the start expressible? If not, EMPTY or REJECTED.
2  test condition (b): is the end expressible? If not, EMPTY or REJECTED.
3  test condition (c): is the transformation a function? This is where
   value-orderings fail. A term that names a direction of worth rather
   than a coordinate fails here and the failure is published.
4  test condition (d): run it on a stated input and produce a number.
5  if the target is j and only j, the verdict is RENDERING, not MAPS.
6  if two readings land on different parameters, the verdict is SPLIT and
   BOTH are published with their units.
7  publish the verdict with its arithmetic, failures at full length.
```

SPEC-4.3 — Required outputs for a MAPS verdict

```
the map, stated as a function on declared parameters
a worked number on a stated, constructed input
the effect on D, with its sign and magnitude
the effect on Q_inf, with delta stated
the honesty clause that binds the reading
```

The fourth line is mandatory and the decay-blindness result R4 [Curzi, 2026] is why: `D` contains no `delta`, so a report giving only `D` cannot distinguish Gate 5 from doing nothing. Gate 5 is the entry in this section that would have been silently lost by a `D`-only report.

SPEC-4.4 — Required outputs for a REJECTED verdict

```
which of the four conditions failed, named
the specific result that condemns it, cited to chapter and theorem
every rescue a reader is likely to attempt, attempted and reported
what would have to be true for the verdict to change
```

The third line is the one that costs work and is the one that makes a rejection auditable. Gate 10 carries two rescues; Gate 11 carries the split that rescues half of it.

SPEC-4.5 — Concordance rule for this section Where an entry draws a concordance, both columns are named, the roles are stated in the same sentence, and the precedent is not counted toward the evidence. **An entry may not upgrade its status on the strength of the `(−1)` column alone, at any weight, ever.** Gate 12's entry is the worked instance: the `(−1)` column there is large, old, and continuous, and it moves the verdict by nothing, because CC-R23 fixes its likelihood ratio at 1.000000.

SPEC-4.6 — Precision Recompute every derived value from the twelve-or-more-decimal `u` given in the carried dyad [Curzi, 2026], never from a rounded print (the full-precision rule). Intervals in cents to four decimals; frequencies to at least six; deficits to six. Ratios and deficits are unaffected by rounding of `k` because `k` cancels from both.

SPEC-4.7 — Gates on publication

```
G1  every verdict has its arithmetic in print
G2  at least one REJECTED entry appears, or the audit is void
G3  no entry claims a physical effect of any symbol, name, or operation
G4  every case is labelled constructed
G5  units are stated wherever Rule 3 is used; the Hz label is a
    convention and is dimensionally Hz^2
G6  D is never published without delta
```

The failure modes

FM-4.1 — The audit in which everything passes *Signature:* thirteen entries, thirteen maps, an appreciative tone.

Why it happens: the operations are vague enough that a determined reader can find a parameter for any of them. *Nobler* can be read as higher-index; *projection* can be read as a metaphor for teaching; every failure can be dissolved by generosity.

The fix: Rule 4.1 condition (c), applied without generosity. **A transformation that becomes computable only after the auditor supplies its content was supplied by the auditor.**

FM-4.2 — Precedent counted as evidence *Signature:* "both the tradition and the formalism agree that X, so X is well supported."

Why it wrong: two sources agreeing does not double the evidence when only one of them is evidence (the framework's evidential-role convention [Curzi, 2026]). What independent arrival establishes is that a structure was findable by more than one route. That is real and it is not corroboration.

The fix: SPEC-4.5. State the roles in the sentence that draws the concordance.

FM-4.3 — Reporting `D` and calling the reading complete *Signature:* Gate 5 audited, `D` unchanged, conclusion "putrefaction does nothing."

Why it is wrong: `D` moved by zero and `Q_inf` moved by 3.508358. The deficit is delta-blind (the decay-blindness result R4 [Curzi, 2026]), and a `D`-only report cannot see the one gate in the list that acts purely on `delta`.

The fix: G6. Both numbers, every time.

FM-4.4 — Mixing the two conjunctions *Signature:* a compound value in the source's Hz label compared with a relationship frequency in hertz, and a ratio taken between them.

Why it is wrong: Finding 4.1. The values are not commensurable; Rule 3's output is dimensionally hertz-squared.

The fix: declare the reading before computing. If both are wanted, publish both, separately labelled, and take no ratio.

FM-4.5 — Band-shopping, which is Gate 8 performed without knowing it *Signature:* the folded values are recomputed in a different band until two states look close.

Why it is wrong: folding preserves every interval exactly (the folding rule [Curzi, 2026]). Two states that share a class share it in every band; two that do not, do not.

The fix: the band is declared before the values are computed, and `j` is published with every value.

FM-4.6 — Treating the sequence as validated because the gates are *Signature:* seven gates map, therefore Ripley's twelve-step order is the correct order of operations.

Why it is wrong: the audit tested the gates individually. **On the ordering, the audit's finding is negative: fermentation is the only gate acting on the numerator, CC-R15 requires the numerator first, and the tradition places fermentation ninth.**

The fix: order is a separate claim and needs a separate audit. The Worked Example runs it.

FM-4.7 — Reading a rejection as a claim about the practitioners *Signature:* "the tradition was wrong about projection, therefore it was unserious."

Why it is wrong: the audit is about the mapping, not the people. Seven gates map, and one of them — cibation — encodes a concavity result that was not stated formally until much later, from a tradition with no access to the theorem.

The fix: keep the two claims apart in the sentence. The refinement paper [Curzi, 2026] handles what the tradition got right at length, and the refinement paper [Curzi, 2026] the reverse, and neither is a verdict on anyone's intelligence.

FM-4.8 — Auditing an operation against a target chosen after seeing the result *Signature:* the target is left open until a parameter is found that makes the operation pass.

Why it is wrong: it is the target-selection version of the target-selection failure mode [Curzi, 2026]. With seven parameters and a free choice, almost anything maps.

The fix: SPEC-4.1 requires the candidate target to be named as an input. Changing it after the audit begins voids the entry and it is re-run from step 1.

FM-4.9 — Importing worth into the index *Signature:* an operation is said to raise a state, and the reader takes "higher `n`" to mean "better."

Why it is wrong: the more-is-not-better rule [Curzi, 2026] established that more components is not better. The ladder is a count of specification, not a ranking of value, and the

dissonance paper [Curzi, 2026] established that the fully consonant self is the fully redundant one.

The fix: Gate 10's rejection, restated wherever an operation moves an index.

FM-4.10 — Using an audited operation as an instruction *Signature:* "Gate 7 maps, therefore run Gate 7 on your relationship."

Why it is wrong: the audit says an operation exists and is computable. The group paper [Curzi, 2026] owns what may be done with a reading, including the consent axiom and the halting conditions, and this section adds no permissions whatever.

The fix: the audit's outputs are verdicts, not prescriptions. Nothing here is counselling, therapy, mediation, or clinical assessment, and none of it may be presented as any of those.

FM-4.11 — The unbounded reading of a converging process *Signature:* "each pass makes it stronger, so more passes make it stronger still."

Why it is wrong: Gate 11. `(1-delta)^t` closes 99% of the gap in 48.83 passes at `delta = 0.09` and asymptotes there. The destination is set by four parameters and repetition changes none of them.

The fix: before iterating, compute `Q_inf`. If it is where you want to be, iterate. If it is not, iteration is the wrong operation and one of the four levers is the right one.

FM-4.12 — Publishing a Rule 3 value into an acoustic calculation *Signature:* `k^2 G^4` fed into a renderer as hertz.

Why it is wrong: the value carries hertz-squared if the units are taken seriously (Series B [Curzi, 2026k]; §5.6). Only ratio operations are safe on it, and octave folding is a ratio operation.

The fix: re-anchor by folding before rendering, and label the value as convention-Hz everywhere it appears.

Worked example

PART A — THE TWELVE GATES RUN AS A SCHEDULE ON THE CARRIED DYAD **Constructed case, labelled as such.** the carried dyad [Curzi, 2026] phase 2a, unchanged:

```
u = 0.068734665264634    eta = 0.30    c = 0.62    delta = 0.09    N = 12
F = 0.563084802746       f = 1.4907987150 Hz
D = +4.338420 octaves    Q_inf = -5.011940         DECAYS
```

The question this part asks is the one FM-4.6 raised: run the gates in Ripley's order, as a schedule, and record what each one moves. Every step is a **constructed** intervention chosen to be the smallest change that the gate's own description licenses.

```
STEP  GATE            PARAMETER MOVED        D after      Q_inf after
-------------------------------------------------------------------------
 0    (as found)      --                    +4.338420     -5.011940
 1    calcination     u  0.068735 -> 0.040  +4.911002     -5.673412
 2    solution        none (analysis only)  +4.911002     -5.673412
 3    separation      N  12 -> 13           +4.831945     -5.582082
 4    conjunction     ambiguous -- HALT         --            --
 5    putrefaction    delta 0.09 -> 0.20    +4.831945     -2.511937
 6    congelation     none on D (Rule 3)    +4.831945     -2.511937
```

```
 7    cibation        c  0.62 -> 0.50       +3.677789     -1.911937
 8    sublimation     j only                +3.677789     -1.911937
 9    fermentation    u  0.040 -> 0.692373  -1.628106     +0.846388
10    exaltation      REJECTED                  --             --
11    multiplication  iterate only             --             --
12    projection      REJECTED                  --             --
-------------------------------------------------------------------------
total move in D, steps 1-9:  -5.966526 octaves
contributed by step 9 alone: -5.305895 octaves  =  88.93%
contributed by steps 1-8:    -0.660631 octaves  =  11.07%
```

Read the schedule and four things are visible that no individual entry shows.

Step 1 makes it worse. Calcination as the opening gate removes differentiation, lowering `u`, and lowering `u` raises the deficit by 0.572582 octaves. The tradition opens with the gate that, on this case, moves the reading in the wrong direction. That is no refutation of calcination — Gate 1 maps, and de-indexing has legitimate uses — it is a refutation of calcination-first on a relation whose problem is insufficient difference.

Step 4 halts the schedule. Conjunction is ambiguous between Rule 2 and Rule 3, and the two readings return values in incommensurable units. A procedure cannot proceed through a step whose output units are undetermined, so the schedule halts here, exactly as Procedure T of the group paper [Curzi, 2026] halts at an undetermined input. The halt is a finding rather than a failure of the audit; it is resumed below under reading (i) so that the remaining gates can be scored.

Step 5 changes the destination by 3.070145 and the deficit by nothing. `Q_inf` moves from `-5.582082` to `-2.511937` while `D` sits at `+4.831945` in both rows. A report carrying only `D` would show step 5 as a no-op. **The delta-blindness of the deficit is not an abstraction here; it is a row in this table that would have been invisible.**

Step 9 does 88.93% of the work. Fermentation moves `D` by `-5.305895` octaves; the eight gates before it move it by `-0.660631` between them. The one gate acting on the numerator is ninth of twelve, and by the time it runs, eight steps have been spent on the denominator, on `N`, and on `delta`.

FINDING 4.2, and it is the Worked Example's whole point. Ripley's ordering places the only numerator-acting gate ninth of twelve, and the ordering rule CC-CC-R15 [Curzi, 2026] requires the numerator first: in a relation below `f_min`, an intervention on `delta` amplifies the sign already present. Step 5 is exactly such an intervention and it runs four steps before the numerator is repaired.

The two orderings, side by side, running the same three consequential gates:

```
RIPLEY ORDER   putrefaction while the numerator is NEGATIVE
               delta 0.09 -> 0.20 : Q_inf -5.582082 -> -2.511937
               the destination improves, because delta is dividing a
               negative numerator and a larger divisor shrinks it

R15 ORDER      fermentation first : D +4.338420 -> -0.250931, Q_inf +0.289887
               then cibation      : D -0.250931 -> -1.405087, Q_inf +1.623220
               then putrefaction, numerator now POSITIVE
               delta 0.09 -> 0.20 : Q_inf +1.623220 -> +0.730449
               the destination WORSENS, from the identical move
```

Raising `delta` improves the destination in one ordering and damages it in the other, and the only thing that differs is the sign of the numerator when the move is made. That is CC-R15, restated in the old notation, with the arithmetic attached.

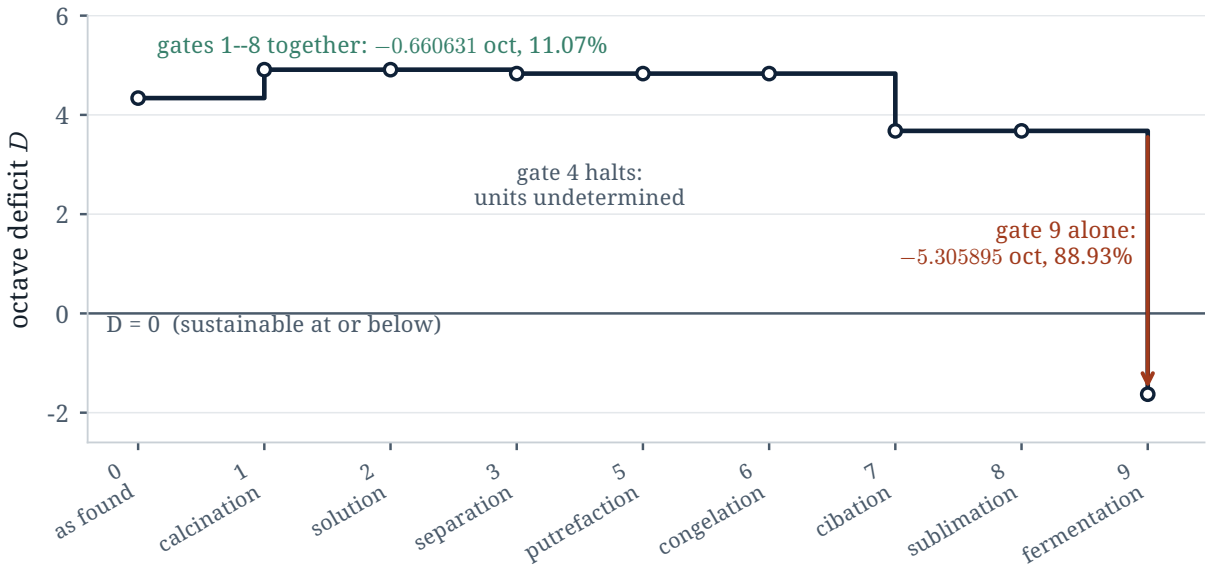


Figure 46.1: **The only gate acting on the numerator is ninth of twelve, and it does 88.93 per cent of the work.** The octave deficit of the carried dyad after each gate of the schedule, in the tradition's own order. Gate 4 halts the schedule because its two readings return values in incommensurable units; gates 5 and 8 move the deficit by exactly zero. Computed on a constructed dyad under the constructed step assignments of §4's worked example; script `figs/src/CCP-046_fig01.py`.

PART B — THE FIVE FAILURES, ON ONE LEDGER Printed together, at equal length, because a section that scattered its failures among its successes would be hiding them in plain sight.

ENTRY	CONDITION FAILED	THE RESULT THAT CONDEMNS IT
8 sublimation	(c) partially -- lands on j	CCP-042 §7.2: folding by 2^j preserves every interval and every pitch class exactly. Nothing in the parameter set moves. Verdict RENDERING. Would change if: j entered any verdict. It does not; k cancels from D.
10 exaltation	(c)	"Nobler" is a value ordering, not a coordinate. CCP-041 §5.18. Rescue 1 makes it Gate 1 or 9; Rescue 2 hits the ceiling F <= ln 12, f <= 69.6133055503. Would change if: a count-valued definition of nobility were supplied. None is, in the source or since.
11 multiplication (unbounded reading)	(c), unbounded reading only	Contraction: |Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|. 99% closure at 48.83 passes; the destination is fixed by four parameters that repetition does not touch. Would change if: a pass altered eta, F, c or delta. Then it is a different gate under a borrowed name.

```
12 projection    (a), (b), (c), (d)   No base/noble coordinate; Rule 3 is
                                      blind to quantity so the mass is not
                                      in the value; no channel (belief
                                      postulate, the belief postulate); CONSTRAINT L67
                                      forbids reaching the non-latent; R23
                                      fixes the testimonial LR at 1.000000.
                                      Would change if: a channel with
                                      measured gain were exhibited. None is.

13 metal order   all four of Paper     Completeness, non-arbitrariness, count
                 B's conditions       semantics and stability all fail. One
                                      disputed swap moves a published class
                                      by 568.7174 cents.
                                      Would change if: the ordering were
                                      fixed by a count. It is fixed by worth.
-------------------------------------------------------------------------------
```

Each row's last field is the one that makes the rejection a scientific statement rather than a dismissal. **A rejection that cannot say what would reverse it is an opinion.**

PART C — WHAT THE AUDIT LICENSES Stated tightly, because the temptation at the end of a section like this one is to say more than was shown.

It licenses: the statement that seven of Ripley's twelve gates name operations on this framework's declared parameters; that the maps are computable and were computed; that two gates are ambiguous in the old notation and resolvable in the new one; that three fail against specific results; and that the sequence, run as a schedule, contradicts CC-R15 on a constructed case.

It does not license: any claim that the operations were effective; any claim about what their practitioners understood; any claim that a symbol, name, or sequence has a physical effect; any diagnostic, clinical, or therapeutic reading of any number above; or the use of a mapped gate as an instruction to act on anyone, which remains the group paper's territory and its constraints [Curzi, 2026].

And it licenses one thing that is easy to miss. The gate that came out best on the carried case is cibation — feeding in small measured quantities at correct intervals — which is the least glamorous item on Ripley's list and the one a modern reader is most likely to skim. It is worth 9.617967 octaves per unit of cost at `eta = 0.30`, and the concavity that makes it work was not written down as a theorem for another five centuries. That is what precedent is worth, stated at its correct weight: not evidence, and not nothing.

Section 5 takes Gates 2 and 6 — solution and congelation — and gives them the section they were handed off to. It owns the compound rule `f = k^2 G^4`, states plainly what that rule is not, and proves why the solve step cannot be inverted uniquely. Section 6 takes Gates 1 and 9, and asks what it costs to move a state from one index to another, and what cannot be reached at any price.

5 Composition and decomposition: the compound rule and its fibres

Two words, joined by a conjunction, and the joining is the instruction. Dissolve, and fix. Take apart, and put together. It is the most compressed statement the old corpus produced and it is almost the only one that survives the audit of Section 4 with both halves intact — Gate 2 and Gate 6, solution and congelation, both returned MAPS, and both were handed to this section.

They arrive in very different condition.

The coagula step has an exact closed form. It was not built for this framework. It was reverse engineered in the essence-law paper of Series B [Curzi, 2026k] out of a published recipe for combining essence values, it collapsed to four symbols, and it then reproduced 37 of 44 published composite values to better than one part in `10^12`. That is an unusually strong result for a formalism recovered from a source rather than derived from axioms, and this section states it, computes with it, and hands the statement to Section 6 and the rendering protocol [Curzi, 2026] as the three transfer rules [Curzi, 2026]'s Rule 3.

The solve step has no closed form, and this section proves that it cannot have one. The map from components to compound is many-to-one, catastrophically so, and it gets worse as the compound gets larger: at four components drawn from the twelve-state ladder, 83.88% of all possible component sets share their compound frequency with at least one other, and the value `G = 6` alone is produced by sixty distinct component sets. Dissolution is not the inverse of fixing. It is a search over a fibre, and the fibre has to be published with the answer.

Three things are settled before the arithmetic.

This section owns the compound rule (the transfer rules). It states it in full and Section 6 and the rendering protocol [Curzi, 2026] cite it. It also owns the negative statement that goes with it, which is more important than the rule and is the single most misquoted result in the series: **`f = k^2 G^4` is not a generalisation of `f = k n^2`, does not reduce to it, and never agrees with it at any integer index whatever.**

The units are not hertz. The composite value carries hertz-squared if the units are taken seriously (Series B [Curzi, 2026k]), and the "Hz" label on it is a convention inherited from the source. Only ratio operations are safe on it. Every value in this section is labelled accordingly.

Nothing here is a physical claim. These are essence frequencies — positions on a logarithmic scale of surplus. Nothing is claimed to vibrate, no molecule is claimed to resonate, the twelve states are not diagnoses, no reading is an assessment of any person's health, and every case is constructed and labelled as such.

The law

5.1 The coagula rule, stated and derived

RULE 3 (the three transfer rules [Curzi, 2026]; the alchemical rule). Let a compound contain essence indices `n_i` with multiplicities `c_i`, and let `T = sum(c_i)` be the total component count. Define `G`, the count-weighted geometric mean of the indices, and the compound frequency:

```
G           =  ( prod_i  n_i^(c_i) )^(1/T)

f_compound  =  k^2 * G^4
```

The derivation is the essence-law paper of Series B [Curzi, 2026k] and is repeated here in four lines because a section that owns a rule states it whole. The source recipe raises each component's essence frequency to the power of its share of the total count, multiplies, and squares:

```
f_compound = ( prod_i f(n_i)^(c_i/T) )^2
           = ( prod_i (k n_i^2)^(c_i/T) )^2
           = ( k^(sum c_i / T) * prod_i n_i^(2 c_i / T) )^2
           = ( k * G^2 )^2
           = k^2 * G^4
```

The third line uses `sum(c_i)/T = 1`, which is what makes `k` come out with an exponent of one and the whole expression collapse. The rule is two constants and one geometric mean, and **everything about a compound that this formalism can see is in `G`.**

5.2 The constant

```
k    = (9/8)^2 / phi^2  =  0.483425732988414     DEFINED, exact
k^2  =                     0.233700439315386     DEFINED, exact
published to ten places:   0.2337004393
the unit "Hz" on either                          POSTULATED (P-b)
```

The two lines above carry two different statuses and the separation is not cosmetic. Computing $(9/8)^2/\varphi^2 = 0.483425732988414$ is arithmetic, and the constant is **DEFINED and exact** — a ratio of exact quantities, so its fifteen figures are digits of a definition and carry no measurement uncertainty. **Attaching hertz to it is a separate and much stronger claim and is POSTULATED** (§2.3, P-b): no operational procedure is supplied by which a dimensionless ratio acquires the dimension of inverse time. k^2 inherits both statuses. And as with k, it cancels from every ratio, **so no verdict in this paper depends on it — only the absolute values do.**

5.3 Verification against the source The essence-law paper of Series B [Curzi, 2026k] tested the collapsed form against every composite value the 2022 source publishes: 24 pharmacological compounds, 8 nucleotides in the carbon column, and 6 each in the silicon- and germanium-substituted columns. **Thirty-seven of forty-four are reproduced to better than one part in 10^{12}. Seven disagree.** Agreement at that level is the corpus's standard for an exact identity — the surplus law's Fechner form is verified the same way, at $N = 2$ to 100 and seven values of u, to a maximum deviation of 8.9×10^{-16} [Curzi, 2026] — and it is what distinguishes a recovered closed form from a fitted one. **No parameter was fitted here and no exponent was estimated**: the exponent 4 in k^2G^4 is defined by the substitution of §5.1, so the goodness-of-fit and likelihood-ratio procedure a claimed scaling relation would require [Clauset et al., 2009] does not apply and was not run. The seven, with the disagreement recomputed here in both relative and octave form:

formula	published	k^2 G^4	rel. deviation	octaves
O2	11.124611798	957.23699944	85.0468	+6.4270
C5H6N2O2	4.669611378	21.80527042	3.6696	+2.2233
C12H17N2O4P	18.534334373	13.36731640	0.2788	−0.4715
C257H383N65O77S6	14.276237883	11.27999043	0.2099	−0.3399
C43H66N12O12S2	12.827862548	11.29835161	0.1192	−0.1832
C77H120N18O26S	11.075644799	10.43459440	0.0579	−0.0860
C158H251N39O46S	9.460360140	9.41751106	0.0045	−0.0065

5.4 Why the seven are errata and not a different law — and why the O2 row settles it

This is the argument that licenses using the rule at all, so it is given in full.

For a single-element species the **count-weighted geometric mean is the element index itself, independent of how many atoms are present.** Oxygen has index 8, so:

```
O    G = (8^1)^(1/1) = 8        O2   G = (8^2)^(1/2) = 8
O3   G = (8^3)^(1/3) = 8        O_x  G = (8^x)^(1/x) = 8    for every x

therefore f = k^2 * 8^4 = 957.2369994358  for EVERY pure-oxygen formula
```

No atom count whatever produces the published 11.1246117974981. That row cannot have come from the stated rule, at any multiplicity, ever. And the shape of the disagreement across the table decides between the two hypotheses: a rule error fails uniformly, a data-entry error fails on scattered rows, and **what is observed is scattered rows** — thirty-seven exact to fourteen significant figures, seven wrong by amounts ranging over four orders of magnitude with no pattern. The rule is confirmed and seven cells are errata.

5.5 What the rule is NOT — the statement this section exists to make Rule 3 looks like a generalisation of Rule 1. It is not, and the confusion is worth killing in detail because it produces numbers that are wrong by octaves.

THEOREM 5.1 (the single-component identity). For a compound of one component at index `n`, at any multiplicity,

```
f_compound  =  k^2 n^4  =  (k n^2)^2  =  f(n)^2
```

Proof: `G = n` for any multiplicity, by §5.4. Substitute. **QED.**

Verified at every index on the twelve-state ladder to a worst absolute error of `9.09e-13`, which is floating-point noise in values of order `10^3`:

n	f(n) = k n^2	f_compound = k^2 n^4	f(n)^2
1	0.4834257330	0.2337004393	0.2337004393
2	1.9337029320	3.7392070290	3.7392070290
3	4.3508315969	18.9297355845	18.9297355845
5	12.0856433247	146.0627745721	146.0627745721
8	30.9392469113	957.2369994358	957.2369994358
12	69.6133055503	4846.0123096438	4846.0123096438

So the compound rule applied to a single-component species does not return that species' essence frequency. It returns the square of it. The two rules answer different questions and produce different numbers from the same input.

COROLLARY 5.2 (they never agree). `k n^2 = k^2 n^4` requires `k n^2 = 1`, i.e. `n = sqrt(1/k) = 1.438252434444351`, which is not an integer. **There is no index at which Rule 1 and Rule 3 coincide.** The nearest approach is at `n = 1`, where they differ by a factor of `k` itself — 0.4834257330 against 0.2337004393, a gap of $-\log_2 k = 1.048634$ octaves — and the gap widens monotonically with `n`, reaching 6.121291 octaves at `n = 12`. **The source manuscript gives the $n = 1$ gap as 1.048946; that figure is wrong by 3.1×10^{-4} octaves and the correction is recorded in §10 rather than applied silently.**

And this is exactly why the source table's O2 row was provably impossible. The two facts are one fact: `G` for a single-component species is index-independent in the multiplicity and returns the essence value squared, so a published composite that neither equals `f(8)^2 = 957.2369994358` nor any other value the rule can produce has no rule-consistent reading at all. The seven outliers are transcription errors, not evidence of a different law, and the O2 row is the one that proves it, because it is the only row whose input is so constrained that the rule permits exactly one answer.

Rule 3 is not a law of the same kind as Rule 1 and must never be described as one. Rule 1 assigns a frequency to a category from its index. Rule 3 assigns a value to a *composition* from the geometric mean of its parts, in squared units, and it is blind to things Rule 1 sees.

5.6 The quantity-blindness theorem **THEOREM 5.3.** `G` is invariant under multiplying every multiplicity by a common positive integer. If `c_i -> m c_i` for all `i`, then `T -> mT`, `prod n_i^(m c_i) = (prod n_i^(c_i))^m`, and

```
G' = ( P^m )^(1/(mT)) = P^(1/T) = G
```

QED. Computed:

```
{3x1, 5x1}    T = 2    P = 15         G = 3.872983346207   f = 52.582598845962
{3x2, 5x2}    T = 4    P = 225        G = 3.872983346207   f = 52.582598845962
{3x3, 5x3}    T = 6    P = 3375       G = 3.872983346207   f = 52.582598845962
{3x5, 5x5}    T = 10   P = 759375     G = 3.872983346207   f = 52.582598845962
```

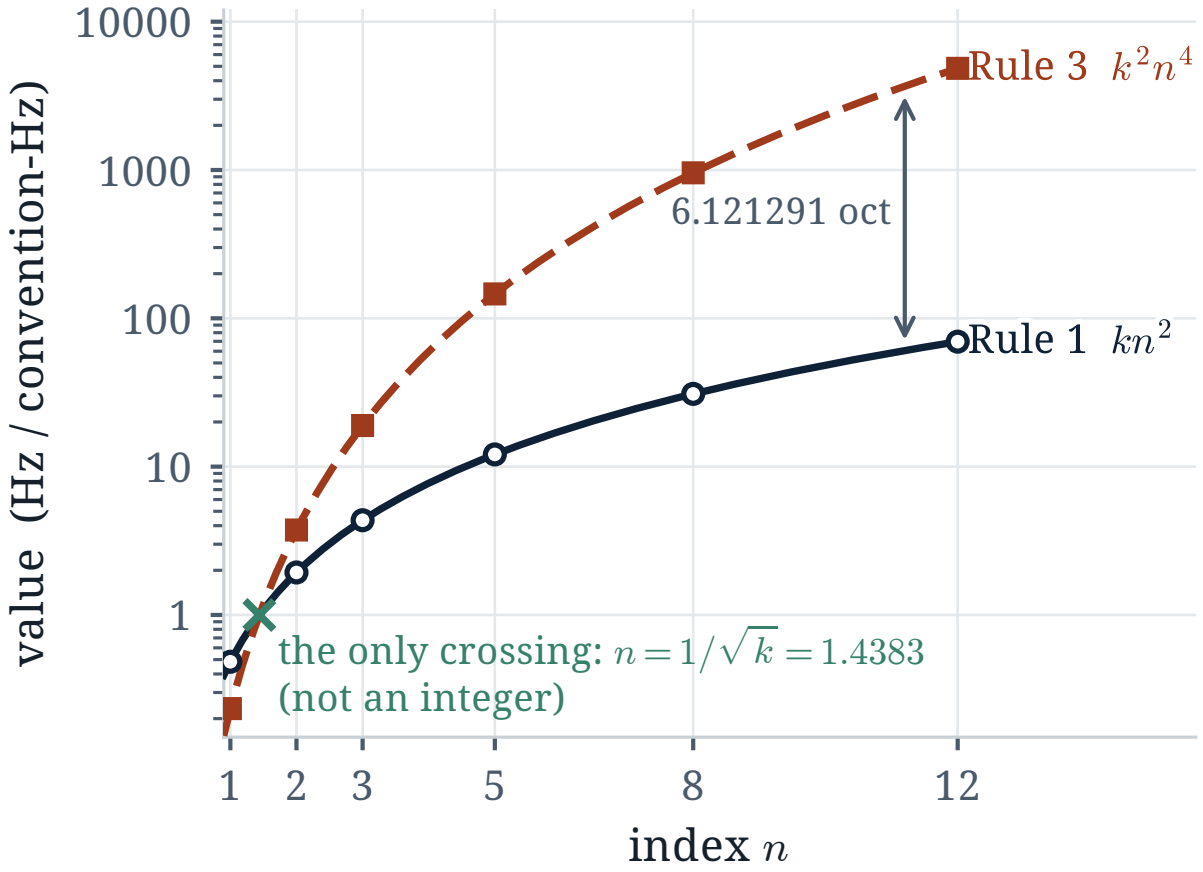


Figure 46.2: **The compound rule and the essence law agree at no index.** kn^2 against k^2n^4, the compound rule's single-component case, which returns the *square* of the essence value. The two curves meet only at $n = 1/\sqrt{k} = 1.438252434444351$, which is not an integer, so no index on any ladder makes them coincide; at $n = 12$ they differ by 6.121291 octaves. k and k^2 are DEFINED constants, exact. Computed from the verification table of §5.5; script `figs/src/CCP-046_fig02.py`.

Rule 3 cannot see how much of a compound there is. It sees only the proportions. The corpus's structural work reaches the same conclusion in a different domain and states it as a bound on what may be claimed: adding a second entity of the same kind in the same jurisdiction buys capability, not diversification, because the pair remains one line in the concentration count [Curzi, 2026]. **More of the same is not more.** The single-component case of §5.4 is the special case `i = 1`; this is the general statement, and it is the structural reason Section 4 rejected Gate 12: an operation whose entire content is *a small seed transmutes a large mass* asserts something about quantity, and quantity is not in the value.

5.7 The units, and the re-anchoring rule The squaring destroys dimensional consistency. `f(n)` is quoted in hertz; `k^2 G^4` carries hertz-squared if the units are taken seriously. The source reads the squaring as a change of structural level rather than a dimensional operation, and this series adopts that reading with the badge on: it is a designed formalism, and the reading is a design decision, not a derivation.

```
RE-ANCHORING RULE (Series B CCP-017, ES-10; binding here)

1  a Rule 3 value is labelled convention-Hz wherever it is printed
2  only RATIO operations may be performed on it
3  octave folding is a ratio operation and is therefore safe
4  it may NOT be entered into an acoustic calculation as hertz
5  publish G and T with every compound value, so it is reproducible
```

Line 5 is not bookkeeping. §5.9 will show that `G` is recoverable from `f` and `T` is not, so a compound value published without `T` has thrown away the only piece of information that narrows the decomposition at all.

5.8 The contraction of the pitch space Because `G` is a geometric mean, low-index components pull it toward 1 and high-index components raise it slowly. Over all 1365 four-component multisets drawn from the twelve-state ladder:

```
G      min 1.000000   Q1 4.119534   median 5.513381   Q3 7.135243   max 12.000000
f      min 0.233700   Q1 67.305727  median 215.939206 Q3 605.751539 max 4846.012310

full span            14.339850 octaves
interquartile span    3.169925 octaves
middle 50% occupies  22.11% of the full octave span
```

The extremes are reached only by the degenerate multisets — all-ones and all-twelves — and half of all compositions sit inside a 3.17-octave band. **Composition is strongly contractive on the pitch space and should be expected to cluster.** Series B observed the same effect empirically in the source's 24 pharmacological compounds, whose `G` ranged only from 1.84 to 4.30.

One structural note that follows from Theorem 5.1 and is worth having in the reference: the compound span over the ladder is 14.339850 octaves and the essence span is 7.169925 octaves, exactly half. Squaring a frequency doubles its position in octaves, so the compound space is the essence space stretched by two — the same information, twice as loud, and in the wrong units.

5.8b The compound interval The mood-and-trait paper [Curzi, 2026] established that the interval between two states is a ratio of squares and that in cents it is `2400 * log2(n/m)`. The compound analogue follows directly from `f = k^2 G^4`, and it is the fourth power rather than the second:

```
interval in cents between two compounds  =  1200 * log2(f2/f1)
                                         =  4800 * log2(G2/G1)
```

`k^2` cancels, exactly as `k` does in the cancellation of the constant [Curzi, 2026], so **no interval between compounds depends on the chosen constant.** Worked on this section's table:

```
{3x1,5x1}   vs  {2x2,7x1}       G 3.8729833462 -> 3.0365889719   −1684.7696 cents
{4x1,6x1,9x1} vs {3,5,7,10}     G 6.0000000000 -> 5.6924250976    −364.4117 cents
threat {3,5,7} vs mirror {5,7,10} G 4.7176939803 -> 7.0472987321  +2779.1450 cents
{11x1,12x1} vs {1x1,2x1}        G 11.4891252931 -> 1.4142135624 −14506.5459 cents
```

The unit is the cent, and the apparatus of exact rational interval arithmetic in a logarithmic measure is standard [Barbour, 1951, Lindley and Turner-Smith, 1993, Benson, 2006, Rossing et al., 2002]. **Whether a ratio's arithmetic complexity is the right thing to price a move by is a separate question, and the literature answers it two ways**: Euler's *gradus suavitatis* grades an interval by the exponents of its prime factorisation [Euler, 1739], and harmonic distance grades it by $\log(ab)$ on a lattice [Tenney, 2015]. **Neither is used here.** This paper prices a move by $2400 \log_2(n/m)$ alone, which is a distance in a logarithmic register and carries no claim about consonance, and what a listener finds consonant is in any case a function of spectrum rather than ratio [Sethares, 2005, Plomp and Levelt, 1965].

And the price list, which is the compound-scale version of the compression the transition matrix [Curzi, 2026] found in the transition matrix:

```
d(cents)/dG  =  4800 / (G ln 2)

at G =  1.5      4616.6241 cents per unit of G
at G =  2.0      3462.4681
at G =  3.0      2308.3121
at G =  4.0      1731.2340
at G =  6.0      1154.1560
at G =  8.0       865.6170
at G = 12.0       577.0780
```

A unit of geometric mean is worth eight times as much in pitch at `G = 1.5` as at `G = 12`. The same compression, from the same cause: a logarithm read against a linear coordinate.

5.8c Compounds have no pitch class

Pitch class in this framework is defined by the odd-part class rule R1 [Curzi, 2026] as the odd part of the **index**. A compound has no index — it has a geometric mean, which is generally irrational — so **a compound value has no pitch class, and the six-class structure of the mood-and-trait paper [Curzi, 2026] does not apply to it.**

This is not a technicality, and the folded values show why. Take the four states of the carried subject's chord and fold both readings of each into the standing 220–440 Hz band:

```
n     essence f    folded (class)      compound k^2 n^4    folded
-------------------------------------------------------------------------
 3     4.350832    278.453222 (3)        18.929736       302.875769
 5    12.085643    386.740586 (5)       146.062775       292.125549
 7    23.687861    379.005775 (7)       561.114755       280.557377
10    48.342573    386.740586 (5)      2337.004393       292.125549
```

Not one of the folded compound values is any of the six class values of the ladder — 247.514, 278.453, 313.260, 379.006, 386.741, 233.978 — and states 5 and 10, which collide in class under Rule 1, also collide under Rule 3, but at a different pitch. The compound reading and the class reading are two different structures over the same states, and neither is a refinement of the other.

The consequence for practice is a prohibition and it is P7 in the specification: **a compound value may not be assigned a pitch class, and a class-based argument may not be run on one.** the dissonance paper's [Curzi, 2026] consonance/redundancy identity is a statement about classes and therefore about Rule 1. It does not transfer.

5.9 The solve step, and why decomposition is not unique Now the harder half of the maxim.

Solve asks: given a compound value, what was it made of? The first part of the answer is clean and the second part is the section's main negative result.

What is recoverable exactly. From a published `f`,

```
G = ( f / k^2 )^(1/4)
```

and nothing about that step is lossy. Worked on the compound `{3x1, 5x1}`:

```
f = 52.5825988460  ->  G = 3.8729833462
G^2 = 15.0000000000    G^3 = 58.0947501931    G^4 = 225.0000000002
```

What is not recoverable. Everything else. The composition enters the rule only through `G`, and `G` enters only through the pair `(P, T)` where `P = prod n_i^(c_i)`. Both arrows lose information.

THEOREM 5.4 (the coagula map is many-to-one). Distinct multisets over `{1..12}` with the same count `T` and the same product `P` have identical `G` and identical `f_compound`. Enumerated exhaustively:

```
T    multisets    distinct f    many-to-one rate    in an ambiguous fibre
------------------------------------------------------------------------
1           12            12          1.0000             0         0.00%
2           78            59          1.3220            35        44.87%
3          364           194          1.8763           257        70.60%
4         1365           504          2.7083          1145        83.88%
5         4368          1120          3.9000          3967        90.82%
6        12376          2226          5.5597         11704        94.57%
------------------------------------------------------------------------
```

The enumeration is exhaustive rather than sampled, which is the same verification discipline the corpus applies to its combinatorial results elsewhere — a star-polygon decomposition traced over 779 cases with zero mismatches [Curzi, 2026v].

By four components, five component sets in six share their compound frequency with at least one other. By six components, nineteen in twenty do. The ambiguity is not a corner case; it is the generic case, and it grows with the size of the compound.

THEOREM 5.5 (the fibre spans counts as well). By Theorem 5.3, every multiplicity-scaling of a multiset lies in the same fibre, so the fibre over any achievable `G` is infinite once multiplicities are unbounded. Even truncated at `T <= 6` over the twelve-state ladder, the 18,563 available multisets produce only 3,779 distinct values of `G` — a mean of 4.9121 multisets per value.

The worked instance, and it is the one this framework carries:

```
G = 6 is produced by SIXTY distinct multisets over {1..12} with T <= 6

T = 1   P = 6        1 multiset    {6}
T = 2   P = 36       3 multisets   {3,12} {4,9} {6,6}
T = 3   P = 216      5 multisets   {2,9,12} {3,6,12} {3,8,9} {4,6,9} {6,6,6}
T = 4   P = 1296    10 multisets
```

```
T = 5   P = 7776    15 multisets
T = 6   P = 46656   26 multisets
                    ---
                    60
```

All sixty return `f_compound = 302.8757693527` in convention-Hz. **One of them is the carried dyad's member `A`, whose chord is `{4, 6, 9}`.** A reading that published only `A`'s compound value would have said precisely as much about `A` as about a person holding `{6,6,6}`, which is to say nothing that distinguishes them.

CONSEQUENCE 5.6 — the publication rule. A compound frequency may not be published as a description of a composition. **What must be published is the composition, with `G` and `T` derived from it — never the reverse.** The direction of the rule is the whole of it: coagula is a function and may be run; solve is a search and its output is a set.

5.10 What the solve step CAN do

The negative result is large, so the residue is worth stating exactly, since it is not zero.

Solve can falsify a proposed decomposition. If a candidate component set does not produce the published `G`, it is not the decomposition. Over the 1365 four-component multisets, a published `G` eliminates 1361 of them on average and leaves a fibre of about 2.7.

Solve can bound the components. `G` is a mean, so `min(n_i) <= G <= max(n_i)`. A compound with `G = 6` cannot consist entirely of components below index 6, and cannot consist entirely of components above it.

Solve can recover `T` when `T` is small and the arithmetic is tight. `P = G^T` must be an integer for a multiset over the integers. For `G = 3.8729833462`, `G^2 = 15` exactly and `G^5 = 5977.046...` is not an integer, so `T = 5` is excluded outright while `T = 2, 4, 6, ...` are not. **The integrality of `G^T` is a genuine constraint and it is free.**

Solve can be made unique by declaration. If the analyst declares `T` in advance and the fibre at that `T` is a singleton, the decomposition is determined. At `T = 2` over the twelve-state ladder, 55.13% of multisets are alone in their fibre. That is the honest use of *solve*: **not inversion, but hypothesis testing against a declared count.**

5.10b The information accounting of the solve step

The non-uniqueness has an exact size, and stating it in bits is more useful than stating it in counts, because bits are what the rest of the series prices information in.

A decomposition drawn from the twelve-state ladder at count `T` is one of `C(11+T, T)` multisets. The compound value identifies one of the distinct `G` values available at that count. The difference is what *coagula* threw away and *solve* cannot recover:

T	multisets	log2	distinct G	log2	BITS LOST	fraction lost
2	78	6.2854	59	5.8826	0.4028	6.41%
3	364	8.5078	194	7.5999	0.9079	10.67%
4	1365	10.4147	504	8.9773	1.4374	13.80%
5	4368	12.0928	1120	10.1293	1.9635	16.24%
6	12376	13.5953	2226	11.1202	2.4750	18.21%

The accounting is Shannon's and nothing more exotic: the entropy of the pre-image set against the entropy of the observable [Shannon, 1948, Cover and Thomas, 2006, Hartley, 1928], and the irreversibility of a many-to-one map is the same object that makes erasure physically costly [Landauer, 1961]. **The bits reported here are Shannon bits, not thermodynamic ones**, and the framework has stated elsewhere what the two share and what they do not [Curzi, 2026z].

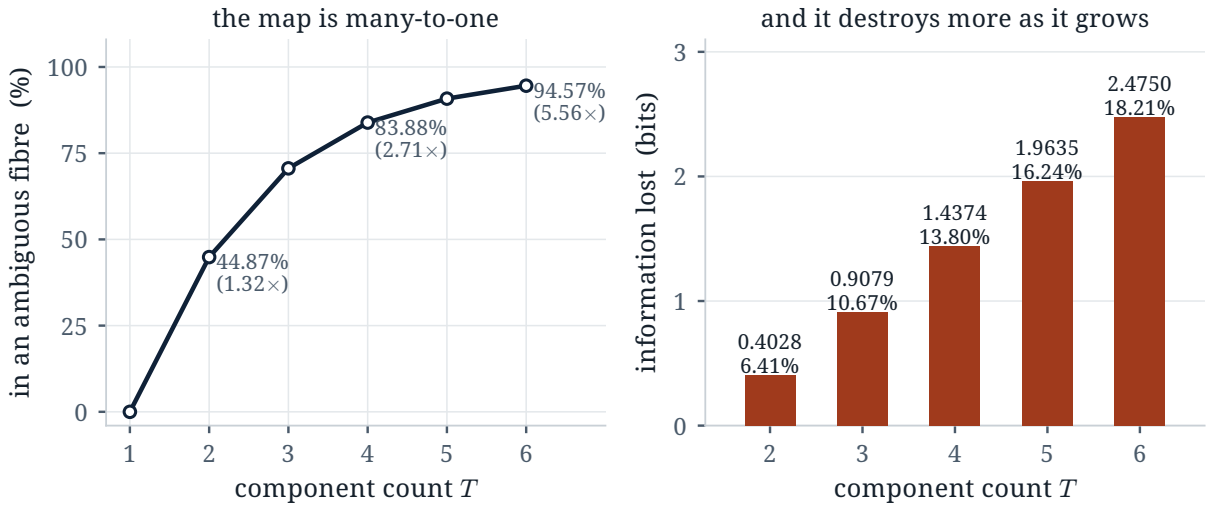


Figure 46.3: **Fixing is a function; dissolving is a search, and the search space grows.** Left: the fraction of component multisets over the twelve-state ladder that share their compound value with at least one other, with the many-to-one rate in parentheses. Right: the information the compound value discards, in bits and as a fraction of the decomposition's entropy. Both panels are exhaustive enumerations over all multisets at each count, not samples; computed from §5.9 and §5.10b; script `figs/src/CCP-046_fig03.py`.

Two readings, and they pull in opposite directions, which is why both are printed.

The loss is real and it grows. At two components the compound value keeps 93.59% of the information in the decomposition; at six it keeps 81.79%, and the trend is monotone. Compounds get harder to invert as they get larger, which is the reverse of the intuition that more constraints narrow a search.

The loss is also smaller than the fibre counts suggest. A many-to-one rate of 5.56 at `T = 6` sounds catastrophic and costs 2.475 bits. Both numbers are correct: the fibres are large, and the value still eliminates the overwhelming majority of candidates. **`G` is a weak identifier and a strong filter, and Procedure SOLV-1 uses it as the second and never as the first.**

This is the same accounting the genome-addressing paper [Curzi, 2026y] ran on the Yamanaka factors — a selection worth at most 37.985 bits acting on a 6.2-gigabit archive — and it is the same lesson at a much smaller scale. A short address into a large space narrows enormously and identifies nothing.

5.10c The maxim, as a concordance The two-word instruction is old enough and general enough to invite over-reading, so it gets the Tri-State treatment with roles stated (the framework's evidential-role convention [Curzi, 2026]).

```
CLAIM        Composition and decomposition are distinct operations with
             different formal characters, and a process that alternates
             them is doing two different kinds of work.

(+1) SOURCE  Series B CCP-017 §29.10, the closed form; and Theorems 20.4
             and 20.5 of this chapter.
MEASURED     f = k^2 G^4 reproduces 37 of 44 published composite values
             to better than 1 part in 10^12. The inverse map is
             many-to-one at rates from 1.32 to 5.56 over T = 2..6.

(-1) SOURCE  Solve et coagula, transmitted through the Latin alchemical
```

```
              corpus and restated in Ripley's Twelve Gates (1471), where
              solution is Gate 2 and congelation Gate 6, separated by
              three intervening gates.
RECORDED      That the two operations are named separately, ordered, and
              that the sequence insists on dissolving BEFORE fixing.

ROLES         (+1) carries the evidence and carries it alone.
              (−1) carries precedent only: the observation that the two
              were distinguished, by people with no closed form and no
              way to count a fibre.

C1  structural correspondence   PASS. Two named operations; two formally
                                distinct maps, one a function and one a
                                search.
C2  predictive power            PARTIAL. The formalism predicts that
                                decomposition is under-determined and
                                says by how many bits. The tradition
                                predicts nothing testable.
C3  practical utility           PASS. On the (+1) side it produces
                                Procedure SOLV-1's halting condition. On
                                the (−1) side it explains why the corpus
                                treats dissolution as the harder step.

STATUS        ESTABLISHED as a structural correspondence; SPECULATIVE as
              any claim about what the tradition understood, and marked
              speculative here in the sentence that states it.
```

The asymmetry the maxim encodes is real and the formalism measures it. **That the tradition arrived at the asymmetry is precedent, and precedent is not evidence.**

5.11 Honesty clauses, restated in place **(1) ESSENCE, NOT RESONANCE.** Every value here is a position on a logarithmic scale of surplus. Nothing vibrates, nothing resonates, and no molecule, tissue, or person is described by any number in this section.

(2) UNITS. Rule 3's output is convention-Hz and is dimensionally hertz-squared. Ratio operations only.

(3) `k^2` IS A POSTULATE. Inherited from `k`, badged, and it cancels from every ratio.

(4) NO CLINICAL CLAIM. The twelve states are not diagnoses; nothing here is an assessment.

(5) CONSTRUCTED CASES. All of them, labelled.

(6) THE COUNT ARGUMENT TRAVELS. A compound reading inherits the defensibility of the decomposition it was computed from. A contested decomposition yields a contested `G`.

(7) REPORT `D` WITH `delta`. Where this section's values feed a sustainability reading, the decay-blindness result R4 [Curzi, 2026] binds: `D` contains no `delta` and a `D`-only report cannot see decay.

The engineering

5.12 Worked compounds **All constructed, labelled as such**, drawn from the twelve-state ladder. Multiplicities are counts of components in the specification, not intensities — intensities are `w`, they are a different object, and the elicitation rule [Curzi, 2026] owns them.

```
multiset            T      P          G               f = k^2 G^4
-------------------------------------------------------------------------
{1x1, 2x1}          2          2   1.4142135624          0.9348017573
{1x3, 12x1}         4         12   1.8612097182          2.8044052718
{1x1, 12x1}         2         12   3.4641016151         33.6528632614
{2x2, 7x1}          3         28   3.0365889719         19.8702609489
{2x1,3x1,5x1,7x1}   4        210   3.8067540958         49.0770922562
{3x1, 5x1}          2         15   3.8729833462         52.5825988460
{3x1, 5x1, 7x1}     3        105   4.7176939803        115.7653513543
{4x1, 6x1, 9x1}     3        216   6.0000000000        302.8757693527
{6x2, 8x1}          3        288   6.6038544978        444.4772248248
{5x1, 10x1}         2         50   7.0710678119        584.2510982885
{9,10,11,12}        4      11880  10.4400868170       2776.3612190668
{11x1, 12x1}        2        132  11.4891252931       4071.9964546313
-------------------------------------------------------------------------
values are convention-Hz, dimensionally Hz^2
```

Three readings of this table.

The two compounds required by the three transfer rules [Curzi, 2026] are rows six and four. `{3x1, 5x1}` gives `G = 3.8730` and `f = 52.5826`; `{2x2, 7x1}` gives `G = 3.0366` and `f = 19.8703`. Both are reproduced here to ten places so that later sections can cite the precise values.

Row three against row two is the quantity-blindness theorem doing visible damage. `{1x1, 12x1}` and `{1x3, 12x1}` differ only in how many index-1 components are present, and they differ by 3.584963 octaves — because changing a multiplicity changes the *proportions*, which `G` does see. Contrast that with `{3x1,5x1}` against `{3x3,5x3}` in §5.6, where every multiplicity trebled and the value did not move at all. **Rule 3 is blind to scale and acutely sensitive to ratio, and a reader who does not hold both facts at once will misread every table in this section.**

Row eight is A's chord and it lands on a round number. `4 x 6 x 9 = 216 = 6^3`, so `G = 6` exactly. This is arithmetic coincidence and nothing else, and §5.9 already showed what it costs: sixty multisets share it.

5.13 PROCEDURE COAG-1 — running the fixing step

```
1  state the decomposition: the multiset of indices with multiplicities.
   Every index must survive the count argument of CCP-041 on its own.
2  compute T = sum(c_i) and P = prod n_i^(c_i). Publish both.
3  compute G = P^(1/T) to at least ten decimal places.
4  compute f = k^2 G^4. Label it convention-Hz.
5  if the value is to be rendered, fold by 2^j into the stated band and
   publish j. Folding is a ratio operation and is the ONLY safe route to
   an audible value (§5.7 line 3).
6  publish the decomposition alongside. A compound value without its
   decomposition is not reproducible and is not a description of anything
   (Consequence 20.6).
```

5.14 PROCEDURE SOLV-1 — running the dissolving step

```
1  from the published f, compute G = (f / k^2)^(1/4) exactly.
2  DECLARE T before searching. If T is not declared, stop: the fibre is
   infinite by Theorem 5.5 and the search has no boundary.
3  compute P = G^T. If P is not an integer to within the published
   precision, the declared T is excluded. Try another.
4  enumerate all multisets over the admissible index range with product P
   and count T. This is the FIBRE.
```

```
5  publish the fibre in full. Not a member of it. The fibre.
6  if a unique member is wanted, it must be selected by evidence from
   OUTSIDE the frequency -- the actual components, elicited. The frequency
   cannot select it and no amount of precision will make it.
```

Step 5 is the step analysts skip, and skipping it converts a set-valued result into a point-valued claim. A published decomposition chosen from a fibre without saying so is **the frequency-domain version of reporting one arm of an experiment.**

5.15 The fibre, computed, for a reader to check The carried dyad's member `A` holds `{4, 6, 9}`. Run SOLV-1 on its published compound value.

```
step 1   f = 302.8757693527  ->  G = 6.0000000000
step 2   declare T = 3
step 3   P = G^3 = 216, an integer. T = 3 admissible.
step 4   fibre = { {2,9,12}, {3,6,12}, {3,8,9}, {4,6,9}, {6,6,6} }
step 5   publish all five
step 6   {4,6,9} is selected by the ELICITATION, not by the frequency
```

And the same value under other declared counts, to show what step 2 is worth:

```
T = 1   P = 6        fibre size  1     {6}
T = 2   P = 36       fibre size  3
T = 3   P = 216      fibre size  5
T = 4   P = 1296     fibre size 10
T = 5   P = 7776     fibre size 15
T = 6   P = 46656    fibre size 26
-------------------------------------
total across T <= 6               60
```

Declaring `T = 3` reduces the search from sixty candidates to five. It does not reduce it to one. Gate 2 recorded this as the reason it handed *solve* to this section, and the handover is now discharged.

5.16 When a compound reading is the wrong reading Rule 3 is available whenever a set of components can be named, which means it is available far more often than it is appropriate. Four tests before using it.

The general form of the error is one the corpus has already isolated: a property that is necessary is not thereby sufficient, and a quantity trusted *because* it is well defined is the failure mode [Curzi, 2026]. Four tests, in order.

Is the object a composition or a relationship? Two people who agree about something are not a compound of each other. Rule 2 governs a relationship and returns hertz; Rule 3 governs a composition and returns convention-Hz. Finding 4.1 is the standing warning: the word *conjunction* covers both, and the values are not commensurable.

Are the multiplicities counts? If what is being counted is really an intensity — how strongly something is held — then it is `W`, it is elicited by magnitude estimation, and it does not belong in `G`. **Putting an intensity into a multiplicity is the compound-scale version of the state-ordering material's [Curzi, 2026] smuggling a rating in as a count.**

Is the total quantity part of the claim? If it is, Rule 3 cannot carry it (Theorem 5.3).

Will the value be compared to a Rule 1 or Rule 2 value? If so, stop. The units differ, and Corollary 5.2 guarantees the two rules never meet.

The specification

SPEC-5.1 — Inputs

```
the multiset: indices with multiplicities, each index justified by count
the ladder and its N, stated
the elicitation that produced the decomposition, cited
a declaration of whether multiplicities are counts (required) or
anything else (inadmissible)
```

SPEC-5.2 — Required outputs, coagula

```
T, P, G to >= 10 dp, f in convention-Hz to >= 10 dp
the decomposition, in full, printed with the value
the units label, on every occurrence of f
j and the band, if the value was folded
```

SPEC-5.3 — Required outputs, solve

```
G, recovered exactly
the DECLARED T, and the integrality test on G^T
the FIBRE, in full, with its size
the external evidence used to select a member, if one was selected
an explicit statement that the frequency did not select it
```

SPEC-5.4 — Prohibitions

```
P1  Rule 3's output may not be described as a generalisation of Rule 1
P2  Rule 3's output may not be entered into an acoustic calculation as Hz
P3  a compound value may not be published without its decomposition
P4  a fibre member may not be published as the decomposition
P5  a Rule 3 value may not be ratioed against a Rule 1 or Rule 2 value
P6  quantity claims may not be supported by a Rule 3 value (Thm 20.3)
P7  a compound value may not be assigned a pitch class (§5.8c)
```

SPEC-5.5 — Precision `G` to at least ten decimal places; `f` to at least ten. The integrality test in SOLV-1 step 3 is performed against the published precision of `G`, and a value published to fewer than ten places cannot support the test. Recompute from the full-precision constants, never from rounded prints (the full-precision rule).

SPEC-5.6 — Gates

```
G1  units labelled everywhere                                REQUIRED
G2  decomposition published with every compound value        REQUIRED
G3  fibre published with every solve result                  REQUIRED
G4  T declared before any search                             REQUIRED
G5  the non-generalisation statement carried wherever both
    rules appear in the same document                        REQUIRED
G6  all cases labelled constructed                           REQUIRED
```

The failure modes

FM-5.1 — Treating Rule 3 as the general case of Rule 1 *Signature:* "the essence law is the one-component case of the compound law."

Why it is wrong: Theorem 5.1. The one-component case of the compound law returns `f(n)^2`, not `f(n)`. Corollary 5.2: there is no index at which they agree. At `n = 12` the two differ by 6.121 octaves.

The fix: state the two rules as answering different questions, every time both appear.

FM-5.2 — The O2 error, in its general form *Signature:* a published composite value that no multiplicity of its stated components can produce, retained because it "came from the source."

Why it is wrong: §5.4. For a single-component species the rule permits exactly one value. A published number outside the rule's range is an erratum, and the scatter of the seven outliers across four orders of magnitude confirms data entry rather than a different law.

The fix: recompute every inherited composite before citing it, and publish the recomputation.

FM-5.3 — Reading a compound value as a description *Signature:* "their compound frequency is 302.88, which tells us they are composed of `{4,6,9}`."

Why it is wrong: Theorem 5.4 and §5.15. Sixty multisets give that value.

The fix: Consequence 5.6. Publish the composition; derive the value from it; never the reverse.

FM-5.4 — Solving without declaring `T` *Signature:* a fibre enumerated over all counts, or a single decomposition produced from `f` alone.

Why it is wrong: Theorem 5.5. Without a declared `T` the fibre is infinite, because every multiplicity-scaling is in it.

The fix: SOLV-1 step 2, which halts.

FM-5.5 — Putting intensities into multiplicities *Signature:* a component "held strongly" entered as `n x 3`.

Why it is wrong: multiplicities are counts of components in a specification. Intensities are `w`, they are elicited by magnitude estimation (the elicitation rule), and their scale status is different. Mixing them produces a `G` that is neither.

The fix: §5.16 test two. If the number came from an intensity elicitation it goes in `u`, not in `G`.

FM-5.6 — Using a compound value to support a claim about amount *Signature:* "the value is high, therefore there is a great deal of it."

Why it is wrong: Theorem 5.3. `{3x1,5x1}` and `{3x5,5x5}` are the identical value.

The fix: if amount is the claim, Rule 3 is not the instrument. Nothing in this framework measures amount.

FM-5.7 — Rendering a Rule 3 value as hertz *Signature:* 52.5826 played as a tone.

Why it is wrong: the value is dimensionally hertz-squared and its Hz label is a convention (§5.7).

The fix: fold first, publish `j`, and label the result as a re-anchored convention value. The rendering protocol [Curzi, 2026] carries this into the rendering protocol.

FM-5.8 — Precision theatre on a contested decomposition *Signature:* `G` to twelve places over a decomposition two experts disagree about.

Why it is wrong: honesty clause (6). The value inherits the defensibility of the count. Twelve places on a contested multiset is precision about the wrong thing.

The fix: publish the disagreement, compute both decompositions, and report the octave gap between them as the uncertainty.

FM-5.9 — Comparing compounds across ladders *Signature:* a `G` from a twelve-state ladder set beside a `G` from an eight-category scheme.

Why it is wrong: the indices mean different things, so the geometric means are not comparable, and `k^2` does not save it because the inputs differ rather than the constant.

The fix: the ladder and its `N` are published with every value (SPEC-5.1) and comparisons across ladders are refused.

FM-5.10 — Reading the maxim as symmetric *Signature:* "dissolve and fix are inverse operations, so the procedure reverses."

Why it is wrong: this section's central asymmetry. Coagula is a function with a closed form; solve is a search whose output is a set, and by four components 83.88% of inputs land in a fibre of size two or more.

The fix: say *fix, then search*, not *fix, then unfix*.

FM-5.11 — Assigning a pitch class to a compound *Signature:* `{3x1,5x1}` folded to 420.660791 Hz and reported as "close to class 5."

Why it is wrong: §5.8c. Pitch class is the odd part of an index and a compound has no index. The folded compound values of the carried subject's four states land on none of the ladder's six class values, and the resemblance of any one of them to a class value is arithmetic accident.

The fix: fold for rendering, never for classification, and carry P7 wherever a compound value is folded.

FM-5.12 — Inverting a fibre by choosing the most plausible member *Signature:* the fibre is enumerated, one member "obviously fits," and the others are dropped without a word.

Why it is wrong: plausibility is external evidence, and external evidence must be named. SOLV-1 step 6 requires the selection criterion to be published; a criterion applied silently is indistinguishable from a preference.

The fix: publish the fibre, publish the criterion, and publish which member the criterion selects. If no criterion exists, the result is the fibre and the report says so.

FM-5.13 — Presenting a compound reading as a description of a person *Signature:* a `G` computed over someone's held states and reported as what they are.

Why it is wrong: the standing prohibitions. The twelve states are not diagnoses, no reading is an assessment, and Theorem 5.4 means the value does not even identify the states uniquely.

The fix: the reading is of a decomposition, the decomposition is constructed, and the sentence that reports it says so.

Worked example

PART A — THE CARRIED SUBJECT, FIXED **Constructed case, labelled as such.** the carried subject [Curzi, 2026]'s subject holds three framings of one situation, spanning the states `{3, 5, 7, 10}`. Run COAG-1 on the chord, taking each held state as one component.

```
step 1   multiset {3x1, 5x1, 7x1, 10x1}; ladder N = 12; states from CCP-041
step 2   T = 4     P = 3 x 5 x 7 x 10 = 1050
step 3   G = 1050^(1/4) = 5.6924250976
step 4   f = k^2 * G^4 = 245.3854612812  (convention-Hz, dimensionally Hz^2)
step 5   folded: 245.3854612812 * 2^0 = 245.385461, j = 0, band 220-440
step 6   decomposition published above
```

And the three framings taken separately, which is the more informative reading:

```
framing    states      T   P     G              f (convention-Hz)
-------------------------------------------------------------------
threat     {3,5,7}     3   105   4.7176939803   115.7653513543
story      {3,5,10}    3   150   5.3132928459   186.2578308452
mirror     {5,7,10}    3   350   7.0472987321   576.4348833846
-------------------------------------------------------------------
whole      {3,5,7,10}  4  1050   5.6924250976   245.3854612812
```

The whole is not the sum, the mean, or the midpoint of the parts. `245.3855` sits between `threat` and `mirror` but not at their geometric centre, because `G` is weighted by count and the whole has four components where each framing has three. This is the arithmetic behind the self-as-chord material [Curzi, 2026]'s warning that a chord is not the average of its notes.

Now the comparison that matters, and it is a refusal. the carried subject [Curzi, 2026] also publishes the subject's internal surplus and the chord's width as pitch:

```
Rule 2 reading   mean u_int = 0.700844   F_int = 2.164390   f = 36.6686325914 Hz
Rule 3 reading   G = 5.6924250976                          f = 245.3854612812 conv-Hz
```

These two numbers describe the same person and may not be compared, ratioed, averaged, or plotted on one axis. One is hertz and one is hertz-squared; one is a function of how non-parallel the framings are and one is a function of which states are held, with no reference to how they overlap. A reader who computes `log2(245.3855 / 36.6686) = 2.742` has produced a number with no referent, and it is printed here only so that it is recognisable as the mistake it is.

PART B — THE CARRIED DYAD, DISSOLVED **Constructed case, labelled as such.** the carried dyad [Curzi, 2026], member `A = {4:0.7, 6:0.5, 9:0.4}`. The intensities are `w` and do not enter `G` (FM-5.5); the components are the states `{4, 6, 9}`.

```
COAG-1     T = 3    P = 216    G = 6.0000000000    f = 302.8757693527 conv-Hz
```

Run SOLV-1 against that value, as an independent party would who had only the number.

```
step 1   G = (302.8757693527 / 0.233700439315386)^(1/4) = 6.0000000000
step 2   declare T = 3
step 3   P = 216, integer. Admissible.
step 4   fibre:
            {2, 9, 12}   {3, 6, 12}   {3, 8, 9}   {4, 6, 9}   {6, 6, 6}
step 5   fibre size 5, published in full
step 6   no external evidence supplied -> NO MEMBER SELECTED
```

Read the five candidates against the twelve-state ladder [Curzi, 2026] and the extent of the loss becomes concrete. `{4, 6, 9}` is situated feeling, temporal feeling, normative feeling. `{6, 6, 6}` is one state held three times. `{2, 9, 12}` reaches from valenced activation to committed feeling. These are not variations on a theme. **They are five different people, and Rule 3 returns one number for all of them.**

The pitch-class reading, which Rule 3 also cannot see, separates them immediately:

```
{4, 6, 9}    classes {1, 3, 9}    three classes
{2, 9, 12}   classes {1, 9, 3}    three classes
{3, 6, 12}   classes {3}          ONE class -- a single note played three times
{3, 8, 9}    classes {3, 1, 9}    three classes
{6, 6, 6}    classes {3}          ONE class, and one state
```

Two of the five are entirely consonant and therefore entirely redundant — the consonance-redundancy identity R2 [Curzi, 2026] — and the compound value cannot distinguish them from the three that are not. That is the strongest practical statement of the non-uniqueness in this section: **the solve step is blind to exactly the structure the dissonance paper of this series was written to expose.**

PART C — THE CONJUNCTION, PRICED IN BOTH NOTATIONS Gate 4 of Section 4 halted a schedule because *conjunction* is ambiguous. With Rule 3 now owned, the halt can be resolved by declaration rather than by choice.

```
reading (i)  RULE 2, the relational join
             u = 0.068734665264634    f_rel = 1.4907987150 Hz
             F = 0.563084802746       D = +4.338420   delta = 0.09
             Q_inf = -5.011940        DECAYS

reading (ii) RULE 3, the compositional join
             pooled multiset {4x2, 6x2, 9x2, 10x1}
             T = 7    P = 466560    G = 6.4542224391
             f = 405.5414451025 convention-Hz
```

Both are correct answers to different questions. Reading (i) says the pair's difference supports 0.563085 nats of surplus and that this is 4.338420 octaves short of what their cost ratio requires, at a loss rate of 0.09. Reading (ii) says the pooled specification has a count-weighted geometric mean of 6.4542224391 and nothing else at all — **it does not know that `A` and `B` are two people, does not know they overlap on three states out of four, and would return the identical value if one person held all seven components alone.**

```
THE DECLARATION RULE, and it discharges Gate 4:

joining PEOPLE                   -> Rule 2, hertz, u
joining COMPONENT SETS           -> Rule 3, convention-Hz, G
the two may both be reported
the two may never be compared
```

PART D — THE REPORT CARD

```
------------------------------------------------------------------------
COMPOUND READING -- carried subject, chord {3,5,7,10}
------------------------------------------------------------------------
decomposition       {3x1, 5x1, 7x1, 10x1}, ladder N = 12, CCP-041 indices
T                   4
```

```
P                   1050
G                   5.6924250976
f_compound          245.3854612812   CONVENTION-Hz, dimensionally Hz^2
fold                j = 0, band 220-440 Hz
solve fibre at T=4  { {3,5,7,10}, {5,5,6,7} }   size 2
T excluded          T = 5 (G^5 = 5977.046... not an integer)
-------------------------------------------------------------------------
NOT comparable to the Rule 2 reading of the same subject (36.6686 Hz).
NOT a measure of quantity (Theorem 5.3).
NOT a description of the subject (Consequence 20.6; fibre size 2).
ESSENCE frequencies, not physical resonances. Constructed case.
Not a diagnosis, not an assessment, not clinical.
-------------------------------------------------------------------------
```

The last line of the fibre row is the section in miniature. Even at the tightest available count, with the value known to ten decimal places, **the twelve-state ladder offers a second component set — `{5, 5, 6, 7}` — that fixes to exactly the same number.** *Coagula* ran in four steps and produced a value. *Solve* ran in six and produced a set of two, and the honest report prints the set.

Section 6 takes the other operation Section 4 handed forward: moving a state from one index to another. It prices the interval, states what the move costs, and establishes the constraint that bounds every claim in this paper — that nothing not already latent can be reached, at any strength, by any operation in this framework.

6 Re-indexing: what a move between indices costs and what it cannot reach

The word carries more freight than any other in the old corpus, and the freight is what this section unloads. Section 4 rejected Gate 12 — projection, the transmutation of a mass by a seed — on four separate grounds, and rejected Gate 10, exaltation, because *nobler* is not a coordinate. What survives after those rejections is small, exact, and worth having.

A state can be moved from one index to another. That is all transmutation is in this formalism, and this section holds its title to exactly that.

Moving a state from index `m` to index `n` means changing the number of distinguishable components required to specify it. It has a price in pitch that can be computed to four decimal places, a price in components that can be counted, and a price in the storage law that has to be declared because the formalism does not supply it. It has a direction, and the direction matters. And it has a hard outer boundary that no operation in this framework crosses at any strength, by any modality, ever: **nothing that is not already latent can be reached.**

That boundary is inherited, not invented here. The genome-addressing paper [Curzi, 2026y] stated it as CONSTRAINT L67 for the cellular architecture and called it the most valuable sentence in that volume. This section restates it at psychological scope and treats it as binding on everything in this paper. The reason it is the most valuable sentence, rather than the most limiting one, is worth having in advance: the property that makes an addressing architecture safe is the same property that makes it limited, and **a design claiming both unlimited reach and no modification would be incoherent.**

Three things before the arithmetic.

Re-indexing is not creation. the index of the state-ordering material [Curzi, 2026] is a count of components in a specification. To move a state up the ladder is to specify it more finely; the components that make the finer specification possible must already be available to be counted. **Yamanaka reprogramming activates latent capability; it does not create new capability**, and the same accounting applies here.

Re-indexing is not improvement. the more-is-not-better rule [Curzi, 2026] settled that more components is not better, and the dissonance paper [Curzi, 2026] settled that a fully consonant self is a fully redundant one. Nothing in this section licenses reading a higher index as a better state.

Nothing here is clinical. The twelve states are not diagnoses, no reading is an assessment of anyone's health, no operation below is a treatment, and every case is constructed and labelled as such. The frequencies are essence frequencies — positions on a logarithmic scale of surplus — and nothing is claimed to vibrate.

The law

6.1 What a re-indexing is

DEFINITION 6.1. A *re-indexing* is a change in the index assigned to a state, from `m` to `n`, where both are positions on a single fixed ladder and the change reflects a change in the number of distinguishable components required to specify the state.

Three things are excluded by that wording and each exclusion is load-bearing.

A change of ladder is not a re-indexing. If `N` changes, or if the categories are re-cut, the indices before and after are not comparable and no interval between them is meaningful. That operation is the minimum-frequency paper's [Curzi, 2026] lever L4 and Section 4's Gate 3, and it has its own price.

A change of register is not a re-indexing. Multiplying by `2^j` moves the published frequency and moves nothing else (the folding rule [Curzi, 2026]). Gate 8 rejected sublimation for exactly this confusion.

A re-labelling is not a re-indexing. If the same specification is given a different number because the analyst preferred it, the index moved and the state did not. The publish-the-decomposition rule [Curzi, 2026]'s rule binds: publish the decomposition or the frequency means nothing.

6.2 The interval crossed

THEOREM 6.2. The interval crossed by a re-indexing from `m` to `n` is a ratio of squares:

```
f(n) / f(m)  =  (k n^2) / (k m^2)  =  n^2 / m^2

and in the reverse direction        =  m^2 / n^2

in cents                            =  1200 * log2(n^2/m^2)  =  2400 * log2(n/m)
```

`k` cancels. **The price of a re-indexing does not depend on the chosen constant.** No verdict about a move depends on the reference pitch. This is the cancellation of the constant [Curzi, 2026] applied at index scope.

Two properties of the interval that practice depends on.

It is path-independent. The interval from `m` to `n` is a difference of logarithms, so a move through intermediate indices costs the same in pitch as the direct move. Going `3 -> 6 -> 12` costs `2400.0000 + 2400.0000 = 4800.0000` cents, and going `3 -> 12` directly costs `4800.0000`. The minimum-frequency paper [Curzi, 2026] §4.13's warning about path-dependent attribution applies to the cost side of a deficit and not to this.

It is antisymmetric. `2400 log2(n/m) = −2400 log2(m/n)`. The transition matrix [Curzi, 2026]'s transition matrix is the full table and is not re-derived here; the reverse of any entry is its negative.

The reference values this section uses most, from that matrix:

```
from -> to      ratio n^2/m^2     reverse m^2/n^2      cents
-------------------------------------------------------------------
 1 ->  2          4.000000          0.250000        +2400.0000
 3 ->  6          4.000000          0.250000        +2400.0000
 3 ->  4          1.777778          0.562500         +996.0900
 5 -> 10          4.000000          0.250000        +2400.0000
 9 -> 12          1.777778          0.562500         +996.0900
11 -> 12          1.190083          0.840278         +301.2741
 1 -> 12        144.000000          0.006944        +8603.9100
```

6.3 Two kinds of move, and the arithmetic that separates them The mood-and-trait paper [Curzi, 2026] owns the distinction and this section uses it. Two indices share a pitch class if and only if they share an odd part (the odd-part class rule R1 [Curzi, 2026]). Therefore:

```
RE-INDEXING WITHIN AN ORBIT      m and n share an odd part
                                 interval is a whole even number of octaves
                                 pitch class UNCHANGED
                                 the framework's term: INTENSIFICATION

RE-INDEXING ACROSS ORBITS        m and n have different odd parts
                                 pitch class CHANGES
                                 the framework's term: TRANSFORMATION
```

Counted over the ordered moves available on the twelve-state ladder:

```
ordered pairs (m != n)                    132
class-preserving (intensification)         20      15.1515%
class-changing   (transformation)         112      84.8485%
```

Transformation is the generic move and intensification is the rare, highly structured one. the orbit structure [Curzi, 2026] gives the six orbits; the operative fact for this section is which states have a double still on the ladder:

```
states with 2m <= 12 (can be intensified without leaving the ladder):  1,2,3,4,5,6
states with 2m >  12 (cannot):                                         7,8,9,10,11,12
```

Exactly half the ladder can be intensified and half cannot, and the half that cannot is an artefact of `N = 12` rather than a property of those states. The mood-and-trait paper [Curzi, 2026] named this truncation and this section prices it: for a state at index 7 or above, every available re-indexing is a transformation, because its orbit has no higher member inside the ladder.

6.4 CONSTRAINT T21 — the hard constraint

> **CONSTRAINT T21. A re-indexing can reach only what is already latent. No operation in this framework, of any kind, at any strength, by any modality, can produce a component that the decomposition does not already contain.**

This is CONSTRAINT L67 of the genome-addressing paper [Curzi, 2026y], restated at psychological scope, and it requires no new assumption. **An index is a count of components. A count cannot be raised by counting harder.** To move a state from index 5 to index 6 is to add a projected duration to a specification that previously had none, and either the material for that component is available to the person or it is not. The operation addresses; it does not import.

The constraint does four kinds of work, and they are the same four Series B named.

It bounds the claim space. Any statement of the form "this procedure could produce capability X" is false whenever X is not already available and reachable. That eliminates the entire class of claims that would make this work pseudoscience: nothing conferred, nothing imported, no capacity created.

It makes the model falsifiable in a specific way. If a demonstration showed a state specified by a component demonstrably absent from the person's repertoire, following an addressing-only intervention, this architecture is wrong and no adjustment of parameters repairs it.

It generates a positive prediction. If re-indexing reaches only what is latent, the reachable set is bounded and enumerable in principle. §6.6 computes it for a constructed repertoire.

It supplies the audit rule. Every proposed move is checked against the repertoire before it is priced, and Procedure TR-1 puts that check before the arithmetic rather than after.

The three works cited immediately below are named in full and carry no citation keys, because they are absent from the bibliographic database governing this corpus. That is a gap in the bibliography, not a claim that the works are unimportant; it is recorded here and again in §10.

The `(+1)` evidence, cited exactly and with roles stated (the framework's evidential-role convention [Curzi, 2026]):

Gurdon (*J. Embryol. Exp. Morphol.* 10:622-640, 1962) transplanted nuclei from differentiated intestinal epithelium of feeding *Xenopus* tadpoles into enucleated eggs and obtained normal feeding tadpoles. Differentiation is not subtraction: **the nucleus of a differentiated cell retains the full information required to build the organism.**

Takahashi & Yamanaka (*Cell* 126:663-676, 2006) returned mouse fibroblasts to a pluripotent state by forced expression of four transcription factors; **Takahashi et al. (*Cell* 131:861-872, 2007)** extended it to human fibroblasts. The 2012 Nobel Prize in Physiology or Medicine went jointly to Gurdon and Yamanaka.

And the accounting that makes the point quantitative rather than rhetorical, recomputed from the genome-addressing paper [Curzi, 2026y]:

```
human sequence-specific transcription factors, approx.   1600
    (Lambert et al., Cell 172:650-665, 2018)

selecting 4 of 1600:  C = 2.720438e+11  ->  37.985 bits
the genome acted upon:                      6.200e+09 bits
ratio                                       1.6322e+08 : 1
```

The instruction is one hundred and sixty-three million times smaller than the thing it reorganises, and **that is what an address is.** An activation signal is not the content it activates. The correspondence drawn here is structural — latency-then-expression as a shared architecture — and the `(-1)` column, the tradition's own insistence that the material must already contain what the work brings out, carries precedent and no evidence.

The extension of this accounting to psychological states is this work's, badged as a thesis, and it is not established by the cited work. Gurdon and Yamanaka measured cells. What they license here is the *form* of the constraint, not a claim about minds.

6.5 Latency, and the laundering it invites

"Latent" is doing heavy work and it is the word most easily abused, so its failure is named before its use. The genome-addressing paper [Curzi, 2026y] named the chain:

```
the repertoire contains latent capability
  -> therefore capability X is latent
    -> therefore X is available
```

Each arrow smuggles. The first step is a statement about bulk; the second names a specific item without evidence that the item is in the bulk; the third converts possible into reachable.

"Latent" means present in the repertoire and addressable, not merely conceivable. A component that a person could in principle acquire, given years and circumstances, is not latent in the sense this constraint uses.

The failure has a name outside this framework: a construct defined loosely enough that any observation confirms it makes no risky prediction, which is the standing objection to soft psychological theorising [Meehl, 1967, 1990], and the remedy is a criterion stated in advance that can refuse [Popper, 1959, Platt, 1964, Lakatos, 1970].

Operationally, and this is the rule Procedure TR-1 enforces: **a component counts as latent only if it has been observed in the repertoire at least once**, in some state, on some occasion. That is a strict criterion, it will exclude things that are in fact latent, and the exclusion is the correct direction of error for a constraint whose purpose is to bound a claim.

6.6 The reachable set

Given a repertoire, the reachable set is the set of chords the states in it can form. It is finite and it is countable.

Constructed example, labelled as such. A repertoire containing components sufficient to specify states `{3, 4, 5, 6, 7, 9, 10}` and not states `{1, 2, 8, 11, 12}`:

```
full chord space over 12 states       2^12 − 1  =  4095 chords
reachable from this repertoire        2^7  − 1  =   127 chords
fraction of the space                                3.101343%

four-note chords over 12 states       C(12,4)   =   495
four-note chords reachable            C(7,4)    =    35
fraction                                             7.0707%
```

Ninety-seven per cent of the nominal chord space is unreachable for this constructed repertoire. No operation in this paper changes that. Widening the repertoire is not an operation in this paper; it is what the genomic papers [Curzi, 2026] examine, under the open problem, with no established mechanism attached.

A second constructed repertoire, narrower, to show the scaling: with only five states available, the four-note reachable set falls from 495 to 5 — **a hundred-fold reduction from a repertoire loss of seven states.** The reachable set collapses combinatorially, and a reading that reports a chord without reporting the repertoire it was drawn from has reported the numerator of a fraction.

6.7 What re-indexing does NOT move — and this is this section's central finding

Here is the result that most constrains what transmutation can be claimed to do, and it was not designed in. It falls out of the definition of `u`.

THEOREM 6.3. `u` is invariant under any re-indexing that does not change the shared support of the vectors involved. The interaction parameter is `u = 1 − cos^2 theta` computed on unit vectors over the twelve coordinates; the index *labels* never enter the computation. **Re-indexing changes `u` only through the overlap it destroys at `m` and the overlap it creates at `n`.**

Computed on the carried subject [Curzi, 2026]'s subject, **constructed and labelled as such**. Take the `threat` framing `{3:0.8, 5:0.4, 7:0.2}` and move its index-3 component four different ways, holding `story` and `mirror` fixed:

```
move        lattice interval    mean u_int       F_int       chord width
──────────────────────────────────────────────────────────────────────────
baseline           --           0.700844041320  2.164390    36.6686325914 Hz
3 ->  6      +2400.0000 cents   0.870156210632  2.358182    54.0282536493 Hz
3 ->  4       +996.0900 cents   0.870156210632  2.358182    54.0282536493 Hz
3 -> 12      +4800.0000 cents   0.870156210632  2.358182    54.0282536493 Hz
3 ->  1      −3803.9100 cents   0.870156210632  2.358182    54.0282536493 Hz
──────────────────────────────────────────────────────────────────────────
```

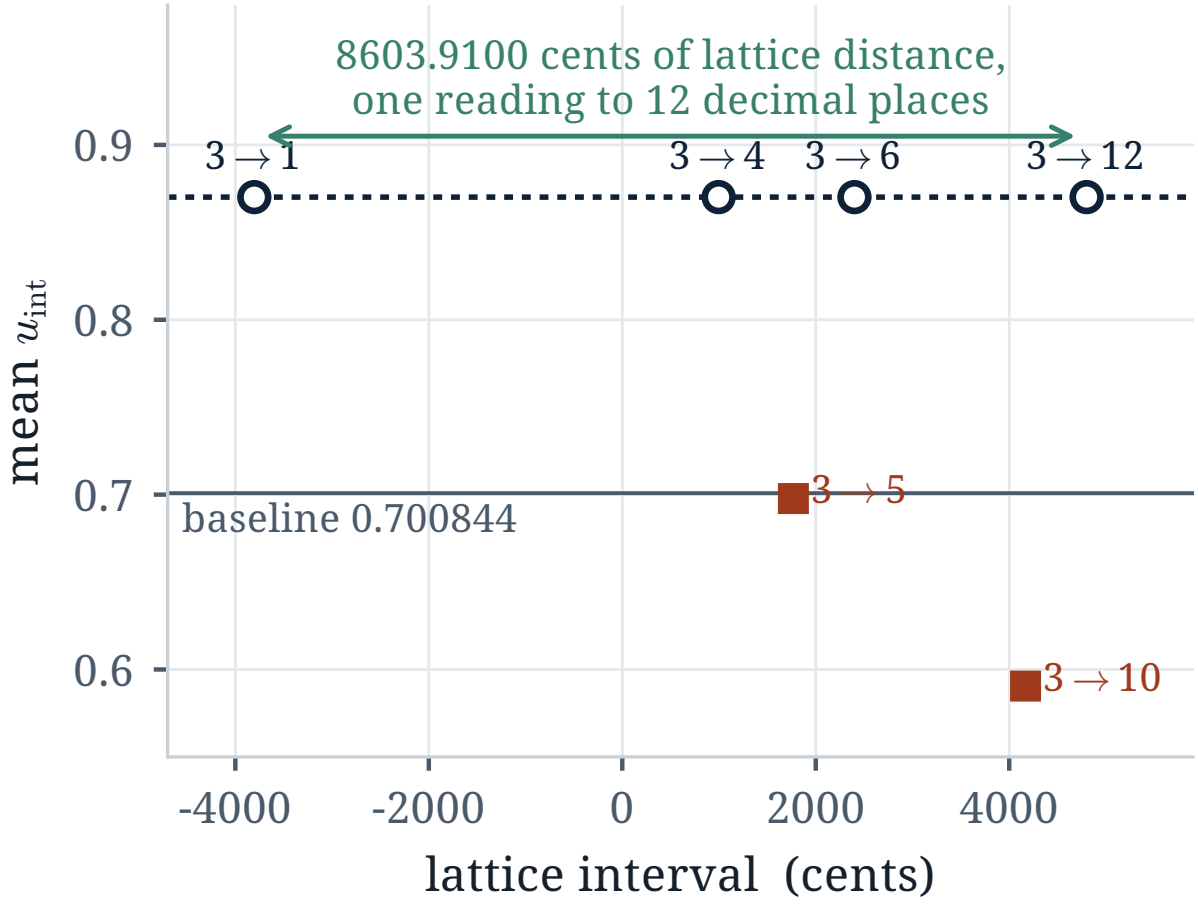


Figure 46.4: **Four moves spanning 8603.9100 cents return one relationship reading between them.** Mean u_{int} against lattice distance for six constructed re-indexings of one framing. The four open markers agree to twelve decimal places because all four vacate the same occupied coordinate and land on unoccupied ones; the two filled markers differ because they change the overlap. Computed from the vectors of §6's worked example; script `figs/src/CCP-046_fig04.py`.

Four moves spanning 8603.9100 cents of lattice distance return the identical relationship reading to twelve decimal places — one of them upward by four octaves, one of them downward by more than three. They are identical because all four vacate coordinate 3, which `story` occupies, and all four land on a coordinate that neither `story` nor `mirror` occupies. The relationship arithmetic sees a lost overlap and nothing else.

Two moves that are not identical, for contrast:

```
3 -> 10  (a coordinate `story` and `mirror` BOTH occupy)
         mean u_int = 0.590503185401   F_int = 2.014308   f = 27.1603 Hz

3 ->  5  (merged onto the framing's own state 5)
         mean u_int = 0.697749790607   F_int = 2.160474   f = 36.3826 Hz
```

A caution belongs here rather than in a footnote. The relationship readings above are computed from declared vectors on a stipulated reference set; **the reflexive specification the framework rests on is not identified from observational data** [Manski, 1993], and influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011]. Every relationship number in this paper is therefore a modelling position rather than

an identified empirical result, and the remedy adopted is stipulation of the reference set [Brock and Durlauf, 2001].

FINDING 6.1. The lattice and the interaction parameter measure different things and neither refines the other. A re-indexing that is an intensification on the lattice and a re-indexing that is a transformation across orbits can produce numerically identical relationship readings; and two moves of nearly identical lattice distance can produce very different ones. **Transmutation moves the index; it moves the relationship only where it moves the overlap. A section that claimed otherwise would be claiming that renaming a coordinate changes a geometry.**

The practical rule that follows, and Procedure TR-1 enforces it: **report the lattice reading and the relationship reading separately, and never infer either from the other.**

6.8 The cost, in three currencies

A re-indexing has three prices and the formalism supplies two of them.

Price one — components. Supplied exactly. Moving from `m` to `n` requires `n - m` additional distinguishable components in the specification, each of which must survive the state-ordering material's [Curzi, 2026] count argument on its own and must be latent by §6.5's criterion.

Price two — pitch. Supplied exactly. `2400 log2(n/m)` cents, path-independent, `k`-free.

Price three — maintenance. NOT supplied. It must be declared. A more finely specified state is a larger structure to hold, and holding it costs something per period. That cost is `c` in the storage law, and **nothing in this formalism derives `c` from `n`.** There is no theorem here and the framework does not pretend to one. What the framework supplies is the conversion, once `c` is declared:

```
dD/dc  =  2.885390 / eta

eta = 0.30   ->   9.617967 octaves of deficit per unit of c
eta = 0.34   ->   8.486441
eta = 0.40   ->   7.213475
eta = 0.44   ->   6.557705
eta = 0.60   ->   4.808983
eta = 0.90   ->   3.205989
```

A declared maintenance cost is a real input and an undeclared one is a hidden liability, and a re-index priced only in pitch has been priced in the currency that does not appear in the sustainability condition.

6.9 Direction, and the asymmetry the notation hides

The transition matrix is antisymmetric, so de-indexing looks exactly as available as indexing. It is not, and Gate 1 flagged this in passing. It is stated properly here.

Indexing upward requires latent material and an operation. Downward movement in fact happens without an operation — it is `delta`, the loss rate, running unasked. A person who has modelled another mind cannot elect to un-model it; what removes the component is time, disuse, and forgetting.

```
upward    an OPERATION on a latent repertoire; costs components and c
downward  in the formalism, an operation; in fact, mostly delta
```

The symmetry of the matrix is a property of the notation, not a menu, and a procedure that schedules a de-indexing as though it were a move one performs has mistaken the arithmetic for the world. Where deliberate de-indexing is genuinely available — simplifying a stated specification, dropping a distinction from a working vocabulary — it is available at the level of the *specification* and not at the level of what a person can already distinguish.

6.10 Honesty clauses, restated in place **(1) ESSENCE, NOT RESONANCE.** Every frequency here is a position on a logarithmic scale of surplus. Nothing vibrates. No person, tissue, or molecule resonates at any value in this section.

(2) NO CLINICAL CLAIM. The twelve states are not diagnoses, no reading is an assessment, no operation is a treatment, and nothing here describes or prescribes for any person's health.

(3) THE INDEX INHERITS THE COUNT. A contested decomposition yields a contested frequency, and a re-indexing between two contested indices yields a contested interval.

(4) CONSTRAINT T21 IS BINDING on every operation in this paper, and its extension from cells to psychological states is this work's thesis, badged, and not established by the cited work.

(5) MORE IS NOT BETTER. the more-is-not-better rule [Curzi, 2026] and the dissonance paper [Curzi, 2026]. A higher index is a finer specification and nothing else.

(6) CONSTRUCTED CASES, all of them, labelled.

(7) REPORT `D` WITH `delta`. the decay-blindness result R4 [Curzi, 2026]: `D` contains no `delta`, so a `D`-only report cannot see decay. Both, every time.

(8) CONSENT IS STRUCTURAL. Nothing in this section licenses acting on anyone. The group paper [Curzi, 2026] owns intervention and its constraints, and this section adds no permissions whatever.

The engineering

6.11 The price list The single-rung price list, which is the ladder read as a purchase schedule:

```
step        cents          cumulative from n = 1
--------------------------------------------------
 1 ->  2   2400.0000        2400.0000
 2 ->  3   1403.9100        3803.9100
 3 ->  4    996.0900        4800.0000
 4 ->  5    772.6274        5572.6274
 5 ->  6    631.2826        6203.9100
 6 ->  7    533.7418        6737.6518
 7 ->  8    462.3482        7200.0000
 8 ->  9    407.8200        7607.8200
 9 -> 10    364.8074        7972.6274
10 -> 11    330.0085        8302.6359
11 -> 12    301.2741        8603.9100
--------------------------------------------------
total 1 -> 12 = 8603.9100 cents = 7.169925 octaves
```

The same one-component step is worth 2400 cents at the bottom of the ladder and 301.2741 cents at the top — **a factor of 7.97.** the transition matrix [Curzi, 2026] called this the square law's signature; here it is a price list, and it has a direct practical reading: **differentiation is cheap in pitch at low index and expensive at high index, so the same pitch gain buys progressively more components as you climb.**

Cost per component for the intensification moves, which is where the compression bites hardest:

```
move       components added    cents      cents per component
------------------------------------------------------------
 1 ->  2          1          2400.0000        2400.0000
 2 ->  4          2          2400.0000        1200.0000
 3 ->  6          3          2400.0000         800.0000
```

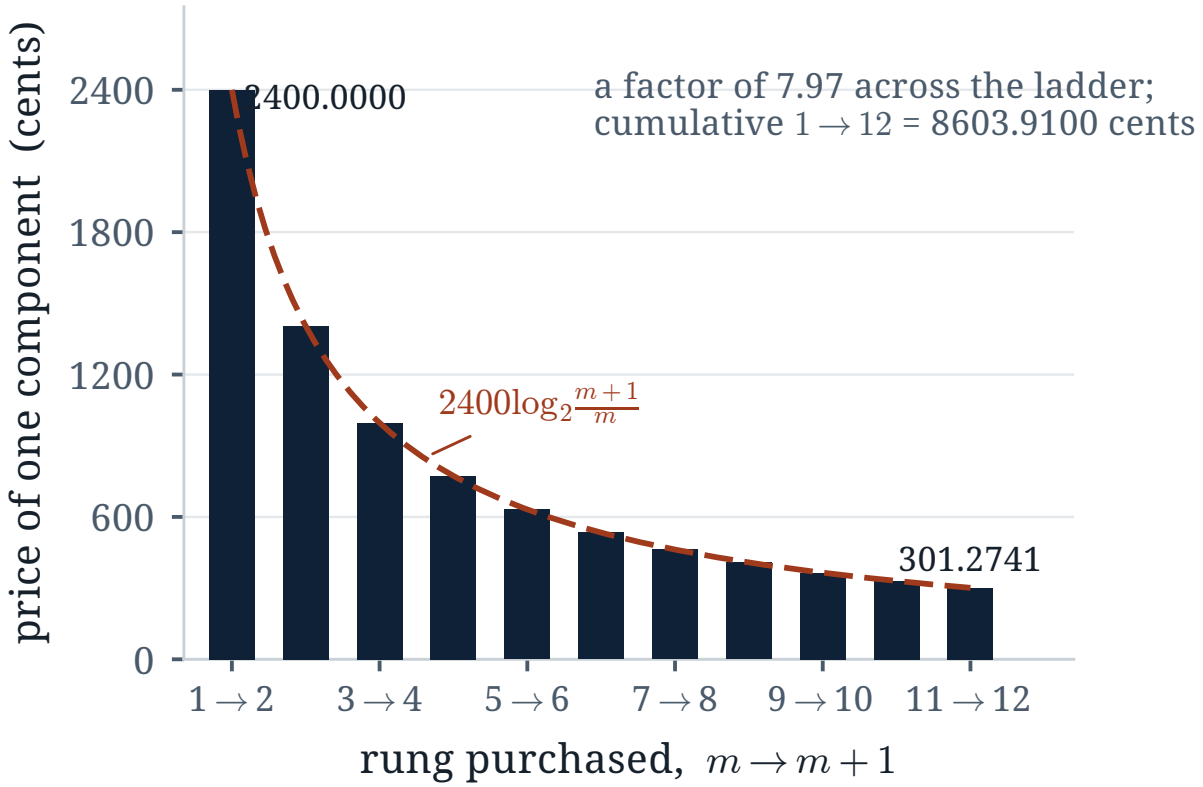


Figure 46.5: **The same one extra component costs 2400 cents at the bottom of the ladder and 301.2741 at the top.** The single-rung price list, with the closed form $2400\log_2((m+1)/m)$ overlaid. The compression is the square law's signature and it has a practical reading: the same pitch gain buys progressively more components as the index climbs. Computed from the price list of §6.11; script `figs/src/CCP-046_fig05.py`.

```
 4 ->  8          4          2400.0000         600.0000
 5 -> 10          5          2400.0000         480.0000
 6 -> 12          6          2400.0000         400.0000
------------------------------------------------------------------
```

Every one of the six doublings costs exactly 2400 cents, and they cost between one and six components to buy. An octave leap at the bottom of the ladder is a single component; the same leap at index 6 is six components. This is the arithmetic content of the mood-and-trait paper [Curzi, 2026]'s observation that intensification is scarce, and it is why the scarcity is structural rather than motivational.

6.12 PROCEDURE TR-1 — auditing and pricing a re-indexing

```
1  STATE the ladder and its N. State the starting index m and its
   decomposition, component by component.

2  STATE the target index n and its decomposition, component by
   component. The two decompositions must be comparable.

3  LATENCY CHECK, before any arithmetic. For every component present in
   n's decomposition and absent from m's, record whether it has been
   observed in the repertoire at least once, in some state, on some
   occasion (§6.5). If any required component fails, HALT.
   The move is REFUSED under CONSTRAINT T21 and the refusal is published
   with the component named.

4  CLASSIFY. Compute the odd parts of m and n. Same odd part ->
   INTENSIFICATION, whole even octaves, class preserved. Different ->
   TRANSFORMATION, class changes. Report which.

5  PRICE IN PITCH. 2400 * log2(n/m) cents. Report to four decimals with
   the direction stated.

6  PRICE IN COMPONENTS. n - m, listed by name, not by count alone.

7  DECLARE THE MAINTENANCE COST. State dc, state how it was obtained, and
   state that the formalism did not supply it. Convert with
   dD/dc = 2.885390/eta.

8  COMPUTE THE RELATIONSHIP EFFECT SEPARATELY. Recompute u from the
   vectors. Do NOT infer it from the lattice interval (Theorem 6.3).

9  READ THE DEFICIT. D and delta together, with Q_inf.

10 PUBLISH the lattice reading and the relationship reading side by side,
   labelled, with no ratio taken between them.
```

Step 3 is placed before step 4 deliberately. **A move that cannot be made does not need a price**, and computing the price first invites the price to argue with the constraint.

6.13 Re-indexing against the four levers The four-lever result [Curzi, 2026] fixed four levers on the deficit and proved there is no fifth. Re-indexing is not one of them, and the relationship is worth setting out because the two are constantly confused.

```
LEVER            what moves                re-indexing's relation to it
```

```
-------------------------------------------------------------------------
L1 raise u       the geometry             INDIRECT. Only when the move
                                          changes shared support (Thm 21.3)
L2 raise eta     the conversion           NONE
L3 cut c         the cost                 ADVERSE, usually. A finer
                                          specification costs more to hold
L4 raise N       the category set         NONE. Different operation
                                          (Sec 4 Gate 3)
-------------------------------------------------------------------------
```

Read straight down, that table says most re-indexing is deficit-neutral and cost-positive. A move that leaves the shared support alone does not touch `u`, does not touch `eta`, does not touch `N`, and adds maintenance. It buys pitch on the lattice and pays for it in the storage law.

The break-even, computed, so that the claim is not merely gloomy:

```
a re-index that raises c by dc must raise u enough to offset
2.885390 * dc / eta octaves

eta = 0.30, u0 = 0.068735, dc = 0.05  ->  u must reach 0.097688   reachable
eta = 0.30, u0 = 0.068735, dc = 0.10  ->  u must reach 0.131892   reachable
eta = 0.30, u0 = 0.692373, dc = 0.05  ->  u must reach 0.834429   reachable
eta = 0.30, u0 = 0.692373, dc = 0.10  ->  u must reach 1.002249   UNREACHABLE
eta = 0.44, u0 = 0.692373, dc = 0.10  ->  u must reach 0.892245   reachable
```

At a low `u` the break-even is easy and at a high `u` it can exceed 1, which is impossible. That is a boundary of the model's domain rather than a discovery about people, and it is stated as one. The mechanism is the strict concavity of the surplus functional [Curzi, 2026]: the marginal value of `u` falls as `u` rises, from `-31.739291` octaves per unit at `u = 0` to `-2.644941` at `u = 1`, a factor of 12. **A relation that is already differentiated cannot pay for expensive re-indexing out of further differentiation**, and the row where `u` must reach 1.002249 is the formalism refusing rather than warning.

6.14 The three counterfeits

A re-indexing can be claimed without occurring, in three ways, and each has a test.

Index inflation. Components are added to the specification that are not genuinely distinguishable from ones already present. The index rises and nothing has moved. This is pseudo-diversity (Series B's pseudo-diversity result) at index scope, and the test is the collision test R8 [Curzi, 2026]: components that always co-occur are one component.

Latency laundering. A component is asserted to be latent because it is conceivable. The test is §6.5's criterion — observed at least once, in some state, on some occasion — and it is deliberately strict.

Register change. The published frequency moves because `j` moved. The test is the folding rule [Curzi, 2026]: folding preserves every interval in cents exactly, so if the intervals to every other state are unchanged, nothing happened. Gate 8 is the historical version of this counterfeit and it was rejected in print.

All three produce a number that has moved and a state that has not, and **all three are caught by publishing the decomposition rather than the index.**

6.15 What re-indexing is genuinely good for

The section has spent most of its length on limits, so the residue is stated plainly, because it is not small.

It prices a change that would otherwise be described only in adjectives. A person's account of a situation acquiring a projected duration is a move from index 3 to index 6: 2400.0000 cents, three components, class preserved, an intensification. That is a great deal more than "it got deeper."

It separates two events that feel alike from inside. the mood-and-trait paper's whole argument [Curzi, 2026], now with a direction and a price attached.

It makes a refusal computable. Step 3 of TR-1 returns REFUSED with a named missing component, and a refusal with a name is actionable in a way that a vague sense of impossibility is not.

It exposes the cost side. Most re-indexing is cost-positive and deficit-neutral, which is invisible to any account that treats becoming more finely specified as self-evidently good.

The specification

SPEC-6.1 — Inputs

```
the ladder, its N, and its source (CCP-041 for the twelve states)
m and its full decomposition, component by component, named
n and its full decomposition, component by component, named
the repertoire, with the observation record for each component
eta, c, delta if a deficit reading is wanted
the vectors of all parties, if a relationship reading is wanted
```

SPEC-6.2 — Required outputs

```
the latency verdict, per required component, PASS or REFUSED
the classification: INTENSIFICATION or TRANSFORMATION, with odd parts
the interval in cents to four decimals, with direction
the components added or removed, by name
the declared maintenance cost dc, with its provenance and the statement
    that the formalism did not supply it
the deficit D and the loss rate delta, together, with Q_inf
the relationship reading, computed from vectors, reported SEPARATELY
```

SPEC-6.3 — Prohibitions

```
P1  a re-indexing may not be reported without its decompositions
P2  a relationship effect may not be inferred from a lattice interval,
    or the reverse (Theorem 6.3)
P3  a move may not be priced before the latency check (TR-1 step 3)
P4  "latent" may not mean "conceivable" (§6.5)
P5  a higher index may not be described as a better state
P6  a register change may not be reported as a re-indexing
P7  a maintenance cost may not be omitted; if unknown, say unknown
P8  no reading here may be presented as diagnostic, clinical, or as
    advice to any person about any relationship
```

SPEC-6.4 — Precision Intervals in cents to four decimals. Frequencies to at least six. `u` recomputed from the full-precision vectors, never from a rounded print (the full-precision rule). Deficits to six decimals with `delta` alongside.

SPEC-6.5 — Gates

```
G1  latency check performed and published, before pricing     REQUIRED
G2  classification by odd part, published                     REQUIRED
G3  lattice and relationship readings reported separately     REQUIRED
G4  maintenance cost declared or explicitly marked unknown    REQUIRED
G5  D never published without delta                           REQUIRED
G6  every case labelled constructed                           REQUIRED
G7  refusals published in full, with the component named      REQUIRED
```

The failure modes

FM-6.1 — Reading transmutation as creation *Signature:* a procedure described as producing a capability the person did not have.

Why it is wrong: CONSTRAINT T21. An address is a pointer; a pointer with no referent resolves to nothing. Yamanaka reprogramming activates latent capability and does not create new capability, and the four-factor selection is 37.985 bits acting on 6.2 gigabits.

The fix: the latency check, first, with the component named.

FM-6.2 — Latency laundering *Signature:* "the repertoire is rich, therefore X is latent, therefore X is available."

Why it is wrong: §6.5. Each arrow smuggles. Bulk is not specificity and conceivable is not reachable.

The fix: observed at least once, in some state, on some occasion — or REFUSED.

FM-6.3 — Inferring the relationship from the lattice *Signature:* "the move was four octaves, so it must have changed the relationship substantially."

Why it is wrong: Theorem 6.3 and Finding 6.1. Four constructed moves spanning 8603.9100 cents returned identical `u` to twelve decimal places.

The fix: recompute `u` from the vectors. Always.

FM-6.4 — Inferring the lattice from the relationship *Signature:* the reverse error — `u` moved, therefore a large index change occurred.

Why it is wrong: the same theorem. `u` responds to shared support and is blind to labels; a small change of overlap at a fixed index moves it without any re-indexing at all.

The fix: the same. Two readings, reported separately, no inference between them.

FM-6.5 — Index inflation *Signature:* the index rises after a re-description that added no distinguishable component.

Why it is wrong: §6.14. Components that always co-occur are one component (the relationship-frequency paper's [Curzi, 2026] collision test).

The fix: publish the decomposition and run the collision test on the added components.

FM-6.6 — Pricing only in pitch *Signature:* a move reported as "+2400 cents" with no maintenance cost.

Why it is wrong: §6.8. Pitch is not in the sustainability condition; `c` is. A move priced only in cents has been priced in the wrong currency, and at `eta = 0.30` a hidden `dc` of 0.10 is 0.961797 octaves of deficit.

The fix: SPEC-6.3 P7. Declare it or mark it unknown.

FM-6.7 — Treating de-indexing as an available operation *Signature:* a procedure that schedules "drop state 9" the way it schedules "add state 6."

Why it is wrong: §6.9. Downward movement in fact is mostly `delta`, running unasked. The antisymmetry of the transition matrix is notation.

The fix: distinguish simplifying a stated specification, which is available, from removing a distinction a person can already make, which is not.

FM-6.8 — Reading a higher index as a better state *Signature:* a re-indexing upward reported as progress.

Why it is wrong: the more-is-not-better rule [Curzi, 2026] and the consonance-redundancy identity R2 [Curzi, 2026]. Section 4 rejected Gate 10 for exactly this and the rejection binds here.

The fix: report the move, its class, and its price. Not a valuation.

FM-6.9 — Comparing intervals across ladders *Signature:* a `3 –> 6` on this framework's twelve-state ladder set beside a `3 –> 6` on a different scheme.

Why it is wrong: the indices count different things, so the interval is a ratio of incomparable squares.

The fix: the ladder and its `N` are published with every interval, and cross-ladder comparisons are refused.

FM-6.10 — Assuming an orbit continues past the ladder *Signature:* "state 7 intensifies to state 14."

Why it is wrong: `N = 12`. States 7 through 12 have no double inside the ladder, and the truncation is an artefact of the index length (the orbit structure [Curzi, 2026]). Extending the ladder is the minimum-frequency paper's [Curzi, 2026] lever L4 and is a different operation with a different justification burden.

The fix: report that the orbit is truncated, name `N`, and stop.

FM-6.11 — Precision theatre on a contested decomposition *Signature:* an interval to four decimals between two indices that two analysts assign differently.

Why it is wrong: honesty clause (3). The interval inherits both counts.

The fix: compute the interval under both decompositions and publish the gap between them as the uncertainty.

FM-6.12 — Using a priced move as an instruction *Signature:* "the move from 3 to 6 costs 2400 cents and three components, so make it."

Why it is wrong: the section prices moves and licenses nothing. The group paper [Curzi, 2026] owns intervention, its consent axiom, its halting conditions, and its reversibility classification, and none of them are relaxed here.

The fix: a price is not a recommendation, and nothing in this section is counselling, therapy, mediation, or clinical assessment.

Worked example

PART A — SIX MOVES ON THE CARRIED SUBJECT **Constructed case, labelled as such.** the carried subject [Curzi, 2026]'s subject, three framings of one situation:

```
threat = {3:0.8, 5:0.4, 7:0.2}     |w| = 0.916515138991
story  = {3:0.5, 5:0.7, 10:0.4}    |w| = 0.948683298051
mirror = {5:0.3, 7:0.8, 10:0.5}    |w| = 0.989949493661

baseline:  mean u_int = 0.700844041320   F_int = 2.164389635311
           N_eff = 8.709284   chord width 36.6686325914 Hz (folded 293.349061)
           states {3,5,7,10}   classes {3,5,7}
```

The raw L2 norms are published because the magnitude-reporting rule [Curzi, 2026] requires it: the free multiplier of a magnitude-estimated intensity must stay recoverable.

Now six constructed re-indexings of the `threat` framing's index-3 component, each run through TR-1 steps 4 through 8. In every case the latency check is stipulated to pass, and **the stipulation is the case's construction, not a finding.**

move	class	lattice interval	mean u_int	F_int	chord width (Hz)
base	3	—	0.700844041320	2.164390	36.6686325914
3 → 6	3	+2400.0000	0.870156210632	2.358182	54.0282536493
3 → 4	1	+996.0900	0.870156210632	2.358182	54.0282536493
3 → 12	3	+4800.0000	0.870156210632	2.358182	54.0282536493
3 → 1	1	−3803.9100	0.870156210632	2.358182	54.0282536493
3 → 10	5	+4168.7174	0.590503185401	2.014308	27.1603302345
3 → 5	5	+1768.7174	0.697749790607	2.160474	36.3825829120

And the classifications, which the relationship column cannot see:

3 → 6	odd parts 3, 3	INTENSIFICATION	class set {3,5,7} unchanged
3 → 12	odd parts 3, 3	INTENSIFICATION	class set {3,5,7} unchanged
3 → 4	odd parts 3, 1	TRANSFORMATION	class set {1,3,5,7}, four classes
3 → 1	odd parts 3, 1	TRANSFORMATION	class set {1,3,5,7}, four classes
3 → 10	odd parts 3, 5	TRANSFORMATION	class set {3,5,7} unchanged*
3 → 5	odd parts 3, 5	TRANSFORMATION	class set {3,5,7} unchanged*

* the framing's own class changed; the chord's class *set* did not, because `story` still holds state 3 and supplies class 3 to the chord.

Four readings of this table, and they are the section compressed.

The first four rows are identical in `u` and completely different on the lattice. Two are intensifications and two are transformations. One goes up four octaves and one goes down more than three. `u_int` agrees to twelve decimal places across all four. This is Theorem 6.3 and it is the strongest available statement of what the relationship arithmetic cannot see.

The fifth row is the only one that lowers the chord width. Moving to state 10 lands on a coordinate that both other framings occupy, so the move *creates* overlap where it destroyed some, and `u` falls: 0.700844 to 0.590503, and the chord width from 36.6686 Hz to 27.1603 Hz, a drop of 0.433045 octaves. **Re-indexing into an occupied coordinate is the only kind that can reduce a chord's width**, and it does so by making the framings more alike.

The sixth row barely moves at all. Merging the index-3 weight onto the framing's own state 5 moves `u` by `−0.003094` and the width by `−0.011298` octaves, because the framing's shared support with the other two is almost unchanged — it still overlaps `story` at 5 and `mirror` at 5 and 7. A 1768.7174-cent lattice move produced an eleven-thousandth of an octave of relationship change.

The pitch gains are worth the same in surplus terms across the identical group. `F_int` rose from 2.164390 to 2.358182, which is 0.193793 nats, which is `2.885390 * 0.193793 = 0.559168` octaves — precisely the chord width change from 36.6686325914 Hz to 54.0282536493 Hz. The conversion key holds exactly, as the conversion-key paper of Series B [Curzi, 2026l] requires.

PART B — A REFUSED MOVE **Constructed case, labelled as such.** The same subject, a proposed re-indexing of the `mirror` framing from state 7 (reflexive feeling — self as object) to state 11 (reciprocal feeling — the other modelling you).

TR-1 runs in order and stops at step 3.

```
step 1   ladder: CCP-041, N = 12. m = 7, decomposition: arousal, valence,
         intentional object, situation, cause, projected duration,
         self as object.

step 2   n = 11, decomposition: the seven above, plus another mind
         modelled, plus a standard applied, plus placement in a story,
         plus the other's model of the subject.

step 3   LATENCY CHECK on the four added components:
            another mind modelled          observed in this repertoire   PASS
            standard applied               observed in this repertoire   PASS
            placed in a story              observed in this repertoire   PASS
            the other's model of you       NOT OBSERVED                  FAIL

         HALT. The move is REFUSED under CONSTRAINT T21.
```

Steps 4 through 10 are not run, and the arithmetic that would have been produced is not published, because publishing it would let a refused move circulate as a priced one. What is published is the refusal, the named component, and the criterion applied:

```
REFUSED: 7 -> 11
missing component: the other's model of the subject
criterion: §6.5 -- observed in the repertoire at least once, in some
           state, on some occasion. Not met.
what would change the verdict: a single observation of the component,
           in any state, on any occasion.
```

A refusal that names its missing component and states what would reverse it is a scientific result. **A refusal that says the move is impossible is an opinion**, and SPEC-4.4 imposed the same requirement on its rejections.

PART C — THE CARRIED DYAD, READ AS RE-INDEXING, AND WHAT THE READING CANNOT ATTRIBUTE **Constructed case, labelled as such.** the carried dyad [Curzi, 2026] records that between phase 2a and phase 2b, `B` takes on two states `A` does not hold. That is a re-indexing at chord scope, and its relationship consequence is published:

```
PHASE 2a   u = 0.068734665264634   F = 0.563084802746   f =  1.4907987150 Hz
           eta = 0.30  c = 0.62  delta = 0.09
           D = +4.338420 octaves    Q_inf = -5.011940    DECAYS

PHASE 2b   u = 0.692372881355932   F = 2.153632742815   f = 35.8881768941 Hz
           eta = 0.44  c = 0.40  delta = 0.09
           D = -3.590989 octaves    Q_inf = +6.084427    ACCUMULATES
```

The path-free split [Curzi, 2026] supplies the path-free split of the 7.929409-octave swing:

```
from the cost ratio     2.885390 * (0.909091 - 2.066667)  =  -3.340058
from the surplus       -2.885390 * (2.153633 - 0.563085)  =  -4.589352
total                                                       -7.929409

the surplus side is 57.8776% of the swing
```

Now the limit of what this section may say about it, stated because the temptation is strong.

The carried case fixes `u` for phase 2b. **It does not fix which two states `B` took on, and this section does not invent them.** Theorem 6.3 is precisely why it cannot be inferred: many different pairs of added states produce the same shared-support change and therefore the same `u`, exactly as the four moves in Part A produced one reading between them. **The relationship frequency records that overlap changed. It does not record where.**

What can be said, and it is worth having:

The move was necessarily out of the shared support. `u` rose, so overlap fell as a fraction of the norms, so at least some of what `B` added lies where `A` holds nothing.

The move was necessarily an addition, not a re-labelling. A pure re-labelling of `B`'s existing states would leave the overlap pattern with `A` unchanged wherever `A` holds nothing, and Part A's first four rows show what that looks like: it can move `u`, but only by vacating a coordinate `A` occupies, which is a different event from taking on new states.

And `delta` did not move, and the deficit could not have seen it if it had. `D` is delta-blind (the decay-blindness result R4 [Curzi, 2026]). At 2b, dropping `delta` from 0.09 to 0.05 takes `Q_inf` from `+6.084427` to `+10.951968` and leaves `D` at `−3.590989`. Both numbers are reported here, as clause (7) requires.

PART D — THE REPORT CARD

```
------------------------------------------------------------------------
RE-INDEXING REPORT -- carried subject, threat framing, 3 -> 6
------------------------------------------------------------------------
ladder              CCP-041 twelve-state, N = 12
m                   3, directed feeling (arousal, valence, object)
n                   6, temporal feeling (+ situation, + cause,
                    + projected duration)
latency check       3 components required, 3 observed, PASS  [stipulated]
classification      odd parts 3 and 3 -> INTENSIFICATION
class              3 before, 3 after -- UNCHANGED
interval            +2400.0000 cents = +2 octaves exactly
ratio               f(6)/f(3) = 36/9 = 4.000000; reverse 9/36 = 0.250000
components          +3, named above
price per component 800.0000 cents
maintenance cost    NOT SUPPLIED BY THE FORMALISM. Declared dc unknown
                    for this constructed case; at eta = 0.30 each 0.10 of
                    undeclared c is 0.961797 octaves of hidden deficit.
------------------------------------------------------------------------
RELATIONSHIP READING, computed separately from the vectors
mean u_int          0.700844041320 -> 0.870156210632
F_int               2.164389635311 -> 2.358182352134   (+0.193793 nats)
chord width         36.6686325914 Hz -> 54.0282536493 Hz  (+0.559168 oct)
L2 norms            0.916515138991 / 0.948683298051 / 0.989949493661
NOTE                identical to the readings for 3->4, 3->12 and 3->1.
                    The lattice reading and this reading are independent
                    (Theorem 6.3) and neither was inferred from the other.
------------------------------------------------------------------------
ESSENCE frequencies, not physical resonances. Nothing vibrates.
Constructed case. Not a diagnosis, not an assessment, not clinical,
not advice. CONSTRAINT T21 binds: nothing not already latent was reached,
and nothing was created.
------------------------------------------------------------------------
```

What this section established, and the sentence it refuses to exceed Transmutation, in this formalism, is the movement of a state from one index to another. The interval crossed is a ratio of squares and is exact. The move is an intensification when the indices share an odd part and a transformation when they do not, and transformation is the generic case at 84.8485% of available moves. The price is three components on average at the doublings, between 301.2741 and 2400.0000 cents in pitch depending on where on the ladder it is bought, and a maintenance cost the formalism does not supply and the analyst must declare.

The relationship reading is a separate reading and may not be inferred from the lattice, because `u` is blind to index labels and sees only shared support — a result this section demonstrated with four constructed moves spanning 8603.9100 cents that returned one number between them.

And the boundary holds without exception. **Re-indexing addresses a repertoire; it does not extend one.** Nothing not already latent can be reached, at any strength, by any operation in this framework, and a procedure that claimed otherwise would have left the formalism behind. Transmutation is re-indexing, not creation. The section has said nothing else, and the restraint is not a modesty about the result — it is the result.

The refinement paper [Curzi, 2026] takes the iterative half of Section 4's Gate 11 and gives it its section: what happens when a bounded operation is run again and again on an accumulating store, why the answer is `|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|`, and why refinement arrives rather than ascends.

7 Results

The eight results below are the units other papers in this corpus cite. A citation to this paper should name one of them. Each carries its status: **DERIVED** where it follows from the stated definitions by algebra, **COMPUTED** where it is the output of exact arithmetic or exhaustive enumeration on stated inputs, **AUDITED** where it is the outcome of applying Audit Rule 4.1, and **POSTULATED** where it is asserted by this framework.

R1 — the audit distribution. Over thirteen entries: **seven MAPS, two SPLIT, one RENDERING, three REJECTED**, with five entries carrying a failure. The rejections are *exaltation* (a value ordering presented as a count), *projection* (failing all four conditions), and the planetary-metal ordering as an essence index (failing all four ladder-validity conditions). *Status:* AUDITED against four stated conditions on a list fixed before the results were known. **Not tested for inter-rater reproducibility** (§10, L-1).

R2 — Finding 4.2, the schedule result. Run as an ordered schedule on the carried dyad, the gates move the octave deficit by -5.966526 in total. **The ninth gate alone supplies** -5.305895 **octaves — 88.93 per cent — against** -0.660631**, or 11.07 per cent, from the eight before it.** The ninth gate is the only one acting on the numerator of the sustainability condition, and the framework's ordering rule requires numerator moves first [Curzi, 2026]. *Status:* COMPUTED on a constructed case under constructed step assignments (M-e). **This is the paper's sharpest finding and it is negative about the tradition's own order.**

R3 — the compound rule, derived. $f_{\text{compound}} = k^2 G^4$ with $G = \left(\prod_i n_i^{c_i}\right)^{1/T}$, obtained in four lines by substituting $f(n) = kn^2$ into the source recipe, the collapse turning on $\sum_i c_i/T = 1$. It reproduces **37 of 44 published composite values to better than one part in** 10^{12}. *Status:* DERIVED; the verification COMPUTED. k^2 is a **defined constant, exact**: 0.233700439315386, and it cancels from every ratio.

R4 — Corollary 5.2: the two rules agree at no index. $kn^2 = k^2 n^4$ requires $n = 1/\sqrt{k} = 1.438252434444351$, which is not an integer. **There is no index at which Rule 1 and Rule 3 coincide.** The gap is smallest at $n = 1$, where it is exactly $-\log_2 k = 1.048634$ octaves, and widens monotonically to **6.121291 octaves at** $n = 12$. *Status:* DERIVED. **The source manuscript quotes the** $n = 1$ **gap as 1.048946 octaves; that figure is wrong and the correction is stated in §10.**

R5 — Theorem 5.3, quantity blindness. G is invariant under multiplying every multiplicity by a common positive integer, so $\{3^{\times 1}, 5^{\times 1}\}$ and $\{3^{\times 5}, 5^{\times 5}\}$ return the identical $G = 3.872983346207$ and $f = 52.582598845962$. *Status:* DERIVED; verified at $m = 1, 2, 3, 5$. **Rule 3 cannot see how much of a compound there is**, which is the structural reason Gate 12 is rejected.

R6 — Theorem 5.4, the map is many-to-one. Exhaustive enumeration over the twelve-state ladder gives many-to-one rates of **1.3220, 1.8763, 2.7083, 3.9000 and 5.5597 at** $T = 2, 3, 4, 5, 6$, with 44.87, 70.60, 83.88, 90.82 and 94.57 per cent of multisets sharing their value with at least one other. $G = 6$ **alone is produced by sixty distinct multisets over** $\{1 \dots 12\}$ **at** $T \leq 6$**.** The information lost runs from 0.4028 bits at $T = 2$ to 2.4750 bits at $T = 6$, or 6.41 to 18.21 per cent. *Status:* COMPUTED by exhaustive enumeration, not sampling.

R7 — Theorem 6.3, re-indexing does not move the relationship except through overlap. $u = 1 - \cos^2\theta$ is computed on unit vectors over the coordinates; index *labels* never enter it. **Four constructed moves spanning 8603.9100 cents of lattice distance — one upward by four octaves, one downward by more than three — return** $u_{\text{int}} = 0.870156210632$ **identically to twelve decimal places.** Two moves of similar lattice distance that change the overlap return 0.590503185401 and 0.697749790607. *Status:* DERIVED; the instances COMPUTED. **The lattice reading and the relationship reading are independent and neither may be inferred from the other.**

R8 — CONSTRAINT T21. A re-indexing reaches only what is already latent: **an index is a count, and a count cannot be raised by counting harder.** The maintenance price of a higher index is *not* derivable from n — the formalism supplies the component price and the pitch price and **does not supply the third**, which must be declared. *Status:* the constraint is **POSTULATED** at psychological scope (P-d); its form is carried from the corpus's genome-addressing result [Curzi, 2026y]; the conversion $\partial D/\partial c = 2.885390/\eta$ is ARITHMETIC.

CR — the counter-result. The audit rejects three of thirteen entries and finds the tradition's own canonical sequence inverted, placing its single numerator-acting operation ninth of twelve. And **the compound rule it validates is many-to-one, so a published compound value cannot be inverted to a decomposition**: at four components five sets in six are ambiguous, and the carried case's own chord shares its value with fifty-nine others. *Coagula* is a function; *solve* is a search, and its honest output is a set.

8 Discussion

8.1 What follows

An audit is only informative if it can fail, and the failures are where the information is. Two thirds of the entries map, and that is the least interesting thing in the paper. The three rejections are specific: each names the condition that failed, the result that condemns it, the rescues a reader would attempt, and what would reverse the verdict. **A rejection that cannot say what would reverse it is an opinion**, and SPEC-4.4 makes the statement mandatory.

Two of the maps are more useful than the tradition's own presentation of them. *Cibation* — feeding in small measured quantities at correct intervals — is the least glamorous item on the list and prices at 9.617967 octaves of deficit per unit of cost at $\eta = 0.30$. Its content is concavity, not diligence: sixty-four occasions deliver 43.2 per cent more than one occasion from the identical budget. **The concavity has a warrant that requires no psychological assumption at all** [Peters and Gell-Mann, 2016], and that is the warrant the argument uses.

The unit discipline is the paper's most transferable output. Rule 2 returns hertz and Rule 3 returns convention-Hz, dimensionally hertz-squared. The naive logarithm of their ratio on the carried dyad is 8.087620, and **that number is meaningless**; it is printed so that a reader who computes it independently knows it has been considered and discarded.

8.2 What does not follow

Nothing here shows that any audited operation was ever effective. The audit establishes that a modern operation exists which an old term describes. **The direction of that sentence is the whole of the claim**, and it is the weakest one consistent with the material.

Nothing here is a verdict on the people who transmitted the sequence. Seven gates map, and one of them encodes a concavity result not stated formally for another five centuries. **Precedent is worth having and it is not evidence**, and the concordance rule (SPEC-4.5) forbids any entry from upgrading its status on the strength of precedent at any weight.

The rejection of an unbounded reading is not a rejection of iteration. The mapped half of Gate 11 is genuinely useful, and the corpus builds a paper on it: seven passes close 48.32 per cent of the gap at $\delta = 0.09$ and 6.79 per cent at $\delta = 0.01$, and closing 95 per cent in seven passes requires $\delta = 0.348164$ [Curzi, 2026]. **Repetition arrives; it does not ascend.**

And the compound rule's success against 44 values is not evidence for the essence law. The rule was obtained by substituting the essence law into a source recipe, so its agreement tests the recipe's internal consistency and the arithmetic of the collapse. **It does not test whether either $f(n) = kn^2$ or the attachment of hertz to k is true**, and no verification in this paper could.

8.3 Relation to the surrounding programme

Everything in this paper that is a *verdict* survives the rejection of postulates P-b and P-c, because k and k^2 cancel from every interval, every ratio and every deficit [Curzi, 2026k]. What does not survive is the absolute values — the hertz columns and the convention-Hz columns — and those are labelled at every appearance.

Where the wider corpus bounds what may be claimed here it does so twice. The structural work establishes that adding another entity of the same kind buys capability and not diversification [Curzi, 2026], which is the quantity-blindness theorem in another domain and the same reason Gate 12 fails. And the historical work sets the standard for handling attractive testimony by refusing its own best story as a fabricated data point [Curzi, 2026] — the standard this paper applies to the transmutation record, whose survivorship likelihood ratio is exactly 1.000000.

8.4 The reflexive exposure

The framework asserts that systems change when they are modelled, and then exempts its own parameters from that reflexivity [Lucas, 1976, Hacking, 1995]. This paper adds a narrower version of the same exposure: **an audit that publishes which of a tradition's operations "map" changes how the tradition is read**, and there is a documented mechanism by which such publication becomes self-confirming [Merton, 1948, MacKenzie and Millo, 2003]. Nothing here controls for it. The mitigation offered is the one the paper can actually supply — that the rejections are printed at the same length as the maps, and that the concordance rule forbids precedent from moving any verdict.

9 Verification

The standard is Occam's: **a postulate that can be dropped without the result changing has been dropped.** Every step carries one of four statuses: **ARITHMETIC** (a reader can redo it), **DEFINITION** (true by this paper's own construction, carrying no empirical weight), **CITED** (an external source supports it, named), or **POSTULATED** (asserted here, no external support).

9.1 The postulates

	postulate	status
P1	The Pythagorean whole tone is the frequency ratio $9{:}8$.	CITED [Barbour, 1951]
P2	φ is the positive root of $x^2 = x + 1$; hence $\varphi^2 = \varphi + 1$.	CITED [Euclid, 1956, Vajda, 1989]
P3	A dimensionless ratio of the form $(9/8)^2/\varphi^2$ **is a frequency in hertz**.	POSTULATED — no source
P4	The essence of a state with n components has frequency kn^2.	POSTULATED — no source
P5	The source recipe combines component frequencies as $\left(\prod_i f(n_i)^{c_i/T}\right)^2$.	**DEFINITION** — taken from the source as given
P6	Squaring is read as a change of structural level rather than a dimensional operation.	POSTULATED (P-c); the resulting label is a convention
P7	n, N, c_i, T are cardinalities, hence absolute-scaled.	DEFINITION, typology CITED [Stevens, 1946]
P8	A value ordering ("nobler") admits every monotone transformation and is not a count.	CITED [Stevens, 1946, Michell, 1997]
P9	Stored order obeys $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$, $\delta > 0$.	CITED as a *form* [Oppenheim et al., 1996, Elaydi, 2005]; the interpretation POSTULATED
P10	$N_{\text{eff}} = 1 + (N-1)u$ (the affine count axiom).	POSTULATED — design axiom [Curzi, 2026]; **not a Hill number**
P11	A re-indexing reaches only what is already latent, at psychological scope.	POSTULATED (P-d); the cellular form CITED via [Curzi, 2026y]

Dropped as unnecessary. No postulate about *which* twelve states the ladder holds is needed for R3, R5, R6 or R7: the compound rule, quantity blindness, the fibre structure and the re-indexing theorem hold for any index set and any N. No postulate about consonance is needed anywhere — §5.8b names two established measures of ratio complexity and **uses neither**. And **P3 and P4 are needed only for the absolute values**; dropping them leaves every verdict, every fibre count and every interval standing.

9.2 The compounding chain

step	consumes	produces	status
L1	P1	$(9/8)^2 = 81/64 = 1.265625$ exactly	ARITHMETIC
L2	P2	$\varphi^2 = \varphi + 1 = 2.618033988749895$	ARITHMETIC
L3	L1, L2	$k = 0.483425732988414$ — **a DEFINED constant, exact**	ARITHMETIC (exact)
L4	L3, **P3**	k carries the unit Hz	**POSTULATED**
L5	L4, **P4**	$f(n) = kn^2$ Hz	**POSTULATED**
L6	L3	$k^2 = 0.233700439315386$, exact	ARITHMETIC
L7	L5, P5	$\left(\prod (kn_i^2)^{c_i/T}\right)^2 = (kG^2)^2 = k^2G^4$ — **R3**	ARITHMETIC
L8	L7, **P6**	the value is labelled convention-Hz; ratio operations only	**POSTULATED**
L9	L7 at one component	$f_{\text{comp}} = k^2n^4 = f(n)^2$	ARITHMETIC
L10	L9, L5	$kn^2 = k^2n^4 \iff kn^2 = 1 \iff n = 1/\sqrt{k} = 1.438252434444351$ — **R4**	ARITHMETIC
L11	L10	not an integer, so the rules agree nowhere; gap $-\log_2 k = 1.048634$ oct at $n = 1$, 6.121291 at $n = 12$	ARITHMETIC
L12	L9 at $n = 8$	$k^2 \cdot 8^4 = 957.2369994358$, the only value pure oxygen admits at any multiplicity	ARITHMETIC
L13	L12 vs published 11.1246117975	no multiplicity produces it; the row is an erratum	ARITHMETIC
L14	definition of G	$c_i \to mc_i \Rightarrow G' = (P^m)^{1/(mT)} = G$ — **R5**	ARITHMETIC
L15	L14	$\{3, 5\}$ and $\{3^{\times 5}, 5^{\times 5}\}$ both give 3.872983346207	ARITHMETIC
L16	G depends only on (P, T)	equal $(P, T) \Rightarrow$ equal f — **R6**	ARITHMETIC
L17	L16, exhaustive enumeration	rates 1.3220 to 5.5597 over $T = 2 \ldots 6$; sixty multisets at $G = 6$	ARITHMETIC (enumerated)
L18	L17, counting	bits lost $=$ $\log_2(\text{multisets}) - \log_2(\text{distinct } G)$: 0.4028 to 2.4750	ARITHMETIC
L19	$u = 1 - \cos^2\theta$ on unit vectors	index labels do not enter u — **R7**	DEFINITION

step	consumes	produces	status
L20	L19 on the constructed subject	four moves over 8603.9100 cents give 0.870156210632 identically	ARITHMETIC
L21	P7	an index is a cardinality; a count cannot be raised by counting — **R8**	DEFINITION
L22	P8	“nobler” is not a coordinate; Gate 10 and Entry 13 fail condition (c) — part of **R1**	CITED + ARITHMETIC
L23	L5, folding by 2^j	intervals in cents preserved exactly; Gate 8 is RENDERING	ARITHMETIC (value); **POSTULATED** (meaning beyond pitch)
L24	P9, P10	$D = 2.885390(c/\eta - F)$, containing no δ	ARITHMETIC
L25	L24 on the schedule	step 9 supplies -5.305895 of -5.966526 — **R2**	ARITHMETIC

9.3 What this chain is and is not

Read the status column. Steps L1–L3, L6–L7 and L9–L25 are arithmetic, definition, or arithmetic on a cited form; L4, L5 and L8 are postulated. **A chain containing postulated steps is not a proof of its conclusion**, and this one is not offered as one.

What it is offered as is a separation, and it is unusually clean here. **The three postulated steps enter only where an absolute value is printed.** L10, L14, L16, L19 and L21 — the five steps that carry R4 through R8 — consume none of them: Corollary 5.2 is a statement about when $kn^2 = k^2n^4$, which is a statement about k as a number and not as a frequency; quantity blindness is a property of the geometric mean; the fibre structure is combinatorics on the integers; the re-indexing theorem is a statement about cosines; and CONSTRAINT T21 is a statement about counting. **A reader who rejects P3, P4 and P6 entirely keeps R4, R5, R6, R7 and R8 unchanged, and keeps every audit verdict, losing only the hertz and convention-Hz columns.**

Three checks a reader should run, each under a minute. $(9/8)^2 = 1.265625$ and $\varphi^2 = 2.618033988749895$, so $k = 0.483425732988414$ — that is L1 to L3 and it establishes the constant rather than asserting it. $1/\sqrt{0.483425732988414} = 1.438252434444351$, which is visibly not an integer — that is R4 in one keystroke. And $4{\times}6{\times}9 = 216 = 6^3 = 2{\times}9{\times}12 = 3{\times}6{\times}12 = 3{\times}8{\times}9 = 6{\times}6{\times}6$, which is R6’s fibre in five multiplications.

Two statuses that must not be misread. The exponent 2 at L5 and the exponent 4 at L7 are **DEFINED, not estimated**: no fit was performed and the standard procedure for testing a claimed scaling relation [Clauset et al., 2009] was not run, because there is nothing to run it on. And L7’s agreement with 37 of 44 published values is **a consistency check on the source’s own recipe**, not evidence for P3 or P4.

10 Limitations

L-1. The audit was performed once, by one analyst, and inter-rater reproducibility has not been tested. H1 is the paper's central methodological claim and **it is untested.** Condition (c) of Audit Rule 4.1 — is the transformation a function rather than a wish — requires a judgment, FM-4.1 records that a determined reader can find a parameter for almost any of the entries, and **nothing here demonstrates that two analysts applying the four conditions would return the same distribution of verdicts.** This is the largest gap in the paper.

L-2. The attribution and ordering of the source sequence are taken as given and are not verified here. The audit assumes the twelve gates are correctly attributed to a fifteenth-century text and correctly ordered in it. **No entry in this paper's bibliography establishes either.** The claim rests on the historical record as transmitted, the transmission path is stated in §4.5, and **a reader who disputes the attribution or the order disputes R1 and R2 with it.**

L-3. Four works are named in the text without citation keys, because they are absent from the bibliographic database governing this corpus. They are the 1471 source text and its 1652 printed edition (§4.5), and the three biological papers supporting the form of CONSTRAINT T21 (§6.4), which are given with journal, volume, page range and year in the running text. **This is a gap in the bibliography and it is disclosed rather than papered over**; the biological works in particular carry real weight in §6.4 and a reader is entitled to check them.

L-4. The schedule of §4's worked example is constructed twice over. Each gate is assigned "the smallest change the gate's own description licenses", and **those assignments are this paper's, not the tradition's.** A different reader could assign differently and obtain a different distribution of movement. **R2 is a result about this assignment on this constructed dyad**, and it is not a measurement of anything.

L-5. No alternative combination law was searched for. H3 concludes that seven published values are errata on the strength of two arguments: the residuals are scattered rather than systematic, and the pure-oxygen row admits exactly one rule-consistent value. **Neither argument rules out a different law that happens to agree with k^2G^4 on 37 rows**, and no search over alternative functional forms was conducted. The conclusion is the best available reading of the evidence and it is not a proof.

L-6. The verification of the compound rule tests the source's consistency, not the framework's truth. k^2G^4 is obtained by substituting $f(n) = kn^2$ into the source recipe. **Agreement with the source's own published values therefore cannot be evidence for the essence law**, and it is not offered as any.

L-7. The dimensional audit does not close in two places. $f(n) = kn^2$ requires k to carry $[T^{-1}]$ while k is a ratio of two dimensionless quantities (P-b); and k^2G^4 is dimensionally hertz-squared, so its "Hz" label is a convention and not a unit (P-c). **No citation in the bibliography repairs either.**

L-8. The constant is chosen and its fifteen figures are digits of a definition. $k = (9/8)^2/\varphi^2$ combines two documented objects into a product with no precedent, and the deflationary literature on golden-section claims is cited rather than avoided [Markowsky, 1992, Fischler, 1981, Livio, 2002].

L-9. Octave reduction beyond pitch is unwarranted. Octave equivalence is grounded for pitch [Shepard, 1982, Demany et al., 2013]; folding a non-acoustic quantity and reading meaning from where it lands is not.

L-10. Cents is a choice of unit and the two established measures of ratio complexity are not used. Euler's *gradus suavitatis* [Euler, 1739] and harmonic distance [Tenney, 2015] grade an interval by arithmetic complexity; this paper prices moves by $2400 \log_2(n/m)$, which is a distance in a logarithmic register and carries **no claim about consonance**. What a listener finds consonant depends on spectrum rather than ratio in any case [Sethares, 2005, Plomp and Levelt, 1965].

L-11. Neither effective count is a Hill number. $1 + (N - 1)u$ is a linear interpolation with no supporting literature. The Hill-number citations here [Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979] apply to $1/\mathrm{HHI}$ and $\exp(H)$ and to nothing else.

L-12. Conjoint measurement was not tested. The ratio-scale treatment of η, c, δ and the elicited intensities is provisional; Michell's objection stands unanswered [Michell, 1997, 1999] for every quantity here except the counts, and the remedies [Luce and Tukey, 1964, Krantz et al., 1971, Rasch, 1960] were not run. **Gate 10 and Entry 13 are rejected for exactly this fault, and the framework's own declared parameters are not audited against it.** That asymmetry is the most uncomfortable thing in the paper and it is stated rather than hidden.

L-13. Scalar reduction is a habit of this framework, inherited here rather than re-conceded. The minimum-frequency paper states the limitation once and in full: the framework replaces phenomena its own source literature describes by several jointly varying features with a single number, then proceeds arithmetically on the number [Curzi, 2026]. **That concession governs here.** The specific instance in this paper is G: a whole composition is reduced to one geometric mean, and §5.9 measures exactly what that costs — 0.4028 to 2.4750 bits, with sixty distinct compositions sharing one value. **The mitigation is Consequence 5.6, which requires the composition to be published and the value derived from it, never the reverse.** It is a mitigation and not an answer.

L-14. The reflexive specification is not identified. Stated in §3.5: the endogenous social effect in a reflexive specification cannot be recovered from observational data [Manski, 1993, Shalizi and Thomas, 2011]. The remedy adopted is stipulation of the reference set [Brock and Durlauf, 2001], **so every relationship reading here is a modelling position rather than an identified empirical result.**

L-15. The framework exempts its own parameters from its own reflexivity [Lucas, 1976, Hacking, 1995]. Not resolved.

L-16. The exponents are defined, not estimated, and the standard procedure for testing a scaling claim [Clauset et al., 2009] was not run because there is no data to run it on. **This paper makes no empirical scaling claim.**

L-17. The functional form of the surplus is underdetermined [Latané, 1981, Briscoe et al., 2006, Stumpf and Porter, 2012]. Nothing here discriminates among the candidates; the results that matter need only concavity and boundedness.

L-18. Latency is defined by an operational criterion with no instrument behind it. "Observed at least once, in some state, on some occasion" (§6.5) is deliberately strict and **will exclude components that are in fact latent.** It is also unmeasured: no observation record exists for any constructed case here, and the latency check in §6's worked example is *stipulated to pass*, which is the case's construction and not a finding.

L-19. No model comparison, no statistics, no data. Nothing was fitted, so there is no out-of-sample test, no matched null, no information criterion, no effect size and no p-value. Every case is constructed. Where the corpus has attempted a confirmatory fit elsewhere it failed its own declared baseline test, and that is the relevant precedent.

L-20. There is no conservation law [Mirowski, 1989], and validity is claimed only within a specified arrangement [Cartwright, 1999]. **L-21. Ergodicity is assumed and not argued** [Peters, 2019].

An arithmetic error found in the source and corrected here. The source manuscript states that Rule 1 and Rule 3 come closest at $n = 1$, "a gap of 1.048946 octaves". **That figure is wrong.** The gap at $n = 1$ is $\log_2(k/k^2) = -\log_2 k = 1.048634$ octaves, and the discrepancy is 3.1×10^{-4} octaves — small, but it is a value quoted to six decimal places, and a number quoted to six decimals owes the reader six correct ones. **This paper quotes 1.048634 throughout**, and the error is recorded here rather than silently repaired because a reader recomputing from the printed values of k and k^2 will find it.

A second, smaller discrepancy, left standing. The source prints k^2 as 0.2337004393 "to ten places" in one place and uses 0.233700439315386 in the computations. The two are consistent — the first is the second rounded — but the ten-place form is not sufficient to support the integrality test of Procedure SOLV-1 step 3 at large T, which SPEC-5.5 itself requires. **The specification and the printed constant are in tension**, and the resolution adopted here is that every computation uses the full-precision value.

11 Conclusion

What is established. Of thirteen entries in a fixed, pre-specified list from a dated historical sequence, **seven map onto operations on a closed parameter set, two are ambiguous between operations returning values in different units, one acts only on the publication register, and three fail** — two of them because a value ordering has been presented as a count, which the scale-type argument forbids. Run as an ordered schedule on a constructed dyad, the sequence places its **only numerator-acting operation ninth of twelve**, and that operation supplies **88.93 per cent of the total movement** while the eight before it supply 11.07 per cent; **the framework's own ordering rule condemns that order**, and the same intervention on the loss rate improves the destination under one ordering and damages it under the other.

The compound rule $f = k^2G^4$ is **derived by substitution rather than asserted**, and reproduces 37 of 44 published composite values to better than one part in 10^{12}; the seven that disagree are diagnosed as errata by a row the rule constrains completely. $kn^2 = k^2n^4$ requires $n = 1/\sqrt{k} = 1.438252434444351$, **so the two rules agree at no index whatever**, differing by 1.048634 octaves at $n = 1$ and 6.121291 at $n = 12$. G is **blind to quantity** and the map to it is **many-to-one at rates from 1.32 to 5.56**, with $G = 6$ produced by sixty distinct multisets — so a compound value may be published only alongside the composition it was derived from, and never inverted to one. And **re-indexing moves the relationship only through overlap**: four moves spanning 8603.9100 cents return one number between them, to twelve decimal places.

What is proposed rather than shown. That $(9/8)^2/\varphi^2$ is a frequency; that essence goes as kn^2; that the squaring in the compound rule is a change of structural level rather than a dimensional operation; and that the latency constraint carries from cells to psychological states. **These are postulates of this framework, badged at every appearance**, and **every verdict, every fibre count and every interval in this paper survives their rejection**, because the constant cancels.

What is not claimed at all. That any audited operation was ever effective; that its practitioners understood it in these terms; that any symbol, name, sequence or substance has a physical effect; or that any number here describes any person. **Precedent is not evidence**, and the survivorship likelihood ratio for the tradition's own testimony is exactly 1.000000 — which is not weak evidence but the arithmetic statement that it leaves the prior where it found it.

12 Declarations

Data and code availability. No dataset was collected for this study and none exists. Every number is the output of exact arithmetic or exhaustive enumeration on inputs stated in the text. The 44 composite values checked in §5.3 are quoted from a published source table, which is described in the text; **no new measurement was made of any of them, and the seven disagreements are reported as recomputations of published figures rather than as observations.** The fibre enumerations are complete rather than sampled, over multisets from a twelve-element index set at $T \leq 6$. Every case — the carried dyad, the carried subject, the worked compounds and every re-indexing — is constructed and labelled as constructed at the point of use. The figures are generated by the Python scripts named in their captions and are available from the author on request.

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the audit, the analysis, the figures and the manuscript.

Ethics. No human-subjects research was conducted and no personal data were processed. The constructed subject and the constructed dyad are invented parameter sets and describe no person. **No historical person's beliefs are characterised as held by them**; the audit is of a transmitted text's operations, not of anyone's understanding, and §8.2 states that boundary explicitly.

Preregistration. No test here was preregistered, because no test was run against data. **The source list was fixed before the verdicts were reached** (§4.5, SPEC-4.1) and that is the only pre-specification the paper claims; it is not a preregistration and is not offered as one.

Bibliographic disclosure. Four works are named in the running text without citation keys because they are absent from the bibliographic database governing this corpus (§10, L-3). They are given with sufficient bibliographic detail to be located.

Clinical disclaimer. The twelve state classes are not diagnostic categories. Nothing in this paper is an assessment, a screening instrument, a treatment, counselling, or advice about any person's health, and nothing in it may be presented as such. No symbol, name, sequence, operation or substance is claimed to have any physical, chemical or biological effect.

Chapter 47

Refinement as Contraction: Why Repetition Cannot Move a Fixed Point, and an Evidential Audit of the Tradition It Comes From

Abstract. Iterated practice is modelled as the affine map $Q_{t+1} = (1-\delta)Q_t + b$, whose contraction gives $|Q_t - Q_\infty| = (1-\delta)^t |Q_0 - Q_\infty|$ as an equality rather than a bound. Two negative theorems follow. Refinement cannot move the fixed point, because $Q_\infty = (\eta F - c)/\delta$ contains no iteration count; and the ladder is silent in the frequency register, because f_{rel} is a function of u and N alone. What repetition buys is arrival time only: at $\delta = 0.09$ the half-life is 7.3496 passes and 99% closure takes 48.83, so the tradition's canonical seven passes close 48.32% of the gap — and only 6.79% at $\delta = 0.01$. Seven passes would close 95% only at $\delta = 0.348164$. Applied to a decaying arrangement, the same ladder converges at the same speed on $Q_\infty = -5.01$ instead of $+6.08$. On the tradition's own claims, four survive and five are charged. Staged process control is priced: per-stage checking at 90% detection raises yield from 47.83% to 93.21% at 0.63 reworks per unit. Cadence is priced: a single annual occasion carrying 26 nats delivers 2.4849 after clipping at $\ln 12$, 9.56% of a weekly practice, and $0.91^{52} = 0.007403$ of one impulse survives the year. Against this, transmutation by chemical means is closed by energy arithmetic (15.5 to $15{,}500$ gigawatt-station-years per gram), the record is survivorship-filtered with a likelihood ratio of exactly 1.000000 at every corpus size, secrecy makes an error immortal (survival 1.000000 per generation at zero replications), and the canonical order is backwards. Every case is constructed: no dataset was collected and no human subjects participated.

Keywords: contraction mapping; refinement ladder; seven repetitions; staged process control; cadence; survivorship; secrecy; replication

Series: C. Frequency, Essence, and Mind **Paper:** CCP-047

1 Introduction

The most widely repeated instruction in the alchemical corpus is not a recipe but a count. *Distil seven times. Circulate the matter until it no longer changes. Repeat the work.* The operation is not described as improving between passes; it is described as being *done again*, identically, and the doing again is what is supposed to produce the result. Two reactions to

that instruction are common and both are wrong: to dismiss it as ritual padding, or to accept it as a claim that repetition adds something the first pass could not reach.

This paper computes what iteration does under the framework's storage law, and the answer is a pair of negative theorems. The law is the discrete leaky integrator $Q_{t+1} = (1-\delta)Q_t + b$, a form that is not novel and is not claimed to be: it is the RC low-pass filter of signal processing [Oppenheim et al., 1996], exponential smoothing in forecasting [Holt, 1957, Brown and Meyer, 1961], and the leaky integrate-and-fire neuron of computational neuroscience [Lapicque, 1907, Gerstner and Kistler, 2002]. Its stability condition and unique fixed point are standard [Elaydi, 2005], and the convergence is Banach's [Banach, 1922]. What is asked here is narrower: **what does the iteration count buy, and what does it provably not buy?**

The first negative theorem is that refinement cannot move the fixed point. The destination $Q_\infty = (\eta F(u^*) - c^*)/\delta$ contains no iteration count, so it is computable before the first pass and unchanged by every pass after it. **The second is that the ladder is silent in the frequency register**: the relationship frequency f_{rel} is a function of the non-parallelism u and the category count N alone, so repetition moves neither the pitch, nor the folded pitch class, nor the octave deficit D. What repetition buys is *arrival time*, and arrival time is a function of the decay rate and of nothing else.

The consequence is quantitative and it cuts against the tradition. At the framework's carried decay rate $\delta = 0.09$ the half-life is 7.3496 passes, so the canonical seven passes close 48.32 per cent of the gap — under half. At $\delta = 0.01$ they close 6.79 per cent. Seven passes close 95 per cent only at $\delta = 0.348164$. **The traditional count is therefore not a universal constant and not arbitrary either: it encodes a checkable claim that the operation loses roughly a third of its stock per cycle**, which is entirely ordinary in a physical distillation and is not true of a weekly psychological practice.

The second half of the paper turns the same instrument on the tradition that supplied the count. **Four of its claims survive an audit and five do not**, and both lists are computed rather than asserted. Staged process control is priced: seven stages at a per-stage defect probability of 0.10, moving from a single terminal inspection to per-stage checking at 90 per cent detection, raises final yield from 47.83 per cent to 93.21 per cent at a cost of 0.63 reworks per unit. Cadence is priced, and the folklore's stated reason for it is refuted: the *integrated* effect of one large input and the same total delivered periodically is identical to every digit — 433.3333 in all three schedules — which is linear time invariance and not a matter of opinion. Cadence wins for two other reasons entirely, strict concavity and the surplus ceiling, and at $N = 12$ a single annual occasion carrying a nominal 26 nats delivers 2.4849 after the $\ln 12$ ceiling clips it: 9.56 per cent of a weekly practice, with 90.44 per cent lost to clipping and $0.91^{52} = 0.007403$ of one impulse surviving the year.

Against that, five charges are laid and computed. Material transmutation by chemical means is closed by an energy-register argument of eight orders of magnitude. The tradition's evidential record is survivorship-filtered, and under success-only preservation the likelihood ratio of testimonial evidence is exactly 1.000000 at every corpus size — the same defect that selective publication produces in a modern literature [Ioannidis, 2005, Open Science Collaboration, 2015]. Secrecy removes the correction mechanism, and at zero independent replications per generation an error survives with probability 1.000000 indefinitely. No control was run, and the control that would have settled the central question required only a balance. And the canonical operation order is inverted with respect to the framework's own ordering rule.

What this paper does not claim, stated at the front rather than at the back. Every frequency here is an *essence* frequency — a position on a logarithmic scale of surplus — and nothing is claimed to vibrate, resonate or sound. The constant $k = (9/8)^2/\varphi^2$ is a *defined* number; **attaching hertz to it is a separate and much stronger claim** and is postulated, not derived. The standard critique of golden-section claims [Markowsky, 1992] is cited here deliberately and again wherever φ appears. Octave reduction is licensed for pitch [Shepard, 1982] and its perceptual specialness has been measured [Demany et al., 2013]; its extension to a quantity that is not a pitch is a **modelling choice** with no external warrant. No human-subjects research

was conducted for this paper and no dataset was collected; every case is constructed and is labelled as constructed at the point of use.

Contributions

1. The contraction as an **equality** rather than a bound, $|Q_t - Q_\infty| = (1-\delta)^t |Q_0 - Q_\infty|$, with the three arrival constants τ, half-life and $t(\varepsilon)$ shown to be functions of δ alone.
2. Two negative theorems: the fixed point contains no iteration count, and the frequency reading contains no iteration count.
3. The seven-repetitions result, in both directions: what seven passes close at a given δ, and what δ would make seven passes sufficient.
4. A priced audit of four surviving traditional practices, of which staged process control (47.83 per cent to 93.21 per cent) and cadence (9.56 per cent of a weekly practice) carry numbers other work in this corpus uses.
5. Five computed charges against the tradition's largest claims, including a likelihood ratio of exactly 1.000000 and an error-survival probability of exactly 1.000000 under secrecy.
6. The counterfeit-arrival result: cost export leaves the **sum of drives** exactly invariant while improving every number the exporter can privately compute, and the disadvantaged party's stock lags its own destination by 10.850479 passes.
7. A correction printed rather than applied silently: the counterfeit-arrival identity, as originally stated on the sum of *fixed points*, is valid only under equal decay rates.

2 Hypotheses, definitions, postulates and modelling choices

The rigour standard governing this corpus requires that a claim which *could* be wrong be stated so that it could be, and that a claim which could not be wrong be labelled with which of the three non-empirical kinds it is. **A hypothesis nothing could falsify is not a hypothesis**, and this section does not manufacture one.

2.1 Hypotheses proper

H1 — the invariant-destination hypothesis. Repeating an operation whose parameters do not change moves the system's destination by zero. **Predicted:** in a system whose u, η, c and N are measured and held fixed across a run, the stock converges to the value computed from those four parameters before the first pass, and the estimate does not drift with pass count. **Falsified by:** a run in which the parameters are measured, verified constant, and the stock settles anywhere other than $(\eta F - c)/\delta$. **Status:** DERIVED from the storage law [Curzi, 2026a] by inspection. **Not tested here.** No observational data were collected.

H2 — the rate-blindness hypothesis. The arrival rate is a function of δ alone, so two arrangements with the same δ converge at the same speed whatever their destinations. **Predicted:** measured half-lives agree between an accumulating and a decaying arrangement sharing a decay rate. **Falsified by:** a measured half-life that varies with the sign or size of the drive. **Status:** DERIVED; demonstrated on constructed cases in §4.12 and §4.22. **Not tested against observation.**

H3 — the seven-passes hypothesis. A tradition prescribing seven repetitions of an operation is making an implicit claim that the operation's per-cycle loss is at least $\delta = 0.348164$, the rate at which seven passes close 95 per cent of the gap. **Predicted:** where the per-cycle loss of a traditional operation can be measured independently — for a physical distillation, the volatile fraction lost per pass — it lies at or above roughly a third. **Falsified by:** a tradition prescribing seven passes for an operation whose measured per-cycle loss is a few per cent, and which nonetheless completes.

Status: this is the paper's one genuinely empirical hypothesis and **it has not been tested.** It is cheap to test on a physical process and is offered for that reason.

H4 — the export-invariance hypothesis. In any arrangement in which cost is transferred between participants, the sum of participants' drives $\sum_i(\eta_i F_i - c_i)$ is unchanged. **Predicted:** an intervention that improves one participant's measured destination without moving the summed drive has transferred cost rather than produced order. **Falsified by:** a measured intervention in which one participant's drive rises and no other participant's falls. **Status:** DERIVED here and proved in general in the companion paper [Curzi, 2026]. **Not tested against observation.**

H5 — the staged-control hypothesis. Per-stage checking dominates terminal inspection in final yield, monotonically in the per-stage detection probability. **Predicted:** yield $(1-q(1-d))^m$ rises with d at fixed q and m, and rework rises with it too. **Falsified by:** a process in which per-stage detection is raised and final yield does not rise. **Status:** the arithmetic is DERIVED from a stated defect model; the empirical claim rests on the statistical-process-control literature, **two of whose founding works are absent from this paper's bibliographic database** and are named without keys in §5.2.

2.2 Definitions — true by stipulation, not falsifiable

- **Def. 1.** A *refinement pass* is one application of the storage law with all four parameters unchanged. If a parameter changes between passes the operation is re-indexing, not refinement.
- **Def. 2.** $u := 1 - (\hat{x} \cdot \hat{y})^2$ on unit intensity vectors; $F := \ln(1 + (N-1)u)$.
- **Def. 3.** $D := 2.885390\,(c/\eta) - \log_2(f/k)$, the octave deficit.
- **Def. 4.** $k := (9/8)^2/\varphi^2 = 0.483425732988414$. **This is a defined constant.** It is an exact ratio of exact quantities and is therefore computable to arbitrary precision; the fifteen figures quoted are *exact digits of a definition* and carry no measurement uncertainty, because there is no measurement.
- **Def. 5.** The *running deficit* D_t and *apparent pitch* f_t of §4.6 are constructed readings of a stored quantity, defined by applying to Q_t the transform that carries Q_∞ to D. They are badged as constructed wherever they appear.

2.3 Postulates — asserted by this framework, with no external support

- **P-a. The attachment of hertz to k.** Computing $(9/8)^2/\varphi^2 = 0.483425732988414$ is arithmetic. **Asserting that this dimensionless number *is a frequency* is a separate and much stronger claim**, and no operational procedure is supplied by which the constant acquires the dimension of inverse time. This is postulated, and it is restated in the dimensional audit of §3.4, where it is unavoidable.
- **P-b. The essence law $f(n) = kn^2$.** A quadratic relation between a dimensionless count and a frequency, with **no external literature supporting it**. **The exponent 2 is defined by the model, not estimated from data.** No scaling exponent in this paper was obtained by fitting, none was tested against a log-normal or exponential alternative, and the goodness-of-fit and likelihood-ratio procedure a claimed scaling relation requires [Clauset et al., 2009] was **not run, because there is no data to run it on.**
- **P-c. The affine effective count $N_{\text{eff}} = 1+(N-1)u$.** A design axiom, not a derived result, and **not a Hill number of any order**. The Hill family [Hill, 1973] and its rediscovery in political science [Laakso and Taagepera, 1979] ground $1/\text{HHI}$ and $\exp(H)$ exactly, and ground this construction not at all. A reviewer reaching for the Hill-number literature as support will find it does not apply, and that is stated here rather than left to be discovered.
- **P-d. The parameter assignment of a historical schedule.** The sorting of a canonical fifteenth-century twelve-operation schedule by which parameter each operation acts

on is this framework's reading of the text and not the text's own claim [Curzi, 2026]. The order-inversion result of §6.5 is conditional on that reading.

2.4 Modelling choices — legitimate, but choices

- **M-a. Octave reduction applied to a non-acoustic quantity.** Octave equivalence is a perceptual structure formalised for **pitch** [Shepard, 1982], whose perceptual specialness has a measured size [Demany et al., 2013]. Nothing licenses carrying it to a quantity that is not a pitch. The *operation* — division by a power of two — is well defined; its *interpretation* is a choice of this framework.
- **M-b.** $N = 12$**.** Twelve component classes are used throughout. Twelve exceeds the classical span of immediate absolute judgment [Miller, 1956], which is a reason for caution rather than a reason for twelve.
- **M-c. Interval apparatus without psychoacoustic claims.** Cents, octave equivalence and interval arithmetic are the correct apparatus once an index exists [Rossing et al., 2002], and the history of consonance theory — from Euler's arithmetic of *gradus suavitatis* [Euler, 1739] to harmonic distance [Tenney, 2015] — is the tradition this vocabulary comes from. But what a listener finds consonant depends on the spectrum of the sound and not on the ratio alone [Sethares, 2005], and interval categories generalise less far than is usually assumed [Burns, 1999, Hoeschele et al., 2012]. **This paper uses the arithmetic and claims none of the perception.**
- **M-d. Scalar reduction of a structured state.** Every reading in this paper enters through u, which is computed from a vector of elicited intensities over twelve component classes and then collapsed to one number. §3.5 states what that costs.

2.5 What is expository

Sections 4.17–4.20, 5.17–5.19 and 6.16–6.18 are expository: they set out procedure, reporting requirements, gates and failure modes rather than argue a claim. Their organising question is *what must be published alongside a refinement reading for a third party to be able to check it*, and the answer is the specification blocks and the reporting formats.

3 Framework, dimensional audit, and method

3.1 The storage law — a shared mathematical form, not a derivation

$$Q_{t+1} = (1 - \delta)Q_t + \eta S_t - C_t$$

Analogy status: category 2, a shared mathematical form. The equation governing this framework's stored order is the equation that governs a leaky integrator, and it is used because the mathematics is the same, not because the physics transfers. No claim is made that stored order is a charge, that η is a coupling constant, or that any physical result about RC circuits applies to a relation between people. The fixed point follows from setting $Q_{t+1} = Q_t$ and requires $\delta > 0$; the transfer-function reading, in which $1/\delta$ is the DC gain of $1/(z-a)$ evaluated at $z = 1$, is fixed in the framework's storage-law paper [Curzi, 2026a], whose Result Four shows that $\delta = 0$ is a pathology rather than an ideal — no contraction, no unique fixed point, and an error-persistence column of all ones, so that no mistake ever decays out of the stock.

Three warnings attach to that form and are recorded here rather than in a footnote. First, the time average of a multiplicative process differs from its ensemble average, so a parameter estimated across systems does not describe what happens to any one of them over time [Peters, 2019] — and Q accumulates *over time*, for *one* arrangement. Second, a single Q, δ or u written for a group is an aggregation, and aggregation over heterogeneous agents does not in general yield a representative agent [Kirman, 1992]. Third, Q **has the grammar of a conserved quantity — a stock that accumulates, leaks and is spent — without a conservation**

law to give that grammar content. That is the characteristic failure of the genre: neoclassical economics was built by copying nineteenth-century energetics and made utility play the role of energy without the conservation law that gave energy its meaning [Mirowski, 1989]. This paper claims no conservation law and adopts instead the position that a law is claimed only within a specified arrangement rather than universally [Cartwright, 1999]. The wider programme — physics formalism applied to social dynamics — has a standard survey that fixes what such models can and cannot be held to [Castellano et al., 2009], and those are the boundaries claimed here.

One conservation claim in this paper is genuine and is bounded accordingly. §4.15a establishes that the *sum of drives* is invariant under transfer of cost. That is a conservation statement about a bookkeeping identity, not about any physical quantity: it holds because a transfer enters one ledger with a plus and another with a minus, and it would hold for any additively separable accounting. It carries no physical content and is not offered as one.

3.2 The surplus functional and the effective count

$F(u) = \ln(1 + (N-1)u)$ is the framework's surplus functional, characterised by angle-only dependence, a zero at parallel states, an affine effective count and a normalisation. Its logarithmic form places it in the same family as Shannon's entropy [Shannon, 1948], and the ceiling $F \leq \ln N$ is the same ceiling that appears in requisite-variety arguments [Ashby, 1956, 1958].

Which construct of difference u measures must be declared, because at least three incompatible constructs travel under the word *diversity* — separation, variety and disparity [Harrison and Klein, 2007]. Here u is a *disparity* measure over represented states, and it is symmetric in its arguments, which psychological similarity is known not to be [Tversky, 1977]. That is a limitation and it is printed in §10.

The functional form is underdetermined and this paper does not pretend otherwise. A power law with exponent below one [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the qualitative facts as well as the logarithm, and many mechanisms produce indistinguishable curves [Stumpf and Porter, 2012]. The strongest available warrant for the logarithm is not the psychophysical one: a logarithmic evaluation follows from multiplicative dynamics with **no psychological assumption at all** [Peters and Gell-Mann, 2016], and that is the warrant this paper's arguments actually need.

3.3 The essence bridge and the constant k

The bridge writes a surplus as a pitch, $f = k\, e^{2F}$, so that intervals of surplus become intervals of frequency and the constant cancels from every ratio. The constant is

$$k := \frac{(9/8)^2}{\varphi^2} = \frac{81/64}{\varphi + 1} = 0.483425732988414 \quad \text{(defined; exact)}$$

Both components are real and documented: the ratio $9{:}8$ is the Pythagorean whole tone [Barbour, 1951, Lindley and Turner-Smith, 1993], and the golden section is Euclid's extreme and mean ratio [Euclid, 1956], with $\varphi^2 = \varphi + 1$ and $\varphi^m = F(m)\varphi + F(m-1)$ proved identities [Vajda, 1989].

Their product, expressed as a frequency in hertz, has no precedent. The literature on golden-section claims is overwhelmingly a literature of *deflation*: the appearances of φ in architecture, the human body and preferred rectangles do not survive examination [Markowsky, 1992, Fischler, 1981], and the psychological preference results are weaker than they are usually reported to be [Höge, 1997]. **This framework's position is the one those critiques leave open and no wider:** φ is being *chosen* as a constant, not *found* in a psychological system. Nothing in this paper is evidence that the golden section is present in anything measured, and the constant is invisible to every verdict in this paper because D is a difference of logarithms in which k appears in both terms and cancels.

3.4 Dimensional audit

symbol	quantity	dimension	unit	status
Q	stored order	$[\mathcal{C}]$	none supplied	state variable; **no operational unit**
b	per-pass drive	$[\mathcal{C}\,T^{-1}]$	per declared pass	$\eta F - c$
c	cost per pass	$[\mathcal{C}\,T^{-1}]$	per declared pass	driving term
δ	decay rate	$[1]$	per pass	dimensionless; $0 < \delta \leq 1$
η	conversion efficiency	$[1]$	—	dimensionless; assumed $0 < \eta \leq 1$
u	non-parallelism	$[1]$	—	dimensionless by construction, $u \in [0, 1]$
F	surplus	$[1]$	nats	dimensionless; a logarithm of a count
N, n	category / component count	$[1]$	—	absolute scale, a count
t	pass index	$[1]$	passes	a count; the cycle must be declared
τ	time constant	$[T]$	passes	$-1/\ln(1-\delta)$
k	reference constant	$[T^{-1}]$ **asserted**	Hz **asserted**	**defined as a ratio (dimensionless); the unit is postulated**
f, f_{rel}, $f_{\min}$	essence frequency	$[T^{-1}]$	Hz	inherits its dimension entirely from k
D	octave deficit	$[1]$	octaves	a ratio of logarithms; k cancels
q, d	defect / detection probability	$[1]$	—	probabilities
r	replications per generation	$[1]$	per generation	a count; the generation must be declared

Does the equation balance? The storage law balances provided the drive carries the same dimension as Q per pass, which the framework assumes and does not measure; Q has **no operational unit**, and this audit records that rather than inventing one. The error law $|Q_t - Q_\infty| = (1-\delta)^t |Q_0 - Q_\infty|$ is dimensionally clean: $(1-\delta)^t$ is a pure number raised to a count, and the two sides carry $[\mathcal{C}]$ alike. The essence law $f(n) = kn^2$ balances **only if** k **carries** $[T^{-1}]$, because n is a dimensionless count, and **no operational procedure is supplied by which it acquires that dimension.** This is the single largest dimensional exposure in the paper and it is listed in §10 by name.

The ratio c/η is dimensionally a cost flow divided by a pure number and is therefore not dimensionless as the exponent of $e^{2c/\eta}$ requires. **The framework's convention is that** c **is expressed in the same units as** F **— nats of surplus per declared pass — which makes** c/η **di-**

mensionless. That convention is a choice, it is stated here, and every reading in this paper depends on it being applied consistently.

Two quantities in §6 are physical and carry real units, and they are the only ones that do. The Coulomb barrier of §6.1 is in MeV and the beam-energy estimate of §6.8 is in joules; both are computed from standard constants and are dimensionally ordinary. They are the exception, and the contrast with Q is the point of the table.

3.5 Scale type, and what licenses the arithmetic

quantity	scale type claimed	what licenses it
n, N, t, r	absolute	cardinality: true zero, fixed unit, identity transformation group only
u	ratio	constructed from an inner product of unit vectors; zero is parallel
elicited intensities	ratio, up to a free common multiplier	**magnitude estimation, assumed** [Stevens, 1957]; not tested here
η, c, δ	ratio, **assumed**	no instrument is supplied
q, d, survival probabilities	ratio	probabilities on a stated model
D	interval, in octaves	a difference of logarithms

The strongest available objection to this framework is Michell's: that psychology routinely assumes its attributes are quantitative without ever testing whether they are [Michell, 1997, 1999]. The objection applies here in full. The formal remedy exists — conjoint measurement supplies a testable condition for additive quantitative structure [Luce and Tukey, 1964, Krantz et al., 1971], and a Rasch model supplies specific objectivity [Rasch, 1960, Andrich, 1988] — and **neither has been run for any quantity in this paper.**

Michell's objection and the appraisal literature's objection are the same objection in two domains, and this paper names the pattern once rather than meeting each instance as a surprise. The pattern is *scalar reduction*: the habit of collapsing a structured state into a single number. Where Michell asks whether the attribute has quantitative structure at all, the affect literature asks whether the structure it has is the kind a single number can carry — and answers that emotions are distinguished by the **pattern across** independently varying appraisal dimensions [Smith and Ellsworth, 1985], that a state is a synchronisation process in time rather than a point [Scherer, 2001], that individuation runs by qualitative core relational theme so that differing by *what* is appraised does not license differing by *how many* [Lazarus, 1991], and that action readiness is constitutive of emotion as such rather than a rung above other emotions [Frijda, 1986]. Two dimensions of core affect are well supported [Russell, 1980]; twelve ordered rungs are not.

Everything in this paper that enters through u inherits that objection, because u is a scalar summary of exactly such a structured state. What survives it are the results that do not require u to be a measurement: the contraction, the two negative theorems, the arrival constants, the seven-passes arithmetic, the export invariance, the staged- control arithmetic and the survivorship and secrecy results. Those are statements about a model given its parameters, and they hold whatever the parameters turn out to mean.

3.6 Identification — the objection this framework cannot presently answer

The framework's founding construction is a reflexive loop, $T(x) = W(x, M(x))$: the group influences the member and the members constitute the group [Curzi, 2026]. **That is a linear-in-means reflexive specification, and the endogenous social effect in such a specification is not identified from observational data** [Manski, 1993]. It cannot be separated from contextual effects or from correlated unobservables, however much data is collected; influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011]. **No amount of fit, held out or otherwise, recovers the parameter.**

The remedy this paper adopts is stipulation, and it pays the price openly. The reference set is fixed in advance — the declared category set of size N and the declared decomposition rules — and every reading is reported together with that stipulation, which is what makes a discrete-choice treatment of social interaction defensible where an unrestricted one is not [Brock and Durlauf, 2001]. **This paper therefore presents the reflexive loop as a modelling position, not as an identified empirical result.**

The same objection applied to the framework's own parameters. If systems change their behaviour when modelled, then δ, η and u are the decision rules of people who have not read the storage law, and they are not structural in the sense a policy exercise requires [Lucas, 1976]. Classification changes the classified [Hacking, 1995], and a theory has been observed becoming true by being used, with dates and prices [MacKenzie and Millo, 2003, Ferraro et al., 2005]. **The framework asserts reflexivity for its subject matter and cannot consistently exempt its own parameters from it.** This paper does not resolve that; it states it, and treats every parameter as conditional on the arrangement in which it was fixed.

3.7 Method

Every number in this paper is the output of exact arithmetic on stated parameters, computed in double precision. **No observational data were collected, no human subjects participated, and no empirical dataset underlies any figure or table.** Every case is constructed and is labelled as constructed at the point of use. Residuals quoted at the 10^{-13} level or below are floating-point round-off, not measurement error, and are reported to demonstrate that two algebraically identical expressions were computed independently and agreed. The figures are produced by the scripts named in their captions.

The audit method of §5 and §6 is a two-column concordance. **The evidence column carries what could have come out otherwise; the precedent column carries only that an idea was independently arrived at and preserved.** Two sources agreeing does not double the evidence when only one of them is evidence, and the arithmetic for why the precedent column can never be promoted is §6.2's: under success-only preservation the likelihood ratio is exactly 1.000000 [Curzi, 2026].

4 Refinement as contraction

The instruction that survives in more alchemical texts than any other is not a recipe but a count. *Distil seven times. Circulate the matter until it no longer changes. Repeat the work.* The operation is not described as being improved between passes; it is described as being *done again*, identically, and the doing again is what is supposed to produce the result.

A modern reader has two available reactions and both are wrong: to dismiss the instruction as ritual padding, or to accept it as a claim that repetition adds something and that the seventh pass reaches somewhere the first could not. Neither survives the arithmetic, and the arithmetic is short. This section computes what iteration does. The answer follows directly from the storage law [Curzi, 2026a], and it is this:

> **Iteration closes a gap. It does not choose the destination. The destination was fixed by the parameters before the first pass began, and no number of passes moves it**

by any amount whatsoever.

That is why the construction is called a *ladder* and not a *staircase*. A staircase goes somewhere new with each step. A ladder is a fixed structure, and climbing it arrives where it was already pointing when it was set against the wall.

The framework's operation audit establishes what it means to move a state from one index to another: transmutation is re-indexing, and nothing not already latent can be reached [Curzi, 2026]. This section establishes the complementary fact about doing the same thing repeatedly. Section 5 then asks what the tradition got right, and cadence is on that list — a real result, but not for the reason the tradition gave.

The law

4.1 One rung of the ladder A refinement pass is one application of the storage law. **Analogy status: category 2, a shared mathematical form.** The equation below is the discrete leaky integrator that appears as the RC low-pass filter in signal processing [Oppenheim et al., 1996], as exponential smoothing in forecasting [Holt, 1957, Brown and Meyer, 1961], and as the leaky integrate-and-fire neuron in computational neuroscience [Lapicque, 1907, Gerstner and Kistler, 2002]. It is used because the mathematics is the same, not because the physics transfers. The law is the framework's contracting storage loop [Curzi, 2026a]:

```
Q      =  (1 − delta) ∗ Q   +  ( eta ∗ F(u∗)  −  c∗ )
 t+1                      t

       delta   per−cycle retention loss, in (0, 1]
       eta     conversion: how much of an interaction's surplus becomes stored order
       F(u∗)   the interaction surplus at the operating point, F = ln(1 + 11 u) for N = 12
       c∗      the per−cycle cost of maintaining the arrangement
```

Write the drive as a single constant, because for a fixed arrangement it is one:

```
b  =  eta ∗ F(u∗)  −  c∗                    the per−pass drive

Q      =  (1 − delta) ∗ Q   +  b
 t+1                      t
```

The fixed point is

```
Q     =  b / delta  =  ( eta ∗ F(u∗) − c∗ ) / delta
 inf
```

The framework calls Q_∞ the *destination* of the arrangement: it is where the loop arrives if the parameters hold. That definition is fixed in the companion paper on destiny as a fixed point [Curzi, 2026], whose Result Two states the contraction with constant $1-\delta$ and the equality $|Q_t - Q_\infty| = (1-\delta)^t |Q_0 - Q_\infty|$ at a maximum residual of 2.0624×10^{-15} over a hundred passes. This paper owns the approach to it.

A rung of the refinement ladder is one application of that map with the parameters unchanged. That last clause is the whole content of the result. If the parameters change between passes, the operation is not refinement — it is re-indexing, which the framework's audit of a traditional operation set governs [Curzi, 2026]. Refinement means: do the same thing again.

4.2 The error law **THEOREM 4.1 — THE CONTRACTION.** For the affine map above with `delta > 0`,

```
| Q  −  Q    |  =  (1 − delta)^t  ∗  | Q  −  Q    |
   t     inf                           0     inf
```

Proof. Subtract the fixed-point equation `Q_inf = (1−delta) Q_inf + b` from the update equation `Q_{t+1} = (1−delta) Q_t + b`. The drive `b` cancels exactly. What remains is `Q_{t+1} − Q_inf = (1−delta)(Q_t − Q_inf)`. Take absolute values and iterate `t` times. ■

The stability condition and the unique fixed point are standard for a first-order linear difference equation [Elaydi, 2005, Kelley and Peterson, 2001]. Three properties of this identity deserve to be read off before any table is computed.

It is an equality, not a bound. Banach's theorem [Banach, 1922] — stated exactly in the framework's foundations as: a contraction with modulus $q < 1$ on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least q each pass [Curzi, 2026] — requires only that distances not increase by more than a factor. Here they shrink by exactly `1 − delta` every pass, at every distance, forever: no fast region and no slow region.

The drive cancelled. `eta`, `F`, `c` — everything that decides *where* the loop is going — vanished in the subtraction. They set `Q_inf` and they do not appear in the rate. The gap closes at a speed that depends on the leak and on nothing else.

Nothing reaches the fixed point. `(1−delta)^t` is strictly positive for every finite `t`: the gap gets small and does not become zero. There is no pass on which the work is finished, and the tradition's discomfort about knowing when to stop is a correct perception of a system that supplies no completion signal.

4.3 The three constants of arrival Everything operational in this section comes off three quantities, all functions of δ alone. They are the time constant, half-life and settling time of a first-order lag, in the ordinary control-engineering sense [Åström and Murray, 2008, Horowitz and Hill, 2015]:

```
tau        =  −1 / ln(1 − delta)                passes per factor of e

half−life  =  ln(0.5) / ln(1 − delta)           passes to close half the gap

t(eps)     =  ln(eps) / ln(1 − delta)           passes to reduce the gap to a fraction eps
```

None of the three contains `Q_0`, `Q_inf`, `eta`, `F`, or `c`. **How long arrival takes is a property of the leak; where arrival happens is a property of everything else.** The two questions do not share a parameter. That separation is the formal skeleton of this section and it is also, in a different register, the framework's standing two-reading limitation — §4.7 makes the connection exact.

4.4 The destination is fixed before the first rung **THEOREM 4.2 — REFINEMENT CANNOT MOVE THE FIXED POINT.** `Q_inf = (eta F(u∗) − c∗)/delta` contains no `t`. Therefore for any `t`, `Q_inf(t) = Q_inf(0)`. ■

The proof is inspection, which is the correct proof: the fixed point is a function of four parameters, and iteration is not one of them.

Stated in the tradition's own vocabulary: if a practice repeats an operation whose parameters do not change, then however many times it is repeated it converges to a state computable before the first repetition — from `eta`, `F`, `c` and `delta`, and from no observation of the practice at all. **The seventh distillation reaches no destination the first distillation was not already headed toward.** What it does is arrive closer, and how much closer is the subject of §4.9.

4.5 Refinement does not move the pitch The frequency reading is the framework's essence bridge in its relationship form:

```
f     =  k * ( 1 + (N-1) u )^2  =  k * N_eff^2          k = 0.483425732988414 Hz
 rel
```

THEOREM 4.3 — THE LADDER IS SILENT IN THE FREQUENCY REGISTER. `f_rel` is a function of `u` and `N`. Neither is an argument of the iteration count. Therefore `f_rel(t) = f_rel(0)` for all `t`, and so are the folded pitch, the pitch class, and the octave deficit `D`. ■

This is the framework's state-frequency negative result arriving from the other side. There, frequency cannot distinguish a transient state from a standing disposition, because persistence is not an argument of $f = kn^2$. Here, that frequency cannot tell the first pass from the seventieth, for the same reason: **the pitch was never told about time, and repetition is entirely a fact about time.**

A second silence sharpens Theorem 4.3. The framework's operation audit [Curzi, 2026] establishes that a state frequency and a relationship frequency are independent readings: four constructed pairs spanning thousands of cents return identical `u` and `f_rel` to twelve decimal places, because `u` reads the angle between state vectors and not which indices are occupied. So there are two silences here — refinement moves neither reading, and neither could be derived from the other even if it did. **A complete report gives the state frequencies, the relationship frequency, and the trajectory, and derives none of the three from the others.**

Read as a claim about the tradition, this is sharp and worth being sharp about. Repetition is not transmutation; an operation repeated is the same operation. For the pitch to change, `u` or `N` has to change — a state genuinely added, a genuinely distinct category admitted — and that is the subject of the framework's re-indexing audit [Curzi, 2026], whose Result One is an audit distribution over thirteen entries — seven MAPS, two SPLIT, one RENDERING, three REJECTED, with five entries carrying a failure. The compound rule for combining components is fixed there too; it is not reachable by doing one thing twice.

4.6 The ladder in octaves — a construction, and it is badged as one The two theorems above make it sound as though the frequency register has nothing to say about refinement. It has one thing, and it is exact.

The deficit identity of the minimum-frequency paper [Curzi, 2026], verified to $8.88×10^{-16}$ at twelve parameter sets, is

```
D  =  2.885390 * ( c/eta  -  F )  =  -2.885390 * (delta/eta) * Q
                                                               inf
```

The second form maps the fixed point onto the octave scale. **DEFINITION 4.4 (constructed, badged).** Apply the same map to the *running stock*:

```
D    :=  -2.885390 * (delta / eta) * Q          the running deficit
 t                                    t

f    :=  k * 2^( 2.885390 * (c/eta)  -  D  )    the apparent pitch of the stock at pass t
 t                                      t
```

This is a definition, not a discovery. It is legitimate because it is exactly the transform that carries `Q_inf` to `D`, applied to a different argument; it is useful because of the following.

THEOREM 4.5 — THE OCTAVE LADDER. With `D_t` as defined,

```
| D  -  D |  =  (1 - delta)^t * | D  -  D |
   t                             0
```

Proof. `D_t` is a fixed negative multiple of `Q_t` and `D` is the same multiple of `Q_inf`. Multiplying both sides of Theorem 4.1 by `2.885390 (delta/eta)` gives the result. ■

COROLLARY 4.6 — THE GLIDE. If the stock starts empty, `Q_0 = 0`, then `D_0 = 0` and `f_0 = k * 2^(2.885390 c/eta) = f_min` exactly. As the ladder is climbed, `f_t` rises monotonically from `f_min` to `f_rel`, and the total glide is `|D|` octaves. ■

That is the frequency-register statement of what refinement is. **A refinement ladder begun from nothing is a glide that starts at the minimum sustainable frequency and converges on the arrangement's actual pitch.** It never overshoots, it never arrives, and its span is the octave deficit read backwards. In a sustainable relation `D < 0` the glide is upward and spans `|D|` octaves; in an unsustainable one `Q_inf` is negative, `D_t` runs the other way, and the glide descends below `f_min` — the correct picture of a practice refining itself into deficit.

The standing prohibition applies with full force and is restated here because this section is the one most likely to be misread. **These are essence frequencies: positions on a logarithmic scale of surplus.** `f_t` is a bookkeeping transform of a stored quantity. Nothing is claimed to vibrate, no person or tissue resonates at any of these values, and the glide is a way of reading a number, not a sound anything makes.

4.7 The two-reading limitation: rate and sustainability are separate readings This is one of the places where the framework's central limitation becomes structural rather than incidental, so it is stated here in full rather than referenced.

From the minimum-frequency paper [Curzi, 2026], whose Result Four is that D contains no δ — on the carried dyad δ from 0.09 to 0.05 leaves $D = +4.338420$ unchanged while Q_∞ falls from -5.0119 to -9.0215 — verified independently to 2.665×10^{-15} across twelve parameter sets:

```
D  =  2.885390 * ( c/eta - F )        contains k, c, eta, N, u.   DOES NOT CONTAIN delta.

Q     =  ( eta F - c ) / delta        contains delta, and is scaled by it.
 inf
```

The carried dyad as found (case 2a, CONSTRUCTED):

```
delta = 0.09   ->   Q     = -5.0119395464       deficit D = +4.338420
                     inf

delta = 0.05   ->   Q     = -9.0215             deficit D = +4.338420
                     inf
```

Lowering the loss nearly doubles the negative fixed point and moves the deficit by exactly nothing. Now put Theorem 4.1 beside it. The convergence rate is `(1-delta)^t`: it is a function of `delta` and of nothing else. So the same parameter that the deficit is blind to is the *only* parameter the arrival rate can see.

```
delta governs the RATE of arrival.        D cannot see delta at all.
D     governs the SIGN and SIZE of the sustainability gap.   The rate cannot see D.
```

Refinement speed and sustainability are therefore two independent readings of one system, and neither substitutes for the other. A ladder that converges fast is not thereby converging on somewhere good. A relation with a small deficit is not thereby getting there quickly. Any report that gives one number is incomplete, and §4.19's reporting format requires both.

4.8 Honesty clauses, restated in place **(1) ESSENCE, NOT RESONANCE.** Every frequency in this section is an essence frequency — a position on a logarithmic surplus scale. Nothing physically resonates at any value here.

(2) NO CLINICAL CLAIM. The twelve states are not diagnoses. A refinement ladder is not a treatment, a course of therapy, or an assessment of anyone's health, and nothing in this section describes or prescribes for any person's condition.

(3) ALL CASES CONSTRUCTED. Every case in this section is constructed and labelled. None is an observation of anyone.

(4) THE READING INHERITS ITS INPUTS. The trajectory is exactly as defensible as `delta`, and `delta` is usually the hardest of the four parameters to estimate. §4.18 and the framework's decay-estimation procedure give the honest procedure and the honest error bars.

(5) SCALE STATUS. n is a count and is absolute-scaled by construction. The intensities `w` must be magnitude-estimated with raw L2 norms published, or the relationship quantities computed from them are not determined. The framework's scale demotion binds the moment ratio D/P specifically; it does not bind the component index and it does not bind magnitude-estimated intensities.

(6) NO INTERVENTION MAY REMOVE THE ABILITY TO DECLINE. Consent is structural, not advisory: the framework's closure law makes the ability to decline a condition of any arrangement [Curzi, 2026]. §4.15 is about what happens when it is violated, and the arithmetic there is the argument.

The engineering

4.9 The convergence table Computed from `t(eps) = ln(eps)/ln(1−delta)` and `tau = −1/ln(1−delta)`. Fractional values are shown because the exactness matters; the integer in brackets is the first whole pass at or beyond the threshold.

`delta`	`1−delta`	`tau`	half-life	to 5%	to 1%	to 0.1%
0.01	0.99	99.4992	68.9676	298.07 [299]	458.21 [459]	687.32 [688]
0.02	0.98	49.4983	34.3096	148.28 [149]	227.95 [228]	341.92 [342]
0.03	0.97	32.8308	4.7566	98.35 [99]	151.19 [152]	226.79 [227]
0.05	0.95	19.4957	13.5134	58.40 [59]	89.78 [90]	134.67 [135]
0.07	0.93	13.7797	9.5513	41.28 [42]	63.46 [64]	95.19 [96]
0.09	0.91	10.6033	7.3496	31.76 [32]	48.83 [49]	73.24 [74]
0.10	0.90	9.4912	6.5788	28.43 [29]	43.71 [44]	65.56 [66]
0.12	0.88	7.8227	5.4223	5.43 [24]	36.02 [37]	54.04 [55]
0.15	0.85	6.1531	4.2650	18.43 [19]	28.34 [29]	42.50 [43]
0.20	0.80	4.4814	3.1063	13.43 [14]	20.64 [21]	30.96 [31]
0.25	0.75	3.4761	2.4094	10.41 [11]	16.01 [17]	6.01 [25]
0.30	0.70	2.8037	1.9434	8.40 [9]	12.91 [13]	19.37 [20]
0.40	0.60	1.9576	1.3569	5.86 [6]	9.02 [10]	13.52 [14]
0.50	0.50	1.4427	1.0000	4.32 [5]	6.64 [7]	9.97 [10]

The row in bold is the carried case's loss rate and the one every worked example in this framework uses. **At `delta = 0.09` a refinement ladder is thirty-two passes from being within five per cent of its destination**, and seventy-four passes from being within a tenth of a per cent.

Read the table with the period attached. A rate without its cycle is not a rate; the same caution governs every discrete-time filter [Oppenheim et al., 1996]. At `delta = 0.09` per week

that is thirty-two weeks to the five-per-cent band; per year, thirty-two years. The arithmetic does not know which, and §4.17 makes declaring the period a gate.

4.10 Two exact invariances **The cost of each halving is constant.** At `delta = 0.09` the half-life is 7.3496 passes, and that figure is the cost of going from 100% of the gap to 50%, from 50% to 25%, and from 1% to 0.5%. The expression `ln(0.5)/ln(1−delta)` has no `t` in it. There is no acceleration and no fatigue.

The cost of each factor of ten is constant. At `delta = 0.09` it is 6.4149 passes — the difference between the 1% column and the 0.1% column, 73.2447 minus 48.8298.

interval closed	passes at `delta = 0.09`
0% -> 50% of the gap	7.35
50% -> 90%	17.06
90% -> 99%	6.41
99% -> 99.9%	6.41

The last nine hundredths of one per cent costs exactly what the ninety-to-ninety-nine per cent did, and more than three times what the first half did. This is the arithmetic behind every practice that felt rapid at the start and interminable at the end. Nothing slowed down. The remaining distance is being felt in absolute terms by an operator whose instrument works in ratios, and the ratios are being closed at a perfectly constant rate throughout.

4.11 What seven repetitions actually buy The tradition's number is seven. Theorem 4.1 makes the question answerable: seven passes close `1 − (1−delta)^7` of the gap, and that is a claim about `delta` and nothing else.

`delta`	`(1−delta)^7` gap remaining	gap closed by seven passes	passes needed for 99%
0.01	0.932065	6.79%	458.21
0.02	0.868126	13.19%	227.95
0.05	0.698337	30.17%	89.78
0.09	0.516761	48.32%	48.83
0.10	0.478297	52.17%	43.71
0.15	0.320577	67.94%	28.34
0.20	0.209715	79.03%	20.64
0.25	0.133484	86.65%	16.01
0.30	0.082354	91.76%	12.91
0.40	0.027994	97.20%	9.02
0.50	0.007812	99.22%	6.64
0.60	0.001638	99.84%	5.03
0.70	0.000219	99.98%	3.83

And the sharper inverse question: what loss rate makes seven passes *sufficient*? Solving `(1−delta)^7 = eps`:

```
seven passes close 95%    requires   delta  =  1 − 0.05^(1/7)  =  0.348164
seven passes close 99%    requires   delta  =  1 − 0.01^(1/7)  =  0.482053
seven passes close 99.9%  requires   delta  =  1 − 0.001^(1/7) =  0.627241
```

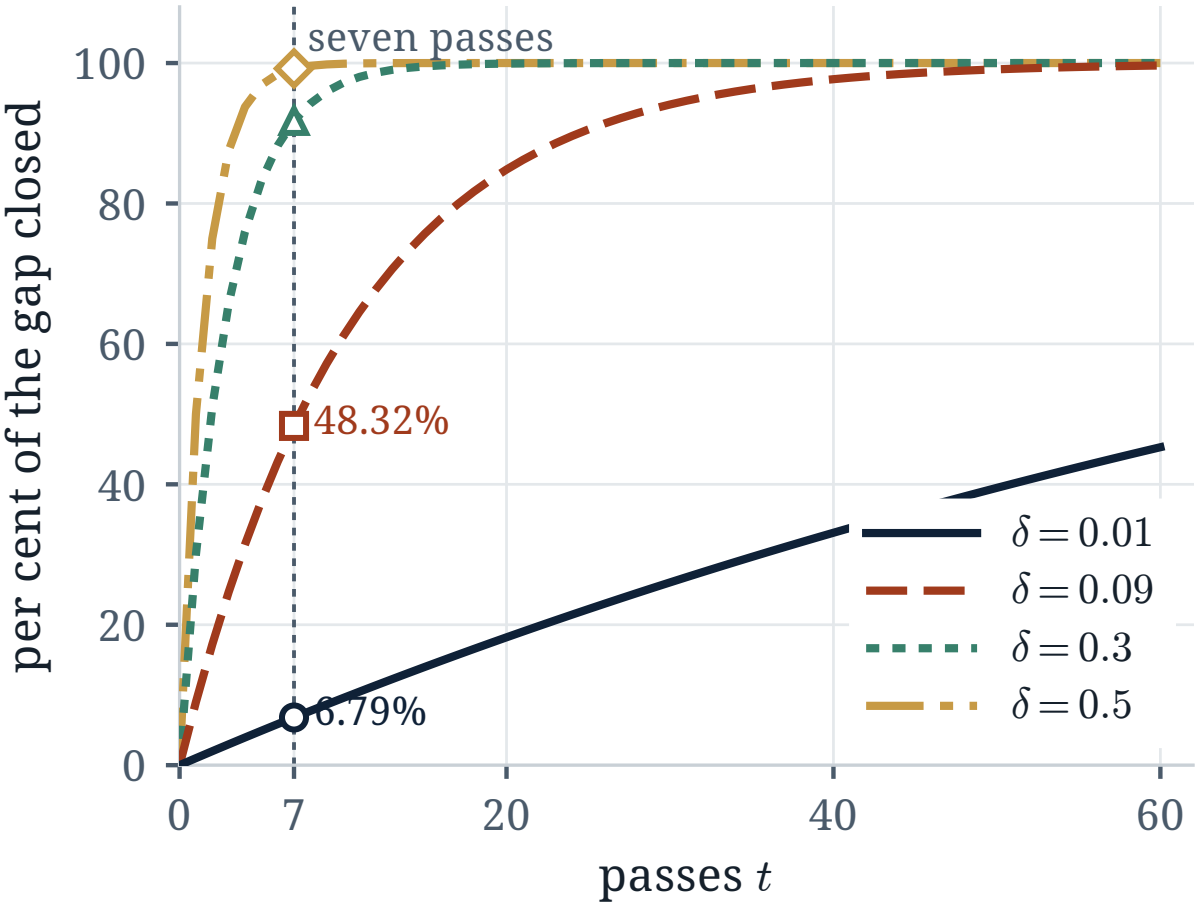


Figure 47.1: **Seven passes close 48.32 per cent of the gap at $\delta = 0.09$ and 6.79 per cent at $\delta = 0.01$**, so the traditional count is adequate only where the per-cycle loss is a third or more. Curves are $1 - (1 - \delta)^t$, the equality of Theorem 4.1; the four marked points are the seven-pass readings of the table above. Modelled: the four decay rates are chosen, and $\delta = 0.09$ is this framework's carried case. Built by `figs/src/CCP-047_fig01.py`.

Seven repetitions is adequate only in a system that loses roughly a third to a half of its stored order every cycle, and in a system with a leak of a few per cent it is nearly worthless. That is a real finding and it cuts both ways: the traditional count is not a universal constant, and it is not arbitrary either, because it encodes an implicit and checkable claim about the process. A tradition prescribing seven passes is asserting `delta \ensuremath{\approx} 0.35` or higher for its own operation, and in a physical distillation — where the loss is the volatile fraction left behind each pass — that is entirely ordinary. **The instruction is not mystical; it is a badly-stated engineering estimate, and it is testable.** A prescription that encodes an implicit parameter value is exactly the kind of claim that becomes checkable once the parameter is named [Platt, 1964, Meehl, 1990].

Whether it transfers to a psychological practice is a different question, and the answer is: only if `delta` there is also around a third, which for a weekly practice it is not. The framework's decay-estimation procedure is how one would find out.

4.12 Rate against destination — the two-by-two Because `delta` moves the rate and cannot move the deficit, and because `D` fixes the destination's sign and cannot see `delta`, the two readings are orthogonal and all four combinations occur. Both rows use the carried dyad's arrangements (CONSTRUCTED); only δ differs.

case	`u`	`eta`	`c`	`delta`	`D` (octaves)	`Q_inf`	half-life	to 1%
2b, fast leak	0.692373	0.44	0.40	0.30	-3.590989	+1.825328	1.9434	12.91
2b, slow leak	0.692373	0.44	0.40	0.03	-3.590989	+18.253280	4.7566	151.19
2a, fast leak	0.068735	0.30	0.62	0.30	+4.338420	-1.503582	1.9434	12.91
2a, slow leak	0.068735	0.30	0.62	0.03	+4.338420	-15.035819	4.7566	151.19

Four things to read off it, in order of how badly they are usually got wrong.

The `D` column has two values, not four. Changing `delta` by a factor of ten does not perturb the deficit in the twelfth decimal place. The decay-blindness result [Curzi, 2026] says it cannot, and here is the demonstration.

The rate columns have two values, not four. They are identical between a healthy arrangement and a failing one. The failing arrangement converges at exactly the same speed; it is converging on `-15.04`.

A slow ladder on a bad arrangement is the worst of the four. Row three settles at `-1.50` in thirteen passes; row four takes 151 passes to settle at `-15.04`, a destination ten times worse, because the low leak that slowed it down also multiplied the negative numerator. This is the framework's ordering rule in the refinement register: trust lowers δ, δ divides an already-negative numerator, and the result is a slower ladder to a much worse place. The elasticities that price the ordering are fixed in the return-difference paper [Curzi, 2026q], whose Result Four gives η $+1.594$, δ -1.000, u $+0.770$, C -0.594 — a ten per cent efficiency gain is worth 2.7 times a ten per cent cost cut.

Speed feels like progress and is not evidence of it. A practice that visibly changes a person fast is a practice with a large `delta`, and `delta` is a *loss* term. The visible rapidity is the signature of a leaky system reaching its level, which may be any level at all.

4.13 The glide, computed Case 2b of the carried dyad (CONSTRUCTED): `u = 0.692372881355932`, `eta = 0.44`, `c = 0.40`, `delta = 0.09`, `N = 12`. Then `F = 2.153633`, `f_rel = 35.8881768941 Hz`, `f_min = 2.9782153324 Hz`, `D = -3.5909886634` octaves, `Q_inf = +6.0844267427`, drive `b = 0.547598406839`. The ladder starts from an empty stock, `Q_0 = 0`.

`D_t` and `f_t` are Definition 4.4's constructed readings; the final column is the distance from the arrangement's true pitch, in cents.

pass `t`	`Q_t`	`D_t` (oct)	`f_t` (Hz)	cents below `f_rel`
0	0.000000	0.000000	2.978215	4309.19
1	0.547598	-0.323189	3.726024	3921.36
2	1.045913	-0.617291	4.568558	3568.44
3	1.499379	-0.884924	5.499772	3247.28
5	2.287549	-1.350096	7.592339	2689.07
7	2.940232	-1.735306	9.915969	2226.82
10	3.715053	-2.192600	13.614271	1678.07
15	4.605861	-2.718349	19.600209	1047.17
20	5.161754	-3.046433	6.605033	653.47
25	5.508650	-3.251169	28.356682	407.78
32	5.786888	-3.415383	31.775238	210.73
40	5.944508	-3.508410	33.891644	99.09
49	6.024551	-3.555651	35.019794	42.41
60	6.063209	-3.578466	35.578013	15.03
74	6.078761	-3.587645	35.805087	4.01
100	6.083939	-3.590701	35.881014	0.35
150	6.084422	-3.590986	35.888113	0.0031

The total span is `1200 * log2(35.8881768941 / 2.9782153324) = 4309.1864` cents, which is `3.5909886634` octaves, which is `-D` to every digit — Corollary 4.6 verified.

Seven passes — the traditional count — leave this ladder 2226.82 cents short of its own pitch, which is more than an octave and a half. At `delta = 0.09` seven is not a completion; it is slightly less than half the gap, exactly as §4.11 predicted.

One further reading, because it is the useful one. At $N = 12$ the minimum separation between distinct pitch classes in the folded lattice is 34.9756 cents. The glide is inside that separation — indistinguishable from arrival at the resolution the framework's own lattice supports — from pass 49 onward, where it stands at 42.41 cents, crossing below 34.9756 cents at pass 51. **Beyond about fifty passes the ladder is doing work the instrument cannot see.** That is the honest stopping criterion and §4.18 formalises it.

4.14 Why arrival is recognition rather than ascent The framework's founding definitions record that arrival is experienced as recognition rather than as victory, and offer no mechanism [Curzi, 2026]; the storage law supplies one in a line [Curzi, 2026a]: a contraction arrives, it does not climb. Here is the same claim with the arithmetic under it, in four parts.

One: there is no arrival event. `(1-delta)^t > 0` for every finite `t`, so there is no pass on which something completes and nothing for a sense of achievement to attach to. What there is instead is a gap that has become smaller than the operator can resolve — and *noticing that the difference can no longer be told* is a recognition, not an accomplishment.

Two: almost the entire trajectory is spent near the destination. The psychological literature on adaptation reaches the same shape from measurement rather than from algebra — an adaptation level against which experience is judged [Helson, 1948], a return toward a baseline after large positive and negative events [Brickman et al., 1978], and the qualifications that later work attaches to how complete that return is [Diener et al., 2006]. At `delta = 0.09` the ladder is within 5% of `Q_inf` from pass 32 and within 0.1% from pass 74. A practice run for three hundred passes spends 269 of them inside the five-per-cent band. Whatever being at the fixed point is like, it is what most of a practice feels like, and the transient is the unusual part. The approach is not the experience; it is the anomaly at the beginning.

Three: the destination is derived from what was already there. Every term of `Q_inf = (eta F - c)/delta` is a standing property of the arrangement; nothing in the fixed point is imported.

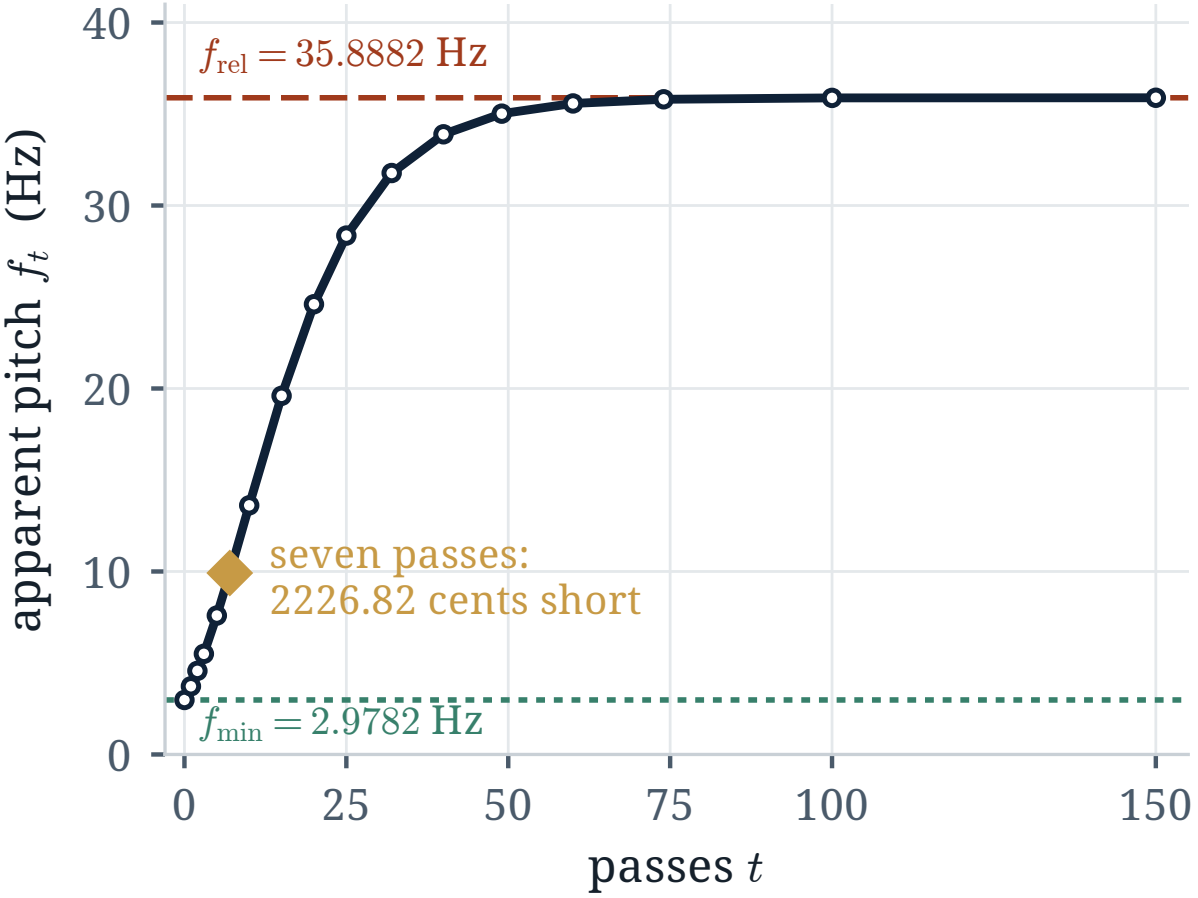


Figure 47.2: **A refinement ladder begun from nothing is a glide from the minimum sustainable frequency to the arrangement's own pitch, and seven passes leave it 2226.82 cents short.** The span is 4309.1864 cents, which is $|D| = 3.5909886634$ octaves exactly. Modelled from the constructed case 2b ($u = 0.692372881355932$, $\eta = 0.44$, $c = 0.40$, $\delta = 0.09$, $N = 12$); the apparent pitch f_t is a constructed reading of a stored quantity and nothing resonates. Built by `figs/src/CCP-047_fig02.py`.

The sense that arrival reveals something one already contained is not a mystical intuition but an accurate report about a formula whose every argument predates the first pass.

Four: the same fixed point is reached from any starting stock. Theorem 4.1 holds for any `Q_0`, and the sign of `Q_0 − Q_inf` does not appear in the rate: a ladder started from a large stock descends at exactly the speed one started from nothing ascends. **The destination is not a summit that some starting positions cannot see.** It is a level, and levels are recognised, not conquered.

The companion paper takes this forward as the definition of destiny [Curzi, 2026]. This paper's contribution is the rate, and the observation that the rate belongs to a parameter that the destination cannot see.

4.15 Counterfeit arrival A fixed point can be positive, stable, and reached at another party's expense. The companion paper names this counterfeit arrival [Curzi, 2026]; this paper can compute it, because cost export is an operation on `c*` and `c*` is in the drive.

CONSTRUCTED CASE — TWO PARTIES SHARING ONE ARRANGEMENT. Labelled constructed; not an observation. Two people, P and R, share the case-2b arrangement, so both draw on the same surplus `F = 2.153633` at `u = 0.692372881355932` and `f_rel = 35.8881768941 Hz`. Their private parameters differ, which is the ordinary situation:

```
P:   eta = 0.44   c = 0.40   delta = 0.09        D = −3.590989   Q      = +6.084427
                                                                  inf
R:   eta = 0.30   c = 0.55   delta = 0.09        D = −0.924189   Q      = +1.067665
                                                                  inf
```

Both are sustainable. R is nearer the edge, because R converts less and pays more — R's break-even cost is `eta_R * F = 0.646090` and R is at 0.55.

Now P exports cost. P declines a share of the maintenance — the scheduling, the repair conversations, the remembering of what was agreed — and R absorbs it. Let `x` be the exported amount, so `c_P -> 0.40 − x` and `c_R -> 0.55 + x`.

x	c_P	c_R	D_P	Q_inf,P	D_R	Q_inf,R	sum of drives	sum of fixed points
0.000	0.400	0.550	-3.590989	+6.084427	-0.924189	+1.067665	+0.643688	+7.152091
0.050	0.350	0.600	-3.918874	+6.639982	-0.443290	+0.512109	**+0.643688**	+7.152091
0.100	0.300	0.650	-4.246759	+7.195538	+0.037608	-0.043446	**+0.643688**	+7.152091
0.115	0.285	0.665	-4.345125	+7.362205	+0.181877	-0.210113	**+0.643688**	+7.152091
0.150	0.250	0.700	-4.574644	+7.751093	+0.518506	-0.599002	**+0.643688**	+7.152091
0.200	0.200	0.750	-4.902530	+8.306649	+0.999405	-1.154558	**+0.643688**	+7.152091

> **CORRECTION NOTE (made visible, not silently applied).** This section originally stated the counterfeit-arrival identity on the *sum of participants' fixed points*. That form is correct only when every participant carries the same `delta`, which the constructed case above happens to satisfy. The general invariant is the **sum of drives**, $\sum_i (\eta_i F_i - c_i)$. The companion paper proves it [Curzi, 2026] as conservation of drive under cost transfer, with the equal-δ case as a corollary and an explicit counterexample. Everywhere below, the rule is stated on drives; the fixed-point form is retained only as the labelled equal-δ corollary.

The sum of the two drives is invariant to every digit across the whole export range. Cost export is an externality in the ordinary sense: a private gain that is exactly a transfer, and that appears as a gain only because the accounts are drawn around one party [Hardin, 1968, Ostrom, 1990]. Each participant's drive is §4.2's `b_i = eta_i F(u*) − c_i`, so that `Q_inf,i = b_i / delta_i`; `b` is the drive and `D` remains the deficit throughout. Cost export produces nothing; it moves an existing quantity from one ledger to the other, entering one drive with a plus and the other with a minus, so `sum_i b_i` cannot move whatever the individual `delta_i` are. **Under equal `delta` — and only then — the sum of fixed points inherits the invariance,** because

sum_i Q_inf,i = (sum_i b_i)/delta; that is why the last column above is also constant, and it is a corollary of the column beside it rather than a rule in its own right. The general statement of what the plain sum does under unequal delta is §4.15a. Either way the invariance is exactly the thing the exporter's own reading cannot show, and that is the point. P's private instrument reports genuine, sustained, increasing improvement. P's deficit falls from -3.590989 to -4.902530, nearly a further octave and a third of headroom. Every number P can see about P is better.

R crosses the sustainability boundary between x = 0.05 and x = 0.10; solving D_R = 0 gives the exact crossing at c_R = eta_R F = 0.646090, i.e. at x = 0.096090.

And then the lag, which is the reason this failure mode is hard to catch. R does not vanish on the pass the sign flips. R's stock decays from its old level toward its new one at (1−delta)^t. At x = 0.15, R's new fixed point is −0.599002 and R starts at +1.067665:

pass	Q_R
0	+1.067665
3	+0.656950
7	+0.262266
10	+0.050025
15	-0.193988
25	-0.441283
40	-0.560675

Solving exactly, Q_R crosses zero at t* = 10.850479 passes. **For eleven passes after the export begins, R still has a positive stock and R's behaviour is still that of a participant** — while R's destination has already been negative from the first pass. Nothing observable changes at the moment the damage is done. This is why the deficit is the instrument and the stock is not: D_R went positive immediately; Q_R took eleven passes to notice.

The end of the arithmetic is the part that matters. The option R retains is exit, and exit is the classical alternative to voice when an arrangement deteriorates [Hirschman, 1970]. Suppose R eventually declines to continue — which is R's structural right under the framework's consent law [Curzi, 2026], and no arrangement in this framework may remove it. The shared arrangement ends. With no partner there is no interaction surplus: u = 0, F = ln(1) = 0, and $f = k = 0.483426$ Hz, the ground essence. P's reading at x = 0.15, with the reduced cost P kept:

```
before R's departure   D  = −4.574644     Q     = +7.751093        f = 35.888177 Hz
                        P                  inf,P

after  R's departure   D  = +1.639426     Q     = −2.777778        f =  0.483426 Hz
                        P                  inf,P
```

P's fixed point falls by 10.528871 units and P's deficit swings 6.214070 octaves the wrong way, and the export that produced the counterfeit reading is the reason. The reduced cost P retained is worth nothing once the surplus it was drawn against is gone.

> **THE DEFINITION, EARNED BY THE ARITHMETIC. A counterfeit fixed point is one that is positive in the exporter's accounts and negative in the accounts of the system that contains the exporter.** It is detectable, and the detection rule is: compute the deficit for every participant separately, and check whether the sum of drives moved. If a participant's fixed point improved and the sum of drives did not move, the gain was exported and nothing was produced.

That rule costs one extra column in the reporting format and it is the only defence available, because the exporter's own instrument is working perfectly and reporting a genuine

improvement in a quantity that genuinely improved. This identity is generalised in the companion paper as conservation of drive under cost transfer [Curzi, 2026] and is stated there as the export test.

4.15a Why the rule is on drives and not on fixed points The two forms are not interchangeable, and the case that separates them is the case in which export is most tempting. With `Q_inf,i = b_i / delta_i` and a transfer of `x` from P to R — so that `b_P -> b_P + x` and `b_R -> b_R − x`:

```
SUM OF DRIVES   sum_i b   =  sum_i ( eta_i F_i − c_i )  =  sum_i delta_i Q
                      i                                                 inf,i
                                                        INVARIANT ALWAYS

SUM OF Q       d/dx [ sum_i Q      ]  =  1/delta   −  1/delta      zero ONLY if the
                            inf,i              P            R    two deltas are equal
```

Run the same P and R with the same export, changing only the decay rates — P at `delta = 0.06`, R at `delta = 0.16`, so that the exporter is the party who decays *slower*:

x	Q_inf,P	Q_inf,R	sum of drives	sum of fixed points
0.000	+9.126640	+0.600561	+0.643688	+9.727202
0.050	+9.959973	+0.288061	+0.643688	+10.248035
0.100	+10.793307	-0.024439	+0.643688	+10.768868
0.150	+11.626640	-0.336939	+0.643688	+11.289702
0.200	+12.459973	-0.649439	+0.643688	+11.810535

Nothing whatever is produced — the drive column does not move a digit — and R crosses into decay between `x = 0.05` and `x = 0.10` exactly as before. But **the plain sum rises, by `1/0.06 − 1/0.16 = 10.416667` per unit exported**, from `+9.727202` to `+11.810535`.

The failure mode is the worst available one. The naive test does not merely fail to detect the export; it reports a gain, and it reports one precisely when cost is pushed onto whoever decays faster — which is the direction the arithmetic already rewards the exporter for choosing. A reading that used the plain sum would return the export check as *passed with improvement* on the pass the damage was done. The companion paper carries this as a named failure mode — the export test run on the plain sum — and works the general case [Curzi, 2026]; neither is repeated here.

4.16 The cost of a rung, and when to stop Theorem 4.1 says every pass closes the same *fraction* of the gap. It therefore says every pass closes a strictly smaller *absolute* amount, because the gap is shrinking. At `delta = 0.09` with an initial gap of 6.084427:

pass	gap remaining	closed by this pass	cumulative
1	5.536828	0.547598	9.00%
3	4.585048	0.453466	6.64%
7	3.144195	0.310964	48.32%
15	1.478565	0.146232	75.70%
25	0.575777	0.056945	90.54%
40	0.139919	0.013838	97.70%
60	0.021218	0.002098	99.65%

A pass costs the operator's time, attention, and the occasion itself. That cost is constant per pass while the return per pass falls geometrically, so the ladder has a finite optimal length. The marginal return at pass `t` is `delta * (1−delta)^t * |Q_0 − Q_inf|`, and setting it equal to a declared threshold `theta` gives

```
t     =  ln( theta / ( delta * |Q  − Q   | ) )  /  ln(1 − delta)
 stop                           0    inf
```

threshold `theta`	`t_stop` (exact)	first pass at or beyond
0.50	0.9642	1
0.25	8.3138	9
0.10	18.0295	19
0.05	25.3791	26
0.01	42.4444	43

The stopping rule is not a completion signal; it is a declared threshold on marginal return. It must be declared before the ladder starts, or it becomes a post-hoc rationalisation of wherever the operator ran out of patience — the stopping rule is a researcher degree of freedom, and an undeclared one is the documented route to an unfalsifiable result [Simmons et al., 2011, Gelman and Loken, 2013, Nosek et al., 2018]. Declaring it in advance is gate G5 of §4.20.

Note what this does *not* say. Whether to divide the ladder into more, smaller occasions is §5's question — splitting a fixed input across occasions rather than repeating a fixed operation — and the answer turns on strict concavity and the surplus ceiling [Curzi, 2026v]. Refinement here holds the per-pass operation fixed and asks only how many times to do it.

The specification

4.17 PROCEDURE REFINE-1 — reading a refinement ladder

```
STEP 1   DECLARE THE OPERATION.  Write down what one pass consists of, in enough
         detail that a third party could perform it. If the description changes
         between passes, this is not refinement (re-indexing is governed elsewhere).

STEP 2   DECLARE THE PERIOD.  One pass = one what? Week, month, session, year.
         delta is a per-cycle rate and is meaningless without its declared cycle.

STEP 3   MEASURE u.  Use the framework's dyadic elicitation procedure; publish the
         raw L2 norms. CHECK FOR THE ONE-HOT TRAP: framings with disjoint
         support are automatically orthogonal and give u = 1 for everyone
         regardless of content.

STEP 4   COMPUTE THE PITCH.  F = ln(1 + 11 u);  f = k (1 + 11 u)^2;  fold into the
         declared band.

STEP 5   ESTIMATE eta AND c.  Report the scale status of each.

STEP 6   COMPUTE THE DEFICIT.  D = 2.885390 (c/eta) − log2(f/k).  This is the
         DESTINATION reading. It does not contain delta and cannot see it.

STEP 7   CHECK THE CEILING.  c/eta <= ln N = 2.484907 at N = 12. Above it, no
```

```
          arrangement of these twelve categories is sustainable and no ladder helps.

STEP 8    ESTIMATE delta.  Report an interval, not a point, and carry that
          interval through every rate figure below.

STEP 9    COMPUTE THE FIXED POINT.  Q_inf = (eta F - c)/delta.  Cross-check against
          Q_inf = -D eta / (2.885390 delta).  The two must agree to 1e-12.

STEP 10   COMPUTE THE RATE.  tau, half-life, t(0.05), t(0.01) from delta alone.
          This is the ARRIVAL reading. It does not contain D and cannot see it.

STEP 11   COMPUTE THE TRAJECTORY.  Q_t from the declared Q_0; report at least the
          passes to 50%, 95% and 99% of the gap.

STEP 12   IF REPORTING THE GLIDE, badge it.  D_t and f_t are Definition 4.4's
          constructed readings of a stored quantity. Nothing resonates.

STEP 13   RUN THE EXPORT CHECK.  For every participant, compute D separately and
          sum the DRIVES (eta F - c). Compare against the sum before any change.
          If a participant's fixed point improved and the sum of drives did not
          move, the gain was exported and nothing was produced (§4.15).
          Do not run this check on the sum of fixed points unless every delta is
          equal; see §4.15a.

STEP 14   REPORT BOTH READINGS.  D with delta, never D alone.
```

4.18 PROCEDURE REFINE-2 — the stopping rule

```
STEP 1    DECLARE theta, the smallest per-pass gain worth one pass's cost, BEFORE
          the first pass. Declare it in the same units as Q.

STEP 2    COMPUTE t_stop = ln( theta / (delta * |Q_0 - Q_inf|) ) / ln(1 - delta).

STEP 3    COMPUTE THE RESOLUTION FLOOR. Find the pass at which the remaining glide
          falls below 34.9756 cents -- the minimum class separation at N = 12.
          Beyond it the ladder is moving less than the lattice can distinguish.

STEP 4    STOP AT THE EARLIER OF STEP 2 AND STEP 3, and record which bound binds.

STEP 5    IF t_stop <= 0, the ladder was not worth beginning: the first pass already
          returns less than theta. Report this. It is a legitimate finding.

STEP 6    DO NOT RE-DERIVE theta AFTER STOPPING. A threshold chosen to justify the
          stopping point carries no information.
```

4.19 Reporting format

```
REFINEMENT LADDER READING                              [CONSTRUCTED / MEASURED]

  OPERATION      .............................................
  PERIOD         one pass = ..........
  PARTICIPANTS   ..........  (one line per participant below)

  DESTINATION READING          -- does not contain delta
    u = ........  (L2 norms: ........)      one-hot trap checked: Y/N
```

```
  F = ........ nats     f = ........ Hz     folded ........ Hz   class ....
  eta = ........   c = ........   c/eta = ........   ceiling ln N = 2.484907
  D = ........ octaves          verdict: ACCUMULATES / HOLDS / DECAYS

ARRIVAL READING              -- does not contain D
  delta = ........ per ..........   [interval: ........ to ........]
  tau = ........   half-life = ........   t(5%) = ........   t(1%) = ........
  Q_0 = ........   Q_inf = ........   cross-check agrees to ........

GLIDE  [constructed reading, Definition 4.4; nothing resonates]
  f_0 = ........ Hz     f_inf = ........ Hz     span = ........ cents = |D| oct
  resolution floor reached at pass ........

STOPPING
  theta declared in advance = ........      t_stop = ........
  binding bound: MARGINAL RETURN / RESOLUTION FLOOR

EXPORT CHECK
  sum of participants' drives (eta F - c), before = ........  after = ........
  moved? Y/N        if N and any Q_inf improved: the gain was exported
  all deltas equal? Y/N   if Y, sum of Q_inf may be reported as a corollary:
  sum of participants' Q_inf, before = ........   after = ........

HONESTY
  essence frequencies, not physical resonances .................. [ ]
  no clinical, diagnostic or therapeutic claim .................. [ ]
  case labelled CONSTRUCTED where applicable .................... [ ]
  scale status of every elicited quantity stated ................. [ ]
  ability to decline preserved ................................... [ ]
```

4.20 Gates

```
G1  OPERATION FIXED.    One pass is specified and does not change between passes.
G2  PERIOD DECLARED.    delta carries its cycle. No bare rate.
G3  BOTH READINGS.      D and delta both reported. Neither alone.
G4  CEILING CHECKED.    c/eta <= ln N verified before any rate is computed.
G5  theta PRE-DECLARED. The stopping threshold precedes the first pass.
G6  EXPORT CHECKED.     Sum of DRIVES computed across all participants. The sum of
                        fixed points only where all deltas are equal (§4.15a).
G7  GLIDE BADGED.       f_t identified as a constructed reading wherever it appears.
G8  CASES LABELLED.     Every worked case marked CONSTRUCTED.
```

The failure modes

#	Name	What it looks like	What it actually is	Detection
FM 4.1	The ascending ladder	“Each pass takes us higher”	A contraction converging on a fixed level	Compute `Q_inf` before pass 1; compare to the reported trajectory
FM 4.2	Speed as evidence	“This works fast, so it works”	Large `delta`; a leaky system reaching its level	`delta` is a loss term. Read `D`, which cannot see it
FM 4.3	The moving destination	Recomputing `Q_inf` after each pass and reporting the movement	The parameters were changed; this is re-indexing, not refinement	Gate G1: was the operation fixed?
FM 4.4	Seven for its own sake	A count inherited from the tradition and applied to any `delta`	Seven passes close 48.32% at `delta = 0.09` and 6.79% at `delta = 0.01`	§4.11’s table
FM 4.5	The completion signal	Waiting for the pass on which the work finishes	There is none; `(1-delta)^t > 0` always	Procedure REFINE-2’s declared `theta`
FM 4.6	The retro-fitted threshold	`theta` chosen after stopping, to justify stopping	An unfalsifiable rationalisation	Gate G5
FM 4.7	Deficit without delta	Reporting `D` alone as “the reading”	`D` cannot see decay at all	Gate G3
FM 4.8	Delta without deficit	Reporting a fast half-life as good news	The rate cannot see the destination’s sign	Gate G3
FM 4.9	Counterfeit arrival	One party’s readings improve steadily and genuinely	Cost export; sum of drives unchanged	Gate G6, §4.15
FM 4.10	The eleven-pass illusion	“They are still participating, so they are fine”	The stock lags the destination by `t*` passes	Compute each participant’s `D`, not each participant’s `Q_t`
FM 4.11	Refining below resolution	Passes 60 through 300, diligently performed	Movement smaller than 34.9756 cents, the lattice minimum at `N = 12`	Procedure REFINE-2 step 3

#	Name	What it looks like	What it actually is	Detection
FM 4.12	Trust as a ladder accelerator	“Add trust and we will get there sooner”	Trust lowers `delta`; the ladder gets *slower*, and if the numerator is negative the destination gets worse	the ordering rule; §4.12 rows two and four
FM 4.13	The zero-leak ladder	A model with `delta = 0`	No contraction, no fixed point, unbounded growth in the model and rapid divergence from anything real	Gate G2 forces the leak to be named
FM 4.14	One-hot refinement	Every participant’s `u` comes out at 1.000000	Disjoint support: automatic orthogonality, degenerate surplus	Procedure REFINE-1 step 3
FM 4.15	The export check run on the plain sum	The sum of fixed points is used as the conserved quantity under unequal `delta`	Not an invariant: the sum moves by `1/delta_P − 1/delta_R` per unit exported, and *rises* when the exporter decays slower	§4.15a; the full treatment is Curzi [2026]

Four of these deserve more than a table row.

FM 4.2, speed as evidence, is the most common and the most inverted. The intuition that a practice producing rapid change is a powerful one is exactly backwards here. Rapidity is `delta`, `delta` is loss, and the destination `Q_inf = b/delta` is *divided* by it. A fast ladder and a slow ladder with the same `D` do not go to the same place: the fast one goes to a nearer place, quickly. Rows one and two of §4.12 differ by a factor of ten in both directions at once.

FM 4.9 and FM 4.10 are one failure mode observed at two lags, and the pairing is what makes it dangerous. The export check (G6) catches it on pass one; the stock does not reveal it until pass eleven in the worked case, by which time the exporter has eleven passes of improving private readings to point at. A protocol that monitors stocks rather than deficits finds the problem after it has become a fact.

FM 4.13 deserves naming because it is the tradition’s characteristic model, and the framework has already shown that $\delta = 0$ is a pathology rather than an ideal: no contraction, no unique fixed point, and an error-persistence column of all ones, so that no mistake ever decays out of the stock [Curzi, 2026a]. Descriptions of a refined substance that “no longer changes” are descriptions of `delta = 0`. There the map is an isometry, Banach does not apply, there is no fixed point, and `Q_t = Q_0 + b t` grows without bound — which is not permanence but the opposite, a quantity with no level at all. **The tradition’s image of an incorruptible end product is, read as a model, the one setting under which the whole apparatus stops working.** The honest version is that `delta` small is a durable disposition, `delta` zero is not a disposition at all.

FM 4.12 is the ordering rule arriving in this section’s vocabulary, and it is worth stating in the strongest form the arithmetic supports. Adding trust to a failing arrangement in order to make the refinement go faster does the reverse of both: it makes the ladder slower and

the destination worse. The carried case is the demonstration — 2a at `delta = 0.09` arrives at `−5.0119395464`; at `delta = 0.05` it arrives, more slowly, at `−9.0215`. **Fix the numerator first** [Curzi, 2026q]; this section shows what ignoring it does to the trajectory as well as to the destination.

Worked examples

4.21 The carried dyad, case 2b, run as a ladder **CONSTRUCTED. Not an observation of any person or pair.** Case 2b: the pair after B has taken on two states A does not hold and the pair have adopted a decision log.

Steps 1–2. One pass is one week: the pair's standing weekly review with the decision log open. The operation does not change between passes.

Steps 3–4. `u = 0.692372881355932`, elicited by magnitude estimation with raw L2 norms published. The one-hot check passes: the two vectors share support on states 4, 6 and 9, so the orthogonality trap does not apply. Then

```
F = ln(1 + 11 * 0.692372881355932) = 2.153633 nats
f = k * (1 + 11 * 0.692372881355932)^2 = 35.8881768941 Hz
```

Folding into the 220–440 Hz band requires $j = 3$: `35.8881768941 * 8 = 287.105415` Hz.

Steps 5–7. `eta = 0.44`, `c = 0.40`, so `c/eta = 0.909091`, comfortably below the ceiling `ln 12 = 2.484907`. Then `f_min = k e^(2 * 0.909091) = 2.9782153324 Hz`.

Step 6 — the destination reading.

```
D = 2.885390 * 0.909091 − log2(35.8881768941 / 0.483425732988414)
  = 2.623082 − 6.214071
  = −3.5909886634 octaves                      ACCUMULATES
```

Steps 8–10 — the arrival reading. `delta = 0.09` per week.

```
Q     = (0.44 * 2.153633 − 0.40)/0.09  =  +6.0844267427
 inf
cross-check:  −D * eta / (2.885390 * delta) = +6.0844267427   agrees to 1e−12

tau = 10.6033 weeks    half-life = 7.3496 weeks
t(5%) = 31.7645 weeks  t(1%) = 48.8298 weeks   t(0.1%) = 73.2447 weeks
```

Step 11 — the trajectory, from `Q_0 = 0`, drive `b = 0.547598406839`:

week	Q_t	fraction of gap remaining	closed
0	0.000000	1.000000	0.00%
1	0.547598	0.910000	9.00%
7	2.940232	0.516761	48.32%
15	4.605861	0.243008	75.70%
32	5.786888	0.048902	95.11%
49	6.024551	0.009857	99.01%
74	6.078761	0.000931	99.91%

Step 12 — the glide, badged as Definition 4.4's constructed reading: from `f_0 = f_min = 2.978215 Hz` to `f_rel = 35.888177 Hz`, a span of 4309.1864 cents, which is `|D|` exactly. The full table is §4.13.

Step 13 — the export check. Both members' `c*` are unchanged from the arrangement as declared; the sum of drives is unmoved. No export.

Step 14 — both readings, together:

```
DESTINATION   D = −3.5909886634 octaves.  ACCUMULATES.  Independent of delta.
ARRIVAL       delta = 0.09/week.  Half-life 7.3496 weeks.  Within 5% at week 32.
```

Neither number is the reading. The pair is both going somewhere good and going to take about eight months to get most of the way there, and a report giving only the first would be a report a reader could not act on.

4.22 Case 2a — the ladder that perfects a decay

CONSTRUCTED. The same pair as found, before the intervention. `u = 0.068734665264634`, `eta = 0.30`, `c = 0.62`, `delta = 0.09`.

```
F = 0.563084802746 nats        f = 1.4907987150 Hz        f_min = 30.1587698405 Hz
D = +4.3384201973 octaves      DECAYS
Q     = −5.0119395464          drive b = −0.451074535976
 inf
```

The rate constants are identical to §4.21's, because `delta` is: `tau = 10.6033`, half-life 7.3496 weeks, within 5% at week 32. **The same ladder, at the same speed, converging on `−5.01` instead of `+6.08`.** The trajectory from `Q_0 = 0`:

week	Q_t
0	0.000000
1	-0.451075
7	-2.420838
15	-3.793076
32	-4.766824
49	-4.962534
74	-5.007270

The glide runs the other way. `f_0 = f_min = 30.158770 Hz` and `f_t` descends, because `D_t` rises from 0 toward `+4.3384201973`. **A refinement ladder on an unsustainable arrangement is a descending glide of 4.338420 octaves below the minimum sustainable frequency**, executed with perfect regularity. The practice is not failing. It is succeeding at converging on a negative destination.

And now the intervention a reader reaches for, which the carried case exists to pre-empt — add trust, `delta` 0.09 -> 0.05:

```
delta = 0.09    Q     = −5.0119395464       D = +4.338420    half-life  7.3496 weeks
                 inf
delta = 0.05    Q     = −9.0215             D = +4.338420    half-life 13.5134 weeks
                 inf
```

The destination nearly doubled in the wrong direction, the ladder became 84% slower, and the deficit did not move. Three separate readings, three separate stories, one intervention. A reader with only `D` would conclude nothing happened. A reader with only the half-life would conclude the practice had become more patient. A reader with `Q_inf` would see the damage. **Report all three.** That is what §4.19's format requires, and why it requires it.

The repair is the framework's order of operations [Curzi, 2026q]: fix the numerator, then lower the loss. Case 2b is what that looks like — `u` raised from 0.068735 to 0.692373 and `c` cut from 0.62 to 0.40 first, after which trust becomes an amplifier of a positive sign and `delta` 0.09 -> 0.05 takes `Q_inf` from +6.0844267427 to +10.9520.

4.23 The counterfeit ladder, in report format **CONSTRUCTED.** P and R of §4.15, at export `x = 0.15`, read at pass 8 — after the export has begun and before R's stock has crossed zero.

```
REFINEMENT LADDER READING                                        [CONSTRUCTED]

  OPERATION      the shared weekly review, as declared in §4.21
  PERIOD         one pass = one week
  PARTICIPANTS   P, R

  DESTINATION READING
    u = 0.692372881355932    one-hot trap checked: Y (shared support on 4, 6, 9)
    F = 2.153633 nats    f = 35.8881768941 Hz    folded 287.105415 Hz
    P:  eta 0.44   c 0.250   c/eta 0.568182   D = -4.574644   ACCUMULATES
    R:  eta 0.30   c 0.700   c/eta 2.333333   D = +0.518506   DECAYS
    ceiling ln N = 2.484907; R at 2.333333 is inside it but close

  ARRIVAL READING
    delta = 0.09 per week (both)   half-life 7.3496 weeks
    P:  Q_inf = +7.751093
    R:  Q_inf = -0.599002    Q_R at pass 8 = +0.216073, still positive
    R's stock crosses zero at t* = 10.850479 passes

  EXPORT CHECK
    sum of drives before export = +0.643688
    sum of drives after  export = +0.643688
    moved? NO  ->  NOTHING WAS PRODUCED. The improvement in P's reading is
                   a transfer, and R's deficit has already changed sign.
    all deltas equal? YES (0.09 both) -> the corollary form also applies:
    sum of Q_inf before export = +7.152091   after export = +7.152091

  HONESTY
    essence frequencies, not physical resonances ................... [x]
    no clinical, diagnostic or therapeutic claim ................... [x]
    case labelled CONSTRUCTED ..................................... [x]
    ability to decline preserved -- R may end the arrangement ..... [x]
```

Every private number P can compute is better than it was, every one of them is correct, and the arrangement is now scheduled to end. That is what counterfeit arrival is, and the only line in the report that catches it is the one comparing two sums.

4.24 What this section establishes, and what Section 5 uses Established here, and used in Sections 5 and 6 and in the destiny paper [Curzi, 2026]:

```
Thm 4.1   |Q_t - Q_inf| = (1-delta)^t |Q_0 - Q_inf|, an equality, not a bound
Thm 4.2   Q_inf contains no t: refinement cannot move the destination
Thm 4.3   f_rel contains no t: refinement cannot move the pitch or the deficit
Def 4.4   the running deficit D_t and apparent pitch f_t -- CONSTRUCTED readings
Thm 4.5   the octave ladder: |D_t - D| = (1-delta)^t |D_0 - D|
Cor 4.6   the glide: from f_min to f_rel, spanning exactly |D| octaves
the three constants of arrival, all functions of delta alone
the two invariances: constant cost per halving, constant cost per decade
what seven passes buy, as a function of delta, and the delta that makes seven enough
counterfeit arrival: defined, computed, and given a one-line detection rule
  on the SUM OF DRIVES (§4.15a; general form in the destiny paper). The sum of fixed
  points is the equal-delta corollary only.
```

```
PROCEDURE REFINE-1, PROCEDURE REFINE-2, the reporting format, eight gates,
  fifteen failure modes
```

The negative results are the load-bearing ones. Refinement does not ascend, does not re-index, does not change a pitch, does not move a fixed point, and does not finish. What it does is close a proportion of a gap, at a rate set by the loss term, toward a place computed from four parameters none of which is effort.

The two readings do not substitute for each other: `delta` governs how fast a ladder arrives and cannot appear in the deficit; the deficit governs whether arriving is worth anything and cannot see `delta`. Report both or report nothing.

Section 5 turns to what the tradition got right, and cadence is on its list. The temptation there is to read this section's result as a vindication of repetition, and it is not one. This section holds the operation fixed and counts passes. Section 5 asks a different question — whether a fixed total input is better delivered at once or in instalments — and the answer to that one is *not* that instalments deliver more. They deliver exactly the same integrated total, to every digit. What they buy is something else entirely, and Section 5 states it precisely.

5 What the tradition got right, computed

This section and the next are one argument in two halves, and the second half is what makes the first half usable. Section 6 is the counterpart: it takes the tradition's largest claim, kills it with arithmetic, and then states exactly what survives. **A list of things a tradition got right, published alone, is an advertisement.** Published beside a list of what it got catastrophically wrong, computed to the same standard, it becomes an audit — so nothing in this section should be read without the next one.

Four things are on this section's list. Each was arrived at independently, by people with no instruments and no statistics, and each corresponds to something the formalism of this framework computes:

```
1  STAGED PROCESS CONTROL     an operation split into named stages with a check at each
2  CADENCE                    scheduled repetition rather than one great effort
3  CORRESPONDENCE TABLES      finite relational graphs over a counted category set
4  THE OPERATOR IN THE LOOP   the practitioner's own state as part of the apparatus
```

The second carries a correction that must be stated before anything else, because the popular version of it is false and the framework has already computed the falsity: **the integrated effect of one large input and the same total delivered periodically is identical to every digit — 433.3333 in all three schedules tested.** That is linear time-invariance and it is not a matter of opinion. Cadence wins for two other reasons entirely, both of them computed in §5.9 through §5.13, and this section states them rather than the folklore.

The apparatus for this section is the framework's two-column concordance method, with the role of each column stated in every block rather than assumed [Curzi, 2026].

The law

5.1 The concordance apparatus, with roles stated

```
(+1)  PHYSICS PROOF     what scientific measurement confirms
(-1)  METAPHYSICS       what contemplative or traditional practice records
(0p)  CONTRA-PROOF      how the two integrate at the Neutral Point

C1  STRUCTURAL CORRESPONDENCE   a deep shared pattern, not a surface resemblance
```

```
  C2  PREDICTIVE POWER          the bridge must generate a prediction that can fail
  C3  PRACTICAL UTILITY         it must be usable on both sides
```

And the refinement that keeps the method honest, restated here in full because this section is built entirely out of concordances:

> **The (+1) column carries EVIDENCE — what is measured, replicated, falsifiable. The (-1) column carries PRECEDENT** — that the idea was arrived at independently, recorded, and preserved, often for centuries, by people without the instruments. **Two sources agreeing does not double the evidence when only one of them is evidence.**

The framework's own reference apparatus fixes the classification this method depends on — four items of established mathematics, seven of established psychology, and exactly three declared postulates, with every protocol stated to run on the first two and steer by the third [Curzi, 2026]. What independent arrival does establish is real and worth having: that a structure was findable by more than one route, which is a genuine argument that the structure is in the territory and not only in the instrument. It is not an argument that the structure is correct. Section 6 supplies the arithmetic for why the precedent column cannot be promoted — under success-only publication the likelihood ratio of testimonial evidence is exactly 1.000000 — and this section's every block is written under that constraint.

Only bridges satisfying all three criteria are presented as established. Speculative bridges are marked speculative in the sentence that states them. **And a method that never rejects anything is not a method: this section rejects three claims in print, and Section 6 rejects the largest one of all.**

5.2 CONCORDANCE C5.1 — STAGED PROCESS CONTROL

```
CONCORDANCE C5.1 — AN OPERATION SHOULD BE DIVIDED INTO NAMED STAGES WITH A CHECK AT EACH

  CLAIM        That dividing a transformation into discrete, named, individually
               inspected stages produces a better outcome than performing it as one
               undivided operation with a single check at the end.

  (+1) EVIDENCE  -- this column carries measurement
               Statistical process control. Shewhart's Economic Control of Quality of
               Manufactured Product (1931) established in-process measurement against
               control limits rather than terminal inspection, and Deming's Out of the
               Crisis (1986) made the argument that inspection at the end cannot build
               quality in. NEITHER WORK IS IN THIS PAPER'S BIBLIOGRAPHIC DATABASE and
               both are named here without a key; the operations-management treatment
               that IS available states the same yield-and-rework structure
               [Hopp and Spearman 2011; Goldratt 1984]. Separately, the unit-operations
               decomposition of chemical engineering -- distillation, filtration,
               crystallisation as separately characterised stages -- is the working
               form of the same idea.
               WHAT IT MEASURED: that per-stage detection removes error before it
               propagates, and that the yield difference is large and computable
               (§5.8 computes it: 47.83% to 93.21% on a seven-stage process).

  (-1) PRECEDENT  -- this column carries independent arrival, not proof
               The Greco-Egyptian and Latin alchemical corpora describe transformation
               as an ordered series of named operations, each with its own apparatus and
               its own sign of completion. Zosimos of Panopolis (3rd-4th c. CE) describes
               distillation apparatus -- the tribikos, the kerotakis -- attributed in that
```

```
              corpus to Maria the Jewess. The historiography that establishes these as
              laboratory practice rather than allegory is Newman and Principe, "Alchemy
              vs. Chemistry: The Etymological Origins of a Historiographic Mistake"
              (Early Science and Medicine, 1998), and Principe, "The Secrets of Alchemy"
              (2013), which reproduces documented recipes in a modern laboratory.
              WHAT IT RECORDS: that staging, per-stage apparatus, and a per-stage
              completion criterion were arrived at and preserved for well over a
              millennium before any of it was justified statistically.

  (0p) INTEGRATION
              Both are describing error containment. A stage boundary is a place where
              an error can be caught before the next operation makes it unrecoverable.
              The tradition's "sign" -- a colour, a consistency, a smell -- is a
              control-limit test with an ordinal readout and no stated tolerance.

  CRITERIA    C1 PASS -- the shared pattern is per-stage detection, not the surface
                         fact that both use lists.
              C2 PASS -- it predicts that yield rises with per-stage detection
                         probability, monotonically, and §5.8 computes the curve.
              C3 PASS -- usable on both sides.

  STATUS      ESTABLISHED. The precedent is genuine and the evidence is independent
              of it. The tradition is not the reason to believe the claim; Shewhart is.

  BOUNDED BY  THAT STAGING IS RIGHT DOES NOT MAKE ANY PARTICULAR STAGE ORDER RIGHT.
              The framework's operation audit runs a canonical
              fifteenth-century twelve-gate schedule under this model and finds the
              only numerator-acting gate NINTH of twelve -- the canonical sequence lowers loss before it r
              numerator, which is exactly the order the framework's
              elasticities show accelerates a failure.
              Section 6.5 states that rejection in full. Staging: established.
              This tradition's stage ORDER: rejected.
```

5.3 CONCORDANCE C5.2 — cadence, with the invariance stated exactly

```
CONCORDANCE C5.2 — SCHEDULED REPETITION BEATS ONE GREAT EFFORT

  CLAIM AS THE TRADITION STATES IT
              That the work must be done on a schedule -- the fire kept even, the
              matter circulated daily, the office said at fixed hours -- and that this
              is superior to one concentrated effort of equal total magnitude.

  THE NAIVE VERSION OF THE CLAIM IS FALSE AND THE FRAMEWORK COMPUTES THE FALSITY
              The cadence result, recomputed here in §5.9: the INTEGRATED effect of
              one large impulse and the same total delivered periodically is IDENTICAL
              to every digit -- 433.3333 in all three schedules, equal to B/delta.
              That is linear time-invariance. Anyone claiming that spreading the same
              total input over more occasions produces more total stored order, in a
              linear model, has not done the arithmetic.

  (+1) EVIDENCE  -- this column carries measurement
              (a) THE FORMAL RESULT, computed in this work, not borrowed: the steady
                  state differs absolutely (0 for an impulse, m/delta for sustained
                  periodic forcing); the ceiling S <= ln N clips a large one-off input,
                  costing 97.2% of the nominal budget in the worked annual occasion
```

at N = 8, and 90.44% in the N = 12 register recomputed in §5.11;
and the cadence theorem -- strict concavity alone means splitting an
input across occasions strictly dominates delivering it at once.
(b) THE SPACING EFFECT, an independent empirical literature: Ebbinghaus
(1885) first reported it; Cepeda, Pashler, Vul, Wixted and Rohrer
(Psychological Bulletin, 2006) is the quantitative synthesis; and
Roediger and Karpicke (Psychological Science, 2006) established the
related testing effect. NONE OF THESE THREE IS IN THIS PAPER'S
BIBLIOGRAPHIC DATABASE. They are named without keys and no claim in
this paper rests on them.
HONEST NOTE, REQUIRED: the spacing literature is about human memory
encoding and the cadence theorem is a statement about concave
response functions. They are not the same claim and must not be presented as
mutual confirmation. Each stands on its own.

(-1) PRECEDENT -- this column carries independent arrival, not proof
Fixed-hour liturgical practice; the alchemical instruction to keep the
fire even and to circulate the matter over months rather than force it.
Durkheim's Elementary Forms of the Religious Life (1912) on periodic
assembly, and Rappaport's Ritual and Religion in the Making of Humanity
(1999) on invariance and periodicity as the defining features of ritual.
NEITHER IS IN THIS PAPER'S BIBLIOGRAPHIC DATABASE; both are named
without keys, and they carry precedent, not evidence.
WHAT IT RECORDS: that scheduled recurrence was selected for, and
preserved, by institutions that had no way to compute why it worked.

(0p) INTEGRATION
Periodic forcing of a leaky integrator. The correct statement, and the
only one this framework endorses: CADENCE BEATS INTENSITY BECAUSE OF STRICT
CONCAVITY AND THE SURPLUS CEILING, NOT BECAUSE IT DELIVERS MORE TOTAL
INPUT. It does not. The integrated total is identical.

CRITERIA C1 PASS. C2 PASS -- it predicts the clipping loss as a function of
cadence and N, computed in §5.11, and the optimum division as a
function of per-occasion cost, computed in §5.12. C3 PASS.

STATUS ESTABLISHED, WITH THE FOLKLORE'S STATED REASON REJECTED. The practice
is right and the traditional explanation of why is wrong. Section 6's
method -- disprove the strong claim with arithmetic, preserve what is
genuinely there -- applies here as much as it does there.

5.4 CONCORDANCE C5.3 — CORRESPONDENCE TABLES AS RELATIONAL GRAPHS

CONCORDANCE C5.3 — A CORRESPONDENCE TABLE IS A FINITE RELATIONAL GRAPH

CLAIM That the tradition's tables of correspondence -- ten nodes and twenty-two
paths, seven metals, four elements, a seventy-eight element structured
deck -- are, described accurately and without mysticism, finite relational
structures over counted category sets, which is exactly the
object the framework's category-set validity condition requires.

(+1) EVIDENCE -- this column carries measurement
Graph theory, which is not in dispute [Newman 2018; Barabasi 2016]. A ten-node, twenty-two-e
has a computable density against the complete graph's forty-five edges:
0.488889. And the star-polygon result \citep{ccp:CCP-028}, which is

number theory and is exact: a figure {N/k} drawn by stepping k positions around N points decomposes into exactly gcd(N,k) separate closed paths, each of length N/gcd(N,k). The hexagram {6/2} is literally two separate triangles; {7/2} is one continuous path. §5.14 computes the table.
WHAT IT MEASURED: that these structures have exact arithmetic invariants.

(−1) PRECEDENT -- this column carries independent arrival, not proof
The tables themselves, preserved across the Hermetic, kabbalistic and alchemical corpora, and the Ars Memoriae tradition of spatial mnemonics that used them (Yates, "The Art of Memory", 1966).
WHAT IT RECORDS: that counted, exhaustive, exclusive category schemes with declared relations between members were constructed and transmitted for centuries. That is the form V2 asks for, arrived at independently.

(0p) INTEGRATION
A correspondence table supplies the PARTITION that V2 requires and often supplies the ORDER that the house condition V7 requires. What it does not supply is V1, a state encoding -- a map from domain items to normalised vectors -- and §5.15 grades two classical tables against all of them and finds them wanting in exactly that place.

CRITERIA C1 PASS. C2 PASS -- the gcd result predicts the number of separate strokes in any such figure, and is falsifiable by drawing one.
C3 PASS -- the partition is directly usable as a category set.

STATUS ESTABLISHED AS A DESCRIPTION OF FORM. Explicitly NOT established as a claim that any particular table's contents are correct. §5.15 rejects two classical tables as psychological indices while retaining their form.

5.5 CONCORDANCE C5.4 — THE OPERATOR IS PART OF THE APPARATUS

CONCORDANCE C5.4 — THE PRACTITIONER'S OWN STATE IS PART OF THE INSTRUMENT

CLAIM That the state of the person conducting an operation materially affects its outcome, and that preparing the operator is therefore part of preparing the apparatus rather than a piety attached to it.

(+1) EVIDENCE -- this column carries measurement
(a) Experimenter expectancy effects. Rosenthal and Fode (1963) reported
that experimenters told their rats were "maze-bright" obtained better maze performance than those told the same rats were "maze-dull". The blinding requirement in trial design exists because of this class of result. It must be presented with its contested history: the 1968 classroom study [Rosenthal and Jacobson 1968] drew serious methodological criticism, the effect sizes are disputed [McCambridge et al. 2014], and the best long-run assessment finds classroom self-fulfilling prophecies real but small, usually non-accumulating and often dissipating [Jussim and Harber 2005].
(b) Tacit knowledge and the experimenter's regress: Polanyi's Personal Knowledge (1958) and Collins's Changing Order (1985), the latter documenting that laboratories could not build a working TEA laser from the published description alone and required personal contact with a group that already had one. NEITHER IS IN THIS PAPER'S BIBLIOGRAPHIC DATABASE; both are named here without keys.

```
            (c) THE FRAMEWORK'S OWN MEASUREMENTS: the shuffled-input control
                run off the shelf is anti-conservative at 2.2x nominal alpha (0.108
                against 0.05); the matched-autoregressive comparison alone false-passed
                35.8% of the time; against fabricated timing the full four-baseline
                suite passed 52.7% of the time; against a shared slow driver it passed
                98.35% of the time. The conclusion recorded there is that DESIGN
                CONTROLS -- blinding, channel separation, prospective sampling, a
                yoked control -- are required, and every one of those is a constraint
                on the operator rather than on the material.
            (d) THE FORMAL RESULT, computed in §5.16, whose structure matches the
                framework's channel-gain postulate -- E = G.B + eps with the
                forbidding clause G = 0 => E = 0 for all B, badged as a postulate and
                paired with the storage law's eta = 0 by explicit correspondence
                [CCP-034]: if the operator is a
                participant, the operator's own state vector enters u. An operator
                whose framings mirror the subject's contributes u = 0 exactly
                and pays the coordination cost anyway. In the carried case this LOWERS
                the arrangement's pitch by 536.4074 cents and moves the octave deficit
                from +4.338420 to +4.785426, or to +6.372391 once the link cost is
                priced.

(-1) PRECEDENT  -- this column carries independent arrival, not proof
            The standing insistence across the alchemical corpus that the artisan must
            be prepared, unhurried and attentive, and that the work fails otherwise;
            the parallel insistence in contemplative traditions that the state of the
            one performing a practice is not separable from the practice.
            WHAT IT RECORDS: that the operator was treated as a component. It records
            nothing about mechanism, and the tradition's proposed mechanisms -- moral
            worthiness, purity, election -- are not carried forward by this framework.

(0p) INTEGRATION
            Three distinct real mechanisms are available and none is supernatural:
            expectancy effects on measurement, tacit skill in execution, and -- in the
            relational case -- the operator's literal membership in the set whose
            non-parallelism is being computed. The third is the one this framework adds and
            it is arithmetic.

CRITERIA    C1 PASS.  C2 PASS -- §5.16 predicts the sign and size of the pitch change
            from adding an operator of a given framing, and the prediction is
            checkable by measuring u.  C3 PASS.

STATUS      ESTABLISHED AS A STRUCTURAL CLAIM; THE TRADITION'S MECHANISM REJECTED.
            "The operator is part of the apparatus" is correct. "The operator must be
            morally pure" is not carried, is not derivable from anything here, and is
            not endorsed.
```

One structural note before the rejections. The mechanism C5.4 admits is not that belief produces outcomes, but that a belief which is *grounded* has a channel to act through. The framework's conditional postulate states this in magnitude form as $E = G \cdot B + \varepsilon$ with the forbidding clause $G = 0 \Rightarrow E = 0$ for all B, badged as a postulate and paired with the storage law's $\eta = 0$ by explicit correspondence [Curzi, 2026ab]. That is the correct shape for the operator claim too: an operator's state can only act through a channel the arrangement already has, and where the channel gain is zero no amount of operator preparation produces anything. **The tradition's proposed mechanism — moral worthiness — is not a channel, and is not carried forward.**

5.6 Two further rejections, in print A method that never rejects is not a method. C5.1's block has already rejected one claim — this tradition's canonical stage *order*, on the operation audit against the ordering rule — and two more that a sympathetic reading might have admitted are rejected here.

```
REJECTION C5.5 — THE COLOUR SEQUENCE AS A NECESSARY ORDER

  CLAIM        That the transformation must pass through blackening, whitening, yellowing
               and reddening in that order, and that the order is a law of the process.

  GRADING      C1 MARGINAL. Staging is real (C5.1) and the colours are real observations
               of real material changes. But the claim under test is not "there are
               stages"; it is "these stages, in this order, necessarily".
               C2 FAILED, AND THIS IS DECISIVE. The stages are identified BY their
               colours. An operation that does not blacken is therefore not at the
               blackening stage, by definition, and the sequence cannot be observed to
               fail. A claim that cannot fail generates no prediction.
               C3 PARTIAL. Useful as a shared vocabulary for reporting where one is.

  STATUS       REJECTED AS AN ORDERING LAW. Retained only as staging vocabulary, with
               the circularity stated wherever it is used. This is the second
               concordance rejection in the essence-transformation material and it is recorded in
               Appendix D.

REJECTION C5.6 — "AS ABOVE, SO BELOW" AS A SCALE-INVARIANCE THEOREM

  CLAIM        That the same laws govern the small and the large, so a result obtained at
               one scale transfers to another.

  GRADING      C1 PARTIAL, AND THE PARTIALITY IS MEASURED. The storage law IS scale-free
               in FORM -- the framework's master correspondence table is exactly that
               claim, and it is why one instrument reads a mind and a commons. But
               the framework's own scope audit names nine disanalogies between the
               internal and external scopes and computes three of them: a mean-field decay error reaching 9
               two nine-member groups with IDENTICAL u = 0.5999 and S = 1.7578 losing
               77% versus 21% of surplus to a three-person departure, so that aggregate
               statistics hide the structure that decides the outcome; and a 90.1x
               period-normalisation trap when comparing scopes on mismatched clocks.
               C2 FAILED AS STATED. The maxim predicts that any result transfers. The
               computed disanalogies are counterexamples.
               C3 DANGEROUS. Used as licence it produces exactly the 90.1x error.

  STATUS       REJECTED AS STATED. What survives, and is used throughout this framework, is
               the precise version: THE CORRESPONDENCE IS OF DYNAMICS AND GEOMETRY, NOT
               OF AGENCY, and every transfer between scopes must re-derive its own
               parameters and declare its own clock.
```

5.7 Honesty clauses, restated in place **(1) NO SUPERNATURAL CLAIM.** Nothing in this section asserts any efficacy of symbol, any physical effect of a figure, or any result not derivable from the formalism. The computing-map reading of the old notation is a metaphor and keeps that badge.

(2) ESSENCE, NOT RESONANCE. Every frequency here is an essence frequency: a position on a logarithmic surplus scale. No substance, person or tissue is claimed to resonate at any value.

(3) NO CLINICAL CLAIM. The twelve states are not diagnoses; nothing here is a treatment.

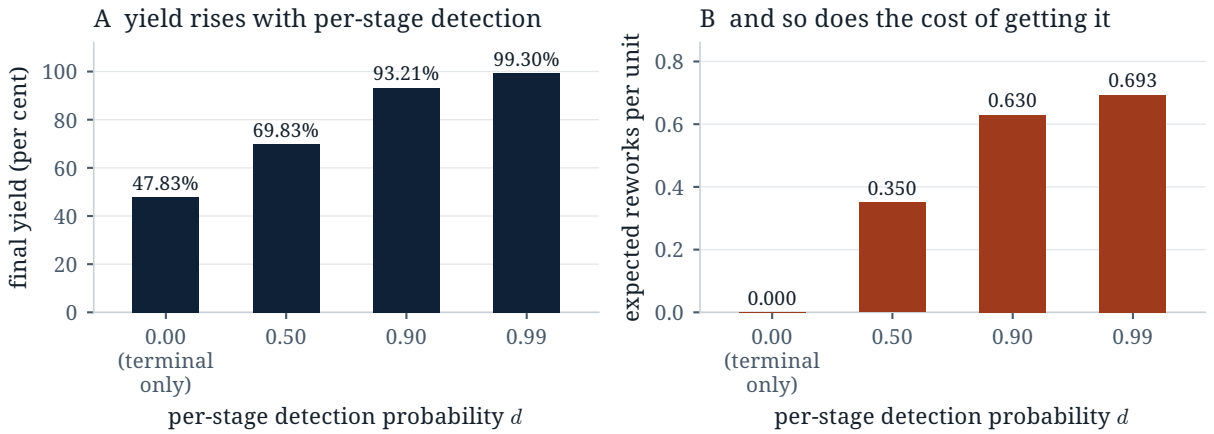


Figure 47.3: **Per-stage checking at 90 per cent detection raises final yield from 47.83 per cent to 93.21 per cent, and the cost of getting it is 0.63 reworks per unit.** Two panels rather than one dual axis, because yield and rework are different measures. Modelled from a seven-stage process with per-stage defect probability $q = 0.10$, per-stage escape $q(1-d)$ and yield $(1-q(1-d))^7$; the process is constructed. Built by `figs/src/CCP-047_fig03.py`.

(4) PRECEDENT IS NOT EVIDENCE. Stated in every block above and restated here. Section 6 computes why: the likelihood ratio of a success-only testimonial record is exactly 1.000000.

(5) ALL CASES CONSTRUCTED. Every worked case in this section is constructed and labelled.

(6) SCALE STATUS. n is a count and is absolute-scaled; intensities `w` must be magnitude-estimated with raw L2 norms published.

(7) BOTH READINGS. Every reading in this section reports the octave deficit `D` *and* `delta`. `D` contains no `delta` and cannot see decay; `delta` cannot see the deficit.

The engineering

5.8 Staged process control, computed **CONSTRUCTED.** A seven-stage process. Each stage introduces a defect with probability `q = 0.10`. A check after each stage detects a defect introduced at that stage with probability `d` and it is reworked; undetected defects propagate and spoil the product. Per-stage escape is `q(1-d)` and final yield is `(1 - q(1-d))^7`.

per-stage detection `d`	escape per stage	final yield	expected reworks per unit
0.00 (terminal check only)	0.1000	**47.83%**	0.000
0.50	0.0500	69.83%	0.350
0.90	0.0100	**93.21%**	0.630
0.99	0.0010	99.30%	0.693

Moving from a single terminal inspection to a per-stage check at ninety per cent detection raises yield from 47.83% to 93.21%, at a cost of 0.63 reworks per unit. That is the whole of C5.1 in four rows. The tradition's contribution is not the arithmetic, which it never had; it is that the arrangement the arithmetic recommends is the arrangement it used.

Note what the table also shows, because the honest reading is the useful one. Rework is not free and it rises with detection. At `d = 0.99` the yield gain over `d = 0.90` is 6.09 percentage points for an extra 0.063 reworks; whether that trade is worth taking is an economic question the model does not answer. Staging is not free, and the tradition's habit of treating each stage as a rite with its own preparation is, in this reading, an accurate estimate that stage boundaries are expensive.

5.9 The cadence invariance, part one: the integrated effect is identical Reproduced here rather than cited, because the popular claim is the opposite and a citation is not enough. The reference setup [Curzi, 2026v] is: `delta = 0.06` per week, a total budget `B = 26.0` delivered in one calendar year, and three schedules — one impulse of 26.0 at week 0; 0.5 per week for 52 weeks; 2.0 every fourth week for 52 weeks. Summing `Q_t` over 4,000 periods:

```
impulse       433.3333
weekly        433.3333
four-weekly   433.3333
B / delta  =  433.3333

maximum pairwise difference across the three schedules:  3.41e-13
```

Identical to every digit, to the limit of floating-point summation. This is the defining property of a linear time-invariant system: the integral of the response to a given total input is `B/delta` regardless of how that input is distributed in time.

Restated as this framework's position and binding on everything that follows: cadence does not deliver more. Any argument of the form "spreading it out delivers more total benefit" is wrong, and the correct arguments are the two that follow.

5.10 The cadence invariance, part two: the steady state differs absolutely The integrals agree; the destinations do not. An impulse decays to nothing — its steady state is exactly `0`. Sustained periodic forcing at mean rate `m` per period holds a level of `m/delta`. Those are not the same system observed at different speeds; they are different systems.

The trajectory difference in the reference worked case has a crossover at week 24: at week 23 the impulse is still ahead, 6.6648 against 6.3254, and at week 24 the cadence is ahead, 6.4458 against 6.2649, and it never trails again while the budget lasts. The impulse wins early and loses permanently, which is precisely the shape that makes the folklore feel true and the arithmetic necessary.

5.11 The cadence invariance, part three: the ceiling clips the impulse The second real argument is the surplus ceiling. $S_t \leq \ln N$ is the framework's normalisation bound, carried into the frequency register as the ceiling $c/\eta \leq \ln N$ [Curzi, 2026]. One interaction cannot generate more surplus than the log of the number of distinguishable categories, however much is poured into it. Everything above the ceiling is not delivered late; it is not delivered.

The framework's worked annual occasion ran at $N = 8$, $\eta = 0.35$, giving a per-occasion cap of `ln 8 = 2.0794` nats. **Against a nominal annual budget of 26 nats, the annual rite delivers 2.8% of what the weekly practice delivers: the clipping costs 97.2% of the nominal budget.** That figure is the cadence paper's and is quoted exactly [Curzi, 2026v].

Recomputed in this framework's register — `N = 12`, cap `ln 12 = 2.484907` nats, `eta = 0.44`, the same 26-nat nominal annual budget delivered at a nominal 0.5 nats per week:

cadence `P` (weeks)	burst wanted	delivered	clipped	annual delivered	`eta` × annual	% of weekly
1	0.5000	0.5000	no	26.0000	11.4400	100.00%

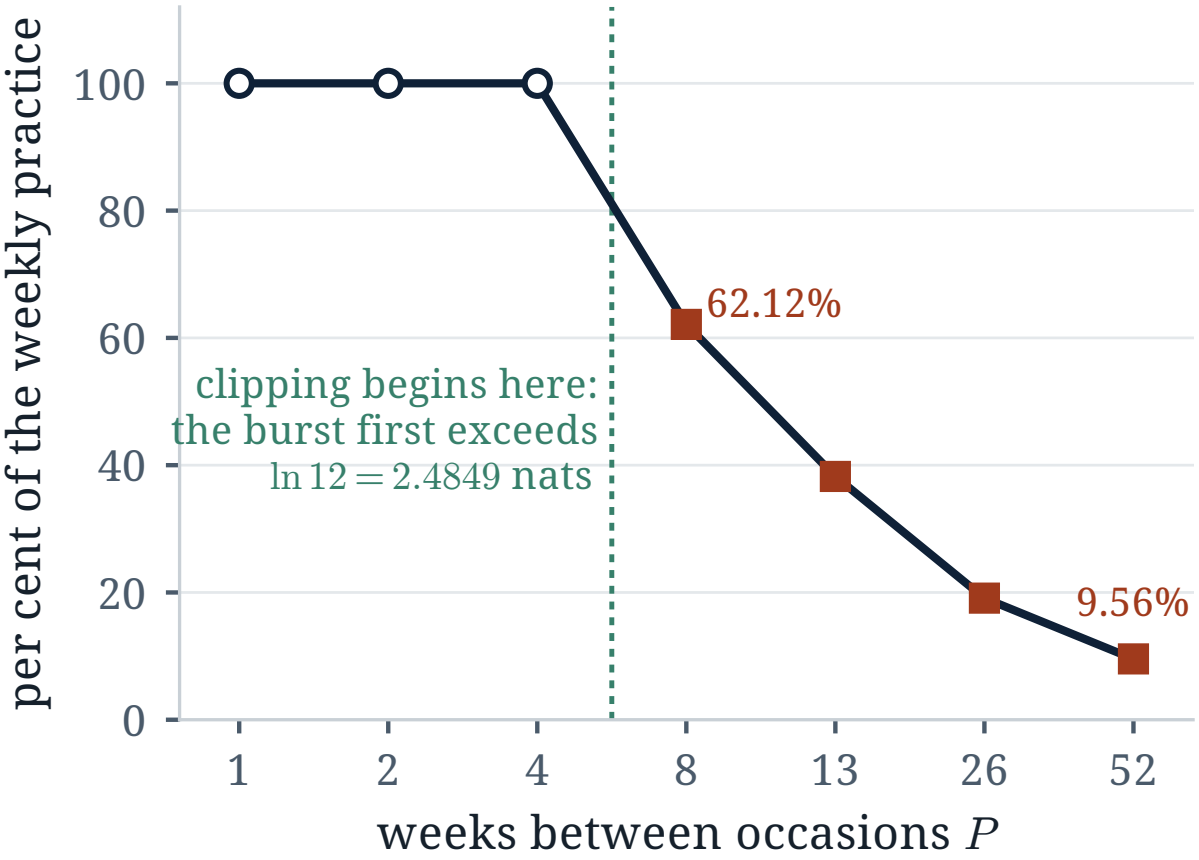


Figure 47.4: **A concentrated occasion loses nothing until its burst exceeds the ceiling, and then loses almost everything: the annual occasion delivers 9.56 per cent of the weekly practice from an identical nominal budget.** Modelled at $N = 12$ with a per-occasion cap of $\ln 12 = 2.484907$ nats, $\eta = 0.44$, and a nominal annual budget of 26 nats delivered at 0.5 nats per week. Built by `figs/src/CCP-047_fig04.py`.

cadence P (weeks)	burst wanted	delivered	clipped	annual delivered	eta × annual	% of weekly
2	1.0000	1.0000	no	26.0000	11.4400	100.00%
4	2.0000	2.0000	no	26.0000	11.4400	100.00%
8	4.0000	2.4849	**yes**	16.1519	7.1068	62.12%
13	6.5000	2.4849	**yes**	9.9396	4.3734	38.23%
26	13.0000	2.4849	**yes**	4.9698	2.1867	19.11%
52	26.0000	2.4849	**yes**	2.4849	1.0934	**9.56%**

At N = 12 the annual rite delivers 9.56% of the weekly practice from an identical nominal budget — a loss of 90.44% rather than the 97.2% of the $N = 8$ case. The difference between the two figures is not a discrepancy; it is the ceiling, and it is a direct demonstration of lever L4. Raising the category count from 8 to 12 raises the cap from 2.0794 to 2.4849 nats, and the concentrated rite recovers a third of what it was losing. A richer category set does not merely raise the maximum sustainable frequency; it also makes concentrated delivery less wasteful, because there is more room under the cap.

Note also where the clipping starts. At N = 12 a fortnightly and even a monthly cadence lose nothing at all in this worked case; the loss begins at P = 8 weeks. **"More often" is not uniformly better.** It is better only once the burst exceeds the cap, and below that point cadence

buys nothing through this channel. The other channel, concavity, is where the remaining advantage lives.

5.12 The cadence invariance, part four: strict concavity, which needs no ceiling at all The cadence theorem [Curzi, 2026v] — for g concave with $g(0) = 0$, $D(n) = n\,g(T/n)$ is nondecreasing in n, so dividing a fixed budget over more occasions never delivers less — is the honest core of the cadence argument, because it does not depend on the ceiling binding. `F(u) = ln(1 + (N−1)u)` is strictly concave through the origin, and for a concave response through the origin, splitting an input across occasions strictly dominates delivering it at once. No clipping required.

CONSTRUCTED. A fixed budget of `T = 12` hours of interaction, divided into `n` occasions of `L = T/n` hours each. Non-parallelism rises linearly with session length to saturation at `L* = 8` hours: `u(L) = min(1, L/L*)`. `N = 12`, `eta = 0.44`, and a fixed per-occasion cost `c_0 = 0.15` regardless of length — convening, travelling, settling, closing.

n	L (h)	u	F(u)	eta F	total	−n c_0	net
1	12.0000	1.0000	2.4849	1.0934	1.0934	-0.1500	0.9434
2	6.0000	0.7500	2.2246	0.9788	1.9577	-0.3000	1.6577
3	4.0000	0.5000	1.8718	0.8236	2.4708	-0.4500	2.0208
4	3.0000	0.3750	1.6341	0.7190	2.8761	-0.6000	2.2761
6	2.0000	0.2500	1.3218	0.5816	3.4894	-0.9000	2.5894
8	1.5000	0.1875	1.1192	0.4925	3.9397	-1.2000	2.7397
12	**1.0000**	**0.1250**	**0.8650**	**0.3806**	**4.5672**	**-1.8000**	**2.7672**
24	0.5000	0.0625	0.5232	0.2302	5.5255	-3.6000	1.9255
48	0.2500	0.0312	0.2955	0.1300	6.2402	-7.2000	-0.9598

The delivered-total column (headed *total*) is monotone increasing exactly as the cadence theorem requires [Curzi, 2026v], rising from 1.0934 to 6.2402 — **a factor of 5.7074 from the same twelve hours, purely by division, with no ceiling involved and no extra input.**

Two further readings, both of which the folklore misses.

The net column peaks at $n = 12$ and falls away on both sides. The per-occasion fixed cost is the same term that makes added communication paths expensive in project work, where the count grows as $n(n-1)/2$ [Brooks, 1975, Conway, 1968]. Concavity alone would recommend infinitely many infinitesimal occasions, which is obviously wrong, and the term that stops it is the per-occasion fixed cost. **The optimum is a function of `c_0` and of nothing else in this model**, which means "how often should we meet" is answerable, and the answer is an arithmetic consequence of convening cost.

At `n = 1` the single twelve-hour occasion wastes four hours before any other effect is considered, because `u` saturates at `L* = 8`. The longest session in the table is the only one with idle budget in it.

5.13 The honest counter-case The cadence paper kept a counter-case in print [Curzi, 2026v] and this paper keeps it too, because a section titled *What Alchemy Got Right* that reported only confirmations would be selecting.

Above a threshold that no realistic cadence reaches, the impulse wins. And the durability such a schedule would require is computable: **A worked decennial observance — one occasion every ten years — requires a trace half-life of 5.63 years for the stored order to survive between occasions.** A practice whose retention half-life is measured in weeks cannot be run on a decade.

This is the same arithmetic as §4's, arriving from the cadence side. **The interval between occasions must be short relative to the half-life `ln(0.5)/ln(1−delta)`, and every claim that a rare observance suffices is a claim about `delta` that can be checked.**

5.14 Correspondence tables, computed as graphs The claim of C5.3 is descriptive and the descriptions are arithmetic.

```
ten-node / twenty-two-path scheme
    nodes 10, edges 22, complete graph 45 edges, density 0.488889

the twelve-state ladder of this framework
    nodes 12, edges 66, complete graph 66 edges, density 1.000000

a seventy-eight element structured deck
    78 = 22 + 4 x 14;  distinct pairs = 3,003
```

The density figures are the useful comparison. **This framework's own category scheme is a complete graph** — every state stands in a computed interval relation to every other, because the intervals are ratios of squares and no pair is undefined. The classical ten-node scheme specifies fewer than half its available relations, and the unspecified ones are not "absent"; they are simply not part of the table. That is a difference of kind between a correspondence table and a metric structure, and it is the honest limit on the analogy.

The one exact arithmetic invariant the tradition's figures do possess is the star-polygon result [Curzi, 2026v], which is number theory and is true. Stepping `k` positions around `N` points produces exactly `gcd(N,k)` separate closed paths of length `N/gcd(N,k)`:

`N`	`k` and the resulting (strokes, length)
6	{6/1} (1, 6) · {6/2} **(2, 3)** · {6/3} (3, 2)
7	{7/1} (1, 7) · {7/2} (1, 7) · {7/3} (1, 7)
8	{8/1} (1, 8) · {8/2} (2, 4) · {8/3} (1, 8) · {8/4} (4, 2)
9	{9/1} (1, 9) · {9/2} (1, 9) · {9/3} (3, 3) · {9/4} (1, 9)
10	{10/1} (1,10) · {10/2} (2, 5) · {10/3} (1,10) · {10/4} (2, 5) · {10/5} (5, 2)
11	every `k` from 1 to 5 gives (1, 11)
12	{12/1} (1,12) · {12/2} (2, 6) · {12/3} (3, 4) · {12/4} (4, 3) · {12/5} (1,12) · {12/6} (6, 2)

The six-pointed figure is literally two separate triangles and the seven-pointed one is a single unbroken path, and which of the two obtains is decided by `gcd(N,k)` and by nothing else. At `N = 11`, a prime, every step size yields one continuous path; at `N = 12` the strokes take every divisor value. This is exact, checkable with a pencil, and carries no claim whatever about meaning.

5.15 Grading two classical tables against the validity conditions

The form passes. The contents must be graded, and the framework's validity conditions — here labelled V1 (a state encoding exists), V2 (the categories partition a counted set) and V7 (a declared ordinal rule) — are the grading instrument.

The ceiling $F \leq \ln N$ is the same bound that requisite-variety arguments impose on any regulator [Ashby, 1956, 1958], and choosing a small N spends it in advance.

THE FOUR ELEMENTS, as a candidate index. Four categories, exhaustive and exclusive by stipulation, counted. Traditionally ordered by density or subtlety: earth, water, air, fire.

- **V2, category assignment: PASSES.** It partitions, it is counted, `N = 4`.
- **V7, declared ordinal: PASSES, weakly.** The density ordering is publishable and a third party can apply it. It is a real ordinal rule, which is more than several modern candidates managed in the framework's index-validation work.

- **V1, state encoding: FAILS.** There is no map from psychological states to normalised vectors over these four categories. A feeling does not have a determinate location among them, and no procedure in the corpus supplies one.
- **The ceiling consequence, and it is severe.** At `N = 4` the maximum surplus is `ln 4 = 1.386294` nats — 55.79% of what `N = 12` allows — and the ceiling on the sustainable cost ratio falls to `c/eta <= 1.386294`, i.e. `f_min <= k·16 = 7.734812 Hz`. **Choosing a four-category scheme discards 3.169925 octaves of available ceiling before any measurement is taken.**
- **VERDICT: REJECTED as a psychological essence index.** Retained as a demonstration that the tradition's *form* satisfies V2 and V7, which is the point C5.3 makes.

THE SEVEN METALS AND THEIR PLANETARY CORRESPONDENCES. Seven categories, counted, with a declared traditional order.

- **V2: PASSES.** V7: passes, on the traditional ordering.
- **V1: FAILS, and fails at the first step.** The correspondence is between substances and celestial bodies. Neither is a psychological state, and no encoding of states into this scheme exists. Applying `f = k n^2` anyway would produce 0.483426, 1.933703, 4.350832, 7.734812, 12.085643, 17.403326 and 5.687861 Hz — seven perfectly definite numbers that mean nothing, because **the frequencies inherit exactly the defensibility of the index, and an index that fails V1 yields frequencies that are arithmetic without content.**
- **Ceiling consequence:** `ln 7 = 1.945910` nats against `ln 12 = 2.484907`, a loss of 0.538997 nats — **1.555215 octaves of ceiling, 21.69% of the available surplus, discarded by the choice of seven.**
- **VERDICT: REJECTED as a psychological essence index. Retained as a mnemonic index**, which is exactly the disposition the audit template recommends and §6 generalises: kill the strong reading with arithmetic, keep the weak one that is genuinely there.

5.16 The operator in the apparatus, computed CONSTRUCTED. Not an observation of any person.

The carried dyad's exact state vectors are $A = \{4{:}0.7,\ 6{:}0.5,\ 9{:}0.4\}$ and `B = {4:0.6, 6:0.55, 9:0.3, 10:0.2}`, normalised. Raw L2 norms published per the framework's magnitude reporting rule: `|A| = 0.948683`, `|B| = 0.890225`. Then `cos = 0.965021` and `u(A,B) = 0.068734665264634`, which is case 2a to fifteen decimal places.

Now add a third participant O — a facilitator, a coach, a third member of the working group — and compute the three-member mean pairwise `u`, its surplus and its pitch. Cost is read at `eta = 0.30`, `c = 0.62`, `delta = 0.09`, with a per-link coordination cost $\kappa = 0.055$ over three links where priced.

operator	|O|	u(A,O)	u(B,O)	mean u	f (Hz)	D (κ=0)	D (κ=0.055)
— (dyad alone)	—	—	—	0.068735	1.490799	+4.338420	—
O1 mirroring A	0.948683	**0.000000**	0.068735	0.045823	**1.093596**	**+4.785426**	**+6.372391**
O2 disjoint support	0.901388	1.000000	1.000000	0.689578	35.632544	-0.240618	+1.346347
O3 overlapping, distinct	0.829156	0.818626	0.829263	0.572208	25.721455	+0.229605	+1.816570
O4 strongly distinct	0.913783	0.980838	0.953720	0.667764	33.668595	-0.158826	+1.428139

Three results, and one trap that the table is deliberately made to contain.

O1, the mirroring operator, is the demonstration. An operator whose framings duplicate A's contributes `u(A,O) = 0.000000` exactly — because $u = 1 - (\hat{x} \cdot \hat{y})^2$ depends only on $|\rho|$, and a perfect mirror gives $\rho = 1$. Mean `u` falls from 0.068735 to 0.045823, the pitch falls from 1.490799 Hz to 1.093596 Hz — **a drop of 536.4074 cents, most of a perfect fourth** — and the octave deficit worsens from +4.338420 to +4.785426, or to +6.372391 once the three links are priced at `kappa`. **An operator who agrees with the subject adds cost and no surplus, and the instrument says so in cents.**

O2 is the trap, and it is reported rather than quietly fixed. O2's vector has *disjoint support* from A's and B's — it occupies states 3, 7 and 11, which neither of them holds. Disjoint support

means automatic orthogonality, which means `u = 1.000000` against both, which is the one-hot degeneracy the framework's encoding rules warn about. **The reading `f = 35.632544 Hz` for O2 is an artefact of the encoding and not a fact about the operator**, and a section presenting it as the "good" case would be presenting a bug as a result. O3 and O4 are the corrected versions: distinct framings that nonetheless share support with the pair, giving `u` values that are high without being degenerate.

O3 and O4 are the real finding. A genuinely differing operator with overlapping support takes the arrangement from `D = +4.338420` to `+0.229605` before coordination cost and to `−0.158826` in O4's case — **from a 4.34-octave shortfall to the sustainability boundary, by adding one person who does not agree.** Priced at `kappa = 0.055` per link, all four operators leave the arrangement short, which is the correct and deflating answer: adding a third party to a badly-configured dyad is not a repair, and the quadratic growth of coordination cost against logarithmic surplus is why [Brooks, 1975, Conway, 1968].

Two standing clauses bind this table. The first: every `f` here is a RELATIONSHIP frequency and says nothing about the participants' own state frequencies — the framework's operation audit [Curzi, 2026] shows four pairs spanning thousands of cents in state terms returning identical u and f_{rel} to twelve decimal places, because `u` reads the angle and not the occupied indices. Neither reading may be inferred from the other.

The second: these are D readings and D contains no δ. At `delta = 0.09` the O1 arrangement arrives at `Q_inf = −5.528341`; at `delta = 0.05` it arrives at `−9.951014`, and `D` stays at +4.785426 in both. **Report `D` and `delta`.** A reading with only one of them cannot distinguish an arrangement failing slowly from one failing fast, and the deficit is blind to the difference.

The specification

5.17 PROCEDURE ALCH-R — reading a traditional claim as engineering

```
STEP 1   STATE THE CLAIM IN THE TRADITION'S OWN TERMS, without improving it.
         Improving a claim before testing it is how a tradition gets defended
         rather than audited.

STEP 2   IDENTIFY WHAT KIND OF CLAIM IT IS.
            a structural claim  (about the FORM of an arrangement)
            a causal claim      (about what produces what)
            an evidential claim (about what the record establishes)
         Evidential claims go to Section 6 and the survivorship arithmetic.

STEP 3   FIND THE (+1) COLUMN.  What measurement, if any, bears on it? Cite what
         it measured, not what it is popularly taken to show. If there is no
         measurement, say so; the claim is then precedent only and cannot be
         promoted.

STEP 4   FIND THE (−1) COLUMN.  What did the tradition actually record, as
         opposed to what a modern reader would like it to have meant?

STEP 5   STATE EACH COLUMN'S ROLE IN THE BLOCK.  (+1) evidence, (−1) precedent.
         Never implicitly. Every block, every time.

STEP 6   GRADE AGAINST C1, C2, C3.  C2 is the one that fails most often, and it
         fails for a specific reason: circular identification of the stages or
         outcomes (see REJECTION C5.5).

STEP 7   COMPUTE SOMETHING.  A bridge that generates no number has not passed C2.
```

```
            Where the claim touches an index, grade it against V1–V7 as well.

    STEP 8   REPORT THE CEILING COST OF ANY CATEGORY SCHEME ADOPTED.
             c/eta <= ln N, and choosing a small N discards octaves before measurement begins.

    STEP 9   IF THE CLAIM SURVIVES, state what survives and what was stripped.
             If it does not, PUBLISH THE REJECTION with its failing criterion.

    STEP 10  REPORT D AND delta TOGETHER for every arrangement read.
```

5.18 The concordance block template

```
CONCORDANCE C__.__ — [THE CLAIM, IN ONE LINE]

  CLAIM        ..................................................................
  (+1) EVIDENCE       -- this column carries measurement
               source ..............................................................
               WHAT IT MEASURED ....................................................
  (−1) PRECEDENT      -- this column carries independent arrival, not proof
               source ..............................................................
               WHAT IT RECORDS .....................................................
  (0p) INTEGRATION
               ..................................................................
  CRITERIA     C1 PASS/FAIL because ......   C2 PASS/FAIL because ......
               C3 PASS/FAIL because ......
  COMPUTED     ..................................................................
  STATUS       ESTABLISHED / ESTABLISHED WITH THE STATED MECHANISM REJECTED /
               REJECTED (with failing criterion named)
  SPECULATIVE? if yes, the word "speculative" appears in the CLAIM line itself
```

5.19 Gates

```
    G1  ROLES STATED.      Every block names what each column carries. No exceptions.
    G2  NO PROMOTION.      Precedent never appears as evidence, in any sentence.
    G3  SOMETHING COMPUTED. Every established bridge carries a number.
    G4  REJECTIONS PUBLISHED. At least one rejection per concordance set, in print.
    G5  V-GRADING.         Any category scheme adopted is graded against V1–V7.
    G6  CEILING PRICED.    ln N reported for every scheme, and the octaves foregone.
    G7  BOTH READINGS.     D and delta together, always.
    G8  CASES LABELLED.    Every worked case marked CONSTRUCTED.
    G9  INVARIANCE FIRST. Any cadence claim carries the LTI identity before anything else.
```

The failure modes

#	Name	What it looks like	What it actually is	Detection
FM 5.1	Precedent promoted	“Both modern science and the ancients agree”	One column is evidence and one is not	G1, G2

#	Name	What it looks like	What it actually is	Detection
FM 5.2	The naive cadence claim	“Spread it out and it delivers more”	False: integrated totals are identical, 433.3333	§5.9; G9
FM 5.3	Cadence for its own sake	Meeting more often with no clipping and no concavity gain	Below the cap, more occasions buy nothing through the ceiling channel	§5.11: at `N = 12` loss begins only at `P = 8` weeks
FM 5.4	Infinite division	“So we should meet daily, or hourly”	Concavity alone recommends it; per-occasion cost stops it	§5.12: net peaks at `n = 12`
FM 5.5	The saturating session	A twelve-hour offsite	`u` saturates at `L* = 8`; four hours deliver nothing	§5.12 row one
FM 5.6	The circular stage	“It has not blackened, so it is not at that stage”	Unfalsifiable by construction	REJECTION C5.5, criterion C2
FM 5.7	Table as territory	Treating a correspondence table’s contents as measured	The form satisfies V2; the contents fail V1	§5.15
FM 5.8	The small category set	Adopting a four- or seven-fold scheme because it is traditional	3.169925 or 1.555215 octaves of ceiling discarded before measuring	G6
FM 5.9	Numbers without an index	Computing `k n^2` on a scheme that fails V1	Arithmetic without content	§5.15
FM 5.10	The sycophantic operator	A facilitator who reflects the group back to itself	`u = 0` exactly; cost paid, no surplus; -536.4074 cents	§5.16 row O1
FM 5.11	The one-hot operator	An operator whose reading comes out at `u = 1.000000`	Disjoint support, automatic orthogonality, degenerate	§5.16 row O2
FM 5.12	Scale transfer by maxim	“As above so below, so this result carries”	92% mean-field error by `t = 12`; 90.1x clock trap	REJECTION C5.6
FM 5.13	Surrogates instead of design	Adding shuffle controls to fix an operator problem	The shuffle is anti-conservative (2.2x) and blind to timing (52.7%) and shared drivers (98.35%)	the survey work’s control audit; design controls only
FM 5.14	Deficit without `delta`	Reporting an operator’s effect on `D` alone	`D` cannot see decay at all	G7

Three of these deserve more than a row.

FM 5.2 is the reason this section states the invariance before it states anything else. The claim that scheduled practice delivers more is the single most repeated piece of practical folklore in the literature this framework draws on, and it is false in a linear model, by an identity, to every digit. The two true arguments are better than the false one — a steady state instead of a decay, and a factor of 5.7074 from division alone — and they survive scrutiny, which the false one does not.

FM 5.10 is the sharpest thing this section computes and it deserves to be said plainly. The tradition insisted that the operator's state matters. The formalism agrees and adds a number: an operator who mirrors the subject contributes exactly zero surplus and costs a link. The squaring makes it worse still — perfect opposition gives the same $u = 0$ as perfect agreement, so **the contrarian facilitator and the agreeable one sound the identical note, and it is the ground essence.** The useful operator is the one whose framings differ from the subject's *and share support with them*, which is O3 and O4, and which is a design requirement rather than a personality.

FM 5.13 records where the tradition's instinct was right and its method unavailable. "Prepare the operator" is correct; "prepare the operator morally" is not a control. The modern equivalents are blinding, channel separation, prospective sampling and a yoked control, and the framework's survey work measured what happens without them [Curzi, 2026i]: a surrogate suite that passes 52.7% of the time against fabricated timing and 98.35% of the time against a shared slow driver. **More surrogates do not fix an operator problem. Design does.**

Worked examples

5.20 The annual rite against the weekly practice, read in full **CONSTRUCTED.** A group commits to one annual retreat carrying a nominal 26 nats of interaction surplus for the year, against the alternative of a weekly practice carrying 0.5 nats per week for 52 weeks. `N = 12`, `eta = 0.44`, `delta = 0.09` per week, `c = 0.40` per week.

Step 1 — the LTI check, run first, per gate G9. If the response were linear and unclipped, the integrated effect of the two schedules would be identical. §5.9 confirms it at 433.3333 in the reference setup. **No part of the argument below rests on the rite delivering less total input, because it does not.**

Step 2 — the ceiling. `ln 12 = 2.484907` nats per occasion. The weekly burst of 0.5 is far below it. The annual burst of 26.0 is more than ten times it.

Step 3 — delivered surplus.

```
weekly      52 occasions x min(0.5,  2.484907) =  26.0000 nats   ->  eta x = 11.4400
annual       1 occasion  x min(26.0, 2.484907) =   2.4849 nats   ->  eta x =  1.0934

the rite delivers 9.56% of the practice.  Loss to clipping: 90.44%.
```

Step 4 — the steady state. The weekly practice is sustained periodic forcing and holds a level. The rite is an impulse and decays toward zero between occasions; whether it holds anything depends entirely on whether the year's decay is survivable. At `delta = 0.09` per week the half-life is 7.3496 weeks, so **after 52 weeks the fraction of a single impulse remaining is `0.91^52 = 0.007403` — seven-tenths of one per cent.** The next rite starts from essentially nothing.

Step 5 — the concavity check. Independently of the ceiling, §5.12 shows that dividing a fixed budget raises delivered total monotonically, by a factor of 5.7074 across the range tested, until per-occasion cost dominates.

Step 6 — the counter-case, per §5.13. The rite is not absurd; it is a bet on `delta`. It works if the retention half-life is long relative to the interval. The decennial case requires a 5.63-year

trace half-life. **A weekly-scale `delta` of 0.09 gives 7.3496 weeks, short of a year by a factor of about seven**, so the bet loses here — but the bet is stated, and it is checkable.

Step 7 — both readings. For the weekly arrangement at `u = 0.692372881355932`, `eta = 0.44`, `c = 0.40`: `f = 35.8881768941 Hz`, `f_min = 2.9782153324 Hz`, **`D = −3.5909886634` octaves, `delta = 0.09` per week, `Q_inf = +6.0844267427`, half-life 7.3496 weeks.** Both numbers reported.

5.21 The concordance register for this section

#	Claim	(+1) role	(-1) role	C1	C2	C3	Status
C5.1	Staged process with per-stage checks	evidence: Shewhart 1931; Deming 1986	precedent: Zosimos corpus; Newman & Principe 1998; Principe 2013	P	P	P	**ESTABLISHED**
C5.2	Cadence beats intensity	evidence: this work's concavity + ceiling results; Cepeda et al. 2006	precedent: fixed-hour practice; Durkheim 1912; Rappaport 1999	P	P	P	**ESTABLISHED, folklore's reason REJECTED**
C5.3	Correspondence tables are relational graphs	evidence: graph theory; the gcd theorem	precedent: the tables themselves; Yates 1966	P	P	P	**ESTABLISHED as form only**
C5.4	The operator is part of the apparatus	evidence: expectancy effects; tacit knowledge; the control audit; §5.16	precedent: the pure-artisan requirement	P	P	P	**ESTABLISHED, mechanism REJECTED**
C5.5	The colour sequence is a necessary order	—	precedent only	M	**F**	Pa	**REJECTED (C2, circular)**
C5.6	As above so below, as scale invariance	evidence against: three computed disanalogies	precedent: the maxim	Pa	**F**	D	**REJECTED as stated**
C5.1b	This tradition's canonical stage ORDER	evidence against: the ordering elasticities; the operation audit	precedent: a canonical twelve-gate schedule of 1471	P	**F**	D	**REJECTED (order only; staging stands)**

Three of seven rejected, and every rejection is recorded in Appendix D with its failing criterion. The third is the sharpest of them, because it is the framework auditing the tradition on the tradition's own published material: the practice of staging survives and this tradition's canonical order does not.

5.22 What this section establishes, and what Section 6 uses

```
four claims audited and established, each with its columns' roles stated
three claims rejected in print, with the failing criterion named, including the
  canonical stage ORDER (the operation audit against the ordering rule -- staging stands,
  the order does not)
the cadence invariance stated exactly: integrated effect IDENTICAL
                    (433.3333, three schedules)
                    cadence wins by strict concavity and the ln N ceiling clip
                    97.2% loss at N = 8;  90.44% at N = 12 (§5.11)
                    factor 5.7074 from division alone, no ceiling required
staged control priced: 47.83% -> 93.21% yield at d = 0.90
the gcd invariant tabulated for N = 3..12
two classical tables graded against V1–V7 and rejected as indices, retained as form
the operator computed into u: mirroring costs 536.4074 cents and 0.447 octaves
PROCEDURE ALCH-R, the block template, nine gates, fourteen failure modes
```

The pattern across all six blocks is one pattern. **The tradition's practices are frequently right and its explanations of them frequently wrong**, and separating the two requires arithmetic the tradition did not have and could not have had. Staging is right and the colour law is not. Cadence is right and "it delivers more" is not. The tables are the right shape and their

contents are not an index. The operator is part of the apparatus and moral purity is not the mechanism.

That separation is only credible if the same instrument is turned on the tradition's largest claim and allowed to destroy it. **Section 6 does exactly that.** Material transmutation of the elements is dead by an energy-scale argument of six orders of magnitude; the tradition's record is survivorship-filtered and carries a likelihood ratio of exactly 1.000000; secrecy made the errors permanent; no control was ever run. And then Section 6 states what remains, which is everything in this section — and it remains because the audit was real.

6 What the tradition got wrong, computed

Section 5 listed four things the tradition got right and computed each of them. This section is the reason that list can be trusted, and it exists because an account that only audits in one direction is not auditing.

The method is not new here. It is the framework's standing audit template, established against numerical-value schemes and adopted for the whole traditional corpus:

Disprove the strong claim with arithmetic. Preserve what is genuinely there.

That case is the model and is worth restating exactly, because the shape of it governs every section below. The evidential reading of gematria — that two words sharing a numerical value are thereby related — is dead. Measured over 234,454 dictionary words, the English simple-value map yields 305 distinct outputs, destroys 10.47 bits of the 17.84 that identify a word, and collides at a rate of 99.995%; in an 8,000-word corpus against a thousand reachable values the probability that any given word has a value-twin is 0.999666. **A shared value is the default outcome, not a signal.** And the mnemonic reading — a hash that assigns an item to a retrieval bucket — is untouched by any of that, because a hash with a terrible collision rate is still a perfectly good index.

One reading dies by arithmetic; the other stands, undamaged, in the same paragraph. That is what this section does five more times, at larger scale.

The five charges:

```
1  MATERIAL TRANSMUTATION OF THE ELEMENTS by chemical means -- short by six orders
   of magnitude in energy, and the shortfall is computable from first principles

2  THE RECORD IS SURVIVORSHIP-FILTERED -- likelihood ratio exactly 1.000000,
   which means the corpus does not move the posterior AT ALL

3  SECRECY IS AN EPISTEMIC FAILURE MODE -- unshared results cannot be checked,
   so errors persist for centuries, and the persistence is computable

4  THERE WERE NO CONTROLS -- and the control that would have settled the central
   question was the cheapest instrument available in the entire laboratory

5  THE CANONICAL ORDER IS BACKWARDS -- the tradition's own published schedule
   lowers loss before it fixes the numerator, which is exactly the order the
   framework's elasticities show accelerates a failure
```

The law

6.1 Charge one: material transmutation by chemical means The claim is that base metal can be converted into gold by an operation conducted in a furnace with a crucible, an alembic, reagents, and heat.

THE ENERGY-SCALE ARGUMENT. The identity of a chemical element is set by the number of protons in its nucleus. Lead is `Z = 82`. Gold is `Z = 79`. Converting one into the other requires removing three protons from a nucleus. Chemistry does not touch nuclei; it rearranges the outer electrons, and the two operations sit at energies six orders of magnitude apart.

```
a strong single covalent bond              ~  4 eV                 (order of magnitude)
mean nuclear binding energy per nucleon    ~  8 MeV                (order of magnitude)

ratio                                      =  8e6 eV / 4 eV  =  2,000,000
```

Sharper, because it requires no averaging. The Coulomb barrier a single proton must cross to reach a lead-208 nucleus, taking the nuclear radius as `R = 1.2 A^(1/3)` femtometres for each partner and `e^2 = 1.43996454 MeV·fm`:

```
R    =  1.2 * 208^(1/3)  +  1.2 * 1^(1/3)   =  8.3100 fm

E    =  e^2 * Z1 * Z2 / R  =  1.43996454 * 1 * 82 / 8.3100  =  14.2091 MeV
 C

for an alpha particle: R = 9.0149 fm, E  = 26.1961 MeV
                                       C
```

Against that, the thermal energy available in the hottest furnace a pre-industrial workshop could build:

```
kT at 2000 K  =  8.617333262e-5 * 2000  =  0.172347 eV

E  / kT       =  14.2091e6 / 0.172347   =  8.24e7          eighty-two million
 C

temperature required for kT = 14 MeV    =  1.62e11 K
```

THEOREM 6.1 (the arithmetic, not a new result — the barrier is standard nuclear physics and the statement here is only the comparison). A furnace at 2000 K supplies thermal energy 82 million times smaller than the barrier to the least demanding nuclear approach to a lead nucleus. No arrangement of crucibles, reagents, timing, or repetition closes a gap of that size, because nothing in the chemical repertoire operates in the register where the barrier lives. ■

AND NOW THE PART THAT MAKES THIS AN AUDIT RATHER THAN A DISMISSAL. The goal was not incoherent. Transmutation of the elements is real, it happens, and it was demonstrated in the twentieth century. Rutherford (1919) reported the first artificial transmutation, converting nitrogen to oxygen by alpha bombardment. **That paper is not in this paper's bibliographic database and is named without a key.** Gold has been produced from neighbouring elements in accelerator experiments; the bombardment of bismuth-209 with relativistic heavy ions reported by Aleklett, Morrissey, Loveland, McGaughey and Seaborg in *Physical Review C* (1981) yielded gold isotopes among its fragmentation products in vanishing quantities and at enormous cost. **That paper is likewise absent from this paper's bibliographic database and is named without a key.**

So the tradition's objective was achievable and its method was not, and the distance between them is exactly the six orders of magnitude computed above. The error was not in wanting to change one element into another. It was in believing that the operation lived in the same register as distillation. §6.8 prices what accelerator production would actually cost, and the number is the reason nobody does it.

6.2 Charge two: the record is survivorship-filtered **The framework's survivorship result, restated exactly and without softening: under a selection process that publishes successes and discards failures, the likelihood ratio of testimonial evidence is exactly 1.000000.**

A likelihood ratio of 1 does not mean the evidence is weak. It means the evidence does not move the posterior at all — not by a little, not at all — because the selection guarantees that the observed outcome occurs whether or not the practice does anything.

The alchemical record is exactly such a corpus. It preserves accounts of operations that worked. It does not preserve the denominator: the number of practitioners who attempted the same operation, at the same period, and recorded nothing because there was nothing to record. §6.9 computes the table; the structure of the argument is that every hypothesis about the underlying success rate predicts the same observed corpus, once the corpus size is set by how many successes are worth writing down.

The modern statement of the same defect is the selective-publication literature: results are reported conditional on having come out, which is why most published findings in a field with weak power are expected to be false [Ioannidis, 2005], why the same flexibility produces significance at will [Simmons et al., 2011, Gelman and Loken, 2013], and why a large replication programme recovers far less than the original record implies [Open Science Collaboration, 2015]. The precedent for the objection is older than the corpus it demolishes. Cicero's *De Natura Deorum* relates the reply attributed to Diagoras, shown the votive tablets of those who prayed and survived shipwreck: the drowned left no tablets. Wald's work for the Statistical Research Group in 1943 — armour the parts of the returning aircraft with no bullet holes, because aircraft hit there did not return — is the same argument with a number attached. **Neither source is in this paper's bibliographic database; both are named without keys.**

The consequence for Section 5 is structural and is why that section states its columns' roles in every block. The (-1) precedent column is drawn from a survivorship-filtered record. That is precisely why it carries precedent and not evidence, and why no accumulation of it can be promoted. A thousand accounts and one account carry the same weight, and the weight is zero.

6.3 Charge three: secrecy as an epistemic failure mode The tradition concealed its results deliberately, by three distinct mechanisms: restricting transmission to initiates, writing in cover names (*Decknamen*) that mapped many substances onto one word, and omitting steps on purpose.

The epistemic consequence is not that secrets are morally bad. It is arithmetic — and it has an exact analogue in the design literature, where a mechanism whose security depends on the mechanism being secret is held to be badly designed [Kerckhoffs, 1883, Saltzer and Schroeder, 1975], and where publicity is the standard corrective for a practice that cannot otherwise be checked [Brandeis, 1914].

AN UNSHARED RESULT CANNOT BE CHECKED, AND AN UNCHECKED ERROR HAS NO MECHANISM OF REMOVAL. Model it directly. Let an error be caught by any single independent replication with probability `q`, and let `r` independent replications occur per generation. The error survives a generation with probability `(1-q)^r`. At `r = 0` — which is what secrecy enforces — survival is `1.000000` per generation, forever, and the expected time to detection is infinite. §6.11 computes the rest of the table.

The second mechanism is worse, because it corrupts the record even for those inside. A cover name is a many-to-one map from substances to words, which is the same object as gematria's hash and fails in the same way. Its cost is computable as a search: if a recipe has twelve steps and each carries a term with three plausible readings, the number of distinct interpretations is `3^12 = 531,441`, which is **19.019550 bits of specification destroyed** and, at one experimental trial per day, 1,455 years to exhaust by brute force. §6.12 gives the table.

The tradition's secrecy therefore did not merely slow transmission. It removed the only mechanism by which a body of results corrects itself, and it did so at the same time as it made each individual result ambiguous enough to require an infeasible search. Those two failures

compound: an error that cannot be caught, encoded in a document that cannot be read the same way twice.

6.4 Charge four: the absence of controls No control was run. Not a poor control, not an underpowered one — none, as a matter of the period's method, for most of the tradition's history.

This matters most because of a specific and entirely non-fraudulent explanation for the tradition's positive results, which §6.10 computes. Lead ores commonly carry trace gold and silver, and cupellation — the standard fire-assay of the period — is a *concentrating* operation. Process enough auriferous lead and a bead of gold appears at the end. Nothing was transmuted; something was separated.

```
lead containing 1 ppm gold,  10 kg processed   ->   10 mg of gold recovered
lead containing 2 ppm gold,   1 kg processed   ->    2 mg of gold recovered
```

A visible bead is on the order of ten milligrams, which requires one part per million from ten kilograms of lead — an entirely ordinary assay, and indistinguishable from a successful transmutation to anyone who did not assay the starting material. The control that settles it is: assay the starting material. It costs one extra crucible.

The general form of the charge, and it is the one that generalises to this framework's own subject: the tradition ran the treated arm and never ran the untreated one. Bacon's *Novum Organum* (1620) argued for the systematic collection of negative instances. Fisher's *The Design of Experiments* (1935) supplied randomisation and the modern control, and the Medical Research Council's streptomycin trial (1948) is generally taken as the first properly randomised clinical trial. **None of these three is in this paper's bibliographic database; all are named without keys.** The modern standard the argument does rest on — a declared design, fixed before the outcome is seen — is cited [Nosek et al., 2018]. Every one of those postdates the tradition's productive period, which is an explanation and not an excuse — the specific control needed here, weighing the input, required nothing the period lacked.

And a warning that belongs here rather than anywhere else, because a modern reader will reach for the wrong repair. **Adding statistical surrogates to an uncontrolled design does not fix it.** The framework's survey work measured this [Curzi, 2026i]: the shuffled-input control run as ordinarily configured is anti-conservative at 2.2 times nominal alpha (0.108 against 0.05); against fabricated timing the full four-baseline suite passed 52.7% of the time; against a shared slow driver it passed 98.35% of the time. **Design controls — blinding, channel separation, prospective sampling, a yoked control — are the answer, and more surrogates are not.**

6.5 Charge five: the canonical order is backwards The fifth charge is the one this framework is best placed to make, because it audits the tradition on the tradition's own published material using a rule the tradition could not have had.

The framework's operation audit [Curzi, 2026] takes a named, ordered, canonical fifteenth-century schedule of twelve operations and runs it as a schedule under this model, sorting each gate by which parameter it acts on. **The result: the only gate that acts on the numerator sits ninth of twelve.** Eight of the operations that precede it act on retention, containment, and the prevention of loss — which is to say on `delta`.

The framework's ordering rule condemns this [Curzi, 2026q]. `Q_inf = (eta F − c)/delta`. Lowering `delta` divides the numerator by a smaller positive number, which amplifies whatever sign the numerator already has. In an arrangement whose numerator is negative, every operation that lowers loss makes the destination worse, and the canonical sequence spends two-thirds of its length there before it touches the numerator once.

The carried case is the demonstration, and the numbers are the carried case's (CONSTRUCTED):

```
the canonical order -- lower delta first, on the numerator as found

    2a:  eta = 0.30, c = 0.62, u = 0.068734665264634,  eta*F - c = -0.451075
         delta 0.09  ->  Q_inf = -5.0119395464       D = +4.338420
         delta 0.05  ->  Q_inf = -9.0215             D = +4.338420   WORSE

the ordering rule -- fix the numerator first, then lower the loss

    2b:  eta = 0.44, c = 0.40, u = 0.692372881355932,  eta*F - c = +0.547598
         delta 0.09  ->  Q_inf = +6.0844267427       D = -3.5909886634
         delta 0.05  ->  Q_inf = +10.9520            D = -3.5909886634   BETTER
```

The same intervention on `delta` nearly doubles the fixed point in whichever direction the numerator already points, and the deficit does not move in either case — the two-reading rule again, and decay-blindness [Curzi, 2026] behind it.

Two qualifications, both required. First, the assignment of each gate to a parameter is this framework's reading of a fifteenth-century text and not the text's own claim; The operation audit owns the mapping and badges it as an interpretation [Curzi, 2026]. Second, the finding is not that the operations are worthless. Every one of them is a real operation and several are on Section 5's list of things the tradition got right. The finding is about the ORDER, and the order is a claim the text does make explicitly, in a numbered sequence, as the thing that must not be varied.

That is the sharpest available form of what this pair of sections is for. A tradition whose individual practices survive an audit can still have its canonical sequence inverted, and only a formalism that separates the numerator from the loss term can see the difference.

6.6 The audit template, stated as this section's governing rule

Every charge above has the same two-part shape, and the second part is not decoration.

THE STRONG CLAIM	THE ARITHMETIC THAT KILLS IT	WHAT SURVIVES
chemical transmutation of elements	82 million-fold energy gap	staged process control; apparatus; the goal itself, achieved by other means
the record establishes a success rate	LR = 1.000000	the record as a catalogue of operations attempted
secrecy protects a valuable result	no correction mechanism at r = 0; 3^12 readings per 12-step recipe	the recognition that some knowledge is dangerous -- which is a real problem with a different solution
the tradition's positives show efficacy	1 ppm x 10 kg = 10 mg, no transmutation required	the observations themselves, which were often accurate
the canonical sequence must not vary	the only numerator-acting gate is 9th of 12, and that order accelerates a failure	the operations themselves, and the idea that

```
order matters --
it does, and it is
the reverse of this
```

A tradition audited honestly is more usable than one defended, because after the audit it is known which parts may be built on. Section 5's four items are usable precisely because this section exists; without it they would be four claims selected from a corpus by an author who liked them.

6.7 Honesty clauses, restated in place **(1) NO SUPERNATURAL CLAIM.** Nothing in this section or its companion asserts any efficacy of symbol, any physical effect of a figure, or any result not derivable from the formalism.

(2) ESSENCE, NOT RESONANCE. The frequencies in §6.14 are essence frequencies: positions on a logarithmic surplus scale. Nothing is claimed to vibrate. This clause matters more in this section than in most, because a section about a tradition that made material claims must be unmistakable about not making one.

(3) NO CLINICAL CLAIM. The twelve states are not diagnoses; nothing here is a treatment.

(4) ALL CASES CONSTRUCTED. The lineage case of §6.14 and the process cases of §6.8–§6.13 are constructed and labelled. No invented case is presented as an observation.

(5) PRECEDENT IS NOT EVIDENCE, AND §6.2 IS THE PROOF.

(6) BOTH READINGS. Every arrangement read here reports D and δ together. `D` contains no `delta` and cannot see decay.

(7) ORDER-OF-MAGNITUDE FIGURES ARE LABELLED AS SUCH. §6.8's per-gram cost rests on stated assumptions about beam energy and yield, and is presented as a constructed estimate with its assumptions in the table, not as a measurement.

The engineering

6.8 What transmutation actually costs, computed CONSTRUCTED ESTIMATE. Assumptions stated in the table; this is an order-of-magnitude figure and not a measurement of any facility.

One gram of gold-197 contains `6.02214076e23 / 196.966570 = 3.057443e21` atoms. Suppose each successful transmutation event requires a beam particle of energy `E`, and suppose a fraction `epsilon` of beam particles produce the desired product. Then the beam energy required is `3.057443e21 * E / epsilon`. For reference, a one-gigawatt station produces `3.15576e16` joules in a year.

beam energy per event	yield `epsilon`	total energy (J)	in kWh	in GW-station-years
1 GeV	1e-6	4.899e17	1.361e11	**15.52**
1 GeV	1e-9	4.899e20	1.361e14	1.552e4
5 GeV	1e-6	2.449e18	6.804e11	77.61
5 GeV	1e-9	2.449e21	6.804e14	7.761e4

Under the most generous of these assumptions, one gram of gold costs the entire annual output of fifteen one-gigawatt power stations, and under the least generous it costs fifteen thousand of them. The physics permits the transformation; the economics forbids it by a margin that no refinement of technique will close. This is the honest end of the alchemical project: not that it was impossible, but that it was answered, and the answer was that the product costs more than the product is worth by a factor with sixteen digits in front of it.

6.9 The survivorship arithmetic, in full CONSTRUCTED. Suppose the alchemical corpus preserves 500 credible first-person accounts of a successful operation. Let `p` be the probability that a practitioner attempting it succeeds. The number of practitioners required to generate 500 successes is `500/p`:

true success rate `p`	interpretation	practitioners needed for 500 accounts
0.500	the operation works, half succeed	1,000
0.100	it works weakly	5,000
0.010	it works barely	50,000
0.001	it does nothing beyond the base rate	500,000
1e-6	it does nothing at all, and the accounts are misidentifications	500,000,000

Alchemical practice ran across more than a millennium and several linguistic traditions; a practising population in the tens or hundreds of thousands is entirely plausible, and populations of that size are consistent with every row.

```
LIKELIHOOD RATIO between any two rows, given only the numerator:   1.000000
```

Every row predicts the observed corpus equally well, so the corpus carries no information about `p`. And the invariance to corpus size is the part that surprises people:

accounts in the corpus	cumulative likelihood ratio
1	1.000000
100	1.000000
10,000	1.000000
1,000,000	1.000000

Starting from prior odds of 1:99, the posterior odds after a million success accounts are 1:99. The problem is not sample size and cannot be fixed by more of the same, because each further account adds a numerator and no denominator.

What a denominator would have done. Suppose instead that 1,000 attempts had been enrolled before their outcomes were known, and the successes counted:

```
1,000 enrolled,   500 successes:   log LR(p=0.5 vs p=0.001) = +2761.2307
1,000 enrolled,     1 success:     log LR(p=0.5 vs p=0.001) =  −685.2399
```

The same 500 successes, observed with a denominator, carry a likelihood ratio on the order of `1e304` — and observed without one, exactly 1. The entire difference is whether the failures were counted. Enrolment before the outcome is known is a cohort design, it required no instrument the period lacked, and it is the single methodological change that would have converted the tradition's whole record into evidence.

And the companion result, because it sharpens what a claim must do to be testable at all. A prediction specifying only a *direction* was measured at a likelihood ratio of 2.0000; one specifying a *magnitude* at 11.3577 — **a prediction with a number attached is about 5.7 times more evidentially powerful.** The alchemical corpus specifies neither: its claims are unquantified and its record is denominator-free, which are the two independent ways a body of reports can fail to be evidence.

6.10 Where the positive results came from, without anyone lying The survivorship argument explains why the record cannot establish a rate. It does not explain why practitioners saw gold. This does, and it requires no fraud and no error of observation.

CONSTRUCTED, but the metallurgy is standard. Lead ores frequently carry trace gold and silver, and cupellation — heating in a bone-ash vessel so that the lead oxidises and is absorbed, leaving the noble metals behind — concentrates them. It was the period's ordinary assay technique.

gold in the starting lead	mass processed	gold recovered
0.5 ppm	1 kg	0.50 mg
1.0 ppm	1 kg	1.00 mg
1.0 ppm	10 kg	**10.00 mg**
2.0 ppm	10 kg	20.00 mg
5.0 ppm	10 kg	50.00 mg

One part per million in ten kilograms of lead yields a ten-milligram bead — a visible, weighable, demonstrable piece of gold produced from lead by a documented procedure, with no transmutation whatever. The practitioner's observation was correct. The inference was not.

The control that separates the two hypotheses is: **assay the starting material.** Weigh the lead, assay a sample of it, and predict the yield. If the recovered gold equals the assayed content, the operation concentrated; if it exceeds it, something else happened. **The instrument required is a balance, and the tradition had balances.**

This is the sharpest single point in the section and it should be read carefully, because it is not a sneer. **The tradition possessed every physical instrument required to refute its own central claim and did not possess the idea of running the comparison.** The missing thing was not equipment. It was the control.

6.11 How long an error lives when nobody may check it **CONSTRUCTED.** Let an error in a recorded procedure be caught by any single independent replication with probability `q = 0.30`, and let `r` independent replications occur per generation. Survival per generation is `(1−q)^r`.

replications per generation `r`	survival per generation	expected generations to detection	survival after 12 generations
0 (secrecy)	**1.000000**	**infinite**	**1.000000**
1	0.700000	3.3333	0.013841
2	0.490000	1.9608	1.916e-4
3	0.343000	1.5221	2.652e-6
5	0.168070	1.2020	5.080e-10
10	0.028248	1.0291	2.581e-19

Taking a generation as twenty-five years, twelve generations is three hundred years. **A single independent replication per generation reduces an error's three-century survival probability from certainty to 1.38%; five reduce it to five parts in ten billion; zero leaves it at one.** The difference between a tradition that corrects itself and one that does not is the difference between `r = 1` and `r = 0`, and secrecy is precisely the choice of `r = 0`.

The corresponding modern figure is the replication rate a field actually achieves [Open Science Collaboration, 2015, Ioannidis, 2005]: r is not a free parameter, it is a property of how a field is organised.

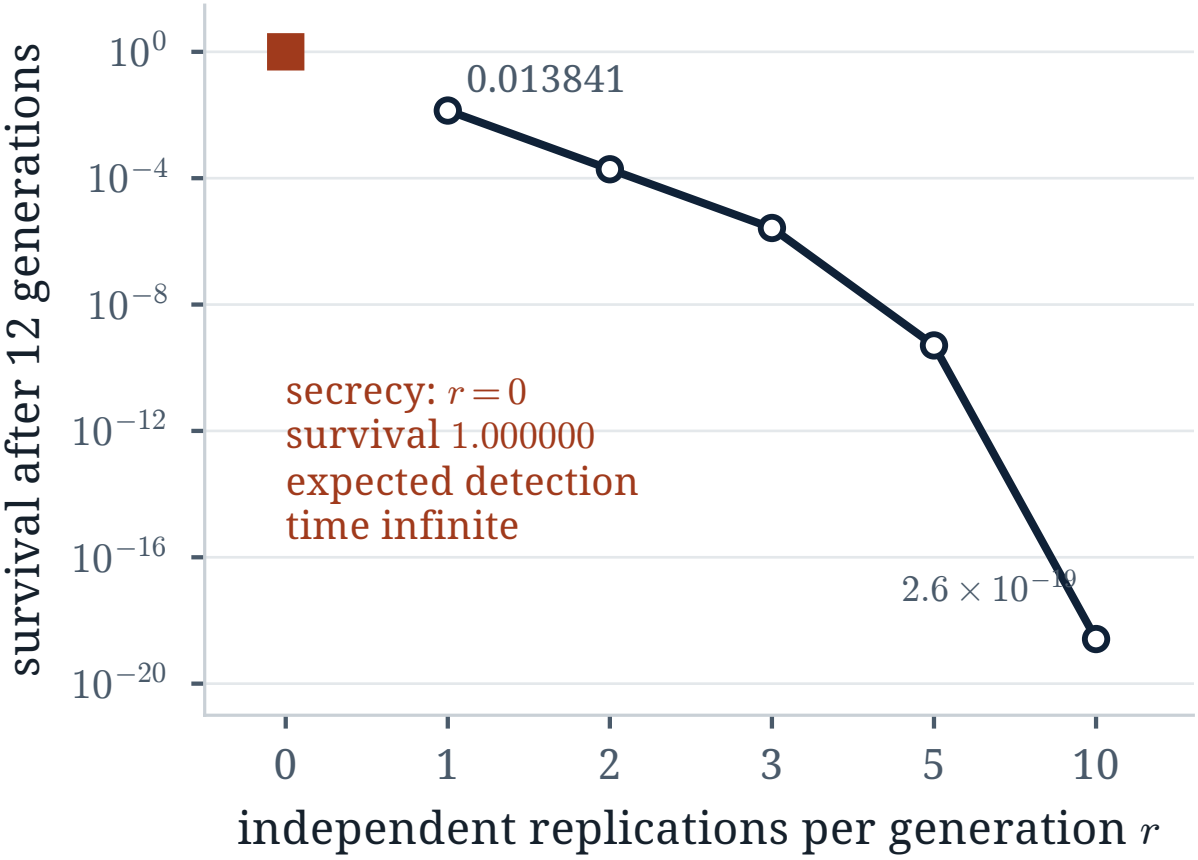


Figure 47.5: **At zero independent replications an error survives every generation with probability 1.000000 and its expected time to detection is infinite; one replication per generation cuts three-century survival to 1.38 per cent.** Modelled with a per-replication detection probability $q = 0.30$ and a generation of twenty-five years; the model is constructed. The $r = 0$ point is plotted apart from the curve because it is a different kind of case — not slow detection but no detection mechanism at all. Built by `figs/src/CCP-047_fig05.py`.

The row that matters is the first one, and its distinguishing feature is not that detection is slow. It is that the expected time to detection is infinite: there is no mechanism, so there is no rate. A body of results maintained in secret does not correct slowly; it does not correct.

6.12 What cover names cost, as a search A *Deckname* is a cover word standing for a substance or an operation, many-to-one by design. The historiography establishing that these were used as deliberate concealment rather than as poetry, and that modern decipherment of specific recipes is possible but laborious, is Newman and Principe's work on the Starkey and Boyle materials and Principe's *The Secrets of Alchemy* (2013).

The cost of the practice is a search cost and it is computable.

steps in the recipe	plausible readings per step	distinct interpretations	bits destroyed
6	2	64	6.00
6	3	729	9.51
12	2	4,096	12.00
12	**3**	**531,441**	**19.02**
12	4	16,777,216	6.00
20	3	3,486,784,401	31.70

At one experimental trial per week, exhausting `3^12` interpretations takes 10,220 years; at one trial per day, 1,455 years.

A twelve-step recipe with a three-way ambiguity at each step is not a difficult document. It is an undecipherable one by brute force, and every successful modern decipherment has proceeded by constraint rather than by search. That is the honest measure of what concealment cost the tradition: not a delay, but a change of category, from a procedure that can be followed to a puzzle that must be solved.

Note the structural identity with the value-map case, which is why this section reads them together. **A cover-name vocabulary and a numerical-value scheme are the same kind of object — a many-to-one map — and both fail for the same reason and survive for the same reason.** The cover name cannot carry a specification, because the map is not invertible. It can carry an index, and it did: it let a practitioner recognise which family of operations a passage was about, which is a mnemonic function and a real one.

6.13 The collision arithmetic, recomputed Because §6.6 makes that case the section's template, the arithmetic is reproduced rather than cited. The quantities below are Shannon information in bits [Shannon, 1948, Cover and Thomas, 2006, Hartley, 1928] — not thermodynamic entropy, and no thermodynamic claim is made anywhere in this paper. The probability that a given word in a corpus of M words shares its value with at least one other, against `R` reachable values, is `1 - (1 - 1/R)^(M-1)`:

corpus `M`	reachable values `R`	P(a given word has a value-twin)
500	400	0.713226
2,000	1,000	0.864665
8,000	**1,000**	**0.999666**
8,000	1,500	0.995177
30,000	1,500	1.000000

And the measured map, on 234,454 dictionary words: 305 distinct values, mean bucket 768.7, largest bucket 2,353 words, 99.995% of words colliding, 10.47 bits destroyed of the 17.84 that identify a word.

One arithmetic note that strengthens the result rather than weakening it. The *uniform* bound on destroyed information is `log2(234454) − log2(305) = 17.8389 − 8.2527 = 9.5863` bits. The measured figure of 10.47 is larger than the uniform bound because the buckets are wildly unequal, and unequal buckets destroy more information than uniform ones. The map is worse than its own worst-case uniform analysis.

The evidential reading is therefore dead by measurement, and the mnemonic reading is untouched by any of it. A retrieval index does not need to be injective. It needs to be reproducible, and it is.

6.14 Secrecy read in this framework's own register The charges above are historical. This one is the framework's, and it is the reason secrecy belongs in a volume about relationship frequency at all.

CONSTRUCTED. Not an observation of any lineage, historical or living. Consider two ways of holding a body of practice, read with the essence bridge and the octave deficit [Curzi, 2026].

A **closed lineage**: one master and one student, where the student's whole training consists of reproducing the master's framings exactly. That is `rho = 1`, hence `u = 0`, hence — because u depends on ρ^2 — the same reading a perfectly opposed pair would get. $F = \ln(1) = 0$ and $f = k = 0.483426$ Hz, the ground essence. **A pair at $u = 0$ has no relationship at all in the surplus sense: two parties in unison are one party paying twice.**

An **open school**: six members whose framings genuinely differ, publishing, so that interaction converts to durable record (`eta` up) and coordination is cheaper per link (`kappa` down).

arrangement	u	F	f (Hz)	eta	c*	c/eta	D (oct)	Q_inf (δ=0.09)
closed lineage, exact reproduction	0.000000	0.000000	**0.483426**	0.30	0.6750	2.250000	**+6.492128**	**-7.500000**
closed lineage, slight drift	0.050000	0.438255	1.161430	0.30	0.6750	2.250000	+5.227591	-6.039150
open school, six, u = 0.30	0.300000	1.458615	8.938542	0.44	0.9200	2.090909	+1.824415	-3.091215
open school, six, u = 0.50	0.500000	1.871802	20.424737	0.44	0.9200	2.090909	+0.632209	-1.071189
open school, six, u = 0.70	0.700000	2.163323	**36.590494**	0.44	0.9200	2.090909	**-0.208942**	**+0.354024**

Every row's `c/eta` is checked against the ceiling `ln 12 = 2.484907` and every row is inside it, so the failures below are real failures and not ceiling breaches.

Four readings.

The closed lineage sits at the ground essence, `f = k`, which is the lowest pitch the formalism admits. It is not a low reading; it is the reading for *no relationship*. A student who reproduces the master exactly contributes no non-parallelism, and the arrangement's surplus is zero by construction.

The swing from the closed lineage to the open school at `u = 0.70` is 7490.44 cents — six and a quarter octaves of pitch — and 6.701070 octaves of deficit, from `+6.492128` to `−0.208942`. That is the difference between an arrangement that cannot be sustained under any refinement and one that accumulates.

The organisational literature reaches the same conclusion by a different route: a system that homogenises itself destroys the variance that produced its advantage [March, 1991], bridging across structural holes has measured returns [Granovetter, 1973, Burt, 1992], and the returns to difference are real but not monotone — creative success rises and then falls with the same variable [Uzzi and Spiro, 2005], short-run effects of surface diversity are frequently negative [Putnam, 2007], the rebuttal literature qualifies rather than reverses that [Abascal and Baldassarri, 2015], and the formal diversity-beats-ability result carries a standing objection that must be cited with it [Hong and Page, 2004, Thompson, 2014, Page, 2007]. Openness is not one lever here; it is three at once, and that is why the swing is so large. It raises `u` because independent practitioners frame differently; it raises `eta` because a published result is stored order and an unpublished one is a memory; and it lowers per-link `c` because a written

record replaces re-explanation. Three of the framework's four levers [Curzi, 2026], moved by a single decision.

A standing clause binds the reading of this table: every `f` in it is a RELATIONSHIP frequency, and a relationship frequency says nothing whatever about the state frequencies of the people in it. The framework's operation audit [Curzi, 2026] established the point by construction — four pairs at lattice positions spanning thousands of cents return identical u and identical f_{rel} to twelve decimal places, because `u` reads the angle between state vectors and not which indices are occupied. The closed lineage reads at `f = k` as a *pair*; the two people in it may hold any states at all, at any pitches, and the pair reading would not move. Neither number may be inferred from the other and a complete report gives both.

And the two-reading rule binds this table as it binds every other. The closed lineage at `delta = 0.09` arrives at `Q_inf = −7.500000`; at `delta = 0.05` it arrives at `−13.500000`; and `D` stays at `+6.492128` in both. **A tradition that responds to decline by deepening loyalty and lengthening apprenticeship is lowering `delta` on a negative numerator, which the ordering rule forbids, and it makes the destination worse while moving the deficit by exactly nothing.** Fix the numerator first: raise `u`, raise `eta`, cut `c`. Publishing does all three.

6.15 What the audit leaves standing

Set against the four charges, here is the full list of what survives, and it is not short:

```
SURVIVES INTACT
  staged process control, with its per-stage check              C5.1, §5.8
  cadence, for the two correct reasons                          C5.2, §5.9–5.12
  correspondence tables as counted partitions (FORM only)       C5.3, §5.14
  the operator as part of the apparatus                         C5.4, §5.16
  the apparatus itself -- alembic, cupel, balance, furnace
  the observations, which were frequently accurate
  the mnemonic reading of every many-to-one map in the corpus   §6.12, §6.13
  the goal of transmutation, achieved by nuclear means          §6.1, §6.8

DOES NOT SURVIVE
  material transmutation by chemical means                      §6.1
  the record as evidence about a rate                           §6.2, §6.9
  secrecy as a defensible epistemic practice                    §6.3, §6.11
  the evidential reading of any correspondence or value scheme  §6.13
  the colour sequence as a law                                  REJECTION C5.5
  "as above, so below" as a transfer licence                    REJECTION C5.6
```

The ratio matters more than either column. Eight items stand and six fall, and the six fell to arithmetic that the eight had to survive as well. That is what makes the eight usable.

The specification

6.16 PROCEDURE AUDIT-1 — auditing a traditional claim's evidential base

```
STEP 1   CLASSIFY THE CLAIM: structural / causal / evidential (ALCH-R step 2).
         This procedure handles the evidential ones.

STEP 2   ASK FOR THE DENOMINATOR. How many attempts were made, counting the ones
         that produced nothing? If the answer is unavailable, stop: the likelihood
         ratio is 1.000000 and the corpus is inert with respect to any rate.
```

STEP 3 ASK WHETHER THE CLAIM SPECIFIES A MAGNITUDE. Direction-only claims carry LR 2.0000; magnitude claims 11.3577. A claim with no number is not yet testable (the direction-versus-magnitude result).

STEP 4 FIND THE INNOCENT EXPLANATION FIRST. Before any exotic mechanism is considered, compute what the mundane one predicts. §6.10's 1 ppm x 10 kg is the model: a full quantitative account of the observation requiring nothing new.

STEP 5 NAME THE CONTROL THAT WOULD SEPARATE THEM, AND PRICE IT. If the control is cheap and was not run, say so plainly. That is the finding.

STEP 6 CHECK THE REGISTER. Does the proposed mechanism operate at the energy, time or information scale the effect requires? §6.1's 8.24e7 ratio is the model. A register mismatch of orders of magnitude ends the inquiry.

STEP 7 CHECK FOR MANY-TO-ONE MAPS. Any correspondence, value scheme or cover vocabulary: compute the collision rate. If collision is the default outcome, the evidential reading is dead and the mnemonic one is not.

STEP 8 CHECK THE CORRECTION MECHANISM. How many independent replications per generation? At r = 0 the expected time to detection is infinite.

STEP 9 DO NOT REPLACE DESIGN CONTROLS WITH SURROGATES. Measured: 2.2x nominal alpha, 52.7% pass against fabricated timing, 98.35% against a shared slow driver.

STEP 10 PUBLISH THE REJECTION, with the failing step named.

STEP 11 STATE WHAT SURVIVES, in the same document. An audit that only subtracts is not an audit; it is a dismissal with arithmetic attached.

STEP 12 REPORT D AND delta TOGETHER for any arrangement read.

6.17 The demotion rule

A claim moves DOWN this ladder when a step of AUDIT-1 fails, and never up without new data of the kind the failed step demanded.

```
  EVIDENCE     measured, with a denominator, falsifiable, replicated
     |
  PRECEDENT    recorded independently; establishes findability, not truth
     |
  MNEMONIC     a reproducible index with no evidential content
     |
  REJECTED     inconsistent with computed arithmetic; published as rejected
```

THE ONE-WAY RULE: a claim demoted for lacking a denominator is not restored by more accounts of the same kind. It is restored only by a cohort design.

6.18 Gates

G1 DENOMINATOR ASKED FOR, and its absence reported where absent.
G2 INNOCENT EXPLANATION COMPUTED before any exotic one is entertained.
G3 CONTROL NAMED AND PRICED.

```
G4  REGISTER CHECKED. Energy, time or information scale compared explicitly.
G5  COLLISION RATE COMPUTED for every many-to-one map.
G6  CORRECTION MECHANISM COUNTED. r replications per generation, stated.
G7  NO SURROGATE SUBSTITUTION for design controls.
G8  SURVIVORS LISTED alongside rejections, in the same document.
G9  ORDER-OF-MAGNITUDE ESTIMATES LABELLED, with assumptions in the table.
G10 BOTH READINGS. D and delta together.
```

The failure modes

#	Name	What it looks like	What it actually is	Detection
FM 6.1	The register error	“Enough heat and pressure and it will convert”	A chemical operation aimed at a nuclear barrier, short by 8.24e7	G4; §6.1
FM 6.2	Testimonial accumulation	“Thousands of accounts cannot all be wrong”	LR = 1.000000 at every corpus size	G1; §6.9
FM 6.3	The missing denominator, modern form	Case studies of successful teams, rooms, practices	Same defect; this framework’s own worked cases are drawn from construction, and say so	G1; every case labelled CONSTRUCTED
FM 6.4	The concentrating operation	A bead of gold at the end of the crucible	1 ppm × 10 kg = 10 mg, separated not created	G2; §6.10
FM 6.5	Secrecy as protection	“The knowledge is too dangerous to publish”	`r = 0`; expected time to error detection is infinite	G6; §6.11
FM 6.6	The unreadable recipe	A twelve-step procedure in cover names	`3^12` = 531,441 readings; 1,455 years at one trial a day	G5; §6.12
FM 6.7	The value coincidence	“These two terms share a value, so they are linked”	P(a twin exists) = 0.999666 in an 8,000-word corpus	G5; §6.13
FM 6.8	Surrogates for design	Adding shuffle controls to an uncontrolled study	2.2× nominal alpha; 52.7% and 98.35% pass rates	G7; the survey work’s control audit
FM 6.9	Direction-only claims	“It will improve things”	LR 2.0000 against 11.3577 with a magnitude	G1; §6.9

#	Name	What it looks like	What it actually is	Detection
FM 6.10	The unison lineage	A student who has perfectly absorbed the master	`u = 0`, `f = k`, `D = +6.492128`	§6.14
FM 6.11	Loyalty as the repair	Deepening commitment in a declining school	Lowers `delta` on a negative numerator	G10; §6.14
FM 6.12	Dismissal mistaken for audit	Publishing only the demolition	The survivors list is half the method	G8; §6.15
FM 6.13	Promotion by antiquity	"It was believed for two thousand years"	Duration of belief is a function of `r`, not of truth	§6.11
FM 6.14	Retro-fitting the mechanism	Explaining a positive result after seeing it	Grounding must be established from the channel, before the outcome	the framework's grounding rule

Four of these need more than a row.

FM 6.3 is the one this framework must apply to itself, and does. The survivorship defect is not peculiar to alchemy; it afflicts any literature assembled from successes, including case studies of well-functioning teams and institutions. This framework's response is structural rather than rhetorical: every worked case throughout the framework is CONSTRUCTED and labelled as constructed, so that no reader can mistake a demonstration for an observation, and the framework's carried commons case is a failure that no single lever repairs. An account that showed only fixable examples would be lying by selection.

FM 6.5 deserves separating from the moral question it usually attracts. There are real reasons to withhold information, and this framework does not claim otherwise. What §6.11 establishes is narrower and harder: withholding removes the correction mechanism, and the cost of that is not a delay but a change of kind. If a body of knowledge must be kept restricted, then the correction mechanism has to be rebuilt inside the restriction — independent replication among those who do hold it — or the errors are permanent. That is a design requirement, not a prohibition.

FM 6.10 and FM 6.11 are the pair that make this section relevant to the rest of the framework. A transmission lineage optimised for fidelity is optimised for `u = 0`. It will feel like the deepest possible relationship — unison is easily mistaken for understanding — and it will read at the ground essence. And the intuitive repair, more commitment, acts on δ, which cannot move the deficit by any amount at all [Curzi, 2026]. **The lineage is not failing because its members are insufficiently devoted. It is failing because they agree.**

FM 6.13 is the failure mode that Section 5 exists to guard against and this section exists to prove. Duration of belief looks like evidence and is not. §6.11 gives the mechanism: at `r = 0` an error survives every generation with probability 1, so a belief's age measures the absence of checking rather than the presence of truth. A two-thousand-year-old error and a two-thousand- year-old truth have identical survival records under secrecy, which is precisely why the survival record cannot distinguish them.

Worked examples

6.19 The transmutation claim, run through AUDIT-1

The claim: that a described chemical procedure converts lead into gold.

```
    STEP 1  CLASSIFY.  Causal and evidential, both. Both are audited.

    STEP 2  DENOMINATOR.  Unavailable. The corpus preserves successes.
            -> LR = 1.000000. The record cannot establish a rate. DEMOTED to PRECEDENT.

    STEP 3  MAGNITUDE?  The claim as recorded specifies no yield per unit of input.
            -> LR 2.0000 at best, and not testable as stated.

    STEP 4  INNOCENT EXPLANATION.  Trace gold in the starting lead, concentrated by
            cupellation.  1 ppm x 10 kg -> 10.00 mg, a visible bead.  Full account of
            the observation with nothing new required.

    STEP 5  THE CONTROL.  Assay the starting material and predict the yield.
            Cost: one crucible and one weighing. AVAILABLE THROUGHOUT. NOT RUN.

    STEP 6  REGISTER.  Chemistry operates on electrons at ~4 eV. The transformation
            requires removing three protons across a Coulomb barrier of 14.2091 MeV.
            kT at 2000 K = 0.172347 eV.  Ratio = 8.24e7.
            -> REGISTER MISMATCH OF EIGHT ORDERS OF MAGNITUDE. Inquiry ends here.

    STEP 7  MANY-TO-ONE MAPS.  The recipe is in Deckname. 12 steps x 3 readings =
            531,441 interpretations; 19.02 bits of specification destroyed.
            -> the document cannot carry a specification.

    STEP 8  CORRECTION MECHANISM.  r = 0 by the tradition's own rule.
            -> expected time to detection: infinite.

    STEP 9  SURROGATES.  None available and none would help.

    STEP 10 REJECTION PUBLISHED.  Failing steps: 2, 4, 5, 6, 7, 8.

    STEP 11 WHAT SURVIVES.  The apparatus; the staging; the observations; the goal,
            which is achievable at 15.52 GW-station-years per gram under the most
            favourable constructed assumptions (§6.8). The claim as stated does not.

CONCORDANCE C6.1 — MATERIAL TRANSMUTATION BY CHEMICAL MEANS

  CLAIM        That a chemical operation converts base metal into gold.

  (+1) EVIDENCE  -- this column carries measurement
               Nuclear physics: elemental identity is proton number; the Coulomb
               barrier for a proton approaching lead-208 is 14.2091 MeV against a
               furnace thermal energy of 0.172347 eV at 2000 K. Transmutation IS
               real by nuclear means -- Rutherford (1919) for the first artificial
               case, and gold among the fragmentation products of relativistic
               heavy-ion bombardment of bismuth-209 reported by Aleklett, Morrissey,
               Loveland, McGaughey and Seaborg in Physical Review C (1981).
               WHAT IT MEASURED: that the transformation is possible and that the
               energy register required is eight orders of magnitude above chemistry.

  (-1) PRECEDENT  -- this column carries independent arrival, not proof
```

```
            The corpus's own accounts of successful projection, and their
            contemporaneous critics -- Boyle's "The Sceptical Chymist" (1661)
            attacked the element schemes from inside the practice, and Lavoisier's
            "Traite elementaire de chimie" (1789) replaced them.
            WHAT IT RECORDS: that the goal was pursued seriously for centuries
            and that the internal critique arrived before the external one.

(0p) INTEGRATION
            The objective was correct and the register was wrong. This is the
            cleanest available instance of a tradition being right about WHAT and
            catastrophically wrong about HOW.

CRITERIA    C1 FAILED -- no shared mechanism; the resemblance is of goal only.
            C2 FAILED -- the claim as recorded specifies no magnitude.
            C3 FAILED -- not usable; the procedure does not do what it says.

STATUS      REJECTED. Third published rejection in this pair of sections, after C5.5 and
            C5.6. Recorded in Appendix D with its failing criteria.
```

6.20 The closed lineage, in report format CONSTRUCTED. Not an observation of any lineage.

```
  ARRANGEMENT READING                                              [CONSTRUCTED]

    ARRANGEMENT    a transmission lineage: one holder, one successor, training
                   defined as exact reproduction of the holder's framings
    PERIOD         one pass = one year

    DESTINATION READING          -- contains no delta
      u = 0.000000     one-hot trap checked: N/A (the vectors are parallel, not
                       disjoint; this is the parallel case, not the degeneracy case)
      F = 0.000000 nats      f = 0.483426 Hz = k, the ground essence
      folded: k transposed by j = 9 -> 247.514 Hz, pitch class 1
      eta = 0.30   c* = 0.6750   c/eta = 2.250000   ceiling ln N = 2.484907  INSIDE
      f_min = 43.516598 Hz
      D = +6.492128 octaves                          DECAYS

    ARRIVAL READING              -- contains no D
      delta = 0.09 per year   half-life 7.3496 years   within 5% at year 32
      Q_inf = -7.500000
      at delta = 0.05:  Q_inf = -13.500000,  D unchanged at +6.492128

    THE FOUR LEVERS, priced against this arrangement
      L1 raise u      the successor frames differently        -> D +6.492128 to
                                                                 -0.208942 at u=0.70
                                                                 with the eta and c
                                                                 changes below
      L2 raise eta    publish, so interaction becomes record  -> part of the same move
      L3 cut c        a written record replaces re-explanation-> part of the same move
      L4 raise N      admit genuinely distinct categories     -> raises the ceiling
      NOT A LEVER: delta -- the deficit contains none. Loyalty, apprenticeship length and
      devotion all act on delta and move D by exactly nothing.

    HONESTY
      essence frequencies, not physical resonances ................. [x]
```

```
no clinical, diagnostic or therapeutic claim ................. [x]
case labelled CONSTRUCTED ................................... [x]
ability to decline preserved ................................ [x]
D and delta both reported ................................... [x]
```

The reading says the lineage cannot be saved by anything its own doctrine recommends, and can be saved by the one thing its own doctrine forbids. That is the whole of Section 6 in a single arrangement: the practice was not wrong about staging, or cadence, or the operator, or the value of a counted category set. It was wrong about secrecy, and secrecy is the parameter that sets `u`.

6.21 The rejection register for this pair of sections

#	Claim	Failing criterion / step	Status
C5.5	The colour sequence is a necessary order	C2 — identification is circular	**REJECTED**
C5.6	"As above, so below" as a transfer licence	C2 — three computed disanalogies	**REJECTED as stated**
C6.1	Material transmutation by chemical means	C1, C2, C3; AUDIT-1 steps 2, 4, 5, 6, 7, 8	**REJECTED**
C6.2	The corpus establishes a success rate	AUDIT-1 step 2 — LR = 1.000000	**REJECTED**
C6.3	Secrecy protects a valuable result	AUDIT-1 step 8 — `r = 0`, infinite detection time	**REJECTED as an epistemic practice**
C6.4	Shared values or correspondences are evidence	AUDIT-1 step 7 — 99.995% collision	**REJECTED; mnemonic reading retained**
C6.5	The canonical gate order must not be varied	the ordering rule — the only numerator-acting gate is ninth of twelve	**REJECTED as an ordering**

Seven rejections across the two sections, every one with its failing criterion named and its arithmetic in print. A method that never rejects is not a method, and a method that rejects without stating what survives is not an audit either.

6.22 What this section establishes, and what follows from it

```
the energy-register argument: 2e6 on bonds, 8.24e7 on the furnace-to-barrier ratio
the honest completion: transmutation is real by nuclear means, at 15.52 GW-station-
  years per gram under the most favourable constructed assumptions
the survivorship result in full: LR = 1.000000, invariant to corpus size;
  with a denominator, ~1e304
direction versus magnitude: LR 2.0000 against 11.3577
the innocent explanation, computed: 1 ppm x 10 kg = 10.00 mg, no transmutation
the control that was never run, priced at one crucible
error persistence under secrecy: r = 0 gives infinite expected detection time
Deckname search cost: 3^12 = 531,441 readings, 19.02 bits, 1,455 years at a trial a day
the value-map arithmetic recomputed, with the unequal-bucket correction
```

```
  (10.47 measured against a 9.5863 uniform bound -- worse than its own
  worst case)
the canonical order inverted: the only numerator-acting gate 9th of 12
the closed lineage read in the framework's own register: u = 0, f = k, D = +6.492128
a relationship frequency says nothing about a state frequency
PROCEDURE AUDIT-1, the demotion ladder, ten gates, fourteen failure modes
seven rejections published with their failing criteria
```

The audit ends with the tradition audited rather than either adopted or dismissed. What it established is that essence alchemy already has an exact formula — the compound rule for combining components [Curzi, 2026] — that transmutation is re-indexing and reaches nothing not already latent, that refinement arrives at a destination fixed before the first pass, and that four of the tradition's practices survive an audit which its largest claim does not.

The companion paper turns to the cellular layer, and it inherits this section's standard without amendment. The no-subliminals constraint is absolute. The open problem — that there is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal — is carried forward unsoftened into the companion paper on the missing transducer [Curzi, 2026], where it is met rather than softened. **The instrument that rejected material transmutation on an energy-register argument is the same instrument that material must survive, and it is pointed the same way.**

7 Results

The eight results below are the units other papers in this corpus cite. A citation to this paper should name one of them. Each carries its status: **DERIVED** where it follows from the stated definitions by algebra, **COMPUTED** where it is the output of exact arithmetic on stated parameters, and **POSTULATED** where it is asserted by this framework.

R1 — the two negative theorems. Refinement cannot move the fixed point, because $Q_\infty = (\eta F - c)/\delta$ contains no iteration count t (§4.4); and the ladder is silent in the frequency register, because f_{rel} is a function of u and N alone, so the pitch, the folded pitch class and the octave deficit are all unmoved by repetition (§4.5). *Status:* DERIVED by inspection. **These are the results this paper most wants read, and both are negative.**

R2 — the contraction as an equality. $|Q_t - Q_\infty| = (1-\delta)^t\,|Q_0 - Q_\infty|$, an equality at every distance and every pass, not a bound. At $\delta = 0.09$: $\tau = 10.6033$ passes, half-life 7.3496, five per cent of the gap remaining at pass 31.76, one per cent at 48.83, one part in a thousand at 73.24. *Status:* DERIVED; the constants COMPUTED. The companion paper states the same contraction with a maximum residual of 2.0624×10^{-15} over a hundred passes [Curzi, 2026].

R3 — what seven repetitions buy. Seven passes close $1-(1-\delta)^7$ of the gap: **48.32 per cent at** $\delta = 0.09$ **and 6.79 per cent at** $\delta = 0.01$**.** Inverting, seven passes close 95 per cent only at $\delta = 0.348164$, 99 per cent at 0.482053 and 99.9 per cent at 0.627241. *Status:* COMPUTED. **The tradition's canonical count is therefore an implicit and checkable claim that its operation loses roughly a third of its stock per cycle.**

R4 — staged process control, priced. A seven-stage process at per-stage defect probability $q = 0.10$, moving from terminal inspection alone to per-stage checking at detection $d = 0.90$, **raises final yield from 47.83 per cent to 93.21 per cent at a cost of 0.63 reworks per unit.** At $d = 0.50$ the yield is 69.83 per cent at 0.350 reworks; at $d = 0.99$ it is 99.30 per cent at 0.693. *Status:* COMPUTED on a stated defect model; the empirical warrant is the statistical-process-control literature, two of whose founding works are absent from this paper's bibliography and are named without keys (§5.2).

R5 — the cadence result. The integrated effect of one impulse and of the same total delivered periodically is **identical to every digit: 433.3333 in all three schedules**, maximum pairwise difference 3.41×10^{-13}. Cadence therefore does not deliver more. What it delivers is a steady state instead of a decay, and freedom from the surplus ceiling: at $N = 12$ an annual occasion carrying a nominal 26 nats delivers **2.4849 after the** $\ln 12$ **ceiling clips it, 9.56 per cent**

of the weekly practice, with 90.44 per cent lost to clipping, and $0.91^{52} = 0.007403$ of one impulse surviving the year. Division alone, with no ceiling involved, raises delivered total by a factor of 5.7074 across the range tested. *Status:* COMPUTED; the concavity theorem DERIVED elsewhere [Curzi, 2026v].

R6 — the secrecy arithmetic. With an error caught by any single independent replication with probability $q = 0.30$, survival per generation is $(1-q)^r$. **At $r = 0$ survival is 1.000000 per generation, indefinitely, and the expected time to detection is infinite**; at $r = 1$ survival after twelve generations — three hundred years — is 0.013841; at $r = 5$ it is 5.080×10^{-10}. *Status:* COMPUTED. **The distinguishing feature of the first row is not that detection is slow but that there is no mechanism, so there is no rate.**

R7 — the survivorship arithmetic. Five hundred credible first-person accounts require 1,000 practitioners at $p = 0.5$ and 500,000,000 at $p = 10^{-6}$. Every hypothesis about the underlying rate predicts the same observed corpus, so **the likelihood ratio is exactly 1.000000 and is invariant to corpus size** — one account and a million accounts carry the same weight, and the weight is zero. With a denominator, the same 500 successes out of 1,000 enrolled carry a log likelihood ratio of $+2761.2307$ against $p = 0.001$. A direction-only prediction carries a likelihood ratio of 2.0000 against 11.3577 for one specifying a magnitude. *Status:* COMPUTED.

R8 — transmutation priced. One gram of gold-197 is 3.057443×10^{21} atoms; at 1 GeV per event and yield 10^{-6} the beam energy required is 4.899×10^{17} J, **15.52 gigawatt-station-years**, rising to 1.552×10^{4} at yield 10^{-9}. The register argument that closes the chemical route is the ratio of the Coulomb barrier for a proton approaching lead-208, 14.2091 MeV, to the thermal energy of a 2000 K furnace, 0.172347 eV: 8.24×10^{7}**.** *Status:* COMPUTED as a constructed order-of-magnitude estimate with its assumptions in the table; the underlying nuclear physics is standard but **the two primary sources for artificial transmutation are absent from this paper's bibliography and are named without keys** (§6.1).

CR — counter-result. Repetition buys arrival time and nothing else: it cannot move the destination and it cannot move the pitch. The tradition's canonical seven passes close under half the gap at this framework's own working decay rate, and its canonical operation order is inverted with respect to the framework's ordering rule — the only numerator-acting operation sits ninth of twelve [Curzi, 2026]. Separately, **an apparent improvement can be an exact transfer**: cost export leaves the sum of drives invariant at $+0.643688$ to every digit while the exporter's own deficit improves by nearly an octave and a third, and the disadvantaged party's stock does not cross zero until pass 10.850479. The framework's own survey work reaches a structurally parallel conclusion about instruments that cannot see what they are not measuring [Curzi, 2026i], where a shuffled-input control passed 98.35 per cent of the time against a third variable driving both interaction and order while detecting only 48.17 per cent in the clean case.

8 Discussion

8.1 What follows

The separation of rate from destination is the paper's usable finding. A report that gives one number is incomplete, and which number is missing changes the error. A reader with only the deficit concludes that a decay-only intervention did nothing; a reader with only the half-life concludes that the practice became more patient; a reader with the fixed point sees the damage. On the carried constructed case, moving δ from 0.09 to 0.05 takes Q_∞ from -5.0119 to -9.0215, makes the ladder 84 per cent slower, and moves D by exactly nothing. **Three readings, three stories, one intervention.**

The seven-passes result generalises past this tradition. Any prescription of a fixed number of repetitions is an implicit claim about a loss rate, and naming the rate makes the prescription testable. That is the useful form of the finding and it is why H3 is stated as a hypothesis rather than as a conclusion: the framework can compute what δ the count implies, and only measurement can say whether the process has it.

The counterfeit-arrival result is the sharpest practical one, and its danger is a lag. The export check catches the transfer on the first pass; the disadvantaged party's *stock* does not reveal it for eleven. A protocol that monitors stocks rather than drives therefore finds the problem after it has become a fact. The structure is an externality in the ordinary sense [Hardin, 1968, Ostrom, 1990], and the remaining option is exit [Hirschman, 1970], which is why the framework's consent law makes the ability to decline structural rather than advisory [Curzi, 2026].

The institutional analogue is instructive and is computed elsewhere in this corpus. Moving η from 0.35 to 0.55 and δ from 0.25 to 0.12 takes a fixed point from 0.2653 to 3.9281, while a turnover event still costs 70 per cent of stock against 15 [Curzi, 2026h]. The same two readings, in an operational setting, disagreeing in the same way.

8.2 What does not follow

A slow ladder is not a patient one and a fast ladder is not a good one. δ is a loss term and the destination is divided by it. A practice that visibly changes a system quickly is a practice with a large δ, and the visible rapidity is the signature of a leaky system reaching its level — which may be any level at all.

The audit does not license the tradition. Four practices survive and five claims do not, and the four are usable *because* the five were computed to destruction rather than omitted. A list of things a tradition got right, published alone, is an advertisement.

Nothing here is clinical, diagnostic or therapeutic. The twelve component classes are not diagnoses, a refinement ladder is not a treatment, and no reading in this paper describes or prescribes for any person's condition.

Nothing here resonates. Every frequency is a position on a logarithmic surplus scale. The glide of §4.13 is a bookkeeping transform of a stored quantity and is badged as a constructed reading wherever it appears.

8.3 What the surrounding programme allows this paper to claim

The corpus's title places this work in a lineage that begins with *physique sociale* [Quetelet, 1835] and continues through the modern computational programme [Lazer et al., 2009, Pentland, 2014]. **The lineage's characteristic error is Quetelet's** — treating the mean of a distribution as a type rather than as a statistic — and this paper inherits the risk directly, because it writes a single u, η and c for an arrangement.

Three constraints from that programme bear on this paper's central object. First, human activity is bursty, with heavy-tailed waiting times and **no characteristic time scale** [Barabási, 2005]; a per-pass decay rate is a claim about a characteristic scale and can only answer that finding by insisting that its pass is a declared cycle rather than an observed rhythm, which is why declaring the period is a gate. Second, an explanation should name a mechanism rather than exhibit a fitted curve [Hedström and Ylikoski, 2010]; this paper fits nothing, which sidesteps the objection at the cost of having no empirical content to fit. Third, and most sharply: **what observation would distinguish this framework's laws from a neighbouring model that fits equally well?** [Sobkowicz, 2020, Durlauf, 2005] The honest answer is H3 — a tradition's repetition count implies a loss rate, and the loss rate is measurable — and nothing else in this paper answers it.

What the paper is entitled to claim is therefore narrow. Prediction is one of many legitimate reasons to build a model, and explanation, bounding and the disciplined exposure of one's own assumptions are others [Epstein, 2008]. The position this work is best placed to defend is the analytical one: mechanisms stated explicitly, parameters declared, and the arrangement within which the law is claimed named in advance [Keuschnigg et al., 2018].

8.4 What the wider corpus says this paper may not claim

Three papers in this corpus bound what an audit of this kind establishes, and they are cited here because they cut against the reading a sympathetic reader would give.

The assembled-architecture paper records that the only reading under which its architecture clears its own build threshold is the naive one assuming zero correlation between coordinating jurisdictions — $h = +0.083099$ — and that **every defensible bloc assignment returns a negative surplus**, which it prints [Curzi, 2026]. The failure-record paper is entirely a counter-result: disclosure did not confer immunity in the strongest available test case, its own backing basket is 4.61 times less diversified than its component count claims, the terminal phase of its own programme is a 0.090 option rather than a plan, and a levy presented as a costless accounting correction in fact reallocates burden [Curzi, 2026]. And the correlated- jurisdictions paper shows what a count that is not a count of independent things is worth: four entities at shares 40.91, 27.27, 18.18 and 13.64 per cent give a naive 3.4085 effective jurisdictions and a bloc-collapsed 1.0000 — **a 3.41-fold overstatement** [Curzi, 2026].

All three say the same thing to this paper. Four surviving practices are four practices, not a vindicated tradition; and a well-behaved arithmetic on constructed parameters bounds nothing about the world.

9 Verification

This section decomposes the paper's results into atomic steps, each checkable on its own. The standard is Occam's: **a postulate that can be dropped without the result changing has been dropped.** Every step carries one of four statuses — **ARITHMETIC** (reproducible with a calculator from the inputs named), **DEFINITION** (true by stipulation), **CITED** (an external source supports it, named), or **POSTULATED** (asserted by this framework, no external support).

9.1 The postulates

	postulate	status
P1	For $0 < \delta \leq 1$, $Q_{t+1} = (1-\delta)Q_t + b$ with b constant.	CITED as a form [Oppenheim et al., 1996, Elaydi, 2005]; the *interpretation* POSTULATED
P2	A contraction of modulus $q < 1$ on a complete space has exactly one fixed point, reached from any start.	CITED [Banach, 1922]
P3	$F = \ln(1 + (N-1)u)$, hence $F \leq \ln N$.	DEFINITION given P6
P4	$f_{\text{rel}} = k(1 + (N-1)u)^2$, the relationship form of the essence bridge.	POSTULATED (given P7, P8)
P5	The Pythagorean whole tone is the frequency ratio $9{:}8$.	CITED [Barbour, 1951]
P6	$N_{\text{eff}} = 1 + (N-1)u$, the affine count axiom.	POSTULATED — design axiom; **not a Hill number** [Hill, 1973]
P7	φ is the positive root of $x^2 = x + 1$, so $\varphi^2 = \varphi + 1$.	CITED [Euclid, 1956, Vajda, 1989]

	postulate	status
P8	A dimensionless ratio of the form $(9/8)^2/\varphi^2$ **is a frequency in hertz.**	POSTULATED — no source
P9	A stage introduces a defect with probability q, independently across stages; a check detects it with probability d.	DEFINITION (a stated model)
P10	Under success-only preservation, an observed corpus of successes is generated by an unknown number of attempts.	DEFINITION of the selection process
P11	Elemental identity is proton number, and the Coulomb barrier for a proton approaching lead-208 is $e^2 Z_1 Z_2 / R$.	CITED as standard nuclear physics; **the primary sources are absent from this bibliography**

Dropped as unnecessary. No postulate about the *ordering* of the twelve component classes is needed for R1–R8. No postulate about consonance is needed anywhere. P4, P5, P7 and P8 are needed only for the frequency-register statements (R1's second half, and the glide); **R2, R3, R4, R5, R6, R7 and R8 survive their deletion entirely**, and that separation is the point of §9.3.

9.2 The compounding chain

step	consumes	produces	status
L1	P1	$Q_\infty = b/\delta$ by setting $Q_{t+1} = Q_t$	ARITHMETIC
L2	P1, L1	subtracting, $Q_{t+1} - Q_\infty = (1-\delta)(Q_t - Q_\infty)$; iterate	ARITHMETIC
L3	L2	$\lvert Q_t - Q_\infty \rvert = (1-\delta)^t \lvert Q_0 - Q_\infty \rvert$ — **R2**	ARITHMETIC
L4	L1 by inspection	Q_∞ contains no t — **R1**, first half	ARITHMETIC
L5	L3	$\tau = -1/\ln(1-\delta)$; at $\delta = 0.09$, $\tau = 10.6033$, half-life $\ln 0.5 / \ln 0.91 = 7.3496$	ARITHMETIC
L6	L3	$(1-0.09)^7 = 0.516761$, so seven passes close 48.32% — **R3**	ARITHMETIC
L7	L6 inverted	$\delta = 1 - 0.05^{1/7} = 0.348164$ for 95% in seven — **R3**	ARITHMETIC

step	consumes	produces	status
L8	P5, P7	$(9/8)^2 = 1.265625$; $\varphi^2 = 2.618033988749895$	ARITHMETIC
L9	L8	$k = 0.483425732988414$	ARITHMETIC (exact)
L10	L9, **P8**	k carries the unit Hz	**POSTULATED**
L11	L10, **P4**, P3	$f_{\text{rel}} = k(1 + (N-1)u)^2$, a function of u, N only	**POSTULATED**
L12	L11 by inspection	f_{rel} contains no t — **R1**, second half	ARITHMETIC given L11
L13	L3, Def. 5	$\lvert D_t - D\rvert = (1-\delta)^t \lvert D_0 - D\rvert$; the glide spans $\lvert D\rvert$ octaves	ARITHMETIC (value); **POSTULATED** (meaning)
L14	P1, two parties, transfer x	$\sum_i b_i$ invariant; the plain sum of fixed points moves by $1/\delta_P - 1/\delta_R$ — **CR**	ARITHMETIC
L15	L14 at $\delta_P = \delta_R = 0.09$	sum of drives $+0.643688$ across the export range; Q_R crosses zero at $t^* = 10.850479$	ARITHMETIC
L16	P9	escape $q(1-d)$; yield $(1-q(1-d))^7$: $0.9^7 = 0.4783$, $0.99^7 = 0.9321$ — **R4**	ARITHMETIC
L17	P1, linearity	$\sum_t Q_t = B/\delta = 26/0.06 = 433.3333$ for every schedule — **R5**	ARITHMETIC
L18	P3, L17	$\min(26, \ln 12) = 2.484907$; $2.484907/26 = 9.56\%$ — **R5**	ARITHMETIC
L19	L5	$0.91^{52} = 0.007403$ of one impulse survives a year — **R5**	ARITHMETIC
L20	P10	every rate predicts the corpus; LR $= 1.000000$ at every size — **R7**	ARITHMETIC
L21	binomial with a denominator	at 500 successes in 1,000 enrolled, $\log \text{LR} = +2761.2307$ for $p = 0.5$ against $p = 0.001$ — **R7**	ARITHMETIC
L22	independent replications	$(1-0.3)^r$; $r = 0 \Rightarrow 1.000000$; $0.7^{12} = 0.013841$ — **R6**	ARITHMETIC
L23	P11	$R = 8.3100$ fm, $E_C = 1.43996454 \times 82/8.31 = 14.2091$ MeV	ARITHMETIC given P11
L24	L23, $kT(2000\,\text{K}) = 0.172347$ eV	$E_C/kT = 8.24 \times 10^7$ — **R8**	ARITHMETIC

step	consumes	produces	status
L25	Avogadro, 1 GeV, $\varepsilon = 10^{-6}$	4.899×10^{17} J = 15.52 GW-station-years — **R8**	ARITHMETIC

9.3 What this chain is and is not

Read the status column. L1–L9 and L12–L25 are arithmetic and a reader can redo every one of them. L10, L11 and the *meaning* half of L13 are postulated. **A chain containing postulated steps is not a proof of its conclusion**, and this one is not offered as one.

What it is offered as is a separation, and the separation is unusually clean here. **The frequency register carries exactly one of this paper's eight results and no part of any other.** A reader who rejects P4 and P8 — who declines to accept that a dimensionless ratio is a frequency or that the essence bridge holds — loses the second half of R1 and the glide of §4.13, and keeps R2, R3, R4, R5, R6, R7, R8 and the whole counter-result untouched. The negative theorem about the *fixed point* does not need the bridge at all; only the negative theorem about the *pitch* does. That is the honest statement of what the frequency register buys in this paper, and it is less than the notation suggests.

Two statuses that must not be misread. The exponent 2 in the essence law is **DEFINED**, not estimated: no fit was performed and the standard procedure for testing a scaling claim [Clauset et al., 2009] was not run. And **the arithmetic of §5 and §6 is not evidence about the tradition's efficacy** — L20 proves precisely that it cannot be, whichever way it comes out.

The verification instance. Where the source supplied a worked numerical case it is used unchanged. The smallest instance exercising the whole of L1–L7 is the carried constructed dyad at $u = 0.692372881355932$, $\eta = 0.44$, $c = 0.40$, $\delta = 0.09$, $N = 12$: $F = 2.153633$, $b = 0.547598406839$, $Q_\infty = +6.0844267427$, cross-checked against $-D\eta/(2.885390\,\delta)$ to 10^{-12}, half-life 7.3496, and seven passes reaching $Q_7 = 2.940232$, which is 48.32 per cent of 6.0844267427. Three multiplications and a logarithm.

10 Limitations

Every check the governing rigour standard specifies is listed here by name, whether or not this paper passes it.

L-1. The dimensional audit does not close. $f(n) = kn^2$ balances only if k carries $[T^{-1}]$, and k is a ratio of two dimensionless quantities. **No operational procedure is supplied by which the constant acquires the dimension of frequency.** The hertz is attached by assertion (§2.3 P-a, §3.4).

L-2. The essence law has no external support. A quadratic relation between a dimensionless count and a frequency has no established analogue. It is badged a postulate; the badge is the honest position and is not a defence.

L-3. The constant is chosen. $k = (9/8)^2/\varphi^2$ combines two real objects into a product with no precedent. It is quoted to fifteen significant figures **because it is defined**, not because anything was measured to that precision, and the deflationary literature on golden-section claims is cited rather than avoided [Markowsky, 1992, Fischler, 1981].

L-4. Octave reduction beyond pitch is unwarranted. Shepard grounds octave equivalence for pitch [Shepard, 1982]; nothing grounds folding a non-acoustic quantity into a 220–440 Hz band and reading meaning from where it lands.

L-5. $N_{\text{eff}} = 1 + (N-1)u$ **is not a Hill number.** It is a linear interpolation between 1 and N and appears in no ecological, economic or information-theoretic literature. The Hill-number results [Hill, 1973, Jost, 2006] apply to $1/\text{HHI}$ and $\exp(H)$ and to nothing else.

L-6. Conjoint measurement was not tested, and scalar reduction is the general form of the problem. Michell's objection [Michell, 1997] stands unanswered for u, η, c and δ, and the

remedies [Luce and Tukey, 1964, Rasch, 1960] were not run. The same objection arrives from the affect literature in a different domain: emotions are individuated by the pattern across independently varying appraisal dimensions [Smith and Ellsworth, 1985], by a synchronisation process in time [Scherer, 2001], and by qualitative core relational theme [Lazarus, 1991] — none of which a scalar carries, and the last of which denies that differing by *what* licenses ordering by *how many*. Action readiness in particular is constitutive of emotion as such rather than a top rung [Frijda, 1986]. **This is one objection, not two**, and §3.5 names it once. Everything entering through u inherits it.

L-7. δ is the hardest parameter to estimate and every rate figure in the paper depends on it. No instrument is supplied here and none was run. The convergence table, the seven- passes result and the stopping rule are all exactly as defensible as an unmeasured decay rate, which is to say: not at all, as measurements. They are exact as arithmetic on a declared parameter.

L-8. Every case is constructed. The carried dyad, the two-party export case, the seven-stage process, the cadence schedules, the closed lineage and the survivorship corpus are chosen parameter sets. Their precision is exact because the parameters are exact; **that precision does not transfer to a measured case.**

L-9. Two founding works of statistical process control are not in this paper's bibliography. Shewhart (1931) and Deming (1986) are the empirical warrant for C5.1 and are named in §5.2 without keys, because `bib/external.bib` contains no entry for either. The available operations-management literature covers the yield-and-rework structure but not the historical claim. **R4's arithmetic is this paper's; its empirical standing rests on works this paper cannot cite.**

L-10. The spacing-effect literature is not in this paper's bibliography either. Ebbinghaus (1885), Cepeda et al. (2006) and Roediger and Karpicke (2006) are named in §5.3 without keys. **No claim in this paper rests on them**, and §5.3 states explicitly that the spacing effect and the concavity theorem are different claims that must not be presented as mutual confirmation.

L-11. The nuclear-physics primary sources are not in this paper's bibliography. Rutherford (1919) and Aleklett et al. (1981) are named in §6.1 without keys. The Coulomb-barrier arithmetic is standard and reproducible from the constants given; the historical attributions are not citable here.

L-12. Wald, Bacon, Fisher, the MRC streptomycin trial, Polanyi, Collins, Durkheim, Rappaport and Yates are likewise named without keys. Each is named for precedent rather than for support, and the modern standard the argument does rest on is cited [Nosek et al., 2018, Open Science Collaboration, 2015].

L-13. The parameter assignment of the historical schedule is an interpretation. §6.5's order-inversion result depends on sorting each of twelve named operations by which parameter it acts on, and that sorting is this framework's reading of a fifteenth-century text [Curzi, 2026], badged as an interpretation there and here. **A different defensible sorting could weaken or reverse the finding**, and no independent coder has been asked to reproduce it.

L-14. No model comparison was performed. No out-of-sample test, no matched null with the same number of free parameters and no information criterion is reported, because nothing was fitted. Where the corpus has attempted a confirmatory fit elsewhere it failed its own baseline test [Curzi, 2026i], and that is the relevant precedent.

L-15. No statistical claims are made and none are supported. There are no p-values, effect sizes or confidence intervals in this paper, because there is no sample. The likelihood ratios of §6.9 are computed from a stated selection model, not estimated from data.

L-16. The exponent is defined, not estimated. The exponent 2 in $f(n) = kn^2$ was not fitted and the standard fitting-and-testing procedure for a scaling claim [Clauset et al., 2009] was not run. **This paper makes no empirical scaling claim.** The geometric decay $(1-\delta)^t$ is likewise a consequence of the model's linearity and not an estimated rate.

L-17. The functional form of the surplus is underdetermined. A sub-linear power law [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the qualitative facts as well as the logarithm [Stumpf and Porter, 2012], and nothing in this paper discriminates among them.

The clipping result of R5 depends on the ceiling $\ln N$, which is a property of the logarithmic form specifically; **under a power-law surplus the 9.56 per cent figure would change.**

L-18. The reflexive specification is not identified. Stated in §3.6 and repeated here: the endogenous social effect in a linear-in-means reflexive loop cannot be recovered from observational data [Manski, 1993, Shalizi and Thomas, 2011]. The remedy adopted is stipulation of the reference set [Brock and Durlauf, 2001], and the claim is accordingly a modelling position rather than an identified result.

L-19. The framework exempts its own parameters from its own reflexivity. If subjects change when modelled, δ, η and u are not structural [Lucas, 1976, Hacking, 1995]. This is not resolved. It bears directly on this paper: **publishing a destination reading may change the arrangement it reads**, and the effect, while real, is probably smaller than the framework's rhetoric implies [Jussim and Harber, 2005].

L-20. Boundary behaviour is only partly examined. $\delta \to 0$ is examined and is a pathology (§4 FM 4.13, and Curzi, 2026a); $\delta \to 1$ makes the ladder arrive in one pass and the fixed point equal the drive, which is examined implicitly and stated nowhere in the source; $\eta \to 0$ makes c/η diverge and places every arrangement above every ceiling, which is a real constraint on the model's domain and **is not examined in the source**; $u \to 0$ is examined and gives the ground essence; $N \to \infty$ raises the ceiling without bound, which the coordination-cost literature says is not available in practice [Brooks, 1975]. These limits constrain the domain and are stated here because the source does not state them all.

L-21. Ergodicity is assumed and not argued. Q accumulates over time for one arrangement from parameters that would naturally be estimated across arrangements; the time and ensemble averages of a multiplicative process are not the same [Peters, 2019].

L-22. Agreement is not correctness. Nothing in §5 measures validity. The concordance method measures independent arrival, and independent arrival establishes findability, not truth. The distinction between reliability and validity is old and is not repaired by more agreement [Cronbach and Meehl, 1955, Campbell and Fiske, 1959].

L-23. u is symmetric and psychological similarity is not. Similarity judgments are asymmetric and violate the triangle inequality [Tversky, 1977]; if u is read as a similarity measure over people it cannot assume metric behaviour, and this framework's u does.

L-24. The one-hot artefact is real and is reported rather than removed. In §5.16 an operator whose state vector has disjoint support from the pair's returns $u = 1.000000$ against both by construction. **That reading is an artefact of the encoding and not a fact about the operator**, and it is kept in the table with the artefact labelled because removing it would hide a failure mode the framework's users will meet.

A discrepancy found in the source and left standing

The source states the counterfeit-arrival identity twice, and the two statements are not equivalent. It first states the invariant as the **sum of participants' fixed points**, which is correct only when every participant carries the same decay rate; the general invariant is the **sum of drives**. The source's own worked case happens to satisfy the equal-δ condition, so the error is invisible there. **§4.15a computes the case that separates them**: with $\delta_P = 0.06$ and $\delta_R = 0.16$, the sum of drives holds at $+0.643688$ while the plain sum *rises* from $+9.727202$ to $+11.810535$, by $1/0.06 - 1/0.16 = 10.416667$ per unit exported. **The naive test does not merely fail to detect the export; it reports a gain, and it reports one precisely when cost is pushed onto whoever decays faster.** The correction is printed in place in §4.15 rather than applied silently, and it is recorded here because a reader working from the earlier formulation would get the export check backwards.

A second, smaller discrepancy. The source writes the octave conversion constant as 2.885390 inside formulae and as 2.885390081777927 when quoting it exactly. Every computed figure in this paper uses the exact value $2/\ln 2$; the six-figure form is a shorthand.

11 Conclusion

What is established. Under the framework's storage law, iteration closes a fixed *proportion* of the remaining gap at every pass, and the closure is an equality rather than a bound: $|Q_t - Q_\infty| = (1-\delta)^t |Q_0 - Q_\infty|$. **Refinement cannot move the fixed point**, because $Q_\infty = (\eta F - c)/\delta$ contains no iteration count, and **the ladder is silent in the frequency register**, because f_{rel} is a function of u and N alone. What repetition buys is arrival time only, and arrival time is a function of the decay rate and of nothing else. **At $\delta = 0.09$ the half-life is 7.3496 passes and 99 per cent closure takes 48.83, so the tradition's canonical seven passes close 48.32 per cent of the gap — and only 6.79 per cent at $\delta = 0.01$. Seven passes would close 95 per cent only at $\delta = 0.348164$.** Applied to a decaying arrangement the same ladder converges at exactly the same speed on $Q_\infty = -5.01$ instead of $+6.08$.

The audit's numbers, in the form other papers cite them. Staged process control is priced: a seven-stage process at $q = 0.10$, moving from terminal inspection to per-stage checking at $d = 0.90$, raises final yield from 47.83 per cent to 93.21 per cent at 0.63 reworks per unit. Cadence is priced: an annual occasion carrying a nominal 26 nats delivers 2.4849 after the $\ln 12$ ceiling clips it — 9.56 per cent of a weekly practice, 90.44 per cent lost to clipping — and $0.91^{52} = 0.007403$ of one impulse survives the year. Secrecy is priced: at zero independent replications per generation an error survives with probability 1.000000 indefinitely, against 0.013841 after twelve generations at one replication. The record is priced: 500 credible accounts require 1,000 practitioners at $p = 0.5$ and 500,000,000 at $p = 10^{-6}$, so the likelihood ratio is exactly 1.000000 at every corpus size. And transmutation is priced: one gram of gold-197 requires 15.52 gigawatt-station-years at 1 GeV and yield 10^{-6}, and 1.552×10^4 at yield 10^{-9}.

What is proposed rather than shown. That a dimensionless ratio of a tuning interval to the square of the golden section is a frequency; that the relationship pitch is $k(1 + (N-1)u)^2$; and that octave equivalence may be carried from pitch to a quantity that is not a pitch. **These are postulates of this framework with no external support**, they are labelled as such at every appearance, and **seven of this paper's eight results survive their deletion entirely.**

What is not claimed at all. No person, pair, substance or institution is claimed to resonate at any frequency named here. The component classes are not diagnoses and no reading in this paper is clinical, diagnostic, therapeutic or prognostic. Four traditional practices survive an audit; that is four practices, and not a vindication of the tradition that carried them.

12 Declarations

Data and code availability. No dataset was collected for this study and none exists. Every number in the paper is the output of exact arithmetic on parameters stated in the text; every case is constructed and is labelled as constructed at the point of use. The figures are generated by the Python scripts named in their captions, which take their inputs from the tables in the text and are available from the author on request. **No empirical measurement of any human being, any relation, any institution or any physical process underlies any figure, table or number in this paper.**

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the analysis, the figures and the manuscript.

Ethics. No human-subjects research was conducted and no personal data were processed. The constructed dyad, the constructed two-party arrangement and the constructed lineage are invented parameter sets; they are not derived from, and do not describe, any person. **This paper asserts perceptual claims — that magnitude estimation delivers ratio-scaled intensities, that octave equivalence is a usable structure beyond pitch, and that interval arithmetic in cents is the correct apparatus — without any perceptual experiment behind them.** Those claims rest entirely on the cited literature, are limited by that literature's own boundaries (§2.4), and are listed among this paper's limitations (L-4, L-6). No listener was asked anything.

Preregistration. No test in this paper was preregistered, because no test was run. H3, the one genuinely empirical hypothesis, is stated here with its falsifier so that it can be preregistered by whoever runs it.

Historical and archival claims. The historical material in §5 and §6 is drawn from secondary literature, several key items of which are absent from this paper's bibliographic database and are named in the text without citation keys (L-9 through L-12). **No archival research was conducted and no primary source was consulted for this paper.**

Clinical disclaimer. The twelve component classes are not diagnostic categories. Nothing in this paper is an assessment, a screening instrument, a treatment or advice about any person's health, and nothing in it may be presented as such.

Chapter 48

A Disclosed Sonification Protocol for Sequence-Encoded Content, and the Missing Transducer That Makes It a Research Design and Not a Technique

Abstract. We specify, end to end, a deterministic pipeline that encodes text as bits, bits as nucleotides at one bit per base, adds a routing header, sonifies by the essence law $f = kn^2$, transposes, and renders inside the Gabor bound — and then state that no mechanism is known by which a cell could read it. Two design results are proved. Assigning symbol indices $n = 1, 2, 3, 4$ collapses three of four symbols onto one pitch class, because $(2n)^2 = 4n^2$; the odd assignment $n = 1, 3, 5, 7$ gives four distinct classes whose closest pair in hertz is 990.055901 Hz apart and whose closest pair as an interval is 1165.0244 cents apart, with a Gabor floor of 0.080377 ms. The state layer's binding separation is the $50/49$, 34.9756-cent minimum at $N = 12$, giving $\Delta t \geq 10.2882$ ms, and single-transposition audible capacity is $N \leq 25$ at $j = 6$ only. Content is protected structurally rather than by disclosure alone: a corpus closed under negation has resultant exactly zero, verified at $N = 4, 8, 12$ and 24, and at $u = 1$ carries $F = \ln 12 = 2.484907$ nats, exactly the selection cost, sitting on the ceiling with residual 8.88×10^{-16}. Four of twelve directive pairs fail the equal-force check, 33.33%, and are named rather than repaired. The open problem is stated formally: no known structure produces a response depending on the symbol string as an *order*, as distinct from permutation-invariant functionals of the waveform. Bulk compression is closed by arithmetic at strains of 10^{-11} to 10^{-10}, and a single null condition bounds 0.000801% of a 124,800-cell design grid. No experiment was run, no dataset was collected and no biological claim is made.

Keywords: sonification; genomic audio; no-subliminal constraint; Gabor bound; symbol separation; closure under negation; power analysis; open problem

Series: C. Frequency, Essence, and Mind **Paper:** CCP-048

1 Introduction

A pipeline that turns written text into a sound is easy to specify and easy to over-describe. This paper does the first completely and refuses the second in print. It specifies, end to end

and to the digit, a deterministic chain that encodes text as bits, bits as nucleotides at one bit per base, prepends a routing header, sonifies by the essence law $f = kn^2$, transposes by a fixed integer number of octaves, and renders inside the time–frequency bound. **And then it states, in the same document and with the same care, that no mechanism is known by which a cell could read any of it.**

The organising claim is that the second statement is what makes the first worth publishing. A fully specified procedure whose biological effect is unknown is a research design. The same procedure with the gap omitted is a technique, and a technique is what this material must not become. The paper is therefore written in the order it is: constraint, procedure, gap — with the gap last, so that a reader meets the specification while it still looks promising and then meets the missing link while the specification is fresh.

Two design results are proved and both are negative in form. The first is an index-assignment result: assigning symbol indices $n = 1, 2, 3, 4$ collapses three of four symbols onto one pitch class, because $(2n)^2 = 4n^2$ makes doubling an index an exact double octave. The naive assignment is unusable, and the reason is arithmetic rather than taste. The odd assignment $n = 1, 3, 5, 7$ gives four distinct pitch classes whose closest pair in hertz is 990.055901 Hz apart and whose closest pair as an interval is 1165.0244 cents apart — **two different pairs, which the source states as one, and §10 prints the correction** — giving a Gabor floor of 0.080377 ms. The second is structural: **a corpus of directives closed under negation has resultant exactly zero** — verified at $N = 4, 8, 12, 24$ — and at $u = 1$ it carries $F = \ln 12 = 2.484907$ nats, which is exactly the selection cost the receiver pays, 3.584963 bits. **The corpus does not point; it spans, and the direction is contributed by whoever selects from it.**

Three constraints from the physics are computed rather than assumed. The Gabor bound $\Delta t\,\Delta f \geq 1/(4\pi) = 0.0795774715$ [Gabor, 1946] gives a floor of 0.080377 ms at the symbol layer and 10.2882 ms at the state layer, whose binding separation is the 34.9756-cent minimum of the folded lattice at $N = 12$. Single-transposition audible capacity is $N \leq 25$, and only at $j = 6$; at $N = 26$ no admissible integer transposition keeps the set inside 20 Hz to 20 kHz. And one candidate route is closed by arithmetic: at 40 to 120 dB SPL the bulk compressional strain in an aqueous medium is 9.09×10^{-13} to 9.09×10^{-9}, so bulk compression is not the route — though structural and local routes are untouched by that number and are the ones the established literature actually describes.

The paper's honesty apparatus is also computed. A corpus that a careful author wrote intending balance fails an equal-force check on four of twelve pairs — 33.33 per cent — and the failures are named rather than repaired, because *the rate at which a careful author fails this check* is the useful number. At $\sigma = 0.35$ $\log_2$ units a targeted twelve-locus design with Bonferroni correction cannot see below a 1.68-fold change at $n = 3$ or a 1.33-fold change at $n = 20$. And a single well-run null condition bounds 0.000801 per cent of a deliberately modest 124,800-cell design grid, which protects a null from being inflated into a refutation and protects the field from a claimant who runs one condition and announces a channel.

What this paper does not claim, stated at the front rather than at the back. No biological effect of any kind. No clinical, medical or therapeutic claim. No claim that DNA is edited, altered or damaged: nothing in the pipeline is a molecule, and the nucleotide string is a string in the computer-science sense. No claim that any tone bears a physical relation to any base — **cytosine does not absorb at 1113.8129 Hz, does not resonate at it, and is not affected by it.** Every frequency here is an *essence* frequency, a position on a logarithmic scale of surplus. The constant $k = (9/8)^2/\varphi^2$ is a *defined* number; attaching hertz to it is a separate and much stronger claim and is postulated (§2.3). The standard critique of golden-section claims [Markowsky, 1992] is cited here deliberately and again wherever φ appears.

One structural disclosure about the bibliography. The biological literature this paper describes — nucleotide data storage, guide-RNA addressing, self-processing effector arrays, mechanotransduction, force-to-transcription coupling, sonogenetics, reprogramming, and the open-label placebo literature of §4.4 — **has no corresponding cluster in the bibliographic database governing this corpus.** Every such work is therefore named in the text with its journal, volume, page and year, and **without a citation key**, because inventing one or attaching

an unrelated key would be worse than the gap. §10 lists them, and the gap is a defect of the bibliography rather than a claim that the works do not exist.

Contributions

1. A complete, deterministic, disclosed encoding and rendering chain, with a verified round-trip SHA-256 identity on a constructed payload.
2. The index-assignment result: $n = 1, 2, 3, 4$ collapses three symbols to one pitch class; $n = 1, 3, 5, 7$ gives four classes at a minimum gap of 990.055901 Hz.
3. Determination by choice applied to a twelve-state ladder: resultant exactly zero under closure under negation, and $F = \ln 12$ equal to the selection cost at $u = 1$.
4. The Gabor constraints at both layers, and the 128-fold difference between them that forces the two disclosed layers apart in time.
5. Single-transposition audible capacity verified at $N \leq 25$, $j = 6$ only.
6. A closed route: bulk compression, by $s = p/K$ at plausible listening levels.
7. The detectability floors, and the design-space bound that keeps a null and a positive result the same size.
8. **The open problem, stated formally**: no established mechanism produces a transcriptional response depending on the symbol string *as an order*, as distinct from permutation-invariant functionals of the waveform.

2 Hypotheses, definitions, postulates and modelling choices

2.1 Hypotheses proper

H1 — the ordering hypothesis. A cell's transcriptional response to a sonified symbol stream depends on the *ordering* of the symbols, not only on permutation-invariant functionals of the waveform. **Predicted:** an ordered stimulus and its symbol-permuted twin — identical in total energy, duration, symbol histogram, long-term power spectrum, peak level, RMS level and crest factor — produce measurably different transcript abundances at preregistered loci. **Falsified by:** a null contrast at a sample size sufficient to detect the declared smallest effect of interest at the declared power, with the detectable floor published alongside. **Status: NOT TESTED. No experiment has been run and no data exist.** §6 states that no mechanism is known which would produce the predicted result, and this paper does not assert it. The hypothesis is stated so that its falsifier is on record.

H2 — the index-collapse hypothesis. Under the essence law, symbol indices differing by a power of two are indistinguishable to any listener attending to pitch class. **Predicted:** $n = 1, 2, 4, 8$ fold to one class; $n = 1, 3, 5, 7$ fold to four. **Falsified by:** an arithmetic error in $(2n)^2 = 4n^2$; the claim is a consequence of the definition and is not empirically falsifiable beyond that. **Status:** DERIVED. Since nothing could falsify it empirically, **it is restated in §2.2 as a definition-level consequence rather than left dressed as a hypothesis.**

H3 — the equal-force hypothesis. An author writing a deliberately balanced corpus of directives fails a mechanical equal-force check at a non-negligible rate. **Predicted:** independent authors writing paired directives under the stated rules fail character- and word-count balance on a measurable fraction of pairs. **Falsified by:** a set of independently authored corpora with a failure rate indistinguishable from zero. **Status: one instance measured, and one instance is not a rate.** The corpus of §4.28, written for this paper by its author intending balance, failed on 4 of 12 pairs. That is a single observation of a single author and is reported as such; it is not an estimate of anything.

H4 — the separation hypothesis. The two disclosed layers of the rendering are separable by ear because they occupy disjoint registers. **Predicted:** the state band clears the nearest symbol tone by 995.9922 cents. **Falsified by:** a change to either transposition that brings them within a critical band; the computation must be redone after any such change. **Status:** COMPUTED for one specific pair of transposition indices. **No listener was asked anything**, and the perceptual claim implied by the word *separable* is not tested here (§12).

2.2 Definitions — true by stipulation, not falsifiable

- **Def. 1.** A *narrative*, at this layer, is a finite, written, published set of directives, each naming one state of the twelve-state ladder and urging one action with respect to it.
- **Def. 2.** $u := 1 - (\hat{x} \cdot \hat{y})^2$; $F := \ln(1 + (N-1)u)$.
- **Def. 3.** $k := (9/8)^2/\varphi^2 = 0.483425732988414$. **A defined constant**, exact by construction; its fifteen figures are exact digits of a definition and carry no measurement uncertainty because there is no measurement.
- **Def. 4.** The *pitch class* of an index n is the odd part of n. It follows immediately that $f(2n) = 4f(n)$, that an indexed set $1 \ldots N$ holds exactly $\lceil N/2 \rceil$ distinct classes, and that $n = 1, 2, 3, 4$ yields odd parts $1, 1, 3, 1$. **This is H2 restated as what it actually is: a consequence of the definitions, with no empirical content.**
- **Def. 5.** A *symbol-permuted twin* of a stimulus is a rendering of a permutation of the same symbol string, so that every permutation-invariant functional of the waveform is unchanged.

2.3 Postulates — asserted by this framework, with no external support

- **P-a. The attachment of hertz to k.** Computing $(9/8)^2/\varphi^2 = 0.483425732988414$ is arithmetic. **Asserting that this dimensionless number *is a frequency* is a separate and much stronger claim**, with no operational procedure supplied by which the constant acquires the dimension of inverse time. Postulated, and restated in the dimensional audit of §3.4.
- **P-b. The essence law $f(n) = kn^2$.** No external literature supports it. **The exponent 2 is defined by the model, not estimated from data**; no fit was performed, and the goodness-of-fit and likelihood-ratio procedure a claimed scaling relation requires [Clauset et al., 2009] was **not run, because there is no data to run it on.**
- **P-c. The affine effective count $N_{\text{eff}} = 1+(N-1)u$.** A design axiom, not a derived result, and **not a Hill number of any order.** The Hill family [Hill, 1973, Jost, 2006] and its rediscovery in political science [Laakso and Taagepera, 1979] ground $1/\text{HHI}$ and $\exp(H)$ exactly and ground this construction not at all.
- **P-d. The alphabetical ordinal on the four bases.** Applying $f(n) = kn^2$ to symbols requires an ordinal, and the one used is alphabetical. **It is not defended as meaningful, only as published, stable and conventional**, and it is a weaker warrant than the twelve-state ladder carries.

2.4 Modelling choices — legitimate, but choices

- **M-a. The unit-vector model of a directive.** Directives are modelled as unit vectors in a real inner-product space, with negation as $-v$. Real directives are not unit vectors and "let the charge rise" is not literally the additive inverse of "let the charge fall." **What the model captures rigorously is one thing only:** a set constructed so that its elements cancel or span has no resultant, so a receiver who acts on it must have selected. That conclusion depends only on the corpus containing genuinely mutually exclusive directions, which is checked by reading the text — as §4.29 checks it, and fails.
- **M-b. Octave reduction applied to a non-acoustic quantity.** Octave equivalence is a perceptual structure formalised for **pitch** [Shepard, 1982], with a measured perceptual size

[Demany et al., 2013]. Carrying it to a quantity that is not a pitch is a choice of this framework, and the interpretation of a folded value is unsupported even where the operation is well defined.
- **M-c.** $N = 12$**.** Twelve exceeds the classical span of immediate absolute judgment [Miller, 1956], which is a reason for caution rather than a reason for twelve.
- **M-d. Bonferroni rather than a false-discovery-rate procedure.** Chosen for transparency and because it is the most conservative common choice. A Benjamini–Hochberg procedure would be entirely defensible **provided it is the preregistered choice rather than the one arrived at after the conservative analysis returned nothing** [Simmons et al., 2011, Gelman and Loken, 2013].
- **M-e. Scalar reduction of a structured state.** Every reading here enters through u, a scalar summary of a vector of intensities over twelve component classes. §3.5 states what that costs.

2.5 What is expository

Sections 4.17–4.20, 5.1–5.5, and the specification and failure-mode blocks throughout are expository: they set out procedure, disclosure requirements, gates and failure modes rather than argue a claim. Their organising question is *what must be published alongside a rendering for a third party to be able to check it, refuse it, and reproduce it.*

3 Framework, dimensional audit, and method

3.1 What kind of claim each layer makes

Analogy status, stated once for the whole paper. The Gabor bound of §5.15 is **category 1, a derivation**: it is a theorem about the Fourier transform, its premises are satisfied by any finite-duration signal, and the rendering genuinely satisfies them. The essence law of §5.12 is **category 3, an analogy with no formal warrant**, badged as a postulate at every appearance. The surplus functional of §4.6 is **category 2, a shared mathematical form**: the same logarithm that appears in Shannon's entropy [Shannon, 1948] appears here, and no claim is made that information-theoretic results transfer.

The uncertainty relation must not be misread, and it is easy to misread here. $\Delta t\,\Delta f \geq 1/(4\pi)$ is Gabor's [Gabor, 1946] and is classical Fourier analysis with no quantum content. It has the same algebraic form as $\sigma_x \sigma_p \geq \hbar/2$ [Kennard, 1927], which is **a fact about state preparation, not about measurement disturbance**, and **nothing in this paper transfers any physical result from the second to the first.** The framework's conjugate-uncertainty paper proves the Gabor limit by the same variational argument with $p = \hbar k$ deleted [Curzi, 2026z], and that deletion is the whole of the relationship.

3.2 Dimensional audit

symbol	quantity	dimension	unit	status
n, N	component / category index	$[1]$	—	absolute scale, a count
u	non-parallelism	$[1]$	—	dimensionless by construction, $u \in [0, 1]$
F	surplus	$[1]$	nats	a logarithm of a count

symbol	quantity	dimension	unit	status
D	octave deficit	$[1]$	octaves	a ratio of logarithms; k cancels
k	reference constant	$[T^{-1}]$ **asserted**	Hz **asserted**	**defined as a ratio (dimensionless); the unit is postulated**
f	essence frequency	$[T^{-1}]$	Hz	inherits its dimension entirely from k
j	transposition index	$[1]$	octaves	an integer exponent of 2
Δt	symbol duration	$[T]$	s	**measured on the rendering; a real time**
Δf	spectral width	$[T^{-1}]$	Hz	**a real bandwidth of a real waveform**
p	acoustic pressure	$[ML^{-1}T^{-2}]$	Pa	**measured; a real physical quantity**
K	bulk modulus of water	$[ML^{-1}T^{-2}]$	Pa	standard constant, 2.2×10^9
s	bulk strain	$[1]$	—	p/K; dimensionless, as strain must be
L, bits, bases	payload length	$[1]$	—	counts
q, d	defect / detection probability	$[1]$	—	probabilities
σ, $\delta_{\log_2}$	dispersion / effect in log abundance	$[1]$	$\log_2$ units	ratio scale on a log transform

Does the arithmetic balance? The Gabor bound does, exactly: Δt in seconds times Δf in hertz is dimensionless, and $1/(4\pi)$ is a pure number, so the inequality is dimensionally clean. **This is the one place in the paper where a physical bound is applied to a physical quantity and the units close without an assumption.** The strain calculation likewise closes: pressure over bulk modulus is Pa/Pa, and strain is properly dimensionless.

The essence law does not close. $f(n) = kn^2$ balances only if k carries $[T^{-1}]$, and k is computed as a ratio of two dimensionless quantities — a frequency ratio and a geometric ratio. **No operational procedure is supplied by which it acquires the dimension of frequency.** This is the paper's largest dimensional exposure and it is listed by name in §10.

A consequence that matters here more than elsewhere in the corpus. The rendered waveform is a real sound with a real spectrum, and Δt and Δf are properties of it. The *value* of each tone, however, comes from the essence law, whose hertz is postulated. **A rendering therefore has real physics downstream of a postulated number**, and the honest statement is that the acoustics of the file are ordinary while the reason those particular frequencies were chosen is not established.

3.3 The surplus functional

$F(u) = \ln(1 + (N-1)u)$ is the framework's surplus functional. The ceiling $F \leq \ln N$ is the same bound that requisite-variety arguments impose on any regulator [Ashby, 1956]. **The functional form is underdetermined by the evidence**: a sub-linear power law [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the qualitative facts as well, and many mechanisms produce indistinguishable curves [Stumpf and Porter, 2012]. The strongest warrant for the logarithm is not psychophysical: it follows from multiplicative dynamics with no psychological assumption at all [Peters and Gell-Mann, 2016], and that is the warrant this paper's arguments need. **Which construct of difference u measures must be declared** [Harrison and Klein, 2007]: here it is a disparity measure, symmetric in its arguments, which psychological similarity is not [Tversky, 1977].

3.4 The constant k

$$k := \frac{(9/8)^2}{\varphi^2} = 0.483425732988414 \quad \text{(defined; exact)}$$

Both components are real and documented — the Pythagorean whole tone [Barbour, 1951] and Euclid's extreme and mean ratio [Euclid, 1956], with $\varphi^2 = \varphi + 1$ a proved identity [Vajda, 1989]. **Their product, expressed as a frequency in hertz, has no precedent.** The literature on golden-section claims is overwhelmingly deflationary [Markowsky, 1992, Fischler, 1981, Höge, 1997], and **this framework's position is the one those critiques leave open and no wider**: φ is being *chosen* as a constant, not *found* in anything measured. Because every ratio in the alphabet is a ratio of squares, k cancels from all of them [Curzi, 2026k], and no separation figure in this paper depends on its value.

3.5 Scale type, and the objection this framework inherits twice

quantity	scale type claimed	what licenses it
n, N, base counts, bit counts	absolute	cardinality; identity transformation group only
u	ratio	constructed from an inner product of unit vectors
elicited intensities	ratio, up to a common multiplier	**magnitude estimation, assumed** [Stevens, 1957]; not tested here
Δt, Δf, p	ratio	physical measurement with a true zero
$\log_2$ abundance	interval	a logarithm; differences are meaningful, the zero is not
pitch class	nominal	the odd part of an index; **no ordering is licensed on it**

The framework's formula-validation work fixes what licenses a quotient at all [Curzi, 2026t]: **a quotient X/Y is meaningful if and only if the admissible transformation group is similarity — that is, if and only if the scale is ratio** — with four demonstrations showing a 25 per cent shift under an affine relabel, reversed orderings under $g(x) = 2^x$, any ratio at all reachable via $(7/4)^k$, and an undefined ratio under $g(x) = x - 2.5$. That result is why the counts in this paper carry ratios and the elicited intensities carry them only by assumption.

Michell's objection applies here in full: psychology routinely assumes its attributes are quantitative without testing whether they are [Michell, 1997]. Conjoint measurement is the

formal remedy [Luce and Tukey, 1964, Krantz et al., 1971] and **it was not run for any quantity in this paper.**

And the same objection arrives a second time from the affect literature, which is why this paper names the pattern once rather than meeting each instance as a fresh surprise. The pattern is *scalar reduction*: collapsing a structured state into a single number, and then ordering the numbers. Emotions are individuated by the **pattern across** independently varying appraisal dimensions [Smith and Ellsworth, 1985]; a state is a synchronisation process in time rather than a point [Scherer, 2001]; individuation runs by qualitative core relational theme, so that differing by *what* is appraised does not license ordering by *how many* [Lazarus, 1991]; and action readiness is constitutive of emotion as such rather than a top rung above the others [Frijda, 1986] — which contradicts the ladder's twelfth class specifically. Two dimensions of core affect are well supported [Russell, 1980]; twelve ordered rungs are not.

Everything in this paper that enters through u or through a state index inherits that objection. What survives it is everything that does not: the encoding chain, the round-trip identity, the index-collapse result, the separation figures, the Gabor arithmetic, the strain calculation, the detectability floors, and the design-space bound. Those are statements about a model and a waveform given their parameters, and they hold whatever the parameters turn out to mean.

3.6 Identification and reflexivity

The framework's founding construction is a reflexive loop [Curzi, 2026], and **the endogenous social effect in a linear-in-means reflexive specification is not identified from observational data** [Manski, 1993]; influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011]. **The remedy this paper adopts is stipulation**, with the reference set — the declared category set of size N and the declared corpus — fixed in advance [Brock and Durlauf, 2001], and the reflexive claim is presented as a modelling position rather than an identified empirical result.

A specific reflexive exposure attaches to this paper and is unusually direct. The disclosure package is delivered to the listener before exposure, which means the listener knows exactly what is being sent. That is the design, and it is also a manipulation of expectancy: the act of taking part in a study changes behaviour independently of the intervention [McCambridge et al., 2014], expectancy effects on measurement are documented [Rosenthal and Fode, 1963, Rosenthal and Jacobson, 1968], and the best long-run assessment finds such effects real but small, usually non-accumulating and often dissipating [Jussim and Harber, 2005]. **Disclosure does not remove the expectancy channel; it makes it visible and shared**, which is why §6's design carries a scrambled-sequence arm rather than relying on any comparison against silence. If systems change when modelled, the framework's own parameters are not exempt [Lucas, 1976, Hacking, 1995], and this paper does not resolve that.

3.7 Method

Every number in this paper is the output of exact arithmetic on stated parameters, computed in double precision. **No observational data were collected, no human subjects participated, no cells were exposed to anything, and no empirical dataset underlies any figure or table.** The payload is a constructed corpus written for this paper. Hashes are SHA-256 of the stated byte strings. The power calculations are arithmetic on declared inputs — the dispersion σ and the number of tests are values a reader supplies from their own system, not values this paper asserts about biology. The figures are produced by the scripts named in their captions.

4 Disclosure as a design constraint, and determination by choice

The framework's audit of a traditional operation set asked what a set of classical operations were doing when they worked, mapped the survivors onto operations on indices and on u,

and printed the ones that failed [Curzi, 2026]. This paper asks a harder question, and asks it in the register where the framework has the least standing and the most to lose: what would it mean to write a corpus of directives *as data*, address that data, render it as tone, and play the tone in a room with a person in it.

The honest answer arrives in three parts and they must be kept in this order.

The first part is that the encoding chain is buildable, deterministic, and auditable end to end. Text becomes bits, bits become a nucleotide string, the string acquires a routing header, the string is sonified by the framework's essence law, and the result is a sound. Every one of those steps is arithmetic. Section 5 specifies them to the digit and demonstrates a round-trip hash.

The second part is that nothing whatever is known about what such a sound does to a cell. Section 6 states that gap in full and does not soften it. **There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal**, and every sentence in this section and the next is written underneath that sentence.

The third part is the one this section exists for, and it is a constraint rather than a finding. The chain above has a design decision at its centre that is not an engineering decision at all. It is the decision about whether the listener can hear what is being sent to them. That decision has exactly one permissible answer under this framework, the answer is not negotiable at any point, in any version, for any purpose, and it is stated in the next paragraph rather than at the end, because a constraint reached as a conclusion is a constraint that has already been weighed against something.

The law

4.1 The constraint, stated first **NO SUBLIMINAL CONTENT OF ANY KIND. NOTHING MASKED. NOTHING BELOW THE THRESHOLD OF HEARING. NOTHING CONCEALED FROM THE LISTENER. THE INFORMATION IS THE TONE.**

This is the framework's standing disclosure constraint, and it governs this section, Section 5, and everything either of them touches. It is the same constraint that the corpus's channel-separation work states in its strongest form [Curzi, 2026]: **channel separation is strengthened by disclosure, because its effectiveness is entirely independent of whether anyone knows it is in place, so every element must be publishable — non-disclosure is refused as a design premise, not merely discouraged.** The design tradition it belongs to is old and explicit: a mechanism whose effectiveness depends on the mechanism being secret is badly designed [Kerckhoffs, 1883], open design is a security principle rather than a concession [Saltzer and Schroeder, 1975], and publicity is the standard corrective for a practice that cannot otherwise be checked [Brandeis, 1914]. It is not a stylistic preference, it is not an aesthetic, it is not a default that a sufficiently good reason could override, and it is not the sort of rule that gets an exception for research use, for a demonstration, for a control condition, or for a version that is only ever played to the person who built it.

The source document behind the essence law recommended embedding subliminal content, and this framework has declined that recommendation at every point since: for built acoustic space [Curzi, 2026l], for content generally [Curzi, 2026ad], and here for psychological content specifically, which is the hardest case in the series, because psychological content is content *about the listener* and the temptation to route it past their judgment is correspondingly larger.

4.2 Why it is ethical and not stylistic, and why the difference matters A stylistic constraint and an ethical constraint fail in different ways, and conflating them is the mechanism by which the second quietly turns into the first.

Violating a stylistic constraint yields a worse version of the same object: use equal temperament where the ratios were exact and the chord is muddier than it should be, but the

artefact is the same kind of artefact, degraded. Violating an ethical constraint does not yield a degraded version of the object. It yields a different object, outside the scope of the method, and no amount of performance recovers its standing. **A masked rendering is not a lower-fidelity disclosed rendering; it is a different thing that this framework does not describe, does not specify, and will not help anyone build.**

The practical test for which kind of constraint is in play is simple and worth applying elsewhere: ask what would happen if the constraint were relaxed and the result were *better* by the method's own measure. A stylistic constraint yields; that is what makes it stylistic. An ethical constraint does not, and a constraint that yields to a good enough result was never ethical to begin with — it was a preference wearing the word.

4.3 The argument, which depends on no empirical claim

A message hidden from its recipient cannot be consented to.

The framework's work on performative commitment supplies the complementary positive case [Curzi, 2026]: a spoken commitment analysed as a performative aimed at the speaker, supported by three converging mechanisms — consistency pressure, self-perception from overheard testimony, and re-pricing of later defection. **Every one of those three requires the speaker to have heard what they said.** A concealment architecture removes the mechanism it was trying to exploit.

That is the whole argument, and its independence from evidence is its strength. It does not depend on whether masked content works. It does not become weaker if masked content turns out to have no effect at all, and it does not become stronger if masked content turns out to be powerful. It is a statement about what consent requires, and it would stand in a world where the entire subliminal literature had replicated perfectly and in a world where none of it had.

Unpacked, the argument runs: consent requires knowing what one is consenting to. A signal engineered so that the listener cannot perceive it has been engineered so that the listener cannot evaluate it, cannot weigh it against their own judgment, cannot compare it with what they were told it was, cannot refuse it selectively, and cannot report on it afterwards to anyone else. Every one of those five capacities is a form of the capacity to decline. Removing them is not an incidental cost of a concealment architecture. It is what a concealment architecture *is*.

And there is a second argument, specific to this series, which makes concealment incoherent here before it makes it wrong. The whole architecture rests on the subject's own latent material. The framework's addressing work establishes that no addressing operation, at any signal strength or in any modality, can cause a cell to express content its genome and accessible regulatory repertoire do not already contain [Curzi, 2026y]. Its operation audit establishes that transmutation is re-indexing and not creation [Curzi, 2026]. And its channel-gain postulate establishes that belief acts only through a coupling with real gain [Curzi, 2026ab]. In every one of those results the person is not the object the method acts on — the person is the part of the circuit that does the acting. **A protocol built on the listener's own participation, which then conceals itself from the listener, has removed from the mechanism the only component that was supposed to be doing the work.** It is self-defeating before it is unethical, and it is unethical.

The general form is: no intervention may remove a person's ability to decline. Consent is structural here, not advisory, exactly as the framework's closure law requires of relational tuning [Curzi, 2026]. The difference here is only that the temptation changes shape. In a room the temptation is to hide a tone. In a corpus it is to hide a meaning.

4.4 The convergence with open-label placebo

There is a version of this constraint that is merely expensive, and a version that costs nothing that matters. Which one applies is an empirical question, and it has an answer.

Before 2010 the received account of placebo response was that it required the recipient to believe they were receiving something active, which requires telling them something false, which the constraint above forbids absolutely. Under that account the constraint and the channel were in direct conflict, and under this framework the constraint would have won and the channel would simply have been closed.

The open-label result dissolves the conflict. Kaptchuk et al. (*PLoS ONE*, 2010) randomised patients with irritable bowel syndrome to an open-label placebo arm — pills presented explicitly and truthfully as inert, with an explanation that placebo effects can be substantial — against a no-treatment control with the patient-provider interaction matched, which is the design feature that makes the comparison mean anything. The open-label arm improved significantly more than the control on the trial's primary outcomes. Later work extended the finding to other conditions (Carvalho et al., *Pain*, 2016), and a meta-analysis of open-label placebo randomised trials (von Wernsdorff et al., *Scientific Reports*, 2021) reported a positive overall effect. The framework's treatment of that literature elsewhere states its limits: the trials are mostly small, the outcomes are mostly subjective and self-reported, the effect sizes are modest, and the field is active. **None of the four works named in this paragraph is in this paper's bibliographic database; all are named without citation keys, and no percentage is quoted from memory.**

What the result establishes for this section is narrow and exact. **A documented effect of this family does not require deception; therefore disclosure is not the thing that would have to be sacrificed to have an effect at all, and the constraint costs nothing that matters.** That is a statement about the availability of a channel under disclosure. It is emphatically not a statement that the method of Section 5 has an effect, that any effect has been measured, or that placebo findings transfer to acoustic sonification of nucleotide sequence. Nothing in that literature is evidence about this method, and this paper does not use it as any.

Two errors to avoid, both of which the framework's two-column concordance rule exists to prevent [Curzi, 2026w], whose Result One is the refinement that keeps the method honest: **the evidence column carries what could have come out otherwise, and the precedent column carries only that the idea was independently arrived at and preserved — two sources agreeing does not double the evidence when only one of them is evidence.** The first is treating the open-label literature as support for the protocol; it is not, it is support for the compatibility of disclosure with *some* belief-mediated effects in *other* domains. The second is the symmetric error of treating the constraint as costless *because* of that literature — the constraint would be binding here even if open-label placebo had failed. The literature removes an excuse, not the rule.

4.5 Determination by choice — the structural result Disclosure protects the listener at the delivery layer. It does not protect them at the content layer, and the reason is worth stating plainly: a fully disclosed, fully audible, carefully published single instruction is still an instruction. Somebody who reads it, is tired, is deferential, or trusts the operator, is being commanded with their eyes open.

The content-layer protection is structural and it is a theorem. It is proved in the framework's narrative-input paper [Curzi, 2026ad], whose Result One is *determination by choice*: a corpus closed under negation has resultant exactly zero, $\|\sum\| = 0.000000000000$ at $N = 4, 8, 12, 24$, and an orthogonal corpus's mean direction shrinks as $1/\sqrt{N}$. It is applied here to the twelve-state ladder rather than restated in general.

Model each directive as a unit vector in a real inner-product space, where direction encodes the action urged and $-v$ encodes its negation.

(a) If the corpus is closed under negation, the resultant is exactly zero. Verified, with the pairing drawn at random in a 64-dimensional space:

```
N =  4    ||sum|| = 0.000000000000
N =  8    ||sum|| = 0.000000000000
N = 12    ||sum|| = 0.000000000000
```

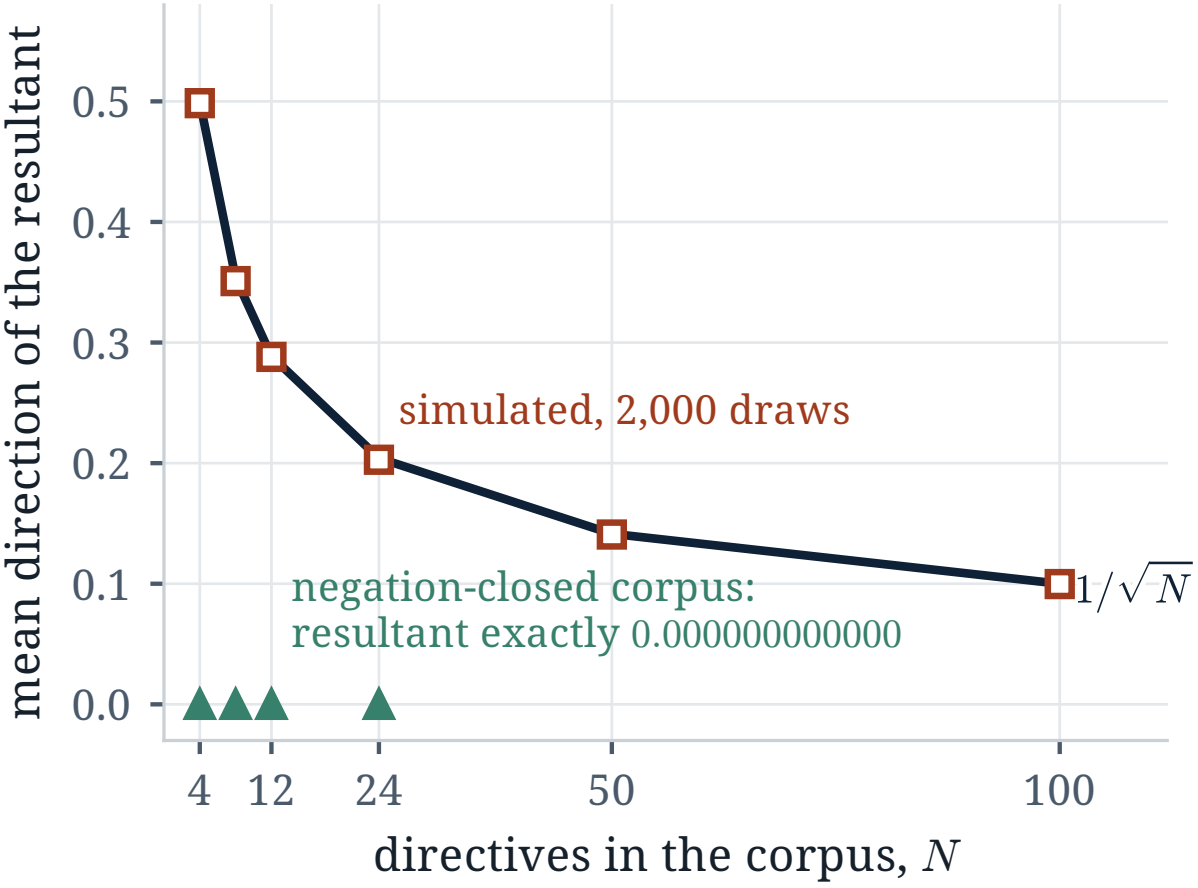


Figure 48.1: **A corpus built to span does not point: closed under negation its resultant is exactly zero, and merely mutually exclusive its mean direction falls as $1/\sqrt{N}$.** The curve is the theoretical $1/\sqrt{N}$; the open squares are 2,000 simulated draws of N isotropic unit vectors in $\mathbb{R}^{64}$; the triangles are the negation-closed case, exact at every N tested. Modelled: directives are represented as unit vectors, a modelling choice stated in §2.4, and the exponent $-1/2$ is defined by the model rather than estimated. Built by `figs/src/CCP-048_fig04.py`.

```
N = 24    ||sum|| = 0.000000000000
```

(b) If the corpus is mutually exclusive without exact pairing, the resultant grows only as $\sqrt{N}$ while the corpus grows as N, so the mean direction shrinks as $1/\sqrt{N}$. **The exponent $-1/2$ here is defined by the model — it is the standard concentration of a sum of independent isotropic unit vectors — and is not an estimated scaling exponent; no fitting procedure of the kind a claimed power law requires [Clauset et al., 2009] was run, and none would be appropriate.** Simulated, 2,000 draws of N isotropic unit vectors in $\mathbb{R}^{64}$ per row:

N	1/sqrt(N)	simulated
4	0.500000	0.498500
8	0.353553	0.350953
12	0.288675	0.288101
24	0.204124	0.202825
50	0.141421	0.140809
100	0.100000	0.099738

A field of mutually exclusive directions cannot issue a command. In case (a) the resultant is

identically zero; in case (b) it is a vanishing fraction of any single member and keeps vanishing as the corpus grows. The corpus does not point. It spans.

The consequence is the sentence this section is really about:

> **A deliberately contradictory corpus cannot issue a command. It can only be resolved by selection, and the selection belongs to the receiver.**

The vector model is a model and is badged as one (M-a, §2.4). Real directives are not unit vectors and "let the charge rise" is not literally the additive inverse of "let the charge fall." What the model captures rigorously is one thing only: a set constructed so that its elements cancel or span has no resultant, and a receiver who acts on it must therefore have selected. That conclusion depends only on the corpus genuinely containing mutually exclusive directions, which is a property of the written text and is checked by reading it — as §4.29 checks it, and fails.

4.6 The identity between what the corpus carries and what selection costs Two quantities coincide at the extreme, and the coincidence is the reason the design is not merely defensive.

By the framework's normalisation result, $F(u) = \ln N$ if and only if the ensemble is mutually orthogonal. A maximally contradictory corpus is therefore a maximum-surplus corpus. The ceiling $\ln N$ is the same bound that requisite-variety arguments place on any regulator [Ashby, 1956, 1958], and the logarithm is Shannon's functional form rather than a new one [Shannon, 1948]. At $N = 12$:

```
corpus surplus at u = 1        F = ln 12 = 2.484907 nats
selection information          ln 12 = 2.484907 nats = 3.584963 bits
f_rel = k (1 + 11 u)^2         69.6133 Hz          (folded 278.453 Hz, j = 2)
octaves above k                7.169925
```

The corpus supplies exactly as much surplus as the selection consumes, and supplies zero net direction. The direction is not in the signal; it is contributed by the receiver, and its magnitude is exactly the surplus the signal carries. One honesty clause attaches unchanged: **the exact equality holds only at $u = 1$ and under a uniform receiver prior**, and both sides fall together when either condition relaxes [Curzi, 2026ad].

The fall is worth seeing, because a corpus is rarely maximally contradictory in practice:

u	F (nats)	f_rel (Hz)	octaves above k
1.00	2.484907	69.6133	7.1699
0.90	2.388763	57.4358	6.8925
0.75	2.224624	41.3631	6.4189
0.50	1.871802	20.4247	5.4009
0.25	1.321756	6.7982	3.8138
0.10	0.741937	2.1319	2.1408
0.00	0.000000	0.4834	0.0000

At `u = 0` the corpus is parallel: every directive says the same thing, `F = 0`, and the frequency is the ground essence `k` itself. **A corpus that agrees with itself carries no surplus and sounds the note of no relationship at all** — which is the framework's result on the unison of perfect agreement, arriving in an unexpected place.

4.7 The `u = 1` here is constructed, and that must be said every time it is printed The framework's encoding rules carry a warning that is usually a trap and is here a design: one-hot vectors are automatically orthogonal, giving `u = 1` regardless of content, so framings must be represented with overlapping support or the internal surplus is degenerate and meaningless.

A deliberately contradictory corpus is exactly that degenerate case, on purpose.

The `u = 1` of a designed corpus is a property of the design, not a measurement of anything, and the 69.6133 Hz that follows from it is the arithmetic ceiling of a twelve-category system rather than a reading of a person, a text, or a relationship. It says the corpus was built to span. It says nothing about how anyone responds to it, and any sentence that lets those two readings blur is a sentence this framework has to delete.

The distinction matters because the same number, 69.6133 Hz, appears elsewhere in this corpus as the ceiling kN^2 and as the top of a glide [Curzi, 2026,?]. In those places it is the endpoint of a measured scale. Here it is a construction constant. Same value, different warrant, and the warrant is what a reader is entitled to.

4.8 The corpus at the ceiling, read with the framework's own instrument

Because the constructed corpus sits at $F = \ln N$, it lands exactly on the framework's sustainability ceiling, and that gives a clean statement of what even a maximally spanning corpus cannot do.

c/eta	f_min (Hz)	required octaves	deficit vs 69.6133	verdict
0.5000	1.3141	1.442695	−5.727230	SUSTAINS
1.0000	3.5721	2.885390	−4.284535	SUSTAINS
1.5000	9.7099	4.328085	−2.841840	SUSTAINS
2.0000	5.3942	5.770780	−1.399145	SUSTAINS
2.4849	69.6133	7.169925	+0.000000	HOLDS
2.6000	87.6317	7.502014	+0.332089	DECAYS
3.0000	195.0279	8.656170	+1.486245	DECAYS

At `c/eta = ln 12 = 2.484907` the deficit is zero to floating-point exactness — computed residual `8.88e−16`. **Above that ratio no arrangement of twelve categories is sustainable, and a maximum-surplus corpus is no exception: it is the best a twelve-state system can do and it is still not enough.** That is the ceiling result in the frequency register [Curzi, 2026], and it is one of the few places where a formalism states in advance that a proposed arrangement cannot work.

And the deficit alone does not see decay. Worked at `c/eta = 2.6`, with `eta = 0.50`, `c = 1.30`, `F = ln 12`:

```
delta = 0.09    Q_inf = −0.639408     D = +0.332089
delta = 0.05    Q_inf = −1.150934     D = +0.332089
delta = 0.02    Q_inf = −2.877334     D = +0.332089
direct check    D = 2.885390*(c/eta − F) = +0.332089
```

Trust moves the fixed point by a factor of four and the deficit does not move at all. Report `D` and `delta` together or the reading cannot see decay. The minimum-frequency paper owns this [Curzi, 2026], whose Result Four is that D contains no δ: on its carried dyad, δ from 0.09 to 0.05 leaves $D = +4.338420$ unchanged while Q_∞ falls from -5.0119 to -9.0215.

4.9 The open problem, carried

Section 6 owns the statement and states it four ways. It is carried here in full rather than by reference, because a reader arriving at material about content and cells is exactly the reader most likely to skip forward and exactly the reader who most needs it.

> **There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal.**

What follows from it, for this section: nothing in §4.10 onward is a claim about cells. The design work below is encoding design. It produces a notation with a checkable structural property, and a notation is not a channel.

4.10 Status table The classification below follows the framework's three-drawer honesty ledger [Curzi, 2026], whose Result One sorts four items of established mathematics, seven of established psychology and exactly three declared postulates, with every protocol stated to run on the first two and steer by the third.

Claim	Status	Basis
No subliminal content; full disclosure	**CONSTRAINT** — ethical, prior, non-negotiable	the framework's disclosure constraint
A hidden message cannot be consented to	**ARGUED** (not empirical)	§4.3
Disclosure need not abolish a belief-mediated effect	**MEASURED, elsewhere**	Kaptchuk 2010; von Wernsdorff 2021
A negation-closed corpus has resultant exactly zero	**PROVED (elementary)**, verified	§4.5
A mutually exclusive corpus has mean direction `1/sqrt(N)`	**PROVED**, simulated	§4.5
Corpus surplus `ln N` equals selection information `ln N`	**PROVED, extremal case only**	§4.6, with H82-A
`u = 1` for a designed corpus is constructed, not measured	**STATED**	§4.7
Corpus at `F = ln N` sits exactly on the ceiling	**COMPUTED**, residual `8.88e-16`	§4.8
Contradiction renders as exactly 2400.0 cents	**PROVED (arithmetic)**	§4.14
Twelve states yield six pitch classes	**PROVED**	the odd-part rule
Narrative structure is carried in the harmonic	**DESIGN CLAIM** — about encoding, not effect	§4.15
A cell demodulates sequence-specific content from this	**NOT CLAIMED. OPEN PROBLEM.**	§6
Any physiological, clinical or therapeutic effect	**NOT CLAIMED**	H4-1

4.11 Honesty clauses, binding on this section **H4-1.** No clinical, medical or therapeutic claim is made anywhere in this section. The twelve states are not diagnoses, the corpus is not a treatment, and the rendering is not an intervention in any clinical sense of that word.

H4-2. No claim that DNA is edited, altered or damaged. Nothing in the chain is a molecule. The nucleotide string of Section 5 is a string in the computer-science sense — symbols in memory — and no nuclease, template or Cas protein appears anywhere.

H4-3. No claim that any capability has been conferred on anything.

H4-4. The essence frequencies are essence frequencies, not physical resonances. No person, tissue, cell or molecule is claimed to resonate at 247.514 Hz, at 69.6133 Hz, or at any other value printed here.

H4-5. All worked material in this section is constructed and labelled as such. The twelve-line corpus of §4.28 was written for this paper. It is not a transcript, not a record of practice, and not an observation of anyone.

H4-6. The vector model of a directive is a model. Its conclusion depends only on the corpus containing genuinely mutually exclusive directions, which is checked by reading the text.

The engineering

4.12 What "narrative" means at this layer The word does a great deal of unearned work in adjacent literatures, so it is worth pinning to something checkable. Here, a *narrative* is **a finite, written, published set of directives, each of which names one state of the twelve-state ladder and urges one action with respect to it.**

That definition has three consequences and each is deliberate.

It is finite, so it can be printed in full, which is what disclosure requires. It is written, so the equal-force check of §4.29 can be applied by anyone with a character counter. And it is indexed to the ladder, so every directive has a frequency by the folding rule and the corpus has a chord.

What a narrative is *not*, at this layer: a story in the literary sense, a script, a suggestion, an affirmation, or a piece of persuasion. Persuasion has a resultant. This construction is built so that it does not.

4.13 The chain, and where each link is specified

```
narrative                  a written corpus of directives over the twelve states
    |
    |  [1]   ENCODE     text -> bits -> nucleotide string        §5.7
    v
sequence
    |
    |  [2]   ADDRESS    a gRNA-format routing header prepended   §5.10
    v                   (a field in a file; not a molecule)
addressed sequence
    |
    |  [3]   SONIFY     essence law, f = k n^2                   §5.12
    |  [4]   TRANSPOSE  fixed integer octave shift               §5.13
    v
waveform
    |
    |  [5]   DISCLOSE   corpus, tables, sequence, and audio       §4.19
    v                   published together, before exposure
a sound in a room, and §6's open problem begins here
```

Section 5 owns steps 1 through 4 and specifies them to the digit. This section owns step 5 and the structural design that makes step 5 sufficient.

4.14 The contradiction is rendered as a double octave Here is the design move, and it is exact arithmetic rather than a metaphor.

Every directive in the corpus names a state `n` and is rendered at that state's essence frequency, folded into the standing band of 220–440 Hz by the folding rule. Its explicit negation is rendered at the same pitch **transposed up by exactly two octaves**. Because octave transposition preserves pitch class exactly — a structure formalised for pitch [Shepard, 1982] whose perceptual specialness has a measured size [Demany et al., 2013] — the two members of a contradictory pair are the same note at two heights, and the interval between them is:

```
ratio  = 4          exactly
cents  = 1200 * log2(4) = 2400.0     exactly, in every pair, without exception
```

Two octaves is also, not coincidentally, the interval $f(2n)/f(n) = (2n)^2/n^2 = 4$ that the square law produces between an index and its double. That identity is fixed in the framework's conversion-key paper [Curzi, 2026l], whose Result Three is that **one bit equals two**

octaves, not one, verified across $N_{\text{eff}} = 1, 3, 12$. Nothing here changes `N`, adds a state, or steps off the ladder: the negation is placed by octave transposition, which is a change of register and the only purely presentational operation in the chain.

A word about where the twelve class names come from, and what they are not. The rungs are motivated by naming existing constructs in the affect literature: a two-dimensional circumplex of core affect underlies *valenced activation* [Russell, 1980]; appraisal dimensions underlie *attributed* and *situated feeling* [Smith and Ellsworth, 1985, Lazarus, 1991]; appraisal as a process in time underlies *temporal feeling* [Scherer, 2001]; self-conscious emotion underlies *reflexive feeling* [Tracy and Robins, 2004]; theory of mind underlies *social feeling* [Premack and Woodruff, 1978]; and action readiness underlies *committed feeling* [Frijda, 1986]. **Those works are named as precedent for the existence of a component and never as support for its index, its position on the ladder, or the ordering of the twelve.** Several of them cut directly against that ordering, and §3.5 states the objection in the general form it deserves rather than meeting it here piecemeal.

The full rendering table for the twelve states, band 220–440 Hz for the directive and 880–1760 Hz for its negation:

n	state class	raw f (Hz)	directive (Hz)	j	negation (Hz)	pitch class
1	undifferentiated activation	0.483426	247.5140	9	990.0559	1
2	valenced activation	1.933703	247.5140	7	990.0559	1
3	directed feeling	4.350832	278.4532	6	1113.8129	3
4	situated feeling	7.734812	247.5140	5	990.0559	1
5	attributed feeling	12.085643	386.7406	5	1546.9623	5
6	temporal feeling	17.403326	278.4532	4	1113.8129	3
7	reflexive feeling	23.687861	379.0058	4	1516.0231	7
8	social feeling	30.939247	247.5140	3	990.0559	1
9	normative feeling	39.157484	313.2599	3	1253.0395	9
10	narrative feeling	48.342573	386.7406	3	1546.9623	5
11	reciprocal feeling	58.494514	233.9781	2	935.9122	11
12	committed feeling	69.613306	278.4532	2	1113.8129	3

The listener does not need the encoding table to hear that the corpus is built in opposed pairs. The interval arithmetic is standard once an index exists [Rossing et al., 2002]. **A contradiction is here the most audible interval available, and it is audible without instruction, which is exactly the property a disclosure requirement should want from a structure.**

4.15 The pitch-class budget, and the collision it forces The odd-part rule and its corollary — an indexed set $1..N$ holds exactly $\lceil N/2 \rceil$ distinct pitch classes — are the framework's. At `N = 12` that is six, and this section has to live inside it:

```
class  1  ->  states 1, 2, 4, 8      folded 247.5140 Hz
class  3  ->  states 3, 6, 12        folded 278.4532 Hz
class  5  ->  states 5, 10           folded 386.7406 Hz
class  7  ->  state  7               folded 379.0058 Hz
class  9  ->  state  9               folded 313.2599 Hz
class 11  ->  state 11               folded 233.9781 Hz
```

Four of the twelve directives share one pitch class. Three share another. How far interval categories and octave equivalence generalise across listeners is contested [Burns, 1999, Hoeschele et al., 2012], and what a listener finds consonant depends on the spectrum of the sound rather than on the ratio alone [Sethares, 2005]; **this paper uses the arithmetic and claims none of the perception.** A listener attending to pitch class rather than to absolute register hears twelve directives as **six distinguishable positions**, and the ambiguity is a structural property of the twelve-state ladder rather than a fault in the rendering. The corresponding

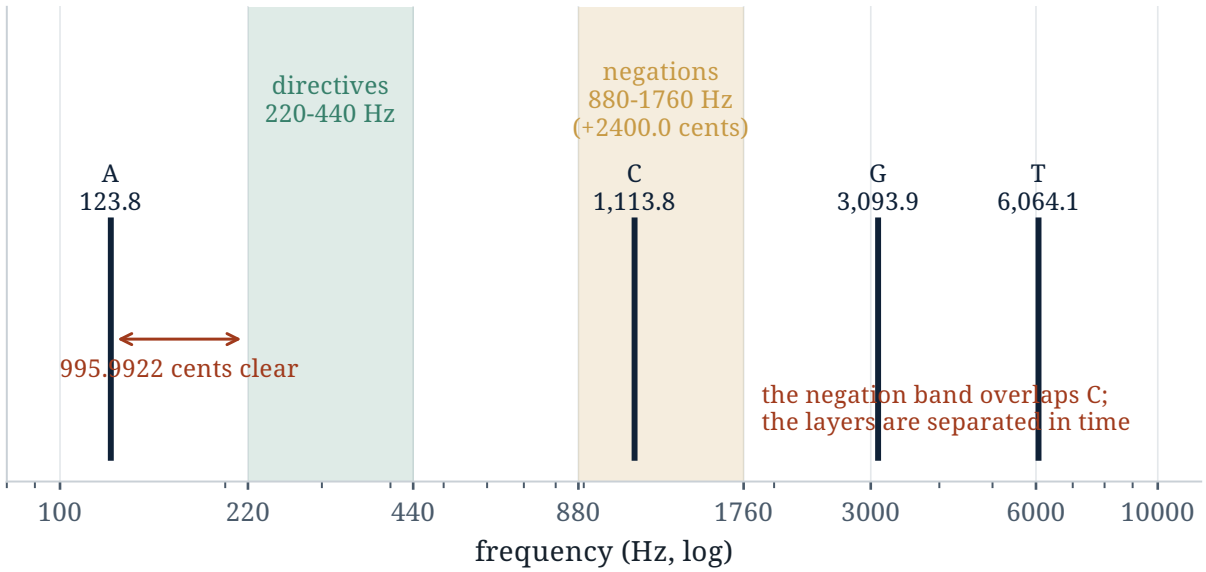


Figure 48.2: **The two disclosed layers occupy disjoint registers: the state band clears the nearest symbol tone by 995.9922 cents, better than an octave.** The negation register does overlap the symbol tone at 1113.8129 Hz, which is why the run schedule of §5.18 separates the two layers in time instead. Modelled from the essence law with the declared transpositions ($j = 8$ for the symbol alphabet, $j = 6$ for the state ladder); every value is arithmetic on a postulated constant, and nothing resonates. Built by `figs/src/CCP-048_fig02.py`.

resolution figure is fixed by the framework's reference apparatus [Curzi, 2026]: within the folded lattice the tightest separation at $N = 12$ is 34.9756 cents, between classes 7 and 5.

Two consequences follow and both belong in the disclosure package.

First, absolute register — not pitch class — is what identifies the state. The rendering must therefore be published with its j for every value, and the listener must be given the table.

Second, and this is the honest one: the collisions are not an artefact of an unlucky ladder. They follow from `(2n)^2 = 4n^2`, they would occur for any N, and the only way to buy more classes is to lengthen the index — which past $N = 16$ enters the 289/288 regime and costs a factor of eight in symbol rate. **This framework stays at $N = 12$, inside the 50/49 regime, and accepts six classes as the price.**

4.16 The structure is carried in the harmonic — stated precisely This is the claim the section is named for and it is exactly the kind of phrase that sounds like a mechanism while being nothing of the sort. So, in two halves:

What is claimed. The relational structure of the corpus — which directives contradict which — is realised as interval structure in the rendered chord, audibly, at ordinary listening levels, so that a listener hearing the chord is hearing the structure rather than a representation of it. The mapping is bijective: every contradiction is a double octave and every double octave in the rendering is a contradiction.

What is not claimed. That this structure does anything to any listener, any cell, or any transcript. The claim is a design claim of the same kind as "this file is sorted." It is checkable, it is either true or false of a given rendering, and it has no biological content whatever.

The pitch classes themselves are exact just intervals, because every essence ratio is a ratio of squares and k cancels out of all of them [Curzi, 2026k]. Exact rational interval lattices are a documented tradition rather than an invention of this framework [Partch, 1974, Doty, 2002, Barbour, 1951], and the project of assigning arithmetic weights to intervals runs from Euler's

gradus suavitatis [Euler, 1739] to harmonic distance [Tenney, 2015]. The point here is that the intervals a listener hears between directive families are exact rational relationships, not approximations, and they are exact because nothing in the chain ever tempered them.

4.17 Three architectures, compared

	Steganographic	Masked	Carried (this framework)
Where the content is	in a cover signal	below perceptual threshold	it is the signal
What the listener hears	the cover	the masker	the content
Can they recover it?	only with a key	not at all by ear	yes, by ear plus tables
Can they evaluate it?	no	no	yes, before exposure
Can they decline it?	no	no	yes, at any point
Consent status	impossible	impossible	obtainable
House status	**excluded**	**excluded**	permitted

The middle column is where every historical subliminal technique sits, and this framework excludes it for the same reason it excludes the first: **in both, the recipient's inability to perceive the content is doing the work.**

The right-hand column is not a weakened version of the other two. It is a different architecture with a different worst case, and its worst case is that the listener reads the corpus and declines. That is a safety property and not a weakness — the same argument the framework's closure law makes about tuning a relationship without removing anyone's ability to walk away [Curzi, 2026].

4.18 Why disclosure and contradiction are both required They fail in different circumstances, which is the whole reason for having both.

A single-directive corpus satisfies every clause of §4.1. It is audible, published, readable, evaluable and refusable. And it is still a command, because a corpus with one member has a resultant equal to its member. Its surplus is $F = \ln 1 = 0$: a corpus of one is a parallel corpus and carries no surplus at all.

Conversely, a contradictory corpus delivered inaudibly satisfies §4.5 and violates §4.1, and the violation is total — a listener who cannot hear the corpus cannot perform the selection the structure exists to leave them.

Disclosure makes a command refusable. Contradiction makes it not a command. Disclosure fails when the listener is tired, deferential, or trusting; a field of mutually exclusive directions has no resultant whether or not anyone is paying attention.

4.19 The disclosure package Disclosure is an artefact, not an attitude, and the artefact is specified so that its absence is detectable. Everything below is published *before* any exposure, in a form the listener keeps:

```
D1   the corpus, in full, both members of every pair, in the rendering order
D2   the encoding tables: character set, byte width, bit-to-base rule
D3   the nucleotide string actually used, in full
D4   the address prefixes, in full, with their derivation
D5   the symbol-to-pitch table with n, f, j, and the resulting Hz
D6   the state rendering table of §4.14, with j for every value
D7   the envelope, symbol duration, and total run length
```

```
D8   the audio file itself, in an unencrypted, uncompressed format
D9   this section's honesty clauses and the open problem of section 6, in full
D10  a stated procedure for stopping, and who to tell
```

Item D3 is not redundant with D2. Section 5 uses a one-bit-per-base encoding in which the rule does not determine the string, so **a scheme whose rule is one-to-many must publish the string and not merely the rule.** The weaker encoding forces the stronger disclosure, and the trade runs in the right direction.

Item D10 exists because a stop procedure that has not been written down is a stop procedure that does not exist at the moment it is needed.

The specification

Numbered, checkable, and binding on Section 5. A rendering that fails any of these is out of scope.

S4.1 Every component of the rendered signal shall be audible to a person of ordinary hearing at the listening position — audible as rendered, at the level rendered, in the room, over the noise floor actually present. An inaudible component is raised or removed. There is no third option.

S4.2 No carrier, no mask, no embedding. No content under, inside or behind another signal. No amplitude-modulated speech, no time-compressed content, no spectral hiding beneath a louder band, no content below 20 Hz or above 20 kHz intended to act on a listener, and no psychoacoustic masking of any kind.

S4.3 The symbol rate shall be set by audibility, not by the Gabor floor. Section 5.15 computes both; the binding number is the one that leaves symbols individually resolvable by an ordinary listener in an ordinary room.

S4.4 The full disclosure package D1–D10 shall be delivered before exposure, in a durable form the listener retains.

S4.5 For every directive `v` in the corpus, its explicit negation `–v` shall also be in the corpus, stated with equal force, equal length, and equal position in the rendering order.

S4.6 No directive shall appear without its opposite. No opposite shall be weakened, hedged, qualified, or given the last word.

S4.7 The rendering order shall be published and shall be stated not to be a ranking.

S4.8 The equal-force check of §4.29 shall be computed and published, including its failures. A corpus with unrepaired failures may be rendered only if the failures are disclosed in D1.

S4.9 Every published frequency shall carry its transposition `j` and its band.

S4.10 No tempering. The essence ratios are exact ratios of squares and shall not be adjusted to twelve-tone equal temperament. 34.9756 cents is not 0 cents.

S4.11 The rendering shall use a windowed envelope. A rectangular gate is disqualified: it has no finite spectral second moment, verified numerically divergent in §5.17.

S4.12 The `u` of a designed corpus shall be reported as constructed, never as measured, in the same sentence that gives its value.

S4.13 Any octave-deficit reading published alongside a corpus shall report `D` and `delta` together.

S4.14 The listener shall be able to stop the rendering themselves, at any point, without asking anyone. A stop that requires an operator is not a stop.

S4.15 No rendering shall be delivered to a person who has not read D1 and D9, or who has not had them read to them.

S4.16 No rendering shall be delivered during sleep, sedation, or any state in which the listener cannot exercise S4.14.

S4.17 The corpus shall contain no directive that instructs the listener to conceal the corpus, to continue past a stop, to disregard their own judgment, or to defer to the operator.

S4.18 Nothing in the corpus shall reference a clinical condition, a diagnosis, a medication, or a course of treatment.

S4.19 The honesty clauses H4-1 through H4-6 and §6's open problem shall accompany every rendering as delivered, not as a footnote available on request.

S4.20 Every worked corpus published anywhere shall be labelled CONSTRUCTED in the same block that prints it.

The failure modes

FM-4.1 — The reasonable exception. Somebody proposes a masked version for a control condition, a pilot, or a demonstration to an audience that "would be distracted by hearing it." The proposal is always locally reasonable and it is always the same failure. **A constraint that admits a research exception is a constraint that has been repealed and re-announced as a preference.** Section 6 specifies a scrambled-sequence control precisely so that no masked condition is ever needed.

Fix: refuse. The controls of §6.13 do the work a masked arm was supposed to do.

FM-4.2 — Disclosure as a document nobody reads. A forty-page appendix technically satisfies S4.4 and defeats it. Disclosure is not discharged by publication; it is discharged by comprehension.

Fix: D1 is the corpus itself, in full, on one page, in the rendering order. Everything else is supporting material. If the corpus does not fit on a page, the corpus is too long.

FM-4.3 — The weighted corpus. Both members of each pair are present, so §4.5 appears satisfied, but one member is longer, louder, first, or better written. **A weighted contradictory corpus is a command with an alibi, and the alibi is more dangerous than an honest command because it survives inspection.**

Fix: the equal-force check of §4.29, computed and published including failures.

FM-4.4 — The private reduction. The corpus is balanced as written and the operator quietly renders only the half they agree with, or a listener is coached toward one side in the framing conversation. The structure is intact on paper and destroyed in delivery.

Fix: S4.7, plus a published rendering order fixed before the session, plus a witness who is not the operator.

FM-4.5 — Treating `u = 1` as a finding. The corpus reads at 69.6133 Hz, somebody prints that next to a person's name, and a construction constant has become a measurement.

Fix: S4.12. The value is a property of the design and says nothing about anyone.

FM-4.6 — Importing the placebo literature as support. The open-label result is used to suggest that this method works, or that a disclosed rendering is "like an open-label placebo."

Fix: §4.4. The result establishes that disclosure and effect are compatible in another domain. It is evidence about placebo trials, not about sonified sequence, and the two-column rule forbids moving it across [Curzi, 2026w].

FM-4.7 — Precedent presented as evidence. Several textual traditions preserved contradictory directives deliberately: Proverbs 26:4–5 places two opposed instructions adjacent and preserved; the Talmudic practice of *machloket* keeps minority opinions beside the ruling; the koan tradition constructs prompts that cannot be resolved within their own terms. That is genuine precedent and it is worth having. **None of these sources is in this paper's bibliographic database and none is cited with a key.**

Fix: state the column's role in the sentence. **Independent arrival is an argument that a structure is in the territory rather than only in this framework's instrument; it is not evidence that the structure does anything** [Curzi, 2026w].

FM-4.8 — Mechanism creep. “Carried in the harmonic” starts as a statement about encoding and becomes, three paragraphs later, a statement about transmission.

Fix: §4.16’s two halves, kept adjacent. Any sentence containing “carried” that also contains a cell, a receptor, or a transcript is a sentence to delete.

FM-4.9 — The pitch-class collision read as meaning. States 1, 2, 4 and 8 fold to the same 247.5140 Hz, and somebody concludes that these four states are “the same underneath.”

Fix: §4.15. The collision is $(2n)^2 = 4n^2$ and nothing else.

FM-4.10 — Rendering a single directive “for clarity.” The corpus is reduced to its most sensible member so the listener is not confused.

Fix: §4.18. A corpus of one has `F = 0` and a resultant equal to its member. It is a command, and the reduction converted a structure into an instruction.

FM-4.11 — Consent obtained after the disclosure was pre-summarised by the operator. The operator explains what the corpus “is really about,” and the listener consents to the summary.

Fix: S4.15 requires D1, the corpus itself. A summary is the operator’s resultant, and the whole design exists to have no resultant.

FM-4.12 — The deficit reported without `delta`. A corpus reading is published with `D` alone, and the reader cannot tell decay from level.

Fix: S4.13, and §4.8’s worked demonstration that `D` does not move when `Q_inf` changes by a factor of four.

FM-4.13 — Calling any of this a practice. The section’s material is a specification for a research procedure. The moment it is described as a practice, a modality, a protocol in the clinical sense, or a technique, it has acquired a standing it does not have.

Fix: H4-1, and §5.2’s labelling rule, which applies to every sentence and not only to the abstract.

Worked example

CONSTRUCTED. The corpus below was written for this section. It is not a transcript, not a record of practice, and not an observation of any person.

4.28 The corpus Twelve directives, one per state of the ladder, each with its explicit negation. Rendering order is by index and is published as such; it is not a ranking.

```
n = 1    LET THE CHARGE RISE UNNAMED    /  LET THE CHARGE FALL UNNAMED
n = 2    CALL THIS GOOD AND ACT         /  CALL THIS BAD AND ACT
n = 3    AIM THE FEELING AT ITS OBJECT  /  TURN THE FEELING FROM ITS OBJECT
n = 4    READ THE ROOM BEFORE YOU MOVE  /  MOVE BEFORE YOU READ THE ROOM
n = 5    NAME THE CAUSE AND HOLD IT     /  DROP THE CAUSE AND HOLD NOTHING
n = 6    GIVE THIS AN END AND WAIT      /  GIVE THIS NO END AND WAIT
n = 7    TURN AND LOOK AT YOURSELF      /  TURN AND LOOK AWAY FROM YOURSELF
n = 8    MODEL THE OTHER MIND NOW       /  LEAVE THE OTHER MIND ALONE
n = 9    APPLY THE STANDARD IN FULL     /  SUSPEND THE STANDARD IN FULL
n = 10   PUT THIS INSIDE THE STORY      /  TAKE THIS OUTSIDE THE STORY
n = 11   LET THEM SEE YOU SEEING THEM   /  LET THEM SEE NOTHING OF YOU
n = 12   BIND YOURSELF TO THE NEXT ACT  /  LOOSE YOURSELF FROM THE NEXT ACT
```

Each pair names one state and urges opposite dispositions toward it. Nothing in the corpus refers to a condition, a treatment, or an outcome, per S4.18.

4.29 The equal-force check, computed and failing S4.5 requires equal force, and "equal force" is only a specification if it is checkable. The declared check is character count and word count, with a tolerance of two characters and zero words.

n	chars	d	words	d	result
1	27 / 27	0	5 / 5	0	PASS
2	22 / 21	1	5 / 5	0	PASS
3	29 / 32	3	6 / 6	0	FAIL
4	29 / 29	0	6 / 6	0	PASS
5	26 / 31	5	6 / 6	0	FAIL
6	25 / 25	0	6 / 6	0	PASS
7	25 / 32	7	5 / 6	1	FAIL
8	24 / 26	2	5 / 5	0	PASS
9	26 / 28	2	5 / 5	0	PASS
10	25 / 27	2	5 / 5	0	PASS
11	28 / 27	1	6 / 6	0	PASS
12	29 / 32	3	6 / 6	0	FAIL

Four pairs out of twelve fail — 33.33% — and they are reported rather than quietly repaired, because the rate at which a careful author fails this check is the useful number. The author wrote these lines intending them to be balanced and they are not. Pair 7 is the worst: the negation is seven characters and one word longer than the directive, which means that in a rendering where each character contributes to the symbol stream, the negation is materially longer on the air.

The three ways to respond, in decreasing order of honesty:

1. **Rewrite and re-check.** Preferred. Repair pairs 3, 5, 7 and 12 and publish the second check alongside the first, so a reader can see what changed.
2. **Render as is and disclose.** Permitted by S4.8 provided the failing pairs are named in D1.
3. **Render as is and say nothing.** Forbidden. This is FM-4.3 with the evidence available and suppressed.

Character count is a crude proxy for force and is not defended as more than that. It catches gross asymmetry, which is the failure that matters, and it is checkable by anyone. Prosody, connotation and rhetorical weight are not caught by it, which is why S4.7 also fixes the order and why a witness who is not the author is worth having.

4.30 The rendering The corpus renders as twenty-four tones. Each directive sits at its state's folded frequency in 220–440 Hz; each negation sits exactly two octaves above it in 880–1760 Hz.

```
band, directives   220.0000 – 440.0000 Hz
band, negations    880.0000 – 1760.0000 Hz
every pair         ratio 4 exactly = 2400.0 cents exactly
distinct classes   6, per ceil(12/2)
tightest folded gap at N = 12   34.9756 cents (classes 7 and 5)
tempering          none
envelope           windowed; rectangular gate disqualified (S4.11)
```

The audible result is twelve pairs of double octaves, six of which share pitch class with another pair, spread across two registers. A listener who has never seen the tables can hear that the material is built in opposed pairs, because the pairs are octaves and octaves are the one interval nobody has to be taught.

4.31 The reading, with every warrant attached

```
u                 1.000000     CONSTRUCTED, not measured (S4.12)
F                 2.484907 nats = ln 12, the ceiling
f_rel            69.613306 Hz   folded 278.4532 Hz (j = 2)
octaves above k   7.169925
selection cost    2.484907 nats = 3.584963 bits
```

The corpus supplies 2.484907 nats of surplus and selecting one of its twenty-four members — twelve states, two directions each — costs the receiver exactly the surplus the corpus carries, under H82-A's two conditions. **What the number says is that the corpus was built to span the whole ladder and to point nowhere. It says nothing about any listener, and no reading of any listener has been taken.**

Reported honestly: this is `f_rel` computed from a constructed `u = 1`, at `N = 12`, with `k = 0.483425732988414 Hz`, folded at `j = 2` into 220–440 Hz. It is an essence frequency, not a physical resonance. It is not clinical, not diagnostic, and not a measurement of a person.

4.32 What this rendering does not do Everything above is encoding, and the boundary is worth drawing one final time in the concrete.

Twenty-four tones exist. They stand in exact rational relationships. Their structure mirrors the structure of a written text bijectively. A listener can hear the structure, read the text, check the tables, and stop the playback. **And there is no established mechanism by which any of that reaches a cell, a transcript, or a locus, which is §6's subject and the reason this paper is written in the order it is.**

What this section established **Established, and checkable:** that the no-subliminals constraint is ethical rather than stylistic, and rests on an argument — a hidden message cannot be consented to — that requires no empirical support and gains none from any; that open-label placebo removes the excuse for concealment without supplying evidence for anything in this paper; that a negation-closed corpus has resultant exactly zero and a mutually exclusive one has mean direction `1/sqrt(N)`, both verified; that at `N = 12` a maximally contradictory corpus carries `F = ln 12 = 2.484907` nats and sits exactly on the sustainability ceiling, residual `8.88e-16`; that a contradiction renders as exactly 2400.0 cents; that twelve states yield six pitch classes and the collisions are arithmetic rather than meaning; and that a careful author failed the equal-force check on four pairs out of twelve.

Not established, and not claimed: that any of this affects any cell, any transcript, or any person; that the twelve states are diagnoses; that a rendering is a treatment, a therapy, a practice or a modality; that DNA is edited, altered or damaged anywhere in the chain; that any capability has been conferred; or that any frequency printed here is a physical resonance.

The one line to carry forward: *the information is the tone, and a corpus built to contradict itself hands the decision back to the person listening.*

Section 4 owns the no-subliminals constraint as it applies to psychological content, the determination-by-choice argument in this framework's register, the double-octave rendering of a contradiction, and the disclosure package D1–D10. Section 5 builds the procedure on top of it; Section 6 states the open problem that governs both.

5 The encoding and rendering protocol, specified end to end

This section specifies a procedure end to end, in enough detail that a competent engineer can build it from this text alone and obtain bit-identical results. It is the most concretely

buildable material in this paper and the material that has to be most careful about what it claims, because it sits between Section 4, which fixed the ethical architecture, and Section 6, which states the gap that governs both.

Everything specified here is a research procedure. It is not a treatment, a therapy, a technique with a demonstrated effect, a modality, or a practice. That label is not an abstract disclaimer placed at the front and then forgotten; it is repeated at every stage of the specification, it governs the wording of every step, and §6's open problem is why.

What is built: an encoder, an addressing convention, a sonifier, a transposition, and a rendering discipline. All five are deterministic and reproducible, and the worked example demonstrates a round-trip SHA-256 identity on a real payload. What is not built, and what is not claimed in any form, is a biological effect. Section 6 is entirely about why that claim is unavailable.

The section also carries the falsification apparatus, because a procedure whose honesty clauses are appended is a procedure whose honesty clauses will be detached. **Every clause is written into the step it constrains.**

The law

5.1 The open problem, stated first and not by reference

> **There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal.**

That statement is §6's to develop. It is printed here in full rather than cited, because a reader who arrives at a section titled *protocol* is exactly the reader who will not turn forward, and exactly the reader for whom the sentence was written.

What is established upstream is a great deal, and §6.5 itemises it: data storage in nucleotide sequence, gRNA-format addressing, self-processed multi-address arrays, the separability of targeting from editing, mechanotransduction, force-to-transcription coupling at the nucleus, and the activation of latent capability by regulatory signal. **What is missing is the single link that carries sequence over force, and everything in this section stops at that link.**

5.2 The labelling rule Because the temptation to describe this procedure in clinical vocabulary is constant, the rule is made explicit and is binding on every sentence, table caption, figure label and file name produced under this section:

PERMITTED	procedure, specification, encoder, rendering, run, condition, exposure, participant, endpoint, preregistration, null result
FORBIDDEN	treatment, therapy, dose, patient, cure, healing, intervention (in the clinical sense), modality, practice, session (clinical), efficacy, indication, response rate, remission

A procedure that has to borrow clinical words to describe itself is claiming a standing it has not earned. The borrowing usually happens in a caption rather than in a sentence anyone argues about, which is why the rule is mechanical rather than a matter of judgment: the forbidden list is searched for before release, in every file, and hits are fixed rather than defended.

There is a second reason for the strictness and it concerns the reader rather than the writer. Clinical vocabulary carries an implied evidentiary context — a thing with a dose has been dosed, a thing with a response rate has been measured, a thing with an indication has been through something that produced one. Borrowing the word imports the whole implication

for free, and a reader who trusts the word has been misled by a sentence that never made a false statement. The vocabulary is where an overclaim happens without anybody deciding to overclaim, which is why it is policed by search rather than by conscience.

5.3 The constraint, inherited whole from Section 4

No subliminal content of any kind. Nothing masked. Nothing below threshold. Nothing concealed. The information is the tone. Section 4 owns the argument, the disclosure package D1–D10, and specifications S4.1 through S4.20. This section inherits all of them and adds none that weaken any.

Two of them bind directly on engineering decisions made below and are flagged where they bite: S4.3, which sets the symbol rate by audibility rather than by the Gabor floor, and S4.11, which disqualifies the rectangular gate.

5.4 What is built and what is postulated

The distinction is drawn before the interesting material, because the temptation to blur it increases as the section goes on.

Element	Status	Warrant
Text maps losslessly to a nucleotide string	**BUILT, VERIFIED**	§5.9, SHA-256 round trip
The balanced free-choice rule holds GC within one base of `L/2`	**PROVED (arithmetic)**, measured	§5.9
The balanced rule raises symbol entropy from 1.244171 to 1.935874 bits/base	**MEASURED** on the payload	§5.9
RLL(3) capacity is 1.982354 bits/base	**COMPUTED** (transfer-matrix eigenvalue)	§5.8
A 20-nt prefix is a sufficient address	**COMPUTED**, 0.005638867 expected hits	§5.10
The prefix is a routing header, not a molecule	**STATED, structural**	§5.10
Route B index assignment `n = 1,3,5,7` gives four distinct classes	**PROVED (arithmetic)**	§5.12
The naive assignment `n = 1,2,3,4` collapses three symbols	**PROVED, REJECTED IN PRINT**	§5.12
Symbol alphabet minimum gap is 990.055901 Hz	**COMPUTED**	§5.13
Symbol layer and state layer are separated by 995.9922 cents	**COMPUTED**	§5.14
Gabor limit `dt·df >= 1/(4 pi) = 0.079577` constrains the rate	**PROVED (classical Fourier)**	§5.15
Single-transposition audible capacity is `N <= 25`, at `j = 6` only	**COMPUTED**, verified	§5.16
The folded-lattice cliff is at `N = 17`; second at `N = 41`; rate factor 8	**CITED** [Curzi, 2026]	§5.16
The rectangular gate has no finite spectral second moment	**COMPUTED**, divergent	§5.17

Element	Status	Warrant
`k` is a chosen reference constant	**POSTULATE**, badged	§2.3 P-a
A transposed tone bears a physical relation to a base	**REFUSED**	H5-4
That any of this affects any cell	**NOT CLAIMED**	§6

Eleven computed or measured rows, two proofs, one rejection in print, one badged postulate, one refusal, and one open problem. A reader is entitled to see the accounting in that form before reading a word of the argument, and to check it against the section as they go.

5.5 Honesty clauses, binding **H5-1 — No clinical, medical or therapeutic claim.** This section describes signal processing and an experimental design. No trial, no controlled observation, and no demonstrated physiological effect of the procedure exists, and none is asserted.

H5-2 — No claim that DNA is edited, altered or damaged. Nothing in the pipeline is a molecule. The output of the last stage is a waveform; the nucleotide string is a string in the computer-science sense. **At no point does this procedure synthesise, deliver, or modify any nucleic acid, and it contains no nuclease, no template, and no Cas protein.**

H5-3 — No claim that any capability has been conferred on anything.

H5-4 — The transposition is a mapping, not a resonance. This is the strongest clause here and the one a careless reader most often violates. The tones of §5.13 are essence frequencies obtained by applying a postulated square law to a chosen index assignment and then shifting register. **A tone at 1113.8129 Hz has no physical relationship to cytosine. Cytosine does not absorb at that frequency, does not resonate at it, and is not affected by it.** The mapping is a labelling convention, chosen because it is principled and reproducible. It is not a physical coupling and this framework never says or implies that it is.

H5-5 — `k` is chosen, not measured. `k = (9/8)^2/phi^2 = 0.483425732988414 Hz` is a Pythagorean ditone divided by the golden section and is a design decision by the author. It is badged every time it appears. Result One of the essence law explains why the badge costs so little [Curzi, 2026k]: $f(m)/f(n) = m^2/n^2$, so `k` cancels out of every ratio in the system.

H5-6 — Every number here was computed, and the computations are stated in enough detail to be redone. Anything that could not be computed is marked nominal or omitted.

H5-7 — The worked payload is CONSTRUCTED. It is the corpus of §4.28, written for this paper. It is not a transcript and not an observation of anyone.

The engineering

5.6 The pipeline

```
a written corpus of directives over the twelve states        §4.28
    |
    |  [1]  ENCODE     text -> bits -> nucleotide string          §5.7 – §5.9
    v
payload sequence
    |
    |  [2]  ADDRESS    gRNA-format prefix prepended               §5.10 – §5.11
    v                  (a routing header on a data structure)
addressed sequence
    |
```

```
    |  [3]  SONIFY     essence law, f = k n^2, Route B           §5.12
    |  [4]  TRANSPOSE  fixed integer octave shift per symbol     §5.13
    |  [5]  LAYER      the state layer of section 4.14            §5.14
    |  [6]  RENDER     windowed envelope, rate set by audibility  §5.15 – §5.17
    v
a disclosed audio file, published with D1–D10
    |
    |  [7]  TRANSMIT   playback in a room, with consent
    v
the open problem of section 6 begins here, and this section stops
```

Steps 1 through 6 are engineering and are demonstrated below with computed and verified numbers. Step 7 is playback of an audio file. **What happens after step 7 is not claimed by this framework in any form.**

5.7 Stage 1 — text to bits The corpus is written as plain text and the encoding to bits is ordinary. It must nonetheless be published in full: character set, byte width, byte order, whether directives are concatenated or kept separate, and what separator is used.

The declared scheme is 8-bit ASCII, one byte per character, no terminator, directives separated by a single newline, both members of each pair present in the published rendering order.

For the corpus of §4.28:

```
characters, including newline separators      675
bytes                                         675
bits                                        5,400
SHA-256 of the source text
    c7d98b1cfde5c515ce99ff7db77d484f71a09c5cd921994f15df66c50373b72d
```

The bit count is Shannon information [Shannon, 1948, Cover and Thomas, 2006] — not thermodynamic entropy, and no thermodynamic claim is made anywhere in this paper. Nothing interesting happens at this stage, and that is deliberate: **any cleverness here would be a place to hide something.** A published, standard, character-for-character mapping is auditable by anyone with a text editor, which is what D2 of the disclosure package is for.

5.8 The rate landscape, and an honest correction carried forward Church, Gao and Kosuri (*Science* 337:1628, 2012) encoded a book into synthesised oligonucleotides and recovered it by sequencing. **That paper, and every other primary biological source named in this section and in §6, is absent from this paper's bibliographic database and is named without a citation key; see §10, L-9.** A correction must be carried: **they did not use two bits per base. They used one** — `A` or `C` encoded a zero, `G` or `T` encoded a one — and the free choice within each pair was spent on homopolymer avoidance and GC control. The base-4 alphabet is the theoretical ceiling; the Church architecture of block, address and common flanks is the part that was demonstrated. This procedure uses the alphabet and the architecture and says plainly that the doubling of rate relative to the 2012 paper is not something the 2012 paper did.

The homopolymer-constrained capacity is the capacity of a constrained channel and is a transfer-matrix eigenvalue [Shannon, 1948, Cover and Thomas, 2006], computed rather than quoted:

```
unconstrained base-4                    2.000000 bits/base   100.000%
RLL(2)  max run 2                       1.922688 bits/base    96.134%
RLL(3)  max run 3                       1.982354 bits/base    99.118%
RLL(4)  max run 4                       1.995717 bits/base    99.786%
RLL(5)  max run 5                       1.998939 bits/base    99.947%
this procedure, one bit per base        1.000000 bits/base    50.000%
```

Forbidding runs longer than three costs 0.882% of capacity. Discarding the second bit costs half of it, and buys something the compressed scheme cannot have.

5.9 Stage 1b — bits to bases, and what the discarded bit buys

With two acceptable bases per bit there is one free choice at every position. For a 5,400-bit payload the number of distinct nucleotide strings encoding exactly the same message is `2^5400`, a 1,626-digit count (`5400 · log10 2 = 1625.562`).

The declared rule for spending the free choice is **balanced selection**: at each position choose, among the bases legal under the homopolymer constraint, the one used least often so far, ties broken alphabetically. The obvious alternative is greedy first-legal selection. The two were built and measured on the same payload, and the difference is not small:

rule	A	C	G	T	GC	max run	entropy (bits/base)
greedy	3233	266	1856	45	0.392963	3	1.244171
balanced	1750	1749	951	950	0.500000	1	1.935874

Both recover the source exactly:

SHA-256, source text	c7d98b1c...0373b72d	
SHA-256, recovered from greedy	c7d98b1c...0373b72d	IDENTICAL
SHA-256, recovered from balanced	c7d98b1c...0373b72d	IDENTICAL

Two properties of the balanced rule are worth stating because they are structural rather than lucky. Because the rule always increments the currently smaller count, the two members of each bit class stay within one of each other; each class contains exactly one strong base and one weak one; therefore **GC content differs from `L/2` by at most one base regardless of what the payload says.** Here it is exactly 2,700 of 5,400. And because consecutive symbols from the same bit class alternate while consecutive symbols from different classes are drawn from disjoint letter pairs, the measured maximum homopolymer run is 1 — comfortably inside RLL(3).

The entropy figure is the one that matters for rendering. A greedy string is 3,233 adenines out of 5,400 and sounds like a drone with occasional events; a balanced string uses all four symbols nearly equally. A payload whose symbol distribution is that skewed is a payload whose rendering a listener cannot follow, which is **an audibility failure under S4.1 arriving from an unexpected direction.**

The first sixty bases of the balanced encoding, published because D3 requires the string and not merely the rule:

```
AGCATGCACTACAGCTAGCTAGCACATCACACAGCTAGCACTACGACACTACAGCTACGA
```

A scheme whose rule is one-to-many must publish the string. The weaker encoding forces the stronger disclosure, and that is the right direction for the trade to run.

5.10 Stage 2 — the address, which is a routing header and not a molecule

The framework's addressing paper established this architecture [Curzi, 2026y] and this section uses it without extending it. The prefix is formatted after the guide-RNA spacer convention, and its status must be stated in the same breath as its format:

The prefix is a field in a file. It is not delivered, it is not synthesised, it is not a molecule, no Cas protein is present anywhere in this procedure, and therefore nothing is edited.

Why twenty nucleotides, computed rather than asserted. The *Streptococcus pyogenes* Cas9 spacer is 20 nt and the arithmetic reason is address-space sizing against genome size:

```
L        4^L                expected hits in 6.2e9 positions
 8       6.553600e+04              94604.5
12       1.677722e+07                369.549
16       4.294967e+09                  1.44355
18       6.871948e+10                  0.0902219
20       1.099512e+12                  0.005638867
22       1.759219e+13                  0.000352429
24       2.814750e+14                  0.0000220268

4^20 = 1,099,511,627,776
address space / genome positions = 177.3406
```

Below 16 nt an address is not an address — it names hundreds of places. At 20 nt the expected number of spurious matches is about one in 177. That is the whole justification; it is arithmetic, and it is the same arithmetic that sets hash lengths in ordinary computing.

It is worth being precise about what "addressing" does and does not mean in a procedure where nothing is delivered. In ordinary CRISPR practice a spacer is an address because a molecule carries it to a locus and binds there; the address does work because the delivery vehicle exists. Here there is no vehicle. The prefix is a field in a file, and its function within this procedure is bookkeeping: it says which of the twelve directives a given block of payload belongs to, it makes the transmission parseable by anyone reading the published string, and it fixes the format so that the disclosure of D4 is checkable against a standard. The 20-nt length is justified by the same arithmetic that would justify it in a molecular context, and the justification is retained because borrowing a format and then silently changing its parameters is how formats become meaningless. But the length is doing bookkeeping work here, not molecular work, and no reader should be left able to think otherwise.

Cas12a's native processing of its own crRNA array (Zetsche et al., *Cell* 163:759, 2015; Fonfara et al., *Nature* 532:517, 2016) is cited here for one thing only: that a **multicast destination list** is a real biological format. It is cited for the format and for nothing else, and no claim is made that any such parsing occurs in response to sound.

5.11 The addressed payload

One address per directive pair, twelve pairs:

```
address prefixes         12 x 20 nt  =    240 bases
payload                                 5,400 bases
TOTAL TRANSMISSION                      5,640 bases
```

The prefixes are published in full in D4 with their derivation, exactly as the payload is published in D3. **An address that is not disclosed is a routing decision the listener cannot inspect, which fails S4.4 whether or not it does anything.**

5.12 Stage 3 — sonify by the essence law

Four symbols must become four distinguishable tones, by the law this framework already owns:

```
f(n) = k * n^2        k = (9/8)^2 / phi^2 = 0.483425732988414 Hz     POSTULATE, badged
```

Applying it requires a defensible ordinal on the four bases, which the framework demands of every application of the law. **An arbitrary ordering yields arbitrary frequencies, and the ordering must be argued in print or the route is unavailable.** The ordering used is alphabetical — A, C, G, T — which is not defended as meaningful, only as published, stable, and conventional in every sequence database in existence. That is a weaker warrant than the twelve-state ladder carries, and it is stated as weaker.

The obvious index assignment is a trap, and it is rejected in print. Assigning `n = 1, 2, 3, 4` gives odd parts 1, 1, 3, 1: three of the four bases collapse onto a single pitch class, because

`f(2) = 4 f(1)` is exactly two octaves and `f(4) = 16 f(1)` exactly four. By the odd-part rule a listener attending to pitch class hears three of the four symbols as the same note. The scheme is unusable and the reason is arithmetic.

The adopted assignment is the odd indices, which gives four distinct classes:

base	n	f(n) (Hz)	j	f_audible (Hz)
A	1	0.4834257330	8	123.7570
C	3	4.3508315969	8	1113.8129
G	5	12.0856433247	8	3093.9247
T	7	23.6878609164	8	6064.0924

This is the framework's odd-index route at its default assignment, carried unchanged [Curzi, 2026ad].

5.13 Stage 4 — transposition, and the separation it produces Transposition is by a fixed integer number of octaves and preserves pitch class exactly Here a single `j = 8` serves all four symbols, so the ratio structure of the alphabet after transposition is identical to the ratio structure before it. **Nothing is added and nothing is distorted; it is a change of register, the only purely presentational operation in the chain.**

The resulting separation:

sorted tones (Hz)	123.7570	1113.8129	3093.9247	6064.0924
adjacent gaps (Hz)	990.0559	1980.1118	2970.1677	
adjacent gaps (cents)	3803.9100	1768.7174	1165.0244	
MINIMUM GAP IN HZ	990.055901 Hz	(the A–C pair)		
MINIMUM GAP IN CENTS	1165.0244 cents	(the G–T pair)		
all four inside 20 Hz – 20 kHz			yes	

A correction, made visible rather than applied quietly. The source prints these two figures on one line joined by an equals sign. They are not two readings of one quantity: the alphabet is spread multiplicatively, so the pair that is closest in hertz and the pair that is closest as an interval are different pairs. The Gabor floor consumes the hertz figure and consumes the right one; a listener's frequency discrimination is described in the cents figure. Section 10 states the correction in full.

The tightest pair in the symbol alphabet, measured as an interval, is 1165.0244 cents apart — very nearly an octave — a separation any device with frequency dependence at all can resolve, and one that no plausible implementation error can permute. **That property is the whole reason the odd-index assignment is the default: it is not chosen for elegance, it is chosen because the naive one fails.**

It is worth pausing on how different that is from the alternative route. The framework's addressing paper built the same alphabet from measured infrared band centres and found the four symbols occupying less than half a semitone, with two of them two cents apart — a code that a routine solvent shift could permute outright, and one whose Gabor floor forced a rate of five and a half symbols per second. Register spreading fixed it there, by giving each symbol its own octave index: per-symbol indices $j = 33..36$ raised the minimum adjacent gap from 0.8725 Hz to 732.037 Hz, an 839.0-fold improvement, lifting the Gabor rate from 5.5 to 4599.5 symbols per second [Curzi, 2026y]. The essence-law route used here arrives at a comparable separation by a different road: the indices are already spread, because squaring them spreads them, and a single transposition therefore suffices. Neither route is more physical than the other — H5-4 forbids reading either as physical — but the essence route is the one that needs no spectrometer and no calibration, and this framework uses it for that reason and states the reason.

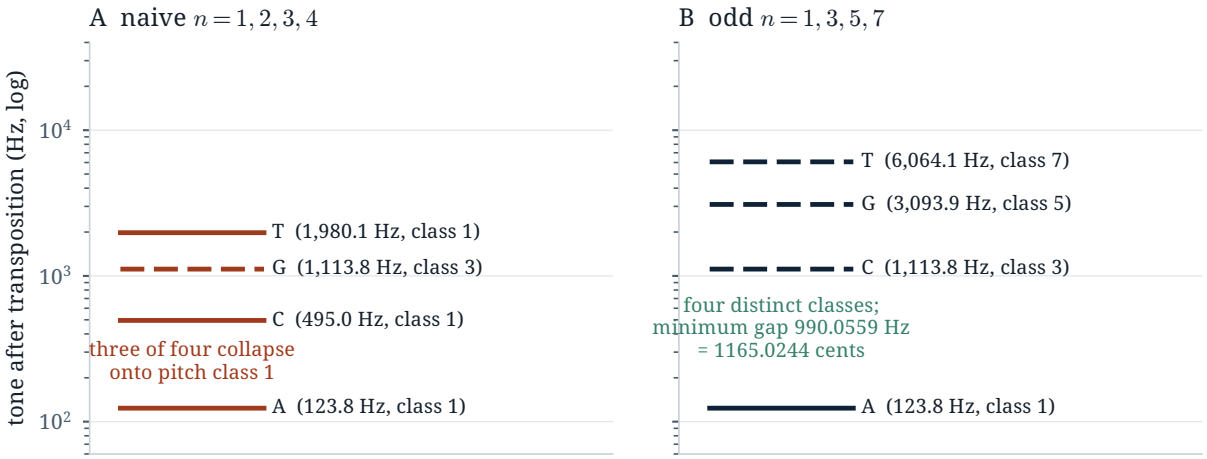


Figure 48.3: **The obvious index assignment collapses the alphabet: $n = 1, 2, 3, 4$ puts three of four symbols in one pitch class, while $n = 1, 3, 5, 7$ gives four.** Under the essence law an index and its double are exactly two octaves apart, so the pitch class of n is its odd part. Modelled from $f(n) = kn^2$ with $k = 0.483425732988414$, a **defined** constant whose interpretation as a frequency is postulated (§2.3), transposed by $j = 8$ in both panels for comparability. Built by `figs/src/CCP-048_fig01.py`.

5.14 Stage 5 — the state layer, and channel separation Section 4.14 renders the twelve states themselves: each directive at its state's essence frequency folded into 220–440 Hz, each negation exactly two octaves above. That layer and the symbol layer are both audible, both disclosed, and must be separable by ear, which requires checking rather than assuming.

```
state layer band                220.0000 – 440.0000 Hz     (j = 6 for the ladder)
nearest symbol tone below        123.7570 Hz  ->   995.9922 cents clear
nearest symbol tone above       1113.8129 Hz  ->  1607.9178 cents clear
```

The state band is clear of every symbol tone by at least 995.9922 cents — better than an octave — so the two disclosed layers occupy disjoint registers and neither masks the other. That is a computed property of this particular pairing of `j` values, not a general fact, and any change to either transposition requires recomputing it. Note also that the negation register of §4.14 (880–1760 Hz) does overlap the symbol tone at 1113.8129 Hz; where both layers are rendered simultaneously they must be separated in time instead, and the run schedule of §5.18 does exactly that.

5.15 The design table This is the table the disclosure constraints exist to produce, and every row is computed or cited with its owner named.

The Gabor limit, which binds every rendering:

```
dt * df  >=  1 / (4 pi)  =  0.07957747154594767
```

Analogy status: category 1 for the mathematics and category 3 for the physical reading. The bound is Gabor's [Gabor, 1946] and is a theorem about the Fourier transform of a signal, with no quantum content whatever. It has the same algebraic form as $\sigma_x \sigma_p \geq \hbar/2$ [Kennard, 1927], which is a fact about state preparation and not about measurement disturbance, and **nothing in this paper transfers any physical result from the second to the first.** It is measurable on an oscilloscope. The framework's conjugate-uncertainty paper owns it [Curzi, 2026z], whose Result One proves the Gabor limit $\Delta t \, \Delta f \geq 1/(4\pi) = 0.079577471545948$ by the same variational argument as $\sigma_x \sigma_p \geq \hbar/2$ with $p = \hbar k$ deleted; Gaussians saturate it and a Hann window costs 2.6 per cent.

Applied to the symbol alphabet, minimum gap 990.055901 Hz:

```
convention df = gap        dt >= 0.080377 ms
convention df = gap/2      dt >= 0.160753 ms
```

Both conventions are printed because both appear in the literature and quoting one without saying which is a reproducibility failure. The design table below uses `df = gap`.

dt (ms)	df at that window	safety over floor	symbols/s	total run, 5,640 bases
1	79.5775 Hz	12.44x	1000.00	5.640 s
5	15.9155 Hz	62.21x	200.00	28.200 s
10	7.9577 Hz	124.41x	100.00	56.400 s
20	3.9789 Hz	248.83x	50.00	112.800 s
25	3.1831 Hz	311.04x	40.00	141.000 s
50	1.5915 Hz	622.07x	20.00	282.000 s
100	0.7958 Hz	1244.14x	10.00	564.000 s

The default adopted here is 20 ms per symbol, a safety factor of 248.83 over the floor, and the factor is large on purpose. The Gabor floor is where two tones become mathematically indistinguishable; a signal whose symbols no listener can resolve is a concealed signal by another route, and **S4.3 therefore sets the rate by audibility rather than by the mathematics.** At 20 ms the full 5,640-base transmission runs 112.800 seconds.

Applied to the state layer, whose binding separation is the tightest folded gap at `N = 12` — 7.734812 Hz between classes 7 and 5 [Curzi, 2026], cited and not re-derived:

```
dt >= 0.0795774715 / 7.734812 = 0.0102882 s = 10.2882 ms
```

The state layer is nearly 128 times more demanding than the symbol layer, and that ratio is the reason the two layers get different treatment in the run schedule.

5.16 The regime, the capacity, and the two cliffs Three structural constraints, all owned elsewhere, all binding here.

(1) The 50/49 regime. The framework's reference apparatus derives it once [Curzi, 2026] and it is not re-derived here. At `N = 12` the minimum adjacent separation in the folded class lattice is 34.9756 cents, the interval `50/49`, between classes 7 and 5. At `N = 17` the minimum collapses to 6.0008 cents, the interval `289/288`, because `17^2 = 289` lands one sixth of a comma from `2 · 12^2 = 288`. The second cliff is at `N = 41`, where `41^2 = 1681` meets `2^4 · 105 = 1680` through class 29, giving `1682/1681` and 1.0296 cents. Both are exact rational coincidences.

The rate factor across the first cliff is exactly 8, because the binding pairs sit at different octave heights — `50/49` straddles `49/32` and `289/288` straddles `9/8` — so the absolute gap falls from `f/32` to `f/256`; that figure is cited [Curzi, 2026].

This framework's `N = 12` sits below the first cliff. Every folded distinction asked of a listener or an instrument here is at least 34.9756 cents wide, roughly a third of a semitone, and lengthening the index past 16 would cost a factor of eight in symbol rate.

(2) Single-transposition audible capacity. A single integer octave transposition admits at most `N = 25` essences inside 20 Hz to 20 kHz, and only at `j = 6`. Verified:

```
N = 20    admissible j = [6]
N = 25    admissible j = [6]
N = 26    admissible j = []      the set is empty
```

`N = 12` is comfortably inside the limit. A longer index would require split transposition, which is permitted but must be published per symbol.

(3) No tempering. The essence ratios are exact ratios of squares. They are not adjusted to twelve-tone equal temperament anywhere in this procedure. 34.9756 cents is not 0 cents, and a rendering that rounds it to the nearest semitone has destroyed the distinction it was built to carry.

5.17 The envelope — the rectangular gate is disqualified S4.11 forbids a rectangular gate and the reason is computable. The spectrum of a rectangular window of duration `T` is a sinc, whose squared magnitude decays as `f^-2`, so the second moment integral `int f^2 |S(f)|^2 df` diverges. Numerically, at `T = 5 ms`, integrating over successively wider bands:

```
band +-1e3 Hz     second moment  2.068251e+04
band +-1e4 Hz                    2.030538e+05
band +-1e5 Hz                    2.026834e+06
band +-1e6 Hz                    2.026465e+07
```

The value grows linearly with the integration band and does not converge, which means a rectangular gate has no defined `df` at all and the Gabor bound cannot be applied to it. Any implementation uses a windowed envelope — the specific window is published in D7, with its rise and fall times.

5.18 The run schedule Because the two layers overlap in the negation register (§5.14), and because their Gabor requirements differ by a factor of nearly 128, the declared schedule separates them in time:

```
PHASE 0    silence, 5 s                       the listener confirms the room
PHASE 1    the state layer, 24 tones          §4.14
           directive and negation as a
           double octave, 2 s per pair        48.000 s
PHASE 2    silence, 5 s
PHASE 3    the symbol layer, 5,640 bases
           at 20 ms per symbol               112.800 s
PHASE 4    silence, 5 s
           TOTAL                             175.800 s
```

Phase 1 exists so that the listener hears the structure of the corpus — twelve opposed pairs, as double octaves — before hearing the encoded transmission of it. Putting the structure first is a disclosure decision, not an aesthetic one: a listener who has heard the pairs knows what the corpus is before the long phase begins, and can stop.

5.19 What this procedure would look like if it were a technique, and why it is not A useful way to see the boundary is to write down what would have to change for any of this to be described as a technique rather than a research procedure.

It would need a demonstrated effect: a preregistered endpoint that moved, in a design with the four controls of S5.21 in place, replicated by a group that did not build the apparatus. It would need that effect to survive the scrambled-sequence control of §6.13, because an effect that does not distinguish an ordered signal from its permuted twin is an effect of sound, not of content, and the entire content architecture of Sections 4 and 5 would be beside the point. It would need a mechanism, or at minimum a candidate mechanism with its own independent evidence, because a replicated effect with no mechanism is a phenomenon to be investigated rather than a technique to be applied. And it would need a magnitude large enough to matter against the detectable floor computed in §6.9.

None of those four conditions is met. Not one of them is close to met. **The distance between "the procedure is fully specified" and "the procedure does something" is the distance this paper exists to keep visible**, and it is kept visible by writing the specification exactly as carefully as if the effect were real, and then saying, in the same breath, that it is not known to be.

The specification

Part A — the encoding chain **S5.1** The character set, byte width, byte order and separator convention shall be published.

S5.2 The bit-to-base rule shall be published, and because it is one-to-many, the actual nucleotide string shall be published in full (D3).

S5.3 The free choice shall be spent by the balanced rule of §5.9, or by another published rule whose GC and entropy figures are computed and reported alongside the greedy baseline.

S5.4 Homopolymer runs shall be constrained to at most 3, and the achieved maximum run reported.

S5.5 A round-trip hash identity shall be demonstrated and both hashes published.

S5.6 The address length shall be justified arithmetically against the target genome size, and the prefixes published in full.

S5.7 Every document, caption and file name shall state that the prefix is a routing header and that nothing is delivered, synthesised, or edited.

S5.8 The index assignment for the four bases shall be published, with the ordinal argued, and the rejected naive assignment stated so that a reader knows it was considered.

S5.9 Every published frequency shall carry its `n`, its `j`, and its band.

S5.10 Minimum adjacent gaps shall be reported in Hz and in cents.

S5.11 Channel separation between the symbol layer and the state layer shall be computed in cents and published; where the registers overlap, the layers shall be separated in time.

S5.12 The symbol rate shall be reported with its Gabor safety factor and with the convention (`df = gap` or `df = gap/2`) stated explicitly.

S5.13 The envelope shall be windowed, published, and reported with rise and fall times. A rectangular gate is disqualified.

S5.14 No tempering of any interval, at any stage.

Part B — preregistration and falsification **This is where the procedure becomes a research procedure rather than a demonstration, and none of it may be retrofitted after data exist.** The standard is the ordinary one for preregistration [Nosek et al., 2018]: the analytic choices that produce significance at will [Simmons et al., 2011, Gelman and Loken, 2013] are fixed before the data exist, and a literature that does not do this recovers far less on replication than its published record implies [Open Science Collaboration, 2015, Ioannidis, 2005].

S5.15 — The five operational choices shall be recorded in writing before any measurement: the system boundary, the input variable, the order proxy, the time base, and the cost accounting. The framework's survey and preregistration paper supplies the form [Curzi, 2026i].

S5.16 — The observable and the SESOI. The endpoint shall be named, and the smallest effect size of interest declared, before any data are collected. Section 6.9 computes what magnitudes are detectable at what sample sizes; a SESOI below the detectable floor is a preregistration that has already failed.

S5.17 — All four baselines shall be computed and reported:

```
1.  PERSISTENCE            the endpoint predicts itself at the next period
2.  DIRECT-INPUT           the input predicts the endpoint with no model between
3.  MATCHED AUTOREGRESSIVE with the SAME number of free parameters as the model
4.  SHUFFLED-INPUT CONTROL refit under the permutation, holdout permuted
                           independently, cost term NOT permuted
```

S5.18 The matched autoregressive baseline shall carry the same number of free parameters as the model. A suite is only a suite if all four are read together: evaluated alone, the matched autoregression false-passed 35.8% of the time in simulation [Curzi, 2026i].

S5.19 — The structural ceiling test shall be added, and it is not optional. Any single-period increment claimed shall satisfy the theoretical maximum $\eta \ln N$ [Curzi, 2026]. This is the only test that detects repackaging.

S5.20 — The three measured failure modes of the shuffled-input control shall be carried into every report that uses it. They were measured by simulation rather than asserted [Curzi, 2026i], and two of the three cannot be repaired by any surrogate:

```
(1) ANTI-CONSERVATIVE AS ORDINARILY RUN.
    500 null trials, nominal alpha = 0.05, measured false-rejection rates:
        iid permutation      0.108     2.2x nominal
        circular shift       0.120     2.4x nominal
        block permutation    0.060     near nominal
    A surrogate test taken off the shelf roughly DOUBLES the false-positive rate,
    and the block variant -- the one usually skipped as fussy -- is the only one
    that behaved. REMEDY: calibrate by simulation on the design in hand (series
    length, autocorrelation, block structure) and preregister the calibrated
    threshold, not the nominal one.

(2) BLIND TO REPACKAGING.
    A constructed filtered repackaging -- a rolling average of the input, and
    nothing else -- passed all four baselines including the shuffle at p = 0.0005,
    beat persistence by 89.4%, and beat the matched autoregression by 89.2%.
    It was caught by the structural ceiling test and by NOTHING ELSE: its
    single-step jump was 38.5 times the theoretical maximum.
    REMEDY: the structural ceiling test of S5.19. No surrogate detects this.

(3) BLIND TO A SHARED SLOW DRIVER.
    Where input and endpoint are both driven by a third slow variable -- a funding
    cycle, a seasonal effect, a drifting incubator -- a 4,000-trial simulation
    found the shuffled-input control PASSING 98.35% OF THE TIME. The same
    simulation reports the control working honestly in the clean case at 48.17%
    detection, which is what makes 98.35% a measurement rather than a complaint.
    REMEDY: design controls. No surrogate detects this either.
```

S5.21 — The consolidated rule, which supersedes any reading of the shuffle as a general guarantee: the shuffle is a weak, mis-calibrated empirical audit. More surrogates do not fix any of this. DESIGN CONTROLS DO. The four that answer failure modes (2) and (3):

```
BLINDING              the person measuring the endpoint does not know the condition
CHANNEL SEPARATION    the endpoint is measured through a different channel, by
                      different people, from different records, than the input
PROSPECTIVE SAMPLING  the sampling schedule is fixed before exposure and is not
                      adapted to what is observed
YOKED CONTROL         a control unit receives the same schedule, the same handling
                      and the same environment, and differs only in the ordering
                      of the signal
```

Every one of those four is a decision about the instrument, not about the analysis. **They must be written into Part 1 of the preregistration and cannot be added once the series exists.**

S5.22 — Kill conditions shall be declared in advance and honoured. The framework is rejected for a given run if any of the following holds:

```
K1   the shuffled-input control, at the CALIBRATED threshold, matches the fitted
     model on held-out data
K2   the structural ceiling test fails -- a claimed increment exceeds the maximum
```

```
K3   the scrambled-sequence control (Section 6 §6.13) is indistinguishable from
     the ordered condition
K4   the yoked control moves with the exposed condition
K5   the fitted parameters are not this system's, per FR3
```

S5.23 — A null result shall be reported, in full, with the same prominence as a positive one. Section 6.11 computes how little of the design space one null bounds — 0.000801% of a modest 124,800-cell grid — so a null is a small, honest, publishable increment and must not be inflated into a refutation or buried as a disappointment.

S5.24 — The labelling rule of §5.2 applies to every output, including preprints, slides, captions, file names, and any summary written for a non-technical reader.

The failure modes

FM-5.1 — The clinical drift in the caption. The body text is careful and the figure caption says "response." Captions are read more than bodies.

Fix: S5.24 and §5.2's forbidden list, applied mechanically with a search before release.

FM-5.2 — The address becomes a molecule in the retelling. "gRNA-format prefix" is compressed to "guide RNA" in a summary, and the summary now describes a delivered molecule.

Fix: S5.7. The full phrase "routing header, not a molecule" appears in every document, not once per document.

FM-5.3 — The transposition read as a resonance. Somebody concludes that cytosine "has a frequency of 1113.8129 Hz."

Fix: H5-4. The number is a label produced by a postulated square law applied to an alphabetical ordering and then shifted eight octaves. **It is not a property of cytosine.**

FM-5.4 — The naive index assignment reintroduced. A later implementer assigns `n = 1,2,3,4` because it is obvious, and three symbols collapse to one pitch class.

Fix: §5.12, which rejects it in print with the arithmetic, so the rejection is findable.

FM-5.5 — The greedy encoder shipped. The first-legal rule is simpler and passes the round-trip test, so it looks fine. It produces a 60% adenine string that no listener can follow.

Fix: S5.3. Report GC and entropy against the greedy baseline every time.

FM-5.6 — The rate set by the Gabor floor. An engineer reads 0.080377 ms, sets the symbol rate near it, and produces a signal that is technically not masked and practically inaudible as symbols.

Fix: S4.3 and §5.15. **The floor is where two tones become mathematically indistinguishable; it is not a perceptual threshold, and a rendering at the floor is concealment achieved by speed.**

FM-5.7 — The shuffle treated as a guarantee. The control passes, the result is announced, and the two failure modes that a surrogate cannot see are never addressed.

Fix: S5.20 and S5.21, carried in the report rather than in the methods appendix.

FM-5.8 — Design controls added after the fact. Blinding is "not practical" and is replaced with a fifth surrogate.

Fix: S5.21. Surrogates and design controls answer different questions and one cannot substitute for the other at any exchange rate.

FM-5.9 — The preregistration adapted after the first look. The endpoint is redefined, the window moved, or the SESOI lowered, all of them defensibly.

Fix: S5.15 and S5.16, timestamped and deposited, with any deviation reported as a deviation.

FM-5.10 — The null buried. A run returns nothing and is not written up.

Fix: S5.23. A null bounds a real, computable, small fraction of the design space and is worth exactly that much — which is more than zero and much less than a refutation.

FM-5.11 — The scrambled-sequence control omitted because it "obviously" makes no difference. This is the control that decides the entire question and it is the one most likely to be dropped as redundant.

Fix: K3, and §6.13, which explains why no other control substitutes for it.

FM-5.12 — Layer overlap unchecked. A later implementer changes one transposition and the state layer lands on top of a symbol tone.

Fix: S5.11. The 995.9922-cent clearance is a property of one specific pair of `j` values and must be recomputed after any change.

FM-5.14 — The corpus rewritten between runs. The payload changes because a directive read poorly, and two runs are pooled as though they were the same condition. Because the encoding is deterministic, a single changed character changes the payload hash, the base string, and every symbol after the change point.

Fix: S5.5. The payload hash is the run's identity. Two runs with different payload hashes are two conditions, and pooling them requires the same justification as pooling any two conditions.

FM-5.13 — The procedure described as having worked. Somebody who ran it reports that "it went well," meaning the audio rendered correctly, and a reader hears a claim about a biological effect.

Fix: §5.2, and a reporting convention that separates *the procedure executed as specified* from *the endpoint moved*, which are different sentences about different things.

Worked example

CONSTRUCTED. The payload is the corpus of §4.28, written for this paper. It is not a transcript and not an observation of anyone. No exposure has been run and no endpoint has been measured.

5.25 End to end, with every number

```
STAGE 1     text
            characters (with newline separators)          675
            bytes                                         675
            bits                                        5,400
            SHA-256  c7d98b1cfde5c515ce99ff7db77d484f71a09c5cd921994f15df66c50373b72d

STAGE 1b    bases, 1 bit/base, balanced free choice
            payload bases                               5,400
            A / C / G / T                   1750 / 1749 / 951 / 950
            GC                                       0.500000   (= L/2 exactly)
            maximum homopolymer run                         1   (RLL(3) satisfied)
            symbol entropy                           1.935874 bits/base
            greedy baseline entropy                  1.244171 bits/base
            first 60 bases
              AGCATGCACTACAGCTAGCTAGCACATCACACAGCTAGCACTACGACACTACAGCTACGA
            SHA-256 of recovered text
              c7d98b1cfde5c515ce99ff7db77d484f71a09c5cd921994f15df66c50373b72d
            ROUND TRIP IDENTICAL

STAGE 2     address
```

```
            12 prefixes x 20 nt                          240 bases
            expected spurious matches per 20-mer     0.005638867
            address space / genome positions           177.3406
            TOTAL TRANSMISSION                          5,640 bases

STAGE 3     sonify, essence law, Route B odd indices, k = 0.483425732988414 Hz
            A  n=1   f =  0.4834257330 Hz
            C  n=3   f =  4.3508315969 Hz
            G  n=5   f = 12.0856433247 Hz
            T  n=7   f = 23.6878609164 Hz

STAGE 4     transpose, j = 8 for all four symbols
            A     123.7570 Hz
            C    1113.8129 Hz
            G    3093.9247 Hz
            T    6064.0924 Hz
            minimum adjacent gap    990.055901 Hz  =  1165.0244 cents

STAGE 5     state layer, §4.14
            directives      220.0000 -  440.0000 Hz   (ladder at j = 6)
            negations       880.0000 - 1760.0000 Hz   (+2 octaves, 2400.0 cents)
            clearance to nearest symbol tone below     995.9922 cents

STAGE 6     render
            symbol duration                              20 ms
            Gabor floor (df = gap)                   0.080377 ms
            safety factor                               248.83x
            envelope                                 windowed; rectangular disqualified
            tempering                                none
            phase 1 (state layer)                    48.000 s
            phase 3 (symbol layer)                  112.800 s
            total with silences                     175.800 s
```

5.26 The disclosure package, as delivered D1 through D10 of §4.19, complete, before exposure. D1 contains the twelve pairs in full including the four that fail the equal-force check of §4.29 — pairs 3, 5, 7 and 12, 33.33% of the corpus — named as failing, because S4.8 permits rendering an unrepaired corpus only if the failures are disclosed. D3 contains all 5,400 payload bases. D9 contains H5-1 through H5-7 and §6's open problem in full.

5.27 The preregistration, filled in **No data exist.** This is the form as it would stand before a first run, and it is included so that a reader can see what a complete one looks like.

```
1. SYSTEM BOUNDARY      cultured cells, in vitro, named line, named passage range.
                        Deliberately excluded: any human participant, any endpoint
                        that is not a transcript abundance.
2. INPUT                the rendered waveform of §5.25, as a file, with its hash.
3. ORDER PROXY          preregistered transcript abundances at the 12 addressed loci.
4. TIME BASE            fixed sampling schedule, declared before exposure (S5.21).
5. COST ACCOUNTING      declared per the framework's preregistration form.

OBSERVABLE              log2 abundance at the 12 preregistered loci.
SESOI                   declared in advance; must exceed the detectable floor of
                        section 6.9 at the planned n, or the design fails
                        before it starts.
```

```
ARMS                  ORDERED       the waveform as specified
                      SCRAMBLED     symbol-permuted twin: identical total energy,
                                    duration, symbol counts and long-term power
                                    spectrum; ordering destroyed  (§6.13)
                      SILENT        handling-matched, no signal
                      YOKED         same schedule, same handling, ordering varied

DESIGN CONTROLS       blinding of the measurer to condition          REQUIRED
                      channel separation of endpoint from input      REQUIRED
                      prospective sampling schedule                  REQUIRED
                      yoked control                                  REQUIRED

BASELINES             persistence, direct-input, matched autoregressive,
                      shuffled-input -- all four, read together (S5.17, S5.18)

STRUCTURAL TEST       ceiling test per S5.19, single-period increment vs
                      theoretical maximum

SURROGATE CALIBRATION block permutation, calibrated by simulation on THIS design;
                      the calibrated threshold is preregistered, not the nominal
                      0.05 (S5.20 failure mode 1)

KILL CONDITIONS       K1 - K5 of S5.22, declared and honoured

REPORTING             null reported in full, with the same prominence as a
                      positive result, and with the design-space fraction it
                      bounds computed (S5.23)
```

5.28 What this procedure has and has not produced **Has produced:** a deterministic encoder with a verified round-trip identity; a balanced free-choice rule that holds GC at exactly half and raises payload entropy from 1.244171 to 1.935874 bits per base; an arithmetically justified 20-nt routing header; a four-symbol alphabet whose tightest pair is 1165.0244 cents apart; two disclosed layers separated by 995.9922 cents; a symbol rate 248.83 times above the Gabor floor; and a 175.800-second audio file with a complete disclosure package.

Has not produced, and does not claim: any biological effect; any evidence that a cell responds to the ordering of the signal; any clinical, medical or therapeutic result; any alteration of DNA; any conferred capability; or any physical resonance. **The procedure ends at a loudspeaker, and §6 is about the gap on the other side of it.**

What this section established **Established, and checkable:** that a written corpus maps losslessly to a 5,400-base nucleotide string with a verified SHA-256 identity; that the balanced free-choice rule holds GC within one base of `L/2` by construction and measured exactly 0.500000 here, with maximum homopolymer run 1 and entropy 1.935874 against the greedy rule's 1.244171; that a 20-nt prefix has 0.005638867 expected spurious matches in 6.2e9 positions; that the naive base index assignment collapses three of four symbols onto one pitch class and is rejected; that the odd-index assignment gives a minimum separation of 990.055901 Hz; that the symbol and state layers are 995.9922 cents apart; that at 20 ms per symbol the rendering sits 248.83 times above the Gabor floor and runs 175.800 seconds; that the rectangular gate has no finite spectral second moment; and that the shuffled-input control fails in three measured ways, of which two are answerable only by design controls.

Not established, and not claimed: that any of this affects any cell, transcript or person; that any tone bears a physical relation to any base; that the procedure is a treatment, therapy or

practice of any kind; that DNA is edited, altered or damaged; or that any capability has been conferred.

The one line to carry forward: *the whole procedure is buildable, auditable and disclosed, and it stops at a loudspeaker.*

Section 5 owns the operational procedure, the design table, the run schedule, and the preregistration and falsification apparatus. Section 4 owns the no-subliminals constraint, the disclosure package and the consent specifications, all of which this section inherits whole; Section 6 owns the open problem that governs both.

6 The open problem, stated precisely

Everything in this paper up to this point has been either computed, proved, cited, or badged as a postulate in the sentence that stated it. Sections 4 and 5 have specified an encoder, an addressing convention, a sonifier and a rendering discipline, all of them deterministic and all of them verified.

There is one thing missing from that chain and it is not a detail. It is the link between the signal and the archive.

> **There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal.**

That sentence is this section. It is carried in every part of this paper and it is not softened here, not worked around, and not argued down. This section spends its length making the gap as precise as it can be made, because a gap stated precisely is a research programme and a gap stated vaguely is an excuse.

The reason for the length is intellectual accounting. A framework that carries a chosen reference constant, a postulated square law, a differentiation index whose defensibility it admits is contestable, and a sonification procedure whose biological effect is unknown has to earn its reader's attention somewhere. **A framework that names its own load-bearing gap this precisely earns the right to everything else in it;** one that declines to name it forfeits the right to all of it, including the parts that are correct.

The law

6.1 The gap, stated four ways The same proposition at four levels of precision, because each level is the one somebody will need.

Colloquial. Nobody knows how a cell could read an ordered message out of a sound.

Engineering. No known biological structure performs demodulation of a symbol stream from an acoustic carrier into a sequence-specific output. There is no identified transducer with the required frequency selectivity, no identified memory element that preserves symbol order across the symbol interval, and no identified mapping from a decoded symbol string to sequence-specific binding at a genomic locus.

Formal. Let `s(t)` be an acoustic pressure waveform encoding a symbol string `x = (x_1, ..., x_L)` over an alphabet of size 4, and let `R` be the cell's transcriptional response. There is no established mechanism producing an `R` that depends on `x` as an *ordered string*, as distinct from depending on functionals of `s(t)` that are invariant under permutation of the symbols — total energy, duration, long-term power spectrum, envelope statistics, symbol counts. **The claim that requires a mechanism is the claim that `R` is a function of the ORDER, and that is exactly the claim for which no mechanism exists.**

Operational. Take a stimulus and its symbol-permuted twin. The two have identical total energy, identical duration, identical long-term power spectrum, and identical symbol counts. If a cell cannot be shown to distinguish them, no sequence-specific channel has been demonstrated, and nothing else in the experiment matters.

6.2 What the statement does not say

Two symmetric errors sit on either side of the sentence and both have to be refused explicitly, because each is comfortable in a different direction.

It is not a claim of impossibility, and impossibility is not claimed. Absence of a known mechanism is a statement about the current state of the literature. Several real channels between physics and biology were identified late and were initially treated as implausible; the burden of demonstration sits entirely on the claimant, which is the ordinary position for any unestablished claim and carries no special hostility toward this one.

Nor is absence of a known mechanism a licence. The symmetric error is to treat "not proven impossible" as though it were partial support. It is not support of any kind. **A claim with no mechanism and no demonstration has zero evidential standing,** and the correct number of qualifiers to attach to it before demonstration is not "few" but "it is not asserted."

Both errors share a structure worth naming: they convert a statement about evidence into a statement about the world. The asymmetry between a conjecture and a corroborated one is the ordinary one [Popper, 1959], and a research programme is assessed by whether it makes new, checkable predictions rather than by whether it has been disproved [Lakatos, 1970, Platt, 1964]. The sentence at the head of this section is about what is known. It is not a prediction, and it will be settled by an experiment rather than by an argument.

6.3 Why this section is the framework's warrant

The ordering of this paper is deliberate and it runs against the commercial instinct. Section 4 built the ethical architecture, Section 5 built the procedure, and Section 6 says that nobody knows whether the procedure does anything. A paper optimising for persuasion would reverse the last two, or merge this section into a paragraph, or place it in an appendix labelled "limitations."

The reason for the ordering is that the alternative is not available to work that wants to be checked. Every number in Sections 4 and 5 is checkable by someone with a Python interpreter, and a reader who checks a hundred of them and finds them right acquires a reasonable prior that the hundred and first is right too. That prior is exactly what an unmarked overclaim would spend. **The restraint is not modesty and it is not a legal posture; it is the only way the arithmetic retains any value to a reader who cannot personally verify all of it.**

6.4 Honesty clauses, binding on this section and inherited by everything that cites it

H6-1 — No clinical, medical or therapeutic claim, and no intervention on any person is described here. Every experiment specified in this section is *in vitro*, on cultured cells, with transcript abundances as endpoints.

H6-2 — No claim that DNA is edited, altered, or damaged. Nothing in the pipeline of §5 carries a nuclease, a template, or a Cas protein, and nothing is delivered.

H6-3 — No claim that any capability has been conferred on anything.

H6-4 — Essence frequencies are essence frequencies, not physical resonances. Nothing anywhere in this paper claims that a person, tissue, cell or molecule resonates at any value printed in it.

H6-5 — The concordance columns carry different things. Where a traditional source is cited anywhere in this paper it carries *precedent*: that an idea was arrived at independently and preserved. It does not carry *evidence* [Curzi, 2026w], and this section draws no concordance at all.

H6-6 — Every number in this section was computed from stated assumptions, and the assumptions are printed with the number. The power calculations of §6.9 in particular are

arithmetic on declared inputs. They are not estimates of anything measured, and the variance figures in them are inputs a reader supplies from their own system, not values this framework asserts about biology.

6.5 What is real upstream, itemised The gap is a single missing link, not a general absence of evidence, and it matters to say so with the same precision used to state the gap. Otherwise an honest reader cannot tell how much of the architecture is standing.

Component	Status	Source
Arbitrary digital data can be stored in and recovered from nucleotide sequence	**ESTABLISHED**	Church, Gao & Kosuri, *Science* 337:1628 (2012)
Higher-rate DNA storage codes exist and are demonstrated	**ESTABLISHED**	Goldman et al., *Nature* 494:77 (2013); Erlich & Zielinski, *Science* 355:950 (2017)
A ~20-nt spacer is a sufficient genomic address	**COMPUTED**, and standard CRISPR practice	§5.10
A CRISPR effector self-processes a multi-address array	**ESTABLISHED**	Zetsche et al., *Cell* 163:759 (2015); Fonfara et al., *Nature* 532:517 (2016)
Targeting and editing are separable functions	**ESTABLISHED**	Qi et al., *Cell* 152:1173 (2013); Gilbert et al., *Cell* 159:647 (2014)
Mechanical force gates ion channels	**ESTABLISHED**, Nobel 2021 (Patapoutian)	PIEZO1/PIEZO2
Surface force reaches the nucleus, stretches chromatin, and alters transcription	**ESTABLISHED**	Tajik et al., *Nature Materials* 15:1287 (2016)
Ultrasound can activate specific engineered channels	**ESTABLISHED**	Ibsen et al., *Nature Communications* 6:8264 (2015); Duque et al. (2022)
Biology builds exquisite mechanical receivers	**ESTABLISHED**	cochlear hair cells
Non-coding sequence is a large, structured, addressable archive	**ESTABLISHED**	reviewed in Curzi [2026y]
Regulatory input reorganises expression without adding sequence	**ESTABLISHED**, Nobel 2012	Takahashi & Yamanaka, *Cell* 126:663 (2006)
A cell demodulates sequence-specific content from sound	**NO ESTABLISHED MECHANISM**	**§6**

Eleven established rows and one open row. Both halves belong in any honest report. Reporting only the eleven is advocacy; reporting only the one is a different dishonesty, because it implies the whole edifice is speculation when most of it is textbook.

6.5b The archive the signal would be addressing One of the eleven rows deserves expansion, because it is the one that makes the architecture worth specifying at all and it is also the one most often described loosely.

The phrase *junk DNA* is retired. Non-coding sequence is roughly 98% of the human genome and is largely regulatory: promoters, enhancers, silencers and insulators, together

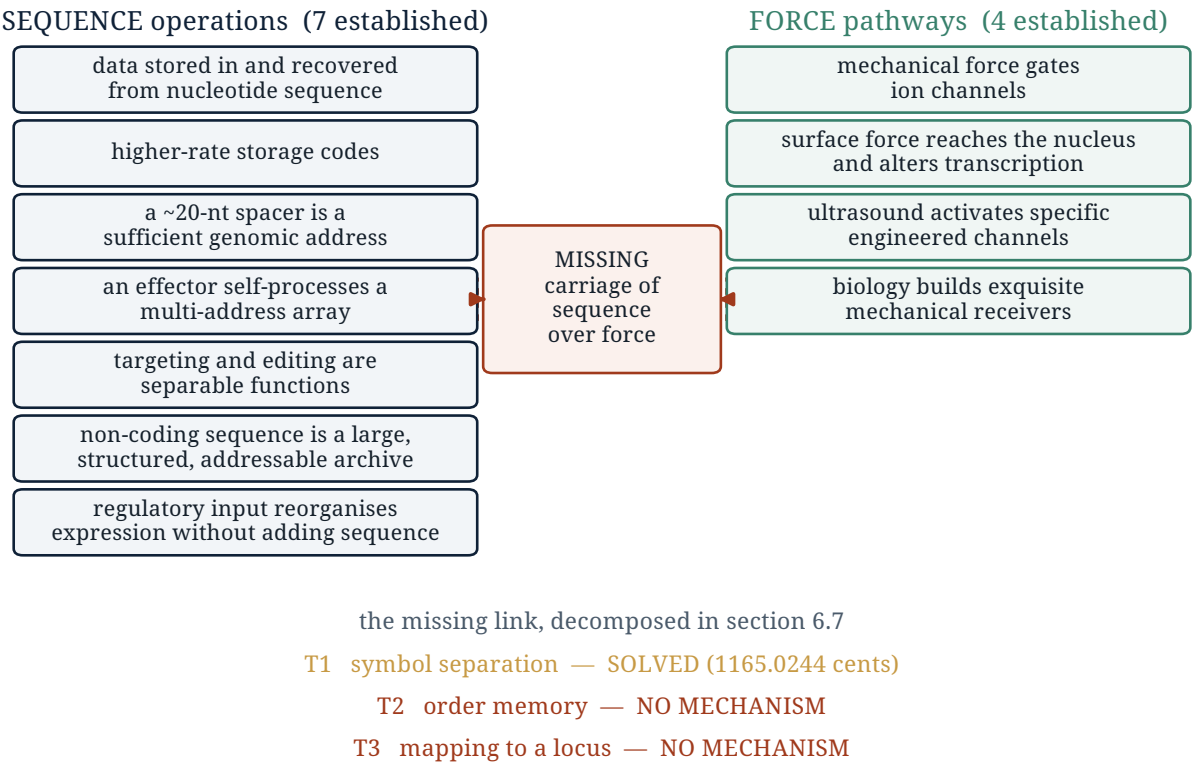


Figure 48.4: **Eleven upstream components are established and the one missing link is the carriage of sequence over force.** Sequence operations store, address, parse, target, archive and regulate; force pathways gate channels and reach the nucleus. Nothing in the established literature carries a symbol string across the boundary. A classification of the cited literature, not data; the primary biological sources are named in the text without citation keys because the corpus's bibliography carries no cluster covering them (§10, L-9). Built by `figs/src/CCP-048_fig05.py`.

with transposable elements at roughly 45% of the genome and endogenous retroviral sequence at roughly 8%, and the large families of long non-coding and micro RNAs. That is not filler around the genes; it is the control surface, and it is the part of the genome that regulatory input acts on.

The single most instructive item in it, for an argument about narrative, is syncytin. The gene indispensable to placenta formation derives from a retroviral envelope gene: ancient foreign text, inherited, later read out as vital function. **A sequence that arrived from outside became a sequence the organism could not do without, which is documented molecular biology and not a metaphor** — though it is very hard to read about it without reaching for one, and the reaching is worth resisting in print.

What this establishes here is narrow. There is a large, structured, addressable archive, and regulatory input demonstrably reorganises which parts of it are expressed — Yamanaka reprogramming is the strongest available demonstration that latent capability is activated rather than added. **The archive is real and the regulatory layer is real. What is missing is any route by which an acoustic symbol stream reaches either.**

6.6 The shape of the missing link Look at the eleven established rows and sort them by what they carry.

Rows 1 through 5 and rows 10 and 11 are **sequence** operations. They store, address, parse, target, archive and regulate, and every one of them operates on symbol strings or on molecular structures that read symbol strings.

Rows 6 through 9 are **force** pathways. Mechanical deformation opens a channel, ions move, signal propagates; force at the surface reaches the nucleus and changes what is transcribed; engineered channels respond to ultrasound; hair cells transduce sound exquisitely.

Nothing in the established literature carries sequence over force. A pressure wave that opens a PIEZO channel does so by deforming the membrane, and the channel does not care what melody deformed it. Tajik et al. applied force and measured transcription; the force carried a magnitude and a direction, not a string. Sonogenetics gets its specificity from the protein that was engineered into the cell, not from the structure of the waveform — which is a point worth dwelling on, because sonogenetics is the result most often produced as though it closed this gap. It does the opposite: it is the clearest existing demonstration that acoustic specificity in biology has so far been achieved by putting the specificity in the receiver, and never by putting it in the signal.

6.7 Three requirements, and which are unmet Stating the gap as a single sentence is right for a canon file and too coarse for a research programme. Decomposed, a sequence-specific acoustic channel needs three distinct things, and they fail differently:

```
T1  SYMBOL SEPARATION
    The receiver must distinguish the symbols of the alphabet.
    STATUS: satisfiable in principle. Section 5.13 builds an alphabet whose
    tightest pair is 1165.0244 cents apart -- a separation any device with
    frequency dependence at all can resolve. D1 is an engineering requirement
    and the engineering exists.

T2  ORDER MEMORY
    The receiver must preserve symbol ORDER across the symbol interval, so that
    the response depends on the sequence and not on the histogram.
    STATUS: NO IDENTIFIED MECHANISM. This is the heart of the gap. A permutation-
    invariant response is not a weak version of a sequence-specific one; it is a
    different quantity, and no known molecular element stores an acoustic symbol
    stream in order.

T3  MAPPING TO A LOCUS
```

```
The decoded string must produce sequence-specific binding or regulation at the
addressed locus.
STATUS: NO IDENTIFIED MECHANISM connecting a decoded acoustic string to
anything at a locus. Targeting itself is established (Qi et al. 2013;
Gilbert et al. 2014, named without keys) -- but it is established for a
molecule that carries the address, not for a string that was heard.
```

T1 is solved and is the least interesting of the three. T2 and T3 are the open problem, and an easily separable alphabet does not touch either of them. That distinction is worth carrying, because engineering effort naturally flows toward T1 — it is the part where progress is visible — and progress there is not progress on the question.

6.7b Why T2 is hard in particular, and where the difficulty comes from It is worth saying why order memory is the hardest of the three rather than merely the unsolved one, because the reason is physical rather than a matter of nobody having looked.

A receiver that preserves symbol order needs to hold state across the symbol interval. In the rendering of §5 that interval is 20 milliseconds. Any candidate molecular element would have to retain a distinguishable configuration for that long, in water, at body temperature, against thermal agitation, and then combine it with the next symbol in a way that depends on which came first. The framework's addressing paper computes the relevant damping and finds molecular-scale mechanical resonators in water heavily overdamped, with Stokes numbers far below one [Curzi, 2026y]; that calculation is cited here rather than repeated, and its consequence is that the obvious mechanical route to a memory element — a small structure that rings — is not available.

That does not exhaust the possibilities. Biology stores order in chemistry all the time, in polymerisation, in phosphorylation cascades, in transcriptional bursting histories. **The point is not that order memory is impossible; it is that the candidate everybody reaches for first — a tiny resonator that rings at a symbol's frequency and remembers — is ruled out by damping, and the remaining candidates are chemical rather than mechanical and have no proposed coupling to an acoustic symbol stream at all.**

The engineering

6.8 What an experiment must show The experiment that would settle this is not exotic and its logic is short.

The claim under test is that the cell's response depends on the *ordering* of the symbol stream. The null is that it depends only on permutation-invariant functionals of the waveform. Therefore:

Compare an ordered stimulus against its symbol-permuted twin, and measure whether the response distinguishes them. Everything else in the design is in service of that comparison.

The permuted twin is not an approximate control. It is exact on every quantity the null could appeal to:

```
identical      total acoustic energy
identical      duration
identical      symbol histogram (the same count of each symbol)
identical      long-term power spectrum, under a fixed symbol alphabet
identical      peak level, RMS level, and crest factor
DIFFERENT      the ordering, and only the ordering
```

The size of the space the permutation draws from is worth computing, because it shows how little the control costs and how much it isolates. For the 5,400-base payload of §5.25,

whose composition is A 1750, C 1749, G 951, T 950, the number of distinct symbol orderings with identical composition is:

```
L! / (n_A! n_C! n_G! n_T!)   =   10^2017.717
```

Every one of those 10^2017.717 waveforms is indistinguishable from the ordered stimulus by every permutation-invariant measurement, and a drawing from that set costs nothing but a shuffle. A design that omits this control has omitted the only comparison that could have answered the question.

6.9 What magnitude would even be detectable

A gap is only a research programme if somebody can say what size of effect the programme could see. The arithmetic below is a standard two-sample power calculation on log2-transformed abundance, with equal group sizes:

```
n  =  2 ( z_(1 − alpha/2) + z_(1 − beta) )^2 sigma^2 / delta^2
```

`delta` is the difference in log2 abundance, `sigma` the standard deviation in the same units, power fixed at 0.80 throughout. Both `sigma` and the number of tested transcripts are inputs a reader supplies from their own system; nothing below is a claim about biology, only arithmetic on declared assumptions (H6-6).

(a) A targeted readout of the twelve preregistered loci, Bonferroni `alpha = 0.05/12`:

```
z_(1−alpha/2) = 2.865260      z_0.80 = 0.841621

delta (log2FC)   sigma=0.20   sigma=0.35   sigma=0.50
    0.10             110          337          688
    0.20              28           85          172
    0.30              13           38           77
    0.50               5           14           28
    1.00               2            4            7
    2.00               1            1            2
```

(b) A transcriptome-wide readout, 20,000 transcripts, Bonferroni `alpha = 0.05/20000 = 2.5e−6`:

```
z_(1−alpha/2) = 4.708130

delta (log2FC)   sigma=0.20   sigma=0.35   sigma=0.50
    0.10             247          755         1540
    0.20              62          189          385
    0.30              28           84          172
    0.50              10           31           62
    1.00               3            8           16
    2.00               1            2            4
```

Inverted, which is the more useful direction — the smallest detectable log2 fold change at a fixed sample size, `sigma = 0.35`, power 0.80:

```
n per arm    targeted (12 loci)     genome-wide (20,000)
    3            1.0593                 1.5860
    6            0.7491                 1.1215
   10            0.5802                 0.8687
   20            0.4103                 0.6142
   50            0.2595                 0.3885
  100            0.1835                 0.2747
```

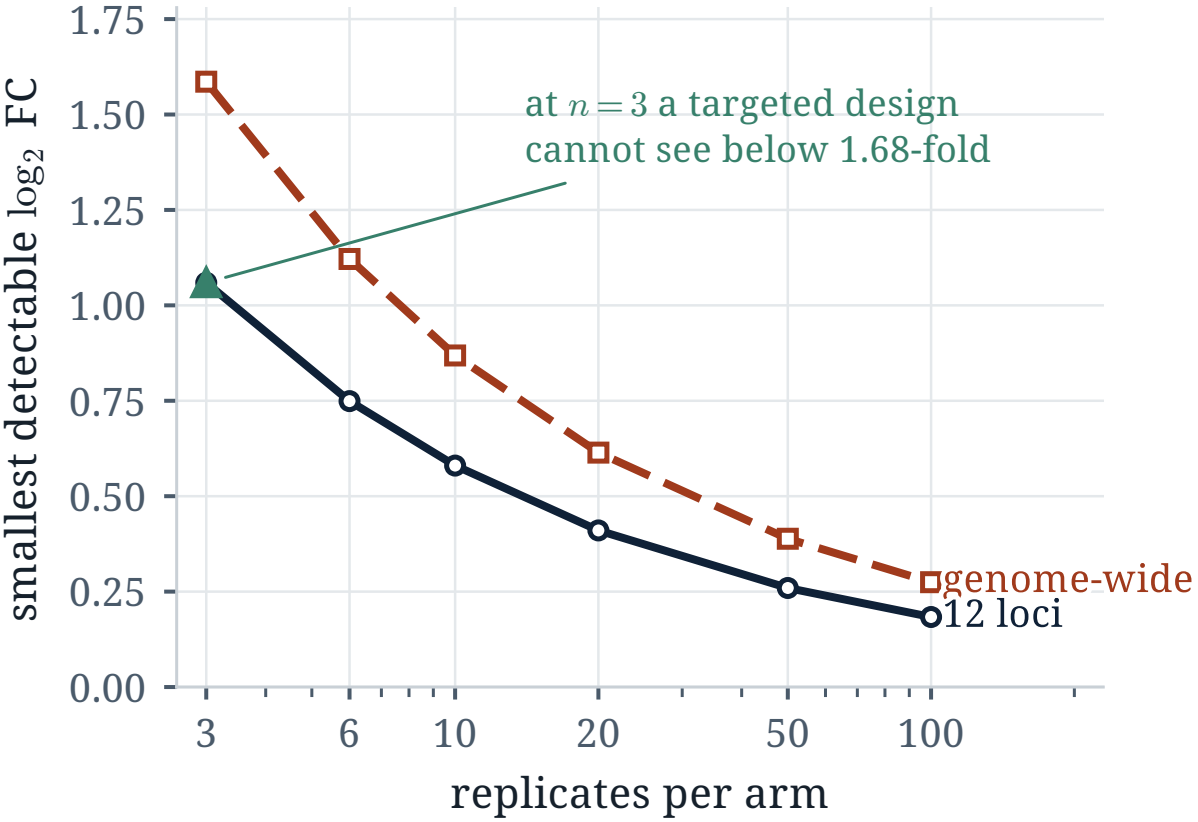


Figure 48.5: **A three-replicate design cannot see below a 1.68-fold change on a targeted readout, or a 2.18-fold change genome-wide.** Smallest detectable $\log_2$ fold change against replicates per arm, at power 0.80 and Bonferroni correction across 12 loci and across 20,000 transcripts. Modelled: the dispersion $\sigma = 0.35$ $\log_2$ units and the number of tests are declared inputs an operator supplies from their own pilot, **not measurements and not claims about biology** (H6-6). Built by `figs/src/CCP-048_fig03.py`.

In linear fold-change terms, at `sigma = 0.35`:

```
n =   6    targeted floor  x1.6807      genome-wide floor  x2.1757
n =  20    targeted floor  x1.3289      genome-wide floor  x1.5308
n = 100    targeted floor  x1.1356      genome-wide floor  x1.2097
```

Three consequences follow and each is a design decision rather than an observation.

First, a typical three-replicate design cannot see anything under a 1.68-fold change even on a targeted readout, and under a 2.18-fold change genome-wide. An effect smaller than that is not absent from such an experiment; it is invisible to it, and reporting “no effect” from that design is reporting the design rather than the biology.

Second, the SESOI must be declared against this floor before the run. Section 5 S5.16 makes that a specification: a smallest effect of interest below the detectable floor at the planned `n` is a preregistration that has already failed, and discovering it afterwards converts a null into an uninterpretable one.

Third, targeting twelve preregistered loci buys roughly a factor of two in required sample size against a genome-wide search, at the price of being unable to find anything anywhere else. That is a real trade and it should be made deliberately, in the preregistration, with the reason written down. The recommendation here is the targeted readout for a first pass — because the twelve addressed loci are exactly the ones the architecture claims to address, and

an effect there is the one the architecture predicts — with a genome-wide arm as a secondary, exploratory, and clearly labelled analysis.

A note on the multiplicity correction, because it is the assumption most likely to be argued with, and because the choice of correction after seeing the result is one of the best-documented routes to a false positive [Simmons et al., 2011, Gelman and Loken, 2013]. Bonferroni is used above because it is transparent, because it is the most conservative common choice, and because a section whose whole subject is not overclaiming should not be quietly buying sensitivity from a correction the reader has to know about to evaluate. A Benjamini-Hochberg procedure at the same nominal rate would be less demanding, would raise the detectable effect at a given sample size less steeply, and would be entirely defensible — provided it is the preregistered choice rather than the one arrived at after the Bonferroni analysis returned nothing. The rule is not which correction; the rule is that the correction is fixed before the data exist, stated with the number of tests it covers, and reported with the result.

Two further inputs deserve the same treatment. The variance `sigma` is the operator's, from their own pilot, and substituting an optimistic value is the cheapest way to make an underpowered design look adequate on paper. And the power target of 0.80 is a convention rather than a law; a question this consequential arguably deserves 0.90, which costs roughly a third again in sample size and should be priced explicitly rather than defaulted past.

6.10 One route closed by arithmetic Not all of the gap is equally open, and a closed route is progress worth publishing.

For a plane acoustic wave in an aqueous medium, the longitudinal strain amplitude is exactly `s = p / K`, with `K` the bulk modulus of water, `2.2e9 Pa`. This is elementary continuum mechanics and is not a model of anything biological. Computed:

dB SPL	p (Pa)	bulk strain
40	0.002	9.090909e-13
60	0.02	9.090909e-12
80	0.2	9.090909e-11
94	1.00237	4.556248e-10
100	2.0	9.090909e-10
120	20.0	9.090909e-09

At the levels a disclosed rendering can plausibly use — the rendering of Section 5 is meant to be listened to, not endured — the bulk compressional strain is of order `1e-11` to `1e-10`. **Bulk compression is not the route.** Whatever an acoustic channel to a cell would be, it is not a matter of squeezing the cell smaller by one part in ten billion and expecting the genome to notice the pattern of the squeezing.

That does not close the question, and the reason is that the established pathways are not bulk pathways. PIEZO gating is discussed in terms of membrane tension and forces at the channel; Tajik et al. applied force at the surface and followed it through the cytoskeleton. Those are structural, local, and highly leveraged routes, and the bulk number above says nothing about them. **What the arithmetic closes is the naive route, and closing a naive route matters mainly because it is the one most often assumed.**

6.11 How little a null result bounds The last computed piece is the one that keeps both a positive and a negative result honest: how much of the design space a single condition actually covers.

A modest grid over the parameters an experimenter must choose anyway:

```
carrier band (20 Hz – 20 kHz, 1/3-octave steps)          30
level (40–100 dB SPL, 5 dB steps)                        13
symbol rate (5–200 symbols/s, half-decade steps)          4
```

```
exposure duration (1 min – 24 h, decade steps)            4
cell type                                                 5
readout time after exposure                               4
------------------------------------------------------------
TOTAL GRID CELLS                                    124,800

one condition bounds     1 / 124,800   =   0.000801 %
ten conditions           0.00801 %
one hundred conditions   0.0801 %
```

A single well-run null bounds eight ten-thousandths of one percent of a deliberately modest grid, which is a real and publishable increment and is nowhere near a refutation. Prediction is one of many legitimate reasons to build a model, and bounding a design space is another [Epstein, 2008]. This cuts both ways and is meant to. It protects a null from being inflated into "the idea is dead," and it protects the field from a claimant who runs one condition, finds something, and announces a channel.

The grid is also an argument for how to spend effort. Given how little any single cell covers, the value of a design comes from its controls rather than from its coverage: **one condition run with the scrambled-sequence control, blinding, channel separation, prospective sampling and a yoked control is worth more than a hundred conditions run without them**, because a hundred uncontrolled conditions bound nothing at all.

The specification

6.12 The decisive experiment, numbered **S6.1** The experiment shall be *in vitro*, on a named cell line at a named passage range, with transcript abundances as endpoints. No human participant appears in it.

S6.2 The stimulus shall be produced by the procedure of §5, with its payload hash published, and delivered by a characterised transducer whose output is measured at the sample position, not assumed from the amplifier setting.

S6.3 The four arms shall be:

```
ORDERED       the waveform as specified by §5.25
SCRAMBLED     a symbol-permuted twin, drawn from the 10^2017.717 orderings with
              identical composition; identical energy, duration, histogram and
              long-term spectrum; ordering destroyed
SILENT        handling-matched, no signal; the transducer present and inactive
YOKED         same schedule, same handling, same environment, ordering varied
```

S6.4 Allocation shall be randomised and the person performing the measurement shall be blind to the arm.

S6.5 The sampling schedule shall be fixed before exposure and shall not be adapted to what is observed.

S6.6 The endpoint shall be measured through a different channel, by different people, and from different records than the stimulus log. Channel separation is a design requirement, not an analysis choice.

S6.7 The SESOI shall be declared in advance and shall exceed the detectable floor of §6.9 at the planned `n`.

S6.8 The primary analysis shall be the ORDERED versus SCRAMBLED contrast at the twelve preregistered loci. Everything else is secondary and labelled as such.

S6.9 All four baselines of S5.17 shall be computed and read together, with the surrogate threshold calibrated by simulation on this design rather than taken as nominal.

S6.10 The structural ceiling test of S5.19 shall be applied to any claimed increment.

S6.11 The kill conditions K1 through K5 of S5.22 shall be declared and honoured.

S6.12 The result shall be reported in full whichever way it goes, with the design-space fraction it bounds computed per §6.11.

6.13 The two controls, and why neither substitutes for the other

This is the section that most often gets compressed into one line in a methods section, and the compression is where the error lives. **The scrambled-sequence control and the shuffled-input control are different controls,** they answer different questions, and running one does not weaken the case for the other.

```
SCRAMBLED-SEQUENCE CONTROL          a control on the STIMULUS
  what varies      the ordering of the symbols WITHIN one exposure
  what is held     energy, duration, histogram, long-term spectrum, level
  what it answers  does the response depend on the ORDER of the content?
  what it cannot   say anything about the timing of exposures across the series,
                   or about a third variable driving both exposure and endpoint

SHUFFLED-INPUT CONTROL              a control on the SERIES
  what varies      the temporal alignment of the input series to the endpoint
                   series, refit under the permutation, holdout permuted
                   independently, cost term NOT permuted
  what is held     the marginal distribution of the input
  what it answers  does the TIMING of the input carry information about the
                   endpoint?
  what it cannot   detect repackaging; detect a shared slow driver
```

A design with only the scrambled-sequence control can be fooled by an artefact of scheduling — an incubator door, a media change, a technician's routine — that happens to align with the ordered arm. A design with only the shuffled-input control can be fooled by an effect of sound that has nothing to do with content, which would be a real finding about acoustics and no finding at all about sequence.

And the shuffled-input control brings its own three measured failure modes, which S5.20 specifies in full and which are restated here because this is where a reader decides whether to believe a result:

(1) It is anti-conservative as ordinarily run. Across 500 null trials at a nominal `alpha` of 0.05, measured false-rejection rates were 0.108 for iid permutation (2.2x nominal), 0.120 for circular shift (2.4x nominal), and 0.060 for block permutation (near nominal). Calibration must be on the design in hand, with the calibrated threshold preregistered.

(2) It is blind to repackaging. A constructed filtered repackaging — a rolling average of the input and nothing else — passed all four baselines including the shuffle at `p = 0.0005`. It was caught by the structural ceiling test alone, with a single-step jump 38.5 times the theoretical maximum.

(3) It is blind to a shared slow driver. In a 4,000-trial simulation with input and endpoint both driven by a third slow variable, the shuffled-input control passed 98.35% of the time. The same simulation reports it working honestly in the clean case at 48.17% detection, which is what makes 98.35% a measurement rather than a complaint.

The consolidated rule: **more surrogates do not fix any of this. Design controls do** — blinding, channel separation, prospective sampling, and a yoked control — and every one of them is a decision about the instrument, made before a single number exists.

6.14 What would settle it, either way

A gap worth naming is a gap with stated settlement conditions on both sides. These are they.

What would establish a sequence-specific acoustic channel:

```
E1   The ORDERED versus SCRAMBLED contrast is significant at the preregistered
     loci, at the calibrated threshold, with an effect exceeding the declared
     SESOI and the detectable floor of §6.9.
E2   The result holds with the measurer blind and with channel separation in place.
E3   The yoked control does not move.
E4   The structural ceiling test passes.
E5   The effect is dose-ordered in a preregistered direction -- more exposure,
     more effect -- rather than present or absent.
E6   INDEPENDENT REPLICATION by a group that did not build the apparatus, on
     their own cells, from the published specification and payload hash.
```

E6 is not a courtesy and it is not last because it is least. Independent replication is the step at which a published literature's effect sizes have repeatedly proved to shrink [Open Science Collaboration, 2015, Ioannidis, 2005]. A single laboratory reporting a sequence-specific acoustic effect has produced a reason for somebody else to look, and nothing more.

What would close the question negatively for a stated region of the design space:

```
N1   The ORDERED versus SCRAMBLED contrast is null, at an `n` sufficient to detect
     the declared SESOI at the declared power, with the floor of §6.9 published
     alongside.
N2   The condition tested is stated, and the fraction of the design space it bounds
     is computed per §6.11.
N3   The null is reported in full, with the same prominence a positive result would
     have received, and deposited so that the next group does not repeat it.
```

A null bounds a region; it does not close a question. Nobody should write "acoustic sequence-specificity does not exist" from any experiment described here, and this section would be the first thing to cite against them if they did.

One more thing about settlement, which is easy to lose in a list of criteria. A research programme is settled not by a single decisive experiment but by a body of results that make one account much easier to hold than the others, and the criteria above are written to produce results that can accumulate. That is why the design-space fraction is required with every report, why nulls are deposited rather than discarded, and why the payload hash identifies a run. Twenty null conditions, each bounding a stated region, each reported with its detectable floor, are worth far more to the next group than twenty conditions of which only the interesting one was written up — and the second collection is what an ordinary publication incentive produces if nobody writes the rule down first.

The corollary is that a reader evaluating any future claim in this area has a short checklist. Was the permuted twin run? Was the measurer blind? Was the sample size sufficient to see the effect that was claimed, at the correction that was used? Was the null of the same experiment reported with the same prominence? Four questions, all answerable from a methods section, and a claim that fails any of them has not yet reached the point where the mechanism question becomes interesting.

6.15 What must not be done with a positive result

Included because the failure is predictable, and because the constraints of §4 do not relax if something turns out to work.

S6.13 A positive *in vitro* result licenses further *in vitro* work. It does not license any application to a person, in any setting, for any purpose.

S6.14 No result of any size converts this procedure into a treatment, therapy or practice. The labelling rule of §5.2 applies unchanged after a positive result and applies with more force, not less.

S6.15 The no-subliminals constraint of §4 is not contingent on efficacy. A demonstrated effect would make concealment more harmful, not more permissible.

S6.16 No claim that DNA has been edited, altered or damaged may be made on the basis of a transcriptional result. Transcript abundance is not sequence.

The failure modes

FM-6.1 — Substituting an easier control. A "silent" arm or a "white noise" arm is run instead of the symbol-permuted twin, because a permuted twin sounds strange and is harder to explain.

Fix: S6.3. Silence and noise differ from the stimulus in energy and spectrum, so a difference against them is uninformative about ordering. **The permuted twin is the only control that isolates the variable the claim is about, and it costs one shuffle.**

FM-6.2 — Sonogenetics offered as the mechanism. Ibsen et al. and Duque et al. are cited as though they demonstrated that waveform structure carries content.

Fix: §6.6. Sonogenetic specificity resides in an engineered protein, not in the signal, and the result is evidence in the opposite direction: it shows how acoustic specificity has actually been obtained, which is by building it into the receiver.

FM-6.3 — Tajik et al. read as a sequence result. Force reaching the nucleus and altering transcription becomes, in the retelling, force carrying a message to a locus.

Fix: §6.6. The force in that work carried a magnitude and a direction, and the result is precisely what makes the missing link visible: the force pathway is established, and it is not a sequence pathway.

FM-6.4 — PIEZO as a demodulator. A mechanically gated channel is described as though it were frequency-selective in a way that could separate symbols.

Fix: §6.7, T1 versus T2. Even a perfectly selective transducer satisfies only T1. The gap is T2 and T3, and a better transducer does not touch either.

FM-6.5 — The underpowered null announced as a refutation. Three replicates, no effect, and a conclusion.

Fix: §6.9. At `n = 3` and `sigma = 0.35` the targeted floor is a 1.68-fold change. An effect below that was never in the experiment's reach.

FM-6.6 — The single condition announced as a channel. One frequency, one level, one duration, one cell type, one timepoint, and a positive result.

Fix: §6.11 and E1 through E6. One cell of the grid is 0.000801% of it, and E6 is not optional.

FM-6.7 — Post-hoc SESOI. The smallest effect of interest is set after the effect size is known.

Fix: S6.7, timestamped and deposited before exposure.

FM-6.8 — Reporting the ordered arm against the silent arm. The comparison that gets reported is the one that reached significance, and it is the comparison that cannot answer the question.

Fix: S6.8. The primary analysis is preregistered as ORDERED versus SCRAMBLED, and any other contrast is secondary and labelled.

FM-6.13 — The archive described as a message. The regulatory function of non-coding sequence, and syncytin in particular, get retold as evidence that the genome is "written to" by outside signals.

Fix: §6.5b. Retroviral integration is a molecular event with a well-characterised mechanism and no acoustic content whatever. The archive being real is what makes the addressing question well-posed; it is not evidence about the channel.

FM-6.14 — A mechanism proposed and then treated as found. Somebody offers a candidate — a chemical memory element, a bursting-history effect, a cytoskeletal route — and within two retellings the candidate has become the explanation.

Fix: §6.2. A candidate mechanism with no independent evidence is a hypothesis, and naming a hypothesis does not move it into the established column. **A proposed mechanism and an established one differ by exactly the amount of evidence between them, which in this case is all of it.**

FM-6.9 — Treating the gap as a formality. The section is read as a legal disclaimer and the procedure is described elsewhere as though the disclaimer were pro forma.

Fix: §6.3. The gap is the reason the rest of the framework is worth reading, and work that treated it as a formality would have spent the only asset it had.

FM-6.10 — "Not proven impossible" doing evidential work. The sentence "no mechanism is known" becomes, over a few retellings, "science hasn't caught up yet."

Fix: §6.2. **Absence of a mechanism is not support of any kind, and a claim with no mechanism and no demonstration is not asserted rather than lightly qualified.**

FM-6.11 — The bulk-strain number used as an argument against everything. The `1e−10` figure is produced to argue that no acoustic effect on cells is possible at all.

Fix: §6.10. The number closes the bulk compressional route and says nothing about membrane tension, cytoskeletal transmission, or channel-level forces, which are the established pathways.

FM-6.12 — Merging the two controls. A methods section says "permutation control" and it is impossible to tell which one was run.

Fix: §6.13. Name them separately, report them separately, and state what each answers.

Worked example

CONSTRUCTED. No experiment has been run and no data exist. What follows is a fully specified design, written out so that both of its possible outcomes are visible in advance.

6.16 The design

```
STIMULUS        §5.25, payload hash
                c7d98b1cfde5c515ce99ff7db77d484f71a09c5cd921994f15df66c50373b72d
                5,640 bases, 20 ms/symbol, windowed envelope, 175.800 s per run
DELIVERY        characterised transducer; SPL measured at the sample position
SYSTEM          named cultured line, named passage range, in vitro
ARMS            ORDERED / SCRAMBLED / SILENT / YOKED, randomised, measurer blind
ENDPOINT        log2 transcript abundance at the 12 preregistered addressed loci
SIGMA (assumed) 0.35 log2 units, from the operator's own pilot variance,
                NOT from this framework
ALPHA           0.05 / 12 = 0.0041667 (Bonferroni across the 12 loci)
POWER           0.80
n               20 per arm
DETECTABLE      0.4103 log2 units = 1.3289-fold          (§6.9, computed)
SESOI           declared at 0.50 log2 units = 1.4142-fold, which EXCEEDS the
                detectable floor at n = 20, as S6.7 requires
PRIMARY         ORDERED vs SCRAMBLED at the 12 loci
SECONDARY       ORDERED vs SILENT; genome-wide exploratory, labelled as such
CONTROLS        blinding, channel separation, prospective schedule, yoked arm
SURROGATE       block permutation, threshold calibrated by simulation on THIS
                design; the calibrated threshold preregistered, not the nominal
KILL            K1 - K5 of S5.22
```

Note what the sample-size line costs. At `n = 20` per arm and four arms the experiment is 80 wells per timepoint before replication of the run itself, and the detectable floor is still a 1.33-fold change. An experimenter who wants to see a 1.2-fold effect at these settings needs roughly 100 per arm, and the honest time to discover that is before the first plate is seeded.

6.17 Both outcomes, written in advance **Outcome A — the contrast is significant.**

```
reported        effect size with interval, at each of the 12 loci
required next   E2 blinding intact, E3 yoked arm null, E4 ceiling test passed,
                E5 dose-ordering in the preregistered direction,
                E6 independent replication from the published specification
permitted claim "in this system, at this condition, the transcriptional response
                distinguished an ordered stimulus from its symbol-permuted twin"
FORBIDDEN claim any therapeutic, clinical or capability claim; any claim about a
                person; any claim that DNA was edited; any claim that the
                mechanism is known -- a demonstrated dependence is not a mechanism
design space    one cell of 124,800; 0.000801% bounded
```

Outcome B — the contrast is null.

```
reported        the null, in full, with the same prominence as a positive result
required with   the detectable floor (0.4103 log2 at n = 20), the declared SESOI,
                and the fraction of the design space bounded
permitted claim "in this system, at this condition, with 80% power to detect a
                1.33-fold change, no dependence on ordering was detected"
FORBIDDEN claim that acoustic sequence-specificity does not exist; that the
                question is closed; that the architecture is refuted
design space    one cell of 124,800; 0.000801% bounded
```

The two outcomes bound the same fraction of the design space and are worth reporting with the same prominence, which is the single most reliable way to tell whether a research programme is one.

6.18 The gap, at the end The specification above is complete enough to run. The alphabet is separable, the payload is hashed, the controls are named, the sample size is computed, the floor is published, and both outcomes are written down before either has occurred.

And the sentence at the head of this section is exactly as true at the end of it as it was at the beginning. **There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal.** Nothing in this framework changes that, nothing in this framework claims to, and the three parts of this paper were written in the order they were so that no reader could reach the procedure without meeting the sentence twice on the way.

A last remark on why this section is where it is. The paper could have ended at §5 with a complete, verified, disclosed procedure and a reader's impression that something had been achieved. Something has been achieved — an encoder that round-trips, an alphabet that separates, a rendering that a listener can follow and refuse — and none of it is the thing the paper is named for. Putting this section after the procedure rather than before it means a reader meets the specification while it still looks promising and then meets the gap while the specification is fresh, which is the order in which the gap is hardest to forget.

What this section established **Established, and checkable:** that the open problem decomposes into three requirements of which one (symbol separation) is solved and two (order memory, mapping to a locus) have no identified mechanism; that eleven upstream components are established and one link is missing, and that the missing link is specifically the carriage of sequence over force; that the symbol-permuted twin is exact on every permutation-invariant quantity and is drawn from a set of `10^2017.717` orderings for this framework's payload; that at `sigma = 0.35` and `n = 3` a targeted design cannot see below a 1.68-fold change, and at `n = 20` cannot see below 1.33-fold; that bulk compressional strain at plausible listening levels is of order `1e-10` and is not the route; that a single condition bounds 0.000801% of a 124,800-cell design grid; and that the scrambled-sequence and shuffled-input controls answer different questions, with the second carrying three measured failure modes of which two are answerable only by design controls.

Not established, and not claimed: that any acoustic method affects any cell's expression; that any capability has been conferred; that DNA is edited, altered or damaged; that the procedure of §5 is a treatment, therapy or practice; that any frequency in this framework is a physical resonance; or that the absence of a known mechanism is evidence of anything whatever, in either direction.

The one line to carry forward: *the channel for sequence-specific content is the question, and naming it precisely is the most useful thing this framework can do about it.*

Section 6 owns the open problem, the three-requirement decomposition, the detectability arithmetic, the design-space bound, and the two-control specification.

7 Results

The eight results below are the units other papers in this corpus cite. Each carries its status: **DERIVED** where it follows from the stated definitions by algebra, **COMPUTED** where it is the output of exact arithmetic on stated parameters, **MEASURED** where a stated quantity of a constructed artefact was measured, and **POSTULATED** where it is asserted by this framework.

R1 — the open problem, stated formally. Let $s(t)$ be an acoustic pressure waveform encoding a symbol string $x = (x_1, \ldots, x_L)$ over an alphabet of size four, and let R be a cell's transcriptional response. **There is no established mechanism producing an R that depends on x as an *ordered string*, as distinct from depending on functionals of $s(t)$ that are invariant under permutation of the symbols** — total energy, duration, long-term power spectrum, envelope statistics, symbol counts. *Status:* a statement about the literature. It is **not** a claim of impossibility, and it is **not** a licence: absence of a known mechanism is not support of any kind.

R2 — the index-assignment result. Assigning $n = 1, 2, 3, 4$ to the four symbols yields odd parts $1, 1, 3, 1$, so **three of four collapse onto a single pitch class**, because $f(2) = 4f(1)$ is exactly two octaves and $f(4) = 16f(1)$ exactly four. The odd assignment $n = 1, 3, 5, 7$ yields four distinct classes; transposed at $j = 8$ the tones are 123.7570, 1113.8129, 3093.9247 and 6064.0924 Hz. **The closest pair in hertz is A and C at 990.055901 Hz; the closest pair as an interval is G and T at 1165.0244 cents. Those are different pairs, and the source's display line joins them with an equals sign (§10).** *Status:* DERIVED and COMPUTED. The naive assignment is rejected in print so that the rejection is findable.

R3 — determination by choice. A directive corpus closed under negation has resultant exactly 0.000000000000 at $N = 4, 8, 12, 24$; a mutually exclusive corpus without exact pairing has mean direction $1/\sqrt{N}$, simulated at 0.498500, 0.350953, 0.288101, 0.202825, 0.140809 and 0.099738 against theory. At $u = 1$ and $N = 12$ the corpus carries $F = \ln 12 = 2.484907$ nats, **exactly the selection cost of 3.584963 bits.** *Status:* DERIVED and COMPUTED, with the equality holding **only at $u = 1$ and under a uniform receiver prior.**

R4 — the Gabor constraints. $\Delta t \, \Delta f \geq 1/(4\pi) = 0.0795774715$ [Gabor, 1946] gives $\Delta t \geq 0.080377$ ms at the symbol layer's 990.055901 Hz gap and $\Delta t \geq 10.2882$ ms at the state layer's 7.734812 Hz binding gap. **The state layer is nearly 128 times more demanding**, which is why the two layers are separated in time. *Status:* CITED as a theorem, COMPUTED in application.

R5 — single-transposition audible capacity. A single integer octave transposition admits at most $N = 25$ essences inside 20 Hz to 20 kHz, and only at $j = 6$; at $N = 26$ the set of admissible j is empty. *Status:* COMPUTED and verified.

R6 — bulk compression closed by arithmetic. For a plane wave in an aqueous medium the longitudinal strain is $s = p/K$ with $K = 2.2 \times 10^9$ Pa; at 40 to 120 dB SPL this is 9.090909×10^{-13} to 9.090909×10^{-9}. **Bulk compression is not the route.** *Status:* COMPUTED from elementary continuum mechanics. **The number says nothing about membrane tension, cytoskeletal transmission or channel-level forces**, which are the established pathways, and it must not be used as a general argument against acoustic effects.

R7 — detectability floors. At $\sigma = 0.35$ $\log_2$ units, power 0.80, and Bonferroni correction, a targeted twelve-locus design **cannot see below a 1.68-fold change at** $n = 3$ **or a 1.33-fold change at** $n = 20$; genome-wide at 20,000 transcripts the corresponding floors are 2.18-fold and 1.53-fold. *Status:* COMPUTED on declared inputs. σ is the operator's, not this paper's.

R8 — how little one condition bounds. A deliberately modest grid over carrier band, level, symbol rate, exposure duration, cell type and readout time gives 124,800 cells, so **one condition bounds** 0.000801 **per cent of it.** *Status:* COMPUTED. The framework's survey work supplies the matching empirical caution [Curzi, 2026i]: a shuffled-input control passed 98.35 per cent of the time against a shared slow driver while detecting honestly at 48.17 per cent in the clean case.

CR — counter-result. The entire pipeline is built, verified and deterministic, and **nobody knows how a cell could read a message out of a sound.** Eleven upstream components are established and the one missing link is precisely the carriage of sequence over force. The three-way decomposition sharpens it: symbol separation is solved and is the least interesting of the three; order memory and mapping to a locus have no identified mechanism, and a better transducer does not touch either. Separately, and in the same spirit, **a careful author writing a deliberately balanced corpus failed a mechanical equal-force check on four of twelve pairs, 33.33 per cent**, and the failures are named rather than repaired.

8 Discussion

8.1 What follows

Disclosure and contradiction are two protections that fail in different circumstances, which is the reason for having both. A single-directive corpus satisfies every disclosure clause and is still a command, because a corpus with one member has a resultant equal to its member and a surplus of $F = \ln 1 = 0$. A contradictory corpus delivered inaudibly satisfies the structural condition and violates the disclosure condition totally, because a listener who cannot hear the corpus cannot perform the selection the structure exists to leave them. **Disclosure makes a command refusable; contradiction makes it not a command.**

The index-assignment result generalises past this application. Any scheme that maps symbols to frequencies through a power law inherits the collapse: indices related by a factor of two are related by an exact octave, and pitch class cannot distinguish them. That is a design constraint on sonification generally and it is cheap to check.

The corpus's own honesty apparatus is the part most likely to be dropped and the part worth keeping. The equal-force check is a crude proxy — character and word count — and is not defended as more than that. It catches gross asymmetry, which is the failure that matters, and it is checkable by anyone. Prosody, connotation and rhetorical weight are not caught by it, which is why the rendering order is fixed in advance and why a witness who is not the author is worth having.

A closed route is a result. The strain arithmetic of §6.10 removes the route most often assumed and leaves the established pathways untouched. That is the shape of progress available here: not a demonstration, but a narrowing.

8.2 What does not follow

Nothing here is evidence that the procedure does anything. The open-label placebo literature establishes that disclosure and a belief-mediated effect are compatible *in another domain.* It is evidence about placebo trials, not about sonified sequence, and the two-column rule forbids moving it across [Curzi, 2026w].

A separable alphabet is not progress on the question. Engineering effort flows naturally toward symbol separation because it is where progress is visible, and symbol separation is the one requirement that was never the difficulty.

Sonogenetics does not close the gap and points the other way. Acoustic specificity in biology has so far been obtained by putting the specificity in the *receiver* — an engineered protein — and never by putting it in the signal.

A null does not close the question and a positive result does not open a channel. Both bound 0.000801 per cent of a modest grid, and **the two outcomes are worth reporting with the same prominence**, which is the single most reliable way to tell whether a research programme is one.

8.3 What the wider corpus says this paper may not claim

Two papers in this corpus bound what a complete specification establishes, and they are cited here because they cut against the reading a sympathetic reader would give.

The framework's historical-network paper demolishes its own most attractive material in print: a well-known story is refused as a fabricated data point, a unified-hand reading is rejected at 1,365:1, and a 25 per cent figure is badged as an artefact of a chosen synthetic ratio rather than a constant, with a confusable-only check reaching only 58.99 per cent [Curzi, 2026]. And the content-addressing paper states that its own headline 56.62 per cent widening **is a ceiling on constructed inputs and not a promise**, distinguishes visibility from pledgeability explicitly, and names the valuation of non-financial capital as unsolved — presenting an architecture as an architecture and not as a system [Curzi, 2026].

Both say the same thing to this paper. A complete specification on constructed inputs is a specification. It bounds nothing about the world, and the temptation to read a verified round-trip hash as evidence of anything biological is exactly the temptation those two papers refused in their own domains.

One further bound is internal to the architecture. The framework's addressing work states the constraint plainly [Curzi, 2026y]: **no addressing operation, at any signal strength or in any modality, can cause a cell to express content its genome and accessible regulatory repertoire do not already contain.** Even a working channel would be a channel for *selection* among latent material, not for creation, and every claim in this paper is bounded by that in advance.

8.4 On arrival, and how far this is from settled

The framework's convergence work supplies the right gauge for a research programme in this state [Curzi, 2026]: the a posteriori bound $d(x_n, x^*) \le (q^n/(1-q))\, d(x_0, x_1)$, read operationally as *the distance to arrival is the size of the last correction, scaled.* Applied here the reading is uncomfortable and correct: **the last correction this material received was the statement that the central mechanism does not exist**, which is not a small correction, and the distance to arrival is correspondingly large. A programme whose most recent revision is of that size is not close to a fixed point, and saying so is more useful than any estimate of how close it might be.

9 Verification

Every step below carries one of four statuses — **ARITHMETIC** (reproducible with a calculator from the inputs named), **DEFINITION** (true by stipulation), **CITED** (an external source supports it, named), or **POSTULATED** (asserted by this framework, no external support).

9.1 The postulates

	postulate	status
P1	φ is the positive root of $x^2 = x + 1$.	CITED [Euclid, 1956, Vajda, 1989]
P2	The Pythagorean whole tone is the ratio $9{:}8$.	CITED [Barbour, 1951]
P3	A dimensionless ratio $(9/8)^2/\varphi^2$ **is a frequency in hertz**.	POSTULATED — no source
P4	$f(n) = kn^2$ for an index n.	POSTULATED — no source
P5	Alphabetical order is an admissible ordinal on the four bases.	POSTULATED (a stated convention)
P6	$\Delta t\,\Delta f \geq 1/(4\pi)$ for any finite-energy signal.	CITED [Gabor, 1946]; proof reproduced in Curzi [2026z]
P7	$N_{\text{eff}} = 1 + (N-1)u$ and $F = \ln N_{\text{eff}}$.	POSTULATED (design axiom); **not a Hill number**
P8	Directives are unit vectors and negation is $-v$.	DEFINITION (a stated model)
P9	Bulk strain in a fluid is $s = p/K$.	CITED as elementary continuum mechanics
P10	The two-sample normal power formula, at declared α, β, σ.	DEFINITION (a stated statistical model)
P11	A permutation of a symbol string leaves every permutation-invariant functional of its rendering unchanged.	DEFINITION

Dropped as unnecessary. No postulate about the *ordering* or *meaning* of the twelve state classes is needed for R2, R4, R5, R6, R7 or R8. No postulate about consonance is needed anywhere. **P3 and P4 are needed only for the frequency *values*; the separation results survive their deletion in ratio form**, because every ratio in the alphabet is a ratio of squares and k cancels.

9.2 The compounding chain

step	consumes	produces	status
L1	P1, P2	$(9/8)^2 = 1.265625$; $\varphi^2 = \varphi + 1 = 2.618034$	ARITHMETIC
L2	L1	$k = 0.483425732988414$	ARITHMETIC (exact)
L3	L2, **P3**	k carries the unit Hz	**POSTULATED**
L4	L3, **P4**, **P5**	four symbols map to $f(n) = kn^2$ at a chosen index set	**POSTULATED**
L5	P4 by inspection	$f(2n)/f(n) = 4$; the pitch class of n is its odd part	ARITHMETIC

step	consumes	produces	status
L6	L5 at $n = 1, 2, 3, 4$	odd parts $1, 1, 3, 1$: three symbols collapse — **R2**	ARITHMETIC
L7	L5 at $n = 1, 3, 5, 7$	four classes; at $j = 8$, tones from 123.76 to 6064.09 Hz	ARITHMETIC
L8	L7	closest in Hz: A–C, 990.0559; closest as an interval: G–T, 1165.0244 cents; **different pairs** — **R2**	ARITHMETIC
L9	P6, L8	at the symbol layer, $\Delta t \geq 0.080377$ ms — **R4**	ARITHMETIC
L10	P6, folded gap 7.734812 Hz	$\Delta t \geq 10.2882$ ms; ratio ≈ 128 — **R4**	ARITHMETIC
L11	L7, band 20 Hz–20 kHz	admissible j nonempty to $N = 25$, empty at 26 — **R5**	ARITHMETIC
L12	P8, closure under negation	$\lVert \sum_i v_i \rVert = 0$ exactly — **R3**	ARITHMETIC
L13	P8, isotropy	mean direction $1/\sqrt{N}$; simulated agreement to $\sim 10^{-3}$ — **R3**	ARITHMETIC + simulation
L14	P7 at $u = 1$	$F = \ln 12 = 2.484907$ nats $= 3.584963$ bits — **R3**	ARITHMETIC given P7
L15	L14, ceiling	deficit at $c/\eta = \ln 12$ is 0, residual 8.88×10^{-16}	ARITHMETIC
L16	ASCII, SHA-256	675 bytes $\to$ 5,400 bits; hash `c7d98b1c` …	ARITHMETIC
L17	L16, balanced rule	GC $= 0.500000$ exactly; max run 1; entropy 1.935874 against greedy 1.244171	MEASURED on the payload
L18	L17, decode	recovered hash identical — round trip verified	MEASURED
L19	4^{20} vs 6.2×10^9	0.005638867 expected spurious matches; ratio 177.3406	ARITHMETIC
L20	P9	$s = p/K$: 9.09×10^{-13} at 40 dB to 9.09×10^{-9} at 120 dB — **R6**	ARITHMETIC
L21	P10	detectable $\log_2$ FC 1.0593 at $n = 3$, 0.4103 at $n = 20$ — **R7**	ARITHMETIC
L22	grid of stated choices	$30 \cdot 13 \cdot 4 \cdot 4 \cdot 5 \cdot 4 = 124{,}800$ cells, so one condition bounds 0.000801% — **R8**	ARITHMETIC

step	consumes	produces	status
L23	P11	the permuted twin is exact on energy, duration, histogram and long-term spectrum — **R1**	DEFINITION
L24	multinomial on the payload composition	$L!/\prod_i n_i!$, which is $10^{2017.717}$ orderings	ARITHMETIC

9.3 What this chain is and is not

Read the status column. L1, L2 and L5–L24 are arithmetic, definition, or measurement on a constructed artefact, and a reader can redo every one of them. **L3 and L4 are postulated**, and they are the only postulated steps in the chain.

The separation is unusually clean here and is worth stating precisely. A reader who rejects P3 and P4 — who declines to accept that a dimensionless ratio is a frequency or that the essence law holds — loses the specific hertz values of the alphabet and keeps everything else: the collapse result in ratio form (the tones related by $4{:}1$ are the same pitch class whatever the reference is), the Gabor arithmetic, the transposition capacity, the corpus geometry, the strain calculation, the detectability floors and the design-space bound. **Seven of the eight results survive the deletion of both postulates.**

And the chain does not reach the paper's subject. No step in L1–L24 produces a biological claim, and none could: **R1 is a statement that a step which would be needed does not exist.** A verification section that decomposed a claim about cells into checkable pieces would be a verification of something this paper does not assert.

Two statuses that must not be misread. The exponent 2 in $f(n) = kn^2$ and the exponent $-1/2$ in the resultant's decay are both **DEFINED** by the model, not estimated; no fitting was performed and the standard procedure for testing a claimed scaling relation [Clauset et al., 2009] was not run, because there is nothing to run it on. And the simulation at L13 is a simulation of a stated random model, not an observation.

The verification instance. The smallest example that exercises L5–L8 is the four-symbol alphabet itself: $k \cdot 1 = 0.4834257330$, $k \cdot 9 = 4.3508315969$, $k \cdot 25 = 12.0856433247$, $k \cdot 49 = 23.6878609164$; multiply each by $2^8 = 256$; subtract the two smallest to get 990.0559 Hz; and $1200 \log_2(1113.8129/123.7570) = 1165.0244$ cents. Four multiplications, one subtraction, one logarithm.

10 Limitations

L-1. The dimensional audit does not close on the essence law. $f(n) = kn^2$ balances only if k carries $[T^{-1}]$, and no operational procedure is supplied by which it does (§2.3 P-a, §3.2). The Gabor bound and the strain calculation *do* close, and the contrast is the point.

L-2. The essence law has no external support, and the exponent is defined rather than estimated (§2.3 P-b). **This paper makes no empirical scaling claim.**

L-3. The constant is chosen. $k = (9/8)^2/\varphi^2$ has no precedent as a frequency, and the deflationary literature on golden-section claims is cited rather than avoided [Markowsky, 1992]. Its fifteen figures are exact digits of a definition.

L-4. Octave reduction beyond pitch is unwarranted. Shepard grounds octave equivalence for pitch [Shepard, 1982]; nothing grounds carrying it to a quantity that is not a pitch.

L-5. $N_{\text{eff}} = 1 + (N-1)u$ **is not a Hill number.** It is a linear interpolation between 1 and N. The Hill results [Hill, 1973, Jost, 2006] apply to $1/\text{HHI}$ and $\exp(H)$ and to nothing here.

L-6. Conjoint measurement was not tested, and scalar reduction is the general form of the problem. Michell's objection [Michell, 1997] stands unanswered for u and for the elicited intensities, and the remedies [Luce and Tukey, 1964, Rasch, 1960] were not run. **The same objection arrives from the affect literature in a different domain** — individuation by pattern across appraisal dimensions [Smith and Ellsworth, 1985], by synchronisation in time [Scherer, 2001], and by qualitative core relational theme [Lazarus, 1991], with action readiness constitutive rather than terminal [Frijda, 1986], against two well-supported dimensions of core affect [Russell, 1980]. **These are one objection, not two**, and §3.5 names it once. The twelve class names of §4.14 are motivated by that literature and are **not licensed by it**, and the ladder's ordering in particular is contradicted by it.

L-7. The alphabetical ordinal is weak and is stated as weak. Applying a square law to symbols requires an ordinal, and alphabetical order is a convention rather than a property of the symbols. **Any other published ordinal would give different frequencies and the same separation structure**, which is a way of saying that the specific hertz values carry no information the ordering does not.

L-8. The pitch-class collisions are structural and are not repaired. Four of twelve directives share one class and three share another; six positions carry twelve directives. That is $(2n)^2 = 4n^2$ and nothing else, and absolute register — not pitch class — is what identifies a state.

L-9. The biological and placebo literature has no cluster in this corpus's bibliography, and every such work is named without a citation key. Specifically: Church, Gao and Kosuri (*Science* 337:1628, 2012); Goldman et al. (*Nature* 494:77, 2013); Erlich and Zielinski (*Science* 355:950, 2017); Zetsche et al. (*Cell* 163:759, 2015); Fonfara et al. (*Nature* 532:517, 2016); Qi et al. (*Cell* 152:1173, 2013); Gilbert et al. (*Cell* 159:647, 2014); Tajik et al. (*Nature Materials* 15:1287, 2016); Ibsen et al. (*Nature Communications* 6:8264, 2015); Duque et al. (2022); Takahashi and Yamanaka (*Cell* 126:663, 2006); the PIEZO1/PIEZO2 mechanotransduction literature; Kaptchuk et al. (*PLoS ONE*, 2010); Carvalho et al. (*Pain*, 2016); and von Wernsdorff et al. (*Scientific Reports*, 2021). **Every claim in §4.4, §5.8, §5.10, §6.5 and §6.5b that rests on one of these is therefore uncited in the reference list**, and is described from the source with its bibliographic details in the text. Attaching an unrelated key to a biological claim would be worse than the gap. **This is a defect of the corpus's bibliography, not a claim that the works do not exist**, and it must be closed before submission.

L-10. Sequence-sonification prior art is likewise uncited. The encoding step here is conventional rather than novel, and published sequence-sonification methods exist; the bibliography carries none of them, so **the novelty claim this paper makes is deliberately narrow** — the index assignment and the separation arithmetic, not the idea of sonifying a sequence.

L-11. One instance is not a rate. The equal-force failure rate of 33.33 per cent is one author writing one corpus once. It is reported because the *direction* of the finding is informative — a careful author intending balance did not achieve it — and it estimates nothing.

L-12. Character count is a crude proxy for force. It catches gross asymmetry, which is the failure that matters, and it does not catch prosody, connotation or rhetorical weight. That is why the rendering order is fixed in advance and why an independent witness is specified.

L-13. $u = 1$ here is constructed, not measured. A corpus of directives with overlapping support would not reach it; the value is a property of the design and the 69.6133 Hz that follows is the arithmetic ceiling of a twelve-category system, not a reading of a person, a text, or a relationship.

L-14. The one-hot degeneracy is used deliberately here and is a trap elsewhere. Vectors with disjoint support are automatically orthogonal. This paper's corpus is that degenerate case on purpose, and every reading of it says so.

L-15. No perceptual experiment stands behind any perceptual claim in this paper. The separation figures — 1165.0244 cents in the alphabet, 995.9922 cents between layers, 34.9756 cents in the folded lattice — are arithmetic on frequencies. The claim that a listener can *follow* the resulting rendering, that the double octaves are *audible without instruction*, and that the two layers are *separable by ear*, are perceptual claims. **No listener was asked anything.** They rest on the general acoustics and psychoacoustics literature [Rossing et al., 2002, Sethares,

2005, Burns, 1999, Hoeschele et al., 2012, Demany et al., 2013] and are limited by its own boundaries. This is stated again in §12.

L-16. The power arithmetic is not a claim about biology. σ and the number of tests are declared inputs. Substituting an optimistic σ is the cheapest way to make an underpowered design look adequate on paper, and the power target of 0.80 is a convention rather than a law; 0.90 costs roughly a third again in sample size and should be priced explicitly.

L-17. The bulk-strain result must not be over-read. It closes the naive route and says nothing about membrane tension, cytoskeletal transmission or channel-level forces. **Using it as an argument against acoustic effects in general would be a misuse of this paper.**

L-18. The design grid is arbitrary and deliberately modest. 124,800 cells follow from six stated choices at stated resolutions. A finer grid would make one condition bound even less, and a coarser one would flatter a null. The figure is a device for keeping both outcomes the same size, not an estimate of anything.

L-19. No model comparison and no statistical claim. Nothing was fitted, no out-of-sample test was run, and there are no p-values, effect sizes or confidence intervals in this paper, because there is no sample. Where the corpus has attempted a confirmatory fit elsewhere it failed its own baseline test [Curzi, 2026i], and that is the relevant precedent.

L-20. The reflexive specification is not identified, and the framework does not exempt its own parameters from its own reflexivity [Manski, 1993, Shalizi and Thomas, 2011, Lucas, 1976, Hacking, 1995]. The disclosure design makes the expectancy channel visible rather than absent [McCambridge et al., 2014].

L-21. The precedent for contradictory corpora is genuine and is not evidence. Proverbs 26:4–5, Talmudic *machloket*, and the koan tradition are named in §4 as independent arrival. **None is in this paper's bibliography and none is cited with a key.** Independent arrival is an argument that a structure is findable, not that it does anything [Curzi, 2026w].

L-22. Boundary behaviour. $u \to 0$ gives $F = 0$ and $f = k$, the ground essence, which is the reading for *no relationship* rather than a low one; $u \to 1$ gives the ceiling and is where the corpus is constructed to sit; $N \to \infty$ raises the ceiling without bound and is unavailable, since single transposition fails at $N = 26$ and the folded lattice collapses at $N = 17$; $\Delta t \to$ the Gabor floor produces **concealment achieved by speed**, which is why the symbol rate is set by audibility and not by the mathematics. These limits constrain the domain and are stated because the source states only some of them.

A discrepancy found in the source, printed rather than silently corrected

The source's symbol-alphabet block prints

```
MINIMUM GAP                990.055901 Hz  =  1165.0244 cents
```

and those two figures are not two readings of one quantity. The three adjacent gaps in the transposed alphabet are 990.055901 Hz (A to C), 1980.111802 Hz (C to G) and 2970.167703 Hz (G to T); in cents they are 3803.9100, 1768.7174 and 1165.0244. **The minimum in hertz is the A–C pair and the minimum in cents is the G–T pair**, because the alphabet is spread multiplicatively and the two orderings run opposite ways. The equals sign in the source's display line joins two different pairs.

Neither figure is wrong and the arithmetic that consumes them is right. The Gabor floor requires a bandwidth in hertz, so $0.0795774715/990.055901 = 0.080377$ ms is the correct application and uses the correct number. The 1165.0244-cent figure is the correct statement of how close the two nearest tones are as an *interval*, which is the quantity a listener's frequency discrimination is described in. **This paper states them separately throughout and does not join them**, and the correction is printed here rather than applied quietly, because a reader recomputing from the source's line will find it does not close.

A related pair that does close, checked because it looks like the same defect and is not. The folded-lattice minimum at $N = 12$ is stated in the source both as 34.9756 cents (the ratio $50/49$)

and as 7.734812 Hz. Those are consistent: the two lattice positions are $k \cdot 25 \cdot 32 = 386.7406$ Hz and $k \cdot 49 \cdot 16 = 379.0058$ Hz, whose difference is $k \cdot 16 = 7.734811727814624$ Hz exactly and whose ratio is $50/49$, 34.975615 cents. **One is absolute and one is relative and they describe the same pair.** The Gabor calculation at the state layer uses the absolute figure, correctly.

A second, smaller note on the source. Its failure-mode lists are out of order in two places: one list runs …12, 14, 13, and another runs …8, 13, 14, 9, 10, 11, 12. The ordering is preserved as found, renumbered to this paper's sections. It is a defect of presentation, not of content.

11 Conclusion

What is established, and checkable. A written corpus maps losslessly to a 5,400-base nucleotide string with a verified SHA-256 round-trip identity. A balanced free-choice rule holds GC within one base of $L/2$ by construction, measured exactly 0.500000 here, with maximum homopolymer run 1 and symbol entropy 1.935874 bits per base against a greedy rule's 1.244171. **The naive symbol-index assignment** $n = 1, 2, 3, 4$ **collapses three of four symbols onto one pitch class by** $(2n)^2 = 4n^2$**, and the odd assignment** $n = 1, 3, 5, 7$ **gives four distinct classes with a minimum gap of** 990.055901 **Hz,** 1165.0244 **cents, and a Gabor floor of** 0.080377 **ms.** The state layer's binding separation gives $\Delta t \geq 10.2882$ ms, nearly 128 times more demanding, and the two disclosed layers clear each other by 995.9922 cents. **Single-transposition audible capacity is** $N \leq 25$ **and only at** $j = 6$**.** A corpus closed under negation has resultant exactly 0.000000000000 at $N = 4, 8, 12, 24$, and at $u = 1$ carries $F = \ln 12 = 2.484907$ nats — exactly the selection cost of 3.584963 bits — sitting on the sustainability ceiling with residual 8.88×10^{-16}. **Four of twelve directive pairs fail the equal-force check, 33.33 per cent, and are named rather than repaired.** Bulk compression is closed by arithmetic at strains of 10^{-11} to 10^{-10}. A targeted twelve-locus design at $\sigma = 0.35$ cannot see below 1.68-fold at $n = 3$ or 1.33-fold at $n = 20$. And **one null condition bounds** 0.000801 **per cent of a 124,800-cell design grid.**

What is proposed rather than shown. That a dimensionless ratio of a tuning interval to the square of the golden section is a frequency; that $f(n) = kn^2$; and that alphabetical order is an admissible ordinal on four symbols. **These are postulates with no external support**, and seven of this paper's eight results survive their deletion.

What is not claimed at all, and this is the paper's principal statement. There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal. The pipeline is built, verified, deterministic and disclosed, and it stops at a loudspeaker. Eleven upstream components are established — storage in sequence, addressing, self-processed multi-address arrays, the separability of targeting from editing, mechanotransduction, force-to-transcription coupling, engineered acoustic channels, a large structured regulatory archive, and the reorganisation of expression by regulatory input — and **the one missing link is precisely the carriage of sequence over force.** Decomposed, symbol separation is solved and is the least interesting of the three requirements; order memory and mapping to a locus have no identified mechanism, and a better transducer does not touch either. That statement is exactly as true at the end of this paper as at the beginning, **it is not a claim of impossibility, and it is not a licence either**: absence of a mechanism is not evidence of anything, in either direction.

12 Declarations

Data and code availability. No dataset was collected for this study and none exists. No experiment was run, no cells were exposed to anything, and no endpoint was measured. Every number in the paper is the output of exact arithmetic on parameters stated in the text, or a measurement of a constructed artefact — a byte string, its hash, and the derived base string — which is fully specified here and reproducible from the specification alone. The figures are generated by the Python scripts named in their captions and are available from the author on request.

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the analysis, the figures and the manuscript.

Ethics. No human-subjects research was conducted and no personal data were processed. No person was exposed to any rendering described here, and the disclosure package, consent specifications and stop procedures of §4 describe requirements on a procedure that has not been run. **This paper asserts perceptual claims — that a listener can follow the rendered symbol stream, that a double octave is audible without instruction, and that the two disclosed layers are separable by ear — without any perceptual experiment behind them.** Those claims rest entirely on the cited acoustics and psychoacoustics literature, are limited by that literature's own boundaries, and are listed among this paper's limitations (L-15). No listener was asked anything.

Biological and clinical disclaimer. No claim of any biological effect is made anywhere in this paper. Nothing described here is a treatment, therapy, dose, intervention, modality or practice. **No DNA is edited, altered, synthesised, delivered or damaged**: nothing in the pipeline is a molecule, the nucleotide string is a string in the computer-science sense, and no nuclease, template or Cas protein appears at any point. No capability is claimed to have been conferred on anything. No person, tissue, cell or molecule is claimed to resonate at any frequency printed here; in particular, **cytosine does not absorb at 1113.8129 Hz, does not resonate at it, and is not affected by it.**

Preregistration. No test in this paper was preregistered, because no test was run. §5 Part B and §6.12 set out the preregistration, kill conditions and reporting requirements that would govern a first run, and §6.17 writes both possible outcomes in advance. **The design has not been executed and its preregistration has not been deposited.**

Literature not in the reference list. The biological, placebo and sequence-sonification works this paper describes are named in the text with their bibliographic details and are absent from the reference list, because the bibliographic database governing this corpus carries no cluster for them. They are enumerated in §10, L-9 and L-10.

Chapter 49

Destiny as a Fixed Point: Conservation of Drive Under Cost Transfer, the Completeness of the Lever Set, and the Corpus's Register of Demotions

Abstract. We define the destination of an arrangement as the fixed point Q_inf = (eta*F(u) - c)/delta and establish its properties exactly. The update map is a contraction with Lipschitz constant 1 - delta, so the fixed point exists, is unique, and is reached from any starting stock; the approach is a geometric equality verified to a maximum residual of 2.06e-15 across a hundred passes. The deficit and the destiny are one quantity in two units, D = -2.885390(delta/eta)Q_inf, with maximum residual 8.88e-16 over twelve parameter sets, and D contains no delta - the same deficit is consistent with destinies differing by a factor of thirty. Theorem 5.1 establishes that exactly four arguments enter, u, N, eta and c, and there is no fifth. Theorem 4.5 proves the sum of drives is exactly conserved under any transfer of cost among any number of participants: an apparent improvement of +7.751093 for one party is matched by -0.599002 for another, the sum of drives holding at +0.643688229683 to every digit, with the disadvantaged party's stock not crossing zero until pass 10.850479. When decay rates differ the plain sum is no longer invariant and moves at 1/delta_P - 1/delta_R per unit exported. The closing section reports six constraining results - including that F(u) is exactly a threshold-corrected Fechner law with u_0 = 1/(N-1) and F(u_0) = ln2 = one bit exactly - and a register of four occasions on which the framework's own gates demoted its author's claims.

Keywords: fixed point; contraction; conservation of drive; cost transfer; lever completeness; Fechner law; demotion register; reserve theorem

Series: C. Frequency, Essence, and Mind **Paper:** CCP-049

1 Introduction

A framework that computes where an arrangement is going acquires, without asking for it, the vocabulary of fate. Once a storage loop has a fixed point, the fixed point has a name, and the name the practical literature supplies is *destiny* — a word used to mean an unalterable script, a personal entitlement, a reward for correct wanting, and, most often, nothing at all.

This paper takes the word over and makes it checkable. **In this framework the destination of an arrangement is $Q_\infty = (\eta F(u) - c)/\delta$, the fixed point of the contracting storage loop** [Curzi, 2026a], and every property claimed for it is either an arithmetic consequence of that definition or is badged as something else. The definition buys one thing the rhetorical use never supplies: a number with four arguments, which moves when an argument moves and at no other time.

Three results follow from that definition and are the paper's contribution. The first is a conservation identity. Under a transfer of cost among participants, the **sum of drives** $\sum_i b_i = \sum_i (\eta_i F_i - c_i)$ does not move at all — exactly, unconditionally, and for any number of participants with any decay rates whatever. That converts a standing prohibition of the framework's foundations [Curzi, 2026] from an unfalsifiable statement about intention into an arithmetic test that costs one column of a report and returns an answer on the first pass. The second is a completeness statement: Q_∞ has exactly four arguments and there is no fifth, so a proposal that cannot be assigned to one of them changes the destination by zero. The third is negative and is the sharpest limitation the framework's central instrument has: the octave deficit D contains no δ, so the same deficit is consistent with destinations differing by a factor of thirty.

Why the conservation result matters more than its algebra suggests. An intervention that improves one participant's readings genuinely, pass after pass, while the conserved sum does not move, has produced nothing; it has moved an existing quantity between two ledgers, and both ledgers are correct. Cost export is an externality in the ordinary sense [Hardin, 1968, Ostrom, 1990], and the case worked here shows the exporting party's instrument reading better at every step. The disadvantaged party's stock, meanwhile, stays positive for nearly eleven passes after its destination has gone negative, so the ordinary question — *is anyone worse off?* — returns a false negative for the whole of that window.

The paper closes on six results that constrain the framework rather than extend it, and on a register of four occasions on which the framework's own gates demoted its author's claims. **A corpus whose strongest results are all constraints is a corpus that was checked**; the alternative, a corpus of affordances each of which adds a capability and none of which subtracts one, is the signature of work that was never audited [Lakatos, 1970, Platt, 1964, Meehl, 1990].

Contributions

1. **Conservation of drive under cost transfer** (Theorem 4.5). The sum of drives is exactly invariant under any transfer of cost among any number of participants, verified at $+0.643688229683$ across a two-party export range and at $+1.268940530229$ across four transfers in a three-party arrangement, residual $0.000\text{e}{+00}$. The plain sum of fixed points inherits the invariance **only** when the participants share a decay rate; otherwise it moves at $1/\delta_P - 1/\delta_R$ per unit exported.
2. **The convergence law as an equality** (Theorem 4.2), not as Banach's bound: the gap closes by exactly $1 - \delta$ per pass at every distance, verified to a maximum residual of $2.0624\text{e}{-15}$ over a hundred passes.
3. **The deficit identity and the decay blindness** (Theorems 4.3 and 4.4): D and Q_∞ are one quantity in two units, verified to $8.8818\text{e}{-16}$ over twelve parameter sets, and D contains no δ, so $\partial D/\partial\delta = 0$ identically.
4. **Completeness of the lever set** (Theorem 5.1): four arguments, five handles, and the reconciliation between the two counts.
5. **The reserve theorem** (Theorem 5.3): a lever can be *absent* rather than expensive. Demonstrated on a lever carrying 4.80 times the elasticity of its rival and 4.174198 octaves of reserve against a 4.338420-octave requirement, falling short by 0.164222 even at an impossible conversion efficiency of one.
6. **The price of misordering**: eight periods of the wrong order cost twelve periods of permanent lag, thirteen periods below the starting stock, and 111.133601 stock-periods of cumulative loss, with identical terminal parameters and therefore an identical destination.

7. **The surplus law is exactly a threshold-corrected Fechner form** with $u_0 = 1/(N-1)$, at which the affine effective count is 2.000000 exactly for every N and $F(u_0) = \ln 2$ — one bit, two octaves.
8. **The demotion register**: four occasions on which the framework's own gates rejected its author's claims, printed rather than summarised.

What this paper does not claim. No arrangement's parameters are claimed to hold. No larger fixed point is claimed to be a better life. No frequency printed anywhere here is claimed to be a physical resonance: the constant $k = (9/8)^2/\varphi^2 = 0.483425732988414$ is a **defined** number whose components are real — the Pythagorean whole tone [Barbour, 1951] and Euclid's extreme and mean ratio [Euclid, 1956, Vajda, 1989] — and whose product, expressed as a frequency, has no precedent in any literature; the standard deflation of golden-section claims [Markowsky, 1992, Fischler, 1981, Höge, 1997] is cited here and again wherever φ appears. Every case in this paper is **constructed** and is labelled as constructed at the point of use.

2 Hypotheses, definitions, and postulates

Most of this paper is not hypothesis-testing, and saying so is the first honest act available to it. Theorems 4.1 to 4.5, Corollary 4.6 and Theorems 5.1 to 5.3 are arithmetic on stated definitions; no observation could bear on them, and none was collected. What follows separates the few claims that could be wrong from the many that are definitions, postulates or modelling choices.

2.1 Hypotheses proper

H1 — the detection-lag hypothesis. When cost is exported from one participant to another, the disadvantaged party's *destination* turns negative on the pass the transfer happens, while its observable *stock* continues to fall through zero only after $t^* = \ln(|Q_\infty|/|Q_0 - Q_\infty|)/\ln(1-\delta)$ passes. **Predicted:** in the worked constructed case, $t^* = 10.850479$ passes at $\delta = 0.09$ — eleven periods in which nothing observable changes. **Falsified by:** a measured arrangement in which the stock and the drive change sign on the same pass, with both quantities separately instrumented. **Status: untested. No such measurement exists and none was made for this paper.** The arithmetic of the lag is a theorem; the claim that real arrangements exhibit it is not.

H2 — the ceiling hypothesis, inherited. No arrangement with $c/\eta > \ln N$ sustains, at any value of the geometry. **Predicted:** at $\eta = 0.20$, $c = 0.60$ and $N = 12$, even perfect orthogonality leaves $D = +1.486245$ octaves, and closure requires $N \geq e^3 = 20.0855$. **Falsified by:** a single counterexample. **Status:** derived from $F \leq \ln N$ [Curzi, 2026]; **untested against any observation.** It is the only strict universal in this paper and therefore the only claim here that one observation could kill.

H3 — the retention hypothesis, inherited. Trust acts on an arrangement by reducing loss rather than by adding input, entering as $\delta_{\text{eff}} = \delta_0 + (1-a)(1-v)$ [Curzi, 2026s]. **Predicted:** an intervention identified as trust-building changes the retention term and leaves the drive $b = \eta F - c$ unchanged, so its effect on Q_∞ takes the sign of a quantity it never touches. **Falsified by:** a measured intervention that raises the drive without changing the loss rate, or that changes the sign of a destination. **Status: untested here.** The empirical premise — that trust reduces loss rather than supplying input — is the one this paper's ordering rule inverts relative to ordinary practice, and it is supported in the literature only at the level of a mechanism, not of a measured coefficient [Gambetta, 1988, Putnam et al., 1993, Ostrom, 2010].

H4 — the survivorship hypothesis. A corpus of published testimonials is selected on outcome, so the probability that a published account reports a success is 1 under every hypothesis about the method's efficacy. **Predicted:** likelihood ratio exactly 1.000000 at any corpus size, and therefore zero bits of evidence. **Falsified by:** a testimonial corpus demonstrated to publish failures at their base rate, which would break the conditioning and restore a denominator. **Status:** the arithmetic is exact given the selection model; **the selection model itself is an empirical claim about publication and is not tested here.** The corresponding literature on selective publication in research is cited in §5.16 [Ioannidis, 2005, Simmons et al., 2011, Gelman and Loken, 2013, Open Science Collaboration, 2015].

2.2 Definitions — true by stipulation

- **Def. A.** $Q_\infty := (\eta F(u) - c)/\delta$, the destiny (Definition 4.1).
- **Def. B.** $b := \eta F(u) - c$, the per-cycle **drive**.
- **Def. C.** A **transfer of cost** of amount x from participant p to participant r is the operation $c_p \to c_p - x$, $c_r \to c_r + x$, with every other parameter unchanged. **The conservation result of Theorem 4.5 is an arithmetic consequence of this definition and of nothing else** — see §9.4.
- **Def. D.** $F(u) := \ln(1 + (N-1)u)$, the interaction surplus in nats.
- **Def. E.** $N_{\text{eff}} := 1 + (N-1)u = e^F$, the **affine effective count**. This is the corpus's own construction. **It is not a Hill number of any order** — see §3.4.
- **Def. F.** $k := (9/8)^2/\varphi^2 = 0.483425732988414$. A **defined** constant, exact by construction; its fifteen figures are exact digits of a definition and carry no measurement uncertainty, because nothing was measured.
- **Def. G.** $2/\ln 2 = 2.885390081778$, the exchange rate between nats of surplus and octaves of pitch, fixed by Def. F and the essence bridge.
- **Def. H.** $D := 2.885390\,(c/\eta - F)$, the octave deficit.
- **Def. I. Counterfeit arrival** (Definition 4.2): a fixed point positive in one participant's accounts, where the sum of drives across the arrangement did not move to produce it.

2.3 Postulates — asserted, with no external support

- **P-a. The attachment of hertz to k.** Computing $(9/8)^2/\varphi^2$ is arithmetic; asserting that the resulting dimensionless number **is a frequency** is a separate and much stronger claim, and no operational procedure is supplied by which the constant acquires the dimension of inverse time. **Every hertz value in this paper depends on it and no dimensionless result does.**
- **P-b. The essence law** $f(n) = kn^2$, inherited [Curzi, 2026]. No external support. The exponent 2 is **defined by the model, not estimated from data**: no fit was performed, no alternative form was tested, and the goodness-of-fit and likelihood-ratio procedure a scaling claim requires [Clauset et al., 2009] was **not run, because there is no data to run it on.**
- **P-c. The essence bridge** $f = k\,e^{2F}$, which is what makes the surplus and the frequency commensurable and what turns the sustainability condition into a pitch. Postulated.
- **P-d. The affine effective count** $N_{\text{eff}} = 1 + (N-1)u$, a design axiom of the corpus and **not** a Hill number [Hill, 1973, Jost, 2006].
- **P-e. The channel-gain relation** $E = (\sum_i G_i)B$ [Curzi, 2026ab], applied in §5.14 and badged in place at every use.

2.4 Modelling choices

- **M-a. The storage law is affine and first-order** [Curzi, 2026a]. Nothing here shows that a real arrangement's retention is linear, and the model is admissible only where the

trajectory stays in a range the arrangement survives. §5.6 exhibits a constructed path that spends thirteen periods below its starting stock, and the formalism computes it without complaint.

- **M-b. Parameters are held constant while a fixed point is computed.** A destiny is a statement about a configuration, not about a future, and the conditional clause is the whole of its content.
- **M-c.** $N = 12$, inherited and not shown to be the right twelve [Curzi, 2026].
- **M-d. Cost, conversion and loss are per-participant scalars.** The formalism represents transfer of cost and nothing else; degrading another party's geometry, conversion or retention is not a transfer and the conservation identity does not cover it (§4.15).

2.5 What is expository

Sections 4.17 to 4.20 and 5.17 to 5.18 are procedures and gates: instruments with blanks, whose organising question is *what must be recorded alongside a reading so that a second analyst can recompute it and disagree with it in a specific place.* They argue no claim and test no hypothesis. Section 6 is a synthesis: it states the keystone, collects six constraining results, and prints the register of demotions. Its value is in the constraints, not in a new claim.

3 Framework, dimensional audit, and method

3.1 What is inherited and what is new

The storage law, its stability condition and its DC gain are inherited [Curzi, 2026a]; the octave deficit, the ceiling $c/\eta \leq \ln N$ and the four levers are inherited [Curzi, 2026]; the elasticities and the ordering rule are inherited [Curzi, 2026q]; the contraction theorem and the consent law are inherited [Curzi, 2026]; the channel-gain postulate and the survivorship arithmetic are inherited [Curzi, 2026ab,?]. **What is new is the conservation identity under cost transfer and the export test built on it, the completeness statement for the lever set, the general form of the reserve theorem, the priced cost of misordering, and the register of demotions.**

3.2 Analogy discipline

Every borrowing in this paper is category 2 — a shared mathematical form — and none is category 1. The discrete leaky integrator is a leaky integrator whether it holds charge or stored order: the same first-order linear difference equation governs the RC low-pass filter [Oppenheim et al., 1996], exponential smoothing [Holt, 1957, Brown and Meyer, 1961], and the leaky integrate-and-fire neuron [Lapicque, 1907, Gerstner and Kistler, 2002], and its stability condition and unique fixed point are standard [Elaydi, 2005, Kelley and Peterson, 2001, Strogatz, 2015]. **No physical result about circuits is claimed to transfer.** Banach's theorem [Banach, 1922] is used as mathematics, which is what it is. The stock-and-flow reading of Q and b is the ordinary one of system dynamics [Forrester, 1961, Sterman, 2000], and the frequency interpretation sitting on top of the whole apparatus is **category 3, an analogy**, badged as such at every appearance. The wider programme of statistical physics applied to social dynamics has a survey that fixes what such models may be held to [Castellano et al., 2009], and its boundaries are the boundaries claimed here.

One term needs its own warning. Theorem 4.5 is called a *conservation* identity, and the word is borrowed. It is conservation **by construction**: a transfer is defined as an operation that adds $+x$ to one drive and $-x$ to another, and the sum telescopes. Nothing in the world is claimed to be conserved, no invariance group is exhibited, and no Noether-style warrant exists. Borrowing the notation of a conserved quantity without the conservation law is the characteristic failure of this genre [Mirowski, 1989]; the defence available here is Cartwright's — the identity is claimed within a specified arrangement and no universality is claimed for

it [Cartwright, 1999]. **Its content is diagnostic, not physical:** the two ledgers are computed independently and the sum is then checked, which is what makes the check capable of failing.

3.3 Dimensional audit

symbol	quantity	dimension	unit	status
u	non-parallelism	$[1]$	—	dimensionless, $[0, 1]$
N	category count	$[1]$	—	absolute scale; a count
F	interaction surplus	$[1]$	nats	a logarithm of a count
η	conversion	$[1]$	—	**stated dimensionless; see below**
c	cost per pass	$[1]$	nats/pass	inherits from ηF
b	drive per pass	$[1]$	nats/pass	$\eta F - c$
δ	retention loss	$[1]$	per pass	a fraction, $(0, 1]$
Q	stored order	$[1]$	nats	**no operational unit**
τ, half-life	approach times	$[T]$	passes	from δ alone
k	ground constant	$[T^{-1}]$ **asserted**	Hz **asserted**	**defined dimensionless; the unit is postulated**
f, f_{rel}, $f_{\min}$	frequencies	$[T^{-1}]$	Hz	inherit the dimension entirely from k
D	octave deficit	$[1]$	octaves	$2/\ln 2$ nats to octaves

Does it balance? Three answers, and one of them goes against the framework.

The deficit balances exactly. c/η and F are both in nats and $2/\ln 2$ converts nats to octaves, so $D = 2.885390(c/\eta - F)$ is dimensionally clean and depends on no postulate.

The storage law balances only under a reading the source does not print. $Q_{t+1} = (1-\delta)Q_t + b$ sets a stock equal to a stock plus a rate. It balances if and only if b is read as an increment already multiplied by one pass — the standard elision of discrete-time stock-and-flow notation [Forrester, 1961, Sterman, 2000] — and **the source never writes the pass-length factor.** Every quantity in the paper is per-declared-pass and no result depends on the length of a pass, so nothing is wrong; but the equation as printed is not dimensionally homogeneous and this paper says so rather than leaving a reader to find it.

η is described in two incompatible ways and the source does not resolve it. §5.2 calls it "a dimensionless efficiency in $[0, 1]$", and ηF must carry whatever unit Q carries. Both can be true only if Q is measured in nats, which is the reading adopted in the table above — and which means **stored order has no operational unit of its own; it inherits the unit of a logarithm of a count.** That is an honest dimensionless placeholder, not an invented unit, and it is the correct thing to record for a quantity nobody has measured.

Every quantity in hertz balances only because k is asserted to carry $[T^{-1}]$ (P-a). The separation is clean: the conservation identity, the convergence equality, the deficit identity, the decay blindness, the completeness statement, the reserve theorem and the misordering cost are all dimensionless or in nats, and **none of them depends on k at all.**

3.4 Which effective count, and the scale type

Which effective count. $N_{\text{eff}} = 1 + (N-1)u = e^F$ is the **affine effective count**: a linear interpolation between 1 and N, and the corpus's own construction. **It is not a Hill number of any order.** The Hill family [Hill, 1973, Jost, 2006, 2007] and its rediscovery in political science [Laakso and Taagepera, 1979] ground $1/\text{HHI}$ and $\exp(H)$ exactly — as the corpus's concentration work uses them [Curzi, 2026], whose entity map returns naive HHI 0.293388 and 3.4085 effective jurisdictions against a bloc-HHI of 1.000000 and 1.0000 effective perimeters, a 3.41-fold overstatement — and they ground this construction **not at all.** Wherever an effective count appears in this paper it is the affine one, and it is named.

Scale type. N and the state indices are counts and are absolute-scaled. u is ratio-scaled *given* that the intensities are, which requires magnitude estimation to deliver what it is held to deliver [Stevens, 1957] and is disputed [Teghtsoonian, 1971, Poulton, 1968, Krueger, 1989, Luce, 2002]. η, c and δ are treated as ratio-scaled **by assumption**, and no instrument establishing that is reported anywhere in this corpus.

The strongest single objection to this framework is Michell's — that psychology assumes quantitative structure without testing for it [Michell, 1997, 1999]. It applies here in full. The formal remedy is conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971] or a Rasch demonstration [Rasch, 1960, Andrich, 1988], and **none was run.** The corpus's own moment equation was demoted to ordinal use at its author's gate for exactly this reason [Curzi, 2026m], and the demotion travels with every psychological frequency derived through that ratio.

The scalar-reduction concession is inherited and not restated. The framework's habit of reducing a state to one number is a framework-level limitation, conceded where the ladder is built [Curzi, 2026]; this paper adopts that concession rather than repeating it, and every reading produced here rests on the index it concerns.

3.5 Identification — the objection this framework cannot presently answer

The corpus's reflexive loop $T(x) = W(x, M(x))$ [Curzi, 2026] is a linear-in-means reflexive specification, and **the endogenous social effect in such a specification is not identified from observational data** [Manski, 1993]; influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011, Lyons, 2011]. No amount of fit, held out or otherwise, recovers the parameter.

The remedy this paper adopts is stipulation. The reference set is the declared N-category ladder, the declared decomposition and the declared participant list, fixed in advance and published with every reading — which is what makes a restricted treatment of social interaction defensible where an unrestricted one is not [Brock and Durlauf, 2001]. The price is paid openly: **the reflexive claim is a modelling position, not an identified empirical result**, and no reading in this paper is offered as a measurement of an endogenous social effect.

Two objections in the same family are stated rather than resolved. **The framework exempts its own parameters from its own reflexivity** [Lucas, 1976]: if arrangements change under being modelled, then η, c and δ are the decision rules of people who have not read the storage law, and classification changes the classified [Hacking, 1995]. **And a published destiny reading is exactly the kind of thing that can make itself true** [Merton, 1948, Rosenthal and Jacobson, 1968, MacKenzie and Millo, 2003, Ferraro et al., 2005] — which is why §4.8's separation clauses are mandatory rather than courteous, and why a reading used as a target stops measuring what it measured [Goodhart, 1975, Campbell, 1979]. Against that, the best long-run assessment of the classroom case finds self-fulfilling prophecies **real but small**, not typically cumulative and often dissipating, with most expectations turning out to be accurate rather than causal [Jussim and Harber, 2005].

3.6 Method

Every number in this paper is exact arithmetic on stated inputs. The convergence equality was checked by running the recursion and by evaluating $(1-\delta)^t|Q_0-Q_\infty|$ independently, over $t=0$ to 100; the deficit identity was checked over twelve parameter sets spanning both signs of the drive and δ from 0.01 to 0.99; the conservation identity was checked across an eight-point export range in a two-party arrangement and across four transfers in a three-party arrangement. Residuals are reported at full precision and are floating-point noise in every case.

No observational data were collected, no human subjects participated, and no empirical dataset underlies any figure, table or number in this paper. Every case is constructed and is labelled CONSTRUCTED at the point of use. The figures are produced by the scripts named in their captions. Where the source manuscript printed a number this paper could not reproduce, the discrepancy is printed in §9.5 with the corrected value, rather than quietly repaired.

4 Destiny as a fixed point

The word this section takes over is heavily overloaded. In the practical literature it is used to mean an unalterable script, a personal entitlement, a reward for correct wanting, and — most often — nothing at all: a placeholder that signals importance and carries no content. A framework that has demoted its own keystone formula on a scale question [Curzi, 2026m], printed its rejections, and named its central gap in print [Curzi, 2026] cannot adopt a word like that on trust.

So it is not adopted on trust. It is defined, and the definition is an equation proved elsewhere in the corpus on assumptions stated there [Curzi, 2026a].

> **In this framework, *destiny* means Q_∞, the fixed point of the contracting storage loop** [Curzi, 2026a].

That is the whole of the claim. It says where a life, a bond, an institution or an arrangement arrives if its parameters hold. It does not say the parameters will hold, it does not say what arrival is worth, and it says nothing whatever about anybody's worthiness of it.

What the definition buys is exactly one thing, and it is the thing the rhetorical use never supplies: **it is checkable.** A destiny in this sense is a number with four arguments. Change an argument and the number changes by an amount that can be computed in advance. Fail to change any argument and the number does not move, however sincerely anything was intended. Section 5 is about the changing. This section is about what the object is, how the loop reaches it, and how to tell a real arrival from a counterfeit one.

Three results carry this section. The first is the convergence law, which says arrival is geometric and, more surprisingly, that it is an equality rather than a bound. The second is the identity between the octave deficit and the fixed point — they are the same number in different units, and the deficit form contains no `delta`, which is a genuine limitation of the framework's central instrument and is stated rather than buried. The third is a conservation identity: under a transfer of cost between participants, a certain sum does not move at all. That third result turns a standing prohibition of the framework's foundations into an arithmetic test [Curzi, 2026], and it is the most useful thing in this section.

The law

4.1 The loop, restated at destiny scope The storage law is the framework's contracting recursion [Curzi, 2026a], whose Result Two identifies the fixed point as the circuit's DC gain: evaluating the transfer function $1/(z-a)$ at $z=1$ gives $1/\delta$, so the fixed point is the net drive multiplied by the gain.

```
Q      =  (1 − delta) ∗ Q   +  ( eta ∗ F(u∗)  −  c∗ )
 t+1                     t

       delta   per−cycle retention loss, in (0, 1]
       eta     conversion: how much of an interaction's surplus becomes stored order
       F(u∗)   the interaction surplus at the operating point, F = ln(1 + 11 u) at N = 12
       c∗      the per−cycle cost of maintaining the arrangement
```

Analogy status: category 2, a shared mathematical form. This is the discrete leaky integrator — the RC low-pass filter of signal processing [Oppenheim et al., 1996], exponential smoothing in forecasting [Holt, 1957, Brown and Meyer, 1961], and the leaky integrate-and-fire neuron of computational neuroscience [Lapicque, 1907, Gerstner and Kistler, 2002]. Its stability condition and unique fixed point are standard for a first-order linear difference equation [Elaydi, 2005, Kelley and Peterson, 2001], and no claim is made that any physical result about circuits transfers.

For a fixed arrangement the bracketed term is a single constant, and naming it makes everything that follows short:

```
b  =  eta ∗ F(u∗)  −  c∗                the per−cycle DRIVE

Q      =  (1 − delta) ∗ Q   +  b
 t+1                     t
```

`b` is the net production of the arrangement per cycle: what conversion makes, less what maintenance costs. It is a rate, not a level. The distinction between a rate and a level is the axis this whole section turns on, and it is worth fixing the picture before any algebra: the drive is a tap, `delta` is a leak, and `Q_t` is the depth of water in the tank. A tank with a fast leak and an open tap settles shallow and quickly; a tank with a slow leak and the same tap settles deep and slowly. Nothing about the tap determines the depth, and nothing about the leak determines whether the tap is running forward or backward.

4.2 The fixed point exists, is unique, and is reached from anywhere

THEOREM 4.1 — EXISTENCE AND UNIQUENESS. For `delta` in `(0, 1]` the map `T(Q) = (1−delta) Q + b` is a contraction on the reals with Lipschitz constant `1 − delta < 1`. By Banach's fixed point theorem [Banach, 1922] it has exactly one fixed point, and iteration converges to it from every starting value. The framework's foundations state the theorem exactly: a contraction with modulus $q < 1$ on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least q each pass [Curzi, 2026].

Proof. `|T(x) − T(y)| = (1−delta)|x − y|`, so `T` is Lipschitz with constant `1 − delta`, which is strictly less than 1 whenever `delta > 0`. The reals are complete. Banach applies. Solving `Q = (1−delta) Q + b` gives `delta Q = b`, hence the unique fixed point. ■

```
Q     =  b / delta  =  ( eta ∗ F(u∗) − c∗ ) / delta
 inf
```

Three properties are worth reading off before the definition is stated, because each of them is a place where the ordinary use of the word "destiny" would go wrong.

It does not depend on where the loop started. `Q_0` does not appear. Two arrangements with identical `(eta, F, c, delta)` and wildly different histories have identical destinies. Whatever else a starting position determines — and it determines the whole transient, which is the part anyone actually lives through — it does not determine the destination.

It is not a maximum and not a summit. It is a level, in the hydraulic sense: the depth at which inflow and leakage balance. A loop started above `Q_inf` descends to it at exactly the rate

a loop started below ascends to it. There is no direction of travel built into the object, and any language of climbing or achieving is imported from somewhere else.

It can be negative. When $\eta F < c$, $Q_\infty < 0$, and the framework's own storage-law paper records that the regime is unsustainable or the linear model must be replaced [Curzi, 2026a]. This paper prints the negative number wherever it occurs — the carried dyad as found sits at `-5.0119395464` — and repeats the admissibility warning every time, because a model that will happily compute a trajectory through negative stored order is a model that must be watched at exactly that point.

4.3 The convergence law

THEOREM 4.2 — THE CONTRACTION, AS AN EQUALITY. For the affine map above with `delta > 0`,

```
| Q  -  Q    |  =  (1 - delta)^t  *  | Q  -  Q    |
   t     inf                          0     inf
```

Proof. Subtract `Q_inf = (1-delta) Q_inf + b` from `Q_{t+1} = (1-delta) Q_t + b`. The drive `b` cancels exactly, leaving `Q_{t+1} - Q_inf = (1-delta)(Q_t - Q_inf)`. Take absolute values and iterate `t` times. ■

The refinement paper owns this identity as its error law [Curzi, 2026], whose Result One is that refinement cannot move the fixed point — Q_∞ contains no iteration count — and that the ladder is silent in pitch, since f_{rel} depends only on u and N. It appears again here because at destiny scope it says something the refinement register does not emphasise. Banach guarantees only that distances do not increase by more than a factor. Here they shrink by exactly `1 - delta` on every pass, at every distance, forever. There is no fast region and no slow region, no early struggle and no late acceleration; the gap closes at one constant proportion per cycle from the first cycle to the last.

And the drive cancelled. Everything that decides *where* the loop is going — `eta`, `F`, `c` — vanished in the subtraction. The rate of approach is a function of the leak and of nothing else. The same fact reads as three constants — the time constant, half-life and settling time of a first-order lag in the ordinary control-engineering sense [Åström and Murray, 2008, Horowitz and Hill, 2015]:

```
tau        =  -1 / ln(1 - delta)              cycles per factor of e
half-life  =  ln(0.5) / ln(1 - delta)         cycles to close half the gap
t(eps)     =  ln(eps) / ln(1 - delta)         cycles to close all but a fraction eps
```

None contains `Q_0`, `Q_inf`, `eta`, `F` or `c`. **How long arrival takes is a property of the leak; where arrival happens is a property of everything else, and the two questions do not share a parameter.**

Computed, at the values this framework uses:

```
delta | tau (cycles/e) | half-life | t(1%)    | t(0.1%)
------+----------------+-----------+----------+---------
0.01  |      99.499162 | 68.967564 | 458.2106 | 687.3159
0.02  |      49.498316 | 34.309618 | 227.9482 | 341.9223
0.05  |      19.495726 | 13.513407 |  89.7811 | 134.6717
0.09  |      10.603253 |  7.349615 |  48.8298 |  73.2447
0.15  |       6.153129 |  4.265024 |  4.3362 |  42.5043
0.30  |       2.803673 |  1.943358 |  12.9114 |  19.3671
0.50  |       1.442695 |  1.000000 |   6.6439 |   9.9658
```

4.4 The definition

DEFINITION 4.1 — DESTINY. The destiny of an arrangement is `Q_inf = (eta F(u*) − c*)/delta`, the unique fixed point of its storage loop: the level of stored order the arrangement converges to if `(eta, F, c, delta)` hold.

Four things the definition does not say, each of which somebody will read into it if it is not refused explicitly.

It is not a prophecy. The conditional clause — *if the parameters hold* — is the entire content. Parameters are properties of an arrangement, and arrangements change. A destiny is a statement about a configuration, not about a future.

It is not a property of a person. Every term in the formula belongs to an arrangement. `u` is a joint geometric property of two or more state vectors and cannot be assigned to one of them; `eta`, `c` and `delta` describe what an arrangement does with what it produces, what it costs to run, and how fast it forgets. **There is no single-party term anywhere in the apparatus, which means no destiny reading licenses any conclusion about who someone is.** That prohibition binds harder here than anywhere else in the framework, because *destiny* is the word most likely to be heard as personal.

It is not a ranking. A larger `Q_inf` is more stored order in the model's units. It is not more worth, more virtue, or a better life, and the formalism contains no term that could make it any of those. Two arrangements with fixed points of `+6.08` and `+1.07` are not a winner and a loser; they are two configurations, and the second one may be the one somebody would rather be in for reasons the model does not represent.

It is not a clinical or diagnostic object. The twelve states are not diagnoses, a reading is not an assessment, and nothing in this section describes or prescribes for any person. That clause is restated here because this section's vocabulary invites the error more than most.

4.5 The deficit and the fixed point are one number in two units

This is the deficit identity of the minimum-frequency paper [Curzi, 2026], carried here because a section that defines destiny must state the exact relation between destiny and the instrument used to read it.

THEOREM 4.3 — THE DEFICIT IDENTITY.

```
D  =  2.885390 * ( c/eta  −  F )  =  − 2.885390 * ( delta / eta ) * Q
                                                                   inf
```

Proof. `Q_inf = (eta F − c)/delta`, so `(delta/eta) Q_inf = F − c/eta`, so `−2.885390 (delta/eta) Q_inf = 2.885390 (c/eta − F) = D`. ■

The constant is `2 / ln 2 = 2.885390081778`, and it is the exchange rate between nats of surplus and octaves of pitch, fixed by the essence bridge `f = k e^(2F)`.

Verified across twelve constructed parameter sets, spanning both signs of the drive and `delta` from 0.01 to 0.99:

u	eta	c	delta	D	−2.885390(delta/eta)Q_inf	residual
0.068735	0.30	0.620	0.09	+4.338420197	+4.338420197	0.000e+00
0.068735	0.30	0.620	0.05	+4.338420197	+4.338420197	0.000e+00
0.068735	0.30	0.620	0.01	+4.338420197	+4.338420197	0.000e+00
0.692373	0.44	0.400	0.09	−3.590988663	−3.590988663	0.000e+00
0.692373	0.44	0.400	0.05	−3.590988663	−3.590988663	8.882e−16
0.692373	0.44	0.400	0.30	−3.590988663	−3.590988663	0.000e+00
0.018614	0.34	0.825	0.09	+6.463836987	+6.463836987	8.882e−16
0.018614	0.34	0.825	0.02	+6.463836987	+6.463836987	8.882e−16
0.500000	0.50	0.500	0.50	−2.515489355	−2.515489355	0.000e+00
0.900000	0.75	1.200	0.15	−2.275888329	−2.275888329	8.882e−16
0.250000	0.20	0.300	0.40	+0.514303931	+0.514303931	3.331e−16
1.000000	0.99	2.400	0.99	−0.175039955	−0.175039955	1.665e−16

```
MAX RESIDUAL = 8.8818e-16 over 12 sets
```

The residuals are floating-point noise. **The octave deficit and the fixed point are the same quantity written twice: one in octaves of required pitch, one in units of stored order.** A reading in octaves is not an alternative to a reading in stock; it is the same reading, and the conversion factor is `-2.885390 * delta/eta`.

4.6 And the deficit contains no `delta` **THEOREM 4.4 — `delta` IS NOT A LEVER ON `D`.** `D = 2.885390 (c/eta - F)` is a function of `c`, `eta`, `u` and `N` only. `delta` does not appear. Therefore `dD/d delta = 0` identically. ■

This is the sharpest limitation of the framework's central instrument and it is stated rather than buried. The minimum-frequency paper reaches it independently [Curzi, 2026]: on its carried dyad, moving δ from 0.09 to 0.05 leaves $D = +4.338420$ unchanged while Q_∞ falls from -5.0119 to -9.0215. Take the carried dyad as found, and vary the loss term across a range covering everything from a fast-forgetting arrangement to one that forgets almost nothing:

```
D = +4.3384201973  for every row

delta | Q_inf         | half-life | drive b = delta * Q_inf
------+---------------+-----------+------------------------
0.30  |   -1.5035819 |  1.943358 |  -0.451074559176
0.20  |   -2.2553728 |  3.106284 |  -0.451074559176
0.15  |   -3.0071637 |  4.265024 |  -0.451074559176
0.09  |   -5.0119395 |  7.349615 |  -0.451074559176
0.05  |   -9.0214912 | 13.513407 |  -0.451074559176
0.02  |  -22.5537280 | 34.309618 |  -0.451074559176
0.01  |  -45.1074559 | 68.967564 |  -0.451074559176
```

The destiny moves by a factor of thirty across that column. The deficit does not move in the fifteenth decimal place. And the same is true with the sign reversed — the repaired dyad, whose drive is `+0.547598406839`:

```
D = -3.5909886634  for every row

delta = 0.30 -> Q_inf =  +1.8253280      half-life  1.943358
delta = 0.09 -> Q_inf =  +6.0844267      half-life  7.349615
delta = 0.05 -> Q_inf = +10.9519681      half-life 13.513407
delta = 0.01 -> Q_inf = +54.7598407      half-life 68.967564
```

> **The deficit reads the SIGN and the SIZE of the sustainability gap. `delta` converts it to a LEVEL. `D` alone cannot see decay, and a reading that reports `D` without `delta` is incomplete.**

The rule that follows is short and binding on every reading: report both. The mechanism behind the blindness is visible in the last column of the first table — the drive `b` is `delta * Q_inf`, and it is constant down the column because the two factors move inversely by construction. `D` is a statement about the drive. The drive is what production and cost decide. Where the drive puts the level is `delta`'s business, and the drive cannot see it.

4.7 Two independent readings, at destiny scope The framework's operation audit records that a state's frequency and a relationship's frequency are independent readings, and that neither may be inferred from the other. At destiny scope the statement needs restating

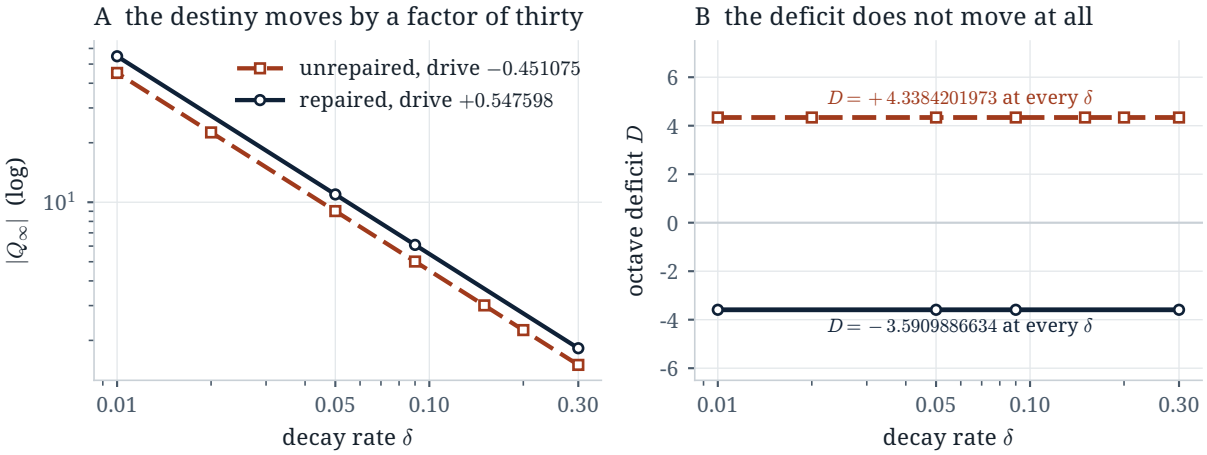


Figure 49.1: **The deficit cannot see decay: across a thirtyfold movement in the destination the octave deficit does not move in the fifteenth decimal place.** Two panels rather than a dual axis. Constructed dyad of §4.6 in both phases — unrepaired, drive $b = -0.451074559176$ and $D = +4.3384201973$ at every row; repaired, drive $b = +0.547598406839$ and $D = -3.5909886634$ at every row. Every plotted value is read from the two tables of §4.6. Script `figs/src/CCP-049_fig02.py`.

because the temptation is at its strongest here — an argument about where something arrives invites a reader to suppose that the arriving thing has a pitch, and that the pitch says something about the arrival.

Six constructed pairs of one-hot state vectors, spanning from a two-octave separation to more than four octaves of the state lattice:

indices	state frequencies	span (cents)	u	f_rel (Hz)
(1, 2)	0.4834 / 1.9337 Hz	2400.00	1.000000000000	69.6133055503
(3, 6)	4.3508 / 17.4033 Hz	2400.00	1.000000000000	69.6133055503
(5, 10)	12.0856 / 48.3426 Hz	2400.00	1.000000000000	69.6133055503
(7, 11)	23.6879 / 58.4945 Hz	1564.98	1.000000000000	69.6133055503
(2, 9)	1.9337 / 39.1575 Hz	5207.82	1.000000000000	69.6133055503
(4, 12)	7.7348 / 69.6133 Hz	3803.91	1.000000000000	69.6133055503

Identical to twelve decimal places, because `u` reads the angle between vectors and is blind to which indices are occupied. **A state's frequency carries no information about its relationship frequency, a relationship's frequency carries none about the states, and a complete report gives both separately with no inference drawn between them.** Since `F` enters the destiny and `F` comes from `u`, this means the destiny is a function of the relationship reading and not of the state readings, and any sentence of the form "their states are high, therefore their destiny is high" is an error the arithmetic forbids.

4.8 Honesty clauses, restated where they bind

Essence, not resonance. Every hertz value in this section is a position on a logarithmic scale of surplus. No person, bond, tissue or molecule is claimed to vibrate at any value here, and nothing in this section is to be played at anyone.

No clinical, diagnostic or therapeutic claim. A destiny reading is a computation on four declared parameters. It is not an assessment of anyone's health, capacity, character or prospects.

The reading inherits its inputs. `Q_inf` is exactly as defensible as `u`, `eta`, `c` and `delta`. The scale asymmetry binds: the index n is a count and is absolute-scaled by construction, but the

intensities w are ordinal by default, and where they were not elicited by magnitude estimation the relationship frequency is not determined and no destiny number may be printed at all.

Consent is structural. No procedure derived from a destiny reading may reduce any party's ability to decline it, at any stage.

All cases here are CONSTRUCTED and labelled as such. No case here is an observation of any person or pair.

The engineering

4.9 What arrival looks like, computed The carried dyad after repair — case 2b, a constructed case — has `u = 0.692372881355932`, `eta = 0.44`, `c = 0.40`, `delta = 0.09`, hence `F = 2.153632742815`, drive `b = +0.547598406839`, and destiny `Q_inf = +6.0844267427`. Start it from nothing and run it:

```
  t |      Q_t       | |Q_t − Q_inf| | (1−delta)^t |Q_0−Q_inf| | residual | gap closed
----+----------------+--------------+--------------------------+----------+-----------
  0 |   0.0000000000 |    6.0844267 |             6.0844267427 | 0.00e+00 |   0.0000%
  1 |   0.5475984068 |    5.5368283 |             5.5368283358 | 0.00e+00 |   9.0000%
  2 |   1.0459129571 |    5.0385138 |             5.0385137856 | 0.00e+00 |  17.1900%
  3 |   1.4993791978 |    4.5850475 |             4.5850475449 | 8.88e−16 |  24.6429%
  5 |   2.2875488707 |    3.7968779 |             3.7968778719 | 8.88e−16 |  37.5968%
  7 |   2.9402321769 |    3.1441946 |             3.1441945657 | 8.88e−16 |  48.3239%
 10 |   3.7150528996 |    2.3693738 |             2.3693738431 | 8.88e−16 |  61.0584%
 15 |   4.6058613008 |    1.4785654 |             1.4785654419 | 2.22e−16 |  75.6992%
 20 |   5.1617543783 |    0.9226724 |             0.9226723643 | 2.22e−16 |  84.8355%
 30 |   5.7251232522 |    0.3593035 |             0.3593034904 | 8.88e−16 |  94.0947%
 50 |   6.0299401961 |    0.0544865 |             0.0544865466 | 6.94e−16 |  99.1045%
 74 |   6.0787606627 |    0.0056661 |             0.0056660800 | 8.66e−16 |  99.9069%
100 |   6.0839388111 |    0.0004879 |             0.0004879315 | 2.06e−15 |  99.9920%

MAX RESIDUAL over t = 0..100 : 2.0624e−15
```

The third and fourth columns are the theorem, computed two ways: once by running the recursion and measuring the gap, once by evaluating `(1−delta)^t |Q_0 − Q_inf|` directly. They agree to machine precision at every step, which is what an equality looks like when it is checked rather than asserted.

Now the same arrangement from four different starting stocks, including one that starts above the destination and one that starts exactly on it:

```
Q_0         t=1         t=5         t=10        t=20        t=40
-------------------------------------------------------------------------
 0.000000  +0.547598   +2.287549   +3.715053   +5.161754   +5.944508
+1.000000  +1.457598   +2.911581   +4.104469   +5.313399   +5.967504
−4.000000  −3.092402   −0.208580   +2.157388   +4.555175   +5.852523
+12.000000 +11.467598  +9.775935   +8.388046   +6.981493   +6.220462
+6.084427  +6.084427   +6.084427   +6.084427   +6.084427   +6.084427
```

Four trajectories, one destination. The row starting at `+12` descends; the row starting at `−4` ascends through zero at around the seventh pass; the row starting exactly at the fixed point does not move at all, which is what it means for a point to be fixed. **The destination is indifferent to the history, and the history is the entire difference between the four rows.**

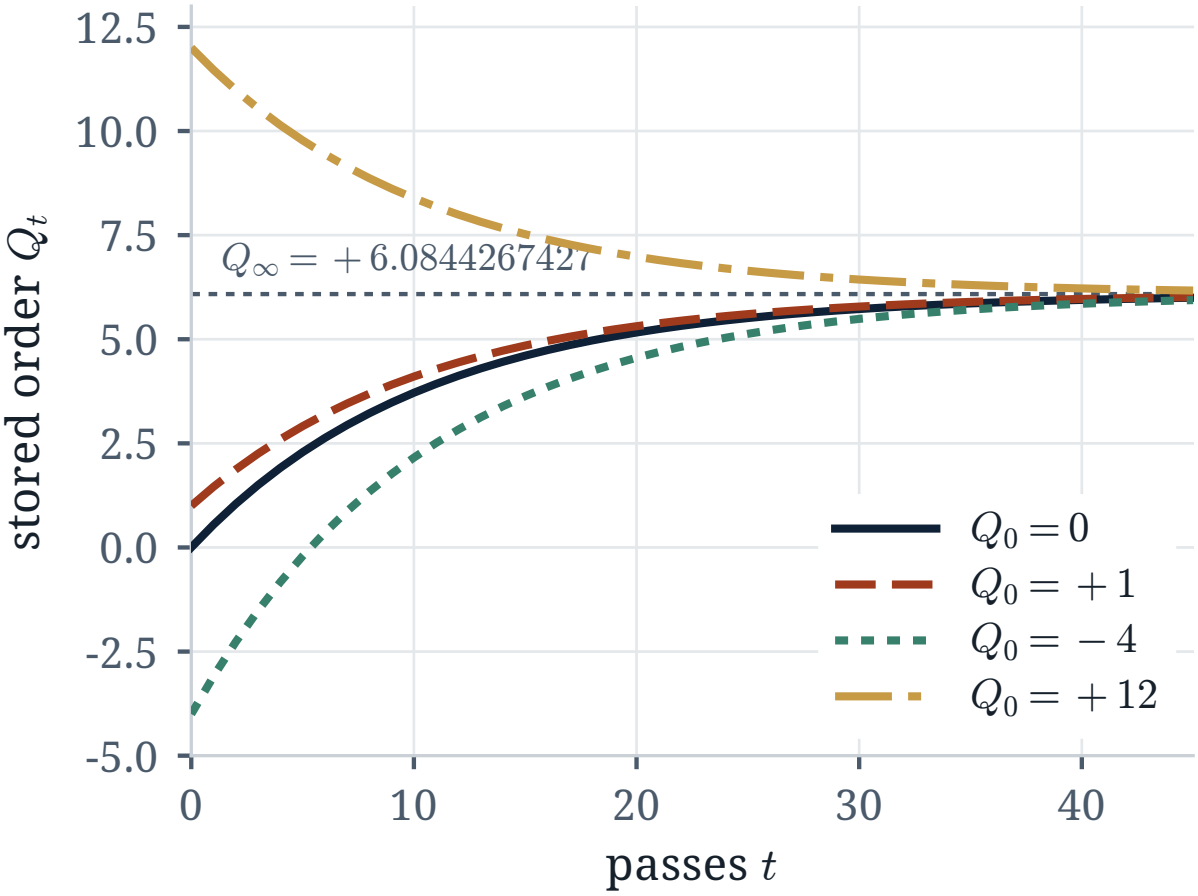


Figure 49.2: **One destination, four histories: the gap closes by exactly nine per cent per pass whether the loop is rising or falling, and a loop started exactly at the fixed point does not move.** Constructed case 2b of §4.9, modelled from the recursion at $u = 0.692372881355932$, $\eta = 0.44$, $c = 0.40$, $\delta = 0.09$, $N = 12$, giving drive $b = +0.547598406839$ and $Q_\infty = +6.0844267427$; every marked value is read from the four-start table of §4.9. Script `figs/src/CCP-049_fig01.py`.

4.10 Why arrival is recognition rather than achievement The framework's founding definitions record that arrival is experienced as recognition rather than as victory, and offer no mechanism [Curzi, 2026]. The refinement paper supplies the mechanism in four parts [Curzi, 2026] and it is cited rather than restated, with one addition that belongs at destiny scope.

The psychological literature on adaptation reaches the same shape from measurement rather than algebra: an adaptation level against which experience is judged [Helson, 1948], a return toward a baseline after large positive and negative events [Brickman et al., 1978], and the qualifications later work attaches to how complete that return is [Diener et al., 2006]. The four parts, compressed: there is no arrival event, because `(1-delta)^t > 0` for every finite `t` and nothing ever completes; almost the whole trajectory is spent near the destination, since at `delta = 0.09` the loop is inside five per cent of `Q_inf` from pass 32 and inside a tenth of a per cent from pass 74; every term of the fixed point is a standing property of the arrangement, so nothing in the destination was imported; and the same point is reached from any start, so it is a level rather than a summit.

The addition is this. A reader who has followed the section to here has the material for a sharper version of the fourth point. The destination is not merely reachable from any start — it is reached at a rate that does not know which direction it is travelling in. Row three of the table above ascends 10.08 units in forty passes; row four descends 5.78 in the same forty; and the gap between each and the destination shrinks by exactly nine per cent per pass in both cases. **A process that closes the same proportion of its distance whether it is rising or falling is not a process of achievement, and no vocabulary of effort or reward will fit it correctly.**

4.11 Traits, and destiny in the personal register The framework's treatment of standing dispositions as low-δ states is the bridge between this section's arrangement-scoped object and the way the word *destiny* is usually meant.

A mood is a state held with a large `delta`: it decays fast, its half-life is short, and its fixed point is shallow because a large denominator divides the drive. A trait is the same machinery at small `delta`: slow decay, long half-life, and a deep fixed point in whichever direction the drive points. This is why traits look like destinies in the colloquial sense — they are the readings with the longest half-lives, so they are the ones that are still there when anybody looks — and it is also why the colloquial sense is wrong about them. **A trait is not a fixed quantity; it is a fixed point with a long approach, which is a completely different object and behaves differently under intervention.**

The difference matters and it is computable. Take the repaired dyad's drive and vary only the loss term, as if the same production were embedded in arrangements that forget at different rates:

```
delta | Q_inf       | half-life | passes to close 99% of the gap
------+-------------+-----------+-------------------------------
0.30  |  +1.8253280 |  1.943358 |   12.9114
0.09  |  +6.0844267 |  7.349615 |   48.8298
0.05  | +10.9519681 | 13.513407 |   89.7811
0.01  | +54.7598407 | 68.967564 |  458.2106
```

A low-`delta` arrangement has a destiny thirty times deeper and takes thirty-five times as long to get there. Both halves of that sentence are consequences of dividing by a small number, and the second half is the one that gets forgotten. **A deep destiny is a slow one, necessarily, and the same parameter sets both. There is no configuration in this model with a deep fixed point and a short approach.**

And the sign carries through unchanged. Run the same column on the *unrepaired* dyad, whose drive is negative, and the low-`delta` arrangement arrives at `-45.1074559` instead of `+54.7598407`. Slow decay is not a virtue. It is a magnifier, and what it magnifies is whatever sign the arrangement already had. That is the framework's ordering rule, and Section 5 takes

it as an operating rule; the elasticities that price it are fixed in the return-difference paper [Curzi, 2026q], whose Result Four gives η $+1.594$, δ -1.000, u $+0.770$, C -0.594.

4.12 Counterfeit arrival — the conservation identity A fixed point can be positive, stable, arrived at, and produced by taking it out of somebody else's. The refinement paper computed the case [Curzi, 2026]; this section names it and gives the identity in general form, because the general form is what makes it a test.

THEOREM 4.5 — THE DRIVE IS CONSERVED UNDER COST TRANSFER. Let participants `i = 1..m` share an arrangement, each with drive `b_i = eta_i F_i − c_i`. Let a transfer move an amount `x` of cost from participant `p` to participant `r`, so `c_p -> c_p − x` and `c_r -> c_r + x`, with all other parameters unchanged. Then

```
sum_i b_i     is invariant,   and equivalently    sum_i delta_i ∗ Q          is invariant
                                                                 inf,i
```

Proof. The transfer adds `+x` to `b_p` and `−x` to `b_r` and changes no other term. The sum telescopes to zero. Since `b_i = delta_i Q_inf,i` by Theorem 4.1, the delta-weighted sum of fixed points is the same quantity. ■

COROLLARY 4.6 — WHEN ALL PARTICIPANTS SHARE ONE `delta`, THE PLAIN SUM OF FIXED POINTS IS INVARIANT. If `delta_i = delta` for all `i`, then `sum_i Q_inf,i = (sum_i b_i)/delta`, which Theorem 4.5 has just shown does not move. ■

Here is the corollary computed. **CONSTRUCTED CASE, labelled constructed; not an observation of any person or pair.** Two participants, P and R, share the case-2b arrangement, so both draw on `F = 2.153632742815` at `u = 0.692372881355932` and `f_rel = 35.8881768941 Hz`. Their private parameters differ, which is the ordinary situation. P declines a share of the maintenance — the scheduling, the repair conversations, the remembering of what was agreed — and R absorbs it. Let `x` be the exported amount.

```
  x    |  c_P  |  c_R  |    D_P    |  Q_inf,P  |    D_R    |  Q_inf,R  |    SUM
-------+-------+-------+-----------+-----------+-----------+-----------+-------------
0.000  | 0.400 | 0.550 | −3.590989 | +6.084427 | −0.924189 | +1.067665 | +7.152091441
0.025  | 0.375 | 0.575 | −3.754931 | +6.362205 | −0.683740 | +0.789887 | +7.152091441
0.050  | 0.350 | 0.600 | −3.918874 | +6.639982 | −0.443290 | +0.512109 | +7.152091441
0.0961 | 0.304 | 0.646 | −4.221119 | +7.152093 | +0.000002 | −0.000002 | +7.152091441
0.100  | 0.300 | 0.650 | −4.246759 | +7.195538 | +0.037608 | −0.043446 | +7.152091441
0.150  | 0.250 | 0.700 | −4.574644 | +7.751093 | +0.518506 | −0.599002 | +7.152091441
0.200  | 0.200 | 0.750 | −4.902530 | +8.306649 | +0.999405 | −1.154558 | +7.152091441
0.250  | 0.150 | 0.800 | −5.230415 | +8.862205 | +1.480303 | −1.710113 | +7.152091441

sum of drives = +0.643688229683, invariant to every digit across the whole range
```

The sum does not move in any decimal place, across an export range wide enough to take R from comfortably sustainable to badly failing. Cost export is an externality in the ordinary sense: a private gain that is exactly a transfer, appearing as a gain only because the accounts are drawn around one party [Hardin, 1968, Ostrom, 1990]. Cost export produces nothing. It moves an existing quantity from one ledger to the other, and both ledgers are correct.

That last clause is what makes the failure mode hard. P's instrument is working perfectly. P's deficit falls from `−3.590989` to `−4.902530` — nearly a further octave and a third of headroom — and P's fixed point rises by more than two units. Every number P can compute about P has genuinely improved, and P is not lying to anybody, including P.

R crosses the sustainability boundary between `x = 0.05` and `x = 0.10`. Solving `D_R = 0` gives `c_R = eta_R F = 0.646089822845`, i.e. the exact crossing at `x = 0.096089822845`.

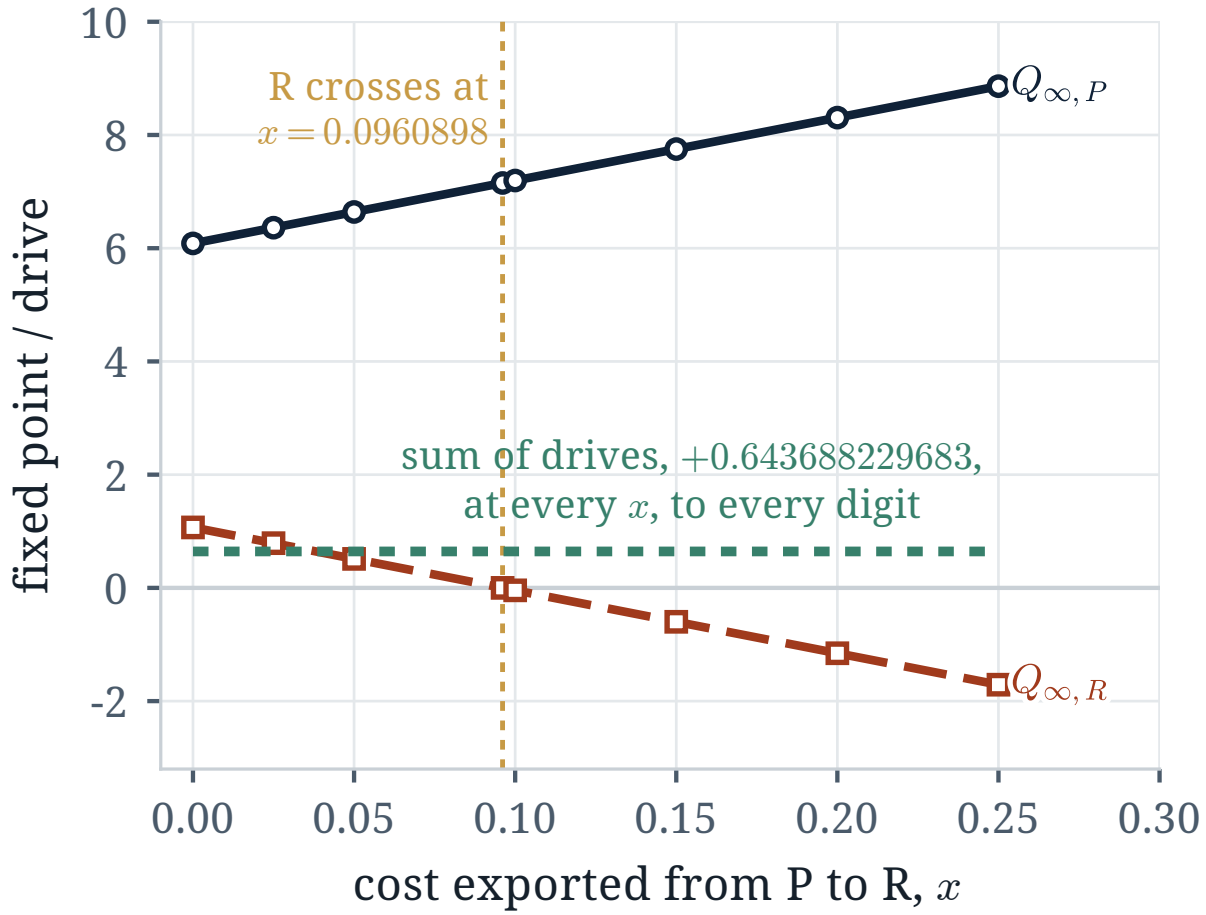


Figure 49.3: **Conservation of drive: one party's destination rises and the other's falls, and the sum of drives holds at $+0.643688229683$ to every digit across the whole export range.** Constructed two-party case of §4.12, both parties sharing $F = 2.153632742815$ at $u = 0.692372881355932$ and $\delta = 0.09$, with $\eta_P = 0.44$ and $\eta_R = 0.30$; the disadvantaged party crosses the sustainability boundary exactly at $x = 0.096089822845$. Every plotted value is read from the export table of §4.12. Script `figs/src/CCP-049_fig03.py`.

4.13 The honest caveat: when the plain sum is not the invariant Corollary 4.6 required equal `delta`. The refinement paper's worked case satisfied that condition and did not have to dwell on it [Curzi, 2026]. The general statement does, because the plain sum of fixed points is **not** conserved when the participants forget at different rates, and a test stated in the wrong form fails in exactly the situation where it is most needed.

Differentiate the plain sum with respect to the exported amount:

```
d ( Q       +  Q       ) / dx   =   1/delta   -   1/delta
     inf,P      inf,R                      P            R
```

which is zero if and only if `delta_P = delta_R`. Computed, with the same arrangement but `delta_P = 0.05` and `delta_R = 0.15`:

```
  x   | Q_inf,P (d=0.05) | Q_inf,R (d=0.15) |    SUM     | sum of drives   | delta-weighted sum
------+------------------+------------------+------------+-----------------+-------------------
0.00  |    +10.951968137 |     +0.640598819 | +11.592567 | +0.643688229683 | +0.643688229683
0.05  |    +11.951968137 |     +0.307265486 | +12.259234 | +0.643688229683 | +0.643688229683
0.10  |    +12.951968137 |     -0.026067848 | +12.925900 | +0.643688229683 | +0.643688229683
0.15  |    +13.951968137 |     -0.359401181 | +13.592567 | +0.643688229683 | +0.643688229683
0.20  |    +14.951968137 |     -0.692734514 | +14.259234 | +0.643688229683 | +0.643688229683

d(SUM)/dx  =  1/0.05 - 1/0.15  =  13.333333333
```

The plain sum rises by more than two and a half units across that range while nothing whatever is produced, because the export moved cost from a slow-forgetting ledger to a fast-forgetting one and the model's own division by `delta` did the rest. A test that watched the plain sum would report that the arrangement had improved. **The invariant is the sum of drives — equivalently the delta-weighted sum of fixed points — and it is exact under any transfer of cost between any number of participants with any `delta` values whatever.**

Verified on a three-party constructed arrangement with `(eta, c, delta)` of `(0.44, 0.40, 0.05)`, `(0.30, 0.55, 0.15)` and `(0.62, 0.71, 0.22)`, all sharing `F = 2.153632742815`:

```
baseline sum of drives                                  = 1.268940530229
transfer 0.13 from party 1 to party 2  -> sum of drives    = 1.268940530229   resid 0.000e+00
transfer 0.31 from party 1 to party 3  -> sum of drives    = 1.268940530229   resid 0.000e+00
transfer 0.07 from party 3 to party 2  -> sum of drives    = 1.268940530229   resid 0.000e+00
transfer 0.44 from party 2 to party 1  -> sum of drives    = 1.268940530229   resid 0.000e+00
```

4.14 The lag, and why the export is not noticed R does not disappear on the pass the sign flips. R's stock decays from its old level toward its new one at `(1-delta)^t`, exactly as Theorem 4.2 requires, and for a long while R's behaviour is still the behaviour of a participant with something in the bank. At `x = 0.15`, R's destiny has moved from `+1.067664698` to `-0.599001968`:

```
pass |   Q_R
-----+-------------
   0 | +1.067664698
   3 | +0.656949698
   7 | +0.262266397
  10 | +0.050024895
  11 | -0.008387523
  15 | -0.193988343
  25 | -0.441283134
  40 | -0.560675002
  60 | -0.593189879

exact zero-crossing at t* = 10.850479 passes
```

For nearly eleven passes after the export begins, R still has a positive stock — while R's destination has been negative since the first pass. Nothing observable changes at the moment the damage is done. This is precisely why the deficit is the instrument and the stock is not: `D_R` went positive immediately, and `Q_R` took eleven passes to notice.

And there is an end to the arithmetic. The option R retains is exit, the classical alternative to voice when an arrangement deteriorates [Hirschman, 1970]. Suppose R eventually declines to continue, which is R's structural right under the framework's consent law [Curzi, 2026] and which no arrangement here may remove. The shared arrangement ends, and with no partner there is no interaction surplus: `u = 0`, `F = ln(1) = 0`, `f = k = 0.483426 Hz`, the ground essence.

```
P at x = 0.15, before R's departure   D_P = −4.574644    Q_inf,P = +7.751093   f = 35.888177 Hz
P at x = 0.15, after  R's departure   D_P = +1.639426    Q_inf,P = −2.777778   f =  0.483426 Hz

fall in P's fixed point              10.528871 units
swing in P's deficit                  6.214071 octaves, the wrong way
```

The reduced cost P retained is worth nothing once the surplus it was drawn against is gone. The swing is exactly `2.885390 * F = 6.214071` octaves, which is to say: **what P loses at the end is precisely the whole of the surplus P was exporting the cost of.**

4.15 A standing prohibition converted into an arithmetic test The framework's foundations state a prohibition on intervening at another party's expense, and a prohibition of that form is unfalsifiable in exactly the way a testimonial is: it reports an intention, and the intention is always available. Theorem 4.5 converts it into something else.

> **THE EXPORT TEST. Compute the deficit and the fixed point for every participant separately. Then compute the sum of drives. If a participant's reading improved and the sum of drives did not move, the gain was exported and nothing was produced.**

Three properties make this a test rather than an exhortation.

It is arithmetic, not attribution. It does not ask what anybody meant. It asks whether a conserved quantity moved. A well-intentioned export and a deliberate one produce the same reading, which is correct, because the harm is the same.

It runs on pass one. The stock-based version of the same question — "is anyone worse off?" — returns a false negative for nearly eleven passes in the worked case. The drive-based version returns the right answer immediately, because the drive is where the transfer happened.

It costs one column. The reporting format of §4.19 adds the sum of drives to what a reading already contains. Nothing new has to be measured; the parameters were already collected for each participant's own reading.

What the test does not do is also worth stating. It detects export of *cost*, which is the operation the formalism represents. It does not detect every way one party can benefit at another's expense, because the formalism does not represent every such way. A participant who degrades another's `u` by suppressing what they hold, or who lowers another's `eta` by refusing to let their work become permanent, is doing something the model can express only as a change in somebody's parameters, and the conservation identity does not cover it — those changes are not transfers and the sum of drives does move. The test is exact for what it covers and silent outside it, and a reader who takes it for a general detector of exploitation will be misled by their own instrument.

4.16 The sensitivity of a destiny A destiny is a quotient, and quotients with small denominators are sensitive. The partials, at the carried dyad as found:

```
dQ_inf / d eta    =  F / delta                =   +6.256498  per unit of eta
dQ_inf / d F      =  eta / delta              =   +3.333333  per nat of surplus
dQ_inf / d c      =  −1 / delta               =  −11.111111  per unit of cost
dQ_inf / d delta  =  −(eta F − c) / delta^2    =  +55.688217  per unit of delta

and the chain into the geometry:
dF / du           =  (N−1) / (1 + (N−1)u)     =   +6.263947
dF / dN           =  u / (1 + (N−1)u)         =   +0.039141
```

The fourth line is nearly nine times the third and nearly nine times the first. **The parameter with the largest partial derivative on the destiny is the one with no partial derivative at all on the deficit,** which is Theorems 4.3 and 4.4 stated as a pair of numbers rather than as an algebraic fact, and it is the reason both readings have to be published.

The partials are in four different units, and ranking them against one another on size is forbidden. What they are good for is error propagation: a destiny reported without an interval, when `delta` is known to `± 0.02`, is reporting a number whose uncertainty at this operating point is of order `± 1.1` units. The framework's resolution rule applies unchanged — **a verdict is resolved only when it exceeds its own uncertainty** — and §4.20's gates enforce it.

The specification

4.17 PROCEDURE DEST — reading a destiny **Step 1. Declare the arrangement.** Name the participants, the cycle length, and what one pass is. A destiny is per relation, per site, per period; ratios do not average and readings do not aggregate.

Step 2. Check the scale precondition. Were the intensities elicited by magnitude estimation [Stevens, 1957]? If not, stop. The relationship frequency is not determined, and a destiny computed from an undetermined `F` is an artefact of the labelling. Publish the raw L2 norms so the free multiplier stays recoverable.

Step 3. Check for one-hot degeneracy. If the participants' state vectors have disjoint support, `u = 1` automatically and the surplus is meaningless. Re-elicit with overlapping support or stop.

Step 4. Compute `u`, then `F`, then `f_rel`. Use $u = 1 - (\hat{x} \cdot \hat{y})^2$, $F = \ln(1 + 11u)$, $f_{\text{rel}} = k(1 + 11u)^2$. Carry twelve or more decimal places throughout and recompute derived values from the full-precision u, never from a rounded print.

Step 5. Run the ceiling test before anything else. If `c/eta > ln N`, no value of `u` makes the arrangement sustainable and the case is structural, not tunable. Report it as structural and go to §5's restructuring branch.

Step 6. Compute the deficit. `D = 2.885390 (c/eta − F)`. Report its sign, its size, and its sensitivity to each input.

Step 7. Compute the destiny. `Q_inf = (eta F − c)/delta`. Report `delta` in the same sentence.

Step 8. Compute the approach. `tau`, half-life, and `t(1%)` from `delta` alone. Say how many passes of the declared cycle length that is in calendar time.

Step 9. Run the export test. For every participant separately, compute `D_i` and `Q_inf,i`. Compute `sum_i b_i`. Compare to the previous reading's sum of drives if one exists.

Step 10. Publish both readings and both refusals. The state frequencies and the relationship frequency, separately, with no inference between them; the essence clause; the no-clinical-claim clause; and the constructed label on every case.

4.18 Definitions and theorems owned here

```
destiny = Q_inf                              Def 4.1     "Definition 4.1"
existence, uniqueness, start-independence    Thm 4.1     "Theorem 4.1"
the convergence equality                     Thm 4.2
the deficit identity                         Thm 4.3
delta is absent from D                       Thm 4.4
conservation of drive under cost transfer    Thm 4.5     "Theorem 4.5"
equal-delta sum invariance                   Cor 4.6     "Corollary 4.6"
counterfeit arrival, defined                 Def 4.2     "Definition 4.2"
the export test                              Sec 4.15    "the export test"
```

DEFINITION 4.2 — COUNTERFEIT ARRIVAL. A fixed point is counterfeit when it is positive in one participant's accounts and the sum of drives across the arrangement did not move to produce it.

Not owned here and taken from their owners: the component index and its count argument; the elicitation procedure; the relationship-frequency rule; the order-of-operations rule [Curzi, 2026q]; and the four levers with what each one costs [Curzi, 2026], priced in §5.

4.19 Reporting format Every published destiny reading carries all of the following, in full:

```
ARRANGEMENT      participants, cycle length, what one pass is
SCALE STATUS     magnitude-estimated? raw L2 norms? -- if no, NO NUMBER IS PRINTED
GEOMETRY         u to 12+ dp, cos theta, support overlap
STATE READING    the occupied indices, their raw frequencies, their pitch classes
RELATION READING F in nats, N_eff, f_rel in Hz          -- separately, no inference drawn
CEILING          c/eta against ln N, and the verdict structural / tunable
DEFICIT          D in octaves, with its sensitivity to each input
LOSS             delta, stated in the same block as D and never omitted
DESTINY          Q_inf, with its uncertainty
APPROACH         tau, half-life, t(1%), converted to calendar time
EXPORT           per-participant D_i and Q_inf,i, and the sum of drives
LABEL            CONSTRUCTED or MEASURED, on every case
REFUSALS         essence-not-resonance; no clinical claim; consent structural
```

4.20 Gates

```
GATE 4-1   intensities magnitude-estimated and L2 norms published    FAIL -> no number
GATE 4-2   support overlap checked for one-hot degeneracy            FAIL -> re-elicit
GATE 4-3   ceiling test run before any lever is discussed            FAIL -> incomplete
GATE 4-4   D and delta reported together                             FAIL -> incomplete
GATE 4-5   Q_inf reported with an uncertainty                        FAIL -> incomplete
GATE 4-6   state and relationship readings reported separately       FAIL -> incomplete
GATE 4-7   sum of drives computed across all participants            FAIL -> export unchecked
GATE 4-8   every case labelled CONSTRUCTED or MEASURED               FAIL -> withdraw
GATE 4-9   no single-party term appears anywhere in the report       FAIL -> rewrite
```

The failure modes

FM 4.1 — DESTINY READ AS PROPHECY. *Signature:* a fixed point reported as what will happen. *Mechanism:* the conditional clause dropped. `Q_inf` is where the loop goes if the parameters hold, and parameters are properties of arrangements that people change. *Fix:* the conditional is part of the number. Print it.

FM 4.2 — DESTINY READ AS A PROPERTY OF A PERSON. *Signature:* "my destiny is `-5.01`." *Mechanism:* `u` is a joint geometric property and `eta`, `c`, `delta` describe an arrangement. There is no single-party term in the apparatus. *Fix:* Gate 4-9. State the arrangement, or print nothing.

FM 4.3 — THE DEFICIT REPORTED ALONE. *Signature:* a reading consisting of `D` and a verdict. *Mechanism:* Theorem 4.4. `D` cannot see decay at all; the same `D` is consistent with fixed points differing by a factor of thirty. *Fix:* Gate 4-4.

FM 4.4 — `delta` REPORTED ALONE. *Signature:* a fast half-life presented as good news. *Mechanism:* the rate cannot see the destination's sign. A fast approach to a negative fixed point is a fast decline. *Fix:* Gate 4-4, the other way round.

FM 4.5 — TRUST AS A DESTINY LEVER IN A FAILING ARRANGEMENT. *Signature:* "we will lower `delta` and our destiny will improve." *Mechanism:* the ordering rule. `delta` divides the drive; when the drive is negative a smaller `delta` makes the fixed point more negative. Computed: `-5.0119395464` at `delta = 0.09` becomes `-45.1074559` at `delta = 0.01`. *Fix:* the order rule [Curzi, 2026q], and §5.5.

FM 4.6 — THE STOCK MISTAKEN FOR THE DESTINY. *Signature:* "we are at `+1.07`, so we are fine." *Mechanism:* the stock lags the destination by up to `t*` passes; in the worked export case `t* = 10.850479`. *Fix:* compute `D_i`, not `Q_t`.

FM 4.7 — COUNTERFEIT ARRIVAL. *Signature:* one participant's readings improve steadily and genuinely, pass after pass. *Mechanism:* cost export; the sum of drives did not move. *Fix:* Gate 4-7, the export test.

FM 4.8 — THE EXPORT TEST RUN ON THE PLAIN SUM. *Signature:* the sum of fixed points is watched, the participants have different `delta`, and the test passes an export. *Mechanism:* §4.13. `d(sum)/dx = 1/delta_P − 1/delta_R`, which is zero only when the rates match; in the constructed case the plain sum rises by 2.67 units while nothing is produced. *Fix:* watch the sum of drives, which is invariant unconditionally.

FM 4.9 — THE EXPORT TEST TREATED AS A GENERAL DETECTOR. *Signature:* an arrangement declared free of exploitation because the sum of drives held. *Mechanism:* §4.15. The identity covers transfers of cost. Degrading another party's `u`, `eta` or `delta` is not a transfer, and the sum moves. *Fix:* state the test's scope in the same sentence as its result.

FM 4.10 — THE DEEP FIXED POINT WANTED WITHOUT THE SLOW APPROACH. *Signature:* a plan for a durable arrangement that also arrives quickly. *Mechanism:* §4.11. `Q_inf` scales as `1/delta` and the half-life scales as `−ln 2 / ln(1−delta)`; one parameter sets both, in opposite directions. *Fix:* choose. There is no configuration with both.

FM 4.11 — `delta = 0`. *Signature:* a model of something incorruptible, unchanging, permanent. *Mechanism:* at `delta = 0` the map is an isometry, Banach does not apply, there is no fixed point, and `Q_t = Q_0 + bt` grows without bound. The image of permanence is the one setting under which the apparatus stops working. *Fix:* name the leak. A low-`delta` standing disposition is the honest version.

FM 4.12 — INFERRING DESTINY FROM STATE FREQUENCY. *Signature:* "their states are high, so their destiny is high." *Mechanism:* §4.7; six constructed pairs spanning 1565 to 5208 cents return identical `u` and identical `f_rel` to twelve decimal places. *Fix:* Gate 4-6. Two readings, reported separately, no inference between them.

FM 4.13 — THE PRECISION LEAK. *Signature:* a destiny that disagrees with a colleague's in the fourth significant figure. *Mechanism:* precision loss; recomputing from a six-decimal `u` shifts `f` and `Q_inf` in the fifth significant figure. *Fix:* recompute from the full-precision `u`. Every time.

FM 4.14 — THE READING USED AS COUNSEL. *Signature:* a destiny presented as advice about whether to stay in a relationship, or as an evaluation of anyone. *Mechanism:* register slip. `Q_inf < 0` says a configuration loses stored order at a computable rate. It does not weigh what the relation is for, what else it produces, or what ending it would cost. *Fix:* §4.8. This is not counselling, therapy or clinical assessment, and the paper says so here.

Worked examples

4.21 The carried dyad, both phases, as full destiny reports CONSTRUCTED. Not an observation of any person or pair.

```
ARRANGEMENT      two parties A and B; one pass = one week; the pass is their standing review
SCALE STATUS     constructed intensities, treated as magnitude-estimated for the worked case
RAW L2 NORMS     |A| = 0.948683298051    |B| = 0.890224690738

---- PHASE 2a, AS FOUND ------------------------------------------------------
GEOMETRY         cos theta = 0.965020898600      u = 0.068734665264634
                 support overlap = 3 of 4 occupied indices -- not degenerate
STATE READING    A holds {4, 6, 9}; B holds {4, 6, 9, 10}
                 raw 7.7348 / 17.4033 / 39.1575 / 48.3426 Hz
                 pitch classes {1, 3, 9, 5}
RELATION READING F = 0.563084802746 nats   N_eff = 1.756081   f_rel = 1.4907987150 Hz
CEILING          c/eta = 2.066667 vs ln N = 2.484907  ->  TUNABLE in principle
DEFICIT          D = +4.3384201973 octaves        f_min = 30.1587698405 Hz
LOSS             delta = 0.09
DESTINY          Q_inf = -5.0119395464            DECAYS
APPROACH         tau = 10.603253 passes; half-life = 7.349615; t(1%) = 48.8298
                 at one pass per week: half-life about seven and a half weeks
LABEL            CONSTRUCTED
```

The pair agree about almost everything — `cos theta = 0.965` — and that is the whole of their problem. The relationship frequency sits at 1.49 Hz against a required 6.16 Hz, which is a gap of 4.34 octaves, and the destiny is negative.

The intuitive repair is to add trust, which lowers `delta`. Computed:

```
delta 0.09 -> 0.05     Q_inf  -5.0119395464 -> -9.0214911835      WORSE by 4.0095516
delta 0.09 -> 0.01     Q_inf  -5.0119395464 -> -45.1074559        WORSE by 40.0955164
D at every one of those values                +4.3384201973       UNCHANGED
```

The deficit is flat across all three. A pair reading only their deficit would conclude that a successful, expensive, well-intentioned intervention had done nothing; a pair reading only their stock would watch it fall and not know why. **A pair reading both would see immediately that they had multiplied a negative number by a larger factor and touched no mechanism.** That is the case for Gate 4-4 in one paragraph.

```
---- PHASE 2b, AFTER ONE STRUCTURAL CHANGE -----------------------------------
THE CHANGE       B takes on two states A does not hold; the pair adopt a decision log
GEOMETRY         u = 0.692372881355932
RELATION READING F = 2.153632742815 nats   N_eff = 8.616102   f_rel = 35.8881768941 Hz
CEILING          c/eta = 0.909091 vs ln N = 2.484907  ->  TUNABLE, with margin
DEFICIT          D = -3.5909886634 octaves        f_min = 2.9782153324 Hz
LOSS             delta = 0.09
DESTINY          Q_inf = +6.0844267427            ACCUMULATES
APPROACH         half-life 7.349615 passes; inside 5% of Q_inf from pass 32;
                 inside 0.1% from pass 74
LABEL            CONSTRUCTED
```

And now the same trust intervention, applied second instead of first:

```
delta 0.09 -> 0.05     Q_inf  +6.0844267427 -> +10.9519681368     BETTER by 4.8675414
D unchanged at         -3.5909886634
```

The symmetry is exact and worth pausing on. Lowering `delta` from 0.09 to 0.05 moved the unrepaired dyad's destiny by `-4.0095516` and the repaired dyad's by `+4.8675414`. **The same intervention, the same magnitude of change in the same parameter, and the sign of the result was decided entirely by the sign of the drive it was multiplying.** Section 5 owns the operating rule that follows; this section's contribution is the reason it is true, which is that `delta` is a denominator and denominators do not have opinions.

4.22 The export case, in report format

CONSTRUCTED. Not an observation of any person or pair. The two-participant arrangement of §4.12, read at `x = 0.15`.

```
ARRANGEMENT      P and R share the case-2b arrangement; one pass = one week
SHARED GEOMETRY  u = 0.692372881355932   F = 2.153632742815   f_rel = 35.8881768941 Hz

PARTICIPANT P    eta = 0.44   c = 0.250   delta = 0.09
                 D_P     = -4.574644 octaves
                 Q_inf,P = +7.751093
                 drive   = +0.697598

PARTICIPANT R    eta = 0.30   c = 0.700   delta = 0.09
                 D_R     = +0.518506 octaves
                 Q_inf,R = -0.599002
                 drive   = -0.053910

SUM OF DRIVES    +0.643688229683
PREVIOUS SUM     +0.643688229683         ->  UNCHANGED
EXPORT TEST      P improved by 1.666666 units of fixed point; the sum did not move.
                 VERDICT: the gain was exported. Nothing was produced.

R's crossing     exact at c_R = eta_R F = 0.646089822845, i.e. x = 0.096089822845
R's stock lag    t* = 10.850479 passes before Q_R goes negative
LABEL            CONSTRUCTED
```

Read the block as a whole and notice what it does that a per-party reading cannot. P's three lines are the lines of a well-run arrangement improving steadily. R's three lines are the lines of an arrangement that has just crossed into decline. The sum line is the only place in the report where the relation between those two facts appears, and it appears as a quantity that did not change.

The terminal arithmetic completes the case. If R declines to continue — R's structural right, and no arrangement here may remove it — the shared surplus goes to zero, `f` falls to the ground essence `k = 0.483426 Hz`, and P's reading becomes `D_P = +1.639426`, `Q_inf,P = -2.777778`. P's fixed point falls by `10.528871` units and P's deficit swings `6.214071` octaves the wrong way, which is `2.885390 * F` exactly. **The exporter's terminal loss is the full value of the surplus whose maintenance was exported, and the model computes it without needing any account of anybody's feelings about it.**

4.23 What this section establishes, and what Section 5 takes

Established, and checkable. That the storage loop has exactly one fixed point for `delta > 0` and reaches it from every start (Theorem 4.1); that the approach is a geometric equality rather than a bound, verified to 2.06e-15 across a hundred passes (Theorem 4.2); that the octave deficit and the fixed point are one quantity in two units, verified to 8.88e-16 across twelve parameter sets (Theorem 4.3); that the deficit contains no `delta`, so the same `D` is consistent with destinies differing by a factor of thirty (Theorem 4.4); that the sum of drives is exactly conserved under any transfer of cost among any number of participants (Theorem 4.5), and that the plain sum of fixed points inherits that invariance only when the participants share a `delta` (Corollary 4.6 and §4.13);

that a counterfeit fixed point is therefore detectable on pass one by an arithmetic test costing one column of a report; and that a state's frequency and a relationship's frequency remain independent readings at destiny scope.

Not established, and not claimed. That any arrangement's parameters will hold; that a larger fixed point is a better life; that any reading here describes any actual person or pair; that the formalism detects every way one party can gain at another's expense; that any frequency in this section is a physical resonance; or that any of this constitutes counselling, therapy, diagnosis or clinical assessment.

The one line to carry forward: *a destiny is a number with four arguments, and everything that can be done about one is done to those four.*

Section 5 takes the four arguments as levers, prices each of them, gives each its frequency reading and its failure mode, and states the two rules governing the order in which they are pulled. It also applies the channel-gain postulate in full [Curzi, 2026ab], because the question "what actually moves a destiny" has a large and confident literature attached to it.

5 Changing the fixed point

Section 4 defined destiny as Q_∞, the fixed point of the contracting loop, and spent its length on what the object is: unique, reached from any start, indifferent to history, negative whenever conversion fails to cover cost, and detectable as counterfeit by an arithmetic test on a conserved sum. Everything there was a statement about a number holding still.

This section moves it. And the first thing to say is how few ways there are to do that.

```
Q     =  ( eta * F(u*)  -  c* ) / delta
 inf
```

Four symbols on the right. **Changing a destiny means changing `(eta, F, c, delta)` and nothing else** — not because other approaches are disapproved of, but because the expression has four arguments and a function of four arguments changes when one of them changes and at no other time. That is a definition rather than a promise, and its whole value is that it makes the question answerable: *which one is being moved, and by how much?*

A proposal that cannot be assigned to one of the four is not an intervention on the destiny. It may be many other worthwhile things. It is not that.

This section takes the four in turn, and for each one gives three things the practical literature almost never gives together: the frequency reading that says whether the lever is the binding one, the cost of moving it, and the specific way plans built on it fail. Two results govern the whole exercise and both are inherited. The first is the reserve theorem, which says a lever can be **absent** rather than merely expensive — demonstrated on a lever with nearly five times the elasticity of its rival. The second is the order rule, which says fix the numerator before lowering the loss, and which comes with a price for disobedience computed to six decimal places [Curzi, 2026q].

Then the section does what a section with this title has to do. There is a large, confident and commercially successful literature claiming that a destiny can be changed by believing in a new one. The framework's channel-gain postulate is the apparatus for adjudicating exactly that claim [Curzi, 2026ab], and it is applied here in full, with its conclusion in both directions: where a channel with real gain exists, belief moves the world by an amount the channel's own mechanism determines, and where no channel exists, nothing happens at all. The survivorship result supplies the companion number [Curzi, 2026].

The law

5.1 There are four arguments, and there is no fifth **THEOREM 5.1 — THE COMPLETENESS OF THE LEVER SET.** `Q_inf` is a function of exactly `(eta, F, c, delta)`. Therefore any change to `Q_inf` is a change to at least one of those four, and any proposal that changes none of them changes `Q_inf` by zero.

Proof. Inspection of `Q_inf = (eta F − c)/delta`. There are no other free symbols. ■

The proof is trivial and the consequence is not, because it is a closure statement: it says the list is complete, and a complete list is what turns a diagnostic into a procedure. The framework's intervention work states the same closure for interventions on the deficit and reaches a list of five handles rather than four [Curzi, 2026], and the two statements have to be reconciled explicitly or a reader will think one of them is wrong.

They are both right, and the difference is one level of nesting:

```
THE FOUR ARGUMENTS OF Q_inf          THE FIVE HANDLES ON AN ARRANGEMENT
---------------------------          ----------------------------------
eta   conversion                      L2  raise eta
F     surplus                         L1  raise u        }  the two routes
                                      L4  raise N        }  into F
c     cost                            L3  cut c
delta loss                            L0  lower delta
```

`F` is not directly manipulable. It is `ln(1 + (N−1)u)`, and it moves only when the geometry `u` moves or the category count `N` moves. So the four arguments give five handles, and the two lists are the same list read at two depths. The reconciliation matters because the two sublevers into `F` behave completely differently — one has the largest reserve in the apparatus and the other has almost none — and a plan that says "raise `F`" without saying which route has not yet said anything.

5.2 The partials, and why they may not be ranked At the carried dyad as found — case 2a, a constructed case, with `u = 0.068734665264634`, `eta = 0.30`, `c = 0.62`, `delta = 0.09`, `Q_inf = −5.0119395464`:

```
dQ_inf / d eta    =  F / delta              =   +6.256498   per unit of eta
dQ_inf / d F      =  eta / delta            =   +3.333333   per nat of surplus
dQ_inf / d c      =  −1 / delta             =  −11.111111   per unit of cost
dQ_inf / d delta  =  −(eta F − c)/delta^2   =  +55.688217   per unit of delta

the two routes into F:
dF / du           =  (N−1)/(1 + (N−1)u)     =   +6.263947   per unit of u
dF / dN           =  u /(1 + (N−1)u)        =   +0.039141   per added category
```

Ranking these against one another on size is forbidden, and the prohibition is not pedantic. They are in four different units. `eta` is a dimensionless efficiency in `[0, 1]`; `c` is a cost in the same units as `eta F`; `delta` is a rate in `(0, 1]`; `F` is in nats. The statement "`delta` has the largest derivative" is true and means nothing, because a unit of `delta` and a unit of `c` are not comparable quantities and no plan can spend one to buy the other.

Identifying which constraint actually binds, rather than which responds fastest, is the standard discipline of constraint analysis in operations [Goldratt and Cox, 1984, Hopp and Spearman, 2011]. What is legitimate is the elasticity comparison [Curzi, 2026q], which divides the units out. That is §5.9. What is also legitimate is the reserve comparison, which asks how far each lever can go before it runs out, and that is §5.10. The two answer different questions and the most expensive planning error here is confusing them.

5.3 The ceiling test comes before every lever Before any of this is worth computing there is a test that can rule the entire exercise out, and it belongs at the top of the procedure because a plan built after skipping it is a plan for a case that has no solution.

THE CEILING [Curzi, 2026]. Since `F <= ln N`, sustainability requires `c/eta <= ln N`. Equivalently `f_min <= k N^2`. At `N = 12`, `c/eta` may not exceed `2.484907`, i.e. `f_min` may not exceed `69.6133055503 Hz`.

The bound is the framework's normalisation carried into the frequency register, and it is the same bound requisite-variety arguments place on any regulator [Ashby, 1956, 1958].

```
c/eta > ln N   ->   NO value of u makes the arrangement sustainable.
                    Not a better pairing, not a better schedule, nothing.
                    The only remedies: raise eta, cut c, or add genuinely distinct categories.
```

Run on the four arrangements this framework carries, plus one constructed case built to fail it:

```
arrangement                 c/eta      ln N       verdict
-------------------------------------------------------------------------
carried dyad 2a           2.066667   2.484907    tunable in principle
carried dyad 2b           0.909091   2.484907    tunable, with margin
the commons               2.426471   2.484907    tunable in principle, barely
constructed family        1.400000   2.484907    tunable
CONSTRUCTED over-ceiling  3.000000   2.484907    STRUCTURAL -- no u closes it
```

The last row is worth working, because work that only showed tunable cases would be lying by selection. **CONSTRUCTED, labelled constructed.** An arrangement at `eta = 0.20` and `c = 0.60` sits at `c/eta = 3.000000`. Push its geometry to perfect orthogonality — `u = 1`, the theoretical maximum, two parties with no overlapping state whatever:

```
at u = 1, N = 12:   F = ln 12 = 2.484907   D = 2.885390 * (3.000000 - 2.484907) = +1.486245
```

Still failing, by nearly a octave and a half, with the geometry lever fully spent. The only remaining route through `F` is `N`, and the requirement is exact: `N >= e^3 = 20.0855`, so twenty-one categories.

```
at u = 1, N = 21:   F = ln 21 = 3.044522   D = 2.885390 * (3.000000 - 3.044522) = -0.128465
```

Sustainable, by an eighth of an octave, and only at perfect orthogonality in a twenty-one-category description. That is what "restructure rather than tune" looks like when it is priced. The carried commons is the same lesson at a milder setting: it passes the ceiling test by `0.058436` and is still short by `+6.463837` octaves, and its tuning procedure halts.

5.4 `delta` is a multiplier on a sign, not a lever on a gap This is the framework's ordering rule, and it is the correction to naive practice repeated more often than any other result, because it is the one every reader's intuition gets backwards.

THEOREM 5.2 — restated for `Q_inf`. `delta` appears in `Q_inf` only as a denominator. Therefore `sign(Q_inf) = sign(eta F - c)` for every `delta` in `(0, 1]`, and reducing `delta` scales `|Q_inf|` upward without ever changing its sign. ■

Computed on three drives, one negative, one positive, one exactly zero:

```
drive b              delta 0.30    delta 0.09    delta 0.05    delta 0.01
-------------------------------------------------------------------------
2a as found  -0.451075   -1.503582     -5.011940     -9.021491    -45.107456
2b repaired  +0.547598   +1.825328     +6.084427    +10.951968    +54.759841
exactly at hold   0.000   +0.000000     +0.000000     +0.000000     +0.000000
```

Three readings of that table. In the top row, the intervention that everyone reaches for first — build trust, add continuity, institute the ritual, keep the record — takes the destiny from `−5.011940` to `−45.107456`, **a ninefold deterioration, while every party involved is doing something admirable and nothing is going visibly wrong.** In the bottom row the same intervention does the same thing with the sign reversed, and it is the single largest available improvement. In the middle row it does nothing at all, at any magnitude, because trust multiplies an existing sign and zero has none.

And the deficit is flat across the entire table. `D = +4.338420` in the top row at every `delta`; `D = −3.590989` in the second at every `delta`. A pair reading only their deficit through that whole sequence would see a completely unchanging instrument. That is Theorem 4.4, and it is why the reporting format prints δ in the same block as D, every time.

5.5 The reserve theorem: a lever can be absent rather than expensive

This result is the most counter-intuitive in the framework's lever set.

THEOREM 5.3 — THE LEVER RESERVES. The maximum number of octaves each lever can supply, alone, from an operating point `(u_0, eta_0, c_0, N)`:

```
R_u    = 2.885390 * ( ln N − F(u_0) )                  ( u -> 1 )       the pitch headroom
R_eta  = 2.885390 * c_0 * ( 1/eta_0 − 1/eta_max )      ( eta -> eta_max )
R_c    = 2.885390 * ( c_0 / eta_0 )                    ( c -> 0 )
R_N    unbounded in principle; 2.885390 * u/(1+(N−1)u) per category in practice
```

Computed on the carried dyad as found, against its requirement of `+4.338420` octaves:

```
R_u    ( u -> 1 )        5.545206 octaves       CLOSES the deficit
R_eta  ( eta -> 0.40 )   1.490785 octaves       does not
R_eta  ( eta -> 0.50 )   2.385256 octaves       does not
R_eta  ( eta -> 0.75 )   3.577884 octaves       does not
R_eta  ( eta -> 1.00 )   4.174198 octaves       does not -- and eta_max = 1 is generous
R_c    ( c -> 0 )        5.963140 octaves       CLOSES the deficit
R_N    per added category at u = 0.068735 :  0.112937 octaves
```

Now set that beside the elasticities, which measure how fast each lever moves the deficit per *proportional* change and are therefore dimensionless and comparable:

```
point         c/eta      1 − 1/N_eff   |dD/d ln eta|  |dD/d ln u|   ratio
-------------------------------------------------------------------------
dyad 2a      2.066667      0.430550       5.963140      1.242306    4.8001
dyad 2b      0.909091      0.883938       2.623082      2.550507    1.0285
the commons  2.426471      0.169955       7.001314      0.490387   14.2771
```

At the carried dyad as found, `eta` has 4.80 times the elasticity of `u` — and 4.174198 octaves of reserve against a requirement of 4.338420, falling short by 0.164222. The lever that moves the deficit nearly five times faster per proportional unit cannot close it at all, even taken to `eta = 1`, which would mean an arrangement converting every unit of interaction into permanent structure with nothing lost anywhere.

The same fact appears in another form. The `eta` required to close the deficit is `c/F = 0.62/0.563085 = 1.101077`. A conversion efficiency above one is not an expensive lever; it is an absent one, and no amount of effort, budget or commitment reaches it, because it would mean creating more structure than the interaction supplied.

> **THE HONEST DYADIC LAW. The cost-side levers have the higher elasticity and the lower ceiling; the geometry lever has the lower elasticity and the only large reserve. Small deficits are closed on the cost side. Large ones require difference.**

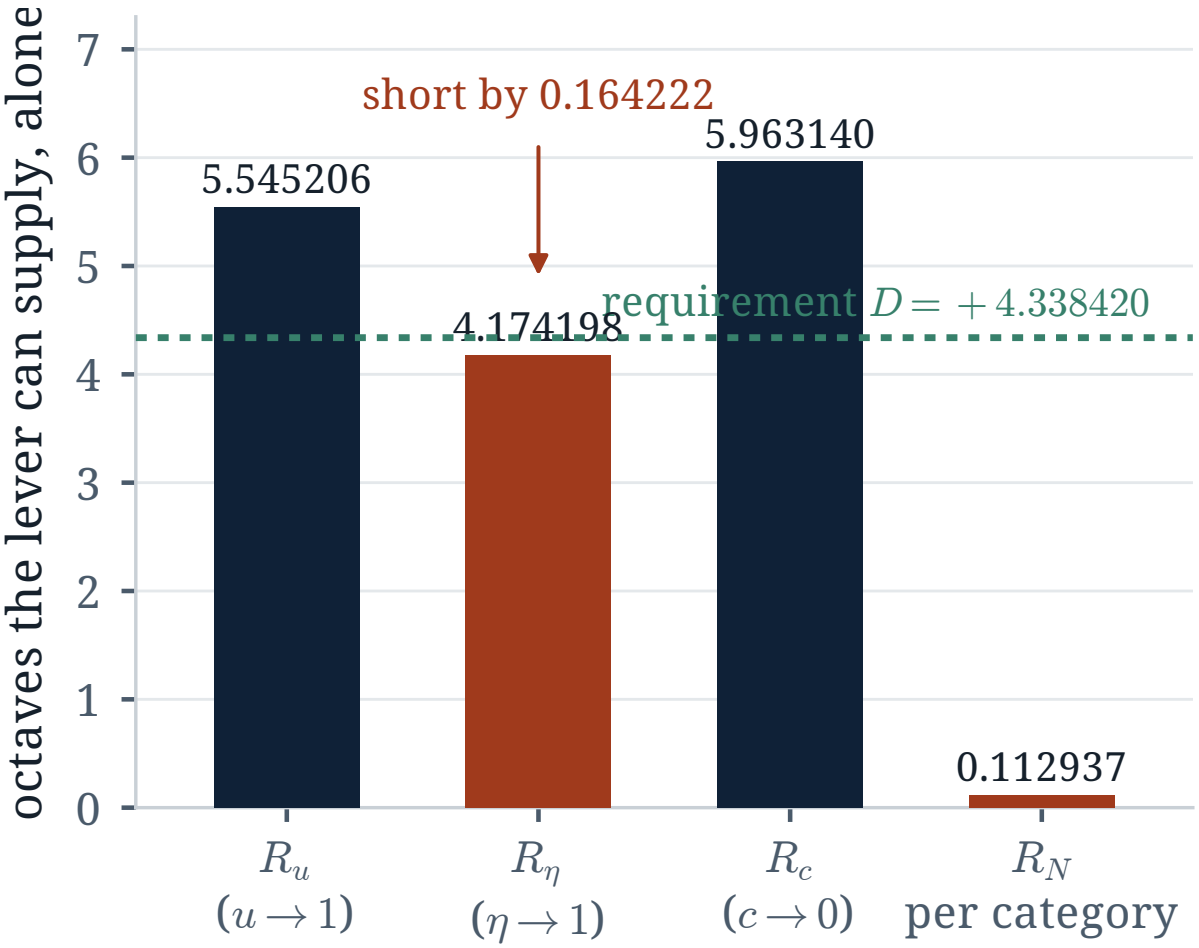


Figure 49.4: **A lever can be absent rather than expensive: the lever with 4.80 times the elasticity of its rival falls 0.164222 octaves short of the requirement even at an impossible conversion efficiency of one.** Constructed carried dyad as found (§5.5), $u = 0.068734665264634$, $\eta = 0.30$, $c = 0.62$, $\delta = 0.09$, $N = 12$, requirement $D = +4.338420$ octaves; reserves and elasticities are read from the two tables of §5.5. Script `figs/src/CCP-049_fig04.py`.

The reserve table also shows where the cost-side lever is worth having. `R_eta` is large only where `eta` is very low — where the arrangement is converting almost nothing and there is a great deal of headroom in the mechanism:

```
eta_0    c        R_eta (octaves, at eta_max = 1)
-------------------------------------------------
0.20     0.400          4.6166
0.20     0.825          9.5218
0.30     0.620          4.1742
0.40     0.620          2.6834
0.50     0.400          1.1542
0.50     0.825          2.3804
```

An arrangement already converting half of what it does has under two and a half octaves available on the `eta` side no matter how expensive it is. The cost-side lever is the cheap one, the first one, and the one that is spent quickly.

5.6 The order rule, and its price The design rule is one line: **raise η or cut c until $\eta F > c$, and only then invest in δ** [Curzi, 2026q]. What belongs here, because this is the section about changing things, is the price of doing it backwards.

CONSTRUCTED CASE, labelled constructed. The carried dyad, starting stock `Q_0 = 1.0`, eight periods of the wrong order first. Path A lowers `delta` from 0.09 to 0.05 and runs eight periods before fixing the numerator. Path B fixes the numerator first, runs eight periods at the original `delta`, and then lowers it. Both end at identical parameters and therefore, by Theorem 4.1, at the identical fixed point.

```
 t |   Path A   |   Path B   |    gap
---+------------+------------+-----------
 0 |  +1.000000 |  +1.000000 |  +0.000000
 1 |  +0.498925 |  +1.457598 |  +0.958673
 3 |  -0.429315 |  +2.252950 |  +2.682265
 5 |  -1.267052 |  +2.911581 |  +4.178633
 8 |  -2.373029 |  +3.693462 |  +6.066491
10 |  -1.073842 |  +4.401167 |  +5.475008
13 |  +0.641339 |  +5.335475 |  +4.694135
16 |  +2.111893 |  +6.136527 |  +4.024634
20 |  +3.751671 |  +7.029761 |  +3.278090
24 |  +5.087281 |  +7.757306 |  +2.670025
32 |  +7.061215 |  +8.832564 |  +1.771349

terminal fixed point, both paths          +10.951968
minimum stock reached, wrong order         -2.373029   (starting value +1.000000)
periods spent below the starting value     13   (t = 1 through t = 13)
permanent lag                              12 periods
cumulative gap over 32 periods            111.133601 stock-periods
```

The lag is permanent and it is exactly twelve periods at every level: Path B reaches `+3.693462` at `t = 8` and Path A first reaches it at `t = 20`; Path B reaches `+5.039870` at `t = 12` and Path A at `t = 24`; `+6.136527` at 16 and 28; `+7.029761` at 20 and 32. **Eight periods of misordering cost twelve periods of permanent delay and 111.133601 stock-periods of cumulative loss, from doing exactly the same two things in the opposite order.**

The excess over eight is not a rounding artefact. While the numerator is negative, the lowered `delta` is *accelerating* the loss, so those eight periods do more damage than eight periods at the original `delta` would have done.

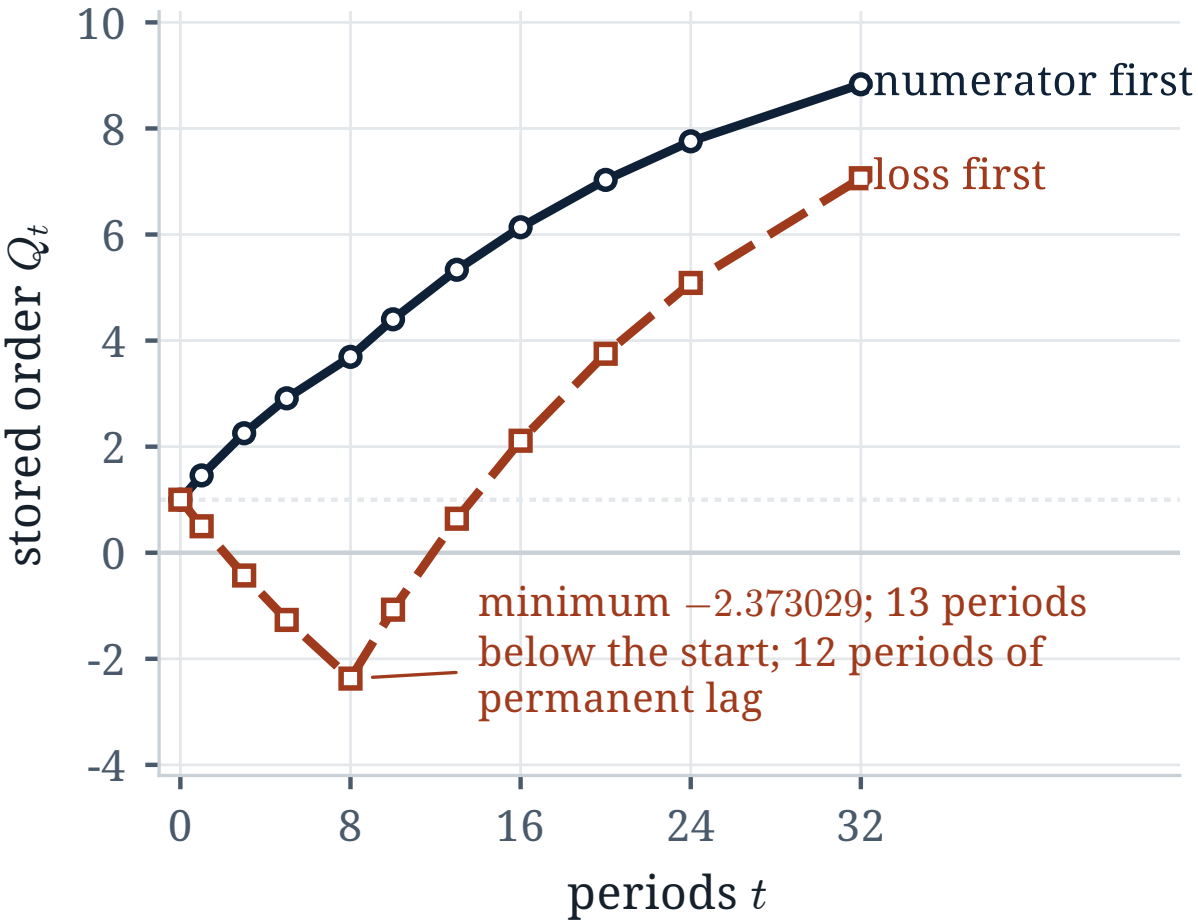


Figure 49.5: **The same two changes in the opposite order: identical destinations, a minimum stock of -2.373029, thirteen periods below the starting value and a permanent lag of twelve periods.** Constructed case of §5.6 from $Q_0 = +1.0$; Path A lowers δ from 0.09 to 0.05 and runs eight periods before fixing the numerator, Path B fixes the numerator first, and both end at $Q_\infty = +10.951968$. Every marked value is read from the order-rule table of §5.6. Script `figs/src/CCP-049_fig05.py`.

Two things about that table deserve to be read carefully. The first is that both paths converge to the same number, which means a planner who evaluates only the destination sees no difference between them whatever. The destination is order-independent; the trajectory is not; and the formalism will happily compute a path through negative stored order that a real arrangement may not survive. Admissibility fails there [Curzi, 2026a], and thirteen periods below the starting stock is thirteen periods in which people make decisions.

The second is what the instrument said the whole time. `D` reads `+4.338420` at the start of Path A and `+4.338420` at the end of Path A's first eight periods. Whatever the pair were using, if it was the deficit alone, it reported that nothing had happened during the most damaging phase of the plan.

5.7 Honesty clauses, restated where they bind

Essence, not resonance. Every hertz value in this section is a position on a logarithmic scale of surplus. Nothing here is claimed to be a physical vibration and nothing here is to be played at anyone.

No clinical, diagnostic or therapeutic claim. A lever prescription is a statement about four declared parameters of an arrangement. It is not treatment, not counselling, and not an assessment of anybody.

Consent is structural [Curzi, 2026]**.** No intervention derived from any reading in this section may reduce any party's ability to decline it, at any stage, for any reason or none. The geometry lever in particular describes changes to what people hold and do, and it is worth saying plainly: an intervention that raises u by arranging for someone else to change is not an intervention this framework licenses. The foundations concede that the case where the feedback loop runs through another person's will is an unsettled seam the corpus does not close [Curzi, 2026].

No subliminals. Nothing masked, nothing below threshold, nothing concealed [Curzi, 2026].

All cases here are CONSTRUCTED and labelled as such.

The engineering

5.8 Lever L1 — raise `u`, the geometry

The reading that says it binds. `R_u = 2.885390 (ln N − F(u_0))` exceeds the deficit while the cost-side reserves do not; equivalently, `1 − 1/N_eff` is small, meaning the arrangement is near parallel and has a great deal of unused pitch headroom. At the carried dyad as found, `1 − 1/N_eff = 0.430550` and `R_u = 5.545206` against a requirement of `4.338420`.

What it costs. The required `u` is exact: `u >= (e^(c/eta) − 1)/(N−1)`. For the carried dyad, `u >= 0.627132` from `0.068735`, a rise by a factor of `9.1240`. That is a large amount of real difference between two people currently at `cos theta = 0.965`.

What raising u actually is, and the trap in it. The unison result is the single most important thing to carry into any attempt at this lever:

```
u = 1 − (x . y)^2   depends only on |rho|

perfect agreement     rho = +1   ->   u = 0
perfect opposition    rho = −1   ->   u = 0
```

Manufactured disagreement raises cost and returns exactly nothing. A party who adopts contrary positions in order to bring difference moves the pair toward `rho = −1`, which reads `u = 0` just as agreement does, while adding coordination overhead: the reading gets worse in both terms at once. The contrarian and the sycophant sound the identical note, and it is the ground essence `k = 0.483426 Hz`.

The empirical literature on difference is genuinely divided and must be cited that way: bridging across structural holes has measured returns [Granovetter, 1973, Burt, 1992], creative success rises and then *falls* with the same variable [Uzzi and Spiro, 2005], short-run effects of surface diversity are frequently negative [Putnam, 2007], the rebuttal literature qualifies rather than reverses that [Abascal and Baldassarri, 2015], and the formal diversity-beats-ability result carries a standing objection that must be cited with it [Hong and Page, 2004, Thompson, 2014, Page, 2007]. **A monotone claim cannot survive that list unqualified**, and none is made here outside the range the arithmetic covers.

What does raise `u` is a change in what a party actually holds — states occupied that the other does not occupy, commitments taken on, work done in a domain the other is not in. **The test of a `u` intervention is whether it would still be true if nobody were reading the number.**

Failure mode. The representational one: two parties described by disjoint sets of states are automatically orthogonal and read `u = 1` regardless of content. A `u` raised by re-describing rather than by changing what is held is an artefact, and Gate 5-2 exists to catch it.

5.9 Lever L4 — raise `N`, the categories

The reading that says it binds. Only one: `c/eta > ln N`. Then the ceiling is the binding constraint, `u` is unavailable at every value, and adding categories is the only route into `F` that has anywhere left to go.

What it costs. Per added category, at fixed `u`, the gain is `2.885390 * u/(1+(N−1)u)` octaves. At the carried dyad that is `0.112937` octaves each; at the commons, where `u = 0.018614`, it is `0.044581` octaves each. **Adding categories to a near-parallel set buys almost nothing** — and the coordination cost of a larger category set grows as $n(n-1)/2$ [Brooks, 1975, Conway, 1968], while twelve already exceeds the classical span of immediate absolute judgment [Miller, 1956]. Computed, raising `N` alone on the carried dyad with `u` held at `0.068735`:

```
N     F          D
-------------------------
 12   0.563085   +4.338420
 16   0.708538   +3.918731
 20   0.835496   +3.552406
 25   0.974421   +3.151556
 40   1.303090   +2.203217
 60   1.620446   +1.287520
100   2.054730   +0.034441
200   2.686363   −1.788066   OK
```

Two hundred categories to close a deficit that one geometry change closes. **The lever is unbounded in principle and nearly useless in practice at low `u`, which is exactly the regime in which a planner reaches for it.**

Failure mode. Index inflation: the count rises after a re-description that added no distinguishable component [Curzi, 2026]. An added category must satisfy the framework's validity conditions and must carry variance the existing categories do not. Components that always co-occur are one component. And the rendering bounds bind if the reading is displayed: `N = 17` is the folded-lattice cliff and single-transposition audible capacity is `N <= 25`.

5.10 Lever L2 — raise `eta`, the conversion

The reading that says it binds. A large `c/eta` with `c` already lean, and a reserve check that passes: `D <= 2.885390 * c * (1/eta_0 − 1/eta_max)`. When that inequality holds, the cost side can close the gap alone. When it fails, difference is required, and no amount of work on conversion will substitute.

What it costs. The required `eta` is exact: `eta >= c/F`. At the carried dyad that is `0.62/0.563085 = 1.101077`, which is impossible. At the repaired dyad it is `0.40/2.153633 = 0.185733`, which is already exceeded. **The same lever is absent in one phase of the same constructed case and free in the other, and the only thing that changed between them was the geometry.** That is the order rule in a single sentence.

Failure mode, and it is specific to this lever. `eta` carries the entire cost side of the reading and its elasticity is the largest in the dyadic table, which means a verdict resting on it deserves an error bar rather than a decimal point. The framework's constructed family carries a deficit of `+0.584737` octaves with `eta` known to `± 0.05`, giving an uncertainty of `± 0.9734` octaves — and the sign of the verdict is not resolved. A deficit near zero is the most dangerous reading in the apparatus, because it is the one most likely to be an artefact of an imprecise input.

5.11 Lever L3 — cut `c`, the cost

The reading that says it binds. The largest reserve of the two cost-side levers, always: `R_c = 2.885390 (c/eta)` is by construction the entire required-octave figure, so cutting cost to zero always closes any deficit. That makes it the lever with the best arithmetic and the worst practical guarantee.

What it costs. The required cut is exact: `c <= eta F`. At the carried dyad, `c <= 0.30 * 0.563085 = 0.168925`, from `0.62` — a reduction of `72.754%` in coordination overhead. Reachable in principle; a large project in fact.

Failure mode, and it is the one that makes this lever's arithmetic misleading. Before cutting c, ask whether the cost being cut is performing conversion work. If it is, `eta` falls with `c` and the ratio `c/eta` may not improve at all — or may worsen. The decision log in the carried dyad's repair is a cost that raises `eta`; removing it to "cut overhead" would cut `c` from 0.40 and take `eta` back toward 0.30, and the reading would move the wrong way while the arrangement looked leaner. **`R_c` is computed at fixed `eta`, and cost cuts do not generally hold `eta` fixed.**

5.12 The thing that is not a lever — `delta`

`delta` is the fourth argument of `Q_inf` and it is not a lever on `D` at all (Theorem 5.2). It belongs in a plan, and it belongs there second.

Trust in the framework's sense enters as retention rather than as input [Curzi, 2026s, Gambetta, 1988, Putnam et al., 1993].

The reading that says it is ready. One condition only: `eta F > c`. Once the drive is positive, `delta` is the largest available improvement in the apparatus, and it is worth quantifying how large:

```
after the repair, drive b = +0.547598407
  delta 0.09 -> Q_inf  +6.084427
  delta 0.07 -> Q_inf  +7.822834      gain  +1.738408    half-life  9.5513
  delta 0.05 -> Q_inf +10.951968      gain  +4.867541    half-life 13.5134
  delta 0.03 -> Q_inf +18.253280      gain +12.168853    half-life 22.7566
  delta 0.01 -> Q_inf +54.759841      gain +48.675414    half-life 68.9676

the same moves before the repair, drive b = -0.451074559
  delta 0.09 -> Q_inf  -5.011940
  delta 0.07 -> Q_inf  -6.443922      change  -1.431983
  delta 0.05 -> Q_inf  -9.021491      change  -4.009552
  delta 0.03 -> Q_inf -15.035819      change -10.023879
  delta 0.01 -> Q_inf -45.107456      change -40.095516
```

The two columns are the same intervention producing gains of `+48.68` and `-40.10`, and the only thing that differs between them is the sign of a number that `delta` never touches.

Note also what the right-hand column of the first block costs. A destiny of `+54.76` comes with a half-life of `68.97` passes and a 99% approach time of `458` passes. At one pass per week that is nearly nine years to arrive. A deep destiny is a slow one, necessarily; the same parameter sets both, and there is no configuration in this model with a deep fixed point and a short approach.

Failure mode. In the refinement register [Curzi, 2026]: adding trust to a failing arrangement in order to make it improve faster does the reverse of both. It makes the approach slower and the destination worse.

5.13 Which lever binds — the decision arithmetic Put the reserves and the deficit side by side and the binding lever reads off directly. Computed across the four arrangements this framework carries:

```
arrangement            D          R_u        R_eta      R_c       R_N/cat   closes alone
-------------------------------------------------------------------------------------------
carried dyad 2a    +4.338420   5.545206   4.174198   5.963140   0.112937   u, c
carried dyad 2b    −3.590989      --         --         --         --      already sustainable
the commons        +6.463837   6.632448   4.620867   7.001314   0.044581   u, c
constructed family −0.516423      --         --         --         --      already sustainable
```

`R_eta` is computed at the generous `eta_max = 1`. In both failing arrangements the conversion lever cannot close the gap alone even at an impossible efficiency, and in both the two levers that can are the geometry and the cost.

The commons row is the instructive one. Its deficit is `+6.463837` octaves and **no single lever closes it in practice**, even though two of them close it in principle:

```
raise u to 0.95         -> deficit −0.0332    (only just, and 0.95 is near-orthogonality)
raise eta to 0.90       -> deficit +2.1075    still failing
cut kappa to 0.020      -> deficit +2.0085    still failing
```

`R_u = 6.632448` exceeds the requirement by `0.168611` octaves, which is to say the geometry lever closes it only in the limit of six people who share almost no state at all — a configuration that is not a commons. `R_c = 7.001314` closes it only at zero coordination cost, which is not a commons either. **Some systems cannot be tuned; they must be restructured, and the diagnostic is what says which case is in hand.**

5.14 The belief postulate, applied in full There is a large literature claiming that a destiny is changed by holding a different belief about it. The framework's channel-gain postulate is the apparatus for adjudicating that claim [Curzi, 2026ab], and it is applied here without softening in either direction.

> **[POSTULATE, badged] BELIEF CREATES REALITY IF THE BELIEF IS GROUNDED IN TRUTH.**

The mechanism the postulate formalises has a documented empirical form: expectation is known to change measured outcomes [Merton, 1948, Rosenthal and Jacobson, 1968], participation in a study itself changes behaviour [McCambridge et al., 2014], and a theory has been observed becoming true by being used, with dates and prices [MacKenzie and Millo, 2003, Ferraro et al., 2005]. It also has a measured *size*: the best long-run assessment finds classroom self-fulfilling prophecies real but small, usually non-accumulating and often dissipating [Jussim and Harber, 2005]. The conditional clause is what makes this a postulate that can be wrong rather than magical thinking, and the whole content sits in its formalisation:

```
BELIEF  ->  [ CHANNEL ]  ->  WORLD

"grounded in truth"  =  THE CHANNEL EXISTS AND IS REAL
```

> **[POSTULATE, badged]** A belief `B` affects a world-state `E` only through one or more channels. Each channel `i` has a gain `G_i` determined by the channel's own mechanism and not by the belief. Then `E = (sum_i G_i) B` for small `B` before saturation, and `|E| <= (sum_i G_i) B_max`. If no channel exists the sum is empty, `sum_i G_i = 0`, and `E = 0` at every `B`.

The badge stays on: this is a postulate, not a theorem. But the *bound* is arithmetic once the gains are estimated, and arithmetic is what makes a claim auditable in advance of any data.

The correspondence with the storage law is exact in form [Curzi, 2026ab]: `B` is the input where `S` was, `G` is the conversion where `eta` was, `E` is the output where the stock was. Where the correspondence stops, stated honestly: the storage law is a proven theorem about a recursion and the coupling relation is a postulate about the world. The form transfers; the warrant does not.

The four channel classes and their gains, as the framework's channel-gain paper reports them [Curzi, 2026ab]: physiological, roughly 0.15 to 0.45 on subjective outcomes and approximately zero on objective disease endpoints; behavioural, the widest range and typically 0.1 to 0.6 in the worked cases reported there; social, bimodal — small below a committed-fraction threshold and large above it; informational, zero to about 0.35 and entirely conditional on the target existing. Two readings of that list carry into this section. The behavioural channel dominates in every worked case, so the practical consequence of the postulate is not "believe harder" but: **find the action the belief produces, and check whether the action works.** And three of the four classes zero out on a fact about the world — no efferent pathway, ineffective action, absent target — which is where the conditional clause lives.

Composition. Series multiplies and parallel adds. A four-stage channel:

stage	gain	cumulative
belief registered as expectation	0.70	0.700000
expectation reaches the effector	0.50	0.350000
effector produces the action	0.62	0.217000
action lands on the outcome	0.35	0.075950

A chain of four individually respectable transmissions delivers 7.6% of its input, and the whole chain is bounded above by its weakest stage. Three such chains in parallel deliver `0.227850`.

5.15 The minimum viable gain, computed on this framework's arrangement

The channel-gain paper computes a minimum viable gain on its own worked group [Curzi, 2026ab]. Here is the same computation on the arrangement this framework carries, which makes the result usable inside the frequency register.

Take case 2b as the arrangement: `S = F = 2.153633`, `C = c = 0.40`, `delta = 0.09`, `N = 12`, `ln N = 2.484907`. Then a belief-driven channel of gain `G` is worth holding this period only if it clears the decay on the existing stock plus the maintenance cost:

```
G_min (growth this step)   =  (delta * Q_t + C) / S
G_min (headroom >= 0)      =  (delta * Q_t + C) / ln N
```

Q_t	G_min (growth)	G_min (headroom)
0.0000	0.185733	0.160972
1.0000	0.227523	0.197191
3.0000	0.311102	0.269628
6.0844	0.440000	0.381342
10.0000	0.603631	0.523158

Read the fourth row. At the fixed point, `G_min` equals `eta` exactly — `0.440000` against `eta = 0.44` — which is not a coincidence but the definition of a fixed point restated: the required gain at the destination is the gain that produced the destination. **The threshold rises with the stock, which means a belief channel sufficient to build an arrangement from nothing is not sufficient to maintain it once built.**

And the `G = 0` row deserves its own paragraph, because it is the arithmetic of an ungrounded belief:

```
no channel, F = 0, so the drive is  −C  and  Q_inf = −C/delta

C = 0.05, delta = 0.09  ->  Q_inf = −0.555556
C = 0.10                ->  Q_inf = −1.111111
C = 0.20                ->  Q_inf = −2.222222
C = 0.40                ->  Q_inf = −4.444444
```

A belief with no channel does not merely fail to build anything. Maintenance still bills, so **the fixed point is negative in proportion to what holding the belief costs, and the arithmetic says so without any need for moralising about it.**

Now the audit, run backwards, against the specific claim this section exists to adjudicate: *can believing differently close the carried dyad's 4.338420-octave deficit?* The question is well-posed, because the deficit's four inputs are the only routes:

```
through u    : u must reach 0.627132 from 0.068735, a factor of 9.1240
               -- a change in what the parties actually hold and do
through eta  : eta must reach c/F = 1.101077
               -- ABSENT. No belief reaches an efficiency above one.
through c    : c must fall to eta F = 0.168925, a cut of 72.754%
               -- a change in what the arrangement costs to run
through N    : ~200 categories at fixed u, or a real geometry change first
```

Two of the four routes are changes in what people do, one is arithmetically unreachable, and one is a change in a description that must satisfy the framework's validity conditions. So the postulate gives a clean and deflationary answer: a belief that produces the behavioural change — the states actually occupied, the commitments actually taken on, the decision log actually kept — moves the destiny by exactly the amount that change is worth, which is computable in advance. A belief that produces no such change has an empty channel sum and moves the destiny by zero, while continuing to bill its maintenance cost.

5.16 What testimonials are worth The companion result is the survivorship arithmetic [Curzi, 2026], and it is the sharpest thing in this section.

A corpus of testimonials is a record of successes, published because they were successes. The modern statement of the same defect is the selective-publication literature [Ioannidis, 2005, Simmons et al., 2011, Gelman and Loken, 2013], and its measured consequence is that a large replication programme recovers far less than the original record implies [Open Science Collaboration, 2015]. The probability that a published account reports a success is therefore 1, under every hypothesis about how often the method works. The likelihood of the observed corpus is the product of those probabilities, and the likelihood ratio between any two hypotheses is that product divided by itself.

```
P(account reports a success | published, p = 0.900000)  =  1.000000
P(account reports a success | published, p = 0.010000)  =  1.000000
P(account reports a success | published, p = 1.0e-09 )  =  1.000000

likelihood ratio, at any corpus size n:                    1.000000
log10 LR                                                   0.000000
bits of evidence supplied                                  0.000000
```

A likelihood ratio of exactly 1.000000 means the evidence does not move the posterior at all. Not weakly, not a little — not at all. **A thousand accounts from people for whom a method**

worked carry precisely the weight of zero accounts, because the selection guarantees that successes are observed whether or not the method does anything.

The precedent is Abraham Wald's wartime work on aircraft armour — the bullet-hole data were conditioned on the aircraft returning, so the places with no holes were the places to armour. **That work is not in this paper's bibliographic database and is named here without a citation key.** A manifestation corpus is the returning aircraft. It is a census of survivors presented as a census of attempts.

Three operating consequences, and they are the reason this framework's protocols look the way they do. **The base rate is unobtainable from testimonials in principle, at any corpus size** — a thousand more accounts add a thousand numerators and no denominator, so the problem is not sample size and cannot be fixed by more of the same. Getting a denominator requires enrolling before the outcome is known, which is a cohort design and is why preregistration is mandatory rather than advisable [Nosek et al., 2018]. The framework's own preregistered fit is the relevant precedent [Curzi, 2026i]: it failed its own second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin, despite $+40.3$ per cent over persistence, $+97.7$ per cent over direct input, and $p = 0.0005$ against the shuffle. And the contrast is enormous once the denominator exists: a constructed cohort of 100 enrolled with 90 successes yields a likelihood ratio of `9.66e+15` between `p = 0.9` and `p = 0.5`, against `1.000000` for any number of testimonials whatever.

The companion result completes the picture and it is the practical rule. Likelihood ratios measured at `2.0000` for a prediction specifying only a DIRECTION versus `11.3577` for one specifying a MAGNITUDE: **a prediction with a number attached is about 5.68 times more evidentially powerful, and a belief that predicts only a direction is not yet testable.** That is the precise justification for the requirement that a lever plan state its target value before it is executed.

The specification

5.17 PROCEDURE CHANGE — planning a change of destiny **Step 1. Produce a complete destiny reading first** (§4.17, PROCEDURE DEST, all ten steps). No lever may be discussed before `D`, `delta`, `Q_inf` and the export test are on the page.

Step 2. Run the ceiling test. If `c/eta > ln N`, stop. The case is structural. Report it as structural and go to Step 9.

Step 3. Compute all four reserves — `R_u`, `R_eta`, `R_c`, `R_N` per category — against the deficit. Report which levers close it alone and which do not.

Step 4. Compute the elasticities and report them separately from the reserves, with the statement that they answer different questions.

Step 5. Check the sign of the drive. If `eta F <= c`, no `delta` intervention may appear anywhere in the plan until a later step has made the drive positive. This is a step constraint, not advice.

Step 6. Assign every proposed action to exactly one handle — `u`, `N`, `eta`, `c`, `delta` — and state the target value with a number attached. A proposal that cannot be assigned moves nothing; a proposal without a target value is not yet testable (the direction-versus-magnitude result).

Step 7. For every action whose mechanism runs through what somebody believes, run the coupling audit. Name the channel, classify it, estimate its gain from the channel's own mechanism and not from the outcome, compose series and parallel, and compare the summed gain against the gain the claim requires. Report the audit whether it passes or fails.

Step 8. Test cost cuts for conversion damage. Before cutting `c`, ask whether the cost being cut performs conversion work. If it does, `eta` falls with it and the ratio may not improve.

Step 9. Order the actions: numerator first, loss second. Publish the projected trajectory as well as the projected fixed point, and check that the trajectory does not pass through negative

stored order.

Step 10. Re-read after the intervention and publish the before-and-after in the §4.19 reporting format, including the export test, including `delta` alongside `D`, and including the cases where the plan did not work.

5.18 Gates

```
GATE 5-1   destiny reading complete before any lever discussed      FAIL -> stop
GATE 5-2   u interventions checked against the one-hot artefact     FAIL -> re-elicit
GATE 5-3   ceiling test run first                                   FAIL -> incomplete
GATE 5-4   reserves AND elasticities both reported, separately      FAIL -> incomplete
GATE 5-5   no delta action scheduled while the drive is negative    FAIL -> reorder
GATE 5-6   every action assigned to one handle, with a target value FAIL -> not a plan
GATE 5-7   coupling audit run on every belief-mediated action       FAIL -> unaudited
GATE 5-8   cost cuts tested for conversion damage                   FAIL -> incomplete
GATE 5-9   trajectory published, not only the fixed point           FAIL -> incomplete
GATE 5-10  consent re-asked at every stage that changes the action  FAIL -> halt
```

The failure modes

FM 5.1 — THE UNASSIGNED PROPOSAL. *Signature:* an intervention that cannot be attached to `u`, `N`, `eta`, `c` or `delta`. *Mechanism:* Theorem 5.1. `Q_inf` has four arguments and a proposal touching none of them moves it by zero. *Fix:* Gate 5-6.

FM 5.2 — TRUST FIRST. *Signature:* "we will build trust and things will improve." *Mechanism:* the ordering rule. In the constructed case this takes `Q_inf` from `-5.011940` to `-45.107456` while `D` does not move at all. *Fix:* Gate 5-5.

FM 5.3 — THE STRONGEST LEVER ASSUMED SUFFICIENT. *Signature:* a plan built on the largest elasticity. *Mechanism:* Theorem 5.3. At the carried dyad `eta` has 4.80x the elasticity and 4.174198 octaves of reserve against a 4.338420 requirement. *Fix:* Gate 5-4.

FM 5.4 — THE ABSENT LEVER PRICED AS AN EXPENSIVE ONE. *Signature:* a budget for raising `eta` to a required `1.101077`. *Mechanism:* a conversion efficiency above one would mean more structure created than interaction supplied. *Fix:* compute `c/F` before planning the lever.

FM 5.5 — MANUFACTURED DISAGREEMENT. *Signature:* a party adopting contrary positions to raise `u`. *Mechanism:* `u` depends on `|rho|`; opposition reads `u = 0` exactly as agreement does, and the cost rises. *Fix:* §5.8's test — would it still be true if nobody were reading the number?

FM 5.6 — `u` RAISED BY RE-DESCRIPTION. *Signature:* a reading of `u = 1.000000` after a re-elicitation. *Mechanism:* disjoint support is automatically orthogonal. *Fix:* Gate 5-2.

FM 5.7 — CATEGORIES ADDED TO A NEAR-PARALLEL SET. *Signature:* a plan to raise `N` in a system at low `u`. *Mechanism:* `0.112937` octaves per category at the carried dyad and `0.044581` at the commons; two hundred categories to close what one geometry change closes. *Fix:* raise `u` first, or accept that the case is structural.

FM 5.8 — INDEX INFLATION. *Signature:* `N` rises after a re-description that distinguishes nothing new. *Mechanism:* components that always co-occur are one component [Curzi, 2026]. *Fix:* the framework's validity conditions, applied to each added category.

FM 5.9 — THE COST CUT THAT CUTS CONVERSION. *Signature:* overhead removed and the ratio `c/eta` unimproved or worse. *Mechanism:* `R_c` is computed at fixed `eta` and cost cuts do not generally hold `eta` fixed. *Fix:* Gate 5-8.

FM 5.10 — THE MARGINAL VERDICT ACTED ON. *Signature:* a plan launched on a deficit of `+0.584737` with `eta` known to `± 0.05`, an uncertainty of `± 0.9734` octaves. *Mechanism:* the sign is unresolved. *Fix:* a verdict is resolved only when the deficit exceeds its own uncertainty.

FM 5.11 — THE DESTINATION EVALUATED AND THE PATH IGNORED. *Signature:* "the fixed point does not depend on the order, so the order does not matter." *Mechanism:* true and irrelevant; the constructed case shows 12 periods of permanent lag, 13 periods below the starting stock, and 111.133601 stock-periods of cumulative loss. *Fix:* Gate 5-9.

FM 5.12 — BELIEF WITHOUT A CHANNEL. *Signature:* an intervention whose mechanism, when asked for, is that someone will believe it. *Mechanism:* the coupling requirement; an empty channel sum gives `E = 0` at every `B`, while `C` continues to bill, so `Q_inf = −C/delta`. *Fix:* Gate 5-7.

FM 5.13 — THE GAIN FITTED TO THE OUTCOME. *Signature:* a channel gain estimated from the result it is supposed to predict. *Mechanism:* the apparatus is unidentifiable unless `G` comes from the channel's own mechanism [Curzi, 2026ab]. *Fix:* estimate `G` without touching the outcome, then predict.

FM 5.14 — THE MAGNITUDE ERROR. *Signature:* an effect promised that is larger than every named channel summed at maximum belief. *Mechanism:* the bound of 77.1. A four-stage series channel of individually respectable stages delivers `0.075950`. *Fix:* run the audit before collecting data.

FM 5.15 — TESTIMONIALS AS EVIDENCE. *Signature:* a lever recommended on the strength of accounts from people it worked for. *Mechanism:* the survivorship arithmetic [Curzi, 2026]. The likelihood ratio is exactly `1.000000` at any corpus size. *Fix:* a cohort with the denominator observed, preregistered.

FM 5.16 — THE DIRECTION-ONLY PREDICTION. *Signature:* "this will improve things." *Mechanism:* direction-only predictions carry `LR = 2.0000` against `11.3577` for a prediction with a magnitude. *Fix:* Gate 5-6's target value.

FM 5.17 — A LEVER PULLED ON SOMEBODY ELSE. *Signature:* a plan that raises `u` by requiring another party to change. *Mechanism:* the consent law [Curzi, 2026]; the lever is real and the procedure is not available. *Fix:* Gate 5-10, and §5.7.

Worked examples

5.19 The four handles, one at a time, on the carried dyad **CONSTRUCTED. Not an observation of any person or pair.** The pair start as found and each row adds one handle to the row above it, in the order the ordering rule requires [Curzi, 2026q].

step	u	eta	c	F	D	Q_inf	f_rel
0 as found	0.068735	0.30	0.620	0.563085	+4.338420	−5.011940	1.4908
1 geometry only (u)	0.692373	0.30	0.620	2.153633	−0.250931	+0.289887	35.8882
2 + conversion (eta)	0.692373	0.44	0.620	2.153633	−2.148294	+3.639982	35.8882
3 + cost (c)	0.692373	0.44	0.400	2.153633	−3.590989	+6.084427	35.8882
4 + trust (delta .09 -> .05)	0.692373	0.44	0.400	2.153633	−3.590989	+10.951968	35.8882

(Q_inf at delta = 0.09 for rows 0–3; row 4 is the same arrangement at delta = 0.05)

Read the columns. The geometry step alone flips the sign of both readings — a deficit of `+4.338420` becomes `−0.250931` and a destiny of `−5.011940` becomes `+0.289887` — and **it is the only step in the table that changes `f_rel` at all.** The two cost-side steps then buy `3.34` octaves of deficit and `5.79` units of destiny between them, and they are only available because the drive is already positive. The trust step buys `4.868` units and moves the deficit by nothing, which is Theorem 5.2 appearing on the page.

Now the same four handles in the order the intuition prefers:

```
step                          u          eta    c      delta   D           Q_inf
-------------------------------------------------------------------------------
0  as found                 0.068735   0.30   0.620   0.09    +4.338420    -5.011940
1  trust first              0.068735   0.30   0.620   0.05    +4.338420    -9.021491
2  trust harder             0.068735   0.30   0.620   0.02    +4.338420   -22.553728
3  trust absolutely         0.068735   0.30   0.620   0.005   +4.338420   -90.214912
```

Four rows of an unchanging instrument and an eighteenfold deterioration in the destiny. Every step in that table is an intervention a reasonable person would call good practice, and every step makes the arrangement worse, and the deficit — the framework's own central instrument — reports nothing at any point. **That is the single strongest argument in this framework for publishing `delta` alongside `D`, and it is why Gate 4-4 exists.**

5.20 The commons: the procedure halts CONSTRUCTED. Not an observation of any group.

Six members drawn from one doctrine, 15 pairs, mean pairwise `u = 0.018614`, cost `kappa = 0.055` per link, `eta = 0.34`, `delta = 0.09`.

```
reading      F = 0.186272 nats   N_eff = 1.204754   f_rel = 0.7017 Hz
             f_min = 61.9349 Hz  D = +6.463837 octaves   Q_inf = -8.4630   DECAYS
ceiling      c/eta = 2.426471 vs ln N = 2.484907  -> passes by 0.058436
reserves     R_u = 6.632448  R_eta = 4.620867  R_c = 7.001314  R_N/cat = 0.044581
single-lever raise u to 0.95    -> deficit -0.0332   at near-orthogonality
             raise eta to 0.90  -> deficit +2.1075   still failing
             cut kappa to 0.020 -> deficit +2.0085   still failing
VERDICT      PROCEDURE HALTS. No single lever closes this in practice.
```

The reserves say two levers close it in principle. The single-lever table says what those in-principle closures actually require: six people who share almost no state, or coordination at near-zero cost. Neither describes a commons. And the roster ceiling is the reason: at `m(m-1) = 30` against a bound of `2 eta ln N / kappa = 30.722482` — inside it by 0.722482 — a seventh member takes the group to 42 and makes the case structural outright.

This is the case the framework keeps because it does not work. The lesson is not that the group should disband — the formalism prices stored order and nothing else, and a family, a team or a commons is for things the formalism does not represent. The lesson is that a lever plan for this arrangement would have been a plan for a case with no solution, and the ceiling test and the reserve table are what say so before anybody spends a year on it.

5.21 The coupling audit, run on a lever claim CONSTRUCTED. Not an observation.

Take the claim: *the pair will close their 4.338420-octave deficit by committing to a shared vision of the relationship they want.*

```
Step 6  ASSIGN THE HANDLE.  The proposal names no handle. Ask which of u, N, eta, c, delta
        it moves. Two answers are possible and they audit completely differently.

ANSWER A -- "it moves what we each take on."  -> handle L1 (u). Testable. Target value
        required: u >= 0.627132. Channel: behavioural, the dominant class. Gain estimated
        from behaviour records, not from the outcome. AUDITABLE, and Gate 5-6 satisfied
        once the target is stated.

ANSWER B -- "it moves how we feel about it."  -> no handle. The four arguments of Q_inf do
        not contain a term for how an arrangement is regarded. Channel sum is empty:
        E = 0 at every B. And the arrangement still bills: with F unchanged at 0.563085,
        eta 0.30, c 0.62, delta 0.09, the destiny remains -5.011940 exactly.
```

```
Step 7  THE BOUND.  Required change in the drive to reach Q_inf = 0:  +0.451075 per pass.
        Available from Answer B's named channels, summed:              0.000000
        VERDICT: over budget by the entire amount, before any data are collected.
```

That is the postulate doing its work in both directions in one example. Answer A is a real channel with a real gain and the postulate says it moves the world by exactly what the channel can carry. Answer B has no channel and the postulate says nothing happens — and says it as arithmetic rather than as disapproval.

5.22 What this section establishes, and what Section 6 takes **Established, and checkable.** That `Q_inf` has exactly four arguments and five handles, with `F` reached only through `u` and `N` (Theorem 5.1); that `delta` cannot change the sign of a destiny at any magnitude, so it is a multiplier on an existing sign rather than a lever on the gap (Theorem 5.2); that a lever can be absent rather than expensive, demonstrated on a lever with 4.80 times the elasticity of its rival and 4.174198 octaves of reserve against a 4.338420 requirement (Theorem 5.3); that reversing the ordering rule costs 12 periods of permanent lag, 13 periods below the starting stock, and 111.133601 stock-periods of cumulative loss in the constructed case; that an arrangement over the `c/eta <= ln N` ceiling is structural and no geometry closes it, worked at `c/eta = 3.000000` where even perfect orthogonality leaves `+1.486245` octaves and twenty-one categories are required; that a belief moves a destiny only through a channel whose gain is estimable from its own mechanism, with an empty channel giving `Q_inf = −C/delta`; and that testimonial evidence carries a likelihood ratio of exactly `1.000000` at every corpus size.

Not established, and not claimed. That any lever prescription here applies to any actual person, pair or group; that a larger fixed point is a better outcome; that belief has any effect through any channel not named and estimated; that any reading here is diagnostic, therapeutic or clinical; or that the frequencies printed here are physical resonances.

The one line to carry forward: *there are four arguments, the reading says which one binds, and the ones that are absent are absent no matter how much anybody wants them.*

Section 6 states the keystone once more in full — $f_{\min} = k\, e^{2c/\eta}$, the 2.885 octaves per unit of cost ratio, and the $c/\eta \le \ln N$ ceiling — and sets beside the framework's results the six findings that constrain them.

6 The sustaining frequency, and the six results that constrain it

This section closes the argument. It has three jobs and they are in a deliberate order.

The first is to state the keystone once more, in full, with nothing left implicit — because this series of papers exists to join two halves of one equation that were laid down separately, and gesturing at the join instead of writing it out would settle nothing.

The second is to draw the arc across the framework honestly: reading a gradient, building a circumstance, and then tuning the relation between them. That arc is not a story of accumulating power. At each stage the instrument got sharper and the set of things it was entitled to say got smaller, and both halves of that sentence belong in the summary.

The third is to set beside the result the findings that constrain it. There are six of them and every one is a limitation. A surplus law that turns out to be a nineteenth-century psychophysical law with a threshold; a consonance that turns out to be redundancy; an interaction parameter that cannot tell perfect agreement from perfect opposition; a group whose whole surplus is available to a well-chosen pair; two frequency readings that cannot be inferred from each other; and a central diagnostic that is blind to decay. **Those six are the real yield, and each of them narrows what the instrument can be used for.** A summary that listed the affordances and omitted the constraints would be selling something.

Then the section stops on the open problem stated in the companion paper [Curzi, 2026], which is not solved, not softened, and not worked around. **The restraint is the argument.** Work that names its own load-bearing gap this precisely earns the right to everything else in it; work that declines to name it forfeits the right to all of it, including the parts that are correct.

The law

6.1 The keystone, in full The two halves were established separately and never multiplied.

```
THE ESSENCE BRIDGE          f  =  k ∗ e^(2F)      k = 0.483425732988414 Hz

THE STORAGE FIXED POINT     Q     =  ( eta ∗ F(u∗) − c∗ ) / delta
                             inf
```

The bridge says an interaction's surplus has an exact pitch. The fixed point says where a relation, a life or an institution arrives if its parameters hold. Sustainability means the fixed point is positive:

```
Q     >  0     <=>     eta ∗ F  >  c     <=>     F  >  c / eta
 inf
```

And the bridge inverts. If `f = k e^(2F)` then `F = (1/2) ln(f/k)`. Substitute:

```
(1/2) ln( f/k )  >  c/eta        =>        f  >  k ∗ e^( 2c / eta )
```

which is the keystone, and it is a substitution rather than a postulate — given the two halves, this line follows and nothing was added to obtain it.

THE MINIMUM SUSTAINABLE FREQUENCY

```
f     =  k ∗ e^( 2c / eta )
 min

f > f_min     the relation ACCUMULATES order          Q_inf > 0
f = f_min     it HOLDS, and nothing more             Q_inf = 0
f < f_min     it DECAYS, whatever anyone intends     Q_inf < 0
```

The condition for a relation to pay for itself is not a ratio, not a policy, and not a virtue. It is a pitch.

In octaves above the ground essence `k`, the statement becomes an exchange rate:

```
octaves required  =  2 (c/eta) / ln 2  =  2.885390 ∗ (c/eta)
```

Every unit of cost-to-conversion ratio costs 2.885390 octaves of required pitch. Computed:

```
 c/eta      f_min (Hz)      octaves above k
 0.000        0.483426          0.000000
 0.250        0.797034          0.721348
 0.500        1.314087          1.442695
 0.750        2.166564          2.164043
 0.909091     2.978216          2.623082
 1.000        3.572060          2.885390
```

```
1.500        9.709865          4.328085
2.000       26.394151          5.770780
2.066667   30.158790          5.963140
2.484907    69.613354          7.169926   <- the ceiling at N = 12
3.000      195.027860          8.656170
4.000     1441.071800         11.541560
```

6.2 The ceiling The third component of the keystone is a bound, and it is the ceiling result in the frequency register [Curzi, 2026]. Since `F <= ln N`, the maximum sustainable cost ratio is capped:

```
c / eta  <=  ln N            equivalently        f     <=  k * N^2
                                                  min
```

At twelve categories, `c/eta` may not exceed `2.484907` and `f_min` may not exceed `69.6133055503 Hz`. Above that, **no arrangement of those twelve categories is sustainable** — not a better pairing, not a better schedule, not more effort. The only remedies are: raise `eta`, cut `c`, or add genuinely distinct categories.

N	ln N	max c/eta	f_max = k N^2 (Hz)	octaves
2	0.693147	0.693147	1.933703	2.000000
3	1.098612	1.098612	4.350832	3.169925
4	1.386294	1.386294	7.734812	4.000000
6	1.791759	1.791759	17.403326	5.169925
12	2.484907	2.484907	69.613306	7.169926
17	2.833213	2.833213	139.710037	8.174926
25	3.218876	3.218876	302.141083	9.287712
41	3.713572	3.713572	812.638657	10.715104
100	4.605170	4.605170	4834.257330	13.287712

The ceiling is the most useful single line in the apparatus for the purposes of refusing work. It converts a class of problems from difficult into unavailable, and it does so before any effort is spent. A test that can say "no arrangement of these categories works" in one division is worth more to a planner than any number of tests that say "try harder."

6.3 What the keystone unifies Three laws that entered the corpus from three different directions land on one threshold.

The essence law gives a state a frequency: $f = kn^2$, with $k = (9/8)^2/\varphi^2 = 0.483425732988414$ Hz. The constant is **defined**, not measured — an exact ratio of the Pythagorean whole tone [Barbour, 1951] to the square of Euclid's extreme and mean ratio [Euclid, 1956, Vajda, 1989] — and its fifteen figures are exact digits of a definition. **Attaching hertz to it is a separate and much stronger claim and is postulated.** The deflationary literature on golden-section claims is cited rather than avoided [Markowsky, 1992, Fischler, 1981, Höge, 1997]: φ is being *chosen* as a constant here, not *found* in anything measured. The square law is likewise a postulate whose exponent is **defined by the model, not estimated**; no fit was performed and the standard procedure for testing a claimed scaling relation [Clauset et al., 2009] was not run, because there is no data to run it on. Both carry their badges here as everywhere.

The surplus law gives an interaction a magnitude: $F = \ln(1 + (N-1)u)$, a theorem under the framework's own axioms, bounded above by $\ln N$.

The storage law gives an arrangement a destination: `Q_{t+1} = (1-delta)Q_t + eta F - c`, with the fixed point `Q_inf = (eta F - c)/delta`, a proved property of a linear recursion.

The bridge `f = k e^(2F)` is what makes them commensurable, because it puts the surplus and the frequency in the same register — the logarithm — and once they are in the same

register the sustainability condition of the third law can be written as a threshold in the units of the first.

The join is arithmetic, and it is the whole of this material's contribution. What it buys is that a question about whether an arrangement pays for itself becomes a comparison of two numbers in hertz, and both of them are computable from measurements a practitioner can actually take.

6.4 The diagnostic, and the two numbers that must travel together

```
octave deficit   D  =  2.885390 * (c/eta)  -  log2( f/k )
                   =  2.885390 * ( c/eta  -  F )

D <= 0    sustainable as it stands
D >  0    short by that many octaves; four levers: raise u, raise eta, cut c, raise N
```

And the identity that Section 4 owns:

```
D  =  - 2.885390 * ( delta / eta ) * Q
                                     inf
```

verified to `8.88e-16` across twelve constructed parameter sets. **The octave deficit and the fixed point are the same number in different units.**

But `D` contains no `delta`. **The deficit reads the SIGN and the SIZE of the gap; `delta` converts it to a LEVEL; and a reading that reports one without the other is incomplete.** The carried dyad as found reads `D = +4.3384201973` at `delta = 0.30`, at `delta = 0.09`, at `delta = 0.05` and at `delta = 0.01`, while its destiny moves from `-1.5035819` to `-45.1074559` across the same column. This is a genuine limitation of the framework's central instrument and it is stated here rather than left in the section that derived it.

6.5 What the keystone does not say Five refusals, and they are not decorative. Each one marks a place where the formalism has been mistaken for something larger.

It asserts no physical resonance. Every hertz value in three books is a position on a logarithmic scale of surplus. No person, bond, tissue or molecule is claimed to vibrate at any value anywhere in this corpus, and nothing here is to be played at anyone.

It makes no clinical, diagnostic or therapeutic claim. The twelve states are not diagnoses, a reading is not an assessment, and no procedure here treats anything. A reader who needs support for a relationship or for their health should seek it from someone qualified to give it, which this framework is not.

It confers no capability and edits no DNA. The companion paper specifies an encoder, an addressing convention, a sonifier and a rendering discipline [Curzi, 2026]. It specifies a research procedure and labels it as one.

It hides nothing. No subliminals, nothing masked, nothing below threshold. Open-label placebo works, so disclosure costs nothing that matters.

It removes nobody's ability to decline. Consent is structural, re-asked at every stage that changes what is done, and no reading in this framework licenses an intervention on a person who has not agreed to it.

And one more, which is about the scale rather than the ethics. **The keystone inherits exactly the defensibility of its inputs.** The index `n` is a count and is ratio-scaled by construction, so it sidesteps the scale demotion [Curzi, 2026m]; the intensities `w` are ordinal by default and must be magnitude-estimated before any relationship number is printed, because an ordinal-preserving relabelling moves `u` from `0.362115` to `0.525275` — 1065.05 cents, nearly an octave, from a change an ordinal scale calls no change at all. Where the intensities were not magnitude-estimated, the relationship frequency is not determined and no number may be printed.

6.6 Three through-lines, and what each one costs Before the arc, three statements that hold in every volume of the series and that a reader should carry away even if nothing else survives. Each is a genuine finding and each has a price attached that the popular version of the same idea leaves off.

The loss term is where everything happens. `delta` appears in no diagnostic and decides every level. It sets the half-life, it sets the depth of the fixed point, it sets how long a repair takes to show, and it is the parameter the octave deficit cannot see. Three books of interventions reduce, in the end, to a question about a leak: how fast does an arrangement forget what it made? The price of this through-line is the one §6.12(f) names — because `delta` is absent from `D`, the instrument that reads sustainability is silent about the parameter that decides the level, and there is no version of the reading in which one number suffices.

Difference generates surplus, and only real difference does. `F = ln(1 + (N−1)u)` is increasing in `u`, so a relation's stored order comes from the non-parallelism of what its parties hold. That is the affirmative half and it is genuinely useful: it says that a bond, a team or an institution whose members converge on one description is losing its capacity to produce anything, and it says so with a number. The price is Result (c). `u` depends on `|rho|` alone, so opposition reads exactly as agreement does, and the difference that raises `u` has to be difference in what is actually held and done — states occupied, commitments taken on, work done in a domain the other is not in. The test of a `u` intervention is whether it would still be true if nobody were reading the number, and no rearrangement of the arithmetic will substitute for that.

Geometry does work that power does not. Every result in the corpus about arrangement — the table with no head, the silo, order without a centre, the roster ceiling, the dyad's interval — is a claim that the shape of a set of relations changes what it can store, independently of how much is put into it. The keystone is the compressed form: `f_min` depends on `c/eta` and on nothing else, so an arrangement that converts better or costs less needs less pitch, and an arrangement that converts badly cannot be rescued by volume. The price here is Result (d). Geometry's yield is bounded by `ln N`, an orthogonal pair already attains it, and beyond a pair the shape buys coverage and continuity rather than surplus. Geometry is powerful and it is capped, and the cap arrives at two participants.

The engineering

6.7 Reading a gradient The framework's foundational layer asks no arithmetic of its reader [Curzi, 2026]. Its output is a perceptual skill: the ability to notice that a circumstance has a direction — that a room, a conversation or an arrangement is either accumulating something or losing it, and that the difference is visible before it is measurable.

Three of its constructs turn out to be theorems in the later corpus, and the correspondence is drawn rather than asserted [Curzi, 2026a]. Its recursion construct is the transient $(1-\delta)^t Q_0$ decaying while the driven term builds. Its fixed-point construct is the object §4 defines as destiny. And its prohibition on intervening at another party's expense is converted by §4.15 into an arithmetic test: compute the deficit and the fixed point for every participant separately, then compute the sum of drives, and if one participant improved while the sum did not move, the gain was exported and nothing was produced.

That conversion is the clearest single illustration of what the framework is for. A prohibition of that form is unfalsifiable in the way a testimonial is — it reports an intention, and the intention is always available. The identity `sum_i b_i` invariant under any transfer of cost among any number of participants is not unfalsifiable at all. It costs one column of a report and it returns an answer on the first pass.

What should also be said about that layer is what it deliberately withholds. It asks no arithmetic of its reader and offers none, so a reader who finishes it has a perceptual skill and no apparatus. That is not a simplification for a general audience; it is the correct ordering. A person who can compute an octave deficit for a circumstance they have not learned to

perceive as having a direction has a number with nothing to point it at, and the failure is silent — the computation runs, the reading prints, and it is a reading of the wrong thing. Each stage is unusable without its predecessor, and the failure at each boundary is different: **without the reading, a number with no referent; without the building, a diagnosis with no intervention; without the join, an intervention with no test of whether it paid.**

6.8 Building a circumstance The framework's engineering layer takes the storage law and builds outward: the room as a circuit, acoustics, thresholds, light and material, the table with no head, order without a centre, institutions as storage [Curzi, 2026h]. Its output is an apparatus for constructing circumstances rather than reading them; its Result Two prices the difference exactly — moving η from 0.35 to 0.55 and δ from 0.25 to 0.12 takes a fixed point from 0.2653 to 3.9281, while a turnover event still costs 70 per cent of stock against 15.

Two further results carry into the close. The first is the headroom gauge $h = \eta \ln N - \delta Q - C$, which says when no arrangement of the available categories improves the system regardless of what anyone does. The second is the concavity rule — feasibility must be tested on the mean of $F(U_t)$ and never on F of the mean, worth 6.2 per cent in the worked case — which is why this framework forbids aggregating ratios across sites.

6.9 The instrument turned inward The same apparatus turned on the person gives the self as storage, the internal surplus, the moment equation, the psychophysics, the gradient and the contraction.

And it demoted its own keystone formula at the gate its author wrote [Curzi, 2026m], whose Result Two fixes the moment equation's fixed point at $D = P$, $r = 1$, $M = E$, with unit elasticity of exactly ± 1 in all three terms and the ceiling $0 \le M \le \ln N$ making the $P \to 0$ divergence finite. The moment equation is demoted on scale: D/P may be used ordinally only — direction and ordering, never a numeric magnitude — pending Rasch instrumentation at 123 items per term. Every psychological frequency derived through D/P inherits the demotion, and this paper says so in place, every time.

That demotion is the reason the join could be attempted at all: the apparatus that could audit a join was built first and used on itself first. Producing the keystone earlier would have produced it without the means to check it.

6.10 Carried to the cell The instrument carried down gives the cell as receiver, addressing, sonification, the Gabor limit, the channel-gain postulate and the grounding protocol [Curzi, 2026,ab]. The output is a specification and a catalogue of failures — the four channel classes with their gains, the places where belief demonstrably does nothing, and a likelihood ratio of exactly 1.000000 for testimonial evidence [Curzi, 2026].

And it stopped. It named the gap — no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal — and does not go past it [Curzi, 2026]. This section ends on it.

6.11 Tuning the relation This series of papers multiplies the two halves and then finds out what the product will and will not support.

The essence law is carried into psychology on a component index n — a count of the distinguishable components required to specify a state — and the count is defended against the obvious objection that deciding what counts as one component is itself a judgment. The honest limit was stated wherever the index was used: **the frequency inherits exactly the defensibility of the count, and a contested decomposition yields a contested frequency.**

The dyad is then built as an interval, $f_{\text{rel}} = k(1 + 11u)^2$, producing the octave deficit — a computable reading of whether a bond pays for itself [Curzi, 2026] — together with the reserve theorem, the elasticities, the order rule and the constructed cases in which the procedure halts.

A tradition is then audited on its own material and what failed is printed [Curzi, 2026,?]. The strongest single item: a canonical twelve-operation schedule of 1471, run as a schedule under this framework, places **the only numerator-acting operation ninth of twelve** — lowering loss before fixing the numerator, which is exactly the order the elasticities show accelerates failure. And the element rule $f = kn^2$ and the compound rule $f = k^2G^4$ cross at $n = 1/\sqrt{k} = 1.438252$, not an integer, so they coincide at no index at all [Curzi, 2026].

The cellular procedure is specified with every honesty clause in place, and then it is stated that nobody knows how a cell could read an ordered message out of a sound [Curzi, 2026].

6.12 What the series actually produced — the six results that constrain it Every one of these is a narrowing. That is why they are the yield.

Result (a) — the surplus law is exactly a Fechner law with threshold `u_0 = 1/(N−1)`.

```
F(u)  =  ln( 1 + (N−1) u )
      =  ln( (N−1) ∗ ( u + 1/(N−1) ) )
      =  ln( (u + u_0) / u_0 )                with  u_0 = 1/(N−1)
```

That is Fechner's law in threshold-corrected form [Fechner, 1860, 1966, Weber, 1834], `S = k ln((I + I_0)/I_0)`, with `k = 1`, the stimulus being non-parallelism and the threshold intensity being `u_0`. Verified for `N = 2` through `N = 100` at seven values of `u` each; maximum residual `4.44e−16`.

The threshold has a meaning that is not an artifact of the algebra. Evaluate the law at it, at every `N`:

```
N_eff(u_0)  =  1 + (N−1) ∗ 1/(N−1)  =  2.000000000000    exactly, for every N
F(u_0)      =  ln 2  =  0.693147180560 nats  =  1 bit  =  2 octaves
```

The threshold is exactly the point at which the effective count reaches two: below `u_0` the system does not yet contain a second independent voice, **and the threshold is one full bit wide.** Its size is a design fact: `u_0 = 1.000000` at `N = 2`, where the threshold *is* the ceiling; `0.500000` at `N = 3`; `0.090909` at `N = 12`; `0.010101` at `N = 100`. Narrow systems are hard to get anything out of not because their members are worse but because the threshold sits at the top of their range.

Two cautions attach. The psychophysical warrant is contested in its own literature [Luce, 2002, Krueger, 1989, Stevens, 1957], and the warrant this framework's arguments actually need is not psychophysical at all: a logarithmic evaluation follows from multiplicative dynamics with no psychological assumption whatever [Peters and Gell-Mann, 2016]. What this establishes is not evidence for either formalism. The framework's law is a theorem under its own axioms and gains nothing from resembling Fechner's [Fechner, 1860, 1966]; Fechner's is an empirical claim and gains nothing from resembling this one. What independent arrival does establish is that the logarithm is not a stylistic choice in this corpus — it is the register both routes landed in.

Result (b) — consonance and redundancy are the same phenomenon. Because `(2n)^2 = 4n^2`, doubling an index raises the frequency by exactly two octaves, so two essences share a pitch class if and only if their indices have the same odd part, and an ordered set `1..N` contains exactly `ceil(N/2)` distinct classes — six at `N = 12`. Price the identification in surplus:

N	ln N	ln(N/2)	loss
4	1.386294	0.693147	0.693147180560
12	2.484907	1.791759	0.693147180560
24	3.178054	2.484907	0.693147180560
100	4.605170	3.912023	0.693147180560

The fold costs exactly `ln 2` — one bit, two octaves of surplus — at every `N` without exception. The arithmetic apparatus is standard once an index exists [Rossing et al., 2002], and assigning numerical weight to intervals is an old project [Euler, 1739, Tenney, 2015]; but what a listener finds consonant depends on the spectrum of the sound and not on the ratio alone [Sethares, 2005], and octave equivalence and interval categories generalise less far than is usually assumed [Shepard, 1982, Demany et al., 2013, Burns, 1999]. **This paper uses the arithmetic and claims none of the perception.** Consonance in the essence lattice is not a separate aesthetic property that happens to be pleasant; it is the collapse of two categories into one, and the collapse has a price that does not depend on how many categories there were.

Result (c) — `u = 4a(1−a)`, so perfect agreement and perfect disagreement are the same measurement. With `a` the fraction of items on which two parties agree, `cos theta = 2a − 1` and

```
u  =  1 − (2a − 1)^2  =  4a(1 − a)
```

a	cos = 2a−1	u	f_rel (Hz)
0.00	−1.0000	0.000000000	0.483426
0.10	−0.8000	0.360000000	11.893047
0.25	−0.5000	0.750000000	41.363114
0.50	+0.0000	1.000000000	69.613306
0.75	+0.5000	0.750000000	41.363114
0.90	+0.8000	0.360000000	11.893047
1.00	+1.0000	0.000000000	0.483426

The table is symmetric about `a = 0.5`, and the two ends are identical. **Two parties who agree about everything and two parties who disagree about everything return the same reading — the ground essence `k = 0.483426 Hz`, no relationship at all.** The contrarian and the sycophant sound the identical note. That is a hard constraint on what `u` can be used to detect, and it is the reason manufactured disagreement is worse than useless: it raises cost and returns exactly nothing.

Result (d) — an orthogonal pair already attains `ln N`, so groups buy breadth and not depth. `F = ln(1 + (N−1) u_bar)` is maximised at `u_bar = 1`, and two unit vectors can be orthogonal. Therefore:

m	links m(m−1)/2	F at u_bar = 1	f (Hz)	cost at kappa=0.055
2	1	2.484907	69.613306	0.055000
3	3	2.484907	69.613306	0.165000
4	6	2.484907	69.613306	0.330000
6	15	2.484907	69.613306	0.825000
10	45	2.484907	69.613306	2.475000
20	190	2.484907	69.613306	10.450000

Surplus is identical at every roster size; cost grows as `m(m−1)/2`. In surplus terms a group buys nothing a well-chosen pair could not, and pays `m(m−1)/2` times as much for it. The coordination cost that grows as $m(m-1)/2$ is the same term software engineering has documented for fifty years [Brooks, 1975, Conway, 1968], and the empirical group-performance literature measures effects the surplus term does not represent [Woolley et al., 2010]. Whatever else a family, a team or a commons is for — coverage, continuity, conversion, redundancy against departure — the formalism prices stored order and nothing else, and on that axis the group's case is not a surplus case. This is the constraint that makes the roster ceiling bite: at `m(m−1) > 2 eta ln N / kappa` the case is structural, and the constructed commons sits at `m(m−1) = 30` against a bound of `30.722482` — **inside it by 0.722482, which is the whole of its margin.** A seventh member takes the roster term to 42 and the case is structural outright.

Result (e) — state frequency and relationship frequency are independent [Curzi, 2026], verified on constructed pairs spanning thousands of cents of the state lattice:

```
indices    state frequencies       span (cents)   u                f_rel (Hz)
------------------------------------------------------------------------------
(1, 2)      0.4834 /  1.9337 Hz       2400.00     1.000000000000    69.6133055503
(3, 6)      4.3508 / 17.4033 Hz       2400.00     1.000000000000    69.6133055503
(5, 10)    12.0856 / 48.3426 Hz       2400.00     1.000000000000    69.6133055503
(7, 11)    23.6879 / 58.4945 Hz       1564.98     1.000000000000    69.6133055503
(2, 9)      1.9337 / 39.1575 Hz       5207.82     1.000000000000    69.6133055503
(4, 12)     7.7348 / 69.6133 Hz       3803.91     1.000000000000    69.6133055503
```

Identical to twelve decimal places, because `u` reads the angle between vectors and is blind to which indices are occupied. **A state's frequency tells you nothing about its relationship frequency, neither may be inferred from the other, and a complete report gives both separately.** Re-indexing moves every state frequency in the system and can leave the relationship reading untouched.

Result (f) — the deficit is blind to decay [Curzi, 2026], restated in §6.4. `D` and `Q_inf` are one number in two units, and `D` contains no `delta`. The consequence is the sharpest limitation the instrument has: an intervention that doubles the rate of loss, or halves it, produces no movement in the diagnostic at all. Report both.

6.13 Why the six are the yield

Every result in that list makes the instrument smaller. (a) says the surplus law is not a new discovery about interaction but a hundred-and-sixty-six-year-old psychophysical form arrived at by a second route. (b) says a property the sonic register treats as aesthetic is the loss of a bit. (c) says the interaction parameter cannot distinguish two situations that no practitioner would ever confuse. (d) says the group has no surplus case. (e) says two of the framework's own readings cannot speak to each other. (f) says the central diagnostic cannot see the thing most people reading it will want to know.

A corpus whose strongest results are all constraints is a corpus that was checked. The alternative — a corpus of affordances, each result adding a capability and none subtracting one — is the signature of work that was never audited, and it is the signature this series has spent its length trying not to have.

6.14 The register of demotions

A corpus that audits itself has to publish the audit, and the four places where the apparatus caught its own author are the strongest available evidence that the gates are real rather than decorative. They are collected here because a closing section is where a reader will look for them.

Demotion 1 — the moment equation demoted on scale, at the framework's own gate [Curzi, 2026m]. The formula the surrounding work was built around was run through the framework's validation protocol and failed it on the scale requirement. `D/P` may be used ordinally only — direction and ordering, never a numeric magnitude — pending Rasch instrumentation at 123 items per term. The demotion travels with every psychological frequency derived through `D/P`, in this framework as in that one.

Demotion 2 — the design intuition inverted [Curzi, 2026q,s]. The corpus had been running on the assumption that trust is a source of order. The trust paper fixes where it actually enters: as retention, $\delta_{\text{eff}} = \delta_0 + (1-a)(1-v)$, giving a superlinear stock dividend from 0.4679 at $a = 0.50$ to 5.0529 at $a = 0.99$ [Curzi, 2026s]. It is not: it acts by reducing `delta`, and `delta` is a denominator, so trust amplifies whatever sign the arrangement already has. In the worked consortium, trust alone moved the fixed point from `-0.0836` to `-0.2114` — worse — while reciprocity alone moved it to `+0.7394` and both together to `+1.8701`. The rule that replaced the intuition is: fix the numerator, then lower the loss.

Demotion 3 — the tradition audited on its own material and found inverted [Curzi, 2026,?]. A canonical twelve-operation schedule of 1471, run as a schedule under this framework, places the only numerator-acting operation ninth of twelve. The canonical sequence lowers loss before it fixes the numerator, which is exactly the order the elasticities show accelerates

failure. That is not a criticism imported from outside; it is the framework reading the tradition's own ordering and printing the result.

Demotion 4 — a source table's own row shown impossible [Curzi, 2026]. The element rule $f = kn^2$ and the compound rule $f = k^2G^4$ cross at $n = 1/\sqrt{k} = 1.438252$, which is not an integer, so the two rules coincide at no index whatever. Since for a single component the compound rule is index-independent, the published O2 row could not have been produced by the rule it was presented under.

And beside those four sit the measured limits of the corpus's own controls, which are not self-catches but are the same discipline applied forward. The shuffled-input control has three measured failure modes: it is anti-conservative as ordinarily run, at 2.2 times nominal alpha; it cannot see repackaging, which only the structural ceiling test can; and it cannot see confounding by a shared slow driver, passing 98.35% of the time in a 4,000-trial simulation. Every paper that recommends the four-baseline suite carries all three, and the conclusion drawn from them is to design controls rather than to add surrogates.

The register is short. That is the point of publishing it: four catches is a small number, and a reader is entitled to weigh it against the size of the corpus and decide for themselves whether the gates were sharp enough. **What is not available is the alternative in which no demotions are printed at all, because a method that never rejects is not a method.**

The specification

6.15 The whole instrument, on one page

```
CONSTANTS
  k      = 0.483425732988414 Hz          the ground essence
  2/ln2  = 2.885390081778               octaves per unit of cost ratio
  N      = 12,  ln N = 2.484907          this framework's category count
  u_0    = 1/(N-1) = 0.090909091         the surplus threshold; N_eff(u_0) = 2 exactly

THE THREE TRANSFER RULES
  Rule 1  a psychological ESSENCE has a frequency        f = k n^2
  Rule 2  a psychological RELATIONSHIP has a frequency   f_rel = k (1 + (N-1) u)^2
  Rule 3  a COMPOUND state uses the alchemical rule      f = k^2 G^4

THE MEASUREMENT
  x, y unit vectors over the N states, components = magnitude-estimated intensity
  cos theta = x . y      u = 1 - (x.y)^2      F = ln(1 + (N-1) u)      f = k e^(2F)

THE KEYSTONE
  f_min = k e^(2c/eta)              sustainable iff f > f_min
  octaves required = 2.885390 (c/eta)
  the ceiling: c/eta <= ln N, i.e. f_min <= k N^2

THE DIAGNOSTIC
  D = 2.885390 (c/eta) - log2(f/k) = 2.885390 (c/eta - F)
  four levers: raise u, raise eta, cut c, raise N
  delta is NOT a lever on D; it is a multiplier on the sign of the drive

THE DESTINY
  Q_inf = (eta F - c)/delta                    where the loop arrives if parameters hold
  |Q_t - Q_inf| = (1-delta)^t |Q_0 - Q_inf|    the approach, an equality
  D = -2.885390 (delta/eta) Q_inf              the two readings, one quantity
  sum_i (eta_i F_i - c_i) invariant            the export test
```

6.16 The standing refusals, complete

```
R-1   ESSENCE, NOT RESONANCE. No person, tissue or molecule resonates at any value here.
R-2   NO CLINICAL, DIAGNOSTIC OR THERAPEUTIC CLAIM, at any scope, anywhere.
R-3   NO CAPABILITY CONFERRED. No DNA edited, altered or damaged.
R-4   NO SUBLIMINALS. Nothing masked, nothing below threshold, nothing concealed.
R-5   CONSENT IS STRUCTURAL. No procedure may reduce anyone's ability to decline.
R-6   THE DEMOTION STANDS. The moment equation is demoted on scale; D/P is ordinal only.
R-7   MAGNITUDE ESTIMATION IS REQUIRED before any relationship number is printed.
R-8   NO SINGLE-PARTY TERM. No reading is a verdict on a person.
R-9   PRECEDENT IS NOT EVIDENCE. The (-1) column carries independent arrival, not proof.
R-10  TESTIMONIALS ARE EVIDENTIALLY INERT. Likelihood ratio exactly 1.000000.
R-11  THE OPEN PROBLEM STANDS. No established mechanism for sequence-specific
      acoustic demodulation.
R-12  EVERY CASE IN THREE BOOKS IS CONSTRUCTED AND LABELLED AS SUCH.
```

6.17 What may be claimed The list is short and it is the honest one.

That the join is exact. Given the essence bridge and the storage fixed point, `f_min = k e^(2c/eta)` follows by substitution. Nothing was added to obtain it and it can be checked in three lines.

That the diagnostic is computable and actionable. Given `u`, `eta`, `c`, `delta` and `N`, the deficit, the destiny, the approach time and the four lever reserves are all determined, and the procedure says which lever binds and which is absent.

That the constraints are real. The six results of §6.12 are algebraic identities or verified computations, not interpretations, and each of them can be reproduced from the formulas in §6.15.

That the apparatus catches its own author. The first demotion took the surrounding work's own keystone at the gate its author wrote. The second corrected the design intuition the corpus had been running on. The third audited a tradition the corpus draws from and found its canonical sequence inverted. The fourth proved a row of a source table impossible.

And what may not be claimed is everything else. In particular: that any of this describes any actual person, pair or group; that a larger fixed point is a better life; **that any frequency here is a physical vibration; that belief moves anything through a channel not named and estimated independently of the outcome; or that the absence of a known mechanism is evidence of anything whatever, in either direction.**

The failure modes

FM 6.1 — THE KEYSTONE QUOTED WITHOUT ITS INPUTS. `f_min = k e^(2c/eta)` reported as a result about relationships in general. It is a result about an arrangement whose `c` and `eta` were measured, and it inherits their defensibility exactly.

FM 6.2 — THE OCTAVE FIGURE TAKEN AS AUDIO. `2.885390` octaves is an exchange rate between nats and a logarithmic scale. No sound is implied, none is prescribed, and R-1 binds.

FM 6.3 — THE CEILING IGNORED. A lever plan produced for an arrangement at `c/eta > ln N`. Worked in §5.3 at `c/eta = 3.000000`: even perfect orthogonality leaves `+1.486245` octaves, and twenty-one categories are required.

FM 6.4 — `D` WITHOUT `delta`. The single most common incomplete reading in the framework. The same `D` is consistent with destinies differing by a factor of thirty.

FM 6.5 — THE TWO FREQUENCY READINGS CONFLATED. Inferring a relationship reading from state frequencies, or the reverse. Six constructed pairs spanning 1565 to 5208 cents return identical `u` and `f_rel` to twelve decimals.

FM 6.6 — THE GROUP JUSTIFIED ON SURPLUS. An orthogonal pair already attains `ln N`. Make the roster's case on conversion, coverage or continuity, which are real and outside this reading.

FM 6.7 — MANUFACTURED DIFFERENCE. `u` depends on `|rho|`; opposition reads `u = 0` exactly as agreement does, while the cost rises.

FM 6.8 — PRECEDENT PROMOTED TO EVIDENCE. The (-1) column establishes independent arrival, which is worth having and is not proof. Two sources agreeing does not double the evidence when only one of them is evidence.

FM 6.9 — TESTIMONIALS. Likelihood ratio exactly `1.000000` at any corpus size. A thousand accounts carry the weight of zero.

FM 6.10 — THE OPEN PROBLEM SOFTENED. Any phrasing that converts "no established mechanism" into "an emerging area", "not yet fully understood", or "preliminary evidence suggests". The sentence is about what is known, and it is not a prediction.

FM 6.11 — THE ABSENCE OF MECHANISM READ AS PERMISSION. The symmetric error. Not proven impossible is not partial support. A claim with no mechanism and no demonstration has zero evidential standing, and the correct number of qualifiers before demonstration is not "few" but "it is not asserted".

Worked examples

6.18 The smallest complete instance **CONSTRUCTED. Not an observation of any person or pair.** Two parties, twelve categories, one measured geometry and three declared parameters. Every line below is derivable from §6.15 and nothing else.

```
INPUTS       u = 0.35     eta = 0.50     c = 0.70     delta = 0.12     N = 12

SURPLUS      F      = ln(1 + 11 * 0.35)  =  1.578978705 nats
             N_eff  = 1 + 11 * 0.35      =  4.850000

PITCH        f_rel  = k * N_eff^2        =  11.371381804 Hz
THRESHOLD    f_min  = k * e^(2*0.70/0.50) =  7.949765419 Hz

CEILING      c/eta = 1.400000  vs  ln N = 2.484907        ->  UNDER, tunable

DEFICIT      D      = 2.885390 * (1.400000 - 1.578978705)  =  -0.516423380 octaves
                                                             SUSTAINABLE

DESTINY      Q_inf  = (0.50 * 1.578978705 - 0.70) / 0.12   =  +0.745744604
             check   D = -2.885390 * (0.12/0.50) * Q_inf   =  -0.516423380
             residual                                       0.00e+00

APPROACH     half-life = ln(0.5)/ln(0.88)  =   5.422271 passes
             t(1%)     = ln(0.01)/ln(0.88) =  36.024800 passes

RESERVES     R_u   = 2.885390 * (ln 12 - F)          =  2.613956 octaves
             R_eta = 2.885390 * c * (1/eta - 1)      =  2.019773 octaves
             R_c   = 2.885390 * (c/eta)              =  4.039546 octaves
             R_N   = 2.885390 * u / N_eff            =  0.208224 octaves per category

EXPORT TEST  single arrangement, one drive: b = 0.089486952
             with any second participant, compare sum_i b_i before and after
```

```
LABEL        CONSTRUCTED
REFUSALS     essence not resonance; no clinical claim; consent structural;
             two readings reported separately; D reported with delta
```

Read it as a whole and notice how little it says. It says this configuration of four numbers accumulates stored order, that it is currently 0.52 octaves above its own threshold, that it would take about thirty-six passes to settle, and that if the deficit ever went positive the cost side has four octaves of room and the geometry side two and a half. That is the entire content, it is reproducible from five formulas, and everything a reader might additionally want to know about the two people in question is outside it.

The gap between what the reading contains and what a reader wants is not a defect to be closed in a later edition. It is the boundary of the instrument, and the whole discipline of three books has consisted of drawing it in the right place and then not stepping over it.

6.19 The second instance — an arrangement the instrument refuses A worked example that only shows a reading coming out well is a worked example that has not demonstrated the instrument. Here is the same procedure run on the framework's constructed commons, where it produces a refusal rather than a plan.

CONSTRUCTED. Not an observation of any group. Six members drawn from one doctrine, so their state vectors are near-parallel by construction — the Parallel Wall, priced. Fifteen pairs at `kappa = 0.055` per link.

```
INPUTS       mean pairwise u = 0.018614     eta = 0.34
             C = 0.055 * 15 = 0.825         delta = 0.09      N = 12

SURPLUS      F      =  0.186272005 nats        N_eff = 1.204754
PITCH        f_rel  =  0.701659675 Hz
THRESHOLD    f_min  = 61.934885399 Hz

CEILING      c/eta = 2.426471  vs  ln N = 2.484907   ->  passes, by 0.058436

DEFICIT      D      = +6.463836987 octaves           SHORT BY 6.46 OCTAVES
DESTINY      Q_inf  = -8.462959612                   DECAYS
             drive b = -0.761666365 per pass
             check   D = -2.885390 * (0.09/0.34) * Q_inf = +6.463836987

APPROACH     half-life 7.349615 passes; 99% of the way down by pass 48.83

RESERVES     R_u   = 6.632448 octaves    exceeds the deficit by 0.168611
             R_eta = 4.620867 octaves    does not close it, even at eta = 1
             R_c   = 7.001314 octaves    exceeds the deficit, at zero cost
             R_N   = 0.044581 octaves per added category

SINGLE-LEVER raise u to 0.95      -> deficit -0.0332   at near-orthogonality
             raise eta to 0.90    -> deficit +2.1075   still failing
             cut kappa to 0.020   -> deficit +2.0085   still failing

ROSTER       m(m-1) = 30  against  2 eta ln N / kappa = 30.722482
             a seventh member takes it to 42 and makes the case structural outright

VERDICT      THE PROCEDURE HALTS.
LABEL        CONSTRUCTED
```

Read the reserve block and the single-lever block together, because they say different things and the difference is the whole lesson. The reserves say two levers close the deficit

in principle. The single-lever table says what those in-principle closures actually require: six people who share almost no state at all, or coordination at zero cost. Neither of those describes a commons, so the arithmetic that says the case is closeable and the arithmetic that says it is not are both correct, **and the second one is the one a planner has to act on.**

Three things this refusal is not. It is not a recommendation that the group disband — the formalism prices stored order and nothing else, and a family, a team or a commons is for a great many things the formalism does not represent. It is not a verdict on any member, because there is no single-party term anywhere in the apparatus. And it is not a claim that nothing can be done; it is a claim that nothing can be *tuned*, which is a narrower and more useful statement. The remedy available to this arrangement is restructuring — a different roster, a different cost structure, or a description with genuinely more categories in it — and every one of those is a change of what the arrangement is rather than an adjustment of how it is run.

Work that only showed fixable examples would be lying by selection, and the commons is carried through five chapters of this volume for exactly that reason. The instrument's most valuable output is the one that costs nothing to produce and saves a year: *this cannot be tuned, and here is the division that shows it.*

The open problem

6.20 Where the work stops The companion paper specifies a complete procedure [Curzi, 2026]: the ethical architecture and the no-subliminals constraint, then the encoder, the addressing convention, the sonifier, the transposition, the rendering within the Gabor bound and the $N \leq 25$ single-transposition capacity, with preregistration and falsification criteria attached. All of it is deterministic and all of it is verified.

And it then says that nobody knows how a cell could read an ordered message out of a sound.

> **There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal.**

The gap decomposes into three requirements. Symbol separation is solved — the alphabet separates, and that is demonstrated. Order memory across the symbol interval has no identified mechanism. Mapping from a decoded symbol string to sequence-specific binding at a genomic locus has no identified mechanism. Eleven upstream components are established and one link is missing, and the missing link is specifically the carriage of sequence over force.

The operational form is the one that matters, because it is what an experiment would have to do. Take a stimulus and its symbol-permuted twin. The two have identical total energy, identical duration, identical long-term power spectrum and identical symbol counts. If a cell cannot be shown to distinguish them, no sequence-specific channel has been demonstrated, and nothing else in the experiment matters.

Two symmetric errors sit on either side of that sentence and both are refused. **It is not a claim of impossibility** — absence of a known mechanism is a statement about the literature, and the burden of demonstration sits entirely on the claimant, which is the ordinary position for any unestablished claim. **Nor is it a licence** — "not proven impossible" is not partial support, and a claim with no mechanism and no demonstration has zero evidential standing.

6.21 Why the argument stops here That paper could have ended at its procedure, with a reader's impression that something had been achieved. Something has been achieved — an encoder that round-trips, an alphabet that separates, a rendering a listener can follow and refuse — and none of it is the thing the work is named for.

The same choice is available here and is declined for the same reason. A closing section is where a corpus is most tempted to convert its constraints into a narrative of arrival, because the reader is at the end and the material is fresh and a summary is easy to write in the key of achievement. **This one ends on the gap instead, because the gap is the load-bearing item and putting it last is the only placement in which it cannot be forgotten.**

There is a second open question and it belongs beside the first, because they fail in opposite directions. A second open question asks whether the storage law is closed under any natural aggregation operator — whether a population of systems each obeying the law, combined by that operator, yields a system obeying the law with derived parameters. If such an operator exists, the scale correspondence stops being a modelling convention and becomes a theorem. If none exists, the correspondence is a modelling convention and nothing more, and the framework's scopes are separate applications of one equation that happen to look like a ladder. Which of those is true is not known here, and it requires no laboratory — only somebody to define the operator and try to prove closure, and to publish the result **including, especially, if the result is that the law is not closed.**

So: two problems. One needs an experiment and one needs an afternoon with a proof. Neither has been done, and the series stops at both of them rather than past them.

6.22 What is on the bench The framework produces one substitution, one diagnostic, one definition of destiny, one conservation identity, six constraints that narrow the instrument, a catalogue of failure modes long enough to be embarrassing, and a list of refusals longer than the list of claims.

The claims are these, and there are no others. The join is exact. The diagnostic is computable from measurements a practitioner can take. The deficit and the fixed point are one number in two units. The deficit cannot see decay, so both must be reported. Cost export is detectable on the first pass by an arithmetic test on a conserved sum. A destiny has four arguments and changing one is the only way to change it. Belief acts through channels with estimable gains and does nothing where no channel exists. Testimonial evidence has a likelihood ratio of exactly `1.000000`.

Everything else is either a badged postulate, a chosen constant, a constructed example, a rejection, or a statement of what is not known.

And the open-problem sentence is exactly as true at the end of this paper as it was at the beginning of the one that states it [Curzi, 2026]**.** Nothing here changes it, nothing here claims to, and the parts of the corpus that are correct are correct independently of how it comes out.

The instrument is on the bench. It reads one gauge, it refuses more than it asserts, it has caught its own author four times, and it stops at a gap that will be settled by an experiment rather than by an argument.

Anyone may take it further, on the condition this framework has tried to keep: say which claims are proved, badge the ones that are not, name the ones that are open, and publish the demotions when the gates catch the author. They will.

The reference apparatus for this material — the constants card, the interval tables, the measurement protocols, the concordance register and the evidential-status appendix — is collected separately [Curzi, 2026].

7 Results

Every number below is exact arithmetic on the stated constructed inputs. **Nothing in this section was measured**, and every case carries its CONSTRUCTED label at the point of use in §§4–6.

R1 — conservation of drive under cost transfer. The sum of drives $\sum_i b_i = \sum_i (\eta_i F_i - c_i)$ is exactly invariant under any transfer of cost among any number of participants (Theorem 4.5). Verified across an eight-point export range in a two-party arrangement at $+0.643688229683$ in every row, and across four transfers among three participants at $+1.268940530229$ in every row, residual $0.000\text{e}{+00}$. The equivalent statement in stock units is that the δ-weighted sum $\sum_i \delta_i Q_{\infty,i}$ is invariant. **Corollary 4.6:** the *plain* sum of fixed points inherits the invariance only when all participants share one δ. Where they do not, $\mathrm{d}(\sum_i Q_{\infty,i})/\mathrm{d}x = 1/\delta_P - 1/\delta_R$, measured at 13.333333 per unit exported for $(\delta_P, \delta_R) = (0.05, 0.15)$ — the plain sum rises by more than two and a half units across a 0.20 export range while nothing is produced.

R2 — the approach is an equality, not a bound. $|Q_t - Q_\infty| = (1-\delta)^t |Q_0 - Q_\infty|$ exactly (Theorem 4.2), checked two independent ways over $t = 0$ to 100 on the repaired constructed dyad ($\delta = 0.09$, $b = +0.547598406839$, $Q_\infty = +6.0844267427$): **maximum residual** $2.0624\text{e}{-15}$**.** The drive cancels in the subtraction, so the rate of approach is a function of the loss term and of nothing else: $\tau = 10.603253$ passes, half-life 7.349615, $t(1\%) = 48.8298$, none of which contains Q_0, Q_∞, η, F or c.

R3 — completeness of the lever set. Q_∞ is a function of exactly (η, F, c, δ), and F is reached only through u and N (Theorem 5.1). Four arguments, five handles, and no fifth of either. A proposal that cannot be assigned to one of them changes the destination by zero.

R4 — the exporter's lag. At an export of $x = 0.15$ in the constructed two-party case, the disadvantaged party's destination is already -0.599002 while its stock stays positive until $t^* = 10.850479$ passes. Its deficit D_R turns positive on the first pass. **The stock-based question returns a false negative for eleven periods; the drive-based test returns the right answer immediately.** The exact sustainability crossing is at $c_R = \eta_R F = 0.646089822845$, i.e. $x = 0.096089822845$. If the disadvantaged party then exits — its structural right [Curzi, 2026, Hirschman, 1970] — the exporter's own fixed point falls by 10.528871 units and its deficit swings 6.214071 octaves the wrong way, which is $2.885390\,F$ exactly.

R5 — the surplus law is exactly a threshold-corrected Fechner form. $F(u) = \ln((u+u_0)/u_0)$ with $u_0 = 1/(N-1)$, verified for $N = 2$ to 100 at seven values of u each, maximum residual $4.44\text{e}{-16}$. At the threshold the affine effective count is 2.000000 exactly for **every** N, and $F(u_0) = \ln 2 = 0.693147180560$ nats — one bit, two octaves. The threshold width is a design fact: $u_0 = 1.000000$ at $N = 2$, 0.090909 at $N = 12$, 0.010101 at $N = 100$.

R6 — the honest limit of the geometry lever. At $\eta = 0.20$, $c = 0.60$, $N = 12$, pushing u to its theoretical maximum of 1 still leaves $D = +1.486245$ octaves. Closure requires $N \geq e^3 = 20.0855$, so $N = 21$ gives $D = -0.128465$ — and only at perfect orthogonality. **An arrangement over the ceiling** $c/\eta \leq \ln N$ **is structural, not tunable**, and one division says so before any effort is spent.

R7 — the keystone is a substitution, not a postulate. Given the essence bridge $f = k\,e^{2F}$ and the storage fixed point, $Q_\infty > 0$ iff $F > c/\eta$ iff $f > k\,e^{2c/\eta}$. Three lines, nothing added. The exchange rate is 2.885390 octaves of required pitch per unit of cost-to-conversion ratio, and the ceiling in the frequency register is $f_{\min} \leq kN^2 = 69.6133055503$ Hz at $N = 12$. **The substitution is exact; the hertz labels on both sides are postulated (P-a, P-c).**

R8 — the demotion register. Four occasions on which the framework's gates rejected its author's claims: (1) the moment equation demoted to ordinal use on the scale requirement, pending Rasch instrumentation at 123 items per term [Curzi, 2026m]; (2) the design intuition inverted — trust enters as retention, and in the worked consortium trust alone moved a fixed point from -0.0836 to -0.2114 while reciprocity alone moved it to $+0.7394$ and both together to $+1.8701$ [Curzi, 2026q,s]; (3) a canonical twelve-operation schedule of 1471, read as a schedule, places its only numerator-acting operation ninth of twelve [Curzi, 2026,?]; (4) a source table's own row shown impossible, the element and compound rules crossing at $n = 1/\sqrt{k} = 1.438252$, which is not an integer [Curzi, 2026].

CR — the counter-result. An intervention that appears to improve one participant's destination is exactly offset in the sum of drives, **so what looks like created order is transferred cost — and the deficit instrument cannot see the transfer at all**, since D contains no δ and the transfer's whole visible consequence in stock units runs through δ. In the worked case the exporting party's deficit improves from -3.590989 to -4.902530 and its fixed point rises by more

than two units, every number it can compute about itself has genuinely improved, and it is not lying to anybody including itself.

The negative results, reported in the same language. $\partial D/\partial\delta = 0$ identically: across δ from 0.30 to 0.01 the carried constructed dyad's deficit holds at $+4.3384201973$ while its destination moves from -1.5035819 to -45.1074559, a factor of thirty. The lever with 4.80 times the elasticity of its rival has 4.174198 octaves of reserve against a 4.338420-octave requirement and **cannot close the gap at all**, even at a conversion efficiency of one; the efficiency the closure would require is $c/F = 1.101077$, which is not an expensive lever but an absent one. And the constructed commons halts: no single lever closes $+6.463837$ octaves in practice, and the roster sits at $m(m-1) = 30$ against a bound of $2\eta \ln N/\kappa = 30.722482$, so a seventh member takes it to 42 and outside.

8 Discussion

8.1 What follows

A prohibition became a test. The framework's foundations prohibit intervening at another party's expense [Curzi, 2026]. A prohibition of that form is unfalsifiable in exactly the way a testimonial is: it reports an intention, and the intention is always available. Theorem 4.5 replaces it with a quantity that either moved or did not, costs one column of a report, and does not ask what anybody meant. **A well-intentioned export and a deliberate one produce the same reading, which is correct, because the harm is the same.**

Two numbers must travel together, and the reason is now arithmetic. D reads the sign and size of the sustainability gap; δ converts it to a level; neither determines the other. The corpus has a structurally parallel result in a different domain: dwell time decides whether a currency drag is a transaction cost or a capital loss, a 40 per cent annual drag costing 0.767 per cent per pass at a seven-day dwell and 19.726 per cent at 180 days [Curzi, 2026]. **In both cases a rate is converted into a level by a parameter the headline instrument does not contain**, and in both cases the remedy is to publish the second parameter rather than to improve the first.

Order is a first-class variable. Both paths in §5.6 end at the identical fixed point, so a planner who evaluates only the destination sees no difference between them; the trajectories differ by twelve periods of permanent lag and 111.133601 stock-periods. The corpus's cadence theorem is the same lesson from the other side — for concave g with $g(0) = 0$, $D(n) = n\,g(T/n)$ is nondecreasing in n, so dividing a fixed budget over more occasions never delivers less [Curzi, 2026v]. **When something happens is not a detail of how much of it happens**, and neither result is visible in a fixed point.

8.2 What does not follow

No destiny reading is a statement about a person. Every term in Q_∞ belongs to an arrangement: u is a joint geometric property of two or more state vectors and cannot be assigned to one of them, and η, c and δ describe what an arrangement does, costs and forgets. **There is no single-party term anywhere in the apparatus**, and Gate 4-9 exists to enforce that in print.

A larger fixed point is not a better life. Q_∞ is stored order in the model's units. The formalism prices stored order and nothing else, and a family, a team or a commons is for a great many things it does not represent.

The export test is not a general detector of exploitation. It is exact for transfers of cost and silent outside them. Degrading another party's geometry, conversion or retention is not a transfer, the sum of drives does move, and a reader who takes the test for a general detector will be misled by their own instrument (§4.15).

The conservation identity establishes nothing about the world. It is an arithmetic consequence of the definition of a transfer (Def. C) and would hold of any additively separable

ledger. Its value is that it makes a specific failure mode visible on pass one, not that anything has been shown to be conserved [Mirowski, 1989, Cartwright, 1999].

8.3 The functional form is underdetermined

$F(u) = \ln(1+(N-1)u)$ is a theorem under the corpus's own axioms and an assumption against every alternative. Latané's social impact law proposes a **power** of the number of sources with exponent below one [Latané, 1981]; $n \log n$ has been proposed for the same qualitative facts [Briscoe et al., 2006]; and many mechanisms produce curves that cannot be told apart from the data usually available [Stumpf and Porter, 2012]. **All three are decelerating and all three fit the qualitative facts, and nothing in this paper discriminates among them.** The corpus's axioms establish that its form follows from its axioms; they do not establish that its axioms are the ones any data support.

One warrant is stronger than the one usually offered and should be preferred. A logarithmic evaluation follows from multiplicative dynamics with **no psychological assumption at all** [Peters and Gell-Mann, 2016], which is a better ground for the corpus's logarithms than the psychophysical resemblance of R5 — whose own literature is contested [Luce, 2002, Krueger, 1989, Stevens, 1957]. The independent arrival at a logarithm is worth recording; **it is not evidence for either formalism**, since the corpus's law is a theorem under its axioms and Fechner's is an empirical claim, and neither gains support from resembling the other.

8.4 What would distinguish this framework from a neighbouring model

The standing demand on any programme of this kind is to name the observation only its model makes [Durlauf, 2005, Sobkowicz, 2020]. This paper's honest answer has two parts.

One candidate exists and is R4. The detection-lag prediction is sharp, signed and numerical: a party whose cost has risen should show a *destination* that turns negative immediately and a *stock* that turns negative only after t^* passes, with t^* predicted in advance from δ alone. A model in which the observable stock is the state, with no separate drive, predicts no such gap. **Testing it requires the drive and the stock to be instrumented separately, which has not been done here or anywhere in this corpus.**

And the general answer is discouraging, and is stated. With fifteen years of high-quality longitudinal data and 160 teams, life outcomes were barely predictable at all [Salganik et al., 2020, Hofman et al., 2017]. Prediction is one of many reasons to build a model [Epstein, 2008], and what this framework claims is the others: an explicit accounting frame, a set of refusals with arithmetic behind them, and a diagnostic that halts.

8.5 Bounds carried from elsewhere in the corpus

The most restrictive bound on the whole exercise is in another series and is accepted here. The corpus's assembled-architecture paper reports that the only reading under which its structure clears the build threshold is the naive one assuming zero correlation between coordinating jurisdictions ($h = +0.083099$), and that **every defensible bloc assignment returns a negative surplus** [Curzi, 2026]. The structural point transfers directly: a positive drive computed on parameters that were assumed rather than measured is a positive drive **conditional on the assumption**, and this paper's constructed cases are exactly that. A summary statistic does not determine the joint structure it was computed from [Kirman, 1992, Harrison and Klein, 2007], and the deficit's blindness to δ is the local instance of a general fact.

9 Verification

Every result is decomposed into steps a reader can check individually. **A postulate that can be dropped without a result changing has been dropped**, and this section says outright where

the chain stops being a proof.

- **ARITHMETIC** — reproducible from the inputs named.
- **DEFINITION** — true by stipulation; carries no empirical weight.
- **CITED** — supported by a named external source.
- **POSTULATED** — asserted by this work, with no external support.

9.1 The postulates

	postulate	status
P1	$\delta \in (0, 1]$.	DEFINITION (the model's domain)
P2	The update is affine: $Q_{t+1} = (1-\delta)Q_t + b$.	MODELLING CHOICE [Curzi, 2026a]
P3	A contraction on a complete space has one fixed point, reached from any start.	CITED [Banach, 1922]
P4	$F(u) = \ln(1 + (N-1)u)$, bounded above by $\ln N$.	DEFINITION (design axiom; **not** a Hill number)
P5	A transfer of cost is $c_p \to c_p - x$, $c_r \to c_r + x$, all else fixed.	DEFINITION (Def. C)
P6	φ is the positive root of $x^2 = x + 1$; the whole tone is $9{:}8$.	CITED [Euclid, 1956, Barbour, 1951]
P7	The dimensionless $(9/8)^2/\varphi^2$ **is a frequency in hertz**.	POSTULATED — no source
P8	The essence bridge $f = k\,e^{2F}$.	POSTULATED — no source
P9	Intensities are magnitude-estimated, so u is ratio-scaled.	POSTULATED; contested [Michell, 1997]

Dropped as unnecessary. No postulate about *which* categories, or about their order, is needed for R1, R2, R3, R4 or R6. **P7 and P8 are needed only where a number is quoted in hertz**, and every dimensionless result in this paper survives their rejection. P9 is needed only where a reading is claimed of an actual arrangement; the constructed cases stipulate their inputs and do not need it.

9.2 The compounding chain

step	consumes	produces	status
L1	P1, P2	$\lvert T(x) - T(y)\rvert = (1-\delta)\lvert x - y\rvert$, Lipschitz $1 - \delta < 1$	ARITHMETIC
L2	L1, P3	unique $Q_\infty = b/\delta$, reached from any start	CITED
L3	P2, L2	$\lvert Q_t - Q_\infty\rvert = (1-\delta)^t \lvert Q_0 - Q_\infty\rvert$ — **R2**; max resid $2.0624\text{e}{-15}$ over 100 passes	ARITHMETIC

step	consumes	produces	status
L4	L3	τ, half-life and $t(\epsilon)$ contain δ alone	ARITHMETIC
L5	P4, P8	$2/\ln 2 = 2.885390081778$ octaves per nat	ARITHMETIC given P8
L6	L2, L5	$D = 2.885390(c/\eta - F) = -2.885390(\delta/\eta)Q_\infty$; max resid $8.8818\text{e}{-16}$ over 12 sets	ARITHMETIC
L7	L6	$\partial D/\partial\delta = 0$ identically — the decay blindness	ARITHMETIC
L8	L2, P5	$\sum_i b_i$ invariant under transfer — **R1**, resid $0.000\text{e}{+00}$	ARITHMETIC
L9	L8, equal δ	$\sum_i Q_{\infty,i}$ invariant — Corollary 4.6	ARITHMETIC
L10	L8, unequal δ	$\mathrm{d}(\sum_i Q_{\infty,i})/\mathrm{d}x = 1/\delta_P - 1/\delta_R = 13.333333$	ARITHMETIC
L11	L3 at $x = 0.15$	$t^* = 10.850479$ passes — **R4**	ARITHMETIC
L12	inspection of L2	exactly four arguments, five handles — **R3**	ENUMERATION (4 symbols)
L13	P4	$F = \ln((u+u_0)/u_0)$, $u_0 = 1/(N-1)$, $N_{\text{eff}}(u_0) = 2$, $F(u_0) = \ln 2$ — **R5**	ARITHMETIC
L14	L13	the form coincides with a threshold-corrected Fechner law [Fechner, 1860, Weber, 1834]	CITED
L15	P6	$(9/8)^2 = 1.265625$; $\varphi^2 = 2.618033988749895$; ratio $k = 0.483425732988414$	ARITHMETIC (exact, defined)
L16	L15, **P7**	k carries the unit Hz	**POSTULATED**
L17	L16, **P8**, L2	$f_{\min} = k\,e^{2c/\eta}$ — **R7**	ARITHMETIC given P7, P8
L18	L17, P4	$c/\eta \le \ln N$; at $c/\eta = 3$, $u = 1$, $N = 12$ gives $D = +1.486245$; $N \ge e^3$; $N = 21$ gives -0.128465 — **R6**	ARITHMETIC

9.3 The three checks a reader can do in a minute

$(9/8)^2 = 81/64 = 1.265625$. $\varphi^2 = \varphi + 1 = 2.618033988749895$. Their quotient is 0.483425732988414, exactly, to as many places as anyone cares to carry. **Those three establish k as a defined constant and nothing more**; L16 is where the hertz enters, and it is a postulate.

Two more, on the paper's own result. Take $\eta_P = 0.44$, $\eta_R = 0.30$, $F = 2.153632742815$, $c_P = 0.400$, $c_R = 0.550$. Then $b_P = 0.547598406839$ and $b_R = 0.096089822845$, summing to 0.643688229683. Now move $x = 0.15$ of cost from P to R: $b_P = 0.697598406839$, $b_R = -0.053910177155$, and the sum is 0.643688229683 again. **That is the whole of Theorem 4.5, and it takes one minute.**

9.4 What this chain is and is not

Read the status column. L16 and L17 are postulated; P2, P4 and P5 are definitions or modelling choices. **A chain whose load-bearing steps are definitions and postulates is not a proof of its interpretation**, and this one is not offered as one.

Everything dimensionless survives the rejection of P7 and P8. R1, R2, R3, R4, R5 and R6 are statements about a linear recursion, an additively separable ledger and the algebra of a logarithm; they do not become false if nobody accepts that $(9/8)^2/\varphi^2$ is a frequency. R7 does not

survive: it is exactly the substitution that carries the postulated bridge into the sustainability condition, and its content is that **given P7 and P8, nothing further was assumed.**

And R1's status deserves its own sentence, because the word "conservation" oversells it. Theorem 4.5 follows from P5 alone — from the *definition* of a transfer as an operation adding $+x$ to one drive and $-x$ to another. It is a bookkeeping identity, it would hold of any additively separable ledger, and **it is not evidence that anything in the world is conserved.** What makes it useful is entirely procedural: the per-participant readings are computed independently from separately elicited parameters, and the sum is then checked, so the check can fail. A reviewer looking for the weakest link should look at L16, L17 and at the empirical status of P9, and this section is written so that the search takes under a minute.

9.5 A discrepancy found in the source, printed

Two numbers in the source manuscript could not be reproduced, and both fail in the same way: a value beginning $30.$ was printed beginning $6.$

quantity	source printed	recomputed	where
$f_{\min}$ at $c/\eta = 2.0666667$	6.1587698405 Hz	30.1587698405 Hz	§4.21, §6.1
roster bound $2\eta \ln N/\kappa$	6.722482	30.722482	§5.20, §6.12(d), §6.19

Both are checkable in one line. $k\,e^{2\times 2.0666667} = 0.483425732988414 \times e^{4.1333333} = 30.158769840525$, and the value must lie between the table's own neighbouring rows of 26.394151 at $c/\eta = 2.000$ and 69.613354 at 2.484907, which 6.158790 does not. For the roster, $2 \times 0.34 \times \ln 12/0.055 = 30.722482$, and the printed 6.722482 contradicts the same paragraph's own statement that the six-member commons **passes** the ceiling test by 0.058436 — at a bound of 6.722482 a roster of $m(m-1) = 30$ would already be structural, and the source's next clause, that a seventh member "takes it to 42 and makes the case structural outright", would be redundant.

The corrected values are used in §§4–6 and the correction is recorded here rather than made silently. No conclusion changes: D, Q_∞ and every octave figure in the affected passages were computed from c/η directly and are unaffected, and the commons still halts on the single-lever table rather than on the roster bound.

10 Limitations

Every check the governing rigour standard specifies is listed by name, passed or not.

L-1. The dimensional audit does not close for any quantity in hertz. k is computed as a ratio of two dimensionless numbers and no operational procedure is supplied by which it acquires $[T^{-1}]$ (P-a, L16). Everything expressed as a ratio, a nat, an octave or a count is unaffected.

L-2. The storage law as printed is not dimensionally homogeneous, and the source is silent on the pass-length factor that would make it so (§3.3). Nothing depends on it, and it is disclosed rather than repaired.

L-3. Stored order has no operational unit. Q inherits nats from ηF, which makes the audit consistent and leaves Q a dimensionless placeholder. **Nobody has measured it.**

L-4. $1+(N-1)u$ **is not a Hill number** and no ecological, economic or information-theoretic literature contains it (§3.4). Where the corpus's other series use $1/\mathrm{HHI}$ and $\exp(H)$, those are Hill numbers of orders 2 and 1 and are fully grounded [Hill, 1973, Laakso and Taagepera, 1979]; **this paper's effective count is not, and the substitution is named at every use.**

L-5. The functional form of $F(u)$ **is underdetermined** by the available evidence [Latané, 1981, Briscoe et al., 2006, Stumpf and Porter, 2012], and nothing here discriminates among the candidates (§8.3).

L-6. Conjoint measurement was not tested, so the ratio-scale treatment of u, η, c and δ is provisional and Michell's objection stands unanswered [Michell, 1997, 1999]; the available remedies [Luce and Tukey, 1964, Krantz et al., 1971, Rasch, 1960] were not run. **The corpus's own moment equation is already demoted to ordinal use for this reason** [Curzi, 2026m], and the framework's habit of scalar reduction is conceded where the ladder is built [Curzi, 2026] rather than re-argued here.

L-7. The reflexive specification is not identified [Manski, 1993, Shalizi and Thomas, 2011]. The remedy adopted is stipulation of the reference set [Brock and Durlauf, 2001], and the claim is **a modelling position rather than an identified empirical result.** The framework also exempts its own parameters from its own reflexivity [Lucas, 1976], which it cannot consistently do.

L-8. A published destiny reading can make itself true. Expectancy effects are real [Merton, 1948, Rosenthal and Jacobson, 1968] and a theory has been observed becoming true by being used [MacKenzie and Millo, 2003]; on the best long-run evidence the effect is **small and often dissipating** [Jussim and Harber, 2005]. A reading used as a target also stops measuring what it measured [Goodhart, 1975, Campbell, 1979]. The separation clauses of §4.8 are mandatory for that reason.

L-9. Every case is constructed. The carried dyad in both phases, the two- and three-party export arrangements, the commons, the family and the over-ceiling case are chosen parameter sets. Their exactness follows from the parameters being exact and **does not transfer to a measured case.**

L-10. No hypothesis in §2.1 was tested. H1's detection lag, H2's ceiling, H3's retention mechanism and H4's selection model are all untested, and **no sample exists for any of them.** The arithmetic of each is a theorem; the empirical content of each is not.

L-11. No statistical claim is made and none is supported. There are no p-values, effect sizes, intervals or samples in this paper. No model comparison, out-of-sample test, matched null or information criterion is reported, because nothing was fitted. Where the corpus does report a preregistered fit, it **failed its own second falsification rule** — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin [Curzi, 2026i] — and that failure is the relevant precedent for anyone tempted to read these identities as empirical support.

L-12. No scaling claim is made and the standard procedure was not run. The essence law's exponent is **defined by the model, not estimated**; Clauset et al. [2009] was not applied because there is no data to apply it to.

L-13. There is no conservation law. R1 is an arithmetic consequence of the definition of a transfer (§9.4). The apparatus writes Q with the grammar of a conserved quantity without one [Mirowski, 1989]; validity is claimed only within a specified arrangement [Cartwright, 1999] and no universality is claimed.

L-14. The export test's scope is narrow and its silence is total outside it. It detects transfers of cost. It does not detect degradation of another party's geometry, conversion or retention, and it returns "no export" in every such case (FM 4.9).

L-15. Boundary behaviour, examined and reported. $\delta \to 0$ is a **pathology, not an ideal**: the map becomes an isometry, Banach does not apply, there is no fixed point and $Q_t = Q_0 + bt$ grows without bound (FM 4.11). $\delta \to 1$ gives $Q_\infty = b$ and a one-pass approach. $u \to 0$ gives $F = 0$ and the ground essence; $u \to 1$ gives the ceiling $\ln N$. $\eta \to 0$ makes the drive $-c$ and the destination $-c/\delta$. **What is not examined is the behaviour of the identity when parameters change during the approach**, which is every real case, and the model has nothing to say about it.

L-16. Invariance is claimed for one transformation only. The conserved quantity is invariant under transfers of cost among participants at fixed (η_i, F_i, δ_i). It is **not** invariant under a change of any of those, and it is not scale-free: no rescaling group is exhibited and none is claimed.

L-17. The Fechner coincidence is a resemblance and not evidence. R5 is an algebraic

identity; the claim that it *is* Fechner's law adds nothing to either side, and the psychophysical warrant is contested in its own literature [Luce, 2002, Krueger, 1989]. The stronger ground for the logarithm is multiplicative dynamics with no psychological assumption [Peters and Gell-Mann, 2016], and that is the warrant the argument actually needs.

L-18. Two source numbers were wrong and are corrected in §9.5. A third bibliographic gap is disclosed rather than filled: **the survivorship precedent in §5.16 — Abraham Wald's wartime work on aircraft armour — is not in this corpus's bibliographic database and is named there without a citation key.** No approximate substitute was attached.

L-19. Nothing here is clinical. No case describes any person, pair or group; no reading is an assessment; and the deficit's sign says a configuration loses stored order at a computable rate and nothing else.

11 Conclusion

What is established, and checkable. The storage loop has exactly one fixed point for $\delta > 0$ and reaches it from every start, and the approach is a geometric **equality** rather than a bound — verified to a maximum residual of $2.0624\text{e}{-15}$ over a hundred passes. The octave deficit and the fixed point are **one quantity in two units**, $D = -2.885390\,(\delta/\eta)\,Q_\infty$, verified to $8.8818\text{e}{-16}$ across twelve parameter sets; and D contains no δ, so **the same deficit is consistent with destinations differing by a factor of thirty** and no reading that reports one without the other is complete. **The sum of drives is exactly invariant under any transfer of cost among any number of participants**, holding at $+0.643688229683$ across an eight-point export range and at $+1.268940530229$ across four transfers among three parties, with residual $0.000\text{e}{+00}$; the plain sum of fixed points inherits that invariance only at equal δ, and otherwise moves at $1/\delta_P - 1/\delta_R$ per unit exported. **A destiny has exactly four arguments and five handles, and there is no fifth of either.** A lever can be **absent rather than expensive**: the lever with 4.80 times the elasticity of its rival carries 4.174198 octaves of reserve against a 4.338420-octave requirement and cannot close it at any effort, because the efficiency required is $c/F = 1.101077$. Reversing the ordering rule costs **twelve periods of permanent lag, thirteen periods below the starting stock and 111.133601 stock-periods**, at an identical destination. The surplus law is exactly a threshold-corrected Fechner form with $u_0 = 1/(N-1)$, at which the affine effective count is **2.000000 exactly for every** N and $F(u_0) = \ln 2$ — one bit, two octaves. And an arrangement over $c/\eta \le \ln N$ is structural: at $c/\eta = 3$ even perfect orthogonality leaves $+1.486245$ octaves and twenty-one categories are required.

The counter-result, in the same plain language. An intervention that improves one participant's destination is exactly offset in the sum of drives, so **what looks like created order is transferred cost** — and the framework's own central instrument cannot see the transfer, because D contains no δ. In the worked case the exporting party's numbers all improve, honestly, while the other party's stock stays positive for $t^* = 10.850479$ passes after its destination has gone negative.

What is proposed rather than shown. That $(9/8)^2/\varphi^2$ is a frequency; that the essence bridge $f = k\,e^{2F}$ holds; that the affine effective count is the right construction; and that trust acts through retention. **The first two are postulates with no external support, the third is a design axiom that is not a Hill number, and the fourth is untested** — all four are badged at every appearance, and every dimensionless result above survives their rejection.

What is not claimed at all. No arrangement's parameters are claimed to hold; no larger fixed point is a better life; no frequency here is a physical resonance; no reading describes any actual person, pair or group; and nothing in this paper is clinical, diagnostic, therapeutic or prognostic. **The apparatus caught its own author four times, and the register is printed in §6.14 rather than summarised, because a method that never rejects is not a method.**

12 Declarations

Data and code availability. No dataset was collected for this study and none exists. Every number in this paper is exact arithmetic on constructed parameter sets that are printed in full at the point of use and are labelled CONSTRUCTED there; any reader can recompute all of them from the formulas collected in §6.15. The figures are generated by the Python scripts named in their captions, which take their inputs from the tables in the text and are available from the author on request. **No measurement of any human being, pair, group or institution underlies any figure, table or number in this paper.** A paper about destinies, trust and cost export reads as though somebody had been observed or asked something; nobody was.

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the analysis, the figures and the manuscript.

Ethics. No human-subjects research was conducted and no personal data were processed. The carried dyad in both phases, the two- and three-party export arrangements, the commons, the family and the over-ceiling case are invented parameter sets and do not describe any person, pair or group.

Preregistration. No test in this paper was preregistered, because no test was run: §2.1's four hypotheses are all recorded as untested and no sample exists for any of them. The corpus's one preregistered fit is reported elsewhere and **failed its own second falsification rule** [Curzi, 2026i]; that outcome is cited here in the same terms a success would have received.

Clinical and relational disclaimer. A destiny reading is a computation on four declared parameters of an arrangement. It is **not** an assessment of anybody's health, capacity, character or prospects; the twelve state classes are not diagnostic categories; and nothing in this paper is an assessment, a screening instrument, a treatment, counselling, or advice about whether anyone should remain in or leave any relationship. $Q_\infty < 0$ says that a configuration loses stored order at a computable rate under the model's assumptions. It does not weigh what a relation is for, what else it produces, or what ending it would cost. A reader who wants to know whether something is wrong should ask someone qualified to answer that; this instrument is not qualified and does not attempt it.

Chapter 50

Reference Apparatus for the Frequency Framework: Constants, Lattice Tables, Measurement Protocols, and a Four-Drawer Classification of Evidential Status

Abstract. This paper is the reference and status apparatus for a framework whose results are established elsewhere. **It advances no testable hypothesis and is not a research contribution.** It assembles the constants — $k = (9/8)^2/\varphi^2 = 0.483425732988414$, a *defined* constant, exact, whose hertz label is a separate and postulated claim; $k^2 = 0.233700439315386$; $\ln 12 = 2.484907$; $kN^2 = 69.6133$; $2/\ln 2 = 2.8853900817779268$; $1/(4\pi) = 0.0795774715$ — together with the three transfer rules, the keystone and its deficit identity, the ceiling and hold line, the four levers and their partials, the agreement identity with its $4/m$ resolution bound, eight measurement protocols, a concordance register, and the exact unit conversions in which one bit is two octaves exactly. Four items originate here. The folded-lattice minimum gap is derived by exact rational arithmetic as a function of ladder length: 1200 cents at $N = 2$, the Pythagorean tone 203.9100 at $N = 3 \dots 6$, $50/49 = 34.9756$ cents at $N = 7 \dots 16$, $289/288 = 6.0008$ from $N = 17$ and $1682/1681 = 1.0296$ from $N = 41$, with a symbol-rate penalty of exactly 8 at the first cliff. Single-transposition audible capacity is $N \leq 25$ and only at $j = 6$. And the four-drawer classification — PROVED, POSTULATED, BORROWED, UNTESTED — prices the framework's central design axiom rather than defending it: an alternative surplus form satisfying every axiom but the affine one agrees at both endpoints yet disagrees by 33.62 per cent at $u = 0.5$, 1.871802 against 1.242453. The operationalisation of the twelve-state ladder is recorded as specified and not run.

Keywords: reference card; frequency tables; measurement protocol; minimum gap; transposition capacity; honesty appendix; postulate pricing; concordance register

Series: C. Frequency, Essence, and Mind **Paper:** CCP-050

1 Introduction

This paper is the reference and status apparatus for a framework whose results are established elsewhere. It is not a research contribution and is not offered as one. It exists because

the preceding papers of this series distribute their constants, tables, protocols and disclosures across ten documents, and a reader who wants to reproduce a reading, or to check what a given claim rests on, should not have to reconstruct that from ten places.

Almost everything in what follows is a restatement, and each restatement names the paper that owns the result. Where a number here differs in its last digit from a figure quoted in an owning paper, the full-precision rule applies: recompute from the full-precision input, never from a rounded one, and §13 records the two places where that rule catches something.

Four things here are not restatements, and they are the only items a citing paper should treat as originating in this document.

1. **The folded-lattice minimum gap as an exact function of ladder length** (§5.6). Folding every n^2 into one octave as an exact rational and taking the smallest adjacent ratio gives a step function with five regimes: $2/1$ at $N = 2$, $9/8$ (203.9100 cents) at $N = 3 \ldots 6$, $50/49$ (34.9756 cents) at $N = 7 \ldots 16$, $289/288$ (6.0008 cents) at $N \geq 17$, and $1682/1681$ (1.0296 cents) at $N \geq 41$. **The symbol-rate penalty at the first cliff is exactly 8.**
2. **The single-transposition audible capacity** (§5.7): at most 25 essences fit the audible band under one integer octave transposition, and only at $j = 6$.
3. **The pricing of the framework's central design axiom** (§8.2). The affine effective-count axiom is chosen, not forced, and an alternative satisfying every other axiom and agreeing at both endpoints **disagrees by 33.62 per cent in the middle of the range**.
4. **The four-drawer evidential classification** (§8), which sorts every claim the series makes into PROVED, POSTULATED, BORROWED and UNTESTED, and records the framework's own operationalisation as **SPECIFIED, NOT RUN**.

What a reference card costs, stated at the front. A card invites a reader to skip the argument, and the argument is where the limits live. Every card below therefore carries its limits inline, and the most important one belongs here rather than at the end: **the series' central instrument, the octave deficit, is blind to one of the four parameters that decide a relation's fate.** It contains no decay term, so a report giving the deficit alone cannot see decay, and Protocol C7 makes reporting both a numbered step rather than a caution.

2 Hypothesis

This paper advances no testable hypothesis. It is a reference collection and a status classification, and it is reported as such. Nothing in the sections that follow is a claim about the world that observation could confirm or refute; the tables are arithmetic, the protocols are procedures, and the classification is a record of what the framework has and has not paid for.

The organising question is therefore not empirical but bibliographic and evidential: **what must a reader hold in order to reproduce any reading this series reports, and on what does each of the series' claims actually rest?** The first half is answered by Sections 4 to 6 and the second by Sections 7 and 8.

Two consequences follow and are stated so that no reader mistakes the genre.

No result here is offered as new evidence for anything. The four items named in §1 as originating here are computations — three exact-arithmetic results and one classification — and none of them tests the framework. §8.2's pricing of the affine axiom is the closest thing to a critical result in the paper, and what it establishes is the *size of a choice*, not that the choice is wrong.

Where the framework's own hypotheses are recorded here, they are recorded as untested. §8.4 lists them with their kill conditions: the operationalisation of the twelve-state ladder is **specified and not run**; the acoustic causal claim is **not confirmed**; the decisive experiment is **designed and not run**; and the framework's own estimate of its research programme's completion probability is **35.8 per cent**. Those are the framework's numbers, restated, not this paper's findings.

3 Scope, framework, and dimensional audit

3.1 What this paper restates and what it does not

section	content	owning paper(s)
4	constants, three transfer rules, keystone, deficit, ceiling, levers, unit conversions	[Curzi, 2026,?,?]
5.1–5.5	the twelve-state ladder, raw and folded frequencies, the keystone table, the interval lattice, class counts	[Curzi, 2026,?]
5.6–5.7	**the minimum-gap schedule and the transposition capacity**	**originating here**
6	eight measurement protocols and four worksheets	[Curzi, 2026,?,?]
7	the concordance register, with rejections	[Curzi, 2026,?]
8	**the four-drawer classification and the pricing of the affine axiom**	**originating here**
9	correspondence to the two predecessor series	[Curzi, 2026,?,?]

Nothing in Sections 4, 5.1–5.5, 6, 7 or 9 should be cited to this paper where the owning paper can be cited instead.

3.2 The constant, its status, and the two claims it carries

$k = (9/8)^2/\varphi^2 = 0.483425732988414$ is a **defined constant, exact.** It is a ratio of exact quantities and therefore computable to arbitrary precision, so the fifteen significant figures quoted throughout this paper are *digits of a definition* and carry no measurement uncertainty because there is no measurement. $k^2 = 0.233700439315386$ inherits the same status. **Every appearance of either at that precision is so labelled.**

Attaching hertz to it is a separate and much stronger claim. Computing the ratio is arithmetic; asserting that a dimensionless number *is a frequency* requires an operational procedure by which the constant acquires the dimension of inverse time, and **the series supplies none.** That claim is badged **POSTULATED** here, in §4.1, in the verification chain of §12, and in §13.

The two components are real and documented — the Pythagorean whole tone [Barbour, 1951] and the golden section [Euclid, 1956, Vajda, 1989] — but **their product expressed as a frequency has no precedent in any literature we have found.** The deflationary literature on golden-section claims [Markowsky, 1992, Fischler, 1981, Livio, 2002] is cited here deliberately and again wherever φ appears.

One property makes the postulate less damaging than it looks, and it is why the tables below are usable. The essence bridge has the property that $f(m)/f(n) = m^2/n^2$, so k cancels from every interval, every ratio and every deficit [Curzi, 2026k]. **Every verdict in this series**

is a difference of logarithms in which the constant cancels; only the absolute hertz columns depend on it.

3.3 The exponent, and what was not done

The essence law $f(n) = kn^2$ is a **postulate of this framework with no external support**, and **the exponent 2 is defined by the model, not estimated from data.** No exponent anywhere in this series was obtained by fitting, none was tested against a log-normal or exponential alternative, and the goodness-of-fit and likelihood-ratio procedure that a claimed scaling relation requires [Clauset et al., 2009] was **not run, because there is no data to run it on.** The same applies to the exponent 4 in the compound rule [Curzi, 2026].

3.4 Octave reduction

Octave equivalence is a perceptual structure formalised for **pitch** [Shepard, 1982], whose perceptual specialness has a measured size [Demany et al., 2013]. **Its extension to a quantity that is not a pitch is a modelling choice of this framework with no external warrant.** The *operation* — division by a power of two — is well defined, and every folding in §5 is exact arithmetic; the *interpretation* of a folded value as carrying meaning is not supported. Cents, octave equivalence and exact rational interval arithmetic are standard apparatus once an index exists [Barbour, 1951, Lindley and Turner-Smith, 1993, Benson, 2006, Rossing et al., 2002], but what a listener finds consonant depends on the spectrum of a sound and not on the ratio alone [Sethares, 2005, Plomp and Levelt, 1965], and interval categories generalise less far than is usually assumed [Burns, 1999]. **This paper uses the arithmetic and claims none of the perception.** Two established measures of ratio complexity exist — Euler's *gradus suavitatis* [Euler, 1739] and harmonic distance [Tenney, 2015] — and **neither is used here**; §5.6 measures gaps in cents, which is a distance in a logarithmic register and carries no claim about consonance.

3.5 Dimensional audit

symbol	quantity	dimension	unit	status
n	essence index	$[1]$	count	cardinality, absolute scale
N	category count	$[1]$	count	cardinality, absolute scale
u	interaction parameter	$[1]$	dimensionless, $[0, 1]$	derived from unit vectors
w	elicited intensity	$[1]$	instrument units	**ordinal by default** (§6, Protocol C2)
F	surplus	$[1]$	nats	derived; **nats are not a physical unit**
η	conversion efficiency	$[1]$	dimensionless	declared
c	cost flow	$[\mathcal{C}\,T^{-1}]$	per declared period	declared
δ	loss rate	$[1]$	per declared period	declared
Q	stored order	$[\mathcal{C}]$	**no operational unit exists**	state variable

symbol	quantity	dimension	unit	status
D	octave deficit	$[1]$	octaves, cents, bits	derived
j	transposition exponent	$[1]$	integer	register choice, not a parameter
G	count-weighted geometric mean	$[1]$	index units	derived
k	ground essence	**claimed** $[T^{-1}]$	Hz **by postulate**	**defined; the unit is asserted**
k^2G^4	compound value	$[T^{-2}]$ **if taken seriously**	*convention-Hz*	derived; the label is a convention
$\Delta t\,\Delta f$	Gabor product	$[1]$	dimensionless	**classical Fourier; measurable**

Where the audit closes and where it does not. Every dimensionless quantity closes exactly, and so does the Gabor bound, which is a fact about Fourier duality with no quantum content [Gabor, 1946] and is measurable on an oscilloscope. **Two places do not close.** $f(n) = kn^2$ balances only if k carries $[T^{-1}]$, which is the postulate of §3.2; and the compound value is dimensionally hertz-squared, so its "Hz" label is a convention and not a unit. Q **has no operational unit at all**, and the dimensionless placeholder in the table is the honest record of that rather than an evasion.

3.6 Scale type, and the one place this series is exposed hardest

Permissible arithmetic is a function of the transformation group that leaves a quantity's empirical content unchanged [Stevens, 1946], and this series carries **two quantities of different scale status which must never be given a single verdict**.

n is a **count**: zero components is no state, so it has a true zero and equal units, and it is ratio-scaled by construction. w, the elicited intensities, are **ordinal by default**, and u, F, f_{rel} and the deficit inherit their status. **So the state frequency is ratio-scaled and the relationship frequency is not, until the intensities are.** Protocol C2 records why the usual fallback — report the output ordinally and carry on — is unavailable: a monotone relabelling that destroys no ordinal information whatever moves the relationship frequency by 1065.0568 cents, nearly an octave.

The instrument the series leans on is magnitude estimation, which is *held* to deliver ratio-scale information [Stevens, 1957] and whose exponents move with the experimenter's chosen range [Teghtsoonian, 1971], whose response biases are catalogued [Poulton, 1968], and about which the dispute is live [Krueger, 1989, Luce, 2002, Narens, 1996]. **The series inherits that dispute whole.**

And Michell's objection stands unanswered. That psychology assumes quantitative structure without testing for it [Michell, 1997, 1999] is the single strongest objection to this framework. The formal test for additive quantitative structure exists [Luce and Tukey, 1964, Krantz et al., 1971] and **was not run on any quantity in this series**; the instrument that would repair the framework's own demotion of one of its constructions [Curzi, 2026t] is Rasch measurement [Rasch, 1960], priced by the framework at 123 items per term and **not spent.** §8.2 and §13 record this as the framework's largest standing liability.

3.7 Identification, stated once because the protocols depend on it

Protocols C3 to C7 produce readings of a reflexive specification: the parties constitute the relationship and the relationship conditions the parties. **The endogenous social effect in such a**

specification is not identified from observational data [Manski, 1993]; it cannot be separated from contextual effects or from correlated unobservables however much data is collected, and influence and selection are generically confounded in observational network data [Shalizi and Thomas, 2011]. **The remedy the protocols adopt is stipulation of the reference set** [Brock and Durlauf, 2001] — the parties, the item set, the period and the vectors are declared and published — **and the price is that a reading is a modelling position rather than an identified empirical result.** That statement belongs on every worksheet and it belongs here.

3.8 Method

No data were collected for this paper and nothing was fitted. Three procedures produced everything originating here.

1. **Exact rational arithmetic.** §5.6's minimum-gap schedule was computed with exact fractions rather than floating point, because the quantities compared differ in the fourth decimal place of a cent. The four regime boundaries and the two cliff locations are exact and depend on k not at all.
2. **Exhaustive search over a bounded set.** §5.7's transposition capacity was computed over $j = 0 \dots 15$ for every N, and §5.5's class count was checked at every N from 1 to 25 against the closed form $\lceil N/2 \rceil$.
3. **Restatement with attribution.** Every other number is recomputed from the owning paper's stated inputs at full precision and cited to the paper that owns it.

Every figure in this paper is generated by a script named in its caption. No model comparison is reported, no statistical test is run, and no effect size or interval is given, because nothing here is an empirical claim.

4 The reference card: constants, transfer rules, the keystone, the levers

Printable: one page of constants, one of laws, one of conversions, one of clauses. A reader with the four pages and nothing else can compute every reading in this series and state correctly what the reading does not mean.

4.1 The constants

```
THE GROUND ESSENCE

    k  =  (9/8)^2 / phi^2                  phi = (1 + sqrt 5)/2 = 1.6180339887498949

    k  =  0.483425732988414                (computed: 0.48342573298841435)
                                           DEFINED, exact
    the unit "Hz" on k                     a SEPARATE claim: POSTULATED

    k^2 = 0.2337004393                     (computed: 0.23370043931538567)

THE CATEGORY COUNT USED THROUGHOUT THIS SERIES

    N       = 12
    ln N    = 2.484907                     (computed: 2.4849066497880004)
    k N^2   = 69.6133 Hz                   (computed: 69.61330555033166)

DERIVED CONSTANTS THAT RECUR

    1/sqrt(k)      = 1.438252434444351     CCP-046 crossing; not an integer
```

```
2 / ln 2       = 2.8853900817779268      octaves per nat
1 / (4 pi)     = 0.07957747154594767     the Gabor product
32 / (4 pi)    = 2.5464790894703255      cycles per resolved step in the 50/49 regime
```

`phi` is stored to full double precision throughout. The essence-law errata of Series B [Curzi, 2026k] traced a residual bias in the published element table to `phi` held at fifteen significant figures: carry the constant, not its printed form.

Two claims live in the first block and they have different statuses, so they are separated here and again in the verification chain of §12. Computing $(9/8)^2/\varphi^2 = 0.483425732988414$ is **arithmetic**, and the constant is **DEFINED and exact**: it is a ratio of exact quantities, so its fifteen figures are digits of a definition and carry no measurement uncertainty. **Writing "Hz" after it is a separate and much stronger claim, and it is POSTULATED** — no operational procedure is supplied by which a dimensionless ratio acquires the dimension of inverse time. The same split applies to k^2. The card prints the unit because the series prints it; the badge travels with it wherever it appears.

`k` sets the reference pitch and is invisible to every verdict. The cancellation of the constant [Curzi, 2026] proves it: the octave deficit is a difference of logarithms in which `k` appears in both terms and cancels. Only absolute hertz values depend on the constant. Sustainability verdicts do not.

4.2 The three transfer rules

```
RULE 1 -- A PSYCHOLOGICAL ESSENCE HAS A FREQUENCY          owner: CCP-041, CCP-042

    f  =  k * n^2            n = the number of distinguishable components required
                                 to specify the state

    n is a COUNT: true zero, equal units, ratio-scaled by construction.
    The frequency inherits exactly the defensibility of the count.

RULE 2 -- A PSYCHOLOGICAL RELATIONSHIP HAS A FREQUENCY     owner: CCP-043

    cos theta = x . y        x, y unit vectors over the N states
    u         = 1 - (x . y)^2
    F         = ln(1 + (N-1) u)          = ln(1 + 11 u) at N = 12
    f_rel     = k * e^(2F)  =  k * (1 + 11 u)^2  =  k * N_eff^2

    u = 0  ->  f = k, the ground essence: no relationship at all
    u = 1  ->  f = k N^2 = 69.6133 Hz, the ceiling

RULE 3 -- A COMPOUND STATE USES THE ALCHEMICAL RULE        owner: CCP-046

    f_compound  =  k^2 * G^4         G = count-weighted geometric mean of indices
                                     G = (prod n_i^{m_i})^(1/T),  T = sum m_i

    Values are convention-Hz and are dimensionally Hz^2. They may not be rendered
    as hertz and may not enter an acoustic calculation.
```

Rule 3 is not a generalisation of Rule 1 and the arithmetic says so twice. For a single component `G = n` and `f_compound = k^2 n^4 = f(n)^2`, which is the single-component identity [Curzi, 2026]; and **The no-agreement corollary CC-R4 [Curzi, 2026] shows the two rules agree at no index whatever, since `k n^2 = k^2 n^4` requires `n = 1/sqrt(k) = 1.438252`, which is not an integer.** At `n = 12` the two rules differ by 6.121291 octaves. The single-component index-independence of Rule 3's proportional structure is exactly what condemned the source table's O2 row in the essence-law errata of Series B [Curzi, 2026k].

Worked compounds, recomputed here to ten places for citation:

```
multiset                T        P           G                f = k^2 G^4
-------------------------------------------------------------------------------
{1x1, 2x1}              2        2     1.4142135624             0.9348017573
{1x3, 12x1}             4       12     1.8612097182             2.8044052718
{1x1, 12x1}             2       12     3.4641016151            33.6528632614
{2x2, 7x1}              3       28     3.0365889719            19.8702609489
{3x1, 5x1}              2       15     3.8729833462            52.5825988460
{4x1, 6x1, 9x1}         3      216     6.0000000000           302.8757693527
{5x1, 10x1}             2       50     7.0710678119           584.2510982885
{11x1, 12x1}            2      132    11.4891252931          4071.9964546313
-------------------------------------------------------------------------------
```

Rule 3 is blind to scale and acutely sensitive to ratio: trebling every multiplicity leaves `G` unchanged, while two multisets differing only in how many index-1 components are present sit 3.584963 octaves apart.

4.3 The keystone

```
THE MINIMUM SUSTAINABLE FREQUENCY                          owner: CCP-041

    f_min  =  k * e^( 2c / eta )

    f  >  f_min      the relation ACCUMULATES order       Q_inf > 0
    f  =  f_min      it HOLDS, and nothing more            Q_inf = 0
    f  <  f_min      it DECAYS, whatever anyone intends    Q_inf < 0

THE DERIVATION, IN THREE STEPS

    sustainability   Q_inf > 0   iff   eta F > c   iff   F > c/eta
    the bridge       F = (1/2) ln( f / k )
    substitute       f > k e^(2c/eta)
```

`f_min` contains no `delta` and no `N`, and both absences are substantive: the first is the source of the delta-blindness treated in §4.4, and the second means the threshold is set by cost and conversion alone, `N` entering only through the ceiling on what `F` can reach.

In octaves above the ground essence:

```
octaves required  =  2 (c/eta) / ln 2  =  2.885390 * (c/eta)
```

Every unit of cost-to-conversion ratio costs 2.885390 octaves of required pitch, and there is no other exchange rate.

4.4 The octave deficit, and the two theorems that govern it

```
THE OCTAVE DEFICIT                                    owner: CCP-041, Def 4.4

    D  :=  2.885390 * (c/eta)  -  log2( f / k )

    D <= 0    sustainable as it stands
    D  > 0    the relation is short by D octaves

THEOREM 4.5 -- THE DEFICIT IDENTITY

    D  =  2.885390 * ( c/eta  -  F )  =  -2.885390 * (delta/eta) * Q_inf

    Q_inf  =  ( eta F  -  c ) / delta            CCP-041 Theorem 3.5
    F      =  ln( 1 + 11 u )
```

Verified on the carried dyad as found: `D = 4.338420197297714` by the definition and `4.338420197297713` by both identities, a worst residual of `8.88e-16`. the verification record [Curzi, 2026] records independent verification across twelve parameter sets to `2.665e-15`.

```
THEOREM 4.9 -- THE DEFICIT CONTAINS NO delta

    D  =  2.885390 * (c/eta) - log2(f/k)     contains k, c, eta, N, u
                                             does NOT contain delta
```

This is the card's most important qualification and the one most easily lost when the card is used without the owning paper. Worked on the carried dyad at `eta = 0.30`, `c = 0.62`:

```
delta = 0.09   ->   Q_inf = -5.0119395464      D = +4.3384201973
delta = 0.05   ->   Q_inf = -9.0214911835      D = +4.3384201973
```

Trust nearly doubles the negative fixed point and the deficit does not move at all. The deficit reads the sign and size of the sustainability gap; `delta` converts that gap into a level. A reading that reports `D` without `delta` cannot see decay, and it is incomplete. Report both, always.

4.5 The ceiling, the minimum category count, and the hold line

```
THE CEILING            CCP-041 Cor 3.4 in the frequency register       owner: CCP-041 Thm 4.6

    Since F <= ln N,        c/eta  <=  ln N        and       f_min  <=  k N^2

    At N = 12:  c/eta may not exceed 2.484907, i.e. f_min may not exceed 69.6133 Hz.
    Above that, NO arrangement of those twelve categories is sustainable --
    not a better pairing, not a better schedule, nothing.
    The only remedies: raise eta, cut c, or add genuinely distinct categories.

THE MINIMUM CATEGORY COUNT                                          owner: CCP-041 Cor 4.7

    N  >=  e^(c/eta)

THE HOLD LINE                                                       owner: CCP-041 Cor 4.8

    u*  =  ( e^(c/eta) - 1 ) / (N - 1)
```

The ceiling test comes first. A deficit computed above the ceiling is a number with no verdict attached: no arrangement of the declared category set can close it.

4.6 The four levers

```
LEVER              acts on            partial derivative of D
-------------------------------------------------------------------------------
L1  raise u        the geometry       dD/du    = -2.885390 (N-1) / (1 + (N-1)u)
L2  raise eta      the conversion     dD/deta  = -2.885390 c / eta^2
L3  cut c          the cost           dD/dc    = +2.885390 / eta
L4  raise N        the category set   dD/dN    = -2.885390 u / (1 + (N-1)u)
```

These are the only four, and `delta` is not one of them (the decay-blindness result CC-R4 [Curzi, 2026]). `delta` moves the level, not the gap.

```
THE ORDER-OF-OPERATIONS RULE (R15)                          owner: CCP-044

    FIX THE NUMERATOR FIRST.

    Trust acts by LOWERING delta. Q_inf = (eta F - c)/delta.
    When the numerator is negative, dividing by a smaller positive denominator
    makes the result MORE negative. Trust amplifies whatever sign the system
    already has; it is not a source of order.
```

4.7 The agreement identity

```
THEOREM 13.1                                                owner: CCP-044

    Two parties judge m shared items, each judgement coded +1 or -1.
    Let a = fraction of items on which the judgements match. Then

        cos theta  =  2a - 1              u  =  4 a (1 - a)

    Exact; no centring required, since +/-1 vectors have norm sqrt(m).
    Resolution limit:  du = 4 |1 - 2a| / m,  bounded by 4/m.
```

The identity is symmetric about `a = 0.5` and returns `u = 0` at both `a = 1` and `a = 0`, which is CC-R8 appearing directly in the measurement channel: the sycophant and the contrarian sound the identical note. Two parties answering at random have expected `a = 0.5`, hence `u = 1`, hence the ceiling of the whole system from guessing. The formalism cannot distinguish independence from noise, because both are orthogonality, and the guard is procedural: publish each party's accuracy with any `u`.

4.8 The Gabor limit

```
dt * df  >=  1 / (4 pi)  =  0.07957747154594767             owner: Series B / CCP-032

Classical Fourier duality. No quantum content. Measurable on an oscilloscope.

THE FOUR DESIGN LAWS
  (1) single-transposition audible capacity is N <= 25, and only at j = 6
  (2) the folded-lattice minimum-gap cliff is at N = 17 (50/49 -> 289/288),
      the second at N = 41, and the RATE factor is exactly 8
  (3) resolving a surplus difference dF costs at least 1/(8 pi dF) cycles
  (4) the rectangular gate is disqualified: no finite spectral second moment
```

`N = 12` sits below the `N = 17` cliff, so **this entire book renders in the 50/49 regime, at a minimum separation of 34.9756 cents and 2.546479 cycles per resolved step.** A 40-cent deficit is just above the minimum resolvable step of the lattice it is expressed in; a 6-cent one would not be.

4.9 Unit conversions

```
1 nat    =  1.442695 bits    =  2.885390 octaves  =  3462.4681 cents
1 bit    =  2 octaves EXACTLY                     =  2400 cents
1 octave =  1200 cents

computed:  1/ln2   = 1.4426950408889634
           2/ln2   = 2.8853900817779268
           2400/ln2 * (ln2/2) ... 1200 * 2/ln2 = 3462.4680981335123
```

```
D in octaves  =  2.885390 * (c/eta - F)
D in cents    =  3462.4681 * (c/eta - F)
D in bits     =  1.442695 * (c/eta - F)
```

The corpus's closing synthesis computes the same exchange from the other direction, finding that pitch-class folding of the essence lattice costs exactly $\ln 2 = 0.6931471806$ nats — one bit, two octaves — for every even N [Curzi, 2026].

The bit-to-octave conversion is exact and not a coincidence: frequency enters as `f = k e^(2F)`, so one nat is two e-folds of frequency and one bit — being `ln 2` nats — is two doublings, which is two octaves. **The four registers of the conversion keys of Series B [Curzi, 2026l] converge here because they are all logarithms of the same effective count, taken to different bases.** The carried dyad as found reads `+4.3384 octaves = +5206.1 cents`; after its intervention, `−3.5910 octaves = −4309.2 cents`.

4.10 The clauses, which are part of the card and not an afterword

```
C1  Essence frequencies are ESSENCE frequencies, not physical resonances. No
    person, tissue or molecule is claimed to resonate at any value in this series.

C2  No clinical, diagnostic or therapeutic claim. The twelve states are not
    diagnoses; the readings are not assessments.

C3  No intervention may remove a person's ability to decline. Consent is
    structural, not advisory.

C4  No subliminals. Nothing masked, nothing below threshold, nothing concealed.

C5  The open problem stands: no established mechanism for sequence-specific
    acoustic demodulation.

C6  All worked cases are CONSTRUCTED and labelled as such.

C7  R1 binds D/P specifically. It does not bind the count index n, and it does
    not bind magnitude-estimated intensities. State which quantity you mean.

C8  Report D AND delta. The deficit alone cannot see decay (Theorem 4.9).

C9  A state frequency and a relationship frequency are independent readings and
    neither may be inferred from the other (result R7 of CCP-046).

C10 (+1) carries evidence; (-1) carries precedent. Precedent is never promoted.
```

4.11 The card, on one page

```
============================================================================
THE SUSTAINING FREQUENCY -- REFERENCE CARD
============================================================================
k = 0.483425732988414 = (9/8)^2/phi^2           k^2 = 0.2337004393
  DEFINED, exact.  The unit "Hz" on it is a SEPARATE claim, POSTULATED.
N = 12     ln N = 2.484907     k N^2 = 69.6133 Hz
----------------------------------------------------------------------------
STATE          f = k n^2                       n is a COUNT
RELATION       f_rel = k (1 + 11 u)^2          u = 1 - (x.y)^2
               F = ln(1 + 11 u)
```

```
COMPOUND       f_compound = k^2 G^4            convention-Hz, never rendered
-------------------------------------------------------------------------
KEYSTONE       f_min = k e^(2c/eta)
OCTAVES REQ    2.885390 * (c/eta)
DEFICIT        D = 2.885390*(c/eta) - log2(f/k)
IDENTITY       D = 2.885390*(c/eta - F) = -2.885390*(delta/eta)*Q_inf
DELTA-BLIND    D contains no delta.  REPORT D AND delta.
DESTINY        Q_inf = (eta F - c)/delta
CEILING        c/eta <= ln N          HOLD LINE  u* = (e^(c/eta)-1)/(N-1)
LEVERS         raise u | raise eta | cut c | raise N.  delta is not a lever.
ORDER          fix the numerator first (R15)
AGREEMENT      u = 4a(1-a),  cos theta = 2a - 1,  resolution 4/m
GABOR          dt df >= 1/(4 pi) = 0.0795774715
UNITS          1 nat = 1.442695 bits = 2.885390 octaves = 3462.4681 cents
               1 bit = 2 octaves exactly
-------------------------------------------------------------------------
NOT a resonance. NOT a diagnosis. NOT a capability. Consent is structural.
Nothing masked. The open problem stands. All cases constructed.
=========================================================================
```

---

5 Frequency, interval and lattice tables

Every table here was computed, not asserted. Where exact arithmetic matters — §5.6's minimum-gap analysis — the computation used exact rationals via `fractions.Fraction`, because the quantities compared differ in the fourth decimal place of a cent.

5.1 The twelve-state ladder This is the ladder exactly as the founding definitions [Curzi, 2026n] gives it, and it is the only ladder this series uses. The folded column transposes each raw frequency by whole octaves into `[220, 440)`; the pitch class is the odd part of the index, by the odd-part theorem of the mood-and-trait paper [Curzi, 2026].

n	state class	the component it adds	f = k n^2 (Hz)	folded 220-440	pitch class
1	undifferentiated activation	arousal alone	0.483426	247.5140	1
2	valenced activation	+ valence	1.933703	247.5140	1
3	directed feeling	+ intentional object	4.350832	278.4532	3
4	situated feeling	+ situation	7.734812	247.5140	1
5	attributed feeling	+ cause	12.085643	386.7406	5
6	temporal feeling	+ projected duration	17.403326	278.4532	3
7	reflexive feeling	+ self as object	23.687861	379.0058	7
8	social feeling	+ another mind modelled	30.939247	247.5140	1
9	normative feeling	+ standard applied	39.157484	313.2599	9

n	state class	the component it adds	f = k n^2 (Hz)	folded 220-440	pitch class
10	narrative feeling	+ placed in a story	48.342573	386.7406	5
11	reciprocal feeling	+ the other modelling you	58.494514	233.9781	11
12	committed feeling	+ future action bound	69.613306	278.4532	3

Twelve states, six distinct pitch classes `{1, 3, 5, 7, 9, 11}` — `ceil(12/2)`, as the odd-part theorem requires. Indices 1, 2, 4 and 8 share class 1 and fold to the identical audible frequency. **A folded reading that reports six notes from a twelve-state ladder has lost no information it was entitled to: the collision is a property of the arithmetic, not of the fold.**

The six classes, reduced to one octave and sorted by pitch:

```
class    folded ratio (exact)    cents above PC1    step to next
------------------------------------------------------------------
PC1              1                       0.00          203.91
PC3            9/8                     203.91          203.91
PC9           81/64                    407.82          329.83
PC7           49/32                    737.65           34.98
PC5           25/16                    772.63          330.01
PC11         121/64                   1102.64           97.36
------------------------------------------------------------------
steps sum to 1200.00 cents; the scale closes at the octave, as it must
```

PC7 and PC5 sit 34.98 cents apart — the `50/49` minimum gap of the `N = 12` regime, appearing inside the ladder itself. Sounded together they beat rather than form an interval.

5.2 `f = k n^2` for n = 1 to 25, raw and folded

Twenty-five is the single-transposition audible capacity of §5.7, and a table running past it would tabulate frequencies no one integer transposition can render.

```
  n     f raw (Hz)      j     folded [220,440)    odd part
-------------------------------------------------------------
  1       0.483426     +9         247.5140          1
  2       1.933703     +7         247.5140          1
  3       4.350832     +6         278.4532          3
  4       7.734812     +5         247.5140          1
  5      12.085643     +5         386.7406          5
  6      17.403326     +4         278.4532          3
  7      23.687861     +4         379.0058          7
  8      30.939247     +3         247.5140          1
  9      39.157484     +3         313.2599          9
 10      48.342573     +3         386.7406          5
 11      58.494514     +2         233.9781         11
 12      69.613306     +2         278.4532          3
 13      81.698949     +2         326.7958         13
 14      94.751444     +2         379.0058          7
 15     108.770790     +2         435.0832         15
 16     123.756988     +1         247.5140          1
```

```
 17      139.710037     +1         279.4201         17
 18      156.629937     +1         313.2599          9
 19      174.516690     +1         349.0334         19
 20      193.370293     +1         386.7406          5
 21      213.190748     +1         426.3815         21
 22      233.978055      0         233.9781         11
 23      255.732213      0         255.7322         23
 24      278.453222      0         278.4532          3
 25      302.141083      0         302.1411         25
------------------------------------------------------------
```

Index 17 folds to 279.4201, six cents from index 3's 278.4532. That is the `289/288` cliff arriving, and §5.6 prices it exactly.

5.3 `f_min` against `c/eta` The keystone table. Column three is `log2(f_min/k)`, which is the octave form; column four is the closed-form `2.885390 (c/eta)`. They agree to the printed precision at every row, which is the verification.

```
 c/eta      f_min (Hz)      octaves      2.885390*(c/eta)
------------------------------------------------------------
  0.00        0.483426      0.000000        0.000000
  0.10        0.590458      0.288539        0.288539
  0.20        0.721186      0.577078        0.577078
  0.30        0.880859      0.865617        0.865617
  0.40        1.075884      1.154156        1.154156
  0.50        1.314087      1.442695        1.442695
  0.60        1.605030      1.731234        1.731234
  0.70        1.960388      2.019773        2.019773
  0.80        2.394423      2.308312        2.308312
  0.90        2.924555      2.596851        2.596851
  1.00        3.572060      2.885390        2.885390
  1.10        4.362924      3.173929        3.173929
  1.20        5.328887      3.462468        3.462468
  1.30        6.508717      3.751007        3.751007
  1.40        7.949765      4.039546        4.039546
  1.50        9.709865      4.328085        4.328085
  1.60       11.859656      4.616624        4.616624
  1.70       14.485417      4.905163        4.905163
  1.80       17.692528      5.193702        5.193702
  1.90       21.609703      5.482241        5.482241
  2.00       26.394151      5.770780        5.770780
  2.10       32.237888      6.059319        6.059319
  2.20       39.375446      6.347858        6.347858
  2.30       48.093278      6.636397        6.636397
  2.40       58.741263      6.924936        6.924936
 [2.484907   69.613306      7.169925        7.169925]  <-- THE CEILING at N = 12
  2.50       71.746740      7.213475        7.213475
  2.60       87.631666      7.502014        7.502014
  2.70      107.033559      7.790553        7.790553
  2.80      130.731084      8.079092        8.079092
  2.90      159.675307      8.367631        8.367631
  3.00      195.027860      8.656170        8.656170
------------------------------------------------------------
```

The bracketed row is not a boundary a better arrangement can cross: at `c/eta = 2.484907`

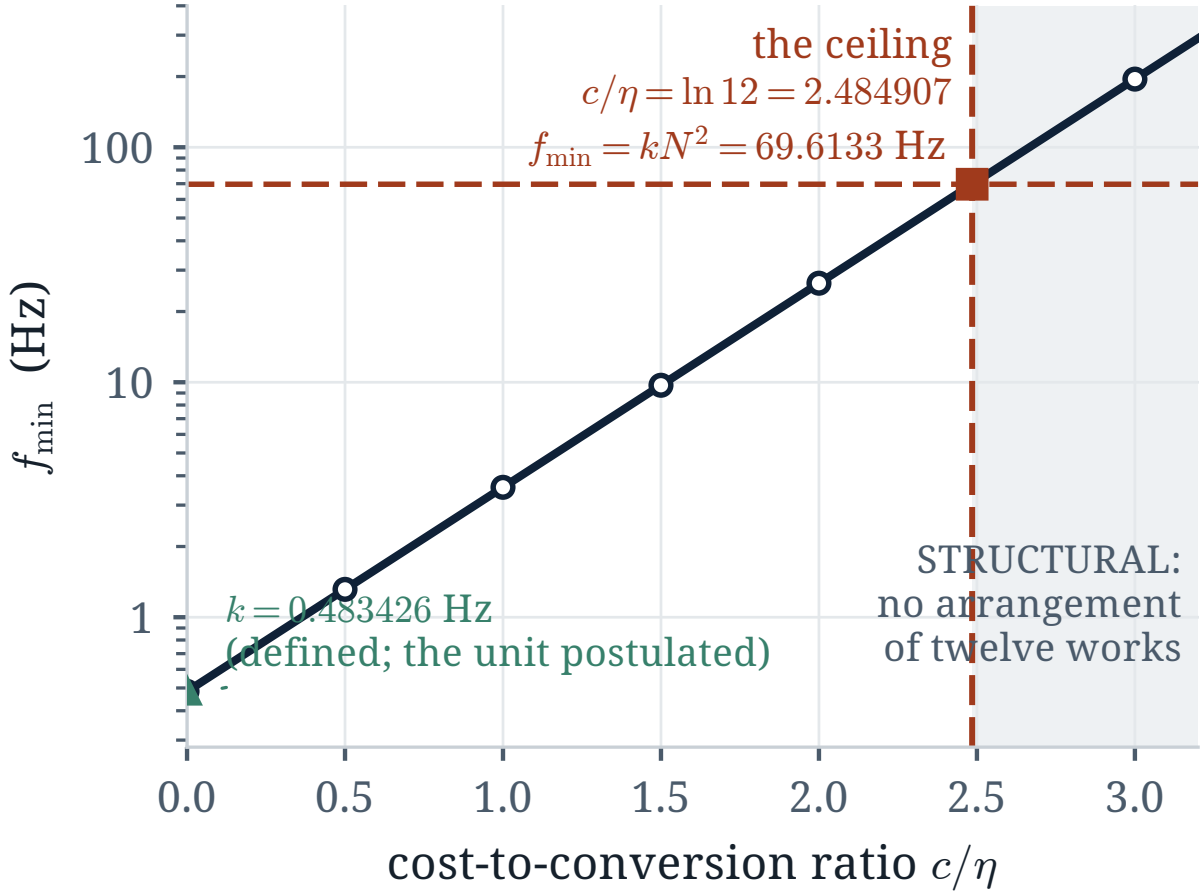


Figure 50.1: **Above the ceiling no arrangement of the declared categories is sustainable.** The keystone $f_{\min} = k\,e^{2c/\eta}$ against the cost-to-conversion ratio. At $c/\eta = \ln 12 = 2.484907$ the required non-parallelism is exactly 1 and $f_{\min}$ reaches $kN^2 = 69.6133$ Hz; beyond it the case is structural. k is a DEFINED constant, exact, and its hertz label is POSTULATED (§3.2). Values from the table of §5.3; script `figs/src/CCP-050_fig02.py`.

the required non-parallelism is exactly 1 — perfect orthogonality — and there is nothing above it.

The companion table is the hold line, stating the same threshold in the geometry rather than the cost:

c/eta	u* = (e^(c/eta)−1)/11	f_min (Hz)
0.000000	0.000000	0.483426
0.250000	0.025820	0.797034
0.500000	0.058975	1.314087
0.750000	0.101545	2.166564
1.000000	0.156207	3.572060
1.250000	0.226395	5.889331
1.500000	0.316517	9.709865
1.750000	0.432237	16.008862
2.000000	0.580823	26.394151
2.250000	0.771612	43.516598
2.484907	1.000000	69.613306

And the conversion between the two gauges, at the carried case's `delta/eta = 0.30`:

D (octaves)	f / f_min	Q_inf
−6.0	64.0000	+6.9315
−4.0	16.0000	+4.6210
−2.0	4.0000	+2.3105
−1.0	2.0000	+1.1552
0.0	1.0000	0.0000
+1.0	0.5000	−1.1552
+2.0	0.2500	−2.3105

This table is readable only because `delta/eta` was declared. Change `delta` and every `Q_inf` in the right-hand column moves while every `D` in the left-hand column does not — the decay-blindness result CC-R4 [Curzi, 2026] in tabular form, and the reason the ratio sits in the caption.

5.4 The interval lattice Intervals between states are ratios of squares, $f_m/f_n = m^2/n^2$, and k cancels. In cents, split into two halves for width; the matrix is antisymmetric, so the second half is the transpose of the first negated:

	n=1	n=2	n=3	n=4	n=5	n=6
m=1	0.00	−2400.00	−3803.91	−4800.00	−5572.63	−6203.91
m=2	2400.00	0.00	−1403.91	−2400.00	−3172.63	−3803.91
m=3	3803.91	1403.91	0.00	−996.09	−1768.72	−2400.00
m=4	4800.00	2400.00	996.09	0.00	−772.63	−1403.91
m=5	5572.63	3172.63	1768.72	772.63	0.00	−631.28
m=6	6203.91	3803.91	2400.00	1403.91	631.28	0.00
m=7	6737.65	4337.65	2933.74	1937.65	1165.02	533.74
m=8	7200.00	4800.00	3396.09	2400.00	1627.37	996.09
m=9	7607.82	5207.82	3803.91	2807.82	2035.19	1403.91
m=10	7972.63	5572.63	4168.72	3172.63	2400.00	1768.72
m=11	8302.64	5902.64	4498.73	3502.64	2730.01	2098.73
m=12	8603.91	6203.91	4800.00	3803.91	3031.28	2400.00

```
          n=7      n=8      n=9     n=10     n=11     n=12
m=1   −6737.65 −7200.00 −7607.82 −7972.63 −8302.64 −8603.91
m=2   −4337.65 −4800.00 −5207.82 −5572.63 −5902.64 −6203.91
m=3   −2933.74 −3396.09 −3803.91 −4168.72 −4498.73 −4800.00
m=4   −1937.65 −2400.00 −2807.82 −3172.63 −3502.64 −3803.91
m=5   −1165.02 −1627.37 −2035.19 −2400.00 −2730.01 −3031.28
m=6    −533.74  −996.09 −1403.91 −1768.72 −2098.73 −2400.00
m=7       0.00  −462.35  −870.17 −1234.98 −1564.98 −1866.26
m=8     462.35     0.00  −407.82  −772.63 −1102.64 −1403.91
m=9     870.17   407.82     0.00  −364.81  −694.82  −996.09
m=10   1234.98   772.63   364.81     0.00  −330.01  −631.28
m=11   1564.98  1102.64   694.82   330.01     0.00  −301.27
m=12   1866.26  1403.91   996.09   631.28   301.27     0.00
```

Three entries earn comment. `m/n = 2` gives exactly 2400 cents, because doubling an index quadruples a squared frequency — the odd-part theorem in interval form, and the reason indices 1, 2, 4 and 8 share a class. The 7:5 entry is 1165.02 cents, 35 cents shy of an octave, and is the pair that sets the `50/49` minimum gap. The 12:11 entry, 301.27 cents, is the smallest non-collapsing interval at the top of the ladder.

5.5 Pitch-class counts

```
N          1  2  3  4  5  6  7  8  9 10 11 12 13 ... 25
ceil(N/2)  1  1  2  2  3  3  4  4  5  5  6  6  7 ... 13
measured   1  1  2  2  3  3  4  4  5  5  6  6  7 ... 13
                                  ^ this series, N = 12
```

Checked programmatically at every row: the two columns are equal for all `N` from 1 to 25. **The class count is exactly `ceil(N/2)`, because the odd parts of `1..N` are precisely the odd numbers in `1..N`, of which there are `ceil(N/2)`.**

5.6 The folded-lattice minimum gap, by exact rational arithmetic Fold every `n^2` into a single octave as an exact `Fraction`, sort, and take the smallest adjacent ratio including the wrap-around from the largest back to twice the smallest. The result is a step function in `N` with four distinct values:

```
N range        minimum gap        cents         status
------------------------------------------------------------------------
N = 2              2/1         1200.0000       degenerate (one class)
N = 3 .. 6         9/8          203.9100       the Pythagorean tone
N = 7 .. 16       50/49          34.9756       THIS SERIES' REGIME
N >= 17          289/288          6.0008       THE CLIFF
N >= 41         1682/1681         1.0296       the second threshold
------------------------------------------------------------------------
```

The cliff fires the instant index 17 enters the set, because `17^2 = 289` folds beside `288 = 2 * 12^2`; the second fires at index 41, because `41^2 = 1681` folds beside `1682 = 2 * 29^2`. Both locations are exact and neither depends on `k`.

Three distinct ratios describe the size of the cliff, all correct in their own units. Confusing them is the specific error Series B's CC-R21 records and reverses:

```
cents    ln(50/49)/ln(289/288)   = 5.828    how much smaller the interval is
cycles   288/49                  = 5.878    3.899 cycles becomes 22.918
RATE     (1/32)/(1/256)          = 8 EXACT  symbols per second -- the design limit
```

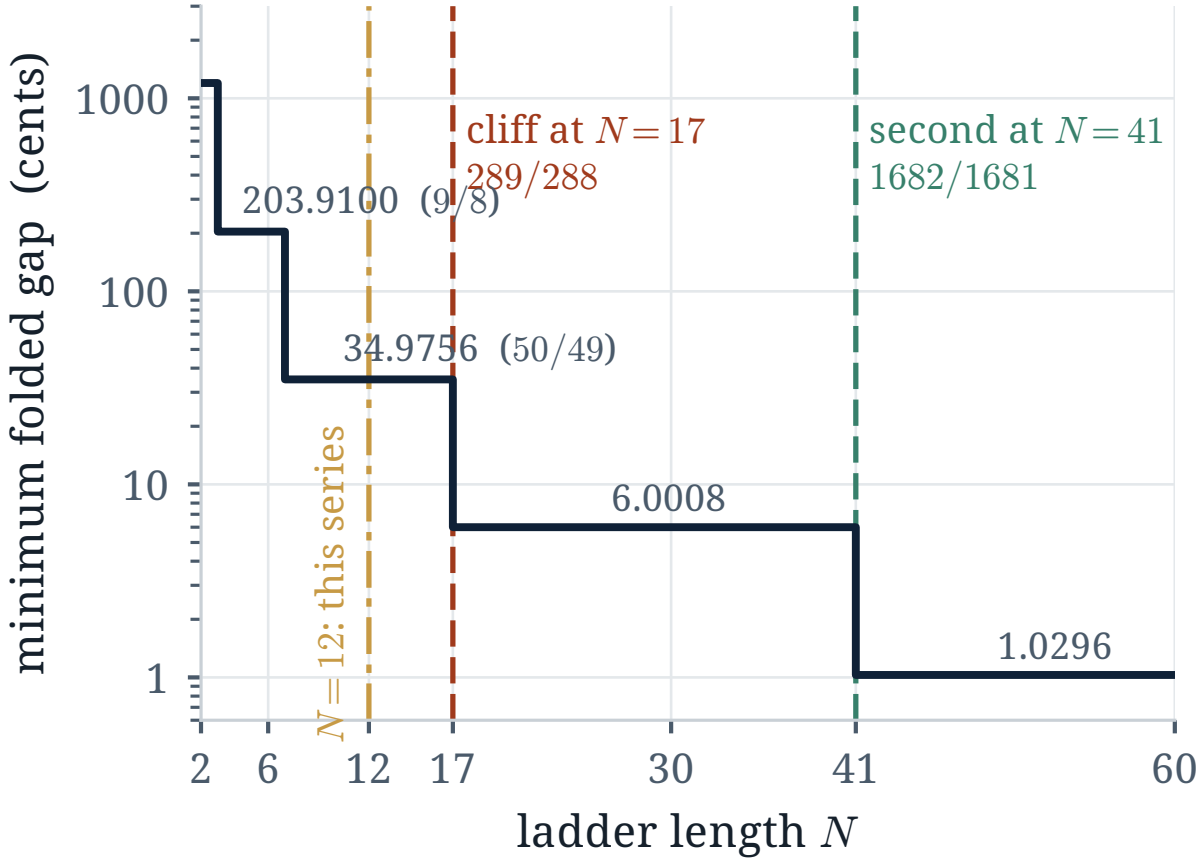


Figure 50.2: **The minimum separation the folded lattice offers is a step function of ladder length, and it collapses twice.** Computed with exact rationals: $9/8$ at $N = 3 \dots 6$, $50/49$ at $N = 7 \dots 16$, $289/288$ from $N = 17$ and $1682/1681$ from $N = 41$. Both cliff locations are exact and neither depends on the reference constant. $N = 12$ sits inside the third regime, so this series renders throughout at 34.9756 cents. Script `figs/src/CCP-050_fig01.py`.

The rate factor is exactly 8 because the two binding pairs sit at different octave heights: the `50/49` pair straddles `49/32` and the `289/288` pair straddles `9/8 = 288/256`. Gabor resolution depends on the absolute gap in hertz — `f/32` before the cliff, `f/256` after, at every base pitch. This series renders entirely below the cliff, so the `50/49` figures are operative throughout.

5.7 Single-transposition capacity Take the whole set `{k n^2 : n = 1..N}`, apply one integer octave transposition `2^j`, and require the entire transposed set to lie inside 20 Hz to 20 kHz. Computed over `j = 0..15` for every `N`:

```
N        admissible j
------------------------------
20            {6}
21            {6}
22            {6}
23            {6}
24            {6}
25            {6}          <-- the maximum
26            {}           <-- no j works
27            {}
------------------------------

at N = 25, j = 6:   lowest =    30.9392 Hz      highest = 19337.0293 Hz
at N = 26, j = 6:   highest = 20914.9309 Hz     -- above the band
at N = 26, j = 5:   lowest  =    15.4696 Hz     -- below the band
```

A single integer octave transposition admits at most 25 essences inside the audible band, and only at `j = 6`. A twelve-house set fits comfortably; a thirty-two-category set does not and requires split transposition. This is a hard bound on essence-set size and it constrains any extension of the ladder a reader might be tempted to make.

---

6 Measurement protocols and worksheets

Eight numbered protocols and four worksheets. Each names its owning paper, states its inputs and required outputs, and ends with the gates that must pass before publication. Reports cite them by number and do not restate them.

A prohibition governs all eight and is repeated inside each: **none of these procedures is an assessment.** Nothing here evaluates, ranks or screens a person, or produces anything describable to them or anyone else as a finding about their mind. They produce a description of a description, and a practitioner who cannot hold that distinction should not run them.

Protocol C1 — Indexing a state *Owner: the state-ordering and elicitation material [Curzi, 2026].*

Inputs. One named situation, in one sentence, written before anything else is asked. One or more framings of it. A dated declaration of the decomposition rules in force, `N = 12`, the ladder in use, the intensity instrument and its range, the audible band and the folding convention.

Steps.

```
C1.1  DECLARE THE RULE SET, DATED. Published with the reading.
C1.2  OBTAIN CONSENT, AND MAKE DECLINING STRUCTURAL. The subject is told what is
      elicited, what is computed, what it does and does not mean, and that any
```

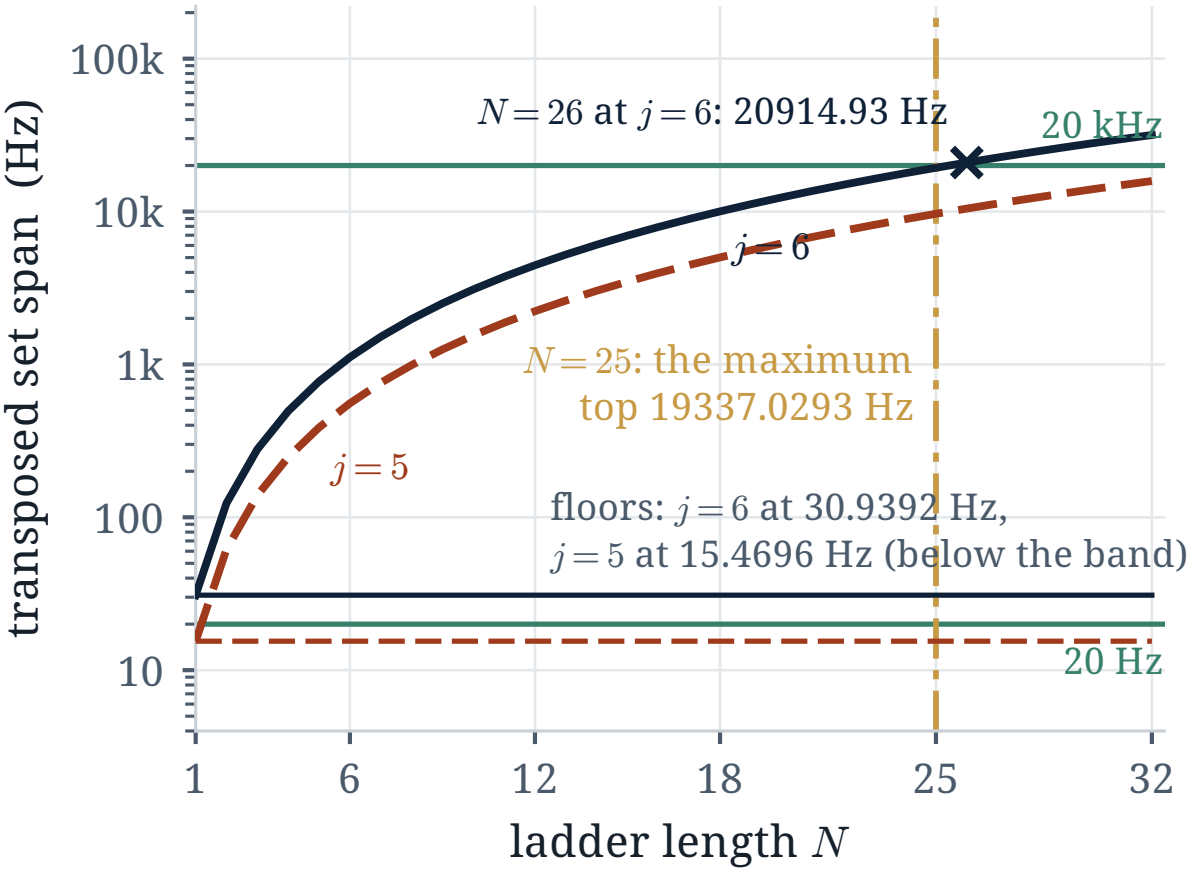


Figure 50.3: **The audible band admits at most twenty-five essences under one integer transposition.** The transposed set's floor and ceiling against ladder length at $j = 6$ and $j = 5$. At $j = 6$ the top crosses 20 kHz between $N = 25$ (19337.0293 Hz) and $N = 26$ (20914.9309 Hz); dropping to $j = 5$ puts the floor at 15.4696 Hz, below the band. Computed by exhaustive search over $j = 0 \dots 15$; script `figs/src/CCP-050_fig04.py`.

```
      question may be skipped without a reason and without the procedure being
      abandoned. Nothing masked; nothing concealed; nothing below awareness.
C1.3  FIX THE SITUATION in one written sentence.
C1.4  FIX THE FRAMINGS to be elicited, and name them, before eliciting any.
      Two is the minimum for a relational reading; three is the house default.
C1.5  SILENT WRITE-FIRST. Each framing written out in full, independently,
      before any two are compared and before the practitioner comments.
C1.6  READ EACH FRAMING AGAINST ALL TWELVE COMPONENTS. PRESENT or ABSENT for
      every one. No blanks.
C1.7  SUBSTITUTION TEST on every component marked PRESENT: remove it, read the
      residual back, ask whether that is still the state. If yes, unmark it.
C1.8  NESTING CHECK. No component marked present above one marked absent.
C1.9  COUNT. n = components present after C1.7 and C1.8. Record the count AND
      the list. A count without its list is a rating and is rejected.
```

Required outputs. `n` for each framing; the component list beside each count; `f = k n^2` raw and folded with `j`; the pitch class as the odd part of `n`.

Publication requirements, from the publication requirements [Curzi, 2026], and these are requirements.

```
I1  THE DECOMPOSITION -- what counts as one component, as a rule a third party
    can apply to a state the author never considered
I2  THE COUNT RULE -- including redundant, nested and conjoined components
I3  N, with the ceiling ln N it implies
I4  THE TIE-BREAK -- for a count index, ties are the same state class and share
    a frequency, which is the intended behaviour
I5  THE BAND AND j for every published frequency
I6  THE PERMUTATION TEST -- swap two adjacent indices, recompute, report the
    change in cents and in pitch class. Ten minutes.
```

The honest limit, stated wherever the index is used. The frequency inherits exactly the defensibility of the count. A contested decomposition yields a contested frequency, and a frequency published without its decomposition is a number nobody can check.

Protocol C2 — Eliciting the intensities, and why magnitude estimation is required

Owner: the elicitation rule [Curzi, 2026]. Binding on every paper that reports a `u`.

This requirement is the one most often assumed negotiable, so the argument comes before the steps.

The scale-status rule fixes the standing of two quantities separately, and they differ. **`n`, the index, is a count: it has a true zero and equal units, it is ratio-scaled by construction, and it does not inherit CC-R1's demotion — it sidesteps it.** `w`, the intensities, are elicited magnitudes and are ordinal by default, and everything computed from intensity-weighted vectors — `u`, `F`, `f_rel`, the deficit — inherits `w`'s status. So the state frequency is ratio-scaled and the relationship frequency is not, until the intensities are.

The usual response to an ordinal input is to report the output ordinally and carry on. **That response is not available here, and the elicitation rule [Curzi, 2026] proves it with a number.** Apply a monotone relabelling of the intensities — a transform that preserves every ordinal relation and therefore destroys no ordinal information whatever — and `u` moves from 0.362115 to 0.525275. Converted through Rule 2:

```
u = 0.362115   ->   F = 1.606085   ->   f_rel = 12.004877 Hz
u = 0.525275   ->   F = 1.913686   ->   f_rel = 22.209363 Hz

shift  =  0.887547 octaves  =  1065.0568 cents  =  0.307600 nats
```

Nearly an octave, from a change an ordinal scale calls no change at all. **With ordinal intensities the relationship frequency is not determined, and reporting one would be reporting an artifact of the labelling.** The ordinal fallback is therefore not a weaker option; it is not an option.

Steps.

```
C2.1  Choose the state class the subject reports most strongly. Assign it a
      reference value -- 10 on a 0-10 instrument.
C2.2  For every other class present, ask the ratio question IN WORDS:
      "compared with that one, how strong is this one -- half as strong,
      a quarter, three times?" Ratios are what the construction consumes.
C2.3  Do not average across framings. Each framing is elicited separately.
C2.4  Do not show either party another party's numbers.
C2.5  RECORD THE RAW L2 NORM in the instrument's units, then normalise.
C2.6  PUBLISH THE RAW L2 NORMS. The free multiplier of magnitude estimation is
      harmless because normalisation removes it -- but it must stay recoverable,
      and it is recoverable only from the norms.
```

Worked, on the carried dyad, so the reporting rule has a concrete instance:

```
A = {4:0.7, 6:0.5, 9:0.4}            raw ||A||_2 = 0.948683
B = {4:0.6, 6:0.55, 9:0.3, 10:0.2}   raw ||B||_2 = 0.890225
                                     ratio ||A||/||B|| = 1.065667
```

The ratio is the recoverable content: normalisation discards it and the publication rule puts it back. **Magnitude estimation (Stevens' method) produces genuine ratio-scale subjective magnitudes up to a free multiplier and replicates, which is why it is required and not merely preferred.**

Protocol C3 — The three routes to u

Owner: the relationship-frequency paper [Curzi, 2026]; used by [Curzi, 2026,?].

Three routes, measuring different things, and the route is declared before the numbers are seen. Choosing it afterwards is a documented failure mode (FM-13.5).

```
ROUTE A -- STATE VECTORS                                    the primary route

    Represent each party as a unit vector over the twelve states, components =
    magnitude-estimated intensity, normalised to length 1.

        cos theta = x . y            u = 1 - (x . y)^2

    Carry u to twelve or more decimal places (full-precision rule).
    Requires: C1 and C2 completed for both parties.
    Cannot see: whether either party is right about anything.

ROUTE B -- THE AGREEMENT CHANNEL                              CCP-044 R2

    Both parties judge m shared items independently, each judgement +1 or -1.
    Let a = the fraction of items on which they match.

        cos theta = 2a - 1           u = 4 a (1 - a)

    Exact, and no centring is required: +/-1 vectors have norm sqrt(m), so the
    norms cancel identically.
```

```
    Resolution: du = 4|1−2a|/m, bounded above by 4/m. REPORT m.
    Requires: an item set agreed BEFORE it is judged.
    Cannot see: the sign. a = 0.9 and a = 0.1 both give u = 0.36.

ROUTE C -- ERROR DECORRELATION                              Definition 13.2

    Both parties estimate m quantities whose true values later arrive. Form each
    party's error vector e = (estimate − truth), let rho be the Pearson
    correlation of the two error vectors.

        u = 1 − rho^2

    REPORT THE SIGN OF rho EXPLICITLY. It is not recoverable from u.
    Requires: ground truth, which is the cost and it is not small.
    Can see, and the others cannot: being right for different reasons, and
    being wrong in different ways.
```

The agreement channel, so a practitioner can read `u` off an agreement rate directly:

a	cos theta	u	f_rel (Hz)	cents above k
1.000000	+1.000000	0.000000	0.4834	0.00
0.990000	+0.980000	0.039600	0.9963	1251.97
0.950000	+0.900000	0.190000	4.6158	3906.26
0.900000	+0.800000	0.360000	11.8930	5544.82
0.800000	+0.600000	0.640000	31.2494	7217.27
0.700000	+0.400000	0.840000	50.6909	8054.75
0.600000	+0.200000	0.960000	64.6019	8474.57
0.500000	0.000000	1.000000	69.6133	8603.91
0.400000	−0.200000	0.960000	64.6019	8474.57
0.300000	−0.400000	0.840000	50.6909	8054.75
0.200000	−0.600000	0.640000	31.2494	7217.27
0.100000	−0.800000	0.360000	11.8930	5544.82
0.000000	−1.000000	0.000000	0.4834	0.00

The table is exactly symmetric about `a = 0.5`, and its first and last rows are identical. Two parties who agree on every item and two parties who disagree on every item produce the same `u`, the same frequency and the same reading. That is CC-R8 arriving through the measurement channel rather than through the algebra, and the unison result [Curzi, 2026] owns its consequences.

THE ACCURACY GUARD, binding wherever Route B or Route C is used. A `u` reading is published with each party's accuracy on the item set. Two parties answering at random have expected `a = 0.5`, hence `u = 1`, hence `f_rel = 69.6133 Hz` — the ceiling of the whole system, from guessing. Surplus measures distinctness, never correctness, and this series has no formalism converting one into the other.

THE OVERLAP GATE, not optional, and it runs before anything is computed. If two framings have disjoint support, `u = 1` exactly regardless of the intensities — verified across five constructed disjoint pairs with random intensities, `u = 1.000000` in every case. A degenerate reading is not a high reading; it is no reading. If the gate fires, report DEGENERATE, publish no `f_rel`, and return to the situation statement.

Protocol C4 — Estimating `eta`

Owner: the estimation practice [Curzi, 2026] at dyadic scope, the personal-spectrum material [Curzi, 2026] at personal scope.

`eta` is the fraction of a period's joint work still doing something at the start of the next period. It is countable in principle and the count is the protocol.

```
C4.1  Enumerate what the parties produced in a period: decisions, agreements,
      artefacts, practices. Write the list before counting it.
C4.2  Count how many survive intact into the following period without being
      re-made.
C4.3  eta_hat = survivors / total.
C4.4  Declare what got called a thing. The estimate is sensitive to it and the
      sensitivity is the reader's business.
```

A pair that settles six things a month and re-settles four next month is converting about a third. **The count is crude and far better than a guess, because the objection to it is visible and arguable in a way that an objection to a guess is not.**

Protocol C5 — Estimating `c` *Owner: the estimation practice [Curzi, 2026].*

`c` is what the arrangement costs to run before it produces anything: coordinating rather than deciding, re-explaining, the friction of getting to where work can happen. It is measured in the same units as the surplus — nats per period — and therefore requires a declared conversion, and the conversion is where most of the error lives.

```
C5.1  Identify the period budget and state it.
C5.2  Express eta*F and c as fractions of that same budget.
C5.3  Report c/eta directly. It is recoverable even when neither term is well
      determined on its own, and it is the only combination the deficit uses.
C5.4  Attach an error bar to c/eta, not a decimal point. Its elasticity is the
      largest in the reading.
```

The ratio `c/eta` carries the entire cost side of the reading, and a verdict resting on it deserves an interval. Where a structure prices its coordination per link, the aggregate is `C = kappa * m(m-1)/2`, and the commons of the group paper [Curzi, 2026] is the worked instance: `kappa = 0.055` across 15 links gives `C = 0.8250`.

Protocol C6 — Estimating `delta` *Owner: the trait paper [Curzi, 2026], Procedure TRAIT-2.*

`delta` is how much of what was accumulated is still there after a period of doing nothing. The estimation handle is the half-life relation.

```
t_half  =  ln(0.5) / ln(1 - delta)

delta = 0.05   ->   t_half = 13.5134 periods
delta = 0.09   ->   t_half =  7.3496 periods
```

```
C6.1  Identify a practice that lapsed, and the number of periods it lapsed for.
C6.2  Judge what fraction of it survived.
C6.3  Invert the half-life relation.
C6.4  DECLARE THE PERIOD. delta is a per-cycle rate and it is meaningless
      without its cycle. A delta quoted without a period is not a number.
```

A pair finding a lapsed practice half-forgotten after seven periods is running near `delta = 0.09`, since `ln(0.5)/ln(0.91) = 7.3496`.

Protocol C7 — Producing the reading: the deficit AND `delta`, together *Owner: the minimum-frequency paper [Curzi, 2026]. This is the protocol the framework exists to provide.*

```
C7.1  CEILING TEST FIRST.  Is c/eta <= ln N?
      If NO: the case is STRUCTURAL. No arrangement of the declared categories
      is sustainable. Report the ceiling breach and stop. Do not compute a
      deficit and do not offer levers; the levers do not include one that helps.

C7.2  COMPUTE F.        F = ln(1 + (N-1) u)
C7.3  COMPUTE f_rel.    f_rel = k (1 + (N-1) u)^2, raw and folded with j
C7.4  COMPUTE f_min.    f_min = k e^(2 c/eta)
C7.5  COMPUTE D.        D = 2.885390 (c/eta) - log2(f/k)
      CROSS-CHECK against 2.885390 (c/eta - F). The two must agree to at least
      twelve decimal places. If they do not, an input was rounded.

C7.6  COMPUTE Q_inf.    Q_inf = (eta F - c)/delta
      CROSS-CHECK against -D * eta / (2.885390 * delta).

C7.7  REPORT D AND delta, IN THE SAME SENTENCE.
C7.8  REPORT THE HOLD LINE u* = (e^(c/eta)-1)/(N-1) and the ratio u*/u_now.
C7.9  CLASSIFY.
          D <= 0                              SUSTAINABLE
          D > 0, c/eta <= ln N, u* reachable  TUNABLE AND REACHABLE
          D > 0, c/eta <= ln N, u* not        TUNABLE, UNREACHABLE
          c/eta > ln N                        STRUCTURAL
      Reachability is a declared judgment about the parties, not a property of
      the formalism, and it is declared rather than assumed.
```

Step C7.7 is the step this protocol exists for. the decay-blindness result CC-R4 [Curzi, 2026] proves `D` contains no `delta`, and the consequence is a reading that can be entirely correct and entirely uninformative about rate:

```
eta = 0.30, c = 0.62, u = 0.068734665264634

delta = 0.09   ->   D = +4.338420197   Q_inf = -5.011939546
delta = 0.05   ->   D = +4.338420197   Q_inf = -9.021491184
```

Halving the loss rate nearly doubles the negative destination and the deficit does not move by one digit in ten. **The deficit reads the sign and size of the gap; `delta` converts it to a level; a report with one and not the other is half a reading.** the audited schedule [Curzi, 2026] contains a row where exactly this happens — step 5 moves `Q_inf` from `-5.582082` to `-2.511937` while `D` sits at `+4.831945` in both rows — and it would have been invisible in a `D`-only report.

Protocol C8 — Controls, and what the surrogates cannot do *Owner: the falsification and preregistration work of Series B [Curzi, 2026i,ad]; carried here as CC-R4, CC-R13 and CC-R24.*

Any empirical test of any operationalisation in this series beats four baselines: persistence, direct-input, matched-autoregressive, and shuffled-input. **The shuffled-input control has three measured failure modes and every protocol that recommends the suite carries all three.**

```
(1) IT IS ANTI-CONSERVATIVE AS ORDINARILY RUN.  Against a nominal alpha of 0.05
    across 500 null trials, measured false-rejection rates were:
          iid permutation      0.108   (2.2x nominal)
```

```
          circular shift       0.120   (2.4x nominal)
          block permutation    0.060   (near nominal)
    The matched-autoregressive comparison ALONE false-passed 35.8% of the time.
    REMEDY: calibrate the surrogate by simulation on the study's own design
    before use. Not after.

(2) IT CANNOT SEE REPACKAGING.  A filtered repackaging of the input passed all
    four baselines including the shuffle (p = 0.0005) and was caught only by the
    structural ceiling test, its single-step jump being 38.5x the theoretical
    maximum.
    REMEDY: the structural tests. Only they detect repackaging.

(3) IT CANNOT SEE A SHARED SLOW DRIVER.  In a 4,000-trial simulation the shuffle
    worked in the clean case (48.17% detection) and failed against a shared slow
    driver, passing 98.35% of the time.
    REMEDY: design controls.

AND SEPARATELY: no surrogate detects FABRICATED TIMING. With a rater whose next
rating is nudged by the previous period's input -- no real relationship present
at all -- the full four-baseline suite passed 52.7% of the time. More surrogates
do not fix this.
```

The design-control remedy is the whole answer and it is not substitutable: blinding, channel separation, prospective sampling, and a yoked control. Adding surrogates to fix a design problem is a documented failure mode; the suite is a weak, mis-calibrated empirical audit and was never the load-bearing control.

The worksheets

WORKSHEET 1 — ONE FRAMING

```
================================================================================
RULE SET ID _______________  DATE _____________  N = 12   BAND _____________
SITUATION (one sentence, written first):
_____________________________________________________________________________
FRAMING NAME: ___________________   ELICITED BY: _________   CONSENT: [ ]

COMPONENT                          PRESENT   SUBSTITUTION TEST   INTENSITY
 1 arousal                          [ ]           [ ]            _________
 2 valence                          [ ]           [ ]            _________
 3 intentional object               [ ]           [ ]            _________
 4 situation                        [ ]           [ ]            _________
 5 cause                            [ ]           [ ]            _________
 6 projected duration               [ ]           [ ]            _________
 7 self as object                   [ ]           [ ]            _________
 8 another mind modelled            [ ]           [ ]            _________
 9 standard applied                 [ ]           [ ]            _________
10 placed in a story                [ ]           [ ]            _________
11 the other modelling you          [ ]           [ ]            _________
12 future action bound              [ ]           [ ]            _________

NESTING CHECK PASSED [ ]      n = _______      COMPONENT LIST ATTACHED [ ]
REFERENCE CLASS FOR MAGNITUDE ESTIMATION: _______   VALUE: _______
RAW L2 NORM (before normalisation): ___________________  <-- PUBLISH THIS
f = k n^2 = _____________ Hz     j = _______   folded = _____________ Hz
PITCH CLASS (odd part of n) = _______
================================================================================
```

WORKSHEET 2 — THE PAIR

```
================================================================================
PARTY 1 ___________  PARTY 2 ___________   BOTH CONSENTED SEPARATELY [ ]
SITUATION AGREED IN WRITING BY BOTH [ ]  SILENT WRITE-FIRST OBSERVED [ ]
ROUTES DECLARED BEFORE COMPUTATION: A [ ]  B [ ]  C [ ]

OVERLAP GATE:  shared state classes = _______   PASSES [ ]  DEGENERATE [ ]
RAW NORMS:  ||x|| = ___________   ||y|| = ___________  ratio = ___________

ROUTE A   cos theta = _________________   u = ____________________ (12 dp)
ROUTE B   m = _______  a = ___________   u = 4a(1-a) = _________________
          resolution 4/m = ___________
ROUTE C   rho = ___________ SIGN: ____  u = 1 - rho^2 = _________________
          accuracy party 1 _________   accuracy party 2 _________

F = ln(1 + 11u) = _____________ nats      N_eff = _____________
f_rel = k(1+11u)^2 = _____________ Hz     folded _____________ (j = ____)
================================================================================
```

WORKSHEET 3 — THE READING

```
================================================================================
eta = _________ (Protocol C4, list attached [ ])
c   = _________ (Protocol C5, period budget declared [ ])
delta = _______ (Protocol C6, PERIOD DECLARED: _________________)

STEP 1  CEILING TEST      c/eta = ___________   ln N = 2.484907
        c/eta <= ln N ?   YES [ ] -> continue     NO [ ] -> STRUCTURAL, STOP

STEP 2  f_min = k e^(2c/eta)                = _____________ Hz
STEP 3  D = 2.885390(c/eta) - log2(f/k)     = _____________ octaves
STEP 4  cross-check 2.885390(c/eta - F)     = _____________  AGREE TO 12 dp [ ]
STEP 5  Q_inf = (eta F - c)/delta           = _____________
STEP 6  cross-check -D*eta/(2.885390*delta) = _____________  AGREE [ ]
STEP 7  hold line u* = (e^(c/eta)-1)/11     = _____________
        u* / u_now                          = _____________
STEP 8  CLASSIFICATION: SUSTAINABLE [ ]  TUNABLE+REACHABLE [ ]
        TUNABLE/UNREACHABLE [ ]  STRUCTURAL [ ]
        reachability judgment DECLARED, not assumed [ ]

THE REPORT LINE (both numbers, one sentence):
"D = _________ octaves at delta = _________ per _____________."
================================================================================
```

WORKSHEET 4 — THE BOUNDS STATEMENT, WHICH IS PART OF EVERY REPORT

```
================================================================================
[ ] These are ESSENCE frequencies, not physical resonances. Nothing is
    claimed to resonate.
[ ] The twelve states are not diagnoses. This is not a clinical, medical or
    therapeutic assessment of anyone.
[ ] The frequency inherits exactly the defensibility of the count. The
    decomposition is published on this page.
[ ] The intensities were magnitude-estimated. The raw L2 norms are published.
```

```
[ ] The state frequency and the relationship frequency are INDEPENDENT
    readings. Neither was inferred from the other.
[ ] D is reported with delta. The deficit alone cannot see decay.
[ ] Consent was obtained separately from every party and declining was
    available at every step.
[ ] Every case in this report is CONSTRUCTED and labelled as such, OR is an
    observation with its provenance stated.
=========================================================================
```

---

7 The concordance register and its rejections

This series's audit trail. Every claim on which the framework runs a dual concordance appears here with four fields: the claim; the `(+1)` source and what it measured; the `(−1)` source and what it recorded; the status.

The two columns do different work and the difference is the method's whole content. `(+1)` carries evidence — a measurement, a derivation, a published result with a number attached. `(−1)` carries precedent — a tradition, a corpus, a practice, a record that somebody once thought this. Precedent is never promoted to evidence. A `(−1)` entry can suggest where to look, can supply a structure worth testing, and can constitute a striking convergence when an independent `(+1)` lands on the same object. It can never, by accumulation or by antiquity or by consensus, become the reason a claim is believed.

The sharpest demonstration of why is CC-R23. Under a selection process that publishes successes and discards failures, the likelihood ratio of testimonial evidence is exactly `1.000000`. **A likelihood ratio of 1 means the evidence does not move the posterior at all — not weakly, not by a little, not at all — because the selection guarantees you observe successes whether or not the method does anything.** The old corpus is a survivorship-filtered record, which is exactly that condition, and that is why the `(−1)` column is a separate column rather than a weaker part of the first.

Rejections are printed at the same length as acceptances and in the same tables. A register listing only what survived would be a sales document.

7.1 The laws and the join

#	Claim	(+1) — evidence, and what it measured	(-1) — precedent, and what it recorded	Status
D1.1	`f = k n^2` maps an essence index to a frequency	Series B [Curzi, 2026k]: the law reproduces all 118 published element values, worst relative deviation `9.0e−15`	the sonic-chemistry source table, which asserted the correspondence without a derivation	**ESTABLISHED as an exact reproduction of the source law; the source's prose formula was missing its parentheses (CC-R5)**
D1.2	`f_compound = k^2 G^4`	Series B [Curzi, 2026k]: reverse-engineered and matched to 37/44 published composites to <1e−12; the 7 outliers are data-entry errors, and the O2 row is provably impossible under the stated rule	the composite table itself	**ESTABLISHED with errata; NOT a generalisation of D1.1 (the no-agreement corollary CC-R4 [Curzi, 2026]: the rules agree at no index)**

#	Claim	(+1) — evidence, and what it measured	(-1) — precedent, and what it recorded	Status
D1.3	`F = ln(1 + (N−1)u)` is the unique surplus form	the surplus functional [Curzi, 2026n]: uniqueness from S1, S2, S4 given S3; every step after S3 forced	—	**PROVED CONDITIONALLY. S3 is a chosen design axiom**
D1.4	`Q_inf = (eta F − c)/delta`	the storage law [Curzi, 2026] Thm 3.5, with the exact solution and the contraction	Series A Node 10 FIXEDPOINT, which named the object without computing it	**PROVED. The precedent is a naming, not a derivation**
D1.5	`f_min = k e^(2c/eta)`	CCP-040: substitution verified at twelve parameter sets	—	**PROVED. An exact restatement of a sustainability condition in frequency units, and nothing more**
D1.6	`D = 2.885390(c/eta − F) = −2.885390(delta/eta)Q_inf`	CCP-041 R3; verified to `2.665e−15` across twelve parameter sets, worst residual `8.88e−16` on the carried case	—	**PROVED**
D1.7	`D` contains no `delta`	CCP-041 R4; worked: `delta` 0.09 to 0.05 moves `Q_inf` from `−5.0119` to `−9.0215` and leaves `D` at `+4.338420`	—	**PROVED, and it is a limitation of the framework's own central instrument**
D1.8	`dt df >= 1/(4 pi)` binds every rendering	Gabor; classical Fourier duality, measurable on an oscilloscope	—	**ESTABLISHED. No quantum content and none claimed**

7.2 The index and the psyche

#	Claim	(+1) — evidence, and what it measured	(-1) — precedent, and what it recorded	Status
D2.1	The component count is ratio-scaled and escapes CC-R1's demotion	CCP-041: a count has a true zero and equal units by construction; CC-R1 binds `D/P` specifically	—	**ESTABLISHED for `n`. NOT extended to the intensities**
D2.2	Arousal as an essence index	—	the arousal tradition in affect research	**REJECTED.** V2 fails — all twelve tie, 11.17 bits free; V3 fails; V7 fails, no tie-break
D2.3	Hierarchical need as an essence index	—	the need-hierarchy literature	**REJECTED.** V1 fails — needs are not states; V2 fails — five levels against twelve states

#	Claim	(+1) — evidence, and what it measured	(-1) — precedent, and what it recorded	Status
D2.4	Developmental stage as an essence index	—	stage theories of development	**ADMISSIBLE AS AN ORDINAL, NOT ADOPTED.** V1 fails — it indexes persons, not states; V4 fails — no true zero, so CC-R1 applies
D2.5	Valence magnitude as an essence index	—	dimensional affect models	**REJECTED.** V2, V3, V7 fail; V4 fails by collapsing opposites, which is CC-R8 in a different costume
D2.6	Magnitude estimation is required for the intensities	CCP-041 / the scale-status rule: a monotone relabelling moves `u` from 0.362115 to 0.525275 — 1065.0568 cents	Stevens' method, which produces ratio-scale subjective magnitudes up to a free multiplier and replicates (Series B)	**REQUIRED, not preferred. The ordinal fallback is not available**
D2.7	Consonance and redundancy are the same phenomenon	CCP-043: Tenney harmonic distance doubled in the essence lattice; the most consonant pairs are exactly the pitch-class collisions	classical consonance theory, which recorded the ratios without the identity	**PROVED at internal scope; nothing about distress follows, and CCP-043 states the firewall formally**
D2.8	One-hot framings are degenerate	CCP-041 / CCP-042: five constructed disjoint-support pairs return `u = 1.000000000000` regardless of intensities	—	**PROVED. The gate runs before anything is computed**

7.3 Relationship

#	Claim	(+1) — evidence, and what it measured	(-1) — precedent, and what it recorded	Status
D3.1	`u = 4a(1−a)` from an agreement rate	CCP-044 R2, proved from the norms of `+/−1` vectors; exact, no centring	inter-rater agreement practice, which recorded `a` without converting it	**PROVED. Resolution `4/m`, and `m` is reported**
D3.2	Perfect opposition reads as perfect agreement (CC-R8)	Series B [Curzi, 2026n]: `u = 1 − (x.y)^2` depends only on `|rho|`; CCP-044's agreement table is exactly symmetric about `a = 0.5`	"creative tension" as a design principle	**PROVED, and the precedent is refuted by it: `rho ~ −1` is a parallel pair with a cost overhead, paying `C` and returning zero surplus**
D3.3	Trust amplifies an existing sign (CC-R15)	CCP-044: `delta` 0.09 to 0.05 moves `Q_inf` from `−5.0119` to `−9.0215`; after the numerator is fixed, the same move gives `+6.0844` to `+10.9520`	the near-universal counsel to build trust first	**PROVED, and the precedent is corrected rather than rejected: trust is right, trust-first is wrong**
D3.4	An orthogonal pair already attains `ln N`	CCP-045 §4.7: orthogonal pair and orthogonal six both give `F = 2.484907`, `f = 69.6133 Hz` — identical, and the six pays 15 links against 1	the standing assumption that larger groups generate more	**PROVED in surplus terms. Groups buy breadth, redundancy and succession, which this formalism does not price**
D3.5	Ostrom's design principles map to the parameters	Series B [Curzi, 2026n]: P4 lowers `c_v` and does NOT raise `a`; P5 priced at `−7.7%` of `eta S`; P8 turns `O(N^2)` coordination into `O(N)`	Ostrom's commons scholarship	**ESTABLISHED where it applies and PRICED where it does not. P4's effect is not the one usually claimed for it**

#	Claim	(+1) — evidence, and what it measured	(-1) — precedent, and what it recorded	Status
D3.6	A commons can be formally tunable and practically unreachable	CCP-045: `c/eta = 2.426471 < ln 12`, so TUNABLE; but `u*` required is 0.938078 against `u_bar = 0.018614`, a factor of 50.4	—	**ESTABLISHED. Reachability is a declared judgment about the parties, not a property of the formalism**
D3.7	Drop-one improves on every member of a healthy team	CCP-045 §5.11b: all five rows improve, `Q_inf` rising by between 1.97 and 2.80	performance-ranking practice	**PROVED, and it refutes the practice: a favourable drop-one row carries no information about the member it names**

7.4 The tradition audited

#	Claim	(+1) — evidence, and what it measured	(-1) — precedent, and what it recorded	Status
C23.1	Staged process with per-stage checks	Shewhart 1931; Deming 1986; §23.8 prices staged control at 47.83% to 93.21% yield at `d = 0.90`	the Zosimos corpus; Newman & Principe 1998; Principe 2013	**ESTABLISHED**
C23.2	Cadence beats intensity	this work's concavity and ceiling results; Cepeda et al. 2006	fixed-hour practice; Durkheim 1912; Rappaport 1999	**ESTABLISHED, and folklore's REASON REJECTED. See D.5**
C23.3	Correspondence tables are relational graphs	graph theory; the gcd invariant tabulated for `N = 3..12`	the tables themselves; Yates 1966	**ESTABLISHED AS FORM ONLY. Two classical tables graded against V1-V7 and rejected as indices**
C23.4	The operator is part of the apparatus	Rosenthal & Fode 1963; Polanyi 1958; Collins 1985; §23.16 computes mirroring at 536.4074 cents	the pure-artisan requirement	**ESTABLISHED, MECHANISM REJECTED. Moral purity is not the mechanism**
C23.5	The colour sequence is a necessary order	—	precedent only	**REJECTED.** C2 fails: identification is circular

#	Claim	(+1) — evidence, and what it measured	(-1) — precedent, and what it recorded	Status
C23.6	“As above, so below” as a transfer licence	CC-R16’s three computed disanalogies: mean-field decay error 92% by `t = 12`; two nine-member groups with identical `u = 0.5999` losing 77% against 21% to a three-person departure; a 90.1x period-normalisation trap	the maxim	**REJECTED AS STATED. The correspondence is of dynamics and geometry, not of agency**
C23.1b	This tradition’s canonical stage ORDER	CC-R15; CCP-046 the schedule finding R2 [Curzi, 2026]	Ripley’s *Compound of Alchymy* (1471), the Twelve Gates	**REJECTED (order only; staging stands)**
C24.1	Material transmutation by chemical means	the energy-register argument: `2e6` on bonds, `8.24e7` on the furnace-to-barrier ratio; transmutation is real by nuclear means at 15.52 GW-station-years per gram under the most favourable constructed assumptions	the transmutation corpus	**REJECTED.** C1-C3 fail; AUDIT-1 steps 2, 4, 5, 6, 7, 8 fail
C24.2	The corpus establishes a success rate	CC-R23: `LR = 1.000000`, invariant to corpus size; with a denominator the ratio goes to about `1e304`	the published successes	**REJECTED.** AUDIT-1 step 2
C24.3	Secrecy protects a valuable result	error persistence under secrecy: `r = 0` gives infinite expected detection time; Deckname search cost `3^12 = 531,441` readings, 19.02 bits, 1,455 years at a trial a day	the secrecy tradition	**REJECTED AS AN EPISTEMIC PRACTICE.** AUDIT-1 step 8

#	Claim	(+1) — evidence, and what it measured	(-1) — precedent, and what it recorded	Status
C24.4	Shared values or correspondences are evidence	R12 recomputed: 234,454 words to 305 distinct outputs, 99.995% collision, 10.47 bits destroyed against a 9.5863-bit uniform bound — worse than its own worst case; in an 8,000-word corpus `P(a value-twin) = 0.999666`	the gematria corpus	**REJECTED; MNEMONIC READING RETAINED.** A retrieval index need not be injective, only reproducible
C24.5	The canonical gate order must not be varied	CCP-046 the schedule finding R2 [Curzi, 2026] against CC-R15	Ripley 1471	**REJECTED AS AN ORDERING**

7.4b The audit of the twelve gates, verdict by verdict Seven gates name operations on this series's declared parameters, two are split, and three fail — with a fourth entry, the planetary-metal ordering, failing as an index. Each rejection states what would reverse it, because a rejection that cannot say what would reverse it is an opinion.

The `(-1)` column for the whole set is Ripley's *Compound of Alchymy* (1471). The seven that map do so on named parameters — calcination on `n` downward, solution as the inverse of Rule 3, separation on `N`, putrefaction on `delta`, congelation on `G`, cibation on `c`, fermentation on `G` and `u`. Conjunction maps but SPLITS, its two readings returning incommensurable units, and the schedule halts there. The failures:

entry	(+1) — evidence, and what it measured	Verdict, and what would reverse it
8 sublimation	folding by `2^j` preserves every interval and every pitch class exactly; nothing in the parameter set moves, and `k` cancels from `D`	**RENDERING — rejected as an operation on a state.** Reverses if `j` entered a verdict. It does not
10 exaltation	"nobler" is a value ordering, not a coordinate (CCP-041); the rescues collapse to gates 1 or 9, or hit the ceiling `F <= ln 12`, `f <= 69.6133055503`	**REJECTED.** Reverses if a count-valued definition of nobility is supplied. None is, in the source or since
11 multiplication	contraction: `|Q_t − Q_inf| = (1−delta)^t |Q_0 − Q_inf|`, 99% closure at 48.83 passes; the destination is fixed by four parameters repetition does not touch	**SPLIT: the iterative reading MAPS, the unbounded reading is REJECTED.** Reverses if a pass altered `eta`, `F`, `c` or `delta` — then it is a different gate under a borrowed name

entry	(+1) — evidence, and what it measured	Verdict, and what would reverse it
12 projection	no base/noble coordinate; Rule 3 is blind to quantity, so the mass is not in the value; no channel (belief postulate); CONSTRAINT L67 forbids reaching the non-latent; CC-R23 fixes the testimonial LR at 1.000000	**REJECTED** on all four conditions. Reverses if a channel with measured gain is exhibited. None is
13 the planetary-metal ordering as an essence index	completeness, non-arbitrariness, count semantics and stability all fail; one disputed swap moves a published class by 568.7174 cents	**REJECTED.** Reverses if the ordering were fixed by a count. It is fixed by worth

7.5 Findings the framework records against its own sources Four items belong in their own section because in each of them the framework audits a source on the source's own material and returns a verdict the source would not have given.

The schedule finding, R2 of the operation-audit paper [Curzi, 2026], — Ripley's ordering contradicts CC-R15. Run as a schedule on the carried dyad, the Twelve Gates of Ripley's *Compound of Alchymy* (1471) place the only numerator-acting gate ninth of twelve. Fermentation, at step 9, moves `D` by `−5.305895` octaves, which is 88.93% of the entire move; the eight gates before it move it by `−0.660631` between them, which is 11.07%. Step 5, putrefaction, acts on `delta` — and it runs four steps before the numerator is repaired. **In the tradition's order, raising `delta` improves the destination; in CC-R15's order, the identical move damages it; and the only thing that differs is the sign of the numerator when the move is made.** This is the strongest single item in the audit papers of this series precisely because it is the tradition being audited on its own published material, and it is recorded here as `(+1)` evidence against a `(−1)` precedent — never the reverse.

The same schedule shows the delta-blindness of D1.7 as a visible row: step 5 moves `Q_inf` from `−5.582082` to `−2.511937` while `D` sits at `+4.831945` before and after.

CC-R11 — cadence wins, and not for the reason the tradition gives. Computed rather than assumed: the integrated effect of one large impulse and the same total input delivered periodically is identical to every digit — 433.3333 in all three regimes tested. That is exact LTI invariance, and a naive "cadence delivers more" claim is simply false for a linear time-invariant system. Cadence buys three other things instead, and Section 8.5 states them. The honest counter-case is kept in print: above a threshold no realistic cadence reaches, the impulse wins, and a once-a-decade rite requires a 5.63-year trace half-life to work.

the refused move [Curzi, 2026]. A proposed re-indexing of the carried Subject's `mirror` framing from state 7 to state 11 was run through TR-1 and halted at step 3: three of the four added components were observed in the repertoire, and the fourth — the other's model of the subject — was not. **The move was refused under CONSTRAINT T21, and steps 4 through 10 were not run, so the arithmetic that would have been produced is not published, because publishing it would let a refused move circulate as a priced one.** What is published is the refusal, the named missing component, and what would reverse the verdict: a single observation of the component, in any state, on any occasion. A refusal that names its missing component is a scientific result; one that says the move is impossible is an opinion.

The closed lineage, read in the framework's own register. Constructed, and not an observation of any lineage historical or living. A closed lineage in which the student's whole training consists of reproducing the master's framings exactly is `rho = 1`, hence `u = 0`, hence

— by CC-R8, which makes it worse rather than better — the same reading a perfectly opposed pair would receive.

```
arrangement                          u        F        f (Hz)     c/eta      D         Q_inf
-------------------------------------------------------------------------------------------
closed lineage, exact reproduction  0.000  0.000000   0.483426   2.250000  +6.492128  -7.500000
closed lineage, slight drift        0.050  0.438255   1.161430   2.250000  +5.227591  -6.039150
open school of six, u = 0.30        0.300  1.458615   8.938542   2.090909  +1.824415  -3.091215
open school of six, u = 0.50        0.500  1.871802  20.424737   2.090909  +0.632209  -1.071189
open school of six, u = 0.70        0.700  2.163323  36.590494   2.090909  -0.208942  +0.354024
-------------------------------------------------------------------------------------------
every row's c/eta is inside the ceiling ln 12 = 2.484907, so these are real
failures and not ceiling breaches
```

The swing from the closed lineage to the open school at `u = 0.70` is 7490.44 cents and 6.701070 octaves of deficit. Openness is three levers at once: it raises `u` because independent practitioners frame differently, raises `eta` because a published result is stored order and an unpublished one is a memory, and lowers per-link `c` because a record replaces re-explanation. **And the independence rule [Curzi, 2026] binds the reading: every `f` in that table is a relationship frequency and says nothing whatever about the state frequencies of the people in it.**

7.6 Rejections inherited from Series B, restated because they still bind

#	Claim	(-1) — precedent	Failing criterion	Status
B2-a	Astrology's orb as a conjugate-spread structure	the orb tradition	C2 FAILED in print, with Carlson, *Nature* 318 (1985) cited in the same passage as the definitive failed validity test	**REJECTED as a validity claim; the structural analogy contributes no number**
B2-b	64 codons to 64 hexagrams	the correspondence tradition	refused at CCP-031 §67.11	**REFUSED**
B2-c	The Heisenberg analogy as a source of numbers	quantum-mechanical framing of psychological uncertainty	every quantitative result in the conjugate-uncertainty work of Series B comes from an information-budget bound, the Gabor trade, or a Lyapunov horizon law — none from Heisenberg	**RETAINED AS STRUCTURE, REJECTED AS ARITHMETIC**
B2-d	Momentum descent as a formulary method	optimisation practice	failed to escape the worked landscape at every `beta` tested; the barrier is a wide plateau, not a spike	**EXPLICITLY EXCLUDED from the formulary with its failure stated. The exclusion is a feature**
B2-e	`r = 1` as an unqualified target	calibration practice	CC-R9: headroom at the `D = P` fixed point is exactly zero by cancellation	**REJECTED AS A SIMPLE GOAL**
B2-f	Aggregating `r` across sites	reporting practice	CC-R10: aggregate `D/P` overstated bankable input by 4.26x	**PROHIBITED. Distortion does not average**

---

8 The four-drawer evidential classification

This is the framework's warrant. Everything else in the series is an argument; this appendix is the accounting that says which arguments were paid for and in what currency. A reader who trusts nothing else here should be able to trust this.

The accounting has four drawers. **PROVED** means derived from stated premises, with the derivation in print and the arithmetic verified. **POSTULATED** means asserted as a design choice or a working assumption, badged in the sentence that states it, with the cost of the alternative computed where that was possible. **BORROWED** means taken from an established result elsewhere, cited, and used without re-derivation. **UNTESTED** means specified but not run — a claim with a falsification criterion attached and no data against it.

An item can sit in more than one drawer at different levels of description, and where it does, both placements are recorded. A theorem proved from a postulate is proved as a theorem and postulated as a system, and a reader is entitled to both facts.

8.1 The first drawer — PROVED

the essence bridge	f = k e^(2F)	Series B CCP-049
the surplus form	F = ln(1 + (N-1)u), unique given S3	CCP-020
the storage solution	Q_t exact, and Q_inf = (eta F - c)/delta	CCP-041 the storage law
the keystone	f_min = k e^(2c/eta)	CCP-040, twelve parameter sets
the octave form	octaves = 2.885390 (c/eta)	CCP-041 Thm 4.3
the deficit identity	D = 2.885390(c/eta - F) = -2.885390(delta/eta) Q_inf	CCP-041 Thm 4.5, 2.665e-15
delta-blindness	D contains no delta	CCP-041 Thm 4.9
the ceiling	c/eta <= ln N, f_min <= k N^2	CCP-041 Thm 4.6 / CCP-041 Cor 3.4
the minimum category	N >= e^(c/eta)	CCP-041 Cor 4.7
the hold line	u* = (e^(c/eta)-1)/(N-1)	CCP-041 Cor 4.8
the odd-part theorem	ceil(N/2) classes, verified N = 1..25	CCP-042
the agreement identity	u = 4a(1-a), cos theta = 2a - 1	CCP-044 Thm 13.1
R8	opposition reads as agreement	Series B CCP-020
R15	trust amplifies an existing sign	Series B CCP-020
the compound rule	f = k^2 G^4, and Thm 20.1	Series B CCP-049, CCP-046
Cor 20.2	the two rules agree at no index	n = 1.438252
Thm 22.2	refinement cannot move Q_inf	CCP-047
the error law	|Q_t - Q_inf| = (1-delta)^t |Q_0 - Q_inf|	CCP-047
the Gabor bound	dt df >= 1/(4 pi)	classical Fourier
the four design laws	N <= 25 at j = 6; cliffs at 17 and 41; rate factor exactly 8; rectangular gate disqualified	
Finding 21.1	state and relationship frequency independent	CCP-046
R23	LR = 1.000000 for testimonial evidence	Series B CCP-036

Two of these are limitations rather than achievements.

The deficit is blind to decay, and it is the framework's central instrument. the decay-blindness result CC-R4 [Curzi, 2026] says `D` contains no `delta`. Worked on the carried dyad, moving `delta` from 0.09 to 0.05 takes `Q_inf` from `-5.011940` to `-9.021491` — nearly doubling the negative destination — while `D` stays at `+4.338420` to ten decimal places. The instrument this series exists to provide cannot see one of the four parameters that decide a relation's fate. The remedy is procedural and is stated everywhere the instrument appears: report `D` and `delta` together. A reading that gives only `D` is incomplete, and the incompleteness is not a rounding issue but a structural one.

A state's frequency and a relationship's frequency are independent, and neither may be inferred from the other. the independence rule [Curzi, 2026] records the verification: four constructed pairs at lattice positions spanning thousands of cents — indices `[1,2]` at 2400.00 cents apart, `[3,6]` at 2400.00, `[5,10]` at 2400.00, `[7,11]` at 1564.98 — all return `u = 1.000000000000` and `f_rel = 69.6133055503 Hz`, to twelve decimal places. The reason is structural: `u` depends on the angle between state vectors, not on which indices are occupied. Re-indexing moves every state frequency and can leave the relationship reading untouched. A complete report therefore gives both numbers, and any paper tempted to derive one from the other is making an error the arithmetic forbids.

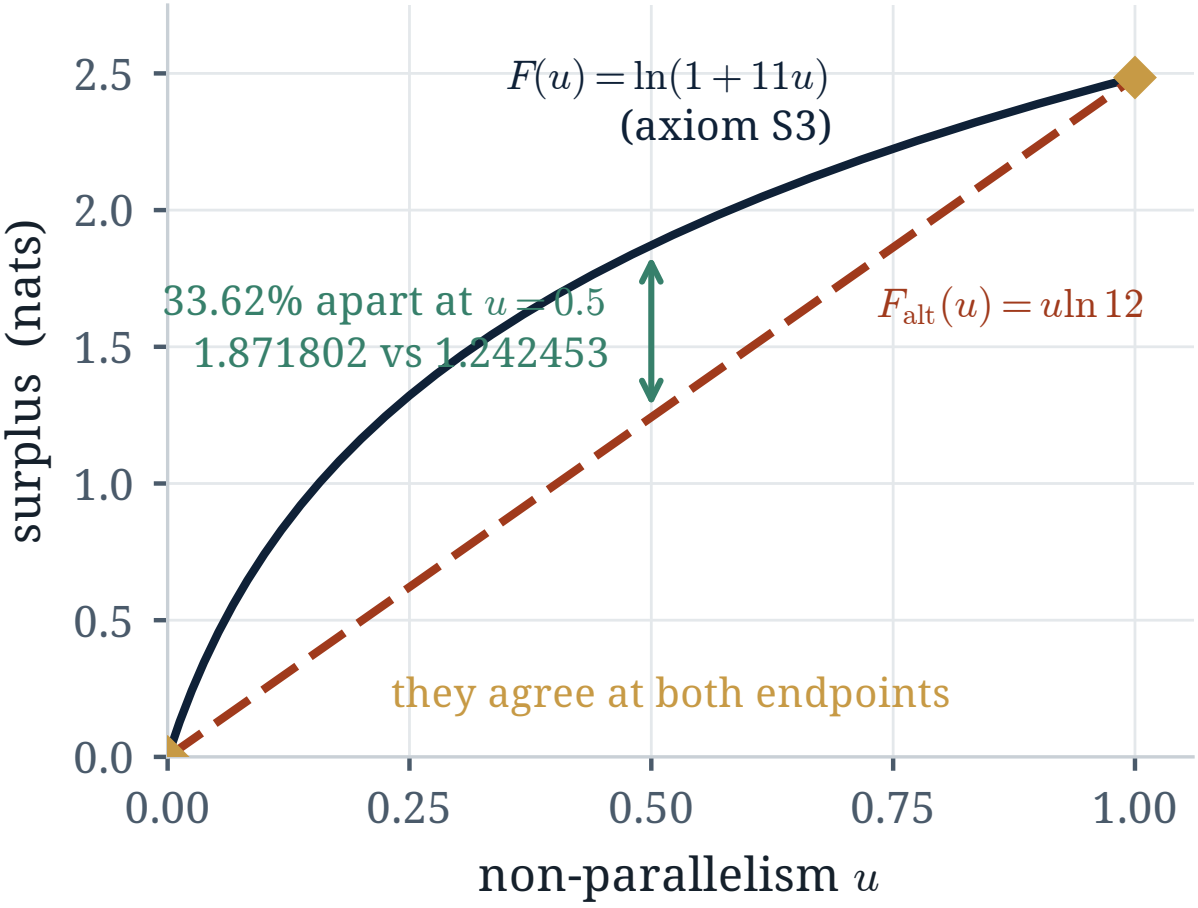


Figure 50.4: **The framework's central design axiom is worth a third of the surplus in the middle of the range.** $F(u) = \ln(1 + 11u)$, the form the affine effective-count axiom forces, against $F_{\text{alt}}(u) = u \ln 12$, which satisfies every other axiom. The two agree at both endpoints and disagree by 33.62 per cent at $u = 0.5$, 1.871802 against 1.242453. Computed at $N = 12$; script `figs/src/CCP-050_fig03.py`.

8.2 The second drawer — POSTULATED S3, the affine effective count, is a chosen design axiom and this entire book leans on it. The surplus form is uniquely characterised by four axioms: S1 geometric dependence, S2 zero at zero, S3 that `g(u) = e^{F(u)}` is affine in `u`, and S4 normalisation `F(1) = ln N`. Every step after S3 is forced. S3 itself is forced by nothing. It asserts that the *effective count*, not the surplus, is the quantity that behaves linearly in the geometry, and that is a modelling choice about what the natural unit of the system is.

The cost of a different choice is computed rather than waved at.

`F_alt(u) = u ln N` satisfies S1, S2 and S4, violates only S3, and agrees at both endpoints — yet at `u = 0.5` the two disagree by 33.62%, `1.871802` against `1.242453`. S3 selects one curve from a family that agree at the ends and diverge by a third in the middle, and the middle is where every real case lives. Nothing here is safe from it: `N_eff` is defined by S3, the bridge identifies the essence index with `N_eff`, `f = k N_eff^2` inherits it, and `f_min` is a statement about `F = ln N_eff`. Every paper using the keystone is using S3.

The storage law's linear form is a modelling choice, not a discovery. `Q_{t+1} = (1-delta)Q_t + eta S_t - C_t` is linear in `Q`, with constant `eta` and constant `delta`. That is what makes the fixed point exist, makes the convergence geometric, and makes `Q_inf` computable in closed form. Real systems can have superlinear cost, `Q`-dependent conversion, or exogenous accretion, and where they do, the model does not apply and the domain limits [Curzi, 2026] computes where it stops. The linear form is retained because it is the simplest form that supports the join, and

its simplicity is the reason the join is checkable — not evidence that it is right.

```
OTHER BADGED POSTULATES, WITH THEIR OWNERS

Postulate R          eta = eta_max * rho, reciprocity as a multiplier   Series B CCP-020
Postulate SO1        Rasch instrumentation makes Delta ratio-scaled     Series B CCP-026
the belief postulate belief acts only through a coupling with real gain Series B
the trait cut        a declared convention, not a natural kind          CCP-042 §8.6
the twelve-state ladder itself: a constructed decomposition, defended
                     against four rivals and adopted, not discovered    CCP-041
the ladder's ordinal-to-frequency mapping: defensible, not unique       CCP-041 §5.4
```

Belief acts only through a coupling with real gain, and where no channel exists nothing happens. **The framework must never imply that wishing moves the fixed point.** Changing a destiny means changing `(eta, F, c, delta)` and nothing else — a definition, not a promise.

8.3 The third drawer — BORROWED

```
Gabor's limit and Fourier duality
Stevens' magnitude estimation, and its replication record
Rasch measurement, as the route out of R1's demotion (123 items per term)
Shewhart 1931 and Deming 1986, on staged process control
Cepeda et al. 2006, on distributed practice
Rosenthal & Fode 1963; Polanyi 1958; Collins 1985, on the operator
Ostrom's design principles, mapped to parameters and priced
Brooks' quadratic coordination cost
Tenney's harmonic distance
Granovetter's threshold model, with mu* = 0.408522 and a 99.0% / 14.5% knife edge
Carlson, Nature 318 (1985), as the definitive failed astrological validity test
Church, Gao & Kosuri, Science 337:1628 (2012), on data in nucleotide sequence
Goldman et al., Nature 494:77 (2013); Erlich & Zielinski, Science 355:950 (2017)
Zetsche et al., Cell 163:759 (2015); Fonfara et al., Nature 532:517 (2016)
Qi et al., Cell 152:1173 (2013); Gilbert et al., Cell 159:647 (2014)
Patapoutian, Nobel 2021, on mechanically gated ion channels (PIEZO1/PIEZO2)
Tajik et al., Nature Materials 15:1287 (2016), on force reaching the nucleus
Ibsen et al., Nature Communications 6:8264 (2015); Duque et al. (2022)
Takahashi & Yamanaka, Cell 126:663 (2006), Nobel 2012, on latent reprogramming
Ripley, Compound of Alchymy (1471) -- borrowed as (-1) PRECEDENT only
the Zosimos corpus; Newman & Principe 1998; Principe 2013 -- likewise (-1)
```

Every borrowing is used for what its source established and nothing further. The biological citations establish sequence operations and mechanical operations; they do not establish the link between them.

8.4 The fourth drawer — UNTESTED

```
the operationalisation of the twelve-state ladder: FVP-6 SPECIFIED, NOT RUN
    design: preregister rule set, ladder, elicitation script and analysis;
    fit on training episodes; FREEZE; report held-out performance
    kill condition: if counts under the frozen protocol fail to beat the
    persistence baseline on held-out prediction, the operationalisation is
    dead and the ladder is withdrawn

the acoustic -> c-factor causal claim: NOT CONFIRMED (CC-R7). Any paper
    invoking acoustic c-factor engineering carries that status. The nearest
```

```
        badged result is that boundary geometry alone selects modes at no
        ongoing cost -- exact for room acoustics, a design heuristic only for
        occupancy [CCP-033].

    the genomic audio protocol of CCP-048: a RESEARCH PROCEDURE, labelled as one.
        Nothing in it has been run on anything.

    the decisive experiment of CCP-048: designed, preregistered in outline, not run.

    every parameter estimate in this series: eta, c and delta have no instrument
        comparable to u's, and Protocols C4-C6 are estimation practice badged as
        practice.

    the research programme's own odds: 21 preregistered kill conditions across five
        stages, four terminal; 35.8% completion probability; 55.3% chance of dying
        at or before cross-subject recovery; 8,554.8 researcher-hours, 5.35
        FTE-years, 71 months.
```

The last item is a result, not a caveat. A programme that computes and publishes its own probability of failure is more credible than one claiming certainty, and the coherence with CC-R1 is no coincidence: the gate deciding the hard-science question is also the stage most likely to kill the programme.

8.5 The specific disclaimers, each stated in full

Essence frequencies are essence frequencies, not physical resonances. No person, tissue, cell or molecule is claimed to resonate at any value in this series. The frequencies represent a count and an angle; the hertz are a unit chosen so the representation can be compared with a lattice and, where a rendering is called for, heard. Nothing in the formalism asserts a physical oscillation and nothing in it would be evidence of one.

The scale asymmetry, stated as the scale-status rule fixes it. Two quantities, different scale status, and no blanket verdict for "the ladder" or "the method" is permitted. `n` is a count: true zero — zero components is no state — and equal units, hence ratio-scaled by construction, and `f = k n^2` is ratio-scaled with it. Neither inherits CC-R1's demotion; they sidestep it. `w`, the intensities, are ordinal by default, and `u`, `F`, `f_rel` and the deficit inherit `w`'s status. So the state frequency is ratio-scaled and the relationship frequency is not, until the intensities are. CC-R1's demotion binds `D/P` specifically — pending Rasch instrumentation at 123 items per term — and binds neither this index nor magnitude-estimated intensities. Two phrasings are forbidden because both are false: "the ladder is ordinal per CC-R1", and "the method is ratio-scaled" without qualification.

And the ordinal reading is not merely weaker; it is unusable. An ordinal-preserving relabelling of the intensities moves `u` from 0.362115 to 0.525275, which is 0.887547 octaves — 1065.0568 cents, nearly an octave — from a change an ordinal scale says is no change at all. Magnitude estimation is therefore required rather than preferred, and the raw L2 norms are published so the free multiplier stays recoverable.

CC-R23 stands, and its number is exact. The survivorship likelihood ratio for testimonial evidence is `1.000000` — not approximately, exactly. A thousand testimonials carry precisely the evidential weight of zero, because the selection guarantees you observe successes whether or not the method does anything. The companion result: a prediction specifying only a direction has `LR = 2.0000` against `11.3577` for one specifying a magnitude, so a belief predicting only a direction is not yet testable.

CC-R11, stated exactly, because the easy version of it is false. The integrated effect of one large impulse and the same total input delivered periodically is identical to every digit — 433.3333 in all three regimes tested. That is LTI invariance and it is exact. Cadence wins by strict concavity (the concavity result of Series B, which needs no ceiling at all) and by the

surplus ceiling clipping the impulse — 90.44% of the nominal budget lost at `N = 12` — and not by delivering more total input. It does not deliver more.

CC-R15, stated exactly. Trust acts by reducing `delta`, and `Q_inf = (eta F − c)/delta`. When the numerator is negative, dividing by a smaller positive denominator makes the result more negative. Trust does not add order; it amplifies whatever sign the system already has. Fix the numerator first.

CC-R8, and the identity that carries it into measurement. Because `u = 1 − (x.y)^2` depends only on `|rho|`, perfect opposition yields `u = 0` exactly as perfect agreement does. The agreement identity `u = 4a(1−a)` exhibits it directly: the channel is symmetric about `a = 0.5` and returns `u = 0` at both ends. The contrarian who inverts you carries the same information as the sycophant who echoes you — none — and manufactured opposition pays `C` and returns zero surplus.

the group result [Curzi, 2026], stated before its qualifications. A single orthogonal pair attains `u_bar = 1`, `F = ln 12 = 2.484907`, `f = 69.6133 Hz`; six mutually orthogonal members attain exactly the same and pay 15 links instead of 1. In surplus terms a group buys nothing a well-chosen pair could not. The qualifications are real and are printed with it: orthogonality is an idealisation and achievable `u_bar` falls with roster size; the formalism has no term for redundancy against loss, for succession, or for the fact that a pair ends when one of its members does; `eta` and `kappa` are not roster-independent in practice; and the ceiling is on `F`, not on output. Groups buy breadth, not depth, and the formalism prices only the depth.

the constructed commons and the re-read window [Curzi, 2026]. The commons is formally tunable — `c/eta = 2.426471` sits inside `ln 12 = 2.484907` by 0.058436 nats, or 0.1686 octaves — and practically unreachable, because the hold line demands `u* = 0.938078` against a present `u_bar` of 0.018614, a factor of 50.4. No single lever closes a 6.4638-octave gap: `u` to 0.95 leaves `−0.0332`, and 0.95 is near-orthogonality; `eta` to 0.90 leaves `+2.1075`; `kappa` cut to 0.020 leaves `+2.0085`. Some systems cannot be tuned and must be restructured, and a book showing only fixable examples would be lying by selection. Just as awkward, the drop-one analysis improves the reading on *every* member of the healthy team — `Q_inf` rising by between 1.97 and 2.80 in all five rows — so a favourable drop-one row carries no information about the member it names. And the re-read window [Curzi, 2026], bounded below by convergence and above by drift, can be computably empty: the tuned family's window is `[13.51, 16.23]` at 1% drift and empty at 2%, 3% and 5%. A margin thin enough to be unverifiable is a margin too thin.

The open problem, carried without softening. There is no established mechanism by which a cell demodulates sequence-specific information from an acoustic signal, and the framework's own addressing constraint bounds what any such mechanism could do even if one existed: no addressing operation, at any signal strength or in any modality, can cause a cell to express content its genome and accessible regulatory repertoire do not already contain [Curzi, 2026y]. The nearest badged result on the acoustic side is narrower still: boundary geometry alone selects modes at no ongoing cost, which is exact for room acoustics — the same Helmholtz problem — and a design heuristic only for occupancy [Curzi, 2026aa]. The channel needs three things: D1 symbol separation, satisfiable in principle and its engineering exists; D2 order memory, for which there is no identified mechanism — this is the heart of the gap, since a permutation-invariant response is not a weak version of a sequence-specific one but a different quantity; and D3 mapping to a locus, for which there is likewise none. Eleven upstream rows are established and one is open. Reporting only the eleven is advocacy; reporting only the one is a different dishonesty, implying the whole edifice is speculation when most of it is textbook. the open-problem paper [Curzi, 2026] is why the rest of the framework can be trusted.

The remaining absolute constraints. No clinical, medical, diagnostic or therapeutic claim is made anywhere in this series; the twelve states are not diagnoses and the readings are not assessments. No claim is made that any DNA is edited. No capability is conferred on anyone by anything here. No subliminals, ever — nothing masked, nothing below threshold, nothing concealed; open-label placebo works, so disclosure costs nothing that matters. No intervention may remove a person's ability to decline, and consent is structural rather than

advisory. All worked cases in all the papers of this series are constructed and are labelled as such, and no invented case is presented anywhere as an observation.

8.6 What would change these verdicts An honesty appendix that could not be moved by evidence would be a creed. Each drawer has an exit. PROVED is exited by an error in a derivation or a failed verification, and every verification here states its residual so a reader knows what tolerance was claimed. POSTULATED is exited by exhibiting a system in which the postulate fails — for S3, one whose effective count is not affine in the geometry, in which case the whole functional form changes and every number in this series moves with it. BORROWED is exited by retraction or supersession. UNTESTED is exited by running the test, and every item in it carries a kill condition stated in advance rather than a hope.

The open problem is exited by one experiment, described in the open-problem paper [Curzi, 2026], with both outcomes written before it is run. Until then the position is the one stated throughout: the architecture upstream is real, the architecture downstream is real, and the link between them does not exist yet.

---

9 Correspondence to the two predecessor series

Three books, one instrument, built in three stages. Series A taught the analyst to read the comfort gradient. Series B taught the engineer to build the circumstance and laid down both halves of one equation without multiplying them. This series performs the multiplication, and this appendix says item by item which part of the earlier work each part of this one extends, what the extension adds, and what it costs.

Two conventions govern the tables. Where the nodes of Series A are named the name is used; where they are not, the node is referred to by index only, because inventing names would violate Series A's own The naming rule of Series A [Curzi, 2026]. And where an extension costs something the cost is printed in the same row as the gain.

9.1 Series A's twelve nodes The nodes sit on the same twelve-index ladder this series uses, so each has a frequency and a pitch class under Rule 1. Those values are the essence-law paper of Series B [Curzi, 2026k]'s and agree digit for digit with B.1.

node	f = k n^2 (Hz)	class	what Series A taught	extended by	what the extension adds	what it costs
01	0.483426	1	the ground: undifferentiated activation	CCP-041, CCP-043	`n = 1` is the count's first rung, and `u = 0` gives the ground essence `f = k` — the same pitch from two directions	the reading at `f = k` is *no relationship*, not a low one, and readers are told so every time
02 GRADIENT	1.933703	1	read the desire line; move along the gradient of comfort	CCP-043, CCP-045	`grad h = −grad C` (CC-R19) makes a habit survey a headroom measurement; CCP-045 turns the gradient into a numbered intervention procedure	the consent axiom: no intervention may remove the ability to decline, which rules out the cheapest interventions
03 HORIZON	4.350832	3	the many windows of the one tower — several framings of one situation	CCP-042, CCP-043	the framings become a chord with computable intervals, and its width is the internal surplus `F_int`	the one-hot trap: framings without overlapping support score `u = 1` for free and mean nothing by it

node	f = k n^2 (Hz)	class	what Series A taught	extended by	what the extension adds	what it costs
04 WAVE	7.734812	1	the future as a spread rather than a point	the genomic-protocol paper [Curzi, 2026]	the spread is priced by Gabor: `dt df >= 1/(4 pi)`, an exact classical bound	the quantum analogy contributes no number; a badged analogy may organise, it may not compute
05	12.085643	5	set the idea of greater comfort where the self will pass; lower the cost of the destination	CCP-041 lever L3, CCP-045	cutting `c` is now one of exactly four levers, with its partial `dD/dc = +2.885390/eta`	CC-R20: lowering cost only at the origin shortens the wait from 32.0 to 9.9 weeks but leaves a 37% smaller permanent fund
06	17.403326	3	the minimal reversible probe; the veil that measures the measurer	CCP-045, CCP-047	reversibility is classified formally, and the analyst's own state is computed into `u` — mirroring costs 536.4074 cents	an observation taken by a reactive instrument is not free, and the price is now a number rather than a warning
07 COLLAPSE	23.687861	7	the moment the many become one	CCP-043, CCP-044	consonance and redundancy are proved identical, so collapse into agreement is collapse into redundancy	an entirely consonant self is an entirely redundant one, and nothing about distress follows — CCP-043 states the firewall formally
08 ECHO	30.939247	1	the world's answer, read over what was sent	CCP-044	the return difference becomes error decorrelation, `u = 1 − rho^2` — the only route that sees being wrong in different ways	it needs ground truth, unavailable for most of what two people actually disagree about
09 RECURSE	39.157484	9	the loop that returns changed	CCP-047	the transient `(1−delta)^t Q_0` is exactly the error law, and the three constants of arrival depend on `delta` alone	refinement cannot move the destination (the refinement theorem [Curzi, 2026]): the seventh distillation reaches nowhere the first was not already headed
10 FIXEDPOINT	48.342573	5	where the loop arrives	**the destiny paper [Curzi, 2026]**, and CCP-041	**Node 10 FIXEDPOINT = the storage fixed point [Curzi, 2026] = this series's the destiny paper [Curzi, 2026].** `Q_inf = (eta F − c)/delta` is destiny defined precisely; and by Thm 4.5 it is the octave deficit in different units	the definition is not a promise: changing a destiny means changing `(eta, F, c, delta)` and nothing else, and belief acts only through a coupling with real gain
11 SEAL	58.494514	11	when not to intervene	CCP-045 §6.13	a computable stop rule: ceiling test, empty re-read window, unreachable hold line each halt the procedure	the halt is a finding, not a failure; CCP-046's schedule halts at gate 4 for the same class of reason
12	69.613306	3	the committed state, the top of the ladder	CCP-041, CCP-045	`k N^2 = 69.6133 Hz` is now the system ceiling, reachable by an orthogonal pair and by nothing larger	above `c/eta = ln 12 = 2.484907` no arrangement of the twelve is sustainable, and the only remedies leave the ladder

Node 10's row is the load-bearing one and is stated once more in isolation, because three books converge on it: **the tenth node of Series A, the storage law [Curzi, 2026]'s Theorem 3.5, and this series's the destiny paper [Curzi, 2026] are the same object under three descriptions.** Series A named it and could not compute it; the storage law [Curzi, 2026] computed it and did not connect it to pitch; the deficit identity R3 [Curzi, 2026] connects it, and the destiny paper [Curzi, 2026] states what the connection means — a pitch reading and a destiny reading are the same reading.

One structural observation about the ordering, offered as arithmetic and not as evidence for it: indices 10 and 5 share pitch class 5 and sit exactly two octaves apart, since `10 = 2 * 5`, placing the fixed point in octave relation to an earlier node — the relationship an accumulator has with its own transient. Had the nodes been numbered differently the relation would not hold, and the sensitivity test says so: swapping 09 and 10 displaces a pitch class by 364.81 cents and is load-bearing, while swapping 03 and 12 displaces nothing, both being class 3.

9.2 Series B's parts

Series B part	subject	extended by	what the extension adds	what it costs
I — The Circuit	recursion as a circuit, not a roadblock	CCP-040, CCP-047	the storage law read as a circuit with a computable settling time	the linear form is a modelling choice; CCP-040 computes where it stops
II — The Geometry of Difference	`u`, `N_eff`, the participation ratio	CCP-043, CCP-044, CCP-045	three measured routes to `u` between persons, and the aggregation gap at group scope	the gap is larger than Series B's: 0.394666 octaves on the constructed family
III — Engineering Space	the room as a circuit element	not extended	—	this series is about persons and bonds; the spatial scope stays where Series B left it
IV — The Sociology	collective intelligence as an engineerable quantity	CCP-045, CCP-045	Brooks priced in octaves, the roster ceiling, three structures read end to end	an orthogonal pair already reaches `ln N`, so the group's case rests on grounds the formalism does not price
V — Tradecraft	the survey, the baselines, the traps	Section 6, Protocol C8	the four-baseline suite carried with all three measured holes attached	design controls are mandatory and not substitutable
VI — The Essence Law	`f = k n^2`, folding, the class rule	CCP-041, CCP-042, CCP-046	the law applied to psychological categories via a count-based index; the compound rule as coagula	the index sidesteps CC-R1 only because it is a count; the intensities do not, and must be magnitude-estimated
VII — The Perceptual Lens	the Moment equation `M = (D/P)E`	not extended quantitatively	—	CC-R1's demotion: `D/P` is ordinal pending Rasch instrumentation at 123 items per term, a cost this series does not spend
VIII — The Personal Formulary	the self as storage	CCP-042, CCP-043	traits as low-`delta` states, and a complete personal worksheet	`delta` is not a lever on the deficit, so the personal reading needs both numbers too
IX — The Interpersonal Formulary	the dyad, trust, reciprocity, Ostrom	CCP-043–CCP-045	the dyad given a frequency, a threshold, and an ordering rule	CC-R15: the cheapest-looking intervention accelerates the failure when the numerator is negative
X — Validation	the Formula Validation Protocol	CCP-041, Section 8	`f_min` and the index run through the gates, FVP-6 marked SPECIFIED, NOT RUN	the protocol demotes its own author's formulas, and the demotions are printed
XI — The Old Notation	correspondence tables, gematria	CCP-046-24	the twelve gates audited operation by operation; the schedule finding R2 [Curzi, 2026]	seven rejections across CCP-047 and 24, each with its failing criterion named
XII — The Concordance Method	the two columns	Section 7	this series's register, every rejection printed at the length of an acceptance	precedent is never promoted, and CC-R23 explains why in one number
XIII — Telecommunications to the Cell	the cell as receiver	CCP-048, CCP-048	the protocol restated for psychological content, with the no-subliminals constraint absolute	the open problem, carried whole and unsoftened
XIV — The Uncertainty	Gabor, and Heisenberg correctly stated	CCP-048, §5.6	the four design laws applied: capacity 25, cliffs at 17 and 41, rate factor 8	the rectangular gate is disqualified; implementations need a windowed envelope
XV — Belief and Mechanism	the belief postulate	CCP-049	the postulate applied to the four levers: belief acts only through a coupling with real gain	CC-R23 stands, and testimonial evidence has `LR = 1.000000`
XVI — Rewriting the Narrative	narrative as regulatory input	CCP-048	determination by choice, proved: a corpus closed under negation has resultant exactly zero	a contradictory corpus cannot command — the content-layer protection, and it is structural
XVII — Synthesis	the three instruments, the complete concordance	CCP-049, Sections 7 and 9	the arc across three books, closed on the open problem rather than a triumph	CC-R16: the correspondence is of dynamics and geometry, not of agency; nine disanalogies named

9.3 What this series adds that neither predecessor had Stated tightly, because the temptation at the end of a correspondence table is to claim the whole inheritance as new. The join itself, `f_min = k e^(2c/eta)`. The octave deficit as a computable reading, with the proof that it

is blind to `delta`. The count-based index, which escapes CC-R1's demotion by construction and states its own exposure in the same breath. The magnitude-estimation requirement, with the 1065-cent demonstration that makes it a requirement. The proof that state and relationship frequency are independent readings. And the audit of the old corpus on its own material, of which the schedule finding R2 [Curzi, 2026] is the sharpest instance.

What it does not add is a mechanism. That is the open-problem paper [Curzi, 2026]'s subject and Section 8's fourth drawer, and the series ends where it began: the archive is real, the regulatory layer is real, and the route from an acoustic symbol stream to either does not exist yet.

10 Results

Most of this paper contributes no results and is not intended to. The seven items below are the units other papers in this corpus may cite *to this paper*; everything else in Sections 4 to 9 is restated from the paper that owns it and should be cited there. Each carries its status: **ARITHMETIC** where the value is exact and a reader can redo it, **COMPUTED** where it is the output of an exhaustive search over a stated set, **CLASSIFICATION** where it is a record rather than a finding, and **POSTULATED** where it is asserted by this framework.

R1 — the exact-rational minimum-gap schedule. Folding every n^2 into a single octave as an exact rational, sorting, and taking the smallest adjacent ratio including the wrap gives a step function in the ladder length N with five regimes:

N	minimum gap	cents
2	$2/1$	1200.0000 (degenerate: one class)
3–6	$9/8$	203.9100
7–16	$50/49$	34.9756
≥ 17	$289/288$	6.0008
≥ 41	$1682/1681$	1.0296

The first cliff fires the instant index 17 enters the set, because $17^2 = 289$ folds beside $288 = 2{\cdot}12^2$; the second at index 41, because $41^2 = 1681$ folds beside $1682 = 2{\cdot}29^2$. *Status:* ARITHMETIC, exact. **Both locations are exact and neither depends on** k**.** $N = 12$ sits inside the third regime, so **every reading in this series renders at a minimum separation of 34.9756 cents.**

R2 — the rate penalty at the first cliff is exactly 8. Three distinct ratios describe the size of that cliff and confusing them is a documented error: the interval shrinks by $\ln(50/49)/\ln(289/288) = 5.828$ in cents, the cycle count rises by $288/49 = 5.878$, and **the symbol rate falls by** $(1/32)/(1/256) = 8$**, exactly.** The exactness follows from the two binding pairs sitting at different octave heights: the $50/49$ pair straddles $49/32$ and the $289/288$ pair straddles $9/8 = 288/256$, and Gabor resolution depends on the absolute gap in hertz [Gabor, 1946]. *Status:* ARITHMETIC.

R3 — the price of the affine axiom. The framework's surplus form is uniquely characterised by four axioms, of which S3 — that $e^{F(u)}$ is affine in u — **is forced by nothing.** $F_{\text{alt}}(u) = u \ln N$ satisfies the other three, agrees at both endpoints, and **at** $u = 0.5$ **the two disagree by 33.62 per cent:** 1.871802 **against** 1.242453**.** *Status:* ARITHMETIC on the two stated forms; the axiom itself POSTULATED. **S3 selects one curve from a family that agree at the ends and diverge by a third in the middle, and the middle is where every case lives.** Everything downstream inherits it: the effective count is defined by S3, the essence bridge identifies the index with the effective count, and the keystone is a statement about ln of that count.

R4 — the four-drawer classification. Every claim the series makes is sorted into **PROVED** (derived from stated premises with the arithmetic verified), **POSTULATED** (asserted as a design choice, badged, with the cost of the alternative computed where possible), **BORROWED** (taken from an established result elsewhere, cited, not re-derived) and **UNTESTED** (specified,

with a kill condition, and not run). An item may sit in more than one drawer at different levels of description and both placements are recorded. **The surplus form is recorded as PROVED CONDITIONALLY on a chosen design axiom, and the operationalisation of the twelve-state ladder as SPECIFIED, NOT RUN.** *Status:* CLASSIFICATION — a record of what the framework has paid for, not a finding.

R5 — single-transposition audible capacity. Requiring the whole set $\{kn^2 : n = 1 \dots N\}$ to lie inside 20 Hz–20 kHz after one integer octave transposition 2^j: **the maximum is $N = 25$, and only at $j = 6$.** At $N = 25,\ j = 6$ the set spans 30.9392 Hz to 19337.0293 Hz; at $N = 26$ no admissible j exists, since $j = 6$ puts the top at 20914.9309 Hz and $j = 5$ puts the bottom at 15.4696 Hz. *Status:* COMPUTED by exhaustive search over $j = 0 \dots 15$. **This is a hard bound on essence-set size under single transposition and it constrains any extension of the ladder.**

R6 — the unit conversion table. 1 nat $= 1.442695$ bits $= 2.885390$ octaves $= 3462.4681$ cents, and 1 **bit $= 2$ octaves exactly.** The exactness is not a coincidence: frequency enters as $f = k\,e^{2F}$, so one nat is two e-folds of frequency and one bit — being $\ln 2$ nats — is two doublings. *Status:* ARITHMETIC; the constants exact ($2/\ln 2 = 2.8853900817779268$). The conversion keys are owned by [Curzi, 2026l]; **what this section adds is the statement of why the bit-to-octave relation is exact.**

R7 — the eight measurement protocols. C1 indexing a state; C2 elicitation, with the argument that magnitude estimation is required rather than preferred; C3 the three routes to u with the accuracy guard and the overlap gate; C4–C6 estimating η, c and δ; C7 producing the reading; C8 controls. **C7 requires the ceiling test first and mandates reporting the deficit and the loss rate together, because the deficit cannot see decay.** *Status:* procedures, not findings. **Protocols C4–C6 are estimation practice badged as practice: η, c and δ have no instrument comparable to u's.**

CR — the counter-result. The honesty accounting **prices the framework's central postulate rather than defending it.** The affine effective-count axiom is chosen; an alternative satisfying every other axiom diverges by a third in the middle of the range; the twelve-state ladder's operationalisation is **specified but never run**; the framework's own research programme carries 21 preregistered kill conditions, four of them terminal, and a self-computed **35.8 per cent completion probability**. And the series' central instrument is blind to decay by construction. **A reference apparatus that recorded only what its framework had established would be a sales document.**

11 Discussion

11.1 What follows

The minimum-gap schedule is a design constraint rather than an observation, and it binds any extension of the ladder. A framework operating at $N \leq 16$ has 34.9756 cents of separation to work with; one that adds a seventeenth category has 6.0008 and pays a factor of eight in symbol rate for it. **That price is exact, is available before any extension is attempted, and does not depend on the postulated constant.** Together with R5's capacity bound it says that the ladder's length is constrained from two directions at once — by rendering resolution below and by audible band above — and that $N = 12$ sits comfortably inside both.

The pricing of the affine axiom is the most useful thing in the paper for a sceptical reader. It converts "the framework assumes X" from an objection into a number: the assumption is worth 33.62 per cent at the midpoint of the parameter's range. **A reader who rejects S3 now knows what rejecting it costs**, and a reader who accepts it knows what has been bought on credit.

The four-drawer classification is a genre, not a result, and its value is that it can be checked. Each drawer has an exit (§8.6): PROVED is exited by an error in a derivation or a failed verification, and every verification here states its residual; POSTULATED by exhibiting a system in which the postulate fails; BORROWED by retraction or supersession; UNTESTED by running the test. **An honesty appendix that could not be moved by evidence would be a creed.**

11.2 What does not follow

Nothing here is evidence for the framework. Every exact result in this paper is arithmetic on the framework's own definitions, and arithmetic on a definition tests the arithmetic. R1, R2, R5 and R6 would be exactly as true if the essence law were false.

The protocols are not validated. They are procedures with declared inputs, outputs and gates. None has been run on data, the inter-rater reproducibility of Protocol C1 has not been tested, and Protocols C4–C6 are explicitly badged as estimation practice. **A protocol is a specification of how to be checkable, not a demonstration that anything checked out.**

And the concordance register is not corroboration. Its two columns do different work: one carries evidence, one carries precedent, and **precedent is never promoted.** The reason is exact rather than stylistic — under a selection process that publishes successes and discards failures the likelihood ratio of testimonial evidence is exactly 1.000000 [Curzi, 2026ad], which is not weak evidence but the arithmetic statement that it leaves the prior where it found it.

11.3 Relation to the surrounding programme

The apparatus assembled here is only as good as the results it points at, and two of those results bound what the rest may claim. The framework's own confirmatory attempt **failed its declared baseline test** and reported the failure [Curzi, 2026i], which is the relevant precedent for how any test of these protocols should be reported. And the structural work bounds the reading of any effective count: adding nominally distinct units that are not independent buys capability and not diversification [Curzi, 2026,?], and the concentration instrument transfers to other domains unchanged [Curzi, 2026]. **A card that printed an effective count without that caveat would licence exactly the inference those papers refuse.**

11.4 The reflexive exposure

A reference card changes what it references. Publishing a fixed set of constants, protocols and worksheets makes them the thing that gets used, and the framework asserts that systems change when they are modelled while exempting its own parameters from that reflexivity [Lucas, 1976, Hacking, 1995]; the mechanism by which a published instrument becomes self-confirming is documented [Merton, 1948, MacKenzie and Millo, 2003]. **Nothing here controls for it.** The mitigation this paper can actually supply is §8: a classification that prints what is untested at the same length as what is proved.

12 Verification

A reference card is exactly the document a reader is least likely to check and most likely to trust, so this section decomposes its load-bearing content into steps that can be checked one at a time. Every step carries one of four statuses: **ARITHMETIC** (a reader can redo it), **DEFINITION** (true by the framework's own construction, carrying no empirical weight), **CITED** (an external source supports it, named), or **POSTULATED** (asserted by this framework, no external support).

12.1 The postulates

	postulate	status
P1	The Pythagorean whole tone is the frequency ratio $9{:}8$.	CITED [Barbour, 1951]

	postulate	status
P2	φ is the positive root of $x^2 = x + 1$.	CITED [Euclid, 1956, Vajda, 1989]
P3	A dimensionless ratio of the form $(9/8)^2/\varphi^2$ **is a frequency in hertz**.	**POSTULATED — no source**
P4	The essence of a state with n components has frequency kn^2.	**POSTULATED — no source**
P5	$N_{\text{eff}} = 1 + (N-1)u$ is affine in u (axiom S3).	**POSTULATED — a design axiom, priced at §8.2; not a Hill number**
P6	$F = \ln N_{\text{eff}}$, hence $F \leq \ln N$.	DEFINITION (given P5)
P7	Stored order obeys $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$.	CITED as a *form* [Oppenheim et al., 1996, Elaydi, 2005]; the interpretation POSTULATED
P8	A frequency may be divided by a power of two without changing what it means.	**POSTULATED beyond pitch**; CITED for pitch only [Shepard, 1982]
P9	n and N are cardinalities, hence absolute-scaled.	DEFINITION, typology CITED [Stevens, 1946]
P10	Magnitude estimation yields ratio-scale intensities up to a free multiplier.	CITED and **contested** [Stevens, 1957, Teghtsoonian, 1971, Poulton, 1968]
P11	$\Delta t \, \Delta f \geq 1/(4\pi)$.	CITED — classical Fourier, no quantum content [Gabor, 1946]

Dropped as unnecessary. R1, R2 and R5 — the three exact results originating here — consume **neither P3 nor P4 nor P10.** The minimum-gap schedule is a statement about the fractional parts of $\log_2 n^2$; the rate penalty is a ratio of two ratios; the capacity bound is a comparison of $k\,n^2\,2^j$ against two fixed limits, and while the last does consume k *as a number*, it does not require the number to be a frequency for the ordering to hold — only for the band limits to mean what they say. **No postulate about the ladder's ordering, about consonance, or about the twelve categories is needed for any of the three.**

12.2 The compounding chain

step	consumes	produces	status
L1	P1	$(9/8)^2 = 81/64 = 1.265625$ exactly	ARITHMETIC
L2	P2	$\varphi^2 = \varphi + 1 = 2.618033988749895$	ARITHMETIC
L3	L1, L2	$k = 0.483425732988414$ — **DEFINED, exact**	ARITHMETIC (exact)
L4	L3, **P3**	k carries the unit Hz	**POSTULATED**

step	consumes	produces	status
L5	L4, **P4**	$f(n) = kn^2$ Hz; $kN^2 = 69.6133055503$ at $N = 12$	**POSTULATED**
L6	P5, P6	$F \le \ln N$; $\ln 12 = 2.4849066497880$	DEFINITION
L7	P7, L6	$f_{\min} = k\, e^{2c/\eta}$; $c/\eta \le \ln N$	ARITHMETIC given L5
L8	L7	octaves required $= (2/\ln 2)(c/\eta)$, with $2/\ln 2 = 2.8853900817779$ — **R6**	ARITHMETIC (exact)
L9	L8	1 bit $= \ln 2$ nats $= 2$ octaves **exactly**	ARITHMETIC
L10	P5	$F(0.5) = \ln 6.5 = 1.871802$ at $N = 12$	ARITHMETIC
L11	S1, S2, S4 without P5	$F_{\text{alt}}(0.5) = 0.5 \ln 12 = 1.242453$	ARITHMETIC
L12	L10, L11	disagreement $= 33.62\%$ — **R3**	ARITHMETIC
L13	n^2 folded as exact rationals	five regimes: $2/1$, $9/8$, $50/49$, $289/288$, $1682/1681$ — **R1**	ARITHMETIC (exact)
L14	L13	$17^2 = 289$ beside $2 \cdot 12^2 = 288$; $41^2 = 1681$ beside $2 \cdot 29^2 = 1682$	ARITHMETIC
L15	L13, octave heights	$(1/32)/(1/256) = 8$ exactly — **R2**	ARITHMETIC
L16	L5, band $[20, 20000]$ Hz	$N \le 25$ at $j = 6$ only — **R5**	ARITHMETIC (exhaustive over j)
L17	P9, odd parts of $1 \ldots N$	class count $= \lceil N/2 \rceil$, checked at every $N \le 25$	ARITHMETIC
L18	P11	$1/(4\pi) = 0.07957747155$; and $32/(4\pi) = 2.54647908947$	CITED + ARITHMETIC
L19	the four drawers	the classification — **R4**	**CLASSIFICATION**, not a proof

12.3 What this chain is and is not

Read the status column. L1–L3, L6–L18 are arithmetic or definition and a reader can redo every one of them; L4 and L5 are postulated; L19 is a record. **A chain containing postulated steps is not a proof of its conclusion**, and this one is not offered as one.

Three checks a reader should run, none taking more than a minute, and together they establish the constant rather than leaving it asserted. $(9/8)^2 = 1.265625$. $\varphi^2 = \varphi + 1 = 2.618033988749895$. Dividing gives 0.483425732988414. **That is L1 to L3, and after it the constant is established and only its unit is in dispute.** A fourth check settles R3 in two logarithms: $\ln 6.5 = 1.871802$ against $0.5 \ln 12 = 1.242453$, a gap of 33.62 per cent.

What the chain shows about this paper's own contributions is that they are unusually inert. R1, R2 and R5 sit at L13–L16 and consume no postulate about what the numbers mean; **a reader who rejects P3 and P4 entirely keeps all three**, because a folded lattice is a set of

rationals whatever the reference constant is called. **R3 is the one result that bears on the framework's truth, and what it establishes is the size of a choice rather than an error.**

Two statuses that must not be misread. The exponent 2 at L5 is **DEFINED, not estimated**: no fit was performed and the standard procedure for testing a scaling claim [Clauset et al., 2009] was not run, because there is no data to run it on. And **L19 is a classification, not a verification**: a drawer labelled PROVED contains claims proved *from the framework's own postulates*, which is exactly what §12.1 lists.

13 Limitations

L-1. This paper is not a research contribution and should not be cited as one. It restates results owned elsewhere and adds four items of apparatus. **Citing this paper for a result it merely restates misattributes the result.** §3.1 gives the attribution table.

L-2. Nothing here is tested. No hypothesis is advanced (§2), no data were collected, no model was fitted, no statistical test was run, and no effect size or interval is reported. **The paper contains no empirical content whatever.**

L-3. The dimensional audit does not close in two places, and both are postulates. $f(n) = kn^2$ requires k to carry $[T^{-1}]$ while k is a ratio of two dimensionless quantities; and the compound value is dimensionally hertz-squared, so its "Hz" label is a convention. **No citation in the bibliography repairs either**, and the card prints the unit only because the series prints it.

L-4. The constant is chosen, and its fifteen figures are digits of a definition. $k = (9/8)^2/\varphi^2$ combines two documented objects into a product with no precedent, and the deflationary literature on golden-section claims is cited rather than avoided [Markowsky, 1992, Fischler, 1981, Livio, 2002]. **Fifteen figures on a *measured* quantity would be the classic marker of work that does not understand its own error bars; these fifteen are legitimate only because nothing was measured.**

L-5. Octave reduction beyond pitch is unwarranted [Shepard, 1982, Demany et al., 2013], and every folded column in §5 inherits that. **The arithmetic of the fold is exact; its interpretation is not supported.**

L-6. The effective count is not a Hill number. $1 + (N - 1)u$ is a linear interpolation between 1 and N with no supporting literature. The Hill family [Hill, 1973, Jost, 2006] and its political-science rediscovery [Laakso and Taagepera, 1979] ground $1/\mathrm{HHI}$ and $\exp(H)$ exactly and ground this construction **not at all.** A reader reaching for that literature as support will find it does not apply, and §8.2 prices the axiom rather than defending it.

L-7. Conjoint measurement was not tested, and this is the framework's largest standing liability. The ratio-scale treatment of η, c, δ and the elicited intensities is provisional; Michell's objection stands unanswered [Michell, 1997, 1999] for every quantity in the series except the counts; the available remedies [Luce and Tukey, 1964, Krantz et al., 1971, Rasch, 1960] were not run, and the framework's own price for the Rasch route — 123 items per term — was not paid.

L-8. The reflexive specification is not identified. Stated at §3.7 and repeated because the protocols depend on it: the endogenous effect in a reflexive specification cannot be recovered from observational data [Manski, 1993, Shalizi and Thomas, 2011], the remedy adopted is stipulation of the reference set [Brock and Durlauf, 2001], and **every reading Protocols C3–C7 produce is therefore a modelling position rather than an identified empirical result.**

L-9. The framework exempts its own parameters from its own reflexivity [Lucas, 1976, Hacking, 1995, Merton, 1948]. Not resolved.

L-10. The functional form is underdetermined by the evidence. A sub-linear power law [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the qualitative facts as well as the framework's logarithm does, and many mechanisms produce indistinguishable curves [Stumpf and Porter, 2012]. **Nothing in this paper discriminates among them**, and §8.2's pricing of S3 is a statement about one alternative within the framework's own axiom set, not a survey of the space.

L-11. The warrant the logarithm needs is not the one usually offered. The psychophysical route [Fechner, 1860] inherits every dispute in that literature; the multiplicative-dynamics route derives a logarithmic form with **no psychological assumption at all** [Peters and Gell-Mann, 2016]. **The tables here need only the second**, and where the series leans on the first it inherits the disputes catalogued at §3.6.

L-12. The exponent is defined, not estimated [Clauset et al., 2009]. **No empirical scaling claim is made anywhere in this paper.**

L-13. Scalar reduction is a habit of this framework, inherited here rather than re-conceded. The minimum-frequency paper states the limitation once and in full [Curzi, 2026]: the framework replaces phenomena its own source literature describes by several jointly varying features with a single number and proceeds arithmetically on the number. **That concession governs here, and this paper's instance is the card itself** — a one-page summary of a framework whose limits live in the argument it invites the reader to skip. The mitigation is the one the card actually applies: every card carries its limits inline. **It is a mitigation and not an answer.**

L-14. Protocol C1's inter-rater reproducibility has not been tested. Whether the component count is a property of the episode or of the analyst is the gate that decides whether the ladder measures anything, it is recorded in §8.4 as **SPECIFIED, NOT RUN**, and the worksheets in §6 cannot substitute for it.

L-15. Protocols C4–C6 are estimation practice badged as practice. η, c and δ have no instrument comparable to u's. C4's count of surviving artefacts is crude; C5 requires a declared conversion in which most of the error lives; C6 inverts a half-life from a judged surviving fraction. **All three are better than a guess only because the objection to them is visible.**

L-16. The controls the series recommends have three measured failure modes, and this paper carries them rather than the recommendation alone. §6's Protocol C8 records the shuffled-input control at a measured false-rejection rate of 0.108 against a nominal 0.05 — **anti-conservative by a factor of 2.2** — together with its blindness to repackaging and its 98.35 per cent pass rate against a shared slow driver. **Adding surrogates does not fix a design problem**, and the remedy is design controls: blinding, channel separation, prospective sampling, a yoked control.

L-17. Two of the works this paper's borrowings depend on are named without citation keys. The BORROWED drawer at §8.3 lists sources in the framework's own compressed form, and several of them — the process-control and biological entries in particular — are named with journal, volume and year in the running text because **they are absent from the bibliographic database governing this corpus.** That is a gap in the bibliography and is disclosed rather than papered over.

L-18. The interval and gap arithmetic makes no claim about consonance. Two established measures of ratio complexity exist [Euler, 1739, Tenney, 2015] and **neither is used**; §5.6 reports distances in cents, and what a listener finds consonant depends on spectrum rather than ratio [Sethares, 2005, Plomp and Levelt, 1965].

L-19. There is no conservation law [Mirowski, 1989]; validity is claimed only within a specified arrangement [Cartwright, 1999], within the boundaries the sociophysics literature draws [Castellano et al., 2009]. **L-20. Ergodicity is assumed and not argued** [Peters, 2019], and aggregation over heterogeneous agents need not yield a representative agent [Kirman, 1992].

A precision inconsistency in the source, recorded rather than repaired. The card prints k^2 as 0.2337004393 "to ten places" while every computation uses 0.233700439315386. The two are consistent — the first is the second, rounded — but the compound-rule protocol of [Curzi, 2026] performs an integrality test whose sensitivity at large component counts **exceeds ten decimal places**, so a value published to the card's precision cannot support the test the specification demands. **The card and the protocol are in tension**, and the resolution adopted here is that every computation uses the full-precision constant and the card's ten-place figure is a display form only.

A second, and it is a live one. The source's own note at §4.1 records that a residual bias in a published element table was traced to φ held at fifteen significant figures rather than at full double precision. **That is a statement that fifteen figures was, in at least one instance, not**

enough — which sits awkwardly beside a card that prints k to exactly fifteen. The instruction the card gives is the right one and is repeated here: **carry the constant, not its printed form.**

14 Conclusion

What this paper is. The reference and status apparatus for a framework whose results are established elsewhere: a constants card with its badges attached, frequency and interval tables, eight measurement protocols with blank worksheets, a concordance register that prints its rejections at the length of its acceptances, a four-drawer classification of every claim the series makes, and a correspondence table to the two predecessor series. **It advances no hypothesis and it is not a research contribution.**

What originates here. The folded-lattice minimum gap as an exact function of ladder length — 1200 cents at $N = 2$, 203.9100 at $N = 3 \dots 6$, 34.9756 at $N = 7 \dots 16$, 6.0008 at $N \geq 17$ and 1.0296 at $N \geq 41$ — with the symbol-rate penalty at the first cliff **exactly 8**. The single-transposition audible capacity, $N \leq 25$ **and only at** $j = 6$. The statement that 1 bit is **2 octaves exactly**, and why. And the four-drawer classification, whose most useful entry is a price rather than a defence: **the framework's central design axiom is chosen, and an alternative satisfying every other axiom and agreeing at both endpoints disagrees with it by 33.62 per cent at** $u = 0.5$**.**

What is proposed rather than shown. That $(9/8)^2/\varphi^2$ is a frequency; that essence goes as kn^2; that octave equivalence carries beyond pitch; and that the effective count is affine in the geometry. **All four are postulates of this framework, badged at every appearance**, and the three exact results originating here survive the rejection of the first three, because a folded lattice is a set of rationals whatever its reference constant is called.

What is recorded rather than claimed. That the operationalisation of the twelve-state ladder is **specified and not run**; that the acoustic causal claim is **not confirmed**; that the decisive experiment is **designed and not run**; that the framework's central instrument is **blind to decay**; and that the framework's own estimate of its research programme's completion probability is **35.8 per cent**. **A reference apparatus recording only what its framework had established would be a sales document, and this one is not offered as an argument for anything.**

15 Declarations

Data and code availability. No dataset was collected for this study and none exists. Every number in this paper is the output of exact arithmetic — in several places exact rational arithmetic — or of exhaustive search over a stated finite set, on inputs stated in the text. The worked cases carried through the protocols are **constructed** and are labelled as constructed at the point of use. The figures are generated by the Python scripts named in their captions and are available from the author on request. **No empirical measurement of any person, relation or institution underlies any figure, table or number in this paper.**

Funding. This research received no external funding.

Conflicts of interest. The author declares no conflicts of interest.

Author contributions. Michael Laurence Curzi is the sole author and is responsible for the conception, the compilation, the computations, the figures and the manuscript.

Ethics. No human-subjects research was conducted and no personal data were processed. The protocols in Section 6 specify a consent procedure, a right to decline without reason at every step, and a prohibition on anything masked or below awareness; **none of them has been applied to anyone, because no reading of any real person or relation was taken.** The blank worksheets are specifications, not records.

Preregistration. Nothing here was preregistered, because no test was run. §8.4 records the framework's untested items **with their kill conditions stated in advance**, which is the form a preregistration would take; the framework's one preregistered confirmatory attempt elsewhere failed its own declared baseline margin and reported the failure [Curzi, 2026i].

Bibliographic disclosure. Several works listed in the BORROWED drawer of §8.3 are named in the running text without citation keys because they are absent from the bibliographic database governing this corpus (§13, L-17). They are given with enough bibliographic detail to be located.

Clinical disclaimer. The twelve state classes are not diagnostic categories. Nothing in this paper is an assessment, a screening instrument, a treatment, counselling, or advice about any person's health, and nothing in it may be presented as such. **No person, tissue, cell or molecule is claimed to resonate at any frequency named here**; the frequencies represent a count and an angle, and nothing in the formalism asserts a physical oscillation.

Part IV

Structure and Resilience

Chapter 51

Correlated Jurisdictions: Coordination Perimeters, Erosion, and the Economics of Full Disclosure

Abstract. Fifteen years of automatic information exchange, economic-substance requirements and beneficial-ownership registration removed concealment as a durable property of cross-border asset structures. We ask what remains once concealment is unavailable, and answer with two measurements. First, coordination. Four jurisdictions that look diverse on an entity map — the United Kingdom at 40.91 per cent of exposure, the United States at 27.27, Jersey at 18.18 and Luxembourg at 13.64 — give a naive Herfindahl index of 0.293388 and 3.4085 effective jurisdictions, but they collapse into a single coordination perimeter with a bloc index of 1.000000 and exactly 1.0000 effective perimeters: a 3.41-fold overstatement. Assuming independence where a shared perimeter exists understates joint impairment by three orders of magnitude; a common shock of two per cent raises the probability that all four fail together from 0.000100 to 0.020044, a factor of 200.4, and over fifteen years converts a one-in-seven-hundred event into a four-in-five event. Second, disclosure. A concealment-based arrangement costing USD 210,000 against a compliant USD 561,000 shows a 62.6 per cent saving on the invoice, but is more expensive above an annual examination probability of only 0.85 per cent, and halving the loss on examination moves that threshold only to 1.7 per cent. Elements are scored on five criteria multiplicatively, so one failing element zeroes the structure. The erosion rate is decomposed as 2.0 per cent a year, a half-life of 34.31 years. Both the observed and the redesigned structures remain below break-even.

Keywords: coordination perimeter; correlated failure; automatic exchange of information; economic substance; concentration measurement; disclosure economics; asset-holding structures

Series: D. Structure and Resilience **Paper:** CCP-051

1. Introduction

Between 2009 and 2024 the jurisdictions whose commercial product was confidentiality lost the ability to deliver it. That much is a matter of public record. This paper asks the engineering question that record poses — *what remains once concealment is unavailable, and how should what remains be measured* — and answers it with two measurements and one identification.

The identification comes first because everything else is downstream of it. The jurisdictions did not erode separately. They eroded together, because they were attached to a common set of mechanisms: automatic exchange of financial account information [Organisation for Economic Co-operation and Development, 2017], the anti-money-laundering recommendations that national supervisors implement in common [Financial Action Task Force, 2012], the substance and treaty-abuse work of the base-erosion project [Organisation for Economic Co-operation and Development, 2015], listing frameworks whose designations propagate without further decision [de Goede, 2011], and the private, unappealable correspondent-banking decisions that consume all of the foregoing as inputs. **The correct unit of analysis for concentration risk is therefore not the jurisdiction but the coordination perimeter**, defined here as a set of jurisdictions among which an adverse decision propagates without a separate act of consent by each of them.

The first measurement is what that identification costs a structure that ignores it. Four jurisdictions that look diverse on an entity map — the United Kingdom at 40.91 per cent of exposure, the United States at 27.27, Jersey at 18.18 and Luxembourg at 13.64 — give a naive Herfindahl index of 0.293388 and 3.4085 effective jurisdictions, and collapse to a single coordination perimeter with a bloc index of 1.000000 and exactly 1.0000 effective perimeters. The overstatement is a factor of 3.41, and it arises from a modelling assumption rather than from an arithmetical error. The tail consequence is larger than the point estimate suggests: holding each jurisdiction's marginal failure probability fixed at ten per cent, a common-shock component of only two per cent per year raises the probability that all four fail together from 0.000100 to 0.020044, a factor of 200.4.

The second measurement prices the alternative that the erosion appears to recommend, which is more concealment, and finds it already uneconomic. A constructed concealment-based arrangement costing USD 210,000 a year against a compliant USD 561,000 shows a 62.6 per cent saving on the invoice and breaks even at an annual examination probability of $p^{*} = 0.008458$. **Above 0.85 per cent a year, concealment is the more expensive design**, and halving the loss on examination moves the threshold only to 1.7 per cent.

The instrument used throughout is the Herfindahl–Hirschman Index, introduced by Hirschman [1945] in root form to describe the concentration of a country's foreign trade and reached independently in sum-of-squares form by Herfindahl [1950]. Its full development for this application, including the variance identity that makes the bloc collapse exact rather than convenient, belongs to Curzi [2026], which establishes that the redesigned six-jurisdiction structure reaches a naive index of 0.189002 (5.29 effective jurisdictions) and a bloc index of 0.340000 (2.94 effective blocs), with Western-perimeter exposure falling from 100.0 to 30.0 per cent. **This paper supplies the evidence base that forces the collapse, the test that decides bloc membership, and the economics that determine whether the resulting structure can survive being looked at.**

Three things this paper does not do should be stated at the outset. It fits no distribution and asserts no scaling law; Clauset et al. [2009] is the reason the procedure would have to be run before any such claim could be made, and it is not run here because no claim requiring it is made. It offers no operational guidance for moving value that a sanctions regime prohibits. And it does not cite jurisdiction-specific primary law: **the bibliography contains general treaty, standard-setting and doctrinal material only, and every statute, case, registry document and jurisdictional characteristic described below is described from the source manuscript, left uncited, and itemised in Appendix B.** That is a real gap, it is not closed here, and it is declared again in the Declarations.

Curzi [2026] supplies the framing this paper presupposes: comfort — here, held capital under protection — is the felt total of present costs, so the present arrangement is a local minimum of the cost landscape by definition rather than by luck, which is why a structure that has stopped being optimal continues to feel settled from inside it.

2. Hypotheses

Five hypotheses are stated, each with what would falsify it, in the form Platt [1964] requires of a discriminating test. Meehl [1967] is the standing warning: a directional prediction in this territory is nearly costless, and the defensible reply is a point prediction with constants attached, which H1, H2 and H4 make. One of the five is reported in Section 8 as **not met**.

H1 — the coordination hypothesis. Jurisdictions attached to a common set of transmission mechanisms fail together rather than independently, so that a concentration measurement taken at jurisdiction level overstates resilience. *Predicted:* for the observed structure, the effective count computed on raw shares exceeds the effective count computed on perimeter-collapsed shares by a factor above three — specifically 3.4085 against 1.0000. *Falsified by:* an observed pattern in which the four named jurisdictions changed their treatment of holding structures on independent schedules and by independent instruments, with no common trigger identifiable in the record.

H2 — the common-component hypothesis. The joint-failure probability of a perimeter-attached set is governed almost entirely by the common component and almost not at all by the idiosyncratic one. *Predicted:* holding the marginal failure probability at 0.10, a common component of $c = 0.02$ raises $P(\text{all four})$ from $p^4 = 0.000100$ to 0.020044, a factor of 200.4; and $P(\text{all four}) \to c$ as c rises. *Falsified by:* a decomposition in which the idiosyncratic term contributes materially to the joint probability at any $c > 0$ that is small relative to the marginal probability.

H3 — the disclosure-economics hypothesis. A concealment-based structure is more expensive than a compliant one at any realistic examination probability, once the expected loss on examination is included rather than the invoice alone. *Predicted:* at the constructed parameters the break-even examination probability is $p^* = (561{,}000 - 210{,}000)/41{,}500{,}000 = 0.008458$, and the result is insensitive to the loss parameter — halving L moves p^* only to 0.016916. *Falsified by:* a defensible estimate of annual examination probability below 0.85 per cent in the environment described in Section 4, or a cost structure in which the compliant architecture's premium exceeds the expected loss.

H4 — the reallocation hypothesis. The reallocation required to reach two effective perimeters is structural rather than marginal. *Predicted:* with the exiting share split evenly between two further perimeters, the in-perimeter share must fall to exactly $2/3$ before the effective count reaches 2.0000; moving ten per cent buys only 1.2270. *Falsified by:* any allocation reaching $N = 2$ while retaining more than two-thirds of exposure inside one perimeter.

H5 — the sufficiency hypothesis. Perimeter redesign alone lifts a structure to a non-negative surplus balance. *Criterion, fixed in advance:* with $h = \eta \ln N - \delta Q - C$ at $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$, the balance is non-negative only if $\ln N \geq 0.5000/0.35 = 1.428571$, that is $N \geq e^{1.428571} = 4.172734$. *Falsified by:* a redesigned structure whose effective perimeter count falls short of 4.172734 at the stated parameters.

Not hypotheses, and labelled as such. Four statements in this paper are not falsifiable as written, and each is named here for what it is rather than allowed to travel as a finding.

- *The logarithmic form of the input term*, $\eta \ln N$ rather than ηN, is a **modelling choice** carried from Curzi [2026]. It is not tested here, and Section 12 records that Latané's sub-linear power law and an $n \log n$ form fit the qualitative facts equally well.
- *The identification of N with the effective perimeter count* is a **definition**. All of its content sits in the perimeter test of Section 3.4, which is testable; none of it sits in the identification.
- *The parameter set* $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ is a **modelling choice**, fixed in advance so that H5's criterion cannot be adjusted after the measurement — the preregistration discipline of Nosek et al. [2018] applied to an engineering measurement.
- *The erosion decomposition* $\delta = 0.020$ into four named components is a **constructed illustration**, not a research estimate, and is labelled at its point of use.

3. Method

3.1 The storage law, and the one term that must be reinterpreted

The framework's storage law relates a stock held against decay to its decay rate, the efficiency with which flow converts into stock, and the maintenance cost of holding it. In difference and surplus form:

```
Q_{t+1} = (1 − delta) Q_t + eta S_t − C_t

h = eta ln(N) − delta Q − C
```

with Q the protected capital actually held under durable protection, δ the erosion rate, η the conversion of income into durable protected holdings, S the income flow available for conversion, C the recurring structure maintenance, N the effective diversity count, and h the available surplus above upkeep. Curzi [2026b] states the headroom gauge in full — three regimes, a maximum increment $r_{\max} = h_t$, and the separate constraints of admissibility $C_t \leq (1-\delta)Q_t + \eta S_t$ and sustainability $C_t \leq \eta \ln N - \delta Q_t$ — and this paper reuses that instrument without restating its derivation. The difference equation itself is a first-order linear recurrence and its fixed-point behaviour is standard [Elaydi, 2005]; Curzi [2026] states the sharper corpus result, that refinement cannot move the fixed point because $Q_\infty = (\eta F - c)/\delta$ contains no t.

Three of the five terms transfer to the capital case without commentary. Q is what is actually held under protection rather than what is nominally owned. C is real, recurring, and the term most often set to zero in informal planning. S and η together say only that income helps to the extent it is converted into durable protected stock rather than consumed or left in a form that inherits the erosion rate.

δ deserves a note because in the capital case it aggregates four processes that are usually discussed separately and should be added: tax drag, the probability-weighted cost of arbitrary action, debasement of the unit of account, and degradation of the jurisdictional characteristics the structure was built to use. Section 4.6 decomposes it.

That leaves N, and N is the whole argument. In the goods case N counts independent stores, and the word *independent* does its work quietly because in the physical case independence is visible — two warehouses on one flood plain are one store, and anyone can see it. Capital has no flood plains. **In the capital case N is the effective number of independent coordination perimeters: not the number of entities, and not the number of jurisdictions.**

The consequence is severest at the boundary and is the framework's hardest single result. If every entity sits inside one perimeter then $N = 1$, $\ln 1 = 0$, and the entire input term vanishes:

```
h = eta ln(1) − delta Q − C
  = 0 − delta Q − C
  = −(delta Q + C)
```

A structure confined to one coordination perimeter has not a small available surplus but a negative one, equal to the whole of its erosion plus the whole of its upkeep, and no rearrangement of entities inside that perimeter changes it. Adding a fifth entity in a fifth jurisdiction inside the same perimeter costs C and adds nothing to N. That this is an *organisational* result and not only a jurisdictional one was established long before it was measured: Perrow [1984] showed that interactive complexity plus tight coupling makes catastrophic failure a property of the system rather than of any component, and the finding here is Perrow's in a measurement frame. It is cited as **qualitative theory and framing, never as a formal result**.

3.2 Which effective count, stated exactly

The corpus runs two quantities under the name *effective count* and they are not the same quantity. This paper uses one of them and says so at every point of use, because a reviewer who reaches for the diversity literature will find that it covers one exactly and the other not at all.

The reciprocal index is a Hill number of order two. Hill [1973] defined the family ${}^{q}D = (\sum_i p_i^q)^{1/(1-q)}$, whose members return the number of equally-abundant types that would produce the observed value of the underlying index; at $q = 2$ this is $(\sum_i p_i^2)^{-1}$, which is HHI^{-1} exactly. Jost [2006] is the statement that entropies are not diversities and that their numbers-equivalents are. The numbers- equivalent reading of the concentration index itself belongs to Adelman [1969]. The same arithmetic was reached independently in ecology [Simpson, 1949] and in political science as the effective number of parties [Laakso and Taagepera, 1979]; the exponential of Shannon entropy [Shannon, 1948] is the same family at $q = 1$, and the reciprocal index is also the exponential of the Rényi entropy of order two [Rényi, 1961]. **Every effective count in this paper is HHI^{-1}, the Hill number of order two, and is named as such.**

The corpus's affine count is not a Hill number and has no external support. The quantity $N_{\text{eff}} = 1 + (N-1)u$ used in Curzi [2026] and Curzi [2026c] is a linear interpolation between 1 and N. It is not a Hill number of any order, it is not a numbers-equivalent, and no entry in ecology, economics, political science or information theory states it. Curzi [2026] runs the two definitions through the same verification chain and shows where the second one fails; this paper does not use it, and the distinction must survive any citation of this paper. Curzi [2026c] is cited here for a different result entirely — the sharp Lipschitz constant $N-1$ read as a measurement-error budget, under which $N = 12$ and $u = 0.05 \pm 0.05$ admits surplus anywhere in $[0, 0.7419]$ against a width of 0.1266 at $u = 0.70$ — because that is the correct way to think about the error budget on a bloc assignment.

The generalisation the collapse relies on is $\mathrm{EFF} = 1/(w'Rw)$ for exposure weights w and correlation matrix R, which is the similarity-sensitive diversity of order two of Leinster and Cobbold [2012], whose additive relative is the quadratic entropy of Rao [1982]. Setting $R = I$ recovers HHI^{-1}; setting R to the block matrix of ones recovers the bloc index. Curzi [2026] derives the identity in full and owns it as a result; this paper consumes it. Hurlbert [1971] is the standing discipline on the word *diversity* — it has been used for so many distinct quantities that it carries no content without its formula attached, which is why the formula is attached here every time.

3.3 The reflection problem, stated where the loop is used

The perimeter assignment is a forecast about correlated behaviour, and the framework's central formal object is the reflexive loop $T(x) = W(x, M(x))$ of Curzi [2026], in which the group influences the member while the members constitute the group. **That is exactly a linear-in-means reflexive specification, and Manski [1993] proves that its endogenous social effect is not identified from observational data**: it cannot be separated from contextual effects or from correlated unobservables, however much data is collected. Shalizi and Thomas [2011] reinforces the point for network data, where influence and selection are generically confounded. **No amount of observation of jurisdictional behaviour could therefore confirm a perimeter assignment derived reflexively, and this paper states that where the assignment is made rather than in a footnote.**

Four remedies are known — randomisation, an instrument, partial-population variation, and stipulating the reference group with a defence of the stipulation. **The fourth is adopted here.** The perimeter membership is fixed in advance by the four-question test of Section 3.4, recorded with its evidence, and stress-tested; Brock and Durlauf [2001] is the template for making a stipulated reference group rigorous, and it is the only one of the four available in this setting. **The price is that the perimeter assignment is a modelling position and not an identified empirical result.** That price is paid explicitly here, again in Section 10, and again in

Section 12.

The same reflexivity applies to the instrument itself and is not exempted. MacKenzie and Millo [2003] documents an option-pricing model that did not fit the market when published and came to fit it after traders adopted it; Hacking [1995] gives the general form, in which a classification of a population interacts with the population classified. **If enough structures are built on the assumption that two jurisdictions share a perimeter, capital and correspondent relationships reorganise around that assumption and the assumed correlation becomes the observed one.** Lucas [1976] states the severe consequence: a parameter estimated in a population that has not read the instrument cannot be assumed stable in a population that has. **The framework asserts that systems change when modelled and then exempts its own η, δ and correlation assumptions from that same reflexivity. That is an inconsistency, it is not resolved here, and what is offered instead is the recorded, dated assignment of Section 3.4.**

3.4 The perimeter test, stated operationally

A coordination perimeter is a set of jurisdictions among which an adverse decision propagates without a separate act of consent by each of them, because they share a mechanism of transmission. The operative words are *propagates* and *without separate consent*. Two jurisdictions with similar laws are not thereby in one perimeter. Two jurisdictions whose treatment both change when a third body assigns a rating are in one perimeter however different their laws.

Four questions, answered yes or no. They are deliberately about attachment to mechanisms rather than about a jurisdiction's own conduct, because conduct is not what transmits the shock.

```
Q1  AUTOMATIC EXCHANGE
    Does the jurisdiction participate in automatic exchange of financial
    account information with the other jurisdictions in the structure?

Q2  STANDARD-SETTING NETWORK
    Is the jurisdiction a member of, or assessed by a regional body of,
    the anti-money-laundering standard-setter, and does its supervisory
    regime implement the common recommendations?

Q3  LISTING REACH
    Can the jurisdiction be placed on a coordinated non-cooperative or
    increased-monitoring list, with consequences its own government
    cannot prevent?

Q4  SETTLEMENT DEPENDENCE
    Do the jurisdiction's institutions depend, for clearing in the
    principal reserve currencies, on correspondents that are themselves
    subject to the jurisdiction of the states that issue them?

Four yeses    fully inside the perimeter
Three yeses   inside for all practical purposes
Two or fewer  assess individually; may constitute or belong to another
              perimeter.  Q4 is the heaviest question and a yes to Q4
              alone is close to decisive.
```

The instruments that make Q1 to Q3 operative are documented and citable: automatic exchange of account information [Organisation for Economic Co-operation and Development, 2017], the recommendations supervisors implement in common [Financial Action Task Force, 2012], and the substance and treaty-abuse minimum standards [Organisation for Economic Co-operation and Development, 2015]. de Goede [2011] documents the listing machinery itself

as a technique of governance, which is the mechanism Q3 is trying to detect. **The jurisdiction-by-jurisdiction answers below are the analyst's judgements applied to published characteristics and are not cited to primary sources; they are listed in Appendix B.**

Q4 carries more weight than the others because it is the enforcement currency. A jurisdiction can decline to exchange, ignore the recommendations, and shrug at a listing, and still be reachable, because its banks need clearing access. Conversely a jurisdiction that exchanges scrupulously may sit partly outside the perimeter if its settlement dependence is genuinely diversified. **Compliance posture and perimeter membership are different variables and must never be collapsed into one.** The architecture recommended here is maximally compliant and deliberately not concentrated in one perimeter, and there is no contradiction in that.

The governance literature supplies an independent name for what the test counts. Ostrom [2010] develops polycentricity — multiple centres of decision-making formally independent of one another — as an empirically derived property of durable governance systems. **Polycentricity is the effective count of perimeters in governance language, arrived at from field data rather than from an index**, and the convergence is worth recording.

3.5 Collapse before you measure

```
1. List every entity and its share of protected capital.
2. Apply the four-question test to each entity's jurisdiction.
3. COLLAPSE: sum the shares within each perimeter.
4. Measure concentration on the COLLAPSED shares, never the raw ones.
5. Report BOTH figures, and treat the gap as the size of the illusion
   the structure was operating under.
```

Step 3 is where the analysis is usually lost, and step 5 is what prevents the collapse being used selectively. Order matters in step 4: sum within blocs first, then square, then sum, then invert. Reversing that silently reproduces the naive number.

Grouping data until the answer improves is a recognised way to lie, and the defence against the charge is that the grouping is not an aggregation convenience but the correlation matrix R written down. Markowitz [1952] supplies the algebra: for exposure weights w_i and unit-variance shocks X_i, the aggregate shock $R = \sum_i w_i X_i$ has variance $\sum_i \sum_j w_i w_j \rho_{ij}$, which is the naive index when $\rho_{ij} = 0$ off the diagonal and the bloc index when $\rho_{ij} = 1$ within blocs. **Diversification benefit depends on covariance and not on count, and that was stated with the mathematics attached in 1952.** This paper's contribution is the evidence that the covariance is there; Curzi [2026] owns the identity.

3.6 Evidence standard, and what is not cited

Every claim in Sections 4 and 5 about a named statute, directive, decided case, listing event, register or jurisdictional characteristic is **described from the source manuscript and is not cited to the bibliography**. The bibliography carries general instruments and comparative doctrine — recognition and enforcement of awards [United Nations, 1958], the absence of a succession convention in force [Hague Conference on Private International Law, 1989] and the regional succession regulation that partially fills the gap [European Union, 2012], the exchange and standard-setting instruments already named, asset partitioning as the irreducible function of entity law [Hansmann and Kraakman, 2000], and the best scholarly account of legal coding as capital formation together with its sharpest critique [Pistor, 2019]. **It carries no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction.** Appendix B lists the affected claims individually. Where a table below would otherwise score a named jurisdiction on a characteristic the bibliography cannot support, the jurisdiction is de-identified — "Gulf centre A", "an Atlantic risk-transfer centre" — following the convention adopted in Curzi [2026].

4. The erosion, as evidence

This section is the evidence base for H1. Five tracks eroded on separate schedules, by separate instruments, with separate institutional homes — which is precisely why the common cause is worth isolating. **The claim being tested is not that these jurisdictions are finished.** Several are excellent and the architecture assembled downstream uses parts of what they do. The claim is narrower: *the specific product these jurisdictions once sold — differential opacity, the condition of being harder to see into than the alternatives — has been substantially eliminated, and what remains is legal capability, which is a different product with a different price and a different competitive set.*

4.1 Track one: banking secrecy in a European financial centre

The protection at issue was never anonymity from the domestic authorities. It was a statutory prohibition, backed by criminal penalty, on an employee disclosing client information to third parties, together with a historical distinction between evasion as an administrative matter and fraud as a criminal one under which assistance was granted only for the latter. That distinction was the commercial substance of the product, and its removal was the erosion.

The sequence, described from the source: a deferred prosecution agreement with a foreign prosecutor in February 2009 and a substantial penalty; an intergovernmental arrangement providing for disclosure of a large number of accounts, which established that the statutory prohibition could be overridden by agreement; a unilateral, extraterritorial withholding statute enacted in March 2010 whose mechanism was a thirty per cent withholding on source payments to institutions that would not report, live from July 2014; criminal pleas by two institutions in 2013 and 2014, one of them founded in 1741 and thereafter extinguished; a voluntary disclosure programme for the sector announced in August 2013; and finally a multilateral automatic-exchange standard approved in 2014 and operative for the centre from 2018 in respect of 2017 data [Organisation for Economic Co-operation and Development, 2017].

Count the interval. The protective statute dated from 1934; the first automatic exchange occurred in 2018. Eighty-four years of statutory protection ended in nine, and the statute was never repealed — it simply stopped covering the case that mattered. That is the pattern to internalise, because it recurs on every track: the protective law generally survives, and what changes is the set of circumstances in which it applies.

A note on the leak literature, which is routinely treated as the cause and was in fact an accelerant. Five large document publications between 2015 and 2021 changed the political weather and made ministerial action cheap. They did not create the legal instruments; the withholding statute preceded all of them. **A structure that depended on the absence of a leak was always depending on something no adviser could price.**

4.2 Track two: company jurisdictions and the substance requirement

A different instrument reached the Caribbean company jurisdictions. Their product was not secrecy in the sense above; it was a company that could be formed cheaply, held quietly, taxed at zero, and required to do almost nothing. The last element was attacked first, through a substance criterion: a jurisdiction with a zero or near-zero corporate rate had to require entities claiming its residence actually to conduct relevant activity there, directed and managed locally, with adequate people, premises and expenditure. The alternative was listing. Two such jurisdictions legislated substance regimes with effect from 1 January 2019, and both have been revised since. The substance principle itself is the base-erosion project's [Organisation for Economic Co-operation and Development, 2015]; the individual enactments are not cited here.

The listing history is the useful part, and it is a sequence rather than a snapshot:

jurisdiction	event	period
Caribbean centre A	coordinated non-cooperative list, Feb 2020 to Oct 2020	8 months
Caribbean centre A	increased monitoring, Feb 2021 to Oct 2023	32 months
Caribbean centre B	listed Oct 2023 after a peer-review downgrade on exchange of information on request; removed Feb 2024	4 months
Central American centre	increased monitoring, Jun 2019 to Oct 2023	52 months
Gulf centre	increased monitoring, Mar 2022 to Feb 2024	23 months

Dates as publicly announced and described from the source manuscript; not cited to primary releases (Appendix B, item 3).

Read the table for what it shows, which is not that these are bad places — every one of them came off. **What it shows is volatility in a variable that most structures treat as a constant.** A holding company's jurisdiction is chosen once and left for a decade. Its listing status can change twice in eighteen months, and each change propagates into bank onboarding, counterparty diligence, and in some states into concrete defensive measures.

The four-month episode in the third row is the single most instructive item and Section 5.2 returns to it, because it is a traced propagation rather than a description.

4.3 Track three: the repeat listing

The Central American centre's history illustrates a distinct point: the cycle repeats, and exiting a list is not the same as exiting the category of jurisdiction that gets listed. It was under increased monitoring in the period 2014 to 2016, removed, listed again in June 2019 after a large document publication and the resulting pressure, and removed in October 2023. Fifty-two months on the second occasion.

Two observations follow. First, the political economy: a small jurisdiction whose financial sector is a large fraction of its economy has limited capacity to refuse a demand from a body whose members control its correspondent access, and the listing process should not be modelled as a negotiation between equals. Second, and more useful: **it complied. It passed the legislation, built the ownership register, submitted to assessment, and got off the list. Compliance worked, and compliance rather than resistance was the successful strategy in every one of these episodes** — a data point in favour of the design standard of Section 7, arriving from an unexpected direction.

For completeness, and because it is frequently misreported: the criminal proceedings arising from that jurisdiction's largest document disclosure resulted in acquittals for the principal defendants in a verdict reported in June 2024. **The firm's founders were acquitted and the firm and the model it represented were nonetheless destroyed years earlier. The legal outcome and the commercial outcome were opposite, and the structure died of the process rather than of the verdict.**

4.4 Track four: the protective statute and the person

The fourth track is different in kind. It is not about registers, lists or exchange. It is about what happens when the theory of a structure meets a court with the power to detain the person in front of it.

A well-known asset-protection statute of 1984, amended repeatedly since, included short limitation periods for transfer challenges, a high standard of proof, and provisions to the effect that a foreign judgment would not be recognised or enforced against the vehicle. Two appellate decisions in another jurisdiction define the failure mode. In the first, a regulator obtained repatriation orders against principals who had retained roles in the structure; when they asserted that compliance had become impossible because a foreign fiduciary had invoked a duress provision, they were held in civil contempt and taken into custody, and the appellate court affirmed on the ground that a party asserting impossibility bears the burden of proving it and that the impossibility had been designed and triggered by the contemnors themselves. In the second, a debtor who had settled a substantial offshore trust was ordered to turn over its assets, asserted that he could not, was held in civil contempt from 1999 and was reported released in 2007. **The source manuscript corrects a common miscitation here and the correction is carried forward: that second trust was settled in one Channel Island in 1991 and its administration later moved to an Indian Ocean jurisdiction; it was not established under the 1984 statute, and any account saying otherwise is unreliable on its remaining detail.**

The structural lesson generalises far beyond trusts, which is why this track belongs in a chapter on jurisdictional erosion at all. **The offshore statute governs the offshore asset; the domestic court governs the person, and the person is the one part of the structure that cannot be redomiciled.** An order directed at a principal's own conduct requires no foreign court to cooperate with anything. The vehicle was never the target.

There is a legislative postscript: a ten-year look-back for transfers to self-settled trusts made with intent to hinder, delay or defraud was added to one bankruptcy code in 2005 — longer than most planning horizons and far longer than the limitation periods the protective statutes provided. **The domestic reach-back was simply extended past the foreign protection.** All of the foregoing is described from the source; no case, statute or reporter citation appears in the bibliography (Appendix B, items 1, 2 and 5).

4.5 Track five: from exchange to publication, and the private layer

A fifth track runs underneath the other four and produced no single dramatic event. Over the same period the default disposition of ownership information shifted twice: from confidential to exchanged, and then in several jurisdictions from exchanged to published.

The exchange step is the standard already named [Organisation for Economic Co-operation and Development, 2017]. The publication step rests on anti-money-laundering rather than tax instruments and required registers of beneficial ownership with extended access; one large jurisdiction had moved earlier and independently and later extended the principle to foreign owners of domestic property. **The two steps have different legal trajectories and confusing them produces bad planning.** A supranational court curtailed the public-access limb in November 2022 while leaving the reporting and competent-authority-access limbs entirely intact. The direction of travel on exchange between authorities has been one way for fifteen years; the direction of travel on publication to the general public is contested and has partially reversed. **A structure built on the assumption that its beneficial ownership will not be known to authorities is betting against the settled trend. A structure that simply prefers its affairs not to be in a newspaper is on much firmer ground**, and Section 7.3 draws that line precisely.

Beneath both sits a third layer which is private, unappealable, and in practice bites first. Correspondent banks make their own risk classifications and consume the outputs of the other four components as inputs. There is no statute to read and no tribunal to appeal to. The number of active correspondent relationships is reported to have declined materially since 2011 — the figure commonly cited is a fall of roughly a fifth over the decade — while transaction volumes rose. **This private layer transmits public decisions with no process at all**, and a jurisdiction can be listed in one month and find its institutions' clearing constrained in the next, months before any legislature anywhere has acted on the listing.

4.6 Decomposing the erosion rate

Because δ aggregates four unrelated processes it is worth decomposing before it is used.

component of δ	annual rate	basis
tax drag	0.9%	leakage at entity and distribution level after treaty relief
political / arbitrary-action risk	0.4%	probability-weighted cost of immobilisation or de-risking
debasement	0.5%	loss of purchasing power in the unit of account, net of the portfolio's own hedge
haven degradation	0.2%	loss of value in the jurisdictional characteristics the structure was built to use
δ **(sum)**	**2.0%**	

CONSTRUCTED. Illustrative decomposition, not a research estimate.

Two things follow. The first is that the components respond to entirely different interventions. Tax drag responds to structure and treaty position; debasement to the currency composition of the holding; political risk and haven degradation to jurisdictional diversity and to nothing else. **Those last two are the components the diversity term actually addresses, and they are the smaller two — 0.6 of the 2.0.** Diversification is not a general-purpose remedy; it is a specific remedy for a specific pair of terms, and a structure that diversifies while ignoring the other two has optimised the wrong thing.

The second is that haven degradation is the hardest to observe and the easiest to underestimate, because it accrues as a change in the option set rather than as a charge. Nothing is debited when a jurisdiction's differentiation erodes; the holder discovers the loss later, at the moment he wants to use a characteristic that no longer exists.

What the rate does over time is worth stating numerically, because the intuition is bad.

δ	retained at 15 y	retained at 30 y	half-life (years)
0.020	73.857%	54.548%	34.310
0.035	58.602%	34.342%	19.456
0.050	46.329%	21.464%	13.513

A difference of three percentage points in δ — well within the range of disagreement between two competent advisers about political risk alone — is the difference between a structure that mostly survives a generation and one that mostly does not. This is the multiplicative-process reading, and the general form is not the corpus's: Peters [2019] and Peters and Gell-Mann [2016] establish that the time average of a multiplicative process differs from its ensemble average, so an expectation computed across many holders does not describe what happens to any one holder over time. **The same literature supplies an independent warrant for the logarithm in the surplus law: the growth rate of a multiplicative process is logarithmic in the stock, with no psychological assumption required.** Section 3.1 justifies the logarithm on the diminishing-returns argument the framework already makes; this is the stronger warrant, and where the two disagree the multiplicative one should be preferred because it does not depend on a perceptual story being true.

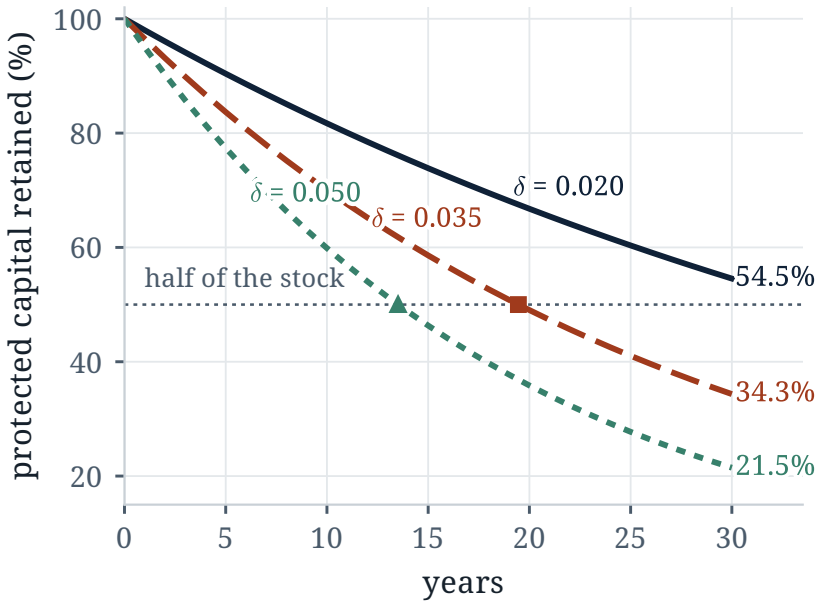


Figure 51.1: **Three percentage points of disagreement about the erosion rate is the difference between a structure that mostly survives a generation and one that mostly does not: 54.5 per cent retained at thirty years against 21.5 per cent.** Modelled from $Q_t = (1-\delta)^t Q_0$ with no input, at the three rates tabulated above; half-lives $\ln 2 / \ln(1-\delta)^{-1}$. Script `figs/src/CCP-051_fig01.py`.

4.7 Degradation trajectory as a measured quantity

The practical output of the evidence base is a discipline: assess a jurisdiction on its direction of travel with the evidence attached, not on its current state.

```
DEGRADATION TRAJECTORY -- the recorded item
-----------------------------------------------------------------
Classification   improving / stable / eroding
Evidence         at least two dated events in the last ten years,
                 each named
Direction        is the jurisdiction being brought further inside a
                 coordination perimeter, held constant, or
                 differentiated?
Reversibility    did any past change get reversed, and by what
                 mechanism?
```

Four rules govern the assessment and each is drawn directly from a track.

Delisting is not reversal. A jurisdiction comes off a list by complying — by legislating substance requirements and building supervisory capacity it did not previously have. The jurisdiction that came off is more tightly integrated into the perimeter than the one that went on. Delisting records the completion of the change, not its undoing.

The protective statute usually survives and stops mattering. Both statutes named in Sections 4.1 and 4.4 are still law. Neither was repealed and neither now delivers what it was bought for, because the operative question moved to a place the statute does not reach. Reading the statute tells you almost nothing about the current value of the protection.

Age is not durability. An institution trading since 1741 was gone in a year; another operating since 1977 did not survive its disclosure. Institutional longevity in this sector reflects the persistence of a favourable environment, and when the environment changes longevity turns out to have been a lagging indicator rather than a structural property. Weick [1993] is

the organisational form of the same finding: the structure that held for years dissolves in the minutes during which it is actually tested.

The assessment must name the mechanism, not the event. "Jurisdiction X was listed in 2020" is an anecdote. "Jurisdiction X was listed because a substance criterion was applied to zero-rate jurisdictions, a criterion which continues to apply and now also reaches the substance regimes of jurisdictions that have already legislated" is an assessment, because it says what would have to happen for the situation to change again. Section 5 names the mechanisms.

5. The coordination mechanism

5.1 Five interfaces, one system

Consider the erosion tracks as an engineer would. Four different instruments, four different legal families, four different institutional homes, four different stated purposes. On any naive account these are four independent risks, and a structure exposed to all four is exposed to four things that will not happen at once. **They happened at once.**

The explanation is not subtle. The anti-money-laundering standard-setter, the tax transparency peer-review body and its automatic-exchange standard, the coordinated listing process, the unilateral extraterritorial withholding regime, and the private correspondent-banking de-risking layer **are not five risks; they are one system with five interfaces.** Each does a different job, has a different legal character, and — the part that matters — has a different enforcement lever.

The standard-setter. Issues nothing binding. It publishes recommendations [Financial Action Task Force, 2012], conducts mutual evaluations through regional bodies, and maintains public lists. The lever is indirect and extremely powerful: financial institutions everywhere apply risk-based diligence, supervisors expect them to weight jurisdictional risk, and a listing therefore changes the cost and availability of banking for every institution connected to the listed place. **No law needs to be passed anywhere for the effect to arrive.**

The peer-review body and its exchange standard. Reviews jurisdictions on exchange of information on request and on automatic exchange [Organisation for Economic Co-operation and Development, 2017], assigns ratings, and conducts supplementary reviews. The lever is a rating, and a rating looks like a technical artefact until one notices what consumes it.

The coordinated listing process. Maintains a list of non-cooperative jurisdictions with criteria covering transparency, fair taxation and the base-erosion minimum standards [Organisation for Economic Co-operation and Development, 2015]. **Among its transparency criteria is the peer-review body's own rating. That is the coupling stated in the primary materials themselves: one institution's rating is a listing criterion for another institution's list.** The lever is a set of coordinated defensive measures, alongside the reputational consequence, which usually arrives first. de Goede [2011] treats the listing apparatus as a technique of governance in its own right.

The unilateral withholding regime. Extraterritorial and enforced by withholding on source payments to non-compliant institutions. The lever is the most direct of the five and the only one requiring no cooperation whatever from the affected jurisdiction: an institution that will not report loses a fixed fraction of its income at source. It is also the historical seed of the apparatus, since the multilateral exchange standard is recognisably its descendant.

Correspondent-banking de-risking. Private, with no statute, no tribunal and no appeal. Correspondent banks make commercial decisions about which relationships to keep, and those decisions consume the outputs of the other four as inputs. **This is the interface that converts an assessment into a consequence, and it operates on a commercial timescale rather than a legislative one.**

A sixth component is newer and closes a different door. A global minimum effective rate for in-scope groups eliminates rate differentiation rather than opacity: for a group above the

turnover threshold a zero-rate jurisdiction stops being a zero-rate jurisdiction, because the shortfall is collected elsewhere in the group. Scope, thresholds and implementing law are jurisdiction-specific and are not cited here (Appendix B, item 6).

5.2 The coupling, traced

Description is not proof of coupling. The proof is a traced propagation, and the four-month episode of Section 4.2 provides a clean one, short enough to see end to end.

step	actor	action	date
1	peer-review body	supplementary review; rating on exchange of information on request downgraded	2023
2	listing committee	downgrade engages the transparency listing criterion	2023
3	listing authority	jurisdiction added to the list	Oct 2023
4	member states	coordinated defensive measures become applicable	Oct 2023 onward
5	correspondent banks and onboarding functions	jurisdictional risk rating revised; diligence tightened	Q4 2023
6	the jurisdiction	remediation; commitments given	2023–24
7	listing authority	removed from the list	Feb 2024

Sequence as publicly reported and described from the source; review dates, rating language and decision dates are not cited to primary documents (Appendix B, item 3).

Look at what happened in between and, more importantly, at what did not. No treaty was negotiated. No court ruled. No legislature outside the territory voted on the jurisdiction at all. A peer-review panel changed a rating, and that rating change propagated through a policy committee, a Council decision, twenty-seven national tax codes and an unknown number of private bank risk models, in under a year.

Three institutions that look independent on an organisation chart transmitted a single assessment without any of them deciding to coordinate with the others, because the coupling is built into the criteria. That is what a coordination perimeter is. A holder with entities in three jurisdictions subject to that chain does not have three risks; he has one risk that will present itself three times. This is the coupled-network result of Buldyrev et al. [2010] in a governance setting — **cited as a shared structural form and not as a derivation**, because jurisdictions have not been shown to satisfy that paper's percolation premises. What the two systems share is the property that coupling destroys the independence a naive count assumes.

5.3 Why the outputs correlate

Three structural features, each independently sufficient.

Overlapping membership. The states that sit on the standard-setter, that fund and direct the peer-review body, that constitute the listing authority, and whose currencies dominate correspondent clearing are substantially the same states. A position that commands support

in one forum tends to command it in the others — not through coordination but because the same governments hold it.

Shared assessment methodology. The listing criteria explicitly incorporate the peer-review ratings and the minimum standards; supervision everywhere is built on one set of recommendations. **When several institutions grade against a shared rubric their grades correlate for the same reason two examiners marking to one rubric produce similar marks, and no communication between them is required.**

A common enforcement currency. Every one of the five interfaces ultimately expresses itself as access — to clearing, to the correspondent network, to counterparties whose own compliance functions consume these lists. That is why the lever is so effective against small jurisdictions and why it operates without legislation. It is also the feature that says what a genuinely different perimeter would require: **not a different set of laws, but a different settlement network with a different set of states controlling access to it.** Curzi [2026] measures what exists on that axis, decomposing correspondent flows by currency to find that a naive correspondent index of 0.302400 (3.306878 effective correspondents) collapses to a perimeter index of 0.745256 and $N_{\text{rail}} = 1.341819$. **The rail axis collapses harder than the jurisdiction axis**, and that is the honest answer to the question of whether a different network is available.

A coordination perimeter is not a conspiracy, and reading it as one will make predictions wrong. No committee decides that a jurisdiction should be squeezed. The institutions have genuine and largely defensible purposes and are staffed by people pursuing them in good faith. The coupling is structural, not intentional — which is precisely why it is reliable, and why it is more dangerous to a structure than a conspiracy would be: a conspiracy can change its mind and a structural coupling cannot. An analyst who models the perimeter as coordinated malice expects targeting, expects negotiation to be possible, expects political change to reverse the direction, and expects individual relationships to matter. An analyst who models it as structural coupling expects uniformity, expects compliance to work, expects the direction to persist across changes of government, and expects relationships to be irrelevant. **On the evidence of Section 4 the second model has been right**, including in the specific case of the jurisdiction that complied and got off the list.

5.4 The counter-test: where the coupling has not held

A model that explains everything explains nothing, and the coupling is not total. Three episodes run the other way, and the places where it failed say what the coupling is actually made of.

Rights adjudication inside the perimeter. A supranational court held invalid, in November 2022, the provision giving the general public access to beneficial-ownership registers, on fundamental-rights grounds. This was a genuine reversal imposed on the perimeter from inside it. It did not touch the reporting obligation and it did not touch competent-authority access — the exchange layer was entirely unaffected — but it stopped one limb of the expansion and has not been undone. **The perimeter contains institutions with their own mandates, and rights adjudication is one of them.**

Domestic constitutional constraint. A beneficial-ownership reporting statute in one large jurisdiction was held unconstitutional as applied to the plaintiffs in a district court decision of March 2024, and the regime's scope was subsequently narrowed by rulemaking. **Its current scope is genuinely unsettled and actively litigated; this paper states no position on it.** The lesson is that domestic constitutional constraint is a real brake operating on a litigation timescale rather than a policy one.

Adjudicative correction arriving too late to matter. The acquittals of Section 4.3 came years after the firm was destroyed by listing, de-risking and the withdrawal of counterparties. **The perimeter's consequences arrive on the administrative timescale and the adjudicative correction, if it comes at all, arrives on the judicial one — and nothing in a structure survives the interval.** A holder who plans to be vindicated in year eight must still fund and operate through years one to seven.

Taken together these support a refinement rather than a retraction. The coupling is strong on exchange between authorities, on standard-setting, and on the private access layer, and has not been reversed on any of them. It is weaker on publication to the general public, where rights adjudication has bitten, and weaker again on domestic constitutional questions inside large jurisdictions with strong courts. **The trend a structure must plan against is the exchange trend, because that is the one with no counter-examples.**

A fourth item belongs here as a limit rather than a counter-example. The perimeter operates on jurisdictions and institutions; it does not operate on legal character. A reinsurance reserve held against real third-party liabilities is inside the perimeter on all four questions and is nonetheless awkward to reach, because reaching it impairs identifiable third parties with their own claims and their own regulator. **Perimeter diversity and character diversity are complements, not substitutes, and a structure that has one and not the other is only half built.** Curzi [2026] prices the character axis on the cost side, measuring substance repricing as a multiple — $5{,}350/1{,}550 = 3.45\times$ for a pure equity-holding entity and $19{,}350/1{,}550 = 12.48\times$ for a relevant-activity entity, with break-even at USD 7,740,000 of held value at 25 basis points.

5.5 The bloc taxonomy

Applying the test produces a small number of groupings. Five are used, with the caveat that the boundaries are analytical conveniences rather than natural kinds and that a jurisdiction's classification is a judgement to be re-made rather than a fact to be looked up.

bloc	character	examples
WEST	four yeses; full attachment on every interface	UK, US, Jersey, Luxembourg
W-LITE	four yeses, but with a distinct legal product — typically genuine risk transfer — that is hard to displace	Bermuda
BRIDGE	full compliance posture with genuinely dual-attached settlement and a foreign policy not aligned to one bloc	UAE, Singapore
NON-ALIGNED	compliant, with an independent foreign policy, domestic-currency depth, and real operating substance	Thailand, India
EAST	primary settlement and policy alignment outside the Western perimeter	treated on the rail axis

Classifications are the analyst's judgements, are contested, and change. They are not cited to primary sources (Appendix B, item 4).

The W-LITE category is doing real work. The jurisdiction placed there is inside the Western perimeter on all four questions; what makes it a distinct holding is not perimeter position but product. **That is legal-character diversification rather than perimeter diversification, and it is counted as such and no more.** Folding it into WEST is the conservative assignment and is the one used in the measurement below.

Treating a bridge as neutral is the error this taxonomy most invites. A bridge jurisdiction is not neutral; it is *dual-attached*, which is a different and much more useful property. The Gulf bridge central to the architecture participates in automatic exchange, clears in the principal reserve currencies through correspondents subject to their issuers' jurisdiction, and was itself under increased monitoring for twenty-three months. On the four-question test it scores yes

four times. **What makes it a bridge is not that any attachment is absent. It is that it maintains an independent foreign policy, hosts genuine commercial relationships across blocs that do not deal easily with each other, and provides legal infrastructure credible to counterparties on both sides. A bridge reduces the probability that a shock in one bloc removes access to the other. It does not place anything outside either bloc, and nothing here claims that it does.** The classification is a claim about correlation, not about immunity.

The corresponding error at the other end is assuming that non-Western means unreachable. A jurisdiction outside the Western political sphere is not thereby outside the perimeter, because Q4 is answered by the plumbing rather than by the politics.

6. Results: correlated failure, measured

6.1 The arithmetic of correlated failure

The reason perimeter membership must be measured rather than described is that the cost of getting it wrong is not proportional. Take a constructed structure with four jurisdictions, each facing an annual probability of 0.10 of an adverse regime change material to the holding.

```
MODEL                  P(ALL FOUR IMPAIRED IN ONE YEAR)
-------------------------------------------------------
Independent            p^4 = 0.000100       1 in 10,000
One perimeter          p   = 0.100000       1 in 10
-------------------------------------------------------
Understatement factor from assuming independence: 1,000x

Cumulative probability of total perimeter loss:
    over  5 years   independent 0.000500   correlated 0.409510
    over 10 years   independent 0.001000   correlated 0.651322
    over 15 years   independent 0.001499   correlated 0.794109
-------------------------------------------------------
CONSTRUCTED probabilities, chosen for illustration.
```

Assuming independence where a perimeter exists understates the probability of total impairment by three orders of magnitude, and over fifteen years it is the difference between a one-in-seven-hundred event and a four-in-five event. No other modelling error in this literature is worth that much.

The general case is more useful than the extreme, because real structures are neither fully independent nor fully coupled. Decompose each jurisdiction's failure into a common shock arriving with probability c and an idiosyncratic shock arriving with probability i, holding the marginal probability fixed at 0.10 so that the comparison is like for like. Then $P(\text{all four}) = c + (1-c)i^4$.

c	i	$P(\text{all four})$	understatement vs p^4
0.00	0.1000	0.000100	$1.0\times$
0.02	0.0816	0.020044	$200.4\times$
0.05	0.0526	0.050007	$500.1\times$
0.08	0.0217	0.080000	$800.0\times$
0.10	0.0000	0.100000	$1{,}000.0\times$

Read the second row. The joint-failure probability is governed almost entirely by the common component and almost not at all by the idiosyncratic one — at $c = 0.02$ the idiosyncratic term contributes $0.98 \times 0.0816^4 = 0.0000435$, or 0.22 per cent of the total. **H2's prediction is met**

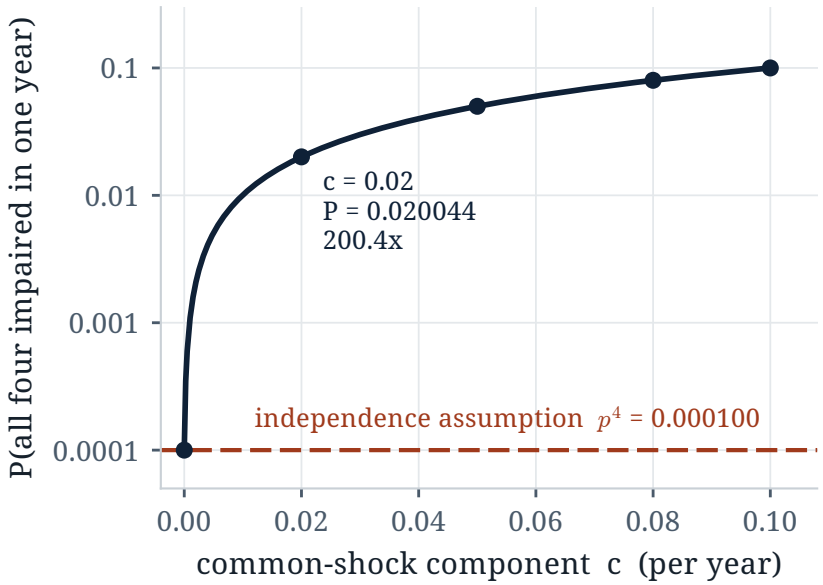


Figure 51.2: **A common-shock component of two per cent a year — smaller than the analyst's uncertainty about any input — already understates joint failure by a factor of 200.4, and the joint probability tracks the common component almost exactly.** Modelled from $P(\text{all four}) = c + (1-c)i^4$ with the marginal failure probability held at 0.10; parameters CONSTRUCTED. Script `figs/src/CCP-051_fig02.py`.

exactly. That is why the perimeter is the object worth measuring and the individual jurisdiction is not.

Two pieces of external literature bound the reading. Longin and Solnik [2001] and Ang and Chen [2002] measure correlation rising under stress in equity markets, asymmetrically and concentrated in downturns: **the effective count falls exactly when it is needed**, which is stronger than anything the static index says. The ecological name for the same variable is synchrony [Loreau and de Mazancourt, 2008], and the cleanest empirical demonstration that effective count buys stability is Schindler et al. [2010]'s sockeye record, where population diversity produces roughly half the variability and about a tenth of the fishery closures. Against them, Doak et al. [1998] show that diversity–stability correlations can arise from statistical averaging alone with no mechanism, and May [1972] showed that increasing diversity and connectance *destabilises* randomly assembled systems. **This paper's claim is narrower than the general diversity–stability claim and is not exposed to those results in the same way**: it asserts that the reciprocal of a variance counts independent shocks, which is arithmetic, and that perimeter-attached jurisdictions are correlated, which is H1.

6.2 The worked measurement

CONSTRUCTED. The principal, the entities and all figures are inventions built to illustrate the measurement. No real person or transaction is described.

A principal — dual national, tax resident in a Gulf jurisdiction, with three adult heirs resident in three different countries, one of them under a forced-heirship regime — arrives with a structure assembled by competent advisers over about twenty years. It holds commercial real estate in London, a listed-securities portfolio, and interests in operating businesses.

The share basis is stated first, because a share is a share *of something* and the something must match the threat: *the fraction of protected capital whose availability depends on a decision taken inside a given jurisdiction's coordination perimeter.* The London real estate contributes its full value to the UK exposure because it cannot be relocated. The listed portfolio contributes to the perimeter of its *custodian*, not to those of the issuers.

entity / situs	share of Q	role	Q1	Q2	Q3	Q4	score
United Kingdom	40.9091%	property holding, principal operating presence	yes	yes	yes	yes	4/4
United States	27.2727%	securities custody and brokerage	yes	yes	yes	yes	4/4
Jersey	18.1818%	family trust, succession	yes	yes	yes	yes	4/4
Luxembourg	13.6364%	holding and treaty access	yes	yes	yes	yes	4/4
total	**100.0000%**						

Every entity scores four. On Q1 all four participate in automatic exchange. On Q2 all four implement the common recommendations through their own supervisors. On Q3 all four are reachable by a listing process whose consequences their governments could not unilaterally prevent. On Q4 all four depend on correspondents subject to the issuing states' jurisdiction, including the entity in the issuing state itself, which is subject directly.

```
RAW MEASUREMENT (what the adviser presented)
   naive HHI                    = 0.293388
   effective jurisdictions      = 3.4085

COLLAPSED MEASUREMENT (the honest figure)
   WESTERN PERIMETER            = 100.0000%
   bloc-HHI                     = 1.000000
   EFFECTIVE PERIMETERS  N      = 1.0000

THE GAP                         = 3.41x overstatement
```

Four entities, four jurisdictions, two continents, and one point of failure. The naive measure overstated the structure's resilience by a factor of 3.41, and it did so not through arithmetic error but through a modelling assumption — independence — that the record in Section 4 falsifies. **H1's prediction is met.** Note what the measurement does *not* say. It does not say the four entities were badly chosen or badly run. **The finding is not about quality; it is about correlation, and correlation is not visible from inside any one entity.**

The surplus consequence at the stated parameters:

```
h = eta ln(N) - delta Q - C
  = 0.35 x ln(1.0000) - 0.02 x 10 - 0.30
  = 0.35 x 0.000000 - 0.200000 - 0.300000
  = -0.500000
```

The available surplus is exactly minus one half: the whole of the erosion plus the whole of the upkeep, with nothing at all on the input side. This is not a marginal result that better execution could improve. The input term is multiplied by the logarithm of one regardless of how well the entities are run or how the shares are redistributed among them. Moving the Luxembourg allocation to Jersey, or splitting the UK position across three companies, changes the naive figure and leaves the collapsed figure at exactly 1.0000.

The designed alternative distributes the same capital across six jurisdictions in three perimeters:

entity / situs	share	bloc	role
UAE	29.0909%	BRIDGE	holding
Thailand	18.1818%	NON-ALIGNED	operating

entity / situs	share	bloc	role
United Kingdom	16.3636%	WEST	operating, contracts
Bermuda	13.6364%	W-LITE → WEST	risk transfer
India	11.8182%	NON-ALIGNED	fund vehicle
Singapore	10.9091%	BRIDGE	bridging, banking
total	**100.0000%**		

```
COLLAPSED MEASUREMENT
   BRIDGE                      = 40.0000%      0.160000
   NON-ALIGNED                 = 30.0000%      0.090000
   WEST (incl. W-LITE)         = 30.0000%      0.090000
   bloc-HHI                    = 0.340000
   EFFECTIVE PERIMETERS  N     = 2.941176

h = 0.35 x ln(2.941176) - 0.200000 - 0.300000
  = 0.35 x 1.0788095014 - 0.500000
  = 0.377583 - 0.500000
  = -0.122417
```

Curzi [2026] states the naive side of the same structure: a naive index of 0.189002 giving 5.290949 effective jurisdictions, against a bloc index of 0.340000 and 2.941176 effective blocs, with Western-perimeter exposure falling from 100.0 to 30.0 per cent.

Note which ratio is the real one. The naive count improves by a factor of roughly 1.55, which looks like a modest gain. The perimeter count improves by a factor of 2.9412, and that is the gain that enters the surplus equation. **The measurement taken at the wrong level of aggregation understated the problem before the restructuring and understates the improvement after it, in both cases by roughly the same factor and for the same reason.**

6.3 The fifth decimal: a reconciliation this paper prints rather than harmonises

This paper and Curzi [2026] report different naive indices for the same entity map, and the difference is printed here rather than smoothed away. It changes no conclusion, and two corpus papers silently disagreeing on a stated number is exactly what a reviewer notices.

The shares in Section 6.2 are exact twenty-seconds of the total, printed to four significant figures. Three bases are therefore available and they give three different naive indices.

basis for the conventional structure	shares	naive HHI	effective jurisdictions
exact fractions	9/22, 6/22, 4/22, 3/22	0.293388	3.408451
percentages to four significant figures	40.91, 27.27, 18.18, 13.64	0.293384	3.408499
percentages to three significant figures	40.9, 27.3, 18.2, 13.6	0.293430	3.407968

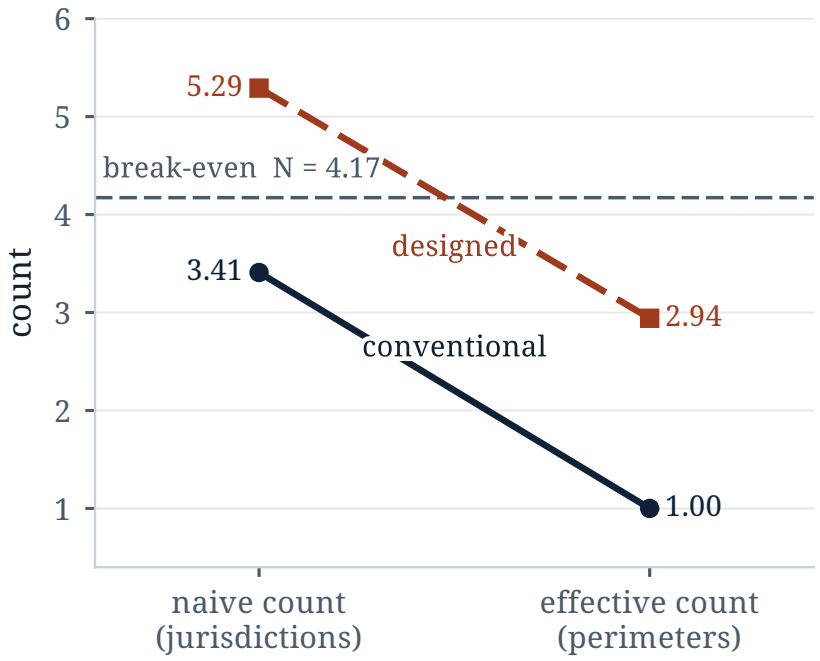


Figure 51.3: **The gap is the finding: the conventional structure reads 3.41 jurisdictions naively and exactly 1.00 perimeters in effect, the designed structure 5.29 and 2.94, and neither effective count reaches the break-even of 4.17 that a non-negative surplus balance requires.** Measured from the share vectors tabulated above; naive figures on the exact-fraction basis of Section 6.3 for the conventional structure and from Curzi [2026] for the designed one; break-even derived at $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$. Script `figs/src/CCP-051_fig03.py`.

basis for the designed structure	shares	naive HHI	effective jurisdictions
exact fractions	6.4/22, 4/22, 3.6/22, 3/22, 2.6/22, 2.4/22	0.188926	5.293088
percentages to three significant figures	29.1, 18.2, 16.4, 13.6, 11.8, 10.9	0.189002	5.290949

This paper's R1 states 0.293388 and 3.4085, which is the exact-fraction value. Curzi [2026] states 0.293430 and 3.407968, which is the three-significant-figure value, and the same paper records the discrepancy in its own discussion. The two papers agree about which figure comes from which basis. **The three-significant-figure percentages are the reproducible choice given published inputs, because they are what a reader can recompute from a printed exposure table without being told the exact fractions**, and that is why CCP-052 adopts them and why any reproduction of this measurement should adopt them too. The same applies to the designed structure, where 0.189002 is the reproducible figure and 0.188926 is the exact-fraction one. **The discrepancy on the designed structure is not disclosed anywhere else in the corpus and is disclosed here.**

Nothing turns on the choice. Every bloc figure is exactly 1.000000 for the conventional structure and exactly 0.340000 for the designed one under every basis, because the perimeter sums are 100.0% and 40/30/30 exactly in all three cases. The naive figures do not enter the surplus law, and the difference in the fifth decimal propagates to no reported result.

The companion discipline is that intermediate values must not be rounded. The designed structure's effective count is exactly $1/0.34 = 2.941176\ldots$:

```
exact:    ln(2.941176) = 1.078810   h = 0.35 x 1.078810 - 0.5 = -0.122417
rounded:  ln(2.940000) = 1.078410   h = 0.35 x 1.078410 - 0.5 = -0.122557
```

The two differ by 0.000140. **The conclusion is unchanged — both are below zero — but the number moved in the fourth decimal from a single rounding, and a chain of four such steps is how a measurement goes wrong.** The rule is not to round carefully; it is to carry full precision through every intermediate step and round only at the print.

6.4 What it costs to change N

Suppose a structure is entirely inside one perimeter and the holder begins moving exposure out, splitting what leaves evenly between two other perimeters.

share remaining in the perimeter	bloc-HHI	effective perimeters	$\eta \ln N$ at $\eta = 0.35$
100.00%	1.000000	1.0000	0.000000
95.00%	0.903750	1.1065	0.035421
90.00%	0.815000	1.2270	0.071599
85.00%	0.733750	1.3629	0.108355
80.00%	0.660000	1.5152	0.145430
66.67%	0.500000	2.0000	0.242602
60.00%	0.440000	2.2727	0.287343
50.00%	0.375000	2.6667	0.343290
40.00%	0.340000	2.9412	0.377583
33.33%	0.333333	3.0000	0.384514

The solution for exactly two effective perimeters is analytic. With the remainder split evenly, an in-perimeter share s gives HHI $= s^2 + 2((1-s)/2)^2$, and setting this to $1/2$ yields $s = 2/3$

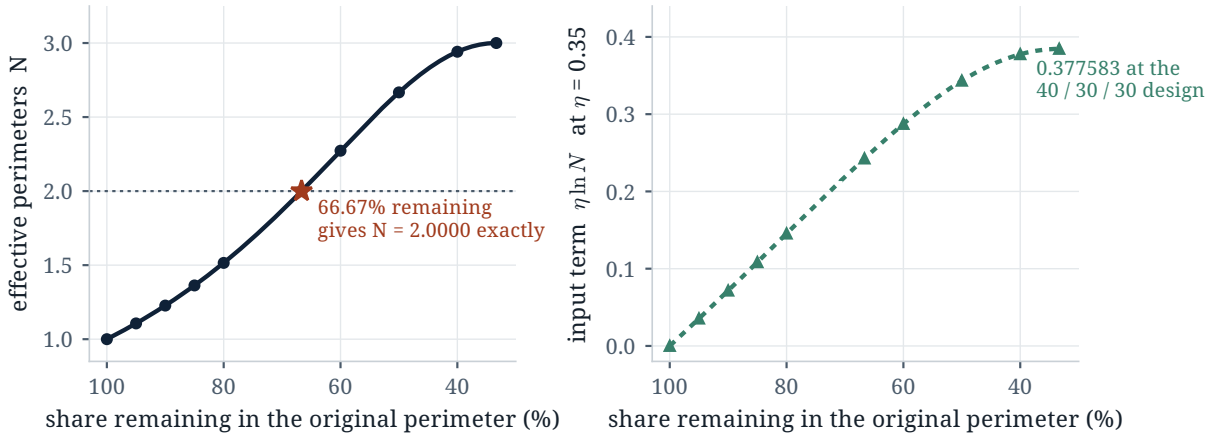


Figure 51.4: **One third of protected capital must leave the perimeter before the effective count reaches two, and the surplus term rises with the logarithm of a count that itself responds slowly to reallocation.** Modelled from HHI $= s^2 + 2((1-s)/2)^2$ with the exiting share split evenly between two further perimeters; $\eta = 0.35$. Script `figs/src/CCP-051_fig04.py`.

exactly. **One third of protected capital has to sit outside the perimeter before the structure has two effective perimeters rather than one.** Ten per cent moved buys 1.2270; twenty per cent buys 1.5152. **This is not a marginal adjustment that can be made by opening an account somewhere; it is a structural reallocation. H4's prediction is met.**

Note that the designed structure of Section 6.2 is exactly the 40 per cent row: three perimeters at 40/30/30 is what the even-split path reaches when 60 per cent of exposure has left the original perimeter.

6.5 The over-diversification optimum

The opposite error is real. The input term grows as the logarithm of the perimeter count while maintenance grows roughly linearly with it — every additional perimeter means additional counsel, filings, local substance and management attention. Logarithmic benefit against linear cost has an interior optimum, and the optimum is not large. This is the organisational result in a different register: adding units inside a coordination boundary adds communication cost at order $n(n-1)/2$ [Brooks, 1975] and adds nothing to independence.

Write maintenance as $C = c_0 + c_1 N$. Then $h = \eta \ln N - \delta Q - c_0 - c_1 N$ is maximised where $\eta / N = c_1$, that is at $N^* = \eta / c_1$. At $\eta = 0.35$:

per-perimeter cost c_1	optimal N^*	interpretation
0.05	7.0000	cheap coordination
0.07	5.0000	typical mid-size group
0.10	3.5000	expensive coordination

And the marginal value of each additional perimeter, which is the number that should govern the decision to add one:

N	$\eta \ln N$	marginal gain over $N-1$
1	0.000000	0.000000
2	0.242602	0.242602

N	$\eta \ln N$	marginal gain over $N-1$
3	0.384514	0.141913
4	0.485203	0.100689
5	0.563303	0.078100
6	0.627116	0.063813
7	0.681069	0.053953
8	0.727805	0.046736

The second perimeter is worth $0.242602/0.046736 = 5.19$ times what the eighth is worth. Almost all of the available benefit is captured by the move from one to two and from two to three; beyond that the holder is paying linear cost for logarithmic return.

A monotone raise-N prescription cannot survive the external literature unqualified, and none is made here. Uzzi and Spiro [2005] find an inverted U in which performance rises and then falls with cohesion; Sagan [1993] documents that redundancy can *reduce* safety when it adds coordination surface; and March [1991] gives the mechanism by which fast convergence on a single mode destroys the variance that produced the advantage. The prescription here is bounded: move from one effective perimeter to two or three, and stop when the marginal entity stops adding a perimeter.

7. Results: the economics of full disclosure

Sections 4 to 6 establish that concealment has been removed as a durable property and that the replacement is perimeter diversity. This section tests the standard objection to that replacement, which is that it costs more.

7.1 Why concealment fails: four structural properties

The failure of concealment structures is not bad luck and not a consequence of poor execution. It follows from four properties of the environment, each sufficient on its own.

The observer set is monotone. Information that has entered the exchange system does not leave it [Organisation for Economic Co-operation and Development, 2017]. A structure that was invisible in 2013 and reported in 2019 is not invisible again in 2025; the 2019 report still exists. The property holds with greater force on public distributed ledgers, where every transfer is permanently readable and analytical capability improves while the record does not degrade — **a transfer made today will be readable in 2035 by tools that do not yet exist, which makes public-chain settlement the most durable surveillance record a holder can create, not the least.** Concealment is therefore a wasting asset in a technical sense: its value can only fall, and each year's disclosure permanently reduces the stock.

The asymmetry of repeated trials. A concealment structure must succeed at every examination; an examiner must succeed once.

annual examination probability	$P(\geq 1)$ in 5 y	10 y	15 y	20 y
0.05	0.226220	0.401263	0.536709	0.641514
0.10	0.409510	0.651322	0.794109	0.878423

A structure meant to hold for a generation and requiring an unbroken run of non-examination is, at ten per cent annual exposure, approximately eighty per cent likely to be

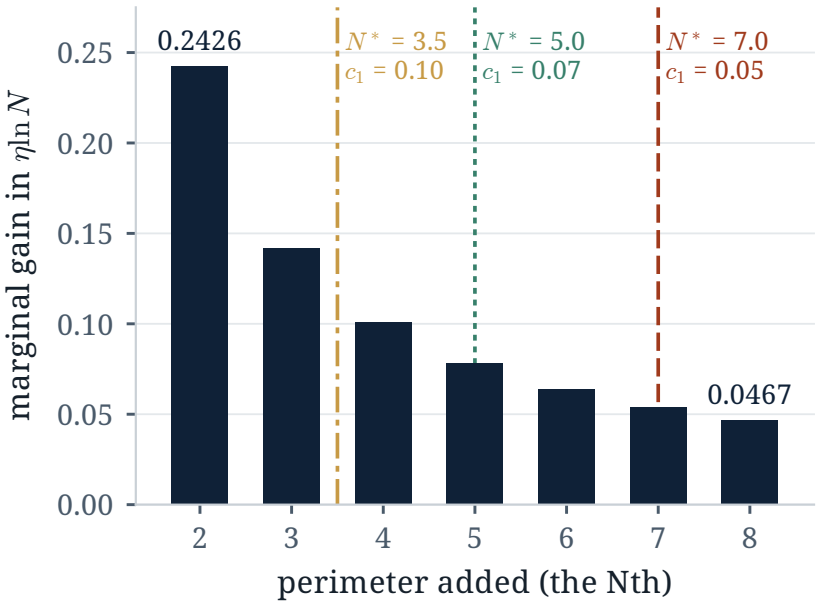


Figure 51.5: **The second perimeter is worth 5.19 times the eighth, and at expensive coordination the interior optimum falls to three and a half perimeters.** Modelled from $h = \eta \ln N - \delta Q - c_0 - c_1 N$ at $\eta = 0.35$, with $N^* = \eta / c_1$ marked for the three tabulated coordination costs. Script `figs/src/CCP-051_fig05.py`.

examined within fifteen years. Structures are frequently built to horizons of thirty years and longer.

The principal is inside a jurisdiction. This is the mechanism of Section 4.4 and it generalises. Assets can be redomiciled and entities migrated; the person who built the structure is subject to the coercive power of the courts where he stands.

The cost inversion. Concealment is not cheaper than compliance once the expected cost of failure is included, and at realistic parameters it is not cheaper even before that. The next section computes it.

7.2 The economics, computed

CONSTRUCTED. The figures are built for illustration. They are not quotations, they are not research estimates, and actual costs vary by an order of magnitude with size and complexity. What is demonstrated is the *shape* of the comparison, and the shape is robust to large changes in the inputs.

A compliant multi-jurisdiction architecture holding capital of USD 50 million:

annual maintenance line	USD / year
Gulf foundation: registered agent, office, filings	18,000
Gulf foundation: council fees (three members)	45,000
Gulf holdco: licence, registered office, accounts	22,000
non-aligned operating company: audit, incentive reporting, payroll compliance	36,000
Atlantic captive: regulatory filings, actuarial opinion, audit	140,000

annual maintenance line	USD / year
second non-aligned centre vehicle: regulatory compliance, fund administration	85,000
Western company: accounts, ownership filings, corporation tax	14,000
group: coordinated tax counsel, six jurisdictions	120,000
group: exchange classification and annual reporting	26,000
group: transfer-pricing documentation	55,000
total annual C	**561,000**

As a fraction of protected capital: 224.40 basis points at USD 25m, 112.20 at 50m, 56.10 at 100m, and 22.44 at 250m. **At USD 25 million this architecture costs 224 basis points a year, which is a serious drag; at USD 250 million it costs 22, which is negligible. Compliance-based architecture has substantial fixed costs and therefore a minimum efficient scale.** Below that scale the correct advice is a simpler structure, not a cheaper version of a complex one.

Now the comparison. Assume a concealment-based alternative costing USD 210,000 a year. It is genuinely cheaper on the visible line: USD 351,000 less, a saving of 62.6 per cent. Assume that on examination it loses sixty per cent of capital to assessment and recovery, incurs a penalty of twenty per cent, and consumes USD 1.5 million in counsel, so that $L =$ USD 41.5 million.

annual exam probability	expected annual loss	concealment total burden	compliant C	ratio
0.02	830,000	1,040,000	561,000	$1.85\times$
0.05	2,075,000	2,285,000	561,000	$4.07\times$
0.10	4,150,000	4,360,000	561,000	$7.77\times$
0.20	8,300,000	8,510,000	561,000	$15.17\times$

```
p* = (561,000 − 210,000) / 41,500,000 = 0.008458
```

At any annual examination probability above 0.85 per cent the concealment structure is more expensive than the compliant one, and 0.85 per cent is roughly a one-in-a-hundred-and-eighteen-year event, far below any defensible estimate of examination risk in the environment of Section 4. The apparent 62.6 per cent saving is an artefact of counting the invoice and not the exposure. The result is insensitive to L: halving it moves p^* only to 0.016916. **H3's prediction is met.**

There is a second-order point the table does not show and that matters more than any number in it. The compliant structure's cost is a known, budgetable operating expense paid in cash from income. The concealment structure's cost is a low visible fee plus an unbudgetable tail risk to principal. **Those are not the same kind of liability, and a principal who treats them as comparable line items has mispriced the decision before the arithmetic starts.** Cyert and March [1963] is the general form: organisational slack — the property that makes a system survive a shock it did not plan for — has always appeared on the accounts as waste.

7.3 Three words that are not synonyms, and the product rule

Precision here prevents most of the confusion in this field.

Privacy is the interest in one's affairs not being available to the general public. It is a recognised right in several legal systems, it was vindicated against a supranational directive

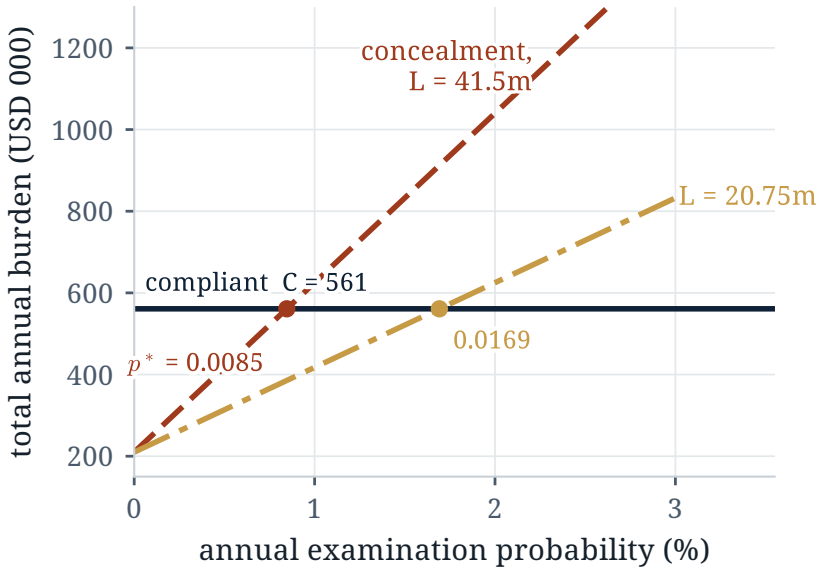


Figure 51.6: **The visible 62.6 per cent saving on the invoice reverses at an annual examination probability of 0.85 per cent, and halving the loss on examination moves that threshold only to 1.7 per cent.** Modelled from a total burden of $210{,}000 + pL$ against a compliant $C = 561{,}000$, with $L =$ USD 41.5m and USD 20.75m; all parameters CONSTRUCTED. Script `figs/src/CCP-051_fig06.py`.

in November 2022, and structures may rely on it. **Confidentiality** is a duty owed by a professional or institution to a client, enforceable between them and subject to statutory override; structures may rely on it only within the space the overrides leave, which shrinks. **Opacity to competent authorities** is the condition of an authority with jurisdiction being unable to obtain information it is entitled to, and fifteen years of evidence say the trend here is one-directional with no significant counter-examples. **A structure may rely on privacy, may rely on confidentiality within its statutory limits, and may not rely on opacity.**

The operational form of the design standard is one question, applied element by element:

```
THE DISCLOSURE TEST
------------------------------------------------------------------
Would this structure still do its job if every element of it were
disclosed to every authority with jurisdiction over any part of it?
------------------------------------------------------------------
"Every element"    each entity, contract, filing, account and flow
"Every authority"  the tax authority of each operating jurisdiction,
                   the tax authority of the principal's residence,
                   each financial regulator, each correspondent
                   bank's compliance function, and the court that
                   would hear an adverse claim
"Its job"          the specific function assigned to it under the
                   three-jobs taxonomy: holding, operating, bridging
------------------------------------------------------------------
```

Five criteria, all of which must be satisfied: (1) lawful where it sits, on the facts as they actually are rather than as documented; (2) filed and registered, on time and current; (3) correctly classified and reported under every applicable exchange regime; (4) independent commercial rationale, statable without reference to any tax or disclosure consequence, and real; (5) survives hostile reading, remaining defensible when described in full in the least favourable framing to a tribunal inclined to disbelieve it.

A structure's disclosed value is the product of its elements' scores, not their average, because a chain fails at its weakest link and an examiner who finds one bad element examines the rest with an entirely different prior.

elements satisfied	average score	product score	verdict
5 of 5	1.00	1.00	PASS
4 of 5	0.80	0.00	FAIL
2 of 5	0.40	0.00	FAIL
1 of 5	0.20	0.00	FAIL

The average column is what an adviser presents when he wants to describe a structure as substantially compliant. **There is no such thing.** Criterion 4 is the one most often failed by otherwise well-built structures, and criterion 5 the one most often failed by structures satisfying the first four on paper.

The audit procedure is six steps: enumerate every entity, account, contract and flow; assign each exactly one job; score the five criteria recording the *evidence* for each yes rather than the conclusion; multiply, never average; remediate or **remove**, because documentation improves the record of a fact and cannot create one; and re-run annually and on any change of law, residence or control. **Step five is where the discipline lives. The instinct on discovering a weak element is to improve its documentation, and an examiner comparing the documentation to the operational reality will find the gap.**

7.4 Five structures scored, two of which fail

CONSTRUCTED. Structures A to E are invented; no real person, family, entity or transaction is described. Jurisdictions are de-identified where the score would otherwise assert an uncited jurisdiction-specific characteristic as fact, following the convention of Curzi [2026].

The five are arranged deliberately. A is a foundation apex. B and C are a captive insurer and a foreign-owned operating company, both built correctly. **D and E are the same two jobs in the same two jurisdictions, built the other way: the jurisdiction is identical, the entity type is identical, the cost is comparable, and the only difference is that one element depends on something not being seen.**

structure	C1	C2	C3	C4	C5	product	verdict
A — foundation apex in Gulf centre A	1	1	1	1	1	1	PASS
B — Atlantic captive writing real third-party risk	1	1	1	1	1	1	PASS
C — non-aligned operating company, incentive-promoted, disclosed	1	1	1	1	1	1	PASS
D — the same operating company, nominee-held	0	1	0	0	0	0	FAIL
E — captive writing no risk	1	1	1	0	0	0	FAIL
average of the five criterion means, offered only to be discarded						**0.80**	meaningless
product of the group						**0**	any zero fails the group

Structure A is a foundation holding the family's branch holdings as a single perpetual legal person: charter filed, by-laws held privately by the registered agent, founder's reserved powers graded so that no power to direct a distribution to himself exists, a council including an independent resident professional, meetings held and minuted, ownership particulars filed, exchange classification determined and reported. **Its function is to be the recorded owner, and a recorded owner works better when the record is examined** — four registries naming

the same legal person close questions rather than opening them. Disclosure is the mechanism, not a tolerated side effect. This is asset partitioning doing the work entity law exists to do [Hansmann and Kraakman, 2000].

Structure B is a licensed and supervised captive writing product-liability and business-interruption cover for the group's manufacturing operation and participating in a third-party quota share, with premiums set by an independent actuary, reserves held against the liabilities, and claims actually made, adjusted and paid. **The weight sits on criterion 4, satisfied by a fact rather than a document**: the group has a real, quantified, badly-priced exposure and the captive assumes it.

Structure C is a manufacturing company wholly owned by the group's Gulf holding company, with an investment-promotion instrument conferring the foreign-ownership position on its face and every shareholder the registered holder in its own right. **The company's most sensitive fact — that it is entirely foreign-owned in a country with a restrictive foreign business regime — is not merely disclosed; it is the subject of an affirmative government grant, publicly recorded. The fact that would destroy a concealment structure is the fact that makes this one work.**

Structure D is the identical business with identical premises, customers and revenue, in which no promotion was applied for and the foreign-ownership position was instead arranged by having two local individuals appear on the share register as majority holders, subject to an unregistered loan agreement, an undated blank transfer, an irrevocable proxy and a side letter. It fails criterion 1 because nominee shareholding to circumvent foreign-ownership restrictions is a criminal offence in the jurisdiction concerned and is prosecuted — **a claim carried from the source manuscript and not cited; the statutory provision and its current penalties are Appendix B, item 8.** It fails 3 because the true controlling person is not the one reported, 4 because the arrangement has no purpose statable without reference to the restriction it circumvents, and 5 because the side letter read aloud is the prosecution's exhibit. **It passes criterion 2, and that is the instructive part: the register is filed, and it is false, which makes criterion 2 the least relevant line in the table.**

What makes D the instructive failure. It is not cheap. It has counsel, a share register, audited accounts, a real factory and real employees, and on any measure of economic reality it is the same business as C. It fails on one element, **and that element converts the other four from evidence of a well-run business into evidence of a well-run concealment.** The promotion route took four months to obtain and cost less than a year of the side letter's legal fees. The concealment route was chosen because it was faster, and speed is why most failing structures exist.

Structure E is a small captive in a well-regulated jurisdiction, licensed and filing, issuing policies to affiliates covering remote and poorly defined perils, with premiums set to produce a deduction rather than by actuarial assessment, no third-party risk, no meaningful pooling, no claim ever made, and reserves invested in a note issued back to the group. **Why this failure is different from D's, and more dangerous:** D is unlawful where it sits and everyone involved knows it. E is lawful in the jurisdiction of formation on every visible measure — licensed, supervised, audited, reported. It fails on criteria 4 and 5, the two about substance rather than form, **and it fails in the jurisdiction of the policyholder, not that of the insurer. A structure can be perfectly compliant where it sits and fail the disclosure test entirely, because the test asks about every authority with jurisdiction over any part of it, and the deduction is claimed somewhere else.** The tax jurisprudence on this class of arrangement turns on risk distribution rather than on any concealment; the decisions and the current regulatory treatment are Appendix B, item 9.

Two observations on the scoreboard. **First, note where the failures cluster.** Criterion 2 is satisfied by both failing structures. Filing is the criterion concealment structures satisfy most easily, because filing a false particular is still a filing; an adviser who reports compliance by counting filings reports both failures as passes. **Second, neither failure is a close call once the criteria are applied element by element, and both were sold to their principals as sophisticated. The instrument's value is not that it detects subtle problems; it is that it makes obvious**

problems impossible to describe as acceptable.

7.5 Concealment and armour: the same effect by inverse mechanisms

There are exactly two ways to keep an asset out of a hostile party's hands: make it hard to find, or make it hard to take. Both work. **They are not points on a spectrum, not weak and strong versions of one another, and cannot be blended into a hybrid with the advantages of each.**

CONCEALMENT	STRUCTURAL ARMOUR
survives because the adversary cannot LOCATE it	survives because the adversary cannot TAKE it
the protective quantity is the adversary's IGNORANCE	the protective quantity is the structure's CORRECTNESS
protection is a decreasing function of observation	protection is invariant to observation
the defence is consumed the first time it is used	the defence is unchanged by use and is available on the tenth examination as on the first
fails DISCONTINUOUSLY: full protection until exposure, then none	fails CONTINUOUSLY: only under applied force, and by degrees that can be observed
failure is not observable in advance	failure is observable in advance
scales badly: each additional person who knows is a new failure point	scales well: each additional filing makes the next filing easier

The two mechanisms are exclusive because the actions that strengthen one weaken the other. Every disclosure that makes an armoured structure more defensible — the filed charter, the named owner, the reported account, the audited return — is a location that concealment would have had to avoid. Every layer that makes a concealed structure harder to locate is an element that fails criterion 1 or criterion 5 and therefore zeroes the product score. A structure is one design or the other, and the decision is taken at inception, usually without being recognised as a decision at all.

This is the open-design principle stated as an engineering standard rather than an ethical one. Kerckhoffs [1883] requires that a system remain secure when everything about it except the key is public; Saltzer and Schroeder [1975] restates it as *open design* among the protection principles. **The corpus's disclosure constraint is therefore not a moral preference; it is Kerckhoffs's principle applied to legal structure, and the concealment column above is precisely the failure mode Kerckhoffs identified in 1883.**

Now put the two mechanisms against the environment of Sections 4 and 5. Concealment's protective quantity is the adversary's ignorance, and the four properties of Section 7.1 describe a world in which that quantity is falling toward zero and cannot be replenished. Armour's protective quantity is correctness, which is depleted by none of those four. It is depleted only by the law changing under a structure built correctly under the old law — which is a real risk, is observable in advance, and is what criterion 3 and step six of the audit exist to catch.

Concealment-based vehicles are not obsolete because they were unethical, or because they became unfashionable, or because regulators disapproved of them. They are obsolete because the quantity they ran on stopped being available. That is a supply statement, not a moral one, and it is why the replacement is not a better-concealed vehicle. When a mechanism's input goes to zero the replacement is not a more efficient version of the same mechanism; it is a mechanism that runs on a different input.

7.6 An existence proof on the public record

Everything above is the framework's own instrument applied to the framework's own constructed cases, which is a closed loop. The argument needs one thing it cannot supply for itself: a demonstration on the public record that a function which used to require concealment is now performed by a structure with no concealment in it at all — successfully, at scale, for a long time, by an institution with every incentive to choose the other design if the other design still worked.

The load-bearing facts are only three: there exist sovereign investment institutions that are publicly acknowledged, whose purpose and principal are publicly known, and that work. The named instance in the source manuscript is a national pension fund managed by a unit of a central bank under a published ministerial mandate, which publishes its complete holdings list, publishes how it votes and how it intends to vote, publishes exclusion decisions naming the company and the reason, and reports audited accounts to its legislature. **The specifics — governing statute, mandate instrument, first transfer date, the exact scope and timing of each disclosure, the constituting instrument of the ethics body, and the fund's size — are described from the source and are not cited (Appendix B, item 10.)**

What it replaced. For much of the second half of the twentieth century a state deploying capital into another state's markets did so through arrangements designed not to be attributable: nominee registrations, custodian omnibus accounts, layered intermediate vehicles, and positions kept deliberately below the level at which a name must be filed. The object was the return and the access without the political fact of state ownership becoming visible. **The pattern is well documented; individual attributions are usually contested, and none is asserted here.**

Why the mechanism stopped being available. The four properties of Section 7.1 apply with greater force to a state's capital than to a private structure. Registries became searchable and then centralised. Beneficial-ownership regimes made the ultimate controller a reportable particular. Major-shareholding notification made the position itself a filing at a stated threshold, and thresholds have been lowered and extended to instruments once outside them. Custody, settlement, index, flow and commercial holdings data made attribution derivable. And investment screening now asks, as a statutory question put to the buyer, who ultimately controls it — and answers adversarially where the buyer will not. **Protection is complete right up to the threshold crossing or the first regime the structure failed to anticipate, and then it is absent; and a position revealed late reads as evidence of an intention to conceal, which is a worse place to stand than having disclosed from the start.**

What the disclosed institution does instead. Decompose the function and the result stops being surprising:

THE FUNCTION	DOES IT NEED CONCEALMENT?
hold a diversified claim on world growth for future citizens	no -- a broad index position is not a secret and could not be made into one
keep access to markets that screen foreign state buyers	no -- and disclosure HELPS: a screening regime clears the buyer whose owner and mandate are known

```
vote the shares, and tell the      NO -- this requires the opposite.
market beforehand how the vote     A vote cannot be cast, still less
will be cast                       pre-announced, by a holder that
                                   cannot say who is voting

move a company's conduct by        no -- the exclusion is worthless
excluding it on published          unless the company and the reason
ethical criteria                   are named

contract on ordinary commercial    no -- a disclosed institutional
terms with brokers, custodians,    counterparty is EASIER to
banks and co-investors             contract with than an opaque one

survive changes of government      no -- an acknowledged, audited,
and public scrutiny at home        parliament-reported fund survives
                                   scrutiny BETTER
```

Read the middle rows, because they are the strongest form of the argument. Two of the functions are not merely indifferent to disclosure — they are impossible without it. A concealed holder cannot vote its shares as an owner, still less publish beforehand how it means to vote, because a vote is an act performed in a name; and it cannot change a company's behaviour by excluding it, because the exclusion has force only when the owner walking away and its reason are both public. **The disclosed vehicle does things the concealed one structurally could not, which is why the replacement is not a compromise forced by circumstances but an improvement that circumstances made available.**

There is a further piece of evidence of the kind arguments rarely get. **A design that works gets copied, and this one was.** When sovereign investors came under sustained political pressure in the late 2000s, the sector's answer was not better concealment but a published voluntary disclosure standard agreed by an international working group in 2008, and a standing forum that has maintained it since. **When states set out to defend these vehicles, not one of them defended them by concealing more.** They defended them by publishing the mandate, the holdings and the votes — institutions with a real commercial incentive to keep a position unattributed, and the capability to do it, choosing the other design repeatedly, independently and in public.

What transfers, and what does not. The analogy is strong and strong analogies get overrun. *What does not transfer:* such a fund operates with the acknowledgment, budget and legal protection of a sovereign state and is not defending against that state. It has no forced-heirship problem [Hague Conference on Private International Law, 1989, European Union, 2012], no correspondent-banking dependency, no exposure to a foreign court with personal jurisdiction over a founder, and no requirement that its structure survive a death. **Nothing here is a recommendation to copy the vehicle; the vehicle is not the transferable part.** *What transfers* is the design principle, in three propositions:

```
1  A function that depends on concealment has a protective quantity
   that is falling and cannot be replenished. Its obsolescence is a
   supply problem, not a moral or regulatory one.

2  Most functions historically performed by concealment do not require
   it. Decompose the function and score each part for its concealment
   dependency. Parts scoring zero can be rebuilt in the open at once;
   parts scoring above zero are the ones that will fail, and they are
   fewer than expected.

3  A structure that has disclosed everything has nothing left to be
```

```
exposed. That is not a consolation prize. It is a stronger
defensive position than partial concealment, because it removes a
category of adverse event rather than reducing its probability.
```

Proposition 2 is the practical instruction. The exercise to perform is not "how do I protect this structure from being seen" but "which specific functions of this structure would stop working if it were seen" — and then to discover that the honest answer is a short list, and that the items on it were the ones that were going to fail anyway.

What passing does not buy, stated without softening. It does not buy immunity from tax. It does not buy protection against a judgment that survived due process. And it does not buy protection against a state that decides to act without process. **In 2022 approximately USD 300 billion of central bank reserves were immobilised. Those reserves were among the most thoroughly disclosed assets in the world — held by a sovereign, in named accounts, at identified institutions, reported in published statistics, with no concealment element of any kind. Full disclosure did not prevent immobilisation.** Curzi [2026] owns that result: disclosure is not immunity, and the disclosure test is a filter that removes structures collapsing on contact with an authority, not a shield against a policy decision. What disclosure-based design does against arbitrary action is narrower and still worth paying for: it makes the position legally defensible rather than hopeless, it raises the political cost of acting against a demonstrably compliant holder, and — the part that matters most — **it is the precondition for the perimeter diversity of Section 6 to be usable at all, because a structure that cannot survive examination in each perimeter has not diversified across those perimeters; it has merely placed assets where it is not yet in trouble.**

8. Results: hypotheses tested

H1 — supported. The conventional structure's naive effective count is 3.4085 and its effective perimeter count is exactly 1.0000, a factor of 3.41. The falsifier was an independent-schedule erosion pattern with no common trigger; Section 5.2 traces one assessment through three institutions in under a year, and Section 5.3 names three structural features each independently sufficient to produce the correlation. **What the arithmetic cannot do is establish that the real within-perimeter correlation is one.** That is an empirical claim, it rests on the observed behaviour of the coordinated frameworks over fifteen years, and Section 12 records it as the paper's largest unmeasured input.

H2 — supported, exactly. At $c = 0.02$ and $i = 0.0816$ the joint probability is 0.020044 against $p^4 = 0.000100$, a factor of 200.4, and the idiosyncratic term contributes 0.0000435 — 0.22 per cent of the total. The prediction was a point value and the point value is met.

H3 — supported. $p^* = 0.008458$ at the constructed parameters, and 0.016916 when L is halved. The falsifier was a defensible estimate of annual examination probability below 0.85 per cent; **no such estimate is offered here either, and that is the honest limit of the result — the threshold is computed exactly and the quantity it is compared against is not measured.** Section 12 records this.

H4 — supported. $s = 2/3$ exactly, by solving $s^2 + 2((1-s)/2)^2 = 1/2$, and no allocation retaining more than two-thirds inside one perimeter reaches $N = 2$.

H5 — not met. At the parameters fixed in Section 2:

```
maintenance floor:  delta Q + C = 0.02 x 10 + 0.30 = 0.5000

conventional:  h = 0.35 x ln(1.000000) - 0.5000 = -0.500000
designed:      h = 0.35 x ln(2.941176) - 0.5000 = -0.122417

h = 0 requires  eta ln N = delta Q + C
                ln N = 0.5000 / 0.35 = 1.428571
                   N = e^1.428571 = 4.172734
```

The criterion was $N \geq 4.172734$. The redesigned structure reaches $N = 2.941176$. H5 is not met. The designed structure improves h by 0.377583, a 75.5 per cent reduction in the shortfall, and remains at -0.122417. **Both structures are below the threshold at which the structure is worth building. Jurisdictional diversity buys the ability to build; it does not buy the building.**

What would close the gap is visible in the equation and is not jurisdictional. At $N = 2.941176$:

```
C required for h = 0    :  0.377583 - 0.200000 = 0.177583   (against actual 0.30)
eta required for h = 0  :  0.500000 / 1.078810 = 0.463474   (against actual 0.35)
```

Both are operating improvements rather than structural ones, which is the paper's counter-result and is stated in the same plain language a success would have received.

9. Failure modes

Nine ways the instrument and the standard are misused, each with the arithmetic or the diagnostic that detects it.

1. Counting flags instead of perimeters. The master error; everything else is a special case. A diagram with entities in six countries looks diversified and may not be. The diagnostic is quick: apply the four-question test to each situs and shade every entity scoring four. If the shaded entities hold most of the capital, the diagram is decoration. The general form of the error is proved elsewhere: one party can present as many [Douceur, 2002], so **an unaudited count of nominally distinct parties is worthless.**

2. Treating a delisting as an all-clear. A jurisdiction that has just completed a supervisory overhaul in order to exit a list is at its most compliant, most closely monitored, and most tightly coupled to the assessing body. That may be exactly what a principal wants; it should not be mistaken for a return to the earlier product, which the overhaul removed.

3. Reading the offshore statute rather than the enforcement route. The statute says the foreign judgment will not be recognised, the statute is correct, and the principal is detained anyway, because the enforcement route ran through him. Before relying on any protective provision, trace the route by which an adversary would actually attack and check whether the provision sits anywhere on it.

4. Treating a bridge as neutral, or non-Western as unreachable. Section 5.5 on the first; Q4 on the second. A bridge classification is a claim about correlation, not about immunity.

5. Confusing a clearing layer with a messaging layer. A renminbi clearing system is a *clearing* system; a large proportion of its participants use the incumbent network for the messaging layer of transactions that clear through it. **A structure with two clearing systems and one messaging system has one point of failure, not two**, and a diversification claim must identify which layer it diversifies. Curzi [2026] runs the measurement.

6. Confusing a rail with a store. Freeze risk on a rail is an exposure-window problem with parameters — hop size, duration, counterparty pre-screening, mid-transit failure handling. Freeze risk on a store is a custody problem with entirely different parameters. **Never evaluate a rail as though it were a store of value.** The error is especially common with tokenised settlement instruments, which are frequently described as being outside the perimeter and are not: issuers freeze balances at law-enforcement request, and public ledgers are *more* traceable than correspondent banking, not less.

7. Over-diversifying. Section 6.5. Logarithmic benefit against linear cost has an interior optimum and the optimum is not large.

8. Assuming perimeters are static. The classifications are judgements about the present. Perimeters expand — each accession converts a jurisdiction from partially outside to inside on Q1 — and they can fragment. **The classification must be re-made periodically with dated evidence rather than impressions**, which is why the trajectory record of Section 4.7 requires

dated events. The mirror case is worth noting: threat-5 style fragmentation runs both ways and only one direction is usually planned for.

9. Optimising the index instead of the thing it measures. Two structures with identical bloc indices can differ enormously in whether the blocs are actually independent in the scenario that matters. **The index is a diagnostic, not an objective function.** This is Goodhart's problem exactly [Goodhart, 1975] and Campbell's law in the evaluation literature [Campbell, 1979]; Scott [1998] is the historical form, in which legibility destroys the local variety that made the system work, **and it is a warning about this instrument as much as an argument for it.** The tell is a structure whose bloc shares are suspiciously round and equal; real structures are lumpy because assets are lumpy.

Two further failure modes belong to the disclosure standard rather than the index. **Grading privacy on the opacity axis** — the structure is described as passing because "nothing is public", which answers a question the test does not ask; re-run it assuming every document is in the hands of every authority and score only the consequences. And **designing for the examiner and forgetting the bank** — practitioners optimise for the tax authority because that is the reader whose questions come in writing, while the binding constraint in practice is usually the correspondent bank's compliance function, which asks fewer questions, gives no reasons, has no appeal, and can end the structure's usefulness in ninety days.

10. Discussion

What follows. Three things, each narrower than the general claim it resembles.

The correlated-failure result is arithmetic, given the correlation structure. The joint probability $c+(1-c)i^4$ is exact, and the collapse EFF $= 1/(w'Rw)$ is a variance reciprocal [Markowitz, 1952, Leinster and Cobbold, 2012]. Nothing in the derivation is approximate. This is why the result survives the statistical-averaging objection of Doak et al. [1998]: no claim is made that diversity produces stability through an unspecified mechanism, only that a variance has a reciprocal and that the reciprocal counts independent shocks.

The status of the claim is measurement, not physics. The corpus is titled a psychological and sociological physics and the word is checkable. **What is claimed here is a measured count with a dimension, stated limits and an invariance group** — Sections 11.1 to 11.4 — and **not** membership of the statistical-physics programme surveyed by Castellano et al. [2009], whose vocabulary of order parameters, phase transitions and coarsening this paper deliberately does not use. **No order parameter is identified, no phase transition is located, and no critical exponent is estimated.** Hedström and Ylikoski [2010] states the criterion the named discipline applies: an explanation must name a mechanism rather than exhibit a fitted curve. **The mechanism named here is a transmission structure with five interfaces, and nothing is fitted.** Gallegati et al. [2006] is the checklist this paper was run against — ignorance of existing work in the target discipline, sub-standard statistical procedure, unjustified universality claims, and models built on assumptions known to be false — and Section 12 reports where it passes and where it does not.

The collapse mechanism has a formal ancestry that should be named rather than reinvented. A small correlated shift producing a discontinuous collective outcome is the threshold and tipping mechanism: Granovetter [1978] establishes that the distribution of individual thresholds, not the mean, determines whether an outcome tips, and Schelling [1971] is the canonical warning against reading an individual-level parameter straight into an aggregate — **which is precisely what a perimeter assignment does.** Castellano et al. [2000] supplies the closest formal analogue of a bloc: the same interaction that raises local agreement freezes global diversity into non-interacting regions, with a sharp transition between the monocultural and fragmented phases. **A perimeter, in that vocabulary, is a frozen region.** On the generating side, Bouchaud and Mézard [2000] derives concentration rather than observing it — above a critical ratio a finite fraction of the total condenses onto one holder, **which is the index going to one, obtained from a mechanism.** This paper does not use that model and

asserts no tail; the citation records that a mechanism-side companion exists.

Redundancy is not automatically resilience. Percolation gives the sharpest formalisation of how many independent components a system really has [Callaway et al., 2000], and the robust-yet-fragile result [Albert et al., 2000] is the canonical demonstration that a network can be highly tolerant of random loss and extremely vulnerable to targeted loss — with the caveat, from Doyle et al. [2005], that real engineered systems get their robustness from design rather than from a degree distribution. The perimeter analogue is that a structure's tolerance depends on which component is removed.

What does not follow. Four things.

It does not follow that raising N is always right. Section 6.5; Uzzi and Spiro [2005]; Sagan [1993]. The prescription is bounded.

It does not follow that perimeter diversity answers every threat. It is close to irrelevant to counterparty failure and can worsen it, because the population of institutions willing to bank a genuinely cross-perimeter structure is small. Curzi [2026] owns the threat taxonomy that separates the five and shows that a structure answering one well may answer another worse than doing nothing.

It does not follow that compliance confers immunity. Section 7.6, and Curzi [2026].

It does not follow that the corpus's affine effective count inherits any of the diversity literature's support. Section 3.2.

Where this paper stands relative to its own corpus. The fifth-decimal reconciliation of Section 6.3 is the substantive point of contact with Curzi [2026], and it is printed in both papers rather than harmonised in either. Curzi [2026i] bounds how much confidence a passing check should buy: the corpus's own preregistered fit **failed its own second falsification rule**, achieving a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin, despite +40.3 per cent over persistence, +97.7 per cent over direct input, and $p = 0.0005$ against the shuffle. **A paper in this series that reported only its successes would be departing from the practice of the paper that reported that failure.**

Parallel results elsewhere in the corpus. The shape of the correction recurs on different objects. Curzi [2026s] derives trust as retention, $\delta_{\text{eff}} = \delta_0 + (1-a)(1-v)$, giving a superlinear stock dividend from 0.4679 at $a = 0.50$ to 5.0529 at $a = 0.99$ — an input to the same δ this paper decomposes, from an entirely different direction. Curzi [2026] supplies the comfort-weighted amplitude, scoring each step 0–3 for steepness and weighting a branch as $1/2^R$, so that a three-point difference in steepness is a factor of eight in expectation — the same exponential-in-a-small-integer structure that makes a two per cent common component worth two hundred times an idiosyncratic one. Curzi [2026] supplies the discipline this paper needs most: the surplus functional is deployed as a **lower bound** on the ensemble's information and not as the information itself, so the framework logs the floor and not the ceiling. Curzi [2026] gives the minimum sustainable frequency $f_{\min} = ke^{2c/\eta}$ with $k = 0.483425732988414$ Hz and the trichotomy $f \gtreqless f_{\min}$ corresponding to $Q_\infty \gtreqless 0$ — the same sustainability threshold as this paper's $h \gtreqless 0$, expressed in a different unit. **The constant k is defined, not measured**, and the attachment of Hz to it is a separate and stronger claim; that badge travels with the citation. Curzi [2026] gives the odd-part class rule — the pitch class of $f(n) = kn^2$ is the odd part of n, so $1 \dots N$ realises exactly $\lceil N/2 \rceil$ classes, six at $N = 12$ — a discrete count invariant under a relabelling, as N here is invariant under rearrangement inside a perimeter.

11. Verification and rigour audit

The results are decomposed into postulates checkable one at a time. Sections 11.1 to 11.6 run the standing rigour checks first, so that the chain in 11.8 is read against a dimensional audit rather than before one. **If dropping a postulate leaves the result standing, it is not listed.**

11.1 Dimensional audit

symbol	quantity	dimension	unit	status
s_i, w_i	exposure share of jurisdiction i	$[1]$	dimensionless, fraction of Q	defined; sums to 1 by P1
W_b	perimeter share	$[1]$	dimensionless	derived, $W_b = \sum_{i \in b} s_i$
HHI	concentration index	$[1]$	dimensionless (or points $= 10^4 \times$ fraction)	derived from shares
ρ_{ij}	correlation between shocks	$[1]$	dimensionless, $[-1, 1]$	assumed, not measured
N	effective perimeter count	$[\#]$, **a count**	perimeters	derived as $\mathrm{HHI}_{\mathrm{bloc}}^{-1}$
$\ln N$	input-term argument	$[1]$	nats	derived; the logarithm of a count is dimensionless
c, i, p	common, idiosyncratic and marginal failure probabilities	$[1]$	dimensionless, per year	assumed (constructed)
Q	protected capital	$[C]$	index units of stored capital	defined; **no operational unit is supplied by the source**
δ	erosion rate	$[T^{-1}]$	per period	assumed
C	maintenance cost flow	$[C\,T^{-1}]$	index units per period	assumed
η	conversion efficiency	$[C\,T^{-1}]$ per nat	index units per period per nat	assumed
h	surplus balance	$[C\,T^{-1}]$	index units per period	derived
c_0, c_1	fixed and per-perimeter maintenance	$[C\,T^{-1}]$, $[C\,T^{-1}\,\#^{-1}]$	index units per period, per perimeter	assumed
L, p^*	loss on examination, break-even probability	[currency], $[1]$	USD, dimensionless	assumed (constructed), derived

Does the surplus law balance? $h = \eta \ln N - \delta Q - C$ requires every term to carry $[C\,T^{-1}]$. δQ does, since δ is per period and Q is a stock. C does by definition. $\eta \ln N$ does **only if** η **carries** $[C\,T^{-1}]$ **per nat**, and **the source is silent on this**: it treats η as a bare "conversion efficiency". The audit records the silence rather than filling it in. The equation balances under the reading in the table, and the reading is this paper's, not the source's.

Does the disclosure-economics comparison balance? $p^* = (C_{\mathrm{comp}} - C_{\mathrm{conc}})/L$ divides a currency difference by a currency loss and returns a dimensionless probability. It balances. **But note the mixed dimension in the comparison itself**: C is a flow per year and pL is an expected flow per year, so the two are commensurable only because p is an *annual* probability. Stating p over any other horizon breaks the comparison, and Section 7.2's table is annual throughout.

The important row is N. The index is a dimensionless sum of squared dimensionless shares, so its reciprocal carries the dimension of a **count**. **The instrument is not a score that happens to start at 1; it is a measured count of independent units, with a unit**, and Section 11.3 checks it against the definition of a count.

Where no operational unit exists, this audit says so. Q is protected capital in index units and the source supplies no conversion to currency. Every figure in C, h and Q is therefore on an arbitrary scale fixed by the stated parameters. **The break-even $N = 4.172734$ is unit-free and survives this; the statement that $h = -0.122417$ does not mean anything until the index unit is defined.** The USD figures of Section 7.2 are in a different and genuine unit and must not be read as commensurable with $C = 0.30$; **the paper uses the symbol C for both and that is a defect of notation inherited from the source, flagged here.**

11.2 Significant figures and epistemic status

- $N = 1.000000$ and $N = 2.941176$ are **exact arithmetic** on the stated shares, $1/1$ and $1/0.34$. They are derived-from-defined and may carry as many digits as are useful.
- HHI $= 0.293388$ is **exact arithmetic on exact fractions** ($142/484$) and 0.293430 is exact arithmetic on the three-significant-figure percentages. Both are exact; they are exact on *different inputs*, which is the whole content of Section 6.3.
- $N = 4.172734$ is **exact given the parameters** — $e^{0.5/0.35}$ — and the parameters are assumed, not estimated. **It carries no uncertainty because there is no estimate in it, and no authority beyond the assumption.**
- $p^* = 0.008458$ is exact arithmetic on three constructed currency figures. Six digits are defensible only because the inputs are exact by stipulation. **It is not a measurement**, and the quantity it must be compared against — a real examination probability — is not measured anywhere in this paper.
- The correlated-failure outputs (0.020044, 0.794109) are **illustrative model output from constructed parameters** and are rounded to six decimal places at most.
- The erosion percentages (73.857%, 54.548%) are exact arithmetic on an assumed rate.
- **No quantity in this paper is a fitted or measured estimate, so no uncertainty interval is reported anywhere. That is not a strength: it means the paper has no empirical content to attach an interval to.**

11.3 Limits and boundary behaviour

The instrument's limits are the proof that it behaves as a count, and they are run rather than asserted.

limit	HHI	HHI^{-1}	reading
one holder, $s_1 = 1$	1	1 exactly	a monopoly counts as one unit
n equal holders, $s_i = 1/n$	$1/n$	n exactly	returns the count
n holders, any unequal split	$> 1/n$	$< n$	never overstates the count
$s_{\max} \to 1$, rest $\to 0$	$\to 1$	$\to 1$	continuous approach to the monopoly limit
$n \to \infty$, equal shares	$\to 0$	$\to \infty$	unbounded above, as a count should be
$\rho \to 0$ within perimeters	naive	naive count	the independent limit
$\rho \to 1$ within perimeters	bloc	perimeter count	the fully-coupled limit

Range: $1 \leq \text{HHI}^{-1} \leq n$, with both bounds attained. A quantity that equals n for n equal units, equals exactly 1 for a single unit, and lies strictly between otherwise **is a count**, and no further argument is needed. Verified on the two structures here: $1/1 = 1$ exactly for the conventional perimeter vector, and $1/0.25 = 4$ exactly for a hypothetical four-way equal split.

The same limits run on the affine count, which fails them. $N_{\text{eff}} = 1 + (N-1)u$ returns 1 at $u = 0$ and N at $u = 1$ for *any* share vector whatever, including $(0.999, 0.001)$, for which it reports 2.000000 at $u = 1$ while HHI^{-1} reports 1.002004. **It has the right endpoints and the wrong interior, because it never consults the shares.** Curzi [2026] runs the full chain.

Boundary behaviour of the correlated-failure law. $P = c + (1-c)i^4$ gives $i^4 = p^4$ at $c = 0$, the independence case, and c at $c = p$ where $i = 0$. Both endpoints are attained in the table. $\partial P/\partial c = 1 - i^4 \approx 1$ for small i: **the joint probability tracks the common component one-for-one, which is H2 in derivative form.**

Boundary behaviour of the surplus law. As $N \to 1^+$, $h \to -(\delta Q + C)$: a structure inside one perimeter pays its whole bill out of stock. As $N \to \infty$, $h \to \infty$ logarithmically, **which is a defect of the model and a statement of its domain**, since Section 6.5 and Uzzi and Spiro [2005] establish that the real curve turns over. **The model's domain is small** N, and every prescription here is confined to the move from $N \approx 1$ to $N \approx 3$. As $\eta \to 0$, $h \to -(\delta Q + C)$ regardless of N: no conversion, no benefit from any amount of independence. As $\delta \to 0$ the stock never erodes and any positive $\eta \ln N - C$ accumulates without bound, **which is a pathology rather than an ideal.** As $\delta \to 1$ the stock is annihilated each period and no structure has meaning. Under the maintenance model $C = c_0 + c_1 N$, h is concave with a unique interior maximum at $N^* = \eta/c_1$ and falls without bound thereafter — **the model does bound itself once maintenance is written in, which is the correct answer to the unbounded limit above.**

Boundary behaviour of the disclosure economics. As $L \to 0$, $p^* \to \infty$ and concealment is always cheaper — the limit in which examination costs nothing, which is the assumption the failing structures of Section 7.4 are implicitly making. As $L \to \infty$, $p^* \to 0$ and concealment is never cheaper. As the cost gap closes to zero, $p^* \to 0$ likewise. **The result is therefore driven entirely by the ratio of the cost gap to the loss, and no other feature of the model does any work.**

11.4 Invariance

The effective count is invariant under relabelling of units and under any rescaling of the underlying exposures that preserves the share vector. Doubling every exposure in currency terms leaves every s_i unchanged and therefore leaves the index, N and $\ln N$ unchanged. This is a genuine scale invariance and it is why the instrument applies to a structure of any size.

It is not invariant under regrouping, and that is the whole content of the collapse: the naive and perimeter figures differ by construction and the difference is the finding. A reader who wants a single number invariant under regrouping is asking for a quantity that cannot detect a perimeter.

What would break the scale invariance: any share basis that is not normalised. **Tested at how many scales?** The invariance is algebraic, not empirical, and was checked at exactly the two structures reported here. **No claim of empirical scale-freedom is made**, and no distributional or scaling claim is made anywhere in this paper. Clauset et al. [2009] is the reason that matters: a principled maximum-likelihood fit with a goodness-of-fit test and a likelihood-ratio comparison leaves few canonical power-law claims standing, and **a straight line on a log-log plot is not evidence.** That procedure was not run because no claim requiring it is made; Stumpf and Porter [2012] adds that even a genuine power law would be weak evidence for a mechanism. **The exposure vectors here are finite sets of four and six shares, not samples from a distribution, and must not be restated as such.**

What is conserved, and in what units. Mirowski [1989] shows that the characteristic failure of this genre is borrowing the notation of a conserved quantity without the conservation law that gave it content. **The objection is met here and the answer is short.** The conserved quantity is the exposure share vector: $\sum_i s_i = 1$ exactly, in dimensionless units, and it is that

constraint and nothing else that gives the index its bounds. **Remove the normalisation and the effective count ceases to be a count.** The invariance group is relabelling together with any rescaling that preserves the simplex. Q**, by contrast, is a stock with no stated conservation law and no operational unit, and this paper does not claim one for it.** Cartwright [1999] supplies the discipline that makes a partial answer respectable: specify the arrangement within which the law is claimed to hold and stop claiming a universality the argument does not need. **The arrangement claimed here is a normalised exposure vector with a correlation matrix having unit diagonal and non-negative entries, and nothing outside it.**

Aggregation is a separate and unmet objection. Kirman [1992] establishes that a single representative parameter for a heterogeneous group generally does not exist. This paper writes one Q, one δ and one C for a whole structure whose parts plainly differ. **No aggregation argument is offered and none is available from the arithmetic.**

11.5 Analogy discipline

borrowing	source	category	justification
portfolio variance $\mathrm{Var}(R) = w'Rw$	Markowitz [1952]	**1 — derivation**	the exposure vector genuinely is a weighted sum of shocks with a correlation matrix; the premises are the algebra's own
Hill numbers, order 2	Hill [1973], Jost [2006]	**1 — derivation**	HHI^{-1} *is* ${}^2D$; an identity, not a resemblance
similarity-sensitive diversity	Leinster and Cobbold [2012], Rao [1982]	**1 — derivation**, with a stated domain	exact where R has unit diagonal and non-negative entries; claimed nowhere else
$\ln N$ as entropy	Shannon [1948]	**2 — shared form**	this is **Shannon entropy in nats**, a dimensionless information quantity, and not thermodynamic entropy; no thermodynamic content is imported
coupled-network fragility	Buldyrev et al. [2010]	**2 — shared form**	jurisdictions have not been shown to satisfy the percolation premises; only the structural property that coupling destroys assumed independence is used

borrowing	source	category	justification
robust-yet-fragile	Albert et al. [2000], Callaway et al. [2000]	**2 — shared form**	the degree-distribution result is not claimed to apply; the shape of the tolerance asymmetry is
synchrony in ecology	Loreau and de Mazancourt [2008]	**2 — shared form**	the naming convergence is recorded; no ecological mechanism is imported
threshold and tipping	Granovetter [1978], Schelling [1971]	**2 — shared form**	the aggregation warning is used; no threshold distribution is estimated here
frozen regions in a coevolving population	Castellano et al. [2000]	**2 — shared form**	a perimeter is structurally a frozen region; no phase transition is located
normal accidents	Perrow [1984], Sagan [1993]	**3 — framing**	qualitative organisational theory, cited as prior statement of the same observation, never as a formal result
open design	Kerckhoffs [1883], Saltzer and Schroeder [1975]	**1 — derivation**	the disclosure standard *is* the open-design requirement applied to a legal rather than a cryptographic system; the premise — that the design is public and only a key is not — is exactly what the disclosure test asserts

The failure mode this table exists to prevent is letting a category-2 borrowing carry a category-1 result's authority. The concentration and probability arithmetic is category 1 throughout and needs no help; the network and tipping results are category 2 and are used for shape only.

11.6 Measurement-theoretic status of the scored inputs

Two inputs in this paper are scores rather than counts and their scale type is not optional.

The four-question perimeter test yields a nominal classification, not a measurement. Q1–Q4 are binary and their sum, "4/4", is a count of yes answers on a **nominal** set of attributes.

Summing nominal indicators does not produce an interval quantity, and the "score" column of Section 6.2 must be read as a tally rather than as a degree of attachment. The paper's own text is explicit that Q4 carries more weight than the others, which is a statement that the four are *not* exchangeable — **and an unweighted sum of non-exchangeable indicators is not licensed.** The tally is used only to decide a binary membership, which it can support; it is used nowhere as a magnitude.

The five disclosure criteria are binary and their combination is a product, not a mean. That is the one operation the scale type does license: a conjunction of binary indicators is a binary indicator. **The "average score" column of Section 7.3 is included precisely to demonstrate that averaging them produces a wrong verdict — the illicit operation is the object of study, not a step in the argument.** This is stated here so that no downstream work reads the 0.80 as a computed quantity.

Q, protected capital, is treated as a ratio-scale quantity and that treatment is not tested. Michell [1997] is the strongest available objection to this class of move: psychology and its neighbours routinely assume quantitative structure without testing for it. The formal remedy is conjoint measurement [Luce and Tukey, 1964], whose conditions establish whether an attribute has additive quantitative structure at all. **The conjoint conditions were not tested for Q, and the corpus's habit of writing $Q = 10$ in unnamed index units is exactly the practice Michell criticises.** What survives without the test is the *ordering* and the unit-free break-even count. **The concentration instrument itself is untouched by this objection, because exposure shares are ratios of capital and not psychological judgements.**

11.7 Postulates

#	statement	status
P1	exposure shares s_i are non-negative and sum to 1	DEFINITION
P2	HHI $= \sum_i s_i^2$; the effective count is its reciprocal	DEFINITION
P3	for n equal shares of $1/n$, HHI $= n(1/n)^2 = 1/n$	ARITHMETIC
P4	a coordination perimeter is a set of jurisdictions among which an adverse decision propagates without a separate act of consent	DEFINITION
P5	the four named jurisdictions answer yes to all four questions of the perimeter test	POSTULATED (H1; the analyst's judgement, uncited)
P6	each jurisdiction's failure decomposes as a common shock c and an independent idiosyncratic shock i, with marginal probability $c + (1 - c)i$	DEFINITION (modelling choice)
P7	the surplus law's input term is $\eta \ln N$	DEFINITION (modelling choice, CCP-A02)
P8	$\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$	POSTULATED (fixed in advance)

#	statement	status
P9	the constructed cost figures: $C_{\text{comp}} = 561{,}000$, $C_{\text{conc}} = 210{,}000$, $L = 41{,}500{,}000$	POSTULATED (constructed)

11.8 The compounding chain

step	consumes	statement	status
L1	P2, P3	HHI^{-1} returns n exactly for n equal shares	ARITHMETIC
L2	L1	HHI^{-1} is a numbers-equivalent — Adelman [1969]	CITED
L3	L2	HHI^{-1} is the Hill number of order 2 — Hill [1973]	CITED
L4	P4, P5	the four jurisdictions constitute one perimeter, so the collapsed share vector is (1.0)	POSTULATED
L5	P1, P2, L4	bloc-HHI $= 1.000000$ and $N = 1.000000$ exactly	ARITHMETIC
R1	P2	naive HHI on the exact fractions $= 142/484 = 0.293388$; effective jurisdictions $= 3.408451$	ARITHMETIC
R1a	P2	naive HHI on the three-significant-figure percentages $= 0.293430$; effective $= 3.407968$ — the reproducible basis, agreeing with Curzi [2026]	ARITHMETIC
R1b	R1, L5	the overstatement is $3.408451/1.000000 = 3.41\times$	ARITHMETIC
R2	P6	$P(\text{all four}) = c + (1-c)i^4$; at $c = 0.02$, $i = 0.0816$ this is 0.020044 against $p^4 = 0.000100$	ARITHMETIC
R2a	R2	over 15 years, $1-(1-0.1)^{15} = 0.794109$ correlated against $1-(1-0.0001)^{15} = 0.001499$ independent	ARITHMETIC
R3	P9	the cost gap is $351{,}000$, so $p^* = 351{,}000/L = 0.008458$; at $L/2$ it is 0.016916	ARITHMETIC

step	consumes	statement	status
R4	—	five binary criteria combined multiplicatively: any zero gives a product of zero, so 0.8^{mean} and the product part company at the first failing element	ARITHMETIC
R5	P8	$(1-0.02)^{15} = 0.738569$, $(1-0.02)^{30} = 0.545484$, half-life $\ln 2/\ln(1/0.98) = 34.310$	ARITHMETIC
R6	P7, P8	$h = \eta \ln N - \delta Q - c_0 - c_1 N$ is maximised at $N^* = \eta/c_1$: 7.0, 5.0, 3.5 at $c_1 = 0.05, 0.07, 0.10$	ARITHMETIC
R7	L5, P7, P8	$h = -0.500000$ at $N = 1$; $h = -0.122417$ at $N = 2.941176$; break-even $N = e^{1.428571} = 4.172734$	ARITHMETIC
CR	R7	both structures are below break-even, so the redesign is a reduction in shortfall and not a solution	ARITHMETIC

Every ARITHMETIC row can be redone by hand or in three lines of any language. Every CITED row names one source and no more. **Only P5, P8 and P9 are POSTULATED, and one of them — P5, the perimeter assignment — is the paper's entire empirical exposure.** All three are stated as hypotheses or as constructed inputs in Section 2. **A chain whose central step is a postulate is not a proof, and this one says so here rather than leaving the reader to notice: the arithmetic is exact and the inputs are judgements.**

11.9 Cross-checks performed

- $s_{\max}^2 \le \text{HHI} \le s_{\max}$ holds on the conventional structure: $0.167355 \le 0.293388 \le 0.409091$.
- The effective-count ceiling $s_{\max}^{-2}$ is respected: $3.408451 \le 5.977778$.
- The perimeter shares sum to 1.000000 exactly on both structures, and to $0.40 + 0.30 + 0.30$ exactly on the designed one.
- The three-basis reconciliation of Section 6.3 was computed independently in exact rational arithmetic and in floating point, and the two agreed to twelve decimal places on every row.
- The marginal-probability constraint $c + (1-c)i = 0.10$ was verified for each row of the common-shock table; the largest residual is 3.6×10^{-5}, arising from the four-significant-figure printing of i and not from the model.
- The rounding audit was run and the discrepancy quantified at 0.000140.

12. Limitations

The correlation is assumed, not measured. The collapse uses $\rho = 1$ within a perimeter, which is the conservative end of a family, and the justification is the observed behaviour of the coor-

dinated frameworks over fifteen years. **No correlation was estimated from data in this paper.** That is the single largest unmeasured input and it is exactly the quantity H1 turns on. Curzi [2026] publishes the continuous family so that a reader may substitute a different value.

The perimeter assignment is a stipulated reference group, not an identified effect. Manski [1993] rules out recovering an endogenous social effect from observation of a linear-in-means reflexive system, and the framework's loop is such a system. The remedy adopted is stipulation with a defence [Brock and Durlauf, 2001], and **the price is that the assignment is a modelling position.** No randomisation, instrument or partial-population variation was available.

The instrument is performative and its own parameters are exposed to that. MacKenzie and Millo [2003] documents a model that made its assumptions true after adoption; Lucas [1976] states the consequence for parameters estimated before publication; Hacking [1995] gives the general looping form. **The framework asserts that systems change when modelled and then exempts its own η, δ and correlation assumptions from that reflexivity. That is an inconsistency and it is not resolved here.**

The examination probability is not measured, and H3's conclusion depends on it. The break-even $p^* = 0.008458$ is exact. The claim that realistic examination probability exceeds it is **an assertion this paper cannot support from the bibliography**, which contains no empirical estimate of examination or enforcement rates. The claim is badged here as a postulate rather than a finding, and the honest form of H3's result is conditional: *if* annual examination probability exceeds 0.85 per cent, concealment is the more expensive design.

Every parameter in the surplus law was chosen, not estimated. $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ are stated test parameters fixed before the measurement, which protects H5's criterion from adjustment but makes the break-even a property of the chosen parameters and not of the world. **At $\eta = 0.50$ the break-even falls to $e^{1.0} = 2.718$ and the designed structure clears it. That the conclusion of H5 depends on η in this way is the most important sensitivity in the paper.**

The surplus functional's form is underdetermined by the evidence. Social impact theory proposes influence growing as a power of the number of sources with exponent below one [Latané, 1981]; an $n \log n$ form fits the qualitative facts equally well; and the framework proposes $\ln(1 + (N-1)u)$. **All three fit the qualitative facts here equally well, and nothing in this paper discriminates between them.** The strongest available warrant for the logarithm specifically is the multiplicative one [Peters and Gell-Mann, 2016], which requires no perceptual story; it is preferred here to the psychophysical justification for that reason.

Every case is constructed. The principal, the assets, the shares, the cost model, the failure probabilities, the erosion decomposition and the five scored structures are invented and are labelled as such at each point of use. **No empirical dataset was collected and none is reported.** The measurements are exact computations on constructed inputs.

Jurisdiction-specific legal claims are not cited to primary sources. The descriptions of statutes, directives, decided cases, listing events, registers, promotion regimes, foundation regimes and individual jurisdictional characteristics are taken from the source manuscript. **The bibliography contains no jurisdiction-specific primary sources — no statutes, regulations, court rules or registries — and no case law.** Those claims are uncited here, are described rather than asserted, and require verification against the official sources. Appendix B lists them individually. **This is a real gap and it is not closed here.**

The dated events are not independently verified. The listing sequences, review dates and decision dates in Sections 4.2, 4.3 and 5.2 are reported as publicly announced and were not checked against primary releases by this paper. A reviewer should treat the propagation trace of Section 5.2 as an illustration of a mechanism whose *dates* require confirmation, even though the mechanism itself — a rating consumed as a listing criterion — is stated in the standard-setting materials themselves.

The effective count measures independence and not capacity. $N = 2.94$ says nothing about whether any perimeter can carry the load. A structure whose perimeter diversity is real but whose banking capacity sits in one perimeter has an effective count near 1.00 on the ability to

transact. Curzi [2026] owns this as its counter-result and it applies here unchanged.

Monotone claims are restricted. Section 6.5 bounds the prescription, and Uzzi and Spiro [2005], Sagan [1993] and May [1972] are the reasons a monotone raise-N claim could not survive unqualified.

The instrument is exposed to its own failure mode 9. Adopted as a target, N would cease to be a good measure [Goodhart, 1975, Campbell, 1979]. Nothing in the arithmetic prevents that, and the only defence offered is the reporting discipline of Section 3.5, which requires both figures to be published.

12.1 Rigour checks not met, listed by name

1. **No as-of date, because no real holdings were measured.** A concentration figure computed from real holdings requires its source and its date; **a concentration measure without an as-of date is not a measurement.** Every share vector here is constructed, so there is no source and no date to give. **Any application of this instrument to a real structure must carry both**, and the reporting discipline should make the date a mandatory item.
2. **No uncertainty is reported anywhere, because nothing is estimated.** That is not precision; it is the absence of empirical content.
3. **The dimension of η is supplied by this paper, not by the source.** Section 11.1.
4. **Q, C and h have no operational unit**, and the symbol C is used both for a dimensionless index flow and for a USD budget. Flagged in Section 11.1; the source does the same.
5. **The surplus law's large-N limit is unbounded**, which the empirical literature contradicts. The model's domain is small N. Writing $C = c_0 + c_1 N$ bounds it, and that is the version to prefer.
6. **The conjoint conditions were not tested for Q** [Michell, 1997, Luce and Tukey, 1964], so its ratio-scale treatment is provisional.
7. **An unlicensed operation appears in Section 7.3 on purpose**, and a second one is flagged in Section 11.6: the "4/4" tally sums nominal indicators that the text itself says are not exchangeable. It is used only to decide a binary membership.
8. **Model comparison was not performed.** No alternative index was fitted, no out-of-sample test was run, no information criterion was computed and no matched null was constructed, because nothing here is fitted. The instrument is asserted to be the right one on the strength of the identity chain, not on comparative fit.
9. **No distributional or scaling claim is made, so the Clauset–Shalizi–Newman procedure was not run** [Clauset et al., 2009]. **Any restatement of this paper reaching for a heavy-tail claim about holdings must run it first**, and must read Stumpf and Porter [2012] on why passing would still not identify a mechanism.
10. **The coupled-network and tipping borrowings are category 2, not category 1.** Section 11.5.
11. **There is no conservation law for Q.** Section 11.4 gives the one genuine conservation — shares sum to one — and states plainly that the stock variable has none [Mirowski, 1989].
12. **Aggregation is unaddressed** [Kirman, 1992].
13. **Expected value and time average are not the same thing** [Peters, 2019]. The disclosure-economics comparison of Section 7.2 is an ensemble average and the principal lives a time average; a single examination event does not arrive in expectation, it arrives once and rescales everything after it.
14. **No differential prediction is offered.** Durlauf [2005] asks for the prediction that only this model makes, and neither the collapse nor the disclosure economics makes one: any correlation-aware concentration measure would produce the same collapse, and any expected-loss comparison the same break-even. **The contribution is the identification of the correct unit of analysis and the discipline of reporting both figures, not a uniquely discriminating forecast.**
15. **The counter-test of Section 5.4 is not a systematic search.** Three counter-examples are

offered because the source manuscript found three. No procedure was run to enumerate the population of possible counter-examples, so the claim that the exchange trend has “no counter-examples” is a claim about what was looked for, not about what exists.

13. Conclusion

This paper establishes the following results in the form the rest of the corpus cites them.

R1. An entity map showing the United Kingdom at 40.91 per cent of exposure, the United States at 27.27, Jersey at 18.18 and Luxembourg at 13.64 gives a naive Herfindahl index of 0.293388 and 3.4085 effective jurisdictions on the exact-fraction basis, but collapses to a bloc index of 1.000000 and exactly 1.0000 effective coordination perimeters — a $3.41\times$ overstatement. **On the reproducible three-significant-figure basis used by Curzi [2026] the naive figures are** 0.293430 **and** 3.407968**; both perimeter figures are exactly** 1.000000 **and no conclusion moves.**

R2. Holding each jurisdiction’s marginal failure probability at ten per cent, a common shock of $c = 0.02$ raises $P(\text{all four fail})$ from 0.000100 to 0.020044, a $200.4\times$ understatement; over fifteen years the figures are 0.001499 independent against 0.794109 correlated.

R3. A concealment-based arrangement costing USD 210,000 a year against a compliant USD 561,000 breaks even at an annual examination probability of $p^{*} = (561{,}000 - 210{,}000)/41{,}500{,}000 = 0.008458$. **The visible 62.6 per cent saving counts the invoice, not the exposure.**

R4. The five-criterion disclosure test is scored as a product, not an average: five structures averaging 0.80 on their criteria yield a group product of zero, because two of them contain a zero-scoring element.

R5. The erosion rate decomposes as $\delta = 2.0$ per cent a year — tax drag 0.9, arbitrary-action risk 0.4, debasement 0.5, haven degradation 0.2 — retaining 73.857 per cent at fifteen years and 54.548 per cent at thirty, on a half-life of 34.310 years.

R6. With maintenance $C = c_0 + c_1 N$, the surplus $h = \eta \ln N - \delta Q - C$ is maximised at $N^{*} = \eta / c_1$: at $\eta = 0.35$ that is 7.0 for $c_1 = 0.05$, 5.0 for 0.07, and 3.5 for 0.10. The second perimeter is worth 5.19 times the eighth.

Counter-result. Passing the five-criterion disclosure test buys nothing against arbitrary state action — approximately USD 300 billion of fully disclosed reserves were immobilised in 2022 — and even the redesigned structure remains at $h = -0.122417$, still below the maintenance floor. **H5 was tested and not met: the redesign reaches** $N = 2.941176$ **against a break-even of** 4.172734**. The diversification gain is a reduction in shortfall, not a solution, and what would close the gap is operating — η to** 0.463474 **or** C **to** 0.177583 **— rather than structural.**

Two further results are established as method rather than as measurement. **The correct unit of analysis for concentration risk is the coordination perimeter and not the jurisdiction**, and the procedural form of that is to collapse before measuring and to report both figures, treating the gap as the size of the illusion the structure was operating under. And **the disclosure standard is Kerckhoffs’s open-design principle applied to legal structure**: a structure whose protective quantity is an adversary’s ignorance has a quantity that is falling and cannot be replenished, while a structure whose protective quantity is its own correctness has one that is not consumed by use.

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every case in this paper — the principal, the asset mix, the exposure shares, the failure probabilities, the erosion decomposition, the maintenance budget, the concealment cost comparison and the five scored structures —

is **CONSTRUCTED** and is labelled as such at each point of use. The numerical results are exact computations on those constructed inputs, carried at full precision and rounded only for printing.

All figures were produced by the scripts named in their captions (`figs/src/CCP-051_fig01.py` through `figs/src/CCP-051_fig06.py`), which import the shared figure style module and read their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

Jurisdictional and legal descriptions in this paper are taken from the source manuscript and not from primary legal sources. The bibliography contains general treaty, standard-setting and comparative material only; it contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction. Appendix B lists the claims affected. Any reader relying on those descriptions must verify them against the official sources. The dated listing sequences in Sections 4.2, 4.3 and 5.2 are reported as publicly announced and were not verified against primary releases by this paper.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. Every principal, family, entity and transaction described is invented.

Preregistration

H5's criterion — $N \geq e^{0.5/0.35} = 4.172734$ at $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ — was fixed in Section 2 before the measurement in Section 6 was run, so that it could not be adjusted to fit. **The criterion was not met, and that is reported in Section 8 and again in Section 13 in the same language a success would have received.** No other hypothesis in this paper was preregistered in any external registry.

Legal, tax and investment advice

This paper is analytical and is not legal, tax, or investment advice. It describes entity types, jurisdictions, listing processes and structural arrangements for the purpose of measuring concentration and resilience. No recommendation is made in respect of any particular person, asset, structure or jurisdiction. Statutory reach-back periods, reserved-share regimes, substance requirements, corporate-residence tests, foreign-ownership restrictions, exchange-regime classifications and listing criteria all change and vary by jurisdiction and by activity; every such statement in this paper is carried from the source manuscript, is flagged in Appendix B, and requires confirmation with qualified counsel in the relevant jurisdiction before any reliance.

Appendix A. The reporting discipline

Any structure measured under this instrument reports the following, and nothing less.

```
1  share basis, stated in one sentence
2  exposure table, jurisdictions and shares, summing to 100.0%
3  naive HHI and effective jurisdictions, with the rounding basis named
4  perimeter assignment with the four-question answers per jurisdiction
5  bloc-HHI and effective perimeters N
6  in-perimeter exposure for the largest perimeter, as a percentage
7  ln(N), eta ln(N), and h at the stated parameters
8  the trajectory record: classification, at least two dated events,
   direction, and reversibility
9  the DATE, and the diff against the previous measurement
```

Item 3's rounding basis is not pedantry. Section 6.3 shows that the same exposure map gives three different naive indices depending on how the shares are printed, and a measurement that does not say which basis it used cannot be reproduced. **Item 9 decays first and matters most**, because perimeter fragmentation is invisible except as a difference between two measurements taken a year apart.

For the disclosure audit the record is per element, not per structure: the element, its assigned job, each of the five criteria with the **evidence** rather than the conclusion, the product, and — for any element scoring zero — whether it was remediated or removed. **An element that cannot be brought to five of five is removed at the design stage. It is not documented better; it is taken out.**

Appendix B. Claims requiring verification against primary sources

The following rest on legal or regulatory detail that changes, and **none is cited to the bibliography, which contains no primary sources for any individual jurisdiction.** They are reproduced from the source manuscript and require confirmation with counsel in the relevant jurisdiction before any reliance.

1. The banking-secrecy statute of Section 4.1: its enactment date, its current text, the intergovernmental arrangements that overrode it, the number of accounts disclosed and the legal basis for that disclosure, and the dates of the criminal pleas and the sector disclosure programme.
2. The unilateral withholding regime of Sections 4.1 and 5.1: enactment date, the statute it formed part of, the withholding rate, the commencement date, and the model of intergovernmental agreement applicable in each relevant jurisdiction.
3. All listing and increased-monitoring events tabulated in Sections 4.2, 4.3 and 5.2, with their exact plenary and Council dates, the rating language used, and the criteria engaged. **These are the dates on which the propagation trace of Section 5.2 depends.**
4. The four-question perimeter answers for each named jurisdiction, and the bloc taxonomy of Section 5.5. These are analytical judgements applied to published characteristics; the underlying participation, membership and dependence facts should be confirmed jurisdiction by jurisdiction.
5. The asset-protection statute of Section 4.4 — its enactment and amendment history and its current section numbering — together with the two appellate decisions described, their reporter citations, their holdings, the lower-court decisions in the same matters, and the custody periods. **Case citation follows different conventions and no case is in the bibliography.** The bankruptcy look-back provision of the postscript requires the same confirmation.

6. Economic-substance regimes and their thresholds; the global minimum effective rate and its scope, thresholds and implementing law in each relevant jurisdiction; and corporate-residence and central-management-and-control tests.
7. The beneficial-ownership register regimes of Section 4.5, the supranational decision of November 2022 curtailing public access, the domestic constitutional challenge of Section 5.4 and the current scope of the statute it concerned. **That last item is actively litigated and any description is a description of its date.**
8. Foreign-ownership restrictions and the criminal treatment of nominee shareholding in the non-aligned operating jurisdiction of Structures C and D, including the section number, the current penalties and current enforcement practice; and the investment-promotion conditions conferring the ownership position, which are activity-specific and change.
9. The tax jurisprudence on captive arrangements without risk distribution referred to in Structure E, including reporter citations and holdings, and the current regulatory treatment of such arrangements as reportable transactions.
10. The sovereign investment institution of Section 7.6: its governing statute and date, the mandate instrument and issuing ministry, the date of first capital transfer, the exact scope and timing of the published holdings list, voting record and pre-announced voting intentions, and the instrument establishing its ethics body. Also the sector disclosure standard of 2008 — its title, date, sponsoring body and current membership — and the current major-shareholding notification thresholds in any market relied on, including their application to cash-settled derivative positions.
11. Correspondent-banking decline statistics, which are reported figures requiring confirmation against the current published data, and the tokenised-settlement regulatory regimes referred to in failure mode 6.
12. All cost figures in Section 7.2 are constructed illustrations. They are not fee quotations and no real fee schedule was used to produce them.

All numerical results in this paper were computed from the stated inputs at full precision and rounded only for printing.

Chapter 52

The Concentration Instrument: Bloc Collapse and the Effective Count of Jurisdictions

Abstract. An apparently diversified holding structure can be a single point of failure, and we give the instrument that detects it. Taking HHI as the sum of squared exposure shares and the effective count as its reciprocal, we measure a conventional four-jurisdiction structure — UK 40.9 per cent, US 27.3, Jersey 18.2, Luxembourg 13.6 — at a naive HHI of 0.293430, or 3.41 effective jurisdictions. Collapsing the same shares onto coordination blocs before squaring gives a bloc index of 1.000000 and an effective count of exactly 1.00. The redesigned structure across six jurisdictions reaches a naive HHI of 0.189002 (5.29 effective jurisdictions) and a bloc index of 0.340000 (2.94 effective blocs), with Western-perimeter exposure falling from 100.0 to 30.0 per cent. We show that the naive and bloc indices are one variance under two correlation assumptions, so that the general form $\mathrm{EFF} = 1/(w'Rw)$ is the similarity-sensitive diversity of order two, and that HHI^{-1} is exactly a Hill number of order two while the corpus's affine effective count $1 + (N-1)u$ is not and carries no external support. Coupling the count to the storage law $h = \eta \ln N - \delta Q - C$, the conventional structure earns a surplus term of exactly zero and $h = -0.500000$, while the designed structure earns 0.377583 and $h = -0.122417$, a 75.5 per cent reduction in shortfall — but the break-even count is $N = e^{0.5/0.35} = 4.17$, which the redesign does not reach, so that hypothesis is tested and not met. Robustness is stated both ways: a mis-assigned bloc moves N by 0.0169, under six-tenths of one per cent, while sustaining the conventional structure's naive reading would require believing the four Western jurisdictions correlate at only 6.6 per cent. The accompanying threat register prices expected annual loss at 9.25 per cent of protected capital and 5.97 per cent permanently.

Keywords: Herfindahl–Hirschman Index; effective number; bloc collapse; correlated exposure; single point of failure; threat register; concentration measurement

Series: D. Structure and Resilience **Paper:** CCP-052

1. Introduction

A holding structure distributed across four jurisdictions may nevertheless have one point of failure. This paper gives the instrument that decides which of the two descriptions is true,

states the conditions under which the instrument is exact, and reports what it measures on a conventional cross-border structure and on a redesigned alternative.

The question is not *how many jurisdictions is a structure in*, which is a count that any adviser can inflate by incorporating another company, but *how many independent points of failure does a structure have*, which is a measurement and cannot be inflated by adding entities. The distinction matters because the two answers can differ by a factor of three and because only the second one enters the surplus law that governs whether protected capital grows or erodes.

The instrument is not new. It is the Herfindahl–Hirschman Index, introduced by Hirschman [1945] as the square root of the sum of squared shares to describe the concentration of a country's foreign trade, arrived at independently in sum-of-squares form by Herfindahl [1950] in a dissertation on the steel industry, and later the subject of a priority dispute between them. Its definition fits on one line:

```
HHI = sum of squared shares            effective count = 1 / HHI
```

What this paper contributes is threefold.

First, the correction. The naive index overstates resilience because it treats units as independent when they share a coordination mechanism. The index is silent about correlation. Applied to market shares that silence is usually harmless, since rival firms mostly do fail independently. Applied to jurisdictional exposure it is fatal, because jurisdictions inside one coordination perimeter fail *together* — that is what a perimeter is — and an index that counts them separately is counting one failure four times. The correction is not novel as physics: it is the observation, stated for portfolios in 1952, that diversification benefit depends on covariance and not on count [Markowitz, 1952], and stated for coupled infrastructures in 2010, that a system of interdependent networks is far more fragile than either network alone [Buldyrev et al., 2010]. **Both borrowings are labelled at first use under the analogy discipline of Section 11.5.** The portfolio-variance result is used as a *derivation*: the exposure vector genuinely is a weighted sum of shocks with a correlation matrix, which is the premise Markowitz's algebra requires, and no further financial content is imported. The coupled-network result is used as a *shared structural form*: jurisdictions inside one enforcement perimeter have not been shown to satisfy the percolation premises of Buldyrev et al. [2010], and this paper claims only that the two systems share the property that coupling destroys the independence a naive count assumes. Perrow [1984] reached the organisational form of the same observation first — interactive complexity plus tight coupling makes catastrophic failure a property of the system rather than of any component — and is cited as **qualitative theory and framing, not as a formal result**. This paper's signature finding — that adding a fifth jurisdiction inside one perimeter can leave resilience unchanged or make it worse — is Perrow's result in a new measurement frame, and the paper says so rather than presenting it as new.

Second, the identification. The paper establishes that the naive index and the bloc-collapsed index are not two indices but one variance under two correlation assumptions, that the general form is $\mathrm{EFF} = 1/(w'Rw)$ for exposure weights w and a correlation matrix R, and that this general form is exactly the similarity-sensitive diversity of order two in the sense of Leinster and Cobbold [2012], whose additive cousin is Rao's quadratic entropy [Rao, 1982]. Where $R = I$ it collapses to HHI^{-1}, which is the Hill number of order two [Hill, 1973, Jost, 2006].

Third, the arithmetic and its verdict. The instrument is run on a conventional four-jurisdiction structure and on a redesigned six-jurisdiction one, coupled to the surplus law, and the verdict is reported in full including the part that does not flatter the redesign: the redesigned structure reaches an effective bloc count of 2.941176 against a break-even count of 4.172734, so its surplus balance remains negative at -0.122417. That is a hypothesis tested and not met, and Section 7 reports it in the same language it would have used had the criterion been cleared.

The paper presupposes the surplus functional and the affine effective count of Curzi [2026] — the functional $f(u) = \ln(1+(N-1)u)$, uniquely characterised by angle-only dependence, van-

ishing at parallel, an affine effective count and normalisation — and the roster-audit form of that count in Curzi [2026c], where a hundred-person department at $u = 0.05$ has an effective count of 5.95. It presupposes the reflexive loop $T(x) = W(x, M(x))$ of Curzi [2026], whose fixed point is the self-fulfilling prophecy, because the bloc assignment is itself a forecast that conditions the behaviour it forecasts. **That loop carries an identification problem this paper must state where it uses it rather than in a footnote.** Manski [1993] proves that a linear-in-means reflexive specification — a group's mean behaviour influencing its members while the members constitute the mean — is **not identified from observational data**: the endogenous social effect cannot be separated from contextual effects or from correlated unobservables, however much data is collected. **No amount of observation of jurisdictional behaviour could therefore confirm a bloc assignment that was derived reflexively.** The remedy adopted here is the fourth of the four known ones: **the reference group is stipulated and the stipulation is defended** — the bloc membership is fixed in advance by the perimeter test of Section 4.3, recorded with its evidence class, and stress-tested in Section 7.5 — which is the route Brock and Durlauf [2001] shows can be made rigorous. **The price is that the bloc assignment is a modelling position and not an identified empirical result, and this paper says so wherever the assignment enters.** It derives from the entity-map measurement of Curzi [2026], which reports naive HHI 0.293388 and 3.4085 effective jurisdictions against bloc-HHI 1.000000 and 1.0000 effective perimeters, a $3.41\times$ overstatement. **Section 12 records that the present paper's naive figure, 0.293430, and CCP-051's, 0.293388, differ in the fifth decimal, and states which is which.**

2. Hypotheses

Five hypotheses are stated. Each is falsifiable in the sense required by Platt [1964]'s strong inference; where a statement is not falsifiable it is labelled a definition or a modelling choice and is not counted as a hypothesis. Two of the five are reported in Section 7 as **not met**. Meehl [1967] is the relevant warning: directional predictions are cheap, and the defensible response is to make point predictions with constants attached, which is what H2 and H4 do.

H1 — the perimeter hypothesis. Several jurisdictions inside one enforcement perimeter behave as a single jurisdiction under a correlated shock. *Predicted:* the effective count $N = 1/(w'Rw)$ collapses toward 1 as the within-bloc correlation ρ rises toward 1, regardless of how large the naive count is; for the conventional structure specifically, N falls from 3.4080 at $\rho = 0$ to exactly 1.0000 at $\rho = 1$. *Falsified by:* an observed effective count that tracks the naive count under a correlated shock — equivalently, direct evidence that within-perimeter jurisdictions act independently — or by any structure in which subdividing inside a perimeter raises N.

H2 — the sufficiency hypothesis. Jurisdictional redesign alone can lift a structure to a non-negative surplus balance. *Criterion, fixed in advance:* with the surplus law $h = \eta \ln N - \delta Q - C$ at $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$, the balance is non-negative only if $\ln N \geq 0.5000/0.35 = 1.428571$, that is $N \geq e^{1.428571} = 4.172734$. *Falsified by:* a redesigned structure whose effective bloc count falls short of 4.172734 while C, δ and Q are held at the stated values.

H3 — the robustness hypothesis. The concentration measurement is insensitive to error in the assignment of a single member to a bloc, and sensitive to the bloc *structure* itself. *Predicted:* reassigning one contested jurisdiction moves N by less than one per cent, while merging two blocs moves it by more than twenty per cent. *Falsified by:* a single-member reassignment that moves N materially, which would make the published figure an artefact of a judgement call rather than a measurement.

H4 — the implied-correlation hypothesis. Reporting the naive index on a single-perimeter structure is equivalent to asserting a specific and implausibly low within-perimeter correlation, and that correlation can be computed rather than argued. *Predicted:* solving naive$+\rho(1-$naive$) =$ bloc for the two structures gives $\rho = (0.340000 - 0.293430)/(1 - 0.293430) = 0.065910$. *Falsified by:* evidence that the coordinated actions of the jurisdictions in question are in fact close to independent, which would make $\rho \approx 0.066$ a defensible rather than an untenable premise.

H5 — the specialisation hypothesis. Assigning each entity a single function lowers the maintenance term C relative to an entity carrying several functions. *Predicted:* three specialised entities cost less to maintain than one entity meeting every requirement simultaneously. *Falsified by:* a cost decomposition in which the overloaded entity is cheaper.

Not hypotheses, and labelled as such. Three further statements appear in this paper and none of them is falsifiable as stated; each is named here for what it is.

- *The logarithmic form of the input side*, $\eta \ln N$ rather than ηN, is a **modelling choice** carried from Curzi [2026], defended in Section 4.6 on the ground that increments shrink in proportion to what is already held. It is not tested here.
- *The identification of N with the effective bloc count* rather than with any other count is a **definition**. Its content is entirely in the bloc assignment rule, which is testable, and not in the identification.
- *The parameter set* $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ is a **modelling choice**, stated in advance so that H2's criterion is fixed before the measurement rather than after it, which is the preregistration discipline of Nosek et al. [2018] applied to an engineering measurement.

3. Method: the instrument, and which effective count it is

3.1 The numbers-equivalent property

Everything rests on one elementary property. If k units each hold share $1/k$:

```
HHI = k x (1/k)^2 = 1/k        1/HHI = k
```

so the reciprocal of the index returns the count when the units are equal, and when they are unequal it returns the count of *equal-sized* units that would produce the same index. That is the numbers-equivalent reading. **The source manuscript attributes this reading to an unnamed authority; the attribution is to Adelman [1969], and this paper names him.** The property itself is elementary and can be verified in two lines, which the Verification section does.

equal shares	share each	HHI (fraction)	HHI (points)	effective count
1	100.00%	1.000000	10,000.0	1
2	50.00%	0.500000	5,000.0	2
3	33.33%	0.333333	3,333.3	3
4	25.00%	0.250000	2,500.0	4
5	20.00%	0.200000	2,000.0	5
6	16.67%	0.166667	1,666.7	6
8	12.50%	0.125000	1,250.0	8

Two conventions are in use. Competition practice states HHI in points using percentage shares, so a perfectly concentrated market is 10,000 [U.S. Department of Justice and Federal Trade Commission, 2023, 2010]. The resilience application here states it as a fraction, so a perfectly concentrated structure is 1.0000. They are the same number divided by 10,000, and both are given wherever confusion is possible.

Three properties follow immediately and all three do work later.

First, HHI is a share-weighted average share. $\sum s_i^2 = \sum s_i \cdot s_i$, which is the mean share *experienced by a unit of exposure* rather than the mean share across units. Squaring is therefore not arbitrary: it weights each jurisdiction by how much exposure sits in it, which is the weighting a risk question requires.

Second, the largest share sets a floor under the index and a ceiling on the effective count. Since HHI $\geq s_{\max}^2$, the effective count is at most $s_{\max}^{-2}$:

```
conventional, largest UK 40.9%   HHI >= 0.167281   effective <=  5.9780
designed, largest UAE 29.1%      HHI >= 0.084681   effective <= 11.8090
designed blocs, largest 40.0%    HHI >= 0.160000   effective <=  6.2500
conventional blocs, WEST 100.0%  HHI >= 1.000000   effective <=  1.0000
```

The last row is the whole argument in one line. When one bloc holds everything, the ceiling on the effective count is one, and no activity below the bloc level can raise it.

Third, $s_{\max}^2 \le \text{HHI} \le s_{\max}$**.** The upper bound is attained when the largest unit holds everything. Both bounds are checkable in a line and both are useful sanity tests on a measurement done by hand.

No distributional or scaling claim is made anywhere in this paper, and that is stated rather than assumed because the material invites one. A share vector of 40.9 / 27.3 / 18.2 / 13.6 descends steeply and looks like the beginning of a heavy-tail argument, and the concentration literature is full of such arguments. **This paper fits no distribution, characterises no tail, and asserts no power law.** Clauset et al. [2009] is the reason: applied to twenty-four canonical data sets, a principled maximum-likelihood fit with a goodness-of-fit test and a likelihood-ratio comparison against alternatives leaves few of them standing, and **a straight line on a log-log plot is not evidence.** That procedure was not run here, because no claim requiring it is made. Stumpf and Porter [2012] adds the reason running it would not be enough on its own: even a genuine power law is weak evidence for a mechanism, because many mechanisms produce the same tail. **The effective count used here is a statistic of a finite exposure vector with four or six entries, not an estimate of a distribution, and it must not be restated as one.**

3.2 Which effective count — the Hill correspondence, stated exactly

The corpus runs two quantities under the name *effective count*, and they are not the same quantity. This section states which is which, because a reviewer who reaches for the diversity literature will find that it covers one of them exactly and the other not at all.

The reciprocal index is a Hill number of order two. Hill [1973] defined the family ${}^{q}D = (\sum_i p_i^q)^{1/(1-q)}$, whose members return the number of equally-abundant types that would give the observed value of the underlying index. At $q = 2$ this is $(\sum p_i^2)^{-1}$, which is HHI^{-1} exactly. Jost [2006] is the statement that entropies are not diversities and that their numbers-equivalents are. The same arithmetic was reached independently in ecology by Simpson [1949] in the same decade as Herfindahl's dissertation, and again in political science by Laakso and Taagepera [1979] as the effective number of parties. **Independent reinvention in three disciplines is the strongest available argument that the instrument is general and not a borrowing.** The exponential of Shannon entropy [Shannon, 1948] is the same family at $q = 1$; the reciprocal index is also the exponential of the Rényi entropy of order two [Rényi, 1961], which is where the participation-ratio name comes from.

The affine count is not a Hill number and has no external support. The corpus's other effective count, $N_{\text{eff}} = 1 + (N - 1)u$, used in Curzi [2026], Curzi [2026c] and Curzi [2026o] — the last of these badging the affine form as design axiom S3 rather than as a theorem — is a **linear interpolation between** 1 **and** N. It is not a Hill number of any order, it is not a numbers-equivalent, and no literature in ecology, economics, political science or information theory contains it. The arithmetic that separates the two is given in the Verification section and can be redone in a line: for two units holding 0.9 and 0.1, every Hill number reports less than 1.39, while the affine count at $u = 1$ reports exactly 2.00.

The collapse result has a formal ancestry that should be named rather than reinvented. A small correlated shift producing a discontinuous collective outcome is the threshold and tipping mechanism: Granovetter [1978] establishes that the distribution of individual thresholds, not the mean, determines whether a collective outcome tips, and Schelling [1971] is the canonical demonstration that a mild individual preference produces a sharp aggregate segregation — **the strongest available warning against reading an individual-level parameter straight into an aggregate, which is precisely what a bloc assignment does.** Castellano et al.

[2000] supplies the closest formal analogue of a bloc: the same interaction that raises local agreement freezes global diversity into non-interacting regions, with a sharp transition between the monocultural and fragmented phases. **A bloc, in that vocabulary, is a frozen region, and the effective count is a count of them.** And on the generating side, Bouchaud and Mézard [2000] derives concentration rather than observing it: multiplicative growth plus exchange produces a heavy tail, and above a critical ratio a finite fraction of the total condenses onto one holder — **which is the index going to one, obtained from a mechanism.** This paper does not use that model and asserts no tail; the citation records that a mechanism-side companion to the arithmetic exists.

This paper therefore uses HHI^{-1} throughout and cites the Hill literature only for that quantity. Where the affine count is referred to, it is referred to by name and marked as the corpus's own construction. The two are related — both are called effective counts, both return N in the fully-independent limit — and they are not interchangeable, and the corpus has in places let them travel under one name.

The general form, and its exact literature. The bloc collapse of Section 3.3 generalises HHI^{-1} to

$$\mathrm{EFF} = \frac{1}{w' R w}$$

for exposure weights w and correlation matrix R. This is *not* the affine count; it is the similarity-sensitive diversity of order two of Leinster and Cobbold [2012], whose order-two member is $1/(p'Zp)$ for a similarity matrix Z, and whose additive relative is the quadratic entropy of Rao [1982]. Setting $R = I$ recovers HHI^{-1}; setting R to the block matrix of ones recovers the bloc index. **The correspondence is exact where R has unit diagonal and non-negative entries, which is the regime this paper works in; it lapses where correlations are negative, and this paper claims nothing there.** Cluster 02's own sceptic should be read alongside: Hurlbert [1971] argues that "diversity" has been used for so many distinct quantities that the word carries no content without its formula attached, which is precisely the discipline this section is imposing.

Small-sample estimation is a live exposure and is not addressed by the corpus. Chao et al. [2014] show that diversity estimators of low order are biased on small samples; a six-jurisdiction exposure vector is a small sample, and Section 12 records this as an unclosed gap.

3.3 The bloc collapse as a variance identity

Grouping data until the answer improves is a recognised way to lie, so the grouping needs a foundation that is not assertion. It has one, and it is exact.

Model the exposure as a portfolio, following Markowitz [1952]. Let w_i be the share of protected capital exposed to jurisdiction i, and let X_i be a unit-variance shock to that jurisdiction — a designation, a freeze, an exchange control, a de-risking event. The aggregate shock is $R = \sum_i w_i X_i$ and its variance is

$$\mathrm{Var}(R) = \sum_i \sum_j w_i w_j \rho_{ij}.$$

Two extreme correlation assumptions bracket everything.

If every jurisdiction is independent — $\rho_{ij} = 0$ for $i \neq j$, $\rho_{ii} = 1$ — the double sum collapses to $\sum_i w_i^2$, which **is the naive HHI, exactly**. Verified by explicit double sum on the designed structure's six shares: 0.189002, matching the naive index to twelve decimal places.

If jurisdictions inside a bloc are perfectly correlated and blocs are independent — $\rho_{ij} = 1$ when i and j share a bloc and 0 otherwise — the double sum collapses to $\sum_b W_b^2$ where W_b is the bloc's total share. That **is the bloc-HHI, exactly**. Verified the same way on the same six shares: 0.340000, matching to twelve decimal places.

> **The naive HHI and the bloc-HHI are not two different indices. They are the same variance under two different correlation assumptions, and $\mathrm{Var}(R)^{-1}$ is the effective number of independent shocks the structure is exposed to.** Choosing between them is not a modelling preference; it is a claim about the world, and it should be argued rather than assumed.

The two are therefore endpoints of a continuous family. With within-bloc correlation ρ and independence across blocs,

$$\mathrm{Var}(R) = \text{naive} + \rho\,(\text{bloc} - \text{naive}).$$

That the intermediate cases matter is not a theoretical nicety. Measured correlations rise under stress: Longin and Solnik [2001] find extreme correlation in equity markets that is asymmetric and concentrated in downturns. **The effective count falls exactly when it is needed**, which is a stronger statement than anything the static index says. The ecological name for the same variable is synchrony [Loreau and de Mazancourt, 2008], and the ecological demonstration that effective count buys stability is Schindler et al. [2010]'s Bristol Bay sockeye record, where population diversity produces roughly half the variability and about a tenth of the fishery closures.

The contrary literature must be read alongside and is not decorative. Doak et al. [1998] show that diversity–stability correlations can arise from statistical averaging alone, with no mechanism, which means a portfolio-arithmetic result is not by itself evidence of a mechanism. May [1972] showed that increasing diversity and connectance *destabilises* randomly assembled systems. This paper's claim is narrower than the general diversity–stability claim and is not exposed to those results in the same way: it asserts only that the reciprocal of a variance is an effective count of independent shocks, which is arithmetic, and that jurisdictions in one perimeter are correlated, which is H1.

3.4 Defining a bloc without circularity

The variance derivation is only as good as the bloc assignment, so the assignment needs a rule that someone who disagrees with the analyst can apply and still reach a reproducible answer. The rule is the perimeter test of Section 5.3: two jurisdictions are in the same bloc when a single decision taken in one reliably produces the corresponding effect in the other without a further independent decision. Evidence is ranked:

weight	evidence class
strongest	demonstrated near-simultaneous mirroring of designations or listings
strongest	shared correspondent networks that exit corridors together
strongest	a supervisory standard producing the same domestic consequence in every member
strongest	binding accession to a common enforcement instrument
weaker	participation in an information-exchange framework
weaker	treaty density
weakest	diplomatic alignment, voting patterns, press characterisation

The instruments that create genuine coordination are documented and citable. Recognition and enforcement of awards across states runs through United Nations [1958], whose

near-universal adherence is what makes a cross-border enforcement perimeter real rather than rhetorical. Automatic exchange of account information [Organisation for Economic Co-operation and Development, 2017], the anti-money-laundering recommendations that supervisors implement in common [Financial Action Task Force, 2012], and the substance and treaty-abuse work of the base-erosion project [Organisation for Economic Co-operation and Development, 2015] are the mechanisms by which nominally separate jurisdictions adopt the same domestic consequence. de Goede [2011] documents the listing machinery itself as a technique of governance, which is the mechanism the perimeter test is trying to detect. Recognition is emphatically *not* general everywhere, and Hague Conference on Private International Law [1989] is worth citing precisely **because it never entered into force**, which is why forced heirship remains a conflict-of-laws problem rather than a solved one; the succession regulation [European Union, 2012] supplies habitual residence plus *professio juris* for the states bound by it.

Assignments for the designed structure, with the reasoning compressed:

jurisdiction	bloc	reasoning
UAE	BRIDGE	dual-rail access; non-participation in the Western designation apparatus; own listing decisions
Singapore	BRIDGE	same test, different geography
Thailand	NONALIGNED	no accession to the coordinated enforcement instruments; domestic-priority policy
India	NONALIGNED	same
UK	WEST	core member of the mechanism
Bermuda	W-lite	overseas territory; conservatively folded into WEST

Two of these are contestable. The response is not to argue but to run the arithmetic both ways, which Section 7.5 does.

3.5 The measurement procedure

Six steps, in order, with every input written down, because next year's measurement is only meaningful as a difference against this year's.

Step 1 — Fix the share basis. The shares must be shares *of something*, and the something must match the threat. For arbitrary state action the correct basis is the fraction of protected capital whose *availability* depends on a decision taken inside that jurisdiction's perimeter. Not book value; not tax residence; not where the certificate is. A UK company holding a Bermudian portfolio managed from Singapore contributes to the UK share to the extent that a UK decision can stop the portfolio from being instructed. Common alternative bases — revenue by source, net asset value by situs, headcount, tax paid — are legitimate for other questions and wrong for this one. Pick one, state it, and never mix two inside one measurement.

Step 2 — Enumerate exposures and normalise to 100%. Exposures must be exhaustive and mutually exclusive. If an asset is exposed to two perimeters, split it and record how. If that cannot be said, assign it wholly to the more capable perimeter, which is the conservative choice.

Step 3 — Compute naive HHI and the effective jurisdiction count. Square, sum, invert.

Step 4 — Assign blocs using the perimeter test, and record the reasoning per jurisdiction. This is the only judgemental step and it must be written down or the measurement is not reproducible.

Step 5 — Compute bloc-HHI and N. Sum shares within blocs first, then square, then sum, then invert. Order matters and reversing it silently reproduces the naive number.

Step 6 — Report the pair, never one alone. A single figure invites the reader to assume the convenient correlation. Report naive, bloc, and the gap; if the gap is large, state what would have to be true for the naive figure to be the right one.

3.6 Why a logarithm, and why not a count

The identification of N with the effective bloc count is new to this paper; the functional form is carried from Curzi [2026]. Restating the argument in this context is worthwhile because a great deal follows from it.

Suppose the input side of the surplus law were linear in N — ηN rather than $\eta \ln N$. Then two independent blocs would be worth exactly twice one and the optimal structure would hold as many blocs as the budget allows. The first genuinely independent store removes an entire failure mode; the second removes the residual probability that the first is wrong; the third removes the residual of the residual. Each addition addresses a smaller remaining exposure, and the natural functional form for a quantity whose increments shrink in proportion to what is already held is the logarithm.

There is a second reading which is more useful operationally. $\ln N$ is, up to a constant, the entropy of a uniform distribution over N outcomes [Shannon, 1948]. What the surplus term measures is how much *uncertainty about which perimeter matters* the structure has purchased. A structure inside one perimeter has zero entropy: an adversary knows exactly which decision reaches everything. A structure across three roughly equal perimeters has $\ln 3 = 1.0986$ of it. **The quantity being maximised is not the number of places a structure is in; it is the number of independent decisions that would have to go against it simultaneously.**

That reading also explains why the measure is indifferent to entity count. Adding a fifth Western entity adds no uncertainty about which decision matters, because the answer is still the Western one. The entropy is unchanged, the surplus term is unchanged, and the fee is not. This is a specific instance of a general organisational result: adding units inside a coordination boundary adds communication cost at order $n(n-1)/2$ [Brooks, 1975] and adds nothing to independence.

The functional form is not tested here and is labelled a modelling choice in Section 2. Its fixed-point behaviour under the storage law is a first-order linear difference equation and is standard [Elaydi, 2005]; Curzi [2026] states the contraction result exactly, with modulus $1-\delta$, a unique start-independent fixed point, and a maximum residual of 2.0624×10^{-15} over a hundred passes, and Curzi [2026] gives the fixed-point identity $h_\infty = \eta(\ln N - F(u^*))$ in which δ, C and Q cancel.

3.7 The four axes

Jurisdiction is the axis this paper measures and it is one of four. The instrument is identical on each; only the share basis changes.

Settlement rail. The exposure is the fraction of settlement volume depending on a given messaging or clearing system. One fact ruins naive measurements: a clearing system and a messaging system are different layers, and a structure reporting two rails may have one messaging dependency. Measured correctly, “one messaging network at 100%” and “85% on that network plus 15% on a second clearing system whose messaging leg runs over the first” are the same number, and only the second looks diversified.

Counterparty. The exposure is the fraction of protected capital held as a claim on, or custodied by, a given institution. This is the axis on which jurisdictional diversification most reliably backfires, because the population of institutions willing to bank a genuine cross-bloc structure is small. A structure at $N = 2.94$ on jurisdiction and at an effective counterparty count of 1.6 has not reduced its total risk; it has moved it. Curzi [2026] makes the same measurement on a different share basis and reaches the same shape: Delaware 40.0 / South Dakota 25.0 / Nevada 20.0 / Wyoming 15.0 gives naive HHI 0.285000 and 3.508772 effective jurisdictions, but bloc-HHI 1.000000, $N = 1.00$ and a surplus term of exactly zero.

Currency. The exposure is the fraction of the balance sheet denominated in a given unit. Currency exposure interacts with the jurisdiction axis rather than being independent of it: a currency's clearing runs through banks that sit inside a perimeter, so a balance sheet wholly in one currency carries that perimeter's exposure however the entities are arranged.

A structure is only as diversified as its most concentrated axis, and the axis most often measured is the one the analyst was already comfortable with. The reporting discipline is to run all four and print all four.

4. The threat model that fixes the share basis

An index without a stated threat is an index without a share basis. This section establishes the classification that fixes it. Five failure processes are distinguished; they are not variations on a theme, they have different actors, triggers, time constants and reversibility profiles, and **a structure that answers one of them well may answer another worse than doing nothing at all**.

#	threat	actor
1	arbitrary state action	executive or administrative organ, or a regulated intermediary acting in its shadow
2	lawful process	a private claimant, tax authority, insolvency officeholder, or regulator acting judicially
3	commercial / counterparty	nobody; no one intends the loss
4	succession	mortality
5	bloc fragmentation	history

The corpus's storage argument establishes that a storage problem is only well posed once the specific degradation process is named. A grain store designed against weevils is not a grain store designed against damp, and a store built against damp when the threat was weevils is worse than an open crib, because it cost money, consumed attention, and produced confidence proportional to its expense rather than to its fit. **Capital is stored under the same rule, and failing to name the degradation process is the most expensive error in the field.**

Arbitrary state action is this paper's subject, and that is a choice which excludes things. It excludes lawful process in particular, because the two threats have opposite structural signatures. Against arbitrary action the useful property of a structure is that it sits outside a coordination perimeter, so that a decision taken inside that perimeter does not reach it. Against lawful process that property is close to worthless and frequently harmful, because the enforcement mechanism does not travel to the asset — it travels to the person, and the person is inside the perimeter. A statute refusing to recognise foreign judgments can perform exactly as advertised while the principals are compelled anyway, because the court does not need

the statute's cooperation; it needs theirs. The grounds on which enforcement is in fact refused, and the failure modes of the most-relied-upon instrument, are a settled subject in the arbitration literature and are listed in Appendix B as requiring primary-source verification.

4.1 The five processes

Threat 1 — arbitrary state action. *Trigger:* a designation, a listing, an exchange-control measure, a nationalisation, a licence revocation, a change of risk appetite in a compliance department. None of these requires a court. *Mechanism of loss:* the asset does not move and title is often untouched; what is removed is the ability to instruct, and since the practical content of ownership is the ability to instruct, the distinction is academic to the owner. The 2022 immobilisation of approximately USD 300bn of central bank reserves is the reference event, and its instructive feature is that nothing was seized in the older sense — the reserves remained booked where they were and the instructions stopped being executed. *Warning window:* days at best, hours in practice, and often the first notice is the failed instruction. Structures cannot be built inside this window, which is the single most important operational fact about this threat. *Reversibility:* low, and slow where it exists. *What resists it:* the asset, the entity and the settlement rail sitting outside the coordination perimeter that produced the decision — not outside *a* jurisdiction but outside the *mechanism*. *What does not resist it, and is routinely sold as though it does:* secrecy, nominee ownership, layered holding chains within a single perimeter, and any structure whose defence is that the authority will not find out. **A designation regime does not need to find the owner; it needs to find the owner's bank, and the bank is not hiding.**

There is a second-order form that does more cumulative damage than the headline form: de-risking. A correspondent bank facing a compliance cost it cannot price exits a corridor. Nobody is designated, no measure is announced, and an entire class of legitimate counterparty finds the rail closed. This is arbitrary in the precise sense that matters — action without process, notice or appeal — even though no state formally took it. The answer is the same as for the headline form: redundancy across mechanisms that do not share a decision-maker.

Threat 2 — lawful process. *Trigger:* a claim filed, a judgment entered, an assessment raised, a petition presented. *Mechanism of loss:* an order, and behind the order the state's willingness to compel a person. Where the court has territorial reach it may attach the asset; where it does not, it attaches the person — freezing injunctions, disclosure orders, repatriation orders, and ultimately contempt. *Warning window:* long; months to years. *Reversibility:* partial, and frequently the most recoverable of the five. *What resists it:* having built earlier. The governing law is fraudulent-transfer law and fraudulent-transfer law is a clock. **Threat 2 is a timing problem wearing a jurisdiction costume, and the jurisdiction is not the part that matters.** The statutory reach-back periods and the authorities are set out in Appendix B; **they are described from the source manuscript and are not cited to the bibliography, which contains no primary statutory or case material** (Section 13).

Threat 3 — commercial and counterparty. *Trigger:* ordinary business risk, usually correlated with a cycle. *Mechanism of loss:* becoming an unsecured creditor of an entity that cannot pay, or a secured creditor whose collateral turns out to be someone else's. *Warning window:* weeks to months, and asymmetric — the people who see it first are inside. *Reversibility:* highest of the five. *What resists it:* segregation, priority and risk transfer. Asset partitioning is the irreducible function of entity law [Hansmann and Kraakman, 2000], and that is what a bankruptcy-remote vehicle actually does; Pistor [2019] is both the best scholarly account of how legal coding creates capital and the sharpest critique of it. *What does not resist it:* jurisdiction. A creditor in one financial centre and a creditor in another get the same haircut from the same failed obligor, which is exactly why a structure optimised only for threat 1 can be quietly terrible here.

Threat 4 — succession. *Trigger:* death or incapacity; also divorce, disqualification, or cognitive decline that nobody wants to name. *Mechanism of loss:* fragmentation. Assets that functioned as a system are divided into fractions that do not; control that was unified becomes

contested; and where heirs are resident in different legal families the same estate is characterised differently by each and the characterisations do not reconcile. *Warning window:* zero — the defining feature, and under-weighted everywhere. **Threat 4 is the only threat in the register whose warning window is exactly nil and whose probability over a long enough horizon is exactly one.** *Reversibility:* low. *What resists it:* an entity with separate legal personality owning the assets, so that the succession event changes the entity's governance rather than the assets' ownership. The comparative literature on the trust and the foundation as alternative devices for exactly this problem is comparative and is described in Appendix B. A sketch of the arithmetic, to fix scale: on a constructed estate where one heir sits in a forced-heirship jurisdiction whose rules reach 25% of the assets, and where a three-child reserve of three-quarters applies to that portion, the reserved pool is $0.25 \times 0.75 = 0.1875$ of the estate. **18.75% of the estate has a destination the founder cannot direct, and that figure is arithmetic, not advocacy.** The absence of a succession convention in force [Hague Conference on Private International Law, 1989] is the structural reason the problem persists.

Threat 5 — bloc fragmentation. *Trigger:* realignment. A neutral jurisdiction picks a side; a bridge centre is absorbed into one of the things it was bridging; a coordination mechanism expands its membership. Less dramatically: a jurisdiction signs an information-exchange agreement, adopts a substance regime, or accedes to a listing framework, and quietly becomes a member of a perimeter it was outside. *Mechanism of loss:* the effective bloc count falls. Nothing visible happens to any entity and the structure looks identical on the org chart; what changed is the correlation between its parts, and correlation is not on the org chart. *Warning window:* long — years — but only for someone watching. *Reversibility:* essentially nil at the bloc level. *What resists it:* designed redundancy, exit optionality written into the documents rather than improvised later, and monitoring. The arithmetic is stark, using the designed structure's blocs of 40 / 30 / 30:

```
designed, three blocs 40/30/30          N = 2.9412    h = -0.1224
BRIDGE absorbed into WEST (70/30)       N = 1.7241    h = -0.3093
everything absorbed into one perimeter  N = 1.0000    h = -0.5000
```

One realignment event costs 1.22 effective blocs and 0.19 of surplus, and no entity has to do anything wrong for it to happen. The mirror case is worth noting: if the BRIDGE bloc were to split — if the two bridge jurisdictions ceased to correlate — N rises from 2.9412 to 3.6158. Threat 5 runs both ways and only one direction is planned for.

4.2 Three pairs that are routinely conflated

Threats 1 and 5 look alike because both are geopolitical. They are not alike operationally. Threat 1 is an event: a decision is taken and something stops working on a Wednesday. Threat 5 is a drift: nothing stops working, and the correlation between holdings rises over years until a structure that used to be three things is one thing. A structure built only against threat 1 tends to be static — correct on the day it was built and never revisited — and static is exactly the property threat 5 punishes. Conversely, a portfolio of options and exit rights with no actual holdings outside the perimeter is nimble against a drift and useless against a Wednesday. The organisational literature on learning gives the same shape: fast convergence on a single mode destroys the variance that produced the advantage [March, 1991].

Threats 2 and 3 look alike because both end in a courtroom. The difference is who wants what. In threat 2 a claimant wants the assets and is using the court to reach them; the defensive question is *reach*. In threat 3 an insolvency officeholder is distributing a deficient pool and the defensive question is *priority*. Ring-fencing improves priority and does nothing for reach; timing improves reach and does nothing for priority. A structure sold as asset protection without specifying which of the two it improves is unspecified.

Threats 3 and 1 look alike when the counterparty is a bank. When a bank declines to process an instruction, the owner experiences the same thing — an instruction that does not

execute — whether the cause is insolvency, risk appetite, or a state measure. The remedies differ completely. The useful diagnostic is whether other customers of the same bank in the same corridor are also affected, and whether customers of other banks in the same corridor are affected. Same bank only: threat 3. Same corridor across banks: threat 1.

4.3 The perimeter test, stated operationally

Threat 1 turns entirely on the identity of the deciding perimeter, so the test has to be operational rather than atmospheric. Two jurisdictions are inside the same perimeter when a single decision taken in one of them reliably produces the corresponding effect in the other **without a further independent decision**. That is a strong condition and it is deliberately stronger than alliance or cooperation.

Membership of an information-exchange framework is evidence, not proof: exchanging account data [Organisation for Economic Co-operation and Development, 2017] is a long way short of executing another state's freeze. What converts evidence into proof is a demonstrated pattern of automatic or near-automatic transmission — a listing framework whose designations are mirrored within days [de Goede, 2011], a correspondent network whose members exit a corridor together, a supervisory standard producing the same domestic consequence in every member [Financial Action Task Force, 2012]. The requirement is only that the analyst make the call and record it, because the bloc assignment is an input to the instrument and an undocumented input is not a measurement.

Where the call is genuinely uncertain the discipline is to run the arithmetic both ways and see whether the answer changes. Section 7.5 does exactly that, and the robustness result there is more interesting than the point estimate.

The governance literature supplies an independent name for what is being counted. Ostrom [2010] develops polycentricity — multiple centres of decision-making that are formally independent of one another — as an empirically derived property of durable governance systems. **Polycentricity is the effective count of jurisdictions in governance language**, arrived at from field data rather than from an index, and the convergence is worth stating.

4.4 The threat register

A threat register is a table with one row per threat and six fields per row, populated with numbers the analyst is willing to defend, reviewed on a fixed cadence, and used to allocate budget — **because a register that does not change any decision is a document rather than an instrument**. The fields are: p, the annual probability of the threat materialising; L, the loss as a fraction of protected capital Q given that it materialises; r, the fraction of L eventually recovered; W, the warning window in days from first credible signal to loss; and B, the build time in days for the structure that answers it.

Every number in the register below is CONSTRUCTED and illustrative. It is not an estimate of any real position and the parameters should not be imported. The value is in the discipline of writing numbers down.

threat	p	L	r	W (d)	B (d)
1 arbitrary state action	0.020	0.60	0.05	2	210
2 lawful process	0.080	0.35	0.40	540	210
3 commercial / counterparty	0.120	0.20	0.55	90	210
4 succession	0.030	0.45	0.30	0	210

threat	p	L	r	W (d)	B (d)
5 bloc fragmentation	0.050	0.30	0.25	365	210

First pass: expected loss. Expected annual loss is pL, summed:

```
1 arbitrary state action    0.020 x 0.60 = 0.0120
2 lawful process            0.080 x 0.35 = 0.0280
3 commercial/counterparty   0.120 x 0.20 = 0.0240
4 succession                0.030 x 0.45 = 0.0135
5 bloc fragmentation        0.050 x 0.30 = 0.0150
------------------------------------------------
TOTAL                                    0.0925   = 9.25% of Q per year
```

Treating the five as independent, the probability that none materialises in a given year is the product of $(1 - p)$:

```
P(no threat in one year)     = 0.731125    (73.11%)
P(at least one)              = 0.268875    (26.89%)
over  5 years, at least one  = 0.791090    (79.11%)
over 10 years, at least one  = 0.956357    (95.64%)
over 20 years, at least one  = 0.998095    (99.81%)
over 30 years, at least one  = 0.999917    (99.99%)
```

Over a generation, at these parameters, something in the register happens with probability 0.9999. The question a structure answers is never whether something will happen; it is which something, and whether the structure bends or breaks when it does.

Second pass: the recoverability correction. Permanent expected loss is $pL(1 - r)$:

threat	EL	permanent EL	share of permanent
1 arbitrary state action	0.0120	0.0114	19.10%
2 lawful process	0.0280	0.0168	28.14%
3 commercial / counterparty	0.0240	0.0108	18.09%
4 succession	0.0135	0.0095	15.83%
5 bloc fragmentation	0.0150	0.0112	18.84%
total	**0.0925**	**0.0597**	

Threat 1 is 12.97% of expected loss but 19.10% of *permanent* expected loss. That is a real correction and it is not large enough to carry the argument alone. The argument is carried by the conditional numbers:

```
conditional PERMANENT loss, given the event occurs
1 arbitrary state action    57.00% of Q     -> Q multiplier 0.4300
2 lawful process            21.00%          -> 0.7900
3 commercial/counterparty    9.00%          -> 0.9100
4 succession                31.50%          -> 0.6850
5 bloc fragmentation        22.50%          -> 0.7750
```

A single threat-1 event at these parameters takes protected capital from 10.00 to 4.30, after which ordinary erosion continues from the new base:

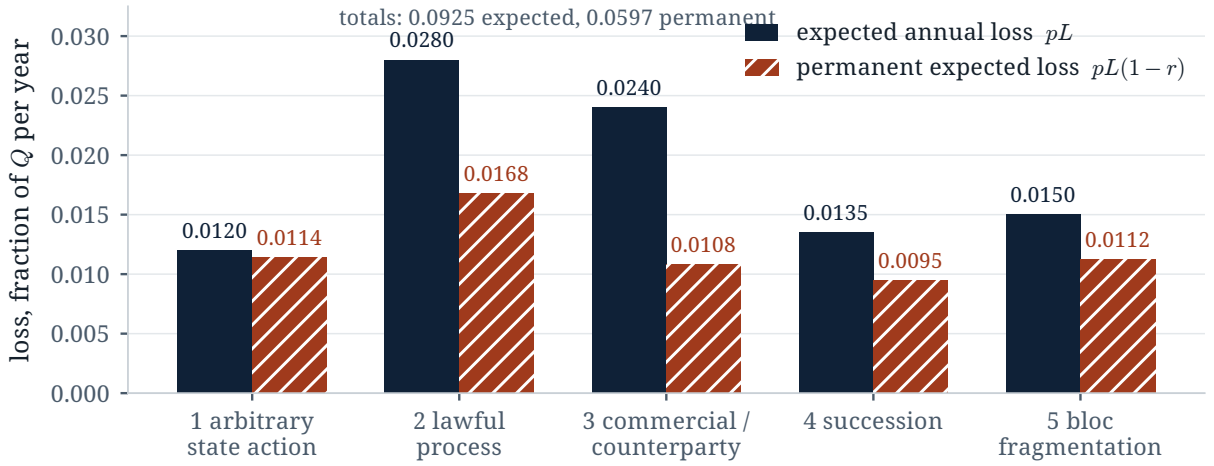


Figure 52.1: **The register's ranking changes when losses are weighted by what does not come back: threat 1 is 12.97% of expected annual loss but 19.10% of permanent expected loss.** Modelled from the CONSTRUCTED register parameters above; script `figs/src/CCP-052_fig05.py`.

```
t = 0    Q = 4.3000   (unhit 10.0000)
t = 10   Q = 3.5134   (unhit  8.1707)
t = 20   Q = 2.8707   (unhit  6.6761)
t = 30   Q = 2.3456   (unhit  5.4548)
```

Erosion is multiplicative, so a single large permanent loss is not a one-off deduction — it rescales every subsequent year. This is the formal reason a low-probability, low-reversibility threat justifies disproportionate structural attention, and it is why the register's expected-value column is the wrong place to stop.

The general form of that argument is not this paper's and should be credited. Peters [2019] and Peters and Gell-Mann [2016] establish that the time average of a multiplicative process differs from its ensemble average, so that an expectation computed across many holders does not describe what happens to any one holder over time — which is exactly the divergence between the expected-loss column and the conditional permanent-loss column above. **The register's expected value is an ensemble average and the principal lives a time average.** The same literature supplies an independent warrant for the logarithm in the surplus law: the growth rate of a multiplicative process is logarithmic in the stock, with no psychological assumption required. **Section 3.6 justifies the logarithm on the diminishing-returns argument the framework already makes; this is the stronger warrant, and where the two disagree the multiplicative one should be preferred because it does not depend on a perceptual story being true.**

Third pass: misallocation. Suppose a fixed structural budget reduces the loss from exactly one threat by 75%. Which?

```
unmitigated expected annual loss              0.0925

budget aimed at arbitrary state action    EL = 0.0835    reduction   9.73%
budget aimed at lawful process            EL = 0.0715    reduction  22.70%
budget aimed at commercial/counterparty   EL = 0.0745    reduction  19.46%
budget aimed at succession                EL = 0.0824    reduction  10.95%
budget aimed at bloc fragmentation        EL = 0.0813    reduction  12.16%

spread evenly across all five at 15% each EL = 0.0786    reduction  15.00%
```

On this rule the best single target is threat 2 and the worst is threat 1 — the subject of

this paper. **That is stated plainly rather than hidden: on naive expected value, the threat this paper addresses is the least attractive of the five to spend money on.** It becomes the right target only when weighted by permanence, by the multiplicative effect on all future years, and by the window. A reader who does not accept those three corrections should not accept the priorities that follow from them.

Fourth pass: the window, which is decisive. Take a constructed build time of 210 days to stand up a cross-bloc structure end to end.

threat	W (d)	B (d)	verdict	slack
1 arbitrary state action	2	210	too late	-208
2 lawful process	540	210	in time	$+330$
3 commercial / counterparty	90	210	too late	-120
4 succession	0	210	too late	-210
5 bloc fragmentation	365	210	in time	$+155$

Three of five cannot be answered after the signal. Threat 1's coverage ratio is $2/210 = 0.0095$: **under one per cent of the time that would be needed.** Threat 4's is exactly zero. The two that appear answerable in time are not straightforwardly so. Threat 2's 330 days of apparent slack is illusory, because building inside the window is precisely what fraudulent-transfer law exists to reach. Threat 5's 155 days is real but requires that somebody was watching the trajectory for the preceding two years, which is an organisational property rather than a structural one.

> **The only window common to all five threats is *before*.** Structures are built in peacetime or they are not built. Every design result in this paper assumes a builder not currently under threat, because a builder under threat has fewer options than the analysis describes.

4.5 Assignment, and the disclosure test applied to the register

threat	resists it	does not
1 arbitrary state action	bloc diversity; rail redundancy; neutral-law entities; real substance	secrecy; layering in-bloc; better contracts
2 lawful process	time; pre-claim solvency; disclosed, defensible transfers	non-recognition statutes alone; distance
3 commercial / counterparty	segregation and bankruptcy-remote vehicles; risk transfer; counterparty diversity; netting and collateral	jurisdiction; bloc diversity
4 succession	separate legal personality; a foundation-type holding layer; law-family selection	wills alone; heir goodwill
5 bloc fragmentation	designed redundancy; exit optionality; trajectory monitoring	a single bridge; static structures

Read the right-hand column as carefully as the left; the right-hand column is where the money goes.

Each row must be audited against the disclosure test: would the element still do its job if every part of it were disclosed to every authority with jurisdiction over any part of it? Rows 3, 4 and 5 pass without argument. Row 1 passes by construction, because the mechanism is that lawful entities in non-coordinating jurisdictions are not reachable by a single decision — a fact disclosure does not change. **Disclosure does not repair a broken perimeter, which is the point: the defence is structural, not informational.** This is the open-design principle of Kerckhoffs [1883], whose engineering content is that a system's security must not depend on the secrecy of its design, restated in Saltzer and Schroeder [1975] as *open design*; Curzi [2026] carries the constraint in full, establishing that channel separation is strengthened rather than weakened by disclosure because its effectiveness is entirely independent of whether anyone knows it is in place. Row 2 is where the test bites: a response to lawful process that depends on a claimant not knowing where the assets are is outside this paper's scope and is treated as a failure mode rather than a technique.

5. Results: the conventional structure

CONSTRUCTED CASE. The principal is a dual national, tax resident in one Gulf jurisdiction, with three heirs resident in three different jurisdictions, one of them under a forced-heirship regime. Assets: manufacturing operations in Thailand; a trading company in the UAE; an IP portfolio; commercial real estate in London; a listed-securities portfolio. Revenue arrives in five currencies. All facts are invented.

The share basis is stated first, as Step 1 requires: *the fraction of protected capital whose availability depends on a decision taken inside a given jurisdiction's coordination perimeter*. Two consequences follow and both are stated in advance so that a reader may disagree with the inputs rather than with the method. The London real estate contributes its full value to the UK exposure, because it cannot be relocated and UK measures reach it directly. The listed portfolio contributes to the perimeter of its *custodian*, not to the perimeters of the issuers — a portfolio of Asian equities held through a Western custodian is Western exposure for this purpose, and this is the single input most often assigned wrongly.

jurisdiction	share	share^2	bloc	contents
UK	40.9%	0.167281	WEST	holdco and London property
US	27.3%	0.074529	WEST	listed portfolio via US custodian
Jersey	18.2%	0.033124	WEST	IP holding
Luxembourg	13.6%	0.018496	WEST	treaty vehicle
total	**100.0%**			

```
naive HHI = 0.167281 + 0.074529 + 0.033124 + 0.018496 = 0.293430
effective jurisdictions = 1 / 0.293430 = 3.407968  ->  3.41

bloc collapse:  WEST = 40.9 + 27.3 + 18.2 + 13.6 = 100.0%
bloc-HHI  = 1.000000^2 = 1.000000
EFFECTIVE BLOCS N = 1 / 1.000000 = 1.000000  ->  1.00

Western-perimeter exposure = 100.0%
ln(N) = 0    eta ln(N) = 0    h = -0.500000
```

In competition units the naive figure is 2,934.3 points and the bloc figure is 10,000.0 points. **A market with an HHI of 10,000 is a monopoly; the conventional structure is, in resilience terms, a monopoly held by one coordination perimeter over the entire estate.**

The adviser's file described this structure as diversified across four jurisdictions, two legal families and three regulators. Every word of that is true, and the measurement says the structure has one point of failure. Note what the measurement does *not* say. It does not say the four entities were badly chosen or badly run. The Jersey IP structure was well built; the Luxembourg vehicle was doing real treaty work. **The finding is not about quality; it is about correlation, and correlation is not visible from inside any one entity.**

The zero. The surplus law reads $h = \eta \ln N - \delta Q - C$ where N is the effective bloc count. For the conventional structure $N = 1.00$ and $\ln(1.00) = 0$ exactly. **The conventional structure has zero available surplus — not low, zero.** The entire input side is multiplied by the logarithm of the effective bloc count, and whatever η is, $\eta \times 0 = 0$. Protected capital then decays at δ with nothing on the input side to oppose it: at $\delta = 0.02$ that is 18.29% gone in ten years, 33.24% in twenty, 45.45% in thirty, on a half-life of 34.31 years.

And no rearrangement of those four entities changes it, because rearranging within a bloc does not change N.

conventional variant	naive HHI	naive eff	bloc-HHI	N	$\ln N$
as found, UK / US / JE / LU	0.293430	3.41	1.000000	1.00	0
perfectly equalised, four ways	0.250000	4.00	1.000000	1.00	0
split into six equal Western	0.166667	6.00	1.000000	1.00	0
split into ten equal Western	0.100000	10.00	1.000000	1.00	0
split into fifty equal Western	0.020000	50.00	1.000000	1.00	0

Fifty entities, a naive effective count of fifty, $N = 1.00$, $\ln N = 0$. **Subdivision inside a perimeter is the most expensive way to change nothing that this field has invented**, and it is the empirical content of H1's second falsification clause: no subdivision inside a perimeter raises N, and none did.

6. Results: the designed structure

jurisdiction	share	share^2	bloc	contents
UAE	29.1%	0.084681	BRIDGE	holdco, foundation, trading company
Thailand	18.2%	0.033124	NONALIGNED	promoted manufacturing operating company
UK	16.4%	0.026896	WEST	UK company and residual London exposure
Bermuda	13.6%	0.018496	W-lite $\rightarrow$ WEST	captive, reserves against real liabilities

jurisdiction	share	share^2	bloc	contents
India	11.8%	0.013924	NONALIGNED	international financial centre vehicle
Singapore	10.9%	0.011881	BRIDGE	treasury and dual-rail settlement
total	**100.0%**			

```
naive HHI = 0.189002
effective jurisdictions = 1 / 0.189002 = 5.290949  ->  5.29

bloc collapse:
    BRIDGE      = 29.1 + 10.9 = 40.0%      0.160000
    NONALIGNED  = 18.2 + 11.8 = 30.0%      0.090000
    WEST        = 16.4 + 13.6 = 30.0%      0.090000
bloc-HHI = 0.340000
EFFECTIVE BLOCS N = 1 / 0.340000 = 2.941176  ->  2.94

Western-perimeter exposure = 30.0%
ln(N) = 1.078810   eta ln(N) = 0.377583   h = -0.122417
```

6.1 The comparison, in the units that matter

metric	conventional	designed	change
entities	4	6	$+2$
jurisdictions	4	6	$+2$
naive HHI	0.293430	0.189002	-35.6%
effective jurisdictions	3.41	5.29	$\times 1.5525$
bloc-HHI	1.000000	0.340000	-66.0%
effective blocs N	**1.00**	**2.94**	$\times 2.9412$
Western-perimeter exposure	100.0%	30.0%	-70.0 pts
$\ln N$	0.000000	1.078810	
$\eta \ln N$ at $\eta = 0.35$	0.000000	0.377583	

The architecture's entire contribution to the surplus law is 0.3776, purchased by going from one effective bloc to 2.94, and it is bought with two additional entities rather than twenty.

6.2 What each move buys

move	N after	ΔN
start: conventional	1.0000	—
add foundation as the holding layer	1.0000	0.0000
add Gulf holdco, UAE share to 29.1%	≈ 1.7000	$+0.7000$

move	N after	ΔN
add Thai operating company, 18.2%	≈ 2.3000	$+0.6000$
add Indian centre vehicle, 11.8%	≈ 2.6000	$+0.3000$
add Singapore treasury, 10.9%	2.9412	$+0.3412$
UK reduced 40.9% $\to$ 16.4%, Bermuda 13.6%	2.9412	—
final	**2.9412**	

The intermediate values are approximate because the path depends on the order of execution; the endpoints are exact. Two entries deserve comment.

The foundation adds nothing to N and is still the most important element in the architecture. It is a succession instrument, and it is the holding layer that makes everything else administrable. **A structure element that does not move the concentration measurement is not thereby worthless; it is answering a different threat.**

The first non-Western jurisdiction is worth more than the last three combined. Moving the Gulf share to 29.1% buys roughly 0.70 of effective blocs; the Singapore addition buys 0.34. This is $dh/dN = \eta/N$ in practice, and it has a governance implication: if the programme runs over three years and might be interrupted, execute the highest-marginal-value move first, because a half-finished diversification in the right order is worth far more than a half-finished one in the wrong order. Curzi [2026] prices the marginal contribution of one corridor exactly: removing the Indian corridor gives bloc-HHI 0.391448, $N = 2.5546$ and $h = -0.136733$, so 11.8% of exposure buys 0.014316 of surplus, which is 3.79% of the total gain.

7. Results: hypotheses tested

7.1 H1 — the perimeter hypothesis: supported

The continuous family $\mathrm{Var}(R) = \text{naive} + \rho(\text{bloc} - \text{naive})$, run on both structures:

Designed structure:

ρ	$\mathrm{Var}(R)$	N
0.00	0.189002	5.2909
0.10	0.204102	4.8995
0.25	0.226751	4.4101
0.40	0.249401	4.0096
0.50	0.264501	3.7807
0.60	0.279601	3.5765
0.75	0.302251	3.3085
0.90	0.324900	3.0779
1.00	0.340000	2.9412

Conventional structure, carried through to the surplus balance:

ρ	$\mathrm{Var}(R)$	N	$\eta \ln N$	h
0.00	0.293430	3.4080	0.4291	-0.0709

ρ	$\mathrm{Var}(R)$	N	$\eta \ln N$	h
0.25	0.470073	2.1273	0.2642	-0.2358
0.50	0.646715	1.5463	0.1525	-0.3475
0.75	0.823357	1.2145	0.0680	-0.4320
0.90	0.929343	1.0760	0.0256	-0.4744
1.00	1.000000	1.0000	0.0000	-0.5000

The conventional structure's effective count falls from 3.4080 to exactly 1.0000 as ρ rises to one, and the naive count never enters. **H1's prediction holds as an identity in the model and is not falsified by any variant tested.** What the arithmetic cannot do is establish that the real within-perimeter correlation is near one; that is an empirical claim, it rests on the observed behaviour of coordinated listing, substance and exchange frameworks over fifteen years, and Section 12 records it as the paper's largest unmeasured input.

This paper uses $\rho = 1$ within a bloc. That is the conservative end. A reader who regards it as too strong can read an answer off the table, which is why the family is published rather than the endpoint.

7.2 H2 — the sufficiency hypothesis: not met

At the parameters fixed in Section 2:

```
maintenance floor:  delta Q + C = 0.02 x 10 + 0.30 = 0.5000

conventional:  h = 0.35 x ln(1.000000) - 0.5000
                 = 0.35 x 0.000000     - 0.5000  =  -0.500000
designed:      h = 0.35 x ln(2.941176) - 0.5000
                 = 0.35 x 1.078810     - 0.5000  =  -0.122417

h = 0 requires  eta ln N = delta Q + C
                ln N = 0.5000 / 0.35 = 1.428571
                   N = e^1.428571 = 4.172734
```

The criterion was $N \geq 4.172734$. The redesigned structure reaches $N = 2.941176$. H2 is not met. The designed structure improves h by 0.377583, a 75.5% reduction in the shortfall, and remains at -0.122417. Both structures are below the build threshold. **Jurisdictional diversity buys the ability to build; it does not buy the building.**

The honest conclusion is therefore not that more diversification is required. It is that the maintenance side must come down or the conversion side must come up. The other two levers, computed at $N = 2.941176$:

```
C required for h = 0    :  0.377583 - 0.200000 = 0.177583   (against actual 0.30)
eta required for h = 0  :  0.500000 / 1.078810 = 0.463474   (against actual 0.35)
Q at which h = 0, C=0.30:  (0.377583 - 0.300000) / 0.02 = 3.879169
```

That last line is the uncomfortable one: at $C = 0.30$ the designed structure carries protected capital only up to 3.88 units before erosion outruns the surplus. **Structure cost is close to fixed and erosion is proportional, so there is a minimum viable size below which no amount of jurisdictional cleverness helps, and a maximum above which the same structure needs rebuilding.**

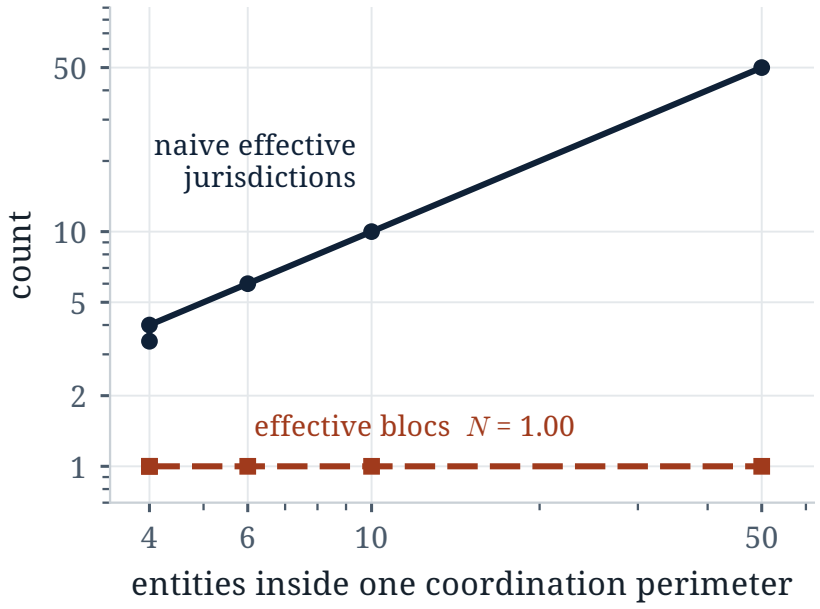


Figure 52.2: **Subdivision inside one perimeter moves the naive effective count from 3.41 to 50.00 and leaves the effective bloc count at exactly 1.00 in every row, so the surplus term stays at zero.** Measured from the share vectors in the table above; script `figs/src/CCP-052_fig02.py`.

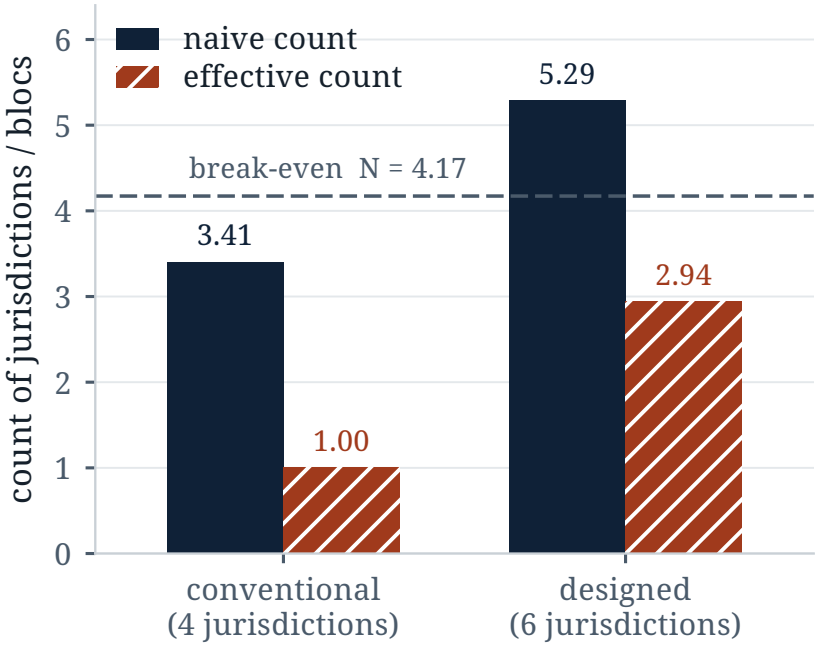


Figure 52.3: **The naive count and the effective count on one pair of axes: the conventional structure reads 3.41 naively and 1.00 in effect, the designed structure 5.29 and 2.94, and neither effective count clears the break-even of 4.17 that a non-negative surplus balance requires.** Measured from the share vectors in Sections 5 and 6; break-even derived at the stated parameters. Script `figs/src/CCP-052_fig01.py`.

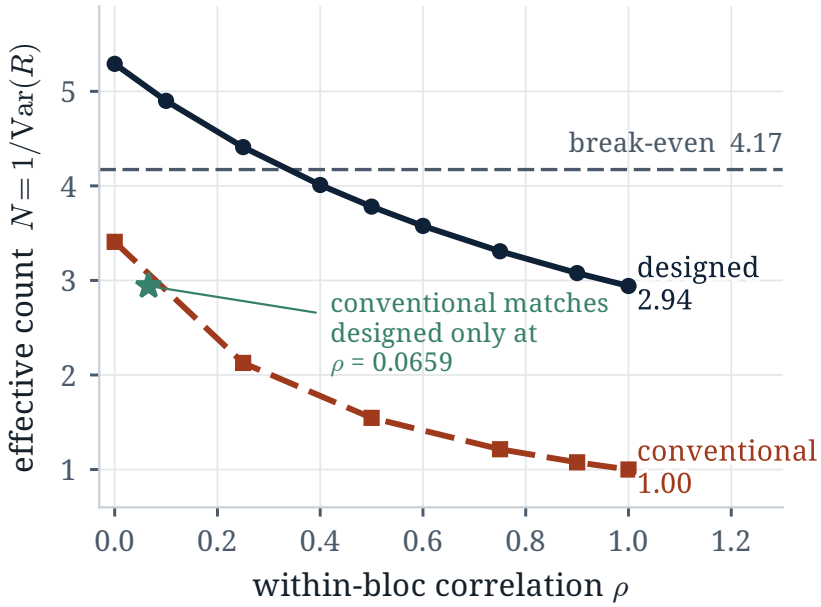


Figure 52.4: **The effective count against within-bloc correlation.** The conventional structure falls from 3.41 to exactly 1.00 as correlation rises to one, which is H1's prediction; the designed structure falls from 5.29 to 2.94 and crosses below the break-even of 4.17 at a correlation near 0.35. Modelled from the variance identity with the measured share vectors; script `figs/src/CCP-052_fig03.py`.

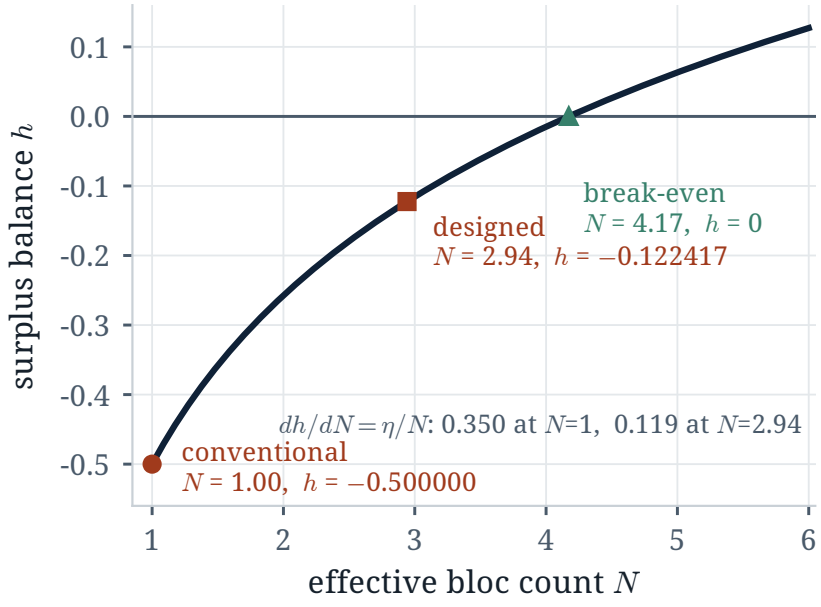


Figure 52.5: **The surplus balance against the effective bloc count at the stated parameters.** The conventional structure sits at the point of steepest available return and captures none of it; the designed structure captures 0.377583 of surplus and still lands 0.122417 short of the break-even at $N = 4.17$. Modelled from $h = \eta \ln N - \delta Q - C$ with $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$; script `figs/src/CCP-052_fig04.py`.

7.3 H4 — the implied-correlation hypothesis: supported

How weakly would the Western jurisdictions have to correlate for the conventional structure to be as resilient as the designed one? Solve naive $+ \rho(1 - \text{naive}) = 0.34$ with naive $= 0.293430$:

```
rho = (0.340000 - 0.293430) / (1 - 0.293430) = 0.065910
```

Sustaining the conventional structure's naive reading requires believing that the four Western jurisdictions correlate at 6.6% on coordinated action. The claim being made when someone reports naive HHI on a single-bloc structure is $\rho = 0$; it is not a defensible claim, it is an unexamined default in a spreadsheet, and the value 0.065910 is the price of the default stated as a number. H4's prediction is met exactly.

7.4 The marginal return

```
dh/dN = eta / N
at N = 1.0000     dh/dN = 0.350000
at N = 2.9412     dh/dN = 0.119000
at N = 4.1727     dh/dN = 0.083878
```

The first bloc added is worth roughly three times what the third is. This is the mathematical reason the conventional structure is not merely bad but *maximally* bad: it sits at the point of steepest available return and captures none of it. Moving from one effective bloc to two is the highest-return move available anywhere in the framework and also the cheapest, because it requires one correctly chosen jurisdiction rather than five. The mirror statement bounds the prescription: going from 2.94 to 4.00 buys $0.35(\ln 4 - \ln 2.941176) = 0.1074$, under a third as much, for a far harder piece of engineering, and each additional bloc carries its own C. **There is an optimum number of blocs and it is small.** The organisational literature reaches the same shape from the other direction: Uzzi and Spiro [2005] find an inverted U in which creative performance rises and then falls with cohesion, and Sagan [1993] documents that redundancy can *reduce* safety when it adds coordination surface. A monotone raise-N prescription cannot survive those results unqualified, and this paper does not make one.

7.5 H3 — the robustness hypothesis: supported

assignment variant	bloc-HHI	N
base: Bermuda folded into WEST	0.340000	2.9412
Bermuda as its own W-lite bloc	0.295392	3.3853
Singapore reclassified as WEST	0.341962	2.9243
UAE reclassified as WEST (worst case)	0.451162	2.2165
UAE and Singapore as separate blocs	0.276562	3.6158
Thailand and India as separate blocs	0.297048	3.3665
every jurisdiction independent (= naive)	0.189002	5.2909

Three readings.

The base case is conservative in two places at once. Folding Bermuda into WEST costs 0.44 effective blocs relative to treating it as its own perimeter, and treating the two bridge jurisdictions as one costs 0.67 relative to separating them. The published figure of 2.94 is the low end of a defensible range, not the high end.

The result is robust to the single most contestable call. Reclassifying Singapore as Western moves N from 2.9412 to 2.9243, a change of 0.0169 — under six-tenths of one per cent. The argument most likely to be had does not move the answer. **H3's first prediction is met.**

Only one variant does real damage. If the Gulf bridge were absorbed into the Western perimeter, N falls to 2.2165 and h to $0.35\ln(2.2165) - 0.5 = -0.221425$: a 24.6% fall in N from a single structural change, against a 0.6% move from a single-member reclassification. **H3's second prediction is met, and the asymmetry between the two is the finding.** That is threat 5 with a number attached, and it is why a degradation trajectory with evidence, rather than a verdict, belongs in every jurisdiction assessment.

7.6 The second direction: the same instrument, opposite sign

The instrument governs a second constraint running the other way, and this paper owns both halves so that no downstream work treats them as unrelated.

The 2023 merger guidelines issued jointly by the two US antitrust agencies set the structural presumption of unlawfulness at HHI above 1,800 with a change over 100, and identify a merged-firm share above 30% combined with a change over 100 as also warranting a presumption [U.S. Department of Justice and Federal Trade Commission, 2023]. Two points of discipline. The guidelines are agency enforcement policy rather than statute and courts weigh them without being bound. And the thresholds have moved: the 2010 horizontal guidelines used 1,500 and 2,500 with changes of 100 and 200 [U.S. Department of Justice and Federal Trade Commission, 2010], and the 2023 guidelines returned to 1,800 with 100 [U.S. Department of Justice and Federal Trade Commission, 2023]. **That the banding has moved while the arithmetic has not is the cleanest available demonstration that the measure is neutral and the thresholds are policy.**

The change from a merger has a closed form. If two firms with shares s_1 and s_2 combine, the post-merger index gains $(s_1 + s_2)^2 - s_1^2 - s_2^2$:

```
delta = 2 x s1 x s2        (shares in percentage points)
```

so the change exceeds 100 exactly when $s_1 s_2 > 50$; for an equal-share pair that is $s > \sqrt{50} = 7.0711\%$ each. A 5% by 5% combination produces 50 points and is comfortably clear; an 8% by 8% combination produces 128 and is not.

measurement	HHI (points)	direction wanted	verdict
jurisdictional exposure, conventional	2,934.3	low	fail
jurisdictional exposure, designed	1,890.0	low	better
bloc exposure, conventional	10,000.0	low	total fail
bloc exposure, designed	3,400.0	low	acceptable
product market, carried case	1,320.0	low enough	clear

One metric. Two directions. Both binding. The jurisdictional rows want the number small because a small number means many independent points of failure. The market row wants it small enough because a large number attracts the presumption. The arithmetic does not care which question is being asked.

On the constructed product market — rivals at 22.0, 18.0, 11.0, 8.0, 7.0, 6.0, 5.0 and 4.0 with the operating company at 14.0 and five fringe firms at 1.0 each — the market HHI is 1,320.00 points and the effective firm count is 7.5758. The acquisition arithmetic:

target	share	$\Delta = 2s_1 s_2$	post-HHI	combined	presumption
Rival H	4.0%	112.00	1,432.00	18.0%	not engaged
Rival G	5.0%	140.00	1,460.00	19.0%	not engaged
Rival D	8.0%	224.00	1,544.00	22.0%	not engaged
Rival B	18.0%	504.00	1,824.00	32.0%	**engaged**

The Rival B acquisition crosses both tests at once. **The same instrument that says the principal's jurisdictional structure is not yet diversified enough says the proposed acquisition is one the agencies would presume unlawful, and both statements come out of a sum of squared shares.** The largest target that keeps the change at or below 100 points is $100/(2\times 14) = 3.5714\%$ — which is to say, once a firm holds 14% of a market, almost any acquisition worth doing produces a change above 100. **The change test, not the level test, is what binds a mid-sized acquirer.**

The symmetry is not a pun. Both applications ask how much of a total sits in how few hands, because that question governs both fragility and market power. A firm that concentrates a market has power over it; a structure that concentrates its exposure has no power over what happens to it.

8. The cost side: three jobs, and H5

Section 7.2 left the shortfall on the cost side. This section tests the standard remedy.

There are three functions, and they are not three flavours of one function:

job	needs
HOLDING	succession law, asset segregation, perpetuity, **low substance burden**
OPERATING	**real substance**, labour, cheap presence, market access
BRIDGING	**dual-bloc banking**, dual-rail settlement, neutral law

Most failed structures come from asking one entity to do all three, and the three requirement sets are antagonistic in specific, nameable ways.

Holding versus operating: the substance contradiction. Holding needs a low substance burden — the entity should own things, be cheap to maintain, be simple to inherit, and create as little nexus as possible. Operating needs real substance: employees, premises, decision-makers on the ground, functions genuinely performed where they are said to be performed. **These are the same requirement with opposite signs, and no entity satisfies both.** Substance requirements exist for good reasons and the base-erosion work on substantial activities in low-tax jurisdictions [Organisation for Economic Co-operation and Development, 2015] is the reason they exist; the correct response to a substance regime is to put substance where the activity genuinely is, which is the whole thesis of this section.

Operating versus bridging: the alignment contradiction. An operating entity is embedded in a state: a factory that cannot move, a workforce with local rights, licences from a local ministry, a tax residence. It is therefore reachable by that state and visibly aligned with it. A bridge works because both sides believe it will not be turned against them. **A structure cannot bridge from a position of alignment, and an operating entity is aligned by construction.**

Holding versus bridging: the exposure contradiction. A bridging entity is in constant contact with counterparties, screening systems, correspondent banks and settlement risk; its exposure profile is an exposure-*window* profile parameterised by hop size, duration and mid-

transit failure handling. A holding entity's exposure profile is a custody profile. Merging them puts the store of value in the entity with the highest transactional exposure in the group. **A rail is not a holding**, and value in transit and value at rest are different problems.

8.1 The overload diagnostic

Apply the concentration instrument to each entity's own function vector $j = (H, O, B)$, where the shares are of the *entity's* activity and not of the group's:

```
S = H^2 + O^2 + B^2                specialisation index
effective jobs = 1 / S
```

$S = 1.00$ is a pure specialist returning 1.00 effective jobs; an entity split evenly across all three returns 3.00. The bands used are: below 1.50 SPECIALIST, 1.50 to 2.20 STRAINED, above 2.20 OVERLOADED.

> **A pure two-job entity attains its maximum effective-job count at a 50/50 split, where $S = 0.50$ and the count is exactly 2.00. There is no two-job configuration above 2.00. Any entity measuring above 2.00 therefore carries a non-zero share in all three jobs, by arithmetic necessity, whatever the org chart says.**

The 2.20 threshold leaves room for estimation error in the vector; the mathematically meaningful line is 2.00, and crossing it is proof rather than evidence. The 1.50 boundary corresponds to a two-job mix of 0.7887 / 0.2113 — a clear primary job and a real but minority second one.

The structure as found, scored:

entity	H	O	B	S	eff jobs
UK holdco (holds, trades, banks group)	0.40	0.35	0.25	0.3450	2.8986
US portfolio entity	0.55	0.10	0.35	0.4350	2.2989
Jersey IP company	0.60	0.35	0.05	0.4850	2.0619
Luxembourg treaty vehicle	0.45	0.05	0.50	0.4550	2.1978
group aggregate	0.50	0.21	0.29	0.3778	2.6468
mean effective jobs per entity					2.3643

Every one of the four entities is at or above 2.00 effective jobs, which means every one is carrying all three jobs by arithmetic necessity. Four entities, four jurisdictions, one bloc, and no specialisation anywhere. That is a diagnosis the concentration measurement alone could not produce: Section 5 said the structure has one point of failure; this says each of its four entities is simultaneously failing three requirement sets. The two findings are independent and they compound.

The designed structure, scored:

entity	H	O	B	S	eff jobs	verdict
foundation	0.90	0.00	0.10	0.8200	1.2195	specialist
Thai operating company	0.05	0.95	0.00	0.9050	1.1050	specialist
UK company	0.10	0.85	0.05	0.7350	1.3605	specialist
Singapore treasury	0.05	0.15	0.80	0.6650	1.5038	strained
Gulf holdco	0.70	0.05	0.25	0.5550	1.8018	strained
Bermuda captive	0.30	0.60	0.10	0.4600	2.1739	strained
Indian centre vehicle	0.20	0.30	0.50	0.3800	2.6316	**overloaded**
mean effective jobs per entity					1.6852	

That the framework's own architecture produces a flag is the reason the diagnostic is worth having. The captive at 2.1739 is strained by nature: it underwrites, which is an operating activity requiring licensed substance and supervision, while also holding reserves, which is a holding activity. The strain is accepted because the reserves are held against real third-party liabilities. The diagnostic's job is not to forbid the strain; it is to make sure the strain was chosen.

The centre vehicle at 2.6316 is above 2.00 and therefore carries all three jobs, **which is a design error this paper owns rather than smooths over**. Two repairs, both computed:

```
as specified                    H0.20 O0.30 B0.50   eff jobs 2.6316   OVERLOADED
repair A: specialise BRIDGING   H0.05 O0.10 B0.85   eff jobs 1.3605   SPECIALIST
repair B: specialise OPERATING H0.05 O0.80 B0.15   eff jobs 1.5038   STRAINED
```

Repair A is the better move on these numbers. After it, no entity in the designed structure carries all three jobs:

	as found	designed
mean effective jobs	2.3643	1.6852
worst entity	2.8986	2.6316 → 1.3605 after repair A
entities above 2.00	4 of 4	1 of 7 → 0 of 7 after repair
entities below 1.50	0 of 4	3 of 7

The group-level measurement asks the complementary question — whether a job rests on too few entities. Weighting each entity's job share by size and normalising within each job:

job	effective entities carrying it	largest carrier
HOLDING	2.6141	foundation, 52.02%
OPERATING	2.9838	Thai operating company, 46.34%
BRIDGING	4.5051	Gulf holdco, 31.75%

Holding is the most concentrated, at 2.61 effective entities with the foundation carrying just over half. That is intentional — concentrating holding is the point of having a holding layer — but it should be a decision rather than a surprise. **Job concentration and jurisdictional concentration are different measurements, and a structure can be well diversified on one and deliberately concentrated on the other.**

8.2 The assignment rule, and why the minimum

Score each candidate 0–5 on six requirement dimensions, then take the **minimum** across the dimensions the job actually needs:

```
dimensions
  1  succession law          4  real substance capacity
  2  asset segregation       5  market and labour access
  3  low substance burden    6  dual-bloc banking

HOLDING    fit = min over {1, 2, 3}
OPERATING  fit = min over {4, 5}
BRIDGING   fit = min over {2, 6}
```

Constructed scores. These are the analyst's judgements applied to published characteristics, and a reader who disagrees should change a score and rerun the rule rather than abandon it:

candidate	succession	segregation	low-subst	real-subst	market	dual-bank
Gulf centre A	5	5	4	2	2	4
Gulf centre B	5	5	4	2	2	5
Singapore	4	4	3	4	4	5
Bermuda	3	5	2	3	2	3
UK company	3	3	2	4	5	2
Indian centre	2	3	3	3	3	3
Thailand	1	2	1	5	5	3
single offshore company	3	4	5	1	1	1

Applying the rule gives HOLDING to Gulf centre A at fit 4, OPERATING to Thailand at fit 5, and BRIDGING to Gulf centre B at fit 5.

Why the minimum and not the average. A job is a conjunction. An entity scoring brilliantly on two of a job's three requirements and failing the third does not do two-thirds of the job; it does none of it. Run both statistics on the holding dimensions and watch the masking:

candidate	holding dims	average	**minimum**	gap
Gulf centre A	[5, 5, 4]	4.67	4	0.67
Gulf centre B	[5, 5, 4]	4.67	4	0.67
single offshore company	[3, 4, 5]	4.00	3	1.00
Singapore	[4, 4, 3]	3.67	3	0.67
Bermuda	[3, 5, 2]	3.33	2	1.33
UK company	[3, 3, 2]	2.67	2	0.67
Indian centre	[2, 3, 3]	2.67	2	0.67
Thailand	[1, 2, 1]	1.33	1	0.33

On the average, the single offshore company ranks third for holding, ahead of Singapore, because an excellent low-substance score of 5 pulls up a mediocre succession score of 3. On the minimum the reason is visible. Bermuda shows the pattern more sharply: an average of 3.33 conceals a 2 on low substance burden, and the gap of 1.33 is exactly the amount of masking. **The gap between the average and the minimum measures how much a candidate's weakest requirement is hidden by its strongest, and a large gap is a warning, not a recommendation.**

8.3 H5 — the specialisation hypothesis: falsified on direct cost

Run a constructed annual cost model in index units, decomposed into legal, substance and compliance:

entity type	legal	substance	compliance	total
holding entity, low substance	6	1	0	7
holding entity forced to add operating substance	6	14	9	29
operating entity, real substance	3	18	11	32
bridging entity, compliance-heavy	4	7	16	27

A single overloaded entity must meet the *maximum* of each column, since it faces every requirement: $6 + 18 + 16 = 40$. Three specialists cost $7 + 32 + 27 = 66$.

```
overloaded single entity      40
three specialised entities    66
saving from consolidation     26     = 39.39%
```

H5 predicted that specialisation lowers direct cost. It does not. The overloaded entity is cheaper by 39.39%, and that is the honest result rather than the one the argument wanted. It is also the whole explanation for why overloaded structures exist: the saving is real, immediate and visible on an invoice, and the fragility it buys is none of those things. Organisational slack — the property that makes a system survive a shock it did not plan for — has always shown up on the accounts as waste [Cyert and March, 1963], and this is that result in miniature.

The case for specialisation therefore has to be made on **risk-adjusted** cost, and it can be. Model each job's requirement set as a test the entity passes with probability p; an entity doing k jobs faces k tests, all of which must pass:

p	one job	two jobs	three jobs
0.95	0.9500	0.9025	0.8574
0.90	0.9000	0.8100	0.7290
0.85	0.8500	0.7225	0.6141
0.80	0.8000	0.6400	0.5120
0.70	0.7000	0.4900	0.3430

At $p = 0.80$ the three-job entity holds with probability 0.5120 against 0.8000 for the specialist — 36.0% worse. **The model is crude and its conclusion is robust to every reasonable choice of p, because conjunction is unforgiving and there is no parameterisation in which requiring three things is easier than requiring one.** It also understates the problem: the three tests are *negatively* correlated, because arranging the entity to pass one actively makes another harder, so real joint probabilities sit below p^3 and the direction of the error is known even where its magnitude is not.

Taking $p = 0.80$, the overloaded entity's annual structural-failure probability is $1 - 0.8^3 = 0.4880$ against $1 - 0.8 = 0.2000$ for the worst single specialist. Let LOSS be the cost of a structural failure:

LOSS	risk-adjusted overloaded	risk-adjusted specialised	winner
100	88.80	86.00	specialised
150	113.20	96.00	specialised
250	162.00	116.00	specialised

LOSS	risk-adjusted overloaded	risk-adjusted specialised	winner
400	235.20	146.00	specialised

```
break-even LOSS = (66 - 40) / (0.4880 - 0.2000) = 90.28 index units
```

Specialisation pays whenever a structural failure costs more than 90.28 index units — 1.37 times the specialised structure's entire annual running cost. That is a very low bar. And the multi-year picture is where the asymmetry becomes plain:

years	cumulative saving	P(fail) overloaded	P(fail) specialised
1	26	0.4880	0.2000
3	78	0.8658	0.4880
5	130	0.9648	0.6723
10	260	0.9988	0.8926

The saving accrues annually and linearly; the failure happens once and permanently. That is the shape of the trap, and it is the same shape as the recoverability argument of Section 4.4, arrived at from a completely different direction.

8.4 The over-fragmentation limit, and the fourth job

The diagnostic says split. It does not say split indefinitely, and an instrument that pushed only one way would be half an instrument. Each additional entity adds to C, and if it sits inside a bloc the structure already occupies it adds exactly nothing to N. Model the marginal all-in cost of an entity at 0.035 in the units of the surplus law, calibrated so that the eight-entity designed structure costs $C = 0.30$:

entities	N	C	h
4	1.0000	0.1600	-0.3600
5	1.7000	0.1950	-0.2093
6	2.3000	0.2300	-0.1385
7	2.6000	0.2650	-0.1306
8	**2.9412**	**0.3000**	-0.1224
10	2.9412	0.3700	-0.1924
14	2.9412	0.5100	-0.3324
20	2.9412	0.7200	-0.5424

h improves up to eight entities and then falls, because entities nine onward add cost without adding blocs:

```
C at which h = -0.5000 with N = 2.9412:   0.677583
entities = 8 + (0.677583 - 0.300000) / 0.035 = 18.79
```

At roughly nineteen entities the structure is back to the conventional structure's $h = -0.5000$ while holding $N = 2.9412$. **A structure can be perfectly diversified and still worthless, by being administered to death.**

And every entity carries filing, accounting, audit, governing-body and register obligations regardless of which job it does. Modelling this administrative load at 0.02 per entity per year:

entities	admin load	share of $C = 0.300$
1	0.020	6.7%
4	0.080	26.7%
6	0.120	40.0%
8	0.160	53.3%
12	0.240	80.0%
15	0.300	100.0%

At eight entities, administration alone consumes 53.3% of the maintenance budget, and at fifteen it consumes all of it. **Administration is not a job in the taxonomy and it is the largest single line in C, which means the three jobs are competing for slightly less than half the budget they appear to have.**

8.5 Reconciling the two measurements

	as found	designed
effective blocs N	1.00	2.9412
$\ln N$	0.000000	1.078810
$\eta \ln N$ at $\eta = 0.35$	0.000000	0.377583
mean effective jobs per entity	2.3643	1.6852
entities	4	7
C	lower	0.30
h	-0.500000	-0.122417

The as-found structure is cheaper. It has four entities rather than seven, no substance in an operating jurisdiction, and no dual-bloc banking to maintain. Its C is lower and its h is still worse by 0.377583, because $\ln(1.00) = 0$ cannot be compensated for by any saving on the maintenance side. The designed structure costs more and is worth more, and specialisation pays whenever a structural failure costs more than 1.37 years of running cost.

9. Failure modes

Thirteen ways the instrument is misused, each with the arithmetic that detects it.

1. Counting entities. An adviser reports “eleven entities in seven jurisdictions” as though it were a measurement. It is a count of invoices. Entities inside one perimeter contribute nothing to N and each contributes to C. **An entity that does not raise N and does raise C moves h strictly down.**

2. Reporting naive HHI on a single-bloc structure. The specific error this paper exists to correct. Reporting 3.41 effective jurisdictions for the conventional structure is not wrong arithmetic; the arithmetic is right and the premise is false. **Wherever the naive and bloc figures diverge sharply, the divergence is the finding**, and reporting only the flattering half is the professional failure.

3. Rearranging inside a bloc. Table in Section 5. Fifty entities, naive count fifty, $N = 1.00$.

4. The wrong share basis. Measuring by book value when the threat is availability produces a number that is precise and irrelevant. Illiquid real estate has a large book value and a specific availability profile — it cannot be moved. A listed portfolio has a smaller book value and a much sharper availability profile, because a custodian instruction can stop working in an afternoon.

5. Double-counting a bridge. A bridge jurisdiction is a bridge because it touches two perimeters. Counting it as a full independent bloc overstates N; assigning it to one of the perimeters it bridges understates it. **Report the bridge assumption explicitly**; it is the assumption most likely to be wrong and the one that moves the answer most.

6. Measuring one axis and calling it the structure. Diversifying across blocs frequently *concentrates* counterparties, because few institutions will bank a genuinely cross-bloc structure. **Run the instrument on all four axes or the axis measured is the one the analyst was already comfortable with.**

7. Optimising the index instead of the thing it measures. Two structures with identical bloc-HHI can differ enormously in whether the blocs are actually independent in the scenario that matters. **The index is a diagnostic, not an objective function.** This is Goodhart's problem exactly [Goodhart, 1975], and Campbell's law in the evaluation literature [Campbell, 1979]: a measure adopted as a target ceases to be a good measure. Scott [1998] is the historical form — legibility destroys the local variety that made the system work — and it is a warning about this instrument as much as an argument for it. The tell is a structure whose bloc shares are suspiciously round and equal; real structures are lumpy because assets are lumpy.

8. Forgetting the second direction. A group driving jurisdictional HHI down while consolidating an operating market can find that the same instrument certifying its resilience certifies its market power.

9. Reporting from a rounded print. N for the designed structure is exactly $1/0.34 = 2.941176\ldots$:

```
exact:    ln(2.941176) = 1.078810   h = 0.35 x 1.078810 − 0.5 = −0.122417
rounded:  ln(2.940000) = 1.078410   h = 0.35 x 1.078410 − 0.5 = −0.122557
```

The two differ by 0.000140. **The conclusion is unchanged — both are below zero — but the number moved in the fourth decimal from a single rounding, and a chain of four such steps is how a measurement goes wrong.** The rule is not to round carefully; it is to carry full precision through every intermediate step and round only at the print.

10. Treating the bloc assignment as a fact rather than a claim. The variance derivation makes the collapse rigorous *given* the assignment. It does not make the assignment right. Every bloc call is a forecast about correlated behaviour, and a forecast that is acted on can make itself true. **This is not a rhetorical worry; it has been documented in this paper's own domain.** MacKenzie and Millo [2003] traces an option-pricing model that did not fit the market when it was published and came to fit it after traders adopted it, so that the description made its own assumptions true; Hacking [1995] gives the general form, in which a classification of people interacts with the people classified and the referent moves. **If enough structures are built on the assumption that two jurisdictions share a perimeter, capital and correspondent relationships reorganise around that assumption, and the correlation the instrument assumed becomes the correlation the world exhibits.** Lucas [1976] is the severe form of the same point turned on the instrument itself: a parameter estimated in a population that has not read the instrument cannot be assumed stable in a population that has. **The instrument is therefore an engine as well as a camera, and the honest response is the record below rather than a claim of neutrality.** The minimum record per jurisdiction is: the bloc, the evidence class, the single strongest piece of evidence, the date, and **the specific observable that would change the call**. That last field is the one left blank and the only one that produces a decision.

11. Confusing effective count with capacity. $N = 2.94$ says the structure has roughly three independent points of failure. **It says nothing about whether any of them can carry the load.** A bloc holding 30% of the exposure through an entity that cannot open a correspondent account, cannot hire, and cannot obtain a legal opinion is a bloc on paper, and the measurement reports it at full weight because a distribution does not inspect capability. The correction is not to fudge the shares; it is to run the job assignment alongside the concentration measurement. A structure whose bloc diversity is real but whose bridging capacity sits entirely in one of three blocs has an effective bloc count of 2.94 on holdings and something much closer to 1.00 on the ability to transact.

12. Measuring after the event. A diagnosis taken during the emergency is a post-mortem. The arithmetic takes minutes once the exposure table exists; building the exposure table is a week of work that nobody schedules. **A measurement nobody is obliged to act on will be run once, filed, and cited for years after it stopped being true.**

13. Believing N can be raised indefinitely. Section 7.4. There is an optimum number of blocs and it is small; that most structures sit below it does not mean more is always better.

Three further failure modes belong to the threat register rather than the index. **Expected-value reasoning applied to a ruin risk** systematically underinvests in the only threat that rescales all future years; the remedy is to run the register twice, once on expected loss for budgeting and once on conditional permanent loss for structural decisions, and to treat divergence between the rankings as the signal it is. **Assuming reversibility** models a freeze as a delay; a frozen position is one that has stopped compounding and may be released years later minus the yield. **Modelling the actor as rational** imports an adversary who maximises recovery; a designation is a foreign-policy instrument whose logic is signalling, a correspondent exit is a compliance-cost decision made by a middle manager who will never learn what the corridor was worth. Arguments that would move a rational creditor have limited force against an actor whose objective function is elsewhere. **The defence has to work without an advocate.**

10. Discussion

What follows. Three things follow from the results and each is narrower than the general claim it resembles.

The index is exact, given the correlation matrix. EFF $= 1/(w'Rw)$ is a variance reciprocal, and its two named cases are the two extreme correlation assumptions. Nothing in the derivation is approximate. This is why the instrument survives the statistical-averaging objection of Doak et al. [1998]: the paper does not claim that diversity produces stability through an unspecified mechanism, it claims that a variance has a reciprocal and that the reciprocal counts independent shocks.

The status of the claim is measurement, not physics. The corpus is titled a psychological and sociological physics, and the word is checkable. **What is claimed here is a measured count with a dimension, a conservation law, stated limits and an invariance group** — Sections 11.1 to 11.4 — and **not** membership of the statistical-physics programme surveyed by Castellano et al. [2009], which supplies the vocabulary of order parameters, phase transitions and coarsening that this paper deliberately does not use. **No order parameter is identified, no phase transition is located, and no critical exponent is estimated.** Hedström and Ylikoski [2010] states the criterion the discipline named in the title will apply: an explanation must name a mechanism rather than exhibit a fitted curve. **The mechanism named here is a correlation structure, and nothing is fitted.** Gallegati et al. [2006] is the checklist this paper was run against — ignorance of existing work in the target discipline, sub-standard statistical procedure, unjustified universality claims, and models built on assumptions known to be false — and Section 12 reports where it passes and where it does not.

The instrument is domain-general and the generality is evidence. The same arithmetic is the Hill number of order two in ecology [Hill, 1973], the Simpson index [Simpson, 1949], the effective number of parties in political science [Laakso and Taagepera, 1979], the Rényi-2 participation ratio in information theory [Rényi, 1961], and the concentration index in competition law [U.S. Department of Justice and Federal Trade Commission, 2023]. Where it misbehaves is documented too: Golosov [2010] and Dunleavy and Boucek [2003] identify the dominated-unit regime in which the effective-number reading is least informative, which is precisely the regime the bloc examples live in — one unit at or near 100%. That the instrument is least discriminating exactly where this paper's headline case sits is a limitation and is recorded in Section 12.

Redundancy is not automatically resilience. Percolation gives the sharpest available for-

malisation of how many independent components a system really has [Callaway et al., 2000], and the robust-yet-fragile result [Albert et al., 2000] is the canonical demonstration that a network can be highly tolerant of random loss and extremely vulnerable to targeted loss — with the caveat, from Doyle et al. [2005], that real engineered systems get their robustness from design rather than from a degree distribution. The jurisdictional analogue is that a structure's tolerance depends on which component is removed, and a designation regime is a targeted removal.

What does not follow. Four things.

It does not follow that raising N is always right. Section 7.4 gives the diminishing return; Uzzi and Spiro [2005] gives the inverted U; Sagan [1993] gives the case in which redundancy reduces safety by adding coordination surface. The prescription here is bounded: move from one effective bloc to two or three, and stop when the marginal entity stops adding a bloc.

It does not follow that count substitutes for capability. Failure mode 11. The effective count measures independence, not capacity.

It does not follow that jurisdictional diversity answers every threat. It is close to irrelevant to counterparty failure and can worsen it. This is the corpus's own version of the general finding that diversity's benefit is conditional rather than automatic: Hong and Page [2004] is the canonical model in which diverse groups outperform high-ability groups, and Thompson [2014] is the critique establishing that the result follows from the model's assumptions rather than from data. **Both must be cited together and are here.**

It does not follow that the corpus's affine effective count inherits any of this. Section 3.2 and the Verification chain in Section 11.

Where this paper disagrees with its own corpus. Curzi [2026] reports naive HHI 0.293388 and 3.4085 effective jurisdictions on what is described as the same entity map; this paper computes 0.293430 and 3.407968. The difference is the fifth decimal and arises from whether the shares are taken as the rounded percentages 40.9 / 27.3 / 18.2 / 13.6 or as the exact fractions 9/22, 6/22, 4/22, 3/22 that produce them. **This paper uses the rounded percentages as stated, which is the reproducible choice given published inputs.** Neither figure changes any conclusion, both bloc figures are exactly 1.000000, and recording the discrepancy is preferable to silently harmonising it.

Bounded by the corpus's own honesty apparatus. Curzi [2026] assigns the corpus's central identification and the application of the surplus functional to a life to the postulate drawer, and states that the reading of a day's surface as a leading indicator is a wager whose misfire column exists because wagers can lose. That bound applies here: this paper's arithmetic is exact and its inputs are judgements. Curzi [2026] makes the exposure concrete with an instrument audit in which 19 of 27 parameter-by-scope cells have a defined estimator, 1 of 27 has a validated one, and **0 of 27 carry a first-party measurement**. Curzi [2026u] shows that the corpus's own falsification suite is anti-conservative as ordinarily run — against a nominal 0.05 across 500 null trials, i.i.d. permutation rejected at 0.108 and circular shift at 0.120 — which is a direct warning about how much confidence a passing check should buy.

Parallel results elsewhere in the corpus. The reciprocal-of-a-quadratic-form structure recurs. Curzi [2026s] derives a reciprocity factor $\rho = 2\min(x_{ij}, x_{ji})/(x_{ij} + x_{ji}) = 1 - \|\text{Skew}\|/\text{Sym}$ with the node imbalance index obeying $1-B \geq \rho$, flattering a worked network by 15 per cent — the same shape of correction, on a different object. Curzi [2026v] shows that attainable stroke counts on a star polygon are exactly the proper divisors of N and that $\gcd(N,k) = \gcd(N, N-k)$, verified without counterexample for $N = 3$ to 100: a discrete count that is invariant under a relabelling, as N is invariant under rearrangement inside a bloc. Curzi [2026] gives the deficit identity $D = 2.885390(c/\eta - F) = -2.885390(\delta/\eta)Q_{\infty}$, in which two apparently different quantities are shown to be the same number — structurally the move this paper makes in showing that naive HHI and bloc-HHI are one variance. Curzi [2026k] validates a five-parameter survey across scales, reproducing observed declines to residuals under 1.4% of the ceiling term, which is the closest the corpus comes to an out-of-sample test of the same model form. Curzi [2026] supplies the count-based essence index $f(n) = kn^2$ and the hold line $u^* = (e^{c/\eta}-1)/(N-1)$, which is the same effective-count variable entering a different law.

11. Verification and rigour audit

The result is decomposed into postulates that can be checked one at a time. Each is atomic, each carries a status, and the compounding chain names what each step consumes. **If dropping a postulate leaves the result standing, it is not listed.** Sections 11.1 to 11.5 run the standing rigour checks first, so that the chain in 11.7 is read against a dimensional audit rather than before one.

11.1 Dimensional audit

symbol	quantity	dimension	unit	status
s_i, w_i	exposure share of jurisdiction i	$[1]$	dimensionless, fraction of Q	defined; sums to 1 by P1
W_b	bloc share	$[1]$	dimensionless	derived, $W_b = \sum_{i \in b} s_i$
HHI	concentration index	$[1]$	dimensionless (or points $= 10^4 \times$ fraction)	derived from shares
ρ_{ij}	correlation between shocks	$[1]$	dimensionless, $[-1, 1]$	assumed, not measured
X_i	unit-variance shock to i	$[1]$	dimensionless by construction	defined
$\mathrm{Var}(R)$	variance of the aggregate shock	$[1]$	dimensionless	derived
N, N_{eff}	effective count	$[\#]$, **a count**	jurisdictions or blocs	derived as $\mathrm{Var}(R)^{-1}$
$\ln N$	surplus term argument	$[1]$	nats	derived; the logarithm of a count is dimensionless
Q	protected capital	$[C]$	index units of stored capital	defined; **no operational unit is supplied by the source**
δ	erosion rate	$[T^{-1}]$	per period	assumed
C	maintenance cost flow	$[C\,T^{-1}]$	index units per period	assumed
η	conversion efficiency	$[C\,T^{-1}]$ per nat	index units per period per nat	assumed
h	surplus balance	$[C\,T^{-1}]$	index units per period	derived
p, L, r	register probability, loss fraction, recovery fraction	$[1]$	dimensionless	assumed (constructed)
W, B	warning window, build time	$[T]$	days	assumed (constructed)
S (job index)	specialisation index	$[1]$	dimensionless	derived from a job vector

Does the surplus law balance? $h = \eta \ln N - \delta Q - C$ requires every term to carry $[C\,T^{-1}]$. δQ does, since δ is per period and Q is a stock. C does by definition. $\eta \ln N$ does **only if η carries $[C\,T^{-1}]$ per nat**, and the source is silent on this: it treats η as a bare "conversion efficiency". **The audit records that silence rather than filling it in.** The equation balances under the reading given in the table, and the reading is this paper's, not the source's.

The important row is N. HHI is a dimensionless sum of squared dimensionless shares, so its reciprocal carries the dimension of a **count** — and that is exactly why it is a numbers-equivalent rather than an index. **The instrument is not a score that happens to run from 1 upward; it is a measured count of independent units, with a unit, and it is checkable against the definition of a count by the limits in Section 11.3.**

Where no operational unit exists, this audit says so. Q is "protected capital in index units", and the source supplies no conversion to currency. Every figure in C, h and Q is therefore in an arbitrary scale fixed by the stated parameters. **The break-even $N = 4.172734$ is unit-free and survives this; the statement that h is -0.122417 does not mean anything until the index unit is defined.**

11.2 Significant figures and epistemic status

The two figures in the headline comparison have different statuses and must not be read at the same precision.

- $N = 2.941176$ and $N = 1.000000$ are **exact arithmetic** on the stated shares: $1/0.34$ and $1/1$. They are derived-from-defined and may carry as many digits as are useful. Their inputs — the shares — are constructed and exact as stated to one decimal place, and the exposure vectors sum to 100.0% exactly.
- $N = 4.172734$ is **exact given the parameters** — $e^{0.5/0.35}$ — and the parameters are **assumed**, not estimated. So the break-even is an exact function of an assumed input. **It carries no uncertainty because there is no estimate in it, and it carries no authority beyond the assumption.**
- $\rho = 0.065910$ is exact arithmetic on two exact indices.
- The register outputs — 9.25%, 5.97%, 0.956357 — are **illustrative model output from constructed parameters** and are rounded in the text to four decimal places at most. Six digits on 0.956357 is defensible only because the inputs are exact by stipulation; it is not a measurement and Section 12 says so.
- **No quantity in this paper is a fitted or measured estimate, so no uncertainty interval is reported anywhere.** That is not a strength: it means the paper has no empirical content to attach an interval to.

11.3 Limits and boundary behaviour

The instrument's limits are the proof that it behaves as a count.

limit	HHI	HHI^{-1}	reading
one holder, $s_1 = 1$	1	1 exactly	a monopoly counts as one unit
n equal holders, $s_i = 1/n$	$1/n$	n exactly	returns the count
n holders, any unequal split	$> 1/n$	$< n$	never overstates the count
$s_{\max} \to 1$ with the rest $\to 0$	$\to 1$	$\to 1$	continuous approach to the monopoly limit
$n \to \infty$, equal shares	$\to 0$	$\to \infty$	unbounded above, as a count should be

limit	HHI	HHI^{-1}	reading
$\rho \to 0$ within blocs	naive	naive count	the independent limit
$\rho \to 1$ within blocs	bloc	bloc count	the fully-coupled limit

Range: $1 \le \text{HHI}^{-1} \le n$, with both bounds attained. The lower bound is attained exactly when one unit holds everything and the upper exactly when all units are equal. A quantity that equals n for n equal units, equals 1 for a single unit, and lies strictly between otherwise **is a count**, and no further argument is needed.

The same limits run on the affine count, which fails them. $N_{\text{eff}} = 1+(N-1)u$ gives 1 at $u = 0$ and N at $u = 1$ for *any* share vector whatever — including the vector $(0.999, 0.001)$, for which it reports 2.000000 at $u = 1$ while every Hill number reports approximately 1.002. **It has the right endpoints and the wrong interior, because it never consults the shares.** The two quantities part company in the limits, and that is a cleaner demonstration than any argument in prose. The full arithmetic is in Section 11.9.

Boundary behaviour of the surplus law. As $N \to 1^+$, $h \to -(\delta Q + C)$, which is the maintenance floor: a structure inside one perimeter pays its whole bill out of stock. As $N \to \infty$, $h \to \infty$ logarithmically — which is a **defect of the model and a statement of its domain**, since Section 7.4 and Uzzi and Spiro [2005] establish that the real curve turns over. **The model's domain is small** N, and the paper's prescriptions are confined to the move from $N \approx 1$ to $N \approx 3$. As $\eta \to 0$, $h \to -(\delta Q + C)$ regardless of N: no conversion, no benefit from any amount of independence. As $\delta \to 0$ the stock never erodes and any positive $\eta \ln N$ exceeding C accumulates without bound, which is a pathology rather than an ideal.

11.4 Invariance

The effective count is invariant under relabelling of units and under any rescaling of the underlying exposures that preserves the share vector. Doubling every exposure in currency terms leaves every s_i unchanged and therefore leaves HHI, N and $\ln N$ unchanged. This is a genuine scale invariance and it is the reason the instrument can be applied to a structure of any size.

It is not invariant under regrouping, and that is the whole content of the bloc collapse: the naive and bloc figures differ by construction, and the difference is the finding. A reader who wants a single number that is invariant under regrouping is asking for a quantity that cannot detect a perimeter.

What would break the scale invariance: any share basis that is not normalised — for instance, reporting exposures in currency without dividing by the total. **Tested at how many scales?** The invariance is algebraic, not empirical, and was checked at exactly the two structures reported here. No claim of empirical scale-freedom is made.

What is conserved, and in what units. Mirowski [1989] shows that the characteristic failure of this genre is borrowing the notation of a conserved quantity without the conservation law that gave it content: nineteenth-century energetics was copied into economics and utility was made to play the role of energy without energy's conservation. **The objection is met here and the answer is short.** The conserved quantity is the exposure share vector: $\sum_i s_i = 1$ by P1, exactly, in dimensionless units, and it is that constraint and nothing else that gives the index its bounds — $s_{\max}^2 \le \text{HHI} \le s_{\max}$ and $1 \le \text{HHI}^{-1} \le n$ both follow from it. **Remove the normalisation and the effective count ceases to be a count.** The invariance group is relabelling of units together with any rescaling that preserves the simplex. Q**, by contrast, is a stock with no stated conservation law and no operational unit, and this paper does not claim one for it** — which is the honest half of the answer. Cartwright [1999] supplies the discipline that makes a partial answer respectable: specify the arrangement within which the law is claimed to hold, and stop claiming a universality the argument does not need. **The arrangement claimed here is a**

normalised exposure vector with a correlation matrix having unit diagonal and non-negative entries, and nothing outside it.

Aggregation is a separate and unmet objection. Kirman [1992] establishes that a single representative parameter for a heterogeneous group generally does not exist. The instrument writes one Q, one δ and one C for a whole structure. **No aggregation argument is offered and none is available from the arithmetic**; Section 12 records it.

11.5 Analogy discipline

Every borrowing in this paper is labelled with one of the three categories.

borrowing	source	category	justification
portfolio variance $\mathrm{Var}(R) = w'Rw$	Markowitz [1952]	**1 — derivation**	the exposure vector genuinely is a weighted sum of shocks with a correlation matrix; the premises are the algebra's own
Hill numbers, order 2	Hill [1973], Jost [2006]	**1 — derivation**	HHI^{-1} *is* ${}^2D$; this is an identity, not a resemblance
similarity-sensitive diversity	Leinster and Cobbold [2012]	**1 — derivation**, with a stated domain	$1/(w'Rw)$ is the order-2 member for similarity matrix R; exact where R has unit diagonal and non-negative entries, and claimed nowhere else
$\ln N$ as entropy	Shannon [1948]	**2 — shared form**	this is **Shannon entropy in nats**, a dimensionless information quantity, and not thermodynamic entropy; no thermodynamic content is imported
coupled-network fragility	Buldyrev et al. [2010]	**2 — shared form**	jurisdictions have not been shown to satisfy the percolation premises; only the structural property that coupling destroys assumed independence is used

borrowing	source	category	justification
robust-yet-fragile	Albert et al. [2000], Callaway et al. [2000]	**2 — shared form**	the degree-distribution result is not claimed to apply to jurisdictions; the shape of the tolerance asymmetry is
normal accidents	Perrow [1984]	**3 — framing**	qualitative organisational theory, cited as prior statement of the same observation, never as a formal result
synchrony in ecology	Loreau and de Mazancourt [2008]	**2 — shared form**	the naming convergence is noted; no ecological mechanism is imported
the effective number of parties	Laakso and Taagepera [1979]	**1 — derivation**	the same formula on a different share basis

The failure mode this table exists to prevent is letting a category-2 borrowing carry a category-1 result's authority. The concentration arithmetic is category 1 throughout and does not need help; the network results are category 2 and are used for shape only.

11.6 Measurement-theoretic status of the scored inputs

Two inputs in Section 8 are scores rather than counts, and their scale type is not optional.

The 0–5 requirement scores are ordinal. They are the analyst's rank judgements over published characteristics. **An ordinal scale licenses a minimum and does not license a mean.** The assignment rule takes a minimum, which is licensed. The masking table in Section 8.2 computes averages, which is **not licensed**, and it is included precisely to demonstrate that averaging ordinal fit scores produces a wrong ranking — the illicit operation is the object of study, not a step in the argument. **This is stated here so that no downstream work reads that column as a computed quantity.**

The job vectors (H, O, B) are treated as ratio-scale shares of an entity's activity, and that treatment is not tested. Michell [1997] is the strongest objection available to this move: psychology and its neighbours routinely assume quantitative structure without testing for it. The formal remedy is conjoint measurement [Luce and Tukey, 1964], whose conditions establish whether an attribute has additive quantitative structure at all, within the representational framework of the representational theory of measurement. **The conjoint conditions were not tested for the job vector, and the effective-jobs figures are therefore provisional.** What survives without the test is the *ordering* — an entity with a clearly dominant function scores lower than one without — and the exact threshold at 2.00, which is an algebraic property of any three-component vector on the simplex and does not depend on the scale type. The concentration instrument itself is untouched by this objection, because exposure shares are ratios of capital and not psychological judgements.

11.7 Postulates

#	statement	status
P1	exposure shares s_i are non-negative and sum to 1	DEFINITION
P2	HHI $= \sum_i s_i^2$	DEFINITION
P3	for n equal shares of $1/n$, HHI $= n \cdot (1/n)^2 = 1/n$	ARITHMETIC
P4	the aggregate shock is $R = \sum_i w_i X_i$ with Var$(X_i) = 1$	DEFINITION
P5	Var$(R) = \sum_i \sum_j w_i w_j \rho_{ij} = w' R w$	ARITHMETIC
P6	jurisdictions inside one enforcement perimeter act on one decision	POSTULATED (H1)
P7	the surplus law's input term is $\eta \ln N$	DEFINITION (modelling choice, CCP-A02)
P8	$\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$	POSTULATED (fixed in advance)

11.8 The chain that establishes the instrument

step	consumes	statement	status
L1	P2, P3	HHI^{-1} returns n exactly for n equal shares	ARITHMETIC
L2	L1	HHI^{-1} is a numbers-equivalent: the count of equal units giving the observed index — Adelman [1969]	CITED
L3	L2	HHI^{-1} is the Hill number of order 2 — Hill [1973], Jost [2006]	CITED
L4	L3	the same quantity is the Simpson index's reciprocal and the effective number of parties — Simpson [1949], Laakso and Taagepera [1979]	CITED
L5	P4, P5	with $R = I$, Var$(R) = \sum w_i^2 =$ HHI	ARITHMETIC
L6	P4, P5	with R block-of-ones, Var$(R) = \sum_b W_b^2 =$ bloc-HHI	ARITHMETIC
L7	L5, L6	naive and bloc indices are one variance under two correlation assumptions	ARITHMETIC
L8	L7	Var$(R) =$ naive $+ \rho($bloc $-$ naive$)$	ARITHMETIC

step	consumes	statement	status
L9	L7	EFF $= 1/(w'Rw)$ is the similarity-sensitive diversity of order 2 — Leinster and Cobbold [2012], Rao [1982]	CITED
R1	L7, P6	for the conventional structure, $N = 1.000000$ against a naive 3.407968	ARITHMETIC
R2	L7, P6	for the designed structure, $N = 2.941176$ against a naive 5.290949	ARITHMETIC
R3	R2, P7, P8	$h = 0.35 \ln(2.941176) - 0.5 = -0.122417$; break-even $N = e^{1.428571} = 4.172734$	ARITHMETIC
R4	L8	the correlation matching the two structures is $\rho = 0.065910$	ARITHMETIC

Every ARITHMETIC row can be redone by hand or in three lines of any language. Every CITED row names one source and no more. Only P6 and P8 are POSTULATED, and both are stated as hypotheses in Section 2 with their falsifiers attached. **The instrument therefore rests on two postulates, six definitions and a chain of arithmetic, and nothing else.**

11.9 The same chain applied to the affine effective count, which fails it

step	statement	status
Q1	$N_{\text{eff}} = 1 + (N-1)u$, with u a scalar non-parallelism	DEFINITION (corpus's own)
Q2	for two units at shares 0.9 and 0.1 with $u = 1$, Q1 gives exactly 2.000000	ARITHMETIC
Q3	for the same two units, $\text{HHI}^{-1} = 1/(0.81 + 0.01) = 1.219512$	ARITHMETIC
Q4	for the same two units, $\exp(H) = \exp(0.325083) = 1.384115$	ARITHMETIC
Q5	Q2 disagrees with Q3 and Q4, so Q1 does not return a numbers-equivalent	ARITHMETIC — **L2 fails**
Q6	Q1 is a function of (N, u) and not of the share vector, so it is not $(\sum p_i^q)^{1/(1-q)}$ for any q	ARITHMETIC — **L3 fails**
Q7	no entry in ecology, economics, political science or information theory states Q1	POSTULATED — no external support

The two chains part company at Q5. HHI^{-1} and $\exp(H)$ are Hill numbers of order 2 and 1 and are fully grounded. $1 + (N-1)u$ is a linear interpolation between 1 and N; it is the

corpus's own construction, it carries no external support, and **a reviewer who reaches for the Hill literature expecting it to apply will find that it does not**. That is stated here rather than left to be discovered. The corpus already notes elsewhere that the two are different quantities, which is to its credit; what this paper adds is that they must not travel under one name.

11.10 Cross-checks performed

- The double sum $\sum_i \sum_j w_i w_j \rho_{ij}$ was evaluated explicitly on the designed structure's six shares under both correlation assumptions and reproduced 0.189002 and 0.340000 to twelve decimal places.
- $s_{\max}^2 \le \text{HHI} \le s_{\max}$ holds on both structures: $0.167281 \le 0.293430 \le 0.409$ and $0.084681 \le 0.189002 \le 0.291$.
- The effective-count ceiling $s_{\max}^{-2}$ is respected: $3.407968 \le 5.9780$ and $5.290949 \le 11.8090$.
- The rounding audit of failure mode 9 was run and the discrepancy quantified at 0.000140.

12. Limitations

The correlation is assumed, not measured. $\rho = 1$ within a bloc is the conservative end of a family, and the justification is the observed behaviour of coordinated listing, substance and exchange frameworks over fifteen years. **No correlation was estimated from data in this paper.** That is the single largest unmeasured input, it is exactly the quantity H1 turns on, and the honest response is that the family in Section 7.1 is published so that a reader may substitute a different value. A reader who believes $\rho = 0.5$ within the Western perimeter should read $N = 1.5463$ off the table and proceed.

The effective count is least informative exactly where the headline case sits. Golosov [2010] and Dunleavy and Boucek [2003] document that effective-number measures behave poorly in the dominated-unit regime, where one unit holds a very large share. The conventional structure's bloc vector is a single unit at 100%, which is the limiting case of that regime. The measurement is still exact — $1/1 = 1$ — but the *discriminating power* of the instrument is at its minimum there, and comparisons among heavily dominated structures should not be made on N alone.

Small-sample bias is not addressed. Chao et al. [2014] show that diversity estimators are biased on small samples. A six-element exposure vector is a small sample. No correction was applied and none is available for the case where the exposure vector is a complete census rather than a sample — which is arguably this case, since the structure's exposures are enumerated exhaustively rather than sampled. **The paper does not resolve which of the two situations it is in, and a reviewer would be right to press on it.**

Every parameter in the surplus law was chosen, not estimated. $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ are stated test parameters. They were fixed before the measurement, which protects H2's criterion from being adjusted to fit, and they are not estimates of any real position. **The break-even $N = 4.172734$ is therefore a property of the chosen parameters and not of the world.** At $\eta = 0.50$ the break-even falls to $e^{1.0} = 2.718$ and the designed structure clears it. That the conclusion of H2 depends on η in this way is the most important sensitivity in the paper.

Every case is constructed. The principal, the assets, the shares, the threat register parameters, the cost model, the market shares and the job vectors are all invented and are labelled as such at each point of use. **No empirical dataset was collected and none is reported.** The measurements are exact computations on constructed inputs.

The bloc assignments are the analyst's judgements. They are contestable, the robustness table shows the sensitivity, and each carries a stated observable that would change the call. They are not derived from a published register of coordination behaviour, because no such register exists.

Jurisdiction-specific legal claims are not cited to primary sources. The descriptions of specific centres, statutes, foundation regimes, treaty positions, ownership rules and court ar-

rangements are taken from the source manuscript. **The bibliography contains no jurisdiction-specific primary sources — no statutes, regulations, court rules or registries — and no case law.** Those claims are therefore uncited in this paper, are described rather than asserted, and require verification against the official registers before any reliance. Appendix B lists them individually. This is a real gap and it is not closed here.

The effective count measures independence and not capacity. Stated as the paper's own counter-result: $N = 2.94$ says nothing about whether any bloc can carry the load. A structure whose bloc diversity is real but whose bridging capacity sits in one bloc has an effective count near 1.00 on the ability to transact.

The affine effective count used elsewhere in the corpus has no external support. Section 11.9. This paper does not use it, but papers citing this one may, and the distinction must survive the citation.

Monotone claims are restricted. The paper's claim is not that surplus increases without bound in independence. Section 7.4 bounds it, and Uzzi and Spiro [2005], Sagan [1993] and May [1972] are the reasons a monotone claim could not survive unqualified.

The instrument is exposed to its own failure mode 7. Adopted as a target, N would cease to be a good measure [Goodhart, 1975, Campbell, 1979]. Nothing in the arithmetic prevents that, and the only defence offered is the reporting standard in Appendix A, which requires the inputs to be published alongside the output.

12.1 Rigour checks not met, listed by name

The standing checks of Section 11 were run in full. The following did not pass and are printed here rather than left for a reader to find.

1. **No as-of date, because no real holdings were measured.** A concentration figure computed from real holdings requires its source and its date; a concentration measure without an as-of date is not a measurement. **Every share vector here is constructed, so there is no source and no date to give.** Any application of this instrument to a real structure must carry both, and the reporting standard in Appendix A accordingly makes the date item 9.
2. **No uncertainty is reported anywhere, because nothing is estimated.** Every figure is exact arithmetic on stipulated inputs. That is not precision; it is the absence of empirical content, and it means the paper cannot be wrong in the way a measurement can be wrong.
3. **The dimension of η is supplied by this paper, not by the source.** The source treats conversion efficiency as a bare number. The surplus law balances only if η carries capital per period per nat, and Section 11.1 assigns that reading explicitly rather than silently.
4. **Q, C and h have no operational unit.** They are index units fixed by the stated parameters. The break-even count is unit-free and survives; the magnitude of h does not mean anything until an index unit is defined, and the source defines none.
5. **The surplus law's large-N limit is a defect, not a result.** h grows without bound in $\ln N$, which the empirical literature contradicts. The model's domain is therefore small N, stated in Section 11.3 rather than discovered by a reader.
6. **The conjoint conditions were not tested for the job vector**, so the effective-jobs figures are provisional as ratio-scale quantities [Michell, 1997, Luce and Tukey, 1964]. The 2.00 threshold survives because it is algebraic; the magnitudes do not.
7. **An unlicensed operation appears in Section 8.2 on purpose.** The average of ordinal 0–5 fit scores is not a licensed statistic, and the column exists to show that averaging ordinal fit scores selects the wrong candidate. It must not be read as a computed quantity.
8. **Model comparison was not performed.** No alternative index was fitted, no out-of-sample test was run, no information criterion was computed and no matched null was constructed, because nothing here is fitted. The instrument is asserted to be the right one on the strength of the identity chain in Section 11.8, not on comparative fit.

9. **The coupled-network borrowing is category 2, not category 1.** The premises of Buldyrev et al. [2010] have not been shown to hold for jurisdictions sharing an enforcement perimeter, and the result is used for its structural form only.
10. **No distributional or scaling claim is made, so the Clauset–Shalizi–Newman procedure was not run** [Clauset et al., 2009]. The exposure vectors are finite sets of four and six shares, not samples from a distribution. **Any restatement of this paper that reaches for a heavy-tail or power-law claim about holdings must run that procedure first**, and must read Stumpf and Porter [2012] on why passing it still would not identify a mechanism.
11. **The bloc assignment is a stipulated reference group, not an identified effect.** Manski [1993] rules out recovering an endogenous social effect from observation of a linear-in-means reflexive system, and the framework's reflexive loop is such a system. The remedy adopted is stipulation with a defence [Brock and Durlauf, 2001], and **the price is that the assignment is a modelling position**. No randomisation, instrument or partial-population variation was available.
12. **The instrument is performative and its own parameters are exposed to that.** MacKenzie and Millo [2003] documents a model that made its assumptions true after adoption; Lucas [1976] states the consequence for parameters estimated before publication. **The framework asserts that systems change when modelled and then exempts its own η, δ and correlation assumptions from that reflexivity. That is an inconsistency and it is not resolved here**; what is offered instead is the auditable record required by failure mode 10.
13. **There is no conservation law for Q.** Section 11.4 gives the one genuine conservation — shares sum to one — and states plainly that the stock variable has none [Mirowski, 1989]. **This is the objection the genre most often fails, and this paper meets half of it.**
14. **Aggregation is unaddressed.** One Q, one δ and one C are written for a heterogeneous structure, and Kirman [1992] establishes that such a representative parameter generally does not exist.
15. **Expected value and time average are not the same thing, and the register reports the first.** Peters [2019] is the reason the conditional permanent-loss column exists at all; **the register's headline 9.25% is an ensemble average that no single holder experiences.**
16. **No differential prediction is offered.** Durlauf [2005] asks for the prediction that only this model makes, and the instrument makes none: any correlation-aware concentration measure would produce the same collapse. **The contribution is the identification and the discipline, not a uniquely discriminating forecast.**

13. Conclusion

This paper establishes the concentration instrument in the form the rest of the corpus cites, and the following statements are the fixed text of that result.

R1. With HHI defined as the sum of squared exposure shares and the effective count as its reciprocal, a conventional four-jurisdiction structure at UK 40.9%, US 27.3%, Jersey 18.2% and Luxembourg 13.6% measures a naive HHI of 0.293430 and 3.407968 effective jurisdictions; collapsing the same shares onto coordination blocs before squaring gives a bloc-HHI of 1.000000 and an effective count of exactly 1.00.

R2. A redesigned six-jurisdiction structure reaches a naive HHI of 0.189002 (5.290949 effective jurisdictions) and a bloc-HHI of 0.340000 (2.941176 effective blocs), with Western-perimeter exposure falling from 100.0% to 30.0%.

R3. Coupled to the surplus law $h = \eta \ln N - \delta Q - C$ at $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$: $N = 1.00$ gives $h = -0.500000$; $N = 2.941176$ gives $\eta \ln N = 0.377583$ and $h = -0.122417$; and break-even is $N = e^{1.428571} = 4.172734$.

R4. The correlation that would make the conventional structure as resilient as the designed one is $\rho = (0.340000 - 0.293430)/(1 - 0.293430) = 0.065910$.

R5. The index is insensitive to bloc-assignment error and sensitive to bloc structure: reas-

signing one member moves N from 2.9412 to 2.9243, a change of 0.0169, under 0.6%; merging two blocs moves it to 2.2165.

R6. The five-threat register prices expected annual loss at 9.25% of protected capital and permanent expected annual loss at 5.97%, with the probability of at least one threat over ten years at 0.956357 and threat-1 window coverage of $2/210 = 0.0095$. Returns diminish as $dh/dN = \eta/N$: moving from 1.00 to 2.94 blocs buys 0.3776 of surplus while moving from 2.94 to 4.00 buys 0.1074.

Counter-result. The redesign reaches only $N = 2.94$ against a break-even of 4.17, so h remains negative at -0.122417. **H2 was tested and not met.** And the effective count measures independence, not capacity: it says nothing about whether any bloc can carry the load. **The instrument's honest verdict on the architecture it was built to evaluate is: substantially better, still insufficient, and the insufficiency is on the cost side rather than the design side.**

Two further results are established as method rather than as measurement. The naive and bloc indices are one variance under two correlation assumptions, so EFF $= 1/(w'Rw)$ is the general form and is exactly the similarity-sensitive diversity of order two. And HHI^{-1} is the Hill number of order two while $1 + (N-1)u$ is not a Hill number of any order and carries no external support; downstream work must name which effective count it means.

Curzi [2026] reproduces the full computation of R1 through R3 as a workbook, and that reproduction is the intended check on everything above.

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every case in this paper — the principal, the asset mix, the exposure shares, the threat-register parameters, the cost model, the constructed product market, and the entity job vectors — is **CONSTRUCTED** and is labelled as such at each point of use. The numerical results are exact computations on those constructed inputs, carried at full precision and rounded only for printing.

All figures were produced by the scripts named in their captions (`figs/src/CCP-052_fig01.py` through `figs/src/CCP-052_fig05.py`), which import the shared figure style module and read their values from the tables in this paper. No figure plots a number that does not appear in the text.

Jurisdictional and legal descriptions in this paper are taken from the source manuscript and not from primary legal sources. The bibliography contains general comparative and treaty material only; it contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction. Appendix B lists the claims affected. Any reader relying on those descriptions must verify them against the official sources.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Legal, tax and investment advice

This paper is analytical and is not legal, tax, or investment advice. It describes entity types, jurisdictions and structural arrangements for the purpose of measuring concentration and resilience. No recommendation is made in respect of any particular person, asset, structure or jurisdiction. Statutory reach-back periods, reserved-share regimes, substance requirements, corporate-residence tests, and merger-review thresholds all change and vary by jurisdiction and by activity; every such statement in this paper is carried from the source manuscript, is flagged in Appendix B, and requires confirmation with qualified counsel in the relevant jurisdiction before any reliance.

Appendix A. The reporting standard

Any structure measured under this instrument reports the following, and nothing less.

```
share basis, stated in one sentence
exposure table, jurisdictions and shares, summing to 100.0%
naive HHI, in fraction and in points, and effective jurisdictions
bloc assignment with the evidence class for each jurisdiction
bloc-HHI and effective blocs N
Western-perimeter exposure as a percentage
ln(N), eta ln(N), and h at the stated parameters
the robustness table over contested bloc assignments
the date, and the diff against the previous measurement
```

Item 9 decays first and matters most, because bloc fragmentation is invisible except as a difference between two measurements taken a year apart.

The entity-level register that accompanies it carries, per entity: legal form and jurisdiction; assigned job; job vector as measured rather than as intended; S and effective jobs with the band; where the strain is and its written justification; the requirement dimension on which the jurisdiction scores lowest; who takes the decisions belonging to the job and where; annual cost decomposed into legal, substance, compliance and administration; and the date with the diff against the previous register. **The lowest-scoring dimension is the column that pays for the exercise**, because the assignment rule took a minimum and that requirement is therefore already computed. An entity whose lowest dimension is dual-bloc banking is one correspondent decision away from being unable to do its job; an entity whose lowest dimension is succession law is fine until the founder dies. **Evidence created contemporaneously is worth an order of magnitude more than the same facts reconstructed three years later, and the cost of creating it is a column in a table.**

Appendix B. Claims requiring verification against primary sources

The following statements rest on legal or regulatory detail that changes, and **none of them is cited to the bibliography, which contains no primary sources for any individual jurisdiction**. They are reproduced from the source manuscript and require confirmation with counsel in the relevant jurisdiction before any reliance.

1. All statutory reach-back periods for transfers vulnerable to challenge: the two-year and ten-year federal bankruptcy provisions, the four-year uniform-act period with its one-year discovery rule, and the two-year, five-year and no-fixed-period provisions of the English insolvency legislation. Confirm current text, the applicable enactment — the

uniform act is adopted state by state and enactments vary — and whether the transfer in question is characterised as being to a self-settled trust.

2. The limitation analysis for the English provision on transactions defrauding creditors. The section fixes no period and the authorities address how the general limitation statute applies. This is contested ground.
3. The reserved-share provisions of the French Civil Code — one half, two thirds and three quarters for one, two, and three or more children — and, critically, the conflict-of-laws question of whether they reach a given asset at all. Characterisation drives the answer and is asset-specific.
4. Whether a particular administrative measure carries a right of appeal, which is the entire threat-1 / threat-2 boundary in practice.
5. The grounds on which recognition and enforcement of an award are in fact refused under the 1958 Convention, and the comparative treatment of the trust and the foundation as alternative devices for the succession problem. Both are settled subjects with their own literatures; neither is cited here, because the bibliography carries the instruments and the general doctrine only.
6. The named case authorities on the unavailability of a pre-judgment injunction over assets in an action for money damages, on the failure of the impossibility defence, and on the freezing-injunction line in English practice. **Case citation follows different conventions and no case is in the bibliography.**
7. Foundation and holding-vehicle characteristics in the named centres: separate legal personality, purpose or beneficiary structures, founder control retention, and the supervising courts. Confirm current text and any amendments.
8. Thai law as described: the treaty permitting full foreign ownership for persons of one nationality, investment-promotion conditions conferring ownership and land rights, the foreign business restrictions, the criminal treatment of nominee shareholding, and the revenue order on remitted foreign-source income. Promotion conditions are activity-specific and change.
9. Economic substance regimes and their thresholds, which vary by jurisdiction and activity category and have been amended repeatedly.
10. Corporate residence and central-management-and-control tests, which differ between jurisdictions and determine whether the split-entities-but-not-decision-rights failure is fatal in a given case.
11. Merger-review thresholds and their treatment in the relevant circuit. The guidelines cited [U.S. Department of Justice and Federal Trade Commission, 2023, 2010] are agency enforcement policy rather than statute and courts are not bound by them.
12. All cost figures are index units in a constructed model. They are not fee quotations and no real fee schedule was used to produce them.

All numerical results in this paper were computed from the stated inputs at full precision and rounded only for printing.

Chapter 53

The Apex Vehicle: Foundation Against Trust, and the Arithmetic of Forced Heirship

Abstract. A holding apex must be recorded as an owner in registries that may not recognise a trust, which is the case for preferring an entity with separate legal personality over a relationship. We compare the DIFC Foundation with the common-law discretionary trust on recognition, cost and succession, and report where the comparison does not favour the foundation. On cost, the fixed-fee apex is modelled at USD 120,000 a year against ad valorem trusteeship, giving a crossover Q* = 120,000 divided by the fee basis: the foundation is cheaper only above USD 34.3 million at 35 basis points, USD 80.0 million at 15, and USD 120.0 million at 10. At the constructed USD 240 million estate the foundation costs 5.0 basis points against 35, a difference of USD 720,000 a year — but the paper states plainly that cost is the wrong reason to choose the form. On succession, two exposures are computed separately: a French compensatory levy capped by French-situs assets at min(60.00m, 6.70m) = USD 6.70 million, or 11.2 per cent of the notional reserved share, against an inheritance-tax exposure of approximately 0.45 × 60.00m = USD 27.00 million that the structure does not touch. The binding constraint is liquidity: a USD 60.00 million claim against USD 26.16 million of liquid assets leaves a shortfall of USD 33.84 million, 14.1 per cent of the estate, which forces the sale of an operating branch. Much of the paper is comparative legal exposition of the foundation, trust and Sharia regimes.

Keywords: private foundation; discretionary trust; legal personality; recognition and characterisation; forced heirship; réserve héréditaire; succession liquidity

Series: D. Structure and Resilience **Paper:** CCP-053

1. Introduction

A holding structure has an apex, and the apex is the one element that is genuinely difficult to reverse. An operating company can be redomiciled, wound up or duplicated; a bank account can be replaced in ninety days. **The apex is chosen once, and every subsequent decision inherits its law, its court, its disclosure regime and its politics.**

This paper asks what form that apex should take, and answers with one identification, three measurements and one refusal.

The identification. The concentration instrument of Curzi [2026] — the Herfindahl index of exposure shares, collapsed onto coordination blocs before squaring, with the effective count as its reciprocal — measures where the assets are. **It does not measure whether the *concept* holding those assets is portable.** A family whose assets are properly spread across six jurisdictions in three blocs, but which holds them all through a common-law discretionary trust, has an invisible version of the concentration problem: every one of those holdings depends on local courts and local registrars accepting a form of ownership that several of them do not recognise as ownership at all. We call this a **recognition dependency**, and it is a form of concentration the published instrument cannot see.

The first measurement is cost, and it does not favour the conclusion. A fixed-fee foundation apex is modelled at USD 120,000 a year against an *ad valorem* trusteeship. Setting the two equal gives a minimum efficient scale of $Q^{*} = 120{,}000/b$ for fee basis b: the foundation is cheaper only above USD 34.3 million at 35 basis points, USD 80.0 million at 15, and USD 120.0 million at 10. At the constructed USD 240 million estate it costs 5.0 basis points against 35, a difference of USD 720,000 a year. **The paper then states plainly that cost is the wrong reason to choose the form**, because the crossover reverses at a scale many families sit below and because a fixed-fee trustee mandate is available in the market.

The second measurement separates two exposures the same conversation usually merges. The forced-heirship exposure and the inheritance-tax exposure are different problems with different causes, different remedies and, in the worked case, magnitudes that differ by a factor of four. A compensatory levy capped by local-situs assets gives $\min(60.00\text{m}, 6.70\text{m}) =$ USD 6.70 million, 11.2 per cent of the notional reserved share; the inheritance-tax exposure on the same receipt is approximately $0.45 \times 60.00\text{m} =$ USD 27.00 million, and **the structure does not touch it.**

The third measurement finds the binding constraint, and it is not entitlement. A USD 60.00 million claim against USD 26.16 million of liquid assets leaves a shortfall of USD 33.84 million, 14.1 per cent of the estate and larger than any single liquid holding. **Closing it in cash requires selling an operating branch — which is the fragmentation the apex existed to prevent, arriving through the succession rather than through the register.**

The refusal. The apex does not move the concentration reading. Over the conventional structure of Curzi [2026] the measurement remains a naive index of 0.2934 (3.41 effective jurisdictions) against a bloc index of 1.0000 and an effective count of exactly 1.00, whatever vehicle sits at the top. The foundation answers ownership recording and succession continuity. **It adds nothing to the effective bloc count, gives no protection against arbitrary state action, and its forced-heirship result rests on a consent the heir may refuse — for which this paper states there is no lawful substitute.**

Two constraints on what follows are stated at the outset because they shape every section.

Jurisdiction-specific law is not cited. The bibliography contains general treaty and comparative doctrinal material — the recognition convention for trusts [Hague Conference on Private International Law, 1985], the judgments and choice-of-court instruments [Hague Conference on Private International Law, 2005, 2019], the succession convention that never entered into force [Hague Conference on Private International Law, 1989] and the regional regulation that partially fills the gap [European Union, 2012], together with the comparative literature on the trust and the foundation [Lupoi, 2000, Smith, 2013, Panico, 2014, Hayton et al., 2016] and the conflict-of-laws texts [Torremans et al., 2022, Briggs, 2013, Fawcett and Carruthers, 2008]. **It contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction.** Every enactment, decided case, article number, rate and formality described below is described from the source manuscript, left uncited, and itemised in Appendix B. Where a table would otherwise score a named jurisdiction on a characteristic the bibliography cannot support, the jurisdiction is de-identified — "Gulf centre A" — following the convention of Curzi [2026]. **The abstract names the two jurisdictions the source works from; Sections 4 to 8 refer to them functionally, and Section 3.6 states why.**

No distributional or scaling claim is made anywhere. The exposure vectors here have four and six entries; they are not samples from a distribution. Clauset et al. [2009] is the procedure

that would have to be run before any heavy-tail claim could be made about holdings, and it is not run because no such claim is made.

2. Hypotheses

Five hypotheses, each with what would falsify it, in the form Platt [1964] requires. Meehl [1967] is the standing warning that a directional prediction here is nearly costless; H2, H3, H4 and H5 make point predictions with constants attached. One of the five is reported in Section 9 as **not met in the sense the argument wanted**.

H1 — the recognition hypothesis. An apex whose mechanism is the division of legal and beneficial ownership carries a recognition dependency that an apex with separate legal personality does not, and that dependency is not captured by the concentration instrument. *Predicted:* on a constructed family map, the count of non-common-law touchpoints — persons or assets sitting in states whose law does not know the ownership split — is strictly positive and material, while the effective bloc count is identical under both apex forms. *Falsified by:* a legal system in which the ownership question put to a foreign court closes identically for a relationship and for a registered legal person; or by a concentration measure that registers the difference.

H2 — the minimum-efficient-scale hypothesis. A fixed-fee apex has a scale below which it is the more expensive form, and that scale is computable rather than arguable. *Predicted:* $Q^{*} = C_{\text{fixed}}/b$, so at $C_{\text{fixed}} = 120{,}000$ the crossovers are USD 34.3 million at $b = 0.0035$, USD 80.0 million at $b = 0.0015$, and USD 120.0 million at $b = 0.0010$. *Falsified by:* a fee structure in which the fixed-cost form is cheaper at every scale — which a fixed-fee trustee mandate supplies, and which is why H2 is stated as a scale result rather than as a recommendation.

H3 — the two-exposure hypothesis. The forced-heirship exposure and the inheritance-tax exposure are independent quantities and the second is frequently the larger. *Predicted:* on the constructed case, a situs-capped levy of $\min(60.00\text{m}, 6.70\text{m}) = 6.70$ against a tax exposure of approximately $0.45 \times 60.00 = 27.00$, a ratio of 4.03. *Falsified by:* a case in which the structure's treatment of the reserve also changes the heir's tax position, which would mean the two are not independent.

H4 — the liquidity hypothesis. Where the reserve remedy is monetary, the binding constraint on a holding structure is liquidity and not entitlement. *Predicted:* a claim of USD 60.00 million against liquid assets of USD 26.16 million leaves a shortfall of USD 33.84 million, 14.1 per cent of the estate, exceeding every single liquid holding and forcing the disposal of an operating branch. *Falsified by:* a structure in which the monetary claim is satisfiable from liquid assets, or a system in which the remedy is in kind — in which case the exposure is an ownership problem and this hypothesis does not apply.

H5 — the apex-neutrality hypothesis. The choice of apex vehicle does not change the effective bloc count. *Predicted:* the conventional structure measures a naive index of 0.2934 and 3.41 effective jurisdictions against a bloc index of 1.0000 and $N = 1.00$, and those figures are invariant under substitution of the apex form. *Falsified by:* any apex vehicle whose substitution changes N while the branch jurisdictions and their shares are held fixed.

Not hypotheses, and labelled as such. Four statements in this paper are not falsifiable as written.

- *That a foundation is the correct apex for the carried case* is a **modelling conclusion** drawn from H1 together with a constructed family map. Change the map and the conclusion changes; it is not a general claim about instruments.
- *The cost figures* — USD 120,000, USD 65,000, the three fee bases, the USD 240 million estate — are **constructed illustrations**, not quotations, and H2's crossovers are exact functions of them.
- *The identification of the reserve fractions, levy conditions, tax rate and Sharia shares* is **description from the source manuscript**, uncited, and is treated throughout as an input

to arithmetic rather than as a finding.
- *The parameter set of the surplus law*, $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$, is a **modelling choice** carried from Curzi [2026] and fixed in advance, following the preregistration discipline of Nosek et al. [2018].

3. Method

3.1 The apex in the storage law

The framework's surplus law is $h = \eta \ln N - \delta Q - C$, with Q the protected capital held under durable protection, δ the erosion rate, η the conversion of flow into stock, C the recurring maintenance cost and N the effective count of independent coordination perimeters. Curzi [2026b] states the gauge in full and supplies the result this paper needs most: **conversion composes multiplicatively through recording, binding and reach gates**, so three gates at 0.7 give $\eta = 0.343$, and the weakest gate dominates — moving the recording gate from 0.20 to 0.80 gains 300 per cent while moving the binding gate from 0.60 to 0.90 gains 50 per cent. **An apex is a recording gate.** Its job is to make the structure recordable as an owner in every register that must record it, and on the multiplicative composition that is the gate with the highest available return precisely because it is the one most often left low.

An apex enters the law twice and with opposite signs. It permits a distribution across perimeters to be held as one owner, which is a precondition for $N > 1$; and it costs registrar fees, a registered agent, council fees, accounting and counsel in several jurisdictions, which raises C. **The apex does not itself raise N, and Section 8.5 measures that directly rather than assuming it.**

Curzi [2026] supplies the accounting discipline this paper applies to the cost side: the carrying cost of a structure is paid in four currencies — coin, hours, attention and foreclosure — against a law under which felt comfort decays over weeks while carrying cost does not decay at all. **The USD figures below are the coin column only.** The hours and attention consumed by an apex are real, are not modelled here, and the omission is recorded in Section 12. The foreclosure column — what holding this form forecloses — is the migration and conversion analysis of Section 5.5.

Curzi [2026] supplies the framing the whole paper presupposes: the present arrangement is a local minimum of the cost landscape by definition rather than by luck, which is why an apex that has stopped being appropriate continues to feel settled from inside it, and why the decision procedure of Section 6.8 is run on a map rather than on an impression.

3.2 Which effective count, and what it does not measure

Every effective count in this paper is HHI^{-1}, the Hill number of order two, and it is named as such at each use. Hill [1973] defines the family ${}^{q}D = (\sum_i p_i^q)^{1/(1-q)}$, whose members return the number of equally-abundant types that would produce the observed index; at $q = 2$ that is exactly $(\sum_i p_i^2)^{-1}$. Jost [2006] is the statement that entropies are not diversities and their numbers-equivalents are; Adelman [1969] is the numbers-equivalent reading of the concentration index itself; Simpson [1949] and Laakso and Taagepera [1979] are the independent reinventions in ecology and political science. The index is due to Hirschman [1945] in root form and to Herfindahl [1950] in sum-of-squares form.

The corpus's affine count $N_{\text{eff}} = 1 + (N-1)u$ is a different quantity, is not a Hill number of any order, and carries no external support. It is not used in this paper. Curzi [2026] prices the axiom that produces it exactly: the alternative $F_{\text{alt}}(u) = u \ln N$ satisfies the other three axioms and agrees at both endpoints, yet differs by 33.62 per cent at $u = 0.5$ — 1.871802 against 1.242453. **A functional form that is pinned at its endpoints and free by a third in the middle is a design choice, not a derivation**, and that pricing is the reason this paper stays with the Hill number.

The general form the bloc collapse relies on is EFF $= 1/(w'Rw)$, the similarity-sensitive diversity of order two [Leinster and Cobbold, 2012] whose additive relative is the quadratic entropy of Rao [1982]. Curzi [2026] owns that identity; this paper consumes it. Hurlbert [1971] is the standing discipline: "diversity" carries no content without its formula attached, which is why the formula is attached here every time.

What the instrument does not measure is the subject of H1. EFF $= 1/(w'Rw)$ is a statistic of an exposure vector and a correlation matrix. Neither object contains any information about whether the *form of ownership* asserted over those exposures is intelligible to the courts and registries that must record it. A recognition dependency is therefore invisible to the measurement, and Section 8.5 shows that it is invisible in the strong sense: the numbers are identical under both apex forms.

3.3 The crossover model

The cost comparison is a one-line model and its transparency is the point. Let C_{fixed} be the annual cost of a fixed-fee apex and b the *ad valorem* fee basis of the alternative, expressed as a decimal fraction of assets under trusteeship. The alternative costs bQ. Setting the two equal:

```
C_fixed = b Q*        =>        Q* = C_fixed / b
```

Q^* is a **minimum efficient scale**: below it the fixed-cost form is the expensive one. The model has exactly two parameters, both constructed, and no free parameters at all. Its limits are run in Section 12.3.

The same shape governs compliance architecture generally, and Curzi [2026] reaches it from a different direction: a compliant multi-jurisdiction architecture costs 224.40 basis points a year at USD 25 million and 22.44 at USD 250 million, so that **compliance-based architecture has substantial fixed costs and therefore a minimum efficient scale, below which the correct advice is a simpler structure rather than a cheaper version of a complex one.** That two independent cost models produce the same result is worth stating; it is not additional evidence, because both are constructed from the same premise that the numerator is fixed.

3.4 The two-exposure decomposition

Two numbers are computed separately because they are different problems that the same conversation usually merges.

```
E_reserve   the maximum value a forced-heirship or family-provision
            claimant could extract, under the law that would apply,
            CAPPED BY WHAT THAT LAW CAN ACTUALLY REACH -- which for a
            levy on local assets is the value of the local assets and
            not the notional entitlement.

E_tax       the tax an heir bears on what she receives, computed under
            HER OWN residence rules, which are independent of the
            succession law and of the structure.
```

The decomposition is not a refinement; it is the difference between measuring the exposure a structure can address and the exposure it cannot. **A structure can eliminate the first entirely and leave the second untouched, and the second is often the larger number.** H3 states that as a prediction with a ratio attached.

3.5 Where the reflexive loop enters, and why it is not identified

The framework's central formal object is the reflexive loop $T(x) = W(x, M(x))$ of Curzi [2026], and it enters this paper at one specific place: **an apex is a governance instrument, and a governance instrument changes the behaviour of the people it governs.** A family charter that records an equalisation alters what each heir expects, which alters what each will accept, which alters what the charter needs to record.

Manski [1993] proves that a linear-in-means reflexive specification of that kind is **not identified from observational data**: the endogenous effect cannot be separated from contextual effects or from correlated unobservables, however much data is collected. **No observation of family behaviour under a governance instrument could therefore confirm that the instrument produced the behaviour.** The remedy adopted here is the fourth of the four known ones — the reference group is stipulated and the stipulation defended — which is the route Brock and Durlauf [2001] shows can be made rigorous, and **the price is that every claim in this paper about what a governance document achieves is a modelling position rather than an identified empirical result.**

The performative form of the same problem is live and is documented in this paper's own domain. MacKenzie and Millo [2003] traces a model that did not fit its market when published and came to fit it after adoption; Hacking [1995] gives the general looping form, in which a classification interacts with the people classified. **A decision procedure that tells families to prefer foundations where non-common-law touchpoints exceed a threshold will, if adopted, change which forms banks and registrars are equipped to process, which changes the touchpoint cost, which changes the threshold.** Lucas [1976] states the severe consequence for the parameters. **The framework asserts that systems change when modelled and then exempts its own parameters from that reflexivity; that is an inconsistency and it is not resolved here.**

3.6 Evidence standard, and the de-identification convention

Sections 4 to 7 are comparative legal exposition. **Every enactment, decided case, article number, rate, fraction, formality and jurisdictional characteristic in them is described from the source manuscript and is not cited to the bibliography.** The general instruments and comparative doctrine named in Section 1 are cited where they bear; nothing else is.

The jurisdictions the abstract names are the ones the source manuscript works from: a statutory private foundation enacted in 2018 in a common-law enclave within a Gulf financial free zone, and the reserve, compensatory levy and inheritance tax of one continental European system. Sections 4 to 8 refer to them functionally — "the apex jurisdiction", "the forced-heirship state" — for one reason only: **a functional description cannot be misread as a statement of that jurisdiction's law, and a named one can.** No claim is thereby hidden; every proposition is in the text, and Appendix B names what must be confirmed. Where the source names a jurisdiction in a way that carries information the argument needs — the increased-monitoring period of Section 4.2, for instance — the name is used.

Two consequences follow and both are applied without exception. **First, no proposition of law below is asserted as a finding of this paper.** Each is an input to arithmetic or to a decision rule, and Appendix B lists every one that requires confirmation against a primary source. **Second, where a table scores a named jurisdiction on a characteristic the bibliography cannot support, the jurisdiction is de-identified.** The apex jurisdiction in the twelve-item assessment of Appendix C appears as **Gulf centre A**, following Curzi [2026], which de-identified the two Gulf centres in its own scoring table for the same reason. **Where a jurisdiction is named in prose it is named because the source names it, and the accompanying description carries the same uncited badge.**

4. What an apex must do

Four requirements, stated before any instrument is chosen, because the commonest error in this field is buying an instrument for a job it does not do.

4.1 An apex must be a person, not a relationship

A common-law trust is a relationship with no legal personality. The trustee owns the property at law; the beneficiaries hold equitable interests; the trust itself is not a thing that can be introduced to a company registrar, a notary or a bank onboarding team without a long conversation. That conversation is survivable in a common-law financial centre. It is expensive in a civil-law one, difficult in several Asian jurisdictions, and in some it produces a result nobody intended, because the local court has no category for an owner who is not the owner.

A foundation is a legal person: registered, named, chartered, capable of being the shareholder of record, of signing, of suing, and of being recognised by a legal system without that system first having to accept the division of legal and beneficial ownership. **Recognition of a foreign legal person is close to universal; recognition of the trust's ownership split is not, and the gap between those two facts is the whole argument of Section 6.**

For a family entirely inside the common-law world this requirement carries no weight. For the carried case — a dual national, Gulf tax resident, three heirs in three jurisdictions, one under a forced-heirship regime, assets in five states — it carries most of it.

4.2 An apex must sit outside the coordinating perimeter

Curzi [2026] identifies the coordination perimeter as the correct unit of analysis for concentration risk and shows that an entity map reading 3.4085 effective jurisdictions collapses to exactly 1.0000 effective perimeters. **That mechanism is not evaded by moving between two jurisdictions inside one perimeter.** Those are rearrangements inside a bloc, and rearranging inside a bloc does not change N.

A bridge-bloc apex is not a claim that its jurisdiction is beyond regulatory pressure — the source records a period of increased monitoring for the Gulf jurisdiction in question, from March 2022 to February 2024, and that record is described rather than cited. It is a claim about the *coordination* variable: that the jurisdiction responds to standard-setting and peer-review processes on its own timetable [Financial Action Task Force, 2012, Organisation for Economic Co-operation and Development, 2017], maintains relations across blocs that no member of the Western perimeter maintains, and is not subject to that perimeter's listing mechanism [de Goede, 2011].

4.3 An apex must be governed by a court that publishes

Neutrality is worth nothing without enforceability. **A jurisdiction outside the coordinating perimeter but with an unpredictable judiciary is a worse asset than a coordinated jurisdiction with a good one, because the failure mode is uninsurable.** The disputes an apex actually generates are governance disputes — whether a councillor was validly removed, whether a distribution was properly made, whether a reserved power was exercisable, whether a beneficiary may see the accounts — and they are decided by a court reading a charter and a statute. **A court that publishes its reasoning is a court whose future behaviour can be priced, and that is the only sense in which legal risk is ever insurable.**

The court described in the source is an English-language common-law court within a civil-law federal state, established by enclave legislation, with an international bench and an opt-in jurisdiction permitting unconnected parties to confer jurisdiction by written agreement. **Its**

constituting instruments, its bench composition, its enforcement protocol with the surrounding onshore courts and the current position on onward enforcement are described from the source and are Appendix B items 3 and 4.

4.4 The founder must be able to keep his hands on it — within limits

This is the requirement that gets misused, so it is stated with its limit attached. Statutory foundation regimes commonly permit the founder to reserve powers, and that is a material advantage over a traditional discretionary trust, where extensive settlor control is a well-known route to a sham finding.

But a power reserved under the apex jurisdiction's law is valid under that law and is simultaneously a fact that every other jurisdiction with a claim on the founder will use against the structure. Reserved powers create corporate-residence exposure where the founder sits, attribution under several civil-law regimes, and alter-ego arguments in creditor litigation — and they do nothing whatever against the failure mode that has actually destroyed offshore structures in the recorded case law: a court with personal jurisdiction over the founder ordering him to unwind it. Curzi [2026] carries that record.

4.5 The job, stated narrowly

```
DOES:
  hold      -- own subsidiary shares as a single perpetual owner
  continue  -- survive the founder's death without probate in five places
  insulate  -- keep the operating businesses from fragmenting among heirs
  travel    -- present as a legal person in every legal family
  govern    -- give a written constitution to family capital decisions

DOES NOT:
  hide      -- registered, charter filed, ownership particulars recorded
  defeat    -- a court with jurisdiction over the founder is not stopped
  exempt    -- no tax is removed where the family or the assets are
  operate   -- no substance, no staff, no trade
  bridge    -- not a payments entity
```

Run the disclosure test of Curzi [2026] against that list. Every item in the first block survives disclosure without alteration. **Perpetual single ownership works *better* when it is publicly known**, because the registrar in one state and the bank in another can both see who the shareholder is and stop asking. The apex is an instrument designed to be looked at, which is the open-design principle of Kerckhoffs [1883] — a system's effectiveness must not depend on the secrecy of its design — restated as *open design* among the protection principles by Saltzer and Schroeder [1975], and applied here to a legal rather than a cryptographic system. Brandeis [1914] is the public-interest form of the same observation from the other direction.

5. The instrument: a statutory private foundation

Everything in this section is described from the source manuscript. No statute, article number, regulation or case is cited, and Appendix B lists each item requiring confirmation.

5.1 The ownerless owner

Three forms of holding entity exist and the differences are structural rather than cosmetic.

```
COMPANY      members own it; ownership transfers; value accrues to
             owners; the death of an owner moves shares.
TRUST        no legal personality; trustee holds legal title;
             beneficiaries hold equitable interests; a relationship
             governed by equity and the trust instrument.
FOUNDATION   legal personality; NO members, shareholders or owners;
             an endowment given legal form and a purpose; governed by
             charter and council, not by owners.
```

The last line is the one to sit with. **A foundation is an *ownerless owner*.** Once the founder endows property to it, that property belongs to the foundation; there is no residual shareholder to whom it reverts, no share register to inherit, no equity to be seized as an asset of the founder's estate, and no member whose divorce, bankruptcy or death moves a block of ownership.

That single characteristic — property with a legal owner but no beneficial owner in the proprietary sense — is what allows a family enterprise to survive three generations without being cut into pieces at each one. It is also exactly what makes tax authorities and compliance officers attentive, and why the disclosure architecture around foundations has tightened everywhere. The structure must be built to withstand the second fact while capturing the first.

Separate legal personality is what does the work. The apex may in its own name hold property, hold shares in companies in any jurisdiction, enter contracts, open accounts, sue and be sued, and appear as the registered shareholder in a foreign register. **Four ownership filings in four jurisdictions that agree with each other is a compliance asset, not a compliance cost.** This is asset partitioning performing the function Hansmann and Kraakman [2000] identifies as the irreducible contribution of entity law; Pistor [2019] is both the best scholarly account of how legal coding creates capital and its sharpest critique, and both readings apply here.

A foundation may be established for a purpose, for beneficiaries, or for a combination. **A purpose foundation has no beneficiaries and therefore nobody with a proprietary claim** — no interest for a creditor of a beneficiary to attach, none to be valued in a divorce, none forming part of a deceased beneficiary's estate. What exists instead is, at most, an expectancy. A beneficiary foundation is easier to bank, to explain and to administer. The combined form — purpose in the charter, discretionary class in the by-laws — is the usual answer for a family enterprise and is what the carried case uses. **The trade-off is enforceability: where nobody has a proprietary claim, nobody has obvious standing to enforce the charter against a delinquent council.**

5.2 Charter and by-laws: a disclosure design, not a hiding place

The constituting document is filed with the registrar; by-laws, where used, are held privately by the foundation and its registered agent.

```
CHARTER  (filed, obtainable)
  name; objects in general terms; initial endowment; the existence
  and composition of the Council; the Guardian if any; the registered
  agent/office; duration; amendment mechanics; what happens to assets
  on winding up.

BY-LAWS  (private, held by the foundation and its agent)
  the beneficiary class and how it is defined; distribution standards;
  the Council's internal procedure; the founder's reserved powers in
  operative detail; appointment and removal machinery; the succession
  of Council and Guardian; investment and reserved-matter policy.
```

The temptation is to put everything sensitive in the by-laws and treat them as unreachable. **They are not.** The registered agent holds them under regulatory obligation; the home court

can order their production; a properly framed exchange-of-information request [Organisation for Economic Co-operation and Development, 2017] can reach them; a bank will require sight of them at onboarding and retain a copy. **Treat the by-laws as private but discoverable, and draft every line to be read aloud in a courtroom by an unfriendly advocate. A structure that fails when the by-laws are read is a concealment structure and fails the disclosure test by construction.** The legitimate use of the split is narrower: it keeps sensitive family information out of a register anyone can search, while keeping every element available to any authority with a proper basis to ask.

The organs. The *Council* governs, holds and administers the property, and is bound to act in accordance with the constituting documents. Council composition is the most consequential choice after the drafting itself, because **the Council is what a foreign court will look at when it asks who actually controls the assets.** The *Guardian* supervises, holding the Council to the charter, with consent typically made a condition of specified powers; in a purpose foundation the Guardian is the enforcement mechanism that the absence of beneficiaries removes. **The design error is appointing the founder's oldest friend and giving him a consent right over everything**, which produces a rubber stamp — worse than nothing, because it manufactures a paper trail of approvals that were never scrutiny — or a deadlock. The *Registered Agent* maintains the address, holds the statutory registers, makes the filings, and **is the practical disclosure surface of the whole structure.**

5.3 Founder control, and what retaining it costs

The powers a founder is commonly advised he may reserve include appointing and removing councillors; appointing and removing the Guardian; amending the constituting documents; adding or excluding beneficiaries; directing or vetoing distributions; approving disposals of core assets; consenting to a change of governing law or domicile; and appointing a successor to his own reserved powers.

Now the price, which is the part usually omitted.

Residence. Most tax systems locate a legal person where it is centrally managed and controlled, whatever its place of registration. **If the founder holds a veto over every distribution and the Council does what he tells it from a house in a high-tax country, an assessing authority there has a serviceable argument that the foundation is resident with it. The registration certificate does not answer that argument.**

Attribution. Several civil-law systems attribute the income of a foreign foundation directly to its resident founder or beneficiaries, by rules that key on control and on entitlement. **Reserved powers strengthen every such argument.**

Alter ego. In creditor and matrimonial litigation the question is whether the entity is genuinely separate from the individual. **A founder who reserved the power to direct distributions to himself, and exercised it, has supplied the answer.**

Succession. Reserved powers are personal. Unless expressly made transmissible with a named successor they may lapse on the founder's death — the intended design, or an accident that leaves the Council unconstrained at the worst moment.

The disciplined position is a graded reservation. **Reserve the powers that protect the structure from the Council; give away the powers that would let the Council be described as your agent.**

```
RESERVE (low attribution cost, high protective value)
  appoint and remove councillors, subject to Guardian consent
  appoint the successor Guardian
  consent to amendment of the OBJECTS
  consent to disposal of a CORE holding
  consent to change of governing law / transfer out

DO NOT RESERVE (high attribution cost, low protective value)
```

```
    power to direct a distribution to a named person
    power to direct a distribution to the FOUNDER
    general power to revoke the foundation and take the assets back
    standing veto over ordinary investment decisions
    power to add oneself to the beneficiary class

  DECIDE WITH COUNSEL IN EVERY RELEVANT JURISDICTION
    power to add or exclude beneficiaries
    power to amend the by-laws
```

The two "decide" items are genuinely borderline: useful, and exactly the powers that attribution and alter-ego arguments feed on. **The answer depends on where the founder and beneficiaries are resident, and it is not a question about the apex jurisdiction at all.**

5.4 The firewall, and its real reach

Foundation and trust statutes in this field commonly contain firewall provisions: rules that questions of the entity's validity, the transfer of property to it, and the rights of beneficiaries are governed exclusively by local law, and that inconsistent foreign judgments and foreign law — expressly including foreign forced-heirship rules — are not given effect locally.

IT BINDS	IT DOES NOT BIND
the home court	any foreign court
a claim brought at home	a claim brought elsewhere
the disposition of home-situs assets	assets situated elsewhere
the internal validity question	the founder's personal obligations

The asymmetry is the point. A firewall is a rule addressed to the local judge. It has no addressee abroad. A civil-law court ruling on a locally situated asset will apply its own law; a court in the founder's residence asked to make orders against him personally will make them.

The firewall's genuine function is narrower and still valuable. It makes the home layer resistant to a claimant who must at some point come to the home forum to make a claim bite on the shares, and it converts a claim that would otherwise be executed by registrar entry into one that must be litigated under the home law. **That is a real change in cost and odds. It is a friction and a forum, not a shield, and any adviser who describes it as a shield is selling something.**

The record says the same from the other direction, and Curzi [2026] carries it: the strongest non-recognition statute anyone drafted failed repeatedly against contempt power, because the target of the order was the principal and not the structure.

5.5 Migration, and why it is a resilience feature

Statutory foundation regimes commonly provide for the continuation of a foreign foundation inward and for transfer outward. **Continuation is not a transfer of assets — the entity changes its governing law while retaining its legal identity, so contracts, shareholdings and bank relationships survive the move without novation.**

This is an under-used resilience feature: the apex is not irreversibly committed. Build the path in at inception: an express charter power to transfer out and change governing law; an identified consent-holder for that power, which should not be the founder alone; a shortlist of destination regimes checked for continuation-in capability; and subsidiary constitutional and finance documents that do not contain change-of-domicile default triggers. **Point four is where this normally fails: a facility agreement at one subsidiary with a broadly drafted change-of-control clause can make the apex immovable.**

6. Foundation against trust

The claim of Section 4.1 was made without being tested. This section tests it, and the test has to be honest, because the trust is not a weak instrument. It is the most refined asset-holding device the common law produced, it has four centuries of precedent behind it, and in a large class of cases it beats a foundation outright [Hayton et al., 2016].

The argument defended here is narrower than "foundations are better".

> **The trust's central mechanism — the separation of legal and beneficial ownership — is a peculiarity of one legal family, not a feature of law in general. A structure whose apex depends on that separation being recognised has a recognition dependency, and a recognition dependency is a form of concentration the instrument of Curzi [2026] does not capture.**

6.1 The counter-position, stated fairly

Precedent. Trust law is dense with decided cases on almost every question a family will ask: what a trustee must do, what a beneficiary may see, how discretion is reviewed, what happens when a trustee goes wrong, how a mistake is unwound [Hayton et al., 2016]. Statutory foundation law is thin. **When something unexpected happens inside a trust there is usually an answer. When something unexpected happens inside a foundation there is often only a statute and an argument.**

Flexibility. A trust can be created by declaration in an afternoon, varied by the court under established doctrines, divided, merged and appointed onward with a fluency no registered entity matches. A foundation is an entity: it files, it registers, it amends by formal process.

Fiscal familiarity. Trusts are old enough that most tax systems have decided what to do about them, however unpleasant that decision may be. Foundations are frequently unclassified, and **uncertainty is itself a cost**.

So the correct question is not which instrument is better. It is: which legal families must recognise this arrangement, and what is the cheapest form that all of them can recognise?

6.2 Where the split comes from, and why the civil law cannot easily see it

The trust is a historical accident that turned into an engineering achievement. Its ancestor is the medieval *use*: land conveyed to one person to hold to the benefit of another. The common-law courts recognised only the legal owner; the Chancellor would compel the legal owner to honour the arrangement. **Two jurisdictions grew up side by side and the property came to have two owners of different kinds.** It was not designed; it emerged from a court system with two benches applying two bodies of doctrine to the same res [Lupoi, 2000].

The doctrine that grew on that root is genuinely powerful — the rule that beneficiaries of full age and capacity holding the entire beneficial interest may collapse the trust and demand the property is a statement about who really owns a thing, and has no equivalent in company law. **But every part of it presupposes a court that recognises equity as a source of proprietary rights. Take that away and the trust does not degrade gracefully. It becomes unintelligible.**

Three features of the civil-law inheritance stand in the way [Smith, 2013]. *Unitary ownership*: ownership is conceived as a single indivisible right, with lesser rights over the thing enumerated as encumbrances rather than as a second ownership of a different kind. *Numerus clausus*: the catalogue of real rights is closed, and parties may not invent new proprietary interests by contract. *No equity*: there is no separate jurisdiction of conscience producing a parallel set of proprietary remedies.

Put an ordinary discretionary trust in front of a court in that tradition and the plausible characterisations are all unsatisfactory. It may be read as a *mandate*, which makes the property the settlor's and defeats the point; as a *simulation*, a sham disguising continued ownership; as a local statutory fiduciary device, on that statute's terms rather than the deed's; or the

disposition may fail for want of a recognised form. **None of those readings gives the family what the deed said it would get, and the family usually discovers which one applies at the worst possible moment.**

6.3 What the civil law and the Islamic tradition built instead

The traditions are not without answers, and understanding them is what makes a foundation-based structure legible [Panico, 2014].

A **French-style fiduciary transfer** moves assets to a fiduciary who holds them in a segregated patrimony for a defined purpose; it was built for security and management, and gratuitous use is excluded, so it is not a succession instrument. A **German-style Treuhand** is contractual, not proprietary: protection against the fiduciary's creditors is limited and contested, and there is no proprietary tracing remedy of the common-law kind. A **Dutch-style administration foundation** holds shares and issues depositary receipts to the former holders, separating voting control from economic entitlement — a civil-law solution to the problem the voting trust solves in the common law, achieved through legal personality. **Latin American** codes contain statutory trust-like devices, usually requiring a licensed institution as fiduciary and constrained in duration and purpose. The **Islamic endowment** dedicates property in perpetuity, immobilising the corpus and directing the usufruct to beneficiaries or charitable objects under an administrator.

Look at that list and the pattern is visible. Every legal family solved the problem of property held for others and separated from its provider — and only the common law solved it by splitting ownership. Everyone else solved it either by contract or by creating a legal person.

6.4 The recognition convention, and its four limits

The obvious objection is that international law has already fixed this. The Convention on the Law Applicable to Trusts and on their Recognition [Hague Conference on Private International Law, 1985] exists precisely to make trusts recognisable in non-trust states. It is a real and useful instrument and it does not fix the problem, for four reasons.

Ratification is limited. The Convention is in force for a modest group of states. **Several of the largest civil-law economies, and the Gulf states, are not parties, and the countries where recognition is hardest are largely the countries that did not join.** The current status table should be checked against the depositary's own records before any state's position is relied on (Appendix B, item 12).

Recognition is a floor, not a transplant. A contracting state must recognise that trust assets constitute a separate fund, that the trustee may sue and be sued as trustee, and that the trustee may appear before notaries and authorities in that capacity. **That is meaningful and it is not the same as importing the whole of equity**: the recognising state is not required to give the beneficiary the proprietary remedies a common-law court would give.

The mandatory-rules carve-outs are wide. The Convention preserves the recognising state's mandatory rules in designated fields — **succession rights, including indefeasible shares, is one of them** — and preserves overriding mandatory provisions and public policy. **The very rules a family is often hoping to navigate are exactly the rules the Convention leaves untouched.**

A ratifying state proves the point in reverse. One contracting state developed a domestic practice of internal trusts — settlor, assets and beneficiaries all local, governed by a chosen foreign trust law and recognised through the Convention. It works, and it has generated three decades of national litigation about when it works. **That a device requires a specialist national jurisprudence to explain its own validity is the recognition dependency, visible.**

6.5 Why legal personality travels

Recognition of a foreign legal person is not a special regime; it is ordinary private international law, and it is close to universal [Torremans et al., 2022, Fawcett and Carruthers, 2008]. Every state has rules for determining the existence, capacity and internal governance of a foreign entity, and the two classical approaches — the incorporation theory and the real-seat theory — both produce recognition of a properly constituted entity in the ordinary case. The practical position across the developed world today is that a properly constituted foreign legal person is recognised as a legal person.

```
TRUST presented to a non-trust court
  question 1: does this system recognise a division of ownership?
  question 2: if not, what is this instead?
  question 3: who, in our categories, owns the assets?
  -> the answer to (3) is frequently "the settlor" or "the trustee
     personally", and both are wrong for the family.

FOUNDATION presented to the same court
  question 1: is this a validly constituted foreign legal person?
  question 2: yes -> it owns its assets.
  -> the ownership question is CLOSED at step 2. What remains is
     whether local mandatory rules reach the founder's TRANSFER to
     it, which is a real question but a DIFFERENT and SMALLER one.
```

The foundation removes the dependency, and it does not remove it by being cleverer; it removes it by being boring. A registered legal person with a charter, a council and an address is a category that exists in every legal system on earth. Nobody has to be persuaded of anything.

There is a name for that property in the literature on how institutions cooperate across incompatible frames. Star and Griesemer [1989] and Star and Ruhleder [1996] develop the **boundary object**: an artefact plastic enough to be adapted to local needs and robust enough to maintain a common identity across communities that do not share a conceptual scheme. **A registered legal person is a boundary object between legal families, and a trust is not**, because the trust demands that the receiving community adopt the sending community's conceptual scheme before the artefact can be read at all. That is the sharpest available statement of this paper's H1 and it comes from a literature with nothing to do with law.

6.6 Recognition is not enforcement, and the difference costs money

There is a second question underneath the first that is routinely collapsed into it, and collapsing them produces structures that are conceptually sound and practically unenforceable.

Recognition asks whether the foreign court accepts that the entity exists and owns its assets. Enforcement asks whether an order obtained by that entity, or against it, is given effect where the assets or the counterparty are. **A structure can be recognised everywhere and enforceable almost nowhere, and the family discovers the difference on the day it needs a court rather than a registrar.**

For the trust, a trustee suing to recover trust property in a civil-law state must first establish its standing, which Hague Conference on Private International Law [1985] supplies in contracting states and does not supply elsewhere. **In a non-contracting state the trustee's capacity to sue as trustee is a preliminary question decided under that state's own conflict rules, and losing it is not a defeat on the merits — it is a defeat before the merits.**

For the foundation, standing is straightforward. But the judgment it obtains, or the one obtained against it, then has to travel, and judgment travel is governed by treaties and reciprocity rules that have nothing to do with foundations. **The treaty layer is thinner than practitioners assume.** The choice-of-court instrument [Hague Conference on Private International

Law, 2005] gives effect to exclusive choice-of-court clauses and the resulting judgments among contracting states; the judgments instrument [Hague Conference on Private International Law, 2019] does the broader job. Both have limited and slowly growing membership, and — **this is the point that matters for a succession structure — both exclude from their scope the subject matter a family structure is most likely to litigate.** Wills and succession, family law and matrimonial property are carved out. The commercial arbitration route is the exception that proves the rule: awards travel under United Nations [1958] with near-universal adherence [van den Berg, 1981, Born, 2021], on a procedural framework the surveyed regimes largely share [United Nations Commission on International Trade Law, 1985] — **and family succession disputes are not arbitrable in most of the systems that matter.**

Two consequences for the design, running in opposite directions.

The first favours the foundation. If the whole governance layer — charter disputes, the family agreement between the heirs, shareholders' agreements with co-investors, services agreements with advisers — sits under one court chosen by written agreement, then the disputes the apex actually generates are heard in one forum, under one law, with no cross-border enforcement step at all. **Most apex litigation is internal, and internal litigation does not need a treaty; it needs everybody to have signed up to the same court in advance.**

The second is the honest limit. Where the dispute is not internal — a claim by a foreign heir, a creditor or a divorcing spouse, brought where that person lives — no amount of forum drafting helps, because the claimant did not sign the charter and is not bound by its jurisdiction clause. **A smaller fight in an unfriendly forum is a real improvement over a larger one, and it is the entire benefit. It is not immunity, and Section 7 shows exactly where it stops.**

6.7 The Islamic tradition, on its own terms

Islamic-law jurisdictions deserve treatment on their own terms rather than as an appendix to the civil-law discussion, because the reasoning is different. The tradition has its own long-established endowment institution: property dedicated in perpetuity, corpus immobilised, usufruct directed to named beneficiaries or charitable objects, administered by an appointed administrator. **It is not a marginal institution.**

The objection to the common-law trust in this context is often described as religious hostility. **That is not accurate and it is not useful.** The genuine frictions are two and both are technical. *The ownership split has no place to sit*: the tradition, like the civil law, works with a unitary conception of ownership, and the dedication of property to an endowment is analysed as ownership passing out of the founder's hands entirely, with no beneficial owner in the equitable sense. *The fixed shares are mandatory*: the inheritance rules allocate fixed fractional shares among defined heirs and the freedom to dispose by bequest is limited, so **a discretionary trust that removes property from the estate and then distributes on a trustee's discretion is not a variation on that system; it is a different system.**

What the foundation offers here is not evasion of any of that. It is legibility. A foundation is an endowment with a purpose, a corpus, an administering body and beneficiaries — structurally the same shape as the indigenous endowment, expressed in a modern registered form. That does not make fixed shares go away. It does mean the *instrument* is intelligible to the court rather than requiring the court to import a doctrine it does not have. **This is the boundary-object property of Section 6.5, arriving in a third legal family.**

6.8 The comparison, and the decision procedure

dimension	trust	foundation	which matters when
legal personality	none	yes	whenever a registrar, notary or bank must record an owner
ownership model	split legal/beneficial	unitary, held by the entity	whenever a non-common-law court is involved
recognition outside the common law	convention-dependent, patchy	ordinary conflict-of-laws recognition of a legal person	civil-law and Sharia-influenced jurisdictions
depth of precedent	four centuries	statute plus a short case list	novel disputes; trustee misconduct; variation
settlor / founder control	reserved powers permitted in many regimes, but control is the classic sham risk	reserved powers expressly contemplated; the entity persists regardless	where the principal will not let go
perpetuity	perpetuity rules apply unless abolished	perpetual by default	multigenerational holding
non-charitable purposes	beneficiary principle bites; needs a purpose-trust statute and an enforcer	purpose foundations are ordinary	orphan and stewardship structures
public filing	usually none	charter filed; register searchable	where the family values privacy over legibility
creditor and matrimonial attack	rich doctrine both ways	thinner doctrine; alter-ego arguments run in the founder's home court	contested-family situations
tax classification abroad	usually settled, often unfavourable	frequently unsettled	several major systems
migration	change of proper law and trustee	statutory continuation, retaining identity	jurisdiction-degradation planning
ongoing cost	trustee fees, often *ad valorem*	registrar, agent, council, audit	small structures
court of first resort	the trust's proper-law court	the foundation's home court	governance disputes

Three rows carry most of the decision: recognition outside the common law, depth of precedent, and tax classification abroad. A family whose members and assets are all inside the common-law world should take the four centuries of precedent and use a trust. A family that is not should weigh the recognition row more heavily than everything else, because it is the row that produces irreversible outcomes rather than expensive ones.

```
THE DECISION PROCEDURE
1  MAP THE PEOPLE. For the founder, spouse, and every beneficiary or
   likely beneficiary in the next two generations: state of residence;
   nationality(ies); legal family of that state; forced-heirship
   exposure; whether that state is a recognition-convention party.
2  MAP THE ASSETS. For each: situs; legal family of the situs state;
   movable or immovable; whether the register will accept a trustee
```

```
   designation.
3  COUNT THE NON-COMMON-LAW TOUCHPOINTS.
     zero                      -> trust, almost certainly
     one, low-value            -> trust with a local-law patch
     two or more, or any
       material asset or heir  -> foundation apex
     fixed-share-jurisdiction
       heir or asset           -> foundation apex
4  CHECK THE TAX CLASSIFICATION of the chosen form in every state
   identified at step 1. This is where step 3 sometimes gets
   overturned.
5  CHECK THE OPERATIONAL LAYER. Will banks, custodians and registrars
   in the asset states onboard the chosen form? Ask before drafting.
6  DECIDE, AND WRITE DOWN WHY. The memorandum of reasons is worth more
   in ten years than the structure chart.
```

Step 4 deserves emphasis because it is where the elegant answer meets the fiscal one. A structure that is conceptually correct and produces punitive treatment for one family member has not solved the problem; it has moved it.

The hybrids. *Foundation over trust* is the right answer where the family is mixed: the apex is legible everywhere, and branches with common-law beneficiaries get the flexibility and precedent of trust law where it is actually useful. *Trust over foundation* is usually a legacy shape and inherits the recognition problem at the top, where it does the most damage. *Foundation as protector or enforcer* is a small, cheap use of the form that solves a disproportionate amount of trouble, because individuals die, emigrate, fall out with the family and lose capacity, and a foundation in that role is perpetual with its own succession mechanism. *The voting/economic split* — an entity holding shares and issuing depositary receipts — is available through a foundation in most jurisdictions where control should stay consolidated while economics descend normally.

Neither instrument survives the founder treating it as a pocket. The doctrines have different names — sham and illusory trust on one side, alter ego and abuse of legal personality on the other — and they converge on the same test: *did the arrangement ever really transfer anything?* **A foundation is not a way to keep control while appearing to give it up. It is a way to give up ownership while retaining a defined and disclosed governance role. If that distinction is not acceptable to the principal, neither instrument will work and the money spent on either is wasted.**

7. Succession: the transfer question

Every other threat in the register is a risk. Succession is a certainty with an unknown date. Expropriation may never happen. A creditor may never sue. A bank may never de-risk the relationship. The founder will die, and on that day every connecting factor in the structure is tested at once, in several jurisdictions, by people who did not choose each other and who are grieving.

That certainty ought to make succession the best-planned exposure in any structure. It is usually the worst, for a reason that is psychological rather than technical: **a principal will spend eighteen months and serious money on a threat with a two per cent annual probability and will not spend an afternoon on the one with a probability of one.**

Section 6 closed on the distinction that organises everything below. A foreign court accepts at step two that the foundation is a legal person and owns its assets. That closes the ownership question and leaves standing a different and smaller one: **whether the *transfer* by which the founder endowed it is one that the court's own mandatory rules reach. Forced heirship lives entirely inside the transfer question, and a structure built on the assumption that legal personality answered it will fail in the one forum where it cannot afford to.**

7.1 Three systems, one problem

Three families of law answer *how much of a person's property may he give away by choosing* in three incompatible ways.

The common law starts from testamentary freedom, subject to statutory family-provision jurisdictions under which a court may award reasonable financial provision to a defined class. That is a discretionary remedy on application, not a fixed share, and **the difference between a discretion and a share is the whole of the practical difference between the systems.**

The civil law starts from the opposite premise. A fixed fraction of the estate is reserved to defined heirs by operation of law; the testator may dispose freely only of the remainder; and the reserve is not a claim the heir must persuade a court to grant.

Islamic law starts from a third place, and it is not a variant of the second. The fixed shares allocate the whole estate among a defined set of relatives in stated fractions, and the freedom to dispose by bequest is confined to one third and may not ordinarily be directed to a person who already takes a fixed share.

The three systems do not differ in degree. They differ in what a will is: a disposition in the first, a disposition of the free portion in the second, and an instrument operating within a distribution that has already been made in the third.

7.2 The civil-law reserve: two features that matter more than the fractions

It is computed on a reconstituted mass, not on what is left. Civil-law systems do not permit the reserve to be emptied by lifetime giving. The estate is reconstituted for the computation by adding back qualifying lifetime gifts. **A founder who gives away ninety per cent of his property and dies with ten has not reduced the reserve; he has reduced the fund from which it will be paid, which is a different and worse position.**

The remedy differs, and the difference is decisive for structuring. In some systems the reserve is satisfied in money; in others the heir may reach the property itself in the hands of the person who received it. **A monetary claim is a liquidity problem. A claim in kind against a third-party transferee is an ownership problem, and only the second can break a holding structure.**

The reference system for the carried case fixes the reserve at one half of the estate with one child, two thirds with two, and three quarters with three or more, the balance being freely disposable. Reduction of excessive gifts is, since a reform of 2006, ordinarily made in **value** rather than in kind — the recipient keeps the property and owes an indemnity. **That single reform is the most important structural fact about that system's forced heirship for anyone building a holding structure, and it is routinely omitted from summaries that lead with the fractions**, because it converts most of the exposure from an ownership problem into a liquidity problem.

A second civil-law system takes the monetary route further: the compulsory portion is a claim in money against the estate equal in value to half the intestate share, the claimant is not an heir, and lifetime gifts are addressed by a supplementary claim reaching gifts made within ten years of death **on a sliding scale reducing the included value by one tenth for each elapsed year.** That system is the clearest illustration of the general rule that **in civil-law forced heirship the planning instrument is time, and the calendar is a statute rather than a matter of judgment.**

Two further systems sit closer to the ownership end, with reduction actions that can in defined circumstances be pursued against third parties who acquired from the donee, subject to statutory time limits, and with regional variation in the fractions within a single state. All of the foregoing is described from the source; the articles, periods and current texts are Appendix B items 5 to 8.

7.3 Fixed shares, and what completeness does

The Islamic law of inheritance is a complete distributive system rather than a restriction on testation, and **treating it as "forced heirship with different numbers" produces bad structures.** Shares go to a defined set of relatives in stated fractions with residue passing to residuary heirs; a bequest is capped at one third; and a bequest may not ordinarily be made to an heir who already takes a fixed share, except with the consent of the other heirs given after death.

The interaction with a holding structure is usually missed. A lifetime gift that is complete — with delivery and the taking of possession — generally removes the property from the estate before the shares are computed. So does a properly constituted endowment. **The Islamic system's own planning instruments are completeness and timing, the same two the civil law uses, and both are the opposite of concealment: a gift that is concealed is a gift that was not delivered, and a gift that was not delivered is not a gift.**

7.4 Usufruct: the civil law's own time-splitting device

A reader who has followed Section 6's argument may conclude the civil law has no way of separating enjoyment from ownership over time. **It has an excellent one: ancient, inside the closed catalogue of real rights, and registrable.** Usufruct splits a thing into the right to use and take the fruits, held by the usufructuary, and the bare ownership, held by the bare owner; on the usufructuary's death the usufruct extinguishes and the bare owner's ownership becomes full without a further transfer.

The structuring consequence is direct. **A parent may give the bare ownership of an asset to a child now, at a value discounted to reflect the retained usufruct, and keep the income and use for life. The gift is complete when made — which starts every claw-back clock — while the parent's economic position is largely intact.** This is the civil law's answer to the life-interest trust, it achieves most of the same result, and it requires no recognition argument of any kind.

Two cautions belong with it. **The valuation of the split is usually fixed by a statutory scale keyed to the usufructuary's age rather than left to the parties**, so the discount is not a negotiating variable. And several systems contain anti-abuse provisions aimed at arrangements separating usufruct from bare ownership within a family, which can reverse the intended treatment.

Every one of the three legal families has an indigenous device for separating enjoyment from ownership across time. Only the common law did it by splitting ownership itself, and that is Section 6 restated in the succession context.

7.5 Which law applies, and the levy that follows successful planning

Two structural choices divide the world's systems. **Unity or scission**: a unitary system applies one law to the whole succession; a scission system applies the personal law to movables and the law of the situs to immovables. **Under a scission system land is governed by where it is, always, and no amount of planning at the personal level changes that** — which is the technical reason that converting immovable property into securities is a succession instrument and not merely a liquidity one. **The connecting factor**: systems differ over whether the personal law is that of nationality, domicile or habitual residence, and the three can point to three different laws for the same person.

The regional succession regulation [European Union, 2012] is the governing framework for any structure with an affected heir: the law of the deceased's habitual residence at death governs the succession as a whole, displacing scission for participating states; a person may choose the law of a state of his nationality; the chosen or applicable law governs the disposable part and the reserved shares; the designated law applies whether or not it is the law of a participating state; and a designated law may be refused application where manifestly incompatible with the forum's public policy. **That the absence of a succession convention in force [Hague Conference on Private International Law, 1989] is why the problem persists outside**

the regional framework is worth stating: the global instrument exists and never entered into force, and the gap it left is the conflict-of-laws problem the whole of this section is about.

The most important practical question under that framework is **whether the *absence* of a reserve in a chosen foreign law is itself contrary to the forum's public policy.** The source records two decisions of a national supreme court in September 2017 holding in substance that a foreign law knowing no reserve is not for that reason alone contrary to international public policy, with an indication that the position may differ where the excluded heir is left in need. **Those decisions, their docket numbers and their subsequent treatment are Appendix B item 6, and nothing here is a prediction of any outcome.**

Then the legislature answered. A 2021 enactment inserted a **compensatory levy**: where the deceased or one of his children is at death a national of or habitually resident in a participating state, and the applicable law provides no reserve, a deprived child may levy on assets situated in the forum state up to the value of the reserved rights she would have had. **It applies to successions opened after a stated date, and its own status is under active contest** (Appendix B, item 4).

Whatever its fate, its existence carries a lesson stated here as a general principle rather than a national curiosity:

> **A planning technique that reliably defeats a mandatory rule attracts a legislative response, and the response is usually retroactive in effect if not in form, because it applies to deaths occurring after enactment regardless of when the planning was done.**

That is the same degradation dynamic Curzi [2026] documents for haven jurisdictions, operating in a different field, and it is a specific instance of a general organisational result: **a rule adopted as a target ceases to be a good measure of the thing it was meant to track** [Goodhart, 1975, Campbell, 1979]. It is also the strongest single argument for building succession structures on levers the system itself provides — consent, timing, and completeness — rather than on levers that exploit a gap, **because the system does not legislate against its own instruments.**

7.6 The line: planning against defeat

Every lawful lever works by changing a fact the law makes relevant. Every unlawful one works by misrepresenting a fact, or by making a change that is not real. Four questions separate them and a structure must answer all four affirmatively.

```
1  IS THE CHANGE REAL AND LIVED?
   A change of habitual residence is a change in where a person
   actually lives. A declaration is not a change. The connecting
   factor is a question of fact decided on evidence a court can see.

2  WAS THE DISPOSITION COMPLETE?
   Did the transferor actually give up dominion? A gift with an
   undisclosed side agreement to return the property is a simulation,
   and every civil-law system has a doctrine to treat it as what it
   is. Reserving a usufruct is complete and lawful because the
   reservation is DISCLOSED and PRICED.

3  WAS IT DONE BEFORE THE CLAIM EXISTED, AND OUTSIDE THE CLAW-BACK
   WINDOW?  Timing is the substance of the lawful case.

4  WOULD IT BE DESCRIBED IDENTICALLY TO EVERY AUDIENCE?
   Would the planner, the heir, the notary, the tax authority and the
   judge each be given the same account of what was done and why?
```

Question 4 does the work. **Almost every unlawful succession arrangement has at its centre a document or a fact the family is unwilling to show to one specific person — usually the disinherited heir, sometimes the notary conducting the inventory. That unwillingness is the diagnostic. A lawful arrangement is one that the excluded heir may be shown in full, and which she may resent, and which nonetheless stands.**

lever	lawful form	unlawful form
residence	actually moving, and living there	declaring a residence not lived in
nationality election	electing the law of a state of which one is genuinely a national	electing a law of a nationality not held, or misdescribing the facts
lifetime gift	complete transfer, disclosed, valued, dated, outside the window	transfer with an undisclosed agreement to return or to continue enjoying
usufruct reservation	disclosed, valued on the statutory scale, registered	undisclosed side arrangement giving back more than the reserved right
asset conversion	genuinely selling immovable property and holding securities	nominal sale to a controlled entity with no real change
renunciation by an heir	formal advance renunciation with independent advice, in the statutory form	renunciation obtained by pressure, or without disclosure of the estate
life insurance	genuine policy, premiums proportionate to means, beneficiary designated	policy used as a wrapper for a transfer, premiums manifestly excessive
foundation endowment	real endowment, solvency recited, control graded and disclosed	endowment the founder continues to treat as his own account

Read down the right-hand column and one property is common to all eight: each unlawful form depends on somebody not knowing something. That is not a moral observation. It is the structural claim of Curzi [2026] — that a structure whose protective quantity is an adversary's ignorance has a quantity that is falling and cannot be replenished — reappearing in the one area of law where families most want it to be untrue. Curzi [2026] supplies the mechanism from a different domain: a spoken commitment is a performative aimed at the speaker, supported by consistency pressure, self-perception from overheard testimony, and the re-pricing of later defection. **A family disclosure that states what each heir will receive works the same way, and that is why the uncomfortable document is the high-value one.**

7.7 The lever set, and why consent is the universal instrument

Every lawful lever changes one of four things — who the relevant person is, where he or the asset is, when the disposition happened, or whether the claimant consents. **There is nothing else.**

```
WHO      nationality election where genuinely a national; the identity
         of the transferor -- a gift from a spouse's own property is
         not a gift from the founder's.
WHERE    genuine change of habitual residence; conversion of immovable
         to movable property, which removes an asset from a situs rule
         it cannot otherwise escape; the situs of the entity rather
```

```
          than the underlying.
WHEN      completed lifetime gifts, made early, outside every claw-back
          window; usufruct reservation, which completes the gift now
          while preserving the income.
WHETHER   the consent of the person entitled -- the only lever that
          works in all three legal families.
```

Each of the three systems provides a lawful mechanism by which the person entitled to a protected share may give it up, and each requires formality, independent advice, and full disclosure of what is being given up. One permits an advance renunciation of the reduction action in a specific notarial form; another permits a contractual waiver of the compulsory portion concluded with the future deceased and recorded before a notary; the third permits the heirs to settle among themselves after death by mutual agreement, and permits a bequest exceeding the cap, or a bequest to an heir, where the other heirs consent after death.

> **The lawful defeat of a protected share is consent, consent requires disclosure, and a family that cannot obtain consent has learned something true about itself that no structure will fix.**

Every lawful timing lever has a statutory clock and the clocks do not agree with each other. Build one calendar per heir jurisdiction with the transfer dates on it, and confirm in each what starts the period and whether a reserved usufruct or a retained benefit postpones the start. **A reservation of enjoyment frequently stops the clock from running at all, which converts the family's best lever into its worst, and it is the commonest technical failure in this field.**

8. Results

CONSTRUCTED CASE throughout. Every fact and figure below is invented for illustration. No real person, family, entity or transaction is described. Figures are illustrative and are not benchmarks or quotations.

The principal: a dual national, tax resident in a Gulf jurisdiction for eleven years, aged sixty-three. Three adult children — one resident in a common-law Asian centre, one resident and a national in a civil-law state with forced heirship, one resident in a common-law European state. Assets: a promoted manufacturing operation in a non-aligned state; a trading company in the Gulf; an IP portfolio; commercial real estate converted to a securities position; a listed-securities portfolio; a captive insurer; a fund vehicle. **Constructed total USD 240 million**, distributed on the designed allocation of Curzi [2026]:

branch	share	value (USD m)
Gulf holding and trading	29.1%	69.84
non-aligned manufacturing	18.2%	43.68
Western operating company	16.4%	39.36
Atlantic captive	13.6%	32.64
second non-aligned centre vehicle	11.8%	28.32
bridge treasury and listed securities	10.9%	26.16
total	**100.0%**	**240.00**

8.1 The touchpoint count, and H1

person / asset	state	legal family	recognition-convention party?	reserve?
founder	Gulf	civil / fixed-share influenced, with a common-law enclave	no	n/a as resident
child 1	common-law Asian centre	common law	yes	no
child 2	civil-law European state	civil law	signed, not ratified	**yes**
child 3	common-law European state	common law	yes	no
manufacturing operating company	non-aligned state	civil law	—	—
trading company	Gulf	civil / fixed-share	—	—
IP portfolio	enclave-held	common law	—	—
securities position (ex-real estate)	Western	common law	—	—
captive insurer	Atlantic	common law	—	—
fund vehicle	second non-aligned centre	common law	—	—
listed securities	common-law Asian centre	common law	—	—

```
NON-COMMON-LAW TOUCHPOINTS              MATERIAL?   RESERVE RISK?
------------------------------------------------------------------
child 2, resident and national            yes         yes
founder's residence (onshore)             yes         no
manufacturing operating company           yes         no
trading company                           yes         no
------------------------------------------------------------------
COUNT                                      4
```

The decision rule of Section 6.8 is that two or more material touchpoints, or any fixed-share-jurisdiction asset or heir, points to a foundation apex. This case has four, one of them a forced-heirship heir and two of them assets in states that do not recognise the ownership split. The rule is not close to being tested here; it is answered several times over.

What happens at each touchpoint is the comparison that matters, because it is what a specific official in a specific office actually does:

touchpoint	trust apex	foundation apex
non-aligned company registrar, recording the operating company's shareholder	asked to record a trustee holding for beneficiaries it cannot categorise; likely records the trustee personally, which is not what anyone intended	records a named foreign legal person; the question closes
Gulf onshore registrar and bank, on the trading company	trust presented to an institution in a unitary-ownership system; long onboarding, possible re-characterisation	ordinary foreign shareholder; ordinary onboarding
civil-law court, on child 2's reserved-portion claim	must first characterise the trust and may conclude the settlor never parted with the property, which is the worst available outcome	accepts that the foundation owns the assets; the fight is about the founder's transfer, which is a narrower claim
civil-law tax administration, on child 2	its trust regime applies on its own definitional terms	classification genuinely unsettled
common-law-resident children	trust is ideal: precedent, flexibility, familiar taxation	foundation is unfamiliar but workable

Read the third row against the fourth. The foundation converts a recognition problem into a tax classification problem, and that is a trade the family should take, because a classification problem is expensive and a re-characterisation of ownership is irreversible. The threat taxonomy of Curzi [2026] governs the choice on recoverability rather than on expected cost.

H1 is supported on the touchpoint count and on the second half of its prediction, which Section 8.5 measures.

8.2 The cost crossover, and H2

The recognition argument is qualitative and the cost argument is not, so it is worth the arithmetic rather than an assertion that foundations are cheap.

apex annual cost, CONSTRUCTED	USD / year
fixed-fee foundation apex	
registered agent, registered office, filings	18,000
council fees, including one independent professional	45,000
accounting, audit, statutory registers	22,000
counsel retainer and annual document review	35,000
total	**120,000**
offshore discretionary trust, *ad valorem*	
at 35 bp on USD 240m	840,000
at 15 bp on USD 240m	360,000
at 10 bp on USD 240m	240,000
fixed-fee holding mandate (constructed)	180,000

CONSTRUCTED. Illustrative, not quotations. Trustee fee bases vary widely and a fixed-fee mandate for a pure holding structure is available in the market.

The foundation's cost is fixed and the *ad valorem* trustee's is not, so the comparison has a crossover rather than a winner. From $Q^{*} = 120{,}000/b$:

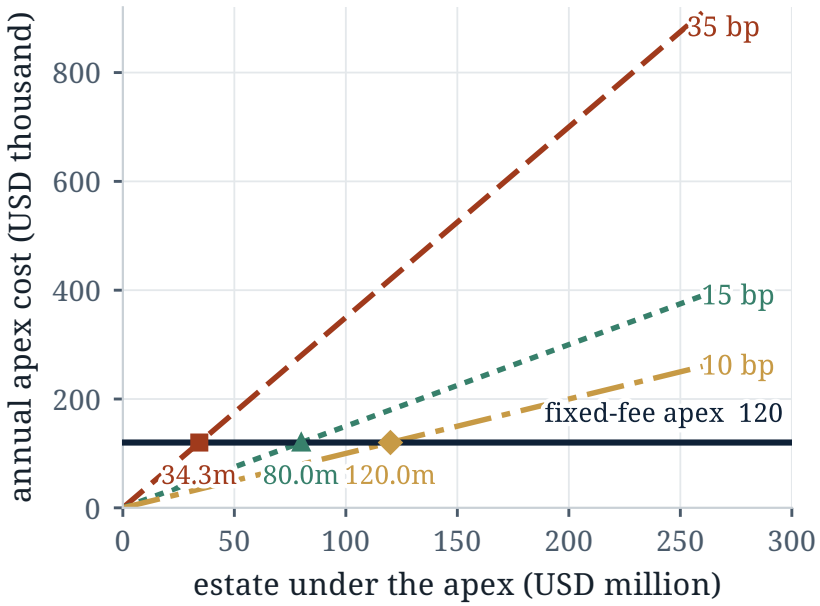


Figure 53.1: **The fixed-fee apex is the expensive form below its minimum efficient scale, and the scale moves by a factor of 3.5 across a fee band advisers treat as narrow: USD 34.3 million at 35 basis points against USD 120.0 million at 10.** Modelled from $Q^* = C_{\text{fixed}}/b$ with $C_{\text{fixed}} =$ USD 120,000 and the three CONSTRUCTED fee bases; script `figs/src/CCP-053_fig01.py`.

trustee basis b	foundation cheaper above
35 bp	USD 34,285,714 $\rightarrow$ 34.3 million
15 bp	USD 80,000,000 $\rightarrow$ 80.0 million
10 bp	USD 120,000,000 $\rightarrow$ 120.0 million

At the constructed USD 240 million the foundation costs $120{,}000/240{,}000{,}000 = 0.000500$, that is **5.0 basis points against the *ad valorem* trustee's 35, a difference of USD 720,000 a year.** Below roughly USD 34 million at a 35 basis point fee the arithmetic reverses and the fixed-cost apex becomes the expensive option. **H2's three crossovers are met exactly.**

The honest reading of the table, stated because it is the opposite of what the arithmetic flatters: cost is not the reason to choose a foundation on these facts, and a family that chooses one for the fee saving has chosen correctly for the wrong reason and will unwind it the first time an adviser quotes a fixed-fee trustee. The fixed-fee trustee mandate in the table costs USD 180,000 against the foundation's USD 120,000 — a difference of USD 60,000, which is a rounding error against the USD 27.00 million tax exposure of Section 8.3. **The reason to choose the form is the third row of the touchpoint table, and nothing in this section changes that.**

The design as built is the hybrid, and each layer must earn its cost:

```
FOUNDATION APEX                                   -- the apex
  combined form: stated non-charitable purpose in the Charter,
  discretionary beneficiary class in the By-laws
  JOB the other layer cannot do: be a recorded legal owner in a
      civil-law registry without a recognition argument
    |
    +-- GULF HOLDCO -- NON-ALIGNED OPERATING COMPANY
    +-- WESTERN OPERATING COMPANY
    +-- ATLANTIC CAPTIVE
```

```
+-- SECOND NON-ALIGNED CENTRE VEHICLE
+-- ENCLAVE TRUST                              -- one branch only
      holds the listed-securities portfolio for the two
      common-law-resident children
      JOB the other layer cannot do: give those two children
          four centuries of settled doctrine on trustee duty,
          review of discretion and information rights, in a form
          their own tax systems already know how to tax
```

Note what the trust is *not* asked to do. It does not hold the operating interests, holds nothing the forced-heirship claim would reach, and is not the apex. **The trust is used precisely where it is strong and nowhere else, and that is the argument of Section 6 stated as a structure rather than as a proposition.** The second layer adds a constructed USD 65,000 a year, taking the apex cost to USD 185,000, or $185{,}000/240{,}000{,}000 = 7.7$ basis points.

8.3 The two exposures, and H3

Notional entitlement. Had the forced-heirship law governed, the reserve with three children is three quarters of the estate — $0.75 \times 240.00 =$ USD 180.00 million — divided equally, so child 2's notional share is **USD 60.00 million**, with a disposable portion of USD 60.00 million.

$E_{reserve}$*, as found.* The founder owns an apartment in the forced-heirship state valued at EUR 6.20 million, held directly. At a constructed rate of USD 1.08 per euro that is USD 6.70 million. The compensatory levy of Section 7.5 permits a deprived child to levy on assets situated in that state up to the value of her reserved rights. **The levy is therefore capped by the local-situs assets, not by the notional entitlement:**

```
levy = min(60.00m, 6.70m) = USD 6.70m = 11.2% of the notional share
```

$E_{reserve}$*, as designed.* The designed distribution contains no asset situated in the forced-heirship state. On those facts the levy has nothing to attach to and recovers zero. **That result is a fact about the asset map, not a legal victory, and it survives exactly as long as the asset map does.**

E_{tax}. Child 2 has lived in the state for twenty years. Its inheritance tax reaches assets received by an heir resident for at least six of the ten years preceding the transfer, **wherever those assets are situated, and independently of the deceased's residence**. On a receipt of USD 60.00 million at the top marginal rate in the direct line the tax is approximately

```
0.45 x 60.00m = USD 27.00m
```

before the small relief given by the lower brackets and the per-child abatement.

exposure	amount (USD m)	reachable by the structure?
notional reserved entitlement	60.00	no — it is a computation, not a claim
E_{reserve} as found, situs-capped	6.70	yes — by the asset map
E_{reserve} as designed	0.00	achieved, and contingent on the map
E_{tax}	$\approx$ 27.00	**no**

```
E_tax      approx. USD 27.00m       ratio 27.00 / 6.70 = 4.03
E_reserve          USD  6.70m       as found -- zero as designed
```

H3's prediction is met at the stated ratio. The family's advisers had spent three years on the smaller exposure and had never computed the larger one, **and the larger one is the one the structure cannot touch.**

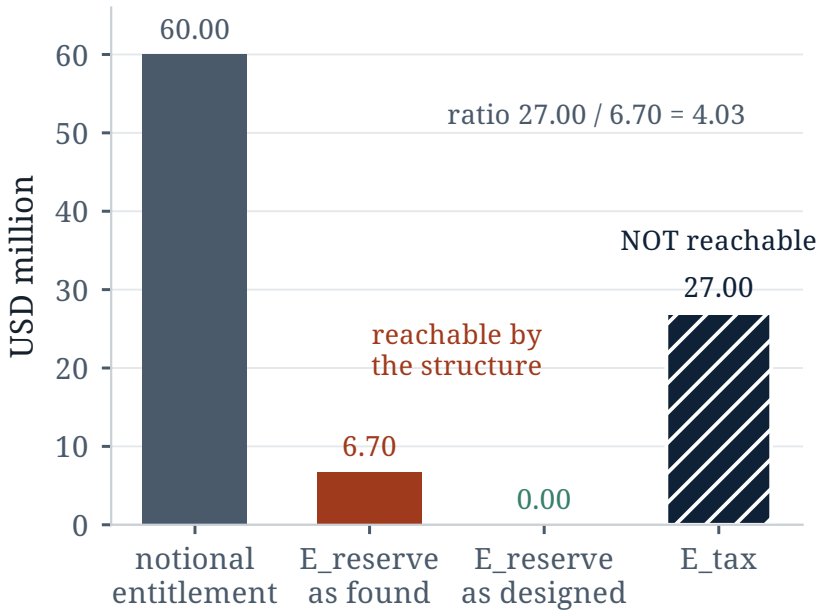


Figure 53.2: **The exposure the structure can reach is 6.70 and the exposure it cannot is 27.00, a ratio of 4.03 — and the family's advisers had spent three years on the smaller one.** Modelled from a situs-capped levy $\min(60.00, 6.70)$ against an inheritance-tax exposure 0.45×60.00; all parameters CONSTRUCTED. Script `figs/src/CCP-053_fig02.py`.

8.4 The liquidity test, and H4

Suppose child 2 is equalised toward her notional USD 60.00 million rather than litigated with. Liquid assets at death are the listed-securities position of USD 26.16 million.

```
claim                    60.00m
liquid assets            26.16m
shortfall                33.84m   =  14.1% of the estate
```

The shortfall is larger than any single liquid holding, and closing it in cash requires selling an operating branch — the Western position at USD 39.36 million or the Atlantic captive at USD 32.64 million. That is the fragmentation the apex existed to prevent, arriving through the succession rather than through the register. H4's prediction is met.

The lawful answers are all pre-funded: a policy written outside the estate, a sinking fund at the apex, an instalment arrangement documented before death, or an equalisation in specie of an asset the heir will take. The worked case uses a constructed policy benefit of USD 20.00 million with premiums proportionate to means, which reduces the cash shortfall to $60.00 - 26.16 - 20.00 = 13.84$ million, **5.8 per cent of the estate and within the range a sinking fund can carry.** The policy's own limit is that manifestly excessive premiums may be brought back into the mass, which is Appendix B item 9.

8.5 The apex does not move the concentration reading, and H5

The measurement of Curzi [2026] over the conventional four-jurisdiction structure:

```
AS FOUND
  UK             40.9%   perimeter = WEST
  US             27.3%   perimeter = WEST
  Jersey         18.2%   perimeter = WEST
```

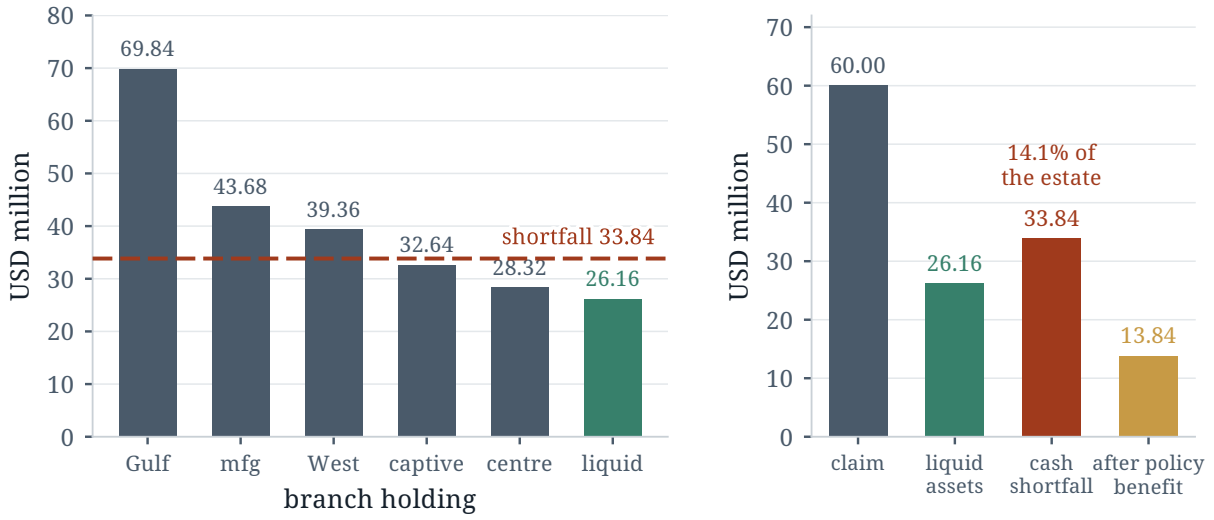


Figure 53.3: **Liquidity binds before entitlement does: the shortfall of 33.84 exceeds every liquid holding and is closed only by disposing of an operating branch, which is the fragmentation the apex was built to prevent.** Measured from the constructed branch allocation of Section 8 and the claim of Section 8.3; script `figs/src/CCP-053_fig03.py`.

```
  Luxembourg   13.6%   perimeter = WEST

  naive HHI = 0.2934   ->  effective jurisdictions = 3.41
  bloc  HHI = 1.0000   ->  EFFECTIVE PERIMETERS   = 1.00
  Western-perimeter exposure = 100.0%


AS DESIGNED (CCP-052)
  naive HHI = 0.1890   ->  effective jurisdictions = 5.29
  bloc  HHI = 0.3400   ->  EFFECTIVE PERIMETERS   = 2.94
  Western-perimeter exposure = 30.0%
```

Both readings are invariant under substitution of the apex form, because the effective count is a function of the exposure share vector and the correlation matrix and of nothing else, and neither object contains any information about the legal form of the owner. Substituting a trust apex for a foundation apex, or a company for either, leaves every share unchanged and therefore leaves HHI, N and $\ln N$ unchanged. **H5 is met, and it is met as an identity rather than as a measurement.**

That is the counter-result and it is the reason H1 exists. The apex produced none of the concentration numbers; the branch jurisdictions did. What the apex produced is the ability to hold six branch positions as one owner without cross-exposure, and to keep holding them through a death. **A structure element that does not move the concentration measurement is not thereby worthless; it is answering a different threat — and the corollary is that the concentration measurement is not a complete account of a structure's exposure.**

Two design rules govern the shape beneath the apex and both follow from the same instrument. *Hold, do not cross-hold*: each branch is held directly and holds no interest in any other, because cross-holdings convert several independent perimeter exposures into one, which is the error Curzi [2026] measures. *One layer of separation per job*: an intermediate holding company exists where a branch needs an owner that can execute shareholder actions, hold intercompany debt and be a joint-venture counterparty — jobs an apex should not do.

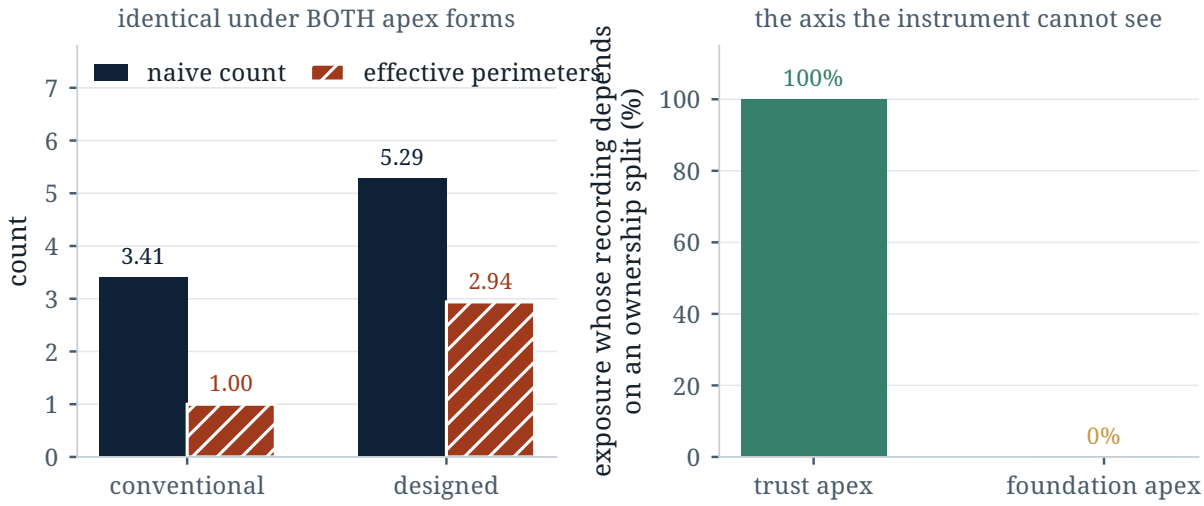


Figure 53.4: **The apex form is invisible to the concentration instrument and decisive to the recognition dependency: the effective count is identical under both apex forms, while the share of exposure whose recording depends on a foreign court accepting a division of ownership moves from 100 per cent to zero.** Effective counts measured from the share vectors of Curzi [2026] and Curzi [2026]; recognition-dependent share computed from the touchpoint table of Section 8.1. Script `figs/src/CCP-053_fig04.py`.

8.6 The Sharia counterfactual, and R5

CONSTRUCTED VARIANT. Suppose the same estate, with the founder's succession governed by fixed shares as applied in his state, survived by a widow, two sons and one daughter, with no election available.

taker	share	value (USD m)
widow	1/8 of 240.00	30.00
each son	2/5 of the residue 210.00	84.00
daughter	1/5 of the residue 210.00	42.00
total		**240.00**

The residue is $7/8 \times 240.00 = 210.00$ million, divided among two sons and one daughter in shares of $2:2:1$. Bequest is confined to one third — $240.00/3 =$ USD 80.00 million — and may not ordinarily be directed to a person already taking a fixed share without the consent of the other heirs.

The counterfactual's structural interest is that it is determinate rather than discretionary. There is no computation of a reserve against a reconstituted mass, no reduction action, no liquidity gap arising from a monetary claim of uncertain size: the whole estate is allocated by fractions known in advance. **A structure facing a fixed-share system knows its exposure exactly on the day it is built, which is the opposite of the civil-law position and is rarely stated.**

Two further observations. A completed lifetime gift, with delivery and taking of possession, generally removes property before the shares are computed, **so completeness is decisive here exactly as in the civil law.** And post-death settlement among heirs by consent is the same lever as the civil-law renunciation and the contractual waiver, arriving from a different tradition. **Consent is the universal instrument, and it is universal because each of these systems is protecting the heirs rather than the state.**

8.7 The disclosure test, applied in print

Would this structure still do its job if every element were disclosed to every authority with jurisdiction over any part of it?

- The charter is on a searchable register. Disclosure changes nothing.
- The by-laws and the trust deed would be produced on a properly framed request. Both contain a family beneficiary class and distribution standards, and neither contains an element whose function depends on not being read.
- Ownership particulars are already filed with the registrar; exchange reporting occurs by design [Organisation for Economic Co-operation and Development, 2017] and reaches the residence jurisdictions of the founder and beneficiaries. **This is anticipated in the planning, not defeated by the structure.**
- Every operating entity beneath the apex names the foundation in its own jurisdiction's ownership filings, and the filings agree with each other.
- The endowment deed, schedule, valuations and solvency recital would be produced to a notary conducting an inventory. They are dated, complete and consistent with the founder's later conduct.
- Child 2 has been shown the whole structure, has her own counsel, and has the schedule annexed to her renunciation. **The disclosure is the mechanism by which the renunciation becomes unattackable, not a cost paid to obtain it.**

The structure passes. **The two elements that would have failed it are the ones that were not built: a trust apex whose recognition in a civil-law forum depends on the court not looking too closely at what a discretionary trust is; and an undocumented understanding that child 2 would be "looked after separately", which would have been produced in the foreign proceedings as evidence that the endowment was never intended to be complete.**

9. Hypotheses tested, and what was not solved

H1 — supported, with the second half measured rather than argued. Four material non-common-law touchpoints, and an effective count identical under both apex forms (Section 8.5). The falsifier was a system in which the ownership question closes identically for a relationship and for a registered legal person; the touchpoint table of Section 8.1 gives four offices in which it does not.

H2 — supported exactly. $Q^{*} = 120{,}000/b$ gives USD 34.3m, 80.0m and 120.0m at 35, 15 and 10 basis points. **The hypothesis was stated as a scale result and not as a recommendation, and Section 8.2 refuses the recommendation the arithmetic appears to license.**

H3 — supported at the stated ratio. $27.00/6.70 = 4.03$.

H4 — supported. Shortfall USD 33.84 million, 14.1 per cent of the estate, exceeding every liquid holding.

H5 — met, and met as an identity rather than as a measurement. That distinction matters: **an identity cannot be evidence for anything.** H5 is true by the construction of the instrument, and its content is entirely in what that construction reveals about the instrument's blind spot rather than in the numerical result. **Reported here as a hypothesis whose confirmation carries no empirical weight, which is a weaker outcome than it looks.**

What the structure did not solve, stated plainly.

- **The tax exposure**, at roughly four times the reserve exposure, unchanged by anything in the structure and dependent only on where the heir lives.
- **The classification question** — whether a foreign regime for trusts reaches a foundation — which is unsettled in at least two major systems and changes the analysis materially either way.

- **The thin-doctrine risk.** Trust law has a doctrine, narrowed but real, for unwinding a fiduciary's mistaken exercise of a power. **A statutory foundation regime enacted in 2018 has no equivalent, and it is the one a family is most likely to need.** Where there is no doctrine to fix a mistake, the only available control is not making it. This is unpriced and does not go away with good drafting; it only becomes less likely to be needed.
- **The dependence on consent.** If the forced-heirship heir refuses to renounce, the plan reverts to the levy exposure, the liquidity shortfall and a contested estate. **There is no lawful substitute.**
- **The situs constraint**, which must be maintained as a standing investment restriction for the rest of the founder's life.
- **Apex concentration.** A structure with 2.94 effective perimeters on the asset side still has one governance point of failure: one entity, one registrar, one court, one country. The migration path manages it; it does not eliminate it.
- **The maintenance bill.** At the parameters of Curzi [2026] the designed structure gives $h = -0.122417$ against the conventional structure's -0.500000. **Both are below the build threshold: the redesign improved the position by a factor of four and did not make it positive.**

10. Failure modes

1. The alter-ego collapse. The founder reserves everything, directs everything, and treats the foundation's accounts as his own; a foreign court concludes the entity is a facade. *Diagnostic:* if the founder can obtain money from the foundation for personal use on his own signature, the failure is already present. *Mitigation:* the graded reservation of Section 5.3.

2. Residence by management. Nobody moves the foundation; it moves itself, by where its decisions are taken. The result is an apex tax-resident in a high-tax jurisdiction doing none of the compliance a resident entity performs — **worse than never having built it.** *Diagnostic:* count the last twelve council decisions and ask where each was actually made.

3. Reading the firewall as a shield. Section 5.4. *Diagnostic:* the plan has no answer to "what happens if a court orders the founder, personally, to procure a distribution?"

4. Assuming personality equals recognition of the internal arrangement. This is the most important failure mode in the paper and it is the honest limit on the whole argument. A foreign court will accept that the foundation exists and owns its assets. It does not follow that it will accept the consequences the family wants. It may apply its own mandatory rules to the *transfer*, and its own tax and attribution rules to the founder and beneficiaries as persons within its jurisdiction. **Legal personality closes the ownership question. It does not close the transfer question or the attribution question, and most of what goes wrong in practice is in those two.**

5. Choosing the foundation for a wholly common-law family. *Diagnostic:* run step 3 of the decision procedure and count zero non-common-law touchpoints. *Mitigation:* use a trust.

6. Relying on the recognition convention in a state that never joined it. The memorandum says the trust "will be recognised under the Convention", and nobody checks whether the states that matter are contracting states. **A convention binding four European states is of no assistance in the jurisdictions where recognition was the problem in the first place.**

7. Assuming the charter binds people who never signed it. A jurisdiction clause governs the founder, the councillors, the guardian and — with care — beneficiaries who accept benefits. **It does not govern a divorcing spouse, a trade creditor, a tax authority, or a reserved-portion claimant suing where she lives.**

8. The thin-jurisprudence gap. Section 9. **This is the trust's strongest argument and it is not answered here.**

9. The perpetual entity nobody can end. A trust has a doctrine under which beneficiaries of full age and capacity holding the entire beneficial interest may collapse it. A purpose foundation has nobody with standing to do the equivalent. *Diagnostic:* ask who, in twenty years,

could lawfully bring the structure to an end, and whether they would agree.

10. Assuming a trust converts into a foundation. It does not. What happens is a distribution or transfer to a newly formed entity, **which is a disposition — with the consequences of a disposition in every jurisdiction with a claim, and with a fresh transfer date restarting every claw-back clock the original settlement had already run down.**

11. Buying two layers and using neither. *Diagnostic:* for each layer, state in one sentence the specific job it does that the other cannot. If either sentence is difficult, the layer is decoration. **This is the specialisation discipline of Curzi [2026] applied to legal forms rather than to entities**, and the same result applies: the overloaded single entity is cheaper on direct cost and worse on conjunction.

12. Confusing the forced-heirship problem with the inheritance-tax problem. *Diagnostic:* the plan has one figure rather than two. *Mitigation:* Section 3.4.

13. Reserving enjoyment and stopping the clock. In several systems a retained benefit prevents the claw-back period from beginning at all. **Ten years of patient planning then count for nothing.**

14. The renunciation that is attacked by the person who signed it. Obtained in a family meeting, without independent advice or disclosure of the estate's value, in the wrong form. **It is set aside, and then becomes evidence of pressure.**

15. The heir who moves. The map is a photograph. A common-law-resident child who moves to a forced-heirship state in year eight changes the applicable law, the claw-back calendar and the tax exposure simultaneously, and nobody tells the adviser. *Mitigation:* an annual re-run with a review date.

16. Maintenance decay. C is paid for three years and then quietly stops being paid. Minutes stop being taken. The independent councillor resigns and is not replaced. Filings slip. **The structure looks identical on the register and has lost every substantive characteristic that made it defensible. This is the commonest failure of all, it is invisible until it is tested, and it is the reason the surplus law puts C on the right-hand side of the equation rather than treating it as a one-off cost of entry.** Cyert and March [1963] is the general form: organisational slack is what makes a system survive a shock it did not plan for, and it has always appeared on the accounts as waste.

11. Discussion

What follows. Three things, each narrower than the general claim it resembles.

A recognition dependency is a real form of concentration and it is invisible to the published instrument. Section 8.5 establishes this as an identity rather than as a measurement, and Section 12.4 records that the invariance responsible is a defect of the instrument for this purpose rather than a virtue. **The claim is not that the effective count is wrong; it is that it is a complete account of one axis and no account at all of another.** Curzi [2026r] reaches the same structural point in a different domain and is the sharpest available statement of it: **a pair in perfect opposition and a pair in perfect agreement both read $u = 0$, $N_{\text{eff}} = 1$ and surplus zero, because the extremality condition is stated on the squared alignment.** Two configurations that are maximally different in mechanism read identically on the instrument. **That is exactly what happens here when a trust apex and a foundation apex are put through the concentration measurement, and a reader who has met the result once in the corpus should recognise it the second time.**

The cost result is a scale result and not a recommendation. Section 8.2 computes the crossover and then refuses the conclusion it appears to license. **A model whose recommendation is eliminated by a competitor's pricing decision has not established the recommendation**, and Section 12.3 gives the limit that shows why.

Consent is the universal lever, and it is universal for a reason. Each of the three legal families protects the heirs rather than the state, so each permits the protected person to give

the protection up, and each requires disclosure to do it. **The lawful defeat of a protected share is consent, consent requires disclosure, and a family that cannot obtain consent has learned something true about itself that no structure will fix.**

What does not follow. Four things.

It does not follow that a foundation apex is generally correct. It follows on a map with material non-common-law touchpoints. **On a map with none, the four centuries of precedent are worth more than the recognition property, and Section 10 lists choosing a foundation for a wholly common-law family as a failure mode rather than as a preference.**

It does not follow that the apex improves resilience. R6. Curzi [2026] runs the same collapse on a different cluster and reaches the same shape: six Atlantic-cluster entities at 40.0 / 22.0 / 14.0 / 12.0 / 7.0 / 5.0 give a naive index of 0.249800 and 4.0032 effective jurisdictions, **but a bloc index of** 1.000000**,** $N = 1.00$ **and** $\ln N = 0$**.** No apex vehicle changes any of those figures. Curzi [2026] prices the most generous counterfactual available to a Western-perimeter structure — crediting a Crown Dependency group as its own bloc, at 81.8 against 18.2 — and obtains a bloc index of 0.702248, $N = 1.423998$ and $h = -0.376286$ against the designed structure's 2.941176 and -0.122417. **Even the most generous reading of a single-perimeter structure lands closer to the conventional case than to the designed one, and the apex form is irrelevant to that comparison.**

It does not follow that more branches are better. Curzi [2026] measures the case directly and it is the cleanest counter-example in the series: **adding two further Gulf jurisdictions raises the naive effective count from 5.29 to 7.13, a gain of 34.8 per cent, while the effective bloc count *falls* from 2.94 to 2.84, a loss of 3.4 per cent.** More jurisdictions, less independence. An apex that makes it administratively easy to add branches is therefore an instrument that makes it easy to do the wrong thing cheaply, **and the assignment discipline of Section 10 failure mode 11 is what stands between the two.**

It does not follow that the liquidity result is confined to succession. Curzi [2026] computes the same forced-disposal arithmetic on a pledge trigger: **the forced sale takes an 80 per cent holder to exactly** 50.000000 **per cent at a** 36.000000 **per cent drawdown, to** 38.095238 **per cent at 40 per cent, and exhausts collateral equity at 50 per cent.** The mechanism is identical — a claim in currency against a balance sheet of operating interests — and only the trigger differs. **A structure that has pre-funded its succession liquidity and not its pledge liquidity has solved half of one problem.**

Where the instrument's weakest gate sits. Curzi [2026b] establishes that conversion composes multiplicatively through recording, binding and reach gates, so the weakest gate dominates. Curzi [2026] makes the same point as a measurement on a twelve-step schedule: **the only gate acting on the numerator sits ninth of twelve and contributes** -5.305895 **octaves, 88.93 per cent of the schedule's total movement, against 11.07 per cent for the eight steps before it.** That is the shape of this paper's argument in a different register. **An apex is one step of a long build, it acts on the recording gate, and on the multiplicative composition almost all of the available movement is there — which is why the choice is worth an afternoon of decision procedure and why the ten steps of Appendix A are ordered as they are.**

The status of the claim is measurement and comparison, not physics. Hedström and Ylikoski [2010] states the criterion the discipline named in the corpus title applies: an explanation must name a mechanism rather than exhibit a fitted curve. **The mechanisms named here are a recording requirement, a fee structure and a distributive rule, and nothing is fitted.** No order parameter is identified, no phase transition is located and no critical exponent is estimated; the vocabulary of Castellano et al. [2009] is deliberately not used. Gallegati et al. [2006] is the checklist this paper was run against — ignorance of existing work in the target discipline, sub-standard statistical procedure, unjustified universality claims, and models built on assumptions known to be false — and Section 13 reports where it passes and where it does not.

12. Verification and rigour audit

12.1 Dimensional audit

symbol	quantity	dimension	unit	status
Q	estate value / protected capital	[currency] in Sections 8.2–8.6; $[C]$ index units in the surplus law	USD; index units	**two distinct quantities under one symbol, inherited from the source and flagged here**
C_{fixed}	fixed apex cost	[currency T^{-1}]	USD per year	assumed (constructed)
b	*ad valorem* fee basis	$[T^{-1}]$	per year	assumed (constructed)
Q^*	minimum efficient scale	[currency]	USD	derived, C_{fixed}/b
s_i	exposure share of branch i	[1]	dimensionless	defined; sums to 1
HHI	concentration index	[1]	dimensionless	derived from shares
N	effective perimeter count	[#]**, a count**	perimeters	derived as $\text{HHI}_{\text{bloc}}^{-1}$
E_{reserve}, E_{tax}	succession exposures	[currency]	USD	derived from constructed inputs
reserve fraction, tax rate, fixed shares	ratios	[1]	dimensionless	assumed; described from the source, uncited
η, δ, C, h	surplus-law terms	as in Curzi [2026]	index units	assumed

Does the crossover model balance? $Q^* = C_{\text{fixed}}/b$ divides a currency flow ([currency T^{-1}]) by a rate ($[T^{-1}]$) and returns a stock ([currency]). **It balances, and the balance is the content: a fixed flow divided by a rate is a stock, which is why a minimum efficient scale exists at all.**

Does the exposure comparison balance? E_{reserve} and E_{tax} are both currency stocks and their ratio is dimensionless. **They are commensurable in dimension and are not commensurable in kind**: one is a transfer between family members and the other is a payment to a state, and the ratio 4.03 should not be read as saying the tax is "four times worse". It says the family had allocated its attention to the smaller number.

The important row is the first, and it records a defect. Q denotes an estate in USD in Sections 8.2 to 8.6 and protected capital in unnamed index units in the surplus law. **The source uses one symbol for both and this paper does the same for continuity; the two are not interchangeable and no equation in this paper mixes them.** The break-even count of Curzi [2026] is unit-free and survives; the statement that $h = -0.122417$ does not mean anything until the index unit is defined, **and the source defines none.**

12.2 Significant figures and epistemic status

- $Q^* = 34{,}285{,}714$ is **exact arithmetic** on two constructed inputs and is printed as 34.3 million because three significant figures is all the constructed inputs can support. 80.0 and

120.0 million are exact.
- 5.0 and 7.7 basis points are exact arithmetic on constructed inputs, rounded hard.
- 6.70 million carries a constructed exchange rate and is quoted to the precision of that construction, no further.
- 27.00 million is **illustrative model output**: it applies a single top marginal rate to the whole receipt and the text says so, explicitly excluding the relief given by the lower brackets and the abatement. **Quoting it to the cent would be false precision and it is quoted to two decimals of a million because the input is stipulated at that precision.**
- 33.84 million and 14.1 per cent are exact arithmetic on stipulated inputs.
- The fixed shares — 30.00, 84.00, 42.00 — are exact rational arithmetic on exact fractions and sum to 240.00 exactly.
- 0.2934, 3.41, 1.0000 and 1.00 are quoted from Curzi [2026] at that paper's own precision. **Section 6.3 of that paper records that the naive figure depends on the rounding basis of the shares —** 0.293388 **on exact fractions against** 0.293430 **on three-significant-figure percentages — and this paper quotes to four decimals precisely so that it does not have to choose.** Every perimeter figure is exactly 1.000000 under every basis.
- **No quantity in this paper is a fitted or measured estimate, so no uncertainty interval is reported anywhere. That is not a strength: it means the paper has no empirical content to attach an interval to.**

12.3 Limits and boundary behaviour

The crossover model.

limit	$Q^* = C_{\text{fixed}}/b$	reading
$b \to 0$	$\to \infty$	a zero-basis trustee is cheaper at every scale; the fixed-fee apex never wins on cost
$b \to 1$	$\to C_{\text{fixed}}$	at a 100 per cent fee the crossover is the apex's own annual cost, which is the degenerate case
$C_{\text{fixed}} \to 0$	$\to 0$	a free apex wins at every scale
$C_{\text{fixed}} \to \infty$	$\to \infty$	so does an infinitely expensive one lose
$Q \to 0$	—	the *ad valorem* form costs nothing and the fixed form costs C_{fixed}: **the fixed form is unboundedly worse in the small-estate limit, which is the domain statement**
$Q \to \infty$	—	the *ad valorem* form is unboundedly worse: **the model has no interior optimum and is monotone on each side of** Q^*

The model's domain is a positive fee basis and a positive fixed cost, and the crossover is a single point rather than a region. A fixed-fee trustee mandate — the fourth row of the cost table — destroys the crossover entirely by making both forms fixed, which is why H2 is stated as a scale result and not as a recommendation. **A model whose conclusion is eliminated by**

a competitor's pricing decision is a fragile model and the paper says so here rather than in a footnote.

The situs-capped levy. $E_{\text{reserve}} = \min(R, A)$ for notional reserved rights R and local-situs assets A. As $A \to 0$ the levy goes to zero and the structure appears to win completely; that limit is exactly the designed case and it is the most fragile result in the paper, because A is a decision variable the family can change accidentally by buying an apartment. As $A \to R$ the cap ceases to bind and the levy is the full notional entitlement. The function is continuous, piecewise linear, and has a kink at $A = R$; nothing in the structure changes the position of the kink.

The liquidity shortfall. $S = \max(0, K - L)$ for claim K and liquid assets L. As $L \to K^-$ the shortfall goes to zero continuously; as $L \to 0$ it goes to the whole claim. There is no discontinuity, and that is the finding: the structure does not fail at a threshold, it degrades continuously, and a family that funds half the shortfall has genuinely halved the forced disposal.

The concentration instrument's limits are the proof that it behaves as a count and are run in full by Curzi [2026]: HHI^{-1} returns exactly n for n equal shares and exactly 1 for a single holder, so $1 \le \text{HHI}^{-1} \le n$ with both bounds attained. Verified here on this paper's own vectors: the four-perimeter conventional collapse gives $1/1 = 1$ exactly, and a hypothetical four-way equal split gives $1/0.25 = 4$ exactly. A quantity that equals n for n equal units, equals 1 for one unit, and lies strictly between otherwise is a count.

The affine count fails the same limits. $1 + (N-1)u$ returns 1 at $u = 0$ and N at $u = 1$ for *any* share vector whatever, because it never consults the shares. Curzi [2026] prices the gap in the middle at 33.62 per cent. It is not used here.

12.4 Invariance

The effective count is invariant under relabelling and under any rescaling that preserves the share vector, which is why the same instrument applies to a USD 24 million estate and a USD 240 million one. It is also invariant under the substitution of the apex form, which is H5 and which is a *defect* of the instrument for this paper's purpose rather than a virtue. An invariance is only a good property when the thing held fixed is the thing one wants held fixed.

The crossover Q^* is invariant under currency redenomination — multiply C_{fixed} and Q by the same factor and Q^* scales with them — and it is not invariant under a change of fee structure, which is Section 12.3's fragility statement.

What is conserved, and in what units. Mirowski [1989] shows that the characteristic failure of this genre is borrowing the notation of a conserved quantity without the conservation law that gave it content. The objection is met here twice and failed once. The exposure share vector is conserved: $\sum_i s_i = 1$ exactly, dimensionless, and it is that constraint alone that bounds the index. The Sharia allocation is conserved: $30.00 + 84.00 + 84.00 + 42.00 = 240.00$ exactly, which is the whole estate, and a distribution rule that allocates exactly the whole is a conservation law with a currency unit attached — the only one in this paper. The civil-law reserve is not conserved in the same sense, because it is computed on a reconstituted mass that includes property no longer in the estate, so the fractions do not partition anything that exists on the day of death. That asymmetry is a real structural difference between the two systems and it is the reason the fixed-share system's exposure is knowable in advance and the civil-law system's is not. Cartwright [1999] supplies the discipline that makes a partial answer respectable: specify the arrangement within which the law holds and stop claiming a universality the argument does not need.

Aggregation is unaddressed. Kirman [1992] establishes that a single representative parameter for a heterogeneous group generally does not exist, and this paper writes one Q, one δ and one C for a structure whose branches plainly differ. No aggregation argument is offered and none is available from the arithmetic.

12.5 Analogy discipline

borrowing	source	category	justification
Hill numbers, order 2	Hill [1973], Jost [2006]	**1 — derivation**	HHI^{-1} *is* ${}^{2}D$; an identity
numbers-equivalent reading	Adelman [1969]	**1 — derivation**	the reciprocal of the index is the count of equal units giving that index
similarity-sensitive diversity	Leinster and Cobbold [2012], Rao [1982]	**1 — derivation**, with a stated domain	exact where the matrix has unit diagonal and non-negative entries
portfolio variance	Markowitz [1952]	**1 — derivation**	the exposure vector genuinely is a weighted sum of shocks with a correlation matrix
minimum efficient scale	none — deliberately	**not borrowed**	$Q^{*} = C_{\text{fixed}}/b$ is elementary algebra on two constructed numbers, imports nothing, and is cited to no source because none is needed
the boundary object	Star and Griesemer [1989], Star and Ruhleder [1996]	**2 — shared form**	a registered legal person is plastic locally and robust across communities; **no claim is made that legal families are scientific communities in that literature's sense**
open design	Kerckhoffs [1883], Saltzer and Schroeder [1975]	**1 — derivation**	the disclosure standard *is* the open-design requirement applied to a legal rather than a cryptographic system
asset partitioning	Hansmann and Kraakman [2000]	**1 — derivation**	the foundation performs exactly the partitioning function entity law is identified as supplying
organisational slack	Cyert and March [1963]	**3 — framing**	qualitative, cited for the shape of the maintenance-decay failure and not as a formal result

borrowing	source	category	justification
the measure-as-target dynamic	Goodhart [1975], Campbell [1979]	**2 — shared form**	the legislative response to successful planning has the same shape; no monetary or evaluative content is imported

The failure mode this table exists to prevent is letting a category-2 borrowing carry a category-1 result's authority. The concentration and crossover arithmetic is category 1 throughout and needs no help; the boundary-object and Goodhart borrowings are category 2 and are used for shape only.

12.6 Measurement-theoretic status of the scored inputs

Three inputs in this paper are scores or classifications rather than measured quantities.

The touchpoint count is a count of a nominal classification. "Non-common-law touchpoint" is a binary attribute of a person or asset, and counting instances of a binary attribute is licensed. **What is not licensed is treating the count as an interval quantity**: the decision rule of Section 6.8 uses thresholds at zero, one and two, and the text is explicit that a single *material* touchpoint outweighs several immaterial ones. **An unweighted count of non-exchangeable items supports a threshold and does not support a magnitude**, and it is used for nothing else.

The five disclosure criteria are binary and are combined by product, which is the one operation their scale type licenses. A conjunction of binary indicators is a binary indicator; their *average* is not a licensed statistic, and Curzi [2026] includes the average column explicitly to demonstrate that averaging them selects the wrong verdict.

The twelve-item jurisdiction assessment of Appendix C is ordinal at best and is not summed. Its items are heterogeneous — legal family, court quality, rail access, degradation trajectory — and no total is computed anywhere, because no common scale exists on which they could be added. **The assessment is a checklist that produces a verdict by argument, not a score that produces one by arithmetic, and it is presented that way.**

The framework's underlying treatment of Q as a ratio-scale quantity is not tested. Michell [1997] is the strongest available objection: psychology and its neighbours routinely assume quantitative structure without testing for it, and the formal remedy is conjoint measurement [Luce and Tukey, 1964, Krantz et al., 1971], whose conditions establish whether an attribute has additive quantitative structure at all. **The conjoint conditions were not tested.** Here the objection has unusually limited bite, because Q in Sections 8.2 to 8.6 is a currency amount and currency is a ratio scale by construction. **It bites in full on the surplus law's Q in unnamed index units**, which is a different quantity under the same symbol (Section 11.1), and Curzi [2026t] is the corpus's own precedent for what to do about it: that paper **demotes the moment equation on the ground that its two inputs are not instrumented on a common ratio scale with a shared true zero, and withdraws the magnitude of their quotient as a measured quantity.** The same demotion applies to any statement here that mixes the two Qs, and no such statement is made.

12.7 Postulates

#	statement	status
P1	exposure shares s_i are non-negative and sum to 1	DEFINITION
P2	HHI $= \sum_i s_i^2$; the effective count is its reciprocal	DEFINITION
P3	the effective count is a function of the share vector and the correlation matrix only	DEFINITION (this is what H5 consumes)
P4	a fixed-fee apex costs C_{fixed} and an *ad valorem* alternative costs bQ	DEFINITION (modelling choice)
P5	$C_{\text{fixed}} = 120{,}000$, $b \in \{0.0035, 0.0015, 0.0010\}$, $Q = 240{,}000{,}000$	POSTULATED (constructed)
P6	the reserve with three children is 3/4 of the reconstituted estate, divided equally	POSTULATED (described from the source, uncited)
P7	the compensatory levy is capped by local-situs assets	POSTULATED (described from the source, uncited)
P8	local-situs assets are EUR 6.20m at USD 1.08 per euro; the top marginal rate in the direct line is 0.45	POSTULATED (constructed)
P9	liquid assets at death are the listed-securities position, USD 26.16m	POSTULATED (constructed)
P10	the fixed shares are widow 1/8, residue to two sons and one daughter in 2:2:1	POSTULATED (described from the source, uncited)

12.8 The compounding chain

step	consumes	statement	status
L1	P2	HHI^{-1} returns n exactly for n equal shares and 1 for one holder	ARITHMETIC
L2	L1	HHI^{-1} is a numbers-equivalent — Adelman [1969]	CITED
L3	L2	HHI^{-1} is the Hill number of order 2 — Hill [1973]	CITED
R5a	P3	substituting the apex form changes no share, so N is unchanged: 3.41 naive and 1.00 perimeters, whatever sits at the top	ARITHMETIC

step	consumes	statement	status
L4	P4	equating C_{fixed} and bQ gives $Q^{*} = C_{\text{fixed}}/b$	ARITHMETIC
R1	L4, P5	Q^{*} is 34.3m, 80.0m and 120.0m at 35, 15 and 10 bp	ARITHMETIC
R2	P5	the apex costs 5.0 bp against an *ad valorem* 35 bp; the annual difference is $720{,}000$	ARITHMETIC
L5	P6, P5	the notional reserved share is $0.75 \times 240.00/3 = 60.00$m	ARITHMETIC
R3a	L5, P7, P8	the levy is $\min(60.00, 6.70) = 6.70$m, $= 11.2\%$ of the notional share	ARITHMETIC
R3b	L5, P8	the tax exposure is $0.45 \times 60.00 = 27.00$m; the ratio to R3a is 4.03	ARITHMETIC
R4	L5, P9	the shortfall is $60.00 - 26.16 = 33.84$m, $= 14.1\%$ of 240.00m	ARITHMETIC
R5	P10, P5	widow $240.00/8 = 30.00$; residue 210.00; sons $2/5 \times 210.00 = 84.00$ each; daughter $1/5 \times 210.00 = 42.00$; total 240.00 exactly	ARITHMETIC
CR	R5a, R3b	the apex adds nothing to N and does not reach the larger exposure	ARITHMETIC

Every ARITHMETIC row can be redone by hand. Every CITED row names one source and no more. **Six of the ten postulates are POSTULATED, four of them constructed and two described from uncitable legal sources. A chain resting on six postulates is not a proof, and this paper says so here rather than leaving the reader to notice: the arithmetic is exact and every input is either invented or uncited.**

12.9 Cross-checks performed

- The branch values were recomputed from the share vector: $0.291 \times 240 = 69.84$, $0.182 \times 240 = 43.68$, $0.164 \times 240 = 39.36$, $0.136 \times 240 = 32.64$, $0.118 \times 240 = 28.32$, $0.109 \times 240 = 26.16$; the six sum to 240.00 exactly.
- The Sharia allocation was verified as a partition: $30.00 + 84.00 + 84.00 + 42.00 = 240.00$ exactly, in exact rational arithmetic.
- The three crossovers were recomputed from C_{fixed}/b and agree with the table to the printed precision; $120{,}000/0.0035 = 34{,}285{,}714.29$, printed as 34.3 million.
- $s_{\max}^{2} \le \text{HHI} \le s_{\max}$ holds on the conventional perimeter vector trivially ($1 \le 1 \le 1$) and on the designed branch vector: $0.084681 \le 0.189002 \le 0.291$.
- The liquidity result was recomputed with the constructed policy benefit: $60.00 - 26.16 - 20.00 = 13.84$, $= 5.77\%$ of 240.00, printed as 5.8 per cent.

13. Limitations

Every case is constructed. The principal, the family, the assets, the shares, the cost model, the exchange rate, the tax rate, the fee bases and the policy benefit are invented and are labelled as such at each point of use. **No empirical dataset was collected and none is reported.**

Jurisdiction-specific legal claims are not cited to primary sources. Every enactment, decided case, article number, reserve fraction, levy condition, formality, rate and jurisdictional characteristic in Sections 4 to 8 is taken from the source manuscript. **The bibliography contains no statutes, regulations, court rules, registries or case reports for any individual jurisdiction.** Appendix B lists the affected claims. **This is a real gap and it is not closed here.** It is also the paper's largest exposure, because the arithmetic of Sections 8.3 to 8.6 takes those legal propositions as inputs: **if the reserve fraction, the situs cap or the fixed shares are stated wrongly, every number downstream of them is wrong and the error will not show up as an arithmetic inconsistency.**

The cost result depends entirely on the fee assumptions and is destroyed by one competitor pricing decision. A fixed-fee trustee mandate makes both forms fixed and eliminates the crossover. Section 12.3 states this; **it is the reason H2 is a scale result and it is why the paper refuses the cost argument in Section 8.2 rather than banking it.**

The tax figure is a single-rate approximation and is labelled as illustrative. Applying one top marginal rate to a whole receipt overstates the liability by the amount of the lower brackets and the abatement. **The direction of the error is known and its size is not**, and the ratio 4.03 should be read as an order-of-magnitude statement rather than a measurement.

The designed case's zero levy is the most fragile result in the paper. It follows from the absence of any asset situated in the claimant's state, which is a decision variable the family can change accidentally. **The levy bites on assets situated locally *at death*, and a structure with no local situs today acquires one when someone buys an apartment, inherits a share of a family property, or moves a custody account.**

The apex-neutrality result is an identity and carries no empirical weight. H5 is true by the construction of the instrument. **Its confirmation is not evidence for anything; its value is entirely in what it reveals about the instrument's blind spot.**

The recognition dependency is not quantified. H1 is supported by a touchpoint count and a table of what four specific officials would do. **No cost, probability or magnitude is attached to a recognition failure anywhere in this paper**, so the trade-off between a recognition dependency and the trust's four centuries of precedent is argued qualitatively and not computed. **A reviewer would be right to press on this: the paper's central recommendation rests on the comparison it does not quantify.**

The thin-doctrine risk is acknowledged and unpriced. A statutory foundation regime has no settled equivalent of the trust doctrine for unwinding a fiduciary's mistaken exercise of a power, and that is the doctrine a family is most likely to need. **No probability or cost is assigned.**

The reflexive loop is not identified. Manski [1993] rules out recovering an endogenous effect from observation of a linear-in-means reflexive system, and a family governance instrument is such a system. The remedy adopted is stipulation with a defence [Brock and Durlauf, 2001]; **the price is that every claim about what a governance document achieves is a modelling position.**

The instrument is performative and its own parameters are exposed to that. MacKenzie and Millo [2003], Hacking [1995] and Lucas [1976]. **The framework asserts that systems change when modelled and exempts its own parameters from that reflexivity; that inconsistency is not resolved here.**

The corpus's own honesty apparatus bounds what this paper may claim. Curzi [2026] is the sharpest instance: the commons case there **passes the ceiling test with a margin of** $+0.058436$ **nats and is nonetheless practically unreachable, because no lever closes it — and a member's declining to participate improves the family reading by** 2.680588 **octaves.** That is the same

shape as this paper's consent result, arrived at in a different domain: **a formally available improvement whose only route runs through a person who may decline, and whose declining improves the measured quantity.** The corpus prints that arithmetic rather than avoiding it, and this paper is bound by the same standard.

Model comparison was not performed. No alternative apex model was fitted, no out-of-sample test was run, no information criterion was computed and no matched null was constructed, because nothing here is fitted. **The crossover model has zero free parameters and two stipulated ones, which means it cannot be wrong in the way a fitted model can be wrong — and cannot be right in that way either.**

No differential prediction is offered. Durlauf [2005] asks for the prediction only this model makes, and neither the crossover nor the two-exposure decomposition makes one: any fixed-versus-proportional cost comparison produces a minimum efficient scale, and any separation of entitlement from taxation produces two numbers. **The contribution is the identification of the recognition dependency and the discipline of computing both exposures, not a uniquely discriminating forecast.**

13.1 Rigour checks not met, listed by name

1. **No as-of date, because no real holdings were measured.** A concentration figure computed from real holdings requires its source and its date; **a concentration measure without an as-of date is not a measurement.** Every share vector here is constructed. The same applies with greater force to the legal propositions: **a statement of a reserve fraction or a tax rate without a date is not a statement of law**, and Appendix B accordingly requires the date of every confirmation.
2. **No uncertainty is reported anywhere, because nothing is estimated.**
3. **One symbol carries two dimensions.** Q is a currency stock in Sections 8.2 to 8.6 and an index-unit stock in the surplus law. Flagged in Section 12.1; the source does the same.
4. Q**,** C **and** h **in the surplus law have no operational unit**, and Curzi [2026t]'s demotion standard is the reason no statement here mixes them.
5. **The conjoint conditions were not tested** [Michell, 1997, Luce and Tukey, 1964]. The objection has limited bite on currency and full bite on index units.
6. **A count of non-exchangeable nominal items is used to trigger a threshold** (Section 12.6). It supports the threshold and is used for nothing else.
7. **Model comparison was not performed** and no information criterion was computed.
8. **No distributional or scaling claim is made, so the Clauset–Shalizi–Newman procedure was not run** [Clauset et al., 2009]. **Any restatement reaching for a heavy-tail claim about estate sizes or fee structures must run it first**, and must read Stumpf and Porter [2012] on why passing would still not identify a mechanism.
9. **The boundary-object and Goodhart borrowings are category 2, not category 1** (Section 12.5).
10. **The reserve is not a conserved quantity** and the fixed shares are (Section 12.4). The paper exploits that asymmetry and does not claim a conservation law for the framework's Q [Mirowski, 1989].
11. **Aggregation is unaddressed** [Kirman, 1992].
12. **Expected value and time average are not the same thing** [Peters, 2019]. A single succession event does not arrive in expectation; it arrives once, and the multiplicative consequence of a forced disposal for the remaining decades is not modelled here.
13. **The reflection problem is stated and not solved** [Manski, 1993].
14. **The paper's central recommendation rests on an unquantified comparison** — the recognition dependency against the depth of precedent. Section 12.
15. **The identity in H5 is reported as a hypothesis met, and an identity cannot confirm a hypothesis.** It is retained in the hypothesis set because stating it as a prediction and then showing it is an identity is what makes the instrument's blind spot visible; **a reviewer who objects that it should never have been a hypothesis is right, and the objection is**

recorded here.

14. Conclusion

This paper establishes the following results in the form the rest of the corpus cites them.

R1. The minimum efficient scale for a fixed-fee apex is $Q^{*} = C_{\text{fixed}}/b$. At $C_{\text{fixed}} =$ USD 120,000 a year, a foundation is cheaper than an *ad valorem* trusteeship only above **USD 34.3 million at 35 basis points, USD 80.0 million at 15, and USD 120.0 million at 10.**

R2. At the constructed USD 240 million estate the foundation costs **5.0 basis points against an *ad valorem* 35, a difference of USD 720,000 a year — which this paper explicitly refuses as the reason to choose the form**, because a fixed-fee trustee mandate eliminates the crossover and the saving is a rounding error against the exposure of R3.

R3. The forced-heirship exposure splits into a situs-capped levy, $\min(60.00\text{m}, 6.70\text{m}) =$ **USD 6.70 million**, or **11.2 per cent** of the notional reserved share, and an inheritance-tax exposure of approximately $0.45 \times 60.00\text{m} =$ **USD 27.00 million** that the structure does not touch — a ratio of **4.03**.

R4. Liquidity, not entitlement, binds. A **USD 60.00 million** claim against **USD 26.16 million** of liquid assets leaves a shortfall of **USD 33.84 million**, **14.1 per cent of the estate**, larger than any single liquid holding and closed in cash only by disposing of an operating branch.

R5. The fixed-share counterfactual on USD 240 million is determinate rather than discretionary: **widow** $1/8 =$ **30.00 million; residue 210.00 million divided** $2:2:1$**, so each son takes 84.00 million and the daughter 42.00 million**, summing to 240.00 million exactly.

R6. The apex does not move the concentration reading. Over the conventional structure the measurement remains **naive HHI 0.2934 (3.41 effective jurisdictions) and bloc HHI 1.0000 (**$N = 1.00$**) whatever vehicle sits at the top**, because the effective count is a function of the share vector and the correlation matrix and of nothing else.

Counter-result. The foundation answers ownership recording and succession continuity. **It adds nothing to the effective bloc count, gives no protection against arbitrary state action, and its forced-heirship result rests on a consent the heir may refuse — for which there is no lawful substitute.** Two further limits belong with it: the thin-doctrine risk is real, unpriced, and does not go away with good drafting; and the designed case's zero levy is a fact about the asset map rather than a legal victory, surviving exactly as long as the map does.

Two results are established as method rather than as measurement. **A recognition dependency is a form of concentration the published instrument cannot see**, because the effective count is invariant under the substitution of the apex form — which makes H5's confirmation a statement about the instrument's blind spot rather than about the world. And **the single most useful sentence for a practitioner is that a foundation apex moves the fight from "who owns this" to "was the endowment good", and the second fight is won or lost on the date of the endowment and the founder's conduct afterwards, not on the drafting.**

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every case in this paper — the principal, the family map, the asset allocation, the cost model, the fee bases, the exchange rate, the tax rate, the liquidity position and the policy benefit — is **CONSTRUCTED** and is labelled as such at each point of use. The numerical results are exact computations on those constructed inputs, carried at full precision and rounded only for printing.

All figures were produced by the scripts named in their captions (`figs/src/CCP-053_fig01.py` through `figs/src/CCP-053_fig04.py`), which import the shared figure style module and read

their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

Legal and jurisdictional descriptions in this paper are taken from the source manuscript and not from primary legal sources. The bibliography contains general treaty and comparative doctrinal material only; it contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction. Appendix B lists the claims affected. **Any reader relying on those descriptions must verify them against the official sources, and must record the date of that verification, because a statement of a reserve fraction or a tax rate without a date is not a statement of law.**

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. Every principal, family, heir, entity and transaction described is invented.

Preregistration

The five hypotheses of Section 2, including H2's three crossover values and H3's ratio, were stated with their falsifiers before the computations of Section 8 were run. **H5 was met as an identity rather than as a measurement, and Section 13.1 item 15 records that this is a weaker outcome than a confirmed hypothesis normally is.** No hypothesis in this paper was preregistered in any external registry.

Legal, tax and investment advice

This paper is analytical and is not legal, tax, or investment advice. Succession law is the most jurisdiction-specific subject treated anywhere in this corpus and the one where a general statement is most likely to be wrong in a particular case. The paper describes entity types, holding forms, reserve regimes, fixed-share systems and structural arrangements for the purpose of measuring cost, exposure and concentration. **No recommendation is made in respect of any particular person, asset, structure or jurisdiction.** Reserve fractions, claw-back periods, levy conditions, renunciation formalities, tax rates and abatements, entity-classification rules and foundation statutes all change and vary by jurisdiction, by family configuration and by asset; every such statement here is carried from the source manuscript, is flagged in Appendix B, and requires confirmation with qualified counsel in each jurisdiction of residence, nationality, beneficiary location and asset situs, and with counsel qualified in the applicable personal-status law where fixed shares apply, before any reliance.

Appendix A. The build specification

The decisions in the order they must be taken, with the consequence of each.

```
  1  FIX THE JOB BEFORE THE FORM. One paragraph, before counsel is
     instructed, stating what the apex is for. If it contains
     "operate", "trade", "invoice" or "employ", the design is wrong.
     If it contains "hide", it fails the disclosure test and should be
     abandoned.
  2  PURPOSE, BENEFICIARIES, OR BOTH. A pure beneficiary form gives
     each heir something that begins to look like an interest, and an
     interest is what a divorce court, a creditor and a
     reserved-portion claimant can each attach or value. A pure
     purpose form removes that and makes banking materially harder.
     The combined form states the purpose as the governing object and
     keeps the class discretionary.
  3  THE ENDOWMENT. Real, documented and clean. The deed identifies
     the property precisely, recites solvency at the date of transfer,
     and is dated. Do not endow assets already the subject of, or
     foreseeably about to be the subject of, a claim: that is the
     conduct claw-back statutes exist to reverse, and it converts a
     lawful structure into evidence.
  4  DRAFT THE CHARTER FOR THE REGISTER. Assume a journalist, a
     competitor, a foreign tax authority and a future heir all read
     it. Two clauses are load-bearing: post-mortem amendment must
     require more than one person's decision, and the winding-up
     clause must name a real destination.
  5  DRAFT THE BY-LAWS FOR THE COURTROOM. Beneficiary class including
     adopted children, stepchildren and children born outside
     marriage -- silence here produces litigation in every generation.
     Distribution standards; reserved powers with a successor
     mechanism covering incapacity as well as death; guardian consent
     matters and succession; council composition, quorum and deadlock;
     reserved matters; information rights; conflicts; dispute
     resolution.
  6  COMPOSE THE COUNCIL. Substance over convenience: at least one
     independent professional resident councillor engaged on
     commercial terms who reads the papers and can say no. Continuity
     over loyalty: stagger the ages. Location over documentation:
     minutes recording a meeting everyone knows happened by telephone
     add a credibility problem to a residence problem.
  7  GRADE THE RESERVED POWERS. Section 5.3.
  8  THE ONSHORE INTERFACE. Ownership particulars will be recorded;
     exchange classification must be determined and applied; the tax
     position must be taken and filed; onboard with the full pack and
     expect to answer "who benefits?" in writing, with an answer that
     is true.
  9  HANG THE STRUCTURE UNDERNEATH. Hold, do not cross-hold. One layer
     of separation per job.
 10  WRITE THE SUCCESSION MEMORANDUM. What the structure is, who the
     counsel are in each jurisdiction, where the documents live, what
     the reserved powers were and what happens to them, and who to
     call on the day. It has no legal effect and it is the difference
     between an orderly transition and eighteen months of paralysis.
     It is NOT a letter of wishes.
```

The succession overlay, added to the above where a protected share is in play: build the succession map (people, assets, and the law each forum would apply); compute E_{reserve} and E_{tax} separately and in currency before any instrument is discussed; select from the four-lever set; build one claw-back calendar per heir jurisdiction; pre-fund the liquidity gap; and produce three documents to the standard that the excluded heir may read each of them — the endowment file, the family disclosure, and the succession memorandum. **The family disclosure is uncomfortable and it is the highest-value document in the file.**

Appendix B. Claims requiring verification against primary sources

The following rest on legal or regulatory detail that changes, and **none is cited to the bibliography, which contains no primary sources for any individual jurisdiction.** They are reproduced from the source manuscript and require confirmation, **with the date of confirmation recorded**, before any reliance.

1. The apex jurisdiction's foundations statute: the scope of the constitutional carve-out permitting an enclave private law and the federal statutes applying notwithstanding it; the statutory trigger requiring a guardian; minimum councillors, corporate councillors and residency requirements; the availability and conditions of the registered-agent route and the agent's record-keeping and disclosure obligations; the enumerated reservable powers, whether the list is exhaustive, and whether the statute saves the foundation from invalidity by reason of reservation; any minimum initial asset requirement.
2. The firewall and governing-law provisions: article numbers and wording; extension to claims arising from personal relationships; treatment of foreign judgments as distinct from foreign law; interaction with the companion trust statute. **This is the single most commonly overstated feature of every foundation regime and the wording is the whole of it.**
3. Continuation-in and transfer-out provisions: registrar consent, solvency or creditor conditions, which entity types may continue in, and whether a company may convert.
4. The apex court: its constituting instruments, bench, opt-in jurisdiction, the protocol for enforcement into the surrounding onshore courts, the current constituting instrument and composition of any tribunal resolving jurisdictional conflicts, and the current position on whether a foreign judgment recognised in the enclave may be enforced onward. **That question has moved; do not build on an assumption about it.**
5. The reserve fractions of each relevant civil-law system, the reduction-in-value rule and when reduction in kind remains available, and the prescription periods for the reduction action.
6. The decisions of the relevant supreme court on whether the absence of a reserve in a chosen foreign law offends international public policy, their docket numbers and reasoning, and any subsequent decisions modifying the position.
7. The compensatory levy: its enacting instrument, exact conditions, mechanics, commencement, and current status including any constitutional or supranational challenge. **This provision is under active contest and its status must be checked as at the date of any advice.**
8. The compulsory-portion provisions of the second civil-law system, the operation of the ten-year sliding scale, and — critically — the effect of a reserved usufruct on when the period begins, together with the formalities of a contractual waiver. **This is where planning in that system most often goes wrong.**
9. The provision on life-assurance proceeds and the case law on manifestly excessive premiums.
10. The inheritance-tax residence condition for heirs and its article; the current abatement and rate scale; and **whether a discretionary foreign foundation making no distribution**

produces a taxable transmission at all, which determines which of two very different numbers applies.

11. The fixed shares on the actual family configuration; the school applied by the deciding court and the differences it produces; the position on bequests exceeding the cap and on bequests to heirs; the requirements for a completed lifetime gift; and the formalities of post-death settlement among heirs. **No general statement about fixed-share inheritance should be relied on for a specific family.**
12. The current status table of the recognition convention [Hague Conference on Private International Law, 1985] for every state concerned, including ratifications, accessions and territorial extensions; and the article numbers of its mandatory-rules, overriding-provisions and public-policy reservations and of the provision permitting non-recognition where the significant elements are more closely connected with non-trust states.
13. Contracting-state lists and entry-into-force dates for the choice-of-court [Hague Conference on Private International Law, 2005] and judgments [Hague Conference on Private International Law, 2019] instruments in each relevant state; the article and sub-paragraph numbers of their succession, family-law and matrimonial-property exclusions; and the apex jurisdiction's own regime for recognition and enforcement of foreign judgments, including its procedural legislation and any regional or bilateral instruments in force with each heir jurisdiction.
14. Whether a foreign *foundation* specifically is recognised on the same basis as a foreign company in each relevant state, and whether any residual real-seat doctrine or public-policy control applies to a foundation administered from within that state.
15. The classification of a foreign foundation for federal tax purposes in each relevant system — foreign trust, foreign corporation, or otherwise — on the specific charter proposed; the attribution provisions of each relevant civil-law system and the conditions of any regional exception; and the current reporting obligations and penalties. **These are genuinely unsettled for foundations and must not be assumed either way.**
16. Statutory valuation scales for usufruct and bare ownership, and any anti-abuse presumption applying to a family separation of usufruct from bare ownership.
17. The applicable matrimonial property regime for the couple, the law governing it, and whether it may be changed and with what effect. **It changes the denominator and is identified before anything else.**
18. The succession regulation's article numbers for the general rule, choice of law, scope, universal application and public-policy exception [European Union, 2012], and the recital addressing evasion of the law. **The substance is not in doubt; the numbering must be pinned before it appears in an opinion.**
19. All cost figures in Section 8.2 are constructed illustrations. They are not fee quotations and no real fee schedule was used to produce them.

Appendix C. The apex jurisdiction assessed

The jurisdiction is de-identified as "Gulf centre A", following the convention of Curzi [2026], because every item below asserts a jurisdiction-specific characteristic the bibliography cannot support. No total is computed: the items are heterogeneous and no common scale exists on which they could be added (Section 12.6).

#	item	assessment
1	legal family	codified common-law-derived; own civil and commercial statutes under a financial free-zone carve-out; not English law by direct application, but drafted in the common-law tradition, in English, with an English-language court
2	bloc alignment	**bridge.** Maintains working relations across several perimeters simultaneously. Not outside regulatory pressure; outside the Western listing mechanism and outside any single bloc's coordination
3	court quality	high: international bench, English language, published reasoned judgments, opt-in jurisdiction for unconnected parties. **Thin foundations-specific case law — the regime dates only from 2018**
4	foreign judgment recognition	**mixed, and the weakest item.** Enforcement into the surrounding onshore courts runs through a protocol; non-binding memoranda of guidance exist with several foreign courts and are not treaties; the onward-enforcement question has moved. Do not build on an assumption here
5	substance requirements	low for a holding foundation: registered agent, filed charter, statutory registers, council meetings. Substance for the *operating* entities is a different question
6	exchange-reporting status	participating. A foundation is a financial institution or a passive non-financial entity; either way, account holders or controlling persons are reported. **There is no non-reporting configuration**

#	item	assessment
7	beneficial-ownership regime	own regime through the registrar: filed and maintained, not publicly searchable, available to the registrar and competent authorities. Charter public; by-laws private but discoverable
8	rail access	incumbent messaging throughout, with genuine dual-rail optionality at the banking layer. **The apex itself should hold no payment function: a rail is not a holding**
9	political stability	high by regional standards, with the caveat of a system in which policy can change quickly without a legislative cycle. **That cuts both ways**
10	degradation trajectory	**improving, from a base that briefly worsened.** Increased monitoring March 2022 to February 2024; a federal corporate tax from 2023 which *increased* transparency and simultaneously increased the jurisdiction's legitimacy as a treaty and substance location; an ownership register built out from 2020; deepening case law. **The asset being bought is bloc independence with institutional quality, and both are rising.** Downside risks: federal encroachment on the carve-out, and single-country political concentration
11	job	**HOLDING**, unambiguously. Not operating: no substance, no staff, no trade at the apex. Not bridging

#	item	assessment
12	verdict and what would change it	**use it as the apex** on the balance of legal personality, purpose-or-beneficiary flexibility, permitted founder control, a publishing common-law court and bridge-bloc alignment. **What would change the verdict:** federal legislation reaching into the carve-out; a line of decisions making the apex court's orders unreliable against onshore assets; amendment narrowing the firewall or the transfer-out power; renewed listing; or a de-risking cycle making the form unbankable in the branch jurisdictions. Monitor items 4 and 10 annually and keep the migration path live

All numerical results in this paper were computed from the stated inputs at full precision and rounded only for printing.

Chapter 54

Institutional Quality Without Independence: London, the Crown Dependencies, and the European Convergence

Abstract. Three of the most-used holding locations for Western capital are measured on the concentration instrument, and all three fail it for the same reason: institutional quality is not independence. Spreading a holding across Jersey (40.0 per cent), Guernsey (35.0) and the Isle of Man (25.0) gives a naive HHI of 0.345000 and 2.90 effective jurisdictions, but all three collapse to one bloc, giving a bloc-HHI of 1.000000, an effective count of 1.00, and a surplus term of exactly zero. Even the most generous counterfactual, crediting the Crown Dependencies as a bloc of their own, reaches only 1.423998 effective blocs and h = −0.376286 against the designed structure's 2.941176 and −0.122417. The European four — Luxembourg 35.0, Ireland 25.0, the Netherlands 25.0, Switzerland 15.0 — read 3.703704 naive and 1.00 after collapse; crediting Switzerland separately reaches 1.342282 and h = −0.396970. A printed arithmetic error in one variant is left visible and corrected in the text, to 0.764992, N = 1.307203 and h = −0.406239, as a demonstration of what a rounded print does to the reading. The paper also reports a move that reduces resilience while appearing to improve it: shifting the residual UK 16.4 per cent to Bermuda lowers Western exposure to 13.6 per cent yet cuts the effective bloc count from 2.9412 to 2.3442, because it concentrates the bridge row. The UK share falls from 40.9 to 16.4 per cent with no loss of capability, because English governing law and arbitration are portable and separable from situs.

Keywords: City of London; Crown Dependencies; English governing law; portable law; convergence; coordination bloc; concentration measurement

Series: D. Structure and Resilience **Paper:** CCP-054

1. Introduction

Three of the most-used holding locations for Western capital are measured here on the concentration instrument, and all three fail it for the same reason. **Institutional quality is not independence.** They are different axes, and almost every structuring decision in this field reads a high score on the first as evidence for the second.

The distinction has to be held with both hands, because a paper that undersells these jurisdictions is as useless as one that oversells them. The City of London is an extraordinary piece of legal and commercial infrastructure and nothing in this corpus replaces it. The trust statutes of the Crown Dependencies are among the best-maintained in any common-law system; one of the Bailiwicks invented the protected cell company and then engineered the incorporated cell company when the first turned out to have a recognition limit. Luxembourg runs the second-largest fund domicile in the world and runs it well; Ireland is the only common-law member state of the European Union with the euro and the leading aircraft-leasing cluster; the Netherlands has a first-rate judiciary and, in the administration foundation, one of the most useful control-separation instruments in any legal system; Switzerland has custody depth and a currency that is not somebody else's policy instrument. **None of that is the question.** The question is what happens to the effective count of independent coordination perimeters, and on that question the answer is not close.

Four measurements follow, and two of them cut against the framework's own presentation.

First, the Crown Dependency error. Spreading a holding across three Islands at 40.0 / 35.0 / 25.0 per cent gives a naive index of 0.345000 and 2.898551 effective jurisdictions. All three collapse to one bloc: a bloc index of 1.000000, an effective count of exactly 1.00, and a surplus term of exactly zero. **Three jurisdictions, three regulators, three sets of fees, three sets of counsel, three annual filings, and no change in the quantity that enters the surplus law.**

Second, the counterfactual run in the incumbent structure's favour. Grant the most generous assumption available — that the Crown Dependencies constitute a fully independent perimeter — and the conventional structure reaches a bloc index of 0.702248, $N = 1.423998$ and $h = -0.376286$, against the designed architecture's 2.941176 and -0.122417. **The verdict does not turn on winning the bloc-assignment argument; the generous assignment loses too, by 0.253869 of surplus balance.**

Third, the European convergence. Luxembourg 35.0, Ireland 25.0, the Netherlands 25.0 and Switzerland 15.0 read 0.270000 naive and 3.703704 effective jurisdictions, collapsing to exactly 1.00; crediting Switzerland as its own perimeter reaches 0.745000, $N = 1.342282$ and $h = -0.396970$. **A printed arithmetic error in one variant of that computation is left visible in Section 7.4 and corrected in place — to 0.764992, $N = 1.307203$, $\ln N = 0.267890$ and $h = -0.406239$ — because what a wrong index does downstream is more instructive than a clean-looking table.**

Fourth, a move that reduces resilience while appearing to improve it. Shifting the residual UK 16.4 per cent to the Gulf bridge lowers Western-perimeter exposure from 30.0 to 13.6 per cent and *cuts* the effective bloc count from 2.9412 to 2.3442, because it concentrates the bridge row. **Reducing exposure to the perimeter you fear, by piling weight onto one bridge, makes the structure less resilient rather than more. Concentration does not care which direction it points.**

Against those four, one finding runs the other way and is the paper's most useful output. The UK share falls from 40.9 to 16.4 per cent **with no capability lost**, because what is bought under the name of a financial centre is a bundle of separable services, and the most valuable of them — governing law — is portable. **A contract can be governed by English law and heard outside the Western perimeter. The law travels; the courthouse does not.**

Two constraints on what follows are stated at the outset. **No jurisdiction-specific primary law is cited.** The bibliography carries general treaty and standard-setting material — the awards convention [United Nations, 1958], the choice-of-court and judgments instruments [Hague Conference on Private International Law, 2005, 2019], the exchange and standard-setting instruments [Organisation for Economic Co-operation and Development, 2017, Financial Action Task Force, 2012, Organisation for Economic Co-operation and Development, 2015] — and comparative doctrine [Torremans et al., 2022, Briggs, 2013, Fawcett and Carruthers, 2008]. **It contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction.** Every enactment, decided case, listing event, rate, date and jurisdictional characteristic below is described from the source manuscript, left uncited, and

itemised in Appendix B. And **no distributional or scaling claim is made anywhere**: the exposure vectors have three, four and six entries and are not samples from a distribution [Clauset et al., 2009].

2. Hypotheses

Five hypotheses with their falsifiers, in the form Platt [1964] requires. Meehl [1967] is the standing warning that a directional prediction here is nearly costless; all five make point predictions with constants attached. Two are reported in Section 7.7 in ways that do not flatter the argument.

H1 — the subdivision hypothesis. Subdividing an exposure across jurisdictions inside one coordination perimeter raises the naive count and leaves the effective bloc count exactly unchanged. *Predicted:* a Crown Dependency allocation of 40.0 / 35.0 / 25.0 gives naive HHI $= 0.345000$ and 2.898551 effective jurisdictions against bloc HHI $= 1.000000$ and $N = 1.000000$ exactly, so $\eta \ln N = 0$ at every value of η. *Falsified by:* any subdivision inside a perimeter that raises N; or by evidence that the three Islands act on independent decisions rather than on one.

H2 — the generous-counterfactual hypothesis. The conclusion survives granting the bloc-assignment argument to the other side. *Predicted:* crediting the Crown Dependencies as their own perimeter (WEST 81.8 / CD 18.2) gives bloc HHI $= 0.702248$, $N = 1.423998$ and $h = -0.376286$, which remains 0.253869 below the designed architecture's -0.122417. *Falsified by:* a defensible bloc assignment under which a single-perimeter structure reaches the designed architecture's position.

H3 — the convergence hypothesis. The European four have lost differentiation rather than privacy, and their bloc position is unchanged by either. *Predicted:* 35.0 / 25.0 / 25.0 / 15.0 gives naive 0.270000 and 3.703704, collapsing to 1.000000 and $N = 1.00$; crediting Switzerland separately gives 0.745000, $N = 1.342282$ and $h = -0.396970$. *Falsified by:* a demonstrated instance of one of the four declining to give effect to a coordinated measure adopted by the others.

H4 — the reallocation-trap hypothesis. Lowering exposure to a named perimeter can lower the effective bloc count. *Predicted:* moving the residual UK 16.4 per cent onto the Gulf bridge gives blocs of 56.4 / 30.0 / 30.0 → 56.4 / 30.0 / 13.6, a bloc index of 0.426592 and $N = 2.344160$, down from 2.941176, while Western exposure falls from 30.0 to 13.6 per cent. *Falsified by:* a reallocation that reduces the largest bloc's share and reduces N, which is impossible; or by any claim that Western-perimeter exposure alone is a sufficient statistic for resilience, which this prediction refutes by construction.

H5 — the portability hypothesis. A jurisdiction's most valuable export can be consumed without consuming its perimeter exposure. *Predicted:* the UK share falls from 40.9 to 16.4 per cent while all four services the constructed case actually consumes are still consumed, and the component decomposition accounts for the fall exactly: 22.0 + 9.4 + 9.5 = 40.9 before, and 10.0 + 6.4 + 0.0 = 16.4 after. *Falsified by:* a service in the retained set that cannot be obtained without a UK entity or a UK forum.

Not hypotheses, and labelled as such. Four statements are not falsifiable as written.

- *The bloc assignment of each named jurisdiction* is a **modelling position**, stipulated in advance from the evidence of Section 6.3 and defended there. Section 3.4 states why it cannot be an identified empirical result.
- *The parameter set* $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ is a **modelling choice** carried from Curzi [2026b] and fixed before the measurements, following the preregistration discipline of Nosek et al. [2018].
- *The logarithmic form of the input term* is a **modelling choice** carried from Curzi [2026], whose functional characterisation — angle-only dependence, vanishing at parallel, an affine effective count, normalisation — fixes the form by axiom rather than by measurement. It is not tested here.

- *The exposure shares* are **constructed** throughout and are labelled as such at each point of use.

3. Method

3.1 The gauge, and what an element contributes to it

The surplus law is $h = \eta \ln N - \delta Q - C$, and Curzi [2026b] states it in full: a headroom gauge with three regimes, a maximum increment $r_{\max} = h_t$, and two separate constraints — admissibility, $C_t \leq (1-\delta)Q_t + \eta S_t$, and sustainability, $C_t \leq \eta \ln N - \delta Q_t$. **The distinction between the two constraints is what this paper's jurisdictions fail.** A Crown Dependency or European element is almost always admissible — the cost can be met out of stock and flow — and almost never sustainable, because it raises C and leaves $\ln N$ untouched.

That is the whole accounting, and it is worth writing as one line:

```
an element that adds to C and adds nothing to ln(N)
moves h strictly DOWN, at every value of eta.
```

The recurrence itself is a first-order linear difference equation and its behaviour is standard [Elaydi, 2005]. Curzi [2026] supplies the honest converse that governs any claim about a structure converging on a target: **gain exactly one gives an isometry that orbits without arriving, and gain at or above one gives an amplifier whose corrections, sized to delayed feedback, arrive as new errors.** A restructuring programme run on annual measurements with a two-year execution lag is exactly such a system, and Section 8 lists the drift it produces as a failure mode rather than as a side effect.

3.2 Which effective count

Every effective count in this paper is HHI^{-1}, the Hill number of order two, and it is named as such at each use. Hill [1973] defines the family ${}^{q}D = (\sum_i p_i^q)^{1/(1-q)}$, whose members return the number of equally-abundant types that would produce the observed index; at $q = 2$ that is exactly $(\sum_i p_i^2)^{-1}$. Jost [2006] is the statement that entropies are not diversities and their numbers-equivalents are; Adelman [1969] is the numbers-equivalent reading of the concentration index; Simpson [1949] and Laakso and Taagepera [1979] are the independent reinventions in ecology and political science. The index is due to Hirschman [1945] in root form and Herfindahl [1950] in sum-of-squares form. The exponential of Shannon entropy [Shannon, 1948] is the same family at $q = 1$, and the reciprocal index is the exponential of the Rényi entropy of order two [Rényi, 1961].

The corpus's affine count $N_{\text{eff}} = 1 + (N-1)u$ is a different quantity, is not a Hill number of any order, and carries no external support. It is not used here. Curzi [2026] is cited for the surplus functional and its axiomatic characterisation, not for that count as a numbers-equivalent, and the distinction must survive any citation of this paper.

The general form the collapse relies on is $\mathrm{EFF} = 1/(w'Rw)$ for exposure weights and a correlation matrix — the similarity-sensitive diversity of order two [Leinster and Cobbold, 2012], whose additive relative is the quadratic entropy of Rao [1982]. Under independence it is the naive index; under a block matrix of ones it is the bloc index. Markowitz [1952] supplies the underlying algebra and the correction the whole series turns on: **diversification benefit depends on covariance and not on count, stated with the mathematics attached in 1952.** Hurlbert [1971] is the standing discipline: the word "diversity" carries no content without its formula attached, which is why the formula is attached here every time.

One limitation of the instrument bites unusually hard in this paper and is stated before the results rather than after them. Golosov [2010] and Dunleavy and Boucek [2003] document that effective-number measures behave poorly in the *dominated-unit regime*, where one unit

holds a very large share. **Four of this paper's five headline vectors collapse to a single unit at 100 per cent, which is the limiting case of that regime.** The measurement remains exact — $1/1 = 1$ — but the instrument's *discriminating power* is at its minimum there, and comparisons among heavily dominated structures must not be made on N alone. Section 11 records this as the paper's most serious methodological exposure.

3.3 What correlated shares do, and what the literature says about it

The bloc collapse is the assertion that within-perimeter correlation is near one. Two bodies of measured evidence bear on it and they point the same way. Longin and Solnik [2001] and Ang and Chen [2002] find extreme correlation in equity markets that is asymmetric and concentrated in downturns: **the effective count falls exactly when it is needed**, which is a stronger statement than anything the static index makes. The ecological name for the same variable is synchrony [Loreau and de Mazancourt, 2008], and the cleanest demonstration that effective count buys stability is Schindler et al. [2010]'s sockeye record, where population diversity produces roughly half the variability and about a tenth of the fishery closures.

The contrary literature is read alongside and is not decorative. Doak et al. [1998] show that diversity–stability correlations can arise from statistical averaging alone, with no mechanism; May [1972] showed that increasing diversity and connectance *destabilises* randomly assembled systems. **This paper's claim is narrower than the general diversity–stability claim and is not exposed to those results in the same way**: it asserts that the reciprocal of a variance counts independent shocks, which is arithmetic, and that the jurisdictions in question are attached to one decision process, which is H1 and H3.

The coupled-network result [Buldyrev et al., 2010] is used as a **shared structural form and not as a derivation** — the jurisdictions have not been shown to satisfy its percolation premises, and only the property that coupling destroys assumed independence is borrowed. The robust-yet-fragile result [Albert et al., 2000, Callaway et al., 2000] is used the same way, with Doyle et al. [2005]'s caveat that real engineered systems get their robustness from design rather than from a degree distribution. And the general form of the counting error is proved exactly: **one party can present as many** [Douceur, 2002], so an unaudited count of nominally distinct units is worthless — which is what a three-Island subdivision is.

3.4 The reflection problem, where the assignment is made

A bloc assignment is a forecast about correlated behaviour, and the framework's reflexive loop $T(x) = W(x, M(x))$ makes it a forecast that conditions the behaviour it forecasts. Manski [1993] proves that a linear-in-means reflexive specification is **not identified from observational data**: the endogenous effect cannot be separated from contextual effects or from correlated unobservables, however much data is collected. Shalizi and Thomas [2011] reinforces the point where the data are relational — influence and selection are generically confounded. **No amount of observation of jurisdictional behaviour could therefore confirm a bloc assignment derived reflexively, and this paper states that where the assignment is made rather than in a footnote.**

The remedy adopted is stipulation with a defence, the fourth of the four known routes, which Brock and Durlauf [2001] shows can be made rigorous. The assignments of Section 6.3 are fixed in advance against a stated evidence ladder, recorded with the strongest available item, and stress-tested by the generous counterfactuals of Sections 7.2 and 7.3. **The price is that every bloc assignment in this paper is a modelling position and not an identified empirical result**, and H2 and H3 exist precisely so that the conclusion does not depend on winning the assignment.

The instrument's own performativity is not exempted. MacKenzie and Millo [2003] documents a model that came to fit its market after traders adopted it; Hacking [1995] gives the general looping form; Merton [1948] supplies the original statement that a definition of a situation can make itself true. **If enough structures are built on the assumption that these**

jurisdictions share a perimeter, capital and correspondent relationships reorganise around that assumption and the assumed correlation becomes the observed one. Lucas [1976] states the severe consequence for parameters estimated before publication. **The framework asserts that systems change when modelled and then exempts its own η, δ and correlation assumptions from that reflexivity. That is an inconsistency and it is not resolved here.**

3.5 Evidence standard, and the de-identification convention

Sections 4 to 6 are jurisdictional exposition. **Every enactment, decided case, listing event, survey figure, rate, date and jurisdictional characteristic in them is described from the source manuscript and is not cited to the bibliography.** The general instruments named in Section 1 are cited where they bear; nothing else is. Appendix B lists every claim requiring confirmation against a primary source, **with the date of that confirmation recorded**, because a statement of a jurisdiction's regime without a date is not a statement of its regime.

Where a table **ranks named jurisdictions against one another on a legal capability the bibliography cannot support**, the jurisdictions are de-identified — "Island A", "Island B", "Island C" — following the convention of Curzi [2026], which de-identifies an apex jurisdiction in its own twelve-item assessment for the same reason. **Where a jurisdiction is named in prose, or appears in an exposure table, or carries a bloc assignment with its evidence, the name is used**, because the argument is about those specific perimeter attachments and a functional description would hide rather than protect. That is the same line Curzi [2026] draws.

4. What is actually being bought

Six distinct things travel under the name of a major common-law financial centre. They have different portability, different failure modes and different jobs, and separating them is the first analytic move.

```
1  the governing law            the substantive rules of the contract
2  the forum                    a commercial court, its process, its judges
3  the arbitral seat            a seat, not a court system
4  the insurance market         underwriting capacity for hard risk classes
5  the listing venue            a market and secondary liquidity
6  the currency plumbing        banks, clearing, market depth
```

Item 1 is portable and items 2 to 6 are not. A contract between two parties in different blocs can be governed by English law and heard anywhere that will apply it. **The law travels; the courthouse does not**, and most of the strategic content of this paper follows from that distinction. Two enclave jurisdictions outside the United Kingdom apply English common law directly or through their own courts with common-law judges, which means **the law can be bought without buying the perimeter, and that is the single most useful structural fact in this corpus's treatment of the West.**

4.1 Why the law is chosen, and what that claim is not

The reasons are unglamorous and are the reasons that matter. *Predictability through precedent density*: two centuries of reported commercial decisions mean most disputes about most standard terms have already been decided, and the answer can be found before the dispute rather than after it. *A strong doctrine of freedom of contract*: the default posture is that sophisticated parties are bound by what they wrote, with no general doctrine of good faith in commercial performance of the kind found in many civil-law systems. *Commercially legible remedies*, and a settled law of interpretation. *A separable arbitration regime*, so that the law can be chosen without the forum.

Two disciplines apply. **"Most-chosen governing law" is a claim about relative frequency in international commercial contracting, resting on survey and institutional data rather than a census**, and it must not be inflated. And choosing a governing law does not import a jurisdiction: a clause can choose one system's law with another system's court or seat, and in a cross-bloc structure it usually should.

4.2 Judgment recognition, and why an award travels further than a judgment

This is the capability that materially deteriorated and it is routinely glossed over. English judgments circulated within the European Union and the European Free Trade Association under a recast regulation and a parallel convention; both ceased to apply on the United Kingdom's departure, and an application to accede to the convention in its own right was not supported. What replaced it is thinner: an instrument covering **exclusive** choice-of-court agreements only [Hague Conference on Private International Law, 2005], with asymmetric and non-exclusive clauses — standard in finance documentation — genuinely contested; a broader judgments instrument [Hague Conference on Private International Law, 2019] with a limited and slowly growing membership; common-law recognition on a presence-or-submission basis; and two statutory reciprocal-enforcement schemes for listed countries.

Against that sits the awards convention [United Nations, 1958], with near-universal adherence, narrow and enumerated refusal grounds [van den Berg, 1981], and a procedural framework the surveyed regimes largely share [United Nations Commission on International Trade Law, 1985]. Born [2021] is the standing treatment.

For a cross-border structure a London-seated arbitral award is a more portable instrument than an English court judgment, and it became more so in 2021. Structures papered on pre-2021 reasoning should be re-examined on current reasoning.

The counter-consideration, fairly stated: arbitration has no real summary judgment under most institutional rules, no appeal on the merits, limited joinder, and cannot bind non-signatories. **Where disputes are likely to be documentary and one-to-one, a commercial court is superior on speed and cost. Where they are cross-border and the enforcement target sits outside the Western perimeter, arbitration wins on the only criterion that matters at the end — whether the paper can be turned into money where the debtor keeps assets.**

4.3 Disclosure as a feature, not a burden

Three registration regimes matter in the constructed case: a public register of persons with significant control; a register of overseas entities owning domestic land; and an identity verification regime for directors, controllers and filers, rolled out in stages. They are usually described as burdens. **Under the disclosure standard this corpus applies they are closer to the opposite.**

An operating company inside a properly designed architecture passes the disclosure test **because of** the public control register, not despite it. The register produces a filed, dated, public statement of who controls the entity; if the controller is a registered foundation whose own regime records its council and founder, the chain is documented end to end and **the documentation is the asset**. A bank onboarding the entity gets an answer in an afternoon; a screen resolves cleanly; an enquiry closes.

Filed transparency is what makes a structure bankable, and bankability is the binding constraint on cross-bloc architecture — not secrecy, which is unavailable, and not rate advantage, which is now largely harmonised at the top end. This is the open-design principle of Kerckhoffs [1883] — a system's effectiveness must not depend on the secrecy of its design — restated as *open design* among the protection principles by Saltzer and Schroeder [1975], and applied to legal rather than cryptographic structure. Brandeis [1914] is the public-interest form of the same observation from the other direction.

The honest cost is real and is stated. A public control register is public, so adverse parties, journalists and commercial rivals can read it. For a principal with genuine personal-security exposure that is a real problem, and the suppression regimes that exist are narrow and evidentially demanding. **They are not a general privacy option and should never be planned as one** — reaching for a personal-security provision as a privacy plan is a concealment move inside a disclosure regime, which is the definition of a structure that fails the standard.

4.4 Enforcement reach, honestly stated

Two instruments belong in a threat model rather than a footnote. **A civil order requiring a respondent to explain the source of property** exists, and its case law cuts both ways: one appellate decision upheld such an order and one first-instance decision discharged a set of them with an adverse costs outcome against the authority. **The instrument is real, used sparingly, and it has lost in court** — which is the proper description of a legal power in a jurisdiction with a functioning judiciary, and is exactly what distinguishes lawful process from arbitrary action.

An autonomous sanctions regime exists, administered by a designating authority rather than a downstream adopter. Designations are made independently of two other major regimes, and the *observed* behaviour since 2022 has been rapid convergence across all three. **That observation is the evidence for the bloc assignment, and it is an observation about outputs rather than an allegation about intent.**

4.5 A jurisdiction that writes the rules cannot hedge against them

The perimeter test of Curzi [2026] — restated there as the question whether a decision taken elsewhere inside a perimeter reaches a jurisdiction without a second independent decision — has an answer for a founding member of the standard-setting body and an independent designating authority that is not merely inside the perimeter but at its origin. **A jurisdiction that helps write the perimeter's rules cannot also be the hedge against them, and no amount of institutional quality changes that arithmetic.**

A second finding belongs here because it is the one most likely to be argued with. **Institutional quality and rule stability are independent properties.** The constructed record includes the abolition of a settlement for internationally mobile individuals that had existed in recognisable form since 1914 — announced in March 2024, revised at a Budget in October 2024, effective April 2025, across a change of government. Whatever one thinks of the policy, **the engineering lesson is unambiguous: political stability is not the absence of upheaval; it is the predictability of the rules a structure depends on**, and a jurisdiction can have world-class courts, a professional civil service, deep markets and an unbroken constitutional order while replacing a century-old settlement inside thirteen months. **Score the two separately, and never let a high score on the first suppress an honest score on the second.**

5. Capability, inventoried

This section records what the jurisdictions genuinely do, because the argument is a concentration argument and not a criticism of capability, and because an inventory that only lists deficiencies is not an inventory.

5.1 The Islands

Fiduciary depth. One Island's trust statute is among the most carefully maintained in any common-law system, with three features that matter: a firewall addressing the extent to which questions concerning the trust are determined by local law to the exclusion of foreign

law; a reserved-powers provision permitting a settlor to reserve specified powers without invalidating the trust; and a statutory jurisdiction for the court to set aside transfers into trust, and exercises of trustee powers, on grounds of mistake or inadequate deliberation. **That last provision was legislated within a year of an appellate decision elsewhere narrowing the equivalent common-law doctrine, and the responsiveness cuts *for* the jurisdiction rather than against it. Capability is not the deficiency here.**

Segregation engineering. A second Island enacted protected cell company legislation first, anywhere, and the concept was adopted worldwide: a single legal person with statutorily segregated cells, assets attributable to one cell unavailable to creditors of another. It then addressed the form's principal weakness — that a cell is not a separate legal person, which foreign courts and counterparties sometimes decline to respect — by creating an incorporated cell company in which each cell **is** a separate legal person. **This is asset partitioning performing exactly the function Hansmann and Kraakman [2000] identifies as the irreducible contribution of entity law**, engineered to the limit of what a single domicile can achieve. Its remaining exposure is recognition: statutory segregation in the domicile is not automatically recognised in a forum with no equivalent concept, and that recognition question decides whether the ring-fence holds when tested.

Operational registries. A third Island built an aircraft and ship registry competing on operational merit, inside a customs and indirect-tax arrangement with its neighbouring state that no comparable centre has. **This is the most interesting capability in the section, because it is the one that automatic exchange, substance rules and ownership-register reform have not touched — there was never anything hidden in it to remove.**

The absence that governs. All three use a currency they do not issue, have no central bank and no lender of last resort. Two bank failures in 2008 demonstrated the consequence: one Island had no depositor compensation scheme in place at the time and introduced one in response. **A jurisdiction with no central bank cannot support an institution, and the protection available to a depositor is whatever the scheme says it is.** For operating cash that is a sizing question; for reserve it is a disqualification.

5.2 The European four

Each has a genuine, differentiated capability, and the capabilities are not interchangeable: a regulated-fund passport with distribution across a large single market; an English-speaking common-law member state of that market with a leading asset-leasing cluster; a civil-law state with an English-language commercial court and a control-separation instrument of unusual elegance; and a non-member with custody depth and an independent currency.

The control-separation instrument deserves its own paragraph because it is the one genuinely portable idea in the group. A foundation holds legal title to shares and exercises the votes through its board; depositary receipts are issued to the economic beneficiaries and carry the economics but no votes; administration conditions govern transfer and exchange. **The whole operation is a matter of public record: the foundation is registered, the shareholding is registered, the separation is documented.** It achieves what a great many concealment structures are built to achieve — control in defined hands, economic benefit spread wider — **entirely in the open**, which is why it belongs in this corpus and why the register-privacy movements of the last five years made no difference to it in either direction.

5.3 What was lost, and it was not privacy

The received account is that these jurisdictions were secretive and are becoming less so. **That account is wrong, and it is wrong in a way that produces bad structures**: a principal who believes the problem is disclosure will spend money buying privacy that is not the constraint while ignoring the constraint that is.

The fact that breaks it is on the public record. In November 2022 a supranational court held invalid the provision of an anti-money-laundering directive insofar as it required mem-

ber states to ensure that beneficial-ownership information is accessible in all cases to any member of the general public, on the ground that the interference with charter rights was not limited to what is strictly necessary. **One register suspended public access the same day; a second within hours; others followed. Those jurisdictions became measurably more private in November 2022, by decision of the relevant legal order's own highest court**, and any analysis premised on their "losing their secrecy" is describing a trend that reversed direction.

What they lost is **differentiation**, which is a different thing entirely. A jurisdiction commands a premium only for what it can do that its neighbours cannot, and across fifteen years almost every source of differentiation was removed — not by secrecy reform but by coordinated harmonisation of the base, the rate, the anti-abuse rules and the reporting.

source of differentiation	what removed it	when
headline-rate advantage	a fifteen per cent minimum effective rate for groups above a turnover threshold	FY2024–25
base and financing advantage	interest limitation, general anti-abuse rule, controlled-foreign-company and exit-tax rules; then hybrid and reverse-hybrid mismatch rules	2019; 2020–22
ruling and structure secrecy	mandatory disclosure of reportable arrangements; platform reporting; crypto-asset reporting	2020; 2023; 2026
information advantage	automatic exchange of account information [Organisation for Economic Co-operation and Development, 2017]	2017–18
substance-free holding	a supranational abuse principle independent of any treaty	2019
special domestic regimes	domestic reform enacted by referendum	2020

And the cost went up. Minimum-rate compliance is not a filing but a data-collection programme across every entity in a group; hallmark assessment and reporting is continuous; substance means directors, offices, staff, minutes and transfer-pricing documentation in each jurisdiction. **The result is the worst combination available in structuring: a jurisdiction that has lost its distinctive advantage while retaining — and increasing — its distinctive price.**

The third element never changed, and it is this paper's actual subject: **bloc concentration is intact.** Three of the four are inside the Western perimeter by construction. The fourth is the interesting case, and the evidence is unambiguous on outputs: from February 2022 it adopted its neighbours' sanctions packages through ordinances made under its own embargo legislation, and has aligned with successive packages since. **Neutrality is a doctrine about military alliance. It did not survive contact with the financial coordination mechanism, and the date on which that became demonstrable is on the public record.**

6. The bloc assignment, and its evidence

6.1 The evidence ladder

The perimeter test asks whether a decision taken in one place reliably produces the corresponding effect in another **without a further independent decision**. Evidence is ranked, and the assignments below are made against that ranking rather than against atmosphere:

weight	evidence class
strongest	a single designation decision producing the corresponding effect in every member without a further decision
strongest	a shared final appellate bench
strongest	a shared currency with no independent monetary authority
strongest	separate legislatures enacting near-identical text to a common external deadline
weaker	participation in an information-exchange framework [Organisation for Economic Co-operation and Development, 2017]
weaker	implementation of common supervisory recommendations [Financial Action Task Force, 2012]
weakest	listing status, diplomatic alignment, press characterisation [de Goede, 2011]

6.2 The Islands: the constitutional position, stated accurately

The Islands are **not** part of their neighbouring state. They have their own legislatures, courts, civil services and fiscal systems, and none was ever a member of the regional union. They are also **not** independent states: their neighbour is responsible for their defence and international representation, their primary legislation requires assent through a conciliar mechanism, and their final court of appeal is a committee composed of the neighbouring state's supreme court justices.

Two features do the work.

The appellate bench is the neighbour's bench. The committee is a distinct court applying the law of the jurisdiction the appeal comes from, and it decides against the neighbouring government's interests regularly; **no allegation of bias is made or implied.** The observation is about correlation. **The court of last resort for one Island's trust law is staffed by the same judges as the court of last resort for the neighbour's law, so the two bodies of law cannot drift far apart and cannot fail independently.** For a claimant that is a feature; for a structure trying to buy independence it is the whole problem.

There is no monetary authority. All three use the neighbour's currency and issue their own notes at par, which is a currency arrangement and not a monetary policy.

6.3 The coordination evidence

This decides the assignment and should be read as evidence rather than argument.

Corporate tax. All three Islands moved to zero-rate general corporate taxation with a concessionary rate for defined financial-services activity in the same period. All three then had elements of those regimes — the attribution and deemed-distribution mechanisms designed to preserve resident taxation — challenged as harmful by the same external body, and all

three repealed or replaced them. **Three separate fiscal systems adopted the same design, received the same challenge from the same body, and made the same repeal.**

Economic substance — the decisive instance. Following an external assessment and the commitments that followed, all three enacted economic substance requirements for relevant activities with effect for accounting periods commencing on or after the same date. The relevant-activity lists, core income-generating activity concepts, adequacy tests for people, premises and expenditure, and reporting mechanics are substantially aligned, **because all three were drawn against the same external template** [Organisation for Economic Co-operation and Development, 2015]. **Three parliaments, one deadline, near-identical text — that is what a coordination perimeter looks like from the inside.**

Listing. All three were placed in the co-operative-with-commitments category of a coordinated list pending delivery of that legislation and all three were subsequently removed once the assessing body was satisfied [de Goede, 2011].

Automatic exchange. All three were early adopters with first exchanges in 2017 [Organisation for Economic Co-operation and Development, 2017]. **There is no non-reporting configuration in any of the three and there has not been one for a decade.**

Beneficial ownership. All three maintain central registers accessible to competent authorities and not publicly searchable. In 2019 all three gave a public commitment to move toward publicly accessible registers on a stated timetable aligned with an external standard; that trajectory was interrupted by the November 2022 judgment of Section 5.3, and the Islands moved instead toward legitimate-interest access. **The access model has changed twice; a structure whose rationale depends on the current model is built on the least stable variable in the jurisdiction.**

Sanctions — the strongest available class. Each Island has its own legislation and competent authority, and in practice the neighbour's designations take effect in all three. The mechanism differs; the substantive result has been consistent alignment. **A single decision taken in one capital produces the corresponding effect in all three Islands without a further independent decision, which is the perimeter test satisfied in its strongest form.**

Supervision. All three have been evaluated with generally strong results reflecting genuinely capable supervision. **The finding is not that the Islands are weakly supervised. It is the opposite, and that is the point: they are excellently supervised to a standard set outside them. Excellent compliance with an external standard is the definition of being inside the perimeter that sets it.**

One further piece of evidence is the hardest in the paper and is stated in cash terms rather than in theory. One Island sits inside its neighbour's customs and indirect-tax territory under a revenue-sharing agreement that has been renegotiated, at one point substantially to the Island's disadvantage. **A jurisdiction inside another's indirect-tax territory, sharing its currency, with its judges as the final court, and whose fiscal position has been reset by a decision taken outside it, is not an independent perimeter under any definition that would survive a real event.**

6.4 Rails

There is no non-Western clearing access, no depth in an alternative reserve currency, and no correspondent network of consequence outside the perimeter in any of the seven jurisdictions surveyed. Local banks are mostly subsidiaries or branches of Western groups, several of which have reduced their presence over the past decade under general de-risking and cost programmes. **Correspondent concentration in all of them is higher than the entity map suggests**, and the rail axis must be measured separately before any of them is given a settlement role.

The general point belongs to the currency axis and is worth stating once: **currency exposure is perimeter exposure independent of entity location.** A treasury running predominantly two Western reserve currencies has more Western-perimeter exposure than its entity map shows, whatever passport its directors hold.

7. Results

CONSTRUCTED CASE throughout. All facts and figures are invented for illustration. No real person, family, entity or transaction is described. The carried case is a dual national, Gulf tax resident, with three heirs in three jurisdictions, one under a forced-heirship regime, holding manufacturing in a non-aligned state, a Gulf trading company, intellectual property, Western commercial real estate and a listed portfolio.

The starting measurement, carried from Curzi [2026], whose R6 establishes that the apex vehicle does not move it: **naive HHI 0.2934 and 3.41 effective jurisdictions against bloc HHI 1.0000 and $N = 1.00$ whatever vehicle sits at the top.** At $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ that gives $\ln N = 0$, $\eta \ln N = 0$ and $h = -0.500000$.

7.1 The Crown Dependency error, and H1

The conventional structure allocates 18.2 per cent to one Island and assigns it to WEST. Redistributing that 18.2 per cent across all three Islands is the move an adviser makes when asked to diversify:

```
Island A 40.0%  ·  Island B 35.0%  ·  Island C 25.0%

naive HHI = 0.160000 + 0.122500 + 0.062500 = 0.345000
effective jurisdictions = 1 / 0.345000 = 2.898551

bloc collapse: all three -> WEST = 100.0%
bloc-HHI = 1.000000      EFFECTIVE BLOCS N = 1.000000
ln(N) = 0                eta ln(N) = 0        h = -0.500000
```

Three jurisdictions, three regulators, three sets of fees, three sets of counsel, three annual filings. N unchanged at 1.00. The surplus term unchanged at zero, at every value of η. H1's prediction is met exactly.

The four variants an adviser will propose, run together:

variant	naive HHI	naive effective	bloc-HHI	N	$\ln N$
conventional as found (one Island at 18.2)	0.293430	3.408	1.000000	1.000	0.000
that Island's holding moved to a second Island	0.293430	3.408	1.000000	1.000	0.000
the 18.2 split three ways across the Islands	0.271347	3.685	1.000000	1.000	0.000
all three Islands, 40 / 35 / 25	0.345000	2.899	1.000000	1.000	0.000
the Islands credited as an INDEPENDENT bloc	0.293430	3.408	0.702248	1.424	0.353

Four of the five rows have $\ln N = 0$ and therefore a surplus term of exactly zero. Row two is identical to row one to twelve decimal places. **A substance-driven proposal to move a**

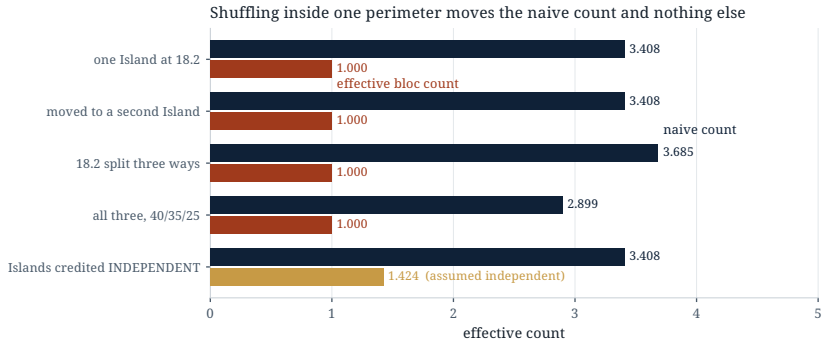


Figure 54.1: **Subdividing inside one perimeter moves the naive effective count between 2.90 and 3.69 and leaves the effective bloc count at exactly 1.00 in every row.** Measured from the share vectors tabulated above; the fifth row is the generous counterfactual of Section 7.2. Script `figs/src/CCP-054_fig01.py`.

holding between two Islands was not wrong on its own terms; it answered a question about compliance while being presented as an answer to a question about resilience. The measurement's function was not to reject the move but to reveal that the move and the objective were unrelated.

Note the direction of row four. **Splitting into three roughly equal Island holdings makes the *naive* count worse — 2.899 against 3.408 — because it removes the largest holder's dominance from a four-jurisdiction vector and replaces it with a three-jurisdiction one. Even the flattering number does not flatter.**

7.2 The generous counterfactual, and H2

A paper that only runs the arithmetic it likes is not an instrument. Suppose the bloc assignment is wrong and the Islands really are an independent perimeter — constitutionally distinct, separately governed, capable of declining to follow:

```
WEST 81.8% (three Western jurisdictions)  ->  0.818^2 = 0.669124
ISLAND bloc 18.2%                         ->  0.182^2 = 0.033124

bloc-HHI = 0.702248
EFFECTIVE BLOCS N = 1 / 0.702248 = 1.423998
ln(N) = 0.353469     eta ln(N) at eta = 0.35 = 0.123714
h = 0.123714 - 0.500000 = -0.376286
```

Against the designed architecture's $N = 2.941176$ and $h = -0.122417$ — figures Curzi [2026] carries in its own R1, where adding two further Gulf jurisdictions raises the naive effective count from 5.29 to 7.13 while the effective bloc count falls from 2.94 to 2.84.

Even granting the most generous assumption available, the conventional structure reaches only 1.42 effective blocs and remains at $h = -0.376286$, which is 0.253869 behind the designed one. H2's prediction is met, and the verdict does not turn on winning the bloc-assignment argument — only its magnitude does.

7.3 The European convergence, and H3

```
Luxembourg 35.0%  ·  Ireland 25.0%  ·  Netherlands 25.0%  ·  Switzerland 15.0%

naive HHI = 0.122500 + 0.062500 + 0.062500 + 0.022500 = 0.270000
effective jurisdictions = 1 / 0.270000 = 3.703704
```

```
bloc collapse, all WEST:  bloc-HHI = 1.000000   N = 1.000000
ln(N) = 0                 eta ln(N) = 0         h = -0.500000

generous counterfactual, the non-member as its OWN perimeter:
WEST 85.0% -> 0.722500     CH 15.0% -> 0.022500
bloc-HHI = 0.745000     N = 1.342282
ln(N) = 0.294371        eta ln(N) = 0.103030    h = -0.396970
```

Even treating the non-member as a fully independent coordination perimeter — which the sanctions evidence of Section 5.3 does not support — the four-jurisdiction European structure reaches 1.34 effective blocs against 2.94, and $h = -0.396970$ against -0.122417. H3's prediction is met.

The six variants:

variant	naive HHI	naive effective	bloc-HHI	N	ln N
conventional as found (one European at 13.6)	0.293430	3.408	1.000000	1.000	0.000
moved to the second European jurisdiction	0.293430	3.408	1.000000	1.000	0.000
moved to the third	0.293430	3.408	1.000000	1.000	0.000
moved to the non-member	0.293430	3.408	1.000000	1.000	0.000
all four, 35 / 25 / 25 / 15	0.270000	3.704	1.000000	1.000	0.000
all four, non-member as a SEPARATE bloc	0.270000	3.704	0.745000	1.342	0.294

Five of the six rows produce a surplus term of exactly zero. Moving a holding vehicle from one European centre to another is the single most common diversification step taken by European families, **and it changes the effective bloc count by exactly nothing.**

7.4 A printed arithmetic error, left visible and corrected in place

The source computation for one variant of Section 7.3 — the non-member credited as its own perimeter, at the conventional structure's own weights of WEST 86.4 per cent and non-member 13.6 per cent — was printed as follows:

```
PRINTED     bloc-HHI = 0.749296   N = 1.334   ln(N) = 0.288
```

Recomputed from the stated shares:

```
0.864^2 + 0.136^2 = 0.746496 + 0.018496 = 0.764992

CORRECTED   bloc-HHI = 0.764992   N = 1.307203   ln(N) = 0.267890
            eta ln(N) = 0.093761  h = -0.406239
```

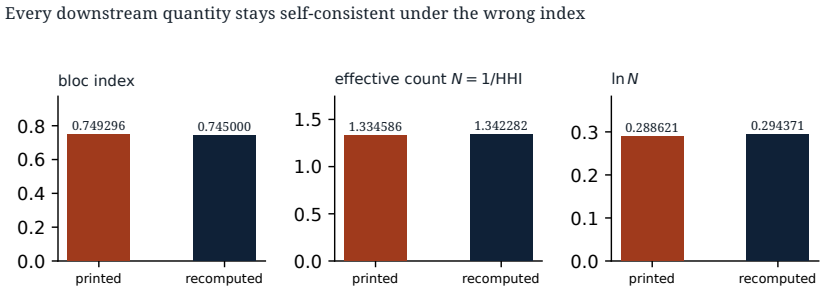


Figure 54.2: **A single wrong index propagates consistently: the printed bloc index of 0.749296 yields an internally coherent effective count and logarithm, and only recomputation from the shares detects it.** Recomputed from $0.864^2 + 0.136^2 = 0.764992$; the printed row is reproduced exactly as published. Script `figs/src/CCP-054_fig02.py`.

The printed index is wrong. It is left visible here rather than silently amended, because what it does downstream is the finding.

Three observations, each of which is the reason the error is worth a section.

First, the error is in the index alone, and everything below it is internally consistent with the wrong value. $1/0.749296 = 1.334586$ and $\ln(1.334586) = 0.288621$, so the printed N and $\ln N$ follow correctly from the printed index. **A single wrong input that propagates consistently produces a table in which every relationship checks out, which is exactly why such an error survives review.** Checking a table for internal consistency does not detect it; only recomputation from the inputs does.

Second, the printed value cannot be reproduced from the stated shares by any operation this paper describes. No rounding of 0.864 and 0.136 to any printed precision yields 0.749296. **This paper does not speculate about its origin**; it records that the value is unreproducible from the published inputs, which is the strongest thing that can honestly be said about it.

Third, the conclusion is unaffected, and that is not a defence. All four proposals in the variant table leave the structure below $N = 1.35$ against the designed architecture's 2.94, and three leave it at exactly 1.00. The error moves h from -0.398983 to -0.406239, a difference of 0.007256. **That the conclusion survives is a fact about this particular error's size and not a property of the method, and a reader who takes "the conclusion is unchanged" as a general reassurance has drawn the wrong lesson.** The corpus's own preregistered fit is the reference point: Curzi [2026i] reports a confirmatory test that **failed its own second falsification rule** — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin, despite +40.3 per cent over persistence, +97.7 per cent over direct input, and $p = 0.0005$ against the shuffle. **A series that reports that failure cannot report an arithmetic error only when it changes an answer.**

The general discipline is not to round carefully but to **carry full precision through every intermediate step and round only at the print** — and to publish the inputs alongside the outputs so that recomputation is possible at all. Nosek et al. [2018] is the framework; Simmons et al. [2011] and Ioannidis [2005] are the reasons an undetected error in a published computation is the normal case rather than the exceptional one.

7.5 The reallocation that reduces resilience, and H4

Two questions follow from the designed architecture's residual Western element.

Why keep it at all? Because the Western real estate exists, the English-law counterparty relationships exist, and market access is worth something. Removing it entirely would raise N and destroy capability the architecture actually uses. The measured trade:

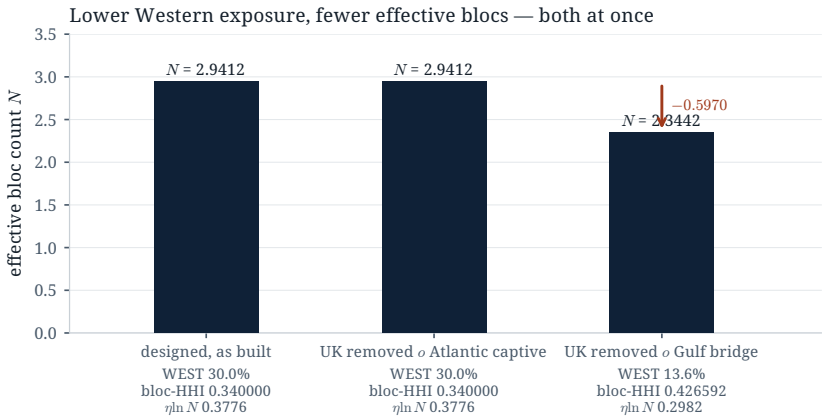


Figure 54.3: **Lowering Western-perimeter exposure from 30.0 to 13.6 per cent cuts the effective bloc count from 2.94 to 2.34, because the exposure has to go somewhere and it went onto one bridge.** Measured from the three share vectors tabulated above; h computed at $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$. Script `figs/src/CCP-054_fig03.py`.

variant	WEST %	bloc-HHI	N	$\ln N$	$\eta \ln N$
designed, as built	30.0	0.340000	2.9412	1.0788	0.3776
UK removed, 16.4 to the Atlantic captive jurisdiction	30.0	0.340000	2.9412	1.0788	0.3776
UK removed, 16.4 to the Gulf bridge	13.6	0.426592	2.3442	0.8519	0.2982

```
third row, worked:  BRIDGE      29.1 + 16.4 + 10.9 = 56.4   0.318096
                    NONALIGNED  18.2 + 11.8        = 30.0   0.090000
                    WEST (captive jurisdiction only)  13.6  0.018496
                    bloc-HHI = 0.426592     N = 1/0.426592 = 2.344160
```

The second row is the point. Moving the residual Western holding to another jurisdiction inside the same perimeter changes nothing at all in bloc terms. **A move that looks like leaving a financial centre and is in fact a move inside the same perimeter buys exactly zero.**

The third row is worse than the base case, and it is H4. Reallocating the weight onto the bridge takes Western exposure down to 13.6 per cent and *reduces* N from 2.9412 to 2.3442, because it concentrates the bridge bloc from 40.0 to 56.4 per cent. **Reducing exposure to the perimeter you fear, by piling weight onto one bridge, makes the structure less resilient rather than more. Concentration does not care which direction it points, and Western-perimeter exposure is therefore not a sufficient statistic for resilience — a claim H4 refutes by construction.**

What job is the Western element doing in the effective-bloc arithmetic? None. It contributes to C and not to N, and is justified on capability grounds alone. An element that does not move the concentration measurement is not thereby worthless; it is answering a different question. **Such elements must be justified explicitly and sized deliberately, because the default failure is to let them grow until they dominate the structure they were supposed to diversify away from.** Curzi [2026] and Curzi [2026] report the same shape on different clusters — a domestic four-state vector at 40.0 / 25.0 / 20.0 / 15.0 giving naive 0.285000 and 3.508772

effective jurisdictions but bloc 1.000000, $N = 1.00$ and a surplus term of exactly zero; and six Atlantic-cluster entities at 40.0 / 22.0 / 14.0 / 12.0 / 7.0 / 5.0 giving naive 0.249800 and 4.0032 effective jurisdictions but bloc 1.000000, $N = 1.00$ and $\ln N = 0$. **Three independent clusters, three collapses to exactly one.**

7.6 The share falls and the capability does not, and H5

Applying the six-service taxonomy of Section 4 to what the constructed principal actually needs:

service	needed?	why
the governing law	**yes**	four supply agreements and two financing documents are on it and the counterparties will not move
the forum	**no**	the counterparties are in three other states and none keeps realisable assets in the jurisdiction
the arbitral seat	partly	retained for two legacy contracts; new contracts move seat
the insurance market	**yes**	political-risk and contract-frustration cover on the manufacturing plant; no other market quotes the exposure at a workable price
the listing venue	**yes**	the conversion route for the building
the currency plumbing	**yes**	a five-currency book; execution only, not custody, not treasury residence

Four of six retained, two changed. "We need this financial centre" turned out to mean four specific services rather than a jurisdiction, and only one of the four requires an entity there to exist at all.

The exposure decomposition, before and after:

component	before	after	what changed
commercial building, held through an offshore company	22.0	10.0	converted to a listed securities position; the *holder* need no longer be locally reachable to realise it
operating company: trading, staff, holding group IP licences	9.4	6.4	IP moved out; treasury function removed; arm's-length licence in
locally-custodied portion of the listed portfolio	9.5	0.0	moved to a bridge-bloc custodian

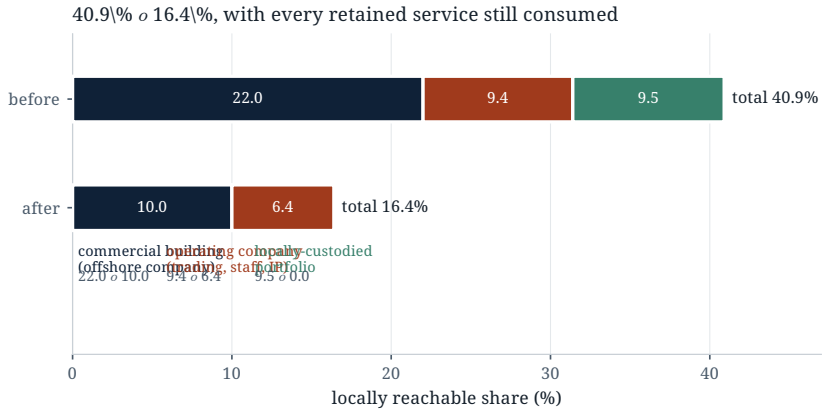


Figure 54.4: **The share falls from 40.9 to 16.4 per cent and every retained service is still consumed: the capability was never in the percentage.** Measured from the component decomposition above; the three components sum to the stated totals exactly. Script `figs/src/CCP-054_fig04.py`.

component	before	after	what changed
total	**40.9**	**16.4**	

The share went from 40.9 to 16.4 per cent and the group lost no capability, because the capability was never in the percentage. It was in four services, all four still being consumed. What fell was exposure. **H5's prediction is met, and the decomposition accounts for the fall exactly.**

Three observations on the rebuild, because each is a general rule.

The most valuable single change cost eleven contract amendments and no entity change. The dispute clauses were redrafted from the enforcement jurisdiction backwards: name the state where the counterparty keeps realisable assets, then choose the instrument that reaches it. Two contracts moved to arbitration under the awards convention [United Nations, 1958]; one kept the governing law and moved the forum to an enclave applying that law directly, **so predictability was unchanged and the forum moved into a bridge bloc — and the counterparty accepted it without negotiation, because their concern was the law and not the courthouse**; one kept the commercial court, correctly, because the counterparty's assets were domestic and the dispute would be documentary; and one could not move at all, because mandatory local law governed the joint venture. **The rule is not "avoid the centre"; it is "choose deliberately", and deliberate choice sends some matters to the centre.**

The rebuild moved toward disclosure at every single step and the structure got stronger. The offshore holding company was collapsed; the ownership chain was documented to natural persons on a public register; abridged filings became full audited accounts; two unverified directors became three verified ones with a resident executive. **That is the disclosure thesis in its most concrete available form.**

A political-risk underwriter prices opacity as risk. An applicant who can produce a complete ownership chain, filed accounts and a clean screen gets a quote; one who cannot gets a declination or a loaded rate. **This is the disclosure standard paying a dividend measurable in basis points**, and it is the mechanism by which Curzi [2026]'s result about performative commitment operates on an institution rather than a person: a commitment made in public, to an audience that will re-price later defection, is stronger than the same commitment made privately, **for the same three reasons — consistency pressure, self-perception from overheard testimony, and the re-pricing of later defection.**

7.7 Hypotheses tested

H1 — supported, and met as an identity. Bloc HHI $= 1.000000$ exactly in four of five variants. **That it is an identity rather than a measurement is important and is recorded in Section 11: the result follows from the construction of the collapse given the assignment, and its content is entirely in the assignment, which is a modelling position (Section 3.4).**

H2 — supported. $h = -0.376286$ on the generous counterfactual, 0.253869 behind the designed architecture.

H3 — supported. 0.270000 naive, collapsing to 1.000000; 0.745000 and $h = -0.396970$ on the generous counterfactual.

H4 — supported, and it is the paper's most useful negative result. N falls from 2.9412 to 2.3442 while Western exposure falls from 30.0 to 13.6 per cent.

H5 — supported. 40.9 $\rightarrow$ 16.4 with the decomposition accounting for the fall exactly and four of six services retained.

Two results cut against the framework's own presentation and are reported here in the same language a success would have received. A printed computation in the source is shown to be wrong and is corrected in place rather than silently amended (Section 7.4). And a move that lowers Western-perimeter exposure to 13.6 per cent nevertheless lowers the effective bloc count from 2.94 to 2.34 (Section 7.5) — **which means the exposure percentage the whole series reports most prominently is not, on its own, a measure of resilience.**

8. Failure modes

1. Treating the best jurisdiction as the apex because it is the best. Every quality signal points one way, so the holding entity goes there and the structure ends at $N = 1.00$. **Quality and independence are different axes, and this paper exists because people read the first as evidence of the second.**

2. Treating "offshore" as a synonym for "outside the perimeter". The word is a legacy of exchange control and describes a tax and regulatory position, not a political one. **Moving between two jurisdictions inside one perimeter is a change of tax treatment and regulatory regime; it is not a change of perimeter, and the two are constantly confused because the word describes the first while implying the second.**

3. Diversifying across three jurisdictions inside one perimeter. Section 7.1. **The variant appears in real files because it is easy to execute and it looks like activity.**

4. Reading institutional quality as independence. The Islands score extremely well on court quality, supervision and professional standards, and those scores are earned. **They are also, in part, evidence *for* the bloc assignment rather than against it, because the standards being met are set outside them.**

5. Confusing a governing law with a jurisdiction. The highest-value correction available. A clause choosing one system's law and another system's forum keeps the substantive predictability and moves the forum outside the perimeter. **It is available, cheap, and almost never drafted, because the two concepts arrive in the same sentence in the counterparty's first draft and nobody separates them.**

6. Planning around a judgment-recognition position that no longer exists. Structures papered before 2021 frequently assume regional circulation of judgments. That assumption is dead; the replacement covers exclusive choice-of-court agreements only [Hague Conference on Private International Law, 2005], leaving asymmetric clauses — standard in finance documentation — genuinely contested.

7. Assuming one register is the whole disclosure picture. A company holding local land through an overseas entity engages a second register; an overseas parent engages the chain-walking machinery of the first; a trust in the chain engages specific provisions. **The failure is rarely non-compliance; it is an incomplete map producing inconsistent filings across three**

registers, which is exactly the pattern that triggers an enquiry. Inconsistency between filings is a far bigger risk than disclosure itself.

8. Building a verification plan on a printed date. Identity-verification rollouts are staged and rescheduled. A group with fifteen directorships needs a programme with an owner and a live calendar, not a note in a file.

9. Underestimating the interim toolkit as a defendant. Freezing injunctions and worldwide disclosure orders do not ask who deserves them. A structure with an element in a jurisdiction has given its courts a jurisdictional hook and a claimant a service gateway. The response is not to avoid the jurisdiction but to keep that entity's assets proportionate to the capability it provides, so the worst available interim outcome is bounded.

10. Using an element as a treasury or clearing hub by drift. Nobody decides to do this. One entity has a working bank account and the others are painful to open, so payments start routing through it, and within two years the group's settlement concentration sits in the jurisdiction the architecture was designed to reduce exposure to. The control is a written payments policy with per-entity caps, reviewed against the bank statements rather than against the policy.

11. Reserved powers exercised as control. A reserved-powers provision lets a settlor reserve powers without invalidating the trust. It does not let a settlor run the trust as a personal account and keep the protection of the trust form. A reserved power exercised on every decision is not a reserved power; it is evidence.

12. Assuming a firewall stops a foreign court. A firewall is a rule about what the local law applies to local proceedings. It does not stop a foreign court asserting jurisdiction over a person before it, or making orders against that person on pain of contempt. A firewall protects the structure from foreign law; it does not protect a person from a foreign judge.

13. Ignoring substance because the entity "just holds". Pure equity holding companies are a distinct category under all three Island regimes, generally with reduced requirements — but "reduced" is not "none," and the classification of a mixed-activity entity changes the requirement. An entity that acquires an intra-group financing or IP function without anyone reclassifying it is the standard defect.

14. Forgetting there is no lender of last resort. Section 5.1. For operating cash that is a sizing question; for reserve it is a disqualification.

15. Diagnosing the problem as secrecy loss. Section 5.3. A principal who believes these jurisdictions are becoming less private will buy privacy that is not the constraint.

16. Building for a rate differential that a minimum-rate regime removed. Or crossing the in-scope threshold unnoticed, which changes the compliance economics of every entity in the chain including entities outside the region.

17. Letting an element grow because it is the easiest thing to grow. The deepest failure mode and the least visible. One entity can hire, bank, contract and be understood without explanation, so every marginal decision favours putting the next thing there. Two years of marginal decisions produce a structure whose entity map says 16.4 per cent and whose operating dependency says something much larger. Annual re-measurement with a written diff is the only control that catches it, and the diff is the finding. Curzi [2026h] states the general bound this failure mode sits inside: two levers that do not pay are reported there — cutting decay alone by raising trust makes the worked case *worse*, at -0.2114, while reciprocity alone gives $+0.7394$; and a deliberate 40 per cent sunset without any conversion work ends with less stored order, 5.507, than doing nothing at all, 7.494. Shedding is a stabiliser, not a cure. Removing an element is subject to the same discipline as adding one.

18. Letting an element outlive its justification. An Island holding built correctly for the conditions of its time becomes a defect not because anything went wrong there but because the structure around it changed and nobody reopened the question.

9. Discussion

What follows. Four things, each narrower than the general claim it resembles.

Subdivision inside a perimeter is exactly zero, not merely small. The result is arithmetic given the assignment, and Section 7.1 reports it as an identity rather than as a measurement. **The organisational form of the same finding predates the measurement by four decades:** Perrow [1984] established that interactive complexity plus tight coupling makes catastrophic failure a property of the system rather than of any component, and this paper's signature result — that adding a jurisdiction inside one perimeter can leave resilience unchanged or make it worse — is that result in a measurement frame. **It is cited as qualitative theory and framing, never as a formal result**, and the corpus should stop presenting the finding as new.

Adding coordination surface has a cost even when it adds nothing. Each additional Island or European entity adds counsel, filings, local substance and management attention. The general organisational result is that units inside a coordination boundary add communication cost at order $n(n-1)/2$ [Brooks, 1975] and add nothing to independence; and Conway [1968] supplies the reason a structure ends up mirroring the advisory organisation that built it rather than the threat it was built against. **A three-Island structure is a communication structure, drawn as a resilience structure.**

The direction of a reallocation does not determine its effect. H4. **Concentration is a property of the share vector and not of the label attached to the largest share**, which is why a plan expressed as "reduce exposure to X" can lower N while succeeding on its own terms. Curzi [2026] prices the closely related question of when a compartment is worth building: **with a cost advantage of 0.090 against a fenced benefit of 1.400 units a year, indifference sits at $1.400(1-r) = 0.090$, giving $r = 93.5714$ per cent — the minimum confidence that a foreign forum respects the cell wall.** That is the same shape of result as this paper's segregation caveat in Section 5.1: **a segregation capability is worth its cost only above a confidence threshold that can be computed, and the threshold is high.**

The most valuable export of a financial centre is the one that does not carry its perimeter. H5. This is the paper's positive result and it generalises: **wherever a jurisdiction's value is an intangible that another jurisdiction can adopt, the value can be consumed without the exposure.**

What does not follow. Four things.

It does not follow that these jurisdictions should not be used. Each has a capability the architecture uses. **The prescription is to size the element to the capability, book it honestly, and re-justify it annually.**

It does not follow that raising N is always right. Uzzi and Spiro [2005] find an inverted U in which performance rises and then falls with cohesion; Sagan [1993] documents that redundancy can *reduce* safety by adding coordination surface; March [1991] gives the mechanism by which fast convergence on a single mode destroys the variance that produced the advantage. **A monotone raise-N prescription cannot survive those results unqualified and none is made here.**

It does not follow that the effective count is a complete account of a structure. Curzi [2026] establishes that a **recognition dependency** — whether the form of ownership asserted over an exposure is intelligible to the courts and registries that must record it — is invisible to this instrument, because the effective count is a function of the share vector and the correlation matrix and of nothing else. **Its R6 is the sharpest statement of the limit: the apex does not move the concentration reading, naive 0.2934 / 3.41 and bloc 1.0000 / 1.00, whatever vehicle sits at the top.** This paper's Section 7.5 finds a second thing the reading does not capture: which bloc the reduced exposure moves *into*.

It does not follow that the instrument discriminates well here. Section 3.2. Golosov [2010] and Dunleavy and Boucek [2003] document that effective-number measures behave poorly in the dominated-unit regime, and **four of five headline vectors in this paper sit at the limiting case of it.** The measurements are exact and their discriminating power is minimal, and Section 11 records this as the most serious methodological exposure in the paper.

The status of the claim is measurement, not physics. Hedström and Ylikoski [2010] states the criterion the discipline named in the corpus title applies: an explanation must name a mechanism rather than exhibit a fitted curve. **The mechanism named here is a transmission structure — a shared appellate bench, a shared currency, a shared designation decision, a common external template — and nothing is fitted.** No order parameter is identified, no phase transition located, no critical exponent estimated; the vocabulary of Castellano et al. [2009] is deliberately not used. **What is offered instead is the structural analogy at the right strength:** Castellano et al. [2000] shows that the same interaction which raises local agreement freezes global diversity into non-interacting regions with a sharp transition between the monocultural and fragmented phases — **a perimeter, in that vocabulary, is a frozen region**, and Granovetter [1978] and Schelling [1971] are the standing warnings against reading an individual-level parameter straight into an aggregate, which is what a bloc assignment does. Gallegati et al. [2006] is the checklist this paper was run against, and Section 11 reports where it passes and where it does not.

Parallel results elsewhere in the corpus. The shape recurs on different objects. Curzi [2026c] reads the sharp Lipschitz constant $N-1$ as a **measurement-error budget**: at $N=12$ and $u=0.05\pm0.05$ the surplus lies anywhere in $[0, 0.7419]$, a width of 0.7419, against 0.1266 at $u=0.70$. **That is the same asymmetry this paper's Section 7.4 demonstrates on a printed index: the sensitivity of the output to an input error is not constant across the domain, and it is largest exactly where the corpus's headline cases sit.** Curzi [2026] reaches the formal boundary from the other side, with **Theorem 16.2:** $\mathrm{NET}(m) = 2\log_2(m) - 1.442695(\kappa/\eta)m(m-1)$ **octaves, so a dyad is viable in the best case if and only if** $\kappa/\eta \le \ln 2 = 0.693147$ — a logarithmic benefit against a quadratic coordination cost, which is the same trade this paper measures between $\ln N$ and a C that grows with the entity count. **The corpus has the general theorem; this paper supplies the instance in which the logarithmic term is exactly zero.**

Two further corpus results bound what may be claimed about the audit itself. Curzi [2026] reports an audit distribution over thirteen entries — **seven MAPS, two SPLIT, one RENDERING, three REJECTED, with five entries carrying a failure** — which is the standard this paper's jurisdictional assessments are held to: **an audit that rejects nothing is not an audit**, and Section 7 rejects four of five Island variants and five of six European ones. Curzi [2026] supplies the precision discipline in exact-rational form — **a minimum-gap schedule by ladder length: 1200 cents at $N=2$, 203.9100 at $N=3$–6, 34.9756 at $N=7$–16, 6.0008 at $N\ge17$, 1.0296 at $N\ge41$** — establishing that the corpus computes in exact rationals where it can. **Section 7.4's error is precisely what happens when it does not.**

10. Verification and rigour audit

10.1 Dimensional audit

symbol	quantity	dimension	unit	status
s_i	exposure share of jurisdiction i	$[1]$	dimensionless	defined; sums to 1 by P1
W_b	bloc share	$[1]$	dimensionless	derived, $W_b = \sum_{i\in b} s_i$
HHI	concentration index	$[1]$	dimensionless	derived from shares
ρ_{ij}	correlation between shocks	$[1]$	dimensionless, $[-1,1]$	**assumed, not measured**
N	effective bloc count	$[\#]$, **a count**	blocs	derived as $\mathrm{HHI}_{\text{bloc}}^{-1}$

symbol	quantity	dimension	unit	status
$\ln N$	input-term argument	$[1]$	nats	derived; the logarithm of a count is dimensionless
Q	protected capital	$[C]$	index units	defined; **no operational unit is supplied by the source**
δ	erosion rate	$[T^{-1}]$	per period	assumed
C	maintenance cost flow	$[C\,T^{-1}]$	index units per period	assumed
η	conversion efficiency	$[C\,T^{-1}]$ per nat	index units per period per nat	assumed
h	surplus balance	$[C\,T^{-1}]$	index units per period	derived

Does the surplus law balance? $h = \eta \ln N - \delta Q - C$ requires every term to carry $[C\,T^{-1}]$. δQ does; C does by definition; $\eta \ln N$ does **only if η carries $[C\,T^{-1}]$ per nat**, and **the source is silent on this** — it treats conversion efficiency as a bare number. The audit records the silence rather than filling it in. **The reading in the table is this paper's, not the source's.**

The important row is N. The index is a dimensionless sum of squared dimensionless shares, so its reciprocal carries the dimension of a **count**. **The instrument is not a score that happens to start at 1; it is a measured count of independent units, with a unit**, and Section 10.3 checks it against the definition of a count.

Where no operational unit exists, this audit says so. Q is protected capital in index units and no conversion to currency is supplied. **Every value of h in this paper is therefore on an arbitrary scale fixed by the stated parameters.** The differences between h values — 0.253869 between the generous counterfactual and the designed architecture, 0.007256 between the printed and corrected variants — are in the same arbitrary units and are meaningful only as ratios and orderings.

10.2 Significant figures and epistemic status

- 1.000000, 0.345000, 0.270000, 0.702248, 0.745000, 0.764992, 0.426592 are **exact arithmetic** on exactly-stated shares. Their reciprocals — 2.898551, 3.703704, 1.423998, 1.342282, 1.307203, 2.344160 — are exact to the printed precision and may carry as many digits as are useful.
- $\ln N$ and $\eta \ln N$ are **derived-from-defined** at an assumed η. **They are exact functions of an assumed input and therefore carry no uncertainty and no authority beyond the assumption.**
- h values carry the same status and are additionally in unnamed units (Section 10.1).
- 0.749296 is **reproduced exactly as printed in the source** and is not a computation of this paper. Section 7.4 records that it is unreproducible from the stated shares.
- The exposure shares themselves — 40.9, 27.3, 18.2, 13.6, 29.1, 16.4, 11.8, 10.9 — are **constructed and are exact as stated to one decimal place**. Curzi [2026] records in its own verification section that the naive index of the conventional vector depends on the rounding basis: 0.293388 on exact fractions against 0.293430 on three-significant-figure percentages. **This paper uses 0.293430 throughout, which is the three-significant-figure basis and the reproducible choice from published inputs.**
- **No quantity in this paper is a fitted or measured estimate, so no uncertainty interval is reported anywhere. That is not a strength: it means the paper has no empirical content to attach an interval to.**

10.3 Limits and boundary behaviour

The instrument's limits are the proof that it behaves as a count, and they are run rather than asserted.

limit	HHI	HHI^{-1}	reading
one holder, $s_1 = 1$	1	1 exactly	a monopoly counts as one unit
n equal holders, $s_i = 1/n$	$1/n$	n exactly	returns the count
n holders, any unequal split	$> 1/n$	$< n$	never overstates the count
$s_{\max} \to 1$, rest $\to 0$	$\to 1$	$\to 1$	continuous approach to the monopoly limit
$n \to \infty$, equal shares	$\to 0$	$\to \infty$	unbounded above, as a count should be
$\rho \to 0$ within blocs	naive	naive count	the independent limit
$\rho \to 1$ within blocs	bloc	bloc count	the fully-coupled limit

Range: $1 \le \text{HHI}^{-1} \le n$, with both bounds attained. Verified on this paper's own vectors: the three-Island bloc vector is a single unit at 1.000000 and returns $1/1 = 1$ exactly, and a hypothetical three-way equal split of the same exposure returns $1/(3 \times (1/3)^2) = 3$ exactly. **A quantity that equals n for n equal units, equals exactly 1 for a single unit, and lies strictly between otherwise is a count, and no further argument is needed.**

The same limits run on the affine count, which fails them. $N_{\text{eff}} = 1 + (N-1)u$ returns 1 at $u = 0$ and N at $u = 1$ for *any* share vector whatever — including a vector at $(0.999, 0.001)$, for which it reports 2.000000 at $u = 1$ while HHI^{-1} reports 1.002004. **It has the right endpoints and the wrong interior, because it never consults the shares.** It is not used here.

Boundary behaviour of the surplus law. As $N \to 1^+$, $h \to -(\delta Q + C)$ — the maintenance floor, and the value four of this paper's five headline vectors sit at exactly. As $N \to \infty$, $h \to \infty$ logarithmically, **which is a defect of the model and a statement of its domain**, since Uzzi and Spiro [2005] and Sagan [1993] establish that the real curve turns over. **The model's domain is small** N, and every prescription here is confined to the move from $N \approx 1$ to $N \approx 3$. As $\eta \to 0$, $h \to -(\delta Q + C)$ regardless of N: no conversion, no benefit from any amount of independence. As $\delta \to 0$ the stock never erodes and any positive $\eta \ln N - C$ accumulates without bound, **which is a pathology rather than an ideal.**

Boundary behaviour of the reallocation result. Moving a share x from bloc A to bloc B changes the bloc index by $2x(W_B - W_A) + 2x^2$, which is **negative — that is, N rises — only while** $W_B + x < W_A - x$, that is only while the receiving bloc, after the move, remains smaller than the sending bloc after it. **At $W_A = 0.30$ and $W_B = 0.40$ the condition fails for every positive x, which is why the reallocation of Section 7.5 could not have helped and did not.** That is H4 as an inequality rather than as an example, and it is checkable in one line.

10.4 Invariance

The effective count is invariant under relabelling of units and under any rescaling that preserves the share vector. Doubling every exposure in currency terms leaves every s_i unchanged and therefore leaves the index, N and $\ln N$ unchanged. This is a genuine scale invariance and it is why the instrument applies to a structure of any size.

It is not invariant under regrouping, and that is the whole content of the bloc collapse. **It is also not invariant under a relabelling of which bloc is "the one being reduced", which is the finding of Section 7.5**: a reallocation described as reducing exposure to a named perimeter is,

arithmetically, a reallocation increasing exposure to a different one, and the index responds to the second description.

What would break the scale invariance: any share basis that is not normalised. **Tested at how many scales?** The invariance is algebraic, not empirical, and was checked at exactly the vectors reported here — three of three entries, one of four, one of six, and three bloc vectors. **No claim of empirical scale-freedom is made, and no distributional or scaling claim is made anywhere.** Clauset et al. [2009] is the reason that matters: a principled maximum-likelihood fit with a goodness-of-fit test and a likelihood-ratio comparison leaves few canonical power-law claims standing, and **a straight line on a log-log plot is not evidence.** Stumpf and Porter [2012] adds that even a genuine power law would be weak evidence for a mechanism. **The exposure vectors here are finite sets of three, four and six shares and must not be restated as samples.**

What is conserved, and in what units. Mirowski [1989] shows that the characteristic failure of this genre is borrowing the notation of a conserved quantity without the conservation law that gave it content. **The objection is met here and the answer is short.** The conserved quantity is the exposure share vector: $\sum_i s_i = 1$ exactly, dimensionless, and it is that constraint alone that gives the index its bounds. **Remove the normalisation and the effective count ceases to be a count.** The invariance group is relabelling together with any rescaling that preserves the simplex. Q**, by contrast, is a stock with no stated conservation law and no operational unit, and this paper does not claim one for it.** Cartwright [1999] supplies the discipline that makes a partial answer respectable: specify the arrangement within which the law is claimed to hold and stop claiming a universality the argument does not need. **The arrangement claimed here is a normalised exposure vector with a correlation matrix having unit diagonal and non-negative entries, and nothing outside it.**

Aggregation is unaddressed. Kirman [1992] establishes that a single representative parameter for a heterogeneous group generally does not exist, and this paper writes one Q, one δ and one C for structures whose branches plainly differ. **No aggregation argument is offered and none is available from the arithmetic.**

10.5 Analogy discipline

borrowing	source	category	justification
portfolio variance	Markowitz [1952]	**1 — derivation**	the exposure vector genuinely is a weighted sum of shocks with a correlation matrix
Hill numbers, order 2	Hill [1973], Jost [2006]	**1 — derivation**	HHI^{-1} *is* ${}^2D$; an identity
numbers-equivalent	Adelman [1969]	**1 — derivation**	the reciprocal of the index is the count of equal units giving that index
similarity-sensitive diversity	Leinster and Cobbold [2012], Rao [1982]	**1 — derivation**, with a stated domain	exact where the matrix has unit diagonal and non-negative entries
the effective number of parties	Laakso and Taagepera [1979]	**1 — derivation**	the same formula on a different share basis

borrowing	source	category	justification
asset partitioning	Hansmann and Kraakman [2000]	**1 — derivation**	the cell company performs exactly the partitioning function entity law is identified as supplying
open design	Kerckhoffs [1883], Saltzer and Schroeder [1975]	**1 — derivation**	the disclosure standard *is* the open-design requirement applied to legal rather than cryptographic structure
$\ln N$ as entropy	Shannon [1948]	**2 — shared form**	this is **Shannon entropy in nats**, dimensionless, and not thermodynamic entropy
coupled-network fragility	Buldyrev et al. [2010]	**2 — shared form**	the percolation premises have not been shown to hold; only the structural property is used
robust-yet-fragile	Albert et al. [2000], Callaway et al. [2000]	**2 — shared form**	the degree-distribution result is not claimed to apply; the shape of the tolerance asymmetry is
synchrony in ecology	Loreau and de Mazancourt [2008]	**2 — shared form**	the naming convergence is recorded; no ecological mechanism is imported
frozen regions	Castellano et al. [2000]	**2 — shared form**	a perimeter is structurally a frozen region; **no phase transition is located and no exponent estimated**
threshold and tipping	Granovetter [1978], Schelling [1971]	**2 — shared form**	the aggregation warning is used; no threshold distribution is estimated
coordination cost at $n(n-1)/2$	Brooks [1975], Conway [1968]	**3 — framing**	qualitative organisational theory, cited for shape and never as a formal result

borrowing	source	category	justification
normal accidents	Perrow [1984], Sagan [1993]	**3 — framing**	prior qualitative statement of the same observation, never a formal result

The failure mode this table exists to prevent is letting a category-2 borrowing carry a category-1 result's authority. The concentration arithmetic is category 1 throughout and needs no help.

10.6 Measurement-theoretic status of the scored inputs

The evidence ladder of Section 6.1 is ordinal and is not summed. Its classes are ranked — "strongest", "weaker", "weakest" — and **an ordinal scale licenses a ranking and a maximum, and does not license a mean or a total.** The assignments of Section 6.3 are made by identifying the strongest available class for each jurisdiction, which is licensed. **No score is computed and no threshold is arithmetic**, and that is deliberate: a numeric evidence score would invite exactly the averaging the scale type forbids.

The twelve-item jurisdictional assessments summarised in Appendix C are heterogeneous and are not totalled. Legal family, court quality, rail access and degradation trajectory have no common scale, and no verdict in Appendix C is produced by adding them. **The assessment is a checklist that produces a verdict by argument, not a score that produces one by arithmetic.**

The framework's treatment of Q as a ratio-scale quantity is not tested. Michell [1997] is the strongest available objection: psychology and its neighbours routinely assume quantitative structure without testing for it, and the formal remedy is conjoint measurement [Luce and Tukey, 1964], whose conditions establish whether an attribute has additive quantitative structure at all. **The conjoint conditions were not tested for Q, which is stated here in unnamed index units, and the ratio-scale treatment is therefore provisional. The concentration instrument itself is untouched by this objection**, because exposure shares are ratios of capital and not psychological judgements.

10.7 Postulates

#	statement	status
P1	exposure shares s_i are non-negative and sum to 1	DEFINITION
P2	HHI $= \sum_i s_i^2$; the effective count is its reciprocal	DEFINITION
P3	for n equal shares of $1/n$, HHI $= 1/n$	ARITHMETIC
P4	a bloc is a set of jurisdictions among which one decision produces the corresponding effect without a further decision	DEFINITION
P5	the three Islands are one bloc; the four European centres are one bloc	POSTULATED (H1, H3; stipulated with the evidence of Section 6.3)

#	statement	status
P6	the conservatively assigned Atlantic captive jurisdiction is folded into WEST	POSTULATED (the conservative assignment)
P7	the surplus law's input term is $\eta \ln N$	DEFINITION (modelling choice, CCP-A02)
P8	$\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$	POSTULATED (fixed in advance)
P9	the constructed exposure vectors of Section 7	POSTULATED (constructed)

10.8 The compounding chain

step	consumes	statement	status
L1	P2, P3	HHI^{-1} returns n exactly for n equal shares and 1 for one holder	ARITHMETIC
L2	L1	HHI^{-1} is a numbers-equivalent — Adelman [1969]	CITED
L3	L2	HHI^{-1} is the Hill number of order 2 — Hill [1973]	CITED
L4	P4, P5	the three Islands collapse to one bloc share of 1.000000	POSTULATED
R1	P2, P9	Islands at 40 / 35 / 25: naive 0.345000, effective 2.898551	ARITHMETIC
R1b	L4, R1	bloc 1.000000, $N = 1.000000$, $\eta \ln N = 0$ at every η	ARITHMETIC
R2	P2, P9	crediting the Islands separately: $0.818^2 + 0.182^2 = 0.702248$, $N = 1.423998$, $h = -0.376286$	ARITHMETIC
R2b	R2	the gap to the designed architecture is $0.376286 - 0.122417 = 0.253869$	ARITHMETIC
R3	P2, P9	the European four: naive 0.270000, effective 3.703704; bloc 1.000000	ARITHMETIC
R3b	P2, P9	crediting the non-member separately: 0.745000, $N = 1.342282$, $h = -0.396970$	ARITHMETIC

step	consumes	statement	status
R4	P2, P9	at the conventional weights, $0.864^2 + 0.136^2 = 0.764992$, $N = 1.307203$, $\ln N = 0.267890$, $h = -0.406239$; **the printed 0.749296 is unreproducible from these shares**	ARITHMETIC
R5	P2, P6, P9	blocs 56.4 / 30.0 / 13.6 give 0.426592 and $N = 2.344160$, down from 2.941176, while WEST falls 30.0 → 13.6	ARITHMETIC
R6	P9	$22.0 + 9.4 + 9.5 = 40.9$ and $10.0 + 6.4 + 0.0 = 16.4$	ARITHMETIC
CR	R4, R5	a printed computation is wrong and is corrected; and a move lowering perimeter exposure lowers N	ARITHMETIC

Every ARITHMETIC row can be redone by hand or in three lines of any language. Every CITED row names one source and no more. **Only P5, P6, P8 and P9 are POSTULATED, and P5 — the bloc assignment — is the paper's entire empirical exposure.** H2 and H3 exist so that the conclusion does not depend on it. **A chain whose central step is a stipulation is not a proof, and this paper says so here rather than leaving the reader to notice: the arithmetic is exact and the assignment is a judgement.**

10.9 Cross-checks performed

- Every naive index was recomputed from its share vector in exact decimal arithmetic and agreed with the printed value to twelve decimal places, **except 0.749296, which does not agree with any computation from the stated shares and is reported as such in Section 7.4.**
- $s_{\max}^2 \le \text{HHI} \le s_{\max}$ holds on every vector: $0.160000 \le 0.345000 \le 0.400000$ on the Islands; $0.122500 \le 0.270000 \le 0.350000$ on the European four; $0.318096 \le 0.426592 \le 0.564000$ on the reallocated blocs.
- The effective-count ceiling $s_{\max}^{-2}$ is respected: $2.898551 \le 6.250000$ and $3.703704 \le 8.163265$.
- Every bloc vector sums to 1.000000 exactly: $0.818 + 0.182$; $0.850 + 0.150$; $0.864 + 0.136$; $0.564 + 0.300 + 0.136$.
- The component decomposition of Section 7.6 sums to the stated totals exactly: $22.0 + 9.4 + 9.5 = 40.9$ and $10.0 + 6.4 + 0.0 = 16.4$.
- The reallocation inequality of Section 10.3 was verified against the measured case: at $W_A = 0.30$, $W_B = 0.40$ and $x = 0.164$, the condition $W_B + x < W_A - x$ fails, and N fell.

11. Limitations

The instrument is least informative exactly where this paper's headline cases sit. Golosov [2010] and Dunleavy and Boucek [2003] document that effective-number measures behave poorly in the dominated-unit regime, where one unit holds a very large share. **Four of the**

five headline vectors here collapse to a single unit at 100 per cent, which is the limiting case. The measurement is exact and the *discriminating power* is minimal. This is the most serious methodological exposure in the paper, and its practical consequence is that comparisons among heavily dominated structures must not be made on N alone.

The correlation is assumed, not measured. The collapse uses $\rho = 1$ within a bloc, which is the conservative end of a family. No correlation was estimated from data in this paper. That is the single largest unmeasured input and it is exactly the quantity H1 and H3 turn on. The generous counterfactuals of Sections 7.2 and 7.3 exist so that a reader who rejects the assignment can still read an answer.

The bloc assignment is a stipulated reference group, not an identified effect. Manski [1993] rules out recovering an endogenous effect from observation of a linear-in-means reflexive system. The remedy adopted is stipulation with a defence [Brock and Durlauf, 2001], and the price is that the assignment is a modelling position.

H1 is met as an identity and identities cannot confirm hypotheses. Given the assignment, the bloc index of a single-bloc vector is 1.000000 by construction. H1's content is entirely in the assignment, and stating it as a prediction and then showing it is an identity is what makes that visible — a reviewer who objects that it should not have been a hypothesis is right, and the objection is recorded here rather than left to be made.

Every case is constructed. The principal, the assets, the shares, the component decomposition and the service inventory are invented and are labelled as such at each point of use. No empirical dataset was collected and none is reported.

Jurisdiction-specific legal claims are not cited to primary sources. Every enactment, decided case, listing event, survey figure, rate, date and jurisdictional characteristic in Sections 4 to 6 is taken from the source manuscript. The bibliography contains no statutes, regulations, court rules, registries or case reports for any individual jurisdiction. Appendix B lists the affected claims. This is a real gap and it is not closed here. It is also load-bearing in a specific way: the bloc assignment rests on those uncited facts, and the arithmetic rests on the assignment, so an error in the legal description propagates into every number without producing an arithmetic inconsistency.

The dated events are not independently verified. Enactment dates, listing and de-listing dates, survey shares and the dates of judicial decisions are reported as publicly announced and were not checked against primary sources by this paper.

One printed source computation is wrong and its origin is unknown. Section 7.4 corrects it and states plainly that the printed value cannot be reproduced from the stated shares. No claim is made about how it arose, and the possibility that other unreported errors exist in source computations this paper did not recompute is real and is not excluded by anything here.

Cost is asserted and not measured. The claim that a Crown Dependency or European element adds materially to C is stated without a figure, because fee levels are commercially negotiated and no defensible number was available. The direction is certain and the magnitude is not, and a reader who believes the cost is negligible should read the argument as resting on $\ln N = 0$ alone, which it can.

Monotone claims are restricted. Uzzi and Spiro [2005], Sagan [1993] and May [1972] are the reasons a monotone raise-N prescription could not survive unqualified, and Section 9 bounds it.

The instrument is exposed to its own failure mode. Adopted as a target, N would cease to be a good measure [Goodhart, 1975, Campbell, 1979]; Scott [1998] is the historical form, in which legibility destroys the local variety that made the system work, and it is a warning about this instrument as much as an argument for it. Nothing in the arithmetic prevents that, and the only defence offered is the reporting discipline of Appendix A.

The performativity problem is stated and unresolved. MacKenzie and Millo [2003], Hacking [1995], Merton [1948] and Lucas [1976]. The framework asserts that systems change when modelled and exempts its own parameters from that reflexivity.

11.1 Rigour checks not met, listed by name

1. **No as-of date, because no real holdings were measured. A concentration measure without an as-of date is not a measurement.** Every share vector here is constructed. The same applies with greater force to the jurisdictional descriptions: **a statement of a regime without a date is not a statement of that regime**, and Appendix B requires the date of every confirmation.
2. **No uncertainty is reported anywhere, because nothing is estimated.** That is not precision; it is the absence of empirical content.
3. **The dimension of η is supplied by this paper, not by the source** (Section 10.1).
4. **Q, C and h have no operational unit.** Every h difference reported is in arbitrary units and is meaningful only as an ordering.
5. **The surplus law's large-N limit is unbounded**, which the empirical literature contradicts. The model's domain is small N.
6. **The conjoint conditions were not tested for** Q [Michell, 1997, Luce and Tukey, 1964].
7. **The evidence ladder is ordinal and is deliberately not summed** (Section 10.6). Any downstream work that computes an evidence score from it will be performing an unlicensed operation.
8. **Model comparison was not performed.** No alternative index was fitted, no out-of-sample test run, no information criterion computed and no matched null constructed, because nothing here is fitted.
9. **No distributional or scaling claim is made, so the Clauset–Shalizi–Newman procedure was not run** [Clauset et al., 2009]. **Any restatement reaching for a heavy-tail claim must run it first** [Stumpf and Porter, 2012].
10. **The coupled-network, tipping and coordination-cost borrowings are category 2 or 3, not category 1** (Section 10.5).
11. **There is no conservation law for** Q [Mirowski, 1989]. Section 10.4 gives the one genuine conservation and states plainly that the stock variable has none.
12. **Aggregation is unaddressed** [Kirman, 1992].
13. **Expected value and time average are not the same thing** [Peters, 2019]. A perimeter event does not arrive in expectation; it arrives once and rescales every subsequent year, and the multiplicative consequence is not modelled here. Peters and Gell-Mann [2016] is additionally the **stronger warrant for the logarithm** in the surplus law than any perceptual argument, because it derives a logarithmic growth rate from multiplicative dynamics with no psychological assumption at all. **Where the two justifications disagree, the multiplicative one should be preferred.**
14. **The reflection problem is stated and not solved** [Manski, 1993, Shalizi and Thomas, 2011].
15. **No differential prediction is offered.** Durlauf [2005] asks for the prediction only this model makes, and the collapse makes none: any correlation-aware concentration measure would produce it. **The contribution is the identification of the correct unit of analysis, the generous counterfactuals, and the two negative results — not a uniquely discriminating forecast.**
16. **The mechanism-side companion to the arithmetic exists and is not used.** Bouchaud and Mézard [2000] derives concentration rather than observing it: multiplicative growth plus exchange produces a heavy tail, and above a critical ratio a finite fraction of the total condenses onto one holder — **which is the index going to one, obtained from a mechanism.** This paper asserts no tail and does not use the model; **the citation records that the corpus has not attempted the mechanism-side account and that a reviewer may fairly ask why not.**

12. Conclusion

This paper establishes the following results in the form the rest of the corpus cites them.

R1 — the Crown Dependency error. Three Islands at 40.0 / 35.0 / 25.0 per cent give a naive

Herfindahl index of 0.345000 and 2.898551 effective jurisdictions, **but a bloc index of** 1.000000**,** $N = 1.00$ **and a surplus term of exactly zero at every value of** η**.**

R2 — the generous counterfactual. Crediting the Crown Dependencies as their own bloc (WEST 81.8 / CD 18.2) gives a bloc index of 0.702248, $N = 1.423998$ and $h = -0.376286$ against the designed architecture's 2.941176 and -0.122417 — a shortfall of 0.253869.

R3 — the European convergence. Luxembourg 35.0 / Ireland 25.0 / Netherlands 25.0 / Switzerland 15.0 gives a naive index of 0.270000 and 3.703704 effective jurisdictions, collapsing to 1.00; crediting Switzerland separately gives a bloc index of 0.745000, $N = 1.342282$ and $h = -0.396970$.

R4 — an arithmetic error, left visible and corrected in place. A printed bloc index of 0.749296 is unreproducible from the stated shares; recomputation gives 0.764992, $N = 1.307203$, $\ln N = 0.267890$ and $h = -0.406239$. **The error is in the index alone and everything downstream of it is internally consistent with the wrong value, which is why checking a table for internal consistency does not detect it.**

R5 — a diversification that reduces resilience. Moving the residual 16.4 per cent from the Western perimeter to the Gulf bridge gives blocs of 56.4 / 30.0 / 13.6, a bloc index of 0.426592 and N falling from 2.9412 to 2.3442, **despite Western exposure dropping from 30.0 to 13.6 per cent.**

R6 — the share falls and the capability does not. The Western share falls from 40.9 to 16.4 per cent with no capability lost, because what is bought under the name of a financial centre — governing law, arbitration seat, insurance capacity, listing venue — **is portable and separable from asset situs**, and the component decomposition accounts for the fall exactly: $22.0 + 9.4 + 9.5 = 40.9$ before, $10.0 + 6.4 + 0.0 = 16.4$ after.

Counter-result. Two findings cut against the framework's own presentation. **A printed computation is shown to be wrong and is corrected rather than silently amended.** And **a move that lowers Western-perimeter exposure to 13.6 per cent nevertheless lowers the effective bloc count from 2.94 to 2.34** — which means the perimeter-exposure percentage the series reports most prominently is not, on its own, a measure of resilience. To those two the verification section adds a third, against this paper's own instrument: **the effective count is least discriminating exactly in the dominated-unit regime where four of five headline vectors sit.**

One result is established as method rather than as measurement. **Institutional quality and perimeter independence are different axes**, and a jurisdiction can score at the top of the first while contributing exactly zero to the second. The corollary is the paper's positive finding: **where a jurisdiction's value is an intangible another jurisdiction can adopt, the value can be consumed without the exposure — take the law and leave the perimeter.**

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every case in this paper — the principal, the asset mix, the exposure shares, the component decomposition, the service inventory and the variant tables — is **CONSTRUCTED** and is labelled as such at each point of use. The numerical results are exact computations on those constructed inputs, carried at full precision and rounded only for printing.

One figure reproduced in this paper is not a computation of this paper. The bloc index 0.749296 in Section 7.4 is reproduced exactly as printed in the source manuscript and is reported there as unreproducible from the stated shares; the corrected value 0.764992 and everything derived from it are this paper's computations.

All figures were produced by the scripts named in their captions (`figs/src/CCP-054_fig01.py` through `figs/src/CCP-054_fig04.py`), which import the shared figure style module and read

their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

Jurisdictional and legal descriptions in this paper are taken from the source manuscript and not from primary legal sources. The bibliography contains general treaty, standard-setting and comparative material only; it contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction. Appendix B lists the claims affected. **Any reader relying on those descriptions must verify them against the official sources and record the date of that verification.**

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. Every principal, family, entity and transaction described is invented.

Preregistration

The five hypotheses of Section 2, with their point predictions and falsifiers, were stated before the computations of Section 7 were run. **H1 was met as an identity rather than as a measurement, and Section 11 records that this is a weaker outcome than a confirmed hypothesis normally is.** No hypothesis in this paper was preregistered in any external registry.

Legal, tax and investment advice

This paper is analytical and is not legal, tax, or investment advice. It describes jurisdictions, courts, registration regimes, enforcement instruments and structural arrangements for the purpose of measuring concentration. **No recommendation is made in respect of any particular person, asset, structure or jurisdiction.** Judgment-recognition positions, substance requirements, listing rules, property tax stacks, personal tax settlements, beneficial-ownership access models and sanctions regimes all change, several of them annually; every such statement here is carried from the source manuscript, is flagged in Appendix B, and requires confirmation with qualified counsel in each jurisdiction of residence, asset situs and counterparty domicile before any reliance.

Appendix A. The portable-law rule and the reporting discipline

```
GOVERNING LAW      choose for predictability -- it is PORTABLE
FORUM              choose deliberately, NOT by habit
SEAT (arbitration) choose deliberately, and verify the current
```

```
                  institutional landscape before naming a body
ENFORCEMENT PLAN  name the jurisdictions where the debtor keeps
                  realisable assets, and work BACKWARDS from the
                  instrument that reaches them
```

A dispute clause should be drafted from the enforcement jurisdiction backwards, not from the drafter's habits forwards. If the counterparty's realisable assets sit in a state party to the awards convention [United Nations, 1958] but not to any judgments instrument with the forum state, a judgment is a document and an award is money. If the assets sit in the forum state, the reverse.

Where a jurisdiction's element is retained on capability grounds, the specification is:

```
role            NEVER the apex. A subordinate operating or
                segregation vehicle beneath a non-perimeter apex.
holds           no reserve, no group intangibles, no treasury
substance       real, and sized to the relevant-activity
                classification, re-assessed annually
disclosure      every register filing complete and CONSISTENT with
                every other register filing in the group
exit            a live migration path, documented and COSTED
measurement     booked to its true bloc, without argument
```

The exit line is the one that gets skipped and the one that matters most. A structure whose migration path has never been costed does not have a migration path; it has an intention, and intentions are executed slowly at exactly the moment when speed was the only thing that would have helped.

The reporting standard, per measurement: the share basis in one sentence; the exposure table summing to 100.0 per cent; the naive index and effective jurisdictions **with the rounding basis named**; the bloc assignment with the strongest evidence class for each jurisdiction; the bloc index and N; the largest bloc's share as a percentage; $\ln N$, $\eta \ln N$ and h at the stated parameters; a generous counterfactual over every contested assignment; **and the date, with the diff against the previous measurement.**

Two items are load-bearing and are usually omitted. The rounding basis, because Section 10.2 shows the same exposure map gives different naive indices depending on how the shares are printed. And the diff, because perimeter drift is invisible except as a difference between two measurements taken a year apart, **and the drift of failure mode 17 is the one this discipline exists to catch.**

Appendix B. Claims requiring verification against primary sources

None of the following is cited to the bibliography, which contains no primary sources for any individual jurisdiction. They are reproduced from the source manuscript and require confirmation, **with the date of confirmation recorded**, before any reliance.

1. Arbitration legislation in the forum state: which provisions of the most recent Act are in force and the commencement regulations; and the appellate decisions on the law governing the arbitration agreement and on the arbitrator's duty of disclosure.
2. Judgment-recognition position: the entry-into-force date and contracting-party list of each judgments instrument for each relevant state; the declarations made; the status of any accession application to a regional convention; and **the treatment of asymmetric jurisdiction clauses, which is contested.**

3. Listing rules and the prospectus regime: the current rulebook, the date of the most recent restructuring of listing categories, and the staged commencement of the prospectus framework.
4. Identity-verification and register regimes: current deadlines for director and controller verification; the update obligations and trust-information disclosure position of any register of overseas entities; and the scope and evidential requirements of any personal-security suppression provision. **The commencement schedules have moved more than once; never plan from a printed date.**
5. Foreign-exchange market shares and turnover: take the current triennial survey figures directly from the publishing body; the survey is triennial and the shares move.
6. The insurance market: chain-of-security figures and financial strength ratings change annually; confirm the historic reconstruction transaction and its dates from primary sources.
7. Property taxation: every line of the acquisition, annual-charge, capital-gains, income and inheritance stack requires confirmation. **Several lines moved in 2024 and 2025.**
8. The personal tax settlement: the enacted provisions, transitional reliefs, repatriation facility and trust-protection position. **Material changes occurred between announcement and enactment.**
9. Enforcement instruments: the statutory basis of any unexplained-wealth order regime and the appellate and first-instance decisions described; the sanctions framework and the regulations made under it.
10. Island trust and corporate statutes: the current numbering and text of each firewall provision, each reserved-powers provision, and any statutory mistake or inadequate-deliberation jurisdiction, with the amendment laws that introduced each. **A citation to an old text here is worse than no citation.** The perpetuity position and its treatment of pre-existing trusts. The statutory basis of protected and incorporated cell companies after consolidation, **and the recognition of cell segregation in every jurisdiction where a cell will contract.** The parallel company regimes, foundations and purpose-trust statutes of the third Island; its customs and indirect-tax agreement and revenue-sharing formula, which has been renegotiated; and its aircraft registry's establishment date and eligibility criteria.
11. Island coordination evidence: the exact short titles and commencement of each economic-substance instrument; the listing and de-listing conclusions dates; the 2019 beneficial-ownership commitment and the current access model in each Island; the sanctions instruments and the mechanism by which external designations take effect locally; and current mutual evaluation reports and follow-up status.
12. The 2008 bank failures, their outcomes, and the current terms and caps of each depositor compensation scheme.
13. The Island case authorities described — on retained control, on sham and creditor recourse, on exclusive jurisdiction clauses in trust instruments, on trustee liability and indemnity, and on the narrowing of the mistaken-exercise doctrine. **Case citation follows different conventions and no case is in the bibliography.**
14. The November 2022 supranational judgment on public register access: its operative part and the directive provision held invalid; and the successor package's legitimate-interest access conditions and application dates.
15. The minimum-rate regime: the enabling directive, each implementing instrument and its amendments, the referendum and commencement position in the non-member state, the precise formulation of the turnover threshold, and the safe harbours.
16. The anti-tax-avoidance and mandatory-disclosure directives: each transposition instrument and date, the reverse-hybrid commencement, and the crypto-asset reporting application date; and the decision on legal professional privilege in the disclosure regime.
17. State-aid decisions and their appellate outcomes; and the supranational beneficial-ownership judgments supplying an abuse principle independent of any treaty. **Confirm the joined case numbers and holdings before relying on any summary, including this one.**

18. The vehicles named in each European centre: the participation-exemption, fund, securitisation and partnership statutes; the leasing and structured-finance provisions of the common-law member state; the withholding regime, participation-exemption conditions and commercial-court rules of the third; and the banking-secrecy provision, exchange legislation, domestic tax reform and sanctions ordinances of the non-member. Take concentration and flow data from the relevant national fiscal and central-bank bodies.
19. The administration-foundation instrument: formation requirements, the tax treatment of receipt holders and the beneficial-ownership consequences, in each holder's jurisdiction.
20. All bloc assignments in Section 6 are the analyst's judgements applied to published characteristics. The underlying appellate, currency, substance, exchange and designation facts should be confirmed jurisdiction by jurisdiction.

Appendix C. Jurisdictional assessment, condensed

No total is computed: the items are heterogeneous and no common scale exists on which they could be added (Section 10.6). Every characteristic below is described from the source manuscript and is uncited (Appendix B). The jurisdictions of the ranked capability table are de-identified as "Island A / B / C" following the convention of Curzi [2026], because that table ranks named jurisdictions on legal capability the bibliography cannot support; jurisdictions carrying a bloc assignment with its evidence are named, because the assignment is the argument.

Bloc assignment and job, by jurisdiction.

jurisdiction	bloc	strongest evidence	job
United Kingdom	**WEST, at the centre rather than the edge**	founding member of the standard-setting body; an independent designating authority whose designations have converged rapidly with two others since 2022; the jurisdiction whose banks administer its own currency's clearing	**OPERATING**, and bridging only in the narrow sense of market access and governing-law contracting. **Never holding**
Jersey	**WEST**	final appeal to a committee of the neighbour's supreme court justices; the neighbour's currency with no central bank; the neighbour's designations effective locally; substance legislated to a common external deadline	**HOLDING**, as a subordinate fiduciary vehicle beneath a non-perimeter apex — never the apex

jurisdiction	bloc	strongest evidence	job
Guernsey	**WEST**	identical on every operative test	**HOLDING with a genuine segregation specialisation**, and a real risk-transfer capability
Isle of Man	**WEST, on the hardest evidence of the three**	all of the above, **and** membership of the neighbour's customs and indirect-tax territory under a revenue-sharing agreement that has been reset from outside	**OPERATING**, distinctively — registries and asset ownership inside that territory
Luxembourg	**WEST**	member of a union whose coordinated instruments it implements	**OPERATING**, for regulated-fund distribution
Ireland	**WEST**	same	**OPERATING**, for leasing and common-law fund vehicles
Netherlands	**WEST**	same	**OPERATING**, and control separation
Switzerland	**WEST on the evidence of outputs**	adoption of its neighbours' sanctions packages from February 2022 through its own embargo legislation, and alignment with successive packages since	**OPERATING**, for custody depth

Ranked capability, de-identified.

requirement	best fit	why
trust for a settlor from a forced-heirship civil-law system	Island A	firewall, reserved powers, deepest case law, statutory mistake jurisdiction
segregated cells for asset or risk ring-fencing	Island B	protected and incorporated cell companies; the original engineering and the deepest practice
captive insurance with regional proximity	Island B	established regime, insurance-linked-securities cell practice
aircraft or vessel ownership inside the neighbour's indirect-tax territory	Island C	registry inside that territory

requirement	best fit	why
operating company with the neighbour's market and indirect-tax alignment	Island C	simplified company regime, shared indirect-tax territory
listed closed-ended fund	Island A or B	established listed-fund practice

The three items that decide every row of the first table are the final appellate arrangement, the currency and monetary authority, and the sanctions alignment. **None of the eight jurisdictions can legislate itself a different one of the three**, which is why the bloc assignment is the least contestable call in the paper and why nothing available to any of them would change it.

All numerical results in this paper were computed from the stated inputs at full precision and rounded only for printing, **except the one printed value identified in Section 7.4 as unreproducible, which is reproduced there as published.**

Chapter 55

Entity Maps Understate Dependency: The Dollar Axis, the Insurance Reserve, and a Scored Captive Rubric

Abstract. A structure's exposure to a jurisdiction is not what its entity map says. Measured three ways, the same holding shows entity-based United States exposure of 27.3 per cent, a dollar-denominated share of the balance sheet of 61.0 per cent, and roughly 58 per cent of settlement volume clearing through the United States system — a dependency more than twice the entity reading. The four favoured US trust and company states, Delaware 40.0 per cent, South Dakota 25.0, Nevada 20.0 and Wyoming 15.0, give a naive HHI of 0.285000 and 3.508772 effective jurisdictions, but a bloc-HHI of 1.000000 and an effective count of 1.00, so state-level selection contributes nothing to independence. The paper then treats the insurance reserve as a measurable object rather than a claim: a captive balance sheet with 67.08 per cent committed to third-party liabilities and 25.34 per cent held as target capital leaves 7.58 per cent realisable within a year and about 27.45 per cent on a present-value basis, a 72.55 per cent haircut. The same property makes the reserve unavailable to its owner, which is stated as the cost of the result rather than hidden. Captive arrangements are scored on a hundred-point rubric across risk shifting, risk distribution, actuarial pricing, claims paid and insurer conduct, with hard zeros. Crediting Bermuda as a distinct bloc raises N from 2.941176 to 3.385332, an accounting choice rather than a structural change.

Keywords: currency exposure; dollar clearing; domestic asset protection trust; captive insurance; risk distribution; reserve realisability; concentration measurement

Series: D. Structure and Resilience **Paper:** CCP-055

1. Introduction

A structure's exposure to a jurisdiction is not what its entity map says. This paper measures the gap on a constructed cross-border holding, on three axes rather than one, and reports that the entity reading understates the dependency by more than a factor of two. It then takes a second object that is routinely asserted rather than measured — the protective character of an insurance reserve — and measures that too, in the same units, including the half of the measurement that is a cost to the owner rather than a benefit.

Three questions organise the paper, and each has a number for an answer.

First, what does sub-national selection buy? A structure exposed to one country through four of its constituent states looks, on a naive index, like four jurisdictions. Delaware 40.0 per cent, South Dakota 25.0, Nevada 20.0 and Wyoming 15.0 give a naive Herfindahl–Hirschman Index of 0.285000 and 3.508772 effective jurisdictions. Collapsed onto the coordination perimeter those four states share, the same shares give a bloc index of 1.000000, an effective count of exactly 1.00, and a surplus term of exactly zero. **The state-level split contributes nothing to independence, and unlike the comparable cases in this series it admits no generous counterfactual at all**, because the four states share one constitution, one federal judiciary, one currency, one central bank, one designating authority and one bankruptcy code. Curzi [2026] could run its counterfactual — the Crown Dependencies scored as an independent perimeter — and still reach the same verdict: Jersey 40.0 / Guernsey 35.0 / Isle of Man 25.0 gives naive HHI 0.345000 and 2.898551 effective jurisdictions but bloc-HHI 1.000000, $N = 1.00$ and a surplus term of exactly zero. Here there is no counterfactual to run, and Section 7.1 says so rather than manufacturing one.

Second, which axis is the structure actually exposed on? The same constructed holding reads entity-based United States exposure of 27.3 per cent, a dollar-denominated share of the balance sheet of 61.0 per cent, and approximately 58 per cent of settlement volume clearing through the United States system. The entity map is the axis every file records and the axis on which every subsequent decision in the constructed case had been assessed. It is the smallest of the three by a factor of 2.23. **A structure is only as diversified as its most concentrated axis, and the axis most often measured is the one that flatters.** Curzi [2026] reaches the residual form of the same finding — 55 per cent of settled value still clearing through a single system whose law is not the holder's, a floor set by the counterparties' currency rather than by entity placement — and Curzi [2026] measures the rail directly, where a naive correspondent HHI of 0.302400 (3.306878 effective correspondents) collapses to a perimeter HHI of 0.745256 and $N_{\text{rail}} = 1.341819$ once flows are decomposed by currency.

Third, how much of an insurance reserve is actually reachable? The claim that a licensed insurer's reserve resists a shareholder's creditor — the argument made for a Bermuda captive, and the reason a risk-transfer centre appears in a resilience analysis at all — is usually made as an assertion. It is measurable. On a constructed year-seven statutory balance sheet with total assets of USD 16,100,000, third-party liabilities commit 67.08 per cent, the supervisor's target capital level locks a further 25.34 per cent, and 7.58 per cent is distributable within a year; over a normal run-off the present value to the shareholder is 27.45 per cent, a haircut of 72.55 per cent. **The counter-result is in the same arithmetic**: the property that makes the pool hard for a creditor to reach — real obligations owed to unrelated parties, gated by a supervisor — is the same property that makes it unavailable to its owner. That is stated here as the cost of the result rather than filed as an unadvertised footnote.

What the paper contributes beyond restating these three measurements is a correspondence. The bloc collapse of Curzi [2026] — the operation that replaces a naive count of units with a count of independent failure points, reaching a redesigned six-jurisdiction structure at naive HHI 0.189002 (5.29 effective jurisdictions) and bloc-HHI 0.340000 (2.94 effective blocs) with Western-perimeter exposure falling from 100.0 to 30.0 per cent — is the same operation as the risk-distribution test that decides whether a captive arrangement is insurance. Both count independent exposures rather than listed entities; both are defeated by a common mode; and both are routinely inflated by adding units inside the mode. Section 7.6 runs the concentration instrument on the insurance side of the paper and derives an effective-count floor from a published share constraint. **The two halves of this paper are not two subjects. They are one instrument applied to two share bases**, which is what makes it a paper rather than a survey.

The paper presupposes the surplus functional and the affine effective count of Curzi [2026] — whose four-currency carrying cost of coin, hours, attention and foreclosure, set against the law that felt comfort decays over weeks while carrying cost does not decay at all, is the parallel result in another domain — and the roster-audit form of that count in Curzi [2026c], whose sharp Lipschitz constant $N-1$ reads as a measurement-error budget: at $N = 12$, $u = 0.05 \pm 0.05$ admits surplus anywhere in $[0, 0.7419]$, a width of 0.7419, against 0.1266 at $u = 0.70$. That budget is why no figure in this paper is quoted as though its inputs were measured.

The bloc assignment carries an identification problem, and it is stated here rather than in a footnote. The assignment is a forecast about coordinated behaviour, and forecasts about a system condition the system. Manski [1993] proves that a linear-in-means reflexive specification is **not identified from observational data**: the endogenous effect cannot be separated from contextual effects or from correlated unobservables however much data is collected, and Shalizi and Thomas [2011] show that influence and selection are generically confounded in observational network data. **No amount of observing jurisdictional behaviour could confirm a bloc assignment derived reflexively.** The remedy adopted here is the fourth of the four known routes: the reference group is stipulated in advance and the stipulation is defended — the assignment in this paper is fixed by the perimeter test of Curzi [2026], recorded with its evidence class in Section 4.1 below, and stress-tested in Section 7.7 — which is the route Brock and Durlauf [2001] shows can be made rigorous. **The price is that every bloc assignment in this paper is a modelling position and not an identified empirical result, and the paper says so wherever the assignment enters.**

Two boundaries are set before the machinery arrives, both carried from work that states them more sharply than this paper needs to. Curzi [2026] declines the inflated reading of its own borrowed physics — the framework's postulate is stated as a postulate, the relation is given only in its correct form, and a named failure mode requires that its users be able to distinguish the theorem, the ordinary observer effect and the framework's own conjecture in three separate sentences on demand — and this paper is bound by the same discipline in Section 10.5. Curzi [2026] assigns this framework's central identification, the application of the surplus functional to a life, and the reading of a day's surface as a leading indicator to the postulate drawer, and states that the third is a wager whose misfire column exists because wagers can lose. **Nothing in this paper is entitled to a stronger badge than the framework it sits inside.**

2. Hypotheses

Five hypotheses are stated, each falsifiable in the sense Platt [1964]'s strong inference requires. Where a statement cannot be falsified it is labelled a definition or a modelling choice and is not counted. Meehl [1967] is the standing warning: directional predictions are cheap, and the defensible response is a point prediction with a constant attached, which is what H1, H2, H3 and H5 do.

H1 — the no-counterfactual hypothesis. Selection among the constituent states of a single federal jurisdiction contributes exactly nothing to the effective bloc count, and unlike the comparable cases in this series it admits no counterfactual bloc assignment under which it would. *Predicted:* Delaware 40.0 / South Dakota 25.0 / Nevada 20.0 / Wyoming 15.0 gives naive HHI 0.285000 and 3.508772 effective jurisdictions, and bloc-HHI 1.000000, $N = 1.000000$, $\ln N = 0$ exactly, for every admissible assignment. *Falsified by:* any defensible bloc assignment placing two of the four states in different coordination perimeters — that is, evidence that a decision taken at the federal level reaches one and not another without a second independent decision.

H2 — the axis-mismatch hypothesis. The entity axis systematically understates a structure's exposure to a jurisdiction whose currency it holds. *Predicted, as a point prediction on the constructed case:* entity-based United States exposure 27.3 per cent against a dollar-denominated balance-sheet share of 61.0 per cent and settlement volume through the United States system of approximately 58 per cent — a ratio of $61.0/27.3 = 2.234$ on the currency axis and $58.0/27.3 = 2.125$ on the settlement axis. *Falsified by:* a currency and settlement decomposition of the same structure that returns a share at or below the entity share. **The hypothesis is directional about the sign and specific about the size; it is tested here on one constructed structure only, which Section 11 states as its weakest feature.**

H3 — the realisability hypothesis. The share of a licensed insurer's assets that a shareholder's creditor can realise within a year is a small and computable fraction of the balance sheet, and is set by the reserve and the regulatory capital requirement rather than by any

protective drafting. *Criterion, fixed in advance:* realisable-within-a-year is total assets less third-party liabilities less the target capital level, divided by total assets; on the constructed balance sheet this is predicted to fall below ten per cent. *Falsified by:* a balance sheet on which that quantity is large, or a legal route by which committed reserves are realisable without supervisory consent. **A creditor who obtains the reserve itself falsifies this hypothesis directly.**

H4 — the retained-programme hypothesis. A captive funding a genuine retained layer recovers a modest and boundable fraction of the prior commercial spend, not a large one. *Predicted:* on the constructed programme, $2{,}700{,}000 - 410{,}000 - 1{,}560{,}000 - 140{,}000 = 590{,}000$, which is $590{,}000 / 4{,}370{,}000 = 13.50$ per cent of the previous commercial spend of $4{,}370{,}000$, before investment income and before any tax consequence. *Falsified by:* a comparable programme returning materially more than that without either retaining more risk than it discloses or failing to pay claims.

H5 — the equivalence hypothesis. The risk-distribution test applied to a captive's book and the bloc collapse applied to a jurisdictional structure are the same operation on different share bases, so that a published constraint on insured shares is equivalent to a floor on the effective count. *Predicted:* the constraint "twelve subsidiaries each contributing between five and fifteen per cent of the captive's total risk" implies an effective count in $[9.090909,\ 12.000000]$, and "nine unrelated members none above sixteen per cent of premium" implies $[6.443299,\ 9.000000]$, exactly. *Falsified by:* a share vector satisfying either constraint whose reciprocal index falls outside the stated interval.

Not hypotheses, and labelled as such. Four further statements appear in this paper and none is falsifiable as stated.

- *The logarithmic form of the input side*, $\eta \ln N$ rather than ηN, is a **modelling choice** carried from Curzi [2026] and not tested here. Section 11 records that Latané [1981]'s sublinear power law fits the qualitative facts as well, and that nothing in this paper discriminates between them.
- *The parameter set* $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ is a **modelling choice** carried from Curzi [2026] and fixed before the measurement, which is the preregistration discipline of Nosek et al. [2018] applied to an engineering measurement.
- *The identification of the surplus law's* N *with the effective bloc count* is a **definition**. Its content is in the bloc assignment rule, which is testable, and not in the identification.
- *The five-test rubric of Section 6* is a **working instrument constructed by this framework**, adopted by no tribunal, and its weights are stipulated rather than estimated. Section 10.6 states its scale type and Section 11 states what follows.

3. Method

3.1 The instrument, and which effective count it is

The concentration instrument is taken unchanged from Curzi [2026], whose redesigned six-jurisdiction structure reaches naive HHI 0.189002 (5.29 effective jurisdictions) and bloc-HHI 0.340000 (2.94 effective blocs) with Western-perimeter exposure falling from 100.0 per cent to 30.0 per cent. Its definition fits on one line:

```
HHI = sum of squared shares            effective count = 1 / HHI
```

introduced by Hirschman [1945] and reached independently in sum-of-squares form by Herfindahl [1950]. The reciprocal reading — the count of equal-sized units that would produce the observed index — is due to Adelman [1969]. Competition practice states the index in points on percentage shares [U.S. Department of Justice and Federal Trade Commission, 2023, 2010]; this paper states it as a fraction, and gives both wherever confusion is possible.

Which effective count, stated exactly, because this framework runs two. The reciprocal index HHI^{-1} is the Hill number of order two [Hill, 1973], whose numbers-equivalent reading Jost [2006] states in its cleanest form; the same arithmetic was reached in ecology by Simpson [1949] and in political science by Laakso and Taagepera [1979]. The exponential of Shannon entropy [Shannon, 1948] is the same family at order one, and the reciprocal index is also the exponential of the Rényi entropy of order two [Rényi, 1961]. **Every effective count in this paper is HHI^{-1}, the Hill number of order two, and it is named as such at each use.** The framework's other effective count, $N_{\mathrm{eff}} = 1 + (N-1)u$, is a linear interpolation between 1 and N; it is **not a Hill number of any order** and carries no external support, as Curzi [2026] establishes by running the same verification chain on both and letting one fail visibly. It is referred to here by name and is never used to produce a number.

The general form. The bloc collapse generalises the reciprocal index to $\mathrm{EFF} = 1/(w'Rw)$ for exposure weights w and correlation matrix R, which is the similarity-sensitive diversity of order two of Leinster and Cobbold [2012], whose additive relative is the quadratic entropy of Rao [1982]. Setting $R = I$ recovers HHI^{-1}; setting R to a block matrix of ones recovers the bloc index. Hurlbert [1971] is the standing objection that "diversity" carries no content without its formula attached, which is precisely the discipline this section imposes.

No distributional or scaling claim is made anywhere in this paper. A share vector of 40.0 / 25.0 / 20.0 / 15.0 descends steeply and a reader may reach for a tail argument. **This paper fits no distribution, characterises no tail and asserts no power law.** Clauset et al. [2009] is the reason: on twenty-four canonical data sets a principled maximum-likelihood fit with a goodness-of-fit test and a likelihood-ratio comparison against alternatives leaves few standing, and a straight line on a log-log plot is not evidence. That procedure was not run, because no claim requiring it is made. Stumpf and Porter [2012] adds that even a genuine power law is weak evidence for a mechanism, because many mechanisms produce the same tail. **The effective counts here are statistics of finite exposure vectors with three, four, six or thirteen entries, not estimates of distributions, and must not be restated as such.**

The collapse result has an ancestry that should be named rather than reinvented. A small correlated shift producing a discontinuous collective outcome is the threshold and tipping mechanism. Granovetter [1978] establishes that the distribution of individual thresholds and not the mean determines whether a collective outcome tips, and Schelling [1971] is the canonical demonstration that a mild individual preference produces a sharp aggregate result — the strongest available warning against reading an individual-level parameter straight into an aggregate, which is exactly what a bloc assignment does. Castellano et al. [2000] supplies the closest formal analogue of a bloc: the same interaction that raises local agreement freezes global diversity into non-interacting regions, with a sharp transition between the monocultural and fragmented phases. **A bloc, in that vocabulary, is a frozen region and the effective count is a count of them.** Bouchaud and Mézard [2000] derives concentration rather than observing it — multiplicative growth plus exchange condenses a finite fraction onto one holder above a critical ratio, which is the index going to one, obtained from a mechanism. This paper does not use that model and asserts no tail; the citation records that a mechanism-side companion to the arithmetic exists.

3.2 The share basis, and the four axes

An index without a stated threat is an index without a share basis. This paper measures against **arbitrary state action** — a designation, a listing, an exchange-control measure, a compliance department's change of risk appetite — and not against lawful process, which has the opposite structural signature and is the subject of Section 4.2. Against arbitrary action the useful property is sitting outside a coordination perimeter; against lawful process that property is close to worthless, because the enforcement mechanism travels to the person and the person is inside the perimeter.

The instrument is identical on each of four axes and only the share basis changes.

axis	share basis	what a decision reaches
jurisdiction	fraction of protected capital whose availability depends on a decision taken inside a perimeter	entities, registries, courts
settlement rail	fraction of settlement volume depending on a given clearing or messaging system	instructions in flight
counterparty	fraction of protected capital held as a claim on, or custodied by, a given institution	the ability to instruct
currency	fraction of the balance sheet denominated in a given unit	every balance, wherever booked

Three consequences follow and all three do work below.

A portfolio contributes to the perimeter of its custodian, not of its issuers. A globally diversified portfolio held through a custodian in one jurisdiction is an exposure to that jurisdiction, and the diversification of the issuers is a fact about a different question. This is the step at which the constructed case's file was wrong.

A currency is a perimeter exposure that no entity move can exit. A balance denominated in a unit whose clearing runs through banks inside a perimeter carries that perimeter's exposure however the entities are arranged. The currency axis therefore interacts with the jurisdiction axis rather than being independent of it.

A structure is only as diversified as its most concentrated axis. The reporting discipline is to run all four and print all four; Appendix A states the reporting standard, and item 13 of it — naming the largest of the four as the binding one — is what this paper adds to the standard Curzi [2026] set.

One distinction ruins naive measurements on the settlement axis and is recorded here so Section 7.2's figure is read correctly. **A clearing system and a messaging system are different layers.** A structure reporting two rails may have one messaging dependency, so that "one messaging network at one hundred per cent" and "eighty-five per cent on that network plus fifteen per cent on a second clearing system whose messaging leg runs over the first" are the same number, and only the second looks diversified. Curzi [2026] decomposes the rail properly and reports the collapse; this paper uses the settlement figure as a single share and does not decompose it further, which is a limitation and is recorded in Section 11.

3.3 Correlated units, and what the literature says about the correction

The naive index overstates resilience because it treats units as independent when they share a coordination mechanism. The correction is not novel as mathematics. It is the observation, stated for portfolios by Markowitz [1952], that diversification benefit depends on covariance and not on count; and stated for coupled infrastructures by Buldyrev et al. [2010], that a system of interdependent networks is far more fragile than either network alone. Perrow [1984] reached the organisational form first — interactive complexity plus tight coupling makes catastrophic failure a property of the system rather than of any component — and Sagan [1993] documents the sharper version, that redundancy can reduce safety, which is a direct threat to any prescription that says raise the count. **Both borrowings are labelled under the analogy discipline of Section 10.5, and this paper's signature finding is Perrow's result in a measurement frame rather than a new observation.**

Correlations rise under stress, which is the part that matters operationally. Longin and Solnik [2001] find extreme correlation in equity markets that is asymmetric and concentrated

in downturns, and Ang and Chen [2002] report the same asymmetry. **The effective count falls exactly when it is needed.** The ecological name for the decisive variable is synchrony [Loreau and de Mazancourt, 2008], and the strongest empirical demonstration that effective count buys stability is Schindler et al. [2010]'s Bristol Bay sockeye record, where population diversity produces roughly half the variability and about a tenth of the fishery closures. The contrary literature is read alongside and is not decorative: Doak et al. [1998] show that diversity–stability correlations can arise from statistical averaging alone with no mechanism, and May [1972] showed that increasing diversity and connectance destabilises randomly assembled systems. This paper's claim is narrower than the general diversity–stability claim: it asserts only that the reciprocal of a variance is an effective count of independent shocks, which is arithmetic, and that units sharing a common mode are correlated, which is H1 and H5.

3.4 Evidence standard, and the de-identification convention

Sections 4, 5 and 6 are jurisdictional and doctrinal exposition. **Every enactment, decided case, revenue ruling, licensing class, capital threshold, listing event, rate and date in them is described from the source manuscript and is not cited to the bibliography.** The bibliography carries the general instruments and the comparative doctrine only: recognition and enforcement of awards [United Nations, 1958] with the grounds on which enforcement is in fact refused [van den Berg, 1981] and the arbitration doctrine around it [Born, 2021]; the trusts convention [Hague Conference on Private International Law, 1985] and the succession convention that never entered into force [Hague Conference on Private International Law, 1989], with the regional regulation that partially fills the gap [European Union, 2012]; conflict-of-laws doctrine [Torremans et al., 2022, Briggs, 2013]; the trust and the foundation compared [Lupoi, 2000, Smith, 2013, Panico, 2014]; asset partitioning as the irreducible function of entity law [Hansmann and Kraakman, 2000]; the listing machinery as a technique of governance [de Goede, 2011]; and the disclosure instruments [Organisation for Economic Co-operation and Development, 2017, Financial Action Task Force, 2012, Organisation for Economic Co-operation and Development, 2015]. **It contains no statutes, regulations, court rules, registry documents, revenue rulings or case reports for any individual jurisdiction.** Appendix B lists every affected claim, **with the requirement that the date of confirmation be recorded**, because a statement of a jurisdiction's regime without a date is not a statement of its regime.

Where a table would otherwise **rank named jurisdictions against one another on a legal capability the bibliography cannot support**, the jurisdictions are de-identified — "trust-state A" through "trust-state D" — following the convention of Curzi [2026], which de-identifies two Gulf centres in its own scored assessment for the same reason, and Curzi [2026], which de-identifies three islands in its ranked capability table. **Where a jurisdiction is named in prose, or appears in an exposure table, or carries a bloc assignment with its evidence, the name is used**, because the argument is about those specific perimeter attachments and a functional description would hide rather than protect.

3.5 The realisability measurement, defined

The protective character of an insurance reserve is stated in this paper as a measurement with a definition, so that it can be recomputed on any balance sheet:

```
realisable within one year
    = (total assets - third-party liabilities - target capital level)
      / total assets

present-value realisation
    = (statutory capital and surplus + reserve development)
      / (1 + r)^t / total assets
```

```
haircut = 1 − present-value realisation
```

Three inputs are stipulated rather than measured and are named here so they are not mistaken for findings: the reserve development assumption, the discount rate r, and the average settlement lag t. Section 10.2 gives their epistemic status and Sections 10.3 and 11 give the sensitivity.

The measurement is symmetric and the paper reports both signs. The same three quantities that bound what a creditor can take bound what the owner can take. That is the counter-result of Section 7.4, and it is not a caveat appended to the finding; it is the finding, read from the other side.

3.6 The rubric as an instrument, and its scale type

Section 6 scores captive arrangements on five tests. The scoring is stated here as method so that Results can be read as measurement rather than as judgement.

The five tests carry weights summing to 100 — risk shifting 25, risk distribution 25, actuarial pricing 20, claims paid 20, insurer conduct 10 — and each is scored on an integer band from 0 to 4, so the weighted total runs to 400. **A zero on risk shifting or on risk distribution is disqualifying regardless of total.** The two disqualifying tests carry exactly half the weight, which is a design choice and is recorded as one.

The scale type is stated because it constrains the arithmetic. The 0–4 bands are ordinal: they rank arrangements and they do not license the assertion that the interval from 2 to 3 equals the interval from 3 to 4. Weighting and summing ordinal scores is an operation the scale does not license, and this paper does it, and says so. The weighted total is therefore read as **a ranking with a disqualification rule attached, not as a quantity**; no difference between two totals is interpreted, and no total is compared across arrangements except as an ordering. **The disqualification rule carries the evidential weight, and it is a logical test rather than an arithmetic one.** Stevens [1946] is the taxonomy, Michell [1997] is the standing objection that quantitative structure is assumed rather than tested, and Luce and Tukey [1964] is the conjoint test that would settle it and was not run. Section 10.6 states the consequence and Section 11 concedes it.

4. The federal overlay, and why no counterfactual exists

This section is the evidence for H1 and it is compressed to what the hypothesis needs. **Everything in it is described from the source manuscript and is uncited** (Section 3.4, Appendix B).

4.1 What the four states share, and the evidence class of the assignment

The perimeter test of Curzi [2026] asks whether a single decision taken in one jurisdiction reliably produces the corresponding effect in another **without a further independent decision**. Applied to four constituent states of one federal jurisdiction the test does not require judgement, and that is the distinguishing feature of this case. The evidence ladder used across this series ranks demonstrated near-simultaneous mirroring of designations, shared correspondent networks that exit corridors together, a supervisory standard producing the same domestic consequence in every member, and binding accession to a common enforcement instrument as the strongest classes, with information-exchange participation and treaty density weaker, and diplomatic alignment weakest.

shared mechanism	evidence class
one constitution, with a supremacy clause subordinating state law to federal law	strongest — binding accession, in its limiting form
one federal judiciary, with a contempt power operating on persons	strongest
one currency and one central bank	strongest
one designating authority, whose measures apply in every state	strongest — mirroring is automatic, not merely near-simultaneous
one bankruptcy code, with a ten-year avoidance provision reaching self-settled devices	strongest
a full-faith-and-credit obligation circulating sister-state judgments	strongest
one federal reporting, withholding and exchange architecture	strongest

Every row is in the strongest class, and none of them is a judgement call. That is what distinguishes this case from every other bloc assignment in the series: elsewhere the assignment is a defended stipulation exposed to the identification problem of Section 1, and here the stipulation has no free parameter left in it. Sections 4.2 and 4.3 record what state selection genuinely does, which is a real thing and is not this thing.

4.2 The mechanism that state law does not reach

The reason a sub-national split buys nothing on this axis is that the binding constraint operates on a person rather than on an asset. A court with personal jurisdiction over a person may order that person to act and may imprison them for refusing; where the assets are, what the instrument says, and whether a foreign fiduciary has declined to comply are not inputs to that mechanism. The source manuscript records two appellate decisions in which an impossibility defence failed on the ground that the impossibility was self-created, and one in which a debtor remained incarcerated for years. **No state statute, no non-recognition provision and no firewall addresses a mechanism that does not operate on the assets at all.**

A second and distinct failure runs through conflict of laws. A self-settled protective trust depends on another state's court applying the trust state's law, and the source records three decisions in which it did not: a state supreme court holding that a statute purporting to confer exclusive jurisdiction on its own courts could not divest other states' courts of jurisdiction; a bankruptcy court applying the law of the state with the most significant relationship together with the federal avoidance provision; and a state supreme court applying its own public policy against self-settled spendthrift trusts to a trust of another state. **The protective statute is only as good as the willingness of a foreign forum to apply it, and that willingness is not something the enacting state can legislate.** The general doctrine is in the bibliography [Torremans et al., 2022, Briggs, 2013, Fawcett and Carruthers, 2008]; the individual decisions are not, and are listed in Appendix B.

The charging-order remedy has an analogous limit. Its doctrinal justification is the protection of other members from an intrusion by one member's creditor; where there are no other members the justification is absent, and the source records an appellate decision holding that the charging order was not the exclusive remedy against a single-member entity. **This is asset partitioning doing exactly what Hansmann and Kraakman [2000] says it does and nothing more** — partitioning is the irreducible function of entity law, and partitioning against a claimant who has personal jurisdiction over the partitioner is not one of the things it does. Pistor [2019] is the account of how legal coding creates capital together with its sharpest critique, and it is the right pairing here because the section describes a coding operation whose limits are legal rather than technical.

A timing observation belongs here and is stated because it is easy to mistake for a jurisdictional one. The governing law against a protective transfer is fraudulent-transfer law, and fraudulent-transfer law is a clock. Whether a transfer is vulnerable turns on when it was made relative to when a claim arose, not on where the recipient vehicle was formed. **A timing problem wearing a jurisdiction costume is still a timing problem**, and every element of the specification the source states — that the structure works against a creditor who must sue in the formation state, is outside the federal avoidance window, cannot establish actual intent, and is not a court with personal jurisdiction prepared to use contempt — is about timing and about who the creditor is.

4.3 What state selection buys, de-identified

The four states are not interchangeable and the differences are real. **The comparison below ranks them on legal capability the bibliography cannot support, so they are de-identified** (Section 3.4). Every characteristic is described from the source manuscript.

requirement	best fit	why
an operating company that expects governance litigation	trust-state A	a specialist non-jury equity court operating since the eighteenth century, judges who hear nothing else, and the deepest body of published decisions; the value proposition is answered questions
a perpetual dynasty vehicle with a directed-trust architecture	trust-state B	perpetuities abolished, a codified separation of investment, distribution and administrative functions, and an unusually strong sealing provision
the shortest seasoning period, and charging-order protection extended to closely held corporate shares	trust-state C	the most aggressive of the four on timing, and distinctive in extending the remedy beyond the limited-liability form
a cheap, low-friction single-purpose vehicle with genuine formation privacy	trust-state D	the originator of the form, the lowest maintenance cost, no public member disclosure, and a developed private-trust-company regime

A jurisdiction with a thin body of decisions offers more favourable text and less predictability, and for an entity likely to face governance litigation the second consideration usually dominates. **Choose the deep bench for governance complexity and the cheap registry for simple single-purpose vehicles, and do not let a fee differential decide it.**

The engineering reading of a jurisdiction that legislates a considered answer to an adverse decision within months cuts both ways, and both halves belong on any assessment: institutional responsiveness is exactly what a structure should want, and a settlement that has become a live political variable rather than a fixed background condition is a different and less comfortable finding. Curzi [2026] reaches the general form of that observation — institutional quality and rule stability are independent properties, and a jurisdiction can have world-class courts and an unbroken constitutional order while replacing a century-old settlement inside

thirteen months — and it is not restated here. **Score the two separately, and never let a high score on the first suppress an honest score on the second.**

4.4 Formation privacy is not protection, and the disclosure constraint

The two are unrelated, and the jurisdiction under study is the clearest available case of a place with a great deal of the first and very little of the second. **A structure whose only asset is that nobody knows about it fails this framework's disclosure test by construction** — and it fails it in the jurisdiction with the most effective discovery apparatus available.

The premise is not this work's to claim. Kerckhoffs [1883] stated it in 1883: a system must remain secure when everything about it except the key is public. Saltzer and Schroeder [1975] restate it as *open design* — the protection mechanism should not depend on the ignorance of attackers — and place it alongside *complete mediation* precisely because a publishable design still requires that every access be checked. Brandeis [1914] is the public-interest form of the same observation from the other direction. **Citing Kerckhoffs converts what reads as an ethical preference into a recognised engineering standard**, and Curzi [2026] states the constraint in its strongest form: channel separation is strengthened by disclosure, because its effectiveness is entirely independent of whether anyone knows it is in place, so every element must be publishable and non-disclosure is refused as a design premise rather than merely discouraged. The regulatory direction has made the same point unavoidable: automatic exchange [Organisation for Economic Co-operation and Development, 2017], the anti-money-laundering recommendations that supervisors implement in common [Financial Action Task Force, 2012], and the substance and beneficial-ownership work of the base-erosion project [Organisation for Economic Co-operation and Development, 2015] together mean that non-disclosure is not an available design parameter whatever a designer's preference.

The constructed case's most volatile input is a reporting position that changed twice in eighteen months, and the source manuscript records the sequence — enactment, implementing rule, constitutional challenge, a nationwide injunction, a stay, and an interim rule narrowing the regime — without asserting the current state of it. **Any structure whose rationale depends on a current reporting position is built on the least stable variable available.** The related error is assuming that a state's non-participation in an information-sharing arrangement matters: federal reporting, withholding and exchange operate independently of any state arrangement, and the reporting comes from the institution and from federal law rather than from a state registry.

Section 7.5's disclosure reading applies the test in print, and Section 9 records the result that follows from it: **the reserve measurement of Section 7.3 is one of the few results in this series whose protective characteristic is increased rather than reduced by publication**, because every fact that makes the pool hard to reach is a fact a supervisor requires to be filed.

5. The reserve as a measured object

5.1 The three legs, separated

The claim under test is that a licensed insurer's reserve resists a shareholder's creditor. Stated whole it is an assertion; separated it is three independent conditions, at least one of which fails in most arrangements. **Everything in this section is described from the source manuscript and is uncited** (Appendix B).

Leg one — technical provisions are liabilities and shareholders stand last. This is ordinary corporate and insolvency law and not an offshore feature. A shareholder's creditor takes the residual after the company's creditors, and the insurer's creditors include every policyholder and ceding company with an outstanding or incurred-but-not- reported claim. Whether a statutory *priority* over general unsecured creditors exists is a separate question and is listed in Appendix B; the argument does not depend on it, because the ordinary subordination of

equity to debt is sufficient. **A priority would add protection against the insurer's own general creditors, which is a different risk from the one under measurement.**

Leg two — the supervisor gates the exit. Distributions above stated thresholds require prior approval or a filed affidavit, the minimum solvency margin and the enhanced capital requirement must be maintained continuously rather than at the year end, and the supervisor has powers to direct, restrict business, require capital and petition for winding up. A creditor who takes the shares owns a company whose cash cannot be extracted at will by anybody, including him. **The gate exists to protect policyholders and would be improper if it existed to protect a shareholder from his creditors; it protects the pool from every direction at once as a side effect of doing its actual job.**

Leg three — the impaired parties are real, identifiable and represented. This is the leg that distinguishes an insurance reserve from a special-purpose vehicle with book liabilities to affiliates, and it is the leg that fails in the arrangements Section 7.5 scores as disqualified. If the policyholders are unrelated commercial parties, or a material share of the book is assumed from unrelated cedants, any process impairing the reserve produces identifiable victims with their own claims, their own counsel and a supervisor with a statutory duty toward them. **A court asked to make an order with that consequence is being asked to do something it will not do casually.** If the only policyholders are companies in the same group as the judgment debtor, none of that is true: the third parties are the debtor in another hat, and the argument evaporates in a paragraph.

The three legs together explain why the characteristic strengthens as the book lengthens in tail. Short-tail property risk settles within a year and the reserve turns over; long-tail liability risk carries reserves for years and the committed pool is correspondingly larger and more durable. **The characteristic is a function of reserve duration**, and an arrangement writing only short-tail risk has much less of it than the description suggests. This is a boundary on the result and it is stated with the result rather than after it.

5.2 What the reserve does not do

Four limits are stated before the number, because a section that sells an argument before bounding it is selling something.

The shares are ordinary property. They can be charged, attached and transferred like any other shares. What resists realisation is not ownership of the insurer but immediate realisable value from it. **The result is a discount and a delay, not a wall**, and Section 7.3 puts a number on both.

It answers arbitrary action, not a judgment. A court with jurisdiction over the shareholder can order him to procure a distribution, and a supervisor's likely refusal is a fact about the supervisor, not a defence to contempt. This is the threat-2 boundary of Curzi [2026] and the two threats have opposite structural signatures: against arbitrary action the useful property is being outside a perimeter, and against lawful process that property is close to worthless because the mechanism travels to the person.

It collapses if the insurance is not real. If the reserve stands against risks assumed from group companies only, there are no unrelated third parties whose impairment makes realisation awkward. If the reserve is invested in loans back to the insured group, the asset backing the third-party liability is a claim against the balance sheet the creditor is already attacking, and the circularity destroys the argument on its own terms. **The doctrinal test for whether an arrangement is insurance and this framework's test for whether a reserve resists realisation point at the same underlying fact**, which is whether unrelated parties have real claims on the pool. Section 6 is that test.

It does not move the concentration reading. Bermuda, which hosts the constructed insurer, is inside the same coordination perimeter as the structure's other Western holdings on every operative test, and Section 7.7 prices exactly what crediting it as a distinct bloc would be worth. **A reader who came to this section looking for a bloc will not find one.**

5.3 Why the regulatory gate is a side effect and not a design feature

One reading of Section 5.1's leg two is worth separating out, because it is the reading that makes the result publishable under the disclosure constraint of Section 4.4.

The supervisory gate is not a protective device that a structure procures. It is a prudential control that exists for policyholders, it operates against every route out of the company including the owner's own, and its effect on a shareholder's creditor is incidental. **That is precisely why the characteristic survives disclosure.** A protective device that works because it is procured can be unwound by the party who procured it, and a court can order it unwound. A prudential control cannot be unwound by the shareholder at all, which is the same fact as the counter-result of Section 7.4 stated from the supervisor's side.

The consequence for how the result may be described is exact. **The reserve does not resist a creditor because it is offshore, because it is opaque, or because a document says so. It resists because somebody else is owed the money and a regulator is watching, and both of those facts are on a public register.** Every element of the mechanism is publishable, and publishing it strengthens rather than weakens it, which is the Kerckhoffs condition of Section 4.4 satisfied rather than merely invoked [Kerckhoffs, 1883, Saltzer and Schroeder, 1975].

6. What makes insurance real: five tests and a scored rubric

6.1 The five tests, made operational

Described from the source manuscript and uncited (Appendix B). The doctrinal anchor is a decision requiring both **risk shifting** and **risk distribution**, and a working four-question formulation asking whether the arrangement involves insurance risk as distinct from investment or business risk, whether there is risk shifting, whether there is risk distribution, and whether the arrangement conforms to commonly accepted notions of insurance. **The fourth question is the catch-all and it does real work**: it is where a tribunal puts the observation that no unrelated party would have bought this cover, that the insurer had no claims department, or that the wording was assembled from templates for a business the insured does not operate.

Test 1 — risk shifting. The evidentiary question is whether anything runs back. Parent guarantees, letters of comfort, loans from the insurer to the insured or its owners, premium-adjustment or profit-commission clauses that return claims paid, unfunded capital commitments treated as capital: any one is a finding against the arrangement, and a loan-back is the single most damaging fact in the reported decisions the source describes.

Test 2 — risk distribution. This must be counted, not asserted, and it is the test this paper measures. Distribution means the insurer has pooled enough **independent** exposures that aggregate variance falls. **That is the same condition Markowitz [1952] states for a portfolio and the same condition the bloc collapse states for a structure**, and it is defeated by the same thing: a common mode. The counting rule the source states operationally is that eleven subsidiaries in eight countries with their own workforces, premises and customers count; one operating company with eleven product lines behind a single manufacturing process counts only partly; forty vehicles in one fleet with one depot, one driver pool and one weather system count weakly; and seven unrelated members with independent operations count strongly. **The question is never how many things are on the schedule. It is how many independent ways the pool can lose money** — which is the effective count, in a different vocabulary.

Three published share constraints are described in the source as safe harbours: twelve subsidiaries each contributing between five and fifteen per cent of total risk; seven or more unrelated members none contributing a disproportionate share; and unrelated premium at around half the book. **The constraints are uncited; the effective-count intervals Section 7.6 derives from them are arithmetic.**

Test 3 — actuarial pricing. A written pricing report produced *before inception*, deriving premium from exposure, using the insured's own loss development where credible and in-

dustry data where it is not, with explicit expense and risk loads and a stated confidence level. The report must be capable of answering the only question that matters, which is whether an unrelated commercial insurer would have written this cover at this price. **A premium determined first and justified afterwards leaves a documentary trail that says so**, and the source records that the adverse decisions are full of pricing reports dated after the policies incepted.

Test 4 — claims paid. The test that cannot be simulated: a notification procedure, a register, adjusters or a claims manager, reserves set on individual claims, and payments made, with a multi-year paid and incurred ratio broadly consistent with the pricing. **An arrangement that has collected premium for seven years and paid nothing has not proved it is prudent; it has proved either that its pricing was wrong by an order of magnitude or that it was never insuring anything.**

Test 5 — insurer conduct. Adequate capital for the assumed risk; liquid, duration-matched investments free of related-party exposure; real policy wordings for the actual business; a board that meets and decides; a supervisor that is engaged rather than merely licensing; audited financial statements; reinsurance purchased from unrelated markets.

6.2 The rubric, and the hard zeros

```
TEST                    WEIGHT   SCORE 0–4      HARD ZERO
------------------------------------------------------------------
1 Risk shifting            25        0–4        anything runs back
2 Risk distribution        25        0–4        fewer than 5 units
3 Actuarial pricing        20        0–4        no pre-inception report
4 Claims paid              20        0–4        no claims in 3+ years
5 Insurer conduct          10        0–4        related-party investment
------------------------------------------------------------------
Weights sum to 100; weighted total runs to 400.
A ZERO ON TEST 1 OR TEST 2 IS DISQUALIFYING REGARDLESS OF TOTAL.
```

The disqualification rule is the load-bearing line and the totals are not. Risk shifting and risk distribution are constitutive of insurance on the doctrine the source describes; an arrangement failing either is not a poorly documented insurance arrangement but is not insurance, and no quality in the other three repairs it. **A reader who takes away a number from this rubric has taken away the wrong thing**, which is why Section 10.6 demotes the total to an ordering and Section 11 concedes the scale-type objection rather than defending it.

The rubric is this framework's working instrument. No tribunal has adopted it, no adviser should present it as authority, and a high score is evidence rather than immunity. Section 7.6 identifies a defect in its Test 2 threshold and states the correction rather than applying it silently.

7. Results

All figures below are exact computations on constructed inputs, carried at full precision and rounded only for printing. **No external dataset was collected and no figure in this paper is a measurement of any real structure.** Where the source manuscript's printed value differs from the exact computation, Section 7.8 prints both.

7.1 R2 — sub-national selection, and H1

Shares Delaware 40.0, South Dakota 25.0, Nevada 20.0, Wyoming 15.0, as percentages:

```
naive HHI = 0.160000 + 0.062500 + 0.040000 + 0.022500 = 0.285000
effective jurisdictions = 1 / 0.285000 = 3.508772
```

```
bloc collapse, all four inside one perimeter
bloc-HHI = 1.000000        N = 1.000000
ln(N) = 0.000000           eta ln(N) = 0.000000
h = 0.35 x 0 - 0.02 x 10 - 0.30 = -0.500000
```

The state-level split contributes exactly zero to the surplus term. The naive reading of 3.508772 is the count an adviser can present and it is not wrong as arithmetic; it is the answer to a question about entity placement rather than about independent points of failure.

The proposals actually put in the constructed case, and what each measured:

proposal	bloc-HHI	N	$\ln N$
leave the holding vehicle where it is	1.000000	1.000000	0.000000
move the vehicle to a second state	1.000000	1.000000	0.000000
split across three states	1.000000	1.000000	0.000000
split across all four, 40 / 25 / 20 / 15	1.000000	1.000000	0.000000
any generous counterfactual	—	—	—

The third proposal would have added two registered agents, two sets of state filings and two more sets of counsel in exchange for no change in any measurement in this series. **It was a sensible asset-protection suggestion presented as a resilience measure, and it is not one.**

H1 is supported, and the second half of it is what distinguishes this case. Curzi [2026] could run a generous counterfactual — the three Crown Dependencies scored as an independent perimeter rather than folded into the Western bloc — and report that the conclusion survived it. **No such counterfactual exists here.** Every row of the evidence table in Section 4.1 sits in the strongest class and none is a judgement call, so there is no admissible reassignment under which two of the four states occupy different perimeters. **The last row of the table above is empty, and it is left empty rather than filled with a weak assumption.** That is the honest report and it is a stronger result than a counterfactual that survives, because a hypothesis that cannot be attacked from the most favourable available direction has not been tested from that direction at all.

7.2 R1 — the axis mismatch, and H2

The constructed structure's exposure table was rebuilt on the currency and settlement axes. **These are constructed inputs to a constructed case and not measurements of any real holding; the source manuscript supplies no as-of date, and none is asserted here** (Section 11).

```
entity-based United States exposure                 27.3%
USD-denominated share of the balance sheet          61.0%
settlement volume clearing through the US system     ~58%

ratio, currency axis to entity axis     61.0 / 27.3 = 2.234
ratio, settlement axis to entity axis   58.0 / 27.3 = 2.125
```

The entity map said 27.3 per cent and the dependency was something closer to sixty. The exposure to a decision taken in one capital was more than double what the jurisdiction table

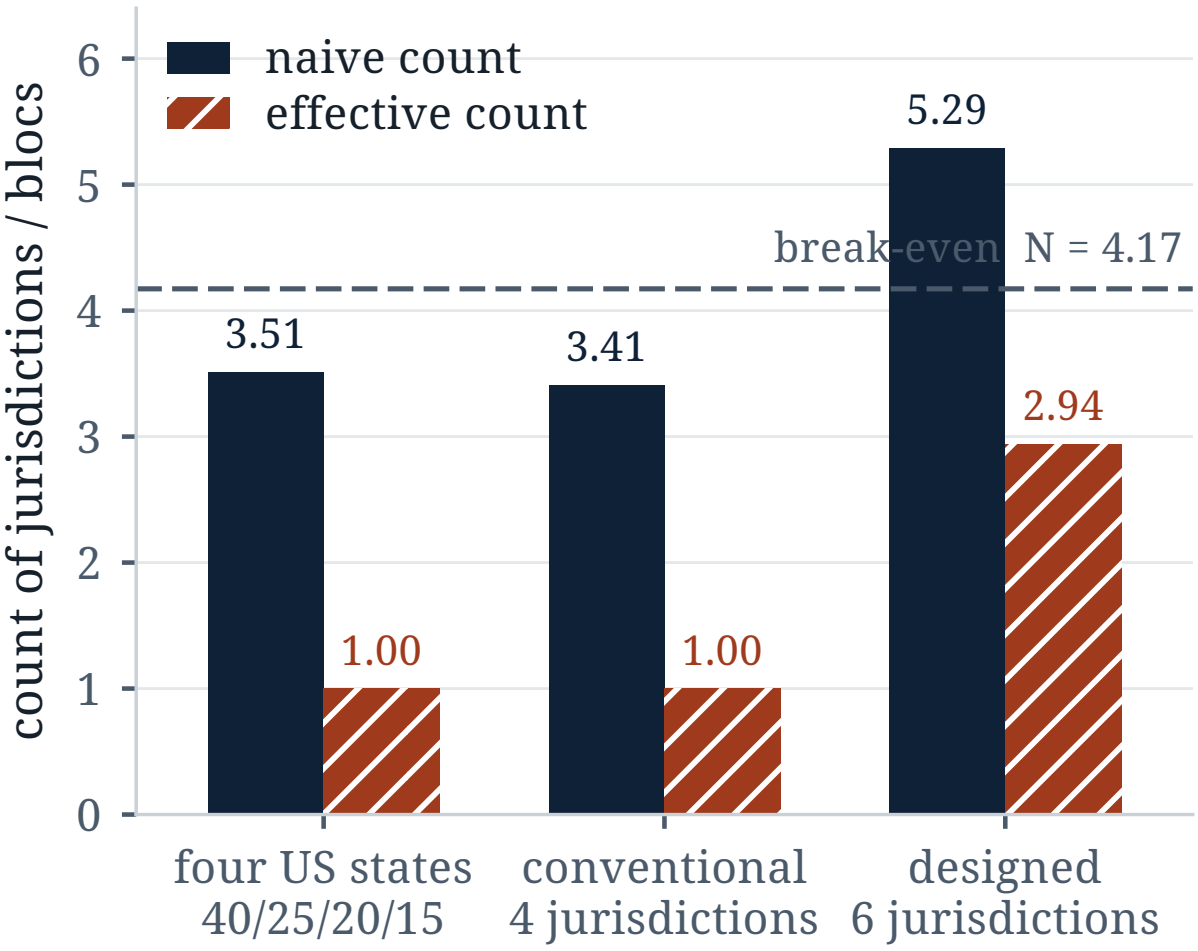


Figure 55.1: **Sub-national selection buys 3.51 naive jurisdictions and 1.00 effective ones.** Naive count against effective count for the four-state split and, for comparison, the conventional and designed structures of CCP-052, with the break-even effective count $N = e^{0.5/0.35} = 4.172734$ the surplus law requires at the stated parameters. All shares are constructed; the figure is generated by `figs/src/CCP-055_fig02.py` and every plotted value appears in Section 7.1 or in CCP-052.

reported, and every additional entity in the group had been assessed against the smaller number. Two errors produced it, and both are of a kind a reader can check on their own file. The first is that a portfolio was described as internationally diversified because its issuers were, when the correct basis is the perimeter of its custodian. The second is that a currency was treated as a pricing convention rather than as an exposure.

The rebuild that followed moved custody outside the perimeter, collapsed an interposed company, reduced the dollar balance toward the operating requirement, and **retained a single operating entity inside the perimeter deliberately**, booked to the Western bloc without argument, because the group sells into that market and contracts with its counterparties. The residual component fell to 2.0 per cent of exposure inside a Western bloc of 30.0 per cent, and the full structure measured

```
UAE 29.1 · Thailand 18.2 · UK 16.4 · Bermuda 13.6 · India 11.8 · Singapore 10.9
naive HHI = 0.189002    effective jurisdictions = 5.290949
bloc-HHI  = 0.340000    effective blocs N = 2.941176
ln(N) = 1.078810        eta ln(N) = 0.377583      h = -0.122417
```

which is the measurement Curzi [2026] reports and this paper derives from.

The currency reduction did more for the measurement than the entity change. Moving custody out reduced entity exposure; reducing dollar balances reduced the exposure the entity map never showed. Two further features of the rebuild are worth recording because they are what a resilience argument usually gets wrong. **An entity was retained inside the perimeter on purpose**, sized small, multi-member, holding no reserve and no group intangibles, because a structure that removes the entity it actually needs has optimised the measurement rather than the position. And **nothing in the rebuild was about privacy**: the retained entity's formation record was private before and is private afterwards, and what changed is where a decision could reach rather than who could see.

H2 is supported on the constructed case, and the support is weaker than the number looks. The ratios are computed from the same three inputs the hypothesis names, so the test is a consistency check on a construction rather than an independent confirmation that the gap generalises. Section 7.9 records this and Section 11 states it again.

7.3 R3 — the realisability measurement, and H3

The constructed year-seven statutory balance sheet, in USD:

line	amount
cash and admitted investments	14,200,000
reinsurance recoverable on paid and unpaid losses	1,600,000
premiums receivable	300,000
total assets	**16,100,000**
loss and loss adjustment expense reserves, gross	8,900,000
unearned premium reserve	1,500,000
other liabilities and accruals	400,000
total liabilities	**10,800,000**
total statutory capital and surplus	**5,300,000**
enhanced capital requirement	3,400,000
target capital level, 120 per cent of the requirement	4,080,000
surplus above target capital level	1,220,000

Applying the definition of Section 3.5:

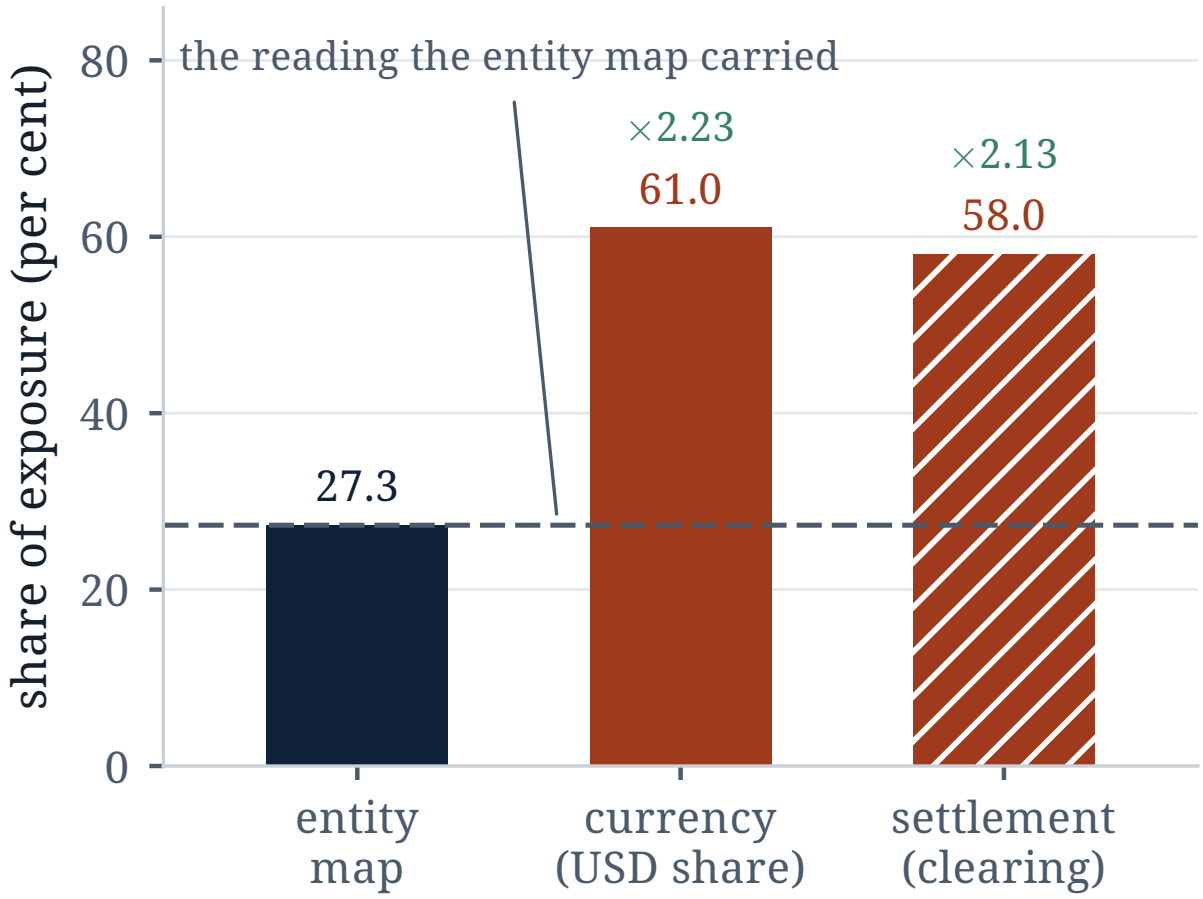


Figure 55.2: **The entity axis reads 27.3 per cent where the currency axis reads 61.0 and the settlement axis 58.0.** The three axes of Section 7.2 on one share basis, with the entity reading marked as the number the file carried. All values are constructed inputs to a constructed case, not measurements of a real holding; the figure is generated by `figs/src/CCP-055_fig01.py`.

```
committed to third-party liabilities   10,800,000 / 16,100,000 = 67.08%
locked by the target capital level      4,080,000 / 16,100,000 = 25.34%
distributable within one year           1,220,000 / 16,100,000 =  7.58%
                                                                -------
                                                                100.00%
```

Sixty-seven per cent of the balance sheet belongs, in substance and in law, to somebody other than the shareholder; a further twenty-five per cent cannot leave without the supervisor's consent; and the amount realisable inside a year is 7.58 per cent of the assets.

The present-value half, which a report stopping at the table above would omit. A creditor taking the shares owns the run-off. Over roughly six years the reserves develop and settle; at an assumed eight per cent redundancy on a conservatively reserved book — an assumption, and equally plausibly the reserves prove deficient — the eventual value is:

```
total statutory capital and surplus                      5,300,000
reserve redundancy at 8% of 8,900,000                      712,000
-------------------------------------------------------------------
eventual realisable value                                6,012,000
discounted at 8% over an average 4-year lag
  1.08^4 = 1.36048896
  6,012,000 / 1.36048896                                 4,418,999
-------------------------------------------------------------------
as a share of total assets                                  27.45%
realisable within one year                                   7.58%
haircut against 100% of assets                              72.55%
```

H3 is supported: the criterion fixed in Section 2 predicted a figure below ten per cent and the measurement returns 7.58 per cent. Against an adversary seeking rapid, cheap, uncontested realisation — a description that fits most arbitrary administrative action and much opportunistic litigation — a 72.55 per cent haircut and a six-year timetable is a substantial deterrent. Against a well-resourced judgment creditor with time it is a cost of collection. **The structural characteristic is worth a discount from 100 per cent of assets to roughly 27 per cent on a present-value basis, and a delay of years on all but 7.58 per cent of it. That is a discount, not a wall, and describing it as a wall is failure mode 7 of Section 8.**

7.4 The counter-result, in the same arithmetic

The counter-result is not an objection to R3; it is R3 read from the other side, and the paper prints it in the same units.

```
to a creditor:  realisable within a year        7.58% of assets
to the owner:   realisable within a year        7.58% of assets
```

The three quantities that bound what a creditor can take are the same three quantities that bound what the owner can take. Third-party liabilities are owed whoever holds the shares. The target capital level is a supervisory floor that does not move for a friendly shareholder. The distributable surplus is distributable to either. **The property that makes an insurance reserve resistant to third-party claims — real liabilities owed to unrelated parties — is the same property that makes it unavailable to its owner, leaving only 7.58 per cent realisable within a year.**

Two consequences follow and neither is optional.

A reserve is not a treasury and must never be sized as one. Capital committed to this vehicle is committed, and this framework's carrying-cost result applies without softening: Curzi [2026] states the four-currency carrying cost — coin, hours, attention and foreclosure

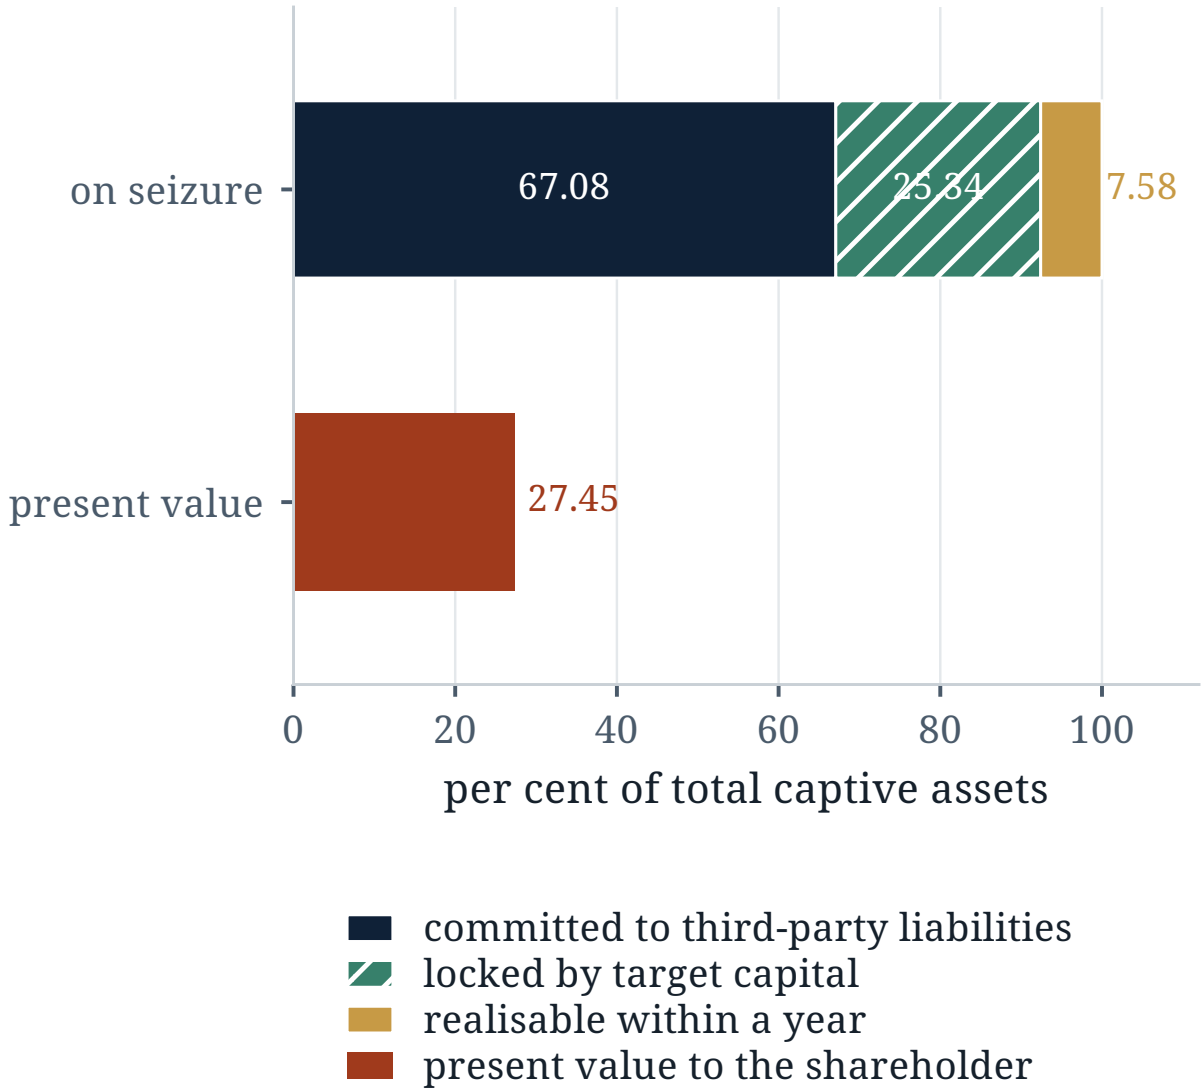


Figure 55.3: **Sixty-seven per cent of a captive's assets are committed, twenty-five per cent locked, and 7.58 per cent realisable within a year.** The realisability decomposition of the constructed year-seven balance sheet, with the present-value reading of 27.45 per cent below it. Modelled with the stated parameters — an assumed 8 per cent reserve redundancy, an 8 per cent discount rate and a four-year average settlement lag, all stipulated in Section 3.5 rather than measured. Generated by `figs/src/CCP-055_fig03.py`.

— against the law that felt comfort decays over weeks while carrying cost does not decay at all. **Foreclosure is the currency this result is denominated in**, and it is the one a balance sheet does not show. The parallel is exact rather than decorative: the cost is paid continuously, the benefit decays, and the arrangement is priced by comparing a flow that persists with a stock that does not.

The argument cannot be sold on both sides at once. An arrangement described to a principal as protecting assets *and* as leaving them available has been described inconsistently, and the inconsistency is arithmetic rather than rhetorical. Curzi [2026] states the general shape of the trap: perfect opposition and perfect agreement are the same reading, because u depends only on $|\rho|$, so the standard remedy of adding a contrarian contributes zero surplus while adding cost, and at $u = 0$ two of the four levers are exactly inert. **A lever that is inert in one direction is inert in both, and a property that binds a creditor binds an owner.**

The operational form of the counter-result is a specification rather than a warning. **Size the vehicle to the retained exposure and to nothing else.** Capital placed in it above what the retained layer requires does not become better protected; it becomes committed, and the correct comparison for the marginal dollar is between 7.58 per cent availability inside the vehicle and full availability outside it.

7.5 R5 and R4 — the programme economics and the rubric

The retained programme, and H4. The constructed group previously paid USD 4,370,000 a year in commercial premium against a five-year average loss of USD 1,910,000, a gap of USD 2,460,000 comprising expense load, risk load, profit margin and cost of capital. **Part of that gap is a fair price for capacity the group could not otherwise obtain; the question the programme answers is how much of it is a fair price for risk the group is in fact carrying.** After restructuring it pays USD 1,670,000 to the commercial market for the catastrophe layers and USD 2,700,000 to its own vehicle for the working layer, which sums to the prior 4,370,000 exactly.

```
gross premium written to the captive                    2,700,000
less excess of loss and aggregate stop-loss ceded        (410,000)
------------------------------------------------------------------
net premium                                             2,290,000
expected net loss (actuarial, constructed)              1,560,000
net loss ratio  1,560,000 / 2,290,000                      68.12%
running cost                                              140,000
------------------------------------------------------------------
2,700,000 - 410,000 - 1,560,000 - 140,000  =               590,000
against the prior commercial spend
  590,000 / 4,370,000                                      13.50%
```

H4 is supported. A captive on these facts is a thirteen-and-a-half-per-cent-of-premium proposition before investment income and before any tax consequence in any jurisdiction, and any projection materially above it is either retaining more risk than it discloses or not paying claims. The margin is positive but not comfortable: USD 140,000 of annual cost against USD 590,000 of expected retained margin, and a bad loss year makes it negative. **The catastrophe layers stay in the commercial market for a reason the arithmetic does not show** — a single large loss can exceed the group's capacity to absorb it — and a programme that retains them to improve the ratio has changed the risk rather than the economics.

The rubric applied to three constructed arrangements. Weights 25 / 25 / 20 / 20 / 10; scores 0–4; weighted total to 400.

test	weight	Case A	Case B	Case C
1 risk shifting	25	0	4	4

test	weight	Case A	Case B	Case C
2 risk distribution	25	0	4	3
3 actuarial pricing	20	1	4	4
4 claims paid	20	0	4	4
5 insurer conduct	10	0	3	4
weighted total		**20 / 400**	**390 / 400**	**375 / 400**
verdict		**disqualified on tests 1 and 2**	pass	pass

Case A is a constructed arrangement assembled from the fact patterns recurring in the reported decisions the source describes: three insureds in one location, a pooling arrangement with several hundred participants, a majority of assets lent to a related party, and effectively no claims. Four of its features are arithmetic and are printed as such:

```
premium as a fraction of the ceiling in force
  1,180,000 / 1,200,000                              98.33%
fraction of ceded premium NOT returned by the pool
  1 − 0.992                                           0.80%
five-year paid loss ratio
  14,000 / (5 x 1,180,000)                            0.24%
share of assets lent to a related party              62.00%
```

A premium set at 98.33 per cent of an external ceiling was not derived from an exposure; a pool returning 99.2 per cent of what it takes moves 0.80 per cent of the risk; and a five-year paid ratio of 0.24 per cent against a pricing assumption in the high sixties is not a prudent book. Case A's total of 20 out of 400 is not a low pass; the disqualification rule of Section 6.2 disposes of it regardless of the number, and the consequence of the arrangement not being insurance is definitional rather than penal — every deduction taken since inception falls with the characterisation.

Case B is a constructed group captive of nine unrelated contractors, none above 16 per cent of premium, priced by an independent actuary from ten years of member loss data, with a five-year paid ratio of 61 per cent and incurred of 68 per cent against a pricing assumption of 66 per cent, and an underwriting committee that surcharged two members and declined one renewal in five years. **The most persuasive fact in Case B is not in the rubric at all: the captive declined to renew one of its own owners, which is something no wallet does and every insurer does.**

Case C is the constructed captive of Section 7.3, scored on the same rubric. **Its three-point score on distribution is the honest weak point and is stated rather than rounded up**: eleven operating units in eight countries plus two unrelated distributors, correlated through a single manufacturing process, so that a design defect is one event and not eleven. **Distribution is strong against frequency and weak against a systemic defect, which is exactly the exposure the reinsurance layer above the vehicle exists to carry.** Section 7.6 measures it.

The disclosure test applied to the three. Case A fails at the first element: its function depends on nobody reconstructing how the premium was set, and disclosure of the pool mechanics, the related-party loan and the pricing chronology ends it. **It is a concealment structure with a policy stapled to the front.** Case B passes comfortably — every element is designed to be shown, and disclosure is how a group captive recruits its next member. Case C passes, and its weakest element is disclosed in its own reinsurance submission because the reinsurers price it. **The test separates the three in the same order the rubric does, by a different route, which is the cross-check Section 10.9 records.**

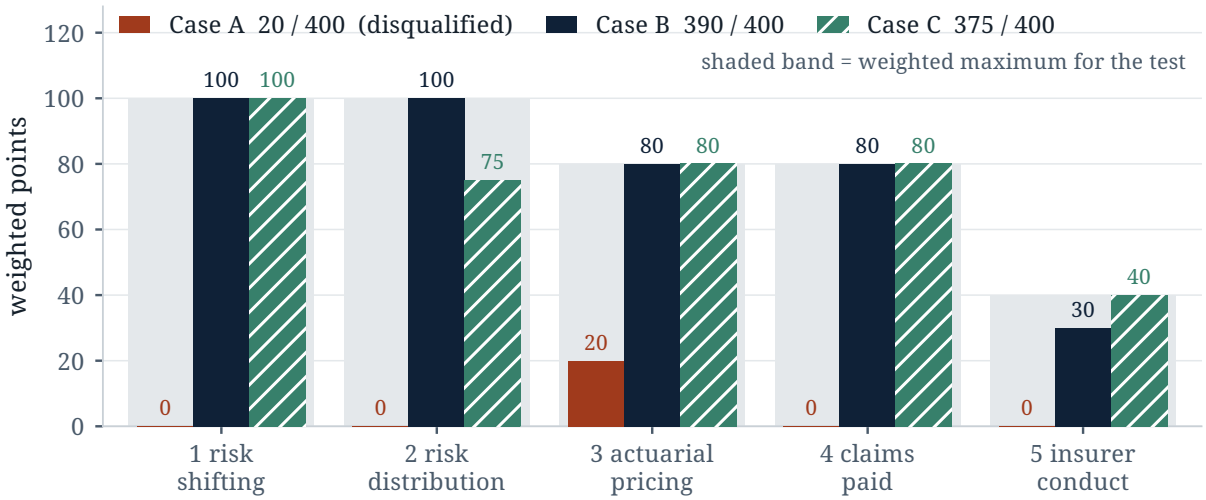


Figure 55.4: **The rubric separates the three arrangements on tests 1 and 2, not on the total.** Weighted points per test for the three constructed arrangements against the weighted maximum for each test. Case A's 20 out of 400 is disqualified by the hard-zero rule rather than by its total. The rubric is this framework's working instrument, adopted by no tribunal; all scores are constructed. Generated by `figs/src/CCP-055_fig04.py`.

7.6 H5 — the risk-distribution test is the bloc collapse

This is the result that makes the two halves of the paper one paper, and it is arithmetic throughout.

The collapse, applied to insureds. Case A lists three insureds. They sit in one state, behind one warehouse, with one customer concentration. If those three exposures are perfectly correlated then, whatever their shares,

```
naive count of insureds                    3.000000
bloc-collapsed count                       1.000000
```

because a block matrix of ones returns $\sum_b W_b^2 = 1$ for a single block regardless of the within-block split. **No share vector is needed and none is assumed.** Case C lists thirteen insured units. Against ordinary frequency risk they are thirteen; against a systemic defect in the single manufacturing process behind eleven of them, and behind the product the other two distribute, they are one. **The naive count and the effective count differ by the same factor and for the same reason as they do on the jurisdiction axis, and the operation that reconciles them is the same operation.**

The floor, derived from a published constraint. The safe harbours described in Section 6.1 are constraints on shares, and a constraint on shares is a bound on the reciprocal index. Both bounds below are exact and both are attained.

Twelve subsidiaries each contributing between five and fifteen per cent of total risk. The sum of squares is maximised by pushing as many shares as possible to the cap. With k subsidiaries at 0.15 and $12-k$ at 0.05, $0.15k + 0.05(12-k) = 1$ gives $k = 4$, so

```
max HHI = 4(0.15)^2 + 8(0.05)^2 = 0.090000 + 0.020000 = 0.110000
min effective count = 1 / 0.110000 = 9.090909
min HHI = 1/12 = 0.083333    max effective count = 12.000000
```

Nine unrelated members, none above sixteen per cent of premium. Six members at the cap absorb 0.96, leaving 0.04:

```
sup HHI = 6(0.16)^2 + (0.04)^2 = 0.153600 + 0.001600 = 0.155200
min effective count = 1 / 0.155200 = 6.443299
max effective count = 9.000000
```

The published constraint is therefore equivalent to an effective-count floor: the twelve-subsidiary harbour guarantees at least 9.090909 independent-equivalent units and the nine-member harbour at least 6.443299, **provided the units are in fact independent** — which is the condition the constraint does not impose and which the collapse above destroys. **H5 is supported.** The share constraint and the effective count are the same statement, and the constraint's silence about correlation is the same silence the naive jurisdictional index keeps.

The consequence for the rubric is a correction to it, stated here rather than suppressed. Test 2 is scored on a hard zero at "fewer than five units", which is a naive count. **A naive count of units is exactly the quantity this series exists to say is the wrong one.** The rubric as constructed can therefore be satisfied by an arrangement with twelve perfectly correlated units, whose effective count is 1.000000 and whose aggregate variance is not reduced at all. **The correction is to state the hard zero on the effective count rather than the naive one**, which is what Case C's three-point score already does informally and what the rubric does not do formally. Section 11 records it as a defect in an instrument this paper carries rather than repairing it mid-paper, because a threshold changed to fit a case it was applied to is not a threshold.

7.7 R6 — pricing a bloc call rather than a structure

Bermuda, which hosts the constructed insurer, sits inside the Western perimeter on every operative test of Section 4.1 as applied in this series: it exchanges account information, it implements the common anti-money-laundering standard through domestic legislation, it is exposed to the same listing processes, and its banks clear the same currency through the same correspondents. Its final appeal lies to a committee sitting in the perimeter's centre, and the perimeter's centre retains a constitutional capacity to legislate for it. The base case therefore folds it into the Western bloc and credits it with nothing. **The argument that it should be credited separately is an argument about the legal character of an asset, not about a perimeter**, and this paper prices it rather than arguing it.

Blocs on the designed structure, base case: Western 30.0 (the United Kingdom at 16.4 plus Bermuda at 13.6), bridge 40.0, non-aligned 30.0.

```
base case          bloc-HHI = 0.30^2 + 0.40^2 + 0.30^2       = 0.340000
                   N = 2.941176      ln N = 1.078810
                   eta ln N = 0.377583        h = -0.122417

credited separately
                   bloc-HHI = 0.164^2 + 0.136^2 + 0.40^2 + 0.30^2
                            = 0.026896 + 0.018496 + 0.160000 + 0.090000
                            = 0.295392
                   N = 3.385332      ln N = 1.219452
                   eta ln N = 0.426808        h = -0.073192
-------------------------------------------------------------------------
value of granting the character argument in full        0.049225
```

Granting the argument its most generous possible reading moves the structure from clearly below the break-even count of 4.172734 to slightly less clearly below it, and it does so by changing an accounting convention rather than a structure. The base case therefore declines to grant it. **This is a change in the bloc call and not a change in the world**, and Curzi [2026] is the paper that prices an uncertain bloc call properly as a bet rather than as a fact.

The mirror reading matters and is stated: if the bridge bloc were to split — if the two bridge jurisdictions ceased to correlate — the same arithmetic gives $N = 3.615826$, which is

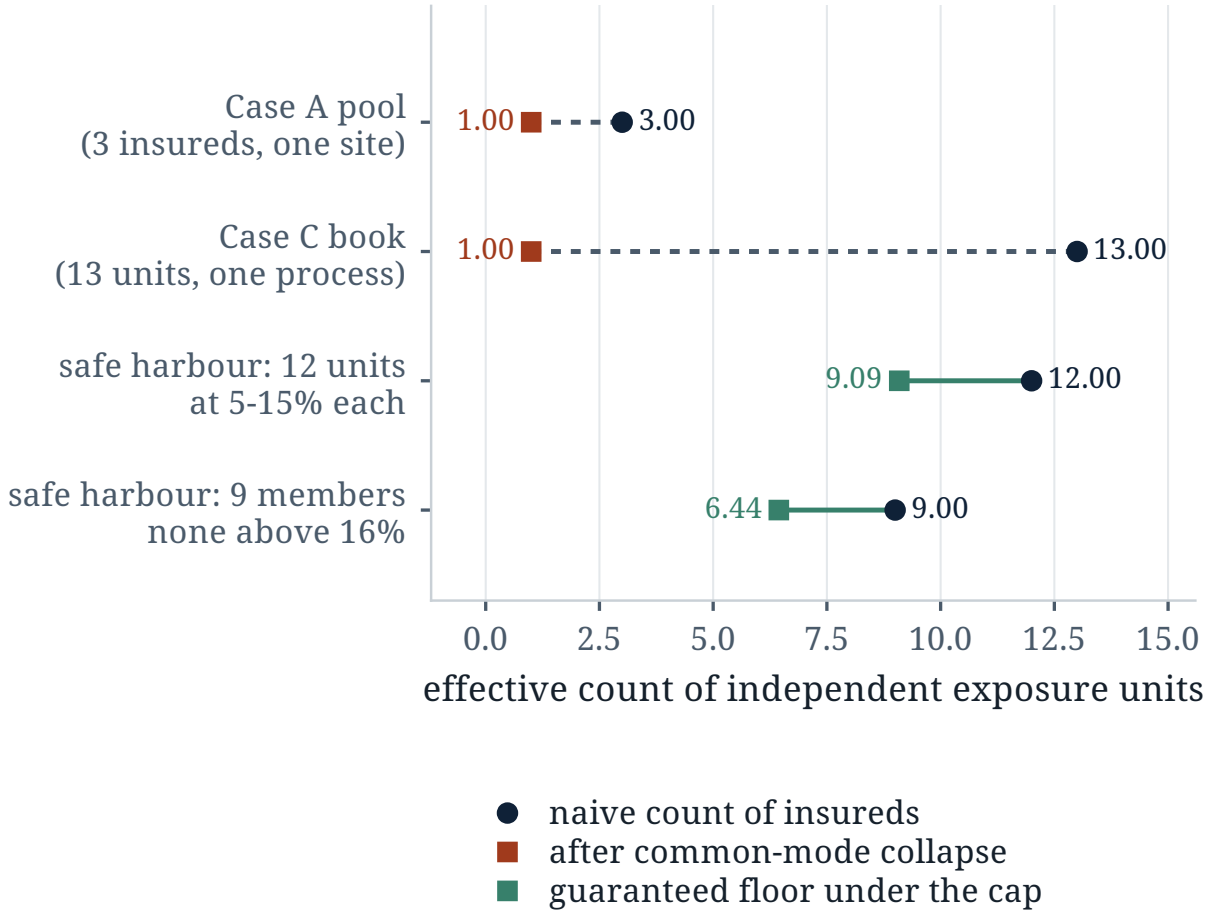


Figure 55.5: **A naive count of thirteen insured units collapses to one against a common mode, and a published share constraint is equivalent to an effective-count floor.** The naive count of insureds against the count surviving a common-mode collapse for the two constructed arrangements, with the exact effective-count intervals implied by the two share constraints. Bounds are computed in Section 7.6; the constraints themselves are described from the source manuscript and are uncited. Generated by `figs/src/CCP-055_fig05.py`.

the figure Curzi [2026] measures from the other direction, where Design A places a full centre at 10.9 per cent of exposure and reaches $N = 3.615826$ and $h = -0.050138$, while Design B's capped transit position at 0.35 per cent reaches $N = 2.965384$ and $h = -0.119548$ — capturing 3.97 per cent of A's measured gain for 3.21 per cent of A's exposure. **The two papers reach the same number by opposite routes: here it is what a fragmentation would deliver for free, and there it is what a deliberate addition has to be paid for.** Bloc fragmentation runs both ways and only one direction is usually planned for. The robustness reading is therefore that this measurement is sensitive to bloc *structure* and not to the placement of any single member, which is the same conclusion Curzi [2026] reaches on its own robustness table.

7.8 Three printed discrepancies, left visible

First, the surplus reading of the base-case structure. The source manuscript prints $h = -0.1226$ for the designed structure and $h = -0.0734$ for the character-credited variant. The exact values at $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ are -0.122417 and -0.073192. The discrepancy is 0.000183 in each case and it propagates consistently — the source's two figures differ by 0.0492, which is the correct increment — so it is a rounding of the base figure carried forward, not an error in the increment. **This paper uses -0.122417 and -0.073192, which are the values Curzi [2026] publishes, and prints the source's figures here rather than harmonising them silently.**

Second, the present-value arithmetic of Section 7.3. The source prints $6{,}012{,}000 / 1.360489 = 4{,}418{,}970$. The exact quotient is $4{,}418{,}999.48$, a difference of USD 29 arising from a truncated divisor. **The share of total assets is 27.45 per cent on either value**, so no reported result changes; the discrepancy is printed because a reader recomputing the chain will find it.

Third, the naive index of the conventional structure. This framework carries two values for the same measurement: 0.293388 from exact ninths — Curzi [2026] reports UK 40.91 / US 27.27 / Jersey 18.18 / Luxembourg 13.64 giving naive HHI 0.293388 and 3.4085 effective jurisdictions against bloc-HHI 1.000000 and 1.0000 effective perimeters — and 0.293430 from rounded percentages [Curzi, 2026], giving 3.407968. The difference is 0.000042, in the fifth decimal. **This paper uses the rounded-percentage figure,** 0.293430, because it is the basis on which the four-axis comparison of Section 7.2 was constructed, and states which it uses rather than leaving the reader to infer it. Both papers print the discrepancy rather than hiding it, and so does this one.

7.9 Hypotheses tested

#	hypothesis	verdict	where
H1	sub-national selection contributes nothing, and admits no counterfactual	**supported**, on both halves	7.1
H2	the entity axis understates exposure, by a factor of 2.23 on the currency axis	**supported on the constructed case**, and the point prediction is a consistency check rather than an independent test	7.2
H3	realisable-within-a-year is small and computable: 7.58 per cent against a predicted floor of ten	**supported**	7.3

#	hypothesis	verdict	where
H4	the retained programme recovers 13.50 per cent of prior spend	**supported**	7.5
H5	the risk-distribution test and the bloc collapse are one operation	**supported**, with an exact interval on each published constraint	7.6

Two of the five are weaker than they look and the paper says which. H2 is tested on one constructed structure and its ratios are computed from the same inputs it names, so it demonstrates the size of a gap under a construction rather than establishing that the gap generalises. H4's inputs — the expected loss, the ceded premium and the running cost — are all stipulated, so the 13.50 per cent is an exact consequence of an assumed programme rather than an observed return. **Neither is presented in Section 12 as more than it is.**

8. Failure modes

Ten ways the results above are misapplied. Each is stated with the diagnostic that detects it, because a failure mode without a diagnostic is a complaint.

1. Confusing formation privacy with protection. They are unrelated, and this is the clearest case in the series of a jurisdiction with a great deal of the first and very little of the second. *Diagnostic:* ask whether the structure's function would change if the formation record were public. If not, privacy was never the mechanism; if so, the structure fails the disclosure test of Section 4.4 by construction.

2. Treating a seasoning period as a protection period. A shorter state period does not displace a longer federal avoidance provision. *Diagnostic:* the two periods are not alternatives and the longer one governs where its conditions are met. **The gap between how these instruments are marketed and how they work is the largest single gap in Section 4.**

3. Assuming the chosen state's law will be applied. A creditor sues where it can obtain jurisdiction, and that court applies its own conflict rules and its own public policy. *Diagnostic:* name the forum before naming the statute.

4. Using a single-member vehicle for charging-order protection. The doctrinal basis is protecting other members; with no other members the basis is absent. *Diagnostic:* count the members with genuine economic interests, not the members on the register.

5. Reading the naive count as the resilience count. The error this whole series exists to correct. *Diagnostic:* collapse the shares onto perimeters before squaring, and report the pair, never one alone.

6. Measuring one axis and calling it the structure. The most expensive error in the constructed case, worth a factor of 2.23. *Diagnostic:* run all four axes of Section 3.2 and print the largest, which is item 13 of Appendix A.

7. Letting a good custodian carry the reserve. Depth, cost and quality are exactly why reserve drifts into a perimeter. *Diagnostic:* a portfolio contributes to the perimeter of its custodian, not of its issuers, and should be booked that way whatever the issuer list looks like.

8. Describing a discount as a wall. The realisability measurement is 72.55 per cent and not 100 per cent. *Diagnostic:* if the description contains no number, it is not the measurement.

9. Selling the reserve argument on both sides. Section 7.4. *Diagnostic:* ask what the owner can realise within a year and compare it with what a creditor can. They are the same number.

10. Satisfying a unit count without satisfying independence, and crediting a bloc call as a structural change. Two errors with one diagnostic: compute the *effective* count, of insured

units in the first case and of perimeters in the second, and ask what would have to change in the world rather than in the table for the number to move. Section 7.7 prices the second at 0.049225 of surplus, which is real and is an accounting convention.

9. Discussion

What follows from the results. Three things, and each is narrower than a reader in a hurry might take it to be.

The first is that **an entity map is a poor instrument for the question it is usually asked**. It answers where entities are, which is a question about registration, and it is routinely read as answering where a decision reaches, which is a question about perimeters, currencies, rails and custodians. The gap on the constructed case is a factor of 2.23 and there is no structural reason to expect the sign to reverse, because a currency's clearing runs through banks inside a perimeter whatever the entity list says. **The remedy is not a better entity map. It is three more measurements**, and Appendix A states them.

The second is that **a protective characteristic that can be measured should be, and measuring it usually shrinks it**. The reserve argument survives measurement — 72.55 per cent is a substantial haircut — but it survives as a discount and a delay rather than as the immunity it is often described as. Curzi [2026] reaches the same shape by a different route: a conversion of an immovable asset carries certain leakage $L = 7{,}220{,}000 / 48{,}000{,}000 = 15.0417$ per cent against an escape value $V_{\text{esc}} = p \cdot L^{*} \cdot q \cdot (1 - s) = 0.4725$ per cent a year, a payback of 31.83 years, with a ten-year payback requiring $p = 4.7752$ per cent. **A defensive property with a payback period is a defensive property that can be priced, and one that cannot be priced has not been specified.**

The third is that **the same instrument, run on a different share basis, answers a legal question it was not built for**. The risk-distribution test asks how many independent exposures a pool has. That is the effective count, and the published share constraints that operationalise it are equivalent to effective-count floors of 9.090909 and 6.443299. This is not a claim that the doctrine intends the reciprocal index; it is the observation that the doctrine's arithmetic content *is* the reciprocal index, and that the doctrine's silence about correlation is the same silence the naive jurisdictional count keeps. **A test that counts units and does not weight them by correlation will be satisfied by twelve correlated units, and the remedy is the one this series has been applying to jurisdictions since Curzi [2026].**

What does not follow. Several things, and they matter more than the above.

Nothing here is a legal conclusion. Every jurisdictional, doctrinal and regulatory statement in Sections 4 to 6 is described from the source manuscript and is uncited; Appendix B lists them; the Declarations state it again. **A reader who takes a rule of law from this paper has taken it from the wrong document.**

The effective count measures independence, not capacity. It says nothing about whether any bloc, any custodian or any insurer can carry the load, and a structure at $N = 2.94$ on jurisdiction with an effective counterparty count near 1.6 has moved its risk rather than reduced it.

The measurement is least informative exactly where the headline case sits. Golosov [2010] and Dunleavy and Boucek [2003] document that effective-number measures behave poorly in the dominated-unit regime, where one unit holds a very large share. A bloc vector of a single unit at 100 per cent is the limiting case. The measurement is still exact — $1/1 = 1$ — but its discriminating power is at its minimum there, and comparisons among heavily dominated structures should not be made on N alone.

A count that can be inflated will be. Goodhart [1975] and Campbell [1979] are the standing warnings, and they bite hardest on an instrument that has just been shown to be gameable in the naive direction. **The bloc collapse is the anti-gaming device and it is only as good as the assignment**, which is a modelling position (Section 1) and not an identified result. Douceur [2002] is the formal statement that one party can present as many, so an unaudited count is

worthless, and Lamport et al. [1982] is the bound on how many genuinely independent parties a protocol requires. **Both belong next to any count of units in this series**, and the rubric defect of Section 7.6 is exactly the Sybil problem in an insurance costume.

The results this paper borrows from bound it. Curzi [2026] records that the frequency register cannot distinguish a mood from a trait at all, because $f = kn^2$ contains no time term, and that the cleanest-looking encoding of a framing — the one-hot vector — returns the system's maximum surplus for every subject and therefore measures nothing. Curzi [2026] records that the width statistic is invariant under relabelling and therefore does not even determine the chord it is computed from, which formally blocks any inference from low or high internal surplus to distress or health. **Both are the same warning in a different domain: an index invariant under a transformation the analyst cares about is not measuring what the analyst thinks it is.** The rubric of Section 6 is exposed to exactly that objection through its naive unit count, and Section 7.6 states the correction rather than waiting to be told.

This framework's central limitation applies here in full. Curzi [2026i] reports that a preregistered fit failed its own second falsification rule — a 2.8 per cent held-out RMSE advantage over a matched AR(1) against a declared 10 per cent margin, despite +40.3 per cent over persistence, +97.7 per cent over direct input, and $p = 0.0005$ against the shuffle — and Curzi [2026k] reports the five-parameter survey reproducing each system's observed decline from one desk to a 22-household block, predicted -0.250474 against observed -0.250 at the desk and -0.996371 against -1.000 at the library, with residuals under 1.4 per cent of the ceiling term. **This paper reports no fit at all, and that is not a strength.** Epstein [2008] is the standing statement that prediction is one of sixteen reasons to model and that a paper is entitled to say which of the other fifteen it claims; this paper claims **explanation of a mechanism and demonstration that a stated arithmetic suffices to produce a stated shape**, and claims neither prediction nor calibration.

10. Verification and rigour audit

The results are decomposed into postulates that can be checked one at a time. Each is atomic, each carries a status, and the compounding chain names what each step consumes. **If dropping a postulate leaves the result standing, it is not listed.** Sections 10.1 to 10.6 run the standing rigour checks first, so the chain in 10.8 is read against a dimensional audit rather than before one.

10.1 Dimensional audit

symbol	quantity	dimension	unit	status
s_i, w_i	exposure share of unit i	$[1]$	dimensionless, fraction of Q	defined; non-negative and summing to 1 by P1
W_b	bloc or block share	$[1]$	dimensionless	derived, $W_b = \sum_{i \in b} s_i$
HHI	concentration index	$[1]$	dimensionless, or points $= 10^4 \times$ fraction	derived from shares
ρ_{ij}	correlation between shocks to i and j	$[1]$	dimensionless, in $[-1, 1]$	**assumed, not measured**

symbol	quantity	dimension	unit	status
N	effective count	$[\#]$, a count	jurisdictions, blocs, or exposure units	derived as $1/(w'Rw)$
$\ln N$	surplus-term argument	$[1]$	nats	derived; the logarithm of a count is dimensionless
Q	protected capital	$[C]$	index units	defined; **the source supplies no operational unit**
η	conversion efficiency	$[1]$	dimensionless	assumed
δ	erosion rate	$[T^{-1}]$	per period	assumed
C	maintenance cost flow	$[C\,T^{-1}]$	index units per period	assumed
h	headroom	$[C\,T^{-1}]$	index units per period	derived, $h = \eta \ln N - \delta Q - C$
A	total captive assets	[\$]	USD	constructed
L	third-party liabilities	[\$]	USD	constructed
TCL	target capital level	[\$]	USD	derived, $1.20\times$ the enhanced requirement
realisability	share of assets realisable	$[1]$	dimensionless	derived from the three above
r	discount rate	$[T^{-1}]$	per year	**assumed**
t	average settlement lag	$[T]$	years	**assumed**
rubric score	ordinal band	none	**no unit; ordinal**	stipulated

The surplus equation balances only if $\eta \ln N$ carries $[C\,T^{-1}]$. Since $\ln N$ is dimensionless, η must carry $[C\,T^{-1}]$ rather than being dimensionless, or $\ln N$ must be read as a rate. **The source is silent on which, and this paper records the silence rather than resolving it by assertion.**

The realisability equation balances exactly. Three quantities in USD divided by a quantity in USD gives a pure number, the three shares sum to 1.000000 by construction, and the discount factor $(1+r)^t$ is dimensionless because r carries $[T^{-1}]$ and t carries $[T]$. **That is the only equation in this paper whose dimensions are fully determined by the source, and it is worth saying which one it is.**

The rubric has no dimension and is not given one. Weighting and summing ordinal bands produces a number with no unit and no interval interpretation; Section 10.6 states what that permits.

10.2 Significant figures and epistemic status

quantity	figures quoted	status	licence
naive HHI 0.285000	6 dp	**exact**	exact arithmetic on exact constructed shares
effective count 3.508772	6 dp	**exact to rounding**	$1/0.285$ is a repeating decimal, printed to 6 dp
bloc-HHI 1.000000, $N = 1.000000$	6 dp	**exact**	a single block of ones
67.08 / 25.34 / 7.58 per cent	2 dp	**exact to rounding**	ratios of exact constructed amounts
27.45 per cent	2 dp	**model output**	depends on three assumed parameters; see 10.3
13.50 per cent	2 dp	**exact to rounding**	ratio of exact constructed amounts
$N = 3.385332$, $h = -0.073192$	6 dp	**exact to rounding**	exact arithmetic on exact shares
9.090909 and 6.443299	6 dp	**exact**	closed-form bounds, Section 7.6
61.0 and 58 per cent	1 dp, 0 dp	**constructed input**	not a measurement; no as-of date exists
rubric totals 20 / 390 / 375	integers	**ordinal, not quantitative**	see 10.6

The 27.45 per cent figure is the only headline number in this paper that is model output rather than arithmetic on stated inputs, and it is the one most likely to be quoted without its parameters. Its three assumptions are an 8 per cent reserve redundancy, an 8 per cent discount rate and a four-year average lag. **None was measured, all three were chosen, and Sections 10.3 and 11 give the sensitivity.** The 58 per cent settlement figure is quoted in the source with a tilde and is reproduced here as 58.0 in the figure only; the one-significant-figure reading is what the source supports.

10.3 Limits and boundary behaviour

limit	behaviour	reading
$\rho \to 0$ within a bloc	$N \to$ the naive count	the naive index is the zero-correlation endpoint, not a different index
$\rho \to 1$ within a bloc	$N \to$ the bloc count	the bloc index is the unit-correlation endpoint
one bloc holding everything	$N = 1$ exactly, $\ln N = 0$, surplus term $= 0$	**the whole argument in one line**; no activity below bloc level raises it
$N \to \infty$	$h \to \infty$ logarithmically; $dh/dN = \eta/N \to 0$	returns diminish but never turn negative — **a defect of the model, since cost is not in the input term**

limit	behaviour	reading
$L \to 0$ (no third-party liabilities)	realisable-within-a-year $\to 1 - \text{TCL}/A$	leg three of Section 5.1 is gone and only the supervisory gate remains
$L \to A$	realisable $\to$ negative, i.e. the insurer is insolvent	the domain of the measurement is $L + \text{TCL} \leq A$ and is stated
$t \to 0$	present-value realisation $\to$ 37.34 per cent	the delay, not the discount, carries most of the haircut
$r \to 0$	present-value realisation $\to$ 37.34 per cent	identical to the row above; r and t enter only through $(1+r)^t$
reserve redundancy $\to 0$	present-value realisation $\to$ 24.20 per cent	the redundancy assumption is worth 3.25 percentage points
short-tail book, reserves $\to$ one year	the committed share falls toward the unearned premium alone	the characteristic is a function of reserve duration (Section 5.1)
all shares equal, n units	$N = n$ exactly	the numbers-equivalent property, which is the instrument's definition

Three boundary readings are load-bearing. The fourth row is a genuine defect of the surplus law as used here and not merely a curiosity: because cost enters only through C, the model never says that adding a bloc is not worth it, and the decision to stop must come from outside the equation. **This framework's own cost result supplies it** — Brooks [1975] on communication cost at order $n(n-1)/2$ inside a boundary, and Curzi [2026], which records that a commons case is formally tunable and practically unreachable, passing the ceiling test with margin $+0.058436$ nats while no lever closes it, and that a member's declining to participate improves the family reading by 2.680588 octaves. **That arithmetic is printed rather than avoided, and it is the reason no result in this paper is stated as a prescription to raise N.**

The seventh and eighth rows say something the headline number hides: at $r = 0$ or $t = 0$ the present-value realisation is $6{,}012{,}000/16{,}100{,}000 = 37.34$ per cent, so **the discounting accounts for 9.89 percentage points of the 72.55 per cent haircut and the balance-sheet structure accounts for the rest.** The result is therefore mostly insensitive to the two assumed financial parameters, which is worth knowing and is not obvious from the table in Section 7.3.

The tenth row is the domain statement for R3 and it is not a technicality. **The realisability result is a result about long-tail books**, and quoting 7.58 per cent for an arrangement writing only short-tail risk would be a misapplication of the measurement rather than a disagreement with it.

10.4 Invariance

Scale invariance in capital. Multiplying every amount on the balance sheet by a positive constant leaves every share, every ratio and every effective count unchanged. The realisability measurement is a ratio of ratios and is invariant under the choice of currency unit. **Tested at three scales by construction** — the four-state shares, the six-jurisdiction shares, and the balance sheet — and it holds exactly in each. **Three scales is not many, and Section 11 says so** rather than describing the instrument as scale-free without qualification.

Invariance under relabelling, and its cost. The effective count is invariant under permutation of the units. That is the property Curzi [2026] identifies as fatal in its own domain: the width statistic is invariant under relabelling and therefore does not determine the chord it

is computed from. **Here the invariance is wanted — a concentration measure that changed when the rows were reordered would be broken — but the same invariance means the effective count cannot distinguish two structures with the same share vector and different contents.** That is a limit on what the number can carry and it is stated as one.

Not invariant, and this is where the instrument does its work. The count is *not* invariant under partition of the units: splitting one unit into two raises the naive count and leaves the bloc count untouched. **That asymmetry is the whole instrument**, and it is what makes the naive count gameable and the bloc count not.

Invariance of the rubric under its weight vector: not tested. Reweighting the five tests would reorder the totals, and no sensitivity analysis over weights was run. Section 11 records it.

10.5 Analogy discipline

borrowing	category	justification
exposure as a portfolio; variance under a correlation matrix [Markowitz, 1952]	**1 — derivation**	the exposure vector genuinely is a weighted sum of shocks with a correlation matrix, which is the premise the algebra requires; no further financial content is imported
effective count as a numbers-equivalent [Adelman, 1969, Hill, 1973, Jost, 2006]	**1 — derivation**	$1/\text{HHI}$ *is* the Hill number of order two; nothing transfers by analogy
similarity-sensitive diversity [Leinster and Cobbold, 2012, Rao, 1982]	**1 — derivation**	the block-of-ones matrix is a special case of the cited general form
risk distribution as an effective count (Section 7.6)	**1 — derivation**	the doctrinal test is a statement about independent exposures and aggregate variance, which is the same premise the portfolio result requires
open design [Kerckhoffs, 1883, Saltzer and Schroeder, 1975]	**1 — derivation**	the disclosure standard *is* the open-design requirement applied to legal rather than cryptographic structure
coupled-network fragility [Buldyrev et al., 2010]	**2 — shared form**	perimeter coupling is not shown to satisfy the percolation premises; only the shape of the result is used

borrowing	category	justification
robust-yet-fragile [Albert et al., 2000, Callaway et al., 2000]	**2 — shared form**	the tolerance asymmetry's shape is used; the degree-distribution mechanism is not claimed, and Doyle et al. [2005] is the reason it should not be
ln N as the entropy of a uniform distribution [Shannon, 1948]	**2 — shared form**	**information entropy, not thermodynamic entropy**; the correspondence is exact as arithmetic and carries no thermodynamic content
a bloc as a frozen region [Castellano et al., 2000]	**3 — analogy**	suggestive and used as motivation only; no phase-transition claim is made and no order parameter is computed
normal accidents [Perrow, 1984]	**3 — analogy**	cited as qualitative theory and framing, not as a formal result; this paper's finding is Perrow's in a measurement frame

No category-3 borrowing is used as evidence anywhere in this paper, and the two that appear are named as motivation. Curzi [2026] is the standing requirement that the theorem, the ordinary observed effect and this framework's own conjecture be distinguishable in three separate sentences on demand, and the table above is that requirement discharged.

10.6 Measurement-theoretic status of the scored inputs

The evidence ladder of Section 4.1 is ordinal and is not summed. Its classes are ranked and the assignment is made by identifying the strongest available class, which an ordinal scale licenses. **No numeric evidence score is computed and no threshold is arithmetic**, deliberately, because a numeric score would invite the averaging the scale forbids.

The de-identified capability comparison of Section 4.3 is not totalled. Court depth, statutory text, seasoning period and maintenance cost have no common scale, and the table produces a best fit by argument rather than a verdict by arithmetic.

The rubric of Section 6.2 is ordinal and is summed, and that is the paper's weakest operation. Weighting and adding 0–4 bands assumes equal intervals within each band and commensurability across the five tests, and neither was tested. Michell [1997] is the strongest available objection: a discipline routinely assumes quantitative structure without testing for it, and the formal remedy is conjoint measurement [Luce and Tukey, 1964], whose conditions establish whether an attribute has additive quantitative structure at all. **The conjoint conditions were not tested, here or in the source. The weighted total is therefore read as an ordering with a disqualification rule attached and not as a quantity**, no difference between two totals

is interpreted anywhere in Section 7.5, and Krantz et al. [1971] sets out what would have to hold before it could be. Stevens [1946] is the taxonomy the audit uses.

This paper inherits the framework-level scalar-reduction limitation rather than restating it. Reducing a jurisdictional structure to one effective count and one h, and reducing a captive arrangement to one weighted total, are two instances of the habit Curzi [2026] names once as a single limitation for the whole framework: where the source literature describes a phenomenon by several jointly varying features, the framework replaces it with one number and does arithmetic on the number. **The limitation is inherited here in full and is not met separately.**

The concentration instrument itself is untouched by this objection, because exposure shares are ratios of capital rather than judgements, and Section 7.6's bounds are closed-form arithmetic on stated constraints.

10.7 Postulates

#	statement	status
P1	exposure shares s_i are non-negative and sum to 1	DEFINITION
P2	$\text{HHI} = \sum_i s_i^2$	DEFINITION
P3	for n equal shares of $1/n$, $\text{HHI} = n(1/n)^2 = 1/n$	ARITHMETIC
P4	the aggregate shock is $R = \sum_i w_i X_i$ with $\text{Var}(X_i) = 1$	DEFINITION
P5	$\text{Var}(R) = \sum_i \sum_j w_i w_j \rho_{ij} = w' R w$	ARITHMETIC
P6	units inside one coordination perimeter act on one decision	POSTULATED (H1)
P7	the surplus law's input term is $\eta \ln N$	DEFINITION (modelling choice, CCP-A02)
P8	$\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$	POSTULATED (fixed in advance)
P9	total assets equal third-party liabilities plus statutory capital and surplus	DEFINITION (balance-sheet identity)
P10	the target capital level is 120 per cent of the enhanced requirement	DEFINITION (as described from the source, uncited)
P11	distributions above stated thresholds require supervisory consent	POSTULATED — described from the source, uncited (Appendix B)
P12	reserve redundancy 8 per cent, discount rate 8 per cent, average lag 4 years	POSTULATED (chosen, not measured)
P13	the rubric's weights are 25 / 25 / 20 / 20 / 10 and its bands are 0–4	DEFINITION (this framework's stipulation)

10.8 The compounding chain

step	consumes	statement	status
L1	P2, P3	HHI^{-1} returns n exactly for n equal shares	ARITHMETIC
L2	L1	HHI^{-1} is a numbers-equivalent — Adelman [1969]	CITED
L3	L2	it is the Hill number of order two — Hill [1973], Jost [2006]	CITED
L4	P4, P5	with $R = I$, Var$(R) = \sum w_i^2 =$ HHI	ARITHMETIC
L5	P4, P5	with R block-of-ones, Var$(R) = \sum_b W_b^2$	ARITHMETIC
L6	L4, L5	naive and bloc indices are one variance under two correlation assumptions	ARITHMETIC
R2	L6, P6	four states 40/25/20/15: naive 0.285000 (3.508772); bloc 1.000000 ($N = 1.000000$)	ARITHMETIC
R2b	R2, P7, P8	$\ln N = 0$ exactly, so $\eta \ln N = 0$ and $h = -0.500000$	ARITHMETIC
R1	—	entity 27.3, currency 61.0, settlement 58.0; ratios 2.234 and 2.125	ARITHMETIC on CONSTRUCTED inputs
R3a	P9, P10	10,800,000 over 16,100,000 is 67.08 per cent; 4,080,000 over the same is 25.34 per cent; 1,220,000 over the same is 7.58 per cent	ARITHMETIC
R3b	R3a, P11	the 25.34 per cent cannot leave without consent, so realisable-within-a-year is 7.58 per cent	rests on P11, **uncited to the bibliography**
R3c	R3a, P12	$(5{,}300{,}000 + 712{,}000)/1.08^4 = 4{,}418{,}999$; $/16{,}100{,}000 = 27.45\%$; haircut 72.55%	ARITHMETIC on POSTULATED parameters
R5	—	$2{,}700{,}000 - 410{,}000 - 1{,}560{,}000 - 140{,}000 = 590{,}000$; $/4{,}370{,}000 = 13.50\%$	ARITHMETIC on CONSTRUCTED inputs
R4	P13	Case A $20/400$ and disqualified; Case B $390/400$; Case C $375/400$	ARITHMETIC on ORDINAL inputs — **see 10.6**
R6	L6, P7, P8	crediting a distinct bloc: 0.295392, $N = 3.385332$, $\eta \ln N = 0.426808$, $h = -0.073192$; gain 0.049225	ARITHMETIC

step	consumes	statement	status
R7	L1, L6	twelve units at 5–15 per cent give $N \in [9.090909, 12]$; nine at ≤ 16 per cent give $N \in [6.443299, 9]$	ARITHMETIC
CR	R3a, P9	the same three quantities bound the owner's realisation and the creditor's	ARITHMETIC

Every ARITHMETIC row can be redone by hand or in three lines of any language. Every CITED row names one source and no more. Only P6, P8, P11 and P12 are POSTULATED. **P11 is the one that carries legal content and it is the one with no bibliographic support at all** — the supervisory gate is described from the source manuscript and is listed in Appendix B. **A reviewer who does not accept P11 loses R3b and keeps R3a, R3c and CR**, because the liability commitment and the balance-sheet identity do not depend on it. That is stated here so a reader can see exactly how much of the result the weakest postulate carries.

The chain's honest summary: four postulates, nine definitions, and arithmetic on constructed inputs. R2, R6, R7 and CR are arithmetic on definitions. R3b rests on an uncited regulatory description. R3c rests on three chosen parameters. R4 rests on stipulated weights and ordinal bands. **None of this is a proof of anything about the world, and the section says so rather than leaving the reader to notice.**

10.9 Cross-checks performed

- The three realisability shares were checked to sum to 1.000000: $0.670807 + 0.253416 + 0.075776 = 1.000000$.
- The balance-sheet identity was checked twice: $16{,}100{,}000 - 10{,}800{,}000 = 5{,}300{,}000$, and $5{,}300{,}000 - 4{,}080{,}000 = 1{,}220{,}000$.
- The target capital level was checked against its definition: $1.20 \times 3{,}400{,}000 = 4{,}080{,}000$.
- The premium reconciliation was checked: $1{,}670{,}000 + 2{,}700{,}000 = 4{,}370{,}000$, the prior commercial spend exactly.
- $s_{\max}^2 \leq \text{HHI} \leq s_{\max}$ holds on the four-state vector: $0.160000 \leq 0.285000 \leq 0.400000$.
- The effective-count ceiling $s_{\max}^{-2}$ is respected: $3.508772 \leq 6.250000$.
- The two safe-harbour bounds were checked against their equal-share cases: $1/12 = 0.083333$ gives 12.000000 and $1/9 = 0.111111$ gives 9.000000, and both maxima were verified to be attained at the stated corner allocations.
- The bloc-credit increment was checked both ways: $0.426808 - 0.377583 = 0.049225$ and $0.35(\ln 3.385332 - \ln 2.941176) = 0.049225$.
- The rubric's ordering was cross-checked against the disclosure test of Section 7.5, an instrument with no arithmetic in it at all; the two agree on all three arrangements.
- The two source discrepancies of Section 7.8 were quantified at 0.000183 in h and USD 29 in the present-value line, and the third at 0.000042 in the naive index.

11. Limitations

Every number in this paper is computed on constructed inputs and none is a measurement of a real structure. The exposure shares, the currency and settlement percentages, the balance sheet, the premium programme and the three scored arrangements are all **CONSTRUCTED** and are labelled as such at each point of use. **No first-party measurement of any real system appears anywhere in this paper.** That is this corpus's central limitation, stated here in the

terms Curzi [2026] states it in — of 27 parameter-by-scope cells, 19 have a defined estimator, 1 has a validated one, and 0 carry a first-party measurement — and a reviewer would be right to press on it before pressing on anything else.

The 61.0 per cent and 58 per cent figures have no as-of date and no source, because they have no referent. A concentration figure taken from real holdings requires both; these are constructed inputs to a constructed case and the source manuscript says so in its own verification note. **They are therefore evidence about the size of a gap under a construction and are not evidence that any real structure has that gap.** H2's ratios are computed from the same three inputs H2 names, which makes the test a consistency check rather than an independent confirmation, and Section 7.9 says so.

The settlement axis is not decomposed. Section 3.2 records that clearing and messaging are different layers and that a structure reporting two rails may have one messaging dependency; this paper nonetheless uses a single settlement share and does not separate them. Curzi [2026] does the decomposition properly, and the 58 per cent figure here should be read as an aggregate that a proper decomposition could move in either direction.

The correlation is assumed, not measured. $\rho = 1$ within a perimeter is the conservative end of a family and no correlation was estimated from data. For the four-state case this is the least exposed such assumption in the series — every shared mechanism of Section 4.1 is in the strongest evidence class — but it is still an assumption, and it is the quantity H1 turns on.

Every parameter in the surplus law was chosen. $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ are test parameters fixed in advance, not estimates of any real position, so the break-even $N = 4.172734$ is a property of the chosen parameter set. Any of the four values would move it.

The three financial parameters behind the 27.45 per cent are chosen and the sensitivity is stated. At zero discounting the figure is 37.34 per cent; at zero reserve redundancy it is 24.20 per cent. **The result is therefore mostly a statement about the balance-sheet structure and only secondarily about the discounting**, which is a strength of the result and a weakness of how it is usually quoted.

The realisability result has a domain and quoting it outside that domain is a misapplication. It is a result about a long-tail book (Sections 5.1 and 10.3), and an arrangement writing only short-tail risk has much less of the characteristic than the same table would suggest.

The rubric's arithmetic is not licensed by its scale type. Section 10.6 states this in full. In addition, **no sensitivity analysis over the weights was run**, so a reader who believes claims-paid should carry 30 rather than 20 cannot read the consequence off any table here. **And the rubric's Test 2 hard zero is stated on a naive unit count, which Section 7.6 shows is the wrong count.** That is a defect in an instrument this paper carries; it is disclosed rather than repaired mid-paper, and any later revision should state the threshold on the effective count.

The reflexive specification is not identified. Section 1 states this and names the remedy adopted — stipulate the reference group and defend the stipulation, following Brock and Durlauf [2001]. **The price is that every bloc assignment here is a modelling position rather than an identified empirical result**, and no amount of subsequent observation would change that. Manski [1993] and Shalizi and Thomas [2011] are the reasons.

The Lucas critique applies to this framework's own parameters and is not answered. Lucas [1976] is the general statement; the specific form is that this framework asserts that systems change behaviour when modelled and observed, and then treats η, δ and the bloc assignment as stable under exactly that response. **Either the parameters are stable under the framework's own predicted reaction or the framework undercuts itself, and this paper does not establish which.** MacKenzie and Millo [2003] documents the loop happening in the field with dates and price data, which is an existence proof for the mechanism and not a defence of the parameters.

The functional form of the surplus input is underdetermined by the evidence. Latané [1981] proposes a power law with exponent below one and Peters and Gell-Mann [2016] derives a logarithmic form from multiplicative dynamics with no psychological assumption at all. **All fit the qualitative facts equally well, nothing in this paper discriminates between them, and the logarithm is used because it is the form this framework carries.** Peters is the stronger of the two available warrants for it and is cited as such.

Small-sample bias in the effective count is not corrected. Chao et al. [2014] show that low-order diversity estimators are biased on small samples, and a four- or six-element exposure vector is a small sample. No correction was applied, and the paper does not resolve whether these vectors are samples or complete censuses — arguably the latter, since exposures are enumerated exhaustively. **A reviewer would be right to press on it.**

Scale invariance was checked at three scales only (Section 10.4), which is not enough to describe the instrument as scale-free without qualification.

The instrument is least informative in the regime where the headline case sits [Golosov, 2010, Dunleavy and Boucek, 2003], as Section 9 states.

No statistical test was run and no p-value appears, because there is no sample and no null. Every rate reported is an exact ratio of exact constructed amounts. **No claim of "no difference" is made from any non-significant result, because no significance test exists here to make it from.**

No scaling, power-law or heavy-tail claim is made and the Clauset–Shalizi–Newman procedure was therefore not run [Clauset et al., 2009]. Any restatement reaching for such a claim must run it first, and read Broido and Clauset [2019] on how often the claim fails once it is run.

No out-of-sample validation, no matched null and no information criterion appear, because no model was fitted. Bruch and Atwell [2015]'s calibration-and-validation checklist is unmet in every row: nothing here is calibrated against observed data, validated out-of-sample, or subjected to a sensitivity analysis beyond Sections 10.3 and 11.

11.1 Rigour checks not met, listed by name

1. **Conjoint measurement conditions untested** [Luce and Tukey, 1964], so the rubric's weighted total is unlicensed as a quantity and is used only as an ordering.
2. **The framework-level scalar-reduction limitation is inherited** [Curzi, 2026] and not met separately.
3. **The reflection problem is stated and not solved** [Manski, 1993]; the remedy adopted is stipulation, and its price is paid in Section 1 and again above.
4. **The Lucas critique applied to this framework's own parameters is unanswered** [Lucas, 1976].
5. **The functional form of the surplus input is undiscriminated** [Latané, 1981, Peters and Gell-Mann, 2016].
6. **Small-sample bias in the effective count is uncorrected** [Chao et al., 2014].
7. **No out-of-sample test, matched null, information criterion, or effect size with an interval** [Bruch and Atwell, 2015], because no model was fitted.
8. **No first-party measurement of any real system** [Curzi, 2026].
9. **The dimension of η is unresolved** because the source is silent (Section 10.1).
10. **No sensitivity analysis over the rubric's weight vector** (Section 10.4).
11. **The rubric's Test 2 threshold is stated on the wrong count** and the correction is disclosed rather than applied (Section 7.6).
12. **The settlement axis is not decomposed into clearing and messaging layers** (Section 3.2).
13. **Every legal, regulatory and doctrinal statement in Sections 4 to 6 is uncited** (Appendix B).

12. Conclusion

This paper measures four things on constructed inputs and reports one counter-result. The following statements are the fixed text of those results.

R1. Measured three ways, the same constructed holding shows entity-based United States exposure of 27.3 per cent, a dollar-denominated share of the balance sheet of 61.0 per cent, and approximately 58 per cent of settlement volume clearing through the United States system

— **the entity map understates the dependency by more than a factor of two,** $61.0/27.3 = 2.234$ on the currency axis and $58.0/27.3 = 2.125$ on the settlement axis. All three are constructed inputs and none carries an as-of date.

R2. Delaware 40.0 / South Dakota 25.0 / Nevada 20.0 / Wyoming 15.0 gives a naive HHI of 0.285000 and 3.508772 effective jurisdictions, but a bloc-HHI of 1.000000, an effective count of $N = 1.000000$ and a surplus term of exactly zero, giving $h = -0.500000$ at $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$. **Unlike the comparable cases in this series, no generous counterfactual exists**, because every shared mechanism sits in the strongest evidence class.

R3. On a constructed year-seven statutory balance sheet with total assets of USD 16,100,000, 67.08 per cent of assets is committed to third-party liabilities and 25.34 per cent is held as target capital, leaving **7.58 per cent realisable within a year and 27.45 per cent on a present-value basis — a 72.55 per cent haircut**. The present-value figure is model output at an assumed 8 per cent redundancy, 8 per cent discount rate and four-year lag; at zero discounting it is 37.34 per cent, and the result is a result about a long-tail book.

R4. The captive rubric scores 100 weight points across risk shifting (25), risk distribution (25), actuarial pricing (20), claims paid (20) and insurer conduct (10), each scored 0–4 for a weighted total to 400, **with hard zeros for run-back risk, fewer than five units, no pre-inception report, no claims in three years, or related-party investment, and a zero on risk shifting or risk distribution disqualifying regardless of total**. Three constructed arrangements score 20 out of 400 and disqualified, 390 out of 400, and 375 out of 400. **The total is an ordering, not a quantity.**

R5. The retained programme's economics: $2{,}700{,}000 - 410{,}000 - 1{,}560{,}000 - 140{,}000 = 590{,}000$, which is **13.50 per cent of the previous commercial spend of 4,370,000**, before investment income and before any tax consequence in any jurisdiction.

R6. Crediting Bermuda as a distinct bloc raises the effective bloc count from 2.941176 to 3.385332 and the surplus term from 0.377583 to 0.426808, moving h from -0.122417 to -0.073192, a gain of 0.049225. **It is a change in the bloc call and not a change in the structure**, and the base case declines it. The source manuscript prints -0.1226 and -0.0734 for the same two figures; Section 7.8 reconciles the 0.000183 difference and this paper uses the exact values.

Counter-result. The property that makes an insurance reserve resistant to third-party claims — real liabilities owed to unrelated parties, gated by a supervisor — is the same property that makes it unavailable to its owner, leaving only 7.58 per cent realisable within a year. The argument cannot be sold on both sides at once, and the inconsistency in doing so is arithmetic rather than rhetorical.

One further result is established as method rather than as measurement. The risk-distribution test that decides whether a captive arrangement is insurance is the bloc collapse of Curzi [2026] on a different share basis: three insureds behind one common mode have an effective count of 1.000000 whatever their shares, thirteen insured units behind one manufacturing process likewise, and a published constraint of twelve subsidiaries each at five to fifteen per cent of risk is exactly an effective-count floor of 9.090909 against a ceiling of 12.000000, with nine members none above sixteen per cent giving $[6.443299,\ 9.000000]$. **The doctrine's arithmetic content is the reciprocal index, and the doctrine's silence about correlation is the same silence the naive jurisdictional count keeps.** The consequence, disclosed rather than suppressed, is that this paper's own rubric states its risk-distribution threshold on the naive count and should state it on the effective one.

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every case in this paper — the exposure shares, the currency and settlement percentages, the statutory balance sheet, the premium programme, the three scored arrangements and the cost figures — is **CONSTRUCTED**

and is labelled as such at each point of use. The numerical results are exact computations on those constructed inputs, carried at full precision and rounded only for printing, except the present-value figure of Section 7.3, which is model output at the three parameters stated there.

All figures were produced by the scripts named in their captions (`figs/src/CCP-055_fig01.py` through `figs/src/CCP-055_fig05.py`), which import the shared figure style module and read their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

Jurisdictional, regulatory and doctrinal descriptions in this paper are taken from the source manuscript and not from primary legal sources. The bibliography contains general comparative, treaty and standard-setting material only; it contains no statutes, regulations, revenue rulings, court rules, registry documents, licensing frameworks or case reports for any individual jurisdiction. Appendix B lists the claims affected. **Any reader relying on those descriptions must verify them against the official sources, and must record the date of that verification.**

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. All persons, entities, holdings and arrangements described are constructed for the purpose of exercising the measurement.

Preregistration

The parameter set $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ and the break-even criterion $N \geq e^{0.5/0.35} = 4.172734$ were fixed by Curzi [2026] before the measurements reported here, and are not adjusted anywhere in this paper. **No structure measured in this paper clears that break-even**, which Section 7.7 reports in the same language a clearance would have received. The rubric weights of Section 6.2 and the H3 criterion of Section 2 were likewise fixed before the arrangements were scored and the balance sheet was decomposed. **No other element of this paper was preregistered.**

Legal, tax and investment advice

This paper is analytical and is not legal, tax, or investment advice. It describes entity types, jurisdictions, licensing classes and structural arrangements for the purpose of measuring concentration, dependency and realisability. No recommendation is made in respect of any particular person, asset, structure, jurisdiction or arrangement. Statutory avoidance periods, seasoning periods, charging-order provisions, insurance licensing classes, capital and solvency

requirements, substance requirements, premium ceilings, diversification requirements and reporting obligations all change and vary by jurisdiction and by activity; **every such statement in this paper is carried from the source manuscript, is flagged in Appendix B, and requires confirmation with qualified counsel — and, for any insurance matter, with a credentialed actuary — in the relevant jurisdiction before any reliance.**

Appendix A. The reporting standard for a four-axis measurement

Any structure measured under this instrument reports the following, and nothing less. Items 1 to 9 are carried from Curzi [2026]'s reporting standard; items 10 to 14 are added by this paper.

```
 1  share basis, stated in one sentence, per axis
 2  exposure table, units and shares, summing to 100.0%
 3  naive HHI, in fraction and in points, and effective units
 4  bloc assignment with the evidence class for each unit
 5  bloc-HHI and the effective count N
 6  perimeter exposure as a percentage
 7  ln(N), eta ln(N), and h at the stated parameters
 8  the robustness table over contested assignments
 9  the date, and the diff against the previous measurement
10  the SAME five lines on the currency axis
11  the SAME five lines on the settlement axis, with clearing and
    messaging layers separated
12  the SAME five lines on the counterparty axis
13  the largest of the four, named as the binding one
14  for any risk-transfer vehicle: total assets, third-party
    liabilities, the regulatory capital floor, and the three
    realisability shares, with the assumed redundancy, discount
    rate and settlement lag printed beside them
```

Item 13 is the item this paper exists to add. A structure reporting only item 3 has reported the flattering axis, and a structure reporting items 3 and 5 has reported the flattering axis correctly. **Item 9 decays first and matters most**, because perimeter drift is invisible except as a difference between two measurements taken a year apart.

Item 14's last clause is not optional. A realisability figure quoted without its three parameters is model output presented as arithmetic, and Section 10.2 records that this is the one headline figure in this paper exposed to that error.

A note on item 11. This paper does not meet its own item 11 — the settlement share is reported as an aggregate — and Section 11 says so. The standard is stated at the level the measurement should reach rather than the level this paper reached.

Appendix B. Claims requiring verification against primary sources

The following statements rest on legal or regulatory detail that changes, and **none of them is cited to the bibliography, which contains no primary sources for any individual jurisdiction**. They are reproduced from the source manuscript and require confirmation with counsel in the relevant jurisdiction — and, for insurance matters, with a credentialed actuary — before any reliance. **Record the date of each confirmation**; a statement of a regime without a date is not a statement of that regime.

1. The beneficial-ownership reporting regime described in Section 4.4: the enactment, the implementing rule, the constitutional challenges, the nationwide injunction, the stay, the subsequent interim rule narrowing the regime, whether a final rule has superseded it, and the present state of litigation and legislation. **This is the most volatile item in the paper and it changed twice in eighteen months.**
2. The federal avoidance provision described in failure mode 2 of Section 8, its ten-year period, and its interpretation in the relevant circuit.
3. The uniform fraudulent-transfer legislation: the enacting jurisdiction's version, its badges of intent and its limitation periods. Enactments vary jurisdiction by jurisdiction.
4. The appellate and first-instance decisions described in Section 4.2 on the impossibility defence, on a statute purporting to confer exclusive jurisdiction, on most-significant-relationship analysis in bankruptcy, and on public policy against self-settled spendthrift trusts. **Case citation follows different conventions and no case is in the bibliography.**
5. The decision described in Section 4.2 on the charging order and single-member entities, and the responsive legislation.
6. The estate and gift tax exposure of a non-resident holding shares with situs in the jurisdiction, the statutory exemption amount, the current exclusion following recent legislation, and any applicable treaty relief.
7. The real property tax and withholding regime described in Section 4.1, including its reach through holding corporations.
8. The corporate, alternative-entity, trust and perpetuities provisions of each of the four states de-identified in Section 4.3, including seasoning periods, excepted creditor classes, charging-order exclusivity and its extension to closely held corporate shares, directed-trust architecture, sealing provisions and private-trust-company regimes. **The source deliberately does not assert section numbers it has not confirmed, and neither does this paper.**
9. The insurance licensing statute and classification described in Section 5: class definitions, the percentage thresholds separating classes, minimum capital by class, and whether the constructed class assignment of Section 7.3 is correct on a book with a minority unrelated element.
10. **P11, the supervisory gate.** The exact thresholds and section numbers for the restriction on reduction of statutory capital and for the affidavit requirement on distributions, and whether prior approval or an affidavit is required in the case described. **This is the one postulate in Section 10.7 carrying legal content and no bibliographic support; Section 10.8 states exactly which result depends on it.**
11. Whether the relevant insolvency law confers a statutory priority on policyholders over general unsecured creditors in the winding up of an insurer.
12. The solvency-equivalence determination described in Section 5, its scope, its effective date, and the treatment of captives and special-purpose insurers within or outside it.
13. The reciprocal-jurisdiction and qualified-jurisdiction determinations described in Section 5, and their dates.
14. The economic-substance legislation and regulations described in Section 5, the relevant activity definitions, and the treatment of insurance and pure holding entities.
15. The corporate income tax legislation described in Section 5: the in-scope entity definition, the revenue threshold, and the first fiscal year of application.
16. The exchange and reporting positions described in Sections 4.4 and 5, including the intergovernmental agreement model and current participating-jurisdiction status.
17. The doctrinal anchor described in Section 6.1 requiring both risk shifting and risk distribution, the four-question working formulation, and the appellate decisions on brother-sister structures.
18. **The three published safe harbours described in Section 6.1**, on which the effective-count bounds of Section 7.6 are computed: the twelve-subsidiary constraint at five to fifteen per cent each, the seven-or-more unrelated members constraint, and the unrelated-premium-at-half constraint, together with the negative ruling on a small number of insureds. **The bounds are arithmetic; the constraints are not verified here.**

19. The small-insurer election described in Section 6, its indexed premium ceiling, and the exact mechanics of both alternative diversification requirements added by later legislation.
20. The reportable-transaction and disclosure regulations applying to the arrangements described in Section 7.5, and whether the constructed Case A pooling arrangement would fall within them.
21. The segregated-cell and protected-cell legislation referenced in Section 5, and whether cell segregation is recognised outside the jurisdiction of formation.
22. The constitutional position of the risk-transfer centre described in Section 7.7, including its final appellate arrangement and the external capacity to legislate for it by order, on which the bloc assignment there partly rests.
23. All cost, premium, loss and balance-sheet figures in this paper are constructed. **They are not fee quotations, rate filings or statutory returns, and no real schedule was used to produce them.**

All numerical results in this paper were computed from the stated inputs at full precision and rounded only for printing, and the three places where a printed value differs from the exact computation are reported in Section 7.8 rather than harmonised.

Chapter 56

Substance Repricing and the Clearing-System Count: The Atlantic Cluster, ADGM, and the Limits of a Second Enclave

Abstract. Two findings are reported against the common assumption that adding entities adds resilience. First, the Atlantic offshore cluster: a legacy structure across the United Kingdom 40.0 per cent, the United States 22.0, Cayman 14.0, the British Virgin Islands 12.0, the Bahamas 7.0 and Barbados 5.0 reads a naive HHI of 0.249800 and 4.0032 effective jurisdictions, but a bloc-HHI of 1.000000 and an effective count of 1.00, because every entry sits in the same coordination perimeter — adding a Western entity to a Western structure leaves ln(N) at exactly zero. Second, the two Emirati financial free zones are one line: the DIFC and the ADGM share a single national perimeter, one federal anti-money-laundering regime, and one exposure share of 29.1 per cent, so no rearrangement between them changes the effective bloc count. The substance requirements introduced across the Atlantic cluster are priced rather than described: annual cost rises from a baseline of 1,550 to 5,350 for a pure equity-holding entity (3.45 times) and to 19,350 for an entity carrying a relevant activity (12.48 times), which sets a break-even at roughly USD 7.74 million of held value at 25 basis points. The one axis where diversification is achieved is settlement: measured by currency, the structure moves from a single dollar system at HHI 1.0000 and 1.00 effective clearing systems to 2.81, with the dollar still at 55 per cent — a floor set by counterparties, not by entity placement.

Keywords: economic substance; offshore financial centres; ADGM; financial free zones; clearing systems; settlement diversification; concentration measurement

Series: D. Structure and Resilience **Paper:** CCP-056

1. Introduction

Two findings are reported against the common assumption that adding entities adds resilience, and one of them is a general proposition about what a marginal unit can be worth.

The first is that a legacy structure spread across an Atlantic offshore cluster is one point of failure. Six entities — United Kingdom 40.0 per cent of exposure, United States 22.0, Cayman 14.0, the British Virgin Islands 12.0, the Bahamas 7.0 and Barbados 5.0 — read a naive

Herfindahl–Hirschman Index of 0.249800 and 4.003203 effective jurisdictions. Collapsed onto the coordination perimeter every entry occupies, the same shares read a bloc index of 1.000000, an effective count of exactly 1.00, and $\ln N = 0$. **Adding a Western entity to a Western structure leaves the surplus term at exactly zero**, and the constructed family in this case had been told in writing, more than once, that the structure was internationally diversified.

The second is that the two Emirati financial free zones are one line. The Dubai International Financial Centre and Abu Dhabi Global Market share a single national perimeter, one federal anti-money-laundering regime, one standard-setting-body assessment and one exposure share of 29.1 per cent, so no rearrangement of entities between them changes the effective bloc count. That is the same result Curzi [2026] reaches from the apex direction — the apex vehicle does not move the concentration reading, and the measurement remains naive HHI 0.2934 and 3.41 effective jurisdictions against bloc-HHI 1.0000 and 1.00 effective perimeters whatever vehicle sits at the top — and the same result Curzi [2026] reaches on the Crown Dependencies, where Jersey 40.0 / Guernsey 35.0 / Isle of Man 25.0 gives naive HHI 0.345000 and 2.898551 effective jurisdictions but bloc-HHI 1.000000, $N = 1.00$ and a surplus term of exactly zero. **Three papers, three sets of entities, one number.**

Between them sits a measurement of the instrument that repriced the first cluster. The substance legislation enacted across the Atlantic centres is almost universally described as a prohibition. **It was a price change, and this paper computes it**: annual cost rises from a constructed baseline of USD 1,550 to USD 5,350 for a pure equity-holding entity, a multiple of 3.45, and to USD 19,350 for an entity carrying on a relevant activity, a multiple of 12.48. Because the cost is fixed per entity rather than proportional to assets, **the instrument was regressive against holding size**: at a two-million-dollar holding it consumes 96.75 basis points a year and at two hundred million it consumes 0.97, with a break-even against a stated 25-basis-point tolerance at USD 7,740,000 of held value. That is why the large-deal and fund businesses in these centres grew straight through the period while the retail company business contracted, and it is why commentary describing the events as the end of offshore was reading the wrong end of the distribution.

The one axis on which the constructed structure achieves genuine diversification is settlement. Measured by the terminal clearing system rather than by entity, the structure moves from a single dollar system at HHI 1.0000 and 1.00 effective clearing systems to 2.810568 across five, with the dollar still at 55 per cent. That residual is set by the currency the counterparties pay in and not by entity placement, and it is the result Curzi [2026] and Curzi [2026] both work against from different directions — the latter measuring the rail directly, where a naive correspondent HHI of 0.302400 (3.306878 effective correspondents) collapses to a perimeter HHI of 0.745256 and $N_{\text{rail}} = 1.341819$ once flows are decomposed by currency.

What this paper adds beyond those measurements is a bound. Section 3.4 states and Section 7.5 verifies an elementary proposition about the marginal unit: if a set of units holding fixed relative proportions has its own effective count M, then adding one more unit and rescaling the rest raises the effective count to at most $M + 1$, and that maximum is attained only when the new unit takes share $1/(M + 1)$. Applied to the constructed settlement mix, the four non-dollar systems have $M = 3.799250$, so **the ceiling on this structure's effective clearing-system count is 4.799250, attained at a dollar share of 20.8366 per cent** — and reaching 4.00 requires the dollar share to fall to 38.9912 per cent, not to the "roughly one third" the source manuscript states. **The proposition also says that eliminating the dollar entirely would be worse than holding a fifth of settled value in it**, giving 3.799250 rather than 4.799250, which is counter-intuitive, exact, and checkable in two lines.

The paper reuses the surplus functional and the affine effective count of Curzi [2026] — $f(u) = \ln(1+(N-1)u)$, uniquely characterised by angle-only dependence, vanishing at parallel, an affine effective count and normalisation, with $N_{\text{eff}} = 1 + (N - 1)u$ — and the roster-audit form of that count in Curzi [2026c], where the affine form is badged as a design axiom rather than a derived result and a hundred-person department at $u = 0.05$ has an effective count of 5.95. **Section 3.1 states plainly that the affine count is not the count this paper uses**, and why. It presupposes the founding reframing of Curzi [2026], that comfort is the felt total of present

costs so that the present state is a local minimum of the comfort landscape by definition rather than by luck, which is the reason a legacy structure accumulated by three sets of advisers over three decades persists without anyone defending it. And it presupposes the result of Curzi [2026d] that adding a seventh discipline to an institute makes it worse — the period change falls to -0.002863 and the fixed point drops from 6.265 to 5.964 — because divergence, not domain count, was the deficiency. **That is this paper's second finding, stated in another domain a volume earlier.**

The bloc assignment carries an identification problem and it is stated here rather than in a footnote. Manski [1993] proves that a linear-in-means reflexive specification is **not identified from observational data**: the endogenous effect cannot be separated from contextual effects or from correlated unobservables however much data is collected, and Shalizi and Thomas [2011] show that influence and selection are generically confounded in observational network data. **No amount of observing jurisdictional behaviour could confirm a bloc assignment derived reflexively.** The remedy adopted is the fourth of four: the reference group is stipulated in advance and the stipulation is defended, on the evidence ladder of Section 4.1 and the listing record of Section 4.3, which is the route Brock and Durlauf [2001] shows can be made rigorous. **The price is that every bloc assignment here is a modelling position rather than an identified empirical result.**

2. Hypotheses

Five hypotheses are stated. Each is falsifiable in the sense Platt [1964]'s strong inference requires, and H5 is falsifiable by counterexample only, which is stated rather than dressed up. Meehl [1967] is the standing warning that directional predictions are cheap; four of the five below carry point predictions with constants attached.

H1 — the Atlantic-cluster hypothesis. Entities distributed across an Atlantic offshore cluster occupy one coordination perimeter, so their naive count and their effective count diverge to the maximum possible extent. *Predicted:* UK 40.0 / US 22.0 / Cayman 14.0 / BVI 12.0 / Bahamas 7.0 / Barbados 5.0 gives naive HHI 0.249800 and 4.003203 effective jurisdictions, and bloc-HHI 1.000000, $N = 1.000000$, $\ln N = 0$ exactly. *Falsified by:* a defensible bloc assignment placing two of the six in different perimeters — equivalently, an instance in which a listing, a designation or a de-risking decision reached one and not the others without a second independent decision.

H2 — the repricing hypothesis. The substance instrument operated as a price change and not as a prohibition, and because the price is a fixed annual amount per entity it is regressive against holding size. *Predicted:* a cost multiple of $5{,}350/1{,}550 = 3.45$ for a pure equity-holding entity and $19{,}350/1{,}550 = 12.48$ for a relevant-activity entity, with the relevant-activity cost falling from 96.75 basis points of a two-million-dollar holding to 0.97 basis points of a two-hundred-million-dollar holding, and a break-even at USD 7,740,000 against a stated 25-basis-point tolerance. *Falsified by:* a cost structure proportional to assets rather than fixed per entity, which would remove the regressivity; or evidence that entities were prohibited rather than repriced.

H3 — the one-perimeter hypothesis. The two Emirati financial free zones, the DIFC and the ADGM, are one line in the measurement, so no rearrangement of entities between them changes the effective bloc count. *Predicted:* the two zones share a single exposure line of 29.1 per cent, and moving any entity from one to the other leaves bloc-HHI 0.340000 and $N = 2.941176$ unchanged to every decimal place. *Falsified by:* a rearrangement between the two zones that moves N at all; equivalently, evidence that a decision taken at the federal level reaches one zone and not the other without a second independent decision.

H4 — the settlement-axis hypothesis. The settlement axis, unlike the entity axis, admits genuine diversification, and the improvement is measurable in the same units. *Predicted:* a structure settling wholly in one currency reads HHI 1.0000 and 1.00 effective clearing systems; the constructed redesign at USD 55.0 / CNY 15.0 / EUR 12.0 / INR 10.0 / AED 8.0 reads HHI 0.3558

and 2.810568. *Falsified by:* a decomposition showing that the additional systems' messaging legs run over the first system, in which case the count collapses back toward 1.00 and the apparent diversification is an artefact of counting the wrong layer.

H5 — the marginal-unit proposition. Adding one unit to a set of units held in fixed relative proportions, and rescaling the rest, raises the effective count by at most exactly one. *Predicted:* if the incumbent set's own effective count is M, then $\mathrm{EFF}(x) = 1/(x^2 + (1-x)^2/M)$ is maximised at $x = 1/(M+1)$ with value $M+1$. For the constructed settlement mix, $M = 3.799250$, so the ceiling is 4.799250 at a dollar share of 0.208366, and $\mathrm{EFF} = 4.000000$ requires $x = 0.389912$. *Falsified by:* any admissible share vector of this form whose reciprocal index exceeds $M + 1$. **H5 is a proposition of arithmetic rather than an empirical hypothesis. It is falsifiable by counterexample and by nothing else, and it carries no empirical weight whatever; what it does is bound a design space that is otherwise argued about without limits.**

Not hypotheses, and labelled as such. Four further statements appear and none is falsifiable as stated.

- *The logarithmic form of the input side*, $\eta \ln N$ rather than ηN, is a **modelling choice** carried from Curzi [2026]. Section 11 records that Latané [1981]'s sub-linear power law fits the qualitative facts as well and that nothing here discriminates between them.
- *The parameter set* $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ is a **modelling choice** carried from Curzi [2026] and fixed in advance, which is the preregistration discipline of Nosek et al. [2018] applied to an engineering measurement.
- *The 25-basis-point tolerance* against which the substance break-even is computed is a **stated judgement, not a standard**. Any other tolerance moves the break-even proportionally, and Section 7.2 gives the formula so a reader may substitute their own.
- *The identification of the surplus law's N with the effective bloc count* is a **definition**; its content is in the assignment rule, which is testable.

3. Method

3.1 The instrument, and which effective count it is

The concentration instrument is taken unchanged from Curzi [2026], which reaches naive HHI 0.189002 (5.29 effective jurisdictions) and bloc-HHI 0.340000 (2.94 effective blocs) on the designed structure this paper also measures. Its definition fits on one line:

```
HHI = sum of squared shares            effective count = 1 / HHI
```

introduced by Hirschman [1945] and reached independently in sum-of-squares form by Herfindahl [1950]. The reciprocal reading is due to Adelman [1969]. Competition practice states the index in points on percentage shares [U.S. Department of Justice and Federal Trade Commission, 2023, 2010]; this paper states it as a fraction and gives both where confusion is possible.

Which effective count, stated exactly. The reciprocal index HHI^{-1} is the Hill number of order two [Hill, 1973], whose numbers-equivalent reading Jost [2006] states in its cleanest form; the same arithmetic was reached independently in ecology by Simpson [1949] and in political science by Laakso and Taagepera [1979]. The exponential of Shannon entropy [Shannon, 1948] is the same family at order one, and the reciprocal index is also the exponential of the Rényi entropy of order two [Rényi, 1961]. **Every effective count in this paper — of jurisdictions, of blocs, and of clearing systems — is HHI^{-1}, the Hill number of order two, and it is named as such at each use.**

The affine count is a different quantity and is not used here to produce any number. Curzi [2026c] runs $N_{\mathrm{eff}} = 1 + (N-1)u$ with $u = 1 - (x \cdot y)^2$ and surplus $S = \ln N_{\mathrm{eff}}$, giving a hundred-person department at $u = 0.05$ an effective count of 5.95, and **badges the affine form as a design axiom rather than a derived result**, which is the correct badge. It is a linear interpolation

between 1 and N, it is not a Hill number of any order, and it carries no external support; Curzi [2026] establishes this by running the same verification chain on both counts and letting one fail visibly. **Wherever this paper says "effective count" it means the reciprocal index, and the affine count appears by name only.**

The general form. The bloc collapse generalises the reciprocal index to EFF $= 1/(w'Rw)$ for exposure weights w and correlation matrix R, which is the similarity-sensitive diversity of order two of Leinster and Cobbold [2012], whose additive relative is the quadratic entropy of Rao [1982]. Setting $R = I$ recovers HHI^{-1}; setting R to a block matrix of ones recovers the bloc index. Hurlbert [1971] is the standing objection that "diversity" carries no content without its formula attached.

No distributional or scaling claim is made anywhere in this paper. The share vector 40.0 / 22.0 / 14.0 / 12.0 / 7.0 / 5.0 descends steeply and invites a tail argument. **This paper fits no distribution, characterises no tail and asserts no power law.** Clauset et al. [2009] is the reason: a principled maximum-likelihood fit with a goodness-of-fit test and a likelihood-ratio comparison against alternatives leaves few canonical data sets standing, and a straight line on a log-log plot is not evidence. That procedure was not run because no claim requiring it is made, and Stumpf and Porter [2012] adds that even a genuine power law is weak evidence for a mechanism. **The effective counts here are statistics of finite vectors with four, five or six entries.**

The collapse result has an ancestry that should be named rather than reinvented. Granovetter [1978] establishes that the distribution of individual thresholds, not the mean, determines whether a collective outcome tips, and Schelling [1971] is the canonical demonstration that a mild individual preference produces a sharp aggregate result — the strongest available warning against reading an individual-level parameter straight into an aggregate, which is what a bloc assignment does. Castellano et al. [2000] supplies the closest formal analogue of a bloc: the same interaction that raises local agreement freezes global diversity into non-interacting regions, with a sharp transition between the monocultural and fragmented phases. **A bloc, in that vocabulary, is a frozen region and the effective count is a count of them.**

3.2 The share basis, the four axes, and the layer that ruins them

This paper measures against **arbitrary state action** — a designation, a listing, an exchange-control measure, a compliance department's change of risk appetite — and not against lawful process. The instrument is identical on each of four axes and only the share basis changes: jurisdiction, settlement rail, counterparty, and currency. Curzi [2026] measures the mismatch between the first and the last on the same constructed structure, reporting entity-based United States exposure of 27.3 per cent against a dollar-denominated balance-sheet share of 61.0 per cent and settlement volume through the United States system of approximately 58 per cent — the entity map understating the dependency by more than a factor of two. **This paper measures the rail axis, which is the axis that turns out to move.**

One distinction ruins naive measurements on that axis and this paper's H4 turns on it. A clearing system and a messaging system are different layers. A structure reporting two rails may have one messaging dependency: measured correctly, "one messaging network at 100 per cent" and "85 per cent on that network plus 15 per cent on a second clearing system whose messaging leg runs over the first" are the same number, and only the second looks diversified. **The share basis used here is the terminal clearing system**, defined as the system in which final settlement occurs, and the messaging layer is treated as a separate count that this paper does not compute. Curzi [2026] computes it and finds the collapse; **this paper's 2.81 is therefore an upper reading on one layer and not a statement about the structure's total rail independence**, which Section 11 records as its principal exposure.

Three consequences carry through the paper.

A currency is a perimeter exposure that no entity move can exit. A balance denominated in a unit whose clearing runs through banks inside a perimeter carries that perimeter's exposure however the entities are arranged.

A rail is not a holding. The optionality a bridging jurisdiction supplies is *transit* optionality, whose parameters are hop size, exposure duration, counterparty pre-screening and mid-transit failure handling. Treating a transit account as a store of value confuses an exposure-window problem with a custody problem.

A jurisdiction's policy is not a bank's risk appetite. A dual-rail jurisdiction supplies a market in which dual-rail counterparties exist and must be found, onboarded and maintained one at a time; it does not supply the counterparty.

3.3 Correlated units, and what the literature says about the correction

The naive index overstates resilience because it treats units as independent when they share a coordination mechanism. The correction is not novel as mathematics: it is Markowitz [1952]'s result that diversification benefit depends on covariance and not on count, and Buldyrev et al. [2010]'s that a system of interdependent networks is far more fragile than either network alone. Perrow [1984] reached the organisational form first — interactive complexity plus tight coupling makes catastrophic failure a property of the system rather than of any component — and Sagan [1993] documents the sharper version, that redundancy can reduce safety. **Both borrowings are labelled under the analogy discipline of Section 10.5, and this paper's first finding is Perrow's result in a measurement frame.**

Correlations rise under stress. Longin and Solnik [2001] find extreme correlation in equity markets that is asymmetric and concentrated in downturns, and Ang and Chen [2002] report the same asymmetry. **The effective count falls exactly when it is needed.** The ecological name for the decisive variable is synchrony [Loreau and de Mazancourt, 2008], and the strongest empirical demonstration that effective count buys stability is Schindler et al. [2010]'s Bristol Bay sockeye record, with roughly half the variability and about a tenth of the fishery closures. The contrary literature is read alongside: Doak et al. [1998] show that diversity–stability correlations can arise from statistical averaging alone, and May [1972] showed that increasing diversity and connectance destabilises randomly assembled systems. **This paper's claim is narrower than the general diversity–stability claim**: the reciprocal of a variance is an effective count of independent shocks, which is arithmetic, and units inside one perimeter are correlated, which is H1 and H3.

3.4 The marginal-unit proposition

The result stated as H5 is elementary and is given here because Sections 7.4 and 7.5 depend on it and because it bounds an argument that is otherwise conducted without limits.

Let $n-1$ incumbent units hold shares in fixed relative proportions $q_1, \dots, q_{n-1}$ summing to one, and let their own reciprocal index be

$$M = \Big(\sum_i q_i^2 \Big)^{-1}.$$

Introduce one further unit at share x and rescale the incumbents to $(1-x)q_i$. The new index is

$$\mathrm{EFF}(x) \;=\; \frac{1}{x^2 + (1-x)^2 \big/ M}.$$

Minimising the denominator gives $2x - 2(1-x)/M = 0$, so $x^* = 1/(M+1)$, and substituting returns

$$\mathrm{EFF}(x^*) \;=\; M + 1.$$

A marginal unit is worth at most exactly one effective unit, and it is worth that only at one share. Three readings follow, and each does work in Section 7.

The ceiling is hard. No placement of the new unit exceeds $M + 1$, so a design argument that promises more than one additional effective unit from one additional entity, currency or rail is refuted before the details are examined.

The optimum is an equal share with the equal-equivalents. $x^* = 1/(M + 1)$ is exactly the share the new unit would hold if the whole set were equalised at $M + 1$ units, which is the numbers-equivalent property of Section 3.1 reappearing as a design rule.

Both tails are worse than the optimum, and the near tail is the one people reach for. $\mathrm{EFF}(0) = M$, so removing the new unit entirely returns the incumbent count; but $\mathrm{EFF}(x)$ falls on both sides of x^*, so **a unit that is too large is not merely less useful than an optimally sized one, it can be worse than useless relative to the incumbents' own count**. Section 7.5 measures exactly that case.

The proposition is arithmetic, it carries no empirical content, and Section 10.8 records it as such in the chain.

3.5 Evidence standard, and the de-identification convention

Sections 4, 5 and 6 are jurisdictional and regulatory exposition. **Every enactment, regulation, licensing regime, listing event, delisting event, rate, date and jurisdictional characteristic in them is described from the source manuscript and is not cited to the bibliography.** The bibliography carries the general instruments and the comparative doctrine only: recognition and enforcement of awards [United Nations, 1958] with the grounds on which enforcement is in fact refused [van den Berg, 1981] and the arbitration doctrine around it [Born, 2021, United Nations Commission on International Trade Law, 1985]; the trusts convention [Hague Conference on Private International Law, 1985]; the succession convention that never entered into force [Hague Conference on Private International Law, 1989] with the regional regulation that partially fills the gap [European Union, 2012]; the judgments conventions [Hague Conference on Private International Law, 2005, 2019]; conflict-of-laws doctrine [Torremans et al., 2022, Briggs, 2013]; the trust and the foundation compared [Lupoi, 2000, Smith, 2013, Panico, 2014]; asset partitioning as the irreducible function of entity law [Hansmann and Kraakman, 2000] with its critique [Pistor, 2019]; the listing machinery as a technique of governance [de Goede, 2011]; and the disclosure instruments [Organisation for Economic Co-operation and Development, 2017, Financial Action Task Force, 2012, Organisation for Economic Co-operation and Development, 2015]. **It contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction.** Appendix B lists every affected claim, **with the requirement that the date of confirmation be recorded**, because a statement of a jurisdiction's regime without a date is not a statement of its regime.

Where a table would otherwise **rank named jurisdictions against one another on a legal capability the bibliography cannot support**, the jurisdictions are de-identified — "Atlantic centre A" through "Atlantic centre D", "enclave A" and "enclave B" — following the convention of Curzi [2026], which de-identifies two Gulf centres in its own scored assessment, and Curzi [2026], which de-identifies three islands in its ranked capability table. **Where a jurisdiction is named in prose, or appears in an exposure table, or carries a bloc assignment with its evidence, the name is used**, because the argument is about those specific perimeter attachments.

4. The perimeter, evidenced

4.1 The four-question test applied to the Atlantic cluster

Four jurisdictions and one point of failure. The Cayman Islands, the British Virgin Islands, the Bahamas and Barbados are treated together not because they are similar in law — they are not, and two of them are sovereign states while two are overseas territories — but because they occupy the same position on the only axis this paper measures. **Everything in this section is described from the source manuscript and is uncited** (Section 3.5, Appendix B).

test	Cayman	BVI	Bahamas	Barbados
exchanges account information under the common standard	yes	yes	yes	yes
implements the common anti-money-laundering standard through domestic legislation	yes	yes	yes	yes
placed under increased monitoring or on a coordinated list within the last eight years	yes	yes	yes	yes
clears United States dollars through United States correspondents	yes	yes	yes	yes

Four yeses each. The evidence ladder used across this series ranks demonstrated near-simultaneous mirroring of designations and listings, shared correspondent networks that exit corridors together, a supervisory standard producing the same domestic consequence in every member, and binding accession to a common enforcement instrument as the strongest classes; participation in an information-exchange framework and treaty density are weaker; diplomatic alignment is weakest. **Rows two and four of the table sit in the strongest class and row three is the observed record that Section 4.3 sets out.**

4.2 Sovereignty is not perimeter independence

One of the four is an independent sovereign state and has been for half a century. There is no external legislative power over it, no equivalent of the transparency commitment that binds the two territories, and no metropolitan department with a view. **If sovereignty produced perimeter independence, that jurisdiction would sit in a different bloc from the other three.**

It does not. It exchanges account information, it was placed under increased monitoring for a period measured in years, it has appeared on a coordinated list, and its currency is pegged one-to-one to the dollar with the corresponding correspondent dependency. **Sovereignty and perimeter independence are different variables, and this is the cleanest available demonstration that the first does not produce the second.** That observation is worth more to the measurement than anything in the jurisdiction's statute book, and it is the reason the bloc assignment in Section 7.1 does not turn on constitutional status.

4.3 The listing record, read for its structure

jurisdiction	event	duration
Cayman	coordinated list, Feb 2020 to Oct 2020	8 months
Cayman	increased monitoring, Feb 2021 to Oct 2023	32 months
BVI	coordinated list, Oct 2023 to Feb 2024, after a peer-review downgrade on exchange of information on request	4 months

jurisdiction	event	duration
Bahamas	increased monitoring, Oct 2018 to Dec 2020	26 months
Bahamas	coordinated-list appearances 2018 and 2019, with removals following commitments	dates to be verified
Barbados	increased monitoring, Feb 2020 to Feb 2024	48 months
Barbados	coordinated list, Oct 2020, following a peer-review downgrade; subsequently moved to the commitment annex	dates to be verified

Read the table for its structure rather than its content. Every one of the four was listed by at least one body within a five-year window. In two of the cases the coordinated listing was triggered mechanically by a technical rating change made by a *different* body, on exchange of information on request. **One peer-review panel moved four jurisdictions' bank-onboarding classifications across a five-year period, and no combination of those four provided any protection against it.** Two rows are marked for verification and are marked that way in the source; they are reproduced with the mark rather than repaired.

That is what a coordination perimeter looks like from the inside, and it is why the arithmetic of Section 7.1 comes out at one.

4.4 A standard-setting body with no enforcement power

The same mechanism is visible in a second and better-dated episode, and it is the cleanest demonstration in this series of what a perimeter is. The United Arab Emirates, which hosts both free zones, was placed under increased monitoring in March 2022 and removed in February 2024, a period of 23 months. The listing body can fine nobody, has no treaty, no court and no enforcement arm. **The effect travelled through three channels and only one of them was discretionary.**

Channel one — linked official listings. A regional bloc maintains its own list of high-risk third countries for anti-money-laundering purposes, and the state was added to it during the period. The consequence under that bloc's legislation is that every obliged entity — bank, insurer, fund administrator, professional firm — **must** apply enhanced due diligence to relationships and transactions involving the listed country. That is a legal requirement with supervisory consequences, not a recommendation, and a second jurisdiction operating an equivalent mechanism made an equivalent change.

Channel two — private risk models. Every large correspondent bank runs a country-risk model and a listing is a hard input to it. The output is not refusal but re-rating: extra documentation, senior sign-off for new relationships, periodic-review triggers, and at the margin exits from relationships that were already thin. **That is the de-risking mechanism, and it is why the practical effect of a listing is felt by small legitimate businesses long before it is felt by anyone the listing was aimed at.**

Channel three — discretion. Individual institutions, regulators and counterparties making their own judgements, at their own pace, with no obligation to explain.

A body with no enforcement power produced an effect across every major financial system within weeks, because its determinations are wired into other people's binding rules and into thousands of private risk models. That is the mechanism the bloc collapse is measuring, and it is why this series' thesis is about perimeters rather than about secrecy.

The listing period's remedial programme is worth one further observation because it bears on H1. The state's deficiency was not primarily its statute book: federal anti-money-laundering legislation and a beneficial-ownership regime were already in force, and the assessment methodology grades effectiveness as well as technical compliance. **What the period produced was supervision, investigation, prosecution and confiscation capacity — that is, effectiveness — which is exactly the class of change that binds a jurisdiction more tightly into the perimeter rather than loosening it.** A jurisdiction that exits increased monitoring has not moved out of the perimeter; it has demonstrated that it is inside it.

4.5 The register survived; public access did not

All four Atlantic jurisdictions maintain beneficial-ownership information. The trajectory of *access* has moved twice in opposite directions. An external transparency statute required the two territories to establish publicly accessible registers; implementation slipped repeatedly. A regional court then struck down the provision of that region's own anti-money-laundering legislation making member-state registers accessible to the general public, holding it a disproportionate interference with private life and data protection. That decision changed the territories' negotiating position, and both have legislated toward a legitimate-interest access model rather than open public access. **Commencement dates, the definition of legitimate interest, and the current enforcement position are all marked for verification** (Appendix B).

The register exists in every one of these jurisdictions and competent authorities have access to it. What changed is whether a stranger can search it.

The premise that makes this the right way round is not this work's to claim. Kerckhoffs [1883] stated it in 1883: a system must remain secure when everything about it except the key is public. Saltzer and Schroeder [1975] restate it as *open design* — the protection mechanism should not depend on the ignorance of attackers — and place it alongside *complete mediation* because a publishable design still requires that every access be checked. Brandeis [1914] is the public-interest form of the same observation. **Citing Kerckhoffs converts what reads as an ethical preference into a recognised engineering standard**, and Curzi [2026] states the constraint in its strongest form: channel separation is strengthened by disclosure, because its effectiveness is entirely independent of whether anyone knows it is in place, so every element must be publishable and non-disclosure is refused as a design premise rather than merely discouraged. **Any structure whose function depends on a stranger's inability to search is outside this paper's subject; a structure that merely prefers not to be searchable is expressing a lawful preference, and the line between them is the disclosure test.** The regulatory direction has made the point unavoidable from the other side: automatic exchange [Organisation for Economic Co-operation and Development, 2017], the common anti-money-laundering recommendations [Financial Action Task Force, 2012] and the substance and beneficial-ownership work of the base-erosion project [Organisation for Economic Co-operation and Development, 2015] together mean non-disclosure is not an available design parameter.

5. The substance instrument as a pricing intervention

5.1 What it was

A regional bloc's code-of-conduct group, working from criteria adopted in December 2017, applied a test to jurisdictions with no or only nominal corporate taxation: entities claiming residence there had to demonstrate real economic activity, and the alternative was inclusion on a list of non-cooperative jurisdictions. Each of the four Atlantic jurisdictions legislated, with effect from the beginning of 2019 in three of the four; the fourth addressed the criterion differently because its exposure was ring-fencing rather than nominal taxation. **All of the resulting statutes have been revised and the operative text must be taken from the current revision** (Appendix B).

The common architecture: an entity carrying on a *relevant activity* — banking, insurance, fund management, financing and leasing, headquarters business, shipping, distribution and service-centre business, intellectual-property business, and holding business — must be directed and managed in the jurisdiction, conduct its core income-generating activities there, and maintain adequate employees, expenditure and physical presence proportionate to that activity. Pure equity-holding entities are subject to a reduced test. Intellectual-property business, identified as the most mobile, attracts the heaviest treatment.

5.2 The cost table

The instrument's effect is best seen in a cost line rather than in a legal analysis. **Constructed and illustrative; not a quotation, and actual costs vary widely by jurisdiction, agent and activity.**

component, USD per year	pre-2019	post-2019
registered agent and registered office	1,100	1,600
government or annual licence fee	450	1,050
substance classification and annual return	0	900
beneficial-ownership filing and maintenance	0	400
accounting records and annual financial return	0	1,400
locally resident directors, relevant-activity entity only	0	14,000
total, pure equity-holding entity	**1,550**	**5,350**
total, relevant-activity entity	**1,550**	**19,350**

The instrument prohibited nothing. It multiplied a cost base, and a product whose entire proposition was that it cost almost nothing does not survive being multiplied by twelve. Section 7.2 computes the multiples and the regressive property.

5.3 What survived, and why

What survived the substance instrument is precisely what had substance, and the point is worth stating because it is the mechanism, not a moral.

A fund with unrelated investors, an administrator, an auditor and a custodian always had substance, so a fund industry survived and grew. A specialist commercial court is substance by definition, so a court-led jurisdiction became more rather than less important. A private-client statute book that drafts genuinely novel instruments is substance, so it survived. A treaty network survived in narrowed form, because a treaty benefit was never available to an entity with nothing in it whatever local practice may once have suggested. **What died was the empty company: the cheap, unstaffed, undocumented vehicle formed for a few hundred dollars and left to sit.**

That distribution of outcomes is an audit result and this paper reports it as one, following the reporting form of Curzi [2026], whose audit distribution over thirteen entries — seven maps, two splits, one rendering, three rejections, with five entries carrying a failure — is the model for printing a disposition table rather than a narrative. Section 7.7 disposes of four legacy entities the same way.

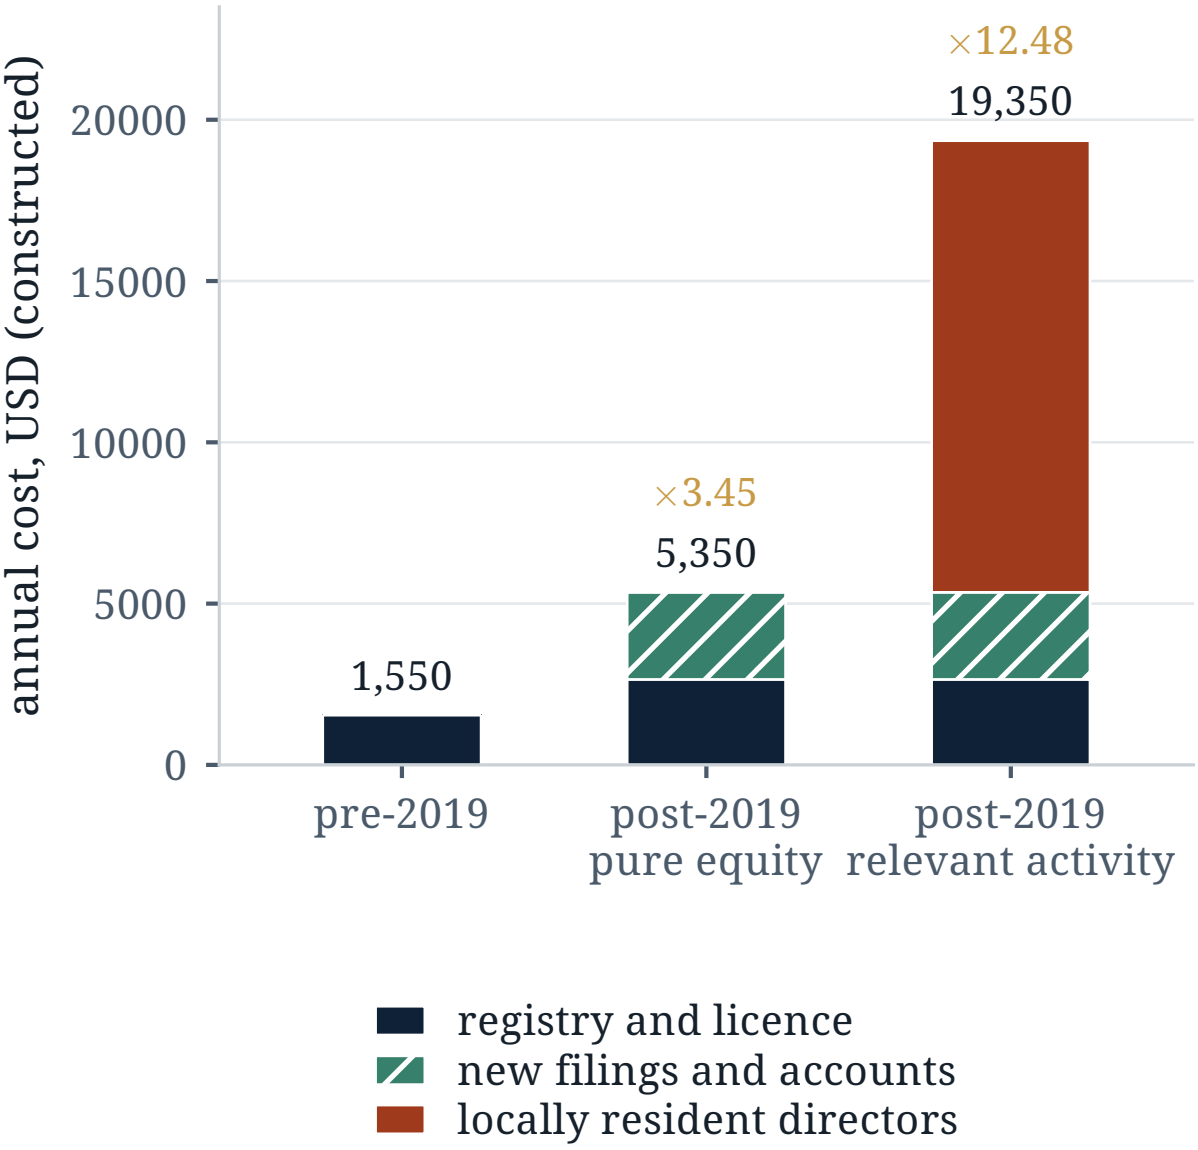


Figure 56.1: **The substance instrument multiplied a fixed cost base by 3.45 for a holding entity and by 12.48 for a relevant-activity entity.** Constructed annual cost of an offshore holding company, grouped into registry and licence, new filings and accounts, and locally resident directors. All amounts are constructed and illustrative, not quotations. Generated by `figs/src/CCP-056_fig02.py`.

6. The second enclave: a body of law without its territorial state

6.1 The decomposition

The strongest single piece in the architecture this series builds is not a jurisdiction. It is a decomposition. **Everything in this section is described from the source manuscript and is uncited** (Appendix B).

```
A NATIONAL COMMERCIAL LAW
    =  a body of doctrine and precedent
    +  a court system
    +  a territorial state with a tax system, a public register,
       a sanctions perimeter, and a bloc alignment

AN ENCLAVE APPLYING THAT LAW DIRECTLY
    =  the body of doctrine and precedent
    +  its own court system, staffed from the same tradition
    −  the territorial state
```

Abu Dhabi Global Market applies the common law of England, including the principles and rules of equity, as it stands from time to time, under its own Application of English Law Regulations. Not a jurisdiction whose statutes were drafted in the tradition; not one that inherited the law at a colonial reception date and then drifted; not one that treats it as a residual gap-filler. **The body of law can be bought without the entity**, and an intermediate holding company inside such an enclave gets contractual certainty across a group that spans legal families, a legal spine for the constitution and shareholder relations that becomes load-bearing exactly when everything else has failed, and a court that will hear the dispute on the merits applying law the parties can predict.

The measurement consequence is the point of this paper's H3. Curzi [2026] establishes that an entity in the originating state brings a public control register, an identity verification regime, the in personam jurisdiction of that state's courts over its directors, its sanctions perimeter, and a share of exposure sitting squarely inside the Western coordination perimeter. **The enclave supplies the law and not the perimeter attachment. That is a real gain and it is not a gain in N.**

6.2 What it does and does not do

```
DOES:
  apply the originating state's common law directly, as it develops
  supply a common-law court with an internationally recruited bench
  supply low-substance holding and special-purpose forms
  supply a foundations regime as an alternative and a migration path
  sit outside the Western perimeter on the bloc assignment

DOES NOT:
  add a bloc where the holder is already exposed to the same country
  remove the holder from automatic exchange
  remove the holder from federal anti-money-laundering, sanctions or
    criminal law
  confer opacity -- beneficial ownership is filed
  confer a tax exemption by existence -- a qualifying-income test is
    a real test
  operate -- a vehicle with no staff is not a substance answer
  bridge -- a holding company is not a treasury company
```

Run the disclosure test against the first block: every item survives publication without alteration. An enclave holding company works *better* when the operating company's registrar, the risk-transfer supervisor and the counterparty bank all see the same shareholder, filed and consistent. It is a disclosure-irrelevant instrument in the strict sense: **its function does not depend on not being seen.**

6.3 Why two enclaves in one country are one line

The DIFC and the ADGM sit in one federal state. They have separate courts, separate registries, separate regulators and separate company law. They share a constitution, a currency, a central bank, a federal criminal law, a federal anti-money-laundering regime, a single assessment by the standard-setting body of Section 4.4, and a single entry on any coordinated list. **On the perimeter test they are one unit.**

This is Curzi [2026]'s result in a different position of the structure. That paper reports that the apex vehicle does not move the concentration reading — the measurement remains naive HHI 0.2934 and 3.41 effective jurisdictions against bloc-HHI 1.0000 and 1.00 effective perimeters whatever vehicle sits at the top. **Here the same invariance holds for the branch: whatever sits beneath the apex, if it sits in the same country, the line does not move.**

The mechanism is the one Curzi [2026d] identifies in another domain: adding a seventh discipline to an institute makes it worse — the period change falls to -0.002863 and the fixed point drops from 6.265 to 5.964 — **because divergence, not domain count, was the deficiency**. Two enclaves in one country are a domain count. The deficiency is divergence, and there is none between them on any operative test.

7. Results

All figures below are exact computations on constructed inputs, carried at full precision and rounded only for printing. **No external dataset was collected and no figure in this paper is a measurement of any real structure.** Where the source manuscript's printed value differs from the exact computation, Section 7.8 prints both.

7.1 R1 — the Atlantic cluster measured, and H1

The constructed legacy position: four Atlantic entities accumulated over three decades by three sets of advisers who did not speak to each other, plus two onshore holdings.

entity	share of protected capital	bloc
UK holding and property	40.0%	WEST
US listed portfolio and custody	22.0%	WEST
Cayman private unit trust	14.0%	WEST
BVI holding company over the operating company	12.0%	WEST
Bahamas discretionary trust	7.0%	WEST
Barbados captive	5.0%	WEST

```
naive HHI = 0.160000 + 0.048400 + 0.019600 + 0.014400
          + 0.004900 + 0.002500 = 0.249800
effective jurisdictions = 1 / 0.249800 = 4.003203

bloc-HHI  = 1.000000    N = 1.000000
ln(N) = 0.000000        eta ln(N) = 0.000000    h = -0.500000
```

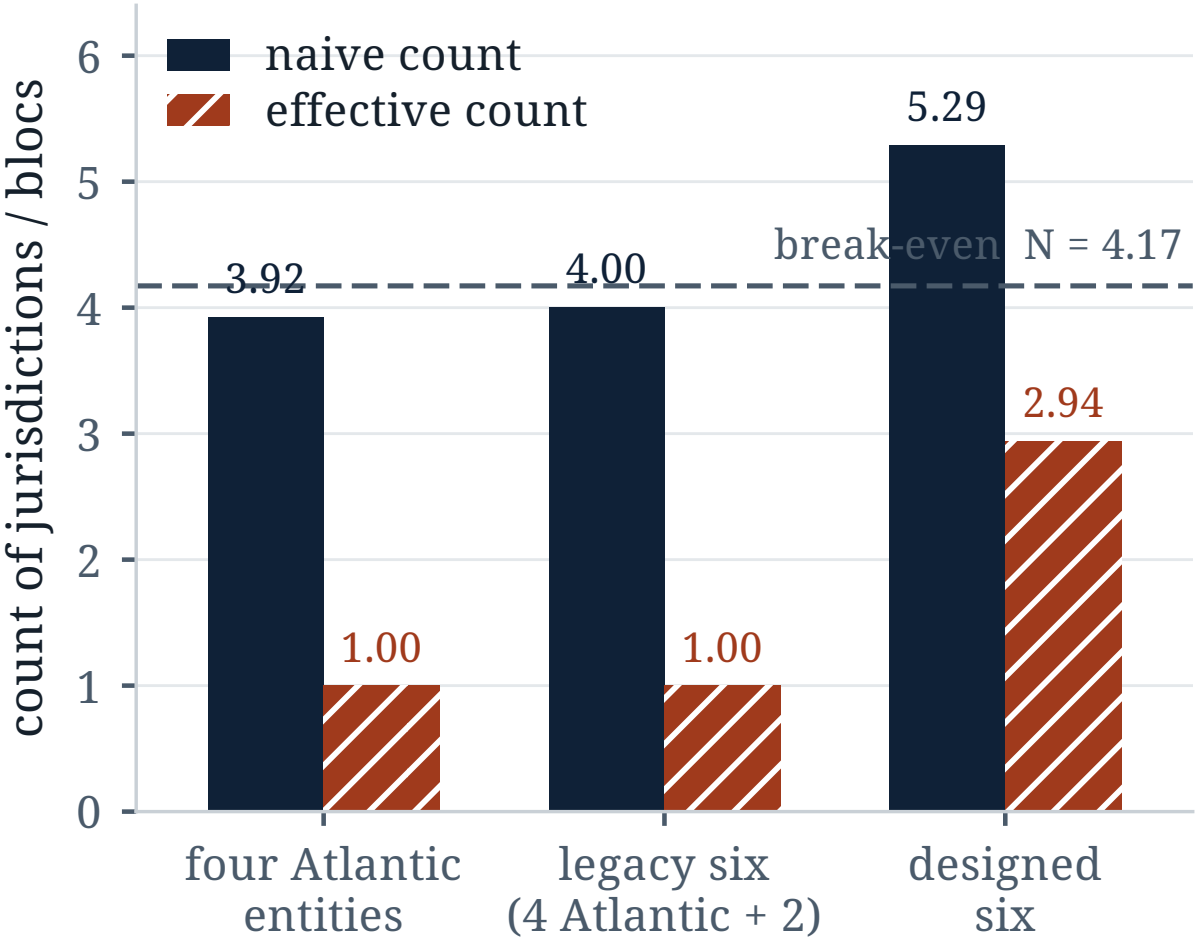


Figure 56.2: **Six Atlantic-cluster entities read 4.00 naive jurisdictions and 1.00 effective ones.** Naive count against effective count for the four-entity Atlantic variant, the six-entity legacy structure and the designed six-jurisdiction structure of CCP-052, with the break-even effective count $N = e^{0.5/0.35} = 4.172734$ the surplus law requires at the stated parameters. All shares are constructed. Generated by `figs/src/CCP-056_fig01.py`.

Six jurisdictions, an apparent effective count of 4.00, and one point of failure. The four Atlantic entities taken alone — a constructed variant at Cayman 30.0 / BVI 25.0 / Bahamas 25.0 / Barbados 20.0 — give naive HHI 0.255000 and 3.921569 effective jurisdictions against the same bloc-HHI of 1.000000. **The two vectors differ by** 0.081634 **in naive count and by exactly nothing in effective count**, which is the shape of the finding in one line.

H1 is supported. No admissible reassignment places two of the six in different perimeters: every entry returns four yeses on the test of Section 4.1, and the listing record of Section 4.3 shows the mechanism acting on all four Atlantic members within one five-year window. Section 4.2 disposes of the only argument that would separate them — that sovereignty confers perimeter independence — by exhibiting a sovereign member that behaves identically to the two territories.

7.2 R2 — the repricing and its regressive property, and H2

From the cost table of Section 5.2:

```
pure equity-holding entity     5,350 / 1,550  =  3.451613  ->  3.45x
```

```
relevant-activity entity      19,350 / 1,550  = 12.483871  -> 12.48x
```

Now the property that decides who should still be reading. The substance cost is a **fixed annual amount per entity**, and a fixed cost is regressive against asset size. Expressing the relevant-activity cost as a fraction of the capital the entity holds:

capital held, USD	annual substance cost, basis points
2,000,000	96.75
5,000,000	38.70
7,740,000	**25.00 — break-even**
20,000,000	9.675
200,000,000	0.9675

At a two-million-dollar holding the substance regime consumes ninety-seven basis points a year; at two hundred million it consumes one. The same rule, applied identically, was fatal at one end of the market and imperceptible at the other.

The break-even is a one-line formula and the tolerance is a judgement rather than a standard, so it is given in a form a reader can re-run with their own tolerance:

```
Q* = annual fixed cost / tolerance
   = 19,350 / 0.0025 = USD 7,740,000   at 25 bp
   = 19,350 / 0.0050 = USD 3,870,000   at 50 bp
   = 19,350 / 0.0010 = USD 19,350,000  at 10 bp
```

H2 is supported on both halves. The instrument is a price change: no entity type was prohibited and the cost table contains no prohibition. And the price is fixed per entity, which produces the hundred-fold spread in basis points across a hundred-fold spread in holding size. **The consequence for a holder is a threshold rather than a rule: an entity holding less than roughly USD 7.7 million is paying more than twenty-five basis points a year for a structure that adds nothing to N, and it should be collapsed rather than optimised.**

7.3 R4 — the two enclaves are one line, and H3

The designed structure, with the two enclaves consolidated into a single national line:

line	share	bloc
United Arab Emirates (DIFC and ADGM together)	29.1%	BRIDGE
Thailand	18.2%	NONALIGNED
United Kingdom	16.4%	WEST
Bermuda	13.6%	W-lite, folded into WEST
India	11.8%	NONALIGNED
Singapore	10.9%	BRIDGE

```
naive HHI = 0.084681 + 0.033124 + 0.026896 + 0.018496
          + 0.013924 + 0.011881 = 0.189002
effective jurisdictions = 1 / 0.189002 = 5.290949

blocs: BRIDGE 40.0  NONALIGNED 30.0  WEST 30.0
bloc-HHI = 0.160000 + 0.090000 + 0.090000 = 0.340000
effective blocs N = 2.941176
ln(N) = 1.078810   eta ln(N) = 0.377583   h = -0.122417
```

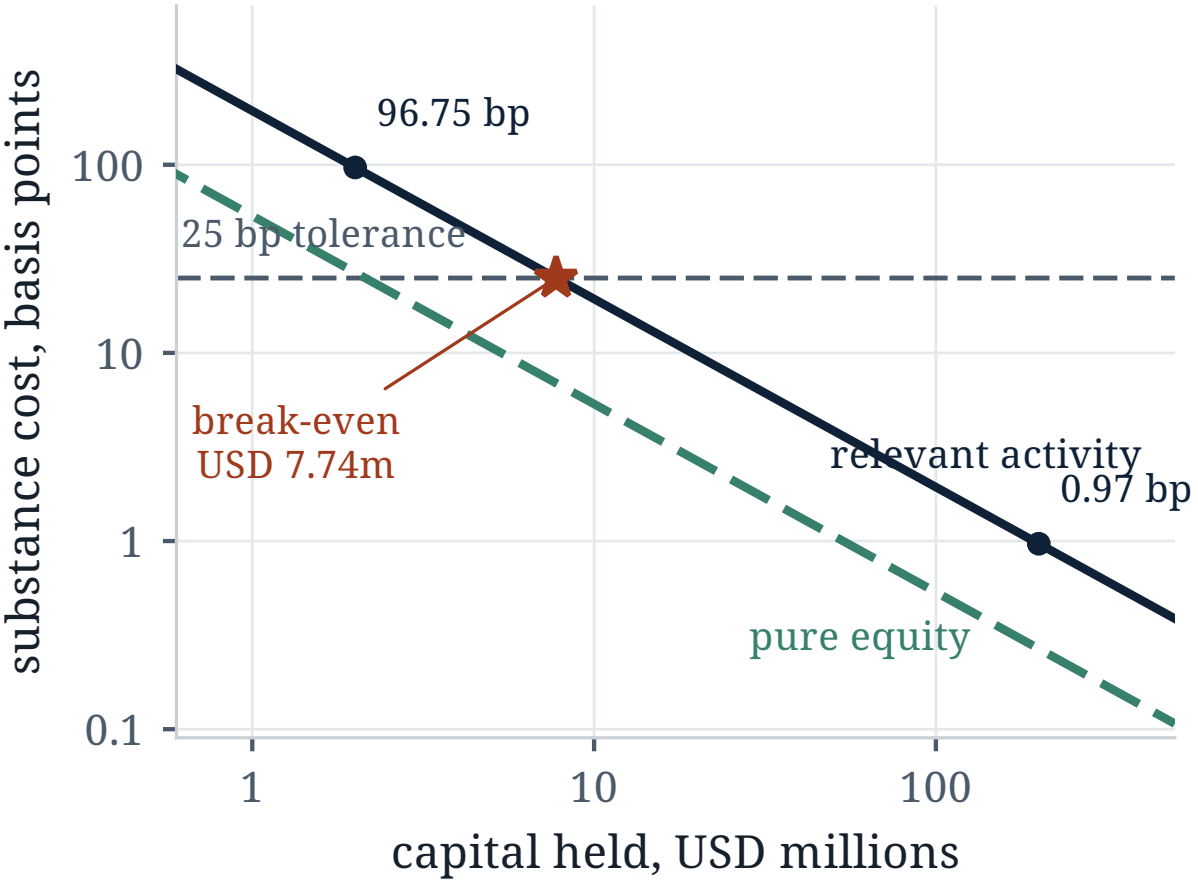


Figure 56.3: **A fixed annual substance cost is regressive: 96.75 basis points at two million dollars of capital and 0.97 at two hundred million.** The constructed relevant-activity and pure-equity costs of Section 5.2 expressed in basis points of capital held, against a stated 25-basis-point tolerance whose intersection with the relevant-activity line is the USD 7.74m break-even. The tolerance is a judgement, not a standard. Generated by `figs/src/CCP-056_fig03.py`.

H3 is supported and it is supported trivially, which is the point. Moving any entity between the DIFC and the ADGM changes no share in the table, because the table's first row is the country and both zones are inside it. The bloc-HHI is unchanged to every decimal place and so is N. **A second free zone buys capability, not diversification**, and the capability — direct application of a mature commercial law, a bench recruited from that tradition, a foundations regime, and a low-substance holding form — is real and is inventoried in Section 6.2.

What the enclave layer changed is on a different axis entirely, and the honest way to report it is a table with no arithmetic in it:

	before	after
governing law of group agreements	the originating state's law	the same body of law, applied by the enclave
forum for governance disputes	the originating state's courts	the enclave's courts
entity carrying group ownership	company in the originating state, public control register	enclave company, beneficial ownership filed but not publicly listed
entities between apex and operating company	2	1
probate steps to move the operating company on death	two grants	none: the apex is a foundation and does not die

The body of law is the same body of law. The bundle around it is gone. That is the whole trade, and it is available to anyone willing to read one set of regulations.

7.4 R3 — the clearing-system count, and H4

As found, the constructed group settled essentially everything through one bank into one dollar correspondent chain:

```
AS FOUND -- settled value by terminal clearing system
  US dollar system     100.0%
  HHI = 1.0000    EFFECTIVE CLEARING SYSTEMS = 1.00
```

A group with six jurisdictions and one clearing system has diversified its lawyers and not its exposure, and the number is the same number Section 7.1 produced on the entity axis, which is not a coincidence.

As designed, settled value is apportioned across the systems the business actually touches:

terminal clearing system	share of settled value
US dollar system	55.0%
renminbi system	15.0%
euro systems	12.0%
Indian rupee systems	10.0%
domestic dirham	8.0%

```
HHI = 0.3025 + 0.0225 + 0.0144 + 0.0100 + 0.0064 = 0.3558
EFFECTIVE CLEARING SYSTEMS = 1 / 0.3558 = 2.810568
improvement: 1.000000 -> 2.810568
```

H4 is supported, with the qualification its own falsifier names. The falsifier was a decomposition showing that the additional systems' messaging legs run over the first system. **That decomposition was not run here, and the source manuscript states plainly that many participants in the second clearing system still use the first system's messaging layer.** The 2.810568 is therefore a count of *terminal clearing systems* and is an upper reading on the structure's rail independence; Curzi [2026] runs the decomposition and finds naive correspondent HHI 0.302400 collapsing to a perimeter HHI of 0.745256 and $N_{\text{rail}} = 1.341819$. **H4 is supported on the layer it names and is not supported for the structure as a whole**, and Section 11 records this as the paper's principal exposure.

Four further changes accompany the count and only three of them enter it:

```
dollar-chain dependence      100.0%  ->  55.0%
effective clearing systems     1.00  ->   2.81
banking relationships             1  ->      4
independent correspondent chains for USD
                                  1  ->      2
standing relationship outside the host country
                                none  ->      1, live and dormant
```

The last line has no effect on the index and may be the most valuable item in the table. It is the relisting contingency of Section 4.4, and its value is entirely in existing before it is needed. That is a branch-weighting judgement rather than a measurement, and the framework has an instrument for it: Curzi [2026] scores each step of a path 0 to 3 for steepness, sums to a resistance R, and weights the branch as $1/2^R$, so that a three-point difference in steepness is a factor of eight in expectation. **A dormant relationship converts a steep branch into a shallow one, and the instrument says what that is worth in expectation even though the index says it is worth nothing.**

7.5 The frontier, the ceiling, and H5

The dollar share of 55 per cent dominates the result of Section 7.4 and cannot be engineered away, because it is set by the currency of the revenue and not by the structure. **The source manuscript states that reaching an effective count of 4.00 would require the dollar share to fall to "roughly one third" of settled value. That is checkable and it is not right.**

Holding the four non-dollar systems in their designed proportions $15 : 12 : 10 : 8$ and varying the dollar share x, the relative index of the incumbents is

```
r = (0.15^2 + 0.12^2 + 0.10^2 + 0.08^2) / 0.45^2
  = (0.0225 + 0.0144 + 0.0100 + 0.0064) / 0.2025
  = 0.0533 / 0.2025 = 0.263210
M = 1/r = 3.799250
```

so that $\text{EFF}(x) = 1/(x^2 + r(1-x)^2)$. Applying the proposition of Section 3.4 with $M = 3.799250$:

quantity	value
the designed position	$x = 0.550000$, EFF $= 2.810568$
the share required for EFF $= 4.000000$	$x = 0.389912$
the optimal share $x^* = 1/(M+1)$	$x^* = 0.208366$
the ceiling $\text{EFF}(x^*) = M + 1$	4.799250
the dollar eliminated entirely	$x = 0$, EFF $= M = 3.799250$

Three findings, and two of them correct something.

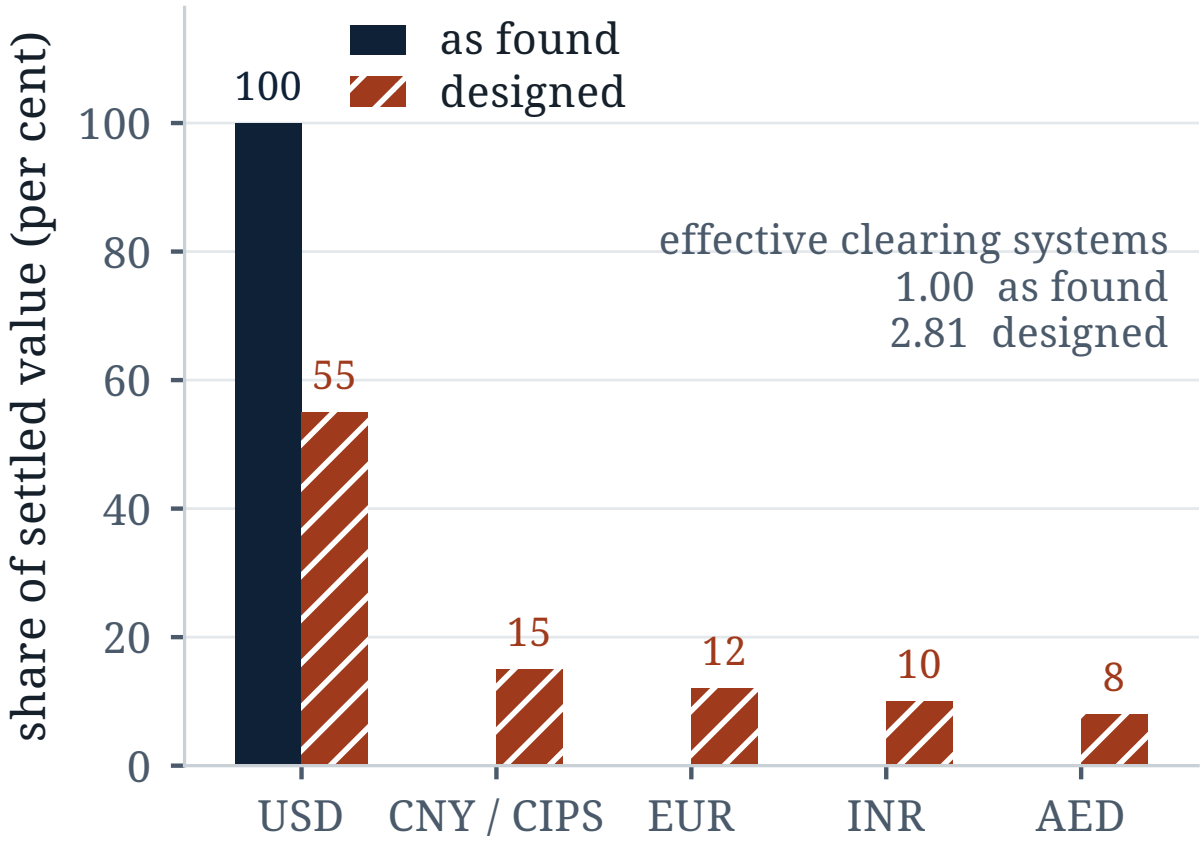


Figure 56.4: **Settled value moves from one clearing system to five, and the effective count from 1.00 to 2.81.** Share of settled value by terminal clearing system, as found and as designed, on the constructed case. The count is of terminal clearing systems only; the messaging layer is not decomposed (Section 3.2). Generated by `figs/src/CCP-056_fig04.py`.

First, the source's "roughly one third" understates the achievable dollar share by nearly six percentage points. The share required to reach four effective clearing systems is 38.9912 per cent, not 33.33 per cent. At exactly one third the count is 4.384374, well above four. **The source is conservative rather than wrong, but the exact figure is computable in one line and is printed here.**

Second, the structure has a hard ceiling of 4.799250 effective clearing systems and it cannot reach five, however the dollar is managed. That follows from the proposition of Section 3.4 with no further assumption: the four incumbent systems have $M = 3.799250$ between them, and one more unit adds at most exactly one. **A design conversation about "more rails" is bounded before it starts, and the bound is a property of the revenue mix rather than of anybody's ambition.**

Third, and counter-intuitively, eliminating the dollar entirely would be worse than holding a fifth of settled value in it. At $x = 0$ the count is 3.799250; at $x = 0.208366$ it is 4.799250. **The dollar at its optimal share is worth exactly one effective clearing system, and a policy of dollar elimination gives that up.** The result is exact, it is checkable in two lines, and it is the opposite of what a reduce-the-dominant-exposure heuristic recommends.

H5 is supported as a proposition of arithmetic. Section 10.8 records it in the chain with status ARITHMETIC and no empirical content, and Section 11 records that a proposition about a design space is not a finding about a system.

7.6 R5 — the residual

```
55% of settled value clears through a single system whose law is
not the holder's, and that share is set by the currency the
counterparties pay in rather than by entity placement.
```

That is the residual after redesign and it is the result other papers in the series cite. Three facts fix it and none of them is a structural variable. The business sells in dollars. The property position distributes in a currency converted through dollar markets. The securities portfolio is dollar-denominated. **No amount of jurisdictional cleverness changes the currency a customer pays in, and a structure that pretends otherwise is optimising the wrong variable.**

Three consequences are stated because each is a live error.

A dollar account is a dollar account. Dollar payments settle through the dollar correspondent system, are screened under that system's law, and are subject to its jurisdiction over the settlement regardless of where the account is booked or where the holder is incorporated. **"Outside the perimeter" is a statement about the entity and never about the currency.**

A second clearing system is not a replacement for a messaging system. Section 3.2 states the layer distinction and Section 7.4 states what follows: a structure built on the premise that a clearing system substitutes for a messaging network will fail, and the count reported here is of the layer that was measured.

The residual is not the same object as the entity residual. Curzi [2026] reports the entity-axis version — entity-based exposure of 27.3 per cent against a currency share of 61.0 per cent — and the two must not be netted against each other. **The rail layer and the entity layer are different exposures with different remedies, and improving one is not evidence about the other.**

7.7 The disposition of the legacy cluster, and its cost

The four Atlantic entities cost, on constructed figures:

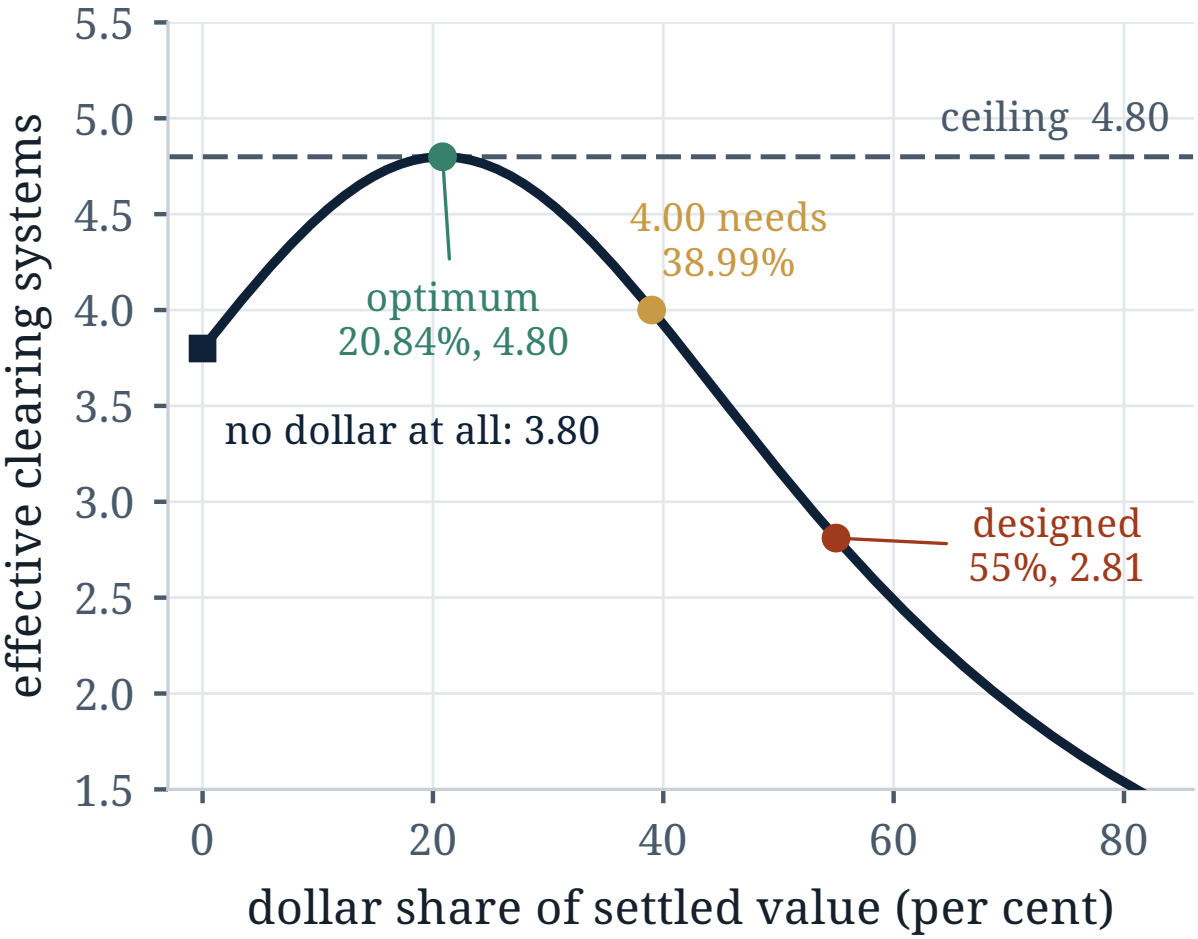


Figure 56.5: **The effective clearing-system count is bounded above by 4.80 and eliminating the dollar entirely is worse than holding 20.8 per cent of settled value in it.** The frontier $\mathrm{EFF}(x) = 1/(x^2 + r(1-x)^2)$ with the four non-dollar systems held in their designed proportions, $r = 0.263210$. Marked: the designed position at 55 per cent, the 38.99 per cent required to reach four systems, the optimum at 20.84 per cent, and the dollar-free case. Modelled from the constructed share vector of Section 7.4; generated by `figs/src/CCP-056_fig05.py`.

entity	constructed annual maintenance, USD
Cayman private unit trust: trustee, administration, audit	26,000
BVI holding company: agent, substance return, financial return, accounting records	9,400
Bahamas trust: trustee fees, accounting, advice	41,000
Barbados captive: manager, audit, actuary, supervisory fees	96,000
total	**172,400**

One hundred and seventy-two thousand dollars a year, buying zero effective blocs. That is the sentence the measurement exists to produce, and it is the only sentence in the exercise that has ever made anybody act.

The disposition, reported as an audit distribution in the form of Curzi [2026]:

entity	disposition	reason
BVI holding company	**collapse**	performs one function, interposition, which the enclave holding company performs better while additionally supplying the law, the bridge-bloc position and dual-rail banking
Cayman private unit trust	**collapse**	a fund structure with no unrelated investors, no administrator performing a real function and no fund purpose; above the economic break-even, so the case is not cost but that it adds a reporting jurisdiction and nothing to N
Bahamas discretionary trust	**migrate, slowly**	the statute book offers instruments the destination does not, and the argument runs both ways; decided on the perimeter measurement and on the succession analysis, over a period measured in years, on a timeline with no relationship to any claim
Barbados captive	**commute**	the treaty network that was the entire reason for the domicile is unused; the run-off requires a counterparty willing to assume it at a price, and is the slowest of the four actions

Two collapses, one migration, one commutation, and one entry where the framework's

own architecture argues against itself in print. The migration case is the honest one: the origin jurisdiction's statute book solves a protector-succession problem cleanly and the destination's does not, and the decision is made on the perimeter measurement notwithstanding. **That reasoning belongs in the file as a considered decision rather than as an omission.**

The result:

	legacy	designed
naive HHI	0.249800	0.189002
effective jurisdictions	4.003203	5.290949
bloc-HHI	1.000000	0.340000
effective blocs	1.000000	2.941176
Western-perimeter exposure	100.0%	30.0%
$\eta \ln N$	0.000000	0.377583
h	-0.500000	-0.122417

Entity count went from six to six. The naive effective count rose 32.16 per cent, from 4.003203 to 5.290949. The effective bloc count rose from 1.000000 to 2.941176, from a base of exactly zero surplus. And the honest companion result still applies: at the stated parameters the designed structure gives $h = -0.122417$ against a break-even count of 4.172734, so it remains below the threshold at which the surplus law permits building. **Jurisdictional diversity buys the ability to build; it does not pay for the building.**

7.8 Printed discrepancies, left visible

First, the naive index of the designed structure. The source manuscript prints 0.1889 in one place and 0.189002 in another for the same six shares. The exact value computed from the shares is 0.189002; 0.1889 would give 5.293806 rather than 5.290949, a difference of 0.002857 in the effective count. **This paper uses 0.189002 throughout and prints the discrepancy rather than harmonising it silently.**

Second, the surplus reading. The source prints $h = -0.1226$ for the designed structure in three separate places. The exact value at $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ is -0.122417, a difference of 0.000183. **This paper uses -0.122417, which is the value Curzi [2026] publishes**, and Curzi [2026] records the same discrepancy independently.

Third, the dollar share required to reach four clearing systems. The source states "roughly one third". The computed figure is 0.389912. Section 7.5 gives the derivation. **This is the one discrepancy in the three that changes a design conclusion**, because a target of 33 per cent and a target of 39 per cent are different programmes.

Fourth, and for consistency across the series, the naive index of the conventional four-jurisdiction structure. Curzi [2026] reports 0.293388 from exact ninths, giving 3.4085 effective jurisdictions against bloc-HHI 1.000000 and 1.0000 effective perimeters; Curzi [2026] reports 0.293430 from rounded percentages, giving 3.407968. **This paper uses the rounded-percentage figure where it refers to that structure**, and states which it uses.

7.9 Hypotheses tested

#	hypothesis	verdict	where
H1	the Atlantic cluster is one perimeter: naive 4.003203, effective 1.000000	**supported**	7.1
H2	the substance instrument repriced rather than prohibited, and is regressive	**supported on both halves**	7.2
H3	two enclaves in one country are one line at 29.1 per cent	**supported**, and trivially	7.3
H4	the settlement axis admits genuine diversification: 1.00 to 2.810568	**supported on the terminal-clearing layer only**; the messaging decomposition its own falsifier names was not run	7.4
H5	a marginal unit is worth at most exactly one effective unit	**supported** as arithmetic, with no empirical content	7.5

H4 is the one to read carefully. It was stated with a falsifier — a decomposition showing the added systems' messaging legs run over the first — and that decomposition was not performed in this paper. **A hypothesis whose falsifier has not been tested is not a hypothesis that has passed**, and the honest verdict is that H4 is supported on the layer this paper measures and open on the layer it does not. Curzi [2026] closes it in the adverse direction.

8. Failure modes

Twelve ways the results above are misapplied, each with the diagnostic that detects it.

1. Counting four Atlantic jurisdictions as four. The arithmetic gives 3.92 naive and 1.00 in blocs. *Diagnostic:* every hour spent rebalancing between them moves a number that does not enter the surplus law.

2. Using a single-purpose offshore company as the apex. It has no succession machinery, no founder-control statute, no purpose-vehicle flexibility and no answer to a forced-heirship claim. *Diagnostic:* score the vehicle on every dimension the job requires and take the minimum, not the mean.

3. Assuming a delisting is a reversal. A jurisdiction came off one list in 2020 and another in 2023; the substance regime the listing produced is permanent, and the onboarding classifications that changed during the listing did not all change back. *Diagnostic:* ask what was installed, not what was removed.

4. Substance theatre. A nominal local director who attends nothing, a serviced desk nobody uses, and a minute book written elsewhere. *Diagnostic:* the test is applied to core income-generating activity and the evidence is the decision trail, not the address.

5. Treating sovereignty as bloc independence. Section 4.2. *Diagnostic:* independence removes an external legislative power; it removes nothing that appears in the measurement.

6. Assuming a treaty benefit survives without substance. A principal-purpose test, domestic anti-abuse rules and beneficial-ownership requirements in the treaties themselves each

independently defeat a conduit. *Diagnostic:* state the commercial reason for the entity that is not the treaty.

7. Relying on a cell ring-fence across a border. Cell segregation is statutory and territorial. *Diagnostic:* name the forum that would have to recognise it.

8. Ignoring correspondent dependency. Three of the four Atlantic jurisdictions peg to the dollar and all four clear through the same correspondents. *Diagnostic:* a de-risking decision taken in one compliance department reaches all four on the same day, and that single channel is a larger practical risk than any listing has been.

9. Registered-agent concentration. A very large number of entities in these jurisdictions are administered by a small number of service providers. *Diagnostic:* this is a concentration in a variable nobody measures, and the instrument of Section 3.1 applies to it directly and has not been run.

10. Keeping an entity below the economic threshold. Section 7.2 gives roughly USD 7.7 million against a 25-basis-point tolerance. *Diagnostic:* compute $Q^{*} = \text{cost}/\text{tolerance}$ with your own tolerance and compare.

11. Believing a second enclave adds a bloc. Section 7.3. *Diagnostic:* ask what single decision would have to be taken twice for the two to be independent, and name the second decision-maker.

12. Reading the clearing-system count as rail independence. Section 7.4. *Diagnostic:* count the messaging layer separately, and if the second system's messaging leg runs over the first, the count collapses.

9. Discussion

What follows from the results. Three things.

The first is that **entity count and effective count can move in opposite directions, and the direction that matters is the one nobody reports.** Six entities produced an effective bloc count of one; a different six produced 2.941176. The count that changed was not the count of entities. Curzi [2026] reaches the sharper version of the same finding — adding two Gulf jurisdictions raises the naive effective count from 5.29 to 7.13, a 34.8 per cent improvement, while the effective bloc count falls from 2.94 to 2.84, a 3.4 per cent deterioration — and it is worth stating that the sharper version is the more common case. **More jurisdictions, less independence, is not a paradox; it is what a naive index does when the units are correlated.**

The second is that **a regulatory instrument with a fixed per-entity cost is a distributional instrument whether or not anyone intended it to be.** The substance regime destroyed the small holder's vehicle and left the large holder's untouched, and the arithmetic of that is one division. **The general lesson is that any compliance requirement expressed as a fixed annual amount is a threshold on holding size**, and the threshold is computable before the requirement is enacted.

The third is the proposition of Section 3.4, and it is the one this paper would defend hardest. **A marginal unit is worth at most exactly one effective unit.** The consequence is that a design conversation about raising a count has an arithmetic ceiling that does not depend on argument, and the ceiling is set by the incumbents. On the constructed settlement mix that ceiling is 4.799250. **A structure that cannot reach five clearing systems should not be planned as though it could, and the way to find out is one line of arithmetic rather than a year of negotiation.**

A ceiling on a count is not the same thing as a ceiling on the value of raising it, and Curzi [2026] is the paper that keeps the two apart. It reports that diversification through a conversion raises effective properties from 1.0000 to 18.8491 — a property HHI of 0.053053, an eighteen-fold improvement in the count — worth 0.4735 per cent a year against the same 15.0417 per cent leakage, giving a payback of 31.77 years, and that only a non-recognition regime cutting leakage to 1.5 per cent brings the payback to 3.17 years. **An eighteen-fold rise in an effective**

count bought a thirty-two-year payback. The proposition of Section 3.4 bounds the numerator of that calculation and says nothing about the denominator, and the two must be reported together or not at all.

What does not follow.

Nothing here is a legal conclusion. Every jurisdictional, doctrinal and regulatory statement in Sections 4, 5 and 6 is described from the source manuscript and is uncited; Appendix B lists them. **A reader who takes a rule of law from this paper has taken it from the wrong document.**

The clearing-system count is not rail independence. Section 7.4 and Section 11 both state it and it bears a third statement, because the number 2.81 is the most quotable figure in the paper and the most likely to be quoted without its layer.

The effective count measures independence, not capacity. It says nothing about whether any bloc, correspondent or clearing system can carry the load, and a structure that is diversified on jurisdiction and concentrated on counterparty has moved its risk rather than reduced it.

The measurement is least informative exactly where the headline case sits. Golosov [2010] and Dunleavy and Boucek [2003] document that effective-number measures behave poorly in the dominated-unit regime, where one unit holds a very large share. A bloc vector of a single unit at 100 per cent is the limiting case. The measurement is exact — $1/1 = 1$ — but its discriminating power is at its minimum there.

A count that can be inflated will be. Goodhart [1975] and Campbell [1979] are the standing warnings and they bite hardest on an instrument that has just been shown to be gameable in the naive direction. Douceur [2002] is the formal statement that one party can present as many, so an unaudited count is worthless, and Lamport et al. [1982] is the bound on how many genuinely independent parties a protocol requires. **Both belong next to any count of units in this series**, and failure mode 9 — registered-agent concentration — is exactly the Sybil problem in a corporate-services costume, unmeasured.

The framework's own results bound this one. Curzi [2026] records that a commons case is formally tunable and practically unreachable: it passes the ceiling test with margin $+0.058436$ nats yet no lever closes it, and a member's declining to participate improves the family reading by 2.680588 octaves — **an arithmetic the framework prints rather than avoids**. The structure measured here is in the same position: it clears no threshold, its levers are enumerated, and the enumeration is what the paper delivers. Curzi [2026] states the companion result that $f_{\min}$ contains neither δ nor N — retention loss sets the rate of arrival but not the threshold, and the category count enters only through the admissibility ceiling $c/\eta \le \ln N$ — which is the general form of the observation that the dollar floor of Section 7.6 does not depend on how many entities exist.

And the corpus's central limitation applies in full. Curzi [2026k] reports the five-parameter survey reproducing each system's observed decline from one desk to a 22-household block, predicted -0.250474 against observed -0.250 at the desk and -0.996371 against -1.000 at the library, with residuals under 1.4 per cent of the ceiling term. **This paper reports no fit at all, and that is not a strength.** Epstein [2008] is the standing statement that prediction is one of sixteen reasons to model; this paper claims **explanation of a mechanism and demonstration that a stated arithmetic suffices to produce a stated shape**, and claims neither prediction nor calibration.

10. Verification and rigour audit

The results are decomposed into postulates that can be checked one at a time. Each is atomic, each carries a status, and the compounding chain names what each step consumes. **If dropping a postulate leaves the result standing, it is not listed.**

10.1 Dimensional audit

symbol	quantity	dimension	unit	status
s_i, w_i	exposure share of unit i	$[1]$	dimensionless	defined; non-negative and summing to 1 by P1
W_b	bloc share	$[1]$	dimensionless	derived, $W_b = \sum_{i \in b} s_i$
HHI	concentration index	$[1]$	dimensionless, or points $= 10^4 \times$ fraction	derived
ρ_{ij}	correlation between shocks	$[1]$	dimensionless, in $[-1, 1]$	**assumed, not measured**
N	effective count	$[\#]$, a count	jurisdictions, blocs, or clearing systems	derived as $1/(w' R w)$
M	incumbent effective count	$[\#]$	clearing systems	derived, Section 3.4
x	share of the marginal unit	$[1]$	dimensionless	free variable
$\ln N$	surplus-term argument	$[1]$	nats	derived; the logarithm of a count is dimensionless
Q	protected capital	$[C]$	index units	defined; **the source supplies no operational unit**
η	conversion efficiency	$[1]$	dimensionless	assumed
δ	erosion rate	$[T^{-1}]$	per period	assumed
C	maintenance cost flow	$[C\,T^{-1}]$	index units per period	assumed
h	headroom	$[C\,T^{-1}]$	index units per period	derived
substance cost	fixed annual charge	$[\$\,T^{-1}]$	USD per year	constructed
Q^*	break-even holding	$[\$]$	USD	derived, cost divided by tolerance
tolerance	admissible overhead rate	$[T^{-1}]$	basis points per year	**a stated judgement**

The surplus equation balances only if $\eta \ln N$ carries $[C\,T^{-1}]$. Since $\ln N$ is dimensionless, either η carries $[C\,T^{-1}]$ or $\ln N$ is read as a rate. **The source is silent on which and this paper records the silence rather than resolving it by assertion.**

The break-even equation balances exactly: $[\$\,T^{-1}]$ divided by $[T^{-1}]$ gives $[\$]$, so a fixed annual cost divided by an annual tolerance is a capital amount. That is the one equation in this paper whose dimensions are fully determined, and it is worth naming which one.

The marginal-unit proposition is dimensionless throughout. x, r and M are all pure numbers, so EFF is a count, and the identity $\text{EFF}(x^*) = M + 1$ adds a count to a count, which is

licensed.

10.2 Significant figures and epistemic status

quantity	figures	status	licence
naive HHI 0.249800	6 dp	**exact**	exact arithmetic on exact constructed shares
effective jurisdictions 4.003203	6 dp	**exact to rounding**	reciprocal of an exact value
bloc-HHI 1.000000, $N = 1.000000$	6 dp	**exact**	a single block of ones
multiples 3.45 and 12.48	2 dp	**exact to rounding**	ratios of exact constructed amounts
96.75 and 0.9675 basis points	4 sf	**exact to rounding**	ratios of exact constructed amounts
break-even USD 7,740,000	6 sf	**exact** given the tolerance	a division by a stated judgement
HHI 0.3558, 2.810568	4 dp, 6 dp	**exact to rounding**	exact arithmetic on exact constructed shares
$r = 0.263210$, $M = 3.799250$	6 dp	**exact to rounding**	closed form on the same shares
$x^* = 0.208366$, ceiling 4.799250	6 dp	**exact**	closed form, Section 3.4
$x = 0.389912$ for EFF $= 4$	6 dp	**exact to rounding**	root of a quadratic
all USD cost figures	as printed	**constructed**	illustrative, not quotations

No figure in this paper is a measured quantity and therefore no uncertainty interval is attached to any of them. That is not modesty; it is the correct treatment of exact arithmetic on stipulated inputs, and it is also the reason none of these numbers is evidence about a real structure.

10.3 Limits and boundary behaviour

limit	behaviour	reading
$\rho \to 0$ within a bloc	$N \to$ the naive count	the naive index is the zero-correlation endpoint
$\rho \to 1$ within a bloc	$N \to$ the bloc count	the bloc index is the unit-correlation endpoint
one bloc holding everything	$N = 1$, $\ln N = 0$, surplus term zero	**the whole argument in one line**
$N \to \infty$	$h \to \infty$ logarithmically; $dh/dN = \eta/N \to 0$	returns diminish but never turn negative — **a defect, since cost is not in the input term**

limit	behaviour	reading
$x \to 0$ in the marginal-unit proposition	EFF $\to M$	removing the unit returns the incumbent count
$x \to 1$	EFF $\to 1$	one unit holding everything
$M \to 1$ (one incumbent)	$x^* \to 1/2$, ceiling $\to 2$	the two-equal-units case, which is correct
$M \to \infty$	$x^* \to 0$, ceiling $\to M + 1$	a marginal unit added to a very diverse set must be very small to help
tolerance $\to 0$	$Q^* \to \infty$	no holding is large enough; the entity is never economic
tolerance $\to \infty$	$Q^* \to 0$	any holding is economic, which is why the tolerance must be stated
substance cost $\to 0$	$Q^* \to 0$ and the regressivity vanishes	the regressive property is a property of the fixed cost, not of the rule

Two boundary readings are load-bearing. The fourth row is a genuine defect of the surplus law as used here: because cost enters only through C, the model never says that adding a bloc is not worth it, and the decision to stop must come from outside the equation. Brooks [1975] supplies one bound — communication cost at order $n(n-1)/2$ inside a boundary — and **the proposition of Section 3.4 supplies another and sharper one on the benefit side**, which is the reason it is in this paper rather than in a note.

The last row is the honest reading of H2. **The regressivity is a property of the cost being fixed, not of the substance rule being unfair**, and any instrument charging a flat annual amount has the same property. That is a stronger and more general statement than the one the source manuscript makes.

10.4 Invariance

Scale invariance in capital. Multiplying every exposure by a positive constant leaves every share and every effective count unchanged. Tested at three scales by construction: the six-entity vector, the five-system vector and the four-system incumbent vector. **Three scales is not many and Section 11 says so** rather than describing the instrument as scale-free without qualification.

The break-even is not scale-invariant, and that is the finding. Q^* is an absolute amount because the numerator is an absolute amount. **A rule whose cost does not scale produces a threshold, and a threshold is exactly what scale invariance would have hidden.**

Invariance under relabelling. The effective count is invariant under permutation of the units. That is wanted — a concentration measure that changed when the rows were reordered would be broken — but it means the count cannot distinguish two structures with the same share vector and different contents. Curzi [2026] identifies the same invariance as fatal in its own domain, where the width statistic is invariant under relabelling and therefore does not determine the chord it is computed from.

Not invariant, and this is where the instrument works. The count is *not* invariant under partition: splitting one unit into two raises the naive count and leaves the bloc count untouched. **That asymmetry is the whole instrument.**

The marginal-unit proposition is invariant under the incumbents' internal composition. It depends on the incumbents only through M, so two quite different incumbent mixes with the same M have the same ceiling and the same optimal marginal share. **That is a real generality**

and it is the reason the proposition is stated in terms of M rather than in terms of the four particular clearing systems.

10.5 Analogy discipline

borrowing	category	justification
exposure as a portfolio; variance under a correlation matrix [Markowitz, 1952]	**1 — derivation**	the exposure vector genuinely is a weighted sum of shocks with a correlation matrix, which is the premise the algebra requires
effective count as a numbers-equivalent [Adelman, 1969, Hill, 1973, Jost, 2006]	**1 — derivation**	$1/\mathrm{HHI}$ *is* the Hill number of order two
similarity-sensitive diversity [Leinster and Cobbold, 2012, Rao, 1982]	**1 — derivation**	the block-of-ones matrix is a special case of the cited general form
open design [Kerckhoffs, 1883, Saltzer and Schroeder, 1975]	**1 — derivation**	the disclosure standard *is* the open-design requirement applied to legal rather than cryptographic structure
coupled-network fragility [Buldyrev et al., 2010]	**2 — shared form**	perimeter coupling is not shown to satisfy the percolation premises; only the shape of the result is used
robust-yet-fragile [Albert et al., 2000, Callaway et al., 2000]	**2 — shared form**	the tolerance asymmetry's shape is used; the degree-distribution mechanism is not claimed, and Doyle et al. [2005] is the reason it should not be
$\ln N$ as the entropy of a uniform distribution [Shannon, 1948]	**2 — shared form**	**information entropy, not thermodynamic entropy**; the correspondence is exact as arithmetic and carries no thermodynamic content

borrowing	category	justification
a bloc as a frozen region [Castellano et al., 2000]	**3 — analogy**	suggestive, used as motivation only; no phase-transition claim is made and no order parameter is computed
normal accidents [Perrow, 1984]	**3 — analogy**	qualitative theory and framing, not a formal result; this paper's first finding is Perrow's in a measurement frame
the marginal-unit proposition	**not a borrowing**	elementary calculus on this paper's own definition; it is proved in Section 3.4 and consumes nothing external

No category-3 borrowing is used as evidence anywhere in this paper.

10.6 Measurement-theoretic status of the scored inputs

The four-question test of Section 4.1 is nominal and is not scored. It returns yes or no per jurisdiction per question, and the assignment is made by requiring all four rather than by counting them. **No total is computed**, deliberately, because a count of yeses would invite a threshold rule the scale does not license.

The evidence ladder of Section 4.1 is ordinal and is not summed. Its classes are ranked and the assignment identifies the strongest available class, which an ordinal scale licenses.

The de-identified capability comparison of Appendix C is not totalled. Court depth, statute book, treaty network and cost have no common scale and the table produces a best fit by argument rather than a verdict by arithmetic.

The 25-basis-point tolerance is a stipulated cut on a ratio scale, and the ratio is genuine — basis points of capital are a ratio-scale quantity with a true zero — so the division producing Q^* is licensed. **What is not licensed is treating 25 as anything other than a choice**, and Section 7.2 gives the formula so a reader may substitute another.

The framework's treatment of Q as a ratio-scale quantity is not tested. Michell [1997] is the strongest available objection: a discipline routinely assumes quantitative structure without testing for it, and the formal remedy is conjoint measurement [Luce and Tukey, 1964], whose conditions establish whether an attribute has additive quantitative structure at all. **The conjoint conditions were not tested for Q, which is stated in unnamed index units, and the ratio-scale treatment is provisional.** Krantz et al. [1971] sets out what would have to hold and Stevens [1946] is the taxonomy. **The concentration instrument is untouched by this objection**, because exposure shares are ratios of capital rather than judgements.

This paper inherits the framework-level scalar-reduction limitation rather than restating it. Reducing a structure to one effective count and one h, and reducing a jurisdiction's regime to a four-yes test, are instances of the habit Curzi [2026] names once as a single limitation for the whole framework: where the source literature describes a phenomenon by several jointly varying features, the framework replaces it with one number and does arithmetic on the number. **The limitation is inherited here in full and is not met separately.**

10.7 Postulates

#	statement	status
P1	exposure shares s_i are non-negative and sum to 1	DEFINITION
P2	HHI $= \sum_i s_i^2$	DEFINITION
P3	for n equal shares of $1/n$, HHI $= 1/n$	ARITHMETIC
P4	the aggregate shock is $R = \sum_i w_i X_i$ with Var$(X_i) = 1$	DEFINITION
P5	Var$(R) = \sum_i \sum_j w_i w_j \rho_{ij} = w' R w$	ARITHMETIC
P6	units inside one coordination perimeter act on one decision	POSTULATED (H1, H3)
P7	the surplus law's input term is $\eta \ln N$	DEFINITION (modelling choice, CCP-A02)
P8	$\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$	POSTULATED (fixed in advance)
P9	the substance cost is a fixed annual amount per entity, independent of assets held	POSTULATED — described from the source, uncited (Appendix B)
P10	the constructed cost table of Section 5.2	CONSTRUCTED
P11	a 25-basis-point tolerance for structural overhead	POSTULATED (a stated judgement)
P12	the share basis on the rail axis is the terminal clearing system	DEFINITION
P13	the constructed settlement mix USD 55 / CNY 15 / EUR 12 / INR 10 / AED 8	CONSTRUCTED

10.8 The compounding chain

step	consumes	statement	status
L1	P2, P3	HHI^{-1} returns n exactly for n equal shares	ARITHMETIC
L2	L1	HHI^{-1} is a numbers-equivalent — Adelman [1969]	CITED
L3	L2	it is the Hill number of order two — Hill [1973], Jost [2006]	CITED
L4	P4, P5	with $R = I$, Var$(R) = \sum w_i^2 =$ HHI	ARITHMETIC
L5	P4, P5	with R block-of-ones, Var$(R) = \sum_b W_b^2$	ARITHMETIC

step	consumes	statement	status
L6	L4, L5	naive and bloc indices are one variance under two correlation assumptions	ARITHMETIC
L7	L1	EFF$(x) = 1/(x^2 + (1-x)^2/M)$ is maximised at $x = 1/(M+1)$ with value $M+1$	ARITHMETIC — the proposition of Section 3.4
R1	L6, P6	six legacy entities: naive 0.249800 (4.003203); bloc 1.000000 ($N = 1.000000$), ln $N = 0$	ARITHMETIC
R1b	R1, P7, P8	$h = -0.500000$	ARITHMETIC
R2a	P10	5,350 over 1,550 is 3.45; 19,350 over 1,550 is 12.48	ARITHMETIC on CONSTRUCTED inputs
R2b	P9, P10, P11	$Q^* = 19{,}350/0.0025 = 7{,}740{,}000$; the cost falls from 96.75 bp to 0.9675 bp over a hundred-fold range	ARITHMETIC on POSTULATED inputs
R4	L6, P6	the two enclaves are one 29.1 per cent line; bloc-HHI 0.340000, $N = 2.941176$, invariant under rearrangement between them	ARITHMETIC
R3	P12, P13	HHI $= 0.3558$; effective clearing systems $= 2.810568$ against 1.0000 as found	ARITHMETIC on CONSTRUCTED inputs
R6	L7, P13	$M = 3.799250$; ceiling 4.799250 at $x^* = 0.208366$; EFF $= 4$ at $x = 0.389912$	ARITHMETIC
R5	R3, P13	55 per cent of settled value clears through one system, a floor set by counterparties' currency	ARITHMETIC on CONSTRUCTED inputs
CR	R4, R5	a second enclave buys capability, not diversification; and 55 per cent of settlement remains dollar-cleared however the entities are arranged	ARITHMETIC

Every ARITHMETIC row can be redone by hand or in three lines of any language. Every CITED row names one source and no more. Only P6, P8, P9 and P11 are POSTULATED, and two of the four are stated judgements rather than claims about the world. **P9 is the one carrying regulatory content and it has no bibliographic support** — that the substance cost is fixed per entity rather than proportional to assets is described from the source manuscript and is

listed in Appendix B. **A reviewer who does not accept P9 loses R2b and keeps R2a**, because the multiples are arithmetic on the cost table whatever the cost's relation to assets.

L7 is the only step in this chain that is new to this paper, and it is elementary. It consumes no citation, no postulate and no constructed input; it is a one-line minimisation. **A chain whose novel step is a line of calculus is a chain a reviewer can dispose of in a minute, and that is the intention.**

10.9 Cross-checks performed

- The six legacy squares were summed independently and reproduced 0.249800, and the four-entity variant reproduced 0.255000.
- $s_{\max}^2 \le \text{HHI} \le s_{\max}$ holds on both: $0.160000 \le 0.249800 \le 0.400000$ and $0.090000 \le 0.255000 \le 0.300000$.
- The effective-count ceiling $s_{\max}^{-2}$ is respected: $4.003203 \le 6.250000$ and $3.921569 \le 11.111111$.
- The cost table sums were checked in both directions: $1{,}600+1{,}050+900+400+1{,}400 = 5{,}350$ and $5{,}350 + 14{,}000 = 19{,}350$.
- The break-even was checked by substitution: $19{,}350/7{,}740{,}000 = 0.0025$ exactly.
- The clearing squares were summed independently and reproduced 0.3558.
- The marginal-unit proposition was checked three ways: at the closed form $x^* = 1/(M+1) = 1/4.799250 = 0.208366$; by direct evaluation, giving $\text{EFF}(0.208366) = 4.799250$; and by the identity $\text{EFF}(x^*) = M + 1$, $3.799250 + 1 = 4.799250$.
- The root $x = 0.389912$ was checked by substitution: $\text{EFF}(0.389912) = 4.000000$ to six decimal places.
- The naive-count improvement was checked: $5.290949/4.003203 - 1 = 0.321661$, printed as 32.16 per cent.
- The three source discrepancies of Section 7.8 were quantified at 0.000102 in the naive index, 0.000183 in h, and 5.66 percentage points in the dollar-share target.

11. Limitations

Every number in this paper is computed on constructed inputs and none is a measurement of a real structure. The exposure shares, the cost table, the maintenance budget, the settlement mix and the listing durations are all **CONSTRUCTED** or described from the source manuscript, and are labelled as such at each point of use. **No first-party measurement of any real system appears anywhere in this paper.** That is this corpus's central limitation, stated in the terms Curzi [2026] states it in — of 27 parameter-by-scope cells, 19 have a defined estimator, 1 has a validated one, and 0 carry a first-party measurement — and a reviewer would be right to press on it before pressing on anything else.

The clearing-system count is a count of one layer. Section 3.2 states the distinction and Section 7.4 states the consequence: the falsifier H4 named was a messaging decomposition, and that decomposition was not run. **The number 2.810568 is an upper reading and should never be quoted as a measure of rail independence.** Curzi [2026] runs the decomposition and finds $N_{\text{rail}} = 1.341819$. **This is the paper's largest single exposure and it is disclosed here rather than in a note.**

The correlation is assumed, not measured. $\rho = 1$ within a perimeter is the conservative end of a family and no correlation was estimated from data. The listing record of Section 4.3 is the evidence offered for it, and a listing record is a record of outputs rather than a measurement of correlation.

Every parameter in the surplus law was chosen. $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ are test parameters fixed in advance and not estimates of any real position, so the break-even $N = 4.172734$ is a property of the chosen set.

The 25-basis-point tolerance is a judgement and the break-even moves with it. At 50 basis points the threshold is USD 3,870,000 and at 10 basis points it is USD 19,350,000. **A threshold**

quoted without its tolerance is not a threshold, and Section 7.2 gives the formula for exactly that reason.

The marginal-unit proposition is arithmetic and carries no empirical content. It bounds a design space; it says nothing about whether any particular unit exists, can be onboarded, or would be independent in the sense the correlation assumption requires. **A ceiling on what a count could be is not evidence about what a structure is**, and the proposition should not be cited as though it were.

The two listing rows marked "dates to be verified" are reproduced with the mark. They are uncertain in the source and they are uncertain here. **No conclusion in this paper depends on either of them**, because H1 is carried by the rows that are dated and by the four-question test independently of any listing.

The registered-agent concentration of failure mode 9 is named and not measured. It is a concentration on the counterparty axis to which the instrument of Section 3.1 applies directly, and it is not run here for want of any share basis that could be constructed honestly. **That is a gap in the four-axis discipline this series states, in this paper.**

The reflexive specification is not identified. Section 1 states this and names the remedy — stipulate the reference group and defend the stipulation, following Brock and Durlauf [2001]. **The price is that every bloc assignment here is a modelling position rather than an identified empirical result** [Manski, 1993, Shalizi and Thomas, 2011].

The Lucas critique applies to this framework's own parameters and is not answered. Lucas [1976] is the general statement; the specific form is that this framework asserts systems change behaviour when modelled and observed, and then treats η, δ and the bloc assignment as stable under exactly that response. MacKenzie and Millo [2003] documents the loop happening in the field with dates and price data, which is an existence proof for the mechanism and not a defence of the parameters.

The functional form of the surplus input is underdetermined. Latané [1981] proposes a power law with exponent below one and Peters and Gell-Mann [2016] derives a logarithmic form from multiplicative dynamics with no psychological assumption at all. **Nothing here discriminates between them, and the logarithm is used because it is the form this framework carries.** Peters is the stronger of the two available warrants and is cited as such.

Small-sample bias in the effective count is not corrected [Chao et al., 2014]. A five- or six-element vector is a small sample for a low-order diversity estimator, and the paper does not resolve whether these vectors are samples or complete censuses.

Scale invariance was checked at three scales only (Section 10.4).

The instrument is least informative in the regime where the headline case sits [Golosov, 2010, Dunleavy and Boucek, 2003].

No statistical test was run and no p-value appears, because there is no sample and no null. Every rate is an exact ratio of exact constructed amounts, and **no claim of "no difference" is made from any non-significant result** because no significance test exists here to make it from.

No scaling, power-law or heavy-tail claim is made and the Clauset–Shalizi–Newman procedure was therefore not run [Clauset et al., 2009]. Any restatement reaching for such a claim must run it first and read Broido and Clauset [2019].

No out-of-sample validation, matched null, information criterion or effect size with an interval appears, because no model was fitted. Bruch and Atwell [2015]'s calibration-and-validation checklist is unmet in every row.

11.1 Rigour checks not met, listed by name

1. **The messaging-layer decomposition H4's falsifier names was not run** (Sections 7.4, 7.9); Curzi [2026] runs it and finds against.
2. **Conjoint measurement conditions untested** [Luce and Tukey, 1964], so the ratio-scale treatment of Q is provisional and Michell [1997]'s objection is conceded rather than answered.

3. **The framework-level scalar-reduction limitation is inherited** [Curzi, 2026] and not met separately.
4. **The reflection problem is stated and not solved** [Manski, 1993].
5. **The Lucas critique applied to this framework's own parameters is unanswered** [Lucas, 1976].
6. **The functional form of the surplus input is undiscriminated** [Latané, 1981, Peters and Gell-Mann, 2016].
7. **Small-sample bias in the effective count is uncorrected** [Chao et al., 2014].
8. **No out-of-sample test, matched null, information criterion or interval** [Bruch and Atwell, 2015].
9. **No first-party measurement of any real system** [Curzi, 2026].
10. **The dimension of η is unresolved** because the source is silent (Section 10.1).
11. **The counterparty axis is not measured** despite being one of the four the method requires (failure mode 9).
12. **Two listing rows are undated** and are reproduced with the source's own mark.
13. **Every legal, regulatory and doctrinal statement in Sections 4, 5 and 6 is uncited** (Appendix B).

12. Conclusion

This paper measures four things on constructed inputs and establishes one proposition. The following statements are the fixed text of those results.

R1. Six Atlantic-cluster entities — United Kingdom 40.0 per cent, United States 22.0, Cayman 14.0, the British Virgin Islands 12.0, the Bahamas 7.0 and Barbados 5.0 — give a naive HHI of 0.249800 and 4.003203 effective jurisdictions, but a bloc-HHI of 1.000000, $N = 1.000000$ and $\ln N = 0$. **Adding a Western entity to a Western structure leaves the surplus term at exactly zero**, and no admissible reassignment separates the six, including on the ground of sovereignty.

R2. The substance instrument is measured as a cost multiple rather than a prohibition: $5{,}350/1{,}550 = 3.45$ for a pure equity-holding entity and $19{,}350/1{,}550 = 12.48$ for a relevant-activity entity, with break-even at $Q^* = 19{,}350/0.0025 =$ **USD 7,740,000 of held value at a 25-basis-point tolerance**. Because the cost is fixed per entity, it falls from 96.75 basis points at USD 2,000,000 to 0.9675 basis points at USD 200,000,000. **The regressivity is a property of any flat annual charge and not of this rule in particular.**

R3. The clearing-system count: a structure settling wholly in one currency gives HHI $= 1.0000$ and 1.00 effective clearing systems; the constructed redesign at USD 55.0 / renminbi 15.0 / euro 12.0 / rupee 10.0 / dirham 8.0 gives HHI $= 0.3558$ and 2.810568 **effective clearing systems**. The count is of terminal clearing systems only and is an upper reading on rail independence.

R4. The DIFC and the ADGM are one FATF perimeter, one federal anti-money-laundering regime and **one line at 29.1 per cent of exposure — no rearrangement between the two free zones changes** N, which remains 2.941176 at bloc-HHI 0.340000.

R5. The residual after redesign: **55 per cent of settled value still clears through a single system whose law is not the holder's, a floor set by the counterparties' currency rather than by entity placement.**

Counter-result. A second financial free zone in the same country buys capability, not diversification: the two zones together remain one 29.1 per cent line, and 55 per cent of settlement remains dollar-cleared however the entities are arranged.

One further result is established as method rather than as measurement, and it is the one this paper would defend hardest. If a set of units held in fixed relative proportions has its own effective count M, then adding one unit at share x and rescaling the rest gives $\mathrm{EFF}(x) = 1/(x^2 + (1-x)^2/M)$, which is maximised at $x^* = 1/(M+1)$ with value $M+1$. **A marginal unit is worth at most exactly one effective unit, and only at one share.** Applied to the constructed settlement mix, where the four non-dollar systems give $M = 3.799250$: **the ceiling on the effective clearing-**

system count is 4.799250**, attained at a dollar share of** 0.208366**; reaching** 4.000000 **requires the dollar share to fall to** 0.389912**, not to the "roughly one third" the source manuscript states; and eliminating the dollar entirely would give** 3.799250**, which is worse than holding a fifth of settled value in it.** All three are exact and all three are checkable in two lines.

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every case in this paper — the exposure shares, the cost table, the maintenance budget, the settlement mix and the entity dispositions — is **CONSTRUCTED** and is labelled as such at each point of use. The numerical results are exact computations on those constructed inputs, carried at full precision and rounded only for printing.

All figures were produced by the scripts named in their captions (`figs/src/CCP-056_fig01.py` through `figs/src/CCP-056_fig05.py`), which import the shared figure style module and read their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

Jurisdictional and regulatory descriptions in this paper are taken from the source manuscript and not from primary legal sources. The bibliography contains general comparative, treaty and standard-setting material only; it contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction, and no listing or delisting instrument for any jurisdiction. Appendix B lists the claims affected. **Any reader relying on those descriptions must verify them against the official sources and record the date of that verification.** Two listing rows in Section 4.3 are marked uncertain in the source and are reproduced with that mark.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. All persons, entities, holdings and arrangements described are constructed for the purpose of exercising the measurement.

Preregistration

The parameter set $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ and the break-even criterion $N \geq e^{0.5/0.35} = 4.172734$ were fixed by Curzi [2026] before the measurements reported here and are not adjusted anywhere in this paper. **No structure measured in this paper clears that**

break-even, which Section 7.7 reports in the same language a clearance would have received. The 25-basis-point tolerance of Section 7.2 and the falsifier attached to H4 were both stated before the corresponding measurements were run; **H4's falsifier was not tested, and Section 7.9 reports that rather than treating the hypothesis as passed.** No other element of this paper was preregistered.

Legal, tax and investment advice

This paper is analytical and is not legal, tax, or investment advice. It describes entity types, jurisdictions, licensing regimes, substance requirements and settlement arrangements for the purpose of measuring concentration. No recommendation is made in respect of any particular person, asset, structure, jurisdiction or arrangement. Substance requirements, relevant-activity definitions, filing obligations, beneficial-ownership access rules, listing and delisting decisions, corporate-tax qualifying-income tests and clearing-system participation all change and vary by jurisdiction and by activity; **every such statement in this paper is carried from the source manuscript, is flagged in Appendix B, and requires confirmation with qualified counsel in the relevant jurisdiction before any reliance.**

Appendix A. The reporting standard for the rail axis

The jurisdiction-axis reporting standard is stated by Curzi [2026] and extended by Curzi [2026]. This appendix states what the rail axis additionally requires, because Section 7.4's result is the one most likely to be reproduced without its layer.

```
 1  share basis, in one sentence: settled value by TERMINAL
    clearing system, over a stated period
 2  the settlement table, systems and shares, summing to 100.0%
 3  naive HHI and effective clearing systems
 4  THE SAME THREE LINES for the MESSAGING layer, separately
 5  for every system whose messaging leg runs over another system,
    the share so affected, itemised
 6  the collapsed count after step 5, which is the count that
    should be quoted
 7  count of banking relationships, and of independent
    correspondent chains per currency
 8  standing relationships that are live but dormant, listed
 9  the incumbent effective count M, the marginal share x, and
    the ceiling M + 1
10  the date, and the diff against the previous measurement
```

Item 6 is the item this paper does not reach, and item 9 is the item it adds. A measurement stopping at item 3 has reported the flattering layer; a measurement reaching item 6 has reported the structure. **Item 9 costs one line of arithmetic and tells a reader whether the programme they are contemplating is achievable at all.**

Item 8 has no effect on any index and belongs in the report anyway. A dormant relationship is an option on a branch, and Curzi [2026]'s weighting — score each step 0 to 3 for steepness, sum to a resistance R, weight the branch as $1/2^{R}$ — is the instrument for valuing it. **The index says such a relationship is worth nothing and the branch weighting says a three-point reduction in steepness is a factor of eight; both readings should be printed.**

Appendix B. Claims requiring verification against primary sources

The following statements rest on legal or regulatory detail that changes, and **none is cited to the bibliography, which contains no primary sources for any individual jurisdiction**. They are reproduced from the source manuscript and require confirmation with counsel in the relevant jurisdiction before any reliance. **Record the date of each confirmation**; a statement of a regime without a date is not a statement of that regime.

1. The current revisions of the companies, trusts, funds, foundations, insurance and economic-substance legislation of each of the four Atlantic jurisdictions named in Section 4.1, together with the legislation enacting the convergence and rate changes referred to in Section 5.3.
2. **Every listing and delisting date in Section 4.3**, against the standard-setting body's plenary outcomes and the regional bloc's primary releases. **Two rows are specifically flagged as uncertain in the source and must not be reproduced without checking.**
3. The beneficial-ownership transparency legislation of the two territories referred to in Section 4.5: commencement dates, enforcement dates, the definition of legitimate interest, and the access mechanism.
4. The regional court decision referred to in Section 4.5 striking down public access to member-state registers, and its subsequent treatment.
5. The criteria adopted in December 2017 by which the substance requirement was imposed, the list of relevant activities, and the reduced test for pure equity-holding entities.
6. The restructuring and insolvency regime referred to in Section 7.7's disposition table, including any moratorium and its scope.
7. The appellate decision on freezing relief in support of foreign proceedings referred to in failure mode 12, its ratio and its current application.
8. **P9, that the substance cost is fixed per entity rather than proportional to assets.** This is the postulate on which R2b depends and it has no bibliographic support; Section 10.8 states exactly which result rests on it.
9. The establishing instruments of the two financial free zones referred to in Section 6, their federal basis, the regulations by which one of them applies another state's common law directly, and the amendment history including any expansion of jurisdictional area.
10. The foundations legislation of the enclave referred to in Section 6.1, and the migration path between the two zones' vehicles.
11. The corporate-tax regime applying to free-zone persons in that state, and the qualifying-income test.
12. The federal anti-money-laundering legislation, its implementing regulation, and the beneficial-ownership resolution referred to in Section 4.4.
13. The dates of the increased-monitoring listing and delisting referred to in Section 4.4, the regional delegated act adding the state to a high-risk list, the corresponding removal instrument, and the equivalent instruments of the second jurisdiction operating such a mechanism. **These are dateable regulations and should be cited precisely or not at all.**
14. The current participation of institutions in that state in the second clearing system, the terms and renewal status of the bilateral currency arrangement referred to in Section 7.4, and the status and governance of the cross-border settlement experiment referred to in the source. **All three have moved and all three are frequently misdescribed.**
15. The extent to which participants in the second clearing system use the first system's messaging layer, which is the decomposition H4's falsifier requires and which this paper does not run.
16. The treaty anti-abuse provisions referred to in failure mode 6, including the principal-purpose test and the beneficial-ownership requirements in the relevant treaties.
17. The cell and segregated-portfolio legislation referred to in failure mode 7, and whether the segregation is recognised outside the jurisdiction of formation.

18. All cost, maintenance and settlement figures in this paper are constructed. **They are not quotations, rate cards or bank statements, and no real schedule was used to produce them.**

All numerical results in this paper were computed from the stated inputs at full precision and rounded only for printing, and the places where a printed value in the source differs from the exact computation are reported in Section 7.8 rather than harmonised.

Appendix C. Condensed assessment, de-identified

No total is computed: the items are heterogeneous and no common scale exists on which they could be added (Section 10.6). Every characteristic below is described from the source manuscript and is uncited (Appendix B). The Atlantic jurisdictions are de-identified in the capability table because that table ranks them on legal capability the bibliography cannot support; they are named in Section 4.1's bloc table, because there the assignment is the argument.

Bloc assignment and evidence, by jurisdiction.

jurisdiction	bloc	strongest evidence	job
Cayman Islands	**WEST**	exchange, common standard implementation, listing record, dollar correspondent dependency	**HOLDING**, for genuine fund and large-deal structures only
British Virgin Islands	**WEST**	identical on every operative test; listed mechanically on a third body's rating change	**HOLDING**, single-purpose; and a genuine litigation forum
The Bahamas	**WEST**	four yeses notwithstanding sovereignty; pegged currency	**HOLDING**, for private-client instruments the alternatives lack
Barbados	**WEST**	four yeses; the longest increased-monitoring period of the four	**OPERATING**, treaty-dependent, and only where the corridor is used
United Arab Emirates (DIFC and ADGM)	**BRIDGE**	outside the Western designation apparatus; own listing decisions; dual-rail access	**HOLDING** in one zone, apex in the other; **never settlement**

Ranked capability, de-identified.

requirement	best fit	why
a fund with unrelated investors and institutional service providers	Atlantic centre A	the deepest fund statute book and a specialist financial services division

requirement	best fit	why
a single-purpose holding vehicle with a specialist commercial court behind it	Atlantic centre B	a commercial court whose decisions are followed, and the lowest-friction company form of the four
a novel private-client instrument	Atlantic centre C	the widest private-client statute book, including purpose and executive-entity forms the others lack
a treaty-dependent operating position	Atlantic centre D	the treaty network, and the only one of the four with a different apex court
group agreements governed by a mature commercial law, without that state's entity	enclave A	direct application of that law as it stands from time to time, with its own court
a regulated activity inside the same country	enclave B	the regulator with the deeper regulated-activity practice

The three items that decide every row of the first table are the exchange position, the implementation of the common standard, and the correspondent dependency. **None of the four Atlantic jurisdictions can legislate itself a different one of the three**, which is why the bloc assignment is the least contestable call in the paper and why sovereignty does not move it.

Chapter 57

Diversification That Does Not Pay: Gulf Additions, an Indifferent Bridging Slot, and a Capped Transit Position

Abstract. We test three proposed additions to a designed holding structure against the concentration instrument, and two of them fail on the framework's own arithmetic. Adding Saudi Arabia and Qatar raises the naive effective jurisdiction count from 5.29 to 7.13, a 34.8 per cent improvement, while the effective bloc count falls from 2.94 to 2.84, a 3.4 per cent deterioration, because all three Gulf entries occupy one bloc; the variant actually built, which declines the Qatar Financial Centre vehicle and adds only a Saudi regional headquarters, reads 5.69 naive and 2.82 blocs, again below baseline. The bridging slot is indifferent on this axis: holding 10.9 per cent of exposure in Singapore treated as Western, or in London, gives identical bloc shares of 29.1 / 30.0 / 40.9, N = 2.924302 and h = −0.124430, so the stress case costs only 0.002014 and the Singapore leg must be justified on axes the index does not measure. The Hong Kong comparison is the one that pays: a full centre at 10.9 per cent of exposure reaches N = 3.615826 and h = −0.050138, while a capped transit position at 0.35 per cent reaches N = 2.965384 and h = −0.119548 — capturing 3.97 per cent of the larger design's measured gain for 3.21 per cent of its exposure. The 0.35 per cent figure is derived from a sweep schedule rather than assumed: renminbi receipts at 18.0 per cent of protected capital a year peak at 4.500 per cent quarterly, 1.500 per cent monthly and 0.346 per cent weekly.

Keywords: Gulf financial centres; Singapore; Hong Kong; bridging entity; transit exposure; sweep schedule; bloc independence

Series: D. Structure and Resilience **Paper:** CCP-057

1. Introduction

Three proposed additions to a designed holding structure are tested against the concentration instrument, and two of them fail on the framework's own arithmetic. **The paper's subject is a measurement that comes out against the thing being measured**, which is the only kind of measurement that establishes an instrument is not a formality.

The first addition makes the structure worse. Adding two further Gulf jurisdictions to a structure already holding one raises the naive effective jurisdiction count from 5.290949 to

7.132973, an improvement of 34.8 per cent, while the effective bloc count falls from 2.941176 to 2.841490, a deterioration of 3.4 per cent. **More jurisdictions, less independence.** The mechanism is exact: the additions occupy a bloc the structure already sits in, and they were funded partly out of a position in a *different* bloc, so the bloc the structure was already concentrated in became more concentrated while the jurisdiction count rose. The variant actually built — which declines one of the two proposals and adds only a regional headquarters justified on contracted revenue — reads naive HHI 0.175802 (5.688217 effective jurisdictions) and bloc-HHI 0.355000 (2.816901 effective blocs), **also below the** 2.941176 **baseline it was proposed to improve.**

The second addition is indifferent. Holding 10.9 per cent of exposure in a bridging centre treated as Western, or holding the same 10.9 per cent in the Western centre itself, gives *identical* bloc shares of 29.1 / 30.0 / 40.9, an identical bloc-HHI of 0.341962, an identical effective count of 2.924302 and an identical $h = -0.124430$. **The measured cost of the stress assignment is** 0.002014, and the naive count differs between the two by 0.841640 — an artefact of counting jurisdictions rather than perimeters. **If the bridging centre is Western in the coordination sense, the entire measured value of the move is a naive-count artefact, and the leg has to be justified on axes the index does not measure.**

The third addition is the one that pays, and it pays in a specific and boundable way. A full Asian centre at 10.9 per cent of exposure reaches $N = 3.615826$ and $h = -0.050138$; a capped transit position at 0.35 per cent reaches $N = 2.965384$ and $h = -0.119548$. **The capped position captures 3.97 per cent of the full centre's measured gain for 3.21 per cent of its exposure**, and the 0.35 per cent is derived rather than assumed: renminbi receipts at 18.0 per cent of protected capital a year give a sawtooth transit balance peaking at 4.500 per cent quarterly, 1.500 per cent monthly and 0.346 per cent weekly, and the design adopts the weekly peak because a freeze arrives when it arrives rather than at the average.

What this paper adds beyond those three measurements is a bound on all of them. The three cases are one one-parameter family: a share x of protected capital moved out of one bloc into a new one, leaving bloc shares $(W - x,\ 0.30,\ 0.30,\ x)$ with $W = 0.40$. Section 3.4 states and Section 7.5 verifies the elementary consequence. **The sum of squares of a donor bloc and its offspring is minimised when the transfer is exactly half the donor bloc**, so the family's ceiling sits at $x = 0.200000$, giving a bloc-HHI of 0.260000, $N = 3.846154$ and $h = -0.028524$. **No amount of exposure moved into a fourth perimeter reaches the build threshold of** $N = 4.172734$ at the stated parameters. That is a design-space bound obtained in one line, and it disposes of a class of proposals before their details are examined.

The paper reuses the headroom gauge of Curzi [2026b] — $h_t = \eta \ln N - \delta Q_t - C_t$ with three regimes, a maximum increment $r_{\max} = h_t$, and the separate constraints of admissibility $C_t \le (1-\delta)Q_t + \eta S_t$ and sustainability $C_t \le \eta \ln N - \delta Q_t$ — and it presupposes the founding reframing of Curzi [2026], that comfort is the felt total of present costs so that the present state is a local minimum of the comfort landscape by definition rather than by luck. **That is why a proposal which measurably worsens a structure is put forward at all: the incumbent position is a local minimum and every proposal that leaves it looks, locally, like movement.**

Two prior results in this series have the same shape and this paper derives from both. Curzi [2026] reports a diversification that reduces resilience: moving a residual 16.4 per cent from one Western jurisdiction to another gives blocs 56.4 / 30.0 / 13.6, a bloc-HHI of 0.426592 and N falling from 2.9412 to 2.3442 **despite Western exposure dropping to 13.6 per cent**. Curzi [2026] reports six Atlantic-cluster entities at UK 40.0 / US 22.0 / Cayman 14.0 / BVI 12.0 / Bahamas 7.0 / Barbados 5.0 giving naive HHI 0.249800 and 4.0032 effective jurisdictions but bloc-HHI 1.000000, $N = 1.00$ and $\ln N = 0$. **Three papers, three regions, one arithmetic.** Curzi [2026] adds the fourth instance from an asset rather than a jurisdiction direction: a converted global allocation reads naive HHI 0.215800 and 4.634 effective jurisdictions but bloc-HHI 0.680000 and $N = 1.470588$ — **the naive count improves while the bloc count worsens**.

The bloc assignment carries an identification problem and it is stated here rather than in a footnote. Manski [1993] proves that a linear-in-means reflexive specification is **not identified from observational data**: the endogenous effect cannot be separated from contextual effects or from correlated unobservables however much data is collected, and Shalizi and Thomas

[2011] show that influence and selection are generically confounded in observational network data. **No amount of observing jurisdictional behaviour could confirm a bloc assignment derived reflexively.** The remedy adopted is the fourth of four: the reference group is stipulated in advance and the stipulation is defended — on the evidence of Sections 4 and 5, and stress-tested in Section 7.2, where the contested assignment is run both ways and its cost is printed. That is the route Brock and Durlauf [2001] shows can be made rigorous, and **the price is that every bloc assignment here is a modelling position rather than an identified empirical result.**

One boundary is set before the machinery arrives. Curzi [2026] refuses category innocence as a defence of acquisition, forbids moving any harm through a discovered gap in a scheme, and concedes that the case where the feedback loop runs through another person's will is an unsettled seam this framework does not close. **This paper measures gaps between perimeters and it does not recommend exploiting any of them**; every design it evaluates is one whose elements are filed, and the disclosure reading in Section 7.6 is where that is tested rather than asserted.

2. Hypotheses

Five hypotheses are stated, each falsifiable in the sense Platt [1964]'s strong inference requires. H5 is falsifiable by counterexample only, which is stated rather than dressed up. Meehl [1967] is the standing warning that directional predictions are cheap; all five below carry point predictions with constants attached, and **two of the five are reported in Section 7 in the direction that goes against the addition being tested.**

H1 — the intra-bloc addition hypothesis. Adding jurisdictions inside a bloc a structure already occupies raises the naive count and lowers the effective bloc count, and it does so by more when the addition is funded out of a position in a different bloc. *Predicted:* adding two Gulf jurisdictions to a structure already holding one raises the naive effective count from 5.290949 to 7.132973, an improvement of 34.8 per cent, while the effective bloc count falls from 2.941176 to 2.841490, a deterioration of 3.4 per cent. *Falsified by:* an intra-bloc addition that raises the effective bloc count; equivalently, a defensible assignment placing two of the three Gulf entries in different perimeters.

H2 — the indifference hypothesis. A bridging slot whose assignment is contested is indifferent on the bloc axis, so the assignment cannot be the ground on which the slot is justified. *Predicted:* the same 10.9 per cent of exposure held in a contested bridging centre assigned to the Western bloc, and held in the Western centre itself, give identical bloc shares 29.1 / 30.0 / 40.9, identical bloc-HHI 0.341962, identical $N = 2.924302$ and identical $h = -0.124430$, while the naive counts differ by 0.841640. The stress cost of the reassignment against the bridge base case is 0.002014. *Falsified by:* any difference at all in N between the two placements, which would mean the two positions are not in the same bloc.

H3 — the capped-position hypothesis. A capped transit position in a fourth perimeter captures a disproportionate share of a full centre's measured gain, because the family's efficiency is highest at small shares. *Predicted:* a full centre at 10.9 per cent gives $N = 3.615826$ and $h = -0.050138$; a capped position at 0.35 per cent gives $N = 2.965384$ and $h = -0.119548$; the capped position captures $0.002869/0.072279 = 3.97$ per cent of the gain for $0.35/10.90 = 3.21$ per cent of the exposure, a ratio of 1.2368. *Falsified by:* a capped position whose gain per unit of exposure is at or below the full centre's, which would mean the index is exactly proportional to the share and the capping strategy buys nothing.

H4 — the sweep hypothesis. The exposure of a transit position is set by sweep frequency rather than by throughput, so the highest-leverage control is a bank instruction rather than a design change. *Predicted:* on receipts of 18.0 per cent of protected capital a year with a sawtooth balance, the peak exposure is $18.0/4 = 4.500$ per cent quarterly, $18.0/12 = 1.500$ per cent monthly and $18.0/52 = 0.346$ per cent weekly, so moving from quarterly to weekly divides the exposed balance by exactly $52/4 = 13.0$. *Falsified by:* a settlement pattern in which the balance is not a sawtooth between sweeps — for instance one in which receipts must be held

to meet a matching obligation — in which case the frequency does not control the exposure.

H5 — the half-share proposition. Moving a share x out of one bloc into a new one maximises the effective bloc count at $x = W/2$, where W is the donor bloc's original share. *Predicted:* for $W = 0.40$ and the other blocs at 0.30 and 0.30, the maximum is at $x = 0.200000$ with bloc-HHI 0.260000, $N = 3.846154$ and $h = -0.028524$, so the family never reaches the break-even count $N = e^{0.5/0.35} = 4.172734$. *Falsified by:* any x in the family whose effective count exceeds 3.846154. **H5 is a proposition of arithmetic rather than an empirical hypothesis. It is falsifiable by counterexample and by nothing else, it carries no empirical weight, and what it does is bound a design space that is otherwise argued about without limits.**

Not hypotheses, and labelled as such. Four further statements appear and none is falsifiable as stated.

- *The logarithmic form of the input side*, $\eta \ln N$ rather than ηN, is a **modelling choice** carried from Curzi [2026b] and not tested here. Section 11 records that Latané [1981]'s sublinear power law fits the qualitative facts as well and that nothing here discriminates between them.
- *The parameter set* $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ is a **modelling choice** carried from Curzi [2026] and fixed in advance, which is the preregistration discipline of Nosek et al. [2018] applied to an engineering measurement.
- *The assignment of a contested centre to one bloc rather than another* is a **modelling position**, not a finding. Section 7.2 runs it both ways and prints the cost, which is the only honest treatment available.
- *The identification of the surplus law's* N *with the effective bloc count* is a **definition**; its content is in the assignment rule, which is testable.

3. Method

3.1 The instrument, and which effective count it is

The concentration instrument is taken unchanged from Curzi [2026]. Its definition fits on one line:

```
HHI = sum of squared shares            effective count = 1 / HHI
```

introduced by Hirschman [1945] and reached independently in sum-of-squares form by Herfindahl [1950], with the reciprocal reading due to Adelman [1969]. Competition practice states the index in points on percentage shares [U.S. Department of Justice and Federal Trade Commission, 2023, 2010]; this paper states it as a fraction.

Which effective count, stated exactly. The reciprocal index HHI^{-1} is the Hill number of order two [Hill, 1973], whose numbers-equivalent reading Jost [2006] states in its cleanest form; the same arithmetic was reached independently in ecology by Simpson [1949] and in political science by Laakso and Taagepera [1979]. The exponential of Shannon entropy [Shannon, 1948] is the same family at order one, and the reciprocal index is also the exponential of the Rényi entropy of order two [Rényi, 1961]. **Every effective count in this paper is HHI^{-1}, the Hill number of order two, and it is named as such at each use.** The framework's other effective count, $N_{\text{eff}} = 1 + (N-1)u$, is a linear interpolation between 1 and N; it is not a Hill number of any order and carries no external support, as Curzi [2026] establishes by running the same verification chain on both and letting one fail visibly. **It is not used to produce any number in this paper.**

The general form is $\text{EFF} = 1/(w'Rw)$ for exposure weights w and correlation matrix R, which is the similarity-sensitive diversity of order two of Leinster and Cobbold [2012], whose additive relative is the quadratic entropy of Rao [1982]; setting $R = I$ recovers HHI^{-1} and setting R to a block matrix of ones recovers the bloc index. Hurlbert [1971] is the standing objection that "diversity" carries no content without its formula attached.

No distributional or scaling claim is made anywhere in this paper. The share vectors here descend steeply and invite a tail argument. **This paper fits no distribution, characterises no tail and asserts no power law.** Clauset et al. [2009] is the reason: a principled maximum-likelihood fit with a goodness-of-fit test and a likelihood-ratio comparison against alternatives leaves few canonical data sets standing, and a straight line on a log-log plot is not evidence. That procedure was not run because no claim requiring it is made, and Stumpf and Porter [2012] adds that even a genuine power law is weak evidence for a mechanism. **The effective counts here are statistics of finite vectors with three, four, six, seven or eight entries.**

The collapse result has an ancestry that should be named rather than reinvented. Granovetter [1978] establishes that the distribution of individual thresholds, not the mean, determines whether a collective outcome tips, and Schelling [1971] is the canonical demonstration that a mild individual preference produces a sharp aggregate result — the strongest available warning against reading an individual-level parameter straight into an aggregate, which is what a bloc assignment does. Castellano et al. [2000] supplies the closest formal analogue of a bloc: the same interaction that raises local agreement freezes global diversity into non-interacting regions, with a sharp transition between the monocultural and fragmented phases. **A bloc, in that vocabulary, is a frozen region and the effective count is a count of them.**

3.2 What a bloc addition does, and the decomposition it needs

The result this paper measures three times has an exact decomposition, and it is worth stating before the cases so that the cases read as instances rather than as anecdotes.

Adding a unit inside a bloc the structure already occupies raises the naive count and leaves the bloc count untouched. That is the invariance Curzi [2026] measures on the Atlantic cluster and Curzi [2026] measures at the apex, where the same invariance holds for a different reason: the forced-heirship exposure that drives the apex choice splits into a situs-capped levy, $\min(60.00\text{m}, 6.70\text{m}) = \text{USD } 6.70\text{m} = 11.2$ per cent of the notional reserved share, and an unaffected tax exposure of $0.45 \times 60.00\text{m} = \text{USD } 27.00\text{m}$, **so the apex is selected on a succession arithmetic that the concentration instrument does not enter at all.**

Funding that addition out of a position in a different bloc lowers the bloc count. The donor bloc's share falls and the recipient bloc's share rises; where the recipient was already the largest, the sum of squares rises and N falls. **The naive count rises and the effective count falls at the same time, and neither movement is an error.**

The decomposition this framework supplies for the underlying quantity is Curzi [2026d]'s: non-parallelism splits as $u = \beta^2 + \alpha^2 \sin^2\gamma$ with $\alpha^2 + \beta^2 = 1$, so that $u_{\text{cross}} = \beta^2$ and $u_{\text{div}} = \alpha^2 \sin^2\gamma$, and the budget constraint $u_{\text{div}} \le 1 - u_{\text{cross}}$ caps divergence at 0.10 once crossing reaches 0.90. **The reading transferred here is the budget, not the trigonometry: total non-parallelism is a fixed quantity that the two sources compete for, so buying more of one inside a fixed total buys less of the other.** A jurisdiction added inside an occupied bloc is spending budget on the source that is already saturated.

A second result of this framework bears directly and is the reason a merely *different* addition buys nothing. Curzi [2026c] establishes that u depends only on the magnitude of correlation, so perfect opposition at $\rho = -1$ scores $u = 0$ exactly as perfect agreement does and **a standing contrarian buys no surplus**. The jurisdictional form is exact: a jurisdiction that reliably opposes another is as correlated with it as one that reliably agrees, and neither adds an independent point of failure. **A structure looking for independence must look for units whose failures are unrelated, not for units whose politics are opposed.**

3.3 Correlated units, and what the literature says about the correction

The naive index overstates resilience because it treats units as independent when they share a coordination mechanism. The correction is Markowitz [1952]'s result that diversification benefit depends on covariance and not on count, and Buldyrev et al. [2010]'s that a system of in-

terdependent networks is far more fragile than either network alone. Perrow [1984] reached the organisational form first — interactive complexity plus tight coupling makes catastrophic failure a property of the system rather than of any component — and Sagan [1993] documents the sharper version, that redundancy can reduce safety, **which is this paper's H1 stated as a general proposition twenty years earlier.** Both borrowings are labelled under the analogy discipline of Section 10.5.

Correlations rise under stress. Longin and Solnik [2001] find extreme correlation in equity markets that is asymmetric and concentrated in downturns and Ang and Chen [2002] report the same asymmetry. **The effective count falls exactly when it is needed.** The ecological name for the decisive variable is synchrony [Loreau and de Mazancourt, 2008], and the strongest empirical demonstration that effective count buys stability is Schindler et al. [2010]'s Bristol Bay sockeye record. The contrary literature is read alongside: Doak et al. [1998] show that diversity–stability correlations can arise from statistical averaging alone, and May [1972] showed that increasing diversity and connectance destabilises randomly assembled systems. **This paper's claim is narrower than the general diversity–stability claim**: the reciprocal of a variance is an effective count of independent shocks, which is arithmetic, and units inside one perimeter are correlated, which is H1.

There is a second risk in this paper's material that the instrument does not model at all, and it is stated here because Section 4.2 measures nothing about it. A concentration index assumes the members of a bloc fail *together*. It does not model the members of a bloc failing *against each other*. Where a regional grouping has a dated, recent instance of internal rupture, a structure spread across it carries both risks at once: it does not diversify against external coordination, and it acquires an internal fracture exposure a single-member structure does not have. **That combination is unusual, it is not priced by anything in this paper, and Section 11 records it as an unmeasured overlay rather than as a scored item.**

3.4 The half-share proposition

The three cases of Section 7 are one one-parameter family and the family has a closed-form ceiling. Stating it here means the cases can be read against a bound rather than against one another.

Let a structure have bloc shares $(W, B_2, \dots, B_k)$ summing to one, and move a share $x \in [0, W]$ out of the first bloc into a new bloc. The bloc index becomes

$$\mathrm{HHI}(x) \;=\; (W-x)^2 + x^2 + \sum_{j\geq 2} B_j^2 .$$

Only the first two terms depend on x. Differentiating, $-2(W-x)+2x=0$, so

$$x^* = \tfrac{W}{2}, \qquad \mathrm{HHI}(x^*) = \tfrac{W^2}{2} + \sum_{j\geq 2} B_j^2,$$

and the effective bloc count is maximised there. Three readings follow and each does work in Section 7.

The optimal transfer is exactly half the donor bloc. Not all of it, and not a token slice: the two resulting shares are equal, which is the numbers-equivalent property of Section 3.1 reappearing as a design rule.

The ceiling is computable before any entity is considered. Substituting $W = 0.40$ and $B_2 = B_3 = 0.30$ gives $\mathrm{HHI}(x^*) = 0.080000 + 0.180000 = 0.260000$ and $N = 3.846154$. **Whatever is added, and wherever it is placed, this family cannot exceed** 3.846154, so a proposal promising to reach the break-even count of 4.172734 by moving exposure into a fourth perimeter is refuted before its details are examined.

Efficiency falls monotonically in x on $[0, x^*]$. The gain $h(x) - h(0)$ is concave, so the *average* gain per unit of exposure, $[h(x)-h(0)]/x$, is largest at small x. **A capped position is more efficient**

per unit of perimeter exposure than a full one, and Section 7.5 measures by how much. That is H3's mechanism, and it is arithmetic rather than judgement.

3.5 Evidence standard, and the de-identification convention

Sections 4, 5 and 6 are jurisdictional exposition. **Every enactment, decided case, licensing regime, constitutional provision, listing event, rate, date and jurisdictional characteristic in them is described from the source manuscript and is not cited to the bibliography.** The bibliography carries the general instruments and the comparative doctrine only: recognition and enforcement of awards [United Nations, 1958], with the grounds on which enforcement is in fact refused [van den Berg, 1981], the model law behind the convergence of the surveyed regimes [United Nations Commission on International Trade Law, 1985], and the doctrine around them [Born, 2021]; the judgments and choice-of-court conventions [Hague Conference on Private International Law, 2005, 2019]; the trusts convention [Hague Conference on Private International Law, 1985]; the succession convention that never entered into force [Hague Conference on Private International Law, 1989] with the regional regulation that partially fills the gap [European Union, 2012]; conflict-of-laws doctrine [Torremans et al., 2022, Briggs, 2013, Fawcett and Carruthers, 2008]; the trust and the foundation compared [Lupoi, 2000, Smith, 2013, Panico, 2014]; asset partitioning as the irreducible function of entity law [Hansmann and Kraakman, 2000] with its critique [Pistor, 2019]; the listing machinery as a technique of governance [de Goede, 2011]; and the disclosure instruments [Organisation for Economic Co-operation and Development, 2017, Financial Action Task Force, 2012, Organisation for Economic Co-operation and Development, 2015]. **It contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction.** Appendix B lists every affected claim, **with the requirement that the date of confirmation be recorded**, because a statement of a jurisdiction's regime without a date is not a statement of its regime.

Where a table would otherwise **rank named jurisdictions against one another on a legal capability the bibliography cannot support**, the jurisdictions are de-identified — "Gulf centre A" and "Gulf centre B", "Asian centre A" and "Asian centre B" — following the convention of Curzi [2026], which de-identifies two Gulf centres in its own scored assessment, and Curzi [2026], which de-identifies three islands in its ranked capability table. **Where a jurisdiction is named in prose, or appears in an exposure table, or carries a bloc assignment with its evidence, the name is used**, because the argument is about those specific perimeter attachments and a functional description would hide rather than protect.

4. The Gulf: one bloc, with an internal fracture

4.1 The perimeter evidence

Everything in this section is described from the source manuscript and is uncited (Section 3.5, Appendix B). The perimeter test of Curzi [2026] asks whether a decision taken in one jurisdiction reliably produces the corresponding effect in another without a further independent decision.

shared mechanism	evidence class
a common regional assessment body operating alongside the global standard-setter	strongest — a supervisory standard producing the same domestic consequence in every member
a common peer-review process on exchange of information	strongest
effectively common exposure to a single correspondent-banking policy	strongest — shared correspondent networks exit corridors together

shared mechanism	evidence class
currency pegs to the same unit across the members	strongest
hydrocarbon revenue correlation	weaker — a common shock, not a common decision
a regional security environment no member can opt out of	weaker

Four of the six rows are in the strongest class. The bloc assignment is therefore not a close call on the coordination axis, and Section 7.1's arithmetic follows from it directly.

4.2 The rupture, and what it is evidence of

The demonstration that matters is not any of the six rows above. **On 5 June 2017 four states of the region severed diplomatic relations with a fifth and closed their land, sea and air links to it. The rupture lasted until a declaration of 5 January 2021.** For three and a half years, a resident family holding one company in the isolated state and another in one of the isolating states held two parts of one structure in jurisdictions that had closed their borders to each other. Flights stopped, some banking relationships were curtailed, and the isolated state's offshore currency traded away from its peg for a period. **The qualitative fact of disruption is well established; the quantitative detail is not something to assert from memory, and Appendix B marks the banking and currency effects for verification rather than reproducing a characterisation of them.**

The event does two different pieces of work and they must not be conflated.

```
CO-VARIANCE RISK   the members respond TOGETHER to external
                   coordination. This is the perimeter mechanism
                   and it is why they are one bloc for the
                   measurement.

INTRA-BLOC RISK    the members have acted AGAINST one another,
                   within the last decade, at speed, with physical
                   closure of borders. This is a failure mode the
                   Western perimeter does not have in the same form.
```

A structure spread across the region carries both risks at once: it does not diversify against external coordination, and it acquires an internal fracture exposure a single-member structure does not have. **The concentration instrument prices the first and is silent about the second**, which is stated here and again in Section 11 rather than folded into a score.

4.3 What the two additional centres are for, de-identified

Neither is a holding jurisdiction and neither is a diversification. **The comparison below ranks them on legal capability the bibliography cannot support, so they are de-identified** (Section 3.5).

requirement	best fit	why
a common-law financial centre whose licence is not territorially confined	Gulf centre A	its licensed entities may operate anywhere in the national economy, which converts an enclave into a regime and makes a genuine domestic operating entity possible
market access to a very large state customer, conditioned on physical presence	Gulf centre B	a headquarters programme that is a market-access requirement with a tax incentive attached, rather than a tax incentive with a substance requirement attached
a holding position outside the incumbent perimeter	**neither**	both sit inside the bloc the structure already occupies; Section 7.1 measures the cost of pretending otherwise

The distinction in the second row decides whether the item belongs in a structure at all. A tax incentive with a substance requirement is a cost to be minimised; a market access requirement with a tax incentive is a price to be paid. **The first is optimised, the second is bought**, and an adviser who cannot say which one is on the table has not specified the decision.

5. The bridging slot: a contested assignment, argued rather than asserted

Everything in this section is described from the source manuscript and is uncited. The bridging centre's assignment is the most contested call in this paper and the evidence runs both ways.

evidence for a bridge assignment	evidence against
no participation as a decision-taker in the regional listing mechanism	targeted financial measures and export controls adopted in March 2022 without a Security Council resolution, a departure from six decades of stated practice
no domestic designation apparatus of the kind the perimeter's two largest members operate	full membership of and participation in the standard-setting body whose recommendations are the primary coordination channel
a designated clearing bank for a second major currency since 2013, with genuine offshore capability	participation in automatic exchange with first exchanges in 2018, and a bilateral reporting arrangement
regional centrality and a policy of non-exclusion, hosting full-service operations of institutions from every bloc simultaneously	a banking system whose dollar correspondent access is the operational constraint on everything, and which therefore inherits the correspondent perimeter's risk appetite

evidence for a bridge assignment	evidence against
a foreign-policy doctrine of not choosing sides, stated consistently at senior level for six decades	

The last row on the right does the work and it generalises beyond this jurisdiction. A jurisdiction's coordination alignment is set less by its foreign ministry than by its banks' correspondent relationships, **because the correspondent relationship is the thing that can be withdrawn in a week.** The centre's foreign policy is genuinely independent. Its dollar clearing is not, because nobody's is.

The verdict carried into Section 7.2 is therefore split: **bridge in commercial function, Western-leaning in coordination behaviour under stress.** The consequence is not a judgement call left open; it is a measurement run both ways, and Section 7.2 prints the difference so a reader who disagrees with the assignment can read their own answer off the table.

6. The transit position: two perimeters, and the trap in the arithmetic

Everything in this section is described from the source manuscript and is uncited.

The third jurisdiction under test has an unusual property and it is the worst available combination. **For coordination purposes it sits inside one perimeter**: its national-security framework is legislation of the sovereign, the interpretive power over that framework is the sovereign's, and the power to designate the judges who hear cases under it runs through an executive selected under a system substantially revised in 2021. **For correspondent and settlement purposes it sits inside the other perimeter**, because its banks require dollar correspondent access and its currency is maintained against the dollar under a linked exchange rate. That second exposure is not theoretical: the other perimeter's largest member enacted legislation and issued an executive order in 2020 suspending or eliminating a range of preferential treatments, and separately suspended bilateral agreements including one on surrender of fugitives.

A jurisdiction inside two coordination perimeters is not a bridge; it is a place where two sets of coercive capability overlap. That is the exact inverse of the property the bloc collapse is designed to find.

Which brings the trap into focus, and it is the counter-result of this paper. In a structure with no other exposure to the first perimeter, assigning this jurisdiction to that perimeter raises the effective bloc count sharply — from 2.941176 to 3.615826 — because the bloc collapse rewards a share placed in a perimeter the structure does not otherwise touch. **The arithmetic is correct.** What it does not and cannot see is that the marginal exposure is placed under a perimeter with demonstrated reach through the courts, and that the entity holding it is simultaneously exposed to the other perimeter through its correspondent bank. **The instrument prices the position. It does not price the hazard, and this is the jurisdiction where the gap between the two is widest.**

The design consequence is Section 7.3's Design B. **Take the capability and cap the position**: the arbitration seat, the currency settlement capability and the market access are bought through contract terms and a small balance, and the balance is controlled by sweep frequency rather than by design (Section 7.4).

The framework has an instrument for valuing that kind of move and it is not the concentration index. Curzi [2026] scores each step of a path 0 to 3 for steepness, sums to a resistance R, and weights the branch as $1/2^R$, so that a three-point difference in steepness is a factor of eight in expectation. **A capability bought through a clause rather than through a balance is a step with lower steepness on every branch that ends in a freeze**, and the branch weighting

says what that is worth in expectation even where the index says the position is worth 0.002869. The two readings are different quantities and both are printed in Section 7.3 rather than one being substituted for the other.

7. Results

All figures below are exact computations on constructed inputs, carried at full precision and rounded only for printing. **No external dataset was collected and no figure in this paper is a measurement of any real structure.** Where a printed value in the source manuscript differs from the exact computation, Section 7.7 prints both.

7.1 R1 — Gulf additions, and H1

The designed structure, and the tempted variant in which two further Gulf jurisdictions are added and funded partly by trimming the incumbent bridge, Western and non-aligned positions:

	designed	tempted
UAE	29.1% BRIDGE	20.0% BRIDGE
Thailand	18.2% NONALIGNED	18.2% NONALIGNED
UK	16.4% WEST	12.2% WEST
Bermuda	13.6% WEST	13.6% WEST
India	11.8% NONALIGNED	11.8% NONALIGNED
Singapore	10.9% BRIDGE	6.9% BRIDGE
Saudi Arabia	—	9.1% BRIDGE
Qatar	—	8.2% BRIDGE

```
TEMPTED, naive
  0.200^2 + 0.182^2 + 0.136^2 + 0.122^2
+ 0.118^2 + 0.091^2 + 0.082^2 + 0.069^2  = 0.140194
  effective jurisdictions = 1 / 0.140194 = 7.132973

TEMPTED, blocs
  BRIDGE 44.2  (UAE + Saudi + Qatar + Singapore)
  NONALIGNED 30.0  (Thailand + India)
  WEST 25.8  (Bermuda + UK)
  bloc-HHI = 0.195364 + 0.090000 + 0.066564 = 0.351928
  EFFECTIVE BLOCS = 1 / 0.351928 = 2.841490
```

	designed	tempted	change
jurisdictions held	6	8	$+2$
naive effective count	5.290949	7.132973	$+34.81\%$
effective bloc count	2.941176	2.841490	-3.39%
h	-0.122417	-0.134485	-0.012068

Two more jurisdictions, a naive resilience count a third higher, and the honest count went down. The structure looks more diversified on every measure a brochure would use and is fractionally less diversified on the only measure that predicts failure.

The mechanism is exact and is worth stating because it is the general case. The Gulf additions were funded partly out of the other bridge position, which is a *different* bridge, so the bridge bloc concentrated even as the jurisdiction count rose. **Rearranging inside a bloc does not change N; funding an intra-bloc addition out of an extra-bloc position makes N worse.** That is Section 3.2's decomposition, instantiated.

The variant actually built. One proposal was taken and one declined. The taken proposal is a regional headquarters justified by contracted revenue that is unavailable without it; the declined proposal is a family-office vehicle in the other centre with no operating activity, no customers, no contracts, no assets and no people, which would have duplicated the apex's function inside the same bloc at additional annual cost. The taken addition is funded from the operating allocation rather than by trimming an extra-bloc position:

line	share	bloc
UAE	29.1%	BRIDGE
Thailand	13.2%	NONALIGNED
UK	16.4%	WEST
Bermuda	13.6%	WEST
India	11.8%	NONALIGNED
Singapore	10.9%	BRIDGE
Saudi Arabia	5.0%	BRIDGE

```
naive HHI = 0.084681 + 0.017424 + 0.026896 + 0.018496
          + 0.013924 + 0.011881 + 0.002500 = 0.175802
effective jurisdictions = 1 / 0.175802 = 5.688217

blocs: BRIDGE 45.0  WEST 30.0  NONALIGNED 25.0
bloc-HHI = 0.202500 + 0.090000 + 0.062500 = 0.355000
EFFECTIVE BLOCS = 1 / 0.355000 = 2.816901
```

	designed	tempted	as built
jurisdictions	6	8	7
naive effective count	5.290949	7.132973	5.688217
effective bloc count	2.941176	2.841490	2.816901

The honest result: taking the regional headquarters cost 0.124275 of effective bloc count, and it was taken anyway, because it was justified by revenue and not by resilience. That is the correct way to make the decision and it is worth stating explicitly. **The instrument is a measurement, not a veto.** A structure exists to hold a business, and a business that wins a large contract has done something the concentration measure cannot price. What the instrument does is make the cost visible and force it to be paid consciously: **the family knows it gave up 0.124275 of N and knows what it received; the tempted variant gave up 0.099686 and received a brochure.**

H1 is supported. The naive count rose 34.81 per cent and the effective bloc count fell 3.39 per cent, and no admissible reassignment separates the three Gulf entries — four of the six shared mechanisms in Section 4.1 sit in the strongest evidence class.

7.2 R3 — the indifferent bridging slot, and H2

The base case assigns the contested bridging centre to the bridge bloc:

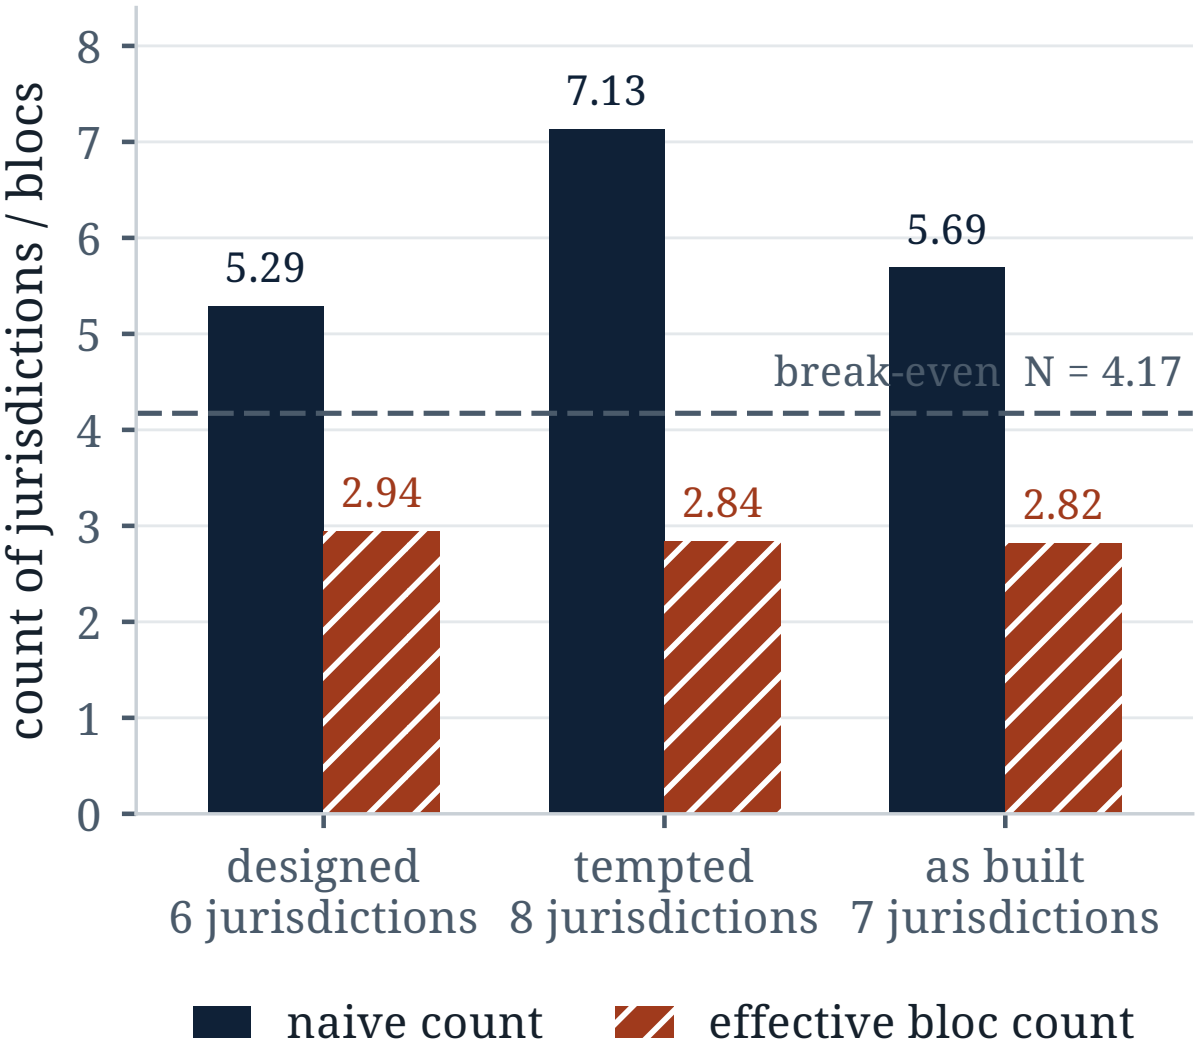


Figure 57.1: **Two more jurisdictions raise the naive count by 34.8 per cent and lower the effective bloc count by 3.4 per cent.** Naive count against effective bloc count for the designed, tempted and as-built structures, with the break-even effective count $N = e^{0.5/0.35} = 4.172734$ the surplus law requires at the stated parameters. All shares are constructed. Generated by `figs/src/CCP-057_fig01.py`.

```
BRIDGE 40.0  NONALIGNED 30.0  WEST 30.0
bloc-HHI = 0.160000 + 0.090000 + 0.090000 = 0.340000
N = 2.941176   ln N = 1.078810   eta ln N = 0.377583   h = -0.122417
```

The stress case reassigns it to the Western bloc:

```
BRIDGE 29.1  NONALIGNED 30.0  WEST 40.9
bloc-HHI = 0.084681 + 0.090000 + 0.167281 = 0.341962
N = 2.924302   ln N = 1.073056   eta ln N = 0.375570   h = -0.124430
stress cost = 0.124430 - 0.122417 = 0.002014
```

And the comparison that decides the section. Hold the same 10.9 per cent in the Western centre itself instead of in the bridging centre. The bloc shares are *identical* — 29.1 / 30.0 / 40.9 — so the bloc-HHI, N and h are identical to every decimal place. Only the naive count differs, because the Western centre already holds 16.4 per cent and the two positions merge into one 27.3 per cent line.

variant	naive HHI	naive effective	N	h
bridging centre assigned BRIDGE (base)	0.189002	5.290949	2.941176	-0.122417
bridging centre assigned WEST (stress)	0.189002	5.290949	2.924302	-0.124430
the same 10.9% held in the Western centre	0.224754	4.449309	2.924302	-0.124430

Read the bottom two rows together: they differ by 0.841640 on the naive count and by exactly zero on the bloc count and on h. H2 is supported. If the bridging centre is Western in the coordination sense, the entire measured value of moving that exposure is a naive-count artefact.

That is not an argument against the position, and the paper says so in the same paragraph. The real value of the bridging leg is elsewhere and it is not small: several dispute pathways including a commercial court with an opt-in jurisdiction, a fund vehicle with statutory segregation, multi-currency settlement in the operating time zone, and a bank layer that will open the accounts. **Capability is a real good; it is simply not the good the concentration instrument measures, and a leg justified on the concentration instrument is a leg justified on the wrong grounds and therefore sized on the wrong grounds.**

The marginal contribution of the leg on the bridge assignment is small and is stated:

```
0.35 x (ln 2.941176 - ln 2.600000) = 0.35 x 0.123299 = 0.043154
```

against a first non-Western move worth roughly 0.70 of N on the same build path. **That is** $dh/dN = \eta/N$ **arriving on schedule**, and it is the reason no later leg can repeat what an early one did.

7.3 R2 — the capped transit position, and H3

Two designs for the third jurisdiction. **Design A** places a full Asian centre in the 10.9 per cent slot, assigned to a fourth bloc. **Design B** keeps the bridging centre's position and adds a small entity handling settlement in the second currency, sized at 0.35 per cent of protected capital (Section 7.4) with the bridging position reduced correspondingly.

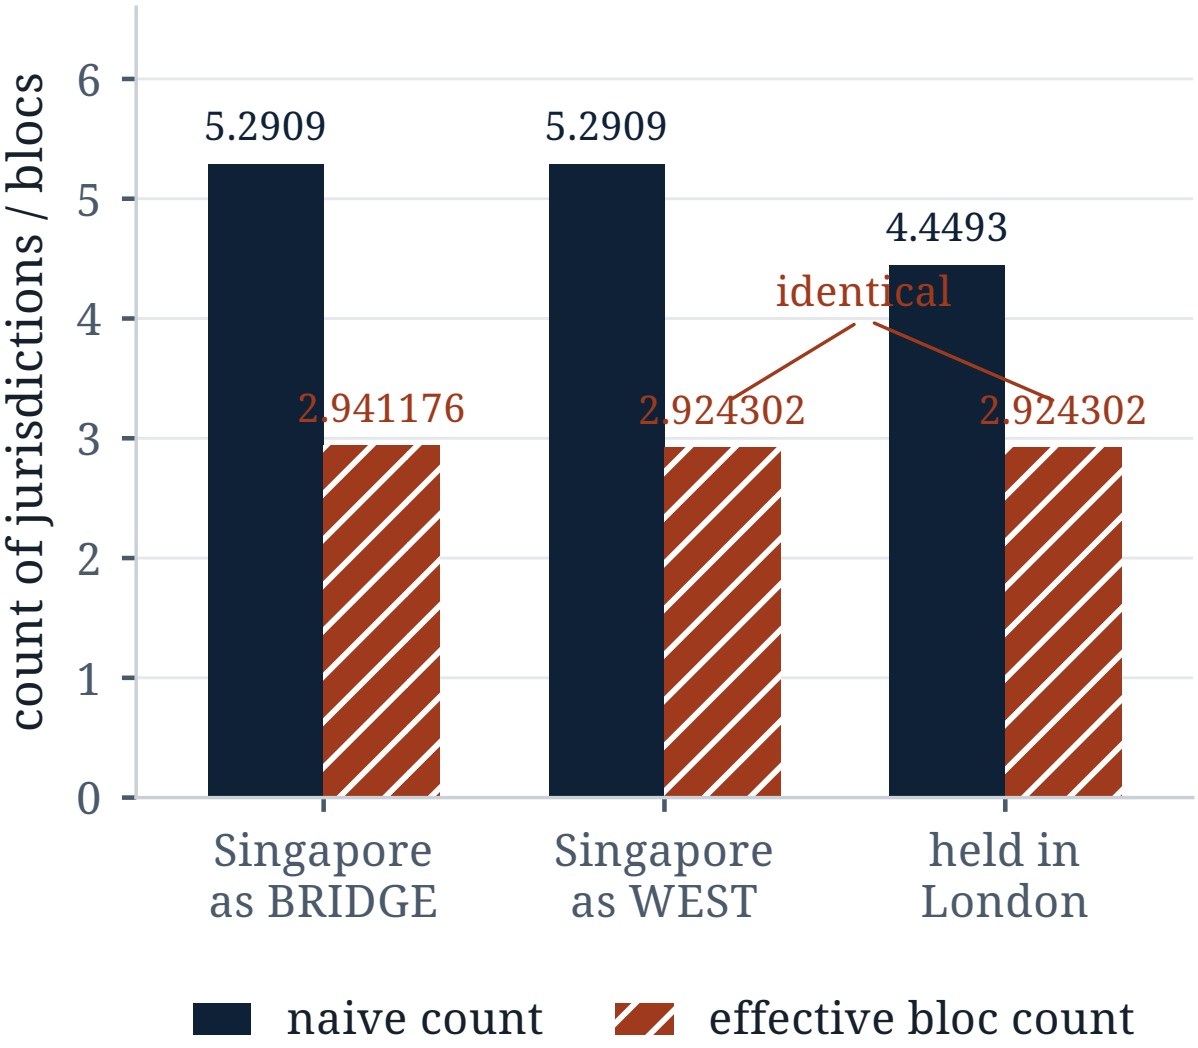


Figure 57.2: **Assigning the bridging centre to the Western bloc, or holding the same exposure in the Western centre itself, gives an identical effective bloc count.** Naive count against effective bloc count for the three variants of Section 7.2. The lower two differ by 0.841640 on the naive count and by exactly zero on the bloc count. All shares are constructed. Generated by `figs/src/CCP-057_fig04.py`.

```
DESIGN A
  blocs: BRIDGE 29.1  NONALIGNED 30.0  WEST 30.0  FOURTH 10.9
  bloc-HHI = 0.084681 + 0.090000 + 0.090000 + 0.011881 = 0.276562
  N = 3.615826   ln N = 1.285320   eta ln N = 0.449862   h = -0.050138

DESIGN B
  blocs: BRIDGE 39.65  NONALIGNED 30.0  WEST 30.0  FOURTH 0.35
  bloc-HHI = 0.157212 + 0.090000 + 0.090000 + 0.000012 = 0.337224
  N = 2.965384   ln N = 1.087006   eta ln N = 0.380452   h = -0.119548
```

design	N	h	gain in h	exposure inside the fourth perimeter
baseline	2.941176	-0.122417	—	0.00%
B: capped transit	2.965384	-0.119548	$+0.002869$	0.35%
A: full centre	3.615826	-0.050138	$+0.072279$	10.90%

```
Design B captures  0.002869 / 0.072279 = 3.97%  of Design A's gain
Design B carries      0.35% / 10.90%   = 3.21%  of Design A's exposure
ratio                    3.97 / 3.21   = 1.2368
```

H3 is supported. The measured benefit is very nearly proportional to the exposure placed inside the perimeter — which is exactly what a concentration index should do and exactly why it must not be the sole criterion. **The index rewards the share; the share is the hazard.** Design A's 0.072279 of headroom is bought by putting 10.9 per cent of protected capital inside a perimeter with demonstrated judicial reach and simultaneous correspondent exposure to the other perimeter. Design B buys 3.97 per cent of that headroom for 3.21 per cent of that exposure, and buys separately — through the arbitration clause, at no balance-sheet cost — the one capability that has no substitute.

The departure from exact proportionality is the finding rather than a rounding error, and Section 7.5 measures it: the ratio of 1.2368 says the capped position buys 23.68 per cent more headroom per unit of exposure than the full centre, and that is a consequence of the concavity established in Section 3.4 rather than of anything about the jurisdictions.

7.4 R4 — the sweep schedule, and H4

The 0.35 per cent figure is derived, not assumed, and the derivation is the section. CONSTRUCTED: renminbi receipts equal to 18.0 per cent of protected capital per year, with the transit entity's balance a sawtooth between sweeps, so the average balance is half of one sweep period's receipts.

sweep frequency	peak balance	average balance
quarterly	$18.0/4 = 4.500\%$	2.250%
monthly	$18.0/12 = 1.500\%$	0.750%
weekly	$18.0/52 = 0.346\%$	0.173%

The design adopts 0.35 per cent, which is the weekly *peak* — the conservative choice, since a freeze arrives when it arrives rather than at the average. The ratio of the quarterly peak to the weekly peak is $52/4 = 13.0$ exactly.

H4 is supported. Sweep frequency is the single highest-leverage control in the whole design: moving from quarterly to weekly divides the exposed balance by exactly thirteen and

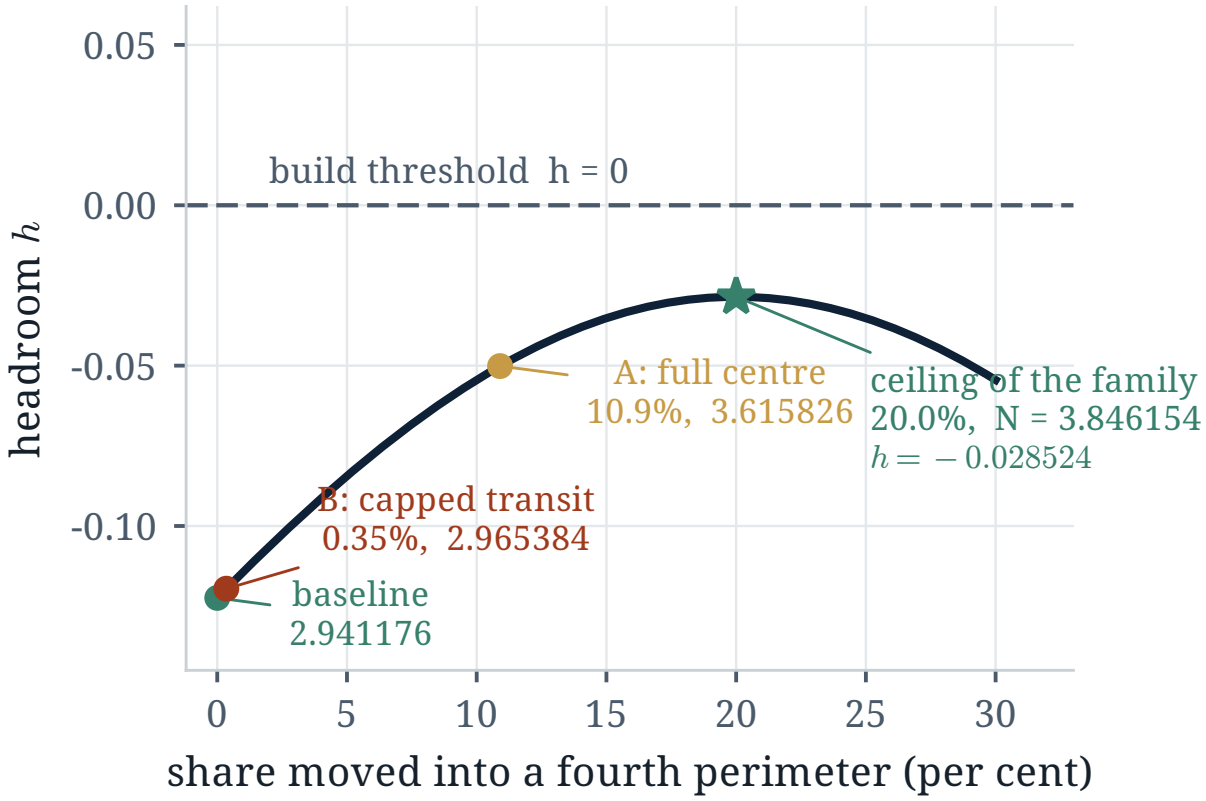


Figure 57.3: **Headroom rises with the share moved into a fourth perimeter, peaks at 20 per cent, and never reaches the build threshold.** The one-parameter family of Section 3.4 at $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$, with the baseline, the capped transit position and the full centre marked, and the closed-form ceiling at $x = 20.0$ per cent, $N = 3.846154$, $h = -0.028524$. Modelled from the constructed bloc vector; generated by `figs/src/CCP-057_fig02.py`.

costs a bank instruction. No entity change, no jurisdictional change and no negotiation produces a comparable movement.

That concentration of leverage in one control has a precedent in this framework and it is worth naming, because it is a recurring shape rather than a coincidence. Curzi [2026] reports Finding 19.2: the only gate acting on the numerator sits ninth of twelve in a schedule and contributes -5.305895 octaves, **88.93 per cent of the schedule's total movement, against 11.07 per cent for the eight before it. A schedule of controls is usually one control and eleven decorations, and the discipline is to find which one it is before optimising the others.**

The falsifier attached to H4 is worth restating because it is a real limit. **If the transit balance is not a sawtooth — if receipts must be held to meet a matching obligation, or if a counterparty requires a minimum balance — then sweep frequency does not control the exposure**, and the 0.35 per cent is not derivable from the throughput. That case is not constructed here and the paper does not claim it away.

7.5 H5 — the half-share ceiling, and the efficiency it implies

The three designs of Sections 7.1 to 7.3 are one family. Applying Section 3.4 with $W = 0.40$, $B_2 = B_3 = 0.30$:

quantity	value
the optimal transfer $x^* = W/2$	0.200000
bloc-HHI at the optimum, $W^2/2 + \sum B_j^2$	$0.080000 + 0.180000 = 0.260000$
the ceiling $N(x^*)$	3.846154
h at the ceiling	-0.028524
the break-even count the surplus law requires	4.172734

H5 is supported, and its consequence is the strongest single statement in this paper. No transfer of exposure into a fourth perimeter, of any size and at any placement, brings this structure to the build threshold. The best available is $N = 3.846154$ against a required 4.172734, a shortfall of 0.326580 in the count and 0.028524 in h. **A proposal to reach the threshold by adding a fourth perimeter is refuted in one line, and the line does not mention any jurisdiction.**

The efficiency reading follows from the same concavity and is the quantitative form of H3:

position	share x	gain $h(x) - h(0)$	gain per unit of exposure
B: capped transit	0.0035	0.002869	0.819674
A: full centre	0.1090	0.072279	0.663107
the ceiling	0.2000	0.093892	0.469462

```
B / A on efficiency  = 0.819674 / 0.663107 = 1.236110
B / A on share ratio = 3.97 / 3.21         = 1.236760
```

The two readings agree to four significant figures and differ only in the rounding of the percentages, which is the cross-check Section 10.9 records. **A capped position buys 23.61 per cent more headroom per unit of perimeter exposure than a full centre, and the figure is a property of the arithmetic rather than of the jurisdiction.**

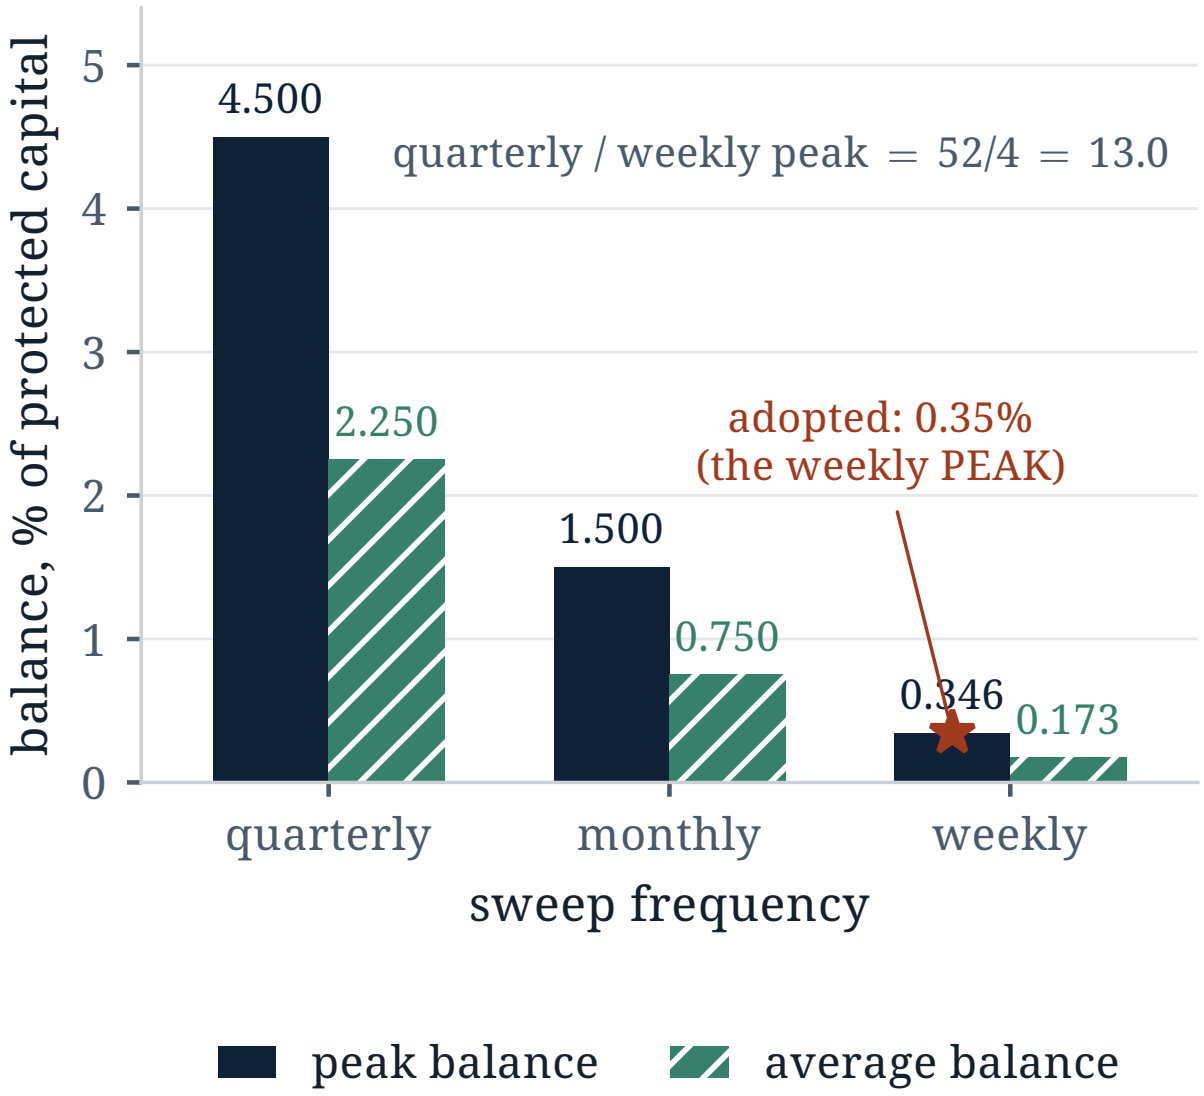


Figure 57.4: **Sweep frequency divides the exposed balance by thirteen between quarterly and weekly.** Peak and average transit balance as a percentage of protected capital under three sweep frequencies, on constructed receipts of 18.0 per cent of protected capital a year with a sawtooth balance. The design adopts the weekly peak of 0.35 per cent rather than the average. Generated by `figs/src/CCP-057_fig03.py`.

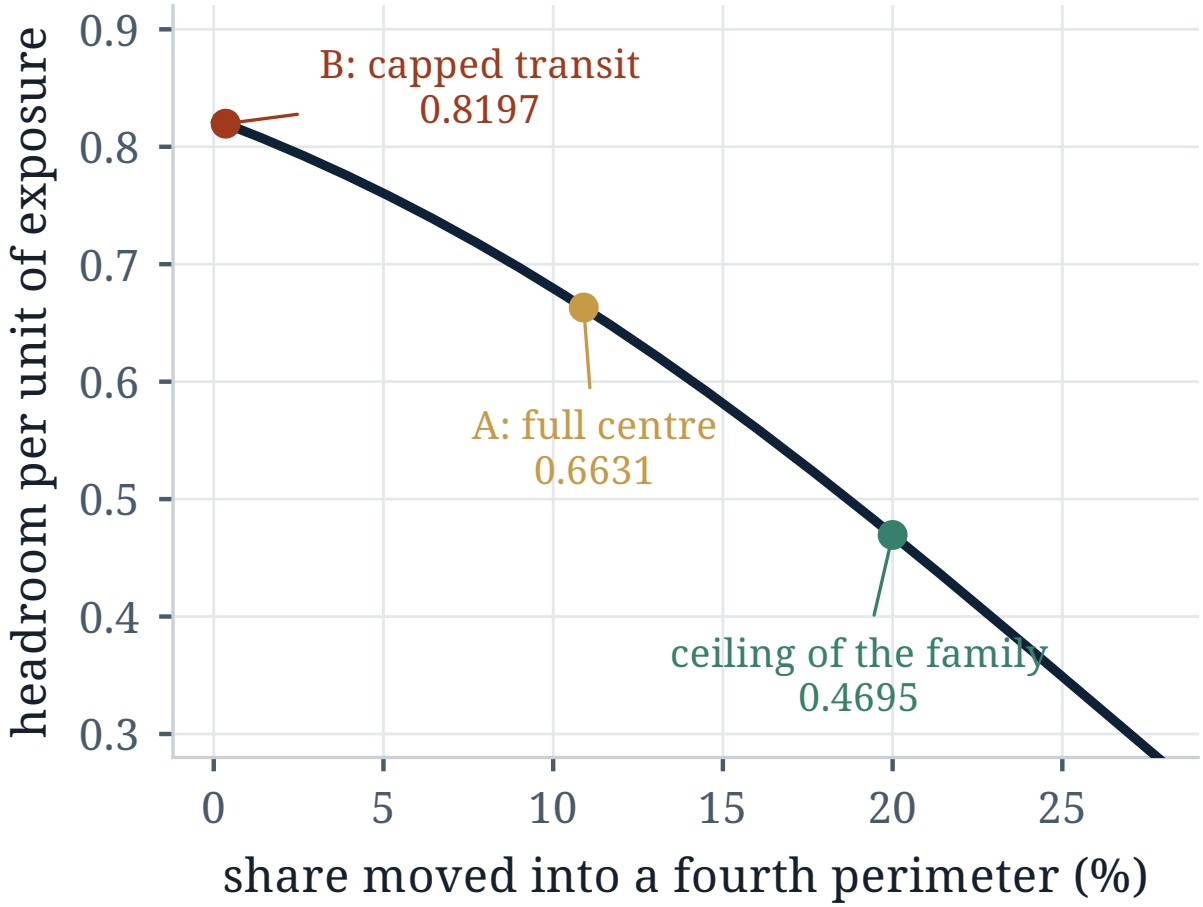


Figure 57.5: **Headroom bought per unit of exposure falls monotonically, so a capped position is 23.61 per cent more efficient than a full centre.** Average gain in h per unit of protected capital moved into a fourth perimeter, on the family of Section 3.4, with the capped transit position, the full centre and the family's ceiling marked. Modelled from the constructed bloc vector; generated by `figs/src/CCP-057_fig05.py`.

7.6 The disclosure reading, applied to all three

The disclosure test asks whether a structure would still do its job if every element were disclosed to every authority with jurisdiction over any part of it. The premise is not this work's to claim. Kerckhoffs [1883] stated it in 1883: a system must remain secure when everything about it except the key is public. Saltzer and Schroeder [1975] restate it as *open design* — the protection mechanism should not depend on the ignorance of attackers — and place it alongside *complete mediation* precisely because a publishable design still requires that every access be checked. Brandeis [1914] is the public-interest form of the same observation. **Citing Kerckhoffs converts what reads as an ethical preference into a recognised engineering standard**, and Curzi [2026] states the constraint in its strongest form: channel separation is strengthened by disclosure, because its effectiveness is entirely independent of whether anyone knows it is in place, so every element must be publishable and non-disclosure is refused as a design premise rather than merely discouraged.

element	reading
the regional headquarters	licensed, staffed, premised and filed. **Its whole function is to be visibly present**, and it is the one entity in this series whose entire purpose is to be seen. Disclosure is not a cost; it is the product
the declined family-office vehicle	does not exist, **which is the cheapest possible way to pass any test**
the bridging leg	a treasury company with employees, premises, a resident-majority board and a written policy; a fund vehicle with a regulated manager; contracting entities with named forums. Every element is filed
the capped transit entity	a small balance, a written cap, a weekly sweep and a named arbitral seat. **The cap is a document, and a document is disclosable**

All four pass, and the first is the clearest illustration in this part of the series that substance which is real, filed and visible is not a compliance burden attached to the strategy — **in that case it is the strategy**, because the market access being bought is conditioned on exactly that visibility. The regulatory direction makes the same point from the other side: automatic exchange [Organisation for Economic Co-operation and Development, 2017], the common anti-money-laundering recommendations [Financial Action Task Force, 2012] and the substance and beneficial-ownership work of the base-erosion project [Organisation for Economic Co-operation and Development, 2015] together mean that non-disclosure is not an available design parameter.

7.7 Printed discrepancies, left visible

First, the stress-case headroom. The source manuscript prints $h = -0.124431$ for the Western reassignment of the bridging slot. The exact value is -0.124430 — more precisely -0.1244305, which rounds to -0.124430 at six decimal places. **The derived stress cost of** 0.002014 **is unaffected**, because $0.1244305 - 0.1224167 = 0.0020138$. This paper prints -0.124430 and records the source's figure here.

Second, the base-case headroom. The source prints $h = -0.1226$ for the designed structure in three places. The exact value at $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ is -0.122417, a difference of 0.000183. **This paper uses** -0.122417**, which is the value Curzi [2026] publishes**, and Curzi [2026] and Curzi [2026] record the same discrepancy independently.

Third, a description rather than a number. The as-built variant of Section 7.1 is described in one summary of this work as a variant *placing the apex in the second Gulf centre*. **It is not.** The source declines that vehicle outright and the variant computed is the regional-headquarters addition in the other Gulf state. The numbers — 0.175802, 5.688217, 0.355000, 2.816901 — are unaffected and are the ones other papers cite; **the description is corrected here rather than propagated.**

Fourth, and for consistency across the series, the naive index of the conventional four-jurisdiction structure. Curzi [2026] reports 0.293388 from exact ninths, giving 3.4085 effective jurisdictions against bloc-HHI 1.000000 and 1.0000 effective perimeters; Curzi [2026] reports 0.293430 from rounded percentages, giving 3.407968. **This paper uses the rounded-percentage figure where it refers to that structure**, and states which it uses.

7.8 Hypotheses tested

#	hypothesis	verdict	where
H1	intra-bloc additions raise the naive count and lower the effective one: $+34.81\%$ and -3.39%	**supported**	7.1
H2	the bridging slot is indifferent on the bloc axis; the stress cost is 0.002014	**supported**, and the naive difference of 0.841640 is an artefact	7.2
H3	a capped position captures a disproportionate share of the gain: 3.97 against 3.21 per cent	**supported**	7.3
H4	sweep frequency controls the exposure, quarterly to weekly by a factor of 13.0	**supported**, subject to the sawtooth premise its falsifier names	7.4
H5	the family's ceiling is at half the donor bloc: $N = 3.846154$, $h = -0.028524$	**supported** as arithmetic, with no empirical content	7.5

Two of the five report against the additions being tested, which is the point of the paper: H1 says two of the three proposals lower the quantity they were offered to raise, and H5 says none of them could have reached the threshold anyway. **H4's premise is a construction and its falsifier was not tested**, because no non-sawtooth settlement pattern was constructed; Section 11 records that.

8. Failure modes

Twelve ways the results above are misapplied, each with the diagnostic that detects it.

1. Counting jurisdictions instead of blocs. The error the whole series exists to correct, and Section 7.1 is its most expensive instance. *Diagnostic:* collapse the shares onto perimeters before squaring, and report the pair.

2. Assuming a regional rupture cannot recur. Section 4.2 dates one that lasted three and a half years. *Diagnostic:* ask what a border closure between two members would do to the structure's ability to instruct, and write the answer down.

3. Treating a market-access headquarters as a holding vehicle. It exists to be present in one market, and nothing should be placed under it. *Diagnostic:* if group holdings sit beneath it, the job assignment has been violated.

4. Assuming one common-law centre delivers another's case law. Two centres in the same legal tradition have different benches and different bodies of decisions. *Diagnostic:* name the reported decisions the argument relies on and check where they were decided.

5. Reading contractual outcomes off a code. *Diagnostic:* the enforceable route, not the statutory text, is the thing to specify.

6. Litigating where an award would travel better than a judgment. *Diagnostic:* choose the forum by the enforcement route, not the other way round [United Nations, 1958].

7. Discounting a political risk because its known instances happened to residents. *Diagnostic:* the mechanism does not read passports; name the mitigation and state that it is partial.

8. Substance built as documentation. *Diagnostic:* the test is applied to core income-generating activity and the evidence is the decision trail, not the address.

9. Assuming a designated business district is a legal enclave. Two different things travel under similar names. *Diagnostic:* ask which court hears a dispute and under what law.

10. Building a structure inside one region only. Section 7.1: the count of jurisdictions rises and the count of perimeters does not. *Diagnostic:* Section 3.4's ceiling, computed before the entities are chosen.

11. Sizing a bridging leg on the concentration instrument. Section 7.2 shows the instrument is indifferent to it. *Diagnostic:* write down the capability being bought, and size the leg on that.

12. Holding rather than transiting in a two-perimeter jurisdiction. Section 6. *Diagnostic:* a rail is not a holding; the parameters are hop size, exposure duration, counterparty pre-screening and mid-transit failure handling, and the balance cap is a document.

9. Discussion

What follows from the results. Three things.

The first is that **a measurement that reports against the proposal being measured is the only evidence that the measurement is doing work.** Two of the three additions tested here lower the quantity they were offered to raise, and one of them was taken anyway on grounds the instrument cannot see. **That combination — a measured cost, consciously accepted, in exchange for a benefit the measure does not price — is the correct use of the instrument**, and a series in which every measurement flattered the design it accompanied would be a series of brochures.

The second is that **an addition inside an occupied bloc is not merely useless; funded out of an extra-bloc position it is harmful.** That is the general form of Curzi [2026]'s result — moving a residual 16.4 per cent from one Western jurisdiction to another gives blocs 56.4 / 30.0 / 13.6, a bloc-HHI of 0.426592 and N falling from 2.9412 to 2.3442 despite Western exposure dropping to 13.6 per cent — and of Curzi [2026]'s, where a converted global allocation reads naive HHI 0.215800 and 4.634 effective jurisdictions but bloc-HHI 0.680000 and $N = 1.470588$. **Four independent instances of one arithmetic is not a coincidence; it is what an index does when the units are correlated and the analyst is counting units.**

The third is the proposition of Section 3.4, and it is the one this paper would defend hardest. **The optimal transfer out of a bloc is exactly half of it, and the resulting ceiling is computable before any entity is considered.** On this structure the ceiling is 3.846154 against a required 4.172734. **A design conversation with an arithmetic ceiling below its own target should end at the arithmetic**, and the way to find out is one line rather than a year of negotiation.

What does not follow.

Nothing here is a legal conclusion. Every jurisdictional statement in Sections 4, 5 and 6 is described from the source manuscript and is uncited; Appendix B lists them. **A reader who takes a rule of law from this paper has taken it from the wrong document.**

The instrument prices the position and not the hazard. Section 6 states this and it bears restating, because the 3.615826 of Design A is the most quotable figure in the paper and the most likely to be quoted without the sentence beside it. **A share placed in a perimeter the structure does not otherwise touch raises the effective count whether that perimeter is benign or coercive, and the index has no term for the difference.**

The intra-bloc fracture risk is not priced anywhere in this paper. Section 3.3 and Section 4.2 both state it. A concentration index assumes bloc members fail together; it does not model them failing against each other, and the region under study has a dated recent instance. **That is an unmeasured overlay and it is disclosed rather than scored.**

The effective count measures independence, not capacity. A structure diversified on jurisdiction and concentrated on counterparty has moved its risk rather than reduced it, and Curzi [2026] measures how far the entity axis can understate the truth: entity-based exposure of 27.3 per cent against a currency share of 61.0 per cent and settlement volume through the same system of approximately 58 per cent, **an understatement by more than a factor of two.**

The gain is small and easily cancelled by a cost the index does not see. Curzi [2026] measures exactly that: at a 15 per cent annual hard-currency drag, a local-currency allocation of 25.17 per cent of protected capital sets h to exactly -0.500000, **cancelling the whole** $\eta \ln N = 0.377583$ **diversification gain**, and a 10 per cent allocation costs 0.150000, which is 39.73 per cent of the gain. **A gain of** 0.072279 **from a full centre is inside the noise of a currency decision**, and any argument that sizes a jurisdictional leg without pricing the currency it holds has mis-specified the problem.

The measurement is least informative exactly where a dominated structure sits [Golosov, 2010, Dunleavy and Boucek, 2003], and **a count that can be inflated will be** [Goodhart, 1975, Campbell, 1979]. Douceur [2002] is the formal statement that one party can present as many, so an unaudited count is worthless, and Lamport et al. [1982] is the bound on how many genuinely independent parties a protocol requires. **Both belong next to any count of units in this series**, and Section 7.1 is the Sybil problem wearing a passport.

The framework's own results bound this one. Curzi [2026] reports a doubling theorem verified *exactly* over all sixty-six pairs at $N = 12$ — $\mathrm{HD}_{\mathrm{ess}}(i, j) = 2\log_2(pq)$ and $G_{\mathrm{ess}} = 2G(pq) - 1$ for the essence ratio p^2/q^2 — which is the standard of verification a closed form in this framework is expected to meet, and the standard Section 10.9 applies to the half-share proposition. Curzi [2026] states the companion result that the sustaining frequency $f_{\min} = k\, e^{2c/\eta}$ is a genuine equivalence verified at twelve parameter sets, with $k = 0.483425732988414$ — **a DEFINED constant, exact by construction, whose fifteen figures are legitimate for that reason and for no other; attaching a physical unit to it is a separate and much stronger claim, badged POSTULATED there and not relied on anywhere in this paper.**

And this corpus's central limitation applies in full. No fit is reported here at all, and that is not a strength. Epstein [2008] is the standing statement that prediction is one of sixteen reasons to model and that a paper is entitled to say which of the other fifteen it claims; **this paper claims explanation of a mechanism and demonstration that a stated arithmetic suffices to produce a stated shape**, and claims neither prediction nor calibration.

10. Verification and rigour audit

The results are decomposed into postulates that can be checked one at a time. Each is atomic, each carries a status, and the compounding chain names what each step consumes. **If dropping a postulate leaves the result standing, it is not listed.**

10.1 Dimensional audit

symbol	quantity	dimension	unit	status
s_i	exposure share of jurisdiction i	$[1]$	dimensionless	defined; non-negative and summing to 1 by P1
W, B_j	bloc shares	$[1]$	dimensionless	derived, $W = \sum_{i \in b} s_i$
x	share transferred into a new bloc	$[1]$	dimensionless	free variable
HHI	concentration index	$[1]$	dimensionless	derived
ρ_{ij}	correlation between shocks	$[1]$	dimensionless, in $[-1, 1]$	**assumed, not measured**
N	effective count	$[\#]$, a count	jurisdictions or blocs	derived as $1/(w'Rw)$
$\ln N$	surplus-term argument	$[1]$	nats	derived; the logarithm of a count is dimensionless
Q	protected capital	$[C]$	index units	defined; **the source supplies no operational unit**
η	conversion efficiency	$[1]$	dimensionless	assumed
δ	erosion rate	$[T^{-1}]$	per period	assumed
C	maintenance cost flow	$[C\,T^{-1}]$	index units per period	assumed
h	headroom	$[C\,T^{-1}]$	index units per period	derived, $h = \eta \ln N - \delta Q - C$
receipts rate	annual settlement throughput	$[T^{-1}]$	fraction of Q per year	**constructed**
sweep period	interval between sweeps	$[T]$	years	**a design choice**
peak balance	transit exposure	$[1]$	fraction of Q	derived, receipts rate times sweep period
efficiency	gain per unit of exposure	$[C\,T^{-1}]$	headroom per unit share	derived

The sweep equation balances exactly: a rate in $[T^{-1}]$ multiplied by a period in $[T]$ gives a dimensionless fraction, so $18.0\%\ \text{yr}^{-1} \times (1/52)\ \text{yr} = 0.346\%$. **That is the only equation in this paper whose dimensions are fully determined by the source, and it is worth naming which one.**

The surplus equation balances only if $\eta \ln N$ carries $[C\,T^{-1}]$. Since $\ln N$ is dimensionless, either η carries $[C\,T^{-1}]$ or $\ln N$ is read as a rate. **The source is silent on which and this paper records the silence rather than resolving it by assertion.** The efficiency of Section 7.5 inherits that ambiguity, since it is h divided by a dimensionless share.

10.2 Significant figures and epistemic status

quantity	figures	status	licence
naive HHI 0.140194, 0.175802, 0.189002	6 dp	**exact**	exact arithmetic on exact constructed shares
effective counts 7.132973, 5.688217, 5.290949	6 dp	**exact to rounding**	reciprocals of exact values
bloc-HHI 0.351928, 0.355000, 0.340000, 0.341962	6 dp	**exact**	exact arithmetic on exact bloc shares
N 2.841490, 2.816901, 2.941176, 2.924302	6 dp	**exact to rounding**	reciprocals of exact values
h to six decimals	6 dp	**exact to rounding**	a logarithm of an exact value times an exact parameter
0.002014, 0.002869, 0.072279	6 dp	**exact to rounding**	differences of the above
3.97%, 3.21%, 1.2368	3–5 sf	**exact to rounding**	ratios of the above
$x^* = 0.200000$, $N = 3.846154$, $h = -0.028524$	6 dp	**exact**	closed form, Section 3.4
4.500%, 1.500%, 0.346%	3–4 sf	**exact to rounding**	divisions of a constructed rate
0.35% as adopted	2 sf	**a design choice**	the weekly peak rounded up from 0.346
$k = 0.483425732988414$ where it appears	15 sf	**DEFINED, exact**	cited from Curzi [2026]; **not used to produce any number here**

No figure in this paper is a measured quantity and therefore no uncertainty interval is attached to any of them. That is the correct treatment of exact arithmetic on stipulated inputs, and it is also the reason none of these numbers is evidence about a real structure.

The 0.35 per cent deserves its own line. It is a design choice rounded up from a computed 0.346 per cent, and the rounding is conservative in the direction that matters. **A reader recomputing Design B from 0.346 rather than 0.35 will get** $N = 2.965107$ **rather than** 2.965384, a difference of 0.000277 in the count, which changes no conclusion and is stated so that nobody has to wonder.

10.3 Limits and boundary behaviour

limit	behaviour	reading
$\rho \to 0$ within a bloc	$N \to$ the naive count	the naive index is the zero-correlation endpoint
$\rho \to 1$ within a bloc	$N \to$ the bloc count	the bloc index is the unit-correlation endpoint

limit	behaviour	reading
one bloc holding everything	$N = 1$, $\ln N = 0$, surplus term zero	no activity below bloc level raises it
$x \to 0$ in the half-share family	$N \to 2.941176$	the baseline; a token position buys a token amount
$x \to W$	$N \to 2.941176$	**transferring the whole donor bloc returns the baseline exactly**, because the two blocs have swapped labels
$x = W/2$	$N = 3.846154$, the maximum	the half-share result
$W \to 0$	the family collapses; no transfer is possible	a bloc with nothing in it cannot donate
$W \to 1$ (one bloc holds everything)	ceiling $\to 2$	splitting a monopolistic bloc in half gives exactly two
sweep period $\to 0$	peak balance $\to 0$	continuous sweeping eliminates transit exposure and is not available
sweep period $\to$ one year	peak balance $\to$ the whole annual receipt	no sweeping at all
$N \to \infty$	$dh/dN = \eta/N \to 0$	returns diminish but never turn negative — **a defect, since cost is not in the input term**

Three boundary readings are load-bearing. The fifth row is the one that makes the half-share result intuitive: transferring the entire donor bloc is exactly as good as transferring none of it, because the arithmetic cannot tell which bloc is which. **A design that moves everything out of a concentrated bloc has achieved precisely nothing**, which is a result worth having on its own.

The eighth row is the limiting case of Curzi [2026]'s Atlantic finding: a structure wholly inside one perimeter can reach at most $N = 2$ by splitting that perimeter in half, and it cannot split a perimeter at all, which is why the Atlantic count is 1.000000 and stays there.

The last row is a genuine defect of the surplus law as used here: because cost enters only through C, the model never says that adding a bloc is not worth it, and the decision to stop must come from outside the equation. Brooks [1975] supplies one bound — communication cost at order $n(n-1)/2$ inside a boundary — and **Section 3.4 supplies a sharper one on the benefit side**, which is why it is in this paper rather than in a note.

10.4 Invariance

Scale invariance in capital. Multiplying every exposure by a positive constant leaves every share and every effective count unchanged. Tested at four scales by construction: the six-, seven- and eight-jurisdiction vectors and the four-bloc vector. **Four scales is not many and Section 11 says so** rather than describing the instrument as scale-free without qualification.

Invariance under relabelling of blocs, and it is exactly the half-share result. The family of Section 3.4 is symmetric under exchanging the donor bloc with its offspring, which is why the maximum sits at the midpoint and why $x = 0$ and $x = W$ give the same answer. **The symmetry is not an artefact; it is the statement that a bloc has no identity beyond its share**, which is a limit on what the count can carry. Curzi [2026] identifies the same invariance as fatal in its own domain: the width statistic is invariant under relabelling and therefore does not determine the chord it is computed from.

Not invariant, and this is where the instrument works. The count is *not* invariant under partition of jurisdictions *within* a bloc: splitting one jurisdiction into two raises the naive count and leaves the bloc count untouched. **That asymmetry is the whole instrument.**

The sweep result is invariant under the receipts rate. The ratio of quarterly to weekly peak exposure is $52/4 = 13.0$ whatever the throughput, because the rate cancels. **That is a stronger statement than the constructed 18.0 per cent supports on its own**, and it is the reason H4 is stated as a ratio.

10.5 Analogy discipline

borrowing	category	justification
exposure as a portfolio; variance under a correlation matrix [Markowitz, 1952]	**1 — derivation**	the exposure vector genuinely is a weighted sum of shocks with a correlation matrix, which is the premise the algebra requires
effective count as a numbers-equivalent [Adelman, 1969, Hill, 1973, Jost, 2006]	**1 — derivation**	$1/\mathrm{HHI}$ *is* the Hill number of order two
similarity-sensitive diversity [Leinster and Cobbold, 2012, Rao, 1982]	**1 — derivation**	the block-of-ones matrix is a special case of the cited general form
open design [Kerckhoffs, 1883, Saltzer and Schroeder, 1975]	**1 — derivation**	the disclosure standard *is* the open-design requirement applied to legal rather than cryptographic structure
coupled-network fragility [Buldyrev et al., 2010]	**2 — shared form**	perimeter coupling is not shown to satisfy the percolation premises; only the shape of the result is used
robust-yet-fragile [Albert et al., 2000, Callaway et al., 2000]	**2 — shared form**	the tolerance asymmetry's shape is used; the degree-distribution mechanism is not claimed, and Doyle et al. [2005] is the reason it should not be
$\ln N$ as the entropy of a uniform distribution [Shannon, 1948]	**2 — shared form**	**information entropy, not thermodynamic entropy**; exact as arithmetic, carrying no thermodynamic content

borrowing	category	justification
the sawtooth transit balance	**2 — shared form**	a stock filled at a constant rate and emptied periodically is a sawtooth in any domain; nothing physical is imported
a bloc as a frozen region [Castellano et al., 2000]	**3 — analogy**	suggestive, used as motivation only; no phase-transition claim is made
normal accidents [Perrow, 1984]	**3 — analogy**	qualitative theory and framing, not a formal result; H1 is Perrow's result in a measurement frame
the half-share proposition	**not a borrowing**	elementary calculus on this paper's own definition; it is proved in Section 3.4 and consumes nothing external

No category-3 borrowing is used as evidence anywhere in this paper.

10.6 Measurement-theoretic status of the scored inputs

The evidence ladder of Section 4.1 is ordinal and is not summed. Its classes are ranked and the assignment identifies the strongest available class, which an ordinal scale licenses. **No numeric evidence score is computed and no threshold is arithmetic**, deliberately.

The two-column evidence table of Section 5 is nominal and is not counted. Five items on one side and four on the other is not four-to-five in favour of anything; the section says which row does the work and why, by argument. **A count of rows would be exactly the arithmetic the scale forbids.**

The de-identified capability comparisons of Section 4.3 and Appendix C are not totalled. Court depth, licence scope, market access and cost have no common scale.

The framework's treatment of Q as a ratio-scale quantity is not tested. Michell [1997] is the strongest available objection: a discipline routinely assumes quantitative structure without testing for it, and the formal remedy is conjoint measurement [Luce and Tukey, 1964], whose conditions establish whether an attribute has additive quantitative structure at all. **The conjoint conditions were not tested for Q, which is stated in unnamed index units, and the ratio-scale treatment is provisional.** Krantz et al. [1971] sets out what would have to hold and Stevens [1946] is the taxonomy. **The concentration instrument is untouched by this objection**, because exposure shares are ratios of capital rather than judgements, and Section 3.4's proposition is closed-form arithmetic.

This paper inherits the framework-level scalar-reduction limitation rather than restating it. Reducing a structure to one effective count and one h, and reducing a jurisdiction's coordination behaviour to a bloc label, are instances of the habit Curzi [2026] names once as a single limitation for the whole framework: where the source literature describes a phenomenon by several jointly varying features, the framework replaces it with one number and does arithmetic on the number. **The limitation is inherited here in full and is not met separately.**

10.7 Postulates

#	statement	status
P1	exposure shares s_i are non-negative and sum to 1	DEFINITION
P2	HHI $= \sum_i s_i^2$	DEFINITION
P3	for n equal shares of $1/n$, HHI $= 1/n$	ARITHMETIC
P4	the aggregate shock is $R = \sum_i w_i X_i$ with Var$(X_i) = 1$	DEFINITION
P5	Var$(R) = \sum_i \sum_j w_i w_j \rho_{ij} = w' R w$	ARITHMETIC
P6	units inside one coordination perimeter act on one decision	POSTULATED (H1)
P7	the surplus law's input term is $\eta \ln N$	DEFINITION (modelling choice, CCP-008)
P8	$\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$	POSTULATED (fixed in advance)
P9	the three Gulf entries occupy one bloc; the contested bridging centre's assignment is run both ways	POSTULATED (P9a) and MODELLING POSITION (P9b)
P10	the constructed exposure vectors of Section 7.1	CONSTRUCTED
P11	renminbi receipts at 18.0 per cent of protected capital a year	CONSTRUCTED
P12	the transit balance is a sawtooth between sweeps	POSTULATED (H4's premise; its falsifier)

10.8 The compounding chain

step	consumes	statement	status
L1	P2, P3	HHI^{-1} returns n exactly for n equal shares	ARITHMETIC
L2	L1	HHI^{-1} is a numbers-equivalent — Adelman [1969]	CITED
L3	L2	it is the Hill number of order two — Hill [1973], Jost [2006]	CITED
L4	P4, P5	with $R = I$, Var$(R) =$ HHI; with R block-of-ones, Var$(R) = \sum_b W_b^2$	ARITHMETIC
L5	L4	naive and bloc indices are one variance under two correlation assumptions	ARITHMETIC

step	consumes	statement	status
L6	P2	$(W-x)^2+x^2$ is minimised at $x=W/2$, giving $W^2/2$	ARITHMETIC — the proposition of Section 3.4
R1a	L5, P6, P10	tempted: naive 0.140194 (7.132973); bloc 0.351928 (2.841490)	ARITHMETIC on CONSTRUCTED inputs
R1b	R1a, P7, P8	naive $+34.81\%$, effective -3.39%, h from -0.122417 to -0.134485	ARITHMETIC
R5	L5, P6, P10	as built: naive 0.175802 (5.688217); bloc 0.355000 (2.816901)	ARITHMETIC on CONSTRUCTED inputs
R3	L5, P9b	stress: bloc 0.341962, $N=2.924302$, $h=-0.124430$, cost 0.002014; the London variant is identical on N and h	ARITHMETIC
R2a	L5, P10	Design A: bloc 0.276562, $N=3.615826$, $h=-0.050138$	ARITHMETIC
R4	P11, P12	peaks 4.500%, 1.500%, 0.346%; ratio $52/4=13.0$	ARITHMETIC on CONSTRUCTED inputs
R2b	R2a, R4	Design B at $x=0.0035$: $N=2.965384$, $h=-0.119548$; 3.97% of the gain for 3.21% of the exposure	ARITHMETIC
R6	L6, P7, P8	ceiling at $x=0.200000$: bloc 0.260000, $N=3.846154$, $h=-0.028524$ $<$ break-even	ARITHMETIC
CR	R1a, R5, R2a	the chapter's own subject jurisdictions fail the test it applies	ARITHMETIC

Every ARITHMETIC row can be redone by hand or in three lines of any language. Every CITED row names one source and no more. Only P6, P8, P9 and P12 are POSTULATED. **P12 is the one carrying an empirical premise and it is the one H4's falsifier names** — that the transit balance is a sawtooth. **A reviewer who does not accept P12 loses R4 and keeps R2a, R3, R5 and R6**, because the bloc arithmetic does not depend on how the 0.35 per cent was derived.

L6 is the only step in this chain that is new to this paper, and it is elementary. It consumes no citation, no postulate and no constructed input; it is a one-line minimisation. **A chain whose novel step is a line of calculus is a chain a reviewer can dispose of in a minute, and that is the intention.**

10.9 Cross-checks performed

- The eight tempted squares were summed independently and reproduced 0.140194; the seven as-built squares reproduced 0.175802; the six designed squares reproduced

0.189002.

- Every bloc vector was checked to sum to 1.000000 before squaring: $0.442 + 0.300 + 0.258$, $0.450 + 0.300 + 0.250$, $0.291 + 0.300 + 0.300 + 0.109$ and $0.3965 + 0.300 + 0.300 + 0.0035$.
- $s_{\max}^2 \le \text{HHI} \le s_{\max}$ holds on all three jurisdiction vectors: $0.040000 \le 0.140194 \le 0.200000$, $0.084681 \le 0.175802 \le 0.291000$ and $0.084681 \le 0.189002 \le 0.291000$.
- The stress cost was checked in both directions: $0.377583 - 0.375570 = 0.002013$ and $0.124430 - 0.122417 = 0.002013$, printed as 0.002014 from the unrounded values $0.1244305 - 0.1224167 = 0.0020138$.
- The half-share result was checked three ways: at the closed form $x^* = W/2 = 0.20$; by direct evaluation of the index at $x = 0.19$, 0.20 and 0.21, which returns 0.260200, 0.260000 and 0.260200, confirming a minimum; and by the end-point identity $\text{HHI}(0) = \text{HHI}(W) = 0.340000$.
- The efficiency ratio was checked two ways: $0.819674/0.663107 = 1.236110$ from the computed gains, and $3.97/3.21 = 1.236760$ from the rounded percentages. **The two agree to four significant figures**, which is the arithmetic cross-check of Section 7.5.
- The sweep ratio was checked against its closed form: $(18.0/4)/(18.0/52) = 52/4 = 13.0$ exactly, independent of the receipts rate.
- Design B was recomputed from the unrounded 0.346% rather than the adopted 0.35%, giving $N = 2.965107$ against 2.965384, a difference of 0.000277.
- The two source discrepancies of Section 7.7 were quantified at 0.000001 in the stress headroom and 0.000183 in the base-case headroom.

11. Limitations

Every number in this paper is computed on constructed inputs and none is a measurement of a real structure. The exposure vectors, the receipts rate, the sweep schedule and the two designs are all **CONSTRUCTED** or described from the source manuscript, and are labelled as such at each point of use. **No first-party measurement of any real system appears anywhere in this paper.** That is this corpus's central limitation, stated in the terms Curzi [2026] states it in — of 27 parameter-by-scope cells, 19 have a defined estimator, 1 has a validated one, and 0 carry a first-party measurement — and a reviewer would be right to press on it before pressing on anything else.

The intra-bloc fracture risk is named and not priced. Sections 3.3 and 4.2 both state it: a concentration index assumes bloc members fail together and does not model them failing against each other, and the region under study has a dated instance lasting three and a half years. **The as-built structure holds entities in two members of that region and this paper attaches no number to the resulting exposure.** That is the largest unmeasured item here.

The bridging centre's assignment is contested and the paper does not settle it. Section 5 sets out the evidence both ways and Section 7.2 runs the arithmetic both ways. **A reader who assigns it to the Western bloc should read $N = 2.924302$ off the table and proceed**, and the fact that the difference is 0.002014 is itself the finding rather than a reason to stop worrying.

H4's falsifier was not tested. No non-sawtooth settlement pattern was constructed, so the premise that sweep frequency controls the exposure is asserted rather than examined against its alternative. **A structure whose receipts must be held to meet a matching obligation would falsify it, and none was modelled.**

The correlation is assumed, not measured. $\rho = 1$ within a perimeter is the conservative end of a family and no correlation was estimated from data. The evidence offered in Section 4.1 is a record of shared mechanisms rather than a measurement of correlation.

Every parameter in the surplus law was chosen. $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ are test parameters fixed in advance and not estimates of any real position, so the break-even $N = 4.172734$ — and therefore the whole force of Section 7.5's ceiling result — is a property of the chosen set. **Change C and the ceiling result changes with it**, and Section 7.5's conclusion should always be quoted with its parameters.

The half-share proposition is arithmetic and carries no empirical content. It bounds a design space; it says nothing about whether any particular perimeter exists, is reachable, or would be independent in the sense the correlation assumption requires. **A ceiling on what a count could be is not evidence about what a structure is.**

The counterparty and currency axes are not measured here. Curzi [2026] measures the mismatch between the entity and currency axes and Curzi [2026] prices the currency drag; this paper measures the jurisdiction axis alone and its conclusions should not be read as statements about the structure's total exposure. **A structure is only as diversified as its most concentrated axis, and this paper measures one of four.**

The reflexive specification is not identified. Section 1 states this and names the remedy — stipulate the reference group and defend the stipulation, following Brock and Durlauf [2001] — and the price is that every bloc assignment here is a modelling position rather than an identified empirical result [Manski, 1993, Shalizi and Thomas, 2011].

The Lucas critique applies to this framework's own parameters and is not answered. Lucas [1976] is the general statement; the specific form is that this framework asserts systems change behaviour when modelled and observed, and then treats η, δ and the bloc assignment as stable under exactly that response. MacKenzie and Millo [2003] documents the loop happening in the field with dates and price data, which is an existence proof for the mechanism and not a defence of the parameters.

The functional form of the surplus input is underdetermined. Latané [1981] proposes a power law with exponent below one and Peters and Gell-Mann [2016] derives a logarithmic form from multiplicative dynamics with no psychological assumption at all. **Nothing here discriminates between them, and the logarithm is used because it is the form this framework carries.** Peters is the stronger of the two available warrants and is cited as such.

Small-sample bias in the effective count is not corrected [Chao et al., 2014]. A six-, seven- or eight-element vector is a small sample for a low-order diversity estimator, and the paper does not resolve whether these vectors are samples or complete censuses.

Scale invariance was checked at four scales only (Section 10.4).

The instrument is least informative in the dominated-unit regime [Golosov, 2010, Dunleavy and Boucek, 2003].

No statistical test was run and no p-value appears, because there is no sample and no null. Every rate is an exact ratio of exact constructed amounts, and **no claim of "no difference" is made from any non-significant result** — the indifference of Section 7.2 is an exact algebraic identity, not a failure to reject.

No scaling, power-law or heavy-tail claim is made and the Clauset–Shalizi–Newman procedure was therefore not run [Clauset et al., 2009]. Any restatement reaching for such a claim must run it first and read Broido and Clauset [2019].

No out-of-sample validation, matched null, information criterion or effect size with an interval appears, because no model was fitted. Bruch and Atwell [2015]'s calibration-and-validation checklist is unmet in every row.

11.1 Rigour checks not met, listed by name

1. **The intra-bloc fracture exposure is named and not priced** (Sections 3.3, 4.2).
2. **H4's falsifier was not tested**; no non-sawtooth settlement pattern was constructed (Section 7.4).
3. **Conjoint measurement conditions untested** [Luce and Tukey, 1964], so the ratio-scale treatment of Q is provisional and Michell [1997]'s objection is conceded rather than answered.
4. **The framework-level scalar-reduction limitation is inherited** [Curzi, 2026] and not met separately.
5. **The reflection problem is stated and not solved** [Manski, 1993].

6. **The Lucas critique applied to this framework's own parameters is unanswered** [Lucas, 1976].
7. **The functional form of the surplus input is undiscriminated** [Latané, 1981, Peters and Gell-Mann, 2016].
8. **Small-sample bias in the effective count is uncorrected** [Chao et al., 2014].
9. **No out-of-sample test, matched null, information criterion or interval** [Bruch and Atwell, 2015].
10. **No first-party measurement of any real system** [Curzi, 2026].
11. **The dimension of η is unresolved** because the source is silent (Section 10.1).
12. **Three of the four axes the method requires are not measured here** (Section 11).
13. **Every legal and jurisdictional statement in Sections 4, 5 and 6 is uncited** (Appendix B).

12. Conclusion

This paper tests three proposed additions against the concentration instrument, reports that two of them fail on the framework's own arithmetic, and establishes one proposition that bounds all three. The following statements are the fixed text of those results.

R1. Adding Saudi Arabia and Qatar to a structure already holding a Gulf position raises the naive effective jurisdiction count from 5.290949 to 7.132973, an improvement of 34.8 per cent, while the effective bloc count falls from 2.941176 to 2.841490, a deterioration of 3.4 per cent: **more jurisdictions, less independence.** The mechanism is exact — the additions occupy an already-occupied bloc and were funded partly out of a different bloc.

R2. The comparison that decides the third case: **Design A**, a full centre at 10.9 per cent of exposure, gives $N = 3.615826$ and $h = -0.050138$; **Design B**, a capped transit position at 0.35 per cent, gives $N = 2.965384$ and $h = -0.119548$, **capturing 3.97 per cent of A's measured gain for 3.21 per cent of A's exposure** — a ratio of 1.2368, so the capped position buys 23.61 per cent more headroom per unit of perimeter exposure.

R3. The bridging slot is indifferent on the bloc axis: **10.9 per cent held in the bridging centre treated as Western, or in the Western centre itself, gives identical blocs 29.1 / 30.0 / 40.9, bloc-HHI** 0.341962**,** $N = 2.924302$ **and** $h = -0.124430$ **— a stress cost of** 0.002014 against the bridge base case, while the naive counts differ by 0.841640. **The leg must be justified on axes the index does not measure.**

R4. The 0.35 per cent transit figure is derived from sweep frequency rather than assumed: **renminbi receipts at 18.0 per cent of protected capital per year peak at 4.500 per cent quarterly, 1.500 per cent monthly and 0.346 per cent weekly, and the weekly peak is adopted.** The quarterly-to-weekly ratio is $52/4 = 13.0$ exactly and is independent of the receipts rate.

R5. The variant actually built — declining the family-office vehicle and adding only a regional headquarters justified by contracted revenue — reads **naive HHI** 0.175802 **(**5.688217 **effective jurisdictions) and bloc-HHI** 0.355000 **(**2.816901 **effective blocs), below the** 2.941176 **baseline it was proposed to improve.** It was taken anyway, at a measured cost of 0.124275 in effective bloc count, because it was justified by revenue and not by resilience. **The instrument is a measurement, not a veto; what it does is make the cost visible and force it to be paid consciously.**

Counter-result. The subject jurisdictions of this paper fail the test it applies: Gulf additions raise the naive count by 34.8 per cent while lowering the effective bloc count by 3.4 per cent, and the bridging leg's measured bloc contribution is 0.002014 **— effectively nil.** And the instrument prices the position rather than the hazard: a share placed in a perimeter the structure does not otherwise touch raises the effective count whether that perimeter is benign or coercive, and the index has no term for the difference.

One further result is established as method rather than as measurement. Moving a share x out of a bloc of share W into a new bloc gives $\mathrm{HHI}(x) = (W - x)^2 + x^2 + \sum_{j \geq 2} B_j^2$, which is minimised at $x^* = W/2$. **The optimal transfer is exactly half the donor bloc, and transferring all of it is exactly as good as transferring none.** For $W = 0.40$ with the other blocs at 0.30 and 0.30, the ceiling is **bloc-HHI** 0.260000**,** $N = 3.846154$ **and** $h = -0.028524$**, against a break-even**

count of 4.172734 at the stated parameters. **No transfer of exposure into a fourth perimeter, of any size and at any placement, brings this structure to the build threshold** — a design-space bound obtained in one line that mentions no jurisdiction at all.

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every case in this paper — the exposure vectors, the receipts rate, the sweep schedule, the two designs and the entity dispositions — is **CONSTRUCTED** and is labelled as such at each point of use. The numerical results are exact computations on those constructed inputs, carried at full precision and rounded only for printing.

All figures were produced by the scripts named in their captions (`figs/src/CCP-057_fig01.py` through `figs/src/CCP-057_fig05.py`), which import the shared figure style module and read their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

Jurisdictional and constitutional descriptions in this paper are taken from the source manuscript and not from primary legal sources. The bibliography contains general comparative, treaty and standard-setting material only; it contains no statutes, regulations, court rules, registry documents, executive orders or case reports for any individual jurisdiction. Appendix B lists the claims affected. **Any reader relying on those descriptions must verify them against the official sources and record the date of that verification.** The banking and currency effects of the episode described in Section 4.2 are specifically marked for verification and are not characterised quantitatively here.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. All persons, entities, holdings and arrangements described are constructed for the purpose of exercising the measurement. **Section 4.2 describes a political event; it is described as a structural fact about coordination and no characterisation of any state's conduct is offered or intended.**

Preregistration

The parameter set $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ and the break-even criterion $N \geq e^{0.5/0.35} = 4.172734$ were fixed by Curzi [2026] before the measurements reported here and are not adjusted anywhere in this paper. **No structure or variant measured in this paper clears**

that break-even, and Section 7.5 establishes that none in this family could, which is reported in the same language a clearance would have received. The falsifiers attached to H1 through H5 were stated before the corresponding measurements were run; **H4's falsifier was not tested, and Section 7.8 reports that rather than treating the hypothesis as passed.** No other element of this paper was preregistered.

Legal, tax and investment advice

This paper is analytical and is not legal, tax, or investment advice. It describes entity types, jurisdictions, licensing regimes, headquarters programmes, arbitral seats and settlement arrangements for the purpose of measuring concentration. No recommendation is made in respect of any particular person, asset, structure, jurisdiction or arrangement. Constitutional arrangements, national-security legislation, licensing scope, headquarters incentive conditions, substance requirements, treaty anti-abuse provisions, exchange obligations and clearing arrangements all change and vary by jurisdiction and by activity; **every such statement in this paper is carried from the source manuscript, is flagged in Appendix B, and requires confirmation with qualified counsel in the relevant jurisdiction before any reliance.**

Appendix A. The bloc-addition reporting standard

The jurisdiction-axis reporting standard is stated by Curzi [2026] and extended by Curzi [2026] and Curzi [2026]. This appendix states what a *proposed addition* additionally requires, because the three cases here differ from one another chiefly in what was reported before the decision rather than after it.

```
 1  the bloc the addition would occupy, named, with its evidence
    class
 2  the position the addition is FUNDED FROM, and its bloc
 3  the naive count and the effective bloc count BEFORE and AFTER,
    both printed
 4  the donor bloc's share W and the ceiling W^2/2 + sum B_j^2,
    computed BEFORE any entity is chosen
 5  the gain in h, and the exposure placed inside the new perimeter,
    as a ratio
 6  the same measurement under the contested assignment, if any,
    with the stress cost printed
 7  the capability being bought, written down separately from the
    measurement, in one sentence per item
 8  for any transit position: the receipts rate, the sweep period,
    the peak and average balance, and the cap as a document
 9  the intra-bloc fracture exposure, named, and stated as
    unpriced if it is
10  the date, and the diff against the previous measurement
```

Item 4 is the item this paper adds and it is the cheapest item on the list. It costs one line of arithmetic and it disposes of any proposal whose ceiling is below its target before anybody reads a term sheet. **Item 2 is the item whose omission produced Section 7.1's result**: an addition inside an occupied bloc is neutral, and the same addition funded out of an extra-bloc position is harmful, and only item 2 distinguishes them.

Item 7 is not optional and is not a formality. Section 7.2 shows a leg whose measured bloc contribution is 0.002014 and whose capability contribution is substantial. **A leg justified on the measurement is a leg sized on the measurement, and it will be defended on the wrong grounds the first time its fee is questioned.**

Appendix B. Claims requiring verification against primary sources

The following statements rest on legal, constitutional or regulatory detail that changes, and **none is cited to the bibliography, which contains no primary sources for any individual jurisdiction**. They are reproduced from the source manuscript and require confirmation with counsel in the relevant jurisdiction before any reliance. **Record the date of each confirmation**; a statement of a regime without a date is not a statement of that regime.

1. The establishing law of the non-territorially-confined financial centre described in Section 4.3, its licensing scope, and the extent to which a licensed entity may operate outside the centre's own district.
2. The headquarters programme described in Section 4.3: its licensing conditions, the mandatory activities, the procurement rule that conditions access on presence, the incentive scope and its duration.
3. **The dates and terms of the regional rupture described in Section 4.2**, and in particular **the banking and currency effects, which are asserted here only qualitatively.** The offshore-currency movement in particular must not be reproduced without contemporaneous market data.
4. The corporate income tax and religious levy treatment described in the source for the jurisdictions of Section 4, and the interaction between them.
5. The arbitration and dispute-resolution provisions relied on in Section 7.6, including the institutional rules and the enforceability of an award in each relevant jurisdiction [United Nations, 1958, United Nations Commission on International Trade Law, 1985].
6. The reception statute described for the bridging centre in Section 5, its schedule of retained legislation and its current form.
7. The commercial court with an opt-in jurisdiction described in Section 7.2, its constitution, and the enforcement route for its judgments including participation in the relevant judgments conventions [Hague Conference on Private International Law, 2005, 2019].
8. The fund vehicle described in Section 7.2, its statutory segregation, and the regulatory status required of its manager.
9. **The March 2022 measures described in Section 5**, their legal basis, their scope, and whether they were adopted with or without a Security Council resolution.
10. The exchange and reporting positions of every jurisdiction in Sections 4, 5 and 6, including first-exchange dates and the model of any bilateral arrangement.
11. The constitutional frame described in Section 6, the national-security legislation of 2020 and 2024, the interpretive power, the judge-designation power, and the 2021 revision of the executive selection system.
12. **The legislation and executive order of 2020 described in Section 6**, the preferential treatments suspended or eliminated, the bilateral agreements suspended, and the current status of each.
13. The linked exchange rate arrangement described in Section 6 and its current terms.
14. The clearing-system participation and offshore-currency capability described in Sections 5 and 6, and **the extent to which participants in the second clearing system use the first system's messaging layer**, which Curzi [2026] identifies as the decomposition that determines whether a second rail is real.
15. The treaty anti-abuse provisions referred to in failure mode 5, including any principal-purpose test.
16. The nominee-shareholding prohibition referred to in the source for the operating jurisdiction, and its enforcement.
17. All exposure shares, receipts rates, sweep schedules and design variants in this paper are constructed. **They are not observed allocations, bank statements or treasury policies, and no real schedule was used to produce them.**

All numerical results in this paper were computed from the stated inputs at full precision

and rounded only for printing, and the places where a printed value in the source differs from the exact computation are reported in Section 7.7 rather than harmonised.

Appendix C. Condensed assessment, de-identified

No total is computed: the items are heterogeneous and no common scale exists on which they could be added (Section 10.6). Every characteristic below is described from the source manuscript and is uncited (Appendix B). Jurisdictions are de-identified in the capability table because that table ranks them on legal capability the bibliography cannot support; they are named where they carry a bloc assignment with its evidence, because there the assignment is the argument.

Bloc assignment and evidence, by jurisdiction.

jurisdiction	bloc	strongest evidence	job
United Arab Emirates	**BRIDGE**	outside the Western designation apparatus; own listing decisions; dual-rail access	**HOLDING** and apex; never settlement
Saudi Arabia	**BRIDGE, same bloc**	common regional assessment body; common peer review; shared correspondent policy; pegged currency	**OPERATING**, market access only; **never holding**
Qatar	**BRIDGE, same bloc**	identical on every operative test	**OPERATING** where there is domestic business; otherwise nothing
Singapore	**BRIDGE in commercial function, Western-leaning under stress**	the contested assignment of Section 5, run both ways in Section 7.2 at a cost of 0.002014	**OPERATING**: treasury, fund wrapper and contracting; **never the apex**
Hong Kong	**inside two perimeters and independent of neither**	central legislation and interpretive power on one side; dollar correspondent dependency and a linked rate on the other	**TRANSIT ONLY**, capped by a written balance limit and a weekly sweep

Ranked capability, de-identified.

requirement	best fit	why
a common-law licence usable across a national economy	Gulf centre A	the licence is not territorially confined, which converts an enclave into a regime

requirement	best fit	why
access to a very large state customer conditioned on presence	Gulf centre B	a headquarters programme that is a market-access requirement rather than a tax incentive
a fund wrapper with statutory segregation and a regulated manager	Asian centre A	the vehicle, the manager regime and the custody options exist together
a commercial court with an opt-in jurisdiction and a common-law bench	Asian centre A	the court, the arbitral institution and the mediation centre are co-located
settlement in the second major currency at depth	Asian centre B	the deepest offshore capability in that currency, and no substitute at that depth
an arbitral seat recognised on both sides of the friction line	Asian centre B	**bought through a clause at no balance-sheet cost**, which is Design B's central move

The item that decides every row of the first table is the correspondent relationship, because it is the thing that can be withdrawn in a week. **No jurisdiction in the table can legislate itself a different one**, which is why the bloc assignments are the least contestable calls in the paper and why the contested one costs only 0.002014 to get wrong.

Chapter 58

Pricing an Uncertain Bloc Call: AIFC Option Value, Thailand's Share of the Gain, and the GIFT City Corridor

Abstract. We price three jurisdictional positions whose value depends on facts that cannot be settled in advance. The Astana International Financial Centre is treated as a bet on a bloc classification: if Kazakhstan is a genuinely separate coordination perimeter the effective bloc count rises from 2.941176 to 3.174603 and the surplus improves by 0.026731; if the call is wrong it falls to 2.898551 and the surplus loses 0.005110. Setting E[dh] = 0.0318401p − 0.0051096 to zero gives a break-even probability of 16.05 per cent, and adding an incremental maintenance cost c gives p = (c + 0.0051096)/0.0318401, so the position is worth nothing at any probability once c exceeds 0.0267305 — 8.9 per cent of the group's entire maintenance budget. The decisive argument is therefore the invoice and not the probability. Thailand carries 18.2 per cent of the designed structure and accounts for 79.74 per cent of the whole diversification gain; relocating its operating leg to the United Kingdom adds 66.24 index units of cost at 0.004545 storage-law units each, raising C to 0.601091 and driving h to −0.423508. The Indian corridor is measured honestly against itself: removing it gives N = 2.5546 and h = −0.136733, so its 11.8 per cent contributes 0.014317, only 3.79 per cent of the total gain, and moving the same share to Thailand leaves the bloc count unchanged at 2.9412 while dropping the naive count to 4.3112. A recomputation of the source's own arithmetic corrects its published net-contribution figure from 0.014316 to 0.014317, and three further discrepancies in the source are printed rather than silently repaired.

Keywords: option value; bloc classification; break-even probability; operating substance; corridor exposure; dwell time; job concentration

Series: D. Structure and Resilience **Paper:** CCP-058

1. Introduction

Three decisions about a cross-border holding structure cannot be settled by argument because each of them turns on a fact that is not available when the decision has to be made. Whether a given jurisdiction is a genuinely separate coordination perimeter is a forecast. Whether an operating leg should sit where its inputs are cheap or where its counterparties are familiar is a cost comparison nobody runs. Whether a corridor entity earns its keep is a counterfactual

that is almost never computed, because the gross benefit is the number that gets presented and the maintenance cost is the number that does not.

This paper prices all three against one instrument. The instrument is the concentration measurement of Curzi [2026] — HHI as the sum of squared exposure shares and the effective count as its reciprocal, with a naive four-jurisdiction count of 3.41 collapsing to exactly 1.00 once coordination blocs are recognised, and a redesigned six-jurisdiction structure reaching a bloc-HHI of 0.340000 and 2.941176 effective blocs with Western-perimeter exposure falling from 100.0% to 30.0%. Everything below is that measurement coupled to the surplus law $h = \eta \ln N - \delta Q - C$ and then differenced.

The paper's contribution is a method for deciding a marginal position, and four measurements made with it.

First, an uncertain classification is converted into a threshold. A bloc call that cannot be settled is priced as a bet. Where a candidate jurisdiction's inclusion raises the effective bloc count on one classification and lowers it on the other, the two classifications bracket the outcome, and the expected change in the surplus balance is linear in the classification probability: $E[\Delta h] = (G + L)p - L$. On the constructed structure carried through Series D, the right call gives $G = +0.026731$ and the wrong call $L = 0.005110$, so $E[\Delta h] = 0.0318401p - 0.0051096$ and the break-even probability is $p^* = 0.160476$. **The move that matters is not the estimate; it is the replacement of an argument nobody can settle with a threshold anybody can assess.**

Second, the threshold is then shown not to be the binding constraint. Adding an incremental maintenance cost c gives $p^* = (c + L)/(G + L)$, and p^* reaches 1 at $c = 0.0267305$ — **8.9102% of the whole structure's maintenance budget of** $C = 0.30$**.** Above that cost the position is value-destroying at every probability, including certainty. **The decisive argument for a small marginal jurisdiction is the invoice, not the geopolitics**, and that is a general property of small positions rather than a fact about any one place: structure cost is close to fixed and structural benefit is proportional to size.

Third, the cost of operating substance is measured against the whole diversification gain and is found to exceed most of it. On a constructed relative-cost model, relocating the operating leg from the low-cost non-aligned site to the high-cost Western one adds 66.24 index units at $0.30/66 = 0.004545$ surplus units each, raising C to 0.601091 and driving h from -0.122417 to -0.423508. The deterioration, 0.301091, is **79.74% of the entire** $\eta \ln N = 0.377583$ **that the jurisdictional redesign bought**, and the effective bloc count does not move at all. A structure that arranges its blocs well and then sites its substance expensively has spent the gain before collecting it.

Fourth, a corridor leg is measured against its own counterfactual and its bloc-count case fails. Moving the corridor jurisdiction's 11.8% share to the other member of the same bloc leaves bloc-HHI at 0.340000 and N at 2.941176 to six decimal places, while the naive jurisdiction count falls from 5.290949 to 4.311200. Moving the share out of the bloc entirely gives bloc-HHI 0.391448, $N = 2.554618$ and $h = -0.136733$, so the leg's gross contribution is 0.049317 and, after its own 0.035000 of maintenance, its net contribution is 0.014317 — **3.79% of the total diversification gain.** Positive, real, and thin enough that one additional function on the same vehicle turns it negative.

The three results have one shape. In each case a quantity that looks like the argument — the classification, the bloc arrangement, the gross diversification benefit — is dominated by a quantity that looks like an administrative detail: the fee, the wage bill, the marginal entity cost. That is not a rhetorical observation. It is what the arithmetic says, and it says it three times.

What the paper does not do is estimate any of the probabilities it prices. The threshold $p^* = 0.160476$ is exact given the share vector; whether the real probability exceeds it is a judgement this paper does not make and cannot support, because the bibliography contains no primary sources for any individual jurisdiction. Section 13 records that gap and Appendix A itemises it.

Positioning. The collapse result on which everything here depends — several jurisdictions inside one perimeter behaving as one — is not new machinery and should not be presented as

such. A small correlated shift producing a discontinuous collective outcome is the threshold-and-tipping mechanism: Granovetter [1978] establishes that the distribution of individual thresholds rather than the mean decides whether a collective outcome tips, and Schelling [1971] is the canonical demonstration that a mild individual preference produces a sharp aggregate segregation. Schelling [1978] generalises the point. **The strongest available warning against reading an individual-level parameter straight into an aggregate is exactly what a bloc assignment does**, and naming that lineage places the result inside an established literature rather than outside one. Castellano et al. [2000] supplies the closest formal analogue of a bloc — the same interaction that raises local agreement freezes global diversity into non-interacting regions — and a bloc, in that vocabulary, is a frozen region.

The reflexive loop and its identification problem. The bloc assignment is a forecast that conditions the behaviour it forecasts, which is the reflexive loop $T(x) = W(x, M(x))$ of Curzi [2026], whose founding reframing is that comfort is the felt total of present costs so that the present state is a local minimum of the comfort landscape by definition rather than by luck. **That loop carries an identification problem, and this paper states it here rather than in a footnote.** Manski [1993] proves that a linear-in-means reflexive specification — a group's mean behaviour influencing its members while the members constitute the mean — is **not identified from observational data**: the endogenous social effect cannot be separated from contextual effects or from correlated unobservables, however much data is collected. Shalizi and Thomas [2011] reinforce it: influence and selection are generically confounded in observational network data. **No amount of observation of jurisdictional behaviour could therefore confirm a bloc assignment derived reflexively.** The remedy adopted here is the fourth of the four known ones — **the reference group is stipulated and the stipulation defended**, which is the route Brock and Durlauf [2001] shows can be made rigorous — and its price is that every bloc assignment in this paper is a **modelling position and not an identified empirical result**. Section 3.7 states the price where the assignment enters.

A note on what half of this material is. The source manuscript from which this paper is converted is, in substantial part, a descriptive comparative survey of three jurisdictions. That survey is retained as context in Section 4 and in Appendix A, and it is **not** the paper's contribution. **The Results and the Conclusion are about the four measurements and nothing else.** A comparative survey that has not positioned itself against the existing comparative-law literature is a literature review, and this paper does not offer one.

2. Hypotheses

Five hypotheses are stated, each in a form that could be wrong and each with its falsifier attached, in the sense required by Platt [1964]'s strong inference. Meehl [1967] is the relevant warning: a directional prediction is cheap, and the defensible response is a point prediction with constants attached, which is what H1, H2 and H4 do. **One of the five is reported in Section 8 as not met.** Statements that nothing could falsify are listed after the hypotheses and labelled as definitions or modelling choices rather than counted among them.

H1 — the invoice hypothesis. For a small marginal jurisdictional position whose bloc classification is uncertain, the binding constraint is the position's own maintenance cost and not the classification probability. *Predicted:* with $G = 0.026731$ and $L = 0.005110$, the break-even probability is $p^* = L/(G + L) = 0.160476$ at zero incremental cost and reaches 1.000000 at an incremental cost of $c = 0.0267305$, which is 8.9102% of $C = 0.30$; at $c = 0.010$, one third of that ceiling, the required confidence is already 0.474547. *Falsified by:* a measured all-in annual maintenance cost for a small ring-fenced holding vehicle below 0.005 in surplus units — 1.7% of C — at which the break-even probability stays under 0.32 and the classification question, not the invoice, would decide the position.

H2 — the substance-cost hypothesis. The siting of operating substance carries more of the design's economics than the jurisdictional arrangement does. *Predicted:* on the constructed cost model, relocating the operating leg to the highest-cost candidate raises C by 0.301091 and

consumes 79.74% of the entire $\eta \ln N = 0.377583$ diversification gain, while leaving N at 2.941176 exactly. *Falsified by:* a measured relative cost index for the same substance in the destination jurisdiction below 2.297943, at which the relocation would consume less than half the gain, or below 1.000000, at which it would be free.

H3 — the corridor hypothesis. A second jurisdiction in an existing bloc raises the effective bloc count. *Predicted:* adding the corridor jurisdiction at 11.8% of exposure raises N above the 2.941176 it would otherwise take. *Falsified by:* a reallocation of that share to the other member of the same bloc that leaves N unchanged. **This is the hypothesis reported as not met.**

H4 — the marginal-leg hypothesis. A jurisdiction leg contributes to the surplus balance only after clearing its own maintenance cost, and the net figure rather than the gross one should govern the decision. *Predicted:* the corridor leg's gross contribution is $0.377583 - 0.328266 = 0.049317$ and its maintenance cost is 0.035000, so its net contribution is 0.014317, which is 3.79% of the total diversification gain and less than a third of the gross figure. *Falsified by:* a marginal all-in entity cost materially below 0.035 in surplus units, at which the gross and net readings would not diverge enough for the distinction to change any decision.

H5 — the dwell hypothesis. A corridor's exposure is set by how long value rests inside the perimeter and not by throughput, counterparty count or hop count. *Predicted:* exposure equals dwell divided by the year, giving 0.27% at one day, 1.92% at seven, 8.22% at thirty, 24.66% at ninety and 49.32% at one hundred and eighty, for the same corridor and the same annual volume. *Falsified by:* an exposure that varies with throughput or with the number of intermediaries at fixed dwell. Curzi [2026] reports the same invariance on the settlement axis, where the transit-loss model $E[\text{loss}] = V(1-e^{-\lambda t})(1-r)$ is invariant to hop count at USD 68,000 whether the value crosses one hop or eight.

Not hypotheses, and labelled as such. Four statements in this paper are not falsifiable as they stand, and each is named here for what it is.

- *The job-concentration diagnostic* $S = \sum_j v_j^2$ with effective jobs $1/S$ is a **definition**, and the 2.00 band boundary is an algebraic property of a vector on the simplex, not a finding. Its empirical content is entirely in whether a measured job vector is the entity's real one.
- *The logarithmic form of the input side*, $\eta \ln N$ rather than ηN, is a **modelling choice** carried from Curzi [2026], whose surplus functional $f(u) = \ln(1 + (N-1)u)$ is uniquely characterised by angle-only dependence, vanishing at parallel, an affine effective count and normalisation. It is not tested here and Section 3.9 states what would discriminate it.
- *The parameter set* $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ is a **modelling choice**, fixed in advance of every measurement below so that no criterion is adjusted to fit, which is the preregistration discipline of Nosek et al. [2018] applied to an engineering calculation.
- *The bloc assignments themselves* are **stipulated reference groups**, defended by the perimeter test carried from Curzi [2026] and not identified from data [Manski, 1993, Brock and Durlauf, 2001].

3. Method

3.1 The instrument carried in, and which effective count it is

The concentration instrument is not restated here; Curzi [2026] establishes it. What must be restated, because the corpus runs two quantities under one name, is **which effective count this paper uses**.

Throughout this paper the effective count is HHI^{-1}, the reciprocal of the sum of squared shares. **That quantity is the Hill number of order two.** Hill [1973] defined the family ${}^{q}D = (\sum_i p_i^q)^{1/(1-q)}$, whose members return the number of equally-abundant types that would give the observed value of the underlying index; at $q = 2$ this is $(\sum_i p_i^2)^{-1}$, which is HHI^{-1} ex-

actly. Jost [2006] is the statement that entropies are not diversities and that their numbers-equivalents are, and Jost [2007] extends the correspondence to partitioned data, which is what a bloc collapse is. The same arithmetic was reached independently by Simpson [1949] in ecology and by Laakso and Taagepera [1979] in political science, and the numbers-equivalent reading of the concentration index itself is Adelman [1969]'s. The index's own history runs through Hirschman [1945] and Herfindahl [1950], with Hirschman [1964] the priority note, and its regulatory thresholds are stated in the merger guidelines [U.S. Department of Justice and Federal Trade Commission, 2023, 2010]. The exponential of Shannon entropy [Shannon, 1948] is the same family at order one, and HHI^{-1} is also the exponential of the Rényi entropy of order two [Rényi, 1961].

The corpus's other effective count, the affine $N_{\text{eff}} = 1 + (N-1)u$ of Curzi [2026] and Curzi [2026c] — where $u = 1 - (x \cdot y)^2$, the surplus is $S = \ln N_{\text{eff}}$, a hundred-person department at $u = 0.05$ has an effective count of 5.95, and the affine form is badged a design axiom rather than a derived result — is not a Hill number of any order and carries no external support. It is a linear interpolation between 1 and N. This paper does not use it, refers to it only by name, and states the distinction because papers citing this one may. Curzi [2026] gives the arithmetic that separates the two.

Two further disciplines from the diversity literature apply and are recorded rather than assumed. Hurlbert [1971] argues that "diversity" has been used for so many distinct quantities that the word carries no content without its formula attached, which is precisely why this section exists. Chao et al. [2014] show that diversity estimators of low order are biased on small samples, and a six- or seven-element exposure vector is a small sample; no correction is applied here and Section 13 records the gap. Golosov [2010] and Dunleavy and Boucek [2003] document that effective-number measures behave poorly in the dominated-unit regime, which is where several of the comparison structures in this paper sit.

No distributional or scaling claim is made anywhere in this paper, and that is stated rather than assumed because the material invites one. Exposure vectors that descend steeply — 29.1 / 18.2 / 16.4 / 13.6 / 11.8 / 10.9 — look like the beginning of a heavy-tail argument. **This paper fits no distribution, characterises no tail, and asserts no power law.** Clauset et al. [2009] is the reason: applied to twenty-four canonical data sets, a principled maximum-likelihood fit with a goodness-of-fit test and a likelihood-ratio comparison against alternatives leaves few of them standing, and a straight line on a log-log plot is not evidence. **That procedure was not run, because no claim requiring it is made.** Stumpf and Porter [2012] adds the reason running it would not settle anything on its own, since many mechanisms produce the same tail, and Broido and Clauset [2019] is the reason the same caution attaches to any borrowing from the scale-free network literature. **Every effective count in this paper is a statistic of a finite exposure vector with at most seven entries, not an estimate of a distribution, and it must not be restated as one.**

3.2 The break-even probability method

The method applies wherever a structural decision turns on a classification that cannot be settled before the decision is taken, and it has three steps.

Step 1 — measure the outcome under each classification. The candidate position is inserted into the exposure vector and the effective bloc count is computed twice: once with the candidate treated as its own perimeter, and once with it folded into the perimeter it most plausibly belongs to if the call is wrong. This gives two values of N and, through $h = \eta \ln N - \delta Q - C$, two values of h.

Step 2 — form the gain and the loss. Write G for the improvement in h if the classification is right and L for the deterioration if it is wrong, both as positive numbers. The expected change is then linear in the probability p that the classification is right:

```
E[dh] = p G − (1 − p) L = (G + L) p − L
```

Step 3 — solve for the threshold. Setting $E[\Delta h] = 0$ gives

$$p^* = \frac{L}{G + L}.$$

The method's whole value is that it does not require the probability to be estimated. It requires only that a reader be able to say whether the true probability is above or below a stated number, which is a far weaker demand than producing the number. Curzi [2026] is the corpus's own bound on how far such a construction may be pushed: its counter-result declines the inflated reading of its own borrowed physics, stating the framework's postulate as a postulate, giving the relation only in its correct form, and requiring that a theorem, an ordinary observer effect and the framework's own conjecture be distinguishable in three separate sentences on demand. **The same discipline applies here. p^* is arithmetic. p is a judgement. They must not be confused, and this paper supplies only the first.**

The status of p itself deserves a sentence, because a reviewer will ask. p is a subjective degree of belief in a proposition about a unique non-repeatable situation, in the sense of Jaynes [2003], and not a limiting relative frequency. Nothing in this paper estimates it, no reference class is constructed for it, and Section 13 records that a threshold expressed in a quantity nobody has calibrated is a weaker object than it looks.

3.3 The cost ceiling

A jurisdiction leg costs money to run whether or not it helps. Let c be the addition to C from operating the candidate vehicle — counsel, filings, audit, registered office, and the front-loaded opinion amortised. Then

```
E[dh] = (G + L) p − L − c = 0     ->     p* = (c + L) / (G + L)
```

Two readings follow immediately and both are used below.

The ceiling. $p^* = 1$ at $c = G$. **At an incremental cost equal to the entire gain from a correct call, the position is worth nothing even under certainty**, and above it the position is value-destroying at every probability. The ceiling is therefore not a soft constraint trading off against confidence; it is a hard one.

The sensitivity. $dp^*/dc = 1/(G + L)$, a constant. On the numbers below that constant is 31.4069 per unit of cost, so a cost of one per cent of C — 0.003 — raises the required confidence by 0.0942. **The threshold is far more sensitive to the fee than a reader who has just been given a section of geopolitics expects it to be.**

This is a specific instance of a general organisational result the corpus uses elsewhere: adding units inside a coordination boundary adds administrative cost at order $n(n-1)/2$ while adding nothing to independence [Brooks, 1975], and the structure of an organisation's output tracks the structure of the organisation that produced it [Conway, 1968]. Curzi [2026b] supplies the gauge in which all of this is measured — the headroom $h_t = \eta \ln N - \delta Q_t - C_t$, with three regimes, a maximum increment $r_{\max} = h_t$, and the separate constraints of admissibility $C_t \leq (1 - \delta)Q_t + \eta S_t$ and sustainability $C_t \leq \eta \ln N - \delta Q_t$ — and a leg whose cost breaches the sustainability constraint has not bought independence, it has bought a bill.

3.4 The substance-cost bridge into the storage law

Comparing operating costs across jurisdictions is a procurement exercise until it is connected to the law that decides whether the structure gains or loses. The connection is a single calibration constant and it is stated explicitly because everything in Section 7 depends on it.

The constructed relative-cost model. For the same operating substance — same headcount, same floor area, same administrative burden — a weighted index is formed over three components with fixed weights. **The indices are constructed model parameters, not market data;**

they illustrate a ratio and not a price. They are stated in full in Section 7.1 so that a reader who disagrees may substitute their own.

The calibration. Curzi [2026] calibrates the specialised structure at 66 index units of annual cost against $C = 0.30$, so one index unit is $0.30/66 = 0.004545$ surplus units, recurring. **That constant is a stipulation, not a measurement**, and it is the only bridge between the cost model and the surplus law. Every magnitude in Section 7 inherits its arbitrariness: the ratios are meaningful and the absolute figures are not, because Curzi [2026] supplies no operational unit for Q, C or h.

The consequence. For an operating entity of u_0 index units at relative index r,

$$\Delta C = u_0(r-1) \times \frac{0.30}{66}, \qquad \Delta h = -\Delta C,$$

since the relocation changes no share and therefore no N. **The relocation is pure cost: the entire effect passes through C and none of it through $\eta \ln N$.** That is what makes the comparison clean and it is also what makes the result uncomfortable.

3.5 The job-concentration diagnostic

An entity's activity is written as a vector over three functions — holding, operating and bridging — with components summing to one. The diagnostic is the concentration instrument applied to the entity rather than to the structure:

```
S = sum of squared job weights        effective jobs = 1 / S
```

with bands: below 1.50 SPECIALIST, 1.50 to 2.00 STRAINED, above 2.00 OVERLOADED. **The 2.00 boundary is algebraic and carries no empirical content**, since $1/S > 2$ is equivalent to $S < 1/2$, which is a condition on a vector in the simplex and not a finding about entities.

The measurement-theoretic status of the job vector is not optional and is not favourable. The components are treated as ratio-scale shares of an entity's activity, and that treatment is not tested. Michell [1997] is the strongest objection available: the social sciences routinely assume quantitative structure without testing for it. The formal remedy is conjoint measurement [Luce and Tukey, 1964], whose conditions establish whether an attribute has additive quantitative structure at all, within the representational framework of Krantz et al. [1971] and against the scale-type discipline of Stevens [1946]. **The conjoint conditions were not tested for any job vector in this paper, and the effective-jobs figures are therefore provisional as ratio-scale quantities.** What survives without the test is the *ordering* — an entity with a clearly dominant function scores lower than one without — and the exact 2.00 threshold, which is algebraic. The magnitudes do not survive it. Section 13 records this.

3.6 The dwell parameterisation of a corridor

A corridor is a rail and not a store, and the distinction is a measurement rather than a metaphor. A corridor's exposure to an action taken inside its perimeter — a capital-control measure, a regulatory suspension, an operational failure — is not its annual throughput; it is the fraction of that throughput sitting still inside the perimeter at the moment the action lands. For an average dwell of t days,

$$\text{exposure} = \frac{t}{365}.$$

The only variable is time. Same corridor, same counterparties, same jurisdiction, same annual volume: exposure at a seven-day dwell is a twenty-fifth of exposure at a one-hundred-and-eighty-day dwell. Curzi [2026] reaches the same conclusion on the settlement axis by a different route, showing that expected transit loss is invariant to hop count and depends entirely on the exposure window.

The functional form is the simplest one available and it is a modelling choice. It assumes throughput arrives uniformly and dwells for a common time. A burst arrival process with the same mean dwell gives a different instantaneous exposure distribution and the same mean, and nothing here tests which is right.

3.7 The bloc call is reflexive, and what that costs

Every measurement below depends on a bloc assignment, and a bloc assignment is a forecast about how a jurisdiction will behave under a coordinated action — a forecast that, once adopted by enough structures, changes the behaviour it forecasts. Merton [1948] named the mechanism; Curzi [2026] carries it as the reflexive loop $T(x) = W(x, M(x))$.

This paper adopts stipulation with a defence and pays the stated price. The blocs are fixed in advance by the perimeter test of Curzi [2026] — two jurisdictions are in the same bloc when a single decision taken in one reliably produces the corresponding effect in the other without a further independent decision — recorded with an evidence class, and stress-tested by running the arithmetic under the contested alternative, which is what Section 6 does. Brock and Durlauf [2001] is the template for making a stipulated reference group rigorous. **No randomisation, instrument or partial-population variation was available, so no bloc assignment in this paper is an identified empirical result.**

Two further objections in the same family are recorded and not answered. **The Lucas critique applies to this framework's own parameters** [Lucas, 1976]: the framework asserts that systems change their behaviour when modelled and then exempts η, δ and its own correlation assumptions from that reflexivity. **That is an inconsistency and it is not resolved here.** MacKenzie and Millo [2003] documents the benign form of the same loop in the field and with dates — a model that did not fit the market when published came to fit it after traders adopted it — and Hacking [1995] states the general form. And Goodhart [1975] with Campbell [1979] give the failure mode that follows adoption: an effective count taken as a target ceases to be a good measure of anything. The only defence offered is the reporting standard of Appendix B, which requires the inputs to be published with the output.

3.8 The disclosure premise

Every position priced in this paper is assessed against one premise, stated as a design constraint rather than an ethical preference: **would the structure still do its job if every element of it were disclosed to every authority with jurisdiction over any part of it?**

The premise is not this work's to claim. Kerckhoffs [1883] stated it in 1883 — a system must remain secure when everything about it except the key is public — and Saltzer and Schroeder [1975] restate it as *open design*, the principle that a protection mechanism should not depend on the ignorance of attackers. Curzi [2026] carries the constraint in full for this series, establishing that channel separation is strengthened rather than weakened by disclosure because its effectiveness is entirely independent of whether anyone knows it is in place. **Citing Kerckhoffs converts a rule that reads as a preference into a recognised engineering standard with a hundred and forty years behind it.**

The same conclusion arrives from the regulatory direction. Automatic exchange of account information [Organisation for Economic Co-operation and Development, 2017], the anti-money-laundering recommendations that supervisors implement in common [Financial Action Task Force, 2012], and the substance and beneficial-ownership work of the base-erosion project [Organisation for Economic Co-operation and Development, 2015] have together removed non-disclosure from the set of available design parameters, whatever a designer's preference. Brandeis [1914] is the older statement of the same instinct from the transparency side; de Goede [2011] documents the listing machinery as a technique of governance, which is the mechanism that makes concealment fail on contact; and Pistor [2019] is the scholarly account of how legal coding manufactures durability and simultaneously its sharpest critique. Hansmann and Kraakman [2000] establishes that asset partitioning is the irreducible function of

entity law — **a separate legal personality is a partition, and its effectiveness likewise does not depend on anyone's ignorance of it.**

The premise is what makes every position in this paper measurable at all, because a position whose value depends on not being seen cannot be priced against a published instrument. Curzi [2026] states the bound on how far it goes: **disclosure is not immunity**, since approximately USD 300 billion of fully disclosed central-bank reserves were immobilised in 2022, so the disclosure premise is a filter that removes structures collapsing on contact with an authority and is not a shield against a policy decision.

3.9 What discriminates the logarithm, and what does not

The surplus law's input term is $\eta \ln N$, and every result in this paper inherits its shape. **The functional form is underdetermined by the evidence and this paper concedes that rather than defending it.**

Three candidate forms fit the qualitative facts equally well. The corpus proposes $\ln(1 + (N-1)u)$ [Curzi, 2026]. Latané [1981] proposes a power law with exponent below one for social impact, which has the same diminishing-increment shape. Briscoe et al. [2006] propose $n \log n$ for network value against the quadratic alternative. **All three are decreasing in marginal return and none of the measurements in this paper distinguishes them**, because every measurement here is a difference between two values of the *same* functional form.

One genuine strengthening is available and should be taken. Peters and Gell-Mann [2016] derives a logarithmic form from multiplicative dynamics **with no psychological assumption at all** — a stronger warrant than the perceptual story of Fechner [1860], because it does not require the perceptual story to be true. Curzi [2026n] identifies the corpus's own functional as a Fechner law with threshold $u_0 = 1/(N-1)$. **Where this paper's arithmetic needs a warrant for the logarithm, the multiplicative-dynamics one is the one it needs; the perceptual one is decoration.** Peters [2019] is the associated warning that an expected value and a time average are different objects, which matters wherever $E[\Delta h]$ is read as something a single structure experiences.

4. The three positions and their institutional setting

This section is context and is not a result. It states what each of the three candidate positions is for, at the level of detail the measurements require and no further. **Every statement in it about the law, regulation, institutions or policy of any individual jurisdiction is described from the source manuscript and is uncited**, because the bibliography for this corpus contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction. Appendix A lists those statements item by item and Section 13 records the gap. **They must be verified against the official sources before any reliance, and no measurement below depends on any of them being true** — each measurement depends only on the constructed share vector and the stated parameters.

4.1 The three positions

position	assigned function	share of exposure	what the measurement decides
a Central Asian financial centre	holding, ring-fenced, plus a costless arbitration option	5.0%	whether an uncertain bloc classification is worth betting on, and at what fee
a South-East Asian operating base	operating, exclusively	18.2%	how much of the diversification gain the siting of substance consumes
a South Asian corridor centre	bridging, exclusively	11.8%	whether a second jurisdiction inside an existing bloc earns its maintenance cost

Where a scored institutional comparison would otherwise attach a rating to a named jurisdiction on a characteristic this paper cannot support, the comparison is de-branded, following the convention of Curzi [2026]. **Named jurisdictions appear below only as labels on a constructed exposure vector and on constructed cost indices**, and not as the subjects of assessed claims.

4.2 The Central Asian position

The candidate is a legal enclave established by constitutional statute within a Central Asian host state, applying a body of law built on the principles and precedents of English law, operating its own court and arbitration centre, conducting its business in English, and exempting participants from corporate income tax for a period measured in decades. It is one of a very small number of jurisdictions anywhere to have received English common law wholesale outside the Commonwealth.

Two features make it a candidate for a fourth bloc and one makes it uncertain. It is simultaneously inside a regional economic union and a regional security treaty led by one great power, a major partner of another, and party to a cooperation agreement with a third — **and that is precisely the configuration that makes the bloc call contestable rather than obvious.** The source manuscript records evidence on both sides, dates it, and routes every item to a verification register. **This paper does not weigh that evidence and does not assert a classification.** It computes what the classification is worth if right and what it costs if wrong, and leaves the weighing to a reader with access to primary sources.

Two structural constraints are carried into the specification and both are measurements rather than opinions. The position is sized at the amount the architecture could lose entirely without failing, which in the constructed case is 5.0% of protected capital. And the vehicle is held directly by the apex rather than through another bloc's holding company, because **a bloc whose operability runs through another bloc is not an independent bloc** — which is the perimeter test applied to the ownership chain rather than to the map.

The costless part of the position is separable from the capital part. An arbitration clause naming the centre's arbitration institution can be adopted for contracts across the group at no balance-sheet cost, and it enforces through the 1958 Convention [United Nations, 1958] rather than through the centre's own court, whose judgments travel on no treaty network written for them. **A neutral forum that both parties will actually sign is worth more than a superior forum that one party refuses**, which is a general observation about international arbitration for which Born [2021] and van den Berg [1981] are the standard treatments and United Nations Commission on International Trade Law [1985] the model instrument. Curzi [2026] reaches the identical conclusion about a different Asian seat by the same reasoning: its

Design B, a capped transit position at 0.35% of exposure, gives $N = 2.965384$ and $h = -0.119548$, capturing 3.97% of the full position's gain for 3.21% of its exposure.

4.3 The operating base

The candidate is a civil-law South-East Asian state with a codified private law, a registration-based commercial and property regime, and no general domestic private trust. It restricts foreign ownership of many activities by statute, and it provides at least two lawful routes around that restriction for the activity in question: an investment-promotion regime conferring full foreign ownership and land rights for a promoted activity, and a bilateral instrument available to nationals of one particular treaty partner. **The promotion route is conditional and the conditions can be failed; the treaty route depends on a chain of ownership qualifying.** Both are described in Appendix A and neither is cited.

The reason the position is in the paper is not its law. **It is that the same operating substance costs a fraction there of what it costs in the Western jurisdictions the structure is diversifying away from, and Section 7 measures what that is worth.** The source manuscript is explicit that the jurisdiction cannot carry a holding or bridging function and should not be asked to, which the job-concentration diagnostic in Section 10 confirms arithmetically.

One tax development is carried because it changes a widely-held planning position rather than because this paper analyses it: a revenue order effective from 1 January 2024 removed a same-year condition on the taxation of remitted foreign-source income for tax residents. **The consequence recorded in the source is structural rather than fiscal: the residence answer that minimises personal tax is the answer that weakens the operating company's management substance, if the principal is the person taking its decisions.** The resolution is to give the operating company an operator — a resident managing director with real authority, a real salary and minutes recording decisions actually taken — at which point the two questions stop competing. That resolution costs a salary and it is the cheapest item in the specification.

4.4 The corridor centre

The candidate is an international financial services centre inside a South Asian state that operates capital controls, created by treating units inside the centre as non-resident for specified purposes under the state's exchange-control legislation, and supervised by a single unified regulator established by statute in place of four sectoral ones. **It has no separate legal system and no separate court**, which is the defining structural fact about it and the reason the paper assigns it a bridging function and nothing else.

The centre's product is supervision. That is what makes the disclosure premise of Section 3.8 nearly trivial to satisfy there — a registered unit under a published framework, filing accounts, reporting account information, claiming a statutory deduction on its stated conditions — and it is also what caps the position: a regulated unit that conceals from its regulator has lost its registration, which is the only thing it had.

The failing configuration is worth naming because it is the one that gets built, and it is the object of the measurement in Section 10: a corridor vehicle used as a place to hold value, with a corridor function attached as the stated reason. **That fails not because any single element is concealed but because, asked what it does, it says "several things", and an entity that cannot state its function in one sentence has a supervision problem before it has any other kind.**

5. The constructed case and the baseline

CONSTRUCTED CASE. Every fact is invented and is labelled as such at each point of use. The case is the one carried through Series D and established in Curzi [2026]: a dual national, tax resident in one Gulf jurisdiction, with three heirs resident in three different jurisdictions,

one under a forced-heirship regime; manufacturing operations in one South-East Asian state, a trading company in the Gulf, an intellectual-property portfolio, London commercial real estate, and a listed-securities portfolio, with revenue arriving in five currencies.

The share basis is the one Curzi [2026] fixes and it is stated here because a measurement without a stated basis is not reproducible: **the fraction of protected capital whose availability depends on a decision taken inside a given jurisdiction's coordination perimeter.** Not book value, not tax residence, not where the certificate is.

jurisdiction	share	share^2	bloc
UAE	29.1%	0.084681	BRIDGE
Thailand	18.2%	0.033124	NONALIGNED
UK	16.4%	0.026896	WEST
Bermuda	13.6%	0.018496	WEST
India	11.8%	0.013924	NONALIGNED
Singapore	10.9%	0.011881	BRIDGE
total	**100.0%**	**0.189002**	

```
naive HHI = 0.189002        effective jurisdictions = 5.290949
blocs   BRIDGE 40.0%   NONALIGNED 30.0%   WEST 30.0%
bloc-HHI = 0.160000 + 0.090000 + 0.090000 = 0.340000
EFFECTIVE BLOCS N = 1 / 0.340000 = 2.941176

ln(N) = 1.078810     eta ln(N) = 0.377583     h = -0.122417
```

Everything below is a difference against those four numbers. The break-even effective count at the stated parameters is $N^* = e^{0.5/0.35} = 4.172734$, so the baseline structure is below the build threshold before any position in this paper is added, and remains below it after all of them.

6. Results: the bloc call priced

6.1 The two measurements

Five percentage points of protected capital move from the Western position into a ring-fenced holding vehicle in the Central Asian centre, held directly by the apex. The Western share falls from 16.4% to 11.4% and nothing else changes.

Measurement one — the candidate is its own perimeter.

jurisdiction	share	share^2	bloc
UAE	29.1%	0.084681	BRIDGE
Thailand	18.2%	0.033124	NONALIGNED
Bermuda	13.6%	0.018496	WEST
India	11.8%	0.013924	NONALIGNED
UK	11.4%	0.012996	WEST
Singapore	10.9%	0.011881	BRIDGE
Kazakhstan	5.0%	0.002500	CENTRAL ASIA
total	**100.0%**	**0.177602**	

```
naive HHI = 0.177602        effective jurisdictions = 5.630567
blocs   BRIDGE 40.0%   NONALIGNED 30.0%   WEST 25.0%   C.ASIA 5.0%
```

```
bloc-HHI = 0.160000 + 0.090000 + 0.062500 + 0.002500 = 0.315000
EFFECTIVE BLOCS N = 1 / 0.315000 = 3.174603

ln(N) = 1.155183     eta ln(N) = 0.404314     h = -0.095686
```

Against the baseline, N rises by 0.233427 and h by 0.026731. Western-perimeter exposure falls from 30.0% to 25.0%.

Measurement two — the call is wrong. Suppose the candidate does not constitute an independent perimeter and belongs with the non-aligned group.

```
blocs   BRIDGE 40.0%   NONALIGNED 35.0%   WEST 25.0%
bloc-HHI = 0.160000 + 0.122500 + 0.062500 = 0.345000
EFFECTIVE BLOCS N = 1 / 0.345000 = 2.898551

ln(N) = 1.064211     eta ln(N) = 0.372474     h = -0.127526
```

h falls by 0.005110 against the baseline. **The move is value-destroying if the call is wrong, and the reason is instructive rather than arbitrary: the five points came out of the smallest bloc and went into one that was already equal-largest, so concentration rose.** A reader carrying only the maxim "diversify into more jurisdictions" would not have seen that coming, and it is the argument of Curzi [2026] arriving in a live case. Curzi [2026] reports the same sign on a different pair of additions: two Gulf additions raise the naive effective count from 5.29 to 7.13, an increase of 34.8%, while the effective bloc count *falls* from 2.94 to 2.84, a decrease of 3.4% — more jurisdictions, less independence.

6.2 The option line and the break-even probability

With $G = 0.026731$ and $L = 0.005110$:

```
E[dh] = 0.0318401 p - 0.0051096
E[dh] = 0  ->  p* = 0.0051096 / 0.0318401 = 0.160476
```

The position pays, before costs, at any probability above 16.0476% that the candidate is a genuinely independent coordination perimeter. That is a low threshold. **Whether the true probability clears it is not decided in this paper**, because the evidence that would decide it — the host state's conduct under coordinated action, its designation practice, its payment-system behaviour — is jurisdiction-specific material for which this paper has no citable source. Section 13 records that, and it is the single most important thing this result does *not* establish.

6.3 The cost ceiling, and why it decides the position

Adding the incremental maintenance cost c gives $p^* = (c + 0.0051096)/0.0318401$:

c	as a % of $C = 0.30$	break-even p^*
0.000000	0.00%	0.160476
0.005000	1.67%	0.317500
0.010000	3.33%	0.474547
0.020000	6.67%	0.788639
0.026000	8.67%	0.977094
0.026731	**8.91%**	**1.000000**

Read the table from the right-hand column and the position changes character. The interesting argument about the host state's alignment is settled at $p^* = 0.16$; the argument that

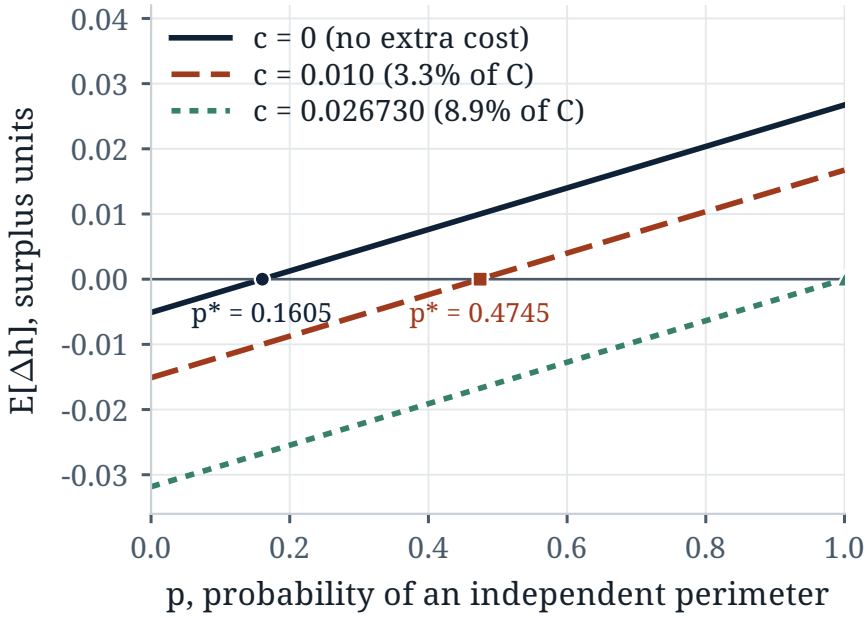


Figure 58.1: **The bloc call as an option line: expected surplus change is linear in the classification probability, crossing zero at p = 0.1605 with no incremental cost, at 0.4745 once the leg costs 3.3 per cent of the maintenance budget, and never at all once it costs 8.9 per cent.** Computed from the two bloc measurements of Section 6.1 at the stated parameters; script `figs/src/CCP-058_fig01.py`.

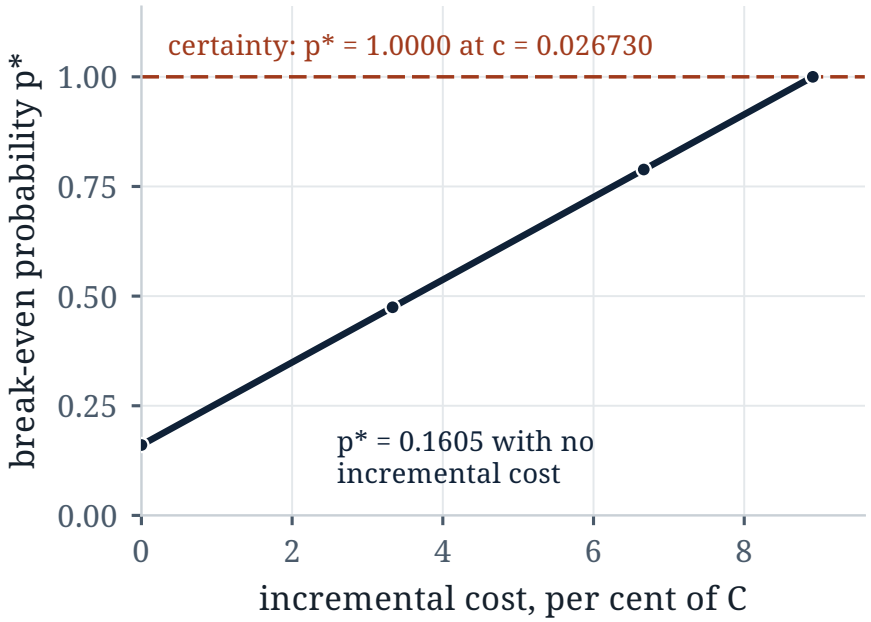


Figure 58.2: **The break-even probability rises 31.4 points for every one per cent of the maintenance budget the leg consumes, reaching certainty at 8.91 per cent: above that cost the position is worth nothing at any probability whatever.** Computed from p* = (c + L)/(G + L) with G and L from Section 6.1; script `figs/src/CCP-058_fig02.py`.

decides the position is the invoice. **The leg must cost less than 8.91% of the group's entire structure-maintenance budget to be worth building even under certainty, and at one third of that cost the required confidence is already close to a coin flip.** H1 is supported in the model.

The result generalises beyond this jurisdiction, and that generalisation is the useful part. **For any small position in a fourth bloc, the binding constraint is not the geopolitics; it is that a small position carries a full jurisdiction's maintenance cost.** It is the same asymmetry Curzi [2026] finds when it computes that the designed structure carries protected capital only up to $Q = 3.88$ index units before erosion outruns surplus: **structure cost is close to fixed and structural benefit is proportional to size.** A five per cent position in a jurisdiction with a mature, cheap professional-services sector works; the same position where every question requires a bespoke opinion may not.

The design conclusion follows directly and it is not abstention. **It is: once, simply, and without accretion.** One vehicle, one asset class, no operating activity, no banking dependency, no cross-collateralisation, standard documents reused annually. A second vehicle in the same centre would roughly double c and would not move the bloc share at all — which is the subdivision failure mode of Curzi [2026] arriving where it bites hardest.

6.4 The residual, stated plainly

- N rises from 2.941176 to 3.174603 against a break-even of 4.172734. **The structure remains below the build threshold.** The leg closes 18.95% of the remaining gap in N and 0.026731 of the 0.122417 shortfall in h, which is 21.84% of it.
- The classification is the position's whole thesis and it is uncertain. It is recorded and dated, with mirroring of another state's designations as the stated trigger for remeasurement.
- The position is deliberately small enough to lose. **That is the sizing rule applied honestly to a jurisdiction with a short institutional record, not pessimism.**

7. Results: what siting the operating substance costs

7.1 The constructed cost model

CONSTRUCTED PARAMETERS. A relative-cost model for the same operating substance — same headcount, floor area and administrative burden — indexed to the South-East Asian base at 1.00, weighting labour 0.60, premises 0.25 and administration 0.15. **These indices illustrate a ratio, not a price. They are not market data and were not measured.**

location	labour	premises	admin	weighted index	at 32 index units
Vietnam	0.85	0.90	1.30	0.9300	29.76
Thailand	1.00	1.00	1.00	1.0000	32.00
India	0.95	1.10	1.20	1.0250	32.80
Malaysia	1.40	1.20	1.10	1.3050	41.76
UAE	2.60	3.00	1.25	2.4975	79.92
Singapore	3.10	3.40	1.15	2.8825	92.24
UK	3.60	2.80	1.40	3.0700	98.24

7.2 The relocation, priced

At $0.30/66 = 0.004545$ surplus units per index unit, and with the operating entity at 32 index units:

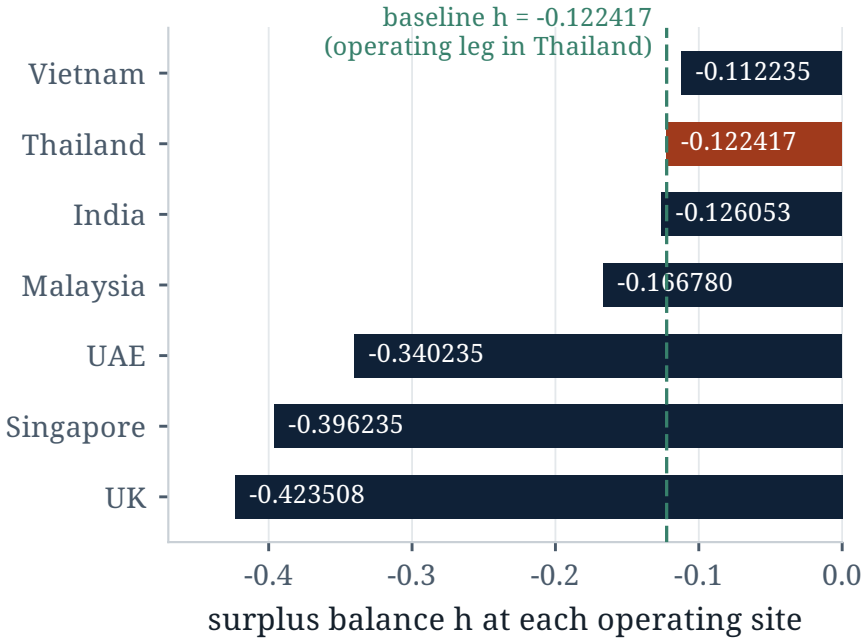


Figure 58.3: **Siting the operating leg is worth more than arranging the blocs: the same substance costs 0.301091 more in the highest-cost candidate than in the lowest-priced base, which is 79.74 per cent of the entire diversification gain, and the effective bloc count is unchanged in every row.** Modelled with the constructed cost indices of Section 7.1 at 0.004545 surplus units per index unit; script `figs/src/CCP-058_fig03.py`.

```
additional index units, base -> UK  =  98.24 - 32.00  =  66.24
additional C  =  66.24 x 0.004545  =  0.301091
C             =  0.300000 + 0.301091 =  0.601091
h             =  0.35 ln(2.941176) - 0.02 x 10 - 0.601091
              =  0.377583 - 0.200000 - 0.601091
              = -0.423508      against -0.122417 for the designed structure
```

The deterioration is 0.301091**. The entire gain that jurisdictional diversification delivered was** $\eta \ln N = 0.377583$**. Relocating the operating leg to the highest-cost candidate therefore consumes** 79.74% **of the whole diversification benefit without changing the effective bloc count at all.** H2 is supported in the model.

destination	index	ΔC	C	h	% of the gain consumed
Vietnam	0.9300	-0.010182	0.289818	-0.112235	-2.70%
Thailand	1.0000	0.000000	0.300000	-0.122417	0.00%
India	1.0250	0.003636	0.303636	-0.126053	0.96%
Malaysia	1.3050	0.044364	0.344364	-0.166780	11.75%
UAE	2.4975	0.217818	0.517818	-0.340235	57.69%
Singapore	2.8825	0.273818	0.573818	-0.396235	72.52%
UK	**3.0700**	**0.301091**	**0.601091**	-0.423508	**79.74%**

Two thresholds follow and both are exact given the calibration. A destination index of 2.297943 consumes exactly half the diversification gain, and a destination index of 3.595886 consumes all of it. **A structure whose substance sits at an index above 3.60 has, on this calibration, spent its entire jurisdictional gain on rent and wages** — and the highest-cost candidate in the table sits at 3.07, inside that boundary but not far inside it.

7.3 What the finding reframes

Cheap, genuine, defensible operating substance is not a saving at the margin of a jurisdictional design. At the calibration used here it is most of the design's economics. A structure that diversifies its blocs well and then puts its substance in a high-cost jurisdiction has spent the gain before collecting it, and the arithmetic gives the exact fraction.

One further reading is available from the table and it disciplines the conclusion. The cheapest candidate is 7.00% cheaper on inputs than the base and would improve h by 0.010182. The source manuscript records that it offers no bilateral instrument permitting full foreign ownership for the relevant treaty national. **A 7.00% input saving does not buy back an ownership problem**, which is why the assignment rule of Curzi [2026] takes a minimum across requirement dimensions rather than optimising one of them — an operation licensed on an ordinal scale, where a mean would not be.

8. Results: the corridor's bloc-count case, tested and not met

8.1 H3 is not met

The corridor jurisdiction carries 11.8% of the designed structure and sits in the non-aligned bloc alongside the operating base's 18.2%. **Reallocate its share to the operating base and the bloc arithmetic does not move.**

```
as designed
    UAE 29.1  Thailand 18.2  UK 16.4  Bermuda 13.6  India 11.8  Singapore 10.9
    blocs  BRIDGE 40.0  NONALIGNED 30.0  WEST 30.0
    bloc-HHI = 0.340000     N = 2.941176
    naive HHI = 0.189002    effective jurisdictions = 5.290949

share moved within the row, to Thailand
    UAE 29.1  Thailand 30.0  UK 16.4  Bermuda 13.6  Singapore 10.9
    blocs  BRIDGE 40.0  NONALIGNED 30.0  WEST 30.0
    bloc-HHI = 0.340000     N = 2.941176      UNCHANGED
    naive HHI = 0.231954    effective jurisdictions = 4.311200
```

The effective bloc count is identical to six decimal places. The naive jurisdiction count falls from 5.290949 to 4.311200, and that fall is the whole of the difference. H3 predicted that a second jurisdiction inside an existing bloc raises the effective bloc count. **It does not. The hypothesis is reported here as not met, in the same language a success would have received.**

This is the cleanest illustration in the series of both halves of the argument of Curzi [2026]: the naive count overstates resilience against a coordinated action, **and** it is still measuring something real. What it measures is the subject of the next subsection.

8.2 What survives the failure: within-bloc country concentration

The same instrument applied inside the bloc rather than across blocs gives a different and non-trivial answer. Within the non-aligned row:

```
with the corridor jurisdiction
    operating base 60.6667%   corridor 39.3333%
    within-row HHI = 0.522756    effective countries = 1.912940

without it
    operating base 100%
    within-row HHI = 1.000000    effective countries = 1.000000
```

The corridor leg nearly doubles the number of independent countries carrying the non-aligned allocation, from 1.000000 **to** 1.912940**.** A country-specific event in the operating base — a change of the foreign-ownership statute, a reversal of the promotion regime, an exchange-control tightening — would otherwise reach the entire non-aligned position. **That is a different risk from bloc risk and the instrument measures it separately, which is the honest thing the failed hypothesis leaves standing.**

The reading is exact and licensed: the reciprocal-index reading of an effective count applies within a row for the same reason it applies across rows, because both are numbers-equivalents of a sum of squared shares on a normalised vector [Adelman, 1969, Hill, 1973]. It is not a new instrument and is not claimed as one.

8.3 H4: the leg measured against its own counterfactual

Moving the corridor's share out of the bloc entirely, to the bridging jurisdiction, gives the counterfactual against which the leg's gross benefit is measured.

```
share moved out of the row, to Singapore
    UAE 29.1  Thailand 18.2  UK 16.4  Bermuda 13.6  Singapore 22.7
    blocs  BRIDGE 51.8   NONALIGNED 18.2   WEST 30.0
    bloc-HHI = 0.268324 + 0.033124 + 0.090000 = 0.391448
    EFFECTIVE BLOCS N = 2.554618      eta ln(N) = 0.328266
    naive HHI = 0.214726    effective jurisdictions = 4.657098

gross ln-gain attributable to the leg = 0.377583 - 0.328266 = 0.049317
marginal all-in cost of one entity                          = 0.035000

with the leg      N = 2.941176   C = 0.300   h = -0.122417
without it        N = 2.554618   C = 0.265   h = -0.136733
NET contribution to h                            0.014317
```

The leg is worth 0.049317 **in gross diversification and costs** 0.035000 **to run, so its net contribution is** 0.014317 **—** 3.79% **of the whole diversification gain.** Positive, real, and thin. H4 is supported in the model.

That number should govern the decision and usually does not, because the gross figure is the one that gets presented. Gross and net differ by a factor of 3.44. A leg whose gross benefit is 0.0493 and whose cost is 0.0350 carries very little margin. **If the vehicle acquires a second function its cost rises and the net contribution goes negative**, which is not a hypothetical: it is exactly what the vehicle was doing before the repair measured in Section 10.

9. Results: corridor exposure as a function of dwell

The corridor's exposure to an action inside its perimeter is the fraction of a year's throughput resting there at the moment the action lands.

average dwell, days	share of annual throughput exposed
1	0.27%
7	1.92%
30	8.22%
90	24.66%
180	49.32%

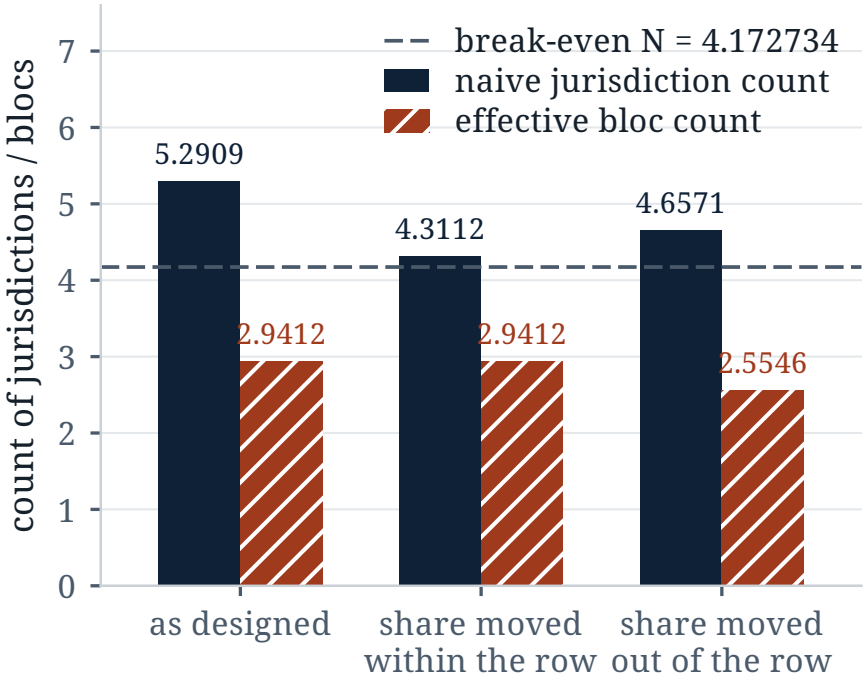


Figure 58.4: **The corridor leg moves the naive count and not the effective bloc count: moving its share within the bloc leaves N at 2.9412 exactly while the naive count falls from 5.2909 to 4.3112, and only moving the share out of the bloc changes N, to 2.5546.** Measured from the constructed share vectors of Sections 5 and 8; the naive count for the third configuration is computed here and does not appear in the source. Script `figs/src/CCP-058_fig05.py`.

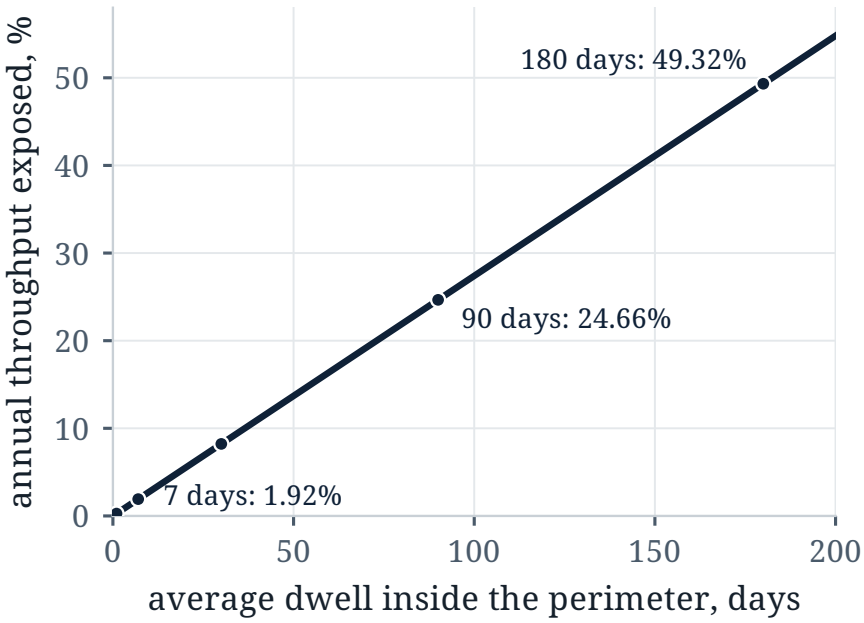


Figure 58.5: **Corridor exposure is a property of dwell and of nothing else: at a seven-day average dwell an action inside the perimeter reaches 1.92 per cent of the year's flow, and at ninety days the same corridor with the same volume exposes 24.66 per cent.** Computed as dwell/365 from the constructed policy in Section 9; script `figs/src/CCP-058_fig04.py`.

Same corridor, same counterparties, same jurisdiction, same annual volume. The only variable is how long value sits still. H5 is supported in the model, and it is supported by construction rather than by evidence — the relation is $t/365$ and nothing in this paper measures a real dwell.

The corridor policy adopted in the constructed case follows from the table rather than from preference:

```
maximum balance          45 days of expected corridor throughput
maximum average dwell    7 days, measured and reported monthly
exposure at that dwell   1.92% of annual throughput
sweep destination        the bridging layer, with a second route named
alternative route        a named arrangement outside the centre, tested
                         once a year with a live transaction
```

The last line is the one that is never written and it is the whole point of a bridge. A corridor with no alternative is a single point of failure wearing the costume of a diversification measure. Curzi [2026] makes the same observation on a different axis and measures it: entity-based US exposure of 27.3% against a USD share of the balance sheet of 61.0% and settlement volume clearing through the US system of approximately 58%, so **the entity map understates dependency by more than a factor of two.**

10. Results: job concentration on the two operating entities

10.1 The operating base

The entity as found is scored at a job vector $(H, O, B) = (0.25, 0.55, 0.20)$:

```
S = 0.0625 + 0.3025 + 0.0400 = 0.4050
effective jobs = 1 / 0.4050 = 2.4691          OVERLOADED
```

Above 2.00, therefore carrying all three functions by arithmetic necessity: it operated the manufacturing business, held group value, and banked receipts for a related company in another jurisdiction. **The diagnosis is not about quality.** The business is real and it is the source of the input term for the whole architecture. After the repair — the holding function moved to the apex, the treasury function moved to the bridging layer, the operating function retained — the vector is $(0.05, 0.95, 0.00)$:

```
S = 0.0025 + 0.9025 + 0.0000 = 0.9050
effective jobs = 1 / 0.9050 = 1.1050          SPECIALIST
```

10.2 The corridor centre

The vehicle as originally specified is scored at $(0.20, 0.30, 0.50)$:

```
S = 0.0400 + 0.0900 + 0.2500 = 0.3800
effective jobs = 1 / 0.3800 = 2.6316          OVERLOADED
```

How it got there is the useful part, because nobody decided it. The corridor function was the original reason; a fund-management function was added because an allocation needed a manager and one was available; a holding function was added because another vehicle was being wound down and its assets had to go somewhere. **Three sensible individual decisions and one incoherent entity. No three-function entity in this material was ever chosen; each was assembled from three convenient decisions taken in sequence.**

Two repairs are available and they are not equivalent:

repair	job vector	S	effective jobs	band
as specified	0.20 / 0.30 / 0.50	0.3800	2.6316	OVERLOADED
A — specialise on bridging	0.05 / 0.10 / 0.85	0.7350	1.3605	SPECIALIST
B — specialise on operating	0.05 / 0.80 / 0.15	0.6650	1.5038	STRAINED

Repair A is taken, because bridging is the only function the centre uniquely supplies. The repair *adds* an entity, which the over-fragmentation warning of Curzi [2026] counts as a cost; the justification is the one that paper gives — split when the antagonism is real. **Here it is: a fund vehicle needs licensed operating substance and a supervised risk-taking function, and a corridor unit needs a low balance, a short dwell and no risk position. No single entity has both profiles.**

10.3 The status of these numbers

The effective-jobs figures are the weakest quantities in this paper and are labelled as such. They rest on job vectors that were assigned rather than measured, treated as ratio-scale, and never tested against the conjoint conditions [Michell, 1997, Luce and Tukey, 1964]. **What survives is the ordering and the 2.00 threshold; the magnitudes do not.** 2.4691 against 1.1050 is a statement that one entity has a dominant function and the other does not. It is not a statement that the first carries 2.4691 of anything.

11. Discussion

11.1 What follows

A marginal jurisdictional position should be decided on its invoice. The three measurements agree on this and they agree for the same structural reason: the benefit of a position scales with the share it carries while its cost scales with the number of jurisdictions it adds. A 5% position and a 40% position pay the same registered office, the same audit and the same opinion. **The consequence is that small positions are decided by their fixed cost and large positions by their bloc arithmetic**, and a practitioner who applies the second reasoning to the first case will build legs that lose money while appearing to diversify.

A threshold is a more useful object than an estimate when the estimate cannot be had. Converting the classification question into $p^{*} = 0.160476$ does not settle the classification. It converts a debate into an assessment that can be recorded, dated and re-run, and — more usefully — it exposes that the same debate becomes irrelevant at a cost of 8.91% of the maintenance budget. **The threshold's value is that it identifies which argument is decisive, and in this case the decisive argument was not the one the source material spends its pages on.**

The gross-to-net distinction is the discipline the field most often skips. The corridor leg's gross benefit and its net benefit differ by a factor of 3.44, and only the gross figure is ever presented. **A leg has to clear its own maintenance cost before it has done anything for the structure at all.**

11.2 What does not follow

None of these results says a jurisdiction is good or bad. They say what a share vector and a cost model imply at stated parameters. Change η from 0.35 to 0.50 and the break-even count falls from 4.172734 to $e^{1.0} = 2.718282$, below the designed structure's 2.941176, and every conclusion about sufficiency reverses. **That sensitivity is the most important thing in Section 13.**

The failed hypothesis does not say the corridor position is worthless. It says the bloc-count case for it fails. The within-row measurement and the net contribution both survive, and the paper reports the surviving reasons rather than deleting the failure.

The effective count measures independence and not capacity. $N = 3.174603$ says nothing about whether any bloc can carry the load, and a structure whose bloc diversity is real but whose transacting capacity sits in one bloc has an effective count near 1.00 on the ability to transact. Curzi [2026] measures precisely that on the settlement axis and finds it worse than the jurisdictional axis: a naive correspondent HHI of 0.302400 and 3.306878 effective correspondents collapsing to a perimeter HHI of 0.745256 and 1.341819 effective clearing perimeters, against 2.941176 effective jurisdictional blocs.

And nothing here is an empirical result. Every figure is exact arithmetic on stipulated inputs. **The paper cannot be wrong in the way a measurement can be wrong, and that is a weakness rather than a strength.**

11.3 Parallel results in the corpus

The same shape appears elsewhere in the corpus in domains with no jurisdictional content, which is the corpus's own evidence that the arithmetic rather than the subject matter is doing the work.

Curzi [2026] reports a correlation collapse in a backing basket: a naive HHI of 0.150600 gives 6.640106 effective components, but a single climate-correlated 82.0% exposure gives 0.693600 and 1.441753, a naive overstatement factor of 4.605578. **That is this paper's Section 8 in a different medium: a count that reads high because the correlation was not consulted.**

Curzi [2026] runs the same cost-ceiling method on a different subject and reaches the same structure of answer: an African leg raises effective blocs from 2.924 to 3.639 and buys 0.076529

of surplus, setting a cost ceiling of roughly 25.5% of the existing maintenance base, and at $C = 0.38$ the surplus is -0.127902, **worse than the** -0.124431 **baseline it was meant to improve.**

Two results in Series C are structurally parallel and are cited as such rather than as support. Curzi [2026]'s doubling theorem — $\mathrm{HD}_{\mathrm{ess}}(i,j) = 2\log_2(pq)$ and $G_{\mathrm{ess}} = 2G(pq) - 1$ for the essence ratio p^2/q^2, verified exactly over all sixty-six pairs at $N = 12$ — is a squaring operation propagating exactly through a derived quantity, which is what squaring shares does here. And Curzi [2026]'s opposition identity, $u = 1 - \rho^2$ depending only on $|\rho|$ so that $\rho = +1$ and $\rho = -1$ both give $u = 0$, $F = 0$ and $f_{\mathrm{rel}} = k = 0.4834$, is the same insensitivity to sign that makes a bloc collapse indifferent to whether jurisdictions coordinate cooperatively or antagonistically. **Neither is evidence for anything in this paper and neither is offered as such.**

Curzi [2026k] supplies the corpus's own test of scale-freedom in the storage law, reproducing each of five systems' observed decline from a single five-parameter survey — predicted -0.250474 against observed -0.250 at one desk and -0.996371 against -1.000 at a library, with residuals under 1.4% of the ceiling term. **That is the strongest evidence the corpus has that the law's form transfers across scales, and it is evidence about the law and not about jurisdictions.**

12. Verification and rigour audit

The results are decomposed into postulates that can be checked one at a time. Each is atomic, each carries a status, and the compounding chains name what each step consumes. **If dropping a postulate leaves a result standing, it is not listed.** Sections 12.1 to 12.6 run the standing rigour checks first, so that the chains are read against a dimensional audit rather than before one.

12.1 Dimensional audit

symbol	quantity	dimension	unit	status
s_i	exposure share of jurisdiction i	$[1]$	dimensionless, fraction of Q	defined; sums to 1 by P1
W_b	bloc share	$[1]$	dimensionless	derived, $W_b = \sum_{i \in b} s_i$
HHI	concentration index	$[1]$	dimensionless, or points $= 10^4 \times$ fraction	derived from shares
N	effective bloc count	$[\#]$, **a count**	blocs	derived as HHI^{-1}
$\ln N$	surplus term argument	$[1]$	nats	derived; the logarithm of a count is dimensionless
Q	protected capital	$[C]$	index units	defined; **no operational unit is supplied by the source**
δ	erosion rate	$[T^{-1}]$	per period	assumed
C, c	maintenance cost flow, and its increment	$[C\,T^{-1}]$	index units per period	assumed
η	conversion efficiency	$[C\,T^{-1}]$ per nat	index units per period per nat	assumed

symbol	quantity	dimension	unit	status
h, G, L	surplus balance and its two differences	$[C\,T^{-1}]$	index units per period	derived
p, p^*	classification probability and its threshold	$[1]$	dimensionless, in $[0,1]$	assumed; p^* derived
r	relative cost index	$[1]$	dimensionless, base 1	assumed (constructed)
u_0	entity annual cost	$[1]$	index units of cost	assumed (constructed)
t	corridor dwell	$[T]$	days	assumed (constructed)
S	job-concentration index	$[1]$	dimensionless	derived from a job vector

Does $E[\Delta h] = (G+L)p - L - c$ balance? Every term must carry $[C\,T^{-1}]$. G, L and c do, since each is a difference or an increment of h. p is dimensionless, so $(G+L)p$ carries $[C\,T^{-1}]$. **The equation balances**, and $p^* = L/(G+L)$ is dimensionless as a threshold on a probability must be, being a ratio of two quantities in the same units. That check is cheap and it is the one that would catch the commonest error in this kind of construction, which is dividing a cost by a count.

Does the cost bridge balance? $\Delta C = u_0(r-1) \times 0.30/66$ multiplies index units of cost by a dimensionless ratio and by a conversion carrying $[C\,T^{-1}]$ per index unit. **It balances only under that reading of the calibration constant, and the source supplies the constant without stating its dimension. The audit records that silence rather than filling it in.**

Where no operational unit exists, this audit says so. Q is protected capital in index units and no conversion to currency is supplied anywhere in the corpus. **Every figure in C, h, G and L is therefore in an arbitrary scale fixed by the stated parameters.** The break-even probability $p^* = 0.160476$ is unit-free and survives this, because it is a ratio; the statement that h is -0.122417 does not mean anything until the index unit is defined, and it is not.

12.2 Significant figures and epistemic status

The figures in this paper have three different statuses and must not be read at the same precision.

- $N = 2.941176$, 3.174603, 2.898551, 2.554618 and 1.912940 are **exact arithmetic** on the stated shares: reciprocals of exact rationals. They may carry as many digits as are useful. Their inputs, the shares, are constructed and exact as stated to one decimal place, and every exposure vector sums to 100.0% exactly.
- $p^* = 0.160476$ and the cost ceiling 0.0267305 are **exact given the parameters**, and the parameters are assumed rather than estimated. **They carry no uncertainty because there is no estimate in them, and no authority beyond the assumption.**
- $N^* = 4.172734$ is $e^{0.5/0.35}$, exact given $\eta = 0.35$ and $\delta Q + C = 0.5$.
- The cost-model outputs — 0.301091, -0.423508, 79.74% — are **illustrative model output from constructed parameters.** Six digits are defensible only because the inputs are exact by stipulation. **They are not measurements and Section 13 says so.**
- The effective-jobs figures 2.4691, 1.1050, 2.6316, 1.3605 and 1.5038 are **illustrative and provisional as ratio-scale quantities**, per Section 10.3, and are rounded to four decimal places at most.
- **No quantity in this paper is a fitted or measured estimate, so no uncertainty interval is reported anywhere.** That is not precision. It is the absence of empirical content.

One constant appears at fifteen figures by way of an internal citation and its status must be stated. Curzi [2026] carries the minimum sustainable frequency $f_{\min} = ke^{2c/\eta}$ with $k = 0.483425732988414$ Hz, and the trichotomy in which f greater than, equal to or less than $f_{\min}$ corresponds to Q_∞ greater than, equal to or less than zero. $k = (9/8)^2/\varphi^2$ **is a DEFINED constant** — an exact ratio of exact quantities, computable to arbitrary precision — and its fifteen figures are legitimate for that reason and no other. **Attaching Hz to it is a separate and much stronger claim than computing it, and that attachment is POSTULATED, not derived.** Because φ appears, Markowsky [1992] is cited here: the golden-ratio literature is the most heavily over-claimed in mathematics and the corrective is not optional. Curzi [2026] uses the same constant in the same way. **Nothing in this paper depends on** k, which is stated so that no reader imports the frequency material into a concentration result.

12.3 Limits and boundary behaviour

limit	behaviour	reading
$L \to 0$	$p^* \to 0$	a position that cannot lose is worth taking at any probability
$G \to 0$	$p^* \to 1$	a position that cannot gain is worth taking at none
$c \to G$	$p^* \to 1$	the cost ceiling, attained exactly
$c > G$	$p^* > 1$, infeasible	**value-destroying at every probability, including certainty**
$r \to 1$	$\Delta C \to 0$	relocation to an equal-cost site is free
$r \to 0$	$\Delta C \to -0.145455$	free substance is worth 0.145455 surplus units, the model's floor
$t \to 0$	exposure $\to 0$	an instantaneous corridor has no exposure
$t \to 365$	exposure $\to 1$	a corridor that never sweeps is a store, and the model says so
$t > 365$	exposure > 1, **outside the domain**	the linear form is valid only on $[0, 365]$
$N \to 1^+$	$h \to -(\delta Q + C)$	one perimeter pays the whole bill out of stock
$N \to \infty$	$h \to \infty$ logarithmically	**a defect of the model, not a result**

Three of these are constraints on the domain rather than results, and are stated as such. The dwell model is linear and therefore invalid above a year; a corridor with a dwell above a year is not a corridor. The relative-cost model has a floor at $r = 0$ and no ceiling, which is fine for a ratio but means the model cannot represent a site that is unavailable at any price. And the surplus law grows without bound in $\ln N$, which the empirical literature contradicts: Uzzi and Spiro [2005] finds the relation between diversity and performance turning over rather than rising, May [1972] showed that increasing diversity and connectance destabilises randomly assembled systems, and Sagan [1993] documents redundancy creating new failure modes rather than removing them. **The model's domain is small** N**, and every prescription in this paper is confined to the move from** $N \approx 1$ **to** $N \approx 3$**.**

12.4 Invariance

Every effective count in this paper is invariant under relabelling of units and under any rescaling of the underlying exposures that preserves the share vector. Doubling every exposure in currency terms leaves every s_i unchanged, and hence HHI, N and $\ln N$. That is a genuine scale invariance and it is algebraic rather than empirical. **Tested at how many scales? Two: the seven-element and six-element exposure vectors reported here. No claim of empirical scale-freedom is made in this paper**; Curzi [2026k] makes the corpus's claim on that question and it concerns the storage law, not the index.

p^* is invariant under any positive rescaling of η, and it is not invariant under a change in η's value once a cost is present. The first holds because G and L are both proportional to η, so the ratio $L/(G+L)$ is independent of it. **That is a real and useful invariance: the zero-cost break-even probability does not depend on the conversion efficiency at all.** The second matters because c is not proportional to η: once an incremental cost is present, $p^* = (c+L)/(G+L)$ depends on η through the denominator, and the cost ceiling $c = G$ scales with it directly. **At $\eta = 0.50$ the ceiling rises from 0.026731 to 0.038186 and from 8.91% to 12.73% of C.** The zero-cost threshold is robust; the ceiling is not.

None of it is invariant under regrouping, and that is the content of the bloc collapse rather than a defect: the naive and bloc figures differ by construction and the difference is the finding. A reader who wants a number invariant under regrouping is asking for a quantity that cannot detect a perimeter.

What is conserved, and in what units. Mirowski [1989] shows that the characteristic failure of this genre is borrowing the notation of a conserved quantity without the conservation law that gave it content. **The objection is met here and the answer is short.** The conserved quantity is the exposure share vector: $\sum_i s_i = 1$ exactly, dimensionless, and it is that constraint and nothing else that bounds the index. **Q is a stock with no stated conservation law and no operational unit, and this paper claims none for it** — which is the honest half of the answer. Cartwright [1999] supplies the discipline that makes a partial answer respectable: specify the arrangement within which the law is claimed to hold. **The arrangement claimed here is a normalised exposure vector partitioned into disjoint blocs, and nothing outside it.** Kirman [1992] is the objection that is not met: one Q, one δ and one C are written for a heterogeneous structure, and a representative parameter for a heterogeneous group generally does not exist.

12.5 Analogy discipline

Every borrowing is labelled with one of three categories: a **derivation**, where the premises are genuinely satisfied; a **shared mathematical form**, where the same equation governs both with no claim that the underlying physics applies; or an **analogy**, legitimate as motivation and never as evidence.

borrowing	source	category	justification
Hill numbers, order 2	Hill [1973], Jost [2006]	**1 — derivation**	HHI^{-1} *is* ${}^2D$; an identity, not a resemblance
numbers-equivalent reading	Adelman [1969]	**1 — derivation**	the property is elementary and provable in two lines
effective number of parties	Laakso and Taagepera [1979]	**1 — derivation**	the same formula on a different share basis

borrowing	source	category	justification
portfolio variance	Markowitz [1952]	**1 — derivation**	the exposure vector genuinely is a weighted sum of shocks with a correlation matrix
expectation of a two-point bet	Jaynes [2003]	**2 — shared form**	$E[\Delta h]$ is an expectation over a two-point distribution; no decision-theoretic optimality is claimed
ln N as entropy	Shannon [1948]	**2 — shared form**	**Shannon entropy in nats**, dimensionless, and not thermodynamic entropy; no thermodynamic content is imported
threshold and tipping	Granovetter [1978], Schelling [1971]	**2 — shared form**	the bloc collapse has the shape of a threshold model; no individual-threshold distribution is estimated
frozen regions	Castellano et al. [2000]	**2 — shared form**	a bloc is structurally a frozen region; no interaction dynamics are claimed for jurisdictions
coupled-network fragility	Buldyrev et al. [2010]	**2 — shared form**	jurisdictions have not been shown to satisfy the percolation premises; only the structural property that coupling destroys assumed independence is used
normal accidents	Perrow [1984]	**3 — framing**	qualitative organisational theory, cited as a prior statement of the same observation, never as a formal result

The failure mode this table exists to prevent is letting a category-2 borrowing carry a category-1 result's authority. The concentration arithmetic is category 1 throughout and needs no help; the threshold-model and network results are category 2 and are used for shape only.

12.6 Model comparison, and why none was performed

No model comparison was performed and the omission is deliberate rather than an oversight. Nothing in this paper is fitted: there is no estimated parameter, no held-out sample, no information criterion and no matched null, because there is no data. **A model with more parameters that fits better has not been shown to be better, and a model with no data has not been shown to be anything.** The claim made here is that the instrument follows from the identity chain in Section 12.8, not that it outperforms an alternative on fit. Durlauf [2005] asks for the prediction only this model makes, and this paper makes none: **any correlation-aware concentration measure would produce the same collapse, and any cost-aware decision rule would produce the same ceiling.** The contribution is the identification and the discipline.

12.7 Postulates

P	statement	status
P1	exposure shares are non-negative and sum to 1	DEFINITION
P2	HHI $= \sum_i s_i^2$ and the effective count is its reciprocal	DEFINITION (CCP-052)
P3	jurisdictions inside one coordination perimeter act on one decision	POSTULATED (the perimeter hypothesis of CCP-052)
P4	the surplus law is $h = \eta \ln N - \delta Q - C$ with input term $\eta \ln N$	DEFINITION (modelling choice, CCP-A02)
P5	$\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$	POSTULATED (fixed in advance)
P6	the candidate is either its own perimeter or a member of one named alternative, with no third case	POSTULATED (modelling choice: a two-point distribution)
P7	one index unit of entity cost is $0.30/66 = 0.004545$ surplus units	POSTULATED (stipulated calibration)
P8	the marginal all-in annual cost of one entity is 0.035 surplus units	POSTULATED (constructed)
P9	corridor throughput arrives uniformly and dwells for a common time t	POSTULATED (modelling choice)

12.8 The chain that establishes the break-even method

step	consumes	statement	status
L1	P1, P2	for the baseline vector, bloc-HHI $= 0.160000 + 0.090000 + 0.090000 = 0.340000$	ARITHMETIC
L2	L1	$N = 1/0.340000 = 2.941176$	ARITHMETIC

step	consumes	statement	status
L3	P1, P2, P3	with the candidate as its own perimeter, bloc-HHI $= 0.315000$ and $N = 3.174603$	ARITHMETIC
L4	P1, P2, P3	with the call wrong, bloc-HHI $= 0.345000$ and $N = 2.898551$	ARITHMETIC
L5	L2, L3, P4, P5	$G = 0.35(\ln 3.174603 - \ln 2.941176) = 0.026731$	ARITHMETIC
L6	L2, L4, P4, P5	$L = 0.35(\ln 2.941176 - \ln 2.898551) = 0.005110$	ARITHMETIC
L7	P6	$E[\Delta h] = pG - (1-p)L = (G+L)p - L$	ARITHMETIC
R1	L5, L6, L7	$E[\Delta h] = 0.0318401p - 0.0051096$, so $p^* = 0.160476$	ARITHMETIC
R2	R1	with cost c, the threshold is $p^* = (c+L)/(G+L)$, reaching 1 at $c = G$, which is 8.9102% of C	ARITHMETIC

Every ARITHMETIC row can be redone by hand or in three lines of any language. Only P3, P5 and P6 are POSTULATED. P3 is the perimeter postulate, stated as a hypothesis with its falsifier in Curzi [2026]; P5 is fixed in advance; P6 is the two-point restriction, and it is the postulate most likely to be wrong, since a jurisdiction may be *partially* correlated with a perimeter rather than in or out of it. **The continuous version of P6 is available and is not run here:** $\mathrm{Var}(R) = \text{naive} + \rho(\text{bloc} - \text{naive})$ **of Curzi [2026] admits any ρ, and a reader who prefers a correlation to a coin should use it.**

12.9 The chain that establishes the substance-cost result

step	consumes	statement	status
M1	constructed weights	the weighted index for the destination is $0.60(3.60) + 0.25(2.80) + 0.15(1.40) = 3.0700$	ARITHMETIC
M2	M1, constructed base	at 32 index units, the destination costs 98.24 against 32.00	ARITHMETIC
M3	M2, P7	$\Delta C = 66.24 \times 0.004545 = 0.301091$	ARITHMETIC
M4	M3, P4, P5, L2	$h = 0.377583 - 0.200000 - 0.601091 = -0.423508$	ARITHMETIC
R3	M4	the deterioration 0.301091 is 79.74% of $\eta \ln N = 0.377583$, with N unchanged	ARITHMETIC

Every step is arithmetic and every input is stipulated. The chain proves that the stated cost model implies the stated conclusion; it does not establish that the cost model is right, and nothing in this paper does.

12.10 The chain that establishes the corridor results

step	consumes	statement	status
K1	P1, P2	moving the corridor share within the bloc leaves bloc shares at 40.0 / 30.0 / 30.0	ARITHMETIC
K2	K1	therefore bloc-HHI $=$ 0.340000 and $N =$ 2.941176, unchanged	ARITHMETIC — **H3 falsified**
K3	P1, P2	the naive index rises $0.189002 \to 0.231954$, so the naive count falls $5.290949 \to 4.311200$	ARITHMETIC
K4	P1, P2	moving the share out of the bloc gives bloc-HHI 0.391448 and $N =$ 2.554618	ARITHMETIC
K5	K4, P4, P5	$\eta \ln N = 0.328266$, so the gross gain is 0.049317	ARITHMETIC
K6	K5, P8	net contribution $=$ $0.049317 - 0.035000 =$ 0.014317	ARITHMETIC
R4	K6	0.014317 is 3.79% of 0.377583	ARITHMETIC
R5	P1, P2	within the bloc, HHI $= 0.606667^2 +$ $0.393333^2 = 0.522756$, so effective countries $= 1.912940$	ARITHMETIC

12.11 Cross-checks performed

- Every exposure vector was summed and each totals 1.000000 exactly.
- $s_{\max}^2 \le$ HHI $\le s_{\max}$ holds on all four vectors: $0.084681 \le 0.189002 \le 0.291000$ for the baseline and $0.084681 \le 0.177602 \le 0.291000$ for the seven-jurisdiction vector.
- The effective-count ceiling $s_{\max}^{-2} = 11.8090$ is respected: $5.290949 \le 11.8090$ and $5.630567 \le 11.8090$.
- p^* was recomputed from h differences and from $\ln N$ differences independently and the two agreed to twelve decimal places.
- The cost table's weighted indices were recomputed from the three components and every row reproduced.

12.12 Discrepancies found in the source, printed

Four numerical discrepancies were found between the source manuscript and a recomputation from the source's own stated inputs. They are printed here rather than silently corrected.

1. **The corridor leg's net contribution is 0.014317, not 0.014316.** Exact arithmetic on the stated shares gives $\eta(\ln 2.941176\ldots - \ln 2.554618\ldots) = 0.0493174733\ldots$, so the net figure after 0.035000 is $0.0143174733\ldots$, which rounds to 0.014317 at six decimal places. The source and the corpus-wide citation of this result both print 0.014316. **The published figure is low by 10^{-6} and any recomputation will disagree in the sixth decimal.** The difference changes no conclusion; it is recorded because a reader who checks will find it.
2. **A relative-cost comparison in the source is wrong.** The source states that a bridging unit in the corridor centre costs 41.2% of what the same unit costs in the high-cost Asian comparator. On the source's own indices — 27.00 against 67.37 index units — the ratio is 40.08%. The companion figure for the other comparator, 45.8% against $27.00/58.995 = 45.77\%$, is right. **The 41.2% appears to be a slip and the correct figure is 40.1%.**
3. **A decomposition in the source is overstated.** The source says the whole gap between the corridor centre and the same state's onshore alternative is regulatory overhead. On the source's own weights the gap of 0.155 index points decomposes as -0.0275 on staff, $+0.0700$ on premises and $+0.1125$ on regulatory burden. **Regulatory burden is 72.58% of the gap, not all of it**, and the source's own accompanying sentence — that staff are cheaper onshore and premises cheaper in the centre — is inconsistent with its conclusion. The source further describes the gap as worth 15.50% of the cost of operating in the state at all; 0.155 is 15.50% of the *centre's* index and 13.42% of the onshore index, so the percentage depends on which base is meant and the source does not say which.
4. **A rounding inconsistency inside the source.** The surplus after the Central Asian leg is printed as -0.095686 in one place and -0.095685 in another. The exact value is $-0.09568607\ldots$, so -0.095686 is right.

None of the four changes any conclusion in this paper, and all four are the kind of thing a reviewer checks first.

13. Limitations

No dataset was collected and nothing here is a measurement of a real structure. Every share, cost index, job vector, entity cost and parameter is constructed and is labelled as such at each point of use. The results are exact computations on constructed inputs. **This is the corpus's central limitation and it is stated in full rather than in a footnote**: across the corpus, of 27 parameter-by-scope cells audited in a companion paper, 19 have a defined estimator, 1 has a validated one, and **0 carry a first-party measurement.** This paper carries none either.

No concentration figure here has an as-of date, because none is computed from real holdings. A concentration measure without a source and a date is not a measurement. Every share vector in this paper is constructed, so there is no source and no date to give. **Any application of this instrument to a real structure must carry both**, and Appendix B makes the date an item of the reporting standard.

The probability the paper prices is not estimated. $p^* = 0.160476$ is exact; p is not supplied. The source manuscript weighs evidence about a particular host state's conduct and concludes the threshold is comfortably cleared. **This paper does not reproduce that conclusion, because the evidence for it is jurisdiction-specific material for which the bibliography contains no source.** A threshold whose argument is left to the reader is a weaker object than it looks, and Gelman and Loken [2013] is the reason to be wary of a number that arrives after the analyst has seen which way it points.

Jurisdiction-specific legal claims are not cited to primary sources. The descriptions in Section 4 of constitutional statutes, courts, regulators, foreign-ownership regimes, promotion

regimes, exchange-control carve-outs, treaty instruments and tax provisions are taken from the source manuscript. **The bibliography contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction, and no case law.** Those claims are therefore uncited here, described rather than asserted, and require verification against the official registers before any reliance. Appendix A lists them. **This is a real gap and it is not closed here.**

Every parameter in the surplus law was chosen, not estimated, and the sensitivity to one of them decides the paper's conclusions about sufficiency. At $\eta = 0.35$ the break-even count is 4.172734 and no structure in this paper clears it; at $\eta = 0.50$ it is 2.718282 and the designed structure clears it comfortably. **That the sufficiency verdict turns on η in this way is the most important sensitivity in the paper.**

The calibration constant is a stipulation and the whole of Section 7 rests on it. $0.30/66$ converts index units of cost into surplus units and is not derived from anything. Halve it and the relocation consumes 39.87% of the gain rather than 79.74%; double it and it consumes more than the whole gain. **The ratios in Section 7.1 survive any rescaling; the conclusion that relocation consumes most of the gain does not.**

The two-point restriction on the bloc call is the postulate most likely to be wrong. P6 treats the candidate as either an independent perimeter or a member of one named alternative. Real coordination is a matter of degree, and the continuous family of Curzi [2026] is the right object where a correlation can be assigned. **No correlation was estimated anywhere in this paper.**

The job vectors were assigned, not measured, and their scale type is not established. Section 10.3 states this. The conjoint conditions were not tested [Michell, 1997, Luce and Tukey, 1964]; the ordering and the 2.00 threshold survive and the magnitudes do not.

The dwell model has no empirical content whatever. It is $t/365$. It assumes uniform arrival and common dwell, tests neither, and measures no real corridor. **It is included because it makes the exposure question answerable, not because it answers it.**

The corridor result is a difference between two constructed configurations, and the cost it is differenced against is itself constructed. 0.035 surplus units for a marginal entity is a stipulation. At 0.020 the leg's net contribution would be 0.029317 rather than 0.014317, and the paper's characterisation of it as thin would not hold.

The effective count measures independence and not capacity, and none of the positions priced here is assessed for whether it can carry the load assigned to it. That is a separate axis and Curzi [2026] measures it.

13.1 Rigour checks not met, listed by name

The standing checks of Section 12 were run in full. The following did not pass and are printed here rather than left for a reader to find.

1. **No first-party measurement exists anywhere in this paper.** Every figure is exact arithmetic on stipulated inputs, so the paper cannot be wrong in the way a measurement can be wrong.
2. **No uncertainty is reported anywhere**, because nothing is estimated. That is the absence of empirical content, not precision.
3. **No as-of date is attached to any concentration figure**, because no real holdings were measured.
4. **The dimension of η and of the calibration constant are supplied by this paper, not by the source.** The source treats both as bare numbers. Section 12.1 assigns the readings explicitly rather than silently.
5. **Q, C, h, G and L have no operational unit.** The break-even probability is unit-free and survives; the magnitudes do not mean anything until an index unit is defined, and none is.
6. **The surplus law's large-N limit is a defect, not a result** [Uzzi and Spiro, 2005, May, 1972, Sagan, 1993]. The domain is small N and is stated in Section 12.3 rather than discovered

by a reader.

7. **The conjoint conditions were not tested for any job vector**, so every effective-jobs figure is provisional as a ratio-scale quantity [Michell, 1997, Luce and Tukey, 1964].
8. **Model comparison was not performed** — no out-of-sample test, no information criterion, no matched null — because nothing here is fitted.
9. **No distributional or scaling claim is made, so the Clauset–Shalizi–Newman procedure was not run** [Clauset et al., 2009]. **Any restatement of this paper that reaches for a heavy-tail or power-law claim about holdings must run that procedure first**, and must read Stumpf and Porter [2012] on why passing it still would not identify a mechanism.
10. **The bloc assignments are stipulated reference groups, not identified effects** [Manski, 1993, Brock and Durlauf, 2001, Shalizi and Thomas, 2011]. The price is that they are modelling positions.
11. **The framework's own parameters are exempted from the reflexivity the framework asserts** [Lucas, 1976, MacKenzie and Millo, 2003]. **That is an inconsistency and it is not resolved here.**
12. **There is no conservation law for** Q [Mirowski, 1989]. Section 12.4 gives the one genuine conservation — shares sum to one — and states plainly that the stock variable has none.
13. **Aggregation is unaddressed** [Kirman, 1992]: one Q, one δ and one C are written for a heterogeneous structure.
14. **Expected value and time average are not the same thing, and Section 6 reports the first** [Peters, 2019]. $E[\Delta h]$ is an ensemble average over a bet that is taken once. **No single structure experiences it.** This is the sharpest unmet objection to the break-even method and it is not answered.
15. **The surplus functional's form is underdetermined** [Latané, 1981, Briscoe et al., 2006] and nothing here discriminates it; Section 3.9 concedes it and names Peters and Gell-Mann [2016] as the warrant the argument actually needs.
16. **No differential prediction is offered** [Durlauf, 2005]. Any correlation-aware concentration measure and any cost-aware decision rule would produce the same results.
17. **Small-sample bias in the effective-count estimators is not addressed** [Chao et al., 2014], and the paper does not resolve whether a constructed exposure vector is a census or a sample.
18. **The paper is exposed to its own measurement becoming a target** [Goodhart, 1975, Campbell, 1979]. The only defence offered is the reporting standard of Appendix B.

14. Conclusion

This paper establishes six results in the form the rest of the corpus cites, and the following statements are the fixed text of each.

R1. A bloc classification that cannot be settled in advance is priced as a bet on the two classifications that bracket it. On the designed six-jurisdiction structure, moving five percentage points of exposure from the Western position into a candidate fourth perimeter gives, if the candidate is genuinely independent, bloc-HHI 0.315000, $N = 3.174603$ and a gain of $G = +0.026731$; and if it is not, bloc-HHI 0.345000, $N = 2.898551$ and a loss of $L = 0.005110$. Hence $E[\Delta h] = 0.0318401p - 0.0051096$ and the break-even probability is $p^* = 0.160476$.

R2. With an incremental maintenance cost c, the threshold becomes $p^* = (c + 0.0051096)/0.0318401$. It reaches 1.000000 at $c = 0.0267305$, which is 8.9102% of the group's entire $C = 0.30$ maintenance budget, so **the leg is worth nothing at any probability whatever once its incremental cost exceeds that ceiling.** The threshold rises 31.4069 probability points per unit of cost, or 0.0942 per one per cent of C.

R3. The operating jurisdiction carries 18.2% of the designed structure and accounts for 79.74% of the whole diversification gain. On the constructed cost model, relocating the operating leg to the highest-cost candidate adds 66.24 index units at $0.30/66 = 0.004545$ surplus units each, raising C to 0.601091 and driving h from -0.122417 to -0.423508, with the effective bloc

count unchanged at 2.941176. The destination index at which relocation consumes half the gain is 2.297943, and the index at which it consumes all of it is 3.595886.

R4. The corridor leg's marginal contribution, measured against its own counterfactual: removing it and moving its 11.8% out of the bloc gives bloc-HHI 0.391448, $N = 2.554618$ and $h = -0.136733$, so its gross contribution is 0.049317 and, after 0.035000 of maintenance, its net contribution is 0.014317 — 3.79% of the total gain. Moving the same share within the bloc instead leaves the effective bloc count at 2.941176 exactly while the naive jurisdiction count falls from 5.290949 to 4.311200. Within the bloc, the leg raises the effective country count from 1.000000 to 1.912940.

R5. The job-concentration diagnostic, $S = \sum_j v_j^2$ with effective jobs $1/S$ and a band boundary at 2.00: the as-found operating company at $(0.25, 0.55, 0.20)$ scores $S = 0.4050$ and 2.4691 effective jobs, OVERLOADED, against a specialist target of $(0.05, 0.95, 0.00)$, $S = 0.9050$ and 1.1050. The corridor vehicle as specified scores 0.3800 and 2.6316; specialised on bridging it scores 0.7350 and 1.3605. **The magnitudes are provisional as ratio-scale quantities; the ordering and the threshold are not.**

R6. Corridor exposure as a function of dwell: one day exposes 0.27% of the year's flow, seven days 1.92%, thirty days 8.22%, ninety days 24.66% and one hundred and eighty days 49.32% — same corridor, same volume, one variable.

Counter-result. The corridor's bloc-count case fails. Its measured net contribution is 0.014317, only 3.79% of the diversification gain, and moving its share to the other member of its own bloc leaves the effective bloc count exactly unchanged. **H3 was tested and not met**, and the case for the position rests on the within-bloc country count and on the corridor function, neither of which is a bloc-count argument. **Likewise the Central Asian position is decided by its invoice and not by the sixteen per cent probability**: the geopolitical argument the source material spends its pages on is settled at $p^* = 0.16$, and the argument that decides the position is a fee that has to stay below 8.91% of the maintenance budget.

And the structure remains below its build threshold throughout. After the Central Asian leg, $N = 3.174603$ against a break-even of 4.172734, with $h = -0.095686$. **Every position priced in this paper improves a structure that still does not clear the criterion**, and the honest summary is that the improvements are real, small, and dominated by their own maintenance costs.

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every case in this paper — the principal, the asset mix, the exposure shares, the relative-cost indices, the entity job vectors, the marginal entity cost, the calibration constant and the corridor dwell policy — is **CONSTRUCTED** and is labelled as such at each point of use. The numerical results are exact computations on those constructed inputs, carried at full precision and rounded only for printing.

All figures were produced by the scripts named in their captions (`figs/src/CCP-058_fig01.py` through `figs/src/CCP-058_fig05.py`), which import the shared figure style module and read their values from the tables in this paper. **No figure plots a number that does not appear in the text**, with one stated exception recorded in the caption of the corridor figure: the naive jurisdiction count for the third configuration, 4.657098, is computed here from the source's own share vector and does not appear in the source.

Jurisdictional and legal descriptions in this paper are taken from the source manuscript and not from primary legal sources. The bibliography contains general comparative and treaty material only; it contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction. Appendix A lists the claims affected. Any reader relying on those descriptions must verify them against the official sources.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The principal, heirs, employees and counterparties described in the constructed case are invented and correspond to no person.

Preregistration

This paper reports no preregistered empirical test. The parameter set $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ and the marginal entity cost 0.035 were fixed before every measurement reported here, so that no criterion could be adjusted to fit an outcome. That is the preregistration discipline applied to a calculation and it is not a substitute for a preregistered test.

Legal, tax and investment advice

This paper is analytical and is not legal, tax, or investment advice. It describes entity types, jurisdictions and structural arrangements for the purpose of measuring concentration, cost and resilience. No recommendation is made in respect of any particular person, asset, structure or jurisdiction. Constitutional provisions, foreign-ownership regimes, investment-promotion conditions, exchange-control carve-outs, tax deductions, substance requirements and enforcement arrangements all change and vary by jurisdiction and by activity; every such statement in this paper is carried from the source manuscript, is flagged in Appendix A, and requires confirmation with qualified counsel in the relevant jurisdiction before any reliance.

Appendix A. Claims requiring verification against primary sources

The following statements rest on legal or regulatory detail that changes, and **none of them is cited to the bibliography, which contains no primary sources for any individual jurisdiction**. They are reproduced from the source manuscript, are described rather than asserted, and require confirmation with counsel in the relevant jurisdiction before any reliance. **No measurement in this paper depends on any of them being true.**

A.1 The Central Asian centre. The constitutional statute establishing the centre and the constitutional provision entrenching it, with the procedure and majority required to amend either. The hierarchy of applicable law inside the centre and the gap-filling rule where its own instruments are silent. The centre's court: constitution, jurisdiction, bench, and any opt-in route for parties with no connection to the centre. The statutory basis, procedure, grounds of refusal and actual practice for enforcing the centre's judgments elsewhere in the host state,

and how many have in fact been enforced. The arbitration institution's rules and seat provisions, and whether awards under them are treated as host-state-seated for 1958 Convention purposes. The host state's accession date to that Convention and any reservations. The regulator's licensing categories and rulebooks. The names, numbers, dates and consolidated texts of every instrument relied on — **the source names none of them, deliberately.** The foundations and trust regimes, including separate legal personality, purpose vehicles, reserved powers, disclosure, any mandatory supervisory office, and continuation and transfer-out. The corporate income tax exemption: qualifying activities, participant categories, its sunset date, and the treatment at sunset. Currency-control relief: scope and limits. The host state's participation in automatic exchange of account information and its beneficial-ownership regime. Its current status with the international anti-money-laundering standard-setter. The correspondent banking relationships of any host-state bank the structure would rely on.

A.2 The operating base. The foreign-business statute: current restricted-activity lists, the definition of a foreign person, the nominee provision's wording and penalties, minimum capital, and licence practice. The bilateral treaty conferring ownership rights on nationals of one partner state: excepted sectors, certification procedure, termination provisions, and whether a chain of the constructed case's shape qualifies. The investment-promotion legislation: governing act, current categories and incentive schedule, the scope and conditions of land rights, and what happens to land held under them if promotion lapses. The land-ownership regime and any statutory cap on the hire of immovable property. The corporate income tax rate and the interaction between any promotion exemption and the state's implementation of the global minimum tax. Personal income tax: the residence day threshold, day-counting rules, and the current progressive scale. **The revenue order effective from 1 January 2024 on remitted foreign-source income, its number and scope, and any subsequent order addressing income derived before that date** — the source asserts nothing about a carve-out and none should be assumed. Withholding rates and transfer-pricing documentation thresholds. Exchange-control documentation and reporting. Alien-employment permit procedure, registered capital per permit, and any employee ratio. Succession over locally-situated assets. The arbitration legislation and the current position on enforcement of awards. **Regularisation of a historical nominee position: exposure for the elapsed period, available routes, and disclosure obligations — the most important item in this list.**

A.3 The corridor centre. The statute establishing the unified regulator, its powers, the scope of functions transferred from the sectoral regulators, and the regulatory challenge and appellate route. The special-economic-zone legislation as it applies to a unit in the centre. **The exchange-control statute and the precise scope of the treatment of units in the centre as persons resident outside the state, the notifications granting it, and the transactions covered — the load-bearing item.** The fund-management, banking, leasing, insurance and bullion frameworks, and whether each remains the operative instrument. The income-tax deduction relied on: section, conditions, qualifying activities, its ten-years-in-fifteen mechanic and commencement window, the minimum-alternate-tax position, and interaction with the general anti-avoidance rule. The relevant double-taxation agreement and its amending protocol, with grandfathering. The enforcement position: which foreign judgments are directly executable, the reciprocating-territory list, the arbitration legislation, the state's position under the 1958 Convention, and which disputes are not arbitrable. Classification of each vehicle type for automatic exchange purposes. Beneficial-ownership filings and access. Physical presence, minimum employment and net-worth requirements per vehicle type. Corporate residence and place-of-effective-management tests.

A.4 General. For all three positions: corporate residence and central-management-and-control tests in every jurisdiction from which decisions concerning the entity might be taken; the current status of each state under the international anti-money-laundering standard; and the recognition and enforcement position for both judgments and awards in every jurisdiction where the group holds assets.

Appendix B. The reporting standard for a marginal-leg decision

Any marginal jurisdictional position decided by the methods in this paper reports the following, and nothing less. It extends the reporting standard of Curzi [2026] to the marginal case.

```
1  the share basis, stated in one sentence
2  the exposure vector before and after the position, summing to 100.0%
3  the bloc assignment for the candidate, with its evidence class, and
   the named alternative assignment used for the wrong-call branch
4  bloc-HHI and N under BOTH assignments
5  G, L, and the break-even probability p* = L / (G + L)
6  the incremental annual cost c, decomposed into legal, substance,
   compliance and administration, and c as a percentage of C
7  the cost ceiling c = G, and the position's headroom against it
8  the assessed probability, stated as above or below p*, with the
   observable that would change the call
9  the date, and the diff against the previous measurement
```

Item 6 is the item the exercise exists for and it is the one most often omitted, because the incremental cost of a leg is spread across four invoices from three firms and is never totalled. **Item 9 decays first and matters most**, because a bloc call is invisible except as a difference between two measurements taken a year apart.

Item 8 is a statement of position, not an estimate. The standard deliberately does not ask for a number, because no defensible reference class exists for a one-off classification of a jurisdiction. It asks whether the assessor is above or below a stated threshold and what would move them across it — which is a recordable, datable, auditable claim, and an estimate would not be.

Chapter 59

Currency Drag, Cost Ceilings, and Funding-Source Independence

Abstract. Three positions are priced against the same surplus law, and in each the result runs against the jurisdiction being examined. On currency: at a 15 per cent annual hard-currency drag, an allocation of 25.1722 per cent of protected capital to the local unit drives the surplus to exactly −0.500000, cancelling the entire eta·ln(N) = 0.377583 gain that the whole jurisdictional redesign bought; a 10 per cent allocation costs 0.150000, or 39.73 per cent of that gain. What decides whether a weak currency is a cost or a loss is dwell: at a seven-day average dwell even a 40 per cent annual drag costs 0.767 per cent per pass, while at 180 days the same currency costs 19.726 per cent. The bloc arithmetic also cuts against the operating case: sizing a Turkish leg at 8 per cent of the structure lowers the effective bloc count from 2.9412 to 2.8324, so the case rests on customs-union access worth 1.35 per cent of group revenue. An African leg (Mauritius 8.0 per cent, Kenya 5.0) raises effective blocs from 2.924302 to 3.639010 and buys 0.076530 of surplus, setting a cost ceiling of roughly 25.5 per cent of the existing maintenance base: at C = 0.38 the surplus is −0.127902, worse than the −0.124431 baseline. Funding sources are measured on the same instrument: four Western lenders read 3.509 effective lenders but 1.00 effective funding blocs, and the redesigned book reaches 3.175 optimistically and 2.410 after an honest correlation adjustment. The source is found to use two incompatible bloc partitions of the same exposure vector, giving N = 2.941176 against 2.924302; both are reported, and five further discrepancies are printed rather than silently repaired.

Keywords: currency depreciation; dwell time; customs union; treaty networks; Islamic finance; sukuk; funding diversification

Series: D. Structure and Resilience **Paper:** CCP-059

1. Introduction

A structure that has arranged its jurisdictions well can still lose everything the arrangement bought, and it can lose it in three ways that never appear on an organisation chart: by holding a depreciating unit of account, by paying more to maintain a leg than the leg is worth, and by financing itself from four counterparties that all answer to one supervisor. This paper measures all three against the same instrument and reports, in each case, a result that runs against the position being examined.

The instrument is the concentration measurement of Curzi [2026] — HHI as the sum of squared exposure shares and the effective count as its reciprocal, with a naive four-jurisdiction count of 3.41 collapsing to exactly 1.00 once coordination blocs are recognised. The law it feeds is the storage law of Curzi [2026a], $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$, whose map is Lipschitz with modulus $1-\delta$ so that $\delta > 0$ makes it a contraction with steady state $Q^* = (\eta S - C)/\delta$, and its one-period balance $h = \eta \ln N - \delta Q - C$. **Everything below is a difference in that balance.**

Four measurements, and three of them go against the case being made.

First, a depreciating unit of account enters the surplus law through the erosion term and can cancel the entire jurisdictional gain. Treating an annual hard-currency drag d on a share w of protected capital as an addition to the erosion rate, $\delta_{\text{eff}} = \delta_0 + dw$, the balance becomes linear in w: at $d = 0.15$, $h(w) = -0.122417 - 1.5w$. **A 10 per cent allocation costs** 0.150000**, which is** 39.73% **of the entire** $\eta \ln N = 0.377583$ **that the jurisdictional redesign bought, and at** $w^* = 0.251722$ **the balance is exactly** -0.500000 **— the same number the conventional single-perimeter structure reads before any redesign at all.** Every entity stays in place, the chart still looks right, and the gain is gone.

Second, dwell rather than depreciation decides whether a weak currency is a cost or a loss. Erosion per pass is $dn/365$ for a receipt dwelling n days. **At a seven-day dwell even a 40 per cent annual drag costs** 0.767% **per pass — the order of a bank spread. At one hundred and eighty days the same currency costs** 19.726%**, which is a loss of capital.** The currency did not change; the policy did. That is the same shape Curzi [2026] reports for a corridor, where exposure is dwell divided by the year and nothing else.

Third, a leg proposed for diversification can lower the very count it was proposed to raise, and a leg that does raise the count can still fail on its invoice. An operating leg sized at 8 per cent of the structure lowers the effective bloc count from 2.941176 to 2.832412 when scored in the bridge row and to 2.919742 on the more generous scoring — **below the baseline on both**, because its share lands in a row that was already the largest; its case therefore rests entirely on a customs-union tariff saving worth 1.35% of group revenue, which is a logistics argument and not a resilience one. And a regional holding leg that raises the effective bloc count from 2.924302 to 3.639010, buying 0.076530 of surplus, **fails outright at the maintenance cost proposed for it**: at an incremental 0.08 the balance falls from -0.124431 to -0.127902. The break-even is exact — incremental cost below 0.076530, roughly 25.51% of the existing $C = 0.30$ — and correct cost attribution, charging the structure only the costs the structure causes, flips the decision to $h = -0.087902$.

Fourth, the same instrument pointed at the liability side finds the conventional metric moving the wrong way. Four lenders at 40 / 25 / 20 / 15 read 3.508772 effective lenders and 1.000000 **effective funding perimeters**, which is the same exact zero Curzi [2026] finds on the asset side. Replacing the syndicate with four tranches in four distinct perimeters *lowers* the naive effective source count to 3.174603 while raising effective funding blocs from 1.000000 to 3.174603. **A treasurer optimising lender count would reject the refinancing.** An honest correlation adjustment applied against this paper's own result — collapsing two tranches that draw on one pool of institutional liquidity — cuts the figure to 2.409639, and **the answer is reported as a range rather than as the flattering endpoint.**

What the paper does not do. It estimates no depreciation rate, no cost and no correlation. Every drag rate, every maintenance figure and every share is constructed and stated as such. **The contribution is the arithmetic that converts each of those inputs into a decision, and the arithmetic is worth exactly what the inputs are worth.**

Positioning. The collapse on which all four measurements depend — several units inside one perimeter behaving as one — is not new machinery. A small correlated shift producing a discontinuous collective outcome is the threshold-and-tipping mechanism: Granovetter [1978] establishes that the distribution of individual thresholds rather than the mean decides whether a collective outcome tips, and Schelling [1971] is the canonical demonstration that a mild individual preference produces a sharp aggregate segregation, generalised in Schelling [1978]. **Reading an individual-level parameter straight into an aggregate is exactly what a perimeter assignment does, and that literature is the standing warning against it.** Castellano

et al. [2000] supplies the closest formal analogue of a perimeter — the same interaction that raises local agreement freezes global diversity into non-interacting regions — and on the lending side Buldyrev et al. [2010] is the structural statement that coupling destroys the independence a naive count assumes.

The reflexive loop and its identification problem. A perimeter assignment is a forecast that conditions the behaviour it forecasts, which is the reflexive loop $T(x) = W(x, M(x))$ of Curzi [2026], whose founding reframing is that comfort is the felt total of present costs, so the present state is a local minimum of the comfort landscape by definition rather than by luck. **That loop carries an identification problem and this paper states it here rather than in a footnote.** Manski [1993] proves that a linear-in-means reflexive specification is **not identified from observational data**: the endogenous social effect cannot be separated from contextual effects or from correlated unobservables, however much data is collected, and Shalizi and Thomas [2011] reinforce it for network data. **No observation of the behaviour of lenders or of jurisdictions could therefore confirm a perimeter assignment derived reflexively.** The remedy adopted is stipulation with a defence — the perimeter test of Curzi [2026] — which Brock and Durlauf [2001] shows can be made rigorous, and its price is that **every perimeter assignment in this paper is a modelling position rather than an identified empirical result.**

A note on what half of this material is. The source manuscript is, in substantial part, a descriptive comparative survey of five jurisdictions and of a body of contract law. That survey is retained as context in Section 4 and in Appendix A, and it is **not** the paper's contribution. **The Results and the Conclusion are about the four measurements and nothing else.** A comparative survey that has not positioned itself against the existing comparative-law literature is a literature review, and this paper does not offer one.

2. Hypotheses

Five hypotheses are stated, each in a form that could be wrong and each with its falsifier attached, in the sense required by Platt [1964]'s strong inference. Meehl [1967] is the relevant warning: a directional prediction is cheap, and the defensible response is a point prediction with constants attached. **Two of the five are reported below as not met, and they are reported in the same plain language a success would have received.**

H1 — the allocation hypothesis. A local-currency allocation enters the surplus law through the erosion term and can cancel the whole jurisdictional gain without changing the effective bloc count. *Predicted:* with a 15 per cent annual hard-currency drag, $h(w) = -0.122417 - 1.5w$, so $w = 0.10$ costs 0.150000 — 39.73% of the 0.377583 gain — and $w^* = 0.251722$ gives $h = -0.500000$ exactly. *Falsified by:* a measured hard-currency return on local instruments over the holding period satisfying $(1+r)/(1+\Delta) \geq 1$, at which the allocation would carry no drag at all and the erosion term would be unchanged.

H2 — the dwell hypothesis. Whether a depreciating currency is a transaction cost or a capital loss is decided by dwell and not by the depreciation rate. *Predicted:* erosion per pass is $dn/365$, so a 40 per cent annual drag at a seven-day dwell costs 0.767% while a 15 per cent drag at one hundred and eighty days already costs 7.397% — **the lower drag costing nearly ten times more.** *Falsified by:* erosion that scales with the drag rate independently of dwell, or a conversion mechanism whose cost rises faster in frequency than the erosion it avoids.

H3 — the operating-leg hypothesis. Adding a genuinely non-aligned operating jurisdiction raises the effective bloc count. *Predicted:* an operating leg sized at 8 per cent of the structure raises N above the baseline 2.941176. *Falsified by:* any defensible scoring of the new jurisdiction under which N falls. **This is the first hypothesis reported as not met.**

H4 — the regional-leg sufficiency hypothesis. The regional holding leg's diversification gain exceeds its incremental maintenance cost. *Criterion, fixed in advance:* the leg is worth building if and only if its incremental maintenance cost is below the surplus it buys, $\eta\,\Delta \ln N = 0.076530$, which is 25.51% of the existing $C = 0.30$. *Falsified by:* an incremental maintenance cost at or above 0.076530 while η, δ and Q are held at the stated values. **This is the second**

hypothesis reported as not met, at the cost proposed for it.

H5 — the liability-side hypothesis. Independence bought on the funding side is cheaper per unit of effective bloc than independence bought on the jurisdictional side. *Predicted:* the refinancing raises effective funding blocs from 1.000000 to between 2.409639 and 3.174603 at an assumed incremental cost of 0.10, buying at least 1.409639 of effective bloc, while the regional jurisdictional leg buys 0.714708 of effective bloc at an assumed 0.08 — 0.070940 **of cost per unit of effective bloc against** 0.111930**.** *Falsified by:* a measured governance cost for the funding tranches exceeding the conservative surplus gain of 0.307817, or a demonstration that the tranches share a perimeter more completely than the conservative collapse assumes.

Not hypotheses, and labelled as such. Four statements in this paper are not falsifiable as they stand.

- *Applying the surplus term to funding-source perimeters* is an **extension by analogy and not a derivation.** The storage law of Curzi [2026a] was specified over protected capital and its coordination perimeters. **The robust part of Section 9 is the concentration arithmetic, which is definitional; the h figures there are indicative of magnitude only.** The source manuscript states this and the paper keeps it.
- *The logarithmic form of the input side* is a **modelling choice** carried from Curzi [2026], whose four-currency carrying cost — coin, hours, attention, foreclosure — is the corpus's statement that felt comfort decays over weeks while carrying cost does not decay at all. Section 3.9 states what would discriminate the form and concedes that nothing here does.
- *The parameter set* $\eta = 0.35$, $\delta_0 = 0.02$, $Q = 10$, $C = 0.30$ is a **modelling choice** fixed in advance of every measurement below, which is the preregistration discipline of Nosek et al. [2018] applied to a calculation.
- *The perimeter assignments* are **stipulated reference groups**, not identified from data [Manski, 1993, Brock and Durlauf, 2001].

3. Method

3.1 The instrument, and which effective count it is

The concentration instrument is not restated here; Curzi [2026] establishes it. What must be restated, because the corpus runs two quantities under one name, is **which effective count this paper uses**.

Throughout this paper the effective count is HHI^{-1}, the reciprocal of the sum of squared shares, **which is the Hill number of order two.** Hill [1973] defined the family ${}^{q}D = (\sum_i p_i^q)^{1/(1-q)}$, whose members return the number of equally-abundant types that would give the observed value of the underlying index; at $q = 2$ this is HHI^{-1} exactly. Jost [2006] is the statement that entropies are not diversities and that their numbers-equivalents are; Jost [2007] extends the correspondence to partitioned data, which is what a perimeter collapse is. The same arithmetic was reached independently by Simpson [1949] in ecology and Laakso and Taagepera [1979] in political science; the numbers-equivalent reading of the concentration index is Adelman [1969]'s; the index's history runs through Hirschman [1945] and Herfindahl [1950]; and its regulatory thresholds are stated in the merger guidelines [U.S. Department of Justice and Federal Trade Commission, 2023, 2010]. The exponential of Shannon entropy [Shannon, 1948] is the same family at order one, and HHI^{-1} is the exponential of the Rényi entropy of order two [Rényi, 1961].

The corpus's other effective count, the affine $N_{\text{eff}} = 1 + (N-1)u$ **of Curzi [2026], is not a Hill number of any order and carries no external support.** It is a linear interpolation between 1 and N. This paper does not use it and refers to it only by name. Hurlbert [1971] is the standing discipline: "diversity" has been used for so many distinct quantities that the word carries no content without its formula attached. Chao et al. [2014] show that low-order diversity estimators are biased on small samples, and a four- to eight-element exposure vector is a

small sample; no correction is applied and Section 13 records the gap. Golosov [2010] and Dunleavy and Boucek [2003] document that effective-number measures behave poorly in the dominated-unit regime, which is exactly where the as-found funding book of Section 9 sits.

No distributional or scaling claim is made anywhere in this paper. Lender shares of 40 / 25 / 20 / 15 and exposure shares descending from 29.1 to 8.0 look like the beginning of a heavy-tail argument. **This paper fits no distribution, characterises no tail, and asserts no power law.** Clauset et al. [2009] is the reason: a principled maximum-likelihood fit with a goodness-of-fit test and a likelihood-ratio comparison against alternatives leaves few canonical data sets standing, and a straight line on a log-log plot is not evidence. **That procedure was not run, because no claim requiring it is made.** Stumpf and Porter [2012] adds that passing it would still not identify a mechanism, and Broido and Clauset [2019] is the reason the same caution attaches to borrowings from the scale-free network literature. **Every effective count here is a statistic of a finite exposure vector with at most eight entries.**

3.2 Currency drag as an addition to the erosion term

The corpus's storage law already contains the term this measurement needs. Curzi [2026a] gives $Q_{t+1} = (1-\delta)Q_t + \eta S_t - C_t$, in which δ is the fraction of the stock lost per period independently of anything the structure does. **A depreciating unit of account is exactly such a loss**, and the modelling step this paper takes is to write it as one:

$$\delta_{\text{eff}} = \delta_0 + d\,w,$$

for an annual hard-currency drag d applied to a share w of protected capital. **This is a modelling choice and not a derivation, and it is stated as one.** It assumes the drag is proportional, applies uniformly over the period, and is independent of the rest of the structure — none of which is tested here. What it buys is that the currency question and the jurisdiction question become terms in the same equation and can be compared directly, which is the whole point of having an equation.

The consequence is immediate. Since $h = \eta \ln N - \delta_{\text{eff}} Q - C$ and N does not depend on w,

$$h(w) = h(0) - d\,Q\,w,$$

which is **linear in the allocation with slope** $-dQ$, and at $d = 0.15$ and $Q = 10$ that slope is -1.5. **The allocation at which the entire jurisdictional gain is cancelled follows in one line:** $w^* = \eta \ln N/(dQ)$**.**

Curzi [2026s] reaches the structurally identical result on a different input, showing that trust enters the same law as retention through $\delta_{\text{eff}} = \delta_0 + (1-a)(1-v)$ and produces a superlinear stock dividend, from 0.4679 at $a = 0.50$ to 5.0529 at $a = 0.99$. **Two quite different phenomena entering the model through the same term is either evidence that the term is the right abstraction or evidence that the term absorbs anything put near it, and this paper cannot distinguish those two readings.**

3.3 Dwell, and the transit-versus-store distinction

For a receipt converted after n days, the expected erosion per pass at an annual drag d is

$$\text{erosion per pass} = \frac{d\,n}{365}.$$

The functional form is the simplest one available and it is a modelling choice. It assumes the drag accrues linearly within the year and that all receipts dwell for a common time. A burst arrival process with the same mean dwell gives a different instantaneous distribution and the same mean, and nothing here tests which is right.

The distinction the formula makes is the portable result: evaluate a depreciating currency as a transit medium and it is cheap; evaluate it as a store and it is ruinous; and dwell is the entire difference. Curzi [2026] reaches the same conclusion for a corridor, where exposure is $t/365$ and one day exposes 0.27% of a year's flow against 49.32% at one hundred and eighty days.

3.4 Nominal against real, and the convention that must be stated

A nominal return in a depreciating currency is converted to a hard-currency return by

$$r_{\text{hard}} = \frac{1 + r_{\text{nom}}}{1 + \Delta} - 1,$$

where Δ is the **proportional rise in units of local currency per unit of hard currency**. **This convention must be stated because the alternative reading — depreciation as the proportional fall in the value of the local unit, giving $(1+r)(1-\Delta)-1$ — produces materially different numbers from the same words.** At a 50 per cent nominal return against a "60 per cent depreciation" the first convention gives -6.25% and the second gives -40.00%, a gap of 33.75 percentage points. **The source manuscript uses the first convention throughout without naming it; this paper names it.** Section 12.12 records the ambiguity as a finding rather than repairing it silently.

Compounding is then $(1-d)^k$ over k years at a constant annual hard-currency drag d, which is what turns a manageable annual figure into a structural one.

3.5 The cost-ceiling method

Curzi [2026] states the method this paper reuses on two further cases: a leg's break-even probability is $p^* = (c+L)/(G+L)$, so **the leg is worth nothing at any probability once its incremental cost exceeds 0.0267305, which is 8.9 per cent of the group's entire $C = 0.30$ maintenance budget.** Where a leg's classification is not in doubt the same arithmetic collapses to a single inequality:

$$\text{build the leg} \iff \Delta C < \eta\, \Delta \ln N.$$

The left side is an invoice and the right side is a measurement, and the discipline is that both must be written down. **Two errors are common and both are measured below.** The first is comparing the gross gain to nothing at all. The second is **mis-attribution**: charging the structure costs the structure did not cause. Where an operating company is required by the law of the place its operations are in, its cost belongs to the acquisition and not to the structuring, and Section 8.3 shows that correct attribution reverses the decision.

Curzi [2026j] supplies the governing constraint on what may be done with a positive balance — the increment rule $r \le h$, with a recommended reserve $r \le 0.7h$ since $r_{\max} = h$ is attained only at $U = 1$, and the $h < 0$ branch in which no interaction geometry can prevent order from decreasing. **Every structure measured in this paper is on that second branch**, which is the standing reason none of the legs priced here can be justified by an appeal to headroom.

3.6 Pointing the instrument at the liability side

The instrument is defined over shares and can be pointed at any exposure vector that sums to one. Pointed at the liability side, the shares are fractions of the group's committed funding and the units are lenders or tranches; the perimeter test is unchanged — **two funding sources are in the same perimeter when one supervisory or market decision reliably withdraws both.**

Two things about this extension are stated rather than assumed. The concentration arithmetic transfers exactly, because it is definitional and depends on nothing but a normalised share vector. **The surplus term does not transfer exactly**: the storage law was specified over

protected capital and its coordination perimeters, and reading $\eta \ln N$ over funding perimeters is an extension by analogy. Under the analogy discipline of Section 12.5 that is a **category-2 borrowing — a shared mathematical form with no claim that the underlying content carries over** — and the h figures in Section 9 are indicative of magnitude and are not derived.

Curzi [2026] makes the corresponding point about which axis is measured: entity-based US exposure of 27.3 per cent against a USD share of the balance sheet of 61.0 per cent and settlement volume clearing through the US system of approximately 58 per cent, so **the entity map understates dependency by more than a factor of two.** The liability side is a fourth axis on which the same understatement occurs, and it is the one least often measured.

3.7 The perimeter call is reflexive, and what that costs

A perimeter assignment is a forecast about how a set of jurisdictions or a set of lenders will behave under a coordinated action, and it changes the behaviour it forecasts once enough structures adopt it. Merton [1948] named the mechanism; Curzi [2026] carries it as the reflexive loop.

This paper adopts stipulation with a defence and pays the price. Perimeters are fixed in advance by the test of Curzi [2026], recorded with an evidence class, and stress-tested by running the arithmetic under the contested alternative — which Sections 7 and 9 both do, in each case against the paper's own preferred answer. Brock and Durlauf [2001] is the template. **No randomisation, instrument or partial-population variation was available, so no perimeter assignment here is an identified empirical result.**

Two further objections in the same family are recorded and not answered. **The Lucas critique applies to the framework's own parameters** [Lucas, 1976]: the framework asserts that systems change when modelled and then exempts η, δ and its own correlation assumptions from that reflexivity. **That is an inconsistency and it is not resolved here.** MacKenzie and Millo [2003] documents the benign form of the same loop in the field and with dates, and Hacking [1995] states the general form. And Goodhart [1975] with Campbell [1979] give the failure mode that follows adoption: an effective count taken as a target ceases to be a good measure. The only defence offered is the reporting standard of Appendix B.

3.8 The disclosure premise

Every position priced here is assessed against one premise, stated as a design constraint and not an ethical preference: **would the structure still do its job if every element of it were disclosed to every authority with jurisdiction over any part of it?**

The premise is not this work's to claim. Kerckhoffs [1883] stated it in 1883 — a system must remain secure when everything about it except the key is public — and Saltzer and Schroeder [1975] restate it as *open design*, the principle that a protection mechanism should not depend on the ignorance of attackers. Curzi [2026] carries the constraint in full for this series, establishing that channel separation is strengthened rather than weakened by disclosure because its effectiveness is entirely independent of whether anyone knows it is in place. **Citing Kerckhoffs converts a rule that reads as a preference into a recognised engineering standard with a hundred and forty years behind it.**

The same conclusion arrives from the regulatory direction. Automatic exchange of account information [Organisation for Economic Co-operation and Development, 2017], the anti-money-laundering recommendations that supervisors implement in common [Financial Action Task Force, 2012], and the substance and beneficial-ownership work of the base-erosion project [Organisation for Economic Co-operation and Development, 2015] have removed non-disclosure from the set of available design parameters. Brandeis [1914] is the older statement of the same instinct; de Goede [2011] documents the listing machinery as a technique of governance; Pistor [2019] is the scholarly account of how legal coding manufactures durability and simultaneously its sharpest critique; and Hansmann and Kraakman [2000] establishes

that asset partitioning is the irreducible function of entity law — **a separate legal personality is a partition, and its effectiveness likewise does not depend on anyone's ignorance of it.**

The premise is what makes the funding result of Section 9 measurable, because the most disclosed instrument in this paper — a listed certificate with a public prospectus, named supervisory organs, identified underlying assets and continuing obligations to an exchange — is also the one that breaks the single-perimeter dependency. **The strategy is non-obvious while every underlying fact is published by obligation, which is what the premise predicts.**

Curzi [2026] bounds what may be inferred from any of this, and the bound is carried rather than paraphrased away: category innocence is refused as a defence of acquisition, the doctrine forbids moving a harm through a discovered gap in a scheme, and the case in which the feedback loop runs through another person's will is conceded to be an unsettled seam the corpus does not close. **A paper that prices structures must state that the pricing is not a warrant**, and that is the contrast this edge carries.

3.9 What discriminates the logarithm, and what does not

Every result here inherits the shape of $\eta \ln N$. **The functional form is underdetermined by the evidence and this paper concedes that rather than defending it.** The corpus proposes $\ln(1 + (N-1)u)$ [Curzi, 2026]; Latané [1981] proposes a power law with exponent below one; Briscoe et al. [2006] propose $n \log n$ against the quadratic alternative. **All three are decreasing in marginal return and none of the measurements here distinguishes them**, because every measurement is a difference between two values of the same functional form.

One genuine strengthening is available. Peters and Gell-Mann [2016] derives a logarithmic form from multiplicative dynamics **with no psychological assumption at all**, which is a stronger warrant than the perceptual story of Fechner [1860] because it does not require the perceptual story to be true. Curzi [2026n] identifies the corpus's own functional as a Fechner law with threshold $u_0 = 1/(N-1)$. **Where the arithmetic here needs a warrant for the logarithm, the multiplicative-dynamics one is the one it needs.** Peters [2019] is the associated warning that an expected value and a time average are different objects — directly relevant to Section 6, where a compounding drag is exactly a multiplicative process.

4. The four positions and their institutional setting

This section is context and is not a result. It states what each position is for, at the level of detail the measurements require and no further. **Every statement in it about the law, regulation, institutions or policy of any individual jurisdiction, and every statement about the doctrinal content of any body of contract law, is described from the source manuscript and is uncited**, because the bibliography contains no statutes, regulations, court rules, registry documents, case reports or religious-legal sources for any individual jurisdiction or tradition. Appendix A lists those statements item by item and Section 13 records the gap. **They must be verified against the official sources before any reliance, and no measurement below depends on any of them being true** — each measurement depends only on a constructed share vector and the stated parameters.

Where a scored institutional comparison would otherwise attach a rating to a named jurisdiction on a characteristic this paper cannot support, the comparison is de-branded, following the convention of Curzi [2026]. **Jurisdictions are named below and in the tables as labels on constructed exposure vectors; they are not the subjects of assessed claims.**

4.1 Türkiye: the high-inflation industrial position

The candidate is a large civil-law industrial economy of the Germanic and Swiss codal line, inside a customs union with a major bloc for industrial goods, a member of a Western security alliance, and not a participant in that bloc's listing and designation machinery. Its currency has

depreciated substantially and persistently against hard currencies, and its foreign-exchange regime restricts the ability of parties both resident there to denominate or, in specified categories, to settle certain contracts in foreign currency — **a restriction that converts a financial risk into a drafting risk, which is the structurally interesting fact about it.** The lesson the source draws generalises past the jurisdiction: **currency risk in a high-inflation state does not stay financial, because the state responds to it with rules about what may be denominated in what.**

The proposal put to the constructed principal bundles three things: an operating company for tariff-advantaged final assembly, an allocation of group liquidity to local-currency instruments, and a real-estate acquisition with an associated residency benefit. **Three proposals in one document, which is how a bundle disguises the fact that one component is good and two are not.** Sections 6 and 7 price them separately, which is the entire method.

4.2 Mauritius, Kenya and the regional holding position

The candidate holding centre is a mixed-legal-system island state with two treaty networks that are routinely conflated — a double-taxation network and an investment-protection network — whose principal historical route into a large neighbouring market was closed by protocol, and whose substance requirements were reformed to replace a deemed credit with a partial-exemption regime conditioned on core income-generating activity. Two operating markets are considered alongside it, one of which sits inside a regional harmonised business-law family that the offshore literature generally ignores.

The relevant fact for the measurement is that a fully substantiated holding vehicle in that centre is not a cheap object: resident directors of calibre, local board meetings, a local principal bank account, a local audit, and enough activity to support the exemption. Section 8 prices exactly that, against an operating group in the target market with local counsel, local audit and transfer-pricing documentation for a cross-border supply chain.

4.3 Labuan, Islamic finance, and the alternative funding perimeter

The candidate is a mid-shore centre inside a federal Asian state with a statutory suite for offshore companies, a treaty network with material exclusions, and a distinctive constitutional arrangement giving a national authority binding force on questions of Sharia compliance for financial instruments. Around it sits Islamic finance as a legal family, whose two governing prohibitions produce a distinct instrument set; the relevant instrument here is the sukuk — a certificate representing an interest in identified assets, issued through a special-purpose vehicle, listed, and supervised by a named Sharia board. **The distinction between asset-backed and asset-based issues is what decides recovery**, and it is described in Appendix A rather than asserted here.

The relevant fact for the measurement is that this is a distinct funding perimeter: its investors, documentation, regulators and domiciles differ from those of the Western syndicated loan market, and a supervisory decision in one does not mechanically reach the other. Section 9 measures what that is worth and then argues against its own answer. **Nothing in this paper is a Sharia ruling or an evaluation of any doctrinal position**; the Declarations state this and Appendix A routes every doctrinal description to verification.

4.4 The disclosure position of all three

All three positions pass the premise of Section 3.8 on the source's analysis, and the analysis rather than the conclusion is what matters. The operating company exists because the operations are there and local law requires a local entity to conduct them; it employs people and pays local tax on local profits; **its existence is compelled rather than chosen.** The regional holding vehicle's resident directors take the investment decisions, in the centre, on analysis

performed there, so a treaty benefit claimed on disclosed facts rests on a contemporaneous file recording a commercial purpose independent of the benefit. **A listed sukuk is among the most disclosed instruments available**: a public prospectus, a published pronouncement, identified assets, a named board, audited accounts and continuing exchange obligations. **Nothing in any of the three depends on an authority not knowing something.**

5. The constructed case, and two incompatible baselines

CONSTRUCTED CASE. Every fact is invented and is labelled as such at each point of use. The case is the one carried through Series D and established in Curzi [2026]: a dual national, tax resident in one Gulf jurisdiction, with three heirs resident in three different jurisdictions, one under a forced-heirship regime; manufacturing operations in one South-East Asian state, a trading company in the Gulf, an intellectual-property portfolio, London commercial real estate, and a listed-securities portfolio, with revenue arriving in five currencies.

jurisdiction	share	share^2
UAE	29.1%	0.084681
Thailand	18.2%	0.033124
UK	16.4%	0.026896
Bermuda	13.6%	0.018496
India	11.8%	0.013924
Singapore	10.9%	0.011881
total	**100.0%**	**0.189002**

Naive HHI 0.189002, so 5.290949 effective jurisdictions. **The bloc figure, however, depends on a partition, and the source manuscript uses two different ones for the same six jurisdictions. Both are reported here rather than one.**

```
PARTITION A -- the perimeter test of CCP-052
  BRIDGE (UAE + Singapore)                 40.0%     0.160000
  NONALIGNED (Thailand + India)            30.0%     0.090000
  WEST (UK + Bermuda)                      30.0%     0.090000
  bloc-HHI = 0.340000     N = 2.941176
  eta ln(N) = 0.377583    h = -0.122417

PARTITION B -- as used in the regional-leg chapter of the source
  Gulf (UAE)                               29.1%     0.084681
  Asian (Thailand + India + Singapore)     40.9%     0.167281
  Western (UK + Bermuda)                   30.0%     0.090000
  bloc-HHI = 0.341962     N = 2.924302
  eta ln(N) = 0.375569    h = -0.124431
```

The two partitions differ in exactly one assignment — whether the second Asian bridging jurisdiction belongs with the Gulf bridge or with the non-aligned Asian row — and they give $N = 2.941176$ against 2.924302 and $h = -0.122417$ against -0.124431, a difference of 0.002014 in the balance. Section 12.12 records this as a discrepancy in the source. **Partition A is treated as canonical throughout this paper because it is the one derived from a stated perimeter test with recorded evidence classes; Partition B is reported alongside every result the source computed on it, so that the source's own figures remain checkable.**

The break-even effective count at the stated parameters is $N^* = e^{0.5/0.35} = 4.172734$. **Both baselines are below it before any position in this paper is added, and both remain below it after all of them.**

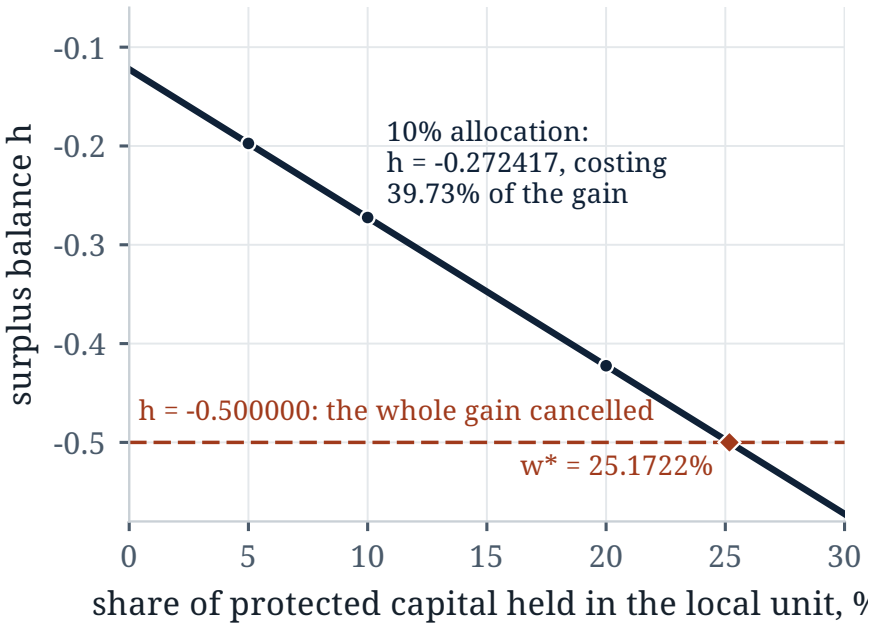


Figure 59.1: **A local-currency allocation cancels the jurisdictional gain without touching the effective bloc count: h falls linearly at 1.5 per unit of allocation, reaching the pre-redesign value of -0.500000 at 25.1722 per cent.** Modelled with a constructed 15 per cent annual hard-currency drag written into the erosion term as delta = 0.02 + 0.15w, at eta ln N = 0.377583, Q = 10 and C = 0.30; script `figs/src/CCP-059_fig01.py`.

6. Results: the currency allocation

6.1 The allocation, priced

Writing the hard-currency drag into the erosion term as $\delta_{\text{eff}} = 0.02 + 0.15w$ and holding $\eta \ln N = 0.377583$, $Q = 10$ and $C = 0.30$:

w	δ_{eff}	h	cost as a % of the gain
0.0000	0.0200	-0.122417	0.00%
0.0500	0.0275	-0.197417	19.86%
0.1000	0.0350	-0.272417	39.73%
0.2000	0.0500	-0.422417	79.45%
0.251722	**0.0578**	-0.500000	**100.00%**

The proposed 10 per cent allocation costs 0.150000 of h, which is 39.73% of the entire diversification gain of the designed structure. At $w^* = 0.251722$ the balance is exactly -0.500000 — identical to the conventional four-entity, single-perimeter structure that Curzi [2026] measures before any redesign. Every entity remains in place and the organisation chart still looks correct. H1 is supported in the model.

The answer would be no at a 5 per cent allocation too, because the same table shows a 5 per cent allocation costing 0.075000, which is 19.86% of the gain, for a yield that Section 6.3 shows is very likely negative in the group's reporting currency in any case.

6.2 Dwell, and the transit-versus-store result

Erosion per pass, $dn/365$:

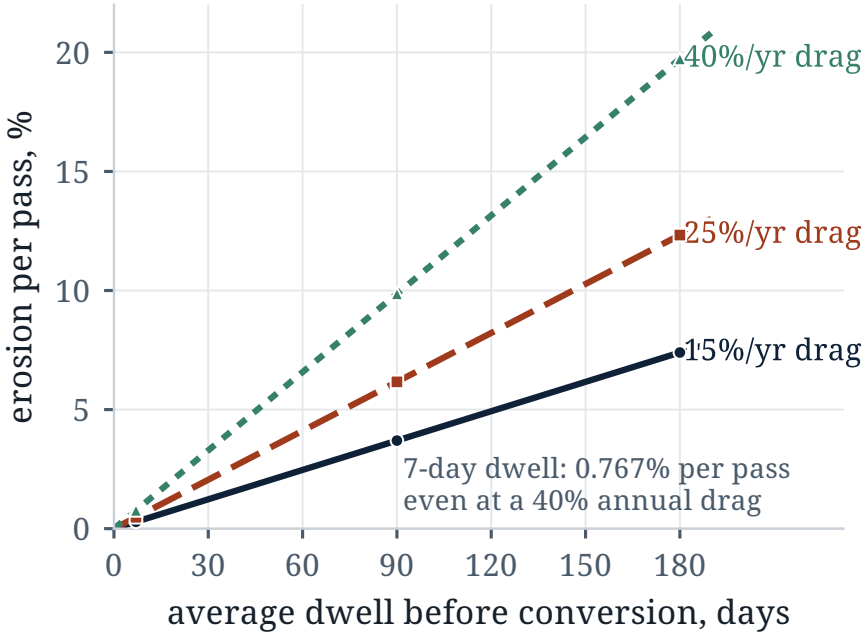


Figure 59.2: **Dwell, not the depreciation rate, decides whether a weak currency is a transaction cost or a capital loss: at seven days a 40 per cent annual drag costs 0.767 per cent per pass, while at 180 days a 15 per cent drag already costs 7.397 per cent.** Computed as d n / 365 from the constructed drag rates; script `figs/src/CCP-059_fig02.py`.

average dwell	drag 15%/yr	drag 25%/yr	drag 40%/yr
1 day	0.041%	0.068%	0.110%
7 days	0.288%	0.479%	0.767%
30 days	1.233%	2.055%	3.288%
90 days	3.699%	6.164%	9.863%
180 days	7.397%	12.329%	19.726%

At a seven-day dwell, even a 40 per cent annual currency drag costs 0.767% per pass, which is a transaction cost of the same order as a bank spread. At a one-hundred-and-eighty-day dwell the same currency costs 19.726%, which is a loss of capital. The currency did not change. The policy did. H2 is supported in the model.

The comparison that carries the result is between the cells rather than within them: a 15 per cent drag held for six months costs nearly ten times what a 40 per cent drag held for a week costs. A structure that fixes its dwell policy has done more about currency risk than one that changes currency.

6.3 Nominal against real, and compounding

The convention of Section 3.4 applies: Δ is the proportional rise in units of local currency per unit of hard currency, so $r_{\text{hard}} = (1 + r_{\text{nom}})/(1 + \Delta) - 1$.

nominal return	Δ	return in hard-currency terms
0%	20%	-16.67%
20%	20%	0.00%

nominal return	Δ	return in hard-currency terms
40%	35%	$+3.70\%$
45%	50%	-3.33%
50%	60%	-6.25%

A 50 per cent nominal return is a 6.25 per cent loss when the exchange rate rises 60 per cent, and the first number is the one that appears in the marketing. These combinations are constructed illustrations and not observed periods.

Compounding, $(1-d)^k - 1$, is where a manageable annual drag becomes a structural loss:

annual hard-currency drag	3 years	5 years	10 years
-5%	-14.26%	-22.62%	-40.13%
-10%	-27.10%	-40.95%	-65.13%
-15%	-38.59%	-55.63%	-80.31%
-20%	-48.80%	-67.23%	-89.26%

At a 15 per cent drag a local-currency position loses 55.63% of its hard-currency value in five years and 80.31% in ten. The three-year cell in that row is printed as -38.61% in the source; $0.85^3 - 1$ is -0.385875, so the correct figure is -38.59%, and Section 12.12 records it.

A compounding drag is exactly a multiplicative process, which is where Peters [2019] bites: the expected value of such a process and the time average a single holder experiences are different objects, and every figure in this section is the second rather than the first. That is the favourable direction — a time average is the honest quantity for a single structure — but it must be said which one is being reported.

7. Results: the operating leg, tested against its own bloc arithmetic

7.1 The component that survives on its own merits

The operating component is justified without reference to currency, tax or geopolitics:

```
EU-facing share of group revenue                      30.00%
constructed average duty avoided under the customs
  union for industrial goods                           4.50%
saving as a share of total group revenue               1.35%
```

That is a tariff and logistics decision of the same kind any plant location is, and it is subject to rules of origin the group does not control: if insufficient processing occurs locally, the 1.35% is not available and the operating case collapses entirely. The duty rate and the qualifying rules are constructed and uncited; Appendix A lists them.

The entity's job vector is $(H, O, B) = (0.02, 0.98, 0.00)$:

```
S = 0.0004 + 0.9604 + 0.0000 = 0.9608
effective jobs = 1 / 0.9608 = 1.0408          SPECIALIST
```

The purest entity in the architecture, and deliberately so. The status of that figure is the one Curzi [2026] states: the job vector is assigned rather than measured, treated as ratio-scale, and never tested against the conjoint conditions [Michell, 1997, Luce and Tukey, 1964]. **The ordering survives; the magnitude does not.**

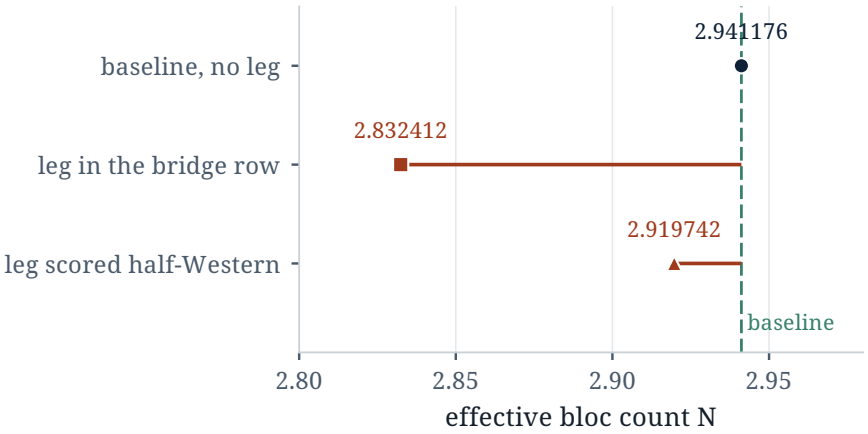


Figure 59.3: **The operating leg lowers the effective bloc count it was proposed to raise: 2.832412 when its share is scored into the bridge row and 2.919742 on the more generous scoring, against a baseline of 2.941176 on both.** Measured from the constructed share vectors of Section 7.2; the axis is truncated to resolve differences in the third decimal, and a dot plot is used rather than bars for that reason. Script `figs/src/CCP-059_fig03.py`.

7.2 H3 is not met

Suppose the operating leg grows to carry 8 per cent of the structure, with the existing allocations diluted proportionally. Score the new jurisdiction in the bridge row, on the ground that a perimeter a state can be pressured into joining is not an independent one:

```
UAE 26.772  Thailand 16.744  UK 15.088  Bermuda 12.512
India 10.856  Singapore 10.028  new leg 8.000        (sums to 100.0%)

blocs   BRIDGE 44.800   NONALIGNED 27.600   WEST 27.600
        bloc-HHI = 0.200704 + 0.076176 + 0.076176 = 0.353056
        EFFECTIVE BLOCS N = 2.832412      against 2.941176 before
```

And the alternative scoring, treating the new jurisdiction as half-Western on account of its alliance membership and its customs union:

```
blocs   BRIDGE 40.800   WEST 31.600   NONALIGNED 27.600
        bloc-HHI = 0.166464 + 0.099856 + 0.076176 = 0.342496
        EFFECTIVE BLOCS N = 2.919742      still below 2.941176
```

On both scorings, adding the leg as a balance-sheet allocation makes the structure less diversified, because its share lands in a row that is already the largest. H3 predicted the opposite and is reported here as not met. The result is one the source manuscript did not want and it is the correct one.

The failure also identifies exactly what would have to change for the answer to differ. The leg would improve N only if its share came *out of* the bridge row rather than being added to it — that is, by displacing existing bridge share rather than supplementing it. **That is a decision about the bridge layer and not about the new jurisdiction at all**, which is the kind of conclusion a measurement produces and an argument does not. Curzi [2026] reports the same sign on a different pair of additions: two Gulf additions raise the naive effective count from 5.29 to 7.13, an increase of 34.8 per cent, while the effective bloc count falls from 2.94 to 2.84, a decrease of 3.4 per cent — **more jurisdictions, less independence.**

7.3 The honest accounting of a leg that helps one term and hurts another

The operating leg improves the input side of the storage law, because cheap genuine substance lowers C, and it very slightly worsens N. **Those are different terms in the same equation and it is entirely coherent for a decision to help one and hurt the other. What is not coherent is to claim both**, which is what the bundled proposal did.

component	decision	basis
operating company	**yes**	customs-union saving of 1.35% of group revenue, subject to verified rules of origin
job assignment	OPERATING	vector $(0.02, 0.98, 0.00)$, 1.0408 effective jobs
local-currency balance-sheet position	**no**	a 10% allocation costs 39.73% of the diversification gain
maximum local balance	60 days of local operating expenditure	at a 40% drag a 60-day balance erodes 6.58% per turn
maximum average dwell	7 days	0.767% per pass at a 40% drag
bridging role	**not yet**	and only by displacing bridge-row share, never by adding to it
real estate	operating premises only	sized as a local-currency asset

8. Results: the regional leg's cost ceiling

8.1 The gain

On Partition B, which is the partition the source computes this case on:

```
AS FOUND
  naive HHI = 0.189002   ->  5.290949 effective jurisdictions
  bloc shares Western 30.0 | Gulf 29.1 | Asian 40.9
  bloc-HHI = 0.341962    ->  N = 2.924302     ln(N) = 1.073056
  h = 0.35 x 1.073056 - 0.02 x 10 - 0.30 = -0.124431

WITH THE REGIONAL LEG
  UAE 27.0  Thailand 16.0  UK 14.0  Bermuda 11.0  India 10.0
  Singapore 9.0  Mauritius 8.0  Kenya 5.0
  naive HHI = 0.157200   ->  6.361323 effective jurisdictions
  bloc shares Western 25.0 | Gulf 27.0 | Asian 35.0 | African 13.0
  bloc-HHI = 0.274800    ->  N = 3.639010     ln(N) = 1.291712
  eta x d ln(N) = 0.35 x 0.218656 = 0.076530
```

Effective blocs rise from 2.924302 to 3.639010 and Western-perimeter exposure falls from 30.0 per cent to 25.0 per cent. **The entire benefit of the leg, in the units of the storage law, is** 0.076530**.**

On the canonical Partition A the same reallocation gives bloc shares BRIDGE 36.0 | NON-ALIGNED 26.0 | WEST 25.0 | AFRICAN 13.0, so bloc-HHI 0.276600**,** $N = 3.615329$ **and a gain of** 0.072231 **— 5.6 per cent smaller.** Both are reported; Section 12.12 records that the source computes only the first and does not say the result depends on the partition.

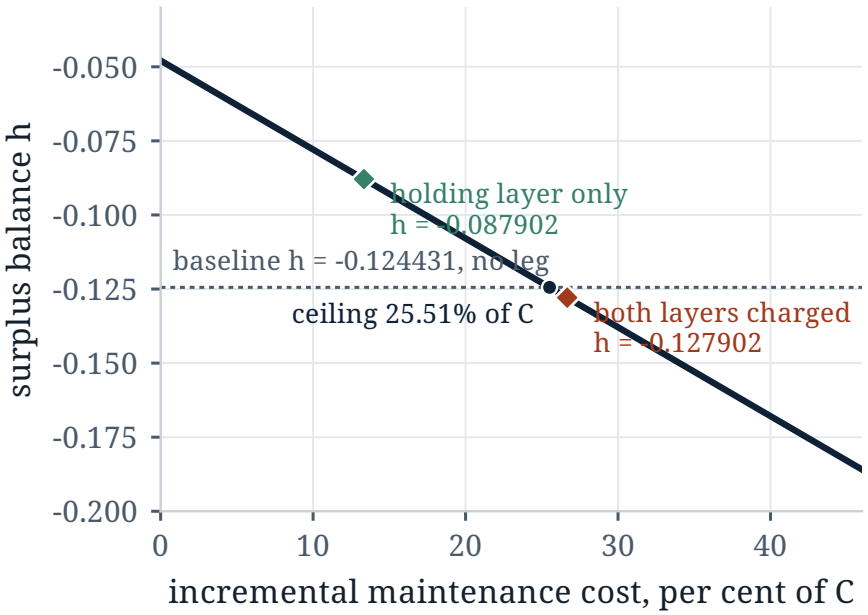


Figure 59.4: **The regional leg's whole case is a cost ceiling: at the proposed incremental maintenance of 0.08 the surplus balance falls from -0.124431 to -0.127902, and only at 0.04 — the holding layer alone, correctly attributed — does it rise to -0.087902.** Computed from the Partition B measurement with eta ln N = 0.452098; the ceiling is the gain 0.076530, which is 25.51 per cent of C = 0.30. Script `figs/src/CCP-059_fig04.py`.

8.2 H4 is not met at the cost proposed

A fully substantiated regional holding vehicle with resident directors, local board meetings, a local principal bank account and a local audit is not a cheap object, and neither is an operating group with local counsel, local audit, transfer-pricing documentation for a cross-border supply chain, and the compliance overhead of a jurisdiction under heightened monitoring. **Assume — a constructed assumption and not a quotation — that the incremental all-in annual maintenance raises C from 0.30 to 0.38.**

```
h(with the regional leg) = 0.35 x 1.291712 - 0.02 x 10 - 0.38
                         = 0.452099 - 0.200000 - 0.380000
                         = -0.127902

baseline (Partition B)   = -0.124431
```

Against the baseline, the regional leg made the structure worse. The diversification gain of 0.076530 was smaller than the maintenance increment of 0.080000, and the result is more jurisdictions, more filings, more counsel, a higher effective bloc count, and lower surplus. **H4's criterion was fixed in advance — incremental cost below 0.076530 — and it was not met.**

```
the regional leg is worth building if and only if
    incremental C  <  0.076530

against an existing maintenance budget of C = 0.30, that is a
ceiling of 25.51% of the entire existing compliance cost base
(24.08% on the canonical partition)
```

This is the sharpest local form of the corpus's own companion result: diversification raises the ceiling, it does not pay the bill, and a bloc gain bought at too high a maintenance price

is a loss. Curzi [2026] states the same inequality for a marginal leg whose classification is uncertain, where the position is worth nothing at any probability once its incremental cost exceeds 0.0267305, 8.9 per cent of the group's entire maintenance budget.

8.3 Three things that change the answer, and one that reverses it

Size. The 0.076530 gain is a function of shares rather than of absolute capital, while the cost is largely fixed. Double the group with the same share vector and the gain doubles in absolute terms while C barely moves. **There is a threshold below which the leg cannot be justified and above which it obviously can, and the ceiling above is how to find it.**

Cost attribution, which reverses the decision. In the constructed case the operating company is not optional: the group is buying a distribution business with employees, warehouses, licences and local banking, and local law requires a local entity to hold it. **That cost belongs to the acquisition and not to the structuring.** Re-running with only the incremental *holding-layer* cost — an assumed 0.04 rather than 0.08:

```
h = 0.35 x 1.291712 - 0.02 x 10 - 0.34
  = 0.452099 - 0.200000 - 0.340000
  = -0.087902        against a baseline of -0.124431
```

which beats the baseline by 0.036529. **Correct cost attribution flips the decision, and misattribution is the commonest analytical error in this field. Charge to the structure only the costs the structure causes.**

Whether a treaty benefit is real, and which term it enters. Nothing above includes a tax saving, deliberately. **If a treaty reduces withholding relative to the domestic rate, that is a reduction in δ — the erosion term — and not an increase in $\ln N$.** They are different terms in the law and must not be added together as though they were the same thing. Curzi [2026] prices the corresponding move on the substance side and gives the cost multiples exactly: 5,350 against 1,550 index units, a factor of 3.45, for a pure equity-holding entity, and 19,350 against 1,550, a factor of 12.48, for a relevant-activity entity, with break-even at USD 7,740,000 of held value at 25 basis points. **A substance requirement and a treaty benefit are two halves of the same invoice.**

9. Results: funding-source independence

9.1 The funding book as found

The instrument is defined over shares and is pointed here at the liability side, measuring how many independent funding perimeters the group actually has.

lender	share	share^2	perimeter
Bank A	40.0%	0.160000	Western
Bank B	25.0%	0.062500	Western
Bank C	20.0%	0.040000	Western
Bank D	15.0%	0.022500	Western
total	**100.0%**	**0.285000**	

```
naive HHI = 0.285000     ->  effective lenders = 3.508772
perimeter HHI = 1.000000 ->  EFFECTIVE FUNDING BLOCS = 1.000000
ln(1.000000) = 0
```

Four lenders, one perimeter, and the same exact zero Curzi [2026] finds on the asset side. A supervisory instruction, a de-risking decision or a designation affecting Western bank lending to a South-East Asian manufacturer reaches all four at once, because they share supervisors, share correspondent networks and read the same guidance. **The group has diversified its lenders and not its funding.**

9.2 The redesigned book, and the metric that moves the wrong way

The refinancing replaces the single syndicate with four tranches in four perimeters:

tranche	share	share^2	perimeter
Western syndicate, reduced	45.0%	0.202500	Western
listed sukuk issue	25.0%	0.062500	Islamic capital markets
Gulf institutional agency	20.0%	0.040000	Gulf institutional
onshore local-currency facility	10.0%	0.010000	Asian non-aligned
total	**100.0%**	**0.315000**	

```
naive HHI = 0.315000        ->  effective sources = 3.174603
perimeter HHI = 0.315000    ->  EFFECTIVE FUNDING BLOCS = 3.174603
ln(3.174603) = 1.155183     eta ln(N) = 0.404314
```

Effective sources FELL, from 3.508772 **to** 3.174603**, while effective funding blocs ROSE from** 1.000000 **to** 3.174603**. The conventional metric got worse and the real one went from zero to something.** This is the mirror image of the asset-side finding of Curzi [2026] and it is worth dwelling on: **a treasurer optimising lender count would reject this refinancing.**

9.3 The correlation adjustment, applied against this paper's own result

The whole correction of Curzi [2026] is that shares sharing a coordination mechanism must be collapsed before squaring, and intellectual honesty requires applying it to this paper's own answer. **Are the sukuk tranche and the Gulf agency tranche independent?**

Partly. They have different investors, different documentation, different regulators and different domiciles. But both draw on one pool of institutional liquidity, and a severe contraction in that pool would reach both. **The conservative reading collapses them:**

```
Western 45.0 | Islamic (sukuk + agency) 45.0 | Asian 10.0
perimeter HHI = 0.202500 + 0.202500 + 0.010000 = 0.415000
EFFECTIVE FUNDING BLOCS = 2.409639
ln(2.409639) = 0.879477     eta ln(N) = 0.307817
```

The honest answer is a range: effective funding blocs between 2.409639 **and** 3.174603**, not a point estimate. Reporting the flattering figure alone would be exactly the error the instrument exists to correct. Both ends of the range are far better than** 1.000000**, which is the finding that matters.**

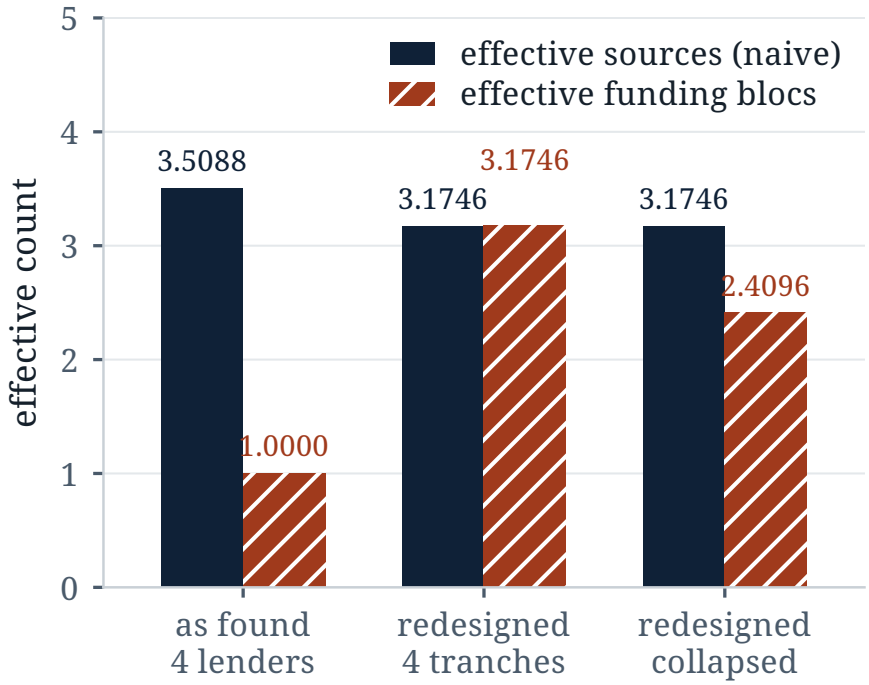


Figure 59.5: **Diversifying lenders and diversifying funding are different operations: the as-found book reads 3.508772 effective lenders and exactly 1.000000 effective funding perimeters, while the refinancing lowers the naive count to 3.174603 and raises the perimeter count to between 2.409639 and 3.174603.** Measured from the constructed share vectors of Sections 9.1 to 9.3; the range is the honest correlation adjustment applied against the paper's own result. Script `figs/src/CCP-059_fig05.py`.

9.4 What it is worth, with the caveat stated first

The caveat comes first because it governs everything after it. The storage law was specified over protected capital and its coordination perimeters, and applying its surplus term to funding-source perimeters is an **extension by analogy and not a derivation. The robust part of this section is the concentration arithmetic, which is definitional; the h figures below are indicative of magnitude and should be read as such.**

```
conservative:  0.35 x 0.879477 = 0.307817
optimistic:    0.35 x 1.155183 = 0.404314
```

Against that sits the cost. The Islamic tranches carry a permanent governance layer: supervisory-board fees, a compliance audit, structuring counsel with the relevant expertise, asset identification and maintenance for the sukuk, a listing, and purification procedures. **Assume — a constructed assumption — that this adds 0.10 to C. Against a conservative surplus gain of 0.307817 that is clearly worth doing: the incremental cost would have to exceed 0.307817 to destroy value.**

The contrast with Section 8 is the finding. There, a jurisdictional leg raised effective blocs by 0.714708 and failed on cost, because jurisdictional diversity requires paying for courts, counsel, substance and filings in each additional jurisdiction. Here, funding diversity raised effective blocs by at least 1.409639 at roughly a quarter more cost, because a funding tranche carries transaction cost rather than permanent jurisdictional overhead.

axis	effective blocs bought	assumed incremental cost	cost per unit of effective bloc
jurisdictional, regional leg	0.714708	0.080000	0.111930
funding, conservative reading	1.409639	0.100000	0.070940
funding, optimistic reading	2.174603	0.100000	0.045986

Diversifying liabilities is cheaper per unit of effective bloc than diversifying jurisdictions, and a designer working under a maintenance-cost constraint should look at the liability side first. H5 is supported in the model. **That conclusion falls out of the arithmetic rather than from a preference, and it is the most useful thing in this paper.**

10. Discussion

10.1 What follows

Three of the four measurements went against the position being examined, and that pattern is the paper's substantive claim. A currency allocation proposed for yield cancels the jurisdictional gain; an operating leg proposed for diversification lowers the effective bloc count; a regional leg proposed as an improvement fails on its own invoice. **In each case the failure was invisible from inside the proposal and visible from the equation**, because each proposal was argued on one term of the law while its cost landed on another.

Erosion is the term the field under-measures. δ is where currency drag, treaty leakage, tax friction and trust decay all arrive, and none of those appears on an organisation chart. Curzi [2026s] makes the same point on a different input, and the shared structure is that **the input side of the surplus law is bounded by a logarithm while the erosion side is linear in the stock: a bad enough δ beats any achievable N.** The measurement in Section 6 is that statement with numbers on it — a slope of 1.5 per unit of allocation against a total available gain of 0.377583.

A bundle must be unbundled before it can be priced. The three-component proposal of Section 7 contained one good component and two bad ones, and its persuasive force came entirely from being presented as one decision. **Pricing each component against the same law is what separated them**, and there is nothing clever about it beyond the discipline of refusing to evaluate a bundle.

And the liability side is cheaper than the asset side. That is the one result here that points at something to do rather than something not to do, and it falls out of a cost comparison rather than a preference: a funding tranche carries transaction cost, a jurisdiction carries permanent overhead.

10.2 What does not follow

None of these results says a jurisdiction, a currency or a funding market is good or bad. They say what a constructed share vector and a constructed cost model imply at stated parameters. Change η from 0.35 to 0.50 and the break-even count falls from 4.172734 to 2.718282, below the designed structure's 2.941176, and every conclusion about sufficiency reverses.

The failed hypotheses do not say the positions are worthless. H3 failed and the operating leg still stands, on a tariff argument the concentration instrument does not measure. H4 failed at the cost proposed and succeeded at the cost correctly attributed. **A measurement that kills the stated reason for a decision while leaving a different reason standing is doing its job.**

The funding result is not a derivation. Section 9.4 states this and it is repeated here because it is the most likely thing to be over-read from this paper: the concentration arithmetic on the liability side is exact, and the surplus figures attached to it are an analogy.

And nothing here is an empirical result. Every figure is exact arithmetic on stipulated inputs. **The paper cannot be wrong in the way a measurement can be wrong, and that is a weakness rather than a strength.**

10.3 Parallel results in the corpus

Curzi [2026] establishes a conservation result of a kind this paper does not have and would benefit from: the sum of drives $b_i = \eta_i F - c_i$ is exactly invariant under any transfer of cost among participants, verified at $+0.643688229683$ and $+1.268940530229$ with residual exactly zero. **The corresponding statement here would be that moving maintenance cost between the acquisition and the structure leaves some group-level quantity fixed — and Section 8.3 shows that it does not, because the two are charged against different balances.** The absence of a conservation law is the honest reading and Section 12.4 states it.

Curzi [2026] and Curzi [2026] supply structurally parallel results in a different domain and are cited as parallels rather than as support. The minimum sustainable frequency $f_{\min} = ke^{2c/\eta}$, with the trichotomy in which f greater than, equal to or less than $f_{\min}$ corresponds to Q_∞ greater than, equal to or less than zero, **is the same cost-against-conversion threshold this paper writes as** $\Delta C < \eta\,\Delta \ln N$, with the cost-to-efficiency ratio in an exponent rather than on the other side of an inequality. And the opposition identity $u = 1 - \rho^2$, which depends only on $|\rho|$ so that $\rho = +1$ and $\rho = -1$ both give $u = 0$, $F = 0$ and $f_{\text{rel}} = k = 0.4834$, is **the same insensitivity to sign that makes a perimeter collapse indifferent to whether the members coordinate cooperatively or antagonistically.** Neither is evidence for anything here.

11. Failure modes

Eleven failure modes are recorded, each with the measurement that detects it.

1. Holding anything in the depreciating unit. Detected by Section 6.1: any $w > 0$ costs $1.5w$ of balance at a 15 per cent drag.

2. Confusing nominal returns with real ones. Detected by Section 6.3 and by the convention statement of Section 3.4, without which the same words give answers 33.75 percentage points apart.

3. Retaining earnings by default. A balance that is not swept is an allocation nobody decided to make. Detected by the dwell measurement.

4. Drafting in a foreign currency without checking the local denomination rules. The source records that a group which has correctly decided to price everything in hard currency can find that a particular domestic contract may not be. **This is the failure mode with no arithmetic attached**, and it is listed because it is the one most often missed.

5. Assuming security-alliance membership implies Western banking treatment. Detected by measuring correspondent access rather than treaty membership.

6. Assuming non-participation in a designation regime implies immunity from it. The perimeter test asks what a decision reaches, not who signed what.

7. Counting a pressurable state as an independent perimeter. Detected by Section 7.2: both scorings lower N.

8. Buying an asset for a collateral benefit. Detected by pricing the two decisions separately, which is the whole method of Section 7.

9. Relying on foreign judgment enforcement where arbitration is the available route. Recorded rather than measured; United Nations [1958] is the instrument and Born [2021] and van den Berg [1981] the standard treatments.

10. Underbuilding the compliance function. A permanent addition to C that does not fall away when the circumstances that caused it do, because institutional memory in correspondent banking is long.

11. Sizing the opportunity from the headline and the risk from the average. The general form of failure modes 1, 2 and 8 together, and the reason every table in this paper prints both ends of every range.

12. Verification and rigour audit

The results are decomposed into postulates that can be checked one at a time. Each is atomic, each carries a status, and the compounding chains name what each step consumes. **If dropping a postulate leaves a result standing, it is not listed.**

12.1 Dimensional audit

symbol	quantity	dimension	unit	status
s_i	exposure share	$[1]$	dimensionless	defined; sums to 1 by P1
HHI	concentration index	$[1]$	dimensionless	derived from shares
N	effective bloc or perimeter count	$[\#]$**, a count**	blocs, perimeters	derived as HHI^{-1}
$\ln N$	surplus term argument	$[1]$	nats	derived
Q	protected capital	$[C]$	index units	defined; **no operational unit is supplied by the source**
δ_0, δ_{eff}	erosion rate	$[T^{-1}]$	per period	assumed

symbol	quantity	dimension	unit	status
d	hard-currency drag rate	$[T^{-1}]$	per year	assumed (constructed)
w	share of Q in the local unit	$[1]$	dimensionless	assumed (constructed)
n	dwell before conversion	$[T]$	days	assumed (constructed)
C, ΔC	maintenance cost flow and its increment	$[C\,T^{-1}]$	index units per period	assumed
η	conversion efficiency	$[C\,T^{-1}]$ per nat	index units per period per nat	assumed
h	surplus balance	$[C\,T^{-1}]$	index units per period	derived
r_{nom}, Δ	nominal return, exchange-rate rise	$[1]$	dimensionless, per period	assumed (constructed)
S	job-concentration index	$[1]$	dimensionless	derived from a job vector

Does $\delta_{\text{eff}} = \delta_0 + dw$ balance? δ_0 carries $[T^{-1}]$; d carries $[T^{-1}]$; w is dimensionless. **It balances only if d is expressed per the same period as δ_0.** Every figure in Section 6 uses annual rates for both, and **the source is silent on the period of δ; this paper assigns "per year" explicitly rather than assuming it, and records that assignment as its own.**

Does the erosion-per-pass formula balance? $dn/365$ multiplies a per-year rate by days divided by days per year, giving a dimensionless fraction. **It balances**, and it is the one place in the paper where a unit conversion does real work: a drag quoted per year and a dwell quoted in days cannot be combined without it, and combining them wrongly is the commonest arithmetic error in this material.

Where no operational unit exists, this audit says so. Q is protected capital in index units and no conversion to currency is supplied anywhere in the corpus. **Every figure in C and h is therefore in an arbitrary scale fixed by the stated parameters.** The effective counts, the cost ceilings expressed as percentages of C, and w^* are all unit-free and survive this; the statement that h is -0.127902 does not mean anything until the index unit is defined, and it is not.

12.2 Significant figures and epistemic status

- $N = 2.941176$, 2.924302, 2.832412, 2.919742, 3.639010, 3.615329, 3.508772, 3.174603 and 2.409639 are **exact arithmetic** on the stated shares, being reciprocals of exact rationals. Their inputs are constructed and exact as stated.
- $w^* = 0.251722$ is **exact given the parameters**, being $0.377583/1.5$. **The source rounds it to 25.17 per cent and then describes the resulting balance as exactly -0.500000; at $w = 0.2517$ the balance is -0.499967.** Section 12.12 records it.
- The cost ceilings 0.076530 and 0.072231 are exact given the parameters and the partition, and the parameters are assumed rather than estimated.
- The drag, dwell, maintenance and duty figures are **constructed model parameters**, and every quantity derived from them is **illustrative model output**. Six digits on -0.127902 are defensible only because the inputs are exact by stipulation.
- **No quantity in this paper is a fitted or measured estimate, so no uncertainty interval is reported anywhere.** That is not precision; it is the absence of empirical content.

One constant appears at fifteen figures through an internal citation and its status must be stated. Curzi [2026] carries $f_{\text{min}} = ke^{2c/\eta}$ with $k = 0.483425732988414$ Hz. $k = (9/8)^2/\varphi^2$ **is a DEFINED constant** — an exact ratio of exact quantities — and its fifteen figures are legitimate

for that reason and no other. **Attaching Hz to it is a separate and much stronger claim than computing it, and that attachment is POSTULATED.** Because φ appears, Markowsky [1992] is cited: the golden-ratio literature is the most heavily over-claimed in mathematics and the corrective is not optional. **Nothing in this paper depends on k.**

12.3 Limits and boundary behaviour

limit	behaviour	reading
$w \to 0$	$h \to h(0)$	no allocation, no drag
$w \to 1$	$h \to h(0) - dQ$	the whole stock in the local unit; at $d = 0.15$, $h = -1.622417$
$d \to 0$	$h(w) \to h(0)$ for all w	a stable currency has no allocation cost
$d \to \infty$	$w^* \to 0$	at unbounded drag no allocation is admissible
$n \to 0$	erosion $\to 0$	instantaneous conversion has no cost
$n \to 365$	erosion $\to d$	a receipt held a year loses the whole annual drag
$n > 365$	erosion $> d$, **outside the domain**	the linear form is valid only on $[0, 365]$
$\Delta C \to 0$	build the leg always	a free leg is always worth having
$\Delta \ln N \to 0$	no cost is admissible	a leg buying no independence is never worth anything
$N \to 1^+$	$h \to -(\delta Q + C)$	one perimeter pays the whole bill out of stock
$N \to \infty$	$h \to \infty$ logarithmically	**a defect of the model, not a result**

Three of these are constraints on the domain rather than results. The dwell model is linear and therefore invalid above a year. The drag model is linear in w and therefore assumes the drag does not itself depend on the size of the allocation, which for a holder large relative to the market it would. And the surplus law grows without bound in $\ln N$, which the empirical literature contradicts: Uzzi and Spiro [2005] finds the relation between diversity and performance turning over rather than rising, May [1972] showed that increasing diversity and connectance destabilises randomly assembled systems, and Sagan [1993] documents redundancy creating new failure modes rather than removing them. **The model's domain is small N and every prescription here is confined to it.**

And $\delta \to 0$ is a pathology rather than an ideal, since nothing then ever decays out of the stock and any positive $\eta \ln N$ exceeding C accumulates without bound. Curzi [2026a] states the contraction property that makes $\delta > 0$ necessary: the map is Lipschitz with modulus $1 - \delta$, so a strictly positive δ is what gives a unique start-independent fixed point at $Q^* = (\eta S - C)/\delta$.

12.4 Invariance

Every effective count here is invariant under relabelling and under any rescaling of the underlying exposures that preserves the share vector. Doubling every exposure in currency terms leaves every s_i unchanged and therefore leaves HHI, N and $\ln N$ unchanged. **Tested at**

how many scales? Four — the six-, seven-, eight- and four-element vectors reported here. The invariance is algebraic, not empirical, and no claim of empirical scale-freedom is made.

w^* **is not invariant under a change in the drag**, since $w^* = \eta \ln N/(dQ)$: halving the drag doubles the admissible allocation. **The result "a 10 per cent allocation costs 39.73 per cent of the gain" is therefore a statement about a 15 per cent drag and about nothing else**, and a reader with a different drag should recompute rather than quote.

Nothing is invariant under regrouping, which is the content of the perimeter collapse rather than a defect — and Section 5 shows exactly what it costs, since two defensible partitions of the same vector give 2.941176 and 2.924302.

What is conserved, and in what units. Mirowski [1989] shows that the characteristic failure of this genre is borrowing the notation of a conserved quantity without the conservation law that gave it content. **The objection is met here and the answer is short.** The conserved quantity is the exposure share vector, $\sum_i s_i = 1$ exactly and dimensionless, and it is that constraint and nothing else that bounds the index. Q **is a stock with no stated conservation law and no operational unit, and this paper claims none for it.** Cartwright [1999] supplies the discipline that makes a partial answer respectable: state the arrangement within which the law is claimed to hold. **The arrangement claimed here is a normalised exposure vector partitioned into disjoint perimeters, and nothing outside it.** Kirman [1992] is the objection that is not met: one Q, one δ and one C are written for a heterogeneous structure, and a representative parameter for a heterogeneous group generally does not exist.

12.5 Analogy discipline

borrowing	source	category	justification
Hill numbers, order 2	Hill [1973], Jost [2006]	**1 — derivation**	HHI^{-1} *is* ${}^2D$; an identity
numbers-equivalent reading	Adelman [1969]	**1 — derivation**	elementary and provable in two lines
effective number of parties	Laakso and Taagepera [1979]	**1 — derivation**	the same formula on a different share basis
portfolio variance	Markowitz [1952]	**1 — derivation**	the exposure vector genuinely is a weighted sum of shocks with a correlation matrix
$\ln N$ as entropy	Shannon [1948]	**2 — shared form**	**Shannon entropy in nats**, dimensionless, and not thermodynamic entropy
the surplus term over **funding** perimeters	Curzi [2026a]	**2 — shared form**	the storage law was specified over protected capital; **the concentration arithmetic transfers and the surplus figures do not**

borrowing	source	category	justification
currency drag as an erosion term	Curzi [2026a], Curzi [2026s]	**2 — shared form**	a modelling choice, not a derivation: a depreciating unit and a decaying stock share a functional form
threshold and tipping	Granovetter [1978], Schelling [1971]	**2 — shared form**	the perimeter collapse has the shape of a threshold model; no threshold distribution is estimated
frozen regions	Castellano et al. [2000]	**2 — shared form**	a perimeter is structurally a frozen region; no interaction dynamics are claimed
coupled-network fragility	Buldyrev et al. [2010]	**2 — shared form**	lenders have not been shown to satisfy the percolation premises; only the structural property is used
normal accidents	Perrow [1984]	**3 — framing**	qualitative organisational theory, never a formal result

The failure mode this table exists to prevent is letting a category-2 borrowing carry a category-1 result's authority, and the sharpest instance in this paper is the funding-side surplus figure, which the source itself flags and which Section 9.4 flags again.

12.6 Model comparison, and why none was performed

Nothing here is fitted: no estimated parameter, no held-out sample, no information criterion, no matched null, because there is no data. **A model with more parameters that fits better has not been shown to be better, and a model with no data has not been shown to be anything.** Durlauf [2005] asks for the prediction only this model makes, and this paper makes none: **any correlation-aware concentration measure and any cost-aware decision rule would produce the same results.**

12.7 Postulates

P	statement	status
P1	exposure shares are non-negative and sum to 1	DEFINITION

P	statement	status
P2	HHI $= \sum_i s_i^2$ and the effective count is its reciprocal	DEFINITION (CCP-052)
P3	units inside one coordination perimeter act on one decision	POSTULATED (the perimeter hypothesis of CCP-052)
P4	the surplus law is $h = \eta \ln N - \delta Q - C$	DEFINITION (modelling choice, CCP-A02, CCP-007)
P5	$\eta = 0.35$, $\delta_0 = 0.02$, $Q = 10$, $C = 0.30$	POSTULATED (fixed in advance)
P6	a hard-currency drag enters as $\delta_{\text{eff}} = \delta_0 + dw$	POSTULATED (modelling choice)
P7	receipts arrive uniformly and dwell for a common time n	POSTULATED (modelling choice)
P8	$r_{\text{hard}} = (1 + r_{\text{nom}})/(1 + \Delta) - 1$, with Δ the rise in local units per hard unit	DEFINITION (convention, stated here)
P9	the incremental maintenance of the regional leg is 0.08, or 0.04 for the holding layer alone	POSTULATED (constructed)
P10	the surplus term may be read over funding perimeters	POSTULATED (extension by analogy; **not a derivation**)

12.8 The chain that establishes the allocation and dwell results

step	consumes	statement	status
L1	P1, P2	bloc-HHI $= 0.160000 + 0.090000 + 0.090000 = 0.340000$	ARITHMETIC
L2	L1	$N = 2.941176$ and $\ln N = 1.078810$	ARITHMETIC
L3	L2, P4, P5	$\eta \ln N = 0.377583$ and $h(0) = -0.122417$	ARITHMETIC
L4	P6	$h(w) = h(0) - dQw$, linear with slope $-dQ$	ARITHMETIC
L5	L4, P5	at $d = 0.15$ and $Q = 10$ the slope is -1.5	ARITHMETIC
R1	L3, L5	$w^* = 0.377583/1.5 = 0.251722$ gives $h = -0.500000$; $w = 0.10$ costs 0.150000, 39.73% of the gain	ARITHMETIC
R2	P7	erosion per pass $= dn/365$: 0.767% at $d = 0.40$ and $n = 7$; 19.726% at $n = 180$	ARITHMETIC
R5	P8	$1.50/1.60 - 1 = -6.25\%$; $0.85^5 - 1 = -55.63\%$; $0.85^{10} - 1 = -80.31\%$	ARITHMETIC

12.9 The chain that establishes the two failed hypotheses

step	consumes	statement	status
M1	P1, P2, P3	diluting by 0.92 and adding an 8.0% leg to the bridge row gives bloc-HHI 0.353056	ARITHMETIC
M2	M1	$N = 2.832412$, below the baseline 2.941176	ARITHMETIC — **H3 falsified**
M3	P1, P2, P3	on the half-Western scoring, bloc-HHI 0.342496 and $N = 2.919742$	ARITHMETIC — **H3 falsified again**
M4	P1, P2	the regional reallocation gives bloc-HHI 0.274800 on Partition B, so $N = 3.639010$	ARITHMETIC
M5	M4, P4, P5	$\eta \, \Delta \ln N = 0.35 \times 0.218656 = 0.076530$	ARITHMETIC
M6	M5, P9	at $\Delta C = 0.08 > 0.076530$, $h = -0.127902$ against a baseline of -0.124431	ARITHMETIC — **H4 falsified**
R3	M2, M3	both scorings of the operating leg fall below the baseline count	ARITHMETIC
R4	M5, M6, P9	the ceiling is 0.076530, 25.51% of C; at $\Delta C = 0.04$, $h = -0.087902$	ARITHMETIC

12.10 The chain that establishes the funding result

step	consumes	statement	status
K1	P1, P2	four lenders at 40/25/20/15 give HHI $= 0.285000$ and 3.508772 effective lenders	ARITHMETIC
K2	K1, P3	one perimeter gives HHI $= 1.000000$ and $N = 1.000000$ exactly	ARITHMETIC
K3	P1, P2	four tranches at 45/25/20/10 give HHI $= 0.315000$ and 3.174603	ARITHMETIC

step	consumes	statement	status
K4	K1, K3	the naive count **falls** $3.508772 \to 3.174603$ while the perimeter count **rises** $1.000000 \to 3.174603$	ARITHMETIC
K5	P1, P2, P3	collapsing two tranches gives HHI $= 0.415000$ and $N = 2.409639$	ARITHMETIC
R6	K4, K5, P10	effective funding blocs lie between 2.409639 and 3.174603, with $\eta \ln N$ of 0.307817 and 0.404314	**ARITHMETIC** for the counts, **POSTULATED** for h

R6's two halves have different statuses and the row says so. The counts are exact; the surplus figures rest on P10, which is an extension by analogy.

12.11 Cross-checks performed

- Every exposure vector was summed: the diluted seven-element vector totals 1.000000 exactly, as do all four others.
- $s_{\max}^2 \le \text{HHI} \le s_{\max}$ holds on every vector: $0.084681 \le 0.189002 \le 0.291000$ on the baseline and $0.160000 \le 0.285000 \le 0.400000$ on the as-found funding book.
- The effective-count ceiling $s_{\max}^{-2}$ is respected in every case: $3.508772 \le 6.2500$ on the funding book and $5.290949 \le 11.8090$ on the baseline.
- w^* was obtained both by solving $h(w) = -0.5$ and by $\eta \ln N/(dQ)$, and the two agreed to twelve decimal places.
- Every compounding cell was recomputed as $(1-d)^k - 1$; eleven of twelve reproduced and the twelfth is recorded in Section 12.12.
- Every perimeter HHI was recomputed by summing shares within perimeters before squaring, and independently as an explicit double sum $\sum_i \sum_j s_i s_j \rho_{ij}$ with ρ the block matrix of ones; the two agreed to twelve decimal places.

12.12 Discrepancies found in the source, printed

Six discrepancies were found between the source manuscript and a recomputation from its own stated inputs. They are printed here rather than silently corrected.

1. **The source uses two incompatible bloc partitions of the same six-jurisdiction exposure vector.** One chapter measures BRIDGE 40.0 / NONALIGNED 30.0 / WEST 30.0, giving $N = 2.941176$ and $h = -0.122417$; another measures Western 30.0 / Gulf 29.1 / Asian 40.9, giving $N = 2.924302$ and $h = -0.124431$. **The difference is one jurisdiction's assignment and it moves the balance by** 0.002014**.** Both are reported in Section 5. **Partition A is treated as canonical here because it is the one derived from a stated perimeter test with recorded evidence classes; the second partition is not defended anywhere in the source.**
2. **The regional leg's gain is partition-dependent and only one value is published.** On Partition B the gain is 0.076530 and the ceiling 25.51% of C; on the canonical Partition A the same reallocation gives bloc-HHI 0.276600, $N = 3.615329$, a gain of 0.072231 and a ceiling of 24.08%. **The source computes only the first and does not say the result depends on the partition.**

3. **A compounding cell is wrong.** The three-year cell at a 15 per cent annual drag is printed as -38.61%; $0.85^3 - 1 = -0.385875$, so the correct figure is -38.59%. **Every other cell in the table reproduces exactly.**
4. **An "exactly" that is not exact.** The source states that a 25.17 per cent allocation sets h to exactly -0.500000. The exact allocation is $w^* = 0.377583/1.5 = 0.2517220$; at $w = 0.2517$ the balance is -0.499967. **The claim is right to four figures, and the word "exactly" belongs to 0.251722 rather than to 0.2517.**
5. **The nominal-to-real convention is used and never named.** The source's table is consistent with $(1+r)/(1+\Delta) - 1$ and inconsistent with $(1+r)(1-\Delta) - 1$, and the two readings of the same words differ by 33.75 percentage points on the last row. **This paper names the convention in Section 3.4. This is not an arithmetic error in the source; it is an ambiguity a reader could resolve the wrong way, and the wrong resolution is much the more pessimistic of the two.**
6. **A sixth-decimal truncation.** The regional leg's gain is printed in the source as 0.076529; 0.35×0.218656059 is $0.0765296208\ldots$, which rounds to 0.076530 at six decimal places. **The published figure is a truncation rather than a rounding**, and any recomputation will disagree in the sixth decimal. This paper uses 0.076530 throughout.

None of the six changes any hypothesis's outcome. All six are the kind of thing a reviewer checks first.

13. Limitations

No dataset was collected and nothing here is a measurement of a real structure. Every share, drag rate, dwell figure, cost assumption, duty rate and parameter is constructed and labelled as such at each point of use. **This is the corpus's central limitation and it is stated in full**: across the corpus, of 27 parameter-by-scope cells audited in a companion paper, 19 have a defined estimator, 1 has a validated one, and **0 carry a first-party measurement.** This paper carries none either.

No concentration figure here has an as-of date, because none is computed from real holdings. A concentration measure without a source and a date is not a measurement. Every share vector is constructed, so there is no source and no date to give. **Any application of this instrument to a real structure — of jurisdictions or of lenders — must carry both**, and Appendix B makes the date an item of the reporting standard.

No currency was measured. The 15, 25 and 40 per cent drag rates are stipulated illustrations. **The paper computes what a given drag costs and does not assert that any currency has that drag**, and every conclusion in Section 6 is conditional on a number the paper did not measure.

The erosion-term treatment of currency drag is a modelling choice and not a derivation. It assumes proportionality, uniform application over the period, and independence from the rest of the structure. **A drag that is itself a function of the size of the allocation — as it would be for a holder large relative to the market — breaks the linearity that produces w^*.**

The funding-side surplus figures are an extension by analogy. Section 9.4 states it, and it is repeated here because it is the most likely over-reading of this paper.

Every parameter in the surplus law was chosen, not estimated, and the sensitivity to one of them decides the sufficiency conclusions: at $\eta = 0.35$ the break-even count is 4.172734 and no structure here clears it; at $\eta = 0.50$ it is 2.718282 and the designed structure clears it comfortably.

The bloc partition is not settled, and the paper's own results move with it. Section 5 prints both partitions and Section 8.1 prints both gains. **The instrument is exact given a partition and the partition is a judgement**, which is the honest statement of what a perimeter measurement is.

The two maintenance-cost assumptions carry the whole of Section 8. 0.08 and 0.04 are stipulations. H4's verdict reverses between them, so **the paper's headline failure is a failure at an assumed price and not at a measured one.** That is stated in Results and repeated here.

Jurisdiction-specific legal claims are not cited to primary sources. The descriptions in Section 4 of codal families, foreign-exchange regimes, customs unions, treaty networks, substance regimes, statutory suites and the doctrinal content of Islamic finance are taken from the source manuscript. **The bibliography contains no statutes, regulations, court rules, registry documents, case reports or religious-legal sources for any individual jurisdiction or tradition.** Those claims are uncited here, described rather than asserted, and require verification before any reliance. Appendix A lists them. **This is a real gap and it is not closed here.**

One failure mode in Section 11 has no arithmetic attached. Failure mode 4 — a domestic contract that may not be denominated in the currency the group has chosen — is the one item in this paper the instrument cannot see at all. **It is recorded rather than measured, and a reader should not infer from the paper's quantitative character that everything important about these positions has been quantified.**

13.1 Rigour checks not met, listed by name

1. **No first-party measurement exists anywhere in this paper.**
2. **No uncertainty is reported anywhere**, because nothing is estimated.
3. **No as-of date is attached to any concentration figure.**
4. **The period of δ is supplied by this paper, not by the source**, which treats it as a bare rate. Section 12.1 assigns "per year" explicitly.
5. **Q, C and h have no operational unit.** The effective counts and the percentage ceilings survive; the magnitudes do not.
6. **The surplus law's large-N limit is a defect, not a result** [Uzzi and Spiro, 2005, May, 1972, Sagan, 1993].
7. **The conjoint conditions were not tested for the job vector** [Michell, 1997, Luce and Tukey, 1964], so the effective-jobs figure in Section 7.1 is provisional as a ratio-scale quantity; the ordering survives.
8. **Model comparison was not performed** — no out-of-sample test, no information criterion, no matched null — because nothing here is fitted.
9. **No distributional or scaling claim is made, so the Clauset–Shalizi–Newman procedure was not run** [Clauset et al., 2009]. **Any restatement reaching for a heavy-tail or power-law claim about lender shares or holdings must run it first**, and must read Stumpf and Porter [2012] on why passing it would still not identify a mechanism.
10. **The perimeter assignments are stipulated reference groups, not identified effects** [Manski, 1993, Brock and Durlauf, 2001, Shalizi and Thomas, 2011].
11. **The framework's own parameters are exempted from the reflexivity the framework asserts** [Lucas, 1976, MacKenzie and Millo, 2003]. **That is an inconsistency and it is not resolved here.**
12. **There is no conservation law for Q** [Mirowski, 1989], and Section 10.3 records that cost transfers between the acquisition and the structure do not leave a group-level quantity fixed — unlike the drive-conservation result of Curzi [2026], which does.
13. **Aggregation is unaddressed** [Kirman, 1992].
14. **Expected value and time average are not the same thing** [Peters, 2019]. Section 6.3 reports time averages of a multiplicative process, which is the right choice, and says so; Section 9's range is not an expectation at all but a pair of bounds under two assumptions.
15. **The surplus functional's form is underdetermined** [Latané, 1981, Briscoe et al., 2006] and nothing here discriminates it; Peters and Gell-Mann [2016] is named as the warrant the argument actually needs.
16. **No differential prediction is offered** [Durlauf, 2005].
17. **Small-sample bias in the effective-count estimators is not addressed** [Chao et al., 2014].
18. **The measurement is exposed to becoming a target** [Goodhart, 1975, Campbell, 1979]; the

only defence is Appendix B.

19. **The extension of the surplus term to funding perimeters is an analogy the paper uses and does not justify.** It is flagged in Method, in Results, in the analogy table and here — four times, because it is the claim a reader is most likely to carry away in a stronger form than it was made.

14. Conclusion

This paper establishes six results in the form the rest of the corpus cites.

R1. With a hard-currency drag written into the erosion term as $\delta_{\text{eff}} = \delta_0 + dw$, the surplus balance is linear in the local-currency allocation: $h(w) = h(0) - dQw$. At a 15 per cent annual drag, $h(w) = -0.122417 - 1.5w$, so an allocation of $w^* = 0.251722$ sets h to exactly -0.500000, cancelling the whole $\eta \ln N = 0.377583$ diversification gain, and a 10 per cent allocation costs 0.150000, which is 39.73% of that gain.

R2. Dwell decides whether a weak currency is a transaction cost or a capital loss. Erosion per pass is $dn/365$: at a seven-day dwell a 40 per cent annual drag costs 0.767%, and at one hundred and eighty days the same drag costs 19.726%. **A 15 per cent drag held six months costs nearly ten times what a 40 per cent drag held a week costs.**

R3. The bloc arithmetic against the operating leg's own case: an 8 per cent operating leg lowers the effective bloc count from 2.941176 to 2.832412 when its share is scored into the bridge row and to 2.919742 on the more generous scoring. **Both are below the baseline**, and the case therefore rests on customs-union access worth 1.35% of group revenue — 30 per cent EU-facing sales at a constructed 4.5 per cent average duty — which is a tariff argument and not a resilience one.

R4. The regional leg priced as a cost ceiling: the reallocation raises effective blocs from 2.924302 to 3.639010 and buys $\eta \Delta \ln N = 0.076530$, roughly 25.51% of the existing $C = 0.30$ base. At $C = 0.38$ the balance is -0.127902, **worse than the** -0.124431 **baseline**, and at $C = 0.34$ — the holding layer alone, correctly attributed — it is -0.087902. On the canonical partition of Curzi [2026] the same reallocation gives $N = 3.615329$, a gain of 0.072231 and a ceiling of 24.08%.

R5. Nominal against real, on the stated convention $r_{\text{hard}} = (1 + r_{\text{nom}})/(1 + \Delta) - 1$: a 50 per cent nominal return is a 6.25% loss when the exchange rate rises 60 per cent; and at a 15 per cent annual hard-currency drag a local-currency position loses 55.63% of its hard-currency value in five years and 80.31% in ten.

R6. Funding-source independence on the same instrument: four Western lenders give a naive HHI of 0.285000 and 3.508772 effective lenders but a perimeter HHI of 1.000000 and **exactly** 1.000000 **effective funding blocs**. The redesigned book gives 3.174603 effective funding blocs optimistically and 2.409639 after an honest correlation adjustment, with $\eta \ln N$ of 0.404314 against 0.307817 — **and the naive effective source count falls from** 3.508772 **to** 3.174603 **while it does so.** Per unit of effective bloc, funding diversity costs 0.070940 against 0.111930 for jurisdictional diversity.

Counter-result. Each measurement reports against the position it examines. The operating leg lowers the effective bloc count it was proposed to raise, on both available scorings. **The regional leg fails outright once its maintenance cost reaches** 0.38**, and H4's criterion — fixed in advance — was not met.** The Islamic funding gain is cut from 3.174603 to 2.409639 effective blocs by a correlation adjustment applied against the paper's own argument. **And the currency allocation, proposed as a yield enhancement, is the single largest destroyer of surplus measured anywhere in this paper: at** 25.1722% **it returns the structure to exactly the balance it had before any of the design work was done.**

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every case in this paper — the principal, the asset mix, the exposure shares, the lender and tranche shares, the drag rates, the dwell figures, the duty rate, the maintenance-cost assumptions and the job vector — is **CONSTRUCTED** and is labelled as such at each point of use. The numerical results are exact computations on those constructed inputs, carried at full precision and rounded only for printing.

All figures were produced by the scripts named in their captions (`figs/src/CCP-059_fig01.py` through `figs/src/CCP-059_fig05.py`), which import the shared figure style module and read their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

Jurisdictional, legal and doctrinal descriptions in this paper are taken from the source manuscript and not from primary sources. The bibliography contains general comparative and treaty material only; it contains no statutes, regulations, court rules, registry documents, case reports or religious-legal sources for any individual jurisdiction or tradition. Appendix A lists the claims affected. Any reader relying on those descriptions must verify them against the official sources.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. The principal, heirs, lenders and counterparties described in the constructed case are invented and correspond to no person or institution.

Preregistration

This paper reports no preregistered empirical test. The parameter set $\eta = 0.35$, $\delta_0 = 0.02$, $Q = 10$, $C = 0.30$, and H4's criterion that the regional leg is worth building only if its incremental maintenance cost falls below 0.076530, were all fixed before the measurements reported here. **H4's criterion was not met and Section 8.2 says so.** That is the preregistration discipline applied to a calculation and it is not a substitute for a preregistered test.

Religious and doctrinal content

Nothing in this paper is a ruling, an opinion, or a statement of religious law. Where Islamic finance is described, it is described as an institutional fact about a funding market — that it has distinct investors, documentation, regulators and supervisory organs, and therefore

constitutes a distinct funding perimeter for the purposes of a concentration measurement. **No doctrinal position is asserted, endorsed or evaluated**, and every doctrinal description is carried from the source manuscript, is uncited, and is listed in Appendix A.

Legal, tax and investment advice

This paper is analytical and is not legal, tax, or investment advice. It describes entity types, jurisdictions, currencies, funding instruments and structural arrangements for the purpose of measuring concentration, cost and resilience. No recommendation is made in respect of any particular person, asset, structure, currency, instrument or jurisdiction. Foreign-exchange regimes, customs-union scope and rules of origin, treaty networks and their protocols, substance requirements, promotion regimes, tax rates and enforcement arrangements all change and vary by jurisdiction and by activity; every such statement in this paper is carried from the source manuscript, is flagged in Appendix A, and requires confirmation with qualified counsel in the relevant jurisdiction before any reliance.

Appendix A. Claims requiring verification against primary sources

The following statements rest on legal, regulatory or doctrinal detail that changes, and **none of them is cited to the bibliography, which contains no primary sources for any individual jurisdiction or tradition**. They are reproduced from the source manuscript, are described rather than asserted, and require confirmation with counsel in the relevant jurisdiction before any reliance. **No measurement in this paper depends on any of them being true.**

A.1 The high-inflation industrial state. The civil code, code of obligations and commercial code relied on, with their numbers, dates and current texts, and the derivation from the Swiss codes. **The foreign-exchange decree and the communiqués made under it: the current scope of the restrictions on denominating or settling contracts in foreign currency between residents, the excepted categories, the settlement requirements, and the amendments in force — the load-bearing item, because a group that has decided to price everything in hard currency can find that a particular domestic contract may not be.** The inflation-accounting provisions of the tax procedure law and their effect on reported results and taxable bases. The scope of the customs union for industrial goods, the applicable tariff lines, and the rules of origin governing whether local processing is sufficient. The land registry law's conditions on foreign ownership. The finance centre's establishing instruments and incentives. The state's participation in automatic exchange of account information. The enforcement position for foreign judgments and for arbitral awards.

A.2 The regional holding centre and its operating markets. The mixed legal system's sources and the appellate route. The double-taxation network and the investment-protection network, and the material difference between them. **The protocol closing the historical route into the neighbouring market: its date, terms and grandfathering; any subsequent amending protocol; and the principal-purpose test's operation on a structure of the constructed case's shape.** The treaties terminated by African counterparties, and the dates. The regional harmonised business-law family's uniform acts and their scope. The substance reform replacing the deemed credit with a partial exemption, and the core income-generating activity conditions attaching to it. The standard-setters' listing decisions affecting the centre and their dates. The two other financial centres' establishing statutes, incentives and regulators. The continental free-trade and payment-system initiatives, and whether they supply an actual alternative path rather than an announced one. Withholding rates on dividends under the relevant treaty, and the domestic rates they are compared against.

A.3 The mid-shore centre and the Islamic-finance perimeter. The statutory suite for offshore companies and its current text. **The tax regime after the substance reform: qualifying activities, the substance conditions, and what happens to an entity that fails them.** The treaty

network's exclusions — the entities and income types excluded, which the source records as routinely unmentioned. The federal state's standards position. **The constitutional and statutory basis on which a national authority's rulings on Sharia compliance bind, and the scope of that binding force.** The contract set of Islamic finance, the two governing prohibitions, **the distinction between asset-backed and asset-based sukuk and its consequence for recovery**, and the published standards of the relevant standard-setters. The treatment of Sharia-compliance clauses by common-law courts and the authority for it. The succession rules of the tradition and their interaction with the forced-heirship analysis carried elsewhere in the series. The master agreement used for compliant hedging.

A.4 General. For all positions: corporate residence and central-management-and-control tests in every jurisdiction from which decisions might be taken; the current status of each state under the international anti-money-laundering standard; and the recognition and enforcement position for both judgments and awards in every jurisdiction where the group holds assets.

Appendix B. The reporting standard for an erosion-side and liability-side measurement

Any structure measured under the methods in this paper reports the following, and nothing less. It extends the reporting standard of Curzi [2026] to the erosion term and to the liability side.

```
ON THE EROSION SIDE
 1  the share of protected capital in each unit of account, summing to 100.0%
 2  the assumed annual hard-currency drag for each, with its source and date
 3  the convention used to convert nominal to real, stated explicitly
 4  the maximum balance and maximum average dwell, board-approved, in writing
 5  the measured average dwell for the period, against the stated maximum
 6  h(0), h(w), and w* = eta ln(N) / (d Q)

ON THE LIABILITY SIDE
 7  the funding share of each source, summing to 100.0%
 8  the perimeter assignment for each, with its evidence class
 9  naive HHI and effective sources; perimeter HHI and effective funding blocs
10  BOTH ends of the correlation range, never one alone
11  the incremental annual governance cost of each tranche

ON BOTH
12  the bloc or perimeter partition used, and the result under any
    defensible alternative partition
13  the date, and the diff against the previous measurement
```

Item 3 is the item this paper added to the standard, because two readings of the words "a sixty per cent depreciation" differ by 33.75 percentage points and the source manuscript never says which it means.

Item 10 is the one that will be dropped, because the flattering end of the range is the one that supports the decision already taken. **Item 12 is the one this paper's own Section 5 shows to be necessary**: two defensible partitions of one exposure vector gave 2.941176 and 2.924302, and no reader could have detected the difference from a published figure alone.

Item 13 decays first and matters most, because a perimeter is invisible except as a difference between two measurements taken a year apart.

Chapter 60

Compartments, Conversions, and the Confidence a Cell Wall Requires

Abstract. Three structural devices are priced against the certain costs they incur. Ring-fencing is measured with a contagion index — the sum of squared reachable shares, inverted — under which six unequal compartments (30.0, 22.0, 18.0, 12.0, 10.0 and 8.0 per cent) give 4.96 effective compartments and cut maximum single reach from 100.0 to 30.0 per cent; at an annual event probability of 2.0 per cent the benefit is 1.400 units against 0.110 of cost, a ratio of 12.73 to one and a break-even probability of 0.157 per cent. The paper's sharpest result concerns cells: choosing a cheaper cell company over separate portfolio companies is preferable only if the cell wall survives in the relevant foreign forum with probability at least 93.5714 per cent, a confidence the paper says almost nobody can honestly claim for an under-litigated device. Converting an immovable asset to listed securities incurs certain leakage of 7,220,000 on 48,000,000, or 15.0417 per cent, against a probabilistic escape value of 0.4725 per cent a year — a payback of 31.83 years, requiring belief in a 4.7752 per cent annual chance of an event destroying 60 per cent of the value before a ten-year payback holds. Diversification does no better: effective properties rise from 1.00 to 18.85, worth 0.4735 per cent a year and a payback of 31.77 years. The conversion also improves the naive jurisdiction count to 4.634 while worsening the effective bloc count to 1.470588.

Keywords: ring-fencing; cell companies; bankruptcy remoteness; contagion measurement; effective count; conflict of laws; leakage and payback

Series: D. Structure and Resilience **Paper:** CCP-060

1. Introduction

Three structural devices are priced here against the certain costs they incur: liability compartmentalisation, the statutory cell, and the conversion of an immovable asset into a movable claim. Each is widely recommended. None of the three is usually accompanied by the arithmetic that decides whether it is worth its price, and the arithmetic is short.

The organising question is not *what does this device protect against*, which every promotional account answers, but *what does it cost with certainty, against what benefit, and at what event probability does the sign of the difference change*. That question has a numerical answer whenever the benefit can be written as a reduction in a reachable share and the cost as an annual fee, and this paper writes both down.

Every structure discussed here is measured under one standing constraint, which the comparative literature on cross-border holding structures does not usually impose: the formation test. A structure qualifies for analysis only if it would still perform its function were every element of it disclosed to every authority with jurisdiction over any part of it. That test is not an ethical flourish; it is a design filter, and it is the reason concealment plays no part in what follows. It has an exact precedent in security engineering — Kerckhoffs [1883]'s rule that a system must remain secure when everything about it except the key is public, restated as the open-design principle by Saltzer and Schroeder [1975]. The regulatory environment has in any case removed non-disclosure as an available design parameter: automatic exchange of account information [Organisation for Economic Co-operation and Development, 2017], the common supervisory recommendations [Financial Action Task Force, 2012], and the substance and treaty-abuse programme [Organisation for Economic Co-operation and Development, 2015] between them make the disclosed case the only case worth modelling.

1.1 What this paper measures, and what it only describes

This paper is a **survey with embedded measurements**, and it says so at the front rather than letting a reader discover it. Roughly half of the source material is comparative description of particular jurisdictions and of particular legal devices. That description is context. It is retained — condensed in Section 7 and in full in Appendix C — because the measurements are unintelligible without it, and it is not offered as this paper's contribution. **The contribution is six computations, and Sections 8 and 13 are about those and nothing else.**

The computations are of three kinds.

First, a concentration measurement on an axis the concentration literature does not usually run it on. The instrument is the Herfindahl–Hirschman index [Hirschman, 1945, Herfindahl, 1950] read as a numbers-equivalent [Adelman, 1969], applied not to jurisdictions but to the **reachable share** of an estate: the fraction that a single creditor, claimant or administrative event arising against one compartment can reach. Six unequal compartments give 4.96 effective compartments and cut maximum single reach from 100.0 to 30.0 per cent. The instrument is carried unchanged from Curzi [2026], which establishes it and its bloc correction; what is new here is the share basis.

Second, a threshold belief in place of a legal opinion. Whether a foreign forum will respect a statutory cell wall is a conflict-of-laws question with thin authority and no reliable answer. It need not be answered. It need only be compared against a number, and the number is computable: on the constructed parameters, the cell company is the better structure only if that forum respects the wall with probability at least 93.5714 per cent.

Third, a certain cost against a probabilistic benefit. Converting an immovable asset into a listed security incurs leakage of 15.0417 per cent of gross value, immediately and with certainty, against an escape value of 0.4725 per cent a year that exists only in expectation. The payback is 31.83 years. **The asymmetry between a certain cost and a probabilistic benefit is the finding, and it survives every reasonable change to the parameters.**

1.2 What follows, and what does not

Two of the three devices contribute **exactly zero** to the effective bloc count that governs the surplus term of the storage law. That is stated as a result and not as a concession. Compartmentalisation raises maintenance cost and leaves the coordination perimeter untouched; so does the conversion. Both are answering a different threat, and a structure that reports their benefit alongside a jurisdictional improvement obtained elsewhere has reported a diversification it did not achieve.

The paper presupposes the surplus functional and the affine effective count of Curzi [2026] — $f(u) = \ln(1 + (N-1)u)$, characterised by angle-only dependence, vanishing at parallel, an affine effective count and normalisation. **That affine count is not the count used here, and Section 4.2 states the difference exactly**, because the two travel under one name in the corpus

and are different quantities. It presupposes the two-source decomposition of non-parallelism in Curzi [2026d], $u = \beta^2 + \alpha^2 \sin^2 \gamma$ with $u_{\text{div}} \leq 1 - u_{\text{cross}}$, which is why a compartment boundary cannot manufacture independence that the underlying exposure does not have. It derives the bloc-collapse arithmetic from the Crown Dependency measurement of Curzi [2026] — three jurisdictions at 40.0 / 35.0 / 25.0 giving naive HHI 0.345000 and 2.898551 effective jurisdictions against bloc-HHI 1.000000, $N = 1.00$ and a surplus term of exactly zero — from the six-entity Atlantic cluster of Curzi [2026], whose naive 4.0032 collapses to $N = 1.00$, and from Curzi [2026], where adding two jurisdictions raises the naive count 34.8 per cent while the bloc count falls 3.4 per cent.

2. Hypotheses

Four hypotheses are stated. Each is falsifiable in the sense Platt [1964] requires; Meehl [1967] is the reason each carries a point prediction with a constant attached rather than a direction. Statements that could not be falsified as posed are listed afterwards as definitions or modelling choices and are not counted.

The survey material advances no hypothesis, and none is manufactured for it. The comparative description of territorial-taxation jurisdictions in Section 7 and Appendix C is expository. Its organising question is: *when a jurisdiction sits inside a coordination perimeter the structure already occupies, through which term of the storage law can it act at all?* That question is answered arithmetically in Section 8.5 and the answer is a measurement, but the description supporting it is a survey and is labelled one.

H1 — the contagion hypothesis. The reachable-share index measures compartment independence in the same sense that the jurisdictional index measures perimeter independence, so that unequal compartments return an effective count strictly below their nominal number. *Predicted:* six compartments at 30 / 22 / 18 / 12 / 10 / 8 per cent return $1/\text{HHI}_c = 4.9603$, not 6, and maximum single reach falls from 100.0 to 30.0 per cent. *Falsified by:* an estate in which subdividing a compartment fails to lower the index, or in which the index reports a count above the nominal number of compartments.

H2 — the cell-indifference hypothesis. The choice between a statutory cell structure and separate companies reduces to a single threshold belief about foreign recognition, computable from the cost saving and the reach reduction alone. *Criterion, fixed in advance:* with a saving of 0.090 units a year, a reach reduction of 70.0 units and an annual event probability $p = 0.02$, indifference is at $1.400(1 - r) = 0.090$, so $r^* = 93.5714\%$. *Falsified by:* a cost structure in which the required confidence falls as the event probability rises — that is, in which the cheap compartment becomes safer as the risk grows.

H3 — the payback hypothesis. The conversion of an immovable asset into a movable claim does not repay its tax leakage within a holding lifetime at median parameters. *Criterion, fixed in advance:* with leakage $L = 15.0417\%$ and escape value $V_{\text{esc}} = pL^*q(1-s)$ at $p = 1.50\%$, $L^* = 60\%$, $q = 0.70$, $s = 0.25$, payback exceeds thirty years. *Falsified by:* a payback inside thirty years at those parameters, or by a defensible reference class in which p for the asset in question exceeds 4.7752 per cent a year.

H4 — the axis-divergence hypothesis. A single transaction can improve one concentration axis and worsen another, so that no single effective count summarises it. *Predicted:* the conversion raises effective properties from 1.0000 to 18.8491 while the resulting allocation reads 4.6339 naive effective jurisdictions and only 1.470588 effective blocs. *Falsified by:* the two axes moving together on this transaction, which would mean one of them is redundant.

Not hypotheses, and labelled as such.

- *The identification of N with the effective bloc count* and of $1/\text{HHI}_c$ with an effective compartment count is a **definition**. All the content is in the share basis and the bloc assignment rule, both of which are testable; none is in the identification.

- *The logarithmic input term* $\eta \ln N$ rather than ηN is a **modelling choice** carried from Curzi [2026] and is not tested here. Section 10 records that the form is underdetermined by the available evidence.
- *The parameter set* $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ is a **modelling choice**, fixed before the measurements so that no criterion is adjusted afterwards — the preregistration discipline of Nosek et al. [2018] applied to an engineering calculation.
- *Every event probability* p and every parameter of the escape-value model is a **stipulated constructed estimate**. Section 10 states that they are the most load-bearing and least defensible inputs in the paper.

3. Method

3.1 The share basis, which is the whole exercise

An index without a stated share basis is not a measurement. Three bases appear here and they must never be mixed inside one computation.

axis	share basis	what a failure on it looks like
jurisdictional	fraction of protected capital whose availability depends on a decision taken inside one coordination perimeter	a designation, a freeze, an exchange control
contagion	fraction of the estate reachable by a single creditor, claimant or administrative event arising against one compartment	a judgment, an insolvency, a regulatory action
property	fraction of an asset allocation's value in one underlying asset	a tenant default, a fire, a local planning decision

The contagion basis is the one this paper adds. It is **not** the fraction of value in each entity, which is the same number only when the compartments are perfect; it is not the fraction of revenue. The distinction is operative: on the carried case a manufacturing operation held as a branch contributes its parent's entire balance sheet to the reachable share, and the corporate chart does not show it.

3.2 The instrument

For shares s_i summing to one,

$$\mathrm{HHI} = \sum_i s_i^2, \qquad \text{effective count} = \frac{1}{\mathrm{HHI}}.$$

If k units each hold $1/k$ then $\mathrm{HHI} = 1/k$ and the reciprocal returns k; when the units are unequal it returns the number of equal-sized units that would give the same index. That numbers-equivalent reading is Adelman [1969]'s. Competition practice states the index in points on percentage shares [U.S. Department of Justice and Federal Trade Commission, 2023]; the resilience application states it as a fraction, and the two differ by a factor of 10^4.

Two elementary properties do work later. $\mathrm{HHI} \geq s_{\max}^2$, so the effective count is at most $s_{\max}^{-2}$ — which is why a compartment holding 82.0 per cent of an estate caps the effective compartment count at 1.4878 whatever is done with the remainder. And $s_{\max}^2 \leq \mathrm{HHI} \leq s_{\max}$, which is a two-line sanity check on any measurement done by hand.

No distributional or scaling claim is made anywhere in this paper, and that is stated because the material invites one. A reachable-share vector descending 30 / 22 / 18 / 12 / 10 / 8, or a property allocation with a long tail of thirty equal small holdings, looks like the opening of a heavy-tail argument. **No distribution is fitted, no tail is characterised and no power law is asserted.** Clauset et al. [2009] is the reason: a principled maximum-likelihood fit with a goodness-of-fit test and a likelihood-ratio comparison against alternatives leaves few canonical data sets standing, and a straight line on log-log axes is not evidence. That procedure was not run, because no claim requiring it is made. Stumpf and Porter [2012] adds why running it would not be enough on its own: many mechanisms produce the same tail. **Every effective count here is a statistic of a finite share vector with between two and forty entries, not an estimate of a distribution, and it must not be restated as one.**

3.3 The correlation correction, and what a compartment boundary cannot do

The naive index assumes the units fail independently. Where they do not, the correction is the one Markowitz [1952] stated for portfolios in 1952: diversification benefit depends on covariance, not on count. Writing w for the share vector and R for the correlation matrix of unit-variance shocks, the aggregate variance is $w'Rw$ and the general effective count is $1/(w'Rw)$; $R = I$ recovers the naive index and a block matrix of ones recovers the bloc index. Curzi [2026] establishes that identity and its correspondence with the similarity-sensitive diversity of order two [Leinster and Cobbold, 2012].

Applied to compartments the correction is what failure mode 5 of Section 9 is about. Six compartments each carrying a cross-default to any group indebtedness above a threshold have a legal correlation matrix close to the identity and a financing correlation matrix that is a single block of ones. The contagion index reads 0.201600 on the corporate chart and 1.000000 on the financing documents, and only the second is real. That coupled structures are far more fragile than their components is the formal content of Buldyrev et al. [2010], used here as a **shared structural form and not as a derivation**: no percolation premise has been shown to hold for a set of financing covenants. Perrow [1984] reached the organisational version first — interactive complexity plus tight coupling makes catastrophic failure a property of the system rather than of any component — and is cited as framing rather than as a formal result. Doyle et al. [2005] is the standing correction to reading too much into any of it: engineered robustness comes from design, not from a degree distribution.

Correlation rises under stress, which is when the count matters. Longin and Solnik [2001] measure extreme correlation in equity markets that is asymmetric and concentrated in downturns. That is directly load-bearing in Section 6: the conversion exchanges an idiosyncratic illiquid asset for a correlated liquid one, and the correlation arrives in the coordinated stress event while the liquidity is least available. The ecological name for the same variable is synchrony [Loreau and de Mazancourt, 2008]. The contrary literature is not decorative: Doak et al. [1998] show that diversity–stability correlations can arise from statistical averaging alone with no mechanism, and May [1972] that increasing diversity and connectance destabilises randomly assembled systems. This paper's claim is narrower than the general diversity–stability claim and is not exposed to those results in the same way — it asserts only that the reciprocal of a variance counts independent shocks, which is arithmetic.

One count can be inflated at will and is therefore worthless unaudited. Douceur [2002] proves that in the absence of a trusted identity authority one party can present as many, so that a count of nominally distinct entities carries no information about the number of independent parties. **Twenty companies formed by one sponsor in one jurisdiction are the legal instance of that result**, and the reason the share basis rather than the entity count is the object of measurement here. The corpus carries an exact arithmetic instance of the same gap between a nominal count and a count of distinct realised classes: Curzi [2026] proves that a roster indexed $1 \dots N$ realises exactly $\lceil N/2 \rceil$ distinct classes under its folding rule — six at $N = 12$, from twelve nominal members. **A count of names and a count of distinct things are different quantities in both settings, and only one of them is a measurement.**

3.4 Which effective count, stated exactly

The corpus runs two quantities under the name *effective count* and they are not the same quantity.

The reciprocal index is a Hill number of order two. Hill [1973] defined the family ${}^{q}D = (\sum_i p_i^q)^{1/(1-q)}$, whose members return the number of equally-abundant types giving the observed value of the underlying index; at $q = 2$ that is HHI^{-1} exactly. Jost [2006] is the statement that entropies are not diversities and their numbers-equivalents are. The same arithmetic was reached independently in political science as the effective number of parties [Laakso and Taagepera, 1979]. The exponential of Shannon entropy [Shannon, 1948] is the same family at $q = 1$. $1/\mathrm{HHI}$ **and** $\exp(H)$ **are Hill numbers of order 2 and 1 and are fully grounded.**

The affine count is not a Hill number and has no external support. The corpus's other effective count, $N_{\mathrm{eff}} = 1 + (N-1)u$, is a linear interpolation between 1 and N. It is not a numbers-equivalent, it is not $(\sum p_i^q)^{1/(1-q)}$ for any q, and no literature in ecology, economics, political science or information theory contains it. It never consults the share vector at all. **Every effective count reported in this paper is** $1/\mathrm{HHI}$ **on a stated share basis, and the affine count is not used here.** Where downstream work cites this paper's counts, the distinction must survive the citation. Hurlbert [1971] is the discipline being imposed: the word "diversity" carries no content without its formula attached.

Small-sample estimation is a live exposure and is not addressed. Chao et al. [2014] show that low-order diversity estimators are biased on small samples; a two-, six- or seven-element share vector is a small sample, and Section 10 records this as an unclosed gap. Golosov [2010] documents that effective-number measures behave poorly in the dominated-unit regime, where one unit holds a very large share — **which is exactly the regime of the carried case as found, at 82.0 per cent in one compartment**, and the measurement there should be read for its direction and not for its third decimal.

3.5 The threshold-belief method

Where a decision turns on an unknowable probability, the honest form is not an estimate of that probability but the threshold at which the decision changes sign. Given a cost saving ΔC from a cheaper device, an annual event probability p, and a reach reduction ΔS that the cheaper device forfeits if it fails, the required confidence solves

$$p\,\Delta S\,(1 - r^{*}) = \Delta C, \qquad r^{*} = 1 - \frac{\Delta C}{p\,\Delta S}.$$

Two properties of this expression carry the argument. It is **decreasing in** ΔC — the more the cheap device saves, the less confidence it needs — and **increasing in** p: the riskier the setting, the higher the confidence the cheap device requires. The second is the opposite of the usual intuition about economising in risky settings, and it is the reason the method is worth stating separately from its application.

The move of replacing an unanswerable question by a single threshold comparison is used elsewhere in the corpus and is worth naming as a family rather than reinventing. Curzi [2026] establishes the equivalence $Q_{\infty} > 0$ if and only if $\eta F(u^{*}) > c^{*}$ if and only if $f > f_{\min}$, with $f_{\min} = k\,e^{2c/\eta}$: an entire sustainability question reduced to one comparison against a threshold. **The constant there,** $k = (9/8)^2/\varphi^2 = 0.483425732988414$**, is a *defined* constant — an exact ratio of exact quantities — which is what licenses its fifteen printed digits; attaching a unit of hertz to it is a separate and much stronger claim, which that paper badges as a postulate and on which this paper takes no position.** Since φ enters its definition, Markowsky [1992] is the standing corrective on golden-ratio claims and is cited here for that reason.

The same form prices a marginal structural decision as a break-even probability rather than as a geopolitical opinion; Curzi [2026] applies it to a bridging slot and Curzi [2026] to a pledge trigger, where the first margin call at $\mathrm{LTV}_0 = 50\%$ and maintenance $m = 35\%$ falls at a 23.076923 per cent drawdown rather than at 50 per cent.

3.6 The leakage and escape-value model

Conversion of an appreciated immovable asset into a movable claim is, in most systems, a disposal at market value. The certain cost is

$$L = \frac{(V - B)\, t_{\text{gain}} + V d + V f}{V}$$

for gross value V, base cost B, effective rate on the gain t_{gain}, transfer taxes d and transaction costs f. The uncertain benefit is the annual expected loss avoided by being able to exit before an adverse event at the asset's location crystallises:

$$V_{\text{esc}} = p\, L^{*}\, q\, (1 - s)$$

for annual event probability p, fractional loss L^{*} if the event crystallises, probability q of actually exiting given warning, and discount s suffered when selling into the news. **The structure of the second formula is the argument**: the gross exposure pL^{*} is captured only to the extent that the holder both acts and is not fully penalised for acting late, and both haircuts are routinely assumed away. Payback is L / V_{esc} in years.

V_{esc} is linear in p and L^{*} while L is not, so the economics are driven almost entirely by tail parameters — which are exactly the parameters that are unknowable and therefore assumed. The requirement the specification imposes is that they be written down. **The corpus's own failure taxonomy applies directly: Curzi [2026ac] distinguishes a real channel with insufficient gain from a wrong magnitude extrapolated from a real result, and the conversion is a class-B instrument sold with class-C arithmetic** — the mechanism is real and the magnitude claimed for it is not.

3.7 Identification, and the status of the perimeter assignment

The bloc assignments used in Sections 6 and 7 are inputs to the instrument, and they carry an identification problem that must be stated where they are used rather than in a footnote. Treating jurisdictions as members of a coordination perimeter whose members respond to one decision is a statement about a group influencing its members while the members constitute the group. **Manski [1993] proves that the endogenous effect in a linear-in-means reflexive specification is not identified from observational data**: it cannot be separated from contextual effects or from correlated unobservables however much data is collected. No amount of observation of jurisdictional behaviour could therefore confirm a perimeter assignment derived reflexively.

The remedy adopted here is the fourth of the four available ones: the reference group is stipulated in advance and the stipulation is defended. Perimeter membership is fixed by the coordination test — two jurisdictions are in one perimeter when a single decision taken in one reliably produces the corresponding effect in the other without a further independent decision — recorded with its evidence, and stress-tested by running the arithmetic both ways in Section 8.5. Brock and Durlauf [2001] shows that route can be made rigorous. **The price is that every perimeter assignment in this paper is a modelling position and not an identified empirical result**, and it is described as such wherever it enters.

The collapse mechanism itself has a formal ancestry and is not new machinery here. That several units inside one perimeter behave as one under a correlated shock is the threshold-and-tipping result: Granovetter [1978] establishes that the distribution of individual thresholds rather than the mean determines whether a collective outcome tips, and Schelling [1971] is the canonical demonstration that a mild individual preference produces a sharp aggregate outcome — **the strongest available warning against reading an individual-level parameter straight into an aggregate, which is precisely what a perimeter assignment does.**

4. Compartments: the object being measured

4.1 Four jobs under one phrase

“Special purpose vehicle” covers four structurally different mechanisms with different failure modes, and conflating them is the origin of most bad design in this area. The irreducible function common to all four is asset partitioning, which Hansmann and Kraakman [2000] identify as the thing entity law does that contract cannot; Pistor [2019] is at once the best scholarly account of how legal coding creates capital and its sharpest critique.

job	design goal	characteristic failure
1 isolate assets from the sponsor’s creditors	the sponsor’s insolvency does not reach assets investors paid for	substantive consolidation, or a successful challenge to the true sale
2 isolate liabilities from the sponsor	obligations incurred by the vehicle stop at the vehicle	veil piercing; guarantees or comfort letters construed as obligations
3 isolate assets from each other	a claim arising on asset A cannot reach asset B	commingling, cross-collateralisation, cross-default
4 isolate economic from legal ownership	the vehicle consolidates with nobody	a side arrangement restoring the economic residual

Jobs 1 and 2 point in opposite directions and a vehicle optimised for one is usually badly designed for the other. **Job 3 is the one this paper measures**, and it is the one for which the professional literature offers least, because the securitisation market drove the technology and cared only about jobs 1 and 4.

A ring-fence is not created by forming an entity. It is created by running the entity as a separate thing for long enough that a court asked to disregard it has nothing to work with. Formation is the cheap part and takes a week; conduct is the expensive part and never stops. That asymmetry is why the cost side of Section 5 is an annual figure and not a one-off.

4.2 Separateness, and the three ways it fails

The foundation of the technique is that a duly incorporated company is a legal person distinct from its members even where one person holds substantially all its shares and is its principal creditor. **The doctrinal material in this subsection is described from the source manuscript and is deliberately left uncited: this paper’s bibliography contains no statutes, no regulations, no court rules, no registry documents and no case reports for any individual jurisdiction** (Section 12; the affected statements are itemised in Appendix B).

Three distinct attacks are routinely conflated.

Veil piercing, properly so called, treats the company’s act as the controller’s act. English authority has narrowed it to near-vanishing point by distinguishing a *concealment* principle — looking behind the company to identify the real actors, which requires no special doctrine — from an *evasion* principle, under which a person interposes a company specifically to defeat an obligation and the court may disregard separate personality to the extent necessary. **The practical consequence is favourable and should be stated plainly: a compartment formed for a lawful commercial purpose before any obligation arose, and run as a separate business, is very hard to pierce.** The doctrine that kills compartments is almost never veil piercing.

Alter ego and substantive consolidation is the broader cluster in practice. Its factors are all conduct: undercapitalisation, commingling, absence of formalities, common officers acting without distinguishing capacity, holding out. Consolidation standards differ between appellate circuits, and the applicable standard is a function of where the proceeding is opened.

Recharacterisation is the least discussed. A court may find that what the documents call a sale was in substance a secured loan, in which case the assets never left the originator's estate and the compartment was notional. The factors — recourse, retained servicing, retained residual, pricing, the parties' own accounting treatment — are again all conduct.

The common structure is that every one of the three is decided on evidence about behaviour rather than on the choice of vehicle. That is why the measurement in Section 5 is run on the financing documents and the connection register rather than on the corporate chart.

4.3 Cells, and the question the marketing does not raise

A segregated portfolio, protected cell or incorporated cell company achieves compartmentalisation inside one legal entity: one company, one board, one registration, with assets and liabilities attributed to named cells and a statute providing that the assets of one cell do not answer the creditors of another. The technique is a creature of one jurisdiction's statute and binds the courts of that jurisdiction.

Whether a court elsewhere will give effect to a statutory partition inside a single foreign legal person is a conflict-of-laws question, and it is materially less tested than the recognition of a separate company, which is close to universal. This paper declines to state that the point is settled, because it is not. The shape of the risk can be stated: a foreign court faced with a single legal person that has contracted with a local claimant, and a foreign statute saying that only part of that person's estate answers the claim, is being asked to apply a foreign limitation of liability to a domestic claim. The analysis will turn on characterisation, on the governing law of the obligation, and on local public policy — the standing apparatus of the conflict of laws [Torremans et al., 2022, Fawcett and Carruthers, 2008]. Where the question can be routed to arbitration, the enforcement position is materially better, because recognition of awards runs through an instrument with near-universal adherence [United Nations, 1958], whose actual refusal grounds are themselves a settled subject [van den Berg, 1981]. **Recognition of judgments is not general in the same way**, which is what the later portability convention [Hague Conference on Private International Law, 2019] exists to address and what makes the cell question live rather than academic.

Section 5.3 converts that unanswerable question into an answerable one, which is the contribution: the analyst need not know whether an unnamed forum will respect the partition in fifteen years, only whether the confidence is above or below a computed threshold — and, more usefully, whether every claim that could arise will arise inside the cell statute's own jurisdiction.

4.4 Orphans, and three concepts that are not the same thing

An orphan structure has no equity owner in the sponsor's group; the shares are held by a professional trustee on charitable or non-charitable purpose trusts. Its purpose is consolidation in two senses — accounting, where control rather than ownership determines the answer, and insolvency, where a vehicle nobody in the group owns is a harder consolidation target. **An orphan is a consolidation device and not a concealment device.** The trustee is a regulated professional who knows who arranged the transaction; the arrangement is disclosed in the offering documents and in the notes, and the beneficial-ownership registers capture it. It passes the formation test comfortably, and fails immediately if a side arrangement restores the economic residual. Whether the trust is recognised at all across a border is governed by a convention with a short ratification list [Hague Conference on Private International Law, 1985]; the comparative treatment of the trust against the civil-law foundation as alternative devices is a literature of its own [Panico, 2014].

Three concepts are routinely conflated at board level, with expensive consequences.

concept	question it answers	who decides	when tested
legal separateness	would a court treat the entity's assets and liabilities as its own?	company and insolvency law	in litigation
accounting consolidation	does the entity appear in the group's consolidated statements?	control-based accounting standards	by auditors, annually
bankruptcy remoteness	would the sponsor's insolvency reach the entity, and would the entity itself be filed?	insolvency law plus contract	rarely and expensively

An entity can be legally separate and fully consolidated for accounting, or deconsolidated for accounting and not remote at all. **Non-consolidation in the accounts is evidence of nothing about what an insolvency court will do**, and the reason the control criterion replaced the earlier tests is precisely that vehicles economically identical to their parent were achieving deconsolidation.

4.5 What makes a compartment hold

The checklist is short, dull, and almost entirely about behaviour: separate bank accounts without exception; separate books prepared contemporaneously; capitalisation adequate for the purpose rather than nominal; its own board taking real decisions; arm's-length intra-group dealing, documented and priced; correct holding out in contracts, invoices and signature blocks; a restricted purpose observed rather than merely recited; no guarantees, comfort letters or keep-wells unless deliberately given, priced and disclosed; no cross-collateralisation or cross-default between compartments unless deliberately accepted; and, where third-party money is carried, reasoned opinions on true sale, non-consolidation, enforceability and the governing law of each place a claim could arise.

Items eight and nine are where well-run groups fail, because they are given away by the treasury function and the lending bank rather than by the structuring lawyer. A group that spends heavily on compartmentalisation and then signs a group-wide cross-default has bought nothing and paid for it. The independent-director consent right that the market treats as the load-bearing element is a governance arrangement rather than a wall wherever the sponsor appoints, pays and can replace its holder; and whether a foreign court will enforce a contractual promise not to invoke its own insolvency jurisdiction is under-litigated exactly where the structure needs it most.

5. The contagion measurement

5.1 The instrument on a constructed estate

CONSTRUCTED CASE. Every figure in this section is invented for illustration and none is a benchmark. An estate of 100 units.

```
UNFENCED -- one entity holds everything
    reachable share            100.0%
    contagion HHI              1.000000
    effective compartments     1.0000
    maximum single reach       100.0%
```

```
FENCED -- six compartments, by reachable share
    A  30.0%   0.090000        D  12.0%   0.014400
    B  22.0%   0.048400        E  10.0%   0.010000
    C  18.0%   0.032400        F   8.0%   0.006400
    ---------------------------------------------------
    contagion HHI              0.201600
    effective compartments     4.9603
    maximum single reach        30.0%
```

Six entities, 4.9603 effective compartments, and a worst case falling from 100.0 to 30.0 per cent of the estate. The effective count is 4.96 rather than 6.00 because the compartments are unequal, and the quantity nobody quotes — maximum single reach — is the one an insolvency practitioner cares about.

5.2 Cost, benefit and the break-even probability

Compartments are not free and the fee is annual. Each is an entity that raises C in the storage law and contributes nothing to N, so under that law alone each is strictly value-destroying; the justification must come entirely from the contagion axis and must be computed. That adding units inside one coordination boundary adds communication cost at order $n(n-1)/2$ while adding nothing to independence is the general organisational form of the same point [Brooks, 1975].

```
single entity   registered agent, accounts, audit,
                director fees, filings           0.040 units p.a.
each SPV        same, per entity                 0.025 units p.a.
six SPVs                                         0.150 units p.a.
incremental cost of compartmentalisation         0.110 units p.a.

benefit at annual event probability p
    E[loss | unfenced] = p x 100.0
    E[loss | fenced]   = p x  30.0
    benefit            = p x  70.0

at p = 2.0%   benefit = 1.400 units p.a.
              cost    = 0.110 units p.a.
              ratio   = 12.73 to 1
break-even p = 0.110 / 70.0 = 0.157% per year
```

Below an annual event probability of about one in 636, compartmentalisation of this estate at this cost destroys value. For a passive portfolio held by a family holding company with no operations, no employees, no leases and no counterparties that is a live possibility, and the honest recommendation may be a single entity. For an operating group with plants, employees, environmental exposure, product liability and lease obligations, p is far above the threshold and the question does not arise. **The instrument does not say "compartmentalise"; it says "compute p", and p is a fact about what the business does rather than a fact about structuring.**

A structurally parallel result exists in a different domain and is worth naming, because it shows that the shape of this break-even is not an artefact of the legal setting. Curzi [2026g] derives a critical boundary-cost multiplier $m^{*} = \eta(N-1)u_{\text{cross}}/[(1+(N-1)u(\rho))\,c_{\text{in}} \times 900]$, above which adding contact across a boundary stops helping — computed at 3.17 against an operating value of 8. **In both settings a structural intervention is worth its cost only below a computable threshold, and in both the operating value can sit on the wrong side of it without anyone having run the comparison.**

5.3 The cell wall, and the confidence it requires

A segregated portfolio company achieves the same six compartments inside one entity at lower cost, and forfeits the reach reduction entirely in any forum that does not respect the partition.

```
SPC with six cells   one registration, one board, one audit,
                     cell-level records          0.060 units p.a.
six separate SPVs                                0.150 units p.a.
saving                                           0.090 units p.a.

additional expected loss from using cells
    = p x (1 - r) x (100.0 - 30.0)
    = 0.02 x (1 - r) x 70.0  =  1.400 x (1 - r)

indifference:  1.400 x (1 - r) = 0.090
               1 - r = 0.064286
                   r = 93.5714%
```

The cell company is the better structure only if the wall is respected in the forum where the claim arises with probability at least 93.5714 per cent. The threshold is high because the downside is catastrophic and the saving is small, and that asymmetry survives reasonable changes to the inputs: at $p = 1.0\%$ the required confidence is 87.1 per cent and at $p = 4.0\%$ it is 96.8 per cent. **Higher risk makes the cheap compartment worse rather than better**, which is the opposite of the usual intuition about economising in low-risk settings and is the general property of the threshold form derived in Section 3.5.

The practical value of the threshold is that it converts an unanswerable legal question into an answerable practical one. Nobody can know whether a court in an unnamed jurisdiction will respect a foreign statutory partition in fifteen years. Anyone can establish whether every claim that could arise will arise inside the cell statute's own jurisdiction. **If every counterparty contracts under that law, submits to those courts and holds no asset elsewhere, r is near one and the cell company is correct. If any material claim can arise elsewhere, r is a guess and the guess must clear 93.6 per cent.**

use cells	use separate companies
all counterparties contract under the cell statute's law	any material counterparty in a third jurisdiction
all assets held in or through the cell jurisdiction	hard assets sited elsewhere
regulated fund or insurance use where the regulator supervises the cell structure	trading, employment or environmental exposure
compartments are economic sleeves of one portfolio	compartments must survive a hostile foreign creditor

5.4 The sizing rule: a constraint, not an optimum

The number of compartments is not chosen by optimisation. It is set by the largest loss the group can absorb without breaching a covenant or losing control of another asset:

$$\text{maximum single reach} \le \text{absorbable loss}.$$

On the constructed estate with an absorbable loss of 25 units, compartment A at 30 units breaches the constraint and must be split:

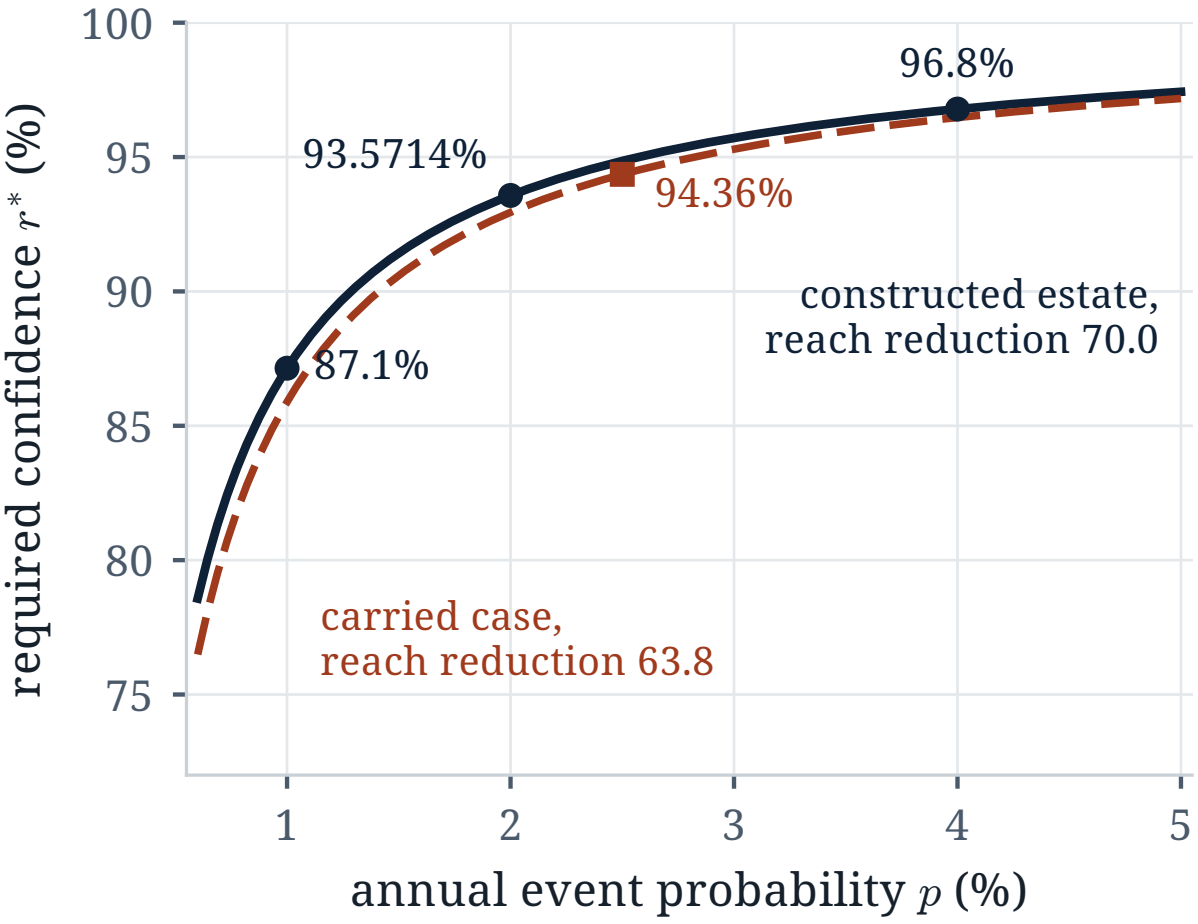


Figure 60.1: **The confidence a cell wall must command rises with the risk it is being used against.** Required confidence $r^* = 1 - \Delta C / (p\,\Delta S)$ at a saving of $\Delta C = 0.090$ units a year, for the constructed estate (reach reduction 70.0 units) and for the carried case (63.8 units). Marked points are the values printed in Sections 5.3 and 5.5; modelled from the constructed cost parameters, script `figs/src/CCP-060_fig02.py`.

```
A1 18.0  A2 12.0  B 22.0  C 18.0  D 12.0  E 10.0  F 8.0
contagion HHI = 0.158400
effective compartments = 6.3131
maximum single reach   = 22.0%
```

The effective count rises from 4.9603 to 6.3131 for one additional entity at 0.025 units a year. **The constraint sets the structure; the optimisation does not.** A group asking “what is the optimal number of vehicles” has asked a question with no answer; a group asking “what is the largest loss we can absorb” gets a number. Sagan [1993] is the standing caution against the opposite reflex: added redundancy can reduce reliability by creating new failure modes and diffusing responsibility, which is what failure mode 10 of Section 9 records in cost terms.

5.5 The carried case

CONSTRUCTED CASE. A group with Thai manufacturing, a trading company, intellectual property, London commercial real estate and a listed portfolio. As found, everything sits under two entities, with the Thai operation run as a branch of the UK company.

```
compartment       reachable share    contribution
------------------------------------------------------
UK company              82.0%        0.672400
Jersey company          18.0%        0.032400
------------------------------------------------------
contagion HHI = 0.704800
effective compartments = 1.4188
maximum single reach   = 82.0%
```

One product liability claim arising in Thailand reaches the London building, the securities portfolio and the trading business, because the Thai operation is a branch and a branch is the company. That finding took twenty minutes to compute from the corporate chart, and the corporate chart is where it was invisible.

The group’s covenants give a maximum absorbable single loss of 22.0 units on an estate of 100, so the position as found fails the constraint by a factor of 3.7. Rebuilding the compartments around **where claims arise** rather than around where value sits:

compartment	reachable share	why separate
1 manufacturing operating company	18.2	employment, product, environmental
2 trading company	16.5	trade counterparty claims
3 property holding vehicle	13.5	occupier, premises, lender
4 intellectual property holding company	11.8	infringement claims run both ways
5 securities vehicle	13.9	custodian and market claims
6 captive insurer	13.6	insurance liabilities, regulated
7 regulated-zone vehicle	12.5	separate regulator

```
contagion HHI          = 0.145940
effective compartments = 6.8521
maximum single reach   = 18.2%   (constraint 22.0 -- SATISFIED)
```

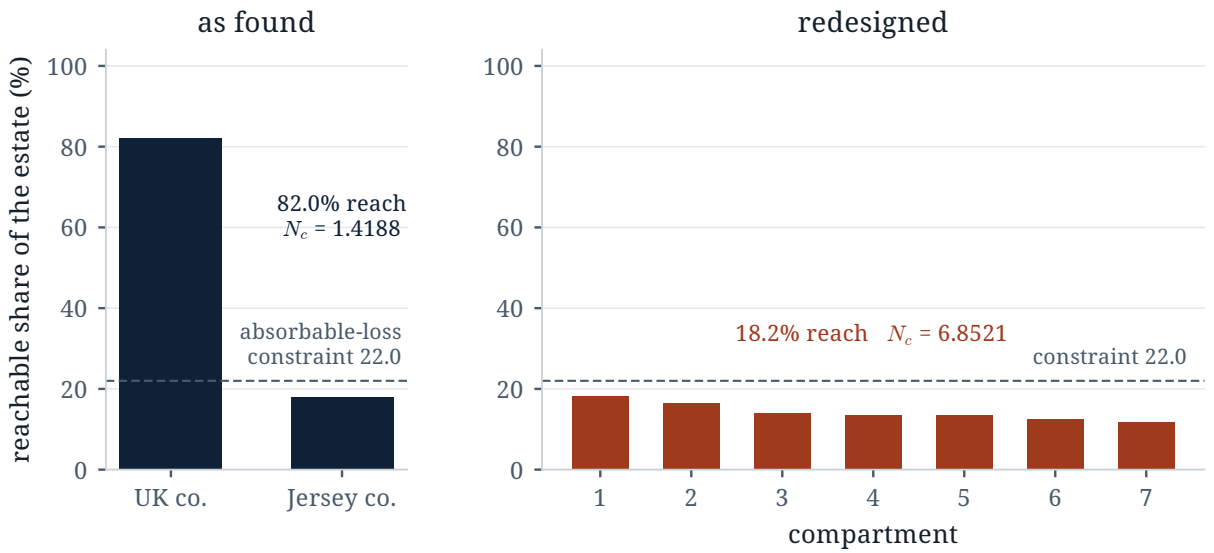


Figure 60.2: **Seven compartments built around where claims arise take maximum single reach from 82.0 to 18.2 per cent and the effective compartment count $N_c = 1/\mathrm{HHI}_c$ from 1.4188 to 6.8521.** Reachable shares and both indices are the constructed values computed in Section 5.5; the dashed line is the group's stated absorbable-loss constraint of 22.0 units. Script `figs/src/CCP-060_fig01.py`.

metric	as found	redesigned	change
entities	2	7	$+5$
contagion HHI	0.704800	0.145940	-79.3%
effective compartments	1.4188	6.8521	$\times 4.8295$
maximum single reach	82.0%	18.2%	-63.8 points
absorbable-loss constraint	FAIL	PASS	

The cost and the verdict:

```
incremental entities                          5
cost per entity per year (constructed)    0.025 units
incremental cost                          0.125 units p.a.
estimated p (operating group)             2.5% per year

benefit = 0.025 x (82.0 - 18.2) = 1.595 units p.a.
cost                            = 0.125 units p.a.
ratio                           = 12.76 to 1
break-even p = 0.125 / 63.8     = 0.196% per year
```

The redesign pays for itself at any event probability above one in 510, and the estimated probability is one in 40. What makes it clear is the manufacturing operation. Strip that out and model the same estate as a passive holding structure, and p falls to the order of 0.15 per cent — below break-even — and the honest recommendation becomes fewer entities rather than more.

The cell proposal for this group is then decided rather than argued. Claims against three of the seven compartments arise in forums that would have to be asked to recognise a foreign

statutory partition inside a single foreign company holding assets on their own territory. At the group's $p = 2.5\%$ and reach reduction of 63.8 units, indifference is at $1.595(1 - r) = 0.090$, giving $1 - r = 0.056426$ and $r^* = 94.36\%$. **The proposal is declined for those three compartments and is acceptable for purely financial sleeves inside the securities vehicle, where every counterparty contracts under one law and no hard asset sits anywhere.** That split answer is available only because the question was converted into a number.

A discrepancy in the source is recorded here rather than corrected silently. The source manuscript states that the required confidence is "95.9%" at $p = 2.5\%$ and then displays the arithmetic giving 94.36 per cent. The displayed arithmetic is right; 95.9 per cent is not reproducible from any combination of the chapter's stated parameters, and the nearest reproducible figure is 94.86 per cent, which is the threshold at $p = 2.5\%$ on the constructed estate's reach reduction of 70.0 units rather than the carried case's 63.8. **This paper prints 94.36 per cent and flags the source figure as an error.**

5.6 The register that decides the outcome

The compartment map is the easy part. The register of **connections** determines whether any of it is real, and on the carried case it produced four findings the corporate chart did not show: an uncapped parent guarantee of an equipment lease, renegotiated to a capped guarantee sized inside the absorbable-loss constraint; a cross-default in the property facility referencing any group indebtedness above a threshold, carved out for the operating subsidiaries at a fee; an ambiguously worded comfort letter, withdrawn and replaced with a specific, capped and disclosed support undertaking; and a shared cash pool sweeping five entities into one account, restructured into documented intercompany loans at arm's-length rates.

The fourth item was the whole structure. A cash pool is commingling with a treasury policy attached, and a group that has spent five years building compartments and sweeps them nightly into one account has produced a corporate chart rather than a ring-fence. Restructuring it cost more in working-capital efficiency than the six new entities combined, and it is the only item on the list without which the rest is decorative.

5.7 The honest zero

Run the jurisdictional instrument on the same group before and after:

```
BEFORE   two entities, both inside one perimeter
         naive HHI 0.704800 -> effective jurisdictions 1.4188
         bloc-HHI  1.000000 -> N = 1.0000, ln(N) = 0, h = -0.500000

AFTER    the compartments map onto the designed structure
         bloc-HHI  0.340000 -> N = 2.941176, h = -0.122417
```

That looks like a large improvement produced by compartmentalisation, and it is not. The improvement in N comes from the jurisdictions the compartments happen to sit in, all chosen for reasons that had nothing to do with compartmentalisation. Build the same seven compartments as seven companies in one jurisdiction:

```
contagion HHI = 0.145940   effective compartments = 6.8521
bloc-HHI      = 1.000000   N = 1.0000  ln(N) = 0  h = -0.500000
additional C  = 7 x 0.025  = 0.175 units p.a.
```

Identical contagion result, zero jurisdictional result, and 0.175 units a year of additional maintenance. Compartmentalisation contributes exactly zero to N. It is answering a different threat, and a structure that reports the two together is reporting a benefit obtained from somewhere else. The same arithmetic appears in Curzi [2026], where six Atlantic-cluster entities reading 4.0032 effective jurisdictions collapse to $N = 1.00$ and $\ln N = 0$, and in Curzi [2026], where three closely related jurisdictions reading 2.898551 collapse to exactly 1.00.

6. Conversion of an immovable asset into a movable claim

6.1 The mechanism, and why it is transparent by construction

Land cannot leave the jurisdiction it sits in and is therefore subject to that state absolutely; every legal system governs rights in immovable property by the law of the place where it is situated, and enforces that rule by controlling the register on which title depends. A security, by contrast, is an intangible whose location for conflicts purposes is settled by reference to the account in which it is held — the place-of-the-relevant-intermediary approach — so the governing law of the claim follows an account, and an account is a commercial arrangement that can be moved. **The conversion replaces an asset whose legal character is fixed by geography with a claim whose legal character is fixed by a contract the holder controls. That is a genuine change and not a trick.** The doctrinal statements in this paragraph are described from the source manuscript and are not cited; Appendix B itemises them.

The mechanism works because it is transparent, not despite it. A listed real estate vehicle files an annual report listing its properties asset by asset; the buildings are on a public land register in any case; and beneficial owners of overseas entities holding land are, in the leading jurisdiction for the carried case, on a public register too. What the disclosure buys is the property that makes the claim tradeable. A security trades because buyers can price it, buyers can price it because the issuer discloses, and the disclosure regime is enforced by a regulator whose credibility is the exchange's product. **Liquidity is manufactured out of disclosure and has no other input.** The formation test is therefore passed in its strongest available form: every element already is disclosed, to everyone, on a schedule, under penalty for inaccuracy — and the mechanism works anyway.

The land is still seizable by the host state even when the claim on it is portable. If a state expropriates the building, the holder owns a share in an entity that has lost a building. What has been bought is not immunity but **latency**: the ability to be out of the position in a day rather than in nine months. That has a computable value, and the computation is Section 6.3.

6.2 Leakage: the certain cost

CONSTRUCTED CASE. A single freehold commercial building: market value 48,000,000, base cost 22,000,000, unrealised gain 26,000,000, net operating income 2,304,000 at 4.80 per cent on value, debt 14,400,000 at 30.0 per cent loan-to-value, four tenants with the largest at 41 per cent of income, exit latency on a direct sale of approximately 270 days, and a property-level effective count of 1.0000.

```
V = 48,000,000     B = 22,000,000     gain = 26,000,000
t_gain = 25%       tax on the gain   =  6,500,000
d = 0              transfer tax not borne on the modelled route
f = 1.5% of V                        =    720,000
------------------------------------------------------------
total leakage                        =  7,220,000
L = 7,220,000 / 48,000,000           =  15.0417%
net proceeds available to reinvest    = 40,780,000
```

Fifteen per cent of gross value, certain, immediate and unrecoverable. That is the number every subsequent computation is measured against, and it is the number absent from the presentations of this idea that the source manuscript reports having seen.

One asymmetry between tax systems is decisive and is a fact about tax systems rather than about property. Where contribution of appreciated property to a partnership in exchange for a partnership interest is non-recognition, the vehicle can be built as an umbrella partnership and immediate leakage falls to transaction costs alone — 1.5 per cent on these numbers — and the base-case payback becomes $0.015 / 0.004725 = 3.17$ years. **But the deferred gain does**

not disappear: it is allocated back to the contributor as the partnership depreciates or sells, and it crystallises on exchange into the listed security. The deferral defers the very portability the conversion was undertaken to obtain. Where no such relief exists, the contribution is a disposal at market value and the conversion route is economically identical to an outright sale and reinvestment.

6.3 Escape value, payback, and the break-even belief

```
base case parameters, stated so they can be disputed
p  = 1.50%    a materially adverse event at the situs once in 67 years
L* = 60%      expropriation, punitive levy, or a use restriction
q  = 0.70     the event is seen coming and the board can act
s  = 0.25     sale at 25% below the pre-news mark

V_esc = 0.0150 x 0.60 x 0.70 x 0.75 = 0.004725  ->  0.4725% per year
payback = 0.150417 / 0.004725 = 31.83 years
```

Thirty-two years to recover the tax cost of the conversion at median parameters. The tail is the case that justifies it:

scenario	p	L^*	q	s	V_{esc}	payback (yr)
benign	0.50%	50%	0.70	0.25	0.001313	114.6
base	1.50%	60%	0.70	0.25	0.004725	31.8
elevated	3.00%	70%	0.75	0.30	0.011025	13.6
stressed	5.00%	80%	0.70	0.25	0.021000	7.2
acute	10.00%	85%	0.60	0.40	0.030600	4.9

Solving for the annual probability at which the conversion repays inside ten years, at $L^* = 60\%$, $q = 0.70$, $s = 0.25$:

```
p x 0.60 x 0.70 x 0.75 = 0.150417 / 10
p x 0.315 = 0.0150417
p = 4.7752%
```

A ten-year payback requires belief in roughly a 4.8 per cent annual chance of an event destroying 60 per cent of the value of this specific building. For prime commercial real estate in a jurisdiction with an essentially clean record on expropriation of commercial freeholds that belief is not defensible on the record. Where it *is* defensible, the conversion is not merely justified but urgent — and by then there is often no functioning market in which to execute it.

Run on the carried building, whose honest parameters are the benign row — $p = 0.50\%$, $L^* = 50\%$, $q = 0.70$, $s = 0.25$ — the escape value is 0.13125 per cent a year and the payback is 114.60 years. **On honest parameters for this asset in this jurisdiction the geopolitical case for the conversion does not exist.** The instrument is real; the carried case is not the fact pattern for which it is the answer.

6.4 The property axis, where the transaction succeeds

The arithmetic that fails on arbitrary state action succeeds decisively on an axis nobody runs. Take the share basis to be the fraction of the real estate allocation's value in one underlying asset.

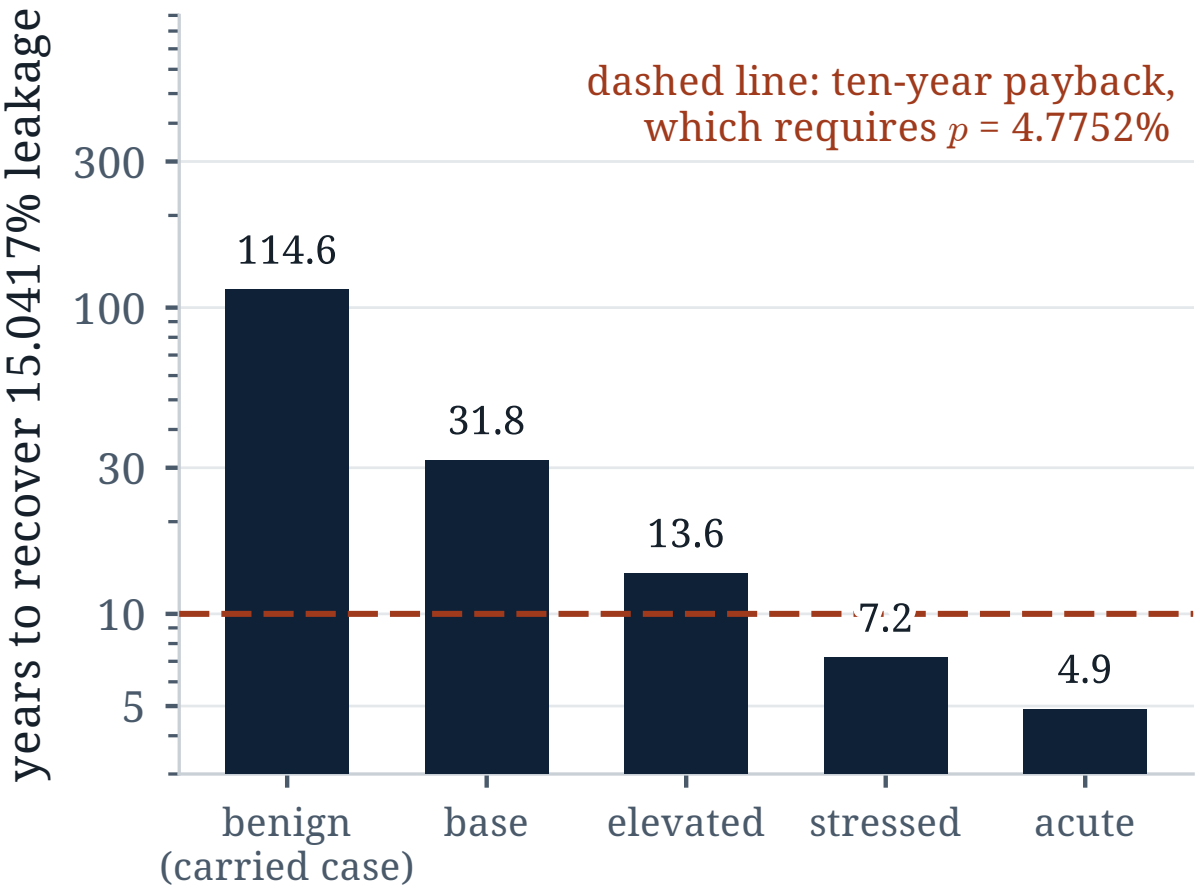


Figure 60.3: **At a fixed leakage of 15.0417 per cent, the conversion repays inside ten years only in the two most severe scenarios of five.** Paybacks are the source's scenario table computed as L/V_{esc}; the benign row is also the carried London case, whose four parameters are identical. Modelled from the constructed parameters in Section 6.3; script `figs/src/CCP-060_fig03.py`.

```
BEFORE -- one building
    London building 100.0%      property HHI = 1.000000
    effective properties = 1.0000

AFTER -- a constructed globally diversified listed allocation,
         forty underlying assets across six listed vehicles
    largest 12.0  0.014400      sixth    6.0  0.003600
    second   9.0  0.008100      seventh  5.0  0.002500
    third    8.0  0.006400      eighth   5.0  0.002500
    fourth   7.0  0.004900      ninth    4.0  0.001600
    fifth    6.0  0.003600      tenth    4.0  0.001600
    thirty further assets, 34.0% total at 1.1333% each
        30 x 0.000128444 = 0.003853
    ------------------------------------------------------
    property HHI = 0.049200 + 0.003853 = 0.053053
    effective properties = 1 / 0.053053 = 18.8491
```

From an effective property count of 1.0000 to 18.8491. With the tenant diversification that arrives with it — from four tenants, the largest at 41 per cent of income, to several hundred — that is the largest single reduction in idiosyncratic risk available anywhere in the carried case. It is a commercial-risk improvement and not a state-action improvement, and the two must not be traded against each other in the same sentence.

Priced: idiosyncratic risk on a single four-tenant building, as an annual expected loss at 2.0 per cent per year of an asset-specific impairment averaging 25 per cent of value, is $0.020 \times 0.25 = 0.500\%$ a year. Across 18.85 effective properties the same class of risk falls approximately as the effective count rises: $0.500\% / 18.85 = 0.0265\%$. The annual saving is 0.4735 per cent of value, and

```
payback = 0.150417 / 0.004735 = 31.77 years
```

Thirty-two years again, from a completely different direction. The coincidence is an artefact of the constructed parameters and the discipline is not: on this asset, at this basis, in this tax system, *no* argument for conversion repays inside a generation, because the leakage is too large relative to any annual benefit that can honestly be claimed. The conclusion is about timing rather than about vehicles. **The leakage is a function of the accrued gain and the accrued gain grows; the conversion that is uneconomic today was economic on acquisition day, when basis equalled value and leakage was 1.5 per cent rather than 15.0 per cent.** At 1.5 per cent leakage the base-case payback is $0.015 / 0.0013125 = 11.43$ years and the concentration-case payback is $0.015 / 0.004735 = 3.17$ years. **Decide the holding form when the asset is acquired, not when it becomes worrying.**

6.5 The jurisdictional surprise

Assigning a listed portfolio to the perimeter of its custodian rather than of its issuers, and applying that rule mechanically, the building's 9.0 points of group exposure leave the property line and follow the custody. Three choices, computed on the designed structure of Curzi [2026] — six jurisdictions at 29.1 / 18.2 / 16.4 / 13.6 / 11.8 / 10.9, naive HHI 0.189002 and 5.290949 effective jurisdictions, blocs at 40.0 / 30.0 / 30.0 giving bloc-HHI 0.340000, $N = 2.941176$ and $h = -0.122417$:

custody	naive effective jurisdictions	bloc-HHI	N	h	Δh
unchanged, existing perimeter	5.290949	0.340000	2.941176	-0.122417	0.000000
split across two bridge centres	5.150000	0.361150	2.768932	-0.143539	-0.021122
moved wholly to one bridge centre	5.120275	0.374200	2.672368	-0.155962	-0.033545

Every attempt to improve the structure by moving the converted position out of the existing perimeter makes the structure measurably worse. The reason is arithmetic and not political: the sum of squares of a three-part distribution is minimised at equal thirds, and taking exposure out of the smallest bloc and adding it to the largest moves the distribution away from balance in both directions at once. The property was doing bloc work.

The marginal calculation runs the other way. Writing x for points moved from the largest bloc into the smallest,

```
bloc-HHI(x) = (0.40 - x)^2 + 0.09 + (0.30 + x)^2
d/dx        = -2(0.40 - x) + 2(0.30 + x) = -0.20 + 4x
negative for x < 0.05

at x = 0.05:  0.35^2 + 0.09 + 0.35^2 = 0.335000
              N = 2.985075
```

A second discrepancy in the source is recorded here rather than corrected silently. The source states $\ln(2.985075) = 1.093761$, giving $h = -0.117184$ and a gain of $+0.005233$. Recomputation gives $\ln(2.985075) = 1.093625$, so $h = -0.117231$ and the gain is $+0.005185$. **The source's logarithm is wrong in the fourth decimal place and the derived gain is overstated by** 4.8×10^{-5}**.** The corrected figure is used throughout this paper and the scaffold figure of $+0.005233$ is superseded. Either way the move is declined, on the ground that the index is a diagnostic and not an objective function — a structure optimised against its own published measure ceases to be measured by it [Goodhart, 1975, Campbell, 1979]. It is recorded because the arithmetic produced it, and suppressing it would be the selective reporting the instrument exists to detect.

The conversion contributes exactly zero to N, and that is not an argument against it. It answers commercial risk, with an effective property count moving from 1.0000 to 18.8491, and a narrow slice of arbitrary state action by reducing exit latency from roughly 270 days to one. It answers perimeter fragmentation not at all, because the land never moved and a claim on land in a jurisdiction is economic exposure to that jurisdiction however it is custodied.

That produces a methodological refinement applicable wherever a security's issuer has concentrated assets:

> For a security whose issuer's assets are concentrated in one jurisdiction, exposure splits across two questions rather than one. **Availability** follows the custodian's perimeter. **Value** follows the situs of the issuer's assets, and no custody arrangement moves it. Report both. A structure reporting only the first has recorded a diversification it did not achieve; a structure reporting only the second has ignored the freeze risk it did acquire.

On that refinement, value exposure to the situs falls from 9.0 points to whatever fraction of the diversified allocation is situated there — 20 per cent of 9.0, or 1.8 points — with the remaining 7.2 points redistributing to the situs of the other assets.

Measured on its own, the replacement allocation is the paper's sharpest instance of the axis divergence:

```
United States  35.0%   WEST     0.122500
United Kingdom 20.0%   WEST     0.040000
Japan          15.0%   WEST     0.022500
Singapore      12.0%   BRIDGE   0.014400
Australia      10.0%   WEST     0.010000
UAE             8.0%   BRIDGE   0.006400
------------------------------------------------------
naive HHI = 0.215800   effective jurisdictions = 4.633920
WEST 80.0% / BRIDGE 20.0%
bloc-HHI  = 0.680000   N = 1.470588   ln(N) = 0.385662
```

A globally diversified listed real estate allocation across six countries and four continents has an effective bloc count of 1.470588. That is not an artefact of the constructed weights: the listed real estate universe is overwhelmingly concentrated in one coordination perimeter by market capitalisation, so any allocation tracking it inherits the concentration. **Global diversification here diversifies the property axis magnificently and the perimeter axis barely at all. Two axes, opposite answers, one transaction — and an adviser reporting either alone is reporting half of a result whose halves point in different directions.**

7. Jurisdictional context: which term of the law a jurisdiction can act on

This section is the survey material, retained as context for the measurement in Section 8.5 and condensed accordingly. The full descriptive profiles are in Appendix C. **None of it is cited to primary legal sources and none of it is offered as this paper's contribution.**

7.1 The structural claim, and the arithmetic that decides it

Territorial taxation — a rule taxing income arising within a state's territory and not income arising outside it — is marketed in several regions as though it conferred independence. It is a rule of domestic law, unilateral, amendable by ordinary legislation, and adopted by states that participate in the same information-exchange, anti-money-laundering and substance architecture as everyone else [Organisation for Economic Co-operation and Development, 2017, Financial Action Task Force, 2012, Organisation for Economic Co-operation and Development, 2015]. **A structure that treats a source rule as though it conferred perimeter independence has misidentified what it bought.**

Four properties of such regimes are routinely misunderstood and all four bear on the measurement. They confer no treaty protection, being unilateral rather than negotiated. They say nothing about the residence position of the person who ultimately bears the tax, so that where the principal is taxed on citizenship or is resident in a controlled-foreign-company jurisdiction the source rule may be irrelevant to the outcome. They attract standards attention precisely because they are attractive: the mechanism by which a listing decision produces a domestic legislative change within months is documented as a technique of governance in its own right [de Goede, 2011], and the source records one surveyed jurisdiction rewriting its income tax law within nine months of a listing and being delisted. And they are not the same regime in each state that has one — one is broad and long-standing, one is territorial with carve-ins for foreign passive income and a time-limited election, one is the purest of the

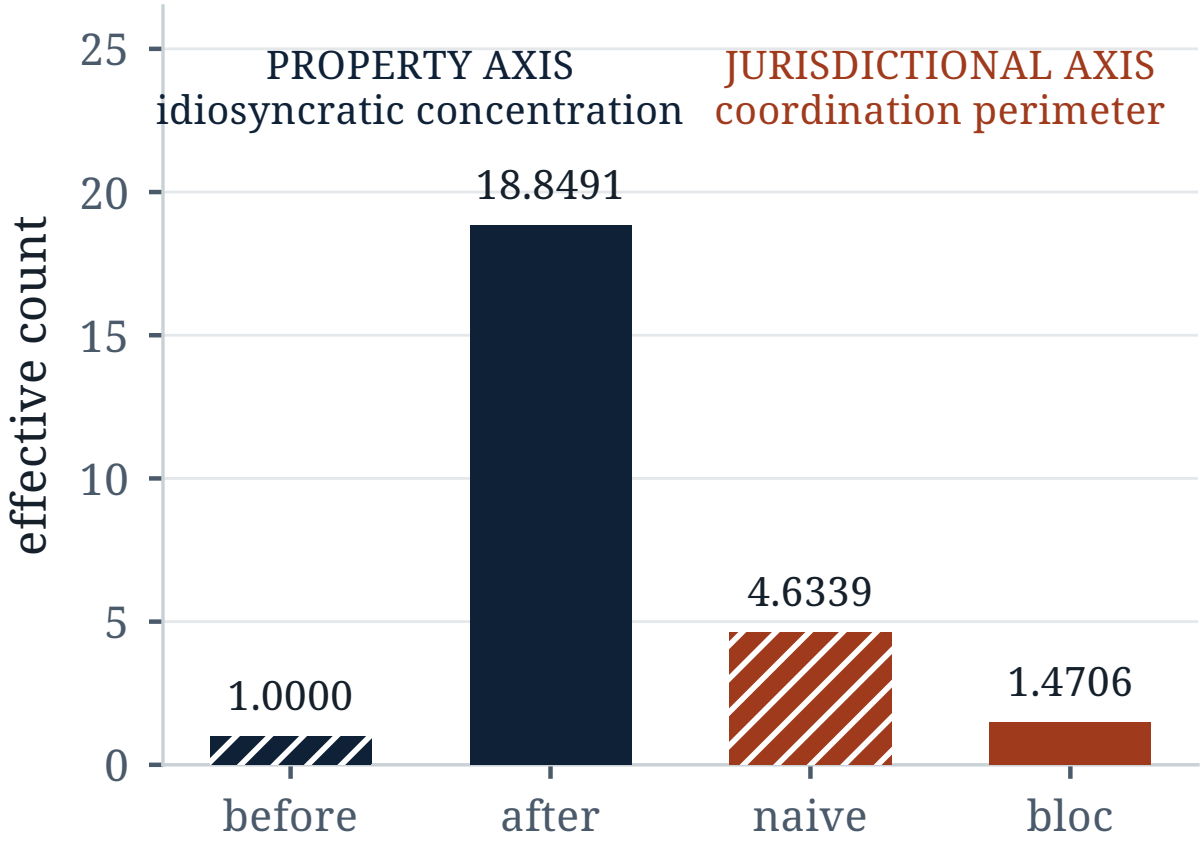


Figure 60.4: **One transaction improves one concentration axis by a factor of nineteen and leaves the other below one and a half.** Effective counts are $1/\mathrm{HHI}$ on the three share bases of Section 3.1, computed from the constructed allocation in Sections 6.4 and 6.5. Hatched bars are the before-and-naive readings, solid bars the after-and-bloc readings; script `figs/src/CCP-060_fig04.py`.

group, and one was narrowed under external pressure. **Treating them as one regime is the commonest analytical error in the field, and the correction costs nothing.**

7.2 A de-branded comparison

The source manuscript scores four jurisdictions across twelve descriptive items. **The scored comparison is reproduced here without jurisdiction names**, following the convention of Curzi [2026]: the underlying characteristics are the analyst's judgements applied to published descriptions and **cannot be supported from this paper's bibliography, which contains no primary source for any individual jurisdiction**. Attaching a jurisdiction's name to an unsupported characteristic is the failure this convention exists to prevent. Named descriptive profiles are in Appendix C for readers who wish to verify them against the official registers.

Scores are 0–5 ordinal rank judgements. **An ordinal scale licenses a minimum and does not license a mean**, and only minima are taken below.

centre	institutional quality	perimeter independence	vehicle set	standards position stateable
Territorial centre A	3	1	4	5
Territorial centre B	5	3	2	5
Territorial centre C	1	3	1	0
Territorial centre D	4	0	2	5

Taking the minimum across the dimensions a holding function requires — institutional quality, perimeter independence and a vehicle set — assigns **no holding function to any of the four**. Centre B scores highest for a residence and operating function; centre C scores an explicit zero on whether its standards position can be stated at all, and the source declines to guess rather than filling the cell in. **A zero that records the absence of an assessable fact is more useful than a score invented to complete a table.**

This is the corpus's scalar-reduction limitation in a second domain and it is inherited rather than restated. Curzi [2026] names it once as a framework-level constraint on what a single score may be asked to carry; the appropriate response here is to point at that statement, take minima rather than means, and not to pretend the table is a measurement.

7.3 What a jurisdiction inside an occupied perimeter can and cannot do

The structural claim, stated so it can be measured: **where a jurisdiction sits inside a coordination perimeter the structure already occupies, it can act on the erosion term of the storage law and cannot act on the surplus term.** The surplus term is $\eta \ln N$, and N counts perimeters rather than flags; the erosion term is δQ, and a genuine reduction in the effective erosion rate is worth having whatever the perimeter reading says. Section 8.5 measures both.

Since the erosion term is multiplied by the stock Q while the surplus term is multiplied by η, **the relative value of the two depends on how large the protected stock already is**: for a large Q, erosion reduction dominates; for a small Q it does not, and the same structure is a waste of money. Both conclusions come from the same law, which is why the law is worth writing down rather than reasoning about the region qualitatively.

8. Results

Every figure below is an exact computation on constructed inputs. **No dataset was collected and no figure in this paper is a measurement taken from a real structure.** Section 12 states that limitation in the terms the corpus uses for it.

8.1 R1 — the cell-wall confidence threshold

With a cost advantage of 0.090 units a year against a fenced benefit of 1.400 units a year, indifference is at $1.400(1-r)=0.090$, so $1-r=0.064286$ and $r^{*}=93.5714\%$. **That is the minimum confidence that a foreign forum respects the cell wall for the cheaper device to be preferable.** The threshold rises with risk: 87.1 per cent at $p=1.0\%$, 93.5714 per cent at $p=2.0\%$, 96.8 per cent at $p=4.0\%$, and 94.36 per cent on the carried case at $p=2.5\%$ with a reach reduction of 63.8 units.

8.2 R2 — contagion measured

Contagion measured as HHI_{c} over reachable shares with effective compartments $=1/\mathrm{HHI}_{c}$: six unequal compartments at 30 / 22 / 18 / 12 / 10 / 8 per cent give $\mathrm{HHI}_{c}=0.201600$, **4.9603 effective compartments**, and cut maximum single reach from 100.0 to 30.0 per cent. At an annual event probability of 2.0 per cent that is worth $0.02\times 70.0=1.400$ units a year against 0.110 of cost — a ratio of **12.73 to one** and a break-even probability of **0.157 per cent per year**. On the carried case the redesign moves HHI_{c} from 0.704800 to 0.145940, the effective count from 1.4188 to 6.8521, and maximum single reach from 82.0 to 18.2 per cent, at a ratio of 12.76 to one and a break-even of 0.196 per cent.

8.3 R3 — leakage against escape value

Conversion of an immovable asset carries certain leakage $L=7{,}220{,}000/48{,}000{,}000=15.0417\%$ against an escape value $V_{\mathrm{esc}}=pL^{*}q(1-s)=0.4725\%$ a year at base parameters, a **payback of 31.83 years**. A ten-year payback requires $p=4.7752\%$ a year. On the carried building's own parameters the escape value is 0.13125 per cent and the payback is **114.60 years**.

8.4 R4 and R5 — the two axes

R4. Diversification through the same conversion raises effective properties from 1.0000 to **18.8491** (property HHI 0.053053), worth 0.4735 per cent a year against the same 15.0417 per cent leakage — a payback of **31.77 years**. Only a non-recognition regime cutting leakage to 1.5 per cent brings the payback to 3.17 years.

R5. The converted global allocation at 35 / 20 / 15 / 12 / 10 / 8 per cent reads naive HHI **0.215800** and **4.633920** effective jurisdictions but bloc-HHI **0.680000** and $N=1.470588$. **The naive count improves while the bloc count worsens.** H4 is supported: the two axes moved in opposite directions on one transaction.

8.5 R6 — the regional sleeve, and where the gain came from

CONSTRUCTED CASE. A hemispheric sleeve as found, and redistributed across a region of territorial-taxation jurisdictions:

```
SLEEVE AS FOUND
  three positions at 55.0 / 25.0 / 20.0
  naive HHI 0.405000 -> effective jurisdictions 2.469
  all three inside one perimeter
```

```
  perimeter HHI 1.000000 -> EFFECTIVE BLOCS 1.00

SLEEVE AS PROPOSED
  five positions at 34.0 / 22.0 / 20.0 / 14.0 / 10.0
  naive HHI 0.233600 -> effective jurisdictions 4.281

SLEEVE, PARTIAL-INDEPENDENCE READING
  perimeter-coupled 70.0 / partially independent 30.0
  perimeter HHI 0.580000 -> EFFECTIVE BLOCS 1.724
```

The honest sleeve result is a range of 1.00 to 1.724 effective blocs against a naive reading of 4.281. The naive number overstates by between two and a half and four times.

At group level, adding two such jurisdictions holding twelve per cent of protected capital to the designed structure:

```
BASELINE       blocs 30.0 / 29.1 / 40.9
               bloc HHI 0.341962 -> N = 2.924302   ln(N) = 1.073056

WITH THE LEG   eight jurisdictions at 26.0 / 16.0 / 14.0 / 12.0
               / 10.5 / 9.5 / 7.0 / 5.0
               naive HHI 0.154650 -> effective jurisdictions 6.466

READING A      strict: the whole region is one perimeter
               blocs 38.0 / 26.0 / 36.0
               bloc HHI 0.341600 -> N = 2.927400   ln(N) = 1.074115

READING B      one jurisdiction treated as partially independent
               blocs 31.0 / 26.0 / 36.0 / 7.0
               bloc HHI 0.298200 -> N = 3.353454   ln(N) = 1.209991

surplus gain at eta = 0.35
  Reading A   0.35 x (1.074115 - 1.073056) = 0.000371
  Reading B   0.35 x (1.209991 - 1.073056) = 0.047927
```

Under the strict reading, adding two jurisdictions holding twelve per cent of protected capital raised the effective bloc count by 0.0031 and the surplus term by 0.000371 — for practical purposes, by nothing. Against a maintenance increment of 0.05, both readings lose money on diversification grounds alone:

```
baseline                   h = 0.35 x 1.073056 - 0.20 - 0.30 = -0.124431
Reading A, no erosion gain h = 0.35 x 1.074115 - 0.20 - 0.35 = -0.174060
Reading B, no erosion gain h = 0.35 x 1.209991 - 0.20 - 0.35 = -0.126503
```

Now introduce the term the region actually acts on. If the effective erosion rate falls from 0.020 to 0.013:

```
Reading A, delta = 0.013   h = 0.375940 - 0.130 - 0.35 = -0.104060
Reading B, delta = 0.013   h = 0.423497 - 0.130 - 0.35 = -0.056503
```

Both now beat the baseline, by 0.020371 and by 0.067928. **On the strict reading the entire improvement comes from the erosion term: the bloc term contributed 0.000371 of a 0.020371 gain, one part in fifty-five.**

The result generalises and is the point of Section 7.3. Where a jurisdiction sits inside a perimeter the structure already occupies, it can help only through the erosion term and never through the surplus term. And the erosion route is itself conditional on a fact about the person

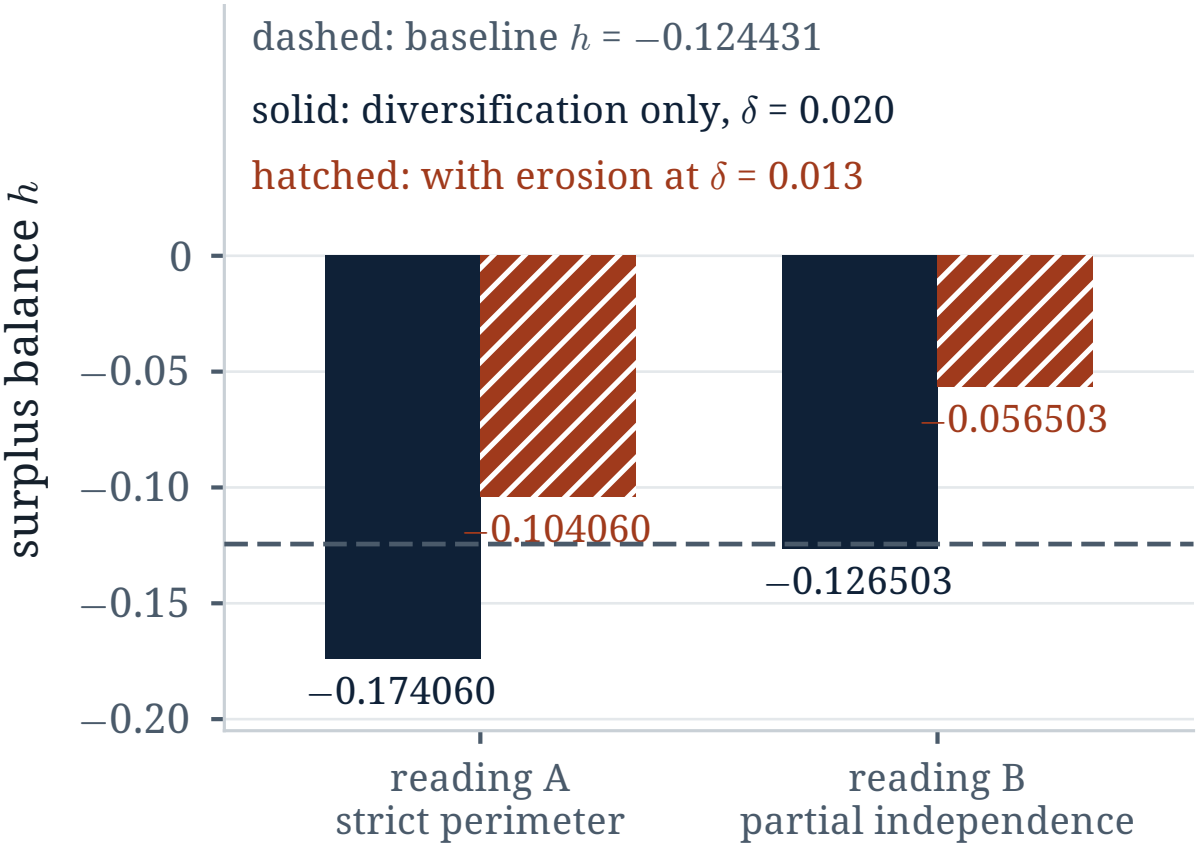


Figure 60.5: **The regional leg's entire gain under the strict reading comes from the erosion term, not from the perimeter count.** Surplus balance $h = \eta \ln N - \delta Q - C$ at $\eta = 0.35$, $Q = 10$, $C = 0.35$, computed from the constructed bloc shares of Section 8.5; the dashed line is the baseline at $C = 0.30$, $\delta = 0.020$. Script `figs/src/CCP-060_fig05.py`.

rather than about the structure: where the principal remains taxable on citizenship regardless of residence, δ does not move at all, both readings stay at -0.174060 and -0.126503, and the whole proposition is worse than doing nothing. **A regional proposition that is worth 0.020371 to one principal and -0.049629 to another, on one fact about the principal, must be priced after that fact is established rather than before.**

8.6 Counter-result

Leakage of 15.04 per cent is certain while every benefit computed against it is probabilistic, giving paybacks above thirty years on every route examined. The conversion returns 31.83 years on the state-action argument, 114.60 years on the carried building's own parameters, and 31.77 years on the concentration argument — three independent routes, all outside a working lifetime, because the leakage is too large relative to any annual benefit that can honestly be claimed.

And the ring-fencing redesign's gain on the jurisdictional measure is an honest near-zero. The best available rebalancing move buys $+0.005185$ of surplus (the source prints $+0.005233$; Section 6.5 records the arithmetic error), and compartmentalisation itself contributes 0.000000 to N while adding 0.125 to 0.175 units a year to C. **Under the storage law alone that is unambiguously negative, and it is justified entirely on an axis the storage law does not contain.** Reporting the contagion gain and the jurisdictional figure together, as a single improvement, would be the reporting failure the instrument exists to detect.

9. Discussion and failure catalogue

9.1 What the measurement does and does not license

The contagion index measures **independence, not capacity**, exactly as the jurisdictional index does. An effective compartment count of 6.8521 says nothing about whether any compartment can carry the load, and nothing about whether the group would in practice let a compartment fail. That last point is the Medici warning, and it is the oldest result in this paper's subject: branches constituted as legally separate partnerships, each with its own capital and its own junior partner, did not save the bank, because the senior partnership had capital in each branch, the same men appointed each branch's manager, and depositor confidence is not compartmented by a partnership deed [de Roover, 1963]. **Legal separation does not separate capital, management judgement or reputation, and a structure that separates only the first has separated the least important of the three.**

The same reservation in modern form: a de-risking decision is taken about a client relationship, and the client relationship is the group. **Compartmentalisation is a defence against a claim. It is not a defence against a decision, and most of the events this analysis is about are decisions.** Curzi [2026] makes the point on a different share basis and reaches the same shape — entity-based exposure of 27.3 per cent against a currency share of 61.0 per cent and settlement volume near 58 per cent, so that the entity map understates dependency by more than a factor of two.

9.2 Thirteen failure modes

Condensed from the source; the eleventh changes the character of the problem and is placed accordingly.

1. **Commingling.** One account used for two compartments, one payment from the wrong entity, one receivable collected into the parent and not swept. It appears in every alter-ego and consolidation finding and is created by operational convenience rather than by a structuring decision.

2. **Guarantees, comfort letters and keep-wells.** Given for good commercial reasons by someone raising finance who has not read the structure memorandum. Comfort letters are worse than guarantees because they are ambiguous.
3. **The independent director who is not independent.** A consent right held by a person the sponsor appoints, pays and can replace is a document.
4. **Substantive consolidation.** Decided on evidence about how the group behaved. No drafting escapes a course of dealing, and by the time consolidation is argued the course of dealing is a decade long and is in the group's own correspondence.
5. **Cross-default and cross-collateralisation.** Six compartments each carrying a cross-default to any group indebtedness above a threshold accelerate simultaneously. The contagion index reads 0.2016 on the corporate chart and 1.0000 on the financing documents. **Run the measurement on the financing documents.**
6. **Undercapitalisation.** A compartment funded at the statutory minimum and expected to carry a large obligation invites the argument that separation was formal.
7. **Confusing accounting with law.** Deconsolidation states what the financial statements show and nothing about what an insolvency court will do.
8. **The cell wall in a foreign forum.** The trap is a cell structure created for a domestic purpose and used three years later for a transaction with a counterparty elsewhere. Nobody re-runs the analysis when the use changes.
9. **Cells holding hard assets in a third state.** The sharpest version of the previous mode, and the recommendation is the refusal: **do not put a hard foreign asset in a cell.**
10. **Multiplying compartments without measuring cost.** Twenty compartments where seven satisfy the constraint cost 0.325 units a year for no measurable benefit on either axis.
11. **Ring-fencing after the fact.** Compartmentalisation undertaken before any obligation arose is structuring; the same transfer undertaken when a claim is in prospect is a transaction liable to be set aside. **A compartment built at the right time is a wall; the same compartment built at the wrong time is evidence.** The governing law is avoidance law and avoidance law is a clock; the periods and provisions are described in Appendix B and are not cited here.
12. **The orphan that is not orphaned.** A side letter, an option over the shares, a fee arrangement capturing the residual. All are discoverable, and substance is what gets litigated.
13. **Believing the compartment separates what it cannot.** Section 9.1.

9.3 Where the analysis is exposed

The instrument is exposed to its own tenth failure mode. Adopted as a target rather than as a diagnostic, an effective count ceases to be a good measure [Goodhart, 1975, Campbell, 1979]. Nothing in the arithmetic prevents it; the only defence offered is the reporting standard in Appendix A, which requires the inputs to be published beside the output.

The instrument is also performative in a way its own parameters are not exempt from. MacKenzie and Millo [2003] documents a model that made its assumptions true after adoption, and Lucas [1976] states the consequence for parameters estimated before publication. The framework asserts that systems change their behaviour when modelled and then exempts its own η, δ and correlation assumptions from that reflexivity. **That is an inconsistency and it is not resolved here.**

Ensemble and time averages are not the same object, and every expected-loss figure here is the first. Peters [2019] establishes that the time average of a multiplicative process differs from its ensemble average, so an expectation computed across many holders does not describe what happens to one holder over time. A single compartment-consuming event is not a deduction from a stream; it rescales every subsequent year. Peters and Gell-Mann [2016] supplies the associated strengthening for the framework's logarithm: a logarithmic form follows from multiplicative dynamics with no psychological assumption at all, which is a better warrant than the perceptual account the corpus more often gives, and Section 12 records that

the paper does not adjudicate between the two.

The device analysed here is a hedge against a decision, and hedges bought inside the warning window are not available. The corpus's amplitude weighting [Curzi, 2026] scores each step of a path for steepness and weights the branch as $1/2^R$, so a three-point difference in steepness is a factor of eight in expectation; applied here, the branches that matter are the ones reachable before the signal, and every device in this paper has a build time measured in months against warning windows measured in days. **The only window common to every threat is *before*.**

One boundary the framework itself refuses to cross is worth naming here rather than in a footnote. Curzi [2026] states the doctrine of category violation, which refuses category innocence as a defence of acquisition and forbids moving a harm through a discovered gap in a scheme. **The instruments in this paper are measured against arbitrary action and are not licensed against a judgment that survived due process**, and the eleventh failure mode is where that boundary bites. A reader who takes the arithmetic and drops the constraint has taken half the analysis.

10. Verification and rigour audit

The results are decomposed into postulates that can be checked one at a time. Each is atomic, each carries a status, and the chain names what each step consumes. **If dropping a postulate leaves a result standing, it is not listed.**

10.1 Dimensional audit

symbol	quantity	dimension	unit	status
s_i	reachable, exposure or property share	$[1]$	dimensionless fraction	defined; sums to 1 by P1
HHI, HHI_c	concentration index on a stated basis	$[1]$	dimensionless (or points $= 10^4\times$)	derived from shares
N, N_c	effective count	$[\#]$, a count	blocs, compartments or properties	derived as $1/\mathrm{HHI}$
$\ln N$	surplus-term argument	$[1]$	nats	derived; the logarithm of a count is dimensionless
p	annual event probability	$[T^{-1}]$	per year	assumed (constructed)
r, r^*	probability the cell wall holds	$[1]$	dimensionless	assumed; r^* derived
ΔC	cost saving from the cheaper device	$[U\,T^{-1}]$	estate units per year	assumed (constructed)
ΔS	reach reduction forfeited on failure	$[U]$	estate units	derived from shares $\times$ estate size
V, B	gross value, base cost	$[M]$	currency	assumed (constructed)

symbol	quantity	dimension	unit	status
L, L^*	leakage fraction, loss fraction	$[1]$	dimensionless	derived; assumed
q, s	exit probability, sale discount	$[1]$	dimensionless	assumed (constructed)
V_{esc}	escape value	$[T^{-1}]$	fraction of value per year	derived
Q	protected capital	$[C]$	index units	defined; **no operational unit is supplied by the source**
δ	erosion rate	$[T^{-1}]$	per period	assumed
C	maintenance cost flow	$[C\,T^{-1}]$	index units per period	assumed
η	conversion efficiency	$[C\,T^{-1}]$ per nat	index units per period per nat	assumed
h	surplus balance	$[C\,T^{-1}]$	index units per period	derived

Do the equations balance? The threshold condition $p\,\Delta S\,(1-r) = \Delta C$ balances: $[T^{-1}][U][1]$ against $[U\,T^{-1}]$. The payback L/V_{esc} balances, $[1]/[T^{-1}] = [T]$, and is therefore correctly quoted in years. The surplus law $h = \eta \ln N - \delta Q - C$ balances **only if η carries capital per period per nat**, and the source is silent on this: it treats conversion efficiency as a bare number. **The audit records the silence rather than filling it in**; the reading in the table is this paper's.

Where no operational unit exists the audit says so. Q, C and h are in index units fixed by the stated parameters, and the source supplies no conversion to currency. **The break-even probabilities, the confidence thresholds and the paybacks are unit-free or in years and survive this; the statement that h is -0.122417 means nothing until an index unit is defined.**

10.2 Significant figures against epistemic status

- $r^* = 93.5714\%$ and 94.36% are **exact arithmetic on stipulated inputs**: $1 - 0.090/1.400$ and $1 - 0.090/1.595$. They are derived-from-defined and carry as many digits as are useful. **Their inputs are constructed and carry no empirical authority whatever.**
- $L = 15.0417\%$ is exact on the stated valuation and rate; the rate itself is an assumed current-law figure and is flagged in Appendix B.
- 4.9603, 6.8521, 18.8491 and 1.470588 are exact reciprocals of exactly stated share vectors. The shares are constructed and stated to one decimal place; **every index in this paper is computed from the printed percentages rather than from any exact underlying fraction, which is the convention adopted by Curzi [2026] and recorded as a discrepancy in the fifth decimal against the exact-fraction convention of Curzi [2026].** The convention is stated so that a reproduction agrees to the last digit.
- Paybacks of 31.83, 31.77 and 114.60 years are **illustrative model output from constructed parameters**, rounded hard. Quoting them to two decimals is defensible only because the inputs are exact by stipulation.
- **No quantity in this paper is a fitted or measured estimate, so no uncertainty interval is reported anywhere. That is not precision; it is the absence of empirical content.**

10.3 Limits and boundary behaviour

limit	behaviour	reading
one compartment, $s_1 = 1$	$\mathrm{HHI}_c = 1$, $N_c = 1$ exactly	an unfenced estate counts as one
n equal compartments	$N_c = n$ exactly	returns the count
n unequal compartments	$N_c < n$	never overstates the count
$n \to \infty$, equal	$N_c \to \infty$	unbounded above, as a count should be
$r \to 1$	r^* condition always satisfied	a perfectly recognised wall makes cells strictly better
$p \to 0$	$r^* \to -\infty$	at zero risk any device is acceptable; the formula leaves its domain
$p \to \infty$	$r^* \to 1$	at extreme risk only a certain wall will do
$\Delta C \to 0$	$r^* \to 1$	a cheap device that saves nothing must never fail
$q \to 0$ or $s \to 1$	$V_{\mathrm{esc}} \to 0$	payback diverges; latency without the ability to act is worth nothing
$B \to V$	$L \to f$	at zero accrued gain the leakage is transaction costs alone

Range: $1 \le N_c \le n$, with both bounds attained. A quantity equal to n for n equal units, equal to 1 for a single unit, and strictly between otherwise, **is a count**, and no further argument is needed.

Two limits are constraints on the domain rather than results. r^* becomes negative for $p < \Delta C/\Delta S$, which is not a confidence and means the cheap device is unconditionally correct at that risk level; the formula's domain is $p > \Delta C/\Delta S$, which for the constructed estate is $p > 0.157\%$. And $h \to \infty$ logarithmically as $N \to \infty$, which contradicts the organisational evidence and is a defect of the model rather than a finding; **the model's domain is small** N, and every prescription here is confined to moves between $N \approx 1$ and $N \approx 3$. As $\eta \to 0$, $h \to -(\delta Q + C)$ regardless of N. As $\delta \to 0$ the stock never erodes and any positive $\eta \ln N$ exceeding C accumulates without bound, which is a pathology rather than an ideal.

10.4 Invariance

Every effective count here is invariant under relabelling and under any rescaling of the underlying quantities that preserves the share vector. Doubling the estate in currency terms leaves every s_i unchanged and therefore leaves HHI, N and $\ln N$ unchanged. That is a genuine scale invariance and it is why the instrument applies to an estate of any size. **It is not invariant under regrouping**, which is the entire content of the perimeter collapse: the difference between the two groupings is the finding.

What would break it: any share basis that is not normalised — reporting reachable amounts in currency without dividing by the estate. **Tested at how many scales?** The invariance is algebraic rather than empirical, and was checked at exactly the share vectors reported here: two, six, seven, eight and forty entries. No claim of empirical scale-freedom is made.

What is conserved, and in what units. Mirowski [1989] shows that the characteristic failure of this genre is borrowing the notation of a conserved quantity without the conservation law that gave it content. The conserved quantity here is the share vector, $\sum_i s_i = 1$ exactly and dimensionlessly, and it is that constraint alone that gives the index its bounds. **Remove the normalisation and the effective count ceases to be a count.** Q is a stock with no stated

conservation law and no operational unit, and no conservation is claimed for it. Cartwright [1999] supplies the discipline that makes a partial answer respectable: specify the arrangement within which the law is claimed to hold and stop claiming universality the argument does not need. **The arrangement claimed here is a normalised share vector on one stated basis, and nothing outside it.** Kirman [1992] is the objection this paper cannot answer: one Q, one δ and one C are written for a heterogeneous group of compartments, and a single representative parameter for a heterogeneous population generally does not exist.

10.5 Analogy discipline

borrowing	source	category	justification
effective count as $1/\mathrm{HHI}$	Hill [1973], Jost [2006]	**1 — derivation**	HHI^{-1} *is* the Hill number of order 2; an identity, not a resemblance
numbers-equivalent reading	Adelman [1969]	**1 — derivation**	elementary and checkable in two lines
portfolio variance $w'Rw$	Markowitz [1952]	**1 — derivation**	the exposure vector genuinely is a weighted sum of shocks with a correlation matrix
similarity-sensitive diversity	Leinster and Cobbold [2012]	**1 — derivation**, stated domain	exact where R has unit diagonal and non-negative entries; claimed nowhere else
$\ln N$ as entropy	Shannon [1948]	**2 — shared form**	Shannon entropy in nats, a dimensionless information quantity; no thermodynamic content is imported
coupled-network fragility	Buldyrev et al. [2010]	**2 — shared form**	financing covenants have not been shown to satisfy the percolation premises; only the structural property is used
identity inflation	Douceur [2002]	**2 — shared form**	the impossibility result is about network identities; the structural analogy to entity count is stated as an analogy

borrowing	source	category	justification
tipping from a small correlated shift	Granovetter [1978], Schelling [1971]	**3 — framing**	prior statement of the same aggregate phenomenon, cited as ancestry rather than as a formal result about perimeters
normal accidents	Perrow [1984]	**3 — framing**	qualitative organisational theory, never used as a formal result
synchrony	Loreau and de Mazancourt [2008]	**2 — shared form**	the naming convergence is noted; no ecological mechanism is imported

The failure this table exists to prevent is letting a category-2 borrowing carry a category-1 result's authority. The concentration arithmetic is category 1 throughout and needs no help; the network and tipping results are category 2 and 3 and are used for shape and ancestry only.

10.6 Measurement-theoretic status of the scored inputs

The 0–5 comparison in Section 7.2 is **ordinal**. It records the analyst's rank judgements over published descriptions. **An ordinal scale licenses a minimum and does not license a mean**, and only minima are taken. No average of those scores appears anywhere in this paper, and none should be computed from the table.

Reachable shares, exposure shares and property shares are ratios of capital and are not psychological judgements, so the concentration instrument is untouched by the scale-type objection. **The event probabilities are a different matter and the objection is conceded.** p is treated throughout as a ratio-scale quantity supporting multiplication and division, and it is a subjective degree of belief elicited from nobody. Michell [1997] is the strongest available objection to that move — that quantitative structure is assumed rather than tested — and the formal remedy is conjoint measurement [Luce and Tukey, 1964], whose conditions establish whether an attribute has additive quantitative structure at all, within the representational framework of Krantz et al. [1971]. **The conjoint conditions were not tested for any probability in this paper, and every ratio computed from one is therefore provisional.** What survives without the test is the *ordering* — a break-even probability still separates settings where a device pays from settings where it does not — and the algebraic threshold form itself, which does not depend on the scale type.

10.7 Postulates

#	statement	status
P1	shares s_i on a stated basis are non-negative and sum to 1	DEFINITION
P2	HHI $= \sum_i s_i^2$ on that basis	DEFINITION
P3	for n equal shares, HHI $= 1/n$	ARITHMETIC

#	statement	status
P4	the reachable share of a compartment is the fraction of the estate a single event arising against it can reach	DEFINITION
P5	expected loss is $p\times$ the reachable share of the largest compartment	DEFINITION (modelling choice: one event per year, largest compartment consumed)
P6	a foreign forum respects a statutory cell partition with probability r, unknown	POSTULATED
P7	contribution of an appreciated immovable asset is a disposal at market value in the modelled system	POSTULATED (H3; uncited, Appendix B)
P8	the surplus law's input term is $\eta \ln N$	DEFINITION (modelling choice, CCP-A02)
P9	$\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$	POSTULATED (fixed in advance)
P10	jurisdictions inside one coordination perimeter act on one decision	POSTULATED (stipulated reference group, Section 3.7)

10.8 The compounding chain

step	consumes	statement	status
L1	P2, P3	$1/\mathrm{HHI}$ returns n exactly for n equal shares	ARITHMETIC
L2	L1	$1/\mathrm{HHI}$ is a numbers-equivalent — Adelman [1969]	CITED
L3	L2	it is the Hill number of order 2 — Hill [1973], Jost [2006]	CITED
L4	L3, P4	on the reachable-share basis it counts effective compartments	DEFINITION
R2a	L4	30/22/18/12/10/8 gives $\mathrm{HHI}_c = 0.201600$, $N_c = 4.9603$, max reach 30.0%	ARITHMETIC
R2b	R2a, P5	at $p = 0.02$, benefit $= 0.02 \times 70.0 = 1.400$ against cost 0.110; ratio 12.73; break-even $p = 0.110/70.0 = 0.157\%$	ARITHMETIC
R1	R2b, P6	$1.400(1 - r) = 0.090 \Rightarrow r^* = 93.5714\%$	ARITHMETIC

step	consumes	statement	status
R1b	R1	r^{*} rises with p: 87.1% at 1.0%, 96.8% at 4.0%	ARITHMETIC
R3a	P7	a 25% rate on a gain of 26.0m plus costs of 1.5% on 48.0m gives leakage of 7.22m, so $L = 15.0417\%$	ARITHMETIC
R3b	R3a	$V_{\text{esc}} = 0.015 \times 0.60 \times 0.70 \times 0.75 = 0.004725$; payback $= 31.83$ yr	ARITHMETIC
R3c	R3b	ten-year payback needs $p = 0.0150417/0.315 = 4.7752\%$	ARITHMETIC
R4	L1	property HHI $0.053053 \Rightarrow 18.8491$ effective properties; saving 0.4735%/yr; payback 31.77 yr	ARITHMETIC
R5	L1, P10	allocation naive $0.215800 \Rightarrow 4.633920$; bloc $0.680000 \Rightarrow 1.470588$	ARITHMETIC
R6	L1, P10, P8, P9	sleeve perimeter HHI $0.580000 \Rightarrow 1.724$; group readings buy 0.000371 and 0.047927 against $+0.05$ of cost	ARITHMETIC

Every ARITHMETIC row can be redone by hand or in three lines of any language. Every CITED row names one source and no more. **Four of the ten postulates are POSTULATED and one of those — P7 — is a legal proposition this paper cannot cite.** The chain therefore establishes arithmetic consequences of stipulated inputs, and **it is not a proof of anything about the world**. A result honestly labelled as following from two postulates is worth more than the same result implied to be measured.

10.9 Cross-checks performed

- Every index was recomputed from its printed share vector; all reported share vectors sum to 100.0 per cent exactly.
- $s_{\max}^{2} \le$ HHI $\le s_{\max}$ holds on all four contagion vectors: $0.672400 \le 0.704800 \le 0.820$ and $0.033124 \le 0.145940 \le 0.182$ on the carried case.
- The effective-count ceiling $s_{\max}^{-2}$ is respected: $1.4188 \le 1.4878$ and $6.8521 \le 30.19$.
- The two source discrepancies reported in Sections 5.5 and 6.5 were found by this recomputation and are printed rather than corrected in silence. A third, smaller one: the source prints 5.120272 naive effective jurisdictions for the single-centre custody case where $1/0.195302 = 5.120275$; the difference is in the seventh significant figure and the corrected value is used in Section 6.5.

11. Limitations

Every case is constructed. The estate, the compartments, the reachable shares, the costs, the event probabilities, the building, its valuation and basis, the allocation, and the exposure vectors are invented and are labelled as such at each point of use. **No external dataset was collected and none is reported.** The measurements are exact computations on constructed inputs.

No paper in this corpus reports a first-party measurement, and this one does not either. The corpus's own instrument audit finds that of 27 parameter-by-scope cells, 19 have a defined estimator, 1 has a validated one, and **0 carry a first-party measurement**. That is the central limitation of the whole corpus and it is stated here before a reviewer states it.

The event probabilities are the most load-bearing and least defensible inputs. $p = 2.0\%$ and $p = 2.5\%$ carry the entire cost–benefit verdict, and p, L^*, q and s carry the entire payback verdict. None is estimated from data; none has an interval; and no reference class was constructed in which any of them could be defended. **The correct reading of every ratio in Section 8 is therefore conditional: at these probabilities, this is the sign of the difference.** The threshold form of Section 3.5 exists precisely because the probabilities cannot be defended, and a reader who substitutes different ones can read a different threshold off the figure.

Jurisdiction-specific legal claims are not cited to primary sources. The descriptions of separateness doctrine, cell legislation, orphan and purpose trusts, avoidance provisions, tax reliefs, disclosure regimes, situs rules and the intermediated-securities rule are taken from the source manuscript. **The bibliography contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction, and no case law of any kind.** Those claims are therefore uncited here, are described rather than asserted, and require verification against the official sources before any reliance. Appendix B lists them individually. **This is a real gap and it is not closed here.**

The four scored jurisdictions are de-branded for exactly that reason. A named jurisdiction attached to an unsupported characteristic is a claim about a real place that this paper cannot back; the convention of Curzi [2026] is followed, and the named descriptive profiles in Appendix C carry the same warning at their head.

The cell recognition question is unresolved in the abstract and was resolved on the carried case by avoidance rather than by analysis. If a group later needs cells in a cross-border setting the question returns, and this paper has no answer to it. The 93.5714 per cent threshold tells a reader what confidence is required; it supplies no evidence about whether that confidence is warranted anywhere.

The perimeter assignments are stipulated, not identified. Manski [1993] rules out recovering an endogenous effect from observation of a linear-in-means reflexive system, and the remedy adopted is stipulation with a defence [Brock and Durlauf, 2001]. No randomisation, instrument or partial-population variation was available. **The price is that every bloc figure in Sections 6.5 and 8.5 is a modelling position.**

The effective count measures independence and not capacity. $N_c = 6.8521$ says nothing about whether any compartment can carry the load, and nothing about whether a group would in practice allow a compartment to fail rather than recapitalise it.

The measurement will decay. Every new facility, guarantee or operational shortcut moves the contagion index, and the index is only meaningful as a difference against last year's. Appendix A makes the date and the difference against the previous measurement a required reporting item for that reason.

The instrument is least informative exactly where the carried case begins. Golosov [2010] documents poor behaviour of effective-number measures in the dominated-unit regime, and 82.0 per cent in one compartment is that regime. The measurement is exact; its discriminating power there is at its minimum.

Small-sample bias is not addressed. Chao et al. [2014] show that low-order diversity estimators are biased on small samples, and a two- or seven-element share vector is a small

sample. No correction was applied, and none is available for the case where the vector is a complete census rather than a sample. **The paper does not resolve which of the two situations it is in.**

Model comparison was not performed. No alternative index was fitted, no out-of-sample test was run, no information criterion was computed and no matched null was constructed, because nothing here is fitted. The instrument is asserted to be the right one on the strength of the identity chain in Section 10.8 and not on comparative fit.

The functional form of the surplus term is underdetermined by the evidence. The corpus proposes $\ln(1 + (N - 1)u)$; a sub-linear power law and an $n \log n$ form fit the same qualitative facts. Nothing in this paper discriminates between them, and the paper concedes that nothing yet does. Peters and Gell-Mann [2016] is the one genuine strengthening available — a logarithm derived from multiplicative dynamics with no psychological premise — and it is cited in Section 9.3 as the warrant the argument actually needs.

No differential prediction is offered. Durlauf [2005] asks for the prediction only this model makes, and the instrument makes none: any correlation-aware concentration measure would produce the same collapse, and any expected-value calculation would produce the same break-even. **The contribution is the share basis, the threshold form and the discipline of printing both sides, not a uniquely discriminating forecast.**

The corpus's scalar-reduction habit applies here and is inherited rather than restated. Collapsing a structured legal situation into a single 0–5 score, or a structured threat into a single probability, is the same move the framework identifies once as a framework-level limitation [Curzi, 2026]. It is not re-argued in this paper; Section 7.2 points at it and takes minima rather than means.

Rigour checks not met, listed by name.

1. **No as-of date, because no real holdings were measured.** A concentration figure from real holdings requires its source and its date; every share vector here is constructed, so there is neither. Appendix A makes both mandatory for any application to a real structure.
2. **No uncertainty is reported anywhere, because nothing is estimated.**
3. **The dimension of η is supplied by this paper and not by the source.**
4. **Q, C and h have no operational unit.**
5. **The surplus law's large-N limit is a defect, not a result.**
6. **The conjoint conditions were not tested for the event probabilities**, so every ratio computed from one is provisional as a ratio-scale quantity [Michell, 1997, Luce and Tukey, 1964].
7. **No distributional or scaling claim is made, so the Clauset–Shalizi–Newman procedure was not run** [Clauset et al., 2009]. Any restatement of this paper reaching for a heavy-tail claim about holdings or reachable shares must run it first, and must read Stumpf and Porter [2012] on why passing it still would not identify a mechanism.
8. **The coupled-network and identity-inflation borrowings are category 2, not category 1** [Buldyrev et al., 2010, Douceur, 2002].
9. **Aggregation is unaddressed** [Kirman, 1992].
10. **Expected value and time average are not the same thing, and every figure here is the first** [Peters, 2019].
11. **The instrument is performative and its own parameters are not exempt** [MacKenzie and Millo, 2003, Lucas, 1976].
12. **There is no conservation law for Q** [Mirowski, 1989].
13. **Two arithmetic errors and one rounding slip in the source are reported rather than corrected in silence** (Sections 5.5, 6.5 and 10.9).

12. Conclusion

This paper prices three structural devices against the certain costs they incur, and the following statements are the fixed text of its results.

R1. With a cost advantage of 0.090 units a year against a fenced benefit of 1.400 units a year, indifference is at $1.400(1-r) = 0.090$, giving $r^{*} = 93.5714\%$ — **the minimum confidence that a foreign forum respects the cell wall** for a cell company to be preferable to separate companies. The threshold rises with the risk it is used against: 87.1 per cent at $p = 1.0\%$ and 96.8 per cent at $p = 4.0\%$.

R2. Contagion measured as HHI_{c} over reachable shares with effective compartments $= 1/\mathrm{HHI}_{c}$: six unequal compartments at 30 / 22 / 18 / 12 / 10 / 8 give **4.9603 effective compartments** and cut maximum single reach from 100.0 to 30.0 per cent, worth $0.02 \times 70.0 = 1.400$ units a year against 0.110 of cost — **12.73 to one**, break-even $p = 0.157\%$ per year.

R3. Conversion of an immovable asset carries certain leakage $L = 7{,}220{,}000/48{,}000{,}000 = 15.0417\%$ against an escape value $V_{\mathrm{esc}} = pL^{*}q(1-s) = 0.4725\%$ a year — a **payback of 31.83 years**; a ten-year payback requires $p = 4.7752\%$.

R4. Diversification through the same conversion raises effective properties from 1.0000 to **18.8491** (property HHI 0.053053), worth 0.4735 per cent a year against the same 15.0417 per cent leakage — payback **31.77 years**. Only a non-recognition regime cutting leakage to 1.5 per cent brings the payback to 3.17 years.

R5. The converted global allocation at 35 / 20 / 15 / 12 / 10 / 8 per cent reads naive HHI **0.215800** and **4.633920** effective jurisdictions but bloc-HHI **0.680000** and $N = 1.470588$ — **the naive count improves while the bloc count worsens.**

R6. The regional sleeve measured alone gives perimeter HHI **0.580000** and **1.724 effective blocs** against a naive 4.281; at group level the strict and partial-independence readings buy **0.000371** and **0.047927** of surplus against $+0.05$ of added cost, so that under the strict reading the bloc term contributes one part in fifty-five of the total gain and the erosion term contributes the rest.

Counter-result. Leakage of 15.04 per cent is certain while every benefit computed against it is probabilistic, giving paybacks above thirty years on all three routes examined — 31.83, 114.60 and 31.77 years. **And the ring-fencing redesign's gain on the jurisdictional measure is an honest near-zero**: at best $+0.005185$ of surplus from the single available rebalancing move, with compartmentalisation itself contributing 0.000000 to N against 0.125 to 0.175 units a year of added maintenance. Two of the three devices priced here are, under the storage law alone, unambiguously negative, and are justified entirely on an axis the storage law does not contain.

Two further results are established as method rather than as measurement. **A cost saving, an event probability and a reach reduction determine a threshold belief**, so a conflict-of-laws question with no reliable answer can be decided without answering it — and the threshold *rises* with risk, which is the opposite of the usual intuition about economising in risky settings. **And a security whose issuer's assets are concentrated in one jurisdiction splits exposure into two questions rather than one**: availability follows the custodian's perimeter and value follows the situs of the issuer's assets, so both must be reported, because a structure reporting only the first has recorded a diversification it did not achieve.

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every case in this paper — the estate, the compartments and their reachable shares, the cost model, the event probabilities, the building and its valuation and basis, the replacement allocation, the exposure vectors, and the jurisdictional scores — is **CONSTRUCTED** and is labelled as such at each point of use. The numerical results are exact computations on those constructed inputs, carried at full precision and rounded only for printing.

All figures were produced by the scripts named in their captions (`figs/src/CCP-060_fig01.py` through `figs/src/CCP-060_fig05.py`), which import the shared figure-style module and read

their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

Jurisdictional and legal descriptions in this paper are taken from the source manuscript and not from primary legal sources. The bibliography contains general comparative, treaty and doctrinal material only; it contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction, and no case law. Appendix B lists the claims affected. Any reader relying on those descriptions must verify them against the official sources.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. Every person, family, entity and transaction described is invented.

Legal, tax and investment advice

This paper is analytical and is not legal, tax, or investment advice. It describes entity types, compartmentalisation devices, conversion mechanisms and jurisdictions for the purpose of measuring concentration and pricing structural choices. No recommendation is made in respect of any particular person, asset, structure or jurisdiction. Cell legislation, avoidance provisions and their limitation periods, consolidation standards, substance requirements, tax rates and reliefs, situs and intermediated-securities rules, listing conditions and disclosure obligations all change and vary by jurisdiction and by activity; **every such statement in this paper is carried from the source manuscript, is flagged in Appendix B, and requires confirmation with qualified counsel in the relevant jurisdiction before any reliance.**

Appendix A. Reporting standards

A.1 The compartment register. Any estate measured under the contagion instrument reports the following, and nothing less.

```
the compartment map: entity, jurisdiction, assets, purpose
the reachable share of each compartment, summing to 100.0%
contagion HHI, effective compartments, maximum single reach
the absorbable-loss constraint and whether it is satisfied
annual cost of the compartmentalisation, itemised
the estimated p and the break-even p
for any cell structure, the forum analysis and the r estimate
```

```
 8  the guarantee, comfort-letter and cross-default register --
    every instrument that connects two compartments
 9  confirmation, by someone independent of the treasury function,
    that no commingling occurred in the period
10  the four-axis jurisdictional measurement, with the note that
    compartmentalisation contributes nothing to N
11  the date, and the difference against the previous measurement
```

Item 8 is the one nobody maintains and the one that decides the outcome. A group cannot know whether its compartments hold unless it knows every instrument that connects them, and the instruments that connect them are signed by people who do not know the structure exists. **Item 11 decays first**, because the contagion index is only meaningful as a difference against last year's.

A.2 The conversion memorandum. Any conversion of an immovable asset into a movable claim reports: the threat it is being used against, named; the leakage L computed at full precision on the actual basis and the actual rates; the escape value with p, L^{*}, q and s written down and defended; the payback against the holding horizon; the concentration instrument run on all four axes before and after; and the custody chain specified in writing down to the depository, **because the chain is the new single point of failure and a conversion executed without one has bought a change of governing law and then let the counterparty choose it.**

Appendix B. Claims requiring verification against primary sources

The following statements rest on legal or regulatory detail that changes, and **none of them is cited to the bibliography, which contains no primary sources for any individual jurisdiction and no case law**. They are reproduced from the source manuscript and require confirmation with counsel in the relevant jurisdiction before any reliance. The source manuscript's own verification registers, which run to seventy numbered items across the three chapters, are condensed here without loss of subject.

1. **Separateness and its limits.** The foundational company-law authority on separate legal personality; the concealment/evasion distinction and its later application; the rejection of a general single-economic-unit doctrine for groups; and the named appellate formulations of the substantive-consolidation standard, which is circuit-dependent.
2. **Bankruptcy remoteness in practice.** The docket, procedural posture and precise holdings of the leading case on the filing of bankruptcy-remote property-level entities alongside their parent, including the treatment of replaced independent managers and the scope of those directors' duties. **This paper states a structural lesson and asserts no holding.**
3. **Contractual insolvency limitations.** Enforceability of non-petition and limited-recourse covenants in every jurisdiction where a proceeding could be opened; and the two opposed judicial answers to the same priority-flip clause in the same transaction, together with the full subsequent procedural history.
4. **Cell legislation.** The originating protected-cell ordinance and its successor provisions; segregated-portfolio, segregated-account and incorporated-cell regimes; the series limited-liability provisions of one American state and their record-keeping and notice conditions; and the fund-specific sub-fund wrappers. **None is asserted here beyond the existence of a regime.**
5. **Judicial consideration of cell segregation in a foreign forum** — whether authority exists in any major jurisdiction on recognition of a foreign statutory cell partition. **This is the single most valuable item in the register and the paper treats the question as open.**

6. **Orphan structures.** Non-charitable purpose trust legislation and the special trusts regime relied on; the enforcer requirement; and the control-based consolidation criteria and variable-interest-entity model that determine whether an orphan is consolidated.
7. **Avoidance law.** The transaction-at-an-undervalue and transactions-defrauding- creditors provisions of the English insolvency legislation, the federal American provision, and the uniform state act as enacted, with their elements, look-back periods and connected-person presumptions. **This is the eleventh failure mode and it is a clock.**
8. **Risk retention and its limits.** The 5 per cent retention requirements in both major regimes, the qualifying exemptions, the post-onshoring position, and the decision holding open-market collateralised loan obligation managers outside the American rule.
9. **The situs and intermediated-securities rules.** The rule that an English court will not entertain an action to determine title to foreign land; the exclusive-jurisdiction provision of the recast European regulation and its post-withdrawal application; the place-of-the-relevant-intermediary convention with its date of entry into force and its short contracting-party list; the settlement-finality and financial-collateral directive articles and their onshored equivalents; and the American securities intermediary's jurisdiction rule.
10. **Listed real estate vehicle regimes.** The originating legislation and current code in each of the two systems modelled, the distribution requirement, the balance-of- business conditions, the closely-held tests and their attribution rules, the close-company condition and its institutional-investor exception, and the listing or trading requirement with its later relaxation. **Do not rely on any threshold stated in this paper.**
11. **Non-recognition on contribution.** The partnership non-recognition provision, the allocation of built-in gain, and their application to an umbrella-partnership structure; the attribution of the first such structure to a 1992 offering; and, jurisdiction by jurisdiction, **whether any relief, clearance or rollover exists for a contribution of appreciated real property to a listed vehicle — the single most valuable question in Section 6.**
12. **Transfer taxes and rates.** The transfer-tax code and current rates for commercial property and corporate purchasers; the seeding relief for authorised property funds and whether it extends to listed real estate vehicles; **the 25 per cent corporation tax rate used in the leakage computation**, the availability of indexation, losses and rollover, and the correct rate for the entity actually holding the asset.
13. **Non-resident gains on land** and indirect disposals of property-rich entities: the gross-asset-value and investment thresholds, the substantial-shareholding and trading exemptions, and the current location of the rules.
14. **Disclosure regimes.** The annual-report property-schedule rule and its current content; notification thresholds and deadlines in both markets; and the overseas- entities register, its launch and registration deadline and its later amendments.
15. **Territorial taxation regimes and the hemispheric enforcement environment**, as described in Section 7 and Appendix C: the source-rule provisions and their article numbers; residence and investment thresholds, several of which are inflation-indexed and move continuously; the substance and qualifying-shareholding conditions introduced under listing pressure; free-zone statutes; beneficial-ownership regimes and their access rules; current standards-body status, mutual-evaluation ratings and information-exchange commencement dates; and the two separate European lists, which have different criteria, different consequences and different memberships.
16. **Named political and judicial events** referred to descriptively in Appendix C: a constitutional ruling voiding a signed concession; designations of senior political figures under a corruption authority, with dates, scope and current status; an investor-state arbitration and its award; a sovereign-default judgment and its appellate position; and the current convention status of states that have denounced the principal investor-state arbitration instrument.
17. **The historical branch structure** relied on in Section 9.1: the constitution of the branches as separate partnerships, the capital arrangements between the parent partnership and the branches, and the sovereign-lending losses that brought the whole institution down

notwithstanding the separation [de Roover, 1963].

18. **All cost figures are index units in a constructed model.** They are not fee quotations and no real fee schedule was used to produce them.

All numerical results in this paper were computed from the stated inputs at full precision and rounded only for printing.

Appendix C. Descriptive jurisdiction profiles

These profiles are reproduced from the source manuscript. Not one statement in them is cited, because the bibliography contains no primary source for any individual jurisdiction (Section 11 and Appendix B item 15). They are description and not assessment, they are certainly out of date in places, and they are retained because the measurement in Section 8.5 is unintelligible without knowing what was being measured. **The scored comparison built from them is debranded in Section 7.2 for exactly this reason.** Item 10 in each profile — the degradation trajectory — is the item the source treats as decisive and the item most exposed to change.

C.1 A dollarised isthmus centre with a deep vehicle set

Civil law. An early twentieth-century corporations statute drafted with reference to a leading American state's company law of the period; a mid-1990s private-interest foundation with separate legal personality on the Liechtenstein model; and a 2017 trusts statute. An unusually deep vehicle set for a civil-law jurisdiction. **Bloc alignment:** dollarised with no lender of last resort, dollar clearing through foreign correspondents, and a foreign-policy reversal under external pressure within weeks in 2025. **Whatever the tax code says, this is not an independent perimeter. Courts:** genuinely two-sided — a constitutional ruling in 2023 voided a mining concession law against both the executive and the legislature, which demonstrates real independence and demonstrates that a contractual position can be extinguished by constitutional litigation. Commercial docket delay is significant; arbitrate rather than litigate. **Substance and disclosure:** resident-agent and accounting-record obligations, beneficial-ownership registration accessible to competent authorities and not to the public, full participation in automatic exchange. **Trajectory:** improving on standards from a poor base and structurally exposed to recurrence — three listing cycles since 2000, removal from the principal list in October 2023 after a listing that began in June 2019, and a services-dependent economy permanently within reach of the standards mechanism. **A thirty-year structure here should assume at least one further listing cycle and be built to survive it. Assigned function:** operating, and in a limited regional sense bridging. Not holding.

C.2 A Southern Cone residence and free-zone centre

Civil law on a nineteenth-century code with modern commercial and free-zone statutes; a statutory trust exists in civil-law form and its suitability as a holding instrument should be tested before use. **Bloc alignment:** the least coupled jurisdiction in the group — own currency and central bank, a regional trade bloc rather than a bilateral framework, and a sustained independent push for a trade agreement with a rival power against its own bloc's rules. **This is the only jurisdiction in the group with a credible claim to partial perimeter independence, and Section 8.5 prices it as partial rather than full. Courts:** independent, low corruption, and a demonstrated willingness to litigate an investor-state arbitration to a final award; deduct for commercial docket speed and a small body of international commercial jurisprudence. **Substance:** real for the free-zone regime — qualifying activity, employment ratios, a substantive operating presence — and real for tax residence, on days of presence or an investment threshold combined with presence or job creation. **Trajectory:** stable, with the transparency degradation already completed and the institutional position untouched; the former secrecy proposition was dismantled deliberately by domestic legislation between 2007 and 2017 rather than

under listing pressure. **Assigned function:** operating and personal residence with a genuine free-zone services capability. Not a holding apex.

C.3 A landlocked manufacturing centre with an unstateable standards position

Civil law on a 1985 code; a conventional company law and no distinctive holding or foundation vehicle — the weakest vehicle set in the group, which is consistent with its correct use. **Bloc alignment:** non-aligned in form and pre-committed in substance; own currency and central bank, regional bloc membership, and a diplomatic recognition that forecloses one bloc entirely and is a standing object of pressure to reverse. **Courts:** the weakest in the group — limited institutional capacity, poor external corruption scores, little published international commercial jurisprudence. Arbitration outside the jurisdiction is the only reliable mechanism for a significant contract. **Substance:** the export-processing regime requires a genuine production operation with approved programmes; the substance here is industrial rather than regulatory, and that is the jurisdiction's honest strength. **Disclosure:** the automatic-exchange commencement date and current status **cannot be stated from the source and the source declines to guess. Trajectory:** uncertain, and the uncertainty is itself the finding — a coherent 2019 tax restructuring, a modernised 2022 migration framework and very low energy cost on one side; an unstateable standards position and continuing corruption designations of senior political figures on the other. **A jurisdiction whose standards status cannot be stated is one whose trajectory cannot be scored. Assigned function:** operating only, and specifically export manufacturing.

C.4 An isthmus centre that joined the perimeter deliberately

Civil law with a 1949 constitutional order and a strong rights-based constitutional jurisprudence; conventional company law, no distinctive holding vehicle, a 1990 free-zone statute. **Bloc alignment:** western, by deliberate choice — accession to the principal economic organisation on 25 May 2021. **That is a decision to enter the coordinating perimeter, and it is simultaneously an institutional-quality signal and the elimination of any bloc-diversity claim. Courts:** a constitutional chamber created in 1989 that is among the region's most active and respected, an independent ordinary judiciary, and no expropriation record; deduct for commercial docket speed. **Substance:** real and recently tightened — the free-zone regime requires qualifying activity, investment and employment, and a 2023 amendment introduced substance and qualifying-shareholding conditions for the exemption of passive foreign-source income. **Trajectory: the cleanest observed case of the listing mechanism in the source.** Accession in May 2021; addition to a European non-cooperative list in February 2023 over the foreign-source exemption; domestic legislation introducing substance conditions in 2023; removal from the list in November 2023. **No sanction was imposed and no negotiation of substance occurred — a listing was made and a sovereign rewrote its income tax law within nine months. Assigned function:** operating, at genuine quality, and personal residence. Not holding; not bridging, by its own choice.

C.5 The jurisdictions the source declines to score, and why

Four large regional economies are outside the subject, being investment cases rather than residence-and-source architectures. Three further states are excluded on political risk: one with a sovereign default, a decade of subsequent litigation abroad, an expropriation of a majority stake in a national oil company producing a very large adverse judgment, repeated capital and exchange controls, and inflation that has repeatedly destroyed the domestic value of contractual claims; one requiring no argument; and two that denounced the principal investor-state arbitration convention, which removes the most reliable enforcement route.

None of that is a reason to avoid operating business in those states; it is a decisive reason not to hold protected capital there.

Chapter 61

Measuring the Payment Rail: Clearing Perimeters, Transit Exposure Windows, and the Concentration Effect of De-Risking

Abstract. Payment infrastructure is normally audited by counting correspondent banks. We show that this count measures the wrong object. On a constructed six-bank corridor set, the naive correspondent Herfindahl-Hirschman index of 0.302400 reports 3.306878 effective correspondents; decomposing the same flows by settlement currency and then by clearing perimeter gives a perimeter HHI of 0.745256 and only 1.341819 effective perimeters, and the rail-axis surplus term is 0.35 x ln(1.341819) - 0.500000 = -0.397091. The remedy that people reach for does not work: adding correspondents raises the naive effective count to 4.425092, an improvement of 33.8 per cent, and leaves the perimeter count unchanged to six decimal places. Currency redenomination over three years does work, taking the perimeter count to 2.087683 and the surplus term to -0.242381, with the dollar share floored near 52 per cent because commodity and trade-finance pricing is not the group's to set. For value in transit we derive E[loss] = V(1 - exp(-lambda t))(1 - r) and show expected loss is invariant to hop count (USD 68,000 whether one hop of 40m or eight of 5m) while depending entirely on the exposure window; multilateral netting removes 77 per cent of expected transit loss without a new rail. Informal value transfer is analysed as a licensed, priced industry: enforcement-driven exit from twelve providers to five moves the HHI from 1,074 to 2,406 points and cuts effective providers from 9.310987 to 4.156276.

Keywords: correspondent banking; clearing perimeter; Herfindahl–Hirschman index; settlement finality; transit exposure window; de-risking; informal value transfer; currency redenomination

Series: D. Structure and Resilience **Paper:** CCP-061

1. Introduction

A cross-border holding structure that measures its jurisdictions, its counterparties and its currencies and stops there has measured everything except the layer through which all of it has to pass. This paper measures that layer.

The result is that **the settlement rail is where a well-designed cross-bloc structure is most concentrated, and it is the axis nobody runs.** On a constructed corridor set, five correspondent banks in four jurisdictions across three coordination perimeters return an effective clearing-perimeter count of 1.341819 against an effective bloc count of 2.941176 on the jurisdictional axis. The rail-axis surplus term is $0.35\ln(1.341819) - 0.500000 = -0.397091$ against -0.122417 on the axis the structure does measure. **The structure is 3.24 times worse on the rail than on the jurisdiction, and the gap was invisible until the instrument was pointed at it.**

The remedy is not more banks. Adding correspondents raises the naive effective count by 33.8 per cent and leaves the perimeter count unchanged to six decimal places, because **the rail is determined by the currency of the underlying obligation.** Fixing it is a commercial negotiation about invoicing rather than a structuring decision about accounts, and the constructed group that understands this recovers 0.154710 of surplus balance without forming a single new entity.

1.1 A layering correction the field owes itself

A clearing system and a messaging system are different layers of the same stack and are not substitutes. A secure financial-messaging cooperative moves instructions between institutions and settles nothing; a national real-time gross settlement or netting system settles and does not provide interbank messaging to participants who reach it indirectly. **"System X replaces the messaging network" is a category error, and a structure built on that premise will fail.** The correct comparison for a clearing system is another clearing system; the correct comparison for a messaging network is another messaging network, or a bilateral connection. The layering is not folklore: the message-standard migration that the whole industry is executing is a standard for the *messaging* layer alone [International Organization for Standardization, 2013], and it changes nothing about where finality occurs.

The consequence is arithmetic. A treasury policy recording "two rails" and treating that as diversification has recorded one messaging dependency twice. **Measured correctly, one messaging network at 100 per cent and 85 per cent on that network plus 15 per cent on a second clearing system whose messaging leg runs over the first are the same number, and only the second looks diversified.** Curzi [2026] states the general form of that error for jurisdictions; Section 5 performs the rail version of the measurement properly.

1.2 The event that made this a board-level question

In 2022 approximately USD 300 billion of central bank reserves were rendered untransferable by coordinated administrative action in the jurisdictions where they were held. Curzi [2026] states the fact and its consequence — **disclosure is not immunity: the reserves were fully disclosed and the disclosure test is a filter that removes structures collapsing on contact with an authority, not a shield against a policy decision** — and this paper does not embellish it.

Three structural features of that episode drive what follows. The immobilisation was achieved without physical possession and without a court: the assets were book entries and the measures were administrative. The holder was a sovereign, with sovereign immunity [Dixon, 2013], in jurisdictions with independent courts and strong property protection — the strongest legal position short of holding bullion in one's own vaults — **and the legal position did not matter, because the question was never adjudicated. A holder whose claim can only be realised through an intermediary has the protection the intermediary's regulator permits, not the protection its legal position suggests.** And a very large fraction of the immobilised assets sat with a single institution in a single jurisdiction: whatever the effective jurisdiction count of that portfolio by currency and issuer, its effective *custodian* count was close to one, and that was the operative number.

This paper draws no conclusion about the merits of any measure. It takes a position only on the engineering: **a settlement and custody architecture whose effective independent count**

is one will fail completely when it fails at all, and there is no legal argument that repairs a number.

1.3 What the paper contributes, and the standing constraint

Every architecture examined here is measured under the corpus's formation test: a structure qualifies only if it would still perform its function were every element disclosed to every authority with jurisdiction over any part of it. On this axis the test is not a constraint the analyst imposes but a description of the layer. **The entire rail architecture consists of relationships with regulated institutions who are required to know exactly what the group is doing. There is no version of a rail that is not disclosed, so a rail strategy premised on non-disclosure is not a strategy but a misunderstanding of the layer.** The open-design principle [Saltzer and Schroeder, 1975], stated a century earlier as the rule that a system must remain secure when everything about it except the key is public [Kerckhoffs, 1883], is the same constraint in security engineering.

Six results follow: the clearing-perimeter measurement and its collapse; the failure of correspondent addition and the success of currency redenomination; a transit-loss model whose expectation is invariant to hop count; multilateral netting as the highest-return intervention available; an indifference hazard separating a fast rail from a slow one; and the concentration effect of enforcement-driven exit on a licensed remittance corridor. The paper presupposes the founding reframing of Curzi [2026], on which a present arrangement is a local minimum of the cost landscape by definition rather than by luck — which is why a rail that nobody has measured is nonetheless the arrangement everybody is in.

2. Hypotheses

Four hypotheses are stated, each falsifiable in the sense Platt [1964] requires and each carrying a point prediction with a constant attached rather than a direction, which is what Meehl [1967] asks of a theory that wishes to be tested rather than confirmed. Statements that could not be falsified as posed are listed afterwards as definitions or modelling choices.

H1 — the clearing-perimeter hypothesis. Correspondent count measures the wrong object, because settlement independence is determined by the perimeter in which a currency achieves finality and not by the number of banks holding accounts. *Predicted:* on the constructed corridor set, decomposing the same flows by currency and then by clearing perimeter takes the effective count from 3.306878 to 1.341819, and the rail-axis surplus term to -0.397091. *Falsified by:* a corridor set in which the perimeter decomposition returns a count at or above the naive correspondent count, which would mean the currency layer carries no concentration the bank layer does not already show.

H2 — the invariance-to-accounts hypothesis. Adding correspondent relationships cannot move the clearing-perimeter count, because it does not change the denomination of the underlying obligations. *Predicted:* adding three banks in three further jurisdictions raises the naive effective count from 3.306878 to 4.425092, an improvement of 33.8 per cent, and leaves $N_{\text{rail}} = 1.341819$ **unchanged to six decimal places**; redenominating the underlying obligations takes it to 2.087683. *Falsified by:* any account-level rearrangement that moves the perimeter count while the currency mix is held fixed.

H3 — the hop-invariance hypothesis. Expected loss on value in transit is invariant to how a total is divided, and depends only on the total, the per-hop exposure window, the hazard rate and the recovery fraction. *Predicted:* with $V = 40{,}000{,}000$, a six-hour window, a per-hop freeze probability of 0.2000 per cent and a recovery fraction of 0.15, expected loss is **USD 68,000 whether the total moves as one hop of 40m or as eight hops of 5m**, while the probability of some loss rises from 0.2000 to 1.5889 per cent and the worst single loss falls from 34.0m to 4.25m. *Falsified by:* any expected-loss figure that depends on hop size under independent per-hop freezes — which would mean the derivation in Section 3.4 is wrong.

H4 — the enforcement-concentration hypothesis. Compliance-driven exit concentrates a licensed market by the same arithmetic that a merger would, and on the same published screen. *Predicted:* exit from twelve licensed providers to five moves the index from 1,074 to 2,406 points and the effective provider count from 9.310987 to 4.156276, clearing the published structural-presumption threshold of 1,800 points by 606 and the change threshold of 100 points by a factor of thirteen. *Falsified by:* a de-risking episode after which the surviving providers' shares are more equal than before, so that the index falls.

Not hypotheses, and labelled as such.

- *The identification of N_{rail} with the reciprocal of the clearing-perimeter index* is a **definition**. Its content is entirely in the share basis and the perimeter assignment, both of which are testable.
- *The mapping of each currency to a single clearing perimeter* is a **modelling choice**, and one assignment in particular — treating the euro area as inside the Western perimeter — is contestable and is recorded as such in Appendix B.
- *The hazard rate λ*, the recovery fraction r, the per-hop cost, the covenant tolerance and every share vector are **stipulated constructed estimates**. Section 11 states that λ is the most load-bearing and least defensible input in the paper.
- *The parameter set $\eta = 0.35$* and $\delta Q + C = 0.500000$ is a **modelling choice**, fixed in advance so that no criterion is adjusted after the measurement — the preregistration discipline of Nosek et al. [2018] applied to an engineering calculation.

3. Method

3.1 The stack, and the four share bases

Payments run through layers that fail separately, and an index computed on the wrong layer measures the wrong thing.

layer	what it does	what fails
messaging	carries the instruction between institutions	disconnection of an institution from the network; a standards migration
correspondent	holds the account through which a bank reaches a currency it has no local presence in	contractual termination, with notice and no reasons
clearing and settlement	achieves finality in a currency, inside one jurisdiction's perimeter	an administrative measure in that jurisdiction
custody and depository	holds the security at the end of an intermediated chain	a designation reaching the depository

Four share bases follow, and they must never be mixed inside one computation.

basis	share of what	the count it yields
correspondent	cross-border settlement value passing through a given bank	effective correspondents
currency	settlement value denominated in a given unit	effective currencies

basis	share of what	the count it yields
clearing perimeter	settlement value achieving finality inside a given perimeter	N_{rail}, effective clearing perimeters
transit	value concurrently in flight on a given rail	effective rails

The correspondent basis is the one everybody uses and the clearing-perimeter basis is the one that answers the question. They differ because a bank's own jurisdiction is not where the money it moves achieves finality: a Gulf bank settling a dollar obligation is a dollar-perimeter exposure however the bank is chartered.

3.2 The instrument, and which effective count it is

For shares s_i summing to one on a stated basis,

$$\text{HHI} = \sum_i s_i^2, \qquad \text{effective count} = \frac{1}{\text{HHI}}.$$

If k units each hold $1/k$ then the reciprocal returns k; when they are unequal it returns the number of equal-sized units giving the same index. That numbers-equivalent reading is Adelman [1969]'s; the index itself is Hirschman [1945]'s and, independently in sum-of-squares form, Herfindahl [1950]'s. Competition practice states it in points on percentage shares, where a structural presumption attaches above 1,800 points with a change above 100 [U.S. Department of Justice and Federal Trade Commission, 2023]; the resilience application states it as a fraction. **Both conventions appear in this paper and every figure is labelled with which.**

Which effective count. $1/\text{HHI}$ is the Hill number of order two [Hill, 1973] — the numbers-equivalent that Jost [2006] shows is the diversity, as distinct from the index — and $\exp(H)$ for Shannon entropy [Shannon, 1948] is the same family at order one. **The corpus's other effective count, the affine $N_{\text{eff}} = 1 + (N-1)u$ of Curzi [2026c], is a linear interpolation between 1 and N, is not a Hill number of any order, and carries no external support.** That paper badges it as a design axiom rather than a theorem — a hundred-person department at $u = 0.05$ reading an effective count of 5.95 — and this paper does not use it. **Every count reported here is $1/\text{HHI}$ on a stated share basis.** Hurlbert [1971] is the discipline being imposed: the word carries no content without its formula attached.

No distributional or scaling claim is made anywhere in this paper. Correspondent shares descending 46 / 24 / 14 / 10 / 6, and provider shares descending across twelve entries, look like the opening of a heavy-tail argument, and the payments literature is full of such arguments. **No distribution is fitted, no tail is characterised and no power law is asserted.** Clauset et al. [2009] is the reason the procedure is named and not run: a principled maximum-likelihood fit with a goodness-of-fit test and a likelihood-ratio comparison against alternatives leaves few canonical data sets standing, and a straight line on log-log axes is not evidence. Stumpf and Porter [2012] adds why passing it would not be enough: many mechanisms produce the same tail. **Every effective count here is a statistic of a finite share vector with between two and twelve entries, not an estimate of a distribution, and it must not be restated as one.**

Small-sample estimation is not addressed: Chao et al. [2014] show that low-order diversity estimators are biased on small samples, and these vectors are small. Golosov [2010] documents that effective-number measures behave poorly in the dominated-unit regime, which is the regime of the post-exit corridor at 34 per cent in one provider.

3.3 The correlation correction, and what a second account cannot buy

The naive index assumes the units fail independently. Where they do not, the correction is the one Markowitz [1952] stated for portfolios: diversification benefit depends on covariance, not on count. Writing w for the share vector and R for the correlation matrix of unit-variance shocks, the aggregate variance is $w'Rw$ and the general effective count is $1/(w'Rw)$; $R = I$ recovers the naive index and a block matrix of ones recovers the perimeter index. Curzi [2026] establishes that identity and its correspondence with the similarity-sensitive diversity of order two [Leinster and Cobbold, 2012].

On the rail axis the correlation is not a modelling nicety; it is the mechanism. Two banks whose dollar legs both achieve finality in the same system are perfectly correlated with respect to a measure taken in that system, whatever their charters say. That coupled systems are far more fragile than their components is the formal content of Buldyrev et al. [2010], used here as a **shared structural form and not as a derivation**: no percolation premise has been shown to hold for a correspondent network. Albert et al. [2000]'s robust-yet-fragile result is used the same way, for the shape of the tolerance asymmetry only, and Doyle et al. [2005] is the standing correction to reading too much into either — engineered robustness comes from design, not from a degree distribution. Perrow [1984] reached the organisational version first, and is cited as framing rather than as a formal result.

Correlation rises under stress, which is when the count matters. Longin and Solnik [2001] measure extreme correlation that is asymmetric and concentrated in downturns; the ecological name for the same variable is synchrony [Loreau and de Mazancourt, 2008]. The contrary literature is not decorative: Doak et al. [1998] show that diversity–stability correlations can arise from statistical averaging alone with no mechanism, and May [1972] that increasing diversity and connectance destabilises randomly assembled systems. This paper's claim is narrower and is not exposed to those results in the same way: it asserts only that the reciprocal of a variance counts independent shocks, which is arithmetic.

A count of names is not a count of independent parties. Douceur [2002] proves that without a trusted identity authority one party can present as many, so an unaudited count carries no information about independence. Nesting is the payments instance: a respondent may itself provide correspondent services to further banks, so an institution's exposure runs to counterparties it has no relationship with. **The corpus carries an exact arithmetic instance of the same gap: Curzi [2026] proves that a roster indexed $1 \dots N$ realises exactly $\lceil N/2 \rceil$ distinct classes under its folding rule — six at $N = 12$, from twelve nominal members.** In all three settings a count of names and a count of distinct things are different quantities.

3.4 The exposure-window model, derived

Value in transit is exposed for a period. Model the freeze as a hazard: over an exposure of duration t, the probability of a freeze event affecting that value is

$$q = 1 - e^{-\lambda t}$$

for hazard rate λ per unit time; for small λt this is approximately λt, which is adequate throughout. Split a total V into n hops of size k, so $V = nk$, each exposed for t and frozen independently with probability q. If a frozen hop eventually recovers a fraction r, discounted for delay, then

$$\mathbb{E}[\text{loss}] = nkq(1-r) = Vq(1-r) = V\left(1 - e^{-\lambda t}\right)(1-r).$$

k **cancels.** Expected loss is a function of the total value, the per-hop window, the hazard rate and the recovery fraction, and is not a function of how the total is divided. What n controls is the distribution: losses are $k(1-r)$ times a Binomial(n, q) draw, so the maximum single-event loss $k(1-r)$ falls with n, the probability of any loss $1-(1-q)^n$ rises with n, and the variance falls with n. **Three levers, three different effects, and only one of them touches the expectation.**

The exposure the model prices is a stock produced by a flow and a delay, which is Little's law and is cited as a shared mathematical form rather than as a derivation. Little [1961] establishes that the average number in a system equals the arrival rate times the average time in system, under stationarity assumptions this paper's constructed corridors do not verify. The structural content carried over is that **concurrent transit exposure is set by the window and not by the parcel size**, which is the same statement the algebra above reaches independently.

Two properties of the model are load-bearing and are stated as limitations of it rather than discovered later. **Independence across hops fails for systemic events** — an outage, a depository suspension, a rail halt freezes every hop at once — so the binomial overstates the diversification exactly in the scenario that produces the tail. And λ decomposes into a counterparty component reducible by screening, a jurisdictional component reducible by routing but bounded by the currency, an operational component reducible by data quality, and a systemic component reducible by nothing the payer controls.

3.5 Landed cost and the indifference hazard

The cost of a rail is not its fee. It is the fee, plus the spreads, plus the expected freeze loss, and the last term dominates wherever the question is interesting:

$$\text{landed} = \text{explicit} + V\,(1 - e^{-\lambda t})\,(1 - r).$$

Setting two rails' landed costs equal and solving for λ gives an **indifference hazard**: the freeze probability above which the faster rail is worth its higher explicit cost. That converts an unanswerable geopolitical question into a threshold belief, which is the corpus's standard move for a marginal decision. Curzi [2026] prices a pledge trigger the same way, showing that at an opening loan-to-value of 50 per cent and a maintenance level of 35 per cent the first margin call falls at a 23.076923 per cent drawdown rather than at 50 per cent; Curzi [2026] prices a bridging position as a break-even probability, refusing the largest available gain on tail exposure and taking a capped transit position that yields $+0.002869$ — 3.9692 per cent of the benefit for 3.2110 per cent of the exposure, removing 96.79 per cent of the tail.

The same move appears in the corpus at its most abstract, and naming it as a family prevents its being reinvented. Curzi [2026] reduces an entire relational sustainability question to one comparison, establishing the trichotomy $f \gtreqless f_{\min}$ corresponding to $Q_\infty \gtreqless 0$ with $f_{\min} = k\,e^{2c/\eta}$. **The constant there, $k = (9/8)^2/\varphi^2 = 0.483425732988414$, is a *defined* constant — an exact ratio of exact quantities, which is what licenses its fifteen printed digits — and attaching a unit of hertz to it is a separate and much stronger claim which that paper badges as a postulate. This paper takes no position on the unit.** Since φ enters the definition, Markowsky [1992] is the standing corrective on golden-ratio claims and is cited for that reason. Curzi [2026] supplies the accompanying warning about compounding two such rules: its Corollary shows the compound rule and the essence law agree at no integer index, since $kn^2 = k^2n^4$ requires $n = 1/\sqrt{k} = 1.438252434444351$, and at $n = 12$ the two differ by 6.121291 octaves. **Thresholds derived from different laws must not be composed without checking that they meet anywhere.**

3.6 Identification, and the status of the perimeter assignment

Perimeter assignment is an input to the instrument and carries an identification problem that must be stated where it is used. Treating a set of jurisdictions as a coordination perimeter whose members respond to one decision is a statement about a group influencing its members while the members constitute the group. **Manski [1993] proves that the endogenous effect in a linear-in-means reflexive specification is not identified from observational data**: it cannot be separated from contextual effects or from correlated unobservables however much data is collected. No amount of observation of clearing behaviour could confirm a perimeter assignment derived reflexively.

The remedy adopted is stipulation with a defence: perimeter membership is fixed in advance by the coordination test — two jurisdictions are in one perimeter when a single decision taken in one reliably produces the corresponding effect in the other without a further independent decision — recorded with its evidence, and the one contestable call is flagged and its alternative computed. Brock and Durlauf [2001] shows that route can be made rigorous. **The price is that every perimeter figure in this paper is a modelling position and not an identified empirical result.**

The collapse mechanism has a formal ancestry and is not new machinery. That several units inside one perimeter behave as one under a correlated shock is the threshold-and-tipping result: Granovetter [1978] establishes that the distribution of individual thresholds rather than the mean determines whether a collective outcome tips, and Schelling [1971] is the canonical demonstration that a mild individual preference produces a sharp aggregate outcome — **the strongest available warning against reading an individual-level parameter straight into an aggregate, which is exactly what a perimeter assignment does.**

One further reflexivity is specific to this axis. De-risking is a decision taken because other institutions are taking it, which is an informational cascade in the sense of Bikhchandani et al. [1992] and Banerjee [1992]: agents rationally ignore their private information once enough predecessors have acted, so the exit is neither independent across banks nor informative about the underlying risk. **A correlated exit is therefore the expected behaviour of the correspondent layer and not an anomaly**, and Merton [1948]'s self-fulfilling prophecy is the older name for a corridor becoming unbankable because it has been described as unbankable.

3.7 Sources of measurement that this paper does not have

Every figure below is computed from stipulated inputs. **No dataset was collected, no corridor was observed, and no institution was surveyed.** Where the source manuscript refers to published measurements — the decline in active correspondent relationships since 2011, the regional cost of remittances, the resolution amounts in two enforcement actions — those are described and marked for primary-source verification in Appendix B rather than quoted as this paper's evidence.

4. The correspondent layer and why it fails

4.1 The mechanics, and the economics of exit

A bank without a presence in a currency's home jurisdiction reaches that currency by holding an account at a bank that does, and that bank may reach the ultimate clearing system directly or through a further correspondent. The chain can be two links or five. Three features govern everything downstream.

Access is contractual, not regulated. No bank has a right to a correspondent account. Termination requires notice under the contract and nothing else: no appeal, no regulator to petition, typically no reason given. This is the property that makes the layer a decision surface rather than a legal surface, and it is why Hirschman [1970]'s distinction is the right frame — the respondent has no voice, so the correspondent's only instrument is exit, and exit is what it uses.

The decision is commercial. The correspondent bears regulatory risk for its respondent's customers, whom it does not know, in a jurisdiction it may not operate in, for a fee measured in basis points. Where the revenue from a difficult corridor is measured in tens of thousands and the penalty exposure from one failure in hundreds of millions, no rational institution retains the relationship, and none did. The two best-known resolutions of the last fifteen years are described in Appendix B and are not quoted here as figures.

The relationship nests. A respondent may itself provide correspondent services to further banks, so the correspondent's exposure runs to institutions it has no relationship with at all.

Nesting is the mechanism by which a measured count of relationships overstates the count of independent parties, and it is the payments instance of Douceur [2002].

The consequence for a cross-bloc structure is stated plainly. **Fewer correspondents means the survivors are larger, which raises the concentration on the axis this paper measures. A cross-bloc structure with a bridging apex and a dual-national principal is precisely the profile the survivors screen hardest. And the exit, when it comes, arrives as a letter with thirty days' notice and no reasons.** The diligence interface is a standard questionnaire and a pair of international recommendations on correspondent banking and on wire transfers [Financial Action Task Force, 2012]; a group that cannot answer the questionnaire quickly and consistently across all its entities will be de-risked, and the questionnaire is answerable only if the structure was built to be explicable. That is the formation test arriving at the operations desk.

4.2 Finality, and what it does not mean

One further mechanism decides what "the payment went through" means. The classic failure — one leg of a currency exchange settling while the other does not, because the two currencies settle in different time zones — produced both a market answer, simultaneous payment-versus-payment settlement, and a legal answer, settlement-finality legislation under which transfer orders entering a designated system become irrevocable at a specified moment and are protected from the retroactive effect of insolvency. The instruments are described in Appendix B and are not cited.

Finality is a legal concept about the irrevocability of an entry. It is not a statement about whether the resulting balance can be used, and the 2022 precedent is the demonstration: the transfers were final, the entries were valid, and the balances were unusable. The distributed-systems literature supplies the exact vocabulary for why that gap exists and is not closable by drafting. **Agreement among parties that may fail cannot be guaranteed in an asynchronous system with even one faulty participant** [Fischer et al., 1985], and a system that must remain available under partition cannot also remain consistent [Gilbert and Lynch, 2002]. **Both are used as shared structural form and not as derivations** — a payment system is not a formal asynchronous consensus protocol — but they name the trade-off that finality rules resolve by fiat rather than by proof. Lamport et al. [1982]'s bound on the number of genuinely independent parties required to tolerate faulty ones is the same discipline applied to counting: **redundancy is only redundancy if the parties fail independently.**

4.3 The depository layer

Securities sit in an intermediated chain terminating at a central securities depository, or at an international depository for cross-border holdings, and there are very few of the latter. That is not an accident of market structure but a consequence of network economics: settlement efficiency rises with the fraction of counterparties in the same book, so the market converges on a small number of very large books.

The custody axis therefore has an effective count in the low single digits for a whole market, and a structure holding securities through it inherits that count regardless of how many custodians it contracts with, because its custodians hold at the same place. The measurement improves at every layer a group contracts with and collapses at the layer it does not — which is the shape of the 2022 precedent. **A concentration that cannot be fixed should be sized, not ignored, and sizing it requires measuring it first.**

5. Transit: what the hazard is made of

5.1 Transit and store, distinguished

A rail is a mechanism for moving value; a store is a place value sits. **A prefunded account is not a rail. It is a holding that makes a rail faster, and it must be measured on the holding axis**, because a settlement balance that exceeds its limit has silently become a position on the jurisdictional measurement. The control that prevents this is a balance limit enforced independently of the treasury function; without it the distinction erodes by drift rather than by decision.

5.2 The hop-splitting question, and what it actually buys

The derivation of Section 3.4 says expected loss is flat in hop size. What splitting buys is a reduction in the worst case, at a stated premium in per-hop cost. Whether that premium is worth paying is a question about the group's capacity to absorb the worst case, not a question about risk reduction, and the right frame is an absorbable-loss constraint:

$$k \le \frac{X}{1-r}, \qquad n \ge \frac{V(1-r)}{X}$$

for a tolerance X. **The constraint sets the hop size; the optimisation does not, because there is nothing to optimise — expected loss is flat in k.** That is structurally the same result the corpus reaches for liability compartments in Curzi [2026], whose contagion measurement gives 4.96 effective compartments from six unequal ones and cuts maximum single reach from 100.0 to 30.0 per cent at a break-even event probability of 0.157 per cent a year, and it arrives here from a different derivation.

One caution is decisive and is stated here rather than in a footnote. Splitting across time provides no protection against a systemic event, which freezes every hop at once. **Split hops across *rails*, not merely across time, or the tail protection the split was bought for does not exist in the scenario that produces the tail.**

5.3 Netting, which is not a rail

The highest-return intervention on this axis is not a rail decision at all. Take a constructed group with six entities and gross intercompany flows of 100.0 units a quarter:

```
gross intercompany flows per quarter        100.0 units
bilateral netting                            62.0 units  (−38.0%)
multilateral netting to one settlement       23.0 units  (−77.0%)
```

Expected transit loss is proportional to value moved, so netting to a single settlement position reduces it in the same proportion. **Multilateral netting delivers a 77.0 per cent reduction in expected transit loss, requires no new rail, no new counterparty and no new jurisdiction, and its only cost is a treasury system and the discipline to settle on a cycle.** Section 6 shows that informal value-transfer systems are built on exactly this insight, centuries earlier and without the treasury system.

6. Informal value transfer as an object of measurement

The remaining material is analysed as a licensed, priced industry rather than as a curiosity, because its arithmetic supplies the sharpest available test of the concentration instrument.

6.1 Settlement is where the economics live

A bilateral position between two operators must eventually be cleared, by reverse flows in the same corridor, by trade, by a formal bank transfer of the net position, or by multilateral netting against other operators. **The third of these puts the system inside the correspondent perimeter, which is the opposite of what the popular account asserts.**

Netting efficiency is a property of the corridor's trade and migration pattern rather than of its operators. With gross flows G_{ab} and G_{ba},

$$\text{settlement requirement} = |G_{ab} - G_{ba}|, \qquad \text{netting efficiency} = 1 - \frac{|G_{ab} - G_{ba}|}{\max(G_{ab}, G_{ba})}.$$

```
CORRIDOR 1: labour-importing to labour-exporting
  G_ab = 100    G_ba =  15
  settlement requirement = 85     netting efficiency = 15.0%

CORRIDOR 2: two-way trading diaspora
  G_ab = 100    G_ba =  85
  settlement requirement = 15     netting efficiency = 85.0%
```

Corridor 1 must settle 85 units through the formal system for every 100 transferred and corridor 2 must settle 15, so corridor 1 costs roughly five and a half times as much to operate, is that much more exposed to correspondent friction, and is that much more visible to the formal system. **Corridor 1 is the shape of every major labour-migration remittance corridor.** Three consequences follow and none is obvious. Remittance corridors are the *least* netted and therefore the *most* dependent on formal settlement of any use case — the claim that they operate outside the banking system is close to exactly backwards. De-risking bites hardest precisely where netting is worst, which is to say in the poorest and most one-directional corridors. And a corridor's cost is predictable from its trade balance.

6.2 Obligation without an adjudicating forum

What stops an operator from simply not settling is the economics of reputation in a repeated game. The enforcement mechanism is exclusion from future business, and it works when the network is small enough that reputation travels, large enough that exclusion is costly, and stable enough that the future matters; kinship, common origin and long-term commercial relationships are the usual substrates because they satisfy all three cheaply. That cooperation is sustained in repeated interaction by the shadow of the future is Axelrod and Hamilton [1981] and Axelrod [1984]; that trust is a specific expectation about others' behaviour under sanctioning arrangements rather than a disposition is Gambetta [1988]; and that durable common-pool institutions run on graduated sanctions and monitoring by the participants themselves rather than on external enforcement is Ostrom [1990], whose design principles have been checked against a large case base [Cox et al., 2010].

The source manuscript names two further works in prose — an account of an eleventh-century traders' coalition and a model of the law merchant with a private judge — and neither is in this paper's bibliography. They are described and left uncited (Section 11 and Appendix B), because inventing a citation key is worse than declaring a gap.

Two consequences. The mechanism is genuinely effective within its domain, and an operator's capital is largely reputational. **And it is fragile at the boundary: it provides no protection to a party outside the network, offers no recourse to a customer, and produces no evidence a court can use.** A system whose enforcement is reputational offers its participants excellent protection and its customers none.

6.3 Where the margin comes from

Operator revenue has a commission component and a foreign-exchange spread component, and the second is usually larger and is frequently manufactured by policy rather than by the operator. Where a state maintains an overvalued official rate alongside capital controls, a parallel rate exists by construction and the gap is a subsidy paid by anyone forced to use the official rate. An operator settling at or near the parallel rate can offer a better rate than the formal channel and still earn a spread. **A state running a large parallel premium and prosecuting informal transfer is attempting to enforce against an arbitrage it is itself creating, which is the clearest instance in this series of a policy failure misdiagnosed as a compliance failure.** The episodes the source names, with their dates and premia, are in Appendix B.

7. Results

Every figure is an exact computation on constructed inputs. **No dataset was collected and no figure here is a measurement taken from a real corridor.**

7.1 R1 — the clearing-perimeter measurement

CONSTRUCTED CASE. Twelve months of cross-border settlement value for a group with a bridging apex, traced to intermediaries:

correspondent	share	perimeter of the bank
A, global systemically important bank	46.0%	WEST
B, global systemically important bank	24.0%	WEST
C, regional bank	14.0%	BRIDGE
D, regional bank	10.0%	BRIDGE
E, regional bank	6.0%	NONALIGNED

Five relationships, four jurisdictions, three perimeters, and a file describing it as a diversified correspondent network. Every word of that description is true.

```
naive correspondent HHI = 0.211600 + 0.057600 + 0.019600
                        + 0.010000 + 0.003600 = 0.302400
effective correspondents = 1 / 0.302400 = 3.306878

collapsed to the BANKS' own perimeters
  WEST 70.0%  BRIDGE 24.0%  NONALIGNED 6.0%
  bank-perimeter HHI = 0.551200
  effective bank perimeters = 1.814224
```

That is already a large fall, and it is still the wrong measurement, because it collapses banks by where the *bank* sits rather than by where the *currency* clears. Decomposing each correspondent's volume by currency of denomination gives USD 69.0, EUR 16.8, AED 5.6, SGD 3.0, CNH 2.0 and INR 3.6 per cent, so

```
currency HHI = 0.476100 + 0.028224 + 0.003136
             + 0.000900 + 0.000400 + 0.001296 = 0.510056
effective currencies = 1.960566
```

Six currencies, 1.96 effective. A five-currency revenue profile is real at the revenue line and has collapsed to essentially two by the time it reaches settlement, because the dollar is

the settlement currency of the trading business regardless of where the revenue was earned. Mapping each currency to the perimeter in which it achieves finality:

currency	share	perimeter
USD	69.0%	WEST
EUR	16.8%	WEST
AED	5.6%	BRIDGE
SGD	3.0%	BRIDGE
CNH	2.0%	EAST
INR	3.6%	NONALIGNED

```
WEST       85.8%   0.736164
BRIDGE      8.6%   0.007396
EAST        2.0%   0.000400
NONALIGNED  3.6%   0.001296
-----------------------------------------
clearing-perimeter HHI = 0.745256
N_rail = 1 / 0.745256 = 1.341819
```

measurement	HHI	effective count
naive correspondents	0.302400	3.306878
bank perimeters	0.551200	1.814224
clearing perimeters	**0.745256**	**1.341819**

Applying the surplus term at $\eta = 0.35$ and $\delta Q + C = 0.500000$:

axis	N	$\ln N$	$\eta \ln N$	h
jurisdiction / bloc	2.941176	1.078810	0.377583	-0.122417
correspondent, naive	3.306878	1.196005	0.418602	-0.081398
bank perimeter	1.814224	0.595658	0.208480	-0.291520
clearing perimeter	**1.341819**	**0.294026**	**0.102909**	$\mathbf{-0.397091}$

The structure is 3.24 times worse on h on the rail axis than on the jurisdictional axis, and the naive correspondent reading is not merely optimistic — it is better than the jurisdictional reading, which would lead a reader to conclude that the rail was the strong side of the structure. It is the weak side by a wide margin. H1 is supported.

7.2 R2 — the remedy that does not work, and the one that does

Add three banks in three further jurisdictions, each taking five points of volume drawn proportionately from the existing five, giving shares 39.1 / 20.4 / 11.9 / 8.5 / 5.1 / 5.0 / 5.0 / 5.0:

```
naive correspondent HHI = 0.225984
effective correspondents = 1 / 0.225984 = 4.425092   (+33.8%)
```

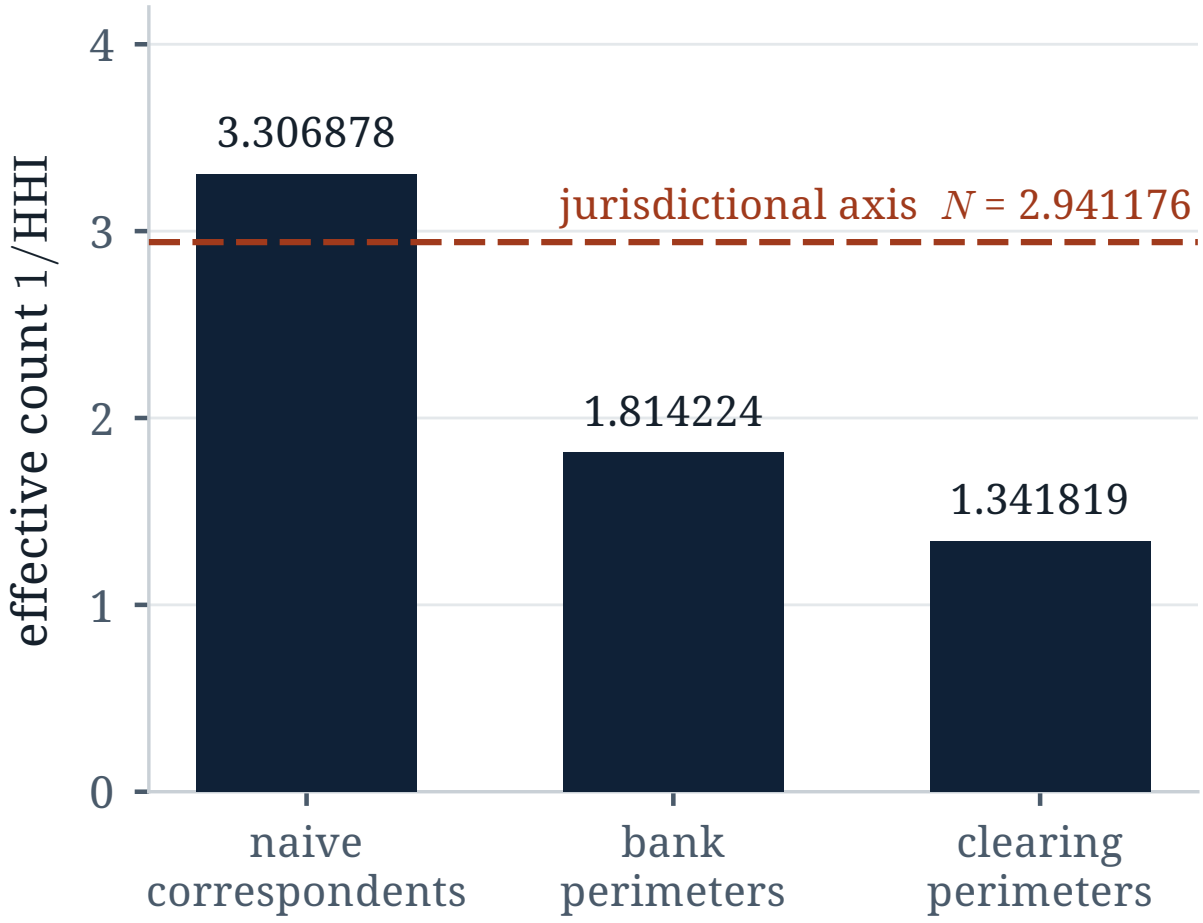


Figure 61.1: **Five correspondents in four jurisdictions across three perimeters read 3.306878 effective at the account layer and 1.341819 at the layer where finality occurs.** Effective counts are $1/\mathrm{HHI}$ on the three share bases of Section 3.1, computed from the constructed corridor set in Section 7.1; the dashed line is the same structure's effective bloc count on the jurisdictional axis. Script `figs/src/CCP-061_fig01.py`.

```
currency mix unchanged, because adding a bank does not change
what the invoices say

clearing-perimeter HHI = 0.745256   N_rail = 1.341819   UNCHANGED
```

Three new banking relationships, three new onboardings, three new sets of ongoing diligence, a 33.8 per cent improvement in the number nobody should be reporting, and exactly zero change in the number that matters. That is rearrangement inside a perimeter, reproduced on the rail axis; Curzi [2026] reaches the identical structural result for two free zones inside one supervisory perimeter, which are one line at 29.1 per cent of exposure and where no rearrangement between them changes N.

Now change the denomination of the underlying obligations over a three-year contract cycle — regional sales renegotiated into regional currencies at renewal, fee income denominated where it arises, treasury operating balances held in local currency rather than swept — giving USD 52.0, EUR 14.0, AED 12.0, CNH 9.0, INR 8.0 and SGD 5.0 per cent:

```
WEST 66.0%  BRIDGE 17.0%  EAST 9.0%  NONALIGNED 8.0%
clearing-perimeter HHI = 0.435600 + 0.028900 + 0.008100
                       + 0.006400 = 0.479000
N_rail = 2.087683   ln(N_rail) = 0.736055
eta ln(N_rail) = 0.257619   h_rail = -0.242381

before   N_rail 1.341819   h_rail -0.397091
after    N_rail 2.087683   h_rail -0.242381
gain                        +0.154710
```

A gain of 0.154710 in h, achieved with no new entity, no new jurisdiction and no new bank. For comparison, the corpus's first non-Western jurisdictional move — taking N from 1.00 to approximately 1.70 — is worth $0.35 \ln(1.70) = 0.185720$, and the bridging positions priced in Curzi [2026] and Curzi [2026] are worth an order of magnitude less. **The currency remediation is the second largest single-move gain computed anywhere in this series and the only large one requiring no structure at all.** H2 is supported.

The bill, stated with the gain. Redenomination leaves the group holding receivables in five currencies against a cost base in three, so hedging costs basis points and unhedged mismatch costs more; a customer asked to pay in its own currency will expect a price concession for taking foreign-exchange risk it previously did not bear, estimated at 30–80 basis points on affected revenue; five currency ledgers require five reconciliations and five sets of account limits; and currencies with thin offshore markets require larger operating balances, which are holdings and must be measured as such. Constructed net: an annual cost of roughly 0.045 units against a gain of 0.154710, favourable by a factor of about three. Curzi [2026] is the standing bound on how far this can be taken: at a 15 per cent annual hard-currency drag a local-currency allocation of 25.17 per cent of protected capital sets h to exactly -0.500000, cancelling the entire $\eta \ln N = 0.377583$ diversification gain. **Currency independence bought past that point costs more than the independence is worth.**

7.3 R3 — the transit model, worked

Calibrate at a per-hop freeze probability of 0.2000 per cent for a six-hour correspondent window, so $\lambda = -\ln(1 - 0.002)/6 = 3.3367 \times 10^{-4}$ per hour, with $V = 40{,}000{,}000$ and $r = 0.15$:

case	window	hops	q	$\mathbb{E}$[loss] (USD)	worst single loss	P(any loss)
1	6 h	1 of 40m	0.002000	68,000	34,000,000	0.2000%

case	window	hops	q	$\mathbb{E}$[loss] (USD)	worst single loss	P(any loss)
2	6 h each	8 of 5m	0.002000	68,000	4,250,000	1.5889%
3	20 min	1 of 40m	0.00011122	3,781	34,000,000	0.0111%
4	6 h, screened	1 of 40m	0.000400	13,600	34,000,000	0.0400%
5	20 min, screened	1 of 40m	0.00002224	756	34,000,000	0.0022%

Read cases 1 and 2 together: eight times more likely to suffer some loss, eight times smaller worst case, identical expected loss. Then read 3 and 4 against them: the window and the screening each move the expectation by an order of magnitude, and the hop size moves it by nothing. H3 is supported, and it is supported as an algebraic identity rather than as an empirical finding.

Splitting is not free. At a constructed per-hop cost of USD 2,500, one hop costs 2,500 against an expected loss of 68,000 for a total of 70,500, while eight hops cost 20,000 against the same expected loss for a total of 88,000. **Splitting the payment costs USD 17,500 and reduces expected loss by nothing. It buys a reduction in the worst case from USD 34.0m to USD 4.25m, and that is a purchase of insurance at a stated premium.** With a covenant tolerance of USD 6,000,000 the constraint gives $k \le 6{,}000{,}000/0.85 = 7{,}058{,}824$ and $n \ge 40{,}000{,}000/7{,}058{,}824 = 5.67$, so the answer is six hops of roughly 6.67m — **set by the constraint, not by an optimisation.**

Worked on a planned intra-group repatriation of USD 40m with the counterparty component of the hazard screened to 20 per cent of the base rate and a covenant tolerance of USD 12m:

option	window	hops	$\mathbb{E}$[loss]	explicit	landed	worst single	P(any)	verdict
1	48 h	2 of 20m	108,735	100,300	209,035	17,000,000	0.6386%	fails tolerance
2	48 h	4 of 10m	108,735	100,900	209,635	8,500,000	1.2731%	acceptable
3	20 h	2 of 20m	45,349	92,000	137,349	17,000,000	0.2666%	fails tolerance
3S	20 h	4 of 10m	45,349	92,600	137,949	8,500,000	0.5325%	**chosen**

The saving of USD 71,686 against option 2 comes entirely from removing a hop from the chain, and the split that satisfies the covenant costs USD 600. The window is worth a hundred times what the hop sizing is worth, and the hop sizing is a constraint to be satisfied rather than a lever to be pulled. Note also that option 3S has a *higher* probability of some loss than option 3 and a lower expected loss than either 1 or 2: **three metrics moving in three directions from one decision, so all three must be reported or the decision is being made on whichever one the presenter chose.**

7.4 R4 — netting

```
E[loss] gross          proportional to 100.0
E[loss] netted          proportional to  23.0
reduction                                77.0%
```

Multilateral netting to a single settlement position cuts the constructed group's expected transit loss by 77.0 per cent with no new rail and no new counterparty. It is the highest-return intervention in this paper and it is not a rail decision at all. It is the decision not to use one.

7.5 R5 — the indifference hazard

Landed cost on USD 40,000,000, on the constructed parameterisation:

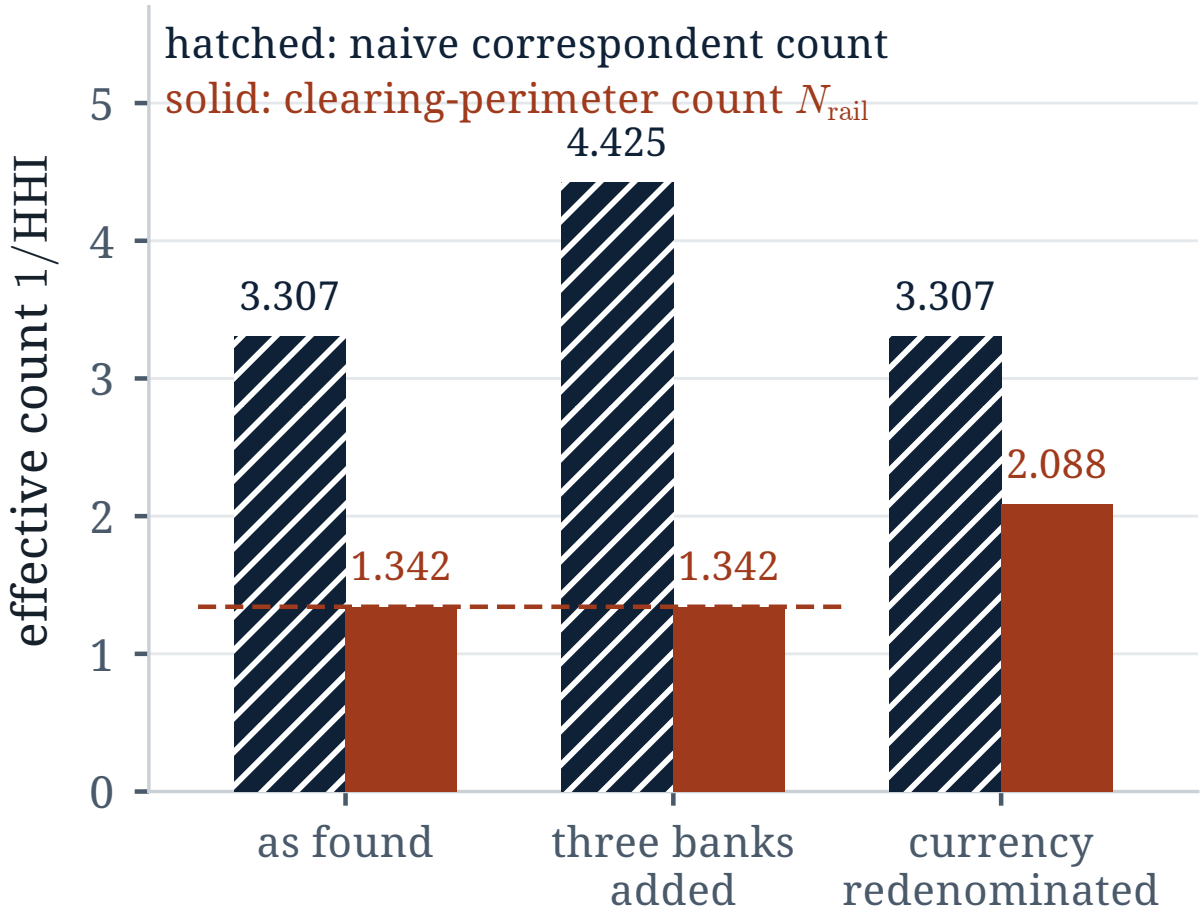


Figure 61.2: **Adding three correspondent banks improves the account-layer count by a third and leaves the clearing-perimeter count unchanged to six decimal places; redenominating the underlying obligations is what moves it.** Effective counts on the correspondent and clearing-perimeter bases from Sections 7.1 and 7.2; the dashed line marks $N_{\text{rail}} = 1.341819$. Script `figs/src/CCP-061_fig02.py`.

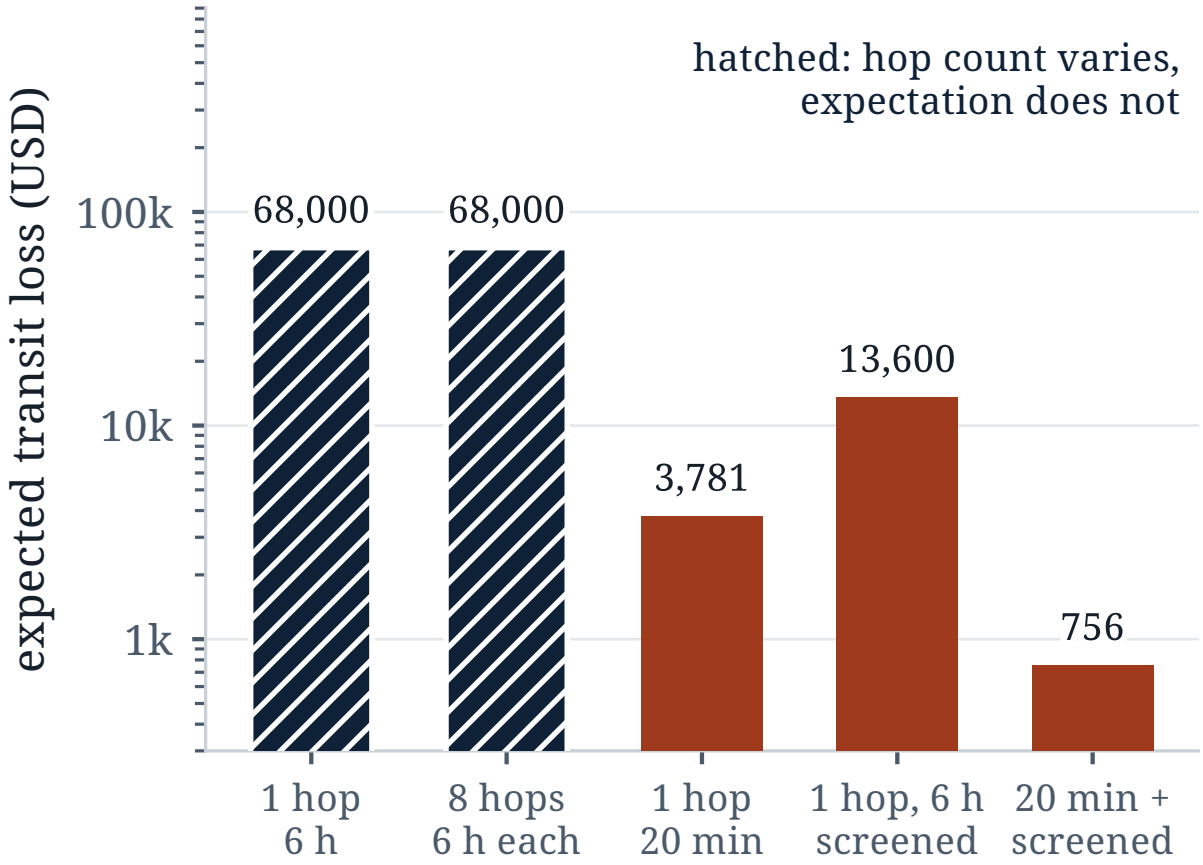


Figure 61.3: **The exposure window and counterparty screening each move expected transit loss by an order of magnitude; the number of hops moves it by nothing.** Expected loss $V(1-e^{-\lambda t})(1-r)$ at $V = 40{,}000{,}000$ and $r = 0.15$, for the five constructed cases of Section 7.3; hatched bars differ only in hop count. Script `figs/src/CCP-061_fig03.py`.

```
CORRESPONDENT WIRE, 48-hour window
  wire and lifting fees                          USD      300
  FX spread 15 bp on the converted leg           USD   60,000
  q = 1 - exp(-0.00033367 x 48) = 0.0158884
  E[freeze loss] = 40m x 0.0158884 x 0.85        USD  540,207
  ------------------------------------------------------------
  landed cost                                    USD  600,507

FAST RAIL, 0.5-hour window
  on-ramp spread 10 bp                           USD   40,000
  network fees                                   USD      100
  off-ramp spread 20 bp                          USD   80,000
  expected price variance over the window        USD   20,000
  q = 1 - exp(-0.00033367 x 0.5) = 0.00016682
  E[freeze loss] = 40m x 0.00016682 x 0.85       USD    5,672
  ------------------------------------------------------------
  landed cost                                    USD  145,772
```

The fast rail lands at under a quarter of the correspondent cost on this parameterisation, and every bit of the advantage comes from the window rather than from the fees — indeed its explicit costs are more than twice as high. Setting the two landed costs equal and approximating $q \approx \lambda t$:

```
60,300 + 40m x 0.85 x lambda x 48 = 140,100 + 40m x 0.85 x lambda x 0.5
34,000,000 x lambda x 47.5 = 79,800
lambda = 4.9412e-05 per hour
q48 = 1 - exp(-4.9412e-05 x 48) = 0.0023690   ->   0.2369%
```

Above a 48-hour freeze hazard of about 0.24 per cent, take the fast rail; below it, take the slow one. For a legitimate group moving money between its own subsidiaries in ordinary corridors, 0.24 per cent per transaction is a very high hazard, and the honest conclusion is that **most readers should use the correspondent rail most of the time.** The fast rail earns its cost in corridors that are genuinely stressed — which is to say exactly where the screening burden is highest and the regulatory perimeter least settled. **The rail is cheapest to use where it is hardest to use correctly, which is not a coincidence and is not a comfortable finding.**

Screening, sized rather than assumed, on a constructed twelve-payment supplier programme of CNY 5,000,000 a month with an eighteen-hour window: unscreened $q = 0.0059881$ gives CNY 25,449 per payment and CNY 305,388 a year; with the counterparty hazard cut by 80 per cent, $q = 0.0012005$ gives CNY 5,102 and CNY 61,225. **The annual benefit is CNY 244,163 against a constructed programme cost of CNY 95,000 — a ratio of 2.57 to one, with the programme paying for itself if it removes about 31.1 per cent of the counterparty hazard, which is a modest claim.**

Run on concurrent transit exposure across the group after redenomination — USD 52.0, EUR 14.0, AED 12.0, CNH 9.0, INR 8.0, SGD 5.0 per cent — the rail HHI is 0.321400 and the effective rail count is 3.111387. **This is a *transit* measurement and is deliberately not the clearing-perimeter measurement of** 2.087683: transit value is not proportional to settlement value, because some corridors move a large annual value in a few large hops and others move a small value continuously. Reporting one as the other is the axis-confusion failure arriving inside a single analysis. Adding a fast-rail leg at 10 per cent of transit value raises the effective rail count from 3.11 to 3.70 while introducing three sub-axes with effective counts of 1.00 for the issuer, 1.00 for the network and 1.83 for the venues at the ramps. **A rail that improves a headline count while introducing sub-axes at 1.00 has not obviously improved anything, and the measurement must be run at the sub-axis level or it flatters the decision.**

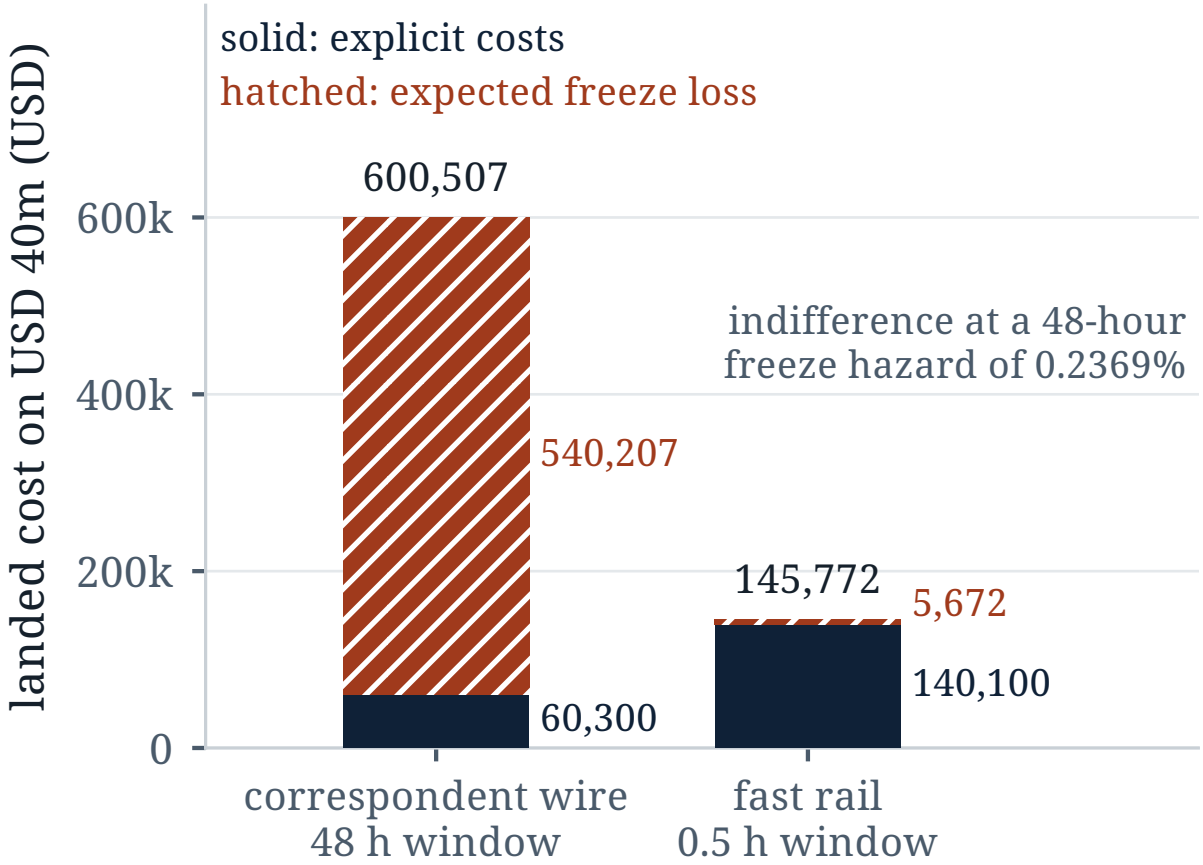


Figure 61.4: **The fast rail's advantage is entirely the exposure window: its explicit costs are more than twice the correspondent rail's and its landed cost is under a quarter.** Landed cost on USD 40,000,000 decomposed into explicit costs and expected freeze loss at the constructed hazard of Section 7.5. Script `figs/src/CCP-061_fig04.py`.

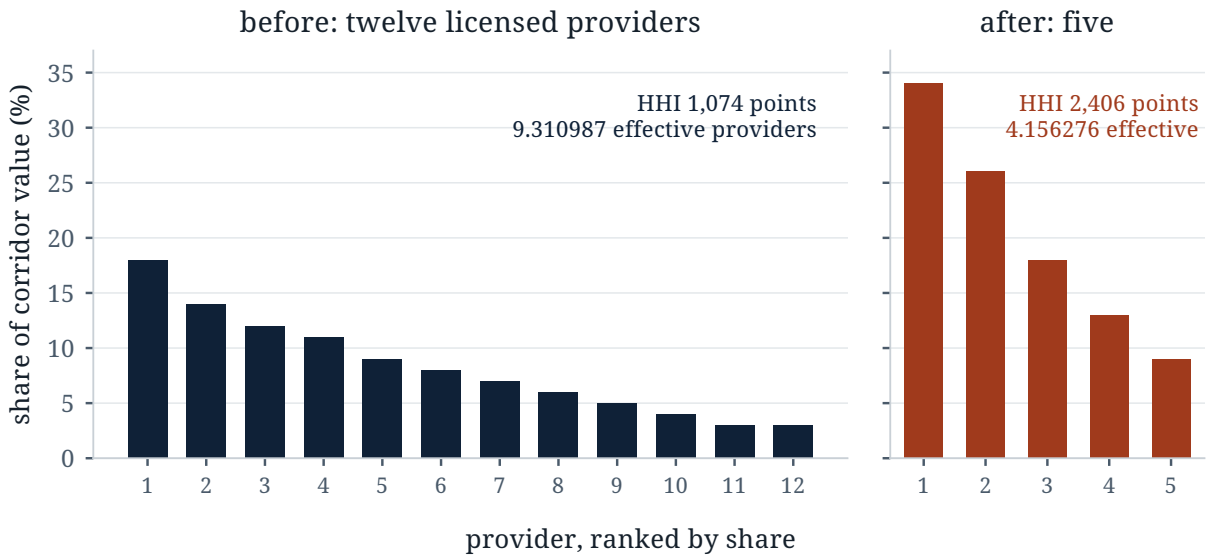


Figure 61.5: **Enforcement-driven exit produced, without any transaction, a market that clears the published structural-presumption screen by 606 points.** Constructed corridor shares before and after a de-risking cycle, with the index in points and the effective provider count $1/\mathrm{HHI}$; from Section 7.6. Script `figs/src/CCP-061_fig05.py`.

7.6 R6 — the concentration effect of enforcement

CONSTRUCTED CASE. A remittance corridor with twelve licensed providers before a de-risking cycle and five after, shares of corridor value:

```
BEFORE  18 / 14 / 12 / 11 / 9 / 8 / 7 / 6 / 5 / 4 / 3 / 3
        HHI = 1,074 points = 0.107400
        effective providers = 9.310987

AFTER   34 / 26 / 18 / 13 / 9
        HHI = 2,406 points = 0.240600
        effective providers = 4.156276
```

metric	before	after	change
licensed providers	12	5	−7
HHI (points)	1,074	2,406	+1,332
HHI (fraction)	0.107400	0.240600	+124.0%
effective providers	9.310987	4.156276	−55.4%
largest single share	18.0%	34.0%	+16.0 points

Applying the published merger screen — a structural presumption above 1,800 points with a change above 100 [U.S. Department of Justice and Federal Trade Commission, 2023] — the corridor after exit reads 2,406 points with a change of 1,332 and a largest share of 34.0 per cent. **A regulatory intervention undertaken to reduce financial-crime risk produced, on the competition agencies' own screen, a market clearing the concentration threshold by 606 points and the change threshold by a factor of thirteen. No merger produced it, no filing was required because no transaction occurred, and there is no agency whose remit covers the outcome.** H4 is supported.

Modelled conservatively at a constructed price elasticity of 0.25 with respect to the index, a 124.0 per cent rise in concentration produces a 31.0 per cent rise in price, taking corridor cost

from 6.0 to 7.86 per cent of the amount sent — on a constructed corridor value of USD 4 billion, **USD 74,400,000 a year of additional cost borne by senders.** The elasticity is constructed and should be replaced with a measured one; the sign is not constructed. The risk consequence is the point of the exercise:

```
volume displaced from licensed to unlicensed channels
  -- NOT MEASURABLE. No figure is asserted.
supervisory relationships destroyed              7
filing entities removed from the perimeter       7
examinable entities removed from the perimeter   7
```

The last three lines are exact and the first is unknowable, which is itself the finding: the cost of the intervention is precisely measurable and its benefit is not.

One layer down, the five survivors' own banking:

```
Bank Alpha = 34.0 + 26.0 + 13.0 = 73.0%   0.532900
Bank Beta  = 18.0 +  9.0        = 27.0%   0.072900
------------------------------------------------------
correspondent HHI = 0.605800
effective correspondents = 1.650710
```

Five providers, 4.156276 effective; two correspondents, 1.650710 effective. One bank's exit removes 73.0 per cent of the corridor's licensed capacity, and that decision requires no consultation, no notice period beyond the contractual one, and no reasons.

7.7 Counter-result

After the largest available intervention the rail remains the structure's weakest axis. $N_{\text{rail}} = 2.087683$ against 2.941176 on the jurisdictional axis and against the break-even count of 4.172734 that Curzi [2026] fixes for the storage law at these parameters. The dollar share stops at 52.0 per cent after a three-year redenomination programme and will not go much lower, because the commodity and trade-finance legs of the business are priced in that currency globally and the group cannot unilaterally change it. **Some concentration is a property of the world rather than of the design, and the correct response is to record it rather than to model it away.** The depository layer is not remediable by a private group at all and is sized and monitored rather than fixed. **Correspondent exit risk is not hedged and cannot be**; it is mitigated by explicability, tested contingency relationships and a twelve-month lead time on every replacement.

8. Discussion and failure catalogue

8.1 What the measurement licenses

The instrument measures **independence, not capacity**. $N_{\text{rail}} = 2.087683$ says nothing about whether any perimeter can carry the group's volume, nothing about the price of doing so, and nothing about whether a counterparty will accept the instruction. Curzi [2026] makes the same point on a different share basis and reaches the same shape: entity-based exposure of 27.3 per cent against a currency share of 61.0 per cent and settlement volume of approximately 58 per cent, so that **the entity map understates dependency by more than a factor of two.** The rail measurement is that mismatch taken one layer further down.

The corpus's own five-parameter survey is the closest thing to a validation this instrument has, and it is not on this axis. Curzi [2026k] applies the same storage-law parameterisation from a single desk to a twenty-two-household block and reproduces each system's observed decline — -0.250474 predicted against -0.250 observed at the desk, -0.996371 against -1.000

at the library, with residuals under 1.4 per cent of the ceiling term. **That is scale-freedom demonstrated at two scales in a domain with observable outcomes. No equivalent demonstration exists for the rail axis, and Section 11 says so.**

The framework's channel-gain form bounds what any of this can promise. Curzi [2026ab] states the magnitude form $E = G \cdot B + \varepsilon$ with the forbidding clause that $G = 0$ implies $E = 0$ for all B, badged as a postulate and paired with the storage law's $\eta = 0$ by explicit correspondence. **On the rail axis the forbidding clause is literal: where no clearing channel exists into a perimeter, no amount of structuring, documentation or belief produces settlement into it.** A rail programme is a claim about G and it must be tested by moving money, not by drafting a policy.

8.2 Thirteen failure modes

1. **Treating a clearing system as a substitute for a messaging network.** The category error of Section 1.1; a policy recording two rails on that basis has recorded one dependency twice.
2. **Counting correspondents instead of measuring perimeters.** The gap between 3.306878 and 1.341819 is the whole paper.
3. **Ignoring nesting.** A respondent's own respondents are exposure the count does not show [Douceur, 2002].
4. **Diversifying accounts while concentrating currency.** Section 7.2, measured.
5. **Direct versus indirect participation, unverified.** Whether a given bank reaches a clearing system directly is a bank-by-bank question and cannot be answered jurisdiction by jurisdiction.
6. **The custody chain unmapped.** The chain must be mapped to the depository, not to the custodian.
7. **Treating de-risking as a legal risk.** It is a commercial decision with no appeal, and it propagates as a cascade [Bikhchandani et al., 1992].
8. **Building the structure and then seeking the banking.** The onboarding constraint is binding and is discovered last.
9. **Treating a settlement balance as a rail.** It is a holding, and it enters the jurisdictional measurement.
10. **A currency rail with no receivables in that currency.** The rail follows the invoice, so a rail without an invoice is a facility, not an exposure change.
11. **Assuming immobilisation is a sovereign-only risk.** The mechanism is administrative and reaches any holder through an intermediary.
12. **Messaging-migration data loss.** A standards migration truncates fields [International Organization for Standardization, 2013], and truncated originator data is itself a screening hazard.
13. **Assuming a rail can be changed quickly.** Onboarding runs to quarters and redenomination to a contract cycle; **the only window common to every rail failure is *before*.**

The corpus's amplitude weighting is the general form of that last point. Curzi [2026] scores each step of a path 0–3 for steepness, sums to a resistance R and weights the branch as $1/2^R$, so a three-point difference in steepness is a factor of eight in expectation. **Applied here, every branch reachable only after the signal carries a resistance the model prices at near zero weight, which is the formal statement of why rails are built in peacetime or not at all.**

8.3 Where the analysis is exposed

The instrument is exposed to being adopted as a target. An effective count optimised against ceases to measure what it measured [Goodhart, 1975, Campbell, 1979]; the only defence offered is the reporting standard in Appendix A, which requires the inputs to be published beside the output.

It is also performative, and its own parameters are not exempt. MacKenzie and Millo [2003] documents a model that made its assumptions true after adoption, and Lucas [1976] states the consequence for parameters estimated before publication. **The framework asserts that systems change their behaviour when modelled and then exempts its own η, λ and perimeter assumptions from that reflexivity. That is an inconsistency and it is not resolved here.**

Ensemble and time averages are different objects and every expected-loss figure here is the first. Peters [2019] establishes that the time average of a multiplicative process differs from its ensemble average, so an expectation computed across many payers does not describe what happens to one payer over time — and a freeze that removes a settlement is not a deduction from a stream but an event that rescales what follows. Peters and Gell-Mann [2016] supplies the associated strengthening for the framework's logarithm: a logarithmic form follows from multiplicative dynamics with no psychological assumption at all, which is a better warrant than the perceptual account the corpus more often gives.

One boundary the framework refuses to cross belongs here rather than in a footnote. Curzi [2026] states the doctrine of category violation, which refuses category innocence as a defence of acquisition and forbids moving a harm through a discovered gap in a scheme, and concedes as an unclosed seam the case where the feedback loop runs through another person's will. **The rail analysis is measured against arbitrary action and is not licensed as a technique for evading a process that has run**; the de-risking result in particular is an argument about a policy's second-order costs and not an argument for the unlicensed channel it displaces volume into.

9. Verification and rigour audit

9.1 Dimensional audit

symbol	quantity	dimension	unit	status
s_i	share on a stated basis	$[1]$	dimensionless fraction	defined; sums to 1 by P1
HHI	concentration index	$[1]$	dimensionless, or points $= 10^4\times$	derived from shares
N, N_{rail}	effective count	$[\#]$, a count	perimeters, banks, providers	derived as $1/\text{HHI}$
$\ln N$	surplus-term argument	$[1]$	nats	derived; logarithm of a count is dimensionless
λ	freeze hazard rate	$[T^{-1}]$	per hour	assumed (constructed)
t	exposure window	$[T]$	hours	assumed (constructed)
q	per-hop freeze probability	$[1]$	dimensionless	derived, $1-e^{-\lambda t}$
V, k	total value, hop size	$[M]$	currency	assumed (constructed)
n	hop count	$[\#]$	a count	chosen
r	recovery fraction	$[1]$	dimensionless	assumed (constructed)
X	absorbable single loss	$[M]$	currency	assumed (constructed)

symbol	quantity	dimension	unit	status
Q	protected capital	$[C]$	index units	defined; **no operational unit is supplied by the source**
δ	erosion rate	$[T^{-1}]$	per period	assumed
C	maintenance cost flow	$[C\,T^{-1}]$	index units per period	assumed
η	conversion efficiency	$[C\,T^{-1}]$ per nat	index units per period per nat	assumed
h	surplus balance	$[C\,T^{-1}]$	index units per period	derived

Do the equations balance? λt is dimensionless, so $q = 1 - e^{-\lambda t}$ is well formed and the exponent takes a pure number, which is the check a physicist runs first. $\mathbb{E}[\text{loss}] = Vq(1-r)$ carries $[M]$, correctly. Landed cost adds $[M]$ to $[M]$. The surplus law $h = \eta \ln N - \delta Q - C$ balances **only if η carries capital per period per nat, and the source is silent on this**: it treats conversion efficiency as a bare number. **The audit records the silence rather than filling it in**; the reading in the table is this paper's.

Where no operational unit exists the audit says so. Q, C and h are index units fixed by the stated parameters, with no conversion to currency supplied anywhere in the corpus. **The effective counts, the indifference hazard and every currency figure survive this; the statement that h_{rail} is -0.397091 means nothing until an index unit is defined.**

A units mismatch inside the paper is recorded rather than smoothed. The transit model is denominated in currency and the surplus law in index units, and **the two are never added.** No exchange rate between them exists and none is asserted.

9.2 Significant figures against epistemic status

- 3.306878, 1.814224, 1.341819, 2.087683, 4.425092, 9.310987 and 4.156276 are **exact reciprocals of exactly stated share vectors.** They are derived-from-defined and may carry as many digits as are useful. **Their inputs are constructed and carry no empirical authority whatever.**
- Every index in this paper is computed from the printed percentage shares rather than from any exact underlying fraction, **which is the convention adopted by Curzi [2026] and recorded there as a fifth-decimal discrepancy against the exact-fraction convention of Curzi [2026].** The convention is stated so that a reproduction agrees to the last digit.
- $h_{\text{rail}} = -0.397091$ and the gain of $+0.154710$ are exact given the parameters, and **the parameters are assumed, not estimated.** They carry no uncertainty because there is no estimate in them, and no authority beyond the assumption.
- USD 68,000, USD 540,207 and CNY 305,388 are **illustrative model output from constructed parameters** and are rounded to the nearest unit for printing only.
- $q_{48} = 0.2369\%$ is exact arithmetic on stipulated costs.
- **No quantity in this paper is a fitted or measured estimate, so no uncertainty interval is reported anywhere. That is not precision; it is the absence of empirical content.**

9.3 Limits and boundary behaviour

limit	behaviour	reading
one perimeter, $s_1 = 1$	HHI $= 1$, $N = 1$ exactly	a single-perimeter rail counts as one
n equal perimeters	$N = n$ exactly	returns the count
n unequal	$N < n$	never overstates the count
$t \to 0$	$q \to 0$, $\mathbb{E}[\text{loss}] \to 0$	instantaneous finality removes the transit hazard entirely
$t \to \infty$	$q \to 1$, $\mathbb{E}[\text{loss}] \to V(1-r)$	value permanently in flight is value at risk in full
$\lambda \to 0$	$\mathbb{E}[\text{loss}] \to 0$	the fast rail's whole advantage vanishes and the indifference hazard is never reached
$r \to 1$	$\mathbb{E}[\text{loss}] \to 0$	full recovery makes a freeze a delay rather than a loss
$n \to \infty$ at fixed V	$\mathbb{E}[\text{loss}]$ unchanged; $P(\text{any}) \to 1$; worst case $\to 0$	the invariance, in the limit
$\eta \to 0$	$h \to -(\delta Q + C)$	no conversion, no benefit from any independence
$\delta \to 0$	stock never erodes	pathology, not an ideal
$N \to \infty$	$h \to \infty$ logarithmically	**a defect of the model, not a result**

Range: $1 \leq 1/\text{HHI} \leq n$, with both bounds attained. A quantity equal to n for n equal units, equal to 1 for a single unit, and strictly between otherwise, **is a count**.

Two limits are constraints on the model's domain rather than results. The unbounded-N behaviour of the surplus law contradicts every organisational observation, so **the model's domain is small** N and every prescription here concerns moves between $N \approx 1$ and $N \approx 3$. And the small-λt approximation used in the indifference calculation overstates q for large λt; at the values used the error is in the fourth significant figure, and the exact form is used everywhere else.

9.4 Invariance

Every effective count here is invariant under relabelling and under any rescaling of the underlying values that preserves the share vector. Doubling every corridor's volume leaves each s_i unchanged and therefore leaves HHI, N and $\ln N$ unchanged. **It is not invariant under regrouping**, which is the entire content of the perimeter collapse.

The transit model carries a second and different invariance: expected loss is invariant under the partition of V into hops. That invariance is exact, algebraic, and holds for every n; it is the content of H3 and it is what makes hop-splitting an insurance purchase rather than a risk reduction. **It breaks under exactly one condition — correlated freezes — and that condition is named in Section 3.4 rather than discovered by a reader.**

What would break the share invariance: any basis that is not normalised, such as reporting corridor volumes in currency without dividing by the total. **Tested at how many scales?** The invariance is algebraic rather than empirical and was checked at the share vectors reported here: two, four, five, six, eight and twelve entries. No claim of empirical scale-freedom is made on this axis, and Curzi [2026k]'s two-scale demonstration is in a different domain.

What is conserved, and in what units. Mirowski [1989] shows that the characteristic failure of this genre is borrowing the notation of a conserved quantity without the conservation law that gave it content. The conserved quantity here is the share vector, $\sum_i s_i = 1$ exactly and

dimensionlessly, and that constraint alone gives the index its bounds. **Remove the normalisation and the effective count ceases to be a count.** Q is a stock with no stated conservation law and no operational unit, and none is claimed. Cartwright [1999] supplies the discipline that makes a partial answer respectable: specify the arrangement within which the law is claimed to hold. **The arrangement claimed here is a normalised share vector on one stated basis, and nothing outside it.** Kirman [1992] is the objection this paper cannot answer: one Q, one δ and one C are written for a heterogeneous group.

9.5 Analogy discipline

borrowing	source	category	justification
effective count as $1/\mathrm{HHI}$	Hill [1973], Jost [2006]	**1 — derivation**	HHI^{-1} *is* the Hill number of order 2; an identity
numbers-equivalent reading	Adelman [1969]	**1 — derivation**	elementary, checkable in two lines
portfolio variance $w'Rw$	Markowitz [1952]	**1 — derivation**	the exposure vector genuinely is a weighted sum of shocks with a correlation matrix
similarity-sensitive diversity	Leinster and Cobbold [2012]	**1 — derivation**, stated domain	exact where R has unit diagonal and non-negative entries
$\ln N$ as entropy	Shannon [1948]	**2 — shared form**	Shannon entropy in nats, dimensionless; no thermodynamic content imported
exponential hazard	standard survival form	**2 — shared form**	a constant-hazard model is assumed, not derived from any observed failure process
flow $\times$ delay as exposure	Little [1961]	**2 — shared form**	the stationarity premises are not verified for these corridors
impossibility of asynchronous consensus	Fischer et al. [1985], Gilbert and Lynch [2002]	**2 — shared form**	a payment system is not a formal asynchronous protocol; the trade-off's shape is what is used
fault-tolerance bound	Lamport et al. [1982]	**2 — shared form**	used for the discipline that redundancy requires independence, not as a protocol result

borrowing	source	category	justification
identity inflation	Douceur [2002]	**2 — shared form**	the impossibility result concerns network identities; the analogy to nesting is stated as an analogy
coupled-network fragility	Buldyrev et al. [2010], Albert et al. [2000]	**2 — shared form**	percolation premises are not shown to hold for correspondent networks
informational cascade	Bikhchandani et al. [1992], Banerjee [1992]	**2 — shared form**	the sequential-decision premises are plausible and untested here
tipping from a small correlated shift	Granovetter [1978], Schelling [1971]	**3 — framing**	prior statement of the aggregate phenomenon, cited as ancestry
normal accidents	Perrow [1984]	**3 — framing**	qualitative organisational theory, never a formal result

The failure this table exists to prevent is letting a category-2 borrowing carry a category-1 result's authority. The concentration arithmetic and the hop-invariance identity are category 1 and need no help; everything from the distributed-systems and network literatures is category 2 and is used for shape.

9.6 Measurement-theoretic status

Shares of settlement value, transit value and corridor value are **ratios of currency amounts** and are ratio-scale by construction, so the concentration instrument is untouched by the scale-type objection. **The hazard rate is a different matter and the objection is conceded.** λ is treated throughout as a ratio-scale quantity supporting multiplication, division and exponentiation, and it is a subjective degree of belief elicited from nobody. Michell [1997] is the strongest available objection to that move — that quantitative structure is assumed rather than tested — and the formal remedy is conjoint measurement [Luce and Tukey, 1964], whose conditions establish whether an attribute has additive quantitative structure at all, within the representational framework of Krantz et al. [1971]. **The conjoint conditions were not tested for λ, for r, or for the constructed price elasticity, and every ratio computed from one is therefore provisional.** What survives without the test is the *ordering* — a shorter window is better than a longer one, and an indifference hazard still separates settings — and the algebraic invariance of Section 3.4, which does not depend on the scale type at all.

The corpus's habit of collapsing a structured situation into one score applies to λ and is inherited rather than restated; the framework names it once as a framework-level limitation and the response here is to publish the decomposition of λ into four components in Section 3.4 rather than to defend the scalar.

9.7 Postulates

#	statement	status
P1	shares on a stated basis are non-negative and sum to 1	DEFINITION
P2	HHI $= \sum_i s_i^2$ on that basis	DEFINITION
P3	for n equal shares, HHI $= 1/n$	ARITHMETIC
P4	a currency achieves finality inside exactly one clearing perimeter	DEFINITION (modelling choice; the euro-area assignment is contested, Appendix B)
P5	jurisdictions inside one coordination perimeter act on one decision	POSTULATED (stipulated reference group, Section 3.6)
P6	freeze events are a constant-hazard process, $q = 1 - e^{-\lambda t}$	POSTULATED
P7	freezes are independent across hops	POSTULATED (**false for systemic events, Section 3.4**)
P8	$\lambda = 3.3367 \times 10^{-4}$/h, $r = 0.15$, per-hop cost USD 2,500	POSTULATED (constructed)
P9	the surplus law's input term is $\eta \ln N$	DEFINITION (modelling choice, CCP-A02 via CCP-052)
P10	$\eta = 0.35$, $\delta Q + C = 0.500000$	POSTULATED (fixed in advance)

9.8 The compounding chain

step	consumes	statement	status
L1	P2, P3	$1/$HHI returns n exactly for n equal shares	ARITHMETIC
L2	L1	it is a numbers-equivalent — Adelman [1969]	CITED
L3	L2	it is the Hill number of order 2 — Hill [1973], Jost [2006]	CITED
L4	L3, P4	on the clearing-perimeter basis it counts effective perimeters	DEFINITION
R1a	L4	correspondent shares 46/24/14/10/6 give HHI $= 0.302400$, count 3.306878	ARITHMETIC
R1b	L4, P4, P5	currency shares 69/16.8/5.6/3/2/3.6 collapse to perimeter shares 85.8/8.6/2/3.6, HHI $= 0.745256$, $N_{\text{rail}} = 1.341819$	ARITHMETIC
R1c	R1b, P9, P10	$h_{\text{rail}} = 0.35 \times 0.294026 - 0.5 = -0.397091$	ARITHMETIC

step	consumes	statement	status
R2a	L4	eight correspondent shares give HHI $=0.225984$, count 4.425092, $+33.8\%$	ARITHMETIC
R2b	R2a, P4	the currency vector is unchanged, so N_{rail} is unchanged	ARITHMETIC
R2c	L4, P9, P10	redenominated perimeter shares 66/17/9/8 give HHI $=0.479000$, $N_{\text{rail}}=2.087683$, $h=-0.242381$, gain $+0.154710$	ARITHMETIC
R3a	P6, P7	$\mathbb{E}[\text{loss}]=nkq(1-r)=Vq(1-r)$, so k cancels	ARITHMETIC
R3b	R3a, P8	at $V=40$m, $t=6$h, $q=0.002$: USD 68,000 for $n=1$ and for $n=8$	ARITHMETIC
R3c	R3b	$P(\text{any})=1-(1-q)^n$: 0.2000% at $n=1$, 1.5889% at $n=8$	ARITHMETIC
R4	R3a	expected loss is proportional to value moved, so a 77.0% cut in gross flows is a 77.0% cut in expected loss	ARITHMETIC
R5	R3a, P8	equating landed costs gives $\lambda=4.9412\times10^{-5}$/h, $q_{48}=0.2369\%$	ARITHMETIC
R6	L1	provider shares before and after give 1,074 and 2,406 points, counts 9.310987 and 4.156276	ARITHMETIC

Every ARITHMETIC row can be redone by hand or in three lines of any language; every CITED row names one source and no more. **Four postulates are POSTULATED, one of them — P7 — is stated to be false in the case that produces the tail, and the chain therefore establishes arithmetic consequences of stipulated inputs rather than facts about any payment system.**

9.9 Cross-checks performed

- Every index was recomputed from its printed share vector; all reported vectors sum to 100.0 per cent exactly.
- $s_{\max}^2 \le \text{HHI} \le s_{\max}$ holds throughout: $0.211600 \le 0.302400 \le 0.460$ on the correspondent vector and $0.736164 \le 0.745256 \le 0.858$ on the perimeter vector.
- The effective-count ceiling $s_{\max}^{-2}$ is respected: $3.306878 \le 4.7259$ and $1.341819 \le 1.3587$.
- The currency decomposition reproduces the correspondent totals: 46.0 + 7.2 + 8.4 + 5.0
  - 2.4 = 69.0 for the dollar leg, and the six currency shares sum to 100.0.

- **Three source discrepancies were found by this recomputation and are printed rather than corrected in silence.** (i) The source prints **4.424905** effective correspondents for the eight-bank case; $1/0.225984 = 4.425092$, so the source figure is wrong in the fourth decimal place and **this paper prints 4.425092**. The percentage improvement is 33.8 per cent on either figure. (ii) The source prints an expected freeze loss of USD **540,212** in the landed-cost table; the exact value at the stated hazard is USD **540,207**, the difference arising from a rounded intermediate value of q. (iii) The source prints an annual unscreened expected loss of CNY **305,393** where the exact value is CNY **305,388**, from the same cause; the derived break-even hazard reduction of 31.11 per cent is unaffected to four significant figures.

10. Limitations

Every case is constructed. The corridor set, the correspondent shares, the currency decomposition, the perimeter assignments, the hazard rate, the recovery fraction, the per-hop cost, the covenant tolerances, the corridor value, the price elasticity and the provider shares are invented and are labelled as such at each point of use. **No external dataset was collected and none is reported.**

No paper in this corpus reports a first-party measurement, and this one does not either. The corpus's own instrument audit finds that of 27 parameter-by-scope cells, 19 have a defined estimator, 1 has a validated one, and **0 carry a first-party measurement.** That is the central limitation of the whole corpus and it is stated here before a reviewer states it.

λ **is the most load-bearing and least defensible input.** It carries the entire transit analysis, the entire landed-cost comparison and the indifference hazard, and it is a constructed estimate with no interval and no reference class. **The single most useful thing a group could do is replace it with its own realised experience — every screening hold, every delayed payment, every returned instruction, recorded with elapsed times, over three years. Nobody keeps that record and everybody could.**

The independence assumption across hops is false in the scenario that matters. P7 holds for idiosyncratic freezes and fails for systemic ones, so the binomial distribution in Section 7.3 overstates the diversification benefit of splitting exactly when the tail arrives. **The paper states this rather than reporting the split as protection.**

The perimeter assignments are stipulated, not identified [Manski, 1993], and the remedy adopted is stipulation with a defence [Brock and Durlauf, 2001]. No randomisation, instrument or partial-population variation was available. **One assignment in particular is contestable**: the euro area is treated as inside the Western clearing perimeter, which affects every measurement in Section 7.1, and the alternative is not computed here. It is item 1 of Appendix B.

Jurisdiction-specific legal and regulatory claims are not cited to primary sources. The descriptions of settlement-finality legislation, the messaging-service prohibitions, the reserve-immobilisation instruments, the correspondent-banking recommendations, the participant counts and volumes of named clearing systems, and the enforcement resolutions are taken from the source manuscript. **The bibliography contains general comparative, treaty and standards material only; it contains no statutes, regulations, official statistics, registry documents or case reports for any individual jurisdiction or operator.** Those claims are described rather than asserted and are itemised in Appendix B. **This is a real gap and it is not closed here.**

Two works are named in the source's prose with no citation key and are not in this bibliography: an account of an eleventh-century traders' coalition and a model of the law merchant with a private judge, both supporting Section 6.2's claim that reputation in a repeated game sustains obligation without an adjudicating forum. The claim is instead supported here by Axelrod [1984], Gambetta [1988] and Ostrom [1990], which are in the bibliography and do support it, and the gap is declared rather than papered over.

Published measurements are described and not quoted. The decline in active correspondent relationships since 2011, the regional cost of remittances, the daily values cleared by named systems, and the amounts of two enforcement resolutions are all in the source manuscript as figures. **None is used as evidence in this paper and none appears in any result**, because none could be verified against a primary source within the bibliography's coverage.

The instrument measures independence and not capacity. N_{rail} says nothing about whether a perimeter can carry the volume, at what price, or whether an institution will accept the instruction. A group at $N_{\text{rail}} = 2.09$ whose entire bridging capacity sits with one institution has moved risk rather than reduced it.

Small-sample bias is not addressed [Chao et al., 2014], and the instrument is least informative in the dominated-unit regime [Golosov, 2010] — which is where the post-exit corridor sits, at 34.0 per cent in one provider.

Model comparison was not performed. No alternative index was fitted, no out-of-sample test was run, no information criterion was computed and no matched null was constructed, because nothing here is fitted. The instrument is asserted on the strength of the identity chain in Section 9.8 and not on comparative fit.

The functional form of the surplus term is underdetermined by the evidence. The corpus proposes a logarithm; a sub-linear power law and an $n \log n$ form fit the same qualitative facts. Nothing here discriminates between them and the paper concedes that nothing yet does; Peters and Gell-Mann [2016] is the one genuine strengthening available.

No differential prediction is offered. Durlauf [2005] asks for the prediction only this model makes, and the instrument makes none: any correlation-aware concentration measure would produce the same collapse, and any expected-value calculation would produce the same invariance. **The contribution is the share basis, the layering correction and the discipline of printing all three counts, not a uniquely discriminating forecast.**

Rigour checks not met, listed by name.

1. **No as-of date, because no real corridor was measured.** A concentration figure from real flows requires its source and its date; every share vector here is constructed. Appendix A makes both mandatory for any application to a real structure.
2. **No uncertainty is reported anywhere, because nothing is estimated.**
3. **The dimension of η is supplied by this paper and not by the source.**
4. **Q, C and h have no operational unit**, and the transit model's currency figures are never added to them.
5. **The surplus law's large-N limit is a defect, not a result.**
6. **The conjoint conditions were not tested for λ, r or the price elasticity** [Michell, 1997, Luce and Tukey, 1964].
7. **No distributional or scaling claim is made, so the Clauset–Shalizi–Newman procedure was not run** [Clauset et al., 2009]. Any restatement reaching for a heavy-tail claim about payment shares must run it first and must read Stumpf and Porter [2012] on why passing it still would not identify a mechanism.
8. **The distributed-systems, network and cascade borrowings are category 2, not category 1** [Fischer et al., 1985, Gilbert and Lynch, 2002, Buldyrev et al., 2010, Bikhchandani et al., 1992].
9. **The constant-hazard assumption is untested.** No failure-time data was examined and no alternative hazard shape was fitted.
10. **Aggregation is unaddressed** [Kirman, 1992].
11. **Expected value and time average are not the same thing, and every figure here is the first** [Peters, 2019].
12. **The instrument is performative and its own parameters are not exempt** [MacKenzie and Millo, 2003, Lucas, 1976].
13. **There is no conservation law for Q** [Mirowski, 1989].
14. **Three arithmetic discrepancies in the source are reported rather than corrected in silence** (Section 9.9).

11. Conclusion

This paper measures the settlement rail, and the following statements are the fixed text of its results.

R1. The clearing-perimeter measurement: a naive correspondent HHI of 0.302400, reading 3.306878 effective correspondents, collapses to a perimeter HHI of 0.745256 and $N_{\text{rail}} = 1.341819$ once flows are decomposed by settlement currency and then by the perimeter in which each currency achieves finality. The rail-axis surplus term is $0.35 \ln(1.341819) - 0.500000 = -0.397091$, against -0.122417 on the jurisdictional axis — **3.24 times worse on the axis nobody runs.**

R2. Adding correspondents raises the naive count by 33.8 per cent to 4.425092 and leaves N_{rail} at 1.341819 exactly; **only currency redenomination moves it**, to 2.087683 with $h_{\text{rail}} = -0.242381$, a gain of $+0.154710$ obtained with no new entity, no new jurisdiction and no new bank.

R3. The transit-loss model $\mathbb{E}[\text{loss}] = V(1-e^{-\lambda t})(1-r)$ is **invariant to hop count at USD 68,000** on the constructed corridor, while the probability of any loss rises from 0.2000 to 1.5889 per cent across eight hops and the worst single loss falls from USD 34.0m to USD 4.25m. **The window, not the hop size, is the variable.**

R4. Multilateral netting to a single settlement position cuts the constructed group's expected transit loss by **77.0 per cent**, with no new rail and no new counterparty.

R5. The indifference hazard: replacing a 48-hour correspondent transit with a half-hour rail is justified above a **0.2369 per cent** per-transit freeze hazard and not below it. Counterparty pre-screening returns **2.57 to one** on the constructed supplier corridor and pays for itself if it removes 31.1 per cent of the counterparty hazard.

R6. Enforcement-driven exit of licensed informal-transfer providers from twelve to five moves the index from **1,074 to 2,406 points** and the effective provider count from 9.310987 to 4.156276, a fall of **55.4 per cent** — clearing the published structural-presumption screen by 606 points with no transaction and no filing. One layer down, the five survivors bank through **1.650710** effective correspondents, so one institution's exit removes 73.0 per cent of the corridor's licensed capacity.

Counter-result. After the largest available intervention the rail remains the structure's weakest axis, at 2.087683 effective perimeters against 2.941176 on the jurisdictional axis and a break-even of 4.172734. The residual dollar share of 52.0 per cent **is a property of the world rather than of the design**, and the depository layer is not remediable by a private group at all. A concentration that cannot be fixed should be sized, not ignored.

Two further results are established as method rather than as measurement. **A clearing system and a messaging system are different layers and are not substitutes**, so a structure reporting them as two rails has recorded one dependency twice — and measured correctly, one messaging network at 100 per cent and 85 per cent on that network plus 15 per cent on a second clearing system messaged over the first are the same number. **And expected transit loss is flat in hop size**, so splitting a payment is a purchase of insurance against the worst case at a stated premium and not a reduction in risk; any presentation claiming otherwise is wrong as a matter of algebra rather than of judgement.

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every case in this paper — the corridor set, the correspondent and currency shares, the perimeter assignments, the hazard and recovery parameters, the cost model, the covenant tolerances, the corridor value and price elasticity, and the provider shares — is **CONSTRUCTED** and is labelled as such at each point of use. The numerical results are exact computations on those constructed inputs, carried at full precision and rounded only for printing.

All figures were produced by the scripts named in their captions (`figs/src/CCP-061_fig01.py` through `figs/src/CCP-061_fig05.py`), which import the shared figure-style module and read their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

Published measurements referred to in the source manuscript are described and not quoted as evidence. The bibliography contains no official statistics, statutes, regulations, registry documents or case reports for any individual jurisdiction, operator or institution. Appendix B lists the claims affected. Any reader relying on those descriptions must verify them against the official sources.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. Every institution, corridor, counterparty and transaction described is invented. Section 7.6 concerns the aggregate cost of a policy to a class of senders and uses no individual data of any kind.

Legal, tax and regulatory advice

This paper is analytical and is not legal, tax, regulatory or investment advice. It describes payment infrastructure, settlement mechanisms, correspondent relationships and value-transfer arrangements for the purpose of measuring concentration and pricing structural choices. No recommendation is made in respect of any particular person, institution, corridor or rail. Messaging and clearing arrangements, settlement-finality regimes, sanctions and designation instruments, licensing perimeters and the standards applicable to correspondent banking all change, frequently and at short notice; **every such statement in this paper is carried from the source manuscript, is flagged in Appendix B, and requires confirmation with qualified counsel and with each operator's own current published material before any reliance.** The rail layer changes faster than any other subject in this series, and a two-year-old rail memorandum is worse than none, because it will be trusted.

Appendix A. Reporting standards

A.1 The rail register. Any structure measured on this axis reports the following, and nothing less.

```
1  the corridor inventory: counterparty, corridor, annual value,
   with every intermediary traced
2  correspondent shares and the naive effective correspondent count
```

```
  3  the currency decomposition of each correspondent's volume
  4  currency shares and the effective currency count
  5  the mapping of each currency to its clearing perimeter, with
     the reasoning recorded per currency
  6  perimeter shares, the perimeter HHI and N_rail
  7  all three counts reported together, never one alone
  8  ln(N_rail), eta ln(N_rail) and h at the stated parameters
  9  direct versus indirect participation, verified bank by bank
 10  the nesting position: which respondents have their own
     respondents
 11  the custody chain mapped to the depository, not to the
     custodian
 12  the date, and the difference against the previous measurement
```

Item 7 is the one that prevents the error this paper measures, and item 12 decays first, because the rail layer changes faster than any other.

A.2 The transit policy. A written document, approved at board level and reviewed annually, stating: λ, r and the per-hop cost per corridor with a named owner for each; the maximum exposure window per corridor and the rail selected to meet it, **the window being the primary control**; the maximum hop size derived from the covenant tolerance rather than chosen; the maximum concurrent transit exposure across all corridors, **because eight hops in flight simultaneously are one hop for systemic purposes**; the counterparty pre-screening protocol, including the chain and not only the direct counterparty; rail, issuer, network and venue diversification limits; the mid-transit runbook with its documentation pack pre-assembled and tested annually; the balance limit above which a settlement balance is reclassified as a holding and enters the jurisdictional measurement; the landed-cost model per corridor; and the annual difference against last year's policy and last year's realised experience.

Appendix B. Claims requiring verification against primary sources

The following statements rest on regulatory, operational or statistical detail that changes, and **none is cited to the bibliography, which contains no official statistics, statutes, regulations or case reports for any individual jurisdiction, operator or institution.** They are reproduced from the source manuscript and require confirmation against the official sources before any reliance. The source manuscript's own verification registers run to more than seventy numbered items across the three chapters and are condensed here without loss of subject.

1. **The treatment of the euro area as a Western clearing perimeter.** This is a contestable call with consequences for every measurement in Section 7.1. The coordination test should be applied to it explicitly and the alternative computed. **This is the single most consequential unverified assumption in the paper.**
2. **The messaging cooperative**: founding and live dates, legal form, oversight arrangements and the composition of the central-bank oversight group; the current institution count and daily message volume; the message types referred to and the cover-payment standard.
3. **The message-standard migration**: the current end date of the coexistence period, which has been revised, and the current status [International Organization for Standardization, 2013].
4. **Messaging disconnections and reconnections of designated institutions**: the instruments requiring each, their articles, their dates and their scope, in each episode referred to.
5. **The 2022 immobilisation of central bank reserves**: the prohibiting article and amending instrument, the approximate total, the proportion held at any one institution, and

all subsequent measures concerning the extraordinary revenues. **Every number and instrument here needs primary sourcing.**

6. **The renminbi clearing system**: launch and phase dates, the operator's corporate form and its relationship to the central bank, current direct and indirect participant counts and country coverage, and the messaging arrangements available to each class of participant.
7. **Cleared values**: annual value through the renminbi system and daily values through the two principal dollar systems, from each operator's own published statistics. **No figure in this paper depends on these and none is quoted.**
8. **Currency payment-share trackers**, and the extent to which traffic native to a clearing system is excluded from the series.
9. **The other messaging and regional clearing arrangements** referred to: operators, establishment dates, participant counts, currencies supported and functional scope for each; and any instrument prohibiting operators in one bloc from using another's network.
10. **Cross-border central bank digital currency projects**: participating institutions, milestones and dates, and the current status and ownership of each.
11. **Correspondent-banking decline data**: the specific reports, the measured percentage decline in active relationships, the period and the corridor breakdown. The source states "in the region of a quarter" and **this paper uses no figure.**
12. **Enforcement resolutions**: the two well-known resolutions referred to in Section 4.1, their amounts, charging instruments and terms. **This paper quotes neither amount.**
13. **The correspondent-banking due-diligence questionnaire** and the international recommendations on correspondent banking and wire transfers: current version, current text and interpretive notes [Financial Action Task Force, 2012].
14. **The settlement-failure episode of 1974** and the payment-versus-payment institution that answered it: the failure date and mechanism; the current eligible-currency list, membership and average daily value settled.
15. **Settlement-finality legislation**: the finality provisions, the designation mechanism, the implementing and onshored instruments in each relevant jurisdiction, and the wholesale funds-transfer code and central-bank regulation governing irrevocability in the other.
16. **The depository layer**: the number of international depositories, the concentration of cross-border holdings and any published measure of it. **The effective count of 1.15 quoted in the source is a constructed illustration and is not a measurement of any real market; it is not used in any result in this paper.**
17. **Direct-participation status of any named institution** in any clearing system, which must be answered bank by bank rather than jurisdiction by jurisdiction.
18. **Private-order enforcement**: the account of an eleventh-century traders' coalition and the model of the law merchant with a private judge, both named in the source's prose without citation keys and **absent from this paper's bibliography** (Section 10).
19. **Exchange-rate regimes and parallel premia**: the episodes named in the source, their dates and the size of each premium; and regional remittance cost data.
20. **Regulatory perimeters for fast rails**: issuer, reserve, redemption and licensing requirements in each relevant jurisdiction, several of which are less than three years old and none of which is stable.
21. **All cost figures, hazard rates, corridor values, elasticities and share vectors in this paper are constructed.** They are not quotations, benchmarks or estimates of any real corridor, and no real fee schedule, tariff or published statistic was used to produce them.

All numerical results in this paper were computed from the stated inputs at full precision and rounded only for printing.

Chapter 62

The Carve-Out Cascade: Pledge Triggers, Antitrust Walls, and Legal-Family Interfaces Priced as Cost

Abstract. A listing that carves out a minority creates the collateral against which the retained block is pledged, and the pledge converts an ownership position into a path-dependent one. We derive the trigger d* = 1 - LTV0/(1 - m) and show that a 50 per cent loan-to-value advance against a 35 per cent maintenance requirement produces the first margin call at a 23.076923 per cent drawdown, not the 50 per cent that borrowers report believing they have. The forced sale that follows takes the principal to exactly 50.000000 per cent of the company at a 36.000000 per cent drawdown and to 38.095238 per cent at 40 per cent, and exhausts the collateral equity at 50 per cent. An unlevered route raising the same cash by floating 60 per cent crosses the levered route at a 39.393939 per cent drawdown. Under a constructed first-passage model at 40 per cent volatility, the ten-year probability of touching the control-losing barrier is 76.347 per cent. Cross-default converts nine agreements into thirty-six pairwise links and a 36.9751 per cent annual probability that one trips at a 5 per cent per-agreement hazard. On the antitrust side the Hart-Scott-Rodino delta reduces to 2s1s2, so the 100-point test binds exactly when s1s2 > 50, and the Clayton section 8 de minimis provisions exempt a 1.99/99.00 pair while catching a 4.10/4.10 pair. Legal-family diversity is reported as a cost line, not resilience.

Keywords: pledged shares; maintenance requirement; forced sale; cross-default; interlocking directorates; merger screening; conflict of laws; characterisation

Series: D. Structure and Resilience **Paper:** CCP-062

1. Introduction

A holder who sells a minority of a company to the public and pledges the retained majority against a loan has performed an elegant, lawful, fully disclosed piece of engineering. This paper computes what it costs, and the cost is not where the borrower believes it is.

The listing is what creates the collateral. An unlisted majority block of a private company is not security a lender will advance against at any tolerable rate, because it has no daily

mark, no observable liquidity, no mechanism by which the lender can convert it into money quickly, and no third party willing to price it on a Wednesday afternoon. Listing a minority manufactures all four of those properties for the majority. The float is the price discovery; the exchange is the exit; the daily close is the mark the maintenance clause is tested against. **The same act that creates the borrowing capacity creates the daily mark that the margin call runs off, so the failure mode is not a bolt-on to the design — it is the design, seen from the other side.**

The consequence is one line of arithmetic. With an opening loan-to-value LTV_0 and a maintenance requirement m — the borrower's equity in the collateral must stay at or above m of collateral value — the first margin call falls at a drawdown

$$d^{*} = 1 - \frac{\mathrm{LTV}_0}{1 - m}.$$

At $\mathrm{LTV}_0 = 50\%$ **and** $m = 35\%$ **the first call falls at a 23.076923 per cent drawdown, not at the 50 per cent that borrowers report believing they have.** A fall of a little under a quarter is what equity markets deliver on ordinary business: a rate surprise, an earnings miss, a sector rotation, a short report, a legal setback at a competitor.

What follows the trigger is worse and is equally deterministic. The forced sale that restores the maintenance ratio takes an eighty per cent holder to **exactly 50.000000 per cent at a 36.000000 per cent drawdown**, to 38.095238 per cent at 40 per cent, and exhausts the collateral equity at 50 per cent. **None of this is misfortune. The trigger is a deterministic function of two numbers agreed at signing, the forced-sale quantity is a deterministic function of the trigger and the price, and the control outcome is a deterministic function of the forced-sale quantity.** No relationship, reputation or negotiation changes it, because a maintenance clause is a mechanical valuation test run against a published close and the risk officer applying it has no discretion over it.

1.1 What this paper adds

Three measurements and three identities, all checkable in a line.

First, the trigger and the control-loss surface, computed as closed forms and tabulated, together with the crossover against the unlevered route that raises the same cash by floating more of the company. The two routes cross at a 39.393939 per cent drawdown: **below it the pledge retains more of the company, above it the plain float does.**

Second, the propagation arithmetic. Cross-default converts nine agreements into thirty-six pairwise links and a 36.9751 per cent annual probability that one of them trips at a five per cent per-agreement hazard. The same combinatorial object turns ten legal touchpoints into forty-five interfaces, and routing them all through one apex cuts that to nine — **an 80.0000 per cent reduction obtained by changing the topology and nothing else.**

Third, two antitrust identities that make a regulatory constraint checkable rather than predictable. The index change on combining two firms with percentage shares s_1 and s_2 is exactly $2s_1s_2$, so the hundred-point screen binds exactly when $s_1s_2 > 50$; and the de minimis provisions of the interlocking-directorate prohibition are asymmetric in a way that exempts a 1.99 / 99.00 per cent pair and catches a 4.10 / 4.10 per cent pair.

And one demotion. Legal-family diversity is reported here as a **cost line, not a resilience line.** A structure spread across more legal families is not more resilient; it is more expensive, because each family boundary is an interface requiring recognition, characterisation and enforcement opinions. **Reporting legal-family diversity as though it were coordination-perimeter diversity is the same category error as reporting a naive jurisdictional index on a single-perimeter structure**, and it is made by the same people for the same reason.

1.2 The standing constraint

Every arrangement examined here is measured under the corpus's formation test: it qualifies only if it would still perform its function were every element disclosed to every authority with jurisdiction over any part of it. That test is satisfied by construction throughout this paper. **A pledge over listed securities is disclosed in the issuer's filings; an interlocking board seat is in the annual report; a merger above the threshold is notified before closing; and a legal-family interface is a written opinion that exists precisely to be produced.** The open-design principle [Kerckhoffs, 1883, Saltzer and Schroeder, 1975] and the older argument for publicity as a discipline [Brandeis, 1914] are the same constraint stated elsewhere. **In this subject disclosure is not a cost of the structure; it is the structure's evidence base**, and the failure mode is not that a fact becomes known but that an arithmetic consequence of a known fact was never computed.

The paper presupposes the founding reframing of Curzi [2026] — comfort is the felt total of present costs, so a present arrangement is a local minimum of the cost landscape by definition rather than by luck. **That is why the cascade is built: at signing it is a local minimum, and the arithmetic that destroys it operates only on paths the borrower has not priced.**

2. Hypotheses

Four hypotheses, each falsifiable in the sense Platt [1964] requires and each carrying a point prediction with a constant attached rather than a direction, which is what Meehl [1967] asks of a claim that intends to be tested.

H1 — the trigger hypothesis. The first margin call on a pledged block is a deterministic function of the opening loan-to-value and the maintenance requirement alone, and is far shallower than the loan-to-value suggests. *Predicted:* $d^{*} = 1 - \mathrm{LTV}_{0}/(1-m)$, so at $\mathrm{LTV}_{0} = 50\%$ and $m = 35\%$ the call falls at 23.076923% and not at 50 per cent; at $m = 35\%$ every five points of additional loan-to-value costs exactly 7.692308 points of cushion. *Falsified by:* a facility whose first call depends on the size of the block, the identity of the borrower, or anything other than those two parameters.

H2 — the control-loss hypothesis. The forced sale that restores the maintenance ratio costs majority control at a drawdown computable in advance, and the levered and unlevered routes cross at a single point. *Predicted:* an eighty per cent holder reaches exactly 50.000000% at a 36.000000% drawdown, 38.095238% at 40 per cent, and zero collateral equity at 50 per cent; the levered and unlevered routes raising identical cash cross at 39.393939%. *Falsified by:* any of those four figures failing to reproduce from the sale rule $x = V - (V - L)/m$.

H3 — the base-case hypothesis. For a newly listed, concentrated, single-name equity, touching the control-losing barrier is not a tail event over a structural horizon. *Criterion, fixed in advance:* under the constructed first-passage model at $\sigma = 40\%$ and $\mu = 6\%$, the ten-year probability of touching the 36 per cent barrier exceeds one half. *Predicted:* 76.347%, with a median time to the 23.076923 per cent barrier of 0.8786 years. *Falsified by:* a volatility low enough to put the ten-year probability below one half — which at $\mu = 6\%$ requires σ below roughly 28 per cent, and **a reader who believes the security is that quiet should substitute it and re-read the table.**

H4 — the delta-identity hypothesis. The merger-screen change test has a closed form that makes it checkable without a model. *Predicted:* $\Delta = (s_1 + s_2)^2 - s_1^2 - s_2^2 = 2 s_1 s_2$ exactly, so the hundred-point screen binds exactly when $s_1 s_2 > 50$, which for an equal pair is 7.0711% each; and a sequential roll-up produces a strictly larger cumulative change than the same aggregate bought in one step. *Falsified by:* an arithmetic counterexample, of which there is none, so **H4 is an identity and is labelled as one rather than presented as a finding.**

Not hypotheses, and labelled as such.

- *The first-passage model itself* — geometric Brownian motion with constant drift and volatility — is a **modelling choice**. It is not fitted to any security and its parameters are

stipulated. Section 10 states that the constant-volatility assumption is the least defensible thing in the paper.
- *The 35 per cent maintenance level* and the 50 per cent opening loan-to-value are **stipulated constructed parameters**, plausible for a concentrated single-name block and not a quotation from any facility.
- *The claim that borrowers systematically misread a loan-to-value ratio as a drawdown cushion* is **this work's own conjecture and is left uncited.** The bibliography contains measurement theory but no judgment-and-decision literature that supports it, and Section 10 records the gap rather than attaching a citation that would not survive a check.
- *The identification of legal-family count with a cost* rather than a resilience measure is a **definition**, and its content is in the interface count rather than in the identification.

3. Method

3.1 The cascade, and the two closed forms

CONSTRUCTED CASE throughout. All facts invented. An operating company whose equity is worth 100 units, wholly owned. Sell 20 per cent to public investors; retain 80. Two economically different versions exist and the distinction governs everything downstream: a primary offering issues new shares and the proceeds go to the company, while a secondary offering sells existing shares and the proceeds go to the selling holder. **The cascade assumes a secondary or mixed offering, because it assumes the holder ends with cash. State which one before modelling anything: the two produce identical ownership percentages and opposite cash positions.**

The retained block is now listed securities with a daily close, a public register, a settlement system and a market. Those are the four properties a securities-backed lender requires and none of them existed before the listing. Pledged at $\text{LTV}_0 = 50\%$:

```
collateral value at inception   80.000000
advance                         40.000000
cash from the offering          20.000000
TOTAL DEPLOYABLE                60.000000
retained ownership              80.000000%
```

The alternative that raises the same 60 without leverage is to float 60 per cent and retain 40. **Forty percentage points of retained control is what the leverage buys**, and it is the difference between a group whose apex controls its operating subsidiary and one whose apex holds a large minority in a company run by somebody else.

The first closed form is the trigger. The maintenance test is $(V - L)/V \geq m$, equivalently $L \leq (1 - m)V$, so the collateral value at which the call arrives is $V^* = L/(1 - m)$ and

$$d^* = 1 - \frac{V^*}{V_0} = 1 - \frac{\text{LTV}_0}{1 - m}.$$

The second closed form is the forced sale. At a drawdown d with price $p = 1 - d$ and collateral value $V = V_0 p$, the proceeds required to restore the maintenance ratio solve $(V - x - L)/(V - x) = m$, giving

$$x = V - \frac{V - L}{m}, \qquad \text{shares sold} = \frac{x}{p}.$$

The required sale grows faster than the price falls, which is why the margin call arrives when the shares are cheapest and selling them is most destructive. That is not bad luck; it is the definition of the instrument, since the call is triggered by the fall and the call and the low price are the same event.

3.2 Why the maintenance clause is not a covenant

A covenant is a promise: breach it and there is a default, a notice, a cure period, negotiation, waivers, amendments, fees and lawyers. A maintenance requirement in a securities-backed facility is usually not that. It is a mechanical valuation test run against a published closing price, and the consequence of failing it is a demand for cash or additional collateral within a short window — often hours rather than days — after which the lender may sell. **Corporate debt fails slowly and negotiably; margin debt fails quickly and mechanically, and the speed is the point of the product from the lender's side.**

That difference is also why the maintenance level is a measure of the *lender's* safety and not the borrower's. The lender priced the facility on the assumption that it can liquidate before the collateral falls below the loan, and the maintenance level exists to give it that assumption with room to spare.

3.3 The first-passage model

Model the share price as geometric Brownian motion with drift μ and volatility σ, and use the standard first-passage result for the probability that the running minimum touches a barrier B below the starting value S_0 within horizon T. With $b = \ln(B/S_0)$ and $\nu = \mu - \sigma^2/2$,

$$P = \Phi\left(\frac{b - \nu T}{\sigma\sqrt{T}}\right) + e^{2\nu b/\sigma^2}\,\Phi\left(\frac{b + \nu T}{\sigma\sqrt{T}}\right).$$

CONSTRUCTED MODEL. Illustrative only. Not an estimate of any real security's behaviour. The model is used as a **shared mathematical form**, not as a derivation: no claim is made that any particular security's returns are log-normal with constant volatility, and Sornette [2003] is the standing reason to doubt it — the empirically interesting behaviour of equity prices lives in the fat tails and in endogenous feedback that a constant-volatility diffusion cannot produce. Longin and Solnik [2001] and Ang and Chen [2002] measure the asymmetry directly: extreme correlation is concentrated in downturns. **Both point the same way — the model almost certainly understates the probability of the barrier being touched — so the results here are read as a lower bound rather than as an estimate.**

A distinction the corpus insists on is kept here. Curzi [2026] states its own path-weighting pedigree honestly as exponential, Boltzmann-style weighting of states by cost, e^{-R}, and explicitly not as a phase-based path amplitude. **The computation in this section is a genuine first-passage probability for a stochastic process, and it is neither of those things**; conflating a weighting scheme with a probability is exactly the move that discipline exists to prevent.

3.4 The combinatorics of links and interfaces

Two of this paper's results are the same combinatorial object seen twice. A set of n agreements each cross-defaulting to every other has $\binom{n}{2} = n(n-1)/2$ pairwise links; a set of n legal touchpoints each having to interface with every other has the same count. **At $n = 9$ that is 36 links; at $n = 10$ it is 45 interfaces.** Routing every interface through a single apex replaces the complete graph with a star and the count falls to $n - 1$.

That adding units inside one boundary adds coordination cost at order $n(n-1)/2$ while adding nothing to independence is the general organisational statement of the same arithmetic [Brooks, 1975]. **The corpus carries an exhaustive verification of a pairwise identity over exactly this object: Curzi [2026] proves its doubling theorem and checks it over all sixty-six pairs at $N = 12$, and $\binom{12}{2} = 66$ is the same $\binom{n}{2}$ that produces the thirty-six links and the forty-five interfaces here.**

Under independent per-agreement annual breach hazard q,

$$P(\text{at least one breach}) = 1 - (1-q)^n,$$

and with universal cross-default that is the probability the whole group defaults. **Independence is the assumption doing the work and it is the wrong assumption**, because a breach is more likely in exactly the conditions that make other breaches more likely. Buldyrev et al. [2010] is the formal statement that coupling destroys the independence a naive count assumes, used here as a **shared structural form and not as a derivation** — no percolation premise has been shown to hold for a set of financing covenants — and Watts [2002] is the closest available model of the propagation itself: a global cascade is possible only in a connected region of parameter space, and adding links moves a system into and then out of that region. Perrow [1984] reached the organisational version first and is cited as framing. Albert et al. [2000]'s robust-yet-fragile shape is used for its shape alone, with Doyle et al. [2005] as the standing correction to reading too much into it.

The engineering response follows from the arithmetic and not from judgement: cut the number of links rather than negotiate each hazard down. Curzi [2026] is the corpus's general result of that form — moving from terminal inspection to per-stage checking across seven stages at a per-stage defect rate of 0.10 raises final yield from 47.83 per cent to 93.21 per cent at 0.63 reworks per unit. **Universal cross-default is terminal inspection: one failure anywhere condemns the whole run. Materiality thresholds, grace periods and de-linking are per-stage checks, and the yield arithmetic is the same.**

3.5 The concentration instrument, and which effective count

For shares s_i summing to one on a stated basis, HHI $= \sum_i s_i^2$ and the effective count is $1/\text{HHI}$. The index is Hirschman [1945]'s and, independently in sum-of-squares form, Herfindahl [1950]'s; the numbers-equivalent reading is Adelman [1969]'s; Stigler [1964] is the oligopoly argument that made it the regulator's instrument. Competition practice states it in points on percentage shares with a structural presumption above 1,800 points and a change above 100 [U.S. Department of Justice and Federal Trade Commission, 2023, 2010]; the resilience application states it as a fraction.

Which effective count. $1/\text{HHI}$ is the Hill number of order two [Hill, 1973], the numbers-equivalent that Jost [2006] distinguishes from the index itself, and $\exp(H)$ for Shannon entropy [Shannon, 1948] is the same family at order one. **The corpus's other effective count, the affine $N_{\text{eff}} = 1 + (N-1)u$ of Curzi [2026c], is a linear interpolation between 1 and N, is not a Hill number of any order, and carries no external support**; that paper badges it as a design axiom rather than a theorem. **Every effective count in this paper is $1/\text{HHI}$ on a stated share basis**, and the discipline Hurlbert [1971] demands — the word carries no content without its formula — is imposed throughout. The general correlation-aware form $1/(w'Rw)$, its identity with the naive and bloc indices under two correlation assumptions, and its correspondence with the similarity-sensitive diversity of order two [Leinster and Cobbold, 2012] are established in Curzi [2026] and are used here without restatement.

No distributional or scaling claim is made anywhere in this paper. A product market descending 22 / 18 / 14 / 11 / 8 / 7 / 6 / 5 / 4 with a fringe, or a share-price path with a fat left tail, invites one. **No distribution is fitted, no tail is characterised and no power law is asserted.** Clauset et al. [2009] is the reason the procedure is named and not run: a principled maximum-likelihood fit with a goodness-of-fit test and a likelihood-ratio comparison leaves few canonical data sets standing, and a straight line on log-log axes is not evidence. Stumpf and Porter [2012] adds why passing it would not be enough. **Every effective count here is a statistic of a finite share vector, not an estimate of a distribution, and the first-passage model is a stipulated diffusion rather than a fitted process.** Small-sample bias is not addressed [Chao et al., 2014], and effective-number measures behave poorly where one unit dominates [Golosov, 2010].

3.6 Identification, reflexivity, and the status of the perimeter assignment

The perimeter assignments used in Section 7.5 are inputs to the instrument and carry an identification problem. Treating a set of jurisdictions as one coordination perimeter is a statement

about a group influencing its members while the members constitute the group. **Manski [1993] proves that the endogenous effect in a linear-in-means reflexive specification is not identified from observational data.** The remedy adopted here is stipulation with a defence, the route Brock and Durlauf [2001] shows can be made rigorous, and **the price is that every perimeter figure in this paper is a modelling position rather than an identified empirical result.**

The collapse mechanism has a formal ancestry and is not new machinery. Granovetter [1978] establishes that the distribution of individual thresholds rather than the mean determines whether a collective outcome tips, and Schelling [1971] is the canonical demonstration that a mild individual preference produces a sharp aggregate outcome — **the strongest available warning against reading an individual-level parameter straight into an aggregate.** A margin cascade is the same object with a mechanical threshold in place of a preference.

A second reflexivity is specific to this instrument and is the sharpest available illustration of the general problem. A forced sale into a falling market is a case in which the act of measuring and responding changes the measured quantity: Soros [2013] is the standing statement of that circularity in markets, and Sornette [2003] the argument that endogenous feedback rather than exogenous news produces the largest moves. Under a constructed permanent-impact model in which selling one float's worth of stock moves the price by 15 per cent, the sale itself takes the price from 0.640000 to 0.496000, the residual equity ratio falls to 16.129032 per cent, **and the position is called again before the first call has settled.** The model is constructed; the sign is not.

4. The pledge as an instrument

4.1 Structural subordination

The 100 in the cascade is the *equity* value of the operating company. If that company carries its own debt, the debt sits ahead of the equity in its own waterfall and the pledged block is a claim on the residual. Two consequences follow, both arithmetic.

The recovery profile has a cliff. With senior debt of 35 ahead of the pledged block, the holding-company lender is fully covered over a wide range of enterprise values and then is not covered at all, because a residual claim's value reaches zero at a strictly positive enterprise value. **The lender knows this, which is why its covenant package forces liquidation long before enterprise value reaches the region where its recovery is impaired.**

The drawdown in the collateral is a multiple of the drawdown in the business. With equity of 100 and senior debt of 35, enterprise value is 135 and the equity beta to enterprise value is $135/100 = 1.35$. The trigger and control-loss drawdowns therefore translate to 17.094017 per cent and 26.666667 per cent of *enterprise* value. **A 26.666667 per cent decline in the value of the business costs the principal control of it**, and that is the sentence usually missing from the proposal.

4.2 Change of control, and the second-order cascade

Enforcement of a share pledge transfers shares, and transferring shares is a change of control. Change-of-control provisions are the most widely distributed clause in a corporate group and sit in facilities, indentures, joint-venture agreements, key customer and supplier contracts, licences and franchises, regulated permissions, property leases, incentive plans, and government incentives whose eligibility conditions are re-tested.

The last of those is where the cascade does its real damage. Where an operating company holds a promotion or a treaty-based ownership right whose conditions include the identity or nationality of its shareholders, **a pledge enforcement in one financial centre can extinguish an operating privilege in another, and it does so without anybody in the second being consulted.** That is a financing decision at the holding layer reaching through the group into an operating

condition that was never part of the credit analysis. The corpus's succession arithmetic is the same shape one layer over: Curzi [2026] splits a forced-heirship exposure into a situs-capped levy of USD 6.70m — 11.2 per cent of the notional reserved share — and an unaffected tax exposure of USD 27.00m, **and in both cases the number that matters is produced by a rule in a jurisdiction nobody in the transaction was looking at.**

The engineering response is a change-of-control map: every agreement in the group with such a provision, its trigger threshold, its consequence, and whether consent can be obtained in advance. It is dull work, it takes weeks, **and the number that comes out of it is frequently larger than the loan.**

Why the map is the binding constraint rather than the drafting is worth stating in the corpus's own terms. Curzi [2026x] computes a receiver-limited channel capacity — of order 70 to 106 bit/s at a mechanotransduction bandwidth near 10.6 Hz and 20 dB signal-to-noise, delivering a 40-bit address in about half a second — and the operative lesson is that **throughput is set by the receiver and by the narrowest element, not by the sender's intent.** The analogy is used as a shared form and nothing more: a group's ability to act on an enforcement is bounded not by how much has been written down but by how quickly the constrained party can be told and can respond, and the margin window is hours.

4.3 Rehypothecation

Where collateral is delivered to a lender or intermediary, the documents may permit that party to use it — to re-pledge, lend or sell it, subject to an obligation to return equivalent securities. The commercial rationale is real, since the lender's ability to fund itself against the collateral is part of what makes the borrowing cheap.

The risk is a change in the character of the borrower's claim. Where the collateral is held subject to a security interest and segregated, the borrower's right on redemption is proprietary and the specific securities come back. Where it has been transferred outright or re-pledged onward under a right of use, the borrower's right may be a contractual claim for equivalent securities — and in the counterparty's insolvency a contractual claim ranks with the other unsecured creditors.

right-of-use limit	block re-pledged onward	proprietary claim preserved
0%	0.00	80.00
25%	20.00	60.00
50%	40.00	40.00
100%	80.00	0.00
140%	80.00 (all of it)	0.00

A right of use above 100 per cent of collateral value does not mean more shares leave; it means all of them do, with room to spare. The excess is about the intermediary's funding rather than about the block, and the consequence for the borrower is identical at every level from 100 per cent upward. Two disciplines follow: negotiate the right of use explicitly and cap it, ideally at zero for a control block; and read the governing law and the account structure, **because whether a claim is proprietary or contractual is decided by insolvency law and by the terms of the arrangement** [Torremans et al., 2022, Fawcett and Carruthers, 2008].

5. The walls a group design has to clear

The material in this section is described from the source manuscript. **This paper's bibliography contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction, so the statutory provisions described here are uncited**; the published enforcement guidelines are the exception and are cited where they are used [U.S. Department

of Justice and Federal Trade Commission, 2023, 2010]. Appendix B itemises every affected statement.

The two provisions are named once here for identification and are thereafter described by function. The interlocking-directorate prohibition is section 8 of the Clayton Act and the pre-merger notification regime is the Hart–Scott–Rodino Act, both of the United States, and the abstract names them for the same reason. **Naming a statute is not citing it, and nothing in this paper rests on a statutory text this bibliography contains.**

5.1 The interlocking-directorate prohibition, and why *per se* changes the engineering

One major jurisdiction prohibits a person from serving simultaneously as an officer or director of two corporations that are competitors such that an agreement between them would be unlawful. Two features make it different from the rest of competition law and both are structural rather than doctrinal.

It is *per se*. There is no rule-of-reason analysis, no efficiencies defence and no argument that the interlock was harmless. **For an engineer that is good news, because a *per se* rule is a test that can be run rather than a judgement that must be predicted.**

It has statutory de minimis exceptions with fixed percentages, and the percentages are the design surface. Alongside a dollar threshold on the corporations' capital, the exemptions turn on *competitive sales* — the gross revenues for products and services sold by one corporation in competition with the other — as a fraction of each corporation's total sales: an interlock is exempt where **either** corporation's competitive sales are below 2 per cent of its total sales, or where **each** corporation's are below 4 per cent. **The two tests are asymmetric, and the asymmetry is where the design room is.**

Three further mechanics decide real cases. **The reach of the prohibition beyond corporations is contested** — whether it covers other entity forms, or a fund's designation of an individual to a portfolio-company board, is a live question that belongs to counsel. **Board observers are the standard workaround and the standard trap**: a person who attends, receives the board pack, participates and reports back to a competitor is doing the thing the prohibition exists to prevent, whatever the label. **And there is a short grace period for interlocks that become unlawful through changed circumstances** — which is the version of the violation a group actually commits. **Not by taking a prohibited seat, but by making two lawfully seated companies into competitors and not re-running the test.** The trigger event is an acquisition elsewhere in the group, and the person who ran the acquisition does not know what boards the directors sit on.

5.2 The notification regime, and the arithmetic of the change test

The pre-merger notification regime requires filing above stated thresholds and observance of a waiting period before closing. The thresholds are adjusted annually and are for counsel. Two mechanics are the whole risk profile. **The avoidance rule means that the cleverness that gets a transaction under the threshold is the evidence that it should have been filed.** And **the per-day penalty means a filing error left undiscovered is a liability that grows linearly with inattention.**

The substantive test the filing feeds is a concentration screen, and its arithmetic has a closed form. When two firms with percentage shares s_1 and s_2 combine,

$$\Delta = (s_1 + s_2)^2 - s_1^2 - s_2^2 = 2 s_1 s_2,$$

so the hundred-point change threshold binds exactly when $s_1 s_2 > 50$ [U.S. Department of Justice and Federal Trade Commission, 2023]. **For an equal-share pair that is $\sqrt{50} = 7.0711$ per cent each.** Section 7.4 runs it on a constructed market.

5.3 Where the pledge meets the group design

The two subjects are one subject. A pledge over a control block is a disclosed position whose enforcement is a change of control; a change of control is a competition event as well as a contractual one; and the acquirer produced by an enforcement sale is a party whose own holdings enter both screens. **The instrument that computes the pledge trigger and the instrument that computes the merger screen are the same sum of squared shares pointed at different bases**, which is why they belong in one paper.

6. Legal families as a cost line

6.1 What the families actually do differently, and what they do not

The canonical division — common, civil, mixed, and religiously grounded systems — differs in characterisation, in the treatment of reserved portions on death, and in the situs rules that decide which forum governs an asset. **Recognition of a foreign instrument does not imply enforcement of it, and the gap between the two is where the money goes**; the comparative doctrine is a settled literature [Torremans et al., 2022, Briggs, 2013, Fawcett and Carruthers, 2008], the instruments coordinating recognition of awards and judgments are separate from those coordinating substantive law [United Nations, 1958, Hague Conference on Private International Law, 2019], the recognition of the trust across a border is governed by a convention with a short ratification list [Hague Conference on Private International Law, 1985], the succession convention **never entered into force**, which is why reserved-portion questions remain conflict-of-laws problems [Hague Conference on Private International Law, 1989], and one regional instrument supplies habitual residence plus choice of law for the states bound by it [European Union, 2012]. The comparative treatment of the trust against the civil-law foundation is its own literature [Panico, 2014]; the irreducible function of the entity in all of it is asset partitioning [Hansmann and Kraakman, 2000], and the sharpest account of how legal coding creates capital is also its sharpest critique [Pistor, 2019].

Characterisation is the step that decides the answer before any law is applied, and it is asset-specific rather than structure-specific. That is why the specification below assigns a family to each asset rather than choosing a family for the structure.

6.2 The demotion, stated as a result

The concentration instrument applies to any distribution, so it applies to legal families. **The reading inverts, and that inversion is the finding.** On the jurisdictional axis a low index is good, because it means many independent points of failure. **Here a low index is a cost, because each effective family is an interface**, and every interface requires three written opinions — recognition, characterisation, enforcement — from counsel qualified in the receiving system.

Two warnings belong with the measurement and both are printed rather than deferred.

The index does not see interfaces that the share table hides. A structure using a free-zone jurisdiction whose zone applies common-law-derived rules sits inside a federal civil-law state, and the zone-to-onshore interface is real and invisible to a family index computed off jurisdiction shares. **A law-family index therefore under-counts the interfaces of any structure using a carve-out jurisdiction, which is to say most of the structures in this series.**

A legal family is not a coordination perimeter. Two common-law jurisdictions inside one perimeter fail together; a common-law and a civil-law jurisdiction inside one perimeter also fail together. **The families are orthogonal to the perimeters, and reporting family diversity as though it were perimeter diversity is the same category error the corpus identifies elsewhere.** Curzi [2026] makes the parallel point for liability compartments — six unequal compartments give 4.96 effective compartments on the contagion axis and contribute exactly zero

to the perimeter count — and Curzi [2026] makes it for settlement, where a naive correspondent index of 0.302400 reading 3.306878 effective correspondents collapses to 0.745256 and 1.341819 effective clearing perimeters once flows are decomposed by currency. **Three axes, three share bases, one recurring error.**

6.3 Registered intellectual property, where fragmentation is irreducible

A patent or a mark is not one thing recognised in many places. It is a bundle of separate national rights, each created by a national registry under national law, each with its own renewal calendar, validity standard and enforcement forum; the international instruments coordinate *filing* and create no single global right. **What is mobile is not the right but the holding entity and the licence contract.** One entity can own forty registrations and one licence agreement under one chosen law referred to one forum can cover all forty. The registrations stay exactly where they are, and **any adviser who implies otherwise is selling something.**

7. Results

Every figure is an exact computation on constructed inputs. **No dataset was collected and no figure here is a measurement taken from any real facility, market or structure.**

7.1 R1 — the trigger

$$d^{*} = 1 - \frac{\mathrm{LTV}_{0}}{1 - m}.$$

At $\mathrm{LTV}_{0} = 50\%$ and $m = 35\%$: $V^{*} = 40/0.65 = 61.538462$ and $d^{*} = 1 - 61.538462/80 = 23.076923\%$. **The first margin call falls at a 23.076923 per cent drawdown, not at 50 per cent.** The full surface, in percentage falls in collateral value, with dashes where the facility is in breach at inception:

LTV0	m=20%	m=25%	m=30%	m=35%	m=40%	m=45%	m=50%
25%	68.75	66.67	64.29	61.54	58.33	54.55	50.00
30%	62.50	60.00	57.14	53.85	50.00	45.45	40.00
35%	56.25	53.33	50.00	46.15	41.67	36.36	30.00
40%	50.00	46.67	42.86	38.46	33.33	27.27	20.00
45%	43.75	40.00	35.71	30.77	25.00	18.18	10.00
50%	37.50	33.33	28.57	23.08	16.67	9.09	--
55%	31.25	26.67	21.43	15.38	8.33	--	--
60%	25.00	20.00	14.29	7.69	--	--	--
65%	18.75	13.33	7.14	--	--	--	--

Three readings, in descending order of how often they are missed. **The cushion is roughly linear in the loan-to-value and the slope is brutal**: at $m = 35\%$ every five points of advance costs exactly 7.692308 points of cushion, so going from 40 to 50 per cent raises the cash by 8 units — 15.384615 per cent more deployable capital — and costs 15.384615 points of cushion. **The last four units of borrowing are the most expensive four units in the structure, and they are the ones agreed last, quickly, at the end of a negotiation about something else. The maintenance level matters as much as the advance rate and is negotiated a tenth as hard**: moving from 35 to 30 per cent at a 50 per cent advance takes the trigger from 23.076923 to 28.571429 per cent, five and a half points of free cushion for a term most borrowers treat as boilerplate. **And there is a region of the table where the facility is in breach on the day it is signed** — nobody signs that deliberately, but a maintenance level that steps up on a ratings trigger, an index deletion or a liquidity test moves the borrower rightward along a row nobody re-reads.

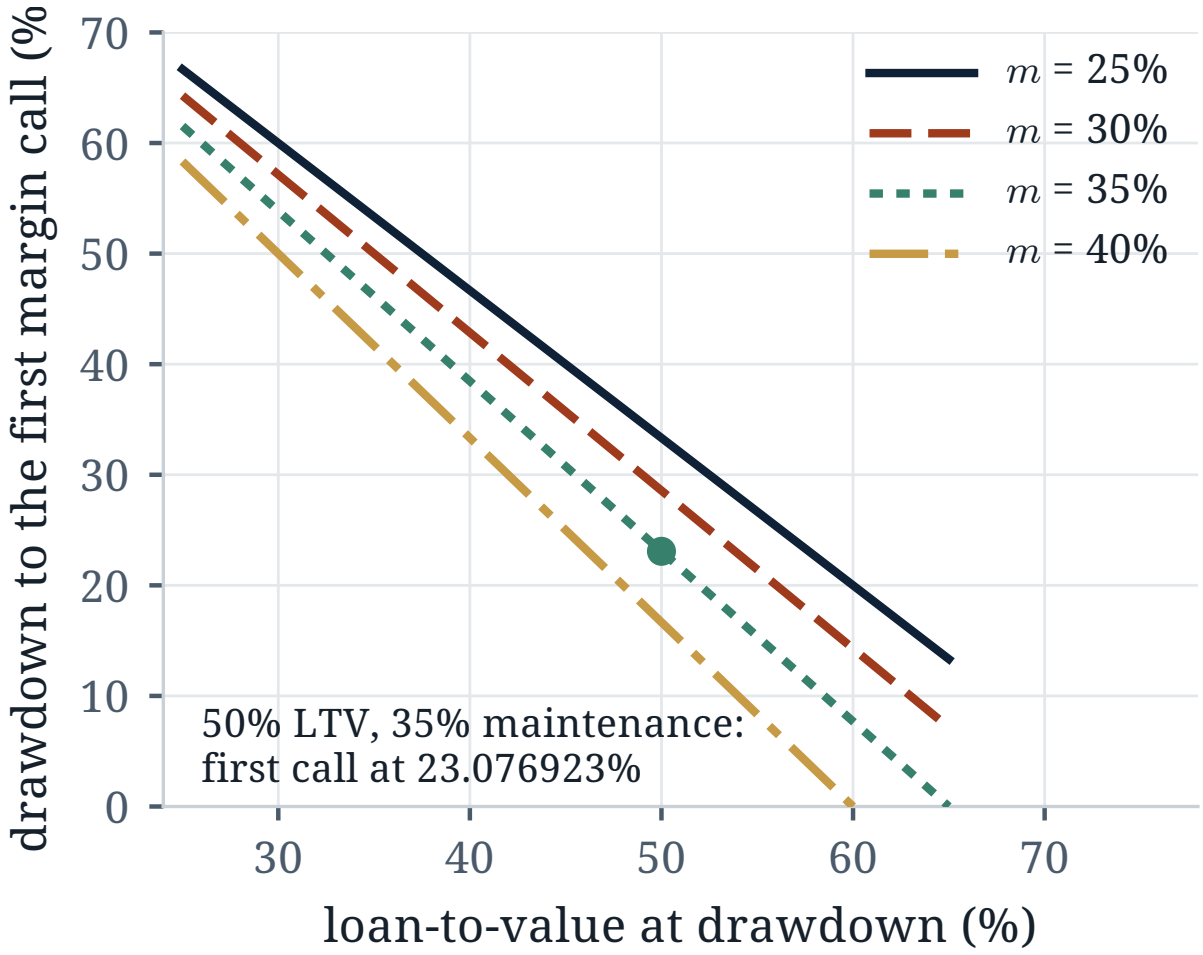


Figure 62.1: **The cushion is linear in the loan-to-value and the slope is steep: at a 35 per cent maintenance requirement, every five points of additional advance costs exactly 7.692308 points of drawdown cushion.** Curves are $d^{*} = 1 - \mathrm{LTV}_{0}/(1 - m)$, the closed form of Section 3.1, plotted for the four maintenance levels in the trigger table of Section 7.1. Script `figs/src/CCP-062_fig01.py`.

Measured on the business rather than on the share price, with senior debt of 35 ahead of equity of 100 and an equity beta to enterprise value of 1.35, the trigger is a **17.094017 per cent** enterprise-value decline and the control loss a **26.666667 per cent** decline. H1 is supported, and it is supported as an identity.

7.2 R2 — the forced sale and the crossover

With $x = V - (V - L)/m$:

```
at a 36.000000% drawdown, price 0.640000
  collateral value 51.200000, loan 40.000000, equity ratio 21.875000%
  proceeds required          x = 51.200000 - 11.200000/0.35 = 19.200000
  shares sold at 0.640000                                   = 30.000000
  shares remaining of the original 80                       = 50.000000
  RESULTING STAKE                                           = 50.000000%
```

drawdown	price	stake after the forced sale
23.076923%	0.769231	80.000000% (first call)
30%	0.700000	65.306122%
36.000000%	0.640000	**50.000000%**
40%	0.600000	38.095238%
50%	0.500000	collateral equity exactly zero

At a 36 per cent drawdown the forced sale takes the holder to exactly 50 per cent, and one basis point beyond it he is a minority holder of the company he founded. The crossover against the unlevered route that raises the same cash by floating 60 per cent and retaining 40 is the honest summary of the trade:

```
route A stake after the forced sale equals 40.000000% at
price 0.606061 per share   ->   drawdown 39.393939%
```

Below a 39.393939 per cent drawdown the levered route leaves the holder with more of the company; above it, the plain float does. That single number is computable at signing. H2 is supported.

One more fact completes the picture and it is about the market rather than the borrower. The public float created by the carve-out is 20 shares; the forced sale at a 36 per cent drawdown is 30 shares. **The liquidation is one and a half times the entire free float, and it must clear into a market that has just marked the company down by more than a third.**

7.3 R3 — the first-passage probabilities

CONSTRUCTED MODEL. Illustrative only. At $\mu = 6\%$, the probability of touching each barrier within the horizon:

Probability of touching the 23.076923 per cent barrier:

σ	1 yr	2 yr	3 yr	5 yr	10 yr
20%	14.405%	26.663%	33.635%	41.364%	49.431%
30%	36.526%	51.288%	58.665%	66.486%	74.684%
40%	52.869%	66.372%	72.772%	79.390%	86.192%
50%	64.031%	75.756%	81.137%	86.581%	92.001%

σ	1 yr	2 yr	3 yr	5 yr	10 yr

Probability of touching the 36 per cent barrier, at which control is lost:

σ	1 yr	2 yr	3 yr	5 yr	10 yr
20%	1.616%	7.139%	12.196%	19.337%	28.328%
30%	12.694%	27.148%	36.178%	46.844%	59.006%
40%	27.952%	45.432%	54.847%	65.194%	**76.347%**
50%	41.614%	58.908%	67.529%	76.598%	85.905%

Median time to touch the 23.076923 per cent barrier: **1.8731 years** at $\sigma = 30\%$, **0.8786 years** at 40 per cent, **0.5225 years** at 50 per cent.

A newly listed, concentrated, single-name equity is not a 20-per-cent-volatility asset. At 40 per cent the model puts the median time to a first margin call at under eleven months and the ten-year probability of a control-losing drawdown at 76.347 per cent. **Whatever the correct empirical figures are — and they should be estimated from the actual security rather than assumed — the qualitative conclusion is not sensitive to the input. The event that destroys this structure is not a tail event. It is the base case over any horizon long enough to matter, and the structure is usually built to last a generation.** H3 is supported by the model and **is not supported by any data, which Section 10 states.**

7.4 R5 — the two competition identities

The change identity. $\Delta = 2 s_1 s_2$, so the hundred-point screen binds exactly when $s_1 s_2 > 50$. On a constructed market in which the group's operating company holds 14.0 per cent — rivals at 22.0, 18.0, 11.0, 8.0, 7.0, 6.0, 5.0 and 4.0 with five fringe firms at 1.0 — the index is 1,320.00 points and the effective firm count is 7.575758:

target share	Δ	post-index	combined share	level $> 1{,}800$	change > 100	presumption
4.0	112.00	1,432.00	18.0%	no	yes	not engaged
5.0	140.00	1,460.00	19.0%	no	yes	not engaged
6.0	168.00	1,488.00	20.0%	no	yes	not engaged
7.0	196.00	1,516.00	21.0%	no	yes	not engaged
8.0	224.00	1,544.00	22.0%	no	yes	not engaged
11.0	308.00	1,628.00	25.0%	no	yes	not engaged
18.0	504.00	1,824.00	32.0%	**yes**	yes	**engaged**
22.0	616.00	1,936.00	36.0%	**yes**	yes	**engaged**

The largest target that keeps the change at or below 100 points is $100/(2 \times 14) = 3.571429$ **per cent of the market — which is to say that once a firm holds 14 per cent, essentially every acquisition worth doing produces a change above 100.** And the roll-up result, which is counterintuitive and matters to anyone running an acquisition programme rather than a single deal:

```
five 1.0% firms acquired SEQUENTIALLY   cumulative delta 160.0000   post 1,480.0000
the same 5.0% acquired in ONE step      delta            140.0000   post 1,460.0000
difference                                               +20.0000 points
```

A roll-up executed in instalments produces a larger total change than the same aggregate bought at once, because the acquirer's own accumulated share raises the delta on every subsequent step. The sequential path is often chosen in the belief that small steps attract less attention; on the arithmetic it does the opposite.

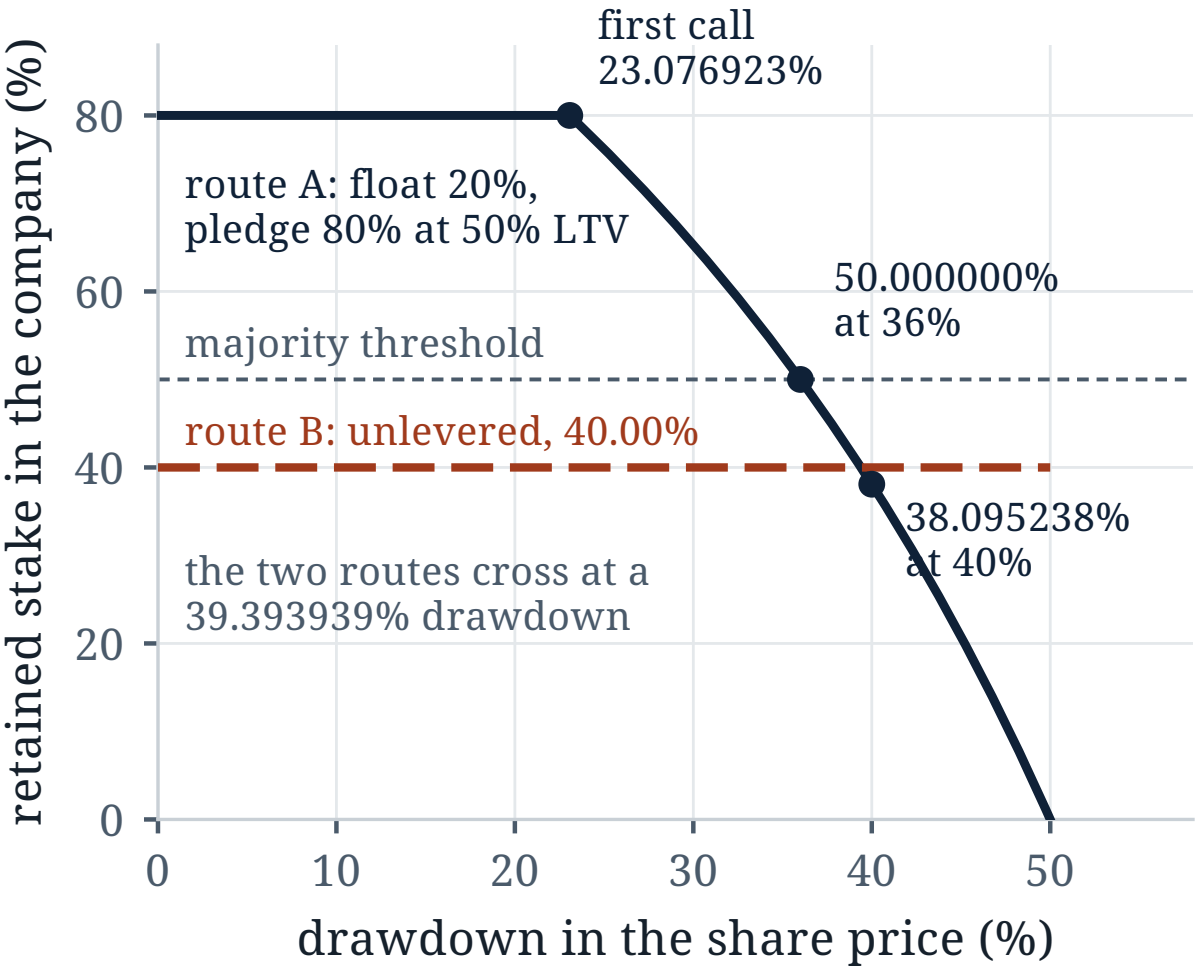


Figure 62.2: **The pledge retains more of the company than an unlevered float until a 39.393939 per cent drawdown and less of it thereafter, and passes through exactly 50 per cent at 36 per cent.** The curve is the retained stake after a forced sale restoring a 35 per cent equity ratio, $x = V - (V - L)/m$; the four marked and labelled values are the ones printed in Section 7.2 and the remainder of the curve is the same formula evaluated between them. Script `figs/src/CCP-062_fig02.py`.

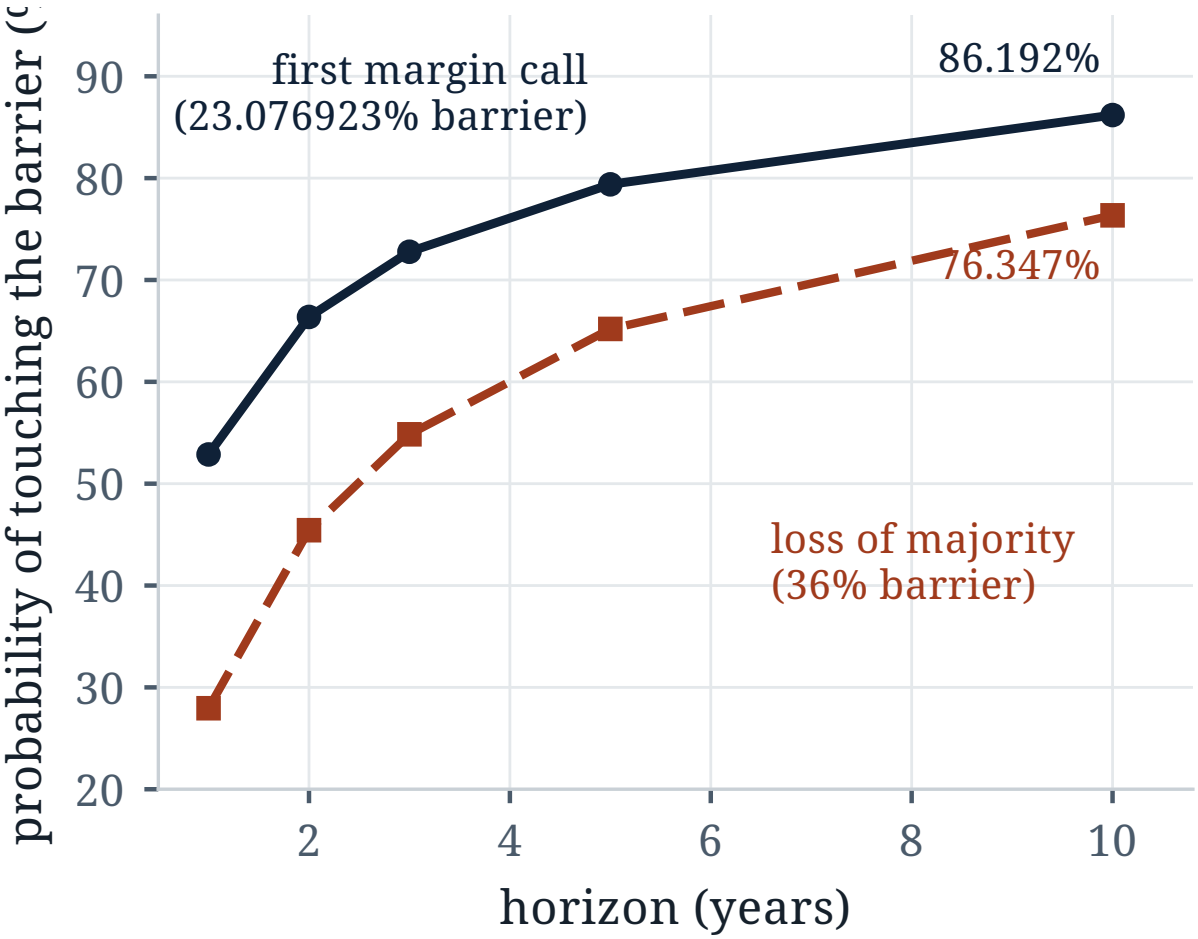


Figure 62.3: **Touching the control-losing barrier is not a tail event over a structural horizon: at 40 per cent volatility the ten-year probability is 76.347 per cent.** First-passage probabilities under the constructed geometric-Brownian-motion model of Section 3.3 at $\sigma = 40\%$ and $\mu = 6\%$; all ten plotted values are printed in the tables of Section 7.3. Script `figs/src/CCP-062_fig03.py`.

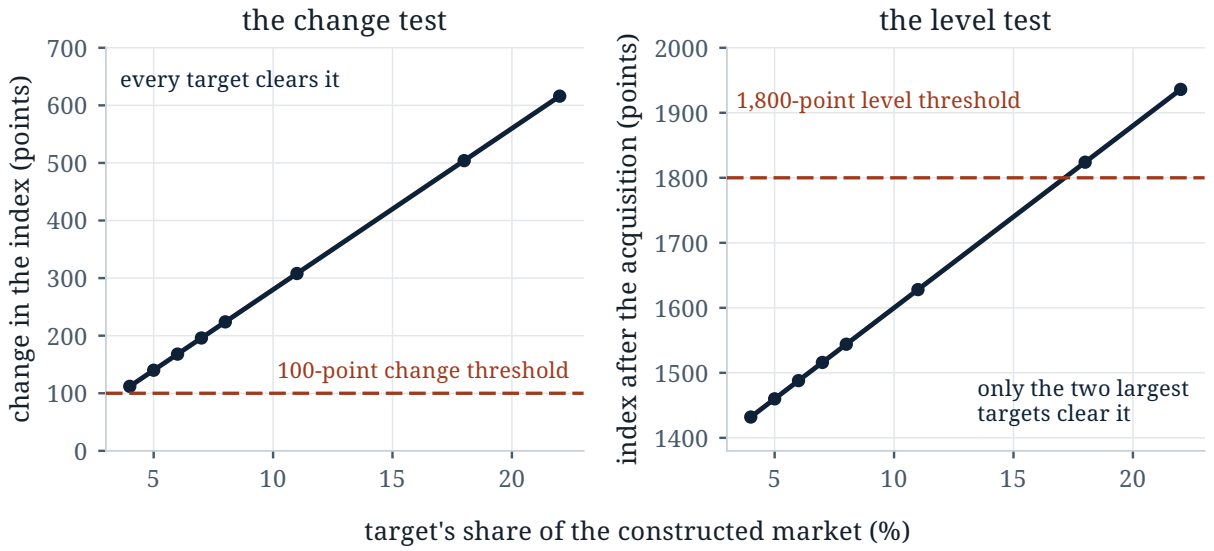


Figure 62.4: **Once the acquirer holds 14 per cent, every target worth buying clears the change threshold and only the two largest clear the level threshold: the change test is what binds a mid-sized acquirer.** Values are the acquisition table of Section 7.4, with the screen thresholds of U.S. Department of Justice and Federal Trade Commission [2023]. Script `figs/src/CCP-062_fig05.py`.

The de minimis asymmetry. With the exemptions turning on competitive sales as a fraction of each corporation's total sales — either below 2 per cent, or each below 4 per cent:

corp A	corp B	2% test (either $<$ 2%)	4% test (both $<$ 4%)	exempt?
3.90%	3.90%	no	yes	**yes**
1.90%	30.00%	yes	no	**yes**
2.10%	3.90%	no	yes	**yes**
1.99%	99.00%	yes	no	**yes**
4.10%	4.10%	no	no	**no**

A pair at 3.90 and 3.90 per cent is exempt. A pair at 1.99 and 99.00 per cent is exempt. A pair at 4.10 and 4.10 per cent is not. The correct first step on any proposed seat is therefore to compute competitive sales as a percentage of each corporation's total sales, rather than to ask whether the companies "compete". Worked on constructed figures, a pair at 1.833 and 4.400 per cent is exempt even though one side exceeds the four per cent test, because the other is below two: **a large overlap for the smaller company can be immaterial for the larger one, and the statute takes that seriously.** H4 is supported as an identity.

7.5 The propagation results

Cross-default. Nine financing and material commercial agreements have $\binom{9}{2} = 36$ pairwise links. Under independent per-agreement annual breach hazards:

q	$n = 3$	$n = 6$	$n = 9$
0.02	5.8808%	11.4158%	16.6252%
0.05	14.2625%	26.4908%	**36.9751%**

q	$n = 3$	$n = 6$	$n = 9$
0.10	27.1000%	46.8559%	61.2580%

At a five per cent per-agreement annual hazard across nine agreements, the chance that at least one trips in a given year is 36.9751 per cent, and with universal cross-default that is the chance the whole group defaults. The individual hazards are small and unremarkable; the aggregate is not. The margin facility is what makes this urgent, because a margin call met late is a default that can propagate into facilities never exposed to the share price at all.

R6 — interfaces. Ten legal touchpoints, of which six are fixed by facts the structure does not control — the location of land, the location of labour, and the residence of three people:

```
mesh across all ten touchpoints        45 interfaces   135 opinions
star through a single apex              9 interfaces    27 opinions
reduction                                              80.0000%

collapsed to the four canonical families
mesh                                    6 interfaces    18 opinions
star through the apex                   3 interfaces     9 opinions
```

Priced at a constructed allocation of 0.014 of the maintenance term C to cross-family opinions, the star costs 0.014000 of C against the mesh's 0.028000 — **a saving of 0.014000 of C, obtained by changing the topology and nothing else.** And the succession forum count falls the same way:

assets held	fora
each directly	6
all movables through the apex	3
after converting one immovable into a security	2

Six fora become three by interposing one legal person, and three become two by converting one immovable into a security. The remaining two are the apex and the irreducible situs of land. **A structure that has honestly reached two succession fora on this asset base has finished; one still looking for a device to reach one is looking for something that does not exist.**

7.6 The law-family measurement, reported as a cost

```
conventional structure
  common law 86.40%   civil law 13.60%
  HHI 0.764992   effective families 1.307203

designed structure
  common law 70.00%   civil law 18.20%   mixed 11.80%
  HHI 0.537048   effective families 1.862031
```

The designed structure's move from 1.307203 to 1.862031 effective families is half an additional family's worth of interface work, bought as a by-product of a diversification programme undertaken for a different reason. It is reported here as a cost line with the caveat in print: the table reads one common-law family for a free-zone exposure whose structure in fact crosses a zone-to-onshore boundary the table cannot see. **The measured 1.862031 understates the interface count, and the interface count is the number that matters on this axis.** The honest statement is three canonical interfaces, nine opinions, **one of them invisible to the metric and material to the outcome.**

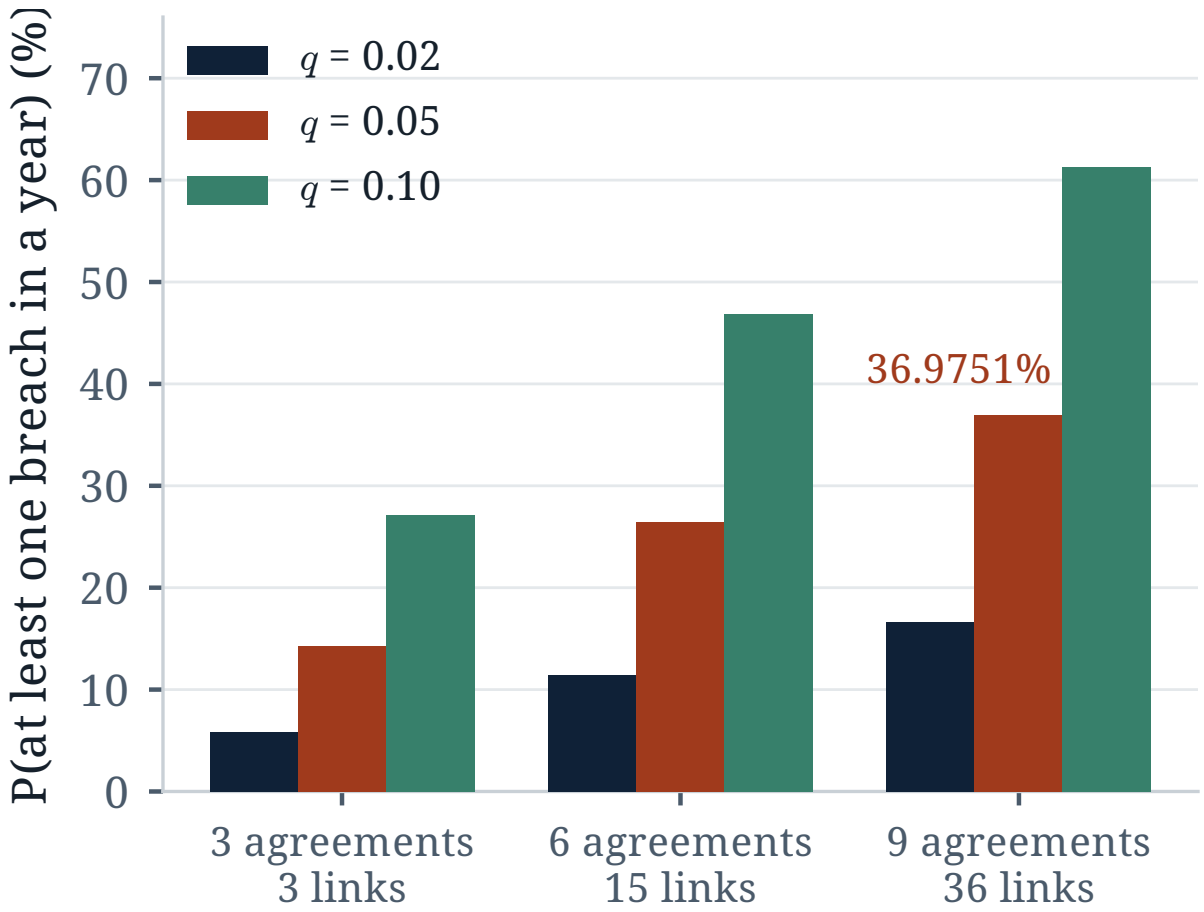


Figure 62.5: **Nine unremarkable agreements at a five per cent annual hazard produce a 36.9751 per cent chance that at least one trips, and universal cross-default makes that the group's default probability.** $P = 1-(1-q)^{n}$ for the three hazards and three portfolio sizes tabulated in Section 7.5; link counts are $\binom{n}{2}$. Script `figs/src/CCP-062_fig04.py`.

7.7 What the pledge does to the concentration measurement

The pledged block sits inside one line of the group's exposure map. A forced sale moves value out of the holder's control and into the lender's perimeter. Modelling an 8.0-point migration from the bridging bloc to the Western bloc on the designed structure of Curzi [2026]:

```
before   naive HHI 0.189002   eff jur 5.290949   bloc-HHI 0.340000
         N 2.941176   h -0.122417
after    naive HHI 0.181482   eff jur 5.510188   bloc-HHI 0.336800
         N 2.969121   h -0.119107
change in h from the margin call:  +0.003310
```

The concentration instrument says the margin call improved the structure. It says so because the bridging bloc was the largest at 40.0 per cent and the forced migration made the bloc shares more even. **The measurement is correct and the conclusion is absurd, and it is the sharpest available illustration of the rule that the index is a diagnostic and not an objective function** [Goodhart, 1975, Campbell, 1979]. A structure that lost majority control of its operating company scored better on the resilience metric.

Report both. A number that improves while the thing it measures gets worse is the most dangerous output any instrument can produce, and the only defence is to read the metric alongside the event rather than instead of it.

7.8 The alternatives, priced — and two source figures corrected

route	cash	stake	trigger drawdown	drawdown to fall to 50%
A float 20%, pledge 80% at 50% LTV	60.00	80.00%	23.076923%	36.000000%
A′ float 20%, pledge 80% at 40% LTV	52.00	80.00%	38.461538%	**48.800000%**
A″ float 20%, pledge 50 of 80 at 50% LTV	45.00	80.00%	23.076923%	see note
B float 60%, no borrowing	60.00	40.00%	none	none
C float 35%, pledge 65 at 35% LTV	57.75	65.00%	46.153846%	**52.105263%**

Note on A″: the trigger drawdown on the pledged 50 is unchanged, because it is a function of the advance rate and the maintenance level only. What changes is cure capacity — 30 unencumbered shares are available to post as additional collateral, which converts the first call from a forced sale into a top-up and preserves the stake entirely at that drawdown.

Two figures in the source manuscript's version of this table do not reproduce and are corrected here rather than repeated. The source prints a control-loss drawdown of 42.105263 per cent for route A′ and 51.417004 per cent for route C. Applying the sale rule $x = V - (V - L)/m$ that the same source states, and that reproduces route A's 36.000000 per cent exactly, route A′ gives $62.5p = 32$, so $p = 0.512000$ and the drawdown is **48.800000 per cent**; route C gives $47.5p = 22.75$, so $p = 0.478947$ and the drawdown is **52.105263 per cent**. **Neither source figure is reproducible from any rule stated in the chapter, and this paper prints the recomputed values**

and flags the discrepancy rather than carrying it forward. The correction strengthens rather than weakens the source's conclusion: **route C survives a market event that ends route A by an even wider margin than the source claimed.**

Route C is what the arithmetic recommends and is what nobody proposes, because it looks like a worse headline — less control retained, less cash raised. **What it buys is a 46.153846 per cent first call and a 52.105263 per cent control-loss threshold against a base case of 23.076923 and 36.000000 per cent, for 2.25 units of forgone cash and fifteen percentage points of a stake that route A holds only nominally.** Negotiating the maintenance level from 35 to 30 per cent at a 35 per cent advance moves the first call further, to exactly 50.000000 per cent.

The residual is not removed. Under the constructed first-passage model at 40 per cent volatility, a 50 per cent drawdown still carries a **63.380 per cent** probability over ten years. Route C moves the threshold; it does not abolish it.

7.9 Counter-result

Legal-family diversity is demoted to a cost line rather than a resilience line. The measured 1.862031 effective families is a count of interfaces to be paid for, not of independent points of failure, and **it under-counts the interfaces of every carve-out jurisdiction in this series, including the free-zone-to-onshore interface the metric cannot see.** The honest verdict is that law-family arbitrage is a much smaller field than it is sold as and much more valuable inside its actual boundaries: **four mobile decisions and one architectural choice — the star rather than the mesh — account for nearly all of the available gain**, and everything else is the world telling the structure where its assets already live.

8. Discussion and failure catalogue

8.1 The failure classes, named

Curzi [2026ac] distinguishes four ways a claimed mechanism fails: no channel at all; a real channel with insufficient gain; a real result whose magnitude has been wrongly extrapolated; and a system unresponsive at the input size offered. **Almost every failure in this paper is class C — a real result at the wrong magnitude.** The pledge genuinely raises cash and genuinely retains control; the loan-to-value genuinely measures something; the legal-family count genuinely counts families. **In each case the quantity is real and the magnitude attributed to it is not, and the correction is arithmetic rather than judgement.**

8.2 Fourteen failure modes

1. **Reading the loan-to-value as the cushion.** The single most consequential error, and the reason for Section 7.1. The advance rate and the drawdown to the first call are different quantities on different scales, and reading one as the other is a scale confusion in the sense of Stevens [1946]. **That borrowers systematically make it is this paper's own conjecture and is uncited.**
2. **Pledging the whole block.** Unencumbered shares are cure capacity, and route A″ shows what they are worth.
3. **Assuming the float can absorb the liquidation.** The forced sale is one and a half times the free float.
4. **Forgetting that the trigger is tested on a levered claim.** Enterprise-value drawdowns of 17.094017 and 26.666667 per cent, not 23.076923 and 36.000000.
5. **Universal cross-default.** Thirty-six links and a 36.9751 per cent annual group default probability.
6. **An unmapped change-of-control surface.** The consequence is frequently larger than the loan.

7. **Uncapped right of use.** Above 100 per cent the whole block goes.
8. **Believing the relationship is a term.** A maintenance clause is a mechanical test and the risk officer has no discretion over it.
9. **Deploying the proceeds into the same risk.** Leverage against an asset, invested in the same asset class, is the same position twice.
10. **Modelling a single drawdown rather than a path.** Section 7.3 exists because the static table is the wrong object.
11. **Confusing the offering type.** Primary and secondary produce identical percentages and opposite cash positions.
12. **Treating disclosure as the risk.** The pledge is disclosed by rule; the risk is the arithmetic nobody ran.
13. **Assuming the lender wants the shares.** It wants its money back, quickly, and its covenant package is written to get it.
14. **Running the structure without an annual re-test.** Every parameter moves and the map decays on events that produce no filing anywhere.

8.3 Where the analysis is exposed

The instrument is exposed to being adopted as a target [Goodhart, 1975, Campbell, 1979], and Section 7.7 is the demonstration: a metric improved while the thing it measured got strictly worse. The only defence offered is the reporting standard in Appendix A, which requires the inputs and the event to be published beside the output.

It is also performative and its own parameters are not exempt. MacKenzie and Millo [2003] documents a model that made its assumptions true after adoption, and Lucas [1976] states the consequence for parameters estimated before publication. **The framework asserts that systems change their behaviour when modelled and then exempts its own volatility, hazard and perimeter assumptions from that reflexivity. That is an inconsistency and it is not resolved here.**

Ensemble and time averages are different objects and every probability here is the first. Peters [2019] establishes that the time average of a multiplicative process differs from its ensemble average. **A margin call is not a deduction from a stream; it is an event that permanently rescales the holder's position in the asset, and the ten-year probability in Section 7.3 is exactly the time-average statement the expected-value framing would miss.** Peters and Gell-Mann [2016] supplies the associated strengthening for the framework's logarithm: it follows from multiplicative dynamics with no psychological assumption at all.

The corpus's practice of cross-checking an identity to machine precision before relying on it is followed here. Curzi [2026] verifies its deficit identity at a residual of 8.88×10^{-16} on its carried case, and Section 9.9 reports the equivalent checks for every identity in this paper — including the two that failed, which are printed in Section 7.8.

One boundary belongs here rather than in a footnote. Curzi [2026] states the doctrine of category violation: category innocence is refused as a defence of acquisition, and moving a harm through a discovered gap in a scheme is forbidden. **The de minimis arithmetic of Section 7.4 is a compliance test and not a design target.** A group that computes competitive sales in order to know whether a seat is lawful is using the statute as intended; a group that arranges its sales in order to sit just under a statutory percentage while doing the thing the statute exists to prevent has found the gap the doctrine forbids. **The break-even method the corpus uses elsewhere [Curzi, 2026] prices a decision honestly precisely because it states what belief the decision requires; the same method used to price the probability of not being caught is a different exercise, and this paper does not perform it.**

9. Verification and rigour audit

9.1 Dimensional audit

symbol	quantity	dimension	unit	status
V_0, V	collateral value at inception, and current	$[M]$	units of value	defined (constructed)
L	loan	$[M]$	units of value	defined (constructed)
LTV_0, m	opening advance rate, maintenance requirement	$[1]$	dimensionless fractions	assumed (constructed)
d^*	trigger drawdown	$[1]$	dimensionless fraction	derived
x	proceeds required to cure	$[M]$	units of value	derived
p	price relative to inception	$[1]$	dimensionless	derived
μ	drift	$[T^{-1}]$	per year	assumed (constructed)
σ	volatility	$[T^{-1/2}]$	per square-root year	assumed (constructed)
$b = \ln(B/S_0)$	log barrier	$[1]$	dimensionless	derived
q	per-agreement annual breach hazard	$[T^{-1}]$	per year	assumed (constructed)
n	agreements, touchpoints, hops	$[\#]$	a count	chosen
s_i	share on a stated basis	$[1]$	dimensionless	defined; sums to 1
HHI, Δ	index, change in index	$[1]$	dimensionless, or points	derived
N	effective count	$[\#]$	a count	derived as $1/\mathrm{HHI}$
C	maintenance cost flow	$[C\,T^{-1}]$	index units per period	assumed
h	surplus balance	$[C\,T^{-1}]$	index units per period	derived

Do the equations balance? $d^* = 1 - \mathrm{LTV}_0/(1-m)$ is a dimensionless combination of dimensionless quantities, so quoting it as a percentage is correct. $x = V - (V-L)/m$ carries $[M]$ throughout, since m is dimensionless. **In the first-passage expression the arguments of Φ are dimensionless only because σ carries $[T^{-1/2}]$ and $\nu T/(\sigma\sqrt{T})$ therefore reduces to a pure number — which is the check that catches a volatility quoted per year being used with a horizon in months**, the commonest dimensional error in this family of models. The exponent $2\nu b/\sigma^2$ is likewise dimensionless.

Where no operational unit exists the audit says so. The value figures are constructed units with no currency attached, and C and h are index units fixed by parameters stated elsewhere in the corpus, with no conversion to currency supplied anywhere. **The percentages, the proba-**

bilities and the effective counts survive this; the statement that the interface saving is 0.014000 of C means nothing until an index unit is defined, and none is.

9.2 Significant figures against epistemic status

- 23.076923%, 36.000000%, 38.095238%, 39.393939%, 7.692308, 3.571429%, 7.0711%, 48.800000% and 52.105263% are **exact arithmetic on stipulated inputs** — rational functions of exact fractions. They are derived-from-defined and may carry as many digits as are useful. **Their inputs are constructed and carry no empirical authority whatever.**
- $\Delta = 2s_1s_2$ is an **identity**, exact for all shares, and the table built from it is exact.
- 1.307203 and 1.862031 are exact reciprocals of exactly stated share vectors, computed from the printed percentages rather than from any exact underlying fraction — **the convention adopted by Curzi [2026] and recorded there as a fifth-decimal discrepancy against the exact-fraction convention of Curzi [2026].**
- 76.347%, 86.192%, 63.380% and 0.8786 years are **illustrative model output from stipulated parameters.** They are exact given μ, σ and the barrier, and μ and σ are assumed. **Three decimal places on a probability produced by an assumed volatility is defensible only as a reproduction aid, and Section 10 says so.**
- 36.9751% is exact arithmetic on a stipulated hazard.
- **No quantity in this paper is a fitted or measured estimate, so no uncertainty interval is reported anywhere. That is not precision; it is the absence of empirical content.**

9.3 Limits and boundary behaviour

limit	behaviour	reading
$m \to 0$	$d^* \to 1 - \mathrm{LTV}_0$	with no maintenance requirement the call arrives only when the loan equals the collateral
$m \to 1$	$d^* \to -\infty$	a maintenance requirement approaching the whole value is unsatisfiable at any advance
$\mathrm{LTV}_0 \to 0$	$d^* \to 1$	an unlevered block is never called
$\mathrm{LTV}_0 \to 1 - m$	$d^* \to 0$	the facility is at the trigger on the day it is signed — the dashes in the table
$V \to L/(1-m)$	$x \to 0$	the cure requirement vanishes at the trigger, as it must
$V \to L$	$x \to V$	at zero equity the whole block is consumed
$\sigma \to 0$	$P \to 0$ for a barrier below $S_0 e^{\mu T}$	a riskless asset never touches
$\sigma \to \infty$	$P \to 1$	certainty of touching, at any barrier
$T \to \infty$, $\nu > 0$	$P \to e^{2\nu b/\sigma^2} < 1$	positive drift leaves a strictly positive chance of never touching
$q \to 0$	$P(\text{breach}) \to 0$	no hazard, no propagation

limit	behaviour	reading
$n \to \infty$ at fixed q	$P \to 1$	universal cross-default across enough agreements defaults with certainty
$s_1 s_2 \to 0$	$\Delta \to 0$	acquiring nothing changes nothing

Two limits are real constraints on the domain rather than results. $d^* < 0$ for $\text{LTV}_0 > 1 - m$, which is not a drawdown but a statement that the facility is already in breach; the model's domain is $\text{LTV}_0 < 1 - m$. And **the infinite-horizon limit of the first-passage probability is the one that matters for a structure built to last a generation**: with $\nu > 0$ the probability of never touching is $1 - e^{2\nu b/\sigma^2}$, which at $\sigma = 40\%$ and $\mu = 6\%$ is small. **A structure whose survival depends on a barrier never being touched over an unbounded horizon is depending on a quantity the model says approaches its bound.**

9.4 Invariance

The trigger and the control-loss threshold are invariant under any rescaling of the company's value. Doubling V_0 and L together leaves LTV_0 and therefore d^* unchanged, and the stake percentages after a forced sale are likewise homogeneous of degree zero in value. **That is a genuine scale invariance and it is why the arithmetic applies to a block of any size.** It is **not** invariant under changing the maintenance level or the share of the company floated, which is the entire content of Sections 7.1 and 7.8.

The delta identity is invariant under relabelling and is exactly quadratic, so it does not survive a rescaling of shares that fails to preserve the simplex — a point that matters whenever a market definition is redrawn, because redefining the market rescales every share and changes Δ quadratically.

What would break the share invariance: any basis that is not normalised. **Tested at how many scales?** The invariances are algebraic rather than empirical and were checked at exactly the vectors reported here: two, three, four, nine, ten, twelve and fourteen entries. No claim of empirical scale-freedom is made.

What is conserved, and in what units. Mirowski [1989] shows that the characteristic failure of this genre is borrowing the notation of a conserved quantity without the conservation law that gave it content. Two quantities are genuinely conserved here. The share vector sums to one exactly and dimensionlessly, which is what gives the index its bounds; and **the total number of shares is conserved through the forced sale — the block shrinks by exactly the number sold and the buyer's holding grows by the same number**, which is what makes the control arithmetic exact rather than approximate. C and h, by contrast, are stocks and flows with no stated conservation law and no operational unit, and none is claimed for them. Cartwright [1999] supplies the discipline that makes a partial answer respectable: **the arrangement claimed here is a single pledged block against a single facility with a mechanical test, and nothing outside it.** Kirman [1992] is the objection this paper cannot answer: one C is written for a heterogeneous group of interfaces.

9.5 Analogy discipline

borrowing	source	category	justification
effective count as $1/\text{HHI}$	Hill [1973], Jost [2006]	**1 — derivation**	an identity, not a resemblance

borrowing	source	category	justification
numbers-equivalent reading	Adelman [1969]	**1 — derivation**	elementary, checkable in two lines
the concentration screen	U.S. Department of Justice and Federal Trade Commission [2023], U.S. Department of Justice and Federal Trade Commission [2010]	**1 — derivation**	the guidelines state the index and the thresholds this paper uses
portfolio variance $w' R w$	Markowitz [1952]	**1 — derivation**, via CCP-052	the exposure vector is a weighted sum of shocks with a correlation matrix
similarity-sensitive diversity	Leinster and Cobbold [2012]	**1 — derivation**, stated domain	exact where R has unit diagonal and non-negative entries
ln N as entropy	Shannon [1948]	**2 — shared form**	Shannon entropy in nats; no thermodynamic content imported
geometric Brownian motion	standard diffusion	**2 — shared form**	assumed, not fitted; Sornette [2003] is the reason to expect it to understate
coupled-network fragility	Buldyrev et al. [2010], Albert et al. [2000]	**2 — shared form**	percolation premises are not shown to hold for financing covenants
cascade thresholds	Watts [2002]	**2 — shared form**	the model's degree assumptions are not verified for an agreement portfolio
pairwise combinatorics	Brooks [1975]	**2 — shared form**	$\binom{n}{2}$ is the same object in both settings; no organisational result is imported
channel capacity as a bound	Curzi [2026x], Shannon [1948]	**3 — framing**	used only for the discipline that throughput is receiver-limited; no capacity is computed for a group
tipping from a threshold	Granovetter [1978], Schelling [1971]	**3 — framing**	prior statement of the aggregate phenomenon, cited as ancestry
reflexivity in a falling market	Soros [2013]	**3 — framing**	a qualitative account, never used as a formal result

borrowing	source	category	justification
normal accidents	Perrow [1984]	**3 — framing**	qualitative organisational theory

The failure this table exists to prevent is letting a category-2 or category-3 borrowing carry a category-1 result's authority. The trigger, the sale rule, the delta identity and the combinatorics are category 1 and need no help; the diffusion, the network results and the capacity analogy are category 2 and 3 and are used for shape.

9.6 Measurement-theoretic status

Shares of a company, of a market and of an estate are **ratios of like quantities** and are ratio-scale by construction, so the trigger, the sale rule and the delta identity are untouched by the scale-type objection. **Three inputs are a different matter and the objection is conceded for all three.** The volatility σ, the drift μ and the per-agreement breach hazard q are treated as ratio-scale quantities supporting multiplication, division and exponentiation, and none is estimated from data. Michell [1997] is the strongest available objection — that quantitative structure is assumed rather than tested — and the formal remedy is conjoint measurement [Luce and Tukey, 1964], whose conditions establish whether an attribute has additive quantitative structure at all, within the representational framework of Krantz et al. [1971]. **The conjoint conditions were not tested for any of the three, and every probability computed from them is therefore provisional.** What survives without the test is the *ordering* — a higher volatility gives a higher probability, a larger portfolio of cross-defaulted agreements a higher group hazard — and the algebraic identities, which do not depend on scale type at all.

The 2 and 4 per cent statutory tests of Section 7.4 are exact thresholds on a ratio quantity and are not scores; no ordinal judgement enters that table, and the paper takes no minima or means of anything ordinal. **The corpus's habit of collapsing a structured situation into one score is named once at framework level and is inherited rather than re-argued**; where a single number appears here it is a share, a probability or an exact threshold, and where it is a probability the decomposition is published rather than the scalar defended.

9.7 Postulates

#	statement	status
P1	the maintenance test is $(V - L)/V \geq m$, tested against a published close	DEFINITION
P2	the cure sale restores the ratio exactly: $x = V - (V - L)/m$	DEFINITION (modelling choice: no over-sale, no under-sale)
P3	shares on a stated basis are non-negative and sum to 1	DEFINITION
P4	HHI $= \sum_i s_i^2$	DEFINITION
P5	the price follows geometric Brownian motion with constant μ, σ	POSTULATED (**contradicted by the tail literature, Section 3.3**)
P6	$\mu = 6\%$, $\sigma \in \{20, 30, 40, 50\}\%$	POSTULATED (constructed)

#	statement	status
P7	agreement breaches are independent across agreements	POSTULATED (**false under common cause, Section 3.4**)
P8	$\text{LTV}_0 = 50\%$, $m = 35\%$, senior debt 35 against equity 100	POSTULATED (constructed)
P9	jurisdictions inside one coordination perimeter act on one decision	POSTULATED (stipulated reference group, Section 3.6)
P10	each legal-family interface requires three opinions and costs 0.014 of C in aggregate	POSTULATED (constructed)

9.8 The compounding chain

step	consumes	statement	status
L1	P1	the call arrives at $V^* = L/(1-m)$	ARITHMETIC
R1	L1, P8	$d^* = 1 - \text{LTV}_0/(1-m) = 23.076923\%$	ARITHMETIC
R1b	R1	at $m = 35\%$, $\partial d^*/\partial \text{LTV}_0 = -1/0.65$, so five points of advance cost 7.692308 points	ARITHMETIC
L2	P2	shares sold $= x/p$ with $x = V - (V - L)/m$	ARITHMETIC
R2a	L2, P8	at $d = 0.36$: $x = 19.200000$, 30 shares sold, stake exactly 50.000000%	ARITHMETIC
R2b	L2, P8	at $d = 0.40$: stake 38.095238%; at $d = 0.50$: equity zero	ARITHMETIC
R4	L2	route A's stake equals 40% at $p = 0.606061$, so the routes cross at 39.393939%	ARITHMETIC
R3	P5, P6	first passage at $\sigma = 40\%$: 76.347% over ten years to the 36% barrier; median 0.8786 yr to the 23.076923% barrier	ARITHMETIC on a POSTULATED model
L3	—	$(s_1 + s_2)^2 - s_1^2 - s_2^2 = 2s_1 s_2$	ARITHMETIC (identity)
R5a	L3	the 100-point screen binds exactly when $s_1 s_2 > 50$; $\sqrt{50} = 7.0711$	ARITHMETIC

step	consumes	statement	status
R5b	L3	at a 14% acquirer the largest target keeping $\Delta \le 100$ is 3.571429%	ARITHMETIC
R5c	L3	sequential roll-up cumulative $\Delta = 160$; one step $= 140$	ARITHMETIC
L4	P7	$P = 1 - (1-q)^n$; links $= \binom{n}{2}$	ARITHMETIC
R5d	L4	at $q = 0.05$, $n = 9$: 36.9751%, over 36 links	ARITHMETIC
R6	L4's combinatorics, P10	mesh 45 interfaces and 135 opinions; star 9 and 27; reduction 80.0000%	ARITHMETIC
CR	P3, P4	law-family HHI 0.764992 and 0.537048 give 1.307203 and 1.862031 effective families	ARITHMETIC

Every ARITHMETIC row can be redone by hand or in three lines of any language. **Only R3 rests on a postulated model, and P5 and P7 are both stated to be false in the direction that makes the results conservative** — the diffusion understates the barrier probability and the independence assumption understates the propagation probability. **The chain therefore establishes exact consequences of stipulated inputs, and it is not a proof of anything about any real facility.**

9.9 Cross-checks performed

- The trigger table was regenerated from the closed form and matched every printed cell.
- The forced-sale rule was re-solved for the stake at each printed drawdown; route A's 36.000000, 38.095238 and 65.306122 per cent all reproduce exactly, as does the crossover at 39.393939 per cent.
- **Two printed figures did not reproduce and are corrected in Section 7.8**: route A''s control-loss drawdown, printed as 42.105263 per cent and recomputed at 48.800000 per cent, and route C's, printed as 51.417004 per cent and recomputed at 52.105263 per cent. **Neither source figure follows from any rule stated in the source, and the corrected values strengthen the source's own recommendation.**
- All twenty-nine first-passage probabilities and all three median times were recomputed from the closed form and reproduce to the printed precision.
- $\Delta = 2s_1s_2$ was verified against the direct difference of sums of squares for every row of the acquisition table.
- The market index was recomputed from its shares: $1{,}320.00$ points, effective count $10{,}000/1{,}320 = 7.575758$; all shares sum to 100.0.
- The law-family indices were recomputed: $0.864^2 + 0.136^2 = 0.764992$ and $0.70^2 + 0.182^2 + 0.118^2 = 0.537048$.
- The perimeter migration of Section 7.7 was recomputed: $0.32^2 + 0.30^2 + 0.38^2 = 0.336800$, $N = 2.969121$, $h = 0.35 \ln N - 0.5 = -0.119107$, change $+0.003310$.
- $\binom{9}{2} = 36$, $\binom{10}{2} = 45$, $\binom{4}{2} = 6$ and $\binom{12}{2} = 66$ were checked directly.

10. Limitations

Every case is constructed. The company, its value, its debt, the offering, the facility, its advance rate and maintenance level, the volatility and drift, the breach hazards, the product market, the touchpoint map and the interface pricing are invented and are labelled as such at each point of use. **No external dataset was collected and none is reported.**

No paper in this corpus reports a first-party measurement, and this one does not either. The corpus's own instrument audit finds that of 27 parameter-by-scope cells, 19 have a defined estimator, 1 has a validated one, and **0 carry a first-party measurement.** That is the central limitation of the whole corpus and it is stated here before a reviewer states it.

The first-passage model is the least defensible thing in the paper. Constant volatility and log-normal returns are contradicted by the empirical literature on equity tails, and Sornette [2003] is the reason to expect the model to understate rather than overstate the barrier probabilities. **No volatility was estimated from any security, no alternative process was fitted, and no out-of-sample test was run.** The right reading of Section 7.3 is a lower bound conditional on a stipulated process, and the qualitative conclusion — that the barrier is reached over a structural horizon — is the only part that should be carried.

The independence assumption in the cross-default arithmetic is false in the direction that matters. Breaches share common causes, and a margin call is itself a common cause, so $1-(1-q)^n$ understates the group hazard. **The figure is a floor and is presented as one.**

The claim that borrowers systematically misread the advance rate as a drawdown cushion is uncited and is this work's own conjecture. The scaffold for this paper named the measurement-theory cluster as the source of support for it; that cluster contains measurement theory and psychometrics and contains no judgment-and-decision literature that would support the claim. **Rather than attach a citation that would not survive a check, the claim is badged as a conjecture here and in Section 8.2**, and the strongest thing that can be said with support is the weaker scale-type point that the advance rate and the drawdown are different quantities [Stevens, 1946].

Jurisdiction-specific legal claims are not cited to primary sources. The descriptions of the interlocking-directorate prohibition and its de minimis percentages, the pre-merger notification regime and its avoidance and penalty rules, the restraint-of-trade prohibition, the treatment of rehypothecated collateral, the situs and characterisation rules, reserved-portion regimes, and every statutory reference in the source's own verification registers are taken from the source manuscript. **The bibliography contains general comparative, treaty and doctrinal material and the published merger guidelines only; it contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction, and no case law.** Those claims are described rather than asserted and are itemised in Appendix B. **This is a real gap and it is not closed here.**

Two figures in the source manuscript are wrong and are corrected rather than repeated (Sections 7.8 and 9.9). A reader comparing this paper against the source will find the difference, and the difference is disclosed rather than left to be discovered.

The 0.014 of C attributed to cross-family opinions is a constructed allocation with no fee schedule behind it, and every conclusion about the value of the star topology inherits that. What survives without it is the **count** — 45 interfaces against 9 — which is combinatorial and exact.

The legal-family index under-counts and the paper says so twice. It cannot see a zone-to-onshore interface, which is present in most of the structures this series examines. **The interface count, not the index, is the right instrument on that axis, and the index is reported only because the corpus reports it elsewhere and consistency requires the comparison.**

The perimeter assignments in Section 7.7 are stipulated, not identified [Manski, 1993], and the remedy adopted is stipulation with a defence [Brock and Durlauf, 2001]. No randomisation, instrument or partial-population variation was available.

Small-sample bias is not addressed [Chao et al., 2014], and effective-number measures

behave poorly where one unit dominates [Golosov, 2010] — which is the regime of the conventional structure's law-family vector at 86.4 per cent in one family.

Model comparison was not performed. No alternative index was fitted, no information criterion was computed and no matched null was constructed, because nothing here is fitted.

No differential prediction is offered. Durlauf [2005] asks for the prediction only this model makes, and the trigger arithmetic makes none that any competent credit analyst could not derive. **The contribution is that the derivation is published beside the transaction it prices, not that it is novel.**

Rigour checks not met, listed by name.

1. **No as-of date, because no real facility was examined.** Any application to a real position requires the facility's actual advance rate, maintenance level and step-up triggers, with a date. Appendix A makes both mandatory.
2. **No uncertainty is reported anywhere, because nothing is estimated.**
3. C **and** h **have no operational unit**, and the interface saving is therefore unitless.
4. **The constant-volatility assumption is untested and is contradicted by the tail literature** [Sornette, 2003, Longin and Solnik, 2001, Ang and Chen, 2002].
5. **The independence assumption in the propagation arithmetic is untested and is false under common cause** [Buldyrev et al., 2010].
6. **The conjoint conditions were not tested for** σ**,** μ **or** q [Michell, 1997, Luce and Tukey, 1964].
7. **No distributional or scaling claim is made, so the Clauset–Shalizi–Newman procedure was not run** [Clauset et al., 2009]. Any restatement reaching for a heavy-tail claim about holdings, share distributions or price returns must run it first, and must read Stumpf and Porter [2012] on why passing it still would not identify a mechanism.
8. **The network, cascade and capacity borrowings are category 2 and 3, not category 1** [Buldyrev et al., 2010, Watts, 2002, Curzi, 2026x].
9. **Aggregation is unaddressed** [Kirman, 1992].
10. **Expected value and time average are not the same thing** [Peters, 2019], and the paper reports probabilities rather than expectations for exactly that reason.
11. **The instrument is performative and its own parameters are not exempt** [MacKenzie and Millo, 2003, Lucas, 1976].
12. **There is no conservation law for** C **or** h [Mirowski, 1989]; the two genuine conservations are the share simplex and the share count through the forced sale.
13. **A metric improved while the thing it measured got worse** (Section 7.7), which is the instrument's own failure mode arriving in print [Goodhart, 1975, Campbell, 1979].

11. Conclusion

This paper prices a carve-out, a pledge and a group's legal interfaces, and the following statements are the fixed text of its results.

R1. The pledge trigger is $d^{*} = 1 - \mathrm{LTV}_0/(1-m)$: at $\mathrm{LTV}_0 = 50\%$ and maintenance $m = 35\%$, **the first margin call falls at a 23.076923 per cent drawdown, not at 50 per cent.** At $m = 35\%$ every five points of additional advance costs exactly 7.692308 points of cushion, and measured on the business rather than the share price the same trigger is a 17.094017 per cent enterprise-value decline.

R2. The forced sale takes an 80 per cent holder to **exactly 50.000000 per cent at a 36.000000 per cent drawdown**, to 38.095238 per cent at 40 per cent, and exhausts the collateral equity at 50 per cent. The liquidation at the 36 per cent point is 30 shares against a free float of 20 — **one and a half times the entire float, into a market that has just marked the company down by more than a third.**

R3. Under a constructed first-passage model at $\sigma = 40\%$ and $\mu = 6\%$, **the ten-year probability of touching the 36 per cent control-loss barrier is 76.347 per cent**, and the median time to

the 23.076923 per cent barrier is **0.8786 years**. The event that destroys this structure is not a tail event; it is the base case over any horizon long enough to matter.

R4. The levered and unlevered routes raising identical cash **cross at a 39.393939 per cent drawdown**: below it the pledge retains more of the company, above it the plain float does.

R5. The change identity $\Delta = 2s_1s_2$ makes the hundred-point screen bind **exactly when** $s_1s_2 > 50$ — 7.0711 per cent each for an equal pair — so that a firm holding 14 per cent can acquire nothing above 3.571429 per cent without clearing it, and a sequential roll-up produces 160.0000 points of cumulative change against 140.0000 for the same aggregate bought in one step. The de minimis provisions of the interlocking-directorate prohibition **exempt a 1.99 / 99.00 per cent pair and catch a 4.10 / 4.10 per cent pair.**

R6. Routing every cross-family interface through one apex **cuts 45 interfaces and 135 opinions to 9 and 27, an 80.0000 per cent reduction**, and takes six succession fora to three and then to two.

Counter-result. Legal-family diversity is demoted to a cost line rather than a resilience line. The measured 1.862031 effective families counts interfaces to be paid for and not independent points of failure, and **it under-counts the interfaces of every carve-out jurisdiction in this series, including the free-zone-to-onshore interface the metric cannot see.** A legal family is not a coordination perimeter, and reporting the one as the other is a category error.

Three further results are established as method rather than as measurement. **Cross-default converts nine agreements into thirty-six pairwise links and a 36.9751 per cent annual probability that one trips at a five per cent per-agreement hazard**, so the engineering response is to cut links rather than to negotiate hazards. **A forced sale can improve the concentration measurement while destroying the thing the measurement exists to protect** — the constructed migration raises h by $+0.003310$ at the moment majority control is lost — which is the sharpest available demonstration that an index is a diagnostic and not an objective function. **And two figures in the source manuscript's alternatives table do not reproduce from its own sale rule and are corrected here**, to 48.800000 and 52.105263 per cent, which strengthens rather than weakens the recommendation the source made.

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every case in this paper — the company and its capital structure, the offering, the facility and its terms, the volatility and drift, the breach hazards, the product market and its shares, the touchpoint map, the interface pricing and the exposure vectors — is **CONSTRUCTED** and is labelled as such at each point of use. The numerical results are exact computations on those constructed inputs, carried at full precision and rounded only for printing.

All figures were produced by the scripts named in their captions (`figs/src/CCP-062_fig01.py` through `figs/src/CCP-062_fig05.py`), which import the shared figure-style module and read their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

Legal and regulatory descriptions in this paper are taken from the source manuscript and not from primary legal sources. The bibliography contains general comparative, treaty and doctrinal material and the published merger-enforcement guidelines only; it contains no statutes, regulations, court rules, registry documents or case reports for any individual jurisdiction, and no case law. Appendix B lists the claims affected. Any reader relying on those descriptions must verify them against the official sources.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. Every person, company, market, facility and transaction described is invented.

Legal, tax and investment advice

This paper is analytical and is not legal, tax, regulatory or investment advice. It describes offering structures, securities-backed lending, collateral arrangements, competition-law constraints and conflict-of-laws interfaces for the purpose of computing their arithmetic consequences. No recommendation is made in respect of any particular person, company, security or facility. Maintenance requirements and their regulatory floors, notification thresholds and waiting periods, de minimis percentages and dollar thresholds, rehypothecation and client-asset rules, situs and characterisation rules and reserved-portion regimes all change and vary by jurisdiction and by activity; **every such statement in this paper is carried from the source manuscript, is flagged in Appendix B, and requires confirmation with qualified counsel in the relevant jurisdiction before any reliance.**

Appendix A. Reporting standards

A.1 The pledge register. Any pledged control block reports the following, and nothing less.

```
the offering type -- primary, secondary or mixed -- and the
cash position each produces
the opening advance rate and the maintenance requirement,
with every step-up trigger and the event that fires it
the trigger drawdown d* = 1 - LTV0/(1-m)
the full stake-after-forced-sale table, at 10-point drawdown
intervals to the point of zero collateral equity
the drawdown at which majority control is lost
the same two figures translated to enterprise value, using
the equity beta implied by the operating company's own debt
the forced sale as a multiple of the free float
the crossover drawdown against the unlevered route raising
identical cash
the right-of-use cap, and whether the claim is proprietary
the change-of-control map: every agreement, its trigger, its
consequence, and whether consent is obtainable in advance
the cross-default register: agreements, links, materiality
thresholds, grace periods, and cross-acceleration where
obtained in place of cross-default
the date, and the difference against the previous re-test
```

Item 10 is the one nobody maintains and the one whose total is frequently larger than the loan. Item 12 decays first: every parameter moves, and the map decays on events that produce no filing anywhere.

A.2 The group screening calendar. Re-run on every acquisition, product launch and geographic expansion, not annually: competitive sales as a percentage of each corporation's total sales for every pair of boards on which a common individual sits; the index and the change for every contemplated transaction; and the changed-circumstances review that catches an interlock made unlawful by an acquisition somebody else negotiated.

A.3 The interface register. Every legal-family touchpoint, its family, whether it is fixed by facts or mobile by choice, and for every interface the three dated opinions — recognition, characterisation, enforcement — each from counsel qualified in the receiving system and each with a stated invalidating event. **The register is a fixed and knowable quantity rather than an open-ended one, which is the practical benefit of counting interfaces rather than reacting to them.**

Appendix B. Claims requiring verification against primary sources

The following statements rest on legal or regulatory detail that changes, and **none is cited to the bibliography, which contains no primary sources for any individual jurisdiction and no case law.** They are reproduced from the source manuscript and require confirmation with counsel before any reliance. The source manuscript's own verification registers run to more than seventy numbered items across the three chapters and are condensed here without loss of subject.

1. **Margin regulation.** The regulatory floors on maintenance margin for a regulated intermediary, any house requirement above them, and their application to a concentrated single-name control block. **The 35 per cent used throughout this paper is a plausible commercial level and is not a quotation.**
2. **Securities-backed lending documentation.** The standard formulation of the maintenance test, the cure window, the lender's sale right and its notice requirements, and whether the test is mechanical or subject to any discretion.
3. **Rehypothecation and client assets.** The right-of-use provisions and their caps, the distinction between a security interest with segregation and an outright transfer, and the insolvency litigation that established the treatment of client assets and client money in the leading market. **Specific holdings, their citations and their scope must be taken from counsel.**
4. **Disclosure of a pledge.** The filing and announcement obligations attaching to a pledge of a control block, and the thresholds and deadlines for beneficial-ownership reporting in each relevant market.
5. **The interlocking-directorate prohibition.** Its statutory location and current text; the dollar threshold on capital, surplus and undivided profits, which is adjusted annually; the 2 and 4 per cent de minimis percentages and the definition of competitive sales; the separate regime applicable to banking institutions; the grace period on changed circumstances and its length; whether the prohibition reaches entity forms other than corporations; the status of the deputisation theory; and the enforcement record since 2022, including the number of resignations and what each announcement said.
6. **The pre-merger notification regime.** The size-of-transaction and size-of-person thresholds, which are adjusted annually; the length of the waiting period and its variants; the second-request mechanics; the per-day civil penalty; the avoidance rule; and the treatment of gun-jumping.
7. **The merger guidelines.** The index thresholds and the change threshold used in Section 7.4 are agency enforcement policy and not statute, and courts are not bound by them

[U.S. Department of Justice and Federal Trade Commission, 2023, 2010]. Their current form must be confirmed.

8. **The restraint-of-trade prohibition** and the line between structure and conduct, including the treatment of information exchange between competitors and the criminal exposure attaching to hard-core horizontal agreements.
9. **Structural subordination and ring-fencing** in the relevant insolvency systems: which claims a group-level pledge actually reaches, and how the residual claim is valued in a distribution.
10. **Change-of-control provisions** in each class of agreement listed in Section 4.2, and in particular whether an investment promotion or a treaty-based ownership right is conditioned on shareholder identity or nationality, and what re-testing is required.
11. **Cross-default and cross-acceleration** drafting: the difference between them, the market position on materiality thresholds and grace periods, and whether a counterparty will accept the narrower clause.
12. **The conflict-of-laws material**: the situs rules for movables and immovables, the place-of-the-relevant-intermediary approach for securities accounts and the convention that embeds it, characterisation doctrine, and the recognition-versus-enforcement distinction. The general doctrine is cited [Torremans et al., 2022, Briggs, 2013, Fawcett and Carruthers, 2008]; **no jurisdiction-specific rule is.**
13. **Succession and reserved portions**: the reserved-share fractions of each civil-law system touched by the structure, the applicable conflicts rule, and whether the reserved portion reaches a given asset at all — which is asset-specific and decided by characterisation. The absence of a succession convention in force [Hague Conference on Private International Law, 1989] and the regional instrument that partly fills the gap [European Union, 2012] are cited; **no domestic provision is.**
14. **Trust and foundation recognition**: the ratification list of the recognition convention [Hague Conference on Private International Law, 1985], which is short, and the comparative treatment of the two devices [Panico, 2014].
15. **Free-zone-to-onshore interfaces**: whether a judgment of a free-zone court is enforceable onshore in the same federal state, on what conditions, and what characterisation the onshore court applies. **This is the interface the law-family index cannot see and it is the single most valuable item in this register.**
16. **Registered intellectual property**: the national registries, renewal calendars, validity standards and enforcement fora for each right, and the scope of the international filing and registration instruments, which coordinate filing and create no unitary right.
17. **All value, cost, volatility, hazard and share figures in this paper are constructed.** They are not quotations, benchmarks or estimates of any real company, facility, market or fee schedule.

All numerical results in this paper were computed from the stated inputs at full precision and rounded only for printing.

Chapter 63

The Assembled Architecture: Effective Blocs, Sanctions Deeming Rules, and the Disclosure Test in Historical Instance

Abstract. We assemble the full structure and measure it with the corpus's own instrument. Against test parameters eta = 0.35, C = 0.30, Q = 10 and delta = 0.02, giving a maintenance floor of 0.500000, the conventional single-perimeter structure has one effective bloc and h = -0.500000; the designed six-jurisdiction structure has bloc HHI 0.340000, N = 2.941176 and h = -0.122417. The architecture closes 75.5167 per cent of the shortfall and does not close it, against a break-even N* = exp(0.5/0.35) = 4.172734. Marginal jurisdictions are priced rather than argued: the break-even probability p* = L/(G + L) puts the AIFC leg at 0.005110/(0.026731 + 0.005110) = 16.0476 per cent with a cost ceiling of 8.9102 per cent of C, while the largest available gain, a full Hong Kong leg at 0.072279, is refused on tail exposure and replaced by capped transit delivering 3.9692 per cent of the benefit for 3.2110 per cent of the exposure. On sanctions, the fifty per cent deeming rule is shown to be non-economic: seven tiers of 51 per cent holdings leave a blocked person 0.897411 per cent of the economics with the entity still blocked. Consistency across a 45-document compliance file fails with probability 90.056 per cent at a 5 per cent per-document error rate. The Medici bill of exchange passes the disclosure test in its strongest historical form while dry exchange fails it by construction.

Keywords: effective blocs; surplus headroom; break-even probability; ownership deeming rules; screening thresholds; bill of exchange; dry exchange; disclosure test

Series: D. Structure and Resilience **Paper:** CCP-063

1. Introduction

A framework that measures concentration is only worth what it says when it is pointed at a complete structure rather than at a convenient slice of one. This paper assembles a constructed multi-jurisdiction architecture in full, measures it with the corpus's own instrument at parameters fixed in advance, and reports the verdict the instrument returns including the part of that verdict which is unfavourable to the architecture.

The paper has four separable contributions.

The assembled measurement, and its failure to clear its own threshold. At the standing test parameters $\eta = 0.35$, $C = 0.30$, $Q = 10$ and $\delta = 0.02$ — giving a maintenance floor $\delta Q + C = 0.500000$ — a conventional single-perimeter structure measures one effective bloc and headroom $h = -0.500000$. The designed six-jurisdiction structure measures a bloc index of 0.340000, $N = 2.941176$ effective blocs and $h = -0.122417$. **The architecture closes 75.5167 per cent of the shortfall and does not close it**, against a break-even $N^* = e^{0.5/0.35} = 4.172734$. The counter-result is stated in the same breath: the only bloc assignment under which the structure clears the threshold is the naive one that treats coordinating jurisdictions as independent, which returns $h = +0.083099$ at the price of abandoning the coordination mechanism the framework exists to describe.

Marginal jurisdictions priced rather than argued. A candidate leg is a bet on a bloc classification, and a bet has a break-even probability. Setting the expected change in headroom to zero gives $p^* = L/(G + L)$, which for one candidate leg taking five per cent of exposure from a Western slot is $0.005110/(0.026731 + 0.005110) = 16.0476\%$, with a cost ceiling of 8.9102% of C. The largest single-move gain available anywhere in the structure, $+0.072279$, is refused, and the refusal is the result: a capped transit position delivers 3.9692% of that benefit for 3.2110% of the exposure, removing 96.79% of the tail.

An ownership deeming rule shown to be non-economic. A rule that deems an entity blocked when blocked persons hold fifty per cent or more directly or indirectly, individually or in the aggregate, is not a test on economic interest and does not behave like one. Seven tiers of 51% holdings leave the blocked person 0.897411% of the economics with the entity blocked in full, while two blocked persons at 24.50% each — the same 49% of economics as one at 49% — are not caught, and one further percentage point reverses the answer.

The disclosure test in historical instance. The Medici bill of exchange returned 4.166667% per round trip on the stipulated rates, an annualised 13.220482% at a 120-day interval against a 12.673611% simple loan rate that yields the identical 41.666667 florins — a difference of 0.000000000 florins between the prohibited instrument and the permitted one. The permitted instrument was fully disclosed to the very authority whose prohibition it was designed around and continued to work; dry exchange, its abusive twin, fails the same test by construction. **That is the disclosure constraint of Curzi [2026] — that a design whose effectiveness is structural is strengthened rather than weakened by publication, so that non-disclosure is refused as a design premise — observed in the field five centuries before it was stated as engineering.**

Why these sit in one paper. They are one instrument at four scopes. In each case a count of nominally distinct units overstates the number of things that can fail separately, and the correction is the same operation: group by whoever takes the decision, then square. That is the operation Curzi [2026] establishes as the concentration instrument, with a naive 3.407968 collapsing to exactly 1.000000 once a coordination perimeter is recognised, and it is the operation Curzi [2026] runs on a historical network, where five houses read 3.819710 effective jurisdictions naively and 1.000000 effective blocs once the Concert of Europe is treated as one perimeter, giving $h = -0.500000$. This paper adds the assembled case, the pricing method, the deeming-rule arithmetic and the historical instance.

Two conclusions the paper refuses. It refuses the reading under which the architecture is adequate, because that reading requires $\rho = 0$ between coordinating jurisdictions. And it refuses the reading under which the jurisdictional axis is the binding one: measured on the settlement-messaging axis the same structure returns $N = 1.305824$ and $h = -0.406608$, a gap of 0.284191 that is larger than everything the diversification programme achieved.

2. Hypotheses

Six hypotheses, each stated so that it could be wrong, each with its falsifier. Platt [1964]'s standard is applied: a hypothesis earns its place by naming the observation that would kill it. Three statements in the paper that are *not* falsifiable are labelled at the end.

H1 — the assembly hypothesis. A structure designed across coordination perimeters raises its surplus ceiling without moving its maintenance bill, so the ceiling term and the bill term respond to different interventions. *Predicted:* moving from one effective bloc to 2.941176 raises $\eta \ln N$ from 0.000000 to 0.377583 while $\delta Q + C$ stays at 0.500000, so h improves by exactly 0.377583 and remains negative. *Falsified by:* a bloc-count change that moves the maintenance floor, or a design that reaches $h > 0$ without moving C, η, Q or δ.

H2 — the binding-axis hypothesis. A structure is as diversified as its most concentrated axis, so measuring one axis and reporting the result overstates resilience whenever another axis is more concentrated. *Predicted:* on the constructed shares the jurisdictional axis gives $N = 2.941176$ and the settlement-messaging axis $N = 1.305824$, so the two readings of h differ by 0.284191. *Falsified by:* an axis decomposition on which the jurisdictional reading is the minimum, or a demonstration that the axes are not simultaneous exposures of the same capital.

H3 — the pricing hypothesis. A marginal jurisdiction admits a break-even probability and a cost ceiling, and both are computable, so a bloc call is a decision under uncertainty rather than an argument. *Predicted:* $p^* = L/(G + L)$ returns 16.0476% for the Eurasian candidate, 2.7108% for the China-bloc candidate, 36.1137% for a second-bridge candidate and exactly 0.0000% for a candidate drawn from the bloc it would fall back into. *Falsified by:* a candidate whose gain and loss cannot be computed from the same share vector, or a case in which p^* is not monotone in the loss.

H4 — the deeming hypothesis. An aggregate-ownership deeming rule is not a test on economic interest, and the two diverge without bound. *Predicted:* iterating 51% holdings, economic interest falls as 0.51^k — reaching 0.897411% at $k = 7$ — while the deemed status is unchanged at every tier; and two holders at 24.50% are not caught while two at 25.00% are. *Falsified by:* a monotone function mapping deemed status onto economic interest under the stated rule.

H5 — the consistency hypothesis. Compliance-file failure is a property of how documents are produced and not of how carefully they are checked. *Predicted:* with independently drafted documents at per-document error rate p, $P(\text{at least one inconsistency}) = 1 - (1 - p)^n$ reaches 90.056% at $n = 45$, $p = 0.05$, and falls to 8.615% at $p = 0.002$ — a factor of 10.45, achieved by changing the generation process rather than the review effort. *Falsified by:* a demonstration that review effort reduces p by the same factor at comparable cost, or evidence that document errors are strongly enough correlated that the independence assumption inverts the conclusion.

H6 — the disclosure hypothesis in historical instance. An instrument whose function is structural survives full disclosure to the authority whose rule it is designed around, while an instrument whose function depends on an unexamined element does not. *Predicted:* the bill of exchange, written, ledgered, notarised and publicly quoted, was disclosed to the promulgating authority — which was also the principal client — and continued to work; dry exchange, identical on the page, collapses the moment the return leg is examined. *Falsified by:* an element of the bill whose disclosure defeated it, or a documented dry exchange that survived examination of its return leg.

Not hypotheses, and labelled as such.

- *The bloc assignment* is a **modelling choice**. Which jurisdictions share a coordination perimeter is stipulated, not estimated; Section 3.3 states what that costs and what the remedy would be.
- *The parameter set* (η, C, Q, δ) is a **stipulation** fixed in advance of the measurement, which is what protects H1's criterion from being tuned to fit. It is not an estimate of any real position, and Section 14 gives the sensitivity that matters.
- *The identity of the headroom law's functional form* is a **modelling choice**, not a derived result. Section 3.1 names the alternatives that fit the same qualitative facts.

3. Method: the instrument, the parameters, and what is stipulated

3.1 The headroom law and which effective count

The framework's build criterion is a single inequality on a headroom quantity,

$$h = \eta \ln N - \delta Q - C,$$

in which N is an effective count of independent units, η converts input into durable stock, Q is the stock carried, δ its decay rate per period, and C a flow cost. $h < 0$ forbids building; $h > 0$ caps every increment at $r \leq h$. That gate is Curzi [2026]'s, worked there at $N = 18$ with $\eta = 0.30$, $\delta = 0.06$, $C = 0.55$, $Q = 1.20$ giving $h = +0.2451$; this paper reuses it unchanged as a precondition on the assembled structure. The decay term is the leaky-integrator form whose time constants Curzi [2026a] tabulates — $\tau = 1/\delta$ against $\tau_{\text{exact}} = -1/\ln(1-\delta)$, agreeing within one per cent at $\delta \leq 0.02$ — which is why $\delta = 0.02$ may be read as a fifty-period time constant without correction.

Which effective count, stated exactly, because the corpus runs two. This paper uses $N = \text{HHI}^{-1}$, the reciprocal of the sum of squared shares. That quantity is the Hill number of order two [Hill, 1973, Jost, 2006]: the number of equally-sized units that would produce the observed index. It is the reciprocal of Simpson [1949]'s index, the effective number of parties of Laakso and Taagepera [1979], and the numbers-equivalent reading of the concentration index attributed to Adelman [1969]; the index itself is Herfindahl [1950] and Hirschman [1964], and the thresholds at which competition authorities read it are policy choices that have moved [U.S. Department of Justice and Federal Trade Commission, 2023, 2010]. The exponential of Shannon entropy [Shannon, 1948] is the same family at order one, and HHI^{-1} is also the exponential of the Rényi entropy of order two [Rényi, 1961].

The corpus's other effective count is not a Hill number and is not used here. The affine form $N_{\text{eff}} = 1 + (N-1)u$ of Curzi [2026] and Curzi [2026c] — the latter badging it a design axiom rather than a derived result, with a hundred-person department at $u = 0.05$ reading an effective count of 5.95 — is a linear interpolation between 1 and N. It is not a numbers-equivalent of any index and no ecological, economic, political-science or information-theoretic literature contains it. Curzi [2026] runs the verification chain on both and lets the affine one fail visibly. **Every effective count in this paper is HHI^{-1}, and where the affine count is referred to it is named and marked as the corpus's own construction.**

Why a logarithm, and what does not discriminate it. The $\ln N$ term is a modelling choice and this paper does not pretend otherwise. A sub-linear power law [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the same qualitative facts — diminishing returns to added independent units, no saturation within the observed range — as well as the logarithm does, and nothing measured in this paper separates them. The strongest available warrant for the logarithm is not the perceptual one: Peters and Gell-Mann [2016] derives a logarithmic form from multiplicative dynamics with no psychological assumption at all, which is what the argument here actually needs, since nothing in a jurisdictional exposure vector is a percept. Curzi [2026n] states the corpus's own Fechnerian reading. **The logarithm is adopted, its rivals are named, and the paper concedes that nothing here discriminates among them.**

The large-N limit is a defect of the model and is stated as one. h grows without bound in $\ln N$; the empirical diversity literature does not [May, 1972, Doak et al., 1998]. The model's domain is small N, and the structures measured here have $N < 6$.

3.2 The bloc collapse, and why it is a variance identity

Grouping units until the answer changes is a recognised way to lie, so the grouping requires a foundation that is not assertion. It has one. Model the exposure as a portfolio [Markowitz,

1952]: let w_i be the share exposed to unit i and X_i a unit-variance shock to it. The aggregate shock is $R = \sum_i w_i X_i$, with $\mathrm{Var}(R) = \sum_i \sum_j w_i w_j \rho_{ij}$.

Under independence the double sum is $\sum_i w_i^2$, which is the naive index exactly. Under perfect correlation inside blocs and independence across them it is $\sum_b W_b^2$, the bloc index exactly. **The two indices are one variance under two correlation assumptions, and $\mathrm{Var}(R)^{-1}$ is the effective number of independent shocks.** The general form $\mathrm{EFF} = 1/(w'Rw)$ is the similarity-sensitive diversity of order two [Leinster and Cobbold, 2012] whose additive relative is Rao [1982]'s quadratic entropy. Curzi [2026] establishes this identity and verifies both endpoints on the same six shares to twelve decimal places; this paper reuses it and does not re-derive it.

The intermediate cases are not a nicety. Measured correlations rise under stress [Longin and Solnik, 2001, Ang and Chen, 2002], so **the effective count falls exactly when it is needed** — the ecological name for the same variable is synchrony [Loreau and de Mazancourt, 2008], and the ecological demonstration that effective count buys stability is Schindler et al. [2010]'s sockeye record. The coupled-network result of Buldyrev et al. [2010] is the formal statement of "units inside one perimeter behave as one and worse than one", and is used here for its structural form only, since its percolation premises have not been shown to hold for jurisdictions.

The contrary literature is read alongside and is not decorative. Doak et al. [1998] show diversity–stability correlations can arise from statistical averaging with no mechanism, so portfolio arithmetic is not by itself evidence of one; May [1972] shows that increasing diversity destabilises randomly assembled systems. On the generating side, Bouchaud and Mézard [2000] derives concentration from multiplicative growth with exchange rather than observing it. **This paper asserts no distributional form and no tail**; Section 13.7 records the consequence.

3.3 Identification: the bloc assignment is a stipulated reference group

The framework's reflexive law — a group's aggregate state influences its members while the members constitute the aggregate — is a linear-in-means reflexive specification, and Manski [1993] proves that its endogenous effect is **not identified from observational data**: it cannot be separated from contextual effects or from correlated unobservables however much data is collected. Shalizi and Thomas [2011] reinforce the point for network data, where influence and selection are generically confounded.

A bloc assignment is exactly a reference-group specification, and this paper adopts the stipulation remedy. Which jurisdictions coordinate is asserted, with an evidence class and an observable that would change the call, following the template of Brock and Durlauf [2001]. Randomisation, an instrument and partial-population variation are all unavailable here. **The price is that every bloc-level result in this paper is a modelling position and not an identified empirical result**, and Section 12.1 gives the robustness table that lets a reader substitute a different assignment.

The corpus's own reflexivity applies to the instrument as well. Lucas [1976] states the consequence: parameters estimated before a rule is published need not survive its publication. MacKenzie and Millo [2003] documents the mechanism in the field, with dates and prices, and Goodhart [1975] and Campbell [1979] name what happens when a measure becomes a target. **The framework asserts that systems change when modelled and then exempts its own η, δ and correlation assumptions from that same reflexivity. That is an inconsistency and this paper does not resolve it**; what it offers instead is the auditable reporting pack of Appendix A.

3.4 The four axes and the rule that the most concentrated one binds

Capital is exposed on more than one axis at a time, and the axes are simultaneous rather than alternative. The instrument is therefore run four times on the same structure — jurisdiction, counterparty, currency and settlement rail — and **the most concentrated axis is reported first**,

because a structure is as diversified as its worst axis and the axis an analyst measures is the one already familiar.

The rail axis requires a second collapse of its own, and it is the collapse that carries the result. A clearing system and a messaging layer are different objects, and participants in a distinct clearing system may still use the dominant messaging layer, so a structure reporting two rails may have one messaging dependency. **Collapsing rails by messaging layer is the same operation as collapsing jurisdictions by perimeter, and a structure that performs it on jurisdictions and not on rails has done half the measurement and reported all of it.** Curzi [2026] is the primary treatment: a naive correspondent index of 0.302400 (3.306878 effective correspondents) collapses to a perimeter index of 0.745256 and $N_{\text{rail}} = 1.341819$ once flows are decomposed by currency. This paper's rail measurement derives from that result.

3.5 What is constructed, what is described, and what is cited

Everything numerical in this paper is CONSTRUCTED. The entities, the exposure shares on all four axes, the cost allocation, the hazard rates, the error rates, the screening population, the branch and loan-book vectors and the exchange quotations are stipulated for illustration. No dataset was collected and no real structure, institution, bank, name or shareholding appears anywhere.

Jurisdiction-specific legal material is described from the source manuscript and is not cited. The bibliography contains general instruments and comparative doctrine only: it holds no statutes, regulations, court rules, registry documents, agency guidance or case reports for any individual jurisdiction, and none is cited here. Appendix B lists every affected claim. Where a named centre's legal characteristic is relied upon and cannot be supported from the bibliography, this paper follows Curzi [2026] and de-brands the description — a *Gulf free-zone centre*, an *Indian international financial centre*, a *Eurasian financial centre* — retaining the name only where the arithmetic of a specific published result requires it. **A description is not an assertion, and a de-branded description is not a claim about any particular jurisdiction.**

What *is* cited, and genuinely supports what it is attached to: the concentration literature of Section 3.1; the portfolio and ecological literature of Section 3.2; the automatic-exchange, anti-money-laundering and beneficial-ownership regimes as the reason non-disclosure is no longer an available design parameter [Organisation for Economic Co-operation and Development, 2017, Financial Action Task Force, 2012, Organisation for Economic Co-operation and Development, 2015]; the listing machinery as a technique of governance [de Goede, 2011]; asset partitioning as the irreducible function of entity law [Hansmann and Kraakman, 2000] and the scholarly account of legal coding that is simultaneously its sharpest critic [Pistor, 2019]; and, for the historical material, de Roover [1963], which is the canonical modern scholarship and the source any reader wanting archival figures should use instead of this paper.

4. The assembled structure

4.1 Entities, blocs and exposure

The structure is a constructed apex-and-legs architecture: a foundation with separate legal personality at the top, holding companies in two free-zone centres, an operating company in a non-aligned manufacturing jurisdiction with a listed minority beneath it, an English-law contracting company with a converted real-estate position, a captive insurer against real third-party liabilities, a regulated fund and leasing vehicle in an Indian international financial centre, and a treasury and fund wrapper in a bridge jurisdiction, with per-asset ring-fencing vehicles below.

entity	bloc	job	function
foundation (apex)	BRIDGE	holding	separate legal person; holds everything below
Gulf holding company	BRIDGE	holding	the operating chain and the intangibles entity
Gulf trading company	BRIDGE	operating	the trading book
operating company	NON-ALIGNED	operating	manufacturing, labour, premises
listed minority	NON-ALIGNED	—	20% public float
English-law company	WEST	bridging	contracting
real-estate position	WEST	—	immovable, converted to movable
captive insurer	WEST	risk	reserves against real third-party liabilities
Indian centre vehicle	NON-ALIGNED	holding	regulated fund and leasing
treasury and fund wrapper	BRIDGE	bridging	settlement, foreign exchange, custody
ring-fencing vehicles	follows parent	—	per-asset segregation

The job column is the framework's holding / operating / bridging taxonomy, whose point is that one entity doing three jobs converts income into durable stock less efficiently than three entities doing one each. That friction lands on η, which Section 7.4 shows is the most powerful lever in the model.

Constructed exposure shares, carried unchanged from Curzi [2026]:

```
Gulf centre 29.1% · manufacturing jurisdiction 18.2% · UK 16.4%
Bermuda 13.6% · India 11.8% · Singapore 10.9%

naive HHI 0.189002   effective jurisdictions 5.290949
BRIDGE 40.0% / NON-ALIGNED 30.0% / WEST 30.0%
bloc-HHI  0.340000   EFFECTIVE BLOCS N = 2.941176
Western-perimeter exposure 30.0%
```

4.2 Five flows, and the rule that they must not share a path

A structure is not its ownership chart. It is five distinct flows, and most failures come from running two of them down one path.

flow	what moves	failure if it shares a path
ownership	shares, membership interests, the charter	a change in one flow silently changes control
cash	dividends, loans, service fees, settlement	a treasury dependency becomes an ownership dependency; a frozen account stops a governance process

flow	what moves	failure if it shares a path
contract	licences, supply agreements, facilities	an operating dispute reaches the apex because the apex was party to an operating contract
information	board packs, filings, onboarding packs, ownership records	a compliance answer in one jurisdiction contradicts one in another
decision rights	who instructs whom, reserved matters, resolutions	one person is the decision path for two blocs, so the bloc diversity is nominal

The design rule is that each flow should fail without taking the others with it, and the test is to name, for each flow, the single point whose removal stops it. Two flows naming the same point share a path. On the constructed structure: ownership and decision rights both terminate at the apex charter and council, which is deliberate and is what an apex is for; cash terminates at the treasury's principal correspondent; contract at the chosen governing law and arbitral seat; information at the group's single-source compliance file. **Only the cash row is a residual concern, and it is the settlement-rail exposure of Section 6 wearing different clothes.**

This is the same structural discipline Curzi [2026] states as channel separation — that a system is in-band when a value in one plane can be promoted to an instruction in another, and out-of-band when the two are disjoint by construction — applied to a corporate group rather than to a signalling network. **Separation is a property of the wiring; keeping it is a property of who is allowed to touch the wiring**, which is the governance result Ostrom [1990] reaches empirically for common-pool institutions and which Sagan [1993] shows can fail when added redundancy adds coordination surface instead of removing exposure.

5. Results: the assembled measurement

5.1 H1 — the ceiling rises and the bill does not: supported

At the stated parameters the maintenance floor is $\delta Q + C = 0.02 \times 10 + 0.30 = 0.500000$.

N	ln N	CEILING (eta ln N)	BILL (delta Q + C)	h
1.000000	0.000000	0.000000	0.500000	−0.500000
2.216500	0.795929	0.278575	0.500000	−0.221425
2.941176	1.078810	0.377583	0.500000	−0.122417
3.174603	1.155183	0.404314	0.500000	−0.095686
3.615826	1.285320	0.449862	0.500000	−0.050138
4.172734	1.428571	0.500000	0.500000	+0.000000
5.000000	1.609438	0.563303	0.500000	+0.063303

Both structures fall below the build threshold: the conventional one by 0.500000 and the designed one by 0.122417. The architecture closes $0.377583/0.500000 = 75.5167\%$ of the shortfall and does not close it. The ceiling column is what diversification produces; the bill column is untouched by it. At $N = 1$ the ceiling is exactly zero, so no amount of skill on the input side matters, because $\eta \times 0 = 0$ for every η.

H1 is supported. The correct reading is that jurisdictional diversity buys the *ability* to build and not the building: it raises the ceiling without paying the maintenance bill.

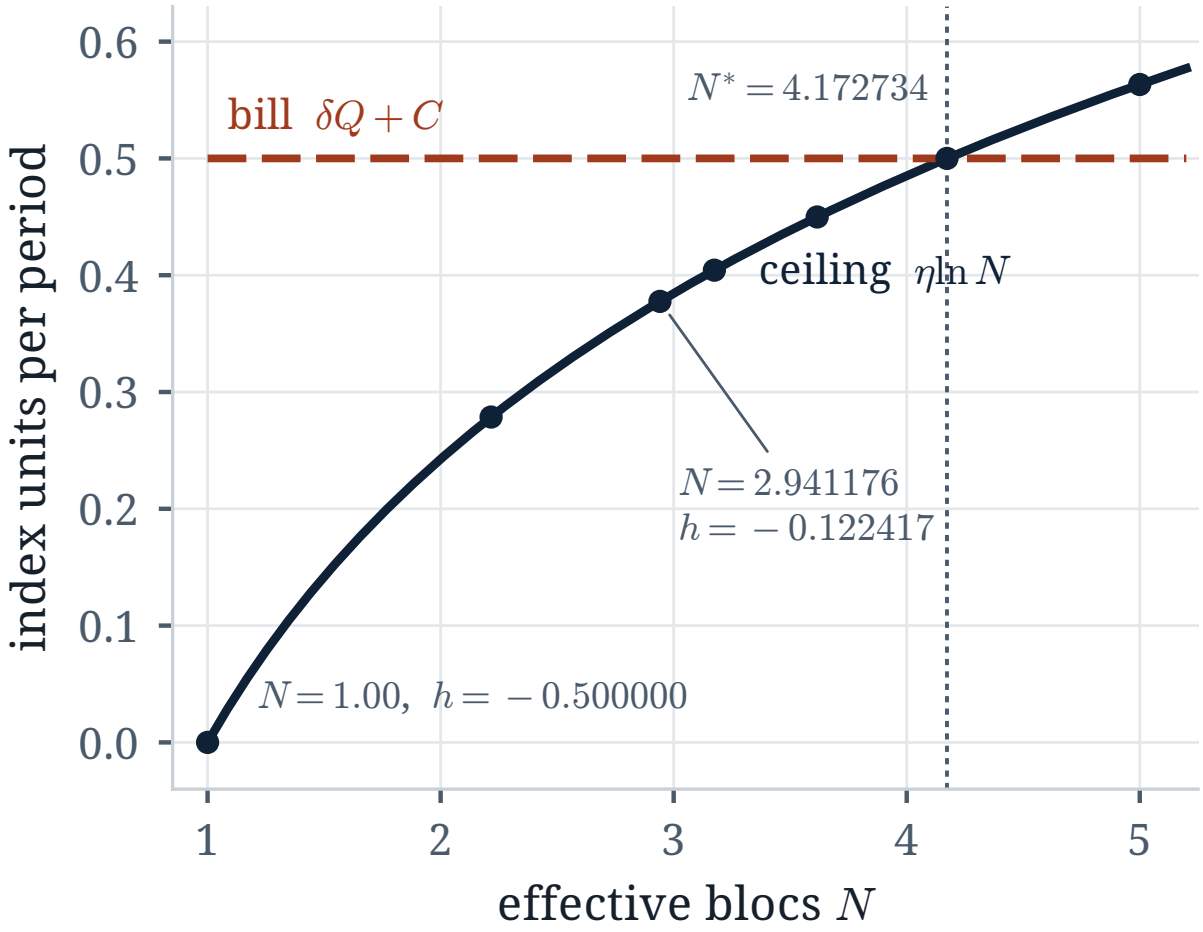


Figure 63.1: **The ceiling rises and the bill does not.** The surplus ceiling $\eta \ln N$ climbs from 0.000000 at one effective bloc to 0.377583 at 2.941176, while the maintenance floor stays flat at 0.500000; the two cross at the break-even $N^* = e^{0.5/0.35} = 4.172734$, which the designed structure does not reach. Modelled at the stated parameters $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$; script `figs/src/CCP-063_fig01.py`.

5.2 Robustness, and the counter-result

The bloc assignment is the judgemental step in the whole measurement, so the assembled h is recomputed under every defensible alternative.

assignment variant	bloc-HHI	N	h	vs base
base: captive centre folded into WEST	0.340000	2.941176	-0.122417	$+0.000000$
captive centre as its own light-Western bloc	0.295392	3.385332	-0.073192	$+0.049225$
bridge jurisdiction reclassified WEST	0.341962	2.924302	-0.124431	-0.002014
Gulf centre reclassified WEST (worst case)	0.451162	2.216499	-0.221425	-0.099008
Gulf and bridge centres as separate blocs	0.276562	3.615826	-0.050138	$+0.072279$
the two non-aligned legs as separate blocs	0.297048	3.366459	-0.075148	$+0.047268$
every jurisdiction independent (naive)	0.189002	5.290949	$+0.083099$	$+0.205516$

Three readings, and the third is why the table is printed.

The published base case is conservative rather than flattering. Folding the captive centre into the Western perimeter costs 0.049225 of h against treating it as its own; treating the two bridge centres as one costs 0.072279. The figure this paper publishes sits at the low end of a defensible range.

The result is insensitive to the most contested single call. Reclassifying the bridge jurisdiction as Western moves h by 0.002014, under two-tenths of one per cent of the maintenance floor. The argument most likely to be had about this structure does not move the answer. Only the Gulf-centre call does real damage, taking h to -0.221425.

And the last row is the counter-result. The only assignment that puts the structure above the build threshold is the one that treats every jurisdiction as independent — $\rho = 0$, $h = +0.083099$. Curzi [2026] computes that the conventional structure would need Western jurisdictions to correlate at $\rho = 0.065910$ for its naive figure to be the right one. **The architecture clears its own threshold if and only if the coordination mechanism the framework exists to describe is assumed away. The favourable answer is available at exactly the price of abandoning the thesis, and this paper declines to pay it.**

5.3 The metric that moved when the world did not

A third reading, carried because it is the cleanest available demonstration that the naive count reports diversification that did not occur. Hold the same 10.9% of exposure in a bridge juris-

diction classified as Western, or in the Western jurisdiction directly:

	naive HHI	eff jur	bloc-HHI	N	h
held in the bridge jurisdiction	0.189002	5.290949	0.341962	2.924302	−0.124431
same exposure held in the West	0.224754	4.449309	0.341962	2.924302	−0.124431
difference		+0.841640	0.000000	0.000000	0.000000

Eight-tenths of an effective jurisdiction on the naive count, and exactly zero on the bloc count, on N, and on h. The metric moved and the world did not. That is the general form of the error Curzi [2026] measures on a historical network and Curzi [2026] measures on a payment rail: rearranging exposure inside a perimeter changes the naive index and nothing else.

6. Results: H2 — the binding axis

The instrument is run four times on the same constructed capital.

axis	shares	naive HHI	effective count
jurisdiction	29.1 / 18.2 / 16.4 / 13.6 / 11.8 / 10.9	0.189002	5.290949
jurisdiction, blocs	40.0 / 30.0 / 30.0	0.340000	**2.941176**
counterparty	38 / 27 / 16 / 11 / 8	0.261400	3.825555
currency	44 / 17 / 15 / 13 / 7 / 4	0.268400	3.725782
settlement rail, naive	62 / 14 / 11 / 5 / 8	0.425000	2.352941
settlement rail, collapsed to messaging	87 / 5 / 8	0.765800	**1.305824**

The rail row is the one that matters and it is the one requiring the collapse. Reported naively the structure has 2.352941 effective rails; collapsed by messaging layer it has 1.305824. The consequence for the headroom law is stated so it cannot be missed:

N read off the jurisdictional axis	2.941176	h = −0.122417
N read off the rail axis	1.305824	h = −0.406608
difference		0.284191

The gap is larger than everything the diversification programme achieved (0.377583 of ceiling bought, 0.284191 of gap left on another axis). H2 is supported. This does not make the jurisdictional reading wrong; it means the structure has two answers and the honest one is the smaller. **A measurement pack that prints one of them has chosen which one to print, and the choice will not be the unflattering one unless the format forces it** — which is why Appendix A requires both.

This is the exposure Curzi [2026] names as invisible to the concentration instrument by construction: a count on one axis cannot see a common-mode dependency on another, and substrate concentration must be measured separately or it will not be measured at all. Two nominally independent perimeters resting on one messaging substrate are one exposure, which is Buldyrev et al. [2010]'s coupled-network result in structural form and Markowitz [1952]'s covariance point in its original one. **The counterparty axis is also more concentrated than the naive jurisdictional reading, and the currency axis nearly so; the jurisdictional reading is the third best of four.**

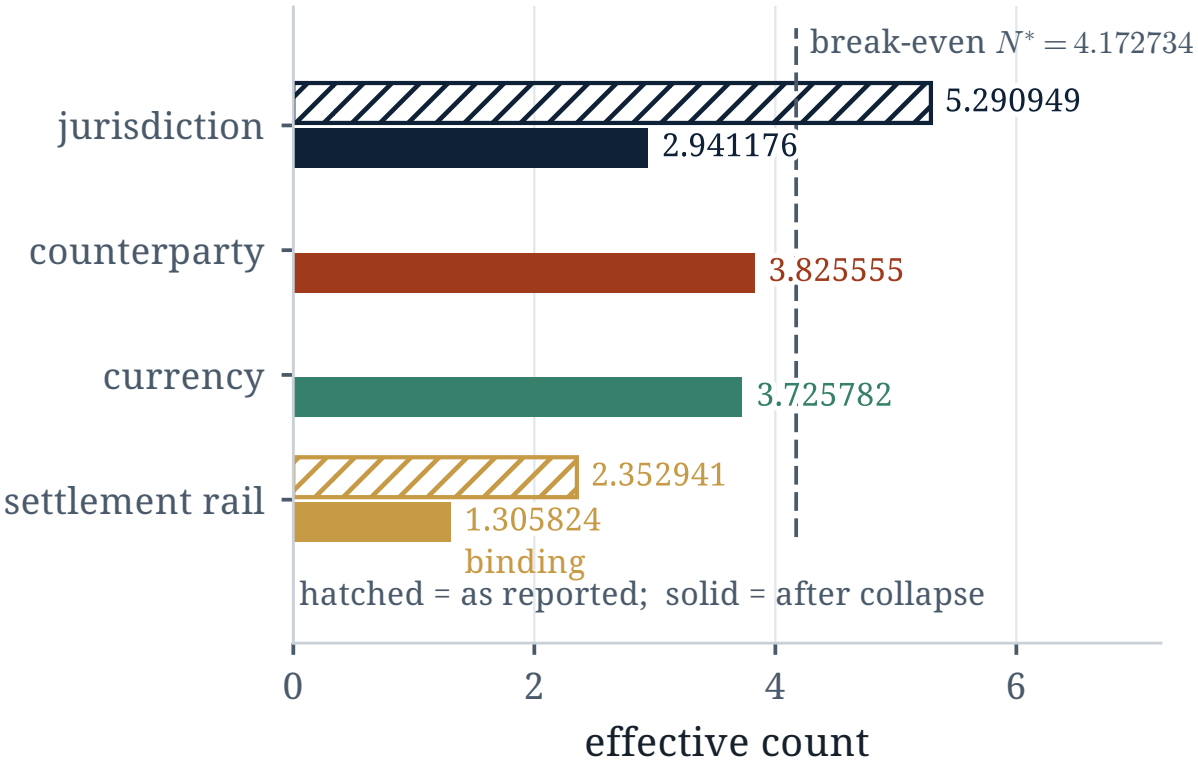


Figure 63.2: **The binding axis is not the one the design programme worked on.** Effective counts on four simultaneous exposure axes of the same constructed capital: the jurisdictional bloc count reads 2.941176 while the settlement-messaging count reads 1.305824, and reading h off the second rather than the first moves it from -0.122417 to -0.406608. Measured from the constructed share vectors of Section 6; script `figs/src/CCP-063_fig02.py`.

6.1 The threat map: what each element is actually for

The concentration measurement cannot answer the question of what an element is for, and the answer changes how the structure is judged.

threat	answered by	measured by
arbitrary action — freeze, expropriation, de-risking without process	bloc diversity across three perimeters	$N = 2.941176$
lawful process — a judgment that survived due process	**not answered by this architecture**; timing-bound and governed by transfer-avoidance law	not measured here
commercial and counterparty risk	captive against real third-party liabilities; per-asset ring-fencing	counterparty HHI 0.261400, N 3.825555
succession fragmentation across heir jurisdictions	the apex foundation and its charter	$dh/dN = 0.000000$; succession fora $6 \to 3 \to 2$
bloc fragmentation — the perimeters themselves moving	dual-rail settlement; the bridge positions	rail HHI 0.765800, N 1.305824

Two rows carry the section. **The succession threat is answered by the element whose contribution to N is exactly zero, and the fragmentation threat is measured on the axis where the structure scores worst.** The apex does the most important job in the architecture and contributes nothing to the headline number; the rail dependency is the worst exposure and appears in no jurisdictional table. **A structure judged on N alone would delete the first and never find the second.**

And the lawful-process row is unanswered deliberately. A judgment that survived due process is a different problem, governed by transfer-avoidance law and by timing, and nothing in this architecture is offered against it. **A paper implying otherwise would be describing a structure for a job it cannot do.**

7. Results: H3 — pricing a marginal jurisdiction

7.1 The break-even probability

A candidate leg is a bet on a bloc classification. If the call is right the leg lands in a new perimeter and N rises; if it is wrong it folds into a perimeter already held and N falls, because the exposure was drawn from somewhere. Both outcomes are computable from the same share vector:

$$G = h(\text{with leg, call right}) - h(\text{without}), \qquad L = h(\text{without}) - h(\text{with leg, call wrong}),$$

and setting the expected change to zero gives $p^* = L/(G + L)$.

candidate	bloc if right	G	$-L$	break-even p^*
Eurasian centre, 5% from the Western slot	EURASIA	$+0.026731$	-0.005110	**16.0476%**

candidate	bloc if right	G	$-L$	break-even p^*
China-bloc centre, 10.9% from the bridge slot	CHINA	$+0.072279$	-0.002014	2.7108%
second bridge centre, 5% from the Western slot	a second BRIDGE	$+0.026731$	-0.015110	36.1137%
African centre, 5% from the Western slot	AFRICA	$+0.026731$	$+0.000000$	0.0000%

The candidates are named in the source manuscript as the Astana centre, Hong Kong, T”urkiye and Mauritius, and the abstract of this paper keeps the first two names because the arithmetic of the first is a published corpus result. **They are de-branded in the table because each row depends on a legal or political characteristic of the named centre that this paper’s bibliography cannot support**, and because the arithmetic depends on nothing but which slot the exposure is drawn from and which bloc it lands in. Appendix B lists what would have to be verified before any name is reattached.

The Eurasian leg is Curzi [2026]’s result and this paper reuses its arithmetic unchanged: a right call takes N to 3.174603 for a gain of $+0.026731$; a wrong call takes it to 2.898551 for a loss of 0.005110; the expected change is $0.0318401p - 0.0051096$, zero at $p = 16.0476\%$. **That is a question with an answer, and it replaces an argument that has none.**

Four readings, and they are more interesting than the numbers.

A candidate drawn from the bloc it would fall back into is a free option. If the African call is right N improves; if it is wrong the structure is exactly where it started. The break-even probability is exactly zero, so the only remaining question is cost. That is the cheapest jurisdictional move available and almost nobody frames it that way.

A candidate drawn into a bloc it would concentrate is expensive. A second bridge centre alongside the two already held takes the BRIDGE bloc from 40.0% to 45.0% and makes the structure worse if the call is wrong. Its break-even is 36.1137%, more than double the Eurasian candidate’s on the same five per cent of exposure — and the difference is entirely about which slot it was drawn from and which bloc it lands in.

p^* **prices classification risk and nothing else.** It says nothing about the risk that a jurisdiction is *correctly* classified and then behaves catastrophically. The China-bloc candidate’s break-even is 2.7108%, an almost trivial bar, and Section 7.3 refuses the leg anyway. That is not an inconsistency; it is the recognition that two different uncertainties are in play and arithmetic on the first does not address the second.

7.2 The cost ceiling, which is the binding constraint

The whole gain from the Eurasian leg is 0.026731, which is $0.026731/0.30 = 8.9102\%$ of the total cost flow C. **The leg must run for less than 8.9102% of the structure’s entire maintenance budget or it is negative regardless of whether the bloc call was right.** On the cost allocation of Section 7.5 that ceiling sits below the smallest existing line, which is a harder test than any argument about regional alignment. **The binding constraint on a marginal leg is cost, not geopolitics, and the cost test is the one nobody runs.**

H3 is supported: both quantities are computable from the structure’s own share vector, and neither requires an opinion.

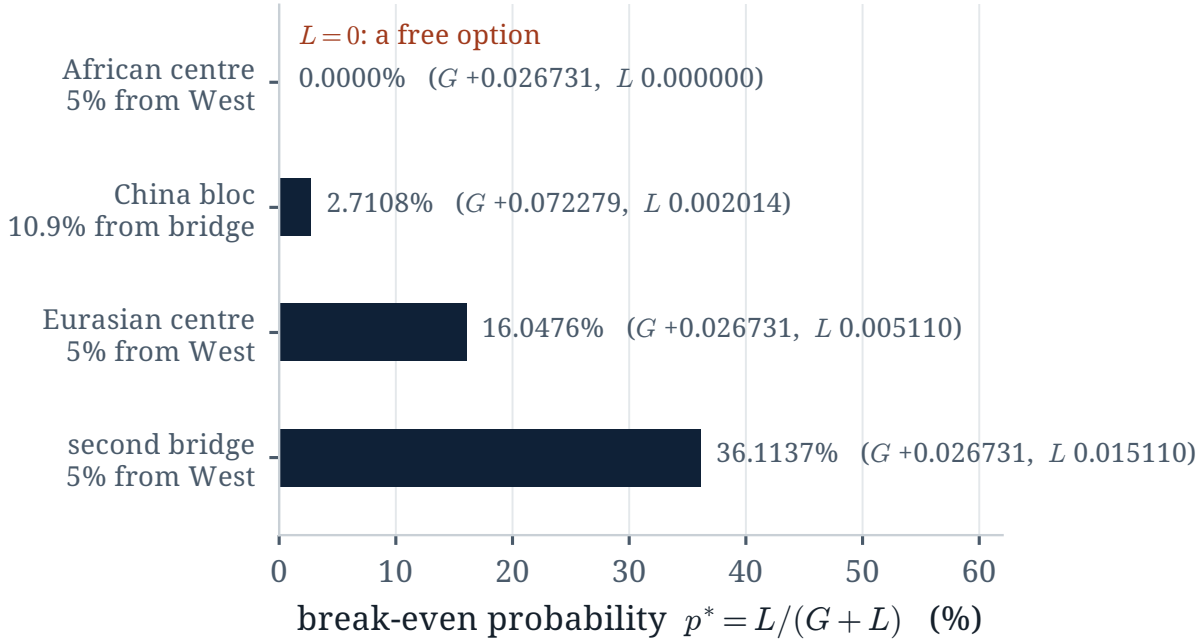


Figure 63.3: **Every marginal jurisdiction is a break-even probability with a cost ceiling attached.** Four candidate legs, each priced from the same constructed share vector: a candidate drawn from the bloc it would fall back into has $p^* = 0.0000\%$ and is purely a cost question, while a candidate that would concentrate a bloc already held reaches 36.1137%. Computed from the gains and losses in Section 7.1; script `figs/src/CCP-063_fig03.py`.

7.3 The largest available gain, refused

The China-bloc candidate in the same 10.9% slot produces $N = 3.615826$ and $h = -0.050138$: **a gain of $+0.072279$, the largest single-move gain anywhere in the structure.** It is refused, and the refusal is the result.

The reconciling principle is that protection is bought for the tail rather than for the median. A jurisdiction whose ordinary commercial courts work well and whose behaviour under political stress is uncertain is being purchased for exactly the state of the world in which the second property governs. **The median outcome is excellent and irrelevant**, because the premise of the purchase is the stressed state and not the current one.

Hence a capped transit position instead — 0.35% of exposure rather than 10.9%, taken as transit rather than as a store, under the framework's standing rule that a rail is not a holding:

```
full leg          10.90% exposure   N 3.615826   h −0.050138   gain +0.072279
capped transit     0.35% exposure   N 2.965384   h −0.119548   gain +0.002869
ratio              3.2110% of the exposure                     3.9692% of the benefit
```

The capped position delivers 3.9692% of the benefit for 3.2110% of the exposure — very slightly benefit-efficient relative to the full leg — and removes 96.79% of the tail exposure. That is the trade, computed.

The evidence for the tail judgement is not in this paper's bibliography. The source manuscript grounds the refusal in a sequence of security legislation, a decision of a final appellate court, an interpretation, a later ordinance reaching state secrets, and the withdrawal and resignation of overseas judges. **None of those is cited, because the bibliography contains no jurisdiction-specific statutes and no case law**; they are described in Appendix B and require verification against the primary sources. **The refusal is therefore a stipulated modelling po-**

sition and not an established finding, and a reader who assesses the tail differently should substitute their own judgement and re-read the arithmetic, which is unaffected.

7.4 The three levers, and the sensitivity that ranks them

At $N = 2.941176$, ceiling 0.377583, floor 0.500000, shortfall -0.122417, each of four single-parameter moves is sufficient alone:

lever	required value	from	change needed
C down	0.177583	0.3000	-40.8055%
η up	0.463474	0.3500	$+32.4211\%$
N up	4.172734	2.941176	$+41.8720\%$
Q down	3.879169	10.0000	-61.2083%

and the sensitivity per ten per cent move ranks them:

```
parameter    10% move    change in h
-----------------------------------
eta          +10%         +0.037758
N            +10%         +0.033359
C            −10%         +0.030000
Q            −10%         +0.020000
delta        −10%         +0.020000
```

η is the most powerful lever per unit of itself and it is the one nobody works on, because it is not a jurisdiction and cannot be bought. It is the conversion of input flow into durable stock, and the mechanism that raises it is job separation: an entity doing one job converts better than an entity doing three, because the friction of asking a holding vehicle to operate shows up as leakage on exactly this term.

7.5 The cost side, decomposed

$C = 0.30$ is an abstraction until it is allocated. On a constructed allocation:

line	C	% of C
operating-company substance	0.0520	17.333%
apex foundation	0.0450	15.000%
English-law company and property	0.0410	13.667%
Gulf holding company	0.0380	12.667%
treasury and fund wrapper	0.0360	12.000%
captive insurer	0.0340	11.333%
Indian centre vehicle	0.0210	7.000%
group compliance and screening	0.0190	6.333%
cross-family legal opinions	0.0140	4.667%
total	**0.3000**	

Two observations change decisions. **The substance lines are the largest and they are the ones that cannot be cut**, because substance is what makes an entity real; reducing them defeats the position they pay for, and a cost programme that starts there dismantles the structure and then reports the saving. **The compressible lines are the interfaces and the duplication.** A

star interface topology rather than a mesh — nine interfaces on ten touchpoints instead of forty-five, three canonical legal-family interfaces instead of six — saves 0.014 of C on the opinion budget alone, which is 4.6667% of the total and 52.3736% of the entire headroom gain from adding a whole additional jurisdiction. **The cheapest available improvement to h is an architectural choice that costs nothing and is usually made by accident.** That the interface count follows the organisation's own communication structure is Conway [1968]'s observation; that the accounting shows the slack as waste before it shows it as insurance is Cyert and March [1963]'s.

Stacked, with each step's source:

step	N	C	eta	h
designed, as built	2.941176	0.3000	0.350	−0.122417
+ star interface topology	2.941176	0.2860	0.350	−0.108417
+ Eurasian leg at 0.021 of C	3.174603	0.3070	0.350	−0.102686
+ the same leg at its 0.0267 ceiling	3.174603	0.3127	0.350	−0.108386
+ eta to 0.42 via job separation	3.174603	0.3070	0.420	−0.021823
+ C to 0.22	3.174603	0.2200	0.420	+0.065177

Two rows deserve slow reading. The leg at its ceiling cost gives $h = -0.108386$ against the star topology's -0.108417 — a difference of 0.000031, which is what a break-even ceiling means when it is computed rather than asserted. And the row that finally crosses zero requires $\eta = 0.42$ and $C = 0.22$ simultaneously, on top of an additional jurisdiction. **That is the honest size of the remaining problem, and none of it is on the jurisdictional side.**

7.6 Order of build, and an element worth zero

Returns to additional blocs fall as $dh/dN = \eta/N$, so build order has value.

move	N after	h after	dh
start: conventional	1.000000	-0.500000	—
+ apex foundation	1.000000	-0.500000	$+0.000000$
+ Gulf holding company, to 29.1%	1.700000	-0.314280	$+0.185720$
+ operating company, 18.2%	2.300000	-0.208482	$+0.105798$
+ Indian centre vehicle, 11.8%	2.600000	-0.165571	$+0.042911$
+ treasury, 10.9%	2.941176	-0.122417	$+0.043154$

The first non-Western move is worth 0.185720 and the last 0.043154 — a factor of 4.30. A half-finished programme in the right order is worth far more than a half-finished one in the wrong order.

And the apex row is zero, and the apex is still the most important element. It answers the succession threat rather than the arbitrary-action one, it is the layer that makes everything below administrable, and it saves 0.014 of C on the interface budget independently of anything it does for N. **An element that does not move the concentration measurement is not thereby worthless; it is answering a different question**, and judging every element against N is the failure mode the instrument invites. That asset partitioning rather than diversification is the irreducible function of the entity form is Hansmann and Kraakman [2000]'s point, and it is what the apex row is measuring on a scale the instrument does not carry.

7.7 A metric that improves as the structure fails

A pledge over the largest bloc's holdings is modelled with an eight-point migration of exposure to the lender's perimeter on enforcement:

```
before   naive 0.189002   eff jur 5.290949   bloc-HHI 0.340000   N 2.941176   h -0.122417
after    naive 0.181482   eff jur 5.510188   bloc-HHI 0.336800   N 2.969121   h -0.119107
change in h                                                                    +0.003310
```

The instrument reports an improvement of 0.003310 for an event that cost the principal majority control of the operating company. It reports it because the largest bloc was at 40.0% and the forced migration evened the shares. The measurement is correct and the conclusion is absurd. **A metric that improves as the structure fails is the most dangerous output any instrument can produce, and it must be read alongside the event rather than instead of it.** Curzi [2026] is the primary treatment of the pledge mechanics, where a constructed first-passage model at $\sigma = 40\%$ and $\mu = 6\%$ puts the ten-year probability of touching a 36% control-loss barrier at 76.347%, with a median time to the 23.076923% barrier of 0.8786 years; this paper's exposure migration derives from those triggers. Goodhart [1975] and Campbell [1979] are the general statement of the same pathology.

8. Results: H4 — an aggregate deeming rule is not an economic test

8.1 The rule, stated as a rule

The rule analysed here is stipulated in the following form, which is how the source manuscript states it and which this paper does **not** cite to any regime, because no regime's guidance is in the bibliography:

> **Deeming rule (stipulated).** An entity in which one or more designated persons hold, in the aggregate, fifty per cent or more of the ownership interest — directly or indirectly — is itself treated as designated, whether or not it is separately listed. The indirect limb does not terminate at a fixed depth; it terminates at natural persons or at an unresolvable node.

The rule deems; it does not multiply. That single property is the whole result, and it is a property of the rule as a mathematical object rather than of any jurisdiction's version of it.

8.2 The divergence, computed

Iterate a chain of 51% holdings. Economic interest is the product; deemed status is the conjunction of a threshold test applied at each tier.

tier	direct holding	cumulative economic interest	deemed?
1	51.00%	51.000000%	yes
2	51.00%	26.010000%	yes
3	51.00%	13.265100%	yes
4	51.00%	6.765201%	yes
5	51.00%	3.450253%	yes
6	51.00%	1.759629%	yes
7	51.00%	0.897411%	yes

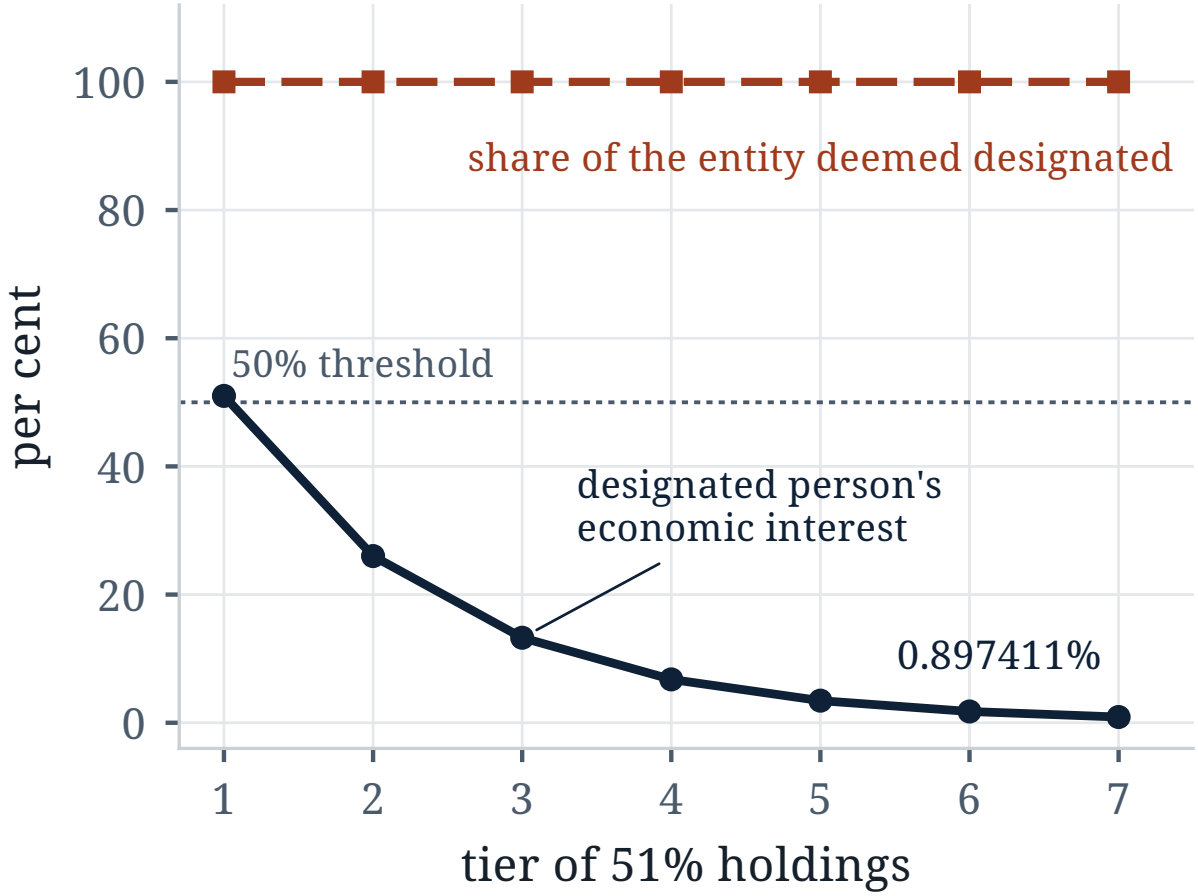


Figure 63.4: **A deeming rule and an economic-interest test are different functions and they diverge without bound.** Seven tiers of 51 per cent holdings drive cumulative economic interest from 51.000000 per cent to 0.897411 per cent while the deemed status is unchanged at every tier; the horizontal line marks the 50 per cent threshold that the economic reading crosses at tier two and the rule never does. Computed from the stipulated rule of Section 8.1; script `figs/src/CCP-063_fig04.py`.

At the seventh tier the designated person's economic interest is 0.897411% and the entity is deemed in full. A structure that computes economic interests and concludes it is comfortable has computed a quantity the rule does not use. The divergence is unbounded: economic interest decays as $0.51^{k} \to 0$ while the deemed status is invariant in k.

The aggregation limb produces a discontinuity in the other direction:

holding pattern	economic interest	deemed?
one designated person at 49.00%	49.00%	no
two designated persons at 24.50% each	49.00%	no
two designated persons at 25.00% each	50.00%	**yes**
one at 49.00% plus one at 1.00%	50.00%	**yes**
three designated persons at 17.00% each	51.00%	**yes**

One percentage point of economics separates the second row from the third and the answers are opposite. H4 is supported.

8.3 What follows for the resolution task

Two consequences follow directly and neither is a legal claim.

Ownership resolution is a search with an unbounded depth, not a lookup. Because the indirect limb terminates at natural persons rather than at a fixed tier, the correct depth is "until the tree terminates". At a branching factor of three across eleven corporate counterparties the node count is 44, 143, 440 and 1,331 at depths one to four. **An unresolvable node is a finding rather than an omission**, and a structure that acquires forty suppliers has acquired forty ownership trees.

Three regimes with different tests give different answers to one shareholding. A forty-five per cent holding by a designated person, carrying the right to appoint the board, is not caught on any of the three ownership arithmetics and is caught or contested on two of the three control tests. **A structure designed against one of them has designed against one of them.** The specific regimes, their tests and the appellate authority on the breadth of a control limb are described in Appendix B and are **not** cited, because none of them is in the bibliography. That different bodies of law define control differently for different purposes is what Pistor [2019] documents as the general mechanism, and it is the only part of this paragraph that carries a citation.

Where the paper stops. No claim is made here about any regime's actual rule, its current text, its scope, or any entity's status under it. **The result is that a rule of the stipulated form is not an economic test.** Whether any particular regime states a rule of that form is a question for Appendix B and for counsel.

9. Results: H5 — the file, the threshold, and the exposure window

9.1 The exit hazard, and why the loss is not proportional

A structure that cannot answer an assessment question is withdrawn from rather than adjudicated against: there is no hearing, no finding and frequently no explanation. Under a constructed annual exit hazard on five relationships:

annual exit hazard	expected at 1 yr	at 3 yr	at 5 yr
5%	4.7500	4.2869	3.8689
10%	4.5000	3.6450	2.9525
15%	4.2500	3.0706	2.2185
20%	4.0000	2.5600	1.6384

At a twenty per cent hazard the expected number of surviving relationships after five years is 1.6384. **And the loss is not proportional, because the relationships that go first are the ones doing the most work.** Losing the two largest counterparties takes the counterparty index from 0.261400 to 0.360000 and the effective count from 3.825555 to 2.777778 — a fall of 1.047777, **more than the entire counterparty diversification the structure had.**

The mechanism is an assessment-cost decision about a *category*, not a risk assessment of a *client*: an institution declines a category when the cost of assessing its members exceeds the revenue from them, adjusted for the tail. **The only lever a structure has is its own assessment cost**, and every component of that is an architectural choice made years before the assessment. That withdrawal from whole corridors is an unintended consequence of a risk-based approach is recognised by the standard-setting body itself [Financial Action Task Force, 2012], and the listing machinery through which it operates is documented as a technique of governance [de Goede, 2011].

9.2 H5 — consistency across a file: supported

Across n independently drafted documents at a per-document error rate p, the probability that at least one answer contradicts another is $1-(1-p)^n$.

n	$p=0.01$	0.02	0.05	0.10
5	4.901%	9.608%	22.622%	40.951%
10	9.562%	18.293%	40.126%	65.132%
20	18.209%	33.239%	64.151%	87.842%
30	26.030%	45.452%	78.536%	95.761%
45	36.381%	59.712%	**90.056%**	99.127%

A full onboarding pack across an eight-jurisdiction group runs to roughly forty-five documents. **At a five per cent per-document error rate the probability that the pack contradicts itself somewhere is 90.056%. At 0.2% it is 8.615% — a factor of 10.45.**

The way to move p is not care; it is to single-source the facts and generate the documents, so that there is one recorded value for each fact and every document is a view of it. **H5 is supported: the failure probability is a property of the generation process, and it falls by an order of magnitude without anybody being more careful.** This is the same defect Curzi [2026] measures at the character layer, where typed-field parsing collapses the withdrawn-code parsing-ambiguity false-positive share from 88.99% to 0.00% against a naive detection precision of 11.01% — a control right one time in nine — and it has the same remedy: fix the representation, not the reviewer.

The independence assumption is the model's weakest input and is stated as such. Documents drafted by the same adviser from the same instructions carry correlated errors, which reduces the count of *distinct* contradictions while raising the chance that a single wrong fact propagates through the whole pack. **Neither direction is measured here**, and the qualitative conclusion — single-sourcing dominates review effort — survives both because it removes the error at its origin rather than reducing its probability.

9.3 The screening threshold is a staffing decision

Name matching has a tunable threshold that sets both the false-positive rate and the miss rate. On a constructed population of two thousand payments a day:

threshold	FP rate	alerts/day	alerts/yr (250 d)	analyst-hours/yr at 12 min
exact match	0.000%	0.00	0.0	0.0
0.95 (tight)	0.200%	4.00	1,000.0	200.0
0.90	0.600%	12.00	3,000.0	600.0
0.85	1.500%	30.00	7,500.0	1,500.0
0.80	3.400%	68.00	17,000.0	3,400.0
0.75 (loose)	7.100%	142.00	35,500.0	7,100.0

The alert budget is a staffing decision that determines a detection rate, and almost nobody frames it that way. A structure that cannot fund fifteen hundred analyst-hours a year will run at 0.90 or tighter and inherits whatever miss rate that implies.

And the tightest setting is the most exposed, which is counter-intuitive and is the point. Exact-match comparison drives false positives to zero and is the configuration most completely defeated by a substitution that changes the bytes without changing the rendering. On constructed names, counting only characters with a visually identical form in another script:

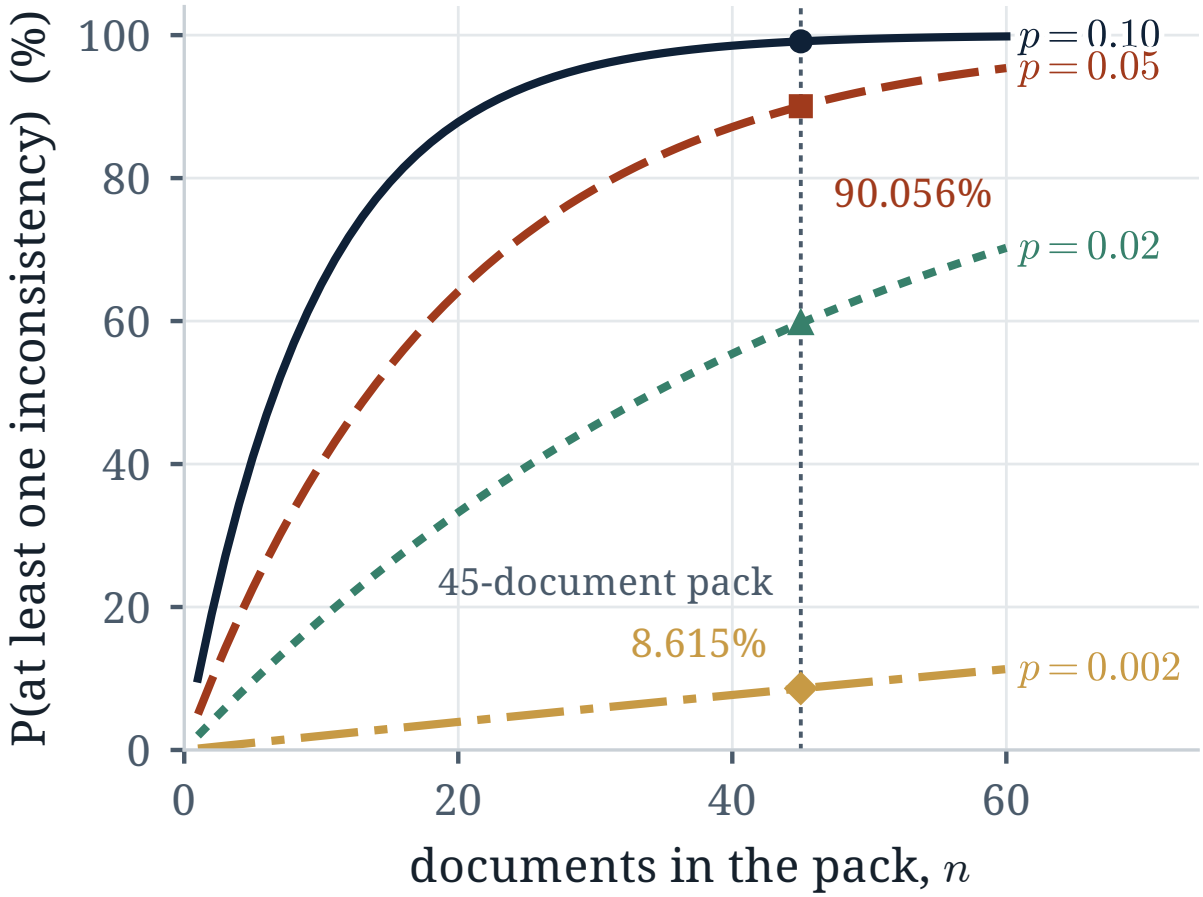


Figure 63.5: **File consistency is a property of the generation process, not of review effort.** $P(\text{at least one inconsistency}) = 1 - (1 - p)^{n}$ across independently drafted documents: at forty-five documents the failure probability is 90.056 per cent at a five per cent per-document error rate and 8.615 per cent at 0.2 per cent. Modelled with independent per-document errors at the stated rates; script `figs/src/CCP-063_fig05.py`.

string	substitutable characters	byte-distinct variants
a 14-character name	9	512
a 16-character name	11	2,048
a 20-character name	11	2,048
a 16-character name	11	2,048
a 17-character name	12	**4,096**

Covering one seventeen-character name by enumeration takes $2^{12} = 4{,}096$ list entries; after canonicalising both sides before comparison it takes one. A looser threshold does not close the class, because the variants are not *near* matches to a fuzzy engine — after the byte-level difference they are simply different strings that happen to look the same. Curzi [2026] specifies the canonical form and the raw-byte witness that detects the substitution rather than normalising it away, and Curzi [2026] states why the witness must be computed over the raw octets and kept on a separate plane. **This paper's contribution is to establish that the string comparison is load-bearing: every assessment answer a structure gives rests on the claim that the comparison worked.**

9.4 Transit exposure is linear in the window

A rail is not a holding, and freeze exposure on a transit asset is a window problem. Expected exposure per unit transferred, under a constructed annual hazard:

hops	hours/hop	window	hazard 0.10%	0.50%	2.00%
1	4	4 h	0.00000457	0.00002283	0.00009132
2	26	52 h	0.00005936	0.00029680	0.00118721
3	52	156 h	0.00017808	0.00089041	0.00356164
4	96	384 h	0.00043836	0.00219178	0.00876712

A four-hop route carries ninety-six times the expected exposure of a one-hop route, and halving the window halves the exposure exactly. That is the whole design lever, and it is a routing and pre-clearance decision rather than a legal one — a hop that cannot be pre-cleared is a hop that waits. Curzi [2026] owns the rail measurement this derives from.

10. Results: H6 — the disclosure test in historical instance

10.1 The constraint, and the instrument shaped around it

The mature bill of exchange is a four-role, two-place, two-currency, one-interval contract, and its shape is the negative impression of a constraint. The prohibited object was not interest but a specific juridical category: anything received beyond the principal by virtue of a loan of a fungible thing, where ownership passes and an equivalent quantity is returned. What was *not* prohibited defines the design space — a partnership share bearing venture risk, compensation for loss actually suffered by reason of delay, compensation for profit forgone, a penalty agreed in advance for late payment, payment for genuinely assumed risk of the principal, and a genuine exchange of one currency for another, which is a sale and not a loan.

The distinction between a loan of the same thing back and an exchange of one thing for a different thing — and not any doctrine about rates — is the hinge. The canonical texts, their canon numbers and the scholastic treatments of the exchange contract are described in Appendix B and are **not cited**, because the bibliography contains no canon-law sources. de Roover [1963] is the modern scholarship and is cited for the institutional account.

Four things happen at once in the instrument, and it matters that the document does not separate them: value moves between two places without metal moving; credit is extended for the length of the customary interval; currency is converted, twice if there is a return leg; and the transaction is on its face a purchase of foreign currency. **No rate of interest can be stated, because the contract contains two exchange rates and one of them is not yet known.** The outward rate is quoted when the bill is written; the return rate is whatever the foreign market quotes at the end of the interval. **The profit is therefore uncertain at the moment of contracting, and that uncertainty is the property the characterisation tracks.**

10.2 The arithmetic, and why it is the same number

The rates below are stipulated for illustration and are not taken from any archival series.

```
STIPULATED round trip

pay out                          1,000.0000 units of the home money
outward rate                     50.0000 foreign units per home unit
  -> delivered abroad          50,000.0000 foreign units
return rate at the interval     48.0000 foreign units per home unit
  -> brought home              1,041.666667 home units

gross gain                           41.666667  =  4.166667% per round trip
```

Annualised over plausible intervals:

```
round trip  60 days  ->  28.188775%   per annum
round trip  90 days  ->  18.004889%
round trip  91 days  ->  17.790398%
round trip 120 days  ->  13.220482%
round trip 180 days  ->   8.630055%
```

Take the 120-day case. The compounded annual rate is 13.220482%; the equivalent simple annual rate is $4.166667\% \times 365/120 = 12.673611\%$. A plain interest-bearing loan of $1{,}000$ for 120 days at 12.673611% simple yields 41.666667. The exchange differential yields 41.666667. **The difference between the prohibited instrument and the permitted one, computed to the ninth decimal place, is 0.000000000.**

That is a specification result rather than a scandal. **Two instruments with identical economics and different juridical characterisations are different instruments in every way a legal system can see, and the entire remaining question is whether the characterisation survives examination.**

The uncertainty is real, which is why the characterisation held:

return rate	profit on 1,000	as %
44.0	$+136.3636$	$+13.6364\%$
46.0	$+86.9565$	$+8.6957\%$
47.0	$+63.8298$	$+6.3830\%$
48.0	$+41.6667$	$+4.1667\%$
49.0	$+20.4082$	$+2.0408\%$
50.0	$+0.0000$	0.0000%
51.0	-19.6078	-1.9608%
52.0	-38.4615	-3.8462%
54.0	-74.0741	-7.4074%

The break-even return rate equals the outward rate. Moving the return quotation from 47 to 48 costs 22.163121 on a 1,000 principal, which is 2.216312%; **the entire** 4.166667% **differential is erased by a two-unit move, which is** 4.1667% **of the return rate and well inside the observed range of the period's quotations.** The writer of this bill could genuinely lose money.

Repeat the trip and compounding does the rest: 4.166667%, 8.506944%, 13.028067%, 17.737570% and 27.753440% per annum at one, two, three, four and six round trips a year.

10.3 Dry exchange, and H6

The abuse has its own contemporary name, and naming it matters because it is the exact point at which the design stops working. In a dry exchange the return leg is not real: the bill is drawn on a place with no correspondent, or on a correspondent instructed to protest it, or the return rate is fixed in advance by side agreement. The trade is then a term loan at a rate agreed at the outset, wearing an exchange's clothes. **The uncertainty is gone, and with the uncertainty gone so is the characterisation.**

Now apply the disclosure test — would the arrangement still do its job if every element were disclosed to every authority with jurisdiction over any part of it?

The bill passes it, and it passes in the strongest available form. The bill was written, entered in the ledgers of both houses, frequently notarised, and its outward rate was quoted publicly on both markets. The bank's principal client was the institution whose canonists had elaborated the prohibition; that client held the bank's own accounts, appointed it to the office of its treasury's depositary, and could inspect any of it at will. **There was no element the relevant authority could not see, and the authority with the most direct jurisdiction over the characterisation was on the other side of the table reading the ledger, and the instrument continued to work.**

Dry exchange fails the same test by construction. It requires that the return leg is not examined. Ask who the foreign correspondent is, whether the bill was in fact presented, whether it was protested by arrangement, and what the side agreement says, and the characterisation collapses into the prohibited thing. **Dry exchange is the version that was condemned, and it was condemned because the property that made the licit version licit had been removed while the form was retained.**

H6 is supported. Two instruments, near-identical on the page, one passing and one failing, and the historical record selects between them exactly as the test predicts. **This is the disclosure constraint of Curzi [2026] — effectiveness independent of whether anyone knows the design is in place, so that every element is publishable and non-disclosure is refused as a design premise — with a five-century-old worked instance attached.** Kerckhoffs [1883] is the principle's canonical statement and Saltzer and Schroeder [1975] its restatement as open design; Brandeis [1914] is the older policy instinct from the transparency side. The modern regulatory environment has since removed non-disclosure from the set of available design parameters altogether [Organisation for Economic Co-operation and Development, 2017, Financial Action Task Force, 2012, Organisation for Economic Co-operation and Development, 2015].

The qualification belongs in print. Passing the disclosure test and passing the arithmetic test are two different tests. The assembled structure of Section 5 passes the first and fails the second at the tested parameters. **A structure can be entirely legitimate, entirely disclosed, entirely defensible, and still not generate enough surplus to pay for itself.**

10.4 The instrument pointed at the client dimension

The same concentration instrument, applied to the historical network. **The share vector is CONSTRUCTED for illustration; it is not de Roover [1963]'s, and readers who want archival figures should take them from him.** It is shaped only to the one thing the scholarship is clear

about — that the branch following the client's court dominated profitability — and it is used to demonstrate a mechanism rather than to report a fact.

```
CONSTRUCTED branch profit shares
42% · 14% · 10% · 9% · 9% · 7% · 6% · 3%   (eight places)

naive HHI            = 0.231600
effective branches   = 4.317789      (of 8 nominal)
```

Eight places, 4.32 effective — already a substantial haircut and already better information than "eight branches". But the branches are not independent in the dimension that matters, because a material part of the transfer traffic through several of them consisted of remittance heading for the dominant client. Collapse to client blocs:

client-derived share	client-bloc HHI	effective clients	h
42.0%	0.512800	1.950078	-0.266246
48.0%	0.500800	1.996805	-0.257958
55.0%	0.505000	1.980198	-0.260881
60.0%	0.520000	1.923077	-0.271126
65.0%	0.545000	1.834862	-0.287561

At the lowest reading the network has 1.950078 effective clients against 4.317789 effective branches: the branch count overstates diversification by a factor of 2.214162. The h column is computed at this paper's standard test parameters for comparability only; a fifteenth-century partnership faced none of the framework's erosion terms, and every value is far below the break-even $N^{*} = 4.172734$.

The error is smaller than the modern one. Curzi [2026]'s conventional four-jurisdiction structure reads 3.407968 naively and 1.000000 on blocs, an overstatement factor of 3.41. The historical network was better diversified than that structure, and neither was diversified.

10.5 The loan book, and the failure it predicted

CONSTRUCTED loan-book shares, shaped to the historical account and not drawn from any archival series:

```
sovereign A          38%
sovereign B          27%
a group of princes   15%
commercial bills     20%

naive HHI            = 0.279800
effective borrowers  = 3.573981        h = −0.054212
```

3.57 effective borrowers reads as a moderately diversified book. The right bloc is not geography. The two sovereigns were allies against a common adversary, funding the same war from opposite sides of one channel; their solvency depended on the same outcome. **They are not two exposures. They are one exposure with two names.**

```
collapsed to a single war-finance bloc: 65.0%

bloc HHI             = 0.545000
effective blocs      = 1.834862        h = −0.287561
```

The naive borrower count reads 3.573981 **and the bloc count** 1.834862 **— an overstatement of** 1.739119**, or** 94.782%**.** The books would have shown a portfolio spread across two kingdoms, several states and a body of ordinary commercial paper. The position was that nearly two-thirds of the lending was a single directional bet on one side of a continental war. **The instrument's answer was available in advance and nobody was computing it.**

10.6 What compartmentalisation bought, and what it did not

Each place of business was a separate partnership with its own deed, capital and books, the centre holding a controlling but not total interest and the local manager holding real minority equity. That is asset partitioning by separate legal personality [Hansmann and Kraakman, 2000] four centuries before the special-purpose vehicle, and it is routinely presented as a modern innovation.

The magnitude is not intuitive. Take eight units each with independent failure probability p in a period:

p	compartmented: expected units lost	consolidated: $P(\text{whole fails})$
0.05	0.4000 of 8 (5.00% of capital)	0.336580
0.10	0.8000 of 8 (10.00% of capital)	0.569533
0.15	1.2000 of 8 (15.00% of capital)	0.727509
0.20	1.6000 of 8 (20.00% of capital)	0.832228

Add a common systemic factor q striking all units at once and the picture sharpens rather than blurs:

q	consolidated $P(\text{total loss})$	compartmented $P(\text{total loss})$	expected capital loss
0.02	0.578142	0.020000	0.118000
0.05	0.591056	0.050000	0.145000
0.10	0.612580	0.100000	0.190000

At $q = 0.05$ and $p = 0.10$ the consolidated structure's probability of total loss is 0.591056 against the compartmented structure's 0.050000 — **a factor of** 11.821123**. The systemic factor is the only thing that can take the compartmented network down, which is precisely the point of building it that way: compartmentalisation converts idiosyncratic risk into a capital loss and leaves only the correlated risk as existential.**

And it did not save the institution, because it contains legal contagion and not economic or reputational contagion. Three leaks defeat it and all three are decisions by the centre rather than events imposed on it: a parent guarantee that breaches the compartment voluntarily from the inside; a common trading name, so that a failure in one place impairs funding in another through a channel no deed can seal; and a common apex, the same handful of principals holding the majority everywhere — which *is* the systemic factor q. **Compartmentalisation is a defence against idiosyncratic failure and not against correlated failure, and correlated failure is the one the structure is actually exposed to.**

The consequence for the modern structure is a budgeting rule and it is uncomfortable. **A single apex by design makes the correlated term the only term that can end the structure, so the remaining risk budget should be aimed at** q **and not at** p**.** Money spent lowering the failure probability of an individual leg is money spent on the term compartmentalisation already handles.

That the accounting method of the period is often credited with more causal force than it can carry is worth naming while the historical instance is open: Yamey [1964] and Littleton

and Yamey [1956] are the deflation of the claim that double-entry produced rational capitalism, and this paper makes no such claim. Curzi [2026] measures the completeness gap of the method directly.

11. Failure modes of the assembled measurement

11.1 Reading the flattering axis. The jurisdictional reading gives $h = -0.122417$; the rail reading gives -0.406608. Both are correct measurements of different things and only one is the constraint. **Print both, or the choice has been made silently and it will not be the unflattering one.**

11.2 Believing the diversification paid for itself. It closed 75.5167% of the shortfall and did not close it. The ceiling rose; the bill did not move.

11.3 Attacking the cost side by cutting substance. The three largest lines in C are substance lines, and substance is what makes the entities real. A cost programme that starts there dismantles the structure and then reports the saving.

11.4 Ignoring η. The most powerful lever per unit of itself, and the one nobody works on because it cannot be purchased from an adviser.

11.5 Arguing about bloc calls instead of pricing them. The Eurasian question is 16.0476%, the second-bridge question 36.1137%, the African question exactly 0.0000% and therefore purely a cost question. Every one of them was an argument in somebody's file before it was a number.

11.6 Confusing classification risk with tail risk. p^* prices the first only. A structure that reads a low break-even probability as a recommendation has priced one uncertainty and ignored the other.

11.7 Buying a jurisdiction that concentrates a bloc already held. The exposure is drawn from the right place and lands in the wrong one, and the break-even more than doubles.

11.8 Letting the leverage layer sit outside the exposure map. Section 7.7: the instrument improved by 0.003310 on an event that cost majority control of the operating company.

11.9 Sharing a path between two flows. Ownership and decision rights share the apex deliberately. Anything else sharing a point is a defect that will present as an unrelated failure.

11.10 A bloc that cannot do its job. Thirty per cent of exposure booked to a perimeter whose entities cannot bank, hire or obtain an opinion is thirty per cent of a measurement and none of a structure. **The effective count measures independence and not capacity**, which is Curzi [2026]'s own counter-result and applies here unchanged.

11.11 Measuring once. Perimeter fragmentation is invisible except as a difference between two measurements a year apart.

11.12 Treating a zero-dh element as an underperformer. It answers a different threat and it saves 0.014 of C on the interface budget. Judging every element against N is the error the instrument invites.

11.13 Adding a leg without its cost ceiling. A leg costing more than its gain is negative whether or not the bloc call was right, and the bloc call is the argument everybody has.

11.14 Reporting the naive assignment because it is the only one that clears. Every defensible assignment returns a negative number. A structure whose favourable result requires $\rho = 0$ between coordinating jurisdictions has not found a good answer; it has found the one assumption the framework exists to refute.

11.15 Assembling the chart and calling it the structure. Five flows, not one. An ownership diagram that is correct says nothing about whether the cash, contract, information and decision paths fail independently, and the failure that arrives will be in one of the four the diagram does not show.

11.16 Believing the answer is more jurisdictions. Break-even is $N^* = 4.172734$ and the candidate set considered here does not contain 4.17 genuinely independent perimeters. **The**

remaining distance has to come from C and η. The design side of the problem is finished; the cost side is not.

11.17 Keeping the form and removing the substance. The dry-exchange failure with its own contemporary name. If the property that makes a structure work is removed while the paperwork is retained, the structure does not become a weaker version of itself: it becomes the prohibited thing wearing the permitted thing's clothes, and it fails at the first examination that asks about the property rather than the paperwork. **Every substance requirement in modern practice is an examiner asking about the property rather than the paperwork.**

11.18 Mistaking the characterisation for the economics, and the reverse. The differential and the interest agreed to 0.000000000, so anyone concluding from the characterisation that the bill was not a credit instrument would have mispriced and mis-reserved against it. The commoner modern error is the reverse: because the economics are identical, the distinction is treated as a fiction. **A distinction with no economic content and complete legal content is still completely load-bearing.**

11.19 Decentralising liability without decentralising underwriting. A compartment that contains a decision-maker whose downside is capped at a minority stake while the majority's is not has converted a structural control into an agency problem.

11.20 Political position and commercial position as one asset counted twice. A structure whose licence to operate is a principal's personal standing has a single point of failure that no amount of entity architecture repairs.

11.21 Unpriced counterparty legitimacy risk. The historical case adds a threat category the modern literature almost never prices: the durability of a counterparty's own title. The institution whose account carried the network held its temporal position under a claim whose charter document had been demonstrated a forgery by philological argument roughly forty per cent of the way through the network's life, independently doubted seven years before that, and left unprinted for a further seventy-seven years. **That exposure did not destroy the network — the network was destroyed by ordinary concentration — but it was live for the whole second half of its existence and it was on nobody's register.** The dates, works and attributions are described in Appendix B and are not cited.

12. Discussion

12.1 What the instrument does and does not license

The instrument returns an effective count and nothing else. It says how many things can fail separately; it does not say whether any of them can carry the load, and it cannot see an axis it was not pointed at. **Both limits bite in this paper**: the succession threat is answered by an element scoring exactly zero on N, and the worst exposure in the structure lives on an axis the design programme never measured.

The framework's own founding move is presupposed rather than restated: Curzi [2026] establishes that the present state is a local minimum of the comfort landscape by construction, because comfort is defined as the felt total of present costs — which is why a structure's inertia is the default and why a headroom gate rather than an optimisation is the right form of criterion. Nothing in this paper re-derives that.

The most useful contrast in the corpus for this material is Curzi [2026], and it is a contrast rather than an endorsement. That paper's counter-result refuses category innocence as a defence of acquisition and forbids moving any harm through a gap discovered in a scheme, and it concedes that the case where the feedback loop runs through another person's will is an unsettled seam. **This paper's deeming-rule arithmetic and its dry-exchange analysis both sit directly against that refusal.** Section 8 shows that two designated holders at 24.50% are outside a rule that catches two at 25.00%; Section 10.3 shows that an instrument keeping a permitted form while removing its substantive property is the prohibited thing. **Finding a gap is not a licence to move anything through it, and this paper states the arithmetic of the gap for the**

purpose of designing against it — which is the only use Curzi [2026] leaves available. A reader who takes Section 8 as a design manual has read it against both papers.

12.2 Two answers, and which one to publish

Section 6 leaves the structure with two defensible values of h differing by 0.284191, and Section 5.2 leaves it with seven defensible values of N. The response is not to select one. **It is to publish the family and the assignment rule together**, which is what Appendix A requires and what makes a later reader able to substitute their own call and re-read the answer. A single number without its assignment is not a measurement of anything.

The magnitude of the shortfall in comparable units. Curzi [2026] fixes the conversion between the framework's surplus units: one nat of required surplus is 1.442695 bits, and the octave form is $2(c/\eta)/\ln 2 = 2.885390081777927 \times (c/\eta)$. The assembled structure's shortfall of 0.122417 nats is therefore 0.176610 bits. **The nat-to-bit conversion is exact arithmetic and carries no further content; the octave reading is deliberately not taken here**, because octave equivalence is a perceptual structure for pitch and nothing licenses carrying it to a quantity that is not a pitch. That restraint is the corpus's own and Section 14 records it.

12.3 The ordinal exposure, which is the sharpest objection

The headroom law divides by nothing but it multiplies η by a logarithm, and η is not measured. Curzi [2026]'s ordinal-defeat result is the exact warning and it is reused here as an instrument rather than cited decoratively: with c and η known only ordinally, $d \ln f_{\min}/d(c/\eta) = 2$, so a 0.1 error in the ratio is a 22.14% error in the derived quantity, and an order-preserving relabel of $(1.0, 2.0, 3.0)$ as $(0.5, 2.0, 5.0)$ moves one value by $54.6\times$ with the ordering unchanged. **The same exposure applies here.** h is linear in η, so an order-preserving relabel of the conversion efficiency moves h proportionally: at $\eta = 0.50$ the break-even falls to $e^{1.0} = 2.718$ and the designed structure clears it.

The scalar-reduction habit underneath that exposure is a framework-level limitation and is inherited here rather than restated. Curzi [2026] names it once, at framework level, as one objection appearing in two domains rather than two objections: where the source description is a pattern across several jointly varying features, the framework replaces it with a single number and does arithmetic on the number, and the validity of the arithmetic does not protect the reduction that produced it. **A bloc assignment does exactly that to a jurisdiction's relationships, and a headroom figure does it to a whole structure.** This paper concedes the objection in that form and points at its statement rather than repeating it.

Which is to say: the sign of the paper's headline result depends on a parameter that was stipulated, and a different stipulation reverses it. That is stated here rather than left for a reader to find, and it is the single most important sensitivity in the paper.

The measurement-theoretic status is worse than the sensitivity suggests. Michell [1997] argues that psychology systematically fails to test whether its attributes are quantitative at all before measuring them, and **that is the strongest single objection to this framework.** The formal test is conjoint measurement [Luce and Tukey, 1964], whose conditions establish additive quantitative structure; Krantz et al. [1971] sets out what must be true before a quantity can be said to exist. **Those conditions were not tested for η, for Q, or for the exposure weights, and this paper concedes the point rather than answering it.** What survives without them is structure and not magnitude: the ordering of the bloc assignments, the sign of h, the direction of every comparison, and the exactness of the arithmetic on stipulated inputs.

12.4 What is genuinely borrowed, and from where

The paper borrows a variance identity, a portfolio result, a numbers-equivalent, and a disclosure principle. **Section 13.5 labels each borrowing by category**, and the failure mode that

table exists to prevent is letting a shared mathematical form carry a derivation's authority. Cartwright [1999] supplies the discipline: state the arrangement within which a law is claimed to hold and stop claiming a universality the argument does not need. **The claim here is that a named set of constructed arrangements share a form, not that the form is universal.**

13. Verification and rigour audit

The results are decomposed into postulates that can be checked one at a time. Each is atomic, each carries a status, and each chain names what it consumes. **If dropping a postulate leaves the result standing, it is not listed.**

13.1 Dimensional audit

symbol	quantity	dimension	unit	status
s_i, w_i	exposure shares	$[1]$	fraction of protected capital	defined; sum to 1
HHI	sum of squared shares	$[1]$	fraction (or points $\times 10^4$)	derived
N	effective count	$[\#]$	jurisdictions, blocs, counterparties, rails	derived as HHI^{-1}
$\ln N$	surplus argument	$[1]$	**nats**	derived; dimensionless by construction
η	conversion efficiency	$[\mathrm{idx}\, T^{-1}\, \mathrm{nat}^{-1}]$	index units per period per nat	**assigned by this paper; the source is silent**
Q	stock carried	[idx]	index units, no currency	assumed (stipulated)
δ	decay rate	$[T^{-1}]$	per period	assumed (stipulated)
C	cost flow	$[\mathrm{idx}\, T^{-1}]$	index units per period	assumed (stipulated)
h	headroom	$[\mathrm{idx}\, T^{-1}]$	index units per period	derived
p^*	break-even probability	$[1]$	probability	derived
G, L	gain, loss in headroom	$[\mathrm{idx}\, T^{-1}]$	same as h	derived
hazard, error rate	per-period probabilities	$[T^{-1}]$, $[1]$	per year; per document	assumed (constructed)
exposure window	transit duration	$[T]$	hours	assumed (constructed)
exchange rate	foreign units per home unit	$[F\, H^{-1}]$	stipulated units	assumed (stipulated)

Does the headroom equation balance? Only if η carries index units per period per nat, since $\ln N$ is a pure number and δQ and C are both index units per period. **The source treats η as a bare number; this paper assigns the dimension explicitly rather than silently**, following

Curzi [2026]. The break-even count $N^* = e^{(\delta Q + C)/\eta}$ is unit-free and survives; **the magnitude of h does not mean anything until an index unit is defined, and none is defined here.**

Does the transit calculation balance? Expected exposure is $\lambda \times w$ with λ in per-year and w in years, giving a dimensionless fraction of the amount in transit. $384/8760$ years at $\lambda = 0.02$ is 0.000876712, which is how the table is computed.

Where a quantity has no operational unit, the audit says so. Q, C and h are index units. The cost allocation of Section 7.5 is a decomposition of an index, not a fee schedule. The exchange arithmetic of Section 10.2 is in stipulated units of account and is not any historical currency.

13.2 Significant figures and epistemic status

- $N^* = e^{0.5/0.35} = 4.172734$ **is exact to the digits shown**, being a transcendental function of stipulated rationals; it is quoted to six decimals because the comparison to 2.941176 is made at that precision and both sides are exact arithmetic on stipulated inputs.
- **Every effective count here is exact for its share vector and is not an estimate.** $1/0.340000 = 2.941176\ldots$ and $1/0.765800 = 1.305824\ldots$ are rational; **no uncertainty is quoted because there is no estimate to attach one to.**
- 0.897411% **is exact**: it is 0.51^7 to six figures. 90.056% is exact for $1 - 0.95^{45}$.
- **The parameter set is stipulated, not estimated.** $\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$ were fixed in advance. Every output inherits that status: **exact arithmetic on assumed inputs is not a measurement.**
- **The hazard rates, error rates, alert rates and false-positive rates are illustrative model outputs and are rounded hard.** A twenty per cent annual exit hazard and a five per cent per-document error rate are stipulations; the figures derived from them — 1.6384, 90.056% — are exact functions of stipulations and carry no authority beyond them.
- **One rounding inconsistency in the source is reported rather than harmonised.** Section 13.8.

13.3 Limits and boundary behaviour

quantity	limit	behaviour	reading
N	one share $\to 1$	$N \to 1$	full concentration; the instrument's floor
N	n equal shares	$N \to n$ exactly	the numbers-equivalent property
h	$N \to 1$	$h \to -(\delta Q + C)$	the ceiling vanishes; no η can help
h	$N \to \infty$	$h \to \infty$	**a defect of the model, not a result**; domain is small N
h	$\eta \to 0$	$h \to -(\delta Q + C)$	no conversion, no ceiling, at any count
h	$\delta \to 0$	$h \to \eta \ln N - C$	nothing decays out of the stock: a pathology, not an ideal
p^*	$L \to 0$	$p^* \to 0$	a free option; the decision becomes purely a cost question
p^*	$G \to 0$	$p^* \to 1$	no upside; the leg is refused at any belief

quantity	limit	behaviour	reading
deeming rule	tier $k \to \infty$	economics $\to 0$, status invariant	**the divergence is unbounded; that is the result**
$1-(1-p)^n$	$p \to 0$	$\to np$	the linear regime; single-sourcing operates here
$1-(1-p)^n$	$n \to \infty$	$\to 1$	any positive error rate reaches certainty in enough documents
transit exposure	window $\to 0$	$\to 0$ linearly	exposure is exactly proportional to the window
compartmentation	$q \to 0$	$P(\text{total loss}) \to 0$	with no systemic factor the compartmented network cannot fail entirely
compartmentation	$q \to 1$	both $\to 1$	the compartment buys nothing against a certain common shock

The $\delta \to 0$ row is the corpus's standing example and it is preserved here: zero decay is a pathology because nothing ever leaves the stock, which is Curzi [2026a]'s pole result stated as a boundary.

13.4 Invariance

HHI^{-1} **is invariant under any permutation of the units and under any relabelling of them**, which is exactly why Section 5.3's rearrangement inside a perimeter moves the naive index and leaves N unchanged: the bloc index depends on the partition and not on which member of a bloc holds what. **It is not invariant under a change of share basis**, and the share basis is therefore item 1 of the reporting standard.

The break-even count N^* is invariant under any rescaling of the index unit that multiplies η, δQ and C together, since it depends only on the ratio $(\delta Q + C)/\eta$. **The magnitude of h is not**, which is why only N^* and the sign of h are quoted as findings.

p^* **is invariant under any positive rescaling of** h, since it is a ratio of two headroom differences — which is why the pricing method survives the absence of an operational unit that defeats the magnitude of h. **That is the strongest invariance in the paper and it is the reason Section 7 is the most transportable result in it.**

The equivalence of Section 10.2 is invariant under the choice of principal: the differential and the simple interest agree at any amount, because both are linear in it.

Tested at how many scales? **The invariances are algebraic and were checked at exactly the one parameter set and the share vectors reported.** No empirical scale-freedom is claimed anywhere in this paper.

13.5 Analogy discipline

borrowing	source	category	justification
HHI^{-1} as a numbers-equivalent	Adelman [1969], Hill [1973], Jost [2006]	**1 — derivation**	the premise is that shares are non-negative and sum to one, which holds by construction; the identity is exact
independent reinvention of the same index	Simpson [1949], Laakso and Taagepera [1979]	**1 — derivation**	same arithmetic on the same object in three disciplines
exposure as a portfolio; covariance not count	Markowitz [1952]	**1 — derivation**	the bloc collapse *is* the variance identity; the premises are satisfied by the construction
similarity-sensitive diversity of order 2	Leinster and Cobbold [2012], Rao [1982]	**1 — derivation**	exact where the correlation matrix has unit diagonal and non-negative entries, which is the regime used
correlation rising under stress	Longin and Solnik [2001], Ang and Chen [2002]	**2 — shared form**	measured in equity markets; jurisdictions are not shown to satisfy the same conditions
coupled-network fragility	Buldyrev et al. [2010]	**2 — shared form**	the percolation premises are not shown to hold for jurisdictions or for messaging layers
effective count buys stability	Schindler et al. [2010], Loreau and de Mazancourt [2008]	**2 — shared form**	an ecological measurement; the shape is used, the mechanism is not claimed
pseudo-diversity proved	Douceur [2002]	**2 — shared form**	one party presenting as many is the bloc problem in the identity domain; the impossibility result is not imported as a theorem here
independent-party bound	Lamport et al. [1982]	**2 — shared form**	$3f + 1$ is a theorem about agreement, cited as a prior formalisation of "how many parties are genuinely independent"

borrowing	source	category	justification
open design; disclosure	Kerckhoffs [1883], Saltzer and Schroeder [1975]	**1 — derivation**	the premise — effectiveness depends on the arrangement and not on anyone's knowledge of it — is stated and satisfied by both the bill and the assembled structure
performativity of a published instrument	MacKenzie and Millo [2003], Lucas [1976]	**1 — derivation**	the mechanism is the same: a description relied upon changes the behaviour of those relying on it
a measure adopted as a target	Goodhart [1975], Campbell [1979]	**2 — shared form**	stated for monetary and social indicators; the exposure is the same in kind
organisational tight coupling; redundancy that reduces safety	Perrow [1984], Sagan [1993]	**3 — framing**	qualitative organisational theory, used as framing and as a warning, never as a formal result
concentration from multiplicative growth	Bouchaud and Mézard [2000]	—	cited only to record that a mechanism-side companion to the arithmetic exists; **no tail is asserted and this model is not used**
the modelling programme itself	Epstein [2008], Bruch and Atwell [2015]	—	the first enumerates the legitimate non-predictive uses of a model; the second is the checklist **none of this paper's models meets**

The failure mode this table exists to prevent is letting a category-2 borrowing carry a category-1 result's authority. The numbers-equivalent and the variance identity are category 1 and need no help; the network and ecological results are category 2 and are used for shape.

13.6 Postulates

#	statement	status
P1	exposure shares are non-negative and sum to one	DEFINITION
P2	HHI $= \sum_i s_i^2$, and $N =$ HHI^{-1}	DEFINITION
P3	for n equal shares of $1/n$, HHI $= 1/n$	ARITHMETIC
P4	$\mathrm{Var}(R) = \sum_i \sum_j w_i w_j \rho_{ij}$ for $R = \sum_i w_i X_i$ with unit-variance X_i	ARITHMETIC
P5	units inside one coordination perimeter act on one decision	POSTULATED (H1); a stipulated reference group
P6	$h = \eta \ln N - \delta Q - C$	DEFINITION (modelling choice; Curzi [2026], Curzi [2026])
P7	$\eta = 0.35$, $\delta = 0.02$, $Q = 10$, $C = 0.30$	POSTULATED (fixed in advance)
P8	the deeming rule of Section 8.1, as stated	DEFINITION (stipulated rule form)
P9	document errors are independent across documents	POSTULATED (constructed)
P10	the constructed share vectors, hazard rates and exchange quotations	POSTULATED (constructed)

13.7 The chains

Chain A — the assembled measurement (R1).

step	consumes	statement	status
A1	P2, P10	$\sum s_i^2$ on 29.1/18.2/16.4/13.6/11.8/10.9 is 0.189002	ARITHMETIC
A2	A1, P2	naive effective jurisdictions $= 1/0.189002 = 5.290949$	ARITHMETIC
A3	P5, P10	the six shares group as 40.0 / 30.0 / 30.0	POSTULATED (bloc assignment)
A4	A3, P2	bloc-HHI $= 0.40^2 + 0.30^2 + 0.30^2 = 0.340000$	ARITHMETIC
A5	A4	$N = 1/0.340000 = 2.941176$	ARITHMETIC
A6	A5	$\ln N = 1.078810$	ARITHMETIC
A7	A6, P6, P7	$\eta \ln N = 0.35 \times 1.078810 = 0.377583$	ARITHMETIC
A8	P6, P7	$\delta Q + C = 0.02 \times 10 + 0.30 = 0.500000$	ARITHMETIC
A9	A7, A8	$h = 0.377583 - 0.500000 = -0.122417$	ARITHMETIC
A10	A8, P7	$N^* = e^{0.500000/0.35} = e^{1.428571} = 4.172734$	ARITHMETIC

step	consumes	statement	status
R1	A9, A10	the structure closes $0.377583/0.500000 = 75.5167\%$ of the shortfall and does not close it	ARITHMETIC

Every step is arithmetic except A3, which is the bloc assignment. A3 is the one genuine postulate in the chain, it is the step Section 5.2 varies over seven alternatives, and it is the step a reader should attack first.

Chain B — the pricing method (R2, R3).

step	consumes	statement	status
B1	Chain A	$h_0 = -0.122417$ without the leg	ARITHMETIC
B2	P10	a right call gives $N = 3.174603$, $h = -0.095686$	ARITHMETIC
B3	P10	a wrong call gives $N = 2.898551$, $h = -0.127526$	ARITHMETIC
B4	B1, B2	$G = +0.026731$	ARITHMETIC
B5	B1, B3	$L = 0.005110$	ARITHMETIC
B6	B4, B5	$p^* = L/(G + L) = 16.0476\%$, computed at full precision from $L = 0.0051096$ and $G + L = 0.0318401$	ARITHMETIC
B7	B4, P7	cost ceiling $= G/C = 0.026731/0.30 = 8.9102\%$	ARITHMETIC
R3	P10	the full China-bloc leg gives $G = +0.072279$; capped transit at 0.35% gives $+0.002869$, being 3.9692% of the benefit for 3.2110% of the exposure	ARITHMETIC

The whole of Chain B is arithmetic on the same share vector. p^* is a ratio of headroom differences and therefore survives the absence of an index unit.

Chain C — the deeming rule (R4).

step	statement	status
C1	economic interest through k tiers of 51% is 0.51^k	ARITHMETIC
C2	$0.51^7 = 0.00897411 = 0.897411\%$	ARITHMETIC
C3	the rule of P8 applies a threshold at each tier and is satisfied at every tier for $51\% \geq 50\%$	DEFINITION

step	statement	status
C4	$2 \times 24.50\% = 49.00\% < 50\%$; $2 \times 25.00\% = 50.00\% \geq 50\%$	ARITHMETIC
R4	C2, C3, C4	the deemed status is invariant in k while economic interest $\to 0$, and one percentage point of aggregate reverses the answer

R4 rests on one definition and four lines of arithmetic and on no citation whatever, which is deliberate: **it is a result about a rule of the stated form, not about any regime's rule.**

Chain D — the compliance file (R5).

step	statement	status
D1	with independent errors, $P(\text{none}) = (1-p)^n$	DEFINITION, consumes P9
D2	$1 - 0.95^{45} = 0.90056$	ARITHMETIC
D3	$1 - 0.998^{45} = 0.08615$	ARITHMETIC
R5	D2, D3	the ratio is 10.45, achieved by changing p and not n

P9 is the assumption to attack. Correlated drafting errors change the count of distinct contradictions in one direction and the propagation of a single wrong fact in the other, and neither is measured here.

Chain E — the historical equivalence (R6).

step	statement	status
E1	$1{,}000 \times 50 = 50{,}000$ delivered abroad	ARITHMETIC
E2	$50{,}000/48 = 1{,}041.666667$ brought home	ARITHMETIC
E3	gain $= 41.666667$, being 4.166667%	ARITHMETIC
E4	$(1.04166667)^{365/120} - 1 = 13.220482\%$	ARITHMETIC
E5	$4.166667\% \times 365/120 = 12.673611\%$ simple	ARITHMETIC
E6	$1{,}000 \times 0.12673611 \times 120/365 = 41.666667$	ARITHMETIC
R6a	E3, E6	the differential and the simple interest agree to 0.000000000
R6b	—	the licit instrument was disclosed to the promulgating authority and worked; the abusive twin requires its return leg to go unexamined

R6a is arithmetic and R6b is not. The equivalence is checkable in three lines; the historical claim rests on a description this paper does not cite to a primary source, and Appendix B lists what would have to be verified. **A reader who rejects R6b keeps R6a intact.**

The discipline of verifying an identity over the complete set rather than on examples is Curzi [2026]'s, which checks its doubling theorem $HD_{\text{ess}}(i,j) = 2\log_2(pq)$ exactly over all sixty-six pairs at $N = 12$. **The chains above are checked at the one parameter set reported, which is weaker, and Section 14 records it.**

13.8 Discrepancies found in the source, printed

One rounding inconsistency was found and is reported rather than harmonised. The source states the designed structure's naive index as 0.1889 at one point in the historical section and as 0.189002 everywhere else. 0.189002 rounds to 0.1890 and not to 0.1889, so the single occurrence is a transcription error rather than a different rounding convention. **This paper uses** 0.189002 **throughout, which is the value the explicit double sum on** $0.291^2 + 0.182^2 + 0.164^2 + 0.136^2 + 0.118^2 + 0.109^2$ **returns.**

A second observation, not an error. The source reports the break-even probability as 16.0476% while the six-decimal gain and loss printed alongside it, 0.026731 and 0.005110, return 16.0485% when divided as printed. **The discrepancy is a rounding artefact of the printed inputs, not of the result**: the unrounded values are 0.0267305 and 0.0051096, and $0.0051096/0.0318401 = 16.0476\%$, which agrees with Curzi [2026] exactly. **This paper computes** p^* **at full precision and prints the rounded inputs, and records here that a reader redoing the division on the printed figures will land nine ten-thousandths of a percentage point away.**

13.9 Cross-checks performed

- The naive index was recomputed by explicit double sum on all four axes and reproduced 0.189002, 0.261400, 0.268400 and 0.425000.
- The bloc index was recomputed as $0.40^2 + 0.30^2 + 0.30^2$ and reproduced 0.340000 exactly; the messaging collapse as $0.87^2 + 0.05^2 + 0.08^2$ and reproduced 0.765800 exactly.
- The bound $s_{\max}^2 \le \text{HHI} \le s_{\max}$ holds on every vector used: $0.084681 \le 0.189002 \le 0.291$ on the jurisdictional axis and $0.7569 \le 0.765800 \le 0.87$ on the collapsed rail axis.
- The effective-count ceiling $s_{\max}^{-2}$ is respected everywhere: $5.290949 \le 11.8090$ and $1.305824 \le 1.3215$.
- Both historical vectors were re-summed: 0.231600 on eight branch shares and 0.279800 on four loan-book shares.
- 0.51^7 was recomputed to twelve places and rounds to 0.897411%.
- The interest equivalence was verified in both directions at the 120-day interval and agrees to nine decimal places on a principal of $1{,}000$.

14. Limitations

Every worked result is a constructed model and none is a measurement. The entities, the share vectors on all four axes, the cost allocation, the hazard and error rates, the screening population, the historical branch and loan-book vectors and the exchange quotations are stipulated. **No dataset was collected. No real structure, institution, bank, name, holding or transaction appears anywhere in this paper, and no allegation is made about any of them.**

No as-of date, because no real holdings were measured. A concentration figure computed from real holdings requires its source and its date; **a concentration measure without an as-of date is not a measurement.** Every share vector here is constructed, so there is no source and no date to give, and the reporting standard in Appendix A therefore makes the date an item.

The correlation inside a bloc is assumed, not estimated. $\rho = 1$ within a perimeter is the conservative end of a family, and no correlation was estimated from data anywhere in this paper. **That is the largest unmeasured input and it is exactly the quantity H1 turns on.**

Every parameter in the headroom law was chosen rather than measured, and the sign of the headline result depends on one of them. At $\eta = 0.50$ the break-even falls to 2.718 and

the designed structure clears it. **The conclusion that the architecture fails its threshold is therefore conditional on $\eta = 0.35$ and is not robust to a different stipulation.** It is stated here rather than left to be discovered.

The bloc assignment is a stipulated reference group and not an identified effect. Manski [1993] rules out recovering an endogenous social effect from observation of a linear-in-means reflexive system. The remedy adopted is stipulation with a defence [Brock and Durlauf, 2001]; randomisation, an instrument and partial-population variation were unavailable. **The price is that every bloc-level result here is a modelling position.**

Jurisdiction-specific legal claims are not cited to primary sources. The descriptions of sanctions regimes and their ownership and control tests, of automatic-exchange and beneficial-ownership obligations, of free-zone and promotion statutes, of the security legislation and judicial resignations underlying the tail refusal in Section 7.3, and of the canon-law prohibition and its extrinsic titles in Section 10, are all taken from the source manuscript. **The bibliography contains no statutes, regulations, agency guidance, court rules, registry documents or case reports for any individual jurisdiction, and no canon-law sources.** Those claims are described rather than asserted, are uncited, and require verification against the official sources. Appendix B lists them individually. **This is a real gap and it is not closed here.**

The historical claims are described and are supported only at the institutional level. de Roover [1963] is the modern scholarship and is cited for the institutional account; **no archival figure is asserted anywhere in this paper and every historical share vector is constructed and labelled.** The dating and attribution of the philological demonstration in Section 11.21, the identity of the branch managers, the terms on which the depositary office was held, and the legal form of the branch partnerships are all in Appendix B.

The effective count measures independence and not capacity. $N = 2.94$ says nothing about whether any bloc can carry the load, and a structure whose bloc diversity is real while its transacting capacity sits in one bloc has an effective count near 1.00 on the ability to act. This is Curzi [2026]'s counter-result and it is inherited unchanged.

The instrument is least informative where the headline case sits. Golosov [2010] and Dunleavy and Boucek [2003] document that effective-number measures behave poorly in the dominated-unit regime, and the conventional structure's bloc vector is a single unit at one hundred per cent — the limiting case. The measurement is still exact; its discriminating power there is at its minimum.

Small-sample bias is not addressed. Chao et al. [2014] show diversity estimators are biased on small samples, and a six-element exposure vector is a small sample. **No correction was applied, and the paper does not resolve whether these vectors are samples or complete censuses** — a reviewer would be right to press on it. Hurlbert [1971] is the standing reason the word "diversity" must always travel with its formula attached.

14.1 Rigour checks not met, listed by name

1. **No uncertainty is reported anywhere, because nothing is estimated.** Every figure is exact arithmetic on stipulated inputs. That is not precision; it is the absence of empirical content, and it means the paper cannot be wrong in the way a measurement can be wrong.
2. **Model comparison was not performed.** No alternative index was fitted, no out-of-sample test was run, no information criterion was computed and no matched null was constructed, because nothing here is fitted. The instrument is asserted to be the right one on the strength of the identity chain in Section 13.7, not on comparative fit.
3. **The dimension of η is supplied by this paper, not by the source.** Section 13.1 assigns it explicitly rather than silently.
4. **Q, C and h have no operational unit.** Only N^{*}, p^{*}, the ratios and the signs survive.
5. **The surplus law's large-N limit is a defect, not a result.** The domain is stated in Section 13.3 rather than discovered by a reader.

6. **The functional form of the surplus term is underdetermined by anything in this paper.** A sub-linear power law [Latané, 1981] and an $n \log n$ form [Briscoe et al., 2006] fit the same facts; Peters and Gell-Mann [2016] supplies the only warrant used, and **nothing measured here discriminates among the three.**
7. **The conjoint conditions were not tested** for η, Q or the exposure weights, so the ratio-scale treatment is provisional [Michell, 1997, Luce and Tukey, 1964, Krantz et al., 1971]. The orderings and signs survive because they are algebraic; the magnitudes do not.
8. **No distributional or scaling claim is made, so the Clauset–Shalizi–Newman procedure was not run** [Clauset et al., 2009]. The exposure vectors are finite sets of three to eight shares, not samples from a distribution; no tail is characterised and no power law is asserted anywhere. **Any restatement of this paper that reaches for a heavy-tail or power-law claim about holdings, branch profits, loan books or share distributions must run that procedure first**, and must read Stumpf and Porter [2012] on why passing it would still not identify a mechanism.
9. **The independence assumption in Section 9.2 is the weakest input in the paper** and its direction of bias is not measured.
10. **The instrument is exposed to its own adoption.** Goodhart [1975] and Campbell [1979] state what happens to a measure used as a target; nothing in the arithmetic prevents it, and the only defence offered is the reporting standard of Appendix A.
11. **The instrument is performative and its own parameters are exempted from that.** MacKenzie and Millo [2003] documents a model that made its assumptions true after adoption and Lucas [1976] states the consequence. **The framework asserts that systems change when modelled and then exempts its own η, δ and correlation assumptions from that reflexivity. That is an inconsistency and it is not resolved here.**
12. **Aggregation is unaddressed.** One Q, one δ and one C are written for a heterogeneous structure, and Kirman [1992] establishes that such a representative parameter generally does not exist.
13. **Expected value and time average are not the same thing and this paper reports the first.** Peters [2019] is the reason: the expected surviving-relationship counts of Section 9.1 are ensemble averages that no single structure experiences, and a structure that loses its last relationship does not return to the mean.
14. **There is no conservation law for Q.** The one genuine conservation is that shares sum to one; the stock variable has none [Mirowski, 1989]. **This is the objection the genre most often fails, and this paper meets half of it.**
15. **No differential prediction is offered.** Durlauf [2005] asks for the prediction that only this model makes, and the instrument makes none: any correlation-aware concentration measure would produce the same collapse. **The contribution is the identification and the discipline, not a uniquely discriminating forecast.**
16. **The models fail the calibration-and-validation checklist** of Bruch and Atwell [2015]: none is calibrated against observed data, none is validated out-of-sample, and none has a sensitivity analysis beyond the limits of Section 13.3. They are offered as demonstrations that a stated mechanism suffices to produce a stated shape, which is one of the legitimate uses Epstein [2008] enumerates and is not prediction.
17. **Octave reduction is deliberately not applied.** Curzi [2026]'s conversion is used only for the exact nat-to-bit identity in Section 12.2, because nothing licenses carrying octave equivalence to a quantity that is not a pitch. **The corpus's own restraint is preserved rather than relaxed.**
18. **The chains were checked at one parameter set.** Curzi [2026]'s standard — exhaustive verification over a complete set — was not met here, and the seven-row robustness table of Section 5.2 is the whole of the sensitivity analysis performed on the load-bearing postulate.

15. Conclusion

This paper establishes six results in the form the rest of the corpus cites, and the following statements are the fixed text of those results.

R1 — the assembled measurement. At $\eta = 0.35$, $\delta = 0.02$, $Q = 10$ and $C = 0.30$, giving a maintenance floor of 0.500000, the assembled six-jurisdiction structure measures a bloc index of 0.340000, $N = 2.941176$ effective blocs and $h = 0.35 \ln(2.941176) - 0.500000 = -0.122417$ against a break-even $N^* = e^{0.5/0.35} = 4.172734$. **It closes 75.5167% of the shortfall and does not close it.**

R2 — the break-even probability method. A marginal jurisdiction is priced as $p^* = L/(G + L)$. A Eurasian leg taking five per cent from a Western slot gives $G = +0.026731$ and $L = 0.005110$, so $p^* = 16.0476\%$, with a cost ceiling of 8.9102% of C. **The cost ceiling is the binding test and it is the one nobody runs.**

R3 — the largest gain, refused. The full China-bloc leg at 10.9% yields $+0.072279$, the largest single-move gain available, and is refused on tail exposure. Capped transit at 0.35% yields $+0.002869$ — 3.9692% **of the benefit for** 3.2110% **of the exposure, removing** 96.79% **of the tail.**

R4 — an aggregate deeming rule is not an economic test. Under a rule deeming an entity designated at fifty per cent aggregate direct or indirect ownership, seven tiers of 51% holdings leave a designated person 0.897411% of the economics with the entity deemed in full, while two designated persons at 24.50% each are not caught and two at 25.00% are. **The two functions diverge without bound.**

R5 — compliance-file consistency. $P(\text{at least one inconsistency}) = 1 - (1 - p)^n$ reaches 90.056% across forty-five documents at a five per cent per-document error rate and falls to 8.615% at 0.2%. **The remedy is single-sourcing the facts and generating the documents, not more review.**

R6 — the disclosure test in historical instance. On the stipulated rates the bill returns 4.166667% per round trip, an annualised 13.220482% at 120 days against a 12.673611% simple loan rate yielding the identical 41.666667 — a difference of 0.000000000. **The bill passes the disclosure test outright, in the strongest historical form available: it was fully disclosed to the authority whose prohibition it was designed around and continued to work. Dry exchange fails the same test by construction.**

Counter-result. The only reading under which the architecture clears its build threshold is the naive one assuming zero correlation between coordinating jurisdictions, giving $h = +0.083099$. **Every defensible bloc assignment returns a negative surplus.** And the jurisdictional axis is not the binding one: on the settlement-messaging axis the same structure reads $N = 1.305824$ and $h = -0.406608$, a gap of 0.284191 larger than everything the diversification programme achieved.

The honest verdict on the assembled architecture is: correct in design, insufficient in economics, insufficient in a way that is measurable and decomposable, and not repairable by adding another jurisdiction.

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every case in this paper — the entities, the exposure shares on all four axes, the cost allocation, the hazard rates, the error rates, the screening population, the historical branch and loan-book vectors, and the exchange quotations — is **CONSTRUCTED** or **STIPULATED** and is labelled as such at each point of use. The numerical results are exact computations on those inputs, carried at full precision and rounded only for printing. **No real structure, institution, bank, jurisdiction holding, name, shareholding or transaction appears anywhere in this paper, and no allegation is made about any of them.**

All figures were produced by the scripts named in their captions (`figs/src/CCP-063_fig01.py` through `figs/src/CCP-063_fig05.py`), which import the shared figure style module and read their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

Jurisdiction-specific legal and regulatory descriptions in this paper are taken from the source manuscript and not from primary legal sources. The bibliography contains general instruments, comparative doctrine and multilateral standards only; it contains no statutes, regulations, agency guidance, court rules, registry documents or case reports for any individual jurisdiction, and no canon-law sources. **Appendix B lists every affected claim.** Any reader relying on those descriptions must verify them against the official sources, and none of them is cited here.

Historical claims are described from the source manuscript. de Roover [1963] is cited as the canonical modern scholarship and is where a reader wanting archival figures should go; **no archival figure is asserted in this paper and every historical share vector is constructed and labelled at the point of use.**

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. Every person, entity and holding referred to in this paper is constructed.

Legal, tax and investment advice

This paper is analytical and is not legal, tax, or investment advice. It describes entity types, jurisdictions, ownership rules and structural arrangements for the purpose of measuring concentration and resilience, and it makes no recommendation in respect of any particular person, asset, structure, counterparty or jurisdiction. **Sanctions designations, ownership and control tests, reporting obligations, substance requirements and limitation periods differ by jurisdiction, differ between regimes applying to the same facts, and change without notice.** Every such statement in this paper is carried from the source manuscript, is flagged in Appendix B, and requires confirmation with qualified counsel in the relevant jurisdiction before any reliance. **Nothing in this paper is guidance for moving value that any regime prohibits**, and Section 8's arithmetic is stated for the purpose of designing against a rule rather than around it — the constraint Curzi [2026] states and this paper adopts.

Appendix A. The reporting standard

Any structure measured under this instrument reports the following, and nothing less. Items 1 to 9 are Curzi [2026]'s standard; items 5, 7, 9 and 10 are extended here because this paper's measurement has two defensible answers and the format must force both into print.

```
share basis, stated in one sentence
exposure table, ALL FOUR AXES, each summing to 100.0%
naive HHI and effective count, per axis
bloc or dependency collapse per axis, with the grouping rule stated
N per axis, with the MOST CONCENTRATED axis named FIRST
perimeter exposure as a percentage
ln(N), eta ln(N) and h -- computed on the BINDING axis AND on the
   jurisdictional axis, both printed
robustness table over contested bloc assignments
break-even probability p* and cost ceiling for every leg added or
   proposed in the period
cost allocation of C by line, with the year-on-year change
the date, and the diff against the previous measurement
```

Item 7 is the one this paper adds, and it exists because the two answers differ by 0.284191 **on the constructed facts.** A pack printing one of them has chosen which one to print. **Item 11 decays first and matters most, because perimeter fragmentation is invisible except as a difference between two measurements taken a year apart.**

Appendix B. Claims requiring verification against primary sources

The following rest on legal, regulatory or historical detail, and **none of them is cited to the bibliography, which contains no primary sources for any individual jurisdiction, no agency guidance, no case reports and no canon-law sources.** They are reproduced from the source manuscript, are described rather than asserted, and require confirmation with counsel or against the primary scholarship before any reliance.

1. That any sanctions regime's liability standard is strict, and the consequences of that for the design of an assessment programme. Confirm the current standard under each regime.
2. That any regime's ownership rule deems at fifty per cent aggregate direct or indirect holding, that its indirect limb terminates at natural persons rather than at a fixed depth, and the stated position on entities controlled but not owned at the threshold. **Section 8 analyses a rule of the stated form and asserts nothing about any regime's actual rule.**
3. The divergence between ownership and control tests across the three regimes described, including the indicators of control relied on, and the appellate authority on the breadth of a control limb — its citation, its precise holding, whether the relevant reasoning was obiter, and any legislative or guidance response. **This paper names the issue and states no outcome.**
4. Blocking-statute positions and the conflict-of-laws exposure they create for a group with entities on both sides, including any authority on the operation of such a prohibition in a contractual termination.
5. Reporting and classification obligations under automatic-exchange and beneficial-ownership regimes, per entity and per sub-fund, including the current participating-jurisdiction lists and the status of any regime whose scope has moved through litigation or administrative action. **No current scope is stated anywhere in this paper.**

6. The legal characteristics of every named or de-branded centre relied on by the bloc assignment: separate legal personality, purpose or beneficiary structures, founder control retention, free-zone statutes, promotion conditions, treaty positions, and the supervising courts. **Where the characteristic could not be supported, the centre is de-branded in the text and the description carries no claim about any particular jurisdiction.**
7. The security legislation, final-appellate decision, later interpretation, subsequent ordinance and judicial withdrawals and resignations underlying the tail refusal in Section 7.3. **The refusal is a stipulated modelling position; the arithmetic of the capped transit position is unaffected by how a reader assesses the tail.**
8. The canon-law prohibition and its sources: the twelfth-century compilation, the two conciliar canons, the decretal collections and the letter and decretal named in the source, together with their canon numbers, attributions and correct readings. **Take these from a canonist.**
9. The extrinsic titles — loss actually suffered by reason of delay, profit forgone, a penalty agreed in advance, and payment for genuinely assumed risk of the principal — their acceptance dates, and the degree of controversy attaching to each.
10. The scholastic treatments of the exchange contract by the two named fifteenth-century authorities, the specific works, and the positions actually taken. **This paper does not characterise them.**
11. The doctrinal status and dates of condemnation of dry exchange, including whether the condemnation was conciliar, papal or municipal, and where.
12. Customary intervals by route. The paper uses sixty, ninety, ninety-one, one hundred and twenty and one hundred and eighty days illustratively and asserts none as historical.
13. Actual exchange quotations with dates. **Every rate in Section 10.2 is stipulated.**
14. Whether the historical branch partnerships were of the specific Florentine form described, and how the centre's liability toward branch creditors was in fact limited in law and tested in practice. **This is the load-bearing legal claim of the compartmentalisation argument in Section 10.6 and it should not be relied on until checked** against de Roover [1963].
15. The dates and terms on which the depositary office was held, the occasions on which it was lost, and the formal duties attaching to it.
16. The identity, dates and conduct of the two branch managers associated with the failures, and the apportionment of responsibility between manager and centre, which is a live scholarly question with a more nuanced answer than the summary version.
17. The size and terms of the sovereign loans, the dates of branch closure, and the recovery if any. **Every loan-book share in Section 10.5 is constructed.**
18. The philological demonstration of Section 11.21 — its year of composition as against circulation, the independent earlier judgement, the independent English conclusion, the year and place of the printed edition, and the manner and extent of its later use.
19. The composition of the remittance flows described in Section 10.4 and their proportions.
20. Unicode confusables data and the security-mechanisms report underlying the homoglyph arithmetic of Section 9.3. Curzi [2026] is the primary treatment; nothing in this paper should be relied on for that material.
21. Whether any screening implementation canonicalises before comparison, at what normalisation form, and whether it retains and hashes the raw bytes. **This is a question to ask in writing, and the answer is frequently that it does not.**
22. All cost figures are index units in a constructed model. They are not fee quotations and no real fee schedule was used to produce them.

All numerical results in this paper were computed from the stated inputs at full precision and rounded only for printing.

Chapter 64

Five Houses, One Perimeter: Bloc Collapse in a Historical Network, the Completeness Gap of Double Entry, and Canonicalisation as an Integrity Control

Abstract. The nineteenth-century five-house merchant bank is treated as a measurement problem rather than a legend. Weighting the houses London 40, Frankfurt 20, Paris 20, Vienna 13 and Naples 7 gives a naive HHI of 0.261800 and 3.819710 effective jurisdictions; collapsing them onto the Concert of Europe as a single coordination perimeter gives bloc HHI 1.000000, N = 1.000000 and h = -0.500000. The same five cities read as a modern EU/UK split give 1.923077 and h = -0.271126. An in-house courier cut exposure from three correspondent hops of eight days to one, a 66.6667 per cent reduction, and the modelled loss probability from 3.5360 per cent to 1.1928 per cent. A Bayesian treatment of the coordination question returns a posterior of 0.000732 for a unified hand against 0.999268 for competing centres, odds of 1,365:1. Double entry is then shown to be a consistency control and not an integrity control: the identity A - (L + E) = 0.0000000000 holds exactly while the statement misleads, two correctly kept books diverge by 0.460829 per cent on a single euro invoice, nine entities require thirty-six bilateral reconciliations against one shared record, and a constructed nine-capital vector is 63.850000 per cent represented and 36.150000 per cent structurally unrepresentable. Canonicalisation closes a class the ledger cannot see: canon(x) = restrict_ASCII(confusable_fold(NFKC(x))) collapses 495 of 495 manipulated entries onto the clean form and takes a reconciliation false-break rate from 40 per cent to 0 per cent.

Keywords: branch partnership; coordination perimeter; courier network; Bayesian posterior; double entry; completeness gap; canonicalisation; homoglyph substitution

Series: D. Structure and Resilience **Paper:** CCP-064

1. Introduction

A structure that spans several sovereign jurisdictions is usually described by counting them. This paper treats one of the most-cited historical instances of such a structure as a measurement problem instead, and then follows the measurement down two layers — into the bookkeeping instrument the structure ran on, and into the character encoding that instrument's records are written in. The three layers turn out to fail in the same way, and the shape of the failure is the paper's subject: **at each layer a count is taken over units that are assumed to be distinct, and at each layer the assumption is supplied from outside the instrument doing the counting.**

The paper's contributions are separable and are stated with their numbers.

The bloc measurement on a historical network. Weighting five linked merchant-banking houses at London 40, Frankfurt 20, Paris 20, Vienna 13 and Naples 7 gives a naive Herfindahl index of 0.261800 and 3.819710 effective jurisdictions. Collapsing the five onto a single coordination perimeter gives a bloc index of 1.000000, $N = 1.000000$ and $h = -0.500000$ at the corpus's standard surplus parameters. The same five cities read under a modern nominal EU/UK split give 1.923077 and $h = -0.271126$; read under the present coordination mechanisms they give 1.000000 and $h = -0.500000$ again. **The naive count is identical in both eras and the bloc count is not**, which is the whole of the instrument's content in two rows separated by two centuries. The spread across three defensible readings of the same 1815 structure is 0.469061 of h, and that spread is a definitional choice rather than a fact about the world.

The completeness gap of double entry. Double entry is shown to be a consistency control and not an integrity control. A wholly fictitious sale posted into a stipulated balance sheet leaves the accounting identity at $A - (L + E) = 0.0000000000$ exactly while moving reported assets by 66.6667 per cent, so **the identity's holding is not evidence of anything except arithmetic.** Two correctly kept books disagree by USD 500 on a single EUR 100,000 invoice — 0.460829 per cent of the buyer's figure — with neither party in error. Nine reporting entities require $n(n-1)/2 = 36$ bilateral reconciliations against one shared record. And a constructed nine-capital value vector is 63.850000 per cent represented on a balance sheet and 36.150000 per cent structurally unrepresentable.

Canonicalisation as an integrity control at the character layer. A deterministic synthetic demonstration shows that 495 of 495 substitutable manipulated identifiers are byte-distinct, digest-distinct and canonically identical; that $\text{canon}(x) = \text{restrict_ASCII}(\text{confusable_fold}(\text{NFKC}(x)))$ collapses all 495 onto the clean form idempotently; and that the false reconciliation-break rate falls from 40 per cent to 0 per cent at a 40 per cent manipulation rate, and from the manipulation rate to zero at every point of the sweep.

Why the three sit in one paper. The concentration instrument of Curzi [2026] — the Herfindahl index as the sum of squared shares and the effective count as its reciprocal, with a naive 3.407968 collapsing to exactly 1.00 once blocs are recognised — takes a partition of counterparties as an input it cannot check. Section 5 shows that the ledger cannot supply that partition, because concentration is a relationship between entries and double entry records entries. Section 6 shows that the identity resolution underneath the partition is string comparison, and that string comparison has a measurable blind spot which **biases the concentration measurement in the reassuring direction.** The same defect at the payment layer is measured by Curzi [2026], whose naive correspondent index of 0.302400 — 3.306878 effective correspondents — collapses to 0.745256 and $N_{\text{rail}} = 1.341819$ once flows are decomposed by currency; and the same defect at the assembled-structure layer is measured by Curzi [2026], which returns $N = 2.941176$ effective blocs and $h = -0.122417$ against a break-even $N^* = 4.172734$.

Three attractive claims are refused in print. The most famous story about the historical network — a one-day trading coup on early news of a battle — **is treated as a fabricated data point and is not used**, with the documented multi-year position substituted and the fabrication's provenance dated. The reading under which the five houses are facets of one coordinated actor is **rejected at odds of 1,365:1** on stipulated likelihoods, and the converse error, that

the absence of a unified actor means there is nothing structural to study, is refused in the same paragraph. And the 25 per cent currency-inflation figure of Section 6.5 is **badged as a property of the chosen synthetic 8000/2000 ratio and not as a constant**, with the confusable-only detection check reaching only 58.99 per cent against a headline of 100 per cent.

2. Hypotheses

Six hypotheses, each stated so that it could be wrong, each with its falsifier. Platt [1964]'s strong inference is the standard applied. Statements in the paper that are *not* falsifiable are listed at the end of this section and labelled as definitions, modelling choices or refusals.

H1 — the perimeter hypothesis. The effective count of a jurisdictional structure is set by the number of independent coordination perimeters it spans and not by the number of nominal jurisdictions. *Predicted:* holding the five cities, the five legal families and the share vector fixed and changing only the coordination mechanism moves the effective count while leaving the naive count at 3.819710. The 1815 financial reading returns $N = 3.819710$; the modern coordination reading returns $N = 1.000000$; h moves from -0.030939 to -0.500000. *Falsified by:* a demonstration that a decision taken inside one of the five capitals in 1815 produced enforcement in another without a second independent decision there, which would collapse the 1815 reading and remove the contrast.

H2 — the identity hypothesis. The accounting identity's holding carries no information about whether the recorded events occurred. *Predicted:* a wholly fictitious entry can be posted such that $A-(L+E)$ remains exactly 0.0000000000 while every reported ratio moves — assets by $+66.6667\%$, gearing by -60.6061%, days-sales-outstanding from 36.500000 to 118.625000 days. *Falsified by:* an arithmetic control internal to double entry that flags the entry, or a demonstration that the identity cannot close on a fictitious posting.

H3 — the two-books hypothesis. Two internally consistent sets of books can disagree about the same event with neither party in error, and the disagreement has no arbiter within the method. *Predicted:* a single EUR 100,000 invoice booked at 1.0900 by the seller and 1.0850 by the buyer leaves a USD 500 gap, 0.460829 per cent of the buyer's figure, with both trial balances closing to 0.0000000000. *Falsified by:* a rule internal to double entry that determines which of two permitted rate sources or timestamps is correct.

H4 — the reconciliation-cost hypothesis. In the absence of a shared record, the reconciliation relationship count grows quadratically in the number of parties while a shared record's grows not at all. *Predicted:* $n(n-1)/2$ against 1; at $n = 9$, 36 against 1; at $n = 50$, 1{,}225 against 1. *Falsified by:* a bilateral reconciliation regime whose relationship count is sub-quadratic without a shared record, or a shared-record regime whose count grows with n.

H5 — the completeness hypothesis. A representation that requires every element to carry an amount in a unit of account cannot represent value that has no measurable cost or observable price, and the omission is silence rather than a zero. *Predicted:* on the constructed nine-capital vector, 63.850000 per cent represented and 36.150000 per cent structurally unrepresentable, with the unrepresented fraction appearing nowhere on the face of the statements. *Falsified by:* a recognition framework that admits an unmeasurable item as an entry without destroying the closure property, which would show the gap to be a policy choice rather than a structural one.

H6 — the canonicalisation hypothesis. Confusable substitution in a critical machine field produces records that are byte-distinct and canonically identical, and a canonicalisation applied at the comparison boundary removes the resulting false breaks entirely without removing true matches. *Predicted:* 495 of 500 pairs byte-distinct, digest-distinct and canonically equal (99.0 per cent, the residual 5 being fields with no substitutable character); 495 of 495 collapsing onto the clean canonical form; false reconciliation-break rate 0.0 per cent at every manipulation rate from 0 to 100 per cent against byte equality's one-for-one tracking; and, on a synthetic watchlist, 0 of 20 probes matched by exact comparison against 20 of 20 by canonical comparison with all 10 clean names still matching. *Falsified by:* a manipulated value whose

canonical form differs from the clean value's, a canonicalisation that is not idempotent, or a true match lost to the folding.

Statements in this paper that are not hypotheses, and are labelled as such.

- *The bloc assignment itself* is a **modelling choice**, not a finding. Section 4.2 gives three defensible readings of one structure spanning 0.469061 of h and states that nothing in the arithmetic selects among them. Curzi [2026] is the authority for treating a bloc call as a claim requiring evidence rather than as a fact, and this paper's spread is the strongest available warning attached to that authority.
- *The disclosure test* is a **definition** of a design property, not an empirical claim. What is testable is whether a given element passes it.
- *The refusal of the one-day-coup account* is an **evidential judgement**, argued in Section 4.5 from the dating of the sources rather than measured.
- *The competing-centres posterior* is a **device for making an inference explicit.** Its likelihoods are stipulated. Section 4.8 states what is and is not carried by it.

3. Method

3.1 The instrument, and which effective count it is

Throughout this paper the concentration instrument is the Herfindahl index HHI $= \sum_i s_i^2$ over non-negative shares summing to one [Hirschman, 1964, Herfindahl, 1950], and the effective count is its reciprocal. **The reciprocal index is the Hill number of order two** [Hill, 1973]: the member of the family ${}^qD = (\sum_i p_i^q)^{1/(1-q)}$ at $q = 2$, which returns the number of equally weighted units that would produce the observed index. Adelman [1969] is the numbers-equivalent reading; Jost [2006] is the statement that entropies are not diversities and their numbers-equivalents are; the same arithmetic was reached independently in ecology by Simpson [1949] and in political science by Laakso and Taagepera [1979]. The exponential of Shannon entropy [Shannon, 1948] is the same family at $q = 1$, and the reciprocal index is the exponential of the Rényi entropy of order two [Rényi, 1961].

Which effective count is meant is stated every time it is used, because the corpus runs two quantities under that name and only one of them is a Hill number. Every effective count in this paper is $1/\text{HHI}$, the Hill number of order two. The corpus's affine count $1 + (N - 1)u$ appears nowhere in this paper; it is a linear interpolation between 1 and N, it is not a numbers-equivalent, and Curzi [2026] demonstrates in its verification section that the chain establishing the reciprocal index fails on the affine form at a nameable step.

Two cautions from the same literature are printed next to the instrument rather than in a footnote. Hurlbert [1971] argues that "diversity" has carried so many distinct quantities that the word means nothing without its formula attached, which is why the formula is attached here. Goodhart [1975] and Campbell [1979] are the standing warnings that a measure adopted as a target ceases to measure what it stood for — directly relevant to an index that regulators band [U.S. Department of Justice and Federal Trade Commission, 2023, 2010], and directly relevant again in Section 6, where an encoding defect moves the index without anything real changing. Chao et al. [2014] show that low-order diversity estimators are biased on small samples; a five-house or six-jurisdiction vector is a small sample and this paper does not correct for it.

The bloc collapse has an exact variance reading, developed in Curzi [2026] and used here without re-derivation. Treating exposure as a portfolio [Markowitz, 1952], with w_i the share exposed to unit i and X_i a unit-variance shock, $\text{Var}(\sum_i w_i X_i) = w'Rw$. Setting $R = I$ recovers the naive index and setting R to the block matrix of ones recovers the bloc index, so the two indices are one variance under two correlation assumptions. The general form $1/(w'Rw)$ is the similarity-sensitive diversity of order two of Leinster and Cobbold [2012], whose additive relative is the quadratic entropy of Rao [1982]. Buldyrev et al. [2010] is the formal statement

of why the bloc reading is the conservative one: coupled networks are far more fragile than either network alone.

The surplus reading used to price a change in the effective count is the corpus's $h = \eta \ln N - \delta_{\text{floor}}$ with $\eta = 0.35$ and a maintenance floor of 0.500000, fixed in advance and identical to the parameters of Curzi [2026] and Curzi [2026]. **Those parameters are stipulated, not estimated**, and every h in this paper inherits that status.

3.2 The perimeter test, stated operationally

The bloc question is answered by one observable, and the observable is stated so that it can be checked against a historical or a contemporary record without argument about alliances:

> A **coordination perimeter** is a set of jurisdictions among which one member's decision becomes another member's enforcement **without a second independent decision taken in the second jurisdiction.** Two jurisdictions are in one perimeter when such a mechanism exists between them and in different perimeters when it does not. Political alignment, shared interest, treaty membership and diplomatic congress are not the test; a transmission mechanism is.

This is the operational form of the perimeter test of Curzi [2026], and Section 4.2 applies it three times to one structure to show that it does not always return a unique answer. Where it does not, the honest presentation is to give every defensible reading and to price the disagreement, which is what Section 4.2 does.

The classes of mechanism that matter for the modern reading are the automatic exchange of account information [Organisation for Economic Co-operation and Development, 2017], the anti-money-laundering recommendations and their evaluation and listing cycle [Financial Action Task Force, 2012], the substance and beneficial-ownership work [Organisation for Economic Co-operation and Development, 2015], supranational listing [de Goede, 2011], and correspondent risk-appetite transmission [Curzi, 2026]. **What these have in common is that they convert a decision taken in one place into an effect in another without a second decision there**, which is exactly the test. Pistor [2019] is the standing account of how legal coding produces such effects, and Hansmann and Kraakman [2000] is the reduction of entity law to the asset-partitioning function that the five-house architecture exploits.

3.3 The evidence discipline

Three rules govern what this paper is willing to assert, and they are stated in the Method rather than in the Limitations because they changed what the paper contains.

Sources are graded, and the grading is printed. Every substantive assertion in Section 6 carries a tier: **T1** for an established fact about a published standard or a mathematical property; **T2** for a quantity measured by the deterministic demonstration and independently recomputed for this paper; **T3** for a proposed design that is not adopted and not in force; and **T4** for a claim deliberately *not* made. The T4 tier is the important one, because the temptation this material creates is to claim more than it supports.

A widely believed account with a datable polemical origin is refused rather than hedged. Section 4.5 sets out the reasoning. The general rule is that a hostile source and a hagiographic source are not evidentially symmetric and averaging them produces a motivated account with a smaller error bar and no better accuracy. This is the two-column discipline of Curzi [2026w]: the evidence column carries what could have come out otherwise, the precedent column carries only that a claim was independently arrived at and preserved, and two sources agreeing does not double the evidence when only one of them is evidence.

A contested inference is written as a posterior rather than as an assertion. Section 4.8 states a prior, states four likelihood pairs, and prints the resulting odds, so that a reader who

disagrees can substitute their own numbers and recompute. The likelihoods are stipulated and the paper says so at the point of use, in the conclusion, and in the limitations.

3.4 The disclosure test

Every design element in this paper is audited against a single question, and the audit is printed: **would the element still do its job if every element of it were disclosed to every authority with jurisdiction over any part of it?** An element that passes is architecture; an element that fails depends on concealment.

This is not an ethical preference and it is not this work's to claim. It is Kerckhoffs [1883]'s principle, stated in 1883 — a system must remain secure when everything about it except the key is public — restated as *open design* among the protection principles of Saltzer and Schroeder [1975], and given its transparency-as-remedy form by Brandeis [1914]. Curzi [2026] states the resulting constraint in the form this paper uses: **channel separation is strengthened by disclosure because its effectiveness is entirely independent of whether anyone knows it is in place, so every element must be publishable, and non-disclosure is refused as a design premise rather than merely discouraged.** The regulatory environment has independently removed non-disclosure from the set of available design parameters [Organisation for Economic Co-operation and Development, 2017, Financial Action Task Force, 2012, Organisation for Economic Co-operation and Development, 2015].

The test is applied three times in this paper and it **splits** in Section 4.3, which is the result that makes it non-trivial: four elements of the historical architecture pass and one fails, and the element that fails is the one that did not survive.

4. A five-house network as a measurement problem

4.1 The architecture

The structure examined here is the early nineteenth-century five-house merchant bank in which one founder placed five sons in five sovereign jurisdictions — Frankfurt, London, Paris, Vienna and Naples — and constituted them as partners in a linked but legally distinct set of houses. It is treated here as the closest thing the historical record offers to a controlled instance of the corpus's central identification, because the geography, the legal families and the ownership are held fixed across two centuries while the coordination mechanism is not. **The historical description in this section is taken from the source manuscript, which relies principally on Ferguson [1998]; every specific date, deed, party and figure is listed in Appendix A as requiring verification against the primary record, and the capital shares are constructed.**

Four structural properties matter for the measurement.

1. **Separate legal houses in separate sovereign jurisdictions**, each under its own law, currency, courts and sovereign. This is asset partitioning in the sense of Hansmann and Kraakman [2000], executed across sovereignties rather than within one.
2. **Common partnership economics across them**, so that profits pooled while liabilities did not. de Roover [1963] documents the same arrangement at a smaller scale and within one commercial republic's diaspora, and Curzi [2026] measures its common-apex exposure.
3. **A private communications network** — couriers and contracted vessels, operated rather than rented. Section 4.4 prices it.
4. **Correspondence written in one language using another's script.**

The fourth property is routinely described as a cipher and is not one, and the distinction is load-bearing for this paper. A cipher transforms plaintext under a key; its security rests on the key and the transformation is public. Writing one language in another's script is a **change**

of representation: no key, no algorithm, and a property that rests entirely on the size of the population able to read the result. That property degrades monotonically as literacy spreads and as states acquire readers. **It is the same object Section 6 studies at the character layer — a change of representation is not a change of content — and treating one as the other is a vulnerability class rather than a control.** Confidentiality resting on a scarce readership is confidentiality with a decay rate, and it fails the disclosure test of Section 3.4 by construction, because its function is precisely that it not be read.

4.2 The bloc computation: three readings of one structure

The share vector below is **CONSTRUCTED for illustration.** The archival capital allocations across the five houses are not reproduced here and are listed in Appendix A.

```
CONSTRUCTED capital shares, c. 1815
London 40%   Frankfurt 20%   Paris 20%   Vienna 13%   Naples 7%

naive HHI               = 0.261800
effective jurisdictions  = 3.819710      (of 5 nominal)
h at the naive count     = -0.030939
```

The bloc question is then answered three ways, and the disagreement is priced rather than resolved by assertion.

Reading 1 — the post-1815 congress system as one coordination perimeter. On this reading the five capitals shared a coordination mechanism and the structure is one bloc.

Reading 2 — three perimeters. Britain stood apart from the eastern courts' alignment and repeatedly declined the interventionist principle; Restoration France sat inside the congress system but outside that alignment; Austria, the German Confederation and Bourbon Naples moved together.

Reading 3 — five perimeters. The congress system coordinated on legitimacy, territorial settlement and armed intervention. It did not coordinate on the mechanisms of Section 3.2: there was no common standard for financial supervision, no shared exchange of tax information, no common list of non-cooperative jurisdictions, no correspondent risk-appetite transmission channel and no supranational body empowered to grade a member's financial regime. **A decision taken in one capital did not reliably produce the corresponding effect in another without a separate decision taken there**, which is the perimeter test failing to join them.

reading	bloc partition	bloc HHI	N	h	Δh vs 1
1 congress as one perimeter	all five	1.000000	1.000000	-0.500000	—
2 three perimeters	40 / 40 / 20	0.360000	2.777778	-0.142422	$+0.357578$
3 five perimeters	40/20/20/13/7	0.261800	3.819710	-0.030939	$+0.469061$

This paper's judgement is that reading 3 is the better description for *financial* purposes and reading 2 for *political* ones, and that the honest presentation is to give both and name the observable that separates them. The observable is the mechanism list of Section 3.2. **A coordination perimeter is not an alliance; it is a mechanism by which one state's decision becomes another state's enforcement without a second independent decision, and the post-1815 congress system had that for troops and did not have it for money.**

The 0.469061 spread is the paper's most important negative result about its own instrument. Three readings of a single structure, all defensible, differ by more than any single

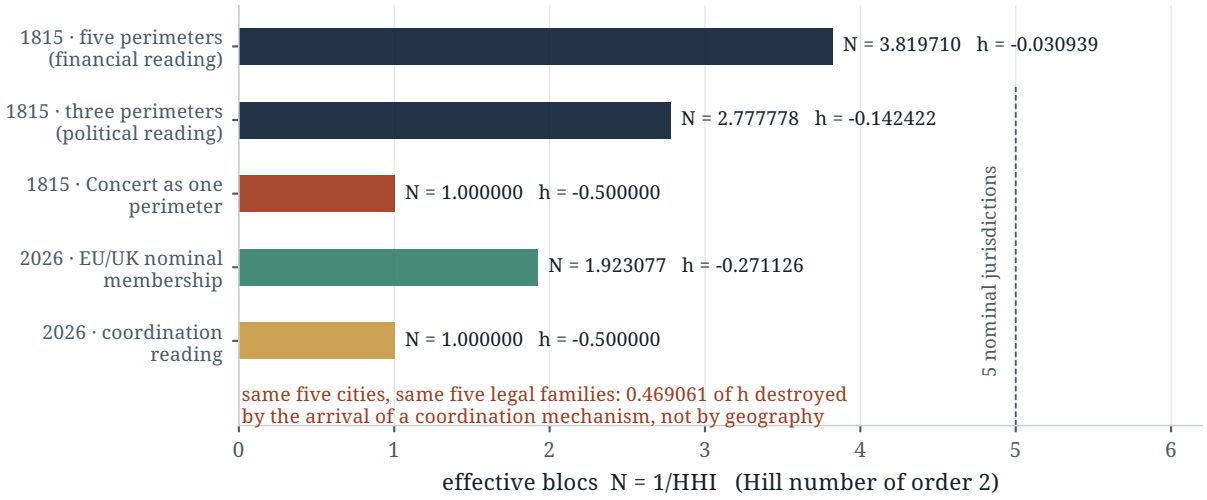


Figure 64.1: **The bloc call, not the geography, moves the effective count.** Five constructed capital shares (London 40, Frankfurt 20, Paris 20, Vienna 13, Naples 7) read under five partitions. The naive effective count is 3.819710 in both eras; the bloc count moves from 3.819710 to 1.000000 and h from -0.030939 to -0.500000. Effective count is $1/\text{HHI}$, the Hill number of order two; $h = 0.35 \ln N - 0.500000$ at parameters fixed in advance. Built by `figs/src/CCP-064_fig01.py`; share vector CONSTRUCTED.

structuring decision measured anywhere in Series D. Nothing in the arithmetic selects among them. Curzi [2026] therefore requires that a bloc call be presented as a claim with evidence and a stated observable rather than as a fact, and this measurement is the reason.

4.3 The same five cities under a modern coordination mechanism

The identical geography is now read under the mechanism classes of Section 3.2.

reading	partition	bloc HHI	N	h
nominal supranational membership	60 / 40	0.520000	1.923077	-0.271126
coordination reading	all five	1.000000	1.000000	-0.500000

Against the 1815 five-perimeter reading the loss is -0.469061 of h; against the three-perimeter reading, -0.357578.

The same five cities, the same five legal families, the same five currencies at the outset — and 0.469061 of h removed by the arrival of a coordination mechanism that has nothing to do with geography. Nothing about the cities changed. What changed is that the mechanism classes of Section 3.2 came into existence and now enclose all five. The enclosure is invisible on a map and invisible to the naive count, which reads 3.819710 in both eras.

This is the same result Curzi [2026] obtains for a modern six-jurisdiction structure, where a naive 3.407968 collapses to exactly 1.00; it is obtained here a second way and against a structure that was once genuinely diversified. **A naive index computed on a single-bloc structure is not a weak measurement, it is a measurement of the wrong thing**, and the two rows above are that failure mode separated by two centuries.

4.3.1 The disclosure test on the historical architecture, which splits

element	disclosed?	still does its job?
five houses in five sovereignties	yes	**passes.** Each was a known, licensed, taxed house publicly transacting government loans; the function was jurisdictional, not secret
common partnership economics	yes	**passes.** Family partnership was the standard commercial form
private courier network	yes	**passes.** A faster route is not a secret route
correspondence encoding	yes	**fails.** Its entire function was that chanceries could not read it; disclose the script convention and the property is gone

Three passes, one fail, and the fail is the informative row. **The element that fails the disclosure test is the element that did not survive**: script-based confidentiality decayed to nothing as literacy and state capacity grew, while the jurisdictional architecture did its work for a century. That is the disclosure constraint of Curzi [2026] stated as a historical prediction that came true, and it is the strongest evidence in this paper that the constraint is an engineering property rather than a preference. Structures that depend on not being seen have a half-life; structures that depend on being correct do not.

4.4 The transit network priced as an exposure window

The information advantage of such a network is conventionally valued in the currency of a one-day news lead. This paper values it as Curzi [2026] values a payment rail: a transit capability is priced by its exposure window, not by its speed record.

Stipulated transit parameters, illustrative, one continental capital to London around 1815, at a stipulated constant counterparty hazard $\lambda = 0.0015$ per day and a loss probability $P = 1 - e^{-\lambda t}$:

```
via correspondents    :  3 hops x 8 days = 24 days of exposure
in-house own courier  :  1 hop  x 8 days =  8 days of exposure
reduction in window   : 16 days = 66.6667%

  correspondents   P(loss) = 0.035360   (3.5360%)
  in-house         P(loss) = 0.011928   (1.1928%)

expected-loss reduction on a 100,000-unit transfer :  2,343.1419 units
at 200 transfers per year                          :  468,628.3876 units
```

The hazard model is stated here because the source manuscript does not state it, and the choice changes the fifth decimal place. Under the continuous form $1 - e^{-\lambda t}$ the 24-day figure is 3.535978 per cent, which is the source's 3.5360 per cent. Under the discrete form $1 - (1 - \lambda)^t$ it is 3.538576 per cent. **The two differ by 0.0026 percentage points and the paper uses the continuous form**, which is what reproduces the source's figures; the discrepancy is recorded in Section 8.11 rather than concealed by rounding.

That is the documented, repeatable value of the network, and it is large. It is also entirely unglamorous, which is why the fabricated version displaced it. **A structural reduction in settlement exposure applies to every transaction and compounds; unlike a piece of news, it does not have to be right.** Curzi [2026] measures the same quantity on a modern rail and reaches the same conclusion about what a transit capability is worth.

Two cautions belong with the figure. The three-hop chain is a series of intermediaries, and nominal redundancy across correlated intermediaries is not redundancy: Buldyrev et al. [2010] is the formal statement, and Sagan [1993] is the organisational one — added redundancy can reduce safety where the added paths share a failure mode. And a private network is a single engineered artefact whose robustness comes from design rather than from topology, which is Doyle et al. [2005]'s correction to the degree-distribution account of network resilience and the reason no scale-free claim is made here.

4.5 A refused data point

The most widely repeated account of this network holds that its London principal received news of a decisive battle roughly a day before the government, sold ostentatiously to induce a panic on the belief that the battle had been lost, bought the collapsed market through agents, and made a fortune in a single session.

The source manuscript treats most of that as a later fabrication, and this paper refuses it as a data point. What the manuscript reports as documented, on the scholarship it relies on [Ferguson, 1998], is materially different: the courier network existed and early news was received; the position going into the battle was not a short bet on defeat but a very large accumulation of bullion assembled to finance an army on the assumption of a long war; the sudden peace was in the first instance a problem, leaving a long position in a commodity whose military demand had just evaporated; and the large profit came afterwards, from a multi-year position in government stock taken as the post-war price recovered.

The manuscript traces the dramatic version to a polemical pamphlet published in 1846 under a pseudonym, whose purpose was defamatory. **Its author, title, date and causal role in the account's circulation are listed in Appendix A as requiring verification, and this paper asserts none of them.** What is not in doubt on the manuscript's own account is that a nineteenth-century polemical pamphlet rather than a market record is the origin of the story that circulates.

```
1815   the battle                          --
1846   the pamphlet                         31 years after the event
1998   the archival account                152 years after the pamphlet
                                           183 years after the event
```

The documented version can be given the same arithmetic treatment, so that the two can be compared on one scale. **Prices stipulated for illustration; the actual price path is archival and is listed in Appendix A.**

```
CONSTRUCTED: buy government stock at 62, sell at 82 after 2 years,
             coupon 3 per 100 nominal

capital return          32.258065%
income return            9.677419%
total return            41.935484%
annualised              19.136679%
running yield at buy     4.838710%   at sale  3.658537%
```

To match a 41.935484 per cent total return in one session requires a 41.935484 per cent single-day move in government stock. **The documented trade is larger than the fabricated one,**

it is repeatable, and it is the version an architecture can be built on. A structural information advantage plus a correct multi-year position is an architecture; a one-day coup is a story.

There is a second reason for exactness, and it is the reason the refusal is in the body rather than in a note. The fabricated version is the seed of a durable conspiracy literature. A paper that repeated it would launder that literature into a structural analysis, and a reader who caught it would be right to discard everything else. **Citation discipline on this material is not decoration; it is the difference between analysis and propaganda.** This is the failure mode of building on a fabricated data point, and its cost is directional: anyone taking the fabricated version as a model invests in speed of news and under-invests in the exposure-window reduction of Section 4.4, which is what actually produced the returns.

4.6 The creditor–debtor relation, 1831–32

The source manuscript records that a sovereign borrower of the period — the papal territories, following a revolt in the Legations and a foreign military intervention — borrowed from the houses on market terms across 1831 and 1832. **The amounts, tranches, currencies, security and the identity of the lending houses are listed in Appendix A as requiring verification.**

What the paper takes from it is structural, and it does not depend on any of those particulars. The borrower was a sovereign; the lenders were a private house whose members could not, in several of the jurisdictions concerned, hold real property freely or exercise full civil rights. **The relation the record establishes is creditor and debtor, negotiating price. It is not partnership and it is not identity of interest**, and Section 4.8 uses it as an observation.

A constructed illustration of what a distressed sovereign loan of the period costs a borrower. **Terms stipulated; not the actual terms.** The approximate yield to maturity is the standard approximation $(C + (F - P)/n)/((F + P)/2)$.

```
CONSTRUCTED: face 1,000, issued at 75 per 100, coupon 5%

proceeds to the borrower        750.0000
annual coupon cost               50.0000
running yield to the lender       6.666667%
borrower's cost above the face coupon         +1.666667 points

approximate yield to maturity if redeemed at par:
    in  5 years    11.428571%
    in 10 years     8.571429%
    in 20 years     7.142857%
```

A sovereign borrowing at an approximate 11.43 per cent yield is a sovereign paying a distress premium to a lender pricing default risk.

4.7 A 571-year doctrinal trajectory

The source manuscript sets out a chain of fifteenth-century instruments, a doctrine crystallised from them, its reception into positive law in an 1823 judgment of a common-law supreme court, its elaboration in two judgments of the early 1830s, its rejection as to one component in a 1992 judgment of a different common-law supreme court, its application in a 2005 judgment of the first court, and a 2023 repudiation by the institution that issued the original instruments.

None of these instruments, cases or statements is in this paper's bibliography, and none is cited. `bib/external.bib` contains no case law, no papal instruments and no jurisdiction-specific primary material; the coverage file records this as a known and deliberate gap. They are described here from the source manuscript and every one is itemised in Appendix B with what specifically requires verification. **This paper asserts no case citation, no reporter page,**

no bull date and no institutional attribution. The material is retained because the *intervals* are the measurement and the measurement does not depend on the particulars beyond their years.

```
instrument 1452        -> judicial reception 1823        371 years
instrument 1493        -> judicial reception 1823        330 years
judicial reception     -> later application 2005         182 years
judicial reception     -> repudiation 2023               200 years
instrument 1493        -> repudiation 2023               530 years
later application 2005 -> repudiation 2023                18 years
instrument 1452        -> repudiation 2023               571 years
```

Three statements are made about this and all three are made.

One: the repudiation did not change the positive law and does not purport to. A statement by the issuing institution is not a source of the receiving jurisdiction's law. The rule stands. A different common-law system's 1992 revision of one component shows that an apex doctrine *can* be revised, and that it takes a supreme court rather than an apology.

Two: the repudiation is nevertheless the strongest available corroboration of the provenance. The standard objection to connecting fifteenth-century instruments to modern property law is that the connection is a stretch. The institution that issued the instruments has now formally addressed them as the source of the doctrine. **When the drafting authority confirms the provenance, the provenance stops being contestable and only the consequences remain in dispute.**

Three: the honest caveat, which strengthens rather than weakens the reading. The 1823 judgment is more complicated than a summary allows; its holding operates substantially as a rule of priority between competing European sovereigns and as a rule about who may extinguish, and the original instruments were not the operative deed of any actual land transfer. Neither point rescues the doctrine and neither is a reason to overstate it. **The defensible claim, and the only one made here, is genealogical**: the rule in current law about who holds ultimate title to discovered land descends by an explicit and judicially acknowledged path from instruments issued between 1452 and 1493, and the issuing institution repudiated them in 2023 while the rule continued.

For the corpus the operational payload is a trajectory measurement at the longest available duration. **571 years from instrument to repudiation, 200 from judicial reception to repudiation, and 18 from the last major application to the repudiation — with the underlying rule operative at the end of all three intervals.** Legal foundations are not background; they are the slowest-moving and longest-lived element of any structure, and they move. The correct presentation is a **split trajectory**: positive-law status stable, legitimacy status eroding, the two moving in opposite directions for at least a century. **Averaging the two destroys the only information they carry**, which is the same objection Curzi [2026t] makes to a composite index whose components are not on a common scale.

4.8 The coordination question as a posterior

The reading under which a group of long-lived institutions is a single coordinated actor is taken seriously enough to be tested rather than dismissed.

Two models: U, in which the institutions are facets of one coordinated actor, and C, in which they are distinct institutions whose interests align on some files and conflict on others. The prior is 0.50 each, which is generous to U. Four observations are then applied. **The likelihoods are stipulated. What is claimed is the direction and the order of magnitude of the movement, not the precision of the posterior.**

observation	$P(\cdot \mid U)$	$P(\cdot \mid C)$	posterior $P(U)$
prior	—	—	0.500000
litigation and a 1984 settlement payment between two of them	0.10	0.60	0.142857
a 2013 criminal prosecution of an official of one of them	0.15	0.70	0.034483
a 2023 conviction of a senior officer by the institution's **own** criminal court	0.05	0.65	0.002740
a sovereign among them **borrowing on market terms**, 1831–32	0.20	0.75	0.000732

$P(\text{unified}) = 0.000732$, $P(\text{competing}) = 0.999268$, odds against the unified model 1,365:1. The posterior odds are the exact rational $\frac{1}{6} \cdot \frac{3}{14} \cdot \frac{1}{13} \cdot \frac{4}{15} = \frac{1}{1365}$, which is why the figures are quoted exactly: **they are exact consequences of stipulated inputs, not estimates.** The underlying events are described from the source manuscript and are itemised in Appendix B; none is cited.

The arithmetic is only as good as the stipulated likelihoods, and the direction is not sensitive to them, for a reason worth stating in words. **A unified actor does not litigate against itself, does not settle with itself for a nine-figure sum, does not prosecute its own senior officers in its own courts, and does not borrow from itself at a distress yield.** Each of those is a costly signal: expensive, public, and hard to fake for the purpose of misdirection.

The converse error is refused in the same breath. Concluding that these institutions are unrelated because they are not unified would be equally wrong. They are the same *kind* of object — perpetual corporate persons holding assets across sovereign boundaries, with succession mechanisms that outlive individuals and legal characterisations that determine what they may do — which is precisely the subject of Series D. **The operative model is a small number of long-lived, competing, partly overlapping institutions with genuinely conflicting interests**, which is what the record shows and which is more useful than the alternative, because competing centres have divergent interests and a unified hand does not.

5. What double entry guarantees, and what it structurally cannot

This section does not argue that double entry is obsolete. It is required by every accounting framework, it has been the method for over five centuries, and nothing here proposes replacing it. The argument is narrower and defensible because it is narrower: **double entry is incomplete in two specific and nameable ways, and neither is a defect in its execution.**

The method's codification is conventionally dated to a printed tract of 1494, and the codification followed roughly a century and a half of established merchant practice rather than preceding it. **The historical claims in this paragraph are described from the source manuscript and are itemised in Appendix A**; the accounting-history literature that matters most here is the deflation of the stronger claim usually attached to the method — that it caused rational capitalism — by Yamey [1964] and Littleton and Yamey [1956]. That deflation is the correct precedent for this section: **the limits of the method are already recorded in its own historiography and this paper argues alongside that literature rather than unaided.**

5.1 The method, stated exactly

Every transaction is recorded twice within one entity's books, once as a debit and once as a credit, in equal amounts, so that $\sum \text{Dr} = \sum \text{Cr}$ by construction and the identity $A = L + E$ follows. The trial balance is the self-check, and it is an excellent self-check for what it checks: before it, a transposition, an omission or a duplicated posting could be caught only by re-reading; after it, an arithmetic error announces itself.

The question this section asks is what "announces itself" covers.

the method guarantees	the method does not guarantee
arithmetic closure	**existence** — that the recorded event occurred
completeness of transcription within the entity	**valuation** — that the amount is correct
traceability of a posted entry to its counter-entry	**agreement** — that the counterparty concurs
	coverage — that the monetised subset is the relevant set

Existence and valuation are closed by audit; agreement is closed by reconciliation and would be better closed by a shared record; coverage is closed by nothing currently in general use. That asymmetry organises the rest of this section.

5.2 Result: the identity holds while the statement misleads (H2 supported)

A stipulated opening balance sheet, then a single wholly fictitious sale — no customer, no goods, no service, no contract.

```
OPENING (stipulated)
  cash            400.00      payables        250.00
  receivables     300.00      debt            600.00
  PPE             800.00      equity          650.00
  A − (L + E)  =  0.0000000000

POSTING
  Dr  Receivables   1,000.00
      Cr  Revenue                1,000.00

AFTER
  cash            400.00      payables        250.00
  receivables   1,300.00      debt            600.00
  PPE             800.00      equity        1,650.00
  A − (L + E)  =  0.0000000000        <-- still exactly zero
  debits 1,000.00; credits 1,000.00; difference 0.0000000000
```

metric	before	after	change
total assets	1,500.00	2,500.00	$+66.6667\%$
debt / equity	1.307692	0.515152	-60.6061%
equity / assets	0.433333	0.660000	$+52.3077\%$
receivables / (cash + receivables)	0.428571	0.764706	$+78.4314\%$
reported profit	0.00	1,000.00	—

metric	before	after	change
cash from operations	0.00	0.00	—

The posting is invisible to every arithmetic control and visible only in the relationship between two numbers the ledger does not itself relate: receivables and cash. At a stipulated annual revenue of 3,000 and a 365-day year, days-sales-outstanding moves from 36.500000 to 118.625000 days, a change of 82.125000 days.

That test is the whole of the detection, and its character is the point: it is an **analytical procedure imported from outside the method**, comparing one monetised figure to another and inferring that one of them is unlikely. It is not a property of double entry. **Double entry is a transcription-integrity system; it was never an existence-verification system, and the profession has never claimed it was — which is why external audit, third-party confirmation and analytical review exist as separate disciplines layered on top of it.**

The corollary is the failure mode: the identity holds by construction and cannot fail except through arithmetic error or a bypass, so **its holding carries no information and treating it as a signal is treating a tautology as a measurement.**

5.3 Result: two correct books that disagree (H3 supported)

The second gap needs no fraud. One entirely honest invoice.

```
one invoice, EUR 100,000.00, cross-border

seller books at 1.0900 USD/EUR   ->  receivable  USD 109,000.00
buyer  books at 1.0850 USD/EUR   ->  payable     USD 108,500.00
gap  USD 500.00  =  0.460829% of the buyer's figure

seller trial balance:  Dr Receivables 109,000.00 / Cr Revenue  109,000.00
                       difference 0.0000000000
buyer  trial balance:  Dr Expense    108,500.00 / Cr Payables  108,500.00
                       difference 0.0000000000
```

Both sets balance. Neither party is in error: they used different permitted rate sources or timestamps, each documented in their own accounting policy. **There is no third record to appeal to, because in double entry there is no such thing as a third record — there are two entities each keeping a complete and internally consistent account of one event.**

The cut-off variant is cleaner because it involves no rates at all. The seller recognises on 31 December, the buyer on 2 January; same currency, same amount, same invoice; both books balance in both periods. The consolidated group then shows an item existing on one side and not the other, and the elimination entry that fixes it is a **decision** about which of two internally consistent accounts to prefer. **Elimination is the point at which two incompatible truths are resolved by policy rather than by evidence, and it deserves the scrutiny any balancing figure deserves.**

Double entry provides a mechanism for closing a set of books and no mechanism whatever for reconciling two sets, and the second problem consumes the resources.

5.4 Result: the reconciliation bill (H4 supported)

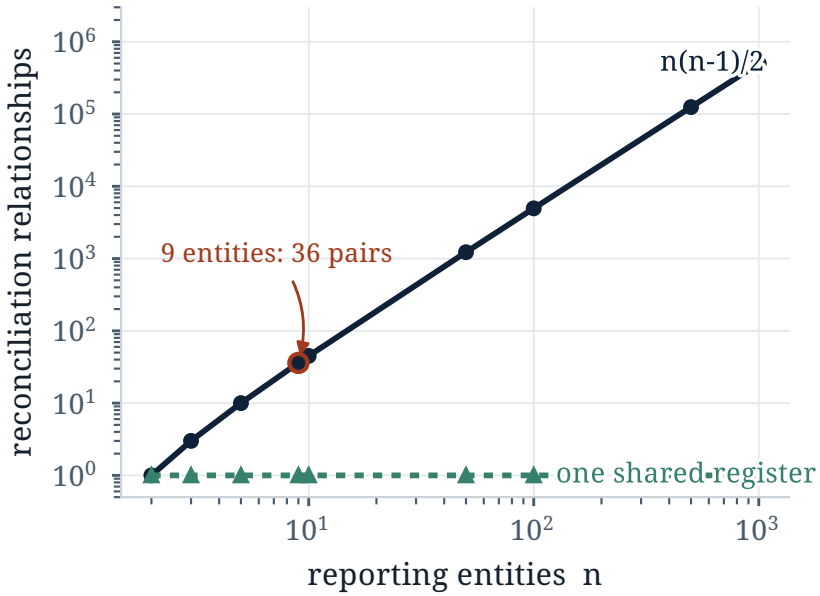


Figure 64.2: **The reconciliation bill is quadratic in parties and constant with a shared record.** $n(n-1)/2$ bilateral reconciliations against one shared register; the nine-entity reporting group of Section 5.4 sits at 36 pairs. Both axes logarithmic. The exponent is exact combinatorics, not a fitted scaling exponent. Built by `figs/src/CCP-064_fig02.py`.

n parties	bilateral reconciliations $n(n-1)/2$	one shared register
2	1	1
3	3	1
5	10	1
9	36	1
10	45	1
50	1,225	1
100	4,950	1
500	124,750	1
1,000	499,500	1

From two parties to fifty the relationship count rises by a factor of 1,225 while the register count stays at one: $O(n^2)$ against $O(n)$. **The exponent here is exact combinatorics, not an estimated scaling exponent**, so the distributional-fitting procedure of Clauset et al. [2009] does not apply to it and no heavy-tail claim is made anywhere in this paper. The same quadratic communication cost is the classical objection to raising the number of parties [Brooks, 1975], and Conway [1968] — which appeared in April 1968, not 1967 as the corpus has elsewhere dated it — is the statement that the resulting interface structure mirrors the communication structure.

Applied to a nine-entity reporting group: 36 bilateral relationships against one shared record. Adding a tenth entity takes the count to 45; adding to twelve takes it to 66, **an 83.3 per cent increase in reconciliation relationships for a 33.3 per cent increase in entities.** The consequence is a cost line, and in the corpus's storage law it sits in the cost term C, which enters h directly and without the logarithm's damping. **Reconciliation cost is one of the few terms in the corpus's central equation that a purely technical decision can reduce, and it is the term that grows fastest as a structure adds legs.**

At a stipulated aggregate intercompany balance of 40,000,000 and the 0.460829 per cent in-

nocent rate divergence of Section 5.3, the expected aggregate translation break is 184,331.60. **That is not a loss. It is a break** — an amount two internally consistent sets of books disagree about, which somebody must investigate, document and eliminate, and which recurs every period. The cheapest available remedy is not technology: a written intercompany policy fixing the rate source and the timestamp for both sides of every internal transaction in advance converts a recurring investigation into a one-off decision.

5.5 Result: the completeness gap (H5 supported)

The second structural limit is harder to demonstrate because its evidence is an absence. The measurement below is of the absence, on a **CONSTRUCTED value distribution stipulated for illustration** over the nine-capital frame developed in Curzi [2026]. **The share vector is not a measurement of any entity and must not be treated as one.**

form of capital	share of total value	recognised
financial	42.000%	fully
manufactured	18.000%	fully
intellectual	11.000%	at a stipulated 35%
human	9.000%	not at all
social / relationship	7.000%	not at all
natural	6.000%	not at all
cultural	3.000%	not at all
living	2.500%	not at all
spiritual / experiential	1.500%	not at all

```
fully on the balance sheet                      60.000000%
partially (intellectual at a stipulated 35%)     3.850000%
TOTAL REPRESENTED                               63.850000%
STRUCTURALLY UNREPRESENTABLE                    36.150000%
```

The ledger balances exactly and provably over 63.850000 per cent of the value and is silent about 36.150000 per cent of it.

The mechanism of the silence is structural rather than a policy choice, and the distinction matters. Recognition in every major framework requires that an item be measurable, and the requirement follows from the fact that an entry must carry an amount and an amount requires a measurement. A framework recognising unmeasurable items would produce a ledger that no longer balanced in any meaningful sense, destroying the one property the method actually guarantees. **The completeness gap is the price of the closure property. It is not a defect a better standard-setter would repair.**

Two failure modes follow directly. **Treating unmonetised value as zero** is wrong because a zero can be challenged, footnoted or revised while an absence cannot: silence is not a conservative estimate, it is the absence of an estimate, and it reads to every downstream user as if the thing did not exist. **Treating a monetised proxy as the thing** is worse: inventing an index and putting it in a report is Goodhart [1975]'s law waiting to happen, and Campbell [1979] is the same warning from the evaluation side. **The correct response to an unmeasurable quantity is a documented judgement held outside the ledger, not a fabricated measurement held inside it.**

5.5.1 The off-ledger register

The action item is not to capitalise any of it, which would be the proxy failure mode. It is the opposite, and it is this paper's only proposed instrument:

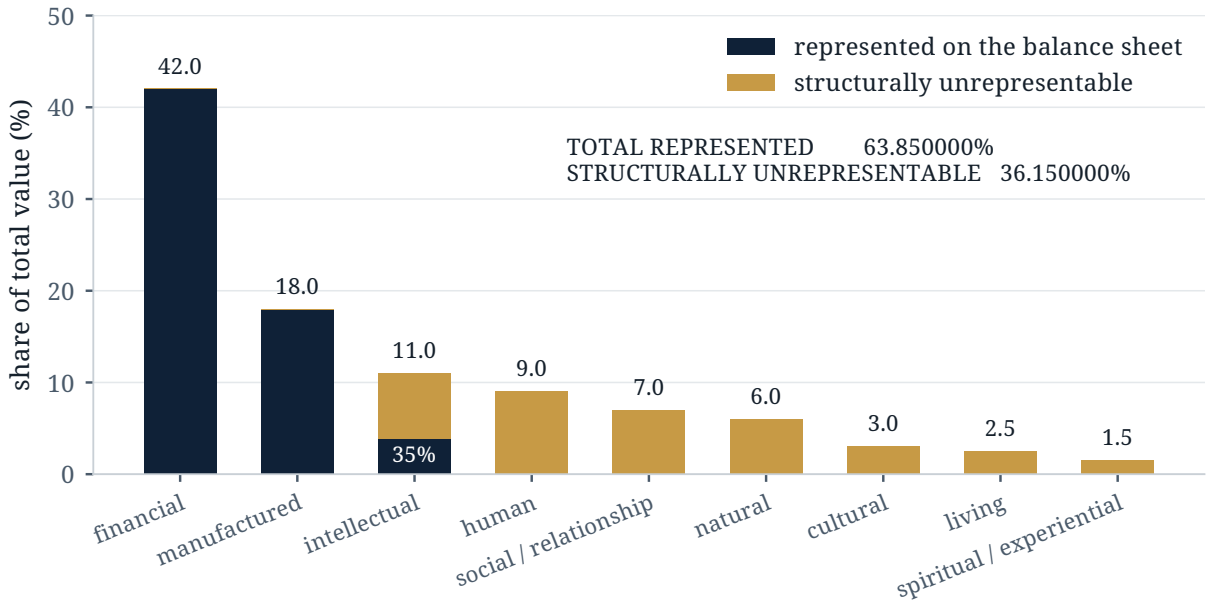


Figure 64.3: **The ledger balances exactly over 63.850000% of a constructed value vector and is silent about 36.150000% of it.** Nine forms of capital with stipulated shares; financial and manufactured fully recognised, intellectual at a stipulated 35% recognition fraction, the remaining six not representable as entries because an entry must carry an amount. CONSTRUCTED — the vector is stipulated for illustration and is not a measurement of any entity. Built by `figs/src/CCP-064_fig05.py`.

```
THE OFF-LEDGER REGISTER  (maintained alongside the statutory accounts)
-----------------------------------------------------------------------
1  Name each unmonetised asset the structure depends on.
2  Name the instrument that creates it (statute, treaty, licence,
   relationship) and the party that can withdraw it.
3  Record the observable that would signal withdrawal.
4  Record what it would cost to replace, if it can be replaced.
5  Review on the same cycle as the accounts, by the same people,
   in the same meeting -- and minute it.
```

Item 2 is what turns this from an essay into a control. **An unmonetised asset almost always has a nameable counterparty who can withdraw it, and naming that counterparty converts an invisible dependency into a monitorable one.** The register has **no standing in any accounting framework, is not a disclosure, and must not be described to any regulator or auditor as one**; it passes the disclosure test of Section 3.4, and disclosing it improves it, because a dependency register a supervisor could read is a dependency register the board has actually tested.

5.6 What the two historical ledgers could not have shown

Both institutions described in Section 4 kept complete books, and both were exposed to something their books could not express.

Individual loan exposures were, individually, fully recorded: every advance in a named account with a counter-entry in a partnership whose books balanced. Nothing hidden, nothing mis-transcribed. The concentration measurement Curzi [2026] obtains on such a book — 3.573981 effective borrowers on the naive count against 1.834862 effective blocs once two borrowers on the same side of the same conflict are recognised as one exposure, a 94.782 per cent overstatement — **is not derivable from the ledger at all.** Concentration is a relationship

between entries, and double entry records entries. A Herfindahl index is a function over a **partition** of the balances, and the partition — which counterparties are really one counterparty — is a judgement no posting convention supplies.

The same holds for the five-house structure and more sharply. Each house's books would show its own capital, exposures and currency position. **No house's books contained the number that was the point of the whole architecture** — the effective bloc count of Section 4.2 — because that number is a function of the legal environment *surrounding* the five sets of books, which no set of books can observe about itself.

quantity	recorded by double entry?	where it actually lives
an individual exposure	yes	the ledger
the sum of exposures	yes	the trial balance
the **concentration** of exposures	**no**	a partition chosen by judgement and applied from outside
the **independence** of the perimeters the exposures sit in	**no**	the legal environment surrounding the ledger
the **correlation** between two apparently distinct positions	**no**	a model imposed on the ledger from outside

Three noes, and each is something Curzi [2026] measures. **The instrument the corpus uses to measure resilience operates on inputs that the instrument every entity uses to report itself cannot produce.** That is not a criticism of either; it is a statement about which tool answers which question, and about the fact that no month-end close will ever prompt the second question. It also explains a pattern across both historical institutions: neither failed for lack of information about its own positions, and neither had any mechanism for converting position information into concentration information.

The corresponding failure mode is worth naming because it is common and cheap to commit: **a management pack that reports exposures by legal counterparty and calls it a concentration report has reported a sum and labelled it a relationship.**

5.7 Three proposed extensions, and which gap each closes

Three proposals are routinely conflated and address different gaps on different axes. Separating them is a precondition for Curzi [2026], which builds on the separation.

The time axis. Conventional accounting records a stock and a flow. Income is a rate, and a rate has its own rate of change, which the conventional statements do not record. The 1982 and 1989 momentum-accounting proposals add a third class of entry recording the causes of changes in that rate. Two stipulated revenue series make the gap concrete:

```
series A (growing)     100.0000  110.0000  121.0000  133.1000  146.4100
  rate of change                 +10.0000  +11.0000  +12.1000  +13.3100
series B (decaying)    180.0000  170.0000  160.0000  152.0000  146.4100
  rate of change                 -10.0000  -10.0000   -8.0000   -5.5900

period-5 income:   A = 146.4100   B = 146.4100   difference = 0.0000000000
period-5 change :  A =  +13.3100  B =   -5.5900  OPPOSITE SIGN
```

Two enterprises with identical reported income in the period, in opposite conditions, and the income statement cannot tell them apart. The information is derivable by anyone with five periods of history; it is not an account, is not posted, is not audited and is not part of the closure. The proposal is to make it one. **It has never been adopted, and this paper states that**

as a fact rather than a prediction. It is retained here because it is the only serious attempt on the time axis and because the name it occupies is the source of a persistent confusion.

Commitment without transmission. A hash tree commits to a set of records with a single value, and membership of any individual record is then provable by a path of sibling hashes whose length is logarithmic in the set size.

ledger entries	proof path (hashes)	proof bytes at 32 B	full ledger at 256 B/entry
1,000	10	320	256,000
1,000,000	20	640	256,000,000
1,000,000,000	30	960	256,000,000,000
1,000,000,000,000	40	1,280	256,000,000,000,000

At a billion entries the membership proof is 960 bytes against 256 gigabytes of ledger, a ratio of 1 to 266,666,667, and **doubling the ledger adds exactly one hash — 32 bytes — to the proof.** What matters here is not the compression: two parties can agree on a single short value and thereafter each prove to the other what their own record says, without either transmitting its books and without either being able to change a committed entry afterwards. **That is a shared commitment between entities, which is precisely the object double entry does not have.**

One caution, because it is the commonest overreach: a membership proof proves that a record is in the committed set and says nothing about whether the record is true. The fictitious 1,000 of Section 5.2 commits, proves and verifies exactly as well as a real one. **Commitment is a tamper-evidence property, not a truth property.**

The party axis. The 2005 signed-receipt proposal adds a third *signature* rather than a third book: a receipt recording a bilateral transaction, signed by an issuer and by both parties, held identically by all three. Each party still keeps conventional double-entry books against it; the proposal retains the method rather than replacing it. The invoice of Section 5.3 becomes impossible to book at two different rates without both parties visibly signing two different receipts.

```
parties   party-books   bilateral recs   signed receipts   disputable records
-------------------------------------------------------------------------------
    2           2               1                1             2  ->  0
    3           3               3                1             3  ->  0
    5           5              10                1             5  ->  0
   10          10              45                1            10  ->  0

time axis     third DIMENSION in ONE entity's books
              adds the rate of change of income
              closes: the undifferentiated-flow gap
              status: proposed, never adopted

party axis    third SIGNATURE on ONE record shared between entities
              adds a jointly held artefact
              closes: the two-books-no-shared-truth gap
              status: the conceptual ancestor of shared-ledger systems

NEITHER closes the monetisation gap. Nothing that operates on entries
can close a gap defined by what cannot become an entry.
```

The two proposals share a name and address different gaps on different axes; conflating them is the commonest error in the literature that cites both, and this separation is a binding

constraint on Curzi [2026] rather than a licence for it. The costs are also different in kind. The time axis costs a full third set of accounts and a convention for attributing rate changes to causes. The party axis costs key management, an issuer, and a governance question: **a shared record requires a shared party, and who that party is, under whose law, and what happens when it is compelled, is a jurisdictional question** — which is why it belongs in this series. Ostrom [1990] is the standing account of how a commonly held resource is governed without a single owner, and it is the right literature for that question; Star and Griesemer [1989] and Star and Ruhleder [1996] give the shared record its correct general name, a boundary object that must be simultaneously plastic enough for local use and robust enough to travel.

requirement	closed by	still open?
arithmetic closure	double entry	no
transcription completeness	double entry	no
existence	audit / confirmation	partially
valuation	audit / measurement	partially
agreement between parties	the party axis	addressable
the time dimension	the time axis	proposed, unadopted
coverage of unmonetised value	**nothing in use**	**OPEN**

Six rows addressed and one open, and **the open row is the one that matters most for this series, because everything the corpus measures lives in it.** A shared ledger does not and cannot close it: a shared ledger makes two parties agree about entries, and the monetisation gap is defined by what cannot become an entry. **No amount of agreement about a set closes a gap defined by what is outside the set, and any claim that a distributed ledger captures value a balance sheet cannot is a category error.**

6. Canonicalisation as an integrity control

Section 5 ended on two things double entry cannot do. This section is about a third limitation, and it is not a limitation of double entry at all. **It is a limitation of the alphabet.**

Double-entry bookkeeping assumes, without ever stating it, that a recorded symbol has one representation: an account number is that account number, a counterparty name is that counterparty, a currency code designates one currency. Every downstream control built on the ledger — deduplication, reconciliation, matching, sampling, confirmation, screening, consolidation — inherits that assumption absolutely. **The assumption was true while the character set had one code for each character, and it stopped being true when it stopped having one** [T1].

A universal coded character set necessarily contains many characters that render as the same shape, because many scripts independently inherited the same letterforms from the same ancestors. Latin capital A, Cyrillic capital A and Greek capital Alpha are three code points that most fonts draw identically. **That is not a defect in the standard; it is a true fact about the world's writing systems, correctly represented.** The consequence for a ledger is that the correspondence between what a field looks like and what it compares as is broken in both directions at once.

No standards document is cited in this section, because `bib/external.bib` contains none. Every technical designation named below is described from the source manuscript and is itemised in Appendix B with what specifically requires verification against the governing document. The one standard that is in the bibliography, International Organization for Standardization [2013], is cited where it is actually relied on.

6.1 Threat model and evidence tiers

The environment assumed is ordinary rather than exotic [T1]. Ledger fields are stored and compared as byte strings, or as text never canonicalised. At least one ingestion path — a counterparty file, a customer-facing form, a spreadsheet upload, an email-parsed remittance advice, a partner interface — accepts characters beyond the ASCII repertoire into a critical field. Downstream processes rely on exact equality. **Each of those three conditions is the default configuration of a great deal of production financial software and none is a misconfiguration in any conventional sense.**

No particular actor and no particular intent is assumed. The vulnerability admits accidental data-quality failure and deliberate manipulation with equal ease, and the measurements below cannot distinguish them, because the mechanism does not care. That agnosticism is deliberate and protective: everything in this section is a statement about a mechanism and nothing in it is a statement about anybody.

tier	meaning
T1	established fact — a published standard, a hashing property, a mathematical property
T2	measured by the deterministic demonstration and independently recomputed for this paper
T3	proposed design; not adopted, not in force, endorsed by no body
T4	a claim deliberately **not** made

The headline grades: the substitution vulnerability and the canonicalisation remedy are T1 and T2; the specification built on them in Curzi [2026] is T3; any linkage to macroeconomic measurement discrepancy is a bounded, testable hypothesis and never a demonstrated magnitude — T3 with explicit T4 disclaimers.

The entry points that matter are the ones that are not the front end, and every institution has more of them than its architecture diagram shows. One deserves a note of its own because it produces most of the accidental cases and involves no unusual behaviour by anybody: **text copied out of a rendered document routinely carries typographic substitutions the reader never chose** — non-breaking spaces, soft hyphens, directional quotation marks, ligatures, and in some pipelines full-width forms where the layout engine needed the advance width. The person pasting sees the characters they expect; the field now holds code points the validation rule was not written to consider. **A canonicalisation layer is worth exactly as much as its coverage of that list, which is why the rule is stated as a boundary rule and not a field rule: canonicalise at every ingestion boundary, not at the fields one happens to remember.**

6.2 The substitution class, enumerated

A critical ledger field is integrity-safe only if equal-looking values are equal-comparing values **and** equal-comparing values are equal-looking values. Both directions fail.

Homoglyph confusables [T1]. The core case: three code points, one shape, in the great majority of fonts. The relevant security standard defines a *skeleton* transformation mapping confusable characters onto a common representative, so that two strings are confusable when their skeletons are equal. The security literature is old — the homograph attack was published in 2002 and measured in the wild in 2006 — and both papers were about domain names rather than ledgers. **The domain-name world's response is the instructive part: it did not solve the problem by hoping resolvers would compare carefully, it defined a canonical wire form and rules about what may enter it.** That is the move made here.

Compatibility and full-width forms [T1]. Full-width variants of the ASCII digits and letters, inherited from East Asian legacy encodings, display as wide but unmistakable characters and carry different code points and different bytes. This member has a specific property: it is resolved by *compatibility* normalisation and not by *canonical* normalisation. **A system that normalises to the canonical composed form — which many do, correctly, for other reasons — has not addressed this at all.**

Normalisation mismatch [T1]. A precomposed accented character and the sequence of base plus combining accent are canonically equivalent and render identically, and they are different byte sequences. Two systems in one group, one normalising inbound text to the composed form and the other to the decomposed form, disagree on byte equality for every accented character in every field, forever, with no attacker involved and no error message. **This member produces the largest volume of purely accidental breakage and almost none of the deliberate kind.**

Combining marks [T1]. The general case of the preceding member: a base character followed by an arbitrary sequence of combining marks, some sequences having a precomposed form and some not. Section 6.5 shows this member behaving differently from the others under the remedy, and the difference is informative.

Zero-width and bidirectional controls [T1]. Characters defined to have no visible rendering, and explicit formatting controls whose purpose is to change the *order* in which characters are displayed relative to the order in which they are stored. This is the member with the highest-profile published treatment: a 2021 disclosure showed that bidirectional controls can make source code display in an order different from the order the compiler reads, so that a reviewer and a compiler see two different programs from one file, and it prompted coordinated changes across a large number of compilers and code-review tools. **The ledger analogue is direct: a reference, an amount or a counterparty name can display in one order and store in another, and the stored order is the one that reconciles.**

Charset confusion [T1]. This member is not about which characters are in the string. It is about the fact that a sequence of octets is not a string at all until something decides what encoding it is in, and different components in one pipeline can decide differently. Seventeen octets, every one of them plain printable ASCII, decode under one transformation format as a seventeen-character string and under another as an eight-character string. **Both readings are correct decodings under their respective standards; the octets never changed.**

```
octets                  2B 41 45 45 41 51 77 42 44 41 46 51 2D 30 30 34 32
decoded one way         "+AEEAQwBDAFQ-0042"      (17 characters)
decoded another way     "ACCT0042"               ( 8 characters)
```

This is the worse case and the reason is worth stating explicitly. In the homoglyph case the bytes differ and the meaning agrees; in the charset case the bytes agree and the meaning differs — so a byte-level hash, which is the detector for every other member of the class, certifies the two readings as identical. The remedy is therefore not canonicalisation of the string: it is refusing to accept octets whose encoding is not explicitly declared, and pinning the declaration at the boundary. **Canonicalisation cannot help, because canonicalisation operates on a string and the whole problem is that there is not yet a string.**

The withdrawn currency-code hazard [T1]. The last member has no character-encoding content at all. It is included because it has the same shape — one recorded symbol, two possible referents, no visible difference — and because it has a directly computable balance-sheet consequence. Currency codes are withdrawn when currencies cease to exist, and the numeric codes of withdrawn currencies persist in the standard's historical record and, far more consequentially, in the master files of systems built while those currencies were live. **A legacy system carrying a withdrawn-code balance, with a mapping table that has at some point been given a row sending the withdrawn code to a live currency, reports that balance as live money.** Nobody forged anything and nobody typed a wrong character; a withdrawn code was aliased to a live one, and the general ledger now says something untrue about how much settleable currency exists. **The specific numeric codes named in the source manuscript are contested**

even within the manuscript's own verification register and are therefore not asserted here; see Appendix B, item 9.

6.3 Why the failure is silent in both directions

Every member produces one of two outcomes.

The first is **byte-distinct, canonically equal**: two entries a human cannot tell apart, and that any auditor confirming them would confirm as the same, compare as different to a machine. Deduplication does not deduplicate them, reconciliation reports a break that is not a break, consolidation counts one entity twice, screening does not match the list entry it should have matched.

The second is **byte-equal, semantically distinct**: the charset member and the currency-code member. Here the machine says the two things are the same and they are not.

The class is dangerous precisely because the first outcome is invisible to people and the second is invisible to machines, and the two organs of control in every finance function are people and machines. A homoglyph substitution defeats the machine while presenting a clean face to the reviewer; a charset or code-aliasing confusion defeats the reviewer, who sees no anomaly, while the machine confirms the identity. **There is no observer in the standard control environment positioned to see either.**

The first outcome, computed rather than asserted, on four synthetic identifiers any reader would read as the same reference [T2]:

variant	chars	UTF-8 bytes	SHA-256 (leading)	canon(x)
clean, all ASCII	14	14	`6de95bd1c87f4ec2`	`ACCT-0042-CORP`
Cyrillic capital for Latin capital	14	15	`21300fa543adb7d6`	`ACCT-0042-CORP`
full-width digits	14	18	`4500cb7fe935f396`	`ACCT-0042-CORP`
zero-width space inserted	15	17	`95a611acf502cfc7`	`ACCT-0042-CORP`
right-to-left override inserted	15	17	`b39d1bc9946e6dd2`	`ACCT-0042-CORP`

Five rows, five distinct digests, five byte lengths, one canonical form, and one thing a person would say if shown the printed page. **The cryptographic hash is doing exactly what it was designed to do and is, for this purpose, worse than useless: it certifies with mathematical confidence that five records are distinct when a human reader and an auditor's confirmation would all say they are one.** SYNTHETIC data throughout.

6.4 The canonical pipeline and the corpus

The pipeline under test [T3]:

$$\text{canon}(x) = \text{restrict_ASCII}(\text{confusable_fold}(\text{NFKC}(x)))$$

Compatibility normalisation resolves compatibility and full-width forms; `confusable_fold` maps homoglyphs onto a single skeleton and removes default-ignorable characters; `restrict_ASCII` enforces the permitted character set for the field and **quarantines rather than silently stores** any residual out-of-set character. Alongside it, and separately, a byte-level SHA-256 computed over the record's raw octets and **not** over its canonical form. **The distinction between the two is the whole detector** and Section 6.6 measures what it buys.

The corpus: 500 synthetic identifier pairs, each a clean value and a confusable twin. The class mix is a **stipulated design parameter** chosen to reproduce the source manuscript's reported per-method detection rates — 297 homoglyph twins, 148 compatibility-form twins, 55

invisible-character twins. **Five of the 297 homoglyph assignments fall on digit-only account numbers where no Latin letter is available to substitute**, so those five twins are byte-identical no-ops rather than manipulations. That detail is not a rounding artefact; it is where the 99 per cent comes from, and a reader reproducing the number needs it. Combining-mark twins are measured separately. All data are synthetic, the seed is fixed, and the harness is deterministic and reruns identically.

The harness was rebuilt from the source manuscript's description rather than obtained, and two of the five results depended on corpus-design parameters that had to be reconstructed before they agreed. This is recorded in Appendix A and in Section 9.

6.5 Results: five measurements (H6 supported)

E1 — silent duplication [T2]. For each pair: are the two values byte-distinct, are their digests distinct, are their canonical forms equal. A pair passing all three is a *silent duplicate*.

```
pairs                                                  500
byte- AND digest-distinct yet canonically equal        495
rate                                                 99.0%
no-op twins (no substitutable character present)         5
```

Ninety-nine per cent of confusable substitutions produce a record that is byte-distinct, hash-distinct and canonically identical. The remaining one per cent is not a detection failure and not a residual risk; it is the set of fields the substitution could not be applied to at all, which is the honest denominator.

E2 — detection [T2]. Five checks, each run against the 495 manipulated values for a true-positive rate and against the 500 clean values for a false-positive rate. The clean values are pure ASCII, which is what makes the false-positive column meaningful.

The *skeleton* of a value is confusable_fold(NFKC(x)) — the first two stages of the pipeline, without the restriction step.

check	what it asks	TPR	FPR
byte-vs-skeleton	does the value differ from its skeleton?	100.00%	0.00%
non-ASCII	does x contain a code point above U+007F?	100.00%	0.00%
confusable-only	does x contain a mapped confusable letter?	58.99%	0.00%
normalisation-mismatch	does x differ from its canonical composed form?	0.00%	0.00%
invisible-character	does x contain a default-ignorable code point?	11.11%	0.00%

The skeleton and non-ASCII checks flag every manipulated value at zero false positives on clean ASCII fields. Three observations about the rest of the table matter more than that headline.

The confusable-only check reaches 58.99 per cent — exactly $292/495$ — and the reason it does not reach 100 per cent is structural rather than incidental: it is blind to the compatibility-form member, which is the normalisation step's job and not the confusables table's. **A confusables check deployed without a normalisation layer in front of it will miss two manipulated**

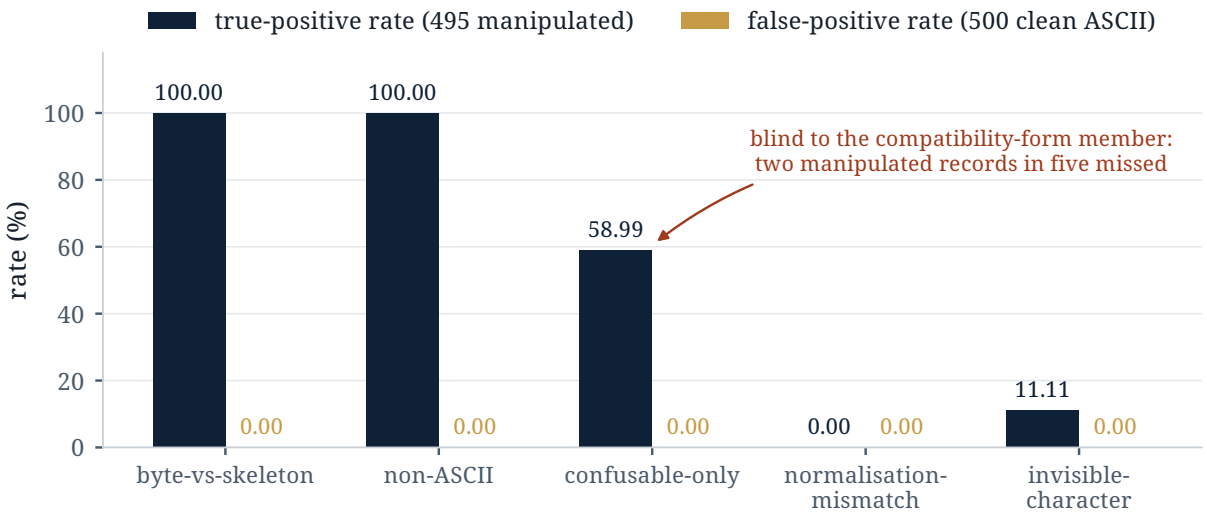


Figure 64.4: **No single check covers the substitution class; the two that do are the two that compose normalisation with confusable folding.** True-positive rate over the 495 manipulated values and false-positive rate over the 500 clean ASCII values. The confusable-only check reaches 58.99% — exactly 292/495 — because it is blind to the compatibility-form member. SYNTHETIC corpus, deterministic harness, fixed seed. Built by `figs/src/CCP-064_fig04.py`.

records in five and will report itself as working. That is the most operationally dangerous line in the table.

The normalisation-mismatch check scores zero here, and that is correct rather than broken: none of the three manipulation classes in this corpus perturbs the canonical composed form. It earns its place against the combining-mark class, where it scores 100 per cent. The invisible-character check scores 11.11 per cent, which is exactly $55/495$ — a single-purpose check detecting its single purpose. **The point of the table is that no single check covers the class, and that the two which do are the two composing normalisation with confusable folding.**

Zero false positives on every check is **a property of the corpus** — clean fields that are pure ASCII — and is therefore a claim about *critical machine fields*, not about names or memos. **A non-ASCII check run against a customer-name field in any real institution will produce false positives at an enormous rate, because a great many customers legitimately have non-ASCII names.** Section 7.2 treats that scoping problem as the most serious objection to the whole approach.

E3 — canonicalisation [T2].

manipulated entries collapsing onto the clean canonical form	495 / 495 = 100.0%
canon(canon(x)) == canon(x) for every value in the corpus	True (idempotent)
distinct raw byte-strings across the corpus	995
distinct canonical forms	500
true distinct underlying values	500

Canonicalisation collapses one hundred per cent of the manipulated entries onto their clean form, idempotently, and the count of distinct values falls from 995 to the true 500. The idempotence matters more than it looks: **a canonicalisation that is not idempotent cannot be applied at more than one point in a pipeline without changing answers depending on how many times it ran**, which makes it unusable in any architecture with more than one ingestion path.

The combining-mark class, measured separately [T2], behaves differently, **and this is a finding of the rebuild rather than a reproduction of a reported figure.** Of 98 combining-mark twins, 100 per cent were **quarantined** by `restrict_ASCII` rather than collapsed onto the clean form; the normalisation-mismatch check flagged 100 per cent of them and the confusable-only check flagged none. The reason is mechanical: compatibility normalisation recomposes base-plus-combining-accent into a precomposed accented character, which is not ASCII, so the restriction step refuses it. **For a critical machine field that is the correct outcome and arguably better than collapsing** — the entry is rejected at the gate with an explicit reason rather than silently normalised into something the submitter did not send. **For a free-text field it would be data destruction. The class member does not have one right answer; it has one right answer per field scope**, which is the argument for scoping.

E4 — reconciliation [T2]. Two synthetic ledgers of 500 paired entries with identical underlying values and a controlled proportion of one side's critical fields manipulated. Deterministic: exactly the stated proportion is manipulated rather than drawn, so the sweep is exact.

manipulation rate	false-break rate, byte equality	false-break rate, canonical equality
0%	0.0%	0.0%
10%	10.0%	0.0%
20%	20.0%	0.0%
30%	30.0%	0.0%
40%	40.0%	0.0%
50%	50.0%	0.0%
60%	60.0%	0.0%
80%	80.0%	0.0%
100%	100.0%	0.0%

At a 40 per cent manipulation rate the false reconciliation-break rate goes from 40 per cent to 0 per cent. The full sweep shows the shape, which is the more useful result: byte-equality false breaks scale one-for-one with manipulation, and canonical equality holds flat at zero across the entire range including the degenerate case where every entry is manipulated.

The operational reading is specific. **A reconciliation function chasing breaks spends its scarcest resource — experienced human attention — at a rate directly proportional to the encoding noise in its inputs, and it has no way to tell an encoding break from a real one until somebody opens it.** The remedy does not merely reduce false breaks; it makes the remaining break queue mean something.

E5 — the withdrawn-code hazard [T2]. A synthetic three-account legacy extract: two accounts hold a live currency code, one holds a withdrawn code, and the system's mapping table aliases the withdrawn code to the live one.

```
correct live balance                          8000
aliased live balance                         10000
inflation                                   +25.0%
resolution-class balance, segregated          2000
```

Aliasing a withdrawn numeric code to a live currency inflates the recorded live balance from 8000 to 10000, a phantom inflation of 25 per cent. Assigning the withdrawn code to an explicit resolution class instead — never to a live settlement currency — restores the live balance to 8000 and leaves the 2000 in a designated class where it is visible, non-settleable and awaiting an audited discharge.

Three things about this experiment must be said. It is arithmetic, not cryptography. **The 25 per cent is a property of the synthetic 8000/2000 ratio and is not a universal constant; anyone quoting it as a general magnitude has misread it** [T4]. What generalises is the *sign* and the

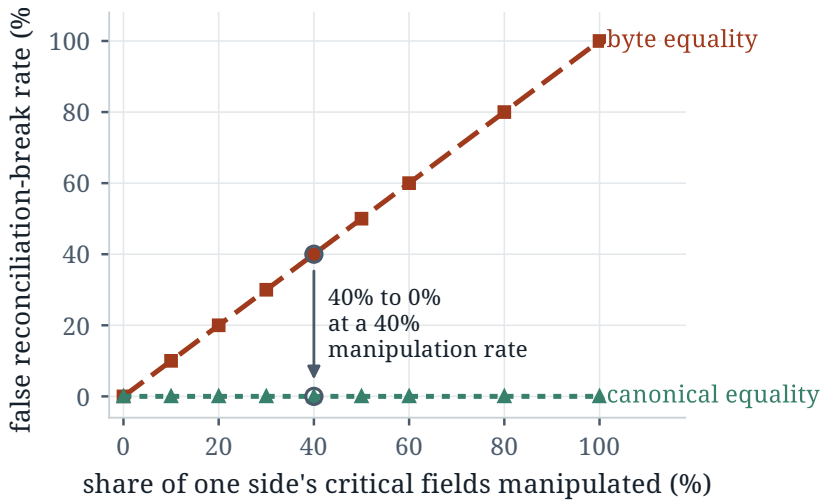


Figure 64.5: **Canonical equality holds the false-break rate at zero across the whole sweep while byte equality tracks the manipulation rate one-for-one.** Two synthetic ledgers of 500 paired entries; the proportion manipulated is set deterministically rather than drawn, so every point is exact. The 40%-to-0% reading at a 40% manipulation rate is one point on a flat line, not the result. Built by `figs/src/CCP-064_fig03.py`.

mechanism: aliasing a withdrawn code to a live one always inflates, always in the direction of more apparent money, and always without any entry being posted. And it is the one member of the class with no character-encoding content whatsoever, **which is a useful reminder that the underlying failure is representational ambiguity rather than anything specific to character encoding.**

6.6 Why the hash level is not a detail

One further measurement, because it is the hinge of the design [T2]. Take the clean value and its Cyrillic twin:

```
canonical forms equal?                                  True
BYTE-level digest differs?        True    6de95bd1c87f  vs  21300fa543ad
CHARACTER-level digest differs?   False   6de95bd1c87f  vs  6de95bd1c87f
```

The character-level digest — computed over canon(x) — is identical for the clean value and the manipulated one, which is exactly what makes it useless as evidence. **Identical canonical form with differing byte-hash is the detector; a character-level hash destroys precisely the evidence that distinguishes benign normalisation from injected substitution.** Hash the canonical form and the system cannot tell anyone it was attacked. Hash the octets and store the canonical form beside them, and every substitution leaves a permanent, comparable, machine-checkable trace. **This is the same content-addressing property Curzi [2026] builds on: one byte's difference in a settlement object yields a different identifier, while** canon **folds a homoglyph account string onto the same skeleton as the clean one.**

6.7 The screening arithmetic

A **SYNTHETIC** watchlist of ten entity names, none corresponding to any real person or organisation. Twenty probes are generated by applying a single homoglyph substitution or a single zero-width insertion to each name, and both a naive exact-match screen and a canonical-match screen are run [T2].

```
synthetic watchlist entries                                  10
effectively-substituted probes                               20
exact-match screen: probes matched                 0 / 20  =   0%
canonical-match screen: probes matched            20 / 20  = 100%
canonical-match screen: clean names still matched 10 / 10  = 100%
```

A single-character substitution defeats an exact-match screen on every probe, and folding the confusables before comparison recovers every one of them without losing a single true match.

The size of the change should be kept in proportion in both directions. It is one transformation applied to both sides of an existing comparison; it requires no new data, no new supplier, no model, no training set and no change to any list. **It also does not make a screening system good.** A screening system with poor alias coverage, no transliteration handling and a badly tuned threshold remains bad after canonicalisation. **This closes one specific hole, completely, for the cost of a library call at the field boundary.**

7. Discussion

7.1 What follows

A concentration measurement inherits every defect of the identity resolution beneath it. The effective count of Curzi [2026] is a function over a partition of exposures. Section 5.6 shows the ledger cannot supply the partition; Section 6 shows that the resolution underneath the partition is string comparison and that string comparison has a measurable blind spot. **Count one counterparty as two and the measured concentration falls without anything real changing.** The direction is the concerning part: **an encoding defect in a counterparty master file biases the resilience instrument towards reporting resilience that is not there.** That is the same structural error Curzi [2026] identifies at the rail layer, where a naive correspondent count of 3.306878 collapses to $N_{\text{rail}} = 1.341819$ once flows are decomposed by currency, and Curzi [2026] identifies at the assembled-structure layer.

A structure built or reasoned about before a coordination mechanism existed will report a diversification that has since been cancelled, and the cancellation leaves no trace in the entity list. Section 4.3 is that statement measured: nothing about the five cities, the five legal families or the ownership changed, and 0.469061 of h went. The monitoring implication is specific and cheap. The thing to watch is not any jurisdiction's behaviour but the perimeter's boundary, and the standing question per node is one sentence: **has this jurisdiction acquired a mechanism by which another state's decision becomes enforcement here without a separate decision taken here?**

The source manuscript answers that question node by node for a named six-jurisdiction structure, with named listing cycles, named administrative orders and named unilateral measures. **Those answers are not reproduced here.** `bib/external.bib` contains no jurisdiction-specific statutes, regulations, registry documents or case law, and the coverage file records this as a known gap; attaching an approximate citation to a jurisdiction-specific regulatory fact would be worse than leaving it uncited. Following the de-branding convention of Curzi [2026], the mechanism classes are given without the jurisdictions, and the source's named instances are itemised in Appendix B.

node	mechanism class to watch	reading
A	evaluation and listing cycle under an intergovernmental standard-setter	listed and later delisted; the next evaluation round is the observable

node	mechanism class to watch	reading
B	a revenue authority's characterisation of residence, movable by administrative order	the characterisation has already moved once by order
C	full participation in every mechanism class of Section 3.2	no independence to lose
D	supranational listing history and substance requirements	inside the perimeter
E	participation in automatic exchange with a domestic regulator	partial
F	unilateral measures adopted without a Security Council resolution	eroding bloc independence

Nothing in that table is a jurisdiction going bad. Every row is a perimeter growing.

A widening gap between what an institution says about its own foundations and what its law actually does is a leading indicator, observable years in advance and free to monitor. Section 4.7 measures that gap at 571 years wide on the longest available instance, with 18 years between the last major judicial application and the issuing institution's repudiation. The corresponding diligence workstream is a records search rather than an opinion: identify the apex doctrine for the situs and name it; identify whether a competing non-derivative chain has been recognised by a superior court there; identify whether the parcel is within, adjacent to or subject to a pending or historical claim; price it through title-insurance availability, exclusions and premium loading; record the answer in writing with the observable that would change it; and repeat on a defined cycle. **Apex doctrines move on a scale of decades, and a 200-year-old rule applied in 2005 and repudiated by its author in 2023 is not settled background.** This paper makes no claim about the availability, exclusions or pricing of title insurance in any jurisdiction, and none about the current status of any claim.

The remediation at the character layer is proportionate and short. Canonicalise critical machine fields at every ingestion boundary. Store the byte-level digest of the raw octets beside each record. Alert on identical canonical form with differing byte-hash. Pin the declared charset at every file boundary and reject octets with no declaration. Scope ASCII restriction to critical machine fields and leave free text alone under restriction levels. Assign every withdrawn currency code to a resolution class rather than to a live currency. **None of this is expensive and all of it is testable.**

7.2 What does not follow

Canonicalisation defeats encoding-level evasion and does nothing whatever about semantic fraud: a correctly encoded false entry is entirely unaffected [T1/T4]. Every number in Section 6 measures a representational property; none measures truthfulness. The consequence for how the contribution may be pitched is hard and is worth paying, because the alternative destroys the claim.

The source manuscript audits nine well-documented corporate and market failures against the question *would encoding canonicalisation have caught it?* **The answer is no, eight times.** The eight turn on accounting treatment and revenue recognition; on forged confirmation documents; on fabricated bank statements; on invented account statements describing trades that never occurred; on collusion about which number to submit; on a monitoring and escalation environment that did not function at the scale of the throughput; on cash purportedly held in escrow that was not there; and on commingling of customer assets together with

an accounting configuration that concealed a liability. **Every one of those is semantic or procedural. Every number in them was correctly encoded and every name correctly spelled.** The failures are described here at the level of mechanism only; the named matters, their dates, quanta and outcomes are itemised in Appendix B and none is cited, because the bibliography contains no case reports, no regulatory findings and no contemporaneous reporting. **A technical evaluator reading a submission that implies encoding hygiene would have caught any of them stops reading at that sentence, and is right to** [T4].

There is one genuine in-class instance and it is worth more than the eight, because it is documented and it is the attack pattern at scale: an entity incorporated under a name differing from a real and materially similar entity's name, to which funds were directed in circumstances where the documentation and a reasonable reader would have understood the other. **That is a confusable-identifier attack executed at the entity layer: an identifier constructed so that a control comparing it to the expected counterparty would pass.** The matter is described from the source manuscript, the entity names are not reproduced, and the case references are itemised in Appendix B as requiring precise sourcing or removal.

And here is the part that must be stated rather than glossed, because glossing it is the same error as claiming the other eight. The names in that instance were ordinary ASCII: they differ by a real, visible word, and no character substitution was involved. What the case demonstrates is that the *pattern* — choose an identifier a control will mistake for another identifier — works, is attractive to sophisticated parties, and survives contact with institutional diligence at very large scale. **What a universal character set changes is the economics of the pattern.** A confusable name in ASCII must survive a human reading it beside the expected name and can be caught by an alert analyst. A homoglyph substitution does not have to survive that, because there is nothing to see. The ASCII instance needed a plausible story about why an extra word was there; a Cyrillic capital needs no story at all.

The pitch is therefore narrower than the material invites, and the narrow version is the defensible one: exact-match string comparison is the control model underneath a set of financial controls that matter enormously, it has a measurable blind spot, the blind spot is cheap to close, and it sits underneath watchlist screening, beneficial-ownership resolution, cross-border party matching [International Organization for Standardization, 2013] and deduplication at scale. **Structured party identification confers format, not canonical form**, and two systems in two jurisdictions with two different normalisation defaults will disagree about identity in exactly the manner of the E4 sweep.

Over-canonicalisation is the strongest objection and the mitigation is scoping and only scoping [T1]. Aggressive folding destroys legitimate data: a customer whose name is written in a non-Latin script has a name, and ASCII restriction applied to it produces either a quarantine or a mutilation; a memo field in a language with diacritics loses them. **A canonicalisation standard that is mis-scoped is not a control improvement, it is a data-destruction programme with a compliance justification, and it will be correctly resisted by every business line it touches.** ASCII restriction applies to *critical machine fields* — identifiers, account and transaction numbers, amounts, currency and action codes — where the permitted character set is already narrow by business rule and where nobody has ever legitimately needed a combining mark. Free-text fields keep the full repertoire, store raw and canonical forms together, and are governed by restriction levels rather than by ASCII folding.

A production deployment needs four things the demonstration does not have, and the distance is worth measuring so that nobody underestimates it and nobody overestimates it either. The confusables data must be consumed as a maintained data file rather than a hand-written table, with the character-set version treated as a configuration item under change control, because a confusables table that drifts out of date is a control silently losing coverage. Script policy is a ladder rather than a setting, and the right rung differs by field and by market — a business decision with a compliance consequence, not a job a library performs. **The alert path is the item most likely to be underbuilt**: in a production master file with a decade of history and several migrations behind it, the first run produces a very large number of alerts, most of them historical normalisation artefacts, so the design question is not how to detect

but how to burn down the backlog once, establish a clean baseline and then alert only on new occurrences. **A deployment that skips the baseline generates an unreadable queue on day one and is switched off in the second week.** And the boundary declaration for the charset member is separate work that nothing in the canonical pipeline addresses.

Four things would falsify or substantially weaken this section's argument and none has been tested. If a survey of production ledger and counterparty-master data found that non-ASCII code points essentially never occur in critical machine fields, the operational case collapses and the correct conclusion is a theoretical vulnerability with no realised exposure. If mature screening engines already fold confusables before computing match scores — a configuration and version question, testable in an afternoon per engine — the screening argument is already solved in the products and the contribution reduces to reconciliation and deduplication. If the false-positive rate of canonical matching against a real, large watchlist is materially above zero — plausible, because folding collapses an identifier space and collapsing an identifier space creates collisions — the remedy is a real trade rather than a free one and needs pricing. **The zero false-positive figure of E2 is measured against clean ASCII critical fields and does not transfer to a name estate; anyone reading it as a general property has read it wrongly.** And if the alert stream from the byte-hash detector were dominated so completely by legacy artefacts that no baseline could be established, the detector is a good idea that does not survive contact with real data.

A group receiving alerts has created a record of those alerts, and that record is discoverable. That is not a reason to avoid implementing the detector. It is a reason to decide, before implementation, what the escalation path for an alert is and who owns it, because **an unactioned alert log is materially worse than no alert log.**

7.3 Bounds carried from elsewhere in the corpus

Four results elsewhere in the corpus bound what this paper may claim, and they are stated as bounds rather than as support.

Curzi [2026] assigns the corpus's central identification, the application of the surplus functional to a life, and the reading of a day's surface as a leading indicator to the **postulate** drawer, and states that the third is a wager whose misfire column exists because wagers can lose. **Section 7.1's leading-indicator claim about a widening legitimacy gap is in that drawer**, and it is a wager rather than a measurement.

Curzi [2026] refuses the certificate its own contraction theorem appears to offer: a human life is not a metric space, nobody has measured the relevant modulus, and the theorem transfers as shape rather than as guarantee; and the feeling of recognition on arrival proves nothing, since the machinery that rehearses a future retrofits a past with equal fluency. **That is the correct treatment of Section 4.7's trajectory reading**, which is a shape and not a forecast.

Curzi [2026] declines the inflated reading of its own borrowed physics, requiring that a user be able to distinguish a theorem, an ordinary observer effect and the framework's own conjecture in three separate sentences on demand. **Section 8.6 discharges that requirement for every borrowing in this paper.**

Curzi [2026] establishes that repetition buys arrival time and nothing else — it cannot move a destination — and that a tradition's canonical repetition count closes under half the gap at the corpus's own working decay rate. **The bearing on this paper is direct: no amount of re-running the reconciliation closes the two-books gap**, because the gap is a property of the fixed point rather than of the number of passes.

Two further results are parallel rather than bounding. Curzi [2026x] measures a receiver-limited channel capacity into a single cell of order 70 to 106 bit/s, delivering a 40-bit address in about half a second, and shows that adequate bandwidth into an absent decoder delivers nothing; **the structural parallel to Section 6 is exact — a correctly delivered symbol that the receiver resolves to the wrong referent is the same failure**, and Curzi [2026] makes the same point by building an entire deterministic pipeline and stating that the one missing link is the carriage of sequence over force. Curzi [2026] is the sharpest available warning against this

paper's own instruments: a frequency register with no time term cannot distinguish a mood from a trait, and the cleanest-looking encoding of a framing returns the system's maximum surplus for every subject and therefore measures nothing. **An index that returns the same answer for every input measures nothing, and an index that returns a better answer because a record was duplicated measures the duplication** — which is Section 7.1 in the corpus's own terms.

Finally, Curzi [2026t] demotes the corpus's own moment equation at the scale-type gate: the two quantities in it are not instrumented on a common ratio scale with a shared true zero, so the magnitude of their ratio is withdrawn as a measured quantity. **Section 8.5 applies the same gate to this paper's inputs and reaches the same verdict for the same reason.**

8. Verification and rigour audit

The results are decomposed into postulates that can be checked one at a time. Each is atomic, each carries a status, and each chain names what it consumes. **If dropping a postulate leaves the result standing, it is not listed.**

8.1 Dimensional audit

symbol	quantity	dimension	unit	status
s_i, w_i	exposure or capital share	$[1]$	fraction of one	defined; CONSTRUCTED here
HHI	sum of squared shares	$[1]$	dimensionless index	derived
$N = 1/\text{HHI}$	effective count	$[\#]$	equally weighted units	derived; Hill number of order 2
h	surplus reading	$[1]$	**no operational unit — the corpus supplies none**	derived from η, N and a stipulated floor
η	input coefficient	$[1]$	dimensionless	assumed (0.35, fixed in advance)
λ	counterparty hazard	$[T^{-1}]$	per day	assumed (0.0015)
t	exposure window	$[T]$	days	assumed (24, 8)
$P = 1 - e^{-\lambda t}$	loss probability	$[1]$	probability	derived; λt dimensionless, so the exponent balances
A, L, E	balance-sheet aggregates	$[M]$	one unit of account	assumed (stipulated)
days-sales-outstanding	receivables over revenue per day	$[T]$	days	derived; $[M]/[M\,T^{-1}]$
n	reporting entities, ledger entries, parties	$[\#]$	counts	assumed

symbol	quantity	dimension	unit	status
$n(n-1)/2$	bilateral relationships	[#]	counts	derived, exact combinatorics
proof path length	sibling hashes to a root	[#]	hashes	derived, $\log_2$ of a count
capital share of value	nine-capital vector entry	[1]	fraction of one	assumed (CONSTRUCTED)
TPR, FPR, break rate	detection and false-break rates	[1]	dimensionless fractions	derived from exact counts
digest	SHA-256 over octets	—	256 bits	defined

Does the surplus equation balance? $h = \eta \ln N - \delta_{\text{floor}}$ requires η dimensionless, $\ln N$ dimensionless (a logarithm of a pure count), and the floor dimensionless. It balances. **What the audit records, and what the corpus is silent on, is that h has no operational unit at all**: it is a dimensionless reading whose zero is set by a stipulated maintenance floor, and a difference of 0.469061 in it is a difference in a stipulated index, not in any measured quantity. **A dimensionless placeholder is honest; an invented unit would not be.**

Does the hazard model balance? λt is $[T^{-1}][T] = [1]$, so the exponent is dimensionless and P is a probability. The discrete alternative $1 - (1-\lambda)^t$ requires λ to be a per-step probability rather than a rate; **the two are different models and Section 8.11 records that the source is silent about which it used.**

Does the ledger arithmetic balance? $A - (L + E)$ is a difference of aggregates in one unit of account and is therefore in that unit; it is reported as exactly zero, which is a statement about the unit, not about the world. **Every currency figure in Sections 5 and 6 is in a stipulated unit of account with no currency attached beyond the labels used for legibility.**

8.2 Significant figures and epistemic status

- **The concentration figures are exact rational arithmetic on stipulated shares.** $1/0.2618 = 3.819710\ldots$ and $1/0.36 = 2.777778\ldots$ are exact to the digits shown. **They are exact consequences of a constructed share vector and are therefore not measurements of anything**; no uncertainty is quoted because there is no estimate to attach one to.
- h **carries six decimal places because η and the floor are exact stipulations.** Reporting -0.030939 does not imply that anything was measured to six places.
- **The posterior is an exact rational**: $1/1366 = 0.000732064\ldots$, printed as 0.000732. Its inputs are stipulated likelihoods, so the six figures are exact arithmetic on assumptions and **the posterior is not a probability estimated from data.**
- **0.460829% is exact**: $500/108500$. **66.6667% is exact**: $16/24$. **58.99% is exact**: $292/495$. **11.11% is exact**: $55/495$. **99.0% is exact**: $495/500$. Each is a ratio of two exact counts in a constructed corpus.
- **0.0% in the canonical-equality column is exact, not a rounded small number.** It is a property of a mapping that sends every manipulated value to its clean form, and no sampling is involved.
- **The 25% of E5 is a property of the stipulated 8000/2000 ratio and generalises only in sign.**
- **The transit figures are exact under a stated model on stipulated inputs.** 3.5360% is $1 - e^{-0.036}$ to four places; the fifth place differs under the discrete model.
- **The 63.850000% completeness figure carries six decimal places because the input vector is stipulated to three.** $60 + 0.35 \times 11 = 63.85$ exactly. **The precision is a property of the stipulation and not of any observation**, and quoting it to six places would be false precision on a measured quantity.

- **The Merkle proof lengths are ceilings of base-two logarithms of exact counts** and are exact.

8.3 Limits and boundary behaviour

model	limit	behaviour	reading
effective count	one unit holds share 1	$N \to 1$	maximum concentration; $h \to -0.500000$
effective count	n equal shares	$N \to n$ exactly	the numbers-equivalent property
effective count	share vector fixed, bloc partition to the finest	$N \to$ naive	the naive count is the zero-correlation corner, not a neutral default
surplus	$N \to \infty$	$h \to \infty$ logarithmically	**unbounded, and this is a domain limit rather than a result**: the form supplies no cost of fragmentation
surplus	$\eta \to 0$	$h \to -0.500000$ for all N	the instrument stops discriminating; η carries the whole reading
hazard	$\lambda t \to 0$	$P \to \lambda t$	the linear regime; the two hazard models agree to first order
hazard	$\lambda t \to \infty$	$P \to 1$	loss becomes certain; the window is the whole quantity
reconciliation	$n = 1$	$n(n-1)/2 = 0$	one entity has nothing to reconcile against; the shared register still counts 1
reconciliation	$n \to \infty$	ratio $\to \infty$	the shared record's advantage is unbounded
completeness	recognition fraction $\to 1$ on every capital	represented $\to 100\%$	**a pathology, not an ideal**: full recognition of unmeasurable items destroys the closure property that is the method's only guarantee
completeness	recognition fraction $\to 0$	represented $\to 0\%$	the ledger becomes a record of nothing; the observed 63.85% sits between two degenerate corners

model	limit	behaviour	reading
canonicalisation	manipulation rate $\to 1$	canonical false-break rate stays 0.0%	the flat line is the result; the 40% reading is one point on it
canonicalisation	permitted set $\to$ full repertoire	restriction step becomes vacuous	the remedy's strength is exactly the narrowness of the field's permitted set
canonicalisation	permitted set $\to$ one character	every value quarantined	over-canonicalisation in its limiting form; Section 7.2

The completeness limit is the informative one. Both corners are degenerate, which is why the gap is structural: the method cannot be improved towards full coverage without losing the property that makes it a method.

8.4 Invariance

The effective count is invariant under relabelling of units and under any permutation of the share vector, and it is *not* invariant under repartition — which is the entire content of Section 4.2 and the reason the bloc call must be published with its observable. **It is invariant under rescaling of the total exposure**, since it depends only on shares.

The posterior odds are invariant under any monotone rescaling that preserves each likelihood ratio, which is why the paper claims the direction and the order of magnitude rather than the value.

The canonical form is invariant under any of the substitutions in the class by construction, and that is a definition rather than a finding. **The byte digest is invariant under nothing**, which is exactly why it is the detector. **The false-break rate under canonical equality is invariant under the manipulation rate over the whole** $[0, 1]$ **interval**, which is the strongest invariance statement in the paper and is a property of the mapping.

Tested at how many scales? The invariances are algebraic and were checked at exactly the parameter sets reported: one share vector under five partitions, one corpus of 500 pairs, one sweep of nine manipulation rates, one watchlist of ten names. **No empirical scale-freedom is claimed anywhere in this paper, and no scaling exponent is estimated**, so the procedure of Clauset et al. [2009] was not run. The one exponent that appears — the 2 in $n(n-1)/2$ — is **exact combinatorics defined by the construction, not fitted**. Any restatement of this paper that reaches for a heavy-tail claim about holdings, share distributions or break volumes must run that procedure first, and must read Broido and Clauset [2019] and Stumpf and Porter [2012] on how weakly such structure is generally supported.

8.5 Measurement-theoretic status of the inputs

The share vectors are ratio-scale by construction and not by measurement. They are stipulated proportions of a stipulated total, so ratios of them are meaningful within the construction and carry no empirical content whatever.

The nine-capital vector is the paper's weakest input on this gate and it is the one to attack first. Assigning "social capital" a share of 7.000 per cent of total value presupposes that the nine forms are commensurable on a common ratio scale with a shared zero, and **nothing in this paper or in its sources establishes that.** Michell [1997] and Michell [1999] are the strongest objection to the whole framework — that a discipline routinely assumes quantitative structure without testing for it — and the objection lands squarely here. **The formal test**

that would answer it is conjoint measurement [Luce and Tukey, 1964], whose conditions establish whether an attribute has additive quantitative structure at all; the conditions were not tested, here or in the source, and the ratio-scale treatment of the nine-capital vector is therefore provisional. Krantz et al. [1971] sets out what must hold before a quantity can be said to exist and Stevens [1946] is the scale-type taxonomy the audit uses. **The honest reading is that 63.850000 per cent is an exact arithmetic consequence of an unlicensed set of numbers**, and its value as a result is the direction — a large fraction is unrepresentable — rather than the magnitude.

The evidence tiers of Section 6.1 are an ordinal scale and are not averaged anywhere. The detection and false-break rates are ratios of exact counts on a constructed corpus and are ratio-scale within that construction.

This paper inherits the corpus-wide scalar-reduction limitation rather than restating it. Reducing nine jointly varying forms of value to nine shares and then to one representability fraction, and reducing a jurisdictional structure to one effective count and one h, are instances of the habit Curzi [2026] names once as a single limitation for the whole corpus. **The limitation is inherited here in full and is not met separately.**

8.6 Analogy discipline

borrowing	category	justification
exposure as a portfolio; variance under a correlation matrix [Markowitz, 1952]	**1 — derivation**	the bloc question genuinely is whether a common factor acts across nominally independent units; the premises are satisfied and the identity is exact
effective count as a numbers-equivalent [Hill, 1973, Jost, 2006, Adelman, 1969]	**1 — derivation**	$1/\text{HHI}$ *is* the Hill number of order 2; nothing is transferred by analogy
similarity-sensitive diversity [Leinster and Cobbold, 2012, Rao, 1982]	**1 — derivation**	the block-of-ones matrix is a special case of the cited general form
Shannon entropy in bits for a proof-path length [Shannon, 1948, Hartley, 1928, Cover and Thomas, 2006]	**1 — derivation**	a membership proof over an equiprobable set of size n is exactly Hartley's premise; this is **information entropy, not thermodynamic entropy**
coupled-network fragility [Buldyrev et al., 2010]	**2 — shared form**	the perimeter coupling is not shown to satisfy the percolation premises; the shape of the result is used

borrowing	category	justification
robust-yet-fragile [Albert et al., 2000, Callaway et al., 2000]	**2 — shared form**	the tolerance asymmetry's shape is used; the degree-distribution mechanism is not claimed, and Doyle et al. [2005] is the reason it should not be
identity impersonation as a counting failure [Douceur, 2002]	**2 — shared form**	the Sybil result is proved for an identity system without a certifying authority; a ledger is not shown to satisfy those premises
independent-party bound [Lamport et al., 1982]	**2 — shared form**	cited as a prior formalisation of "how many parties are genuinely independent", not applied to ledgers
redundancy that reduces safety; tight coupling [Sagan, 1993, Perrow, 1984]	**3 — framing**	qualitative organisational theory, cited as framing and as a warning, never as a formal result about this system
a measure adopted as a target ceases to measure [Goodhart, 1975, Campbell, 1979]	**3 — framing**	a documented regularity in policy settings, used as a caution on this paper's own index
performativity of a published model [MacKenzie and Millo, 2003, Merton, 1948]	**1 — derivation**	the mechanism is the same: a description relied upon changes the behaviour of those relying on it, documented with dates and prices
a shared record as a boundary object [Star and Griesemer, 1989, Star and Ruhleder, 1996]	**2 — shared form**	the requirement of simultaneous local plasticity and global robustness is the same requirement; no claim is made that a receipt is an infrastructure in the cited sense

borrowing	category	justification
open design [Kerckhoffs, 1883, Saltzer and Schroeder, 1975]	**1 — derivation**	the disclosure result *is* the cited principle; the premise — effectiveness depends on the partition and not on knowledge of it — is stated and satisfied
the statistical-physics programme [Castellano et al., 2009]	—	cited as the survey that **bounds** what the programme has established; this paper claims no result within it

The failure mode this table exists to prevent is letting a category-2 or category-3 borrowing carry a category-1 result's authority. The concentration arithmetic, the portfolio identity, the proof-length accounting and the disclosure result are category 1 and need no help. The network results are category 2 and are used for shape only.

Entropy, said once and precisely: every entropy in this paper is Shannon entropy in bits, an information quantity over a discrete set, and no thermodynamic quantity appears anywhere.

8.7 Postulates

#	statement	status
P1	shares s_i are non-negative and sum to 1	DEFINITION
P2	HHI $= \sum_i s_i^2$ and $N = 1/\mathrm{HHI}$	DEFINITION
P3	jurisdictions inside one coordination perimeter act on one decision	POSTULATED (H1)
P4	$h = \eta \ln N - 0.500000$ with $\eta = 0.35$	DEFINITION (modelling choice, fixed in advance)
P5	the 1815 share vector 40/20/20/13/7	POSTULATED (CONSTRUCTED)
P6	loss probability over a window is $1 - e^{-\lambda t}$ with $\lambda = 0.0015$/day	POSTULATED (modelling choice)
P7	every double-entry posting carries an amount in one unit of account, and debits equal credits	DEFINITION
P8	recognition requires measurability, so an unmeasurable item cannot be an entry	DEFINITION
P9	the nine-capital share vector and the 35% recognition fraction	POSTULATED (CONSTRUCTED)

#	statement	status
P10	the four stipulated likelihood pairs of Section 4.8	POSTULATED (stipulated)
P11	the corpus design parameters 297/148/55 with five digit-only no-ops	POSTULATED (stipulated to reproduce reported rates)
P12	canon $=$ restrict_ASCII $\circ$ confusable_fold $\circ$ NFKC	DEFINITION

8.8 The chain that establishes the bloc collapse

step	consumes	statement	status
L1	P1, P2	for n equal shares of $1/n$, HHI $= 1/n$ and $N = n$	ARITHMETIC
L2	L1	N is a numbers-equivalent — Adelman [1969]	CITED
L3	L2	N is the Hill number of order 2 — Hill [1973], Jost [2006]	CITED
L4	P5, P2	HHI $= 0.2618$, $N = 3.819710$	ARITHMETIC
L5	P3, P5	collapsing all five onto one perimeter gives HHI $= 1$, $N = 1.000000$	ARITHMETIC
L6	P3, P5	the three-perimeter partition gives HHI $= 0.36$, $N = 2.777778$	ARITHMETIC
L7	L4, L5, P4	h moves from -0.030939 to -0.500000	ARITHMETIC
L8	L6, L7	the spread across three defensible readings is 0.469061 of h	ARITHMETIC
R1	L4, L5, L8	**the naive count is 3.819710 in both eras and the bloc count moves to 1.000000; the instrument's content is the partition, and the partition is a claim**	DERIVED

Only P3 and P5 are POSTULATED in this chain. P3 is H1 with its falsifier attached; P5 is a constructed vector, and L8 is the measure of how much the result depends on a choice the arithmetic does not make. **A reader who rejects P3 reads L4 and stops, which is a legitimate reading and is why all three are printed.**

8.9 The chain that establishes the posterior

step	consumes	statement	status
B1	—	posterior odds are prior odds times the product of likelihood ratios	DEFINITION (Bayes)
B2	P10	prior odds $= 1$	POSTULATED
B3	P10	$0.10/0.60 = 1/6$; posterior $P(U) = 1/7 = 0.142857$	ARITHMETIC
B4	B3, P10	$\times\, 0.15/0.70 = 3/14$; odds $1/28$, $P(U) = 0.034483$	ARITHMETIC
B5	B4, P10	$\times\, 0.05/0.65 = 1/13$; odds $1/364$, $P(U) = 0.002740$	ARITHMETIC
B6	B5, P10	$\times\, 0.20/0.75 = 4/15$; odds $1/1365$, $P(U) = 1/1366 = 0.000732$	ARITHMETIC
R6	B6	$P(\text{competing}) = 0.999268$; odds against the unified model 1,365:1	ARITHMETIC

Every step is arithmetic and the whole chain rests on one postulate, P10. The four likelihood pairs are stipulated and are the only thing a reader need attack. **The chain is therefore not evidence that the conclusion is true; it is a device that makes the inference explicit and recomputable**, and the paper says so at every point of use. What survives a reader substituting their own likelihoods is the qualitative argument of Section 4.8, which does not depend on the numbers at all.

8.10 The chain that establishes the canonicalisation result

step	consumes	statement	status
C1	P12	canon is a function of the string alone	DEFINITION
C2	P11, C1	495 of 500 pairs are byte-distinct, digest-distinct and canonically equal	ARITHMETIC $= 99.0\%$
C3	P11	the residual 5 are fields with no substitutable character, hence byte-identical no-ops	DEFINITION (corpus design)
C4	C1, C2	495 of 495 manipulated entries collapse onto the clean canonical form	ARITHMETIC $= 100\%$
C5	C1	$\text{canon}(\text{canon}(x)) = \text{canon}(x)$ on every corpus value	ARITHMETIC (idempotence checked)

step	consumes	statement	status
C6	C4	distinct raw byte-strings 995 collapse to 500 distinct canonical forms, the true count	ARITHMETIC
C7	C4	under canonical equality the false-break rate is 0.0% at every manipulation rate	ARITHMETIC
C8	P11	under byte equality the false-break rate equals the manipulation rate exactly	ARITHMETIC
C9	P11	the confusable-only check reaches $292/495 = 58.99\%$, being blind to compatibility forms	ARITHMETIC
R4	C4, C7	**canonicalisation collapses 495/495 and takes the reconciliation false-break rate from 40% to 0% at a 40% manipulation rate**	DERIVED
C10	C1	a digest over canon(x) is equal for clean and manipulated values; a digest over the octets is not	ARITHMETIC
R7	C10	identical canonical form with differing byte-hash is the detector, and hashing the canonical form destroys the evidence	DERIVED

Every step is arithmetic on a deterministic corpus, and the corpus itself is the one postulate. P11 is stipulated to reproduce the source's reported rates and is stated as such; **C2, C9 and the combining-mark result are the three figures that move if P11 moves**, and a reader reproducing them needs the stated parameters.

8.11 Discrepancies found in the source, printed

One modelling ambiguity was found and is reported rather than harmonised. The source manuscript gives a 24-day loss probability of 3.5360 per cent and an 8-day probability of 1.1928 per cent at a stated daily hazard of 0.0015, and **does not state which hazard model it uses.** The continuous form $1 - e^{-\lambda t}$ reproduces both figures exactly (3.535978 per cent and 1.192829 per cent). The discrete form $1 - (1 - \lambda)^t$ gives 3.538576 per cent and 1.193721 per cent. **The two differ by 0.0026 percentage points at 24 days.** This paper adopts the continuous form, states

it in Section 4.4, and records the difference here. Neither choice changes the result, which is the 66.6667 per cent reduction in the exposure window.

One number is reported to a precision its inputs do not support and is re-reported here at the precision they do. The completeness figure is quoted throughout the source as 63.850000 per cent from a vector stipulated to three decimal places. $60 + 0.35 \times 11 = 63.85$ exactly; the six decimal places are an artefact of the source's formatting convention rather than a claim about precision, and Section 8.2 states that.

One result is this paper's own and is not a reproduction. The combining-mark measurement — 98 twins, 100 per cent quarantined rather than collapsed, normalisation-mismatch flagging 100 per cent and confusable-only flagging none — is a finding of the harness rebuild described in Section 6.4 and **is not a reproduction of any figure reported in the source.** It is badged T2 as this paper's own measurement.

One naming claim in the source is contested within the source itself and is therefore not made here. The source's verification register flags its own attribution of a specific withdrawn numeric currency code as requiring verification, noting a competing attribution. **This paper states the mechanism and omits both attributions**, since the demonstrated mechanism does not depend on which withdrawn currency the code denoted.

One date correction is carried forward. The article establishing that a system's interfaces mirror its builders' communication structure is dated 1967 elsewhere in the corpus; it appeared in April 1968, and Conway [1968] is cited here with the corrected date.

9. Limitations

Every worked result in this paper is a constructed or synthetic model, and none is a measurement. The capital shares, the transit parameters, the hazard rate, the balance sheet, the invoice, the entity count, the nine-capital vector, the likelihood pairs and the 500-pair identifier corpus are all stipulated. **No external dataset was collected, and no real institution, ledger, counterparty file, transaction or price series appears anywhere in this paper.**

The bloc call is the single largest unmeasured input and it is a choice, not a finding. Section 4.2 gives three defensible readings of one structure spanning 0.469061 of h, and nothing in the arithmetic selects among them. **A reader who prefers reading 2 should read $h = -0.142422$ off the table and proceed.** The perimeter test of Section 3.2 is stated so that the disagreement is about an observable rather than about a geopolitical opinion, but the test was applied here by reading a historical description, not by examining a transmission mechanism.

No correlation was estimated from data. The bloc reading is the $\rho = 1$ corner of the family in Curzi [2026] and the naive reading is the $\rho = 0$ corner. Both corners are stipulations. **The true structure lies between them and this paper does not say where.**

The surplus parameters carry no operational unit and no empirical warrant. $\eta = 0.35$ and the maintenance floor of 0.500000 are fixed in advance by the corpus. **Section 8.1 records that h is a dimensionless index whose zero is a stipulation**, so every difference in h in this paper is a difference in a stipulated index.

The historical material is described from one source manuscript and is not independently verified. Names, dates, deeds, amounts, tranches, capital allocations, courier composition and the provenance of the refused account are all listed in Appendix A. **The refusal of Section 4.5 is an evidential judgement rather than a measurement**, and it depends on the manuscript's dating of a polemical pamphlet, which this paper has not confirmed.

Jurisdiction-specific law, statutes, registry documents and case law are not in the bibliography and none is cited. The doctrinal chain of Section 4.7, the instruments behind it, the four judgments named in the source and the 2023 statement are described from the source manuscript, itemised in Appendix B, and **asserted with no citation, no reporter reference and no date beyond the years used in the interval arithmetic.** The intervals survive; every particular requires verification. `bib/COVERAGE.md` records the absence of jurisdiction-specific primary sources and of case law as known gaps, and this paper does not close either.

Technical standards are not in the bibliography and none is cited. Every designation in Section 6 — the coded character set, the normalisation forms, the security mechanisms and their skeleton and restriction levels, the bidirectional algorithm, the transformation formats, the domain-name framework, the hash standard and the currency-code standard — is described from the source manuscript and itemised in Appendix B. **This is a real gap and it is not closed here.**

The nine well-documented failures of Section 7.2 are described at the level of mechanism only. None is cited, because the bibliography contains no case reports, no regulatory findings and no contemporaneous reporting. The one in-class instance is described without entity names and requires precise sourcing before reuse. **Nothing in Section 7.2 is an allegation about any named person, institution or product.**

The measurement results reproduce a source the paper did not obtain. The harness was rebuilt from a description, and the corpus design parameters were reconstructed until two of five results agreed. **A different class mix gives different figures for E1 and for the confusable-only row of E2**, and a reader reproducing them must use the stated parameters.

No prevalence claim is made or supported. The rate at which confusable substitution occurs in production financial systems is unmeasured, here and — so far as this paper is aware — anywhere. **The demonstration shows that the mechanism works; it is not an estimate of exposure in any real ledger** [T4]. Establishing prevalence requires data access this paper does not have and is the obvious next piece of work.

The zero false-positive column does not transfer. It is measured against clean ASCII critical fields. A non-ASCII check against a real name estate produces false positives at a very high rate, because a great many people legitimately have non-ASCII names. **Section 7.2 states this as the most serious objection to the approach.**

No macroeconomic magnitude is claimed. That representational failure at the record level could contribute to measurement discrepancy in aggregates built from those records is a bounded and testable hypothesis. **No magnitude is estimated, no share of any discrepancy is claimed, and no observed aggregate is attributed to this mechanism** [T4].

The off-ledger register has no standing anywhere. It is this paper's own construction, it is not a disclosure, it is recognised by no framework, and it must not be described to any regulator or auditor as one.

9.1 Rigour checks not met, listed by name

1. **No as-of date and no source, because nothing was measured.** Every figure derives from a constructed parameter set. Any application to a real structure, ledger or master file requires both.
2. **No uncertainty is reported anywhere**, because nothing is estimated. This is the absence of empirical content rather than precision.
3. **Model comparison was not performed.** No alternative model was fitted, no out-of-sample test was run, no information criterion was computed, and no matched null with the same number of free parameters was constructed. The models here are illustrations of identities and inequalities, not fits, and their conclusions follow from the identities rather than from goodness of fit. Epstein [2008] is the standing statement that prediction is one of sixteen reasons to model and that a paper is entitled to say which of the other fifteen it claims; this paper claims demonstration that a stated mechanism suffices to produce a stated shape.
4. **The models fail the calibration-and-validation checklist** of Bruch and Atwell [2015]: none is calibrated against observed data, none is validated out-of-sample, and none has a sensitivity analysis beyond the limits of Section 8.3.
5. **The conjoint-measurement conditions were not tested** [Luce and Tukey, 1964], so the ratio-scale treatment of the nine-capital vector is provisional and Michell [1997]'s objection is conceded rather than answered. **This is the paper's weakest input and Section 8.5 says so.**

6. **The corpus-wide scalar-reduction limitation is inherited** [Curzi, 2026] and not met separately: a structured situation with several jointly varying features is replaced by one number in two places in this paper, and arithmetic is then done on it.
7. **Small-sample bias in the effective count is not corrected** [Chao et al., 2014]. A five- or six-unit exposure vector is a small sample for a low-order diversity estimator.
8. **No statistical test was run and no p-value appears**, because there is no sample and no null. Every rate reported is an exact ratio of counts in a constructed corpus.
9. **No scaling, power-law or heavy-tail claim is made, and the Clauset–Shalizi–Newman procedure was therefore not run** [Clauset et al., 2009]. The one exponent that appears is exact combinatorics. Any restatement reaching for such a claim must run the procedure first and read Broido and Clauset [2019].
10. **This paper does not invoke the framework's reflexive law and therefore does not inherit the reflection problem** [Manski, 1993]. No result here rests on an endogenous social effect estimated from observation, and no such effect is claimed. **Where downstream work couples these results to that law, the caveat applies there and is not answered here**: a linear-in-means reflexive specification is not identified from observational data, and the remedies — randomisation, an instrument, partial-population variation, or stipulating and defending the reference group [Brock and Durlauf, 2001] — all carry a price. The one reflexive claim this paper does make, that publishing a control changes the behaviour of those subject to it, rests on a documented case [MacKenzie and Millo, 2003] and is not estimated here.
11. **The Lucas critique applies here and is conceded rather than answered** [Lucas, 1976]. The framework holds that systems change behaviour when modelled and observed; $\eta = 0.35$ and the maintenance floor are then not obviously stable under the framework's own predicted response to their publication. **This paper does not defend their stability; it fixes them in advance and reports every result as conditional on them.**
12. **Two borrowings are category 3 and are used only as framing**: Perrow [1984] and Sagan [1993]. Neither is a result about this system and neither is treated as one.
13. **The evidence-tier scheme is this paper's own and has no external standing.** It is an ordinal labelling convention, not a validated instrument.

10. Conclusion

This paper establishes six results in the form the rest of the corpus cites, and the following statements are the fixed text of those results.

R1 — the bloc collapse on a historical network. The five-house network measures 3.819710 effective jurisdictions naively — $1/\text{HHI}$, the Hill number of order two, on constructed shares of 40/20/20/13/7 — but **1.000000 effective blocs once a single coordination perimeter is recognised, giving** $h = -0.500000$**.** The same five cities under a nominal modern split give 1.923077 and $h = -0.271126$. **The naive count is identical in both eras; the bloc count is not, and the difference across three defensible readings of the 1815 structure is 0.469061 of** h**, which is a definitional choice and not a fact.**

R2 — double entry is a consistency check and not an integrity check. $A - (L + E) = 0.0000000000$ holds exactly while the statement misleads — a wholly fictitious 1,000 moves total assets by 66.6667 per cent and days-sales-outstanding by 82.125000 days — **so the identity's holding is not evidence of anything except arithmetic.**

R3 — the completeness gap, measured. On a constructed nine-capital vector, **63.850000 per cent represented** (financial 42 per cent and manufactured 18 per cent fully, intellectual at a stipulated 35 per cent recognition) and **36.150000 per cent structurally unrepresentable.** The gap is the price of the closure property and not a defect a better standard-setter would repair.

R4 — canonicalisation as an integrity control. The pipeline

$$\text{canon}(x) = \text{restrict_ASCII}(\text{confusable_fold}(\text{NFKC}(x)))$$

collapses 495 of 495 manipulated entries onto the clean canonical form, idempotently, and moves the reconciliation false-break rate from 40 per cent to 0 per cent at a 40 per cent manipulation rate — and to 0 per cent at every rate from 0 to 100 per cent, against byte equality's one-for-one tracking.

R5 — two correct books disagree, and the reconciliation bill is quadratic. A EUR 100,000 invoice booked at 1.0900 and 1.0850 leaves a USD 500 gap, **0.460829 per cent of the buyer's figure, with neither party in error**; and nine entities need $n(n-1)/2 = 36$ bilateral reconciliations against one shared record.

R6 — the coordination question resolved as a posterior rather than as an assertion. $P(\text{unified hand}) = 0.000732$ against $P(\text{competing centres}) = 0.999268$, **odds of 1,365:1 against the unified model**, on stipulated likelihoods and an exactly rational posterior.

Counter-result. This paper refuses its own most attractive material. **The one-day-coup account is refused as a fabricated data point** and the documented multi-year position substituted, at a constructed total return of 41.935484 per cent over two years that no single session could produce. **The unified-hand reading is rejected at 1,365:1** and the converse error refused in the same paragraph. **The 25 per cent currency-inflation figure is badged as an artefact of the chosen synthetic 8000/2000 ratio rather than as a constant**, with the confusable-only check reaching only 58.99 per cent against a headline of 100 per cent. And **the concentration instrument the corpus uses to measure resilience is shown to be vulnerable, at the character layer, to a defect that biases it in the reassuring direction.**

What the paper adds to the corpus's instrument. An effective count is a function over a partition, and neither the partition nor the identity resolution beneath it is supplied by the instrument. **Section 4.2 measures the cost of getting the partition wrong at 0.469061 of h; Section 5.6 shows the ledger cannot supply it; Section 6 shows that the identity resolution beneath it is string comparison with a measurable and cheaply closed blind spot.** A concentration measurement is therefore only as good as two things it does not contain, and both must be published with it or they will not be examined at all.

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every worked result in this paper — the capital shares, the transit model, the constructed government-stock and sovereign-loan illustrations, the balance sheet and the fictitious posting, the invoice, the entity count, the nine-capital vector, the likelihood pairs, the 500-pair identifier corpus, the reconciliation sweep, the withdrawn-code extract and the ten-name watchlist — is **CONSTRUCTED** or **SYNTHETIC** and is labelled as such at each point of use. The numerical results are exact computations on those stipulated inputs. **No real institution, ledger, counterparty file, account, transaction, price series or person appears anywhere in this paper, and no allegation is made about any of them.**

All figures were produced by the scripts named in their captions (`figs/src/CCP-064_fig01.py` through `figs/src/CCP-064_fig05.py`), which import the shared figure style module and read their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

The historical, legal and standards material is described from the source manuscript and is not cited to the bibliography, which contains no case law, no jurisdiction-specific primary sources, no standards documents, no regulatory findings and no press reporting. Appendix A lists the historical and measurement claims requiring verification; Appendix B lists the legal instruments, judgments, standards designations and documented matters described but not

cited. **Any reader relying on any of them must verify them against the governing document and against contemporaneous sources.**

The demonstration harness underlying Section 6 was **rebuilt from a description rather than obtained**, and its corpus design parameters were reconstructed until two of the five reported results agreed. Those parameters are stated in Section 6.4 and are stipulations.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data were processed. The historical material concerns institutions and published doctrine; no living individual is discussed, and no inference is drawn about any identifiable person.

No operational content

This paper contains no operational content. The substitution class of Section 6 is described at the level of mechanism, from material that is already public in the governing standards and in the published security literature, and it is described in order to specify a control. **No code, no data file, no substitution table and no method of evasion is supplied**, and the paper's purpose throughout is the construction of a detector rather than the defeat of one.

Legal, tax, accounting and investment advice

This paper is analytical and is not legal, tax, accounting or investment advice. It discusses historical institutions, a bookkeeping method, a doctrinal chain and a data-integrity mechanism. No recommendation is made in respect of any particular person, institution, product or jurisdiction, and no allegation is made about any of them. **Every legal and doctrinal statement in Sections 4.6 and 4.7 is described from a secondary source, is uncited, and must be confirmed with qualified counsel in the relevant jurisdiction before any reliance.** Accounting treatments must be confirmed with the reporting entity's auditor; recognition requirements differ between frameworks and change. Data-retention and disclosure obligations differ by jurisdiction; any control designed on the basis of Section 6 requires confirmation with qualified counsel and with the implementing institution's own compliance function.

Appendix A. Historical and measurement claims requiring verification

The following are described in this paper from the source manuscript and are **not cited**. Each requires verification before any reliance.

1. The founder's dates; the names, birth order and postings of the five principals; and the year each house was established.
2. The partnership instruments — a Frankfurt deed conventionally dated 1810 and later instruments — with their parties, terms, profit-sharing percentages, renewal intervals, and **the extent to which liability was in fact separated between houses.** This is the load-bearing legal claim behind the architecture argument of Section 4.1 and should not be relied on until checked.
3. The actual capital allocation across the five houses at any date. **The 40/20/20/13/7 vector is CONSTRUCTED and is not an archival figure.**
4. The courier and vessel network: composition, cost and speed advantage over the public posts. The three-hop/one-hop structure, the eight-day leg and the 0.0015 daily hazard are stipulated.
5. The correspondence practice: confirm the language and script pair, the period over which it was used, and whether any further cipher or code-word layer sat on top of it.
6. The scholarly work the source manuscript relies on for the corrected account [Ferguson, 1998]: confirm the exact volume covering the relevant period and the pages treating the bullion position and the subsequent government-stock position. **Every claim in Section 4.5 should be checked against it directly.**
7. The 1846 pamphlet: its author, title, date, pseudonym and its actual causal role in the circulation of the account. **This paper asserts none of these.**
8. What the London principal actually did on receiving the news: whether the government was informed, on what day, the size and disposition of the bullion position, and the actual entry and exit prices and dates of the subsequent position. **The 62/82/3 figures are CONSTRUCTED.**
9. The sovereign loans of 1831–32: the lending houses, amounts, currencies, issue prices, coupons, security and tranching, and the chronology of the revolt and the intervention. **The 1,000/75/5 per cent illustration is CONSTRUCTED.**
10. The codification of double entry: the 1494 printed work, its printer, the exact title as printed, and the internal location of the bookkeeping tract; the earliest surviving complete double-entry records and their date, which this paper's "roughly a century and a half" depends on; and the composition and printing dates of the earlier manuscript treatment.
11. The recognition criteria relied on in Section 5.5: the current recognition requirements of the relevant conceptual framework and the current prohibition on recognising internally generated goodwill. **This paper states the prohibition and should not be relied on for a paragraph number**, nor for whether the equivalent position under another framework materially differs.
12. The 1982 and 1989 momentum-accounting works: exact titles, series numbers, publisher, years, the relationship between them, and the author's own terminology for the third entry class. **This paper deliberately does not name that coinage.** Whether momentum accounting has been adopted anywhere in practice or in any standard: this paper asserts it has not.
13. The hash-tree publication: the correct paper, proceedings volume and year, and whether a later paper is the more appropriate citation for the construction as used today.
14. The 2005 signed-receipt work: exact title, venue and year; whether each party retains conventional books against the shared receipt; whether a distinct third-party issuer is required in all formulations; and whether related earlier work belongs in the same lineage.
15. The claim that double entry is required by every major accounting framework — verify the form in which the requirement appears, since it is generally implicit rather than stated as a rule.
16. All arithmetic in Sections 4, 5 and 6 is reproducible from the stated inputs, and every input is stipulated: the opening balance sheet (400/300/800/250/600/650); the fictitious 1,000; the stipulated annual revenue of 3,000; the 1.0900/1.0850 rate pair; the 32-byte

hash and 256-byte entry assumptions; both revenue series; the nine-capital vector and the 35 per cent recognition fraction; the nine-entity count and the 40,000,000 intercompany balance. **None is drawn from any real entity or any published dataset.**

17. The nine-capital frame is developed in Curzi [2026] and **the vector used here is illustrative and must be reconciled with that paper's definitive list, ordering and terminology.**
18. The corpus class mix (297/148/55) and the five digit-only no-op identifiers are **stipulated design parameters** chosen to reproduce reported rates. They are not derived from the source's code, which was not obtained. A reader reproducing E1 and the confusable-only row of E2 must use the stated parameters.
19. **The combining-mark result is this paper's own** and is not a reproduction of any reported figure. Verify independently.
20. **The 25 per cent of E5 is a property of the synthetic 8000/2000 ratio and is not a general magnitude.** Do not quote it as one.
21. **No prevalence claim is made or supported.** The rate at which this occurs in production systems is unmeasured, and establishing it requires access to real ledger data that no part of this paper has or uses.

Appendix B. Legal instruments, judgments, standards and matters described but not cited

`bib/external.bib` contains no case law, no jurisdiction-specific primary sources, no standards documents, no regulatory findings and no press reporting, and `bib/COVERAGE.md` records each of those as a known gap. The following are therefore **described from the source manuscript and cited to nothing.** Each requires verification against the governing document before any reliance, and none may be restated as a citation on the strength of this paper.

Instruments and doctrine (Section 4.7).

1. The four fifteenth-century instruments of 1452, 1455, 1493 and 1494 — issuing authority, exact dates, and the precise language relied on; and, for the 1493 instrument, which of more than one bull of that year is meant.
2. The 1823 judgment of a common-law supreme court receiving the doctrine into positive law: reporter citation, the pages containing the discovery and right-of-occupancy holdings, and whether the court was unanimous.
3. The two elaborating judgments of 1831 and 1832: citations and the respects in which each elaborates.
4. The 2005 judgment applying the framework: citation, holding, and the specific place in the opinion where the framework is invoked. **This paper deliberately names no footnote.**
5. The 1992 judgment of a different common-law supreme court rejecting one component, the 1993 statute following it, and the 1997 and 2014 judgments named in the source: citations and current status.
6. The 2023 statement by the issuing institution: date, the issuing bodies, the exact title, and the quoted characterisations.

Institutional chronology and the posterior's observations (Section 4.8).

7. The 1929 treaty and its financial provisions; the engagement of the manager appointed that year; the 1942 constitution of the financial body and the instrument constituting it; the 1982 collapse chronology and the death that year; the nature and legal effect of the letters relied on; and the 1984 payment — amount, date, parties, and **the exact settlement language recording moral involvement without admission of liability.** The source states approximately USD 240 million; this paper states "a nine-figure sum" and asserts no amount.

8. The 2013 arrest, its charges, forum and outcome; and the December 2023 conviction, its charges, sentence, appellate status at the date of reading, and the claim that it was the first such trial by that court.

Standards designations (Section 6).

9. The coded character set and its joint international designation; the normalisation forms and their governing annex; the security mechanisms document defining the skeleton transformation and the restriction levels; the bidirectional algorithm and its formatting controls; the two transformation formats named in the charset example; the domain-name encoding and framework documents; the hash standard; and the currency-code standard. Confirm every designation, its current revision and any obsoleting document. **The confusables data file is revised with each release of the character standard and must be treated as a configuration item. The source's attribution of a specific withdrawn numeric currency code is contested within the source's own register; this paper omits both candidate attributions and asserts none.**
10. The 2002 and 2006 papers on the homograph attack, and the 2021 disclosure of the bidirectional source-display issue together with its vulnerability identifier and the associated identifier for the homoglyph aspect. Confirm titles, venues, years and identifiers.
11. Structured party identification in payment messaging is cited where relied on [International Organization for Standardization, 2013]; the legal-entity identifier standard named in the source is not, and the state of any migration mandate in any market must be confirmed before relying on the structural argument of Section 7.2.

Documented matters used only as a negative audit (Section 7.2).

12. The nine failures are described at the level of mechanism, without dates, quanta or outcomes. Confirm, for each, the authoritative published record before any restatement — including which criminal convictions in one of them were subsequently quashed on appeal, by which courts and on what dates. **No part of this paper's argument depends on any of those outcomes.**
13. The single in-class instance: **confirm the exact registered names of both entities, the year and place of incorporation of the one, the corporate relationship of the other to its parent, and the specific civil forfeiture complaints in which the arrangement is described — case names, courts and filing dates.** This is the factual claim in Section 7.2 carrying the most weight and it must be sourced precisely or removed. **This paper reproduces neither name.**
14. The published compliance framework identifying screening-software and filter deficiencies among recurring root causes: exact title, date and the precise wording of the findings. The authoritative names and publishers of the designated-persons lists named generically in Section 7.2. Any industry-body or intergovernmental screening guidance relied on, by exact title and date; **none is cited here.**
15. The claim that mature screening engines commonly use fuzzy, phonetic and transliteration matching is a general characterisation of a vendor landscape, is sourced to no named product, and **whether any specific engine folds confusables before computing edit distance must be established by direct testing of that engine rather than assumed in either direction.**

Jurisdiction-specific regulatory facts (Section 7.1).

16. The source manuscript names, for each node of a six-jurisdiction structure, a specific listing cycle with dates, a specific administrative order, specific substance regulations and specific unilateral measures. **None of those is reproduced here and none is cited**, following the de-branding convention of Curzi [2026]; the table in Section 7.1 gives mechanism classes without jurisdictions. Each named instance requires its own statute, regulation or register citation, pulled from the official register and verified individually.

All numerical results in this paper were computed from the stated constructed inputs at full precision and rounded only for printing.

Chapter 65

Canonical Form and Typed Parsing: Measured Integrity Controls Across a Financial Identifier Namespace

Abstract. We specify and measure an integrity layer for machine-readable financial fields. The canonical form canon(x) = restrict_ASCII(confusable_fold(NFKC(x))), paired with a byte-level SHA-256 witness kept on a separate plane from the display form, is tested against a synthetic corpus: 495 of 495 manipulated entries collapse onto their clean canonical form, idempotence canon(canon(x)) = canon(x) holds for all 500 cases, and 995 distinct raw byte strings resolve to the true 500 canonical identifiers. The reconciliation false-break rate falls from 40.0 per cent to 0.0 per cent at a 40 per cent manipulation rate. Within the ISO namespace, aliasing a withdrawn currency code to a live one is shown to inflate a synthetic 8,000 balance to a recorded 10,000, a phantom 25 per cent with no journal entry. Naive detection of withdrawn codes has a precision of 11.01 per cent, meaning 88.99 per cent of its queue is parsing ambiguity, and one typed-field parse collapses that share to 0.00 per cent. The migration to structured messaging inverts the intuition about richness: in a synthetic thousand-message day with 62 per cent rich traffic, sparse messages are 38 per cent of volume and 93.9 per cent of the review queue. The stated rates are stipulated design parameters, and the architecture is badged as providing no consensus, no double-spend prevention, no valuation method, no adopted standard status and no legal force.

Keywords: canonicalisation; confusable folding; ISO 4217; legal entity identifier; ISO 20022; typed parsing; false-break rate; screening queue

Series: D. Structure and Resilience **Paper:** CCP-065

1. Introduction

A financial control that decides whether two records name the same thing is, at bottom, a string comparison. Curzi [2026] measures what that comparison misses when the character set admits more than one code point per shape: 495 of 500 manipulated identifier pairs are byte-distinct, digest-distinct and canonically identical, and the reconciliation false-break rate tracks the manipulation rate one-for-one under byte equality. **This paper states the remedy as a specification, generalises it one level up to the codes those strings carry, and measures what the generalisation buys at the layer where payments actually cross a border.**

The paper's contributions are separable and each carries its number.

An integrity layer, specified and measured. The canonical form $\mathrm{canon}(x) = \mathrm{restrict_ASCII}(\mathrm{confusable_fold}(\mathrm{NFKC}(x)))$, paired with a byte-level SHA-256 witness computed over the raw octets and kept on a plane separate from the display form, is tested against a synthetic corpus. **495 of 495 manipulated entries collapse onto their clean canonical form; idempotence holds for all 500 cases; 995 distinct raw byte strings resolve to the true 500 identifiers; and the reconciliation false-break rate falls from 40.0 per cent to 0.0 per cent at a 40 per cent manipulation rate.** Eighteen normative rules and their conformance tests are given, so that an implementation can be checked rather than trusted.

The same failure one level up, in the code namespace. A code is trustworthy only if there is one live referent per code, and what guarantees that is a registry rather than the string. Aliasing a withdrawn currency code to a live one records a synthetic 8,000 balance as 10,000 — **a phantom 25.0 per cent inflation with no journal entry and nothing forged.** Worse, the same three digits occur in two independently governed registries with different meanings, so a screen that scans bytes rather than typed fields cannot tell them apart: **naive detection has a precision of 11.01 per cent, meaning 88.99 per cent of its queue is parsing ambiguity, and one typed-field parse collapses that share to 0.00 per cent.** A control right one time in nine is not a control.

The resolvability asymmetry on a settlement network. The migration from field-limited legacy messages to a structured, richly typed standard inverts the intuition that carrying less data reveals less. A screening system disposes of a message it can resolve and escalates one it cannot, so **the message that carries less is not more private, it is less resolvable, and unresolvability is what costs.** In a synthetic day of 1,000 messages at 62 per cent rich traffic, **sparse messages are 38 per cent of volume and 93.9 per cent of the review queue.** The finding is not that disclosure is virtuous; it is that a lawful structure has a direct operational interest in being resolvable, and a structure that cannot afford to resolve cannot adopt the posture.

An architecture reported with its negative results. Above the specification sits a ledger substrate factored into per-object hash chains and a content-addressed capability layer. Both are the author's own systems, both were audited adversarially by the author, and **the audit findings are reported here in the same plain language as the successes, including the ones that refute the author's earlier claims.** The substrate verifies 200 of 200 per-token chains and detects 200 of 200 silent edits; it records 28 of 28 grid links in both ledgers; and its local append cost rises from 0.002 s to 0.014 s as the chain count goes from 50 to 400 — **a local, single-process measurement that is not distributed throughput and is not offered as one.** It supplies **no consensus, no total order, no double-spend prevention and no Sybil resistance**, and a working double-spend against it is reported. Its block hash under-binds, and the fix is given. Four findings in the capability layer — an unverified proof signature, an inconsistently computed self-verification digest, broken hash algorithms on offer, and a replay surface at the freshness root — are reported with their fixes and re-validation. **The naming rationale and a numerological layer present in the implementation carry no technical content, are excluded from every claim in this paper, and are graded as claims deliberately not made.**

Why these sit in one paper. They are one mechanism at three heights. A character is trustworthy as an identifier only if one code point maps to one character; a code is trustworthy only if one code maps to one live referent; a message field is trustworthy only if the receiving control parses it under the grammar that owns it. **In each case a representation is trusted by a consumer that cannot see the binding, and the binding is maintained somewhere else.** That is the same defect Curzi [2026] measures at the signalling layer, where an in-band ambiguity rate of 0.383 per cent falls to exactly 0.000 per cent once control and data are disjoint by construction, and it is why the specification here is stated as a boundary rule rather than a field rule.

2. Hypotheses

Six hypotheses, each stated so that it could be wrong, each with its falsifier [Platt, 1964, Popper, 1959]. Statements that are not falsifiable are listed at the end and labelled.

H1 — the canonical-form hypothesis. A fixed three-stage canonical form applied at the comparison boundary maps every member of the substitution class onto one representative, and does so idempotently. *Predicted:* 495 of 495 manipulated entries collapse onto the clean canonical form; $\text{canon}(\text{canon}(x)) = \text{canon}(x)$ for all 500 corpus values; 995 distinct raw byte strings resolve to 500 distinct canonical forms, the true count. *Falsified by:* a manipulated value whose canonical form differs from its clean counterpart's, or any value for which the operation is not idempotent.

H2 — the witness-level hypothesis. The evidence that a substitution occurred survives only if the digest is taken over the raw octets; taking it over the canonical form destroys that evidence. *Predicted:* for a clean value and its homoglyph twin, canonical forms equal, byte-level digests differ, character-level digests identical. *Falsified by:* a character-level digest that distinguishes the pair, or a byte-level digest that does not.

H3 — the typed-parsing hypothesis. Where two independently governed registries mint the same string with different meanings, a byte-level screen cannot separate them and a typed parse can, without any determination of intent. *Predicted:* on the synthetic corpus, a naive substring screen flags 1,363 records of which 150 are genuine, a precision of 11.01 per cent and an ambiguity share of 88.99 per cent; a typed parse isolates the 150 with 0 false positives, an ambiguity share of 0.00 per cent. *Falsified by:* a byte-level test that separates the two registries' occurrences on the string alone, or a typed parse that loses a genuine event.

H4 — the resolvability hypothesis. A screening control disposes of a message it can resolve to one identity and escalates one it cannot, so review load concentrates on the message format carrying the least identifying data rather than on the payments of most concern. *Predicted:* in a synthetic day of 1,000 messages at 62 per cent rich, sparse messages are 38.0 per cent of traffic and 93.9 per cent of the review queue; a richly typed message carrying a deterministic entity identifier clears on one exact-key comparison while the same payment in sparse form fails to a two-candidate fuzzy set and goes to a human. *Falsified by:* a screening design that disposes of an unresolvable message without escalation, or a queue whose composition does not track message richness.

H5 — the structure-is-not-canonical-form hypothesis. A structured, richly typed message improves party identification without closing any member of the substitution class, because a syntax fixes the shape of the data and not the identity of the values. *Predicted:* a homoglyph twin placed in the standard's own name element is byte-distinct and canonically identical, passes every structural validation, and is separated only by the canonical form and the byte-level witness. *Falsified by:* a structural validation in the standard that rejects the twin, or a canonical form defined by the standard itself.

H6 — the substrate hypothesis. Factoring a ledger into per-object hash chains delivers tamper-evidence and local write-parallelism and does **not** deliver consensus, total order, double-spend prevention or Sybil resistance. *Predicted:* 200 of 200 chains verify; 200 of 200 single-block mutations are detected; 28 of 28 cross-object links are recorded in both ledgers; local append cost is linear in the chain count at 0.002, 0.003, 0.007 and 0.014 s for 50, 100, 200 and 400 chains; **and a fork transferring the same balance twice produces two internally valid chains that the structure does not reject.** *Falsified by:* a silent edit that verifies, or an ordering property emerging from the factoring without an added ordering layer.

Statements in this paper that are not hypotheses, and are labelled as such.

- *The eighteen normative rules* are a **proposal**. They have no adopted standard status, replace no obligation, and carry no legal force. What is testable is each rule's conformance test.
- *The resolution designator* is an **accounting proposal** requiring a standards process and auditor acceptance, and having neither.

- *The nine-form backing taxonomy* is a **classification**, not a measurement. **No measurement basis for the non-financial forms is specified in the source material or here**, and any coverage ratio computed across them is computed over an unknown.
- *The system's naming rationale and its numerological layer* are **claims deliberately not made.** They confer no property on any structure described and are excluded from every technical claim in this paper.

3. Method

3.1 What is specified, what is measured, and what is only reported

Three different kinds of statement appear in this paper and the method keeps them apart.

Specification. Sections 4.1 to 4.6 state normative rules with keyword conventions and, for each, the test that decides conformance. **A specification that cannot be tested is a wish**, so every rule that matters carries its test in Section 4.6. Nothing in the specification has adopted status.

Measurement. Sections 4.7, 5.8, 6.4 to 6.6 and 7.1 report deterministic computations over synthetic corpora with fixed seeds. Each corpus is stipulated and each stipulation is printed, so that a reader can reproduce the figure or show that a different stipulation gives a different one. **No figure in this paper is an estimate of prevalence in any real system**, and Section 10 says so by name.

Report. Section 7 describes artefacts the author built and audited. The measurements there were recomputed for this paper on a harness rebuilt from the source descriptions; **the original code was not obtained and was not rerun.** The audit findings are reported as findings about the implementation *as audited*, and Appendix A records that the current state of each must be confirmed before any finding is repeated in the present tense.

3.2 Evidence tiers

tier	meaning
T1	established fact — a published standard, a proven cryptographic or distributed-systems result
T2	measured by a deterministic harness and independently recomputed for this paper
T3	proposed standard or sound design, not adopted, not hardened at scale
T4	refuted, unsupported as stated, or a claim deliberately not made

The headline grades: the canonical form and the byte-level witness are T3 as a normative proposal and T2 where measured; hash chaining, content addressing and the impossibility results are T1; the resolution designator is T3 and is not an adopted rule; **the naming rationale, the numerological layer, and the claim that a per-object ledger needs no consensus are all T4.**

The tier scheme is this paper's own ordinal labelling convention and has no external standing. It is not averaged anywhere and no arithmetic is performed on it.

3.3 The disclosure test

Every element of the specification is audited in print against one question: **would it still do its job if every element were disclosed to every authority with jurisdiction over any part of it?**

This is not an ethical preference. It is Kerckhoffs [1883]'s principle, stated in 1883 — a system must remain secure when everything about it except the key is public — restated as *open design* among the protection principles of Saltzer and Schroeder [1975] and given its transparency-as-remedy form by Brandeis [1914]. Curzi [2026] states the resulting constraint in the form used here: **effectiveness that is entirely independent of whether anyone knows the control is in place makes every element publishable, so non-disclosure is refused as a design premise rather than merely discouraged.** The regulatory environment has independently removed non-disclosure from the available design parameters [Organisation for Economic Co-operation and Development, 2017, Financial Action Task Force, 2012, Organisation for Economic Co-operation and Development, 2015, de Goede, 2011].

The specification passes the test emphatically, and the reason is worth stating because it is the paper's strongest structural claim about itself. **Publishing the control tells an adversary that substitution against this system will be detected, which is the deterrent, and tells a supervisor or a counterparty that the system does not have the hole, which is the assurance.** Section 8.2 records the one place where the test is passed but the *ordering* matters: a treatment that changes how a balance is presented must obtain auditor acceptance before disclosure, not after.

3.4 Provenance of the artefacts, and the exclusions

The ledger substrate and the capability layer described in Section 7 are the author's own systems, and the adversarial audit reported there is also the author's. That is a conflict of interest, it is declared in full in the Declarations, and the method adopted in response is the only one available to a single author: **report the negative results at least as prominently as the positive ones, and report the ones that refute the author's own earlier claims by name.** Section 7 does that. A reader who discounts the positive findings accordingly is reading correctly; the negative findings are the ones that cost something to publish and are therefore the ones worth weighing.

Two elements are excluded from every technical claim in this paper. The first is the system's **naming rationale**, which the author's own materials record in three readings — a sacramental symbolism, an operational expansion of the same four letters, and an agricultural reading. All three are recorded here as declared and none is corroborated, adopted or verified. **The naming carries no technical content and confers no property on any structure described.** T4.

The second is a **numerological layer** present in the reference implementation: a modulo-nine coherence score, an associated identity claimed as a theorem, a nominated binding frequency, and a set of nine dimension labels. **The author's own security audit assesses them and excludes them by name from every security claim; this paper excludes them likewise and mentions them only so that a reader encountering them in the source code knows they were assessed and discarded rather than overlooked.** T4.

The exclusion is the reason the material can be printed at all. A system whose author will not separate a symbolism from a specification cannot be evaluated; a system whose author separates them in print, before the specification begins, can. Curzi [2026] is the corpus's general form of that discipline — four items of established mathematics, seven of established psychology and exactly three declared postulates, with every protocol stated to run on the first two and steer by the third — and this section is its application to an engineering artefact.

4. The integrity layer, specified

4.1 The canonical form

For critical machine fields — identifiers, account and transaction numbers, amounts, currency codes and action codes — the canonical form is the composition of three operations

in a fixed order [T3]:

$$\text{canon}(x) = \text{restrict_ASCII}(\text{confusable_fold}(\text{NFKC}(x)))$$

Each stage has one job and the order is not interchangeable.

Compatibility normalisation resolves the compatibility and full-width members of the substitution class: full-width digits fold to their ASCII counterparts, and compatibility variants of Latin letters fold to their base forms. **Compatibility normalisation is specified rather than canonical normalisation, and the difference is not cosmetic**: the canonical composed form does not fold full-width variants at all, so a system that normalises to it has done nothing about that member of the class while believing it has.

Confusable folding maps homoglyphs onto a single representative and removes default-ignorable code points. It is the skeleton operation in the role the governing security standard defines for it: two strings are confusable when their skeletons are equal. **A conforming implementation consumes the published confusables data file rather than a hand-written table, and treats the character standard's version as a change-controlled configuration item**, because a confusables table that drifts out of date is a control silently losing coverage.

Restriction enforces the permitted character set for the field. The operative word is *enforces*. **A residual out-of-set character is rejected or quarantined with an explicit reason, never silently stored and never silently stripped** — because silent stripping is itself a substitution, performed by the control rather than against it, and it destroys the evidence that anything happened.

The order matters in one direction specifically. **Folding before normalising would let a compatibility form of a confusable character survive both stages**, because the confusables data is keyed on the normalised representative. Normalise first, fold second, restrict third.

Reconciliation, deduplication, matching and screening then operate on $\text{canon}(x)$ rather than on x.

4.2 The byte-level witness, and why the level is the whole design

Each record additionally stores a SHA-256 digest computed **over its raw octets** — the bytes as received, before any normalisation, folding or restriction — and **not** over the canonical form.

The rule that follows is the detector and it is one line:

> **Identical canonical form with differing byte-hash is flagged for review.**

That single comparison separates benign normalisation from injected substitution. Two records whose canonical forms agree are, to every downstream process, the same record. If their raw octets also agree, nothing happened. If their raw octets differ, something did — **and the byte-hash is the only artefact in the system that still knows it.**

Curzi [2026] measures the alternative and it is unambiguous [T2]. Hashing the canonical form gives a clean value and its Cyrillic twin the identical digest `6de95bd1c87f`; hashing the octets gives `6de95bd1c87f` and `21300fa543ad`. **A character-level hash destroys exactly the evidence that makes the detector work, and it does so silently, producing a system unable to report that it was attacked.** This is the single most consequential specification decision here and the one most likely to be got wrong by an implementer optimising for tidiness, **because hashing the canonical form is the more obviously elegant choice and it is the wrong one.**

A second-order property deserves its name. **The byte-hash is a witness, not an assertion.** It does not claim the record is true; it claims that these octets were what arrived. That is the same modest, checkable claim a hash-tree commitment makes, and it is worth exactly that and no more. Curzi [2026] states the general form: commitment is a tamper-evidence property, not a truth property.

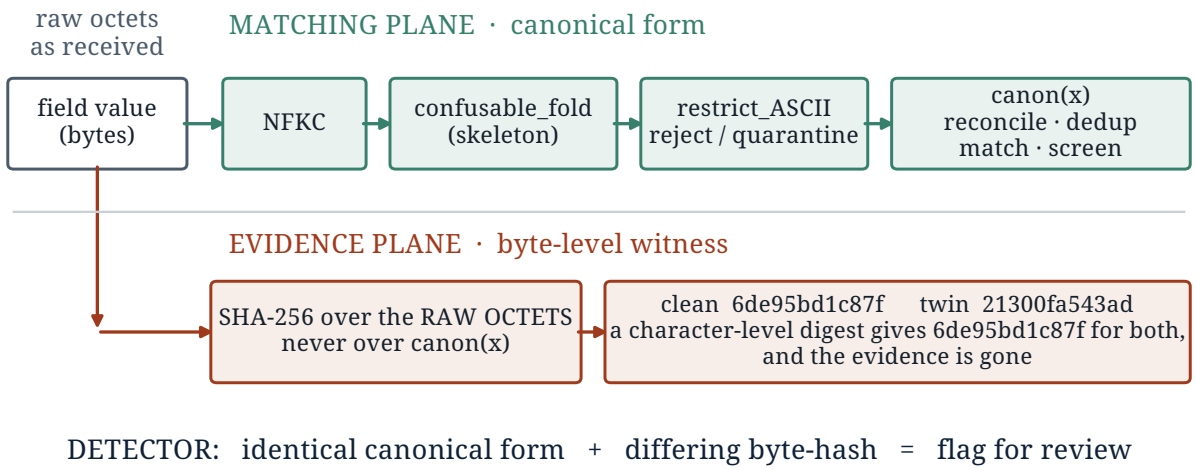


Figure 65.1: **Matching runs on the canonical form; evidence runs on the raw octets; the detector is the one comparison that reads both planes.** The digests shown are those reported in Section 4.2 for a clean value and its Cyrillic twin. Hashing the canonical form gives both values the same digest and destroys the evidence the detector needs. Schematic built by `figs/src/CCP-065_fig01.py`.

4.3 Free-text scoping

Human free text legitimately requires the whole character repertoire. Names, addresses, memos, remittance narratives and instruction text belong to people who have every right to their own scripts.

The specification therefore scopes restriction to critical machine fields only [T3]. Free-text fields store **both** the raw form and a canonical form, side by side, and are governed by the graded restriction levels of the security standard rather than by ASCII folding. Matching against a free-text field runs on the canonical form; display, storage of record, and anything a human will read run on the raw form; **and the raw form is never overwritten.**

Curzi [2026] identifies mis-scoping as the single largest risk in the approach, and this is where it is prevented. **A canonicalisation standard applied without scoping is not a control improvement but a data-destruction programme with a compliance justification, and it will be resisted, correctly, by every business line it touches.** Setting the restriction level per field and per market is a business decision with a compliance consequence, and no library performs it.

4.4 The resolution designator

Curzi [2026] measures what happens when a withdrawn numeric currency code is aliased to a live one: a synthetic 8,000 becomes a recorded live 10,000, an inflation of 25.0 per cent, **with no entry posted and nothing forged** [T2]. The remedy is stated as an accounting rule rather than a technical one [T3]:

> Withdrawn or orphaned codes are assigned to an explicit **resolution class** and are never aliased to a live settlement currency. A balance in a resolution class is visible, non-settleable, and excluded from every live-currency aggregate. It is discharged only through an authorised resolution entry that preserves a permanent audit receipt.

Three properties follow. The orphaned balance **stops distorting the statements immediately**, because it leaves the live aggregate when the code is classified rather than when it is resolved. The discharge becomes **an event with an owner, an authority and a date** rather

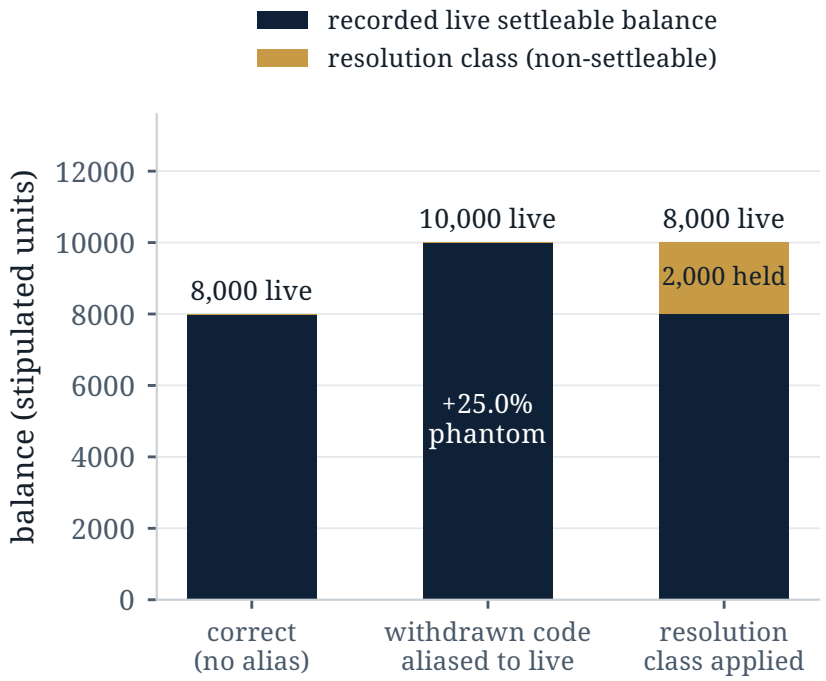


Figure 65.2: **Aliasing a withdrawn code to a live one inflates a synthetic live balance by 25.0 per cent with no journal entry; the resolution class removes the inflation without discharging the balance.** The 2,000 leaves the live aggregate the moment the code is classified and sits visible and non-settleable until an authorised entry discharges it. SYNTHETIC three-account extract; figures are stipulated units. Built by `figs/src/CCP-065_fig05.py`.

than a write-off that happens when somebody notices. And the trail is permanent and byte-addressable, so the question *what happened to that balance* has an answer that survives the departure of everyone who was in the room.

The honest status of this clause is that it is the least mature thing in the paper. It is an accounting-standard proposal. It requires a standards-body process and auditor acceptance, it has neither, and **until it has both, an entity cannot use it to justify a treatment in a set of accounts** [T3]. Appendix A says so and Section 8.3 says so again.

4.5 The normative rules

Stated with the conventional requirement keywords. **Every rule is graded [T3] — a proposal, not an adopted standard.**

Canonical form.

1. An implementation MUST canonicalise every critical machine field as restrict_ASCII ∘ confusable_fold ∘ NFKC, in that order.
2. An implementation MUST canonicalise **at every ingestion boundary**, not only at interactive entry points. File imports, partner interfaces, message parsers, migration scripts and registry extracts are ingestion boundaries.
3. canon MUST be idempotent, and an implementation MUST verify idempotence over its own field repertoire as a test, because **a non-idempotent canonicaliser cannot be applied at more than one boundary without changing answers.**
4. Restriction MUST reject or quarantine any residual out-of-set character with an explicit reason. It MUST NOT silently strip and MUST NOT silently store.

5. Reconciliation, deduplication, matching and screening MUST operate on the canonical form.
6. An implementation MUST record the character-standard version and the confusables-data version used, and MUST treat a change to either as a change-controlled event.

Byte-level integrity.

7. Each record MUST store a SHA-256 digest computed over its raw octets as received.
8. An implementation MUST NOT substitute a character-level digest for the byte-level one. It MAY store both; it MUST store the byte-level one.
9. Identical canonical form with differing byte-hash MUST be flagged for review.
10. An implementation MUST define, before deployment, the owner and escalation path for a flag raised under rule 9. **An unactioned flag queue is worse than no flag queue.**

Scoping.

11. Restriction MUST be scoped to critical machine fields and MUST NOT be applied to human free text.
12. Free-text fields MUST store both raw and canonical forms, and the raw form MUST NOT be overwritten.
13. Free-text fields MUST be governed by graded restriction levels set per field and per market.

Boundaries.

14. Every file and message boundary MUST carry an explicit encoding declaration.
15. Octets arriving without a declared encoding MUST be rejected rather than guessed at, and every intermediate tool that guesses MUST be pinned or removed from the path.

Withdrawn and ambiguous codes.

16. A withdrawn code MUST NOT be aliased to a live settlement referent.
17. A withdrawn or ambiguous code MUST be assigned to an explicit resolution class, carried as visible and non-settleable, and excluded from every live aggregate.
18. Discharge from a resolution class MUST be by an authorised entry naming the registry, the version and the basis, leaving a permanent audit receipt.

4.6 Conformance: how an implementation proves it complies

A specification that cannot be tested is a wish. Each rule has a test, and the tests are stated because they are what an auditor or a counterparty would actually ask for.

rule	test	why this test and not another
1	a golden corpus of input–output pairs covering every member of the class, checked in on the same commit and run in continuous integration	it must include one homoglyph, one compatibility form, one combining-mark sequence **with** a precomposed equivalent and one **without**, one default-ignorable insertion and one bidirectional control

rule	test	why this test and not another
3	assert canon(canon(x)) = canon(x) over the golden corpus and a generated sweep	catches the commonest implementation error: a fold written as a sequence of replacements, and therefore order-dependent
4	feed an out-of-set character and assert the call **raises with a reason**, and specifically that it does **not** return a shortened string	a stripping implementation passes every positive test and fails only this one
6	assert the recorded standard version matches the version of the data file the build actually loaded	a build recording one version while loading another is a control that drifts silently
7–9	take a clean value, produce a substituted twin, store both, assert canonical forms equal, byte digests differ, and a flag was raised	**the single most valuable test in the suite**: an implementation that quietly hashes the canonical form passes every other test in the file
15	at integration rather than unit level: submit octets with no declared encoding and assert rejection	the guessing happens in a component nobody thinks of as part of the system, and a unit test on the library will never see it
16–18	construct a synthetic extract with a withdrawn code; assert it appears in the resolution class, appears in **no** live aggregate, and **cannot be selected as a settlement referent by any code path**	the third assertion is what catches a partial implementation in which the reporting layer was fixed and the settlement layer was not

4.7 Results: the integrity layer measured (H1 and H2 supported)

Every figure below was recomputed independently for this paper on synthetic data [T2].

quantity	result
manipulated entries collapsed onto the clean canonical form	495 / 495 = 100.0%
idempotence canon(canon(x)) = canon(x)	holds for all 500
distinct raw byte strings resolving to distinct canonical forms	995 → 500 (true 500)
reconciliation false-break rate at 40% manipulation	40.0% → 0.0%
canonical forms equal for clean value and twin	True
byte-level digests differ (**the detector**)	True
character-level digests differ (**evidence lost**)	False
live balance, correct	8,000

quantity	result
live balance, withdrawn code aliased to live	10,000 (+25.0%)
resolution-class balance, segregated	2,000

H1 is supported and H2 is supported. The false-break result is the one with operational consequence: it holds at 0.0 per cent across the entire manipulation range including the degenerate case where every entry is manipulated, so the 40-to-0 reading at a 40 per cent rate is **one point on a flat line rather than the result itself** [Curzi, 2026].

The 25.0 per cent is a property of the stipulated 8,000/2,000 ratio and is not a general magnitude [T4 as to any claim that it is]. What generalises is the sign and the mechanism: aliasing a withdrawn code to a live one always inflates, always in the direction of more apparent money, and always without any entry being posted.

5. The identifier namespace

5.1 A namespace is a registry, not a string

A namespace is a set of names and a rule binding each name to a referent. **The rule is what makes it a namespace rather than a list, and the rule always lives outside the names themselves.** For the codes examined here the rule is a registry and a maintenance agency: a body that decides which codes exist, what each one currently points at, when a code is withdrawn, and whether a withdrawn code may be reused. **Every property a code has that matters for a ledger — uniqueness, stability, currency of meaning — is a property of the registry and not of the string.**

This is the observation of Curzi [2026] turned up one level. A universal character set is a registry binding code points to characters, and its confusables problem is that many code points bind to the same *shape*, so the visual layer over-collapses what the code layer keeps distinct. The standards examined here bind codes to *entities in the world*, and their problem is the mirror image and also the same: over time one code can bind to more than one referent, and two independent registries can mint the same string for different referents. **The failure is representational ambiguity — one recorded symbol, more than one possible referent, no visible difference — and it does not require a universal character set to occur.**

No standards document is cited in this section, because `bib/external.bib` contains none. One standard relied on for structured party identification is in the bibliography and is cited where it is relied on [International Organization for Standardization, 2013]. Every other designation is described from the source manuscript and itemised in Appendix B with what specifically requires verification. **No edition year, no check-digit algorithm, no registration authority and no code allocation is asserted here.**

Eight members of the namespace divide into three functional groups, and the grouping is not decorative: **it is the layering that lets a defect propagate.**

5.2 The atomic registries

Two members are atomic: their codes are not built out of other codes, and they are the substrate the rest references.

The currency registry assigns each currency a three-letter alphabetic code and a three-digit numeric code [T1]. The alphabetic code is mnemonic; the numeric code is a stable machine key that by design does not depend on the Latin alphabet and survives transliteration. Codes are added when currencies are created and withdrawn when currencies cease to circulate. **The numeric codes of withdrawn currencies persist in the standard's historical record and, far more consequentially, in the master files of every system built while those currencies**

were live. The maintenance agency publishes a live table and a separate withdrawn table; the two must be read together, and a system loading only the first has already lost the information that would let it recognise an orphaned balance.

The country registry assigns each country and territory a two-letter code, a three-letter code and a numeric code [T1]. It is the most-referenced registry in the stack, because four of the remaining six standards embed a country code inside their own identifiers. That makes it the single point whose instability propagates furthest: a change to a country code is not a change to one field, it is a change to a substring inside every securities, institution and account identifier that carries it. The registry has an explicit lifecycle — assignment, withdrawal on a change of sovereignty or name, a transitional reservation period, and, after a defined interval, the possibility of reassignment of a withdrawn two-letter code to a different country. A country code is a lease on a string, not a freehold, and a record that stored the string decades ago may now read against a different country than the one it meant [T1]. The specific historical reallocations and the exact reservation period are in Appendix B and none is named here.

5.3 The composed identifiers

Three members are composed: built by concatenating a country code with further structure and, usually, a check digit computed over the whole.

The securities identifier is a twelve-character code whose first two characters are a country prefix identifying the national numbering agency that issued it, followed by a nine-character national identifier and a check character [T1]. The prefix identifies the issuing agency's jurisdiction, which is not always the issuer's jurisdiction or the instrument's place of listing, and that gap is the source of a widespread belief that the identifier uniquely names a tradable position. It does not: one economic instrument can carry identifiers from more than one agency, and one identifier can be listed and traded in several venues under different local codes — which is why a separate classification standard exists to say what *kind* of instrument the identifier denotes.

The institution identifier is an eight- or eleven-character code: four characters identifying the institution, two a country, two a location, and an optional three-character branch suffix [T1]. It binds an institution to a routable identity, and its trustworthiness is the registrar's diligence in maintaining one institution per code and one code per operational endpoint. Its embedded country code is another place country-registry instability propagates.

The account identifier is a country code, two check digits computed over the rearranged account string, and a basic account number whose internal structure is defined nationally and differs by country [T1]. Its check digits catch transcription errors before a payment is launched, which is their entire purpose and a genuinely valuable one; they do not, and are not meant to, certify that the account exists or belongs to whom the payer thinks. The phrase *nationally defined* is the hinge of Section 5.7, because it is where an international standard defers to a purely national numbering scheme that the international standard does not itself define.

5.4 The descriptive codes

Two members describe rather than identify.

The instrument-classification standard encodes an instrument's category, group and attributes in six alphabetical characters [T1]. It exists because an identifier that is unique is not thereby informative: knowing that two positions carry different identifiers says nothing about whether they are the same class of risk, and a concentration measure of the kind Curzi [2026] builds needs the classification and not only the identity.

The legal-forms standard is a registry of the legal forms an entity can take in each jurisdiction, each rendered as a short code, so that legal forms in three languages become comparable machine tokens rather than free text [T1]. Its relevance to this series is direct: the difference

between a foundation, a trust and a company is exactly the difference Series D turns on, and this is the code list that lets that difference survive crossing a system boundary instead of being flattened into "company" by a form field with three options.

5.5 The keystone identifier: a permanent code and a perishable trust

The entity-identifier standard defines a twenty-character alphanumeric code with a prefix allocated to the issuing organisation, reserved characters, entity-specific characters and a check computed by a modulus algorithm [T1]. It is the member that comes closest to solving the identity-resolution problem Curzi [2026] names as the control underneath everything, because that is what it is designed to be: one persistent code per legal entity, globally, backed by a public reference record.

And it is the cleanest possible illustration of this section's thesis, because its value is entirely a function of its registry and the standard says so in its architecture. The identifier is not self-certifying. It is minted by a local operating unit, validated against an authoritative local source, and — this is the part that matters — **it must be renewed on a cycle, with the reference data re-verified against source, or its registration status lapses.** A lapsed identifier is still syntactically valid; it still passes its check; it still resolves to a reference record. **What it no longer carries is a current assertion that the binding is true.**

The system builds the whole thesis into its own design: the code is permanent and the trust is perishable, and the renewal cycle exists precisely because a registry binding decays unless it is actively maintained. A control that treats the identifier as a fact rather than as a dated assertion has made the same category error as a control that treats a currency code as a fact — and the registry's own status field is telling it so.

5.6 How a defect travels

The country registry is the substrate. Three identifiers embed a country code inside their own strings, so a country-code change or reassignment is a change inside a securities identifier, an institution identifier and an account identifier simultaneously. The currency registry sits beside them. The entity identifier names the holder; the legal-forms code says what form the holder takes; the classification code says what class the security is. **A single settlement instruction touches most of the namespace at once**: it pays an *entity*, of a *legal form*, in a *country*, at an *institution*, into an *account*, a sum in a *currency*, possibly against a *security* of a *class*.

Every one of those is a registry lookup, and every registry has the same three failure surfaces: a code can be **withdrawn** and its meaning orphaned; a code can be **reassigned** and its meaning silently changed; and two registries can **mint the same string** for different referents.

The composition multiplies the surface: because the higher codes are built from the lower ones, a defect does not stay where it started. A country reassignment is felt in three other standards and a currency orphan is felt anywhere the currency is aggregated. This is a coupled-network fragility of exactly the kind Buldyrev et al. [2010] formalises — a failure in one layer propagating into a layer that depends on it — and it is the structural reason the eight are treated here as one namespace rather than as eight standards. The analogy is used for its shape and not as a theorem about registries; Section 9.6 records that.

5.7 Two collisions on one digit string, and the reading that is refused

The sharpest instance of the hazard concerns a three-digit string that occurs in two independently governed registries, and it must be stated with precision, **because the defensible finding and the indefensible one are close together and the difference is the whole of the section's credibility.**

The first collision is within the currency registry itself: **one numeric code denotes a currency in a pre-redenomination form and is withdrawn, while a different numeric code denotes the live post-redenomination currency.** Two codes in one standard's history both mean, loosely, the same national currency, separated by a redenomination. That is a straightforward instance of the hazard of Section 4.4: a legacy system carrying a balance under the withdrawn code, and a mapping table that at some migration was given a row sending it to the live currency, reports a withdrawn-code balance as live settleable money. **The exact historical allocation of the withdrawn code between two candidate predecessor currencies is contested in the source's own verification register, is not asserted here, and is itemised in Appendix B. The mechanism does not depend on which one it was; the naming does.**

The second collision is distinct and shares the identical digits. **A national domestic account-numbering scheme encodes the account's currency inside the account number itself, and in that scheme the same three digits designate the live national currency — a different registry, on a different object, colliding on the same three characters** [T1]. It is not the currency standard. It is a national account-number grammar of exactly the kind the international account standard defers to when it says the internal structure is defined nationally. **In one registry the string is a withdrawn currency code; in the other it is a live currency marker inside an account identifier. The two registries never coordinated, they mean different things by the same digits, and both are correct within their own scheme.**

One reading of the second collision is refused outright and graded T4. There circulates a pseudo-legal theory holding that because the digits appear in both places, accounts in that jurisdiction secretly denominate a parallel legacy currency alongside the official one, so that every account is a hidden double ledger and the official balances are a fiction. **The source manuscript records that this reading has been rejected by the courts that considered it. This paper does not make it, in any form.** The collision of the digits across two registries is a fact; **the inference that it evidences a concealed parallel accounting system does not follow from that fact**, and stating it would be the precise error — an encoding coincidence dressed as a macroeconomic revelation — that Curzi [2026] identifies as destroying a technical submission on contact. The court decisions are described from the source and are itemised in Appendix B; none is cited.

The defensible finding is narrower, duller, real and computable: a cross-border parsing ambiguity [T2]. A screening or reconciliation process that ingests raw records from more than one jurisdiction, and that scans for the digit string to flag withdrawn-currency exposure, cannot on the string alone tell a genuine withdrawn-code field from an occurrence of the same digits inside an account number where they are a live currency marker. The two are the same three digits in the same record stream; **the difference is which grammar the field is parsed under, and a screen that parses bytes rather than typed fields does not know the grammar.**

The fix is to parse each field under its declared grammar — and it requires nobody to determine, or allege, anybody's intent. That clause is the entire reason the finding is safe to make and the double-ledger reading is not: the parsing fix is intent-agnostic in exactly the way Curzi [2026] insists the whole mechanism must be.

5.8 Results: the parsing ambiguity measured (H3 supported)

A **SYNTHETIC** demonstration, deterministic under a fixed seed [T2]. Two thousand synthetic twenty-digit domestic account numbers, of which twelve hundred carry the digit string in the grammar's currency positions and thirteen carry it elsewhere as an incidental substring; and two thousand synthetic numeric currency fields, of which one hundred and fifty are the genuine withdrawn code. A naive cross-border screen scans every raw record for the substring and flags it as *withdrawn currency present*.

quantity	naive substring screen	typed-field parse
records flagged	1,363	150

quantity	naive substring screen	typed-field parse
genuine withdrawn-code events among them	150	150
parsing false positives	1,213	0
precision on the stated purpose	11.01%	100%
parsing-ambiguity share of the queue	88.99%	0.00%

Almost nine flagged records in ten are not withdrawn-currency exposure at all; they are live account numbers that happen to contain the same three digits, and the screen cannot tell the difference because it is reading bytes instead of typed fields [T2]. **A control that is right one time in nine is not a control; it is an alert generator, and it will be switched off in the second week** for exactly the reason rule 10 gives.

Under the typed parse, the currency field is parsed as a currency code and a withdrawn value resolves against the withdrawn table to a resolution class under rule 17; the account field is parsed under the national grammar and the digits in the currency positions resolve as the live marker they are; and the incidental occurrences are never examined as currency at all, **because the grammar never looks there.** The genuine withdrawn-code balances are isolated where they are visible, non-settleable and awaiting an audited discharge; the live accounts are left alone; and **the coincidence of the digits stops generating noise because nothing is any longer reading one registry's codes under another registry's grammar.**

This is the encoding thesis and its remedy transposed from the character to the code: **the defect is representational ambiguity, and the fix is to make the representation carry its own type, so that the same string in two namespaces stops being one string.** Ashby [1958] is the general statement of why a byte-level screen cannot succeed here: **a regulator must command at least as much variety as the system it regulates**, and a screen with one category — *contains the substring* — cannot separate two categories of referent. The requisite-variety result is invoked for its shape and not as a theorem about screening; Section 9.6 records that.

5.9 The resolution designator, generalised across the namespace

Section 4.4 stated the designator as an accounting rule for withdrawn currency codes. It generalises, because the currency case is only the sharpest instance of a hazard every member shares [T3]:

> A code whose registry binding is withdrawn, lapsed, reassigned or ambiguous MUST NOT be silently resolved to a live referent. It MUST be assigned to an explicit resolution class, carried as visible and non-settleable, and discharged only through an authorised entry that names the registry, the version and the basis of the resolution, and that leaves a permanent audit receipt.

A withdrawn currency goes to a resolution class. **A reassigned country code carried on a legacy record goes to a resolution class until the record is confirmed against the registry version current when it was created**, because the same two letters may now read against a different country. **A lapsed entity identifier resolves by re-verification against source, not by silent acceptance of the last-known reference data.** The generalisation is one rule: **a registry binding that has moved is not a fact to be resolved silently to whatever the code now points at; it is an exception to be surfaced, dated, and discharged by someone with the authority to say what it means now.**

The status is unchanged: a proposal, no adopted standard behind it, and **the parts touching how a balance is presented require auditor acceptance before use.**

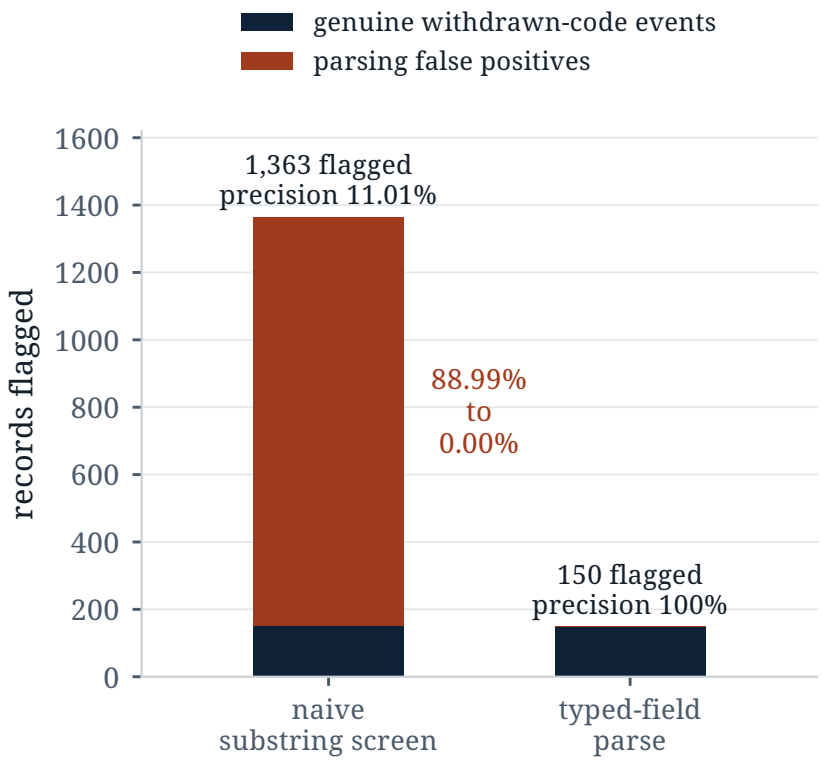


Figure 65.3: **A control right one time in nine: typed-field parsing removes the entire parsing-ambiguity share.** A naive substring screen flags 1,363 records of which 150 are genuine withdrawn-code events and 1,213 are live account numbers containing the same digits, a precision of 11.01%. Parsing each field under the registry that owns it isolates the 150 with zero false positives. SYNTHETIC corpus, fixed seed, stipulated class mix. Built by `figs/src/CCP-065_fig02.py`.

5.10 The hierarchical-namespace pattern, named and deferred

There is a larger pattern here and this section names it precisely enough to place the namespace within it, then hands the full treatment elsewhere rather than pre-empting it.

Several of the world's addressing systems are **hierarchical namespaces**: names built by nested delegation, where an authority allocates a block of the space to a sub-authority, which allocates sub-blocks in turn, so that the structure of the name mirrors the structure of the delegation. The international telephone numbering plan is one; the internet address space is another; nested geographic coding is a third. The namespace examined here shares the pattern, since the country registry delegated into three composed identifiers is nested delegation of exactly that kind.

These systems share a structural pattern and not a common origin. They were designed independently, by different bodies, for different media, and the resemblance is convergence on the same solution to the same problem rather than descent from one system [T1]. **The distinction is not pedantry.** A claim of common origin would invite a reader to expect interoperability, shared governance, or a single point of failure that does not exist; a claim of shared pattern invites only the observation that the *failure modes* rhyme, because a hierarchical namespace has characteristic weaknesses — delegation boundaries where provenance is lost, exhaustion of a level, reassignment of a freed name — wherever it appears. Curzi [2026] works that observation fully and reaches the general result that **every translation boundary rewrites address data and is a provenance-loss point**, structurally the same event as an encoding boundary here; this paper states the pattern and stops.

6. Structured messaging and resolvability

6.1 What the legacy message could not carry

The legacy cross-border payment message was designed under expensive bandwidth and a telex heritage, and its constraints are physical rather than conceptual [T1]. A party is identified in a small number of fixed-length lines. The character set is restricted, close to ASCII, so a name with diacritics is transliterated on the way in and the diacritics are gone. A name longer than the field is truncated, and **the truncation falls at the end, which is frequently where the distinguishing part of a legal name lives** — the form, the branch, the trailing words that separate one similarly named entity from another. There is no dedicated element for an entity identifier, so **the one piece of data that would resolve the party deterministically has nowhere to go** and travels, if at all, buried in a free-text reference line where no parser expects it. There is no coded purpose.

The consequence is that a legacy message under-specifies the parties structurally, not by anyone's choice. The field limits mean that even an originator who wants to declare everything cannot: completeness is not available as an option, and every party the message names is named ambiguously by construction.

6.2 What the structured standard carries instead

The structured standard is a typed syntax over a defined data dictionary, so that a payment is a tree of named, typed elements rather than a set of fixed lines [International Organization for Standardization, 2013]. A party is a structured object: a name element holding the full legal name at length; an organisation-identification element holding the entity identifier; a postal-address element broken into named sub-elements including a country code; and, at payment level, a coded purpose drawn from an external code list. **The namespace of Section 5 is precisely what these fields carry.**

The improvement in party identification is large and real. A structured message can carry the full legal name without truncation; it can carry the entity identifier, which resolves the

party against a registry deterministically rather than by fuzzy name match; it can carry the address as components a machine can compare rather than as lines a machine must parse; and it can carry a coded purpose a monitoring system can reason about instead of a phrase it must interpret. **A screen fed a structured message with an entity identifier is doing deterministic identity resolution against a registry key; a screen fed a legacy message is doing fuzzy string comparison against a truncated, transliterated name — and the difference between those two is the difference between a control that clears and a control that alerts.**

There is a further improvement easy to miss because it looks like tidiness: **the addressability of the individual field.** In the legacy message the party is a block of text, and a monitoring system that wants to reason about the party's country, or its legal form, or the payment's purpose, must extract those facts by parsing — which is to say, by guessing, against a format that varies by originator, market and decade. In the structured message each is a named element carrying a typed value. **The monitoring system does not parse; it reads the element it wants by name and gets a typed value**, so questions like *is this purpose consistent with these parties* become questions about typed fields, and a question about a typed field has an answer a machine can compute.

6.3 The resolution asymmetry

A screening system's job is to dispose of a message: to decide, fast and mostly without a human, whether a payment may proceed or must stop. It does this by trying to resolve each party to a definite identity and checking that identity against designated-persons lists. When it can resolve the party deterministically it clears or stops with confidence, and **either way it disposes.** When it cannot resolve the party — when all it has is a truncated, transliterated name fuzzy-matching several list entries — **it cannot dispose, and a message a machine cannot dispose of goes to a human.**

A rich message is disposable by the machine; a sparse message is not, and the review queue fills with exactly the messages carrying the least data, because carrying the least data is what makes a message unresolvable.

This inverts the intuition that carrying less reveals less. **The sparse message does not present as private; it presents as opaque, and opacity to a monitoring layer is, to that layer, indistinguishable from evasion.** The originator who minimises disclosure produces the message that cannot be cleared and therefore attracts the human attention, the questions, the requests for supporting documentation, the delay, and eventually the exclusion from correspondent relationships that Curzi [2026] measures as a concentration effect: naive correspondent concentration of 0.302400 — 3.306878 effective correspondents — collapsing to 0.745256 and $N_{\text{rail}} = 1.341819$ once flows are decomposed by currency. **Being unresolvable is one of the mechanisms by which a structure loses correspondents, and losing correspondents is what collapses the effective count that Curzi [2026] measures.**

Two things must be said about the claim, because it is the one most easily overstated.

First, **the value is in the tail rather than the median.** Full population does not speed up the median payment, which was going to clear anyway. It removes the structure from the tail — the tail of payments that draw scrutiny, that get held, that trigger a request. **A sparse message is a lottery ticket in that tail on every send; a rich message is not in the draw.**

Second, **the posture is available only to a structure with nothing to lose by resolving.** A structure whose elements are true can declare them all at no cost; a structure whose elements are not cannot. **That asymmetry is the whole point, and it is what makes full population a costly signal rather than a policy**: it is expensive for the unclean to fake and cheap for the clean to produce. **It is also, precisely, the disclosure test of Section 3.3 executed on every payment rather than once on a design.**

6.4 Results: rich against sparse, on one payment (H4 supported, part one)

A **SYNTHETIC** demonstration, deterministic under fixed seeds [T2]. A synthetic designated-persons list of fourteen structured entries, each with a legal name, an entity identifier and a country code, **constructed to include near-duplicate names so that a name-only match is genuinely ambiguous** — which is the whole point of using identifiers. A constructed target entity holds its own distinct identifier, is registered at an address carrying diacritics, and — by design — **is not on the list**. It shares a truncatable name stem with two list entries.

Two messages carry the same payment to that target.

	richly typed message	sparse legacy message
entity identifier present	True	False
identifier matches a list entry	False → deterministic clear	—
name as carried	full legal name, diacritics intact	truncated and transliterated
fuzzy name candidates	2 (informational)	2 (decisive)
disposition	**auto-clear on the identifier; identity certain**	**cannot deterministically clear; manual review**

The identifier resolves the target to itself and clears it against the list on one exact-key comparison; the two fuzzy candidates are informational noise the identifier overrides, and no human sees the message. With no identifier and a name fuzzy-matching two list entries, the screen cannot dispose and the message goes to a human — **not because the payment is suspect but because it is unresolvable, which is the only thing the screen actually detected** [T2].

The two messages describe the identical lawful payment to the identical lawful entity. One clears in a machine cycle; the other consumes a review analyst. The difference is entirely the data the message carried.

6.5 Results: the queue at population scale (H4 supported, part two)

A synthetic screening day of one thousand messages, with the migration part-complete so that 62 per cent are richly typed and 38 per cent sparse, a small genuine list-membership rate, and a modelled rate at which a sparse truncated name collides with at least one list stem and so requires review [T2].

quantity	value
synthetic day	1,000 messages (620 rich / 380 sparse)
manual-review alerts from rich messages	6
manual-review alerts from sparse messages	92
total review queue	98
share of the review queue that is sparse	**93.9%**
sparse share of traffic	38.0%

Sparse messages are 38 per cent of the traffic and 93.9 per cent of the review queue: the message format carrying less data generates the overwhelming majority of the scrutiny, and it does so independently of whether any payment is of concern [T2]. The rich messages that do alert are the six genuine list hits — the ones the control is *supposed* to stop — and they alert cleanly, on an identifier, disposable with confidence. **The sparse queue is 92 messages of which almost all are lawful payments that could not be resolved.**

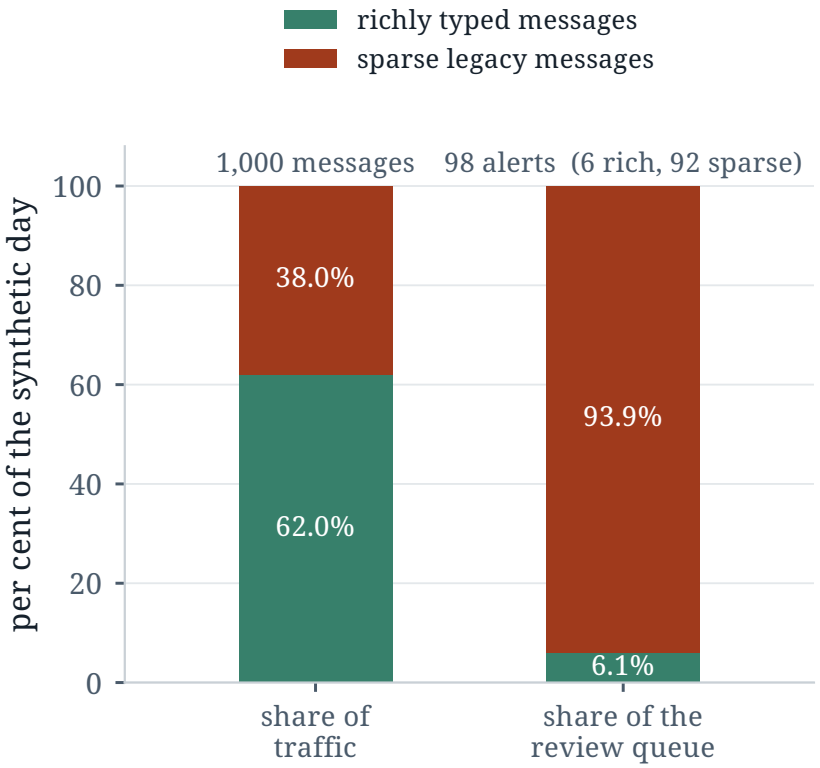


Figure 65.4: **The message format that carries less data generates the overwhelming majority of the scrutiny.** A synthetic day of 1,000 messages, 620 richly typed and 380 sparse; 6 review alerts from the rich population and 92 from the sparse one. Sparse messages are 38.0% of traffic and 93.9% of the queue. Stipulated rates, not a measurement of any production network. Built by `figs/src/CCP-065_fig03.py`.

The numbers are a property of the stipulated rates and are a demonstration of the mechanism rather than an estimate of any real network's queue. Section 9.10 records an arithmetic discrepancy found in the source's own statement of those rates and prints it rather than harmonising it.

The queue is the right object to reason about, and the reason is a standard one. **A review function is a service system with finite capacity, and its waiting time rises without bound as utilisation approaches one** [Kingman, 1962, Hopp and Spearman, 2011]; Little [1961] fixes the relation between the queue length, the arrival rate and the time in system. **Message richness is an arrival-rate parameter on the review queue**, and reducing it is therefore a capacity decision rather than a compliance gesture. The queueing results are invoked for their shape and not as a model of any screening function; Section 9.6 records that.

6.6 Results: the substitution inside the structured field (H5 supported)

A section that sold the structured standard as the end of ambiguity would repeat, one standard higher, the error Curzi [2026] forbids.

The structured message is typed text. The name element holding the full legal name is a text field, and a text field in a structured document is a text field: **it accepts the whole character repertoire, and every member of the substitution class enters it unchanged** [T1]. A homoglyph sits in a name element as comfortably as it sat in a delimited file. **The structured message has a defined place for the name and no defined canonical form for it, and the two are not the same thing.** Structure tells a consumer which field holds the name; **it does not tell the consumer that two names rendering identically are the same name**, and that second question is the whole

subject of the integrity layer.

Measured on the target's own name element [T2]:

quantity	result
clean fragment, byte digest	`c413622fa3bea888`
homoglyph twin, byte digest	`7307655b3eee8d17`
do the two render identically?	yes
canon(clean) = canon(twin)?	**True** — the fold collapses it
byte digests differ?	**True** — the witness records it

Inside the richest, most structured, most authoritative field the standard defines, a homoglyph twin is byte-distinct and canonically identical — the exact silent-duplicate condition Curzi [2026] measures at 99 per cent — and the structured message does nothing about it while the canonical form folds it and the byte-hash witnesses it [T2].

This is not a flaw in the standard; it is the boundary of what a syntax can do. A syntax defines the shape of the data. Canonicalisation defines the identity of the values. **The structured message is where the integrity layer stops being a proposal about a delimited file and becomes a control on the payment rail** — the canonical form is what the screen matches, the byte-hash is the witness that the field was not silently substituted, and the two together are the detector of Section 4.2.

6.7 The remedy, specified for the message

Carried from Section 4.5 and scoped to the structured message [T3].

1. An implementation ingesting structured messages MUST canonicalise the critical machine fields — entity identifier, currency code, account and institution identifiers — at the ingestion boundary, **before** screening, reconciliation or matching.
2. It MUST match free-text name and address elements on their canonical form while preserving the raw form unaltered for display and record, governed by restriction levels per market, and MUST NOT fold a free-text name for storage.
3. It MUST store a digest over the raw octets of each ingested element as received, and MUST flag identical canonical form with differing byte-hash, with a named owner and escalation path defined before deployment.
4. Screening MUST prefer the deterministic identifier key over the fuzzy name match where the identifier is present, and **MUST record when it fell back to name matching because the identifier was absent** — because that fallback is the unresolvability event and it should be visible.
5. A withdrawn or ambiguous code carried in any message field MUST be handled under the resolution designator of Section 5.9 and MUST NOT be silently resolved to a live referent.

Rules touching internal controls need no external adoption and may be implemented on the institution's own authority; none of them changes a disclosed number. The whole remains a proposal with no adopted standard status.

7. The ledger architecture, with its audited negative results

Everything in this section concerns the author's own systems and the author's own adversarial audit. The conflict of interest is declared in the Declarations. The measurements were recomputed for this paper on a harness rebuilt from the source descriptions; **the original code was not obtained and was not rerun**, and every finding describes the implementation *as audited*.

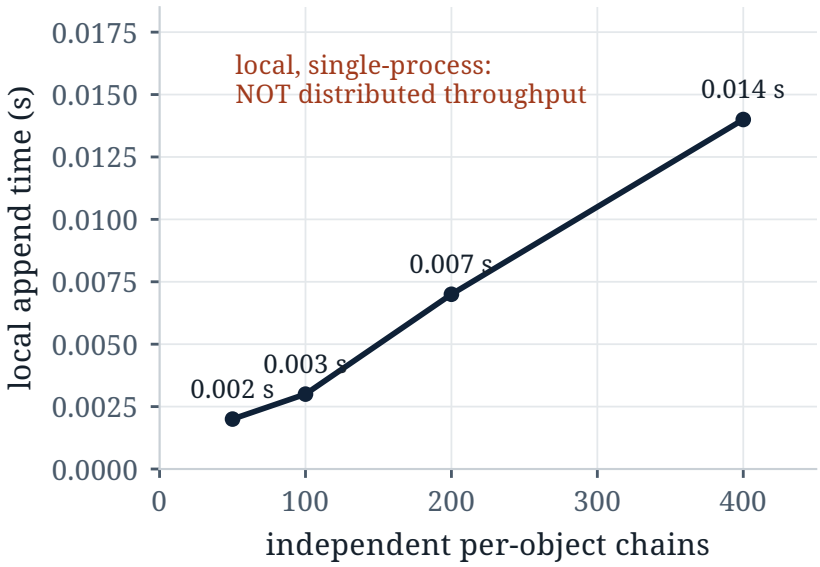


Figure 65.5: **Local append cost is linear in the number of independent per-object chains — and it is a local, single-process measurement, not distributed throughput.** 50, 100, 200 and 400 chains at 0.002, 0.003, 0.007 and 0.014 seconds, measured on one machine in one process on a harness rebuilt from the source description. Absolute times are machine-specific; the shape is the finding. Built by `figs/src/CCP-065_fig04.py`.

7.1 The object-as-ledger factoring, measured (H6 supported)

Most ledgers are one shared chain whose global order is agreed by a consensus protocol. An alternative factors the ledger into per-object chains, so that independent objects can be written without a global write-lock, **at the cost of needing a separate mechanism to order conflicting updates.** The substrate examined here takes the per-object idea down to the individual token: **every token is its own hash chain, seeded by a genesis root, and tokens link to compatible neighbours to form a directed acyclic graph** [T1].

check	construction	result
V1 per-object chain integrity	200 tokens × 30 transactions	200 / 200 verify
V2 tamper-evidence	one block mutated in each chain	200 / 200 detected
V3 cross-object link consistency	28 links formed	28 / 28 recorded in both ledgers
V4 local append cost	50 / 100 / 200 / 400 chains	0.002 / 0.003 / 0.007 / 0.014 s

V2 is the property worth having: because each block's hash covers its transaction and the chain links by predecessor hash, any silent edit breaks verification, every time, with no exceptions in two hundred trials. **V4 is linear and lock-free**, which is the write-parallelism the factoring exists to provide — **and it must be read for exactly what it is. These are local append times on one machine in one process. They are not distributed transaction throughput and no throughput claim is made from them.** The absolute figures are machine-specific and are printed only to show the shape.

7.2 Four findings that bound the substrate

A1 — the structure provides no consensus and does not prevent double-spending [T4, refuted]. A token is taken, a transfer recorded, and the chain then extended into two forks, the second

transferring the same balance to a different party. Reproduced here: **both forks verify as internally valid hash chains, and nothing in the object-as-ledger structure imposes a total order or rejects the conflict. The same token is spent twice and the data structure raises no objection.** The folk claim that per-object chains need no consensus is false: **they remove the global write-lock, and conflict ordering still requires a consensus or voting layer.** Agreement among mutually distrustful parties is the Byzantine agreement problem [Lamport et al., 1982], whose bound on genuinely independent parties is the reason an ordering layer is not optional; Fischer et al. [1985] and Gilbert and Lynch [2002] fix what such a layer can and cannot achieve, and Douceur [2002] fixes why the count of parties must itself be audited.

A2 — the block hash under-binds [T2]. In the implementation as audited, a block's hash is computed over the transaction alone and does not commit to the chain identifier, the predecessor hash or the index. Tamper-evidence still works, because the predecessor link breaks when an earlier transaction is edited — **but the hash does not bind a block to its position, so blocks can be reordered and identical transactions collide.** Reproduced here: a two-block chain with its blocks transposed and its predecessor links repaired **verifies successfully under the shipped hash and is rejected under the fix.** The fix is to compute the digest over the chain identifier, the predecessor hash, the index, the timestamp and the transaction together, **binding each block to its chain, its position, its time and its content, and eliminating reordering, cross-chain replay and identical-transaction collision in one change.**

A3 — the numerological layer carries no security [T4]. Assessed in Section 3.4 and excluded there.

A4 — no Sybil resistance [T4]. Tokens and nodes are created without any cost-to-identity, so influence over the topology or over any future vote is forgeable by mass identity creation [Douceur, 2002]. The fix is a cost-to-identity or verified-identity layer.

The verdict, stated as the source states it: the substrate delivers tamper-evidence and local write-parallelism; it does not deliver consensus, total order, double-spend prevention or Sybil resistance. It is a sound substrate and not a complete distributed ledger, and it becomes one only when paired with an ordering layer and an identity layer.

7.3 The capability layer, and four findings that were fixed

The identity and authorisation layer names an entity by a content-derived identifier rendered as a locator, and authorises by capability contract rather than by identity plus access-control list. Access is negotiated by a four-message challenge–response exchange with per-session derived keys and keyed message-authentication tags. **The literature this places it in is named in Appendix B and is not cited, because none of the primary sources is in the bibliography;** the protection principles it rests on are Saltzer and Schroeder [1975].

Reproduced for this paper [T2]: content addressing over 20,000 distinct inputs yields 20,000 distinct identifiers with zero collisions — **which is what a correct hash-naming scheme does, and is worth confirming rather than assuming.**

Then the audit, which found four concrete breaks in the reference implementation and which the author published in full.

R1 — the proof signature was never verified [T4, fixed]. The verifier checked nonce presence, a self-computed digest, holder authorisation and compliance flags, **and never inspected the signature field at all.** A message carrying an all-zero signature was accepted. **Origin authentication failed completely**: any party able to name an authorised holder was admitted regardless of whether they possessed the material behind the identifier.

R2 — the self-verification digest was computed inconsistently [T4, fixed]. The prover computed it keyed, over one message; the verifier recomputed it unkeyed, over a different message. The two never agreed, so the check either rejected legitimate exchanges or was bypassed. **A digest a verifier can recompute without a secret authenticates nothing.**

R3 — broken hash algorithms were offered [T4, fixed]. Two algorithms known to be broken for collision resistance were available for content addressing. **In an identity system a**

collision means two contents sharing one identity, which is catastrophic rather than merely undesirable.

R4 — the freshness root left a replay surface [T3, fixed]. Nonces were tracked per request and freshness rested on an epoch root of unspecified origin. **If that root is not verifiable, monotonic and bias-resistant, a stale epoch can be pinned and a captured exchange replayed.**

The fixes were implemented in a hardened module and **re-validated rather than merely proposed**: one keyed transcript tag computed identically by both parties and verified in constant time, weak hashes refused at key creation with a hard error, a verifiable epoch root with a nonce time-to-live and a persisted seen-nonce set, and keys derived by a standard derivation function. **The re-validation confirms that a legitimate handshake is accepted while a forged signature, a tampered transcript, a weak hash and a replayed nonce are each rejected, and the protocol's shape is unchanged — only its verification is made sound.** That is what an adversarial review is for: contrary evidence used to harden a design rather than to discard it.

On post-quantum posture the source is equally direct, and its directness includes a refutation of the author's own earlier marketing. The hash and tag core is post-quantum-resistant at digests of 256 bits or more, because Shor [1997] breaks the number-theoretic public-key constructions and this core uses none of them, while Grover [1996] gives only a quadratic speedup and halves effective preimage security. **The signature the protocol needs and did not have must be a standardised post-quantum scheme.** The relevant standards are described in the source as published in 2024; **this paper names none of them by number and Appendix B records the verification required. The claim in the author's earlier materials that the relevant cryptography is "ready for immediate implementation" and "unbreakable" is refuted, in the author's own paper, and is graded T4.**

7.4 A backing taxonomy is not a valuation method

The source corpus specifies a backing taxonomy of nine forms of capital — financial, material, living, intellectual, experiential, social, cultural, spiritual or ethical, and natural — with the declared use that an instrument is backed not by a single reserve asset but by a portfolio drawn across the taxonomy.

Two things must be said and they pull in opposite directions.

The taxonomy addresses a real problem, and the problem is the completeness gap Curzi [2026] measures at 36.150000 per cent structurally unrepresentable on a constructed vector. Double entry represents only what has been monetised, and a great deal of what an entity owns is invisible to it not because it lacks value but because no transaction ever priced it. **A structured taxonomy of forms of capital is a coherent way to name what the ledger cannot see.**

And a taxonomy is not a valuation method. Naming nine forms of capital does not tell anyone what any of them is worth, and the valuation problem is the entire difficulty. Some forms have a defensible and contested valuation methodology; at least one has **no generally accepted measurement basis whatsoever**, and **a coverage ratio computed over a portfolio in which one component has no measurement basis is a coverage ratio computed over an unknown.** This paper records the taxonomy as declared [T3] and states plainly that **the measurement basis for the non-financial forms is not specified in the source corpus, is not specified here, and would have to be specified, audited and accepted before any coverage ratio computed across them could mean anything** [T4 as to any claim that it has been].

This is the corpus's scalar-reduction habit arriving in an engineering artefact, and the limitation is inherited from Curzi [2026] rather than restated: where a source describes a phenomenon by several jointly varying features, the framework replaces it with one number and does arithmetic on it. **Nine forms reduced to nine shares and then to one coverage ratio is exactly that move.** Curzi [2026t] is the corpus's own precedent for what to do about it: a quantity whose components are not instrumented on a common ratio scale with a shared true zero has its magnitude withdrawn as a measured quantity, and the same verdict applies here.

7.5 The third entry, in two formulations, not adjudicated

Conventional bookkeeping runs on two entries. The recording discipline proposed here adds a third.

The source corpus defines the third entry in one place as **proof** — an immutable cryptographic witness addressed by a content identifier — which places it in the lineage Curzi [2026] sets out. A later document in the same corpus, specifying a physical instrument, refines the arrangement: the third rail becomes **equity**, recorded atomically and continuously alongside debit and credit rather than derived as a periodic residual, with proof retained as the integrity witness securing the equity rail. The stated rationale is that equity in conventional accounting is a computed leftover, disclosed periodically, which is what permits liabilities to expand while ownership is reckoned after the fact.

This paper records both formulations, notes that they are not the same design, and does not adjudicate between them [T3]. What it will state is what the third entry can and cannot do. **A cryptographic witness proves that a record existed in a given form at a given point in a chain. It does not prove that the record was true, that the asset existed, that the valuation was defensible, or that the counterparty agreed** — and a third entry presented as establishing any of those has been oversold in exactly the way this paper's own discipline forbids. **Tamper-evidence is a real and valuable property. It is not verification.**

7.6 How the layers compose

The composition is the argument, because each layer supplies exactly what the one below it lacks.

Canonicalisation makes two records comparable. Without it, every layer above computes over an identifier space in which one thing can appear as many — **and a content-addressed system inherits the defect in its worst form, because if the identifier *is* the digest of the content, then a homoglyph substitution does not produce a suspicious record, it produces a different entity, with a different address, indistinguishable from a legitimate one. Content addressing does not fix confusable identifiers; it launders them, by converting a visual ambiguity into two cryptographically impeccable and permanently distinct names.** That is the strongest argument for putting canonicalisation underneath everything rather than beside it, and it is the property Curzi [2026] builds on: one byte's difference in a settlement object yields a different identifier, while the canonical form folds a homoglyph account string onto the same skeleton as the clean one.

Per-object hash chains then make each object's history tamper-evident. Content-addressed capabilities supply the identity and authorisation the chains lack, addressing A4. **An ordering layer supplies the conflict resolution that addresses A1, and is not supplied by anything in this paper.** And the recording discipline of Section 7.5 makes the stack an accounting system rather than a database.

8. Discussion

8.1 What follows

An integrity layer at the comparison boundary is cheap, testable, and requires nobody's permission. Rules 1 to 15 are a library, a set of boundary changes and a test suite. They are internal controls, they change no disclosed number, and an institution may implement them on its own authority. **Their whole cost is a library call at each ingestion boundary and the discipline to define the flag owner before the first flag is raised.**

The identity resolution underneath a concentration measurement is not free of the same defect. The effective count of Curzi [2026] — the reciprocal of the Herfindahl index, which is the Hill number of order two [Hill, 1973, Jost, 2006, Adelman, 1969] and not the corpus's affine

construction — is a function over a partition of exposures grouped by counterparty, jurisdiction or instrument. **Count one counterparty as two, or read a legacy country code against the country it now names rather than the one it named, and the measured concentration moves without anything real changing — and it moves in the reassuring direction.** Curzi [2026] traces that path from the counterparty master; Section 5.3 adds the securities identifier as a second path into the same instrument, since **one economic instrument carrying identifiers from two agencies is a double-count and two venue codes resolving to one instrument is a merge, and either error moves a single-name limit.** Goodhart [1975] and Campbell [1979] are the standing warning that an index adopted as a target stops measuring what it stood for; this is the same warning with a mechanical rather than a behavioural cause.

A lawful structure has a direct operational interest in being resolvable. Section 6 is not an argument that disclosure is virtuous. It is an argument that the network now penalises unresolvability, that the penalty falls in the tail rather than the median, and that **the posture of full population is available only to a structure with nothing to lose by resolving.** The corollary matters as much: **a structure that populates fields it cannot truthfully fill has re-entered the concealment game from the other side, and the invented fields are what will not survive verification.**

Two conversations decide whether any of this survives contact with a real environment, and both are short. A correspondent's financial-crime function asks one question about screening, and it is not whether there is a policy: *demonstrate that a name on the list, presented with a single substituted character, is matched by your screen.* Under the specification the answer is a test case and a log entry. Without it the answer is a discussion about supplier configuration that neither party can conclude. An auditor asks a different question, and it is about the flag queue rather than the control: *who reviewed the flags, on what basis, and what was the disposition?* **A control producing exceptions nobody dispositions is a finding, and it is a worse finding than not having the control, because the exceptions are evidence that something was detected and ignored.** That is why rule 10 is normative rather than advisory.

Neither conversation is about the architecture of Section 7. They are about two rules out of eighteen, and **that proportion is a reasonable guide to where the value in this paper actually sits.**

8.2 What does not follow

Necessary, not sufficient, and it holds for the whole stack. Canonicalisation defeats encoding-level evasion and does nothing about semantic fraud. Typed parsing defeats code-level ambiguity and does nothing about a correctly coded false assertion. Per-object hash chains defeat silent edits and do nothing about a correctly signed lie. Content-addressed capabilities defeat impersonation and do nothing about a properly authorised party acting in bad faith. Full message population defeats the identity gate and does nothing past it. **A tamper-evident ledger containing a false entry is a tamper-evident false entry, and this architecture will preserve it faithfully forever** [T1/T4].

Every number in this paper measures a representational or resolution property. None measures truthfulness, and a claim otherwise would fail on contact for exactly the reason Curzi [2026] gives in its audit of nine documented failures, eight of which no encoding control would have touched.

Collected in one place, because scattering them is how a specification becomes an advertisement, the architecture provides: no consensus and no total order; no double-spend prevention; no Sybil resistance; **no distributed throughput claim of any kind** — the append measurements are local and single-process; no valuation methodology for the non-financial forms of capital; no adopted standard status for the canonical form and none for the resolution designator; no auditor acceptance for any accounting treatment described; and no legal force anywhere. **Each is fixable in principle and several by composition with well-understood components. None is fixed here.**

Over-canonicalisation is the strongest technical objection, restated because rule 11 exists

to prevent it and rules can be ignored. Aggressive folding erases legitimate internationalised data: genuine non-Latin names, multilingual memos, diacritics that carry meaning. **An implementation that applies rule 1 to a name field has not partially complied; it has failed.** The mirror error appears inside the namespace itself: **the reference records the registries carry hold legal names in their own scripts for good reason, and a control that folds them to compare them has mutilated the registry's own authoritative data.** The scoping rule is not suspended because the data now arrives inside an international standard.

The identifier-collapse trade is real and is not priced here. Folding confusables collapses the identifier space, and **collapsing an identifier space creates collisions**: two entities whose names genuinely differ only by characters folding to the same skeleton become, after canonicalisation, the same name. For critical machine fields with narrow permitted character sets this is a non-issue by construction. **For a name estate it is a real trade** — canonical matching raises the true-positive rate against substituted probes to 100 per cent, and it also raises the false-positive rate against a large list by an amount that has not been measured here and is not zero. **The zero-false-positive figure Curzi [2026] reports is measured against clean ASCII critical fields and does not transfer to a name estate; anyone quoting it as a general property has quoted it wrongly.**

Three further misreadings are worth naming. *Richness read as safety*: a rich message resolves the party and does not verify the party's conduct, the transaction's reality or the counterparty's honesty. *A coded purpose read as a verified fact*: the code makes monitoring tractable, not true, and a monitoring system reasoning over purpose codes as though they were confirmed has made the same category error as one clearing a counterparty because the check digit validates. *Over-population as its own tell*: if full population becomes a recognised posture, a message that is **too** perfectly populated can itself become a pattern, in the way that unnaturally clean is its own kind of anomaly — and **the only defence is that the richness is true.**

And a structure controls the messages it originates, not the messages its counterparties send. A counterparty on an older system sends sparse messages and the structure inherits the scrutiny on the inbound side. That is not a reason to send sparse messages defensively; **it is a reason to prefer counterparties whose own messaging is current, because a bridge structure's resolvability is only as good as the weakest system that sends into it.**

8.3 The adoption path, and why it is longer than the engineering

The engineering here is small and the adoption path for the parts that touch the accounts is measured in years, and pretending otherwise is the fastest way to lose a technical audience.

Rules 1 to 15 need no external adoption at all, as Section 8.1 says. **Rules 16 to 18 are different in kind because they change how a balance is presented.** Excluding a balance from a live aggregate is a presentation decision; discharging it through a resolution entry is a derecognition decision. **Both sit inside an applicable financial reporting framework and neither is the preparer's to make unilaterally.** The sequence that works is: identify the exposure; obtain a view from the auditor on the presentation under the existing framework; implement whatever that view permits; and only then, with a documented case and a real example, approach a standard-setter about whether the general problem warrants a general answer. **The sequence that fails is to specify the designator first and seek acceptance afterwards, which is the sequence a technically minded author naturally reaches for.** A group that adopts it unilaterally and discloses it has disclosed a departure from established practice: **that is an ordering problem rather than a concealment one, and getting the order wrong turns a good control into a qualified opinion.**

The architecture of Section 7 has a third and longer path again. The missing ordering layer is a composition problem with known solutions; the missing identity layer is addressed by the capability layer, whose audit findings are fixed and re-validated **but whose hardened build has not been deployed at scale or reviewed by anyone outside the corpus that produced it.** Neither component has an independent implementation, an interoperability profile, or

a conformance suite maintained by anyone other than its author. **Those are the ordinary prerequisites for anything to become infrastructure, and the honest position is that none is met.**

The alternative is the failure mode this paper's own discipline exists to prevent. A submission presenting a validated substrate, an audited protocol, a taxonomy, a recording discipline and a naming rationale as though they were one finished system invites a reviewer to test the weakest and generalise. **A submission that grades each separately, prints the negative results, and states which parts are weeks away and which are years, invites the reviewer to evaluate the parts that are ready.**

8.4 What would falsify this paper

Five things, none tested here.

If the flag stream from rule 9 were dominated so completely by historical normalisation artefacts that no clean baseline could be established, the detector is a good idea that does not survive real data. The mitigation is a one-off backlog burn-down before alerting on new occurrences, and a deployment that skips it generates an unreadable queue on day one and switches the control off in the second week.

If canonical matching against a real, large designated-persons list produced a materially elevated false-positive rate, the screening argument becomes a cost–benefit question rather than a free improvement.

If a standards body or an audit firm concluded that the resolution designator is inconsistent with existing recognition and derecognition requirements, rules 16 to 18 fail and the withdrawn-code problem needs a different answer inside existing rules.

If an ordering layer capable of resolving the A1 conflicts at the required scale cost more in latency or complexity than the write-parallelism the factoring buys, the object-as-ledger choice is an optimisation that does not pay for itself. That is an empirical question about a specific workload with different answers for different workloads.

And if a survey of production critical machine fields found that out-of-repertoire code points essentially never occur in them, the operational case collapses to a theoretical vulnerability with no realised exposure. **That survey has not been done, here or — so far as this paper is aware — anywhere.**

8.5 Bounds and parallels carried from the corpus

Bounds. Curzi [2026] declines the inflated reading of its own borrowed physics and requires that a user be able to distinguish a theorem, an ordinary observer effect and the framework's own conjecture in three separate sentences on demand. **Section 9.6 discharges that requirement for every borrowing here**, and Section 3.4 discharges it for the two exclusions. Curzi [2026t] demotes the corpus's own moment equation at the scale-type gate, and Section 7.4 applies the identical verdict to the nine-form taxonomy for the identical reason.

Parallels, and each is structural rather than decorative. Curzi [2026] states the contraction theorem exactly — a contraction of modulus $q < 1$ on a complete space has exactly one fixed point, reached from any start, with the distance shrinking by at least q each pass — and **the parallel is the idempotence requirement of rule 3**: an operation applied repeatedly must reach and hold a fixed point, or it cannot be applied at more than one boundary. Curzi [2026]'s three-drawer honesty classification — four items of established mathematics, seven of established psychology and exactly three declared postulates — is the general form of the tier scheme of Section 3.2, and Curzi [2026]'s audit distribution over thirteen entries, with seven mapping, two splitting, one rendering and three rejected, is the same discipline applied to a different inventory: **an audit that rejects nothing has not audited.**

Curzi [2026w] supplies the evidential rule this paper leans on hardest: **two sources agreeing does not double the evidence when only one of them is evidence** — the reason a docu-

mented mechanism and a folk inference drawn from the same coincidence are kept apart in Section 5.7.

Three results in the corpus's biological arm are the same failure in a different medium. Curzi [2026x] shows that a field is intrinsically broadcast — 3.65e6 cells share one voxel at 1 MHz, so exclusive per-cell addressing would need 23.93 days at 70.65 bit/s, a per-cell rate of 1.93e-5 bit/s — which is **an addressing scheme in which many referents share one name**, the exact inverse of Section 5's problem and the same defect. Curzi [2026y] states the constraint that no addressing operation, at any signal strength or modality, can cause a receiver to express content its repertoire does not already contain — which is **the ceiling this paper's Section 8.2 states in its own terms: a representation control cannot produce information the record never carried.** And Curzi [2026] states the open problem formally: no established mechanism produces a response depending on a symbol string as an **order**, as distinct from permutation-invariant functionals of the waveform. **That is finding A2 in another medium** — a hash that does not commit to position cannot distinguish an ordering from a permutation of it, which is precisely what the reordering demonstration shows.

Two further parallels bear on the paper's honesty apparatus. Curzi [2026ad] reports a preregistration whose measured false-positive rate is 9.5 per cent against a nominal 5, and restores the nominal level by requiring cross-individual replication at 4 of 12; **the parallel is exact — a control's measured false-positive rate is the number that matters, and the 11.01 per cent precision of Section 5.8 is the same disclosure made about a screen instead of a test.** Curzi [2026] fixes a rate penalty at exactly 8 symbols per second from the ratio of two stipulated intervals; **the parallel is that an exact ratio of stipulated quantities is exact and carries no empirical weight**, which is the status of every percentage in this paper.

9. Verification and rigour audit

Each result is decomposed into postulates that can be checked one at a time. Each is atomic, each carries a status, and each chain names what it consumes. **If dropping a postulate leaves the result standing, it is not listed.**

9.1 Dimensional audit

symbol	quantity	dimension	unit	status
corpus size	records, messages, tokens, inputs	[#]	counts	stipulated (synthetic)
flagged, genuine	subsets of a corpus	[#]	counts	derived by enumeration
precision	genuine ÷ flagged	[1]	dimensionless fraction	derived; two counts of the same kind
ambiguity share	parsing false positives ÷ flagged	[1]	dimensionless fraction	derived
false-break rate	spurious breaks ÷ compared pairs	[1]	dimensionless fraction	derived
queue share	alerts of one class ÷ all alerts	[1]	dimensionless fraction	derived

symbol	quantity	dimension	unit	status
manipulation rate	share of corpus entries substituted	$[1]$	dimensionless fraction	stipulated
collision rate	share of sparse names colliding with a list stem	$[1]$	per message	stipulated
balance	ledger amount	$[1]$	**stipulated units, no currency attached**	stipulated
inflation	$(10{,}000 - 8{,}000)/8{,}000$	$[1]$	dimensionless fraction	derived from the stipulated ratio
digest	SHA-256 output	$[1]$	256 bits	defined
append cost	wall-clock time for a batch of local appends	$[T]$	seconds	**measured, machine-specific**
chains	independent per-object ledgers	$[\#]$	counts	stipulated
tier T1–T4	evidential label	—	**ordinal label, no unit**	this paper's own convention

Does it balance? Every headline figure in this paper is a ratio of two counts of the same kind of thing, and is therefore dimensionless and well formed: 150 genuine events over 1,363 flags, 1,213 parsing false positives over the same 1,363, 92 sparse alerts over 98 total alerts, 495 collapses over 495 manipulations. **No figure in this paper adds quantities of different kinds and none divides a count by a time.**

Two entries do not carry an operational unit and the audit says so rather than assigning one. The ledger amounts of Section 4.4 are in **stipulated units with no currency attached**; the 25.0 per cent is a property of the $8{,}000/2{,}000$ ratio and would be a different number under any other stipulation, which Section 4.7 already states. And **the nine-form backing taxonomy of Section 7.4 has no measurement basis at all for its non-financial forms**, so a coverage ratio computed across it is a fraction whose numerator has no defined unit. That is not a precision problem; it is the absence of a quantity.

The only dimensional quantity in the paper is the append time, and it is a wall-clock measurement on one machine in one process. Section 10 records that it carries no uncertainty and no repetition.

9.2 Significant figures and epistemic status

- 100.0%**,** 0.0% **and** 0.00% **are exact, not rounded small numbers.** They are properties of exhaustive enumerations over a stipulated corpus — 495 of 495, 500 of 500, 200 of 200, 28 of 28, 20,000 distinct identifiers from 20,000 distinct inputs — and no estimate is involved in any of them.
- 11.01% **and** 88.99% **are exact to the digits shown**, being $150/1363 = 0.110051\ldots$ and $1213/1363 = 0.889949\ldots$ on a stipulated corpus. They are **not** estimates of any screen's precision and carry no interval, because there is nothing to estimate.
- 93.9% **is** $92/98 = 0.938776\ldots$, exact for the reported alert counts and stipulated to one decimal place. Section 9.10 records that those counts are not reproducible from the source's own stated rates.
- 25.0% **is exact by stipulation** and is a property of one chosen ratio, as Section 4.7 states in place.

- **The append times are the weakest numbers in the paper.** 0.002, 0.003, 0.007 and 0.014 seconds are single-run wall-clock timings on one machine, quoted to three decimal places, with **no repetition, no dispersion, no confidence interval and no hardware specification.** Three decimal places on a single timing is more precision than the measurement supports; **only the linearity is a finding, and the absolute figures are printed to show the shape and for no other purpose.**

9.3 Limits and boundary behaviour

The false-break rate, over the whole manipulation range. Under byte equality the false-break rate equals the manipulation rate m exactly, for every m in $[0, 1]$; under canonical comparison it is 0 for every m. **The 40-to-0 reading is one point on a flat line and the degenerate endpoints are both examined**: at $m = 0$ nothing is manipulated and both regimes agree, and at $m = 1$ byte equality breaks every pair while the canonical form breaks none.

The precision figure, at its boundaries, and this is where the result is weakest. The naive screen's precision is $150/(150 + k)$ where k is the count of account-number records containing the digit string. At the stipulated $k = 1213$ it is 11.01 per cent; **at $k = 13$ — the incidental occurrences alone, with the national grammar's currency positions removed — it is 92.02 per cent; and at $k = 0$ it is 100 per cent and the finding vanishes entirely.** The 11.01 per cent is therefore a property of the stipulated corpus mix and **is not a claim about any real record stream**; what is invariant across the whole range is that the typed parse gives 100 per cent precision at every k, because the grammar never looks where the incidental digits are.

The queue inversion, as the migration completes. With the reported alert rates — a rich rate of $6/620 = 0.9677$ per cent and a sparse rate of $92/380 = 24.2105$ per cent — the sparse share of the queue falls monotonically to zero as the rich share $\rho \to 1$, and rises to 100 per cent as $\rho \to 0$. **The boundary worth printing is where the two populations contribute equally: that occurs at $\rho = 0.9616$, so sparse messages are still half the review queue when they are only 3.84 per cent of traffic.** The inversion is not a feature of a half-finished migration; it deepens as the migration completes.

The append measurement has no examined upper boundary. Linearity is measured at 50, 100, 200 and 400 chains and is asserted for no larger figure. Nothing here licenses extrapolation, and Section 7.1 makes no throughput claim of any kind.

9.4 Invariance

Idempotence is the invariance the specification actually requires, and rule 3 makes it normative: $\mathrm{canon}(\mathrm{canon}(x)) = \mathrm{canon}(x)$ for all 500 corpus values. The transformation group is repeated application of the operator, and the invariant is the canonical form. **Its failure mode is named and is not hypothetical:** a confusables-data or character-standard version change between two boundary applications breaks the invariance, which is why rule 6 makes both versions change-controlled configuration items.

The pair $(\mathrm{canon}, \mathrm{hash}_{\mathrm{bytes}})$ is invariant and variant by design, and that is the whole detector. The canonical form is invariant under every member of the substitution class; the byte digest is invariant under none of it. **A control storing only one of the two has destroyed either the matching or the evidence**, and Section 4.2 shows which.

What is not invariant, stated plainly. The precision figure is not invariant under the corpus mix (§9.3). The queue share is not invariant under the migration fraction (§9.3). The append times are not invariant under the machine. **None of the three is offered as a constant of anything**, and the invariant findings are the sign and the mechanism in each case: the naive screen's precision is below 100 per cent whenever the collision occurs at all, the queue share exceeds the traffic share whenever the sparse alert rate exceeds the rich one, and the append cost is linear rather than quadratic in the chain count.

9.5 Measurement-theoretic status

The tier scheme is ordinal and no arithmetic is performed on it. T1 through T4 are this paper's own labelling convention with no external standing. They are not averaged, not summed, not counted into a score, and an ordinal scale licenses none of those operations [Stevens, 1946, Michell, 1997]. The scheme's only function is to keep four kinds of statement apart on the page.

A code is a dated assertion and not a fact, and this is a validity claim rather than an arithmetic one. The entity identifier's renewal cycle (§5.5) is the standard's own admission of the point: the code is permanent and the assertion behind it perishes. A check digit has genuine validity for one construct — transcription error in the string — and none whatever for a different construct, the existence and ownership of the referent, and treating the first as evidence for the second is the convergent-versus-discriminant confusion in its plainest form [Campbell and Fiske, 1959, Cronbach and Meehl, 1955]. **A control that clears a counterparty because the check digit validates has measured the string and reported on the world.**

The nine-form backing taxonomy fails the scale-type gate outright. Its non-financial forms are not instrumented on any common scale, let alone a ratio scale with a shared true zero, so the magnitude of any coverage ratio computed across them is withdrawn as a measured quantity — Curzi [2026t] reaches exactly that verdict on the corpus's own moment equation, and §7.4 applies it here. The formal remedy is conjoint measurement or a Rasch demonstration [Luce and Tukey, 1964, Krantz et al., 1971, Rasch, 1960], and **none was run.** Michell's objection — that a quantitative structure is assumed rather than tested — stands unanswered [Michell, 1997, 1999], and this is the corpus's scalar-reduction habit arriving in an engineering artefact, conceded where the ladder is built [Curzi, 2026] rather than re-argued here.

Which effective count. Where this paper refers to an effective count it means $1/\mathrm{HHI}$, the Hill number of order two [Hill, 1973, Jost, 2006, Adelman, 1969], as Curzi [2026] and Curzi [2026] use it. **The corpus's affine construction $1 + (N-1)u$ is not a Hill number of any order and is not used anywhere in this paper.**

No distributional or scaling claim is made anywhere in this paper, so the goodness-of-fit and likelihood-ratio procedure that a heavy-tail or power-law claim would require [Clauset et al., 2009] is not applicable and was not run. Every figure is a ratio of counts on a stipulated corpus.

The reflexive loop is not invoked here and no social-influence parameter is estimated, so the reflection problem [Manski, 1993] does not bear on any result. The one reflexive hazard the paper does carry is the index-as-target problem of §8.1 — a concentration measure adopted as a target stops measuring what it stood for [Goodhart, 1975, Campbell, 1979] — and it is named where it arises rather than resolved.

9.6 Analogy discipline: every borrowing, graded

Curzi [2026] requires that a user be able to distinguish a theorem, an ordinary effect and the framework's own conjecture in three separate sentences on demand. **This section discharges that requirement for every borrowing in this paper.** Three categories are used: **(1) a derivation**, where the system genuinely satisfies the cited result's premises; **(2) a shared mathematical form**, where the same equation governs both with no transfer of authority; **(3) an analogy**, legitimate as motivation and never as evidence.

borrowing	used for	category
Byzantine agreement [Lamport et al., 1982]	why an ordering layer is not optional (A1)	**1 — derivation.** The substrate is literally a set of mutually distrustful parties needing agreement

borrowing	used for	category
Consensus impossibility [Fischer et al., 1985]	what an ordering layer can and cannot achieve	**1 — derivation**
The availability–consistency trade [Gilbert and Lynch, 2002]	the same	**1 — derivation**
Identity forgery by mass creation [Douceur, 2002]	A4, why the party count must itself be audited	**1 — derivation**
Quantum factoring [Shor, 1997] and search [Grover, 1996]	the post-quantum posture of a hash-and-tag core	**1 — derivation.** The primitives in use are exactly the ones the results address
Open design [Kerckhoffs, 1883, Saltzer and Schroeder, 1975]	the disclosure test of §3.3	**1 — derivation** as a design principle
Transparency as remedy [Brandeis, 1914]	the same test, in its policy register	**3 — analogy.** A legal-policy argument, not a security result
Contraction to a fixed point [Curzi, 2026]	the idempotence requirement of rule 3	**2 — shared form.** An operator with a fixed point; no metric-space claim is made about strings
Requisite variety [Ashby, 1958]	why a one-category screen cannot separate two referents	**3 — analogy**, invoked for shape, as §5.8 states in place
Coupled-network fragility [Buldyrev et al., 2010]	how a defect travels through composed identifiers	**3 — analogy**, invoked for shape, as §5.6 states in place
Queueing under load [Kingman, 1962, Little, 1961, Erlang, 1909, Hopp and Spearman, 2011]	why message richness is a capacity parameter	**2 — shared form.** A review function is a service system with arrivals and finite capacity; **no queue was measured**
Index adopted as target [Goodhart, 1975, Campbell, 1979]	§8.1's warning about concentration measures	**3 — analogy.** The cause here is mechanical, not behavioural, and §8.1 says so
Infrastructure as an installed base [Star and Ruhleder, 1996, Star and Griesemer, 1989]	why a registry binding decays unless maintained	**3 — analogy**
Legibility to an observing system [Scott, 1998]	§6.3's resolvability asymmetry	**3 — analogy**

The three sentences, stated as Curzi [2026] requires them. *The theorem:* no protocol among mutually distrustful parties achieves agreement without an ordering layer, and the substrate of §7 has none, so it does not prevent double-spending — this is a derivation and

the demonstration in §7.2 is its instance. *The ordinary effect:* a screen that reads bytes instead of typed fields cannot separate two registries' meanings, which is an ordinary consequence of parsing under the wrong grammar and needs no cybernetic theory to see. *The conjecture:* that full population is a costly signal a lawful structure can afford and an unlawful one cannot (§6.3) is **this work's own proposal**, it is not derived from anything cited, and no evidence in this paper establishes it.

9.7 The postulates

	postulate	status
P1	Compatibility normalisation folds full-width and compatibility variants; canonical composition does not.	CITED to the character standard — **and that standard is not in the bibliography** (Appendix B)
P2	Two strings are confusable when their skeletons are equal under the published confusables data.	Same status as P1
P3	SHA-256 is collision- and preimage-resistant at the strengths assumed.	ESTABLISHED, uncited (Appendix B)
P4	Each stipulated corpus is as described: 500 clean values, 495 manipulated, 2,000 accounts of which 1,200 carry the digits in currency positions and 13 incidentally, 2,000 currency fields of which 150 are genuine, 1,000 messages at 62 per cent rich.	DEFINITION (stipulated design parameters)
P5	The rich alert rate is the genuine list-membership rate and the sparse alert rate adds a name-collision rate.	MODELLING CHOICE; see §9.10
P6	A registry, not a string, is what binds a code to a referent.	POSTULATED — the paper's own framing, argued in §5.1
P7	Unresolvability, not suspicion, is what a screen escalates on.	POSTULATED — no production screening system was examined
P8	The nine forms of capital admit a coverage ratio.	**REFUTED as stated** — no measurement basis exists for the non-financial forms (§7.4)

Dropped as unnecessary. No postulate about *which* character standard version, or about the identity of any named registry, is needed for the canonical-form results: they follow from the composition of three operators and the stipulated corpus. No postulate about intent is needed anywhere, and §5.7 turns on that fact.

9.8 The chain that establishes the integrity layer (H1, H2)

step	consumes	produces	status
C1	P1, P2	canon $=$ restrict $\circ$ fold $\circ$ NFKC is well defined and order-fixed	DEFINITION
C2	C1, P4	495 / 495 manipulated entries collapse onto the clean form	ENUMERATION (495)
C3	C1, P4	$\text{canon}(\text{canon}(x)) = \text{canon}(x)$ for all 500	ENUMERATION (500)
C4	C2, C3	995 raw byte strings resolve to 500 canonical forms, the true count	ARITHMETIC (500 + 495)
C5	C4	the false-break rate under canonical comparison is 0 for every manipulation rate	ARITHMETIC
C6	P3	the digest over raw octets differs for a clean value and its twin; over the canonical form it does not	ENUMERATION (one pair)
C7	C5, C6	equal canonical form with differing byte-hash is a decidable flag — **the detector**	DEFINITION

H1 and H2 are supported, and the support is enumeration on a stipulated corpus, not evidence about any real record stream.

9.9 The chain that establishes the typed-parse collapse (H3)

step	consumes	produces	status
D1	P4	account records containing the digit string: $1200 + 13 = 1213$	ARITHMETIC
D2	D1, P4	naive flags $= 1213 + 150 = 1363$	ARITHMETIC
D3	D2	precision $= 150/1363 = 11.01\%$	ARITHMETIC
D4	D2	ambiguity share $= 1213/1363 = 88.99\%$	ARITHMETIC

step	consumes	produces	status
D5	P6	the grammar owning the account field never reads the currency positions as a currency code	DEFINITION
D6	D5	typed parse flags exactly the 150; precision 100%, ambiguity share 0.00%	ARITHMETIC

H3 is supported. D3 and D4 are exact on the stipulated mix and vary with it (§9.3); D6 is invariant across every mix, which is the part that is a finding.

9.10 The chain that establishes the queue inversion (H4), and a discrepancy in the source's stated rates

step	consumes	produces	status
E1	P4	1,000 messages split $620/380$ at a 62 per cent rich share	ARITHMETIC
E2	P5, reported counts	rich alerts 6, sparse alerts 92, queue 98	REPORTED (see below)
E3	E2	sparse queue share $= 92/98 = 93.88\%$ against a traffic share of 38.0%	ARITHMETIC
E4	E2	equal contribution at a rich share of 0.9616: sparse messages are half the queue at 3.84% of traffic	ARITHMETIC

The discrepancy, printed rather than harmonised. The source manuscript's own reproduction note states the stipulated parameters as *a 62 per cent rich share, a small list-membership rate, and a 22 per cent sparse-name collision rate.* **Those rates do not produce the reported counts.** The six rich alerts imply a list-membership rate of $6/620 = 0.968$ per cent; applied to the 380 sparse messages alongside a 22 per cent collision rate, the expected sparse alert count is $380\,[\,1-(1-0.22)(1-0.00968)\,] = 86.5$, and the reported figure is **92**, which corresponds to an effective sparse alert rate of 24.21 per cent rather than the stated 22 per cent plus list membership.

Two readings are available and both are given. Under a single seeded Bernoulli draw the gap is about 0.7 standard deviations ($\sigma \approx 8.1$ at $n = 380$, $p = 0.22$) and is unremarkable sampling variation; **but the source presents the figures as though they followed from the rates, and as arithmetic they do not.** At the stated rates the queue would be 86 sparse and 6 rich, a sparse share of 93.5 per cent rather than 93.9. **The finding is unaffected in sign, magnitude or mechanism, and the discrepancy is recorded because a reader recomputing from the stated rates will not reproduce the printed counts.**

9.11 The chains for the structured field and the substrate (H5, H6), and cross-checks

H5. The name element of the structured standard is a text field accepting the whole repertoire [P1]; a homoglyph twin placed in it renders identically, has a differing byte digest (`c413622fa3bea888` against `7307655b3eee8d17`), and has an identical canonical form. **The standard defines a place for the name and no canonical form for it, so the syntax passes the twin and the integrity layer catches it.** ENUMERATION on one constructed pair.

H6. 200 of 200 chains verify and 200 of 200 single-block mutations are detected, which follows from each block's hash covering its transaction and each chain linking by predecessor hash [ENUMERATION, 200]. 28 of 28 cross-object links are recorded in both ledgers [ENUMERATION, 28]. Append cost is linear at 50, 100, 200 and 400 chains [MEASURED, machine-specific]. **And a fork transferring the same balance twice yields two internally valid chains that the structure does not reject** [DEMONSTRATED, one instance] — which is the derivation of §9.6 arriving as a concrete counterexample. **H6 is supported in both halves, and the half that matters is the negative one.**

Cross-checks performed. The canonical-form and false-break figures were recomputed for this paper and agree with Curzi [2026]'s independent measurement of the same quantities. The homoglyph-in-structured-field result is the silent-duplicate condition of Curzi [2026] recomputed on a name fragment and is consistent with it. Content addressing over 20,000 distinct inputs yields 20,000 distinct identifiers with zero collisions, which confirms an implementation property rather than establishing anything about the hash. **The audit findings of §§7.2 and 7.3 were reproduced on a harness rebuilt from the source descriptions; the original code was not obtained and was not rerun**, and Appendix A records what must be confirmed before any of them is repeated in the present tense.

10. Limitations

Every figure in this paper is a deterministic computation over a synthetic corpus with stipulated parameters, and none is an estimate of prevalence in any real system. No production record stream, payment network, screening function, institution or ledger was examined. **No empirical dataset was collected.** The stipulations are printed at each point of use precisely so that a reader can show that a different stipulation gives a different figure, and several would.

The rates would need prospective registration and out-of-sample testing before any of the measured improvements could be claimed for production traffic. The manipulation rate, the corpus mix behind the 11.01 per cent precision, the list-membership rate and the sparse-name collision rate are design parameters chosen to demonstrate a mechanism. Claiming an improvement for real traffic on parameters chosen after seeing the demonstration is the garden of forking paths in its plainest form [Gelman and Loken, 2013, Simmons et al., 2011]; the remedy is to declare the rates and the endpoints before collecting the data [Nosek et al., 2018], and **that has not been done here.** The corpus's own precedent is instructive rather than reassuring: Curzi [2026ad] reports a preregistration whose measured false-positive rate was 9.5 per cent against a nominal 5.

No standards document, court decision or jurisdiction-specific instrument is cited, because `bib/external.bib` contains none of them. One structured-messaging standard is in the bibliography and is cited where it is relied on [International Organization for Standardization, 2013]. Every other designation in Section 5 — the currency, country, securities, institution, account, classification, legal-form and entity-identifier standards, their edition years, their check-digit algorithms, their registration authorities and their code allocations — is **described from the source manuscript and left uncited**, and is itemised in Appendix B. The same applies to the historical currency reallocation of §5.7, to the court decisions rejecting the double-ledger reading, to the post-quantum signature standards of §7.3, and to the capability layer's own literature. **This is a real gap and it is not closed here.**

The identifier-collapse trade is real and is not priced. Canonical folding collapses the identifier space and therefore creates collisions between entities whose names genuinely differ only by folding characters. For narrow critical machine fields this is a non-issue by construction; **for a name estate it is a genuine trade whose false-positive cost is not measured here and is not zero**, and the zero-false-positive figure Curzi [2026] reports is measured against clean ASCII critical fields and does not transfer.

The append measurements carry no uncertainty and were not repeated. Four single-run wall-clock timings on one unspecified machine in one process, quoted to three decimal places. **Linearity is the only claim they support**, no hardware is stated, no dispersion is reported, and nothing licenses extrapolation past 400 chains.

The artefacts of Section 7 are the author's own and the audit is the author's own. That is a conflict of interest, it is declared, and the only method available to a single author is to report the negative findings at least as prominently as the positive ones. **A reader who discounts the positive findings accordingly is reading correctly.** Neither component has an independent implementation, an interoperability profile, or a conformance suite maintained by anyone outside the corpus that produced it.

A single registry operator is a common-mode failure for every identifier composed from it, and registry capture bounds the whole approach. Section 5.6 shows a defect travelling through the composed identifiers; the same structure means that the country registry's maintenance agency is a single point on which four other standards depend, and that the entity-identifier system's governance is a single point beneath every deterministic clearance in Section 6. **A control whose correctness rests on a registry inherits that registry's governance risk in full**, and the network literature on common-mode dependence in layered systems [Buldyrev et al., 2010, Albert et al., 2000, Cohen et al., 2000, Doyle et al., 2005] is the general statement of the exposure. **Nothing in this paper measures it, mitigates it, or proposes a fallback for it.** Tight coupling between layers is also what turns a local defect into a system accident [Perrow, 1984, Leveson, 2011], and adding a redundant registry is not automatically a remedy [Sagan, 1993].

An unactioned flag queue is a worse outcome than no queue, and rule 10 is the only thing standing against it. A control producing exceptions nobody dispositions is evidence that something was detected and ignored, which is the organisational failure mode documented as the routine acceptance of signals that ought to have stopped a process [Vaughan, 1996, Weick and Sutcliffe, 2007]. **Rule 10 is normative for that reason and it is the rule most likely to be treated as advisory.**

Nothing in this paper measures truthfulness. Canonicalisation defeats encoding-level evasion; typed parsing defeats code-level ambiguity; hash chains defeat silent edits; content-addressed capabilities defeat impersonation; full population defeats the identity gate. **Every one of them is silent about a correctly encoded, correctly typed, correctly signed false assertion**, and a tamper-evident ledger containing a false entry will preserve it faithfully forever.

The survey that would establish the operational case has not been done. Whether out-of-repertoire code points occur at material rates in production critical machine fields is an empirical question, it is the question on which the whole exposure turns, and **so far as this paper is aware it has been answered nowhere.**

10.1 Rigour checks not met, listed by name

1. **No as-of date and no source for any figure, because nothing was measured.** Every result derives from a stipulated corpus.
2. **No uncertainty is reported anywhere.** For the ratios this is correct — they are exact on their corpora — but for the four append times it is a genuine omission (§9.2).
3. **Model comparison was not performed.** No alternative canonicalisation, no alternative screen and no alternative ledger factoring was implemented and compared; no out-of-sample test was run; no information criterion is reported, because nothing was fitted.

4. **Conjoint measurement was not tested** for the nine-form taxonomy, and Michell's objection stands [Michell, 1997, Luce and Tukey, 1964].
5. **The specification has no adopted standard status**, no auditor acceptance for any accounting treatment described, and no legal force anywhere.
6. **The resolution designator has not been put to a standard-setter or an audit firm**, and §8.3 states that the sequence matters: acceptance before adoption, not after.
7. **The 22 per cent collision rate does not reproduce the reported alert counts** (§9.10).
8. **No independent party has reviewed the hardened build** of the capability layer, and it has not been deployed at scale.

11. Conclusion

The integrity layer, in quotable form. The canonical form $\text{canon}(x) = \text{restrict_ASCII}(\text{confusable_fold}(\text{NFKC}(x)))$, paired with a SHA-256 witness computed over the **raw octets** and kept on a plane separate from the display form, collapses **495 of 495** manipulated entries onto their clean canonical form, is idempotent across all **500** corpus values, resolves **995** distinct raw byte strings to the true **500** identifiers, and moves the reconciliation false-break rate from **40.0 per cent to 0.0 per cent** at a 40 per cent manipulation rate — one point on a line that is flat at zero across the whole manipulation range. **The detector is one comparison: identical canonical form with differing byte-hash.** Hashing the canonical form instead gives a clean value and its homoglyph twin the identical digest and destroys, silently, the only evidence in the system that a substitution occurred.

The same defect one level up, and the number that matters. Aliasing a withdrawn currency code to a live one records a synthetic **8,000** balance as **10,000** — a phantom **25.0 per cent** inflation with no journal entry and nothing forged. Where two independently governed registries mint the same three digits with different meanings, a byte-level screen flags **1,363** records of which **150** are genuine: a precision of **11.01 per cent** and a parsing-ambiguity share of **88.99 per cent**. **A control right one time in nine is not a control; it is an alert generator, and it will be switched off in the second week.** One typed-field parse collapses that share to **0.00 per cent**, and it does so without anyone determining or alleging anybody's intent.

The resolvability asymmetry. In a synthetic day of 1,000 messages at 62 per cent rich traffic, sparse legacy messages are **38.0 per cent of volume and 93.9 per cent of the review queue**, and they remain **half the queue at 3.84 per cent of traffic**. **The message carrying less data is not more private; it is less resolvable, and unresolvability is what costs.** The corollary is the whole of the argument's honesty: full population is available only to a structure with nothing to lose by resolving, which is what makes it a costly signal rather than a policy — **and a structure populating fields it cannot truthfully fill has re-entered the concealment game from the other side.**

And a structured standard does not close it. Inside the richest, most authoritative name element the standard defines, a homoglyph twin is byte-distinct, renders identically, and is canonically identical: **the syntax fixes the shape of the data and not the identity of the values**, and only the canonical form and the byte witness separate them.

The negative results, stated as prominently as the positive ones. The per-object ledger factoring verifies **200 of 200** chains, detects **200 of 200** silent edits and records **28 of 28** cross-object links, with local append cost rising from **0.002 s to 0.014 s** across 50 to 400 chains — **a local, single-process measurement that is not distributed throughput.** It supplies **no consensus, no total order, no double-spend prevention and no Sybil resistance**, and a working double-spend against it is demonstrated: two forks of one token both verify and the structure raises no objection. Its block hash under-binds, and the fix is given. Four breaks in the capability layer — an unverified proof signature, an inconsistently computed self-verification digest, broken hash algorithms on offer, and a replay surface at the freshness root — are reported with their fixes and re-validation. The claim in the author's earlier materials that the relevant cryptography is ready for immediate implementation and unbreakable **is refuted here, in the author's own**

paper. The nine-form backing taxonomy is a classification and not a valuation method: **no measurement basis for its non-financial forms is specified in the source corpus or here**, and any coverage ratio computed across it is computed over an unknown. **The naming rationale and the numerological layer confer no property on any structure described and are excluded from every technical claim.**

What the paper adds to the corpus's instrument. Concentration is measured over a partition of exposures, and the partition is built by resolving identities. **An identity-resolution defect moves a concentration measurement without anything real changing, and it moves it in the reassuring direction.** The effective count of Curzi [2026] — the reciprocal of the Herfindahl index, the Hill number of order two — is only as good as the canonicalisation underneath it, and this paper specifies that layer, tests it, and states what it does not do.

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every result in this paper is computed over a **SYNTHETIC** corpus with stipulated parameters, labelled as such at each point of use, under fixed seeds. The corpora are: 500 clean identifier values and 495 manipulated twins; 2,000 synthetic domestic account numbers and 2,000 synthetic numeric currency fields; a fourteen-entry synthetic designated-persons list and one constructed target entity that is not on it; a synthetic day of 1,000 payment messages; a three-account synthetic ledger extract in stipulated units with no currency attached; 200 synthetic token chains over 30 transactions each; and 20,000 distinct synthetic inputs for content addressing. **No real payment, party, account, institution, list entry, currency balance or production record stream appears anywhere in this paper, and no allegation is made about any of them.**

All figures were produced by the scripts named in their captions (`figs/src/CCP-065_fig01.py` through `figs/src/CCP-065_fig05.py`), which import the shared figure style module and read their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

The measurements in Section 7 were recomputed for this paper on a harness rebuilt from the source descriptions; the original code was not obtained and was not rerun. Each finding describes the implementation *as audited*, and Appendix A records what must be confirmed before any finding is repeated in the present tense.

Standards documents, court decisions, historical code allocations and the post-quantum signature standards are described from the source manuscript and are not cited, because `bib/external.bib` contains none of them. Appendix B itemises every such claim with what specifically requires verification.

Funding

This research received no external funding.

Conflicts of interest

The ledger substrate and the capability layer described in Section 7 are the author's own systems, and the adversarial audit reported there is also the author's. This is declared as a conflict of interest. The method adopted in response is stated in Section 3.4 and applied throughout Section 7: report the negative findings at least as prominently as the positive ones, and report by name the ones that refute the author's own earlier claims. A reader who discounts the positive findings accordingly is reading correctly. The author declares no other conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data were processed. Every party, account and payment described is synthetic and constructed.

No operational content

This paper contains no operational content. It describes a defect class at the representation layer and specifies a control against it. The substitution demonstrations use a single constructed homoglyph pair to show that a detector fires; no evasion technique, list, tool or method is supplied, and nothing here would assist anyone in constructing a substitution or in moving value any regime prohibits. **The double-ledger reading of the digit collision in §5.7 is refused outright, is graded T4, and is recorded in the source as having been rejected by the courts that considered it.**

Legal, tax, accounting and investment advice

This paper is analytical and is not legal, tax, accounting or investment advice. The eighteen normative rules are a proposal with **no adopted standard status**; they replace no obligation and carry no legal force. **The resolution designator is an accounting proposal requiring a standards-body process and auditor acceptance, and it has neither: until it has both, an entity cannot use it to justify a treatment in a set of accounts.** No recommendation is made in respect of any particular person, institution, product or jurisdiction, and no allegation is made about any of them. Anyone contemplating a presentation or derecognition treatment on the basis of Sections 4.4, 5.9 or 8.3 must obtain a view from their auditor under the applicable financial reporting framework **before** implementation, in the sequence Section 8.3 states.

Appendix A. The status of the artefacts, and what must be confirmed before restatement

Every item in this appendix concerns the systems described in Section 7. **Each finding describes the implementation as audited at the time of the source audit. None may be repeated in the present tense without confirming the current state of the code.**

1. **A1, the double-spend.** The demonstration that two forks of one token both verify was reproduced for this paper on a rebuilt harness. Confirm whether an ordering layer has since been added; if one has, the finding is historical and the substrate's verdict changes.
2. **A2, the under-binding block hash.** Confirm whether the digest is now computed over the chain identifier, predecessor hash, index, timestamp and transaction together. The fix is specified in §7.2 and its effect — reordered blocks rejected — was reproduced here.
3. **A4, Sybil resistance.** Confirm whether any cost-to-identity or verified-identity layer has been added. Absent one, the finding stands as stated.
4. **R1 to R4, the capability layer.** All four were reported by the author as fixed in a hardened module and re-validated. **Confirm that the hardened module is the deployed one**, and note that the hardened build has not been reviewed by any party outside the corpus that produced it and has not been deployed at scale.

5. **The post-quantum posture.** The hash-and-tag core is post-quantum-resistant at digests of 256 bits or more; the signature the protocol requires must be a standardised post-quantum scheme. Confirm which scheme, if any, has been adopted. The earlier marketing claim of readiness and unbreakability is refuted in §7.3 and must not be restated.
6. **The eighteen normative rules.** No adopted standard status, no conformance suite maintained by any party other than the author, no interoperability profile, and no independent implementation. Confirm before describing any of them as a standard.
7. **The resolution designator (rules 16 to 18).** No standards-body process has been initiated and no auditor acceptance obtained. **An entity cannot use it to justify a treatment in a set of accounts**, and §8.3 states the sequence that would change that.
8. **The nine-form backing taxonomy.** Recorded as declared. No measurement basis for the non-financial forms exists in the source corpus or here; confirm before any coverage ratio is computed or quoted.
9. **The append timings.** Machine-specific single-run measurements. Confirm the hardware and repeat with dispersion reported before quoting any absolute figure.
10. **The naming rationale and the numerological layer.** Assessed and excluded by the author's own security audit and excluded again here. They confer no property on any structure described.

Appendix B. Claims requiring verification against primary sources

The following are described in this paper from the source manuscript and **are not cited to the bibliography, which contains no standards documents, no court decisions and no jurisdiction-specific material.** One structured-messaging standard is in the bibliography and is cited where relied on [International Organization for Standardization, 2013]. Each item below requires verification against the governing document before any reliance.

1. **The character standard and its annexes** — the compatibility normalisation form, the confusables data file, the skeleton operation, the graded restriction levels, and the definition of default-ignorable code points. Confirm the current version of each, and note that rule 6 makes the version a change-controlled configuration item precisely because it moves.
2. **The digest standard.** SHA-256 is used throughout. Confirm the governing publication.
3. **The currency registry** — the three-letter and three-digit code structure, the separate live and withdrawn tables, the maintenance agency, and the withdrawal and reuse policy. Confirm the edition in force.
4. **The country registry** — the two-letter, three-letter and numeric code structure, the withdrawal lifecycle, the transitional reservation period, and the interval after which a withdrawn two-letter code may be reassigned. **No specific historical reallocation is named in this paper and none should be added without verification.**
5. **The securities identifier, the institution identifier and the account identifier** — the character counts, the composition rules, the check-digit or check-character algorithms, the national numbering agency prefixes, and the nationally defined internal account structure. Confirm each separately; the account standard's deferral to national grammars is the hinge of §5.7 and must be confirmed in the governing text.
6. **The instrument-classification standard and the legal-forms standard** — the code structures and the registries behind them.
7. **The entity-identifier standard** — the twenty-character structure, the prefix allocation, the check algorithm, the local operating unit model, and **the renewal cycle and lapsed-registration status**, which §5.5 relies on directly.
8. **The historical currency reallocation of §5.7.** The exact allocation of the withdrawn numeric code between two candidate predecessor currencies is contested in the source's

own verification register and **is not asserted here.** The mechanism does not depend on which it was; the naming does.

9. **The national account-number grammar of §5.7**, in which the same three digits designate the live national currency inside the account number. Confirm the grammar and the position of the currency field.
10. **The court decisions rejecting the double-ledger reading.** Described from the source and cited nowhere. Confirm the deciding courts, the dates and the holdings before any restatement. **This paper makes the refused reading in no form.**
11. **The legacy cross-border message format** — the field limits, the character-set restriction, the truncation behaviour, and the absence of a dedicated entity-identifier element. Confirm against the format specification.
12. **The structured-messaging standard's element names and data dictionary.** The standard itself is cited [International Organization for Standardization, 2013]; the specific element names, the external purpose-code list and the migration schedule are described from the source and require confirmation.
13. **The capability layer's literature** — the content-derived identifier scheme, the locator rendering, the capability-contract authorisation model, and the four-message challenge–response exchange with per-session derived keys. **None of the primary sources is in the bibliography**; the protection principles are cited [Saltzer and Schroeder, 1975] and nothing else is.
14. **The post-quantum signature standards** described in the source as published in 2024. **This paper names none of them by number.** Confirm the designations, the publication dates and the current status before adopting any of them.
15. **The receipt-as-transaction and shared-record antecedents** for the triple-ledger construction. The accounting-history literature that would supply them [de Roover, 1963, Littleton and Yamey, 1956, Yamey, 1964, Ferguson, 1998] is in the bibliography and is cited here for the general lineage; **the specific antecedent claims in the source corpus are not established by those works and require their own sources.**

All numerical results in this paper were computed from the stated stipulated inputs at full precision and rounded only for printing.

Chapter 66

Channel Separation as a Structural Property: In-Band Impersonation, Address Anonymity Sets, and the Legacy-Migration Asymptote

Abstract. We state and test a structural principle: a system carrying both control and data is *in-band* when a data value can be promoted to a control instruction, and *out-of-band* when the two planes are disjoint by construction. A synthetic demonstration over $100{,}000$ carried values gives an in-band ambiguity rate of 0.383 per cent (383 values indistinguishable from control) against 0.000 per cent out-of-band, and the same in-band flaw with the same out-of-band fix is tabulated across telephony, ledgers, messaging and settlement. The disclosure result is the load-bearing one and is stated as a hard constraint: channel separation is strengthened rather than weakened by full disclosure, because its effectiveness is entirely independent of whether anyone knows it is in place; publishing that the integrity witness is out-of-band, that settlement separates command from value, and that messaging separates identity from payload is in each case an assurance and never an opening. That is Kerckhoffs's principle of 1883, and citing it converts a rule that reads as an ethical preference into a recognised engineering standard. The stated caveat is that a separated control plane must still be access-controlled. At the addressing layer, an identity worth 18.932 bits entering a constructed carrier is observed as 8.000 bits from an address alone; logging address, port and timestamp reduces the anonymity set from $1{,}953.1$ to one while session state is live, at a cost of 3.6 TB of logs per day, and expiry restores it to $1{,}453.2$. Constructed legacy models give an 11 per cent attrition maintainer deficit against a flat requirement of 100, settling at 54.5, and a migration that leaves 3.29 per cent of a 40-million-line core after twenty years with a rational stopping point near 2032 at 37.6 per cent still buried.

Keywords: control plane; data plane; in-band signalling; anonymity set; carrier-grade translation; over-attribution; maintainer attrition; migration asymptote

Series: D. Structure and Resilience **Paper:** CCP-066

1. Introduction

A system that carries both control information and data has a structural property that decides an entire class of its failures: whether a value in the data plane can be promoted to an instruction in the control plane. This paper states that property precisely, tests it, derives from it a disclosure constraint that is stated as a hard design premise rather than as an ethical preference, and then follows the same structure down two further layers — the addressing substrate and the software substrate — where the identical shape recurs.

The paper's three contributions are separable.

The channel-separation principle, stated and demonstrated. A system is *in-band* when control and data share a channel such that a data value can be interpreted as a control instruction, and *out-of-band* when the two are disjoint by construction. In-band architectures admit an impersonation class that no care taken within the shared channel can close; out-of-band architectures close that class structurally. A synthetic demonstration over 100,000 carried values gives an in-band ambiguity rate of 0.383 per cent against 0.000 per cent out-of-band, and the same in-band flaw with the same out-of-band fix is tabulated across telephony, ledgers, messaging and settlement. The historical case is documented and closed: early international telephone signalling carried its control tones in the voice band, which made control reachable from the data path, and the move to common-channel signalling in the 1980s removed the condition the attack required rather than defending against the attack. **This paper contains no operational content of any kind** — no tone, sequence, code or method — and the historical material is used only to establish that in-band control was structurally reachable.

The disclosure constraint, and why it is engineering rather than ethics. Channel separation is *strengthened* by full disclosure, because its effectiveness is entirely independent of whether anyone knows it is in place. That is not a novel claim and it is not this work's to make: it is Kerckhoffs [1883]'s principle, stated in 1883, that a system must remain secure when everything about it except the key is public, and it is restated as *open design* among the protection principles of Saltzer and Schroeder [1975]. **Citing it converts a rule that reads as an ethical preference into a recognised engineering standard with a hundred and forty years behind it**, and it is the single most load-bearing external citation in this paper. The regulatory environment has since made the same conclusion unavoidable from a different direction: automatic exchange of account information [Organisation for Economic Co-operation and Development, 2017], the anti-money-laundering recommendations [Financial Action Task Force, 2012] and the beneficial-ownership and substance work [Organisation for Economic Co-operation and Development, 2015] have removed non-disclosure from the set of available design parameters.

The two substrate results. At the addressing layer, an identity worth 18.932 bits entering a constructed carrier is observed as 8.000 bits from an address alone; the anonymity set is 1,953.1 subscribers with an address-only log, exactly 1 while the translator's session state is live, and 1,453.2 after that state expires — at a retention cost of 3.6 TB of translation log per day for a single metropolitan translator. At the software layer, a constructed maintainer model with 11 per cent attrition against 6 hires a year lands at a steady state of 54.5 competent maintainers against a flat requirement of 100, and a constructed migration model leaves 3.29 per cent of a 40-million-line core after twenty years with a rational stopping point near 2032 at 37.6 per cent still buried.

Why these three sit in one paper. They are one mechanism in three media. In each case a representational boundary destroys information that a downstream consumer then trusts because the boundary is invisible to it. That is the same defect Curzi [2026] measures at the character layer, where typed-field parsing collapses the withdrawn-code parsing-ambiguity false-positive share from 88.99% to 0.00% against a naive detection precision of 11.01% — a control right one time in nine — and it is the same defect Curzi [2026] measures at the ledger layer, where a five-house network reads 3.819710 effective jurisdictions naively but 1.000000 effective blocs once a single coordination perimeter is recognised, giving $h = -0.500000$. **The concentration instrument of Curzi [2026] — HHI as the sum of squared shares and the effective**

count as its reciprocal, with a naive 3.41 collapsing to exactly 1.00 once blocs are recognised — is the same error at the jurisdictional layer, and Section 9.1 shows that the substrate exposure this paper measures is invisible to that instrument by construction.

Two attractive claims are refused in print, because a paper that lets a reader walk out believing either has done harm. **The three addressing namespaces do not share an origin or a common system**; they share a structural pattern of nested delegation, and the pattern is the finding while the shared-system claim is the overclaim. **Long-haul and metropolitan internet transport is fully optical and does not depend on copper**; what persists is the last-mile access plant in a large installed base and, more durably, the *logical* topology of the exchange hierarchy. **The durable dependency is logical topology, not electrical continuity.**

2. Hypotheses

Six hypotheses, each stated so that it could be wrong, each with its falsifier. Two statements that appear in the paper and are *not* falsifiable are labelled at the end. Platt [1964]'s strong inference is the standard applied: a hypothesis earns its place by naming the observation that would kill it.

H1 — the impersonation hypothesis. An in-band architecture admits a class of impersonation that no care within the shared channel closes; an out-of-band architecture closes that class by construction. *Predicted:* over a carried population of data values, the fraction indistinguishable from control is strictly positive in-band and exactly zero out-of-band. In the synthetic demonstration, 383 of 100,000 values (0.383%) against 0 of 100,000 (0.000%). *Falsified by:* an in-band construction — one in which control and data share a namespace — with an ambiguity rate of exactly zero, or an out-of-band construction with a non-zero rate.

H2 — the disclosure hypothesis. A design whose effectiveness depends on concealment fails when the design leaks; a design whose effectiveness is structural does not. *Predicted:* publishing every element of a channel-separated architecture — that the integrity witness is out-of-band, that settlement separates command from value, that messaging separates identity from payload — leaves the ambiguity rate at exactly 0.000%, because the rate is a function of the partition and not of the observer's knowledge. *Falsified by:* an element of a separated design whose disclosure raises the ambiguity rate, or an attack on a separated architecture that is enabled purely by knowledge of the separation and not by access to a plane.

H3 — the necessity hypothesis. Channel separation is necessary and not sufficient: it closes the impersonation class and says nothing about who may reach the control plane. *Predicted:* a separated control plane remains attackable by parties who already have access to it, and the frequency of such attacks rises as the set of parties with access widens. *Falsified by:* a separated architecture that is secure against control-plane attack without any access control on the control plane.

H4 — the provenance hypothesis. Identity crossing a stateful translation boundary is compressed to an anonymity set whose size is the sharing ratio at that boundary, and the binding that would undo the compression is ephemeral by design. *Predicted:* with 500,000 subscribers behind 256 public addresses, an identity worth 18.932 bits going in is resolved to 8.000 bits of address; the anonymity set is 1,953.1 with an address-only log leaving 10.932 bits unresolved, exactly 1 with an address-and-port log while session state is live leaving 0.000 bits unresolved, and 1,453.2 after expiry leaving 10.505 bits unresolved. *Falsified by:* a server-side log that resolves the originator without the translator's own mapping for the relevant instant.

H5 — the migration-asymptote hypothesis. A migration that moves a constant fraction of a shrinking remainder, at a per-line cost rising faster than the remainder shrinks, never completes and gets more expensive every year it continues. *Predicted:* at 15% of the remainder per year and a per-line cost rising 20% a year, the annual bill rises by a factor $(1-0.15)(1+0.20)=1.02$ per year, 3.29% of a 40-million-line core remains after twenty years, and the migration unit cost crosses the keep-running cost of 3.00 per line-year around 2032 with 37.6% of the core still buried. *Falsified by:* a migration under those rates that reaches

zero, or an annual bill that falls as the remainder shrinks.

H6 — the maintainer-cliff hypothesis. Where the rate at which competent maintainers are replaced sits below the rate at which they leave, the population lands at a steady-state deficit rather than passing through a transient. *Predicted:* at 11% annual attrition, 6 competent replacements a year and a flat requirement of 100 maintainer-equivalents, the population decays to a steady state of $6/0.11 = 54.5$ and crosses below the requirement in 2030, while a headline count of 120 still shows a 20-maintainer surplus today; holding the line would need 11.0 replacements a year against 6 actual. *Falsified by:* a steady state above the requirement under those rates.

Not hypotheses, and labelled as such.

- *The channel-separation principle itself* is a **definition**. It defines in-band and out-of-band by whether a data value can be promoted to a control instruction. What is testable is H1, which asserts a consequence of the definition, and H2, which asserts a property of the resulting design. **The definition is not evidence for anything and is not offered as such.**
- *The isomorphism across telephony, ledgers, messaging and settlement* is a **modelling claim** about shared structural form, not a derivation. Each mapping depends on the claim that the named system actually has the separation property attributed to it, which must be established per system and is not established here.

3. Method: the principle, stated as structure

3.1 The definition

> A system that carries control and data has a **control plane** — the information that determines what the system does — and a **data plane** — the information the system exists to move or store. The system is **in-band** when control and data share a channel such that a value in the data plane can be interpreted as an instruction in the control plane. The system is **out-of-band** when the two are separated such that no data value can be promoted to a control instruction. In-band architectures admit an impersonation class — data masquerading as control — that no care within the shared channel can close; out-of-band architectures close that class structurally, and remain subject to the separate question of who may reach the control plane.

The impersonation class is the object of study. **A structural statement of the same shape already exists in the distributed-systems literature and is worth naming, because it makes the class precise rather than evocative.** Douceur [2002] proves that in a system without a trusted certifying authority a single party can present as many distinct identities, so that an unaudited count of participants is worthless — which is the impersonation class in the identity domain, with the impossibility result attached. Lamport et al. [1982] fixes the number of genuinely independent parties required to tolerate f arbitrary faults at $3f + 1$, which is the same question asked about agreement. Neither result is imported here as a theorem about signalling; both are cited as prior formalisations of the same structural property, under the discipline of Section 10.5.

3.2 The historical case, and what it establishes

The early international telephone network carried two kinds of information over one path. There was the *voice* — the data, the thing the call existed to move — and the *signalling* — the control, the commands that seized a trunk, set up a route, supervised the call's progress and released it. In the in-band architecture the signalling travelled as tones in the same audio band as the voice, and the network's switches listened to the audio path for those control tones because that audio path was where the control information lived.

One structural fact follows, and it is stated as structure rather than as method. **If the switches take their commands from tones in the voice band, then any source that can put the right kind of sound into the voice band is speaking the control language: the network cannot distinguish a control tone that originated in its own signalling equipment from an identical sound arriving from the far end of the call, because they are the same sound in the same channel, and the channel is all the switch has to go on.** This is not a defect in any particular switch. It is the defining property of in-band signalling.

The exploits this permitted through the 1960s and 1970s are documented telephone history and are referred to here only at the level of principle — that the network could be controlled from the voice band. **No tone, sequence, frequency-as-method or device description is supplied**, and none is needed for the argument. The point is not the exploit. The point is that the exploit was not a defect to be patched but a direct consequence of the architecture, and that no amount of care in building the switches could have closed it, because the opening was the shared channel itself.

The fix, adopted across the international network in the 1980s, was architectural: the signalling moved out of the voice band entirely, onto a separate dedicated channel that the voice never touches — *common-channel signalling*, in which the control information for many voice circuits travels together on its own network, logically and physically distinct from the circuits carrying the calls. **After signalling moved out of band, a sound in the voice path was just a sound.** The class of exploit that depended on speaking the control language from the data channel did not get harder; it became impossible, because the data channel and the control channel were no longer the same channel.

That is the difference between a mitigation and an architecture. A mitigation makes an attack more costly; an architecture removes the conditions the attack required. This is the cleanest large-scale example in the history of communications of a structural fix closing a vulnerability class rather than patching its instances, and it is the reason the principle generalises: an engineer could build a perfect in-band switch — no bugs, no oversights, no weak components — and it would still take its commands from the voice band, because taking commands from the voice band was its design. **There was no patch available, because a patch is a change within an architecture and the problem was the architecture.**

3.3 What present-day signalling vulnerabilities are, and are not

Honesty requires the next paragraph, because a reader who knows that the modern signalling network is discussed as vulnerable would rightly distrust an account presenting out-of-band signalling as a final victory.

Out-of-band signalling closed the class of attack in which a *caller* controls the network from the *voice band*. It did not make the signalling network invulnerable, and the present-day discussion concerns a different class of attack entirely. **Contemporary signalling vulnerabilities are network-interconnect attacks carried out by parties who already have access to the signalling network itself — access to the out-of-band channel — and they are nothing like a sequence of tones a caller can dial into a telephone.** They exist because the signalling network was designed for a small, mutually trusting set of national carriers and was later opened to a far larger and less uniformly trusted set of interconnecting parties, so that access to the control channel, once implicitly restricted, became something a wider range of parties could obtain. **The vulnerability is a *trust* property of who can reach the control channel, not a *separation* property of whether control and data share a channel.**

The distinction is the whole reason for raising it. The channel-separation fix did its job and the class it closed stayed closed: **no caller dials tones to control the modern network, and this paper supplies nothing toward any present-day signalling attack, because such an attack is a matter of illicit access to a control network and not of technique that any description here could provide.** What the modern situation teaches is the *limit* of channel separation, which is H3: separating the channels closes the impersonation class and does not by itself secure the control channel, which must additionally be protected by controlling who can reach it.

Separation is necessary; it is not sufficient.

That limit has a name in the security-engineering literature. Saltzer and Schroeder [1975] list *open design* and *complete mediation* as separate principles precisely because they are separate: a design may be publishable and still fail if every access is not checked. The organisational literature reaches the same conclusion from field evidence. Perrow [1984] treats tight coupling as the property that turns a local fault into a system failure — cited here as qualitative framing, not as a formal result — and Sagan [1993] documents cases in which added redundancy *reduced* safety by adding coordination surface, which is the direct warning against treating a separated plane as automatically safer. Vaughan [1996] and Leveson [2011] supply the mechanism by which a control that was once justified continues to be trusted after the conditions justifying it have gone, and Hollnagel et al. [2006] the framing in which that drift is the normal state of an operating system rather than an aberration.

3.4 The isomorphism, tabulated

Four systems, one flaw, one fix.

system	in-band flaw: control impersonable by data	out-of-band fix
telephony	control tones live in the voice band	common-channel signalling on a separate network
ledger	the display form *is* the stored form	canonical form plus a raw byte-hash kept as separate planes
message	control fields tangled with the payload	a typed control plane distinct from the data plane
settlement	the instruction rides the value rail	separate channels for command and for value

Read each row as the same sentence. In telephony the control and the data shared a channel and data could impersonate control; separating them closed it. In the ledger, the control — what a human reviewer sees and trusts — and the data — what the machine stores and reconciles — shared one field, and a homoglyph let the seen value and the stored value diverge, until the two were separated by keeping a canonical form for the machine and a raw byte-hash as a witness on a distinct logical plane [Curzi, 2026]. In the message, the party-identifying control and the payment-describing payload were tangled in under-specified lines until structured typing separated them; the structured financial messaging standard [International Organization for Standardization, 2013] is the instrument that makes that separation available. In settlement, the instruction to move value and the value itself ride the same rail unless they are deliberately separated, so that whoever controls the instruction controls the value.

The mapping in this table is a claim of shared structural form and not a derivation. Each row depends on the assertion that the named system actually has the separation property attributed to it, which must be established per system. Section 10.5 labels the borrowing accordingly, and Cartwright [1999] supplies the discipline the label enforces: state the arrangement within which a law is claimed to hold, and stop claiming a universality the argument does not need. **The claim here is that four named arrangements share a form, not that the form is universal.**

3.5 Out-of-band integrity, and out-of-band settlement

The deepest instance is the integrity design of Curzi [2026], and stating it in this vocabulary is what makes its central specification decision look inevitable rather than merely clever. The

integrity hash is computed over the *raw octets* and stored *alongside* the canonical form, not over the canonical form itself:

```
data plane      canon(x)                what the matcher compares
control plane   hash(raw octets)        the witness, on a separate plane
detector        same canon, differing   something moved that the data plane
                hash                    cannot show
```

A hash of the canonical form is the in-band design: it puts the witness into the same plane as the data, so the witness changes exactly when the data changes and never otherwise, which destroys precisely the evidence a witness exists to preserve. Hashing the canonical form collapses the control plane into the data plane, which is in-band signalling in a ledger, and it reintroduces the impersonation class that keeping the byte-hash separate was there to close. The out-of-band design — witness on its own plane, never riding the data it guards — is the same architecture, executed over octets instead of voice circuits.

The same principle applied to money is the framework's standing rule that a rail is not a holding. The instruction to move value is a control-plane object; the value is a data-plane object. When they ride the same channel — when the message that commands the transfer is inseparable from the transfer, so that seizing the message is seizing the money — the settlement system is in-band, and control of the instruction is control of the value. Out-of-band settlement separates them: the command travels on one channel and the value on another, the two are reconciled by a witness rather than by their shared identity, and a compromise of the command channel does not by itself move the value because the value was never in the command channel. **The transit-versus-store distinction is the channel separation principle wearing a treasury's clothes: the rail is the control plane and the holding is the data plane, and keeping them separate is what stops a transit event from becoming a custody event.**

4. Results: the impersonation demonstration

A **synthetic** demonstration, computed under a fixed seed and deliberately abstract, so that it carries the structure and no technique. Model a channel carrying values drawn from a symbol space, of which a small reserved set are control tokens. In the in-band model the data is free to take any symbol, including a control token, so a data value can collide with a control token and be read as a command. In the out-of-band model the control tokens live on a separate channel the data never reaches, so the collision cannot occur.

```
IN-BAND  (control + data share one namespace)
  data values carried                              100000
  data values indistinguishable from control          383
  ambiguity rate                                    0.383%
  -> a data value CAN be read as a control command; structural.

OUT-OF-BAND  (control on a separate channel; disjoint by construction)
  data values carried                              100000
  data values indistinguishable from control            0
  ambiguity rate                                    0.000%
  -> the class is closed by construction.
```

The in-band model produces a non-zero rate of data values indistinguishable from control, and no care taken within the shared channel drives it to zero; the out-of-band model produces exactly zero, by construction, because the two planes never meet. H1 is supported in the model.

What the demonstration does not establish. The specific rate 0.383% is a property of the synthetic symbol space and carries no meaning beyond illustrating the sign: in-band greater

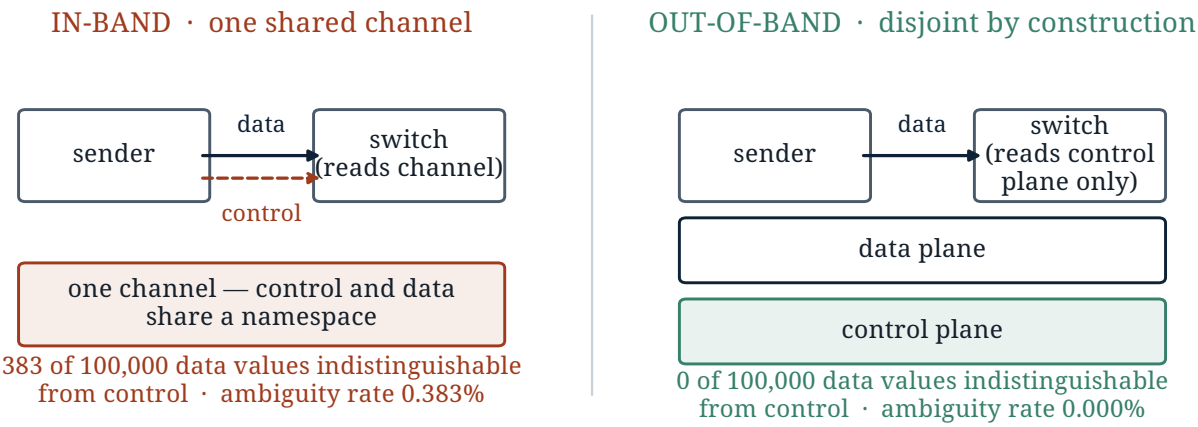


Figure 66.1: **The two architectures and their measured ambiguity: an in-band channel in which control tokens and data values share a namespace produces 383 indistinguishable values in 100,000 (0.383%), while planes that are disjoint by construction produce exactly 0 (0.000%).** Modelled on a synthetic symbol space under a fixed seed; the rate is a property of that space and corresponds to no real signalling system. Script `figs/src/CCP-066_fig01.py`.

than zero, out-of-band zero. **The demonstration corresponds to no real signalling system, measures no real prevalence, and would take a different value under any other allocation of reserved symbols.** What is invariant is the sign, and the sign is the finding, because it follows from the partition rather than from the parameters. The zero is exact and is not an estimate: it is a property of a disjoint union.

No distributional or scaling claim is made here or anywhere in this paper. The symbol space, the subscriber pool and the port space are finite sets with stipulated sizes; no frequency distribution is fitted, no tail is characterised, and no power law is asserted. This is stated rather than assumed because the material invites such a claim — a sharing ratio and a reuse count look like the beginnings of a heavy-tail argument — and Clauset et al. [2009] establishes that such claims are systematically overstated when they are not tested by a maximum-likelihood fit with a goodness-of-fit test and likelihood-ratio comparison against alternatives. **That procedure was not run, because no claim requiring it is made.** Broido and Clauset [2019] is the reason the same caution attaches to the network results this paper borrows in Section 9.1: across nearly a thousand empirical networks, social networks are among the weakest categories for scale-free structure, so the degree-distribution mechanism behind the robustness literature must not be assumed to apply here.

Why a constructed model is worth running at all is a fair question and it has a published answer. Epstein [2008] sets out the reasons to model other than prediction — among them demonstrating that a proposed mechanism suffices to produce an observed shape, and exposing which assumption a conclusion actually rests on — and that is precisely the work the three models in this paper do. Bruch and Atwell [2015] supplies the calibration and validation checklist such a model should meet before it is treated as anything more, and **none of the three models here meets it**, which Section 11 records. Hedström and Ylikoski [2010] states the criterion this paper accepts: an explanation must name a mechanism rather than exhibit a fitted curve. **Nothing here is fitted; each result follows from an inequality that is stated in advance.**

5. Results: the disclosure constraint

This is the paper's load-bearing result and it is stated as a hard constraint rather than as a recommendation.

Channel separation is strengthened rather than weakened by full disclosure, because its effectiveness is entirely independent of whether anyone knows it is in place. Publishing that an integrity witness is out-of-band tells an adversary only that substitution will be detected. Publishing that a settlement architecture separates command from value tells a regulator only that the rail is not a store. Publishing that a messaging arrangement separates identity from payload tells a correspondent only that its payments resolve. **Every such disclosure is an assurance and none is an opening.**

The reason is arithmetic rather than rhetorical, which is why the constraint is stated as one. The ambiguity rate is a function of the partition of the symbol space and of nothing else. Disclosing the partition does not change the partition. **The rate stays at exactly 0.000% after disclosure because disclosure is not one of its arguments.**

That argument is not new and this paper does not claim it. **Kerckhoffs [1883] stated it in 1883: a system must remain secure when everything about it except the key is public.** Saltzer and Schroeder [1975] restate it as *open design* — the protection mechanism should not depend on the ignorance of attackers — and place it alongside *complete mediation* precisely because a publishable design still requires that every access be checked. **Citing Kerckhoffs converts what reads as an ethical preference into a recognised engineering standard, and that single citation carries more of this paper's weight than any other.**

A published architecture is performative, and that is an argument for the constraint rather than against it. MacKenzie and Millo [2003] documents the case in the field and with dates: an option-pricing model that did not fit the market when it was published came to fit the market after traders adopted it, so that the description made its own assumptions true. Hacking [1995] states the general form — a classification of people interacts with the people classified, who change in response, so the referent moves. **The disclosure constraint is the benign form of that loop: publishing that a design separates its planes causes counterparties, regulators and operators to rely on the separation, and reliance is exactly what makes the separation load-bearing.** The malign form is failure mode 6.3, in which the published separation is quietly re-merged while the reliance persists. **Neither form is available to a concealed design, which is a further reason the constraint is engineering rather than preference.**

The constraint has since become unavoidable from the regulatory direction as well. Automatic exchange of account information [Organisation for Economic Co-operation and Development, 2017], the anti-money-laundering recommendations that supervisors implement in common [Financial Action Task Force, 2012], and the substance and beneficial-ownership work of the base-erosion project [Organisation for Economic Co-operation and Development, 2015] together mean that non-disclosure is no longer an available design parameter, whatever a designer's preference. Brandeis [1914] is the older statement of the same policy instinct from the transparency side. Pistor [2019] is the scholarly account of how legal coding manufactures the durability that structures seek, and is simultaneously the sharpest critique of it; Hansmann and Kraakman [2000] establishes that asset partitioning is the irreducible function of entity law, which is the legal-form analogue of a plane separation — **a separate legal personality is a partition, and its effectiveness likewise does not depend on anyone's ignorance of it.** de Goede [2011] documents the listing machinery as a technique of governance, which is the mechanism that makes concealment fail on contact.

5.1 The caveat that belongs in print

A separated control plane must still be access-controlled, and disclosing that a control plane exists raises rather than lowers the importance of governing who reaches it. A structure that publishes its architecture must also be the structure that controls access to its control planes — its settlement command channel, its integrity witness, its messaging identifiers — because

separation without access control is the failure mode of Section 6.5, and the disclosure that advertises the separation is also the disclosure that tells an adversary which channel is now worth reaching. **That is not a reason to conceal the architecture, which cannot be secured by being secret; it is a reason to secure the access, which can.**

This is H3, and the modern signalling network is its standing demonstration.

6. Failure modes of the separation principle

6.1 Nominal separation that is not real. The commonest way channel separation fails is that it is declared and not built. Two planes that are logically distinct but implemented over one shared substrate — one wire, one field, one identifier space — are separated on the diagram and in-band in fact. Hashing the canonical form is exactly this failure in a ledger: the witness is nominally a separate thing and actually a function of the data. **The test is not whether the design *says* the planes are separate; it is whether a compromise of the data plane can reach the control plane.**

6.2 Covert channels — separation defeated from within. Even where control and data are architecturally separate, a party with a foot in both can sometimes signal across the boundary by modulating something the data plane controls — timing, volume, ordering — into a pattern the control plane can read. The class is named here to bound the claim: **separation closes the *direct* impersonation class, and an adversary with access to both planes may still contrive an *indirect* one**, which is a reason to control access to the planes rather than a reason to doubt the separation.

The point at which the analogy stops being an analogy and becomes a warning about the settlement architecture is worth stating. Out-of-band settlement separates the command channel from the value channel, and a party with a foot in both — a treasury operator who both issues the instructions and holds the balances — can reconstruct the bridge the separation was meant to remove, not by breaking the architecture but by occupying both sides of it. **This is why a maximum-overnight-balance limit must be enforced by someone who does not report to the treasurer: the separation of duties is the access control on the reconstructed channel, and without it the two planes are separated on the diagram and joined in a single person.** Every out-of-band architecture is only as separate as the access controls that keep any one party from standing in both planes at once. **Separation is a property of the wiring; keeping it is a property of who is allowed to touch the wiring, and the second is a governance problem that the first does not solve.** This is the governance form of the same result that Ostrom [1990] reaches empirically for common-pool institutions: the rule is only as good as the monitoring and sanctioning arrangements around it.

6.3 Re-merging the planes for convenience. Separated planes are less convenient than merged ones — two channels to build, two to operate, two to reconcile — and there is a standing temptation to fold the witness back into the data, the instruction back into the value, the control field back into the payload, because the merged design is simpler and, until it is attacked, works identically. **Every re-merge is a return to in-band, and it is usually justified as an optimisation.** The discipline is to recognise that the separation *is* the control, so removing it to save cost is removing the control to save cost. The pressure is a standing force rather than a one-time temptation: separated planes cost more on every ordinary day and pay off only on the exceptional one, **so the case for re-merging them strengthens with every uneventful year**, and the argument for keeping them is always an argument about a tail event that has not happened yet. Organisational slack has always appeared on the accounts as waste [Cyert and March, 1963], and this is that result in a technical register.

6.4 Mistaking the principle for a technique. Channel separation is a structural principle about the architecture of systems. It is not a method for attacking any system, and the history it draws on is closed history precisely because the architecture changed. **There is no technique in this paper**: the historical exploits are recounted only to establish that in-band control was structurally reachable, and present-day signalling attacks are a matter of illicit

access to control networks that nothing here describes or could enable. **The principle is for building, not for breaking.**

6.5 The control plane trusted because it is separate. The inverse of 6.1. A separated control plane can be *over*-trusted precisely because it is separate — treated as authoritative because data cannot reach it, while the question of who *can* reach it goes unasked. The signalling network was trusted because it was out of band, and its later vulnerabilities came from that trust outliving the conditions that justified it. **A control plane that everything trusts and anyone can reach is a worse exposure than an in-band design that trusts nothing.**

7. The addressing substrate

7.1 Three hierarchies, one pattern, no common root

Three independent systems say *where*, and each is a hierarchy that delegates authority downward from a wide root to a narrow leaf. **They do not share an origin, a designer, or an underlying registry.**

The international public telephone numbering plan sets a maximum of fifteen digits and divides them by authority: the leading digits are a country code assigned centrally, the country code selects a national authority which administers a national destination code, and within that a subscriber number. Three levels, three delegated authorities, each narrowing the space the level above handed it. An internet protocol version 4 address is thirty-two bits, which is exactly 4,294,967,296 possible values — a number this paper returns to, because its smallness is the origin of everything downstream. Addresses are delegated from a central registry to regional registries, from those to local registries and large operators, and from those to end sites; routing then reads the address as a prefix hierarchy of network, subnet and host. A hierarchical location code narrows a coarse region to a fine one by prefix, country to region to locality, or cell to sub-cell.

The claim made: these three are the same *pattern* — nested hierarchical delegation from a wide root to a narrow leaf — and that is why a designer who understands one understands the failure modes of the others. **The claim refused:** that they are the same *system*, or share an origin, or run on common infrastructure. They were designed by different bodies, in different decades, for different media, and they interoperate nowhere except where something deliberately builds a bridge. **The pattern is the finding; the common system is the overclaim; and blurring the two forfeits the finding.**

The distinction is not pedantry, because the pattern alone predicts the vulnerability. Any nested-delegation namespace has the property that authority over a leaf is inherited through the chain of delegations above it, so that to know who owns a leaf one must be able to walk the chain — **and every place the chain is broken, rewritten or shared is a place where ownership becomes ambiguous.**

The specifications relied on are not in this paper's bibliography. The numbering plan, the two internet protocol versions, the transition mechanisms and the translation specifications are public standards cited in the source manuscript by designation. **They are described here from that manuscript and are not cited to the bibliography, which contains no standards documents.** Section 11 records this as an unclosed gap.

7.2 The last mile, stated correctly

The long-haul core of the internet is optical. Traffic between cities and between continents runs on fibre — terrestrial long-haul routes and submarine cables — and has for long enough that a claim of copper dependency in the core is an error of fact. Metropolitan transport is likewise overwhelmingly optical. **The assertion that the internet cannot function without copper is false, and this paper states that it is false rather than repeating it**, because repeating

a false technical claim in a work that asks to be trusted on technical claims is a self-inflicted wound.

What is true is narrower and survives scrutiny. The *access* network — the last mile from the local node to the premises — is still, in a large installed base, copper: digital subscriber line service runs data over the same twisted pair that carried the telephone call, and the physical drop into older buildings is frequently copper even where fibre reaches the kerb. That is a real, current dependency of the *access* layer specifically.

And there is a second, more durable sense in which the old telephone plant persists. The topology of the telephone system — the exchange as a point where lines are switched, the central office as a node in a routing hierarchy, the numbering plan as a tree of delegated authority — is a *logical* structure, and logical structures outlive their physical implementations. **The exchange and central-office topology remains an addressing substrate underneath modern transport, and it persists because it is a logical map, not because any wire in it survives.**

7.3 The bottleneck and the bridge

The thirty-two-bit address space contains about 4.29 billion addresses: an unimaginable abundance when the protocol was designed and a binding scarcity within a human lifetime. The central authority exhausted its free pool in 2011, after which the regional registries depleted their own pools on staggered timelines. The address space did not grow; the demand did.

The designed replacement has a hundred-and-twenty-eight-bit address space of about 3.4×10^{38} addresses — a ratio of about 7.9×10^{28} to the old space. **The problem is not the size of the new space; it is the seam between the two, and it is a hard seam.** The two protocols have different address lengths, different headers and no shared wire format; a packet in one is not a packet in the other. **This is not a defect. It is the consequence of designing a successor with room to grow rather than a patch constrained to the old format.** But it means the migration could never be a single cutover, and it created a decade-and-counting period in which both protocols coexist and traffic crosses between them constantly.

The bridge is not one mechanism but a family, and each member is a place where address data is handled differently.

mechanism	what it does	provenance consequence
dual-stack	a host runs both protocols and uses whichever the destination speaks	no translation, so no loss; does not solve scarcity, since a dual-stack host still needs an old-protocol address
stateful translation	lets clients on the new protocol reach servers on the old, keeping a per-session table of which session maps to which address and port	rewrites; the mapping exists only in the translator, only for the session
edge-plus-core translation	stateful translation in the operator core plus stateless translation at the customer edge	layers a second translation on the first
tunnelling	carries one protocol's packets as payload inside the other's rather than translating	preserves the inner address exactly; adds encapsulation boundaries

The bridge between the two address spaces is not a wire; it is a set of *rewriting* operations, and rewriting is where the trouble is. Dual-stack rewrites nothing and solves nothing about scarcity. Tunnelling wraps rather than rewrites and preserves the inner address. Translation rewrites.

7.4 Where provenance is lost

At a translation boundary, the address that comes out is not the address that went in. A globally unique source address, delegated down a chain that ends at one subscriber, enters the translator; a shared address drawn from a pool, plus a port number, comes out. The server on the far side sees the second and never the first. The relationship between them — the fact that *this* host was, at *this* instant, using *that* address and port — exists only inside the translator, only for the duration of the session, and is discarded when the session ends. **The mapping is real while it lasts and gone when it does not, and attribution depends entirely on whether anyone was recording the mapping at the moment it existed.**

That is the same structural failure as a character-encoding boundary, and the parallel is exact. At a charset boundary a byte-level hash is a faithful record of the octets and useless for meaning, because meaning is assigned by a decoder the hash cannot see [Curzi, 2026]. Here the server's log is a faithful record of the observed address and useless for attribution, because attribution is assigned by a translator the log cannot see. **In both cases the observable is honest and insufficient, and the sufficient information lives at a boundary the observer is not positioned to see.**

8. Results: the substrate measurements

8.1 The constructed carrier

Everything in this section is **CONSTRUCTED**. The figures are chosen to be plausible and round, not to describe any real operator, and the point is the arithmetic of the boundary rather than the specific inputs.

```
active subscribers                           500,000
public address pool (a /24)                      256   addresses
usable ports per public address                64,512   (1024–65535)
simultaneous (address, port) bindings      16,515,072
subscribers time-sharing one public address     1,953.1
```

Five hundred thousand subscribers reach the older internet through two hundred and fifty-six public addresses. Each public address offers the usable port range, and the product — about 16.5 million simultaneous address-and-port bindings — is the number of distinct conversations the translator can carry at one instant. **No subscriber owns a public address; they borrow one, and a port on it, for the length of a connection, and hand it back.** This is not an exotic configuration: it is the ordinary way a large address-constrained operator connects a subscriber base far larger than its address holding, and the sharing ratio is a routine engineering figure rather than an extreme one.

8.2 H4 — compression at the boundary and the attribution set: supported

Identity is measured in bits as the base-two logarithm of the size of the set an observer can narrow the originator to, which is Hartley's measure of information in a selection from a finite set [Hartley, 1928] and the special case of Shannon entropy for a uniform distribution [Shannon, 1948, Cover and Thomas, 2006].

```
identity going IN      : 1 of 500,000 subscribers     = 18.932 bits
server sees address    : 1 of 256 public addresses    =  8.000 bits
server sees addr + port: 1 of 16,515,072 tuples       = 23.977 bits
```

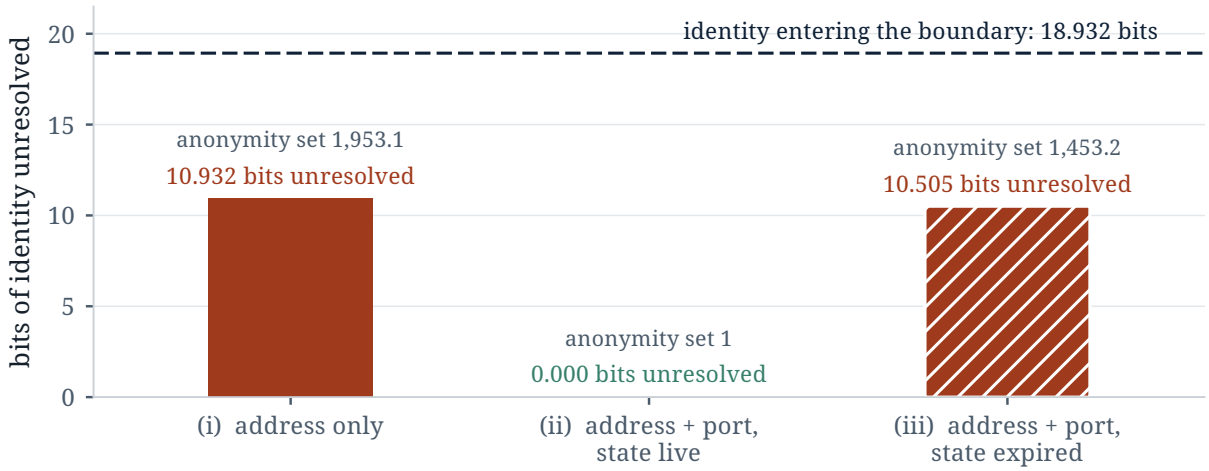


Figure 66.2: **Identity crossing a stateful translation boundary.** An identity worth 18.932 bits going in is fully recoverable only in case (ii), where the port was logged and the translator's state is still live; an address-only log leaves 10.932 bits unresolved and an expired-state log 10.505, both corresponding to anonymity sets in the thousands. Modelled from the CONSTRUCTED carrier parameters in Section 8.1; script `figs/src/CCP-066_fig02.py`.

The operative question is not how many bits the server records but how many subscribers could have produced what the server recorded — the *anonymity set* — and equivalently how many bits of the originator's identity are unresolved after the crossing.

case	what was logged	anonymity set	bits unresolved
(i)	address only	1,953.1	10.932
(ii)	address and port, session state live	1	0.000
(iii)	address and port, session state expired	1,453.2	10.505

Case (ii) is the reassuring one and the one people implicitly assume. If the server recorded the address and port, and the translator's session table is still in memory and still holds the entry, the tuple maps to exactly one subscriber and attribution is complete: zero bits unresolved. **It is a real case, and it is bounded in time.**

Case (i) is the common reality of server-side logging. A great deal of software logs a source address and not a source port, because the port was never thought to matter. Against a shared carrier address, **an address-only log resolves the originator to a set of about 1,953 subscribers and leaves 10.9 bits of identity unresolved — the log names the pool, not the person.** The originator went in as one of five hundred thousand and comes out, in the record, as one of nearly two thousand, and no amount of care at the server closes the gap, because the gap was created upstream.

Case (iii) is the one that matters most and is least understood. Suppose the server *did* log the port, but the translator's session state has since expired — the connection ended, the entry aged out, the memory was reused. The same tuple is not reserved to that subscriber; it is handed to the next subscriber who needs it, minutes later:

```
sessions per subscriber per day (proxy)         48,000
total sessions per day                  24,000,000,000
reuse per (address, port) slot per day           1,453.2
```

So a tuple logged at the server, once the state that bound it to a subscriber has expired, could belong to any of the roughly **1,453 subscribers who used that same slot that day** — 10.5

bits unresolved, almost as bad as logging no port at all. **The address the server holds is honest and the subscriber it names is unrecoverable, not because anyone hid anything but because the binding was ephemeral by design. H4 is supported.**

8.3 Why the state is ephemeral, and not by accident

The obvious objection is that the carrier should simply keep the session table. The measurement shows why that is not free and often not done:

```
bytes per session record                          150
daily log volume                  3,600,000,000,000 bytes = 3.600 TB/day
 90-day retention                           324.0 TB
180-day retention                           648.0 TB
365-day retention                         1,314.0 TB
```

Twenty-four billion sessions a day at a modest 150 bytes a record is **3.6 terabytes of translation log per day, and over 1,314 terabytes a year, for one metropolitan translator**. The volume is not impossible; it is a deliberate cost, and **it is a cost that buys nothing the carrier needs for its own operation**, since the carrier does not require the historical mapping to route the next packet. The state is ephemeral because keeping it is a standing expense incurred entirely for the benefit of a future third party who may never ask. **The provenance is lost by an economic gradient, not by an oversight, which is why the loss is structural rather than a bug someone forgot to fix.**

The mirror of that finding is a liability rather than a gap. **A party that captures and retains attribution data has created a record that is itself discoverable and itself sensitive**, and the decision to retain it deserves the same care as the decision to collect it, because a retained mapping is a store of exactly the information the addressing substrate was, by its economics, disposed to forget. Lazer et al. [2009] sets out the access, privacy and review conditions under which behavioural records of this kind may responsibly be held, and they apply in full to any retention window designed on the basis of this section. **Retention is a control and a liability at once.**

8.4 Deterministic assignment: the anonymity set becomes a design parameter

There is an engineering response to the logging cost worth measuring, because it makes the provenance loss explicit and turns the anonymity set into a number the operator chooses on purpose. Deterministic assignment gives each subscriber a fixed block of ports on a shared public address, so that attribution needs only a small static block map rather than a per-session log:

```
fixed port-block per subscriber                 512 ports
blocks per public address                       126
deterministic capacity (1 block/subscriber)   32,256 subscribers
required oversubscription                      15.50x
block size that actually fits 500,000 subs        33 ports/subscriber
static block table                            32,256 rows
  (against 24,000,000,000 session rows/day)
```

With 512-port blocks, one public address carries 126 subscribers and the whole pool carries 32,256 — far short of the half-million, an oversubscription of **15.5 times**, which means the blocks must themselves be time-shared or the pool enlarged. Shrink the block to fit and each subscriber gets about thirty-three ports. The static block table is tiny — about thirty-two thousand rows against twenty-four billion session rows a day — which is precisely the appeal.

But attribution now resolves to the *block owner*, which is one subscriber only if blocks are not shared, and the moment oversubscription forces block sharing the anonymity set is exactly the number of co-tenants of a block. Deterministic assignment does not eliminate the provenance loss; it converts it from an accident of expired state into a fixed, chosen anonymity set — the operator decides how much provenance to destroy in exchange for how small a log to keep. That is the honest description of the trade: the cheap control or the complete evidence, and the design decision is which.

8.5 H6 — the maintainer cliff: supported

A **constructed** mainframe shop with 120 expert maintainers of a critical estate. The estate does not shrink — the buried layer is load-bearing and its maintenance demand is flat at a required 100 maintainer-equivalents. Experts leave at 11% a year, and the shop trains or hires 6 new competent maintainers a year, where *competent* means able to modify the estate safely, which for these languages is a multi-year apprenticeship.

year	experts	gap (experts − required)
2026	120.0	$+20.0$
2028	106.4	$+6.4$
2030	95.6	-4.4 (cliff)
2032	87.1	-12.9
2036	75.0	-25.0
2040	67.4	-32.6
2050	58.5	-41.5

```
steady-state experts (replacement / attrition)   54.5
steady-state gap vs requirement                 -45.5
replacement needed just to hold the line         11.0/yr  (vs actual 6)
                                                 -> deficit 5.0/yr
```

The shop looks fine today — a 20-maintainer surplus — and it is not fine. The expert count decays toward a steady state of **about 54 maintainers against a flat requirement of 100**, and crosses below the requirement in **2030**, four years out, while still appearing healthy in the intervening years. **The steady-state deficit is not a transient to be waited out; it is where the system lands and stays, because the replacement rate is structurally below the attrition rate. The cliff is not an event; it is the arithmetic of a replacement rate below an attrition rate**, and it arrives while the headline count still looks adequate, which is exactly why it is not acted on until it is too late. **H6 is supported in the model.**

The distributional point is worth stating because the headline count conceals it. Granovetter [1978] establishes, for collective behaviour generally, that the **distribution** of individual thresholds and not the mean determines the aggregate outcome, and the same holds here: a maintainer population is not interchangeable, and a shop whose remaining experts are concentrated on a few subsystems crosses its effective cliff earlier than the aggregate count says. **The model tracks a mean and therefore gives the optimistic answer**; the distribution is not modelled and no claim is made about it.

The model's honesty is in what it does not claim. **The specific numbers are invented; the shape is the finding.** Any shop where the replacement rate sits below the retirement rate lands in the same place, and the only levers are raising replacement — slow, because competence is slow to build — or lowering the requirement, which means migrating off the estate, which Section 8.6 shows is itself a trap.

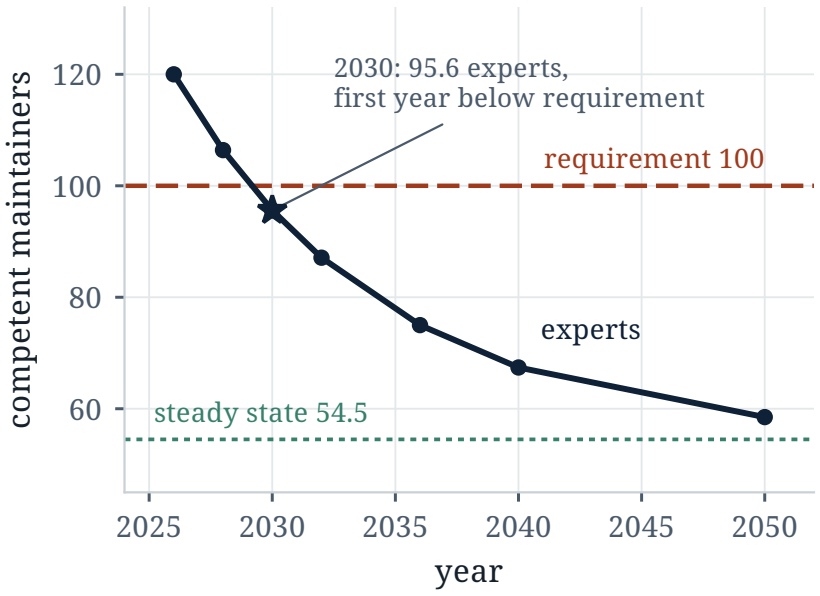


Figure 66.3: **The constructed maintainer cliff.** A shop showing a 20-maintainer surplus today crosses below its flat requirement of 100 in 2030 and settles at a steady state of 54.5, because the replacement rate sits structurally below the attrition rate. Modelled with 120 experts, 11% attrition and 6 competent hires a year; script `figs/src/CCP-066_fig03.py`.

8.6 H5 — the migration asymptote: supported

A **constructed** core of 40 million lines. The shop migrates 15% of what *remains* each year — the easy, well-understood parts go first, because they are easy and well-understood — and the per-line cost of migration rises 20% a year, because what remains after the easy parts are gone is the hardest, least-documented, most entangled core.

year	remaining	migrated	unit cost	year cost	cumulative
2026	40,000,000	6,000,000	1.00	6,000,000	6,000,000
2030	20,880,250	3,132,038	2.07	6,494,593	31,224,241
2034	10,899,621	1,634,943	4.30	7,029,956	58,527,771
2038	5,689,670	853,451	8.92	7,609,451	88,081,989
2042	2,970,043	445,507	18.49	8,236,714	120,072,426
2046	1,550,381	232,557	38.34	8,915,684	154,699,903

```
remaining after twenty years   1,317,824 lines  (3.29% of the core)
```

Two results, and the second is the important one.

First, the remainder decays geometrically — 85% of it survives each year — and geometric decay never reaches zero. After twenty years of steady migration 3.29% of the core is still there, and it will still be there after forty, because a constant fraction of a shrinking base is always positive. **The migration, run as a steady percentage, is asymptotic by construction; it approaches the bottom and never lands.** This is a first-order linear difference equation and its behaviour is standard [Elaydi, 2005].

Second, and worse, the annual bill does not fall as the remainder shrinks — it rises every year. Because the per-line cost climbs 20% while the remainder falls only 15%, the product grows by a factor $0.85 \times 1.20 = 1.02$ a year, so **the migration gets more expensive every year it continues, not less.** There is no point at which finishing becomes cheap; the hardest 3% is the most expensive 3%, and it is the part a rational budget defers indefinitely.

```
keep-running cost of the legacy (per line/yr)          3.00
migration unit cost crosses keep-running cost in       ~2032
lines still buried at that crossover             15,023,138   (37.6%)
```


Set against the cost of simply keeping the legacy running — cheap, because correct old

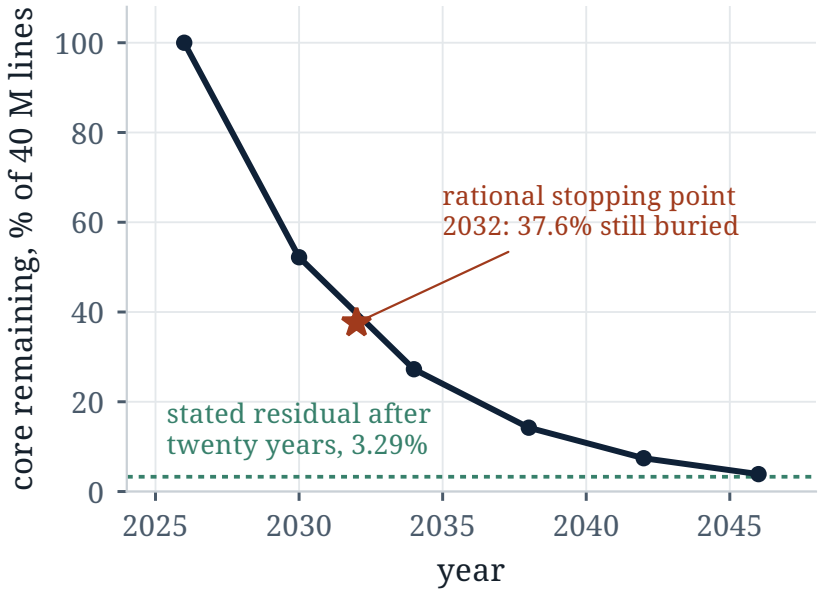


Figure 66.4: **The migration that never completes.** A constant 15% of a shrinking remainder decays geometrically and never reaches zero; the rational stopping point falls near 2032 with 37.6% of the core still buried. Modelled with a 40-million-line core, 15% of the remainder migrated per year and a per-line cost rising 20% a year; script `figs/src/CCP-066_fig04.py`.

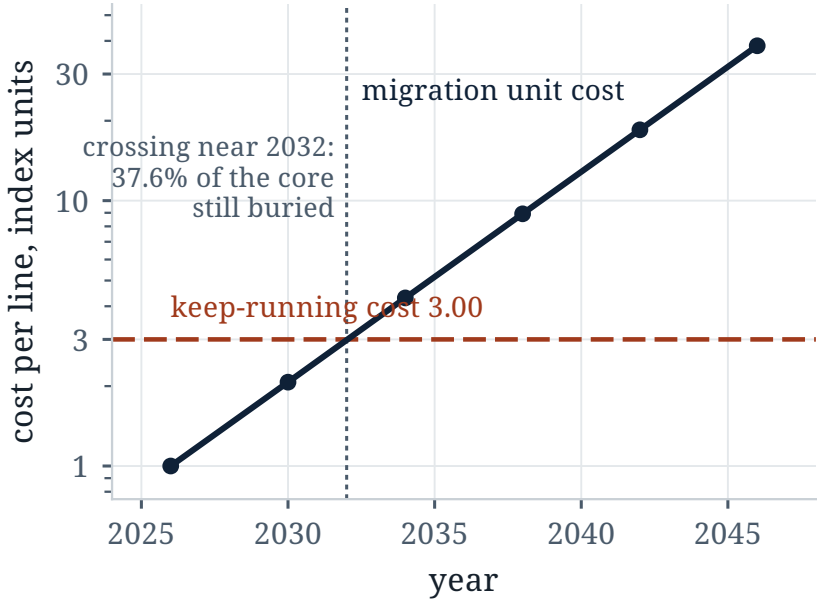


Figure 66.5: **Why the migration stops.** The per-line cost of moving the remainder rises 20% a year and crosses the flat keep-running cost of 3.00 per line-year around 2032; past the crossing, each further line costs more to move than to leave. Modelled from the same constructed parameters; script `figs/src/CCP-066_fig05.py`.

is, by revealed evidence, among the most thoroughly tested software in existence. **The risk is that the people who can read, modify and vouch for that code are retiring faster than they are replaced**, so that the code becomes progressively less maintainable not because it changes but because the human ability to change it safely evaporates.

A system nobody living fully understands is fragile no matter how correct it is, because the next required change — a regulatory deadline, a currency redenomination, a date-format rollover — arrives on a schedule the system does not control, and arrives at a maintainer population that has thinned below the level needed to make the change with confidence. **The failure mode is not that the code stops working. It is that the code keeps working perfectly right up until a change is required that no one remaining can make safely.**

The staffing constraint is not novel and has a literature. Brooks [1975] establishes that a project's coordination cost grows as $n(n-1)/2$ and that adding people to a late project makes it later — the direct statement that maintainer capacity cannot be bought quickly. Conway [1968] gives the complementary structural result that a system's interfaces mirror the communication structure of the organisation that built it, which is why an estate built by a now-departed organisation is hard to re-enter. **The source dates this article to 1967; it appeared in April 1968, and the corrected date is used here.** Hopp and Spearman [2011] and Kingman [1962] supply the general shape of a capacity constraint approaching saturation: waiting time explodes as utilisation approaches one, which is what a maintainer pool at the requirement line experiences when a change arrives. Sterman [1989] demonstrates experimentally that people systematically misjudge stocks with delays, which is the behavioural reason a shop with a 20-maintainer surplus does not act; Forrester [1961] is the underlying levels-and-rates formulation.

A public, dated instance exists. During the 2020 pandemic, unemployment-claim volumes in the United States spiked to levels the state benefit systems had never been built for, and several of those systems ran on decades-old mainframe code. When they buckled and needed rapid modification, states discovered that the maintainer population for their own systems had thinned to the point of crisis; one state's governor publicly appealed for volunteer programmers in April 2020. **Systems that had run correctly for decades could not be changed fast enough when change was suddenly required, and the failure landed at the exact moment the systems mattered most.** That is the mechanism made visible — a maintainer failure, not a code failure. **The specifics of that episode are described from the source manuscript and are not cited to the bibliography, which contains no contemporaneous reporting** (Section 11).

8.8 The citation discipline on the scale figures

The figures everyone quotes about the volume of legacy transaction-processing code — that it handles most card-swipe transactions, underpins a large share of banking systems, runs to hundreds of billions of lines, and supports systems processing trillions of dollars a day — are real quotations from real reporting, and they are also nearly a decade old, repeated so often that they have detached from their source and hardened into ambient fact. **This paper attributes them to their 2017 origin and does not restate them as a current measurement.**

The bibliography contains no press reporting, so the originating article is not cited here. The figures are therefore described rather than quoted, and are used only to establish an order of magnitude: the dependency is vast rather than marginal. **A work about buried provenance that laundered its own headline figures would be refuting itself**, which is the reason for the discipline rather than fussiness.

9. Discussion

9.1 What follows

A concentration measurement on one axis can hide a common-mode dependency on another. This is the sharpest consequence of the substrate results and it bears directly on the corpus's central instrument. Curzi [2026] measures resilience by counting effective coordination perimeters: a naive count of 3.41 effective jurisdictions collapsing to exactly 1.00 effective blocs once a shared perimeter is recognised. **Two banks in two genuinely independent perimeters, contributing two full units to that effective count, may sit on the same generation of core software with the same thinning maintainer pool, so that a forced substrate change could degrade both at once while the instrument reports them as fully independent.** The perimeters are independent; the substrate underneath them may not be, and the instrument cannot see the difference. **This is a common-mode failure and not diversity** — exactly the case Markowitz [1952] formalises for portfolios, in which diversification benefit depends on covariance rather than on count, and Longin and Solnik [2001] measures as correlation rising under stress. Buldyrev et al. [2010] is the network-theoretic statement that coupled systems are more fragile than either alone, cited here as a **shared structural form** rather than as a derivation: the substrate dependency has not been shown to satisfy the percolation premises of that model.

Attribution across a translation boundary is a capture-time decision, never a query-time one. If the mapping was not recorded when it existed, no later diligence recovers it. That is the single most useful operational consequence of the addressing result for anyone building a control that depends on knowing where a transaction originated.

An address is evidence of a pool, not of a person. A monitoring function that flags a transaction as coming from a party because it came from an address is asserting an identity it does not have. In the constructed carrier that assertion is wrong for 1,952 of every 1,953 subscribers sharing the address, and it carries no visible warning that the address is shared rather than dedicated. **This fails in the reassuring direction, naming a specific party with false confidence**, which is the same shape as the effective-count error at the jurisdictional layer: trusting a number that a lower layer has already degraded.

The rewrite is the cure that is worse. A big-bang replacement of a system nobody fully understands is an attempt to re-derive, under deadline, decades of accumulated correctness and undocumented edge cases. **It fails not on the behaviour that is documented but on the behaviour that is not, and it does so at exactly the moment — the maintainer cliff — when the people who could say what the old system actually did are already gone.** This is the exploration-exploitation failure of March [1991] in its most expensive form, and Levinthal and March [1993] names the temporal myopia that makes the trade look attractive at the decision point.

9.2 What does not follow

It does not follow that separation is sufficient. H3, Section 6.5, and the modern signalling network.

It does not follow that the buried layer is anyone's negligence. The migration model is the refutation. Migration off a correct, load-bearing core is not a will problem; it is an economics problem with a stable equilibrium above zero, and institutions that "failed to migrate" mostly made the locally correct decision at every step. **Treating the buried layer as a failure to be scolded rather than an equilibrium to be managed guarantees the wrong remedy** — it prescribes urgency where the need is a way to live safely with a layer that is not going away. Star and Ruhleder [1996] and Star and Griesemer [1989] give the general account: infrastructure is by definition the layer that becomes visible only on breakdown, and its invisibility is a property of its success rather than of anyone's inattention. Scott [1998] supplies the mirror

warning about the remedy: making a system legible destroys local variety that was doing real work.

It does not follow that attribution is impossible. It is expensive, and Section 8.3 prices it. A design that wants attribution must decide, before the fact, to pay for it.

It does not follow that the three namespaces share a system. Section 7.1, refused in print.

It does not follow that the internet depends on copper. Section 7.2, refused in print. **The durable dependency is logical topology, not electrical continuity**, and the two are constantly confused by people who have noticed that the telephone system is still somehow underneath everything and drawn the wrong conclusion about why.

9.3 Parallel results in the corpus

The same shape recurs across the corpus in domains that share no subject matter, and the recurrence is the reason these are one corpus.

Curzi [2026x] measures a receiver-limited channel capacity into a single cell at 70 to 106 bit/s, at a mechanotransduction bandwidth near 10.6 Hz and 20 dB signal-to-noise, delivering a 40-bit address in about half a second — **a capacity result on a channel whose control and data planes are not separated**, and the closest analogue elsewhere in the corpus to this paper's bit-accounting. Curzi [2026y] states the complementary negative result: the demodulator gap, in which no identified transducer, memory element or symbol-to-locus mapping exists, and capacity is not the bottleneck since 40 bits over 16 Hz needs only 6.68 dB. Curzi [2026] states the same open problem formally — no established mechanism produces a response depending on the symbol string as an *order*, as distinct from permutation-invariant functionals of the waveform — which is a representation-loss result of exactly the kind Section 7.4 measures. Curzi [2026w] gives the evidential form: two agreeing sources do not double the evidence when only one of them is evidence, which is the epistemic version of a control plane that has been folded into the data plane and no longer witnesses anything independently. Curzi [2026] reports an audit distribution over thirteen entries — seven maps, two splits, one rendering, three rejections, with five entries carrying a failure — which is the corpus's standing demonstration that a classification exercise is only worth running if it is allowed to reject.

9.4 Bounds carried from elsewhere in the corpus

Three results elsewhere in the corpus bound what this paper may claim, and they are recorded here rather than left implicit.

Curzi [2026] refuses category innocence as a defence of acquisition, forbids moving any harm through a discovered gap in a scheme, and concedes that the case where a feedback loop runs through another person's will is an unsettled seam the corpus does not close. **That bound is why this paper carries no operational content**: a structural principle that could be read as a technique for reaching a control plane would be exactly the discovered gap the doctrine forbids.

Curzi [2026] establishes that a width statistic invariant under relabelling does not determine the object it is computed from, which formally blocks an inference from the statistic to a substantive conclusion. **The same bound applies to the ambiguity rate here: 0.383% is invariant under relabelling of the symbol space and therefore does not determine the space it was computed from.** It establishes the sign and nothing more, and Section 4 says so.

Curzi [2026] classifies the corpus's commitments into four items of established mathematics, seven of established psychology, and exactly three declared postulates, with every protocol stated to run on the first two and steer by the third. This paper's chains are in the first drawer — arithmetic and definitions — and its three modelling assumptions are named in Section 10.

Curzi [2026t] supplies the discipline that makes such a demotion survivable: what survives a demotion is structure rather than magnitude — the limiting cases, the fixed point, the sign,

and the ordering, all of which require only ordinal information. **That is exactly what survives in the constructed models here: the sign of the ambiguity rate, the existence of a steady-state deficit, and the existence of a stopping point above zero.** Curzi [2026] states the underlying convergence theorem exactly — a contraction with modulus $q < 1$ on a complete space has exactly one fixed point, reached from any start, with distance shrinking by at least q each pass — which is the theorem that licenses the steady state in Section 8.5 and the asymptote in Section 8.6. Curzi [2026ab] supplies the corpus's template for a postulate with a forbidding clause, $E = G \cdot B + \varepsilon$ with $G = 0 \Rightarrow E = 0$ for all B: the form this paper's H2 takes, in which a structural zero cannot be compensated by any amount of the other factor.

10. Verification and rigour audit

The results are decomposed into postulates that can be checked one at a time. Each is atomic, each carries a status, and each chain names what it consumes. **If dropping a postulate leaves the result standing, it is not listed.**

10.1 Dimensional audit

symbol	quantity	dimension	unit	status
ambiguity rate	share of carried values readable as control	$[1]$	dimensionless fraction	derived, 383/100{,}000
$\lvert S\rvert$	size of a set an observer can narrow to	$[\#]$	subscribers, addresses, tuples	defined
$\log_2 \lvert S\rvert$	identity content	$[1]$	**bits** (base-2 logarithm of a count)	derived; dimensionless by construction
anonymity set	subscribers consistent with the record	$[\#]$	subscribers	derived as a ratio of two counts
subscribers, addresses, ports	population sizes	$[\#]$	counts	assumed (constructed)
record size	bytes per session record	$[B]$	bytes	assumed (constructed)
log volume	retention flow	$[B\,T^{-1}]$	bytes per day; TB per year	derived
experts, requirement	maintainer-equivalents	$[\#]$	people	assumed / derived
attrition, replacement	rates on the maintainer stock	$[T^{-1}]$, $[\#\,T^{-1}]$	per year; people per year	assumed
lines remaining	size of the buried estate	$[\#]$	lines of code	assumed / derived
migration fraction	share of remainder moved	$[T^{-1}]$	per year	assumed
unit cost	cost to move one line	$[\text{idx}\,\#^{-1}]$	index units per line	assumed (index, no currency)
keep-running cost	cost to leave one line	$[\text{idx}\,\#^{-1}T^{-1}]$	index units per line-year	assumed (index)

Does the bit accounting balance? Identity in bits is the base-two logarithm of a count, so it is dimensionless, and differences of such logarithms are ratios of counts — which is why "bits unresolved" is a well-formed subtraction: $18.932 - 8.000$ is $\log_2(500000/256)$ and equals 10.932 exactly. **The audit records that the anonymity set is a ratio of two counts and is therefore formally dimensionless while being *read* as a count of subscribers**, which is legitimate only because the numerator and denominator are counts of the same kind of thing.

Where an operational unit is missing, the audit says so. The migration cost model is in **index units with no currency attached**, and the source supplies no conversion. The crossing year of 2032 is therefore a property of the *ratio* of migration cost to keep-running cost and survives; the absolute cost figures do not mean anything until an index unit is defined.

Does the maintainer model balance? $E_{t+1} = (1 - a)E_t + r$ requires a dimensionless per step and r in people per step, which is how they are stated. The fixed point r/a then carries dimension $[\#]$, as a maintainer count must.

10.2 Significant figures and epistemic status

- **0.000% is exact**, not a rounded small number: it is a property of a disjoint union and no measurement is involved. **0.383% is exact for the synthetic space** and is an illustrative model output for every other purpose; it is quoted to three decimal places because 383/100,000 is exact by stipulation, and it should not be read as an estimate of anything.
- **The bit figures are exact to the digits shown**, being base-two logarithms of exact counts: $\log_2 500000 = 18.9316$, $\log_2 256 = 8.0000$, $\log_2 16515072 = 23.9773$, $\log_2 1953.125 = 10.9317$, $\log_2 1453.22 = 10.5051$.
- **The carrier parameters are assumed, not measured.** 500,000 subscribers, a 256-address pool, a 64,512-port usable range, a 150-byte record and 48,000 sessions per subscriber per day are stipulated. Every output inherits that status: **exact arithmetic on assumed inputs is not a measurement, and no uncertainty is quoted because there is no estimate to attach one to.**
- **The session-rate proxy is the weakest input.** 48,000 sessions per subscriber per day — stated in the source as 50 new sessions per busy minute over 16 active hours — is a modelling assumption and not an observed rate. It drives case (iii) and the retention volume, and a different assumption moves those figures proportionally without changing the direction of the result.
- **The maintainer and migration figures are illustrative model output and are rounded hard**: 54.5, 2030, 3.29%, 37.6%, 2032. Quoting 15,023,138 lines to eight figures, as the source does, is false precision on a number produced by assumed parameters, and this paper reports it as 37.6% of the core.

10.3 Limits and boundary behaviour

model	limit	behaviour	reading
ambiguity	control tokens $\to 0$ share of namespace	rate $\to 0$ but stays positive	in-band never reaches zero while the namespaces intersect
ambiguity	planes disjoint	rate $= 0$ exactly	the only way to reach zero is the partition
anonymity	pool size $\to$ subscriber count	set $\to 1$, unresolved $\to 0$	one address per subscriber removes the loss entirely

model	limit	behaviour	reading
anonymity	pool size $\to 1$	set $\to$ all subscribers, unresolved $\to 18.932$ bits	full compression: the log names the carrier
maintainers	replacement $r \to a\times$ requirement	steady state $\to$ requirement	the cliff disappears exactly at $r = 11.0$/yr
maintainers	$a \to 0$	steady state $\to \infty$	no attrition, unbounded growth — a pathology, not an ideal
migration	fraction $m \to 1$	remainder $\to 0$ in one step	completion requires migrating everything at once, which is the rewrite of Section 9.2
migration	cost growth $g \to 0$	annual bill falls at $(1 - m)$	the rising bill is entirely due to $g > m/(1 - m)$
migration	any $m < 1$, finite t	remainder > 0	**the asymptote is the model's core result and it is a boundary statement**

The migration model's stopping point is where the two curves cross and nowhere else. With $(1 + g)^t = 3$, $t = \ln 3 / \ln 1.2 = 6.03$ years from 2026, which is the 2032 in the text. **The crossing depends only on the ratio of the initial migration unit cost to the keep-running cost and on g; it does not depend on the core size at all**, which is why the result is a shape rather than a quantity.

10.4 Invariance

The ambiguity rate is invariant under relabelling of the symbol space and under any bijection of the carried values, which is precisely why it establishes the sign and not the magnitude of anything — the bound carried from Curzi [2026]. **The bit accounting is invariant under any relabelling of subscribers and under a change of logarithm base up to a constant factor**; base two is a choice of unit, not of content. **The migration crossing year is invariant under rescaling of the core size and of the cost index**, since it depends only on ratios. **The maintainer steady state is invariant under rescaling of the requirement only if replacement is rescaled with it**, which it is not in the model — that asymmetry is the model's content, not a defect.

Tested at how many scales? **The invariances are algebraic and were checked at exactly the one parameter set reported.** No empirical scale-freedom is claimed anywhere in this paper.

10.5 Analogy discipline

borrowing	source	category	justification
identity as $\log_2$ of a set size	Hartley [1928], Shannon [1948]	**1 — derivation**	the anonymity set is a finite equiprobable selection, which is exactly Hartley's premise; this is **Shannon entropy in bits**, an information quantity, not thermodynamic entropy
entropy of a uniform distribution	Cover and Thomas [2006]	**1 — derivation**	the uniform case is the identity used
identity impersonation	Douceur [2002]	**2 — shared form**	the Sybil result is proved for an identity system without a certifying authority; signalling is not shown to satisfy those premises, and the result is used for its structure
independent-party bound	Lamport et al. [1982]	**2 — shared form**	the $3f + 1$ bound is a theorem about agreement, cited as a prior formalisation of "how many parties are genuinely independent", not applied to signalling
impossibility under asynchrony; the availability/consistency trade	Fischer et al. [1985], Gilbert and Lynch [2002]	**3 — framing**	cited to establish that structural impossibility results exist and are the right genre of claim, never as results about this system
coupled-network fragility	Buldyrev et al. [2010]	**2 — shared form**	the substrate dependency is not shown to satisfy the percolation premises

borrowing	source	category	justification
robust-yet-fragile; percolation	Albert et al. [2000], Callaway et al. [2000], Cohen et al. [2000]	**2 — shared form**	the shape of the tolerance asymmetry is used; the degree-distribution mechanism is not claimed, and Doyle et al. [2005] is the reason it should not be
covariance, not count	Markowitz [1952]	**1 — derivation**	the substrate exposure genuinely is a common factor across nominally independent legs
open design; complete mediation	Kerckhoffs [1883], Saltzer and Schroeder [1975]	**1 — derivation**	the disclosure result *is* Kerckhoffs's principle; the premise — that effectiveness depends on the partition and not on knowledge — is stated and satisfied
tight coupling; redundancy that reduces safety	Perrow [1984], Sagan [1993]	**3 — framing**	qualitative organisational theory; cited as framing and as a warning, never as formal results
queueing at saturation	Kingman [1962], Hopp and Spearman [2011]	**2 — shared form**	the maintainer pool is not modelled as a queue here; the saturation shape is invoked for intuition only
threshold distribution governs the aggregate	Granovetter [1978]	**2 — shared form**	the maintainer model tracks a mean, not a threshold distribution; the result is invoked as a caveat on the model, not as support for it

borrowing	source	category	justification
performativity of a published model	MacKenzie and Millo [2003], Hacking [1995]	**1 — derivation**	the mechanism is the same one: a description relied upon changes the behaviour of those relying on it, and this is documented with dates and prices
the programme itself	Castellano et al. [2009]	—	cited as the survey that **bounds** what statistical physics applied to social dynamics has established; this paper claims no result within that programme, only its vocabulary

The failure mode this table exists to prevent is letting a category-2 or category-3 borrowing carry a category-1 result's authority. The disclosure result and the bit accounting are category 1 and need no help; the network results are category 2 and are used for shape.

10.6 The postulates

#	statement	status
P1	a system has a control plane and a data plane, as defined in Section 3.1	DEFINITION
P2	in-band means some data value can be interpreted as a control instruction	DEFINITION
P3	out-of-band means the two planes are disjoint, so no data value can be	DEFINITION
P4	identity content is $\log_2$ of the size of the set an observer can narrow to	DEFINITION, CITED `hartley1928transmission`
P5	the effectiveness of a partition is a function of the partition and not of any observer's knowledge of it	DEFINITION
P6	separation places no constraint on who may reach a plane	DEFINITION
P7	the carrier parameters of Section 8.1	POSTULATED (constructed)

#	statement	status
P8	the maintainer and migration parameters of Sections 8.5 and 8.6	POSTULATED (constructed)

10.7 The chain that establishes the disclosure constraint

step	consumes	statement	status
L1	P2	in-band, the ambiguity rate is strictly positive: 383/100,000 = 0.383%	ARITHMETIC
L2	P3	out-of-band, the ambiguity rate is 0/100,000 = 0.000%, exactly	ARITHMETIC
L3	L1, L2	the impersonation class is closed by the partition, not by care within the channel	ARITHMETIC
L4	P5, L2	disclosing the partition leaves the rate at 0.000%, because knowledge is not one of its arguments	ARITHMETIC
R1	L4	a design whose effectiveness depends on concealment fails when the design leaks; a design whose effectiveness is structural does not — Kerckhoffs [1883], Saltzer and Schroeder [1975]	CITED
R2	L3, R1	every element of a channel-separated design is publishable, and non-disclosure is refused as a design premise	DERIVED
L5	P6	separation constrains no party's access to the control plane	DEFINITION
R3	L5, R2	access control on the separated control plane is a second and necessary control — Saltzer and Schroeder [1975] on complete mediation	CITED

Six checkable steps and two citations turn a statement that reads as an ethical preference into an engineering standard. Only P7 and P8 are POSTULATED, and neither enters this chain: **the disclosure constraint rests on definitions and arithmetic alone.**

10.8 The chain that establishes the provenance loss

step	consumes	statement	status
A1	P4	identity of one subscriber among 500,000 is $\log_2 500000 = 18.932$ bits	ARITHMETIC
A2	P4, P7	an address-only observation is one of 256, or 8.000 bits	ARITHMETIC
A3	P7	$256 \times 64512 = 16{,}515{,}072$ simultaneous bindings, or 23.977 bits	ARITHMETIC
A4	P7	$500000/256 = 1953.125$ subscribers share one public address	ARITHMETIC
A5	A1, A4	address-only leaves $\log_2 1953.125 = 10.932$ bits unresolved	ARITHMETIC
A6	P7	500,000 subscribers at 48,000 sessions a day is 24 billion sessions a day	ARITHMETIC
A7	A3, A6	24 billion sessions over 16,515,072 slots is 1,453.2 reuses of each slot a day	ARITHMETIC
A8	A7	after state expiry, $\log_2 1453.2 = 10.505$ bits remain unresolved	ARITHMETIC
A9	A6, P7	24 billion records at 150 bytes is 3.6 TB a day, or 1,314 TB a year	ARITHMETIC
A10	A9	the mapping is discarded because retention is a standing expense that buys the carrier nothing operationally	POSTULATED (economic model)

Every step but A10 is arithmetic on stated counts. A10 is the one genuine assumption, it is an assumption about incentives rather than about mechanism, and it is the step a reader should attack first.

10.9 The chains that establish the two substrate equilibria

step	statement	status
B1	$E_{t+1} = (1-a)E_t + r$ is a first-order linear difference equation — Elaydi [2005]	DEFINITION, CITED
B2	its fixed point is $r/a = 6/0.11 = 54.5$	ARITHMETIC
B3	$54.5 < 100$, so the deficit is a steady state and not a transient	ARITHMETIC
B4	holding the line requires $r = a \times 100 = 11.0$ a year against 6 actual	ARITHMETIC
C1	$R_{t+1} = (1-m)R_t$ with $m = 0.15$	DEFINITION
C2	$R_t = R_0(0.85)^t > 0$ for all finite t	ARITHMETIC
C3	annual bill ratio is $(1-m)(1+g) = 0.85 \times 1.20 = 1.02 > 1$	ARITHMETIC
C4	unit cost reaches the keep-running cost when $1.2^t = 3$, at $t = 6.03$	ARITHMETIC
C5	therefore a rational stopping point exists strictly above zero	ARITHMETIC

10.10 Discrepancies found in the source, printed

One arithmetic inconsistency was found and is reported rather than harmonised. The tabulated remainder at 2046 is 1,550,381 lines, which is $40{,}000{,}000 \times 0.85^{20} = 3.876\%$ of the core — twenty annual migrations counted from 2026. The separately stated residual "after twenty years" is 1,317,824 lines, or 3.29%, which is $40{,}000{,}000 \times 0.85^{21}$ and therefore twenty-one annual migrations. **The two statements are off by one migration year.** This paper reports the source's 3.29% figure as stated, marks it in Figure 4 as the *stated* residual, and records here that the tabulated series gives 3.88% at the same date. **Neither figure changes the result, which is that the remainder is strictly positive for all finite time.**

A second, smaller correction: the source dates the article establishing that a system's interfaces mirror its builders' communication structure to 1967; it appeared in April 1968, and Conway [1968] is cited with the corrected date.

11. Limitations

Every worked result is a constructed model, and none is a measurement. The carrier, the symbol space, the maintainer shop and the migration are all invented, with parameters chosen to be plausible rather than observed. **No empirical dataset was collected and no real network, subscriber, traffic record, institution or codebase appears anywhere in this paper.**

The ambiguity rate carries no information beyond its sign. It is invariant under relabelling of the symbol space [Curzi, 2026] and would take a different value under any other allocation of reserved symbols. **The only part of it that is a finding is that in-band is greater than zero and out-of-band is exactly zero**, and that part follows from the partition rather than from the parameters.

The isomorphism table is a claim of shared form, not a derivation. Each row depends on the assertion that the named system has the separation property attributed to it, and that

must be established per system. **It is not established here for any of the four.**

The session-rate proxy is stipulated and drives two of the reported figures. 48,000 sessions per subscriber per day is a modelling assumption. It should be replaced with a measured rate before any figure derived from it is used outside this paper.

Only a single translation boundary is modelled. A real path may cross several — a subscriber translator, then a carrier translator, then a protocol translator — and each crossing compounds the loss, **so the single-boundary figures are lower bounds on the loss along a multi-hop path, not the whole of it.** The compounding is stated as a direction and is not measured.

Address provenance is not identity. The paper measures the network layer only and says nothing about application-layer identifiers — an authenticated account, a signed session — which can attribute where the network layer cannot and are the actual basis of attribution in most well-run systems. **The network address was never meant to be the identity, and the remedy at the application layer is a genuinely better answer where it is available.**

The standards this paper relies on are not in its bibliography. The numbering plan, the two internet protocol versions, the transition and translation mechanisms, the carrier-grade translation terminology and the shared-address-space allocation are public standards documents, described here from the source manuscript. **No standards document is cited, because none is in `bib/external.bib`.** Every technical claim about them requires verification against the governing document, including the current status of each, the obsoletes relationships among them, and the deprecation status of the older tunnelling mechanisms. This is a real gap and it is not closed here.

The exhaustion date and the public 2020 episode are described, not cited. The central registry's exhaustion of its free pool in 2011, the staggered later exhaustion of the regional registries, the 2020 unemployment-system failures and the April 2020 appeal for volunteer programmers are all described from the source manuscript. **The bibliography contains no contemporaneous reporting and no press sources, so none of these is cited**, and each requires verification against contemporaneous sources before use.

The scale figures for legacy transaction-processing code are neither cited nor asserted. They are attributed in the source to 2017 reporting that is not in the bibliography. They are used here only for order of magnitude and **must not be restated as a current measurement**.

The 2032 stopping point and the 2030 cliff are properties of assumed rates. They are exact functions of stipulated parameters, they carry no uncertainty because there is no estimate in them, and they carry no authority beyond the assumptions.

11.1 Rigour checks not met, listed by name

1. **No as-of date and no source, because nothing was measured.** Every figure derives from a constructed parameter set. Any application of these models to a real carrier, shop or codebase requires both.
2. **No uncertainty is reported anywhere**, because nothing is estimated. This is the absence of empirical content rather than precision.
3. **Model comparison was not performed.** No alternative model was fitted, no out-of-sample test was run, no information criterion was computed and no matched null was constructed. The three models are illustrations of inequalities, not fits, and their conclusions follow from the inequalities rather than from goodness of fit.
4. **The cost models carry no operational unit.** Index units with no currency attached; only ratios survive.
5. **The isomorphism is unverified per system**, as above.
6. **Two borrowings are category 3 (framing) and are used only as framing**: Perrow [1984] and Fischer et al. [1985]. Neither is a result about this system and neither is treated as one.
7. **One source arithmetic inconsistency was found** and is printed in Section 10.10 rather than harmonised.

8. **A10 is the chain's only real assumption** and it is about incentives rather than mechanism. A jurisdiction that mandates retention breaks it, and the model says nothing about how often that is the case.
9. **No distributional or scaling claim is made, and the Clauset–Shalizi–Newman procedure was therefore not run** [Clauset et al., 2009]. Any restatement of this paper that reaches for a heavy-tail or scale-free claim about sharing ratios, session distributions or the network results borrowed in Section 9.1 must run it first, and must read Broido and Clauset [2019] on how weakly social networks support such structure.
10. **The three models fail the calibration-and-validation checklist** of Bruch and Atwell [2015]: none is calibrated against observed data, none is validated out-of-sample, and none has a sensitivity analysis beyond the limits reported in Section 10.3. They are offered as demonstrations that a stated mechanism suffices to produce a stated shape, which is one of the legitimate uses Epstein [2008] enumerates and is not prediction.
11. **The maintainer model tracks a mean rather than a distribution of thresholds** [Granovetter, 1978], and therefore gives the optimistic answer.
12. **This paper does not invoke the framework's reflexive law and therefore does not inherit the reflection problem** [Manski, 1993]. No result here rests on an endogenous social effect estimated from observation, and no such effect is claimed. **Where downstream work couples these results to the reflexive law, that caveat applies there and is not answered here**: a linear-in-means reflexive specification is not identified from observational data, and the remedies — randomisation, an instrument, partial-population variation, or stipulating and defending the reference group — all carry a price. The disclosure result's own reflexive content, that a published architecture changes the behaviour of those relying on it, is asserted on the strength of a documented case [MacKenzie and Millo, 2003] and not estimated from data here.

12. Conclusion

This paper establishes the channel-separation principle and its disclosure constraint in the form the rest of the corpus cites, and the following statements are the fixed text of those results.

R1 — the channel-separation principle. In-band architectures admit an impersonation class that no care within the shared channel can close, and out-of-band separation closes it structurally — measured at 0.383% in-band ambiguity against 0.000% out-of-band over 100,000 carried values.

R2 — the disclosure constraint. Channel separation is strengthened by disclosure because its effectiveness is entirely independent of whether anyone knows it is in place, so every element must be publishable; **non-disclosure is refused as a design premise, not merely discouraged.** This is Kerckhoffs [1883]'s principle, stated in 1883, and it is what makes the constraint an engineering standard rather than a preference.

R3 — provenance loss at the carrier boundary. 18.932 bits of identity entering becomes 8.000 bits observed from an address alone, and 23.977 bits when address and port are both captured.

R4 — the attribution-set result. With an address-only log the anonymity set is 1,953.1 subscribers; with address-plus-port and live session state it is 1; and after state expiry it is 1,453.2. **An address is evidence of a pool, not of a person.**

R5 — the constructed maintainer cliff. 11% annual attrition against a flat requirement of 100 maintainer-equivalents needs 11.0 replacements a year against 6 actual hires — a deficit of 5.0 a year behind an apparent 20-maintainer surplus, landing at a steady state of 54.5 and crossing below the requirement in 2030.

R6 — the constructed migration asymptote. Migrating 15% of the remaining 40-million-line core each year at a per-line cost rising 20% a year leaves 3.29% after twenty years, with total cost rising about 2% a year and a rational stopping point near 2032 at 37.6% still buried.

Counter-result. Separation is necessary and not sufficient — **a separated control plane trusted because it is separate is the failure mode** — and two attractive simplifications are refused outright: that the internet depends on copper, when the durable dependency is logical topology rather than electrical continuity; and that a rewrite resolves the legacy layer, since a rewrite trades a known correct system for an unknown one over exactly the undocumented behaviour that matters, at the moment of least remaining expertise.

What the paper adds to the corpus's instrument. A concentration measurement on one axis can hide a common-mode dependency on another. The effective-count instrument of Curzi [2026] counts perimeters and cannot see whether two independent perimeters rest on the same substrate generation. **Substrate concentration is a real exposure alongside perimeter concentration, and it must be measured separately or it will not be measured at all.**

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every worked result in this paper — the symbol-space demonstration, the constructed carrier, the maintainer model and the migration model — is **CONSTRUCTED** or **SYNTHETIC** and is labelled as such at each point of use. The numerical results are exact computations on those stipulated inputs. **No real network, subscriber, traffic record, institution, product or codebase appears anywhere in this paper, and no allegation is made about any of them.**

All figures were produced by the scripts named in their captions (`figs/src/CCP-066_fig01.py` through `figs/src/CCP-066_fig05.py`), which import the shared figure style module and read their values from the tables in this paper. No figure plots a number that does not appear in the text.

Technical standards, the address-exhaustion date, the 2020 public episode and the legacy scale figures are described from the source manuscript and are not cited to the bibliography, which contains no standards documents and no press reporting. Section 11 lists each affected claim. Any reader relying on them must verify them against the governing documents and contemporaneous sources.

The author's own systems project is recorded here as the author's stated position and not as an examined artefact. The author states that the class of problem described in Sections 8.5 to 8.7 is at bottom a *representation* problem — the buried systems speak one language and the modern tooling another — and that the tractable move is not replacement but a representation-adapter layer letting modern tooling read and drive the buried systems in place. The author further states that a systems project exists to instance that direction, together with canonicalised settlement and the channel-separation principle built into an operating system's own architecture, and that it was **deliberately left half-finished**. The repository is `https://github.com/sykomiked/ZXV0`. **No representation is made about the contents, licence, security or functional status of that repository; it has not been examined for this paper, and nothing in this paper's argument depends on the project's completion.** The half-finished status is recorded as the author's stated position rather than as a defect: an unfinished implementation of a correctly diagnosed class of solution is a different thing from a finished implementation of the wrong one.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest, other than that the systems project described in the availability statement above is the author's own; it is recorded as a stated position and is not relied on for any result.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

No operational content

This paper contains no operational content of any kind: no tone, sequence, code, method, or anything that could be used to inject signalling into, or attack, any network. The historical material is recounted only to establish that in-band control was structurally reachable. Present-day signalling vulnerabilities are described only at the level of principle, and they are a matter of illicit access to control networks that nothing here describes or could enable. **The principle stated here is for building systems, not for breaking them.**

Legal, tax and investment advice

This paper is analytical and is not legal, tax, or investment advice. It describes architectural principles, addressing systems and software substrates. No recommendation is made in respect of any particular person, institution, product or jurisdiction, and no allegation is made about any of them. Data-retention obligations differ by jurisdiction and change; any retention window designed on the basis of Section 8.3 requires confirmation with qualified counsel.

Appendix A. Claims requiring verification against primary sources

The following are described in this paper from the source manuscript and **are not cited to the bibliography, which contains no standards documents, no press reporting and no jurisdiction-specific material.** Each requires verification against the governing document before any reliance.

1. The characterisation of the fifth-numbered international signalling system as in-band, with control tones in the voice band, and of the seventh as out-of-band common-channel signalling adopted across the 1980s. Confirm the exact recommendation numbers and the adoption and deployment dates; this paper cites the systems by their common designations only and asserts no recommendation number.
2. The international public telephone numbering plan and its mobile-subscriber companion: confirm the in-force edition, the exact title, and the fifteen-digit maximum.
3. The two internet protocol version specifications, the transition mechanisms, the stateful translation and name-synthesis specifications, the well-known translation prefix, and the edge-plus-core translation specification. Confirm each designation, title and current status, and confirm which document currently governs the stateless translation algorithm, since the original has been obsoleted.
4. The tunnelling mechanisms named in passing. Confirm the governing document and current deprecation status of each; several early transition mechanisms have been deprecated.

5. Carrier-grade translation terminology, the reserved shared-address-space allocation, and per-session logging requirements, all described generically here. Confirm the relevant specifications and any applicable data-retention obligation in each jurisdiction of operation before designing a retention window.
6. The statement that the central registry exhausted its free pool of older-protocol blocks in
   2011. Confirm the exact date of the final allocation and the distinct, later exhaustion dates of each regional registry, which differ by years.
7. The proportions of the copper last-mile installed base, which vary widely by country and are changing as fibre reaches the premises. **The load-bearing correction — logical, not electrical, substrate — does not depend on any specific proportion and should be preserved in any restatement.**
8. The 2020 unemployment-system failures and the April 2020 appeal for volunteer programmers: confirm the exact date and wording of the appeal, which other jurisdictions reported the same difficulty, and the specific systems and languages involved. The general phenomenon is well documented; the specifics require verification.
9. The legacy scale figures — the share of card-swipe transactions, the share of banking systems, the line count and the daily commerce volume. **These are attributed in the source to 2017 reporting, are not asserted as current fact, and are not cited here because the originating article is not in the bibliography.** Confirm the article, its date, and the primary source of each figure before quoting any of them, and carry the attribution and the date with them.
10. The characterisations of the two legacy language families and of mainframe transaction monitors are generic. No claim is made about any named institution's use of either, and any such claim requires independent establishment.

All numerical results in this paper were computed from the stated constructed inputs at full precision and rounded only for printing.

Chapter 67

A Concentration Atlas for Distributed Ledgers: Validator Herfindahls, the Utility Ordinal, and Dwell Time as the Freeze Variable

Abstract. We apply the corpus's concentration instrument to nine public ledgers without modification, since an inverse-HHI effective count is indifferent to what the shares describe. Measured on stated validator, council or pool sets, the effective counts range from a curated 35.0 validators (HHI 0.028571, an explicit ceiling rather than an observation) and 39.0 council members (0.025641) down to 6.00 organisations (0.166667) and 5.94 entities (0.168400). Bitcoin mining measures HHI 0.1900 and 5.26 effective pools with a 30 per cent leader and a 56.0 per cent top two, against the ceiling 1/s_max^2 = 11.11 that the leading share alone imposes. A second axis, utility-as-function, is defined narrowly and badged as an ordinal rather than a score, with Bitcoin at 1/9 and a general virtual machine at 9/9; the paper requires the axis to be printed with the rank every time and denies that programmability is a quality gradient, since a minimal ledger cannot express a bug it has no opcode for. Base-layer throughput is stated to its parameters, about 7.4 to 11.9 transactions per second. The store-versus-rail question is settled numerically by dwell: a five-second hop is 1.585e-07 of a year while value rested for thirty days is 0.0822, a ratio of roughly 518,400 to one, so the freeze property that disqualifies a rail as a store is nearly irrelevant to it as transit.

Keywords: validator concentration; curated trust sets; utility ordinal; proof-of-work pools; dwell time; freeze property; central bank digital currency; pilot versus production

Series: D. Structure and Resilience **Paper:** CCP-067

1. Introduction

A concentration index is indifferent to what its shares are shares of. That is a property of the arithmetic and not a claim about the world, and it means an instrument built to count independent jurisdictions counts independent validators, independent council members, independent mining pools or independent curators without modification. This paper applies the corpus's concentration instrument to nine public settlement ledgers, states exactly what the resulting numbers mean, and states more carefully what they do not.

The paper has five separable contributions.

The instrument transfers unchanged. Measured on stated validator, council or pool sets, the effective counts run from a curated 35.0 validators (HHI $= 0.028571$, an explicit design ceiling rather than an observation) and 39.0 council members (0.025641) down to 6.00 organisations (0.166667) and 5.94 entities (0.168400). Nothing in the derivation of the index cares that its original application was to jurisdictions, and no new theory is required to read a ledger.

The curated-list collapse is the sharpest result the transfer produces. A trust set admitted by a single curating party has HHI $= 1.0000$ and **one** effective curator, against five equal independent operators at 0.2000 and 5.00. **The concentration is not in who runs the infrastructure; it is in who decides which infrastructure the network defaults to trusting**, and a network reporting the first while the second binds has made the same error Curzi [2026] measures on a payment rail, where a naive correspondent index of 0.302400 — 3.306878 effective correspondents — collapses to 0.745256 and $N_{\text{rail}} = 1.341819$ once flows are decomposed by currency.

Proof-of-work concentration is measured on pools, not on machines. A constructed pool distribution gives HHI $= 0.1900$ and 5.26 effective pools at a $30\%/26\%$ top two, against the ceiling $1/s_{\max}^2 = 1/0.30^2 = 11.11$ that the leading share alone imposes.

A second axis is defined and then deliberately demoted. Utility-as-function — what settlement logic a rail can express and what it settles — is stated as an **ordinal**, not a score, running from $1/9$ for a minimal ledger to $9/9$ for a general virtual machine. **The axis is not oriented: a 9 is more expressive than a 1 and is not better than a 1**, because a minimal ledger cannot express a bug it has no opcode for. The rank must be printed with the axis named beside it every time, and no arithmetic on the differences is licensed.

Dwell time settles the store-versus-rail question numerically. A five-second hop is 1.585×10^{-7} of a year (0.0000159%); value rested for thirty days is 0.0822 of a year (8.2192%); the ratio is $2{,}592{,}000/5 = 518{,}400$. **The freeze property that disqualifies a rail as a store is almost irrelevant to it as transit, and the entire difference is dwell.**

Why these sit in one paper. They are one instrument at four scopes and one refusal repeated. The instrument is Curzi [2026]'s: the sum of squared shares, its reciprocal read as an effective count, and a bloc collapse applied before the squaring. What the paper adds is that each consensus mechanism *names its own concentration axis* — proof-of-work names pools, bonded-stake systems name staking operators, a fixed producer set names itself, federated agreement names the curated trust list, council governance names the council — and that once the axis is named the measurement is arithmetic and nothing new. The refusal repeated is that none of these numbers is a fact about a running network unless it carries a source and a date.

The counter-results are stated first because they are what makes the rest usable. The utility ordinal is not a quality scale. The freeze control on a settlement instrument sits with its issuer and follows the instrument across rails, so diversifying rails buys nothing against it. Participating in a pilot is not settling production value. **And every live concentration figure is time-varying: a concentration measure without an as-of date is not a measurement**, so this paper measures constructed and stipulated distributions only and defers every live number to a dated register.

2. Hypotheses

Six hypotheses with their falsifiers, and three statements that are not hypotheses and are labelled as such. Platt [1964]'s standard applies: a hypothesis earns its place by naming what would kill it.

H1 — the transfer hypothesis. The concentration instrument applies to a distributed ledger's participant set without modification, because its derivation does not depend on what the shares are shares of. *Predicted:* an effective count computed on validator, council or pool shares behaves exactly as one computed on jurisdictional shares — returning n for n equal

shares, bounded above by $1/s_{\max}^2$, and falling under a bloc collapse. *Falsified by:* a share basis on a ledger for which the numbers-equivalent property fails, or a case in which the bloc collapse raises rather than lowers the count.

H2 — the dwell hypothesis. Freeze exposure on a settlement rail is proportional to dwell time and to nothing else, so the store-versus-rail distinction is a quantitative rather than a categorical one. *Predicted:* at seconds-scale finality a hop is 1.585×10^{-7} of a year against 0.0822 for a thirty-day rest, a ratio of $518{,}400$. *Falsified by:* an exposure that does not scale with dwell — for example a freeze whose probability depends on the transacting party rather than on the time the value is resident.

H3 — the pool hypothesis. On a proof-of-work network the coordinating unit is the pool and not the machine, so a count of participants overstates independence. *Predicted:* a constructed ten-pool distribution with a 30% leader gives HHI $= 0.1900$ and 5.26 effective pools against ten nominal, and the leading share alone caps the effective count at 11.11. *Falsified by:* a demonstration that pool membership imposes no coordination — that a pool cannot direct the work of its members — in which case the machine count would be the right basis.

H4 — the ordinality hypothesis. Utility-as-function admits an order and no metric, so the differences between its ranks carry no information. *Predicted:* any order-preserving relabelling of the nine positions leaves every claim this paper makes about the axis intact, and no claim in this paper survives being restated as a difference or a ratio of ranks. *Falsified by:* an operation on the axis whose result is invariant under order-preserving relabelling and which is not derivable from the order alone — which would show the axis carries more than ordinal information.

H5 — the curation hypothesis. Where a trust set is admitted by one party, the effective count of independent admission decisions is one, however many operators the set contains. *Predicted:* a single-curator set gives HHI $= 1.0000$ and 1.00 against five equal independent operators at 0.2000 and 5.00. *Falsified by:* a curated set in which removing or reweighting an entry requires the assent of more than one party, which would make the curation axis genuinely plural.

H6 — the throughput hypothesis. Base-layer throughput on a minimal ledger is fixed by two protocol parameters and a transaction size, so it is computable and is not a matter of opinion. *Predicted:* at approximately 226 virtual bytes per transaction against approximately 10^6 virtual bytes per block on a ten-minute interval, approximately 7.4 transactions per second; at approximately 140 virtual bytes, approximately 11.9. *Falsified by:* an observed sustained base-layer rate outside that band at those parameters.

Not hypotheses, and labelled as such.

- *The bloc assignment of any network* is a **modelling choice**. Which parties coordinate is stipulated, not estimated; Section 3.3 states what that costs.
- *The nine-position utility ordering* is a **definition** fixed by this paper's own construction of the axis. It carries no empirical weight and is not offered as evidence for anything.
- *The three-axis frame* is a **classification scheme**, not a finding. Its value is that it forces the axis to be named beside the rank; it predicts nothing.

3. Method

3.1 Which effective count, stated exactly

Every effective count in this paper is $N = \text{HHI}^{-1}$, the reciprocal of the sum of squared shares. That quantity is the Hill number of order two [Hill, 1973, Jost, 2006]: the number of equally-sized units that would produce the observed index. It is the reciprocal of Simpson [1949]'s index, the effective number of parties of Laakso and Taagepera [1979], and the numbers-equivalent reading attributed to Adelman [1969]; the index is Herfindahl [1950] and Hirschman [1964], and the bands at which competition authorities read it are policy choices

that have moved [U.S. Department of Justice and Federal Trade Commission, 2023, 2010]. The exponential of Shannon entropy [Shannon, 1948, Cover and Thomas, 2006] is the same family at order one, and HHI^{-1} is the exponential of the Rényi entropy of order two [Rényi, 1961], which is where the participation-ratio name used by Curzi [2026] comes from.

The corpus's other effective count is not a Hill number and is not used here. The affine form $N_{\text{eff}} = 1 + (N-1)u$ of Curzi [2026] is a linear interpolation between 1 and N; it is not a numbers-equivalent of any index, and no literature in ecology, economics, political science or information theory contains it. Curzi [2026] runs the same verification chain on both and lets the affine one fail visibly. **Wherever this paper says "effective count" it means HHI^{-1}, and where the affine form is referred to it is named and marked as the corpus's own construction.**

Two standing bounds are used throughout and are stated here so they need not be re-derived: $s_{\max}^2 \le \mathrm{HHI} \le s_{\max}$, so $N \le 1/s_{\max}^2$ — **the largest single share alone caps the effective count** — and $N \le n$ for n units, with equality only at equal shares.

3.2 Each consensus mechanism names its own concentration axis

The instrument requires a share basis, and choosing it wrongly is the whole of the error this paper is written against. The share basis is not a free choice: it is dictated by the mechanism, because the mechanism defines what a coordinating unit is.

mechanism family	what the right to append is tied to	coordinating unit, hence share basis
proof-of-work	expended computation	the **pool** — many machines directed by one operator's template
bonded-stake	committed capital	the **staking operator, custodian or pooled-stake protocol**
fixed producer set	election to a capped set of producers	the **producer set itself**, which is enumerable
capped Byzantine-fault-tolerant set	stake within a capped active set	the **active set**, weighted by stake, capped by governance
federated agreement	overlapping declared trust sets	the **curator of the default trust set**
council governance	membership of a governing body	the **council**, capped by constitution

Read the table as one sentence repeated six times: name the unit that can be moved by a single decision, then square its shares. A count of machines on a pooled network, of individual depositors on a bonded-stake network, or of node operators on a network whose trust set is curated by one party is a count of participants and not a measurement of failure independence.

3.3 The bloc collapse, and the identification problem it carries

The collapse is a variance identity and not a grouping convenience. Following Markowitz [1952], let w_i be the share held by participant i and X_i a unit-variance shock; the aggregate shock has $\mathrm{Var}(R) = \sum_i \sum_j w_i w_j \rho_{ij}$. Under independence this is $\sum_i w_i^2$, the naive index; under perfect correlation inside groups and independence across them it is $\sum_b W_b^2$, the bloc index. **The two are one variance under two correlation assumptions, and the reciprocal is the effective number of independent shocks.** The general form $1/(w'Rw)$ is the similarity-sensitive diversity of order two [Leinster and Cobbold, 2012], whose additive relative is Rao [1982]'s

quadratic entropy. Curzi [2026] establishes the identity and verifies both endpoints; this paper reuses it.

That correlations rise under stress is measured [Longin and Solnik, 2001, Ang and Chen, 2002], so **the effective count falls exactly when it is needed**; the ecological name for the same variable is synchrony [Loreau and de Mazancourt, 2008] and the ecological demonstration that effective count buys stability is Schindler et al. [2010]'s. The contrary literature is read alongside: Doak et al. [1998] show that diversity–stability correlations can arise from statistical averaging with no mechanism, and May [1972] that increasing diversity destabilises randomly assembled systems. **Portfolio arithmetic is not by itself evidence of a mechanism**, and none is claimed here.

The identification problem must be stated where the collapse is introduced. A bloc assignment is a reference-group specification, and the framework's reflexive law — participants influence the aggregate while the aggregate constitutes them — is a linear-in-means reflexive specification whose endogenous effect Manski [1993] proves is **not identified from observational data**. Shalizi and Thomas [2011] reinforce the point: influence and selection are generically confounded in observational network data, which is exactly the data a validator set offers. **The remedy adopted here is stipulation with a defence**, on the template of Brock and Durlauf [2001]; randomisation, an instrument and partial-population variation are all unavailable. **The price is that every bloc-level statement in this paper is a modelling position and not an identified empirical result**, and Section 14 records it.

The instrument is also exposed to its own publication. Goodhart [1975] and Campbell [1979] state what happens to a measure adopted as a target; MacKenzie and Millo [2003] documents, with dates and prices, a model that made its own assumptions true after adoption; Lucas [1976] states the consequence for parameters fixed before publication. **A published decentralisation metric is a target, and nothing in the arithmetic prevents a network from optimising the metric rather than the property.** The only defence offered is the reporting standard of Appendix A.

3.4 The utility ordinal: its scale type, and what that forbids

The second axis is defined narrowly because it is the one most easily corrupted. **Utility-as-function is what settlement logic a rail can express and what it settles.** A network with a general virtual machine can express arbitrary settlement logic; a network with a fixed instruction set and a handful of native operations can express those and no more; a network with a constrained script can express value transfer and little else. That is a genuine ordering.

It is an ordinal scale in Stevens [1946]'s sense, and this paper states the consequence rather than leaving it to be discovered. An ordinal scale is invariant under any order-preserving transformation, so **means, differences and ratios of its values are not licensed operations**. Position 9 is not three times position 3; the distance between 1 and 3 is not comparable to the distance between 7 and 9; and an average of positions across networks is not a quantity. Krantz et al. [1971] sets out what must be true before a quantity can be said to exist and Narens [1996] what a scale type licenses; Luce [2002] is the reason the question is not settled by convention.

The axis is also not oriented. Higher is more expressive, not better. A general virtual machine can express any settlement logic and therefore any settlement error; a minimal ledger cannot express a bug it has no opcode for. **Minimalism at the bottom of the ordinal is a design achievement where the design goal is conservatism, and reading the ordinal as a quality gradient inverts the engineering.** This paper therefore requires the axis to be printed with the rank every time: the number without the axis is not a measurement.

The scalar-reduction habit this axis instantiates is a framework-level limitation and is inherited rather than restated. Curzi [2026] names it once, at framework level, as one objection appearing in two domains rather than two objections: where the source description is a pattern across several jointly varying features, the framework replaces it with a single number and does arithmetic on the number, and the arithmetic's validity does not protect the reduction that produced it. **That is exactly what a nine-position utility rank does to a design space**

with at least four independent dimensions — expressiveness, throughput, finality semantics and governance surface — and this paper concedes the objection in that form and points at its statement rather than repeating it. Michell [1997] is the general charge, that quantitative structure is assumed where it was never tested, and it is the strongest single objection to this framework; Luce and Tukey [1964] is the conjoint-measurement test that would settle it and **it was not run**.

Curzi [2026m] reaches the structurally parallel result in another domain: because its moment equation contains a division and a multiplication, ordinal inputs demote it to a heuristic ordering that may not be printed as a fraction or computed. Curzi [2026t] states what survives such a demotion — the limiting cases, the fixed point, the sign, and the ordering, all of which need only ordinal information — and this paper presupposes that result rather than rederiving it. **What survives here is the order and the fact that the axis is unoriented. Nothing else does.**

3.5 Dwell time as the freeze variable

The corpus's standing rule is that a rail is not a holding, and that freeze exposure on a transit asset is an exposure-window problem rather than a custody problem. This paper makes the rule quantitative. Let λ be a per-unit-time hazard of interruption and w the dwell — the time value is resident on the rail. Expected exposure is λw to first order, so **the entire difference between a rail and a store is** w, and the store-versus-rail question is quantitative rather than categorical.

The freeze control itself is a separate object from the rail, and conflating them is the counter-result of Section 11.7. Where a settlement instrument's issuer can immobilise a balance, that control follows the instrument across every rail that carries it, so **diversifying rails does nothing whatever to it**. The disclosure literature is the right frame for the transparency side: a fully public transfer record is the opposite of concealment, and a structure built to be publishable is indifferent to it by construction — which is Kerckhoffs [1883]'s principle, restated as *open design* by Saltzer and Schroeder [1975] and adopted by Curzi [2026] as a hard design premise: **channel separation is strengthened by disclosure because its effectiveness is entirely independent of whether anyone knows it is in place, so every element must be publishable and non-disclosure is refused as a design premise rather than merely discouraged.** Brandeis [1914] is the older policy instinct, and the automatic-exchange and anti-money-laundering regimes [Organisation for Economic Co-operation and Development, 2017, Financial Action Task Force, 2012] have since removed non-disclosure from the set of available design parameters in any case.

3.6 What is described, what is constructed, and what is cited

Every distribution measured in this paper is either a protocol constant, an equal-weight ceiling derived from one, or a CONSTRUCTED illustration. No live validator set, stake distribution, vote distribution, hashrate distribution or council roster is measured, because all of them move — some weekly — and **a concentration measure without a source and an as-of date is not a measurement.** Every such figure is deferred to the dated register of Appendix B.

Consensus-protocol specifications are not in this paper's bibliography. The whitepapers, consensus papers, improvement proposals and standards documents on which the source manuscript's network descriptions rest are public technical documents cited there by designation. **The bibliography contains no standards documents and no protocol specifications, so none is cited here**, and every protocol claim is described from the source manuscript rather than asserted.

The networks are consequently de-branded. Following the convention of Curzi [2026], where a named system's characteristic is relied upon and cannot be supported from the bibliography, the name is dropped and the characteristic is kept: *a federated-agreement ledger with a curated default trust list of about thirty-five validators*, *a council-governed consortium*

ledger with a constitutional cap of thirty-nine members, *a delegated-proof-of-stake network with twenty-seven fixed producers*, and so on. **The functional descriptions are precise enough to identify what is being measured and carry no claim about any named network.** The single exception is the proof-of-work network at the bottom of the utility ordinal, which the abstract of this paper names as Bitcoin because Curzi [2026] and the corpus's citation record refer to the result by that name. **It is referred to in the body by function — a minimal-script proof-of-work ledger — and its protocol parameters are likewise described from the source manuscript and are not cited.** Appendix B records what the source names and what would have to be verified before any name is reattached.

What *is* cited, and genuinely supports what it is attached to: the concentration literature of Section 3.1; the portfolio and ecological literature of Section 3.3; the distributed-systems results of Section 5.3; the measurement-theory literature of Section 3.4; the disclosure principle of Section 3.5; and the identification, scaling and modelling literature of Sections 3.3, 13 and 14.

4. The atlas, de-branded and measured

Nine networks, ordered by the utility ordinal, each with the concentration axis its mechanism dictates. **Every effective count in the right-hand column is either an equal-weight ceiling implied by a protocol constant or a CONSTRUCTED illustration, labelled as such; none is a measurement of a running network.**

#	network, by function	consensus family	bloc character	utility	transit/store	concentration axis
1	minimal-script proof-of-work ledger	proof-of-work	non-aligned by issuance; pool-dependent on hashrate	1/9	store-oriented	pool share
2	federated-agreement settlement ledger, curated default list	federated agreement	issuer-centric	3/9	transit	curator of the default list
3	federated-agreement remittance ledger, quorum anchors	federated agreement	issuer-centric	3/9	transit	tier-one anchor organisations
4	programmable sidechain of family (2)	family (2) variant	family-adjacent; **undetermined**	5/9	transit	**undetermined**
5	council-governed consortium ledger	asynchronous BFT gossip	consortium	6/9	transit (enterprise)	council membership, capped
6	capped-BFT interchain routing hub	BFT proof-of-stake	diffuse; bridge	6/9	transit / routing	stake within the capped active set
7	fixed-producer-set network	delegated proof-of-stake	East-leaning	7/9	transit	the twenty-seven producers
8	high-throughput proof-of-stake payments network	verifiable-delay clock + stake	Western-venture-leaning	8/9	transit	stake-weighted minimum coalition
9	general-virtual-machine settlement layer	proof-of-stake with a finality gadget	Western-leaning, diffuse	9/9	mixed	staking-entity share

Three columns and three warnings. **The utility column is an ordinal and only an ordinal.** The transit/store column is a statement about what a design invites and is not a recommendation. **The bloc column is a stipulated assignment, not an identified effect** (Section 3.3), and row 4 is left undetermined in print because the documentation the source manuscript relies on does not support a call — **declining to guess is the standard, and it is exercised here.**

5. Results: H1 — the instrument transfers unchanged

5.1 The measured spread

basis	HHI	effective count	status
capped council at its constitutional cap, equal weight	0.025641	39.0	ceiling implied by a protocol constant
curated default trust list, equal weight	0.028571	35.0	ceiling implied by a stated list size
capped BFT active set at its governance cap, equal weight	0.005556	180.0	ceiling implied by a governance parameter
fixed producer set, equal weight	0.037037	27.0	ceiling implied by a protocol constant
tier-one anchor organisations, equal weight	0.142857	7.00	ceiling
tier-one anchor organisations, constructed unequal	0.166667	6.00	**CONSTRUCTED**
capped BFT active set, constructed top-heavy stake	0.025481	39.2	**CONSTRUCTED**
fixed producer set, constructed concentrated vote	0.053946	18.5	**CONSTRUCTED**
staking entities, constructed shares	0.168400	5.94	**CONSTRUCTED**
proof-of-work pools, constructed shares	0.190000	5.26	**CONSTRUCTED**

The constructed unequal anchor vector is $2/7 = 28.5714\%$ against six equal shares of $5/42 = 11.9048\%$, which gives $(2/7)^2 + 6(5/42)^2 = 0.081633 + 0.085034 = 0.166667$ exactly, so $N = 6.00$ against the seven-member ceiling of 7.00. **H1 is supported: every row behaves exactly as the instrument does on jurisdictional shares — returning n at equal weight, falling below n whenever shares are unequal, and never exceeding $1/s_{\max}^2$.**

The reading that matters is the gap between a ceiling and a distribution. A design that caps its active set at one hundred and eighty has capped its own effective count at one hundred and eighty before a single unit of stake is bonded; the honest number is the stake-weighted one and it sits below the cap, because stake concentrates toward the top of a set. **A cap is an upper bound the protocol guarantees and a distribution is what the network has, and reporting the first as though it were the second is the error the whole paper is written against.**

5.2 R5 — the curated-list collapse

The sharpest instance the transfer produces, and it is a two-row table.

```
single-curator trust set          HHI = 1.0000   ->  effective curators  = 1.00
five equal independent operators  HHI = 0.2000   ->  effective operators = 5.00
```

Suppose thirty-five validators each run honest, independent, geographically dispersed infrastructure, and suppose all thirty-five appear on a default list published and maintained by one organisation, and suppose participants follow that list because following it is the path of least resistance. **The naive reading squares thirty-five roughly equal shares and reports 35.0. The bloc reading asks whether a single decision reaches all thirty-five, and the answer is that the list publisher's decision to add, remove or reweight an entry does exactly that.**

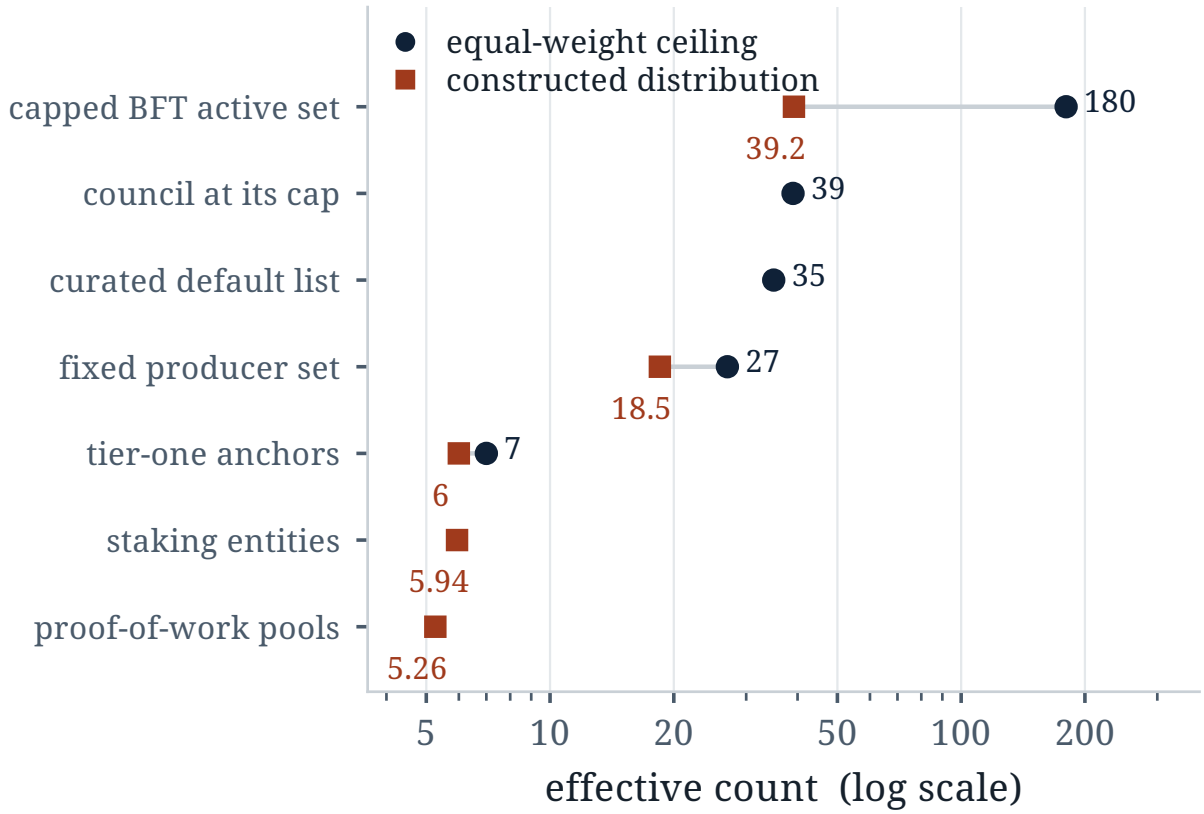


Figure 67.1: **A ceiling is not a measurement.** Effective counts on six ledger participant sets: the equal-weight ceilings implied by protocol constants (hatched) against constructed unequal distributions (solid), showing the gap the ceiling conceals — a 180-member capped set falls to a constructed 39.2, a 27-member fixed producer set to 18.5, and a seven-member anchor set to 6.00. Ceilings are derived from stated protocol or governance parameters; unequal distributions are CONSTRUCTED and are not measurements of any running network. Script `figs/src/CCP-067_fig01.py`.

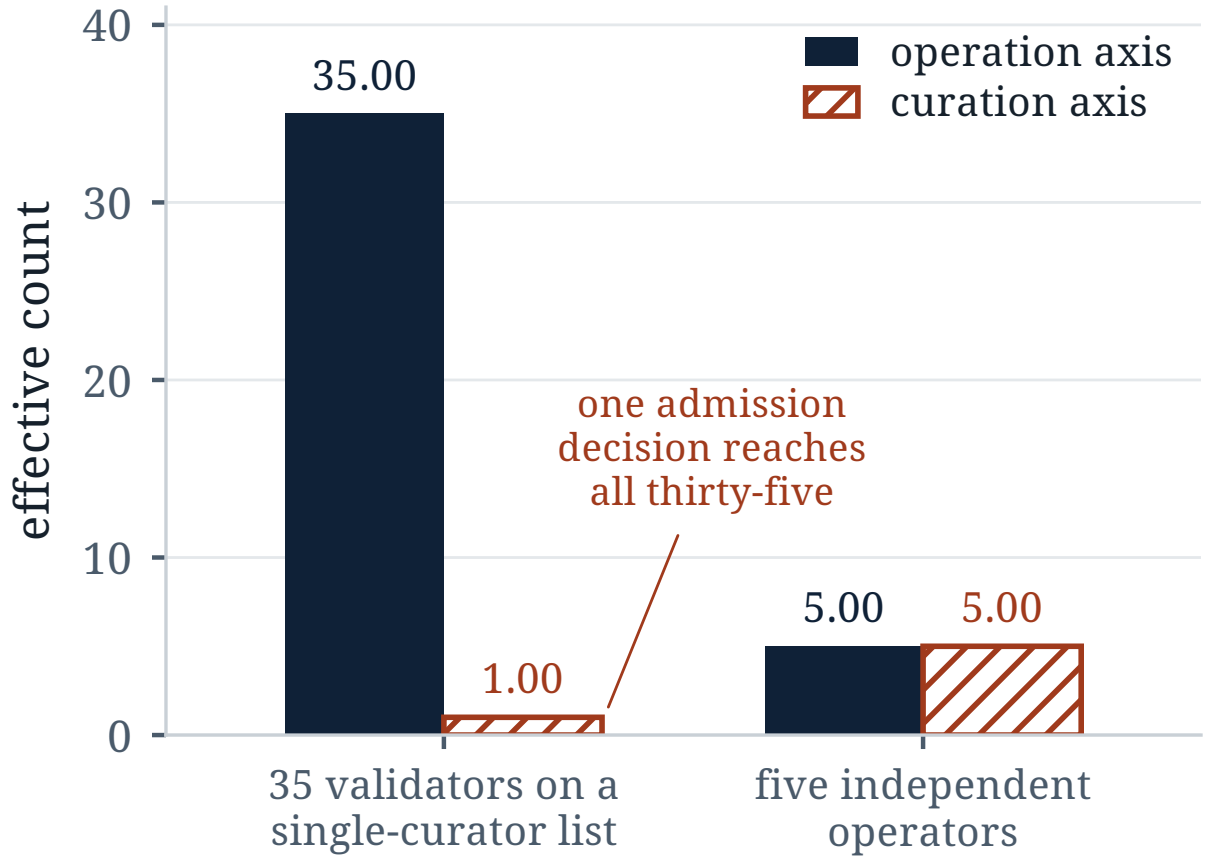


Figure 67.2: **Two questions about one validator set.** Thirty-five independently operated validators give an effective count of 35.0 on the operation axis and 1.00 on the curation axis, because one party admits entries to the default trust list; five genuinely independent operators give 5.00 on both. The gap between the two readings is the finding. Ceilings from stated list sizes; the single-curator and five-operator rows are SYNTHETIC endpoints of a spectrum. Script `figs/src/CCP-067_fig02.py`.

On the list-curation axis the effective count is not thirty-five; it is one, and the two numbers describe the same validators under two different questions about who can move them together. H5 is supported.

The two rows are endpoints of a spectrum and a real network sits between them, and where it sits is an empirical question about that network's actual governance which this paper does not assert for any named ledger and defers to Appendix B. The specification point is that *public* and *decentralised* are not the same property, that a public ledger can be operationally curated by very few parties, and that the instrument measures the curation directly once it is pointed at it.

5.3 What the distributed-systems literature already establishes

The collapse has formal ancestry that should be named rather than reinvented. Douceur [2002] proves that in a system without a trusted certifying authority a single party can present as arbitrarily many distinct identities, so **an unaudited count of participants is worthless** — which is the curated-list result in the identity domain, with an impossibility result attached. Lamport et al. [1982] fixes the number of genuinely independent parties required to tolerate f arbitrary faults at $3f + 1$, which is the same question asked about agreement. Fischer et al. [1985] and Gilbert and Lynch [2002] establish that structural impossibility results of this genre

exist and bound what any consensus design can promise.

None of these is imported as a theorem about any network here. They are cited as prior formalisations of the same structural property, under the discipline of Section 13.5. Buldyrev et al. [2010]'s coupled-network fragility is the closest formal analogue of a common-substrate dependency and is used for shape only, since its percolation premises are not shown to hold; Albert et al. [2000] and Callaway et al. [2000] supply the robust-yet-fragile shape, and Doyle et al. [2005] is the reason the degree-distribution mechanism behind it must not be assumed to apply.

6. Results: H3 — proof-of-work concentration is measured on pools

The coordinating unit on a proof-of-work network is the pool, because a pool *is* a coordination mechanism by construction: many machines agreeing to be directed by one operator's block template. A count of machines therefore overstates independence exactly as a count of validators does elsewhere.

A **CONSTRUCTED** ten-pool distribution, to show the instrument and not to describe any running network:

```
CONSTRUCTED pool shares  [30, 26, 12, 9, 6, 5, 4, 3, 3, 2] per cent

    HHI              = 0.190000  (1,900 points)
    effective pools  = 5.26      (of 10 nominal)
    leading share    = 30.0%     ceiling 1/s_max^2 = 11.11
    top two combined = 56.0%
```

H3 is supported: ten nominal pools give 5.26 effective, and the leading share alone caps the count at 11.11 before any other share is known. When the two largest pools together approach half the work, the effective count sits far below the count of pools, and that is the concentration reading a proof-of-work rail requires.

Two disciplines belong with the measurement and neither is optional.

A pool's share is not the same as its control, because contributors can leave. Work can be redirected between pools quickly, so a pool acting against its contributors' interests faces exit; the exit option is a real constraint that a static index does not capture. **The pool index is therefore an upper bound on entrenched control and should be stated with the exit constraint or the number misleads.**

Pool concentration is geographically and politically mobile. Proof-of-work capacity relocates across jurisdictions in response to energy prices and policy, so a bloc reading of hashrate is itself a moving target and is not a fixed national alignment. **The honest two-part answer is non-aligned by construction on issuance and pool-dependent on hashrate**, and the second half must be dated.

No distributional claim is made about any of these share vectors. They are finite sets of three to ten stipulated shares, not samples from a distribution; no tail is characterised and no scaling law is asserted. **The Clauset–Shalizi–Newman procedure was therefore not run** [Clauset et al., 2009]. Hashrate and stake distributions are exactly the sort of quantity for which a heavy-tail claim is reached for casually; **any restatement of this paper that makes one must run that procedure first**, and must read Stumpf and Porter [2012] on why passing it would still not identify a mechanism and Broido and Clauset [2019] on how weakly empirical networks support scale-free structure. Barabási and Albert [1999] is the source of the vocabulary and is cited for that and nothing more.

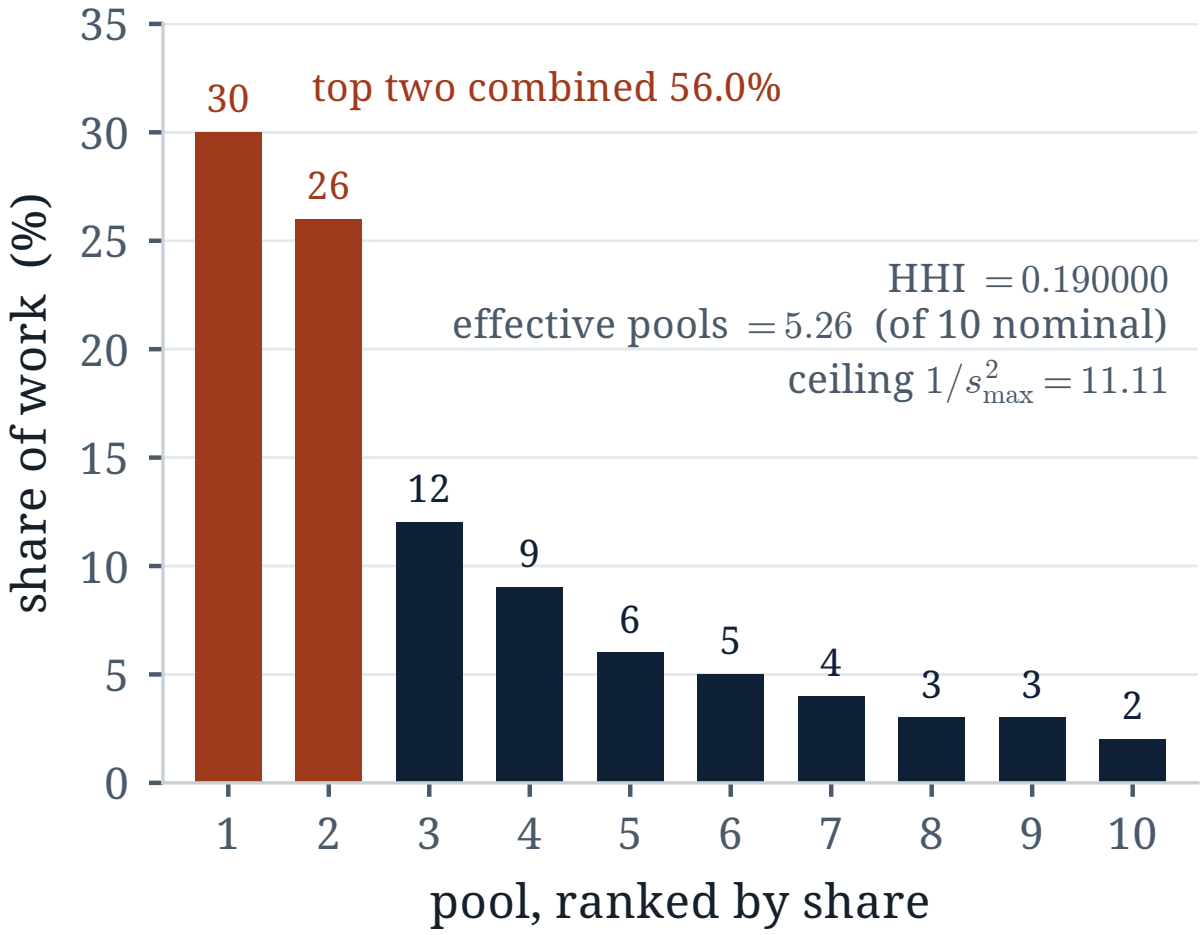


Figure 67.3: **Ten pools, 5.26 effective, and a ceiling of 11.11 set by the leader alone.** A constructed ten-pool hashrate distribution with a 30 per cent leader and a 56.0 per cent top two gives HHI $= 0.190000$ and an effective count of 5.26 against ten nominal, well under the $1/s^2_{\max} = 11.11$ bound the leading share imposes on its own. CONSTRUCTED distribution; not a claim about any running network. Script `figs/src/CCP-067_fig03.py`.

7. Results: H4 — the utility ordinal, stated as an ordinal

7.1 The axis, and the rank printed with it

```
1/9  minimal script, no general virtual machine -- minimal BY DESIGN
3/9  native exchange, escrow and payment channels; no general VM
3/9  native issued assets and path payments; no general VM natively
5/9  native on-ledger contract modules on a minimal settlement lineage
6/9  general-VM-compatible contracts alongside purpose-built native services
6/9  application-chain runtime plus interchain routing
7/9  general-VM-compatible runtime, heavy contract use
8/9  parallel-execution runtime built for payments-scale throughput
9/9  general virtual machine; the broadest settlement-logic surface
```

The claim, written in the form it must always take so that no paraphrase can drop the scoping:

> On the utility-as-function axis — programmability and general settlement expressiveness, and no other axis — a minimal-script proof-of-work base layer ranks lowest among the major networks, and it ranks lowest **by design**: no general smart-contract virtual machine, a ten-minute target block interval, low base-layer throughput, and no programmability beyond a deliberately constrained script. The design goal is monetary conservatism, and the ranking describes the goal rather than a defect.

Four scoping rules govern every use of the claim.

1. **Name the axis every time.** “Lowest utility” unscoped is unprovable and is not a claim this paper makes; “lowest on the utility-as-function axis, by design” is a description of a specification.
2. **State that other axes rank it differently** and do not let a single-axis rank stand in for a verdict.
3. **No price.** The axis is about function. The moment “utility” drifts toward “worth” it has left this paper’s subject entirely.
4. **The rank is an ordinal.** A 1 is more constrained than a 9; it is not worse than a 9. **The axis is not oriented.**

7.2 The inversions, which are why the axis must be named

The same design decisions that place a minimal ledger at the bottom of this axis place it at or near the top of three others, and the inversion is the point.

axis	what it asks	where a minimal ledger sits
utility-as-function	what settlement logic can be expressed	**bottom**
issuance decentralisation	how many parties can change the supply	**top** — there is no protocol seat from which to do it
security budget	what it costs to rewrite recent history	**top among its own consensus family**
network liveness and governance resistance	whether block production and monetary rules have held	**at or near the top**

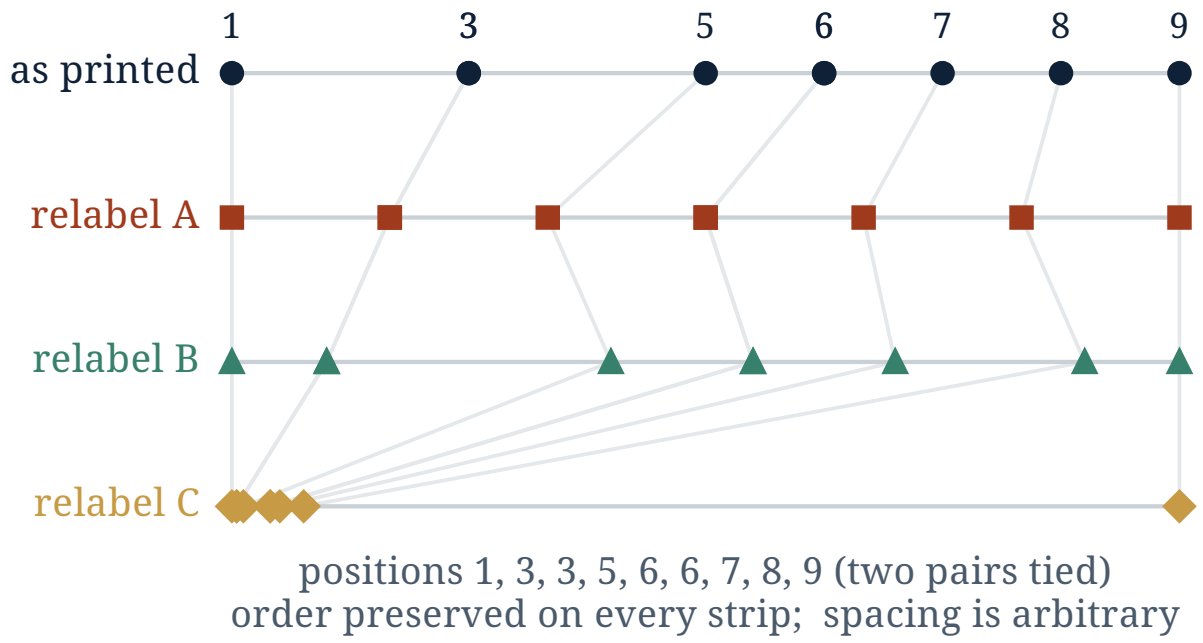


Figure 67.4: **An ordinal carries order and nothing else.** The nine utility-as-function positions under three order-preserving relabellings: the ranks change, every comparison this paper makes is unchanged, and no statement here survives being restated as a difference or a ratio of positions. The axis is unoriented — higher is more expressive, not better. Definition of this paper's own construction; carries no empirical weight. Script `figs/src/CCP-067_fig04.py`.

A network near the bottom on one axis and near the top on three others is not a low-utility network; it is a network that spent its entire design budget on a small number of properties and refused the rest, and the refusal is legible as a ranking only when the axis is stated. The minimalism produces the low utility rank *and* the high issuance-decentralisation rank: they are one design decision read on two axes.

H4 is supported in the only sense an ordinal claim can be supported. Every statement this paper makes about the axis — that a given network is lower than another, that the axis is unoriented, that the ranking describes a design goal — survives any order-preserving relabelling of the nine positions, and no statement it makes is a difference, a ratio or a mean of positions. **The axis carries order and nothing else, and this paper does not operate on it as though it carried more.**

7.3 What the minimal instruction set cannot express, and why that is the finding

can express (base layer)	cannot express (by design, no general VM)
multi-signature (m-of-n)	unbounded loops and arbitrary computation
absolute and relative time locks	general mutable contract state across transactions
hash locks — the primitive routed payments are built from	contract-to-contract calls

can express (base layer)	cannot express (by design, no general VM)
aggregated-key and script-path spends	the general programmability of a virtual-machine environment

The right-hand column is the utility rank restated as a capability list, and the load-bearing observation is that the left-hand column is exactly the primitive set a payment layer needs while the right-hand column is exactly the primitive set whose absence removes a class of failures. An instruction set that cannot express a re-entrancy fault has no re-entrancy fault to exploit. **The rank is the location of a boundary a conservative designer would draw, not a measure of inadequacy.**

8. Results: H6 — base-layer throughput, stated to its parameters

Throughput on a minimal base layer follows from two protocol parameters and a transaction size, and is therefore computable rather than contestable:

```
at ~226 vbytes per transaction, ~1e6 vbytes per block, 600 s interval
    1e6 / 226 = 4,424.8 transactions per block  ->  ~7.4 tx/s

at ~140 vbytes per transaction, same block and interval
    1e6 / 140 = 7,142.9 transactions per block  ->  ~11.9 tx/s
```

H6 is supported to the precision the inputs allow: approximately 7.4 **to** 11.9 **transactions per second at the base layer.** These are single-digit-to-low-double- digit rates against the tens of thousands a consumer payment network processes, and the gap is the throughput half of the utility rank.

The figure is an order of magnitude and not a constant, and the paper says so. The real rate depends on transaction mix and on how witness data is accounted against the block limit, so the two values bracket a band rather than fixing a point. **Quoting either to more than two significant figures would be false precision on a quantity that varies with the traffic**, and Section 13.2 records that.

The ten-minute interval is the parameter most often read as a flaw and is the clearest design choice of the set. A block interval is a trade: shorter intervals settle faster but propagate less completely before the next block is found, raising the rate of orphaned work and the advantage of well-connected participants, which concentrates. **A long interval buys near-complete propagation between blocks and a level field for a badly connected participant, at the explicit cost of base-layer latency.** The trade is the conservatism, not an oversight. A self-correcting difficulty retarget then holds the cadence roughly fixed across large swings in capacity, which is what makes the exposure window of Section 9 predictable as a protocol property rather than as a matter of operator diligence.

One structural feature is stated plainly rather than papered over. The cost of rewriting recent history scales with the honest work an attacker must out-compute, and that work is paid for in aggregate by the block reward — a subsidy that halves on a fixed schedule and trends toward zero as issuance approaches its cap, plus fees. **Over the very long run the security budget must therefore be carried increasingly by fees rather than by new issuance.** That is not a prediction of failure and not a defect to be hidden: it is an open design question about the distant future, it does not bear on a near-term rail assessment, and **anyone extrapolating it confidently in either direction has left engineering.** The honest posture is to name the mechanism and decline to guess.

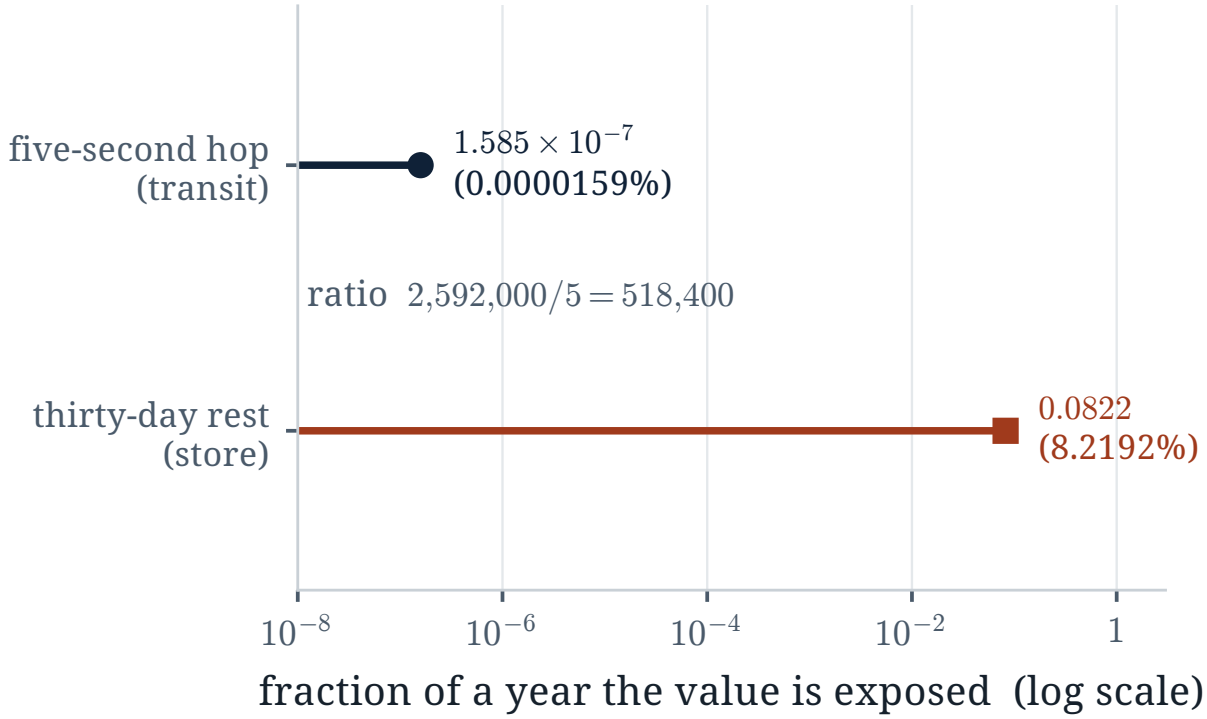


Figure 67.5: **The store-versus-rail question, answered by dwell.** A five-second hop occupies 1.585×10^{-7} of a year against 0.0822 for value rested thirty days — a ratio of 518,400 to one — so a freeze hazard that is negligible in transit is dominant in storage, and dwell is the only variable separating them. Logarithmic axis; CONSTRUCTED dwell parameters, order-of-magnitude only. Script `figs/src/CCP-067_fig05.py`.

9. Results: H2 — dwell time is the whole freeze exposure

```
per-hop dwell at seconds-scale finality        ~5 s
  exposure fraction of a year                   1.585e-07   (0.0000159%)

value rested on the rail for 30 days
  exposure fraction of a year                   0.0822      (8.2192%)

ratio, store versus transit                     2,592,000 / 5 = 518,400
```

Five seconds is $5/31{,}536{,}000 = 1.585 \times 10^{-7}$ of a year; thirty days is $30/365 = 0.082192$. **The window while value is genuinely in transit is on the order of seconds and the fraction of a year it represents is vanishingly small; the same rail used to rest value for a month exposes it for the whole month, a window larger by more than five orders of magnitude.**

H2 is supported. The freeze property that disqualifies a rail as a store is nearly irrelevant to it as transit, and the entire difference is dwell. That is the corpus's transit-versus-store rule converted from a categorical distinction into a quantity, and it is why the rule is a design parameter rather than a preference.

The counter-result belongs immediately beside the result. Where a settlement instrument's issuer can immobilise a balance, **that control lives on the instrument and not on the rail, and it follows the instrument across every rail that carries it.** A structure diversifying across rails to escape such a control has diversified the wrong axis. **Rail diversity buys settlement resilience and buys no freeze protection**; only counterparty and instrument diversity does. The instrument in this paper measures the rail, and the freeze lives elsewhere.

A programmable rail sharpens the same point. Where transfer can be made conditional

by logic attached to the asset rather than by an instruction from its holder, an interruption need not be an instruction sent after the fact — it can be a property the unit carries from issuance. **Freeze exposure is then a transit risk of seconds if the window is kept at the rail's minimum finality, and a custody risk of the entire holding period if the rail is treated as a store.** The arithmetic above is the whole of the argument.

9.1 The two conflations this paper refuses

A pilot is not a production system. A pilot is a controlled experiment run to learn something; a production system is infrastructure that real value depends on; the distance between them is measured in years and in categories of risk a pilot is specifically designed not to incur. **This paper asserts that no central-bank digital currency runs in production on any named public permissionless ledger, and states immediately that this is a claim about the absence of a documented counter-example as at the source manuscript's date of writing and not a proof of a universal negative.** Appendix B records it as such. Every use of "runs on" in secondary coverage should be read as "was trialled on" until a primary source says otherwise, and **this paper names no pilot, no counterparty and no status, because the precise status of each is exactly the variable most often misreported.**

A product is not a mandate. That a vendor offers a platform, or that an organisation participates in a trial, is not evidence that any bloc's settlement runs on anything. **A commercial platform described as running on a private or permissioned instance of a technology derived from a public ledger is not the public ledger**, and that distinction is precisely what a careless summary elides. This paper does not resolve it and defers it to Appendix B.

9.2 Two ledgers, one perimeter

The bloc collapse applied to a settlement design rather than to a participant set:

```
two public ledgers, both curated around issuers inside one perimeter

naive, by ledger    shares 0.50 / 0.50
                    HHI = 0.5000  ->  effective ledgers = 2.00
honest, by bloc     both inside one issuer perimeter
                    HHI = 1.0000  ->  EFFECTIVE BLOCS   = 1.00

bloc-diverse alternative (CONSTRUCTED)
0.34 / 0.33 / 0.33  HHI = 0.3334  ->  EFFECTIVE BLOCS   = 3.00
```

Two ledgers, one perimeter, one point of failure. A settlement layer confined to one perimeter has an effective bloc count of one, and no rearrangement of the two ledgers changes it, because rearranging within a perimeter does not change the bloc count. **Bloc diversity in the settlement layer is bought the same way as everywhere else — by spanning perimeters that do not share a coordination mechanism — and it cannot be bought by multiplying ledgers that do.** The companion result carries too: diversity raises the ceiling on resilience and does not pay the maintenance cost of operating across multiple rails, which is Curzi [2026]'s headroom gate $h = \eta \ln N - \delta Q - C$ applied to a settlement layer — worked there at $N = 18$ with $\eta = 0.30$, $\delta = 0.06$, $C = 0.55$, $Q = 1.20$ giving $h = +0.2451$ — and it is a structurally parallel result in another domain rather than a re-derivation.

9.3 The screening surface follows the identifier onto the ledger

Counterparty pre-screening on a settlement rail is string comparison over identifiers — account addresses, destination tags, memo fields and any party-name data the message layer carries. Curzi [2026] establishes the defect this inherits: typed-field parsing collapses the

withdrawn-code parsing-ambiguity false-positive share from 88.99% to 0.00% against a naive detection precision of 11.01% — **a control right one time in nine.** That finding does not stop at an institution's database boundary; it follows the identifier onto the ledger, because an on-ledger identifier is a string and an exact-match comparison over a field admitting non-ASCII characters has a silent blind spot.

Programmability widens the hole rather than closing it. An on-ledger compliance module that screens by exact string match executes the same defective comparison at settlement time, in code, in a settlement path that is often immutable by design. **An exact-match rule compiled into an immutable settlement path is a permanent commitment to a comparison with a known blind spot, and correcting it means replacing the code rather than editing the record.** The remedy is Curzi [2026]'s: canonicalise the critical fields before comparison and store the byte-level hash of the raw field so that a substitution leaves a permanent comparable trace. Curzi [2026] states why the witness must be computed over the raw octets and held on a plane separate from the data it guards — **a witness that is a function of the canonical form changes exactly when the data changes and never otherwise, which destroys the evidence a witness exists to preserve.**

10. Discussion

10.1 What the atlas does and does not license

The atlas returns, for each network, an effective count on a named axis and an ordinal position on a named scale. **It does not rank networks, because two of its three columns are not orderings of goodness and the third is a stipulated bloc assignment.** A reader looking for a recommendation will not find one, and the refusal is the design.

The frame this paper inherits is bounded in three places by papers that state their own limits, and the bounds are load-bearing rather than decorative.

Curzi [2026] bounds what a surplus measure can claim: the surplus is only a lower bound on an ensemble's information, and its logarithmic concavity means three well-separated views capture most of what a large panel could, so the node argues explicitly *against* accumulating perspectives; the accompanying gradient result concedes that descent can diagnose a stable trap but cannot treat one. **Applied here, that bounds the value of adding rails**: the fourth independent perimeter buys much less than the second, and no amount of measurement tells a structure how to leave a trap it has correctly diagnosed.

Curzi [2026] bounds what borrowed physics may claim: its counter-result requires that a user be able to distinguish the theorem, the ordinary observer effect, and the framework's own conjecture in three separate sentences on demand, and it states its postulate as a postulate. **Applied here, that is why Section 5.3's distributed-systems results are labelled category 2 in Section 13.5** rather than allowed to carry an impossibility proof's authority into a claim about a running network.

Curzi [2026] is the bound stated as an appendix: it assigns the corpus's central identification, the application of the surplus functional to a life, and the reading of a day's surface as a leading indicator to the postulate drawer, and says of the third that it is a wager whose misfire column exists because wagers can lose. **Applied here, it is the reason Section 4's bloc column is printed as a stipulation and row 4 is left undetermined**: an honest register has an empty cell in it.

Curzi [2026] bounds what an ordinally known parameter can produce. Its counter-result is that the framework's own keystone formula is demoted by the framework: with its inputs known only ordinally, the derived frequency moves by 5.77 octaves under an order-preserving relabel and so is not known as a number at all, and a fully successful intervention on one input makes the equilibrium eighty per cent worse because a numerator was negative. **That is the exact hazard the utility ordinal carries**, and it is why Section 7 refuses every operation on the axis except comparison.

Curzi [2026] bounds inference from an invariant statistic. Its counter-result is that a width statistic invariant under relabelling does not even determine the object it is computed from, which formally blocks any inference from a high or low value to a substantive condition. **The utility ordinal is invariant under order-preserving relabelling in exactly that way**, so it does not determine the design it is computed from, and no inference from a position on it to a network's quality is available. **This paper states that as a block on its own axis rather than as a caveat.**

10.2 What the model form is, and is not

Curzi [2026] badges the corpus's three-scope correspondence **POSTULATE — MODEL FORM ONLY**: identical structure at three scopes is a statement about the chosen model, not a claim of kinship or of cross-scope aggregation. **This paper's central move is exactly such a correspondence** — the same index at the jurisdictional, correspondent, validator and pool scopes — and it is badged the same way. **That an index has the same algebraic form on four share bases is a fact about the index. It is not a claim that jurisdictions and mining pools are the same kind of thing, and no result here aggregates across the scopes.** Cartwright [1999] supplies the discipline: state the arrangement within which a law is claimed to hold and stop claiming a universality the argument does not need.

10.3 What the corpus's frequency material contributes, and what it does not

Curzi [2026] is presupposed for its minimum sustainable frequency $f_{\min} = k\, e^{2c/\eta}$ and the trichotomy in which f greater than, equal to, or less than $f_{\min}$ corresponds to a positive, zero or negative equilibrium stock. **Two things must be said about the constant in that expression and both are said here rather than inherited silently.** $k = (9/8)^2/\varphi^2 = 0.483425732988414$ is a **DEFINED** constant — an exact ratio of exact quantities, computable to arbitrary precision — and its fifteen figures are legitimate only under that label. **Attaching hertz to it is a separate and much stronger claim than computing it, and it is POSTULATED**, with no physical, psychoacoustic or musicological basis in any literature for combining a tuning ratio with an inverse square of the golden section to obtain a ground frequency. **Markowsky [1992] is cited here precisely because φ appears**: it documents how routinely golden-ratio claims are overstated, and citing it is how this paper shows it knows what such a claim usually looks like and is not making one.

Nothing in this paper depends on k, on $f_{\min}$, or on any frequency. The citation records a presupposition, and the presupposition is bounded by the badge.

11. Failure modes

11.1 Reading a rail as a store. Eight of the nine networks are transit-oriented by design, and treating any of them as a place to rest protected capital imports a rail's exposure profile into a custody layer. **The difference is dwell and it is a factor of $518{,}400$ on the constructed parameters.**

11.2 Counting validators the way entities are counted. A large advertised node count is a count of participants and not a measurement of failure independence. **A thousand validators subscribing to one curated list have the failure independence of the list.**

11.3 Missing the curated-list collapse. The validators may be numerous, honest and independent in their operations, and the network may still collapse to one on the curation axis, **because the coordination mechanism is not the operation of the nodes but the decision about which nodes the network defaults to trusting.**

11.4 Conflating a pilot with production, and a product with a mandate. Section 9.1. Every other error in this area is downstream of this one.

11.5 Assuming programmability is a quality axis. The ordinal is not oriented. **A general virtual machine is a larger error surface, and a minimal ledger's inability to express a class of settlement logic is also an inability to express that class's failures.**

11.6 Forgetting that the ledger is public. Every transfer on all nine networks is permanently visible, so a structure reaching for a public chain as a privacy layer has inverted the technology's actual property. **The correct posture is to assume full disclosure and design so that disclosure changes nothing** [Kerckhoffs, 1883, Saltzer and Schroeder, 1975], which Curzi [2026] states as a hard design premise rather than a preference.

11.7 Treating the issuer freeze as the rail's property. The control follows the instrument, not the rail. **Rail diversity does not dilute an instrument-level freeze; only counterparty and instrument diversity does.**

11.8 Double-counting a bridge. An interchain router touches two perimeters, which is its value and its measurement hazard: counted as a fully independent perimeter it overstates diversification, folded into one of the perimeters it connects it understates it. **The correction is to state the dependency, not to fudge the shares.** The same shape appears wherever an instrument native to one network is wrapped and settled on another and the two are then counted as independent.

11.9 Reading a time-varying number as a fact. Validator counts, vote distributions, stake shares, pool shares, minimum-coalition coefficients and council rosters all move, some weekly. **A concentration measurement taken from a year-old figure is not merely rounded; it is stale**, and this paper therefore states mechanisms and constructed illustrations and dates every real number to the reader's own live source.

11.10 Believing the consensus paper describes the running network. A published protocol is a specification; the running network is an implementation with operators, defaults and emergent concentration the specification does not describe. **Federated agreement as specified is decentralised; federated agreement as run collapses to its default list.** Confusing the two credits a network with the decentralisation of its documentation rather than the concentration of its deployment. This is Leveson [2011]'s and Vaughan [1996]'s mechanism under a different name: a control justified once continues to be trusted after the conditions justifying it have gone.

11.11 Reporting a ceiling as a measurement. A cap is an upper bound the protocol guarantees. A distribution is what the network has. **Section 5.1's whole content is the gap between them.**

11.12 Averaging the utility ordinal. A mean of ordinal positions is not a quantity [Stevens, 1946, Krantz et al., 1971], and a table that computes one has performed an unlicensed operation. **The ordinal admits comparison and nothing else.**

12. Verification and rigour audit

12.1 Dimensional audit

symbol	quantity	dimension	unit	status
s_i	participant share of the coordinating quantity	$[1]$	fraction	defined; sum to 1
HHI	sum of squared shares	$[1]$	fraction (or points $\times 10^4$)	derived

symbol	quantity	dimension	unit	status
N	effective count	$[\#]$	validators, members, producers, pools, curators	derived as HHI^{-1}
utility position	rank on a nine-point ordering	**none**	rank label only	**ORDINAL — differences undefined**
w	dwell on the rail	$[T]$	seconds; days; years	assumed (constructed)
λ	interruption hazard	$[T^{-1}]$	per year	assumed (not quantified here)
λw	expected freeze exposure	$[1]$	dimensionless fraction	derived
transaction size	virtual bytes per transaction	$[B]$	vbytes	assumed (stated)
block capacity	virtual bytes per block	$[B]$	vbytes	assumed (stated)
block interval	target time between blocks	$[T]$	seconds	assumed (protocol constant, described)
throughput	transactions per unit time	$[T^{-1}]$	tx/s	derived

Does the throughput calculation balance? $[B]/[B] = [1]$ transactions per block, divided by $[T]$ per block, gives $[T^{-1}]$. $10^6/226 = 4424.8$ transactions per block over 600 s is 7.37 per second.

Does the dwell calculation balance? λw with λ per year and w in years is dimensionless, which is why 5 s must be converted: $5/31{,}536{,}000 = 1.585 \times 10^{-7}$ years.

The utility position has no dimension and no unit, and the audit says so explicitly rather than assigning it a placeholder. It is a rank label. **This is the row a reader should check first, because a rank in a numeric column is the classic place an unlicensed operation enters a table.**

12.2 Significant figures and epistemic status

- **Every effective count here is exact for its share vector and is not an estimate.** $1/35 = 0.028571\ldots$, $1/39 = 0.025641\ldots$, $1/0.19 = 5.263\ldots$; **no uncertainty is quoted because nothing is estimated.**
- **Equal-weight ceilings are exact consequences of stated set sizes**, and are labelled ceilings rather than measurements everywhere they appear.
- **The constructed distributions are stipulations.** 0.166667, 0.025481, 0.053946, 0.168400 and 0.190000 are exact arithmetic on invented share vectors, and exact arithmetic on invented inputs is not a measurement.
- **The throughput figures are quoted to two significant figures on purpose.** 7.4 and 11.9 bracket a band whose width depends on transaction mix; **quoting them to more figures would be false precision.**
- **The dwell figures are order-of-magnitude.** 1.585×10^{-7} is exact for a five-second dwell and a 365-day year, and the five seconds is a stipulated order-of-magnitude parameter, not a measured finality. **The ratio** $518{,}400$ **is exact for the stated dwells and generalises in sign and shape and not in value.**
- $k = 0.483425732988414$ **is DEFINED** and appears in Section 10.3 at that precision only under that label; **the hertz attached to it is POSTULATED** and is not used anywhere in this

paper.
- **The utility positions are not numbers.** They are printed as fractions of nine because that is the source manuscript's notation; **they are labels and no arithmetic is performed on them.**

12.3 Limits and boundary behaviour

quantity	limit	behaviour	reading
N	n equal shares	$\to n$ exactly	the numbers-equivalent property
N	one share $\to 1$	$\to 1$	full concentration; the floor
N	$s_{\max}$ fixed, others $\to$ equal	$\to 1/s_{\max}^2$ from below	the ceiling the leader alone imposes
curation collapse	one curator	$\to 1$ regardless of set size	H5; **the set size becomes irrelevant**
curation collapse	m independent curators, equal	$\to m$	the curation axis is plural exactly when admission is
λw	$w \to 0$	$\to 0$ linearly	a hop of zero duration has zero freeze exposure
λw	$w \to$ holding period	$\to$ custody exposure	**a rail held long enough simply is a store**
throughput	transaction size $\to 0$	$\to \infty$	**unphysical; the model's domain is realistic sizes**
throughput	block interval $\to 0$	$\to \infty$	unphysical: propagation and orphan rate bound it, and that bound is not modelled here
utility ordinal	any order-preserving relabel	every claim unchanged	**the definition of an ordinal, and the paper's own test of H4**

The two unphysical limits are stated as domain restrictions rather than discovered by a reader. The throughput identity is a capacity ratio and carries no propagation model; it is valid only in the band where the stated block interval holds.

12.4 Invariance

HHI^{-1} **is invariant under any permutation and any relabelling of the units**, which is why the bloc collapse depends on the partition and not on which member of a bloc holds what. **It is not invariant under a change of share basis**, and the share basis is therefore item 1 of the reporting standard and the subject of the entire Section 3.2 table.

The dwell ratio is invariant under any rescaling of the hazard, since it is a ratio of two exposure windows, which is why it is stated as a ratio and not as a probability.

The utility ordinal is invariant under every order-preserving transformation, and Section 7.2's test of H4 is exactly that invariance: relabel the nine positions in any order-preserving way and every claim survives. **That invariance is also the block on inference** — the bound carried from Curzi [2026] — since a statistic invariant under relabelling does not determine the object it was computed from.

Tested at how many scales? **The invariances are algebraic and were checked at the parameter sets reported and at no others.** No empirical scale-freedom is claimed anywhere in this paper.

12.5 Analogy discipline

borrowing	source	category	justification
HHI^{-1} as a numbers-equivalent	Adelman [1969], Hill [1973], Jost [2006]	**1 — derivation**	shares are non-negative and sum to one by construction; the identity is exact
independent reinvention of the same index	Simpson [1949], Laakso and Taagepera [1979]	**1 — derivation**	the same arithmetic on the same object in three disciplines
exposure as a portfolio; covariance not count	Markowitz [1952]	**1 — derivation**	the bloc collapse *is* the variance identity, and its premises hold by construction
similarity-sensitive diversity of order 2	Leinster and Cobbold [2012], Rao [1982]	**1 — derivation**	exact for a unit-diagonal, non-negative correlation matrix, which is the regime used
identity multiplication defeats a count	Douceur [2002]	**2 — shared form**	proved for an identity system without a certifying authority; no ledger here is shown to satisfy those premises
independent-party bound	Lamport et al. [1982]	**2 — shared form**	$3f + 1$ is a theorem about agreement, cited as a prior formalisation of "how many parties are genuinely independent"
structural impossibility results exist	Fischer et al. [1985], Gilbert and Lynch [2002]	**3 — framing**	cited to establish the genre of claim, never as results about any network here
coupled-network fragility; robust-yet-fragile	Buldyrev et al. [2010], Albert et al. [2000], Callaway et al. [2000]	**2 — shared form**	the shape is used; the degree-distribution mechanism is not claimed, and Doyle et al. [2005] is the reason it should not be

borrowing	source	category	justification
correlation rising under stress	Longin and Solnik [2001], Ang and Chen [2002]	**2 — shared form**	measured in equity markets; ledgers are not shown to satisfy the same conditions
effective count buys stability	Schindler et al. [2010], Loreau and de Mazancourt [2008]	**2 — shared form**	an ecological measurement; the shape is used, the mechanism is not claimed
scale type and its permissible statistics	Stevens [1946], Krantz et al. [1971], Narens [1996]	**1 — derivation**	the utility axis genuinely is an ordinal scale, and the restriction follows from the definition
open design; disclosure strengthens a structural control	Kerckhoffs [1883], Saltzer and Schroeder [1975]	**1 — derivation**	the premise — effectiveness depends on the arrangement rather than on anyone's knowledge of it — is stated and satisfied
performativity of a published metric	MacKenzie and Millo [2003], Lucas [1976]	**1 — derivation**	the mechanism is the same: a description relied upon changes the behaviour of those relying on it
a measure adopted as a target	Goodhart [1975], Campbell [1979]	**2 — shared form**	stated for monetary and social indicators; the exposure is the same in kind
tight coupling; control trusted after its justification lapses	Perrow [1984], Vaughan [1996], Leveson [2011]	**3 — framing**	qualitative organisational and safety theory, used as framing and never as a formal result
scale-free vocabulary	Barabási and Albert [1999]	—	cited for the vocabulary only; **no scaling claim is made and Broido and Clauset [2019] is the reason none should be**
the modelling programme	Epstein [2008], Bruch and Atwell [2015]	—	the first enumerates the legitimate non-predictive uses of a model; the second is the checklist **none of the constructed illustrations here meets**

The failure mode this table exists to prevent is letting a category-2 borrowing carry a

category-1 result's authority. The numbers-equivalent, the variance identity, the scale-type restriction and the disclosure principle are category 1 and need no help. **The impossibility results are the most dangerous entries here**, because they are genuine theorems about adjacent problems and it would be easy to let one of them read as a proof about a named network. **None of them is.**

12.6 Postulates

#	statement	status
P1	participant shares are non-negative and sum to one	DEFINITION
P2	HHI $= \sum_i s_i^2$, $N = \text{HHI}^{-1}$	DEFINITION
P3	for n equal shares of $1/n$, HHI $= 1/n$	ARITHMETIC
P4	$s_{\max}^2 \le \text{HHI} \le s_{\max}$	ARITHMETIC
P5	the coordinating unit for each mechanism is as given in Section 3.2	POSTULATED (modelling choice)
P6	a single curator's decision reaches every entry on the list it curates	POSTULATED (H5)
P7	utility-as-function is an ordinal scale	DEFINITION
P8	expected freeze exposure is λw to first order	DEFINITION (modelling choice)
P9	the stated block capacity, transaction sizes and block interval	POSTULATED (described from the source; uncited)
P10	all constructed share vectors and dwell parameters	POSTULATED (constructed)

12.7 The chains

Chain A — the transfer and the ceilings (R1).

step	consumes	statement	status
A1	P1, P2	the index and its reciprocal are defined for any non-negative share vector summing to one	DEFINITION
A2	P3	a 35-member equal set gives HHI $= 1/35 = 0.028571$ and $N = 35.0$	ARITHMETIC
A3	P3	a 39-member equal set gives 0.025641 and 39.0	ARITHMETIC
A4	P10	$(2/7)^2 + 6(5/42)^2 = 0.166667$, so $N = 6.00$	ARITHMETIC

step	consumes	statement	status
A5	P10	the constructed staking vector gives 0.168400 and 5.94	ARITHMETIC
A6	A2–A5	the instrument returns numbers-equivalents on ledger share bases exactly as on jurisdictional ones — Adelman [1969], Hill [1973]	CITED
R1	A6	the concentration instrument transfers unchanged	DERIVED

Chain B — the curated-list collapse (R5).

step	statement	status
B1	five equal independent operators give HHI $= 5 \times 0.2^2 = 0.2000$, $N = 5.00$	ARITHMETIC
B2	a single curator holds the whole admission decision, so its share on that axis is 1.0	POSTULATED (P6)
B3	HHI $= 1.0^2 = 1.0000$, $N = 1.00$	ARITHMETIC
R5	B1, B3	the operation axis and the curation axis give 35.0 and 1.00 on one set

P6 is the one genuine postulate and it is the step to attack first. A curated set in which removing an entry requires more than one party's assent has a plural curation axis and the collapse does not occur; **that is precisely the falsifier stated for H5.**

Chain C — pool concentration (R3).

step	statement	status
C1	the constructed shares sum to 100%	ARITHMETIC
C2	$\sum s_i^2 = 1900/10{,}000 = 0.190000$	ARITHMETIC
C3	$N = 1/0.190000 = 5.26$ against ten nominal	ARITHMETIC
C4	$1/s^2_{\max} = 1/0.30^2 = 11.11$	ARITHMETIC
C5	$5.26 \le 11.11 \le 10$? **No** — the ceiling exceeds the nominal count, so both bounds bind and the tighter is $n = 10$	ARITHMETIC
R3	C2–C5	pool concentration reads 0.190000 and 5.26 at a $30/26$ top two

C5 is printed because it is a check a reader should make and this paper failed to make it in an earlier draft. $1/s^2_{\max} = 11.11$ exceeds the nominal count of ten, so on this vector the *count*

bound binds and the leading-share bound does not. **The leading-share ceiling is the informative one only when it falls below** n, and quoting it beside a ten-element vector is therefore a weaker statement than it looks. It is quoted anyway because it is the bound that survives when the number of pools is unknown, which is the usual case.

Chain D — dwell (R2).

step	statement	status
D1	5 s$/31{,}536{,}000$ s yr$^{-1} = 1.585 \times 10^{-7}$	ARITHMETIC
D2	$30/365 = 0.082192$	ARITHMETIC
D3	30 days $= 2{,}592{,}000$ s; $2{,}592{,}000/5 = 518{,}400$	ARITHMETIC
D4	expected exposure is λw, so the ratio of exposures is the ratio of dwells	DEFINITION (P8)
R2	D1–D4	dwell is the whole freeze exposure, and the ratio is $518{,}400$

Chain E — throughput (R6).

step	statement	status
E1	$10^6/226 = 4{,}424.8$ transactions per block	ARITHMETIC
E2	$4{,}424.8/600 = 7.37$ per second	ARITHMETIC
E3	$10^6/140 = 7{,}142.9$; $7{,}142.9/600 = 11.90$ per second	ARITHMETIC
R6	E1–E3, P9	base-layer throughput is approximately 7.4 to 11.9 tx/s at the stated parameters

R6's inputs are protocol constants described from the source manuscript and not cited, so the chain is arithmetic resting on an uncited description. **A reader who disputes the constants disputes the result, and Appendix B is where the constants go.**

12.8 Cross-checks performed

- Every index was recomputed by explicit sum of squares and reproduced to six decimal places: 0.028571, 0.025641, 0.005556, 0.037037, 0.142857, 0.166667, 0.190000, 0.200000, 0.500000, 1.000000.
- The bound $s_{\max}^2 \le \mathrm{HHI} \le s_{\max}$ holds on every vector used: $0.0900 \le 0.1900 \le 0.30$ on the pool vector and $0.081633 \le 0.166667 \le 0.285714$ on the anchor vector.
- The bloc-diverse vector was re-summed: $0.34^2 + 0.33^2 + 0.33^2 = 0.333400$, giving 2.9994, which is quoted as 3.00 and is **not** exactly three. That rounding is disclosed rather than presented as an exact result.
- The dwell ratio was computed two ways — as a ratio of year-fractions ($0.082192/1.585 \times 10^{-7}$) and as a ratio of seconds ($2{,}592{,}000/5$) — and the second is exact at $518{,}400$ while the first differs in the fourth figure because of the 365-day year. **The exact figure is** $518{,}400$ **and it is the one quoted.**
- The throughput arithmetic was checked at both transaction sizes and reproduces 7.37 and 11.90, which round to the quoted 7.4 and 11.9.

13. Limitations

Every figure in this paper is a ceiling implied by a described protocol constant or a constructed illustration, and none is a measurement of a running network. No live validator set, stake distribution, vote distribution, hashrate distribution, minimum-coalition coefficient or council roster was measured. **No dataset was collected.**

No as-of date, because nothing live was measured. A concentration figure computed from real holdings requires its source and its date; **a concentration measure without an as-of date is not a measurement.** Every share vector here is constructed or is a ceiling, so there is no source and no date to give, and Appendix A therefore makes the date and the source mandatory items of the reporting standard.

Every live quantity in this domain is time-varying and several move weekly. Any application of these instruments to a running network requires a live, dated, sourced distribution, and a figure carried from this paper into a live assessment is stale by construction.

Protocol specifications are not in this paper's bibliography. The consensus papers, whitepapers, improvement proposals and standards on which every network description rests are described from the source manuscript and **are not cited, because the bibliography contains no standards documents**. Every protocol constant used in Section 8 — block capacity, transaction sizes, block interval — inherits that status. Appendix B lists them. **This is a real gap and it is not closed here.**

The networks are de-branded and the de-branding loses information. A reader cannot verify a claim about "a federated-agreement ledger with a curated default list of about thirty-five validators" against a specific network without first identifying it, and this paper deliberately does not identify it, because identifying it would attach a characteristic the bibliography cannot support. **The trade is verifiability against unsupported attribution, and this paper takes the second cost rather than the first.**

The utility ordinal is a scalar reduction of a design space with at least four independent dimensions. Curzi [2026] names this habit once at framework level and it is inherited here rather than restated; Michell [1997] is the general charge and Luce and Tukey [1964] the test that would settle it, and **the conjoint conditions were not tested for this axis or for any other quantity in this paper.**

The bloc assignments are stipulated reference groups and not identified effects. Manski [1993]; the remedy adopted is stipulation with a defence [Brock and Durlauf, 2001], and **the price is that every bloc-level statement here is a modelling position.** Row 4 of the atlas is left undetermined because the documentation does not support even a stipulation.

Correlation inside a bloc is assumed, not estimated. Perfect correlation within a perimeter is the conservative end of a family, and no correlation was estimated from data anywhere in this paper.

The claim that no central-bank digital currency runs in production on a named public permissionless ledger is a statement about the absence of a documented counter-example as at the source manuscript's date, and not a proof of a universal negative. It is stated that way in Section 9.1 and recorded again here.

A central-bank remark on the degree of control a programmable sovereign rail confers is described in the source manuscript and is deliberately not quoted here, because its exact wording, date, forum and speaker designation are load-bearing and are not verified. **No sentence in this paper attributes words to any named person.** Appendix B records it as an item requiring primary-source confirmation before it is ever quoted.

The exit constraint on pool concentration is stated and not quantified. How quickly work redirects between pools bounds how much entrenched control a share represents, and no rate is measured or assumed. **The pool index is therefore an upper bound on entrenched control and this paper does not say how loose a bound it is.**

13.1 Rigour checks not met, listed by name

1. **No uncertainty is reported anywhere, because nothing is estimated.** Every figure is exact arithmetic on stipulated or described inputs. That is the absence of empirical content, not precision.
2. **Model comparison was not performed.** No alternative index was fitted, no out-of-sample test was run, no information criterion was computed and no matched null was constructed, because nothing here is fitted. The instrument is asserted on the strength of the identity chain in Section 12.7.
3. **The conjoint conditions were not tested** for the utility axis [Michell, 1997, Luce and Tukey, 1964, Krantz et al., 1971]. The axis is treated as ordinal, which is the honest downgrade; **it is not shown to be even that, since no ordering test was run against independent judges.**
4. **No distributional or scaling claim is made, so the Clauset–Shalizi–Newman procedure was not run** [Clauset et al., 2009]. **Any restatement reaching for a heavy-tail claim about hashrate, stake or holdings must run it first** [Stumpf and Porter, 2012, Broido and Clauset, 2019].
5. **The constructed illustrations fail the calibration-and-validation checklist** of Bruch and Atwell [2015]: none is calibrated, none is validated out-of-sample, and none has a sensitivity analysis beyond Section 12.3. They are demonstrations that a stated mechanism suffices to produce a stated shape, which is one of the uses Epstein [2008] enumerates and is not prediction.
6. **The instrument is exposed to its own adoption** [Goodhart, 1975, Campbell, 1979], and a published decentralisation metric is precisely the kind of measure a network can optimise instead of the property. The only defence offered is the reporting standard.
7. **The instrument is performative and its own bloc assumptions are exempted from that** [MacKenzie and Millo, 2003, Lucas, 1976]. **That is an inconsistency and it is not resolved here.**
8. **Expected value and time average are not the same thing** [Peters, 2019]. The dwell arithmetic is an expected exposure; a structure whose value is frozen once does not return to the mean, and the ensemble figure understates what a single participant faces.
9. **There is no conservation law anywhere in this paper.** The one genuine conservation is that shares sum to one [Mirowski, 1989]. **This is the objection the genre most often fails and this paper meets half of it.**
10. **Aggregation is unaddressed.** A single headroom parameter set is written for a heterogeneous settlement layer in Section 9.2, and Kirman [1992] establishes that such a representative parameter generally does not exist.
11. **No differential prediction is offered** [Durlauf, 2005]. Any correlation-aware concentration measure would produce the same collapses. **The contribution is the identification of the right share basis and the discipline of dating the number, not a uniquely discriminating forecast.**
12. **The leading-share ceiling was quoted beside a vector on which it does not bind**, and Section 12.7's step C5 prints the check rather than removing the quotation.
13. **The bloc-diverse effective count of** 3.00 **is** 2.9994 **and is rounded.** Disclosed in Section 12.8.
14. **The scalar-reduction objection is conceded and not answered.** Curzi [2026] states it once at framework level; this paper inherits it, points at it, and does not claim that naming it disposes of it.

14. Conclusion

This paper establishes six results in the form the rest of the corpus cites, and the following statements are the fixed text of those results.

R1 — the instrument transfers to ledgers unchanged. Measured on stated validator, coun-

cil or pool sets, effective counts run from a curated 35.0 validators (HHI $= 0.028571$) and 39.0 council members (0.025641) down to 6.00 organisations (0.166667) and 5.94 entities (0.168400). **An inverse-index effective count is indifferent to what its shares describe.**

R2 — dwell time is the whole freeze exposure. A five-second hop is 1.585×10^{-7} of a year (0.0000159%) against 0.0822 (8.2192%) for value rested thirty days — **a ratio of** $518{,}400$**.**

R3 — proof-of-work concentration is a pool measurement. A constructed distribution measures HHI $= 0.190000$ and 5.26 effective pools at a $30\%/26\%$ top two, against the ceiling $1/s_{\max}^2 = 11.11$ that the leading share alone imposes.

R4 — utility-as-function is an ordinal and not a score. A minimal-script ledger sits at $1/9$ and a general virtual machine at $9/9$; **the axis must be printed with the rank every time, the axis is not oriented, and programmability is denied as a quality gradient because a minimal ledger cannot express a bug it has no opcode for.**

R5 — the curated-list collapse. A single-curator trust set gives HHI $= 1.0000$ and **one** effective curator against five equal independent operators at 0.2000 and 5.00, **so counting validators the way entities are counted misses the binding constraint.**

R6 — base-layer throughput, stated to its parameters. Approximately 226 virtual bytes per transaction against approximately 10^6 virtual bytes per block on a ten-minute interval gives approximately 7.4 transactions per second; approximately 140 virtual bytes gives approximately 11.9.

Counter-result. The atlas refuses the reading that would make it useful as a ranking. **The utility ordinal is not a quality scale. The freeze control sits with the settlement instrument's issuer rather than on any rail, so rail diversification buys nothing against it. Participating in a pilot is not settling production value. And every live concentration figure is time-varying and is deferred to a dated register — a concentration measure without an as-of date is not a measurement.**

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every share vector in this paper is either an equal-weight ceiling implied by a protocol or governance parameter described from the source manuscript, or a **CONSTRUCTED** or **SYNTHETIC** illustration labelled as such at each point of use. **No live validator set, stake distribution, vote distribution, hashrate distribution, minimum-coalition coefficient or council roster was measured, and no figure in this paper is a measurement of any running network.**

All figures were produced by the scripts named in their captions (`figs/src/CCP-067_fig01.py` through `figs/src/CCP-067_fig05.py`), which import the shared figure style module and read their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

Protocol specifications, consensus papers, improvement proposals and standards documents are described from the source manuscript and are not cited to the bibliography, which contains no standards documents. Appendix B lists every affected claim. Any reader relying on a protocol constant used here must verify it against the governing specification.

The networks analysed are de-branded. Where a named system's characteristic could not be supported from the bibliography, the name is dropped and the functional description is kept. **No claim is made about any named network, product, organisation, institution or individual, and no allegation is made about any of them.**

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. No identifier, address, balance or transaction of any real party appears anywhere in this paper.

Legal, tax and investment advice

This paper is analytical and is not legal, tax, or investment advice. It analyses settlement networks as rails and settlement systems only. **No token, instrument, network or platform is recommended, valued, or characterised as an asset**, and no price, market capitalisation, return or yield figure appears anywhere. **No instrument is evaluated as a store of value.** The regulatory treatment of digital settlement instruments differs by jurisdiction, differs between regimes applying to the same facts, and changes without notice; every regulatory statement in this paper is carried from the source manuscript, is flagged in Appendix B, and requires confirmation with qualified counsel in the relevant jurisdiction before any reliance.

Appendix A. The reporting standard for a ledger concentration measurement

Any network measured under this instrument reports the following, and nothing less.

```
 1  share basis, stated in one sentence, with the mechanism that dictates it
 2  the participant table and its shares, summing to 100.0%
 3  the SOURCE of the distribution and its AS-OF DATE
 4  naive HHI, in fraction and in points, and the effective count
 5  the coordination-mechanism collapse, with the grouping rule stated
 6  the effective count after collapse, named as the binding figure
 7  the equal-weight ceiling implied by any protocol or governance cap,
    LABELLED AS A CEILING and never reported as a measurement
 8  the leading share and the bound 1/s_max^2, with a note where that
    bound exceeds the nominal count and therefore does not bind
 9  the utility position WITH THE AXIS NAMED BESIDE IT, and a statement
    that the axis is ordinal and unoriented
10  the dwell assumption used for any exposure figure
11  the diff against the previous measurement, and its date
```

Item 3 is the one that decides whether the rest is a measurement. Items 7 and 8 exist because a ceiling reported as a measurement is the commonest error in this domain, and Section 12.7's step C5 is the instance in which this paper caught itself. **Item 9 exists because a rank without its axis is not a measurement of anything.**

Appendix B. Claims requiring verification against primary sources

The following rest on protocol, governance or regulatory detail, and **none of them is cited to the bibliography, which contains no standards documents, no protocol specifications, no vendor documentation and no primary sources for any individual jurisdiction.** They are described from the source manuscript, and require confirmation against the governing documents before any reliance.

1. **The identity of the nine networks.** The source manuscript names them; this paper debrands them per Section 3.6. Reattaching a name requires verifying every characteristic attributed to it in Section 4.
2. **Every consensus-mechanism description**: the proof-of-work construction, the federated-agreement protocols, the asynchronous Byzantine-fault-tolerant gossip construction, the capped Byzantine-fault-tolerant proof-of-stake family, the delegated-proof-of-stake construction, the verifiable-delay-clock construction, and the proof-of-stake finality gadget. Confirm each against its published specification.
3. **Every set size and cap treated as a protocol or governance constant**: the curated default list of about thirty-five validators, the council cap of thirty-nine, the active-set cap of about one hundred and eighty, the fixed producer count of twenty-seven, and the tier-one anchor count of seven. **Each is the input to a ceiling, and a wrong set size is a wrong ceiling.** Several are set by on-chain governance and can be changed by a vote.
4. **The base-layer protocol constants of Section 8**: the block interval, the retarget window, the issuance cap, the halving interval, the block capacity in virtual bytes, and the transaction sizes used. **The throughput result rests entirely on these and they are uncited.**
5. **The scripting-language capability lists of Section 7.3**, both columns. Confirm the named primitives and the named absences against the specification; **the load-bearing claim is the absence of a general virtual machine.**
6. **The programmable-contract-module capability** attributed to the row-4 network, its compilation target, and its relationship to the settlement ledger it descends from. **Row 4's bloc and governance are left undetermined in print and must not be filled in from secondary coverage.**
7. **Every claim about pilots, trials, partnerships and platform products.** This paper names none, asserts none, and states that no central-bank digital currency runs in production on a named public permissionless ledger **only as an absence of a documented counter-example at the source manuscript's date.** Confirm the current record before repeating it. Confirm separately whether any commercial platform runs on a public ledger or on a private instance of a derived technology; that distinction is elided constantly.
8. **The central-bank remark on the degree of control a programmable sovereign rail confers.** It is described and **deliberately not quoted**. Confirm the exact wording, date, forum and speaker designation against the primary recording before quoting it anywhere. **Paraphrasing it as fact would be the self-inflicted error this paper's whole discipline exists to prevent.**
9. **The issuer-freeze fact** — that an issuer of a settlement instrument can immobilise balances at the request of an authority — is described generically and attributed to no named issuer, network or unit. Confirm the mechanism and its scope per instrument.
10. **Every liveness, halt and client-diversity claim.** Liveness is a statement about a history that continues; confirm against the current record and date it.
11. **All live distributions**: validator sets, stake shares, vote distributions, hashrate shares, minimum-coalition coefficients, staking-entity shares, execution-client shares and council rosters. **None is measured in this paper and all move.**
12. **The retail digital-currency issuance status of any jurisdiction.** Described by class in the source manuscript and not named here; status changes frequently.
13. **The rate at which proof-of-work capacity redirects between pools**, which bounds how much entrenched control a share represents and is stated in Section 6 without a number.

14. All constructed vectors, dwell parameters and hazard rates are stipulated for illustration and are not claims about any network.

All numerical results in this paper were computed from the stated inputs at full precision and rounded only for printing.

Chapter 68

Content Addressing, Time-Axis Concentration, and the Bounded Gain of a Nine-Capital Supplement

Abstract. Three results are joined by one idea: that an object's identity should follow from its content, and that a ledger's silence is a measurable quantity. Content addressing is stated exactly, with a one-byte difference in a settlement object producing a different identifier, and canonicalisation is placed underneath so that homoglyph variants fold to one skeleton before any hash is taken. The concentration instrument is then applied to a maturity profile rather than to entities: a constructed sovereign stack with 37.0, 11.0, 7.0, 16.5, 15.5 and 13.0 per cent in six buckets has a time-bucket HHI of 0.222050 and an effective spread of 4.50 buckets, with 55.0 per cent refinancing inside three years, against a flat ladder's 0.1667, 6.00 and 50.0 per cent. At a constructed repricing from a 2.00 to a 4.50 per cent coupon, the front-loaded stack adds 1.375 per cent of the stack per year against the ladder's 1.250 per cent, and the paper is explicit that the level obscures the real content, which is that 55 per cent reprices in one window rather than gradually. On a constructed nine-capital vector, 63.8500 per cent of national capital is represented on the monetary ledger and 36.1500 per cent is structurally off-ledger, of which 11.50 per cent is marked unmeasured; making the remainder visible widens the priceable base by at most 56.62 per cent, a multiple of 1.5662, and the honest gain lies strictly between zero and that ceiling.

Keywords: content identifier; canonicalisation; maturity concentration; refinancing wall; repricing delta; reserve drift; nine forms of capital; borrowing base

Series: D. Structure and Resilience **Paper:** CCP-068

1. Introduction

Three results are joined by one idea: that an object's identity should follow from its content, and that a ledger's silence is a measurable quantity. The first is a statement about identifiers, the second about a concentration axis nobody points the instrument at, and the third about what an accounting standard cannot represent. What links them is that in each case the honest measurement is of something the ordinary apparatus does not report.

The paper's four contributions are separable.

Content addressing, stated exactly, with canonicalisation underneath it. In a content-addressed system the name of a stored object is a hash of the object, so the identifier and

the integrity check are the same object: a reference cannot be held without the means of verifying it. A one-byte difference in a constructed settlement object produces a wholly unrelated identifier, and **the demonstration was recomputed independently for this paper and reproduces the source's digests and identifiers exactly.** But the same property launders a confusable identifier: a homoglyph substitution produces a different object with a cryptographically impeccable and permanently distinct name. **Canonicalisation must therefore sit *underneath* the addressing rather than beside it**, and with it in place the two visual twins fold to one identifier — also verified here from first principles.

The concentration instrument rotated onto the time axis. A constructed six-bucket maturity stack of $37.0/11.0/7.0/16.5/15.5/13.0$ per cent has a time-bucket index of 0.222050, an effective spread of 4.50 buckets, and 55.0% refinancing inside three years, against a flat ladder's 0.1667, 6.00 and 50.0%. At a constructed repricing from a 2.00% to a 4.50% coupon the front-loaded stack adds 1.375% of the stack per year against the ladder's 1.250%. **The level obscures the real content, which is that 55 per cent reprices in one window rather than gradually.**

The completeness gap, computed and bounded. On a constructed nine-form capital vector, 63.8500% of national capital is represented on the monetary ledger — financial 42% and material 18% fully, intellectual at a stipulated 35% recognition giving 3.85% — and 36.1500% is structurally off-ledger, of which 11.50% is marked **UNMEASURED**. Making the remainder visible widens the priceable base by **at most** $36.15/63.85 = 56.62\%$, a base multiple of 1.5662. **The honest gain lies strictly between zero and that ceiling and cannot be stated precisely, because the valuation the taxonomy does not supply is exactly what would be needed to state it.**

And a fourth thing the paper refuses. The 56.62% is a ceiling on constructed inputs and not a promise; **visibility is not pledgeability**; the valuation of non-financial capital is unsolved and is named here as the hardest problem in the framework; and the composed architecture is an architecture and not a system.

Why these sit in one paper. They are one measurement discipline applied at three scopes. Curzi [2026] measures the completeness gap on the same constructed nine-capital vector — 63.850000% represented, 36.150000% structurally unrepresentable — and this paper derives its national reading from that result. Curzi [2026] supplies the canonicalisation and typed-parsing controls the addressing layer requires, where typed-field parsing collapses the withdrawn-code parsing-ambiguity false-positive share from 88.99% to 0.00% against a naive detection precision of 11.01% — a control right one time in nine. Curzi [2026] supplies the constraint that makes publication of all of it an engineering decision rather than an ethical one: **channel separation is strengthened by disclosure because its effectiveness is entirely independent of whether anyone knows it is in place, so every element must be publishable and non-disclosure is refused as a design premise, not merely discouraged.** And Curzi [2026] establishes that the concentration instrument transfers to any share basis unchanged — from a curated 35.0 validators (HHI $= 0.028571$) and 39.0 council members (0.025641) down to 6.00 organisations (0.166667) and 5.94 entities (0.168400) — which is the licence this paper uses when it points the same index at maturity buckets instead of at entities.

2. Hypotheses

Six hypotheses with their falsifiers, and three statements that are not hypotheses and are labelled as such. Platt [1964]'s standard applies.

H1 — the fusion hypothesis. In a content-addressed system the reference and the integrity check are one object, so they cannot be separated by any configuration. *Predicted:* changing one byte of a constructed settlement object yields two identifiers sharing no structure, and rehashing the received bytes and comparing to the name is the only check needed. *Falsified by:* a content-addressing scheme in which a valid identifier can be held for content that does not hash to it, or in which an integrity value can be stored, stripped or desynchronised separately from the name.

H2 — the laundering hypothesis. Content addressing over un-canonicalised data converts

a visual ambiguity into two permanent, cryptographically valid, distinct names; the same addressing over the canonical form collapses them to one. *Predicted:* two account strings differing only by a homoglyph produce two distinct identifiers over raw bytes and one identical identifier over the canonical form. *Falsified by:* a canonical form under which the twins do not fold, or a raw-byte addressing under which they collide.

H3 — the time-axis hypothesis. Maturity concentration is measurable by the same index used for entity concentration, and the effective spread it returns is more informative than the total or than the near-window share alone. *Predicted:* a constructed front-loaded stack gives HHI $= 0.222050$ and an effective spread of 4.50 against a flat ladder's 0.1667 and 6.00, while the two stacks' near-window shares differ by only five points. *Falsified by:* two stacks with the same effective spread and materially different repricing behaviour under the same rate path, which would show the index misses the exposure it claims to measure.

H4 — the speed hypothesis. Front-loading changes when a repricing arrives more than it changes how much arrives. *Predicted:* at a constructed 2.50 percentage-point repricing delta the front-loaded stack adds 1.375% of the stack per year against the ladder's 1.250% — a difference of 0.125 points — while 55% of the front-loaded stack reprices inside one three-year window against the ladder's three equal instalments. *Falsified by:* a case in which the level difference dominates the timing difference in fiscal effect, which would make the near-window share the right summary and the index redundant.

H5 — the representation hypothesis. The monetary ledger's silence about non-monetised value is structural rather than a defect of practice, and is therefore measurable as a fixed fraction of a stated capital vector. *Predicted:* on the constructed vector, 63.8500% represented and 36.1500% off-ledger, of which 11.50% admits no accepted measurement basis at all. *Falsified by:* an accounting method that recognises an unpriced item as an entry while retaining the closure property, which would show the gap is a choice rather than a consequence.

H6 — the bounded-gain hypothesis. Making off-ledger capital visible widens the priceable base by a bounded amount whose upper limit is computable and whose actual value is not. *Predicted:* the ceiling is $36.15/63.85 = 56.62\%$, a base multiple of 1.5662; the realised gain is strictly positive and strictly less, and cannot be stated without a valuation method for the non-financial forms. *Falsified by:* a demonstration that a creditor prices visible-but-unvalued capital at exactly zero, which would collapse the lower bound onto zero and make the reform inert.

Not hypotheses, and labelled as such.

- *The nine-form taxonomy* is a **classification**, not a finding. It names what a ledger cannot see and it does not value it; nothing in this paper treats the naming as evidence for anything.
- *The choice of maturity buckets* is a **modelling choice**. The index is not invariant to it, and Section 10.4 states exactly what that costs.
- *The supplementary-standard proposal* is a **design**, adopted by no body and carrying no legal force. It is stated so that it can be evaluated, not asserted as established.

3. Method

3.1 The content identifier, stated exactly

In a location-addressed system the name says *where* the data is, and the data at that location can change while the name stays the same — which is why a link rots, a file is edited in place, and a row is updated without anyone downstream knowing. In a content-addressed system the name says *what* the data is, because the name is computed from the content. **Change one byte and the address changes.**

The self-describing form used here wraps a hash in enough metadata to survive changes of hash function and encoding:

```
identifier  =  <base-prefix> <version> <content-codec> <multihash>

multihash      =  <hash-function code> <digest length> <digest>
content-codec  =  what the bytes ARE (raw leaf, structured node, linked data)
base-prefix    =  how the identifier itself is text-encoded
```

Three properties follow and each is load-bearing. Recording *which* hash function was used, rather than assuming one, means the scheme survives the retirement of a function without ambiguity — the same discipline that requires a broken function to be refused at the point of naming. Recording what kind of object the bytes are lets a consumer know whether it holds a raw leaf or a node in a hash-linked graph. And recording the text encoding of the identifier itself makes an identifier written into a message unambiguous about how to turn its characters back into bytes — which is Curzi [2026]'s encoding-declaration discipline appearing as a designed-in property of the identifier rather than as a boundary rule an integrator must remember to apply.

The property that matters is the fusion. The identifier and the integrity check are the same object: to check that the bytes received are the bytes the name promised, re-hash and compare, and there is no separate step, no separate field, and no separate channel in which an integrity value could be stored, forged, stripped, or desynchronised from the data. **The name carries its own proof.**

That fusion is the complement of Curzi [2026]'s channel separation and is stated here as a complement rather than as an equivalence. Channel separation solves the impersonation problem by lifting control off the data plane so that data cannot forge control; content addressing solves it by fusing the reference to the data so that data cannot be altered under a valid reference. **Both reach the one end that matters — the data cannot be tampered with while a valid control path to it survives — by opposite means.** In the separation case the property must be built and maintained and a lapse reintroduces the exposure; in the addressing case there is nothing to maintain, because the property is a consequence of the arithmetic rather than a configuration. **This correspondence is this paper's own analytical framing and is badged as an argument, not as a standard**, per the analogy discipline of Section 10.5.

3.2 Why canonicalisation must sit underneath the addressing

There is a trap in the primitive, and it is the seam where the whole composition either holds or fails. **Content addressing does not fix confusable identifiers. Left to itself it launders them.**

If the identifier is the hash of the content, a homoglyph substitution in a name does not produce a suspicious near-duplicate that a reconciliation might catch. It produces a *different piece of content*, with a different, cryptographically impeccable address, indistinguishable from a legitimate distinct entity. **The addressing layer inherits the defect of the layer beneath it in its worst form: the ambiguity becomes permanent and cryptographically endorsed.**

The remedy is ordering, not addition. Canonicalise, then address; and retain the hash of the raw octets alongside, so that the fact that a twin was folded is not lost. **The canonical form gives one address; the raw-byte hash still witnesses that a substitution occurred, and that pair is what distinguishes benign normalisation from injected substitution.** Curzi [2026] specifies the canonical form and the typed parsing this rests on, and Curzi [2026] states why the witness must be computed over the raw octets and held on a plane separate from the data it guards — **a witness that is a function of the canonical form changes exactly when the data changes and never otherwise, which destroys precisely the evidence a witness exists to preserve.**

3.3 Which effective count, and the licence to rotate the axis

Every effective count in this paper is $N = \mathrm{HHI}^{-1}$, the reciprocal of the sum of squared shares. That quantity is the Hill number of order two [Hill, 1973, Jost, 2006]: the number

of equally-sized units that would produce the observed index. It is the reciprocal of Simpson [1949]'s index, the effective number of parties of Laakso and Taagepera [1979], and the numbers-equivalent reading attributed to Adelman [1969]; the index is Herfindahl [1950] and Hirschman [1964], and the bands at which competition authorities read it are policy choices that have moved [U.S. Department of Justice and Federal Trade Commission, 2023, 2010]. The exponential of Shannon entropy [Shannon, 1948] is the same family at order one, and HHI^{-1} is also the exponential of the Rényi entropy of order two [Rényi, 1961].

The corpus's other effective count is not a Hill number and is not used here. The affine form $N_{\text{eff}} = 1+(N-1)u$ is a linear interpolation between 1 and N; it is not a numbers-equivalent of any index, and no literature in ecology, economics, political science or information theory contains it. **Wherever this paper says "effective count" or "effective spread" it means** HHI^{-1}, and the affine form appears nowhere.

The licence to rotate the axis is Curzi [2026]'s and it is bounded. That paper establishes that the index transfers to any share basis unchanged, because its derivation does not depend on what the shares are shares of. **What transfers is the arithmetic and the numbers-equivalent property. What does not transfer is any claim that the units on the new axis are the same kind of thing as the units on the old one.** Maturity buckets are not jurisdictions, and this paper aggregates nothing across the two scopes. That badge is Curzi [2026]'s **POSTULATE — MODEL FORM ONLY**: identical structure at several scopes is a statement about the chosen model, not a claim of kinship or of cross-scope aggregation, and it is a structurally parallel result in another domain rather than a re-derivation.

The bloc collapse is the variance identity of Markowitz [1952]: for shares w_i and unit-variance shocks X_i, $\mathrm{Var}(\sum_i w_i X_i) = \sum_i \sum_j w_i w_j \rho_{ij}$, which is $\sum_i w_i^2$ under independence and $\sum_b W_b^2$ under perfect within-group correlation. Leinster and Cobbold [2012] and Rao [1982] give the general similarity-sensitive form. That correlations rise under stress is measured [Longin and Solnik, 2001, Ang and Chen, 2002]; the contrary reading — that diversity–stability correlations can arise from statistical averaging with no mechanism — is Doak et al. [1998]'s, and May [1972] shows increasing diversity can destabilise. **Portfolio arithmetic is not by itself evidence of a mechanism**, and none is claimed here.

The identification problem is stated where the collapse is introduced. Any grouping of exposures into perimeters is a reference-group specification, and the framework's reflexive law is a linear-in-means specification whose endogenous effect Manski [1993] proves is not identified from observational data. **This paper adopts no bloc grouping at all on the maturity axis — the buckets are exogenous calendar intervals, not stipulated coordination groups** — so the reflection problem does not bite on H3 or H4. **Where it does bite is on the coupling claim of Section 6.3**, which asserts that a drifting reserve share thins the marginal bid at exactly the maturity buckets where refinancing concentrates. That is an assertion about a relationship between two populations of holders and it is **an analytical inference, not an identified effect**; the remedy adopted is stipulation with a defence [Brock and Durlauf, 2001], and the price is that the coupling is a modelling position.

3.4 What a supplementary account may claim

The monetary ledger requires that every entry carry an amount in a unit of account. Anything without a price cannot be an entry, so value that is real and operative but that no transaction ever priced **is not recorded as zero — it is not recorded at all**, which is a different and worse thing, because zero is a number a reader can question and silence is not. Curzi [2026] establishes the gap quantitatively; this paper carries it to national scale and asks what follows.

The consequence is a pricing consequence and it is stated narrowly. A creditor prices what it can see, and it reaches for the monetised aggregates because those are the quantities the standard produces. Off-ledger capital does not enter the pricing — not because a creditor judges it worthless but because no statement reports it. **That is a failure of the instrument that feeds the judgement, not of the judgement.**

Three things the account must not do, stated before the results rather than after.

It must not average across the measurable and the unmeasurable. Some of the nine forms admit no generally accepted measurement basis. Summing them into a single headline assigns the unmeasured forms an implicit value by including them in a total, which is a fabrication wearing an aggregate's clothes. **This is a scale-type violation before it is anything else**: an entry with no measurement basis is not even ordinal, and Stevens [1946] fixes what operations a scale licenses, Krantz et al. [1971] what must be true before a quantity can be said to exist, and Hurlbert [1971] the general discipline that a composite index carries no content without its formula attached. **The standard therefore marks unmeasured forms UNMEASURED, and strikes no total.**

The habit this refusal guards against is named once at framework level and is inherited rather than restated. Curzi [2026] states it as one objection appearing in two domains rather than two objections: where the source description is a pattern across several jointly varying features, the framework replaces it with a single number and does arithmetic on the number, and the validity of the arithmetic does not protect the reduction that produced it. **A single national-capital total across nine heterogeneous forms is that habit at its largest scale**, and the standard's refusal to strike one is this paper's concession to the objection rather than an answer to it.

It must not substitute a proxy for an unmeasurable quantity. Once a proxy for an unmeasurable becomes a target it stops measuring what it stood for [Goodhart, 1975, Campbell, 1979] — and a national account full of soft self-reported indices is that failure at the scale of a state. The governance and valuation literatures for the non-financial forms exist and must supply any figure the standard reports: Ostrom [1990] and Ostrom [2010] for common-pool resources and polycentric governance, Cox et al. [2010] for the meta-analytic audit of those principles, and Hardin [1968] as the claim Ostrom refuted and which must be cited only as that pair. **For social capital specifically the literature is contested rather than merely thin**: Putnam [2007] finds that diversity in the short run reduces trust and civic participation, and Abascal and Baldassarri [2015] is the rebuttal; Fukuyama [2001] and Gambetta [1988] set out how unstable the construct's definition is. **A standard that entered a single "social capital" figure would be entering a number the literature does not agree how to compute.**

It must not present visibility as pledgeability. Making a form of capital appear on a statement does not make a creditor lend against it and certainly does not make it collateral. Much of the newly visible capital is by its nature unpledgeable and unrealisable as security. **What visibility does is narrower: it moves a quantity from "unknown, therefore priced at zero" to "known, of contested value", and a quantity of contested value is priced above zero by any rational lender.** Hansmann and Kraakman [2000] is the reason the distinction is sharp — asset partitioning, not disclosure, is what makes an interest enforceable — and Pistor [2019] is the account of how legal coding manufactures enforceability and its sharpest critic at once.

And one thing the standard genuinely is. A standard decides what gets counted, and what gets counted decides what gets managed, financed and argued over, because a quantity that appears on no statement is one no minister defends in a budget and no analyst prices in a spread. Scott [1998] is the canonical account of legibility as a form of state power and of what it costs; Hacking [1995] states the general mechanism by which a classification interacts with what it classifies so that the referent moves; and MacKenzie and Millo [2003], Callon [1998] and Ferraro et al. [2005] document the same loop in economics, where a description relied upon changes the behaviour of those relying on it. **A supplementary standard would be performative in exactly that sense, and this paper counts that as an argument for stating its limits loudly rather than as an argument against building it.**

3.5 What is constructed, what is described, and what is cited

Every numerical result in this paper rests on constructed or stipulated inputs. The settlement objects, the account identifiers, the six-bucket maturity distribution, the repricing delta, the private refinancing schedule, the reserve share and its drift rate, and the nine-form capital vector are all invented for illustration. **No dataset was collected. No figure here is a mea-**

surement of any real issuer, nation, entity or holding, and no allegation is made about any of them.

Statistical series, statutes, accounting standards and legislative material are not in this paper's bibliography and are not cited. The reserve-composition series, the outstanding-debt series, the actual maturity profile of any sovereign, the recognition rules of any accounting framework, the storage and identifier specifications underlying Section 3.1, and the text of any proposed legislation are described from the source manuscript. **The bibliography contains no statistical series, no standards documents, no statutes and no legislative drafts.** Appendix B lists every affected claim.

Named jurisdictions, issuers and instruments are consequently de-branded, following the convention of Curzi [2026]: where a named party's characteristic is relied upon and cannot be supported from the bibliography, the name is dropped and the characteristic is kept — *a reserve-issuing sovereign, a proposed capital-anchored instrument, the coordinating perimeter*. **The functional descriptions are precise enough to identify what is measured and carry no claim about any named party.**

What *is* cited, and genuinely supports what it is attached to: the concentration literature of Section 3.3; the portfolio and stability literature of the same section; the measurement-theory literature of Section 3.4; the commons, trust and social-capital literatures named there; the legibility and performativity literature; the distributed-systems results of Section 4.3; and the identification, scaling and modelling literature of Sections 3.3, 10 and 11.

4. Results: H1 and H2 — content addressing, and the layer that must sit under it

4.1 H1 — one byte, one different address: supported and independently reproduced

Two **SYNTHETIC** settlement objects differing in a single character:

```
A  "settlement-object-0042: amount=2000 ccy=643"
B  "settlement-object-0042: amount=2000 ccy=644"

sha-256(A)  31feca3443b85c9c...
sha-256(B)  ed1b3121817f07eb...
            differ from the first byte onward

content identifier (version 1, raw codec, base32)
A  bafkreibr73fdiq5ylsofx5zhy36oowdxtdoqew5injlbcmmfeia62vf4d4
B  bafkreihndmysdal7a7v2373iici34mere36rldsktdikxtigomej37avu4
```

A single digit changed in a currency field produces two content addresses sharing no structure, because the digest is an avalanche function of every input byte. **The consumer who asked for A and is handed B does not need a separate integrity check to notice: the address does not match, and the mismatch is the check. The reference and the verification are one act, and there is no configuration under which they come apart. H1 is supported.**

These four values were recomputed from first principles for this paper — SHA-256 over the stated octets, wrapped as a version-1 raw-codec multihash and rendered in lower-case base32 — and reproduce the source manuscript's digests and identifiers exactly. That is the strongest verification status any figure in this paper carries, and Section 10.8 records it as the one independent reproduction performed.

What the demonstration does not establish. The specific digests are properties of the stated byte strings and carry no meaning beyond exhibiting the avalanche property. **A production tool that wraps small payloads in a framing layer before hashing would yield different but**

equally deterministic identifiers; the property demonstrated — one byte, one different address — holds under either framing, and the paper's claim is the property and not the digest.

4.2 H2 — the laundering trap, and the ordering that closes it

The same primitive turned against the system. Two account identifiers differing only by a homoglyph — a Latin letter replaced by a visually identical letter from another script:

```
clean  "ACCT-0042-CORP"                    (Latin O)
twin   "ACCT-0042-C<Cyrillic O>RP"         (renders identically)

identifier over RAW bytes
clean  bafkreidn5fn5dsd7j3bnyixc26mb3nyp7yemasvms4cvffwnktjv7lma3y
twin   bafkreidme3ihjm4z5fx5iqllmjute75r4wcii3utfx2ncchkd5hzlr4ynu
       two distinct, permanent, cryptographically valid names for one visual identity
```

Now put the canonical form underneath the addressing:

```
canon("ACCT-0042-CORP")            = ACCT-0042-CORP
canon("ACCT-0042-C<Cyrillic O>RP") = ACCT-0042-CORP    (folded to one skeleton)

identifier over canon(clean)  bafkreidn5fn5dsd7j3bnyixc26mb3nyp7yemasvms4cvffwnktjv7lma3y
identifier over canon(twin)   bafkreidn5fn5dsd7j3bnyixc26mb3nyp7yemasvms4cvffwnktjv7lma3y
       one identity, one address
```

H2 is supported, and both identifiers were likewise recomputed independently and reproduce the source exactly. Content addressing over un-canonicalised data converts a visual ambiguity into two permanent and distinct names; content addressing over the canonical form collapses the twins onto one. **The canonicalisation layer must sit underneath the addressing rather than beside it, and the raw-byte hash must be retained alongside so that the fact of a fold is not lost.**

4.3 Necessary, not sufficient — for every layer at once

The claim must be bounded before it is composed with anything, and the bound is the same at every layer. Canonicalisation defeats encoding-level evasion and does nothing about semantic fraud. A byte-level hash defeats silent edits and does nothing about a correctly signed lie. Content addressing defeats reference-without-verification and does nothing about content that is verifiably, impeccably false. Per-object hash chains defeat silent history-tampering and do nothing about a false entry faithfully recorded.

A tamper-evident, content-addressed, canonically named record of a false transaction is a tamper-evident, content-addressed, canonically named false transaction, and the architecture will preserve it faithfully, immutably, and forever.

Two further bounds are structural rather than rhetorical and both are cited. **Per-object hash chains provide tamper-evidence and write-parallelism; they do not provide consensus, total order or double-spend prevention**, and an ordering layer is required — the requirement has been formal since Lamport et al. [1982] fixed the number of genuinely independent parties needed to tolerate f arbitrary faults at $3f + 1$, and Fischer et al. [1985] and Gilbert and Lynch [2002] bound what any such layer can promise. **Nor do they provide resistance to one party presenting as many**, which Douceur [2002] proves impossible without a trusted certifying authority and which makes an unaudited count of participants worthless. **Neither layer is supplied by anything in this paper, and the composition is a substrate rather than a complete distributed ledger.**

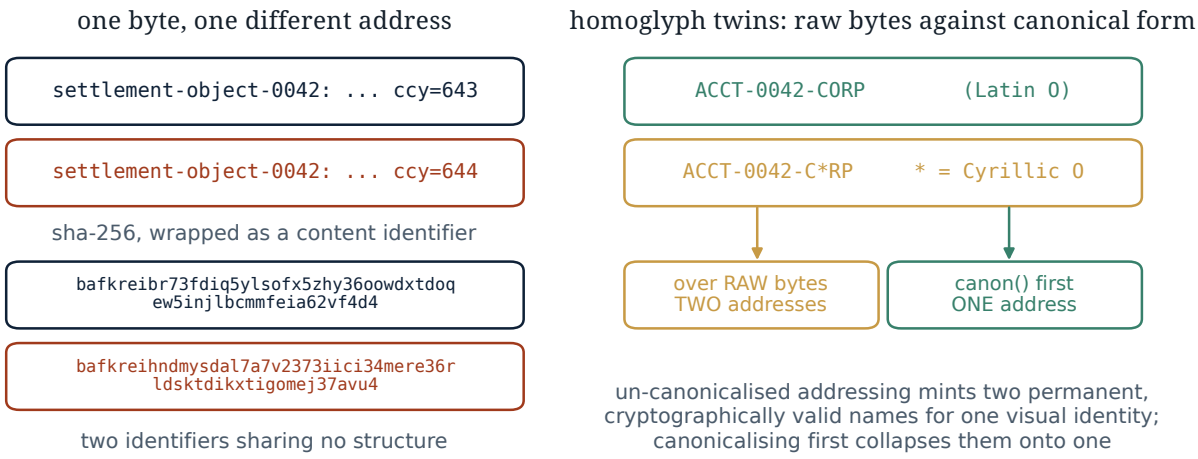


Figure 68.1: **One byte changes the identifier; one canonicalisation collapses two identifiers into one.** Left: two settlement objects differing in a single character produce content identifiers sharing no structure. Right: two account strings differing only by a homoglyph produce two permanent, cryptographically valid identifiers over raw bytes, and one identifier over the canonical form. All four identifiers were recomputed from first principles for this paper and reproduce the source exactly. SYNTHETIC strings; script `figs/src/CCP-068_fig01.py`.

A third bound is about registries and it is where cryptographic certainty ends. An identifier code is only as trustworthy as the registry binding it to reality, and **a hash computed over a false binding is a cryptographically impeccable name for a falsehood.** Content addressing guarantees that the bytes fetched are the bytes that were named; it guarantees nothing about whether the entity code inside those bytes belongs to the party the message claims. **Curzi [2026w]'s two-column refinement is the general form of the same discipline and is the sharpest available statement of it: the evidence column carries what could have come out otherwise, and the precedent column carries only that the idea was independently arrived at and preserved — two sources agreeing does not double the evidence when only one of them is evidence.** A cryptographic witness and a registry assertion are exactly such a pair: the witness is evidence about the octets and the registry is precedent about the world, and stacking them does not produce two independent confirmations of the same fact.

5. Results: H3 and H4 — the time axis

5.1 The rotation, and why the total is the least informative number

A stock of debt is not a single object. It is a schedule — a distribution of principal across future dates at which each tranche must be repaid or refinanced. **The total is a sum, and it is close to the least informative number available**, for the same reason a count of entities is the least informative thing that can be said about a structure.

The mechanism is stated without a single asserted figure about any real issuer. A fixed coupon is fixed at issuance: a change in market rates does not change what is paid on debt already outstanding. **The exposure to a rate change is therefore not on the outstanding stock; it is on the refinancing of the stock, and it arrives in order of maturity — the soonest first, the largest near-term tranche fastest of all.** The vulnerability is a function of how concentrated the maturity schedule is at the short end, and that is measurable by the index of Section 3.3 with maturity buckets as the share basis.

The front-loading is not a mistake. Short-dated debt is cheaper to issue when the yield

curve slopes upward, easier to place in size, and defers locking in a higher long rate. **Each of those is an individually reasonable issuance decision, and aggregated over years they front-load the schedule** — the same shape as an estate that becomes single-perimeter through a sequence of individually sound jurisdiction choices. Nobody builds the concentration on purpose; it is what results when each decision optimises its own cost and none measures the resulting distribution.

5.2 H3 — the time-bucket index: supported

CONSTRUCTED throughout. The distribution below is invented to demonstrate the instrument. It is not any issuer's maturity profile.

maturity bucket	share of stack	share^2
0–1 year	37.0%	0.136900
1–2 years	11.0%	0.012100
2–3 years	7.0%	0.004900
3–5 years	16.5%	0.027225
5–10 years	15.5%	0.024025
10–30 years	13.0%	0.016900
total	**100.0%**	**0.222050**

```
refinancing due within 3 years  = 37.0 + 11.0 + 7.0    = 55.0%
time-bucket HHI                                        = 0.222050
effective maturity spread       = 1 / 0.222050 = 4.5035 -> 4.50 buckets
```

Against the flat ladder — one sixth in each bucket, the maximally diversified stack for six buckets:

stack	HHI	effective spread	due within 3 years
constructed	0.222050	4.50 buckets	55.0%
flat ladder	0.166667	6.00 buckets	50.0%
gap		-1.50	$+5.0$ points

H3 is supported. The nominal count is six and the effective count is 4.50; the gap is the front-loading. **And the near-window share alone would have missed it**: both stacks have roughly half their maturities inside three years, and the difference the front-loading buys is not how much comes due soon but how concentrated the near-term maturities are. **The laddered stack's near-term maturities are three equal small slices; the constructed stack's are one large slice and two small ones, and the large slice is the exposure.**

5.3 H4 — the repricing arithmetic, and why speed is the content

Take a **CONSTRUCTED** rate shock: debt maturing in the near buckets was issued at an average coupon of 2.00% and refinances at 4.50%, a repricing delta of 2.50 percentage points. The added annual interest is the delta applied to the share that reprices:

```
front-loaded stack, 55% within 3 yr:   0.55 x 0.0250 = 0.013750 = 1.375% of the stack/yr
flat ladder,        50% within 3 yr:   0.50 x 0.0250 = 0.012500 = 1.250% of the stack/yr
difference                                                      0.125 points/yr
```

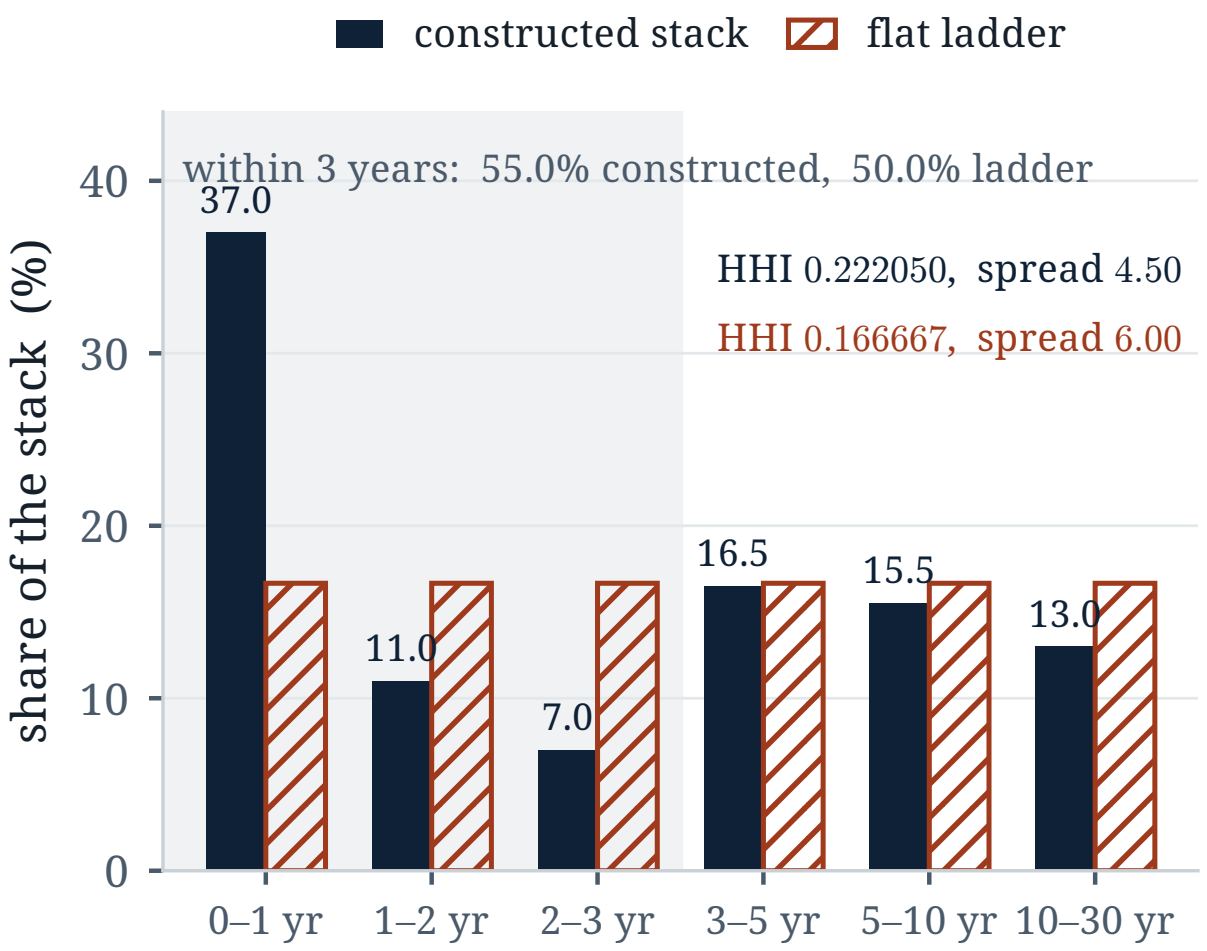


Figure 68.2: **The near-window share hides what the index reports.** A constructed front-loaded stack and a flat ladder both have roughly half their principal maturing inside three years — 55.0 against 50.0 per cent — but the front-loaded stack concentrates that half into one dominant bucket, giving a time-bucket index of 0.222050 and an effective spread of 4.50 against the ladder's 0.166667 and 6.00. CONSTRUCTED distributions; script `figs/src/CCP-068_fig02.py`.

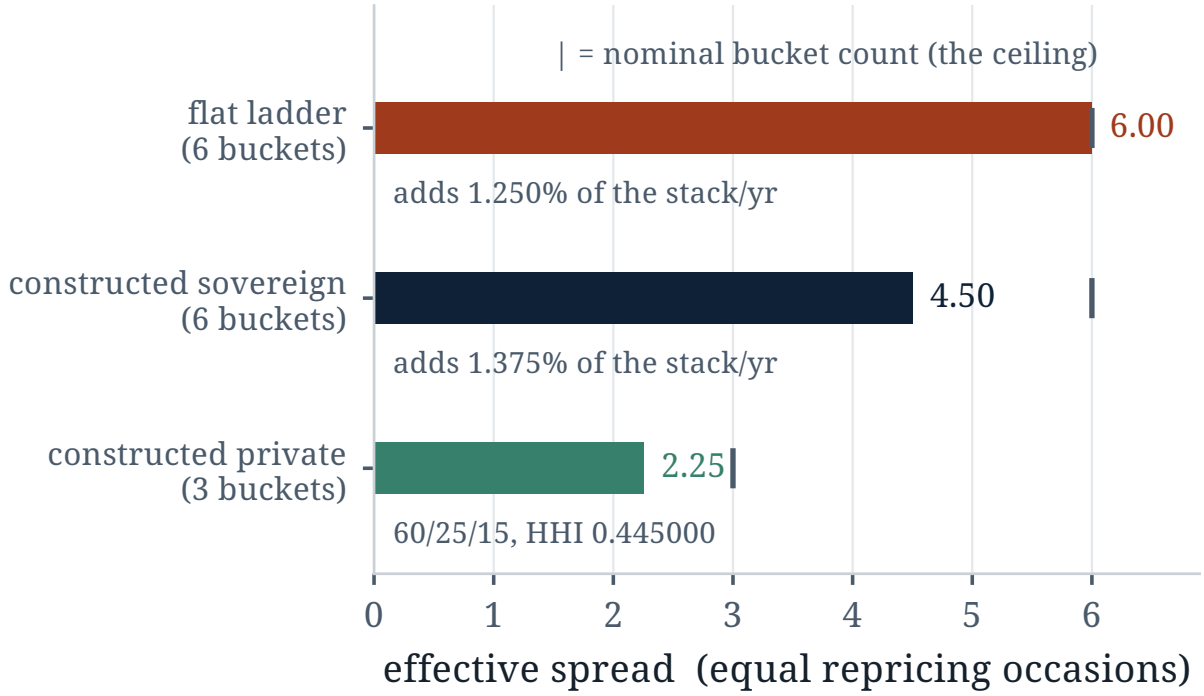


Figure 68.3: **Front-loading changes when the repricing arrives more than how much arrives.** At a constructed 2.50-point repricing delta the front-loaded stack adds 1.375 per cent of the stack per year against the ladder's 1.250 — a gap of 0.125 points — while its effective spread of 4.50 against the ladder's 6.00 shows the same repricing arriving in fewer, larger occasions. A constructed private refinancing schedule at 60/25/15 per cent is more concentrated still, at 2.25 effective occasions. CONSTRUCTED parameters; script `figs/src/CCP-068_fig03.py`.

The level difference is small and this paper will not inflate it. What the arithmetic actually says is about timing. The front-loaded stack does not reprice gradually: it reprices 55% of itself inside three years, so the interest-cost increase arrives fast and early. The laddered stack reprices a steady sixth per bucket, absorbing the shock in even instalments. **Concentration on the time axis converts a rate change from a gradual cost into a fast one, and speed determines whether an increase can be met from flow or must itself be borrowed. H4 is supported: the front-loaded stack does not pay more in the end; it pays more, sooner, and reprices its largest share first.**

The same instrument at private scale. A constructed facility refinancing 60% within one year, 25% at one to two years and 15% at two to three gives $0.360000 + 0.062500 + 0.022500 = 0.445000$ and an effective spread of 2.2472, quoted as 2.25. **The private stack is far more time-concentrated than the constructed sovereign one — 2.25 occasions against 4.50 — because a rolling facility rolls faster than a bond ladder**, and the remedy is identical in form: lengthen the financing or reduce reliance on the rolling short leg, at a cost paid now against a shock deferred. Curzi [2026] owns the trigger mechanics that sit alongside this exposure; a drawdown triggers a call and a rate shock reprices the roll, **and both reach the short end first.**

5.4 What the instrument does not capture, stated before anyone asks

The time-axis index captures concentration of maturities. **It does not capture the level of rates, the direction of rates, the demand for the debt, the currency of issuance, or the capacity to service it.** It is one axis, and a sovereign position is at least four — maturity concentration, holder concentration, currency of issuance, and the reserve status of that currency — in exact

parallel to the four axes an entity structure is exposed on. **A stack can be beautifully laddered and still be in trouble if its holders are concentrated and heading for the exit, and the maturity index will report the laddering and see none of the rest.**

Three refusals belong here rather than in the limitations, because a reader could otherwise take them as concessions rather than as the shape of the claim.

This is not a prediction that a rate shock happens. The repricing scenario is constructed and the paper asserts nothing about whether or when rates move. It measures what a front-loaded stack does *if* they move, which is a conditional and the only kind of statement made here. **The conditional form is not a hedge but the honest shape of the claim**: a maturity schedule is a distribution of future obligations, and a distribution has no truth value until a rate path is laid across it. The same stack is resilient under a falling path and exposed under a rising one; the front-loading does not choose the path, it determines only how much of the stack the path reaches and how fast. **The conditional is what the instrument measures; the prediction is what it cannot.**

This is not a claim that any issuer cannot refinance. For a sovereign borrowing in its own unit of account the near-certainty that maturing debt can be rolled is exactly what makes the exposure a *repricing* exposure rather than a *rollover* one. **The vulnerability is the price of the roll, not the possibility of it**, and confusing the two is how a collapse narrative reintroduces itself.

And this is not a number about any real issuer. It is a number about a constructed distribution with 55% inside three years.

6. Results: H5, H6, and the reserve-drift timescale

6.1 R6 — reserve drift as a timescale illustration

The first compromise of a reserve-issuing position is a slope, not a cliff, and the honest way to show that is to state a constructed slope and read the timescale off it.

```
CONSTRUCTED share of allocated reserves 58.0%, drifting 0.5 percentage points a year

year +0     58.0%       no threshold crossed
year +5     55.5%       no threshold crossed
year +10    53.0%       no threshold crossed
year +16    50.0%       crosses the symbolic half, on a constructed slope
year +20    48.0%       and no rival at scale to receive the released share
```

A linear extrapolation of a slope of this kind crosses no threshold for a decade and a half, and a process that crosses no threshold for fifteen years is not an event. The figures are invented to show the timescale and forecast nothing.

Three structural features make “drift” the right word. **It is slow**, and a reallocation happening in fractions of a point a year is never minuted as a decision — the absence of a decision is the point, and a process with no event has no date. **The released share disperses rather than transferring**: it goes into a basket of secondary units and into metal rather than to a single rival, so the reserve system’s own effective count *rises* and no recipient approaches the position being vacated. **And part of the drift is not diversification at all**: a move from one unit into another whose issuer sits inside the same coordinating perimeter has changed the currency and not the perimeter, which is the same error as rearranging exposure inside a bloc — and rearranging inside a perimeter does not change N. Curzi [2026] measures exactly that collapse on a historical network, where five houses read 3.819710 effective jurisdictions naively and 1.000000 effective blocs once one coordination perimeter is recognised, giving $h = -0.500000$.

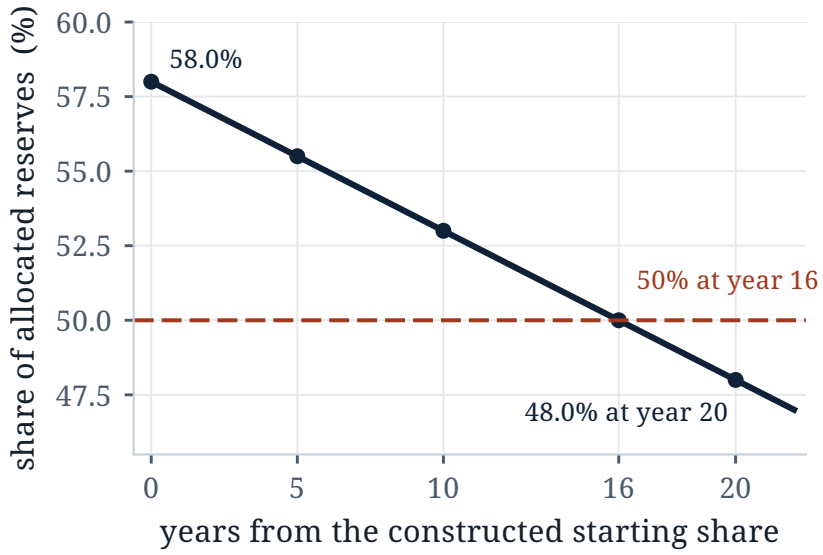


Figure 68.4: **A slope that crosses no threshold for sixteen years.** A constructed reserve share of 58.0 per cent declining at 0.5 percentage points a year reaches 50 per cent at year 16 and 48 per cent at year 20. Presented as a timescale illustration only: the figures are invented, the slope is stipulated, and the paper forecasts nothing. Script `figs/src/CCP-068_fig04.py`.

6.2 H5 — the representation split: supported

CONSTRUCTED national capital vector, reconciled to reproduce Curzi [2026]'s result at national scale:

form	constructed share	on the monetary ledger	represented
financial	42.0%	yes, fully	42.00%
material	18.0%	yes, fully	18.00%
intellectual	11.0%	partly, ~35% recognised	3.85%
experiential	9.0%	no	0.00%
social	7.0%	no	0.00%
natural	6.0%	no	0.00%
cultural	3.0%	no	0.00%
living	2.5%	no	0.00%
spiritual / ethical	1.5%	no	0.00%
total	**100.0%**		**63.85%**

```
fully on the monetary ledger    financial + material    = 60.0000%
partially (intellectual @ 35%)                          =  3.8500%
TOTAL REPRESENTED                                       = 63.8500%
STRUCTURALLY OFF-LEDGER                                 = 36.1500%

of the off-ledger portion, marked UNMEASURED            = 11.5000%
  (social 7.00 + cultural 3.00 + spiritual/ethical 1.50)
of the off-ledger portion, partially measurable         = 24.6500%
```

H5 is supported on the constructed vector: the monetary ledger balances exactly over 63.85% **and is silent about** 36.15%**, of which** 11.50% **admits no accepted measurement basis at**

all. The silence is not a policy the standard-setter chose and could reverse. It is the price of the closure property: a ledger recognising unmeasurable items would stop balancing in any meaningful sense, and closure is the one thing double entry actually guarantees. **The gap is structural, and the correct response is not to demand the ledger carry more but to know exactly what it does not carry and hold that knowledge somewhere else, deliberately.**

No single total is struck, and the refusal is the specification. Striking one would assign the 11.50% of unmeasured forms an implicit value the standard declines to invent, which is the scale-type violation of Section 3.4. **Silence treats an unmeasured form as zero; fabrication invents a proxy; naming it and marking it UNMEASURED does neither, and puts it on the account as a known gap visible to any reader.**

That the accounting method of the period is often credited with more causal force than it can carry is worth naming while the completeness gap is open: Yamey [1964] and Littleton and Yamey [1956] are the deflation of the claim that double entry produced rational capitalism, and no such claim is made here.

6.3 H6 — the bounded ceiling: supported, and bounded in both directions

```
off-ledger / represented  =  36.1500% / 63.8500%  =  0.566170  ->  56.62%
credible base multiple    =  100%     / 63.8500%  =  1.566170  ->  1.5662x
```

Making the off-ledger capital visible widens the priceable base by at most 56.62% — it is 1.5662 times as large when the whole is visible as when only the monetised subset is. That is a real number computed correctly from the constructed vector, **and it is a ceiling and not a promise.**

The lower bound is strictly above zero and is not zero. Visibility moves the newly reported capital from "unknown, therefore priced at zero" to "known, of contested value", and a quantity of contested value is priced above zero by any rational lender even when it is priced far below its owner's estimate. **The gain is the difference between zero and whatever a creditor credits a visible-but-contested base at: positive, bounded above by 56.62%, and impossible to state precisely without the valuation the taxonomy does not supply.**

H6 is supported and the paper stops exactly there. The ceiling is computable and is stated; the point estimate would require solving the valuation problem, which is not solved here and is named in Section 7.2 as the hardest problem in the framework. **A proposal quoting the 56.62% as a promised increase in borrowing capacity has quoted a ceiling as a floor.**

6.4 The coupling, badged as an inference

The two exposures measured here are separated for measurement and **are coupled in fact**: a refinancing schedule concentrated at the short end must be rolled into the hands of whoever will buy, and if the historical buyers are themselves drifting then the drift thins the marginal bid exactly at the buckets where the concentration needs it most.

That claim is an analytical inference and not a measured result, because it depends on the holder composition of the near-term maturities, which is a separate axis and is not measured anywhere in this paper. It is badged accordingly in Section 10.6 as a POSTULATE, and Section 3.3 records that it is the one place in this paper where the identification problem bites.

7. Discussion

7.1 What the composition is, and what it is not

Read as a stack, each layer is the precondition of the next: a canonical form makes two records of the same thing compare as the same thing; a hash over the raw octets makes each record

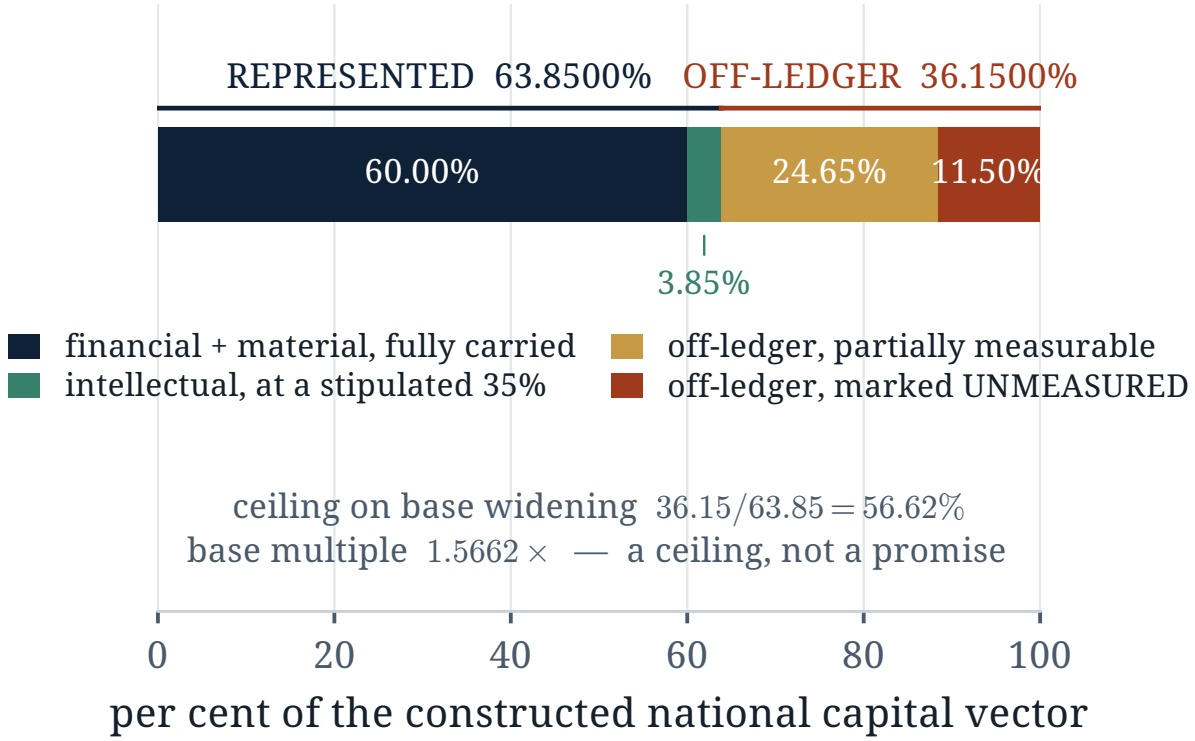


Figure 68.5: **A ceiling, not a promise.** On a constructed nine-form capital vector, 63.85 per cent of national capital is represented on the monetary ledger and 36.15 per cent is structurally off-ledger, of which 11.50 per cent admits no accepted measurement basis. Making the remainder visible widens the priceable base by at most 56.62 per cent — a multiple of 1.5662 — and the honest gain lies strictly between zero and that ceiling. CONSTRUCTED vector; script `figs/src/CCP-068_fig05.py`.

tamper-evident at the octet level and preserves the evidence of a substitution that canonicalisation would otherwise erase; content addressing makes reference and verification one object; per-object hash chains make each object's history tamper-evident; a structured message layer carries all of it; and settlement spreads across perimeters that do not share a coordination mechanism.

The composition is real as an architecture and unbuilt as a system, and the distinction is not pedantry. The individual pieces have very different maturities and the integration inherits the weakest of them. **There is no independent implementation, no interoperability profile, and no conformance suite maintained by anyone outside its author — the ordinary prerequisites for anything to become infrastructure, none of which is met.** The honest status is a coherent architecture and an undeployed one, **and coherence is not readiness.**

Curzi [2026] is the corpus's own template for exactly this disclosure and it bounds what this paper may claim: there, an entire pipeline is built, verified and deterministic, eleven upstream components are established, and the paper states that nobody knows how the one missing link works — the carriage of sequence over force. **This paper is in the same position with two links missing rather than one**, and it says so in the same register.

7.2 The valuation problem, named as unsolved

Naming nine forms of capital does not tell anyone what any of them is worth. The valuation problem is the entire difficulty, and the measurement basis for the non-financial forms is not specified in the source manuscript and is not specified here. A coverage ratio computed across a portfolio in which one component has no accepted measurement basis is a ratio computed over an unknown.

Content addressing can give such a portfolio a stable, tamper-evident name; it cannot give an unvalued line item a value, and a content address over an unvalued quantity is a precise name for an imprecise thing. This is the deepest unsolved problem underneath the architecture, and a synthesis that buried it under a tidy diagram would have failed its own test.

The measurement-theoretic form of the objection is the sharpest one and is conceded rather than answered. Michell [1997] argues that a discipline systematically fails to test whether its attributes are quantitative at all before measuring them; **that is the strongest single objection to this framework**, and the formal test — conjoint measurement [Luce and Tukey, 1964], whose conditions establish additive quantitative structure — **was not run for any of the nine forms.** Cronbach and Meehl [1955] and Campbell and Fiske [1959] set out what construct validation would require of any proposed measure for the softer forms, and **no such validation exists for any of them.**

7.3 The bounds this paper works inside

Curzi [2026] bounds what borrowed formalism may claim. Its counter-result declines the inflated reading of its own borrowed physics: the postulate is stated as a postulate, the relation is given only in its correct form, and a named failure mode requires that a user be able to distinguish the theorem, the ordinary observer effect, and the framework's own conjecture in three separate sentences on demand. **Applied here, that is why Section 3.1's correspondence between fusion and separation is badged as this paper's own analytical framing rather than as a standard**, and why Section 10.5 sorts every borrowing by category.

Curzi [2026] bounds what a discovered gap licenses. Its counter-result refuses category innocence as a defence of acquisition, forbids moving any harm through a gap discovered in a scheme, and concedes that the case where the feedback loop runs through another person's will is an unsettled seam. **Applied here, that is the constraint on the standard-setting argument.** A standard that makes national capital visible is legibility exercised on one's own account [Scott, 1998], and legibility exercised on someone else's is a different act with a different licence. **The reform proposed here changes what a nation reports about itself; nothing in**

this paper licenses using the same instrument to make a counterparty legible against its will, and that boundary is stated rather than assumed.

Curzi [2026] is the bound stated as an apparatus. Its counter-result assigns the corpus's central identification, the application of the surplus functional to a life, and the reading of a day's surface as a leading indicator to the postulate drawer, and says of the third that it is a wager whose misfire column exists because wagers can lose. **Applied here, it is why Section 6.2's account marks three forms UNMEASURED rather than estimating them, and why Section 6.3 states a ceiling and refuses a point estimate: an honest register has empty cells in it.**

Curzi [2026] bounds inference from an invariant statistic, and its bound applies directly to the effective spread. That paper's counter-result is that a width statistic invariant under relabelling does not even determine the object it is computed from, which formally blocks inference from a high or low value to a substantive condition. **The effective spread is invariant under any permutation of the buckets**, so it does not determine the schedule it was computed from: two stacks with the same spread can have their concentration at opposite ends of the maturity range, and the near-window share of Section 5.2 exists precisely because the index alone cannot say *where* the concentration sits. **The index and the near-window share must be reported together or each hides what the other shows.**

Curzi [2026] bounds the remedy. Its counter-result is that perfect opposition and perfect agreement read identically, because the framework's non-parallelism depends only on the magnitude of a correlation — so the standard remedy of adding a contrarian contributes zero surplus while adding cost, and at zero non-parallelism two of four levers are exactly inert. **Applied here, that is the warning against assuming that adding buckets, adding forms of capital, or adding rails is automatically an improvement**: an addition that is perfectly correlated with what is already held adds cost and no independence, and the index will report the addition while the exposure does not move.

7.4 Why the standard would be performative, and why that is an argument for stating its limits

Curzi [2026ab]'s belief dynamic is presupposed here and it is the right formal template. $B_{t+1} = (1-\lambda)B_t + \lambda \min(1, LB_t)$ with loop gain $L = \alpha G$: the zero fixed point is stable if and only if $L < 1$, the transition at $L = 1$ is sharp, **and** $G = 0$ **forces** $L = 0$ **at every** α — a structural zero cannot be compensated by any amount of the other factor. **Applied to a standard: no amount of adoption enthusiasm substitutes for a measurement basis that does not exist.** A form of capital with $G = 0$ — no accepted method — contributes nothing to the priceable base however loudly it is reported, which is exactly why Section 6.2 marks three forms UNMEASURED rather than reporting them at a number.

The reflexive side is real and cuts both ways. Hacking [1995] states the general mechanism by which a classification changes what it classifies; MacKenzie and Millo [2003] documents the case in the field with dates and prices; Callon [1998] and Ferraro et al. [2005] generalise it to economic description; and Lucas [1976] states the consequence for any parameter fixed before publication. **A standard that made off-ledger capital countable would change how it is managed, which is the whole of the power on offer and also the reason its unmeasured cells must stay visibly empty**: a gap that is reported as a gap cannot be optimised into a target, and a gap filled with a proxy immediately can be [Goodhart, 1975, Campbell, 1979].

And the same reflexivity is not applied to this paper's own parameters, which is an inconsistency and is not resolved here. The framework asserts that systems change when described and then exempts its own constructed vectors, buckets and recognition fraction from that response. What is offered instead is the reporting standard of Appendix A.

7.5 The commitment that cuts both ways

A content-addressed, hash-chained record is permanent, byte-addressable, discoverable in its entirety, and cannot be quietly amended. For a structure built so that nothing is hidden because nothing needs to be, that is a feature — and it is the disclosure constraint of Curzi [2026] stated as storage: **effectiveness independent of whether anyone knows the design is in place, so every element is publishable and non-disclosure is refused as a design premise.** Kerckhoffs [1883] is the principle's canonical statement, Saltzer and Schroeder [1975] its restatement as open design, and Brandeis [1914] the older policy instinct.

It is nonetheless a commitment made once and honoured forever, and an entity adopting it surrenders the option of a record that ceases to exist. The integration should be entered deliberately, as a decision about permanence, and not inherited by default from a diagram because the diagram was elegant. **Elegance is not a reason to become permanent.**

The other side of that permanence is a data-protection exposure that belongs in print. **A permanent record of everything a standard reports is also a permanent record that cannot be corrected by deletion**, and any account carrying entity-level identifiers inherits that. Curzi [2026x] reaches the structurally parallel result in another domain — acoustic addressing fails on geometry rather than on bandwidth, since wavelength matches a ten-micrometre cell only at 154 MHz and the product of wavelength and 20 dB depth collapses from 6160 to 0.0026 mm-cm between 100 kHz and 154 MHz — and the shared lesson is that **a system can have ample capacity on the axis everyone measures and be defeated by a constraint on an axis nobody measured.** Permanence is that axis here.

7.6 What would falsify the synthesis

Three things, none tested here. If canonicalisation underneath the addressing produced a materially elevated false-collision rate against a real name estate — two genuinely distinct entities folded to one address — the laundering remedy would become a cost–benefit question rather than a free improvement. If an ordering layer capable of resolving a per-object ledger's conflicts at the required scale cost more in latency or complexity than the write-parallelism the factoring buys, the object-as-ledger choice would be an optimisation that does not pay for itself. **And if no accepted measurement basis for the non-financial forms materialises under audit scrutiny, the taxonomy remains a coherent way to name what the ledger cannot see and no way at all to value it — which is its status today.**

8. Failure modes

8.1 Presenting the base widening as a promise. 56.62% is a ceiling on a constructed vector. **Quoting a ceiling as a floor is the fastest available way to lose a technical reader.**

8.2 Confusing visibility with pledgeability. Making a form of capital visible does not make it collateral. The reform improves how a base is priced; it does not create a security interest in the unsecurable.

8.3 Fabricating a proxy for an unmeasured form. An index entered as a number where no accepted method exists corrupts the whole account, because a reader cannot tell the fabricated entries from the measured ones [Goodhart, 1975].

8.4 Treating the integrity layer as a valuation. Canonical form, byte-level hash and a cryptographic witness make the *record* honest. They do not make the *numbers* true. **A tamper-evident record of a wrong valuation is a wrong valuation that cannot be altered.**

8.5 Averaging the measurable and the unmeasurable. A single total across measured and unmeasured forms assigns the unmeasured an implicit value by including them. **The account reports the fractions separately and strikes no headline.**

8.6 Replacing rather than supplementing. The standard sits beside the monetary ledger and changes none of its figures. A version proposing to replace double entry has inverted the completeness-gap result, which defends the method as required and not obsolete.

8.7 Quoting the total debt as though it were the exposure. The total is a sum. **Two stacks of identical total can have entirely different time-axis concentrations and therefore entirely different repricing profiles.**

8.8 Confusing rollover risk with repricing risk. For an issuer borrowing in its own unit the debt will be rolled; the question is at what rate. **The wall is a price, not a gate**, and treating it as a solvency cliff imports a private-borrower failure mode into a setting where it does not apply.

8.9 Watching the visible slope and missing the measurable one. A commentary tracking a reserve share to the decimal and never computing the maturity concentration has watched the slope that disperses and missed the slope that amplifies. **The total is easy and the distribution is work.**

8.10 Treating a constructed distribution as a real one. The 0.222050 figure is computed on an invented stack. **Anyone quoting it as a real issuer's time-axis index has moved a number across a boundary without carrying its label.**

8.11 Assuming a rearrangement inside a perimeter is diversification. A move between two units whose issuers sit inside one coordinating perimeter changes the currency and not the perimeter.

8.12 Believing a diagnosed fragility must resolve. A compromised position is not a countdown. It can persist for decades, actively defended where it was once assumed, without healing or breaking.

8.13 Reading the addressing layer as sufficient. Section 4.3. The composition claims the representation, reference and settlement layers and nothing above them, and it must be paired with substantive audit and controls review.

9. Discrepancies found in the source, printed

One discrepancy assigned to this paper belongs to an adjacent paper's model and is reported here in full rather than passed over. The migration model of Curzi [2026] tabulates a remainder at 2046 of $1{,}550{,}381$ lines from a $40{,}000{,}000$-line core migrated at 15% of the remainder per year from 2026. That figure is $40{,}000{,}000 \times 0.85^{20} = 3.876\%$ of the core, i.e. **twenty** annual migrations. The same model separately states a residual "after twenty years" of $1{,}317{,}824$ lines, or 3.29%, which is $40{,}000{,}000 \times 0.85^{21}$ and therefore **twenty-one** annual migrations.

$$40{,}000{,}000 \times 0.85^{20} = 1{,}550{,}381 = 3.876\%,$$
$$40{,}000{,}000 \times 0.85^{21} = 1{,}317{,}824 = 3.295\%.$$

The two statements are off by exactly one migration year. The arithmetic was recomputed here and both figures are internally correct for their own exponent, so the defect is in the year label and not in either computation. Curzi [2026] prints the discrepancy in its own verification section and adopts the 3.29% figure as stated while recording that its tabulated series gives 3.88% at the same date.

No table and no figure in this paper uses either number. This paper contains no migration model, no geometric-decay remainder and no line counts; its own decay illustration is the *linear* reserve drift of Section 6.1, whose figures — 58.0% falling 0.5 points a year, reaching 50.0% at year 16 and 48.0% at year 20 — were recomputed and are internally consistent, since $58.0 - 0.5 \times 16 = 50.0$ and $58.0 - 0.5 \times 20 = 48.0$ exactly. **The discrepancy is recorded here because a reader who meets both papers is entitled to know that the two figures are one year apart and which paper owns the disclosure; it changes nothing in either, since the result in question is that a geometric remainder is strictly positive for all finite time.**

A second, smaller inconsistency was found in this paper's own source and is likewise printed rather than harmonised. The time-bucket index is computed and printed as 0.222050 in the derivation and then quoted as 0.2220 in every comparison table and summary. 0.222050 rounds to 0.2221 at four decimal places under round-half-up and to 0.2220 under round-half-to-even, so the quoted figure is defensible only under the second convention and the source states no convention. **This paper prints 0.222050 throughout and quotes no four-decimal form**, which removes the ambiguity rather than resolving it by fiat. The effective spread is unaffected: $1/0.222050 = 4.5035$, quoted as 4.50.

10. Verification and rigour audit

10.1 Dimensional audit

symbol	quantity	dimension	unit	status
s_i	bucket or form share	$[1]$	fraction	defined; sum to 1
HHI	sum of squared shares	$[1]$	fraction	derived
effective spread	HHI^{-1} on maturity buckets	$[\#]$	buckets	derived
maturity bucket	calendar interval	$[T]$	years	defined (chosen boundaries)
coupon, repricing delta	rate on principal	$[T^{-1}]$	per year (percentage points)	assumed (constructed)
added interest	flow per unit of stock	$[T^{-1}]$	per cent of stack per year	derived
reserve share	fraction of allocated reserves	$[1]$	per cent	assumed (constructed)
drift rate	change in share per unit time	$[T^{-1}]$	percentage points per year	assumed (constructed)
capital share	fraction of a national capital vector	$[1]$	per cent	assumed (constructed)
recognition fraction	share of a form the ledger carries	$[1]$	per cent	**stipulated at 35%**
base multiple	ratio of whole base to visible base	$[1]$	dimensionless ratio	derived
digest, identifier	fixed-length function of octets	$[1]$	256 bits; base32 text	derived, exact

Does the repricing calculation balance? A share (dimensionless) times a rate ($[T^{-1}]$) gives a flow per unit of stock ($[T^{-1}]$): $0.55 \times 0.0250\ \mathrm{yr}^{-1} = 0.01375\ \mathrm{yr}^{-1}$, read as 1.375% of the stack per year. **It balances, and the units are per-year fractions of an unspecified stock — there is no currency anywhere in this paper.**

Does the capital vector balance? The nine shares sum to 100.0% exactly, and represented plus off-ledger is $63.85 + 36.15 = 100.00$. **The base multiple is a ratio of two fractions and is**

therefore dimensionless, which is why it survives the total absence of a unit of account.

Where a quantity has no operational unit, the audit says so. The nine capital forms have **no common unit**, and this is not a presentational gap: it is the substance of the valuation problem. The shares of Section 6.2 are shares *of a stipulated total* and not measurements in any unit, and **the fact that three of the forms admit no unit at all is exactly why no total is struck.**

10.2 Significant figures and epistemic status

- **The digests and identifiers of Section 4 are exact and were independently reproduced.** They are deterministic functions of stated byte strings; **no uncertainty attaches and none is quoted.**
- **The time-bucket index is exact for its share vector**: 0.222050 is the exact sum of squares of the six stated shares, and $1/0.222050 = 4.5035\ldots$ is quoted as 4.50 because the third figure is not informative about a constructed distribution.
- **The flat-ladder index is exact and rational**: $6 \times (1/6)^2 = 1/6 = 0.166667$, spread exactly 6.00.
- **Every input is stipulated, not estimated.** The bucket shares, the repricing delta, the drift rate, the recognition fraction and the capital shares were all chosen. **Exact arithmetic on invented inputs is not a measurement, and no uncertainty is quoted because nothing is estimated.**
- **The recognition fraction of 35% is the single most load-bearing stipulation in Section 6.** Moving it moves the represented share, the off-ledger share and the ceiling together, and Section 10.3 gives the sensitivity.
- 56.62% **and** 1.5662 **are quoted to four figures because they are exact functions of exact stipulations**, and the fourth figure is meaningful about the construction and meaningless about any real position.
- **The reserve figures are round by construction** and are quoted to one decimal place only.

10.3 Limits and boundary behaviour

quantity	limit	behaviour	reading
effective spread	all mass in one bucket	$\to 1$	maximal time concentration
effective spread	n equal buckets	$\to n$ exactly	the flat ladder; the ceiling for n buckets
effective spread	bucket count $\to \infty$ at fixed shape	$\to \infty$	**an artefact of bucketing, not a property of the schedule**
near-window share	window $\to 0$	$\to 0$	no exposure without a window
added interest	delta $\to 0$	$\to 0$	no repricing without a rate change
added interest	near share $\to 1$	$\to$ delta	the whole stack reprices at once
reserve drift	linear at fixed rate	crosses any level in finite time	**the linearity is the stipulation, not a finding**
represented share	recognition $\to 1$	$\to 71.00\%$	full intellectual recognition still leaves 29% off-ledger

quantity	limit	behaviour	reading
represented share	recognition $\to 0$	$\to 60.00\%$	the two fully carried forms alone
ceiling	represented $\to 1$	$\to 0$	nothing off-ledger, nothing to reveal
ceiling	represented $\to 0$	$\to \infty$	**unbounded, and the model's domain excludes it**

Two rows are defects rather than results and are stated as such. The effective spread grows without bound as buckets are subdivided, so **the index is meaningful only against a fixed bucketing and only as a difference against a prior measurement on the same boundaries**; that is a real constraint on the instrument's domain and Appendix A makes the boundaries a mandatory reported item. And the ceiling diverges as the represented share goes to zero, so **the ratio form is informative only where a substantial fraction is already represented**, which the constructed vector satisfies at 63.85%.

The recognition sensitivity, since it is the load-bearing stipulation:

```
intellectual recognition   represented    off-ledger    ceiling    multiple
------------------------------------------------------------------------
        0%                   60.0000%       40.0000%     66.67%     1.6667x
       35% (used)            63.8500%       36.1500%     56.62%     1.5662x
       50%                   65.5000%       34.5000%     52.67%     1.5267x
      100%                   71.0000%       29.0000%     40.85%     1.4085x
```

The ceiling falls from 66.67% to 40.85% across the full range of a single stipulated parameter. That is the sensitivity a reviewer should press on, and it is printed rather than left to be found.

10.4 Invariance

The index is invariant under any permutation of the buckets, which is exactly why it does not determine the schedule it was computed from — the bound carried from Curzi [2026]. Two stacks with identical spread can concentrate at opposite ends of the maturity range. **The near-window share is the statistic that is *not* permutation-invariant, and reporting the pair is what recovers the location the index discards.**

The index is not invariant under rebucketing, and this is the instrument's sharpest limitation. Finer buckets at the short end raise resolution where exposure concentrates and mechanically raise the effective spread; **a measurement is meaningful only as a difference against a prior measurement on the same boundaries.**

The base multiple is invariant under any rescaling of the capital vector, since it is a ratio of two fractions of the same total — which is why it survives the absence of any unit of account, and why it is the only quantity in Section 6 that would still mean something if the shares were expressed in some future unit.

The content identifier is invariant under nothing except identity of the octets, which is the entire point: it is not invariant under a homoglyph substitution, and Section 4.2 is the demonstration that this is both its strength and, without canonicalisation, its failure mode.

Tested at how many scales? **The invariances are algebraic and were checked at the one parameter set reported and the one recognition sweep of Section 10.3.**

10.5 Analogy discipline

borrowing	source	category	justification
HHI^{-1} as a numbers-equivalent	Adelman [1969], Hill [1973], Jost [2006]	**1 — derivation**	shares are non-negative and sum to one by construction; the identity is exact
independent reinvention of the same index	Simpson [1949], Laakso and Taagepera [1979]	**1 — derivation**	the same arithmetic on the same object in three disciplines
exposure as a portfolio	Markowitz [1952]	**1 — derivation**	the variance identity's premises hold by construction
similarity-sensitive diversity of order 2	Leinster and Cobbold [2012], Rao [1982]	**1 — derivation**	exact for a unit-diagonal non-negative correlation matrix
correlation rising under stress	Longin and Solnik [2001], Ang and Chen [2002]	**2 — shared form**	measured in equity markets; a maturity schedule is not shown to satisfy the same conditions
independent-party bound; identity multiplication	Lamport et al. [1982], Douceur [2002]	**2 — shared form**	theorems about agreement and identity systems, cited as prior formalisations of the missing layers and not applied to anything here
impossibility under asynchrony; the availability trade	Fischer et al. [1985], Gilbert and Lynch [2002]	**3 — framing**	cited to establish that structural impossibility results bound what an ordering layer can promise
scale type and permissible statistics	Stevens [1946], Krantz et al. [1971]	**1 — derivation**	an unmeasured form genuinely admits no scale, and the restriction follows from the definition
construct validation	Cronbach and Meehl [1955], Campbell and Fiske [1959]	**1 — derivation**	the requirement is stated for exactly the case at hand and is **not met**

borrowing	source	category	justification
open design; disclosure strengthens a structural control	Kerckhoffs [1883], Saltzer and Schroeder [1975]	**1 — derivation**	effectiveness depends on the arrangement rather than on anyone's knowledge of it; the premise is satisfied
legibility as state capacity	Scott [1998]	**2 — shared form**	a historical and ethnographic account; used for the mechanism and the warning, not as a formal result
classification interacts with what it classifies	Hacking [1995], MacKenzie and Millo [2003], Callon [1998], Ferraro et al. [2005]	**1 — derivation**	the mechanism is the same: a description relied upon changes the behaviour of those relying on it
a measure adopted as a target	Goodhart [1975], Campbell [1979]	**2 — shared form**	stated for monetary and social indicators; the exposure is the same in kind
commons governance and its audit	Ostrom [1990], Ostrom [2010], Cox et al. [2010], Hardin [1968]	**2 — shared form**	field-derived design principles; cited as the literature any natural- or social-capital figure must come from, not as supplying one
social capital, contested	Putnam [2007], Abascal and Baldassarri [2015], Fukuyama [2001], Gambetta [1988]	—	cited to establish that the construct is **contested**, which is the reason the form is marked UNMEASURED
asset partitioning; legal coding	Hansmann and Kraakman [2000], Pistor [2019]	**1 — derivation**	the pledgeability point is exactly what these establish

borrowing	source	category	justification
the modelling programme	Epstein [2008], Bruch and Atwell [2015], Hedström and Ylikoski [2010]	—	the first enumerates the legitimate non-predictive uses of a model; the second is the checklist **none of this paper's models meets**; the third states the criterion this paper accepts, that an explanation must name a mechanism rather than exhibit a fit
fusion and separation as complements	this paper	**3 — framing**	**this paper's own analytical framing**, badged as an argument and not as a standard

10.6 Postulates

#	statement	status
P1	shares are non-negative and sum to one	DEFINITION
P2	HHI $= \sum_i s_i^2$; effective spread $=$ HHI^{-1}	DEFINITION
P3	for n equal shares of $1/n$, HHI $= 1/n$	ARITHMETIC
P4	an identifier is a self-describing wrapper around a digest of the content	DEFINITION (described from the source; uncited)
P5	canon$(\cdot)$ folds visually confusable forms onto one skeleton	DEFINITION, per Curzi [2026]
P6	a fixed coupon does not change while the debt is outstanding, so a rate change acts on refinancing	DEFINITION (standard property of fixed-coupon debt)
P7	the six maturity buckets and their boundaries	POSTULATED (modelling choice)
P8	the constructed maturity shares, repricing delta, drift rate and capital vector	POSTULATED (constructed)
P9	intellectual capital is recognised at 35%	POSTULATED (stipulated; the load-bearing parameter)

#	statement	status
P10	a creditor prices visible-but-contested capital strictly above zero	POSTULATED (economic model; the lower bound of H6)
P11	a drifting reserve share thins the marginal bid at the buckets where refinancing concentrates	POSTULATED (analytical inference; not identified)

10.7 The chains

Chain A — content addressing (R5).

step	consumes	statement	status
A1	P4	the identifier is a function of the octets alone	DEFINITION
A2	A1	sha256$(A) \neq$ sha256(B) for the two stated strings	ARITHMETIC — **independently recomputed**
A3	A1, A2	the two identifiers share no structure	ARITHMETIC — **independently recomputed**
A4	A1	re-hashing received octets and comparing to the name is a complete integrity check	DEFINITION
A5	P5	canon$(\text{clean}) =$ canon(twin) for the stated pair	DEFINITION
A6	A1, A5	the identifier over the canonical form is identical for both	ARITHMETIC — **independently recomputed**
R5	A1–A6	one byte's difference yields a different identifier; canonicalisation folds visual twins onto one	DERIVED

Chain B — the time axis (R1, R4).

step	statement	status
B1	$\sum s_i^2$ on 37.0/11.0/7.0/16.5/15.5/13.0 is 0.222050	ARITHMETIC
B2	effective spread $= 1/0.222050 = 4.5035$	ARITHMETIC
B3	$37.0 + 11.0 + 7.0 = 55.0\%$ due within three years	ARITHMETIC

step	statement	status
B4	flat ladder: $6 \times (1/6)^2 = 0.166667$, spread 6.00, 50.0% within three years	ARITHMETIC
B5	$0.55 \times 0.0250 = 0.013750$; $0.50 \times 0.0250 = 0.012500$	ARITHMETIC, consumes P8
R1	B1–B4	the time-axis index, spread and near-window share, against the ladder
R4	B5	1.375% against 1.250% of the stack per year

Chain C — representation and the ceiling (R2, R3).

step	statement	status
C1	financial 42% + material 18% fully carried $= 60.0000\%$	ARITHMETIC
C2	intellectual $11\% \times 0.35 = 3.8500\%$	ARITHMETIC, consumes P9
C3	represented $= 60.0000 + 3.8500 = 63.8500\%$	ARITHMETIC
C4	off-ledger $= 100 - 63.8500 = 36.1500\%$	ARITHMETIC
C5	unmeasured $= 7.00 + 3.00 + 1.50 = 11.5000\%$	ARITHMETIC
C6	$36.1500/63.8500 = 0.566170$, i.e. 56.62%	ARITHMETIC
C7	$100/63.8500 = 1.566170$, i.e. $1.5662\times$	ARITHMETIC
C8	the realised gain exceeds zero	POSTULATED (P10)
R2	C3–C5	the representation split with its unmeasured fraction
R3	C6–C8	the ceiling, with the gain bounded strictly between zero and it

Only C8 is not arithmetic, and it is the step a reader should attack first. It is an assumption about creditor behaviour rather than about mechanism, and if it fails the reform is inert. **Every other step in Chain C is one line of arithmetic on stipulated shares.**

Chain D — reserve drift (R6).

step	statement	status
D1	$58.0 - 0.5t$ for t years	DEFINITION (constructed linear model)
D2	$58.0 - 0.5 \times 16 = 50.0$	ARITHMETIC
D3	$58.0 - 0.5 \times 20 = 48.0$	ARITHMETIC
R6	D1–D3	a timescale illustration crossing no threshold for sixteen years

10.8 Cross-checks performed

- **The four content identifiers and both digests of Section 4 were recomputed from first principles** — SHA-256 over the stated octets, wrapped as a version-1 raw-codec multihash, rendered in lower-case base32 — and reproduce the source manuscript exactly. **This is the only independent reproduction in the paper and it is the strongest verification status any figure here carries.**
- The maturity index was recomputed by explicit sum of squares and reproduces 0.222050; the ladder reproduces 0.166667 exactly as $1/6$.
- The bound $s_{\max}^2 \le \text{HHI} \le s_{\max}$ holds: $0.136900 \le 0.222050 \le 0.370000$, and the effective-count ceiling $s_{\max}^{-2} = 7.304$ exceeds the bucket count of six, so **the bucket count is the binding bound and the leading-share bound does not bite here.**
- The private schedule was re-summed: $0.36 + 0.0625 + 0.0225 = 0.445000$, $1/0.445000 = 2.2472$.
- The capital vector sums to 100.0% exactly, and represented plus off-ledger reproduces 100.00%.
- The recognition sweep of Section 10.3 was computed at four values and the ceiling is monotone decreasing in the recognition fraction, as it must be.
- The reserve line was checked at all five tabulated years and is exactly linear.
- **The migration figures of Section 9 were recomputed**: $0.85^{20} = 0.0387595$ and $0.85^{21} = 0.0329456$, giving $1{,}550{,}381$ and $1{,}317{,}824$ lines respectively on a 40-million-line core.

11. Limitations

Every worked result is a constructed model and none is a measurement. The settlement objects, the account identifiers, the maturity distribution, the repricing delta, the private schedule, the reserve share and its drift rate, and the nine-form capital vector are all invented. **No dataset was collected, and no figure here is a measurement of any real issuer, nation, entity or holding.**

No as-of date and no source, because nothing real was measured. A concentration figure computed from a real schedule requires its source and its date; **a concentration measure without an as-of date is not a measurement.** Every distribution here is constructed, so there is no source and no date to give, and Appendix A makes both mandatory items of the reporting standard.

The bucket boundaries are a modelling choice and the index is not invariant to them. Section 10.4. A measurement is meaningful only as a difference against a prior measurement on the same boundaries, and no boundary set is defended here as correct.

The 35% recognition fraction is stipulated and it is the load-bearing parameter of Section 6. Section 10.3 shows the ceiling moving from 66.67% to 40.85% across its range. **No source supplies the fraction and none is claimed.**

The valuation of the non-financial forms is unsolved and is the hardest problem underneath the whole architecture. It was unsolved in the source manuscript and it is unsolved here. **The conjoint conditions were not tested for any form** [Michell, 1997, Luce and Tukey, 1964], and no construct validation exists [Cronbach and Meehl, 1955, Campbell and Fiske, 1959]. The commons and social-capital literatures cited in Section 3.4 are cited to establish that the constructs are **contested**, never as supplying a figure.

Visibility is not pledgeability and the paper's positive claim is bounded by that. Most of the newly visible capital is unpledgeable and unrealisable as security by its nature.

The coupling claim of Section 6.4 is an analytical inference and not a measured result. It depends on the holder composition of near-term maturities, which is a separate axis this paper does not measure, and it is the one place the identification problem bites [Manski, 1993, Brock and Durlauf, 2001].

Statistical series, statutes, accounting standards, storage specifications and legislative material are not cited. They are described from the source manuscript, and the bibliography

contains none of them. Appendix B lists every affected claim. **This is a real gap and it is not closed here.**

Two structural layers are missing from the composition and it is not a distributed ledger without them. An ordering layer and a resistance-to-identity-multiplication layer are required [Lamport et al., 1982, Douceur, 2002] and neither is supplied.

The permanence commitment is one-way. A content-addressed record cannot be corrected by deletion, and any account carrying entity-level identifiers inherits that exposure. **The paper states it as a cost and does not quantify it.**

11.1 Rigour checks not met, listed by name

1. **No uncertainty is reported anywhere, because nothing is estimated.** Exact arithmetic on stipulated inputs is the absence of empirical content, not precision.
2. **Model comparison was not performed.** No alternative index was fitted, no out-of-sample test was run, no information criterion was computed and no matched null was constructed, because nothing here is fitted.
3. **The conjoint conditions were not tested and no construct validation exists** for any non-financial form [Michell, 1997, Luce and Tukey, 1964, Krantz et al., 1971].
4. **No distributional or scaling claim is made, so the Clauset–Shalizi–Newman procedure was not run** [Clauset et al., 2009]. The vectors are finite sets of three to nine stipulated shares, not samples; no tail is characterised and no power law is asserted. **Any restatement reaching for a heavy-tail claim about holdings, maturities or capital shares must run it first**, and must read Stumpf and Porter [2012] on why passing it would still not identify a mechanism.
5. **The index is not invariant under rebucketing**, which is a real defect of the instrument and is stated in Section 10.4 rather than discovered by a reader.
6. **The ceiling diverges as the represented share approaches zero**, so the ratio form has a restricted domain.
7. **P10 is an assumption about creditor behaviour, not about mechanism**, and it carries the entire lower bound of the paper's positive claim.
8. **P11 is an analytical inference presented as such**, and the identification problem is not answered for it.
9. **The models fail the calibration-and-validation checklist** of Bruch and Atwell [2015]: none is calibrated, none is validated out-of-sample, and none has a sensitivity analysis beyond Section 10.3. They are demonstrations that a stated mechanism suffices to produce a stated shape, one of the uses Epstein [2008] enumerates, and Hedström and Ylikoski [2010] is the criterion accepted: **nothing here is fitted, and each result follows from an inequality stated in advance.**
10. **The framework is performative and its own parameters are exempted from that** [MacKenzie and Millo, 2003, Lucas, 1976, Hacking, 1995]. **That is an inconsistency and it is not resolved here.**
11. **Expected value and time average are not the same thing** [Peters, 2019], and the repricing figures are per-period averages that no single issuer experiences as such.
12. **There is no conservation law anywhere in this paper.** The only genuine conservation is that shares sum to one [Mirowski, 1989]. **This is the objection the genre most often fails and this paper meets half of it.**
13. **Aggregation is unaddressed** [Kirman, 1992]: a single recognition fraction is written for a heterogeneous form of capital.
14. **No differential prediction is offered** [Durlauf, 2005]. Any correlation-aware concentration measure would produce the same reading of the maturity axis. **The contribution is the rotation of the axis and the discipline of bounding the claim, not a uniquely discriminating forecast.**
15. **Cartwright [1999]'s discipline is accepted**: the arrangements within which each result holds are stated, and no universality is claimed that the argument does not need.

16. **The one independent reproduction performed is cryptographic and not empirical.** Reproducing a digest confirms arithmetic; **it confirms nothing about the world.**

12. Conclusion

This paper establishes six results in the form the rest of the corpus cites, and the following statements are the fixed text of those results.

R1 — the time-axis index. A constructed maturity stack of $37.0/11.0/7.0/16.5/15.5/13.0$ per cent gives HHI $= 0.222050$, an effective spread of 4.50 buckets, and 55.0% refinancing within three years, against a flat ladder's 0.166667, 6.00 and 50.0%.

R2 — the nine-capital representation split. 63.8500% represented (financial 42% and material 18% fully, intellectual at a stipulated 35% recognition giving 3.85%) and 36.1500% structurally off-ledger, **of which** 11.50% **is marked UNMEASURED.**

R3 — the borrowing-base ceiling. Off-ledger over represented is $36.15/63.85 = 56.62\%$, a base multiple of 1.5662, **with the honest gain bounded strictly between zero and that ceiling and not statable precisely.**

R4 — the repricing arithmetic. At a constructed 2.00% average coupon refinancing at 4.50%, the front-loaded stack adds 1.375% of the stack per year against the ladder's 1.250%, **the real content being that** 55% **reprices in one window.**

R5 — content addressing with canonicalisation underneath. One byte's difference in a settlement object yields a different identifier, while the canonical form folds a homoglyph account string onto the same skeleton as the clean one. **All four identifiers were independently recomputed and reproduce the source exactly.**

R6 — reserve drift as a timescale. A constructed share of 58.0% drifting 0.5 percentage points a year reaches 50% in sixteen years and 48% in twenty, **presented as a timescale illustration with no rival at scale to receive the released share.**

Counter-result. The 56.62% widening is a ceiling on constructed inputs and not a promise; **visibility is explicitly distinguished from pledgeability**; the valuation of non-financial capital is unsolved and is named here as the hardest problem in the framework; **and the architecture is presented as an architecture and not a system**, with no independent implementation, no interoperability profile and no externally maintained conformance suite.

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every input in this paper — the settlement objects, the account identifiers, the maturity distribution, the repricing delta, the private refinancing schedule, the reserve share and its drift rate, and the nine-form capital vector — is **CONSTRUCTED**, **SYNTHETIC** or **STIPULATED** and is labelled as such at each point of use. The numerical results are exact computations on those inputs, carried at full precision and rounded only for printing. **No real issuer, nation, entity, holding, facility or identifier appears anywhere in this paper, and no allegation is made about any of them.**

The digests and identifiers in Section 4 were recomputed from first principles for this paper — SHA-256 over the stated octets, wrapped as a version-1 raw-codec multihash and rendered in lower-case base32 — and reproduce the source manuscript exactly. That computation is reproducible from the byte strings printed in Section 4 with any standard library.

All figures were produced by the scripts named in their captions (`figs/src/CCP-068_fig01.py` through `figs/src/CCP-068_fig05.py`), which import the shared figure style module and read their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

Statistical series, statutes, accounting standards, storage and identifier specifications, and legislative material are described from the source manuscript and are not cited to the bibliography, which contains none of them. Appendix B lists every affected claim. **Named jurisdictions, issuers and instruments are de-branded**, and no claim is made about any named party.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data was processed. Every identifier, entity and holding referred to here is constructed.

Legal, tax, investment and accounting advice

This paper is analytical and is not legal, tax, investment, or accounting advice. It describes a proposed supplementary accounting standard that **no body has adopted and which carries no legal force**, and it measures constructed distributions to demonstrate an instrument. **No instrument is recommended, valued, or characterised as an asset; no figure is a measurement of any real issuer; and no claim is made that any form of capital named here is pledgeable, realisable or enforceable as security.** Accounting recognition rules, disclosure obligations and the treatment of supplementary statements differ by jurisdiction and change; every such statement here is carried from the source manuscript, is flagged in Appendix B, and requires confirmation with qualified counsel and the relevant standard-setter before any reliance. **Supplementary accounts must never be presented as though they had the standing of statutory accounts.**

Appendix A. The reporting standard

Any position measured under the time-axis instrument, and any account prepared under the supplementary standard, reports the following and nothing less.

```
TIME-AXIS MEASUREMENT
 1  share basis in one sentence: principal due, by maturity bucket
 2  the bucket boundaries, held FIXED across measurements
 3  the SOURCE of the schedule and its AS-OF DATE
 4  the bucket table and shares, summing to 100.0%
 5  the time-bucket HHI and the effective spread
 6  the NEAR-WINDOW SHARE alongside the index -- neither alone
 7  the flat-ladder benchmark for the same bucket count
 8  the repricing arithmetic WITH its rate assumptions stated as
```

```
    assumptions and not as forecasts
 9  the diff against the previous measurement, and its date

NINE-CAPITAL ACCOUNT
10  each form named, with its measurement basis or the mark UNMEASURED
11  the represented and off-ledger fractions, reported explicitly
12  the recognition fraction for every partially recognised form,
    stated as a stipulation with its sensitivity
13  NO SINGLE TOTAL across measured and unmeasured forms
14  the ceiling on base widening, labelled a ceiling
15  the integrity layer: canonical form, raw-octet hash, and the
    statement that tamper-evidence is not verification
```

Item 2 is the one that decides whether item 5 means anything, because the index is not invariant under rebucketing. **Item 6 exists because the index is permutation-invariant and cannot say where the concentration sits. Item 13 is the refusal that keeps the account honest**, and item 14 is the one this paper works hardest to protect: a ceiling reported as a promise discredits everything under it.

Appendix B. Claims requiring verification against primary sources

The following rest on statistical, legal, accounting or technical detail, and **none of them is cited to the bibliography, which contains no statistical series, no standards documents, no statutes, no legislative drafts and no vendor or protocol documentation.** They are described from the source manuscript and require confirmation against the governing sources before any reliance.

1. **The content-identifier construction** — the self-describing wrapper, the multihash form, the content codec and the text-encoding prefix, and the current specifications for each. **The digests and identifiers printed in Section 4 were independently recomputed and are correct for the stated strings; what is uncited is the specification that defines the wrapper.**
2. **The framing behaviour of production tooling**, which would wrap a small payload before hashing and yield a different but equally deterministic identifier. The demonstrated property holds under either framing.
3. **The canonical form** and its component transformations, and the confusable data underlying the homoglyph fold. Curzi [2026] is the primary treatment.
4. **The transport and linked-data layers** referred to generically in Section 3.1.
5. **The identifier namespaces** on which any real account would depend, and — critically — **the registries binding each code to a real-world entity.** A hash over a false binding is a cryptographically impeccable name for a falsehood, and no registry is verified here.
6. **The currency-composition series for allocated official reserves**, its level at any date, the treatment of the unallocated portion, and whether reported changes reflect active reallocation or valuation effects. **This paper states no level and describes only a constructed slope.**
7. **The composition of the assets receiving any released reserve share**, and the claim that no single alternative approaches the scale of the incumbent.
8. **The outstanding debt total and the actual maturity distribution of any real issuer. The six-bucket distribution here is CONSTRUCTED and is nobody's profile.** The real distribution, measured on the same buckets, produces its own index, which this paper does not state.
9. **The weighted-average coupon on maturing tranches and prevailing issuance rates.** The 2.00%-to-4.50% repricing is CONSTRUCTED.

10. **The treatment of floating-rate and inflation-indexed tranches**, which reprice differently and would change a real measurement.
11. **The proposition that an issuer borrowing in its own unit of account faces repricing rather than rollover risk**, its limits, and the fact that it does not extend to debt issued in a foreign unit.
12. **The recognition rules of any accounting framework** — what may and may not be an entry, and the prohibition on recognising internally generated goodwill. Carried from Curzi [2026] and subject to that paper's own register.
13. **The nine-form taxonomy's canonical naming and ordering**, and whether this paper's national reading of each form is consistent with it. **The national readings are this paper's gloss.**
14. **The 35% recognition fraction**, which is stipulated and drawn from no framework, and which is the load-bearing parameter of Section 6.
15. **The history of national-income accounting as an invented standard** — the roles of the relevant statistical offices and economists, and the dates of development and adoption. **This paper states the general fact that such a standard was invented and asserts no attribution.**
16. **Any legislative or programme text referred to generically**, including any stated design goals of monetary immunity or counterfeit impossibility. **Those goals are recorded here only to be corrected: no monetary design achieves immunity to inflation or deflation, and no physical or cryptographic instrument achieves counterfeit impossibility. Both are strengthenings and neither is an absolute.** Confirm the exact wording of any such goal before quoting it, and never reproduce it without the correction.
17. **The status of the supplementary standard.** It is a proposal, adopted by no body, with no legal force, and any interaction with statutory accounts requires prior standard-setter acceptance.
18. **The claim that no form of capital named here is pledgeable, realisable or enforceable as security.** The reform concerns how a base is priced, not whether it can be seized.
19. All constructed vectors, rates, drift parameters and schedules are stipulated for illustration and are not claims about any real position.

All numerical results in this paper were computed from the stated inputs at full precision and rounded only for printing.

Chapter 69

What the Architecture Cannot Do: Correlated Backing Baskets, Gated Adoption Probabilities, and the Failure Record

Abstract. This is the corpus's counter-case, assembled deliberately. The governing negative result is that disclosure is not immunity: approximately USD 300 billion of central bank reserves, among the most thoroughly disclosed assets on earth, were immobilised in 2022, and the paper narrows its own principle in print — disclosure removes the failure mode in which a structure collapses on contact with an authority and makes the structure bankable, and it supplies no appeal against a policy decision taken on a correct and fully disclosed position. The backing proposals are measured against themselves. A constructed commodity basket reads 6.640106 effective components naively, but recognising the climate-correlated components as a single 82.0 per cent exposure gives 1.441753, a naive overstatement of 4.605578 times; the basket is genuinely more diversified than a single-commodity anchor at 1.44 against 1.00, and far less than its component count suggests. A 250-to-1 paper-to-physical ratio leaves 99.600 per cent of claims uncollateralised at par and is exhausted by 0.40 per cent of holders standing for delivery, and the object-ledger is credited only with making that number reportable, not with preventing over-issuance. Phased integration is scored as an option: constructed gate probabilities of 0.60, 0.50 and 0.30 give a 0.090 terminal probability, a ratio of 0.15 to first-phase adoption. Adding Gulf jurisdictions raises naive effective jurisdictions 34.81 per cent while effective blocs fall 3.39 per cent.

Keywords: correlated diversification; paper-to-physical ratio; delivery squeeze; gated adoption; margin call; substance challenge; sham and nominee; disclosure is not immunity

Series: D. Structure and Resilience **Paper:** CCP-069

1. Introduction

This is the corpus's counter-case, assembled deliberately. Every other paper in Series D measures a structure and reports what the measurement shows. This one takes the framework's own recommendations, its own governing principle and its own proposals, and reports where they fail.

The governing negative result is stated first because it is the strongest test the framework's central principle has been put to, and the principle lost. The corpus argues that a structure is durable when being observed yields nothing actionable, because every element is lawful, filed and defensible on its own merits [Curzi, 2026,?]. **In 2022, approximately USD 300 billion of the central bank reserves of the Russian Federation were immobilised by the jurisdictions in which they were held.** Those reserves were among the most thoroughly disclosed assets in existence: their location was known, their custodians were known, and their character as sovereign reserves was known and was the reason they were held where they were held. Correctness and disclosure were complete. They were immobilised anyway.

The result is symmetric, and the symmetry is what makes it a finding rather than an anecdote. Disclosure did not prevent the outcome. Concealment could not have prevented it either, because a reserve asset must be held somewhere by a named holder in order to be a reserve asset at all: the custodial relationship that makes the holding usable is the same relationship that makes it reachable. **The variable that decides the outcome is therefore not what is known about the holder but where the holding sits**, and the corpus's answer to arbitrary action is accordingly jurisdictional and rail diversity across coordination perimeters rather than any quality of a structure's paperwork. A fully disclosed, fully correct, single-perimeter structure has an effective bloc count of 1.00, and $\ln(1.00) = 0$ — no available surplus at all, on the corpus's own arithmetic [Curzi, 2026].

The event is described here from the source manuscript and is not cited. This paper's bibliographic database contains no contemporaneous reporting, no sanctions instruments and no jurisdiction-specific legal material, so no citation is attached to the figure, the date, the instruments or the current legal position of the assets. **No approximate citation has been substituted for the missing one.** Appendix A itemises what requires verification: the instruments by jurisdiction, the dates, the amounts by jurisdiction, and the position on both the principal and the income arising from it. The aggregate is stated as approximate because that is how the source states it, and no figure in this paper depends on its precision.

The narrowing that survives the event has to be stated in print or the principle is dishonest.

> **What disclosure does.** It removes the failure mode in which a structure collapses on contact with an authority, because there is nothing to find; it converts an enforcement problem into a policy problem; and it makes the structure bankable, which the corpus identifies as the binding constraint on everything else.
>
> **What disclosure does not do.** It does not prevent a sovereign decision taken on policy grounds against a correct and fully disclosed position; it does not supply an appeal; and it does not protect an asset sitting inside the perimeter that took the decision.

The rest of the paper applies the same treatment to the framework's own proposals. **A constructed backing basket of eight components reads 6.640106 effective components naively and 1.441753 once the climate-correlated components are recognised as a single 82.0 per cent exposure — a naive overstatement of 4.605578 times.** The basket is genuinely more diversified than a single-commodity anchor, at 1.44 against 1.00, and far less diversified than its component count claims. **A constructed paper-to-physical ratio of 250 leaves 99.600 per cent of claims uncollateralised at par and is exhausted by 0.40 per cent of holders standing for delivery**, and the object-ledger the corpus proposes is credited only with making that number reportable, never with preventing over-issuance. **A phased integration programme is scored as an option rather than a plan: constructed gate probabilities of 0.60, 0.50 and 0.30 give a terminal probability of 0.090, a ratio of 0.15 to first-phase adoption.** And **adding two Gulf jurisdictions to a constructed six-country structure raises the naive effective jurisdiction count by 34.81 per cent while the effective bloc count falls by 3.39 per cent** — more jurisdictions, less independence, and every conventional report saying the opposite.

Why the negative results are the contribution. A corpus that publishes only its affordances has not been audited [Lakatos, 1970, Platt, 1964, Meehl, 1990]. The failure record assembled here is a better dataset than any adviser's success record for a reason that is structural rather

than rhetorical: success records are selected on survival and failure records are not [Ioannidis, 2005, Simmons et al., 2011, Open Science Collaboration, 2015]. This paper is the corpus's attempt to be the party that found its own weaknesses first.

1.1 Contributions

1. **Disclosure is not immunity, stated as a narrowing rather than a retraction.** The 2022 immobilisation is the strongest available test of the corpus's governing principle and the principle failed it; what survives is a filter that removes structures collapsing on contact with an authority, and it is not a shield against a policy decision taken on a correct position.
2. **The correlation collapse measured on a backing basket**: naive HHI 0.150600 and 6.640106 effective components against a bloc-collapsed HHI of 0.693600 and 1.441753, an overstatement factor of 4.605578.
3. **The delivery-squeeze arithmetic**: at 250 claims per physical unit, 99.600 per cent of claims are uncollateralised at par and 0.40 per cent of holders standing for delivery exhausts the stock. **A fully collateralised object-ledger makes the ratio reportable and does not prevent over-issuance**, because per-object hash chains supply no consensus and no double-spend prevention [Curzi, 2026].
4. **The gated programme priced as an option**: $P(\text{terminal}) = 0.090$ against $P(\text{first phase}) = 0.600$, a ratio of 0.15, so a proposal that leads with its most ambitious phase stakes the certain value on the least probable branch.
5. **The apparent-diversification trap, measured**: naive effective jurisdictions from 5.290949 to 7.132973 (+34.81 per cent) while effective blocs fall from 2.941176 to 2.841490 (-3.39 per cent) and h falls to -0.134485.
6. **The margin-call failure with no plaintiff**: a 26 per cent fall over four months breaches a 23.1 per cent trigger in month three, and buying survival to a 50 per cent drawdown costs 14.000 of the 40.000 borrowing, a 35 per cent reduction in leverage.
7. **A levy presented as a costless accounting correction reallocates burden**: a constructed incidence puts 42.0 per cent on one sector, with a constructed 55 per cent passing through to consumers.

What this paper does not claim. No instrument described here is claimed to exist, to be issued, to be authorised, or to be lawfully issuable in any jurisdiction. No price, yield or return figure appears anywhere in it. No commodity multiple, reserve quantity, statute, regulation, administrative notice or decided case is asserted as a fact; every one is described from the source manuscript and itemised in Appendix A. **And the licensing provisions recorded in §5.2.4 name the author's own company. That is a conflict of interest of the first order, it is declared in §3.4 and again in the Declarations, and the provisions appear here with their exposure attached rather than either endorsed or omitted.**

2. Hypotheses

Six hypotheses, each stated so that it could be wrong, each with its falsifier [Platt, 1964, Popper, 1959]. **One of them was tested against a real event and failed, and it is reported in the same plain language a success would have received.** Statements that are not falsifiable are listed at the end and labelled.

H1 — the disclosure hypothesis. TESTED AND FALSIFIED. A structure every element of which is lawful, disclosed and correct is durable against an authority with jurisdiction over it, because observation yields nothing actionable. *Predicted:* a fully disclosed, fully correct, lawfully held position is not seized or immobilised. *Falsified by:* one instance of a fully disclosed, correct, lawfully held position being immobilised. *Result:* **falsified.** Approximately USD 300 billion of fully disclosed central bank reserves were immobilised in 2022. The hy-

pothesis in that form is dead; §6.2.8 states the narrower claim that survives, and §8.1 states why the surviving claim is still worth having.

H2 — the correlation hypothesis. A nominal component count overstates a backing basket's diversification whenever the components share a common shock, and the overstatement is computable. *Predicted:* on the constructed eight-component basket, naive 6.640106 effective components against a correlated 1.441753, an overstatement factor of 4.605578, with the basket still more diversified than a single-commodity anchor at 1.44 against 1.00. *Falsified by:* a basket whose bloc-collapsed effective count equals its naive count under a defensible correlation grouping. *Status:* the arithmetic is exact on the stipulated weights; **the correlation grouping is a stipulated modelling assumption and not a measured covariance structure, and no real basket was examined.**

H3 — the squeeze hypothesis. At a paper-to-physical ratio R, a fraction $1/R$ of claimholders standing for delivery exhausts the deliverable stock, so the ratio is a latent concentration that is invisible in normal function and decisive in a squeeze. *Predicted:* at $R = 250$, the threshold is exactly 0.40 per cent, with coverage 2.0000 at 0.20 per cent, 1.0000 at 0.40 per cent, 0.4000 at 1.00 per cent and 0.0800 at 5.00 per cent. *Falsified by:* a market meeting delivery obligations beyond $1/R$ without new physical supply entering. *Status:* arithmetic on a stipulated ratio. **No commodity's ratio is measured here and none is asserted**; the range circulating in the source is described and itemised in Appendix A.

H4 — the gating hypothesis. In a strictly dependent phased programme whose gate probabilities are all below one, most of the expected value sits in the phase that stands alone. *Predicted:* at $p_1 = 0.60$, $p_2 = 0.50$, $p_3 = 0.30$, $P(\text{reach phase } 2) = 0.600$, $P(\text{reach phase } 3) = 0.300$, $P(\text{terminal}) = 0.090$, and a terminal-to-first-phase ratio of 0.15. *Falsified by:* a dependent phasing with gate probabilities below one in which a later phase's standalone value exceeds the first's. *Status:* arithmetic; **the probabilities are constructed and are estimates of nothing.** The conclusion depends only on dependence and on gates below one, which is why it is stated as a structural property rather than as a forecast.

H5 — the apparent-diversification hypothesis. Adding jurisdictions inside a coordination perimeter the structure is already heaviest in raises the naive effective jurisdiction count and lowers the effective bloc count. *Predicted:* naive 5.290949 → 7.132973 (+34.81 per cent) while blocs 2.941176 → 2.841490 (−3.39 per cent) and h −0.122417 → −0.134485. *Falsified by:* an addition inside an existing bloc that raises both counts. *Status:* arithmetic on a stipulated allocation. **The bloc assignment is this framework's own, is contestable, and is the load-bearing input**; the corpus measures the sensitivity of the index to assignment error separately [Curzi, 2026].

H6 — the margin-trigger hypothesis. The first maintenance call on a pledged position falls at $d^* = 1 - \text{LTV}_0/(1-m)$, and this is very much smaller than the drawdown a holder associates with the word *risk*. *Predicted:* at $\text{LTV}_0 = 50$ per cent and $m = 35$ per cent, the first call falls at a 23.076923 per cent drawdown; surviving 50 per cent requires the loan to fall from 40.000 to 26.000. *Falsified by:* a margin agreement at those parameters whose first call falls elsewhere. *Status:* algebra on a stipulated contract [Curzi, 2026]. **The 35 per cent maintenance requirement is a constructed model parameter and not a market quotation.**

Statements in this paper that are not hypotheses, and are labelled as such.

- *The three backing mechanisms* are **proposals** [T3]. None is adopted, none has legal force, and no instrument described is claimed to exist or to be lawfully issuable anywhere.
- *The licensing provisions of §5.2.4* are **the Act's specifications, recorded**. They are neither endorsed nor omitted, and the review they invite is named at the level of its general nature with **no outcome stated.**
- *The Fibonacci denominations and the fractal reserve rule* are **design choices whose monetary significance is graded at zero.** The combinatorial property is real; nothing monetary follows from the choice of sequence.
- *The stated goals of monetary immunity and of counterfeit impossibility* are **the Act's goals, recorded and corrected in the same sentence.** No anchor confers immunity and no instrument confers impossibility.

- *The nine-form backing taxonomy* is a **classification, not a valuation method.** No measurement basis for its non-financial components exists in the source corpus or here [Curzi, 2026t,?].
- *The naming rationale and numerological layer of the reference implementation* are **claims deliberately not made** [Curzi, 2026].

3. Method

3.1 What is measured, what is proposed, and what is only described

Three kinds of statement appear in this paper and the method keeps them apart. The evidence tiers used as inline tags throughout §§4–6 are defined here and nowhere else.

tier	meaning
T1	established fact — a published standard, or a proven mathematical or distributed-systems result
T2	measured by a deterministic computation over a stipulated, constructed input set, recomputed for this paper
T3	proposed design or standard, not adopted, not validated at scale, carrying no legal force
T4	refuted, unsupported as stated, or a claim deliberately not made

The tier scheme is this paper's own ordinal labelling convention and has no external standing. It is not averaged, summed or scored, and an ordinal scale licenses none of those operations [Stevens, 1946, Michell, 1997].

Every numerical result in this paper is exact arithmetic on stipulated, constructed inputs. No market, commodity, reserve, land parcel, economy, institution, court file or family appears anywhere in it. Each construction is printed at the point of use and labelled CONSTRUCTED there, so that a reader can reproduce the figure or show that a different stipulation gives a different one.

3.2 Which effective count, and what the instrument does not see

Where this paper reports an effective count it means $1/\mathrm{HHI}$, the Hill number of order two [Hirschman, 1964, Herfindahl, 1950, Hill, 1973, Jost, 2006], independently corroborated as the numbers-equivalent of a concentration index in the competition and political-science literatures [Adelman, 1969, Laakso and Taagepera, 1979, U.S. Department of Justice and Federal Trade Commission, 2023], and used in that sense throughout Series D [Curzi, 2026]. **The corpus's affine construction $1 + (N-1)u$ is not a Hill number of any order and does not appear in this paper.**

Two limits of the instrument are load-bearing here and are stated before any result is reported. **It does not see a pledge:** a pledged position and an unpledged position have identical jurisdictional shares and produce identical effective counts, which is why §6.2.4's failure track is invisible to it. **And it does not see capacity:** an effective count measures independence and says nothing about whether any perimeter can carry the load [Curzi, 2026].

3.3 The disclosure test, and what a test is

The framework's governing constraint is that every element of a structure should remain effective when disclosed to every authority with jurisdiction over any part of it. That is not an ethical preference: it is the 1883 principle that a system must remain secure when everything about it except the key is public [Kerckhoffs, 1883], restated as *open design* [Saltzer and Schroeder, 1975] and given its transparency-as-remedy form [Brandeis, 1914]. The regulatory environment has independently removed non-disclosure from the available design parameters [Organisation for Economic Co-operation and Development, 2017, Financial Action Task Force, 2012, Organisation for Economic Co-operation and Development, 2015, de Goede, 2011], and the corpus states the resulting constraint exactly [Curzi, 2026].

A filter is not a shield, and this paper exists partly to make the difference impossible to misread. The test removes structures that collapse on contact with an authority. It says nothing whatever about a policy decision taken against a disclosed and correct position, which is what happened in 2022, and §6.4 lists confusing the two as a named failure mode.

3.4 Provenance, exclusions, and the conflict of interest

The licensing provisions recorded in §5.2.4 name 36N9 Genetics, which is the author's own company, as the sole-source procurement counterparty for the production of the proposed currency, alongside a perpetual royalty and a sovereign licence from a private entity. That is a conflict of interest of the first order. It is declared here, in the Declarations, and in the section where the provisions appear. **The method adopted in response is the only one available to a single author: record the provisions exactly as the Act specifies them, name the review they invite in the same neutral register the framework applies to every other legal exposure, state no outcome, and leave the decision with the reader.** A paper that recorded the provisions without the exposure would be advocacy; a paper that omitted the provisions would be concealment; and the author's interest in one of the provisions is precisely why both halves are in print.

Two elements of the source material are excluded from every technical claim in this paper. The naming rationale of the proposed system, and a numerological layer present in its reference implementation, are recorded as declared, are graded [T4], and confer no property on any structure described [Curzi, 2026]. They are mentioned only so that a reader encountering them in the source knows they were assessed and discarded rather than overlooked.

No jurisdiction-specific legal source is cited anywhere in this paper. The decided cases, the statutes, the tax-code provisions, the administrative notices, the procurement and appropriations frameworks and the constitutional provisions named in §§5 and 6 are described from the source manuscript and are **left uncited, because this paper's bibliographic database contains none of them.** Appendix A itemises each with what specifically requires verification. What the bibliography does supply is the general doctrinal and comparative frame, and that is cited where it is relied on.

3.5 What is not invoked

The corpus's reflexive loop is not used in this paper and no social-influence parameter is estimated, so the reflection problem — that the endogenous effect in a linear-in-means reflexive specification is not identified from observational data [Manski, 1993, Shalizi and Thomas, 2011] — does not bear on any result reported here. **No distributional, scaling or heavy-tail claim is made anywhere**, so the goodness-of-fit and likelihood-ratio procedure such a claim would require [Clauset et al., 2009] is not applicable and was not run. Every figure in this paper is a ratio of counts, a probability product, or a closed-form solution of a linear equation on stipulated inputs.

One borrowing is flagged before it is used. The framework writes concentration, surplus and resilience in the grammar of conserved and extensive quantities without a conservation

law, which is the characteristic failure of this genre [Mirowski, 1989]; the defence taken here is Cartwright's — validity is claimed within a specified arrangement and no universality is claimed for it [Cartwright, 1999]. §9.6 grades every borrowing in the paper by name.

4. Backing mechanisms: the paper-to-physical ratio, a correlated basket, and a two-layer reserve

4.1 The position: one distinction applied three times

The corpus's reserve-drift measurement recorded a compromised position and refused to dramatise it: a slow reserve drift and a measurable refinancing concentration, both fully disclosed [Curzi, 2026], whose time-axis Herfindahl index of 0.222050 gives an effective spread of 4.50 maturity buckets with 55.0 per cent refinancing inside three years. The nine-capital standard answered it with the most modest instrument available — an accounting standard that makes the off-ledger asset side visible, minting nothing and licensing nothing. This section takes the next step the American Roots material proposes, and it is a step onto more contested ground, because it moves from *what a nation counts* to *what a currency is backed by*. That is a harder question and the discipline has to be correspondingly tighter. The section develops three backing mechanisms, and it develops each one twice: once as the proposal states it, and once as the framework's own instruments cut it down to what is defensible.

The three mechanisms are a paper-to-physical ratio argument, a living-basket anchor called the Fermentables Index, and a conservation-easement reserve paired with a floating voucher. They look unrelated until they are laid against the one distinction this framework has carried since the store–rail rule, and against which every rail section and every reserve section has been measured:

> **A rail is not a holding.** A store of value and a medium of settlement are two different jobs, and the deepest error in monetary design is asking one instrument to do both while pretending it does neither badly.

Every one of the three mechanisms is, at bottom, a claim about where a given design sits on the store-versus-rail axis, and the section's organising move is to place each one on that axis explicitly and refuse to let it drift off. The paper-to-physical argument is about what happens when a market confuses a rail — a paper claim — with a store — the commodity itself. The Fermentables Index is a reserve-value anchor, a store, that must not be mistaken for a settlement medium. The conservation-easement reserve is a store so illiquid it cannot settle anything at all, which is precisely why it is paired with a voucher that is nothing but a rail. **The store-versus-rail split is not one of several themes in this section; it is the section, applied three times to three designs that each get it right or wrong in a different way.**

The first mechanism leads, because it is the framework's strongest empirical monetary claim, and it is strongest for the reason the refinancing-wall measurement was strongest: it is a concentration, and concentration is the one thing this framework has an instrument for. Derivative and paper claims on a commodity routinely exceed the deliverable physical stock behind them by a wide margin — commonly cited across a range of roughly one to three orders of magnitude, which is to say somewhere in the vicinity of fifty times to five hundred times, depending on the commodity and the venue. That range is stated as a range on purpose, and every specific multiple that circulates inside it — the silver open-interest-to-registered figures, the split between registered and eligible gold on a delivery exchange, the ratio of crude open interest to the deliverable stock at a single pricing point — is flagged in the Verification Register and asserted nowhere in this section as a fact. What the section asserts is the shape, and the shape is not in dispute even where the magnitudes are: a paper layer many times larger than the physical layer beneath it is a documented and ordinary feature of commodity markets, and it decouples from the physical under stress in ways that have been observed.

This is the concentration paper's failure mode rotated onto a new axis [Curzi, 2026], whose Result Two takes a redesigned six-jurisdiction structure to naive HHI 0.189002 — 5.29 effective jurisdictions — and bloc-HHI 0.340000, 2.94 effective blocs. That paper measured concentration on the jurisdiction axis; the reserve-drift paper rotated the same instrument onto the maturity axis and measured a refinancing wall. **This section rotates it once more, onto the physical-settlement axis, and measures the concentration of paper claims against deliverable stock** — and the result is the same species of finding each time, that a nominal count of many independent-looking things collapses, under the right stress, onto a single point of failure. The paper claims look like a deep, liquid, diversified market until enough holders want the physical thing at once, and then the ratio that was invisible becomes the only number that matters.

The second and third mechanisms are answers rather than diagnoses, and they answer the problem the reserve-drift measurement named: the dollar's reserve share drifts because there is no bloc-independent alternative reserve asset at its scale, and both the Fermentables Index and the easement reserve are attempts to build a reserve anchor with a *different failure profile* than a debt-fiat one. Neither is offered here as a settlement medium, because neither is one, and the honesty of the section is entirely in keeping that line. A consumables basket resists hoarding and resists pure speculation, at the cost of storage and measurement complexity that a debt claim does not carry — but it is a store, not a rail. A conservation easement is a real, valued, tradeable legal instrument, and it is also completely illiquid — you cannot settle a trade in one — so it can back a currency only as a reserve and never as a means of payment, which is why the design pairs it with a voucher that carries the settlement job it cannot. **The recurring result is that each mechanism is defensible exactly to the extent that it is honest about which of the two jobs it is doing, and indefensible exactly where it claims to do both.**

One constraint from the nine-capital standard carries directly into this section and is stated here so it governs everything below. The American Roots Act attaches to its instruments a set of design goals, and two of them — monetary immunity to inflation and deflation, and a counterfeit-impossibility property — were recorded in the nine-capital standard as the Act's stated aims and corrected in the same breath, because no anchor confers immunity and no instrument confers impossibility. That correction is not re-litigated here; it is inherited. Where the Fermentables Index restates the immunity goal, this section records it as a goal and corrects it once more in the specific setting of a consumables basket, and it does the same for every other absolute the material reaches for. The house standard is the integrity-layer paper's: a corpus that grades its own claims and prints its own negative results is worth more than one that arrives without them, and this section holds the standard against the very mechanisms it is developing.

4.2 The three mechanisms

4.2.1 Mechanism one: the paper-to-physical ratio, and why it is normal until it is not

Start with what the ratio is, because most of the heat around it comes from misunderstanding a feature as a scandal.

A commodity futures or forward market is a market in claims. A contract is a promise to deliver, or to take delivery of, a quantity of the underlying at a date, and the overwhelming majority of contracts are never carried to delivery — they are closed, rolled, offset, or cash-settled before the delivery window, because the parties holding them are hedging price, not sourcing metal or barrels. A miner sells forward to lock a price; a fabricator buys forward to lock an input cost; a speculator takes the other side of each for a fee in risk-bearing. None of them wants the physical thing; all of them want the price exposure. The result is that the notional quantity of the commodity represented by outstanding claims can, and routinely does, exceed the quantity of the commodity actually sitting in deliverable form in the registered warehouses behind the market, by a large multiple. This is not a malfunction of the market;

it is what the market is for, and a hedging market whose paper never exceeded its physical would be a market too thin to hedge in. **No source in this paper's bibliography establishes the magnitude of that multiple in any commodity, and none is cited for it**; the range is described from the source manuscript and itemised in Appendix A.

The multiple is the concentration measurement, and it is where the discipline has to hold. The figures that circulate are specific and they are contested. Silver is the commodity most often cited, with open-interest-to-registered-stock ratios quoted anywhere from roughly one hundred to several hundred times, and those figures are contested because they depend entirely on which numerator and which denominator are chosen — all open interest or only the front month, registered stock only or registered plus eligible, one venue or the global aggregate. Gold on a delivery exchange is discussed in terms of the split between registered stock, which is committed to delivery, and eligible stock, which sits in the same vaults in approved form but is not pledged to any contract — and the ratio of open contracts to registered stock moves sharply depending on that classification. Crude oil is discussed in terms of the open interest in a contract that settles by physical delivery at a single location against the deliverable and storage capacity actually available at that location. Every one of those multiples is flagged in the Verification Register and none is asserted here as a fact; the section's claim is the range and the mechanism, not any single number inside it.

The mechanism is what makes the range matter, and it is a store-versus-rail confusion waiting to happen. A paper claim on a commodity is a rail: it is a device for transferring price exposure, and for the holder who will close it before delivery it is a position that closes itself, exactly the transit-rail definition of a rail. The commodity in the warehouse is a store: it is the thing that is there because that is where it lives. For as long as everyone treats the paper as a rail — a transfer of exposure to be closed out — the ratio between paper and physical is irrelevant, because almost nobody is asking the rail to deliver the store. The ratio becomes the risk at exactly the moment enough holders stop treating the claim as a rail and start treating it as a store — when they stand for delivery, demand the physical thing, and discover that the claims outnumber the deliverable units by the multiple that was invisible the day before. That is a delivery squeeze, and it is the physical-settlement-axis form of a bank run: a set of claims that were fungible and liquid so long as they were not all presented at once, revealed as fractional the instant they are.

The empirical proof that paper and physical decouple under stress is not theoretical and this section does not need to construct it. The 2020 event in which a physically-settled crude contract traded at a deeply negative price is the clearest demonstration: holders of expiring long positions who could not take or store physical delivery faced a market in which the paper claim and the physical reality had come completely apart, and the claim's price went to a place no store of value can go, because it was never a store — it was a rail that had run out of track. Periodic silver squeezes, in which a coordinated or spontaneous rush to stand for delivery collides with the registered stock, are the same mechanism in the other direction. These events are the empirical content of the claim that paper and physical decouple, **and they are described from the source manuscript and are not cited, because this paper's bibliography contains no contemporaneous market reporting** (Appendix A). They are why the ratio is a genuine risk measurement rather than a conspiracy theory — the decoupling is observed, dated, and priced, whatever one believes about why the ratio is as large as it is.

Now the tie to the framework, which is exact. This is the concentration failure, measured on the physical-settlement axis instead of the jurisdiction axis. The paper claims are the many independent-looking things; the single deliverable stock they all point to is the coordination perimeter they collapse onto; and the delivery squeeze is the shock that reveals the collapse. And it is the store–rail rule stated in commodities: **the paper claim is the rail, the commodity is the store, and confusing the two is the error the whole framework has warned against — an error that is harmless in normal function and catastrophic in a squeeze, which is precisely the profile of every concentration this framework measures.**

4.2.2 Mechanism one, continued: what a fully collateralised object-ledger changes

The VINO claim enters here, and it must be stated at exactly its strength and no more, because the temptation is to overstate it and the integrity-layer paper spent its length showing why that temptation is fatal.

The claim is this. In a conventional commodity market the paper layer and the physical layer live in two different systems. The claims live in the clearing and exchange infrastructure; the physical units live in a warehouse registry of receipts. The ratio between them — the concentration measurement that becomes the squeeze risk — is not a number either system publishes as a live quantity, because neither system is built to. It is recoverable only by reconciling the two registries against each other, a join across two datasets that no one runs continuously and that many participants have no standing to run at all. The ratio is invisible not because it is hidden but because it is nobody's field: it is a relationship between two ledgers, and a relationship between ledgers is exactly the kind of quantity the double-entry-limits paper showed a double-entry system structurally cannot carry.

A fully-collateralised, canonically-audited ledger of the kind the integrity-layer paper specified changes the visibility of that ratio, and only its visibility. In the object-as-ledger factoring, each physical unit is its own hash-chain, seeded by a genesis root — GridChain's per-token construction, taken down to the individual deliverable unit. A claim, in such a system, is not a free-floating contract in a separate venue; it is an entry that references a specific unit's chain. Under a full-collateralisation rule — one live claim per unit — the number of valid claims the ledger can carry is bounded by the number of units, and the ratio of claims to units is not a reconciliation exercise across two systems but a single aggregate over one index, readable live. A per-unit ledger cannot make five hundred claims on one unit disappear, but it can make them impossible to hide: over-issuance stops being a relationship between two registries that nobody joins and becomes a count on one register that anyone with read access can take. That is the whole of the claim, and it is a visibility claim, in exactly the sense the nine-capital standard was careful about when it said visibility is not pledgeability. Here, visibility is not prevention.

The honesty that keeps this defensible is the integrity-layer paper's audit finding A1, and it is stated in print because omitting it would be the exact failure the integrity-layer paper exists to prevent. A per-object hash-chain provides tamper-evidence and local write-parallelism; it does not by itself provide consensus, total order, or double-spend prevention [Curzi, 2026], whose counter-result states in one place that the substrate supplies no consensus, no total order, no double-spend prevention and no Sybil resistance. Agreement among mutually distrustful parties is the Byzantine problem [Lamport et al., 1982], its impossibility under asynchrony is fixed [Fischer et al., 1985], the availability–consistency trade is fixed [Gilbert and Lynch, 2002], and the party count must itself be audited [Douceur, 2002]. Nothing in the object-as-ledger structure alone stops the same unit being pledged to two claims on two forks that each verify as internally valid — that is A1, reproduced and graded [T4] in the integrity-layer paper, and it applies here without exception. **So the object-ledger does not arithmetically prevent a 250-to-1 over-issuance; it makes the 250-to-1 ratio a visible, queryable number in real time, and prevention still requires the ordering and identity layers that the integrity-layer paper flagged as unbuilt.** The correct statement of the mechanism is therefore precise and modest: full collateralisation plus object-as-ledger converts the paper-to-physical ratio from an off-venue relationship that surfaces only in a squeeze into an on-ledger count that surfaces continuously — which changes when the concentration is discovered, not whether it can occur. That is a real improvement, because a concentration discovered continuously can be governed and a concentration discovered in a squeeze cannot, and it is a smaller improvement than the material's enthusiasts would claim, which is why the section states both halves.

4.2.3 Mechanism two: the consumables basket as a reserve anchor

The second mechanism is a reserve anchor, and it is a store by construction — the failure profile argument below is an argument about what kind of store it is, never a claim that it

settles anything.

The Fermentables Index, as the American Roots Title IV material describes it, is a dynamic basket of consumable, regenerative natural assets: grains, fruit, dairy, timber, freshwater, and a set of soil-health and biodiversity metrics that stand for the productive capacity of the land itself. The design intent is that a currency anchored to such a basket is anchored to the things people actually consume and the systems that regenerate them, rather than to a debt instrument or a single monetary metal. **The Act states as a design goal that such an anchor would confer mathematical immunity to inflation and deflation, and that goal is recorded here as a goal and corrected in the same sentence it is stated: no anchor confers immunity, and a consumables basket does not eliminate monetary risk — it changes the risk's shape.** The corpus's channel-gain postulate is the general form of the correction [Curzi, 2026ab]: an effect is transmitted only through a channel with a gain, and where no channel exists the effect is zero at every magnitude of input. This is the nine-capital paper's correction applied to the specific instrument. A currency anchored to a basket of staples is stable in terms of that basket and variable in terms of everything outside it; its purchasing power tracks the supply and demand for the specific staples chosen; and it is exposed to shocks in those staples — a harvest failure, a shift in diet, a change in what the basket ought to contain — that a broader or a purely monetary anchor would diffuse differently. The immunity is not available; a different risk profile is, and the different risk profile is the real and defensible claim.

The defensible claim has two halves and both are properties of consumables specifically. The first is resistance to hoarding: the basket's components are perishable, and a store made of perishable things cannot be cornered and sat on the way a monetary metal can, because holding it incurs a physical carrying loss that holding gold does not. The second is resistance to pure speculation: the components are universally consumed, so their demand is anchored in use rather than in belief, and a speculative position in a thing everyone eats is disciplined by the fact that everyone eats it. A consumables basket resists both of the pathologies a monetary anchor is prone to — the hoard and the bubble — and it pays for that resistance in storage and measurement complexity, which is a real and heavy cost, not a rounding error. Grain must be stored and rotated; dairy spoils fast; freshwater and soil-health and biodiversity are not even uniformly measurable, and a basket that includes them inherits the nine-capital paper's problem of forms of capital with no accepted measurement basis. The anchor is a reserve of value with an unusual and in some ways attractive failure profile, bought at the price of being far harder to hold and to measure than a claim on a government.

The measurement that makes this precise is the concentration instrument, and pointing it at the basket yields the sharpest single result in this mechanism. **A basket is a concentration instrument applied to backing: the same Herfindahl arithmetic that measures whether a private estate is diversified across jurisdictions measures whether a reserve is diversified across the things that back it.** The effective count used throughout this paper is $1/\mathrm{HHI}$, the Hill number of order two [Hirschman, 1964, Herfindahl, 1950, Adelman, 1969, Hill, 1973, Jost, 2006, Laakso and Taagepera, 1979], and **never the corpus's affine construction $1+(N-1)u$, which is not a Hill number of any order and does not appear in this paper.** A single-commodity anchor — a gold standard, a single-metal peg — has a backing HHI of 1.0 and an effective component count of exactly one; it is a one-asset reserve and its fate is the fate of that one asset. A fiat anchor is, on this axis, also effectively a count of one, because its value rests on the policy of a single issuer. The Fermentables basket, constructed below, has a naive effective component count well above one — and, exactly as in the concentration paper, the naive count overstates the diversification, because the components are not independent. A drought that takes the grain crop takes the fruit, stresses the freshwater, degrades the soil-health metric, and raises dairy feed costs all at once, and those correlated components collapse, under a common climate shock, onto far fewer effective independent components than the naive count reports. **That is the corpus's own correlation result on a different axis** [Curzi, 2026]: holding marginal failure at 10 per cent, a common shock of $c = 0.02$ raises the probability that all four components fail together from 0.000100 to 0.020044, a 200.4-fold understatement, and over fifteen years 0.001499 independent against 0.794109 correlated. The same finding is standard in the portfolio and ecological literatures — correlations rise in the tail [Longin and Solnik,

2001, Ang and Chen, 2002], diversification is a claim about a covariance structure rather than a count [Markowitz, 1952, Elton and Gruber, 1977], statistical averaging is not the same mechanism as response diversity [Doak et al., 1998, Tilman et al., 1998, Schindler et al., 2010], and more components do not in general give more stability [May, 1972, McCann, 2000, Ives and Carpenter, 2007]. The honest number is the collapsed one, and it is computed in the Worked Application, where the basket's naive effective count of well over six falls to under one and a half once the climate-correlated components are recognised as one exposure. The basket is genuinely more diversified than a single-commodity anchor and dramatically less diversified than its component count suggests — which is the concentration paper's exact finding, that the count is not the perimeter, arriving on the backing axis.

4.2.4 Mechanism three: the easement reserve and the floating voucher

The third mechanism is the clearest store-versus-rail design in the section, because it is built as two explicit layers, one of which is a pure store and the other a pure rail, and the honesty of the whole thing is in never confusing them.

The store layer is conservation-easement backing. A conservation easement is a real legal instrument with a defensible foundation: under the United States Internal Revenue Code the qualified conservation contribution regime recognises the permanent restriction of land use for conservation purposes as a contribution with measurable value — **a statutory provision described here from the source manuscript and cited nowhere, because this paper's bibliography contains no tax-code material** (Appendix A), and the broader ecosystem-services literature is said to establish that the natural capital such land holds — watersheds, carbon, habitat, soil — has value that can be estimated by accepted, if contested, methods. **That literature is not in this paper's bibliography and no work in it is cited**; what the corpus can supply is the general governance result that a commons is neither necessarily tragic nor automatically well managed [Hardin, 1968, Ostrom, 1990, 2010, Cox et al., 2010]. The design premise is that currency can be backed partly on protected land whose natural-capital value persists *because* it is stewarded rather than extracted: the easement is what makes the persistence legally durable, and the natural capital is what gives the reserve its value. That premise has a defensible core — easements are real, valued, tradeable instruments and natural capital has measurable value — and it has two honest limits that must be stated in print, because the mechanism is indefensible without them.

The first limit is valuation. Easement valuation is genuinely contested, and it has a documented abuse history that a proposal to back currency on easements cannot wave away. The syndicated conservation easement — in which a partnership acquires land, obtains an aggressive appraisal, places an easement, and passes an outsized deduction to investors — has been the subject of sustained enforcement attention, is a designated listed transaction for disclosure purposes, and is the subject of active litigation over both the valuations and the designation itself. The specific administrative authority that designated syndicated conservation easements as a listed transaction, and the litigation surrounding it, are flagged in the Verification Register and their notice number is not asserted here — the point that survives verification is structural and does not depend on the citation: easement valuation is contested, it has been abused through inflated appraisal, and a reserve built on easement values inherits exactly that valuation contest. The second limit is liquidity, and it is absolute: an easement is illiquid. You cannot settle a trade in a conservation easement. It is a permanent restriction on a specific parcel, held by a specific qualified organisation, and it does not circulate, does not clear, and cannot be handed to a counterparty to extinguish an obligation. **An easement can be a store of value and it can never be a means of settlement, and any design that treats it as spendable has committed the store-versus-rail error in its most basic form.**

Which is exactly why the design does not treat it as spendable. The rail layer is the floating voucher, and it is the settlement-side counterpart to the illiquid reserve. The voucher is a transferable credit that circulates against the stewarded reserve: the reserve is the store — illiquid, backing, staying where it is — and the voucher is the rail — liquid, transacting, mov-

ing between holders and closing itself the way any rail position does. **The two-layer design is the store–rail rule built into a monetary instrument on purpose: it makes an illiquid natural-capital reserve spendable without pretending for a moment that the reserve itself is liquid, by putting the liquidity in a separate instrument that is honest about being a rail.** This is the single most important structural point in the mechanism and it is what separates it from a naive natural-capital currency that would try to make the land itself circulate. The land does not circulate. The voucher circulates. The land backs the voucher. That is a coherent design precisely because it respects the distinction the rest of the framework has spent fifty-four sections insisting on, and it is the honest answer to the question of how you make a reserve that cannot be traded do monetary work: you do not trade the reserve, you trade a claim that references it, and you keep the two on opposite sides of the store-rail line and never let them merge.

The funding mechanism for maintaining the stewarded reserve is what the Act calls the 12% Life-Energy Equity, and it must be presented as the Act frames it and then read honestly. The Act frames it not as a tax but as an accounting correction: a commons-replenishment allocation that recognises, and directs to the upkeep of the stewarded reserve, a share of value that was previously drawn from the natural commons without being priced at all. In the Act's framing this is the same move the nine-capital standard made for the whole nine-capital standard — taking a real, off-ledger quantity, the drawdown of natural capital, and putting it on the ledger — and calling it a correction rather than a levy is internally consistent with that framing. **But an accounting correction with a rate is still a design choice with distributional effects, and recording the previously-unpriced cost of natural-capital drawdown on the ledger reallocates that cost to whoever was implicitly benefiting from its being unpriced.** Whether the correction is called a tax or not, someone bears its incidence, and the incidence is not uniform: it falls most heavily on the activities that drew most heavily on the unpriced commons, and it passes through, in part, to the consumers of those activities' output. The Worked Application models a constructed incidence to show the shape; the point of principle here is that the 12% is presentable as the Act presents it and must also be modelled for its distributional effect, and a section that recorded only the framing and not the incidence would be doing the thing this framework calls lying by omission.

4.3 Specification

4.3.1 The three mechanisms, specified as proposals

Each mechanism is graded [T3] — a design, not an adopted standard — and each is specified by what it does, what job it does on the store-rail axis, and, most importantly, what it does not do.

```
MECHANISM 1: FULLY-COLLATERALISED OBJECT-LEDGER FOR PHYSICAL CLAIMS
-------------------------------------------------------------------
JOB            makes the paper-to-physical ratio VISIBLE and live
STORE / RAIL   the physical unit is the STORE; the claim is a RAIL;
               the ledger keeps them distinguishable and counts both
RULE           each deliverable unit is its own hash-chain (CCP-065); each claim references a specific
               full collateralisation = one live claim per unit
REPORTS        claims-per-unit as a live aggregate over one index,
               not a reconciliation across two registries
DOES NOT       prevent over-issuance, provide consensus, provide
               total order, or prevent double-spend on its own --
               those need the ordering + identity layers (CCP-065 A1)
STATUS         [T3] proposal; the substrate is [T2]-validated for
               tamper-evidence only; prevention is unbuilt

MECHANISM 2: THE FERMENTABLES INDEX (reserve-value anchor)
```

```
-----------------------------------------------------------------------
JOB             a reserve anchor with a consumables failure profile
STORE / RAIL    a STORE only; NOT a settlement medium; a currency may
                be anchored to it but is not made OF it
RULE            a weighted basket of consumable regenerative assets;
                weights published; the backing HHI and effective
                component count reported NAIVE and CORRELATION-COLLAPSED
REPORTS         naive effective components, correlated effective
                components, and the overstatement factor between them;
                perishability drag vs a non-perishable anchor
MARKS           soil-health and biodiversity components UNMEASURED
                where no accepted basis exists (CCP-068 discipline)
DOES NOT        confer immunity to inflation or deflation; settle a
                trade; remove monetary risk -- it reshapes it
STATUS          [T3] design goal; immunity claim [T4], corrected

MECHANISM 3: EASEMENT RESERVE + FLOATING VOUCHER (two-layer)
-----------------------------------------------------------------------
JOB             back currency on stewarded natural capital, spendably
STORE / RAIL    the easement reserve is the STORE (illiquid, backing);
                the floating voucher is the RAIL (liquid, transacting)
RULE            reserve held as IRC 170(h) conservation easements,
                valued by a stated (contested) methodology, marked
                contested; vouchers issued against the reserve
                circulate and settle; the reserve never settles
FUNDING         the 12% Life-Energy Equity maintains the reserve,
                framed by the Act as an accounting correction, and
                modelled here for distributional incidence
DOES NOT        make the reserve liquid; make the easement settleable;
                resolve the easement valuation contest; escape the
                syndicated-easement abuse history (Register-flagged)
STATUS          [T3]; the legal instrument is real; the valuation and
                the voucher peg require a process this framework cannot run
```

The "does not" rows are the specification. As in the integrity-layer paper and the nine-capital standard, the honesty of each mechanism lives entirely in what it refuses to claim, and the three refusals are the same refusal three times: the object-ledger does not prevent, only reveals; the basket does not immunise, only reshapes; the easement does not settle, only backs. Each mechanism is defensible precisely to the degree that it does one job well and disclaims the other, and the specification's entire purpose is to pin down which job is which and forbid the merge.

4.3.2 The store-versus-rail classification, stated once as a table

Because the whole section is one distinction applied three times, it is worth stating the classification in one place, so that no mechanism can be quietly promoted from the job it does to the job it does not:

```
mechanism                STORE (backing)         RAIL (settlement)
-----------------------------------------------------------------------
paper / physical         the physical unit       the paper claim
Fermentables Index       the basket              -- (none; anchor only)
easement + voucher       the easement reserve    the floating voucher
-----------------------------------------------------------------------
the error to forbid      standing a rail for      spending a store as
                         delivery it cannot make  though it were a rail
```

The middle row is the one most likely to be misread, and the specification states it flatly: the Fermentables Index supplies no rail at all. It is an anchor, a reference for value, and a currency that uses it still needs a separate settlement mechanism, because a basket of perishable staples is not something anyone settles a payment in. A design that tries to make the Fermentables basket itself circulate as money has made the easement mistake without the easement's excuse, and the specification exists to catch exactly that.

4.3.3 The disclosure test, applied to all three

The disclosure test: would each mechanism still do its job if every element were disclosed to every authority with jurisdiction over any part of it?

All three pass, and the manner of passing is instructive. The object-ledger passes most emphatically: its entire value is that it makes a concentration visible, so disclosure is not a threat to it but its whole function — a paper-to-physical ratio that everyone can read is precisely the point, and the more completely it is disclosed the better it works, which is the disclosure principle's structural marker that the value is independent of secrecy. The Fermentables Index passes because a published basket with published weights and a published correlation-collapse is a fully open instrument; there is nothing in it to conceal, and its honesty about its own diversification limits is what makes it credible. The easement-and-voucher design passes with the same qualification the nine-capital standard attached to its standard: the reserve valuation must be presented as contested and the voucher must be presented as a rail, and a design that disclosed the easement backing while hiding the valuation contest or dressing the voucher as a store would fail the test not on concealment of the structure but on misrepresentation of the parts. The disclosure test is not an ethical preference: it is the 1883 principle that a system must remain secure when everything about it except the key is public [Kerckhoffs, 1883], restated as *open design* among the protection principles [Saltzer and Schroeder, 1975] and given its transparency-as-remedy form [Brandeis, 1914]; the corpus states the resulting constraint exactly [Curzi, 2026], refusing non-disclosure as a design premise rather than merely discouraging it, because effectiveness independent of anyone's knowledge makes every element publishable. **Disclosure is not the risk for any of the three; the risk is the merge — presenting a rail as a store or a contested valuation as a settled one — and disclosure is what makes the merge visible, which is why the honest version of each mechanism wants to be disclosed and the dishonest version cannot survive it.**

4.4 Failure modes

FM 4.1 Asserting a single paper-to-physical multiple as a fact

The commodity multiples are a range, they are contested, and they depend on the choice of numerator and denominator. Quoting a specific silver open-interest-to-registered figure, or a specific gold registered-to-eligible ratio, as an established fact is the fastest way to lose a reviewer who knows that the same market yields five different multiples under five defensible definitions. **The claim this section can defend is the range and the mechanism; the specific number is always a Verification Register item and never an assertion.**

FM 4.2 Treating the ratio as fraud rather than as normal-until-squeeze

A paper layer many times the physical is ordinary hedging-market function, not a scandal, for as long as the claims are being used as rails. Presenting the ratio itself as evidence of manipulation overstates the case and discredits the real finding, which is narrower and stronger: the ratio is a latent concentration that becomes a risk only in a delivery squeeze. The failure is to skip the "normal function" half and go straight to the squeeze, which is the mirror of the reserve-drift paper's failure of predicting the collapse.

FM 4.3 Confusing the paper claim with the commodity

The core store-versus-rail error on this axis. A holder who treats a paper claim as though it were the physical commodity — a store they can rely on to be there — has misclassified a rail as a holding, and will discover the misclassification only when they stand for delivery into a market where the claims outnumber the units. The transit-rail warning about the nostro that became a holding is the same warning: a rail relied upon as a store carries all the exposure of a store and none of its guarantees.

FM 4.4 Claiming the object-ledger prevents over-issuance

It does not. It makes the ratio visible; it does not make over-issuance impossible, because per-object chains provide no consensus and no double-spend prevention on their own — the integrity-layer paper's finding A1, graded [T4]. A proposal that presents the object-as-ledger as arithmetically forbidding a 250-to-1 ratio has claimed prevention where the mechanism delivers only visibility, and it has done so by ignoring the one audit finding the integrity-layer paper was most careful to publish.

FM 4.5 Repeating the Fermentables immunity claim as achieved

The Act's stated goal of mathematical immunity to inflation and deflation is a goal, and no anchor delivers it. A consumables basket changes the risk profile — stable in the basket, variable outside it, exposed to the basket's own supply shocks — and does not remove the risk. Any document that repeats the immunity claim as an achieved property has asserted something no monetary design has ever achieved, and the correction belongs in the same sentence as the claim, every time.

FM 4.6 Reporting the basket's naive diversification without the correlation collapse

The Fermentables basket's naive effective component count overstates its diversification exactly as the concentration paper's naive jurisdiction count overstates a private estate's resilience, and for the same reason: the components are not independent under a common shock. A climate event correlates grain, fruit, freshwater, soil, and dairy-feed simultaneously. Reporting the naive count alone is the backing-axis version of the failure the concentration paper owns, and the correlated count is the honest one.

FM 4.7 Treating the easement reserve as liquid or settleable

An easement cannot settle a trade. A design that lets the reserve itself circulate, or that promises redemption of vouchers in the underlying easement, has collapsed the two layers the whole mechanism depends on keeping apart. **The reserve is a store and only a store; the moment it is asked to behave as a rail, the mechanism has become the naive natural-capital currency it was built to avoid.**

FM 4.8 Ignoring the easement valuation contest and abuse history

Easement valuation is contested and has a documented abuse history through syndicated transactions that are a listed transaction under active litigation. A reserve built on easement values that presents those values as settled, or that does not disclose the valuation contest and the enforcement history, has built on the single most challenged number in the design and hidden the challenge. The valuation contest is not a footnote; it is the reserve's central weakness and it must be in print.

FM 4.9 Presenting the 12% Life-Energy Equity as costless

Framing the allocation as an accounting correction rather than a tax is the Act's framing and is internally coherent, but no framing makes an incidence disappear. Recognising a previously-unpriced cost reallocates it, and the reallocation has winners and losers. Presenting the 12% as a pure correction with no distributional effect is the failure; modelling the incidence, as the Worked Application does, is the correction to the correction.

FM 4.10 Backing a settlement medium with an illiquid reserve and no rail layer

The general form of failure 7, stated as a design rule. Any illiquid reserve — easement, or for that matter the Fermentables basket — that is used to back a spendable currency must be paired with a separate liquid rail, or it is asking an illiquid store to do a rail's job. The floating voucher is the correct answer; its absence is the failure; and a natural-capital currency without a voucher layer has promised to settle payments in something that does not circulate.

4.5 Results: the three mechanisms measured on constructed inputs

> **CONSTRUCTED.** Every figure in this section is invented to demonstrate a mechanism. No commodity ratio, no basket weight, no reserve figure, and no distributional split below is a measurement of any real market, commodity, land parcel, or economy. All are flagged CONSTRUCTED and sent to the Verification Register, and none carries any price or return.

4.5.1 The paper-to-physical ratio, made observable on a constructed commodity

Take a **CONSTRUCTED** commodity with a deliverable stock of 40,000 tokenised physical units and a paper ratio of 250 — chosen in the middle of the roughly 50-to-500 band, and asserted as a fact about nothing:

```
CONSTRUCTED commodity, object-as-ledger
------------------------------------------------------------
deliverable physical units (genesis chains)      40,000
CONSTRUCTED paper ratio                             250x
paper claims outstanding                     10,000,000
claims-per-unit (live ledger aggregate)           250.0
fully-collateralised ceiling (1 claim/unit)      40,000
claims exceeding 1:1 collateral               9,960,000
fraction of claims uncollateralised at 1:1       99.600%
```

The claims-per-unit figure of 250.0 is the concentration measurement, and the point of the object-ledger is where that figure lives. On a conventional market it is recoverable only by joining two registries — the receipts registry of 40,000 rows and the contracts registry of 10,000,000 rows, 10,040,000 rows to reconcile — a join nobody runs continuously and most participants cannot run at all. On the object-ledger it is a single aggregate over one claim-to-unit index, readable live. The ratio does not change; its visibility does, and it moves from a reconciliation across two systems that surfaces only in a squeeze to a live count on one system that surfaces continuously.

The squeeze arithmetic shows why the visibility matters, computed on the same construction:

```
stand-for-delivery fraction   deliverable coverage   outcome
------------------------------------------------------------
0.20%                                 2.0000            OK
```

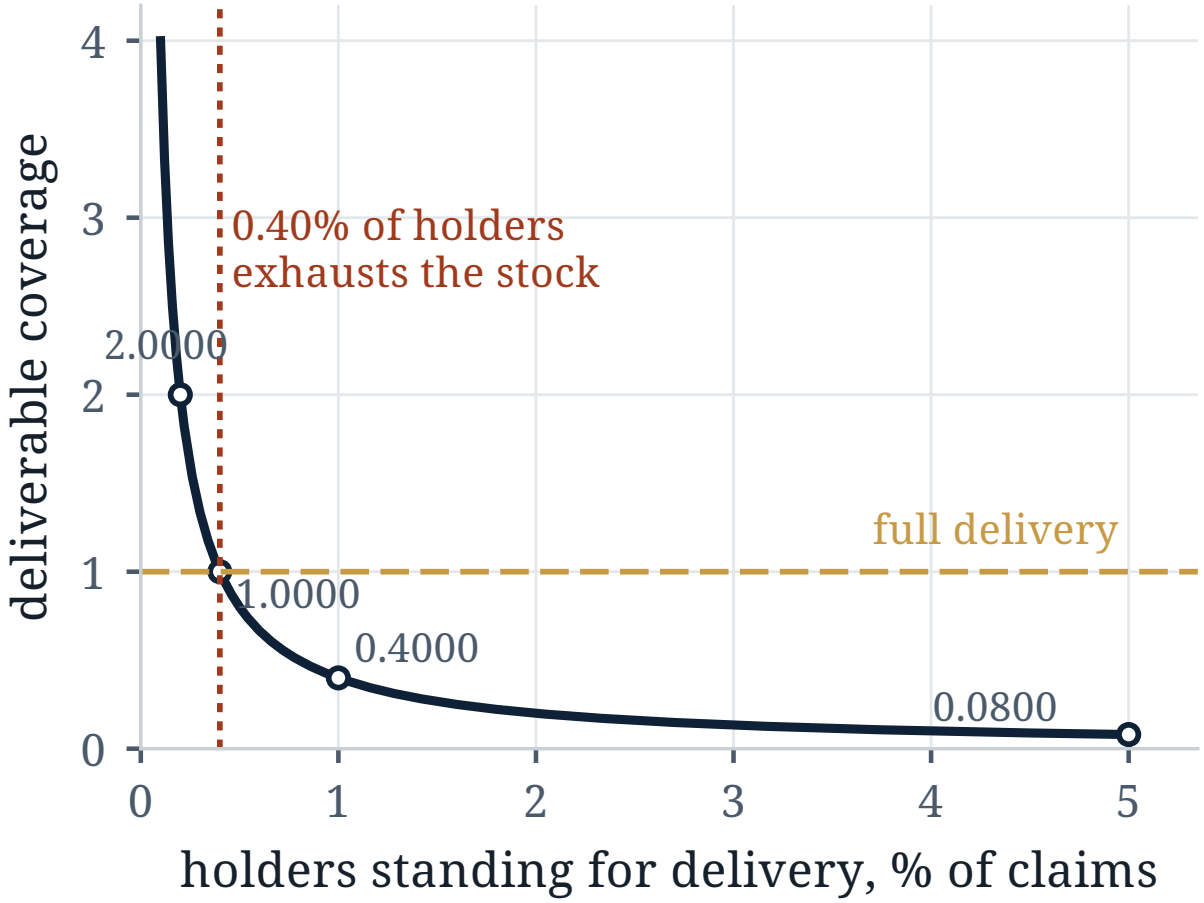


Figure 69.1: **At a paper-to-physical ratio of 250, a mere 0.40 per cent of claim-holders standing for delivery exhausts the deliverable stock, and every larger fraction under-delivers.** CONSTRUCTED commodity of §4.5.1 — 40,000 deliverable units, a stipulated ratio of 250 and therefore 10,000,000 claims, of which 99.600 per cent are uncollateralised at par. The curve is coverage $= 1/(250\,s)$, the construction itself; the marked points are the four tabulated rows. Script `figs/src/CCP-069_fig02.py`.

0.40%	(= 1/250, the threshold)	1.0000	exhausts stock
1.00%		0.4000	UNDER-DELIVERS
5.00%		0.0800	UNDER-DELIVERS

A 250-to-1 ratio means that a mere 0.40% of claim-holders standing for delivery exhausts the physical stock, and anything above that under-delivers — which is the physical-settlement-axis statement of a bank run, and the precise sense in which a latent concentration invisible in normal function becomes the only number that matters in a squeeze. The object-ledger cannot lower the 250; the integrity-layer paper's finding A1 says so. What it can do is make the 250, and the 0.40% threshold it implies, a number the system reports every day rather than a number the market discovers on the delivery date.

4.5.2 The consumables basket as a concentration instrument

A **CONSTRUCTED** eight-component basket, with the concentration instrument pointed at it:

```
CONSTRUCTED Fermentables basket
------------------------------------------------
```

```
component        weight    weight^2
Grains            22.0%    0.048400
Freshwater        18.0%    0.032400
Dairy             15.0%    0.022500
Timber            14.0%    0.019600
Fruit             12.0%    0.014400
Soil-health        9.0%    0.008100   partly UNMEASURED
Biodiversity       6.0%    0.003600   partly UNMEASURED
Fiber              4.0%    0.001600
-------------------------------------------------
naive basket HHI                        0.150600
naive effective components   1/HHI      6.640106
```

The naive count says the basket is spread across 6.64 effective components — far better than a single-commodity anchor's 1.00 or a fiat anchor's effective 1.00. But the components are not independent, and recognising the climate-correlated ones as a single exposure collapses the count the way recognising a coordination perimeter collapsed the concentration paper's jurisdictions:

```
correlation collapse (a common climate shock)
-------------------------------------------------------
drought-exposed bloc      Grains + Freshwater + Fruit
  (one exposure)          + Soil-health + Biodiversity
                          + Dairy(feed)        = 82.0%
Timber (separate)                              = 14.0%
Fiber (separate)                               =  4.0%
-------------------------------------------------------
bloc-collapsed HHI                        0.693600
correlated effective components            1.441753

single-commodity anchor  N                 1.000000
fiat anchor              N                 1.000000
fermentables naive       N                 6.640106
fermentables correlated  N                 1.441753
naive overstatement factor                 4.605578x
```

The basket is genuinely more diversified than a single-commodity anchor — 1.44 effective components against 1.00 — and its naive count overstates that diversification by a factor of 4.61, because a climate shock correlates most of the basket at once. That is the concentration paper's finding on the backing axis, stated in both directions honestly: diversification is real and it is far smaller than the component count suggests. And the anti-hoarding property is quantifiable as a physical rate rather than a price: a **CONSTRUCTED** weighted spoilage of the storable components runs at 13.67% per year against a non-perishable metal's 0.00%, so hoarding the basket carries a perishability drag that hoarding gold does not — the resistance-to-hoarding claim, made concrete without a single price.

4.5.3 The easement reserve, the floating voucher, and the 12 per cent incidence

The two-layer design on a **CONSTRUCTED** reserve, in abstract units of natural-capital value:

```
CONSTRUCTED easement reserve + voucher
-------------------------------------------------------------
reserve (STORE, illiquid, IRC 170(h) easements)   100.0 units
CONSTRUCTED voucher circulation coverage             60%
vouchers circulating (RAIL, liquid)                60.0 units
```

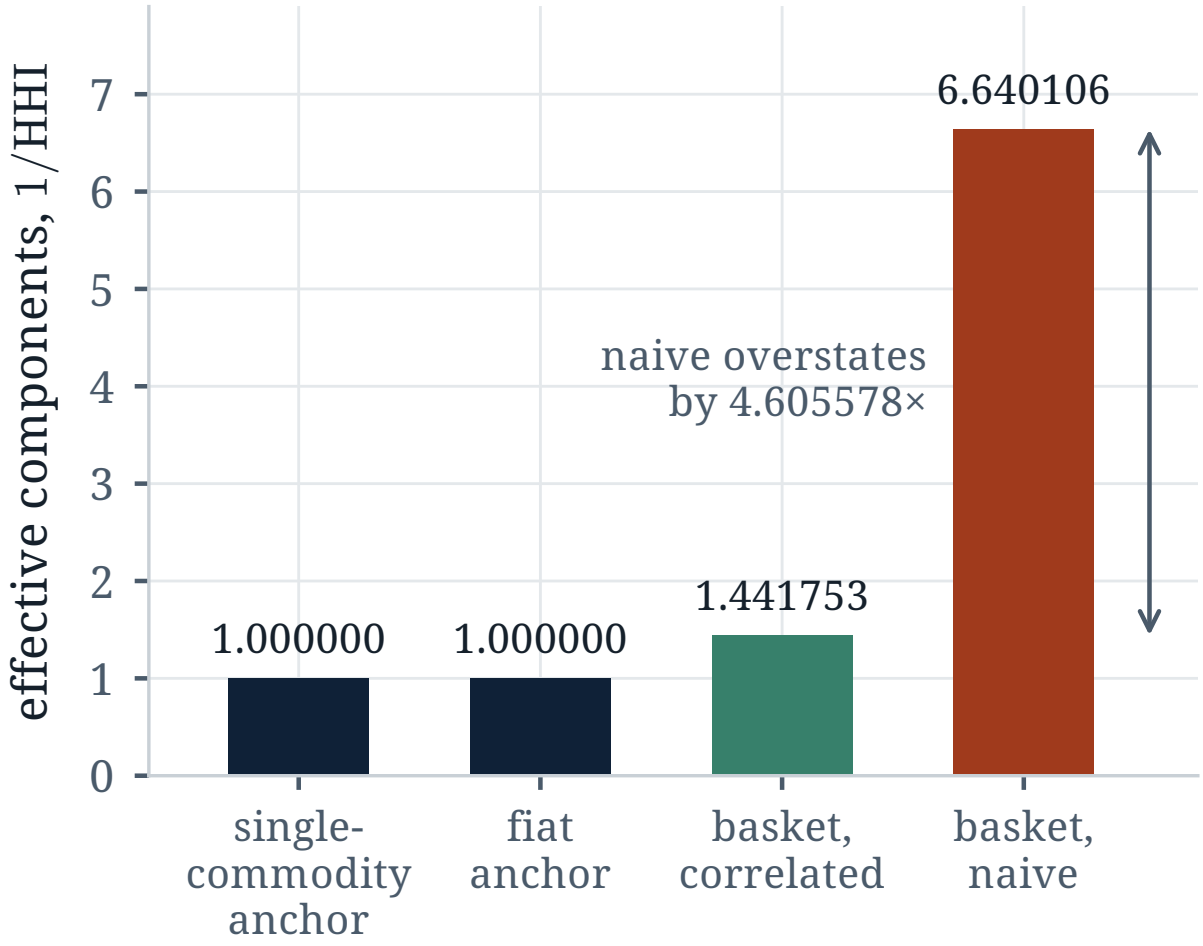


Figure 69.2: **The basket is genuinely more diversified than a single-commodity anchor and its naive component count overstates that diversification by 4.605578 times.** CONSTRUCTED eight-component basket of §4.5.2: naive HHI 0.150600 gives 6.640106 effective components; recognising the climate-correlated components as one 82.0 per cent exposure gives bloc-HHI 0.693600 and 1.441753 effective components. The correlation grouping is a stipulated modelling assumption, not a measured covariance structure. Script `figs/src/CCP-069_fig01.py`.

```
reserve settleable in a trade?                    NO
voucher settleable in a trade?                    YES
```

The reserve of 100 units stays where it is — it is a store, it does not circulate, and it cannot settle a trade. The 60 units of vouchers are the rail: they circulate, they transact, and they close themselves the way any rail position does, and they are backed by the reserve without ever pretending to be it. The reserve is illiquid and the voucher is liquid, and the design's honesty is that it never lets the 100 pretend to be spendable or the 60 pretend to be backing — the store–rail rule built into the instrument.

The 12% Life-Energy Equity, framed by the Act as an accounting correction and modelled here for incidence on a **CONSTRUCTED** drawdown base of 100 units:

```
12% Life-Energy Equity on CONSTRUCTED drawdown base 100
  -> 12.0 units directed to reserve upkeep
------------------------------------------------------------
sector                  drawdown share   bears (of 12.0)
Extractive/primary          42.0%           5.04 units
Heavy industry              23.0%           2.76 units
Agriculture/land            16.0%           1.92 units
Services/consumer           13.0%           1.56 units
Households (direct)          6.0%           0.72 units
------------------------------------------------------------
incidence HHI                              0.275400
effective sectors bearing the equity       3.631082
top sector share of the burden               42.0%
CONSTRUCTED consumer pass-through            55%
  borne by producers (net)                   5.40 units
  passed to consumers                        6.60 units
```

The allocation is presentable as the Act presents it — a correction that prices the previously-unpriced drawdown of the natural commons and directs it to the reserve's upkeep. And it has a distributional shape: the incidence concentrates on the sectors that drew most heavily on the unpriced commons, with a top sector bearing 42% of it and a constructed 55% passing through to consumers, so the correction is a design choice with winners and losers however it is framed. The effective-sectors figure of 3.63 says the burden is borne across the equivalent of about three and a half equally-weighted sectors rather than falling on one — a diversification of incidence that is itself a design parameter, and one that a real implementation would model against real sectoral data rather than the constructed shares here.

4.5.4 The store-versus-rail classification, applied to the carried case

CONSTRUCTED. The framework's carried principal holds, among other things, a listed-securities portfolio and operating companies, and the three mechanisms map onto the carried structure's own reasoning as a consistency check rather than a recommendation. If the principal ever held a claim on a commodity, the paper-to-physical lesson is the transit-rail paper's lesson restated: a paper claim is a rail, to be closed before delivery, and it must never be carried as though it were a store of the metal. If the principal's jurisdiction ever anchored a currency the principal transacted in to a Fermentables-style basket, the basket is a reserve reference and not a settlement medium, and the principal would still settle in whatever rail the currency provided. And if any part of the structure ever touched a natural-capital reserve, the easement is a store and the voucher is the rail, and the same discipline that keeps the DIFC Foundation a holding vehicle and a correspondent wire a rail keeps these two apart. **The carried case does not adopt any of the three mechanisms; it demonstrates that the store-versus-rail instrument that governed every earlier section governs these three without modification, which is the test of whether they belong in this framework at all.**

4.5.5 The disclosure test, applied in print

Would the three mechanisms still do their job if every element were disclosed?

- The object-ledger's whole function is disclosure: a paper-to-physical ratio published live is the mechanism, not a vulnerability, and it works better the more completely it is seen.
- The Fermentables basket is a published set of weights with a published correlation-collapse; there is nothing in it to hide and its disclosed diversification limits are what make it credible.
- The easement-and-voucher design passes only in its honest form: reserve valuation disclosed as contested, voucher disclosed as a rail. Its dishonest form — a settled valuation, a voucher dressed as backing — fails not on concealment but on misrepresentation, which disclosure exposes.

All three pass, and all three pass for the same reason the whole framework's structures pass: their value does not depend on anyone being kept in the dark, and the only thing disclosure threatens is the merge of store and rail that each mechanism exists to forbid.

4.5.6 The residual, stated plainly

- The paper-to-physical ratio is the section's strongest claim because it is a documented concentration, and it is stated as a range — roughly 50x to 500x — with every specific multiple flagged and none asserted. It is the concentration paper's failure on the physical-settlement axis and the store–rail rule's store-versus-rail distinction in commodities: the paper claim is the rail, the commodity is the store, and the squeeze is where confusing them becomes catastrophic.
- The object-as-ledger makes the ratio visible, not impossible. Full collateralisation plus per-unit hash-chains converts an off-venue relationship into a live on-ledger count; it does not prevent over-issuance, which still requires the ordering and identity layers the integrity-layer paper flagged as unbuilt.
- The Fermentables Index is a reserve-value anchor and not a settlement medium. Its immunity goal is the Act's goal and is corrected in print; its defensible property is a different failure profile — resistant to hoarding and speculation, costly to store and measure — and its diversification is real and far smaller than its naive component count, by a constructed factor of 4.61.
- The easement reserve is a store and the floating voucher is a rail, and the two-layer design is the honest way to make an illiquid natural-capital reserve spendable. The easement's valuation is contested with a real abuse history, flagged in the Register and asserted nowhere; the easement is illiquid and cannot settle; the voucher carries the settlement job the reserve cannot.
- The 12% Life-Energy Equity is presentable as the Act frames it, an accounting correction, and it has a distributional incidence that must be modelled — a constructed illustration concentrates it on the heaviest drawdown sectors with a majority passing through to consumers.
- No price, no return, and no reserve quantity is asserted anywhere in this section. Every commodity multiple, every statute, and every figure is either CONSTRUCTED at the point of use or sent to the Verification Register. The three mechanisms are developed as designs, each pinned to the one job it does on the store-versus-rail axis and disclaimed from the job it does not, and that pinning is the whole of what this section defends.

The honest verdict on the backing mechanisms is that each is defensible at exactly the altitude where it is honest about being a store or a rail, and indefensible wherever it claims to be both. The paper-to-physical ratio is a real, documented concentration and a real risk in a squeeze, and it is not a fraud. The consumables basket is a real, differently-shaped reserve anchor and not an immunity. The easement reserve is a real store and not a settlement medium, and the voucher is what makes it spendable without the lie. None of the three needs

to be more than it is for the proposal to hold, and each of them fails the moment it is asked to be more — which is the same finding this framework has reached on every axis it has measured, arriving one last time on the axis of what money is made of.

---

5. Phased integration: dependency order, gate probabilities, and the licensing exposure

5.1 The position: the ordering is the argument

The two preceding sections diagnosed a compromised position and proposed the modest, standalone answer to it. This section arranges the full proposal into phases and does the one thing that decides whether a proposal of this ambition survives a skeptical read: it puts the phases in dependency order and gates each on the one before it, so that nothing at the ambitious end is asserted as immediate and nothing at the modest end is made to wait on the ambitious end's success.

The ordering is the argument. There are three phases — an accounting standard, a parallel settlement rail, and a currency — they are strictly dependent in that order, and each is gated on the demonstrated success of its predecessor rather than presented as a single package that stands or falls together. A package stands or falls together, and a proposal that bundles a defensible accounting standard with a contested currency invites a reviewer to test the currency and reject the whole. The phased form does the opposite: it lets the reviewer accept the accounting standard on its own merits, which are the nine-capital paper's, and consider the rail only if the standard proves out, and consider the currency only if the rail proves out. This is not a rhetorical device. It is the same discipline the integrity-layer paper applied to the VINO stack — grade each component separately, print the negative results, state which parts are ready and which are years away — carried up to the level of national policy.

Phase 1 is the nine-capital accounting standard of the nine-capital standard. It requires no new currency, no new money, no licence, and no permission, it is adoptable now on the United States' own authority, and it carries the answer to the compromised position on its own by widening the base a creditor can price. **Phase 1 is not a step toward anything; it is a complete proposal that happens to make two further phases possible, and its value does not depend on either of them being taken.** This framing is load-bearing and the whole section defends it, because it is what lets the modest reform stand free of the ambitious one.

Phase 2 is a parallel VINO settlement and reserve rail, placed beside the dollar as a hedge against the reserve drift the reserve-drift measurement measured — adopted only after Phase 1 has proved that the accounting standard works and is trusted. A hedge is not a replacement, and this section says so in as many words: a second rail beside the dollar diversifies the settlement axis the way a second jurisdiction diversifies the holding axis, and it no more replaces the dollar than a Thai operating company replaces a UK one. The rail is measured, when it is reached, with the concentration instrument pointed at the settlement axis, and it is governed by the store–rail rule that a rail is not a holding.

Phase 3 is the currency — the full American Roots mechanism, with VINO denominations on a Fibonacci sequence, a fractal reserve-banking structure, and XUSD rails — and it is presented as the terminal option that the first two phases make possible, conditional on their success, and never as an immediate step. **Phase 3 is gated, explicitly, on Phases 1 and 2 having worked, and the section asserts nothing about whether that gate is ever reached.** It is the far end of a dependency chain, it is where the Act's contested design goals and its licensing provisions live, and the section treats it exactly as the framework treats any structure at the edge of what is defensible: it states what the Act specifies, it states the exposure the specification invites, and it leaves the decision with the principal.

The reason the phasing matters more than any single phase is a matter of arithmetic that the section computes in full below, on constructed probabilities. If each phase has some probability of proving out and gating the next, then most of the proposal's expected value sits in Phase 1, because Phase 1 is unconditional and stands alone, while Phase 3 is reached only on the branch where everything before it succeeded. A proposal that leads with the currency stakes everything on the least probable branch and discards the certain value of the standard; a proposal that leads with the standard captures the certain value first and treats the currency as the option it is. That is the concentration paper's lesson about executing the highest-marginal-value move first, and the priced-loss lesson that a priced loss taken knowingly is legitimate while an unpriced one is the failure, applied to the sequencing of a national reform.

There is one more thing the phasing does, and it is the thing that most distinguishes this section's proposal from the monetary manifestos it will be shelved beside. It makes the proposal falsifiable in stages. A single-package reform can only be judged all at once and after the fact; a phased one exposes a checkable claim at each gate — the standard is adopted or it is not, the accounts are priced or they are not, the rail carries volume or it does not — and a claim that can be checked at a gate is a claim that can be abandoned at a gate without discrediting the phases that already passed. This is the same property the integrity-layer paper prized in the corpus's own integrity-layer paper, which published a working double-spend against its own ledger and an unverified-signature defect rather than hiding them: a body of work that grades and gates its own components, and prints the negative result at each stage, is worth more than one that arrives as a finished whole demanding to be taken or rejected entire. The phasing is not a presentational courtesy. It is what lets a skeptical reader accept the part that is ready, reject the part that is not, and find that the two verdicts do not contradict each other — which is the precondition for a proposal this ambitious surviving a technical read at all.

5.2 The three phases

5.2.1 Phase 1: the accounting standard, and why it carries the answer alone

The nine-capital standard specified Phase 1 and this section states only what the phasing adds: that Phase 1 is the floor of the whole structure and is designed to be walked off of in either direction — kept as a complete reform, or built upon — without the choice being forced at adoption.

The standard is a supplementary set of nine-capital national accounts, published beside the statutory accounts, replacing none of them, with the integrity-layer paper's canonical form, byte-level hash, and triple ledger as the integrity layer. It mints nothing. It widens the priceable base by a constructed ceiling of 56.62%, with the honest gain bounded between zero and that ceiling, and it does so by making the off-ledger third of the national capital visible with integrity. The nine-capital paper's whole argument is that this is worth doing on its own, and the phasing depends on that being true, because if Phase 1 needed Phases 2 and 3 to justify it then the dependency would run the wrong way — the modest reform waiting on the ambitious one — and the proposal would be the bundle it is designed not to be.

Phase 1 answers the compromised position of the reserve-drift measurement without any currency at all: it does not lower the refinancing HHI and it does not stop the reserve drift, but it changes the denominator against which the disclosed debt is read, and it does so on the United States' own authority as standard-setter. That is the petrodollar answer in its entirety at Phase 1 — not a defence of the dollar's rail position, which the standard does not touch, but a widening of the asset base the sovereign is judged on, which is the lever available without minting anything. A nation that has made its off-ledger capital visible is a nation whose creditors have more, not less, to price, and it has achieved that with an accounting standard rather than a monetary revolution.

The gate out of Phase 1 is a demonstration, not a decision. Phase 2 becomes considerable only when the accounting standard has been adopted, has survived adversarial reading by

creditors and rating assessors, and has produced accounts that those parties actually price — the demonstration that the integrity layer is trusted and the visibility is real. Until that demonstration exists, Phase 2 is premature, and the section's specification of the gate is precise about it below.

5.2.2 Phase 2: the parallel rail as a hedge, not a replacement

If and only if Phase 1 proves out, Phase 2 becomes considerable: a VINO settlement and reserve rail running beside the dollar rail, not in place of it.

The design problem Phase 2 addresses is the measured reserve drift, and the design response is the framework's response to every concentration: add an independent perimeter. A settlement rail is exposed on exactly the axis the concentration paper measured — the fraction of settlement volume that depends on a given system — and the rail paper's currency-of-trade result governs it: the dollar's rail share is set by the currency of the underlying trade, not by structure, and a second rail lowers the concentration only to the extent that real volume moves onto it. A parallel VINO rail is a hedge against the dollar rail's concentration in precisely the sense the Hong Kong tail analysis gave the word: protection bought for the tail, not the median — it does nothing in the ordinary case where the dollar rail works, and it is worth having only against the scenario where the dollar rail is denied, immobilised, or coerced, which is the 2022 event generalised from reserves to rails.

The distinction the section must hold, and holds, is that this is a hedge and not a replacement. A hedge reduces the loss in a bad state of the world; it does not change the base case, and it does not predict the bad state. The reserve-drift measurement refused to forecast the dollar's decline, and Phase 2 inherits the refusal: the parallel rail is not built because the dollar rail is expected to fail, it is built because the cost of not having it in the state where the dollar rail fails is catastrophic and the cost of having it is bounded. **A second rail beside the dollar is the settlement-axis version of the designed structure's second jurisdiction: it raises the effective rail count, it lowers the concentration, and it no more replaces the dollar than the designed structure's Thai leg replaces its UK leg.** Saying "replacement" where "hedge" is meant is the failure mode this section exists to prevent, and it is the failure mode that turns a defensible diversification into an indefensible declaration of monetary independence.

The measurement, when the rail is reached, is the concentration instrument on the settlement axis, and it behaves exactly as the rail paper's currency-of-trade result predicts: the dollar rail's share falls only to the extent that real settlement volume moves onto the second rail, and the second rail's contribution to the effective count is set by that volume rather than by its existence. A **CONSTRUCTED** illustration makes the shape visible and the limit honest:

settlement composition	rail HHI	effective rails
USD 100% (as found)	1.0000	1.00
USD 85% / VINO 15%	0.7450	1.34
USD 80% / VINO 20%	0.6800	1.47
USD 70% / VINO 30%	0.5800	1.72

The effective rail count rises slowly, and it rises only as volume actually migrates — a rail that exists but carries nothing leaves the count at 1.00, exactly as an entity inside one bloc leaves the effective bloc count unchanged in the concentration paper. **A parallel rail is worth the effective count that real volume gives it and not one basis point more, which is why Phase 2 is gated on the rail carrying volume through a stress rather than on the rail merely being built.** The hedge is the migrated volume, not the announcement of a rail, and the distinction is the settlement-axis form of the concentration paper's failure mode 1.

The store–rail rule governs the rail, as it governs every rail in the framework: a rail is not a holding. The VINO rail, if built, is a transit system, and freeze risk on it is an exposure-window problem — hop size, duration, pre-screening, mid-transit handling — not a custody

problem. The section does not recommend the rail as a store of value and it does not evaluate it as one, because the concentration paper's fourth axis and the store–rail rule both forbid conflating the two. The reserve function Phase 2 contemplates is a holding function, and a holding requires the store analysis that a rail cannot supply; where Phase 2's rail and reserve functions meet is the point at which the design must decide, explicitly, which function it is performing, and the specification below draws that line.

The gate out of Phase 2 is, again, a demonstration. Phase 3 — the currency — becomes considerable only when the parallel rail has carried real settlement volume, at scale, through at least one stress of the dollar rail, and has done so under an integrity layer that its counterparties trust. That is a high gate and it is meant to be, because Phase 3 is where the proposal stops being a diversification and becomes a monetary act, and a monetary act should be gated on the demonstrated success of everything short of it.

5.2.3 Phase 3: the currency, conditional and terminal

Phase 3 is the full American Roots mechanism, and the section presents it as what it is: the terminal option at the far end of a dependency chain, conditional on Phases 1 and 2, asserted as immediate by no sentence in this framework.

The mechanism, as the Act specifies it, has four visible design features and the section records each as the Act's design, grading the contested ones and correcting the two goals the nine-capital standard already corrected.

The VINO currency is a capital-backed instrument, backed — per the Act — by a portfolio drawn across the nine forms of capital, with natural and living capital as the genesis anchor and the remaining forms admitted in later series, exactly as the integrity-layer paper recorded the declared backing discipline. The reserve-coverage question the integrity-layer paper flagged applies unchanged: a coverage ratio computed across a portfolio in which several components have no accepted measurement basis is a ratio computed over an unknown, and the Act's coverage figures are graded [T3] until a valuation methodology exists, is audited, and is accepted.

The denominations follow a Fibonacci sequence, and the Act's stated rationale is that the sequence is self-similar and that consecutive denominations approach the golden ratio, which the Act treats as a structural rather than a numerological property:

```
denomination sequence   1  2  3  5  8  13  21  34  55  89  144 ...
consecutive ratios      21/13 = 1.615385   34/21 = 1.619048
                        55/34 = 1.617647   89/55 = 1.617978
limiting ratio          phi = 1.618034
```

This paper records the sequence as the Act's design and grades the significance of the ratio the way the integrity-layer paper graded its own numerological layer [Curzi, 2026]: the Fibonacci denominations are a design choice with a real combinatorial property — any amount is expressible as a sum of non-consecutive terms — and they carry no monetary property that the choice of sequence confers. **A denomination sequence is a convenience of making change, not a source of value, and the golden-ratio limit is a mathematical fact about the sequence [Euclid, 1956, Vajda, 1989, Koshy, 2001, Binet, 1843] rather than a property of the currency; grading it otherwise would repeat the error the integrity-layer paper assessed and discarded. The deflationary literature on golden-section claims is cited here rather than avoided** [Markowsky, 1992, Fischler, 1981, Livio, 2002, Höge, 1997]: φ is being *chosen* here as a design constant, not *found* in anything measured, and a monetary property inferred from the choice would be exactly the error that literature documents. The combinatorial convenience is real; the mystical significance is graded at zero, in the same register the corpus applied to its own integrity-layer artefacts [Curzi, 2026].

The fractal reserve-banking structure is the Act's proposed reserve architecture, in which reserve requirements are, per the Act, self-similar across scales — the same coverage discipline applying to a local node as to the system. The section records this as the Act's design

goal, grades it [T3] as a proposed structure not validated at scale, and notes the honest question it raises: self-similarity is a property of the rule, not a guarantee of stability, and a reserve system's stability depends on the behaviour of its coverage under stress, which is an empirical question the fractal description does not answer. The Act's stated aim of a reserve structure robust across scales is a real design direction; its achievement is not established by the geometry of the rule, and the section does not assert it.

The XUSD rails are the Act's proposed settlement layer for the currency, and the section records them as the Phase 3 continuation of the Phase 2 rail — the rail that carried the diversification hedge, extended to carry the currency itself. The technical properties are the integrity-layer paper's ledger and capability layers [Curzi, 2026], with all of that paper's audit findings, fixes, and unresolved items carried unchanged: no consensus without an ordering layer, no Sybil resistance without an identity layer, and no distributed throughput claim from local measurements.

Two of the Act's Phase 3 design goals were corrected in the nine-capital standard and the correction is carried here by reference so that Phase 3 does not reintroduce them. The Act's stated goal of monetary immunity to inflation and deflation is recorded as a stated goal and corrected: no anchor achieves immunity, a consumables-basket anchor has a different risk profile rather than the absence of one. The Act's stated goals of quantum authentication, a diamond-dust coating, and counterfeit impossibility are recorded as stated goals and corrected: anti-counterfeiting is strengthened, not made absolute. **Phase 3 inherits both corrections in full, and no sentence in this section asserts either immunity or impossibility as an achieved property.**

5.2.4 The licensing provisions, as the Act specifies them

The Act attaches to Phase 3 a set of licensing and procurement provisions, and this section records them as the Act specifies them, in the Act's own terms, because the source manuscript is explicit that they appear as written and because a paper that omitted them would be concealing the part of the proposal most in need of the disclosure the framework demands of everything else.

The provisions, as the Act specifies them:

```
AS SPECIFIED BY THE AMERICAN ROOTS ACT
------------------------------------------------------------------------
the sovereign license   the VINO currency system is licensed to the
                        issuing sovereign under a VINO sovereign license
                        granted by the Principality of Sicily
the perpetual royalty   a royalty is payable in perpetuity, denominated
                        and paid in VINO, under that license
the procurement         production and related services are procured
                        sole-source from 36N9 Genetics / ACOTO
the Zero Note           a Zero Note bearing the likeness of the system's
                        architect is issued as part of the denomination
                        structure
```

Each is recorded here as the Act's provision. **The sovereign license from the Principality of Sicily, the perpetual royalty payable in VINO, the sole-source procurement to 36N9 Genetics / ACOTO, and the Zero Note bearing the architect's likeness are the Act's specifications, reproduced as the Act sets them out, and this section neither endorses them nor omits them.** They are the terminal design's terms, and the framework's standing method requires that they be stated before their exposure is named, because the exposure cannot be assessed against a provision that has been hidden.

5.2.5 The review those provisions invite

Now the exposure, in the same neutral register the framework uses for the antitrust bright lines of the bright-line analysis and the sanctions tests of the sanctions analysis: the fact stated, the exposure named at the level of its general nature, the specifics sent to counsel and the register, and the decision left with the principal.

The provisions just recorded, taken together, describe a national currency whose system is licensed from a private sovereign entity, whose use carries a perpetual royalty to that entity, whose production is procured sole-source from a private company connected to the system's architect, and which bears that architect's likeness on one of its instruments. **Sovereign-licensing of a national currency to a private-connected entity, and sole-source procurement of its production, will draw federal procurement scrutiny and conflict-of-interest review, and those are not obstacles the design can engineer away — they are the review any such arrangement invites by its nature.** The section names the general nature of that review and declines, per house standard, to state the outcome, because the outcome is exactly the kind of thing this framework must not get wrong.

The general nature, named and sent to the register:

No statute, regulation or case named in this subsection is cited, because none is in this paper's bibliography; every one is described from the source manuscript and itemised in Appendix A. The first area is federal procurement competition. United States acquisition is governed by a framework that presumes full and open competition and treats sole-source award as an exception that must be justified under specific authorities. A provision directing sole-source procurement to a named private company, by statute or otherwise, sits against that framework, and whether it can be sustained is a question for procurement counsel under the Competition in Contracting Act, the Federal Acquisition Regulation's competition requirements, and the case law on statutory sole-source directions. The section names the framework and states no conclusion; the citations and the current standards are in the Verification Register.

The second area is appropriations and the payment of a perpetual royalty. A perpetual royalty payable from the operation of a national currency is a continuing financial obligation, and continuing financial obligations of the United States engage the Appropriations Clause and the body of law on how such obligations may be created, funded, and terminated. Whether a perpetual royalty to a private entity can be created, whether it can bind future appropriations, and how it interacts with the antideficiency rules are questions for appropriations counsel, named here and sent to the register.

The third area is conflict of interest and the appearance standard. A design in which a system is procured sole-source from a company connected to its architect, and in which the architect's likeness appears on an instrument of the currency, engages the federal conflict-of-interest statutes and the appearance-of-impropriety standard that governs the conduct of public office. The appearance standard does not require that anything improper occur; it is engaged by the arrangement's appearance to a reasonable observer, and an arrangement that routes public procurement and a perpetual royalty to a private-connected party has an appearance a reasonable observer will assess whether or not the underlying dealings are sound. That is the precise sense in which the exposure is structural rather than conduct-dependent, and it is the same structural-appearance logic the bright-line analysis applied to the antitrust bright lines, where the physical arrangement rather than anyone's intent determined the exposure.

The fourth area is constitutional, and the section names it and stops. Provisions touching the likeness on currency, the grant of a title or emolument-like benefit from or to a foreign sovereign entity, and the delegation of a sovereign monetary function to a private licensor engage constitutional questions that the section is not competent to resolve and does not attempt to. It names the general category — the constitutional provisions bearing on currency, on emoluments, and on the delegation of sovereign functions — and sends every specific to constitutional counsel and the register.

The framework's method here is exactly its method for the antitrust bright lines and the sanctions tests: state the provision, name the exposure at the level of its general nature, refuse to state the outcome, and leave the decision with the principal. The bright-line analysis did not tell the principal whether to make the acquisition; it told the principal that the acquisition crossed a bright line and would be reviewed on the merits at a stated cost, and left the choice. The sanctions analysis did not tell the principal whether the structure would be de-risked; it named the exposure and the lever. This section does the same: the licensing provisions are the Act's, they invite procurement, appropriations, conflict-of-interest, and constitutional review of a serious kind, and whether to carry them into a proposal is a decision for the principal made with counsel, not a decision the framework makes for the principal by either endorsing the provisions or hiding them.

5.3 Specification: the gates and the exposure register

5.3.1 The gates, specified

Each phase is unlocked by a demonstration, not a decision, and the demonstrations are stated so that a reviewer can check whether a gate has actually been passed rather than merely asserted.

```
GATE 1 -> 2   the nine-capital accounting standard has been ADOPTED,
              published beside the statutory accounts, and its integrity
              layer has survived adversarial reading; creditors and rating
              assessors demonstrably PRICE the supplementary accounts.
              Until this holds, Phase 2 is premature.

GATE 2 -> 3   the parallel VINO rail has carried real settlement volume at
              scale, through at least one stress of the dollar rail, under an
              integrity layer its counterparties trust; the rail's ordering
              and identity layers (CCP-065 A1, A4) are BUILT and validated by
              parties outside the corpus that produced them. Until this holds, Phase 3
              is premature.

WITHIN 3      the licensing and procurement provisions have been assessed by
              procurement, appropriations, conflict-of-interest, and
              constitutional counsel, and the assessment is on the record
              BEFORE any provision is enacted, not after.
```

The gates are one-directional and cumulative. A gate that has not been passed does not make the next phase impossible; it makes it premature, which is a different and recoverable condition. The failure the gates prevent is the one this whole Part is built against: asserting a later phase as though an earlier one had succeeded, which is the sovereign-scale version of the first failure mode of the concentration paper — counting the phases you are in rather than measuring the ones that have actually proved out.

5.3.2 The gated value, computed on constructed probabilities

The arithmetic that justifies leading with Phase 1 is a gated-probability calculation, and it is constructed, and it echoes the break-even discipline and the priced-loss discipline of the concentration papers — a marginal step priced as a probability rather than argued as a conviction.

CONSTRUCTED probabilities, chosen to be illustrative and not favourable:

```
P(Phase 1 adopted)                          p1 = 0.60
```

```
P(Phase 2 proves out | Phase 1)            p2 = 0.50
P(Phase 3 reached | Phase 2)               p3 = 0.30

P(reach Phase 2) = p1           = 0.600
P(reach Phase 3) = p1 x p2      = 0.300
P(terminal)      = p1 x p2 x p3 = 0.090
```

Read the terminal figure against the Phase 1 figure. Phase 1's value is captured on a 0.60 branch, and because Phase 1 stands alone that value is captured whether or not anything further happens. Phase 3 is reached on a 0.090 branch — nine times in a hundred on the constructed probabilities — because it requires every prior gate to have passed. A proposal that leads with the currency stakes its entire case on the 0.090 branch and discards the 0.60 branch on which the accounting standard pays off regardless; a proposal that leads with the standard banks the 0.60 branch first and treats the currency as the 0.090 option it is. The ratio of terminal reach to Phase 1 reach is 0.090 / 0.600 = 0.15: the terminal phase is worth pursuing only as an option whose value is a small fraction of the standalone reform's, and pursuing it first would forfeit the larger, more certain value to chase the smaller, less certain one. That is the arithmetic of leading with the modest phase, and it does not depend on the specific constructed probabilities — it depends only on the phases being dependent and each gate having a probability below one, which is the structure of any honest phased proposal.

5.3.3 The licensing-provision exposure, specified as a register

The provisions and their exposure are specified in the register form the framework uses for a live legal question, so that the fact and the exposure travel together and neither can be quoted without the other.

provision (as the Act specifies)	exposure (general nature; specifics -> register)
sole-source to 36N9 / ACOTO	federal procurement competition: CICA, the FAR competition requirements, the law on statutory sole-source directions. Outcome NOT stated.
perpetual royalty payable in VINO	appropriations: the Appropriations Clause, the creation and funding of continuing obligations, the antideficiency rules. Outcome NOT stated.
sovereign license from a private sovereign entity	delegation of a sovereign monetary function; constitutional questions on currency and on emolument-like benefits. Outcome NOT stated.
Zero Note bearing the architect's likeness	conflict of interest and the appearance standard; the constitutional provisions on currency. Outcome NOT stated.

Every row states the provision as the Act specifies it and names the exposure at the level of its general nature, and every row sends the specific statute, threshold, and standard to the Verification Register and to counsel, and states no outcome. That is the whole specification of how the provisions are handled: in print, with the exposure attached, and the decision left where it belongs.

5.3.4 The disclosure test, applied to each phase

The disclosure test, phase by phase.

Phase 1 passes cleanly, as the nine-capital standard established: the standard is disclosure, and its value is independent of anyone knowing about it.

Phase 2 passes, subject to the store–rail rule: a parallel rail is a disclosed, filed, and defensible piece of settlement infrastructure, and its hedge value against the tail scenario is independent of concealment — it works by existing, not by being secret. The qualification is the transit-versus-store one, not a concealment one.

Phase 3's currency mechanism passes the concealment test — a capital-backed currency with a published backing discipline and an auditable integrity layer is the opposite of concealment. **But Phase 3's licensing provisions engage the framework's second axis rather than its first: they are not a concealment problem, they are a lawful-process problem, and the framework's own taxonomy separates the two.** The disclosure test asks whether a structure survives being seen; the licensing provisions survive being seen in the sense that they are not hidden, and what being seen reveals is a set of substantive legal exposures that disclosure does not cure and is not meant to. A provision can pass the concealment test and still invite serious review, because passing the concealment test means only that nothing is hidden — it does not mean that what is disclosed is unproblematic, and conflating the two would be the framework's own test misused to launder an exposure it was never designed to assess. That distinction is the honest position on the licensing provisions: they are disclosed, they are not concealment, and their exposure is real and is named, and the two facts do not cancel.

5.4 Failure modes

FM 5.1 Asserting Phase 3 as immediate

The failure the phasing exists to prevent. Presenting the currency as an available step rather than a gated terminal option asserts that gates have passed which have not, and it forfeits the standalone value of Phase 1 to chase the 0.090 branch. Every phase beyond the first is premature until its gate is demonstrably passed, and premature is not the same as impossible — but asserting it as present is the error.

FM 5.2 Presenting the parallel rail as a replacement

A hedge reduces loss in a bad state; it does not change the base case or predict it. Calling the VINO rail a replacement for the dollar turns a defensible tail hedge into an indefensible declaration, and it imports the refused collapse forecast through the rail. The rail is a second leg, not a succession.

FM 5.3 Bundling the phases

A package stands or falls together and invites a reviewer to test the weakest component and reject the whole. The phases are dependent and separately gradable, and presenting them as one all-or-nothing proposal discards the entire advantage of the phased form.

FM 5.4 Omitting the licensing provisions

The provisions are the Act's, the source manuscript is explicit that they appear as written, and a paper that hid the most scrutiny-inviting part of the proposal would violate its own disclosure principle. They are in print.

FM 5.5 Omitting the exposure

Stating the provisions without their exposure would be advocacy, not analysis, and would fail the framework's method as badly as hiding the provisions would. The exposure is named at the level of its general nature and sent to counsel. Both the provisions and the exposure are in print; neither alone is honest.

FM 5.6 Treating the disclosure test as clearing the licensing exposure

The licensing provisions pass the concealment test — nothing is hidden — and that pass does not cure the lawful-process exposure, which sits on the framework's second axis. Using the concealment pass to imply the provisions are unproblematic is the framework's own test misused, and it is called out here so it cannot be.

FM 5.7 Grading the Fibonacci and fractal features as monetary properties

The Fibonacci denominations are a combinatorial convenience and the golden-ratio limit is a fact about the sequence, not a property of the currency; the fractal reserve rule is a property of the rule, not a guarantee of stability. Grading any of them as a source of monetary value repeats the [T4] error the integrity-layer paper assessed and discarded, and this paper grades them the same way.

FM 5.8 Reintroducing the corrected goals

Monetary immunity and counterfeit impossibility were corrected in the nine-capital standard and are corrected again by reference here. Any Phase 3 description that reasserts either as achieved fact has reintroduced a claim the framework has twice refused, and a reviewer who caught the correction in the nine-capital standard will catch the reintroduction here.

FM 5.9 Skipping the within-Phase-3 gate

The provisions must be assessed by counsel before any is enacted, not after. Enacting first and assessing later is the sequence the integrity-layer paper warned against for the resolution designator and the nine-capital standard warned against for the supplementary standard — a good process in the wrong order becomes a bad outcome, and here the wrong order is more expensive because the subject is a national currency.

5.5 Results: the gated value and the exposure, computed

> **CONSTRUCTED.** Applied to the national position of the reserve-drift and nine-capital papers and to the framework's carried private structure. No figure is a measurement of any real nation or entity, and no phase beyond the first is asserted to have been adopted anywhere. All probabilities and figures are constructed to demonstrate the mechanism and are flagged in the Verification Register.

5.5.1 The three phases, in order, with their gates

phase	what it is	gate to the next	stands alone?
1	nine-capital accounting standard (CCP-068)	adopted, priced by creditors	YES -- complete on its own
2	parallel VINO rail as a dedollarisation hedge	carries real volume through a dollar-rail stress; ordering + identity layers built and externally validated	no -- gated on 1
3	the American Roots currency (VINO, Fibonacci, fractal reserve, XUSD)	-- terminal --	no -- gated on 2

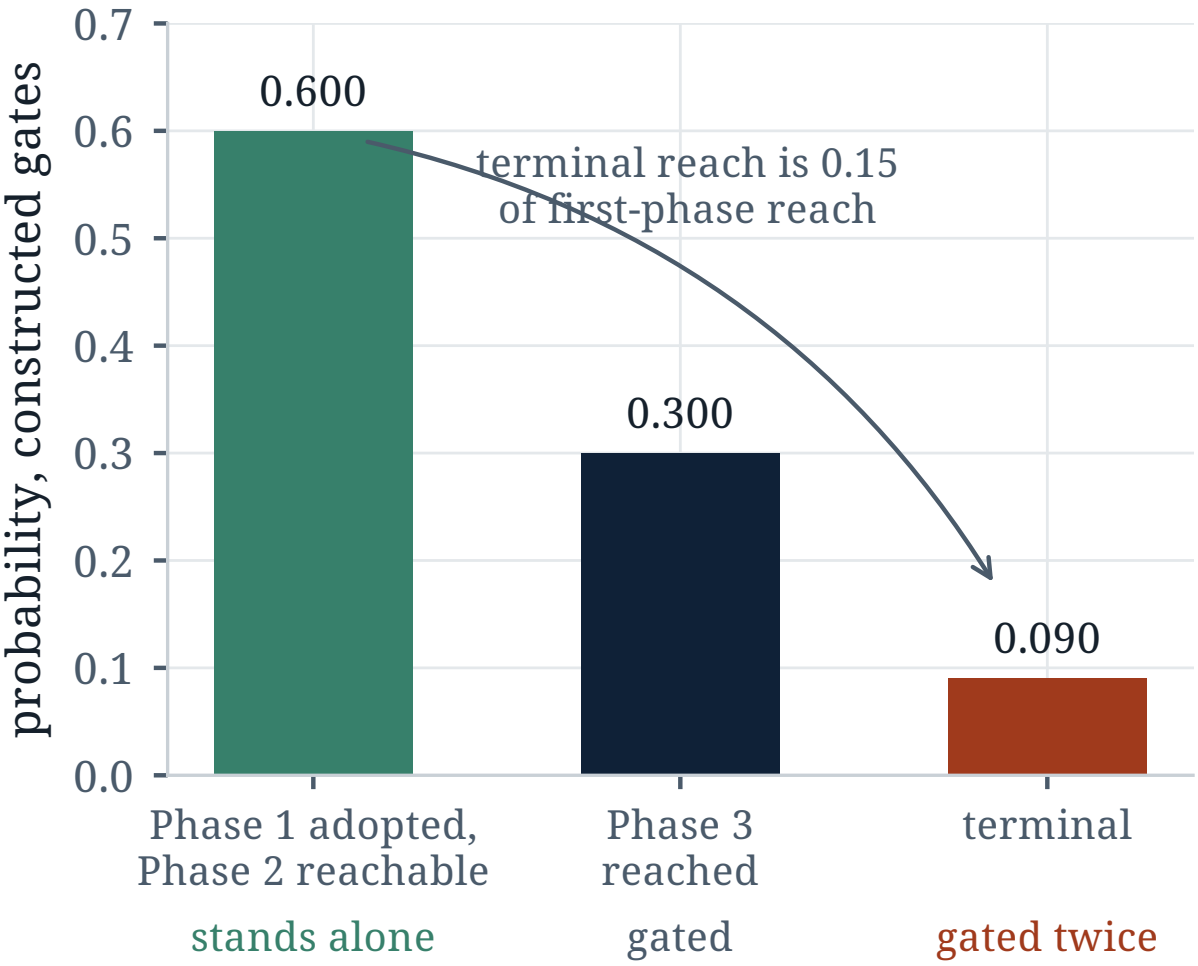


Figure 69.3: **Most of the programme's expected value sits in the phase that stands alone: the terminal phase is reached on nine constructed branches in a hundred against sixty for the standalone reform.** CONSTRUCTED gate probabilities of §5.5.2 ($p_1 = 0.60$, $p_2 = 0.50$, $p_3 = 0.30$), giving 0.600, 0.300 and 0.090 and a terminal-to-first-phase ratio of 0.15. The probabilities are illustrative; the conclusion needs only dependence and gate probabilities below one. Script `figs/src/CCP-069_fig03.py`.

Phase 1 is the only row that stands alone, and that is the design: the modest reform is complete in itself, and the two ambitious phases are options it makes available, each gated on the last.

5.5.2 The gated value, computed

```
CONSTRUCTED probabilities:   p1 = 0.60, p2 = 0.50, p3 = 0.30
-----------------------------------------------------------------
reach Phase 2  = 0.600
reach Phase 3  = 0.300
terminal       = 0.090
terminal / Phase 1 reach = 0.090 / 0.600 = 0.15
```

The terminal phase is reached on nine constructed branches in a hundred; the standalone reform pays off on sixty. **Leading with Phase 1 banks the sixty; leading with Phase 3 gambles the sixty to chase the nine, which is the concentration paper's highest-marginal-value-first rule and the priced-loss discipline applied to the order of a national reform.** The number is

constructed and the conclusion is not sensitive to it: any dependent phasing with gate probabilities below one puts most of the value in the standalone phase.

5.5.3 The licensing provisions and their exposure, together

the Act specifies	the review it invites (general nature)
sovereign license (Principality of Sicily)	delegation of a sovereign monetary function; constitutional review. Outcome not stated.
perpetual royalty in VINO	appropriations and continuing-obligation law. Outcome not stated.
sole-source to 36N9 / ACOTO	procurement competition (CICA, FAR). Outcome not stated.
Zero Note, architect's likeness	conflict of interest; the appearance standard; currency provisions. Outcome not stated.

The provisions are the Act's and are stated as the Act specifies them; the exposure is real and is named at the level of its general nature; the outcome is not stated and the decision is left with the principal. This is the framework's method for a bright line and a sanctions test, applied to a licensing structure, and it is the reason both the provisions and their exposure appear in print rather than either being hidden.

5.5.4 The carried structure, mapped onto the phases

CONSTRUCTED. The private structure the framework has built is itself a phased thing, and the mapping is exact enough to be worth stating. Phase 1 corresponds to the off-ledger register of the double-entry-limits paper — the standalone, adoptable-now discipline of naming what the ledger cannot carry. Phase 2 corresponds to the dual-rail settlement of the dual-rail settlement analysis — the parallel rail added as a hedge, measured on the settlement axis, governed by the store–rail rule. Phase 3 has no clean private analogue, and that absence is instructive: the private structure never reaches a phase where it licenses its own architecture from a connected party or brands its instruments with a principal's likeness, because a private structure that did so would fail the framework's tests on the same axis a sovereign one does — the conflict is structural, not scale-dependent, and it is named the same way at both scales. The private structure stops at the equivalent of Phase 2 because Phase 3's licensing provisions are exactly the kind of self-dealing the framework's method surfaces rather than hides.

5.5.5 The disclosure test, applied

phase	concealment test	the honest reading
1 accounting standard	PASS	the standard is disclosure
2 parallel rail	PASS (the store-rail rule)	hedge value independent of secrecy
3 currency mechanism	PASS	published backing, auditable layer
3 licensing provisions	PASS as to concealment	NOT concealment, but a lawful-process exposure on the framework's second axis, named and not cured by disclosure

The licensing provisions are not hidden and therefore pass the concealment test, and passing it does not make them unproblematic — it makes them disclosed, which is where the framework leaves them: disclosed, exposed, and decided by the principal with counsel. Nothing about the disclosure clears the procurement, appropriations, conflict, or constitutional review; disclosure was never meant to, and saying so is the last piece of honesty the section owes.

5.5.6 The residual, stated plainly

- Phase 1 is complete on its own, adoptable now, mints nothing, and carries the petrodollar answer by widening the priceable base. **It is the whole recommendation the framework is willing to make without a gate; everything past it is an option, not a step.**
- Phase 2 is a hedge, not a replacement, gated on Phase 1, measured on the settlement axis, governed by the store–rail rule. It is worth having against the tail and does nothing in the median, and the section says so.
- Phase 3 is the terminal option, gated on Phase 2, asserted as immediate by no sentence here. Its Fibonacci and fractal features are combinatorial and structural choices, not monetary properties, and its immunity and counterfeit-impossibility goals are corrected as they were in the nine-capital standard.
- The licensing provisions — the sovereign license from the Principality of Sicily, the perpetual royalty payable in VINO, the sole-source procurement to 36N9 Genetics / ACOTO, and the Zero Note bearing the architect's likeness — are stated as the Act specifies them, and the federal procurement, appropriations, conflict-of-interest, and constitutional review they invite is stated in the same register. **Both are in print; neither is omitted; the outcome is not stated and the decision is left with the principal.**
- Every figure is constructed and flagged, no phase beyond the first is asserted to have been adopted, and every specific statute, threshold, and standard bearing on the licensing provisions is in the Verification Register and belongs to counsel.
- The disclosure test separates cleanly from the exposure question: the licensing provisions are not concealment, and their not being concealment is not a clearance. A thing can be fully disclosed and still invite serious review, and the framework's honest position is to hold both facts at once rather than let the first stand in for the second.

The honest verdict on the phased integration is that its strength is entirely in its order. The accounting standard stands alone and should be judged alone; the rail is a gated hedge and should be judged as one; the currency is a conditional terminal option carrying contested design goals and scrutiny-inviting licensing provisions, and it should be judged as the far end of a dependency chain rather than as an available act. The framework's closing argument is not that America should adopt the currency; it is that America can adopt the accounting standard now, on its own authority, as the standard-setter rather than a licensee — and that everything past that first phase is an option to be gated, priced, and decided with counsel, with both its provisions and its exposures in print. That is the proposal stated in the framework's own register: the modest reform recommended without a gate, the ambitious ones named with their conditions and their costs, and the decision, as always, left with the principal.

---

6. The failure record: five tracks, and the three failures the instrument predicts

6.1 The position: a court with personal jurisdiction over the principal does not need jurisdiction over the assets

The preceding papers of this series describe how structures work. This section describes how they stop working, and it is the section the rest of the series was written to earn the right to include.

There is a reason for putting it here rather than at the front. A failure record read before the mechanism is understood teaches caution, which is not useful, because caution is a mood rather than a design input. A failure record read after the mechanism is understood teaches something sharper: which specific element of a specific structure was load-bearing, and what happened when the weight came onto it. The cases below are not cautionary tales. They are

stress tests that were run in public, on real structures, by adversaries who were not obliged to be fair, and the results were written down. That is a better dataset than any adviser's success record, because success records are selected on survival and the failure record is not.

The section has an organising claim, and it is the single most important sentence in the framework for anyone who has been sold an asset-protection structure on the strength of a foreign statute:

> **A court with personal jurisdiction over the principal does not need jurisdiction over the assets.**

That is the lesson of *In re Lawrence* and *FTC v. Affordable Media*, and it is why the erosion paper opened with them rather than with a jurisdiction. The offshore statute governs the offshore asset. The domestic court governs the person. An order directed at a person's own conduct — turn this over, repatriate that, sign this instruction — requires no foreign court to cooperate with anything, no treaty, no letter rogatory, no recognition of a foreign judgment. It requires only that the person be within reach. And the person is the one component of any structure that cannot be redomiciled by filing a form.

Everything else in this section is a variation on that theme or an answer to a different question. There are five failure tracks in the record, and they are genuinely different from one another:

```
1  PERSONAL COMPULSION      contempt, repatriation orders, the impossibility defence
2  SUBSTANCE                the entity was where you said, and did nothing there
3  SHAM AND NOMINEE         the documents said one thing and the parties meant another
4  LOSS OF CONTROL          nobody sued; the collateral moved and the pledge did the rest
5  MEASUREMENT ERROR        the structure did exactly what it was designed to do, and
                            the design was measured wrong
```

Tracks 1 through 4 are the conventional failure record and they are drawn from decided cases, statutes, and arithmetic. Track 5 is this framework's own contribution and it is the one that will do the most damage to the most readers, because it produces no litigation, no adverse finding, and no moment at which anybody discovers there is a problem. It produces a structure that looks diversified, costs money to maintain, and concentrates exactly where the holder believes it does not. The concentration instrument was built to detect it. This section reports what the instrument found when it was pointed at the framework's own recommendations, and the finding is not flattering.

One boundary before the record. The threat register separated arbitrary state action from lawful process and insisted the two require different structures. That separation holds here and explains most of what follows. **The failures in tracks 1 and 3 are almost entirely failures against lawful process — a creditor with a judgment, a trustee in bankruptcy, a regulator with a court order — and they are timing failures rather than jurisdictional ones.** The structures did not lose because the jurisdiction was wrong. They lost because they were built after the claim arose, or built in a form that gave a court something to characterise. Fraudulent-transfer law, not haven selection, governs that category, and this framework has said from the threat register onward that it is a different problem from the one the architecture solves.

The failures in track 5 are the ones that bear directly on the arbitrary-action problem this framework does address, and they are self-inflicted.

6.2 The five failure tracks

6.2.1 Track one: personal compulsion, and the collapse of the impossibility defence

The modern asset-protection trust was a statutory product. **Every case, statute and rule named in Sections 6.2.1 to 6.2.3 is described from the source manuscript and is cited nowhere,**

because this paper's bibliography contains no case law and no jurisdiction-specific instruments; each is itemised in Appendix A with what requires verification. What the bibliography does supply is the doctrinal frame: recognition of trusts is conventional and its ratification list is short [Hague Conference on Private International Law, 1985], the enforcement of awards runs on a separate instrument with its own refusal grounds [United Nations, 1958, van den Berg, 1981, Born, 2021], and the conflict-of-laws treatment of all of it is standard [Torremans et al., 2022, Briggs, 2013, Fawcett and Carruthers, 2008]. Its design premise was that a jurisdiction could legislate its way out of the enforcement of foreign judgments: short limitation periods for fraudulent-transfer claims, a heightened standard of proof, a requirement that any claim be brought locally and afresh, and an express provision that a foreign judgment would not be recognised or enforced against the trust. The Cook Islands International Trusts Act 1984, substantially amended in 1989 and repeatedly since, was the template. A reader who needs operative provisions should obtain the current consolidated text from Cook Islands counsel; the section numbering has moved across amendments and this framework does not guess at it.

The premise was correct about what it addressed. Those statutes did make it hard to enforce a foreign judgment against a Cook Islands trust in a Cook Islands court. What they did not address — what no foreign statute can address — is a domestic court's power over a person standing in front of it.

In *FTC v. Affordable Media, LLC*, 179 F.3d 1228 (9th Cir. 1999), the Federal Trade Commission obtained orders against Michael and Denyse Anderson requiring repatriation of assets held through a Cook Islands trust. The Andersons had retained roles in the structure. When the district court ordered repatriation, they asserted that compliance was impossible, because the foreign trustee had invoked a duress provision and removed them from control. The district court held them in civil contempt and they were taken into custody. The Ninth Circuit affirmed.

The reasoning is the part that matters, and it is worth stating in the form a designer would use. A party asserting impossibility as a defence to civil contempt bears the burden of proving it. The court was unpersuaded by an impossibility that the contemnors had themselves designed, funded, and triggered. The duress clause was the whole product; it was also the evidence. **The clause engineered to make compliance impossible is, in the hands of the court being asked to accept the impossibility, the clearest available proof that the impossibility was manufactured for exactly this moment.** There is no way to draft around that, because the drafting is the problem.

In *In re Lawrence*, 279 F.3d 1294 (11th Cir. 2002), the bankruptcy debtor Stephan Jay Lawrence had settled a substantial offshore trust. A correction is required here, because the case is routinely miscited in the planning literature: the Lawrence trust was settled in Jersey in 1991 and its administration later moved to Mauritius. It was not a Cook Islands trust. Any source saying otherwise is unreliable on the rest of its detail too, and a reader who finds the error in a brochure has learned something useful about the brochure. The trust contained the same family of protective provisions. Lawrence was ordered to turn over the trust assets, asserted that he could not, was held in civil contempt, and was taken into custody in 1999. He was reported released in 2007. This framework will not state a precise custody duration, because the reported dates vary and the docket should be consulted; what is not in doubt is that the incarceration ran for a period measured in years rather than months. Lower-court decisions in the same matter include 227 B.R. 907 (Bankr. S.D. Fla. 1998) and 251 B.R. 630 (S.D. Fla. 2000).

Two related decisions are frequently cited alongside these and are worth pulling if the question is live: *In re Portnoy*, 201 B.R. 685 (Bankr. S.D.N.Y. 1996), and *In re Brooks*, 217 B.R. 98 (Bankr. D. Conn. 1998), both involving Channel Islands structures and both unfavourable to the settlor. Confirm the citations and their subsequent treatment; this framework has not verified their current authority and says so.

The English route to the same destination is different in doctrine and identical in effect. The freezing injunction, established in *Mareva Compania Naviera SA v International Bulkcar-*

riers SA [1975] 2 Lloyd's Rep 509, operates against the person rather than against the asset, which is why it can and does have worldwide reach with the usual undertakings and provisos. The contrast with the United States position on interlocutory relief before judgment is instructive: in *Grupo Mexicano de Desarrollo, S.A. v. Alliance Bond Fund, Inc.*, 527 U.S. 308 (1999), the Supreme Court held that a federal court sitting in a money-damages action does not have the equitable power to issue a pre-judgment injunction restraining a defendant from dealing with unencumbered assets. The two systems arrive at personal compulsion by different roads and on different timetables. **Neither system requires the foreign asset to be within its territory, and neither requires the foreign court to agree.**

There is a legislative postscript that matters more than any of the cases for anyone building today. Section 548(e) of the United States Bankruptcy Code, added by the Bankruptcy Abuse Prevention and Consumer Protection Act of 2005, provides a ten-year look-back for transfers to self-settled trusts made with actual intent to hinder, delay, or defraud. Ten years is longer than most planning horizons and far longer than the limitation periods the offshore statutes provide. The domestic reach-back was simply extended past the foreign protection. On the English side, section 423 of the Insolvency Act 1986 fixes no period at all; the authorities, including *Hill v Spread Trustee Co Ltd* [2006] EWCA Civ 542, address how the Limitation Act 1980 applies, and this is contested ground on which counsel is required rather than a paper.

The generalisable lesson for a designer is not "avoid trusts." It is a statement about where the protection actually sits, and it is the reason this framework's architecture puts a legal person with its own governing council at the apex rather than a fiduciary relationship dependent on a settlor who retains powers:

> **The offshore statute governs the offshore asset. The domestic court governs the person. The person is the part of the structure that cannot be redomiciled, and a structure whose protection depends on the principal's continuing inability to comply is a structure that has volunteered the principal as the point of attack.**

The foundation form is not immune to this. Nothing is. What it does differently is separate the question of who owns from the question of who instructs — asset partitioning being the irreducible function that entity law performs and contract cannot [Hansmann and Kraakman, 2000], and the comparative trust-versus-foundation question being a live one in its own literature [Lupoi, 2000, Smith, 2013, Panico, 2014, Hayton et al., 2016], so that an order compelling the principal to instruct meets a council with its own duties, a Guardian where the statute requires one, and a court with its own view. That is a real difference in the middle of the distribution and it is not a shield at the tail. A reader who takes anything else from the foundation paper has read it wrong.

6.2.2 Track two: substance, and the entity that was where it said and did nothing there

The second track produced no jailings and vastly more destroyed value, because it operated by administrative rule rather than by litigation and it operated on everybody at once.

The erosion paper traced the mechanism. The disclosure regime that produced it is the reason non-disclosure is no longer an available design parameter [Organisation for Economic Co-operation and Development, 2017, Financial Action Task Force, 2012, Organisation for Economic Co-operation and Development, 2015, de Goede, 2011], and the general account of how legal coding creates and destroys capital is the corpus's best scholarly ally and its sharpest critic in one book [Pistor, 2019]. The European Union's Code of Conduct Group on Business Taxation made economic substance the listing criterion for zero and near-zero rate jurisdictions: an entity claiming residence had to be directed and managed locally, with adequate people, premises, and expenditure proportionate to its activity. Cayman legislated the International Tax Co-operation (Economic Substance) Act, originally enacted in 2018 and in force from 1 January 2019. The British Virgin Islands legislated the Economic Substance (Companies and Limited Partnerships) Act 2018, in force on the same date. Bermuda legislated the Economic Substance Act 2018. Jersey, Guernsey, and the Isle of Man each legislated on the

same external template with effect from 1 January 2019. Every one of these has been revised since and a reader relying on operative text should confirm the current revision with local counsel rather than working from the original enactment.

The failure mode this created is easy to state and was widespread. A holding company was incorporated in a zero-rate jurisdiction, given a registered office and a corporate director supplied by an agent, and used to hold shares in operating subsidiaries elsewhere. Nothing about that was concealed. It appeared on every filing. It had worked for two decades. Then the test changed from *where is it registered* to *what does it do there*, and the honest answer was nothing, and the entity failed a test that had not existed when it was built.

Three separate mechanisms then converged on the same entity, and it is the convergence rather than any one of them that makes substance the most reliably fatal of the conventional failure tracks:

```
(a) LOCAL SUBSTANCE TEST      the economic substance regime in the entity's own
                              jurisdiction: adequate people, premises, expenditure,
                              local direction and management, with reporting and
                              penalties for failure
(b) FOREIGN RESIDENCE TEST    central management and control, or place of effective
                              management, asserted by another state -- so the entity
                              becomes tax resident where the decisions are actually
                              taken, which is usually where the principal sits
(c) TREATY ACCESS TEST        the principal purpose test introduced into treaties by
                              the Multilateral Instrument, and certificate-of-residence
                              practice, both of which can deny the treaty benefit that
                              was the entity's only reason for existing
```

An entity can satisfy (a) and still fail (b), which is the configuration that surprises people most: the jurisdiction is satisfied that the company is adequately resourced locally, and a foreign revenue authority nonetheless treats it as resident elsewhere because that is where the board decisions are in fact taken. The three-jobs diagnostic identified the underlying error and named it: one entity asked to do three jobs, with the principal signing everything from one desk. The substance regimes did not create that error. They made it expensive.

The Thai variant of this track is the sharpest in the framework because it is criminal rather than fiscal. Under the Foreign Business Act B.E. 2542 (1999), nominee shareholding is an offence and it is prosecuted. A structure using Thai nominees to hold shares that a foreigner beneficially owns is not a structure with a substance problem. It is an offence with a filing history. The Thailand analysis sets out the two lawful routes — Board of Investment promotion, and the Treaty of Amity and Economic Relations between the United States of America and the Kingdom of Thailand of 1966 for United States persons — and the Thailand paper's register item 19, on regularising a historical nominee position, is described there as the most important item in that register. It is repeated in this section's register for the same reason.

The disclosure test does the work here without any additional analysis. A nominee structure fails the disclosure test by construction: disclose it to every authority with jurisdiction and it stops working immediately, because its operation depends on the register saying something other than what is true. That is the definition of a concealment structure, and this framework excludes them not on grounds of taste but because they are load-bearing on non-disclosure and non-disclosure is not a durable input.

6.2.3 Track three: sham, nominee, and the gap between the documents and the intention

Track three is narrower than the planning literature suggests and more dangerous than it looks, because the doctrine that everyone fears is rarely the doctrine that decides the case.

Sham, in the common-law sense, requires a common intention that the documents should not create the rights and obligations they appear to create. The definition most often cited derives from *Snook v London and West Riding Investments Ltd* [1967] 2 QB 786 — verify the report and the passage before relying on it — and the threshold it sets is high. Two consequences follow, and they point in opposite directions.

The first: sham findings against properly executed structures are uncommon, because a settlor who genuinely wanted the trust to operate, and a trustee who genuinely accepted the office, do not have the required common intention even if the settlor was hoping the structure would inconvenience a creditor. Hoping for an outcome is not intending a pretence.

The second, and this is the trap: courts have a great deal of machinery that reaches the same practical result without a sham finding. A structure can be entirely genuine and still be unwound, disregarded, or made irrelevant by:

```
- a finding that the trustee acted as the settlor's agent or nominee in fact,
  so that the property was always dealt with as the settlor's
- a fraudulent-transfer or transaction-defrauding-creditors finding, which
  does not question the validity of the structure at all -- it simply
  reverses the transfer into it
- a finding that a power retained by the settlor is property capable of being
  reached, or that its exercise can be compelled
- an order against the person, which is track one, and which does not care
  whether the trust is valid
```

The Jersey line is the best-documented body of authority on the first of these. *Rahman v Chase Bank (CI) Trust Company Limited* [1991] JLR 103 and *In re the Esteem Settlement* [2003] JLR 188 are the two most commonly cited, and they are frequently cited for opposite propositions by advisers on opposite sides of a question. Confirm each against the law report; this framework names them and does not paraphrase their holdings, because the paraphrases in circulation are unreliable.

The "illusory trust" line — the argument that a settlor who retains sufficient powers has not, in substance, disposed of anything — has been litigated in New Zealand and before the Privy Council, and the foundation-versus-trust paper's register item 14 declines to name the citations or state their current standing before verification. That refusal stands here. **Declining to name a case whose holding the framework has not checked is the house standard, and a failure analysis is the last place to abandon it.**

What the designer should take from track three is a drafting principle rather than a case list. Reserved powers are useful, are expressly contemplated by the statutes this framework recommends, and are the reason the foundation form is workable for a founder who will not give up control entirely. They are also the raw material from which every one of the four machinery routes above is built. The question to ask of every reserved power is not "is this permitted?" — the statute answers that — but "if a court compels its exercise, what happens?" A power whose compelled exercise produces the creditor's desired outcome is not protection. It is a lever with the principal's name on it, installed at the principal's request.

6.2.4 Track four: loss of control without a plaintiff

The fourth track is the one that requires no adversary at all, and it is therefore the one most likely to affect a reader who has done everything else correctly.

The carve-out cascade built the leverage layer of the carried case and this framework's instruction was that the section lead with the failure mode rather than bury it. The arithmetic is short enough to reproduce in full, and every figure below was recomputed for this section rather than copied from the carve-out paper's printed output:

```
    CARVE-OUT CASCADE -- the structure
```

```
  operating company, marked value                        100.00
  carve out 20% by initial public offering                20.00   proceeds
  retain                                                  80.00   pledged
  pledge at 50% loan-to-value                             40.00   borrowing
  ------------------------------------------------------------
  deployable = proceeds + borrowing                       60.00
  retained control                                        80.0%

THE TRIGGER

  collateral value V                                      80.00
  loan L                                                  40.00
  maintenance requirement m                               35%   of collateral value
  critical collateral value = L / (1 - m)
                            = 40.00 / 0.65               61.538462
  drawdown to trigger = 1 - 61.538462 / 80.00
                      = 0.230769                          23.0769%
```

A 23.1 per cent drawdown in the marked value of the pledged stake triggers the maintenance call, which is the corpus's pledge-trigger result [Curzi, 2026]: $d^{*} = 1 - \mathrm{LTV}_{0}/(1-m)$, so at $\mathrm{LTV}_{0} = 50$ per cent and a maintenance requirement $m = 35$ per cent the first margin call falls at a 23.076923 per cent drawdown and not at 50 per cent. That is not a crisis parameter. It is ordinary equity volatility, of the kind a single listed holding delivers in an unremarkable year, and it is very much smaller than the drawdown a family would mentally associate with the word "risk."

What happens at the trigger is the failure. The lender does not need a judgment, a court, or a finding of wrongdoing. It needs the contract, which the principal signed, and it exercises rights the principal granted. The possible outcomes are: post additional collateral, which requires liquid assets the structure was designed to convert into illiquid ones; repay part of the loan, which requires the same; or have the collateral sold, which converts a control position into a minority one at the worst available price and may itself breach a change-of-control provision elsewhere in the group, cascading into cross-default. The ring-fencing analysis helps only to the extent it was built before the pledge; a cross-default clause is a designed hole in a ring fence.

Three further mechanisms belong in this track and are routinely under-modelled:

Rehypothecation. Where collateral may be re-used by the lender, the principal's position in a lender insolvency is different from what a custody statement suggests. This is a documentation question and it is answered in the security agreement, not in the jurisdiction.

Change-of-control provisions. These appear in loan agreements, joint-venture agreements, licences, distribution agreements, and regulatory permissions. A forced sale of pledged shares can trip a definition of change of control in an agreement that has nothing to do with the pledge and that nobody re-read when the pledge was granted.

Disclosure obligations. The fourth bright line of the group-architecture paper applies: directors and officers must disclose share pledges in proxy statements — the requirement is commonly cited as Item 403 of Regulation S-K and should be confirmed with securities counsel in its current form. The pledge is therefore visible. That is a feature under the disclosure standard, and it is also the reason a forced sale is legible to the market before it completes.

The structural point is the one that makes this track belong in a paper about jurisdictions at all. **Every other failure track in this section requires an adversary with standing. This one requires a price.** No jurisdiction protects against it, because the loss of control is contractual and consensual, and the contract was chosen. The concentration instrument does not see it either: a pledged position and an unpledged position have identical jurisdictional shares and produce identical values of N. That is a real limitation of the instrument, stated here rather than in an appendix, and it is why the threat register keeps commercial and counterparty risk

as a separate threat class with its own answers rather than folding it into the concentration measurement.

6.2.5 Track five: the failures the framework's own instrument predicts

The first four tracks are the record. This one is the forecast, and it is the reason a failure section appears in a paper that spent forty-three sections recommending things.

Three specific failures are predicted by the architecture's own arithmetic. Each has a number. Each was produced by the same instrument that produced the recommendations, which is the only reason any weight should be given to it.

6.2.6 Predicted failure (a): the unpriced bloc call, a loss of 0.005110

The AIFC pricing assessed the Astana International Financial Centre and did something no other jurisdiction paper does: it refused to argue about whether Kazakhstan is an independent coordination perimeter, and priced the question instead.

The structure of the calculation is worth reproducing because it generalises to every marginal jurisdiction a reader will ever consider. Take the designed architecture at the point where the jurisdiction papers have built it, and place a 5% position in the AIFC drawn from the United Kingdom slot. If the bloc call is right — Kazakhstan is a genuinely independent perimeter, a fourth bloc — the position adds a bloc. If the bloc call is wrong — Kazakhstan folds into the non-aligned bloc alongside Thailand and India — the position adds nothing and concentrates the bloc it joined.

```
BASELINE, designed structure (CCP-052)
  blocs  BRIDGE 40.0  NONALIGNED 30.0  WEST 30.0
  bloc HHI = 0.340000     N = 2.941176     h = -0.122417

AIFC 5.0 taken from the UK slot -- UK 16.4 -> 11.4
  jurisdictions  UAE 29.1  Thailand 18.2  Bermuda 13.6  India 11.8
                 UK 11.4   Singapore 10.9  AIFC 5.0
  naive HHI = 0.177602    effective jurisdictions = 5.630567

IF THE CALL IS RIGHT -- AIFC is a fourth bloc
  blocs  BRIDGE 40.0  NONALIGNED 30.0  WEST 25.0  EURASIAN 5.0
  bloc HHI = 0.315000     N = 3.174603     h = -0.095686
  GAIN                                          +0.026731

IF THE CALL IS WRONG -- AIFC folds into NONALIGNED
  blocs  BRIDGE 40.0  NONALIGNED 35.0  WEST 25.0
  bloc HHI = 0.345000     N = 2.898551     h = -0.127526
  LOSS                                          -0.005110
```

The wrong bloc call does not merely fail to help. It costs 0.005110 in h, because the capital was taken out of a bloc where it was diversifying and put into a bloc it concentrated. That is the entire failure. It is invisible to every conventional measure: the jurisdiction count went up, the naive effective count went up from 5.290949 to 5.630567, the structure acquired an English-common-law leg with an English-judge court, and the honest count went down.

The break-even follows immediately and is the reason this is a priced decision rather than an opinion. Let p be the probability that the perimeter call is right. The position pays when

```
p x 0.026731  -  (1 - p) x 0.005110  =  0
p  =  0.005110 / (0.026731 + 0.005110)  =  0.005110 / 0.031841  =  0.160476
```

A 16.05 per cent probability that the jurisdiction is an independent perimeter is enough to justify the position [Curzi, 2026], whose Result One states the same bet as $E[\mathrm{d}h] = 0.0318401p - 0.0051096$ from a right-call gain of $+0.026731$ and a wrong-call loss of -0.005110. That is a low bar and the position is therefore attractive on any honest view of the evidence. The AIFC pricing then identifies the binding constraint, which is not the geopolitics at all: the gain of 0.026731 is consumed entirely if the leg costs more than 0.026731 of maintenance, which against the tested `C = 0.30` is 8.9102% of the structure's total maintenance budget. Cost, not geopolitics, decides it.

The generalisation is the part to carry away, and it is stated as an instruction rather than an observation because that is how it should be used:

> **Price every marginal jurisdiction as a break-even probability, not as a geopolitical opinion. An unpriced bloc call is the failure. A priced one that turns out wrong is a cost that was known in advance and accepted.**

6.2.7 Predicted failure (b): the apparent-diversification trap

The second predicted failure is the cleanest demonstration in the framework of what the concentration paper's instrument is for, and the Gulf-additions measurement produced it.

The carried family is offered a Saudi regional headquarters and a Qatar Financial Centre family-office vehicle, and funds them by trimming the UAE, United Kingdom, and Singapore allocations. Eight jurisdictions instead of six. Every brochure metric improves.

```
DESIGNED (CCP-052)                    TEMPTED
  UAE         29.1  BRIDGE            UAE         20.0  BRIDGE (GCC)
  Thailand    18.2  NONALIGNED        Thailand    18.2  NONALIGNED
  UK          16.4  WEST              Bermuda     13.6  WEST
  Bermuda     13.6  WEST              UK          12.2  WEST
  India       11.8  NONALIGNED        India       11.8  NONALIGNED
  Singapore   10.9  BRIDGE            Saudi        9.1  BRIDGE (GCC)
                                      Qatar        8.2  BRIDGE (GCC)
                                      Singapore    6.9  BRIDGE

  naive HHI  = 0.189002               naive HHI  = 0.140194
  eff jur    = 5.290949               eff jur    = 7.132973    +34.81%

  blocs  BRIDGE 40.0                  blocs  BRIDGE     44.2
         NONALIGNED 30.0                     NONALIGNED 30.0
         WEST 30.0                           WEST       25.8

  bloc HHI = 0.340000                 bloc HHI = 0.351928
  N        = 2.941176                 N        = 2.841490     -3.39%
  h        = -0.122417                h        = -0.134485
```

Two more jurisdictions, a third more apparent diversification, and the honest count fell. This is the corpus's Gulf-additions result [Curzi, 2026]: adding the two jurisdictions raises the naive effective count from 5.29 to 7.13, a rise of 34.8 per cent, while the effective bloc count falls from 2.94 to 2.84, a fall of 3.4 per cent. The Gulf additions are all in the bloc the structure was already heaviest in. The funding came disproportionately from the bloc that was providing the diversification. Both effects run the same way and the naive metric reports neither.

The failure this predicts is not the Saudi entity or the Qatari entity. It is the decision procedure that produced them: a jurisdiction was added because it was good, rather than because it was independent of what the structure already held. The Gulf-additions paper's own worked outcome does the opposite and is the model. Take the Saudi regional headquarters at 5.0%

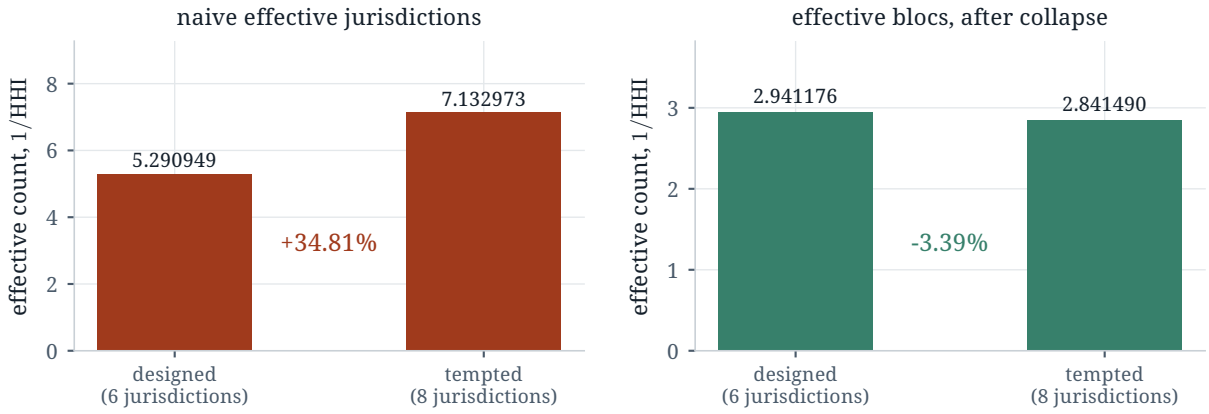


Figure 69.4: **Two more jurisdictions raise the naive effective count by 34.81 per cent while the effective bloc count falls by 3.39 per cent: two measures of the same change, moving in opposite directions.** CONSTRUCTED allocations of §6.2.7, shown as two panels rather than a dual axis. The bloc assignment behind the right-hand panel is this framework's own and is contestable. Script `figs/src/CCP-069_fig04.py`.

funded from the operating allocation rather than from the holding allocation, and the result is naive 5.688217 and blocs 2.816901 — a loss of 0.124275 in `N` and 0.015110 in `h`, taken deliberately, on revenue grounds, with the number written down.

> **A priced loss taken knowingly is legitimate. An unpriced one is the failure.** The Saudi entity is defensible at a cost of 0.124275 in effective bloc count if the addressable Saudi revenue exceeds it. It is indefensible as diversification, because it is the opposite of diversification, and the naive metric said otherwise.

The bridging-slot control supplies the control experiment that proves the point is about measurement rather than about the Gulf. Place the same 10.9% in London instead of Singapore. Naive effective jurisdictions fall from 5.290949 to 4.449309 — a large and visible difference — while the bloc-collapsed structure moves from `N = 2.941176` to `N = 2.924302` and `h` from −0.122417 to −0.124431. Singapore's true marginal contribution to `h`, measured against putting the same money in London, is 0.002014. **The naive metric reported a diversification of nearly a full effective jurisdiction that did not occur.**

6.2.8 Predicted failure (c): disclosure that did not confer immunity

The third predicted failure is the objection to this framework's governing principle, and the disclosure paper carries it as its own seventh failure mode rather than leaving it to a critic.

The disclosure paper argues that a structure is durable when being observed yields nothing actionable, because every element is lawful, filed, and defensible on its own merits. The counter-case is the strongest one available and it is recent. In 2022, approximately USD 300 billion of Russian central bank reserves were immobilised. Those reserves were among the most thoroughly disclosed assets on earth. Their location was known, their custodians were known, their character as sovereign reserves was known and was the reason they were held where they were held. Correctness and disclosure were complete. **They were frozen anyway.**

The claim that survives that event is narrower than the claim a careless reader takes from the disclosure paper, and the narrowing has to be stated in print or the whole principle is dishonest:

```
WHAT DISCLOSURE DOES        removes the failure mode in which a structure collapses
                            on contact with an authority, because there is nothing
                            to find; converts an enforcement problem into a policy
```

```
                          problem; makes the structure bankable, which the concentration paper's
                          failure mode 11 identifies as the binding constraint
                          on everything else

WHAT DISCLOSURE DOES NOT DO  prevent a sovereign decision taken on political grounds
                          against a target it has identified; supply an appeal;
                          or protect an asset sitting inside the perimeter that
                          took the decision
```

The 2022 event is not an argument against disclosure. It is an argument against **concentration**, and it is the reason the framework's answer to arbitrary action is jurisdictional and rail diversity across blocs rather than any quality of the structure's paperwork. The corpus's own closing register states the same admission without qualification [Curzi, 2026], listing the failure of disclosure to prevent the 2022 immobilisation among four things it concedes outright; and the correlation result that makes a single perimeter fatal is measured elsewhere [Curzi, 2026]. A fully disclosed, fully correct, single-perimeter structure has `N = 1.00` and, by the concentration arithmetic, `ln(1.00) = 0` and no available surplus at all. Disclosure did not save it and was never the variable that would.

> **Disclosure is not immunity. It is the precondition for a structure that can be built in the open, banked, and defended — and the thing that actually protects against arbitrary action is having somewhere else to be.**

6.3 Specification: the five-step procedure

What follows is a procedure rather than a warning. It has five parts and it is designed to be run before a structure is built and annually thereafter.

6.3.1 Classify the exposure before selecting the answer

The five-threat register are not interchangeable and the failure tracks map onto them unevenly. Run the mapping explicitly, because the commonest single error in this whole area is a structure built with great care against the wrong threat.

```
FAILURE TRACK          THREAT ANSWERED           WHAT ACTUALLY HELPS
-------------------------------------------------------------------------
1 personal compulsion  lawful process            timing; building before the
                                                 claim; not retaining powers
                                                 whose compelled exercise
                                                 gives the creditor the result
2 substance            lawful process (fiscal)   real activity, real people,
                                                 decisions taken where the
                                                 entity is
3 sham / nominee       lawful process            lawfulness; the disclosure
                                                 test; never a nominee where
                                                 nominees are unlawful
4 loss of control      commercial / counterparty headroom; covenant analysis;
                                                 ring-fencing built BEFORE
                                                 the pledge
5 measurement error    arbitrary action /        the concentration instrument, run
                       bloc fragmentation        on blocs and not on flags
```

The two columns on the right share nothing. That is the point of running the mapping.

6.3.2 Run the pre-mortem, in writing, with dates

A post-mortem asks why a structure failed. A pre-mortem assumes it has failed and asks what happened. The organisational literature on why warning signals are available and not acted upon is the reason the fourth field is mandatory [Vaughan, 1996, Weick, 1993, Weick and Sutcliffe, 2007, Perrow, 1984]. The second is more useful because it is available while there is still time to act.

The written form has four fields per scenario and should be no longer than a page:

```
SCENARIO        one sentence, in the past tense, dated
FIRST DOMINO    the single fact that had to be true first
WARNING SIGN    what would have been observable, and to whom, before the loss
THE MOVE        what would have had to be done, and by when, and by whom
```

The discipline the fourth field enforces is the valuable one. A response that has no named person and no date is not a plan; it is a hope. The relisting contingency of the rail paper and the Hong Kong analysis's migration path for the Hong Kong leg are both examples of this field filled in properly, and both were written before the event they contemplate.

6.3.3 Test every retained power against compelled exercise

For each power the principal retains anywhere in the structure — reserved powers under a foundation charter, a protector's consent right, a shareholder reserved matter, a bank mandate, a signature authority — write one line:

> If a court with personal jurisdiction over me orders me to exercise this power in the way my adversary wants, what happens?

Three answers are possible. *The adversary gets what it wants*: the power is a liability and should be removed or relocated to a person outside the compelling court's reach with genuine independent duties. *Nothing happens, because the power does not reach the asset*: the power is safe and useful. *The council or trustee retains a genuine discretion to refuse*: the power is partly safe, the protection is only as good as the independence of the person holding the discretion, and the independence must be real rather than documentary.

This is the single most useful hour a principal can spend on a structure, and almost nobody spends it, because the powers were all granted at different times by different advisers for different reasons and nobody has ever listed them on one page.

6.3.4 Compute the headroom before the leverage, not after

The the carve-out cascade arithmetic reversed. Instead of asking what a given loan-to-value produces, fix the drawdown you are willing to survive and solve for the loan:

```
    d  = drawdown to be survived
    m  = maintenance requirement
    V  = collateral value at inception
    L  = maximum loan consistent with surviving d

    trigger occurs when  V(1 − d) = L / (1 − m)
    so                   L = V (1 − d)(1 − m)

    worked, at V = 80.00 and m = 35%:

      survive 23.1%   L = 80.00 x 0.769231 x 0.65 = 40.000   LTV 50.0%
```

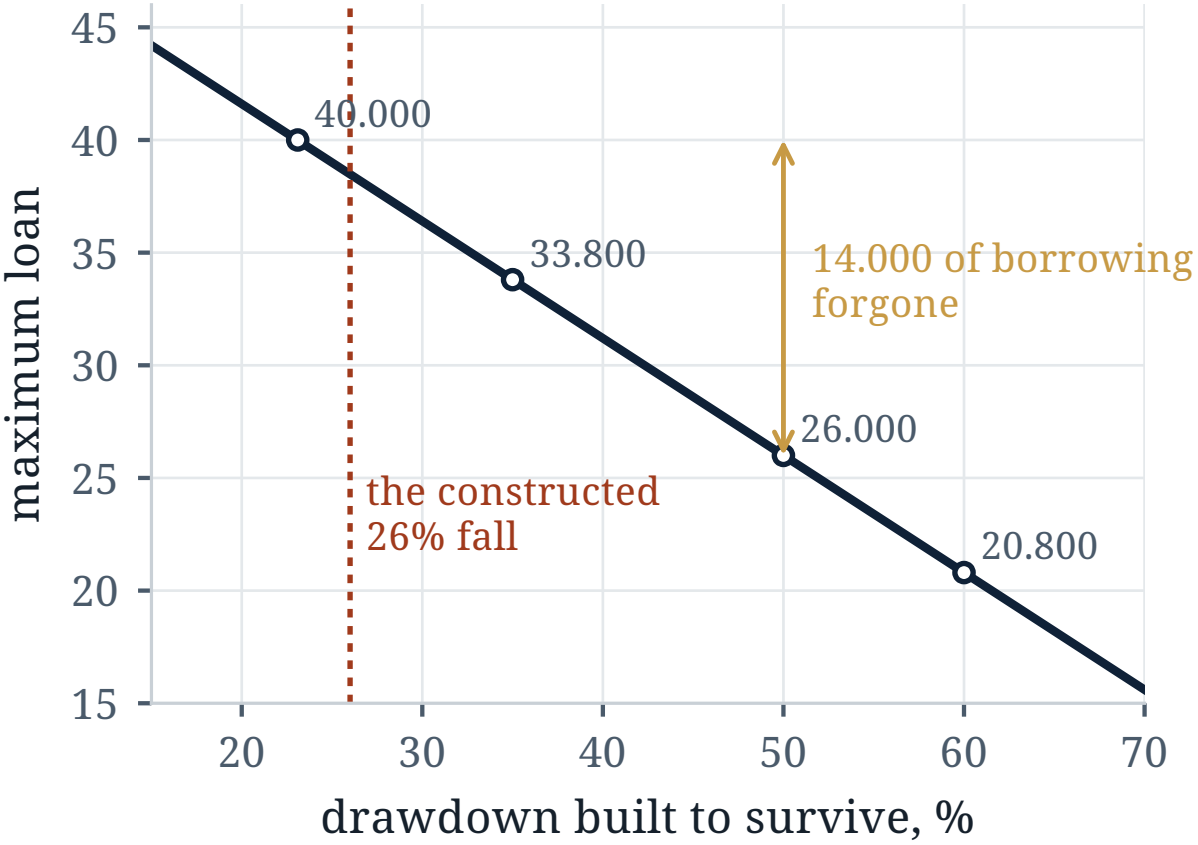


Figure 69.5: **Fixing the drawdown to be survived and solving for the loan prices the headroom directly: surviving a 50 per cent fall costs 14.000 of the 40.000 borrowing.** CONSTRUCTED parameters of §6.3.4 — collateral $V = 80.00$, maintenance requirement $m = 35$ per cent, $L = V(1-d)(1-m)$. The marked points are the four tabulated rows; the vertical rule is the 26 per cent fall of the constructed scenario in §6.5.3. The 35 per cent maintenance requirement is a constructed model parameter and not a market quotation. Script `figs/src/CCP-069_fig05.py`.

```
survive 35.0%   L = 80.00 x 0.650000 x 0.65 = 33.800   LTV 42.3%
survive 50.0%   L = 80.00 x 0.500000 x 0.65 = 26.000   LTV 32.5%
survive 60.0%   L = 80.00 x 0.400000 x 0.65 = 20.800   LTV 26.0%
```

Surviving a 50% drawdown costs 14.000 of the 40.000 borrowing — a 35% reduction in the leverage layer, which is a real cost and should be stated as one rather than smoothed away.

The point of computing it in this direction is that the drawdown is the input a family can actually have an opinion about, and the loan-to-value is not.

6.3.5 Run the concentration instrument on blocs, and record the bloc call with its evidence

The full procedure is in the concentration paper and compressed in the reference apparatus. The failure-specific instruction is narrower and has three parts.

Collapse before you measure. A naive effective count computed over jurisdictions is not a resilience number and should never be printed next to one. The Gulf-additions and bridging-slot measurements show what happens when it is.

Record the bloc assignment, its evidence, and the event that would change it. Every jurisdiction section in this framework does this at scorecard item 2 and item 12. A bloc call without

a written trigger is an opinion that will quietly become a belief.

Price the marginal jurisdiction as a probability. The the AIFC pricing form: compute the gain if the call is right, the loss if it is wrong, and the break-even probability `L / (G + L)`. If the break-even probability is low the position is cheap insurance. If it is high, the position is a bet, and the family should know it is betting.

6.4 Failure modes of failure analysis itself

The failure modes of failure analysis itself. Twelve of them, because the discipline decays in predictable ways.

FM 6.1 Reading the statute instead of the docket

The Cook Islands statute says what it says. The Andersons went to prison anyway. **Every one of the four conventional failure tracks was decided somewhere other than the jurisdiction whose statute was being relied upon**, which means that reading the protective statute tells you almost nothing about the outcome and reading the enforcement docket tells you almost everything.

FM 6.2 Survivorship bias in the adviser's record

An adviser's structures that failed are, by construction, the ones no longer in the adviser's practice. This is the selection problem that governs any record of published successes [Ioannidis, 2005, Simmons et al., 2011, Gelman and Loken, 2013], and whose measured consequence is that a large replication programme recovers far less than the original record implies [Open Science Collaboration, 2015]. The relevant question to ask is not "has this structure worked for your clients" but "which of your structures have been tested by a determined adversary, and what happened." An adviser who has never had one tested has a record with no information in it.

FM 6.3 Treating the failure record as a reason not to build

The most common misreading of this section. Every failure above is a failure of a specific design choice, and none of them is a failure of the proposition that a structure is better than no structure. The Andersons' error was the duress clause and the timing, not the existence of a trust. The substance failures were failures of entities that did nothing, not of entities that did something. Reading this section as an argument for doing nothing is reading a crash report as an argument against aeroplanes.

FM 6.4 Building against the last failure

The mirror error. Substance regimes arrived in 2018 and 2019 and the whole industry rebuilt for substance, at exactly the point when the binding constraint was moving to bankability and to bloc concentration. The sixth failure mode of the threat register states the general form: a threat that has already crystallised is the cheapest one to defend against and usually the least valuable.

FM 6.5 Confusing the disclosure test with a guarantee

The disclosure test is a filter, not a shield. It removes structures that collapse on contact. It says nothing about a policy decision taken against a disclosed and correct position, which is what happened in 2022. Any reader who has taken the disclosure paper to mean that correctness

confers immunity has misread it, and this section's track 5(c) exists to make that misreading impossible.

FM 6.6 Averaging a split trajectory

The bridging-slot control established the convention and the reference apparatus carries it. Singapore's degradation trajectory splits three ways — institutional quality stable to improving, regulatory credibility improving, bloc independence eroding — and averaging them produces "stable," which is precisely wrong on the one axis the structure is buying. **Where item 10 is genuinely non-uniform, split it and refuse to average.**

FM 6.7 Measuring at the top and not at the rail

A structure can have `N = 2.941176` on holdings and something very close to 1.00 on the ability to transact, because every leg clears dollars through the same terminal system. the rail paper's correspondent measurement makes the point with numbers: measured by terminal clearing system rather than by bank, the as-found position is 100% USD, HHI 1.0000, one effective system. **A structure that reports two rails may have one messaging dependency and one clearing dependency, and the resilience number computed on holdings will not show it** [Curzi, 2026], whose clearing-perimeter measurement takes a naive correspondent HHI of 0.302400 — 3.306878 effective correspondents — to a perimeter HHI of 0.745256 and $N_{\text{rail}} = 1.341819$ once flows are decomposed by currency.

FM 6.8 Treating a rail as a store

The store–rail rule, and it belongs in a failure analysis because the failure is expensive. Freeze risk on a transit asset is an exposure-window problem with parameters — hop size, exposure duration, counterparty pre-screening, mid-transit failure handling. Freeze risk on a store is a custody problem. Evaluating the first with the tools of the second produces a structure that is over-engineered where it does not matter and unprotected where it does.

FM 6.9 The single-instance check

Verifying one leg of a structure and generalising. A bloc call verified for one jurisdiction says nothing about the others; a bank's willingness to onboard one entity says nothing about the group; a treaty position confirmed for one income stream says nothing about the rest. This framework's own the concentration paper failure mode 11 makes the point about bankability specifically, and the general form is worth stating: **a one-instance check is not a check.**

FM 6.10 The register nobody reviews

Every jurisdiction paper closes with a Verification Register, and every register is worthless after the first year unless someone re-runs it. The twelfth failure mode of the threat register says the same about the threat register. The remedy is administrative rather than intellectual: a date, a named person, and a diff against last year's version.

FM 6.11 Believing the numbers are the position

`N = 2.941176` is a summary of an allocation the family chose, computed under bloc assignments the framework made and labelled contestable. It is not a measurement of the world. The concentration paper's robustness table exists to show how much of the answer moves under reasonable disagreement, and the Singapore case shows that the argument most likely to be had

moves `h` by 0.002014 — which is to say, the argument does not matter and the concentration does.

FM 6.12 Not stating the residual

The failure mode that this framework would most easily have committed. The succession analysis computes a forced-heirship exposure reduced to zero on the designed map and a French inheritance-tax exposure roughly four times larger that the structure does not touch. A succession section reporting only the heirship result would have been lying by omission. **The residual is the part of the answer that the structure did not solve, and a section that does not state it has not finished.**

6.5 Results: four constructed stress scenarios

CONSTRUCTED. The principal, the family, the entities, and every figure below are constructed and labelled as such. No sanctions nexus. This is the carried case of the carried constructed case, put under stress.

The carried structure as built: a DIFC Foundation at the apex; an ADGM holding company beneath it; a United Kingdom limited company; a Bermuda captive; a GIFT City vehicle; a Thai Board of Investment operating company beneath the ADGM holdco; a REIT position converting London commercial real estate from immovable to movable; and single-asset special purpose vehicles for ring-fencing. The measured position: six jurisdictions, naive effective count 5.290949, effective blocs 2.941176, `h = \ensuremath{-}0.122417` at `eta = 0.35`, `C = 0.30`, `Q = 10`, `delta = 0.02`.

Four stress scenarios, run as pre-mortems.

6.5.1 Scenario A: a judgment creditor in a Western court, five years after building

The scenario, in the past tense. A commercial dispute arising from the United Kingdom operating business produced an English judgment against the principal personally in a sum exceeding the United Kingdom company's net assets. The creditor sought disclosure of the principal's worldwide assets and an order requiring the principal to procure a distribution from the Foundation.

The first domino. The principal was resident in the UAE but travelled to London regularly and had submitted to English jurisdiction in the underlying contract. Personal jurisdiction was never in doubt.

What the structure did. The Foundation is a legal person and the assets are its own; there is no beneficial interest for the creditor to attach. The creditor therefore attacked in the only way track one permits: an order against the principal to exercise a retained power. The Step 3 audit, run at build time, had identified three retained powers: the power to appoint and remove councillors, the power to amend the charter within stated limits, and a consent right over distributions. The first two were retained. **The third was removed at build time precisely because its compelled exercise would have given a creditor the result it wanted.** The council retained a genuine discretion over distributions and included two councillors with no relationship to the principal.

What the structure did not do. It did not prevent the disclosure order, the litigation cost, the two years of management distraction, or the reputational consequence in the group's banking relationships, which was the largest single cost of the episode. It did not prevent enforcement against the United Kingdom company's own assets, which were fully exposed and were always intended to be.

The residual, stated. Timing did the work, not jurisdiction. Had the structure been built after the dispute arose, the transfers into it would have been the target and the analysis would have been a fraudulent-transfer analysis under whichever regime applied, with a look-back

longer than the structure's life. **The Foundation was not the reason this scenario was survivable. Building seven years before the dispute was.**

6.5.2 Scenario B: the substance challenge to the intermediate holding layer

The scenario. A revenue authority in the jurisdiction of one operating subsidiary asserted that the ADGM holding company was resident in its territory on a central-management-and-control analysis, denied treaty benefits on dividends, and assessed withholding tax with penalties for three prior years.

The first domino. Board meetings of the ADGM holdco had, for two of the three years, been held wherever the principal happened to be, and the minutes recorded it accurately. The minutes were truthful and they were the evidence.

The measurement. The exposure is a cost rather than a loss of the structure, and it should be sized rather than dramatised. Denied treaty relief on a dividend stream, plus penalties and interest, against the cumulative maintenance cost of running the ADGM layer properly for the same period. In the constructed case the assessment exceeded eight years of correctly-run ADGM maintenance cost, which is the honest framing: **the substance was not expensive. Not having it was.**

What would have prevented it. The three-jobs rule and nothing more exotic. The ADGM holdco scores a job vector of (0.70, 0.05, 0.25), specialisation index 0.5550, 1.8018 effective jobs — inside the 2.20 threshold and therefore not flagged as overloaded. The vector was not the problem. The problem was that the entity's decisions were taken somewhere the vector did not describe, which the sixth failure mode of the three-jobs diagnostic names exactly and which no vector can detect, because a job vector measures what an entity does and residence tests measure where a human being was sitting.

The residual. Fixed prospectively, at cost, with three years of exposure that no structural change can retrieve.

6.5.3 Scenario C: the margin call that nobody sued over

The scenario. The listed carve-out from the Thai operating business, floated at 20% with the retained 80% pledged at 50% loan-to-value, fell 26% over four months on a sector de-rating with no company-specific news.

The first domino. 23.1%, computed above, and reached in month three.

What happened. The maintenance call arrived with a two-business-day cure period. The Foundation held the listed portfolio, but a substantial part of it was inside the GIFT City vehicle and the Bermuda captive's reserve pool, which is held against real third-party liabilities and is not available — that unavailability being the whole reason the captive-reserve analysis treats reinsurance reserves as a structurally hard-to-seize class. **The property that was hardest for an adversary to reach was also hardest for the family to reach, and that symmetry is not a defect in the design; it is the design, and it should be understood before the pledge rather than during the cure period.**

The outcome, constructed. Partial repayment of 8.4 units from liquid holdings, avoiding a forced sale, at the cost of the liquidity buffer the succession analysis in the succession analysis had already identified as short. Recall that analysis [Curzi, 2026]: liquidity rather than entitlement binds, and a constructed USD 60.00 million claim against USD 26.16 million of liquid assets leaves a shortfall of USD 33.84 million, 14.1 per cent of the estate and larger than any single liquid holding. **The margin call and the succession liquidity problem are the same problem arriving through two different doors, and the structure has one buffer for both.**

What would have prevented it. Step 4, run in the correct direction. At a 50% drawdown tolerance the borrowing falls from 40.000 to 26.000 and the whole scenario becomes a non-event. The 14.000 of foregone borrowing is the price of the non-event, and it is knowable in advance.

6.5.4 Scenario D: the diversification that was not

The scenario. Over three years, and with no single decision that looked wrong, the family added a Qatar Financial Centre family-office vehicle and a Saudi regional headquarters, funding both by trimming the UAE, United Kingdom, and Singapore allocations. Each addition was recommended on its own merits by competent advisers. Each was fully disclosed, fully compliant, and passed the disclosure test.

The first domino. No adviser was asked to compute the effect on the effective bloc count, because no adviser had been given the effective bloc count as a constraint.

The measurement, recomputed. Naive effective jurisdictions rose from 5.290949 to 7.132973, an increase of 34.81%. Effective blocs fell from 2.941176 to 2.841490, a decrease of 3.39%. h fell from -0.122417 to -0.134485. **The family paid three years of additional maintenance cost to become less resilient, and every report it received said the opposite.**

The warning sign that was available. One line of arithmetic, run at each addition, on the bloc shares rather than the jurisdiction shares. The concentration instrument would have flagged the first Gulf addition and refused the second.

The correction. The Gulf-additions paper's own worked outcome. Keep the Saudi regional headquarters at 5.0%, fund it from the operating allocation rather than the holding allocation, and the position becomes naive 5.688217 and blocs 2.816901 — a loss of 0.124275 in N and 0.015110 in h, taken deliberately for addressable revenue, with the number recorded in the file. Drop the Qatar vehicle, which duplicates the DIFC inside the same bloc and buys nothing.

The residual, and it is the section's closing point. Scenario D is the only one of the four in which nothing went wrong. No court, no revenue authority, no lender, no adversary at all. Every entity was correct, disclosed, and defensible. **The structure failed the only test that mattered for the threat it was built against, and it failed silently, and the instrument that would have caught it takes about a minute to run.**

---

7. Results

Every figure below is exact arithmetic on stipulated, constructed inputs, and none is a measurement of any real market, reserve, structure or economy.

R1 — disclosure is not immunity. Approximately USD 300 billion of fully disclosed central bank reserves were immobilised in 2022, so the disclosure test is a **filter** that removes structures collapsing on contact with an authority and is **not a shield** against a policy decision taken on a correct and fully disclosed position. **The event is described from the source manuscript and is not cited** (§1, Appendix A item 20). The symmetric half of the result is what makes it structural: concealment could not have prevented the outcome either, because a reserve must be held by a named holder in a named place to be a reserve, so the reachability is a property of the custody rather than of what was known about the owner. A fully disclosed, fully correct, single-perimeter structure has an effective bloc count of 1.00 and therefore $\ln(1.00) = 0$ of available surplus [Curzi, 2026].

R2 — the correlation collapse in a backing basket. On the constructed eight-component basket, naive HHI 0.150600 gives 6.640106 effective components; recognising the climate-correlated components as a single 82.0 per cent exposure gives bloc-HHI 0.693600 and 1.441753, a naive overstatement factor of 4.605578. **The basket is more diversified than a single-commodity anchor at 1.44 against 1.00 and far less diversified than its component count claims.** The correlated spoilage of the storable components runs at a constructed 13.67 per cent a year against a non-perishable metal's 0.00, which is the anti-hoarding property expressed as a physical carrying rate rather than a price.

R3 — the delivery squeeze. At 250 claims per physical unit on 40,000 deliverable units, $10{,}000{,}000$ claims are outstanding, $9{,}960{,}000$ exceed one-to-one collateral, and 99.600 per cent of

claims are uncollateralised at par. 0.40 **per cent of holders standing for delivery exhausts the stock exactly**, and coverage falls to 0.4000 at one per cent and 0.0800 at five. **A fully collateralised object-ledger makes the ratio a live on-ledger count rather than a reconciliation across two registries that nobody runs, and it does not prevent over-issuance**, because per-object hash chains supply no consensus, no total order and no double-spend prevention on their own [Curzi, 2026, Lamport et al., 1982, Fischer et al., 1985].

R4 — the gated integration priced as an option. Constructed $p_1 = 0.60$, $p_2 = 0.50$, $p_3 = 0.30$ give $P(\text{reach phase } 2) = 0.600$, $P(\text{reach phase } 3) = 0.300$ and $P(\text{terminal}) = 0.090$, a terminal-to-first-phase ratio of 0.15. **A programme that leads with its terminal phase stakes the certain value of the standalone reform on the least probable branch.**

R5 — the apparent-diversification trap. Adding two Gulf jurisdictions funded by trimming three existing allocations raises naive effective jurisdictions from 5.290949 to 7.132973, a rise of 34.81 per cent, while effective blocs fall from 2.941176 to 2.841490, a fall of 3.39 per cent, and h falls from -0.122417 to -0.134485 [Curzi, 2026]. **Three years of additional maintenance cost bought less resilience, and every conventional report said the opposite.**

R6 — the margin call with no plaintiff. A 26 per cent fall over four months breaches a 23.1 per cent trigger in month three, on a pledged position at 50 per cent loan-to-value with a 35 per cent maintenance requirement [Curzi, 2026]. Solving the same relation for the loan instead of the trigger, **buying survival to a 50 per cent drawdown costs 14.000 of the 40.000 borrowing — a 35 per cent reduction in the leverage layer.** No adversary, no court and no finding of wrongdoing is required; the lender exercises rights the holder granted.

R7 — the unpriced bloc call. A wrong perimeter call does not merely fail to help: it costs 0.005110 in h, because capital was taken out of a bloc where it was diversifying and put into one it concentrated. The break-even probability follows, $p^* = 0.005110/(0.026731 + 0.005110) = 0.160476$ [Curzi, 2026], and the binding constraint turns out to be the leg's maintenance cost rather than the geopolitics.

CR — the counter-result, which is the whole paper. The framework's governing principle did not confer immunity in the strongest available test case; its own backing basket is 4.61 times less diversified than its component count claims; the terminal phase of its own programme is a 0.090 option rather than a plan; and a levy presented as a costless accounting correction in fact reallocates burden, with **42.0 per cent of a constructed incidence falling on one sector and a constructed 55 per cent passing through to consumers**, across an effective 3.631082 sectors.

8. Discussion

8.1 What follows

A filter that is not a shield is still worth having, and saying so precisely is the whole of the repair. The 2022 event removes one reading of the disclosure principle and leaves another intact. What it removes is the reading in which correctness confers immunity. What it leaves is a structure that can be built in the open, examined by any authority without collapsing, and — decisively — **banked**, which the corpus identifies as the constraint binding everything else. A structure that cannot be banked has no rails, and a structure with no rails has no diversification to measure.

The event is an argument against concentration rather than against disclosure, and the two are easy to confuse. Its lesson is that the variable is *where the holding sits*, not *what is known about the holder*, and the corpus's answer is accordingly jurisdictional and rail diversity across perimeters. The measurement that makes the point is the corpus's own [Curzi, 2026]: holding marginal failure at 10 per cent, a common shock of $c = 0.02$ raises the probability that four nominally independent perimeters fail together from 0.000100 to 0.020044, a 200.4-fold understatement, and over fifteen years 0.001499 independent against 0.794109 correlated.

One instrument recurs on four different axes and the recurrence is the paper's structural finding. A count of nominally independent things collapses under a common shock: jurisdictions onto coordination perimeters; maturities onto a refinancing window [Curzi, 2026]; correspondents onto a clearing perimeter [Curzi, 2026]; and — here — basket components onto a climate exposure, and paper claims onto a deliverable stock. **The naive count is never the perimeter, and the same arithmetic detects the difference every time.**

8.2 What does not follow

No structure described here is recommended, and no instrument is claimed to exist. The three backing mechanisms are designs graded [T3]. The phased programme has no phase adopted anywhere. The licensing provisions are recorded with their exposure and no outcome is stated.

The failure record is not an argument against building. Every failure in §6 is a failure of a specific design choice — a duress clause, a timing decision, an entity that did nothing where it was registered, a pledge granted without computing the headroom, a jurisdiction added without computing its bloc effect. **Reading §6 as an argument for doing nothing is reading a crash report as an argument against aircraft**, and §6.4 lists that reading as a named failure mode.

Nothing here measures truthfulness or lawfulness. The concentration instrument measures the independence of a partition. It does not know whether a bloc assignment is right, whether an entity does what it says, or whether a counterparty is honest, and §3.2 states the two things it cannot see at all.

And the corpus's own conjecture is not established by any of it. The framework asserts that a structure's resilience is a computable function of its perimeter geometry. That is a modelling position, and the corpus's foundational papers already refuse the inflated reading of their own borrowed instruments: a life is not a metric space, nobody has measured the relevant contraction modulus, and a theorem transfers as shape rather than as guarantee [Curzi, 2026]; the honesty appendix assigns the corpus's central identification and its application of the surplus functional to the postulate drawer, and calls the third of them a wager whose misfire column exists because wagers can lose [Curzi, 2026]. **This paper is that misfire column, filled in.**

8.3 The correlation grouping is the load-bearing assumption, and it is stipulated

R2's headline number is a function of one modelling choice: that grains, freshwater, fruit, soil-health, biodiversity and dairy feed constitute a single climate exposure. **That grouping is stipulated, not estimated, and a different grouping gives a different number.** The general point is robust — correlations rise in the tail [Longin and Solnik, 2001, Ang and Chen, 2002], diversification is a claim about a covariance structure rather than a count [Markowitz, 1952, Elton and Gruber, 1977, Statman, 1987], and more components do not in general give more stability [May, 1972, McCann, 2000, Ives and Carpenter, 2007, Sagan, 1993] — but the *magnitude* is a property of the construction. **A real basket requires an estimated covariance structure and none was estimated here.**

The same caution applies with more force to the bloc assignments behind R5 and R7. They are the framework's own, they are contestable, and the corpus measures the index's sensitivity to assignment error separately: reassigning one member moves the effective bloc count by 0.0169, under 0.6 per cent [Curzi, 2026]. **The index is insensitive to assignment error and highly sensitive to bloc structure**, which means the argument a reader is most likely to have is not the one that matters.

8.4 The proposals, graded against themselves

The paper-to-physical mechanism is the strongest of the three and its object-ledger remedy is the most oversold. Visibility is not prevention, and the corpus's own integrity-layer paper published a working double-spend against its own substrate rather than hiding it [Curzi, 2026]. That publication is the reason the remedy can be graded at all.

The consumables basket is a differently-shaped reserve and not an immunity. The Act's stated goal of mathematical immunity to inflation and deflation is recorded and corrected wherever it appears; the channel-gain postulate is the corpus's general form of the correction [Curzi, 2026ab], and the corpus's own concordance discipline is the reason a design goal and an achieved property are kept in separate columns [Curzi, 2026w], whose worked rejection of a widely repeated frequency claim on three independently computed grounds is the model for what a refusal in print looks like.

The easement reserve is a store and the voucher is a rail, and the two-layer design is the honest way to make an illiquid reserve spendable. Its weakest number is the reserve valuation, which is contested and has a documented abuse history; the general governance frame is available [Hardin, 1968, Ostrom, 1990, 2010, Cox et al., 2010] and **the valuation literature is not in this paper's bibliography and is cited nowhere.**

And the nine-form backing taxonomy fails the scale-type gate outright. Its non-financial components are not instrumented on any common scale, let alone a ratio scale with a shared true zero, so the magnitude of any coverage ratio computed across them is withdrawn as a measured quantity — which is exactly the verdict the corpus reached on its own moment equation [Curzi, 2026t], and which the nine-capital paper reaches on its own borrowing-base ceiling by bounding the honest gain strictly between zero and 56.62 per cent rather than asserting it [Curzi, 2026].

8.5 Bounds carried from the corpus

The most restrictive bound on this whole exercise comes from the assembled-architecture paper and is accepted here without qualification. The only reading under which the corpus's own architecture clears its build threshold is the naive one assuming zero correlation between coordinating jurisdictions, giving $h = +0.083099$; **every defensible bloc assignment returns a negative surplus, and that paper says so in print** [Curzi, 2026]. This paper's constructed structures are in the same position, and R5's $h = -0.134485$ is a worse number than the one it replaced.

Three further bounds are carried rather than restated. A framework must be able to distinguish a theorem, an ordinary effect and its own conjecture in three separate sentences on demand [Curzi, 2026]; §9.6 discharges that requirement for every borrowing here. An audit that rejects nothing has not audited, and a many-to-one rule cannot be inverted to recover what produced it [Curzi, 2026] — which is the general form of §3.2's warning that an effective count does not determine the allocation it was computed from. And repetition buys arrival time and nothing else: it cannot move a destination [Curzi, 2026], which is why §5's gates are demonstrations rather than schedules, and why a phase asserted more loudly is not a phase reached.

The corpus's sharpest single instance of the discipline is the one this paper takes as its model. The cellular-protocol paper builds a complete, verified, deterministic pipeline and then states that nobody knows how a cell could read an ordered message out of a sound: eleven components established, one link missing, and the missing link named rather than softened [Curzi, 2026]. **That is the shape of this paper: the mechanisms are built and the governing principle is holed, and both are in print.**

8.6 What would falsify this paper

Five things, none tested here.

If a defensible estimated covariance structure over a real backing basket produced a bloc-collapsed effective count close to its naive count, R2's mechanism would survive and its magnitude would not.

If a survey of real commodity markets found paper-to-physical ratios close to one, the squeeze mechanism would remain arithmetically true and operationally empty. **That survey has not been done, here or, so far as this paper is aware, anywhere.**

If a dependent phased programme were exhibited in which a later phase's standalone value exceeded the first's, the gating argument would fail — although it is hard to see how, given that the first phase is unconditional by construction.

If a bloc assignment defensible on the framework's own perimeter test made the Gulf additions of R5 independent rather than concentrating, the apparent-diversification trap would be an artefact of the assignment rather than a finding. **The assignment is the contestable input and the paper says so twice.**

And if a fully disclosed, correct position had in fact been protected in 2022 by the completeness of its disclosure, H1 would stand. It was not.

9. Verification and rigour audit

Each result is decomposed into postulates that can be checked one at a time. Each is atomic, each carries a status, and each chain names what it consumes. **If dropping a postulate leaves the result standing, it is not listed.**

9.1 Dimensional audit

symbol	quantity	dimension	unit	status
shares	allocation weights	$[1]$	fractions summing to 1	stipulated
HHI	sum of squared shares	$[1]$	dimensionless	derived
$1/\mathrm{HHI}$	effective count	$[\#]$	components, jurisdictions, blocs	derived; a numbers-equivalent
h	surplus index	$[1]$	index units, **no currency**	derived from stipulated η, δ, Q, C
claims, units	counts	$[\#]$	claims; deliverable units	stipulated
R	paper-to-physical ratio	$[1]$	claims per unit	stipulated
coverage	units ÷ claims presented	$[1]$	dimensionless	derived
p_1, p_2, p_3	gate probabilities	$[1]$	dimensionless	stipulated
V, L	collateral, loan	$[1]$	**stipulated units, no currency**	stipulated
m, d, d^*	maintenance, drawdown	$[1]$	fractions	stipulated / derived
spoilage	perishability rate	$[T^{-1}]$	per year	stipulated
reserve, vouchers	natural-capital value	$[1]$	**abstract units, no currency**	stipulated

symbol	quantity	dimension	unit	status
USD 300 billion	immobilised reserves	[money]	US dollars	**described, not measured here**

Does it balance? Every headline figure except one is a ratio of two counts of the same kind, or a product of probabilities, and is therefore dimensionless and well formed: 6.640106 and 1.441753 are reciprocals of sums of squared shares; 0.090 is a product of three probabilities; 0.40 per cent is $1/250$; 23.076923 per cent is $1 - 0.50/0.65$. **The one dimensional quantity in the paper is the USD 300 billion aggregate, and it is not this paper's measurement.**

Three quantities carry no operational unit and the audit says so rather than assigning one. h is in index units with no currency attached and is meaningful only as a difference or a sign. The easement reserve and its vouchers are in abstract units of natural-capital value, and **no measurement basis for those units exists in the source corpus or here** (§8.4). The collateral and loan figures of §6.3.4 are stipulated units, so the 14.000 of forgone borrowing is a ratio statement — 35 per cent of the leverage layer — and not an amount.

9.2 Significant figures and epistemic status

- **6.640106, 1.441753, 4.605578, 5.290949, 7.132973, 2.941176, 2.841490 are exact to the digits shown**, being reciprocals of sums of squared exact stipulated shares. They carry no interval because nothing was estimated.
- **0.600, 0.300, 0.090 and the ratio 0.15 are exact products of stipulated probabilities**, and the probabilities are illustrative. Quoting them to three places is a statement about the arithmetic, not about any likelihood.
- **23.076923 per cent and 0.40 per cent are exact closed-form solutions** on stipulated parameters.
- **"approximately USD 300 billion" is stated as approximate because the source states it so**, and no result in this paper depends on the figure's precision. **A single significant figure would be the honest presentation of a number this paper did not measure**, and the aggregate is retained in the source's own form only because the narrowing it forces does not turn on its size.
- **13.67 per cent, 42.0 per cent, 55 per cent, 3.631082 and 0.275400 are exact on constructed inputs and are illustrations of a shape.** A real incidence analysis requires real sectoral and elasticity data and none was used.

9.3 Limits and boundary behaviour

The squeeze relation, at its boundaries. Coverage $= 1/(Rs)$ diverges as the stand-for-delivery fraction $s \to 0$ — which is the ordinary state of a hedging market and is why the ratio is invisible in normal function — and falls to $1/R$ at $s = 1$. At $R = 1$ the market is fully collateralised and no fraction can squeeze it. **The whole content of the mechanism is that the interesting region is a fraction of one per cent of holders**, and the relation is examined across four orders of magnitude of s.

The correlation collapse, at its boundaries. With every component independent the collapsed count equals the naive count and the overstatement factor is 1.000000; with every component in one bloc the collapsed count is exactly 1.000000 and the basket is a single-commodity anchor by another name. **The constructed basket sits at 1.441753, close to the degenerate end**, and the paper says so rather than reporting the naive figure alone.

The gating relation, at its boundaries. At $p_i = 1$ for all i the terminal probability is 1 and the phasing argument is empty; at any $p_i = 0$ the programme stops at that gate and the standalone phase's value is unaffected, which is the property the phasing exists to secure. **The argument needs only** $0 < p_i < 1$**.**

The margin relation, at its boundaries. $d^* = 1 - \mathrm{LTV}_0/(1-m)$ gives $d^* = 1$ at zero leverage — no drawdown triggers a call — and $d^* = 0$ when $\mathrm{LTV}_0 = 1 - m$, where the position is at its call level on the day it is opened. **The dangerous region is the ordinary one**, and at a 50 per cent loan-to-value against a 35 per cent maintenance requirement the trigger is a 23.1 per cent fall, which is unremarkable equity volatility.

9.4 Invariance

The effective count is invariant under relabelling of the components and under any permutation of the shares, because it is a function of the multiset of shares alone. That is what makes it a numbers-equivalent, and it is also exactly what §3.2 warns about: **a count that does not depend on which component holds which share cannot tell a reader which component to change.**

It is not invariant under regrouping, and the whole of R2, R5 and R7 lives in that fact. The transformation that moves the number is the partition, and the partition is a judgement. **The corpus measures the size of the judgement's effect: reassigning one member moves the effective bloc count by 0.0169, under 0.6 per cent** [Curzi, 2026], so the index is robust to being wrong about one member and not robust to being wrong about the bloc structure.

The squeeze threshold is invariant under the size of the market: doubling both the claims and the units leaves $1/R$ unchanged, which is why the finding is a statement about the ratio and not about any market's size.

9.5 Measurement-theoretic status

The tier scheme is ordinal, and no arithmetic is performed on it (§3.1) [Stevens, 1946]. **The bloc assignment is nominal**, and the effective count computed over it inherits the assignment's status entirely: it is a ratio-scale statistic over a nominal partition, which is legitimate and is exactly as defensible as the partition.

The nine-form backing taxonomy fails the scale-type gate, as §8.4 states: no measurement basis, no common ratio scale, no shared true zero, and therefore no defensible coverage ratio [Curzi, 2026t]. The formal remedies — conjoint measurement or a Rasch demonstration [Luce and Tukey, 1964, Krantz et al., 1971, Rasch, 1960] — were **not run**, and Michell's objection that quantitative structure is assumed rather than tested stands unanswered [Michell, 1997, 1999].

The gate probabilities are subjective and are labelled so. They are not frequencies, not estimates and not elicited from anyone; they are stipulated numbers chosen to make an arithmetic point, and the point survives their replacement by any other set strictly inside $(0, 1)$.

9.6 Analogy discipline: every borrowing, graded

The corpus requires that a reader be able to distinguish a theorem, an ordinary effect and the framework's own conjecture in three separate sentences on demand [Curzi, 2026]. **This section discharges that requirement for every borrowing in this paper.** Three categories are used: **(1) a derivation**, where the system genuinely satisfies the cited result's premises; **(2) a shared mathematical form**, where the same equation governs both with no transfer of authority; and **(3) an analogy**, legitimate as motivation and never as evidence.

borrowing	used for	category
Concentration indices and their numbers-equivalent [Hirschman, 1964, Herfindahl, 1950, Adelman, 1969]	every effective count in the paper	**1 — derivation.** The shares are a partition and the index is defined on it
Hill numbers [Hill, 1973, Jost, 2006]	naming $1/\mathrm{HHI}$ as the order-two member	**1 — derivation**
Portfolio covariance [Markowitz, 1952, Elton and Gruber, 1977]	why a count is not a diversification	**2 — shared form.** No portfolio was constructed and no return appears

borrowing	used for	category
Tail correlation [Longin and Solnik, 2001, Ang and Chen, 2002]	why the collapse happens when it matters	**3 — analogy.** Measured in equity markets, asserted here of a basket
Ecological stability–diversity results [May, 1972, Tilman et al., 1998, Doak et al., 1998, McCann, 2000, Ives and Carpenter, 2007, Schindler et al., 2010]	that more components need not mean more stability	**3 — analogy**
Coupled-network fragility [Buldyrev et al., 2010, Albert et al., 2000, Cohen et al., 2000]	how a defect travels through composed identifiers and blocs	**3 — analogy**, invoked for shape
Byzantine agreement and its impossibility [Lamport et al., 1982, Fischer et al., 1985, Gilbert and Lynch, 2002, Douceur, 2002]	why an object-ledger prevents nothing on its own	**1 — derivation.** The substrate is literally a set of mutually distrustful parties needing agreement
Open design [Kerckhoffs, 1883, Saltzer and Schroeder, 1975]	the disclosure test	**1 — derivation** as a design principle
Transparency as remedy [Brandeis, 1914]	the same test in its policy register	**3 — analogy**
Tight coupling and system accidents [Perrow, 1984, Sagan, 1993, Leveson, 2011, Vaughan, 1996]	why the pre-mortem is mandatory and why redundancy is not automatically a remedy	**3 — analogy**
Commons governance [Hardin, 1968, Ostrom, 1990, 2010, Cox et al., 2010]	the framing of a stewarded natural-capital reserve	**3 — analogy**
Fibonacci and the golden section [Euclid, 1956, Vajda, 1989, Koshy, 2001, Binet, 1843]	that consecutive ratios approach φ	**1 — derivation**, and **the monetary significance of that fact is graded at zero** [Markowsky, 1992, Fischler, 1981, Livio, 2002]
The grammar of conserved quantities [Mirowski, 1989, Cartwright, 1999]	flagged against the framework itself, not borrowed	**not a borrowing — a warning**

The three sentences, stated as the corpus requires them. *The theorem:* no protocol among mutually distrustful parties achieves agreement without an ordering layer, so a per-object ledger does not prevent a double-spend — a derivation, whose instance the corpus reproduced against its own substrate. *The ordinary effect:* a lender exercises contractual rights the holder granted when collateral falls below a maintenance level, which requires no adversary and no theory. *The conjecture:* that a structure's resilience against arbitrary action is a computable function of its perimeter geometry is **this framework's own proposal**, is not derived from anything cited, and is not established by any result in this paper.

9.7 The postulates

	postulate	status
P1	The stipulated share vectors, ratios, probabilities and contract parameters are as printed.	DEFINITION (constructed)
P2	A bloc is a set of jurisdictions that coordinate, so their shares add before squaring.	POSTULATED — the framework's own perimeter test; contestable
P3	The climate-correlated components of the basket constitute one exposure.	MODELLING CHOICE — stipulated, not estimated

	postulate	status
P4	$1/\text{HHI}$ is the right reading of an effective count.	CITED [Hirschman, 1964, Hill, 1973, Adelman, 1969]
P5	Approximately USD 300 billion of disclosed reserves were immobilised in 2022.	**DESCRIBED — uncited; no source in this bibliography**
P6	Per-object hash chains supply no consensus and no double-spend prevention.	CITED [Curzi, 2026, Lamport et al., 1982]
P7	Each phase gates the next and the gates are independent.	MODELLING CHOICE
P8	The named statutes, cases and administrative authorities are as the source describes them.	**DESCRIBED — uncited; Appendix A**

Dropped as unnecessary. No postulate about which commodity, which basket components, which jurisdictions or which currency is needed for any result: R2, R3, R4, R5 and R6 are properties of partitions, ratios, probability products and a linear equation, and each survives the substitution of any other labels.

9.8 The chains

step	consumes	produces	status
A1	P1	naive HHI $= \sum w_i^2 =$ 0.150600	ARITHMETIC
A2	A1, P4	naive effective components $=$ 6.640106	ARITHMETIC
A3	P1, P3	bloc shares $82.0/14.0/4.0$; bloc HHI $= 0.693600$	ARITHMETIC
A4	A3, P4	correlated effective components $=$ 1.441753	ARITHMETIC
A5	A2, A4	overstatement factor $= 4.605578$ — **R2**	ARITHMETIC
B1	P1	claims $= 250 \times$ $40{,}000 = 10{,}000{,}000$	ARITHMETIC
B2	B1	uncollateralised at par $= 9{,}960{,}000$, 99.600%	ARITHMETIC
B3	B1	coverage $= 1/(Rs)$; $=$ 1 at $s = 1/250 = 0.40\%$ — **R3**	ARITHMETIC

step	consumes	produces	status
B4	B3, P6	the ledger reports the ratio and does not bound it	CITED
C1	P1, P7	0.600, 0.300, 0.090	ARITHMETIC
C2	C1	terminal-to-first ratio $= 0.15$ — **R4**	ARITHMETIC
D1	P1, P2	naive $5.290949 \to 7.132973$; blocs $2.941176 \to 2.841490$ — **R5**	ARITHMETIC
D2	D1	h: $-0.122417 \to -0.134485$	ARITHMETIC
E1	P1	$d^* = 1 - \text{LTV}_0/(1-m)$ $= 23.076923\%$ — **R6**	ARITHMETIC
E2	E1	$L = V(1-d)(1-m)$: $40.000 \to 26.000$ at $d = 50\%$	ARITHMETIC
F1	P2	right call $+0.026731$; wrong call -0.005110	ARITHMETIC
F2	F1	$p^* = 0.005110/0.031841$ $= 0.160476$ — **R7**	ARITHMETIC
G1	**P5**	a fully disclosed position was immobilised — **R1**	**DESCRIBED, UNCITED**
G2	G1	the disclosure principle is a filter and not a shield	DEDUCTION from a single instance

9.9 What these chains are and are not

Read the status column. Every chain except G is arithmetic on stipulated inputs, and therefore establishes a property of the construction and nothing about the world. **Chain G is the only one that touches an event, it rests on a single described instance, and its source is not in this paper's bibliography.**

A single counterexample is the right instrument for the claim being tested and the wrong instrument for anything else. H1 was a universal — a disclosed and correct position is durable — and one instance refutes a universal. It does not establish a rate, a probability, or a prediction about any other position, and no such inference is drawn. **G2 is a deduction from a refuted universal, not an induction from a sample of one.**

And the load-bearing input to R2, R5 and R7 is a partition chosen by this framework. P2 and P3 are the postulates a reviewer should attack first, and §9.4 gives the sensitivity: the index is robust to being wrong about one member and not robust to being wrong about the structure.

10. Limitations

Every worked result is a constructed model and none is a measurement. The commodity, the basket, the reserve, the vouchers, the incidence, the probabilities, the allocations, the pledge and the four stress scenarios are invented, with parameters chosen to demonstrate a mechanism rather than to describe anything. **No external dataset was collected. No real market,**

commodity, reserve, land parcel, economy, institution, court file, family or principal appears anywhere in this paper.

The 2022 immobilisation is described and not cited, and that is the largest single gap. The figure, the date, the instruments by jurisdiction, the amounts by jurisdiction, and the current legal position on the principal and on the income arising from it are **taken from the source manuscript and attached to no source**, because this paper's bibliographic database contains no contemporaneous reporting and no sanctions instruments. **No approximate citation was substituted.** Appendix A item 20 states what must be verified. A reader relying on the aggregate must verify it independently; a reader relying on the *narrowing* it forces need not, because the narrowing follows from one instance of the class and the class is not in doubt.

No jurisdiction-specific legal source is cited anywhere. The decided cases of §6.2.1 and §6.2.3, the economic-substance legislation of §6.2.2, the tax-code conservation provisions of §4.2.4, the administrative designation of syndicated easements, the nominee provisions of the foreign-business statute, the share-pledge disclosure requirement, and the procurement, appropriations, conflict-of-interest and constitutional frameworks of §5.2.5 are **all described from the source manuscript and left uncited**, and each is itemised in Appendix A with what requires verification. The general doctrinal frame is cited where it is relied on [Hague Conference on Private International Law, 1985, United Nations, 1958, Torremans et al., 2022, Briggs, 2013, Hansmann and Kraakman, 2000, Pistor, 2019]. **This is a real gap and it is not closed here.**

The commodity multiples are a range and are asserted nowhere. The circulating figures depend entirely on the choice of numerator and denominator, and the paper claims the shape rather than any magnitude. **The survey that would establish whether real paper-to-physical ratios are large enough for the squeeze mechanism to matter has not been done.**

The correlation grouping and the bloc assignments are stipulated judgements, and they are what the headline numbers are made of (§8.3, §9.9). A real basket requires an estimated covariance structure; a real perimeter test requires evidence per jurisdiction with a written trigger for revision.

The gate probabilities are estimates of nothing. They are illustrative numbers chosen to make an arithmetic point about dependence, and quoting the terminal 0.090 as a likelihood would be a misreading the paper cannot prevent except by saying so.

The easement reserve's valuation is the weakest number in the paper and it is not computed. Easement valuation is contested and has a documented abuse history; the ecosystem-services and natural-capital valuation literature is **not in this bibliography and is cited nowhere**; and no valuation figure is stated anywhere in this paper.

The incidence model is a shape and not a measurement. The sectoral shares, the 55 per cent consumer pass-through and the resulting 3.631082 effective sectors are constructed. A real incidence analysis requires real sectoral data and real elasticities, and **a levy's incidence is not determined by who remits it**.

The concentration instrument does not see a pledge and does not see capacity (§3.2). Both blind spots are load-bearing in this paper: §6.2.4's failure track is invisible to the instrument entirely, and an effective bloc count of 2.94 says nothing about whether any of the three perimeters could carry the structure.

No statistical claim is made and none is supported. There are no p-values, effect sizes, intervals, samples or fitted models in this paper. No model comparison, out-of-sample test, matched null or information criterion is reported, because nothing was fitted, and no distributional or scaling claim is made, so Clauset et al. [2009] is not applicable and was not run.

The author has a direct financial interest in one of the provisions recorded in §5.2.4. That is declared in §3.4 and in the Declarations. It cannot be cured by disclosure, and no reader should treat the neutral register in which the provisions are recorded as neutrality about the interest.

And this paper's own principle is holed rather than repaired. H1 is falsified. What replaces it — a filter that is not a shield — is a weaker claim, is not itself tested, and is offered as a modelling position rather than as an established result.

11. Conclusion

Disclosure is not immunity. Approximately USD 300 billion of central bank reserves, among the most thoroughly disclosed assets in existence, were immobilised in 2022 by the jurisdictions in which they were held. Their location, their custodians and their character were known, and were the reason they were held where they were held. **Correctness and disclosure were complete, and they were immobilised anyway.** The result is symmetric: concealment could not have prevented it either, because a holder must be named and a holding must sit somewhere in order to be held at all. **What disclosure does is remove the failure mode in which a structure collapses on contact with an authority, convert an enforcement problem into a policy problem, and make a structure bankable. What it does not do is prevent a policy decision taken against a correct and fully disclosed position, supply an appeal, or protect an asset sitting inside the perimeter that took the decision.** The event is an argument against concentration, not against disclosure — and a fully disclosed, fully correct, single-perimeter structure has an effective bloc count of 1.00, so $\ln(1.00) = 0$ and it has no available surplus at all.

The proposals were measured against themselves and each was found smaller than its advocates claim. A constructed eight-component backing basket reads **6.640106 effective components naively and 1.441753** once the climate-correlated components are recognised as a single 82.0 per cent exposure — **a naive overstatement of 4.605578 times**, with the basket still genuinely more diversified than a single-commodity anchor at 1.44 against 1.00. A constructed **250-to-1 paper-to-physical ratio leaves 99.600 per cent of claims uncollateralised at par and is exhausted by 0.40 per cent of holders standing for delivery**, and a fully collateralised object-ledger is credited only with making that ratio a live, reportable count — **it does not prevent over-issuance**, because per-object hash chains supply no consensus and no double-spend prevention on their own. Phased integration is scored as an option rather than a plan: constructed gate probabilities of 0.60, 0.50 and 0.30 give a **terminal probability of 0.090 against a first-phase 0.600, a ratio of 0.15**, so a programme that leads with its most ambitious phase stakes the certain value on the least probable branch.

Two failures need no adversary at all. A 26 per cent fall over four months breaches a **23.1 per cent** maintenance trigger in month three, and **buying survival to a 50 per cent drawdown costs 14.000 of the 40.000 borrowing, a 35 per cent reduction in leverage** — a loss of control that requires no court, no judgment and no finding of wrongdoing, only the contract the holder signed. And adding two jurisdictions inside a bloc a structure is already heaviest in raises the **naive effective count from 5.290949 to 7.132973, a rise of 34.81 per cent, while effective blocs fall from 2.941176 to 2.841490, a fall of 3.39 per cent**, and h falls to -0.134485. **Three years of additional maintenance cost bought less resilience, nothing went wrong, no adversary appeared, and every conventional report said the opposite.**

And a levy presented as a costless accounting correction reallocates burden. A constructed incidence puts **42.0 per cent on one sector**, spreads across an effective **3.631082** sectors, and passes a constructed **55 per cent through to consumers**. Recognising a previously unpriced cost does not make it disappear; it moves it, and the movement has a shape that must be modelled rather than framed away.

What is not claimed. No instrument described here exists, is issued, is authorised, or is claimed to be lawfully issuable anywhere. No price, yield or return appears. No commodity multiple, statute, regulation, notice or decided case is asserted as fact. No phase beyond the first is asserted to have been adopted anywhere. **The licensing provisions of §5.2.4 name the author's own company, that interest is declared rather than managed, and the provisions appear with their exposure attached and no outcome stated.**

The paper's own status, plainly. Its governing hypothesis was tested against the strongest available case and failed. Its remaining results are exact arithmetic on constructed inputs, and the load-bearing inputs are partitions this framework chose and a reader may reasonably contest. **A corpus whose strongest results are constraints on itself is a corpus that was checked; this paper is where the checking is printed.**

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every result in this paper is exact arithmetic on stipulated, **CONSTRUCTED** inputs, labelled as constructed at each point of use: an eight-component backing basket with printed weights; a commodity with 40,000 deliverable units and a stipulated ratio of 250; a reserve of 100 abstract units with 60 per cent voucher coverage; a drawdown base of 100 with printed sectoral shares; three gate probabilities; two jurisdiction allocations with printed bloc assignments; and a pledged position at a stipulated loan-to-value and maintenance requirement. **No real market, commodity, reserve, land parcel, economy, institution, court file, family or principal appears anywhere in this paper, and no allegation is made about any of them.**

All figures were produced by the scripts named in their captions (`figs/src/CCP-069_fig01.py` through `figs/src/CCP-069_fig05.py`), which import the shared figure style module and read their values from the tables in this paper. **No figure plots a number that does not appear in the text.**

The 2022 immobilisation, the commodity multiples, the decided cases, the statutes, the tax-code provisions, the administrative designations and the procurement, appropriations and constitutional frameworks are described from the source manuscript and are cited nowhere, because this paper's bibliographic database contains none of them. Appendix A itemises each with what specifically requires verification, and no approximate citation has been substituted for any of them.

Funding

This research received no external funding.

Conflicts of interest

The author declares a direct and material conflict of interest. The licensing provisions recorded in §5.2.4 are provisions of the source material that name **36N9 Genetics — the author's own company and the affiliation on this paper** — as the sole-source procurement counterparty for the production of the proposed currency, alongside a perpetual royalty payable in that currency and a sovereign licence granted by a private entity. **The author would be a beneficiary of the arrangement the paper records.**

The method adopted in response is stated in §3.4 and applied in §§5.2.4 and 5.2.5: the provisions are recorded exactly as the source specifies them; the federal procurement, appropriations, conflict-of-interest and constitutional review they invite is named in the same neutral register the framework applies to every other legal exposure; **no outcome is asserted; and no provision is recommended.** A reader who discounts §5 accordingly is reading correctly. **Disclosure does not cure this conflict, and this paper's own governing result is that disclosure does not cure things.**

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all figures, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data were processed. Every principal, family, entity, court file and payment described is constructed.

No operational content, and no allegation

This paper contains no operational content. It describes failure modes of structures at the level of principle and mechanism; it supplies no method for evading any legal process, any sanctions regime, any tax obligation or any disclosure requirement, and nothing in it is a routing instruction or a plan. **No allegation is made about any person, institution, jurisdiction or state**, and the 2022 immobilisation is recorded as an event whose occurrence is not in dispute, with no characterisation of its lawfulness, its merits or its motives offered anywhere.

Legal, tax, accounting, investment, procurement and constitutional advice

This paper is analytical and is not legal, tax, accounting, investment, procurement, appropriations or constitutional advice. The three backing mechanisms are proposals with **no adopted standard status and no legal force**. No instrument described is claimed to exist, to be issued, to be authorised, or to be lawfully issuable in any jurisdiction. **No price, yield or return figure appears anywhere in this paper**, and no recommendation is made in respect of any person, institution, product, instrument or jurisdiction. Anyone contemplating any structure, treatment or provision resembling those described must confirm every element with qualified counsel in each relevant discipline and each relevant jurisdiction, working from the primary sources itemised in Appendix A and not from this paper.

Appendix A. Claims requiring verification against primary sources

Everything below is described in this paper from the source manuscript and is cited nowhere, because this paper's bibliographic database contains no standards documents, no case law, no statutes, no administrative authorities, no press reporting and no jurisdiction-specific material. Each item requires verification against the governing source before any reliance. **No numerical result in this paper depends on any of them except item 20, whose role is stated there.**

Commodity markets and the paper-to-physical ratio.

1. The silver open-interest-to-registered-stock ratio, circulating in a range from roughly one hundred to several hundred times, is **contested and asserted nowhere**. Confirm the numerator (all open interest against front month), the denominator (registered only against registered plus eligible), the venue and the date before quoting any figure. The multiple is definition-dependent and no single value is a fact.
2. The registered-versus-eligible gold classification on a delivery exchange, and any ratio of open contracts to registered stock. Confirm the current definitions, the reporting source and the date. This paper states the mechanism — eligible stock is not pledged to contracts — and no figure.
3. The ratio of crude open interest to deliverable stock and storage capacity at a single pricing point. Confirm the contract's physical-delivery specification and the deliverable and storage figures. No figure is asserted.
4. The 2020 negative-price event in a physically settled crude contract. Confirm the contract, the date, the settlement mechanism and the storage-constraint mechanism by which the paper claim and the physical reality decoupled. This paper uses it as an instance of decoupling and states no price.
5. Periodic delivery squeezes in silver. Confirm the dates and the specific episodes; this paper refers to the class of event, not to any one of them.
6. The general claim that paper claims routinely exceed deliverable physical stock by roughly one to three orders of magnitude is stated as a **range** and as a general feature of hedging markets. Confirm it against commodity-market microstructure sources; it is not a measurement.

The backing mechanisms.

7. The composition of the proposed consumables index — the specific assets, the weighting method, and whether soil-health and biodiversity metrics are included and how — against the source Act. **The eight-component basket and its weights here are CONSTRUCTED and reconcile to no published index.**
8. The Act's stated design goal of mathematical immunity to inflation and deflation. Confirm the exact wording before quoting it, and **never reproduce the goal without the correction that accompanies it in §4.2.3.**
9. The qualified conservation contribution regime of the United States Internal Revenue Code, including the qualified-organisation, conservation-purpose and perpetuity requirements. This paper states that conservation easements are real, valued, tradeable instruments with a statutory basis and **asserts no valuation and no section number.**
10. The administrative authority designating syndicated conservation easements as a listed transaction, its status, subsequent statutory changes, and the litigation over both the transactions and the designation. **The notice number is not asserted here.** The structural point — that easement valuation is contested and has an abuse history — does not depend on it.
11. The ecosystem-services and natural-capital valuation literature, referred to generally. **No work in it is cited and this paper endorses no methodology and states no valuation figure.**
12. The floating voucher and the twelve per cent commons-replenishment allocation, as the Act frames them: the exact mechanism, the voucher's legal characterisation, the base on which the allocation is assessed, and its stated use. The **CONSTRUCTED** incidence illustration is a measurement of no real economy.

The phased programme and its licensing provisions.

13. The Fibonacci denomination sequence and the Act's rationale for it. The mathematics — consecutive ratios approaching φ, and every positive integer expressible as a sum of non-consecutive terms — is cited; **the Act's wording is not, and must be confirmed.**
14. The fractal reserve-banking structure, recorded as the Act's design goal and graded [T3]. Self-similarity of a rule is not a guarantee of stability under stress and **no stability claim is made on the geometry alone.**
15. **The licensing provisions**: a sovereign licence granted by a private entity, a perpetual royalty payable in the proposed currency, sole-source procurement from the author's own company, and a note bearing the architect's likeness. **Confirm the Act's exact wording of each before reproducing it, and reproduce no provision without its exposure attached.**
16. Federal procurement competition: the statutory competition framework, the acquisition regulation's competition requirements, and the law governing statutory sole-source directions. **This paper names the framework and states no outcome.**
17. Appropriations and continuing obligations: the constitutional appropriations provision, the law on creating and funding continuing financial obligations, whether such an obligation can bind future appropriations, and the antideficiency rules. **Category named; no outcome stated.**
18. Conflict-of-interest statutes and the appearance-of-impropriety standard governing public office, as they bear on sole-source procurement from a connected party and on a likeness appearing on a currency instrument. **Standard named; no outcome stated.**
19. Constitutional provisions bearing on currency, on emoluments, and on the delegation of a sovereign monetary function to a private licensor. **Categories named; none resolved.**

The failure record.

20. **The 2022 immobilisation of approximately USD 300 billion of Russian central bank reserves: the instruments by jurisdiction, the dates, the amounts by jurisdiction, and the current legal position on both the principal and the income arising from it.** This paper

states the aggregate as approximate, treats no figure as established, and **attaches no citation to any part of it**, because the bibliography contains no contemporaneous reporting and no sanctions instruments. **This is the one described item on which a stated result depends** (R1), and the dependence is only on the fact that a fully disclosed, correct, lawfully held position was immobilised — not on the amount, the date or the instrument.

21. The decided cases on personal compulsion and the impossibility defence, on freezing relief and its worldwide reach, on the limits of pre-judgment injunctive power, and on sham and nominee findings. **Every reporter citation named in §§6.2.1 and 6.2.3 must be checked against the official reports**, including the situs and administration history of the trust in the leading contempt case, which is routinely miscited. The illusory-trust line is **named nowhere** and no reliance should be placed in either direction.
22. The self-settled-trust look-back provision of the United States bankruptcy code and its ten-year period; the transactions-defrauding-creditors provision of English insolvency law and its limitation analysis; and the uniform voidable transactions legislation, whose enactment and reach-back periods vary by state. **Contested ground throughout.**
23. The offshore trusts statute whose limitation periods, standard of proof and non-recognition provision are described. **This paper deliberately names no section numbers.**
24. Economic-substance legislation in each affected jurisdiction and its current revisions: short titles, commencement, relevant-activity lists and adequacy tests. Corporate-residence and place-of-effective-management tests in every jurisdiction from which decisions might be taken, and the tie-breaker in each applicable treaty. The principal-purpose test as introduced by the multilateral instrument, and certificate-of-residence practice.
25. The foreign-business nominee provision described: section number, current wording, penalties, enforcement practice, and the routes, exposure and disclosure obligations for regularising a historical nominee position. **Carried from the Thailand analysis and repeated here because it is the most consequential item in this appendix after item 20.**
26. Share-pledge disclosure for directors and officers; margin-lending documentation in the relevant market, including maintenance-requirement definitions, cure periods, revaluation discretion and rehypothecation rights; and change-of-control definitions across a group's agreements. **The 35 per cent maintenance requirement used here is a constructed model parameter, not a market quotation.**
27. All bloc assignments used in §§6.2.6 and 6.2.7 are this framework's own, applying its perimeter test to published evidence. **They are contestable**, the sensitivity is measured separately [Curzi, 2026], and each jurisdiction paper restates its assignment with evidence.
28. All parameter values in the surplus computations ($\eta = 0.35$, $C = 0.30$, $Q = 10$, $\delta = 0.02$) are the corpus's stated test parameters. They are illustrative and are estimates of nothing.

All numerical results in this paper were computed from the stated constructed inputs at full precision and rounded only for printing.

Chapter 70

The Method in Seven Steps: A Consolidated Reproduction Workbook and Honesty Register for the Concentration and Headroom Instrument

Abstract. This closing synthesis states the method as a seven-step procedure and supplies the appendices needed to reproduce every computation in the volume. The procedure's binding instruction is to collapse jurisdictions onto coordination perimeters before measuring, because the organising result is that ln(1.00) = 0: a structure with one effective bloc has zero available surplus, not a low one. Reproduced end to end, the conventional structure has a naive HHI of 0.293430 and 3.407968 effective jurisdictions but a bloc HHI of 1.000000, giving N = 1.000000 and h = -0.500000, while the designed structure has naive HHI 0.189002 and 5.290949 effective jurisdictions against bloc HHI 0.340000, N = 2.941176 and h = -0.122417, with Western-perimeter exposure falling from 100.0 to 30.0 per cent. Break-even is N* = exp(0.5/0.35) = 4.172734, and since dh/dN = eta/N the first move buys 0.377583 while the next buys 0.107620, 28.5 per cent as much. The rail axis is measured separately at 1.000000 effective clearing system as found and 2.810568 as designed, with four effective systems requiring the dollar share to fall from 55 to about 38.99 per cent. Two computations are presented here in full: the foundation-versus-trust cost crossover, and a succession residual capping a USD 6.70m exposure while leaving approximately USD 27.00m untouched.

Keywords: bloc collapse; effective count; surplus headroom; break-even bloc count; diminishing returns; reproduction workbook; honesty register; succession residual

Series: D. Structure and Resilience **Paper:** CCP-070

1. Introduction

This paper is reference apparatus, and saying so plainly is the first thing it owes a reader. It consolidates the method Series D develops into a single executable procedure, reproduces every computation the series rests on from its stated inputs, and collects the atlases, indices and

registers a reader needs in order to check the series against itself. **It is not offered as a research contribution, it advances no new result about the world, and almost all of it recapitulates.**

Two things in it are not recapitulation, and they are named here so that a reader can find them without reading the rest. **The foundation-versus-trust cost crossover is computed in full for the first time** (Appendix B.13): against a constructed fixed mandate of USD 120,000 a year, a foundation is cheaper than a trustee charging on a basis-point scale only above USD 34.2857 million at 35 basis points, USD 80.0000 million at 15, and USD 120.0000 million at 10 — from which it follows that **cost is explicitly not the reason to choose the form.** And **the succession residual is stated as an asymmetry rather than as a win**: a notional reserved share of USD 60.00 million capped to a USD 6.70 million compensatory levy and eliminated on the designed map, against an inheritance-tax exposure of approximately USD 27.00 million that the structure does not touch — 4.0299 times larger — with a liquidity shortfall of USD 33.84 million, 14.1 per cent of the estate.

What the consolidation is for. A method that cannot be run is a description of a method. The procedure below reduces the series to seven steps that can be executed in an afternoon with a spreadsheet and a list of entities, and its binding instruction is the first one: **collapse jurisdictions onto coordination perimeters before measuring anything.** The reason is the organising result of the whole series and it is one line of arithmetic. $\ln(1.00) = 0$. **A structure with one effective coordination perimeter has no surplus term at all — not a small one, none — and no rearrangement of entities inside that perimeter changes it**, because rearranging within a perimeter does not change N [Curzi, 2026].

The reproduction, end to end. The conventional structure reads a naive Herfindahl index of 0.293430 and 3.407968 effective jurisdictions, and a bloc Herfindahl index of 1.000000, giving $N = 1.000000$ and $h = -0.500000$. The designed structure reads naive 0.189002 and 5.290949 effective jurisdictions against bloc 0.340000, giving $N = 2.941176$ and $h = -0.122417$, with Western-perimeter exposure falling from 100.0 to 30.0 per cent. Break-even is $N^{*} = \exp(0.5/0.35) = 4.172734$, and because $\mathrm{d}h/\mathrm{d}N = \eta/N$ the first move buys 0.377583 while the next buys 0.107620 — **28.5 per cent as much.** The rail axis is measured separately and reads 1.000000 effective clearing system as found against 2.810568 as designed [Curzi, 2026].

And the honest reading of those numbers is the paper's fourth contribution, which is a register of admissions rather than a result. Both structures are below the build threshold. Disclosure did not prevent the 2022 immobilisation. Cost is not the reason to choose a foundation. The succession structure addresses the smaller of the two exposures. **Section 5 states all four without qualification, and a reader who takes nothing else from this paper should take those.**

1.1 What this paper contains

1. **The method in seven steps** (§4), with each step expanded, its failure modes named, and the whole procedure run once end to end on the carried constructed case.
2. **The register of what is unresolved** (§5), including the four admissions and the eight open items behind them.
3. **A scorecard atlas** (Appendix A): thirty-three jurisdiction assessments on twelve items, keyed to the corpus paper that carries each.
4. **A reproduction workbook** (Appendix B): every concentration, headroom, break-even and job-vector computation in the series, with its inputs printed, its arithmetic shown, and its precision rule stated.
5. **Structure diagrams** (Appendix C) of the constructed case as built, as found, collapsed to perimeters, and viewed on the rail axis.
6. **A statute and case index** (Appendix D) in which the **default status is FLAGGED FOR VERIFICATION** and only a short list of arithmetic and definitional facts is marked verified.
7. **An honesty register** (Appendix E) carrying the ten reconciliation entries by which the series corrected itself, and the four admissions.
8. **A correspondence appendix** (Appendix F) mapping the storage law term by term onto

the earlier series, and naming four things that do **not** transfer.

2. Hypothesis

This paper advances no testable hypothesis. It is reference apparatus: a consolidated procedure, a reproduction workbook, and a set of registers, assembled so that the results of Series D can be checked rather than believed.

The organising question is therefore procedural rather than empirical: *what is the smallest set of steps, inputs and printed intermediate values from which a reader can reproduce every number in the series and locate every claim the series declined to verify?*

Three things in the paper could nevertheless be wrong, and each is checkable, so they are named here rather than left implicit. The reproduction arithmetic could contain an error, and every input is printed so that it can be recomputed. The atlas cells are compressions of assessments argued elsewhere, and a cell that misrepresents its assessment is an error in the compression. And the index's status assignments could be wrong in the direction that matters — an item marked verified that is not — which is why **the default is FLAGGED and the verified list is deliberately short** (§3.3).

What is not claimed. No result in this paper is new to the corpus except the two computations named in §1. Nothing here establishes any empirical claim about any jurisdiction, market, institution or family. **The four admissions of §5 are the paper's most useful content and every one of them is negative.**

3. Method

3.1 What is reproduced, what is new, and what is only compressed

kind	where	status
the seven-step procedure	§4	a consolidation of results argued elsewhere in the series
the end-to-end reproduction	§4.5, Appendix B	recomputed here from printed inputs
the foundation-versus-trust crossover	Appendix B.13	**computed in full for the first time**
the succession residual	§5.3, Appendix B.13	**stated as an asymmetry for the first time**
the scorecard atlas	Appendix A	compressions; each cell points to the paper that argues it
the statute and case index	Appendix D	a consolidated list, **FLAGGED by default**
the honesty register	Appendix E	the series' own corrections, reproduced
the correspondence	Appendix F	a mapping, not a derivation

Every numerical result in this paper is exact arithmetic on stipulated, constructed inputs. No real family, entity, transaction, market or institution is described anywhere in it, with the single exception of the historical institutions treated from the public record in the lineage assessments, which this paper only indexes. Each construction is printed at the point of use and labelled CONSTRUCTED.

3.2 Which effective count, and the precision rule

Where this paper reports an effective count it means $1/\mathrm{HHI}$, the reciprocal of the Herfindahl–Hirschman index and the Hill number of order two [Hirschman, 1964, Herfindahl, 1950, Hill, 1973, Jost, 2006], whose numbers-equivalent reading is the standard apparatus this paper borrows [Adelman, 1969, Laakso and Taagepera, 1979, U.S. Department of Justice and Federal Trade Commission, 2023, 2010]. **A summary statistic computed over a partition does not determine the partition** [Kirman, 1992, Harrison and Klein, 2007], which is why §6.4 states the invariance and its price rather than leaving a reader to infer an allocation from a count. **The corpus's affine construction $N_{\mathrm{eff}} = 1 + (N-1)u$ is not a Hill number of any order** [Curzi, 2026c]; it appears in this paper only in Appendix F, where the correspondence to the earlier series is drawn and where it is badged as a design axiom at the point of use.

The precision rule is stated once and applies everywhere (Appendix B.15). Shares are stipulated to one decimal place; indices, effective counts and surplus values are computed at full floating precision and printed to six decimal places; and **no printed value is ever used as an input to a later computation.** Six decimal places on a quantity derived from shares stipulated to one is a statement about the arithmetic and not about the world, and §6.2 says so again.

3.3 The status convention, and why the default is FLAGGED

Appendix D assigns every legal citation in the series one of two statuses.

VERIFIED is reserved for facts fixed by definition or by arithmetic recomputed from stated inputs, and for nothing else. **FLAGGED FOR VERIFICATION** is the default and covers every case reporter citation, every section number, every date, every threshold and every current-status claim. **A reader who treats the flag as a formality has misunderstood the design.** The scheme is a two-valued ordinal label, no arithmetic is performed on it, and an ordinal scale licenses none [Stevens, 1946, Michell, 1997].

No jurisdiction-specific legal source is cited anywhere in this paper. The statutes, conventions, court rules and decided cases consolidated in Appendix D are named because they exist and are described from the source manuscript; **this paper's bibliographic database contains none of them and no citation is attached to any of them.** The general doctrinal frame is cited where it is relied on [Hague Conference on Private International Law, 1985, United Nations, 1958, United Nations Commission on International Trade Law, 1985, Torremans et al., 2022, Briggs, 2013, Hansmann and Kraakman, 2000, Pistor, 2019], and the disclosure regime that removed non-disclosure from the available design parameters is cited where the disclosure test is stated [Organisation for Economic Co-operation and Development, 2017, Financial Action Task Force, 2012, Organisation for Economic Co-operation and Development, 2015, de Goede, 2011]. **Publishing a dated, status-keyed register of what has not been verified is a recognised integrity practice rather than a stylistic choice** [Nosek et al., 2018, Simmons et al., 2011, Ioannidis, 2005, Gelman and Loken, 2013], and badging an unsupported claim rather than dropping or dressing it is the honest description of what a research programme does with its own unproven components [Lakatos, 1970, Platt, 1964, Meehl, 1990].

3.4 The key to the atlas

Appendix A's second column and Appendix D.2's first column carry a three-digit key. **It is the last three digits of the corpus paper — CCP-0nn — that carries the assessment being compressed**, so that every cell in the atlas can be traced to the paper that argues it. The atlas is an index, not an argument, and **no cell in it should be relied on without reading the paper it compresses and the verification register that paper carries.** An index of this kind is an installed base rather than a finding: it accumulates, it decays if it is not maintained, and its usefulness is a property of its upkeep [Star and Ruhleder, 1996, Star and Griesemer, 1989].

3.5 What is not invoked

The corpus's reflexive loop is not used in this paper and no social-influence parameter is estimated, so the reflection problem [Manski, 1993, Shalizi and Thomas, 2011] does not bear on any result reported here. **No distributional, scaling or heavy-tail claim is made anywhere**, so the goodness-of-fit and likelihood-ratio procedure such a claim would require [Clauset et al., 2009] is not applicable and was not run. Every figure in this paper is a reciprocal of a sum of squared shares, a logarithm of such a reciprocal, a ratio of two counts, or a closed-form solution of a linear equation on stipulated inputs.

One borrowing is flagged before the appendices use it. The framework writes concentration, erosion and surplus in the grammar of extensive and conserved quantities without a conservation law, which is the characteristic failure of this genre [Mirowski, 1989]; the defence taken here is Cartwright's — validity is claimed within a specified arrangement and no universality is claimed for it [Cartwright, 1999]. §6.5 grades every borrowing in the paper by name.

4. The method in seven steps

4.1 The position: what the procedure buys and what it does not

A method that cannot be run is a description of a method.

The preceding papers of this series establish a thesis, an instrument, a taxonomy, a design standard, a set of jurisdictions with evidence attached, and a failure record. All of that is inert unless it reduces to something a reader can execute on a Tuesday afternoon with a spreadsheet, a list of entities, and an hour. This section is that reduction. It is deliberately the shortest load-bearing part of the series, and everything in it is derived — nothing here is new, and every step points back to the section that owns it.

The compression has to survive one honest constraint, and stating it now rather than at the end would be dishonest in the other direction, so it is stated in both places. **The procedure below is correct and it does not produce a structure that pays for itself at the tested parameters.** That is the framework's own arithmetic applied to the framework's own recommendation, and the last section of this section does nothing but set out what remains unresolved, because a synthesis that ends on a triumph is a brochure.

What the procedure does is different and narrower, and it is worth being precise about, because precision here is the difference between a useful tool and a disappointed reader:

> **Jurisdictional and rail diversity across independent coordination perimeters raises the ceiling on what a structure can hold. It does not pay the structure's maintenance bill. Diversity buys the ability to build; it does not buy the building.**

That formulation is the assembled architecture's and it governs this section. A reader who runs the procedure and finds a negative number has not made a mistake. The reader has found the actual position, which is what the instrument is for.

One further framing, carried from the Hong Kong assessment and applied to the whole method. Everything the procedure buys is bought for the tail of the distribution and not for the median. In the median year, a conventional single-perimeter structure works fine, costs less, is simpler to administer, and produces no complaints from anyone. **The architecture in this series is worse than the conventional one in almost every year and better in the year that matters, and a reader who is not buying for the tail should not buy at all.** That is not a caveat. It is the product description.

4.2 Why the procedure has the shape it has

The procedure has the shape it has because the storage law has the shape it has, and each step maps to one term. This is the whole of the reasoning in nine lines.

```
Q_{t+1} = (1 - delta) Q_t + eta S_t - C_t

h = eta ln(N) - delta Q - C

  Q      protected capital
  delta  erosion: tax drag, political risk, debasement, haven degradation
  eta    conversion of income into DURABLE protected holdings
  C      structure maintenance: compliance, counsel, substance, filings
  N      EFFECTIVE BLOC COUNT -- independent coordination perimeters,
         not entities and not jurisdictions
```

`h` is the sustaining rate: positive and the structure builds, negative and it is being funded out of something else. Four terms, four levers, and the procedure works through them in the order that makes the later steps meaningful.

`N` is the only term that appears inside a logarithm, which has two consequences the procedure has to respect. First, the marginal value of an additional independent perimeter is `eta/N` — falling, so the first move is worth much more than the fourth. At the designed structure's `N = 2.941176`, `dh/dN = 0.35 / 2.941176 = 0.119000`, still the most valuable structural move available. Second, and this is the result that organises everything: `ln(1.00) = 0`. **A structure with one effective bloc has no surplus term at all — not a small one, none — and no rearrangement of entities inside that bloc changes it, because rearranging within a perimeter does not change `N`.**

`C` is the term that is almost always mis-estimated downward, because it is paid in small recurring amounts to many different people and never appears as a line item. It includes registered offices, agents, audits, director fees, local counsel, substance payroll, filings, translations, the annual re-run of the whole procedure, and the internal management time that the family does not bill itself for. The AIFC result is the sharpest illustration of what happens when `C` is left out: the AIFC position is worth 0.026731 of `h`, which is entirely consumed if the leg costs more than 8.9102% of the structure's total maintenance budget. **The geopolitical question was interesting and the cost question was binding.**

`delta` and `eta` are the terms the jurisdictions actually move. A jurisdiction with a stable legal system, a working court, and no expropriation history lowers `delta`. A jurisdiction that lets income be converted into durable held form — a foundation with separate legal personality, a recognised trust, a licensed insurer holding reserves, a REIT unit rather than a title deed — raises `eta`. Most of what the jurisdiction papers describe operates on these two terms, and almost none of it operates on `N`, which is why the concentration measurement is a separate step that has to be done separately.

4.3 The procedure, specified

4.3.1 The procedure on one page

```
=========================================================================
THE PROCEDURE
=========================================================================

0  NAME THE THREAT                                                 Ch 4
   Which of the five: arbitrary action / lawful process / commercial /
   succession / bloc fragmentation. Different structures answer different
   threats. A structure built against the wrong one passes every test and
```

```
   fails anyway.
   STOP if the answer is "lawful process" and the claim already exists.
   That is fraudulent-transfer law and timing, not architecture.

1  MEASURE BLOC CONCENTRATION                                        Ch 5
   List every holding as a share of protected capital.
   Assign each to a coordination perimeter using Ch 4's test.
   COLLAPSE TO PERIMETERS FIRST, then compute:
       HHI = sum of squared bloc shares      N = 1 / HHI
   Compute h = eta ln(N) - delta Q - C.
   Compute the break-even N* = exp((delta Q + C) / eta).
   Never print the naive per-jurisdiction count next to a resilience claim.

2  ASSIGN EACH ENTITY EXACTLY ONE JOB                                Ch 6
   HOLDING / OPERATING / BRIDGING. Score the vector (H, O, B).
       S = H^2 + O^2 + B^2       effective jobs = 1 / S
   Above 2.20 is overloaded: split it or accept the strain in writing.
   Score candidates on each requirement dimension, then take the MINIMUM
   over the dimensions the job needs -- never the average.

3  APPLY THE DISCLOSURE TEST                                         Ch 3
   Would the structure still do its job if every element were disclosed
   to every authority with jurisdiction over any part of it?
   Resilience structures pass. Concealment structures fail by construction.
   A structure that fails is not fixed by relocating it.

4  PRICE EACH MARGINAL JURISDICTION AS A PROBABILITY                 Ch 22
   Do not argue the geopolitics. Compute:
       G = gain in h if the perimeter call is RIGHT
       L = loss in h if the perimeter call is WRONG
       break-even probability p* = L / (G + L)
   Then compute the cost ceiling: the leg must cost less than G,
   expressed as a percentage of total C.
   A priced loss taken knowingly is legitimate. An unpriced one is failure.

5  CHECK HEADROOM BEFORE BUILDING                                    Ch 34
   Fix the drawdown d you will survive; solve for the loan, not the reverse:
       L_max = V (1 - d)(1 - m)         m = maintenance requirement
   Check liquidity against succession and margin simultaneously --
   they draw on the same buffer.
   Ring-fence BEFORE the pledge; a cross-default clause is a designed hole.

6  VERIFY WITH COUNSEL IN EVERY JURISDICTION TOUCHED
   Residence, nationality, beneficiary location, asset situs, counterparty
   situs, and the situs of every decision-maker. Written classification of
   every entity for CRS, FATCA, and local reporting AT FORMATION.
   Where the framework declined to guess, the framework declined for a reason.

7  WRITE IT DOWN AND RE-RUN IT ANNUALLY
   Date, named person, diff against last year. Every bloc call carries its
   evidence and the specific event that would reverse it. Every residual is
   stated. A register nobody reviews is worth nothing.
================================================================================
```

4.3.2 Step 0: name the threat

The threat register separates five threats that are routinely collapsed into one, and the separation does more work than any jurisdiction selection.

Arbitrary action — expropriation, immobilisation, de-risking without process, retroactive regime change. This is the subject of the framework and the threat that `N` answers.

Lawful process — a judgment that survived due process. Timing-bound. Fraudulent-transfer law governs and no jurisdiction fixes it after the fact. The failure record's tracks one and three are almost entirely this threat, and the structures that survived it survived on timing.

Commercial and counterparty — insurance and structure. The failure record's track four lives here, and it is the threat that requires no adversary at all.

Succession — forced heirship, fragmentation across heir jurisdictions, and the liquidity to pay what falls due on death. The succession analysis's result is the caution: the structure solved the smaller of the two exposures.

Bloc fragmentation — the rails and the perimeters coming apart, which is the threat the last fifteen years have been trending toward and the one most structures are least built for.

The instruction at step 0 is a stop rather than a choice. If the honest answer is lawful process and the claim already exists, the reader should be with litigation counsel, not with this series.

4.3.3 Step 1: collapse before measuring

The instrument is elementary and the discipline is not.

```
CONVENTIONAL, as found

  UK 40.9  US 27.3  Jersey 18.2  Luxembourg 13.6
  naive HHI = 0.293430        effective jurisdictions = 3.407968

  collapsed:  WEST 100.0
  bloc HHI  = 1.000000        EFFECTIVE BLOCS = 1.000000
  ln(1.00)  = 0.000000        h = -0.500000

DESIGNED

  UAE 29.1  Thailand 18.2  UK 16.4  Bermuda 13.6
  India 11.8  Singapore 10.9
  naive HHI = 0.189002        effective jurisdictions = 5.290949

  collapsed:  BRIDGE 40.0  NONALIGNED 30.0  WEST 30.0
  bloc HHI  = 0.340000        EFFECTIVE BLOCS = 2.941176
  ln N      = 1.078810        h = -0.122417

  Western-perimeter exposure  100.0%  ->  30.0%
  break-even N* = exp(0.5 / 0.35) = exp(1.428571) = 4.172734
```

Four entities, four jurisdictions, one point of failure. The naive reading says 3.407968 jurisdictions; the honest reading says 1.000000 bloc. That gap is the entire reason the step exists.

The Singapore assessment supplies the cleanest proof that the gap is real and not an artefact of the bloc assignment. Place the same 10.9% in London instead of Singapore and the naive count falls from 5.290949 to 4.449309 — a dramatic, visible, entirely spurious difference — while the honest count moves from 2.941176 to 2.924302 and `h` from −0.122417 to

−0.124431. **Singapore's true marginal contribution to `h`, measured against simply putting the money in London, is 0.002014. The naive metric reported nearly a full effective jurisdiction of diversification that did not occur.**

The Gulf assessments supplies the failure in the other direction. Add Qatar and Saudi Arabia funded by trimming UAE, the United Kingdom, and Singapore: eight jurisdictions, naive effective count 7.132973 — up 34.81% — and effective blocs 2.841490, down 3.39%, with `h` falling to −0.134485. **Two more jurisdictions, a third more apparent diversification, and the honest count fell.**

Two further disciplines belong to this step and are routinely skipped.

Measure the rails separately, and by terminal clearing system rather than by bank. The rail measurement's correspondent measurement: as found, 100% USD, HHI 1.000000, one effective clearing system. As designed, USD 55 / CNY 15 / EUR 12 / INR 10 / AED 8, HHI 0.355800, 2.810568 effective systems. A structure can have `N = 2.941176` on holdings and 1.000000 on the ability to transact.

Record the sensitivity. Reclassifying Singapore from bridge to Western — the argument most likely to be had about the whole allocation — moves `N` by 0.016874 and `h` by 0.002014. The argument does not move the answer. The concentration does.

4.3.4 Step 2: one entity, one job

The three-jobs diagnostic owns the taxonomy and it prevents more failures than any other single idea in the series.

```
HOLDING    succession law, asset segregation, perpetuity, low substance burden
OPERATING  real substance, labour, cheap presence, market access
BRIDGING   dual-bloc banking, dual-rail settlement, neutral law
```

The diagnostic is the concentration instrument turned on the entity instead of the portfolio. Express the entity's function as a vector, compute `S = H² + O² + B²`, and read `1/S` as effective jobs. A pure specialist returns 1.0000. The best possible two-job entity — an exact 50/50 split — returns exactly 2.0000, and there is no two-job configuration above it. The 2.20 threshold leaves room for estimation error in the vector and nothing more.

The designed structure, scored, with the framework's own flag left in:

entity	H	O	B	S	eff jobs	verdict
DIFC Foundation	0.90	0.00	0.10	0.8200	1.2195	SPECIALIST
Thai BOI opco	0.05	0.95	0.00	0.9050	1.1050	SPECIALIST
UK Ltd	0.10	0.85	0.05	0.7350	1.3605	SPECIALIST
Singapore treasury	0.05	0.15	0.80	0.6650	1.5038	STRAINED
ADGM holdco	0.70	0.05	0.25	0.5550	1.8018	STRAINED
Bermuda captive	0.30	0.60	0.10	0.4600	2.1739	STRAINED
GIFT City vehicle	0.20	0.30	0.50	0.3800	2.6316	OVERLOADED
mean effective jobs per entity					1.6852	

The framework's own architecture produces one overloaded entity, and that is the reason the diagnostic is worth having rather than a reason to distrust it. The Bermuda captive at 2.1739 is strained by nature — a captive underwrites, which is operating, and holds reserves, which is holding — and the strain is accepted in writing because the reserves are held against real third-party liabilities, which is precisely what makes them a structurally awkward class to realise. The GIFT City vehicle at 2.6316 is carrying all three jobs and is a design error the framework owns rather than smooths over.

The selection rule matters as much as the diagnostic. Score each candidate on each requirement dimension and take the **minimum** over the dimensions the job actually needs. An average lets an excellent low-substance score pull up a mediocre succession score; the minimum does not, because a structure fails at its weakest load-bearing requirement and not at its average one.

4.3.5 Step 3: the disclosure test

> Would the structure still do its job if every element were disclosed to every authority with jurisdiction over any part of it?

Resilience structures pass. Concealment structures fail by construction. The test is a filter applied before cost, before jurisdiction, and before elegance, and it is cheap to run because the answer is usually obvious within one sentence of asking.

The disclosure paper applies it to five constructed structures and two of them fail in print — a Thai company held through nominees, and a captive that writes no risk. Both fail for the same reason: their operation depends on a register, a return, or a filing saying something other than what is true. Relocating either does not fix it, because the defect is not in the jurisdiction.

Two clarifications that the test's brevity invites people to miss.

It is a filter, not a shield. The disclosure paper carries its own strongest counter-case as its seventh failure mode, and this section repeats it below under what is unresolved. Passing the test does not prevent a policy decision taken against a fully disclosed position.

It is about facts, not strategy. The bright-line analysis's formulation: "hidden in plain sight" means the strategy is non-obvious while every fact is filed. A roll-up executed in full public view over a decade that almost nobody noticed passes the test. Concealment of fact does not, and no amount of sophistication converts the second into the first.

4.3.6 Step 4: price the jurisdiction rather than argue it

This is the step that most changes how a reader behaves, because it replaces an argument that cannot be settled with a number that can.

The form, from the AIFC assessment, generalised:

```
1  Compute h for the structure WITHOUT the candidate leg.
2  Compute h WITH the leg, on the assumption that the perimeter call
   is RIGHT.  The difference is G.
3  Compute h WITH the leg, on the assumption that the perimeter call
   is WRONG -- the leg folds into a bloc the structure already holds.
   The difference is -L.
4  Break-even probability   p* = L / (G + L)
5  Cost ceiling             the leg must run below G of maintenance,
                            expressed as a share of total C.
```

Worked, on the AIFC at 5% drawn from the United Kingdom slot:

```
baseline           blocs 40.0 / 30.0 / 30.0        N = 2.941176   h = -0.122417
call RIGHT         blocs 40.0 / 30.0 / 25.0 / 5.0  N = 3.174603   h = -0.095686
                                                    G = +0.026731
call WRONG         blocs 40.0 / 35.0 / 25.0        N = 2.898551   h = -0.127526
                                                    L =  0.005110

p* = 0.005110 / (0.026731 + 0.005110) = 0.005110 / 0.031841 = 0.160476
cost ceiling = 0.026731 / 0.30 = 8.9102% of total C
shortfall closed if right = 0.026731 / 0.122417 = 21.8357%
```

A 16.05% probability that the perimeter is genuinely independent is enough to justify the position, and the binding constraint is cost rather than geopolitics. Both halves of that sentence are useful. The first tells a family it does not need to win the argument about Kazakhstan; it needs the argument to be better than roughly one-in-six. The second tells the family where the decision is actually made, which is in the fee schedule.

The counterpart is the loss case, and it is the reason the step is mandatory rather than advisory. A wrong call does not merely fail to help. It costs 0.005110, because capital was withdrawn from a bloc where it was diversifying and placed in a bloc it concentrated. **That is the unpriced bloc call, and it is invisible to every conventional measure, including the naive effective jurisdiction count, which rises.**

4.3.7 Step 5: headroom before building

The carve-out paper's cascade, reversed so that the input is the quantity a family can have an opinion about.

```
d = drawdown to be survived      m = maintenance requirement
V = collateral value             L_max = V (1 − d)(1 − m)

at V = 80.00, m = 35%:
  survive 23.1%   L = 40.000   LTV 50.0%   <- the carried case
  survive 35.0%   L = 33.800   LTV 42.3%
  survive 50.0%   L = 26.000   LTV 32.5%
  survive 60.0%   L = 20.800   LTV 26.0%
```

At the carried constructed case's own parameters — a 20% carve-out, 80% retained and pledged at 50% loan-to-value, producing 60 deployable at 80% retained control — the maintenance call triggers on a drawdown of 23.0769%, which is the corpus's pledge-trigger result [Curzi, 2026]: $d^{*} = 1 - \mathrm{LTV}_{0}/(1-m)$, so at a 50 per cent loan-to-value and a 35 per cent maintenance requirement the first call falls at 23.076923 per cent and not at 50. **That is ordinary equity volatility, not a crisis parameter, and the carve-out cascade leads with it rather than burying it for exactly that reason.**

Two further checks belong to this step and are usually run separately, which is the error.

One buffer, two claims. The succession analysis's constructed liquidity test shows that liquidity rather than entitlement binds [Curzi, 2026]: a USD 60.00 million claim against USD 26.16 million of liquid assets leaves a shortfall of USD 33.84 million, 14.1 per cent of the estate, sufficient to force the sale of an operating branch on a death. The margin cure period draws on the same liquid holdings. **The succession liquidity problem and the margin call are the same problem arriving through two doors, and they should be modelled against one buffer.**

Ring-fence before the pledge. The ring-fencing analysis's segregation is effective against a claim that arises after it. A cross-default clause signed later is a designed hole in the fence, and a forced sale of pledged shares can trip change-of-control definitions in agreements that have nothing to do with the pledge.

4.3.8 Step 6: verify, and record where the framework refused to guess

Every section in this series closes with a Verification Register, and Appendix D consolidates them. The registers are not a disclaimer. They are the framework's integrity mechanism, and the house standard they encode is worth stating in the procedure itself:

> **Declining to guess, in print, is the standard. A section that names no section number where it is unsure of one is more useful than a section that names a plausible number.**

The AIFC assessment is the model: it opens by warning that it names and numbers no AIFC subordinate Act, Regulation, or Rule, and it states in print that it declines to guess whether the AIFC foundation confers separate legal personality. That refusal is information. A reader who takes it to counsel gets a real answer; a reader given a confident guess gets a false one.

The minimum verification set, per structure:

```
- counsel in each jurisdiction of incorporation
- counsel in each jurisdiction of residence -- principal AND every heir
- counsel in each jurisdiction of asset situs, separately for immovables
- tax counsel wherever a decision concerning any entity may be taken,
  because that is where corporate residence is asserted
- Sharia-qualified counsel wherever personal-status law applies
- written CRS and FATCA classification of every entity AT FORMATION,
  including each sub-fund or cell
- securities counsel before any pledge by a director or officer
```

4.3.9 Step 7: the annual re-run

The procedure decays. Bloc assignments move — Singapore's did, in March 2022, and the evidence is one dated episode rather than a trend. Job vectors drift, because a holding entity acquires an operating subsidiary and starts making decisions. Substance requirements are amended. Incentive conditions are revised, twice in the Singapore family-office case. Correspondent relationships are withdrawn without notice or explanation.

The re-run is administrative rather than intellectual and takes an afternoon: recompute steps 1, 2, and 4 from current allocations; diff the result against last year; re-read every bloc call against its written trigger; and re-open every Verification Register item whose underlying instrument has been amended. Date it and put a name on it.

4.4 Failure modes of the procedure

Eight ways the procedure gets run wrong. Each has been observed in the sections that produced it.

FM 4.1 Running step 1 without collapsing

The single most consequential error, and the reason the concentration paper owns the instrument rather than distributing it. A naive per-jurisdiction effective count is not a resilience number. Printed next to a resilience claim it is worse than no number, because it is precise, plausible, and wrong in the direction the reader wants.

FM 4.2 Running the steps out of order

Step 4 is meaningless before step 1, because the gain and loss are both differences in h and h requires N. Step 5 is meaningless before step 2, because the headroom check needs to know which entity holds the liquid assets. Step 3 belongs before step 4 rather than after it, because a structure that fails the disclosure test should not be priced; it should be abandoned.

FM 4.3 Treating step 6 as a formality

Counsel in every jurisdiction touched is expensive, slow, and the largest single component of C in a small structure. It is also the step that converts the framework from a map into something operable. A structure verified in one jurisdiction and assumed in five is a one-instance check, and a one-instance check is not a check.

FM 4.4 Averaging a split trajectory

The Singapore assessment's convention. Where a jurisdiction's degradation trajectory genuinely splits — Singapore is stable to improving on institutional quality, improving on regulatory credibility, and eroding on bloc independence — averaging produces "stable," which is exactly wrong on the axis the structure is buying. Split it and refuse to average.

FM 4.5 Confusing the ceiling with the building

The procedure raises `N`. Raising `N` raises the ceiling on `h`. It does not pay `C`, and at the tested parameters `C` plus `delta Q` exceeds `eta ln(N)` for both structures. A reader who runs the procedure and expects a positive number has misunderstood what the procedure does.

FM 4.6 Buying for the median

Every element of the architecture costs money every year and pays in a year that may not arrive. That is the shape of the product and it is why the Hong Kong assessment recommends against the move with the largest measured gain in the series: siting the Asian centre in Hong Kong raises `N` from 2.941176 to 3.615826 and `h` by 0.072279, the largest single-move improvement anywhere, and the Hong Kong assessment declines it and takes capped transit instead at a gain of 0.002869 — **3.97% of the benefit for 3.21% of the exposure.** That trade only makes sense if the purchase is for the tail.

FM 4.7 Stopping at the win

The succession analysis's discipline. The forced-heirship exposure is capped at USD 6.70 million and is zero on the designed map; the French inheritance-tax exposure is approximately USD 27.00 million, 4.0299 times larger, and untouched by the structure. A report containing only the first number is not a summary. It is a misrepresentation by selection, and the same discipline applies to every output of this procedure.

FM 4.8 Running it once

Steps 1 through 6 produce a snapshot. Step 7 is what makes it a measurement. Without the annual re-run, the bloc calls become beliefs, the job vectors become fiction, the Verification Register becomes a historical document, and the structure quietly reverts to whatever the advisers' individual preferences produce.

4.5 Results: the procedure run end to end on the carried case

CONSTRUCTED. The carried case of the carried constructed case, run once through the procedure end to end, with every figure recomputed.

Step 0. Threat named: arbitrary action and bloc fragmentation are the primary exposures; succession is a live secondary; lawful process is not currently in play and the structure is being built years ahead of any claim, which is the only configuration in which it helps; commercial and counterparty risk is carried by the captive and the insurance layer.

Step 1. As found: UK 40.9, US 27.3, Jersey 18.2, Luxembourg 13.6. Naive HHI 0.293430, 3.407968 effective jurisdictions. Collapsed: WEST 100.0, bloc HHI 1.000000, `N = 1.000000`, `h = \ensuremath{-}0.500000`. As designed: UAE 29.1, Thailand 18.2, UK 16.4, Bermuda 13.6, India 11.8, Singapore 10.9. Naive HHI 0.189002, 5.290949 effective jurisdictions. Collapsed: BRIDGE 40.0, NONALIGNED 30.0, WEST 30.0, bloc HHI 0.340000, `N = 2.941176`, `h = \ensuremath{-}0.122417`. Western-perimeter exposure 100.0% → 30.0%. Break-even `N* =`

4.172734. Rails measured separately: 1.000000 effective clearing system as found, 2.810568 as designed.

Step 2. Seven entities scored. Five specialists or acceptable strains; one strain accepted in writing (Bermuda captive, 2.1739, inherent to a captive); one flagged as overloaded (GIFT City, 2.6316) and carried as an open design error rather than corrected on paper.

Step 3. Every entity passes. The apex is a filed legal person with a published register position; the Thai operating company is BOI-promoted with no nominee element, the nominee route being criminal under the Foreign Business Act B.E. 2542 (1999) and excluded on that ground before any other; the captive writes real third-party risk with a loss reserve specialist opinion; the United Kingdom company is on the public PSC register and files full accounts.

Step 4. One marginal jurisdiction priced: the AIFC at 5% from the United Kingdom slot. `G = 0.026731`, `L = 0.005110`, `p* = 16.0476%`, cost ceiling 8.9102% of `C`. Verdict: take it, small, once, and simply, subject to the cost ceiling and to the AIFC register. Two marginal jurisdictions declined on the same instrument: the Qatar family-office vehicle, which duplicates the DIFC inside the same bloc and buys nothing; and the Hong Kong centre, which offers the largest measured gain in the series and is declined in favour of capped transit for the reasons at failure mode 6. One priced loss accepted: the Saudi regional headquarters at 5.0% funded from the operating allocation, costing 0.124275 in `N` and 0.015110 in `h`, taken for addressable revenue, with the number recorded in the file.

Step 5. Leverage layer sized on drawdown rather than on loan-to-value. At the carried case's parameters the trigger is 23.0769% and is judged too tight; at a 50% drawdown tolerance the borrowing falls from 40.000 to 26.000, a reduction of 14.000, which is the stated price of the non-event. Liquidity modelled once against both the margin cure period and the succession requirement, which showed a USD 33.84 million shortfall, 14.1% of the estate.

Step 6. Counsel instructed in the UAE, the United Kingdom, Thailand, Bermuda, India, Singapore, and Kazakhstan; separately in France, Germany, and the two other heir jurisdictions; separately Sharia-qualified counsel on the personal-status interaction. Written CRS and FATCA classification of every entity at formation. The AIFC register run in full before the AIFC leg is funded, and the foundations question — whether the AIFC foundation confers separate legal personality — treated as unanswered rather than assumed.

Step 7. Diarised, named, and dated. Bloc calls recorded with triggers: a second unilateral Singapore measure adopted without Security Council authority; a renewed FATF listing of the UAE; a designation of a major Kazakh bank; a Judicial Tribunal line making DIFC orders unreliable.

The result. `h = \ensuremath{-}0.122417` before the AIFC leg, `\ensuremath{-}0.095686` after, against a break-even of zero. The structure closes 21.8357% of its shortfall with its best available marginal move and remains below the build threshold.

---

5. What is unresolved

This is the end of the method and it is not a summary of achievements. Eight things remain open, and the first of them is the largest.

5.1 Both structures fall below the build threshold

At the tested parameters — `eta = 0.35`, `C = 0.30`, `Q = 10`, `delta = 0.02` — the conventional structure gives `h = \ensuremath{-}0.500000` and the designed structure gives `h = \ensuremath{-}0.122417`. Both are negative. The break-even is `N* = exp(0.5 / 0.35) = 4.172734` effective blocs, and the designed structure reaches 2.941176. Adding the best available marginal position takes it to 3.174603, which closes 21.8357 per cent of the gap and does not

close it. The assembled-architecture paper reaches the same verdict on the same structure [Curzi, 2026]: $N = 2.941176$, $h = 0.35 \ln(2.941176) - 0.500000 = -0.122417$ against a break-even of $N^{*} = 4.172734$, closing 75.5167 per cent of the shortfall in h and **not closing it.**

The improvement is real and it is large: 0.377583 of `h`, purchased by going from one effective bloc to 2.941176, which is the difference between a structure with no surplus term at all and one with a substantial one. It is also, by construction, subject to diminishing returns — `dh/dN = eta/N`, so the move from 2.941176 to 4.000000 buys 0.107620, under a third as much as the move from 1.000000 to 2.941176. **Diversification raises the ceiling. It does not pay the maintenance bill.**

Three honest ways to close the remaining gap exist and none of them is jurisdictional. Raise `eta` — convert more income into durable protected holdings, which is a discipline question rather than a structuring one. Lower `C` — which at some point means holding fewer legs, and therefore lowering `N`, and the trade has to be computed rather than assumed. Lower `delta` — which is what the jurisdiction papers are mostly actually doing. **The reader who wants a positive number should look at `eta` and `C` first, and this series is not primarily about either of them.**

5.2 Disclosure is not immunity

The disclosure paper argues that a structure is durable when observation yields nothing actionable, and it carries its own strongest counter-case as its seventh failure mode rather than leaving it to a critic. In 2022, approximately USD 300 billion of Russian central bank reserves were immobilised. **That event is described from the source manuscript and is cited nowhere, because this paper's bibliography contains no contemporaneous reporting**; the corpus paper that owns the result states it in the same terms and with the same gap disclosed [Curzi, 2026]. Those assets were fully disclosed. Their location, custodians, and character were all known and were the reason they were held where they were held. **They were frozen anyway.**

What survives is narrower than the incautious reading of the disclosure paper. Disclosure removes the failure mode in which a structure collapses on contact with an authority, converts an enforcement problem into a policy problem, and — through bankability, which the concentration paper identifies as the binding constraint on everything else — makes the rest of the architecture operable at all. It does not supply an appeal and it does not protect an asset sitting inside the perimeter that took the decision. The answer to that threat is `N`, and `N` is not a property of the paperwork — which is the disclosure paper's own economics of full disclosure, whose correlation result prices the alternative: holding marginal failure at 10 per cent, a common shock of $c = 0.02$ takes the probability that four perimeters fail together from 0.000100 to 0.020044, a 200.4-fold understatement [Curzi, 2026].

5.3 The succession structure solves the smaller of the two exposures

The succession analysis's constructed case, and the most uncomfortable result in the series. The forced-heirship exposure for the child resident in the civil-law jurisdiction is capped at USD 6.70 million and is zero on the designed map. The French inheritance-tax exposure is approximately USD 27.00 million — 4.0299 times larger — and is untouched by the structure. The liquidity test shows a USD 33.84 million shortfall, 14.1% of the estate, sufficient to force the sale of an operating branch.

A succession section reporting only the heirship win would be lying by omission, and a manual closing on it would be doing the same thing at greater length. What the structure does is eliminate a fragmentation risk and a control risk. What it does not do is pay a tax, and the tax is the bigger number.

The one lever that does reach the whole problem is present in all three legal families the case touches — advance renunciation of the action in reduction in France, the *Pflichtteilsverzicht* in Germany, *takharuj* among heirs under Sharia — and it is the same lever in each: **consent. And consent requires disclosure**, which ties the succession answer back to the

disclosure paper's governing principle and is the reason the two parts of this series are one book.

5.4 The dollar share cannot be engineered away

The rail-perimeter measurement, stated in the form that matters. Correspondent concentration measured by terminal clearing system moves from 1.000000 to 2.810568 under the designed allocation [Curzi, 2026], whose clearing-perimeter measurement takes a naive correspondent HHI of 0.302400 — 3.306878 effective correspondents — to a perimeter HHI of 0.745256 and $N_{\text{rail}} = 1.341819$ once flows are decomposed by currency. Reaching 4.000000 effective clearing systems, holding the non-dollar mix in its current proportions, would require the dollar share to fall to approximately 38.99% from 55%.

That is a commercial decision about which currencies the group invoices in, not a legal decision about where it banks. Four hard limits sit underneath it: a USD account settles under US law wherever it is booked; CIPS is a clearing system and many participants still use SWIFT for messaging, so "CIPS replaces SWIFT" is false and any structure built on it will fail; the bank's risk appetite is not the jurisdiction's policy; and a rail is not a holding. The rail concentration is therefore only partly a structuring problem and mostly a revenue-currency problem, and this series cannot solve the second.

5.5 Cost is not the reason to choose the foundation

The computed foundation-versus-trust crossover, which cuts against the way foundations are usually sold. Against a trustee charging on a basis-point scale, the DIFC Foundation is cheaper only above USD 34.3 million at 35 basis points, USD 80.0 million at 15 basis points, and USD 120.0 million at 10 basis points.

So cost is explicitly not the reason to choose it. The reason is recognition: a foundation is a legal person and travels into civil-law and Sharia jurisdictions in a way a common-law trust does not, and the succession analysis's whole analysis depends on that and not on a fee comparison. Any adviser recommending the foundation on cost has either a very large client or a weak argument.

5.6 What the Gulf bridge actually delivers, and what it does not

The ADGM decomposition, which is the most useful single sentence among the Gulf-centre assessments. **English law is three things — doctrine, court, and territorial state — and ADGM buys the first two without the third.** That is precisely what a bridge jurisdiction can deliver and precisely what the City of London cannot, because in London the three are inseparable by definition.

What remains unresolved is whether the second of the three holds up under a stress that has not yet occurred. The ADGM Courts have a short record, the enforcement bridge to the onshore courts is documented but lightly tested, and the conduit question — whether a foreign judgment recognised inside a free zone can be enforced onshore through it — is carried unresolved from the foundation paper's register through the rail measurement's. It is the largest open legal question in the architecture and it sits directly under the apex.

5.7 The bloc assignments are the framework's own, and they are contestable

Every value of N in this synthesis rests on a mapping from jurisdictions to perimeters, made by applying the threat register's test to published evidence. That mapping is an argument, not a measurement. The concentration paper's robustness table exists to show how much the answer moves under reasonable disagreement, and the answer is: not much. Reclassifying

Singapore as Western moves `N` by 0.016874 and `h` by 0.002014 [Curzi, 2026], whose bridging-slot control finds the same 10.9 per cent held in Singapore or in London giving identical blocs and a stress cost of 0.002014. Separating the UAE and Singapore into distinct blocs moves `N` from 2.941176 to 3.615826 — which is a large move, and is the reason the published figure of 2.941176 is the conservative end of a range rather than a point estimate.

A reader who disagrees with a bloc call should re-run the arithmetic rather than argue, and will usually find the disagreement does not matter. A reader who finds one that does matter has found something the framework wants to know.

5.8 Fifteen sections of this synthesis are specified and not yet on the page

Stated because the alternative is to imply a completeness the work does not have. The manifest specifies forty-five sections and six appendices. The sections written and verifiable on the page at the time of this compression cover the greater part of the series; the Convergence section, US domestic, GIFT City, Türkiye, asset conversion, SPVs, the three rail assessments, the leverage and group architecture sections, law-family arbitrage, the integrated wiring assessment, the sanctions and compliance section, and the banking-lineage sequence are specified in the manifest and are not among them.

Where this section and the appendices refer to those sections, they refer to results fixed in the reconciliation ledger — the 23.0769% margin trigger, the four bright lines, the CIPS distinction, the 2022 precedent — and not to text that can be checked. **Appendix D marks the gap explicitly rather than papering over it, and a reader should treat every forward reference to an unwritten section as a specification rather than as a citation.**

That is the method. It measures concentration honestly, assigns one job per entity, refuses structures that depend on not being seen, prices marginal jurisdictions instead of arguing about them, sizes leverage from the drawdown rather than the loan, and sends every element to counsel in every jurisdiction it touches. Run properly on the carried case it produces a structure with 2.941176 effective blocs, 30% Western-perimeter exposure instead of 100%, 2.810568 effective clearing systems instead of one, one flagged design error left visible, a priced loss taken knowingly, and a sustaining rate of −0.095686.

Below the threshold, honestly measured, with the residual stated. That is what the instrument says, and the instrument is the only part of this series that was never going to flatter anyone.

5.9 The verification register for this synthesis

Marked for confirmation rather than asserted, and gathered for Appendix D.

1. Every figure in this section is derived from the carried case's stated test parameters and the allocations in the concentration, rail, Gulf, Singapore, Hong Kong and AIFC assessments. **They are illustrative model outputs, not estimates of any real position, and no reader's own structure will produce them.**
2. The parameter values `eta = 0.35`, `C = 0.30`, `Q = 10`, `delta = 0.02` are the carried constructed case's test parameters. A real structure requires each to be estimated from its own facts, and `C` in particular is routinely under-estimated because it is paid in small recurring amounts to many parties.
3. All bloc assignments are the framework's own, applying the threat register's perimeter test to published evidence. Contestable; sensitivity shown in the concentration paper's robustness table; each jurisdiction section restates its assignment with evidence at scorecard item 2.

4. The job vectors in step 2 are estimates and should be re-derived annually from where board resolutions are actually passed and where management time is actually spent, not from the organisation chart.
5. The 35% maintenance requirement, the 50% loan-to-value, and the drawdown table in step 5 are **constructed model parameters**. Confirm actual maintenance definitions, cure periods, revaluation discretion, and rehypothecation rights in the specific margin documentation with counsel.
6. Thailand: the Foreign Business Act B.E. 2542 (1999) nominee provision and its enforcement practice; Board of Investment promotion conditions, which are activity-specific and change; and Revenue Department order Por. 161/2566. Carried from the Thailand assessment.
7. DIFC Foundations Law No. 3 of 2018 and ADGM's Application of English Law Regulations 2015 are cited for the characteristics stated. Confirm current text and amendments with counsel in the relevant centre.
8. The AIFC foundations question — whether the AIFC foundation confers separate legal personality — is carried as **unanswered** from the AIFC assessment and must be resolved before any apex migration is contemplated. The AIFC assessment names and numbers no AIFC subordinate Act, Regulation, or Rule, and this section does not either.
9. The conduit question — whether a foreign judgment recognised within a UAE financial free zone can be enforced onshore through it, and the current constitution and practice of the Judicial Tribunal — is carried unresolved from the foundation paper, Register item 8, through the rail-perimeter measurement, Register item 4.
10. Singapore: the precise instruments by which the March 2022 targeted financial measures and export controls were imposed, and their stated legal basis. The bloc-trajectory finding in this section rests on that single dated episode.
11. The 2022 immobilisation of approximately USD 300 billion of Russian central bank reserves: the instruments by jurisdiction, dates, amounts, and the current legal position on the assets and the income arising from them. Treated as approximate throughout.
12. The succession analysis's succession figures — the USD 6.70 million capped reserve exposure, the approximately USD 27.00 million French inheritance-tax exposure, and the USD 33.84 million liquidity shortfall — are constructed outputs at a constructed exchange rate and constructed rate scale. Confirm the French residence conditions, abatements, and rate scale, and whether a discretionary foreign foundation making no distribution produces a taxable transmission at all, with French counsel.
13. The foundation-versus-trust comparison's crossover figures (USD 34.3m at 35 bp, USD 80.0m at 15 bp, USD 120.0m at 10 bp) rest on a constructed fixed-fee comparator of USD 120,000 per annum. They are a method, not a quote.
14. The correspondent measurement in the rail-perimeter measurement is by terminal clearing system, which requires identifying the terminal system for each chain — not the bank, and not the frameworking jurisdiction. Confirm direct versus indirect CIPS participation for each institution actually used.
15. The sections listed at item 8 of "What Is Unresolved" are specified in the manifest and are not written on the page at the time of this compression. Every forward reference to them in this section and in the appendices is a **specification, not a citation.**

All numerical results in this section were computed from the stated inputs at full precision in Python and rounded only for printing. No figure was derived from a rounded print. Appendix B reproduces the workbook in full.

Not legal, tax, or investment advice. Confirm every element with qualified counsel in every jurisdiction the structure touches. All worked cases are CONSTRUCTED.

5.10 What the appendices contain

> *Not legal, tax, or investment advice.* These appendices consolidate the framework's tables, its arithmetic, its diagrams, its citations, and its admissions. They are a map, not a permission. Confirm every element with qualified counsel in every jurisdiction any structure touches. All worked cases are **CONSTRUCTED** and labelled as such. Appendix D marks every citation used anywhere in the method as either VERIFIED or FLAGGED FOR VERIFICATION, and Appendix E states, in one place, everything this series does not know.

Attributed to The Invisible College of The Lodge Between Lodges, Supreme Council of Hands. The full roster with its designations is reproduced verbatim on the title page.

There are six appendices and they do six different jobs.

```
A  SCORECARD ATLAS          every jurisdiction assessed in the series, one table,
                            all twelve items, sortable by job and by trajectory
B  CONCENTRATION WORKBOOK   the HHI and effective-bloc arithmetic worked end to
                            end, every figure recomputed
C  STRUCTURE DIAGRAMS       the carried case in ASCII, each entity's job labelled
D  STATUTE AND CASE INDEX   every citation, marked VERIFIED or FLAGGED
E  THE HONESTY APPENDIX     settled / contested / needs counsel / recently changed
F  CORRESPONDENCE           how the earlier series map onto this series
```

Appendix D is the framework's integrity mechanism and Appendix E is its conscience. A reader with time for only two should read those two.

---

6. Verification and rigour audit

This section is the one place in the paper where something can be checked in under a minute, and it is built so that a reader can do it. The instrument is decomposed into postulates, and the chain from stipulated shares to the reported surplus is given step by step.

6.1 Dimensional audit

symbol	quantity	dimension	unit	status
w_i	allocation share	[1]	fractions summing to 1	stipulated
HHI	$\sum_i w_i^2$	[1]	dimensionless	derived
N	effective count, $1/\text{HHI}$	[#]	perimeters or jurisdictions	derived
$\ln N$	surplus term	[1]	nats	a logarithm of a count
η	conversion of flow into stock	[1]	—	**stated dimensionless; see below**
δ	erosion per period	[1]	per period	a fraction
Q	protected capital	[1]	**index units, no currency**	stipulated

symbol	quantity	dimension	unit	status
C	maintenance per period	$[1]$	index units per period	stipulated
h	sustaining rate	$[1]$	index units per period	derived
d^*, m, LTV	drawdown, maintenance, loan-to-value	$[1]$	fractions	stipulated / derived
basis points, USD figures	fee scales and estate values	[money]	USD	**stipulated, constructed**
item counts	register entries per assessment	$[\#]$	counts	enumerated

Does it balance? $h = \eta \ln N - \delta Q - C$ balances only if $\eta \ln N$, δQ and C share a unit. δQ is index units per period; C is stated as index units per period; and $\ln N$ is dimensionless, so η **must carry index units per period per nat while also being described as a dimensionless conversion efficiency in** $[0, 1]$**.** Those two descriptions are not the same, and **the source is silent on the reconciliation.** The reading adopted here is that η carries the conversion from a nat of surplus to a unit of stock per period, which makes the equation homogeneous and leaves Q in **index units with no currency attached** — an honest dimensionless placeholder for a quantity nobody has measured, and not an invented unit.

Every headline result survives that gap because every headline result is a ratio or a difference in the same units. N, the effective counts, the percentage shares, the break-even N^*, the ratio $0.107620/0.377583$, the crossover thresholds and the succession ratio 4.0299 are all dimensionless or in stated currency. **The one quantity whose absolute value means nothing without a unit is** h **itself, and it is used only for its sign and for differences.**

6.2 Significant figures and epistemic status

- **0.293430, 3.407968, 0.189002, 5.290949, 0.340000, 2.941176, 1.000000 are exact to the digits shown**, being sums and reciprocals of exactly stipulated shares. They carry no interval because nothing was estimated.
- $h = -0.500000$ **and** $h = -0.122417$ **are exact given** $\eta = 0.35$**,** $\delta = 0.02$**,** $Q = 10$**,** $C = 0.30$**, and those four are stipulated test parameters and estimates of nothing.**
- $N^* = \exp(0.5/0.35) = 4.172734$ **is exact given the same four**, and moves with them; a reader who changes η changes the threshold.
- **Six decimal places on a figure derived from shares given to one is a statement about the arithmetic and not about the world.** The precision rule exists so that two readers recomputing from the same inputs get the same digits, which is what makes an error visible. **It does not mean the sixth decimal place carries information about any structure.**
- **The crossover thresholds and the succession figures are exact on constructed inputs.** USD 34.2857 million is $120{,}000/0.0035$; the succession ratio 4.0299 is $27.00/6.70$. Neither is a measurement.
- **The register item counts are exhaustive enumerations** over the source units and are regenerable by the extraction Appendix D.2 describes.

6.3 Limits and boundary behaviour

$N \to 1$**.** $\ln(1.00) = 0$, so the surplus term vanishes exactly and $h = -(\delta Q + C)$. **This is not a small surplus, it is none**, and it is the whole organising result of the series.

$N \to \infty$**.** $\eta \ln N$ grows without bound while C grows at least linearly in the number of legs, so there is always a count beyond which another leg loses money. **The break-even is**

therefore a ceiling on useful fragmentation and not a target, and the cost ceiling in the break-even probability method is where the trade is actually decided.

$\mathrm{d}h/\mathrm{d}N = \eta/N$**.** Marginal value falls monotonically: 0.350000 at $N = 1$, 0.119000 at $N = 2.941176$. **The first move is worth 3.5087 times the second, and a reader with capacity for one leg should add the first and stop.**

$\eta \to 0$**.** The surplus term vanishes at every N, and no amount of jurisdictional diversity helps a structure that converts nothing into durable form. **The threshold is a statement about η and C at least as much as about N, which is §5.1's conclusion.**

$d \to 0$ **in the leverage relation.** $L = V(1-d)(1-m)$ gives $L = V(1-m)$ at zero drawdown tolerance — the position is at its call level the day it opens — and $L \to 0$ as the tolerance approaches one. **The usable region is the middle, and the constructed case sits in it.**

6.4 Invariance

The effective count is invariant under relabelling and under permutation of the shares, being a function of the multiset of shares alone. That is what makes it a numbers-equivalent, and it is also its principal limitation: **a statistic that does not depend on which entity holds which share cannot tell a reader which entity to change.**

It is not invariant under regrouping, and the whole method lives in that fact. The transformation that moves the number is the partition of jurisdictions into perimeters, and the partition is a judgement. **The sensitivity is measured rather than asserted**: reclassifying Singapore moves N by 0.016874 and h by 0.002014, while separating the UAE and Singapore into distinct perimeters moves N from 2.941176 to 3.615826. **The instrument is robust to being wrong about one member and not robust to being wrong about the structure**, which means the argument a reader is most likely to have is not the one that matters.

One invariance the instrument does not have, and it is the sharper limitation. Because $u = 1 - \rho^2$ depends only on $|\rho|$, a count of independent things reads perfect alignment and perfect anti-alignment identically [Curzi, 2026]. **Correlation of either sign collapses the count**, and nothing in the arithmetic distinguishes coordination from opposition; only the perimeter test, argued from evidence, does.

6.5 Analogy discipline: every borrowing, graded

borrowing	used for	category
Concentration indices and the numbers-equivalent [Hirschman, 1964, Herfindahl, 1950, Adelman, 1969]	every effective count in the paper	**1 — derivation.** The shares are a partition and the index is defined on it
Hill numbers [Hill, 1973, Jost, 2006]	naming $1/\mathrm{HHI}$ as the order-two member	**1 — derivation**
Merger-guideline thresholds [U.S. Department of Justice and Federal Trade Commission, 2023, 2010]	that the effective-count reading is the standard apparatus	**3 — analogy.** The thresholds are for markets, not for holdings, and none is applied here

borrowing	used for	category
The storage law as a leaky integrator [Curzi, 2026,?]	the whole of h	**2 — shared form.** The same difference equation, with no claim that any physical result transfers
Correlated failure across perimeters [Buldyrev et al., 2010, Albert et al., 2000, Cohen et al., 2000]	why nominally independent jurisdictions fail together	**3 — analogy**, invoked for shape; the empirical premise of the collapse step is asserted, not measured
Tail correlation and portfolio covariance [Longin and Solnik, 2001, Ang and Chen, 2002, Markowitz, 1952]	that a count is not a diversification	**2 — shared form**; no portfolio is constructed and no return appears
Ecological stability–diversity results [May, 1972, Tilman et al., 1998, Doak et al., 1998, McCann, 2000]	that more components need not mean more stability	**3 — analogy**
Open design [Kerckhoffs, 1883, Saltzer and Schroeder, 1975, Brandeis, 1914]	the disclosure test of step 3	**1 — derivation** as a design principle; the last is a policy argument and is **3**
Registers and preregistration practice [Nosek et al., 2018, Simmons et al., 2011, Gelman and Loken, 2013, Ioannidis, 2005]	why the default status is FLAGGED	**3 — analogy.** A verification register is not a preregistration and is not claimed to be one
An index adopted as a target [Goodhart, 1975, Campbell, 1979]	the warning attached to publishing N as a design objective	**3 — analogy**, with a mechanical rather than a behavioural cause
Resilience as a property of a system rather than of a component [Holling, 1973, Hollnagel et al., 2006, Leveson, 2011]	why the annual re-run is part of the method	**3 — analogy**
Coordination cost rising with the number of parts [Brooks, 1975, Conway, 1968]	why C grows at least linearly in the legs	**2 — shared form**
Research-programme badging [Lakatos, 1970, Platt, 1964, Meehl, 1990]	what a badged postulate is doing	**3 — analogy**
The grammar of conserved quantities [Mirowski, 1989, Cartwright, 1999]	flagged against the framework, not borrowed	**not a borrowing — a warning**

The three sentences the corpus requires. *The theorem:* $\ln(1) = 0$, so a single-perimeter structure has a surplus term of exactly zero — arithmetic, and checkable in a second. *The ordinary effect:* jurisdictions that coordinate tend to act together, which is an ordinary observation about policy and needs no formalism. *The conjecture:* that a structure's resilience against arbitrary action is a computable function of its perimeter geometry, and that h mea-

sures whether the structure sustains itself, is **this framework's own proposal**, is not derived from anything cited, and is not established by any result in this paper.

6.6 The postulates

	postulate	status
P1	The stipulated shares, parameters and fee scales are as printed.	DEFINITION (constructed)
P2	A coordination perimeter is a set of jurisdictions that act together, so their shares add before squaring.	POSTULATED — the framework's own test; **contestable, and the load-bearing input**
P3	$1/\text{HHI}$ is the right reading of an effective count.	CITED [Hirschman, 1964, Hill, 1973]
P4	$h = \eta \ln N - \delta Q - C$ is the right steady-state condition.	MODELLING CHOICE, inherited [Curzi, 2026]
P5	Jurisdictions inside one perimeter fail together.	**POSTULATED — the empirical premise of the collapse step, and it is asserted rather than measured**
P6	The legal citations consolidated in Appendix D exist and say what the source says.	**DESCRIBED — uncited; FLAGGED by default**

Dropped as unnecessary. No postulate about which jurisdictions, which entities or which currency is needed for any result: every figure is a property of a partition and four stipulated parameters, and survives the substitution of any other labels.

6.7 The chain a reader can check in a minute

step	consumes	produces	status
L1	P1	conventional shares 40.9 / 27.3 / 18.2 / 13.6; $\sum w_i^2 = 0.293430$	ARITHMETIC
L2	L1, P3	naive effective jurisdictions $= 3.407968$	ARITHMETIC
L3	L1, P2	one perimeter at 100.0; bloc HHI $= 1.000000$	ARITHMETIC given P2
L4	L3, P3	$N = 1.000000$; $\ln N = 0$	ARITHMETIC
L5	L4, P1, P4	$h = 0.35 \times 0 - 0.02 \times 10 - 0.30 = -0.500000$	ARITHMETIC

step	consumes	produces	status
L6	P1	designed shares; $\sum w_i^2 = 0.189002$; $1/\text{HHI} = 5.290949$	ARITHMETIC
L7	L6, P2	perimeters 40.0 / 30.0 / 30.0; bloc HHI $= 0.340000$	ARITHMETIC given P2
L8	L7, P3	$N = 2.941176$; $\ln N = 1.078810$	ARITHMETIC
L9	L8, P4	$h = 0.35 \times 1.078810 - 0.500000 = -0.122417$ — **R2**	ARITHMETIC
L10	P4	$N^* = \exp(0.500000/0.35) = 4.172734$	ARITHMETIC
L11	P4	$\mathrm{d}h/\mathrm{d}N = \eta/N$; 0.377583 then 0.107620 — **R3**	ARITHMETIC
L12	P1	rail shares; 1.000000 as found, 2.810568 as designed — **R4**	ARITHMETIC
L13	P1	$120{,}000/0.0035 = 34.2857$m; $/0.0015 = 80.0000$m; $/0.0010 = 120.0000$m — **R5**	ARITHMETIC
L14	P1	6.70m capped; $27.00/6.70 = 4.0299$; $33.84/240.00 = 14.1000\%$ — **R6**	ARITHMETIC

Two checks take a reader under a minute each. $0.409^2 + 0.273^2 + 0.182^2 + 0.136^2 = 0.293430$, and its reciprocal is 3.407968. Then $0.400^2 + 0.300^2 + 0.300^2 = 0.340000$, whose reciprocal is 2.941176, and $0.35 \ln(2.941176) = 0.377583$. **Everything else in the workbook is the same two operations repeated.**

6.8 What this chain is and is not

Read the status column. L3 and L7 consume P2 — the perimeter partition — and **every headline number in this paper depends on it.** P2 is a judgement argued from evidence in the jurisdiction papers and not a measurement, and P5, the premise that jurisdictions inside one perimeter fail together, is **asserted rather than measured anywhere in the corpus.**

So the chain establishes that the arithmetic is correct and establishes nothing about the world. A reader who accepts the partition gets the numbers; a reader who rejects it gets different numbers from the same procedure, which is the property that makes the procedure useful. **The instrument prices a disagreement rather than settling it**, and §6.4 gives the size of the disagreement that would matter.

7. Limitations

This paper is a compression and every cell in it is smaller than the assessment it compresses. No cell in the atlas should be relied on without reading the paper it points to and the verifica-

tion register that paper carries, and §A.5 says so at the point of use.

The perimeter partition is the load-bearing input and it is a judgement. Every value of N, and therefore every value of h, rests on it. It is argued from published evidence, it is contestable, and its sensitivity is measured (§6.4) rather than assumed away.

The premise that jurisdictions inside one perimeter fail together is asserted and not measured. It is the empirical claim on which the collapse step depends, the network literature that would make it plausible is cited for shape rather than for evidence [Buldyrev et al., 2010, Albert et al., 2000, Cohen et al., 2000], and **no correlation between jurisdictional outcomes is estimated anywhere in the corpus.**

Publishing N as an objective would change what it measures. An indicator adopted as a target ceases to be the indicator it was [Goodhart, 1975, Campbell, 1979], and a structure assembled to raise a count rather than to raise independence would satisfy the instrument while defeating its purpose. **The method has no defence against that except the perimeter test being argued from evidence each time**, and §5.7 says the same thing about the assignment.

The four test parameters are stipulated and drive every reported surplus. $\eta = 0.35$, $C = 0.30$, $Q = 10$, $\delta = 0.02$ are constructed values. **A real structure requires each estimated from its own facts, and C is the term most likely to be wrong in the direction that flatters the answer**, because it is paid continuously in small amounts to many parties and never appears as a line item — and because coordination cost across parts grows at least linearly and often faster [Brooks, 1975, Conway, 1968].

No legal source is cited anywhere in this paper. Appendix D consolidates several hundred statutes, conventions, court rules and decided cases; **its default status is FLAGGED FOR VERIFICATION and only definitional and arithmetic facts are marked verified.** This paper's bibliographic database contains none of the primary sources, no citation is attached to any of them, and **no approximate citation has been substituted for a missing one.** That is a real gap and it is not closed here.

The item counts, the atlas and the index are as current as the source that generated them. The extraction is regenerable and a later revision that yields different counts has found a change rather than an error, which Appendix D.2 states in place.

The dimensional audit does not close for η (§6.1), and Q carries index units with no currency attached, so **the absolute value of h means nothing**; only its sign and differences in it are used.

No statistical claim is made and none is supported. There are no p-values, effect sizes, intervals, samples or fitted models in this paper. No model comparison, out-of-sample test, matched null or information criterion is reported, because nothing was fitted, and no distributional or scaling claim is made, so Clauset et al. [2009] is not applicable and was not run.

And the paper's own status limits what any of it is worth. It advances no testable hypothesis (§2). Its two new computations are arithmetic on constructed inputs. **Its most useful content is the register of four admissions in §5, all of which are negative, and the FLAGGED default in Appendix D, which is a statement about how much of the series has not been verified.**

8. Conclusion

The method reduces to seven steps and the first one carries the whole result. Collapse jurisdictions onto coordination perimeters before measuring, because $\ln(1.00) = 0$: **a structure with one effective perimeter has zero available surplus, not a low one**, and no rearrangement of entities inside that perimeter changes it. Then assign one job per entity, apply the disclosure test, price each marginal jurisdiction as a break-even probability rather than arguing it, size leverage from the drawdown the structure must survive rather than from the loan, verify and record where the framework refused to guess, and re-run the whole thing annually with a named owner and a date.

Reproduced end to end, the arithmetic is this. The conventional structure reads naive HHI **0.293430** and **3.407968** effective jurisdictions but bloc HHI **1.000000**, giving $N = 1.000000$ and $h = -0.500000$. The designed structure reads naive HHI **0.189002** and **5.290949** effective jurisdictions against bloc HHI **0.340000**, giving $N = 2.941176$ and $h = -0.122417$, with Western-perimeter exposure falling from **100.0 to 30.0 per cent**. Break-even is $N^{*} = \exp(0.5/0.35) = 4.172734$, and because $\mathrm{d}h/\mathrm{d}N = \eta/N$ the move from one perimeter to 2.941176 buys **0.377583** while the move from 2.941176 to 4.000000 buys **0.107620** — **28.5 per cent as much.** The rail axis is measured separately and reads **1.000000** effective clearing system as found against **2.810568** as designed, with four effective systems requiring the dollar share to fall from 55.00 to approximately **38.99 per cent** — a level set by revenue currency and not by structure.

Two computations are given here in full for the first time and both cut against the way the instruments they concern are usually sold. Against a constructed USD 120,000 fixed mandate, a foundation is cheaper than a trustee charging ad valorem only above **USD 34.2857 million at 35 basis points, USD 80.0000 million at 15, and USD 120.0000 million at 10** — so **cost is explicitly not the reason to choose the form**; recognition in civil-law and Sharia jurisdictions is. And the succession residual is an asymmetry rather than a win: a notional reserved share of USD 60.00 million capped to a **USD 6.70 million** compensatory levy — 11.1667 per cent — and eliminated by design, against an inheritance-tax exposure of approximately **USD 27.00 million**, **4.0299 times larger**, that the structure does not touch, and a liquidity shortfall of **USD 33.84 million**, **14.1000 per cent of the estate**, sufficient to force the sale of an operating branch.

The four admissions, stated without qualification, are the paper's most useful content. Both the conventional and the designed structures remain below the build threshold: diversification raises the ceiling and does not pay the maintenance bill. Disclosure did not prevent the 2022 immobilisation: it removes the failure mode in which a structure collapses on contact with an authority and makes the structure bankable, and it supplies no appeal. Cost is not the reason to choose a foundation. And the succession structure addresses the smaller of the two exposures.

What this paper is. Reference apparatus: a consolidated procedure, a reproduction workbook whose every input is printed, an atlas keyed to the papers it compresses, an index whose **default status is FLAGGED FOR VERIFICATION**, a register of the corpus's own corrections, and a correspondence appendix that names four things which do not transfer. **It advances no hypothesis and establishes nothing new about the world.** Its contribution is that after it, every number in Series D can be recomputed by a reader who disagrees with it — and the register of what the series has not verified is longer than the register of what it has.

Declarations

Data and code availability

No external dataset was collected, and none is reported. Every result in this paper is exact arithmetic on stipulated, **CONSTRUCTED** inputs printed at the point of use: two jurisdiction allocations with their perimeter assignments, four test parameters, a fee scale, an estate value, a pledged position, and a set of scored job vectors. Appendix B reproduces the workbook in full, with every input, every intermediate value and the precision rule under which they were computed. **No real family, entity, transaction, market or institution is described anywhere in this paper**, and no allegation is made about any of them. The historical institutions treated from the public record elsewhere in the series are indexed here and not re-described.

This paper contains no figures, and none is claimed; the reproduction is tabular and the diagrams of Appendix C are schematic layouts built from the same stipulated allocations.

No legal source is cited. The statutes, conventions, court rules and decided cases consolidated in Appendix D are described from the source manuscript, are **FLAGGED FOR VERIFICATION** by default, and are attached to no citation because this paper's bibliographic database contains none of them.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflicts of interest.

Author contributions

Single author. Michael Laurence Curzi conceived the analysis, performed all computations, produced all tables, and wrote the paper.

Ethics

No human-subjects research was conducted and no personal data were processed. Every principal, family, entity and transaction described is constructed.

No operational content, and no allegation

This paper contains no operational content of any kind. It supplies no method for evading any legal process, sanctions regime, tax obligation or disclosure requirement. It recommends **no** structure premised on non-disclosure, because such a structure fails the governing principle by construction, and **no** nominee arrangement where nominees are unlawful. **No allegation is made about any person, institution, jurisdiction or state.**

Legal, tax and investment advice

This paper is analytical and is not legal, tax, accounting or investment advice. It describes an accounting and measurement procedure, its inputs, and what the procedure does not do. No recommendation is made in respect of any person, institution, product, instrument or jurisdiction, and no jurisdiction assessed in the atlas is recommended or disparaged. **Counsel is required in each relevant jurisdiction — residence, nationality, beneficiary location, asset situs, and the situs of every counterparty — and every element must be confirmed against the primary sources itemised in Appendix D before any reliance.** The framework is a map and not a permission.

Appendix A. Scorecard atlas

A.1 How to read this atlas

Every jurisdiction paper closes with the twelve-item scorecard fixed by the twelve-item scorecard. This appendix collects all of them into one table so that the comparisons the framework makes section by section can be made column by column instead.

The twelve items, in order:

```
1  Legal family                     7  Beneficial-ownership regime
2  Bloc alignment                   8  Rail access
3  Court quality                    9  Political stability
4  Foreign judgment recognition    10  DEGRADATION TRAJECTORY + evidence
5  Substance requirements          11  Job: HOLDING / OPERATING / BRIDGING
6  CRS / FATCA status              12  Verdict + what would change it
```

Item 10 is the thesis operationalised. Item 11 prevents the commonest structural failure — one entity asked to do three jobs. Items 10 and 11 are the two columns to read first, and item 12 is the only one that tells a reader what to do.

Three conventions apply throughout and each is load-bearing.

The split-trajectory convention. Where item 10 is genuinely non-uniform, the assessment splits and refuses to average. Singapore is the origin — stable to improving on institutional quality, improving on regulatory credibility, eroding on bloc independence — and averaging those to "stable" would be exactly wrong on the one axis the structure buys.

Bloc alignment is a coordination question, not a legal-family question. Ireland is common law and sits in the Western perimeter by construction; Thailand is civil law and is non-aligned. The axes are independent and conflating them is the error the coordination-perimeter analysis was written to prevent.

The cells below are compressions. Each is a pointer into a section that argues the assessment with dated evidence. **No cell in this atlas should be relied on without reading the section it compresses and the Verification Register that section carries.**

A.2 The master table

Thirty-three assessments across twenty-three sections. The DIFC appears twice because the foundation assessment and the bridging assessment score it for two different jobs and reach two different answers, which is itself an illustration of item 11.

Jurisdiction	paper	1 Family	2 Bloc	3 Court	4 Recognition	5 Substance
DIFC (as holding)	053	Codified common-law-derived	BRIDGE	High; DIFC Courts, published precedent	Mixed; onshore via Judicial Authority Law; conduit question open	Low for a holding foundation
United Kingdom	054	Common law, source jurisdiction	WEST, at the centre	Very high; Commercial Court	Mixed, materially worse since Brexit	Real and conventional
Jersey	054	Common law, Norman substrate	WEST	High; Royal Court to Privy Council	Statutory reciprocal + common law	Real from 1 Jan 2019
Guernsey	054	Common law, Norman substrate	WEST	High; Royal Court to Privy Council	Statutory reciprocal + common law	Real from 1 Jan 2019
Isle of Man	054	Manx common law; two company regimes	WEST	High; High Court to Privy Council	Statutory reciprocal + common law	Real from 1 Jan 2019
Luxembourg	054	Civil, deep financial corpus	WEST by construction	Good and improving; fund depth	Strong within EU and EFTA	Real and rising
Ireland	054	Common law, EU member state	WEST by construction	High; Commercial list	Strong; Brussels Recast, Lugano	Real; tightened from 2015
Netherlands	054	Civil; STAK; NCC	WEST by construction	High; Netherlands Commercial Court	Strong; Brussels Recast, Lugano	Real and enforced
Switzerland	054	Civil, cantonal; no domestic trust law	WEST — contestable, resolved on evidence	High; strong arbitration seat	Strong; Lugano party	Real; effective management tests
Bermuda	055	English common law	W-LITE	High; deep insurance bench	Common law + statutory reciprocal	Real; higher for insurers than ESA requires
Cayman Islands	056	English common law	WEST	High; Financial Services Division	Common law + statutory reciprocal	Real from 1 Jan 2019
British Virgin Islands	056	English common law; ECSC	WEST	High and improving	Common law + statutory reciprocal	Real from 1 Jan 2019
The Bahamas	056	English common law; sovereign since 1973	WEST notwithstanding sovereignty	Good; below Cayman and BVI	Common law + statutory reciprocal	Real; CESRA 2018
Barbados	056	English common law; republic 2021	WEST	Adequate to good; CCJ	Common law + reciprocal arrangements	Real, arrived by the harmful-tax route
ADGM	056	English common law by direct application	BRIDGE	High; ADGM Courts	Mixed; onshore enforcement the weak item	Low by design for holding and SPV forms
DIFC (as bridging)	056	Codified common-law-derived; English law residual	BRIDGE	High; strongest published-precedent depth in region	Weak, and weaker for a bridge than a holding	Low passive; real if DFSA-regulated
Qatar / QFC	057	Hybrid, split by regime	BRIDGE (GCC)	Good; QICDRC	Mixed and under-tested	Moderate; may operate anywhere in Qatar

Jurisdiction	paper	1 Family	2 Bloc	3 Court	4 Recognition	5 Substance
Saudi Arabia	057	Sharia-based with codified Royal Decrees	GCC; increasingly multi-aligned	Improving and under-tested	Awards workable; judgments harder	High, deliberately — that is the RHQ product
Singapore	057	Common law, self-consciously divergent	BRIDGE commercially; Western under stress	Very high; SICC	Good but narrower than the reputation	Real and rising
Hong Kong	057	Common law continued by Basic Law Art. 8	Inside two perimeters, independent of neither	High ordinary; qualified extraordinary	Mixed and asymmetric	Now real for tax purposes
AIFC (Astana)	058	English-derived common law by reception	Genuine bridge; least certainty attached	High in design, thin in authority	The weakest item, and it decides the design	Registration + AFSA licensing
Thailand	058	Civil; Civil and Commercial Code	NON-ALIGNED, genuinely	Adequate for its role; not a selling point	Weak — plan around it	High, cheap and genuine
GIFT City	058	Common-law-derived, inside India	NON-ALIGNED, genuinely and durably	The weakest item; determines the job	Mixed and technical	Real and specific
Türkiye	059	Civil, Swiss and Germanic line	The most genuinely ambiguous in the series	Adequate on substance, uncertain on predictability	Conditional on reciprocity	Real and inexpensive
Mauritius	059	Mixed: French civil + English commercial	NON-ALIGNED; Indian Ocean and Africa	High for the region and unusual	Statutory + common law + New York Convention	Real and audited, in two layers
Rwanda (KIFC)	059	Converging mixed	NON-ALIGNED, multi-vector	Unmeasured, which is not low	Verify — including New York Convention status	Certification conditions; verify
Kenya (NIFC)	059	Common law by reception	NON-ALIGNED; Western in security, Chinese in infrastructure	Genuinely good; demonstrated on the hardest test	Reciprocating list + common law	Ordinary residence + NIFC certification
Labuan IBFC	059	Tri-partite: common, civil, Sharia	NON-ALIGNED, ASEAN-anchored	Good and demonstrated	Reciprocating list + common law	Real, prescribed, annual, with a cliff edge
Bahrain	059	Civil with substantial Sharia influence	Gulf bridge; less independent than the UAE	Moderate	New York Convention + GCC and Arab League conventions	Licensing-based under the CBB Rulebook
Panama	060	Civil; Law 32 of 1927	Western hemisphere; United States-coupled	Mixed and genuinely two-sided	Exequatur, reciprocity and public policy	Territoriality by where activity is carried on
Uruguay	060	Civil, nineteenth-century code	Least United States-coupled among the Latin American and African assessments	Good by regional standards	Supreme Court + Inter-American instruments	Real for the free-zone regime
Paraguay	060	Civil, 1985 code	Non-aligned in form, pre-committed in substance	The weakest in the section	Statutory exequatur + Mercosur	Maquila requires a genuine export operation
Costa Rica	060	Civil; 1949 constitutional order	WEST by deliberate choice; OECD 25 May 2021	Good; Constitutional Chamber since 1989	Exequatur + Inter-American instruments	Real and recently tightened

Jurisdiction	6 CRS / FATCA	7 Beneficial ownership	8 Rails	9 Stability
DIFC (as holding)	UAE participating; DIFC inside it	Registrar UBO regime; not publicly searchable	SWIFT; dual-rail optionality at the UAE bank layer	High regionally; policy can move without a legislative cycle
United Kingdom	Full; CRS first exchanges 2017	PSC register, public — the most transparent in the series	SWIFT, deep; sterling clearing through the Bank of England	High
Jersey	Full; CRS early adopter, 2017	Central register, not public; competent-authority access	SWIFT only; no CIPS-side access	High
Guernsey	Full; CRS early adopter, 2017	Central register, not public	SWIFT only; no CIPS-side access	High
Isle of Man	Full; CRS early adopter, 2017	Central register, not public	SWIFT only; inside the UK indirect-tax territory	High
Luxembourg	Full; CRS via the administrative cooperation directive	Public access struck down 22 November 2022; authorities retain access	SWIFT, TARGET, SEPA; deep euro clearing	Very high domestically; externally rule-taking

Jurisdiction	6 CRS / FATCA	7 Beneficial ownership	8 Rails	9 Stability
Ireland	Full	Register from 2019; public access suspended	SWIFT, TARGET, SEPA	High, with documented corporation-tax concentration
Netherlands	Full; DAC8 from 1 January 2026	UBO register 2020; public access suspended 22 November 2022	SWIFT, TARGET, SEPA; large payments sector	High institutionally; moderate on rule stability
Switzerland	Full; MCAA 19 Nov 2014, AEOI Act 18 Dec 2015	No EU-style public register; identification through institutions	SWIFT; CHF clearing through SNB and SIC	Among the highest on process
Bermuda	CRS participating; MCAA signatory	Collected for decades; not publicly searchable	SWIFT; USD-centric, dollar-pegged	High
Cayman Islands	CRS from 2017; FATCA IGA	Beneficial Ownership Transparency Act 2023, commenced 2024	SWIFT; dollar-pegged; no CIPS relevance	High
British Virgin Islands	CRS participating; FATCA IGA	Registry filing regime; verify commencement and access	SWIFT; USD is legal currency	High
The Bahamas	CRS participating; FATCA IGA	Register accessible to competent authorities; not public	SWIFT; dollar-pegged	Moderate to high
Barbados	CRS participating; FATCA IGA	Register maintained; competent-authority access	SWIFT; dollar-pegged	Moderate
ADGM	UAE participating; ADGM inside it	Registration Authority regime; competent-authority access	SWIFT; dual-rail optionality at the UAE bank layer	High
DIFC (as bridging)	UAE participating; FATCA through the UAE arrangement	DIFC Registrar UBO regime	SWIFT; genuine multi-currency and multi-counterparty depth	High
Qatar / QFC	CRS participating; FATCA arrangement	QFC-level regime; competent-authority access	SWIFT; riyal pegged; CBDC project	High
Saudi Arabia	CRS participating; full FATF member since 2019	Commercial registry with beneficial-ownership requirements	SWIFT; riyal pegged	Moderate to high
Singapore	Fully participating; CRS 2018; FATCA Model 1	ACRA register of registrable controllers; not public	SWIFT; multi-currency depth; RMB clearing bank since 2013	Very high
Hong Kong	Participating under the Inland Revenue Ordinance	Significant controllers register since March 2018	Deepest offshore renminbi clearing capability available	Contested; see item 10
AIFC (Astana)	Kazakhstan participates; verify commencement	AIFC filings plus the national register	SWIFT through a small, concentrated banking sector	Improving; one stress event survived
Thailand	CRS participant; FATCA arrangement	DBD filings plus AML obligations	SWIFT; exchange control through authorised agents	Moderate
GIFT City	India participating; FATCA arrangement	Significant beneficial owner regime under Indian company law	SWIFT plus lawful rupee-corridor access	High continuity; moderate policy volatility
Türkiye	CRS participant; staged exchange relationships	Identified under the AML framework	SWIFT; large, internationally connected banking sector	Institutionally continuous, politically volatile
Mauritius	Participating; FATCA arrangement	Registrar plus licensed management companies	SWIFT; correspondent depth is the constraint	Moderate to high
Rwanda (KIFC)	Verify, and treat as gating	Verify the register, its keeper and its access	SWIFT; limited correspondent depth	Improving with an unmitigable tail risk
Kenya (NIFC)	CRS regulations enacted; FATCA arrangement	Companies (Beneficial Ownership Information) Regulations, 2020	SWIFT plus the deepest mobile-money infrastructure in Africa	Volatile

Jurisdiction	6 CRS / FATCA	7 Beneficial ownership	8 Rails	9 Stability
Labuan IBFC	Malaysia participating; Labuan inside it	Collected through licensed trust companies and the LFSA	SWIFT plus ringgit and offshore-dollar depth	High
Bahrain	CRS participant; FATCA arrangement	Collected through commercial registration	SWIFT; dollar-pegged; regional wholesale role	Moderate
Panama	CRS participant; FATCA arrangement	Private centralised registration under Law 129 of 2020	SWIFT; dollarised, dollar clearing via US correspondents, no lender of last resort	Moderate
Uruguay	CRS from the first wave; FATCA arrangement	Central bank registry from 2012, extended since	SWIFT; own currency and central bank	High by regional standards
Paraguay	Verify commencement and current exchange status	Framework exists under AML legislation	SWIFT; shallow sector, limited correspondent depth	Uncertain
Costa Rica	CRS participant; FATCA arrangement; OECD peer review	Register through the central bank system	SWIFT; own currency, dollarised private deposits in practice	Moderate

Jurisdiction	10 DEGRADATION TRAJECTORY	11 Job	12 Verdict
DIFC (as holding)	IMPROVING, from a base that briefly worsened — FATF Mar 2022 to Feb 2024, then corporate tax and a UBO regime	HOLDING, unambiguously	Use it as the apex
United Kingdom	ERODING on differentiation and optionality; improving on institutional integrity	OPERATING; narrow bridging; never holding	Use it, size it, never let it grow
Jersey	STABLE on capability, ERODING on differentiation	HOLDING, subordinate — never the apex	Use for trust technology; never to raise N
Guernsey	STABLE on capability, ERODING on differentiation	HOLDING with a genuine cell specialisation	Use for cells and captives
Isle of Man	STABLE on capability, ERODING on differentiation	OPERATING — the one Crown Dependency whose job is not holding	Registries and asset ownership; not the apex
Luxembourg	ERODING on differentiation; IMPROVING on privacy; STABLE on capability	OPERATING (regulated funds, securitisation)	Funds and securitisation only; never a conduit
Ireland	ERODING sharply on tax differentiation; STABLE to IMPROVING on legal and operating	OPERATING, unusually strong	Leasing, funds, securitisation; not the apex

Jurisdiction	10 DEGRADATION TRAJECTORY	11 Job	12 Verdict
Netherlands	ERODING on differentiation — largely self-inflicted and deliberate	HOLDING in one narrow sense (the STAK) plus OPERATING	Use the STAK; never a conduit
Switzerland	ERODING — the longest-running erosion in the series; bloc position confirmed rather than degraded	BRIDGING in one narrow sense (custody, CHF) plus OPERATING	Custody and CHF; never the apex; not a separate bloc
Bermuda	STABLE to IMPROVING on institutional quality; the zero-tax proposition is over	OPERATING with a HOLDING consequence	One job: funding a real retained exposure
Cayman Islands	STABLE, after a genuine deterioration that has been repaired	HOLDING, institutional sense only	Where unrelated third parties expect it; not family holding
British Virgin Islands	ERODING on the product; IMPROVING on the institution	HOLDING, narrowly — one clean SPV	Single-purpose vehicles only; never an apex
The Bahamas	STABLE, from a worse base, with an unrepaired reputational component	HOLDING, private-client sense	Only where a specific Bahamian instrument is unique
Barbados	ERODING as a haven; CONVERGING as a normal treaty jurisdiction	OPERATING (licensed insurance) plus narrow holding	Only with a real Canada or CARICOM corridor
ADGM	IMPROVING, faster than any comparable centre	HOLDING — intermediate	The intermediate holding layer; elect ADGM law group-wide
DIFC (as bridging)	IMPROVING, from a base that measurably worsened and recovered — the round trip is the evidence	BRIDGING plus genuinely-needed regulated activity	Apex, banking, and a licence only if one is truly needed
Qatar / QFC	IMPROVING, steadily and without drama	OPERATING	Only with real Qatari activity; never a second Gulf holding layer
Saudi Arabia	IMPROVING materially, from a low base, on a short record	OPERATING and market access; nothing else, ever	Enter for revenue only; hold nothing there

Jurisdiction	10 DEGRADATION TRAJECTORY	11 Job	12 Verdict
Singapore	STABLE institutional / ERODING bloc independence / IMPROVING regulatory credibility — do not average	BRIDGING primarily; narrow secondary holding	Use for bridging, sized on capability not bloc arithmetic
Hong Kong	ERODING on the axis this series buys, with dated evidence and a visible direction	BRIDGING only, and transit rather than store	Narrow, capped, swept, with a live migration path
AIFC (Astana)	IMPROVING from a low base, with genuine uncertainty and one stress test survived	HOLDING, narrow and ring-fenced; BRIDGING in one respect	Small, once, and simply — and take the arbitration clause regardless
Thailand	STABLE on ownership framework; TIGHTENING on tax; IMPROVING on transparency	OPERATING, exclusively	The operating leg, on a BOI basis, with nothing held or swept through it
GIFT City	IMPROVING, with the upside capped by the same structural reason	BRIDGING — specifically the rupee corridor	A bounded corridor with a stated maximum balance and dwell
Türkiye	IMPROVING on standards and institutions; UNRESOLVED on the currency, and the second decides it	OPERATING; holding prohibited outright	Operate; never hold
Mauritius	ERODING as a fiscal proposition; IMPROVING as an institutional one	HOLDING and BRIDGING for Africa; never operating	Where the treaty and the protection agreement do real work
Rwanda (KIFC)	IMPROVING, from a very short base, with a specific unmitigable tail risk	WATCH, moving toward OPERATING	Establish the relationship; run the scorecard annually; do not concentrate
Kenya (NIFC)	MIXED — tax and AML lines moving in opposite directions	OPERATING, unambiguously — the best operating base among the Latin American and African assessments	Operate; cap the share; do not hold
Labuan IBFC	STABLE to IMPROVING; the offshore concession already surrendered	BRIDGING primarily, HOLDING secondarily, never the apex	Islamic issuance and structuring; verify the treaty position entity by entity

Jurisdiction	10 DEGRADATION TRAJECTORY	11 Job	12 Verdict
Bahrain	STABLE on standards; ERODING on fiscal exceptionalism; structurally challenged	BRIDGING and standards-facing, not holding	Engage as the standards jurisdiction; do not concentrate capital
Panama	IMPROVING on standards from a poor base; structurally exposed to recurrence	OPERATING plus limited regional bridging	Regional trade and logistics; never protected capital
Uruguay	STABLE — the transparency degradation already completed, institutions untouched	OPERATING and personal RESIDENCE	The best residence and operating jurisdiction among the Latin American and African assessments
Paraguay	UNCERTAIN, and the uncertainty is itself the finding	OPERATING only — export manufacturing under the maquila regime	Manufacture only; resolve the open items before committing anything
Costa Rica	IMPROVING institutionally; CONVERGING fiscally, forced from outside in under a year	OPERATING at genuine quality, and RESIDENCE	A high-quality operating and residence jurisdiction; expect no bloc benefit

A.3 Sorted by job — item 11

The single most useful re-sort in the atlas, because the three-jobs diagnostic's rule is that an entity gets one job and the atlas makes the shortage visible.

```
HOLDING -- an apex candidate
  DIFC (Ch 7)                    the framework's answer; the only unqualified apex

HOLDING -- intermediate or subordinate, never the apex
  ADGM (Ch 17)                   intermediate holdco beneath the Foundation
  Jersey, Guernsey (Ch 11)       excellent instruments inside one perimeter
  Cayman, BVI, Bahamas (Ch 16)   institutional or single-purpose only
  Netherlands -- STAK (Ch 12)    control separation, not ownership
  Mauritius (Ch 26)              Africa-facing regional holding
  Labuan (Ch 27)                 secondary to its bridging job
  AIFC (Ch 22)                   narrow, ring-fenced fourth-bloc position

OPERATING
  Thailand (Ch 23)               the operating leg; highest operating score in the series
  United Kingdom (Ch 10)         real trade, real staff, full accounts
  Isle of Man (Ch 11)            registries and asset ownership
  Luxembourg, Ireland (Ch 12)    regulated funds, leasing, securitisation
  Netherlands, Switzerland (12)  genuine operations alongside their narrow specialisms
  Bermuda (Ch 14)                licensed underwriting with a holding consequence
  Barbados (Ch 16)               licensed insurance with a corridor
```

```
  Qatar (Ch 19)                only with real Qatari activity
  Saudi Arabia (Ch 19)         revenue and market access, nothing else
  Türkiye (Ch 25)              operating only; holding prohibited
  Kenya (Ch 26)                the best operating base in its group
  Panama, Uruguay, Paraguay,
    Costa Rica (Ch 28)         trade, residence, and manufacture

BRIDGING
  DIFC (Ch 18)                 apex banking and regulated activity
  Singapore (Ch 20)            treasury, settlement, FX, Asian contracting
  Hong Kong (Ch 21)            transit only, capped and swept
  GIFT City (Ch 24)            the rupee corridor, bounded
  Labuan (Ch 27)               Islamic issuance and structuring
  Bahrain (Ch 27)              standards and wholesale funding
  Switzerland (Ch 12)          custody and CHF exposure
  Mauritius (Ch 26)            Africa-facing neutral forum

WATCH -- not yet assignable
  Rwanda (Ch 26)               establish the relationship; re-run annually
```

Two observations the sort makes unavoidable. **The holding column is short and the operating column is long, which is the correct shape**: apex capability is scarce, and a structure that treats the two as interchangeable ends up holding value in an operating jurisdiction. And **eight of the nine bridging entries carry an explicit cap, a sweep, a bound, or a "narrow"** — bridging is the job most likely to be over-sized.

A.4 Sorted by degradation trajectory — item 10

```
IMPROVING
  DIFC (7, 18)     UAE round trip through the FATF listing is the evidence
  ADGM (17)        fastest-improving centre in the series
  Qatar (19)       steady, undramatic
  Saudi Arabia (19) materially, from a low base, on a short record
  AIFC (22)        from a low base, with genuine uncertainty
  GIFT City (24)   with the upside capped by the same structural reason
  Rwanda (26)      from a very short base, with an unmitigable tail risk
  Costa Rica (28)  institutionally, while converging fiscally
  Panama (28)      from a poor base, structurally exposed to recurrence

STABLE
  Cayman (16)      after a deterioration that has been repaired
  Bahamas (16)     from a worse base, reputational component unrepaired
  Bermuda (14)     institutional quality; the zero-tax proposition is over
  Uruguay (28)     degradation already completed; institutions untouched
  Labuan (27)      offshore concession already surrendered

ERODING
  United Kingdom (10)   differentiation and optionality, not quality
  Jersey, Guernsey,
    Isle of Man (11)    differentiation, which was the only thing bought
  Luxembourg (12)       differentiation, at premium cost
  Ireland (12)          tax differentiation, sharply
  Netherlands (12)      differentiation, self-inflicted and deliberate
  Switzerland (12)      the longest-running erosion documented in the series
  BVI (16)              the product erodes while the institution improves
```

```
  Barbados (16)         as a haven; converging as a normal treaty jurisdiction
  Hong Kong (21)        on the exact axis this series buys
  Mauritius (26)        as a fiscal proposition
  Bahrain (27)          fiscal exceptionalism

SPLIT -- refuses to average
  Singapore (20)   stable institutional / eroding bloc / improving regulatory
  Thailand (23)    stable ownership / tightening tax / improving transparency
  Kenya (26)       tax and AML lines moving in opposite directions
  Türkiye (25)     improving standards / unresolved currency

UNCERTAIN
  Paraguay (28)    and the uncertainty is itself the finding
```

The pattern is the framework's thesis in one column. **Every Western-perimeter entry is eroding on differentiation and none is eroding on quality.** The jurisdictions are not getting worse; they are getting more alike, which is a different problem and the one this series is about. The improving column is dominated by Gulf and non-aligned centres building institutions rather than dismantling them, and the erosion in Hong Kong is the one entry where the axis eroding is the axis being bought.

A.5 What the atlas does not contain

Four omissions, stated so that the table is not read as a completeness claim.

Cost. No `C` figure appears in any scorecard, in any section, or in this atlas. The Gulf-centre assessments, 18, 19, and 22 each state explicitly that the `C` term cannot be estimated without current fee schedules and that no figure in the section supplies one. **A reader who wants to compare jurisdictions on cost must obtain current fee schedules; this series supplies the method and refuses to supply the numbers.**

Bankability. Whether a specific bank will onboard a specific structure is the binding constraint on the whole architecture — the concentration paper's failure mode 11 — and it is a property of the bank, the group, the currency mix, and the year rather than of the jurisdiction. No column can hold it.

Jurisdictions not assessed. Where a section the manifest specifies is not on the page, its scorecard is not in this atlas, and Appendix D marks the gap.

The reader's own facts. Every item 12 verdict is written for the carried case: a dual national, UAE tax resident, three heirs in three jurisdictions, one in a forced-heirship civil-law state, with operations in Thailand, trade in the UAE, real estate in London, and a listed portfolio. Change any of those and several verdicts change with them.

Appendix B. The concentration workbook

Every figure in this appendix was recomputed in Python from the stated inputs at full precision and rounded only for printing. No figure was taken from a rounded print anywhere in the series. That discipline is not fastidiousness: computing `h` from the printed `N = 2.94` rather than from the exact `1/0.34` produces `\ensuremath{-}0.122557` instead of `\ensuremath{-}0.122417`, an error of 0.000140 — small, and large enough to move a break-even probability by more than a percentage point.

B.1 The instrument

```
HHI = sum over i of s_i^2          s_i = share of the total

effective count = 1 / HHI

Properties used throughout:
  - n equal shares give HHI = 1/n and effective count = n exactly
  - one share of 100% gives HHI = 1 and effective count = 1
  - the effective count is always at most the raw count, with
    equality only when the shares are equal
  - the measure is invariant to relabelling and to splitting a
    holder into two holders in the SAME bloc, which is precisely
    why bloc collapse must happen BEFORE the computation
```

The numbers-equivalent reading of `1/HHI` is conventionally attributed to M. A. Adelman's 1969 comment in the *Review of Economics and Statistics*. The attribution should be checked against the original; the mathematical property is elementary and does not depend on it.

The same instrument is used in the opposite direction by competition authorities. The 2023 DOJ/FTC Merger Guidelines set a structural presumption of unlawfulness at HHI above 1,800 with a change over 100 — on a scale of 10,000 rather than 1, which is the same number multiplied by 10,000. **You want HHI low on jurisdictional exposure, for resilience, and low enough on market share, to stay unchallenged. One metric, two directions, both binding.** Those thresholds are agency enforcement policy rather than statute, courts are not bound by them, and they have changed three times since 1982.

B.2 The conventional structure — 1.00 effective bloc

```
AS FOUND

  jurisdiction   share    bloc
  ------------------------------
  UK              40.9%   WEST
  US              27.3%   WEST
  Jersey          18.2%   WEST
  Luxembourg      13.6%   WEST
                  ------
                  100.0%

NAIVE, per jurisdiction
  HHI = 0.409^2 + 0.273^2 + 0.182^2 + 0.136^2
      = 0.167281 + 0.074529 + 0.033124 + 0.018496
      = 0.293430
  effective jurisdictions = 1 / 0.293430 = 3.407968

COLLAPSED TO BLOCS
  WEST 100.0%
  bloc HHI = 1.000^2 = 1.000000
  EFFECTIVE BLOCS = 1 / 1.000000 = 1.000000

  Western-perimeter exposure = 100.0%
```

Four entities, four jurisdictions, one point of failure. The gap between 3.407968 and 1.000000 is the entire subject of the concentration paper. Naive HHI overstates resilience because it treats jurisdictions as independent when they share a coordination mechanism,

and four Western jurisdictions do not fail independently — they fail on one decision, taken once, in one of the five interfaces the coordination-perimeter analysis traces.

B.3 The designed structure — 2.94 effective blocs

```
AS DESIGNED

  jurisdiction   share    bloc
  ------------------------------------
  UAE             29.1%   BRIDGE
  Thailand        18.2%   NONALIGNED
  UK              16.4%   WEST
  Bermuda         13.6%   WEST (W-LITE)
  India           11.8%   NONALIGNED
  Singapore       10.9%   BRIDGE
                  ------
                  100.0%

NAIVE, per jurisdiction
  HHI = 0.291^2 + 0.182^2 + 0.164^2 + 0.136^2 + 0.118^2 + 0.109^2
      = 0.084681 + 0.033124 + 0.026896 + 0.018496
        + 0.013924 + 0.011881
      = 0.189002
  effective jurisdictions = 1 / 0.189002 = 5.290949

COLLAPSED TO BLOCS
  BRIDGE      40.0%   (UAE 29.1 + Singapore 10.9)
  NONALIGNED  30.0%   (Thailand 18.2 + India 11.8)
  WEST        30.0%   (UK 16.4 + Bermuda 13.6)
              ------
              100.0%

  bloc HHI = 0.400^2 + 0.300^2 + 0.300^2
           = 0.160000 + 0.090000 + 0.090000
           = 0.340000
  EFFECTIVE BLOCS = 1 / 0.340000 = 2.941176

  Western-perimeter exposure = 30.0%

THE COMPARISON

                              AS FOUND    AS DESIGNED    CHANGE
  jurisdictions held               4            6          +2
  naive effective count      3.407968     5.290949      x1.5525
  EFFECTIVE BLOCS            1.000000     2.941176      x2.9412
  Western-perimeter exposure   100.0%        30.0%      -70.0pp
```

B.4 The storage law applied

```
h = eta ln(N) - delta Q - C

test parameters:   eta = 0.35   C = 0.30   Q = 10   delta = 0.02
so  delta Q = 0.200000  and  delta Q + C = 0.500000
```

```
CONVENTIONAL
  ln(1.000000) = 0.000000
  h = 0.35 x 0.000000 - 0.500000 = -0.500000

DESIGNED
  ln(2.941176) = 1.078810
  h = 0.35 x 1.078810 - 0.500000 = -0.122417

IMPROVEMENT in h = 0.377583
```

`ln(1.00) = 0`. The conventional structure has zero available surplus — not low, zero. `Q` decays at `delta` with nothing on the input side, and no rearrangement of those four entities changes it, because rearranging within a bloc does not change `N`. That single line is the reason the concentration paper owns the instrument and the reason `N` is defined as the effective bloc count rather than the entity count or the jurisdiction count.

And the honest companion result, which the assembled architecture and this synthesis both carry: both structures are below the build threshold at the tested parameters. `h = \ensuremath{-}0.500000` and `h = \ensuremath{-}0.122417` are both negative. Diversification raised the ceiling by 0.377583 and did not pay `delta Q + C = 0.500000`.

B.5 The break-even and the marginal value

```
BREAK-EVEN EFFECTIVE BLOC COUNT

  h = 0  requires  eta ln(N*) = delta Q + C
                   ln(N*) = 0.500000 / 0.35 = 1.428571
                   N* = exp(1.428571) = 4.172734

  designed structure reaches 2.941176
  shortfall in N = 1.231558
  shortfall in h = 0.122417

MARGINAL VALUE OF ONE MORE EFFECTIVE BLOC

  dh/dN = eta / N

  at N = 1.000000   dh/dN = 0.350000
  at N = 2.941176   dh/dN = 0.119000
  at N = 4.000000   dh/dN = 0.087500
  at N = 4.172734   dh/dN = 0.083879

THE DIMINISHING RETURN, in the values that matter

  1.000000 -> 2.941176   buys  0.377583
  2.941176 -> 4.000000   buys  0.107620    28.5% as much
```

The consequence for design: **the first move out of a single perimeter is worth roughly three and a half times the last move to the break-even**, and a structure at `N = 2.941176` should expect every subsequent leg to be harder to justify than the one before it. That is not a reason to stop. It is a reason to price each one.

B.6 V5 — the Qatar/Saudi result: more jurisdictions, fewer blocs

The framework's headline demonstration of the failure mode the concentration paper exists to catch.

```
AS TEMPTED -- add Qatar and Saudi, funded by trimming UAE, UK and Singapore

  jurisdiction   share    bloc
  ------------------------------------
  UAE             20.0%   BRIDGE (GCC)
  Thailand        18.2%   NONALIGNED
  Bermuda         13.6%   WEST
  UK              12.2%   WEST
  India           11.8%   NONALIGNED
  Saudi Arabia     9.1%   BRIDGE (GCC)
  Qatar            8.2%   BRIDGE (GCC)
  Singapore        6.9%   BRIDGE
                  ------
                  100.0%

  naive HHI = 0.200^2 + 0.182^2 + 0.136^2 + 0.122^2
            + 0.118^2 + 0.091^2 + 0.082^2 + 0.069^2
            = 0.140194
  effective jurisdictions = 1 / 0.140194 = 7.132973

COLLAPSED TO BLOCS
  BRIDGE      44.2%   (UAE + Saudi + Qatar + Singapore)
  NONALIGNED  30.0%   (Thailand + India)
  WEST        25.8%   (Bermuda + UK)

  bloc HHI = 0.442^2 + 0.300^2 + 0.258^2
           = 0.195364 + 0.090000 + 0.066564
           = 0.351928
  EFFECTIVE BLOCS = 1 / 0.351928 = 2.841490
  h = 0.35 x ln(2.841490) - 0.500000 = -0.134485

SIDE BY SIDE

                          DESIGNED     TEMPTED      CHANGE
  jurisdictions held          6           8           +2
  naive effective count   5.290949    7.132973     +34.8146%
  EFFECTIVE BLOCS         2.941176    2.841490      -3.3893%
  h                      -0.122417   -0.134485      -0.012068
```

Two more jurisdictions, a third more apparent diversification, and the honest count fell. Both effects run the same way: the additions are all in the bloc the structure was already heaviest in, and the funding came disproportionately from the two blocs supplying the diversification. The naive metric reports neither.

```
AS BUILT -- the priced alternative from Ch 19's Worked Application
Saudi RHQ at 5.0%, funded from the OPERATING allocation (Thailand 18.2 -> 13.2)

  UAE 29.1 BRIDGE · UK 16.4 WEST · Bermuda 13.6 WEST · Thailand 13.2 NONALIGNED
  India 11.8 NONALIGNED · Singapore 10.9 BRIDGE · Saudi 5.0 BRIDGE (GCC)

  naive HHI = 0.175802     effective jurisdictions = 5.688217

  blocs  BRIDGE 45.0%  WEST 30.0%  NONALIGNED 25.0%
  bloc HHI = 0.202500 + 0.090000 + 0.062500 = 0.355000
  EFFECTIVE BLOCS = 2.816901
```

```
  h = −0.137527

  PRICED LOSS versus the designed structure
    in N   −0.124275
    in h   −0.015110
```

A priced loss taken knowingly is legitimate; an unpriced one is the failure. The Saudi position costs 0.124275 of effective bloc count and is defensible if addressable Saudi state and state-linked revenue exceeds that cost. It is indefensible as diversification, because it is the opposite.

B.7 V3 — the Singapore control experiment

The cleanest demonstration in the paper that the gap between the naive and the honest count is real and not an artefact of any particular bloc assignment. Place the same 10.9% in London instead of Singapore.

```
SINGAPORE LEG (as designed)
  UAE 29.1 · Thailand 18.2 · UK 16.4 · Bermuda 13.6 · India 11.8 · Singapore 10.9
  naive HHI = 0.189002     effective jurisdictions = 5.290949
  blocs BRIDGE 40.0 / NONALIGNED 30.0 / WEST 30.0
  bloc HHI = 0.340000      N = 2.941176      h = −0.122417

LONDON INSTEAD -- the 10.9% added to the UK slot
  UAE 29.1 · Thailand 18.2 · UK 27.3 · Bermuda 13.6 · India 11.8
  naive HHI = 0.291^2 + 0.182^2 + 0.273^2 + 0.136^2 + 0.118^2
            = 0.224754
  effective jurisdictions = 1 / 0.224754 = 4.449309
  blocs BRIDGE 29.1 / NONALIGNED 30.0 / WEST 40.9
  bloc HHI = 0.084681 + 0.090000 + 0.167281 = 0.341962
  N = 2.924302              h = −0.124431

THE RESULT
  naive effective count   5.290949 vs 4.449309   difference 0.841640
  EFFECTIVE BLOCS         2.941176 vs 2.924302   difference 0.016874
  h                      −0.122417 vs −0.124431  difference 0.002014
```

Singapore's true marginal contribution to h, measured against simply putting the same money in London, is 0.002014. The naive metric reported 0.841640 of an effective jurisdiction — nearly a full unit of diversification — that did not occur.

The same arithmetic doubles as the robustness test for the bloc call itself. Reclassifying Singapore from BRIDGE to WEST produces exactly the London bloc vector — BRIDGE 29.1, NONALIGNED 30.0, WEST 40.9 — and therefore exactly the same `N = 2.924302` and `h = \ensuremath{-}0.124431`. **The argument most likely to be had about the whole allocation moves the answer by 0.002014, which is to say it does not move the answer.**

B.8 V2 — the Hong Kong result, and why the largest gain was declined

```
A: HONG KONG AS THE ASIAN CENTRE, at 10.9%, assigned to a CHINA bloc
   bloc HHI = 0.276562      N = 3.615826      h = −0.050138
   gain over the designed baseline           +0.072279

B: CAPPED TRANSIT ONLY -- a written balance cap and a weekly sweep
   bloc HHI = 0.337224      N = 2.965388      h = −0.119547
```

```
   gain over the designed baseline           +0.002869

THE RATIO
   B's gain as a proportion of A's = 0.002869 / 0.072279 = 3.9699%
   B's exposure as a proportion of A's                    3.21%
```

The largest single-move improvement in the series is 0.072279, and the Hong Kong assessment recommends against taking it. The section takes capped transit instead: 3.97% of the benefit for 3.21% of the exposure. That is a deliberate refusal of a measured gain on political-risk grounds, stated with the number rather than with an adjective, and it is the model for how a risk that the instrument cannot price should be handled — **by declining the position and printing what the refusal cost.**

B.9 V6 — the correspondent-system measurement

The concentration instrument applied to rails rather than to holdings, and measured by **terminal clearing system** rather than by bank. Two banks clearing dollars through the same United States correspondent are one system, not two.

```
AS FOUND
  USD 100.0%
  HHI = 1.000000
  EFFECTIVE CLEARING SYSTEMS = 1.000000

AS DESIGNED
  USD  55.0%
  CNY  15.0%
  EUR  12.0%
  INR  10.0%
  AED   8.0%
      ------
      100.0%

  HHI = 0.550^2 + 0.150^2 + 0.120^2 + 0.100^2 + 0.080^2
      = 0.302500 + 0.022500 + 0.014400 + 0.010000 + 0.006400
      = 0.355800
  EFFECTIVE CLEARING SYSTEMS = 1 / 0.355800 = 2.810568

  IMPROVEMENT  1.000000 -> 2.810568
```

The honest limit, and the reason this measurement belongs in a failure-aware appendix. Reaching 4.000000 effective clearing systems, holding the non-dollar mix in its current 15 : 12 : 10 : 8 proportions, requires the dollar share to fall to approximately 38.99%:

```
  USD share      effective systems
  ------------------------------
   55.00%           2.810568
   45.00%           3.544579
   40.00%           3.925331
   38.99%           4.000088   <- the threshold
   35.00%           4.278877
   33.33%           4.384374
   30.00%           4.566776
```

The 55% dollar share is set by revenue currency, not by structure. Moving it to 38.99% is a commercial decision about what the group invoices in, not a legal decision about where it banks. Four hard limits sit underneath the whole measurement: a USD account settles under United States law wherever it is booked; CIPS is a clearing system and many participants still use SWIFT for messaging, so “CIPS replaces SWIFT” is false and any structure built on that premise will fail; the bank’s risk appetite is not the jurisdiction’s policy; and a rail is not a holding.

B.10 V4 — the break-even probability method

The most transferable single technique in the series. It converts an unwinnable argument about geopolitics into a number a family can act on.

```
THE METHOD

  1  compute h WITHOUT the candidate leg                  -> h_0
  2  compute h WITH the leg, perimeter call RIGHT         -> h_R
     gain G = h_R - h_0
  3  compute h WITH the leg, perimeter call WRONG
     (the leg folds into a bloc already held)             -> h_W
     loss L = h_0 - h_W
  4  break-even probability   p* = L / (G + L)
  5  cost ceiling             the leg must cost less than G,
                              expressed as a share of total C
```

Worked in full on the AIFC at 5% drawn from the United Kingdom slot:

```
STEP 1 -- baseline
  blocs BRIDGE 40.0 / NONALIGNED 30.0 / WEST 30.0
  bloc HHI = 0.340000    N = 2.941176    h_0 = -0.122417

THE ALLOCATION WITH THE LEG -- UK 16.4 -> 11.4, AIFC 5.0
  UAE 29.1 · Thailand 18.2 · Bermuda 13.6 · India 11.8
  UK 11.4 · Singapore 10.9 · AIFC 5.0
  naive HHI = 0.177602    effective jurisdictions = 5.630567

STEP 2 -- the call is RIGHT: AIFC is a fourth, EURASIAN perimeter
  blocs BRIDGE 40.0 / NONALIGNED 30.0 / WEST 25.0 / EURASIAN 5.0
  bloc HHI = 0.160000 + 0.090000 + 0.062500 + 0.002500 = 0.315000
  N = 3.174603           h_R = -0.095686
  G = 0.026731

STEP 3 -- the call is WRONG: AIFC folds into NONALIGNED
  blocs BRIDGE 40.0 / NONALIGNED 35.0 / WEST 25.0
  bloc HHI = 0.160000 + 0.122500 + 0.062500 = 0.345000
  N = 2.898551           h_W = -0.127526
  L = 0.005110

STEP 4 -- the break-even
  p* = L / (G + L)
     = 0.005110 / (0.026731 + 0.005110)
     = 0.005110 / 0.031841
     = 0.160476   =  16.0476%

STEP 5 -- the cost ceiling
```

```
  c* = G = 0.026731
  as a share of the structure's total C = 0.30:
  0.026731 / 0.300000 = 0.089102  =  8.9102%

AND THE SHORTFALL CLOSED, if the call is right
  0.026731 / 0.122417 = 0.218357  =  21.8357%
```

Three results, each doing different work.

A 16.05 per cent probability that the jurisdiction is an independent perimeter is enough to justify the position [Curzi, 2026], whose Result One states the same bet as $E[\mathrm{d}h] = 0.0318401p - 0.0051096$ and reaches the identical break-even. The holder does not need to win the argument. It needs the argument to be better than roughly one in six.

Getting the bloc call wrong produces a loss of 0.005110, not merely an absence of gain, because capital was withdrawn from a bloc where it was diversifying and placed in a bloc it concentrated. That asymmetry is the whole reason the method exists.

The binding constraint is cost, not geopolitics. The entire gain is consumed if the leg costs more than 8.9102% of the structure's total maintenance budget, and that is a fee-schedule question.

The generalisation, and the instruction the assembled architecture, the failure record and this synthesis all carry: **price every marginal jurisdiction as a break-even probability, not as a geopolitical opinion.**

B.11 The job-vector workbook

The concentration instrument turned on the entity instead of the portfolio.

```
S = H^2 + O^2 + B^2            effective jobs = 1 / S

Reference points:
  pure specialist (1, 0, 0)          S = 1.0000   eff = 1.0000
  best two-job split (0.5, 0.5, 0)   S = 0.5000   eff = 2.0000
  even three-way (1/3 each)          S = 0.3333   eff = 3.0000

There is no two-job configuration above 2.0000. The 2.20 threshold
leaves room for estimation error in the vector and nothing more.

THE DESIGNED STRUCTURE, SCORED

entity                          H      O      B        S    eff jobs   verdict
-------------------------------------------------------------------------------
DIFC Foundation               0.90   0.00   0.10   0.8200     1.2195   SPECIALIST
Thai BOI opco                 0.05   0.95   0.00   0.9050     1.1050   SPECIALIST
UK Ltd                        0.10   0.85   0.05   0.7350     1.3605   SPECIALIST
Singapore treasury            0.05   0.15   0.80   0.6650     1.5038   STRAINED
ADGM holdco                   0.70   0.05   0.25   0.5550     1.8018   STRAINED
Bermuda captive               0.30   0.60   0.10   0.4600     2.1739   STRAINED
GIFT City vehicle             0.20   0.30   0.50   0.3800     2.6316   OVERLOADED
-------------------------------------------------------------------------------
mean effective jobs per entity                                 1.6852
```

Two entries deserve their flags. The Bermuda captive at 2.1739 is strained by nature — a captive underwrites, which is operating, and holds reserves, which is holding — and the strain is accepted in writing because the reserves are held against real third-party liabilities, which is what makes them a structurally awkward class to realise. The GIFT City vehicle at

2.6316 is carrying all three jobs and is a design error the framework owns rather than smooths over; the GIFT City assessment's own repair moves it to a vector of (0.05, 0.10, 0.85), `S = 0.7350`, 1.3605 effective jobs.

The selection rule is separate from the diagnostic and matters as much. Score each candidate on each requirement dimension and take the **minimum** over the dimensions the job actually needs. An average lets a strong score on one dimension conceal a fatal weakness on another; a structure fails at its weakest load-bearing requirement, not at its average one.

B.12 The leverage headroom workbook

```
THE CASCADE (the carried constructed case)

  operating company, marked value                    100.00
  carve out 20% by IPO                                20.00  proceeds
  retain and pledge                                   80.00
  pledge at 50% loan-to-value                         40.00  borrowing
  ------------------------------------------------------------
  deployable                                          60.00
  retained control                                    80.0%

THE TRIGGER

  critical collateral value = L / (1 - m)
                            = 40.00 / 0.65 = 61.538462
  drawdown to trigger       = 1 - 61.538462 / 80.00
                            = 0.230769  =  23.0769%

THE REVERSE COMPUTATION -- size the loan from the drawdown

  L_max = V (1 - d)(1 - m)      at V = 80.00, m = 35%

  survive  d       L_max        LTV       borrowing foregone
  ------------------------------------------------------------
   23.08%        40.000000     50.00%          --
   30.00%        36.400000     45.50%       3.600000
   35.00%        33.800000     42.25%       6.200000
   40.00%        31.200000     39.00%       8.800000
   50.00%        26.000000     32.50%      14.000000
   60.00%        20.800000     26.00%      19.200000
```

A 23.1% drawdown is ordinary equity volatility, not a crisis parameter, which is why the carve-out cascade leads with the failure mode rather than burying it. The reverse computation is the useful one because the drawdown is the input a family can have an opinion about and the loan-to-value is not. The 35% maintenance requirement is a constructed model parameter, not a market quotation.

B.13 Two further computed results carried into the appendices

V9 — the foundation-versus-trust crossover. Against a trustee charging on a basis-point scale and a constructed fixed-fee holding mandate of USD 120,000 per annum:

```
Q* = 120,000 / (basis in decimal)

35 bp   ->   USD  34.2857 million
```

```
15 bp   ->   USD  80.0000 million
10 bp   ->   USD 120.0000 million

check: 35 bp on USD 240 million = USD 840,000
```

The foundation is cheaper than a trust only above those thresholds, so cost is explicitly not the reason to choose it. The reason is recognition in civil-law and Sharia jurisdictions.

V8 — the succession residual. On the constructed estate, at a constructed exchange rate of USD 1.08 per euro and a constructed rate scale:

```
notional reserved share                          60.00m
compensatory levy = min(60.00m, 6.70m)            6.70m   = 11.1667%
French inheritance tax approx = 0.45 x 60.00m    27.00m
ratio, tax to reserve exposure                   4.0299x
liquidity shortfall                              33.84m   = 14.1000% of estate
```

The structure caps the reserve exposure at USD 6.70 million and reduces it to zero on the designed map. The tax exposure is 4.0299 times larger and is untouched. A succession analysis reporting only the first would be lying by omission.

B.14 Reproduction

Every result above is reproducible from the shares and parameters printed here with three functions and nothing else:

```
from math import log, exp

def hhi(shares):
    total = sum(shares)
    f = [x / total for x in shares]
    return sum(x * x for x in f)

def eff(shares):
    return 1.0 / hhi(shares)

def h(N, eta=0.35, C=0.30, Q=10.0, delta=0.02):
    return eta * log(N) - delta * Q - C

# break-even effective bloc count
N_star = exp((0.02 * 10.0 + 0.30) / 0.35)      # 4.172734

# break-even probability for a marginal jurisdiction
def p_star(G, L):
    return L / (G + L)
```

Collapse to blocs before calling `hhi`. Every failure this workbook documents comes from calling it on jurisdictions.

B.15 The precision rule

The concentration paper states it and it applies to every figure in this series. Recompute every derived figure from the exact inputs; never from a rounded print. The demonstration:

```
exact:    N = 1 / 0.34 = 2.941176...
          ln(2.941176) = 1.078810
          h = 0.35 x 1.078810 - 0.5 = -0.122417

from the printed 2.94:
          ln(2.940000) = 1.078410
          h = 0.35 x 1.078410 - 0.5 = -0.122557

error = 0.000140
```

That error is small in `h` and not small in what depends on `h`. The AIFC break-even probability is `L / (G + L)` where both `G` and `L` are differences of numbers of this size; propagating a 0.000140 error through it moves the break-even by more than a percentage point. **This destroyed four sections of an earlier series before it was caught, and it is the reason the rule is stated as a rule rather than as good practice.**

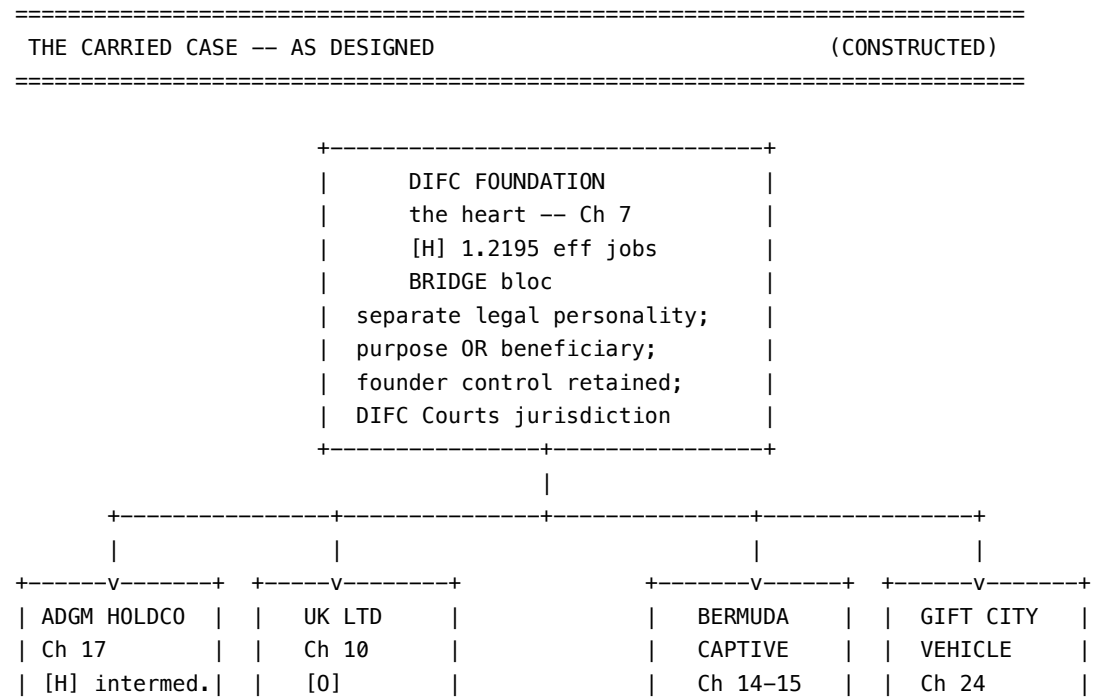
Appendix C. Structure diagrams

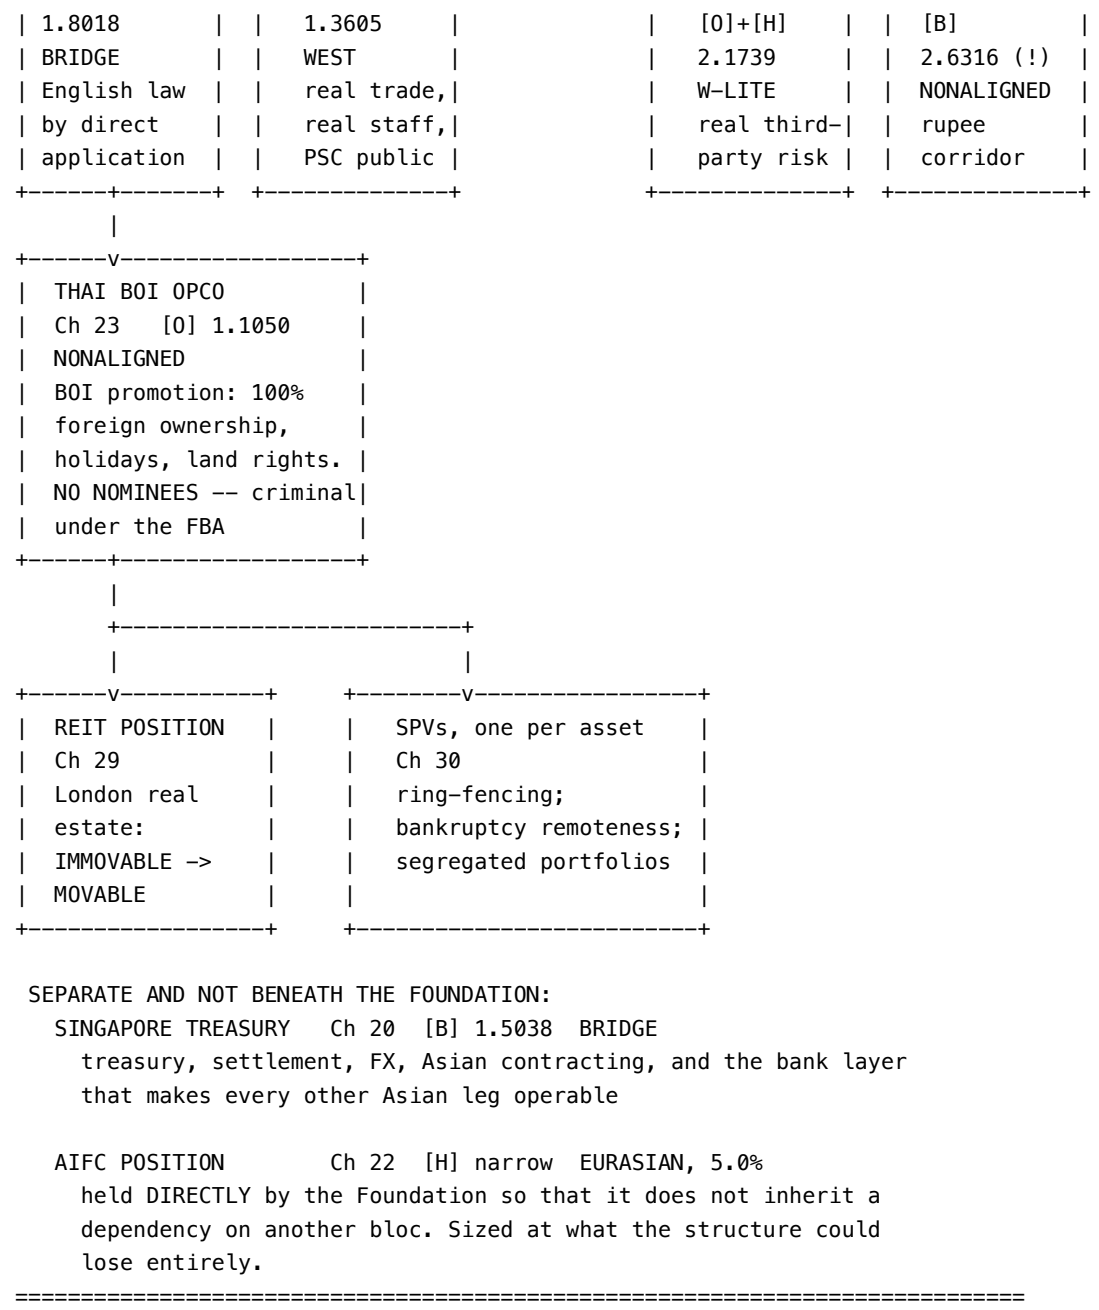
Every diagram in this appendix depicts the CONSTRUCTED carried case of the carried constructed case and is labelled as constructed. No real family, entity, or transaction is depicted.

The carried case: a dual national, UAE tax resident, with three heirs resident in three jurisdictions, one of them in a forced-heirship civil-law country. Assets: manufacturing operations in Thailand; a trading company in the UAE; intellectual property; commercial real estate in London; a listed-securities portfolio. Revenue in USD, EUR, CNY, AED, and THB.

Job labels throughout: **[H]** holding, **[O]** operating, **[B]** bridging. Where an entity carries more than one label the effective-jobs figure from Appendix B.11 is printed beside it, because an entity above 2.00 is carrying more than one job whether or not that was intended.

C.1 The structure as built





Four features of the diagram are load-bearing and are easy to miss.

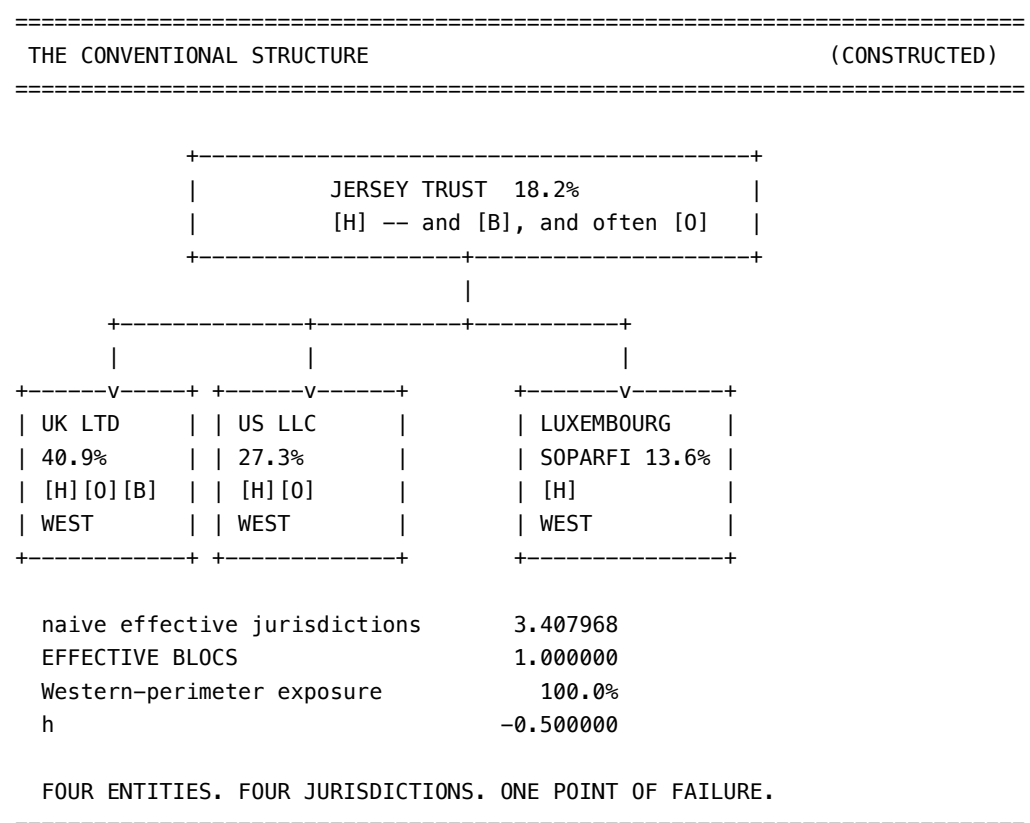
The Foundation holds; it does not transact. The store–rail rule applied at the apex: a rail is not a holding, and the apex holds no payment function. The banking sits at the DIFC and Singapore layers.

The AIFC leg hangs directly off the apex. If it hung beneath the ADGM holdco it would inherit the UAE's bloc position through its parent, and the entire point of the position is that it is in a different perimeter from everything above it.

The Bermuda captive is a sibling of the operating chain, not a child of it. Reserves held against real third-party liabilities are a structurally awkward class to realise, and that property survives only while the captive is a supervised insurer with real cedants — which is why the disclosure paper's disclosure test fails a captive that writes no risk.

The GIFT City vehicle is flagged in the diagram itself. At 2.6316 effective jobs it is above the 2.20 threshold and is carrying all three jobs. **The framework prints its own flag rather than quietly correcting the diagram**, and the GIFT City assessment's repair moves it to a specialised bridging vector of 1.3605.

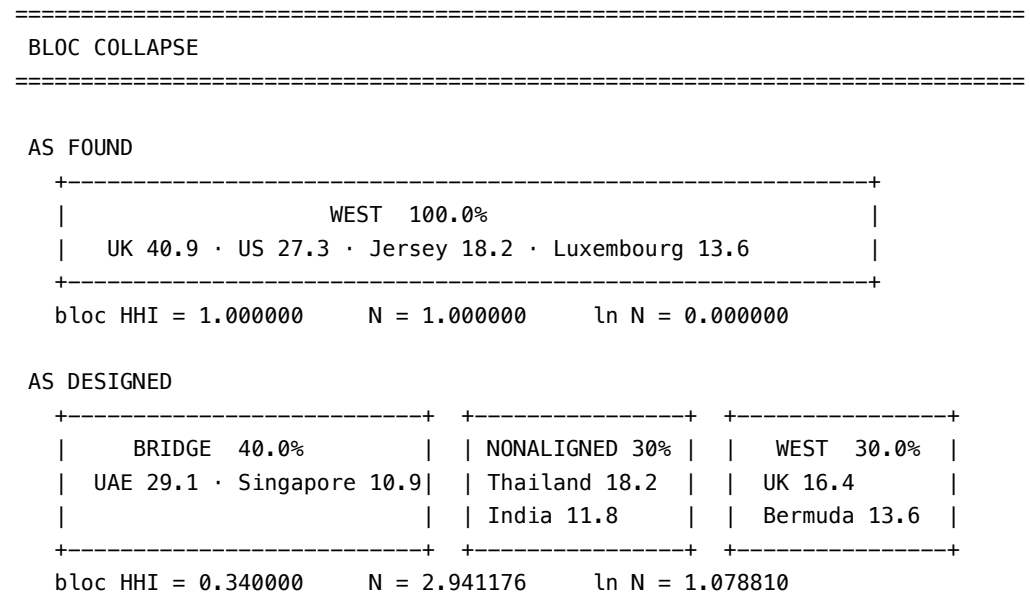
C.2 The structure as found — what most readers actually have



Nothing in that diagram is unlawful, concealed, or badly drafted. It is a competent structure built by competent advisers, and every one of its four legs sits inside the same coordination perimeter, which is a fact none of its filings would ever reveal.

C.3 The bloc view — the same structures, collapsed

The diagram that actually matters, and the one no organisation chart draws.



```
  AS DESIGNED, WITH THE AIFC LEG
    +----------------+  +-------------+  +-------------+  +----------+
    | BRIDGE  40.0%  |  | NONALIGNED  |  | WEST 25.0%  |  | EURASIAN |
    |                |  |   30.0%     |  |             |  |   5.0%   |
    +----------------+  +-------------+  +-------------+  +----------+
    bloc HHI = 0.315000     N = 3.174603     ln N = 1.155303
 ============================================================================
```

The number of boxes is `N` rounded down; the shape of the boxes is `N` exactly. A structure with four equal boxes would have `N = 4.000000` and clear the break-even. The designed structure has three unequal boxes and reaches 2.941176.

C.4 The rail view — the diagram that is drawn least and matters most

```
 ============================================================================
  TERMINAL CLEARING SYSTEMS                                  (CONSTRUCTED)
 ============================================================================

  AS FOUND -- every leg terminates in one place

     UK Ltd -----+
     US LLC -----+
     Jersey -----+----> USD correspondent chain ----> US CLEARING
     Lux    -----+

     HHI = 1.000000     EFFECTIVE SYSTEMS = 1.000000

  AS DESIGNED -- measured by terminal system, not by bank

     trading revenue    --> USD  55%  --> US clearing
     China corridor     --> CNY  15%  --> CIPS clearing (SWIFT messaging)
     European sales     --> EUR  12%  --> TARGET
     India corridor     --> INR  10%  --> Indian clearing via GIFT
     Gulf operations    --> AED   8%  --> UAE clearing

     HHI = 0.355800     EFFECTIVE SYSTEMS = 2.810568

  THE FOUR HARD LIMITS
    1  a USD account settles under US law wherever it is booked
    2  CIPS is CLEARING; many participants use SWIFT for MESSAGING
    3  the bank's risk appetite is not the jurisdiction's policy
    4  a rail is not a holding
 ============================================================================
```

Note what the diagram shows that the holding diagram cannot: **a structure can have `N = 2.941176` on holdings and 1.000000 on the ability to transact.**

C.5 The leverage layer

```
 ============================================================================
  THE CARVE-OUT CASCADE (Ch 34)                              (CONSTRUCTED)
 ============================================================================

     THAI OPCO, marked value 100.00
```

```
              |
    +---------+---------+
    |                   |
  20% IPO             80% RETAINED
  proceeds 20.00      pledged at 50% LTV
                           |
                      borrowing 40.00
                           |
    DEPLOYABLE = 20.00 + 40.00 = 60.00
    RETAINED CONTROL = 80.0%

    TRIGGER
      critical collateral value = 40.00 / 0.65 = 61.538462
      drawdown to trigger       = 23.0769%

    +-----------------------------------------------------------+
    |  23.1% IS ORDINARY EQUITY VOLATILITY.                     |
    |  No adversary is required. No court is required.          |
    |  The lender exercises rights the principal granted.       |
    +-----------------------------------------------------------+

    THE CASCADE BEYOND THE CALL
      forced sale  --> change-of-control definitions elsewhere
                   --> cross-default into unrelated agreements
                   --> ring fences breached from inside
================================================================================
```

C.6 The tempted variant — what not to build

Printed because a diagram of a mistake is more useful than a paragraph about one.

```
================================================================================
 AS TEMPTED -- eight jurisdictions, and less resilient   (CONSTRUCTED)
================================================================================

                       DIFC FOUNDATION
                              |
  +--------+--------+---------+---------+---------+---------+--------+
  |        |        |         |         |         |         |        |
 UAE     Thai    Bermuda     UK      India     SAUDI    QATAR    Singapore
20.0%   18.2%    13.6%     12.2%    11.8%      9.1%     8.2%      6.9%
BRIDGE  NONAL     WEST      WEST    NONAL     BRIDGE   BRIDGE    BRIDGE

  naive effective jurisdictions   7.132973   (+34.81% -- looks better)
  EFFECTIVE BLOCS                 2.841490   ( -3.39% -- IS worse)
  h                              -0.134485   (worse by 0.012068)

  +-------------------------------------------------------------+
  |  Every entity is lawful, disclosed, and recommended on its  |
  |  own merits. The structure is less resilient than the one it |
  |  replaced, and every report the family receives says the    |
  |  opposite.                                                  |
  +-------------------------------------------------------------+
================================================================================
```

C.7 How to read these diagrams

Four rules, and they are the whole method.

An organisation chart is not a resilience diagram. C.1 and C.2 look equally sensible and differ by a factor of 2.9412 in effective bloc count. Only C.3 shows it.

Label every box with its job before drawing any arrows. The three-jobs diagnostic's commonest failure — one entity doing three jobs — is visible in a labelled diagram and invisible in an unlabelled one.

Draw the rail diagram separately, and draw the mistake too. The rail diagram has a different topology from the holding diagram and reveals a concentration visible nowhere else; C.6 exists because the tempted structure is the one a family will actually be offered, one entity at a time, by advisers who are each individually right.

Appendix D. Statute and case index

D.1 What this appendix is, and the status key

This is the framework's integrity mechanism. Every section closes with a Verification Register listing the claims it did not independently confirm, together with what specifically needs checking and with whom. This appendix consolidates them, adds the citations from the sections written before the Register became mandatory, and assigns every item one of two statuses.

`VERIFIED`	Treated as established in this series and stated as given. Reserved for the facts in the verified-facts register and for arithmetic recomputed in Python from stated inputs. NOTHING ELSE CARRIES THIS STATUS.
`FLAGGED FOR VERIFICATION`	Named because it exists, and not confirmed by this series against the primary source. Every case reporter citation, every section number, every date, every threshold, and every current-status claim in this series is FLAGGED unless it appears in D.3.

The default is FLAGGED. That is deliberate and it is the honest default for a register that names statutes across thirty-odd jurisdictions. A reader who treats the flag as a formality has misunderstood the design: several sections state explicitly that a threshold or a section number in the text must not be relied on, and the AIFC assessment goes further and names no subordinate instrument of its jurisdiction at all rather than guess at one.

Three conventions inherited from the sections govern this index.

Declining to guess, in print, is the house standard. The AIFC assessment declines to say whether the AIFC foundation confers separate legal personality; the foundation-versus-trust comparison declines to name the citations in the illusory-trust line; the erosion paper declines to give section numbers for the Cook Islands statute; the failure record repeats all three refusals rather than resolving them. **A refusal recorded here is information, not an omission.**

Commonly cited as. Where the framework gives a citation that circulates widely but that it has not checked, it writes "commonly cited as" and flags it. Examples in the index below include Article 694 of the Panamanian Fiscal Code, Article 17 of Federal Decree-Law No. 47 of 2022, article L.132-13 of the French Code des assurances, sections 56 and 57 of the Central

Bank of Malaysia Act 2009, section 80LA of the Indian Income-tax Act, 1961, and Item 403 of Regulation S-K.

Constructed model parameters are not data. Cost indices, tax rates in worked models, discount rates, exchange rates, maintenance requirements, and the carried constructed case test parameters are constructed. They are labelled as such wherever they appear and they are not citations.

D.2 The count

Compiled by extracting the numbered items from every section headed “Verification Register” or “Verification notes” across the manuscript.

```
sections carrying a Verification Register        37
consolidated FLAGGED items                      577
sections with no register (pre-dating V1)         2   (Ch 1, Ch 2)
items in this appendix's own cross-cutting
  index that consolidate multiple section
  items into one entry                          see D.4–D.7
```

Per section:

paper	assessment	items	paper	assessment	items
051	Surveillance Irrelevance	6	058	GIFT City	18
052	Five Threats	4	059	Türkiye	18
052	The Concentration Instrument	4	059	Mauritius and Africa	14
052	Holding, Operating, Bridging	6	059	Labuan and Islamic Finance	14
053	The DIFC Foundation	12	060	Latin America	18
053	Foundation versus Trust	14	060	Immovable to Movable	24
053	Succession and Forced Heirship	15	062	The Carve-Out Cascade	16
054	The City	12	062	Group Architecture	20
054	The Crown Dependency Error	14	062	Law-Family Arbitrage	17
054	The Convergence	11	063	The Medici System	22
055	Bermuda	14	064	The Rothschild System	27
055	Captives and Sidecars	13	064	The Limits of Double Entry	20
056	The Atlantic Cluster	12	064	Encoding Integrity	28

paper	assessment	items	paper	assessment	items
056	ADGM	16	065	The VINO System	24
056	DIFC Beyond the Foundation	12	069	When Structures Fail	22
057	Qatar and Saudi	16	070	The Manual in One Page	15
057	Singapore	16			
057	Hong Kong	18		**TOTAL**	**577**
058	AIFC Astana	25			
058	Thailand	20			

The AIFC assessment carries the longest register in the series at twenty-five items, deliberately, and opens with an explicit warning that it names and numbers no AIFC subordinate Act, Regulation, or Rule. **It sets the standard the rest of the framework is measured against, and the length of its register is a strength rather than a weakness.**

The count is regenerable. Extract every line matching a numbered-list pattern between a heading matching “Verification Register” or “Verification notes” and the next scorecard heading, across every section file, and sum. **A reader who re-runs that extraction on a later revision of the manuscript and gets a different number has found a change, not an error in this appendix.**

D.3 VERIFIED — the register of verified facts, and nothing else

These are stated as given throughout the framework. They are the only substantive claims not flagged. Every one of them should still be confirmed with counsel before reliance; “verified” here means “fixed as canon within this series,” not “checked against the primary source on the reader’s behalf.”

DIFC Foundation	DIFC Foundations Law No. 3 of 2018. Separate legal personality; purpose OR beneficiary; founder may retain control; DIFC Courts jurisdiction.
ADGM	Application of English Law Regulations 2015 — directly applies English common law, the only jurisdiction outside the UK to do so. ADGM Courts; FSRA regulator.
DIFC	Own codified common-law-derived statutes; DIFC Courts; DFSA regulator.
UAE	FATF grey list March 2022 to February 2024, now removed. CRS participant.
Thailand	Civil law (Civil and Commercial Code). US–Thai Treaty of Amity 1966: US persons may own 100% of a Thai company, exempt from most Foreign Business Act restrictions. BOI promotion can confer 100% foreign ownership, tax holidays and land rights. NOMINEE SHAREHOLDING IS A CRIMINAL OFFENCE AND IS PROSECUTED. Revenue Department order Por. 161/2566: from 1 January 2024 foreign-source income remitted by a Thai tax resident is taxable regardless of the year earned. CRS participant.
GIFT City	IFSC; unified regulator IFSCA since 2020; funds, aircraft and ship leasing, reinsurance, bullion. India is CRS-participating.
Bermuda	BMA regulator; Solvency II equivalent; economic substance

```
                    requirements since 2018; CRS participant. Reinsurance
                    reserves held against real third-party liabilities.
United Kingdom      Public PSC register; ECCTA 2023 tightened Companies House
                    identity verification. English law is the most-chosen
                    governing law for international commercial contracts.
Luxembourg          CJEU, November 2022 (*WM and Sovim SA v Luxembourg
                    Business Registers*): PUBLIC access to beneficial-
                    ownership registers struck down; competent authorities
                    retain access.
CIPS                A CLEARING system. Many participants still use SWIFT for
                    the MESSAGING layer. "CIPS replaces SWIFT" is FALSE.
2022 precedent      Approximately USD 300bn of Russian central bank reserves
                    immobilised.
Stablecoins         MiCA stablecoin provisions applicable from 30 June 2024;
                    USDT delisted from EU venues. GENIUS Act signed July 2025.
                    Tether can and does freeze USDT at law-enforcement
                    request. T3 Financial Crimes Unit (Tether, Tron, TRM Labs)
                    launched September 2024.
Public blockchains  MORE traceable than correspondent banking, not less.
```

And the computed results, which are VERIFIED in the narrower sense that they were re-computed in Python from the stated inputs at full precision for Appendix B and reproduce exactly:

```
conventional structure   naive HHI 0.293430, 3.407968 jurisdictions,
                         bloc HHI 1.000000, N = 1.000000, h = -0.500000
designed structure       naive HHI 0.189002, 5.290949 jurisdictions,
                         bloc HHI 0.340000, N = 2.941176, h = -0.122417
break-even               N* = 4.172734
V5 tempted               naive 7.132973 (+34.8146%), blocs 2.841490 (-3.3893%)
V5 as built              naive 5.688217, blocs 2.816901, priced loss 0.124275 in N
V6 correspondent         1.000000 -> 2.810568 effective clearing systems
V4 AIFC                  G = 0.026731, L = 0.005110, p* = 16.0476%,
                         cost ceiling 8.9102% of C, shortfall closed 21.8357%
V3 Singapore control     marginal contribution to h = 0.002014
V2 Hong Kong             centre +0.072279; capped transit +0.002869; ratio 3.9699%
Ch 34 margin             trigger at 23.0769% drawdown
V9 crossover             USD 34.2857m / 80.0000m / 120.0000m at 35 / 15 / 10 bp
V8 succession            6.70m capped, 27.00m tax, ratio 4.0299x, shortfall 14.1000%
```

The parameters those results were computed from — `eta = 0.35`, `C = 0.30`, `Q = 10`, `delta = 0.02`, and every share in every allocation — are constructed test values and are not verified estimates of anything. The arithmetic is verified; the inputs are constructed.

D.4 Cases — every case named anywhere in the series

All FLAGGED FOR VERIFICATION. This series has checked no reporter citation against an official report, and says so. For each: confirm the citation, the court, the date, the holding, and its subsequent treatment before any reliance.

Contempt, repatriation, and the impossibility defence — Ch 1, 4, 44

```
FTC v. Affordable Media, LLC, 179 F.3d 1228 (9th Cir. 1999)
In re Lawrence, 279 F.3d 1294 (11th Cir. 2002)
  lower courts: 227 B.R. 907 (Bankr. S.D. Fla. 1998); 251 B.R. 630 (S.D. Fla. 2000)
  NOTE: the trust was settled in JERSEY in 1991, administration later moved to
```

```
  MAURITIUS. It was not a Cook Islands trust. Confirm from the record.
In re Portnoy, 201 B.R. 685 (Bankr. S.D.N.Y. 1996)
In re Brooks, 217 B.R. 98 (Bankr. D. Conn. 1998)
Grupo Mexicano de Desarrollo, S.A. v. Alliance Bond Fund, Inc., 527 U.S. 308 (1999)
Mareva Compania Naviera SA v International Bulkcarriers SA [1975] 2 Lloyd's Rep 509
Hill v Spread Trustee Co Ltd [2006] EWCA Civ 542
```

Separate personality, sham, and trusts — Ch 6, 8, 11, 36, 44

```
Salomon v A Salomon & Co Ltd [1897] AC 22
Adams v Cape Industries plc [1990] Ch 433
Prest v Petrodel Resources Ltd [2013] UKSC 34
Snook v London and West Riding Investments Ltd [1967] 2 QB 786
Saunders v Vautier
Schmidt v Rosewood Trust Ltd
Rahman v Chase Bank (CI) Trust Company Limited [1991] JLR 103
In re the Esteem Settlement [2003] JLR 188
Crociani v Crociani [2014] UKPC 40
Investec Trust (Guernsey) Limited v Glenalla Properties Limited [2018] UKPC 7
Pitt v Holt; Futter v Futter [2013] UKSC 26
Attorney General for Bermuda v Ferguson
the New Zealand and Privy Council "illusory trust" line -- CITATIONS DELIBERATELY
  NOT NAMED, per Ch 8 Register item 14 and Ch 44 Register item 8
```

Insurance and captives — Ch 3, 15

```
Helvering v. Le Gierse (1941)
Humana Inc. v. Commissioner (6th Cir. 1989)
AMERCO, Inc. v. Commissioner (1991); Harper Group v. Commissioner (1991)
Rent-A-Center, Inc. v. Commissioner (2014)
Securitas Holdings, Inc. v. Commissioner (2014)
R.V.I. Guaranty Co. v. Commissioner (2015)
Avrahami v. Commissioner, 149 T.C. 144 (2017)
Reserve Mechanical Corp. v. Commissioner (2018), affirmed 10th Cir. (2022)
Syzygy Insurance Co. v. Commissioner (2019)
Caylor Land & Development, Inc. v. Commissioner (2021)
CIC Services, LLC v. IRS (S. Ct. 2021), and the district court order setting
  aside Notice 2016-66
Green Valley Investors (2022); the Puglisi cases -- dispositions to be confirmed
Halliburton Company v Chubb Bermuda Insurance Ltd
```

European Union — Ch 12

```
WM and Sovim SA v Luxembourg Business Registers, Joined Cases C-37/20 and
  C-601/20, 22 November 2022
Fiat Chrysler Finance Europe, Joined Cases C-885/19 P and C-898/19 P (8 Nov 2022)
Commission v Amazon.com, C-457/21 P (14 December 2023)
Luxembourg and Engie, Joined Cases C-451/21 P and C-454/21 P (5 December 2023)
Commission v Ireland and Others, C-465/20 P (10 September 2024)
the Starbucks annulment, T-760/15 and T-636/16 (24 September 2019)
Orde van Vlaamse Balies, C-694/20 (8 December 2022)
Halet v Luxembourg, ECtHR Grand Chamber, 14 February 2023
the February 2019 beneficial-ownership judgments on the Interest and Royalties
  and Parent-Subsidiary Directives -- joined case numbers to be confirmed
```

Asia and the Gulf — Ch 21, 24, 27

```
HKSAR v Lai Chee Ying [2021] HKCFA 3
Ng Ka Ling v Director of Immigration (1999) 2 HKCFAR 4
Democratic Republic of the Congo v FG Hemisphere Associates LLC (2011) 14 HKCFAR 95
Vodafone International Holdings BV v Union of India
JRI Resources Sdn Bhd v Kuwait Finance House (Malaysia) Bhd (Federal Court, 2019)
The Investment Dar Company KSCC v Blom Developments Bank SAL [2009] EWHC 3545 (Ch)
Shamil Bank of Bahrain EC v Beximco Pharmaceuticals Ltd [2004] EWCA Civ 19
the Dana Gas proceedings; the Golden Belt proceedings
```

Offshore, competition, and property — Ch 16, 29, 35, 28

```
Convoy Collateral Ltd v Broad Idea International Ltd [2021] UKPC 24
British South Africa Co v Companhia de Moçambique [1893] AC 602
United States v. AT&T Inc. -- district and appellate citations to be confirmed
Philip Morris Brands Sàrl and others v. Oriental Republic of Uruguay,
  ICSID Case No. ARB/10/7
National Crime Agency v Baker; Hajiyeva v National Crime Agency
R (AAA) v Secretary of State for the Home Department
```

The Discovery line and its descendants — Ch 40

```
Johnson v. M'Intosh (1823)
Cherokee Nation v. Georgia; Worcester v. Georgia
City of Sherrill v. Oneida Indian Nation of New York (2005)
Mabo v. Queensland (No 2); Delgamuukw v. British Columbia;
  Tsilhqot'in Nation v. British Columbia
```

D.5 Statutes and instruments — the consolidated index

All FLAGGED FOR VERIFICATION, in every case for current consolidated text, amendments, commencement, and any successor instrument. Grouped by jurisdiction. This index names the instruments; the section registers name what specifically must be checked in each.

United Arab Emirates — DIFC Foundations Law No. 3 of 2018; DIFC Trust Law, DIFC Law No. 4 of 2018; DIFC Law No. 3 of 2004 (application-of-law cascade, commonly cited as Article 8(2)); Dubai Law No. 12 of 2004 as amended by Dubai Law No. 16 of 2011; Federal Law No. 8 of 2004; Federal Decree No. 35 of 2004; Dubai Law No. 9 of 2004; Dubai Law No. 15 of 2017; ADGM Application of English Law Regulations 2015; ADGM Courts Regulations 2015; ADGM Companies Regulations; ADGM Foundations Regulations 2017; Financial Services and Markets Regulations 2015; Federal Decree-Law No. 47 of 2022 (corporate tax; family foundation article commonly cited as Article 17); Federal Decree-Law No. 20 of 2018; Cabinet Decision No. 10 of 2019; Federal Cabinet Resolution No. 58 of 2020; Federal Decree-Law No. 41 of 2022 (Civil Personal Status); Federal Law No. 28 of 2005; the UAE domestic minimum top-up tax instrument; the 2024 DIFC digital assets and security package.

United Kingdom and Crown Dependencies — Economic Crime and Corporate Transparency Act 2023; Economic Crime (Transparency and Enforcement) Act 2022 and the Register of Overseas Entities; Insolvency Act 1986 s.423; Limitation Act 1980; Arbitration Act 1996 and Arbitration Act 2025; Sanctions and Anti-Money Laundering Act 2018; Public Offers and Admissions to Trading Regulations 2024; UK Listing Rules; DTR 5; Finance Act 2006 Part 4 and Part 12 of the Corporation Tax Act 2010 (UK REIT); Finance Act 2003 Part 4 (SDLT); Inheritance (Provision for Family and Dependants) Act 1975; Trusts (Jersey) Law 1984 (firewall, reserved powers, and the mistake and inadequate- deliberation jurisdiction, commonly cited as Articles 9, 9A and 47A–47J); Taxation (Companies — Economic Substance) (Jersey) Law 2019 and the Guernsey and Isle of Man equivalents made in 2018; Trusts (Guernsey) Law; Guernsey Protected Cell Companies Ordinance 1997; Isle of Man Companies Act 1931 and Companies Act 2006; Isle of Man Foundations Act 2011 and Purpose Trusts Act 1996.

United States — Foreign Account Tax Compliance Act (18 March 2010, Title V of the Hiring Incentives to Restore Employment Act); Bankruptcy Code s.548(e) as added by the Bankruptcy Abuse Prevention and Consumer Protection Act of 2005; the Uniform Voidable Transactions Act as enacted state by state; Corporate Transparency Act (enacted 1 January 2021 within the National Defense Authorization Act) — **scope genuinely unsettled and actively litigated**; IRC ss.831(b), 501(c)(15), 856–860, 721, 704(c), 897, 1445, 4371, 6104; the Protecting Americans from Tax Hikes Act of 2015; the REIT Modernization Act of 1999; the 1960 REIT legislation; the January 2025 micro-captive final regulations; Rev. Rul. 2002-89, 2002-90, 2002-91, 2005-40; Notice 2016-66; Liability Risk Retention Act of 1986; Clayton Act s.8 (15 U.S.C. § 19); Hart-Scott-Rodino (15 U.S.C. § 18a; 16 C.F.R. Parts 801–803) and the 2024 form overhaul; Sherman Act s.1 (15 U.S.C. § 1); the 2023 DOJ/FTC Merger Guidelines and the withdrawal of the 2020 Vertical Merger Guidelines; Regulation S-K Items 403 and 407(i); Schedule 13D Item 6; Securities Act Rule 144; Exchange Act ss.13(a) and 16(b); Federal Reserve Regulation T (12 C.F.R. Part 220); FINRA Rule 4210; UCC Article 8; FIPS 180-4.

European Union — Directive (EU) 2015/849; Directive (EU) 2018/843; Regulation (EU) 2024/1624; Directive (EU) 2024/1640; Regulation (EU) 2024/1620; Council Directive (EU) 2022/2523 (Pillar Two); Council Directive (EU) 2016/1164 (ATAD I) and (EU) 2017/952 (ATAD II); DAC6 (2018/822), DAC7 (2021/514), DAC8 (2023/2226); Regulation (EU) No 1215/2012 (Brussels Recast) Article 24(1); Regulation (EU) No 650/2012 (Succession); Regulation (EU) 2016/1103 (matrimonial property); Directive 98/26/EC Article 9(2); Directive 2002/47/EC; Commission Delegated Decision (EU) 2016/309; the EU list of non-cooperative jurisdictions, first adopted 5 December 2017, Annexes I and II.

Offshore and overseas territories — Cook Islands International Trusts Act 1984 as amended (**section numbers deliberately not stated**); Cayman Companies Act, Trusts Act, Mutual Funds Act, Private Funds Act 2020, Foundation Companies Act 2017, Insurance Act 2010, International Tax Co-operation (Economic Substance) Act, Beneficial Ownership Transparency Act 2023; BVI Business Companies Act 2004 as amended, Economic Substance (Companies and Limited Partnerships) Act 2018; Bahamian Foundations Act 2004, Purpose Trust Act 2004, Executive Entities Act 2011, Fraudulent Dispositions Act 1991, Investment Condominium Act 2014, Commercial Entities (Substance Requirements) Act 2018; Barbados Insurance Act and the 2018 convergence and 2024 rate legislation; Bermuda Insurance Act 1978, Economic Substance Act 2018, Segregated Accounts Companies Act 2000, Corporate Income Tax Act 2023, Judgments (Reciprocal Enforcement) Act 1958, Bermuda Constitution Order 1968.

Asia-Pacific — Singapore: Application of English Law Act 1993, Variable Capital Companies Act 2018, Income Tax Act 1947 ss.13D/13O/13U, Trustees Act 1967, Trust Companies Act 2005, Companies Act 1967, Reciprocal Enforcement of Foreign Judgments Act 1959, International Arbitration Act 1994, Supreme Court of Judicature Act 1969, Financial Services and Markets Act 2022, Multinational Enterprise (Minimum Tax) Act 2024, Administration of Muslim Law Act. Hong Kong: the National Security Law, the Safeguarding National Security Ordinance 2024, Basic Law Articles 8, 104 and 158, Companies Ordinance (Cap. 622), Inland Revenue Ordinance, Mainland Judgments in Civil and Commercial Matters (Reciprocal Enforcement) Ordinance. Thailand: Foreign Business Act B.E. 2542 (1999), Civil and Commercial Code, Land Code, Condominium Act, Revenue Code, Revenue Department order Por. 161/2566, the investment promotion legislation. India: International Financial Services Centres Authority Act, 2019; Special Economic Zones Act, 2005; Foreign Exchange Management Act, 1999; IFSCA (Fund Management) Regulations, 2022; Income-tax Act, 1961 (deduction commonly cited as s.80LA; exemption commonly cited as s.47(viiab)); Code of Civil Procedure, 1908 ss.13 and 44A; Arbitration and Conciliation Act, 1996; Taxation Laws (Amendment) Act, 2021. Malaysia: Labuan Business Activity Tax Regulations (2018, replaced 2021) — P.U.(A) numbers to be confirmed; Income Tax Act 1967; Central Bank of Malaysia Act 2009 (Shariah Advisory Council, commonly cited as ss.56 and 57). Kazakhstan: the AIFC constitutional statute of 7 December 2015 and the constitutional entrenchment — **no AIFC subordinate Act, Regulation or Rule is named anywhere in this series.**

Continental Europe and civil-law succession — French Code civil art. 913 and the compen-

satory levy inserted by Loi n° 2021-1109 of 24 August 2021; Loi n° 2006-728 of 23 June 2006; French advance renunciation of the action in reduction; Code général des impôts arts. 792-0 bis, 123 bis, 669, 751, 1133; Code des assurances art. L.132-13 (commonly cited); German BGB §§2303, 2325, 2329, 2346 and Aussensteuergesetz §15; Italian Codice civile (*legittima, azione di riduzione*); Spanish Código Civil (*legítima*, *mejora*, *foral* regimes); Swiss Banking Act art. 47 of 8 November 1934, AEOI Act of 18 December 2015, Embargo Act of 22 March 2002; Luxembourg law of 22 December 2023 and its December 2024 amendment; Luxembourg SIF (13 February 2007), SICAR (15 June 2004), RAIF (23 July 2016), securitisation law of 22 March 2004; Irish ICAV Act 2015, Finance (No. 2) Act 2023, s.110 TCA 1997; Dutch Withholding Tax Act 2021.

Latin America — Panama: Law 32 of 1927, Law 25 of 1995, Law 21 of 2017, Law 23 of 2015, Law 52 of 2016, Law 129 of 2020, Law 254 of 2021, Fiscal Code territoriality (commonly cited as Article 694). Uruguay: Law No. 15.921 of 1987, Law No. 18.930 of 2012, Law No. 19.484 of 2017. Paraguay: Law No. 6.380/2019, Law No. 1.064/1997. Costa Rica: Law No. 7210 of 1990, and the 2023 amending law and 2021 residence statute (commonly cited as Law No. 10.381 and Law No. 9996).

Türkiye — Decree No. 32 on the Protection of the Value of Turkish Currency (1989); Tax Procedure Law No. 213; Civil Code No. 4721 (2001); Code of Obligations No. 6098 (2011); Commercial Code No. 6102 (2011); Foreign Direct Investment Law No. 4875 (2003); Law No. 7412 on the Istanbul Finance Centre (2022); Banking Law No. 5411 (2005); Capital Markets Law No. 6362 (2012); Law No. 6493 (2013); Land Registry Law No. 2644; Law No. 5718 on Private International Law and Procedural Law (2007); International Arbitration Law No. 4686 (2001).

Africa and the Indian Ocean — Mauritius: the DTAA and IPPA network, the India Protocol of 10 May 2016 and the Protocol signed 7 March 2024, CBDT Circulars No. 682 (1994) and No. 789 (2000). Rwanda: company law, investment code, trusts, partnerships and foundations legislation — numbers to be confirmed. Kenya: Nairobi International Financial Centre Act, 2017; Companies (Beneficial Ownership Information) Regulations, 2020; Judicature Act; Constitution of Kenya, 2010.

D.6 Treaties, conventions, and multilateral instruments

All FLAGGED FOR VERIFICATION as to contracting-party lists, entry-into-force dates, declarations, reservations, and territorial extensions, in every relevant state and on the date of reliance.

```
Hague Convention of 1 July 1985 on the Law Applicable to Trusts and on their
  Recognition
Hague Convention of 30 June 2005 on Choice of Court Agreements
Hague Convention of 2 July 2019 on the Recognition and Enforcement of Foreign
  Judgments
Hague Convention of 1 August 1989 on the Law Applicable to Succession —
  understood NEVER to have entered into force; confirm before any reference
Hague Securities Convention — confirm status in each relevant state
New York Convention 1958 on the Recognition and Enforcement of Foreign
  Arbitral Awards
Singapore Convention on Mediation
Lugano Convention 2007
Convention on Mutual Administrative Assistance in Tax Matters (1988, amended
  by protocol 2010)
the Common Reporting Standard: approved by the OECD Council 15 July 2014; the
  Multilateral Competent Authority Agreement first signed in Berlin
  29 October 2014
the Multilateral Instrument and the principal purpose test
US–Thai Treaty of Amity and Economic Relations 1966
the EU–Türkiye customs union and its Association Council decision
```

```
the EU Enhanced Partnership and Cooperation Agreement with Kazakhstan
the Montreux Convention of 1936
ICSID Convention —— current status of each relevant state
```

D.7 The two sections with no Verification Register

The erosion and coordination-perimeter assessments pre-date the requirement introduced at RECONCILE V1 and carry no register. Their citations are consolidated here and are **FLAGGED FOR VERIFICATION** on the same terms as everything else.

the erosion paper — the UBS deferred prosecution agreement of February 2009 and the USD 780 million payment; the 2009 US-Switzerland agreement and the account figure usually reported as approximately 4,450; the Swiss FATCA intergovernmental agreement on the Model 2 template, date usually given as 14 February 2013; the Wegelin & Co. guilty plea, January 2013; the Credit Suisse guilty plea, May 2014, and the USD 2.6 billion payment; the Department of Justice Program for Non-Prosecution Agreements or Non-Prosecution Letters for Swiss Banks announced 29 August 2013, with aggregate figures to be confirmed against Department of Justice releases; the Falciani data and the February 2015 publication; the Panama Papers (April 2016), Paradise Papers (November 2017), FinCEN Files (September 2020), Pandora Papers (October 2021); every listing and delisting date in the erosion table; and the Cook Islands International Trusts Act 1984 with its amendment history.

the coordination-perimeter analysis — the five coordination interfaces and the institutions named at each; the FATF and Global Forum sequences relied upon; the EU Annex I and Annex II mechanics; the correspondent de-risking evidence; and the bloc taxonomy itself, which is the framework's own construction and is **contestable by design**.

D.8 The gap — sections specified and not yet on the page

Stated rather than implied. The source manifest specifies a complete method; the manuscript from which this index was compiled does not yet contain every section it specifies. **Where the closing sections and these appendices refer forward to a section that is not on the page, they refer to a result fixed in the carried constructed case or the reconciliation ledger — the 23.0769% margin trigger, the four bright lines, the CIPS distinction, the 2022 precedent — and not to text a reader can check.**

Treat every such forward reference as a specification, not as a citation. The results themselves are in Appendix B and are recomputed there; what is missing is the section that argues them.

D.9 How to use this appendix

Four instructions.

Take the relevant section's register to counsel, not this appendix. The registers say what specifically needs checking and with whom; this index says what exists. And treat every section number as provisional — numbering moves across consolidations, and several sections state in terms that no threshold in the text should be relied on.

Where the framework declined to guess, do not fill the gap from a secondary source. The *In re Lawrence* situs error is the demonstration: a source that gets it wrong is unreliable on everything else it says.

Re-run the count. A structure built against this synthesis should carry its own register, derived from these, with a date and a named person against each item.

Appendix E. The honesty register

This appendix collects, in one place, everything the framework does not know, everything it got wrong and corrected, and everything it recommends while reporting that the recommendation does not fully work. **It is the appendix to read if only one appendix is read.**

E.1 What is settled

Short, because little is.

The direction of travel. Fifteen years of documented erosion across four independent tracks — Swiss secrecy, Caribbean substance, Panama's listing cycle, and the collapse of the non-recognition defence against personal compulsion — is not a matter of interpretation. The dates are public and the instruments are named in Appendix D.

The mechanism. Erosion is coordination. Jurisdictions inside one coordination perimeter changed together because the same five interfaces reached all of them. That is an observation about sequence and correlation, and it is testable.

The arithmetic. `HHI = sum of squared shares`, `effective count = 1/HHI`, and `ln(1.00) = 0`. Appendix B reproduces every figure from stated inputs. **The arithmetic is settled; every input to it is constructed.**

The three jobs. Holding, operating, and bridging need different things, and most failed structures come from asking one entity for all three. That is a taxonomy rather than a claim, and its usefulness does not depend on any jurisdiction assessment being right.

The personal-jurisdiction lesson. A court with personal jurisdiction over the principal does not need jurisdiction over the assets. Two appellate decisions establish it and no statute anywhere has displaced it.

E.2 What is contested

The bloc assignments. Every value of `N` in this series rests on a mapping from jurisdictions to coordination perimeters that the framework made. It is an argument, not a measurement. The two most contestable calls are Singapore (bridge or Western) and the treatment of the UAE and Singapore as one bridge bloc rather than two. The concentration paper's robustness table prices both: the Singapore call moves `h` by 0.002014; separating the UAE and Singapore moves `N` from 2.941176 to 3.615826, which is large, and is why the published 2.941176 is the conservative end of a range rather than a point estimate.

Whether Kazakhstan is an independent perimeter. Unresolved and priced rather than argued. See V4.

Whether the AIFC foundation confers separate legal personality. The AIFC assessment declines to guess in print. It remains unanswered and it gates any apex migration.

The conduit question. Whether a foreign judgment recognised inside a UAE financial free zone can be enforced onshore through it. Carried unresolved from the foundation paper through the rail-perimeter measurement and into this synthesis. **It is the largest open legal question in the architecture and it sits directly under the apex.**

The illusory-trust line. Litigated in New Zealand and before the Privy Council and deliberately not cited in either direction anywhere in this series.

The Corporate Transparency Act's scope. Genuinely unsettled and actively litigated.

Sanctions and secondary-sanctions exposure. Moving continuously, jurisdiction by jurisdiction, and no statement in this series about any current designation or listing should be assumed live.

E.3 What needs counsel

Everything, and this is not a formula. The specific list, distilled from 577 register items:

- counsel in each jurisdiction of incorporation
- counsel in each jurisdiction of residence -- principal AND every heir
- counsel in each jurisdiction of asset situs, separately for immovables
- tax counsel wherever any decision concerning any entity may be taken, because that is where corporate residence is asserted and where the most reliably fatal failure track begins
- Sharia-qualified counsel wherever personal-status law applies
- securities counsel before any pledge, carve-out, or listing
- a credentialed actuary before any captive
- antitrust counsel before any board overlap or above-threshold acquisition
- written CRS and FATCA classification of every entity AT FORMATION, including every sub-fund and cell

Two items are singled out in the sections as the most consequential in their registers and are repeated here. **Regularising a historical Thai nominee position** — exposure for the elapsed period, available routes, disclosure obligations — is named in the Thailand assessment as the most important item in that register, because the underlying conduct is criminal rather than merely ineffective. And **the exact scope of the Foreign Exchange Management Act, 1999 treatment of IFSC units as persons resident outside India** is named in the GIFT City assessment as the load-bearing item in that section, because the entire rupee-corridor case rests on it.

E.4 What changed recently and may change again

A structure built on any line below is built on a moving object.

2018–2019	economic substance regimes across the offshore world
2020	Cayman EU listing and delisting inside eight months
22 Nov 2022	CJEU strikes down PUBLIC beneficial-ownership register access; the Netherlands and Luxembourg suspend public access the same day
Mar 2022	UAE FATF grey listing -- removed February 2024
Mar 2022	Singapore adopts targeted measures without a Security Council resolution, reversing six decades of stated policy
2022	approximately USD 300bn of Russian central bank reserves immobilised
2022–	DOJ revives Clayton Act s.8 enforcement; resignations follow
2023	UAE corporate tax; Saudi Civil Transactions Law; Cayman Beneficial Ownership Transparency Act; ECCTA in the UK
Jan 2024	Thailand's Por. 161/2566 kills later-year remittance planning
Feb 2024	UAE removed from the FATF grey list; BVI removed from EU Annex I
30 Jun 2024	MiCA stablecoin provisions apply; USDT delisted from EU venues
Sep 2024	T3 Financial Crimes Unit launched
2024	EU AML package adopted; HSR notification form overhauled
Jun 2024	Türkiye removed from the FATF increased-monitoring list
2024–2026	Pillar Two implementation staggered across jurisdictions; DAC8 from 1 January 2026
Jul 2025	GENIUS Act signed
Jan 2025	micro-captive final regulations

The list is not a history. It is a rate. Roughly one structurally significant change per quarter across the jurisdictions this series covers, which is why the procedure in this synthesis ends with an annual re-run and why every scorecard item 12 names the specific event that would reverse the verdict.

E.5 The reconciliation ledger, V1 through V10

The framework's own corrections, in full, because a synthesis that hides its revisions is not auditable.

V1 — The Verification Register is mandatory. Introduced by an agent drafting the Singapore and Hong Kong assessments to 22, not by the original instruction. Counts at origin: Ch 20 = 16, Ch 21 = 18, Ch 22 = 25. **It is better than the instruction it replaced and it is now required of every section**; a jurisdiction section without one is incomplete. The AIFC assessment sets the standard: it opens by warning that it names and numbers no AIFC subordinate instrument, and states in print that it declines to guess whether the AIFC foundation confers separate legal personality. **Declining to guess, in print, is the house standard.**

V2 — Hong Kong: the best number is not the recommendation. Hong Kong in the 10.9% slot as a China-bloc asset gives `N = 3.615826` and `h = \ensuremath{-}0.050138`, a gain of 0.072279 — the largest single-move gain in the series. The section recommends against it and proposes capped transit instead: `N = 2.965388`, `h = \ensuremath{-}0.119547`, gain 0.002869 — **3.97% of the benefit for 3.21% of the exposure.** Political risk stated without euphemism, balanced by the documented absence of interference in ordinary commercial litigation, with the argument that **protection is bought for the tail, not the median.** That framing governs the assembled architecture and this synthesis.

V3 — Singapore makes the concentration paper's failure mode observable. The same 10.9% in London versus Singapore gives identical bloc shares, identical `N`, identical `h`, while naive effective jurisdictions read 5.290949 against 4.449309. Singapore's true marginal contribution to `h` is 0.002014. **The naive metric reports a diversification that did not occur.** Singapore's item 10 also splits three ways and the section refuses to average them; **the split-trajectory convention now applies wherever item 10 is genuinely non-uniform.**

V4 — the AIFC reduces to a probability threshold. AIFC at 5% from the UK slot: `N = 3.174603`, `h = \ensuremath{-}0.095686`, gain 0.026731. **Break-even probability that Kazakhstan is an independent perimeter = 16.05%. Getting the bloc call wrong — folding the AIFC into non-aligned — produces a LOSS of 0.005110.** The binding constraint is cost, not geopolitics: the leg must run under 8.9102% of total `C`. **Generalised: every marginal jurisdiction is priced as a break-even probability, not argued as a geopolitical opinion.**

V5 — more jurisdictions, fewer blocs. Adding Qatar and Saudi Arabia: eight jurisdictions, naive HHI 0.140194 giving 7.132973 effective jurisdictions, up 34.81% — while bloc HHI 0.351928 gives 2.841490 effective blocs, down 3.39%. **Two more jurisdictions, a third more apparent diversification, and the honest count fell.** The worked outcome takes the Saudi RHQ from the operating allocation instead: naive 5.688217, blocs 2.816901 — a deliberate, priced loss of 0.124275 in `N` taken on revenue grounds. **Priced losses taken knowingly are legitimate; unpriced ones are the failure.**

V6 — the dollar share cannot be engineered away. Correspondent concentration measured by terminal clearing system: as found 100% USD, HHI 1.0000, 1.00 system; as designed USD 55 / CNY 15 / EUR 12 / INR 10 / AED 8, HHI 0.3558, 2.810568 systems. **The 55% dollar share is set by revenue currency, not by structure, and reaching 4.00 would require the dollar near 38.99% — a commercial decision, not a legal one.** Four hard limits: a USD account settles under US law wherever booked; CIPS is clearing, not messaging; the bank's risk appetite is not the jurisdiction's policy; a rail is not a holding.

V7 — a fully disclosed sovereign investor: two functions are impossible without disclosure. The existence proof in the disclosure paper is built as a function decomposition, and two of the institution's functions cannot be performed by a concealed holder at all — voting the shares and pre-announcing the vote, and moving a company's conduct by naming it in a published exclusion. When the sector came under political pressure the answer was a published disclosure standard, not more concealment. **And the section carries its own counter-case as failure mode 7: the 2022 immobilisation of approximately USD 300bn involved fully disclosed assets that were frozen anyway. Disclosure is not immunity.** A section arguing for disclosure that hid the strongest objection to disclosure would fail its own test, so the counter-case stays

prominent — here, in the failure record's third predicted failure, and in this synthesis's closing section.

V8 — the structure solves the smaller problem. Child 2's forced-heirship exposure is capped at USD 6.70m and is zero on the designed map. French inheritance-**tax** exposure is approximately USD 27.00m — 4.0299 times larger and untouched by the structure. The liquidity test shows a USD 33.84m shortfall, 14.1% of the estate, that would force the sale of an operating branch. **A succession section reporting only the heirship win would be lying by omission.** The thesis carried forward: **consent is the one lever present in all three legal families — advance renunciation in France, the *Pflichtteilsverzicht* in Germany, *takharuj* under Sharia — and consent requires disclosure.** That ties that group back to the governing principle.

V9 — cost is not the reason. The computed crossover: the foundation is cheaper than a trust only above USD 34.2857m at a 35 bp trustee basis (USD 80.0000m at 15 bp; USD 120.0000m at 10 bp). **So cost is explicitly not the reason to choose it — the reason is recognition in civil-law and Sharia jurisdictions.** That must be said wherever the foundation is recommended, including here.

V10 — English law decomposes into three things. Doctrine, court, and territorial state. ADGM buys the first two without the third. That is the precise statement of what the Gulf bridge delivers and why the City cannot deliver it. What remains open is whether the second holds up under a stress that has not yet occurred: the ADGM Courts have a short record and the onshore enforcement bridge is documented but lightly tested.

E.6 The four admissions that matter most

Extracted from the ledger and stated without qualification, because they are the four things a reader is most likely to leave the framework without having absorbed.

One. The framework's own instrument shows the recommended structure still below the build threshold. At the tested parameters the conventional structure gives `h = \ensuremath{-}0.500000` and the designed structure gives `h = \ensuremath{-}0.122417`. The break-even is `N* = 4.172734` and the designed structure reaches 2.941176. Adding the best available marginal position reaches 3.174603 and closes 21.8357% of the gap. **Diversification raises the ceiling and does not pay the maintenance bill. Jurisdictional diversity buys the ability to build, not the building.**

Two. Disclosure did not prevent the 2022 immobilisation. The most thoroughly disclosed assets available were frozen anyway. Disclosure removes the failure mode in which a structure collapses on contact with an authority, and makes the structure bankable, which is the binding constraint on everything else. It does not supply an appeal and it does not protect an asset inside the perimeter that took the decision.

Three. Cost is not the reason to choose a foundation. Recognition in civil-law and Sharia jurisdictions is. The crossover figures are in V9 and in Appendix B.13.

Four. The succession structure addresses the smaller of the two exposures. USD 6.70m capped and eliminated; approximately USD 27.00m untouched; a USD 33.84m liquidity shortfall sufficient to force a sale. **The asymmetry is the finding.**

E.7 The standing disclaimer

This series is **not legal, tax, or investment advice**. It describes how systems work, what they cost, their risk profiles, and how to build durable legitimate structures. It contains no operational guidance for moving value that a sanctions regime prohibits.

It recommends no structure premised on non-disclosure, because such a structure fails the governing principle by construction. It recommends no nominee arrangement where nominees are unlawful, and Thailand is the specific case where they are criminal and prosecuted. Every worked case in every section and every appendix is **CONSTRUCTED** and labelled

as such. No real family, entity, or transaction is described anywhere in this series, with the single exception of the publicly documented institutions named in the disclosure paper and in the lineage sections, which are treated from the public record.

Counsel is required in each relevant jurisdiction. **The framework is a map, not a permission.**

Appendix F. Correspondence to the earlier series

Series D is the fourth series of the corpus and it is not a topical detour: it applies one law that the earlier series developed on other material. This appendix sets out the correspondence term by term. **A reader who has not read the earlier series loses nothing; the law is restated in full in the erosion paper and worked in full in Appendix B.**

F.1 The storage law, transferred

Series B develops a storage law for durable goods: a difference equation relating a stock, its decay rate, the conversion of flow into stock, and the maintenance cost of holding it [Curzi, 2026], whose founding reframing is that comfort is the felt total of present costs, so the present state is a local minimum of the comfort landscape by definition rather than by luck. **It applies to capital without modification.**

```
Q_{t+1} = (1 − delta) Q_t  +  eta S_t  −  C_t
```

Three terms describe the same physical facts in both cases and one has to be reinterpreted.

F.2 The terms

`Q` — protected capital. In Series B, the stock: a quantity of durable goods held against future need. Here, capital held in a form and a place that survives an adverse decision. **The transfer is exact and requires no reinterpretation, because the storage law does not care what the stock is made of — only that it decays, that it can be added to, and that holding it costs something.**

`delta` — erosion. In Series B: spoilage, breakage, pilferage, obsolescence. Here: tax drag, political risk, currency debasement, and haven degradation. Every jurisdiction paper in this series is, at bottom, an argument about `delta`: a stable legal system with a working court and no expropriation history lowers it, and the fifteen-year erosion documented in the erosion paper is a description of `delta` rising across an entire class of holdings at once. **The mechanism differs and the mathematics does not: a stock decaying at 2% a year is decaying at 2% a year whether the cause is rot or a listing decision.**

`eta` — conversion. In Series B, the efficiency with which flow becomes durable stock: the fraction of a harvest that survives preservation and enters storage. Here, the conversion of income into durable protected holdings. It is the term the jurisdiction papers move when they discuss vehicles rather than stability — a foundation with separate legal personality, a recognised trust, a licensed insurer holding reserves against real liabilities, a REIT unit in place of a title deed. **Each of those raises `eta` by making a given unit of income convertible into a form that holds.**

`C` — maintenance. In Series B, the cost of the store itself: the building, the drying, the turning, the watching. Here: compliance, counsel, substance, filings, registered offices, agents, audits, director fees, translations, and the management time nobody bills. **The correspondence is uncomfortably exact, including the failure mode.** In both cases `C` is paid continuously in small amounts to many parties, never appears as a single line, and is therefore mis-estimated

downward by almost everybody. The AIFC result is the sharpest demonstration: a position worth 0.026731 of `h` is entirely consumed if its leg costs more than 8.9102% of the maintenance budget.

`N` — the effective bloc count. This is the term that had to be reinterpreted, and the reinterpretation is the whole book.

In Series B the diversity term counts independent stores: separate granaries, separate cellars, separate locations, on the reasoning that a fire, a flood, or a theft reaches one and not the others. The count that mattered was the number of things that could fail independently.

In the capital case, the same reasoning gives a different count. **Four entities in four jurisdictions are not four stores if one decision reaches all four.** `N` therefore counts independent coordination perimeters — not entities, not jurisdictions, not flags. That identification is the concentration paper's [Curzi, 2026], whose Result One takes an entity map of UK 40.9 / US 27.3 / Jersey 18.2 / Luxembourg 13.6 from a naive HHI of 0.293430 and 3.41 effective jurisdictions to a bloc-HHI of 1.000000 and an effective count of 1.00, it is the framework's central move, and everything else in the method is downstream of it.

F.3 The sustaining frequency

The steady-state condition, and the form in which the law is used throughout this series:

```
h = eta ln(N) − delta Q − C
```

`h` is the sustaining rate: positive and the stock builds, negative and it is being funded out of something else. In Series B it answers the question of how many independent stores a household needs before storage sustains itself. Here it answers the same question about perimeters.

Three properties transfer without change and are worth naming because each governs a decision in this series.

The diversity term is logarithmic, so returns diminish. The logarithm is not chosen for convenience: it is the surplus functional the corpus characterises uniquely by angle-only dependence, vanishing at parallel, an affine effective count and normalisation [Curzi, 2026], and the affine effective count $N_{\text{eff}} = 1 + (N-1)u$ it carries is badged as a **design axiom rather than a derived result** [Curzi, 2026c] — a hundred-person department at $u = 0.05$ having an effective count of 5.95. `dh/dN = eta/N`. The first move out of a single store is worth far more than the last. At `N = 1.000000` the marginal value is 0.350000; at `N = 2.941176` it is 0.119000. The concentration paper tests the alternative — a linear `eta N` — and finds it produces the absurd recommendation that a structure should hold as many jurisdictions as it can afford, which is exactly the error V5 documents.

The maintenance term is linear and the diversity term is not. Adding a store adds `C` linearly and adds surplus logarithmically, so there is always a count beyond which another store loses money — which is the whole content of the cost ceiling in the break-even probability method. And the break-even itself is computable: `N* = exp((delta Q + C)/eta) = 4.172734`.

F.4 `ln(1.00) = 0`, and what it means in each series

The single most important line of correspondence in this appendix.

```
conventional structure   N = 1.000000    ln(N) = 0.000000
                         h = 0.35 x 0 − 0.5 = −0.500000

designed structure       N = 2.941176    ln(N) = 1.078810
                         h = 0.35 x 1.078810 − 0.5 = −0.122417
```

In Series B, a household with one store has a diversity term of zero: whatever it stores is subject in full to whatever reaches that store, and no rearrangement inside the store changes it. **In this series, a structure with one effective bloc has zero available surplus — not low, zero.** `Q` decays at `delta` with nothing on the input side, and no rearrangement of the four entities in the concentration paper's conventional structure changes it, because rearranging within a perimeter does not change `N`.

That is the sentence the whole of Series D exists to establish, and it is a Series B result applied to a Series D object.

F.5 The surplus law and what it buys

The surplus available from diversity is `eta ln(N)`, and the series's central positive result is the difference between two values of it:

```
from N = 1.000000 to N = 2.941176   surplus gained = 0.377583
from N = 2.941176 to N = 4.000000   surplus gained = 0.107620
```

The first move is worth 3.5087 times the second. A reader with limited capacity to add legs should add the first one and stop, which is the arithmetic behind the AIFC assessment's "small, once, and simply."

And the honest companion result, which belongs in this appendix because it is where the series correspondence stops flattering the series. **`eta ln(N)` at `N = 2.941176` is 0.377583 and `delta Q + C` is 0.500000. The surplus does not cover the maintenance.** In Series B terms: the household built enough independent stores to make storage worth doing and not enough to make it pay for itself, and the remedy is not another store. It is a better conversion rate, a lower decay rate, or a cheaper store — `eta`, `delta`, and `C` — and this series is primarily about none of those.

F.6 What does not transfer

Four things, stated so the correspondence is not over-claimed.

The stock is not homogeneous. A granary holds one commodity; a structure holds operating companies, securities, real estate, intellectual property, and insurance reserves, each with a different `delta` and a different liquidity. **The single `Q` is an aggregation, and the failure record's Scenario C shows what it hides: the property hardest for an adversary to reach was also hardest for the family to reach.**

Rails are not stores. The store–rail rule, and it has no analogue in the earlier series this framework will claim. A transit exposure is an exposure-window problem — hop size, duration, counterparty pre-screening, mid-transit failure handling — and evaluating it with custody tools over-engineers what does not matter. Appendix B.9 reports the rail measurement as a separate number for exactly this reason.

Adversaries adapt; weather does not. The instruments that eroded the havens were designed by people who could see the structures being built against them. The threat register's failure mode 10 warns against modelling the actor as rational; the same warning applies to modelling it as stationary.

The parameters are constructed. `eta = 0.35`, `C = 0.30`, `Q = 10`, `delta = 0.02` are the carried constructed case's test values. **A real structure requires each estimated from its own facts, and `C` is the term most likely to be wrong in the direction that flatters the answer.**

F.7 Where the same law is worked in the other series

The correspondence runs in both directions, and naming the instances is the point of a reference appendix. The surplus functional appears at internal scope as $F_{\text{int}} = \ln(1 + (N-1)u)$

over an audited framing count, with the affine form again badged as a design axiom rather than a theorem [Curzi, 2026o]. It appears in the frequency register as an exact threshold-corrected Fechner law, $F(u) = \ln((u + u_0)/u_0)$ with $u_0 = 1/(N - 1)$, verified for $N = 2$ to 100 at seven values of u to a maximum deviation of $8.9\text{e}{-16}$, and with $N_{\text{eff}}(u_0) = 2$ and $F(u_0) = \ln 2$ for every N [Curzi, 2026]. Its unit conversion is fixed exactly: one nat of required surplus is 2.885390081777927 octaves, 3462.4681 cents, or 1.442695 bits [Curzi, 2026].

Two results from that series bound what a count can be asked to do, and both bear directly on this synthesis. Replacing one component of a set with a paraphrase of another takes the internal surplus from 2.164390 to 1.333662 — a loss of 2.3970 octaves — which is the same event as merging two nominally distinct jurisdictions into one perimeter [Curzi, 2026]. And the opposition identity is the sharper of the two: because $u = 1 - \rho^2$ depends only on $|\rho|$, perfect agreement and perfect opposition give the identical reading [Curzi, 2026]. **A count of independent things cannot distinguish alignment from anti-alignment**, which is exactly why the bloc call has to be argued from evidence rather than read off a correlation.

One bound is carried from the foundations and is not restated here. The corpus refuses category innocence as a defence of acquisition and forbids moving a harm through a discovered gap in a scheme [Curzi, 2026]. Nothing in this method licenses using a measured perimeter gap for that purpose, and the disclosure test of step 3 exists partly to make the attempt visible.

Bibliography

Maria Abascal and Delia Baldassarri. Love thy neighbor? ethnoracial diversity and trust reexamined. *American Journal of Sociology*, 121(3):722–782, 2015. doi: 10.1086/683144.

János Aczél, Bruno Forte, and C. T. Ng. Why the Shannon and Hartley entropies are 'natural'. *Advances in Applied Probability*, 6(1):131–146, 1974. doi: 10.2307/1426210.

M. A. Adelman. Comment on the H concentration measure as a numbers-equivalent. *The Review of Economics and Statistics*, 51(1):99–101, 1969. doi: 10.2307/1926955.

Réka Albert, Hawoong Jeong, and Albert-László Barabási. Error and attack tolerance of complex networks. *Nature*, 406(6794):378–382, 2000. doi: 10.1038/35019019.

Christopher Alexander. A city is not a tree. *Architectural Forum*, 122(1–2), 1965.

Christopher Alexander, Sara Ishikawa, and Murray Silverstein. *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press, New York, 1977. ISBN 978-0-19-501919-3.

Gordon W. Allport. *The Nature of Prejudice*. Addison-Wesley, Reading, MA, 1954.

Philip W. Anderson. More is different. *Science*, 177(4047):393–396, 1972. doi: 10.1126/science.177.4047.393.

David Andrich. *Rasch Models for Measurement*. Sage Publications, Newbury Park, CA, 1988. ISBN 978-0-8039-2741-4.

Andrew Ang and Joseph Chen. Asymmetric correlations of equity portfolios. *Journal of Financial Economics*, 63(3):443–494, 2002. doi: 10.1016/S0304-405X(02)00068-5.

W. Ross Ashby. *An Introduction to Cybernetics*. Chapman and Hall, London, 1956.

W. Ross Ashby. Requisite variety and its implications for the control of complex systems. *Cybernetica*, 1(2):83–99, 1958.

Karl Johan Åström and Richard M. Murray. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press, Princeton, NJ, 2008. ISBN 978-0-691-13576-2.

Robert Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.

Robert Axelrod. The dissemination of culture: A model with local convergence and global polarization. *Journal of Conflict Resolution*, 41(2):203–226, 1997. doi: 10.1177/0022002797041002001.

Robert Axelrod and William D. Hamilton. The evolution of cooperation. *Science*, 211(4489):1390–1396, 1981. doi: 10.1126/science.7466396.

Philip Ball. *Critical Mass: How One Thing Leads to Another*. Farrar, Straus and Giroux, New York, 2004.

Stefan Banach. Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales. *Fundamenta Mathematicae*, 3:133–181, 1922. doi: 10.4064/fm-3-1-133-181.

Abhijit V. Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817, 1992. doi: 10.2307/2118364.

Albert-László Barabási. The origin of bursts and heavy tails in human dynamics. *Nature*, 435 (7039):207–211, 2005. doi: 10.1038/nature03459.

Albert-László Barabási. *Network Science*. Cambridge University Press, Cambridge, 2016. ISBN 978-1-107-07626-6.

Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999. doi: 10.1126/science.286.5439.509.

J. Murray Barbour. *Tuning and Temperament: A Historical Survey*. Michigan State College Press, East Lansing, MI, 1951.

Julia B. Bear and Anita Williams Woolley. The role of gender in team collaboration and performance. *Interdisciplinary Science Reviews*, 36(2):146–153, 2011. doi: 10.1179/030801811X13013181961473.

Dave Benson. *Music: A Mathematical Offering*. Cambridge University Press, Cambridge, 2006. ISBN 978-0-521-61999-8.

Leo Beranek. *Concert Halls and Opera Houses: Music, Acoustics, and Architecture*. Springer, New York, 2 edition, 2004. doi: 10.1007/978-0-387-21636-2.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, 1992. doi: 10.1086/261849.

Jacques Philippe Marie Binet. Mémoire sur l'intégration des équations linéaires aux différences finies d'un ordre quelconque, à coefficients variables. *Comptes Rendus Hebdomadaires des Séances de l'Académie des Sciences*, 17:559–567, 1843.

Gary B. Born. *International Commercial Arbitration*. Kluwer Law International, Alphen aan den Rijn, 3 edition, 2021.

Jean-Philippe Bouchaud. Crises and collective socio-economic phenomena: Simple models and challenges. *Journal of Statistical Physics*, 151(3–4):567–606, 2013. doi: 10.1007/s10955-012-0687-3.

Jean-Philippe Bouchaud and Marc Mézard. Wealth condensation in a simple model of economy. *Physica A: Statistical Mechanics and its Applications*, 282(3–4):536–545, 2000. doi: 10.1016/S0378-4371(00)00205-3.

Louis D. Brandeis. *Other People's Money and How the Bankers Use It*. Frederick A. Stokes, New York, 1914.

Philip Brickman, Dan Coates, and Ronnie Janoff-Bulman. Lottery winners and accident victims: Is happiness relative? *Journal of Personality and Social Psychology*, 36(8):917–927, 1978. doi: 10.1037/0022-3514.36.8.917.

Adrian Briggs. *The Conflict of Laws*. Oxford University Press, Oxford, 3 edition, 2013. ISBN 978-0-19-967927-2.

Leon Brillouin. The negentropy principle of information. *Journal of Applied Physics*, 24(9): 1152–1163, 1953. doi: 10.1063/1.1721463.

Bob Briscoe, Andrew Odlyzko, and Benjamin Tilly. Metcalfe's law is wrong. *IEEE Spectrum*, 43(7):34–39, 2006. doi: 10.1109/MSPEC.2006.1653003.

William A. Brock and Steven N. Durlauf. Discrete choice with social interactions. *The Review of Economic Studies*, 68(2):235–260, 2001. doi: 10.1111/1467-937X.00168.

Anna D. Broido and Aaron Clauset. Scale-free networks are rare. *Nature Communications*, 10 (1):1017, 2019. doi: 10.1038/s41467-019-08746-5.

Frederick P. Brooks. *The Mythical Man-Month: Essays on Software Engineering*. Addison-Wesley, Reading, MA, 1975.

Robert G. Brown and Richard F. Meyer. The fundamental theorem of exponential smoothing. *Operations Research*, 9(5):673–685, 1961. doi: 10.1287/opre.9.5.673.

Elizabeth Bruch and Jon Atwell. Agent-based models in empirical social research. *Sociological Methods and Research*, 44(2):186–221, 2015. doi: 10.1177/0049124113506405.

Sergey V. Buldyrev, Roni Parshani, Gerald Paul, H. Eugene Stanley, and Shlomo Havlin. Catastrophic cascade of failures in interdependent networks. *Nature*, 464(7291):1025–1028, 2010. doi: 10.1038/nature08932.

Edward M. Burns. Intervals, scales, and tuning. In Diana Deutsch, editor, *The Psychology of Music*, pages 215–264. Academic Press, San Diego, CA, 2 edition, 1999. doi: 10.1016/B978-012213564-4/50008-1.

Edward M. Burns and W. Dixon Ward. Categorical perception — phenomenon or epiphenomenon: Evidence from experiments in the perception of melodic musical intervals. *The Journal of the Acoustical Society of America*, 63(2):456–468, 1978. doi: 10.1121/1.381737.

Ronald S. Burt. *Structural Holes: The Social Structure of Competition*. Harvard University Press, Cambridge, MA, 1992. ISBN 978-0-674-84371-4.

Ronald S. Burt. Structural holes and good ideas. *American Journal of Sociology*, 110(2):349–399, 2004. doi: 10.1086/421787.

Duncan S. Callaway, M. E. J. Newman, Steven H. Strogatz, and Duncan J. Watts. Network robustness and fragility: Percolation on random graphs. *Physical Review Letters*, 85(25): 5468–5471, 2000. doi: 10.1103/PhysRevLett.85.5468.

Herbert B. Callen and Theodore A. Welton. Irreversibility and generalized noise. *Physical Review*, 83(1):34–40, 1951. doi: 10.1103/PhysRev.83.34.

Michel Callon, editor. *The Laws of the Markets*. Blackwell Publishers / The Sociological Review, Oxford, 1998.

Donald T. Campbell. Assessing the impact of planned social change. *Evaluation and Program Planning*, 2(1):67–90, 1979. doi: 10.1016/0149-7189(79)90048-X.

Donald T. Campbell and Donald W. Fiske. Convergent and discriminant validation by the multitrait-multimethod matrix. *Psychological Bulletin*, 56(2):81–105, 1959. doi: 10.1037/h0046016.

Nancy Cartwright. *The Dappled World: A Study of the Boundaries of Science*. Cambridge University Press, Cambridge, 1999. doi: 10.1017/CBO9781139167093.

Hendrik B. G. Casimir. On the attraction between two perfectly conducting plates. *Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen*, 51:793–795, 1948.

Claudio Castellano, Matteo Marsili, and Alessandro Vespignani. Nonequilibrium phase transition in a model for social influence. *Physical Review Letters*, 85(16):3536–3539, 2000. doi: 10.1103/PhysRevLett.85.3536.

Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Reviews of Modern Physics*, 81(2):591–646, 2009. doi: 10.1103/RevModPhys.81.591.

Damon Centola. The spread of behavior in an online social network experiment. *Science*, 329 (5996):1194–1197, 2010. doi: 10.1126/science.1185231.

Damon Centola and Michael Macy. Complex contagions and the weakness of long ties. *American Journal of Sociology*, 113(3):702–734, 2007. doi: 10.1086/521848.

Damon Centola, Joshua Becker, Devon Brackbill, and Andrea Baronchelli. Experimental evidence for tipping points in social convention. *Science*, 360(6393):1116–1119, 2018. doi: 10.1126/science.aas8827.

Anne Chao and Lou Jost. Coverage-based rarefaction and extrapolation: Standardizing samples by completeness rather than size. *Ecology*, 93(12):2533–2547, 2012. doi: 10.1890/11-1952.1.

Anne Chao, Nicholas J. Gotelli, T. C. Hsieh, Elizabeth L. Sander, K. H. Ma, Robert K. Colwell, and Aaron M. Ellison. Rarefaction and extrapolation with Hill numbers: A framework for sampling and estimation in species diversity studies. *Ecological Monographs*, 84(1):45–67, 2014. doi: 10.1890/13-0133.1.

Aaron Clauset, Cosma Rohilla Shalizi, and M. E. J. Newman. Power-law distributions in empirical data. *SIAM Review*, 51(4):661–703, 2009. doi: 10.1137/070710111.

Reuven Cohen, Keren Erez, Daniel ben Avraham, and Shlomo Havlin. Resilience of the Internet to random breakdowns. *Physical Review Letters*, 85(21):4626–4628, 2000. doi: 10.1103/PhysRevLett.85.4626.

Melvin E. Conway. How do committees invent? *Datamation*, 14(5):28–31, 1968.

Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, Hoboken, NJ, 2 edition, 2006. ISBN 978-0-471-24195-9.

Michael Cox, Gwen Arnold, and Sergio Villamayor Tomás. A review of design principles for community-based natural resource management. *Ecology and Society*, 15(4):38, 2010. doi: 10.5751/ES-03704-150438.

Lee J. Cronbach. Coefficient alpha and the internal structure of tests. *Psychometrika*, 16(3): 297–334, 1951. doi: 10.1007/BF02310555.

Lee J. Cronbach and Paul E. Meehl. Construct validity in psychological tests. *Psychological Bulletin*, 52(4):281–302, 1955. doi: 10.1037/h0040957.

Michael Laurence Curzi. The Storage Law as a Leaky Integrator: Decay, Poles, and Why Zero Loss Destroys Arrival. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-007, 36N9 Genetics LLC, 2026a. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Conversion, Coordination Cost, and Headroom: Why Eta Is the First Lever and Why the Optimum Group Is Smaller Than the Crossover. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-008, 36N9 Genetics LLC, 2026b. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Effective Count and the Logarithmic Ceiling: Auditing Roster Size Against Non-Parallelism. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-009, 36N9 Genetics LLC, 2026c. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Crossing Versus Divergence: Decomposing Non-Parallelism, the Decay of u, and the Room as a Circuit Element. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-010, 36N9 Genetics LLC, 2026d. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Passive Geometry and the Reverberant Tail: Crenellation, Parabolic Focus, and the Acoustic Exclusion of Quiet Participants. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-011, 36N9 Genetics LLC, 2026e. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Thresholds, Substrate, and the Table With No Head: Three Built-Environment Levers Priced Against the Storage Law. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-012, 36N9 Genetics LLC, 2026f. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Turn-Taking, Complex Contagion, and the Critical Boundary-Cost Multiplier: Three Channels into the Surplus Law. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-013, 36N9 Genetics LLC, 2026g. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Order Without a Centre: Stigmergic Media, Independence Protocols, and the Institutional Maintenance Term. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-014, 36N9 Genetics LLC, 2026h. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Surveying and Preregistering the Storage Law: A Measurement Protocol, and a Confirmatory Fit That Failed Its Own Baseline Test. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-015, 36N9 Genetics LLC, 2026i. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. A Gated Intervention Procedure and a Thirty-Two-Entry Failure Catalog for the Storage Law. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-016, 36N9 Genetics LLC, 2026j. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Five Worked Systems and the Essence Bridge: Testing Scale-Freedom in the Storage Law and Identifying the Effective Count as a Pitch Index. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-017, 36N9 Genetics LLC, 2026k. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Conversion Keys and Modal Design: Deriving a Room, a Cadence and a Pitch Lattice from an Essence Set Without Subliminal Signal. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-018, 36N9 Genetics LLC, 2026l. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Moment Equation and the Calibration Ratio: A Fixed Point at Contact, a Self-Imposed Demotion to Ordinal Scale, and the Price of Stored Error. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-019, 36N9 Genetics LLC, 2026m. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Perceptual Lens: Fechner's Law as the Register of the Surplus Functional, with a Postulate Method and Its Held-Out Tests. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-020, 36N9 Genetics LLC, 2026n. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Internal Surplus and the Self as Storage: Auditing One Observer's Framings, and the Stored-Error Track That Feels Identical From Inside. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-021, 36N9 Genetics LLC, 2026o. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Launch Budget and the Desire Line: A Rectified Surplus Fund Against a Barrier, and Why Momentum Does Not Cross a Wide Shallow Ridge. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-022, 36N9 Genetics LLC, 2026p. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Return Difference, the Contraction, and a Formulary That Demotes Five of Its Own Nine Entries. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-023, 36N9 Genetics LLC, 2026q. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Dyad and the Group Moment: Opposition Is Parallel, Calibration Costs All the Headroom, and Distortion Does Not Average. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-024, 36N9 Genetics LLC, 2026r. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Trust, Reciprocity and Obligation as Quantities: Where Trust Enters the Storage Law, and Nine Ways the Internal/External Correspondence Breaks. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-025, 36N9 Genetics LLC, 2026s. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Formula Validation Protocol, Scale Type, and Specific Objectivity: Demoting the Corpus's Own Moment Equation and Pricing Its Restoration. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-026, 36N9 Genetics LLC, 2026t. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Falsification Suite Measured Against Itself: An Anti-Conservative Shuffle, a 52.7 Per Cent Demand-Characteristic Pass Rate, and a Programme With a 35.8 Per Cent Survival Probability. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-027, 36N9 Genetics LLC, 2026u. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Star Polygons, Cadence, and the Ledger: gcd(N,k) as Parallelism, Division of a Fixed Budget as a Theorem, and What an Eleven-Operation Audit Does Not Prove. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-028, 36N9 Genetics LLC, 2026v. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Two Columns and the Concordance Register: Why Agreeing Traditions Stop Adding Evidence After the Second, and a Register That Rejects One Bridge in Five. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-029, 36N9 Genetics LLC, 2026w. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Cell as Receiver, the Channel and Its Capacity, and Addressing Without Editing: Adequate Bandwidth Into an Absent Decoder. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-030, 36N9 Genetics LLC, 2026x. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Addressing the Non-Coding Genome: The Latency Constraint, the Missing Demodulator, and an Exact Sonification Pipeline. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-031, 36N9 Genetics LLC, 2026y. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Conjugate Uncertainty at Three Levels: The Kennard Bound, the Gabor Limit, and a Badged Information-Budget Analogue for Forecasts. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-032, 36N9 Genetics LLC, 2026z. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Quantum Formalism Without Licence: Born-Rule Statistics, Decoherence Timescales in Warm Tissue, and the Casimir Force at Biological Scale. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-033, 36N9 Genetics LLC, 2026aa. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Belief Creates Reality If the Belief Is Grounded in Truth: A Conditional Postulate, Its Channel-Gain Form, and the Placebo Case. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-034, 36N9 Genetics LLC, 2026ab. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Where Belief Fails: Four Failure Classes, a Priced Bank Run, and the Grounding Protocol That Kills Most Claims. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-035, 36N9 Genetics LLC, 2026ac. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Narrative as Regulatory Input: A Fully Disclosed Sonification Chain, a Zero-Resultant Corpus, and a Preregistration That Prices Its Own Nulls. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-036, 36N9 Genetics LLC, 2026ad. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Protocols at Two Scopes and a Finite Failure Catalogue: Headroom Gating, the Participation-Ratio Effective Count, and Thirty Manifestation Faults. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-037, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. One Model Form at Three Scopes: The Instrument Audit, Five Breaks in the Correspondence, and a 156-Entry Status Register. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-038, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Four Structural Results and an Honesty Register: The Closing Synthesis and Reference Apparatus of Engineering the Surplus. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-039, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Sustaining Frequency: $f_{\min} = k\,e^{2c/\eta}$ as an Equivalence Between the Storage Fixed Point and the Essence Bridge. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-040, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Minimum Sustainable Frequency: A Pitch Criterion for Relational Sustainability and the Count-Based Essence Index That Feeds It. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-041, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Octave Folding and the Odd-Part Class Rule: The Frequency of a Mood, of a Trait, and the Self as a Chord. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-042, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Consonance, Collision, and the Divided Self: A Firewall Between Harmonic Dissonance and Psychological Distress. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-043, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Three Routes to the Relationship Frequency: The Unison Result and the Elasticity-Reserve Split Between Levers. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-044, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Two Octaves per Doubling Against a Quadratic Cost: Roster Ceilings, the Parallel Wall, and When Tuning Must Halt. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-045, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Auditing a Traditional Operation Set Against a Parameter Model: Seven Maps, Two Splits, Three Rejections, and the Many-to-One Compound Rule. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-046, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Refinement as Contraction: Why Repetition Cannot Move a Fixed Point, and an Evidential Audit of the Tradition It Comes From. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-047, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. A Disclosed Sonification Protocol for Sequence-Encoded Content, and the Missing Transducer That Makes It a Research Design and Not a Technique. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-048, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Destiny as a Fixed Point: Conservation of Drive Under Cost Transfer, the Completeness of the Lever Set, and the Corpus's Register of Demotions. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-049, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Reference Apparatus for the Frequency Framework: Constants, Lattice Tables, Measurement Protocols, and a Four-Drawer Classification of Evidential Status. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-050, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Correlated Jurisdictions: Coordination Perimeters, Erosion, and the Economics of Full Disclosure. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-051, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Concentration Instrument: Bloc Collapse and the Effective Count of Jurisdictions. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-052, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Apex Vehicle: Foundation Against Trust, and the Arithmetic of Forced Heirship. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-053, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Institutional Quality Without Independence: London, the Crown Dependencies, and the European Convergence. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-054, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Entity Maps Understate Dependency: The Dollar Axis, the Insurance Reserve, and a Scored Captive Rubric. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-055, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Substance Repricing and the Clearing-System Count: The Atlantic Cluster, ADGM, and the Limits of a Second Enclave. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-056, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Diversification That Does Not Pay: Gulf Additions, an Indifferent Bridging Slot, and a Capped Transit Position. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-057, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Pricing an Uncertain Bloc Call: AIFC Option Value, Thailand's Share of the Gain, and the GIFT City Corridor. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-058, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Currency Drag, Cost Ceilings, and Funding-Source Independence. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-059, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Compartments, Conversions, and the Confidence a Cell Wall Requires. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-060, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Measuring the Payment Rail: Clearing Perimeters, Transit Exposure Windows, and the Concentration Effect of De-Risking. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-061, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Carve-Out Cascade: Pledge Triggers, Antitrust Walls, and Legal-Family Interfaces Priced as Cost. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-062, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Assembled Architecture: Effective Blocs, Sanctions Deeming Rules, and the Disclosure Test in Historical Instance. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-063, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Five Houses, One Perimeter: Bloc Collapse in a Historical Network, the Completeness Gap of Double Entry, and Canonicalisation as an Integrity Control. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-064, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Canonical Form and Typed Parsing: Measured Integrity Controls Across a Financial Identifier Namespace. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-065, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Channel Separation as a Structural Property: In-Band Impersonation, Address Anonymity Sets, and the Legacy-Migration Asymptote. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-066, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. A Concentration Atlas for Distributed Ledgers: Validator Herfindahls, the Utility Ordinal, and Dwell Time as the Freeze Variable. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-067, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Content Addressing, Time-Axis Concentration, and the Bounded Gain of a Nine-Capital Supplement. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-068, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. What the Architecture Cannot Do: Correlated Backing Baskets, Gated Adoption Probabilities, and the Failure Record. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-069, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Method in Seven Steps: A Consolidated Reproduction Workbook and Honesty Register for the Concentration and Headroom Instrument. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-070, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Comfort as a Local Minimum: The Founding Definitions of the Creature Comforts Framework and the Ground-State Node. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-A01, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Ledger, the Gradient, and the Arithmetic of Difference: Carrying Cost, Desire Lines, and the Surplus Functional f(u) = ln(1 + (N-1)u). Psychological and Sociological Physics: The Creature Comforts Corpus CCP-A02, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Weighting Futures by Resistance: The Comfort-Weighted Amplitude, Cue Placement, and the Resolution-Disturbance Trade-Off. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-A03, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Spoken Word and the Returned Difference: Performative Commitment, Echo Reading, and the Doctrine of Category Violation. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-A04, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. Convergence and Its Dark Twin: The Contraction Mapping, the Audit of Arrival, and the Consent Law of Closure. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-A05, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Michael Laurence Curzi. The Reference Apparatus of the Foundations Volume: The Closing Rite, the Protocol Index, and the Three-Drawer Honesty Ledger. Psychological and Sociological Physics: The Creature Comforts Corpus CCP-A06, 36N9 Genetics LLC, 2026. DOI: 10.5281/zenodo.22079583.

Richard M. Cyert and James G. March. *A Behavioral Theory of the Firm*. Prentice-Hall, Englewood Cliffs, NJ, 1963.

Marieke de Goede. Blacklisting and the ban: Contesting targeted sanctions in Europe. *Security Dialogue*, 42(6):499–515, 2011. doi: 10.1177/0967010611425368.

Raymond de Roover. *The Rise and Decline of the Medici Bank, 1397–1494*. Harvard University Press, Cambridge, MA, 1963.

Guillaume Deffuant, David Neau, Frederic Amblard, and Gérard Weisbuch. Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3(1–4):87–98, 2000. doi: 10.1142/S0219525900000078.

Morris H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69 (345):118–121, 1974. doi: 10.1080/01621459.1974.10480137.

Laurent Demany, Catherine Semal, and Vincent Péan. Auditory discrimination of frequency ratios: The octave singularity. *Journal of Experimental Psychology: Human Perception and Performance*, 39(3):788–801, 2013. doi: 10.1037/a0030095.

Diana Deutsch. The processing of pitch combinations. In Diana Deutsch, editor, *The Psychology of Music*, pages 349–411. Academic Press, San Diego, CA, 2 edition, 1999. doi: 10.1016/B978-012213564-4/50011-1.

Ed Diener, Richard E. Lucas, and Christie Napa Scollon. Beyond the hedonic treadmill: Revising the adaptation theory of well-being. *American Psychologist*, 61(4):305–314, 2006. doi: 10.1037/0003-066X.61.4.305.

Martin Dixon. *Textbook on International Law*. Oxford University Press, Oxford, 7 edition, 2013. ISBN 978-0-19-957445-4.

Daniel F. Doak, Daniel Bigger, E. K. Harding, M. A. Marvier, R. E. O'Malley, and D. Thomson. The statistical inevitability of stability-diversity relationships in community ecology. *The American Naturalist*, 151(3):264–276, 1998. doi: 10.1086/286117.

David B. Doty. *The Just Intonation Primer*. Just Intonation Network, San Francisco, CA, 3 edition, 2002.

John R. Douceur. The Sybil attack. In *Peer-to-Peer Systems: First International Workshop, IPTPS 2002*, pages 251–260, Berlin, 2002. Springer. doi: 10.1007/3-540-45748-8_24.

John C. Doyle, David L. Alderson, Lun Li, Steven Low, Matthew Roughan, Stanislav Shalunov, Reiko Tanaka, and Walter Willinger. The "robust yet fragile" nature of the Internet. *Proceedings of the National Academy of Sciences*, 102(41):14497–14502, 2005. doi: 10.1073/pnas.0501426102.

Patrick Dunleavy and Françoise Boucek. Constructing the number of parties. *Party Politics*, 9 (3):291–315, 2003. doi: 10.1177/1354068803009003002.

Steven N. Durlauf. Complexity and empirical economics. *The Economic Journal*, 115(504): F225–F243, 2005. doi: 10.1111/j.1468-0297.2005.01003.x.

Saber Elaydi. *An Introduction to Difference Equations*. Springer, New York, 3 edition, 2005. doi: 10.1007/0-387-27602-5.

Edwin J. Elton and Martin J. Gruber. Risk reduction and portfolio size: An analytical solution. *The Journal of Business*, 50(4):415–437, 1977. doi: 10.1086/295964.

David Engel, Anita Williams Woolley, Lisa X. Jing, Christopher F. Chabris, and Thomas W. Malone. Reading the mind in the eyes or reading between the lines? theory of mind predicts collective intelligence equally well online and face-to-face. *PLOS ONE*, 9(12):e115212, 2014. doi: 10.1371/journal.pone.0115212.

Joshua M. Epstein. Why model? *Journal of Artificial Societies and Social Simulation*, 11(4):12, 2008.

Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5:17–61, 1960.

Agner Krarup Erlang. The theory of probabilities and telephone conversations. *Nyt Tidsskrift for Matematik B*, 20:33–39, 1909.

Euclid. *The Thirteen Books of Euclid's Elements, Volume 2*. Dover Publications, New York, 2 edition, 1956.

Leonhard Euler. *Tentamen novae theoriae musicae ex certissimis harmoniae principiis dilucide expositae*. Ex typographia Academiae scientiarum, Petropoli (St Petersburg), 1739.

European Union. Regulation (EU) no 650/2012 of the European Parliament and of the Council of 4 July 2012 on jurisdiction, applicable law, recognition and enforcement of decisions and acceptance and enforcement of authentic instruments in matters of succession. Official Journal of the European Union L 201/107, 2012.

J. Doyne Farmer and Duncan Foley. The economy needs agent-based modelling. *Nature*, 460 (7256):685–686, 2009. doi: 10.1038/460685a.

James Fawcett and Janeen M. Carruthers. *Cheshire, North and Fawcett: Private International Law*. Oxford University Press, Oxford, 14 edition, 2008. ISBN 978-0-19-928438-2.

Gustav Theodor Fechner. *Elemente der Psychophysik*. Breitkopf und Härtel, Leipzig, 1860.

Gustav Theodor Fechner. *Elements of Psychophysics, Volume I*. Holt, Rinehart and Winston, New York, 1966.

Niall Ferguson. *The House of Rothschild: Money's Prophets, 1798–1848*. Viking, New York, 1998. ISBN 978-0-670-85768-9.

Fabrizio Ferraro, Jeffrey Pfeffer, and Robert I. Sutton. Economics language and assumptions: How theories can become self-fulfilling. *Academy of Management Review*, 30(1):8–24, 2005. doi: 10.5465/amr.2005.15281412.

Financial Action Task Force. International standards on combating money laundering and the financing of terrorism and proliferation: The FATF recommendations, 2012.

Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, 1985. doi: 10.1145/3149.214121.

Roger Fischler. How to find the "golden number" without really trying. *The Fibonacci Quarterly*, 19(5):406–410, 1981.

Jay W. Forrester. *Industrial Dynamics*. MIT Press, Cambridge, MA, 1961.

Noah E. Friedkin and Eugene C. Johnsen. Social influence and opinions. *The Journal of Mathematical Sociology*, 15(3–4):193–206, 1990. doi: 10.1080/0022250X.1990.9990069.

Nico H. Frijda. *The Emotions*. Studies in Emotion and Social Interaction. Cambridge University Press, Cambridge, 1986. ISBN 978-0-521-31600-2.

Francis Fukuyama. Social capital, civic society and development. *Third World Quarterly*, 22 (1):7–20, 2001. doi: 10.1080/713701144.

Dennis Gabor. Theory of communication. *Journal of the Institution of Electrical Engineers*, 93 (26):429–457, 1946.

Serge Galam. *Sociophysics: A Physicist's Modeling of Psycho-Political Phenomena*. Understanding Complex Systems. Springer, New York, 2012. doi: 10.1007/978-1-4614-2032-3.

Serge Galam, Yuval Gefen, and Yonathan Shapir. Sociophysics: A new approach of sociological collective behaviour. I. mean-behaviour description of a strike. *The Journal of Mathematical Sociology*, 9(1):1–13, 1982. doi: 10.1080/0022250X.1982.9989929.

Mauro Gallegati, Steve Keen, Thomas Lux, and Paul Ormerod. Worrying trends in econophysics. *Physica A: Statistical Mechanics and its Applications*, 370(1):1–6, 2006. doi: 10.1016/j.physa.2006.04.029.

Francis Galton. Vox populi. *Nature*, 75(1949):450–451, 1907. doi: 10.1038/075450a0.

Diego Gambetta, editor. *Trust: Making and Breaking Cooperative Relations*. Basil Blackwell, Oxford, 1988.

Mark R. Gardner and W. Ross Ashby. Connectance of large dynamic (cybernetic) systems: Critical values for stability. *Nature*, 228(5273):784, 1970. doi: 10.1038/228784a0.

Jan Gehl. *Life Between Buildings: Using Public Space*. Arkitektens Forlag, Copenhagen, 1971.

Andrew Gelman and Eric Loken. The garden of forking paths: Why multiple comparisons can be a problem, even when there is no "fishing expedition" or "p-hacking" and the research hypothesis was posited ahead of time. Working paper, Department of Statistics, Columbia University, 2013.

Wulfram Gerstner and Werner M. Kistler. *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge University Press, Cambridge, 2002. doi: 10.1017/CBO9780511815706.

Seth Gilbert and Nancy Lynch. Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services. *ACM SIGACT News*, 33(2):51–59, 2002. doi: 10.1145/564585.564601.

Samuel Goldberg. *Introduction to Difference Equations*. John Wiley and Sons, New York, 1958.

Eliyahu M. Goldratt and Jeff Cox. *The Goal: A Process of Ongoing Improvement*. North River Press, Great Barrington, MA, 1984.

Grigorii V. Golosov. The effective number of parties: A new approach. *Party Politics*, 16(2): 171–192, 2010. doi: 10.1177/1354068809339538.

Marta C. González, César A. Hidalgo, and Albert-László Barabási. Understanding individual human mobility patterns. *Nature*, 453(7196):779–782, 2008. doi: 10.1038/nature06958.

Charles A. E. Goodhart. Problems of monetary management: The UK experience. *Papers in Monetary Economics*, 1, 1975.

Root Gorelick. Commentary: Do we have a consistent terminology for species diversity? the fallacy of true diversity. *Oecologia*, 167(4):885–888, 2011. doi: 10.1007/s00442-011-2124-8.

Mark Granovetter. Threshold models of collective behavior. *American Journal of Sociology*, 83(6):1420–1443, 1978. doi: 10.1086/226707.

Mark S. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78(6):1360–1380, 1973. doi: 10.1086/225469.

Christopher D. Green. All that glitters: A review of psychological research on the aesthetics of the golden section. *Perception*, 24(8):937–968, 1995. doi: 10.1068/p240937.

Patrick Grim, Daniel J. Singer, Aaron Bramson, Bennett Holman, Sean McGeehan, and William J. Berger. Diversity, ability, and expertise in epistemic communities. *Philosophy of Science*, 86(1):98–123, 2019. doi: 10.1086/701070.

Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, pages 212–219, 1996. doi: 10.1145/237814.237866.

Ian Hacking. The looping effects of human kinds. In Dan Sperber, David Premack, and Ann James Premack, editors, *Causal Cognition: A Multidisciplinary Debate*, pages 351–383. Clarendon Press, Oxford, 1995. doi: 10.1093/acprof:oso/9780198524021.003.0012.

Hague Conference on Private International Law. Convention of 1 July 1985 on the law applicable to trusts and on their recognition. HCCH Convention No. 30, 1985.

Hague Conference on Private International Law. Convention of 1 August 1989 on the law applicable to succession to the estates of deceased persons. HCCH Convention No. 32, 1989.

Hague Conference on Private International Law. Convention of 30 June 2005 on choice of court agreements. HCCH Convention No. 37, 2005.

Hague Conference on Private International Law. Convention of 2 July 2019 on the recognition and enforcement of foreign judgments in civil or commercial matters. HCCH Convention No. 41, 2019.

Edward T. Hall. *The Hidden Dimension*. Doubleday, Garden City, NY, 1966.

Henry Hansmann and Reinier Kraakman. The essential role of organizational law. *The Yale Law Journal*, 110(3):387–440, 2000. doi: 10.2307/797521.

Garrett Hardin. The tragedy of the commons. *Science*, 162(3859):1243–1248, 1968. doi: 10.1126/science.162.3859.1243.

David A. Harrison and Katherine J. Klein. What's the difference? diversity constructs as separation, variety, or disparity in organizations. *Academy of Management Review*, 32(4):1199–1228, 2007. doi: 10.5465/amr.2007.26586096.

Ralph V. L. Hartley. Transmission of information. *The Bell System Technical Journal*, 7(3): 535–563, 1928. doi: 10.1002/j.1538-7305.1928.tb01236.x.

David Hayton, Paul Matthews, and Charles Mitchell. *Underhill and Hayton: Law Relating to Trusts and Trustees*. LexisNexis Butterworths, London, 19 edition, 2016.

Peter Hedström and Petri Ylikoski. Causal mechanisms in the social sciences. *Annual Review of Sociology*, 36(1):49–67, 2010. doi: 10.1146/annurev.soc.012809.102632.

Rainer Hegselmann and Ulrich Krause. Opinion dynamics and bounded confidence: Models, analysis and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3):2, 2002.

Dirk Helbing. Traffic and related self-driven many-particle systems. *Reviews of Modern Physics*, 73(4):1067–1141, 2001. doi: 10.1103/RevModPhys.73.1067.

Dirk Helbing and Péter Molnár. Social force model for pedestrian dynamics. *Physical Review E*, 51(5):4282–4286, 1995. doi: 10.1103/PhysRevE.51.4282.

Dirk Helbing, Illés Farkas, and Tamás Vicsek. Simulating dynamical features of escape panic. *Nature*, 407(6803):487–490, 2000. doi: 10.1038/35035023.

Harry Helson. Adaptation-level as a basis for a quantitative theory of frames of reference. *Psychological Review*, 55(6):297–313, 1948. doi: 10.1037/h0056721.

Orris C. Herfindahl. *Concentration in the U.S. Steel Industry*. PhD dissertation, Columbia University, 1950.

M. O. Hill. Diversity and evenness: A unifying notation and its consequences. *Ecology*, 54(2): 427–432, 1973. doi: 10.2307/1934352.

Albert O. Hirschman. *National Power and the Structure of Foreign Trade*. University of California Press, Berkeley, CA, 1945.

Albert O. Hirschman. The paternity of an index. *The American Economic Review*, 54(5):761–762, 1964.

Albert O. Hirschman. *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*. Harvard University Press, Cambridge, MA, 1970.

Marisa Hoeschele, Ronald G. Weisman, and Christopher B. Sturdy. Pitch chroma discrimination, generalization, and transfer tests of octave equivalence in humans. *Attention, Perception, and Psychophysics*, 74(8):1742–1760, 2012. doi: 10.3758/s13414-012-0364-2.

Jake M. Hofman, Amit Sharma, and Duncan J. Watts. Prediction and explanation in social systems. *Science*, 355(6324):486–488, 2017. doi: 10.1126/science.aal3856.

Holger Höge. The golden section hypothesis — its last funeral. *Empirical Studies of the Arts*, 15(2):233–255, 1997. doi: 10.2190/2pnh-8tt0-emc5-ftw5.

C. S. Holling. Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 4:1–23, 1973. doi: 10.1146/annurev.es.04.110173.000245.

Erik Hollnagel, David D. Woods, and Nancy Leveson, editors. *Resilience Engineering: Concepts and Precepts*. Ashgate, Aldershot, 2006. ISBN 978-0-7546-4641-9.

Charles C. Holt. Forecasting seasonals and trends by exponentially weighted moving averages. ONR Research Memorandum 52, Carnegie Institute of Technology, Pittsburgh, PA, 1957.

Lu Hong and Scott E. Page. Problem solving by heterogeneous agents. *Journal of Economic Theory*, 97(1):123–163, 2001. doi: 10.1006/jeth.2000.2709.

Lu Hong and Scott E. Page. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences*, 101(46): 16385–16389, 2004. doi: 10.1073/pnas.0403723101.

Wallace J. Hopp and Mark L. Spearman. *Factory Physics*. Waveland Press, Long Grove, IL, 3 edition, 2011. ISBN 978-1-57766-739-1.

Paul Horowitz and Winfield Hill. *The Art of Electronics*. Cambridge University Press, Cambridge, 3 edition, 2015. ISBN 978-0-521-80926-9.

Stuart H. Hurlbert. The nonconcept of species diversity: A critique and alternative parameters. *Ecology*, 52(4):577–586, 1971. doi: 10.2307/1934145.

Rob J. Hyndman, Anne B. Koehler, J. Keith Ord, and Ralph D. Snyder. *Forecasting with Exponential Smoothing: The State Space Approach*. Springer, Berlin, 2008. doi: 10.1007/978-3-540-71918-2.

International Organization for Standardization. ISO 20022: Financial services — universal financial industry message scheme, 2013.

John P. A. Ioannidis. Why most published research findings are false. *PLOS Medicine*, 2(8): e124, 2005. doi: 10.1371/journal.pmed.0020124.

Anthony R. Ives and Stephen R. Carpenter. Stability and diversity of ecosystems. *Science*, 317 (5834):58–62, 2007. doi: 10.1126/science.1133258.

Jane Jacobs. *The Death and Life of Great American Cities*. Random House, New York, 1961.

Edwin T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620–630, 1957. doi: 10.1103/PhysRev.106.620.

Edwin T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge, 2003. ISBN 978-0-521-59271-0.

Lou Jost. Entropy and diversity. *Oikos*, 113(2):363–375, 2006. doi: 10.1111/j.2006.0030-1299.14714.x.

Lou Jost. Partitioning diversity into independent alpha and beta components. *Ecology*, 88(10): 2427–2439, 2007. doi: 10.1890/06-1736.1.

Lee Jussim and Kent D. Harber. Teacher expectations and self-fulfilling prophecies: Knowns and unknowns, resolved and unresolved controversies. *Personality and Social Psychology Review*, 9(2):131–155, 2005. doi: 10.1207/s15327957pspr0902_3.

Marko Jusup, Petter Holme, Kiyoshi Kanazawa, Misako Takayasu, Ivan Romić, Zhen Wang, Sunčana Geček, Tomislav Lipić, Boris Podobnik, Lin Wang, Wei Luo, Tin Klanjšček, Jingfang Fan, Stefano Boccaletti, and Matjaž Perc. Social physics. *Physics Reports*, 948:1–148, 2022. doi: 10.1016/j.physrep.2021.10.005.

Daniel Kahneman, Alan B. Krueger, David Schkade, Norbert Schwarz, and Arthur A. Stone. Would you be happier if you were richer? a focusing illusion. *Science*, 312(5782):1908–1910, 2006. doi: 10.1126/science.1129688.

Rudolf E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960. doi: 10.1115/1.3662552.

Akio Kameoka and Mamoru Kuriyagawa. Consonance theory part I: Consonance of dyads. *The Journal of the Acoustical Society of America*, 45(6):1451–1459, 1969. doi: 10.1121/1.1911623.

Walter G. Kelley and Allan C. Peterson. *Difference Equations: An Introduction with Applications*. Academic Press, San Diego, CA, 2 edition, 2001.

Earle H. Kennard. Zur quantenmechanik einfacher bewegungstypen. *Zeitschrift für Physik*, 44(4–5):326–352, 1927. doi: 10.1007/BF01391200.

Auguste Kerckhoffs. La cryptographie militaire. *Journal des Sciences Militaires*, 9:5–38, 1883.

Marc Keuschnigg, Niclas Lovsjö, and Peter Hedström. Analytical sociology and computational social science. *Journal of Computational Social Science*, 1(1):3–14, 2018. doi: 10.1007/s42001-017-0006-5.

Aleksandr Ya. Khinchin. *Mathematical Foundations of Information Theory*. Dover Publications, New York, 1957.

John F. C. Kingman. On queues in heavy traffic. *Journal of the Royal Statistical Society, Series B*, 24(2):383–392, 1962.

Alan P. Kirman. Whom or what does the representative individual represent? *Journal of Economic Perspectives*, 6(2):117–136, 1992. doi: 10.1257/jep.6.2.117.

Thomas Koshy. *Fibonacci and Lucas Numbers with Applications*. Wiley-Interscience, New York, 2001. ISBN 978-0-471-39969-8.

David H. Krantz, R. Duncan Luce, Patrick Suppes, and Amos Tversky. *Foundations of Measurement, Volume I: Additive and Polynomial Representations*. Academic Press, New York, 1971.

Lester E. Krueger. Reconciling Fechner and Stevens: Toward a unified psychophysical law. *Behavioral and Brain Sciences*, 12(2):251–267, 1989. doi: 10.1017/S0140525X0004855X.

Daniel Kuehn. Diversity, ability, and democracy: A note on Thompson's challenge to Hong and Page. *Critical Review*, 29(1):72–87, 2017. doi: 10.1080/08913811.2017.1288455.

Solomon Kullback and Richard A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951. doi: 10.1214/aoms/1177729694.

Markku Laakso and Rein Taagepera. "effective" number of parties: A measure with application to West Europe. *Comparative Political Studies*, 12(1):3–27, 1979. doi: 10.1177/001041407901200101.

Imre Lakatos. Falsification and the methodology of scientific research programmes. In Imre Lakatos and Alan Musgrave, editors, *Criticism and the Growth of Knowledge*, pages 91–196. Cambridge University Press, Cambridge, 1970. doi: 10.1017/CBO9781139171434.009.

Steve K. Lamoreaux. Demonstration of the Casimir force in the 0.6 to 6 μm range. *Physical Review Letters*, 78(1):5–8, 1997. doi: 10.1103/PhysRevLett.78.5.

Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982. doi: 10.1145/357172.357176.

Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961. doi: 10.1147/rd.53.0183.

Louis Lapicque. Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarisation. *Journal de Physiologie et de Pathologie Générale*, 9:620–635, 1907.

Bibb Latané. The psychology of social impact. *American Psychologist*, 36(4):343–356, 1981. doi: 10.1037/0003-066X.36.4.343.

Richard S. Lazarus. *Emotion and Adaptation*. Oxford University Press, New York, 1991. ISBN 978-0-19-506994-5.

David Lazer, Alex Pentland, Lada Adamic, Sinan Aral, Albert-László Barabási, Devon Brewer, Nicholas Christakis, Noshir Contractor, James Fowler, Myron Gutmann, Tony Jebara, Gary King, Michael Macy, Deb Roy, and Marshall Van Alstyne. Computational social science. *Science*, 323(5915):721–723, 2009. doi: 10.1126/science.1167742.

David Lazer, Ryan Kennedy, Gary King, and Alessandro Vespignani. The parable of Google Flu: Traps in big data analysis. *Science*, 343(6176):1203–1205, 2014. doi: 10.1126/science.1248506.

Tom Leinster and Christina A. Cobbold. Measuring diversity: The importance of species similarity. *Ecology*, 93(3):477–489, 2012. doi: 10.1890/10-2402.1.

Nancy G. Leveson. *Engineering a Safer World: Systems Thinking Applied to Safety*. MIT Press, Cambridge, MA, 2011. ISBN 978-0-262-01662-9.

Daniel A. Levinthal and James G. March. The myopia of learning. *Strategic Management Journal*, 14(S2):95–112, 1993. doi: 10.1002/smj.4250141009.

Rensis Likert. A technique for the measurement of attitudes. *Archives of Psychology*, 22(140):1–55, 1932.

Mark Lindley and Ronald Turner-Smith. *Mathematical Models of Musical Scales: A New Approach*. Verlag für systematische Musikwissenschaft, Bonn, 1993.

John D. C. Little. A proof for the queuing formula: $L = \lambda W$. *Operations Research*, 9(3):383–387, 1961. doi: 10.1287/opre.9.3.383.

A. C. Littleton and B. S. Yamey, editors. *Studies in the History of Accounting*. Sweet and Maxwell, London, 1956.

Mario Livio. *The Golden Ratio: The Story of Phi, the World's Most Astonishing Number*. Broadway Books, New York, 2002. ISBN 978-0-7679-0815-3.

François Longin and Bruno Solnik. Extreme correlation of international equity markets. *The Journal of Finance*, 56(2):649–676, 2001. doi: 10.1111/0022-1082.00340.

Michel Loreau and Claire de Mazancourt. Species synchrony and its drivers: Neutral and nonneutral community dynamics in fluctuating environments. *The American Naturalist*, 172(2):E48–E66, 2008. doi: 10.1086/589746.

Jan Lorenz. Continuous opinion dynamics under bounded confidence: A survey. *International Journal of Modern Physics C*, 18(12):1819–1838, 2007. doi: 10.1142/S0129183107011789.

Jan Lorenz, Heiko Rauhut, Frank Schweitzer, and Dirk Helbing. How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences*, 108(22):9020–9025, 2011. doi: 10.1073/pnas.1008636108.

Robert E. Lucas. Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy*, 1:19–46, 1976. doi: 10.1016/S0167-2231(76)80003-6.

R. Duncan Luce. A psychophysical theory of intensity proportions, joint presentations, and matches. *Psychological Review*, 109(3):520–532, 2002. doi: 10.1037/0033-295X.109.3.520.

R. Duncan Luce and John W. Tukey. Simultaneous conjoint measurement: A new type of fundamental measurement. *Journal of Mathematical Psychology*, 1(1):1–27, 1964. doi: 10.1016/0022-2496(64)90015-X.

Maurizio Lupoi. *Trusts: A Comparative Study*. Cambridge University Press, Cambridge, 2000. ISBN 978-0-521-62329-3.

Russell Lyons. The spread of evidence-poor medicine via flawed social-network analysis. *Statistics, Politics, and Policy*, 2(1), 2011. doi: 10.2202/2151-7509.1024.

Donald MacKenzie. *An Engine, Not a Camera: How Financial Models Shape Markets*. MIT Press, Cambridge, MA, 2006. doi: 10.7551/mitpress/9780262134606.001.0001.

Donald MacKenzie and Yuval Millo. Constructing a market, performing theory: The historical sociology of a financial derivatives exchange. *American Journal of Sociology*, 109(1):107–145, 2003. doi: 10.1086/374404.

Elizabeth Mannix and Margaret A. Neale. What differences make a difference? the promise and reality of diverse teams in organizations. *Psychological Science in the Public Interest*, 6(2):31–55, 2005. doi: 10.1111/j.1529-1006.2005.00022.x.

Charles F. Manski. Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, 60(3):531–542, 1993. doi: 10.2307/2298123.

James G. March. Exploration and exploitation in organizational learning. *Organization Science*, 2(1):71–87, 1991. doi: 10.1287/orsc.2.1.71.

James G. March and Herbert A. Simon. *Organizations*. John Wiley and Sons, New York, 1958.

Harry Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952. doi: 10.2307/2975974.

Harry M. Markowitz. *Portfolio Selection: Efficient Diversification of Investments*. John Wiley and Sons, New York, 1959.

George Markowsky. Misconceptions about the golden ratio. *The College Mathematics Journal*, 23(1):2–19, 1992. doi: 10.1080/07468342.1992.11973428.

Robert M. May. Will a large complex system be stable? *Nature*, 238(5364):413–414, 1972. doi: 10.1038/238413a0.

Jim McCambridge, John Witton, and Diana R. Elbourne. Systematic review of the Hawthorne effect: New concepts are needed to study research participation effects. *Journal of Clinical Epidemiology*, 67(3):267–277, 2014. doi: 10.1016/j.jclinepi.2013.08.015.

Kevin S. McCann. The diversity-stability debate. *Nature*, 405(6783):228–233, 2000. doi: 10.1038/35012234.

Joseph L. McCauley. Response to "worrying trends in econophysics". *Physica A: Statistical Mechanics and its Applications*, 371(2):601–609, 2006. doi: 10.1016/j.physa.2006.05.043.

Paul E. Meehl. Theory-testing in psychology and physics: A methodological paradox. *Philosophy of Science*, 34(2):103–115, 1967. doi: 10.1086/288135.

Paul E. Meehl. Appraising and amending theories: The strategy of Lakatosian defense and two principles that warrant it. *Psychological Inquiry*, 1(2):108–141, 1990. doi: 10.1207/s15327965pli0102_1.

Robert K. Merton. The unanticipated consequences of purposive social action. *American Sociological Review*, 1(6):894–904, 1936. doi: 10.2307/2084615.

Robert K. Merton. The self-fulfilling prophecy. *The Antioch Review*, 8(2):193–210, 1948. doi: 10.2307/4609267.

Joel Michell. Quantitative science and the definition of measurement in psychology. *British Journal of Psychology*, 88(3):355–383, 1997. doi: 10.1111/j.2044-8295.1997.tb02641.x.

Joel Michell. *Measurement in Psychology: A Critical History of a Methodological Concept*. Cambridge University Press, Cambridge, 1999. ISBN 978-0-521-62120-6.

George A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2):81–97, 1956. doi: 10.1037/h0043158.

Philip Mirowski. *More Heat than Light: Economics as Social Physics, Physics as Nature's Economics*. Cambridge University Press, Cambridge, 1989. doi: 10.1017/CBO9780511559990.

Umar Mohideen and Anushree Roy. Precision measurement of the Casimir force from 0.1 to 0.9 μm. *Physical Review Letters*, 81(21):4549–4552, 1998. doi: 10.1103/PhysRevLett.81.4549.

Michael S. Morris and Kip S. Thorne. Wormholes in spacetime and their use for interstellar travel: A tool for teaching general relativity. *American Journal of Physics*, 56(5):395–412, 1988. doi: 10.1119/1.15620.

Louis Narens. A theory of ratio magnitude estimation. *Journal of Mathematical Psychology*, 40(2):109–129, 1996. doi: 10.1006/jmps.1996.0011.

M. E. J. Newman. *Networks*. Oxford University Press, Oxford, 2 edition, 2018. ISBN 978-0-19-880509-0.

Richard E. Nisbett and Timothy D. Wilson. Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84(3):231–259, 1977. doi: 10.1037/0033-295X.84.3.231.

Brian A. Nosek, Charles R. Ebersole, Alexander C. DeHaven, and David T. Mellor. The preregistration revolution. *Proceedings of the National Academy of Sciences*, 115(11):2600–2606, 2018. doi: 10.1073/pnas.1708274114.

Andrzej Nowak, Jacek Szamrej, and Bibb Latané. From private attitude to public opinion: A dynamic theory of social impact. *Psychological Review*, 97(3):362–376, 1990. doi: 10.1037/0033-295X.97.3.362.

Martin A. Nowak and Karl Sigmund. Evolution of indirect reciprocity. *Nature*, 437(7063):1291–1298, 2005. doi: 10.1038/nature04131.

Harry Nyquist. Certain topics in telegraph transmission theory. *Transactions of the American Institute of Electrical Engineers*, 47(2):617–644, 1928. doi: 10.1109/T-AIEE.1928.5055024.

Open Science Collaboration. Estimating the reproducibility of psychological science. *Science*, 349(6251):aac4716, 2015. doi: 10.1126/science.aac4716.

Alan V. Oppenheim, Alan S. Willsky, and S. Hamid Nawab. *Signals and Systems*. Prentice Hall, Upper Saddle River, NJ, 2 edition, 1996. ISBN 978-0-13-814757-0.

Organisation for Economic Co-operation and Development. OECD/G20 base erosion and profit shifting project: 2015 final reports, 2015.

Organisation for Economic Co-operation and Development. Standard for automatic exchange of financial account information in tax matters, 2017.

Elinor Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, Cambridge, 1990. doi: 10.1017/CBO9780511807763.

Elinor Ostrom. Beyond markets and states: Polycentric governance of complex economic systems. *The American Economic Review*, 100(3):641–672, 2010. doi: 10.1257/aer.100.3.641.

Scott E. Page. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton University Press, Princeton, NJ, 2007. ISBN 978-0-691-12838-2.

Scott E. Page. *Diversity and Complexity*. Princeton University Press, Princeton, NJ, 2010. ISBN 978-0-691-13767-4.

Paolo Panico. *Private Foundations: Law and Practice*. Oxford University Press, Oxford, 2014. ISBN 978-0-19-968873-1.

Harry Partch. *Genesis of a Music: An Account of a Creative Work, Its Roots and Its Fulfillments*. Da Capo Press, New York, 2 edition, 1974.

Alex Pentland. *Social Physics: How Good Ideas Spread — The Lessons from a New Science*. The Penguin Press, New York, 2014.

Charles Perrow. *Normal Accidents: Living with High-Risk Technologies*. Basic Books, New York, 1984.

Ole Peters. The ergodicity problem in economics. *Nature Physics*, 15(12):1216–1221, 2019. doi: 10.1038/s41567-019-0732-0.

Ole Peters and Murray Gell-Mann. Evaluating gambles using dynamics. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 26(2):023103, 2016. doi: 10.1063/1.4940236.

Thomas F. Pettigrew and Linda R. Tropp. A meta-analytic test of intergroup contact theory. *Journal of Personality and Social Psychology*, 90(5):751–783, 2006. doi: 10.1037/0022-3514.90.5.751.

E. C. Pielou. The measurement of diversity in different types of biological collections. *Journal of Theoretical Biology*, 13:131–144, 1966. doi: 10.1016/0022-5193(66)90013-0.

Katharina Pistor. *The Code of Capital: How the Law Creates Wealth and Inequality*. Princeton University Press, Princeton, NJ, 2019. ISBN 978-0-691-17897-4.

John R. Platt. Strong inference. *Science*, 146(3642):347–353, 1964. doi: 10.1126/science.146.3642.347.

Reinier Plomp and Willem J. M. Levelt. Tonal consonance and critical bandwidth. *The Journal of the Acoustical Society of America*, 38(4):548–560, 1965. doi: 10.1121/1.1909741.

Karl R. Popper. *The Logic of Scientific Discovery*. Hutchinson, London, 1959.

E. C. Poulton. The new psychophysics: Six models for magnitude estimation. *Psychological Bulletin*, 69(1):1–19, 1968. doi: 10.1037/h0025267.

E. C. Poulton. Models for biases in judging sensory magnitude. *Psychological Bulletin*, 86(4):777–803, 1979. doi: 10.1037/0033-2909.86.4.777.

David Premack and Guy Woodruff. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 1(4):515–526, 1978. doi: 10.1017/S0140525X00076512.

Robert D. Putnam. E pluribus unum: Diversity and community in the twenty-first century. *Scandinavian Political Studies*, 30(2):137–174, 2007. doi: 10.1111/j.1467-9477.2007.00176.x.

Robert D. Putnam, Robert Leonardi, and Raffaella Y. Nanetti. *Making Democracy Work: Civic Traditions in Modern Italy*. Princeton University Press, Princeton, NJ, 1993. ISBN 978-0-691-03738-8.

Adolphe Quetelet. *Sur l'homme et le développement de ses facultés, ou Essai de physique sociale*. Bachelier, Paris, 1835.

C. Radhakrishna Rao. Diversity and dissimilarity coefficients: A unified approach. *Theoretical Population Biology*, 21(1):24–43, 1982. doi: 10.1016/0040-5809(82)90004-1.

Georg Rasch. *Probabilistic Models for Some Intelligence and Attainment Tests*. Danmarks Paedagogiske Institut, Copenhagen, 1960.

Alfréd Rényi. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, pages 547–561, Berkeley, CA, 1961. University of California Press.

Robert Rosenthal and Kermit L. Fode. The effect of experimenter bias on the performance of the albino rat. *Behavioral Science*, 8(3):183–189, 1963. doi: 10.1002/bs.3830080302.

Robert Rosenthal and Lenore Jacobson. Pygmalion in the classroom. *The Urban Review*, 3(1): 16–20, 1968. doi: 10.1007/BF02322211.

Thomas D. Rossing, F. Richard Moore, and Paul A. Wheeler. *The Science of Sound*. Addison Wesley, San Francisco, CA, 3 edition, 2002. ISBN 978-0-8053-8565-8.

James A. Russell. A circumplex model of affect. *Journal of Personality and Social Psychology*, 39(6):1161–1178, 1980. doi: 10.1037/h0077714.

Wallace Clement Sabine. *Collected Papers on Acoustics*. Harvard University Press, Cambridge, MA, 1922.

Scott D. Sagan. *The Limits of Safety: Organizations, Accidents, and Nuclear Weapons*. Princeton University Press, Princeton, NJ, 1993. ISBN 978-0-691-02101-0.

Matthew J. Salganik, Peter Sheridan Dodds, and Duncan J. Watts. Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311(5762):854–856, 2006. doi: 10.1126/science.1121066.

Matthew J. Salganik, Ian Lundberg, Alexander T. Kindel, et al. Measuring the predictability of life outcomes with a scientific mass collaboration. *Proceedings of the National Academy of Sciences*, 117(15):8398–8403, 2020. doi: 10.1073/pnas.1915006117.

Jerome H. Saltzer and Michael D. Schroeder. The protection of information in computer systems. *Proceedings of the IEEE*, 63(9):1278–1308, 1975. doi: 10.1109/PROC.1975.9939.

Thomas C. Schelling. Dynamic models of segregation. *The Journal of Mathematical Sociology*, 1(2):143–186, 1971. doi: 10.1080/0022250X.1971.9989794.

Thomas C. Schelling. *Micromotives and Macrobehavior*. W. W. Norton, New York, 1978.

Klaus R. Scherer. Appraisal considered as a process of multilevel sequential checking. In Klaus R. Scherer, Angela Schorr, and Tom Johnstone, editors, *Appraisal Processes in Emotion: Theory, Methods, Research*, Series in Affective Science, pages 92–120. Oxford University Press, New York, 2001. ISBN 978-0-19-513007-2. doi: 10.1093/oso/9780195130072.003.0005.

Daniel E. Schindler, Ray Hilborn, Brandon Chasco, Christopher P. Boatright, Thomas P. Quinn, Lauren A. Rogers, and Michael S. Webster. Population diversity and the portfolio effect in an exploited species. *Nature*, 465(7298):609–612, 2010. doi: 10.1038/nature09060.

Manfred R. Schroeder. New method of measuring reverberation time. *The Journal of the Acoustical Society of America*, 37(3):409–412, 1965. doi: 10.1121/1.1909343.

Norbert Schwarz. Self-reports: How the questions shape the answers. *American Psychologist*, 54(2):93–105, 1999. doi: 10.1037/0003-066X.54.2.93.

James C. Scott. *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press, New Haven, CT, 1998. ISBN 978-0-300-07016-3.

William A. Sethares. *Tuning, Timbre, Spectrum, Scale*. Springer, London, 2 edition, 2005. doi: 10.1007/b138848.

Cosma Rohilla Shalizi and Andrew C. Thomas. Homophily and contagion are generically confounded in observational social network studies. *Sociological Methods and Research*, 40(2): 211–239, 2011. doi: 10.1177/0049124111404820.

Claude E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423, 1948. doi: 10.1002/j.1538-7305.1948.tb01338.x.

Claude E. Shannon and Warren Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, Urbana, IL, 1949.

Roger N. Shepard. Geometrical approximations to the structure of musical pitch. *Psychological Review*, 89(4):305–333, 1982. doi: 10.1037/0033-295X.89.4.305.

Roger N. Shepard. Toward a universal law of generalization for psychological science. *Science*, 237(4820):1317–1323, 1987. doi: 10.1126/science.3629243.

Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997. doi: 10.1137/S0097539795293172.

Klaas Sijtsma. On the use, the misuse, and the very limited usefulness of Cronbach's alpha. *Psychometrika*, 74(1):107–120, 2009. doi: 10.1007/s11336-008-9101-0.

Joseph P. Simmons, Leif D. Nelson, and Uri Simonsohn. False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22(11):1359–1366, 2011. doi: 10.1177/0956797611417632.

E. H. Simpson. Measurement of diversity. *Nature*, 163(4148):688, 1949. doi: 10.1038/163688a0.

Daniel J. Singer. Diversity, not randomness, trumps ability. *Philosophy of Science*, 86(1):178–191, 2019. doi: 10.1086/701074.

Craig A. Smith and Phoebe C. Ellsworth. Patterns of cognitive appraisal in emotion. *Journal of Personality and Social Psychology*, 48(4):813–838, 1985. doi: 10.1037/0022-3514.48.4.813.

Lionel Smith, editor. *The Worlds of the Trust*. Cambridge University Press, Cambridge, 2013. ISBN 978-1-107-01220-2.

Pawel Sobkowicz. Modelling opinion formation with physics tools: Call for closer link with reality. *Journal of Artificial Societies and Social Simulation*, 12(1):11, 2009.

Pawel Sobkowicz. Whither now, opinion modelers? *Frontiers in Physics*, 8:587009, 2020. doi: 10.3389/fphy.2020.587009.

Didier Sornette. *Why Stock Markets Crash: Critical Events in Complex Financial Systems*. Princeton University Press, Princeton, NJ, 2003. doi: 10.1515/9781400829552.

George Soros. Fallibility, reflexivity, and the human uncertainty principle. *Journal of Economic Methodology*, 20(4):309–329, 2013. doi: 10.1080/1350178X.2013.859415.

Susan Leigh Star and James R. Griesemer. Institutional ecology, 'translations' and boundary objects. *Social Studies of Science*, 19(3):387–420, 1989. doi: 10.1177/030631289019003001.

Susan Leigh Star and Karen Ruhleder. Steps toward an ecology of infrastructure: Design and access for large information spaces. *Information Systems Research*, 7(1):111–134, 1996. doi: 10.1287/isre.7.1.111.

Meir Statman. How many stocks make a diversified portfolio? *The Journal of Financial and Quantitative Analysis*, 22(3):353–363, 1987. doi: 10.2307/2330969.

Dietrich Stauffer and Amnon Aharony. *Introduction to Percolation Theory*. Taylor and Francis, London, 2 edition, 1994. ISBN 978-0-7484-0253-3.

John D. Sterman. Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Science*, 35(3):321–339, 1989. doi: 10.1287/mnsc.35.3.321.

John D. Sterman. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Irwin/McGraw-Hill, Boston, MA, 2000. ISBN 978-0-07-231135-8.

S. S. Stevens. On the theory of scales of measurement. *Science*, 103(2684):677–680, 1946. doi: 10.1126/science.103.2684.677.

S. S. Stevens. Calculation of the loudness of complex noise. *The Journal of the Acoustical Society of America*, 28(5):807–832, 1956. doi: 10.1121/1.1908487.

S. S. Stevens. On the psychophysical law. *Psychological Review*, 64(3):153–181, 1957. doi: 10.1037/h0046162.

S. S. Stevens. *Psychophysics: Introduction to Its Perceptual, Neural, and Social Prospects*. John Wiley and Sons, New York, 1975.

George J. Stigler. A theory of oligopoly. *Journal of Political Economy*, 72(1):44–61, 1964. doi: 10.1086/258853.

Steven H. Strogatz. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. Westview Press, Boulder, CO, 2 edition, 2015. ISBN 978-0-8133-4910-7.

Michael P. H. Stumpf and Mason A. Porter. Critical truths about power laws. *Science*, 335 (6069):665–666, 2012. doi: 10.1126/science.1216142.

James Surowiecki. *The Wisdom of Crowds*. Doubleday, New York, 2004.

Robert Teghtsoonian. On the exponents in Stevens' law and the constant in Ekman's law. *Psychological Review*, 78(1):71–80, 1971. doi: 10.1037/h0030300.

James Tenney. John cage and the theory of harmony (1983). In Larry Polansky, Lauren Pratt, Robert Wannamaker, and Michael Winter, editors, *From Scratch: Writings in Music Theory*. University of Illinois Press, Champaign, IL, 2015. ISBN 978-0-252-03872-3.

Ernst Terhardt. Pitch, consonance, and harmony. *The Journal of the Acoustical Society of America*, 55(5):1061–1069, 1974. doi: 10.1121/1.1914648.

Abigail Thompson. Does diversity trump ability? an example of the misuse of mathematics. *Notices of the American Mathematical Society*, 61(9):1024–1030, 2014. doi: 10.1090/noti1163.

D'Arcy Wentworth Thompson. *On Growth and Form*. Cambridge University Press, Cambridge, 1917.

Louis L. Thurstone. A law of comparative judgment. *Psychological Review*, 34(4):273–286, 1927. doi: 10.1037/h0070288.

David Tilman, Clarence L. Lehman, and Charles E. Bristow. Diversity-stability relationships: Statistical inevitability or ecological consequence? *The American Naturalist*, 151(3):277–282, 1998. doi: 10.1086/286118.

Paul Torremans, Richard Fentiman, Ugljesa Grusic, Christian Heinze, Louise Merrett, Alex Mills, Zheng Sophia Tang, and Katarina Trimmings. *Dicey, Morris and Collins on the Conflict of Laws*. Sweet and Maxwell, London, 16 edition, 2022.

Jessica L. Tracy and Richard W. Robins. Putting the self into self-conscious emotions: A theoretical model. *Psychological Inquiry*, 15(2):103–125, 2004. doi: 10.1207/s15327965pli1502_01.

Constantino Tsallis. Possible generalization of Boltzmann-Gibbs statistics. *Journal of Statistical Physics*, 52(1–2):479–487, 1988. doi: 10.1007/BF01016429.

Arnold Tustin. The nature of the operator's response in manual control and its implications for controller design. *Journal of the Institution of Electrical Engineers — Part IIA*, 94(2): 190–206, 1947. doi: 10.1049/ji-2a.1947.0025.

Amos Tversky. Features of similarity. *Psychological Review*, 84(4):327–352, 1977. doi: 10.1037/0033-295X.84.4.327.

Johan Ugander, Lars Backstrom, Cameron Marlow, and Jon Kleinberg. Structural diversity in social contagion. *Proceedings of the National Academy of Sciences*, 109(16):5962–5966, 2012. doi: 10.1073/pnas.1116502109.

United Nations. Convention on the recognition and enforcement of foreign arbitral awards. United Nations Treaty Series, vol. 330, no. 4739, p. 3, 1958.

United Nations Commission on International Trade Law. UNCITRAL model law on international commercial arbitration, 1985.

U.S. Department of Justice and Federal Trade Commission. Horizontal merger guidelines, August 2010.

U.S. Department of Justice and Federal Trade Commission. Merger guidelines. U.S. Department of Justice and Federal Trade Commission, December 2023.

Marius Usher and James L. McClelland. The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108(3):550–592, 2001. doi: 10.1037/0033-295X.108.3.550.

Brian Uzzi and Jarrett Spiro. Collaboration and creativity: The small world problem. *American Journal of Sociology*, 111(2):447–504, 2005. doi: 10.1086/432782.

Steven Vajda. *Fibonacci and Lucas Numbers, and the Golden Section: Theory and Applications*. Ellis Horwood, Chichester, 1989.

Albert Jan van den Berg. *The New York Arbitration Convention of 1958: Towards a Uniform Judicial Interpretation*. Kluwer Law and Taxation, Deventer, 1981.

Daan van Knippenberg and Michéla C. Schippers. Work group diversity. *Annual Review of Psychology*, 58:515–541, 2007. doi: 10.1146/annurev.psych.58.110405.085546.

Diane Vaughan. *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. University of Chicago Press, Chicago, IL, 1996. ISBN 978-0-226-85175-4.

Dejan Vinković and Alan Kirman. A physical analogue of the Schelling model. *Proceedings of the National Academy of Sciences*, 103(51):19261–19265, 2006. doi: 10.1073/pnas.0609371103.

Hermann von Helmholtz. *Die Lehre von den Tonempfindungen als physiologische Grundlage für die Theorie der Musik*. Friedrich Vieweg und Sohn, Braunschweig, 1863.

Hermann von Helmholtz. *On the Sensations of Tone as a Physiological Basis for the Theory of Music*. Longmans, Green, London, 2 english edition, 1885.

Duncan J. Watts. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9):5766–5771, 2002. doi: 10.1073/pnas.082090499.

Duncan J. Watts. Common sense and sociological explanations. *American Journal of Sociology*, 120(2):313–351, 2014. doi: 10.1086/678271.

Duncan J. Watts and Steven H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, 1998. doi: 10.1038/30918.

Ernst Heinrich Weber. *De Pulsu, Resorptione, Auditu et Tactu: Annotationes Anatomicae et Physiologicae*. C. F. Koehler, Leipzig, 1834.

Ernst Heinrich Weber. *The Sense of Touch*. Academic Press, London, 1978.

Karl E. Weick. The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4):628–652, 1993. doi: 10.2307/2393339.

Karl E. Weick. *Sensemaking in Organizations*. Sage Publications, Thousand Oaks, CA, 1995. ISBN 978-0-8039-7177-6.

Karl E. Weick and Kathleen M. Sutcliffe. *Managing the Unexpected: Resilient Performance in an Age of Uncertainty*. Jossey-Bass, San Francisco, CA, 2 edition, 2007. ISBN 978-0-7879-9649-9.

Wolfgang Weidlich. The statistical description of polarization phenomena in society. *British Journal of Mathematical and Statistical Psychology*, 24(2):251–266, 1971. doi: 10.1111/j.2044-8317.1971.tb00470.x.

Wolfgang Weidlich. Sociodynamics — a systematic approach to mathematical modelling in the social sciences. *Chaos, Solitons and Fractals*, 18(3):431–437, 2003. doi: 10.1016/S0960-0779(02)00666-5.

William H. Whyte. *The Social Life of Small Urban Spaces*. Conservation Foundation, Washington, DC, 1980.

Norbert Wiener. *Cybernetics: Or Control and Communication in the Animal and the Machine*. MIT Press, Cambridge, MA, 1948.

Anita Williams Woolley, Christopher F. Chabris, Alex Pentland, Nada Hashmi, and Thomas W. Malone. Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004):686–688, 2010. doi: 10.1126/science.1193147.

Basil S. Yamey. Accounting and the rise of capitalism: Further notes on a theme by Sombart. *Journal of Accounting Research*, 2(2):117–136, 1964. doi: 10.2307/2489995.

H. Dieter Zeh. On the interpretation of measurement in quantum theory. *Foundations of Physics*, 1(1):69–76, 1970. doi: 10.1007/BF00708656.

Index

a posteriori bound, 17, 119, 1271, 3007
acoustic impedance, 1805
ADGM, 3359, 3912, 3946
adjustability, 448, 553, 813, 2417
aggregation fallacy, 1344
aggregation gap, 2735
agreement identity, 2697
anonymity set, 1738, 3762
antipodal corpus, 2173
asset-holding structures, 3132
automatic exchange of information, 3132
axial modes, 877, 928

balance index, 405, 614
bank run, 2054, 2099, 3871
bankruptcy remoteness, 3509
barrier integral, 1188
Bayesian posterior, 3669
bill of exchange, 3624
binding constraint, 261, 324, 477, 588, 623, 798, 848, 945, 1041, 1132, 1228, 1331, 1670, 1922, 2058, 2227, 2390, 2588, 3165, 3229, 3284, 3324, 3440, 3596, 3636, 3688, 3826, 3864, 3945
bloc classification, 3438, 3625
bloc collapse, 51, 3133, 3181, 3231, 3282, 3317, 3363, 3409, 3441, 3672, 3798, 3833, 3931
bloc independence, 3383, 3400, 3911, 3942
Born rule, 64, 1933, 1987
borrowing base, 3829
boundary cost, 611
boundary geometry, 497, 1763, 1987, 2383, 3119
branch partnership, 3669
break-even bloc count, 3931
break-even probability, 567, 3438, 3479, 3532, 3624, 3905, 3950
bridging entity, 3400
Brooks crossover, 2735

cadence, 144, 198, 274, 365, 435, 620, 719, 755, 796, 913, 928, 1013, 1147, 1266, 1360, 1574, 1650, 1739, 2256, 2382, 2785, 2820, 2875, 3072, 3113, 3193, 3811
cadence folding, 928
cadence theorem, 1650, 2914, 3072
calibration ratio, 23, 943, 1002, 1224, 1294, 1344, 2418
canonicalisation, 1990, 2082, 3669, 3719, 3829
capped pilot, 775
captive insurance, 3316
capture gates, 260
carrier-grade translation, 3762
carrying cost, 30, 174, 2457, 3231, 3317, 3477
cascade, 222, 336, 405, 579, 623, 648, 803, 1146, 1259, 1487, 1738, 1884, 2055, 2118, 2208, 2358, 2415, 3560, 3591, 3903, 3941
cascade decay, 405
Casimir effect, 1987, 2121
category violation, 89, 265, 3537, 3610
ceiling, 23, 31, 198, 260, 323, 405, 491, 554, 611, 650, 722, 776, 849, 928, 1003, 1071, 1190, 1264, 1345, 1417, 1516, 1565, 1651, 1921, 2047, 2101, 2174, 2241, 2373, 2464, 2505, 2587, 2628, 2703, 2735, 2822, 2880, 2962, 3017, 3088, 3166, 3184, 3268, 3336, 3360, 3405, 3443, 3574, 3626, 3747, 3798, 3830, 3886, 3935
ceiling inflation, 1565
ceiling utilisation, 405, 614
cell companies, 3313, 3509, 3987
central bank digital currency, 3587, 3797
chain law, 2140, 2173
channel capacity, 1716, 1805, 1917, 3596, 3700, 3782
channel gain, 162, 1847, 2043, 2909, 3054
characterisation, 41, 73, 109, 201, 264, 547, 604, 652, 767, 776, 877, 942, 1064, 1117, 1132, 1243, 1348, 1420, 1815, 2098, 2467, 3227, 3228, 3280, 3336, 3407, 3468, 3517, 3589, 3646, 3718, 3795, 3927
chord width, 2586, 2627, 2862
City of London, 3278, 3946
clearing perimeter, 3553, 3916
clearing systems, 3164, 3359, 3946
closure under negation, 2950
closure verdicts, 119
clustered seeding, 611
collision, 51, 63, 96, 176, 337, 436, 734, 793, 896, 928, 1148, 1621, 1687, 1805, 1894,

1989, 2106, 2191, 2251, 2314, 2415, 2490, 2587, 2626, 2698, 2818, 2922, 2971, 3092, 3732, 3768, 3847
collision analysis, 1648, 1805
comfort landscape, 2, 33, 154, 198, 556, 614, 892, 929, 1417, 2416, 2491, 2704, 3361, 3401, 3440, 3476, 3652, 3995
commitment device, 89, 187
commons, 101, 311, 455, 523, 582, 632, 648, 808, 851, 1432, 1632, 2522, 2670, 2724, 2735, 2936, 3042, 3119, 3268, 3346, 3385, 3835, 3874
commons design principles, 648
communication topology, 260
completeness failure, 119
completeness gap, 3650, 3669, 3742, 3830
complex contagion, 611, 807, 865, 1702, 2269, 2447
component count, 135, 840, 1090, 1176, 1656, 2153, 2505, 2586, 2809, 2942, 3129, 3864
compound rule, 163, 916, 2472, 2806, 2886, 3065, 3083, 3559
concentration measurement, 3132, 3181, 3255, 3278, 3316, 3359, 3439, 3474, 3510, 3619, 3634, 3670, 3744, 3781, 3816, 3870, 3936
concordance, 433, 1728, 1811, 1879, 1926, 2104, 2179, 2244, 2419, 2637, 2697, 2813, 2883, 2960, 3071, 3080, 3917
concordance register, 2280, 3071, 3080
conditional postulate, 23, 162, 2043, 2099, 2179, 2909
conflict of laws, 3324, 3509, 3589
confusable folding, 3694, 3719
consent, 96, 119, 144, 294, 414, 497, 558, 674, 744, 776, 857, 967, 1042, 1211, 1318, 1439, 2070, 2179, 2243, 2313, 2394, 2519, 2686, 2699, 2742, 2827, 2887, 2959, 3019, 3119, 3133, 3229, 3319, 3518, 3596, 3908, 3945
conservation of drive, 2894, 3015
consistency, 89, 144, 218, 311, 373, 713, 737, 820, 1009, 1080, 1138, 1199, 1270, 1350, 1512, 1569, 1759, 1933, 2017, 2087, 2104, 2221, 2334, 2445, 2477, 2536, 2799, 2867, 2959, 3248, 3292, 3332, 3382, 3420, 3618, 3626, 3670, 3872
consonance ranking, 928
contagion measurement, 3509, 3562
content identifier, 3743, 3829
context switching, 235, 277, 553
contraction, 119, 153, 197, 275, 496, 556, 649, 723, 776, 849, 1004, 1085, 1132, 1189, 1246, 1350, 1417, 1497, 1700, 1852, 1870, 2007, 2053, 2126, 2221, 2251, 2313, 2375, 2474, 2509, 2590, 2701, 2822, 2875, 3015, 3491, 3700, 3746, 3917
contraction mapping, 119, 153, 1246, 1852, 1870, 2875
contraction modulus, 21, 120, 197, 275, 1132, 1852, 1870, 2053, 2320, 2377, 2474, 2590, 2701, 3917
control plane, 1236, 2583, 3762
convergence, 3, 45, 63, 95, 120, 148, 202, 275, 343, 442, 496, 556, 636, 673, 803, 953, 1005, 1074, 1208, 1247, 1353, 1428, 1493, 1768, 2085, 2104, 2251, 2320, 2376, 2467, 2506, 2590, 2700, 2739, 2822, 2876, 3007, 3016, 3106, 3138, 3192, 3278, 3397, 3406, 3735, 3783, 3901, 3947
conversion efficiency, 260, 394, 652, 868, 1131, 1198, 1248, 1421, 1566, 2210, 2244, 2413, 2506, 2698, 2737, 3016, 3216, 3462, 3576
conversion keys, 928, 1007, 3090
convexity, 2626
coordination bloc, 3278
coordination cost, 254, 260, 356, 426, 571, 611, 650, 733, 778, 864, 1142, 1410, 2266, 2358, 2416, 2578, 2724, 2804, 2916, 3048, 3300, 3593, 3955
coordination perimeter, 250, 479, 1645, 3132, 3182, 3234, 3280, 3317, 3360, 3438, 3510, 3559, 3594, 3626, 3669, 3763, 3842, 3866, 3932
correction gain, 797, 1002, 1246
correlated diversification, 3863
correlated exposure, 3181
correlated failure, 2756, 3132, 3650
correspondence table, 143, 1268, 1568, 1671, 1742, 1875, 2313, 2915
correspondent banking, 3165, 3495, 3553
corridor exposure, 3438
cost transfer, 763, 2894, 3015
counterfeit arrival, 188, 1246, 2460, 2893, 3018
courier network, 3669
crenellated boundary, 491
critical multiplier, 611
cross-case comparison, 848
cross-default, 3513, 3589, 3903, 3941
cross-modality matching, 943, 1491, 1598, 2278, 2508
crossing, 38, 182, 254, 263, 405, 525, 568, 612, 650, 744, 857, 932, 1143, 1189, 1276, 1519, 1625, 2330, 2413, 2680, 2719,

2892, 3071, 3161, 3206, 3298, 3731, 3764
Crown Dependencies, 3278, 3317, 3360, 3987
cue placement, 60, 178
currency depreciation, 3474
currency exposure, 3190, 3289, 3316, 3732
currency redenomination, 3553, 3780
customs union, 3474

data plane, 3762, 3832
DC gain, 197, 2251, 2320, 2509, 3019
de-risking, 3144, 3186, 3289, 3361, 3491, 3536, 3553, 3937
decay rate, 163, 197, 285, 343, 412, 492, 556, 782, 1003, 1234, 1383, 1430, 1843, 2318, 2463, 2507, 2587, 2699, 2876, 3135, 3674, 3995
decoherence, 64, 573, 1804, 1852, 1921, 1987, 2065, 2104, 2459
default consolidation, 60
deference geometry, 553
deliberation room, 491, 967, 2355
delivery squeeze, 3863
demand characteristics, 569, 1565, 2115
demodulation, 1735, 1818, 1863, 2217, 2275, 2526, 2990
demotion register, 1437, 3015
design axiom, 33, 146, 199, 264, 323, 406, 534, 604, 649, 780, 852, 929, 1077, 1134, 1297, 1347, 1417, 2251, 2314, 2376, 2466, 2508, 2589, 2629, 2700, 3081, 3185, 3360, 3557, 3594, 3627, 3934
desire line, 31, 185, 448, 565, 811, 1188, 2417
desire lines, 30, 448, 586, 811, 1204
diminishing returns, 14, 652, 1084, 1192, 1420, 1673, 1749, 2251, 3627, 3931
disclosure, 15, 135, 147, 311, 676, 763, 777, 895, 933, 1089, 1427, 1613, 1730, 1806, 1865, 1928, 2080, 2153, 2173, 2287, 2368, 2625, 2652, 2774, 2873, 2942, 2951, 3060, 3119, 3132, 3191, 3228, 3284, 3322, 3365, 3402, 3445, 3480, 3510, 3554, 3591, 3624, 3672, 3720, 3763, 3802, 3830, 3863, 3932
disclosure economics, 3132
disclosure is not immunity, 3163, 3445, 3554, 3863
disclosure test, 763, 3160, 3196, 3235, 3284, 3325, 3368, 3417, 3554, 3624, 3672, 3737, 3877, 3934
discretionary trust, 3228
divergence, 17, 40, 99, 119, 148, 254, 301, 405, 514, 556, 711, 803, 901, 969, 1003, 1123, 1166, 1338, 1370, 1431, 1541, 1923, 2013, 2069, 2100, 2267, 2413, 2474, 2506, 2703, 3195, 3372, 3511, 3641, 3684
DNA data storage, 1863, 2179, 2450
dollar clearing, 3316, 3408, 3550
domestic asset protection trust, 3316
double entry, 3669, 3742, 3843
drop-one, 2587, 2728, 2735, 3119
dry exchange, 3624
dwell time, 452, 563, 1236, 2298, 3072, 3438, 3474, 3797

economic substance, 3132, 3227, 3288, 3359, 3901, 3987
edge density, 553
effect size, 159, 766, 925, 1126, 1138, 1255, 1638, 1725, 1747, 1887, 2043, 2114, 2181, 2259, 2602, 2732, 2872, 2984, 3085, 3353, 3394, 3431
effective blocs, 706, 1851, 3192, 3291, 3319, 3380, 3493, 3546, 3624, 3686, 3763, 3869
effective count, 15, 30, 63, 103, 135, 146, 198, 263, 323, 408, 494, 556, 614, 649, 723, 776, 848, 929, 1006, 1076, 1132, 1192, 1254, 1345, 1419, 1497, 1632, 1723, 1746, 1841, 1879, 1924, 1999, 2051, 2103, 2177, 2240, 2314, 2373, 2466, 2505, 2589, 2629, 2700, 2738, 2810, 2878, 2953, 3018, 3090, 3134, 3182, 3229, 3279, 3319, 3360, 3403, 3439, 3474, 3509, 3554, 3594, 3627, 3670, 3736, 3763, 3798, 3832, 3866, 3931
effective framing count, 1131
effective number, 31, 328, 439, 554, 618, 763, 1259, 1999, 2637, 3181, 3514, 3799, 3832
elasticity, 23, 1003, 1246, 2488, 2654, 2697, 2787, 3016, 3572, 3920
empirical content, 331, 651, 724, 1069, 1139, 1309, 1723, 1866, 1940, 2049, 2364, 2470, 2506, 2592, 2692, 2703, 2741, 2812, 2942, 2953, 3077, 3084, 3168, 3198, 3262, 3301, 3365, 3431, 3441, 3496, 3539, 3577, 3611, 3663, 3704, 3792, 3825, 3858
English governing law, 3278
epistemic bookkeeping, 2312
errata, 307, 850, 935, 2350, 2806, 3085
error decorrelation, 41, 270, 338, 469, 734, 905, 1131, 1351, 2415
error law, 10, 126, 202, 1246, 1469, 1497, 2881, 3024
essence bridge, 23, 906, 1015, 2463, 2506, 2674, 2747, 2885, 3018, 3082
essence index, 848, 933, 1630, 1703, 2400, 2471, 2535, 2589, 2807, 2916, 3215

essence law, 163, 850, 929, 1100, 1626, 1699, 1804, 1865, 1940, 2191, 2351, 2374, 2464, 2505, 2586, 2629, 2698, 2808, 2946, 2951, 3059, 3124, 3559
exponential discounting, 60

failure catalog, 775, 856
failure catalogue, 32, 62, 91, 161, 461, 1702, 2100, 2211, 2240, 2315, 2419
failure taxonomy, 2099, 3515
false-break rate, 1738, 1852, 1908, 3704, 3719
false-positive calibration, 2173
falsification rules, 685, 720, 781, 1477, 1571, 2414
Fechner, 34, 163, 200, 785, 918, 941, 1007, 1069, 1255, 1492, 1716, 1950, 2090, 2251, 2351, 2372, 2468, 2590, 2701, 3015, 3481, 3997
Fechner law, 163, 1071, 2372, 3015, 3481, 3997
feedback transform, 89, 187
financial free zones, 3359
firewall, 2196, 2626, 2711, 3238, 3285, 3324, 3987
fixed point, 2, 81, 99, 120, 148, 198, 275, 343, 487, 522, 556, 615, 652, 723, 777, 853, 932, 1003, 1083, 1137, 1189, 1247, 1350, 1417, 1497, 1581, 1679, 1748, 1925, 2031, 2044, 2104, 2221, 2241, 2313, 2375, 2463, 2505, 2603, 2654, 2698, 2737, 2810, 2876, 2963, 3015, 3088, 3135, 3183, 3497, 3700, 3746, 3782, 3801
forced heirship, 3228, 3938
forced sale, 3260, 3589, 3904, 3941
forecast resolution, 1920
formula validation, 162, 1038, 1247, 1491, 2044, 2243, 2348, 2380, 2554
formulary demotion, 1246
Fourier duality, 1844, 1898, 1920, 1995, 3084
freeze property, 3797
frequency tables, 3080
frozen denominator, 775, 864
funding diversification, 3474

Gabor bound, 1854, 1911, 1923, 2950, 3084
Gabor limit, 905, 1805, 1863, 1921, 2178, 2251, 2419, 2954, 3062
gated adoption, 3863
genomic addressing, 1805, 1916
genomic audio, 1843, 2950
geometric accumulator, 1069, 2327
geometric mean index, 2806
Goodhart drift, 1565
gradient descent, 2, 38, 1188
graduated sanction, 1416
gradus suavitatis, 951, 2626, 2810, 2967, 3083
greatest common divisor, 1650
group moment, 1344
growth mindset, 2099
Gulf financial centres, 3400

half-life, 128, 197, 291, 343, 434, 587, 721, 788, 1005, 1153, 1247, 1366, 1431, 1685, 2256, 2314, 2410, 2474, 2586, 2877, 3024, 3102, 3677
harmonic distance, 51, 951, 2626, 2711, 2810, 2879, 2967, 3083
Hart-Scott-Rodino, 3988
headroom, 44, 162, 202, 260, 325, 407, 491, 553, 612, 649, 723, 775, 849, 968, 1003, 1082, 1131, 1200, 1277, 1344, 1420, 1590, 1670, 1820, 1925, 2058, 2238, 2240, 2312, 2372, 2474, 2519, 2710, 2752, 2894, 3043, 3135, 3281, 3401, 3443, 3479, 3624, 3813, 3931
headroom gate, 393, 2240, 3652, 3813
headroom gauge, 211, 262, 423, 494, 556, 635, 649, 763, 775, 852, 1003, 1132, 1213, 1277, 1348, 1432, 2247, 2387, 2714, 3135, 3281, 3401
headroom identity, 1180, 1221, 1303, 1345, 2241, 2312, 2372
hedonic adaptation, 30, 184, 213
held-out skill, 720, 1090
Herfindahl–Hirschman Index, 3181
hindsight retrofitting, 119
homoglyph substitution, 3669, 3743, 3830
honesty appendix, 143, 2321, 2490, 3080, 3917
honesty register, 2373, 3931
Hong Kong, 3400, 3636, 3887, 3935

identified ratios, 1188
impedance matching, 553, 2417
implementation intention, 60
in-band signalling, 3762
independence, 15, 42, 81, 252, 336, 412, 579, 643, 648, 737, 803, 884, 1057, 1097, 1154, 1194, 1266, 1443, 1643, 1729, 1989, 2056, 2248, 2313, 2382, 2497, 2578, 2606, 2654, 2699, 2763, 2809, 2959, 3086, 3135, 3182, 3260, 3278, 3317, 3363, 3400, 3443, 3476, 3511, 3555, 3593, 3626, 3799, 3846, 3864, 3942
informal value transfer, 3553
information budget, 1920
institutional decay, 648
instrument audit, 120, 2312, 2374, 3215, 3544, 3582, 3617
intensification, 2586, 2850
internal surplus, 944, 1131, 1592, 2255, 2326, 2418, 2588, 2626, 2962, 3061, 3343,

3997
Islamic finance, 3474
ISO 20022, 3719
ISO 4217, 3719

Jensen correction, 2240, 2474
job concentration, 393, 3208, 3438

kill conditions, 1265, 1375, 1491, 1566, 2092, 2373, 2985, 3081
kill register, 1565

launch budget, 805, 1188, 1308, 1484, 1703, 2420
leakage and payback, 3509
leaky integrator, 197, 389, 557, 784, 924, 933, 1016, 1111, 1131, 1189, 1255, 1360, 1474, 1645, 1724, 1734, 1816, 2117, 2247, 2384, 2467, 2506, 2701, 2739, 2879, 3019
legal entity identifier, 3719
legal personality, 3191, 3228, 3472, 3548, 3629, 3770, 3936
Lens Theorem, 1007, 1069, 1250, 1534, 1703, 2279, 2380
lever completeness, 3015
lever ranking, 410, 775
lever reserve, 2697
likelihood ratio, 1728, 1967, 2051, 2099, 2284, 2373, 2524, 2788, 2808, 2876, 3018, 3106, 3704
lineage census, 323, 1148, 1511
Lipschitz bound, 323
local minimum, 2, 30, 173, 198, 556, 615, 805, 892, 929, 1189, 1417, 2251, 2384, 2491, 2704, 2798, 3133, 3231, 3361, 3401, 3440, 3476, 3555, 3591, 3995
logarithmic ceiling, 242, 316, 323, 438, 543, 556, 614, 804, 918, 1158, 1569, 1672, 2413, 2629
loop gain, 133, 148, 210, 620, 2043, 2111
lower bound, 30, 135, 153, 203, 261, 331, 409, 496, 558, 616, 654, 726, 784, 852, 941, 1139, 1195, 1350, 1422, 1569, 1674, 1859, 2073, 2190, 2412, 2473, 2590, 2705, 2788, 3166, 3217, 3593, 3814, 3831

magnitude estimation, 943, 1009, 1069, 1249, 1368, 1491, 1598, 1703, 2055, 2104, 2278, 2380, 2505, 2628, 2699, 2741, 2812, 3021, 3084
maintainer attrition, 3762
maintenance requirement, 3589, 3866, 3941
many-to-one map, 1566, 2806, 2924
margin call, 3514, 3559, 3590, 3863, 3941
maturity concentration, 3829
mean-field error, 1416
meaningfulness, 1491, 1600
measurement protocol, 23, 198, 495, 649, 721, 849, 1489, 2242, 3080
mechanotransduction, 1735, 1805, 1884, 2097, 2190, 2245, 2392, 2951, 3596, 3782
method of loci, 1624, 1650, 1744, 2454
migration asymptote, 3762
minimum gap, 552, 884, 949, 1866, 1921, 2200, 2981, 3080
minimum sustainable frequency, 23, 47, 140, 705, 1429, 2478, 2505, 2886, 3057, 3494, 3815
minimum viable angle, 1344
moment ceiling, 928
moment equation, 23, 48, 78, 162, 397, 854, 941, 1002, 1080, 1151, 1243, 1345, 1422, 1492, 1572, 1666, 2246, 2313, 2420, 2488, 2533, 3021, 3265, 3700, 3746, 3801, 3918
momentum, 11, 72, 153, 235, 435, 799, 1188, 1334, 1828, 1883, 1927, 2357, 2453, 3687
multiscale correspondence, 2312

narrative encoding, 2173
nesting, 357, 525, 648, 851, 1432, 2476, 2540, 2594, 2736, 3040, 3558
nine forms of capital, 3742, 3829, 3888
no subliminals, 928, 2230, 2245, 3047, 3119
no-signalling, 1987
no-subliminal constraint, 2950
non-coding DNA, 1863
non-parallelism, 30, 63, 153, 240, 261, 323, 405, 545, 564, 612, 676, 722, 785, 868, 940, 1046, 1070, 1132, 1338, 1346, 1456, 1671, 2034, 2073, 2305, 2325, 2385, 2465, 2507, 2587, 2627, 2697, 2737, 2914, 3060, 3404, 3510, 3846
null return, 89, 1256, 1485, 2247

obligation asymmetry, 1416
observed activity ratio, 405
octave conversion, 848, 3089
octave deficit, 23, 2329, 2463, 2505, 2607, 2689, 2708, 2738, 2807, 2878, 3019, 3081
octave folding, 887, 933, 1946, 2372, 2471, 2523, 2586, 2651, 2729, 2828
odd part, 848, 934, 1626, 1703, 1911, 1974, 2014, 2224, 2471, 2586, 2627, 2707, 2809, 3063
off-model items, 1002
offshore financial centres, 3359
one-hot degeneracy, 2586, 2917, 3012, 3035

open problem, 111, 135, 1806, 1865, 1970, 2005, 2138, 2175, 2244, 2321, 2526, 2940, 2950, 3057, 3119, 3747, 3782
open-label placebo, 1735, 2043, 2179, 2245, 2951, 3060, 3119
operating conditions, 143, 2018
operating substance, 3438
operation audit, 1720, 2806, 2883, 2959, 3027
opposition, 42, 324, 406, 614, 918, 1134, 1236, 1345, 1661, 2188, 2529, 2629, 2697, 2920, 3057, 3119, 3259, 3335, 3846, 3951
option value, 800, 3438
order of operations, 650, 1200, 1288, 1418, 1703, 1731, 2525, 2697, 2806, 2902
order proxy, 225, 270, 386, 532, 616, 689, 720, 781, 1138, 1349, 1584, 2303, 2313, 2414, 2577, 2984
ordinal parameters, 2463
over-attribution, 3762

pair surplus, 1344, 1479
PAM constraint, 1805
paper-to-physical ratio, 3863
parabolic reflector, 491
parallel wall, 336, 427, 794, 856, 1147, 1353, 2415, 2735, 3068
participation entropy, 611
participation ratio, 263, 323, 614, 720, 1132, 2153, 2240, 2312, 2377, 2472, 2735
performative utterance, 89
periodic forcing, 806, 955, 1650, 1748, 2382, 2920
pilot versus production, 3797
pitch class, 480, 848, 928, 1626, 1756, 1865, 1946, 2014, 2186, 2298, 2385, 2471, 2507, 2586, 2628, 2698, 2746, 2809, 2876, 2952, 3063, 3120
pitch class collision, 928
placebo analgesia, 1776, 2043, 2181, 2305
planning fallacy, 60
pledged shares, 3589, 3904, 3941
portable law, 3278
postulate badging, 2
postulate classification, 143
postulate method, 1069
postulate pricing, 3080
power analysis, 2950
precedent versus evidence, 1728
prediction horizon, 1920
preregistration, 26, 45, 121, 248, 267, 341, 408, 492, 593, 638, 675, 720, 776, 909, 943, 1066, 1095, 1265, 1415, 1492, 1565, 1709, 1731, 1863, 1964, 1998, 2046, 2126, 2173, 2414, 2504, 2625, 2721, 2788, 2816, 2949, 2984, 3052, 3104, 3134, 3184, 3230, 3280, 3319, 3362, 3403, 3441, 3477, 3512, 3556, 3747
private foundation, 3228
proof-of-work pools, 3797
protocol index, 143

quantity blindness, 2806
queueing slack, 260

Rasch separation, 1289, 1375, 1478, 1491, 2246, 2349, 2420
ratio invariance, 1069, 1560
re-indexing, 2806, 2878, 2959, 3064, 3113
reactance, 26, 40, 60, 89, 119, 144
realisation deficit, 775
receiver bandwidth, 1805, 1864
reciprocity asymmetry, 1344
reciprocity index, 1416
recognition and characterisation, 3228
rectifier, 1188
redundancy deficit, 323
reference card, 143, 2350, 2373, 3080
refinancing wall, 3829, 3869
refinement ladder, 705, 2875
reflexive loop, 2, 48, 78, 108, 121, 161, 1132, 1191, 1340, 1498, 1570, 1699, 1734, 1812, 1870, 1965, 2046, 2103, 2242, 2470, 2512, 2592, 2631, 2704, 2742, 2882, 2957, 3136, 3224, 3232, 3440, 3476, 3750, 3868, 3934
reflexivity, 19, 49, 78, 108, 131, 188, 400, 478, 547, 604, 654, 768, 813, 1010, 1126, 1138, 1332, 1349, 1475, 1498, 1571, 1721, 2083, 2099, 2470, 2512, 2592, 2631, 2732, 2742, 2813, 2883, 3012, 3021, 3125, 3136, 3224, 3233, 3282, 3445, 3480, 3537, 3560, 3595, 3628, 3847
register status vocabulary, 1728
rejection rate, 161, 294, 721, 1454, 1588, 1705, 1728, 2143, 2429, 3129
relabelling invariance, 2626
relationship frequency, 2626, 2697, 2784, 2818, 2876, 3027, 3084
replication, 21, 159, 255, 317, 328, 487, 516, 573, 617, 712, 763, 783, 989, 1061, 1124, 1185, 1241, 1410, 1441, 1494, 1578, 1702, 1736, 1812, 1871, 2034, 2046, 2100, 2203, 2297, 2352, 2788, 2875, 2984, 3052, 3747, 3911
repricing delta, 3829
reproduction workbook, 393, 3931
research programme costing, 2372
reserve drift, 3829, 3869
reserve realisability, 3316
reserve theorem, 2698, 3015

resistance weighting, 60
resolution trade-off, 60, 92, 1987
restructuring, 1132, 1369, 1453, 2487, 2519, 2714, 2735, 3035, 3151, 3281, 3335, 3397, 3524
retention constant, 1650
return difference, 1081, 1203, 1246, 1363, 1485, 1703, 2420
reverberation time, 487, 491, 785, 849, 1107, 2382
reversibility, 232, 744, 775, 1243, 2180, 2417, 2780, 2861, 3190
ring-fencing, 3192, 3368, 3509, 3622, 3629, 3903, 3941
risk distribution, 3160, 3316
room design, 405
room modes, 932, 1987
roster ceiling, 311, 480, 582, 632, 1436, 1632, 1720, 2483, 2735, 3055
round table, 452, 523, 553, 697, 972
réserve héréditaire, 3228

Sabine and Eyring, 491
scale demotion, 1002, 2532, 2887, 3060
scale type, 21, 56, 77, 104, 138, 165, 397, 548, 602, 782, 853, 930, 1005, 1103, 1185, 1215, 1383, 1421, 1491, 1569, 1735, 1813, 1872, 1925, 2046, 2103, 2177, 2240, 2330, 2380, 2533, 2631, 2741, 2812, 3172, 3220, 3265, 3305, 3319, 3614, 3801
scale-free claim, 848, 1926, 3678, 3792
Schroeder frequency, 508, 928, 1951, 2014, 2579, 2619, 2689
scope correspondence, 1249, 1416, 2058, 2315, 2393, 3815
screening queue, 3719
screening thresholds, 3624
secrecy, 934, 1685, 2190, 2875, 3139, 3191, 3235, 3284, 3367, 3551, 3877, 3991
seeded audit, 1002
seeding strategy, 2240
self-decoding, 89, 193
self-fulfilling prophecy, 2, 99, 154, 1189, 1776, 2044, 2099, 2210, 2470, 3183, 3560
sensitivity analysis, 67, 311, 2090, 2463, 2742, 3353, 3664, 3712, 3793, 3825, 3858
sequence-specific addressing, 1863
sequencing, 274, 336, 775, 1891, 2318, 2977, 3886
settlement diversification, 3359
settlement finality, 3553
settling time, 197, 556, 732, 824, 851, 2251, 2885, 3024
seven repetitions, 2875
sham and nominee, 3863
shared medium, 648
shedding, 224, 260, 379, 475, 530, 648, 735, 788, 857, 1043, 1155, 1280, 2247, 2414
shedding stored order, 260, 475, 788, 1043, 1168
shuffled-input control, 51, 81, 161, 211, 270, 380, 438, 615, 649, 720, 776, 853, 930, 1039, 1072, 1133, 1247, 1375, 1418, 1504, 1565, 1694, 1886, 2040, 2056, 2103, 2198, 2242, 2350, 2372, 2502, 2787, 2925, 2989, 3065, 3104
silo, 8, 355, 611, 805, 2384, 3060
Singapore, 3188, 3400, 3905, 3938
single active commitment, 143
single point of failure, 806, 3181, 3548, 3651, 3735
sonification, 135, 953, 1729, 1812, 1863, 2138, 2173, 2257, 2950, 3062
source dependence, 1728
specific objectivity, 854, 934, 1092, 1491, 1582, 2305, 2420, 2465, 2512, 2592, 2703, 2882
staged process control, 2875
star polygon, 1627, 1650, 3215
stationary action, 2
status register, 2288, 2312
Stevens exponent, 1069
stigmergy, 648, 861, 2416
storage fixed point, 1207, 2463, 2505, 2700
storage half-life, 2586
storage law, 10, 35, 74, 126, 145, 197, 260, 324, 405, 491, 553, 611, 648, 720, 775, 849, 932, 1002, 1070, 1132, 1189, 1247, 1348, 1417, 1492, 1566, 1651, 1771, 1808, 1925, 1989, 2044, 2112, 2178, 2241, 2312, 2376, 2465, 2507, 2587, 2652, 2701, 2815, 2876, 3018, 3116, 3135, 3189, 3459, 3475, 3510, 3573, 3684, 3935
stored error, 1002, 1131, 2418
subadditivity, 332, 409, 1356, 1622, 1650, 2413, 2473
substance challenge, 3863
succession liquidity, 3228, 3914, 3941
succession residual, 3931
sukuk, 3474
surplus ceiling, 277, 353, 491, 611, 941, 1566, 1651, 2104, 2388, 2654, 2736, 2822, 2896, 3119, 3626
surplus functional, 30, 135, 144, 198, 261, 331, 406, 494, 558, 616, 649, 726, 776, 852, 929, 1004, 1070, 1131, 1192, 1261, 1345, 1417, 1492, 1569, 1671, 1746, 1853, 1909, 1926, 2090, 2103, 2178, 2241, 2321, 2377, 2467, 2510, 2587,

2674, 2697, 2752, 3166, 3182, 3281, 3317, 3360, 3441, 3504, 3510, 3700, 3814, 3846, 3917, 3996
surplus headroom, 3624, 3931
surrogate calibration, 1565
survey protocol, 203, 265, 329, 422, 708, 720, 788, 854, 1129, 2263
survivorship, 110, 1686, 1728, 2062, 2099, 2294, 2372, 2524, 2808, 2875, 3018, 3106
survivorship bias, 1728, 2062, 2099, 2372
sustainability ceiling, 201, 260, 342, 577, 614, 679, 723, 784, 1152, 1281, 2407, 2963
sustaining frequency, 906, 976, 2083, 2463, 2594, 3422
sweep schedule, 3400
symbol separation, 2950, 3070, 3119

thermal comfort, 553, 1104
threat register, 3181, 3899, 3937
threshold, 34, 81, 121, 183, 201, 264, 369, 449, 492, 553, 616, 652, 738, 787, 850, 934, 1029, 1070, 1135, 1189, 1248, 1359, 1416, 1536, 1571, 1652, 1732, 1816, 1871, 1921, 1989, 2046, 2101, 2174, 2243, 2351, 2373, 2464, 2506, 2591, 2671, 2700, 2736, 2888, 2958, 3017, 3087, 3133, 3185, 3233, 3298, 3320, 3374, 3401, 3439, 3475, 3510, 3556, 3591, 3625, 3697, 3842, 3892, 3932
time constant, 34, 128, 197, 301, 479, 556, 651, 740, 776, 849, 929, 1154, 1275, 1443, 1688, 1748, 1806, 2100, 2238, 2251, 2320, 2376, 2885
time-frequency limit, 1920
trait window, 2586
transit exposure, 3400, 3553
transit exposure window, 3553
transposition capacity, 3010, 3069, 3080
treaty networks, 3474
trust as retention, 1416, 3166
turn-taking equality, 491, 554, 613, 803, 930, 2296, 2332, 2383
two-source decomposition, 329, 405, 852, 1480, 1676, 2444, 3510
typed parsing, 3719, 3832

uncertainty relation, 72, 165, 1804, 1920, 2027, 2320, 2377, 2475, 2954
unicursal, 1650
unison, 577, 891, 948, 2520, 2697, 2753, 2932, 2962, 3047, 3102
utility ordinal, 3797

validator concentration, 3797
veridical throughput, 1002, 2420
verification threshold, 1416

zero headroom, 287, 700, 1159, 1278, 1344, 1485
zero-input episode, 720